Ultrastar SN100 Series NVMe PCIe x4 Lane
2.5-Inch Small Form Factor (SFF)
Solid-State Drive Product Manual

Model(s): HUSPR3280ADP301
HUSPR3295ADP301
HUSPR3216ADP301
HUSPR3219ADP301
HUSPR3232ADP301
HUSPR3238ADP301

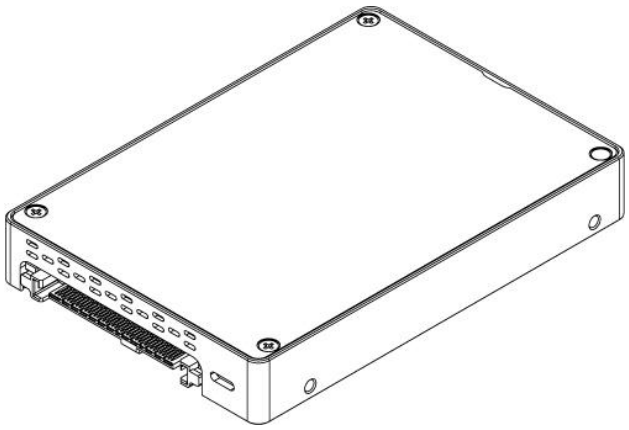

Document Number: 61000-08040-306

Document Version: 3.06

Revision Date: March 22, 2016

3RD Edition (Revision 3.06)

March 22, 2016

Conventions

The following icon and text conventions are used throughout this document to identify additional information of which the reader should be aware.

Conventions		Description
SHOCK HAZARD		***This icon indicates the danger of an electrical shock that may harm or otherwise prove fatal to the user***.
ESD		This icon indicates the possible presence of Electrostatic Discharge (ESD or "static electricity") that may harm the internal electronic components. The user is advised to handle the device only after discharging any possible electrostatic buildup that may be present.
CAUTION		This icon indicates the existence of a hazard that could result in equipment or property damage, or equipment failure if the safety instructions are not observed.
NOTE		This icon identifies information that relates to the safe operation of the equipment or related items.
Bold	**Text**	Used to indicate **important technical notes**.
Bold Italic	***Text***	Used to indicate ***critical instructions***.
Light Blue	Text	Used to indicate a hyperlink or "jump" to a related topic or subtopic. In addition, the text may be **bold**.
Dark Blue Bold	**Text**	Used to indicate a hyperlink or "jump" to a related topic or subtopic. In addition, the text may be ***bold italic***.

Revision History

Revision	Date	Page(s)	Description
-300	02/01/2015	All	Initial release.
-301	06/08/2015	All	Preliminary release.
-302	07/30/2015	40-43	Section 8; Environmental Characteristics. Complete re-write to align environmental conditions with similar SSD products.
		40	Table 22; consolidated operating and non-operating temperature, humidity and altitude conditions.
		41	Section 8.3; added Case Temperature hyperlink.
			Table 23 developed to document temperature check point.
			Figure 12 updated to reflect accurate location of temperature check point area above ASIC.
		43	Section 8.7.2; Operating Vibration and subsections developed.
			Section 8.7.4; Operating Shock topic and subsection developed; operating shock, non-operating shock and half-sine wave shock pulse data updated.
		63-92	Section 16. Appendix A: Log Pages developed.
		93-96	Section 17. Appendix B: Inquiry VPD Pages developed (Inquiry through SMBus and NVMe MI).
-303	09/02/2015	30	Table 14, Section 6.7.1. Clarified Pin P11 as an “Activity Output” signal, added “Notes” column to matrix and re-wrote subsequent notes to reflect design.
		97	Section 18, Contact Information. WWW URLs, email and telephone support information updated.
-304	09/24/2015	15	Section 2.7, Temperature Monitoring via SMBus. Clarified statement to indicate that both the SMBus and NVMe MI specifications and temperature monitoring are supported according to the power rail circuit implementation.
		93	Section 17.2, Enterprise SSD Form Factor Vital Product (VPD) Table; listed address offset of 0xA6.
		95	Section 17.4, NVMe Management Interface Table; listed address offset of 0xD4. Table 31; updated the NVMe MI Command Code 32 table.
		96	Table 32; corrected “Power” field description.
		97	Section 18, Contact Information. WWW URLs, email and telephone support information updated.
-305	12/14/2015	54	Table 28; struck row Subpage Name “Media Log” and Subpage Identifier “0x11”. Obsolete specification.
		75-77	Section 16.3.5, Subpage 0x11-Media Log Page struck from product manual. Obsolete specification.

