



Solidigm™ D7-P5520/D7-P5620

(Formerly Intel®)
Product Specification

April 2024
Revision 007

Solidigm Confidential

SOLIDIGM™

Solidigm™ D7-P5520/D7-P5620

- Capacities:
 - D7-P5520 - 1.92TB, 3.84TB, 7.68TB, 15.36TB
 - D7-P5620 - 1.6TB, 3.2TB, 6.4TB, 12.8TB
- Performance^{1,2}
 - Seq. R/W 128K QD256: Up to 7100/4200 MB/s³
 - 4K QD1 Random Latency (typ.) R/W: 75/20 μ s
 - 4K QD1 Sequential Latency (typ.) R/W: 10/13 μ s³
 - D7-P5520: Random 4KB⁴ R/W: Up to 1100K/220K IOPS
 - D7-P5620: Random 4KB⁴ R/W: Up to 1100K/390K IOPS
- Solidigm™ 144 layer TLC 3D NAND Flash Memory
- Supported Operating Systems
 - Windows Server 2016, Windows Server 2019, Windows Server 2022
 - RHEL 7.8, 7.9, & 8.3
 - SLES 12 SP5 & 15 SP2
- Certifications
 - VMware IOVP - ESXi 8.0, ESXi 7.0
 - VMware vSAN - ESXi 8.0, ESXi 7.0
 - Windows Hardware Compatibility Program (Inbox Driver) - Server 2022, Server 2019, Server 2016, Windows 10
 - Microsoft SDDC Premium AQ - Server 2022, Server 2019, Server 2016
 - Microsoft Azure Stack - Server 2022, Server 2019, Server 2016
 - UEFI 2.7
- Hardware-based AES-256 XTS Mode Encryption
- Compliance
 - NVM Express 1.4
 - NVM Express Management Interface 1.1
 - PCI Express Base Specification Rev 4.0
 - TCG SIIS Version 1.08, Revision 1.0
 - TCG Storage SSC Opal Specification Version 2.01, Revision 1.00
 - SFF-8639 Module Specification Rev 3.0, Ver 1.05⁵
 - PCI Express Card Electro-Mechanical Spec. Rev 4.0
- Certifications and Declarations
 - UL, CE, C-Tick, BSMI, KCC, Microsoft WHQL, Microsoft WHCK, VCCI
- Power
 - 12V Supply Rail & Optional 3.3V_{aux} Supply Rail for SMBus access over I2C
 - Consumption: Up to 20W
 - Enhanced power-loss data protection
- Endurance Rating⁶
 - D7-P5520: Up to 28PBW
 - D7-P5620: Up to 65.4PBW
- Reliability
 - Uncorrectable Bit Error Rate (UBER): 1 sector per 10¹⁷ bits read
 - Mean Time Between Failure (MTBF): 2 million hours
 - Host Sector Size: 512, 520, 4096, 4104, 4160 Bytes
- Temperature Specification
 - Operating Temperature: 0° C to 70° C (SMART) with specified airflow
 - Non-Operating Temperature⁷: -40° C to 85° C
 - Temperature monitoring (In-band and out of band)
 - SSD Enclosure Touch Temperature can be extended up to 80° C
 - Thermal Throttling at 70° C (SMART)
 - Thermal Shutdown at 80° C (SMART)
- Form Factor and Approximate Drive Weight
 - U.2 15mm: 152g +/- 5g
 - E1.S 9.5mm: 74g +/- 5g
 - E1.S 15mm: 82g +/- 5g
 - E1.L 9.5mm: 196g +/- 5g
- Shock
 - Operating: 1000G (0.5ms)
 - Non-operating: 1000G (0.5ms)
- Vibration
 - Operating: 2.17 GRMS (5-700 Hz)
 - Non-Operating: 3.13 GRMS (5-800 Hz)
- Altitude
 - Operating: -1,000 to 10,000ft.
 - Non-Operating: -1,000 to 40,000ft.
- Product Ecological Compliance
 - RoHS, REACH, EPEAT

Note:

1. Performance values vary by capacity and form factor
2. Performance specifications apply to both compressible and incompressible data
3. MB/s = 1,000,000 bytes/second, Transfer size 128K QD256
4. 4 KB = 4,096 bytes; 8 KB = 8,192 bytes
5. Complete SFF-8639 compliance with latch slot and L-shaped key
6. 1PB = 10^{15} Bytes
7. Please contact your Solidigm representative for details on the non-operating temperature range

Ordering Information

Contact your local Solidigm sales representative for ordering information.

Revision History

Revision	Description	Date
001	<ul style="list-style-type: none"> Initial Release 	January 2022
002	<ul style="list-style-type: none"> Updated Thermal Operating Domain graph Corrected notes within Appendix A Updated Endurance table within Appendix A Updated Appendix B: Power Metrics Updated Vendor Unique Feature Identifier list 	March 2022
003	<ul style="list-style-type: none"> Updated Thermal Operating Domain graph Added EDSFF form factor information 	March 2022
004	<ul style="list-style-type: none"> Updated E1.S 9.5mm and E1.S 15mm Thermal Operating Domain graph Set/Get Max LBA (C1h) section updated Appendix A updated 	April 2022
005	<ul style="list-style-type: none"> Added E1.L 9.5mm Thermal Operating Domain graph Updated U.2 15.36TB Thermal Operating Domain graph Updated byte 5 factory default value within Table 81 - Common Header Updated byte 1 factory default value within Table 82 - Product Info Area Updated Table 87 - Command Response 0x6A (Intel Specific Vendor Unique Commands) Added default value to bit 3 within Byte Offset 11 of Table 84 - NVMe PCIe Port MultiRecord Area 	June 2022
006	<ul style="list-style-type: none"> Updated document template; no update to content 	April 2023
007	<ul style="list-style-type: none"> Updated the Product Specification to reflect the new Solidigm™ branding alongside the former Intel® branding in the following locations: <ul style="list-style-type: none"> Table 21 on page 45 Table 73 on page 76 Table 82 on page 120 PCIe ID on page 146 Updated mechanical drawing and drive weights in Mechanical Information on page 24 Added ESXi 8.0 to VMware IOVP and vSAN certifications Adjusted drive weight measurements Updated Table 19 on page 38 Increased the electrical noise immunity for 100kHz - 20MHz from 50mV to 240mV for U.2 and EDSFF form factors in Electrical Characteristics on page 9 	April 2024

All product plans, roadmaps, specifications, and product descriptions are subject to change without notice.

Nothing herein is intended to create any express or implied warranty, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, or any warranty arising from course of performance, course of dealing, or usage in trade.

The products described in this document may contain design defects or errors known as "errata," which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your Solidigm representative or your distributor to obtain the latest specifications before placing your product order.

For copies of this document, documents that are referenced within, or other Solidigm literature, please contact your Solidigm representative.

All products, computer systems, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.

© Solidigm. "Solidigm" is a trademark of SK hynix NAND Product Solutions Corp (d/b/a Solidigm). "Intel" is a registered trademark of Intel Corporation. Other names and brands may be claimed as the property of others.

Solidigm may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Solidigm reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

Contents

1	Overview.....	6
1.1	References.....	6
1.2	Terms and Acronyms.....	7
2	Product Specifications.....	9
2.1	Capacity.....	9
2.2	Performance.....	9
2.3	Electrical Characteristics.....	9
2.4	Product Features and Availability.....	11
2.5	Endurance.....	12
2.6	Environmental Conditions.....	12
2.7	Product Regulatory Compliance.....	13
2.8	Reliability Specifications.....	15
2.9	Thermal Specifications.....	16
2.9.1	Composite Temperature.....	16
2.9.2	Thermal Throttling.....	16
2.9.3	Airflow Requirements.....	18
2.10	Power Loss Capacitor Test.....	21
2.11	Hot Plug Support.....	21
2.12	Out-of-Band Management (SMBus).....	22
2.12.1	VPD Page Readout over SMBus.....	22
2.13	Variable Sector Size (Extended LBA Format).....	23
3	Mechanical Information.....	24
4	Pin and Signal Descriptions.....	28
4.1	Pin Signal Definitions.....	28
5	Supported Command Sets.....	32
5.1	NVMe Admin Command Set.....	32
5.1.1	Device Self-Test Operations - D7-P5520/D7-P5620.....	32
5.2	NVMe I/O Command Set.....	33
5.3	NVMe Management Interface 1.1 (MI 1.1) Command Set.....	33
5.4	NVMe and Vendor Unique Log Page Support.....	34
5.5	SMART Attributes.....	47
5.6	SET Features Identifiers.....	50
5.7	Vendor Unique Opcodes.....	51
5.7.1	Clear Assert (Opcode C8h).....	51
5.8	Vendor Unique Feature Identifiers (Set/Get).....	52
5.8.1	Set/Get Endurance Management Disable Setting (C0h).....	52
5.8.2	Set/Get Max LBA (C1h).....	52
5.8.3	Set/Get Native Max LBA (C2h).....	53
5.8.4	Extend Temp Setting (C3h).....	53
5.8.5	Set/Get Power Governor Setting (C6h).....	54
5.8.6	Set/Get SMB ASIC Address (C8h).....	54

5.8.7	Set/Get Blink activity for LED (C9h).....	55
5.8.8	Set SMART Execute Self Test (D4h).....	56
5.8.9	Get Reset Timed Workload Counters (D5h).....	56
5.8.10	Set/Get Enable Latency Tracking (E2h).....	56
5.8.11	Set Read/Write IO Latency Threshold (F4h).....	57
5.8.12	Set/Get Feature NAND Location for Vt Histo (F6h).....	57
6	NVMe Driver Support.....	58
7	Other Compliance and Certifications.....	59
	Appendix.....	60
	Appendix A:Performance and Endurance Metrics.....	60
	Appendix B:Power Metrics.....	75
	Appendix C:IDENTIFY Data Structure.....	76
	Appendix D:NVMe Management Interface 1.1 (NVMe MI 1.1) Supported Command Set.....	117
	Appendix E:Vital Product Data Structure (0x53).....	119
	Appendix F:Out-of-Band Command Response Using SMBus (0x6A).....	132
	Appendix G:Out-of-Band Command Response Using SMBus (0x6A Solidigm Specific).....	136
	Appendix H:Persistent Event Log Events.....	139
	Appendix I:SCSI Command Translation.....	145
	Appendix J:PCIe ID.....	146
	Appendix K:LED Behavior - D7-P5520/D7-P5620.....	148
L	LED Behavior (SFF).....	148
M	LED Behavior (EDSFF) - D7-P5520/D7-P5620 ¹	148

List of Figures

Figure 2.1:	Thermal Throttling Behavior D7-P5520/D7-P5620.....	17
Figure 2.2:	Airflow Approach Curve - Solidigm™ D7-P5520/D7-P5620 - U.2 15mm1.....	19
Figure 2.3:	Airflow Approach Curve - Solidigm™ D7-P5520 - E1.S 9.5mm1.....	19
Figure 2.4:	Airflow Approach Curve - Solidigm™ D7-P5520 - E1.S 15mm1.....	20
Figure 2.5:	Airflow Approach Curve - Solidigm™ D7-P5520 - E1.L 9.5mm1.....	21
Figure 3.1:	Solidigm™ D7-P5520/D7-P5620 U.2 15mm Weight and Dimensions.....	24
Figure 3.2:	Solidigm™ D7-P5520 E1.S 9.5mm Weight and Dimensions.....	25
Figure 3.3:	Solidigm™ D7-P5520 E1.S 15mm Weight and Dimensions.....	26
Figure 3.4:	Solidigm™ D7-P5520 E1.L 9.5mm Dimensions.....	26

List of Tables

Table 1:	Standard Information Referenced in this Document.....	6
Table 2:	Glossary of Terms and Acronyms.....	7
Table 3:	Input Supply Rails - U.2 15mm - D7-P5520/D7-P5620.....	9
Table 4:	Input Supply Rails - EDSFF- D7-P5520.....	10
Table 5:	Product Features - D7-P5520/D7-P5620.....	11
Table 6:	Temperature, Shock, Vibration.....	12

Table 7:	Product Regulatory Compliance Specifications.....	13
Table 8:	Reliability Specifications.....	15
Table 9:	Thermal Throttling Settings -- Solidigm™ D7-P5520/D7-P5620 - SFF.....	17
Table 10:	Thermal Throttling Settings -- Solidigm™ D7-P5520/D7-P5620 - EDSFF.....	18
Table 11:	Out-of-Band Readout Address.....	22
Table 12:	Extended LBA Formats.....	23
Table 13:	Pin Definition (SFF-8639 PCIe Connector Pin-Out).....	28
Table 14:	Pin Definition for EDSFF PCIe Connector Pin-Out1.....	30
Table 15:	Drive Self-Test Operation	33
Table 16:	OCF Extended SMART Log Page(Log Identifier C0h).....	34
Table 17:	Read/Write Command Latency Log (Log Identifier C1h/C2h).....	35
Table 18:	Temperature1 Statistics (Log Identifier C5h).....	37
Table 19:	Vendor Unique SMART Log (Log Identifier CAh).....	38
Table 20:	NVMe IO Queue Metrics Log Page (Log Identifier CBh).....	44
Table 21:	Drive Marketing Name Log (Log Identifier DDh).....	45
Table 22:	Get Power Usage Log Page (Log Identifier F2h).....	45
Table 23:	Vt Histo Get Log Page (Log Identifier F6h).....	45
Table 24:	Garbage Control Collection Log (Log Identifier FDh).....	46
Table 25:	Latency Outlier Log Page (Log Identifier FEh).....	46
Table 26:	SMART Attributes (Log Identifier 02h).....	47
Table 27:	Get Log Page - Temperature Sensor Data Structure.....	50
Table 28:	Clear Assert - Command DWORD 10.....	51
Table 29:	C0h - Endurance Management Disable Setting - Command Dword 11.....	52
Table 30:	Set Max LBA Setting - Command Dword 11 and Command Dword 12.....	53
Table 31:	Status Code - Set Max LBA Command Specific Status Values.....	53
Table 32:	C3h - Extend Temp Setting - Command Dword 11.....	53
Table 33:	Status Codes - Extend Temp Setting Command Specific Status Values.....	54
Table 34:	Set/Get Power (Typical) Governor Setting (C6h).....	54
Table 35:	Status Codes - Power Governor Setting Command Specific Status Value.....	54
Table 36:	C8h - Get/Set SMB ASIC Address.....	55
Table 37:	C9h - Set/Get Blink activity for LED.....	55
Table 38:	Status Code - Set/ Get P11 Blink Activity for LED Specific Status Values.....	55
Table 39:	D4h - Set SMART Execute Self Test - Command Dword 11.....	56
Table 40:	D5h - Reset Timed Workload Counters - Command Dword 11.....	56
Table 41:	E2h - Set/Get Enable Latency Tracking.....	57
Table 42:	F4h - Set Read/Write IO Latency Threshold.....	57
Table 43:	F6h - Set/Get NAND Location for Vt Histo - Command Dword 11 = 20 bytes.....	57
Table 44:	Solidigm™ D7-P5520 NVMe Driver Support.....	58
Table 45:	Other Compliance and Certifications.....	59
Table 46:	Solidigm™ D7-P5520/D7-P5620 User Addressable Sectors.....	60
Table 47:	Random Read/Write (IOPS) - D7-P5520 - U.2 15mm.....	60
Table 48:	Random Read/Write (IOPS) - D7-P5520 - EDSFF.....	61
Table 49:	Random Read/Write (IOPS) - D7-P5620.....	62

Table 50:	Random Read/Write IOPS Consistency (%) - D7-P5520.....	63
Table 51:	Random Read/Write IOPS Consistency (%) - D7-P5620.....	63
Table 52:	Sequential Read and Write Bandwidth (MB/s) - D7-P5520 - U.2 15mm.....	64
Table 53:	Sequential Read and Write Bandwidth (MB/s) - D7-P5520 - EDSFF.....	64
Table 54:	Sequential Read and Write Bandwidth (MB/s) - D7-P5620.....	65
Table 55:	Typical Latency - D7-P5520- U.2 15mm.....	66
Table 56:	Typical Latency - D7-P5520- EDSFF.....	66
Table 57:	Typical Latency - D7-P5620- U.2 15mm.....	67
Table 58:	Quality of Service - D7-P5520 - U.2 15mm.....	67
Table 59:	Quality of Service - D7-P5520 - EDSFF.....	68
Table 60:	Quality of Service - D7-P5620.....	69
Table 61:	Endurance - Drive Writes Per Day (DWPD) - U.2 15mm.....	69
Table 62:	Endurance - Drive Writes Per Day (DWPD) - EDSFF.....	70
Table 63:	Drive Shutdown (RTD3 Entry) and Time To Ready (TTR) - U.2 15mm.....	71
Table 64:	Drive Shutdown (RTD3 Entry) and Time To Ready (TTR) - EDSFF.....	71
Table 65:	Format NVM Secure Erase and Sanitize Completion Times - U.2 15mm.....	72
Table 66:	Format NVM Secure Erase and Sanitize Completion Times - EDSFF.....	72
Table 67:	Dataset Management Deallocate (TRIM) Command Completion Time - U.2 15mm.....	73
Table 68:	Dataset Management Deallocate (TRIM) Command Completion Time - EDSFF.....	73
Table 69:	Opal Revert1 Method Completion Time - U.2 15mm.....	73
Table 70:	Opal Revert1 Method Completion Time - EDSFF.....	74
Table 71:	Power Consumption - U.2 15mm.....	75
Table 72:	Power Consumption - EDSFF.....	75
Table 73:	Identify Controller.....	76
Table 74:	Vendor Specific Data Structure.....	100
Table 75:	Power State Descriptor Data Structure.....	101
Table 76:	Identify Namespace Data Structure.....	104
Table 77:	LBA Format Data Structure.....	115
Table 78:	NVMe MI 1.1 Supported Command Set.....	117
Table 79:	Vital Product Data Structure (VPD Elements).....	119
Table 80:	Type/Length Byte Format.....	119
Table 81:	Common Header.....	119
Table 82:	Product Info Area (Offset 8 Bytes).....	120
Table 83:	NVMe MultiRecord Area.....	122
Table 84:	NVMe PCIe Port MultiRecord Area.....	124
Table 85:	Topology MultiRecord Area.....	127
Table 86:	System Management Data Structure (NVMe-MI Commands).....	133
Table 87:	Command Response 0x6A (Solidigm Specific Vendor Unique Commands).....	136
Table 88:	Persistent Events.....	139
Table 89:	SMART / Health Log Snapshot Event Data Format (Event Type 01h).....	139
Table 90:	Firmware Commit Event Data Format (Event Type 02h).....	139
Table 91:	Timestamp Change Event (Event Type 03h).....	140
Table 92:	Power-on or Reset Event (Event Type 04h).....	140

Table 93:	NVM Subsystem Hardware Error Event Format (Event Type 05h).....	140
Table 94:	Change Namespace Event Data Format (Event Type 06h).....	141
Table 95:	Format NVM Start Event Data Format (Event Type 07h).....	142
Table 96:	Format NVM Completion Event Data Format (Event Type 08h).....	142
Table 97:	Sanitize Start Event Data Format (Event Type 09h).....	143
Table 98:	Sanitize Completion Event Data Format (Event Type 0Ah).....	143
Table 99:	PCIe ID - U.2 15mm.....	146
Table 100:	PCIe ID - EDSFF.....	146
Table 101:	LED Behavior (U.2 Form Factor).....	148
Table 102:	LED Behavior (EDSFF Form Factor).....	149

1 Overview

This document describes the specifications and capabilities of the Solidigm™ D7-P5520/D7-P5620.

Solidigm™ D7-P5520/D7-P5620 is a PCIe Gen4 SSD architected with Solidigm's 144 layer 3D NAND technology and high performance controller interface – NVMe (Non-Volatile Memory Express) delivering leading performance, low latency and QoS (Quality of Service).

Matching the performance with world-class reliability and endurance,

- Solidigm™ D7-P5520 is available in 1.92TB, 3.84TB, 7.68TB, & 15.36TB capacities in the U.2 15mm form factor and offers up to 1.0 DWPD.
- Solidigm™ D7-P5520 is available in 1.92TB, 3.84TB, & 7.68TB capacities in the E1.S 9.5mm and E1.S 15mm form factor and offers up to 1.0 DWPD.
- Solidigm™ D7-P5520 is available in 15.36TB capacity in the E1.L 9.5mm form factor and offers up to 1.0 DWPD.
- Solidigm™ D7-P5620 is available in 1.6TB, 3.2TB, 6.4TB, & 12.8TB capacities in the U.2 15mm form factor and offers up to 3.0 DWPD.

With PCIe Gen4 support and NVMe queuing interface, the D7-P5520/D7-P5620 delivers excellent sequential read performance of up to 7.2GB/s and sequential write speeds of up to 4.3GB/s. The D7-P5520/D7-P5620 delivers very high random read performance up to 1100K IOPS and random write IOPS of 341K for 4KB operations. Taking advantage of the direct path from the storage to the CPU by means of NVMe, the D7-P5520/D7-P5620 exhibits low latency as low as 10µs for sequential access to the SSD.

The Solidigm™ D7-P5520/D7-P5620 include these key features:

- Consistently High IOPS and throughput over PCIe 4.0 Interface
- Sustained low latency
- NVMe Sanitize
- Device Self-test
- Out of band management
- Host and Controller Initiated Telemetry
- Extended host sector sizes
- Scatter-Gather List¹
- Persistent Event Logs
- TCG Storage SSC Opal Specification Version 2.01, Revision 1.00 (Availability limited to OPAL SKU)

Note:

1. SGL bitbucket descriptors are not supported.

1.1 References

Table 1: Standard Information Referenced in this Document

Date	Title	Location
June 2019	NVMe Revision 1.4	http://www.nvmexpress.org
Oct 2017	PCIe Base Specification Revision 4.0	http://pcisig.com
Jan 2013	Enterprise SSD Form Factor Version 1.0a	http://www.ssdformfactor.org

Table 1: Standard Information Referenced in this Document

Date	Title	Location
July 2012	Solid State Drive (SSD) Endurance Workloads (JESD219)	http://www.jedec.org/standards-documents/results/jesd219
Sept 2010	Solid State Drive (SSD) Requirements and Endurance Test Method (JESD218)	http://www.jedec.org/standards-documents/docs/jesd218/
Dec 2008	VCCI	http://www.vcci.jp/vcci_e/
June 2009	RoHS	https://www.solidigm.com/products/data-center.html Navigate to product, configure SKU, MDDS (Material Description Datasheet). It can be found under "PCN/MDDS" tab.
1995 1996 1995 1995 1997 1994	International Electrotechnical Commission EN 61000 4-2 (Electrostatic discharge immunity test) 4-3 (Radiated, radio-frequency, electromagnetic field immunity test) 4-4 (Electrical fast transient/burst immunity test) 4-5 (Surge immunity test) 4-6 (Immunity to conducted disturbances, induced by radio-frequency fields) 4-11 (Voltage Variations, voltage dips, short interruptions and voltage variations immunity tests)	http://www.iec.ch/

1.2 Terms and Acronyms

Table 2: Glossary of Terms and Acronyms

Term	Definition
ATA	Advanced Technology Attachment
CRC	Cyclic Redundancy Check
DAS	Device Activity Signal
DMA	Direct Memory Access
ECC	Error Correction Code
EEPROM	Electrically Erasable Programmable Read Only Memory
GB	Gigabyte Note: The total usable capacity of the SSD may be less than the total physical capacity because a small portion of the capacity is used for NAND flash media management and maintenance purposes.
Gb	Gigabit
I/O	Input/ Output

Table 2: Glossary of Terms and Acronyms

Term	Definition
IOPS	Input/ Output Operations Per Second
KB	Kilobyte
LBA	Logical Block Address
MB	Megabyte (1,000,000 bytes)
MI	Management Interface
MTBF	Mean Time Between Failures
NOP	No Operation
NVMe	Non-Volatile Memory Express
PB	Petabyte
PCB	Printed Circuit Board
RDT	Reliability Demonstration Test
RMS	Root Mean Square
SMART	Self-Monitoring, Analysis and Reporting Technology. This is an open standard for developing hard drives and software systems that automatically monitors the health of a drive and reports potential problems.
SRIS	Separate Refclk Independent Spread
SSD	Solid State Drive
TB	Terabyte
TBD	To Be Determined
TLC	Triple Layer Cell
TYP	Typical
UBER	Uncorrectable Bit Error Rate
VPD	Vital Product Data
Vt	Threshold Voltage

2 Product Specifications

2.1 Capacity

Total user addressable sectors (LBAs) per drive capacity information is provided in Performance and Endurance Metrics [on page 60](#).

2.2 Performance

Performance specifications are highly dependent on the product operating within the operating requirements as listed in the Environmental Conditions section. Additional information is provided in Performance and Endurance Metrics [on page 60](#).

2.3 Electrical Characteristics

Table 3: Input Supply Rails - U.2 15mm - D7-P5520/D7-P5620

Electrical Characteristics	12V Host	3.3V Aux ⁷
Voltage Tolerance	+10/-10%	+/-15%
Inrush Current (Typical Peak) ¹	1.5A	-
Maximum Sustained Device Power	25W	-
Maximum Device Power	37.5W	-
Min Off-Time ²	500ms	500ms
Rising Time ⁵	1ms - 100ms (off to nominal)	1ms - 100ms (off to nominal)
Falling Time ⁵	1ms - 100ms (nominal to off)	1ms - 100ms (nominal to off)
Maximum Continuous Current	See Note 3	5mA active
Absolute Maximum	14.4V for 40 μ s	4.6V for 40 μ s
Absolute Minimum	9.6V for 30 μ s	2.7V for 30 μ s
Noise Immunity ⁶	(10Hz - 100kHz) 1000 mV	(10Hz - 100kHz) 300 mV
	(100kHz - 20 MHz) 240 mV	(100kHz - 20 MHz) 240 mV
Maximum Cap Load ⁴	5uF	5uF

Note:

1. Measured during initial power supply application. Typically this will be seen within 2 seconds of initial power up. Inrush specified for 12V, not the 3.3Vaux
2. Off is considered to be time below 10% of target voltage. If this requirement is not met, drive may not be detected (non-

enumeration at both PCIe & host level) by host system. To recover from this scenario, drive might be powered off for indicated time.

3. 12V - see maximum sustained device power
4. Maximum capacitance presented by the module on the power rail at the receptacle
5. Maximum slew rate not exceeding 12kV/s or Hot Plug / Open Circuit.
6. Within voltage tolerance
7. 3.3Vaux is optional, not needed for power up or functionality. 3.3Vaux is needed for accessing VPD page by means of SMBus when 12V host power is off.

Table 4: Input Supply Rails - EDSFF- D7-P5520

Electrical Characteristics	12V Host	3.3V Aux ⁹
Voltage Tolerance	+10/-10%	+/-10%
Inrush Current (Typical Peak) ¹	2A	-
Maximum Sustained Device Power ³	20W(9.5mm)/25W(15mm)	-
Maximum Device Power ⁴	37.5W	-
Min Off-Time ²	500ms	500ms
Rising Time ⁷	1ms - 100ms (off to nominal)	1ms - 100ms (off to nominal)
Falling Time ⁷	1ms - 100ms (nominal to off)	1ms - 100ms (nominal to off)
Maximum Continuous Current	See Note 5	5mA active
Absolute Maximum	14.4V for 40μs	4.6V for 40μs
Absolute Minimum	9.6V for 30μs	2.7V for 30μs
Noise Immunity ⁸	(10Hz - 100kHz) 1000 mV	(10Hz - 100kHz) 300 mV
	(100kHz - 20 MHz) 240 mV	(100kHz - 20 MHz) 240 mV
Maximum Cap Load ⁶	5uF	5uF

Note:

1. Measured during initial power supply application. Typically this will be seen within 2 seconds of initial power up. Inrush specified for 12V, not the 3.3Vaux
2. Off is considered to be time below 10% of target voltage. If this requirement is not met, drive may not be detected (non-enumeration at both PCIe & host level) by host system. To recover from this scenario, drive might be powered off for indicated time.
3. Measured over any 1s window per EDSFF spec. See Appendix B for per SKU/per workload power consumption.
4. Measured over any 100μs window per EDSFF spec. See Appendix B for per SKU/per workload power consumption.
5. 12V - see maximum sustained device power
6. Maximum capacitance presented by the module on the power rail at the receptacle
7. Maximum slew rate not exceeding 12kV/s or Hot Plug / Open Circuit.
8. Within voltage tolerance
9. 3.3Vaux is optional, not needed for power up or functionality. 3.3Vaux is needed for accessing VPD page by means of SMBus when 12V host power is off.

2.4 Product Features and Availability

Table 5: Product Features - D7-P5520/D7-P5620

Features	Availability
Common Clock (RefClk)	Production Release
Sanitize	Production Release
SMART	Production Release
Firmware commit with NVMe Controller Reset (Commit Action 1)	Production Release
Firmware activate without reset (Commit Action 3)	Production Release
U.2 Hot Plug	Production Release
EDSFF Hot Plug	Production Release
Thermal shutdown	Production Release
Composite Temperature	Production Release
Dataset Management (Deallocate)	Production Release
NVMe 1.4	Production Release
NVMe-MI 1.1 (SMBus)	Production Release
Multiple Power Modes	Production Release
Multiple Namespace Support Namespace Attachment Namespace Management Namespace Locking (up to 16)	Production Release
TCG Storage SSC Opal Specification Version 2.01, Revision 1.00	Production Release
TCG Storage SSC Ruby Specification Version 1.00	Production Release
Multiple Firmware Slots	Production Release
Extended LBA Format (Variable Sector Size)	Production Release
Scatter-Gather List ¹	Production Release
Host / Controller Initiated Telemetry (NVMe 1.4 Compliant)	Production Release
Host Data Recovery	Production Release
Persistent Event Logs	Production Release

Note:

1. SGL bitbucket descriptors are not supported.

2.5 Endurance

Additional information is provided in Performance and Endurance Metrics [on page 60](#).

2.6 Environmental Conditions

Power information is provided in Appendix B: Power Metrics.

Table 6: Temperature, Shock, Vibration

Specification		Form Factor			
		U.2 15mm	E1.S 9.5mm	E1.S 15mm	E1.L 9.5mm
Temperature	Operating ¹ Non-operating ²	0° C to 70° C -40° C to 85° C			
Temperature Gradient ³	Operating Non-operating	20° C/hr (Typical) 30° C/hr (Typical)			
Humidity	Operating Non-operating	5 - 90% 5 - 95%			
Shock ⁴	Operating Non-operating	1,000 G (Max) at 0.5 msec 1,000 G (Max) at 0.5 msec			
Vibration ⁵	Operating Non-operating	2.17 GRMS (5 - 700 Hz) Max 3.13 GRMS (5 - 800 Hz) Max			
Altitude ⁶	Operating Non-operating	-1,000 to 10,000 ft -1,000 to 40,000 ft			

Note:

1. Operating temperature (SMART) implies SSD temperature under defined airflow.
2. Non-operating temperature refers to the ambient air temperature.
3. Temperature gradient measured without condensation.
4. Shock specifications assume the SSD is mounted securely with the input vibration applied to the drive-mounting screws. Stimulus may be applied in the X, Y or Z axis.
5. Vibration specifications assume the SSD is mounted securely with the input vibration applied to the drive-mounting screws. Stimulus may be applied in the X, Y or Z axis. Vibration specification is measured using RMS value.
6. Atmospheric pressure only, not radiation.

