



NVM Express™
Management Interface

Revision 1.1

April 29, 2019

Please send comments and questions to info@nvmexpress.org

NVM Express™ Management Interface revision 1.1 specification available for download at <http://nvmexpress.org>. NVM Express Management Interface revision 1.1 specification consists of the NVM Express Management Interface revision 1.0a specification ratified on April 8th, 2017 along with NVMe-MI 1.0a ECN 001, ECN 002, and ECN 003, TP 6001, TP 6002, TP 6003, TP 6004, TP 6005, TP 6006, TP 6007, and editorial changes including applying the NVM Express™ trademark guidelines.

SPECIFICATION DISCLAIMER

LEGAL NOTICE:

© Copyright 2007 to 2019 NVM Express, Inc. ALL RIGHTS RESERVED.

This NVM Express Management Interface revision 1.1 specification 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 Management Interface revision 1.1 specification 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: “© 2007 to 2019 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
1.2.1	Outside of Scope	7
1.3	Theory of Operation	7
1.3.1	Out-of-Band Theory of Operation	7
1.3.2	In-Band Theory of Operation	8
1.4	NVM Subsystem Architectural Model	8
1.5	NVMe Storage Device Architectural Model	10
1.6	NVMe Enclosure Architectural Model	13
1.7	Conventions	18
1.8	Definitions	19
1.8.1	Carrier	19
1.8.2	Command Message	19
1.8.3	Command Slot	20
1.8.4	Control Primitive	20
1.8.5	NVMe Controller (Controller)	20
1.8.6	NVMe Controller Management Interface (Controller Management Interface)	20
1.8.7	Enclosure Management	20
1.8.8	Enclosure Services Process	20
1.8.9	Expansion Connector	20
1.8.10	Field-Replaceable Unit (FRU)	20
1.8.11	FRU Information Device	20
1.8.12	In-Band	21
1.8.13	Management Controller	21
1.8.14	Management Endpoint or NVMe Management Endpoint	21
1.8.15	Management Endpoint Buffer	21
1.8.16	NVMe Enclosure	21
1.8.17	NVMe Processing	21
1.8.18	NVMe Storage Device	21
1.8.19	NVMe Storage Device FRU	21
1.8.20	NVMe Subenclosure (Subenclosure)	22
1.8.21	NVMe-MI Message	22
1.8.22	NVM Subsystem	22
1.8.23	Out-of-Band	22
1.8.24	Process	22
1.8.25	Request Message	22
1.8.26	Requester	23
1.8.27	Responder	23
1.8.28	Response Message	23
1.8.29	SMBus/I2C Mux	23
1.8.30	Upstream Connector	23
1.8.31	VPD or Vital Product Data	23
1.9	Keywords	23
1.9.1	mandatory	23
1.9.2	may	23
1.9.3	optional	24
1.9.4	R	24
1.9.5	reserved	24
1.9.6	shall	24
1.9.7	should	24
1.10	Byte, Word, and Dword Relationships	24
1.11	References	25
2	PHYSICAL LAYER	27

2.1	PCI Express.....	27
2.2	SMBus/I2C.....	27
2.3	Error Handling.....	30
3	MESSAGE TRANSPORT	31
3.1	NVMe-MI Messages	31
3.1.1	Message Fields	31
3.2	Out-of-Band Message Transport.....	35
3.2.1	MCTP Packet	35
3.2.2	Out-of-Band Error Handling.....	38
3.3	In-Band Tunneling Message Transport	38
4	MESSAGE SERVICING MODEL.....	39
4.1	NVMe-MI Messages	39
4.1.1	Request Messages.....	39
4.1.2	Response Messages.....	40
4.2	Out-of-Band Message Servicing Model.....	43
4.2.1	Control Primitives	45
4.2.2	Out-of-Band Error Handling.....	53
4.2.3	Management Endpoint Buffer.....	54
4.3	In-Band Tunneling Message Servicing Model.....	55
4.3.1	NVMe-MI Send Command	55
4.3.2	NVMe-MI Receive Command.....	62
5	MANAGEMENT INTERFACE COMMAND SET	68
5.1	Configuration Get.....	71
5.1.1	SMBus/I2C Frequency (Configuration Identifier 01h).....	72
5.1.2	Health Status Change (Configuration Identifier 02h).....	73
5.1.3	MCTP Transmission Unit Size (Configuration Identifier 03h)	73
5.2	Configuration Set.....	73
5.2.1	SMBus/I2C Frequency (Configuration Identifier 01h).....	74
5.2.2	Health Status Change (Configuration Identifier 02h).....	74
5.2.3	MCTP Transmission Unit Size (Configuration Identifier 03h)	76
5.3	Controller Health Status Poll	76
5.3.1	Filtering by Controller Type	80
5.3.2	Filtering by Controller Health Status Changed Flags.....	80
5.4	Management Endpoint Buffer Read	83
5.5	Management Endpoint Buffer Write	84
5.6	NVM Subsystem Health Status Poll	85
5.7	Read NVMe-MI Data Structure.....	88
5.8	Reset	94
5.9	SES Receive.....	94
5.10	SES Send	95
5.11	VPD Read.....	96
5.12	VPD Write	97
6	NVM EXPRESS ADMIN COMMAND SET	99
6.1	Request and Response Data	103
6.2	Status.....	104
6.3	Sanitize Operation	104
7	PCIe COMMAND SET (OPTIONAL).....	106
7.1	PCIe Configuration Read.....	108

7.2	PCIe Configuration Write	109
7.3	PCIe I/O Read	110
7.4	PCIe I/O Write.....	110
7.5	PCIe Memory Read	111
7.6	PCIe Memory Write	112
8	NVM EXPRESS MANAGEMENT ENHANCEMENTS.....	114
8.1	Identify Controller.....	114
8.2	Management Interface Specific Features.....	115
8.2.1	Controller Metadata	115
8.2.2	Namespace Metadata	117
9	MANAGEMENT ARCHITECTURE	119
9.1	Out-of-Band Operational Times.....	119
9.2	Vital Product Data.....	120
9.2.1	Common Header	121
9.2.2	Product Info Area (offset 8 bytes).....	121
9.2.3	NVMe MultiRecord Area.....	123
9.2.4	NVMe PCIe Port MultiRecord Area	124
9.2.5	Topology MultiRecord Area	126
9.3	Reset	140
9.3.1	NVM Subsystem Reset	140
9.3.2	Controller Level Reset.....	140
9.3.3	Management Endpoint Reset.....	141
9.3.4	SMBus Reset	141
9.4	Security	141
	APPENDIX A – TECHNICAL NOTE: NVM EXPRESS BASIC MANAGEMENT COMMAND	142
	APPENDIX B – EXAMPLE MCTP MESSAGES & MESSAGE INTEGRITY CHECK	147
	APPENDIX C – EXAMPLE NVME-MI MESSAGES OVER SMBUS/I2C	148

1 Introduction

1.1 Overview

The NVM Express™ (NVMe™) interface is a register-level interface that allows in-band host software to communicate with an NVM Subsystem. Since this specification builds on the NVM Express specification, knowledge of the NVM Express specification is assumed.

This specification defines several mechanisms to manage NVMe Storage Devices (refer to section 1.8.18) or NVMe Enclosures (refer to section 1.8.16). One mechanism allows a Management Controller to communicate out-of-band with an NVMe Storage Device or NVMe Enclosure over one or more external interfaces. Another mechanism is the in-band tunneling mechanism which allows the NVMe-MI Management Interface Command Set to be tunneled in-band via the NVMe Admin Commands NVMe-MI Send and NVMe-MI Receive to an NVMe Storage Device or NVMe Enclosure. Refer to the NVM Express specification and section 4.3 of this specification for additional details on the NVMe-MI Send and NVMe-MI Receive commands.

NVMe Storage Devices and NVMe Enclosures that comply with this specification are allowed to support only the out-of-band mechanism, only the in-band tunneling mechanism, or both the out-of-band mechanism and in-band tunneling mechanism.

1.2 Scope

This specification defines an architecture and command set for out-of-band and in-band management of an NVMe Storage Device as well as an architecture and mechanisms for monitoring and controlling the elements of an NVMe Enclosure.

This specification defines the following key capabilities for NVMe Storage Devices:

- Discover NVMe Storage Devices that are present and learn capabilities of each NVMe Storage Device;
- Store data about the host environment enabling a Management Controller or other entity to query the data later;
- Health and temperature monitoring;
- Multiple concurrent commands to prevent a long latency command from blocking monitoring operations;
- Out-of-band mechanism is host processor and operating system agnostic;
- A standard format for VPD and defined mechanisms to read/write VPD contents; and
- Preserves data-at-rest security.

This specification defines the following key capabilities for NVMe Enclosures:

- Discover NVMe Enclosures and learn their capabilities;
- Manage and sense the state of NVMe Enclosure elements such as power supplies, cooling devices, displays, and indicators;
- Multiple concurrent commands to prevent a long latency command from blocking monitoring operations;
- Out-of-band mechanism is host processor and operating system agnostic;
- Discover NVMe Storage Devices that are present in Enclosure slots; and
- Preserves data-at-rest security.

1.2.1 Outside of Scope

The architecture and command set are specified apart from any usage model. This specification does not specify whether the NVMe interface is used to implement a solid-state drive, a main memory, a cache memory, a backup memory, a redundant memory, etc. Specific usage models are outside the scope of this specification.

This interface is NVM technology agnostic and is specified at a level that abstracts implementation details associated with any specific NVM technology. For example, NAND wear leveling, block erases, and other management tasks are abstracted.

The implementation or use of other published specifications referred to in this specification, even if required for compliance with the specification, are outside the scope of this specification (e.g., PCI Express, SMBus/I2C and MCTP).

The management of NVMe over Fabrics is outside the scope of this specification. This specification does not define new security mechanisms.

This specification does not cover management of non-transparent bridges or PCIe switches. Co-ordination between multiple Requesters or a Requester and a device other than a Responder is outside the scope of this specification. Refer to section 1.8 for the definitions of Requester and Responder.

Coordinating concurrency resulting from operations associated with multiple Responders or between host and Management Endpoint operations is outside the scope of this specification.

The specification of specific Enclosure elements that make up an NVMe Enclosure is outside the scope of this specification. Support for cards or modules that connect to a device slot element (slot) of an NVMe Enclosure, that are not NVMe Storage Devices (e.g., GPUs or FPGAs) is outside the scope of this specification.

An enclosure may support comprehensive management capabilities using SCSI Enclosure Services, basic management capabilities using transport specific mechanisms, or no management capabilities. An example of basic enclosure management capabilities is Native PCIe Enclosure Management (NPEM) specified by the PCI-SIG for PCI Express. The specification of such transport specific basic management capabilities is outside the scope of this specification. This specification only defines comprehensive management using SCSI Enclosure Services.

An NVMe Enclosure may contain multiple Enclosure Services Processes. Communication and coordination between the Enclosure Services Processes that manage NVMe Enclosure elements is outside the scope of this specification.

1.3 Theory of Operation

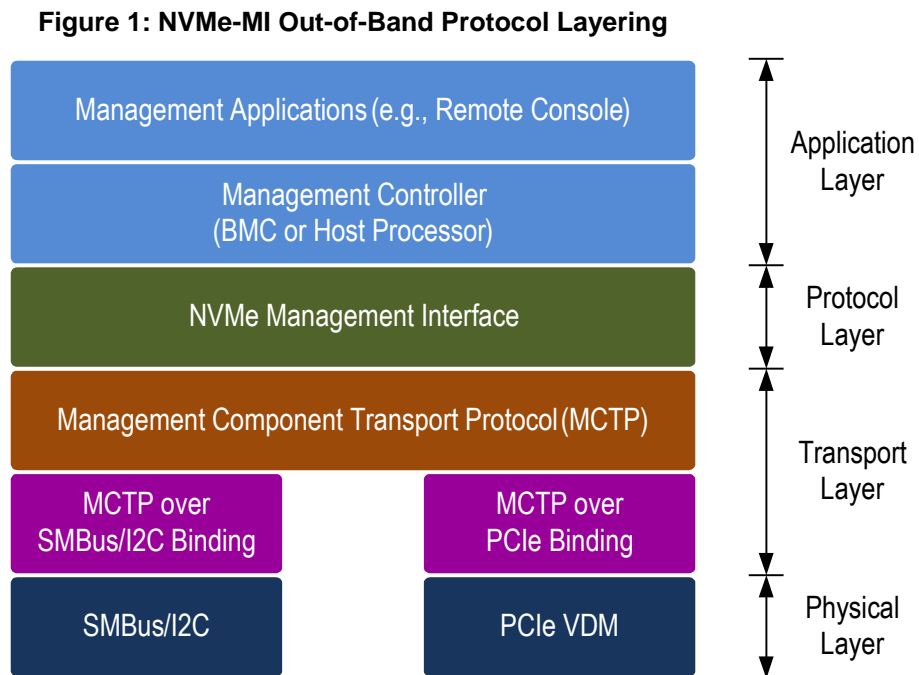
This specification is designed to provide a common interface over multiple physical layers (i.e., PCI Express, SMBus/I2C) for inventory, monitoring, configuration, and change management. This specification provides the flexibility necessary to manage NVMe Storage Devices or NVMe Enclosures using an out-of-band mechanism or in-band tunneling mechanism in a variety of host environments and systems. This specification also defines a FRU Information Device that contains Vital Product Data (refer to section 1.3.1.2).

1.3.1 Out-of-Band Theory of Operation

This specification defines a mechanism for managing NVMe Storage Devices and NVMe Enclosures out-of-band via the Management Component Transport Protocol (refer to section 1.3.1.1).

1.3.1.1 Management Component Transport Protocol

The out-of-band mechanism utilizes the Management Component Transport Protocol (MCTP) as the transport and utilizes existing MCTP SMBus/I2C and PCIe bindings for the physical layer. Command Messages are submitted to one of two Command Slots associated with a Management Endpoint contained in an NVM Subsystem. Figure 1 shows the NVMe-M out-of-band protocol layering.



1.3.1.2 FRU Information Device

This specification defines a mechanism to access a FRU Information Device either via SMBus/I2C as defined by the IPMI Platform Management FRU Information Storage Definition specification or via the VPD Read and VPD Write commands. The data stored in the FRU Information Device is referred to as Vital Product Data (refer to section 9.2). A FRU Information Device may be implemented in a variety of ways (e.g., a serial EEPROM, one-time programmable memory in an NVMe Controller ASIC, etc.).

1.3.2 In-Band Theory of Operation

This specification defines an in-band tunneling mechanism that utilizes the NVMe Admin Commands NVMe-MI Send and NVMe-MI Receive. Refer to the NVM Express specification and section 4.3 of this specification for additional details on the NVMe-MI Send and NVMe-MI Receive commands.

1.4 NVM Subsystem Architectural Model

This specification defines an interface that may be used to manage NVM Subsystems contained within an NVMe Storage Device or NVMe Enclosure.

Management of an NVM Subsystem in the in-band tunneling mechanism and the out-of-band mechanism consists of sending Command Messages and receiving corresponding Response Messages. Command Messages consist of standard NVMe Admin Commands that target a Controller within the NVM Subsystem; commands that provide access to the PCI Express configuration, I/O, and memory spaces of a Controller in the NVM Subsystem; and Management Interface specific commands for inventorying, configuring, and monitoring of the NVM Subsystem. The Command Messages supported by an NVM Subsystem are dependent on the mechanism used to send the NVMe-MI Message (i.e., in-band tunneling mechanism or out-of-band mechanism) and whether the NVM Subsystem is contained within an NVMe Storage Device or NVMe Enclosure.

The architecture and behavior of an NVM Subsystem when the in-band tunneling mechanism is used as defined by the NVM Express specification with extensions defined by this specification. The remainder of this section describes the architecture and behavior of an NVM Subsystem when the out-of-band mechanism is used.

The PCIe ports and SMBus/I2C port of an NVM Subsystem may optionally each contain a single NVMe Management Endpoint (hereafter referred to as simply Management Endpoint). A Management Endpoint is an MCTP endpoint that is the terminus and origin of MCTP packets/messages and is responsible for implementing the MCTP Base Protocol, processing MCTP Control Messages, and internal routing of Command Messages. Each Management Endpoint in an NVM Subsystem has a Port Identifier that is less than or equal to the Number of Ports (NUMP) field value in the NVM Subsystem Information Data Structure.

Management Interface Request Messages and Response Messages are transported as MCTP messages with the Message Type set to NVM Express Management Messages over MCTP (refer to the MCTP IDs and Codes specification). All out-of-band mechanism Command Messages originate with the Management Controller and result in a Response Message from a Management Endpoint.

Each Management Endpoint advertises its unique capabilities. All Management Endpoints may support the same commands even though PCIe ports are full duplex with much higher data rates than SMBus (i.e., both SMBus/I2C and PCIe VDM are capable of providing the same functionality).

Each NVMe Controller in the NVM Subsystem shall provide an NVMe Controller Management Interface (hereafter referred to as simply Controller Management Interface). The Controller Management Interface executes Controller operations on behalf of any Controller (in-band tunneling mechanism) or Management Endpoint (out-of-band mechanism) in the NVM Subsystem. Controllers or Management Endpoints may route commands to any NVMe Controller in the NVM Subsystem. A Controller Management Interface logically executes one operation at a time. A Controller Management Interface is not precluded from executing two or more operations in parallel; however, there shall always be an equivalent pattern of sequential operations with the same results.

Figure 2 illustrates an example NVM Subsystem. The NVM Subsystem contains a single Controller and there is a Management Endpoint associated with the PCIe port.

Figure 2: NVM Subsystem Associated with Single PCIe Port

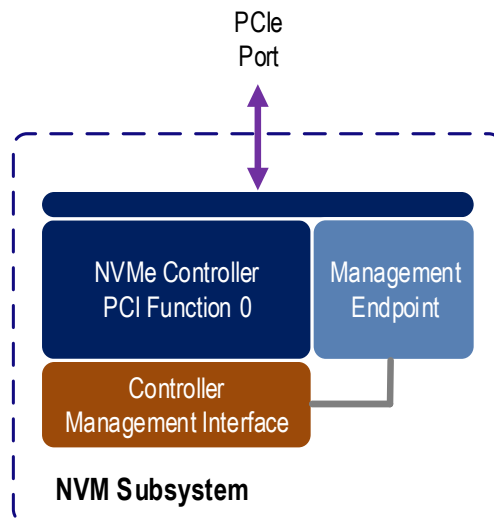
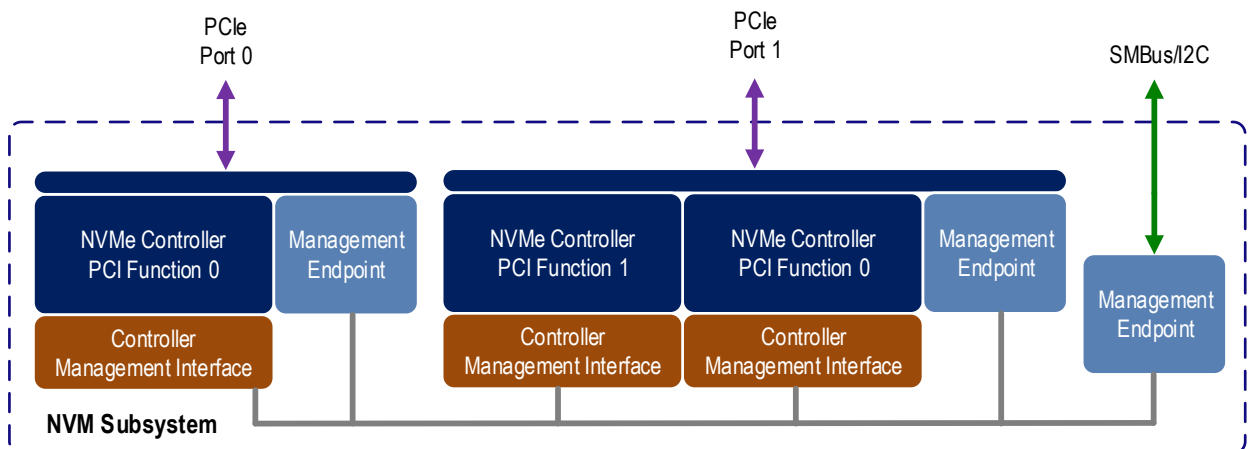


Figure 3 illustrates an example NVM Subsystem that is associated with a dual ported PCIe SSD. The NVM Subsystem contains one Controller associated with PCIe Port 0 and two Controllers associated with PCIe Port 1. There is a Management Endpoint associated with the each PCIe port and the SMBus/I2C port. Since the NVM Subsystem contains a Management Endpoint, all Controllers have an associated Controller Management Interface.

Figure 3: NVM Subsystem with Dual Ported PCIe Ports and an SMBus/I2C Port



1.5 NVMe Storage Device Architectural Model

The architectural model for NVMe Storage Devices that support the in-band tunneling mechanism follows the architectural model defined in the NVM Express specification.

An NVMe Storage Device that implements the out-of-band mechanism but not the in-band tunneling mechanism defined in this specification consists of zero or more NVM Subsystems. An NVMe Storage

Device that implements the in-band tunneling mechanisms defined in this specification consists of one or more NVM Subsystems. Each NVM Subsystem includes one or more Management Endpoints.

An NVMe Storage Device that is a Field-Replaceable Unit (FRU) is a physical component, device, or assembly that is able to be removed and replaced (e.g., by an end user or technician) without having to replace the entire system in which it is contained. Examples of NVMe Storage Device Field-Replaceable Units include a U.2 PCIe SSD, a PCI Express Card Electromechanical (CEM) add-in card, and an M.2 module. The FRU referenced by the FRU Globally Unique Identifier (FGUID) field in the NVM Express specification shall be an NVMe Storage Device Field-Replaceable Unit.

There are many variants of an NVMe Storage Device. One example is an NVMe Storage Device that only contains a single NVM Subsystem. Another example may contain no NVM Subsystems and instead have one or more Expansion Connectors for adding additional NVMe Storage Device FRUs. Such an NVMe Storage Device is referred to as a Carrier. In another example, the NVMe Storage Device may contain one or more NVM Subsystems and one or more Expansion Connectors. NVMe Storage Devices may contain PCIe switches which connect to one or more NVM Subsystems or Expansion Connectors. NVMe Storage Devices may contain SMBus/I2C Muxes that connect to one or more NVM Subsystems or Expansion Connectors.

This specification defines Vital Product Data (VPD) that utilizes the format defined in the IPMI Platform Management FRU Information Storage Definition and is stored in a FRU Information Device. VPD is accessible over any port that supports the out-of-band mechanism or in-band tunneling mechanism. If the NVMe Storage Device has an SMBus/I2C port, then the VPD is accessible using the access mechanism over I2C as defined in the IPMI Platform Management FRU Information Storage Definition.

If an NVMe Storage Device contains multiple NVM Subsystems, then the FRU Information Device associated with each NVM Subsystem is optional since the required FRU Information Device accessible via the Upstream Connector describes the entire NVMe Storage Device. The contents of these additional FRU Information Devices is out of scope for this specification. Refer to section 9.2 for more information

Figure 4 illustrates an NVMe Storage Device that is a single-port PCIe SSD with the FRU Information Device implemented by the NVM Subsystem. Figure 5 illustrates an NVMe Storage Device that is a dual-port PCIe SSD with an SMBus/I2C port and a FRU Information Device implemented using a Serial EEPROM.

Figure 4: Single-Port PCIe SSD

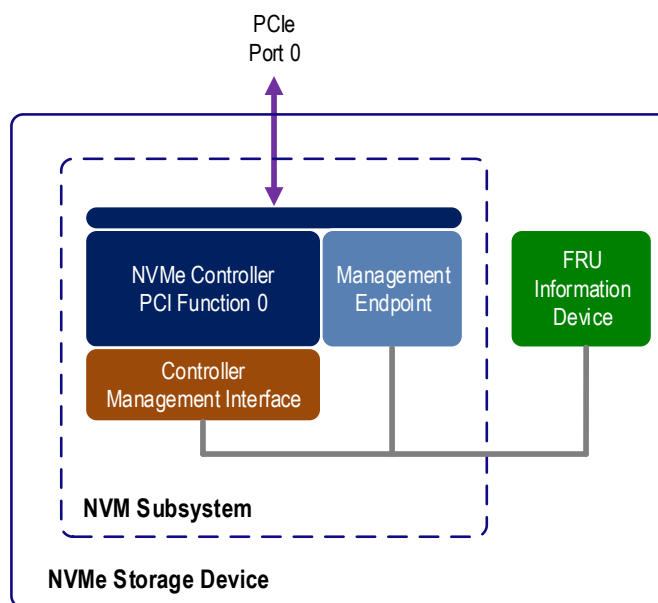
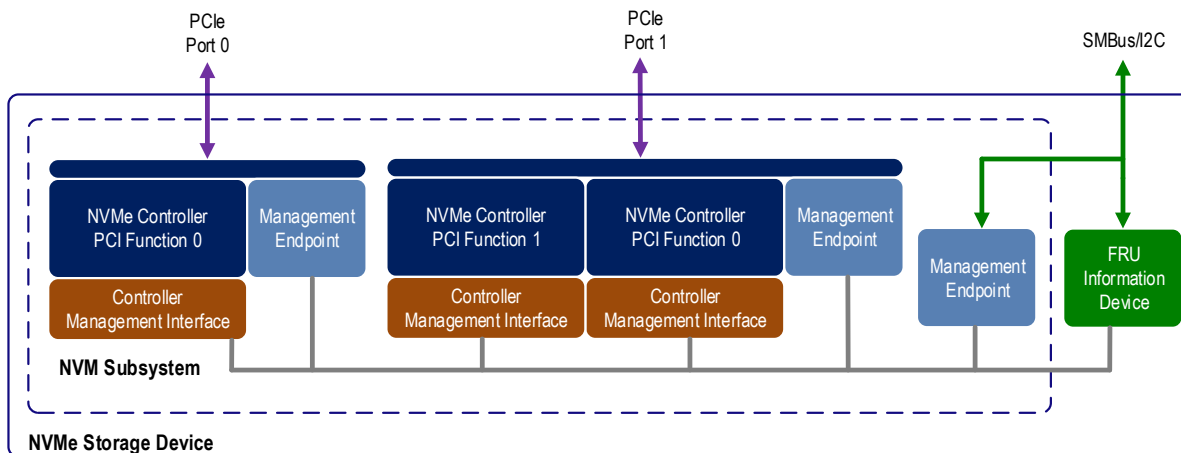


Figure 5: Dual-Port PCIe SSD with SMBus/I2C



An example U.2 form factor NVMe Storage Device with Expansion Connectors (i.e., a Carrier) is shown in Figure 6. This Carrier has two M.2 Expansion Connectors for connecting two M.2 NVMe Storage Device FRUs. The Carrier and each M.2 NVMe Storage Device are separate NVMe Storage Device FRUs, each with their own FRU Information Device. The FRU Information Device on the Carrier is at address A4h and the FRU Information Devices on each M.2 NVMe Storage Device has a default address of A6h and supports the SMBus Address Resolution Protocol (ARP). ARP is used after power is applied to reassign the conflicting A6h addresses before the M.2 FRU Information devices are read. ARP would also be used to reassign the conflicting MCTP addresses and potentially additional elements.

Figure 6: NVMe Storage Device with Expansion Connectors (i.e., a Carrier)

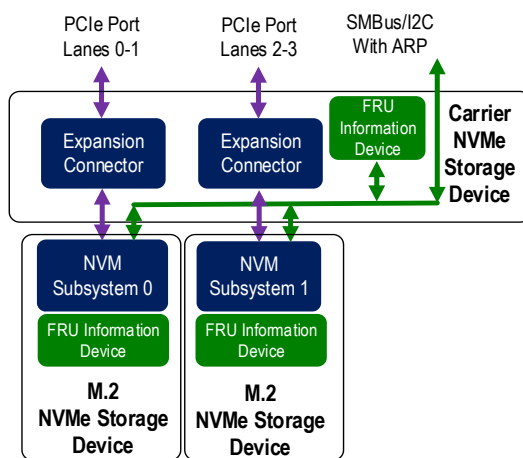
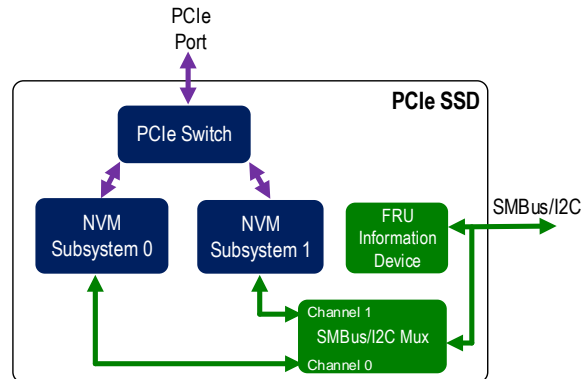


Figure 7 shows an NVMe Storage Device that contains two NVM Subsystems implemented using soldered down BGA packages and a FRU Information Device at address A6h . An NVMe Storage Device without Expansion Connectors that implements an SMBus/I2C port always contains a FRU Information Device at address A6h directly connected to the Upstream Connector. An SMBus/I2C Mux is used in this example instead of ARP to eliminate SMBus/I2C address collisions. The SMBus/I2C Mux is configured by a

Management Controller prior to communications with the selected NVM Subsystem. The FRU Information Device contains the details necessary to configure the SMBus/I2C Mux.

Figure 7: NVMe Storage Device with two NVMe Subsystems and an SMBus/I2C Mux



1.6 NVMe Enclosure Architectural Model

An NVMe Enclosure is a platform, card, module, box, rack, or set of boxes that may provide power, cooling, and mechanical protection for one or more NVM Subsystems. These NVM Subsystems may be part of the NVMe Enclosure itself and/or may be contained in NVMe Storage Devices FRUs that connect to the NVMe Enclosure through one or more NVMe Enclosure slots. An NVMe Enclosure contains one or more NVM Subsystems. NVM Subsystems that are part of an NVMe Enclosure may support just the in-band tunneling mechanism, just the out-of-band mechanism, or both.

An NVMe Enclosure may contain elements that support operation of the NVMe Enclosure (e.g., power supplies, fans, locks, temperature sensors, current sensors, and voltage sensors). An NVMe Enclosure may also contain displays and/or indicators that indicate the state of the NVMe Enclosure (e.g., state of elements, NVM Subsystems, or RAID volumes) and/or NVMe Storage Devices that connect to the NVMe Enclosure. Some of the elements that make up an NVMe Enclosure may be removable and replaceable while the NVMe Enclosure continues to operate normally.

SCSI Enclosure Services - 3 (SES-3) is a standard developed by the American National Standards Institute T10 committee for management of enclosures using the SCSI architecture. While the NVMe and SCSI architectures differ, the elements of an NVMe Enclosure and a SCSI enclosure are similar and the capabilities required to manage elements of an NVMe Enclosure and a SCSI enclosure are similar. Thus, this specification leverages SES for Enclosure Management. SES manages the elements of an enclosure using control and status diagnostic pages transferred using SCSI commands (refer to Enclosure Control and Enclosure Status diagnostic pages in SES-3). This specification uses these same control and status diagnostic pages but transfers them using the SES Send and SES Receive commands. This specification supports only the standalone Enclosure Services Process model as defined in SES.

A Requester manages an NVMe Enclosure using SES Send and SES Receive commands that are part of the Management Interface Command Set (refer to section 5). The SES Send command provides the functionality of the SES-3 SCSI SEND DIAGNOSTIC command and is used by a Requester to send SES control type diagnostic pages to modify the state of the NVMe Enclosure. The SES Receive command provides the functionality of the SES-3 SCSI RECEIVE DIAGNOSTIC RESULTS command and is used by a Requester to retrieve SES status type diagnostic pages that contain various status and warning information available from the NVMe Enclosure.

Refer to SES-3 for a list and description of SES control type diagnostic pages and SES status type diagnostic pages. The mapping of bytes in SES pages to NVMe-MI Request and Response Data is one-to-one where byte x of the SES page maps to byte x in the NVMe-MI Request or Response Data (e.g., byte zero of the SES control type diagnostic page corresponds to byte zero of NVMe-MI Request Data). The

NVMe firmware update process is used (i.e., Firmware Image Download and Firmware Commit commands) to update NVMe firmware. Download Microcode Control and Status diagnostic pages, if supported, shall only be supported on NVMe Enclosure elements.

An Enclosure Services Process, that is logically part of the NVMe Enclosure, is responsible for managing NVMe Enclosure elements and participates in servicing SES Send and SES Receive commands issued by a Requester. Unlike the SES-3 Enclosure Services Process model that maintains state for each I_T nexus (refer to SES-3), unless otherwise noted, this specification requires an NVMe Enclosure to maintain a single global state regardless of the Requester or path used to access that state.

An NVMe Enclosure may contain one or more Subenclosures (refer to SES-3). Each Subenclosure is identified by an SES-3 defined one-byte Subenclosure identifier. If multiple Subenclosures are present, then one of the Subenclosures is designated as the primary Subenclosure and the remaining Subenclosures are secondary Subenclosures. When an NVMe Enclosure consists of only a single Subenclosure, then that Subenclosure is the primary Subenclosure. The Enclosure Services Process associated with the primary Subenclosure is the one that provides access to NVMe Enclosure services information for all Subenclosures. Refer to SES-3 for more information.

Associated with each NVMe Enclosure slot is an SES element that may be used to manage the slot. Refer to SES-3 for more information.

Figure 8 illustrates an example NVMe Enclosure that contains one NVM Subsystem. This NVMe Enclosure has multiple ports that Requesters may use to communicate with the NVMe Enclosure. It also has multiple slots that are used to connect NVMe Storage Devices to the NVMe Enclosure (e.g., PCIe). The mapping of NVMe Enclosure ports to NVM Subsystems, NVMe Controllers within these NVM Subsystems, and NVMe Storage Devices is vendor specific and outside the scope of this specification. An NVMe Enclosure shall contain one or more NVM Subsystems used for Enclosure Management. The NVMe Enclosure in this example may be managed using the out-of-band mechanism via the Responder (Management Endpoint in Figure 8) or using the in-band tunneling mechanism via the NVMe Controller.

Figure 8: Example NVMe Enclosure

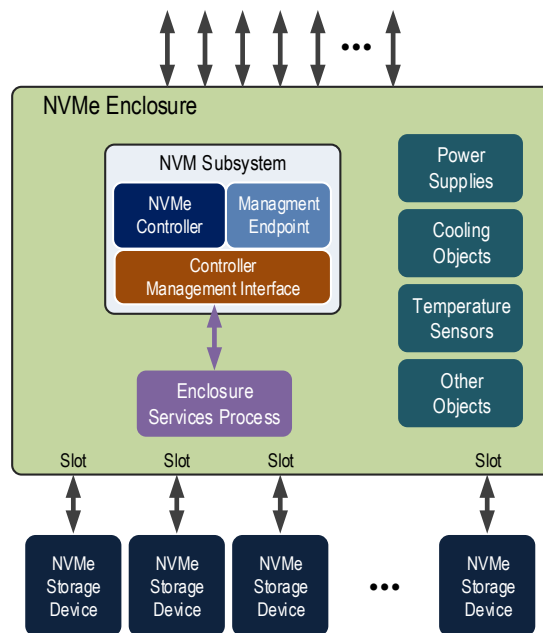


Figure 9 illustrates an example NVMe Enclosure that contains multiple NVM Subsystems and no NVMe Storage Devices. This may represent a software storage appliance. The NVM Subsystems and Controllers contained within these NVM Subsystems may be real or emulated in software. Not all Controllers within these NVM Subsystems are required to have the same capabilities. Some of the possible capability configurations are illustrated in this example. Some Controllers in this example simply provide access to namespaces; others provide access to namespaces and support for the in-band tunneling mechanism; and others provide access to namespaces and support for the out-of-band mechanism.

Figure 9: Example NVMe Enclosure with Multiple NVM Subsystems

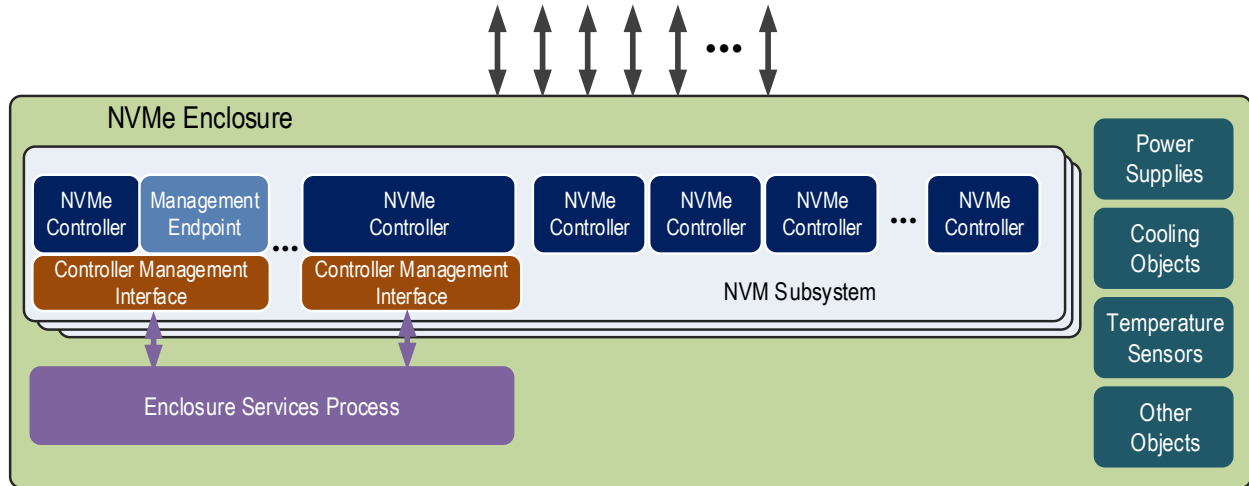


Figure 10 shows an Enclosure that supports two Enclosure Services Processes. Elements of the NVMe Enclosure may be accessible by one or both of these Enclosure Services Processes. The coordination of access to elements by multiple Enclosure Services Processes is outside the scope of this specification.

Figure 10: Example NVMe Enclosure with Multiple Enclosure Services Processes

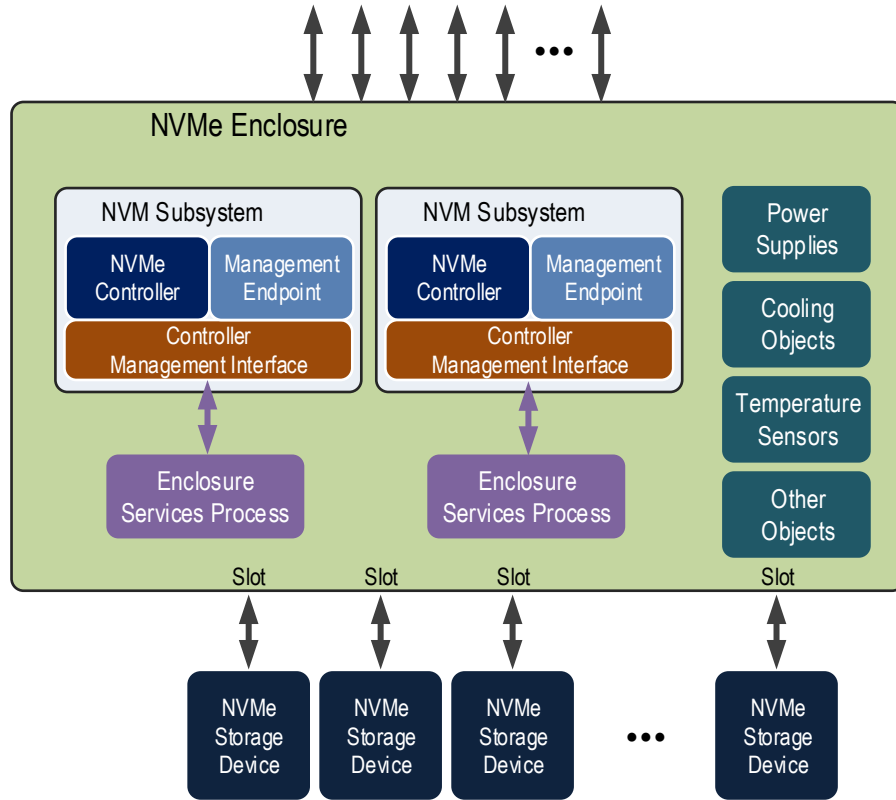
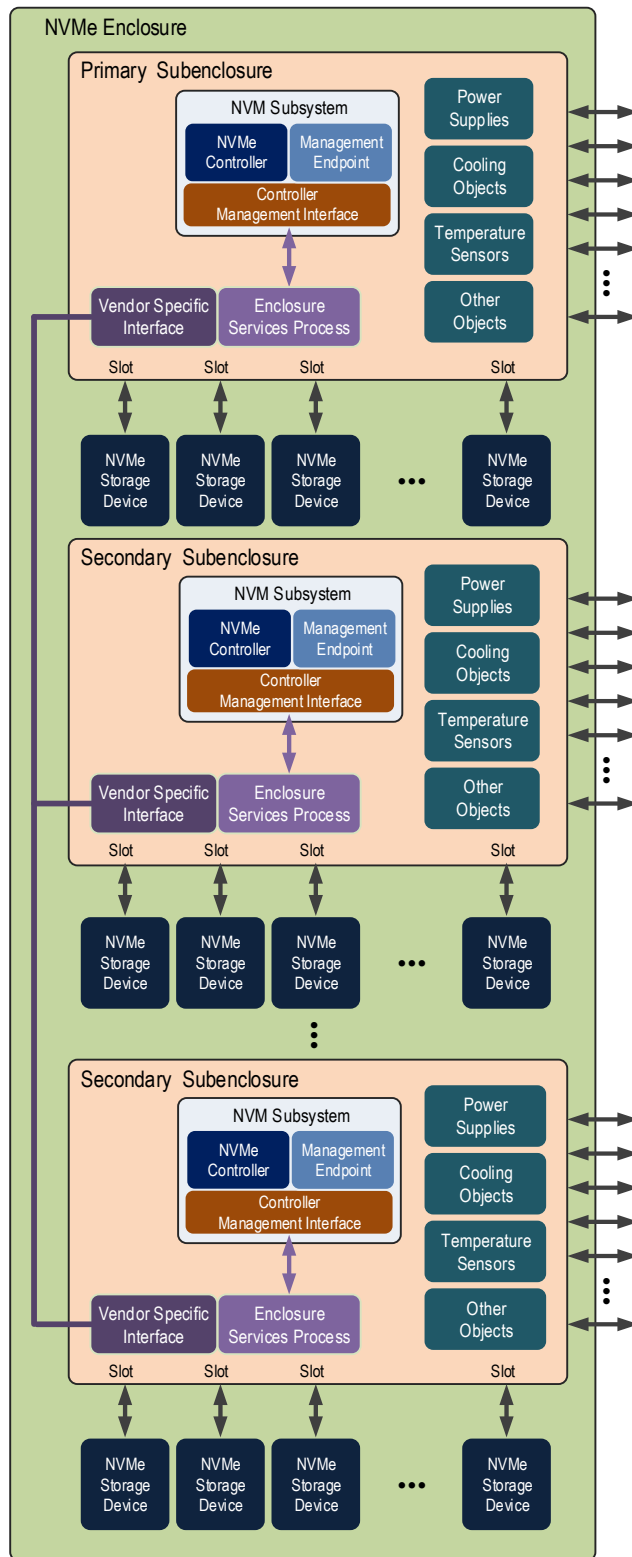


Figure 11 shows an NVMe Enclosure that consists of multiple Subenclosures. Each Subenclosure in this example contains an Enclosure Services Process. NVMe Enclosure services information from Subenclosures is combined into a single set of SES diagnostic pages by the primary Subenclosure. A Subenclosure identifier is used to distinguish from which Subenclosure the information was obtained. Refer to SES-3 for more information. A primary Subenclosure may access NVMe Enclosure services information in Subenclosures using the out-of-band mechanism, the in-band tunneling mechanism, or both; or may use a vendor specific interface. This example illustrates the use of a vendor specific interface.

Figure 11: Example NVMe Enclosure with Subenclosures



Certain NVMe Enclosure behaviors are managed by setting controls and testing status of elements within an NVMe Enclosure. An Enclosure Services Process may monitor a variety of warning and error conditions. These conditions may be communicated to the Requester through polling by the Requester (refer to Enclosure Services Management mode page in SES-3 for details).

The mapping of SES-3 sense keys and additional sense codes associated with CHECK CONDITION status to NVMe-MI Response Message Status values is shown in Figure 12. The asynchronous event notification reporting mechanism described in SES-3 is not supported by this specification.

Figure 12: Mapping of SES-3 Sense Keys and Additional Sense Codes to Response Message Status

Response Message Status Values	SES-3	
	Sense Key	Additional Sense Code
Enclosure Services Failure	HARDWARE ERROR	ENCLOSURE SERVICES FAILURE
Enclosure Services Transfer Failure		ENCLOSURE SERVICES TRANSFER FAILURE
Enclosure Failure		ENCLOSURE FAILURE
Enclosure Services Transfer Refused	HARDWARE ERROR or ILLEGAL REQUEST	ENCLOSURE SERVICES TRANSFER REFUSED
Unsupported Enclosure Function	ILLEGAL REQUEST	UNSUPPORTED ENCLOSURE FUNCTION
Enclosure Services Unavailable	NOT READY	ENCLOSURE SERVICES UNAVAILABLE
Enclosure Degraded	RECOVERED ERROR	WARNING – ENCLOSURE DEGRADED

1.7 Conventions

Hardware shall return zero for all bits, fields, and registers that are marked as reserved. The Requester should not rely on a value of zero being returned as future revisions of this specification may contain non-zero values. The Requester should write all reserved bits and registers with the value of zero. Future revisions of this specification may rely on a zero value being written for backward compatibility.

Hexadecimal (i.e., base 16) numbers are written with a lower case “h” suffix (e.g., 0FFFh, 80h). Hexadecimal numbers larger than eight digits are represented with an underscore character dividing each group of eight digits (e.g., 1E_DEADBEEFh).

Binary (i.e., base 2) numbers are written with a lower case “b” suffix (e.g., 1001b, 10b). Binary numbers larger than four digits are written with an underscore character dividing each group of four digits (e.g., 1000_0101_0010b).

All other numbers are decimal (i.e., base 10). A decimal number is represented in this specification by any sequence of digits consisting of only the Western-Arabic numerals 0 to 9 not immediately followed by a lower-case b or a lower-case h (e.g., 175). This specification uses the following conventions for representing decimal numbers:

- the decimal separator (i.e., separating the integer and fractional portions of the number) is a period;
- the thousands separator (i.e., separating groups of three decimal digits in a portion of the number) is a comma;
- the thousands separator is used in only the integer portion of a number and not the fractional portion of a number; and
- the decimal representation for a year does not include a comma (e.g., 2018 instead of 2,018).

SMBus/I2C addresses are written as 8-bit hex values where bits 7:1 contain the 7-bit SMBus/I2C address and bit 0 is cleared to ‘0’.

When a register field is referred to in the document, the convention used is “Register Symbol.Field Symbol” (e.g., the Controller Status (CSTS) register Shutdown Status (SHST) field is referred to by the name

CSTS.SHST). If the register field is an array of bits, the field is referred to as “Register Symbol.Field Symbol (array offset to element)”.

A 0’s based value is a numbering scheme for which the number 0h represents a value of 1h and thus produces the pattern of 0h represents 1h, 1h represents 2h, 2h represents 3h, etc. In this numbering scheme, there is not a method for specifying the value of 0h. Values in this specification are 1-based (i.e., the number 1h represents a value of 1h, 2h represents 2h, etc.) unless otherwise specified.

Some parameters are defined as a string of ASCII or UTF-8 characters. ASCII data fields shall contain only code values 20h to 7Eh. UTF-8 is backwards compatible with ASCII encoding and supports additional characters with variable length encoding. For the string “Copyright”, the character “C” is the first byte, the character “o” is the second byte, etc. The string is left justified and shall be padded with spaces (ASCII character 20h) to the right if necessary.

A range of numeric values is represented in this specification in the form “a to z”, where a is the first value included in the range, all values between a and z are included in the range, and z is the last value included in the range (e.g., the representation “0h to 3h” includes the values 0h, 1h, 2h, and 3h).

Size values are shown in binary units or decimal units. The symbols used to represent these values are as shown in Figure 13.

Figure 13: Decimal and Binary Units

Decimal		Binary	
Symbol	Power (base-10)	Symbol	Power (base-2)
kilo / k	10 ³	kibi / Ki	2 ¹⁰
mega / M	10 ⁶	mebi / Mi	2 ²⁰
giga / G	10 ⁹	gibi / Gi	2 ³⁰
tera / T	10 ¹²	tebi / Ti	2 ⁴⁰
peta / P	10 ¹⁵	pebi / Pi	2 ⁵⁰
exa / E	10 ¹⁸	exbi / Ei	2 ⁶⁰
zetta / Z	10 ²¹	zebi / Zi	2 ⁷⁰
yotta / Y	10 ²⁴	yobi / Yi	2 ⁸⁰

Implementation Specific (Impl Spec) – the controller has the freedom to choose its implementation.

Hardware Initialize (Hwlnit) – The default state is dependent on NVM Express controller and system configuration. The value is initialized at reset (e.g., by an expansion ROM, or in the case of integrated devices, by a platform BIOS).

1.8 Definitions

1.8.1 Carrier

An NVMe Storage Device FRU with one or more Expansion Connectors and zero or more NVM Subsystems.

1.8.2 Command Message

A type of Request Message that contains an NVMe Admin Command, PCIe Command, or NVMe-MI Command.

1.8.3 Command Slot

A logical target within a Management Endpoint where a Management Controller sends a Request Message. Each Management Endpoint has exactly two Command Slots.

1.8.4 Control Primitive

Single-packet Request Messages sent from a Management Controller to a Management Endpoint to affect the servicing of a previously issued Command Message or get the state of a Command Slot and Management Endpoint. Control Primitives are applicable only in the out-of-band mechanism and are prohibited in the in-band tunneling mechanism.

1.8.5 NVMe Controller (Controller)

Refer to the NVM Express specification.

1.8.6 NVMe Controller Management Interface (Controller Management Interface)

An interface associated with each NVMe Controller in the NVM Subsystem that is responsible for executing management operations on behalf of a Management Endpoint.

1.8.7 Enclosure Management

The discovery, monitoring and control of elements that make up an NVMe Enclosure.

1.8.8 Enclosure Services Process

A process that implements Enclosure services for an NVMe Enclosure that supports Enclosure Management. Refer to SCSI Enclosure Services - 3 (SES-3) for more information.

1.8.9 Expansion Connector

A connector that allows an NVMe Storage Device FRU or cable to be attached or removed from a Carrier. Expansion Connectors may be empty or populated. A connector to a non-removable NVMe Storage Device is not an Expansion Connector.

1.8.10 Field-Replaceable Unit (FRU)

A physical component, device, or assembly in a system that is able to be removed and replaced (e.g., by an end user or technician) without having to replace the entire system in which it is contained. The Field-Replaceable Unit described in this specification is an NVMe Storage Device Field-Replaceable Unit (refer to section 1.8.19).

1.8.11 FRU Information Device

A logical or physical device used to hold the VPD. A FRU Information Device may be implemented in a variety of ways (e.g., a serial EEPROM, one-time programmable memory in silicon, etc.).

1.8.12 In-Band

Per the Management Component Transport Protocol (MCTP) Overview White Paper, in-band management is management that operates with the support of hardware components that are critical to and used by the operating system. The in-band communication path defined by this specification is via the NVMe Admin Queue using the NVMe-MI Send and NVMe-MI Receive commands from host software to an NVMe Controller. Refer to the NVM Express specification and section 4.3 of this specification for additional details on the NVMe-MI Send and NVMe-MI Receive commands.

1.8.13 Management Controller

A device (e.g., Baseboard Management Controller) responsible for platform management that uses the NVM Express Management Interface to communicate to Management Endpoints.

1.8.14 Management Endpoint or NVMe Management Endpoint

An MCTP endpoint associated with an NVM Subsystem (e.g., an NVMe SSD or NVMe Enclosure) that is the terminus and origin of MCTP packets/messages and which processes Request Messages.

1.8.15 Management Endpoint Buffer

An intermediate buffer defined by this specification to allow servicing out-of-band NVMe-MI Messages that have a Message Body that is larger than the 4,224-byte limit that is specified by the NVMe Management Messages over MCTP Binding Specification.

1.8.16 NVMe Enclosure

A platform, card, module, box, rack, or set of boxes that may provide power, cooling, mechanical protection and/or external interfaces for zero or more NVMe Storage Device FRUs. An NVMe Enclosure contains one or more NVM Subsystems and one or more Enclosure Services Processes.

1.8.17 NVMe Processing

NVMe command processing as defined by the NVM Express specification. The term NVMe Processing is used to distinguish command processing as defined by the NVM Express specification from the Command Message processing defined by this specification (refer to section 1.8.24).

1.8.18 NVMe Storage Device

A logical or physical component, device, or assembly that contains at least one NVM Subsystem or Expansion Connector, at least one Upstream Connector, and at least one FRU Information Device. An NVMe Storage Device that implements the out-of-band mechanism contains at least one Management Endpoint and a Controller Management Interface per Controller. An NVMe Storage Device contains zero or more PCIe switches and SMBus/I2C Muxes. An NVMe Storage Device shall comply with the NVM Express specification. In this specification, NVMe Storage Devices shall also comply with this specification.

1.8.19 NVMe Storage Device FRU

An NVMe Storage Device that is able to be removed and replaced (e.g., by an end user or technician) without having to replace the entire system in which it is contained. Examples of NVMe Storage Device

Field-Replaceable Units include a U.2 PCIe SSD, a PCI Express Card Electromechanical add-in card, or an M.2 module.

1.8.20 NVMe Subenclosure (Subenclosure)

A portion of an NVMe Enclosure accessed through a primary NVMe Enclosure's Enclosure Services Process. Refer to SCSI Enclosure Services - 3 (SES-3) for more information.

1.8.21 NVMe-MI Message

A type of MCTP Message that is defined by this specification in sections 3.1 and 4.1. See the MCTP IDs and Codes specification and the NVMe Management Messages over MCTP Binding Specification for more details on this type of MCTP Message (note that NVMe-MI Messages are referred to as NVM Express Management Messages over MCTP in these specifications).

1.8.22 NVM Subsystem

This specification extends the definition of an NVM Subsystem defined in the NVM Express specification (e.g., by adding a Management Endpoint, Controller Management Interface, etc.). NVMe Enclosures and NVMe Storage devices that are not Carriers have one or more NVM Subsystems. Carriers have zero or more NVM Subsystems.

1.8.23 Out-of-Band

Per the Management Component Transport Protocol (MCTP) Overview White Paper, out-of-band management is management that operates with hardware resources and components that are independent of the operating system's control. The out-of-band communication paths supported by this specification are via MCTP over SMBus/I2C or MCTP over PCIe VDM from a Management Controller to a Management Endpoint. In addition, this specification supports the out-of-band access mechanism defined by the IPMI Platform Management FRU Information Storage Definition specification for accessing a FRU Information Device from a Management Controller over SMBus/I2C.

1.8.24 Process

This is the state when a Command Message is processed. Processing of a Command Message consists of checking for errors with the Command Message and performing the actions specified by the Command Message. This state is applicable in both the out-of-band mechanism and the in-band tunneling mechanism. Refer to section 4.2 for additional details on the Process state in the out-of-band mechanism. Refer to section 4.3 for additional details on the Process state in the in-band tunneling mechanism.

This specification uses the terms process/processing/processed to refer to actions performed in the Process state. These terms are distinct from the NVMe Processing term used to describe NVMe command processing as defined by the NVM Express specification (refer to section 1.8.17 in this specification).

1.8.25 Request Message

An NVMe-MI Message originating from a Requester. A Request Message may be a Command Message or a Control Primitive. Request Messages may be used in both the out-of-band mechanism and the in-band tunneling mechanism.

1.8.26 Requester

The entity that sends Request Messages and receives Response Messages. For the out-of-band mechanism, the Requester is a Management Controller. For the in-band tunneling mechanism, the Requester is host software.

1.8.27 Responder

The entity that receives Request Messages and sends back Response Messages. For the out-of-band mechanism, the Responder is a Management Endpoint. For the in-band tunneling mechanism, the Responder is an NVMe Controller.

1.8.28 Response Message

An NVMe-MI Message originating from a Responder in response to a Request Message. Response Messages may be used in both the out-of-band mechanism and the in-band tunneling mechanism.

1.8.29 SMBus/I2C Mux

A bidirectional SMBus/I2C fan-out multiplexer with one upstream channel and one or more downstream channels configured by an I2C command from a Management Controller to connect zero or more downstream channels to the upstream channel. Each downstream channel may be connected to devices with SMBus/I2C ports. This multiplexer permits multiple devices to use the same SMBus/I2C addresses as long as they are on separate channels.

1.8.30 Upstream Connector

A connector on the NVMe Storage Device or NVMe Enclosure to which a Requester attaches. It may be a physical connector as in U.2 form factors, solder balls as in a BGA form factor, or PCB trace fingers as in a CEM Add in Card or EDSFF form factor. An Upstream Connector may include multiple communications ports, control signals, and power supply rails.