Revision	Date	Page(s)	Description
-306	03/22/2016	Title	Title page updated with new model numbers for 956GB, 1.91TB and 3.82TB drives.
		12	Table 1 updated with new model numbers for 956GB, 1.91TB and 3.82TB drives.
		13	Section 1.4, Features; User Capacities item updated.
		16	Section 2.4, Drive Capacities; minor edit to include 956GB, 1.91TB and 3.82TB in first sentence.
		16	Table 3, Performance Characteristics updated.
			Table 4, Device Time to Ready table updated.
		20	Section 3.8, Thermal Throttling and Performance; the throttling threshold has been updated from 68°C to 74°C to reflect the firmware implementation and hardware characteristics.
		22	Section 5.6, Minimum Off-Time developed; the drive must be powered off for a minimum of 2.0 seconds before being powered on again.
		40	Table 24; calculated capacities for 956GB, 1.91TB and 3.82TB for both 512-byte and 4096-aligned data.
		64	Model Number Decoder updated with 956GB, 1.91TB and 3.82TB designators.

Table of Contents

List of Figures

List of Tables

1. Scope

1.1 Overview

This document describes the applications, specifications and installation of the Ultrastar SN100 Series NVMe PCIe 3.0 x4 Lane 2.5-Inch SFF SSD.

1.2 Audience

This manual is intended for system engineers or system designers employed by an Original Equipment Manufacturer (OEM). This document was therefore written specifically for a technically advanced audience; it is not intended for end-users that will eventually purchase the commercially available product. The ***user***, as referenced throughout the manual, is primarily concerned with industrial, commercial or military computing applications.

This device can be damaged by Electrostatic Discharge (ESD). When handling the device, always wear a grounded wrist strap and use a static dissipative surface.

Any damage to the SSD that occurs after its removal from the shipping package and ESD protective bag is the responsibility of the user.

1.3 Model Numbers

Table 1 lists the model numbers, capacities, NAND type used, available form factors and case heights.

Table 1: Capacities and Form Factors

Model Number	Capacity	NAND	Form Factor	Case Height
HUSPR3280ADP301	800GB	A19nm eMLC	2.5-Inch	15mm
HUSPR3295ADP301	956GB	A19nm eMLC	2.5-Inch	15mm
HUSPR3216ADP301	1.6TB (1600GB)	A19nm eMLC	2.5-Inch	15mm
HUSPR3219ADP301	1.91TB (1920GB)	A19nm eMLC	2.5-Inch	15mm
HUSPR3232ADP301	3.2TB (3200GB)	A19nm eMLC	2.5-Inch	15mm
HUSPR3238ADP301	3.82TB (3820GB)	A19nm eMLC	2.5-Inch	15mm

1.4 Features

Feature	Description
Form Factor	2.5-Inch, SFF-8223 Specification
Case Z-Height	15mm +0.0/0.5
Interface	PCI Express (PCIe) 3.0 x4 Lane
NVMe	NVMe Revision 1.0a Specification
UEFI Boot	UEFI Specification, Version 2.3; OS must be UEFI-aware.
User Capacities	800GB, 956GB, 1.6TB, 1.91TB, 3.2TB, 3.82TB
Sector Sizes	512-byte, 4096-byte
Operating Voltage	12V±15% DC
Operating Temperature	0°C to 70°C
Temperature Monitoring	In-band / Out-band using SMBus
SMBus Voltage	Optional +3.3Vaux±15% DC for SMBus/VPD implementation.
Thermal Throttling	Supported
Power Backup	10ms at EOL; 25-Watts
PowerSAFE™ Technology	*See PowerSAFE™ Technology*
Secure Array of Flash Elements™ (S.A.F.E.)	*See Secure Array of Flash Elements™ (S.A.F.E.) Technology*
EDC/ECC	100 bits per 4kB.
Event Logging	Non-volatile event logs.