2.7 Product Regulatory Compliance

Solidigm™ D7-P5520/D7-P5620 aims to meet or exceeds the regulatory or certification requirements in the following table. The following table applies to the production unit only.

Table 7: Product Regulatory Compliance Specifications


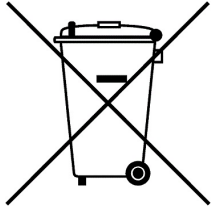





Certification	Description
CE Compliant 	<p>European Economic Area (EEA): Compliance with the essential requirements of EC Council Directives Low Voltage Directive (LVD) 2014/35/EU and EMC Directive 2014/30/EU.</p> <p>Compliance with Directive 2011/65/EU of the European Parliament and of the Council of 8 June 2011 on the restriction of the use of certain hazardous substances in electrical and electronic equipment</p>
EU WEEE  	<p>Compliance with Directive 2012/19/EU of the European Parliament and of the Council of 4 July 2012 on waste electrical and electronic equipment (WEEE)</p>
UL Recognized 	<p>Certified Underwriters Laboratories, Inc. Bi-National Component Recognition; UL 60950-1, 2nd Edition, 2014-10-14 [Information Technology Equipment - Safety - Part 1: General Requirements].</p> <p>CSA C22.2 No. 60950-1-07, 2nd Edition, 2014-010 (Information Technology Equipment - Safety - Part 1: General Requirements)</p> <p>These products have been Complimentary Recognized to UL/CSA 62368-1, 2nd Edition [Audio/video, information and communication technology equipment - Part 1: Safety requirements]</p>  <p>"CAUTION: Hot Surface, Do not touch the SSD surface" "Attention: Surface Chaude. Ne touchez pas la surface"</p>
UKCA 	<p>Great Britain (England, Wales and Scotland): compliance with UK S.I. No.1101, Electrical Equipment (Safety) Regulations 2016, UK S.I. No. 1091, Electromagnetic Compatibility Regulations 2016, and UK S.I. No. 3032, The Restriction of the Use of Certain Hazardous Substances in Electrical and Electronic Equipment Regulations 2012.</p>
Australia / New Zealand: RCM 	<p>Compliance with the Australia/New Zealand Standard(s) AS/NZ CISPR 32:2015 and AS/NZ CISPR 22:2009 +A1:2010, In compliance with the Radiocommunications Act 1992 as part of the ACMA's Electromagnetic Compatibility (EMC) Regulatory Arrangement and RSM Radiocommunications (EMC Standards) Notice 2015.</p>

Table 7: Product Regulatory Compliance Specifications








Certification	Description
<p>Taiwan BSMI</p> 	<p>Compliance to the Taiwan EMC standard CNS 13438: Information technology equipment - Radio disturbance Characteristics - limits and methods of measurement, as amended on June 1, 2006, is harmonized with CISPR 22: 2005.04.</p> <p>Compliance to the Taiwan CNS 15663 [Guidance to reduction of the restricted chemical substances in electrical and electronic equipment (EEE)].</p>
<p>Korea KCC</p> 	<p>Compliance with paragraph 1 of Article 11 of the Electromagnetic Compatibility Control Regulation and meets the Electromagnetic Compatibility (EMC) Framework requirements of the Radio Research Laboratory (RRL) Ministry of Information and Communication Republic of Korea.</p>
<p>Morocco Maghreb</p> 	<p>Compliant with Decree # 2574-14 (EMC) on electromagnetic compatibility</p>
<p>Canada ICES CAN ICES-3 (B)/NMB-3(B)</p>	<p>Compliance with Innovation, Science and Economic Development Canada standard ICES-003.</p>
<p>Japan VCCI</p> 	<p>Voluntary Control Council for Interface to cope with disturbance problems caused by personal computers or facsimile.</p>
<p>China EFUP</p> 	<p>China Environmentally Friendly Use Period (EFUP) symbol. Compliance with GB/T 26572: Requirements on concentration limits for certain restricted substances in electrical and electronic products.</p>

Table 7: Product Regulatory Compliance Specifications

Certification	Description
Ukraine RoHS 	Compliant with TECHNICAL REGULATION on the restriction of the use of certain hazardous substances in electrical and electronic equipment APPROVED by Resolution of the Cabinet of Ministers of Ukraine of 10 March 2017 No. 139.
Low Halogen	Applies only to brominated and chlorinated flame retardants (BFRs/CFRs) and PVC in the final product. Solidigm components as well as purchased components on the finished assembly meet JS-709 requirements, and the PCB/substrate meet IEC 61249-2-21 requirements. The replacement of halogenated flame retardants and/or PVC may not be better for the environment.
Microsoft WHCP 	The windows Hardware Compatibility Program is designed to help Microsoft partners to deliver compatible and reliable systems, software, and hardware products. End users trust the logo as an assurance of compatibility and reliability. This program is intended to help partners develop systems and devices that have been tested to ensure that they meet Microsoft standards for Windows as well as the quality level that ensures a great Windows experience for end users.

2.8 Reliability Specifications

Solidigm™ D7-P5520/D7-P5620 meets or exceeds SSD endurance and data retention requirements as specified in the JESD218-B standard. Reliability specifications are listed in the table below.

Table 8: Reliability Specifications

Parameter	Value
Uncorrectable Bit Error Rate (UBER) <ul style="list-style-type: none"> Uncorrectable bit error will not exceed one sector in the specified number of bits read. In the unlikely event of a non-recoverable read error, the SSD will report it as a read failure to the host; the sector in error is considered corrupt and is not returned to the host. 	$< 1 \text{ sector per } 10^{17} \text{ bits read}$
Mean Time Between Failures (MTBF) <ul style="list-style-type: none"> Mean Time Between Failures is estimated based on Telcordia methodology and demonstrated through Reliability Demonstration Test (RDT). 	2 million hours
Data Retention <ul style="list-style-type: none"> The time period for retaining data in the NAND at maximum rated endurance 	3 months power-off retention once SSD reaches rated write endurance at 40° C

Note: Refer to JESD218-B standard table 1 for UBER, FFR, and other Enterprise SSD requirements.

2.9 Thermal Specifications

2.9.1 Composite Temperature

The Solidigm™ D7-P5520/D7-P5620 follows the composite temperature scheme as part of NVMe Health log (SMART attribute log page identifier 02h, bytes 1 & 2).

The composite temperature calculation incorporates temperature sensors from the ASIC, Board and Media to provide an aggregate value. This value represents all components with temperature monitoring capability, resolving all components to a single critical temperature upon which throttling is based. Thermal throttle engages when the composite temperature reaches 70° C (as reported by SMART). For more information on sensor reading see SMART attributes section.

This thermal management system utilizes temperature sensors now integrated onto the media die and other parts of the drive PCB as a robust method of monitoring of key component temperatures in the varied environment.

In addition, drive will provide out-of-band access to temperature via SMBUS. The SMBus slave address to read SMART data structure is the same address we use for MCTP, and defaults to 0x6Ah.

The NVMe 1.4 spec recommends a Warning Composite Temperature Threshold (WCTEMP) of 70° C.

The SSD Enclosure Temperature Limit on the Solidigm™ D7-P5520/D7-P5620 can be extended from 70° C up to 80° C using the Vendor Unique Extend Temp Setting Set Features commands (see Section 5.8.1).

2.9.2 Thermal Throttling


D7-P5520/D7-P5620 provides performance throttling during high temperature scenarios to mitigate thermal challenges. Thermal throttling works by moving the power ceiling in N number of linear steps, see Figure 1.

- Thermal shutdown of the drive is set to 80° C (as reported by SMART, Composite Temperature)
- Drive will throttle performance when Composite temperature reaches 70° C.

Key definitions for important parameters are:

- T_{start} - throttling starts
- T_{max} - max throttling is applied
- $P_{maxdrop}$ - maximum drop allowed determined by the internal setting
- Hysteresis - 2° C comes into play during the cooling phase to prevent rapid oscillations between the throttle states. The temperature must hit the hysteresis level to return the power ceiling back to previous levels.

Figure 2.1: Thermal Throttling Behavior D7-P5520/D7-P5620



Shown above is the thermal throttling behavior for the D7-P5520/D7-P5620, T_{start} set to 70° C, $T_{max_throttle}$ is set to 73° C and N is set to 4.

The table below shows the thermal throttling settings for D7-P5520/D7-P5620 across all SKUs.

Table 9: Thermal Throttling Settings -- Solidigm™ D7-P5520/D7-P5620 - SFF

Form Factor	Capacity (TB)	Throttle Start - T_{start}	Throttle Max - $T_{max_throttle}$	Throttle Steps
D7-P5520 - U.2 - 15mm	1.92TB	70	73	4
	3.84TB	70	73	4
	7.68TB	70	73	4
	15.36TB	70	73	4

Table 9: Thermal Throttling Settings -- Solidigm™ D7-P5520/D7-P5620 - SFF


Form Factor	Capacity (TB)	Throttle Start - T _{start}	Throttle Max - T _{max_throttle}	Throttle Steps
D7-P5620 - U.2 - 15mm	1.6TB	70	73	4
	3.2TB	70	73	4
	6.4TB	70	73	4
	12.8TB	70	73	4

Table 10: Thermal Throttling Settings -- Solidigm™ D7-P5520/D7-P5620 - EDSFF

Form Factor	Capacity (TB)	Throttle Start - T _{start}	Throttle Max - T _{max_throttle}	Throttle Steps
D7-P5520 - E1.S - 9.5mm	1.92TB	70	73	4
	3.84TB	70	73	4
	7.68TB	70	73	4
D7-P5520 - E1.S - 15mm	1.92TB	70	73	4
	3.84TB	70	73	4
	7.68TB	70	73	4
D7-P5520 - E1.L - 9.5mm	15.36TB	70	73	4


2.9.3 Airflow Requirements

Figure 2.2: Airflow Approach Curve - Solidigm™ D7-P5520/D7-P5620 - U.2 15mm¹




The figure above shows the Approach LFM vs. Ambient Air Temperature for the D7-P5520 U.2 1.92TB, 3.84TB, 7.68TB, and 15.35TB capacities, as well as the D7-P5620 U.2 1.6TB, 3.2TB, and 6.4TB, and 12.8TB capacities at the Standard (70° C) and the Extended (80° C) case temperature, where the drive will begin to throttle based on the SMART temperature.

Figure 2.3: Airflow Approach Curve - Solidigm™ D7-P5520 - E1.S 9.5mm¹




The figure above shows the Approach LFM vs. Ambient Air Temperature for the D7-P5520 E1.S 9.5mm 1.92TB, 3.84TB, and 7.68TB, capacities at the Standard (70° C) and the Extended (80° C) case temperature, where the drive will begin to throttle based on the SMART temperature.

Figure 2.4: Airflow Approach Curve - Solidigm™ D7-P5520 - E1.S 15mm¹



The figure above shows the Approach LFM vs. Ambient Air Temperature for the D7-P5520 E1.S 15mm 1.92TB, 3.84TB, and 7.68TB, capacities at the Standard (70° C) and the Extended (80° C) case temperature, where the drive will begin to throttle based on the SMART temperature.

Figure 2.5: Airflow Approach Curve - Solidigm™ D7-P5520 - E1.L 9.5mm¹



The figure above shows the Approach LFM vs. Ambient Air Temperature for the D7-P5520 E1.L 9.5mm 15.36TB, capacity at the Standard (70° C) and the Extended (80° C) case temperature, where the drive will begin to throttle based on the SMART temperature.

Note:

1. The displayed airflow curves are based on empirical data collected in our Thermal Validation Lab.
2. The above plots have dual-axes. The data can be interpreted in either CFM per SSD or LFM
3. Curves are based on a longitudinal flow condition with the trailing edge at the connector.
4. These curves were measured at Sea Level with a drive pitch of 18mm center-to-center..

2.10 Power Loss Capacitor Test

The Solidigm™ D7-P5520/D7-P5620 supports testing of the power loss capacitor, which can be monitored using SMART attribute critical warning in log page identifier 02h, byte 0, and bit 4.

2.11 Hot Plug Support

The U.2 form factor SSDs support orderly hot insertion and removal and surprise hot insertion by means of presence detect and link-up detect in capable platforms and OSs. On surprise hot removal during IOs, the drive will ensure the integrity of already committed data on the media and commit acknowledged writes to the media.

The EDSFF form factor SSDs support orderly hot insertion and removal and surprise hot insertion by means of presence detect and link-up detect in capable platforms and OSs. On surprise hot removal during IOs, the drive will ensure the integrity of already committed data on the media and commit acknowledged writes to the media.

2.12 Out-of-Band Management (SMBus)

The D7-P5520/D7-P5620 provides out-of-band management by means of an SMBus interface on two addresses; 0x53/0xA6 (7bit/8bit) provides a Vital Product Data (VPD) page; and 0x6A/0xD4 (7bit/8bit) provides an NVMe-MI 1.1 interface.

For the U.2 form factor, data over SMBus (VPD or Drive Status) can only be read with 3.3V_{aux}. It does not require the drive to have 12V applied nor a properly configured PCIe link. Only a limited functionality is available when only 3.3V_{aux} is present. More details on the structure of the VPD page at address 0x53/0xA6 (7bit/8bit) can be found in the [Vital Data Structure](#) appendix.

Table 11: Out-of-Band Readout Address

Description	Address	
	7-bit Address	8-bit Address
VPD Page	0x53	0xA6
NVMe Basic Management Command	0x6A	0xD4
NVMe MI over MCTP	0x1D	0x3A

The D7-P5520/D7-P5620 provides additional drive information via address 0x6A/0xD4 (7bit/8bit) as outlined in the NVMe Basic Management Command (see [here](#)). This interface requires the drive to have power supplied, a valid PCIe link, an initialized NVMe configuration space, and the controller must be enabled. Available commands on this interface are detailed in Appendices [Out-of-Band Command Response Using SMBus \(0x6A\)](#) and [Out-of-Band Command Response Using SMBus \(0x6A Solidigm Specific\)](#), which have details on the Out-of-Band Management data structure.

Note:

- In certain tools the address for the VPD and temperature sensor will appear as 0xA6 and 0xD4 respectively, due to bit shift.
- BMC should not access the SMBUS address within 100msec of device power up, it may experience some glitch on the bus. Some of the drive information (NVMe MI Appendix A per NVMe MI Specification) drive functional status, temperature read out, etc. might have stale value right after power ON and host should wait until drive TTR recommendation. To see the Time to Ready (TTR) refer to the Latency table in Appendix A.

2.12.1 VPD Page Readout over SMBus

D7-P5520/D7-P5620 should support simple Reads to Vital Product Data (VPD). Please refer to [Vital Product Data Structure \(0x53\)](#) for details on VPD Data Structure. VPD contains:

- Basic inventory information such as type and size of Enterprise PCIe SSD, manufacturer, date, revision, and GUID
- Power management data such as power level and power modes
- Vendor specific data

VPD is stored in an SMBus device with a slave address of [0xA6](#) (i.e., slave address bits 7-1 correspond to 1010_011). The VPD page can be read via SMBus through address 0x 53 (7- bit address with bit 8 set to 1) or 0xA6 (8-bit address).

2.13 Variable Sector Size (Extended LBA Format)

D7-P5520/D7-P5620 supports 512, 520, 4096, 4104, 4160 Bytes host sector sizes using the following LBA formats:

- **LBA Format 0 (512B):** LBA Data Size = 512B, Metadata Size = 0
- **LBA Format 1 (520B):** LBA Data Size = 512B, Metadata Size = 8B
- **LBA Format 2 (4096B):** LBA Data Size = 4096B, Metadata Size = 0
- **LBA Format 3 (4104B):** LBA Data Size = 4096B, Metadata Size = 8B
- **LBA Format 4 (4160B):** LBA Data Size = 4096B, Metadata Size = 64B

Table 12: Extended LBA Formats

Sector Size	LBA Format			Metadata Setting	Protection Information (PI)	
	Index	LBA Data Size	Metadata Size	DIF/DIX	Type	PI Location
512B	0	512B	0	-	Disabled	-
520B	1	512B	8B	DIF	Type 2	Last 8B
					Disabled	-
				DIX	Type 2	Last 8B
					Disabled	-
4096B	2	4096B	0	-	Disabled	-
4104B	3	4096B	8B	DIF	Type 2	Last 8B
					Disabled	-
				DIX	Type 2	Last 8B
					Disabled	-
4160B	4	4096B	64B	DIF	Disabled	-
				DIX	Disabled	-

Drives will be shipped with 512B LBA Data Size with Metadata Size = 0. The standard Format NVM Command can be used to format a drive to another supported LBA Format. The drive can be formatted to an LBA format with Metadata Size > 0. At 70% of the drive's usage, drives formatted to an LBA Format with Metadata Size > 0 will have its endurance reduced by up to 30% for the remainder of the drive's life.


Note:

- The first time the drive is formatted to an LBA Format with Metadata Size > 0, the Format NVM Command is expected to take a longer time to complete due to the one-time endurance change.
- This product does not support different LBA Formats in multiple Namespaces. All Namespaces must have identical LBA format, Protection Information Type, Protection Information Location within Metadata, and Metadata Settings. When more than one Namespace exists, an attempt to format or create a Namespace with a different LBA Format will result in command completion status Invalid Format.

3 Mechanical Information


The following figures show the physical package information for the Solidigm™ D7-P5520/D7-P5620 in the U.2 15mm form factor. All dimensions are in millimeters.

Figure 3.1: Solidigm™ D7-P5520/D7-P5620 U.2 15mm Weight and Dimensions




Capacity	Weight (grams)	X - Length	Y - Width	Z - Height
1.6TB/1.92TB	142g +/-5g	100.38 Max	69.85 +/- 0.16	14.75 +/-0.25
3.2TB/3.84TB	152g +/-5g			
6.4TB/7.68TB	152g +/-5g			
12.8TB/15.36TB	152g +/-5g			

Figure 3.2: Solidigm™ D7-P5520 E1.S 9.5mm Weight and Dimensions




Capacity	Weight (grams)	X - Length	Y - Width	Z - Height
1.92TB	68g +/-5g	118.7+/-0.20 Max	33.85 +/-0.10	9.5+/-0.35
3.84TB	74g +/-5g			
7.68TB	74g +/-5g			

Figure 3.3: Solidigm™ D7-P5520 E1.S 15mm Weight and Dimensions



Capacity	Weight (grams)	X - Length	Y - Width	Z - Height
1.92TB	82g +/-5g	112.35+/-0.26 Max	33.85 +/-0.10	15.10 +/- 0.25
3.84TB	82g +/-5g			
7.68TB	82g +/-5g			

Figure 3.4: Solidigm™ D7-P5520 E1.L 9.5mm Dimensions



Capacity	Weight (grams)	X - Length	Y - Width	Z - Height
15.36TB	196g +/-5g	317.3+/-0.55 Max	38.4 +/-0.25	9.5 +/- 0.35

4 Pin and Signal Descriptions

4.1 Pin Signal Definitions

Table 13: Pin Definition (SFF-8639 PCIe Connector Pin-Out)

Pin	Name	Description	Pin	Name	Description
S1	GND	Ground	E7	REFCLK+	Reference clock port 0
S2	S0T+/ PETp0	Not used (SATA/SAS)	E8	REFCLK-	Reference clock port 0
S3	S0T-/PETn0	Not used (SATA/SAS)	E9	GND	Ground
S4	GND	Ground	E10	PETp0 ⁵	Transmitter differential pair, Lane 0
S5	S0R-/PERn0	Not used (SATA/SAS)	E11	PETn0 ⁵	Transmitter differential pair, Lane 0
S6	S0R+/ PERp0	Not used (SATA/SAS)	E12	GND	Ground
S7	GND	Ground	E13	PERn0	Receiver differential pair, Lane 0
E1	REFCLKB+	Reference clock port 1	E14	PERp0	Receiver differential pair, Lane 0
E2	REFCLKB-	Reference clock port 1	E15	GND	Ground
E3	3.3Vaux ⁷	3.3V auxiliary power	E16	RSVD	Reserved
E4	CLKREQ# / PERSTB#	Fundamental reset port 1	S8	GND	Ground
E5	PERST# ⁶	Fundamental reset port 0	S9	S1T+ /PETp1	Not used (SATAe/SAS)
E6	RSVD	Reserved	S10	S1T- /PETn1	Not used (SATAe/SAS)
P1	Wake#	Not used (SATAe/SAS)	S11	GND	Ground
P2	Dev_Reset	Not used (SATAe/SAS)	S12	S1R- /PERn1	Not used (SATAe/SAS)
P3	PWRDIS	Power Disable	S13	S1R+ /PERp1	Not used (SATAe/SAS)
P4	IfDet#	Interface detect (drive type)	S14	GND	Ground
P5	GND	Ground	S15	RSVD	Reserved
P6	GND	Ground	S16	GND	Ground
P7		Not used (SATA/SAS)	S17	PETp1 ⁵	Transmitter differential pair, Lane 1
P8		Not used (SATA/SAS)	S18	PETn1 ⁵	Transmitter differential pair, Lane 1
P9		Not used (SATA/SAS)	S19	GND	Ground

Table 13: Pin Definition (SFF-8639 PCIe Connector Pin-Out)

Pin	Name	Description	Pin	Name	Description
P10	PRSNT_N ³	Presence detect (also used for drive type)	S20	PERn1	Receiver differential pair, Lane 1
P11	Activity/Spinup ¹	Activity signal from the drive	S21	PERp1	Receiver differential pair, Lane 1
P12	GND	Ground	S22	GND	Ground
P13	+12V	12V power	S23	PETp2 ⁵	Transmitter differential pair, Lane 2
P14	+12V	12V power	S24	PETn2 ⁵	Transmitter differential pair, Lane 2
P15	+12V	12V power	S25	GND	Ground
			S26	PERn2	Receiver differential pair, Lane 2
			S27	PERp2	Receiver differential pair, Lane 2
			S28	GND	Ground
			E17	PETp3 ⁵	Transmitter differential pair, Lane 3
			E18	PETn3 ⁵	Transmitter differential pair, Lane 3
			E19	GND	Ground
			E20	PERn3	Receiver differential pair, Lane 3
			E21	PERp3	Receiver differential pair, Lane 3
			E22	GND	Ground
			E23	SMCLK ²	SMBUS clock
			E24	SMDAT ²	SMBUS data
			E25	DualPortEn_N _{8, 9}	Dual port enable

Note:

1. Active Low output indicates read or write activity of the drive. 3.3V tolerant. Power should be supplied from outside of the drive thru an LED and current limit resistor. P11 is used for activity. When idle, logic level is low (LED Solid On). During IO activity and formatting, pin toggles 250msec high, 250msec low signal.
2. SMCLK and SMDAT routes to an MIC which contains Vital Product Data (VPD)
3. PRSNT_N is kept open
4. IfDet_N is grounded
5. Transmit differential pair lanes have 220nF of AC coupling capacitance
6. U.2 15mm form factor only uses PERST0# as a fundamental reset
7. 3.3Vaux is only needed during SMBUS access to the MIC, when 12V is not present, applicable to U.2 15mm. When 12V is present, drive generates 3.3V internally for MIC access

8. DualPortEn_N is pulled up to 1.8V for single port drives. The drive has an internal pull resistor of 10K to 3.3V
9. If the drive is plugged into a Dual Port host, it will enumerate as a x2 device, and Lanes 2 and 3 will be disabled.

Table 14: Pin Definition for EDSFF PCIe Connector Pin-Out¹

Pin	Name	Description	Pin	Name	Description
A28	GND	Ground	B28	GND	Ground
A27	PERp3	Receiver differential pair, Lane 3	B27	PETp3	Transmitter differential pair, Lane 3
A26	PERn3	Receiver differential pair, Lane 3	B26	PETn3	Transmitter differential pair, Lane 3
A25	GND	Ground	B25	GND	Ground
A24	PERp2	Receiver differential pair, Lane 2	B24	PETp2	Transmitter differential pair, Lane 2
A23	PERn2	Receiver differential pair, Lane 2	B23	PETn2	Transmitter differential pair, Lane 2
A22	GND	Ground	B22	GND	Ground
A21	PERp1	Receiver differential pair, Lane 1	B21	PETp1	Transmitter differential pair, Lane 1
A20	PERn1	Receiver differential pair, Lane 1	B20	PETn1	Transmitter differential pair, Lane 1
A19	GND	Ground	B19	GND	Ground
A18	PERp0	Receiver differential pair, Lane 0	B18	PETp0	Transmitter differential pair, Lane 0
A17	PERn0	Receiver differential pair, Lane 0	B17	PETn0	Transmitter differential pair, Lane 0
A16	GND	Ground	B16	GND	Ground
A15	REFCLKp1	Reference clock port 1	B15	REFCLKp0	Reference clock port 0
A14	REFCLKn1	Reference clock port 1	B14	REFCLKn0	Reference clock port 0
A13	GND	Ground	B13	GND	Ground
A12	PRSNT0#	Presence detect	B12	PWRDIS	Power disable
A11	CLKREQ#/ PERST1#	Fundamental reset port 1	B11	+3.3Vaux	3.3V auxiliary power
A10	LED/ACTIVITY	LED from host / Activity from SSD	B10	PERST0#	Fundamental reset port 0
A9	SMBRST#	SMBus reset	B9	DUALPORTEN#	Dual port enable

Table 14: Pin Definition for EDSFF PCIe Connector Pin-Out¹

A8	SMBDAT	SMBus data	B8	RFU	Reserved
A7	SMBCLK	SMBus clock	B7	MFG	
A6	GND	Ground	B6	+12V	12V power
A5	GND	Ground	B5	+12V	12V power
A4	GND	Ground	B4	+12V	12V power
A3	GND	Ground	B3	+12V	12V power
A2	GND	Ground	B2	+12V	12V power
A1	GND	Ground	B1	+12V	12V power

Note:

1. Pin and signal definition follows SFF-TA-1009, Revision 2.0 Enterprise and Datacenter SSD Pin and Signal Specification. For mechanical details of the card edge and mating connector, please refer to SFF-TA-1002 Card Edge Multilane protocol agnostic connector specification
2. SMCLK and SMDAT routes to an internal EEPROM which contains Vital Product Data (VPD)
3. PRSNT_N is pulled to ground by the D7-P5520/D7-P5620
4. Transmit differential pair lanes have 220nF of AC coupling capacitance
5. A10 is used as LED/Activity, an input to the drive to turn on the Amber LED as described by the SFF-TA-1009 specification.
6. D7-P5520/D7-P5620 only uses REFCLK0+ and REFCLK0- as reference clock pair
7. D7-P5520/D7-P5620 only uses PERST0# as a fundamental reset
8. 3.3Vaux is only needed during SMBUS access to the VPDROM

5 Supported Command Sets

The Solidigm™ D7-P5520/D7-P5620 supports all mandatory Admin and I/O commands defined in NVMe 1.4 specification.

5.1 NVMe Admin Command Set

The Solidigm™ D7-P5520/D7-P5620 supports all mandatory NVMe commands, which are:

- Delete I/O Submission Queue
- Delete I/O Completion Queue
- Create I/O Submission Queue
- Create I/O Completion Queue
- Get Log Page
- Identify
- Abort
- SET Features
- GET Features
- Asynchronous Event Request

D7-P5520/D7-P5620 also supports the following optional I/O commands defined in NVMe specification rev 1.4

- Device Self-test
- Firmware Activate
- Firmware Image Download
- Directive Send
- Directive Receive
- Virtualization Management
- Format NVM²
- Namespace Management
- Namespace Attachment
- Sanitize²
- Security Send³
- Security Receive³

Note:

1. See Appendix IDENTIFY Data Structure for details on commands and capabilities.
2. See Appendix Power and Performance Metrics for command completion times.
3. Security Send and Security Receive commands are available only for OPAL enabled SKUs

5.1.1 Device Self-Test Operations - D7-P5520/D7-P5620

The D7-P5520/D7-P5620 supports both the short and extended device self-test operations defined in the NVMe 1.4 Specification. The two test operations perform the same set of tests in the segments shown below in Table 11, with the short device-test operation having a completion time of two minutes or less. The extended device self test will not be aborted by a Controller Level Reset, as noted in NVMe 1.4, Section 8.11.2, while the short device-self test will be

aborted by a Controller Level Reset. The table below lists the segments and failure criteria of the tests performed during the short and extended device self-test operations, based on the informative example outlined in NVMe 1.4, Section 8.11.

Table 15: Drive Self-Test Operation

Solidigm D7-P5520/D7-P5620		
Segment Number - Segment	Test Performed	Failure Criteria
1 - SMART Check	Check SMART or health status for Critical Warning bits set to '1' in SMART / Health Information Log.	Any Critical Warning bit set to '1' fails this segment
2 - Volatile Memory Backup	Validate volatile memory backup solution health (e.g., measure backup power source charge and/or discharge time).	Significant degradation in backup capability
4 - NVM Integrity	Write/read/compare to reserved areas of each NVM. Ensure also that every read/write channel of the controller is exercised.	Data miscompare
7 - Drive Life	End-of-life condition: Assess the drive's suitability for continuing write operations.	The Percentage Used is set to 255 in the SMART / Health Information Log or an analysis of internal key operating parameters indicates that data is at risk if writing continues.
8 - SMART Check	Same as 1 - SMART Check	

5.2 NVMe I/O Command Set

D7-P5520/D7-P5620 supports all of the mandatory NVMe I/O command set defined in NVMe 1.4 specification, which includes:

- Flush
- Write
- Read

Additionally, the following optional commands are supported:

- Write Uncorrectable
- Write Zeros
- Dataset Management (Deallocate only)

Note: See Appendix [SCSI Command Translation](#) for details on SCSI supported commands and capabilities.

5.3 NVMe Management Interface 1.1 (MI 1.1) Command Set

D7-P5520/D7-P5620 devices supports NVMe-MI 1.1. Please refer to [Appendix D](#) for complete set of supported command sets.

Status Flags and Smart Warnings readout: Status flags and Smart Warnings like Temperature can be read out of band using the NVMe-MI basic from the PIC using the address 0x6A/0xD5 (7bit/8bit) as shown in Appendix [Out-of-Band Command Response Using SMBus \(0x6A\)](#).

5.4 NVMe and Vendor Unique Log Page Support

D7-P5520/D7-P5620 supports the following mandatory log pages defined in NVMe 1.4 specification:

- Supported Log Pages (Log Identifier 00h)
- Error Information (Log Identifier 01h)
- SMART/ Health Information (Log Identifier 02h)
- Firmware Slot Information (Log Identifier 03h)
- Command Effects Log (Log Identifier 05h)
- Device Self-Test (Log Identifier 06h)
- Telemetry Host-Initiated (Log Identifier 07h) ¹
- Telemetry Controller-Initiated (Log Identifier 08h)
- Persistent Event Log (Log Identifier 0Dh)
- Sanitize (Log Identifier 81h) ²

Additionally, D7-P5520/D7-P5620 will support the following Vendor Unique (C0h-FFh) log pages:

- OCP Extended SMART Log Page (Log Identifier C0h)
- Read Command Latency Statistics Log Page (Log Identifier C1h)
- Write Command Latency Statistics Log Page (Log Identifier C2h)
- Temperature Statistics (Log Identifier C5h)
- Vendor Unique SMART Log (Log Identifier CAh)
- Vendor Unique NVMe IO Queue Metrics Log Page (Log Identifier CBh)
- Marketing Description Log (Log Identifier DDh)
- Get Power Usage Log Page (Log Identifier F2h)
- Garbage Control Collection Log (Log Identifier FDh)
- Latency Outlier Log (Log Identifier FEh)

Note:

1. Please contact your Solidigm representative for the Telemetry Customer Observable Data Map
2. See version 1.4 of the NVMe specification for log page content.