1.8.31 VPD or Vital Product Data

Field-Replaceable Unit (FRU) Information which is stored in a FRU Information Device. This specification defines a standard VPD format for NVMe Storage Devices.

1.9 Keywords

Several keywords are used to differentiate between different levels of requirements.

1.9.1 mandatory

A keyword indicating items to be implemented as defined by this specification.

1.9.2 may

A keyword that indicates flexibility of choice with no implied preference.

1.9.3 optional

A keyword that describes features that are not required by this specification. However, if any optional feature defined by the specification is implemented, the feature shall be implemented in the way defined by the specification.

1.9.4 R

“R” is used as an abbreviation for “reserved” when the figure or table does not provide sufficient space for the full word “reserved”.

1.9.5 reserved

A keyword indicating reserved bits, bytes, words, fields, and opcode values that are set-aside for future standardization. Their use and interpretation may be specified by future extensions to this or other specifications. A reserved bit, byte, word, field, or register shall be cleared to zero, or in accordance with a future extension to this specification. The recipient shall not check the value of reserved bits, bytes, words, or fields.

1.9.6 shall

A keyword indicating a mandatory requirement. Designers are required to implement all such mandatory requirements to ensure interoperability with other products that conform to the specification.

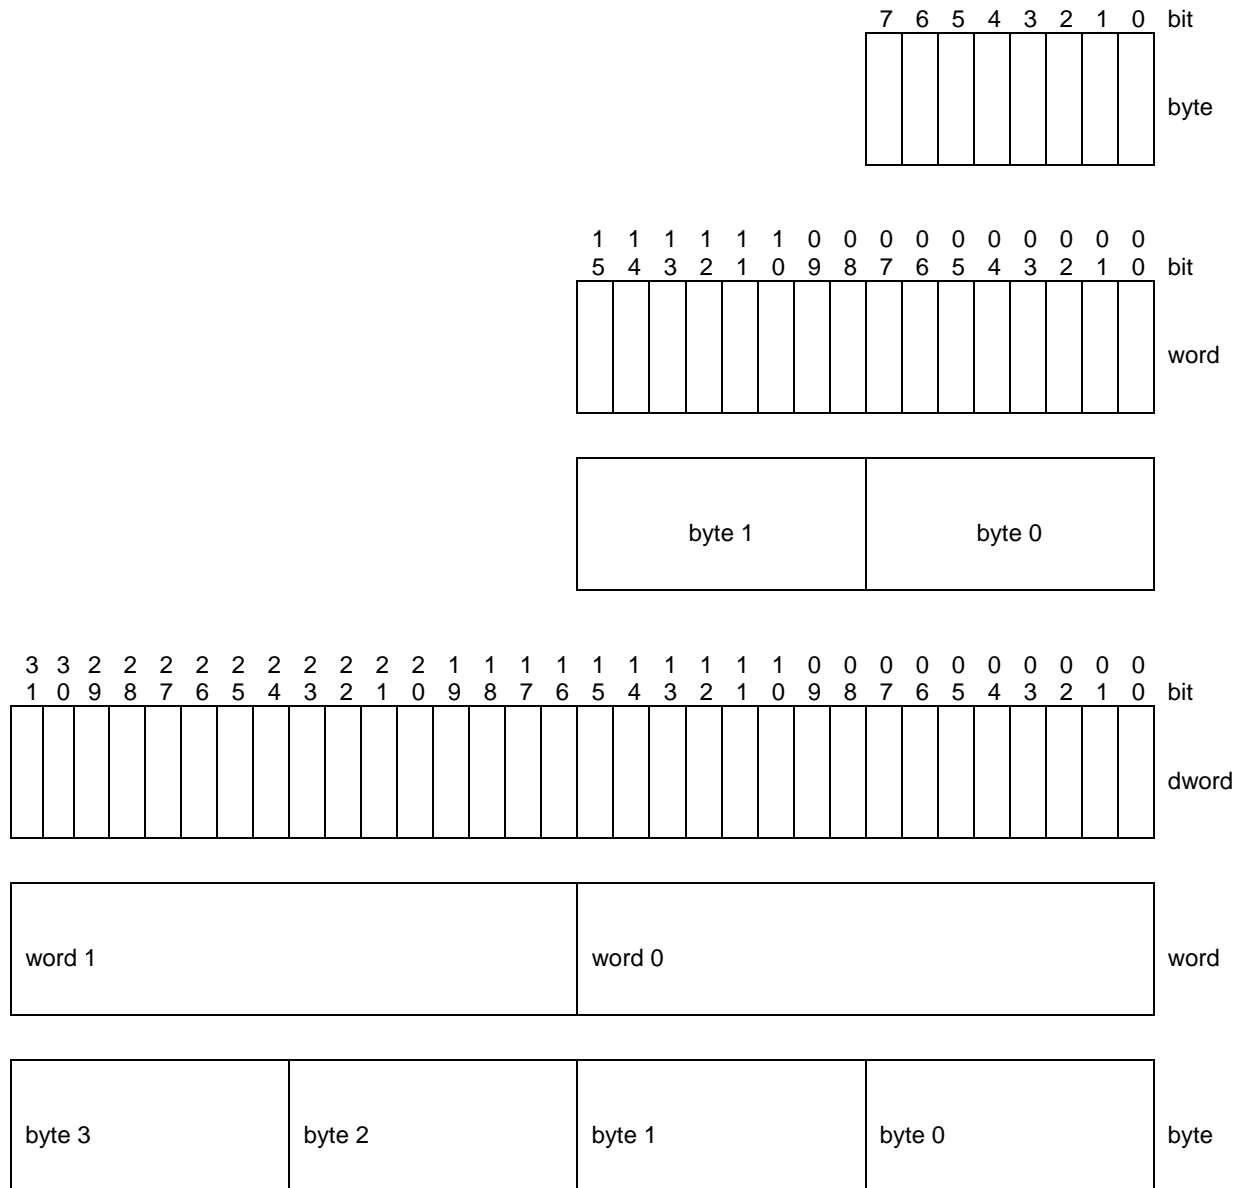
1.9.7 should

A keyword indicating flexibility of choice with a strongly preferred alternative. Equivalent to the phrase “it is recommended”.

1.10 Byte, Word, and Dword Relationships

Figure 14 illustrates the relationship between bytes, words, and dwords. This specification specifies data in a little-endian format.

Figure 14: Byte, Word, and Dword Relationships



1.11 References

- I2C Bus specification, revision 6.0. Available from <http://www.i2c-bus.org>
- IPMI Platform Management FRU Information Storage Definition 1.0, Version 1.3. Available from <http://www.intel.com>.
- INCITS 518-2017 Information Technology – SCSI Enclosure Services – 3 (SES-3). Available from <https://webstore.ansi.org/>.
- MCTP Base Specification (DSP0236), version 1.3.0. Available from <http://www.dmtf.org>.

MCTP IDs and Codes (DSP0239), version 1.5.0. Available from <http://dmf.org>.

MCTP Overview White Paper (DSP2016), version 1.0.0. Available from <http://dmf.org>.

MCTP PCIe VDM Transport Binding Specification (DSP0238), version 1.1.0. Available from <http://www.dmf.org>.

MCTP SMBus/I2C Transport Binding Specification (DSP0237), version 1.1.0. Available from <http://www.dmf.org>.

NVM Express specification, revision 1.3a. Available from <http://www.nvmexpress.org>.

NVMe™ (NVM Express™) Management Messages over MCTP Binding specification (DSP0235), revision 1.0.1. Available from <http://www.dmf.org>.

PCI Express Base Specification, revision 4.0. Available from <http://www.pcisig.com>.

System Management Bus (SMBus) Specification, revision 3.1. Available from <http://www.smbus.org>.

2 Physical Layer

This section describes the physical layers supported by this specification for NVMe Storage Devices or NVMe Enclosures.

2.1 PCI Express

PCI Express is used as a physical layer in both the out-of-band mechanism and the in-band tunneling mechanism in this specification.

For the out-of-band mechanism, a PCIe port in an NVMe Storage Device or NVMe Enclosure may implement a Management Endpoint. If the PCIe port implements a Management Endpoint, the PCIe port shall support MCTP over PCIe Vendor Defined Messages (VDMs) as specified by the Management Component Transport Protocol (MCTP) PCIe VDM Transport Binding Specification.

For the in-band tunneling mechanism, host software issues NVMe Admin Commands (NVMe-MI Send and NVMe-MI Receive) to the NVMe Admin Queue over PCI Express. Refer to the NVM Express specification and section 4.3 of this specification for additional details on the NVMe-MI Send and NVMe-MI Receive commands.

2.2 SMBus/I2C

This section defines the requirements for an NVMe Storage Device or NVMe Enclosure that implements an SMBus/I2C port. The SMBus/I2C physical layer is only applicable for the out-of-band mechanism.

If an NVMe Storage Device or NVMe Enclosure implements an NVM Subsystem with a Management Endpoint associated with an SMBus/I2C port, then that port shall comply to the Management Component Transport Protocol (MCTP) SMBus/I2C Transport Binding Specification.

An NVM Subsystem may also support the NVMe Basic Management Command for health and status polling. The NVMe Basic Management Command is defined as an informative technical note in Appendix A, though it is not recommended for new designs.

Figure 15 lists SMBus/I2C elements that are supported on an NVMe Storage Device or NVMe Enclosure. For each SMBus/I2C element, the default SMBus/I2C address is provided as well as the conditions under which the SMBus/I2C element is required on an NVMe Storage Device or NVMe Enclosure. The presence or absence of Expansion Connectors on an NVMe Storage Device determines which of the two mutually exclusive SMBus/I2C addresses is used for the FRU Information Device. Using a different SMBus/I2C address for the FRU Information Device on NVMe Storage Devices that are Carriers versus non-Carriers avoids SMBus/I2C address conflict when Expansion Connectors are populated with NVMe Storage Devices.

ARP support on SMBus/I2C elements is optional unless there are multiple SMBus/I2C elements in the NVMe Storage Device or NVMe Enclosure with the same default SMBus/I2C address are present on the same SMBus/I2C channel.

Figure 15: SMBus/I2C Elements and Requirements

SMBus/I2C Element	Default SMBus/I2C Address		SMBus ARP Support	Required Element Presence
	Hex Format	Binary Format ¹		
FRU Information Device	A6h	1010_011xb	Optional	Required on an NVMe Storage Device with <i>no</i> Expansion Connectors. Undefined on NVMe Enclosures.
FRU Information Device	A4h	1010_010xb	Optional	Required on Carriers (i.e., an NVMe Storage Device with <i>one or more</i> Expansion Connectors). Undefined on NVMe Enclosures.

Figure 15: SMBus/I2C Elements and Requirements

SMBus/I2C Element	Default SMBus/I2C Address		SMBus ARP Support	Required Element Presence
	Hex Format	Binary Format ¹		
SMBus/I2C Management Endpoint	3Ah	0011_101xb	Optional	Required if an NVMe Storage Device or NVMe Enclosure an SMBus/I2C Management Endpoint.
SMBus/I2C Mux	E8h	1110_100xb	Optional	For NVMe Storage Devices, required if there is more than one SMBus/I2C element on any SMBus/I2C channel with the same SMBus/I2C address that does not support ARP. Undefined on NVMe Enclosures.
Basic Management Command ²	D4h	1101_010xb	Optional	For NVMe Storage Devices, not recommended for new designs. Undefined on NVMe Enclosures.
NOTES: 1. The x represents the SMBus/I2C read/write bit. 2. The NVMe Basic Management Command is defined in Appendix A as an informative technical note.				

Host platforms expecting to be used with one or more Management Endpoints (e.g., data center platforms and workstations) should isolate SMBus/I2C channels to avoid a Management Endpoint conflicting with the address of another SMBus/I2C element. An SMBus/I2C address conflict may occur when a Management Endpoint that does not support ARP is used with platforms that do not isolate SMBus/I2C channels (e.g., some client platforms). ARP may be used to dynamically reassign SMBus/I2C addresses in a system when supported by both the Management Controller and the NVMe Storage Devices or NVMe Enclosure.

SMBus/I2C elements that support ARP should be implemented as Default Slave Address (DSA) devices as defined by the SMBus specification. These devices should not issue “Notify ARP Master” commands.

If ARP is supported by an NVM Subsystem, then all SMBus/I2C elements associated with that NVM Subsystem shall use the SMBus Address Resolution Protocol Unique Device Identifier (UDID) shown in Figure 16. The ARP UDID is a unique identifier. The UDID Vendor ID bits 30 and 31 allow up to four SMBus/I2C elements to be grouped together with the same NVM Subsystem. The only difference within this group of UDIDs is the most significant two bits of the Vendor Specific ID. This fact may be used by the Management Controller to associate an SMBus/I2C Management Endpoint with its corresponding FRU Information Device.

If there are multiple NVM Subsystems in an SMBus ARP-capable NVMe Storage Device or NVMe Enclosure, then the Unique NVM Storage Device ID field of the UDID shall be incremented by one for each NVM Subsystem. If the Upstream Connector has an SMBus/I2C port, then the FRU Information Device associated with that connector shall be present directly on the SMBus/I2C channel connected to the Upstream Connector.

Clock stretching is allowed by the Management Controller, Management Endpoint, and the FRU Information Device. However, implementations are strongly discouraged from using clock stretching so that communications are more predictable with higher throughput.

When a NACK is received, a Management Endpoint shall follow the MCTP SMBus/I2C Transport Binding Specification for a non-bridge endpoint. The Management Endpoint treats a STOP condition due to excessive SMBus NACKs as an implicit Pause Control Primitive. Refer to section 4.2.1.1.

Figure 16: SMBus/I2C Element UDID

Bits	Field	Description										
127:120	Device Capabilities	This field describes the device capabilities.										
		Bits	Description									
		7:6	Address Type: This field describes the type of address contained in the device. Refer to the MCTP SMBus/I2C Transport Binding Specification.									
		5:1	Reserved									
0	PEC Supported: All MCTP transactions shall include a Packet Error Code (PEC) byte. This bit shall be set to '1' to indicate support for PEC.											
119:112	Version and Revision	This field is used to identify the UDID version and silicon revision.										
		Bits	Description									
		7:6	Reserved									
		5:3	UDID Version: This field specifies the UDID version and shall be set to 001b.									
2:0	Silicon Revision ID: This field is used to specify a vendor specific silicon revision level.											
111:96	Vendor ID	This field contains the PCI-SIG vendor ID for the Management Endpoint.										
95:80	Device ID	This field contains a vendor assigned device ID for the Management Endpoint.										
79:64	Interface	This field defines the SMBus version and the Interface Protocols supported.										
		Bits	Description									
		15:08	Reserved									
		07	ZONE: This bit shall be cleared to '0'.									
		06	IPMI: This bit shall be cleared to '0'.									
		05	ASF: This bit shall be set to '1'. Refer to the MCTP SMBus/I2C Transport Binding Specification.									
04	OEM: This bit shall be set to '1'.											
03:00	SMBus Version: This field shall be set to 4h for SMBus Version 2.0, or to 5h for SMBus Version 3.0 and 3.1.											
63:48	Subsystem Vendor ID	This field contains the PCI-SIG vendor ID for the Management Endpoint.										
47:32	Subsystem Device ID	This field contains a vendor assigned device ID for the Management Endpoint.										
31:00	Vendor Specific ID	This field ensures all UDIDs from a vendor are unique and is used to associate elements implemented within an NVMe Storage Device or NVMe Enclosure.										
		Bits	Description									
		31:30	UDID Type: This field distinguishes which NVM Subsystem that implements multiple SMBus elements is providing the UDID. Note that Management Controllers implemented prior to NVMe-MI 1.1 may be incompatible with devices using values 1h and 3h.									
		<table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>FRU Information Device</td> </tr> <tr> <td>1h</td> <td>SMBus/I2C Mux</td> </tr> <tr> <td>2h</td> <td>Management Endpoint</td> </tr> <tr> <td>3h</td> <td>Vendor Specific Devices</td> </tr> </tbody> </table>	Value	Description	0h	FRU Information Device	1h	SMBus/I2C Mux	2h	Management Endpoint	3h	Vendor Specific Devices
Value	Description											
0h	FRU Information Device											
1h	SMBus/I2C Mux											
2h	Management Endpoint											
3h	Vendor Specific Devices											
29:00	UDID Device ID: This field contains a unique vendor assigned ID for the NVM Subsystem. The ID is different in each NVM Subsystem instance and remains static during the life of the device.											

2.3 Error Handling

Physical layer errors are handled as specified by the corresponding physical layer specification and MCTP transport binding specification. This specification does not require any physical layer specific error handling requirements beyond those outlined in the MCTP transport binding specifications.

3 Message Transport

This specification defines an interface that supports multiple message transports. The message format is the same for the out-of-band mechanism and the in-band tunneling mechanism and is described in section 3.1. The out-of-band message transport is described in section 3.2. The in-band tunneling message transport is described in section 3.3.

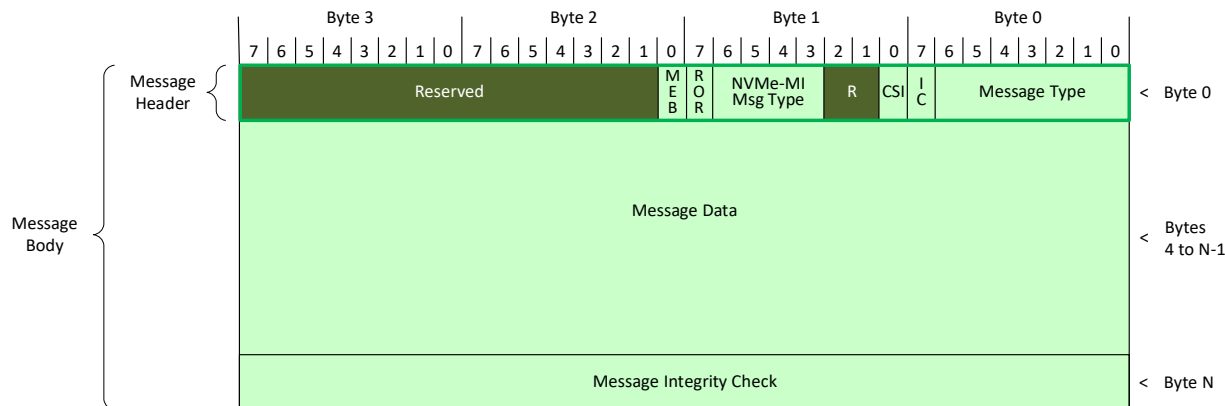
3.1 NVMe-MI Messages

NVMe-MI Messages are used in both the out-of-band mechanism and the in-band tunneling mechanism. The format of an NVMe-MI Message is shown in Figure 17 and Figure 18.

In the out-of-band mechanism, an NVMe-MI Message consists of the payload of one or more MCTP packets. The maximum sized NVMe-MI Message is 4,224 bytes (i.e., 4 KiB + 128 bytes). Refer to the NVMe Management Messages over MCTP Binding Specification. NVMe-MI Messages with lengths greater than 4,224 bytes are considered invalid NVMe-MI Messages. See section 4.2 for details on how NVMe-MI Messages are used in the out-of-band mechanism.

In the in-band tunneling mechanism, NVMe-MI Messages are not split into MCTP packets and the maximum NVMe-MI message size is equal to the Maximum Data Transfer Size (refer to the NVM Express specification). See section 4.3 for details on how NVMe-MI Messages are used in the in-band tunneling mechanism.

Figure 17: NVMe-MI Message



3.1.1 Message Fields

The format of an NVMe-MI Message consists of a Message Header in the first dword, followed by the Message Data. If the Integrity Check (IC) bit is set to '1', then the NVMe-MI Message ends with the Message Integrity Check as shown in Figure 17.

The Message Header contains a Message Type (MT) field and an Integrity Check (IC) bit that are defined by the MCTP Base Specification. The Message Type field specifies the type of payload contained in the message body and is required to be set to 4h in all NVMe-MI Messages (refer to the MCTP IDs and Codes specification). The Integrity Check (IC) field indicates whether the NVMe-MI Message is protected by a Message Integrity Check. All NVMe-MI Messages in the out-of-band mechanism shall be protected by a 32-bit CRC computed over the Message Body contents. The IC field shall be set to '1' in all NVMe-MI Messages in the out-of-band mechanism. The Integrity Check (IC) bit shall be cleared to '0' in all NVMe-MI Messages in the in-band tunneling mechanism.

The Request or Response (ROR) bit in the Message Header specifies whether the NVMe-MI Message is associated with a Request Message or a Response Message. The NVMe-MI Message Type (NMIMT) field specifies whether the Request Message is a Control Primitive or a specific type of Command Message (refer to Figure 23). Finally, the Command Slot Identifier (CSI) bit specifies the Command Slot with which the NVMe-MI Message is associated in the out-of-band mechanism. Refer to section 4.2 for additional information about Command Slots.

The Management Endpoint Buffer (MEB) bit in the Message Header specifies whether Message Data is contained in the associated Message Data field of an NVMe-MI Message or in the Management Endpoint Buffer. This bit should only be set in Command Messages that support Management Endpoint Buffer operation (i.e., those listed in the Management Endpoint Buffer Supported Command List data structure). It is an error to set this bit in any other Command Message and when this occurs it causes the Command Message to complete with an Invalid Parameter Error Response.

Figure 18: NVMe-MI Message Fields

Byte	Description	
0	MCTP Data (MCTPD): This field contains the Message Type and Integrity Check fields as defined by the MCTP Base Specification.	
	Bits	Description
	7	Integrity Check (IC): This bit is defined by the MCTP Base Specification and indicates whether the MCTP message is covered by an overall MCTP Message Integrity Check. All NVMe-MI Messages in the out-of-band mechanism shall be protected by a CRC and thus this bit shall be set to '1' in all out-of-band NVMe-MI Messages. All NVMe-MI Messages in the in-band tunneling mechanism shall not be protected by a CRC and thus this bit shall be cleared to '0' in all in-band NVMe-MI Messages.
6:0	Message Type (MT): This field is defined by the MCTP Base Specification for the message type. This field shall be set to 4h in all NVMe-MI Messages. Refer to the MCTP IDs and Codes specification.	

Figure 18: NVMe-MI Message Fields

Byte	Description																																					
1	<p>NVMe-MI Message Parameters (NMP): This field contains parameters applicable to the NVMe-MI Message.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7</td> <td> <p>Request or Response (ROR): This bit indicates whether the message is a Request Message or Response Message. This bit is cleared to '0' for Request Messages. This bit is set to '1' for Response Messages.</p> </td> </tr> <tr> <td rowspan="6">6:3</td> <td> <p>NVMe-MI Message Type (NMIMT): This field specifies the NVMe-MI Message Type. Refer to the sections referenced in the table below for details about each NVMe-MI Message Type and whether they apply to the out-of-band mechanism, the in-band tunneling mechanism, or both.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> <th>Reference Section</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Control Primitive</td> <td>4.2.1</td> </tr> <tr> <td>1h</td> <td>NVMe-MI Command</td> <td>5</td> </tr> <tr> <td>2h</td> <td>NVMe Admin Command</td> <td>6</td> </tr> <tr> <td>3h</td> <td>Reserved</td> <td>-</td> </tr> <tr> <td>4h</td> <td>PCIe Command</td> <td>7</td> </tr> <tr> <td>5h to Fh</td> <td>Reserved</td> <td>-</td> </tr> </tbody> </table> </td> </tr> <tr> <td>2:1</td> <td>Reserved</td> </tr> <tr> <td>0</td> <td> <p>Command Slot Identifier (CSI): This bit indicates the Command Slot with which the NVMe-MI Message is associated. For Request Messages this bit indicates the Command Slot with which the Request Message is associated. For Response Messages, this bit indicates the Command Slot associated with the Request Message with which the Response Message is associated. This bit is only applicable to NVMe-MI Messages in the out-of-band mechanism. This bit is reserved for NVMe-MI Messages in the in-band tunneling mechanism.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0b</td> <td>Command Slot 0</td> </tr> <tr> <td>1b</td> <td>Command Slot 1</td> </tr> </tbody> </table> </td> </tr> </tbody> </table>	Bits	Description	7	<p>Request or Response (ROR): This bit indicates whether the message is a Request Message or Response Message. This bit is cleared to '0' for Request Messages. This bit is set to '1' for Response Messages.</p>	6:3	<p>NVMe-MI Message Type (NMIMT): This field specifies the NVMe-MI Message Type. Refer to the sections referenced in the table below for details about each NVMe-MI Message Type and whether they apply to the out-of-band mechanism, the in-band tunneling mechanism, or both.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> <th>Reference Section</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Control Primitive</td> <td>4.2.1</td> </tr> <tr> <td>1h</td> <td>NVMe-MI Command</td> <td>5</td> </tr> <tr> <td>2h</td> <td>NVMe Admin Command</td> <td>6</td> </tr> <tr> <td>3h</td> <td>Reserved</td> <td>-</td> </tr> <tr> <td>4h</td> <td>PCIe Command</td> <td>7</td> </tr> <tr> <td>5h to Fh</td> <td>Reserved</td> <td>-</td> </tr> </tbody> </table>	Value	Description	Reference Section	0h	Control Primitive	4.2.1	1h	NVMe-MI Command	5	2h	NVMe Admin Command	6	3h	Reserved	-	4h	PCIe Command	7	5h to Fh	Reserved	-	2:1	Reserved	0	<p>Command Slot Identifier (CSI): This bit indicates the Command Slot with which the NVMe-MI Message is associated. For Request Messages this bit indicates the Command Slot with which the Request Message is associated. For Response Messages, this bit indicates the Command Slot associated with the Request Message with which the Response Message is associated. This bit is only applicable to NVMe-MI Messages in the out-of-band mechanism. This bit is reserved for NVMe-MI Messages in the in-band tunneling mechanism.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0b</td> <td>Command Slot 0</td> </tr> <tr> <td>1b</td> <td>Command Slot 1</td> </tr> </tbody> </table>	Value	Description	0b	Command Slot 0	1b	Command Slot 1
	Bits	Description																																				
	7	<p>Request or Response (ROR): This bit indicates whether the message is a Request Message or Response Message. This bit is cleared to '0' for Request Messages. This bit is set to '1' for Response Messages.</p>																																				
	6:3	<p>NVMe-MI Message Type (NMIMT): This field specifies the NVMe-MI Message Type. Refer to the sections referenced in the table below for details about each NVMe-MI Message Type and whether they apply to the out-of-band mechanism, the in-band tunneling mechanism, or both.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> <th>Reference Section</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Control Primitive</td> <td>4.2.1</td> </tr> <tr> <td>1h</td> <td>NVMe-MI Command</td> <td>5</td> </tr> <tr> <td>2h</td> <td>NVMe Admin Command</td> <td>6</td> </tr> <tr> <td>3h</td> <td>Reserved</td> <td>-</td> </tr> <tr> <td>4h</td> <td>PCIe Command</td> <td>7</td> </tr> <tr> <td>5h to Fh</td> <td>Reserved</td> <td>-</td> </tr> </tbody> </table>	Value	Description	Reference Section		0h	Control Primitive	4.2.1	1h	NVMe-MI Command	5	2h	NVMe Admin Command	6	3h	Reserved	-	4h	PCIe Command	7	5h to Fh	Reserved	-														
		Value	Description	Reference Section																																		
		0h	Control Primitive	4.2.1																																		
1h		NVMe-MI Command	5																																			
2h		NVMe Admin Command	6																																			
3h		Reserved	-																																			
4h	PCIe Command	7																																				
5h to Fh	Reserved	-																																				
2:1	Reserved																																					
0	<p>Command Slot Identifier (CSI): This bit indicates the Command Slot with which the NVMe-MI Message is associated. For Request Messages this bit indicates the Command Slot with which the Request Message is associated. For Response Messages, this bit indicates the Command Slot associated with the Request Message with which the Response Message is associated. This bit is only applicable to NVMe-MI Messages in the out-of-band mechanism. This bit is reserved for NVMe-MI Messages in the in-band tunneling mechanism.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0b</td> <td>Command Slot 0</td> </tr> <tr> <td>1b</td> <td>Command Slot 1</td> </tr> </tbody> </table>	Value	Description	0b	Command Slot 0	1b	Command Slot 1																															
Value	Description																																					
0b	Command Slot 0																																					
1b	Command Slot 1																																					
2	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:1</td> <td>Reserved</td> </tr> <tr> <td rowspan="3">0</td> <td> <p>Management Endpoint Buffer (MEB): This bit indicates whether the Message Data is contained in the Message Data field of this NVMe-MI Message or in the Management Endpoint Buffer. Refer to section 3.1.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0b</td> <td>The Message Data is contained in the Message Data of this NVMe-MI Message.</td> </tr> <tr> <td>1b</td> <td>The Message Data is contained in the Management Endpoint Buffer.</td> </tr> </tbody> </table> </td> </tr> </tbody> </table>	Bits	Description	7:1	Reserved	0	<p>Management Endpoint Buffer (MEB): This bit indicates whether the Message Data is contained in the Message Data field of this NVMe-MI Message or in the Management Endpoint Buffer. Refer to section 3.1.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0b</td> <td>The Message Data is contained in the Message Data of this NVMe-MI Message.</td> </tr> <tr> <td>1b</td> <td>The Message Data is contained in the Management Endpoint Buffer.</td> </tr> </tbody> </table>	Value	Description	0b	The Message Data is contained in the Message Data of this NVMe-MI Message.	1b	The Message Data is contained in the Management Endpoint Buffer.																									
	Bits	Description																																				
	7:1	Reserved																																				
0	<p>Management Endpoint Buffer (MEB): This bit indicates whether the Message Data is contained in the Message Data field of this NVMe-MI Message or in the Management Endpoint Buffer. Refer to section 3.1.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0b</td> <td>The Message Data is contained in the Message Data of this NVMe-MI Message.</td> </tr> <tr> <td>1b</td> <td>The Message Data is contained in the Management Endpoint Buffer.</td> </tr> </tbody> </table>	Value	Description	0b	The Message Data is contained in the Message Data of this NVMe-MI Message.	1b	The Message Data is contained in the Management Endpoint Buffer.																															
	Value	Description																																				
	0b	The Message Data is contained in the Message Data of this NVMe-MI Message.																																				
1b	The Message Data is contained in the Management Endpoint Buffer.																																					
3	Reserved																																					
N-1:4	<p>Message Data (DATA): This field contains the NVMe-MI Message payload. The format of this field depends on the NVMe-MI Message Type.</p>																																					
N+3:N	<p>Message Integrity Check (MIC): If the Integrity Check (IC) bit is set to '1', then this field contains a CRC computed over the contents of the NVMe-MI Message. Refer to section 3.1.1.1.</p> <p>If the IC bit is cleared to '0', then this field is not included in the NVMe-MI Message.</p>																																					

3.1.1.1 Message Integrity Check

If the Integrity Check (IC) bit is set to '1', then the Message Integrity Check field contains a 32-bit CRC computed over the contents of the NVMe-MI Message. The 32-bit CRC required by this specification is

CRC-32C (Castagnoli) which uses the generator polynomial 1EDC6F41h. The Message Integrity Check is calculated using the following Rocksoft™ Model CRC Algorithm parameters:

```
Name      : "CRC-32C"
Width     : 32
Poly      : 1EDC6F41h
Init      : FFFFFFFFh
RefIn     : True
RefOut    : True
XorOut    : FFFFFFFFh
Check     : E3069283h
```

When sending a message, the Message Integrity Check shall be calculated using the following procedure or a procedure that produces an equivalent result:

1. Initialize the CRC register to FFFFFFFFh. This is equivalent to inverting the lowest 32 bits of the NVMe-MI Message (Dword 0 in Figure 17);
2. Append 32 bits of 0's to the end of the Message Data to allow room for the Message Integrity Check (Dword N in Figure 17). This results in the Message Body shown in Figure 17 with the Message Integrity Check field cleared to 0h;
3. Map the bits in the Message Body from step 2 to the coefficients of the message polynomial $M(x)$. Assume the length of $M(x)$ is Y bytes. Bit 0 of byte 0 in the Message Body is the most significant bit of $M(x)$, followed by bit 1 of byte 0, on through to bit 7 of byte Y - 1. Note that the bits within each byte are reflected (i.e., bit n of each byte is mapped to bit (7 - n) resulting in bit 7 to bit 0, bit 6 to bit 1, and so on);

Figure 19: Message Integrity Check Example

Message Body (Length = Y bytes)																											
Byte 0								Byte 1								...	Byte Y - 1										
M(x) =	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	...	0	1	2	3	4	5	6	7		

4. Divide the polynomial $M(x)$ by the generator polynomial 1EDC6F41h to produce the 32-bit remainder polynomial $R(x)$;
5. Reflect each byte of $R(x)$ (i.e., bit n of each byte is mapped to bit (7 - n) resulting in bit 7 to bit 0, bit 6 to bit 1, and so on) to produce the polynomial $R'(x)$;
6. Invert $R'(x)$ to produce the polynomial $R''(x)$; and
7. Store $R''(x)$ in the Message Integrity Check field of the Message Body.

Upon receipt of an NVMe-MI Message, the Message Integrity Check may be validated as follows:

1. Save the received Message Integrity Check;
2. Initialize the CRC register to FFFFFFFFh. This is equivalent to inverting the lowest 32 bits of the NVMe-MI Message (Dword 0 in Figure 17);
3. Clear the Message Integrity Check field to 0h;
4. Map the bits in the Message Body to the coefficients of the message polynomial $M(x)$ as described in step 3 in the Message Integrity Check calculation procedure above;
5. Divide the polynomial $M(x)$ by the generator polynomial 1EDC6F41h to produce the 32-bit remainder polynomial $R(x)$;
6. Reflect each byte of $R(x)$ (i.e., bit n of each byte is mapped to bit (7 - n) resulting in bit 7 to bit 0, bit 6 to bit 1, and so on) to produce the polynomial $R'(x)$;
7. Invert $R'(x)$ to produce the polynomial $R''(x)$; and

8. Compare $R''(x)$ from step 7 to the Message Integrity Check value saved in step 1. If both values are equal, the Message Integrity Check passes.

Refer to Appendix B for artificial messages and their corresponding Message Integrity Check values.

See Section 4.2.1.5 for special requirements on how to construct the Response Message when the Management Controller issues a Replay of a Response Message with a non-zero Response Replay Offset.

3.2 Out-of-Band Message Transport

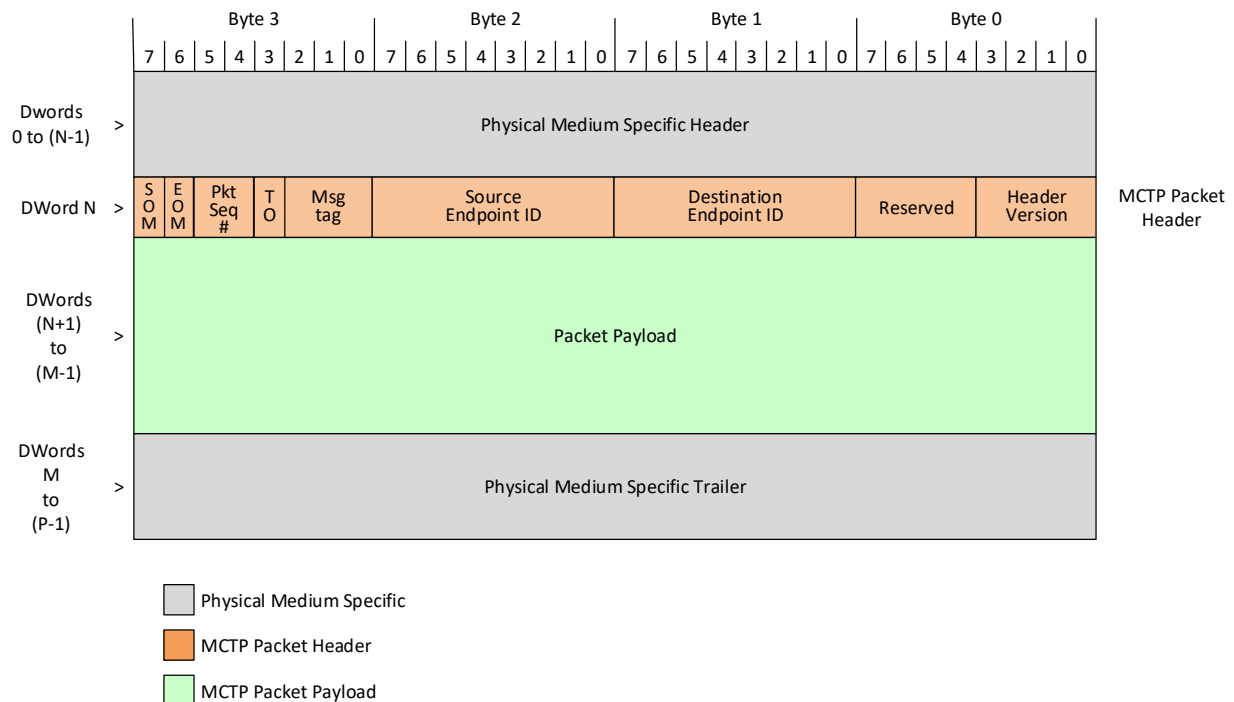
The out-of-band mechanism defined in this specification utilizes MCTP as a reliable in-order message transport between a Management Controller and a Management Endpoint.

This section summarizes the NVMe-MI MCTP packet format. A Management Endpoint compliant to this specification shall implement all required behaviors detailed in the Management Component Transport Protocol (MCTP) Base Specification and corresponding transport binding specification in addition to the requirements outlined in this specification (e.g., the Message Integrity Check algorithm).

3.2.1 MCTP Packet

In the MCTP Base Specification, the smallest unit of data transfer is the MCTP packet. One or more packets are combined to create an MCTP message. In this specification, the MCTP messages are referred to as NVMe-MI Messages (refer to section 1.8.21). Refer to section 3.2.1.1 for details on how MCTP packets are assembled into NVMe-MI Messages. A packet always contains at least 1 byte of payload but the total length shall never exceed the negotiated MCTP Transmission Unit Size. The format of an MCTP packet is shown in Figure 20.

Figure 20: MCTP Packet Format



MCTP specifications use big endian byte ordering while NVM Express specifications use little endian byte ordering. All figures in this specification are illustrated with little endian byte ordering. Note that this pictorial representation does not change the order that bytes are sent out on the physical layer.

The Physical Medium-Specific Header and Physical Medium-Specific Trailer are defined by the MCTP transport binding specification utilized by the port. Refer to the MCTP transport binding specifications.

The Management Component Transport Protocol (MCTP) Base Specification defines the MCTP packet header (refer to DSP0236 for field descriptions). The fields of an MCTP Packet are shown in Figure 21.

Figure 21: MCTP Packet Fields

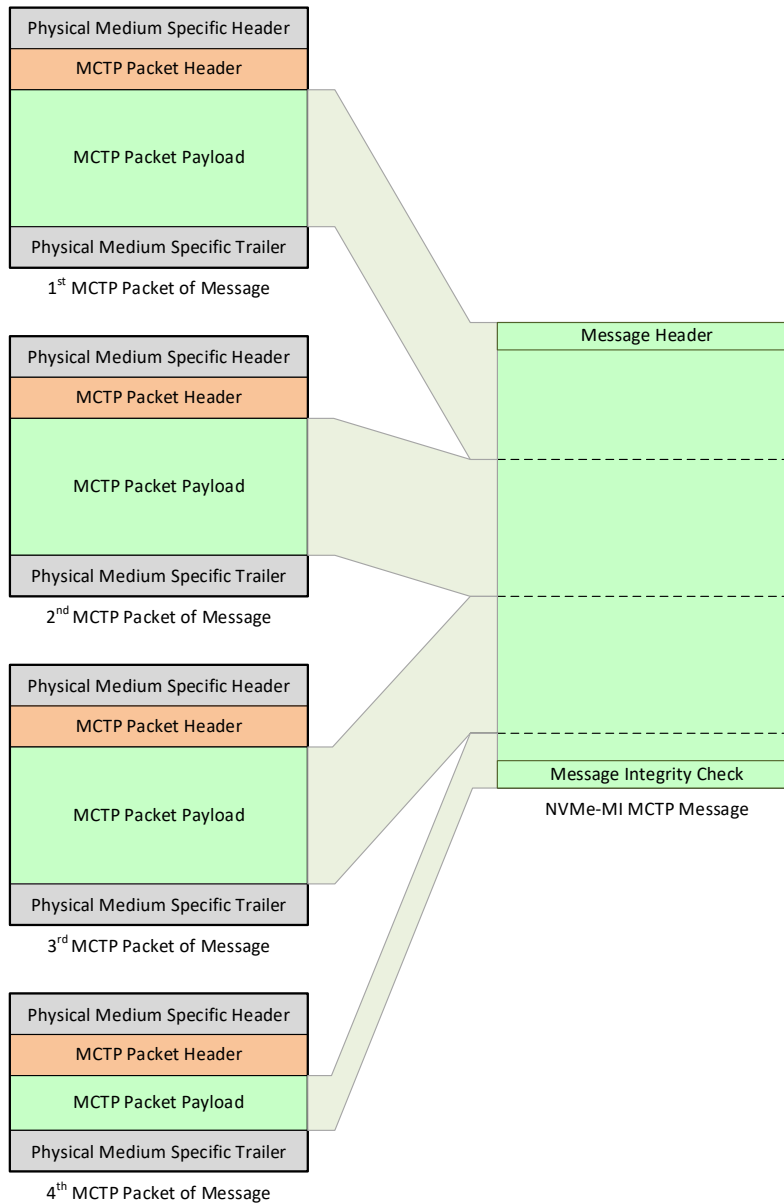
Field Name	Field Size
Medium-Specific Header	varies
Header Version	4 bits
Reserved	4 bits
Destination Endpoint ID	8 bits
Source Endpoint ID	8 bits
Msg tag (Message Tag)	3 bits
TO	1 bit
Pkt Seq #	2 bits
EOM	1 bit
SOM	1 bit
Packet Payload	varies
Medium-Specific Trailer	varies

A compliant Management Endpoint shall implement all MCTP required features defined in the MCTP Base Specification. Optional features may be supported.

3.2.1.1 Packet Assembly into Messages

An NVMe-MI Message may be split into multiple MCTP Packet Payloads and sent as a series of packets. An example NVMe-MI Message whose contents are split across four MCTP packets is shown in Figure 22. Refer to the MCTP Base Specification for packetization and message assembly rules.

Figure 22: NVMe-MI Message Spanning Multiple MCTP Packets



In addition to the requirements outlined in the MCTP Base Specification and transport binding specifications, this specification has the following additional requirements:

- With the exception of the last packet in a message, the MCTP Transmission Unit size of all packets in a given message shall be equal to the negotiated MCTP Transmission Unit Size;
- The MCTP Transmission Unit size of the last packet in a Request Message or Response Message (i.e., the one with the EOM bit set in the MCTP header) shall be the smallest size needed to transfer the MCTP Packet Payload for that Packet with no additional padding beyond any padding required by the physical medium-specific trailer; and
- Once a complete NVMe-MI Message has been assembled, the Message Integrity Check is verified. If the Message Integrity Check passes, then the NVMe-MI Message is processed. If the Message Integrity Check fails, then the NVMe Message is discarded. Refer to section 4.2.

3.2.2 Out-of-Band Error Handling

The Management Endpoint shall drop (silently discard) packets for error conditions as specified in the MCTP Base Specification. Some example conditions which result in discarding packets include unexpected middle or end packets. Silently discarded packets also cause the corresponding bit in the Get State Control Primitive Success Response Fields to be set to '1' (refer to Figure 39).

3.3 In-Band Tunneling Message Transport

The in-band tunneling mechanism in this specification utilizes the NVMe Admin Commands NVMe-MI Send and NVMe-MI Receive as a message transport. Refer to the NVM Express specification and section 4.3 of this specification for additional details on the NVMe-MI Send and NVMe-MI Receive commands.

4 Message Servicing Model

NVMe-MI Messages are used for communication in both the out-of-band and in-band tunneling message servicing models and are described in section 4.1. This specification defines multiple message servicing models. The out-of-band message servicing model is described in section 4.2. The in-band tunneling message servicing model is described in section 4.3.

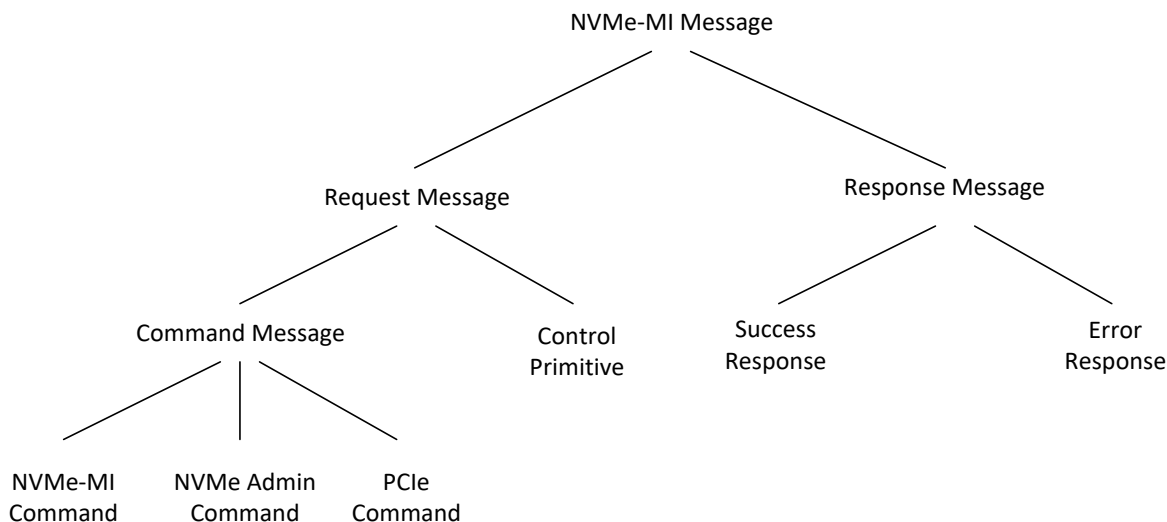
4.1 NVMe-MI Messages

Figure 23 illustrates the taxonomy of NVMe-MI Messages. The two main categories of NVMe-MI Messages are Request Messages and Response Messages. Request Messages are sent by a Management Controller to a Management Endpoint when using the out-of-band mechanism. Request Messages are sent by host software to an NVMe Controller when using the in-band tunneling mechanism. The entity sending the Request Message is collectively referred to as the Requester and the entity receiving the Request Message is collectively referred to as the Responder. After receiving a Request Message, the Responder processes the Request Message. When processing is complete, the Responder sends a Response Message back to the Requester.

A Request Message may be classified as a Command Message or a Control Primitive. Command Messages specify an operation to be performed by the Responder and may be further classified as an NVMe-MI Command, an NVMe Admin Command, or a PCIe Command. Control Primitives are used in the out-of-band mechanism to affect the servicing of a previously issued Command Message or get the state of a Command Slot and Management Endpoint (refer to section 4.2.1).

A Response Message may be classified as a Success Response or an Error Response.

Figure 23: NVMe-MI Message Taxonomy



4.1.1 Request Messages

Request Messages are NVMe-MI Messages that are generated by a Requester to send to a Responder.

Request Messages specify an action to be performed by the Responder. Request Messages are either Control Primitives (refer to section 4.2.1) or Command Messages. The format of the Message Body for a Command Message is command set specific and is specified by the NMIMT field in the Message Header.

The NVMe Management Interface supports three command sets:

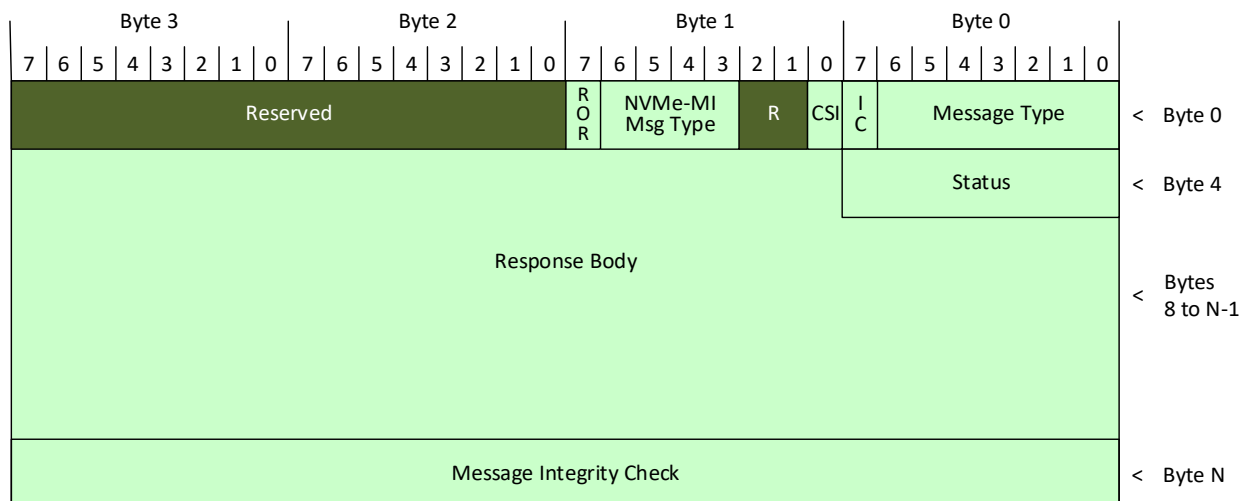
- The Management Interface Command Set is described in section 5;
- The NVM Express Admin Command Set is described in section 6; and
- The PCIe Command Set is described in section 7.

4.1.2 Response Messages

Response Messages are NVMe-MI Messages that are generated when a Responder completes processing of a previously issued Request Message.

The format of a Response Message is shown in Figure 24 and Figure 25. The first dword contains the Message Header. The Status field encodes the status associated with the Response Message. This is followed by the Response Body whose format is NVMe-MI Message Type and Response Message Status specific. Finally, if the Integrity Check (IC) bit is set to ‘1’, then the Response Message ends with the NVMe-MI Message Integrity Check field.

Figure 24: Response Message Format



In the out-of-band mechanism, the CSI bit in the Message Header specifies the Command Slot of the Request Message with which the Response Message is associated. The NVMe-MI Message Type (Msg Type) field contains the value from the same field in the corresponding Request Message.

In the in-band tunneling mechanism, the CSI bit in the Message Header is reserved.

Figure 25: Response Message Fields

Byte	Description
3:0	NVMe-MI Message Header (NMH): Refer to section 3.1.
4	Status (STATUS): This field indicates the status associated with the Response Message. Response Message Status values are summarized in Figure 26.
N-1:5	Response Body (RESPB): This field contains response specific fields whose format is dependent on the NVMe-MI Message Type and Status field.
N+3:N	Message Integrity Check: Refer to section 3.1.

Response Message Status values are summarized in Figure 26. A Response Message Status of Success indicates that the corresponding Request Message completed successfully and that the Response Message is a Success Response. The format of the Response Body for a Success Response is dependent on the NVMe-MI Message Type and is described later in this specification.

A Response Message Status other than Success indicates that an error occurred during servicing of the corresponding Request Message and that the Response Message is an Error Response. The format of the Response Body is dependent on the Response Message Status as shown in Figure 26. If multiple errors are present, a Responder may choose which error status to report.

Figure 26: Response Message Status Values

Value	Description	Error Response Format Section
00h	Success: The command completed successfully.	4.1.2.1
01h	More Processing Required: The Command Message is in progress and requires more time to complete processing. When this Response Message Status is used in a Response Message, a subsequent Response Message contains the result of the Command Message. This Response Message Status shall not be sent more than once per Request Message.	4.1.2.1
02h	Internal Error: The Request Message could not be executed due to a vendor specific internal error.	4.1.2.1
03h	Invalid Command Opcode: The associated command opcode field is not valid. Invalid opcodes include reserved and optional opcodes that are not implemented.	4.1.2.1
04h	Invalid Parameter: Invalid parameter field value. Request Messages received with reserved values in defined fields shall be completed with an Invalid Parameter Error Response. Request Messages received with reserved or unimplemented values in defined fields shall be completed with an Invalid Parameter Error Response. Other error conditions that result in Invalid Parameter Error Response are noted elsewhere in this specification.	4.1.2.2
05h	Invalid Command Size: The Message Body of the Command Message was larger or smaller than that expected due to a reason other than too much or too little Request Data (e.g., the Command Message did not contain all the required parameters or no Request Data was expected but the Request Data is larger than that needed to contain the required parameters). The expected Message Body size is determined by the NVMe-MI Message Type and opcode assuming no other errors are detected (e.g., Invalid Command Opcode or Invalid Parameter).	4.1.2.1
06h	Invalid Command Input Data Size: The Command Message requires Request Data and contains too much or too little Request Data.	4.1.2.1
07h	Access Denied: A Request Message was prohibited from being processed due to a vendor specific protection mechanism.	4.1.2.1
08h to 1Fh	Reserved	-
20h	VPD Updates Exceeded: More updates to the VPD are attempted than allowed.	4.1.2.1
21h	PCIe Inaccessible: The PCIe functionality is not available at this time.	4.1.2.1
22h	Management Endpoint Buffer Cleared Due to Sanitize: An attempt was made access data in the Management Endpoint Buffer that was zeroed due to a sanitize operation.	4.1.2.1
23h	Enclosure Services Failure: The Enclosure Services Process has failed in an unknown manner.	4.1.2.1
24h	Enclosure Services Transfer Failure: Communication with the Enclosure Services Process has failed.	4.1.2.1
25h	Enclosure Failure: An unrecoverable enclosure failure has been detected by the Enclosure Services Process.	4.1.2.1
26h	Enclosure Services Transfer Refused: The NVM Subsystem or Enclosure Services Process indicated an error or an invalid format in communication.	4.1.2.1

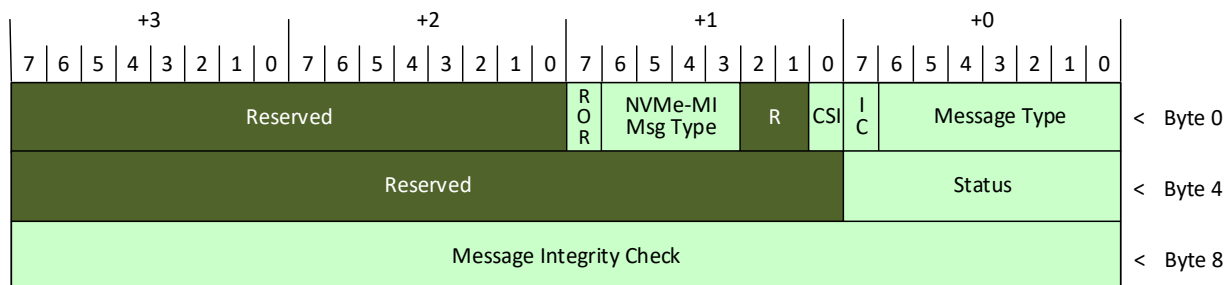
Figure 26: Response Message Status Values

Value	Description	Error Response Format Section
27h	Unsupported Enclosure Function: An SES Send command has been attempted to a simple Subenclosure.	4.1.2.1
28h	Enclosure Services Unavailable: The NVM Subsystem or Enclosure Services Process has encountered an error but may become available again.	4.1.2.1
29h	Enclosure Degraded: A noncritical failure has been detected by the Enclosure Services Process.	4.1.2.1
2Ah	Sanitize In Progress: The requested command is prohibited while a sanitize operation is in progress. Refer to section 9.1.	4.1.2.1
2Bh to DFh	Reserved	-
E0h to FFh	Vendor Specific	Vendor Specific

4.1.2.1 Generic Error Response

A Generic Error Response is generated for errors in which no additional information is provided beyond the Response Message Status. Bytes 7:5 are reserved. The format of a Generic Error Response is shown in Figure 27.

Figure 27: Generic Error Response



4.1.2.2 Invalid Parameter Error Response

An Invalid Parameter Error Response is generated for Error Responses where the Status field is set to Invalid Parameter. The format of an Invalid Parameter Error Response is shown in Figure 28 and the response specific fields are summarized in Figure 29.