1.5 Glossary and Acronyms

Table 2: Glossary and Acronyms

Term	Definition
AER	Advanced Error Recovery
AES	Advanced Encryption Standard
ASIC	Application-Specific Integrated Circuit
ATA	Advanced Technology Attachment
BER	Bit Error Rate, or percentage of bits that have errors relative to the total number of bits received.
BIOS	Basic Input / Output System
BIST	Built-In Self-Test
BOL	Beginning-of-Life
Chipset	A term used to define a collection of integrated components required to make a PC function.
CRC	Cyclic Redundancy Check
DAS	Device Activity Signal
DMA	Direct Memory Access
ECC	Error Control Coding (according on this definition, CRC is a subset). Can also refer to Error Correction Code.
EOL	End-of-Life
Gb	Gigabit
GB	Gigabyte
GUI	Graphical User Interface
HDD	Hard Disk Drive
Hot Plug	A term used to describe the removal or insertion of a hard drive while the system is powered.
I/O or IO	Input / Output
IOPS	Input / Output Operations Per Second
ISO	International Standards Organization
LBA	Logical Block Address
LED	Light Emitting Diode
MB	Megabyte
MLC	Multi-Level Cell (NAND)
MP	Multi-Processor
MTBF	Mean Time Between Failure
NOP	No Operation
NVMe	NVM Express. NVMe was developed by the Non-Volatile Memory Host Controller Interface (NVMHCI) Workgroup.

Term	Definition
OS	Operating System
PB	Petabyte
PCB	Printed Circuit Board
PCIe	Peripheral Component Interconnect Express. The Peripheral Component Interconnect Special Interest Group (www.pcisig.com) oversees the development and ongoing enhancement of the PCI standard.
Port	The point at which a drive physically connects to the system.
PQI	PCIe Queuing Interface; a T10 standard in collaboration with SOP.
RAID	Redundant Array of Independent Disks
RAM	Random Access Memory
RMS	Root Mean Square
ROM	Read Only Memory
SAS	Serial-Attached SCSI
SCSI	Small Computer System Interface
SFF	Small Form Factor
SMART	Self-Monitoring, Analysis and Reporting Technology: An open standard for developing hard drives and software systems that automatically monitors the health of a drive and reports potential problems.
SOP	SCSI Over PCIe
SSD	Solid State Drive
TB	Terabyte
UBER	Uncorrectable Bit Error Rate
UI	User Interface
VPD	Vital Product Data
VUC	Vendor-Unique Command
WHQL	Windows Hardware Quality Labs
Write Cache	A memory device within a SSD which is allocated for the temporary storage of data before that data is copied to its permanent storage location.
XOR	Exclusive OR
XTS-AES	A method for Data at Rest encryption according to the AES as standardized by IEEE P1619.

2. Product Description

2.1 Overview

The Ultrastar® SN100 Series NVMe PCIe 3.0 x4 Lane 2.5-Inch SFF SSD is a non-volatile, low-latency, high-performance device and is intended for cloud, hyperscale, enterprise data centers and computing environments that require high levels of application acceleration. The SSD can be configured as a boot device, data storage device, or caching device.

2.2 Interface

The SSD can be installed in any system that conforms to the PCI Express 3.0 Specification. The modular design allows direct access to the system bus, thereby reducing CPU overhead. The design specifically excludes the use of host resources, with no use of the host CPU, host interrupts or host memory.

2.3 Reliability

The solid-state design eliminates electromechanical noise and delay inherent in traditional magnetic rotating media. The wear-leveling and bad-block mapping algorithms ensure data integrity. The embedded Error Detection Code and Cyclic Redundancy Checking (EDC/CRC) ensures data reliability. *See Reliability Characteristics*.

2.4 Drive Capacities

The SSD is available in unformatted capacities of 800GB, 956GB, 1.6TB (1600GB), 1.91 (1910GB), 3.2TB (3200GB) and 3.82TB (3820GB). The memory is comprised of enterprise Multi-Level Cell (eMLC) 19nm NAND components. The device supports the use of 512-byte and 4096-byte sector sizes.

As used for capacity:

One (1) megabyte (1MB) = One (1) million bytes.

One (1) gigabyte (1GB) = One (1) billion bytes.

One (1) terabyte (1TB) = One (1) trillion bytes.

The available capacity is therefore dependent upon the operating environment and formatting.