Table 16: OCP Extended SMART Log Page(Log Identifier C0h)

Solidigm™ D7-P5520/D7-P5620		
Byte	# of Bytes	Log Page Content
0-15	16	Physical Media Units Written
16-31	16	Physical Media Units Read
32-39	8	Bad User NAND blocks
40-47	8	Bad System NAND blocks
48-55	8	XOR Recovery Count

Table 16: OCP Extended SMART Log Page(Log Identifier C0h)

Solidigm™ D7-P5520/D7-P5620		
Byte	# of Bytes	Log Page Content
56-63	8	Uncorrectable Read Error Count
64-71	8	Soft ECC error count
72-79	8	End to End Correction Counts
80	1	System data % Used
81-87	7	Refresh Counts
88-95	8	User data erase counts
96-97	2	Thermal throttling status & count
104-111	8	PCIe Correctable Error Count
112-115	4	Incomplete Shutdowns
116-119	4	Reserved
120	1	% Free blocks
121-127	7	Reserved
128-129	2	Capacitor Health
131-135	5	Reserved
136-143	8	Unaligned I/O
144-151	9	Security Version Number
152-159	8	NUSE
160-175	16	PLP Start Count
176-191	16	Endurance Estimate
208-493	285	Reserved
494-495	2	Log Page Revision
496-511	16	Log Page GUID

Table 17: Read/Write Command Latency Log (Log Identifier C1h/C2h)

Solidigm™ D7-P5520/D7-P5620		
Byte	# of Bytes	Log Page Content
0-1	2	Major Version
2-3	2	Minor Version

Table 17: Read/Write Command Latency Log (Log Identifier C1h/C2h)

Solidigm™ D7-P5520/D7-P5620		
Byte	# of Bytes	Log Page Content
4-259	256	1st group of buckets: range 0-63μs, step 1μs, each bucket size is 4 bytes, total 64 buckets
260-515	256	2nd group of buckets: range 63-127μs, step 1μs, each bucket size is 4 bytes, total 64 buckets
516-771	256	3rd group of buckets: range 127-255μs, step 2μs, each bucket size is 4 bytes, total 64 buckets
772-1027	256	4th group of buckets: range 255-510μs, step 4μs, each bucket size is 4 bytes, total 64 buckets
1028-1283	256	5th group of buckets: range 510μs-1.02ms, step 8μs, each bucket size is 4 bytes, total 64 buckets
1284-1539	256	6th group of buckets: range 1.02-2.04ms, step 16μs, each bucket size is 4 bytes, total 64 buckets
1540-1795	256	7th group of buckets: range 2.04-4.08ms, step 32μs, each bucket size is 4 bytes, total 64 buckets
1796-2051	256	8th group of buckets: range 4.08-8.16ms, step 64μs, each bucket size is 4 bytes, total 64 buckets
2052-2307	256	9th group of buckets: range 8.16-16.32ms, step 128μs, each bucket size is 4 bytes, total 64 buckets
2308-2563	256	10th group of buckets: range 16.32-32.64ms, step 256μs, each bucket size is 4 bytes, total 64 buckets
2564-2819	256	11th group of buckets: range 32.64-65.28ms, step 512μs, each bucket size is 4 bytes, total 64 buckets
2820-3075	256	12th group of buckets: range 65.28-130.56ms, step 1.024ms, each bucket size is 4 bytes, total 64 buckets
3076-3331	256	13th group of buckets: range 130.56-261.12ms, step 2.048ms, each bucket size is 4 bytes, total 64 buckets
3331-3587	256	14th group of buckets: range 261.12-522.24ms, step 4.096ms, each bucket size is 4 bytes, total 64 buckets
3588-3843	256	15th group of buckets: range 522.24ms-1.04s, step 8.192ms, each bucket size is 4 bytes, total 64 buckets
3844-4099	256	16th group of buckets: range 1.04-2.09s, step 16.384ms, each bucket size is 4 bytes, total 64 buckets
4100-4355	256	17th group of buckets: range 2.09-4.18s, step 32.768ms, each bucket size is 4 bytes, total 64 buckets

Table 17: Read/Write Command Latency Log (Log Identifier C1h/C2h)

Solidigm™ D7-P5520/D7-P5620		
Byte	# of Bytes	Log Page Content
4356-4611	256	18th group of buckets: range 4.18-8.36s, step 65.536ms, each bucket size is 4 bytes, total 64 buckets
4612-4867	256	19th group of buckets: range 8.36s-, step 131.072ms, each bucket size is 4 bytes, total 64 buckets
4868-4875	8	Average latency statistics

These log pages will show zero contents until latency tracker is enabled using the Set Features command E2h (Set/Get Enable Latency Tracking). Enabling latency tracker adds a performance penalty and must be disabled upon completing the debug.

Table 18: Temperature¹ Statistics (Log Identifier C5h)

Solidigm™ D7-P5520/D7-P5620		
Byte	# of Bytes	Log Page Content
0-7	8	Current Internal Temperature in Celsius
8-15	8	SSD Overtemp Shutdown flag for last power on. Clears after subsequent reset.
16-23	8	SSD Overtemp Shutdown flag for lifetime. Sticky/persistent once set.
24-31	8	Highest (Lifetime) Composite Temperature in Celsius
32-39	8	Lowest (Lifetime) Composite Temperature in Celsius
40-79	40	Reserved
80-87	8	Max Warning Normalized Threshold in Celsius (TX)
88-95	8	Reserved
96-103	8	Specified Minimum Operating Temp in Celsius (TM)
104-111	8	Estimated Offset in Celsius (TC)
111-511	400	Reserved

Note:

1. All temperature values indicate internal composite temperature values. The log page will read 00h for reserved bytes.

Table 19: Vendor Unique SMART Log (Log Identifier CAh)

Byte Offset	Attribute	Description	Updated	Saved to NAND
00h	AB (Program Fail Count)	Raw value: shows total count of program fails. Normalized value: beginning at 100, shows the percent remaining of allowable program fails.	Upon Event	Upon Event
01h	Reserved			
03h	Normalized Value			
04h	Reserved			
05h	Current Raw Value			
0Ch	AC (Erase Fail Count)	Raw value: shows total count of erase fails. Normalized value: beginning at 100, shows the percent remaining of allowable erase fails.	Upon Event	Upon Event
0Dh	Reserved			
0Fh	Normalized Value			
10h	Reserved			
11h	Current Raw Value			
18h	AD (Wear Leveling Count)	Raw value: Min, max and average values of NAND erase cycles for all blocks. Bytes 1-0: Min. erase cycle Bytes 3-2: Max. erase cycle Bytes 5-4: Avg. erase cycles Normalized value: decrements from 100 to 0.	Upon Event	Upon Event
19h	Reserved			
1Bh	Normalized Value			
1Ch	Reserved			
1Dh	Current Raw Value			
24h	B8 (End to End Error Detection Count)	Raw value: reports number of End-to-End detected and corrected errors by hardware. Normalized value: always 100.	Upon Event	Upon Event
25h	Reserved			
27h	Normalized Value			
28h	Reserved			
29h	Current Raw Value			

Table 19: Vendor Unique SMART Log (Log Identifier CAh)

Byte Offset	Attribute	Description	Updated	Saved to NAND
30h	C7 (CRC Error Count)	Raw value: total number of PCIe Interface CRC errors encountered, as specified in PCIe Link Performance Counter Parameter for "Bad TLP". Count is preserved across power cycles. Normalized value: always 100.	Upon Event	Upon Event
31h	Reserved			
33h	Normalized Value			
34h	Reserved			
35h	Current Raw Value			
3Ch	E2 (Timed Workload, Media Wear)	Raw value: measures the wear seen by the SSD (since reset of the workload timer, attribute E4h), as a percentage of the maximum rated cycles. Divide the raw value by 1024 to derive the percentage with 3 decimal points. Example: if the raw value is 4450, the percentage is $4450/1024 = 4.345\%$. Normalized value: always 100.	After 60 Minutes operation	Every 60 seconds and upon shut-down notification
3Dh	Reserved			
3Fh	Normalized Value			
40h	Reserved			
41h	Current Raw Value			
48h	E3 (Timed Workload, Host Reads %)	Raw value: shows the percentage of I/O operations that are read operations (since reset of the workload timer, attribute E4h). Reported as integer percentage from 0 to 100. Normalized value: always 100.	After 60 Seconds operation	Every 60 seconds and upon shut-down notification
49h	Reserved			
4Bh	Normalized Value			
4Ch	Reserved			
4Dh	Current Raw Value			

Table 19: Vendor Unique SMART Log (Log Identifier CAh)

Byte Offset	Attribute	Description	Updated	Saved to NAND
54h	E4 (Timed Workload, Timer)	Raw value: measures the elapsed time (number of minutes since starting this workload timer). Example: if the raw value is 500, the timer has been running for 500 minutes. Normalized value: always 100.	After 60 Seconds operation	Every 60 seconds and upon shut-down notification
55h	Reserved			
57h	Normalized Value			
58h	Reserved			
59h	Current Raw Value			
60h	EA (Thermal Throttle Status)	Raw value: reports Percent Throttle Status and Count of events Byte 0: Throttle status reported as integer percentage. Bytes 1-4: Throttling event count. Number of times thermal throttle has activated. Preserved over power cycles. Byte 5: Reserved. Normalized value: always 100.	When Requested	Upon Event
61h	Reserved			
63h	Normalized Value			
64h	Reserved			
65h	Current Raw Value			
6Ch	F0 (Retry Buffer Overflow Counter)	Raw Value: Counter to indicate the number of times Retry Buffer has overflowed. Normalized Value: always 100.	Upon Event	Upon Event
6Dh	Reserved			
6Fh	Normalized Value			
70h	Reserved			
71h	Current Raw Value			
78h	F3 (PLL Lock Loss Count)	Raw Value: Counter to indicate the number of times PCIe Refclock PLL has unlocked. Normalized Value: always 100.	Upon Event	Upon Event
79h	Reserved			
7Bh	Normalized Value			
7Ch	Reserved			
7Dh	Current Raw Value			

Table 19: Vendor Unique SMART Log (Log Identifier CAh)

Byte Offset	Attribute	Description	Updated	Saved to NAND
84h	F4 (NAND Bytes Written)	Current Value: NAND sectors written divided by 65536 (1 count = 32 MiB) Normalized value: always 100	When Requested, but no more than 60s +/- 5s	When Requested, but no more than 60s +/- 5s
85h	Reserved			
87h	Normalized Value			
88h	Reserved			
89h	Current Value			
90h	F5 (Host Bytes Written)	Current Value: Host sectors written divided by 65536 (1 count = 32 MiB) Normalized value: always 100	When Requested, but no more than 60s +/- 5s	When Requested, but no more than 60s +/- 5s
91h	Reserved			
93h	Normalized Value			
94h	Reserved			
95h	Current Value			
9Ch	F6 (System Area Life Remaining)	Current value is normalized, representing the amount of system area writes that have been utilized. A value of 0x64 / 100d reflects that the user's allocation of system area writes has been exhausted. Normalized value: always 100	When Requested	When Requested
9Dh	Reserved			
9Fh	Normalized Value			
A0h	Reserved			
A1h	Current Value			
DEh	F8 (NAND Bytes Read)	Raw Value: Host sectors written divided by 65536 (1 count = 32MiB) Normalized value: always 100	Upon Event	Upon Event
DFh	Reserved			
E0h	Normalized Value			
E1h	Reserved			
E2h	Current Value			
E3h	Reserved			

Table 19: Vendor Unique SMART Log (Log Identifier CAh)

Byte Offset	Attribute	Description	Updated	Saved to NAND
E4h	F9 (Available Firmware Downgrades)	Raw Value: Number of FW downgrades available for the user via the standard/spec admin Normalized value: always 100	Upon Event	Upon Event
E5h	Reserved			
E7h	Normalized Value			
E8h	Reserved			
E9h	Current Raw Value			
EFh	Reserved			
F0h	FA (Read Host Collision Count)	Raw Value: Collision count of host read to trim Normalized value: always 100	Upon Event	Upon Event
F1h	Reserved			
F3h	Normalized Value			
F4h	Reserved			
F5h	Current Raw Value			
FBh	Reserved			
FCh	FB (Write Host Collision Count)	Raw Value: Collision count of host write to trim Normalized value: always 100	Upon Event	Upon Event
FDh	Reserved			
FFh	Normalized Value			
100h	Reserved			
101h	Current Raw Value			
107h	Reserved			

Table 19: Vendor Unique SMART Log (Log Identifier CAh)

Byte Offset	Attribute	Description	Updated	Saved to NAND
108h	FC (XOR Pass Count)	Raw Value: Shows total XOR Pass Count when requested Normalized value: always 100	When Requested	Upon Event
109h	Reserved			
10Bh	Normalized Value			
10Ch	Reserved			
10Dh	Current Raw Value			
113h	Reserved			
114h	FD (XOR Fail Count)	Raw Value: Shows total XOR Fail Count when requested Normalized value: always 100	When Requested	Upon Event
115h	Reserved			
117h	Normalized Value			
118h	Reserved			
119h	Current Raw Value			
11Fh	Reserved			
120h	FE (XOR Invoked Count)	Raw Value: Shows total XOR Invoked Count when requested Normalized value: always 100	When Requested	Upon Event
121h	Reserved			
123h	Normalized Value			
124h	Reserved			
125h	Current Raw Value			
12Bh	Reserved			

Table 19: Vendor Unique SMART Log (Log Identifier CAh)

Byte Offset	Attribute	Description	Updated	Saved to NAND
12Ch	E5 (In-flight Read IO Commands)	Raw Value: Shows total in-flight host read IO commands counter when requested Normalized value: always 100	When Requested	Upon Event
12Dh	Reserved			
12Fh	Normalized Value			
130h	Reserved			
131h	Current Raw Value			
138h	E6 (In-flight Write IO Commands)	Raw Value: Shows total in-flight host write IO commands counter when requested Normalized value: always 100	When Requested	Upon Event
139h	Reserved			
13Bh	Normalized Value			
13Ch	Reserved			
13Dh	Current Raw Value			

Note: It is recommended for the host to wait at least 90 seconds after power ON to retrieve the most current counter values from SMART Log Attributes (Log CAh)

Table 20: NVMe IO Queue Metrics Log Page (Log Identifier CBh)

Offset	# of Bytes	Log Page Content
0-1	2	Log Version (current is 1)
2-3	2	IOSQ Count
4-5	2	IOCQ Count
6-389	384	IOSQ Structs (12B each @ 32 queues)
390-709	320	IOCQ Structs (10B each @ 32 queues)
710-1023	314	Reserved padding to reach 1024B
IOSQ Structure Definition		
0-1	2	IOSQ ID
2-3	2	Associated IOCQ ID

Table 20: NVMe IO Queue Metrics Log Page (Log Identifier CBh)

Offset	# of Bytes	Log Page Content
4-5	2	Head Pointer
6-7	2	Tail Pointer
8-9	2	Outstanding Commands
10-11	2	Queue Size (Max Depth)
IOCQ Structure Definition		
0-1	2	IOCQ ID
2-3	2	Head Pointer

Table 21: Drive Marketing Name Log (Log Identifier DDh)

Solidigm™ D7-P5520 & D7-P5620			Intel® SSD D7-P5520 & D7-P5620		
Byte	# of Bytes	Log Page Content	Byte	# of Bytes	Log Page Content
0	12	Solidigm(TM)	0	8	Intel(R)
12	1	Space	8	1	Space
13	2	D7	9	3	SSD
15	1	Space	12	1	Space
16	5	Product Name	13	2	D7
21	3	Space (three)	15	1	Space
24	6	Series	16	5	Product Name
30-511	482	Reserved (0x0)	21	3	Space (three)
			24	6	Series
			30-511	482	Reserved (0x0)

Table 22: Get Power Usage Log Page (Log Identifier F2h)

Byte	# of Bytes	Log Page Content
0-3	4	VIN1 in uW
4-7	4	VIN2 in uW

Table 23: Vt Histo Get Log Page (Log Identifier F6h)

Byte	# of Bytes	Log Page Content
0	2	Version Major

Table 23: Vt Histo Get Log Page (Log Identifier F6h)

Byte	# of Bytes	Log Page Content
2	2	Version Minor
4	4	Die
8	4	Block
12	4	Page
16	1	HistoCompleteFlag
17	1	Lower Page Uncorrectable
18	1	Upper Page Uncorrectable
19	1	Extra Page Correctable
20	1	Top Page Correctable
21	9600	Histo Data
9621	1	Read
9622	1	Pad

Table 24: Garbage Control Collection Log (Log Identifier FDh)

Byte	# of Bytes	Log Page Content		
0-1	2	Log Page Major Version		
2-3	2	Log Page Minor Version		
4-1203	1200	GC Log Timer Type	0-3	This part contains all of the aggressive GC log entries. Every entry contains a 12-byte log entry information. The maximum log entry index is 100.
		GC Log Timestamp	4-11	
1204-4095	2892	Reserved		

Table 25: Latency Outlier Log Page (Log Identifier FEh)

Byte	# of Bytes	Log Page Content
1-0	2	majorVersion 1
3-2	2	minorVersion 0
4-7	4	Reserved
15-8	8	logIdx - total log numbers can be read

Table 25: Latency Outlier Log Page (Log Identifier FEh)

Byte	# of Bytes	Log Page Content	
23-16	8	Timestamp of ticks (logs generated to log dump)	This is the entire log entry that can be duplicated up to 50 entries.
27-24	4	Cmd type - 0 for read or 1 for write	
31-28	4	Latency in μ s	
39-32	8	LBA of the outlier location	

5.5 SMART Attributes

The following tables list the SMART attributes supported by the D7-P5520/D7-P5620 in accordance with NVMe specification revision 1.4.

Table 26: SMART Attributes (Log Identifier 02h)

Byte	# of Bytes	Attribute	Description
0	1	Critical Warning	These bits if set, flag various warning sources. Bit 0: Available Spare is below Threshold Bit 1: Temperature has exceeded Threshold Bit 2: Reliability is degraded due to excessive media or internal errors Bit 3: Media is placed in Read-Only Mode Bit 4: Volatile Memory Backup System has failed (e.g., enhanced power loss capacitor test failure) Bits 5-7: Reserved Any of the critical warning can be tied to asynchronous event notification.
1	2	Temperature	Device SMART temperature in Kelvin.
3	1	Available Spare	Starts from 100 and decrements.
4	1	Available Spare Threshold	Contains a normalized percentage (0 to 100%) of the remaining spare capacity available Threshold is set to 10%.
5	1	Percentage Used Estimate (Value allowed to exceed 100%)	A value of 100 indicates that the estimated endurance of the device has been consumed, but may not indicate a device failure. The value is allowed to exceed 100. Percentages greater than 254 shall be represented as 255. This value shall be updated once per power-on hour (when the controller is not in a sleep state). At drive life of 105%, write performance to the drive is throttled. Drive enters into read only mode.

Table 26: SMART Attributes (Log Identifier 02h)

Byte	# of Bytes	Attribute	Description
32	16	Data Units Read (in LBAs)	Contains the number of 512 byte data units the host has read from the controller; this value does not include meta-data. This value is reported in thousands (i.e., a value of 1 corresponds to 1000 units of 512 bytes read) and is rounded up. When the LBA size is a value other than 512 bytes, the controller shall convert the amount of data read to 512 byte units.
48	16	Data Units Write (in LBAs)	Contains the number of 512 byte data units the host has written to the controller; this value does not include meta-data. This value is reported in thousands (i.e., a value of 1 corresponds to 1000 units of 512 bytes written) and is rounded up. When the LBA size is a value other than 512 bytes, the controller shall convert the amount of data written to 512 byte units. For the NVM command set, logical blocks written as part of Write operations shall be included in this value. Write Uncorrectable commands shall not impact this value.
64	16	Host Read Commands	Contains the number of read commands issued to the controller.
80	16	Host Write Commands	Contains the number of write commands issued to the controller.
96	16	Controller Busy Time (in minutes)	Contains the amount of time the controller is busy with I/O commands. The controller is busy when there is a command outstanding to an I/O Queue (specifically, a command was issued by way of an I/O Submission Queue Tail doorbell write and the corresponding completion queue entry has not been posted yet to the associated I/O Completion Queue). This value is reported in minutes. Note: Controller Busy Time may not be equal to I/O time, as controller may release the operation to SW routines.
112	16	Power Cycles	Contains the number of power cycles.
128	16	Power On Hours	Contains the number of power-on hours. This does not include time that the controller was powered and in a low power state condition.
144	16	Unsafe shutdowns	Contains the number of unsafe shutdowns. This count is incremented when a shutdown notification (CC.SHN) is not received prior to loss of power.

Table 26: SMART Attributes (Log Identifier 02h)

Byte	# of Bytes	Attribute	Description
160	16	Media Errors	Contains the number of occurrences where the controller detected an unrecovered data integrity error. Errors such as uncorrectable ECC, CRC checksum failure, or LBA tag mismatch are included in this field.
176	16	Number of Error Information Log Entries	Contains the number of Error Information log entries over the life of the controller.
195:192	4	Warning Composite Temperature Time	Contains the amount of time in minutes that the controller is operational and the Composite Temperature is greater than or equal to the Warning Composite Temperature Threshold (WCTEMP) field and less than the Critical Composite Temperature Threshold (CCTEMP) field in the Identify Controller data structure in Appendix C. If the value of the WCTEMP or CCTEMP field is 0h, then this field is always cleared to 0h regardless of the Composite Temperature value.
199:196	4	Critical Composite Temperature Time	Contains the amount of time in minutes that the controller is operational and the Composite Temperature is greater than the Critical Composite Temperature Threshold (CCTEMP) field in the Identify Controller data structure in Figure 90. If the value of the CCTEMP field is 0h, then this field is always cleared to 0h regardless of the Composite Temperature value.
201:200	2	Temperature Sensor 1	Currently not implemented on D7-P5520/D7-P5620
203:202	2	Temperature Sensor 2	Currently not implemented on D7-P5520/D7-P5620
205:204	2	Temperature Sensor 3	Currently not implemented on D7-P5520/D7-P5620
207:206	2	Temperature Sensor 4	Currently not implemented on D7-P5520/D7-P5620
209:208	2	Temperature Sensor 5	Currently not implemented on D7-P5520/D7-P5620
211:210	2	Temperature Sensor 6	Currently not implemented on D7-P5520/D7-P5620
213:212	2	Temperature Sensor 7	Currently not implemented on D7-P5520/D7-P5620
215:214	2	Temperature Sensor 8	Currently not implemented on D7-P5520/D7-P5620
219:216	4	Thermal Management Temperature 1 Transition Count	Contains the number of times the controller transitioned to lower power active power states or performed vendor specific thermal management actions while minimizing the impact on performance in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature. This counter shall not wrap once it reaches its maximum value. A value of zero indicates that this transition has never occurred. This is not supported on D7-P5520/D7-P5620.

Table 26: SMART Attributes (Log Identifier 02h)

Byte	# of Bytes	Attribute	Description
223:220	4	Thermal Management Temperature 2 Transition Count	Contains the number of times the controller transitioned to lower power active power states or performed vendor specific thermal management actions regardless of the impact on performance (e.g., heavy throttling) in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature. This is not supported on D7-P5520/D7-P5620.
227:224	4	Total Time For Thermal Management Temperature 1	Contains the number of seconds that the controller had transitioned to lower power active power states or performed vendor specific thermal management actions while minimizing the impact on performance in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature. This is not supported on D7-P5520/D7-P5620.
231:228	4	Total Time For Thermal Management Temperature 2	Contains the number of seconds that the controller had transitioned to lower power active power states or performed vendor specific thermal management actions regardless of the impact on performance (e.g., heavy throttling) in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature. This is not supported on D7-P5520/D7-P5620.
511:232	N/A	Reserved	

Note: Solidigm recommends the host wait at least 90 seconds after power ON before retrieving the most current counter values. From SMART Log Attributes (Log 02h)

Table 27: Get Log Page - Temperature Sensor Data Structure

Bits	Description
15:00	Temperature Sensor Temperature (TST): Contains the current temperature in degrees Kelvin reported by the temperature sensor. The physical point in the NVM subsystem whose temperature is reported by the temperature sensor and the temperature accuracy is implementation specific. An implementation that does not implement the temperature sensor reports a temperature of zero degrees Kelvin. The temperature reported by a temperature sensor may be used to trigger an asynchronous event

5.6 SET Features Identifiers

In addition to the SMART attribute structure, features pertaining to the operation and health of the D7-P5520/D7-P5620 can be reported to the host on request through the Get Features command. Get Features support is per device and not

available for a per namespace basis. The user can change settings using SET Features on the following items as defined in NVMe 1.4 specification.

Please refer to NVMe 1.4 specification for details of the following SET features:

- Arbitration (Feature Identifier 01h)
- Power Management (Feature Identifier 02h)¹
- Temperature Threshold (Feature Identifier 04h)
- Error Recovery (Feature Identifier 05h)¹
- Number of Queues (Feature Identifier 07h)
- Interrupt Coalescing (Feature Identifier 08h)
- Interrupt Vector Configuration (Feature Identifier 09h)
- Write Atomicity (Feature Identifier 0Ah)¹
- Asynchronous Event Configuration (Feature Identifier 0Bh)
- Timestamp (Feature Identifier 0Eh)
- Controller Metadata (Feature Identifier 7Eh)
- Namespace Metadata (Feature Identifier 7Fh)

Note:

1. Commands to Get/Set features Power Management, Error Recovery, Write Atomicity not supported during Sanitize operations.

D7-P5520/D7-P5620 also supports the following Vendor Unique Opcodes:

- Clear Assert (Opcode C8h)

5.7 Vendor Unique Opcodes

5.7.1 Clear Assert (Opcode C8h)

This vendor-unique Opcode will:

- Clear drive assert, but preserve SMART data, remove user data.

Note: These two features can be supported when NCAT feature is disabled

Table 28: Clear Assert - Command DWORD 10

Bits	Description
31:00	<p>Write Usage:</p> <p>Bit [15:0] Reserved</p> <p>Bit [31:16] reserved for Assert Error Clear feature</p> <p>0x10000 = Clear Assert, Preserve SMART data, Remove User Data (System Data Recovery)</p> <p>0x20000 = Read Only, drive remains asserted (Host Data Recovery)</p> <p>Read Usage: None</p>

5.8 Vendor Unique Feature Identifiers (Set/Get)

Solidigm™ D7-P5520/D7-P5620 supports the following Vendor Unique FIDs (VU) SET Features:

- Set/Get Endurance Management Disable Setting (C0h)
- Set/Get Max LBA (C1h)
- Set/Get Native Max LBA (C2h)
- Extend Temp Setting (C3h)
- Set/Get Power Governor Setting (C6h)
- Set/Get SMB ASIC Address (C8h)
- Set/Get Blink activity for LED (C9h)
- Set SMART Execute Self Test (D4h)
- Reset Timed Workload Counters (D5h)
- Set/Get Enable Latency Tracking (E2h)
- Set Read/Write IO Latency Threshold (F4h)
- Set/Get Feature NAND Location for Vt Histo (F6h)

5.8.1 Set/Get Endurance Management Disable Setting (C0h)

This Vendor Unique command returns the Endurance Management setting.

Table 29: C0h - Endurance Management Disable Setting - Command Dword 11

Bit	Description
31:01	Reserved
00	Read Usage: 0 = Enabled, 1 = Disabled

If a Get Features command is issued for this Feature, the attributes specified are returned in Dword 0 of the completion queue entry for that command.

5.8.2 Set/Get Max LBA (C1h)

This Vendor Unique command sets/ gets the maximum addressable Logical Block for the entire Drive. Command Dword 11 and Command Dword 12 are used to transfer up to 64 bits of LBA information. This is a global setting that applies to the entire drive that namespaces may be allocated from.

- All namespaces must be deleted before using the SetMaxLBA command
- LBA size that this command sets and reports is in 512B units regardless of the logical LBA size
- Setting a value higher than device specified IDEMA capacity will result in a returned NVMe error status, see Table 31 on page 53
- Setting a value that is lower than 1GB will result in a returned NVMe error status
- Drive firmware rounds up Set Max LBA to the nearest 1GB boundary

Device will be shipped as default single namespace with IDEMA mode out of factory. Once this default namespace is deleted, it will remain in namespace mode throughout the life of the device even though 1 namespace is created with maximum capacity. Single namespace can be created with native factory capacity.

With multiple namespaces, each new namespace created is aligned to 1GB granularity rounding up by the firmware. This is controller limitation, Solidigm does not expect this to go away in next generation product. It is recommended for customers not to utilize SetMaxLBA command in combination with Multiple Namespace. If drive is over-provisioned with SetMaxLBA, drive will disable support for multiple namespaces (MNS). To re-enable support for MNS drive should be reverted back to native capacity using SetMaxLBA command. With multiple namespace, each namespace can have single locking range spanning the allocated LBA capacity. On a single namespace range, multiple locking range can be initiated, but the namespace should be configured to native LBA capacity.

Table 30: Set Max LBA Setting - Command Dword 11 and Command Dword 12

Bit	Description
63:00	Maximum User LBA: Write Usage: This field sets the 64-bit maximum LBA addressable by the Drive. Read Usage: This field contains the 64-bit maximum LBA addressable by the Drive. Command Dword 11 contains bits 31:00; Command Dword 12 contains bits 63: 32.

The return status code type (SCT) 0x1 (Command Specific Status), status codes (SC) for Set Max LBA command are listed in the table below.

Table 31: Status Code - Set Max LBA Command Specific Status Values

Value	Description
C0h	Requested MAX LBA exceeds Available capacity
C1h	Requested MAX LBA smaller than minimum allowable limit.
C2h	Requested MAX LBA is smaller than allocated Namespace requirements

5.8.3 Set/Get Native Max LBA (C2h)

This vendor unique command enables the user to read Native max LBA value, which is the highest LBA that the device accepts. This command would be of value to an end user attempting to restore the full capacity of the drive, without having first recorded the max value, or using trial/error to determine the effective max LBA size.

Set Features of the Native Max LBA FID shall always return error "Invalid Command 01h"

LBA size that this command reports is in 512B units no matter what the current logical format.

5.8.4 Extend Temp Setting (C3h)

Table 32: C3h - Extend Temp Setting - Command Dword 11

Bit	Description
31:08	Reserved
07:00	Extend Temp setting: Write Usage - Turns on/off extended case temperature settings Read Usage: Reads the current setting whether it is turned on/off

If a Get Features command is issued for this Feature, the attributes specified are returned in Dword 0 of the completion queue entry for that command.

The return status code type (SCT) 0x1 (Command Specific Status), status codes (SC) for Extend thermal setting command are:

Table 33: Status Codes - Extend Temp Setting Command Specific Status Values

Value	Description
C0h	Invalid Setting

5.8.5 Set/Get Power Governor Setting (C6h)

Table 34: Set/Get Power (Typical) Governor Setting (C6h)

Bit	Description				
31:08	Reserved				
07:00	Power Mode	1.92TB/1.6TB	3.84TB/3.2TB	7.68TB/6.4TB	15.36TB/12.8TB
	PM0	15W	15W	18W	20W
	PM1	11W	13W	15W	15W
	PM2	9W	11W	11W	11W

The return status code type (SCT) 0x1 (Command Specific Status), status codes (SC) for Power Governor setting command are:

Table 35: Status Codes - Power Governor Setting Command Specific Status Value

Value	Description
C0h	Invalid Setting

5.8.6 Set/Get SMB ASIC Address (C8h)

This FID enables customers to disable or change SMB Address, in the field, if address conflict arises.