Figure 28: Invalid Parameter Error Response

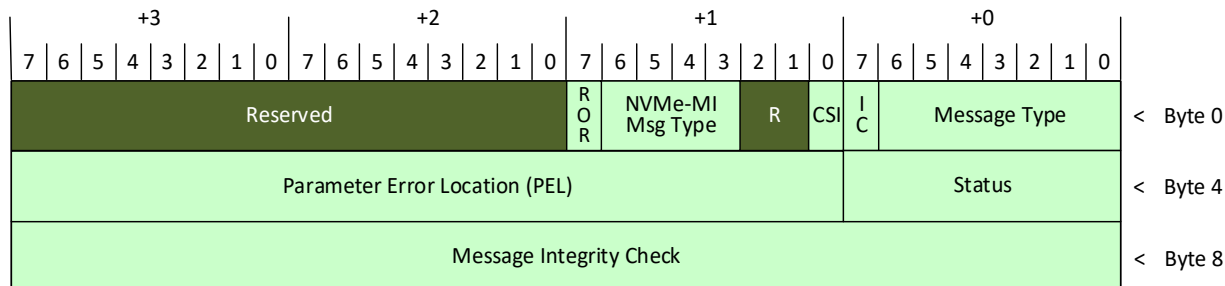


Figure 29: Invalid Parameter Error Response Fields

Byte	Description								
7:5	<p>Parameter Error Location (PEL): This field indicates the byte and bit of the request parameter within the Request Message that contains the first invalid parameter (i.e., the invalid parameter with the lowest byte and bit).</p> <p>If the invalid parameter spans multiple bytes or bits, then the location indicates the first byte and bit of the parameter.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>23:08</td> <td>Byte in the Request Message of the parameter that contained the error. If the error is beyond byte 65,535, then the value 65,535 is reported in this field.</td> </tr> <tr> <td>07:03</td> <td>Reserved</td> </tr> <tr> <td>02:00</td> <td>Bit in the Request Message of the parameter that contained the error. Valid values are 0 to 7.</td> </tr> </tbody> </table>	Bits	Description	23:08	Byte in the Request Message of the parameter that contained the error. If the error is beyond byte 65,535, then the value 65,535 is reported in this field.	07:03	Reserved	02:00	Bit in the Request Message of the parameter that contained the error. Valid values are 0 to 7.
Bits	Description								
23:08	Byte in the Request Message of the parameter that contained the error. If the error is beyond byte 65,535, then the value 65,535 is reported in this field.								
07:03	Reserved								
02:00	Bit in the Request Message of the parameter that contained the error. Valid values are 0 to 7.								

4.2 Out-of-Band Message Servicing Model

The out-of-band mechanism in this specification utilizes a request and response servicing model. A Management Controller sends a Request Message to a Management Endpoint, the Management Endpoint processes the Request Message, and when processing has completed, sends a Response Message back the Management Controller. Under no circumstances does a Management Endpoint generate an unsolicited Response Message (i.e., a Response Message that does not correspond to a previously received Request Message).

Unlike other NVMe-MI Messages that may span multiple MCTP packets, NVMe-MI Messages containing a Control Primitive shall consist of exactly one MCTP packet.

This specification utilizes Command Slots for Command Message servicing. A Management Controller should not send a new Command Message to a Command Slot until the Response Message for the previously issued Command Message to that Command Slot has been received. Each Management Endpoint contains two Command Slots that each include state information and a Pause flag (refer to section 4.2.1.4).

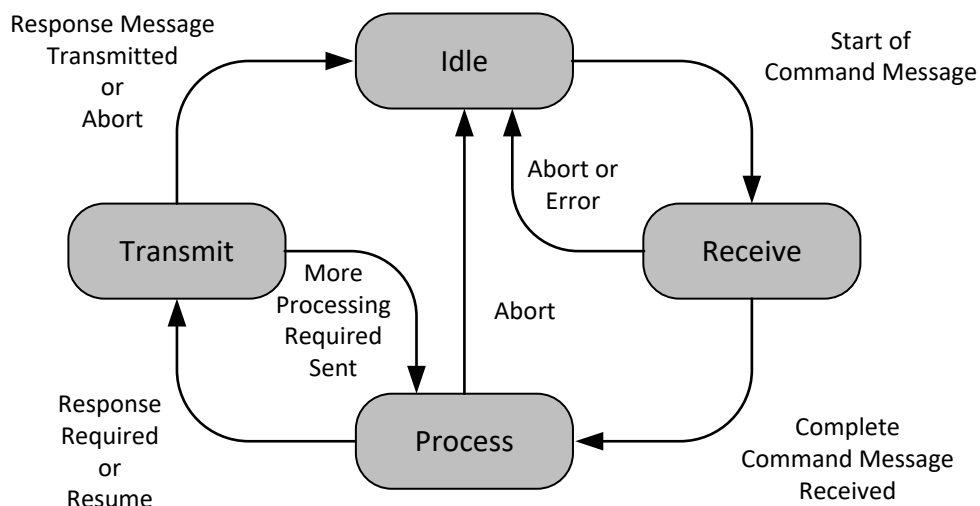
A Management Controller sends a Command Message to a Management Endpoint that targets a specific Command Slot in the Management Endpoint. The Management Endpoint assembles MCTP packets into Command Messages targeting a Command Slot. The Command Slot remains allocated to the Command Message until servicing of the Command Message has completed and command servicing transitions back to the Idle state.

A Command Message is the only type of multi-packet NVMe-MI Message that may be received by a Management Endpoint. The maximum number of Command Messages in flight to a Management Endpoint

is equal to the number of Command Slots. The operation of each Command Slot is independent, allowing a Management Controller to have two independent streams of Command Messages to a Management Endpoint. The Command Message associated with each Command Slot is serviced in parallel. If the NVM Subsystem implements multiple Management Endpoints, then command servicing of each Management Endpoint occurs in parallel. An NVM Subsystem that implements N Management Endpoints may have up to $2N$ Command Messages serviced in parallel.

The Command Servicing State Diagram in Figure 30 is used to describe functional requirements and does not mandate an implementation.

Figure 30: Command Servicing State Diagram



1. **Idle:** This is the default state of the command servicing state machine (e.g., following a reset). Command servicing transitions from Idle to the Receive state when the first MCTP packet of a Command Message is received (i.e., an MCTP packet with the SOM bit in the MCTP packet header set to '1' and the Message Type set to 4h).
2. **Receive:** The state when the first packet of a Command Message has been received and the message is being assembled and/or validated. Command servicing transitions from Receive to the Idle state when an Abort Control Primitive is received, an error is detected in message assembly (refer to section 3.2.1.1), or the Message Integrity Check fails (refer to section 3.1.1.1). Command servicing transitions from Receive state to the Process state when a Command Message is assembled and the message integrity check is successful.
3. **Process:** The state when a Command Message is processed. Processing of a Command Message consists of checking for errors with the Command Message and performing the actions specified by the Command Message or aborting the Command Message. Command servicing transitions from Process to the Transmit state when a Response Message is required to be sent (i.e., the Pause Flag is cleared to '0' and either of the following are true: all processing of the Command Message has completed or command processing is expected to exceed the corresponding transport binding specification response timeout). Command servicing transitions from the Process state to the Idle state due to an Abort Control Primitive (refer to section 4.2.1.3).

4. **Transmit:** The state in which a Response Message for the Command Message is transmitted to the Management Controller. Command servicing transitions from the Transmit to the Idle state once the entire NVMe-MI Message associated with the response to the Command Message has been transmitted on the physical medium or due to an Abort Control Primitive (refer to section 4.2.1.3). If command servicing did not complete in the Process state, then the Management Endpoint transmits a Response Message with status More Processing Required and the command servicing transitions back to the Process state.

The behavior of receiving two or more overlapping Command Messages to the same Command Slot is undefined. If this results in the Management Endpoint discarding a Command Message, then this is considered receiving a Command Message to a non-Idle Command Slot (CMNICS). Refer to section 4.2.1.4.

4.2.1 Control Primitives

Control Primitives are Request Messages sent from a Management Controller to a Management Endpoint to affect the servicing of a previously issued Command Message or get the state of a Command Slot and Management Endpoint. Control Primitives are applicable only in the out-of-band mechanism and are prohibited in the in-band tunneling mechanism.

Control Primitives may target a Command Slot. Unlike Command Messages, Control Primitives may be sent while the Command Slot is in any command servicing state and are processed immediately by the Management Endpoint. Unless otherwise indicated, Control Primitives do not change the command servicing state of the Command Slot.

The format of a Control Primitive is shown in Figure 31 and the fields are described in Figure 32.

Figure 31: Control Primitive Request Message Format

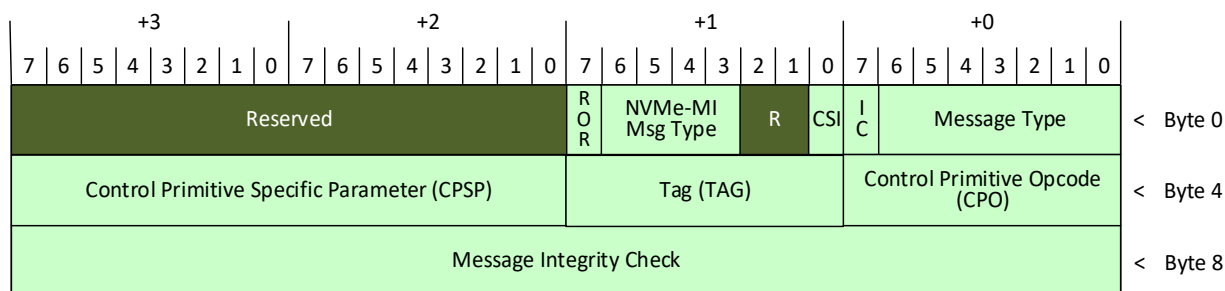


Figure 32: Control Primitive Fields

Byte	Description
03:00	NVMe-MI Message Header (NMH): Refer to section 3.1.
04	Control Primitive Opcode (CPO): This field specifies the opcode of the Control Primitive to be executed. Refer to Figure 33.
05	Tag (TAG): This field contains an opaque value that is sent from the Management Controller in the Control Primitive and returned by the Management Endpoint in to the associated Response Message. A Management Controller may use any value in this field.
07:06	Control Primitive Specific Parameter (CPSP): This field is used to pass Control Primitive specific parameter information.
11:08	Message Integrity Check (MIC): Refer to section 3.1.

Figure 33: Control Primitive Opcodes

Opcode	O/M ¹	Command
00h	M	Pause
01h	M	Resume
02h	M	Abort
03h	M	Get State
04h	M	Replay
05h to EFh		Reserved
F0h to FFh	O	Vendor Specific
NOTES:		
1. O/M: O = Optional, M = Mandatory.		

The format of a Success Response associated with a Control Primitive is shown in Figure 34 and the fields are described in Figure 35.

Figure 34: Control Primitive Success Response Format

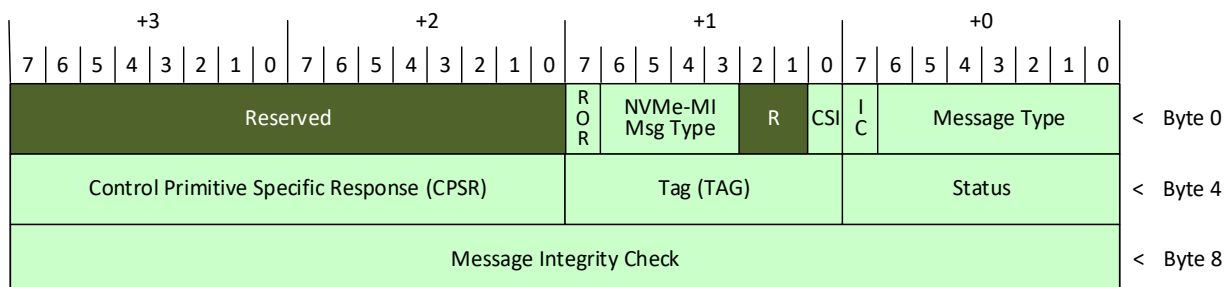


Figure 35: Control Primitive Success Response Fields

Byte	Description
03:00	NVMe-MI Message Header (NMH): Refer to section 3.1.
04	Status (STATUS): Refer to section 4.1.2.
05	Tag (TAG): This field contains an opaque value that is passed by the Management Endpoint from the Control Primitive to the associated Response Message. The Response Message contains the same value in this field as the corresponding Request Message.
07:06	Control Primitive Specific Response (CPSR): This field is used to return Control Primitive specific status.
11:08	Message Integrity Check (MIC): Refer to section 3.1.

A Management Endpoint transmits a Response Message to the Management Controller when the actions associated with that Control Primitive have completed.

Unlike Command Messages, a Management Controller may issue a Control Primitive to a Command Slot without waiting for a response for previously issued Control Primitives to that Command Slot. If multiple Control Primitives are sent without waiting for responses from the Management Endpoint, only the actions and response associated with the last Control Primitive are guaranteed (i.e., the actions associated with previously issued but unacknowledged Control Primitives may or may not be performed and the Response

Messages for previously issued but unacknowledged Control Primitives may or may not be transmitted). Receipt of a Control Primitive never corrupts a previous Control Primitive associated with the Command Slot. The Response Message is either entirely transmitted or discarded.

The TAG field is an opaque value copied from the Control Primitive Request Message into the Response Message. By using unique TAG values, it is possible for the Management Controller to link Response Messages with Request Messages.

4.2.1.1 Pause

The Pause Control Primitive is used to suspend response transmission and suspend the timeout waiting for packet for both Command Slots in a Management Endpoint. The CSI bit in a Pause Control Primitive is not used and shall be cleared to 0h.

Associated with each Command Slot is a Pause Flag that determines whether the slot is 'paused.' The Pause Flag status is included with a Success Response and may also be read using the Get State primitive.

The CPSP field for the Pause primitive is reserved.

The format of the CPSP field in the Control Primitive Success Response is shown in Figure 36.

Figure 36: Pause Control Primitive Success Response Fields

Byte	Description								
07:06	Control Primitive Specific Response (CPSR): This field is used to return Control Primitive specific status.								
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:02</td> <td>Reserved</td> </tr> <tr> <td>01</td> <td>Pause Flag Status Slot 1 (PFSS1): This bit indicates whether or not Command Slot 1 is paused after completing the Pause primitive. This bit set to '1' indicates the Command Slot is paused. This bit cleared to '0' indicates the Command Slot is not paused.</td> </tr> <tr> <td>00</td> <td>Pause Flag Status Slot 0 (PFSS0): This bit indicates whether or not Command Slot 0 is paused after completing the Pause primitive. This bit set to '1' indicates the Command Slot is paused. This bit cleared to '0' indicates the Command Slot is not paused.</td> </tr> </tbody> </table>	Bits	Description	15:02	Reserved	01	Pause Flag Status Slot 1 (PFSS1): This bit indicates whether or not Command Slot 1 is paused after completing the Pause primitive. This bit set to '1' indicates the Command Slot is paused. This bit cleared to '0' indicates the Command Slot is not paused.	00	Pause Flag Status Slot 0 (PFSS0): This bit indicates whether or not Command Slot 0 is paused after completing the Pause primitive. This bit set to '1' indicates the Command Slot is paused. This bit cleared to '0' indicates the Command Slot is not paused.
	Bits	Description							
	15:02	Reserved							
01	Pause Flag Status Slot 1 (PFSS1): This bit indicates whether or not Command Slot 1 is paused after completing the Pause primitive. This bit set to '1' indicates the Command Slot is paused. This bit cleared to '0' indicates the Command Slot is not paused.								
00	Pause Flag Status Slot 0 (PFSS0): This bit indicates whether or not Command Slot 0 is paused after completing the Pause primitive. This bit set to '1' indicates the Command Slot is paused. This bit cleared to '0' indicates the Command Slot is not paused.								

The result of a Pause Control Primitive on a Command Slot is dependent on the command servicing state of the Command Slot when the Pause Control Primitive is received, as described below:

Idle: The Pause primitive has no effect, and the Pause Flag is not changed (i.e., remains cleared to '0'). Refer to section 4.2.1.4.

Receive: The Pause primitive sets the Pause Flag to '1' (refer to section 4.2.1.4) and alerts the Management Endpoint that remaining MCTP packets associated with the command may be delayed. Further packets sent to this Command Slot while the Pause Flag is set are received normally.

Process: The Pause primitive sets the Pause Flag to '1' (refer to section 4.2.1.4) causing the Command Slot to remain in the Process state until a Resume Control Primitive is received. Pause has no effect on the command processing in the Command Slot. Though command processing may complete, the Command Slot shall not transition to the Transmit state.

Transmit: The Pause primitive sets the Pause Flag to '1' (refer to section 4.2.1.4) suspending transmission of MCTP response packets associated on a packet boundary with the Command in the Command Slot.

The Management Endpoint shall transmit a Response Message with success status after receiving the Pause primitive. It is not an error to issue a Pause Control Primitive when a Command Slot is already paused.

While the Pause Flag is set to '1', the Management Endpoint disables the timeout waiting for packet timer and does not transmit responses to commands. The timeout waiting for a packet is the lesser of 100 ms or the time defined in the appropriate MCTP transport binding specification. The Management Controller should not send commands while a Management Endpoint is paused.

4.2.1.2 Resume

The Resume Control Primitive is used to resume from a paused condition. This is the complement to the Pause Control Primitive.

Like the Pause Control Primitive, the Resume Control Primitive affects both slots and the CSI bit in a Resume Control Primitive shall be cleared to '0'. If a Command Slot was not paused before receiving the Resume primitive, the Resume primitive completes successfully and has no effect.

The CPSP field for the Resume primitive is reserved. The CPSR field in the Control Primitive Success Response is reserved.

The result of a Resume Control Primitive is based on the state of a Command Slot when the Resume Control Primitive is received, as described below:

Idle: The Resume primitive has no effect.

Receive: The Resume primitive alerts the Management Endpoint that transmission of any remaining MCTP packets associated with the command is resuming. The Pause Flag is cleared to '0' (refer to section 4.2.1.4).

Process: The Resume primitive allows a previously paused Command Slot to transition to the Transmit state once processing is complete and starts transmitting a response after responding to the Resume primitive. The Pause Flag is cleared to '0' (refer to section 4.2.1.4).

Transmit: The Management Endpoint resumes transmission of the response corresponding to the command associated with that slot after responding to the Resume primitive. The Pause Flag is cleared to '0' (refer to section 4.2.1.4).

The Management Endpoint shall transmit a Control Primitive Response Message with success status after receiving the Resume primitive.

4.2.1.3 Abort

The Abort Control Primitive is used to re-initialize a Command Slot to the Idle state, clear the Pause Flag associated with that Command Slot to '0', and attempt to abort command servicing associated with that Command Slot.

Aborting a Command Message shall have no effect on the other Command Slot of the Management Endpoint, other Management Endpoints, or NVMe Controllers in the NVM Subsystem. Subsequent command servicing in the Command Slot is not affected by the Abort.

A Management Controller may issue an Abort primitive to clean-up resources associated with a Command Slot in an unknown state.

The CPSP field for the Abort primitive is reserved. The format of the CPSR field in the Control Primitive Success Response is shown in Figure 37.

Figure 37: Abort Control Primitive Success Response Fields

Byte	Description				
07:06	Control Primitive Specific Response (CPSR): This field is used to return Control Primitive specific status.				
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:02</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	15:02	Reserved
	Bits	Description			
15:02	Reserved				
01:00	Command Processing Abort Status (CPAS): This field indicates the effect of the Abort primitive on the processing of the Command Message associated with the Command Slot. 0h – Command aborted after processing completed or no command to abort. 1h – Command aborted before processing began. 2h – Command processing partially completed. 3h – Reserved				

The result of an Abort primitive is based on the command servicing state of the specified Command Slot when the Abort primitive is received, as described below:

Idle: The Abort primitive has no effect. The Management Endpoint shall transmit a Response Message with success status and the CPAS field cleared to 0h.

Receive: The Management Endpoint discards the contents of the Command Slot and transitions to the Idle state. The Management Endpoint shall transmit a Response Message with success status and the CPAS field set to 1h.

Process: The Abort primitive causes processing of the command in the Command Slot to be aborted:

- a) If the Abort primitive was received before command processing started, the Management Endpoint discards the contents of the Command Slot and transitions to the Idle state. The Management Endpoint shall transmit a Success Response and the CPAS field set to 1h; or
- b) If the Abort primitive was received while the command is being processed, the Management Endpoint discards the contents of the Command Slot and transitions to the Idle state. The Management Endpoint attempts to abort the command:
 - If the command is aborted and had no effect on the NVM Subsystem, then the Management Endpoint shall transmit a Success Response and the CPAS field set to 1h;
 - If the Management Endpoint is not able to abort the command, then the Management Endpoint shall transmit a Success Response and set the CPAS field to 2h; or
 - If the command has completed processing (e.g., the Management Endpoint is paused), then the Management Endpoint shall transmit a Success Response and the CPAS field is cleared to 0h.

Transmit: The Management Endpoint discards the contents of the Command Slot and transitions to the Idle state. The Management Endpoint transmits a Response Message with success status and the CPAS field cleared to 0h.

It is not an error to issue an Abort Control Primitive to a slot that is paused. The state of slot is reinitialized clearing the Pause Flag to '0'.

4.2.1.4 Get State

The Get State Control Primitive is used to get the state of a Command Slot and Management Endpoint.

The format of the CPSP field in the Control Primitive Request Message is shown in Figure 38.

Figure 38: Get State Control Primitive Request Message Fields

Byte	Description						
07:06	Control Primitive Specific Parameter (CPSP): This field is used to pass Control Primitive specific parameter information.						
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:01</td> <td>Reserved</td> </tr> <tr> <td>00</td> <td>Clear Error State Flags (CESF): This bit specifies whether or not to clear the error state flags when completing this command.</td> </tr> </tbody> </table>	Bits	Description	15:01	Reserved	00	Clear Error State Flags (CESF): This bit specifies whether or not to clear the error state flags when completing this command.
	Bits	Description					
15:01	Reserved						
00	Clear Error State Flags (CESF): This bit specifies whether or not to clear the error state flags when completing this command.						

The Management Endpoint shall transmit a Response Message with success status after receiving the Get State primitive. The format of the CPSP field in the Control Primitive Success Response is shown in Figure 39.

Figure 39: Get State Control Primitive Success Response Fields

Byte	Description																											
07:06	Control Primitive Specific Response (CPSR): This field is used to return Control Primitive specific status.																											
	<table border="1"> <thead> <tr> <th>Bits</th> <th>CS Specific ¹</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15</td> <td>Yes</td> <td>Pause Flag (PFLG): This bit indicates whether or not the Command Slot is paused. This bit set to '1' indicates the Command Slot is paused. This bit cleared to '0' indicates the Command Slot is not paused. While the Pause Flag is set, the Management Endpoint disables the timeout waiting for packet timer, as defined in the MCTP Base Specification, for the Command Slot and does not transmit responses to Command Messages.</td> </tr> <tr> <td>14</td> <td>No</td> <td>NVM Subsystem Reset Occurred (NSSRO): This bit indicates when an NVM Subsystem Reset occurs while main power is applied. This bit is set to '1' if the last occurrence of an NVM Subsystem Reset occurred while main power was applied to the NVM Subsystem. This bit is cleared to '0' following a power cycle and following a Get State primitive with the CESF bit set to '1'.</td> </tr> <tr> <td>13</td> <td>No</td> <td>Bad Packet or Other Physical Layer (BPOPL): This bit is set to '1' if a packet sent to the Management Endpoint failed a transport specific packet integrity check since the last time Get State primitive was executed with the CESF bit set to '1'.</td> </tr> <tr> <td>12</td> <td>No</td> <td>Bad, Unexpected, or Expired Message Tag (BUEMT): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.</td> </tr> <tr> <td>11</td> <td>No</td> <td>Out-of-Sequence Packet Sequence Number (OSPSN): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.</td> </tr> <tr> <td>10</td> <td>No</td> <td>Unexpected Middle or End of Packet (UMEP): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.</td> </tr> <tr> <td>09</td> <td>No</td> <td>Incorrect Transmission Unit (ITU): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.</td> </tr> <tr> <td>08</td> <td>No</td> <td>Unknown Destination ID (UDSTID): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.</td> </tr> </tbody> </table>	Bits	CS Specific ¹	Description	15	Yes	Pause Flag (PFLG): This bit indicates whether or not the Command Slot is paused. This bit set to '1' indicates the Command Slot is paused. This bit cleared to '0' indicates the Command Slot is not paused. While the Pause Flag is set, the Management Endpoint disables the timeout waiting for packet timer, as defined in the MCTP Base Specification, for the Command Slot and does not transmit responses to Command Messages.	14	No	NVM Subsystem Reset Occurred (NSSRO): This bit indicates when an NVM Subsystem Reset occurs while main power is applied. This bit is set to '1' if the last occurrence of an NVM Subsystem Reset occurred while main power was applied to the NVM Subsystem. This bit is cleared to '0' following a power cycle and following a Get State primitive with the CESF bit set to '1'.	13	No	Bad Packet or Other Physical Layer (BPOPL): This bit is set to '1' if a packet sent to the Management Endpoint failed a transport specific packet integrity check since the last time Get State primitive was executed with the CESF bit set to '1'.	12	No	Bad, Unexpected, or Expired Message Tag (BUEMT): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.	11	No	Out-of-Sequence Packet Sequence Number (OSPSN): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.	10	No	Unexpected Middle or End of Packet (UMEP): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.	09	No	Incorrect Transmission Unit (ITU): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.	08	No	Unknown Destination ID (UDSTID): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.
	Bits	CS Specific ¹	Description																									
	15	Yes	Pause Flag (PFLG): This bit indicates whether or not the Command Slot is paused. This bit set to '1' indicates the Command Slot is paused. This bit cleared to '0' indicates the Command Slot is not paused. While the Pause Flag is set, the Management Endpoint disables the timeout waiting for packet timer, as defined in the MCTP Base Specification, for the Command Slot and does not transmit responses to Command Messages.																									
	14	No	NVM Subsystem Reset Occurred (NSSRO): This bit indicates when an NVM Subsystem Reset occurs while main power is applied. This bit is set to '1' if the last occurrence of an NVM Subsystem Reset occurred while main power was applied to the NVM Subsystem. This bit is cleared to '0' following a power cycle and following a Get State primitive with the CESF bit set to '1'.																									
	13	No	Bad Packet or Other Physical Layer (BPOPL): This bit is set to '1' if a packet sent to the Management Endpoint failed a transport specific packet integrity check since the last time Get State primitive was executed with the CESF bit set to '1'.																									
	12	No	Bad, Unexpected, or Expired Message Tag (BUEMT): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.																									
	11	No	Out-of-Sequence Packet Sequence Number (OSPSN): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.																									
	10	No	Unexpected Middle or End of Packet (UMEP): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.																									
	09	No	Incorrect Transmission Unit (ITU): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.																									
08	No	Unknown Destination ID (UDSTID): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.																										

Figure 39: Get State Control Primitive Success Response Fields

Byte	Description											
07	No	Bad Header Version (BHVS): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.										
06	No	Unsupported Transmission Unit (UTUNT): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.										
05	No	Timeout Waiting for a Packet (WPTT): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.										
04	No	Bad Message Integrity Check Error (BMICE): This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State primitive was executed with the CESF bit set to '1'.										
03	No	Command Message to non-Idle Command Slot (CMNICS): This bit is set to '1' if the Management Endpoint discarded one or more Command Messages due to overlapping Command Messages to a Command Slot since the last time Get State primitive was executed with the CESF bit set to '1'.										
02		Reserved										
01:00	Yes	<p>Slot Command Servicing State (SSTA): This field indicates the current command servicing state of the Command Slot. An implementation may choose to indicate only the Idle and Process states in this field. Refer to Figure 30.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Idle</td> </tr> <tr> <td>1h</td> <td>Receive</td> </tr> <tr> <td>2h</td> <td>Process</td> </tr> <tr> <td>3h</td> <td>Transmit</td> </tr> </tbody> </table>	Value	Description	0h	Idle	1h	Receive	2h	Process	3h	Transmit
Value	Description											
0h	Idle											
1h	Receive											
2h	Process											
3h	Transmit											
<p>NOTES:</p> <p>1. Command Slot Specific. A 'Yes' in this column indicates the value of the field is independent per Command Slot within a Management Endpoint.</p>												

4.2.1.5 Replay

The Replay Control Primitive is used to retransmit the Response Message for the last Command Message processed in a Command Slot. The replayed Response Message forms a new MCTP Response Message with Message Data starting from Response Replay Offset of the original Response Message and continuing to the end of the Response Message, including the original MIC. The first packet shall have SOM set to '1' and shall include the Message Header of the original Response Message even if the Response Replay Offset is not zero.

Note that the Management Controller will need extensions to the MCTP Base Specification in its MCTP layer in order to Replay a Response Message using a non-zero Response Replay Offset. No extensions to the MCTP Base Specification are needed to Replay with Response Replay Offset equal to zero. For the case where a Management Controller chooses to use a non-zero Response Replay Offset, the MCTP Base Specification requires terminating message assembly for certain errors (i.e., receiving a packet with bad packet data integrity).

If a Management Controller receives a number of packets with no errors in a Response Message and then gets an error on a packet that causes termination of message assembly, the Management Controller will need extensions in its MCTP layer to forward the packets it received with no errors to its NVMe-MI layer prior to terminating message assembly. The Management Controller can then issue a Replay to get the second part of the Response Message using a non-zero Response Replay Offset. The Management Controller's NVMe-MI layer can then assemble the two partial Response Messages to create the whole

Response Message. The MIC can then be validated across the whole Response Message as described in Section 3.1.1.1.

The format of the CPSP field in the Control Primitive Request Message is shown in Figure 40.

Figure 40: Replay Control Primitive Request Fields

Byte	Description						
07:06	Control Primitive Specific Parameter (CPSP): This field is used to pass Control Primitive specific parameter information.						
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:08</td> <td>Reserved</td> </tr> <tr> <td>07:00</td> <td> <p>Response Replay Offset (RRO): This field specifies the starting packet number from which the Response Message associated with the last Command Message processed in the Command Slot should be replayed.</p> <p>This is a 0's based value. When this field is cleared to 0h, the first packet of the associated Response Message is the first packet replayed.</p> <p>If this field specifies an offset that is beyond the length of the Response Message, then processing of the Control Primitive is aborted and the Management Endpoint transmits an Invalid Parameter Error Response.</p> </td> </tr> </tbody> </table>	Bits	Description	15:08	Reserved	07:00	<p>Response Replay Offset (RRO): This field specifies the starting packet number from which the Response Message associated with the last Command Message processed in the Command Slot should be replayed.</p> <p>This is a 0's based value. When this field is cleared to 0h, the first packet of the associated Response Message is the first packet replayed.</p> <p>If this field specifies an offset that is beyond the length of the Response Message, then processing of the Control Primitive is aborted and the Management Endpoint transmits an Invalid Parameter Error Response.</p>
	Bits	Description					
15:08	Reserved						
07:00	<p>Response Replay Offset (RRO): This field specifies the starting packet number from which the Response Message associated with the last Command Message processed in the Command Slot should be replayed.</p> <p>This is a 0's based value. When this field is cleared to 0h, the first packet of the associated Response Message is the first packet replayed.</p> <p>If this field specifies an offset that is beyond the length of the Response Message, then processing of the Control Primitive is aborted and the Management Endpoint transmits an Invalid Parameter Error Response.</p>						

The format of the CPSR field in the Control Primitive Success Response is shown in Figure 41.

Figure 41: Replay Control Primitive Success Response Fields

Byte	Description						
07:06	Control Primitive Specific Response (CPSR): This field is used to return Control Primitive specific status.						
	<table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:01</td> <td>Reserved</td> </tr> <tr> <td>00</td> <td> <p>Response Replay (RR): This bit indicates if a previous Response Message is retransmitted. This bit is set to '1' if the requested Response Message is retransmitted by the Management Endpoint. This bit is cleared to '0' if the requested Response Message is not retransmitted.</p> </td> </tr> </tbody> </table>	Bit	Description	15:01	Reserved	00	<p>Response Replay (RR): This bit indicates if a previous Response Message is retransmitted. This bit is set to '1' if the requested Response Message is retransmitted by the Management Endpoint. This bit is cleared to '0' if the requested Response Message is not retransmitted.</p>
	Bit	Description					
15:01	Reserved						
00	<p>Response Replay (RR): This bit indicates if a previous Response Message is retransmitted. This bit is set to '1' if the requested Response Message is retransmitted by the Management Endpoint. This bit is cleared to '0' if the requested Response Message is not retransmitted.</p>						

The result of a Replay primitive is based on the command servicing state of the specified Command Slot when the Replay primitive is received, as described below:

Idle: The Replay primitive requests retransmission of the completion at the offset specified by the RRO field if such a completion is available:

- a) If the Replay primitive was received following an Abort primitive or a reset (refer to section 9.3) before any Command Messages are processed, then there is no Response Message available to retransmit. The Management Endpoint shall transmit a Response Message with success status with the RR bit cleared to '0'; or
- b) If the Replay primitive was received following the processing of one or more Command Messages, then the Management Endpoint shall transmit a Response Message with success status with the

RR bit set to '1'. The Management Endpoint transmits the MCTP packets associated with the requested Response Message after the Control Primitive Success Response.

Receive: The Management Endpoint transmits a Response Message with success status with the RR bit cleared to '0'.

Process: The Replay primitive requests retransmission of the last response transmitted for the command in this Command Slot:

- a) If a Response Message has not been transmitted for the Command Message (i.e., the slot never entered the Transmit state for the Command Message), then the Management Endpoint transmits a Response Message with success status and the RR bit cleared to '0'; or
- b) If a Response Message has been transmitted for the Command Message (i.e., a Response Message was transmitted indicating that more processing was required), then the Management Endpoint transmits a Response Message with success status with the RR bit set to '1'. The Management Endpoint retransmits the response indicating that more processing is required.

Transmit: The Management Endpoint stops transmitting response packets for the Command Slot and then transmits a Response Message with success status with the RR bit set to '1'. The Management Endpoint transmits a Response Message containing the packets starting at the packet offset specified in the Response Replay Offset field of the Replay after the Control Primitive Success Response. The Command Slot remains in the Transmit state until retransmission is complete.

It is not an error to issue a Replay primitive to a Command Slot that is paused. The response is retransmitted even if the Command Slot is paused (i.e., there is an implicit Resume primitive affecting both Command Slots when processing the Replay primitive) at any time during the response, including before the first packet was transmitted. After successful completion of the Replay primitive, neither Command Slot is paused.

4.2.2 Out-of-Band Error Handling

This section describes error handling specific to the NVMe-MI out-of-band message processing model.

4.2.2.1 Command Timeouts

MCTP defines a maximum response time for MCTP control messages (refer to the appropriate MCTP transport binding specification).

If a Management Endpoint determines that command processing may not complete within the lesser of 100 ms or the request-to-response time specified in the appropriate MCTP transport binding specification, the Management Endpoint shall utilize the More Processing Required response mechanism. The Response Message from the Management Endpoint is allowed to be delayed beyond this timeout while the transport is busy or unavailable.

A Management Endpoint should only use the More Processing Required response for commands that are expected to take longer than the required time (e.g., Format NVM). Implementations are strongly discouraged from using this response while processing Commands Messages that take less than or the required time.

4.2.2.2 Control Primitive Timeouts

A Management Endpoint shall attempt to respond to a Control Primitive within the lesser of 100 ms or the request-to-response time specified in the appropriate MCTP transport binding specification. The Response Message from the Management Endpoint is allowed to be delayed beyond this timeout while the transport is busy or unavailable.

4.2.3 Management Endpoint Buffer

Since the maximum size of the NVMe-MI Message is 4,224 bytes, the maximum possible amount of out-of-band Request Data that may be contained in a Request Message is 4,216 bytes (i.e., 4,224 bytes minus 4-byte Message Header and 4-byte Message Integrity Check field) and the maximum possible amount of out-of-band Response Data that may be contained in a Response Message is 4,215 bytes (i.e., 4,224 bytes minus 4-byte Message Header, 1-byte Status field, and 4-byte Message Integrity Check field). The amount of supported Request or Response Data is Command Message specific due to the presence of additional command specific fields. In some cases, it is desirable to service Command Messages that contain more Request Data or Response Data than allowed to be transferred in an NVMe-MI Message. For example, one may wish to issue an NVM Express Admin Command Set Get Log Page command to transfer a log page that is greater in size than that allowed in the Response Data.

A Management Endpoint may support an optional Management Endpoint Buffer that facilitates Request Data and Response Data transfers that exceed that maximum size allowed by an NVMe-MI Message. A Management Endpoint Buffer is exclusive to one Management Endpoint and shall not be shared. Support for the Management Endpoint Buffer and its size in bytes is indicated by the Management Endpoint Buffer Size field in the Port Information Data Structure of the port with which the Management Endpoint is associated. Management Endpoints are not required to all have the same Management Endpoint Buffer support. For example, a subset of Management Endpoints may support a Management Endpoint Buffer and the size of each of these Management Endpoint Buffers may be different.

If a Management Endpoint supports a Management Endpoint Buffer, then all Command Messages or a subset of Command Messages supported by the Management Endpoint may support use of the Management Endpoint Buffer. A list of commands that support the use of the Management Endpoint Buffer is contained in the Management Endpoint Buffer Command Support List data structure that is retrieved using the Read NVMe-MI Data Structure command. If a Management Endpoint supports a Management Endpoint Buffer, then the Management Endpoint shall support the Management Endpoint Buffer Read and Management Endpoint Buffer Write commands.

The contents of a Management Endpoint Buffer may be read or written by a Management Controller by issuing Management Endpoint Buffer Read and Management Endpoint Buffer Write commands. The Management Endpoint Buffer is permitted to be read or written in an arbitrary manner. For example, the contents of the Management Endpoint Buffer may be written sequentially using a sequence of Management Endpoint Buffer Write commands or the partial contents of the Management Endpoint Buffer may be written in any order with gaps using these commands. Furthermore, Management Endpoint Buffer Read and Write commands may be interleaved allowing a portion of the Management Endpoint Buffer to be read while another portion of the Management Endpoint Buffer is written.

If the Management Endpoint Buffer (MEB) bit is set to '1' in a Command Message that normally contains Request Data, then no Request Data is transferred in the Command Message itself and the required Request Data is instead transferred from the Management Endpoint Buffer. The Request Data starts at a zero offset from the start of the Management Endpoint Buffer. If the MEB bit is set to '1' in a Command Message that normally contains Request Data, then the Command Message shall contain no Request Data. If the Command Message contains Request Data or is one that does not support Request Data, then the Management Endpoint responds with an Invalid Parameter Error Response. The parameter with the error in this case is the Request Data field.

If the Management Endpoint Buffer (MEB) bit is set to '1' in a Command Message that normally results in Response Data, then no Response Data is transferred in the corresponding Response Message itself and the Response Data is instead transferred to the Management Endpoint Buffer. The Response Data starts at a zero offset from the start of the Management Endpoint Buffer.

The contents of the Management Endpoint Buffer are set to 0h when the corresponding Management Endpoint is reset. The contents of the Management Endpoint Buffer are modified by the Management Endpoint Buffer Write command and by Command Messages that generate Response Data with the MEB

bit set to '1'. When the Management Endpoint Buffer is updated with Response Data, the contents of the Management Endpoint Buffer that are not updated are set to zero (i.e., the Message Data from previous Command Messages is not preserved). The same contents of the Management Endpoint Buffer may be used as Request Data for multiple Command Messages. Similarly, the Management Endpoint Buffer allows the use of Response Data generated by one Command Message to be used as the Request Data for a subsequent Command Message.

Since it is possible to have two out-of-band Command Messages, one associated with each of the two Command Slots, being simultaneously serviced that use the Management Endpoint Buffer, the Management Controller must comprehend and manage any possible race conditions. Updates to the Management Endpoint Buffer are not guaranteed to be atomic. Therefore, when a race condition involving two operations that update the Management Endpoint Buffer occurs, the final contents of the Management Endpoint Buffer may be an arbitrary mixture of the updates.

The Management Endpoint Buffer is considered a cache in the context of sanitize operations performed in an NVM Subsystem. The MCTP Management Endpoint Buffer may contain Response Data associated with a previously executed command that is not allowed during a sanitize operation. When a sanitize operation is initiated, the contents of the Management Endpoint Buffer shall be cleared to 0h. An attempt to access this zeroed data by a Management Endpoint Buffer Read command or any Command Message that uses the Management Endpoint Buffer, then the Management Endpoint responds with a Response Message Status of Management Endpoint Buffer Cleared Due to Sanitize. This Response Message Status is commonly associated with a Management Endpoint Buffer Read command but may be associated with any command that uses the Management Endpoint Buffer as Request Data.

4.3 In-Band Tunneling Message Servicing Model

The in-band tunneling mechanism in this specification utilizes two NVMe Admin Commands (NVMe-MI Send and NVMe-MI Receive). The NVMe-MI Send command is used to tunnel an NVMe-MI Command from host software to an NVMe Controller that transfers data from the host to the NVMe Controller (similar to a write operation) or to instruct the Management Endpoint to perform an action (e.g., to reset the NVM Subsystem using the Reset command). The NVMe-MI Receive command is used to tunnel an NVMe-MI Command from a host to an NVMe Controller that transfers data from the NVMe Controller to the host (similar to a read operation). Figure 57 specifies whether an NVMe-MI Command is tunneled via the NVMe-MI Send command or the NVMe-MI Receive command.

Refer to the NVM Express specification for additional details on the NVMe-MI Send and NVMe-MI Receive commands. Additional details on NVMe-MI Send are in section 4.3.1 and additional details on NVMe-MI Receive are in section 4.3.2.

4.3.1 NVMe-MI Send Command

The NVMe-MI Send command is an NVMe Admin Command as defined by this specification and the NVM Express specification. It is used to tunnel an NVMe-MI Command in-band from host software to an NVMe Controller that transfers data from a host to an NVMe Controller (similar to a write operation) or to instruct the Management Endpoint to perform an action (e.g., to reset the NVM Subsystem using the Reset command). The data being transferred or action to be performed is in one or more of the following locations: Request Data, NVMe Management Dword 0, NVMe Management Dword 1. Figure 57 specifies which NVMe-MI Commands are tunneled via the NVMe-MI Send command.

NVMe-MI Commands may apply to the NVM Subsystem, Controllers, and/or Namespaces. If the tunneled NVMe-MI Command applies to one or more Controllers, then the applicable Controller(s) are specified by fields in the tunneled NVMe-MI Command. Note that unlike some other Admin Commands, the Controller to which the NVMe-MI Send command is issued is not used to determine which Controller the tunneled NVMe-MI Command applies to. If the tunneled NVMe-MI Command applies to one or more Namespaces, then the applicable Namespace(s) are specified by fields in the tunneled NVMe-MI Command. Note that

the Namespace Identifier (NSID) field of the NVMe-MI Send command (bytes 7:4 of the Submission Queue Entry) is not used and should be cleared to 0h by host software.

The mapping of how an NVMe-MI Command is tunneled inside of NVMe-MI Send commands is described in section 4.3.1.1. The NVMe-MI Send command servicing model is described in section 4.3.1.2.

4.3.1.1 NVMe-MI Send Command Request Message to NVMe Admin Command SQE Mapping

In order to tunnel an NVMe-MI Command in-band via NVMe-MI Send, an NVMe-MI Request Message is mapped onto an NVMe Submission Queue Entry (SQE) as shown pictorially in Figure 42 and in table form in Figure 43. An NVMe-MI Response Message is mapped on to an NVMe Completion Queue Entry (CQE) as shown pictorially in Figure 44 and in table form in Figure 45. Refer to the NVM Express specification for details on an NVMe Submission Queue Entry and an NVMe Completion Queue Entry.

Figure 42: NVMe-MI Send Command Request Message to NVMe Admin Command SQE Mapping Diagram

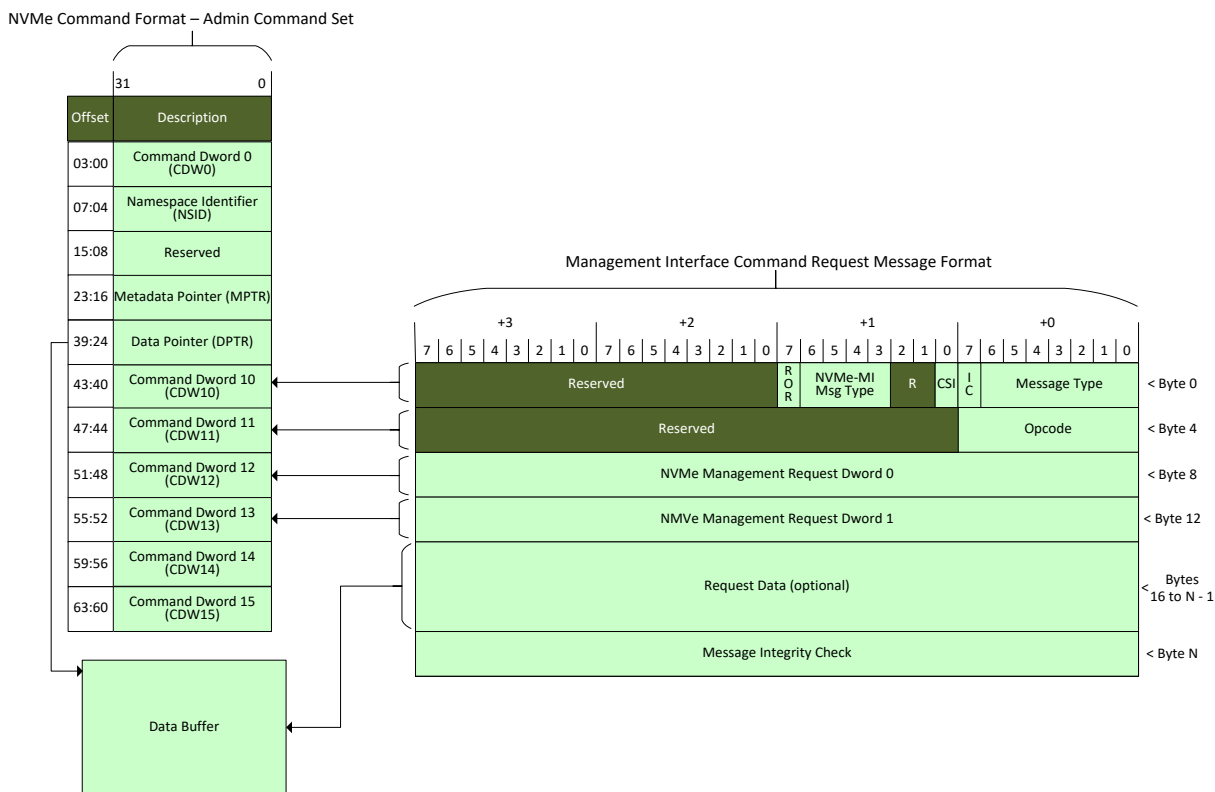


Figure 43: NVMe-MI Send Command Request Message to NVMe Admin Command SQE Mapping Table

NVMe-MI Command Request Message		NVMe Admin Command SQE Mapping	
Byte	Description	Byte	Description
Not applicable (n/a)	This field has no equivalent in this specification.	03:00	Command Dword 0 (CDW0): Refer to the NVM Express specification.

Figure 43: NVMe-MI Send Command Request Message to NVMe Admin Command SQE Mapping Table

NVMe-MI Command Request Message		NVMe Admin Command SQE Mapping	
Byte	Description	Byte	Description
n/a	If the tunneled NVMe-MI Command requires one or more Namespaces to be specified, then the applicable Namespace Identifiers are specified by the tunneled NVMe-MI Command.	07:04	Namespace Identifier (NSID): This field should be cleared to 0h by host software. Refer to the NVM Express specification for more details.
n/a	These bytes have no equivalent in this specification.	23:08	Refer to the NVM Express specification.
n/a	There is no equivalent of DPTR in this specification. In NVMe-MI Send, the Request Data is included in the Request Data portion of the Request Message.	39:24	Data Pointer (DPTR): This field contains a pointer to the start of the data buffer that contains the Request Data portion of the NVMe-MI Command that is being tunneled. If there is no Request Data for this command, then this field is ignored. Refer to the NVM Express specification for the definition of this field.
03:00	NVMe-MI Message Header (NMH)	43:40	Command Dword 10 (CDW10): Dword 0 of the Request Message (NMH) that is being tunneled maps to CDW10 of the SQE. The byte ordering within CDW10 is little endian (i.e., NMH byte 0 maps to CDW10 byte 0, NMH byte 1 maps to CDW10 byte 1, etc.).
04	Opcode (OPC)	47:44	Command Dword 11 (CDW11): Dword 1 of the Request Message (OPC and Reserved bytes 7:5) that is being tunneled maps to CDW11 of the SQE. The byte ordering within CDW11 is little endian (i.e., OPC maps to CDW11 byte 0, the LSB of the Reserved field (NVMe-MI Command Request Message byte 5) maps to CDW11 byte 1, etc.).
07:05	Reserved		
11:08	NVMe Management Dword 0 (NMD0)	51:48	Command Dword 12 (CDW12): Dword 2 of the Request Message (NMD0) that is being tunneled maps to CDW12 of the SQE. The byte ordering within CDW12 is little endian (i.e., NMD0 byte 0 maps to CDW12 byte 0, NMD0 byte 1 maps to CDW12 byte 1).
15:12	NVMe Management Dword 1 (NMD1)	55:52	Command Dword 13 (CDW13): Dword 3 of the Request Message (NMD1) that is being tunneled maps to CDW13 of the SQE. The byte ordering within CDW13 is little endian (i.e., NMD1 byte 0 maps to CDW13 byte 0, NMD1 byte 1 maps to CDW13 byte 1).
n/a	This field has no equivalent in this specification.	59:56	Command Dword 14 (CDW14): Reserved.
n/a	This field has no equivalent in this specification.	63:60	Command Dword 15 (CDW15): Reserved.
N-1:16	Request Data (REQD)	n/a	Request Data is placed by host software into the data buffer pointed to by DPTR. If the Request Data is not dword granular, then the Request Data shall be padded with the minimum number of bytes of zeroes to make the Request Data dword granular. The byte ordering within the data buffer pointed to by DPTR is little endian (i.e., REQD byte 0 maps to byte 0 of the data buffer pointed to by DPTR, REQD byte 1 maps to byte 1 of the data buffer pointed to by DPTR, etc.).

Figure 43: NVMe-MI Send Command Request Message to NVMe Admin Command SQE Mapping Table

NVMe-MI Command Request Message		NVMe Admin Command SQE Mapping	
Byte	Description	Byte	Description
N+3:N	Message Integrity Check (MIC)	n/a	The Message Integrity Check is not used in the in-band tunneling mechanism.

Figure 44: NVMe-MI Send Command Response Message to NVMe Admin Command CQE Mapping Diagram

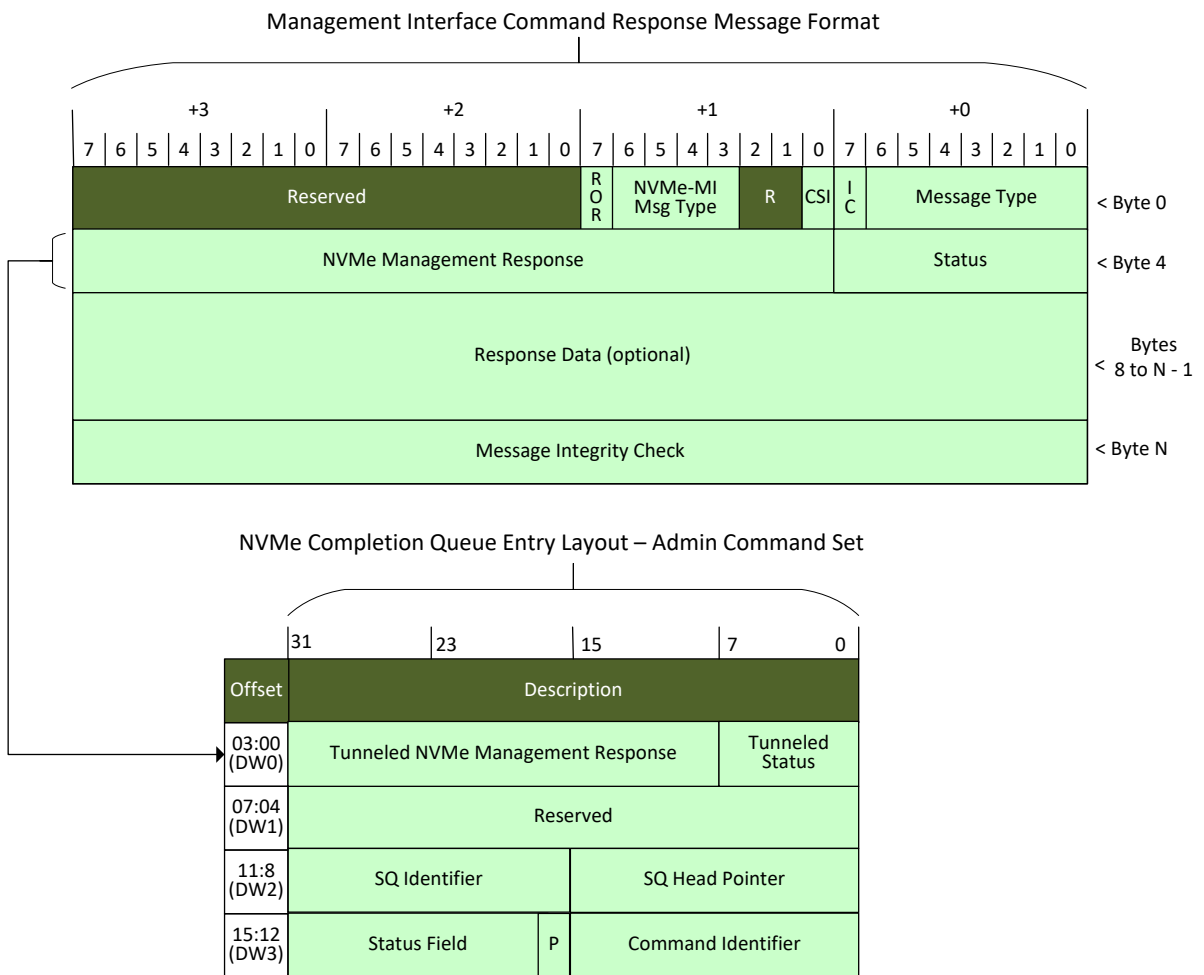


Figure 45: NVMe-MI Send Command Response Message to NVMe Admin Command CQE Mapping Table

NVMe-MI Command Response Message		NVMe Admin Command CQE Mapping	
Byte	Description	Byte	Description
00	MCTP Data (MCTPD)	n/a	This field has no equivalent in the NVMe Admin Command CQE.
01	NVMe-MI Message Parameters (NMP)	n/a	This field has no equivalent in the NVMe Admin Command CQE.
03:02	Reserved	n/a	This field has no equivalent in the NVMe Admin Command CQE.
04	Status (STATUS)	03:00	Command Specific (DW0): Dword 1 of the Response Message (STATUS and NMRESP) that is being tunneled maps to DW0 of the CQE. The byte ordering within DW0 is little endian (i.e., STATUS maps to DW0 byte 0, the LSB of the NMRESP field (NVMe-MI Command Response Message byte 5) maps to DW0 byte 1, etc.). Refer to Figure 46 for additional details on this field.
07:05	NVMe Management Response (NMRESP)		
N-1:8	Response Data (RESPD)	n/a	There is no Response Data for NVMe-MI Send.
N+3:N	Message Integrity Check (MIC)	n/a	The Message Integrity Check is not used in the in-band tunneling mechanism.
n/a	These bytes have no equivalent in this specification.	15:04	Refer to the NVM Express specification.

The definition of Dword 0 of the completion queue entry is in Figure 46.

Figure 46: NVMe-MI Send – Completion Queue Entry Dword 0 (NSCQED0)

Byte	Description
31:08	Tunneled NVMe Management Response (TNMRESP): This field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band. If any errors are detected in the NVMe Context as described in section 4.3.1.2, then this field shall be cleared to 0h.
07:00	Tunneled Status (TSTAT): This field contains the Status field from the NVMe-MI Command that is being tunneled in-band. If any errors are detected in the NVMe Context as described in section 4.3.1.2, then this field shall be cleared to 0h.

4.3.1.2 NVMe-MI Send Command Servicing Model

The NVMe-MI Send command servicing model is illustrated in Figure 47 as a series of phases and NVMe/NVMe-MI Contexts. The phases of the NVMe-MI Send command servicing model are further described in this section. The behavior of the portions of the figure in the NVMe Context are specified by the NVM Express specification. The behavior of the portions of the figure in the NVMe-MI Context are specified by this specification. The phases and NVMe/NVMe-MI Contexts are logical constructs that illustrate the NVMe-MI Send command servicing model and do not mandate a particular implementation.

This section describes the NVMe-MI Send command servicing model starting at NVMe Processing as shown in phase 1 of Figure 47. In phase 1, CDW0 to CDW9 are checked for errors per the NVM Express specification. If any errors are encountered in CDW0 to CDW9, then the NVMe-MI Send command is completed with an error status code in the Status Field as per the NVM Express specification and the Tunneled Status and Tunneled NVMe Management Response fields shall be cleared to 0h.