2.5 Data Security

The SSD offers erase and sanitization (purge) features. Erase times vary according to the capacity of the drive. The drive can also be “sanitized”, thereby making data recovery impossible. *See Erase Times and Secure Purge.*

2.6 NVMe Support

The SSD is designed as an NVMe device and complies with the latest NVMe specification. The user should consult the latest *NVM Express Revision 1.0a Specification*, published by the *NVM Express Working Group*, *http://www.nvmeexpress.org*.

2.7 Temperature Monitoring via SMBus

Temperature monitoring is implemented using the SMBus interface (In-band/Out-band). The SMBus may switch between the Enterprise SSD Form Factor VPD table and the NVMe Management Interface tabe depending upon the 3.3Vaux/12V power rail configuration. The temperature data is contained in VPD Inquiry Page 0xA6 or NVMe MI 0xD4 respectively. *See Appendix B: Inquiry VDP Pages*.

2.8 UEFI Boot

The SSD supports the use of UEFI (Unified Extensible Firmware Interface). The *UEFI Specification, Version 2.1*, describes the interface between the operating system and platform firmware. The interface as defined by the EFI specification is installed as a set of files (data tables that contain platform-related data, and boot and runtime service calls) on a designated “boot partition” of the drive that are made available to the OS and OS loader, providing a standard environment for booting an operating system and running pre-boot applications.

2.9 Variable Sector Size Support

The SSD supports native 512-byte or 4096 sector sizes with a metadata size of either 0 bytes or 8 bytes to emulate a 520-byte or 4104 sector size emulation for T10 DIF protection.

3. Performance Characteristics

3.1 Overview

The performance rates of the device are dependent upon several factors. The capacity, configuration parameters, traffic patterns and NAND type all impact the performance results. The user should be aware that a device of a specific capacity and configuration may result in lower or higher performance characteristics.

Table 3: Performance Characteristics

	Capacity					
Operation	**800GB**	**956GB**	**1.6TB**	**1.91TB**	**3.2TB**	**3.82TB**
Read Throughput (MB/s, Sequential 128k)	2,600	2,600	3,000	3,000	3,000	3,000
Write Throughput (MB/s, Sequential 128k)	1,400	1,400	1,600	1,600	1,600	1,600
Read IOPS (Max. IOPS, Random 4k)	634K	634K	743K	743K	743K	743K
Write IOPS (Max. IOPS, Random 4k)	80K	24K	140K	38K	140K	38K
Mixed IOPS (70:30 R/W, Random 4k)	190K	77K	310K	115K	310K	115K
Note: The performance values reflect the default 25W power mode option.						

3.2 Mount Time

The time required to initialize and mount the solid-state device varies according to the operating system environment, capacity, and device configuration.

3.3 Latency

The device has no actuator mechanism, read/write heads, or platter. There is no access time or rotational latency. The latency may be affected by the operating system, capacity, and device configuration. Mechanical shock, vibration or gravitational forces do not affect transaction throughput, provided the drive is operating within the environmental specifications. See *Environmental Characteristics*.

3.4 Device Time to Ready

The time to ready performance statistics are dependent on the following factors: 1.) The total drive capacity; 2.) The charge state of the backup capacitors; 3.) The prior failure state according to post graceful and post ungraceful power cycling; 4.) The preconditioned state of the device. Table 4 lists the average time to ready values for the SSDs.

Table 4: Device Time to Ready

Capacity	NAND	Post Graceful Power Cycle	Post Ungraceful Power Cycle
800GB	eMLC	Up to ~2 seconds	Up to ~127 seconds
956GB	eMLC	Up to ~2 seconds	Up to ~127 seconds
1.6TB (1600GB)	eMLC	Up to ~2 seconds	Up to ~127 seconds
1.91TB (1910GB)	eMLC	Up to ~2 seconds	Up to ~127 seconds
3.2TB (3200GB)	eMLC	Up to ~2 seconds	Up to ~127 seconds
3.82TB (3820GB)	eMLC	Up to ~2 seconds	Up to ~127 seconds

3.5 Erase Times

Table 5 lists the maximum time(s) required to erase the unit(s) according to capacity. Every physical memory storage location that can be accessed by the host interface will be erased. It will also destroy any saved data. Once initiated, the SSD will be offline and the host cannot access the device. The SSD will be restarted after the erase. If power ceases during an erase operation, the operation will restart when power is restored.