Setting an SMB address greater than 7Fh (7bit slave address corresponding to 1111 111) will disable SMB ASIC access and set the address to 0x80. If Bits[9:1] is set to 0xFF, FW disabled SMB ASIC access

If a Get Features command is issued for this Feature, the attributes specified are returned in Dword 0 of the completion queue entry for that command. Refer to Section 2.12 for details on SMBus address.

Table 36: C8h - Get/Set SMB ASIC Address

Bit	Description
0	Reserved
09:01	SMB Controller Address
31:10	Reserved

5.8.7 Set/Get Blink activity for LED (C9h)

Table 37: C9h - Set/Get Blink activity for LED

Feature Option	Feature Value Range	Description	Default
0	0-1	LED state while host is inactive. (0 = OFF, 1 = ON)	1 (ON)
1	0-1	LED duration increment size. (0 = 50ms, 1 = 25ms)	0 (50ms)
2	0-15	Off duration during IO activity in 25ms/50ms increments. 0 = solid ON	5 (ON)
3	0-15	On duration during IO activity in 25ms/50ms increments. 0 = match OFF duration	0
4	0-15	Off duration during format activity in 25/50ms increments. 0 = solid ON	5 (250ms)
5	0-15	On duration during format activity in 25/50ms increments. 0 = match OFF duration	0 (250ms)

Note:

- C9- Set Features Command Dword 11 will be divided into following sections
 - Bits[31:24] - Feature options and Bits[23:0] - Feature Value
 - Feature options and value ranges are defined above
- C9- Get Features Command Dword 11 will be divided into following sections
 - Bits[31:24] - Feature options and Bits[23:0] - Reserved. Must be 0.
 - Current value for the requested LED feature option will be returned in DW0[31:0]

The return status code type (SCT) 0x1 (Command Specific Status), status codes (SC) for "Get/Set Blink Activity for LED" interval setting are:

Table 38: Status Code - Set/ Get P11 Blink Activity for LED Specific Status Values

Value	Description
C0h	Invalid Setting

5.8.8 Set SMART Execute Self Test (D4h)

This vendor unique setting initiates a SMART Self-Test in various modes.

Table 39: D4h - Set SMART Execute Self Test - Command Dword 11

Bit	Description
31:08	Reserved
07:00	<p>Execute Offline Test:</p> <p>Write Usage:</p> <p>0: Reserved</p> <p>1: Execute Short Self Test in Offline Mode</p> <p>2: Execute SMART Extended self-test in Offline mode</p> <p>3-63: Reserved</p> <p>64: Reset Timed Workload Attributes (E1, E2, E3 legacy SATA SMART Attributes)</p> <p>65-126: Reserved</p> <p>127: Abort off-line mode Self-test routine</p> <p>128: Reserved</p> <p>129: Execute SMART Short self-test routine immediately in captive mode</p> <p>130: Execute SMART Extended Self-test routine immediately in captive mode</p> <p>131 - 255: Reserved</p> <p>Read Usage: Returns the last written command.</p>

If a Get Features command is issued for this Feature, the attributes specified are returned in Dword 0 of the completion queue entry for that command.

The return status code type (SCT) 0x1 (Command Specific Status), status codes (SC) for SMART Execute Self-Test command are:

Value	Description
C0h	Invalid Test/Setting

5.8.9 Get Reset Timed Workload Counters (D5h)

Table 40: D5h - Reset Timed Workload Counters - Command Dword 11

Bit	Description
31:01	Reserved
00	Timed Workload Reset Settings: Write Usage: 00 = NOP, 1 = Reset E2, E3,E4 counters; Read Usage: Not Supported

Note: Get Features will not work for "Reset Timed Workload Counters," and status code is same as in the table above.

5.8.10 Set/Get Enable Latency Tracking (E2h)

Table 41: E2h - Set/Get Enable Latency Tracking

Bit	Description
31:01	Write Usage: 00h = Disable Latency Tracking (Default) 01h = Enable Latency Tracking

5.8.11 Set Read/Write IO Latency Threshold (F4h)

This vendor unique setting enables setting of read/write IO latency threshold. This is used to trigger the slow IO logging (Log Page 0xFD). The host command latency threshold values will be set in DW11.

Table 42: F4h - Set Read/Write IO Latency Threshold

Bit	Description
00	0 for read limit and 1 for write limit
31:01	Latency limit in μ s

5.8.12 Set/Get Feature NAND Location for Vt Histo (F6h)

This feature allows the customer to set and verify the NAND location needed for NAND debug. They can then use the F6h LID to pull the histo data from defined set Feature NAND location. The set/get feature command must be Dword11 size of 20 bytes.

Table 43: F6h - Set/Get NAND Location for Vt Histo - Command Dword 11 = 20 bytes

Bit	Description
00:15	Version Major
16:31	Version Minor
32:63	Die
64:95	Block
96:127	Page
128:159	Media Bank ID

6 NVMe Driver Support

The following table describes the NVMe Driver Support for D7-P5520/D7-P5620. The support includes releasing and validating NVMe drivers for certain operating systems and validating functionality for open source drive, inbox, or native drivers for select operating systems.

It is recommended to use an NVMe Driver that supports Namespace Optimal IO Boundary (NOIOB) with this product. If an NVMe Driver is used that does not support NOIOB issuing IO commands that do not align to 128KB boundaries may result in a significant drop in performance and Quality of Service (QoS). If the driver being used does not support NOIOB, it is recommended that the host maintains splitting of 128KB IO boundaries in order to avoid a performance impact. NOIOB is supported on Linux kernel v4.13 or greater.

Table 44: Solidigm™ D7-P5520 NVMe Driver Support



Support Level	Operating System Description
Solidigm Data Center NVMe Windows Driver	The Solidigm Data Center NVMe Windows Driver has reached end-of-life and is no longer available as of February 6th, 2023.
In-box Driver or external package ¹	Microsoft Windows Inbox NVMe Driver VMware Inbox Driver (ESXi) Linux Upstream Inbox Driver (tested on kernel v5.15.24) Linux OSV Native Driver

Note:

1. NVMe driver support list is subject to change in future product specification releases.

7 Other Compliance and Certifications

Table 45: Other Compliance and Certifications

<div>NVMe</div> <div></div>	Indicates compliance with UNH-IOL testing for NVMe compliance
<div>PCIe</div> <div></div>	Indicates compliance with PCI-SIG Organization testing requirements

Appendix

Appendix A: Performance and Endurance Metrics

Table 46: Solidigm™ D7-P5520/D7-P5620 User Addressable Sectors

	Capacity	Unformatted Capacity ¹ (Total User Addressable Sectors in LBA Mode)
D7-P5520	1.92TB	3,750,748,848
	3.84TB	7,501,476,528
	7.68TB	15,002,931,888
	15.36TB	30,005,842,608
D7-P5620	1.6TB	3,125,627,568
	3.2TB	6,251,233,968
	6.4TB	12,502,446,768
	12.8TB	25,004,872,368

Note:

1. The total usable capacity of the SSD may be less than the total physical capacity because a small portion of the capacity is used for NAND media management and maintenance. IDEMA standard is used. 1 sector = 512 bytes LBA count shown represents total user storage capacity and will remain the same throughout the life of the drive.
2. 1TB = 10¹² bytes.

Table 47: Random Read/Write (IOPS) - D7-P5520 - U.2 15mm

Capacity	PCIe Gen	Solidigm™ D7-P5520					
		Specification ^{1, 2, 3}					
		4KB Random Write ⁴	8KB Random Write ⁴	4KB Random Read ⁵	8KB Random Read ⁵	4KB Random 70/30 Read/Write ⁴	8KB Random 70/30 Read/Write ⁴
1.92TB	3.0	114,000	58,000	700,000	425,000	177,000	110,000
	4.0	114,000	58,000	700,000	612,000	177,000	110,000
3.84TB	3.0	200,000	97,000	747,000	425,000	305,000	188,000
	4.0	200,000	97,000	1,000,000	755,000	305,000	188,000
7.68TB	3.0	220,000	100,000	806,000	425,000	400,000	242,000
	4.0	220,000	100,000	1,100,000	800,000	400,000	242,000

Table 47: Random Read/Write (IOPS) - D7-P5520 - U.2 15mm

Capacity	PCIe Gen	Solidigm™ D7-P5520					
		Specification ^{1,2,3}					
		4KB Random Write ⁴	8KB Random Write ⁴	4KB Random Read ⁵	8KB Random Read ⁵	4KB Random 70/30 Read/Write ⁴	8KB Random 70/30 Read/Write ⁴
15.36TB	3.0	200,000	109,000	806,000	425,000	390,000	230,000
	4.0	200,000	109,000	1,000,000	800,000	390,000	235,000

Note:

1. 4KB = 4,096 bytes; 8KB = 8,192 bytes
2. For 100% read workloads, drive should be pre-conditioned with the equivalent 100% write workload transfer size before measurement.
3. Solidigm expects up to 5% variation in performance between drive to drive runs. Any change in the system or drive configuration may impact drive performance.
4. Performance measured with Queue Depth 64 x 4 workers (Universal QD=256). Measurements are performed on a full Logical Block Address (LBA) span of the drive. Power mode set at PM0.
5. Performance measured with Queue Depth 32 x 8 workers (Universal QD=256). Measurements are performed on a full Logical Block Address (LBA) span of the drive. Power mode set at PM0.

Table 48: Random Read/Write (IOPS) - D7-P5520 - EDSFF

Capacity	PCIe Gen	Queue Depth	Workers	Solidigm™ D7-P5520					
				Specification ^{1,2,3}					
				4KB Random Write ⁴	8KB Random Write ⁴	4KB Random Read ⁵	8KB Random Read ⁵	4KB Random 70/30 Read/Write ⁴	8KB Random 70/30 Read/Write ⁴
1.92TB	3.0	32	8	114,000	58,000	706,000	612,000	177,000	110,000
	4.0	32	8	114,000	58,000	706,000	612,000	177,000	110,000
3.84TB	3.0	32	8	147,000	74,000	747,000	425,000	250,000	155,000
	4.0	32	8	147,000	74,000	1,000,000	727,000	250,000	155,000
7.68TB	3.0	32	8	200,000	93,000	806,000	425,000	400,000	242,000
	4.0	32	8	220,000	100,000	1,100,000	751,000	400,000	242,000

Table 48: Random Read/Write (IOPS) - D7-P5520 - EDSFF

Capacity	PCIe Gen	Queue Depth	Workers	Solidigm™ D7-P5520					
				Specification ^{1,2,3}					
				4KB Random Write ⁴	8KB Random Write ⁴	4KB Random Read ⁵	8KB Random Read ⁵	4KB Random 70/30 Read/Write ⁴	8KB Random 70/30 Read/Write ⁴
15.36TB	3.0	32	8	151,000	79,000	806,000	425,000	390,000	235,000
	4.0	32	8	200,000	109,000	1,000,000	800,000	390,000	235,000

Note:

1. 4KB = 4,096 bytes; 8KB = 8,192 bytes
2. For 100% read workloads, drive should be pre-conditioned with the equivalent 100% write workload transfer size before measurement.
3. Solidigm expects up to 5% variation in performance between drive to drive runs. Any change in the system or drive configuration may impact drive performance.
4. Performance measured with Queue Depth 64 x 4 workers (Universal QD=256). Measurements are performed on a full Logical Block Address (LBA) span of the drive. Power mode set at PM0.
5. Performance measured with Queue Depth 32 x 8 workers (Universal QD=256). Measurements are performed on a full Logical Block Address (LBA) span of the drive. Power mode set at PM0.

Table 49: Random Read/Write (IOPS) - D7-P5620

Capacity	PCIe Gen	Solidigm™ D7-P5620					
		Specification ^{1,2,3}					
		4KB Random Write ⁴	8KB Random Write ⁴	4KB Random Read ⁵	8KB Random Read ⁵	4KB Random 70/30 Read/Write ⁴	8KB Random 70/30 Read/Write ⁴
1.6TB	3.0	200,000	105,000	700,000	425,000	241,000	157,000
	4.0	200,000	105,000	700,000	611,000	241,000	157,000
3.2TB	3.0	341,000	180,000	747,000	425,000	402,000	270,000
	4.0	341,000	180,000	1,000,000	755,000	402,000	270,000
6.4TB	3.0	390,000	215,000	806,000	425,000	514,000	276,000
	4.0	390,000	215,000	1,100,000	800,000	514,000	359,000
12.8TB	3.0	374,000	210,000	806,000	425,000	595,000	307,000
	4.0	374,000	210,000	1,000,000	800,000	595,000	410,000

Note:

1. 4KB = 4,096 bytes; 8KB = 8,192 bytes
2. For 100% read workloads, drive should be pre-conditioned with the equivalent 100% write workload transfer size before measurement.
3. Solidigm expects up to 5% variation in performance between drive to drive runs. Any change in the system or drive configuration may impact drive performance.
4. Performance measured with Queue Depth 64 x 4 workers (Universal QD=256). Measurements are performed on a full Logical Block Address (LBA) span of the drive. Power mode set at PM0.
5. Performance measured with Queue Depth 32 x 8 workers (Universal QD=256). Measurements are performed on a full Logical Block Address (LBA) span of the drive. Power mode set at PM0.

Table 50: Random Read/Write IOPS Consistency (%) - D7-P5520

Capacity	Solidigm™ D7-P5520			
	Specification ^{1,2,3}			
	4KB Random Read	4KB Random Write	8KB Random Read	8KB Random Write
1.92TB	90%	85%	90%	85%
3.84TB	90%	85%	90%	85%
7.68TB	90%	85%	90%	85%
15.36TB	90%	85%	90%	85%

Note:

1. 4KB = 4,096 bytes; 8KB = 8,192 bytes
2. Performance measured with Queue Depth 32 x 8 workers (Universal QD=256).
3. Solidigm expects up to 10% variation in consistency between drive to drive runs. Any change in the system or drive configuration may impact drive performance.

Table 51: Random Read/Write IOPS Consistency (%) - D7-P5620

Capacity	Solidigm™ D7-P5620			
	Specification ^{1,2,3}			
	4KB Random Read	4KB Random Write	8KB Random Read	8KB Random Write
1.6TB	90%	85%	90%	85%
3.2TB	90%	85%	90%	85%
6.4TB	90%	85%	90%	85%
12.8TB	90%	85%	90%	85%

Note:

1. 4KB = 4,096 bytes; 8KB = 8,192 bytes

2. Performance measured with Queue Depth 32 x 8 workers (Universal QD=256).
3. Solidigm expects up to 10% variation in consistency between drive to drive runs. Any change in the system or drive configuration may impact drive performance.

Table 52: Sequential Read and Write Bandwidth (MB/s) - D7-P5520 - U.2 15mm

Capacity	PCIe Gen	Queue Depth	Workers	Solidigm™ D7-P5520	
				Specification ¹	
				Sequential Read	Sequential Write
1.92TB	3.0	256	1	3,400	1,900
	4.0	256	1	5,300	1,900
3.84TB	3.0	256	1	3,500	3,400
	4.0	256	1	6,700	3,600
7.68TB	3.0	256	1	3,500	3,400
	4.0	256	1	7,100	4,200
15.36TB	3.0	256	1	3,500	3,400
	4.0	256	1	7,100	3,700

Note:

1. Performance measured using FIO Linux CentOS 8.2.2004 kernel 5.4.49 with 128KB (131,072 bytes) of transfer size with Queue Depth 256 (1 worker). Measurements are performed on a full Logical Block Address (LBA) span of the drive once the workload has reached steady state but including all background activities required for normal operation and data reliability. Power mode set at PM0
2. Solidigm expects up to 5% variation in performance between drive to drive runs. Any change in the system or drive configuration may impact drive performance.

Table 53: Sequential Read and Write Bandwidth (MB/s) - D7-P5520 - EDSFF

Capacity	PCIe Gen	Queue Depth	Workers	Solidigm™ D7-P5520	
				Specification ^{2,3}	
				Sequential Read	Sequential Write
1.92TB	3.0	256	1	3,400	1,900
	4.0	256	1	5,300	1,900
3.84TB	3.0	256	1	3,500	3,250
	4.0	256	1	6,500	3,400
7.68TB	3.0	256	1	3,500	3,400
	4.0	256	1	6,600	4,200

Table 53: Sequential Read and Write Bandwidth (MB/s) - D7-P5520 - EDSFF

Capacity	PCIe Gen	Queue Depth	Workers	Solidigm™ D7-P5520	
				Specification ^{2,3}	
				Sequential Read	Sequential Write
15.36TB	3.0	256	1	3,500	3,400
	4.0	256	1	7,100	3,700

Note:

1. Performance measured using FIO Linux Centos 7.2 kernel 4.8.6 with 128KB (131,072 bytes) of transfer size with Queue Depth 256 (1 worker). Measurements are performed on a full Logical Block Address (LBA) span of the drive once the workload has reached steady state but including all background activities required for normal operation and data reliability. Power mode set at PM0
2. Solidigm expects up to 5% variation in performance between drive to drive runs. Any change in the system or drive configuration may impact drive performance.

Table 54: Sequential Read and Write Bandwidth (MB/s) - D7-P5620

Capacity	PCIe Gen	Queue Depth	Workers	Solidigm™ D7-P5620	
				Specification ¹	
				Sequential Read	Sequential Write
1.6TB	3.0	256	1	3,400	1,900
	4.0	256	1	5,300	1,900
3.2TB	3.0	256	1	3,500	3,400
	4.0	256	1	6,700	3,600
6.4TB	3.0	256	1	3,500	3,400
	4.0	256	1	7,100	4,200
12.8TB	3.0	256	1	3,500	3,400
	4.0	256	1	7,100	3,700

Note:

1. Performance measured using FIO Linux Centos 7.2 kernel 4.8.6 with 128KB (131,072 bytes) of transfer size with Queue Depth 256 (1 worker). Measurements are performed on a full Logical Block Address (LBA) span of the drive once the workload has reached steady state but including all background activities required for normal operation and data reliability. Power mode set at PM0
2. Solidigm expects up to 5% variation in performance between drive to drive runs. Any change in the system or drive configuration may impact drive performance.

Table 55: Typical Latency - D7-P5520- U.2 15mm

Capacity	Solidigm™ D7-P5520			
	Specification			
	4K Sequential ¹ Read QD1	4K Sequential ¹ Write QD1	4K Random ¹ Read QD1	4K Random ¹ Write QD1
1.92TB	10µs	13µs	75µs	15µs
3.84TB	10µs	13µs	75µs	15µs
7.68TB	10µs	13µs	75µs	15µs
15.36TB	10µs	13µs	75µs	20µs

Note:

1. Latency measured using 4 KB (4,096 bytes) transfer size with Queue Depth equal to 1.
2. For 100% read workloads, drive should be pre-conditioned with the equivalent 100% write workload transfer size before measurement.

Table 56: Typical Latency - D7-P5520- EDSFF

Capacity	Solidigm™ D7-P5520			
	Specification			
	4K Sequential ¹ Read QD1	4K Sequential ¹ Write QD1	4K Random ¹ Read QD1	4K Random ¹ Write QD1
1.92TB	10µs	13µs	75µs	15µs
3.84TB	10µs	13µs	75µs	15µs
7.68TB	10µs	13µs	75µs	15µs
15.36TB	10µs	13µs	75µs	20µs

Note:

1. Latency measured using 4 KB (4,096 bytes) transfer size with Queue Depth equal to 1.
2. For 100% read workloads, drive should be pre-conditioned with the equivalent 100% write workload transfer size before measurement.

Table 57: Typical Latency - D7-P5620- U.2 15mm

Capacity	Solidigm™ D7-P5620			
	Specification			
	4K Sequential ¹ Read QD1	4K Sequential ¹ Write QD1	4K Random ¹ Read QD1	4K Random ¹ Write QD1
1.6TB	10µs	13µs	75µs	15µs
3.2TB	10µs	13µs	75µs	15µs
6.4TB	10µs	13µs	75µs	15µs
12.8TB	10µs	13µs	75µs	15µs

Note:

1. Latency measured using 4 KB (4,096 bytes) transfer size with Queue Depth equal to 1.
2. For 100% read workloads, drive should be pre-conditioned with the equivalent 100% write workload transfer size before measurement.

Table 58: Quality of Service - D7-P5520 - U.2 15mm

Capacity	Queue Depth	Unit	Solidigm™ D7-P5520			
			Specification ^{1,2,3}			
			4K Random Read	4K Random Write	8K Random Read	8K Random Write
Quality of Service ¹ (99%)						
1.92TB	QD1	µs	99	60	100	90
3.84TB	QD1	µs	99	60	100	70
7.68TB	QD1	µs	99	60	100	70
15.36TB	QD1	µs	99	70	100	95
Quality of Service ¹ (99.99%)						
1.92TB	QD1	µs	143	150	150	1800
3.84TB	QD1	µs	143	130	150	150
7.68TB	QD1	µs	143	130	150	160
15.36TB	QD1	µs	143	170	150	200

Note:

1. Measured as the time taken for 99.0 or 99.99 percentile of commands to finish the round-trip from host to drive and back to host.

2. Data measured using FIO Linux Centos 7.5.1804 kernel 4.14.74 Quality of Service measured using 4KB (4,096 bytes) transfer size on a random workload on a full Logical Block Address (LBA) span of the drive once the workload has reached steady state but including all background activities required for normal operation and data reliability.
3. For 100% read workloads, drive should be pre-conditioned with the equivalent 100% write workload transfer size before measurement.
4. Solidigm expects up to 10% variation in performance between drive to drive runs. Any change in the system or drive configuration may impact drive performance.

Table 59: Quality of Service - D7-P5520 - EDSFF

Capacity	Queue Depth	Unit	Solidigm™ D7-P5520			
			Specification ^{1,2,3}			
			4K Random Read	4K Random Write	8K Random Read	8K Random Write
Quality of Service ¹ (99%)						
1.92TB	QD1	µs	99	60	100	90
3.84TB	QD1	µs	99	60	100	80
7.68TB	QD1	µs	99	60	100	70
15.36TB	QD1	µs	99	70	100	95
Quality of Service ¹ (99.99%)						
1.92TB	QD1	µs	143	150	150	1800
3.84TB	QD1	µs	143	170	150	170
7.68TB	QD1	µs	143	130	150	160
15.36TB	QD1	µs	143	170	150	200

Note:

1. Measured as the time taken for 99.0 or 99.99 percentile of commands to finish the round-trip from host to drive and back to host.
2. Data measured using FIO Linux Centos 7.5.1804 kernel 4.14.74 Quality of Service measured using 4KB (4,096 bytes) transfer size on a random workload on a full Logical Block Address (LBA) span of the drive once the workload has reached steady state but including all background activities required for normal operation and data reliability.
3. For 100% read workloads, drive should be pre-conditioned with the equivalent 100% write workload transfer size before measurement.
4. Solidigm expects up to 10% variation in performance between drive to drive runs. Any change in the system or drive configuration may impact drive performance.

Table 60: Quality of Service - D7-P5620

Capacity	Queue Depth	Unit	Solidigm™ D7-P5620			
			Specification ^{1,2,3}			
			4K Random Read	4K Random Write	8K Random Read	8K Random Write
Quality of Service ¹ (99%)						
1.6TB	QD1	µs	99	50	100	60
3.2TB	QD1	µs	99	50	100	60
6.4TB	QD1	µs	99	50	100	50
12.8TB	QD1	µs	99	60	100	60
Quality of Service ¹ (99.99%)						
1.6TB	QD1	µs	143	150	150	160
3.2TB	QD1	µs	143	120	150	130
6.4TB	QD1	µs	143	110	150	120
12.8TB	QD1	µs	143	120	150	130

Note:

1. Measured as the time taken for 99.0 or 99.99 percentile of commands to finish the round-trip from host to drive and back to host.
2. Data measured using FIO Linux Centos 7.5.1804 kernel 4.14.74 Quality of Service measured using 4KB (4,096 bytes) transfer size on a random workload on a full Logical Block Address (LBA) span of the drive once the workload has reached steady state but including all background activities required for normal operation and data reliability.
3. For 100% read workloads, drive should be pre-conditioned with the equivalent 100% write workload transfer size before measurement.
4. Solidigm expects up to 10% variation in performance between drive to drive runs. Any change in the system or drive configuration may impact drive performance.

Table 61: Endurance - Drive Writes Per Day (DWPD) - U.2 15mm

Form Factor	Capacity	JEDEC Workload			64K Sequential Workload		
		DWPD ^{1,4} (3 years)	DWPD ^{1,4} (5 years)	PBW ^{2,3,4}	DWPD ^{1,4} (3 years)	DWPD ^{1,4} (5 years)	PBW ^{2,3,4}
P5520 - U.2 15mm	1.92TB	1.6	1	3.5	9.5	5.7	20
	3.84TB	1.6	1	7	9.6	5.7	40.6
	7.68TB	1.6	1	14	9	5.4	76.5
	15.36TB	1.6	1	28	8	4.8	134

Table 61: Endurance - Drive Writes Per Day (DWPD) - U.2 15mm

Form Factor	Capacity	JEDEC Workload			64K Sequential Workload		
		DWPD ^{1,4} (3 years)	DWPD ^{1,4} (5 years)	PBW ^{2,3,4}	DWPD ^{1,4} (3 years)	DWPD ^{1,4} (5 years)	PBW ^{2,3,4}
P5620 - U.2 15mm	1.6TB	4.8	3	8.7	11.4	6.8	20
	3.2TB	4.8	3	17.5	11.5	6.9	40.6
	6.4TB	4.8	3	35	10.9	6.5	76.5
	12.8TB	4.4	2.8	65.4	9.6	5.7	134

Note:

1. Refer to JESD219 standard table 1 for UBER, FFR and other Enterprise SSD endurance verification requirements. Endurance verification acceptance criterion based on establishing <1E-17
2. 1PB = 10¹⁵ bytes
3. Petabytes Written (PBW). Refer to JESD219 standard table 1 for UBER, FFR and other Enterprise SSD requirements. The number of drive writes such that the SSD meets the requirements according to the JESD218 standard. Endurance rating verification is defined to establish UBER <1E-16 at 60% upper confidence limit.
4. If the drive is being used in VSS Mode (LBA Format with Metadata size > 0 Bytes) the Drive Endurance will be reduced by ~30%.

Table 62: Endurance - Drive Writes Per Day (DWPD) - EDSFF

Form Factor	Capacity	JEDEC Workload			64K Sequential Workload		
		DWPD ^{1,4} (3 years)	DWPD ^{1,4} (5 years)	PBW ^{2,3,4}	DWPD ^{1,4} (3 years)	DWPD ^{1,4} (5 years)	PBW ^{2,3,4}
P5520 - E1.S 9.5mm &15mm	1.92TB	1.6	1	3.5	9.5	5.7	20
	3.84TB	1.6	1	7	9.6	5.7	40.6
	7.68TB	1.6	1	14	9.1	5.4	76.5
P5520 - E1.L 9.5mm	15.36TB	1.6	1	28	8	4.8	134

Note:

1. Refer to JESD219 standard table 1 for UBER, FFR and other Enterprise SSD endurance verification requirements. Endurance verification acceptance criterion based on establishing <1E-17
2. 1PB = 10¹⁵ bytes
3. Petabytes Written (PBW). Refer to JESD219 standard table 1 for UBER, FFR and other Enterprise SSD requirements. The number of drive writes such that the SSD meets the requirements according to the JESD218 standard. Endurance rating verification is defined to establish UBER <1E-16 at 60% upper confidence limit.
4. If the drive is being used in VSS Mode (LBA Format with Metadata size > 0 Bytes) the Drive Endurance will be reduced by

~30%.

Table 63: Drive Shutdown (RTD3 Entry) and Time To Ready (TTR) - U.2 15mm

Form Factor	Capacity	Safe Shutdown Time or RTD3 Entry (seconds)	Time to Ready (TTR ¹) after Surprise Shutdown (seconds)	Time to Ready (TTR) after Planned Shutdown or RTD3 Resume (seconds)
P5520 - U.2 15mm	1.92	5s	17s	3s
	3.84	5s	19s	4s
	7.68	5s	20s	10s
	15.36	5s	30s	15s
P5620 - U.2 15mm	1.6	5s	17s	3s
	3.2	5s	19s	4s
	6.4	5s	20s	10s
	12.8	5s	30s	15s

Note:

1. Time to Ready (TTR) is based on time from power on to when the drive can begin receiving PCIe commands from host after a single #PERST, when CSTS.RDY bit toggles from '0' to '1'.

Table 64: Drive Shutdown (RTD3 Entry) and Time To Ready (TTR) - EDSFF

Form Factor	Capacity	Safe Shutdown Time or RTD3 Entry (seconds)	Time to Ready (TTR ¹) after Surprise Shutdown (seconds)	Time to Ready (TTR) after Planned Shutdown or RTD3 Resume (seconds)
P5520 - E1.S 9.5mm & 15mm	1.92	5s	17s	3s
	3.84	5s	19s	4s
	7.68	5s	20s	10s
P5520 - E1.L 9.5mm	15.36	5s	30s	15s

Note:

1. Time to Ready (TTR) is based on time from power on to when the drive can begin receiving PCIe commands from host after a single #PERST, when CSTS.RDY bit toggles from '0' to '1'.

Table 65: Format NVM Secure Erase and Sanitize Completion Times - U.2 15mm

Form Factor	Capacity	Format NVM, No Secure Erase (SES=0) (seconds)	Format NVM, Cryptographic Erase ¹ (SES=1) (seconds)	Format NVM, Cryptographic Erase (SES=2) (seconds)	Sanitize, Block Erase (Command DWord10 - Bits [02:00] = 010b) (minutes)	Sanitize, Cryptographic Erase (Command DWord10 - Bits [00:02] = 100b) (seconds)
P5520 - U.2 15mm	1.92TB	<=10s	<=10s	<=10s	<=5min	<=30s
	3.84TB	<=10s	<=10s	<=10s	<=5min	<=30s
	7.68TB	<=10s	<=10s	<=10s	<=5min	<=30s
	15.36TB	<=10s	<=10s	<=10s	<=5min	<=30s
P5620 - U.2 15mm	1.6TB	<=10s	<=10s	<=10s	<=5min	<=30s
	3.2TB	<=10s	<=10s	<=10s	<=5min	<=30s
	6.4TB	<=10s	<=10s	<=10s	<=5min	<=30s
	12.8TB	<=10s	<=10s	<=10s	<=5min	<=30s

Note:

- These User Data Erase completion times assume the command is issued across all namespaces on the drive, or that the drive is configured with a single namespace. On a drive configured with multiple namespaces, the User Data Erase may take significantly longer if only some of the namespaces are targeted by the command instead of the full drive. It is recommended that Cryptographic Erase (SES=2) is used instead when issuing a Secure Erase to a single namespace in a multiple namespace environment.