If there are no errors in CDW0 to CDW9, then command servicing enters phase 2 where the portion of the tunneled NVMe-MI Command in CDW10 to CDW15 is checked for errors. Note that if there is no Request

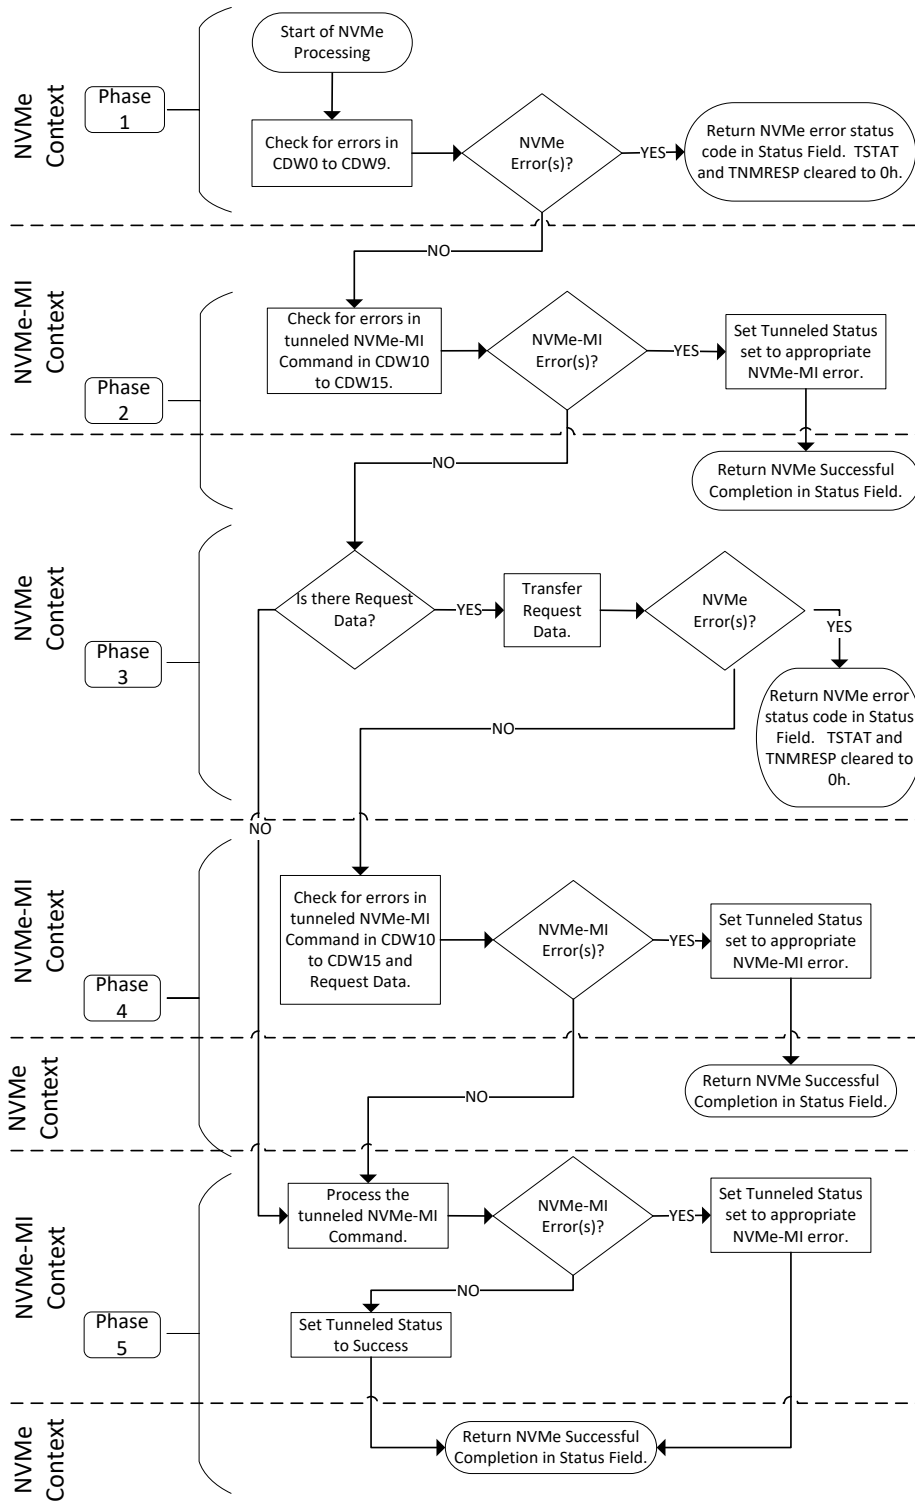
Data, then CDW10 to CDW15 contain the entire tunneled NVMe-MI Command. If any errors are encountered in the portion of the tunneled NVMe-MI Command in CDW10 to CDW15, then the NVMe-MI Send command is completed with a status code of Successful Completion in the Status Field as defined in the NVM Express specification. The Tunneled Status field contains the error Response Message Status for the portion of the tunneled NVMe-MI Command in CDW10 to CDW15 and the Tunneled NVMe Management Response field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band.

If there are no errors in phase 2, then command servicing enters phase 3 where there is a check to determine if there is any Request Data for the tunneled NVMe-MI Command. If there is no Request Data for the tunneled NVMe-MI Command, then command servicing skips to phase 5. If there is Request Data, then the Request Data is transferred from the buffer pointed to by DPTR. If any errors are encountered transferring the Request Data, then the command is completed with an error status code in the Status Field as per the NVM Express specification and the Tunneled Status and Tunneled NVMe Management Response fields shall be cleared to 0h.

If there are no errors transferring the data, then command servicing enters phase 4 where the whole tunneled NVMe-MI Command is constructed from CDW10 to CDW15 and the Request Data that was transferred. If any errors are encountered in the tunneled NVMe-MI Command, then the NVMe-MI Send command is completed with a status code of Successful Completion in the Status Field as defined in the NVM Express specification. The Tunneled Status field contains the appropriate error Response Message Status and the Tunneled NVMe Management Response field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band.

If there are no errors in phase 4, then command servicing enters phase 5 where the tunneled NVMe-MI Command finishes processing. If any errors are encountered processing the tunneled NVMe-MI Command, then the NVMe-MI Send command is completed with a status code of Successful Completion in the Status Field as defined in the NVM Express specification and the Tunneled Status field contains the appropriate error Response Message Status. If the tunneled NVMe-MI Command is processed successfully, then the NVMe-MI Send command is completed with a status code of Successful Completion in the Status Field as defined in the NVM Express specification. The Tunneled Status field contains a Response Message Status of Success for the tunneled NVMe-MI Command and the Tunneled NVMe Management Response field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band.

Figure 47: NVMe-MI Send Command Servicing Model



4.3.2 NVMe-MI Receive Command

The NVMe-MI Receive command is an NVMe Admin Command as defined by this specification and the NVM Express specification. It is used to tunnel an NVMe-MI Command in-band from host software to an NVMe Controller that transfers data from an NVMe Controller to a host (similar to a read operation). The data being transferred is in one or more of the following locations: Response Data, NVMe Management Response. Figure 57 specifies which NVMe-MI Commands are tunneled via the NVMe-MI Receive command.

NVMe-MI Commands may apply to the NVM Subsystem, Controllers, and/or Namespaces. If the tunneled NVMe-MI Command applies to one or more Controllers, then the applicable Controller(s) are specified by fields in the tunneled NVMe-MI Command. Note that unlike some other Admin Commands, the Controller to which the NVMe-MI Receive command is issued is not used to determine which Controller the tunneled NVMe-MI Command applies to. If the tunneled NVMe-MI Command applies to one or more Namespaces, then the applicable Namespace(s) are specified by fields in the tunneled NVMe-MI Command. Note that the Namespace Identifier (NSID) field of the NVMe-MI Receive command (bytes 7:4 of the Submission Queue Entry) is not used and should be cleared to 0h by host software.

The mapping of how an NVMe-MI Command is tunneled inside of an NVMe-MI Receive command is described in section 4.3.2.1. The NVMe-MI Receive command servicing model is described in section 4.3.2.2.

4.3.2.1 NVMe-MI Receive Command Request Message to NVMe Admin Command SQE Mapping

In order to tunnel an NVMe-MI Command in-band via NVMe-MI Receive, an NVMe-MI Request Message is mapped onto an NVMe Submission Queue Entry (SQE) as shown pictorially in Figure 48 and in table form in Figure 49. An NVMe-MI Response Message is mapped on to an NVMe Completion Queue Entry (CQE) as shown pictorially in Figure 48 and in table form in Figure 50. Refer to the NVM Express specification for details on an NVMe Submission Queue Entry and NVMe Completion Queue Entry.

Figure 48: NVMe-MI Receive Command Request/Response Message to NVMe Admin Command SQE/CQE Mapping Diagram

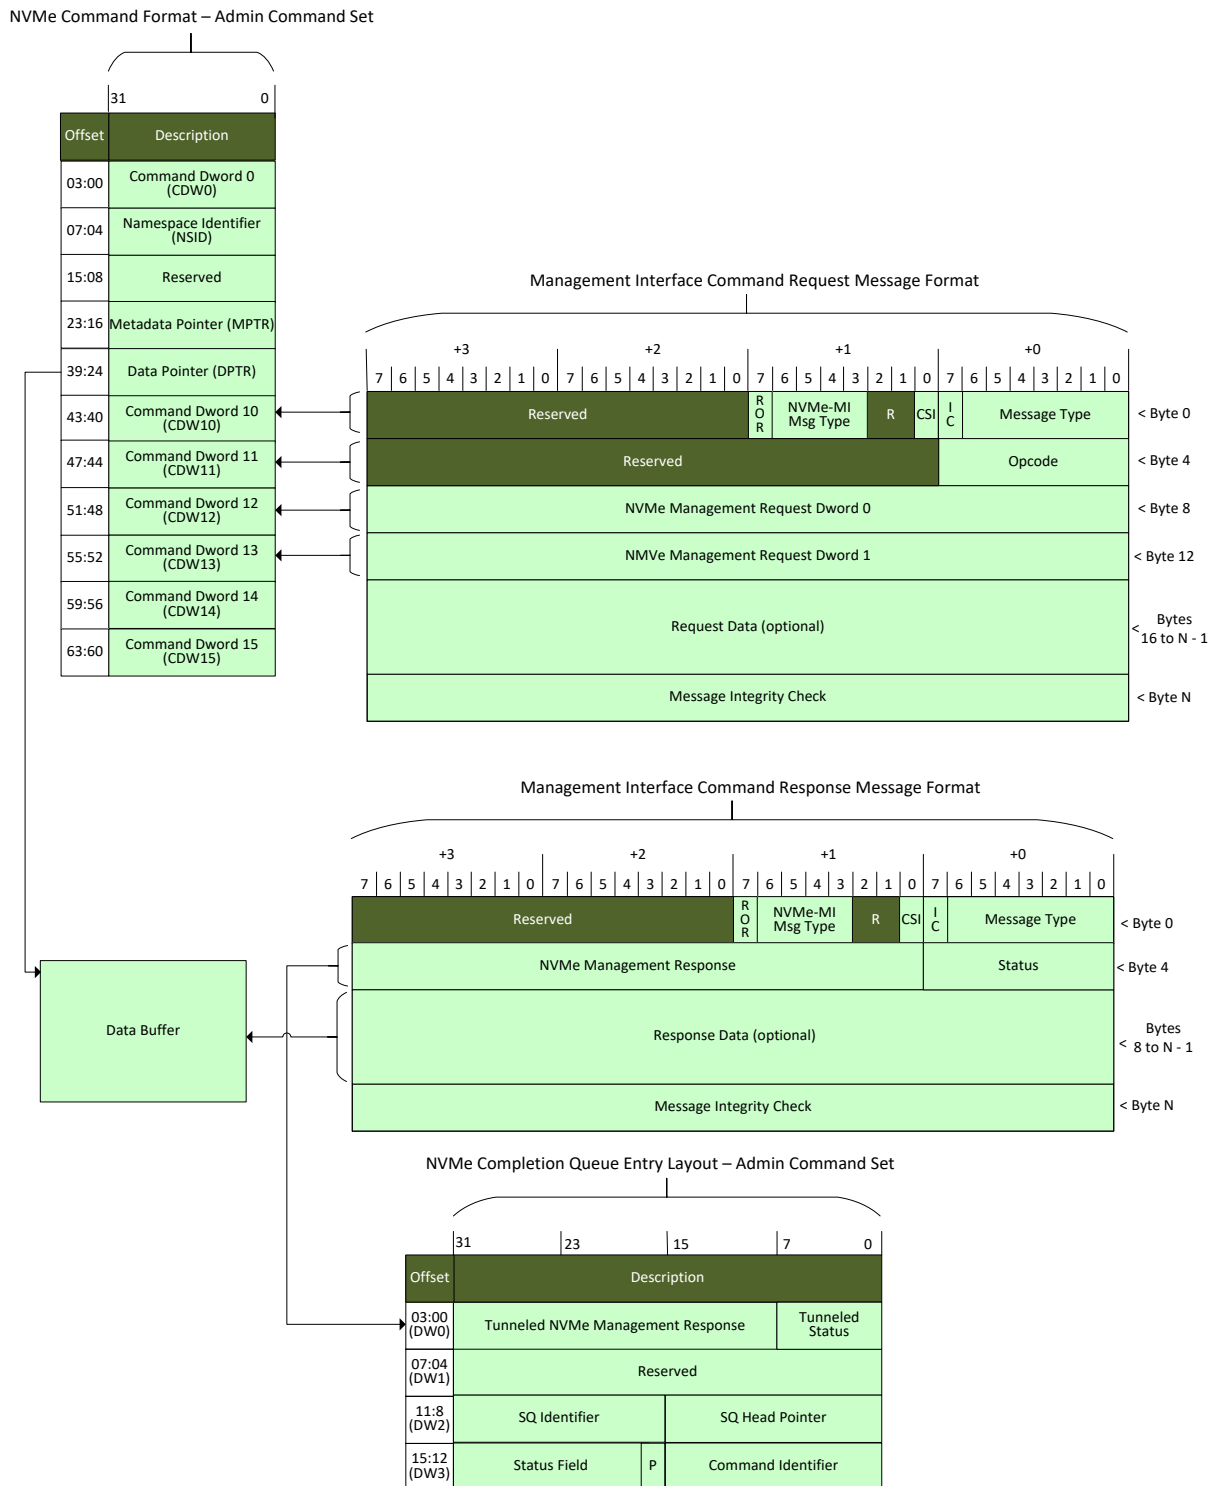


Figure 49: NVMe-MI Receive Command Request/Response Message to NVMe Admin Command SQE/CQE Mapping Table

NVMe-MI Command Request Message		NVMe Admin Command SQE Mapping	
Byte	Description	Byte	Description
n/a	This field has no equivalent in this specification.	03:00	Command Dword 0 (CDW0): Refer to the NVM Express specification.
n/a	If the tunneled NVMe-MI Command requires one or more Namespaces to be specified, then the applicable Namespace Identifiers are specified by the tunneled NVMe-MI Command.	07:04	Namespace Identifier (NSID): This field should be cleared to 0h by host software. Refer to the NVM Express specification for more details.
n/a	These bytes have no equivalent in this specification.	23:08	Refer to the NVM Express specification.
n/a	There is no equivalent of DPTR in this specification. In NVMe-MI Receive, the Response Data is included in the Response Data portion of the Response Message.	39:24	Data Pointer (DPTR): This field contains a pointer to the start of the data buffer that contains the Response Data portion of the NVMe-MI Command that is being tunneled. If there is no Response Data for this command, then this field is ignored. Refer to the NVM Express specification for the definition of this field.
03:00	NVMe-MI Message Header (NMH)	43:40	Command Dword 10 (CDW10): Dword 0 of the Request Message (NMH) that is being tunneled maps to CDW10 of the SQE. The byte ordering within CDW10 is little endian (i.e., NMH byte 0 maps to CDW10 byte 0, NMH byte 1 maps to CDW10 byte 1, etc.).
04	Opcode (OPC)	47:44	Command Dword 11 (CDW11): Dword 1 of the Request Message (OPC and Reserved bytes 7:5) that is being tunneled maps to CDW11 of the SQE. The byte ordering within CDW11 is little endian (i.e., OPC maps to CDW11 byte 0, the LSB of the Reserved field (NVMe-MI Command Request Message byte 5) maps to CDW11 byte 1, etc.).
07:05	Reserved		
11:08	NVMe Management Dword 0 (NMD0)	51:48	Command Dword 12 (CDW12): Dword 2 of the Request Message (NMD0) that is being tunneled maps to CDW12 of the SQE. The byte ordering within CDW12 is little endian (i.e., NMD0 byte 0 maps to CDW12 byte 0, NMD0 byte 1 maps to CDW12 byte 1).
15:12	NVMe Management Dword 1 (NMD1)	55:52	Command Dword 13 (CDW13): Dword 3 of the Request Message (NMD1) that is being tunneled maps to CDW13 of the SQE. The byte ordering within CDW13 is little endian (i.e., NMD1 byte 0 maps to CDW13 byte 0, NMD1 byte 1 maps to CDW13 byte 1).
n/a	This field has no equivalent in this specification.	59:56	Command Dword 14 (CDW14): Reserved.
n/a	This field has no equivalent in this specification.	63:60	Command Dword 15 (CDW15): Reserved.
N-1:16	Request Data (REQD)	n/a	There is no Request Data for NVMe-MI Receive.
N+3:N	Message Integrity Check (MIC)	n/a	The Message Integrity Check is not used in the in-band tunneling mechanism.

Figure 50: NVMe-MI Receive Command Response Message to NVMe Admin Command CQE Mapping Table

NVMe-MI Command Response Message		NVMe Admin Command CQE	
Byte	Description	Byte	Description
00	MCTP Data (MCTPD)	n/a	This field has no equivalent in the NVMe Admin Command CQE.
01	NVMe-MI Message Parameters (NMP)	n/a	This field has no equivalent in the NVMe Admin Command CQE.
03:02	Reserved	n/a	This field has no equivalent in the NVMe Admin Command CQE.
04	Status (STATUS)	03:00	Command Specific (DW0): Dword 1 of the Response Message (STATUS and NMRESP) that is being tunneled maps to DW0 of the CQE. The byte ordering within DW0 is little endian (i.e., STATUS maps to DW0 byte 0, the LSB of the NMRESP field (NVMe-MI Command Response Message byte 5) maps to DW0 byte 1, etc.). Refer to Figure 51 for additional details on this field.
0 7:05	NVMe Management Response (NMRESP)		
N-1:8	Response Data (RESPD)	n/a	Response Data is placed by the NVMe Controller into the data buffer pointed to by DPTR. If the Response Data size is not dword granular, then the Response Data shall be padded with the minimum number of bytes of zeroes to make the Response Data dword granular. The byte ordering within the data buffer pointed to by DPTR is little endian (i.e., RESPD byte 0 maps to byte 0 of the data buffer pointed to by DPTR, RESPD byte 1 maps to byte 1 of the data buffer pointed to by DPTR, etc.).
N+3:N	Message Integrity Check (MIC)	n/a	The Message Integrity Check is not used in the in-band tunneling mechanism.
n/a	These bytes have no equivalent in this specification.	15:04	Refer to the NVM Express specification.

The definition of Dword 0 of the completion queue entry is in Figure 51.

Figure 51: NVMe-MI Receive – Completion Queue Entry Dword 0 (NRCQED0)

Byte	Description
31:08	Tunneled NVMe Management Response (TNMRESP): This field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band. If any errors are detected in the NVMe Context as described in section 4.3.2.2, then this field shall be cleared to 0h.
07:00	Tunneled Status (TSTAT): This field contains the Status field from the NVMe-MI Command that is being tunneled in-band. If any errors are detected in the NVMe Context as described in section 4.3.2.2, then this field shall be cleared to 0h.

4.3.2.2 NVMe-MI Receive Command Servicing Model

The NVMe-MI Receive command servicing model is illustrated in Figure 52 as a series of phases (described in this section) and NVMe/NVMe-MI Contexts. The phases of the NVMe-MI Receive command servicing model are further described in this section. The behavior of the portions of the figure in the NVMe Context are specified by the NVM Express specification. The behavior of the portions of the figure in the NVMe-MI Context are specified by this specification. The phases and NVMe/NVMe-MI Contexts are logical constructs that illustrate the NVMe-MI Receive command servicing model and do not mandate a particular implementation.

This section describes the NVMe-MI Receive command servicing model starting at NVMe Processing as shown in phase 1 of Figure 52. In phase 1, CDW0 to CDW9 are checked for errors per the NVM Express specification. If any errors are encountered in CDW0 to CDW9, then the command is completed with an error status code in the Status Field as per the NVM Express specification and the Tunneled Status and Tunneled NVMe Management Response fields shall be set to 0h.

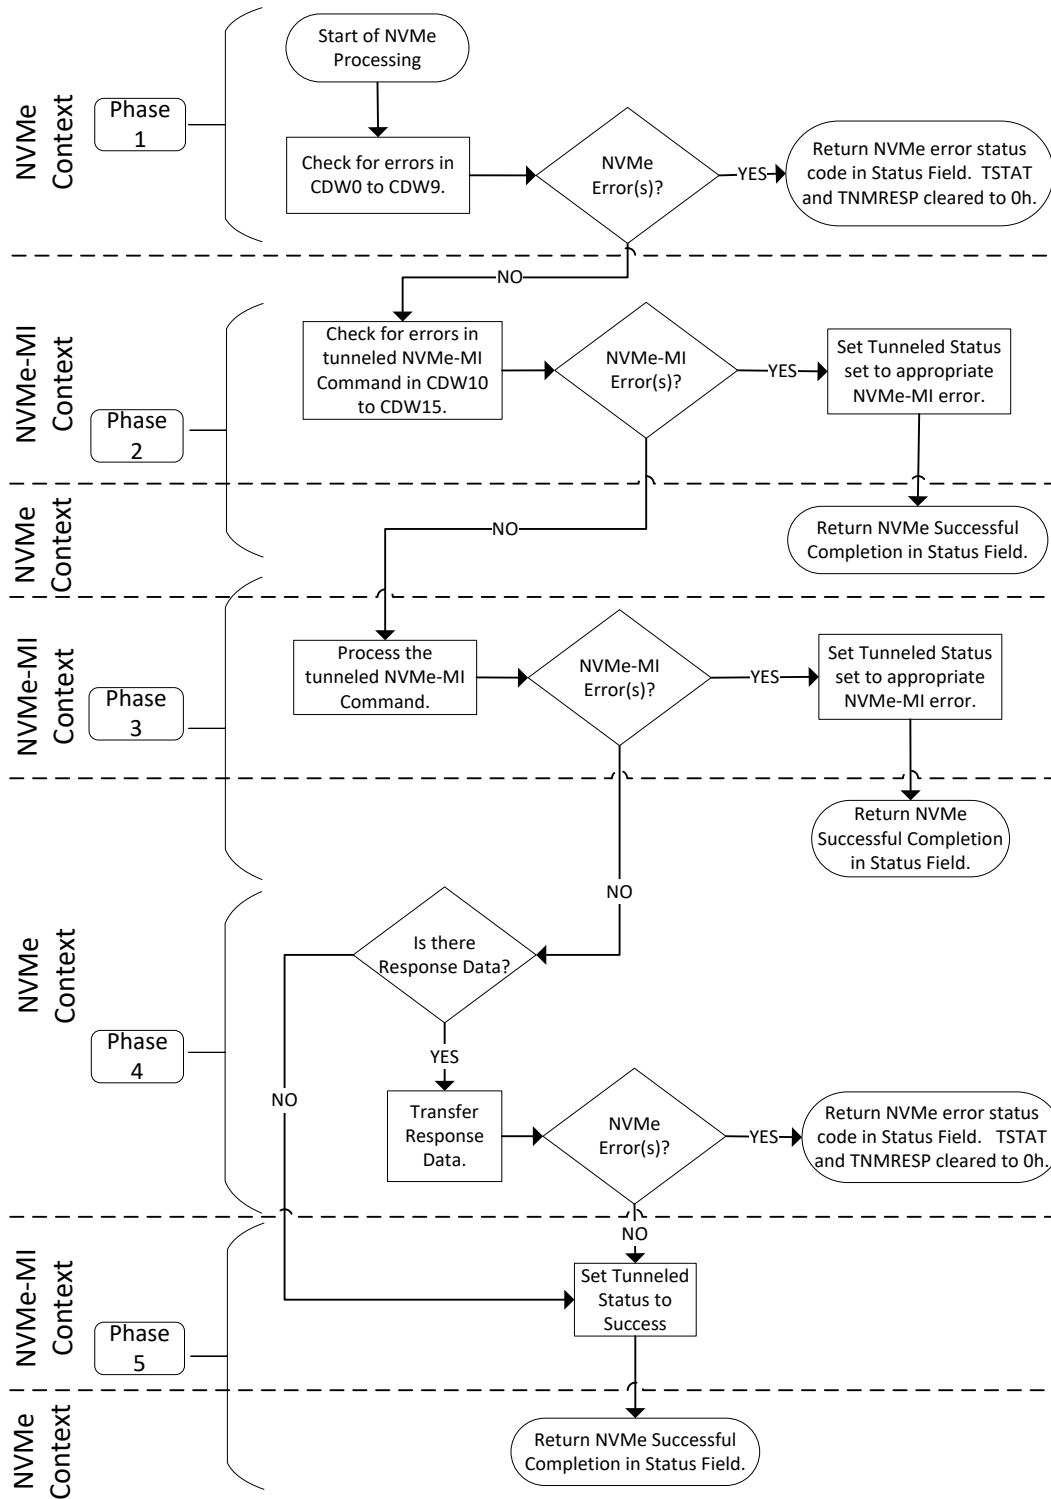
If there are no errors in CDW0 to CDW9, then command servicing enters phase 2 where the tunneled NVMe-MI Command in CDW10 to CDW15 is checked for errors. If any errors are encountered in the tunneled NVMe-MI Command in CDW10 to CDW15, then the NVMe-MI Receive command is completed with a status code of Successful Completion in the Status Field as defined in the NVM Express specification. The Tunneled Status field contains the appropriate error Response Message Status and the Tunneled NVMe Management Response field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band.

If there are no errors in phase 2, then command servicing enters phase 3 where the tunneled NVMe-MI Command finishes processing. If any errors are encountered processing the tunneled NVMe-MI Command, then the NVMe-MI Receive command is completed with a status code of Successful Completion in the Status Field as defined in the NVM Express specification. The Tunneled Status field contains the appropriate error Response Message Status and the Tunneled NVMe Management Response field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band.

If there are no errors in phase 3, then command servicing enters phase 4 where there is a check to determine if there is any Response Data for the tunneled NVMe-MI Command. If there is no Response Data for the tunneled NVMe-MI Command, then command servicing skips to phase 5. If there is Response Data, then the Response Data is transferred to the buffer pointed to by DPTR. If any errors are encountered transferring the Response Data then the command is completed with an error status code in the Status Field as per the NVM Express specification and the Tunneled Status and Tunneled NVMe Management Response fields shall be set to 0h.

If there are no errors in phase 4, then command servicing enters phase 5 where the NVMe-MI Receive command is completed with a status code of Successful Completion in the Status Field as defined in the NVM Express specification. The Tunneled Status field contains a Response Message Status of Success for the tunneled NVMe-MI Command and the Tunneled NVMe Management Response field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band.

Figure 52: NVMe-MI Receive Command Servicing Model



5 Management Interface Command Set

The Management Interface Command Set defines the Command Messages that may be submitted by a Requester when the NMIMT value is set to NVMe-MI Command. The Management Interface Command Set is applicable to both the out-of-band mechanism and the in-band tunneling mechanism.

The NVMe-MI Message structure with all fields that are common to all NVMe-MI Messages are defined in section 3.1. The Response Message structure for the Management Interface Command Set is defined in section 4.1.2. The Message Body for the Management Interface Command Set is shown in Figure 54. Command specific fields for the Management Interface Command Set are defined in this section.

Figure 53: NVMe-MI Command Request Message Format

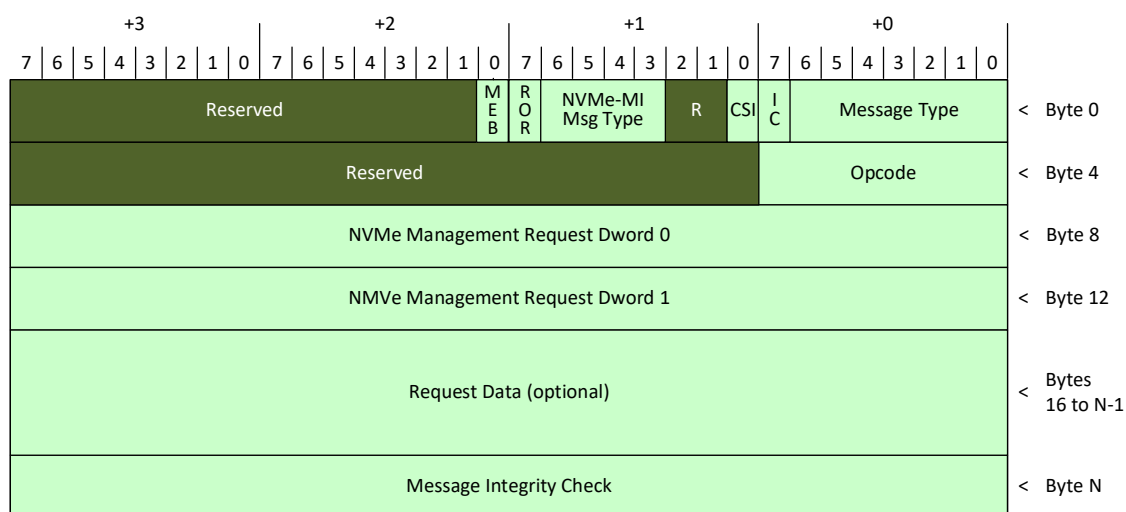


Figure 54: NVMe-MI Command Request Message Description (NCREQ)

Byte	Description
03:00	NVMe-MI Message Header (NMH): Refer to section 3.1.
04	Opcode (OPC): This field specifies the opcode of the NVMe-MI Command to be executed. Refer to Figure 55.
07:05	Reserved
11:08	NVMe Management Dword 0 (NMD0): This field is command specific Dword 0.
15:12	NVMe Management Dword 1 (NMD1): This field is command specific Dword 1.
N-1:16	Request Data (REQD): (Optional)
N+3:N	Message Integrity Check (MIC): Refer to section 3.1.

The Request Data field is an optional field included in some NVMe-MI Commands. If the size of the Request Data does not match the specified Data Length of the Command Message, then the Responder responds with a Generic Error Response and Invalid Command Input Data Size status.

Figure 55 defines the Management Interface Command Set opcodes.

Figure 55: Opcodes for Management Interface Command Set

Opcode	Command
00h	Read NVMe-MI Data Structure
01h	NVM Subsystem Health Status Poll
02h	Controller Health Status Poll
03h	Configuration Set
04h	Configuration Get
05h	VPD Read
06h	VPD Write
07h	Reset
08h	SES Receive
09h	SES Send
0Ah	Management Endpoint Buffer Read
0Bh	Management Endpoint Buffer Write
0Ch to BFh	Reserved
C0h to FFh	Vendor specific

Figure 56 shows the Management Interface Command Set commands that are mandatory, optional, and prohibited for an NVMe Storage Device as well as for an NVMe Enclosure using the out-of-band mechanism. Figure 57 shows Management Interface Command Set commands that are mandatory, optional, and prohibited for an NVMe Storage Device as well as for an NVMe Enclosure using the in-band tunneling mechanism.

Figure 56: Management Interface Command Set Support using an Out-of-Band Mechanism

NVMe Storage Device O/M/P ¹	NVMe Enclosure O/M/P ¹	Command
M	M	Read NVMe-MI Data Structure
M	O ³	NVM Subsystem Health Status Poll
M	O ³	Controller Health Status Poll
M	M ²	Configuration Set
M	M ²	Configuration Get
M	O ³	VPD Read
M	O ³	VPD Write
M	O ³	Reset
P	M	SES Receive
P	M	SES Send
O	M	Management Endpoint Buffer Read
O	M	Management Endpoint Buffer Write

Figure 56: Management Interface Command Set Support using an Out-of-Band Mechanism

NVMe Storage Device O/M/P ¹	NVMe Enclosure O/M/P ¹	Command
O	O	Vendor specific
<p>NOTES:</p> <ol style="list-style-type: none"> O/M/P definition: O = Optional, M = Mandatory, P = Prohibited from being supported. An NVMe Enclosure that is also an NVMe Storage Device (i.e., implements namespaces) shall implement mandatory commands required by either an NVMe Storage Device or an NVMe Enclosure and may implement optional commands allowed by either an NVMe Storage Device or an NVMe Enclosure. This command was architected for an NVMe Storage Device. The mapping of Health Status Change Configuration Identifier to an NVMe Enclosure is outside the scope of this specification. This command was architected for an NVMe Storage Device. The mapping of this command to an NVMe Enclosure is outside the scope of this specification. 		

Figure 57: Management Interface Command Set Support using In-Band Tunneling Mechanism

NVMe Storage Device		NVMe Enclosure		Command
O/M/P ¹	NVMe-MI Send/Receive Mapping ³	O/M/P ¹	NVMe-MI Send/Receive Mapping ³	
M	NVMe-MI Receive	O ²	NVMe-MI Receive	Read NVMe-MI Data Structure
M	NVMe-MI Receive	O ²	NVMe-MI Receive	NVM Subsystem Health Status Poll
M	NVMe-MI Receive	O ²	NVMe-MI Receive	Controller Health Status Poll
M	NVMe-MI Send	O ²	NVMe-MI Send	Configuration Set
M	NVMe-MI Receive	O ²	NVMe-MI Receive	Configuration Get
M	NVMe-MI Receive	O ²	NVMe-MI Receive	VPD Read
M	NVMe-MI Send	O ²	NVMe-MI Send	VPD Write
M	NVMe-MI Send	O ²	NVMe-MI Send	Reset
P	n/a	M	NVMe-MI Receive	SES Receive
P	n/a	M	NVMe-MI Send	SES Send
P	n/a	P	n/a	Management Endpoint Buffer Read
P	n/a	P	n/a	Management Endpoint Buffer Write
O	Vendor Specific	O	Vendor Specific	Vendor specific
<p>NOTES:</p> <ol style="list-style-type: none"> O/M/P definition: O = Optional, M = Mandatory, P = Prohibited from being supported. An NVMe Enclosure that is also an NVMe Storage Device (i.e., implements namespaces) shall implement mandatory commands required by either an NVMe Storage Device or an NVMe Enclosure and may implement optional commands allowed by either an NVMe Storage Device or an NVMe Enclosure. This command was architected for an NVMe Storage Device. The mapping of this command to an NVMe Enclosure is outside the scope of this specification. This column indicates whether the NVMe-MI Command is tunneled in-band using the NVMe-MI Send or NVMe-MI Receive command. 				

Figure 58: NVMe-MI Command Response Message Format

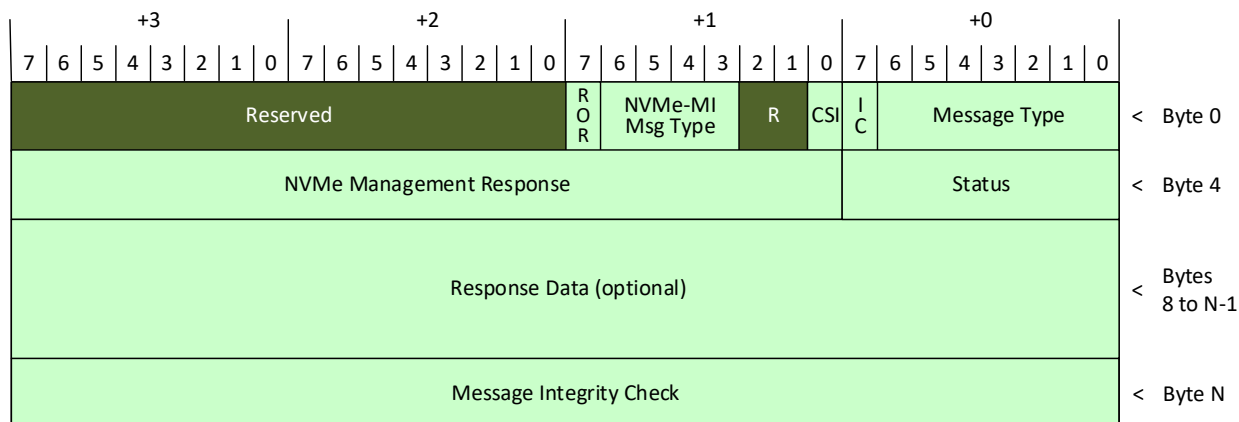


Figure 59: NVMe-MI Command Response Message Description (NCRESP)

Byte	Description
03:00	NVMe-MI Message Header (NMH): Refer to section 3.1.
04	Status (STATUS): This field indicates the status of the NVMe-MI Command. Refer to section 4.1.2.
07:05	NVMe Management Response (NMRESP): This field is command specific.
N-1:08	Response Data (RESPD): (Optional)
N+3:N	Message Integrity Check (MIC): Refer to section 3.1.

5.1 Configuration Get

The Configuration Get command allows the Requester to read the current configuration of a Responder.

The command uses NVMe Management Dwords 0 and 1. The format of NVMe Management Dwords 0 and 1 are shown in Figure 60 and Figure 61 respectively. There is no Request Data included in a Configuration Get command.

Figure 60: Configuration Get – NVMe Management Dword 0

Bit	Description
31:08	Configuration Identifier specific
07:00	Configuration Identifier: This field specifies the identifier of the Configuration that is being read. Refer to Figure 62.

Figure 61: Configuration Get – NVMe Management Dword 1

Bit	Description
31:00	Configuration Identifier specific

NVMe-MI Configuration Identifiers are listed in Figure 62.

Figure 62: NVMe Management Interface Configuration Identifiers

Configuration Identifier	Out-of-Band Mechanism O/M/P ¹	In-Band Tunneling Mechanism O/M/P ¹	Description
00h	-	-	Reserved
01h	M	P	SMBus/I2C Frequency
02h	M	M	Health Status Change
03h	M	P	MCTP Transmission Unit Size
04h to BFh	-	-	Reserved
C0h to FFh	O	O	Vendor Specific
NOTES:			
1. O/M/P definition: O = Optional, M = Mandatory, P = Prohibited from being supported.			

The NVMe Management Response field is configuration specific.

5.1.1 SMBus/I2C Frequency (Configuration Identifier 01h)

The SMBus/I2C Frequency configuration indicates the current frequency of the SMBus port, if applicable.

The configuration specific fields in NVMe Management Dword 0 are shown in Figure 63. The configuration specific fields in NVMe Management Dword 1 are reserved. The current SMBus/I2C Frequency configuration is returned in the NVMe Management Response field as shown in Figure 64.

Figure 63: SMBus/I2C Frequency – NVMe Management Dword 0

Bit	Description
31:24	Port Identifier: This field specifies the port whose SMBus/I2C Frequency is indicated.
23:08	Reserved
07:00	Configuration Identifier: This field specifies the identifier of the Configuration that is being read. Refer to Figure 62.

Figure 64: SMBus/I2C Frequency – NVMe Management Response

Bit	Description												
23:04	Reserved												
03:00	SMBus/I2C Frequency: The current frequency of the SMBus/I2C. The default value for this field following a reset or power cycle is 1h, if SMBus is supported. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>SMBus is not supported or is disabled</td> </tr> <tr> <td>1h</td> <td>100 kHz</td> </tr> <tr> <td>2h</td> <td>400 kHz</td> </tr> <tr> <td>3h</td> <td>1 MHz</td> </tr> <tr> <td>4h to Fh</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Description	0h	SMBus is not supported or is disabled	1h	100 kHz	2h	400 kHz	3h	1 MHz	4h to Fh	Reserved
		Value	Description										
		0h	SMBus is not supported or is disabled										
		1h	100 kHz										
		2h	400 kHz										
		3h	1 MHz										
4h to Fh	Reserved												

5.1.2 Health Status Change (Configuration Identifier 02h)

The Health Status Change configuration is used to clear the selected status bits in the Composite Controller Status field using Configuration Set. A Requester should not use Configuration Get for this Configuration Identifier.

The configuration specific fields in NVMe Management Dwords 0 and 1 are reserved. A Responder shall complete a Configuration Get command on this Configuration Identifier with a Success Response. The NVMe Management Response field is reserved and there is no Response Data.

5.1.3 MCTP Transmission Unit Size (Configuration Identifier 03h)

The MCTP Transmission Unit Size configuration indicates the current MCTP Transmission Unit Size of the Port Identifier specified in Dword 0.

The configuration specific fields in NVMe Management Dword 0 are shown in Figure 65. The configuration specific fields in NVMe Management Dword 1 are reserved. The current Transmission unit size of the specified port is returned in the NVMe Management Response field as shown in Figure 66.

Figure 65: MCTP Transmission Unit Size – NVMe Management Dword 0

Bit	Description
31:24	Port Identifier: This field specifies the port whose MCTP Transmission Unit Size is indicated.
23:08	Reserved
07:00	Configuration Identifier: This field specifies the identifier of the Configuration that is being read. Refer to Figure 62.

Figure 66: MCTP Transmission Unit Size – NVMe Management Response

Bit	Description
23:16	Reserved
15:00	MCTP Transmission Unit Size: This field contains the MCTP Transmission Unit Size in bytes to be used by the port. The default value for this field following a reset or power cycle is 40h (64).

5.2 Configuration Set

The Configuration Set command allows the Requester to modify the current configuration of a Responder.

The command uses NVMe Management Dwords 0 and 1. The format of NVMe Management Dwords 0 and 1 are shown in Figure 67 and Figure 68 respectively. There is no Request Data included in a Configuration Set command.

Figure 67: Configuration Set – NVMe Management Dword 0

Bit	Description
31:08	Configuration Identifier specific
07:00	Configuration Identifier: This field specifies the identifier of the Configuration that is being written. Refer to Figure 62.

Figure 68: Configuration Set – NVMe Management Dword 1

Bit	Description
31:00	Configuration Identifier specific

NVMe-MI Configuration Identifiers are listed in Figure 62. Specifying a reserved identifier in the Configuration Identifier field causes the command to complete with an Invalid Parameter Error Response.

The NVMe Management Response field is configuration Identifier specific.

5.2.1 SMBus/I2C Frequency (Configuration Identifier 01h)

The SMBus/I2C Frequency configuration specifies a new frequency for the SMBus port.

The configuration specific fields in NVMe Management Dword 0 are shown in Figure 69. The configuration specific fields in NVMe Management Dword 1 are reserved. NVMe Management Response field is reserved.

After successful completion of this command, the SMBus/I2C frequency is updated to the specified frequency. A Management Controller should not change this configuration while there are other Command Messages outstanding.

If the specified frequency is not supported or the Port Identifier specified is not an SMBus/I2C port, the Management Endpoint shall respond with an Invalid Parameter Error Response.

Figure 69: SMBus/I2C Frequency – NVMe Management Dword 0

Bit	Description		
31:24	Port Identifier: This field specifies the port whose SMBus/I2C Frequency is specified.		
23:12	Reserved		
11:08	SMBus/I2C Frequency: This field specifies the new frequency for the specified SMBus/I2C port.		
		Value	Description
		0h	Reserved
		1h	100 kHz
		2h	400 kHz
		3h	1 MHz
4h to Fh	Reserved		
07:00	Configuration Identifier: This field specifies the identifier of the Configuration that is being written. Refer to Figure 62.		

5.2.2 Health Status Change (Configuration Identifier 02h)

This Configuration Identifier is used to clear selected status bits in the Composite Controller Status field of the NVM Subsystem Health Data Structure (refer to Figure 87) returned by the NVM Subsystem Health Status Poll command.

The Composite Controller Status field of the NVM Subsystem Health Data Structure is used to report the occurrence of health and status events associated with the NVM Subsystem. When a bit in this field is set to '1', it remains a '1' until cleared by a Requester.

A Configuration Set command that selects Health Status Change may be used to clear corresponding bits selected in NVMe Management Dword 1 of the Composite Controller Status field to '0'.

A Configuration Set command that selects Health Status Change operates independently in the out-of-band mechanism and the in-band tunneling mechanism.

An NVMe Storage Device or NVMe Enclosure supporting the Health Status Change Configuration Identifier in the out-of-band mechanism shall have an independent copy of the Composite Controller Status dedicated to the out-of-band mechanism. In the out-of-band mechanism, a Configuration Set command that selects Health Status Change only applies to the copy of the Composite Controller Status dedicated to the out-of-band mechanism. Refer to section 5.4 for more details on Composite Controller Status.

An NVMe Storage Device or NVMe Enclosure supporting the Health Status Change Configuration Identifier in the in-band tunneling mechanism shall have an independent copy of the Composite Controller Status dedicated to the in-band tunneling mechanism. In the in-band tunneling mechanism, a Configuration Set command that selects Health Status Change only applies to the copy of the Composite Controller Status dedicated to the in-band tunneling mechanism.

Figure 70: Health Status Change - NVMe Management Dword 0

Bit	Description
31:08	Reserved
07:00	Configuration Identifier: This field specifies the identifier of the Configuration that is being written. Refer to Figure 62.

Figure 71: Health Status Change – NVMe Management Dword 1

Bit	Description
31:12	Reserved
11	Critical Warning (CWARN): When this bit is set to '1', bit 12 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
10	Available Spare (SPARE): When this bit is set to '1', bit 11 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
09	Percentage Used (PDLU): When this bit is set to '1', bit 10 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
08	Composite Temperature (CTEMP): When this bit is set to '1', bit 9 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
07	Controller Status Change (CSCHNG): When this bit is set to '1', bit 8 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
06	Firmware Activated (FA): When this bit is set to '1', bit 7 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
05	Namespace Attribute Changed (NAC): When this bit is set to '1', bit 6 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
04	Controller Enable Change Occurred (CECO): When this bit is set to '1', bit 5 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
03	NVM Subsystem Reset Occurred (NSSRO): When this bit is set to '1', bit 4 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.

Figure 71: Health Status Change – NVMe Management Dword 1

Bit	Description
02	Shutdown Status (SHST): When this bit is set to '1', bit 2 in the Composite Controller Status field of the NVMe Subsystem Health Data Structure is cleared to '0'.
01	Controller Fatal Status (CFS): When this bit is set to '1', bit 1 in the Composite Controller Status field of the NVMe Subsystem Health Data Structure is cleared to '0'.
00	Ready (RDY): When this bit is set to '1', bit 0 in the Composite Controller Status field of the NVMe Subsystem Health Data Structure is cleared to '0'.

5.2.3 MCTP Transmission Unit Size (Configuration Identifier 03h)

The MCTP Transmission Unit Size configuration specifies a new MCTP Transmission Unit Size for the specified Port Identifier. A Management Controller should check the maximum MCTP Transmission Unit Size for the port reported by the Management Endpoint using the Read NVMe-MI Data Structure command (refer to Figure 91).

The configuration specific fields in NVMe Management Dwords 0 and 1 are shown in Figure 72 and Figure 73 respectively. The NVMe Management Response field is reserved.

After successful completion of this command, the MCTP Transmission Unit Size for MCTP packets on the specified port is updated to the specified size for future Command Messages. A Management Controller should not change this configuration while there are other commands outstanding.

If the specified MCTP Transmission Unit Size is not supported or the Port Identifier specified is not valid, the Management Endpoint shall abort the command and send a Response Message with an Invalid Parameter Error Response.

Figure 72: MCTP Transmission Unit Size – NVMe Management Dword 0

Bit	Description
31:24	Port Identifier: This field specifies the port whose MCTP Transmission Unit Size is specified.
23:08	Reserved
07:00	Configuration Identifier: This field specifies the identifier of the Configuration that is being written. Refer to Figure 62.

Figure 73: MCTP Transmission Unit Size – NVMe Management Dword 1

Bit	Description
31:16	Reserved
15:00	MCTP Transmission Unit Size: This field contains the MCTP Transmission Unit Size in bytes to be used by the port.

5.3 Controller Health Status Poll

The Controller Health Status Poll command is used to efficiently determine changes in health status attributes associated with one or more Controllers in the NVMe Subsystem.

The Controller Health Status Poll command operates independently in the out-of-band mechanism and the in-band tunneling mechanism.

An NVMe Storage Device or NVMe Enclosure supporting the Controller Health Status Poll command in the out-of-band mechanism shall have an independent copy of the Controller Health Data Structure (refer to Figure 77) and the Controller Health Status Changed Flags (refer to Figure 78) dedicated to the out-of-band mechanism. In the out-of-band mechanism, a Controller Health Status Poll command only applies to the copy of the Controller Health Data Structure and the Controller Health Status Changed Flags dedicated to the out-of-band mechanism.

An NVMe Storage Device or NVMe Enclosure supporting the Controller Health Status Poll command in the in-band tunneling mechanism shall have an independent copy of the Controller Health Data Structure and the Controller Health Status Changed Flags dedicated to the in-band tunneling mechanism. In the in-band tunneling mechanism, a Controller Health Status Poll command only applies to the copy of the Controller Health Data Structure and the Controller Health Status Changed Flags dedicated to the in-band tunneling mechanism.

The Controller Health Status Poll command uses NVMe Management Dwords 0 and 1. The format of NVMe Management Dword 0 is shown in Figure 74 and the format of NVMe Management Dword 1 is shown in Figure 75.

Figure 74: Controller Health Status Poll – NVMe Management Dword 0

Bit	Description
31	Report All (ALL): When this bit is set to '1', a Controller Health Data Structure is returned regardless of the status of the Controller Health Status Changed Flags. The Controller selection fields (SCTLID, MAXRENT, INCF, INCPF, and INCVF) still apply even when this bit is set to '1' but the error selection bits (CWARN, SPARE, PDLU, and CTEMP in Figure 75) do not apply. When this bit is cleared to '0', a Controller Health Data Structure is returned based on the Controller selection fields (SCTLID, MAXRENT, INCF, INCPF, and INCVF) and error selection fields (CWARN, SPARE, PDLU, and CTEMP in Figure 75).
30:27	Reserved
26	Include SR-IOV Virtual Functions (INCVF): When this bit is set to '1', a Controller Health Data Structure is returned for Controllers associated with SR-IOV Virtual Functions (VFs).
25	Include SR-IOV Physical Functions (INCPF): When this bit is set to '1', a Controller Health Data Structure is returned for Controllers associated with SR-IOV Physical Functions (PFs).
24	Include PCI Functions (INCF): When this bit is set to '1', a Controller Health Status Data Structure is returned for Controllers associated with non-SR-IOV PCI Functions.
23:16	Maximum Response Entries (MAXRENT): This field specifies the maximum number of Controller Health Data Structure entries that may be returned in the completion. This is a 0's based field. The maximum number of entries is 255. Specifying 256 entries is interpreted as an Invalid Parameter.
15:00	Starting Controller ID (SCTLID): This field specifies the Controller ID of the first Controller whose Controller Health Data Structure may be returned.

Figure 75: Controller Health Status Poll – NVMe Management Dword 1

Bit	Description
31	Clear Changed Flags (CCF): When this bit is set to '1', the Controller Health Status Changed Flags are cleared in Controllers whose Controller Health Data Structure is contained in the Response Data.
30:05	Reserved
04	Critical Warning (CWARN): When this bit is set to '1', a Controller Health Data Structure is returned for Controllers with the Critical Warning bit set to '1' in their Controller Health Status Changed Flags.
03	Available Spare (SPARE): When this bit is set to '1', a Controller Health Data Structure is returned for Controllers with the Available Spare bit set to '1' in their Controller Health Status Changed Flags.
02	Percentage Used (PDLU): When this bit is set to '1', a Controller Health Data Structure is returned for Controllers with the Percent Used bit set to '1' in their Controller Health Status Changed Flags.

Figure 75: Controller Health Status Poll – NVMe Management Dword 1

Bit	Description
01	Composite Temperature Changes (CTEMP): When this bit is set to '1', a Controller Health Data Structure is returned for Controllers with the Composite Temperature bit set to '1' in their Controller Health Status Changed Flags.
00	Controller Status Changes (CSTS): When this bit is set to '1', a Controller Health Data Structure is returned for Controllers with the Ready, Controller Fatal Status, Shutdown Status, NVM Subsystem Reset Occurred, Controller Enable Change Occurred, Namespace Attribute Changed, or Firmware Activated bit set to '1' in their Controller Health Status Changed Flags.

The Controller Health Status Poll Response Messages use the NVMe Management Response field with the format shown in Figure 76.

The Response Data field size may vary based on the number of Controllers whose Controller Health Data Structure has changed and based on the number of Controllers whose Controller Health Data Structure is filtered out by Controller type (refer to section 5.3.1) or Controller Health Status Changed Flags (refer to section 5.3.2). The Response Entries field indicates the number of Controller Health Data Structures that are contained in the Response Data.

Figure 76: Controller Health Status Poll – NVMe Management Response

Bit	Description
23:16	Response Entries (RENT): This field specifies the number of Controller Health Data Structure Entries present in the Response Data for this Response Message.
15:00	Reserved

The Controller Health Data Structure, shown in Figure 77, contains the health status attributes that are tracked for each Controller. When the command is executed, health status is returned for up to 255 Controllers starting at or above the Starting Controller ID (SCTLID). Controllers are returned in ascending order of Controller Identifier starting at offset 0h of the Response Data.

Figure 77: Controller Health Data Structure (CHDS)

Bytes	Description
01:00	Controller Identifier (CTLID): This field specifies the Controller Identifier with which the data contained in this data structure is associated.

Figure 77: Controller Health Data Structure (CHDS)

Bytes	Description																											
03:02	<p>Controller Status (CSTS): This field reports the Controller status.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Reset</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:08</td> <td>0</td> <td>Reserved</td> </tr> <tr> <td>07</td> <td>Hwlnit</td> <td> <p>Firmware Activated (FA): This bit is set to '1' when a new firmware image is activated. Firmware activation is described in the NVM Express specification.</p> <p>The reset value of this bit is set to '1' if a reset caused a new firmware image to be activated.</p> </td> </tr> <tr> <td>06</td> <td>0</td> <td> <p>Namespace Attribute Changed (NAC): This bit is set to '1' under the same conditions that causes the Namespace Attribute Changed asynchronous event to be sent if Namespace Attribute Notices are enabled as specified in the NVM Express specification. This bit may be set to '1' regardless of whether Namespace Attribute Notices are enabled or not.</p> </td> </tr> <tr> <td>05</td> <td>0</td> <td> <p>Controller Enable Change Occurred (CECO): This bit is set to '1' when the Enable bit (refer to CC.EN in the NVM Express specification) changes state.</p> </td> </tr> <tr> <td>04</td> <td>Hwlnit</td> <td> <p>NVM Subsystem Reset Occurred (NSSRO): This bit corresponds to the value of the NVM Subsystem Reset Occurred (refer to CSTS.NSSRO in the NVM Express specification) bit.</p> </td> </tr> <tr> <td>03:02</td> <td>0</td> <td> <p>Shutdown Status (SHST): This field corresponds to the value of the Shutdown Status (refer to CSTS.SHST in the NVM Express specification) field.</p> </td> </tr> <tr> <td>01</td> <td>Hwlnit</td> <td> <p>Controller Fatal Status (CFS): This bit corresponds to the value of the Controller Fatal Status (refer to CSTS.CFS in the NVM Express specification) bit.</p> </td> </tr> <tr> <td>00</td> <td>0</td> <td> <p>Ready (RDY): This bit corresponds to the value of the Ready (refer to CSTS.RDY in the NVM Express specification) bit.</p> </td> </tr> </tbody> </table>	Bit	Reset	Description	15:08	0	Reserved	07	Hwlnit	<p>Firmware Activated (FA): This bit is set to '1' when a new firmware image is activated. Firmware activation is described in the NVM Express specification.</p> <p>The reset value of this bit is set to '1' if a reset caused a new firmware image to be activated.</p>	06	0	<p>Namespace Attribute Changed (NAC): This bit is set to '1' under the same conditions that causes the Namespace Attribute Changed asynchronous event to be sent if Namespace Attribute Notices are enabled as specified in the NVM Express specification. This bit may be set to '1' regardless of whether Namespace Attribute Notices are enabled or not.</p>	05	0	<p>Controller Enable Change Occurred (CECO): This bit is set to '1' when the Enable bit (refer to CC.EN in the NVM Express specification) changes state.</p>	04	Hwlnit	<p>NVM Subsystem Reset Occurred (NSSRO): This bit corresponds to the value of the NVM Subsystem Reset Occurred (refer to CSTS.NSSRO in the NVM Express specification) bit.</p>	03:02	0	<p>Shutdown Status (SHST): This field corresponds to the value of the Shutdown Status (refer to CSTS.SHST in the NVM Express specification) field.</p>	01	Hwlnit	<p>Controller Fatal Status (CFS): This bit corresponds to the value of the Controller Fatal Status (refer to CSTS.CFS in the NVM Express specification) bit.</p>	00	0	<p>Ready (RDY): This bit corresponds to the value of the Ready (refer to CSTS.RDY in the NVM Express specification) bit.</p>
	Bit	Reset	Description																									
	15:08	0	Reserved																									
	07	Hwlnit	<p>Firmware Activated (FA): This bit is set to '1' when a new firmware image is activated. Firmware activation is described in the NVM Express specification.</p> <p>The reset value of this bit is set to '1' if a reset caused a new firmware image to be activated.</p>																									
	06	0	<p>Namespace Attribute Changed (NAC): This bit is set to '1' under the same conditions that causes the Namespace Attribute Changed asynchronous event to be sent if Namespace Attribute Notices are enabled as specified in the NVM Express specification. This bit may be set to '1' regardless of whether Namespace Attribute Notices are enabled or not.</p>																									
	05	0	<p>Controller Enable Change Occurred (CECO): This bit is set to '1' when the Enable bit (refer to CC.EN in the NVM Express specification) changes state.</p>																									
	04	Hwlnit	<p>NVM Subsystem Reset Occurred (NSSRO): This bit corresponds to the value of the NVM Subsystem Reset Occurred (refer to CSTS.NSSRO in the NVM Express specification) bit.</p>																									
	03:02	0	<p>Shutdown Status (SHST): This field corresponds to the value of the Shutdown Status (refer to CSTS.SHST in the NVM Express specification) field.</p>																									
	01	Hwlnit	<p>Controller Fatal Status (CFS): This bit corresponds to the value of the Controller Fatal Status (refer to CSTS.CFS in the NVM Express specification) bit.</p>																									
00	0	<p>Ready (RDY): This bit corresponds to the value of the Ready (refer to CSTS.RDY in the NVM Express specification) bit.</p>																										
05:04	<p>Composite Temperature (CTEMP): This field contains a value corresponding to a temperature in degrees Kelvin that represents the current composite temperature of the Controller and namespace(s) associated with that Controller. The value of this field corresponds to the value in the Controller's SMART / Health Information Log.</p>																											
06	<p>Percentage Used (PDLU): This field contains a vendor specific estimate of the percentage of NVM Subsystem life used based on the actual usage and the manufacturer's prediction of NVM life. The value of this field corresponds to the value in the Controller's SMART / Health Information Log.</p>																											
07	<p>Available Spare (SPARE): This field contains a normalized percentage (0% to 100%) of the remaining spare capacity available. The value of this field corresponds to the value in the Controller's SMART / Health Information Log.</p>																											
08	<p>Critical Warning (CWARN): This field indicates critical warnings for the state of the Controller. The value of this field corresponds to the value in the Controller's SMART / Health Information Log.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:5</td> <td>Reserved</td> </tr> <tr> <td>4</td> <td> <p>Volatile Memory Backup Failed (VMBF): This bit is set to '1' when the volatile memory backup device has failed.</p> </td> </tr> <tr> <td>3</td> <td> <p>Read Only (RO): This bit is set to '1' when the media has been placed in read only mode.</p> </td> </tr> <tr> <td>2</td> <td> <p>Reliability Degraded (RD): This bit is set to '1' when NVM Subsystem reliability has been degraded due to significant media related errors or an internal error.</p> </td> </tr> <tr> <td>1</td> <td> <p>Temperature Above or Under Threshold (TAUT): This bit is set to '1' when a temperature is above an over temperature threshold or below an under-temperature threshold.</p> </td> </tr> <tr> <td>0</td> <td> <p>Spare Threshold (ST): This bit is set to '1' when the available spare has fallen below the available spare threshold.</p> </td> </tr> </tbody> </table>	Bit	Description	7:5	Reserved	4	<p>Volatile Memory Backup Failed (VMBF): This bit is set to '1' when the volatile memory backup device has failed.</p>	3	<p>Read Only (RO): This bit is set to '1' when the media has been placed in read only mode.</p>	2	<p>Reliability Degraded (RD): This bit is set to '1' when NVM Subsystem reliability has been degraded due to significant media related errors or an internal error.</p>	1	<p>Temperature Above or Under Threshold (TAUT): This bit is set to '1' when a temperature is above an over temperature threshold or below an under-temperature threshold.</p>	0	<p>Spare Threshold (ST): This bit is set to '1' when the available spare has fallen below the available spare threshold.</p>													
	Bit	Description																										
	7:5	Reserved																										
	4	<p>Volatile Memory Backup Failed (VMBF): This bit is set to '1' when the volatile memory backup device has failed.</p>																										
	3	<p>Read Only (RO): This bit is set to '1' when the media has been placed in read only mode.</p>																										
	2	<p>Reliability Degraded (RD): This bit is set to '1' when NVM Subsystem reliability has been degraded due to significant media related errors or an internal error.</p>																										
1	<p>Temperature Above or Under Threshold (TAUT): This bit is set to '1' when a temperature is above an over temperature threshold or below an under-temperature threshold.</p>																											
0	<p>Spare Threshold (ST): This bit is set to '1' when the available spare has fallen below the available spare threshold.</p>																											
15:09	Reserved																											

Associated with each Controller in the NVM Subsystem is a set of Controller Health Status Changed Flags shown in Figure 78. The Controller Health Status Changed Flags are set when the corresponding field in the Controller Health Data Structure changes state as described in Figure 78. Figure 79 shows a graphical representation of which field(s)/bit(s) in the Controller Health Data Structure are associated with each bit in the Controller Health Status Changed Flags. When a bit in the Controller Health Status Changed Flags for any Controller transitions from '0' to '1', then the corresponding bit in the Composite Controller Status is also set to '1'. The Controller Health Status Changed Flags are cleared in Controllers whose Controller Health Data Structure is returned in the Success Response to a Controller Health Status Poll Command Message with the Clear Changed Flags bit set to '1'.