Table 66: Format NVM Secure Erase and Sanitize Completion Times - EDSFF

Form Factor	Capacity	Format NVM, No Secure Erase (SES=0) (seconds)	Format NVM, Cryptographic Erase ¹ (SES=1) (seconds)	Format NVM, Cryptographic Erase (SES=2) (seconds)	Sanitize, Block Erase (Command DWord10 - Bits [02:00] = 010b) (minutes)	Sanitize, Cryptographic Erase (Command DWord10 - Bits [00:02] = 100b) (seconds)
P5520 - E1.S 9.5mm & 15mm	1.92TB	<=10s	<=10s	<=10s	<=5min	<=30s
	3.84TB	<=10s	<=10s	<=10s	<=5min	<=30s
	7.68TB	<=10s	<=10s	<=10s	<=5min	<=30s
P5520 - E1.L 9.5mm	15.36TB	<=10s	<=10s	<=10s	<=5min	<=30s

Note:

1. These User Data Erase completion times assume the command is issued across all namespaces on the drive, or that the drive is configured with a single namespace. On a drive configured with multiple namespaces, the User Data Erase may take significantly longer if only some of the namespaces are targeted by the command instead of the full drive. It is recommended that Cryptographic Erase (SES=2) is used instead when issuing a Secure Erase to a single namespace in a multiple namespace environment.

Table 67: Dataset Management Deallocate (TRIM) Command Completion Time - U.2 15mm

Form Factor	Capacity	Dataset Management Deallocate completion (milliseconds)
P5520 - U.2 15mm	1.92TB	<100ms
	3.84TB	<100ms
	7.68TB	<100ms
	15.36TB	<100ms
P5620 - U.2 15mm	1.6TB	<100ms
	3.2TB	<100ms
	6.4TB	<100ms
	12.8TB	<100ms

Table 68: Dataset Management Deallocate (TRIM) Command Completion Time - EDSFF

Form Factor	Capacity	Dataset Management Deallocate completion (milliseconds)
P5520 - E1.S 9.5mm & 15mm	1.92TB	<100ms
	3.84TB	<100ms
	7.68TB	<100ms
P5520 - E1.L 9.5mm	15.36TB	<100ms

Table 69: Opal Revert¹ Method Completion Time - U.2 15mm

Form Factor	Capacity	Opal Revert completion (seconds)
P5520 - U.2 15mm	1.92TB	<=8s
	3.84TB	<=8s
	7.68TB	<=8s
	15.36TB	<=8s

Table 69: Opal Revert¹ Method Completion Time - U.2 15mm

Form Factor	Capacity	Opal Revert completion (seconds)
P5620 - U.2 15mm	1.6TB	<=8s
	3.2TB	<=8s
	6.4TB	<=8s
	12.8TB	<=8s

Note:

1. Revert method only supported on Opal enabled SKUs

Table 70: Opal Revert¹ Method Completion Time - EDSFF

Form Factor	Capacity	Opal Revert completion (seconds)
P5520 - E1.S 9.5mm & 15mm	1.92TB	<=8s
	3.84TB	<=8s
	7.68TB	<=8s
P5520 - E1.L 9.5mm	15.36TB	<=8s

Note:

1. Revert method only supported on Opal enabled SKUs

Appendix

Appendix B: Power Metrics

Table 71: Power Consumption - U.2 15mm

Specification	P5520 - U.2 15mm				P5620 - U.2 15mm			
	1.92TB	3.84TB	7.68TB	15.36TB	1.6TB	3.2TB	6.4TB	12.8TB
Max Average Active Write Power ¹	13W	15W	18W	20W	13W	15W	18W	20W
Max Average Active Read Power ²	13W	15W	18W	20W	13W	15W	18W	20W
Max Burst Power ³	20W	20W	23W	30W	20W	20W	23W	25W
Idle	5W	5W	5W	5.3W	5W	5W	5W	5.3W

Note:

1. The workload equates QD256/128KB Sequential Writes. Average power is measured over a 100ms sample period
2. The workload equates QD256/128KB Sequential Read. Average power is measured over a 100ms sample period
3. The workload equates QD256/128KB Sequential Writes. Burst power is measured over a 500µs sample period.

Table 72: Power Consumption - EDSFF

Specification	P5520 - E1.S 9.5mm & 15mm			P5520 - E1.L 9.5mm
	1.92TB	3.84TB	7.68TB	15.36TB
Max Average Active Write Power ¹	13W	15W	18W	20W
Max Average Active Read Power ²	13W	15W	18W	20W
Max Burst Power ³	20W	20W	23W	30W
Idle	5W	5W	5W	5.3W

Note:

1. The workload equates QD256/128KB Sequential Writes. Average power is measured over a 100ms sample period
2. The workload equates QD256/128KB Sequential Read. Average power is measured over a 100ms sample period
3. The workload equates QD256/128KB Sequential Writes. Burst power is measured over a 500µs sample period.

Appendix

Appendix C: IDENTIFY Data Structure

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
Controller Capabilities and Features				
1:00	M	PCI Vendor ID (VID)	Contains the company vendor identifier that is assigned by the PCI SIG. This is the same value as reported in the ID register in section 2.1.1.	See PCIe ID on page 146
3:02	M	PCI Subsystem Vendor ID (SSVID)	Contains the company vendor identifier that is assigned by the PCI SIG for the subsystem.	See PCIe ID on page 146
23:04	M	Serial Number (SN)	Contains the serial number for the NVM subsystem that is assigned by the vendor as an ASCII string.	<variable>
63:24	M	Model Number (MN)	Contains the model number for the NVM subsystem that is assigned by the vendor as an ASCII string.	<variable>
71:64	M	Firmware Revision (FR)	Contains the currently active firmware revision for the NVM subsystem. This is the same revision information that may be retrieved with the Get Log Page command.	<variable>
72	M	Recommended Arbitration Burst (RAB)	This is the recommended Arbitration Burst size. The value is in commands and is reported as a power of two (2^n). This is the same units as the Arbitration Burst size.	0x00
75:73	M	IEEE OUI Identifier (IEEE)	Contains the Organization Unique Identifier (OUI) for the controller vendor. The OUI shall be a valid IEEE/RAC assigned identifier that may be registered at http://standards.ieee.org/develop/regauth/oui/public.html .	<extracted from WWN>

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
76	O	Controller Multi-Path I/O and Namespace Sharing Capabilities (CMIC)	This field specifies multi-path I/O and namespace sharing capabilities of the controller and NVM subsystem.	0x00
			Bits 7:3 are reserved	
			Bit 2: If set to '1' then the controller is associated with an SR-IOV Virtual Function. If cleared to '0' then the controller is associated with a PCI Function.	0
			Bit 1: If set to '1' then the NVM subsystem may contain two or more controllers. If cleared to '0' then the NVM subsystem contains only a single controller.	0
			Bit 0: If set to '1' then the NVM subsystem may contain two or more physical PCI Express ports. If cleared to '0' then the NVM subsystem contains only a single PCI Express port.	0
77	M	Maximum Data Transfer Size (MDTS)	This field indicates the maximum data transfer size between the host and the controller. The host should not issue a command that exceeds this transfer size. If a command is processed that exceeds the transfer size, then the command is aborted with a status of Invalid Field in Command. The value is in units of the minimum memory page size (CAP.MPSMIN) and is reported as a power of two (2^n). A value of 0h indicates no restrictions on transfer size. The restriction includes metadata if it is interleaved with the logical block data.	0x05
				(128KB)
79:78	M	Controller ID (CNTLID)	Contains the NVM subsystem unique controller identifier associated with the controller.	0x00
83:80	M	Version (VER)	This field contains the value reported in the Version register defined in section 3.1.2. Implementations compliant to revision 1.2 or later of this specification shall report a non-zero value in this field.	0X00 01 04 00
87:84	M	RTD3 Resume Latency (RTD3R)	This field indicates the typical latency in microseconds resuming from Runtime D3 (RTD3). A value of 0h indicates RTD3 Resume Latency is not reported.	0x00989680
91:88	M	RTD3 Entry Latency (RTD3E)	This field indicates the typical latency in microseconds to enter Runtime D3 (RTD3). A value of 0h indicates RTD3 Entry Latency is not reported.	0x004C4B40

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
95:92	M	Optional Asynchronous Events Supported (OAES)		
			This field indicates the optional asynchronous events supported by the controller. A controller shall not send optional asynchronous events before they are enabled by host software.	0x300
			Bit 31:10 is reserved	0
			Bit 9 is set to '1' if the controller supports the Firmware Activation Notices event. If cleared to '0', then the controller does not support the Firmware Activation Notices event.	1
			Bit 8 is set to '1' if the controller supports the Namespace Attribute Notices event and the associated Changed Namespace List log page. If cleared to '0', then the controller does not support the Namespace Attribute Notices event nor the associated Changed Namespace List log page.	1
			Bits 7:0 are reserved.	0
99:96	M	Controller Attributes (CTRATT)	This field indicates attributes of the controller.	0x0
			Bits 31:2 are reserved.	0
			Bit 1 (Non-Operational Power State Permissive Mode): If set to '1' then the controller supports host control of whether the controller may temporarily exceed the power of a non-operational power state for the purpose of executing controller initiated background operations in a non-operational power state (i.e., Non-Operational Power State Permissive Mode supported). If cleared to '0' then the controller does not support host control of whether the controller may exceed the power of a non-operational state for the purpose of executing controller initiated background operations in a non-operational state (i.e., Non-Operational Power State Permissive Mode not supported). Refer to section 5.22.1.17.	0
			Bit 0 if set to '1' then the controller supports a 128-bit Host Identifier. Bit 0 if cleared to '0' then the controller does not support a 128-bit Host Identifier.	0
110:100			Reserved	

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value								
111	M	Controller Type (CNTRL-TYPE)	<p>This field specifies the controller type. A value of 0h indicates that the controller type is not reported. Implementations compliant to version 1.4 or later of this specification shall report a controller type (i.e., the value 0h is reserved and shall not be used). Implementations compliant to an earlier specification version may report a value of 0h to indicate that a controller type is not reported.</p> <p>Value Controller Type 0h Reserved (controller type not reported) 1h I/O Controller 2h Discovery Controller 3h Administrative Controller 4h to FFh Reserved</p>	1h								
127:112	O	FRU Globally Unique Identifier (FGUID)	<p>This field contains a 128-bit value that is globally unique for a given Field Replaceable Unit (FRU). Refer to the NVM Express Management Interface (NVMe-MITM) specification for the definition of a FRU. This field remains fixed throughout the life of the FRU. This field shall contain the same value for each controller associated with a given FRU.</p> <p>This field uses the EUI-64 based 16-byte designator format. Bytes 122:120 contain the 24-bit Organizationally Unique Identifier (OUI) value assigned by the IEEE Registration Authority. Bytes 127:123 contain an extension identifier assigned by the corresponding organization. Bytes 119:112 contain the vendor specific extension identifier assigned by the corresponding organization. See the IEEE EUI-64 guidelines for more information. This field is big endian (refer to section 7.10).</p> <p>When not implemented, this field contains a value of 0h.</p>	Varies								
254:128			Reserved									
255	M	Management Endpoint Capabilities (MEC)	<p>This field indicates the capabilities of the Management Endpoint in the Controller.</p> <table><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.	0x1
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.											
Admin Command Set Attributes												

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
257:256	M	Optional Admin Command Support (OACS):	This field indicates the optional Admin commands supported by the controller.	0x001E
			Bits 15:9 are reserved.	
			Bit 8 if set to '1' then the controller supports the Doorbell Buffer Config command. If cleared to '0' then the controller does not support the Doorbell Buffer Config command.	0
			Bit 7 if set to '1' then the controller supports the Virtualization Management command. If cleared to '0' then the controller does not support the Virtualization Management command.	0
			Bit 6 if set to '1' then the controller supports the NVMe-MI Send and NVMe-MI Receive commands. If cleared to '0' then the controller does not support the NVMe-MI Send and NVMe-MI Receive commands.	0
			Bit 5 if set to '1' then the controller supports Directives. If cleared to '0' then the controller does not support Directives. A controller that supports Directives shall support the Directive Send and Directive Receive commands.	0
			Bit 4 if set to '1' then the controller supports the Device Self-test command. If cleared to '0' then the controller does not support the Device Self-test command.	1
			Bit 3 if set to '1' then the controller supports the Namespace Management and Namespace Attachment commands. If cleared to '0' then the controller does not support the Namespace Management and Namespace Attachment commands.	1
			Bit 2 if set to '1' then the controller supports the Firmware Activate and Firmware Download commands. If cleared to '0' then the controller does not support the Firmware Activate and Firmware Download commands.	1
			Bit 1 if set to '1' then the controller supports the Format NVM command. If cleared to '0' then the controller does not support the Format NVM command.	1
			Bit 0 if set to '1' then the controller supports the Security Send and Security Receive commands. If cleared to '0' then the controller does not support the Security Send and Security Receive commands.	1 - Opal/Ruby Enabled SKU 0 - Opal/Ruby Disabled SKU

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
258	M	Abort Command Limit (ACL)	This field is used to convey the maximum number of concurrently outstanding Abort commands supported by the controller. This is a 0's based value. It is recommended that implementations support a minimum of four Abort commands outstanding simultaneously.	0x3
				4
259	M	Asynchronous Event Request Limit (AERL)	This field is used to convey the maximum number of concurrently outstanding Asynchronous Event Request commands supported by the controller. This is a 0's based value. It is recommended that implementations support a minimum of four Asynchronous Event Request Limit commands outstanding simultaneously.	0x03
				4
260	M	Firmware Updates (FRMW)	This field indicates capabilities regarding firmware updates.	0x18
			Bits 7:5 are reserved.	
			Bit 4 if set to '1' indicates that the controller supports firmware activation without a reset. If cleared to '0' then the controller requires a reset for firmware to be activated.	0
			Bits 3:1 indicate the number of firmware slots that the device supports. This field shall specify a value between one and seven, indicating that at least one firmware slot is supported and up to seven maximum. This corresponds to firmware slots 1 through 7.	4
			Bit 0 if set to '1' indicates that the first firmware slot (slot 1) is read only. If cleared to '0' then the first firmware slot (slot 1) is read/write. Implementations may choose to have a baseline read only firmware image.	0

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
261	M	Log Page Attributes (LPA)	This field indicates optional attributes for log pages that are accessed via the Get Log Page command.	0x1E
			Bits 7:4 are reserved.	
			Bit 3 if set to '1' then the controller supports the Telemetry Host-Initiated and Telemetry Controller-Initiated log pages and sending Telemetry Log Notices. If cleared to '0' then the controller does not support the Telemetry Host-Initiated and Telemetry Controller-Initiated log pages and Telemetry Log Notice events.	1
			Bit 2 if set to '1' then the controller supports extended data for Get Log Page (including extended Number of Dwords and Log Page Offset fields). Bit 2 if cleared to '0' then the controller does not support extended data for Get Log Page.	1
			Bit 1 if set to '1' then the controller supports the Commands Supported and Effects log page. Bit 1 if cleared to '0' then the controller does not support the Commands Supported and Effects log page.	1
			Bit 0 if set to '1' then the controller supports the SMART / Health information log page on a per namespace basis. If cleared to '0' then the controller does not support the SMART / Health information log page on a per namespace basis.	0
262	M	Error Log Page Entries (ELPE)	This field indicates the maximum number of Error Information log entries that are stored by the controller. This field is a 0's based value.	0xFF
				63
263	M	Number of Power States Support (NPSS)	This field indicates the number of NVM Express power states supported by the controller. This is a 0's based value.	0x02
			Power states are numbered sequentially starting at power state 0. A controller shall support at least one power state (i.e., power state 0) and may support up to 31 additional power states (i.e., up to 32 total).	
264	M	Admin Vendor Specific Command Configuration (AVSCC)	This field indicates the configuration settings for Admin Vendor Specific command handling.	0x00
			Bits 7:1 are reserved.	
			Bit 0 if set to '1' indicates that all Admin Vendor Specific Commands use the format defined in Figure 8 of NVMe Spec 1.2. If cleared to '0' indicates that the format of all Admin Vendor Specific Commands are vendor specific.	

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
265	O	Autonomous Power State Transition Attributes (APSTA)	This field indicates the attributes of the autonomous power state transition feature.	0x00
			Bits 7:1 are reserved.	
			Bit 0 if set to '1' then the controller supports autonomous power state transitions. If cleared to '0' then the controller does not support autonomous power state transitions.	
267:266	M	Warning Composite Temperature Threshold (WCTEMP)	This field indicates the minimum Composite Temperature field value (reported in the SMART / Health Information log that indicates an overheating condition during which controller operation continues. Immediate remediation is recommended (e.g., additional cooling or workload reduction). The platform should strive to maintain a composite temperature below this value.	0x0157
			A value of 0h in this field indicates that no warning temperature threshold value is reported by the controller. Implementations compliant to revision 1.2 or later of this specification shall report a non-zero value in this field.	
			It is recommended that implementations report a value of 0157h in this field.	
269:268	M	Critical Composite Temperature Threshold (CCTEMP)	This field indicates the minimum Composite Temperature field value (reported in the SMART / Health Information log) that indicates a critical overheating condition (e.g., may prevent continued normal operation, possibility of data loss, automatic device shutdown, extreme performance throttling, or permanent damage).	0x0161
			A value of 0h in this field indicates that no critical temperature threshold value is reported by the controller. Implementations compliant to revision 1.2 or later of this specification shall report a non-zero value in this field.	
271:270	O	Maximum Time for Firmware Activation (MTFA)	Indicates the maximum time the controller temporarily stops processing commands to activate the firmware image. This field shall be valid if the controller supports firmware activation without a reset. This field is specified in 100 millisecond units. A value of 0h indicates that the maximum time is undefined.	0x32
275:272	O	Host Memory Buffer Preferred Size (HMPRE)	This field indicates the preferred size that the host is requested to allocate for the Host Memory Buffer feature in 4KB units. This value shall be larger than or equal to the Host Memory Buffer Minimum Size. If this field is non-zero, then the Host Memory Buffer feature is supported. If this field is cleared to 0h, then the Host Memory Buffer feature is not supported.	0x00

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
279:276	○	Host Memory Buffer Minimum Size (HMMIN)	This field indicates the minimum size that the host is requested to allocate for the Host Memory Buffer feature in 4KB units. If this field is cleared to 0h, then the host is requested to allocate any amount of host memory possible up to the HMPRE value.	0x00
295:280	○	Total NVM Capacity (TNVMCAP)	This field indicates the total NVM capacity in the NVM subsystem. The value is in bytes. This field shall be supported if Namespace Management and Namespace Attachment commands are supported.	varies
311:296	○	Unallocated NVM Capacity (UNVMCAP)	This field indicates the unallocated NVM capacity in the NVM subsystem. The value is in bytes. This field shall be supported if Namespace Management and Namespace Attachment commands are supported.	varies

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value																
315:312	O	Replay Protected Memory Block Support (RPMBS)	This field indicates if the controller supports one or more Replay Protected Memory Blocks (RPMBs) and the capabilities.	0x00																
			<table><tr><th>Value</th><th>Definition</th></tr><tr><td>31:24</td><td>Access Size: If the Number of RPMB Units field is non-zero, then this field indicates the maximum number of 512B units of data that may be read or written per RPMB access by Security Send or Security Receive commands for the controller. This is a 0's based value. A value of 0h indicates support for one unit of 512B of data. If the Number of RPMB Units field is 0h, then this field shall be ignored.</td></tr><tr><td>23:16</td><td>Total Size: If the Number of RPMB Units field is non-zero, then this field indicates the number of 128 KiB units of data in each RPMB supported in the controller. This is a 0's based value. A value of 0h indicates support for one unit of 128 KiB of data. If the Number of RPMB Units field is 0h, this field shall be ignored.</td></tr><tr><td>15:06</td><td>Reserved</td></tr><tr><td>05:03</td><td>Authentication Method: This field indicates the authentication method used to access all RPMBs in the controller. The values for this field are:<table><tr><th>Value</th><th>Definition</th></tr><tr><td>000b</td><td>HMAC SHA-256 (refer to RFC 6234)</td></tr><tr><td>001b to 111b</td><td>Reserved</td></tr></table></td></tr></table>		Value	Definition	31:24	Access Size: If the Number of RPMB Units field is non-zero, then this field indicates the maximum number of 512B units of data that may be read or written per RPMB access by Security Send or Security Receive commands for the controller. This is a 0's based value. A value of 0h indicates support for one unit of 512B of data. If the Number of RPMB Units field is 0h, then this field shall be ignored.	23:16	Total Size: If the Number of RPMB Units field is non-zero, then this field indicates the number of 128 KiB units of data in each RPMB supported in the controller. This is a 0's based value. A value of 0h indicates support for one unit of 128 KiB of data. If the Number of RPMB Units field is 0h, this field shall be ignored.	15:06	Reserved	05:03	Authentication Method: This field indicates the authentication method used to access all RPMBs in the controller. The values for this field are: <table><tr><th>Value</th><th>Definition</th></tr><tr><td>000b</td><td>HMAC SHA-256 (refer to RFC 6234)</td></tr><tr><td>001b to 111b</td><td>Reserved</td></tr></table>	Value	Definition	000b	HMAC SHA-256 (refer to RFC 6234)	001b to 111b	Reserved
			Value		Definition															
			31:24		Access Size: If the Number of RPMB Units field is non-zero, then this field indicates the maximum number of 512B units of data that may be read or written per RPMB access by Security Send or Security Receive commands for the controller. This is a 0's based value. A value of 0h indicates support for one unit of 512B of data. If the Number of RPMB Units field is 0h, then this field shall be ignored.															
			23:16		Total Size: If the Number of RPMB Units field is non-zero, then this field indicates the number of 128 KiB units of data in each RPMB supported in the controller. This is a 0's based value. A value of 0h indicates support for one unit of 128 KiB of data. If the Number of RPMB Units field is 0h, this field shall be ignored.															
15:06	Reserved																			
05:03	Authentication Method: This field indicates the authentication method used to access all RPMBs in the controller. The values for this field are: <table><tr><th>Value</th><th>Definition</th></tr><tr><td>000b</td><td>HMAC SHA-256 (refer to RFC 6234)</td></tr><tr><td>001b to 111b</td><td>Reserved</td></tr></table>	Value	Definition	000b	HMAC SHA-256 (refer to RFC 6234)	001b to 111b	Reserved													
Value	Definition																			
000b	HMAC SHA-256 (refer to RFC 6234)																			
001b to 111b	Reserved																			

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value				
			<table><tr><th>Value</th><th>Definition</th></tr><tr><td>02:00</td><td>Number of RPMB Units: This field indicates the number of RPMB targets the controller supports. All RPMB targets supported shall have the same capabilities as defined in the RPMBs field. A value of 0h indicates the controller does not support Replay Protected Memory Blocks. If this value is non-zero, then the controller shall support the Security Send and Security Receive commands.</td></tr></table>	Value	Definition	02:00	Number of RPMB Units: This field indicates the number of RPMB targets the controller supports. All RPMB targets supported shall have the same capabilities as defined in the RPMBs field. A value of 0h indicates the controller does not support Replay Protected Memory Blocks. If this value is non-zero, then the controller shall support the Security Send and Security Receive commands.	
Value	Definition							
02:00	Number of RPMB Units: This field indicates the number of RPMB targets the controller supports. All RPMB targets supported shall have the same capabilities as defined in the RPMBs field. A value of 0h indicates the controller does not support Replay Protected Memory Blocks. If this value is non-zero, then the controller shall support the Security Send and Security Receive commands.							
317:316	O	Extended Device Self-test Time (EDSTT)	If the Device Self-test command is supported, then this field indicates the nominal amount of time in one minute units that the controller takes to complete an extended device self-test operation when in power state 0. If the Device Self-test command is not supported, then this field is reserved.	0x78				
318	O	Device Self-test Options (DSTO)	This field indicates the optional Device Self-test command or operation behaviors supported by the controller or NVM subsystem. Bit 0 if set to '1' then the NVM subsystem supports only one device self-test operation in progress at a time. If cleared to '0' then the NVM subsystem supports one device self-test operation per controller at a time.	0x1				
319	M	Firmware Update Granularity (FWUG)	This field indicates the granularity and alignment requirement of the firmware image being updated by the Firmware Image Download command (refer to section 5.12). If the values specified in the NUMD field or the OFST field in the Firmware Image Download command do not conform to this granularity and alignment requirement, then the firmware update may fail with status of Invalid Field in Command. For the broadest interoperability with host software, it is recommended that the controller set this value to the lowest value possible. The value is reported in 4KB units (e.g., 1h corresponds to 4KB, 2h corresponds to 8KB). A value of 0h indicates that no information on granularity is provided. A value of FFh indicates there is no restriction (i.e., any granularity and alignment in Dwords is allowed).	0x1				
321:320	O	Keep Alive Support (KAS)	This field indicates the granularity of the Keep Alive Timer in 100 ms units (refer to section 7.12). If this field is cleared to 0h then Keep Alive is not supported. Keep Alive shall be supported for NVMe over Fabrics implementations.	0x0				

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
323:322	M	Host Controlled Thermal Management Attributes (HCTMA)	This field indicates the attributes of the host controlled thermal management feature. Refer to section 8.4.5. Bits 15:1 are reserved. Bit 0 if set to '1' then the controller supports host controlled thermal management. If cleared to '0' then the controller does not support host controlled thermal management. If this bit is set to '1' then, the controller shall support the Set Features command and Get Features command with the Feature Identifier field set to 10h.	0x0
325:324	O	Minimum Thermal Management Temperature (MNTMT)	This field indicates the minimum temperature, in degrees Kelvin, that the host may request in the Thermal Management Temperature 1 field and Thermal Management Temperature 2 field of a Set Features command with the Feature Identifier field set to 10h. A value of 0000h indicates that the controller does not report this field or the host controlled thermal management feature (refer to section 8.4.5) is not supported.	0x0
327:326	O	Maximum Thermal Management Temperature (MXTMT)	This field indicates the maximum temperature, in degrees Kelvin, that the host may request in the Thermal Management Temperature 1 field and Thermal Management Temperature 2 field of the Set Features command with the Feature Identifier set to 10h. A value of 0000h indicates that the controller does not report this field or the host controlled thermal management feature is not supported.	0x0

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value										
331:328	O	Sanitize Capabilities (SANI-CAP)	This field indicates attributes for sanitize operations. If the Sanitize command is supported then this field shall be non-zero. If the Sanitize command is not supported, then this field shall be cleared to 0h. Bits 31:30 are reserved.	0x60000003										
			<table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Additional media modification after sanitize operation completes successfully is not defined. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return this value.</td></tr><tr><td>01b</td><td>Media is not additionally modified by the NVMe controller after sanitize operation completes successfully</td></tr><tr><td>10b</td><td>Media is additionally modified by the NVMe controller after sanitize operation completes successfully. The Sanitize Operation Completed event does not occur until the additional media modification associated with this field has completed.</td></tr><tr><td>11b</td><td>Reserved</td></tr></table>		Value	Definition	00b	Additional media modification after sanitize operation completes successfully is not defined. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return this value.	01b	Media is not additionally modified by the NVMe controller after sanitize operation completes successfully	10b	Media is additionally modified by the NVMe controller after sanitize operation completes successfully. The Sanitize Operation Completed event does not occur until the additional media modification associated with this field has completed.	11b	Reserved
			Value		Definition									
			00b		Additional media modification after sanitize operation completes successfully is not defined. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return this value.									
			01b		Media is not additionally modified by the NVMe controller after sanitize operation completes successfully									
			10b		Media is additionally modified by the NVMe controller after sanitize operation completes successfully. The Sanitize Operation Completed event does not occur until the additional media modification associated with this field has completed.									
11b	Reserved													

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
			<p>Bit 29: No-Deallocate Inhibited (NDI): If set to '1' and the NoDeallocate Response Mode bit is set to '1', then the controller deallocates after the sanitize operation even if the No- Deallocate After Sanitize bit is set to '1' in a Sanitize command. If:</p> <p>a) set to '1';</p> <p>b) the No-Deallocate Response Mode bit (refer to Figure 314) is cleared to '0'; and</p> <p>c) the No-Deallocate After Sanitize bit is set to '1' in a Sanitize command, then the controller aborts the Sanitize command with a status of Invalid Field in Command. If cleared to '0', then the controller supports the No-Deallocate After Sanitize bit in a Sanitize command. If bits 2:0 of the SANICAP field are cleared to 0h, then the controller shall clear this bit to '0'</p> <p>Bits 28:03 are reserved.</p> <p>Bit 2 if set to '1' then the controller supports the Overwrite sanitize operation. If cleared to '0' then the controller does not support the Overwrite sanitize operation.</p> <p>Bit 1 if set to '1' then the controller supports the Block Erase sanitize operation. If cleared to '0' then the controller does not support the Block Erase sanitize operation.</p> <p>Bit 0 if set to '1' then the controller supports the Crypto Erase sanitize operation. If cleared to '0' then the controller does not support the Crypto Erase sanitize operation.</p>	
335:332	O	Host Memory Buffer Minimum Descriptor Entry Size (HMMINDS)	This field indicates the minimum usable size of a Host Memory Buffer Descriptor Entry in 4 KiB units. If this field is cleared to 0h, then the controller does not indicate any limitations on the Host Memory Buffer Descriptor Entry size.	0x0
337:336	O	Host Memory Maximum Descriptors Entries (HMMAXD)	This field indicates the number of usable Host Memory Buffer Descriptor Entries. If this field is cleared to 0h, then the controller does not indicate a maximum number of Host Memory Buffer Descriptor Entries.	0x0
339:338	O	NVM Set Identifier Maximum (NSETIDMAX)	This field defines the maximum value of a valid NVM Set Identifier for any controller in the NVM subsystem. The number of NVM Sets supported by the NVM subsystem is less than or equal to NSETIDMAX.	0x0
341:340	O	Endurance Group Identifier Maximum (ENDGIDMAX)	This field defines the maximum value of a valid Endurance Group Identifier for any controller in the NVM subsystem. The number of Endurance Groups supported by the NVM subsystem is less than or equal to ENDGIDMAX.	0x0

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
342	O	ANA Transition Time (ANATT)	This field indicates the maximum amount of time, in seconds, for a transition between ANA states or the maximum amount of time, in seconds, that the controller reports the ANA change state. If the controller supports Asymmetric Namespace Access Reporting (refer to the CMIC field), then this field shall be set to a non-zero value. If the controller does not support Asymmetric Namespace Access Reporting, then this field shall be cleared to 0h. Refer to section 8.21.4.	0x0
343	O	Asymmetric Namespace Access Capabilities (ANACAP)	<p>This field indicates the capabilities associated with Asymmetric Namespace Access Reporting (refer to section 8.20).</p> <p>Bit 7 if set to '1', then the controller supports a non-zero value in the ANAGRPID field of the Namespace Management command. If cleared to '0', then the controller does not support a non-zero value in the ANAGRPID field of the Namespace Management command. If the Namespace Management command is not supported, then this bit shall be cleared to '0'.</p> <p>Bit 6 if set to '1', then the ANAGRPID field in the Identify Namespace data structure (refer to Figure 245) does not change while the namespace is attached to any controller. If cleared to '0', then the ANAGRPID field may change while the namespace is attached to any controller. Refer to section 8.20.2.</p> <p>Bit 5 is reserved.</p> <p>Bit 4 if set to '1', then the controller is able to report ANA Change state (refer to section 8.20.3.5). If cleared to '0', then the controller does not report ANA Change state.</p> <p>Bit 3 if set to '1', then the controller is able to report ANA Persistent Loss state (refer to section 8.20.3.4). If cleared to '0', then the controller does not report ANA Persistent Loss state.</p> <p>Bit 2 if set to '1', then the controller is able to report ANA Inaccessible state (refer to section 8.20.3.3). If cleared to '0', then the controller does not report ANA Inaccessible state.</p> <p>Bit 1 if set to '1', then the controller is able to report ANA Non-Optimized state (refer to section 8.20.3.2). If cleared to '0', then the controller does not report ANA NonOptimized state.</p> <p>Bit 0 if set to '1', then the controller is able to report ANA Optimized state (refer to section 8.20.3.1). If the controller supports Asymmetric Namespace Access Reporting, then this bit is set to '1'</p>	0x0
347:344	O	Number of ANA Group Identifiers (NANAGRPID)	This field indicates the maximum value of a valid ANA Group Identifier for any controller in the NVM subsystem. If the controller supports Asymmetric Namespace Access Reporting (refer to the CMIC field), then this field shall be set to a non-zero value. If the controller does not support Asymmetric Namespace Access Reporting, then this field shall be cleared to 0h	0x0