A Controller Health Status Poll response may return the Controller Health Data Structure for up to 255 Controllers in the Response Data field. An NVM Subsystem may contain up to 64 Ki Controllers, so a method is needed to limit the size of the Response Message. The Starting Controller ID field in the Command Message specifies the Controller ID of the first Controller whose Controller Health Data Structure may be returned in the Response Data field. The Maximum Response Entries field specifies the maximum number of Controllers whose Controller Health Data Structure may be returned in the Response Data field.

The Response Data field contains the Controller Health Status Data Structure for up to the first M Controllers starting with Controller N, where M is equal to the Maximum Response Entries field and N is equal to the Starting Controller ID field. The Response Data field shall contain the Controller Health Status Data Structure for all Controllers that do not match the filtering criteria in Controller Health Status Poll - NVMe Management Dword 0 (refer to section 5.3.1) and that have one or more Controller Health Status Changed Flags that are: a) set and b) do not match the filtering criteria in Controller Health Status Poll - NVMe Management Dword 1 (refer to section 5.3.2). The Response Data field shall not contain the Controller Health Status Data Structure for any Controllers that meet the filtering criteria in sections 5.3.1 or 5.3.2.

5.3.1 Filtering by Controller Type

The Controller Health Data Structures that are returned by Controller Health Status Poll may be filtered (i.e., excluded from being included in the Response Data field regardless of the state of the Controller Health Status Changed Flags) by Controller type (i.e., non SR-IOV PCI Function, SR-IOV PF, and SR-IOV VF). Controller type filtering is controlled by the Include PCI Functions, Include SR-IOV PFs, and Include SR-IOV VFs fields in NVMe Management Dword 0. When one of these bits is set to '1', Controller Health Data Structures for Controllers corresponding to that type of PCI Function are included in the Response Data field; else, the Controller Health Data Structure for that Controller is excluded from the Response Data field.

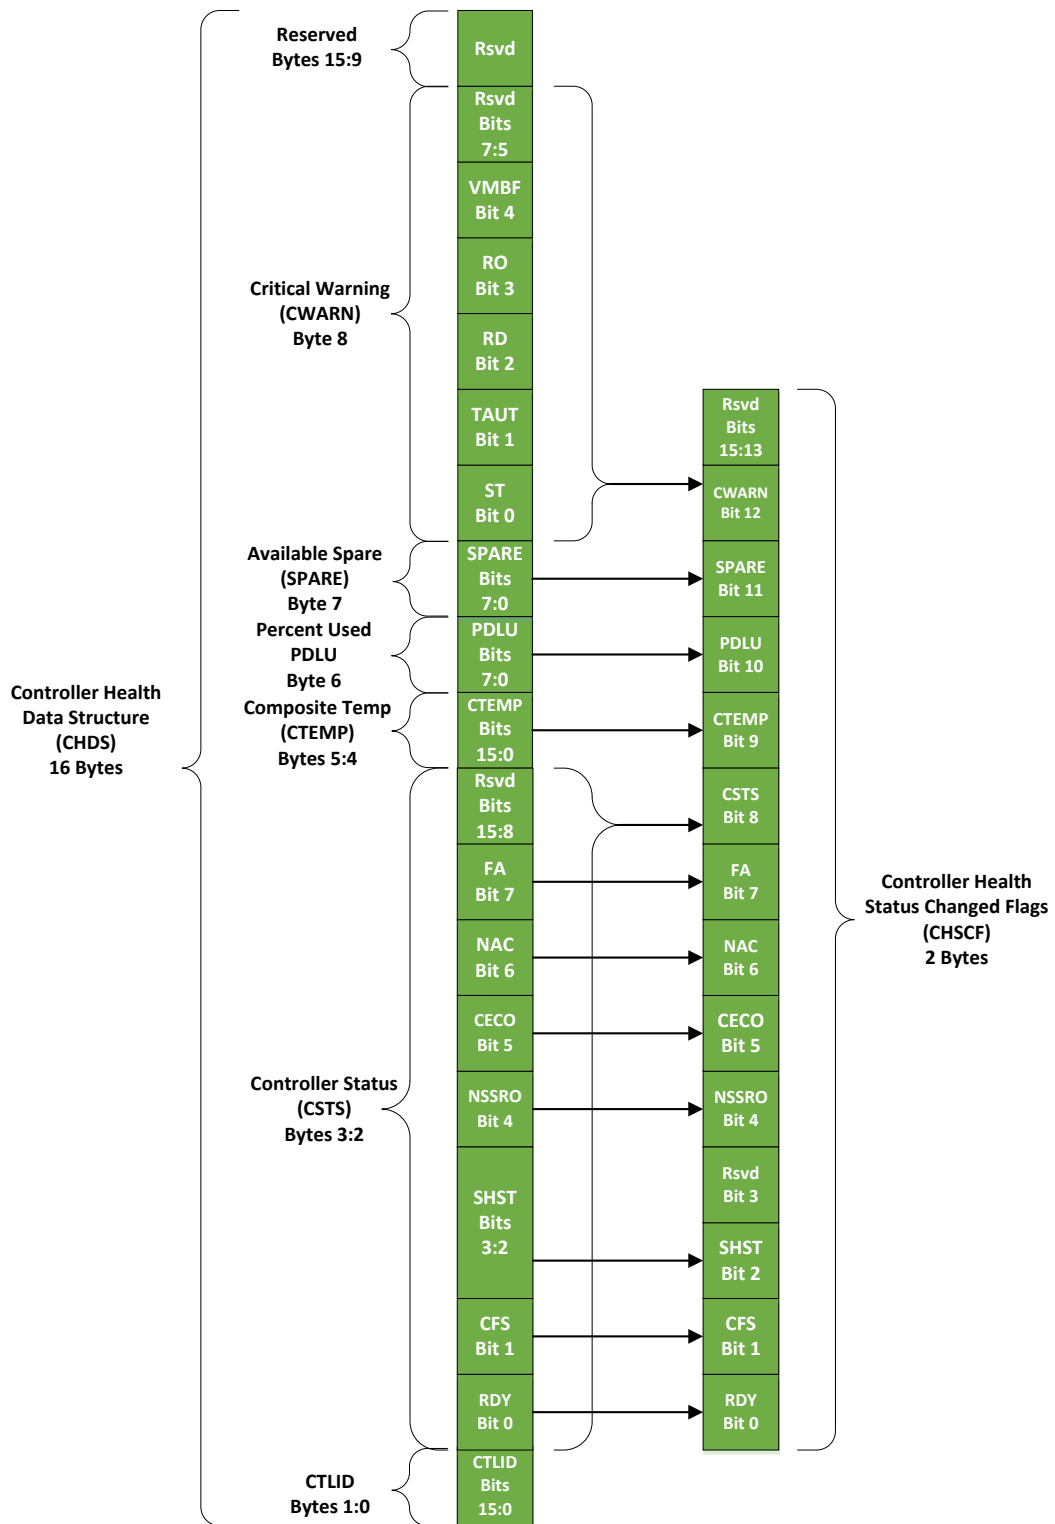
5.3.2 Filtering by Controller Health Status Changed Flags

The Controller Health Data Structures that are returned by Controller Health Status Poll may also be filtered by the Controller Health Status Changed Flags. Filtering of changes by Controller Health Status Changed Flags is controlled by some of the bits in NVMe Management Dword 1. When one or more of these bits is set and any of the corresponding bit(s) in the Controller Health Status Changed Flags for the Controller are also set (refer to Figure 75 for Controller Health Status Changed Flags associated with each bit in NVMe Management Dword 1), then the entire Controller Health Data Structure (including any filtered fields) for that Controller is returned in the Response Data field; else, the Controller Health Data Structure for that Controller is excluded from the Response Data field. The contents returned in the Controller Health Data Structure for filtered fields are undefined.

Figure 78: Controller Health Status Changed Flags (CHSCF)

Bit	Reset	Description
15:13	0	Reserved
12	0	Critical Warning (CWARN): This bit is set to '1' when any of the Critical Warning bits in the Controller Health Data Structure transition from '0' to '1'.
11	0	Available Spare (SPARE): This bit is set to '1' when the Available Spare field in the Controller Health Data Structure changes state.
10	0	Percentage Used (PDLU): This bit is set to '1' when the Percentage Used field in the Controller Health Data Structure changes state.
09	0	Composite Temperature Change (CTEMP): This bit is set to '1' when the Composite Temperature field in the Controller Health Data Structure changes state.
08	HwInit	Controller Status Change (CSTS): This bit is set to '1' when the Shutdown Status field in the Controller Health Data Structure changes state or when the Ready, Controller Fatal Status, NVM Subsystem Reset Occurred, Controller Enable Change Occurred, Namespace Attribute Changed, or Firmware Activated bit in the Controller Health Data Structure transitions from '0' to '1'.
07	HwInit	Firmware Activated (FA): This bit is set to '1' when the Firmware Activated bit in the Controller Health Data Structure transitions from '0' to '1'.
06	0	Namespace Attribute Changed (NAC): This bit is set to '1' when the Namespace Attribute Changed bit in the Controller Health Data Structure transitions from '0' to '1'.
05	0	Controller Enable Change Occurred (CECO): This bit is set to '1' when the Controller Enable Change Occurred bit in the Controller Health Data Structure transitions from '0' to '1'.
04	HwInit	NVM Subsystem Reset Occurred (NSSRO): This bit is set to '1' when the NVM Subsystem Reset Occurred bit in the Controller Health Data Structure transitions from '0' to '1'.
03	0	Reserved
02	0	Shutdown Status (SHST): This bit is set to '1' when the Shutdown Status field in the Controller Health Data Structure changes state.
01	HwInit	Controller Fatal Status (CFS): This bit is set to '1' when the Controller Fatal Status bit in the Controller Health Data Structure transitions from '0' to '1'.
00	0	Ready (RDY): This bit is set to '1' when the Ready bit in the Controller Health Data Structure transitions from '0' to '1'.

Figure 79: Controller Health Data Structure to Controller Health Status Changed Flags Mapping

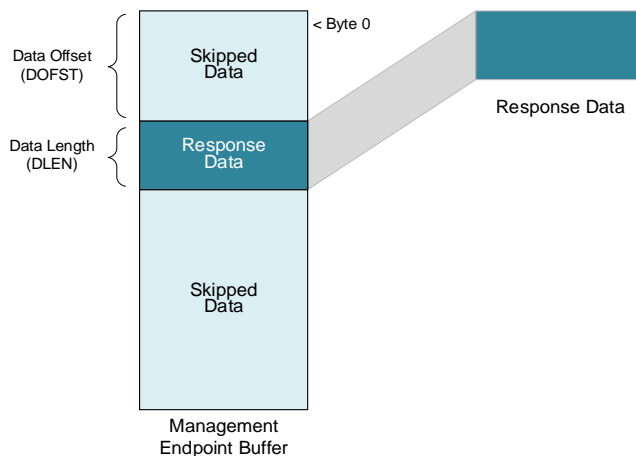


5.4 Management Endpoint Buffer Read

The Management Endpoint Buffer Read command allows the Management Controller to read the contents of the Management Endpoint Buffer. This data is returned in the Response Data.

The command uses NVMe Management Dwords 0 and 1. The format of NVMe Management Dwords 0 and 1 are shown in Figure 81 and Figure 82 respectively. There is no Request Data included in a Management Endpoint Buffer Read command. The NVMe Management Response field is reserved.

Figure 80: Management Endpoint Buffer Read Response Data



If the Data Offset (DOFST) field is greater than or equal to the size of the Management Endpoint Buffer, then the Management Endpoint responds with an Invalid Parameter Error Response. The parameter with the error in this case is the DOFST field. If the DOFST field is less than the size of the Management Endpoint Buffer and the sum of the DOFST and DLEN fields is greater than or equal to size of the Management Endpoint Buffer, then the Management Endpoint responds with an Invalid Parameter Error Response. The parameter with the error in this case is the DLEN field.

When an attempt is made to read Management Endpoint Buffer contents that were zeroed due to a sanitize operation, then the Management Endpoint responds with a Response Message Status of Management Endpoint Buffer Cleared Due to Sanitize.

Figure 81: Management Endpoint Buffer Read – NVMe Management Dword 0

Bit	Description
31:00	Data Offset (DOFST): This field specifies the starting offset, in bytes, into the Management Endpoint Buffer.

Figure 82: Management Endpoint Buffer Read – NVMe Management Dword 1

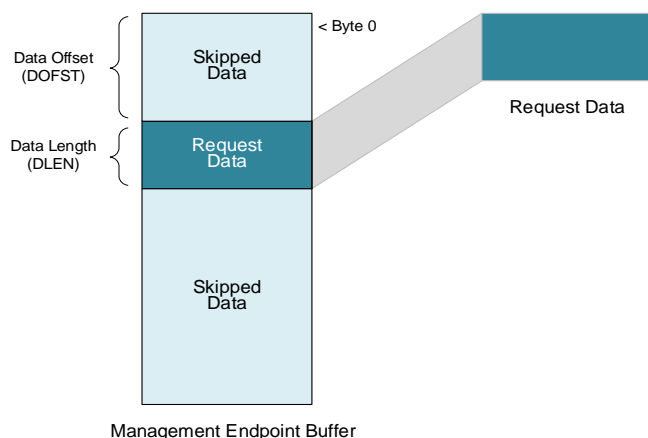
Bit	Description
31:16	Reserved
15:00	Data Length (DLEN): This field specifies the length, in bytes, to be transferred from the Management Endpoint Buffer starting at the byte offset specified by DOFST and returned in the Response Data. Specifying a DLEN field value that is greater than the maximum supported Response Data size results in an Invalid Parameter Error Response. A Data Length value of 0h and no data is valid. The Management Endpoint responds with a Success Response and no Response Data.

5.5 Management Endpoint Buffer Write

The Management Endpoint Buffer Write command allows the Management Controller to update the contents of the optional Management Endpoint Buffer. The data used to update the Management Endpoint Buffer is transferred in the Request Data included in a Management Endpoint Buffer Write command.

The command uses NVMe Management Dwords 0 and 1. The format of the NVMe Management Dwords 0 and 1 are shown in Figure 84 and Figure 85 respectively. The NVMe Management Response field is reserved and there is no Response Data.

Figure 83: Management Endpoint Buffer Write Request Data



If the Data Offset (DOFST) field is greater than or equal to the size of the Management Endpoint Buffer, then the Management Endpoint responds with an Invalid Parameter Error Response. The parameter with the error in this case is the DOFST field. If the DOFST field is less than the size of the Management Endpoint Buffer and the sum of the DOFST and DLEN fields is greater than or equal to size of the Management Endpoint Buffer, then the Management Endpoint responds with an Invalid Parameter Error Response. The parameter with the error in this case is the DLEN field.

Figure 84: Management Endpoint Buffer Write – NVMe Management Dword 0

Bit	Description
31:00	Data Offset (DOFST): This field specifies the starting offset, in bytes, into the Management Endpoint Buffer.

Figure 85: Management Endpoint Buffer Write – NVMe Management Dword 1

Bit	Description
31:16	Reserved
15:00	Data Length (DLEN): This field specifies the length, in bytes, to be transferred from the Request Data to the Management Endpoint Buffer starting at the byte offset specified by DOFST. Specifying a DLEN field value that is greater than the maximum supported Response Data size results in an Invalid Parameter Error Response. A Data Length value of 0h specifies that no data shall be transferred. This condition shall not be considered an error.

5.6 NVM Subsystem Health Status Poll

The NVM Subsystem Health Status Poll command is used to efficiently determine changes in health status attributes associated with the NVM Subsystem.

The NVM Subsystem Health Status Poll command operates independently in the out-of-band mechanism and the in-band tunneling mechanism.

An NVMe Storage Device or NVMe Enclosure supporting the NVM Subsystem Health Status Poll command in the out-of-band mechanism shall have an independent copy of the NVM Subsystem Health Data Structure (refer to Figure 87) dedicated to the out-of-band mechanism. In the out-of-band mechanism, an NVM Subsystem Health Status Poll command only applies to the copy of the NVM Subsystem Health Data Structure dedicated to the out-of-band mechanism.

An NVMe Storage Device or NVMe Enclosure supporting the NVM Subsystem Health Status Poll command in the in-band tunneling mechanism shall have an independent copy of the NVM Subsystem Health Data Structure dedicated to the in-band tunneling mechanism. In the in-band tunneling mechanism, an NVM Subsystem Health Status Poll command only applies to the copy of the NVM Subsystem Health Data Structure dedicated to the in-band tunneling mechanism.

The NVM Subsystem Health Status Poll command uses NVMe Management Dword 1 as shown in Figure 86.

Figure 86: NVM Subsystem Health Status Poll - NVMe Management Dword 1

Bit	Description
31	Clear Status (CS): When this bit is set to '1', the state of reported Composite Controller Status is cleared.
30:00	Reserved

All other command specific fields are reserved.

The NVM Subsystem Health Data Structure, shown in Figure 87, is returned in the Response Data of a Successful Response Message. NVM Subsystem Health Status Poll Command responses do not use the NVMe Management Response field and this field is reserved. The Response Data field contains the NVM Subsystem Health Data Structure and is always the size of the NVM Subsystem Health Data Structure.

Figure 87: NVM Subsystem Health Data Structure (NSHDS)

Byte	Description																
0	<p>NVM Subsystem Status (NSS): This field indicates the status of the NVM Subsystem.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:6</td> <td>Reserved</td> </tr> <tr> <td>5</td> <td>Drive Functional (DF): This bit is set to '1' to indicate an NVM Subsystem is functional. If cleared to '0', then there is an unrecoverable failure detected in the NVM Subsystem.</td> </tr> <tr> <td>4</td> <td>Reset Not Required (RNR): This bit is set to '1' to indicate the NVM Subsystem does not need a reset to resume normal operation. If cleared to '0', then the NVM Subsystem has experienced an error that prevents continued normal operation. A Controller Level Reset is required to resume normal operation.</td> </tr> <tr> <td>3</td> <td>Port 0 PCIe Link Active (P0LA): This bit is set to '1' to indicate the first port's PCIe link is up (i.e., the Data Link Control and Management State Machine is in the DL_Active state). If cleared to '0', then the PCIe link is down.</td> </tr> <tr> <td>2</td> <td>Port 1 PCIe Link Active (P1LA): This bit is set to '1' to indicate the second port's PCIe link is up. If cleared to '0', then the second port's PCIe link is down or not present.</td> </tr> <tr> <td>1:0</td> <td>Reserved</td> </tr> </tbody> </table>	Bit	Description	7:6	Reserved	5	Drive Functional (DF): This bit is set to '1' to indicate an NVM Subsystem is functional. If cleared to '0', then there is an unrecoverable failure detected in the NVM Subsystem.	4	Reset Not Required (RNR): This bit is set to '1' to indicate the NVM Subsystem does not need a reset to resume normal operation. If cleared to '0', then the NVM Subsystem has experienced an error that prevents continued normal operation. A Controller Level Reset is required to resume normal operation.	3	Port 0 PCIe Link Active (P0LA): This bit is set to '1' to indicate the first port's PCIe link is up (i.e., the Data Link Control and Management State Machine is in the DL_Active state). If cleared to '0', then the PCIe link is down.	2	Port 1 PCIe Link Active (P1LA): This bit is set to '1' to indicate the second port's PCIe link is up. If cleared to '0', then the second port's PCIe link is down or not present.	1:0	Reserved		
Bit	Description																
7:6	Reserved																
5	Drive Functional (DF): This bit is set to '1' to indicate an NVM Subsystem is functional. If cleared to '0', then there is an unrecoverable failure detected in the NVM Subsystem.																
4	Reset Not Required (RNR): This bit is set to '1' to indicate the NVM Subsystem does not need a reset to resume normal operation. If cleared to '0', then the NVM Subsystem has experienced an error that prevents continued normal operation. A Controller Level Reset is required to resume normal operation.																
3	Port 0 PCIe Link Active (P0LA): This bit is set to '1' to indicate the first port's PCIe link is up (i.e., the Data Link Control and Management State Machine is in the DL_Active state). If cleared to '0', then the PCIe link is down.																
2	Port 1 PCIe Link Active (P1LA): This bit is set to '1' to indicate the second port's PCIe link is up. If cleared to '0', then the second port's PCIe link is down or not present.																
1:0	Reserved																
1	<p>Smart Warnings (SW): This field contains the Critical Warning field (byte 0) of the NVMe SMART / Health Information log. Each bit in this field is inverted from the NVM Express specification definition (i.e., the management interface shall indicate a '0' value while the corresponding bit is set to '1' in the log page). Refer to the NVM Express specification for bit definitions.</p> <p>If there are multiple Controllers in the NVM Subsystem, the Responder shall combine the Critical Warning field from every Controller in the NVM Subsystem such that a bit in this field is:</p> <ul style="list-style-type: none"> • Cleared to '0' if any Controller in the NVM Subsystem indicates a critical warning for that corresponding bit; or • Set to '1' if all Controllers in the NVM Subsystem do not indicate a critical warning for the corresponding bit. 																
2	<p>Composite Temperature (CTEMP): This field indicates the current temperature in degrees Celsius. If a temperature value is reported, it should be the same temperature as the Composite Temperature from the SMART log of hottest Controller in the NVM Subsystem. The reported temperature range is vendor specific and shall not exceed the range -60 °C to +127 °C. The 8-bit format of the data is shown below.</p> <p>This field should not report a temperature that is older than 1 s. If recent data is not available, the Responder should indicate a value of 80h for this field.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h to 7Eh</td> <td>Temperature is measured in degrees Celsius (0 °C to 126 °C)</td> </tr> <tr> <td>7Fh</td> <td>127 °C or higher</td> </tr> <tr> <td>80h</td> <td>No temperature data or temperature data is more the 5 s old.</td> </tr> <tr> <td>81h</td> <td>Temperature sensor failure</td> </tr> <tr> <td>82h to C3h</td> <td>Reserved</td> </tr> <tr> <td>C4h</td> <td>Temperature is -60 °C or lower</td> </tr> <tr> <td>C5h to FFh</td> <td>Temperature measured in degrees Celsius is represented in two's complement (-1 °C to -59 °C)</td> </tr> </tbody> </table>	Value	Description	00h to 7Eh	Temperature is measured in degrees Celsius (0 °C to 126 °C)	7Fh	127 °C or higher	80h	No temperature data or temperature data is more the 5 s old.	81h	Temperature sensor failure	82h to C3h	Reserved	C4h	Temperature is -60 °C or lower	C5h to FFh	Temperature measured in degrees Celsius is represented in two's complement (-1 °C to -59 °C)
Value	Description																
00h to 7Eh	Temperature is measured in degrees Celsius (0 °C to 126 °C)																
7Fh	127 °C or higher																
80h	No temperature data or temperature data is more the 5 s old.																
81h	Temperature sensor failure																
82h to C3h	Reserved																
C4h	Temperature is -60 °C or lower																
C5h to FFh	Temperature measured in degrees Celsius is represented in two's complement (-1 °C to -59 °C)																
3	<p>Percentage Drive Life Used (PDLU): Contains a vendor specific estimate of the percentage of NVM Subsystem NVM life used based on the actual usage and the manufacturer's prediction of NVM life. If an NVM Subsystem has multiple Controllers, the highest value is returned. A value of 100 indicates that the estimated endurance of the NVM in the NVM Subsystem has been consumed but may not indicate an NVM Subsystem failure. The value is allowed to exceed 100. Percentages greater than 254 shall be represented as 255. This value should be updated once per power-on hour and equal the Percentage Used value in the NVMe SMART Health Log Page.</p>																

Figure 87: NVM Subsystem Health Data Structure (NSHDS)

Byte	Description																																													
5:4	<p>Composite Controller Status (CCS): This field reports the composite status of all Controllers in the NVM Subsystem.</p> <p>The bits in this field are cleared after the NVM Subsystem Health Data Structure (refer to Figure 87) is returned in a Success Response associated with an NVM Subsystem Health Status Poll command where the Clear Status bit set. A Configuration Set command that selects Health Status Change may be used to clear selected bits to '0'.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Reset</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:13</td> <td>0</td> <td>Reserved</td> </tr> <tr> <td>12</td> <td>0</td> <td>Critical Warning (CWARN): This bit is set to '1' when the Critical Warning bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</td> </tr> <tr> <td>11</td> <td>0</td> <td>Available Spare (SPARE): This bit is set to '1' when the Available Spare bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</td> </tr> <tr> <td>10</td> <td>0</td> <td>Percentage Used (PDLU): This bit is set to '1' when the Percentage Used field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</td> </tr> <tr> <td>09</td> <td>0</td> <td>Composite Temperature Change (CTEMP): This bit is set to '1' when the Composite Temperature field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</td> </tr> <tr> <td>08</td> <td>Hwlnit</td> <td>Controller Status Change (CSTS): This bit is set to '1' when the Controller Status field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</td> </tr> <tr> <td>07</td> <td>Hwlnit</td> <td>Firmware Activated (FA): This bit is set to '1' when the Firmware Activated bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</td> </tr> <tr> <td>06</td> <td>0</td> <td>Namespace Attribute Changed (NAC): This bit is set to '1' when the Namespace Attribute Changed bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</td> </tr> <tr> <td>05</td> <td>0</td> <td>Controller Enable Change Occurred (CECO): This bit is set to '1' when the Controller Enable Change Occurred bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</td> </tr> <tr> <td>04</td> <td>Hwlnit</td> <td>NVM Subsystem Reset Occurred (NSSRO): This bit is set to '1' when the value of the NVM Subsystem Reset Occurred (CSTS.NSSRO) bit transitions from a '0' to a '1' in one or more Controllers in the NVM Subsystem.</td> </tr> <tr> <td>03</td> <td>0</td> <td>Reserved</td> </tr> <tr> <td>02</td> <td>0</td> <td>Shutdown Status (SHST): This bit is set to '1' when the Shutdown Status bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</td> </tr> <tr> <td>01</td> <td>Hwlnit</td> <td>Controller Fatal Status (CFS): This bit is set to '1' when the Controller Fatal Status bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</td> </tr> <tr> <td>00</td> <td>0</td> <td>Ready (RDY): This bit is set to '1' when the Ready bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</td> </tr> </tbody> </table>	Bit	Reset	Description	15:13	0	Reserved	12	0	Critical Warning (CWARN): This bit is set to '1' when the Critical Warning bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.	11	0	Available Spare (SPARE): This bit is set to '1' when the Available Spare bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.	10	0	Percentage Used (PDLU): This bit is set to '1' when the Percentage Used field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.	09	0	Composite Temperature Change (CTEMP): This bit is set to '1' when the Composite Temperature field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.	08	Hwlnit	Controller Status Change (CSTS): This bit is set to '1' when the Controller Status field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.	07	Hwlnit	Firmware Activated (FA): This bit is set to '1' when the Firmware Activated bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.	06	0	Namespace Attribute Changed (NAC): This bit is set to '1' when the Namespace Attribute Changed bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.	05	0	Controller Enable Change Occurred (CECO): This bit is set to '1' when the Controller Enable Change Occurred bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.	04	Hwlnit	NVM Subsystem Reset Occurred (NSSRO): This bit is set to '1' when the value of the NVM Subsystem Reset Occurred (CSTS.NSSRO) bit transitions from a '0' to a '1' in one or more Controllers in the NVM Subsystem.	03	0	Reserved	02	0	Shutdown Status (SHST): This bit is set to '1' when the Shutdown Status bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.	01	Hwlnit	Controller Fatal Status (CFS): This bit is set to '1' when the Controller Fatal Status bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.	00	0	Ready (RDY): This bit is set to '1' when the Ready bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.
	Bit	Reset	Description																																											
	15:13	0	Reserved																																											
	12	0	Critical Warning (CWARN): This bit is set to '1' when the Critical Warning bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.																																											
	11	0	Available Spare (SPARE): This bit is set to '1' when the Available Spare bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.																																											
	10	0	Percentage Used (PDLU): This bit is set to '1' when the Percentage Used field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.																																											
	09	0	Composite Temperature Change (CTEMP): This bit is set to '1' when the Composite Temperature field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.																																											
	08	Hwlnit	Controller Status Change (CSTS): This bit is set to '1' when the Controller Status field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.																																											
	07	Hwlnit	Firmware Activated (FA): This bit is set to '1' when the Firmware Activated bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.																																											
	06	0	Namespace Attribute Changed (NAC): This bit is set to '1' when the Namespace Attribute Changed bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.																																											
	05	0	Controller Enable Change Occurred (CECO): This bit is set to '1' when the Controller Enable Change Occurred bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.																																											
	04	Hwlnit	NVM Subsystem Reset Occurred (NSSRO): This bit is set to '1' when the value of the NVM Subsystem Reset Occurred (CSTS.NSSRO) bit transitions from a '0' to a '1' in one or more Controllers in the NVM Subsystem.																																											
	03	0	Reserved																																											
	02	0	Shutdown Status (SHST): This bit is set to '1' when the Shutdown Status bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.																																											
01	Hwlnit	Controller Fatal Status (CFS): This bit is set to '1' when the Controller Fatal Status bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.																																												
00	0	Ready (RDY): This bit is set to '1' when the Ready bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.																																												
7:6	Reserved																																													

5.7 Read NVMe-MI Data Structure

The Read NVMe-MI Data Structure command requests data that describes information about the NVM Subsystem, the Management Endpoint, or the NVMe Controllers.

The command uses NVMe Management Dword 0. The format of NVMe Management Dword 0 is shown in Figure 88. NVMe Management Dword 1 is reserved. There is no Request Data included in a Read NVMe-MI Data Structure command.

Figure 88: Read NVMe-MI Data Structure – NVMe Management Dword 0

Bit	Description																
31:24	Data Structure Type (DTYP): This field specifies the data structure to return.																
	<table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>NVM Subsystem Information</td> </tr> <tr> <td>01h</td> <td>Port Information</td> </tr> <tr> <td>02h</td> <td>Controller List</td> </tr> <tr> <td>03h</td> <td>Controller Information</td> </tr> <tr> <td>04h</td> <td>Optional Commands Supported</td> </tr> <tr> <td>05h</td> <td>Management Endpoint Buffer Command Support List</td> </tr> <tr> <td>06h to FFh</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Definition	00h	NVM Subsystem Information	01h	Port Information	02h	Controller List	03h	Controller Information	04h	Optional Commands Supported	05h	Management Endpoint Buffer Command Support List	06h to FFh	Reserved
	Value	Definition															
	00h	NVM Subsystem Information															
	01h	Port Information															
	02h	Controller List															
	03h	Controller Information															
	04h	Optional Commands Supported															
05h	Management Endpoint Buffer Command Support List																
06h to FFh	Reserved																
23:16	Port Identifier (PORTID): This field contains the identifier of the port whose data structure is returned.																
	If the DTYP field value corresponds to Port Information, then this field contains the Port Identifier whose information is requested.																
	If the DTYP field value corresponds to Management Endpoint Buffer Command Support List, then this field contains the Port Identifier whose information is requested.																
15:00	Controller Identifier (CTRLID): This field contains the Controller identifier whose data structure is returned.																
	If the DTYP field value corresponds to Controller List or Controller Information, then this field contains the Controller identifier in the NVM Subsystem whose information is requested.																
	For all other values of the DTYP field, this field is reserved.																

Upon successful completion of the Read NVMe-MI Data Structure, the NVMe Management Response field is shown in Figure 89 and the specified data structure is returned in the Response Data.

Figure 89: Read NVMe-MI Data Structure – NVMe Management Response

Bit	Description
23:16	Reserved
15:00	Response Data Length: The length, in bytes, of the Response Data field in this Response Message.

The NVM Subsystem Information data structure contains information about the NVM Subsystem. The Port Identifier and Controller Identifier fields are reserved. The format is shown in Figure 90.

Figure 90: NVM Subsystem Information Data Structure

Byte	Description
00	Number of Ports (NUMP): This field specifies the maximum number of ports of any type supported by the NVM Subsystem. This is a 0's based value.

Figure 90: NVM Subsystem Information Data Structure

Byte	Description
01	NVMe-MI Major Version Number (MJR): This field shall be set to 1h to indicate the major version number of this specification.
02	NVMe-MI Minor Version Number (MNR): This field shall be set to 1h to indicate the minor version number of this specification.
31:03	Reserved

The Port Information data structure contains information about a port within the NVM Subsystem. The Port Identifier specifies the port. The Controller Identifier fields are reserved. The format is shown in Figure 91.

Figure 91: Port Information Data Structure

Byte	Description										
00	<p>Port Type: Specifies the port type.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Inactive</td> </tr> <tr> <td>1h</td> <td>PCIe</td> </tr> <tr> <td>2h</td> <td>SMBus</td> </tr> <tr> <td>3h to FFh</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Definition	0h	Inactive	1h	PCIe	2h	SMBus	3h to FFh	Reserved
Value	Definition										
0h	Inactive										
1h	PCIe										
2h	SMBus										
3h to FFh	Reserved										
01	Reserved										
03:02	<p>Maximum MCTP Transmission Unit Size: The maximum MCTP Transmission Unit size the port is capable of sending and receiving.</p> <p>If the port does not support MCTP, then this field shall be cleared to 0h.</p> <p>If the port type is PCIe and the port supports MCTP, then this field shall be set to a value between 64 bytes and the PCIe Max Payload Size supported minus 4, inclusive. All PCIe ports within an NVM Subsystem should report the same value in this field.</p> <p>If the port type is SMBus and the port supports MCTP, then this field shall be set to a value between 64 bytes and 250 bytes, inclusive.</p>										
07:04	<p>This field specifies the size of the Management Endpoint Buffer in bytes when a Management Endpoint Buffer is supported.</p> <p>A value of 0h in this field indicates that the Management Endpoint does not support a Management Endpoint Buffer.</p>										
31:08	Port Type Specific (refer to Figure 92 and Figure 93)										

Figure 92: PCIe Port Specific Data

Byte	Description																
08	<p>PCIe Maximum Payload Size: This field indicates the Max Payload Size for the specified PCIe port. If the link is not active, this field should be cleared to 0h.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>128 bytes</td> </tr> <tr> <td>1h</td> <td>256 bytes</td> </tr> <tr> <td>2h</td> <td>512 bytes</td> </tr> <tr> <td>3h</td> <td>1 KiB</td> </tr> <tr> <td>4h</td> <td>2 KiB</td> </tr> <tr> <td>5h</td> <td>4 KiB</td> </tr> <tr> <td>6h to FFh</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Definition	0h	128 bytes	1h	256 bytes	2h	512 bytes	3h	1 KiB	4h	2 KiB	5h	4 KiB	6h to FFh	Reserved
Value	Definition																
0h	128 bytes																
1h	256 bytes																
2h	512 bytes																
3h	1 KiB																
4h	2 KiB																
5h	4 KiB																
6h to FFh	Reserved																

Figure 92: PCIe Port Specific Data

Byte	Description																														
09	<p>PCIe Supported Link Speeds Vector: This field indicates the Supported Link Speeds for the specified PCIe port.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:4</td> <td>Reserved</td> </tr> <tr> <td>3</td> <td>This bit shall be set to '1' if the link supports 16.0 GT/s.</td> </tr> <tr> <td>2</td> <td>This bit shall be set to '1' if the link supports 8.0 GT/s.</td> </tr> <tr> <td>1</td> <td>This bit shall be set to '1' if the link supports 5.0 GT/s.</td> </tr> <tr> <td>0</td> <td>This bit shall be set to '1' if the link supports 2.5 GT/s.</td> </tr> </tbody> </table>	Bit	Description	7:4	Reserved	3	This bit shall be set to '1' if the link supports 16.0 GT/s.	2	This bit shall be set to '1' if the link supports 8.0 GT/s.	1	This bit shall be set to '1' if the link supports 5.0 GT/s.	0	This bit shall be set to '1' if the link supports 2.5 GT/s.																		
Bit	Description																														
7:4	Reserved																														
3	This bit shall be set to '1' if the link supports 16.0 GT/s.																														
2	This bit shall be set to '1' if the link supports 8.0 GT/s.																														
1	This bit shall be set to '1' if the link supports 5.0 GT/s.																														
0	This bit shall be set to '1' if the link supports 2.5 GT/s.																														
10	<p>PCIe Current Link Speed: The port's PCIe negotiated link speed using the same encoding as the PCIe Supported Link Speed Vector field. A value of 0h in this field indicates the PCIe Link is not available.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Link not active</td> </tr> <tr> <td>1h</td> <td>The current link speed is the speed indicated in the supported link speed bit 0.</td> </tr> <tr> <td>2h</td> <td>The current link speed is the speed indicated in the supported link speed bit 1.</td> </tr> <tr> <td>3h</td> <td>The current link speed is the speed indicated in the supported link speed bit 2.</td> </tr> <tr> <td>4h</td> <td>The current link speed is the speed indicated in the supported link speed bit 3.</td> </tr> <tr> <td>5h</td> <td>The current link speed is the speed indicated in the supported link speed bit 4.</td> </tr> <tr> <td>6h</td> <td>The current link speed is the speed indicated in the supported link speed bit 5.</td> </tr> <tr> <td>7h</td> <td>The current link speed is the speed indicated in the supported link speed bit 6.</td> </tr> <tr> <td>8h to FFh</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Definition	0h	Link not active	1h	The current link speed is the speed indicated in the supported link speed bit 0.	2h	The current link speed is the speed indicated in the supported link speed bit 1.	3h	The current link speed is the speed indicated in the supported link speed bit 2.	4h	The current link speed is the speed indicated in the supported link speed bit 3.	5h	The current link speed is the speed indicated in the supported link speed bit 4.	6h	The current link speed is the speed indicated in the supported link speed bit 5.	7h	The current link speed is the speed indicated in the supported link speed bit 6.	8h to FFh	Reserved										
Value	Definition																														
0h	Link not active																														
1h	The current link speed is the speed indicated in the supported link speed bit 0.																														
2h	The current link speed is the speed indicated in the supported link speed bit 1.																														
3h	The current link speed is the speed indicated in the supported link speed bit 2.																														
4h	The current link speed is the speed indicated in the supported link speed bit 3.																														
5h	The current link speed is the speed indicated in the supported link speed bit 4.																														
6h	The current link speed is the speed indicated in the supported link speed bit 5.																														
7h	The current link speed is the speed indicated in the supported link speed bit 6.																														
8h to FFh	Reserved																														
11	<p>PCIe Maximum Link Width: The maximum PCIe link width for this NVM Subsystem port. This is the expected negotiated link width that the port link trains to if the platform supports it. A Requester may compare this value with the PCIe Negotiated Link Width to determine if there has been a PCIe link training issue.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Reserved</td> </tr> <tr> <td>1</td> <td>PCIe x1</td> </tr> <tr> <td>2</td> <td>PCIe x2</td> </tr> <tr> <td>3</td> <td>Reserved</td> </tr> <tr> <td>4</td> <td>PCIe x4</td> </tr> <tr> <td>5 to 7</td> <td>Reserved</td> </tr> <tr> <td>8</td> <td>PCIe x8</td> </tr> <tr> <td>9 to 11</td> <td>Reserved</td> </tr> <tr> <td>12</td> <td>PCIe x12</td> </tr> <tr> <td>13 to 15</td> <td>Reserved</td> </tr> <tr> <td>16</td> <td>PCIe x16</td> </tr> <tr> <td>17 to 31</td> <td>Reserved</td> </tr> <tr> <td>32</td> <td>PCIe x32</td> </tr> <tr> <td>33 to 255</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Definition	0	Reserved	1	PCIe x1	2	PCIe x2	3	Reserved	4	PCIe x4	5 to 7	Reserved	8	PCIe x8	9 to 11	Reserved	12	PCIe x12	13 to 15	Reserved	16	PCIe x16	17 to 31	Reserved	32	PCIe x32	33 to 255	Reserved
Value	Definition																														
0	Reserved																														
1	PCIe x1																														
2	PCIe x2																														
3	Reserved																														
4	PCIe x4																														
5 to 7	Reserved																														
8	PCIe x8																														
9 to 11	Reserved																														
12	PCIe x12																														
13 to 15	Reserved																														
16	PCIe x16																														
17 to 31	Reserved																														
32	PCIe x32																														
33 to 255	Reserved																														

Figure 92: PCIe Port Specific Data

Byte	Description																														
12	<p>PCIe Negotiated Link Width: The negotiated PCIe link width for this port.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr><td>0</td><td>Link not active</td></tr> <tr><td>1</td><td>PCIe x1</td></tr> <tr><td>2</td><td>PCIe x2</td></tr> <tr><td>3</td><td>Reserved</td></tr> <tr><td>4</td><td>PCIe x4</td></tr> <tr><td>5 to 7</td><td>Reserved</td></tr> <tr><td>8</td><td>PCIe x8</td></tr> <tr><td>9 to 11</td><td>Reserved</td></tr> <tr><td>12</td><td>PCIe x12</td></tr> <tr><td>13 to 15</td><td>Reserved</td></tr> <tr><td>16</td><td>PCIe x16</td></tr> <tr><td>17 to 31</td><td>Reserved</td></tr> <tr><td>32</td><td>PCIe x32</td></tr> <tr><td>33 to 255</td><td>Reserved</td></tr> </tbody> </table>	Value	Definition	0	Link not active	1	PCIe x1	2	PCIe x2	3	Reserved	4	PCIe x4	5 to 7	Reserved	8	PCIe x8	9 to 11	Reserved	12	PCIe x12	13 to 15	Reserved	16	PCIe x16	17 to 31	Reserved	32	PCIe x32	33 to 255	Reserved
Value	Definition																														
0	Link not active																														
1	PCIe x1																														
2	PCIe x2																														
3	Reserved																														
4	PCIe x4																														
5 to 7	Reserved																														
8	PCIe x8																														
9 to 11	Reserved																														
12	PCIe x12																														
13 to 15	Reserved																														
16	PCIe x16																														
17 to 31	Reserved																														
32	PCIe x32																														
33 to 255	Reserved																														
13	PCIe Port Number: This field contains the PCIe port number. This is the same value as that reported in the Port Number field in the PCIe Link Capabilities Register.																														
31:14	Reserved																														

Figure 93: SMBus Port Specific Data

Byte	Description												
08	Current VPD SMBus/I2C Address: This field indicates the current VPD SMBus/I2C address. A value of 0h indicates there is no VPD.												
09	<p>Maximum VPD Access SMBus/I2C Frequency: This field indicates the maximum SMBus/I2C frequency supported on the VPD interface.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr><td>0h</td><td>Not supported</td></tr> <tr><td>1h</td><td>100 kHz</td></tr> <tr><td>2h</td><td>400 kHz</td></tr> <tr><td>3h</td><td>1 MHz</td></tr> <tr><td>4h to FFh</td><td>Reserved</td></tr> </tbody> </table>	Value	Definition	0h	Not supported	1h	100 kHz	2h	400 kHz	3h	1 MHz	4h to FFh	Reserved
Value	Definition												
0h	Not supported												
1h	100 kHz												
2h	400 kHz												
3h	1 MHz												
4h to FFh	Reserved												
10	Current Management Endpoint SMBus/I2C Address: This field indicates the current MCTP SMBus/I2C address. A value of 0h indicates there is no Management Endpoint on this port.												
11	<p>Maximum Management Endpoint SMBus/I2C Frequency: This field indicates the maximum SMBus/I2C frequency supported by the Management Endpoint.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr><td>0h</td><td>Not supported</td></tr> <tr><td>1h</td><td>100 kHz</td></tr> <tr><td>2h</td><td>400 kHz</td></tr> <tr><td>3h</td><td>1 MHz</td></tr> <tr><td>4h to FFh</td><td>Reserved</td></tr> </tbody> </table>	Value	Definition	0h	Not supported	1h	100 kHz	2h	400 kHz	3h	1 MHz	4h to FFh	Reserved
Value	Definition												
0h	Not supported												
1h	100 kHz												
2h	400 kHz												
3h	1 MHz												
4h to FFh	Reserved												
12	NVMe Basic Management: If bit 0 in this field is set to '1', then the port implements the NVMe Basic Management Command. If bit 0 in this field is cleared to '0', then the port does not implement the NVMe Basic Management Command. It is strongly recommended that implementations clear this bit to '0'. The NVMe Basic Management Command is included in Appendix A for information purposes only and is not a part of the standard NVMe-MI protocol. All other bits in this field are reserved.												

Figure 93: SMBus Port Specific Data

Byte	Description
31:13	Reserved

The Controller List data structure contains a list of NVMe Controllers in the NVM Subsystem greater than or equal to the value specified in the Controller Identifier (CTRLID) field. A Controller List may contain up to 2,047 Controller identifiers. Refer to the NVM Express specification for a definition of the Controller List.

Figure 94: Controller Information Data Structure

Byte	Description	
00	Port Identifier (PORTID): This field specifies the PCIe Port Identifier with which the Controller is associated.	
04:01	Reserved	
05	PCIe Routing ID Information (PRII): This field provides additional data about the PCI Express Routing ID (PRI) for the specified Controller.	
		Bit Description
		7:1 Reserved
07:06	PCIe Routing ID (PRI): This field contains the PCIe Routing ID for the specified Controller.	
		Bit Description
		15:08 PCI Bus Number: The Controller's PCI Bus Number.
		07:03 PCI Device Number: The Controller's PCI Device Number.
		02:00 PCI Function Number: The Controller's PCI Function Number.
	Note: For an ARI Device, bits 7:0 represents the (8-bit) Function Number, which replaces the (5-bit) Device Number and (3-bit) Function Number fields above.	
09:08	PCI Vendor ID: The PCI Vendor ID for the specified Controller.	
11:10	PCI Device ID: The PCI Device ID for the specified Controller.	
13:12	PCI Subsystem Vendor ID: The PCI Subsystem Vendor ID for the specified Controller.	
15:14	PCI Subsystem Device ID: The PCI Subsystem Device ID for the specified Controller.	
31:16	Reserved	

The Optionally Supported Command List data structure contains a list of optional commands that a Responder supports. The Optionally Supported Command List data structure may contain up to 2,047 commands, and shall be minimally sized (i.e., if there is one optionally supported command, the data structure is 4 bytes total).

Figure 95: Optionally Supported Command List Data Structure

Byte	Description
01:00	Number of Commands (NUMCMD): This field contains the number of optionally supported commands in the list. A value of 0h indicates there are no commands in the list.
03:02	Command 0 (CMD0): This field contains the Command Type and Opcode for the first optionally supported command or 0h if the list is empty (i.e., no optional commands are supported). Refer to Figure 96.

Figure 95: Optionally Supported Command List Data Structure

Byte	Description
05:04	Command 1 (CMD1): This field contains the Command Type and Opcode for the second optionally supported command, if applicable. Refer to Figure 96.
...	
(N*2 +3): (N*2 + 2)	Command N (CMDN): This field contains the Command Type and Opcode for the N+1 optionally supported command, if applicable. Refer to Figure 96.

Figure 96: Optionally Supported Command Data Structure

Byte	Description									
00	Command Type: This field specifies the command set used by the optionally supported command.									
		<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7</td> <td>Reserved</td> </tr> <tr> <td>6:3</td> <td>NVMe-MI Message Type (NMIMT): This field specifies the NVMe-MI Message Type. Refer to Figure 18.</td> </tr> <tr> <td>2:0</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	7	Reserved	6:3	NVMe-MI Message Type (NMIMT): This field specifies the NVMe-MI Message Type. Refer to Figure 18.	2:0	Reserved
		Bits	Description							
		7	Reserved							
6:3	NVMe-MI Message Type (NMIMT): This field specifies the NVMe-MI Message Type. Refer to Figure 18.									
2:0	Reserved									
01	Opcode: This field specifies the opcode used for the optionally supported command.									

If the Management Endpoint Buffer Size field in the Port Information Data Structure is not 0h, then returning of the Management Endpoint Buffer Command Support List data structure shall be supported by the Management Endpoint. If the Management Endpoint Buffer Size field in the Port Information Data Structure is 0h, then the Data Structure Type value for Management Endpoint Buffer Command Support List is reserved.

The Management Endpoint Buffer Command Support List data structure contains a list of commands that support the use of the Management Endpoint Buffer. The data structure may contain up to 2,047 commands, and shall be minimally sized (i.e., if there is 1 optionally supported command, the data structure is 4 bytes total).

The list of commands that support the Management Endpoint Buffer may be different among Management Endpoints within the NVM Subsystem. The Port Identifier (PORTID) field in NVMe Management Dword 0 of the Read NVMe-MI Data Structure specifies the port of the Management Endpoint whose Management Endpoint Buffer Command Support List data structure is returned.

Figure 97: Management Endpoint Buffer Supported Command List Data Structure

Byte	Description
01:00	Number of Commands (NUMCMD): This field contains the number of commands in the list. A value of 0h indicates there are no commands in the list.
03:02	Command 0 (CMD0): This field contains the Management Endpoint Buffer Supported Command Data Structure (refer to Figure 98) for the first command that supports the use of the Management Endpoint Buffer associated with the Management Endpoint.
05:04	Command 1 (CMD1): This field contains the Management Endpoint Buffer Supported Command Data Structure (refer to Figure 98) for the second command that supports the use of the Management Endpoint Buffer associated with the Management Endpoint.
...	
(N*2 + 3): (N*2 + 2)	Command N (CMDN): This field contains the Management Endpoint Buffer Supported Command Data Structure (refer to Figure 98) for the N+1 command that supports the use of the Management Endpoint Buffer associated with the Management Endpoint.

Figure 98: Management Endpoint Buffer Supported Command Data Structure

Byte	Description								
00	<p>Command Type: This field specifies the command set that supports the Management Endpoint Buffer.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7</td> <td>Reserved</td> </tr> <tr> <td>6:3</td> <td>NVMe-MI Message Type (NMIMT): This field specifies the NVMe-MI Message Type. Refer to Figure 18.</td> </tr> <tr> <td>2:0</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	7	Reserved	6:3	NVMe-MI Message Type (NMIMT): This field specifies the NVMe-MI Message Type. Refer to Figure 18.	2:0	Reserved
Bits	Description								
7	Reserved								
6:3	NVMe-MI Message Type (NMIMT): This field specifies the NVMe-MI Message Type. Refer to Figure 18.								
2:0	Reserved								
01	Opcode: This field specifies the opcode of the command that supports the Management Endpoint Buffer.								

5.8 Reset

The Reset command may be used to initiate a reset.

The Reset command uses NVMe Management Dword 0. The format of NVMe Management Dword 0 is shown in Figure 99. All other command specific fields in the Request Message and Response Message are reserved.

Figure 99: Reset - NVMe Management Dword 0

Bit	Description									
31:24	<p>Reset Type: This field specifies the type of reset to be performed.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>O/M¹</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>O/M²</td> <td>Reset NVM Subsystem</td> </tr> <tr> <td>01h to FFh</td> <td>-</td> <td>Reserved</td> </tr> </tbody> </table>	Value	O/M ¹	Description	00h	O/M ²	Reset NVM Subsystem	01h to FFh	-	Reserved
Value	O/M ¹	Description								
00h	O/M ²	Reset NVM Subsystem								
01h to FFh	-	Reserved								
23:00	Reserved									
<p>NOTES:</p> <ol style="list-style-type: none"> O/M definition: O = Optional, M = Mandatory The Reset Type for Reset NVM Subsystem is required if the NVM Subsystem Reset feature is supported in-band as defined in the NVM Express specification; else, it is optional. 										

When a Reset command is completed successfully, the NVM Subsystem Reset is immediately initiated (refer to section 9.3). No Success Response is transmitted.

5.9 SES Receive

The SES Receive command is used to retrieve SES status type diagnostic pages. Upon successful completion of the SES Receive command, the SES status type diagnostic page is returned in the Response Data.

The SES Receive command uses NVMe Management Dwords 0 and 1. The format of NVMe Management Dword 0 is shown in Figure 100 and the format of NVMe Management Dword 1 is shown in Figure 101. There is no Request Data sent in the Request Message.

The Page Code (PCODE) field specifies the SES status type diagnostic page to be retrieved. Refer to SES-3 for a list and description of SES diagnostic pages. If the PCODE field specifies a reserved value, an

unsupported value, or a value that only corresponds to an SES control type diagnostic page, then the Responder responds with an Invalid Parameter Error Response.

The Allocation Length (ALENGTH) field specifies the maximum length of the Response Data field in the Response Message and is used to limit the maximum amount of SES diagnostic page data that may be returned. The length of the Response Data field shall be the total length of the SES diagnostic page specified by the PCODE field or the number of bytes specified by the ALENGTH field (i.e., the SES diagnostic page is truncated), whichever is less. When the SES diagnostic page is truncated, the value of fields within the SES diagnostic page are not altered to reflect the truncation.

All errors are detected and reported while servicing the SES Receive command and reported via an Error Response. If an invalid field is detected in an SES Receive command, then the Responder responds with an Invalid Parameter Error Response. If a condition occurs that in SES-3 results in a CHECK CONDITION, then the Responder responds with an Error Response. The mapping of Error Response Status values to SES-3 sense keys and additional sense codes is shown in Figure 12.