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
351:348	O	Persistent Event Log Size (PELS)	This field indicates the maximum reportable size for the Persistent Event Log (Refer to section 5.14.1.13) in 64 KiB units. If the Persistent Event Log is not supported, then this field is reserved	0x50
511:332			Reserved	
NVM Command Set Attributes				
512	M	Submission Queue Entry Size (SQES)	This field defines the required and maximum Submission Queue entry size when using the NVM Command Set.	0x66
			Bits 7:4 define the maximum Submission Queue entry size when using the NVM Command Set. This value is larger than or equal to the required SQ entry size. The value is in bytes and is reported as a power of two (2^n). The recommended value is 6, corresponding to a standard NVM Command Set SQ entry size of 64 bytes. Controllers that implement proprietary extensions may support a larger value.	
			Bits 3:0 define the required Submission Queue Entry size when using the NVM Command Set. This is the minimum entry size that may be used. The value is in bytes and is reported as a power of two (2^n). The required value shall be 6, corresponding to 64.	
513	M	Completion Queue Entry Size (CQES)	This field defines the required and maximum Completion Queue entry size when using the NVM Command Set.	0x44
			Bits 7:4 define the maximum Completion Queue entry size when using the NVM Command Set. This value is larger than or equal to the required CQ entry size. The value is in bytes and is reported as a power of two (2^n). The recommended value is 4, corresponding to a standard NVM Command Set CQ entry size of 16 bytes. Controllers that implement proprietary extensions may support a larger value.	
			Bits 3:0 define the required Completion Queue entry size when using the NVM Command Set. This is the minimum entry size that may be used. The value is in bytes and is reported as a power of two (2^n). The required value shall be 4, corresponding to 16.	
515:514	M	Maximum Outstanding Commands (MAX-CMD)	Indicates the maximum number of commands that the controller processes at one time for a particular queue (which may be larger than the size of the corresponding Submission Queue). The host may use this value to size Completion Queues and optimize the number of commands submitted at one time to a particular I/O Queue. This field is mandatory for NVMe over Fabrics and optional for NVMe over PCIe implementations. If the field is not used, it shall be cleared to 0x0.	0x0

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
519:516	M	Number of Namespaces (NN)	This field defines the number of valid namespaces present for the controller. Namespaces shall be allocated in order (starting with 1) and packed sequentially.	0x80
				(128 decimal)
521:520	M	Optional NVM Command Support (ONCS)	This field indicates the optional NVM commands supported by the controller.	0x5E
			Bits 15:7 are reserved.	
			Bit 6 if set to '1' then the controller supports the Timestamp feature. If cleared to '0', then the controller does not support the Timestamp feature.	0
			Bit 5 if set to '1' then the controller supports reservations. If cleared to '0' then the controller does not support reservations. If the controller supports reservations, then it shall support the following commands associated with reservations: Reservation Report, Reservation Register, Reservation Acquire, and Reservation Release.	0
			Bit 4 if set to '1' then the controller supports the Save field in the Set Features command and the Select field in the Get Features command. If cleared to '0' then the controller does not support the Save field in the Set Features command and the Select field in the Get Features command.	1
			Bit 3 if set to '1' then the controller supports the Write Zeroes command. If cleared to '0' then the controller does not support the Write Zeroes command.	1
			Bit 2 if set to '1' then the controller supports the Dataset Management command.	1
			If cleared to '0' then the controller does not support the Dataset Management command.	
			Bit 1 if set to '1' then the controller supports the Write Uncorrectable command.	1
			If cleared to '0' then the controller does not support the Write Uncorrectable command.	
			Bit 0 if set to '1' then the controller supports the Compare command. If cleared to '0' then the controller does not support the Compare command.	0

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
523:522	M	Fused Operation Support (FUSES)	This field indicates the fused operations that the controller supports.	0x00
			Bits 15:1 are reserved.	
			Bit 0 if set to '1' then the controller supports the Compare and Write fused operation. If cleared to '0' then the controller does not support the Compare and Write fused operation. Compare shall be the first command in the sequence.	
524	M	Format NVM Attributes (FNA)	This field indicates attributes for the Format NVM command.	0x04
			Bits 7:3 are reserved.	
			Bit 2 indicates whether cryptographic erase is supported as part of the secure erase functionality. If set to '1', then cryptographic erase is supported. If cleared to '0', then cryptographic erase is not supported.	1
			Bit 1 indicates whether secure erase functionality applies to all namespaces or is specific to a particular namespace. If set to '1', then a secure erase of a particular namespace as part of a format results in a secure erase of all namespaces. If cleared to '0', then a secure erase as part of a format is performed on a per namespace basis.	0
			Bit 0 indicates whether the format operation applies to all namespaces or is specific to a particular namespace. If set to '1', then all namespaces shall be configured with the same attributes and a format of any namespace results in a format of all namespaces. If cleared to '0', then the controller supports format on a per namespace basis.	0

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
525	M	Volatile Write Cache (VWC)	This field indicates attributes related to the presence of a volatile write cache in the implementation.	0x6
			Bits 7:1 are reserved.	
			Bits 2:1 indicate Flush command behavior (refer to section 6.8) if the NSID value is set to FFFFFFFFh as follows: Value 00b; Definition: Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return this value. Value: 01b; Definition: Reserved. Value: 10b; Definition: The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format. Value: 11b; Definition: The Flush command supports the NSID field set to FFFFFFFFh	11
			Bit 0 if set to '1' indicates that a volatile write cache is present. If cleared to '0', a volatile write cache is not present. If a volatile write cache is present, then the host may issue Flush commands and control whether it is enabled with Set Features specifying the Volatile Write Cache feature identifier. If a volatile write cache is not present, the host shall not issue Flush commands nor Set Features or Get Features with the Volatile Write Cache identifier.	
527:526	M	Atomic Write Unit Normal (AWUN)	This field indicates the atomic write size for the controller during normal operation. This field is specified in logical blocks and is a 0's based value. If a write is issued of this size or less, the host is guaranteed that the write is atomic to the NVM with respect to other read or write operations. A value of FFFFh indicates all commands are atomic as this is the largest command size. It is recommended that implementations support a minimum of 128KB (appropriately scaled based on LBA size).	0x00
529:528	M	Atomic Write Unit Power Fail (AWUPF)	This field indicates the atomic write size for the controller during a power fail condition. This field is specified in logical blocks and is a 0's based value. If a write is issued of this size or less, the host is guaranteed that the write is atomic to the NVM with respect to other read or write operations.	0x00

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
530	M	NVM Vendor Specific Command Configuration (NVSCC)	This field indicates the configuration settings for NVM Vendor Specific command handling.	0x00
			Bits 7:1 are reserved.	
			Bit 0 if set to '1' indicates that all NVM Vendor Specific Commands use the format defined in Figure 8. If cleared to '0' indicates that the format of all NVM Vendor Specific Commands are vendor specific.	0
531	M	Namespace Write Protection Capabilities (NWPC)	<p>This field indicates the optional namespace write protection capabilities supported by the controller. Refer to section 8.19.</p> <p>Bits 7:3 are reserved.</p> <p>Bit 2 if set to '1', then the controller supports the Permanent Write Protect state. If cleared to '0', then the controller does not support the Permanent Write Protect state. If this bit is set to '1', then the controller shall support the Namespace Write Protection Authentication field (refer to section 8.10).</p> <p>Bit 1 if set to '1', then the controller supports the Write Protect Until Power Cycle state. If cleared to '0', then the controller does not support Write Protect Until Power Cycle state. If this bit is set to '1', then the controller shall support the Namespace Write Protection Authentication field (refer to section 8.10).</p> <p>Bit 0 if set to '1', then the controller shall support the No Write Protect and Write Protect namespace write protection states and may support the Write Protect Until Power Cycle state and Permanent Write Protect namespace write protection states (refer to section 8.19). If cleared to '0', then the controller does not support Namespace Write Protection and bits 2:1 shall be cleared to 00b.</p>	0
533:532	O	Atomic Compare & Write Unit (ACWU)	This field indicates the size of the write operation guaranteed to be written atomically to the NVM across all namespaces with any supported namespace format for a Compare and Write fused operation. If a specific namespace guarantees a larger size than is reported in this field, then this namespace specific size is reported in the NACWU field in the Identify Namespace data structure.	0x00
			This field shall be supported if the Compare and Write fused command is supported.	
			This field is specified in logical blocks and is a 0's based value. If a Compare and Write is submitted that requests a transfer size larger than this value, then the controller may fail the command with a status code of Invalid Field in Command. If Compare and Write is not a supported fused command, then this field shall be 0h.	
535:534	M		Reserved	

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value								
539:536	O	SGL Support (SGLS)	This field indicates if SGLs are supported for the NVM Command Set and the particular SGL types supported.	0xD0002								
			Bit [31:21]: Reserved									
			Bit 20: If set to '1', then the controller supports the Address field in SGL Data Block, SGL Segment, and SGL Last Segment descriptor types specifying an offset. If cleared to '0' then the Address field specifying an offset is not supported.	0								
			Bit 19: If set to '1', then use of a Metadata Pointer (MPTR) that contains an address of an SGL segment containing exactly one SGL Descriptor that is Qword aligned is supported. If cleared to '0', then use of a MPTR containing an SGL Descriptor is not supported.	1								
			Bit 18: If set to '1', then the controller supports commands that contain a data or metadata SGL of a length larger than the amount of data to be transferred. If cleared to '0', then the SGL length shall be equal to the amount of data to be transferred.	1								
			Bit 17: If set to '1', then use of a byte aligned contiguous physical buffer of metadata (the Metadata Pointer field in Figure 12) is supported. If cleared to '0', then use of a byte aligned contiguous physical buffer of metadata is not supported.	0								
			Bit 16: If set to '1', then the SGL Bit Bucket descriptor is supported. If cleared to '0', then the SGL Bit Bucket descriptor is not supported.	0								
			Bit [15:3]: Reserved									
			Bit 2: If set to '1', then the controller supports the Keyed SGL Data Block descriptor. If cleared to '0', then the controller does not support the Keyed SGL Data Block descriptor.	0								
			Bit [1:0]: This field is used to determine the SGL support for the NVM Command Set. Valid values are shown in the table below.	10								
			<table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>SGLs are not supported.</td></tr><tr><td>01b</td><td>SGLs are supported. There is no alignment nor granularity requirement for Data Blocks.</td></tr><tr><td>10b</td><td>SGLs are supported. There is a dword alignment and granularity requirement for Data Blocks (refer to section 4.4).</td></tr><tr><td>11b</td><td>Reserved</td></tr></table>		Value	Definition	00b	SGLs are not supported.	01b	SGLs are supported. There is no alignment nor granularity requirement for Data Blocks.	10b	SGLs are supported. There is a dword alignment and granularity requirement for Data Blocks (refer to section 4.4).
Value	Definition											
00b	SGLs are not supported.											
01b	SGLs are supported. There is no alignment nor granularity requirement for Data Blocks.											
10b	SGLs are supported. There is a dword alignment and granularity requirement for Data Blocks (refer to section 4.4).											
11b	Reserved											

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
453:540	M	Maximum Number of Allowed Namespaces (MNAN)	This field indicates the maximum number of namespaces supported by the NVM subsystem. If this field is cleared to 0h, then the maximum number of namespaces supported by the NVM subsystem is less than or equal to the value in the NN field. If the controller supports Asymmetric Namespace Access Reporting, then this field shall be set to a non-zero value that is less than or equal to the NN value	0x0
767:544	M		Reserved	
1023:768	M	NVM Subsystem NVMe Qualified Name (SUB-NQN)	This field specifies the NVM Subsystem NVMe Qualified Name as a UTF-8 null-terminated string.	Varies
1791:1024			Reserved	
2047:1792	M		Refer to the NVMe over Fabrics Specification.	
Power State Descriptors				

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
2079:2048	M	Power State 0 Descriptor (PSD0)	This field indicates the characteristics of power state 0.	Data Structure in Table 75 on page 101
2111:2080	O	Power State 1 Descriptor (PSD1)	This field indicates the characteristics of power state 1.	
2143:2112	O	Power State 2 Descriptor (PSD2)	This field indicates the characteristics of power state 2.	
2175:2144	O	Power State 3 Descriptor (PSD3)	This field indicates the characteristics of power state 3.	
2207:2176	O	Power State 4 Descriptor (PSD4)	This field indicates the characteristics of power state 4.	
2239:2208	O	Power State 5 Descriptor (PSD5)	This field indicates the characteristics of power state 5.	
2271:2240	O	Power State 6 Descriptor (PSD6)	This field indicates the characteristics of power state 6.	
2303:2272	O	Power State 7 Descriptor (PSD7)	This field indicates the characteristics of power state 7.	
2335:2304	O	Power State 8 Descriptor (PSD8)	This field indicates the characteristics of power state 8.	
2367:2336	O	Power State 9 Descriptor (PSD9)	This field indicates the characteristics of power state 9.	
2399:2368	O	Power State 10 Descriptor (PSD10)	This field indicates the characteristics of power state 10.	
2431:2400	O	Power State 11 Descriptor (PSD11)	This field indicates the characteristics of power state 11.	

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
2463:2432	O	Power State 12 Descriptor (PSD12)	This field indicates the characteristics of power state 12.	
2495:2464	O	Power State 13 Descriptor (PSD13)	This field indicates the characteristics of power state 13.	
2527:2496	O	Power State 14 Descriptor (PSD14)	This field indicates the characteristics of power state 14.	
2559:2528	O	Power State 15 Descriptor (PSD15)	This field indicates the characteristics of power state 15.	
2591:2560	O	Power State 16 Descriptor (PSD16)	This field indicates the characteristics of power state 16.	
2623:2592	O	Power State 17 Descriptor (PSD17)	This field indicates the characteristics of power state 17.	
2655:2624	O	Power State 18 Descriptor (PSD18)	This field indicates the characteristics of power state 18.	
2687:2656	O	Power State 19 Descriptor (PSD19)	This field indicates the characteristics of power state 19.	
2719:2688	O	Power State 20 Descriptor (PSD20)	This field indicates the characteristics of power state 20.	
2751:2720	O	Power State 21 Descriptor (PSD21)	This field indicates the characteristics of power state 21.	
2783:2752	O	Power State 22 Descriptor (PSD22)	This field indicates the characteristics of power state 22.	
2815:2784	O	Power State 23 Descriptor (PSD23)	This field indicates the characteristics of power state 23.	

Table 73: Identify Controller

Bytes	O/M	Title	Description	Expected Value
2847:2816	O	Power State 24 Descriptor (PSD24)	This field indicates the characteristics of power state 24.	
2879:2848	O	Power State 25 Descriptor (PSD25)	This field indicates the characteristics of power state 25.	
2911:2880	O	Power State 26 Descriptor (PSD26)	This field indicates the characteristics of power state 26.	
2943:2912	O	Power State 27 Descriptor (PSD27)	This field indicates the characteristics of power state 27.	
2975:2944	O	Power State 28 Descriptor (PSD28)	This field indicates the characteristics of power state 28.	
3007:2976	O	Power State 29 Descriptor (PSD29)	This field indicates the characteristics of power state 29.	
3039:3008	O	Power State 30 Descriptor (PSD30)	This field indicates the characteristics of power state 30.	
3071:3040	O	Power State 31 Descriptor (PSD31)	This field indicates the characteristics of power state 31.	

Note:

O = Optional. The content of the word is optional

M = Mandatory. The content of the word is mandatory

Table 74: Vendor Specific Data Structure

Vendor Specific			Expected Value
3072	Solidigm	Auto Detected PCIe Refclk source	varies
3073		Reserved.	0x00

Table 74: Vendor Specific Data Structure

Vendor Specific			Expected Value
3074	m Ven- dor Spe- cific	Namespace Grow and Shrink support Bits 7:2 are reserved. Bit 1 if set to '1' indicates the controller supports Namespace Shrink operation. If cleared to '0', there is no support for Namespace Shrink operation. Bit 0 if set to '1' indicates the controller supports Namespace Grow operation. If cleared to '0', there is no support for Namespace Grow operation	0x0
3075		Stripe size: Specifies the size of the stripe, value is read only and cannot be changed by SW. The value is in units of the minimum memory page size (CAP.MPSMIN) and is reported as a power of two (2^n) 0: Driver assisted striping not supported by FW 1: 8 KiB of user data (i.e. 16 512 byte or 2 4096 byte and extended sectors) 2: 16 KiB of user data (i.e. 32 512 byte or 4 4096 byte and extended sectors) 3: 32 KiB of user data (i.e. 64 512 byte or 8 4096 byte and extended sectors) 4: 64 KiB of user data (i.e. 128 512 byte or 16 4096 byte and extended sectors) 5: 128 KiB of user data (i.e. 256 512 byte or 32 4096 byte and extended sectors)	0x05
3095:3076		Standardized Failure mode String	<variable>
3096		Current PCIe Link Speed field (CLS)	Healthy link would report Gen4 0x04
3097		Negotiated Link Width (NLW)	Healthy link would report 4 lanes 0x04
3098		Bit[31:0] Reserved	N/A
3099		Bit[31:0] Reserved	N/A
3107:3100		VS_BootloaderVersion	0x00
3115:3108		Reserved	N/A
4095:3116		Vendor Specific: This range of bytes is allocated for vendor specific usage	Reserved

Table 75: Power State Descriptor Data Structure

Bytes	Description	Expected Value
255:184	Reserved	

Table 75: Power State Descriptor Data Structure

Bytes	Description	Expected Value										
183:182	Active Power Scale (APS): This field indicates the scale for the Active Power field. If an Active Power Workload is reported for a power state, then the Active Power Scale shall also be reported for that power state.	0x00										
	<table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Not reported for this power state</td></tr><tr><td>01b</td><td>0.0001 W</td></tr><tr><td>10b</td><td>0.01 W</td></tr><tr><td>11b</td><td>Reserved</td></tr></table>		Value	Definition	00b	Not reported for this power state	01b	0.0001 W	10b	0.01 W	11b	Reserved
	Value		Definition									
	00b		Not reported for this power state									
	01b		0.0001 W									
	10b		0.01 W									
11b	Reserved											
181:179	Reserved											
178:176	Active Power Workload (APW): This field indicates the workload used to calculate maximum power for this power state. This field shall not be “No Workload” unless ACP is 0000h.	0x00										
175:160	Active Power (ACP): This field indicates the largest average power consumed by the NVM sub-system over a 10 second period in this power state with the workload indicated in the Active Power Workload field. The power in Watts is equal to the value in this field multiplied by the scale indicated in the Active Power Scale field. A value of 0000h indicates Active Power is not reported.	0x00										
159:152	Reserved											
151:150	Idle Power Scale (IPS): This field indicates the scale for the Idle Power field.	0x00										
	<table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Not reported for this power state</td></tr><tr><td>01b</td><td>0.0001 W</td></tr><tr><td>10b</td><td>0.01 W</td></tr><tr><td>11b</td><td>Reserved</td></tr></table>		Value	Definition	00b	Not reported for this power state	01b	0.0001 W	10b	0.01 W	11b	Reserved
	Value		Definition									
	00b		Not reported for this power state									
	01b		0.0001 W									
	10b		0.01 W									
11b	Reserved											
149:144	Reserved											
143:128	Idle Power (IDLP): This field indicates the typical power consumed by the NVM subsystem over 30 seconds in this power state when idle (i.e., there are no pending commands, register accesses, nor background processes). The measurement starts after the NVM subsystem has been idle for 10 seconds. The power in Watts is equal to the value in this field multiplied by the scale indicated in the Idle Power Scale field. A value of 0000h indicates Idle Power is not reported.	0x00										
127:125	Reserved											
124:120	Relative Write Latency (RWL): This field indicates the relative write latency associated with this power state. The value in this field shall be less than the number of supported power states (e.g., if the controller supports 16 power states, then valid values are 0 through 15). A lower value means lower write latency.	0x00										

Table 75: Power State Descriptor Data Structure

Bytes	Description	Expected Value
116:112	Relative Write Throughput (RWT): This field indicates the relative write throughput associated with this power state. The value in this field shall be less than the number of supported power states (e.g., if the controller supports 16 power states, then valid values are 0 through 15). A lower value means higher write throughput.	0x00
111:109	Reserved	
108:104	Relative Read Latency (RRL): This field indicates the relative read latency associated with this power state. The value in this field shall be less than the number of supported power states (e.g., if the controller supports 16 power states, then valid values are 0 through 15). A lower value means lower read latency.	0x00
103:101	Reserved	
100:96	Relative Read Throughput (RRT): This field indicates the relative read throughput associated with this power state. The value in this field shall be less than the number of supported power states (e.g., if the controller supports 16 power states, then valid values are 0 through 15). A lower value means higher read throughput.	0x00
95:64	Exit Latency (EXLAT): This field indicates the maximum exit latency in microseconds associated with entering this power state.	0x00
63:32	Entry Latency (ENLAT): This field indicates the maximum entry latency in microseconds associated with entering this power state.	0x00
31:26	Reserved	
25	Non-Operational State (NOPS): This field indicates whether the controller processes I/O commands in this power state. If this field is cleared to '0', then the controller processes I/O commands in this power state. If this field is set to '1', then the controller does not process I/O commands in this power state.	0x00
24	Max Power Scale (MXPS): This field indicates the scale for the Maximum Power field. If this field is cleared to '0', then the scale of the Maximum Power field is in 0.01 Watts. If this field is set to '1', then the scale of the Maximum Power field is in 0.0001 Watts.	0x00
23:16	Reserved	
15:00	Maximum Power (MP): This field indicates the maximum power consumed by the NVM subsystem in this power state. The power in Watts is equal to the value in this field multiplied by the scale specified in the Max Power Scale field.	<variable>

Table 76: Identify Namespace Data Structure

Identify Namespace Data Structure, NVM Command Set Specific			
Bytes	O/M	Description	Expected Value
7:00	M	Namespace Size (NSZE): This field indicates the total size of the namespace in logical blocks. A namespace of size n consists of LBA 0 through $(n - 1)$. The number of logical blocks is based on the formatted LBA size. This field is undefined prior to the namespace being formatted.	Varies
15:08	M	Namespace Capacity (NCAP): This field indicates the maximum number of logical blocks that may be allocated in the namespace at any point in time. The number of logical blocks is based on the formatted LBA size. This field is undefined prior to the namespace being formatted. This field is used in the case of thin provisioning and reports a value that is smaller than or equal to the Namespace Size. Spare LBAs are not reported as part of this field.	Equal to NSZE if NS is formatted and attached to the controller; Else undefined. Equal to 0 if NS is valid but deleted or detached from controller
		A value of 0h for the Namespace Capacity indicates that the namespace ID is an inactive namespace ID	
		A logical block is allocated when it is written with a Write or Write Uncorrectable command. A logical block may be deallocated using the Dataset Management command, Sanitize, or Write Zeros command.	
23:16	M	Namespace Utilization (NUSE): This field indicates the current number of logical blocks allocated in the namespace. This field is smaller than or equal to the Namespace Capacity. The number of logical blocks is based on the formatted LBA size.	Equal to NSZE
		When using the NVM command set: A logical block is allocated when it is written with a Write or Write Uncorrectable command. A logical block may be deallocated using the Dataset Management command.	

Table 76: Identify Namespace Data Structure

Identify Namespace Data Structure, NVM Command Set Specific			
Bytes	O/M	Description	Expected Value
24	M	Namespace Features (NSFEAT): This field defines features of the namespace.	0x0
		Bits 7:5 are reserved.	
		Bit 4 if set to '1': Indicates that the fields NPWG, NPWA, NPDG, NPDA, and NOWS are defined for this namespace and should only be used by the host for I/O optimization; and NOWS defined for this namespace shall adhere to Optimal Write Size field setting defined in NVM Sets Attributes Entry for the NVM Set with which this namespace is associated. If Cleared to '0' then: the controller does not support the field NPWG, NPWA, NPDG, NPDA and NOWS for this namespace; and Optimal Write Size field in NVM Sets Attributes Entry for the NVM Set with which this namespace is associated should be used by the host for I/O optimization	0x10
		Bit 3 if set to '1' indicates that the non-zero NGUID and non-zero EUI64 fields for this namespace are never reused by the controller. If cleared to '0', then the NGUID and EUI64 values may be reused by the controller for a new namespace created after this namespace is deleted. This bit shall be cleared to '0' if both NGUID and EUI64 fields are cleared to 0h.	0
		Bit 2 if set to '1' indicates that the controller supports the Deallocated or Unwritten Logical Block error for this namespace. If cleared to '0', then the controller does not support the Deallocated or Unwritten Logical Block error for this namespace.	0
		Bit 1 if set to '1' indicates that the fields NAWUN, NAWUPF, and NACWU are defined for this namespace and should be used by the host for this namespace instead of the AWUN, AWUPF, and ACWU fields in the Identify Controller data structure. If cleared to '0', then the controller does not support the fields NAWUN, NAWUPF, and NACWU for this namespace. In this case, the host should use the AWUN, AWUPF, and ACWU fields defined in the Identify Controller data structure in Figure 90.	0
		Bit 0 if set to '1' indicates that the namespace supports thin provisioning. Specifically, the Namespace Capacity reported may be less than the Namespace Size. When this feature is supported and the Dataset Management command is supported then deallocating LBAs shall be reflected in the Namespace Utilization field. Bit 0 if cleared to '0' indicates that thin provisioning is not supported and the Namespace Size and Namespace Capacity fields report the same value.	0

Table 76: Identify Namespace Data Structure

Identify Namespace Data Structure, NVM Command Set Specific			
Bytes	O/M	Description	Expected Value
25	M	Number of LBA Formats (NLBAF): This field defines the number of supported LBA size and metadata size combinations supported by the namespace. LBA formats shall be allocated in order (starting with 0) and packed sequentially. This is a 0's based value. The maximum number of LBA formats that may be indicated as supported is 16. The supported LBA formats are indicated in bytes 128 - 191 in this data structure. The LBA Format fields with an index beyond the value set in this field are invalid and not supported. LBA Formats that are valid, but not currently available may be indicated by setting the LBA Data Size for that LBA Format to 0h.	0x4
		The metadata may be either transferred as part of the LBA (creating an extended LBA which is a larger LBA size that is exposed to the application) or it may be transferred as a separate contiguous buffer of data. The metadata shall not be split between the LBA and a separate metadata buffer.	
		It is recommended that software and controllers transition to an LBA size that is 4KB or larger for ECC efficiency at the controller. If providing metadata, it is recommended that at least 8 bytes are provided per logical block to enable use with end-to-end data protection.	
26	M	Formatted LBA Size (FLBAS): This field indicates the LBA size & metadata size combination that the namespace has been formatted with.	Varies
		Bits 7:5 are reserved.	0
		Bit 4 if set to '1' indicates that the metadata is transferred at the end of the data LBA, creating an extended data LBA. Bit 4 if cleared to '0' indicates that all of the metadata for a command is transferred as a separate contiguous buffer of data.	1 for DIF and 0 for DIX based on Format NS
		Bits 3:0 indicates one of the 16 supported combinations indicated in this data structure. This is a 0's based value.	0, 1, 2, 3, or 4 based on Format
27	M	Metadata Capabilities (MC): This field indicates the capabilities for metadata.	0x3
		Bits 7:2 are reserved.	
		Bit 1 if set to '1' indicates the namespace supports the metadata being transferred as part of a separate buffer that is specified in the Metadata Pointer. Bit 1 if cleared to '0' indicates that the controller does not support the metadata being transferred as part of a separate buffer.	1
		Bit 0 if set to '1' indicates that the namespace supports the metadata being transferred as part of an extended data LBA. Specifically, the metadata is transferred as part of the data PRP Lists. Bit 0 if cleared to '0' indicates that the namespace does not support the metadata being transferred as part of an extended data LBA.	1

Table 76: Identify Namespace Data Structure

Identify Namespace Data Structure, NVM Command Set Specific			
Bytes	O/M	Description	Expected Value
28	M	End-to-end Data Protection Capabilities (DPC): This field indicates the capabilities for the end-to-end data protection feature. Multiple bits may be set in this field. Protection Information is not supported for format with Metadata size > 8 Bytes. Attempt to Format (or create Namespace) to LBA Format 4 (LBA Data Size = 4096 bytes, Metadata Size = 64 bytes) with Protection Information set to non-zero value will fail with command completion status value Ah - Invalid Format.	Varies
		Bits 7:5 are reserved.	
		Bit 4 if set to '1' indicates that the namespace supports protection information transferred as the last eight bytes of metadata. Bit 4 if cleared to '0' indicates that the namespace does not support protection information transferred as the last eight bytes of metadata.	1
		Bit 3 if set to '1' indicates that the namespace supports protection information transferred as the first eight bytes of metadata. Bit 3 if cleared to '0' indicates that the namespace does not support protection information transferred as the first eight bytes of metadata.	0
		Bit 2 if set to '1' indicates that the namespace supports Protection Information Type 3. Bit 2 if cleared to '0' indicates that the namespace does not support Protection Information Type 3.	0
		Bit 1 if set to '1' indicates that the namespace supports Protection Information Type 2. Bit 1 if cleared to '0' indicates that the namespace does not support Protection Information Type 2.	1
		Bit 0 if set to '1' indicates that the namespace supports Protection Information Type 1. Bit 0 if cleared to '0' indicates that the namespace does not support Protection Information Type 1.	0

Table 76: Identify Namespace Data Structure

Identify Namespace Data Structure, NVM Command Set Specific															
Bytes	O/M	Description	Expected Value												
29	M	End-to-end Data Protection Type Settings (DPS): This field indicates the Type settings for the end-to-end data protection feature.	Varies												
		Bits 7:4 are reserved.													
		Bit 3 if set to '1' indicates that the protection information, if enabled, is transferred as the first eight bytes of metadata. Bit 3 if cleared to '0' indicates that the protection information, if enabled, is transferred as the last eight bytes of metadata.													
		Bits 2:0 indicate whether Protection Information is enabled and the type of Protection Information enabled. The values for this field have the following meanings:													
		<table><tr><th>Value</th><th>Definition</th></tr><tr><td>000b</td><td>Protection Information is not enabled</td></tr><tr><td>001b</td><td>Protection Information is enabled, Type 1</td></tr><tr><td>010b</td><td>Protection Information is enabled, Type 2</td></tr><tr><td>011b</td><td>Protection Information is enabled, Type 3</td></tr><tr><td>100b - 111b</td><td>Reserved</td></tr></table>		Value	Definition	000b	Protection Information is not enabled	001b	Protection Information is enabled, Type 1	010b	Protection Information is enabled, Type 2	011b	Protection Information is enabled, Type 3	100b - 111b	Reserved
		Value		Definition											
		000b		Protection Information is not enabled											
001b	Protection Information is enabled, Type 1														
010b	Protection Information is enabled, Type 2														
011b	Protection Information is enabled, Type 3														
100b - 111b	Reserved														
30	O	Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC): This field specifies multi-path I/O and namespace sharing capabilities of the namespace. Bits 7:1 are reserved Bit 0: If set to '1' then the namespace may be accessible by two or more controllers in the NVM subsystem (i.e., may be a shared namespace). If cleared to '0' then the namespace is a private namespace and may only be accessed by the controller that returned this namespace data structure.	0x00												

Table 76: Identify Namespace Data Structure

Identify Namespace Data Structure, NVM Command Set Specific			
Bytes	O/M	Description	Expected Value
31	O	<p>Reservation Capabilities (RESCAP): This field indicates the reservation capabilities of the namespace. A value of 00h in this field indicates that reservations are not supported by this namespace.</p> <p>Bit 7 is reserved</p> <p>Bit 6 if set to '1' indicates that the namespace supports the Exclusive Access - All Registrants reservation type. If this bit is cleared to '0', then the namespace does not support the Exclusive Access - All Registrants reservation type.</p> <p>Bit 5 if set to '1' indicates that the namespace supports the Write Exclusive - All Registrants reservation type. If this bit is cleared to '0', then the namespace does not support the Write Exclusive - All Registrants reservation type.</p> <p>Bit 4 if set to '1' indicates that the namespace supports the Exclusive Access - Registrants Only reservation type. If this bit is cleared to '0', then the namespace does not support the Exclusive Access - Registrants Only reservation type.</p> <p>Bit 3 if set to '1' indicates that the namespace supports the Write Exclusive - Registrants Only reservation type. If this bit is cleared to '0', then the namespace does not support the Write Exclusive - Registrants Only reservation type.</p> <p>Bit 2 if set to '1' indicates that the namespace supports the Exclusive Access reservation type. If this bit is cleared to '0', then the namespace does not support the Exclusive Access reservation type.</p> <p>Bit 1 if set to '1' indicates that the namespace supports the Write Exclusive reservation type. If this bit is cleared to '0', then the namespace does not support the Write Exclusive reservation type.</p> <p>Bit 0 if set to '1' indicates that the namespace supports the Persist Through Power Loss capability. If this bit is cleared to '0', then the namespace does not support the Persist Through Power Loss Capability.</p>	0x00
32	O	<p>Format Progress Indicator (FPI): If a format operation is in progress, this field indicates the percentage of the namespace that remains to be formatted.</p> <p>Bit 7 if set to '1' indicates that the namespace supports the Format Progress Indicator defined by bits 6:0 in this field. If this bit is cleared to '0', then the namespace does not support the Format Progress Indicator and bits 6:0 in this field shall be cleared to 0h.</p> <p>Bits 6:0 indicate the percentage of the namespace that remains to be formatted (e.g., a value of 25 indicates that 75% of the namespace has been formatted and 25% remains to be formatted). A value of 0 indicates that the namespace is formatted with the format specified by the FLBAS and DPS fields in this data structure.</p>	0x00

Table 76: Identify Namespace Data Structure

Identify Namespace Data Structure, NVM Command Set Specific											
Bytes	O/M	Description	Expected Value								
33	O	Deallocate Logical Block Features (DLFEAT): This field indicates information about features that affect deallocating logical blocks for this namespace.	0x19								
		Bits 7:5 are reserved	0								
		Bit 4 if set to '1' indicates that the Guard field for deallocated logical blocks that contain protection information is set to the CRC for the value read from the deallocated logical block and its metadata (excluding protection information). If cleared to '0' indicates that the Guard field for the deallocated logical blocks that contain protection information is set to FFFFh	1								
		Bit 3 if set to '1' indicates that the controller supports the Deallocate bit in the Write Zeros command for this namespace. If set to '0' indicates that the controller does not support the Deallocate bit in the Write Zeros command for this namespace. This bit shall be set to the same value for all namespaces in the NVM subsystem.	1								
		Bits 2:0 indicate the values read from a deallocated logical block and its metadata (excluding protection information). The values for this field have the following meanings: <table><tr><th>Value</th><th>Definition</th></tr><tr><td>000b</td><td>Not reported</td></tr><tr><td>001b</td><td>All bytes set to 00h</td></tr><tr><td>010b</td><td>All bytes set to FFh</td></tr><tr><td>100b - 111b</td><td>Reserved</td></tr></table>	Value	Definition	000b	Not reported	001b	All bytes set to 00h	010b	All bytes set to FFh	100b - 111b
Value	Definition										
000b	Not reported										
001b	All bytes set to 00h										
010b	All bytes set to FFh										
100b - 111b	Reserved										
35:34	O	Namespace Atomic Write Unit Normal (NAWUN): This field indicates the namespace specific size of the write operation guaranteed to be written atomically to the NVM during normal operation. A value of 0h indicates that the size for this namespace is the same size as that reported in the AWUN field of the Identify Controller data structure. All other values specify a size in terms of logical blocks using the same encoding as the AWUN field.	0x00								
37:36	O	Namespace Atomic Write Unit Power Fail (NAWUPF): This field indicates the namespace specific size of the write operation guaranteed to be written atomically to the NVM during a power fail or error condition. A value of 0h indicates that the size for this namespace is the same size as that reported in the AWUPF field of the Identify Controller data structure. All other values specify a size in terms of logical blocks using the same encoding as the AWUPF field.	0x00								

Table 76: Identify Namespace Data Structure

Identify Namespace Data Structure, NVM Command Set Specific			
Bytes	O/M	Description	Expected Value
39:38	O	<p>Namespace Atomic Compare & Write Unit (NACWU): This field indicates the namespace specific size of the write operation guaranteed to be written atomically to the NVM for a Compare and Write fused command.</p> <p>A value of 0h indicates that the size for this namespace is the same size as that reported in the ACWU field of the Identify Controller data structure. All other values specify a size in terms of logical blocks using the same encoding as the ACWU field.</p>	0x00
41:40	O	<p>Namespace Atomic Boundary Size Normal (NABSN): This field indicates the atomic boundary size for this namespace for the NAWUN value. This field is specified in logical blocks. Writes to this namespace that cross atomic boundaries are not guaranteed to be atomic to the NVM with respect to other read or write commands.</p> <p>A value of 0h indicates that there are no atomic boundaries for normal write operations. All other values specify a size in terms of logical blocks using the same encoding as the AWUN field.</p>	0x00
43:42	O	<p>Namespace Atomic Boundary Offset (NABO): This field indicates the LBA on this namespace where the first atomic boundary starts.</p> <p>If the NABSN and NABSPF fields are cleared to 0h, then the NABO field shall be cleared to 0h. NABO shall be less than or equal to NABSN and NABSPF.</p>	0x00
45:44	O	<p>Namespace Atomic Boundary Size Power Fail (NABSPF): This field indicates the atomic boundary size for this namespace specific to the Namespace Atomic Write Unit Power Fail value. This field is specified in logical blocks. Writes to this namespace that cross atomic boundaries are not guaranteed to be atomic with respect to other read or write commands and there is no guarantee of data returned on subsequent reads of the associated logical blocks.</p> <p>A value of 0h indicates that there are no atomic boundaries for power fail or error conditions. All other values specify a size in terms of logical blocks using the same encoding as the AWUPF field.</p>	0x00
47:46	O	<p>Namespace Optimal IO Boundary (NOIOB): This field indicates the optimal IO boundary for this namespace. This field is specified in logical blocks. The host should construct read and write commands that do not cross the IO boundary to achieve optimal performance. A value of 0h indicates that no optimal IO boundary is reported.</p>	0x100 (512B) or 0x20 (4096B)
63:48	O	<p>NVM Capacity (NVMCAP): This field indicates the total size of the NVM allocated to this namespace. The value is in bytes. This field shall be supported if Namespace Management and Namespace Attachment commands are supported.</p> <p>Note: This field may not correspond to the logical block size multiplied by the Namespace Size field. Due to thin provisioning or other settings (e.g., endurance), this field may be larger or smaller than the Namespace Size reported.</p>	Varies

Table 76: Identify Namespace Data Structure

Identify Namespace Data Structure, NVM Command Set Specific			
Bytes	O/M	Description	Expected Value
65:64	O	Namespace Preferred Write Granularity (NPWG): This field indicates the smallest recommended write granularity in logical blocks for this namespace. This is a 0's based value. The size indicated should be less than or equal to Maximum Data Transfer Size (MDTS) that is specified in units of minimum memory page size. The value of this field may change if the namespace is reformatted. The size should be a multiple of Namespace Preferred Write Alignment (NPWA). Refer to section 8.25 for how this field is utilized to improve performance and endurance	0x0 (4096B) or 0x7 (512B)
67:66	O	Namespace Preferred Write Alignment (NPWA): This field indicates the recommended write alignment in logical blocks for this namespace. This is a 0's based value. The value of this field may change if the namespace is reformatted. Refer to section 8.25 for how this field is utilized to improve performance and endurance.	0x0 (4096B) or 0x7 (512B)
69:68	O	Namespace Preferred Deallocate Granularity (NPDG): This field indicates the recommended granularity in logical blocks for the Dataset Management command with the Attribute - Deallocate bit set to '1' in Dword 11. This is a 0's based value. The value of this field may change if the namespace is reformatted. The size should be a multiple of Namespace Preferred Deallocate Alignment (NPDA). Refer to section 8.25 for how this field is utilized to improve performance and endurance.	0x0 (4096B) or 0x7 (512B)
71:70	O	Namespace Preferred Deallocate Alignment (NPDA): This field indicates the recommended alignment in logical blocks for the Dataset Management command with the Attribute - Deallocate bit set to '1' in Dword 11. This is a 0's based value. The value of this field may change if the namespace is reformatted. Refer to section 8.25 for how this field is utilized to improve performance and endurance	0x0 (4096B) or 0x7 (512B)
73:72	O	Namespace Optimal Write Size (NOWS): This field indicates the size in logical blocks for optimal write performance for this namespace. This is a 0's based value. The size indicated should be less than or equal to Maximum Data Transfer Size (MDTS) that is specified in units of minimum memory page size. The value of this field may change if the namespace is reformatted. The value of this field should be a multiple of Namespace Preferred Write Granularity (NPWG). Refer to section 8.25 for how this field is utilized to improve performance and endurance.	0x0 (4096B) or 0x7 (512B)
103:74		Reserved	

Table 76: Identify Namespace Data Structure

Identify Namespace Data Structure, NVM Command Set Specific			
Bytes	O/M	Description	Expected Value
119:104	O	<p>Namespace Globally Unique Identifier (NGUID): This field contains a 128-bit value that is globally unique and assigned to the namespace when the namespace is created. This field remains fixed throughout the life of the namespace and is preserved across namespace and controller operations (e.g., controller reset, namespace format, etc.).</p> <p>This field uses the EUI-64 based 16-byte designator format. Bytes 114:112 contain the 24-bit Organizationally Unique Identifier (OUI) value assigned by the IEEE Registration Authority. Bytes 119:115 contain an extension identifier assigned by the corresponding organization. Bytes 111:104 contain the vendor specific extension identifier assigned by the corresponding organization. See the IEEE EUI-64 guidelines for more information.</p> <p>The controller shall specify a globally unique namespace identifier in this field or the EUI64 field when the namespace is created.</p>	Varies based on WWN & NSID
127:120		<p>IEEE Extended Unique Identifier (EUI64): This field contains a 64-bit IEEE Extended Unique Identifier (EUI-64) that is globally unique and assigned to the namespace when the namespace is created. This field remains fixed throughout the life of the namespace and is preserved across namespace and controller operations (e.g., controller reset, namespace format, etc.).</p> <p>The EUI-64 is a concatenation of a 24-bit or 36-bit Organizationally Unique Identifier (OUI or OUI-36) value assigned by the IEEE Registration Authority and an extension identifier assigned by the corresponding organization. See the IEEE EUI-64 guidelines for more information.</p> <p>The controller shall specify a globally unique namespace identifier in this field or the NGUID field when the namespace is created. If the controller is not able to allocate a globally unique 64-bit identifier then this field shall be cleared to 0h.</p>	Varies based on OUI +WWN+ #times NS is created

Table 76: Identify Namespace Data Structure

Identify Namespace Data Structure, NVM Command Set Specific			
Bytes	O/M	Description	Expected Value
131:128	M	LBA Format 0 Support (LBAF0): This field indicates the LBA format 0 that is supported by the controller.	LBA Format Data Structure in Table 77 on page 115
135:132	O	LBA Format 1 Support (LBAF1): This field indicates the LBA format 1 that is supported by the controller.	
139:136	O	LBA Format 2 Support (LBAF2): This field indicates the LBA format 2 that is supported by the controller.	
143:140	O	LBA Format 3 Support (LBAF3): This field indicates the LBA format 3 that is supported by the controller.	
147:144	O	LBA Format 4 Support (LBAF4): This field indicates the LBA format 4 that is supported by the controller.	
151:148	O	LBA Format 5 Support (LBAF5): This field indicates the LBA format 5 that is supported by the controller.	
155:152	O	LBA Format 6 Support (LBAF6): This field indicates the LBA format 6 that is supported by the controller.	
159:156	O	LBA Format 7 Support (LBAF7): This field indicates the LBA format 7 that is supported by the controller.	
163:160	O	LBA Format 8 Support (LBAF8): This field indicates the LBA format 8 that is supported by the controller.	
167:164	O	LBA Format 9 Support (LBAF9): This field indicates the LBA format 9 that is supported by the controller.	
171:168	O	LBA Format 10 Support (LBAF10): This field indicates the LBA format 10 that is supported by the controller.	
175:172	O	LBA Format 11 Support (LBAF11): This field indicates the LBA format 11 that is supported by the controller.	
179:176	O	LBA Format 12 Support (LBAF12): This field indicates the LBA format 12 that is supported by the controller.	
183:180	O	LBA Format 13 Support (LBAF13): This field indicates the LBA format 13 that is supported by the controller.	
187:184	O	LBA Format 14 Support (LBAF14): This field indicates the LBA format 14 that is supported by the controller.	
191:188	O	LBA Format 15 Support (LBAF15): This field indicates the LBA format 15 that is supported by the controller.	
383:192		Reserved	
4095:384	O	Vendor Specific (VS): This range of bytes is allocated for vendor specific usage.	

Note:

O = Optional. The content of the word is optional

M = Mandatory. The content of the word is mandatory

Table 77: LBA Format Data Structure

Bits	Description	512B	520B	4096B	4104B	4160B										
31:26	Reserved															
25:24	<p>Relative Performance (RP): This field indicates the relative performance of the LBA format indicated relative to other LBA formats supported by the controller. Depending on the size of the LBA and associated metadata, there may be performance implications. The performance analysis is based on better performance on a queue depth 32 with 4KB read workload. The meanings of the values indicated are included in the following table.</p> <table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Best performance</td></tr><tr><td>01b</td><td>Better performance</td></tr><tr><td>10b</td><td>Good performance</td></tr><tr><td>11b</td><td>Degraded performance</td></tr></table>	Value	Definition	00b	Best performance	01b	Better performance	10b	Good performance	11b	Degraded performance	0x02	0x02	0x02	0x02	0x02
Value	Definition															
00b	Best performance															
01b	Better performance															
10b	Good performance															
11b	Degraded performance															
23:16	<p>LBA Data Size (LBADS): This field indicates the LBA data size supported. The value is reported in terms of a power of two (2^n). A value smaller than 9 (i.e. 512 bytes) is not supported.</p> <p>Note: If the value reported is 0h then the LBA format is not supported / used.</p>	0x09	0x09	0x0C	0x0C	0x0C										
15:00	<p>Metadata Size (MS): This field indicates the number of metadata bytes provided per LBA based on the LBA Size indicated. The namespace may support the metadata being transferred as part of an extended data LBA or as part of a separate contiguous buffer. If end-to-end data protection is enabled, then the first eight bytes or last eight bytes of the metadata is the protection information.</p>	0x00	0x08	0x00	0x08	0x40										

Appendix

Appendix D: NVMe Management Interface 1.1 (NVMe MI 1.1) Supported Command Set

Table 78: NVMe MI 1.1 Supported Command Set

	Features	O/M	Supported
I2C Access	Refer to NVMe MI Basic Specification Appendix A	M	Yes
	VPD Read	M	Yes
MCTP over SMBus	MCTP Command Set Endpoint ID	M	Yes
	MCTP Command Get Endpoint ID	M	Yes
	MCTP Command Get Endpoint UUID	M	Yes
	MCTP Command Get MCTP Version	M	Yes
	MCTP Command Get Message Type	M	Yes
	MCTP Command Prepare for Endpoint Discovery	O	Yes
	MCTP Command Endpoint Discovery	O	Yes
	Fairness arbitration	M	No
	MCTP packet timing requirements (DSP0237 Page 32)	O	Yes
MI Commands MCTP over SMBus (from NVMe MI 1.1 specification)	Get Configuration	M	Yes
	Set Configuration	M	Yes
	Controller Health Status Poll	M	Yes
	NVMe Subsystem Health Status Poll	M	Yes
	Read NVMe-MI Data Structure	M	Yes
	VPD Read	M	Yes
	VPD Write	M	Yes
	Reset	O	No
MI Primitives MCTP over SMBus (from NVMe MI 1.1 specification)	Resume	M	Yes
	Abort	M	Yes
	Get State	M	Yes
	Replay	M	Yes
	Pause	M	Yes

Table 78: NVMe MI 1.1 Supported Command Set

	Features	O/M	Supported
Admin Over NVMe-MI MCTP over SMBus (from NVMe MI 1.1 specification)	Set/Get Features	M	Yes
	Get Log Page	M	Yes
	Identify	M	Yes
	Device Self-test	O	Yes
	Firmware Activate/Commit	O	Yes
	Firmware Image Download	O	Yes
	Format NVM	O	Yes
	Namespace attach/detach	O	Yes
	Namespace management (Create/Delete)	O	Yes
	Sanitize	O	Yes
	Security Receive	O	No
	Security Send	O	No
MetaData MCTP over SMBus (from NVMe MI 1.1 specification)	NVMe Management Enhancement: Identify Controller	M	Yes
	NVMe Management Enhancement: Controller Meta-data	M	No
	NVMe Management Enhancement: Namespace Meta-data	M	No
PCIe MCTP over SMBus (from NVMe MI 1.1 specification)	PCIe Configuration Write	O	No
	PCIe Configuration Read	O	Yes
	PCIe I/O Read	O	No
	PCIe I/O Write	O	No
	PCIe Memory Write	O	No
	PCIe Memory Read	O	No
Spec 1.1	NVMe-MI over SMBus	O	Yes
	NVMe-MI over MCTP/VDM/PCIe	O	Yes
Spec NVMe 1.4 MI 1.1	NVMe Admin Commands over NVMe-MI	O	No

Note:

- O is optional per NVMe MI 1.1 specification
- M is mandatory per NVMe MI 1.1 specification

Appendix

Appendix E: Vital Product Data Structure (0x53)

Each NVM Subsystem with one or more Management Endpoints shall have a FRU Information Device which is compliant with the IPMI Platform Management FRU Information Storage Definition. The VPD shall contain the required elements defined in table below. The size of the VPD is 256 bytes as defined by the IPMI Platform Management FRU Information Storage Definition.

Refer to Appendix A of the NVMe-MI 1.1 Specification on www.nvme.org for Basic Command Management description.

Table 79: Vital Product Data Structure (VPD Elements)

Byte	Name
07:00	Common Header
79:08	Product Info Area
Vendor Specific:120	MultiRecord Info Area
Vendor Specific	Internal Use Area (optional)
Vendor Specific	Chassis Info Area (optional)
Vendor Specific	Board Info Area (optional)

The VPD shall be accessible using the VPD Read command. The entire contents of the VPD may be updated using the VPD Write command.

If the NVM Subsystem has a 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.

VPD records utilize the Type/Length byte format defined in the IPMI Platform Management FRU Information Storage Definition. Type/Length byte encodings utilized in this specification are summarized in the following table

Table 80: 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 000000b indicates that th field is empty

Table 81: Common Header

Byte	Factory Default	Description
0	01h	IPMI Format Version Number (IPMIVER): This field indicates the IPMI Format Version.

Table 81: Common Header

Byte	Factory Default	Description
1	00h	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 00h may be used to indicate the Internal Use Area is not present.
2	00h	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 00h may be used to indicate the Chassis Info Area is not present.
3	00h	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 00h may be used to indicate the Board Info Area is not present.
4	01h	Product Info Area Starting Offset (PIAOFF): This field indicates the starting offset in multiples of 8 bytes for the Product Info Area.
5	0Ah	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		Common Header Checksum (CHCHK): Checksum computed over bytes 0 through 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.

Table 82: Product Info Area (Offset 8 Bytes)

Solidigm™ D7-P5520/D7-P5620		Intel® SSD D7-P5520/D7-P5620		Description
Byte Offset	Factory Default	Byte Offset	Factory Default	
0	01h	0	01h	IPMI Format Version Number (IPMIVER): This field indicates the IPMI Format Version.
1	09h	1	09h	Product Info Area Length (PALEN): This field indicates the length of the product info area in multiples of 8 bytes. 112 bytes/8 = 14 = 0x0Eh
2	19h	2	19h	Language Code (LCODE): This field indicates the language used. A value of 19h is used to indicate English.
3	C8h	3	C5h	Manufacturer Name Type/Length (MNTL): This byte indicates the type and length of the Manufacturer Name field.
11:4	534F4C4 9444947 4Dh	8:4	494E544 54Ch	Manufacturer Name (MNAME): This field indicates the Manufacturer name in 8-bit ASCII. Unused bytes should be NULL characters. The MNAME value shall correspond to "SOLIDIGM" or "INTEL" in ASCII padded with 0x00.
12	C5h	9	C5h	Product Name Type/Length (PNTL): This byte indicates the type and length of the Product Name field.

Table 82: Product Info Area (Offset 8 Bytes)

Solidigm™ D7-P5520/D7-P5620		Intel® SSD D7-P5520/D7-P5620		Description
Byte Offset	Factory Default	Byte Offset	Factory Default	
17:13	Varies depending on SKU	14:10	Varies depending on SKU	Product Name (PNAME): This field indicates the Product name in 8-bit ASCII. Unused bytes should be NULL characters.
18	D7h	15	D8h	Product Part/Model Number Type/Length (PPMNNTL): This byte indicates the type and length of the Product Part/Model Number field.
41:19	Varies depending on SKU	39:16	Varies depending on SKU	Product Part/Model Number (PPMN): This field indicates the Product Part/Model Number in 8-bit ASCII. Unused bytes should be NULL characters.
42	C2h	40	C2h	Product Version Type/Length (PVTL): This byte indicates the type and length of the Product Part/Model Number field.
44:43	3100h	42:41	3100h	Product Version (PVER): This field indicates the Product Version in 8-bit ASCII. Unused bytes should be NULL characters.
45	D2h	43	D4h	Product Serial Number Type/Length (PSNTL): This byte indicates the type and length of the Product Serial Number field.
63:46	Serial Number	63:44	Serial Number	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
64	0h	64	0h	Asset Tag Type/Length (ATTL): This byte indicates the type and length of the Asset Tag field. A value of 00h may be used to indicate an Asset Tag is not present.
65	0h	65	0h	Asset Tag (AT): This field indicates the asset tag.
66	0h	66	0h	FRU File ID Type/Length (ATTL): This byte indicates the type and length of the FRU File ID field. A value of 00h may be used to indicate a FRU File ID is not present.
67	0h	67	0h	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
68	0h	68	0h	Custom Product Info Area (CPIA): This optional field allows for the addition of custom Product Info Area fields that shall be proceeded with a Type/Length field
69	C1h	69	C1h	End of Record (EOR): A value of C1h in this byte indicates the end of record

Table 82: Product Info Area (Offset 8 Bytes)

Solidigm™ D7-P5520/D7-P5620		Intel® SSD D7-P5520/D7-P5620		Description
Byte Offset	Factory Default	Byte Offset	Factory Default	
70		70		Reserved
71		71		Product Info Area (PICCHK): Checksum computed over bytes 0 through 70. 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.

Table 83: NVMe MultiRecord Area

Byte Offset	Factory Default	Description
0	0Bh	NVMe Record Type
1	02h	Bit 7 – end of list; record format version = 2h
2	3Bh	Record Length (RLEN): This field indicates the length of the MultiRecord Area in bytes.
3		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 through the end of this record plus this checksum byte equals zero)
4		Header Checksum: This field is used to give the record header a zero checksum (i.e., the modulo 256 sum of the preceding record bytes starting with the first byte of the header plus this checksum byte equals zero.
5	0h	NVMe MultiRecord Area Version Number: This field indicates the version number of this multirecord. This field shall be set to 0h in this version of the specification.

Table 83: NVMe MultiRecord Area

Byte Offset	Factory Default	Description																																																											
06	11h	Management Endpoint Form Factor (MEFF): This field indicates the form factor of the Management Endpoint.																																																											
		Value	Definition	0	Other - unknown	1 - 15	Reserved	16	2.5" Form Factor - unknown	17	2.5" Form Factor - U.2 (SFF-8639) 15mm	18	2.5" Form Factor - U.2 (SFF-8639) 7mm	19 - 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)	36-47	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	54-63	Reserved	64	BGA SSD - unknown	65	BGA SSD - 16 x 20mm (M.2 Type 1620)	66	BGA SSD - 16 x 20mm (M.2 Type 1113)	67-79	Reserved	80	Enterprise and Datacenter SSD Form Factor - unknown	81	E1.S Form Factor - (SFF-TA-1006) 5.9mm	82	E1.S Form Factor - (SFF-TA-1006) 8mm	83	E1.L Form Factor - (SFF-TA-1007) 9.5mm	84	E1.L Form Factor - (SFF-TA-1007) 18mm	85	E1.S Form Factor - (SFF-TA-1008) 7.5mm	86	3" Short Form Factor - (SFF-TA-1008) 16.8mm
		Value	Definition																																																										
		0	Other - unknown																																																										
		1 - 15	Reserved																																																										
		16	2.5" Form Factor - unknown																																																										
		17	2.5" Form Factor - U.2 (SFF-8639) 15mm																																																										
		18	2.5" Form Factor - U.2 (SFF-8639) 7mm																																																										
		19 - 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)																																																										
		36-47	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																																																										
		54-63	Reserved																																																										
		64	BGA SSD - unknown																																																										
		65	BGA SSD - 16 x 20mm (M.2 Type 1620)																																																										
		66	BGA SSD - 16 x 20mm (M.2 Type 1113)																																																										
		67-79	Reserved																																																										
		80	Enterprise and Datacenter SSD Form Factor - unknown																																																										
		81	E1.S Form Factor - (SFF-TA-1006) 5.9mm																																																										
		82	E1.S Form Factor - (SFF-TA-1006) 8mm																																																										
		83	E1.L Form Factor - (SFF-TA-1007) 9.5mm																																																										
		84	E1.L Form Factor - (SFF-TA-1007) 18mm																																																										
		85	E1.S Form Factor - (SFF-TA-1008) 7.5mm																																																										
		86	3" Short Form Factor - (SFF-TA-1008) 16.8mm																																																										

April 2024
349727-007US

Product Specification
12

Table 83: NVMe MultiRecord Area

Byte Offset	Factory Default	Description
12:07		Reserved
13	00h	Initial 1.8V Power Supply Requirements: This field specifies the initial 1.8V power supply requirements in Watts prior to receiving a Set Slot Power message.
14	00h	Maximum 1.8V Power Supply Requirements: This field specifies the maximum 1.8V power supply requirements in Watts. A value of zero indicates that the power supply voltage is not used.
15	00h	Initial 3.3V Power Supply Requirements: This field specifies the initial 3.3V power supply requirements in Watts prior to receiving a Set Slot Power message. 00h - U.2, Add-in card 08h - M.2
16	00h	Maximum 3.3V Power Supply Requirements: This field specifies the maximum 3.3V power supply requirements in Watts. A value of zero indicates that the power supply voltage is not used.
17	0h	Reserved
18	02h	Maximum 3.3Vaux Power Supply Requirements: This field specifies the maximum 3.3V power supply requirements in 10 mW units. A value of zero indicates that the power supply voltage is not used.
19	00h	Initial 5V Power Supply Requirements: This field specifies the initial 5V power supply requirements in Watts prior to receiving a Set Slot Power message.
20	00h	Maximum 5V Power Supply Requirements: This field specifies the maximum 5V power supply requirements in Watts. A value of zero indicates that the power supply voltage is not used.
21	12h	Initial 12V Power Supply Requirements: This field specifies the initial 12V power supply requirements in Watts prior to receiving a Set Slot Power message.
22	19h	Maximum 12V Power Supply Requirements: This field specifies the maximum 12V power supply requirements in Watts. A value of zero indicates that the power supply voltage is not used.
23	19h	Maximum Thermal Load: This field specifies the maximum thermal load from the NVM Subsystem in Watts.
36:24	Same as TNVM-CAP	Total NVM Capacity: This field indicates the total NVM capacity of the Management Endpoint 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

Table 84: NVMe PCIe Port MultiRecord Area

Byte Offset	Factory Default	Description
0	0Ch	NVMe PCIe Port Record Type ID
1	82h	Bit 7 - end of list; record format version = 2h

Table 84: NVMe PCIe Port MultiRecord Area

Byte Offset	Factory Default	Description												
2	0Bh	Record Length (RLEN): This field indicates the length of the MultiRecord Area in bytes.												
3		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 through the end of this record plus this checksum byte equals zero)												
04		Header Checksum: This field is used to give the record header a zero checksum (i.e., the modulo 256 sum of the preceding record bytes starting with the first byte of the header plus this checksum byte equals zero.												
5	01h	NVMe PCIe Port MultiRecord Area Version Number: This field indicates the version number of this multirecord. This field shall be set to zero in this version of the specification.												
6	00h	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.												
7	01h	Port Information: This field indicates information about the PCIe Ports in the device. Bits 7 to 1 are reserved. 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). Single port.												
08	0Fh	PCIe Link Speed: This field indicates a bit vector of link speeds supported by the PCIe port. <table><tr><th>Bit</th><th>Definition</th></tr><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></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'.													

Table 84: NVMe PCIe Port MultiRecord Area

Byte Offset	Factory Default	Description																													
09	04h	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 Management Controller may compare this value with the PCIe Negotiated Link Width to determine if there has been a PCIe link training issue.																													
		Value	Definition	0	Reserved	1	PCIe x1	2	PCIe x2	3	Reserved	4	PCIe x4	5-7	Reserved	8	PCIe x8	9-11	Reserved	12	PCIe x12	13-15	Reserved	16	PCIe x16	17-31	Reserved	32	PCIe x32	33-255	Reserved
		Value	Definition																												
		0	Reserved																												
		1	PCIe x1																												
		2	PCIe x2																												
		3	Reserved																												
		4	PCIe x4																												
		5-7	Reserved																												
		8	PCIe x8																												
		9-11	Reserved																												
		12	PCIe x12																												
		13-15	Reserved																												
		16	PCIe x16																												
		17-31	Reserved																												
32	PCIe x32																														
33-255	Reserved																														
10	01h	MCTP Support: This field contains a bit vector that specifies the level of support for the NVMe Management Interface. Bits 7 to 1 are reserved. Bit 0, if set to '1' indicates that MCTP based management commands are supported on the PCIe port.																													