If the SES Receive command is supported in the out-of-band mechanism, then the Management Endpoint Buffer shall support the use of the Management Endpoint Buffer with SES Receive command and the size of the Management Endpoint Buffer shall be greater than or equal to the maximum supported SES status type diagnostic page. This allows a Requester to retrieve an SES status type diagnostic page whose size exceed the maximum size allowed by one NVMe-MI Message.

The amount of data returned in the Response Data or transferred to the Management Endpoint Buffer is dependent on the SES status diagnostic page that is returned. The Response Data Length field in the NVMe Management Response contains the length of the Response Data.

Figure 100: SES Receive – NVMe Management Dword 0

Bit	Description
31:8	Reserved
07:00	Page Code (PCODE): This field specifies the SES status diagnostic page to be transferred.

Figure 101: SES Receive – NVMe Management Dword 1

Bit	Description
31:16	Reserved
15:00	Allocation Length (ALENGTH): This field specifies the maximum length in bytes of the Response Data field in the Response Message.

Figure 102: SES Receive – NVMe Management Response

Bit	Description
23:16	Reserved
15:00	Response Data Length (RDL): The length, in bytes, of the Response Data field in this Response Message or transferred to the Management Endpoint Buffer.

5.10 SES Send

The SES Send command is used to transfer SES control type diagnostic pages to an SES Enclosure Service Process. Upon successful completion of the SES Send command, the Request Data, containing an SES control type diagnostic page, is transferred by the Request Message or to the Management Endpoint Buffer.

Unlike the SES Receive command that specifies the page code of the SES status diagnostic page being retrieved, the SES Send command specifies the page code of the SES control type diagnostic page that is

being transferred in the SES control type diagnostic page itself. Refer to SES-3 for a list and description of SES control type diagnostic pages. If the page code in the SES control type diagnostic page specifies a reserved value, an unsupported value, or a value that only corresponds to an SES status diagnostic page, then the Responder responds with an Invalid Parameter Error Response.

The SES Send command does not use NVMe Management Dword 0 or the NVMe Management Response field. All of these are reserved.

All errors are detected and reported while processing the SES Send command and reported via an Error Response. If an invalid field is detected in the SES control type diagnostic page data transferred by an SES Send command, then the Responder responds with an Invalid Parameter Error Response. If a condition occurs that in SES-3 results in a CHECK CONDITION, then the Responder responds with an Error Response. The mapping of Response Message Status values to SES-3 sense keys and additional sense codes is shown in Figure 12.

The length in bytes of the Request Data field is specified in the Data Length (DLEN) field in NVMe Management Dword 1. An SES Send command with DLEN equal to 0h and no data is valid, and results in a Success Response. If the DLEN field specifies a value that is greater than PAGE LENGTH field in the SES control type diagnostic page plus four, then the extra data in the Request Data field following the page is ignored. If the DLEN field specifies a value that is less than PAGE LENGTH field in the SES control type diagnostic page plus four, then the page is processed using the data contained in the Request Data field.

If the SES Send command is supported in the out-of-band mechanism, then the Responder shall support the use of the Management Endpoint Buffer with the SES Send command and the size of the Management Endpoint Buffer shall be greater than or equal to the maximum supported SES control type diagnostic page. This allows a Requester to transfer an SES control type diagnostic page whose size exceeds the maximum size allowed by one NVMe-MI Message.

Figure 103: SES Send – NVMe Management Dword 1

Bit	Description
31:16	Reserved
15:00	Data Length (DLEN): This field specifies the Request Data field in bytes.

5.11 VPD Read

The VPD Read command is used to read the Vital Product Data described in section 9.2. Upon successful completion of the VPD Read command, the specified portion of the VPD contents is returned in the Response Data.

The VPD Read command uses NVMe Management Dword 0 and 1. The format of NVMe Management Dwords 0 and 1 are shown in Figure 104 and Figure 105 respectively. There is no Request Data sent in the Request Message.

A VPD Read command with length 0 and no data is valid. The Responder responds with a Success Response and no Response Data. If the Data Length plus Data Offset fields are greater than the size of the VPD, then the Responder does not return the VPD contents and responds with an Invalid Parameter Error Response.

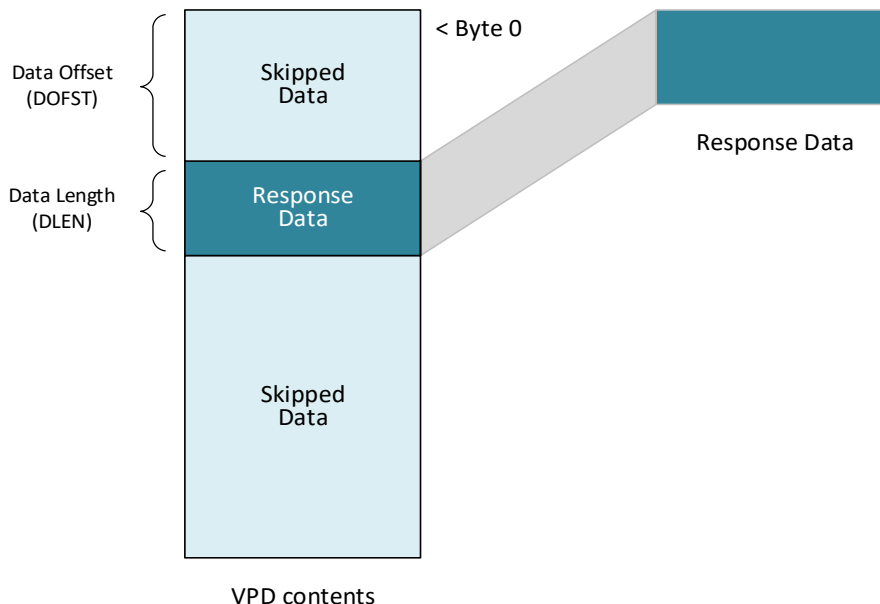
Figure 104: VPD Read NVMe Management Dword 0

Bit	Description
31:16	Reserved
15:00	Data Offset (DOFST): This field specifies the starting offset, in bytes, into the VPD data that is contained in the Response Message.

Figure 105: VPD Read NVMe Management Dword 1

Bit	Description
31:16	Reserved
15:00	Data Length (DLEN): This field specifies the length, in bytes, to be read from the VPD starting at the byte offset specified by DOFST.

Figure 106: VPD Read Response Data



5.12 VPD Write

The VPD Write command is used to update the Vital Product Data described in section 9.2.

After the VPD Write command completes successfully, reading the contents of the FRU Information Device directly or executing a VPD Read command shall return the new VPD contents (i.e., those supplied with the VPD Write command). The data to be written to the VPD is specified in the Request Data field. VPD Write uses NVMe Management Dwords 0 and 1 as shown in Figure 107 and Figure 108.

The VPD contents should be capable of being updated at least 8 times using the VPD Write command¹. If the initial value of the VPD Write Cycles Remaining field is less than 100, then the VPD Write Cycle Remaining Valid bit should be set to '1' (Refer to the VPD Information field in Figure 136). If there is an error preventing update of the VPD contents, then the Responder responds with a Generic Error Response and VPD Writes Exceeded status.

A VPD Write command with length 0h and no data is valid. The Responder responds with a Success Response.

¹ NVM Express Management Interface Specification, Revision 1.0a and prior recommended that VPD contents should be capable of being updated at least 100 times using the VPD Write command.


Figure 107: VPD Write – NVMe Management Dword 0

Bit	Description
31:16	Reserved
15:00	Data Offset (DOFST): This field specifies the starting offset, in bytes, into the VPD data that is written.

Figure 108: VPD Write – NVMe Management Dword 1

Bit	Description
31:16	Reserved
15:00	Data Length (DLEN): This field specifies the length, in bytes, to be written to the VPD starting at the byte offset specified by DOFST.

Figure 109: VPD Write Request Data



The Requester should not read the contents of the VPD while this command is servicing. Reading the contents of the VPD or executing a VPD Read command while a VPD Write command is executing may return incorrect data as a result of the read.

If the Data Length plus Data Offset fields are greater than the size of the VPD, then the Responder does not write to the VPD and responds with an Invalid Parameter Error Response.

6 NVM Express Admin Command Set

The NVM Express Admin Command Set allows NVMe Admin Commands to be issued to any Controller in the NVM Subsystem using the out-of-band mechanism. Figure 110 shows NVM Express Admin Commands that are mandatory, optional, and prohibited for an NVMe Storage Device and an NVMe Enclosure using the out-of-band mechanism. All NVM Express Admin Commands are prohibited using the in-band tunneling mechanism. The commands are defined in the NVM Express specification. If an NVMe Admin Command is issued in a Request Message that is a prohibited command in Figure 110, the Management Endpoint shall return an Invalid Parameter Error Response with Parameter Error Location pointing to the NVMe opcode. Future revisions of this specification may add additional commands to Figure 110. The NVM Express Admin Command Set is only applicable in the out-of-band mechanism and is prohibited in the in-band tunneling mechanism.

Figure 110: List of NVMe Admin Commands Supported using the Out-of-Band Mechanism

Command	NVMe Storage Device O/M/P ¹	NVMe Enclosure O/M/P ¹
Abort	P	P
Asynchronous Event Request	P	P
Create I/O Completion Queue	P	P
Create I/O Submission Queue	P	P
Delete I/O Completion Queue	P	P
Delete I/O Submission Queue	P	P
Device Self-test	O	O
Directive Receive	P	P
Directive Send	P	P
Doorbell Buffer Config	P	P
Firmware Activate/Commit	O	O
Firmware Image Download	O	O
Format NVM	O	P
Get Features	M	O
Get Log Page	M	O
Identify	M	O
Keep Alive	P	P
Namespace Management	O	P
Namespace Attachment	O	P
NVMe-MI Receive	P	P
NVMe-MI Send	P	P
Sanitize	O	O
Security Send	O	P
Security Receive	O	P
Set Features	O	O
Vendor Specific	O	O
Virtualization Management	O	O
NOTES:		
1. O/M/P definition: O = Optional, M = Mandatory, P = Prohibited from being supported. An NVMe Enclosure that is also an NVMe Storage Device (i.e., implements namespaces) shall implement mandatory commands required by either an NVMe Storage Device or an NVMe Enclosure and may implement optional commands allowed by either an NVMe Storage Device or an NVMe Enclosure. Mandatory commands shall be supported if the NVMe Controller specified by the Controller ID field supports the command.		

NVMe Admin Commands over the out-of-band mechanism may interfere with host software. A Management Controller should coordinate with the host or issue only NVMe Admin Commands that do not interfere with host software or in band NVMe commands (e.g., Identify). Coordination between a Management Controller and host is outside the scope of this specification.

NVMe Admin Commands over the out-of-band mechanism may target a controller that is disabled or held in reset by the host. When this occurs, the NVMe Admin Command is processed normally.

The Request Message format for NVMe Admin Commands is shown in Figure 111 and is described Figure 112.

Figure 111: NVMe Admin Command Request Format

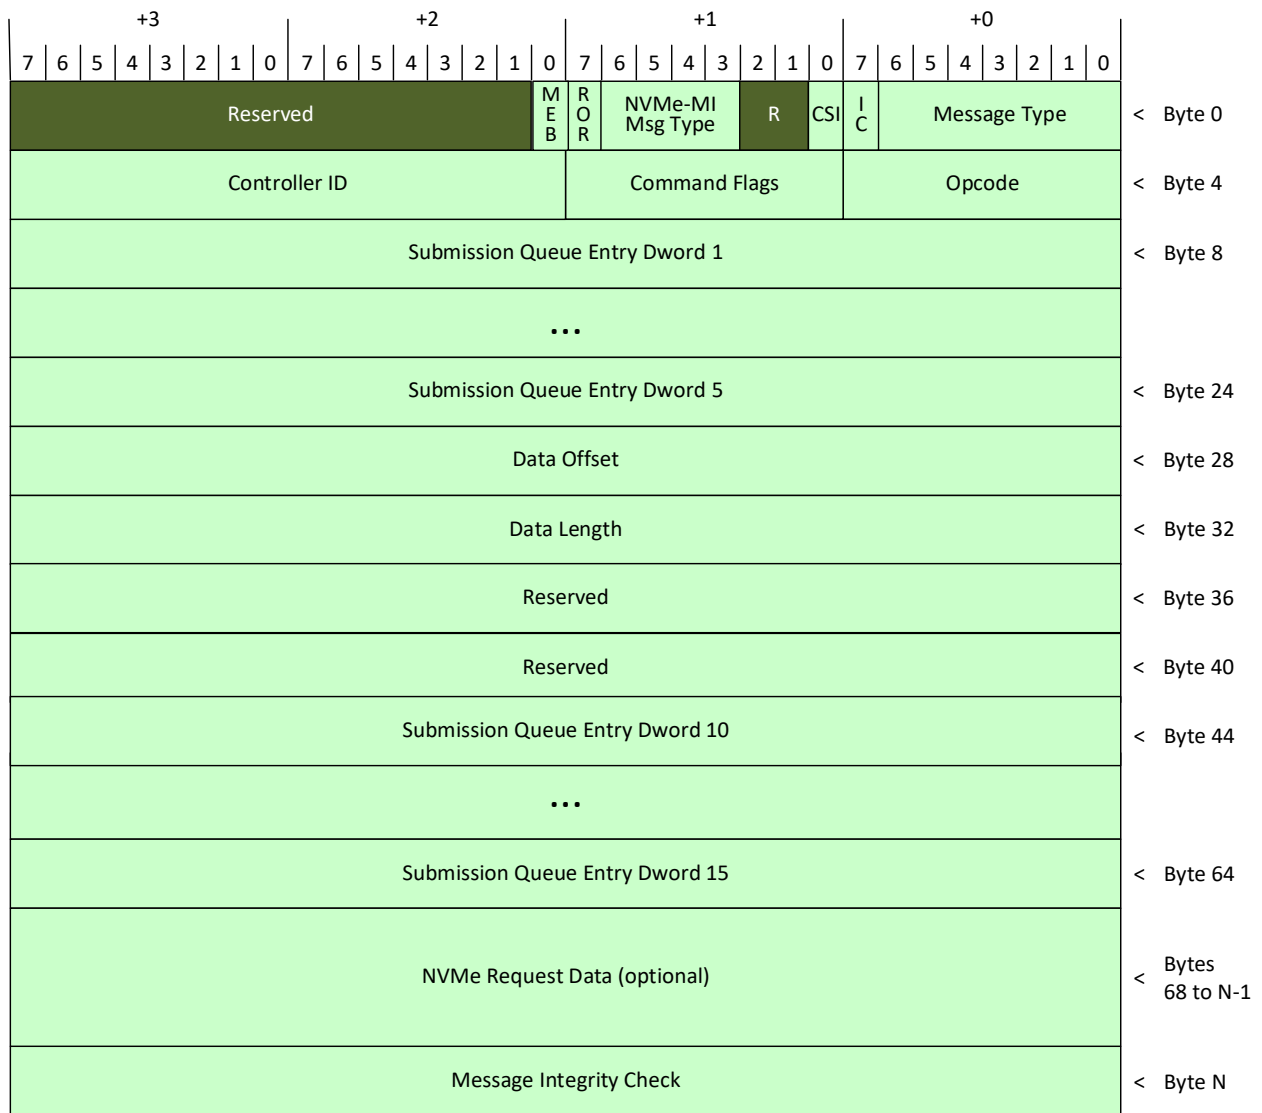


Figure 112: NVMe Admin Command Request Description

Byte	Description
03:00	NVMe-MI Message Header: Refer to section 3.1.
04	Opcode (OPC): This field specifies the opcode of the command to be executed. Refer to the NVM Express specification.
05	Command Flags (CFLGS): This field specifies flags for the command. Bits 7:2 are reserved. Bit 1, if set to '1', then the command contains an offset value in bytes 28-31. If cleared to '0', then the DOFST field shall be cleared to 0h. Bit 0, if set to '1', then the command contains a length value in bytes 32-35. If cleared to '0', then the DLEN field shall be cleared to 0h.
07:06	Controller ID (CTLID): This field specifies the Controller ID of the Controller that this command targets.
11:08	Submission Queue Entry Dword 1 (SQEDW1): Submission Queue Entry Dword 1 as defined in the NVM Express specification.
15:12	Submission Queue Entry Dword 2 (SQEDW2): Submission Queue Entry Dword 2 as defined in the NVM Express specification.
19:16	Submission Queue Entry Dword 3 (SQEDW3): Submission Queue Entry Dword 3 as defined in the NVM Express specification.
23:20	Submission Queue Entry Dword 4 (SQEDW4): Submission Queue Entry Dword 4 as defined in the NVM Express specification.
27:24	Submission Queue Entry Dword 5 (SQEDW5): Submission Queue Entry Dword 5 as defined in the NVM Express specification.
31:28	Data Offset (DOFST): For commands that transmit data from the Management Controller to the Management Endpoint (i.e., the Request Data field in the Request Message has non-zero length) or do not transmit data, this field shall be cleared to 0h. If this field is not 0h, then the Management Endpoint shall return an Invalid Parameter Error Response. For commands that transmit data from the Management Endpoint to the Management Controller (i.e., the Response Data field in the Response Message has non-zero length), this field specifies the starting offset, in bytes, of the portion of data contained in the NVMe Admin Command completion data that is returned starting at byte offset 0h of the Response Data field in the Response Message. Bits 0 and 1 of this field shall be cleared to '0'.
35:32	Data Length (DLEN): For commands that do not transmit data in either the Request Message or Response Message, this field shall be cleared to 0h. If this field is not 0h, then the Management Endpoint shall return an Invalid Parameter Error Response. For commands that transmit data from the Management Controller to the Management Endpoint (i.e., the Request Data field in the Request Message has non-zero length), this field specifies the length, in bytes, of the data contained in the Request Data field in the Request Message. For commands that transmit data from the Management Endpoint to the Management Controller (i.e., the Response Data field in the Response Message has non-zero length), this field specifies the length, in bytes, of the portion of data contained in the NVMe Admin Command completion data that is returned in the Response Data field in the Response Message. Bits 0 and 1 of this field shall be cleared to '0'. This field shall be less than or equal to 4,096.
43:36	Reserved
47:44	Submission Queue Entry Dword 10 (SQEDW10): Submission Queue Entry Dword 10 as defined in the NVM Express specification.
51:48	Submission Queue Entry Dword 11 (SQEDW11): Submission Queue Entry Dword 11 as defined in the NVM Express specification.

Figure 112: NVMe Admin Command Request Description

Byte	Description
55:52	Submission Queue Entry Dword 12 (SQEDW12): Submission Queue Entry Dword 12 as defined in the NVM Express specification.
59:56	Submission Queue Entry Dword 13 (SQEDW13): Submission Queue Entry Dword 13 as defined in the NVM Express specification.
63:60	Submission Queue Entry Dword 14 (SQEDW14): Submission Queue Entry Dword 14 as defined in the NVM Express specification.
67:64	Submission Queue Entry Dword 15 (SQEDW15): Submission Queue Entry Dword 15 as defined in the NVM Express specification.
N-1:68	NVMe Request Data (Optional)
N+3:N	Message Integrity Check (MIC): Refer to section 3.1.

The Response Message contains the corresponding format for NVMe Admin Commands is shown in Figure 113 and is described in Figure 114.

Figure 113: NVMe Admin Command Response Format

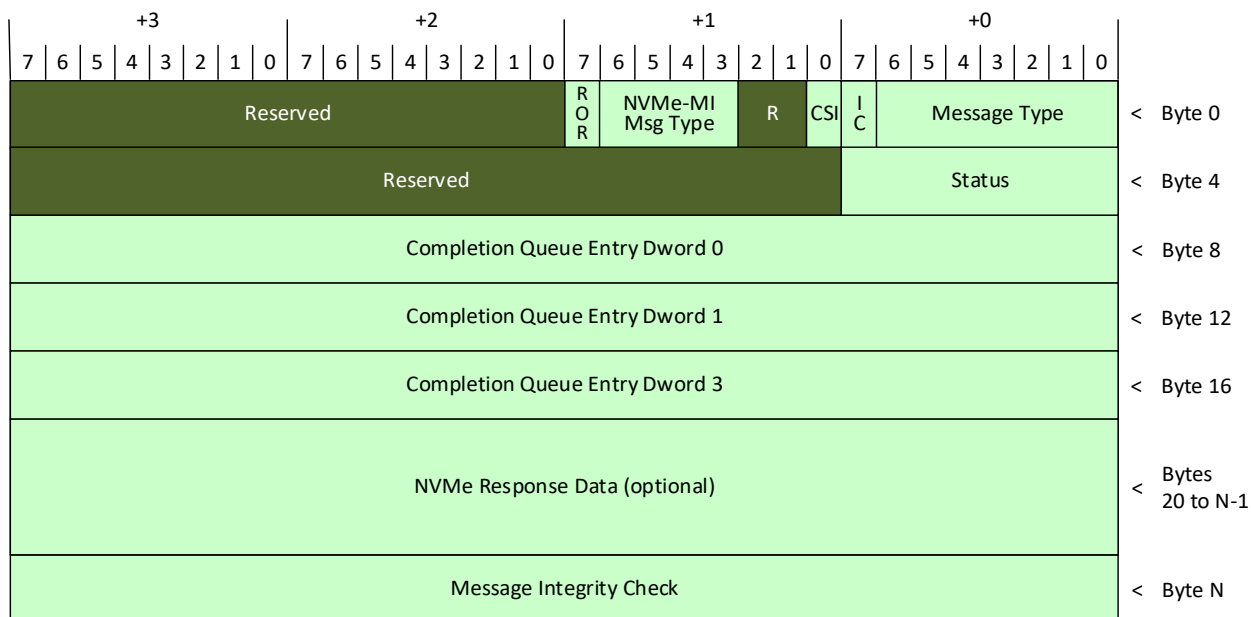


Figure 114: NVMe Admin Command Response Description

Byte	Description
03:00	NVMe-MI Message Header: Refer to section 3.1.
04	Status: This field indicates the status of the NVMe Admin command. Refer to section 4.1.2.
07:05	Reserved
11:08	Completion Queue Entry Dword 0 (CQEDW0): Completion Queue Entry Dword 0 as defined in the NVM Express specification.

Figure 114: NVMe Admin Command Response Description

Byte	Description
15:12	Completion Queue Entry Dword 1 (CQEDW1): Completion Queue Entry Dword 1 as defined in the NVM Express specification.
19:16	Completion Queue Entry Dword 3 (CQEDW3): Completion Queue Entry Dword 3 as defined in the NVM Express specification. The Command ID field shall be cleared to 0h.
N-1:20	NVMe Response Data (Optional)
N+3:N	Message Integrity Check: Refer to section 3.1.

6.1 Request and Response Data

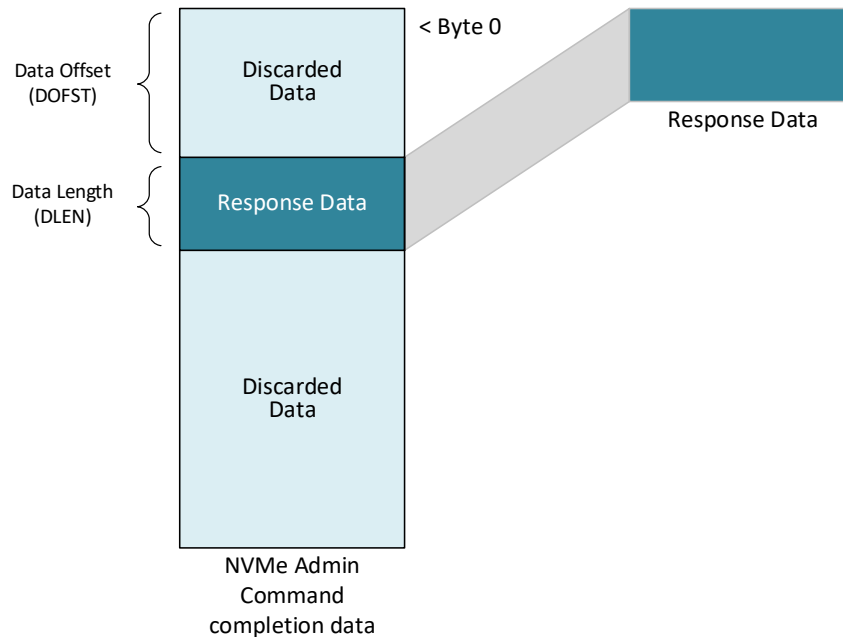
NVMe Admin Commands may contain data as part of the Command Message. This data is passed in the Request Data field instead of using PRP Lists or SGL segments. The PRP Entry 2 (PRP2) and Metadata Pointer (MPTR) fields within the NVMe Admin Commands are reserved.

If there is no data sent with the NVMe Admin Command (e.g., the Data Transfer subfield for the opcode is 00b), then the Data Offset and Data Length fields shall be cleared to 0h.

If there is data sent with the NVMe Admin Command (i.e., the Data Transfer subfield for the opcode is 01b), then the Data Offset field shall be 0h and the Data Length field shall be set to the length of the Request Data required by the command. If the Data Length field does not correspond to the required length, the Management Endpoint shall respond with an Invalid Parameter Error Response.

If there is Response Data expected in the Response Message in the completion of the NVMe Admin Command (i.e., the Data Transfer subfield in the corresponding NVMe Admin Command for the opcode is 10b), then the Data Offset and Data Length fields describe the portion of the NVMe Admin Command completion data that is transferred in the Response Message. Any remaining data not transferred in the Response Message is discarded by the Management Endpoint as shown in Figure 115. If the Data Length plus Data Offset fields are greater than the size of the NVMe Admin Command completion data, the Management Endpoint should respond with an Invalid Parameter Error Response.

Figure 115: NVMe Admin Command Response Data Example



6.2 Status

A Response Message for an NVMe Admin Command may contain two status fields. The first status field, contained in Byte 4 of the Response Message, is defined by this specification, and the second Status Field, if present, is contained in Completion Queue Entry Dword 3 and defined in the NVM Express specification.

An NVMe Admin Command Request Message is well formed if it does not contain any of the following errors:

- Invalid Opcode (e.g., the opcode is not listed in Figure 110);
- Invalid Parameter (e.g., the Controller ID field specifies a Controller ID not implemented in the NVM Subsystem);
- Invalid Command Size (e.g., the Request Message does not contain a complete command); or
- Invalid Command Input Data Size (e.g., the Request Data field is larger than the size specified in the Data Length field).

If the NVMe Admin Command Request Message is well formed, then a Success Response is transmitted. The Success Response contains the status associated with NVMe Admin Command in the Status Field of Completion Queue Entry Dword 3. The Status Field contains any NVM Express specific status codes (e.g., Success or Invalid Parameter).

6.3 Sanitize Operation

Figure 116 specifies the Command Messages allowed during a sanitize operation. Refer to the NVM Express specification for the definition of a sanitize operation.

Figure 116: Command Messages Allowed During Sanitize Operation

Command Set	Command Message	Allowed During Sanitize Operation ¹
Management Interface Command Set	Configuration Get	Yes
	Configuration Set	
	Controller Health Status Poll	
	Management Endpoint Buffer Read	
	Management Endpoint Buffer Write	
	NVM Subsystem Health Status Poll	
	Read NVMe-MI Data Structure	
	Reset	
	SES Receive	
	SES Send	
	VPD Read	
	VPD Write	
	NVMe Admin Command Set ²	
Firmware Activate/Commit		
Firmware Image Download		
Format NVM		
Get Features		
Get Log Page		
Identify		
Namespace Attachment		
Namespace Management		
Sanitize		
Security Receive/Send		
Security Send		
Set Features		
Vendor Specific		
Virtualization Management		
PCIe Command Set	PCIe Configuration Read	Yes
	PCIe Configuration Write	
	PCIe I/O Read	
	PCIe Memory Read	
	PCIe Memory Write	
<p>NOTES:</p> <ol style="list-style-type: none"> 1. Refer to the NVM Express specification for the definition of a sanitize operation. 2. NVMe Admin Commands that are prohibited via the out-of-band mechanism (refer to Figure 110) are not listed since they are always prohibited including during a sanitize operation. 		

7 PCIe Command Set (Optional)

The PCIe Command Set defines optional commands that a Management Controller may submit to access the memory, I/O, and configuration addresses spaces associated with a Controller in the NVM Subsystem. Only addresses mapped to the specified Controller may be accessed (e.g., these commands do not directly access memory on a host). The NMIMT field in the message header for PCIe Command Messages and Response Messages is set to 4h (PCIe Command). The PCIe Command Set is only applicable in the out-of-band mechanism and is prohibited in the in-band tunneling mechanism.

PCIe Commands over the out-of-band mechanism may interfere with host software. A Management Controller should coordinate with the host or issue only PCIe Commands that do not interfere with host software or in-band NVMe commands (e.g., PCIe Configuration Read). Coordination between a Management Controller and a host is outside the scope of this specification.

The Request Message format for PCIe Commands is shown in Figure 117 and described in Figure 118.

Figure 117: PCIe Command Request Format

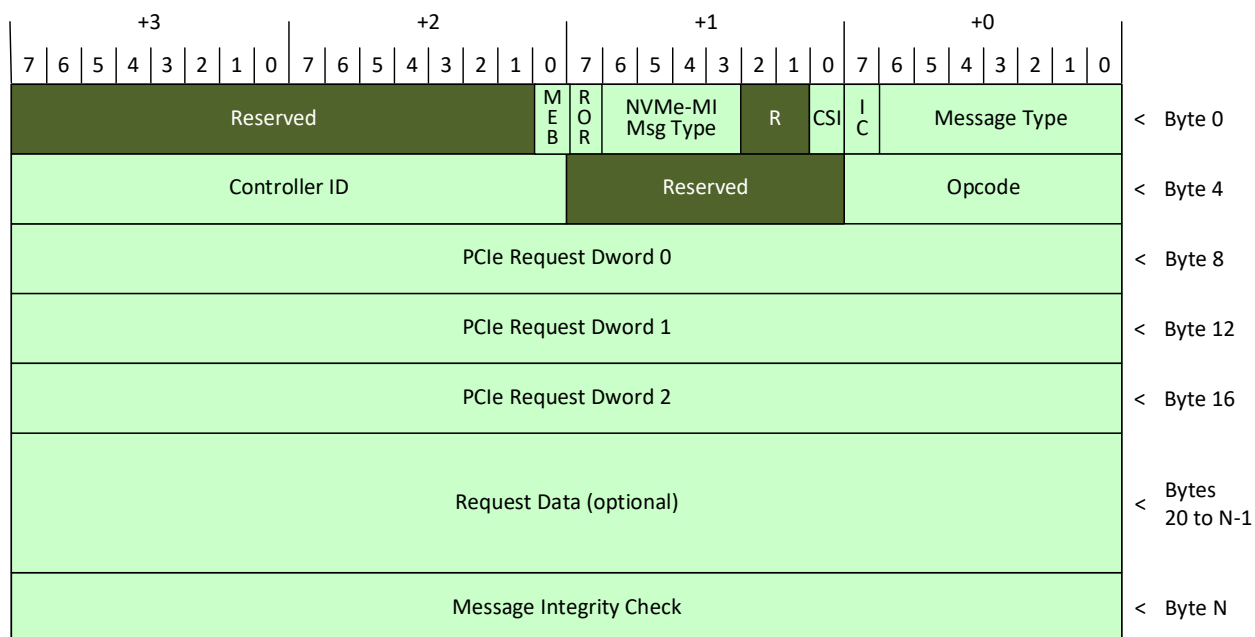


Figure 118: PCIe Command Request Description

Byte	Description
03:00	NVMe-MI Message Header (NMH): Refer to section 3.1.
04	Opcode (OPC): This field specifies the opcode of the command to be executed. Refer to Figure 119
05	Reserved
07:06	Controller ID (CTLID): This field specifies the Controller ID of the NVMe Controller that this command targets.
11:08	PCIe Request Dword 0 (NMD0): This field is command specific Dword 0.
15:12	PCIe Request Dword 1 (NMD1): This field is command specific Dword 1.
19:16	PCIe Request Dword 2 (NMD2): This field is command specific Dword 2.
N-1:20	Request Data (REQD): (Optional)
N+3:N	Message Integrity Check (MIC): Refer to section 3.1.

Figure 119 defines the PCIe Command opcodes. It also shows PCIe Commands that are mandatory, optional, and prohibited for an NVMe Storage Device and an NVMe Enclosure using the out-of-band mechanism. All PCIe Commands are prohibited using the in-band tunneling mechanism.

Figure 119: Opcodes for PCIe Commands using an Out-of-Band Mechanism

Opcode	NVMe Storage Device O/M/P ¹	NVMe Enclosure O/M/P ¹	Command
00h	O	O	PCIe Configuration Read
01h	O	O	PCIe Configuration Write
02h	O	O	PCIe Memory Read
03h	O	O	PCIe Memory Write
04h	O	O	PCIe I/O Read
05h	O	O	PCIe I/O Write
06h to FFh	-	-	Reserved

NOTES:
 1. O/M/P definition: O = Optional, M = Mandatory, P = Prohibited from being supported. An NVMe Enclosure that is also an NVMe Storage Device (i.e., implements namespaces) shall implement mandatory commands required by either an NVMe Storage Device or an NVMe Enclosure and may implement optional commands allowed by either an NVMe Storage Device or an NVMe Enclosure.

The Response Message for PCIe Command is shown in Figure 120 and described in Figure 121.

Figure 120: PCIe Command Response Format

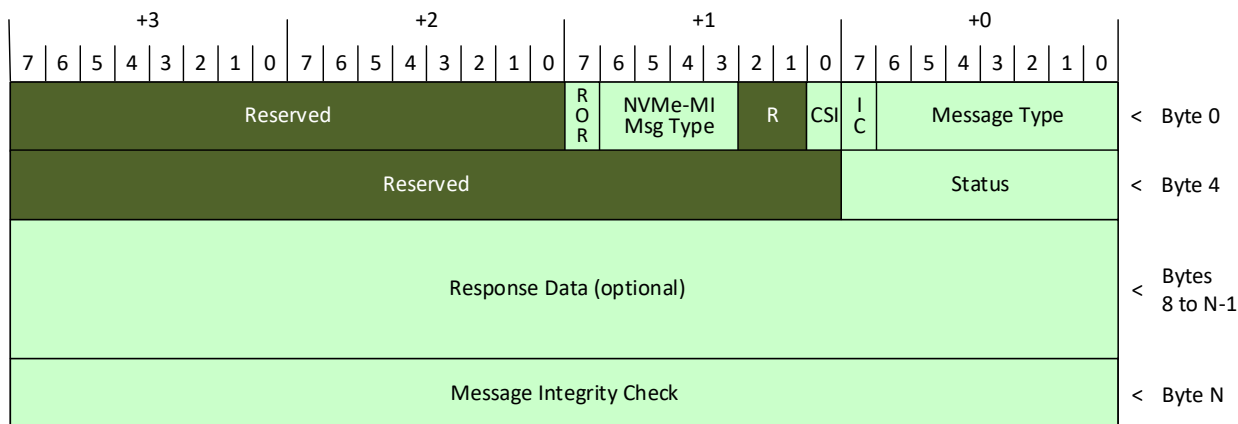


Figure 121: PCIe Command Response Description

Byte	Description
03:00	NVMe-MI Message Header (NMH): Refer to section 3.1.
04	Status (STATUS): This field indicates the status of the PCIe Command. Refer to section 4.1.2.
07:05	Reserved
N-1:08	Response Data (RESPD): (Optional)

Figure 121: PCIe Command Response Description

Byte	Description
N+3:N	Message Integrity Check (MIC): Refer to section 3.1.

PCIe Commands allow the Management Controller to access PCI Express configuration, I/O, and memory spaces of any Controller in the NVM Subsystem. Support for PCIe Commands is optional and indicated by the Optionally Supported Commands data structure. Refer to Figure 95.

An implementation may support a subset of the PCIe Commands. For supported commands, an implementation may block access to certain address space ranges (e.g., due to security concerns). A PCIe Command that attempts to access such a blocked address range is aborted with the Status field set to Access Denied.

It is recommended that PCIe Commands provide access to all non-blocked address spaces whenever MCTP access is supported. In some implementations, it may not be possible to access PCIe resources in certain states. A PCIe Command executed when a Controller is in one of these states may be aborted with the Status field set to PCIe Inaccessible. Refer to section 9.1.

A PCIe Command that is not well-formed results in an Error Response. A PCIe Command is well formed if it does not contain any of the following errors:

- Invalid Opcode (e.g., the Opcode is not listed in Figure 119);
- Invalid Parameter (e.g., the Controller ID field specifies a Controller ID not implemented in the NVM Subsystem);
- Invalid Command Size (e.g., the Request Message does not contain a complete command); or
- Invalid Command Input Data Size (e.g., the NVMe Request Data field is larger than the size expected by the command).

7.1 PCIe Configuration Read

The PCIe Configuration Read command allows the Management Controller to read the contents of the PCIe configuration address space associated with an NVMe Controller in the NVM Subsystem. The Controller ID field in the Command Message specifies the Controller ID that is being accessed.

The command uses PCIe Request Dwords 0 and 1. PCIe Request Dword 2 is not used and is reserved. The format of PCIe Request Dwords 0 and 1 are shown in Figure 122 and Figure 123 respectively.

Figure 122: PCIe Configuration Read – PCIe Request Dword 0

Bit	Description
31:16	Reserved
15:00	Length (LENGTH): This field specifies the number of bytes to be read.

Figure 123: PCIe Configuration Read – PCIe Request Dword 1

Bit	Description
31:12	Reserved
11:00	Offset (OFFSET): This field specifies the offset in bytes into the 4 KiB configuration space associated with the NVMe Controller at which the read begins.

When this command is completed successfully, PCI configuration space associated with the NVMe Controller specified by Controller ID is read and returned in the Response Data field. The Offset field specifies the starting read offset in PCIe configuration address space and the Length field specifies the number of bytes to be read. The Response Data field is always an integral number of dwords and is equal to the Length field rounded up to the next dword. If Length is not an integral number of dwords, then zero padding follows read data.

If the sum of the Offset and Length fields fall outside of PCI configuration space, then the Management Endpoint responds with an Invalid Parameter Error Response. The parameter with the error in this case is always the Offset field.

A Management Endpoint shall support the PCIe Configuration Read command if any of the other PCIe Command Set commands are supported. Access to the BAR offsets shall not return an Access Denied Response Message Status (i.e., the correct data shall be provided).

7.2 PCIe Configuration Write

The PCIe Configuration Write command allows the Management Controller to write the contents of the PCIe configuration address space associated with an NVMe Controller in the NVM Subsystem. The Controller ID field in the Command Message specifies the Controller ID that is being accessed.

The command uses PCIe Request Dwords 0 and 1. PCIe Request Dword 2 is not used and is reserved. The format of PCIe Request Dwords 0 and 1 are shown in Figure 124 and Figure 125 respectively.

Figure 124: PCIe Configuration Write – PCIe Request Dword 0

Bit	Description
31:16	Reserved
15:00	Length (LENGTH): This field specifies the number of bytes to be written.

Figure 125: PCIe Configuration Write – PCIe Request Dword 1

Bit	Description
31:12	Reserved
11:00	Offset (OFFSET): This field specifies the offset in bytes into the 4,096B configuration space associated with the NVMe Controller at which the write begins.

When this command is completed successfully, PCI configuration space associated with the NVMe Controller specified by Controller ID is written with the data contained in the Request Data field. The Offset field specifies the starting write offset in PCIe configuration address space and the Length field specifies the number of bytes to be written. The Request Data field is always an integral number of dwords and is equal to the Length field rounded up to the next dword. If Length is not an integral number of dwords, then unused padding bytes are discarded.

If the sum of the Offset and Length fields fall outside of PCI configuration space, then the Management Endpoint responds with an Invalid Parameter Error Response. The parameter with the error in this case is always the Offset field.

7.3 PCIe I/O Read

The PCIe I/O Read command allows the Management Controller to read the contents of PCIe I/O space associated with an NVMe Controller in the NVM Subsystem. The Controller ID field in the Command Message specifies the Controller ID that is being accessed.

The command uses PCIe Request Dwords 0 and 1. PCIe Request Dword 2 is not used and is reserved. The format of PCIe Request Dword 0 and 1 are shown in Figure 126 and Figure 127 respectively.

Figure 126: PCIe I/O Read – PCIe Request Dword 0

Bit	Description																
31:19	Reserved																
18:16	<p>Base Address Register (BAR): This field specifies the PCI Base Address Register (BAR) of the I/O space to be read. BARs are located beginning at 10h in PCI Configuration space and the value of this field specifies the starting offset of the associated BAR. For a 64-bit BAR, this field should correspond to the lower 32-bits of the BAR.</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>Value</th> <th>BAR Offset</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>10h</td> </tr> <tr> <td>1h</td> <td>14h</td> </tr> <tr> <td>2h</td> <td>18h</td> </tr> <tr> <td>3h</td> <td>1Ch</td> </tr> <tr> <td>4h</td> <td>20h</td> </tr> <tr> <td>5h</td> <td>24h</td> </tr> <tr> <td>6h to 7h</td> <td>Reserved</td> </tr> </tbody> </table>	Value	BAR Offset	0h	10h	1h	14h	2h	18h	3h	1Ch	4h	20h	5h	24h	6h to 7h	Reserved
Value	BAR Offset																
0h	10h																
1h	14h																
2h	18h																
3h	1Ch																
4h	20h																
5h	24h																
6h to 7h	Reserved																
15:00	Length (LENGTH): This field specifies the number of bytes to be read.																

Figure 127: PCIe I/O Read – PCIe Request Dword 1

Bit	Description
31:00	Offset (OFFSET): This field specifies the offset in bytes into the PCI BAR associated with the NVMe Controller at which the read begins.

When this command is completed successfully, PCI I/O space associated with the NVMe Controller specified by Controller ID is read and returned in the Response Data field. The Offset field specifies the starting read offset in PCIe I/O address space specified by the Base Address Register field. The Length field specifies the number of bytes to be read. The Response Data field is always an integral number of dwords and is equal to the Length field rounded up to the next dword. If Length is not an integral number of dwords, then zero padding follows read data.

If the Base Address Register field does not correspond to an I/O BAR implemented by the specified NVMe Controller, then the Management Endpoint responds with an Invalid Parameter Error Response.

If the sum of the Offset and Length fields fall outside the address range of the BAR specified by the Base Address Register field, then the Management Endpoint responds with an Invalid Parameter Error Response. The parameter with the error in this case is always the Offset field.

7.4 PCIe I/O Write

The PCIe I/O Write command allows the Management Controller to write the contents of PCIe I/O space associated with an NVMe Controller in the NVM Subsystem. The Controller ID field in the Command Message specifies the Controller ID that is being accessed.

The command uses PCIe Request Dwords 0 and 1. PCIe Request Dword 2 is not used and is reserved. The format of PCIe Request Dword 0 and 1 are shown in Figure 128 and Figure 129 respectively.

Figure 128: PCIe I/O Write – PCIe Request Dword 0

Bit	Description																
31:19	Reserved																
18:16	<p>Base Address Register (BAR): This field specifies the PCI Base Address Register (BAR) of the I/O space to be written. BARs are located beginning at 10h in PCI Configuration space and the value of this field specifies the starting offset of the associated BAR. For a 64-bit BAR, this field should correspond to the lower 32-bits of the BAR.</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>Value</th> <th>BAR Offset</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>10h</td> </tr> <tr> <td>1h</td> <td>14h</td> </tr> <tr> <td>2h</td> <td>18h</td> </tr> <tr> <td>3h</td> <td>1Ch</td> </tr> <tr> <td>4h</td> <td>20h</td> </tr> <tr> <td>5h</td> <td>24h</td> </tr> <tr> <td>6h to 7h</td> <td>Reserved</td> </tr> </tbody> </table>	Value	BAR Offset	0h	10h	1h	14h	2h	18h	3h	1Ch	4h	20h	5h	24h	6h to 7h	Reserved
Value	BAR Offset																
0h	10h																
1h	14h																
2h	18h																
3h	1Ch																
4h	20h																
5h	24h																
6h to 7h	Reserved																
15:00	Length (LENGTH): This field specifies the number of bytes to be written.																

Figure 129: PCIe I/O Write – PCIe Request Dword 1

Bit	Description
31:00	Offset (OFFSET): This field specifies the offset in bytes into the PCI BAR associated with the NVMe Controller at which the write begins.

When this command is completed successfully, PCI I/O space associated with the NVMe Controller specified by Controller ID is written with the data contained in the Request Data field. The Offset field specifies the starting write offset in PCIe I/O address space specified by the Base Address Register field. The Length field specifies the number of bytes to be written. The Request Data field is always an integral number of dwords and is equal to the Length field rounded up to the next dword. If Length is not an integral number of dwords, then unused padding bytes are discarded.

If the Base Address Register field does not correspond to an I/O BAR implemented by the specified NVMe Controller, then the Management Endpoint responds with an Invalid Parameter Error Response.

If the sum of the Offset and Length fields fall outside the address range of the BAR specified by the Base Address Register field, then the Management Endpoint responds with an Invalid Parameter Error Response. The parameter with the error in this case is always the Offset field.

7.5 PCIe Memory Read

The PCIe Memory Read command allows the Management Controller to read the contents of PCIe memory associated with an NVMe Controller in the NVM Subsystem. The Controller ID field in the Command Message specifies the Controller ID that is being accessed.

The command uses PCIe Request Dwords 0, 1, and 2. The format of PCIe Request Dword 0, 1, and 2 are shown in Figure 130, Figure 131, and Figure 132 respectively.

Figure 130: PCIe Memory Read – PCIe Request Dword 0

Bit	Description
31:19	Reserved

Figure 130: PCIe Memory Read – PCIe Request Dword 0

Bit	Description																
18:16	<p>Base Address Register (BAR): This field specifies the PCI Base Address Register (BAR) of the memory space to be read. BARs are located beginning at 10h in PCI Configuration space and the value of this field specifies the starting offset of the associated BAR. For a 64-bit BAR, this field should correspond to the lower 32-bits of the BAR.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>BAR Offset</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>10h</td> </tr> <tr> <td>1h</td> <td>14h</td> </tr> <tr> <td>2h</td> <td>18h</td> </tr> <tr> <td>3h</td> <td>1Ch</td> </tr> <tr> <td>4h</td> <td>20h</td> </tr> <tr> <td>5h</td> <td>24h</td> </tr> <tr> <td>6h to 7h</td> <td>Reserved</td> </tr> </tbody> </table>	Value	BAR Offset	0h	10h	1h	14h	2h	18h	3h	1Ch	4h	20h	5h	24h	6h to 7h	Reserved
Value	BAR Offset																
0h	10h																
1h	14h																
2h	18h																
3h	1Ch																
4h	20h																
5h	24h																
6h to 7h	Reserved																
15:00	Length (LENGTH): This field specifies the number of bytes to be read.																

Figure 131: PCIe Memory Read – PCIe Request Dword 1

Bit	Description
31:00	Offset (OFFSET): This field specifies the lower 32-bits (i.e., bit 0 to bit 31) of the offset in bytes into the PCI BAR associated with the NVMe Controller at which the read begins.

Figure 132: PCIe Memory Read – PCIe Request Dword 2

Bit	Description
31:00	Offset (OFFSET): This field specifies the upper 32-bits (i.e., bit 32 to bit 63) of the offset in bytes into the PCI BAR associated with the NVMe Controller at which the read begins.

When this command is completed successfully, PCI memory space associated with the NVMe Controller specified by Controller ID is read and returned in the Response Data field. The Offset field specifies the starting read offset in PCIe memory address space specified by the Base Address Register field. The Length field specifies the number of bytes to be read. The Response Data field is always an integral number of dwords and is equal to the Length field rounded up to the next dword. If Length is not an integral number of dwords, then zero padding follows read data.

If the Base Address Register field does not correspond to one implemented by the specified NVMe Controller, or the address range specified by the Base Address Range is not a memory region, then the Management Endpoint responds with an Invalid Parameter Error Response.

If the sum of the Offset and Length fields fall outside the address range specified by the Base Address Register field, then the Management Endpoint responds with an Invalid Parameter Error Response. The parameter with the error in this case is always the Offset field.

7.6 PCIe Memory Write

The PCIe Memory Write command allows the Management Controller to write the contents of PCIe memory associated with an NVMe Controller in the NVM Subsystem. The Controller ID field in the Command Message specifies the Controller ID that is being accessed.

The command uses PCIe Request Dwords 0, 1, and 2. The format of PCIe Request Dword 0, 1, and 2 are shown in Figure 133, Figure 134, and Figure 135 respectively.

Figure 133: PCIe Memory Write – PCIe Request Dword 0

Bit	Description																
31:19	Reserved																
18:16	<p>Base Address Register (BAR): This field specifies the PCI Base Address Register (BAR) of the memory space to be written. BARs are located beginning at 10h in PCI Configuration space and the value of this field specifies the starting offset of the associated BAR. For a 64-bit BAR, this field should correspond to the lower 32-bits of the BAR.</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>Value</th> <th>BAR Offset</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>10h</td> </tr> <tr> <td>1h</td> <td>14h</td> </tr> <tr> <td>2h</td> <td>18h</td> </tr> <tr> <td>3h</td> <td>1Ch</td> </tr> <tr> <td>4h</td> <td>20h</td> </tr> <tr> <td>5h</td> <td>24h</td> </tr> <tr> <td>6h to 7h</td> <td>Reserved</td> </tr> </tbody> </table>	Value	BAR Offset	0h	10h	1h	14h	2h	18h	3h	1Ch	4h	20h	5h	24h	6h to 7h	Reserved
Value	BAR Offset																
0h	10h																
1h	14h																
2h	18h																
3h	1Ch																
4h	20h																
5h	24h																
6h to 7h	Reserved																
15:00	Length (LENGTH): This field specifies the number of bytes to be written.																

Figure 134: PCIe Memory Write – PCIe Request Dword 1

Bit	Description
31:00	Offset (OFFSET): This field specifies the lower 32-bits (i.e., bit 0 to bit 31) of the offset in bytes into the PCI BAR associated with the NVMe Controller at which the write begins.

Figure 135: PCIe Memory Write – PCIe Request Dword 2

Bit	Description
31:00	Offset (OFFSET): This field specifies the upper 32-bits (i.e., bit 32 to bit 63) of the offset in bytes into the PCI BAR associated with the NVMe Controller at which the write begins.

When this command is completed successfully, PCI memory space associated with the NVMe Controller specified by Controller ID is written with the data contained in the Request Data field. The Offset field specifies the starting write offset in PCIe memory address space specified by the Base Address Register field. The Length field specifies the number of bytes to be written. The Request Data field is always an integral number of dwords and is equal to the Length field rounded up to the next dword. If Length is not an integral number of dwords, then unused padding bytes are discarded.

If the Base Address Register field does not correspond to one implemented by the specified NVMe Controller, or the address range specified by the Base Address Range is not a memory region, then the Management Endpoint responds with an Invalid Parameter Error Response.

If the sum of the Offset and Length fields fall outside the address range of the BAR specified by the Base Address Register field, then the Management Endpoint responds with an Invalid Parameter Error Response. The parameter with the error in this case is always the Offset field.

8 NVM Express Management Enhancements

This section describes NVMe Management Interface enhancements to the NVM Express specification.

8.1 Identify Controller

The NVMe Identify Controller data structure contains information about an NVMe Controller. Byte 240 to byte 255 have been allocated by the NVM Express specification for this specification and are defined below.

Figure 136: NVMe Management Interface Identify Controller

Bytes	O/M ¹	Description								
252:240		Reserved								
253	M	<p>NVM Subsystem Report (NVMSR): This field reports information associated with the NVM Subsystem. At least one bit in this field shall be set to '1'.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:2</td> <td>Reserved</td> </tr> <tr> <td>1</td> <td>NVMe Enclosure (NVMEE): If set to '1', then the NVM Subsystem is part of an NVMe Enclosure. If cleared to '0', then the NVM Subsystem is not part of an NVMe Enclosure.</td> </tr> <tr> <td>0</td> <td>NVMe Storage Device (NVMESD): If set to '1', then the NVM Subsystem is part of an NVMe Storage Device. If cleared to '0', then the NVM Subsystem is not part of an NVMe Storage Device.</td> </tr> </tbody> </table>	Bits	Description	7:2	Reserved	1	NVMe Enclosure (NVMEE): If set to '1', then the NVM Subsystem is part of an NVMe Enclosure. If cleared to '0', then the NVM Subsystem is not part of an NVMe Enclosure.	0	NVMe Storage Device (NVMESD): If set to '1', then the NVM Subsystem is part of an NVMe Storage Device. If cleared to '0', then the NVM Subsystem is not part of an NVMe Storage Device.
Bits	Description									
7:2	Reserved									
1	NVMe Enclosure (NVMEE): If set to '1', then the NVM Subsystem is part of an NVMe Enclosure. If cleared to '0', then the NVM Subsystem is not part of an NVMe Enclosure.									
0	NVMe Storage Device (NVMESD): If set to '1', then the NVM Subsystem is part of an NVMe Storage Device. If cleared to '0', then the NVM Subsystem is not part of an NVMe Storage Device.									
254	M	<p>VPD Write Cycle Information (VWCI): This field indicates information about remaining number of times that VPD contents are able to be updated using the VPD Write command.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7</td> <td>VPD Write Cycle Remaining Valid (VWCRV): If this bit is set to '1', then the VPD Write Cycle Remaining field is valid. If this bit is cleared to '0', then the VPD Write Cycles Remaining field is invalid and cleared to 0h.</td> </tr> <tr> <td>6:0</td> <td>VPD Write Cycles Remaining (VWCR): If the VPD Write Cycle Remaining Valid bit is set to '1', then this field contains a value indicating the remaining number of times that VPD contents are able to be updated using the VPD Write command. If this field is set to 7Fh, then the remaining number of times that VPD contents are able to be updated using the VPD Write command is greater than or equal to 7Fh. If the VPD Write Cycle Remaining Valid bit is cleared to '0', then this field is not valid and shall be cleared to a value of 0h.</td> </tr> </tbody> </table>	Bits	Description	7	VPD Write Cycle Remaining Valid (VWCRV): If this bit is set to '1', then the VPD Write Cycle Remaining field is valid. If this bit is cleared to '0', then the VPD Write Cycles Remaining field is invalid and cleared to 0h.	6:0	VPD Write Cycles Remaining (VWCR): If the VPD Write Cycle Remaining Valid bit is set to '1', then this field contains a value indicating the remaining number of times that VPD contents are able to be updated using the VPD Write command. If this field is set to 7Fh, then the remaining number of times that VPD contents are able to be updated using the VPD Write command is greater than or equal to 7Fh. If the VPD Write Cycle Remaining Valid bit is cleared to '0', then this field is not valid and shall be cleared to a value of 0h.		
Bits	Description									
7	VPD Write Cycle Remaining Valid (VWCRV): If this bit is set to '1', then the VPD Write Cycle Remaining field is valid. If this bit is cleared to '0', then the VPD Write Cycles Remaining field is invalid and cleared to 0h.									
6:0	VPD Write Cycles Remaining (VWCR): If the VPD Write Cycle Remaining Valid bit is set to '1', then this field contains a value indicating the remaining number of times that VPD contents are able to be updated using the VPD Write command. If this field is set to 7Fh, then the remaining number of times that VPD contents are able to be updated using the VPD Write command is greater than or equal to 7Fh. If the VPD Write Cycle Remaining Valid bit is cleared to '0', then this field is not valid and shall be cleared to a value of 0h.									
255	M	<p>Management Endpoint Capabilities (MEC): This field indicates the capabilities of the Management Endpoint in the Controller.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:2</td> <td>Reserved</td> </tr> <tr> <td>1</td> <td>PCIe Port Management Endpoint (PCIEME): If set to '1', then the NVM Subsystem contains a Management Endpoint on a PCIe port.</td> </tr> <tr> <td>0</td> <td>SMBus/I2C Port Management Endpoint (SMBUSME): If set to '1', then the NVM Subsystem contains a Management Endpoint on an SMBus/I2C port.</td> </tr> </tbody> </table>	Bits	Description	7:2	Reserved	1	PCIe Port Management Endpoint (PCIEME): If set to '1', then the NVM Subsystem contains a Management Endpoint on a PCIe port.	0	SMBus/I2C Port Management Endpoint (SMBUSME): If set to '1', then the NVM Subsystem contains a Management Endpoint on an SMBus/I2C port.
Bits	Description									
7:2	Reserved									
1	PCIe Port Management Endpoint (PCIEME): If set to '1', then the NVM Subsystem contains a Management Endpoint on a PCIe port.									
0	SMBus/I2C Port Management Endpoint (SMBUSME): If set to '1', then the NVM Subsystem contains a Management Endpoint on an SMBus/I2C port.									

NOTES:
1. O/M definition: O = Optional, M = Mandatory.

8.2 Management Interface Specific Features

The NVMe Get Features and Set Features Admin commands are used to retrieve and modify Feature values. Feature Identifiers 78h to 7Fh have been allocated by the NVMe Express specification for this specification and are defined below.