Table 84: NVMe PCIe Port MultiRecord Area

Byte Offset	Factory Default	Description												
11	01h	Ref Clk Capability: This field contains a bit vector that specifies the PCIe clocking modes supported by the port.												
		<table><tr><th>Bit</th><th>Definition</th></tr><tr><td>7:4</td><td>Reserved</td></tr><tr><td>3</td><td>Default value set to '0'. 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>Default value set to '0'. Set to '1' if the PCIe link supports Separate RefClk with SSC (SRIS). Otherwise cleared to '0'.</td></tr><tr><td>1</td><td>Default value set to '0'. 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></table>	Bit	Definition	7:4	Reserved	3	Default value set to '0'. Set to '1' if the device automatically uses RefClk if provided and otherwise uses SRIS. Otherwise cleared to '0'.	2	Default value set to '0'. Set to '1' if the PCIe link supports Separate RefClk with SSC (SRIS). Otherwise cleared to '0'.	1	Default value set to '0'. 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	Default value set to '0'. Set to '1' if the device automatically uses RefClk if provided and otherwise uses SRIS. Otherwise cleared to '0'.											
		2	Default value set to '0'. Set to '1' if the PCIe link supports Separate RefClk with SSC (SRIS). Otherwise cleared to '0'.											
		1	Default value set to '0'. 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	0h	Port Identifier: This field contains the NVMe-MI Port Identifier.												
15:13	00h	Reserved												

Table 85: Topology MultiRecord Area

Byte Offset	Factory Default	Description
0	0Dh	Topology Record Type ID
1	82h	Bit 7 - end of list: Bit 7 - last record in list; record format version = 82h
2	24h	Record Length (RLEN): This field indicates the length of the MultiRecord Area in bytes.
3		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 through the end of this record plus this checksum byte equals zero)
4		Header Checksum: This field is used to give the record header a zero checksum (i.e., the modulo 256 sum of the preceding record bytes starting with the first byte of the header plus this checksum byte equals zero.
5	00h	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
6	00h	Reserved
7	02h	Element Count (N): This field indicates the number of element descriptors in this topology multirecord. The value of 0h is reserved.
8	02h	Upstream Connector Element Descriptor:Type: This field indicates the type of the Element Descriptor

Table 85: Topology MultiRecord Area

Byte Offset	Factory Default	Description
9	00h	Upstream Connector Element Descriptor: Revision: This field indicated the revision of the Element Descriptor. The upstream Connector Element Descriptor Revision is 0h for this specification
10	15h	Upstream Connector Element Descriptor: Length: Number of bytes in the Element Descriptor
11	11h	Upstream Connector Element Descriptor: Form Factor: This field indicates the Form Factor of the NVMe Storage Device
12	00h	Upstream Connector Element Descriptor: 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.
14:13	00h	Reserved
15	02h	Upstream Connector Element Descriptor: 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.
17:16	0019h	Upstream Connector Element Descriptor: Maximum Power: This field specifies the maximum power in Watts consumed by the NVMe Storage Device.
18	02h	Upstream Connector Element Descriptor: 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.
19	00h	Upstream Connector Element Descriptor: SMBus/I2C Upstream Port Descriptor: Type: This field indicates the type of the Port Descriptor. The SMBus/I2C Port Descriptor Type is 0h.
20	4h	Upstream Connector Element Descriptor: SMBus/I2C Upstream Port Descriptor: Length: This field indicates the length of the SMBus/I2C Port Descriptor in bytes
21	01h	Upstream Connector Element Descriptor: SMBus/I2C Upstream Port Descriptor: 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
22	01h	Upstream Connector Element Descriptor: SMBus/I2C Upstream Port Descriptor: 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
23	01h	Upstream Connector Element Descriptor: PCIe Upstream Port Descriptor: Type: This field indicates the type of Upstream Port Descriptor. The PCIe Upstream Port Descriptor Type is 1h.
24	6h	Upstream Connector Element Descriptor: PCIe Upstream Port Descriptor: Length: This field indicates the length of the PCIe Upstream Port Descriptor in bytes.
25	0h	Upstream Connector Element Descriptor: PCIe Upstream Port Descriptor: Starting Lane: This field indicates first PCIe lane (i.e., lane 0) of the port from the Upstream Connector.
26	03h	Upstream Connector Element Descriptor: PCIe Upstream Port Descriptor: Ending Lane: This field indicates the ending PCIe lane of the port from the Upstream Connector.

Table 85: Topology MultiRecord Area

Byte Offset	Factory Default	Description
27	01h	Upstream Connector Element Descriptor: PCIe Upstream Port Descriptor: PCIe pointer: This field contains the child index of the Element Descriptor whose PCIe port is connected to this PCIe Upstream Port
28	00h	Upstream Connector Element Descriptor: PCIe Upstream Port Descriptor: 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.
29	07h	NVM Subsystem Element Descriptor: Type: This field indicates the type of the Element Descriptor. The NVM Subsystem Element Descriptor Type is 7h.
30	00h	NVM Subsystem Element Descriptor: Revision: This field indicates the revision of the Element Descriptor. The NVM Subsystem Element Descriptor Revision is 0h for this specification.
31	0Ch	NVM Subsystem Element Descriptor: Length: This field indicates the length of the NVM Subsystem Element Descriptor in bytes.
32	3Bh	NVMe Subsystem Element Descriptor: 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 Bits 7:1 - SMBus/I2C Address: This field contains the 7-bit SMBus/I2C address Bit 0 - ARP Capable: This bit is set to '1' if SMBus ARP is supported, else it is cleared to '0'
33	01h	NVMe Subsystem Element Descriptor: 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 Bit 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 Bit 6:2: Reserved Bit 1:0 : Maximum Speed - This field is set to the highest supported SMBus/I2C clock speed Value of 0 = 100kHz Value of 1 = 400kHz Value of 2 = 1 MHz Value of 3 = Reserved
34	01h	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
35	00h	NVM Subsystem Port Descriptor 0: Type: This field indicates the type of an NVM Subsystem Port Descriptor. The NVM Subsystem Port Descriptor Type is 0.
36	06h	NVM Subsystem Port Descriptor 0: Length: This field indicates the length of the NVM Subsystem Port Descriptor in bytes.

Table 85: Topology MultiRecord Area

Byte Offset	Factory Default	Description																														
37	0Fh	NVM Subsystem Port Descriptor 0: PCIe Link Speed: This field indicates a bit vector of link speeds supported by the PCIe port Bit 0 - Set to '1' if the PCIe link supports 2.5 GT/s. Otherwise cleared to '0' Bit 1 - Set to '1' if the PCIe link supports 5.0 GT/s. Otherwise cleared to '0' Bit 2 - Set to '1' if the PCIe link supports 8.0 GT/s. Otherwise cleared to '0' Bit 3 - Set to '1' if the PCIe link supports 16.0 GT/s. Otherwise cleared to '0' Bits 7:4 - Reserved																														
38	04h	NVM Subsystem Port Descriptor 0: 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 Management Controller may compare this value with the PCIe Negotiated Link Width to determine if there has been a PCIe link training issue. <table><tr><th>Value</th><th>Definition</th></tr><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>5-7</td><td>Reserved</td></tr><tr><td>8</td><td>PCIe x8</td></tr><tr><td>9-11</td><td>Reserved</td></tr><tr><td>12</td><td>PCIe x12</td></tr><tr><td>13-15</td><td>Reserved</td></tr><tr><td>16</td><td>PCIe x16</td></tr><tr><td>17-31</td><td>Reserved</td></tr><tr><td>32</td><td>PCIe x32</td></tr><tr><td>33-255</td><td>Reserved</td></tr></table>	Value	Definition	0	Reserved	1	PCIe x1	2	PCIe x2	3	Reserved	4	PCIe x4	5-7	Reserved	8	PCIe x8	9-11	Reserved	12	PCIe x12	13-15	Reserved	16	PCIe x16	17-31	Reserved	32	PCIe x32	33-255	Reserved
Value	Definition																															
0	Reserved																															
1	PCIe x1																															
2	PCIe x2																															
3	Reserved																															
4	PCIe x4																															
5-7	Reserved																															
8	PCIe x8																															
9-11	Reserved																															
12	PCIe x12																															
13-15	Reserved																															
16	PCIe x16																															
17-31	Reserved																															
32	PCIe x32																															
33-255	Reserved																															

Table 85: Topology MultiRecord Area

Byte Offset	Factory Default	Description
39	01h	NVM Subsystem Port Descriptor 0: Ref Clk Capability: This field contains a bit vector that specifies the PCIe clocking modes supported by the port.
40	00h	NVM Subsystem Port Descriptor 0: Port Identifier: Port Identifier: This field contains the NVMe-MI Port Identifier.

Appendix

Appendix F: Out-of-Band Command Response Using SMBus (0x6A)

The SMBus slave address to read Subsystem Management Data is 0x6Ah

Refer to Appendix A of the NVMe-MI 1.1 Specification on www.nvme.org for Basic Command Management description.

Table 86: System Management Data Structure (NVMe-MI Commands)

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 Controller 1 in the NVM subsystem.</p> <p>SMBus Arbitration - Bit 7 is set '1' after a SMBus block read is completed all the way to the stop bit without bus contention and cleared to '0' if a 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 cannot process 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 segment.</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 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 NVMe 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. <p>Set to '1' if all controllers in the NVM subsystem do not indicate a critical warning for the corresponding bit.</p>

Table 86: System Management Data Structure (NVMe-MI Commands)

Command Code	Offset (byte)	Description																
	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 to +127° C. The 8 bit format of the data is shown below. This field should not report a temperature when that is older than 5 seconds. If recent data is not available, the NVMe management endpoint should indicate a value of 80h for this field.																
		<table><tr><th>Value</th><th>Description</th></tr><tr><td>00h-7Eh</td><td>Temperature is measured in degrees Celsius (0° 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 than 5 seconds old.</td></tr><tr><td>81h</td><td>Temperature sensor failure</td></tr><tr><td>82h-C3h</td><td>Reserved</td></tr><tr><td>C4</td><td>Temperature is -60° C or lower</td></tr><tr><td>C5-FFh</td><td>Temperature measured in degrees Celsius is represented in twos complement (-1° to -59° C)</td></tr></table>	Value	Description	00h-7Eh	Temperature is measured in degrees Celsius (0° to 126° C)	7Fh	127° C or higher	80h	No temperature data or temperature data is more than 5 seconds old.	81h	Temperature sensor failure	82h-C3h	Reserved	C4	Temperature is -60° C or lower	C5-FFh	Temperature measured in degrees Celsius is represented in twos complement (-1° to -59° C)
		Value	Description															
		00h-7Eh	Temperature is measured in degrees Celsius (0° to 126° C)															
		7Fh	127° C or higher															
		80h	No temperature data or temperature data is more than 5 seconds old.															
		81h	Temperature sensor failure															
		82h-C3h	Reserved															
	C4	Temperature is -60° C or lower																
	C5-FFh	Temperature measured in degrees Celsius is represented in twos complement (-1° to -59° C)																
04	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.																	
06:05	Reserved: Shall be set to 0000h																	
07	PEC: An 8 bit CRC calculated over the slave address, command code, second slave address and returned data. Algorithm is in SMBus Specifications.																	
8	08	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.																
	10:09	Vendor ID: The 2 byte vendor ID, assigned by the PCI SIG. Should match VID in the Identify Controller command response. MSB is transmitted first.																
	30:11	Serial Number: 20 characters that match the serial number in the NVMe Identify Controller command response. First character is transmitted first																
	31	PEC: An 8 bit CRC calculated over the slave address, command code, second slave address and returned data. Algorithm is in SMBus Specifications.																

Table 86: System Management Data Structure (NVMe-MI Commands)

Command Code	Offset (byte)	Description
32+	255:32	Vendor Specific - This data structure shall not exceed the maximum read length of 255 specified in the SMBus version 3 specification. Preferably length is not greater than 32 for compatibility with SMBus 2.0, additional blocks shall be on 8 byte boundaries.

Appendix

Appendix G: Out-of-Band Command Response Using SMBus (0x6A Solidigm Specific)

Table 87: Command Response 0x6A (Solidigm Specific Vendor Unique Commands)

Command Code	Offset (byte)	Description	Expected Value
32	32	Length of Solidigm Corporation's Block: shall be E1h until this spec is updated	1Eh
	33	Reserved	0h
	34	Power Measurement: Value = 0xff if not implemented or not functional	Feature not functional
	35	WCTEMP: Minimum Composite Temperature field value in Celsius that indicates an overheating condition during which controller operation continues	157h (343K, 70°C)
	36	CCTEMP: This field indicates the minimum Composite Temperature field value that indicates a critical overheating condition.	161h (353K, 80°C)
	38:37	Reserved	0h
	46:39	Firmware Version: 8 characters, ASCII representation	Varies per firmware release version. FW Revision (FR) from Identify Controller.
	54:47	Bootloader Version: 8 characters, ASCII representation	Varies per firmware release version
	62:55	Worldwide Number: Unique 8 Byte of WWN; factory programmed. This field matches Upper and Lower bytes of DSN register in PCIe space	Varies
	63	PEC: An 8 bit CRC calculated over the slave address, command code, second slave address and returned data. Algorithm is in SMBus Specifications.	Varies

Table 87: Command Response 0x6A (Solidigm Specific Vendor Unique Commands)

Command Code	Offset (byte)	Description	Expected Value
64	64	Length: Indicates number of additional bytes to read before encountering PEC.	1Eh
	65	Reserved: Shall be set to 0000h.	0000h
	67:66	Device ID	See PCIe ID PCIe ID on page 146
	69:68	Subsystem Vendor ID	See PCIe ID PCIe ID on page 146
	71:70	Subsystem Device ID	See PCIe ID PCIe ID on page 146
	83:82	Reserved: Shall be set to 0000h.	0000h
	84	Reserved: Shall be set to 0000h.	0000h
	85	Throttle State	
	86	Fault State	
	94:87	Reserved: Shall be set to 0000h.	0000h
	95	PEC: An 8 bit CRC calculated over the slave address, command code, second slave address and returned data. Algorithm is in SMBus Specifications.	Varies
96	96	Length of Status: Indicates number of additional bytes to read before encountering PEC. This value should always be 6 in implementations of this version of the spec.	1Eh
	97	Reserved: Shall be set to 0000h.	0000h
	107:98	FW Version Current	Varies
	117:108	FW Version next	Varies
	128:118	Reserved: Shall be set to 0000h.	0000h
	126	Reserved: Shall be set to 0000h.	0000h
	127	PEC: An 8 bit CRC calculated over the slave address, command code, second slave address and returned data. Algorithm is in SMBus Specifications.	Varies

Table 87: Command Response 0x6A (Solidigm Specific Vendor Unique Commands)

Command Code	Offset (byte)	Description	Expected Value
128	128	Length: Number of bytes expected from CC 128.	2Bh
	129	Reserved: Shall be set to 0000h.	0000h
	130	Throttle: Hardcoded temperature value that throttling begins.	157h (343K, 70°C)
	131	TShutdown: Hardcoded temp value that drive will shutdown.	161h (353K, 80°C)
	171:132	Model Number	Varies depending on SKU
	172	PEC: An 8 bit CRC calculated over the slave address, command code, second slave address and returned data. Algorithm is in SMBus Specifications.	Varies

Appendix

Appendix H: Persistent Event Log Events

The values that may be reported in the Event Type field (refer to section 5.14.1.13 of the NVMe 1.4 Specification) of the event header for events in the Persistent Event log are defined in the following table.

Table 88: Persistent Events

Type	Event	NVMe 1.4 Specification Reference Section
00h	Reserved	
01h	SMART / Health Log Snapshot	5.14.1.13.1.1
02h	Firmware Commit	5.14.1.13.1.2
03h	Timestamp Change	5.14.1.13.1.3
04h	Power-on or Reset	5.14.1.13.1.4
05h	NVM Subsystem Hardware Error	5.14.1.13.1.5
06h	Change Namespace	5.14.1.13.1.6
07h	Format NVM Start	5.14.1.13.1.7
08h	Format NVM Completion	5.14.1.13.1.8
09h	Sanitize Start	5.14.1.13.1.9
0Ah	Sanitize Completion	5.14.1.13.1.10

Table 89: SMART / Health Log Snapshot Event Data Format (Event Type 01h)

Bytes	Description
511:0	Event Data: Contains a snapshot of the SMART/Health Information Log data.

Table 90: Firmware Commit Event Data Format (Event Type 02h)

Bytes	Description
07:00	Old Firmware Revision: Contains the firmware revision of the active firmware before this firmware commit event.
15:08	New Firmware Revision: Contains the firmware revision for the firmware that was requested to become active.
16	Firmware Commit Action: Contains the value from the Commit Action field in the Firmware Commit command.
17	Firmware Slot: Contains the value from the Firmware Slot field in the Firmware Commit command.
18	Status Code Type for Firmware Commit Command: Contains the status code type from the completion queue entry for the Firmware Commit command.
19	Status Returned for Firmware Commit Command: Contains the status code from the completion queue entry for the Firmware Commit command.

Table 90: Firmware Commit Event Data Format (Event Type 02h)

Bytes	Description
21:20	Vendor Assigned Firmware Commit Result Code: Contains a vendor specific value that provides more information about the result of the firmware commit. A value of 0h indicates that no vendor assigned firmware commit result code is provided.

Table 91: Timestamp Change Event (Event Type 03h)

Bytes	Description
07:00	Previous Timestamp: Contains a timestamp using the format Timestamp – Data Structure for Get Features as defined in Figure 300 containing the timestamp for the time immediately before the timestamp was changed (i.e., the old timestamp).
15:08	Milliseconds Since Reset: Contains the time since the last Controller Level Reset reported in milliseconds.

Table 92: Power-on or Reset Event (Event Type 04h)

Bytes	Description
7:0	Firmware Revision: Contains the firmware revision that becomes effective when CC.EN transitions from '0' to '1'.
EL-VSIL-1:8	Reset Information List: Contains a list of one or more Controller Reset Information descriptors (refer to NVMe 1.4 Specification Figure 218). If virtualization management is not implemented, then the list shall contain a Controller Reset Information descriptor for every controller in the NVM subsystem. If virtualization management is implemented, then the list shall contain a Controller Reset Information descriptor for every primary controller. The Controller Reset Information descriptor is shown in NVMe 1.4 Specification Figure 218.

Table 93: NVM Subsystem Hardware Error Event Format (Event Type 05h)

Bytes	Description
1:00	NVM Subsystem Hardware Error Event Code: This field contains a code (refer to NVMe 1.4 Specification Figure 220) indicating the type of NVM subsystem hardware error that is being reported.
3:02	Reserved
M+3:04	Additional Hardware Information: This field contains additional information about the hardware error event indicated in the NVM Subsystem Hardware Error Event Code field (refer to NVMe 1.4 Specification Figure 220). Where M is the number of bytes of additional hardware error information. This field is omitted if the subsystem hardware error being reported does not contain additional hardware error information.

Table 94: Change Namespace Event Data Format (Event Type 06h)

Bytes	Value
3:00	Namespace Management CDW10: Contains the value from command Dword 10 of the Namespace Management command that initiated the namespace change event (refer to NVMe 1.4 Specification Figure 264).
7:04	Reserved
15:08	Namespace Size (NSZE): For a create operation, contains the NSZE value from the Identify Namespace data structure in the Namespace Management command (refer to NVMe 1.4 Specification Figure 265). For a delete operation that specifies a single namespace this field contains the value from the NSZE field of the Identify Namespace data (refer to NVMe 1.4 Specification Figure 249) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved.
23:16	Reserved
31:24	Namespace Capacity (NCAP): For a creation operation, contains the NCAP value from the Identify Namespace data structure in the Namespace Management command (refer to NVMe 1.4 Specification Figure 265). For a delete operation that specifies a single namespace this field contains the value from the NCAP field of the Identify Namespace data (refer to NVMe 1.4 Specification Figure 245) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved.
32	Formatted LBA Size (FLBAS): For a create operation, contains the FLBAS value from the Identify Namespace data structure in the Namespace Management command (refer to Figure 265). For a delete operation that specifies a single namespace this field contains the value from the FLBAS field of the Identify Namespace data (refer to NVMe 1.4 Specification Figure 245) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved.
33	End-to-end Data Protection Type Settings (DPS): For a create operation, contains the DPS value from the Identify Namespace data structure in the Namespace Management command (refer to NVMe 1.4 Specification Figure 265). For a delete operation that specifies a single namespace this field contains the value from the DPS field of the Identify Namespace data (refer to NVMe 1.4 Specification Figure 245) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved.
34	Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC): For a create operation, contains the NMIC value from the Identify Namespace data structure in the Namespace Management command (refer to Figure 265). For a delete operation that specifies a single namespace this field contains the value from the NMIC field of the Identify Namespace data (refer to NVMe 1.4 Specification Figure 245) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved.
35	Reserved
39:36	ANA Group Identifier (ANAGRPID): For a create operation, contains the ANAGRPID value from the Identify Namespace data structure in the Namespace Management command (refer to Figure 265), if specified, or contains the ANAGRPID value from the Identify Namespace data (refer to Figure 245) after the namespace was created if an ANA Group Identifier was not specified in the command. For a delete operation that specifies a single namespace this field contains the value from the ANAGRPID field of the Identify Namespace data (refer to Figure 245) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved. If ANA Groups are not supported, then the ANAGRPID field shall be cleared to 0h.

Table 94: Change Namespace Event Data Format (Event Type 06h)

Bytes	Value
41:40	NVM Set Identifier (NVMSETID): For a create operation, contains the NVMSETID value from the Identify Namespace data structure in the Namespace Management command (refer to NVMe 1.4 Specification Figure 265). For a delete operation that specifies a single namespace this field contains the value from the NVMSETID field of the Identify Namespace data (refer to Figure 245) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved.
43:42	Reserved
47:44	Namespace ID (NSID): For a create operation, contains the NSID for the namespace that was created. For a delete operation, contains the NSID from CDW1.NSID (i.e., the NSID for the namespace being deleted or FFFFFFFFh for a delete operation specifying all namespaces)

Table 95: Format NVM Start Event Data Format (Event Type 07h)

Bytes	Description
03:00	Namespace Identifier: Contains the namespace identifier specified in the Format NVM command.
04	Format NVM Attributes (FNA): Contains the value from the identify controller FNA field.
07:05	Reserved
11:08	Format NVM CDW10: Contains the value from command Dword 10 of the Format NVM command (refer to NVMe 1.4 Specification figure 328).

Table 96: Format NVM Completion Event Data Format (Event Type 08h)

Bytes	Description
03:00	Namespace Identifier: Contains the namespace identifier specified in the Format NVM command.
04	Smallest Format Progress Indicator: For a Format NVM command that formats a single namespace this field contains the lowest numerical value that was available for reporting in the FPI field of the Identify Namespace data structure (i.e., if the format did not complete successfully and the FPI field is supported then this field contains the percentage of the namespace that remained to be formatted at the time the format NVM command completed, refer to NVMe 1.4 Specification Figure 245) during the format operation. For a Format NVM command that formats all namespaces this field shall be cleared to 0h.

Table 96: Format NVM Completion Event Data Format (Event Type 08h)

Bytes	Description								
05	Format NVM Status: Contains the status of the format operation.								
	<table><tr><th>Bits</th><th>Definition</th></tr><tr><td>7:2</td><td>Reserved</td></tr><tr><td>1</td><td>Incomplete Format: If set to '1', then the format operation modified one or more LBAs but did not complete successfully. If set to '1', then the Format NVM Error bit shall be set to '1'. If cleared to '0' then then the format operation either did not modify any LBAs or the format operation completed successfully.</td></tr><tr><td>0</td><td>Format NVM Error: If set to '1', then the format operation did not complete successfully. If cleared to '0', then the format operation completed successfully.</td></tr></table>	Bits	Definition	7:2	Reserved	1	Incomplete Format: If set to '1', then the format operation modified one or more LBAs but did not complete successfully. If set to '1', then the Format NVM Error bit shall be set to '1'. If cleared to '0' then then the format operation either did not modify any LBAs or the format operation completed successfully.	0	Format NVM Error: If set to '1', then the format operation did not complete successfully. If cleared to '0', then the format operation completed successfully.
	Bits	Definition							
	7:2	Reserved							
	1	Incomplete Format: If set to '1', then the format operation modified one or more LBAs but did not complete successfully. If set to '1', then the Format NVM Error bit shall be set to '1'. If cleared to '0' then then the format operation either did not modify any LBAs or the format operation completed successfully.							
0	Format NVM Error: If set to '1', then the format operation did not complete successfully. If cleared to '0', then the format operation completed successfully.								
07:06	Completion Information: Contains a vendor specific value that may provide more information about the completion of the format operation (e.g., if the format operation did not complete successfully, then this field may contain a vendor specific code that indicates a vendor specific reason).								
11:08	Status Field: Contains the value that was reported in the status code field for the completion queue entry, if any, for the Format NVM command associated with this event. If no completion queue entry was reported, then this field shall be cleared to 0h.								

Table 97: Sanitize Start Event Data Format (Event Type 09h)

Bytes	Description
03:00	SANICAP: Contains the contents of the SANICAP field from the Identify Controller data structure.
07:04	Sanitize CDW10: Contains the value from command Dword 10 of the Sanitize command (refer to NVMe 1.4 Specification Figure 330).
11:08	Sanitize CDW11: Contains the value from command Dword 11 of the Sanitize command (refer to NVMe 1.4 Specification Figure 331).

Table 98: Sanitize Completion Event Data Format (Event Type 0Ah)

Bytes	Description
1:0	Sanitize Progress: Contains the sanitize progress at the time of this event using the format specified for the SPROG field in the Sanitize Status log page (refer to section NVMe 1.4 Specification 5.14.1.16.2).
3:2	Sanitize Status: Contains the sanitize status for the time of this event using the format specified for the SSTAT field in the Sanitize Status log page. (e.g., the Global Data Erase bit indicates the status at the time of this event).
5:4	Completion Information: Contains a vendor specific value that may provide more information about the completion of the sanitize operation (e.g., if the sanitize operation did not complete successfully, then this field may contain a vendor specific code that indicates a vendor specific reason).
7:6	Reserved

Appendix

Appendix I: SCSI Command Translation

Following SCSI commands are supported:

- Inquiry
- Request Sense
- Mode Sense 6,10
- Mode Select 6,10
- Read Capacity 10,16
- Read 6,10,16
- Write 6,10,16
- Report LUNs
- Unmap
- Format Unit
- Log Sense
- Security Protocol In/Out
- Synchronize Cache 10,16
- Test Unit Ready
- Write Buffer

Note: Refer to NVM Express: SCSI translation reference doc under nvmexpress.org

Appendix

Appendix J: PCIe ID

Table 99: PCIe ID - U.2 15mm

ID Name	Description	D7-P5520 U.2 15mm	D7-P5620 U.2 15mm	PCIe Register Location	Identify Controller Location	Vital Product Data Location
Vendor ID (VID)	Vendor ID assigned by PCI-SIG	0x025E ¹	0x025E ¹	PCI Header Offset 00h (bits 15:00)	Bytes 01:00h	Address 3, (size 2B)
Device ID (DID)	Device ID assigned by vendor	0x0B60	0x0B60	PCI Header Offset 00h (bits 31:16)	NA	NA
Subsystem Vendor ID	Indicates Sub- system vendor ID	0x025E ¹	0x025E ¹	PCI Header Offset 2Ch (bits 15:00)	Bytes 03:02h	NA
Subsys- tem ID	Sub-system identifier	0x9008	0x9108	PCI Header Offset 2Ch (bits 31:16)	NA	NA

Note:

1. Solidigm™ D7-P5520/D7-P5620 SKUs have the Vendor ID and Subsystem Vendor ID set to 0x025E while Intel® SSD D7-P5520/D7-P5620 SKUs Vendor ID and Subsystem Vendor ID are set to 0x8086.

Table 100: PCIe ID - EDSFF

ID Name	Description	D7-P5520 - E1.S 9.5mm	D7-P5520 - E1.S 15mm	D7-P5520 - E1.L 9.5mm	PCIe Register Location	Identify Controller Location	Vital Product Data Location
Vendor ID (VID)	Vendor ID assigned by PCI-SIG	0x025E ¹	0x025E ¹	0x025E ¹	PCI Header Offset 00h (bits 15:00)	Bytes 01:00h	Address 3, (size 2B)
Device ID (DID)	Device ID assigned by vendor	0x0B60	0x0B60	0x0B60	PCI Header Offset 00h (bits 31:16)	NA	NA
Subsys- tem Ven- dor ID	Indicates Sub-system vendor ID	0x025E ¹	0x025E ¹	0x025E ¹	PCI Header Offset 2Ch (bits 15:00)	Bytes 03:02h	NA

Table 100: PCIe ID - EDSFF

ID Name	Description	D7-P5520 - E1.S 9.5mm	D7-P5520 - E1.S 15mm	D7-P5520 - E1.L 9.5mm	PCIe Register Location	Identify Controller Location	Vital Product Data Location
Subsys- tem ID	Sub-system identifier	0x900C	0x900D	0x901C	PCI Header Offset 2Ch (bits 31:16)	NA	NA

Note:

1. Solidigm™ D7-P5520/D7-P5620 SKUs have the Vendor ID and Subsystem Vendor ID set to 0x025E while Intel® SSD D7-P5520/D7-P5620 SKUs Vendor ID and Subsystem Vendor ID are set to 0x8086.

Appendix

Appendix K: LED Behavior - D7-P5520/D7-P5620

Appendix L: LED Behavior (SFF)

There are no LEDs on the drives for the U.2 form factor. The activity GPIO output is routed to Pin11 on the connector for driving an LED on the host system. Activity blink rate is Set Features C9h as described in Set/Get Blink activity for LED (C9h) [on page 55](#).

Activity LED shall be OFF upon drive Shutdown or Not-Ready/Disabled.

Commands over SMBus will not drive any LED activity.

Table 101: LED Behavior (U.2 Form Factor)

Drive State	Condition	LED Status (Pin 11 status)
Not-Ready (Drive Status)	CSTS.RDY is 0	OFF
Idle	CSTS.RDY is 1	Solid Green (ON) (LED color is defined by the host platform)
IO	Host Outstanding IO count > 0	Blink with blink rate set as 250ms
Admin Command	NVMe Admin Commands	OFF
Format	NVMe Format Command	Blink w/ blink rate set as 250ms
Shutdown	CC.SHN, both Normal and Abrupt Shutdown, or D3 Power State	OFF
Asserted/Fault State	N/A	OFF
Commands over SMBus	CSTS.RDY is 1 (Drive is ready)	Solid Green (ON)
Commands over SMBus	CSTS.RDY is 0 (Drive not ready)	OFF

Note:

1. Blink rate is Set Features C9h as described in Section 5.8.3
2. When idle, logic level is low (LED Solid On).
3. During IO activity, pin toggles 250msec high, 250msec low signal

Appendix M: LED Behavior (EDSFF) - D7-P5520/D7-P5620¹

The table below describes the LED behavior EDSFF form factor. There are 2 LEDs Green for I/O Activity and amber which can be driven by host and the drive as described in the table.

Table 102: LED Behavior (EDSFF Form Factor)

LED Color	Function
Green	Solid ON - Drive is idle Blinking - I/O activity (default 2 Hz blink rate)
Amber	Primary Function Driven by host via pin A10 on EDSFF header as described by the SFF-TA-1009 specification Additional Function Drive will turn LED Solid ON for a drive in disable logical mode; this will override any host-driven behavior. Note that pin A10 must be in a valid high or low state for LED point intensity to be spec-compliant.

Note:

1. These results are preliminary and provided for information purposes only. These values and claims are neither final nor official.