Figure 137: NVMe Management Interface Feature Identifiers

Feature Identifier	NVMe Storage Device O/M ¹	NVMe Enclosure O/M ¹	Persistent Across Power States and Reset ²	Uses Memory Buffer for Attributes	Description
78h to 7Dh	-	-	-	-	Reserved
7Eh	M	M	No	Yes	Controller Metadata
7Fh	M	O	No	Yes	Namespace Metadata

NOTES:

- O/M definition: O = Optional, M = Mandatory. Mandatory features shall be supported if the NVMe Subsystem implements a Management Endpoint. These features are not mandatory if the NVMe Subsystem does not implement a Management Endpoint.
- This column is only valid if bit 4 in the Optional NVMe Command Support field of the Identify Controller Data Structure is cleared to '0'. Refer to the NVMe Express specification.

8.2.1 Controller Metadata

This feature is used to store metadata about the host platform in an NVMe Subsystem for later retrieval. The values stored in the Controller Metadata Feature do not modify Controller behavior.

The Controller Metadata feature uses NVMe Set Feature Command Dword 11 as shown in Figure 138.

Figure 138: Host Metadata – Command Dword 11

Bit	Description								
31:15	Reserved								
14:13	<p>Element Action (EA): This field specifies the action to perform on the Metadata Element Descriptor data structure. This field shall be cleared to 0h for a Get Features.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Add/Update Entry</td> </tr> <tr> <td>01b</td> <td>Delete Entry</td> </tr> <tr> <td>10b to 11b</td> <td>Reserved</td> </tr> </tbody> </table> <p>If the Element Action field is cleared to 00b (Add/Update Entry) and a Metadata Element Descriptor with the specified Element Type (refer to Figure 140) does not exist in Controller Metadata, then the Controller creates a new descriptor with the value in the Controller Metadata structure. This operation is performed in an atomic manner.</p> <p>If the Element Action field is cleared to 00bh (Add/Update Entry) and a Metadata Element Descriptor with the specified Element Type exists in the Controller Metadata, then the Controller updates the descriptor with the value in the Controller Metadata structure. This operation is performed in an atomic manner.</p> <p>If the Element Action field is set to 01b (Delete Entry) and a Metadata Element Descriptor with the specified Element Type does not exist in the Controller Metadata, then no operation is performed, and the command completes successfully.</p> <p>If the Element Action field is set to 01b (Delete Entry) and a Metadata Element Descriptor with the specified Element Type exists in the Controller Metadata, then the Controller deletes the specified Metadata Element Descriptor. This operation is performed in an atomic manner.</p>	Value	Definition	00b	Add/Update Entry	01b	Delete Entry	10b to 11b	Reserved
Value	Definition								
00b	Add/Update Entry								
01b	Delete Entry								
10b to 11b	Reserved								
12:00	Reserved								

New Metadata Element Descriptors may be added, updated, or deleted based on the action specified in the Element Action field.

If a Set Features command is submitted for this Feature, a Host Metadata data structure, defined in Figure 139, is transferred in the data buffer for the command. The Host Metadata data structure is 4 KiB in size and contains zero or more Metadata Element Descriptors. If host software attempts to add or update a Metadata Element that causes the stored Host Metadata data structure to grow larger than 4 KiB, the Controller shall abort the command with an Invalid Parameter Error Response. The Host Metadata Data Structure for this feature is independent of the Host Metadata data structure for the Namespace Metadata feature described in section 8.2.2.

Figure 139: Host Metadata Data Structure

Byte	Description
00	Number of Metadata Element Descriptors: This field contains the number of Metadata Element Descriptors in the data structure.
01	Reserved
x:02	Metadata Element Descriptor 0: This field contains the first Metadata Element descriptor.
y:x+1	Metadata Element Descriptor 1: This field contains the second Metadata Element descriptor or 0h if there is only 1 entry.
...	...
4095:z	Metadata Element Descriptor N: This field contains the (N+1)th Metadata Element descriptor or 0h if there are fewer than N+1 entries.

A Host Metadata data structure may contain at most one Metadata Element Descriptor of each element type. Each Metadata Element Descriptor contains the data structure shown in Figure 140.

Figure 140: Metadata Element Descriptor

Bit	Description								
31 + (Element Length*8) :32	Element Value (EVAL): This field specifies the value for the element.								
31:16	Element Length (ELEN): This field specifies the length of the Element Value field in bytes. This field shall be 0h when deleting an entry (EA = 01b). This field should be non-zero when adding/updating and entry (EA = 00b).								
15:12	Reserved								
11:08	Element Revision (ER): This field specifies the revision of this element value. Unless specified otherwise elsewhere in this specification, all Metadata Element Descriptors compliant with this version of the NVMe-MI Specification shall clear this field to a value of 0h.								
07:06	Reserved								
05:00	Element Type (ET): This field specifies the type of metadata stored in the descriptor.								
	<table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>Reserved</td> </tr> <tr> <td>01h to 017h</td> <td>Element types defined by this specification. Controller Metadata Element types are defined in Figure 141. Namespace Metadata Element types are defined in Figure 142.</td> </tr> <tr> <td>18h to 1Fh</td> <td>Vendor Specific</td> </tr> </tbody> </table>	Value	Definition	00h	Reserved	01h to 017h	Element types defined by this specification. Controller Metadata Element types are defined in Figure 141. Namespace Metadata Element types are defined in Figure 142.	18h to 1Fh	Vendor Specific
	Value	Definition							
	00h	Reserved							
01h to 017h	Element types defined by this specification. Controller Metadata Element types are defined in Figure 141. Namespace Metadata Element types are defined in Figure 142.								
18h to 1Fh	Vendor Specific								

If a Get Features command is issued for this Feature, all Controller Metadata associated with the specified Controller is added to a Host Metadata Data Structure specified in Figure 139 and returned in the data

buffer for that command. The data buffer size is equal to the size of the Host Metadata Data Structure and is 4,096 bytes in size.

Figure 141: Controller Metadata Element Types

Value	Definition
00h	Reserved
01h	Operating System Controller Name: The name of the Controller in the operating system as a UTF-8 string.
02h	Operating System Driver Name: The name of the driver in the operating system as a UTF-8 string.
03h	Operating System Driver Version: The version of the driver in the operating system as a UTF-8 string.
04h	Pre-boot Controller Name: The name of the Controller in the pre-boot environment as a UTF-8 string.
05h	Pre-boot Driver Name: The name of the driver in the pre-boot environment as a UTF-8 string.
06h	Pre-boot Driver Version: The version of the driver in the pre-boot environment as a UTF-8 string.
07h to 17h	Reserved
18h to 1Fh	Vendor Specific

Controller Metadata is reset on a Controller Level Reset (i.e., the number of stored Metadata Element Descriptors is zero). Executing a Get Features command while the Controller is disabled returns zero Metadata Element Descriptors.

8.2.2 Namespace Metadata

This feature is used to store metadata about a namespace associated with a Controller in the NVM Subsystem for later retrieval. The values stored in the Namespace Metadata Feature do not modify Controller behavior on the namespace. This feature is namespace specific.

The Namespace Metadata feature uses Command Dword 11 as shown in Figure 138.

New Metadata Element Descriptors may be added, updated, or deleted based on the action specified in the Element Action field.

If a Set Features command is submitted for this Feature, a Host Metadata data structure, defined in Figure 139, is transferred in the data buffer for the command. The Host Metadata data structure is 4 KiB in size and contains zero or more Metadata Element Descriptors. If host software attempts to add or update a Metadata Element Descriptor that causes the stored Host Metadata data structure to grow larger than 4 KiB, the Controller shall abort the command with an Invalid Parameter Error Response. The Host Metadata structure for this feature is independent of the Host Metadata data structure for the Controller Metadata feature described in section 8.2.1.

A Host Metadata data structure may contain up to one Metadata Element Descriptor of each element type. Each Metadata Element Descriptor contains the data structure shown in Figure 140.

If a Get Features command is issued for this Feature, all Namespace Metadata associated with the specified Controller is added to a Host Metadata Data Structure specified in Figure 139 and returned in the data buffer for that command. The data buffer size is equal to the size of the Host Metadata Data Structure and is 4 KiB in size.

Namespace Metadata is reset on a Controller Level Reset (i.e., the number of stored Metadata Element Descriptors is zero). Executing a Get Features command while the Controller is disabled returns zero Metadata Element Descriptors.

Figure 142: Namespace Metadata Element Types

Value	Definition
00h	Reserved
01h	Operating System Namespace Name: The name of the namespace in the operating system as a UTF-8 string.
02h	Pre-boot Namespace Name: The name of the namespace in the pre-boot environment as a UTF-8 string.
03h to 17h	Reserved
18h to 1Fh	Vendor Specific

9 Management Architecture

9.1 Out-of-Band Operational Times

In the out-of-band mechanism, the ability of a Management Endpoint to receive and process Request Messages outlined in this specification is dependent on the state of the Management Endpoint. This section enumerates Management Endpoint operational times and the operations supported in each of these operational times.

The NVM Subsystem power state is defined by the state of main power and auxiliary power. Main power consists of one or more voltage rails as defined by form factor. When main power consists of multiple voltage rails, main power is considered “on” when power is good on all main voltage rails. Auxiliary power is optionally supported by a form factor and enables SMBus/I2C communications in the absence of main power. Only the Powered On and Powered Off states are applicable in form factors and platforms that do not support auxiliary power. Figure 143 defines the power states of a Management Endpoint. Note that auxiliary power is described from the perspective of the NVM Subsystem and could be provided by any appropriate power rail in a host platform.

The operations supported in each NVM subsystem power state are summarized in Figure 143. VPD SMBus/I2C access consists of processing read operations to the FRU Information Device. SMBus/I2C MCTP access consists of processing and responding to MCTP messages on the NVM Subsystem SMBus/I2C port. PCIe MCTP access consists of processing and responding to MCTP messages issued on any NVM Subsystem PCIe port. The behavior of an operation that is “Not Supported” in Figure 143 is undefined.

Figure 143: Operations Supported During NVM Subsystem Power States

Operation	Powered Off -All Power Rails Off	Powered On -All Power Rails On	Auxiliary Power Only ² -Main Power Off -Auxiliary Power On	Main Power Only ² -Main Power On -Auxiliary Power Off
SMBus/I2C VPD and SMBus/I2C Mux Access	Not Supported	Supported	Supported	Implementation Specific
SMBus/I2C MCTP Access	Not Supported	Supported	Optional ¹	Implementation Specific
PCIe MCTP Access	Not Supported	Supported	Not Supported	Supported
NOTES: 1. An implementation that supports SMBus/I2C MCTP Access during Auxiliary Power may support a subset of commands during this power state. The commands that are supported are implementation specific. 2. Auxiliary Power Only and Main Power Only columns are not applicable to form factors that do not define Auxiliary power.				

When an NVM Subsystem transitions from a power state in which accesses are not supported to one where accesses are supported, accesses shall be processed 1 s after entering the power state in which accesses are supported. For example, an SMBus/I2C MCTP access issued 1 s after transitioning from a “Powered

Off” to a “Main Power” state is guaranteed to be processed. The behavior of accesses prior to this 1 s time interval is undefined. For example, the behavior of an SMBus/I2C MCTP access issued 50 ms after transitioning from a “Powered Off” to a “Main Power” state is undefined.

When transitioning between power states in which accesses are supported in both states (i.e., the state before and after the transition), there is no interruption in access processing (i.e., accesses are processed prior to the state transition, during the state transition, and immediately after entering the new power state).

Request Messages are processed whenever MCTP access is supported on an interface (i.e., SMBus/I2C or PCIe). Although not recommended, an implementation may not support PCIe and SMBus/I2C MCTP accesses during a PCI Express conventional reset on any PCI Express port in the NVM Subsystem. Although not recommended, an implementation may choose not to support processing of PCIe Commands that target a Controller in the NVM Subsystem that is in any of the following states:¹

- Controller Level Reset;
- SR-IOV virtual function is not enabled;
- During any type of PCI Express Conventional Reset;
- During a PCI Express Function Level Reset (FLR);
- When the PCI Express Function is in a non-D0 power D-state; or
- When the PCI Express link is down (i.e., not in the DL_Active state).

If a PCIe Command is received that targets a Controller in one of these states and the implementation does not support processing of PCIe Commands in that state, then the PCIe command is completed with status PCIe Inaccessible. Processing of supported PCIe Commands is required in all other Controller states.

If a PCIe Command is received that targets a Controller whose corresponding PCIe link is in a low power state (i.e., PCIe ASPM), then processing of the command may cause the link to temporarily exit the low power state.

9.2 Vital Product Data

The Vital Product Data (VPD) is information describing an NVMe Storage Device. Each NVMe Storage Device FRU shall have a FRU Information Device with a size of 256 bytes to hold the VPD as defined in the IPMI Platform Management FRU Information Storage Definition. The VPD for NVMe Storage Device FRUs shall contain the required elements defined in Figure 144. The VPD and FRU Information Device are optional for a) NVMe Storage Devices that are not FRUs (e.g., NVMe Storage Devices with a Form Factor type of Integrated per Figure 156), and b) NVMe Enclosures. The VPD contents for these optional use cases is outside the scope of this specification.

The VPD shall be accessible using the VPD Read command on all Management Endpoints on the NVMe Storage Device FRU. The entire contents of the VPD may be updated using the VPD Write command.

If the NVM Subsystem has an SMBus/I2C interface, the VPD shall be accessible at the SMBus/I2C address of the FRU Information Device using the access mechanism over I2C as defined in the IPMI Platform Management FRU Information Storage Definition. Updating the VPD by writing to the FRU Information Device directly on SMBus/I2C shall not be supported if the VPD Write command is supported.

¹ A Management Controller shall only send these commands using SMBus/I2C or another PCIe port since the link associated with the PCIe port and controller is down in these states.

Figure 144: VPD Elements

Byte	Name
7:0	Common Header
Vendor Specific	Product Info Area (Optional)
Vendor Specific	MultiRecord Info Area
Vendor Specific	Internal Use Area (Optional)
Vendor Specific	Chassis Info Area (Optional)
Vendor Specific	Board Info Area (Optional)

9.2.1 Common Header

The fields that make up the VPD Common Header are shown in Figure 145.

Figure 145: Common Header

Byte	Factory Default	Description
0	01h	IPMI Format Version Number (IPMIVER): This field indicates the IPMI Format Version.
1	Impl Spec	Internal Use Area Starting Offset (IUAOFF): This field indicates the starting offset in multiples of 8 bytes for the Internal Use Area. A value of 0h may be used to indicate the Internal Use Area is not present.
2	Impl Spec	Chassis Info Area Starting Offset (CIAOFF): This field indicates the starting offset in multiples of 8 bytes for the Chassis Info Area. A value of 0h may be used to indicate the Chassis Info Area is not present.
3	Impl Spec	Board Info Area Starting Offset (BIAOFF): This field indicates the starting offset in multiples of 8 bytes for the Board Info Area. A value of 0h may be used to indicate the Board Info Area is not present.
4	Impl Spec	Product Info Area Starting Offset (PIAOFF): This field indicates the starting offset in multiples of 8 bytes for the Product Info Area.
5	Impl Spec	MultiRecord Info Area Starting Offset (MRIOFF): This field indicates the starting offset in multiples of 8 bytes for the MultiRecord Info Area.
6	00h	Reserved
7	Impl Spec	Common Header Checksum (CHCHK): Checksum computed over byte 0 to byte 6. The checksum is computed by adding the 8-bit value of the bytes modulo 256 and then taking the 2's complement of this sum. When the checksum and the sum of the bytes module 256 are added, the result should be 0h.

9.2.2 Product Info Area (offset 8 bytes)

The optional Product Info Area shall have the same format and conventions as the Product Info Area Format as defined by the IPMI Platform Management FRU Information Storage Definition. Therefore, all fields within the Product Info Area shall not follow the conventions defined in section 1.8. The Product Info Area factory default values shall be set to the values defined in Figure 147. The Type/Length bytes use the format shown in Figure 146.

Figure 146: Type/Length Byte Format

Bits	Field Name	Description
7:6	Type Code	Specifies field encoding 11b – Always corresponds to ASCII in this specification
5:0	Number of Data Bytes	Specifies field length 0h indicates that the field is empty

Figure 147: Product Info Area Factory Default Values

Factory Default	Description
01h	IPMI Format Version Number (IPMIVER): This field indicates the IPMI Format Version.
Impl Spec	Product Info Area Length (PALEN): This field indicates the length of the Product Info Area in multiples of 8 bytes.
19h	Language Code (LCODE): This field indicates the language used. A value of 19h is used to indicate English.
Impl Spec	Manufacturer Name Type/Length (MNTL): This field indicates the type and length of the Manufacturer Name field. The maximum length is 8.
Impl Spec	Manufacturer Name (MNAME): This field indicates the Manufacturer name in 8-bit ASCII. Unused bytes should be NULL characters. The Manufacturer name in this field should correspond to that in the PCI Subsystem Vendor ID (SSVID) and IEEE OUI Identifier fields in the Identify Controller Data Structure.
Impl Spec	Product Name Type/Length (PNTL): This field indicates the type and length of the Product Name field. The maximum length is 24.
Impl Spec	Product Name (PNAME): This field indicates the Product name in 8-bit ASCII. Unused bytes should be NULL characters.
Impl Spec	Product Part/Model Number Type/Length (PPMNNTL): This field indicates the type and length of the Product Part/Model Number field. The maximum length is 40.
Impl Spec	Product Part/Model Number (PPMN): This field indicates the Product Part/Model Number in 8-bit ASCII. Unused bytes should be NULL characters. This field should contain the same value as the Model Number (NM) field in the NVMe Identify Controller Data Structure.
Impl Spec	Product Version Type/Length (PRTL): This field indicates the type and length of the Product Part/Model Number field. The maximum length is 2.
Impl Spec	Product Version (PVER): This field indicates the Product Version in 8-bit ASCII. Unused bytes should be NULL characters.
Impl Spec	Product Serial Number Type/Length (PSNTL): This field indicates the type and length of the Product Serial Number field. The maximum length is 20.
Impl Spec	Product Serial Number (PSN): This field indicates the Product Serial Number in 8-bit ASCII. Unused bytes should be NULL characters. This field should contain the same value as the Serial Number (SN) field in the NVMe Identify Controller Data Structure.
Impl Spec	Asset Tag Type/Length (ATTL): This field indicates the type and length of the Asset Tag field. A value of 0h may be used to indicate an Asset Tag is not present.
Impl Spec	Asset Tag (AT): This field indicates the asset tag.
Impl Spec	FRU File ID Type/Length (ATFL): This field indicates the type and length of the FRU File ID field. A value of 0h may be used to indicate a FRU File ID is not present.
Impl Spec	FRU File ID (FFI): This field provides manufacturing aid for verifying the file that was used during manufacture or field update to load the FRU information.
Impl Spec	Custom Product Info Area (CPIA): This optional field allows for the addition of custom Product Info Area fields that shall be preceded with a Type/Length field.
C1h	End of Record (EOR): A value of C1h in this field indicates the end of record.
0h	Zero or more bytes of value 0h that are used to pad the size of the Product Info Area to a multiple of 8 bytes.
Impl Spec	Product Info Area (PICCHK): Checksum computed over all bytes in the Product Info Area excluding this field. The checksum is computed by adding the 8-bit value of the bytes modulo 256 and then taking the 2's complement of this sum. When the checksum and the sum of the bytes module 256 are added, the result should be 0h.

9.2.3 NVMe MultiRecord Area

This MultiRecord is used to describe the form factor, power requirements, and capacity of NVMe Storage Devices with a single NVM Subsystem. Starting with version 1.1 of this specification, this MultiRecord has been superseded by the Topology MultiRecord (refer to section 9.2.5). For backwards compatibility with Management Controllers designed to versions of this specification prior to 1.1, the NVMe MultiRecord and the NVMe PCIe Port MultiRecord (refer to section 9.2.4) should both be included in the VPD in addition to the Topology MultiRecord unless the NVMe Storage Device FRU has Expansion Connectors, has more than one NVM Subsystem, or if including both this MultiRecord and the NVMe PCIe Port MultiRecord would extend the size of the VPD beyond the 256-byte limit. If both the NVMe MultiRecord and NVMe PCIe Port MultiRecord cannot fit within the size limit of the VPD then neither MultiRecord should be included.

Figure 148: NVMe MultiRecord Area

Byte Offset	Factory Default	Description						
00	0Bh	NVMe Record Type ID						
01	2h or 82h	Record Format:						
		<table border="1"> <thead> <tr> <th>Bit</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>7</td> <td>Set to '1' if last record in list.</td> </tr> <tr> <td>6:0</td> <td>Record format version = 2.</td> </tr> </tbody> </table>	Bit	Definition	7	Set to '1' if last record in list.	6:0	Record format version = 2.
		Bit	Definition					
7	Set to '1' if last record in list.							
6:0	Record format version = 2.							
02	3Bh	Record Length (RLEN): This field indicates the length of the MultiRecord Area in bytes without including the first 5 bytes that are common to all MultiRecords.						
03	Impl Spec	Record Checksum: This field is used to give the record data a zero checksum (i.e., the modulo 256 sum of the record data bytes from byte offset 05 to the end of this record plus this checksum byte equals zero).						
04	Impl Spec	Header Checksum: This field is used to give the record header a zero checksum (i.e., the modulo 256 sum of the first byte of the header through this checksum byte equals zero).						
05	0h	NVMe MultiRecord Area Version Number: This field indicates the version number of this NVMe MultiRecord. This field shall be cleared to 0h in this version of the specification.						
06	Impl Spec	Form Factor (FF): This field indicates the form factor of the Management Endpoint. Refer to the values in Figure 156.						
12:07	0h	Reserved						
13	Impl Spec 1	Initial 1.8 V Power Supply Requirements: This field specifies the initial 1.8 V power supply requirements in Watts prior to receiving a Set Slot Power message.						
14	Impl Spec 1	Maximum 1.8 V Power Supply Requirements: This field specifies the maximum 1.8 V power supply requirements in Watts.						
15	Impl Spec 1	Initial 3.3 V Power Supply Requirements: This field specifies the initial 3.3 V power supply requirements in Watts prior to receiving a Set Slot Power message.						
16	Impl Spec 1	Maximum 3.3 V Power Supply Requirements: This field specifies the maximum 3.3 V power supply requirements in Watts.						
17	0h	Reserved						
18	Impl Spec 1	Maximum 3.3 V aux Power Supply Requirements: This field specifies the maximum 3.3 V power supply requirements in 10 mW units.						
19	Impl Spec 1	Initial 5 V Power Supply Requirements: This field specifies the initial 5 V power supply requirements in Watts prior to receiving a Set Slot Power message.						
20	Impl Spec 1	Maximum 5 V Power Supply Requirements: This field specifies the maximum 5 V power supply requirements in Watts.						
21	Impl Spec 1	Initial 12 V Power Supply Requirements: This field specifies the initial 12 V power supply requirements in Watts prior to receiving a Set Slot Power message.						
22	Impl Spec 1	Maximum 12 V Power Supply Requirements: This field specifies the maximum 12 V power supply requirements in Watts.						

Figure 148: NVMe MultiRecord Area

Byte Offset	Factory Default	Description
23	Impl Spec	Maximum Thermal Load: This field specifies the maximum thermal load from the NVM Subsystem in Watts.
36:24	Impl Spec	Total NVM Capacity: This field indicates the total NVM capacity of the NVM Subsystem in bytes. If the NVM Subsystem supports Namespace Management, then this field should correspond to the value reported in the TNVMCAP field in the NVMe Identify Controller Data Structure. A value of 0h may be used to indicate this feature is not supported.
63:37	0h	Reserved
NOTES:		
1. Power supply requirements shall be set to the smallest integer value which fully supplies the necessary power to the NVMe Storage Device. A value of 0h indicates that the power supply voltage is not used.		

9.2.4 NVMe PCIe Port MultiRecord Area

This MultiRecord is used to describe the PCIe connectivity for NVMe Storage Devices with a single NVM Subsystem. Starting with version 1.1 of this specification, this MultiRecord has been superseded by the Topology MultiRecord (refer to section 9.2.5). For backwards compatibility with Management Controllers designed to versions of this specification prior to 1.1, the NVMe PCIe Port MultiRecord and the NVMe MultiRecord (refer to section 9.2.3) should both be included in the VPD in addition to the Topology MultiRecord unless the NVMe Storage Device FRU has Expansion Connectors, has more than one NVM Subsystem, or if including both this MultiRecord and the NVMe MultiRecord would extend the size of the VPD beyond the 256-byte limit. If both the NVMe MultiRecord and NVMe PCIe Port MultiRecord cannot fit within the size limit of the VPD then neither MultiRecord should be included.

Figure 149: NVMe PCIe Port MultiRecord Area

Byte Offset	Factory Default	Description
00	0Ch	NVMe PCIe Port Record Type ID
01	2h or 82h	Record Format:
		Bit Definition
		7 Set to '1' if last record in list.
6:0 Record format version = 2.		
02	0Bh	Record Length (RLEN): This field indicates the length of the MultiRecord Area in bytes without including the first 5 bytes that are common to all MultiRecords.
03	Impl Spec	Record Checksum: This field is used to give the record data a zero checksum (i.e., the modulo 256 sum of the record data bytes from byte offset 05 to the end of this record plus this checksum byte equals zero).
04	Impl Spec	Header Checksum: This field is used to give the record header a zero checksum (i.e., the modulo 256 sum of the first byte of the header through this checksum byte equals zero).
05	1h	NVMe PCIe Port MultiRecord Area Version Number: This field indicates the version number of this NVMe PCIe Port MultiRecord. This field shall be set to 1h in this version of the specification.
06	Impl Spec	PCIe Port Number: This field contains the PCIe port number. This is the same value as that reported in the Port Number field in the PCIe Link Capabilities Register.

Figure 149: NVMe PCIe Port MultiRecord Area

Byte Offset	Factory Default	Description																														
07	Impl Spec	<p>Port Information: This field indicates information about the PCIe Ports in the device.</p> <p>Bits 7:1 are reserved.</p> <p>Bit 0, if set to '1' indicates that all PCIe ports within the device have the same capabilities (i.e., the capabilities listed in this structure are consistent across each PCIe port).</p>																														
08	Impl Spec	<p>PCIe Link Speed: This field indicates a bit vector of link speeds supported by the PCIe port.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>7:4</td> <td>Reserved</td> </tr> <tr> <td>3</td> <td>Set to '1' if the PCIe link supports 16.0 GT/s, otherwise cleared to '0'.</td> </tr> <tr> <td>2</td> <td>Set to '1' if the PCIe link supports 8.0 GT/s, otherwise cleared to '0'.</td> </tr> <tr> <td>1</td> <td>Set to '1' if the PCIe link supports 5.0 GT/s, otherwise cleared to '0'.</td> </tr> <tr> <td>0</td> <td>Set to '1' if the PCIe link supports 2.5 GT/s, otherwise cleared to '0'.</td> </tr> </tbody> </table>	Bit	Definition	7:4	Reserved	3	Set to '1' if the PCIe link supports 16.0 GT/s, otherwise cleared to '0'.	2	Set to '1' if the PCIe link supports 8.0 GT/s, otherwise cleared to '0'.	1	Set to '1' if the PCIe link supports 5.0 GT/s, otherwise cleared to '0'.	0	Set to '1' if the PCIe link supports 2.5 GT/s, otherwise cleared to '0'.																		
Bit	Definition																															
7:4	Reserved																															
3	Set to '1' if the PCIe link supports 16.0 GT/s, otherwise cleared to '0'.																															
2	Set to '1' if the PCIe link supports 8.0 GT/s, otherwise cleared to '0'.																															
1	Set to '1' if the PCIe link supports 5.0 GT/s, otherwise cleared to '0'.																															
0	Set to '1' if the PCIe link supports 2.5 GT/s, otherwise cleared to '0'.																															
09	Impl Spec	<p>PCIe Maximum Link Width: The maximum PCIe link width for this NVM Subsystem port. This is the expected negotiated link width that the port link trains to if the platform supports it. A Requester may compare this value with the PCIe Negotiated Link Width to determine if there has been a PCIe link training issue.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Reserved</td> </tr> <tr> <td>1</td> <td>PCIe x1</td> </tr> <tr> <td>2</td> <td>PCIe x2</td> </tr> <tr> <td>3</td> <td>Reserved</td> </tr> <tr> <td>4</td> <td>PCIe x4</td> </tr> <tr> <td>5 to 7</td> <td>Reserved</td> </tr> <tr> <td>8</td> <td>PCIe x8</td> </tr> <tr> <td>9 to 11</td> <td>Reserved</td> </tr> <tr> <td>12</td> <td>PCIe x12</td> </tr> <tr> <td>13 to 15</td> <td>Reserved</td> </tr> <tr> <td>16</td> <td>PCIe x16</td> </tr> <tr> <td>17 to 31</td> <td>Reserved</td> </tr> <tr> <td>32</td> <td>PCIe x32</td> </tr> <tr> <td>33 to 255</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Definition	0	Reserved	1	PCIe x1	2	PCIe x2	3	Reserved	4	PCIe x4	5 to 7	Reserved	8	PCIe x8	9 to 11	Reserved	12	PCIe x12	13 to 15	Reserved	16	PCIe x16	17 to 31	Reserved	32	PCIe x32	33 to 255	Reserved
Value	Definition																															
0	Reserved																															
1	PCIe x1																															
2	PCIe x2																															
3	Reserved																															
4	PCIe x4																															
5 to 7	Reserved																															
8	PCIe x8																															
9 to 11	Reserved																															
12	PCIe x12																															
13 to 15	Reserved																															
16	PCIe x16																															
17 to 31	Reserved																															
32	PCIe x32																															
33 to 255	Reserved																															
10	Impl Spec	<p>MCTP Support: This field contains a bit vector that specifies the level of support for the NVMe Management Interface.</p> <p>Bits 7:1 are reserved.</p> <p>Bit 0, if set to '1' indicates that MCTP-based management commands are supported on the PCIe port.</p>																														
11	Impl Spec	<p>Ref Clk Capability: This field contains a bit vector that specifies the PCIe clocking modes supported by the port.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>7:4</td> <td>Reserved</td> </tr> <tr> <td>3</td> <td>Set to '1' if the device automatically uses RefClk if provided and otherwise uses SRIS, otherwise cleared to '0'.</td> </tr> <tr> <td>2</td> <td>Set to '1' if the PCIe link supports Separate RefClk with SSC (SRIS), otherwise cleared to '0'.</td> </tr> <tr> <td>1</td> <td>Set to '1' if the PCIe link supports Separate RefClk with no SSC (SRNS), otherwise cleared to '0'.</td> </tr> <tr> <td>0</td> <td>Set to '1' if the PCIe link supports common RefClk, otherwise cleared to '0'.</td> </tr> </tbody> </table>	Bit	Definition	7:4	Reserved	3	Set to '1' if the device automatically uses RefClk if provided and otherwise uses SRIS, otherwise cleared to '0'.	2	Set to '1' if the PCIe link supports Separate RefClk with SSC (SRIS), otherwise cleared to '0'.	1	Set to '1' if the PCIe link supports Separate RefClk with no SSC (SRNS), otherwise cleared to '0'.	0	Set to '1' if the PCIe link supports common RefClk, otherwise cleared to '0'.																		
Bit	Definition																															
7:4	Reserved																															
3	Set to '1' if the device automatically uses RefClk if provided and otherwise uses SRIS, otherwise cleared to '0'.																															
2	Set to '1' if the PCIe link supports Separate RefClk with SSC (SRIS), otherwise cleared to '0'.																															
1	Set to '1' if the PCIe link supports Separate RefClk with no SSC (SRNS), otherwise cleared to '0'.																															
0	Set to '1' if the PCIe link supports common RefClk, otherwise cleared to '0'.																															
12	Impl Spec	<p>Port Identifier: This field contains the NVMe-MI Port Identifier.</p>																														

Figure 149: NVMe PCIe Port MultiRecord Area

Byte Offset	Factory Default	Description
15:13	00h	Reserved

9.2.5 Topology MultiRecord Area

This MultiRecord describes an NVMe Storage Device’s architectural elements and their connections. It is required on all NVMe Storage Device FRUs.

The Topology MultiRecord consists mainly of a list of Element Descriptors as shown in Figure 150. Element Descriptors are used to describe the architectural elements that make up an NVMe Storage Device such as NVM Subsystems, Upstream Connectors, Expansion Connectors, SMBus/I2C elements, and PCIe elements. Each architectural element has an Element Descriptor Type. The format of an Element Descriptor is shown in Figure 152 and Element Descriptor Types are listed in Figure 153.

Element Descriptors may have fields that are used to point to other Element Descriptors. When an Element Descriptor contains a pointer to another Element Descriptor, then the Element Descriptor containing the pointer is called the parent and the Element Descriptor pointed to by the parent is called the child. An Element Descriptor may be both a child and a parent.

An Element Descriptor pointer is either populated with an index of the child or 0h to indicate that there is no child. The index is a logical construct that indicates the position of an Element Descriptor in the VPD. The Element Descriptor at the lowest byte offset in the VPD has an index of 0, the Element Descriptor at the second lowest byte offset has an index of 1, and so on. A child may have an index that is higher or lower than its parent. The Element Descriptor at the lowest byte offset (i.e., index 0) shall be an Upstream Connector Element Descriptor. Some Element Descriptors use indexes in a similar manner to select a Port from a list of Ports.

Figure 150: Topology MultiRecord

Byte Offset	Factory Default	Description
00	0Dh	Topology Record Type ID
01	2h or 82h	Record Format:
		Bit Definition
		7 Set to '1' if last record in list.
6:0	Record format version = 2.	
02	Impl Spec	Record Length (RLEN): This field indicates the length of the MultiRecord Area in bytes without including the first 5 bytes that are common to all MultiRecords.
03	Impl Spec	Record Checksum: This field is used to give the record data a zero checksum (i.e., the modulo 256 sum of the record data bytes from byte offset 05 to the end of this record plus this checksum byte equals zero).
04	Impl Spec	Header Checksum: This field is used to give the record header a zero checksum (i.e., the modulo 256 sum of the first byte of the header through this checksum byte equals zero).
05	0h	Version Number: This field indicates the version number of this Topology MultiRecord. This field shall be cleared to 0h in this version of the specification.
06	0h	Reserved
07	Impl Spec	Element Count (N): This field indicates the number of Element Descriptors in this Topology MultiRecord. The value of 0h is reserved.
Impl Spec	Impl Spec	Element Descriptor 0: This field contains the first Element Descriptor in this Topology MultiRecord.

Figure 150: Topology MultiRecord

Byte Offset	Factory Default	Description
Impl Spec	Impl Spec	Element Descriptor 1: This field contains the second Element Descriptor in this Topology MultiRecord if Element Count is greater than one, otherwise this field is not present.
...

The VPD may contain more than one Topology MultiRecord only when the list of required Element Descriptors is too large to fit into a single Topology MultiRecord. If there is more than one Topology MultiRecord, then the index associated with Element Descriptors continues to increment sequentially across Topology MultiRecord instances. Figure 151 illustrates multiple Topology MultiRecords where Index 0 is at the lowest byte offset of any Element Descriptor in the VPD. Parent Element Descriptors may be in different Topology MultiRecords from their Child Element Descriptors.

Figure 151: Indexing Across Extended MultiRecords

Index	Topology Multi Record Instance	Element Descriptors	Child Indices
0	0	Element Descriptor 0, parent of 2, 3, 5	2, 3, 5
1		Element Descriptor 1, child of 5	
2		Element Descriptor 2, child of 0	
3		Element Descriptor 3, child of 0	
4	1	Element Descriptor 0 ¹	
5		Element Descriptor 1, child of 0, parent of 1	1
NOTES: 1. This Element Descriptor is an Extended Element Descriptor that extends the preceding Element Descriptor at index 3. Extended Element Descriptors are further detailed in section 9.2.5.1.			

Figure 152: Element Descriptor

Byte Offset	Factory Default	Description
00	Impl Spec	Type: This field indicates the type of the Element Descriptor. Values are defined in Figure 153.
01	Impl Spec	Revision: This field indicates the revision of the Element Descriptor.
02	Impl Spec	Length: Number of bytes in the Element Descriptor.
Length - 1:03	Impl Spec	This area contains the Type-specific information associated with the Element Descriptor. Type-specific information is defined for each Element Descriptor Type in the subsections below.

Element Descriptor Types, fields, and bits in the VPD that are defined as reserved should be ignored by Requesters to ensure forward and backward compatibility. Extra trailing bytes in an Element Descriptor should be treated as reserved in order to tolerate the Length of an Element Descriptor increasing as new fields are appended in future revisions of the Element Descriptor.

Element Descriptor Types are defined in Figure 153. Subsequent sections define the details associated with each Element Descriptor Type.

Figure 153: Element Descriptor Types

Value	Name	Reference Section
0	Reserved	-
1	Extended Element Descriptor	9.2.5.1
2	Upstream Connector Element Descriptor	9.2.5.2
3	Expansion Connector Element Descriptor	9.2.5.3
4	Label Element Descriptor	9.2.5.4
5	SMBus/I2C Mux Element Descriptor	9.2.5.5
6	PCIe Switch Element Descriptor	9.2.5.6
7	NVM Subsystem Element Descriptor	9.2.5.7
8 to 239	Reserved	-
240 to 255	Vendor specific	9.2.5.8

9.2.5.1 Extended Element Descriptor

The Extended Element Descriptor is shown in Figure 154. This Element Descriptor Type shall only be used when an Element Descriptor spans across more than one Topology MultiRecord. Extended Element Descriptors shall not be the children of other Element Descriptors.

If an Element Descriptor causes the maximum size of a Topology MultiRecord to be exceeded, then that Element Descriptor is truncated so that the non-truncated portion of the Element Descriptor fits into the Topology MultiRecord. The truncated portion of the Element Descriptor forms the contents of the Extended Content field in an Extended Element Descriptor. That Extended Element Descriptor is the first Element Descriptor in the next Topology MultiRecord. If the truncated portion of the Element Descriptor does not fit into a single Topology MultiRecord, then two or more Extended Element Descriptors are required, each in subsequent Topology MultiRecords.

An example is shown in Figure 151 where the Element Descriptor at index 4 is an Extended Element Descriptor that extends the Element Descriptor at index 3. Element Descriptor 3 is the child of Element Descriptor 0 and Element Descriptor 4 is not the child of any parent Element Descriptor.

Figure 154: Extended Element Descriptor

Byte Offset	Factory Default	Description
00	01h	Type: This field indicates the type of the Element Descriptor. The Extended Element Descriptor Type is 1h.
01	00h	Revision: This field indicates the revision of the Element Descriptor. The Extended Element Descriptor Revision is 0h for this specification.
02	Impl Spec	Length: This field indicates the length of the Extended Element Descriptor in bytes.
Length - 1:03	Impl Spec	Extended Content: This field extends the content of the Element Descriptor at the immediately preceding index.

9.2.5.2 Upstream Connector Element Descriptor

The Upstream Connector Element Descriptor is shown in Figure 155 and is used to describe an Upstream Connector (i.e., a connector through which a Requester communicates with the NVMe Storage Device). Upstream Element Descriptors are always a parent and never a child.

Figure 155: Upstream Connector Element Descriptor

Byte Offset	Factory Default	Description
00	02h	Type: This field indicates the type of the Element Descriptor. The Upstream Connector Element Descriptor Type is 2h.
01	00h	Revision: This field indicates the revision of the Element Descriptor. The Upstream Connector Element Descriptor Revision is 0h for this specification.
02	Impl Spec	Length: This field indicates the length of the entire Upstream Connector Element Descriptor in bytes.
03	Impl Spec	Form Factor: This field indicates the Form Factor of the NVMe Storage Device. See Figure 156 for a list of defined values.
04	Impl Spec	Label Pointer: If the Upstream Connector has a label, then this field shall contain the index of a Label Element Descriptor that contains the label. The value 0h indicates there is no associated label.
06:05	00h	Reserved
07	Impl Spec	Maximum Auxiliary Power: This field specifies the maximum auxiliary power supply requirements in 10 mW increments consumed by the NVMe Storage Device. A value of 0h indicates that auxiliary power is not used from this Upstream Connector.
09:08	Impl Spec	Maximum Power: This field specifies the maximum power in Watts consumed by the NVMe Storage Device.
10	Impl Spec	Upstream Port Descriptor Count: This field indicates the number of Upstream Port Descriptors associated with this Upstream Connector Element Descriptor. The permitted range of values is 1 to 64.
Impl Spec	Impl Spec	Upstream Port Descriptor 0: This field contains the first Upstream Port Descriptor.
Impl Spec	Impl Spec	Upstream Port Descriptor 1: This field contains the second Upstream Port Descriptor in this Upstream Connector Element Descriptor if Port Descriptor Count is greater than one, otherwise this field is not present.
...

The value of the Form Factor field indicates the NVMe Storage Device’s form factor. Figure 156 lists the NVMe Storage Device’s Form Factor values.

Figure 156: Form Factors

Value	Description
0	Other – unknown
1	Integrated
2 to 15	Reserved
16	2.5” Form Factor – unknown
17	2.5” Form Factor – U.2 (SFF-8639) 15 mm
18	2.5” Form Factor – U.2 (SFF-8639) 7 mm
19	2.5” Form Factor – (SFF-TA-1001) 15 mm
20	2.5” Form Factor – (SFF-TA-1001) 7 mm
21 to 31	Reserved
32	CEM add in card – unknown
33	CEM add in card – Low Profile (HHHL)
34	CEM add in card – Standard Height Half Length (FHHL)
35	CEM add in card – Standard Height Full Length (FHFL)
36 to 47	Reserved
48	M.2 module – unknown
49	M.2 module – 2230
50	M.2 module – 2242
51	M.2 module – 2260
52	M.2 module – 2280
53	M.2 module – 22110

Figure 156: Form Factors

Value	Description
54 to 63	Reserved
64	BGA SSD – unknown
65	BGA SSD – 16 x 20mm (M.2 Type 1620)
66	BGA SSD – 11.5 x 13mm (M.2 Type 1113)
67 to 79	Reserved
80	Enterprise & Datacenter SSD Form Factor – unknown
81	1U Short Form Factor - (SFF-TA-1006) 5.9 mm
82	1U Short Form Factor - (SFF-TA-1006) 8 mm
83	1U Long Form Factor - (SFF-TA-1007) 9.5 mm
84	1U Long Form Factor - (SFF-TA-1007) 18 mm
85	3" Short Form Factor - (SFF-TA-1008) 7.5 mm
86	3" Short Form Factor - (SFF-TA-1008) 16.8 mm
87	3" Long Form Factor - (SFF-TA-1008) 7.5 mm
88	3" Long Form Factor - (SFF-TA-1008) 16.8 mm
89 to 239	Reserved
240 to 255	Vendor Specific

The Upstream Connector may have an associated label, such as silk-screened text on the printed circuit board. If the Upstream Connector has a label, then the Label Pointer may contain the index of the associated Label Element Descriptor.

The Upstream Connector Element Descriptor contains a list of the Upstream Port Descriptors that are ports through which a Requester communicates with the NVMe Storage Device. Each Upstream Port Descriptor has a type. The types defined in this specification are SMBus/I2C Upstream Port Descriptor and PCIe Upstream Port Descriptor.

An SMBus/I2C Upstream Port Descriptor is shown in Figure 157. It contains a list of pointers to child Element Descriptors whose SMBus/I2C port is directly connected to the Upstream Connector.

A PCIe Upstream Port Descriptor is shown in Figure 158. It indicates the starting and ending PCIe lane numbers on the Upstream Connector that make up a PCIe Upstream Port. The PCIe Upstream Port Descriptor contains a single pointer to a child Element Descriptor connected to this PCIe Upstream Port. The Destination Port field of the PCIe Upstream Port Element Descriptor specifies which port of the child is connected to this Upstream Connector. The Destination Port value is an index into the child Element Descriptor's list of Port Descriptors.

The PCIe lanes associated with a PCIe Upstream Connector may be organized as a single large port or subdivided into multiple ports. Each of these ports is described with its own PCIe Upstream Port Descriptor. The PCIe Upstream Port Descriptors may be listed in any order. A form factor specific mechanism, such as the U.2 Dual Port Enable signal, may be used to determine which of the listed PCIe Upstream Port Descriptors are currently applicable. These form factor specific mechanisms are outside the scope of this specification.

For example, a U.2 NVMe Storage Device capable of running in either single-port mode or dual-port mode based on the Dual Port Enable signal would have three PCIe Upstream Port Descriptors describing PCIe ports on the following PCIe Lanes:

1. PCIe lanes 0 to 3 (single-port mode);
2. PCIe lanes 0 to 1 (dual-port mode); and
3. PCIe lanes 2 to 3 (dual-port mode).

In the example above, if the U.2 NVMe Storage Device is only capable of running in single-port mode, then only the PCIe Upstream Port Descriptor describing the single-port mode (item 1 in the list above) shall be included in the Upstream Connector Element Descriptor. And if the U.2 NVMe Storage Device is only capable of running in dual-port mode, then only the two PCIe Upstream Port Descriptors describing the dual-port mode (items 2 and 3 in the list above) shall be included in the Upstream Connector Element Descriptor.

In another example, consider a x16 CEM add-in card Upstream Connector that is subdivided into four x4 PCIe ports, also referred to as bifurcation. Each of these x4 PCIe Upstream Ports may connect to different elements on the NVMe Storage Device. The Upstream Connector in this example shall contain four PCIe Upstream Port Descriptors describing the four PCIe ports:

1. PCIe lanes 0 to 3;
2. PCIe lanes 4 to 7;
3. PCIe lanes 8 to 11; and
4. PCIe lanes 12 to 15.

Figure 157: SMBus/I2c Upstream Port Descriptor

Byte Offset	Factory Default	Description
00	00h	Type: This field indicates the type of the Port Descriptor. The SMBus/I2C Port Descriptor Type is 0h.
01	Impl Spec	Length: This field indicates the length of the SMBus/I2C Port Descriptor in bytes.
02	Impl Spec	Count: This field indicates the number of SMBus/I2C Pointers in the SMBus/I2C Upstream Port Descriptor. The permitted range of values is 1 to 32.
03	Impl Spec	SMBus/I2C Pointer 0: This field contains the child index of the first Element Descriptor whose SMBus/I2C port is connected to this SMBus/I2C port.
04	Impl Spec	SMBus/I2C Pointer 1: If Count is greater than one, then this field is present and contains the child index of the second Element Descriptor whose SMBus/I2C port is connected to this SMBus/I2C Upstream Port. If Count is not greater than one, then this field is not present.
...

Figure 158: PCIe Upstream Port Descriptor

Byte Offset	Factory Default	Description
00	01h	Type: This field indicates the type of Upstream Port Descriptor. The PCIe Upstream Port Descriptor Type is 1h.
01	Impl Spec	Length: This field indicates the length of the PCIe Upstream Port Descriptor in bytes.
02	Impl Spec	Starting Lane: This field indicates first PCIe lane (i.e., lane 0) of the port from the Upstream Connector.
03	Impl Spec	Ending Lane: This field indicates the ending PCIe lane of the port from the Upstream Connector.
04	Impl Spec	PCIe Pointer: This field contains the child index of the Element Descriptor whose PCIe port is connected to this PCIe Upstream Port.
05	Impl Spec	Destination Port: This field contains the index of the Port Descriptor in the child Element Descriptor. If the child Element Descriptor has one PCIe upstream port (i.e., a PCIe Switch Element Descriptor) this field shall be cleared to 0h.
...

9.2.5.3 Expansion Connector Element Descriptor

The Expansion Connector Element Descriptor is shown in Figure 159 and is used to describe the form factor, label, and port configurations for Expansion Connectors on a Carrier. The Expansion Connector Element Descriptor shall be a child Element Descriptor.

Figure 159: Expansion Connector Element Descriptor

Byte Offset	Factory Default	Description
00	03h	Type: This field indicates the type of the Element Descriptor. The Expansion Connector Element Descriptor Type is 3h.
01	00h	Revision: This field indicates the revision of the Element Descriptor. The Expansion Connector Element Descriptor Revision is 0h for this specification.
02	Impl Spec	Length: This field indicates the length of the Expansion Connector Element Descriptor in bytes.
03	Impl Spec	Form Factor: This field indicates the Form Factor of the NVMe Storage Device FRU that plugs into the Expansion Connector. Refer to Figure 156 for a list of defined values.
04	Impl Spec	Label Pointer: If the Upstream Connector has a label, then this field shall contain the index of a Label Element Descriptor that contains the label. The value 0h indicates there is no associated label.
05	Impl Spec	Expansion Connector Port Descriptor Count: This field indicates the number of Expansion Port Descriptors associated with this Expansion Connector Element Descriptor. The permitted range of values is 1 to 64.
Impl Spec	Impl Spec	Expansion Connector Port Descriptor 0: This field contains the first Expansion Connector Port Descriptor.
Impl Spec	Impl Spec	Expansion Connector Port Descriptor 1: This field contains the second Expansion Connector Port Descriptor in this Expansion Connector Descriptor if Expansion Connector Port Descriptor Count is greater than one, otherwise this field is not present.
...

In a manner similar to the PCIe Upstream Connector, the Expansion Connector Element Descriptor's PCIe lanes may support one or more PCIe ports for connecting to external NVMe Storage Device FRUs. The PCIe ports have a starting and ending PCIe lane number on the Expansion Connector that are determined by the external NVMe Storage Device FRU's form factor's lane numbering.

The Expansion Connector Element Descriptor holds the list of Expansion Connector PCIe Port Descriptors. Each PCIe port is described by an Expansion Connector PCIe Port Descriptor whose format is shown in Figure 160. Parent Element Descriptors, such as Upstream Connectors and PCIe Switches, contain Port Descriptors that point to Expansion Connectors. The Destination Port field of the parent Port Descriptor contains an index to the specific Expansion Connector PCIe Port Descriptor instance to which the Port Descriptor is connected. Each Expansion Connector PCIe Port Descriptor is the destination of exactly one pointer from a parent Element Descriptor.

Figure 160: Expansion Connector PCIe Port Descriptor

Byte Offset	Factory Default	Description
00	00h	Type: This field indicates the type of Expansion Connector Port Descriptor. The Expansion Connector PCIe Port Descriptor Type is 0.
01	Impl Spec	Length: This field indicates the length of the Expansion Connector PCIe Port Descriptor in bytes.
02	Impl Spec	Starting Lane: This field indicates first PCIe lane (i.e., lane 0) of the port on the Expansion Connector PCIe Port Descriptor.
03	Impl Spec	Ending Lane: This field indicates the ending PCIe lane of the port on the Expansion Connector PCIe Port Descriptor.

9.2.5.4 Label Element Descriptor

The Label Element Descriptor is shown in Figure 161 and is used to store text strings in the VPD for Element Descriptors that have a label. A Label Element Descriptor shall be a child Element Descriptor.

Figure 161: Label Element Descriptor

Byte Offset	Factory Default	Description
00	04h	Type: This field indicates the type of the Element Descriptor. The Label Element Descriptor Type is 4h.
01	00h	Revision: This field indicates the revision of the Element Descriptor. The Label Element Descriptor Revision is 0h for this specification.
02	Impl Spec	Length: This field indicates the length of the Label Element Descriptor in bytes including the null termination.
Length - 1:03	Impl Spec	Label String: This field contains a null-terminated UTF-8 string used to identify the parent Element Descriptor.

9.2.5.5 SMBus/I2C Mux Element Descriptor

The SMBus/I2C Mux Element Descriptor is shown in Figure 162 and is used to describe an SMBus/I2C multiplexor element that connects a single upstream SMBus/I2C channel to zero or more downstream SMBus/I2C channels. This Element Descriptor contains the address and capabilities of the SMBus/I2C Mux followed by a list of SMBus/I2C Mux Channel Descriptors that describe SMBus/I2C Mux downstream channel connections. The SMBus/I2C Mux shall be compatible with the industry standard PCA9542/45/48 family of SMBus/I2C multiplexors and may be extended to support ARP, error detection, and additional downstream channels as defined below.

Figure 162: SMBus/I2C Mux Element Descriptor

Byte Offset	Factory Default	Description						
00	05h	Type: This field indicates the type of the Element Descriptor. The SMBus/I2C Mux Element Descriptor Type is 5h.						
01	00h	Revision: This field indicates the revision of the Element Descriptor. The SMBus/I2C Mux Element Descriptor Revision is 0h for this specification.						
02	Impl Spec	Length: This field indicates the length of the SMBus/I2C Mux Element Descriptor in bytes.						
03	E8h or E9h	SMBus/I2C Address Info: This field indicates the SMBus/I2C address and whether or not ARP is supported.						
		<table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:1</td> <td>SMBus/I2C Address: This field contains the 7-bit SMBus/I2C address. Refer to Figure 15 for requirements.</td> </tr> <tr> <td>0</td> <td>ARP Capable: This bit is set to '1' if SMBus ARP is supported, else it is cleared to '0'. Refer to Figure 15 for requirements.</td> </tr> </tbody> </table>	Bit	Description	7:1	SMBus/I2C Address: This field contains the 7-bit SMBus/I2C address. Refer to Figure 15 for requirements.	0	ARP Capable: This bit is set to '1' if SMBus ARP is supported, else it is cleared to '0'. Refer to Figure 15 for requirements.
		Bit	Description					
7:1	SMBus/I2C Address: This field contains the 7-bit SMBus/I2C address. Refer to Figure 15 for requirements.							
0	ARP Capable: This bit is set to '1' if SMBus ARP is supported, else it is cleared to '0'. Refer to Figure 15 for requirements.							

Figure 162: SMBus/I2C Mux Element Descriptor

Byte Offset	Factory Default	Description																				
04	Impl Spec	SMBus/I2C Capabilities: This field indicates the SMBus/I2C Mux capabilities.																				
		<table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7</td> <td>Form Factor Reset: This bit is set to '1' if all of the SMBus/I2C reset mechanisms are supported as defined by the associated form factor specification. This bit is cleared to '0' if the form factor does not define SMBus Reset or the NVMe Storage Device does not support all of the SMBus/I2C reset mechanisms defined in the specification for the Form Factor in the Host Connector Element Descriptor.</td> </tr> <tr> <td>6</td> <td>Packet Error Code (PEC) Support: This bit is set to '1' if PEC is supported by the SMBus/I2C Mux. This bit is cleared to '0' if PEC is not supported.</td> </tr> <tr> <td>5:2</td> <td>Reserved</td> </tr> <tr> <td>1:0</td> <td>Maximum Speed: This field is set to the highest supported SMBus/I2C clock speed by the SMBus/I2C Mux. <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>100 kHz</td> </tr> <tr> <td>1</td> <td>400 kHz</td> </tr> <tr> <td>2</td> <td>1 MHz</td> </tr> <tr> <td>3</td> <td>Reserved</td> </tr> </tbody> </table> </td> </tr> </tbody> </table>	Bit	Description	7	Form Factor Reset: This bit is set to '1' if all of the SMBus/I2C reset mechanisms are supported as defined by the associated form factor specification. This bit is cleared to '0' if the form factor does not define SMBus Reset or the NVMe Storage Device does not support all of the SMBus/I2C reset mechanisms defined in the specification for the Form Factor in the Host Connector Element Descriptor.	6	Packet Error Code (PEC) Support: This bit is set to '1' if PEC is supported by the SMBus/I2C Mux. This bit is cleared to '0' if PEC is not supported.	5:2	Reserved	1:0	Maximum Speed: This field is set to the highest supported SMBus/I2C clock speed by the SMBus/I2C Mux. <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>100 kHz</td> </tr> <tr> <td>1</td> <td>400 kHz</td> </tr> <tr> <td>2</td> <td>1 MHz</td> </tr> <tr> <td>3</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Description	0	100 kHz	1	400 kHz	2	1 MHz	3	Reserved
		Bit	Description																			
		7	Form Factor Reset: This bit is set to '1' if all of the SMBus/I2C reset mechanisms are supported as defined by the associated form factor specification. This bit is cleared to '0' if the form factor does not define SMBus Reset or the NVMe Storage Device does not support all of the SMBus/I2C reset mechanisms defined in the specification for the Form Factor in the Host Connector Element Descriptor.																			
6	Packet Error Code (PEC) Support: This bit is set to '1' if PEC is supported by the SMBus/I2C Mux. This bit is cleared to '0' if PEC is not supported.																					
5:2	Reserved																					
1:0	Maximum Speed: This field is set to the highest supported SMBus/I2C clock speed by the SMBus/I2C Mux. <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>100 kHz</td> </tr> <tr> <td>1</td> <td>400 kHz</td> </tr> <tr> <td>2</td> <td>1 MHz</td> </tr> <tr> <td>3</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Description	0	100 kHz	1	400 kHz	2	1 MHz	3	Reserved											
Value	Description																					
0	100 kHz																					
1	400 kHz																					
2	1 MHz																					
3	Reserved																					
05	Impl Spec	SMBus/I2C Mux Channel Descriptor Count: This field indicates the number of downstream channels listed for this SMBus/I2C Mux. Each channel has a corresponding SMBus/I2C Channel Descriptor in the list below. The permitted range of values is 1 to 64. The value of this field may be less than the actual number of Channels implemented by the SMBus/I2C Mux if the truncated SMBus/I2C Mux Channel Descriptors are not connected to anything.																				
Impl Spec	Impl Spec	SMBus/I2C Mux Channel Descriptor 0: This field contains the first SMBus/I2C Mux Channel Descriptor.																				
Impl Spec	Impl Spec	SMBus/I2C Mux Channel Descriptor 1: This field contains the second SMBus/I2C Mux Channel Descriptor in this SMBus/I2C Mux Element Descriptor if SMBus/I2C Mux Channel Descriptor Count is greater than one, otherwise this field is not present.																				
...																				

An SMBus/I2C Mux Channel Descriptor is shown in Figure 164. SMBus/I2C Mux Channel Descriptors that are not connected to anything have a value 0h in the Count field and contain no SMBus/I2C Mux Channel Descriptors. Unconnected SMBus/I2C Mux Channel Descriptors at the end of the list in Figure 162 may be truncated unless they are needed to position the optional Packet Error Code (PEC).

Writing to an SMBus/I2C Mux configures the SMBus/I2C Mux and reading from an SMBus Mux returns its current configuration. Figure 163 shows the protocol for reading and writing an SMBus/I2C Mux configuration. The white background blocks are transmitted by a Management Controller and the grey background blocks are transmitted in response by the SMBus/I2C Mux. The first byte sent or received is the SMBus/I2C Mux address followed by one or more channel bytes. Each channel byte has eight channel bits that are set to '1' for connecting the corresponding downstream channel to the upstream channel or cleared to '0' for disconnecting the corresponding downstream channel from the upstream channel.

The first channel byte sent or received represents channels 0 to 7, the second channel byte sent or received represents channels 8 to 15, and so on. Within each channel byte the least significant bit in the byte that is transmitted or received represents the lowest numbered channel. Bits for channels exceeding the SMBus/I2C Mux Channel Descriptor Count are reserved.

Figure 163: SMBus/I2C Mux Read and Write Command Format



The minimum number of channel bytes are read or written to reach all the channels specified in the SMBus/I2C Mux Channel Descriptor Count field. Thus, SMBus/I2C Muxes with one to eight downstream channels would have one channel byte while an SMBus/I2C Mux with 57 downstream channels would have 8 channel bytes. In the example shown in Figure 163, the SMBus/I2C Mux has 16 downstream channels that require 2 bytes. In this example, channels 1 and 8 are being connected while all others are being disconnected.

An SMBus/I2C Mux may also protect communications with an optional Packet Error Code (PEC) that is appended after sufficient channel bytes have been read or written to satisfy the SMBus/I2C Mux Channel Descriptor Count value. If the write command includes a PEC byte and the PEC byte is incorrect, then the entire command shall be ignored by the SMBus/I2C Mux, otherwise the actions associated with the write command take place after the STOP condition is received. Write commands with insufficient channel bytes shall be accepted with truncated channel bytes having an implied value of zero. Bytes beyond the size needed for the number of channels and PEC are reserved.

Multiple downstream channels may be simultaneously connected to the upstream channel to bridge them together. All downstream channels shall be disconnected when the NVMe Storage Device is powered off (refer to Figure 143) or by an SMBus Reset (refer to section 9.3.4). Connecting or disconnecting channels while they are active is strongly discouraged and results in undefined behavior.

Figure 164: SMBus/I2C Mux Channel Descriptor

Byte Offset	Factory Default	Description
00	00h	Type: This field indicates the type of the Descriptor. The SMBus/I2C Mux Channel Descriptor Type is 0h.
01	Impl Spec	Length: This field indicates the length of the SMBus/I2C Mux Channel Descriptor in bytes.
02	Impl Spec	Count: This field indicates the number of SMBus/I2C Pointers in the SMBus/I2C Mux Channel Descriptor. The permitted range of values is 0 to 32.
03	Impl Spec	SMBus/I2C Pointer 0: This field contains the child index of the first Element Descriptor whose SMBus/I2C is connected to this channel.
04	Impl Spec	SMBus/I2C Pointer 1: If Count is greater than one, then this field is present and contains the child index of another Element Descriptor whose SMBus/I2C is connected to this channel. If Count is not greater than one, then this field is not present.
...

9.2.5.6 PCIe Switch Element Descriptor

The PCIe Switch Element Descriptor is shown in Figure 165 and is used to describe a PCIe switch. This Element Descriptor is the child of a single parent and the parent of one or more children.

Figure 165: PCIe Switch Element Descriptor

Byte Offset	Factory Default	Description
00	06h	Type: This field indicates the type of the Element Descriptor. The PCIe Switch Element Descriptor Type is 6h.
01	Impl Spec	Revision: This field indicates the revision of the Element Descriptor. The PCIe Switch Element Descriptor Revision is 0h for this specification.

Figure 165: PCIe Switch Element Descriptor

Byte Offset	Factory Default	Description
02	Impl Spec	Length: This field indicates the length of the PCIe Switch Element Descriptor in bytes.
03	Impl Spec	Upstream Switch Port Descriptor: This field contains the PCIe Switch Port Descriptor that describes the upstream switch port.
Impl Spec	Impl Spec	Downstream Switch Port Descriptor Count: This field indicates the number of PCIe Port Descriptors associated with downstream switch ports.
Impl Spec	Impl Spec	Downstream Switch Port Descriptor 0: This field contains the PCIe Switch Port Descriptor associated with the first downstream port.
Impl Spec	Impl Spec	Downstream Switch Port Descriptor 1: This field contains the PCIe Switch Port Descriptor associated with the second downstream port if Downstream Switch Port Descriptor Count is greater than one, otherwise this field is not present.
...

The PCIe Switch Element Descriptor consists of a list of PCIe Switch Port Descriptors. There is an Upstream Switch Port Descriptor that describes the upstream port and is the child of exactly one parent Element Descriptor. A variable length list of Downstream Switch Port Descriptors describes the downstream ports.

The format of a PCIe Switch Port Descriptor is shown in Figure 166. It describes the PCIe port's supported PCIe link speeds, PCIe maximum link width, reference clock capabilities, and PCIe Port Number. Downstream ports also have a child Element Descriptor and its Destination Port index value.

Figure 166: PCIe Switch Port Descriptor

Byte Offset	Factory Default	Description												
00	00h	Type: This field indicates the type of Port Descriptor. The PCIe Switch Port Descriptor Type is 0.												
01	Impl Spec	Length: This field indicates the length of the PCIe Switch Port Descriptor in bytes.												
02	Impl Spec	<p>PCIe Link Speed: This field indicates a bit vector of link speeds supported by the PCIe port.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:4</td> <td>Reserved</td> </tr> <tr> <td>3</td> <td>Set to '1' if the PCIe link supports 16 GT/s, otherwise cleared to '0'.</td> </tr> <tr> <td>2</td> <td>Set to '1' if the PCIe link supports 8.0 GT/s, otherwise cleared to '0'.</td> </tr> <tr> <td>1</td> <td>Set to '1' if the PCIe link supports 5.0 GT/s, otherwise cleared to '0'.</td> </tr> <tr> <td>0</td> <td>Set to '1' if the PCIe link supports 2.5 GT/s, otherwise cleared to '0'.</td> </tr> </tbody> </table>	Bit	Description	7:4	Reserved	3	Set to '1' if the PCIe link supports 16 GT/s, otherwise cleared to '0'.	2	Set to '1' if the PCIe link supports 8.0 GT/s, otherwise cleared to '0'.	1	Set to '1' if the PCIe link supports 5.0 GT/s, otherwise cleared to '0'.	0	Set to '1' if the PCIe link supports 2.5 GT/s, otherwise cleared to '0'.
Bit	Description													
7:4	Reserved													
3	Set to '1' if the PCIe link supports 16 GT/s, otherwise cleared to '0'.													
2	Set to '1' if the PCIe link supports 8.0 GT/s, otherwise cleared to '0'.													
1	Set to '1' if the PCIe link supports 5.0 GT/s, otherwise cleared to '0'.													
0	Set to '1' if the PCIe link supports 2.5 GT/s, otherwise cleared to '0'.													

Figure 166: PCIe Switch Port Descriptor

Byte Offset	Factory Default	Description																														
03	Impl Spec	<p>PCIe Maximum Link Width: The maximum PCIe link width for this port.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr><td>0</td><td>Reserved</td></tr> <tr><td>1</td><td>PCIe x1</td></tr> <tr><td>2</td><td>PCIe x2</td></tr> <tr><td>3</td><td>Reserved</td></tr> <tr><td>4</td><td>PCIe x4</td></tr> <tr><td>5 to 7</td><td>Reserved</td></tr> <tr><td>8</td><td>PCIe x8</td></tr> <tr><td>9 to 11</td><td>Reserved</td></tr> <tr><td>12</td><td>PCIe x12</td></tr> <tr><td>13 to 15</td><td>Reserved</td></tr> <tr><td>16</td><td>PCIe x16</td></tr> <tr><td>17 to 31</td><td>Reserved</td></tr> <tr><td>32</td><td>PCIe x32</td></tr> <tr><td>33 to 255</td><td>Reserved</td></tr> </tbody> </table>	Value	Definition	0	Reserved	1	PCIe x1	2	PCIe x2	3	Reserved	4	PCIe x4	5 to 7	Reserved	8	PCIe x8	9 to 11	Reserved	12	PCIe x12	13 to 15	Reserved	16	PCIe x16	17 to 31	Reserved	32	PCIe x32	33 to 255	Reserved
Value	Definition																															
0	Reserved																															
1	PCIe x1																															
2	PCIe x2																															
3	Reserved																															
4	PCIe x4																															
5 to 7	Reserved																															
8	PCIe x8																															
9 to 11	Reserved																															
12	PCIe x12																															
13 to 15	Reserved																															
16	PCIe x16																															
17 to 31	Reserved																															
32	PCIe x32																															
33 to 255	Reserved																															
04	Impl Spec	<p>RefClk Capability: This field contains a bit vector that specifies the PCIe clocking modes supported by the port.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr><td>7:4</td><td>Reserved</td></tr> <tr><td>3</td><td>Set to '1' for upstream ports that automatically use RefClk if provided and otherwise uses SRIS, otherwise, cleared to '0'. Reserved for downstream ports.</td></tr> <tr><td>2</td><td>Set to '1' if the PCIe port supports Separate RefClk with SSC (SRIS), otherwise cleared to '0'.</td></tr> <tr><td>1</td><td>Set to '1' if the PCIe port supports Separate RefClk with no SSC (SRNS), otherwise cleared to '0'.</td></tr> <tr><td>0</td><td>Set to '1' if the PCIe port supports common RefClk, otherwise cleared to '0'.</td></tr> </tbody> </table>	Bit	Description	7:4	Reserved	3	Set to '1' for upstream ports that automatically use RefClk if provided and otherwise uses SRIS, otherwise, cleared to '0'. Reserved for downstream ports.	2	Set to '1' if the PCIe port supports Separate RefClk with SSC (SRIS), otherwise cleared to '0'.	1	Set to '1' if the PCIe port supports Separate RefClk with no SSC (SRNS), otherwise cleared to '0'.	0	Set to '1' if the PCIe port supports common RefClk, otherwise cleared to '0'.																		
Bit	Description																															
7:4	Reserved																															
3	Set to '1' for upstream ports that automatically use RefClk if provided and otherwise uses SRIS, otherwise, cleared to '0'. Reserved for downstream ports.																															
2	Set to '1' if the PCIe port supports Separate RefClk with SSC (SRIS), otherwise cleared to '0'.																															
1	Set to '1' if the PCIe port supports Separate RefClk with no SSC (SRNS), otherwise cleared to '0'.																															
0	Set to '1' if the PCIe port supports common RefClk, otherwise cleared to '0'.																															
05	Impl Spec	Port Number: This field indicates the PCIe Port Number, as defined by the PCI Express Base Specification, associated with this port.																														
06	Impl Spec	PCIe Pointer: In downstream ports this field contains the child index of the Element Descriptor that has a PCIe port connected to this PCIe port. In upstream ports this field is cleared to 0h.																														
07	Impl Spec	Destination Port: This field contains the index of the Port Descriptor in the child Element Descriptor. If the child Element Descriptor has one PCIe upstream port (i.e., a PCIe Switch Element Descriptor), this field shall be cleared to 0h.																														

9.2.5.7 NVM Subsystem Element Descriptor

The NVM Subsystem Element Descriptor is shown in Figure 167 and is used to describe an NVM Subsystem contained in the NVMe Storage Device.

Figure 167: NVM Subsystem Element Descriptor

Byte Offset	Factory Default	Description
00	07h	Type: This field indicates the type of the Element Descriptor. The NVM Subsystem Element Descriptor Type is 7h.
01	00h	Revision: This field indicates the revision of the Element Descriptor. The NVM Subsystem Element Descriptor Revision is 0h for this specification.

Figure 167: NVM Subsystem Element Descriptor

Byte Offset	Factory Default	Description																		
02	Impl Spec	Length: This field indicates the length of the NVM Subsystem Element Descriptor in bytes.																		
03	3Ah or 3Bh	<p>SMBus/I2C Address Info: If the NVM Subsystem supports an MCTP over SMBus/I2C port, then this field indicates the SMBus/I2C address for MCTP over SMBus/I2C port and whether or not SMBus ARP is supported; otherwise this field has a value of 0h.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:1</td> <td>SMBus/I2C Address: This field contains the 7-bit SMBus/I2C address. Refer to Figure 15 for requirements.</td> </tr> <tr> <td>0</td> <td>ARP Capable: This bit is set to '1' if SMBus ARP is supported, else it is cleared to '0'. Refer to Figure 15 for requirements.</td> </tr> </tbody> </table>	Bit	Description	7:1	SMBus/I2C Address: This field contains the 7-bit SMBus/I2C address. Refer to Figure 15 for requirements.	0	ARP Capable: This bit is set to '1' if SMBus ARP is supported, else it is cleared to '0'. Refer to Figure 15 for requirements.												
Bit	Description																			
7:1	SMBus/I2C Address: This field contains the 7-bit SMBus/I2C address. Refer to Figure 15 for requirements.																			
0	ARP Capable: This bit is set to '1' if SMBus ARP is supported, else it is cleared to '0'. Refer to Figure 15 for requirements.																			
04	Impl Spec	<p>SMBus/I2C Capabilities: If the NVM Subsystem supports an SMBus/I2C port then this field indicates the SMBus/I2C capabilities; otherwise this field has a value of 0h.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7</td> <td>Reset: This bit is set to '1' if all of the SMBus/I2C reset mechanisms are supported as defined by the associated form factor specification. This bit is cleared to '0' if the form factor does not define SMBus Reset or the NVMe Storage Device does not support all of the SMBus/I2C reset mechanisms defined by the specification for the Form Factor in the Host Connector Element Descriptor.</td> </tr> <tr> <td>6:2</td> <td>Reserved</td> </tr> <tr> <td>1:0</td> <td> <p>Maximum Speed: This field is set to the highest supported SMBus/I2C clock speed.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>100 kHz</td> </tr> <tr> <td>1</td> <td>400 kHz</td> </tr> <tr> <td>2</td> <td>1 MHz</td> </tr> <tr> <td>3</td> <td>Reserved</td> </tr> </tbody> </table> </td> </tr> </tbody> </table>	Bit	Description	7	Reset: This bit is set to '1' if all of the SMBus/I2C reset mechanisms are supported as defined by the associated form factor specification. This bit is cleared to '0' if the form factor does not define SMBus Reset or the NVMe Storage Device does not support all of the SMBus/I2C reset mechanisms defined by the specification for the Form Factor in the Host Connector Element Descriptor.	6:2	Reserved	1:0	<p>Maximum Speed: This field is set to the highest supported SMBus/I2C clock speed.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>100 kHz</td> </tr> <tr> <td>1</td> <td>400 kHz</td> </tr> <tr> <td>2</td> <td>1 MHz</td> </tr> <tr> <td>3</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Description	0	100 kHz	1	400 kHz	2	1 MHz	3	Reserved
Bit	Description																			
7	Reset: This bit is set to '1' if all of the SMBus/I2C reset mechanisms are supported as defined by the associated form factor specification. This bit is cleared to '0' if the form factor does not define SMBus Reset or the NVMe Storage Device does not support all of the SMBus/I2C reset mechanisms defined by the specification for the Form Factor in the Host Connector Element Descriptor.																			
6:2	Reserved																			
1:0	<p>Maximum Speed: This field is set to the highest supported SMBus/I2C clock speed.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>100 kHz</td> </tr> <tr> <td>1</td> <td>400 kHz</td> </tr> <tr> <td>2</td> <td>1 MHz</td> </tr> <tr> <td>3</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Description	0	100 kHz	1	400 kHz	2	1 MHz	3	Reserved									
Value	Description																			
0	100 kHz																			
1	400 kHz																			
2	1 MHz																			
3	Reserved																			
05	Impl Spec	NVM Subsystem Port Descriptor Count: This field indicates the number of NVM Subsystem Port Descriptors associated with the NVM Subsystem. The permitted range of values is 1 to 64.																		
Impl Spec	Impl Spec	NVM Subsystem Port Descriptor 0: This field contains the NVM Subsystem Port Descriptor associated with the first NVM Subsystem port.																		
Impl Spec	Impl Spec	NVM Subsystem Port Descriptor 1: This field contains the NVM Subsystem Port Descriptor associated with the second NVM Subsystem port if NVM Subsystem Port Descriptor Count is greater than one, otherwise this field is not present.																		
...																		

Each upstream port is described by an NVM Subsystem Port Descriptor as shown in Figure 168. It describes the PCIe port's supported PCIe link speeds, PCIe max link width, RefClk capabilities, and PCIe Port Identifier. Each NVM Subsystem Port Descriptor should be the child of exactly one parent Element Descriptor.

Figure 168: NVM Subsystem Port Descriptor

Byte Offset	Factory Default	Description
00	00h	Type: This field indicates the type of an NVM Subsystem Port Descriptor. The NVM Subsystem Port Descriptor Type is 0.

Figure 168: NVM Subsystem Port Descriptor

Byte Offset	Factory Default	Description																														
01	Impl Spec	Length: This field indicates the length of the NVM Subsystem Port Descriptor in bytes.																														
02	Impl Spec	<p>PCIe Link Speed: This field indicates a bit vector of link speeds supported by the PCIe port.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:4</td> <td>Reserved</td> </tr> <tr> <td>3</td> <td>Set to '1' if the PCIe link supports 16 GT/s, otherwise cleared to '0'.</td> </tr> <tr> <td>2</td> <td>Set to '1' if the PCIe link supports 8.0 GT/s, otherwise cleared to '0'.</td> </tr> <tr> <td>1</td> <td>Set to '1' if the PCIe link supports 5.0 GT/s, otherwise cleared to '0'.</td> </tr> <tr> <td>0</td> <td>Set to '1' if the PCIe link supports 2.5 GT/s, otherwise cleared to '0'.</td> </tr> </tbody> </table>	Bit	Description	7:4	Reserved	3	Set to '1' if the PCIe link supports 16 GT/s, otherwise cleared to '0'.	2	Set to '1' if the PCIe link supports 8.0 GT/s, otherwise cleared to '0'.	1	Set to '1' if the PCIe link supports 5.0 GT/s, otherwise cleared to '0'.	0	Set to '1' if the PCIe link supports 2.5 GT/s, otherwise cleared to '0'.																		
Bit	Description																															
7:4	Reserved																															
3	Set to '1' if the PCIe link supports 16 GT/s, otherwise cleared to '0'.																															
2	Set to '1' if the PCIe link supports 8.0 GT/s, otherwise cleared to '0'.																															
1	Set to '1' if the PCIe link supports 5.0 GT/s, otherwise cleared to '0'.																															
0	Set to '1' if the PCIe link supports 2.5 GT/s, otherwise cleared to '0'.																															
03	Impl Spec	<p>PCIe Maximum Link Width: The maximum PCIe link width for this NVM Subsystem port.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Reserved</td> </tr> <tr> <td>1</td> <td>PCIe x1</td> </tr> <tr> <td>2</td> <td>PCIe x2</td> </tr> <tr> <td>3</td> <td>Reserved</td> </tr> <tr> <td>4</td> <td>PCIe x4</td> </tr> <tr> <td>5 to 7</td> <td>Reserved</td> </tr> <tr> <td>8</td> <td>PCIe x8</td> </tr> <tr> <td>9 to 11</td> <td>Reserved</td> </tr> <tr> <td>12</td> <td>PCIe x12</td> </tr> <tr> <td>13 to 15</td> <td>Reserved</td> </tr> <tr> <td>16</td> <td>PCIe x16</td> </tr> <tr> <td>17 to 31</td> <td>Reserved</td> </tr> <tr> <td>32</td> <td>PCIe x32</td> </tr> <tr> <td>33 to 255</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Description	0	Reserved	1	PCIe x1	2	PCIe x2	3	Reserved	4	PCIe x4	5 to 7	Reserved	8	PCIe x8	9 to 11	Reserved	12	PCIe x12	13 to 15	Reserved	16	PCIe x16	17 to 31	Reserved	32	PCIe x32	33 to 255	Reserved
Value	Description																															
0	Reserved																															
1	PCIe x1																															
2	PCIe x2																															
3	Reserved																															
4	PCIe x4																															
5 to 7	Reserved																															
8	PCIe x8																															
9 to 11	Reserved																															
12	PCIe x12																															
13 to 15	Reserved																															
16	PCIe x16																															
17 to 31	Reserved																															
32	PCIe x32																															
33 to 255	Reserved																															
04	Impl Spec	<p>RefClk Capability: This field contains a bit vector that specifies the PCIe clocking modes supported by the port.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:4</td> <td>Reserved</td> </tr> <tr> <td>3</td> <td>Set to '1' if the device automatically uses RefClk if provided and otherwise uses SRIS, otherwise cleared to '0'.</td> </tr> <tr> <td>2</td> <td>Set to '1' if the PCIe link supports Separate RefClk with SSC (SRIS), otherwise cleared to '0'.</td> </tr> <tr> <td>1</td> <td>Set to '1' if the PCIe link supports Separate RefClk with no SSC (SRNS), otherwise cleared to '0'.</td> </tr> <tr> <td>0</td> <td>Set to '1' if the PCIe link supports common RefClk, otherwise cleared to '0'.</td> </tr> </tbody> </table>	Bit	Description	7:4	Reserved	3	Set to '1' if the device automatically uses RefClk if provided and otherwise uses SRIS, otherwise cleared to '0'.	2	Set to '1' if the PCIe link supports Separate RefClk with SSC (SRIS), otherwise cleared to '0'.	1	Set to '1' if the PCIe link supports Separate RefClk with no SSC (SRNS), otherwise cleared to '0'.	0	Set to '1' if the PCIe link supports common RefClk, otherwise cleared to '0'.																		
Bit	Description																															
7:4	Reserved																															
3	Set to '1' if the device automatically uses RefClk if provided and otherwise uses SRIS, otherwise cleared to '0'.																															
2	Set to '1' if the PCIe link supports Separate RefClk with SSC (SRIS), otherwise cleared to '0'.																															
1	Set to '1' if the PCIe link supports Separate RefClk with no SSC (SRNS), otherwise cleared to '0'.																															
0	Set to '1' if the PCIe link supports common RefClk, otherwise cleared to '0'.																															
05	Impl Spec	Port Identifier: This field contains the NVMe-MI Port Identifier associated with this port.																														

9.2.5.8 Vendor-Specific Element Descriptors

The Vendor-Specific Element Descriptor is shown Figure 169.

Figure 169: Vendor-Specific Element Descriptors

Byte Offset	Factory Default	Description
00	Impl Spec	Type: This field indicates the type of the Element Descriptor. Vendor-Specific Types have a value in the range of F0h to FFh.

Figure 169: Vendor-Specific Element Descriptors

Byte Offset	Factory Default	Description
01	Impl Spec	Revision: This field indicates the revision of the Element Descriptor. The Vendor-Specific Element Descriptor Revision is determined by the Vendor.
02	Impl Spec	Length: This field indicates the length of the Vendor-Specific Element Descriptor in bytes.
04:03	Impl Spec	PCI Vendor ID: This field indicates PCI-SIG assigned vendor identifier.
Impl Spec	Impl Spec	Vendor Specific: Vendor-specific information.

9.3 Reset

This section describes the reset architecture defined by this specification that are applicable to NVMe Storage Devices and NVMe Enclosures.

9.3.1 NVM Subsystem Reset

An NVM Subsystem Reset is initiated under the conditions outlined in the NVM Express specification (e.g., when main power is applied to the NVM Subsystem). In addition to these conditions, if NVM Subsystem Reset is supported, then it may be initiated by executing a Reset command.

An NVM Subsystem Reset initiated via the out-of-band mechanism may interfere with host software. A Management Controller should coordinate with the host. Coordination between a Management Controller and a host are outside the scope of this specification.

When an NVM Subsystem Reset is initiated, the entire NVM Subsystem is reset. This includes all NVM Subsystem ports (PCIe and SMBus/I2C), Management Endpoints, and Controller Management Interfaces. All state is returned to its default condition.

9.3.2 Controller Level Reset

A Controller Level Reset is initiated under the conditions outlined in the NVM Express specification.

A Controller Level Reset initiated via the out-of-band mechanism may interfere with host software. A Management Controller should coordinate with the host. Coordination between a Management Controller and a host are outside the scope of this specification.

The actions performed on a Controller Level Reset are outlined in the NVM Express specification. A Controller Level Reset has no effect on the Controller Management Interface associated with that Controller, the PCI Express port associated with that Controller, or a Management Endpoint associated with that port. A Controller Level Reset also does not stop the servicing of the Management Interface Command Set or NVM Express Admin Command Set commands that target that Controller (i.e., the NVM Express Admin Command Set is still available even though the NVMe Controller may be disabled or held in reset) or Control Primitives. A Controller Level Reset may affect PCIe Command Set commands executing on that Controller (refer to section 9.1). If a PCIe Command is affected, then the command is completed with status PCIe Inaccessible.

A Controller Level Reset that causes a new firmware image to activate is considered a special event and may impact the operation of the Controller Management Interface associated with one or more Controllers,

servicing of NVMe-MI Messages, or Management Endpoints within an NVM Subsystem. This impact is unspecified and vendor specific. The Management Controller and host should coordinate the activation of a new firmware image. Coordination between a Management Controller and a host are outside the scope of this specification.

9.3.3 Management Endpoint Reset

A Management Endpoint reset is initiated under the conditions outlined in the MCTP Base Specification or the associated MCTP transport binding specifications.

In addition to these conditions, a Management Endpoint associated with a PCI Express port is reset when the PCI Express port is in any of the following states:

- A PCI Express conventional reset; or
- When the PCI Express link is down (i.e., not in the DL_Active state).

When a Management Endpoint Reset is initiated, the state of that Management Endpoint is returned to its default condition and any commands associated with that Management Endpoint being processed are aborted. A reset of a Management Endpoint in an NVM Subsystem has no effect on any other Management Endpoint in the NVM Subsystem or any other NVM Subsystem entity.

9.3.4 SMBus Reset

All SMBus/I2C elements should support the recommendation for SMBus reset when the SMBus/I2C clock is low for longer than $t_{\text{TIMEOUT,MIN}}$.

Some form factors may also specify one or more separate SMBus reset mechanisms. If such mechanisms are supported by an NVM Subsystem, then the NVM Subsystem shall propagate the reset to all SMBus/I2C elements on the NVM Subsystem and translate the reset, if needed, to Expansion Connector form factors.

If the SMBus/I2C element on an NVM Subsystem is in master mode, then an SMBus Reset shall cause it to generate a STOP condition as defined in the SMBus specification within or after the current data byte in the transfer process. The NVM Subsystem shall remain idle on SMBus for the remainder of the SMBus Reset assertion even if other SMBus/I2C elements attempt to address it. An NVM Subsystem shall be ready to receive a START condition as defined in the SMBus Specification within 10 ms after SMBus Reset de-assertion.

An SMBus Reset shall not modify ARP assigned addresses. Management Controllers may send an ARP reset after the SMBus Reset if addresses need to be reinitialized.

An SMBus Reset shall cause SMBus/I2C Management Endpoints to drop the MCTP packet in flight. If the MCTP Command Servicing State is in Transmit, then it shall change to Idle as if transmit completed. An SMBus Reset does not reset other MCTP state information or abort NVMe-MI Message Servicing.

9.4 Security

The Responder may respond with a Response Message Status of Access Denied in an Error Response. While a drive is in an unlocked state, this mechanism shall not be used for the Management Interface Command Set or the NVMe Admin Command Set.

The commands and the times at which such a response is generated is vendor specific. The mechanism used to lock a drive is outside the scope of this specification.

Appendix A – Technical Note: NVM Express Basic Management Command

This appendix describes the NVMe Basic Management Command and is included here for informational purposes only. The NVMe Basic Management Command is not formally a part of the NVMe-MI specification and its features are not tested by the NVMe Compliance program. No further enhancements to the NVMe Basic Management Command are planned, and it is strongly recommended that any consumers of the NVMe Basic Management Command transition to using the standard NVMe-MI protocol.

This specification utilizes Management Component Transport Protocol (MCTP) messages. The NVMe Basic Management Command does not use MCTP. Support for the NVMe Basic Management Command is optional.

This command does not provide any mechanism to modify or configure the NVMe device. Such features use the more capable MCTP protocol rather than this command's simpler SMBus Block Read. The host may reuse existing SMBus or FRU Information Device read subroutines for this read and is not required to switch the SMBus between master and slave modes as in MCTP.

The block read protocol is specified by the SMBus specification which is available online at www.smbus.org. First slave address write and command code bytes are transmitted by the host, then a repeated start and finally a slave address read. The host keeps clocking as the drive then responds in slave mode with the selected data. The command code is used as a starting offset into the data block shown in Figure 149, like an address on a serial EEPROM.

The offset value increments on every byte read and is reset to zero on a stop condition. A read command without a repeated start is permissible and starts transmission from offset zero. Reading more than the block length with an I2C read is also permissible and these reads continue into the first byte in the next block of data. The Packet Error Code (PEC) accumulates all bytes sent or received after the start condition and the current value is inserted whenever a PEC field is reached.

Blocks of data are packed sequentially. The first 2 blocks are defined by the NVMe-MI workgroup. The first block is the dynamic host health data. The second block includes the Vendor ID (VID) and serial number of the drive. Additional blocks of data may be defined by the owner of the VID. Reading past the end of the vendor defined blocks shall return zeros.

The SMBus slave address to read this data structure defaults to D4h. After the Management Controller successfully assigns the MCTP UDID to D4h using ARP, then the Basic Management Command may track and respond to slave reads at future ARP assigned MCTP addresses. This method of changing the Basic Command address is optional and does not persist through power cycles. Interleaved MCTP and block read traffic is permissible and neither command type shall disturb the state of the other commands.

Here are a few example reads from an NVMe drive at 30 °C, no alarms, VID=1234h, serial number is AZ123456 using the format defined in Figure 149. Host transmissions are shown in white blocks and drive responses are shown in grey blocks:

Example 1: SMBus block read of the drive's status (status flags, SMART warnings, temperature):

Start	Addr	W	Cmd Code	Ack	Restart	Addr	R	Length	Ack	Status Flags	Ack	SMART Warnings	Ack	Temp	Ack	Drive Life Used	Ack	Warning Temp	Ack	Power State	Ack	PEC	MACK	Stop
	D4h		00h			D5h		06h		BFh		FFh		1Eh		01h		3Ch		08h		10h		

Example 2: SMBus block read of the drive's static data (VID and serial number):

Start	Addr	W	Cmd Code	Ack	Restart	Addr	R	Length	Ack	VID	Ack	VID	Ack	Serial #	Ack	Serial #	Ack	Serial #	Ack	Serial #	Ack	Serial #	Ack	Serial #	Ack	
	D4h					08h	D5h			16h		12h		34h		'A'		'Z'		'1'		'2'		'3'		'4'
Serial #	Ack	'5'	Serial #	Ack	'6'	Serial #	Ack	' '	Serial #	Ack	' '	Serial #	Ack	' '	Serial #	Ack	' '	Serial #	Ack	' '	Serial #	Ack	' '	Serial #	Ack	' '
35h		36h	20h		20h	20h		20h	20h		20h	20h		20h	20h		20h	20h		20h	20h		20h	20h		20h
Serial #	Ack	' '	Serial #	Ack	' '	PEC	NACK	Stop	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	
20h		20h	DAh																							

Example 3: SMBus send byte to reset Arbitration bit:

Start	Addr	W	Cmd Code	Ack	Stop
	D4h				

Example 4: I2C read of status and vendor content, I2C allows reading across SMBus block boundaries:

Start	Addr	W	Cmd Code	Ack	Restart	Addr	R	Length	Ack	Status	Ack	SMART	Ack	Temp	Ack	Drive Life	Ack	Warning	Ack	Power	Ack	PEC	Ack	Length	Ack	
	D4h					00h	D5h			06h		BFh		FFh		1Eh		01h		3Ch		08h		10h		16h
VID	Ack		VID	Ack		Serial #	Ack	'A'	Serial #	Ack	'Z'	Serial #	Ack	'1'	Serial #	Ack	'2'	Serial #	Ack	'3'	Serial #	Ack	'4'	Serial #	Ack	'5'
12h		34h	41h		5Ah	31h		32h	33h		34h	35h		36h	20h		20h									
Serial #	Ack	' '	Serial #	Ack	' '	Serial #	Ack	' '	Serial #	Ack	' '	Serial #	Ack	' '	Serial #	Ack	' '	Serial #	Ack	' '	Serial #	Ack	' '	Serial #	Ack	' '
20h		20h	20h		20h	20h		20h	20h		20h	20h		20h	20h		20h	20h		20h	20h		20h	20h		20h
Serial #	Ack	' '	Serial #	Ack	' '	PEC	NACK	Stop	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	
20h		20h	20h		20h	20h																				20h

The SMBus Arbitration bit may be used for simple arbitration on systems that have multiple drives on the same SMBus channel without ARP or muxes to separate them. To use this mechanism, the host follows this 3 step process to handle collisions for the same slave address:

1. The host does an SMBus byte write to send byte FFh which clears the SMBus Arbitration bit on all listening Management Endpoints at this slave address;
2. The host does an I2C read starting from offset 0h and continuing at least through the serial number in the second block. The drive transmitting a '0' when other drives sent a '1' wins arbitration and sets the arbitration bit to '1' upon read completion to give other drives priority on the next read;
3. Repeat step 2 until all drives are read, host receiving the Arbitration bit as a '1' indicates loop is done; and
4. Sort the responses by serial number since the order of drive responses varies with health status and temperatures.

Be careful that there are no short reads of similar data between steps 1 and 3. If the read data is exactly the same on multiple drives, then all these drives set the arbitration bit. After that a new send byte FFh is required to restart the process.

The logic levels were intentionally inverted to normally high in the bytes 1 and 2. This is an additional mechanism to assist systems that do not have ARP or muxes. Since '0' bits win arbitration on SMBus, a drive with an alarm condition is prioritized over healthy drives in the above arbitration scheme. A single I2C

read of byte of two bytes starting at offset one from an array of drives detects alarm conditions. Note that only one drive with an alarm may be reliably detected because drives without the same alarm stop transmitting once the bus contention is detected. For this reason the bits are sorted in order of priority. Continuing to read further provides the serial number of the drive that had the alarm.

Figure 170: Subsystem Management Data Structure

Command Code	Offset (byte)	Description
0	00	Length of Status: Indicates number of additional bytes to read before encountering PEC. This value should always be 6 (06h) in implementations of this version of the spec.
	01	<p>Status Flags (SFLGS): This field indicates the status of the NVM Subsystem.</p> <p>SMBus Arbitration – Bit 7 is set to ‘1’ after an SMBus block read is completed all the way to the stop bit without bus contention and cleared to ‘0’ if an SMBus Send Byte FFh is received on this SMBus slave address.</p> <p>Drive Not Ready – Bit 6 is set to ‘1’ when the subsystem is not capable of processing NVMe management commands, and the rest of the transmission may be invalid. If cleared to ‘0’, then the NVM Subsystem is fully powered and ready to respond to management commands. This logic level intentionally identifies and prioritizes powered up and ready drives over their powered off neighbors on the same SMBus channel.</p> <p>Drive Functional – Bit 5 is set to ‘1’ to indicate an NVM Subsystem is functional. If cleared to ‘0’, then there is an unrecoverable failure in the NVM Subsystem and the rest of the transmission may be invalid. Note that this bit may default to ‘0’ after reset and transition to ‘1’ after the NVM Subsystem has completed initialization and this case should not be considered an error.</p> <p>Reset Not Required - Bit 4 is set to ‘1’ to indicate the NVM Subsystem does not need a reset to resume normal operation. If cleared to ‘0’, then the NVM Subsystem has experienced an error that prevents continued normal operation. A Controller Level Reset is required to resume normal operation.</p> <p>Port 0 PCIe Link Active - Bit 3 is set to ‘1’ to indicate the first port’s PCIe link is up (i.e., the Data Link Control and Management State Machine is in the DL_Active state). If cleared to ‘0’, then the PCIe link is down.</p> <p>Port 1 PCIe Link Active - Bit 2 is set to ‘1’ to indicate the second port’s PCIe link is up. If cleared to ‘0’, then the second port’s PCIe link is down or not present.</p> <p>Bits 1:0 shall be set to ‘1’.</p>
	02	<p>SMART Warnings: This field shall contain the Critical Warning field (byte 0) of the NVMe SMART / Health Information log. Each bit in this field shall be inverted from the NVMe definition (i.e., the management interface shall indicate a ‘0’ value while the corresponding bit is ‘1’ in the log page). Refer to the NVM Express specification for bit definitions.</p> <p>If there are multiple Controllers in the NVM Subsystem, the management endpoint shall combine the Critical Warning field from every Controller such that a bit in this field is:</p> <ul style="list-style-type: none"> • Cleared to ‘0’ if any Controller in the subsystem indicates a critical warning for that corresponding bit. • Set to ‘1’ if all Controllers in the NVM Subsystem do not indicate a critical warning for the corresponding bit.
	03	Composite Temperature (CTemp): This field indicates the current temperature in degrees Celsius. If a temperature value is reported, it should be the same temperature as the Composite Temperature from the SMART log of hottest Controller in the NVM Subsystem. The reported temperature range is vendor specific, and shall not exceed the range -60 °C to +127°C. The 8 bit format of the data is shown below.

Figure 170: Subsystem Management Data Structure

Command Code	Offset (byte)	Description																
		<p>This field should not report a stale temperature, which means that it was sampled more than 5 s prior. If recent data is not available, the Management Endpoint should indicate a value of 80h for this field.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h to 7Eh</td> <td>Temperature is measured in degrees Celsius (0 °C to 126 °C)</td> </tr> <tr> <td>7Fh</td> <td>127 °C or higher</td> </tr> <tr> <td>80h</td> <td>No temperature data or temperature data is more the 5 s old.</td> </tr> <tr> <td>81h</td> <td>Temperature sensor failure</td> </tr> <tr> <td>82h to C3h</td> <td>Reserved</td> </tr> <tr> <td>C4h</td> <td>Temperature is -60 °C or lower</td> </tr> <tr> <td>C5h to FFh</td> <td>Temperature measured in degrees Celsius is represented in two's complement (-1 °C to -59 °C)</td> </tr> </tbody> </table>	Value	Description	00h to 7Eh	Temperature is measured in degrees Celsius (0 °C to 126 °C)	7Fh	127 °C or higher	80h	No temperature data or temperature data is more the 5 s old.	81h	Temperature sensor failure	82h to C3h	Reserved	C4h	Temperature is -60 °C or lower	C5h to FFh	Temperature measured in degrees Celsius is represented in two's complement (-1 °C to -59 °C)
Value	Description																	
00h to 7Eh	Temperature is measured in degrees Celsius (0 °C to 126 °C)																	
7Fh	127 °C or higher																	
80h	No temperature data or temperature data is more the 5 s old.																	
81h	Temperature sensor failure																	
82h to C3h	Reserved																	
C4h	Temperature is -60 °C or lower																	
C5h to FFh	Temperature measured in degrees Celsius is represented in two's complement (-1 °C to -59 °C)																	
	04	<p>Percentage Drive Life Used (PDLU): Contains a vendor specific estimate of the percentage of NVM Subsystem NVM life used based on the actual usage and the manufacturer's prediction of NVM life. If an NVM Subsystem has multiple Controllers the highest value is returned. A value of 100 indicates that the estimated endurance of the NVM in the NVM Subsystem has been consumed, but may not indicate an NVM Subsystem failure. The value is allowed to exceed 100. Percentages greater than 254 shall be represented as 255. This value should be updated once per power-on hour and equal the Percentage Used value in the NVMe SMART Health Log Page.</p>																
	05	<p>Current Over Temperature Warning Threshold (Optional): This field indicates the composite temperature over temperature warning threshold in degrees Celsius. This is intended to initially match the temperature reported in the WCTEMP field in the NVMe Identify Controller data structure. If the Over Temperature threshold for Composite Temperature is modified with set features, then the most recent value should be reported. The data format should match the same single byte format as the CTemp field with a range from -60 °C to 127 °C. A value of zero means that this field is not reported or that the threshold is set to 0 °C.</p>																
	06	<p>Current Power (Optional): This field reports the current NVM subsystem power consumption. If both bit mapped fields are zero it means that this field is not reported.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>7</td> <td>NVM Subsystem Idle (NVMSI): This bit is set to '1' when the NVM subsystem is idle and has been idle for at least 5 s. Refer to the NVMe Idle Power (IDLP) definition.</td> </tr> <tr> <td>6:0</td> <td>NVM Subsystem Power (NVMSP): This field reports the ceiling function of the power consumed by the NVM subsystem in Watts. If this value is greater than 127 W, then 127 W is reported. Power reported by the NVM subsystem is determined in the following manner. If NVMSI bit is set to '1', then the value returned is equal to that reported in the Idle Power (IDLP) field in the Power State Descriptor Data Structure for the corresponding NVMe power state. If NVMSI bit is cleared to '0', then the value returned is equal to that reported in the Active Power Workload (APW) field in the Power State Descriptor Structure for the corresponding NVMe power state. The Maximum Power (MP) field value is substituted for IDLP or APW if these are not for reported in the Power State Descriptor Structure for the current NVMe power state.</td> </tr> </tbody> </table>	Bit	Definition	7	NVM Subsystem Idle (NVMSI): This bit is set to '1' when the NVM subsystem is idle and has been idle for at least 5 s. Refer to the NVMe Idle Power (IDLP) definition.	6:0	NVM Subsystem Power (NVMSP): This field reports the ceiling function of the power consumed by the NVM subsystem in Watts. If this value is greater than 127 W, then 127 W is reported. Power reported by the NVM subsystem is determined in the following manner. If NVMSI bit is set to '1', then the value returned is equal to that reported in the Idle Power (IDLP) field in the Power State Descriptor Data Structure for the corresponding NVMe power state. If NVMSI bit is cleared to '0', then the value returned is equal to that reported in the Active Power Workload (APW) field in the Power State Descriptor Structure for the corresponding NVMe power state. The Maximum Power (MP) field value is substituted for IDLP or APW if these are not for reported in the Power State Descriptor Structure for the current NVMe power state.										
Bit	Definition																	
7	NVM Subsystem Idle (NVMSI): This bit is set to '1' when the NVM subsystem is idle and has been idle for at least 5 s. Refer to the NVMe Idle Power (IDLP) definition.																	
6:0	NVM Subsystem Power (NVMSP): This field reports the ceiling function of the power consumed by the NVM subsystem in Watts. If this value is greater than 127 W, then 127 W is reported. Power reported by the NVM subsystem is determined in the following manner. If NVMSI bit is set to '1', then the value returned is equal to that reported in the Idle Power (IDLP) field in the Power State Descriptor Data Structure for the corresponding NVMe power state. If NVMSI bit is cleared to '0', then the value returned is equal to that reported in the Active Power Workload (APW) field in the Power State Descriptor Structure for the corresponding NVMe power state. The Maximum Power (MP) field value is substituted for IDLP or APW if these are not for reported in the Power State Descriptor Structure for the current NVMe power state.																	
	07	<p>PEC: An 8 bit CRC calculated over the slave address, command code, second slave address, and returned data. The algorithm is defined in the SMBus specification.</p>																
8	08	<p>Length of identification: Indicates number of additional bytes to read before encountering PEC. This value should always be 22 (16h) in implementations of this version of the spec.</p>																

Figure 170: Subsystem Management Data Structure

Command Code	Offset (byte)	Description
	10:09	Vendor ID: The 2 byte vendor ID, assigned by the PCI-SIG. Should match VID in the Identify Controller command response. Note the MSB is transmitted first.
	11:30	Serial Number: 20 characters that match the serial number in the NVMe Identify Controller command response. Note the first character is transmitted first.
	31	PEC: An 8 bit CRC calculated over the slave address, command code, second slave address, and returned data. The algorithm is defined in the SMBus specification.
32+	32:255	Vendor Specific – These data structures shall not exceed the maximum read length of 255 specified in the SMBus version 3 specification. Preferably their lengths are not greater than 32 for compatibility with SMBus 2.0.

Appendix B – Example MCTP Messages & Message Integrity Check

Below are artificial MCTP Messages with their corresponding Message Integrity values. Figure 173 shows an example where the message is not an even number of dwords and the MIC spans Dwords 7 and 8. The contents of the messages listed below should be used for reference and do not correspond to valid MCTP messages.

Figure 171: MIC Example 1 – 32 Bytes of 0's

	3	2	1	0
Dword 0	00h	00h	00h	00h
...
Dword 7	00h	00h	00h	00h
Dword 8 (MIC)	8Ah	91h	36h	AAh

Figure 172: MIC Example 2 – 32 Bytes of 1's

	3	2	1	0
Dword 0	FFh	FFh	FFh	FFh
...
Dword 7	FFh	FFh	FFh	FFh
Dword 8 (MIC)	62h	A8h	ABh	43h

Figure 173: MIC Example 3 – 30 Incrementing Bytes from 00h to 1Dh

	3	2	1	0
Dword 0	03h	02h	01h	00h
...
Dword 7 (MIC)	92h	D7h	1Dh	1Ch
Dword 8 (MIC)	<unused>		1Eh	05h

Figure 174: MIC Example 4 – 32 Decrementing Bytes from 1Fh to 00h

	3	2	1	0
Dword 0	1Ch	1Dh	1Eh	1Fh
...
Dword 7	00h	01h	02h	03h
Dword 8 (MIC)	11h	3Fh	DBh	5Ch

Appendix C – Example NVMe-MI Messages over SMBus/I2C

This section contains example NVMe-MI Messages over SMBus/I2c between a Management Controller (e.g., a Baseboard Management Controller) and a Management Endpoint. The Request Messages are sent from the Management Controller to the Management Endpoint and the corresponding Response Messages are sent back from the Management Endpoint to the Management Controller.

The examples assume the following:

- Management Endpoint SMBus/I2C address is 3Ah;
- Management Controller SMBus/I2C address is 20h;
- Management Endpoint MCTP Endpoint ID is 0, examples only use SMBus/I2C address;
- Management Controller MCTP Endpoint ID is 0, examples only use SMBus/I2C address;
- MCTP Transmission Unit Size is 64 bytes;
- NVMe Storage Device Composite Temperature (CTEMP) is 30 °C;
- NVMe Storage Device Controller ID is 1; and
- NVMe Storage Device Serial Number is AZ123456.

The first 4 bytes and the last byte of each packet (shown in orange in the examples below) are defined by the MCTP SMBus/I2C Transport Binding Specification. Bytes 4 to 7 of each packet and the Message Integrity Check (green) are defined by the MCTP Base Specification. The CRC-32C algorithm and the NVMe-MI Message Header (blue) are defined in section 3.1.1.1. Management Controller transmission bytes are shown in white blocks and Management Endpoint transmission bytes are shown in grey blocks. All messages are sent in SMBus master mode and received in slave mode so both sides must reconfigure SMBus between commands and responses.

Example 1: In this example, a Management Controller issues an Identify Command to read the Serial Number (bytes 23:04 of the Identify Controller Data Structure) of an NVMe Storage Device. The NVMe Storage Device's response is shown in the Example 2.

The Request Message is longer than the default 64-byte MCTP Transmission Unit Size and thus spans two MCTP packets. The NVMe-MI Message Type (NMIMT) field specifies that this is an NVMe Admin Command. The NVMe Opcode 06h specifies that this is an Identify Command. This NVMe Opcode and the required values for Dwords 1 to 15 are defined in the NVM Express specification for the Identify Command. The Data Offset of 00000004h skips the first 4 bytes of the Identify Controller Data Structure response. The Data Length of 00000014h limits the response to 20 bytes.

Notice that the blue header is only present in the first packet of a message. The MCTP packet sequence number is incremented from 0 for the first packet to 1 for the second packet. The SMBus PEC is calculated per packet and includes every byte sent. The Message Integrity Check is calculated across both packet payloads but skips all orange and green bytes. The value for SMBus Length field (Byte 2) is the number of bytes following it in the packet, not including the SMBus PEC field per the SMBus Specification.

Start	SSD Addr	0	Protocol=MCTP	Length	BMC Addr	1	MCTP Version	SSD EID	BMC EID	flags,seq,own,tag	Type = NVMe-MI	NVMe Admin	Rsvd	Rsvd
	3Ah		0Fh	45h	21h		01h	00h	00h	8Bh	84h	10h	00h	00h
Opcde=Identify	Flags=Len+ Off	Cntrl Id LSB	Cntrl Id MSB	Dword1 LSB	Dword1	Dword1	Dword1	Dword1 MSB	Dword2 LSB	Dword2	Dword2	Dword2	Dword2 MSB	
06h	03h	01h	00h	00h	00h	00h	00h	00h	00h	00h	00h	00h	00h	
Dword3 LSB	Dword3	Dword3	Dword3 MSB	Dword4 LSB	Dword4	Dword4	Dword4	Dword4 MSB	Dword5 LSB	Dword5	Dword5	Dword5	Dword5 MSB	
00h	00h	00h	00h	00h	00h	00h	00h	00h	00h	00h	00h	00h	00h	
Offset LSB	Offset	Offset	Offset MSB	Length LSB	Length	Length	Length	Length MSB	Dword8 LSB	Dword8	Dword8	Dword8	Dword8 MSB	
04h	00h	00h	00h	14h	00h	00h	00h	00h	00h	00h	00h	00h	00h	
Dword9 LSB	Dword9	Dword9	Dword9 MSB	Dword10 LSB	Dword10	Dword10	Dword10	Dword10 MSB	Dword11 LSB	Dword11	Dword11	Dword11	Dword11 MSB	
00h	00h	00h	00h	01h	00h	00h	00h	00h	00h	00h	00h	00h	00h	
Dword12 LSB	Dword12	Dword12	Dword12 MSB	Dword13 LSB	Dword13	Dword13	Dword13	Dword13 MSB	Dword14 LSB	Dword14	Dword14	Dword14	Dword14 MSB	
00h	00h	00h	00h	00h	00h	00h	00h	00h	00h	00h	00h	00h	00h	
PEC														
B2h														
Start	SSD Addr	0	Protocol=MCTP	Length	BMC Addr	1	MCTP Version	SSD EID	BMC EID	flags,seq,own,tag	Dword15 LSB	Dword15	Dword15	Dword15 MSB
	3Ah		0Fh	0Dh	21h		01h	00h	00h	5Bh	00h	00h	00h	00h
CRC32C LSB	CRC32C	CRC32C	CRC32C MSB	PEC										
4Ah	C3h	2Ch	FAh	EFh										

Example 2: This example shows an NVMe Storage Device’s Response Message to the Identify Command from Example 1. This message is small enough to fit in a single packet so both MCTP SOM and EOM flags are set. The NVMe Express specification defines the format (Dwords 0, 1, and 3) of the Identify Controller Data Structure bytes that are returned.

Note that the SMBus/I2C addresses and MCTP Endpoint IDs in the Response Message are swapped from their order in the Request Message. Also note that the incrementing MCTP packet sequence number for the Management Endpoint is independent from the Management Controller’s MCTP packet sequence number.

Start	BMC Addr	0	Protocol=MCTP	Length	31h	SSD Addr	1	MCTP Version	01h	BMC EID	00h	SSD EID	00h	flags, seq own, tag	C3h	Type= NVMe-MI	84h	NVMe Admin	90h	Rsvd	00h	Rsvd	00h	
Status= Success	Rsvd	00h	Rsvd	00h	Rsvd	00h	Dword0 LSB	00h	Dword0	00h	Dword0	00h	Dword0	00h	Dword0 MSB	00h	Dword1	00h	Dword1	00h	Dword1	00h	Dword1 MSB	00h
Dword3 LSB	Dword3	00h	Dword3	00h	Dword3 MSB	00h	Response Data 'A'	41h	Response Data 'Z'	5Ah	Response Data '1'	31h	Response Data '2'	32h	Response Data '3'	33h	Response Data '4'	34h	Response Data '5'	35h	Response Data '6'	36h		
Response Data ''	Response Data ''	20h	Response Data ''	20h	Response Data ''	20h	Response Data ''	20h	Response Data ''	20h	Response Data ''	20h	Response Data ''	20h	Response Data ''	20h	Response Data ''	20h	Response Data ''	20h	Response Data ''	20h	Response Data ''	20h
CRC32C LSB	CRC32C	7Ah	CRC32C	1Fh	CRC32C	C4h	CRC32C MSB	7Bh	PEC	48h													Stop	

Example 3: In this example, a Management Controller issues an NVMe Subsystem Health Status Poll command and clears the Composite Controller Status. Note that the MCTP packet sequence number is incremented from the last packet the Management Controller sent in Example 1. The NVMe-MI Message Type value of 08h with Opcode 01h makes this an NVMe Subsystem Health Status Poll command. Bit 31 of Dword1 set to '1' clears the Composite Controller Status after preparing the response. Only the first non SR-OV PCI function with any of the trigger able changes is requested.

Start	SSD Addr	0	Protocol=MCTP	Length	19h	BMC Addr	1	MCTP Version	01h	SSD EID	00h	BMC EID	00h	flags, seq own, tag	EBh	Type = NVMe-MI	84h	Cmd = NVMe-MI	08h	Rsvd	00h	Rsvd	00h	
Opcode= SubSys	Rsvd	00h	Rsvd	00h	Rsvd	00h	Dword0 LSB	00h	Dword0	00h	Dword0	00h	Dword0	00h	Dword0 MSB	00h	Dword1 LSB	00h	Dword1	00h	Dword1	00h	Dword1 MSB	80h
CRC32C LSB	CRC32C	AAh	CRC32C	EFh	CRC32C	81h	CRC32C MSB	B4h	PEC	48h													Stop	

Example 4: This example shows an NVMe Storage Device's response to the NVMe Subsystem Health Status Poll command from Example 3. Note that the MCTP packet sequence number is incremented from the last packet the NVMe Storage Device sent in Example 2. Controller ID 0 had a reportable trigger due to its composite temperature change.

Start	BMC Addr	0	Protocol=MCTP	Length	19h	SSD Addr	1	MCTP Version	01h	BMC EID	00h	SSD EID	00h	flags, seq own, tag	D3h	Type= NVMe-MI	84h	Cmd= NVMe-MI	88h	Rsvd	00h	Rsvd	00h
Status= Success	Rsvd	00h	Rsvd	00h	Rsvd	00h	Subsystem Status	38h	SMART Warnings	FFh	Composite Temp.	1Eh	Percent Life Used	05h	Ctrl Stat LSB	01h	Ctrl Stat MSB	00h	Rsvd	00h	Rsvd	00h	
CRC32C LSB	CRC32C	C8h	CRC32C	3Bh	CRC32C	3Bh	CRC32C MSB	57h	PEC	DAh													Stop

Example 5: This example shows a Management Controller issuing a Replay Control Primitive. The Management Controller may choose to replay an entire Response Message if, for example, the Message Integrity Check failed on the initial Response Message. Or the Management Controller may choose to replay a partial message starting at a specified MCTP Transmission Unit Size boundary if, for example, the SMBus PEC failed on an individual packet. The Control Primitive Tag is arbitrarily set to 45h and

remembered by the Management Controller to match response packets to the correct Control Primitives. The MCTP Tag is also modified for this example to show the effect on the replayed packet.

Start	SSD Addr	0	Protocol=MCTP	Length	BMC Addr	1	MCTP Version	SSD EID	BMC EID	flags, seq own, tag	Type = NVMe-MI	Cmd = Primitive	Rsvd	Rsvd
	Ack		Ack	Ack	Ack		Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack
		3Ah	0Fh	11h	21h	01h	00h	00h	00h	FCh	84h	00h	00h	00h
Opcode= Replay	Tag	CPSR Packet#	CPSR Rsvd	CRC32C LSB	CRC32C	CRC32C	CRC32C MSB	PEC						
Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Stop
		04h	45h	00h	00h	CDh	21h	ECh	1Eh	C1h				

Example 6: This example shows an NVMe Storage Device sending an acknowledgement Response Message to the Replay Control Primitive and then sending a second Response Message that replays the previous Response Message from specified offset of zero. Note that the previous command is not reissued because that could return different data after having the Composite Controller Status cleared.

Start	BMC Addr	0	Protocol=MCTP	Length	SSD Addr	1	MCTP Version	BMC EID	SSD EID	flags, seq own, tag	Type= NVMe-MI	Cmd= Primitive	Rsvd	Rsvd
	Ack		Ack	Ack	Ack		Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack
		20h	0Fh	11h	3Bh	01h	00h	00h	00h	E4h	84h	80h	00h	00h
Status= Success	Tag	CPSR Response	CPSR Rsvd	CRC32C LSB	CRC32C	CRC32C	CRC32C MSB	PEC						
Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Ack	Stop
		00h	45h	01h	00h	BDh	86h	02h	83h	94h				