Table 5: Erase Times

Capacity	NAND	Erase Time
800GB	eMLC	20 minutes
956GB	eMLC	TBD
1.6TB (1600GB)	eMLC	40 minutes
1.91TB (1910GB)	eMLC	TBD
3.2TB (3200GB)	eMLC	80 minutes
3.82TB (3820GB)	eMLC	TBD

3.6 NAND Block Erase Time vs. NAND Wear-Out Erase Time

The typical block erase time for A19nm eMLC NAND is ~7ms at BOL and can increase up to ~20ms EOL due to the expected NAND wear-out process. This erase time, and the increase in erase time with wear, does not result in any performance degradation as the system is designed to handle the worst case scenario.

3.7 Secure Purge

The SSD supports a Secure Purge (Sanitize Erase-Fill) feature. The process erases the data from the drive and then overwrites (fills) each addressable block of memory with a predetermined pattern.

Notes:

1. The security erase algorithms monitor and confirm the completion of the sanitization process.
2. In the event of a power loss, the status of the executing purge is maintained and the purge process must be manually restarted.
3. The interface is ignored while a purge command is being executed and any attempt to reset the drive (hard or soft) will fail.

3.8 Thermal Throttling and Performance

The SSD implements a thermal throttling feature to prevent overheating of the controller and components. The thermal throttling will reduce the IO performance of the device should the temperature exceeds the default threshold of 74°C.

Notes:

1. If the temperature continues to increase and exceeds 74°C, then throttling is enabled and the concurrent operations are reduced proportionally until the SSD begins to return its normal operating temperature range.
2. As the temperature of the SSD begins return to its normal temperature operating range, the throttling algorithm will attempt to proportionally increase the concurrent operations to minimize the impact on the performance.
3. If the temperature continues to increase and reaches 80°C, then the device undergoes a complete shutdown and manual intervention is required to restore operation.

See Component Temperatures and Thermal Throttling.

See Cooling Requirements for the Airflow Guidelines.

4. Reliability Characteristics

4.1 Overview

The following factors affect the reliability statistics:

- Operating/storage temperature(s) will greatly affect power-fail capacitor longevity.
- Operating/storage temperature(s) will greatly affect device longevity.
- Component temperatures are maintained as specified in this document.
- DC power is maintained as specified in this document.
- Errors caused by the host are excluded from the rates.
- Errors from the same causes are counted as 1 block.
- Data stream is assumed random.

4.2 PowerSAFE™ Technology

PowerSAFE™ Technology is an integrated power failure protection system that provides data persistence across any power failure event. If the system power fails, the SSD will transfer the cache contents to the non-volatile flash and restore the contents upon power restoration. The SSD is capable of data persistence regardless of the duration of the power failure event.

A backup circuit monitors the system power state and power-fail capacitors provide the backup power. If a power loss occurs, the data in the cache is flushed to the NAND. The SSD will restore the tables to the cache upon power application; however, any correctable flash errors will require the engagement of the ECC, thereby resulting in a longer restore time. The media commands will be processed after the SSD has rebuilt the tables.

4.3 Secure Array of Flash Elements™ (S.A.F.E.) Technology

Secure Array of Flash Elements™ or S.A.F.E. Technology is a proprietary system implemented in the controller that is designed to overcome the risk of page or block failures over the given lifetime of the flash. An integrated parity protection engine will recover data and automatically relocate it to good known blocks in the flash array when the flash controller encounters uncorrectable data errors.

NAND flash can suffer cascading page and block failures that can produce uncorrectable data errors. These data errors are further affected by flash die geometry reductions that result in potentially higher bit errors in the flash media. The problem is further compounded by the number of flash components that are actually installed in the SSD (e.g. the reliability of the SSD decreases as the number of flash devices increases).

4.4 Bad-Block Management

The bad-block mapping scheme will detect faulty blocks during operation. Bad blocks are flagged in a defect list. Blocks within the defect list are excluded and never used for data storage.

4.5 End-Of-Life Data Retention

End-Of-Life Data Retention is defined as the period for which there is data retained on the device meeting the specified reliability after the maximum rated endurance has been exhausted. The End-Of-Life Data Retention on these devices is three (3) months when stored at 40°C.

4.6 Endurance

The endurance is calculated according to a 100% 4K-aligned transfer random workload, with User/Application Full-Capacity Drive Writes per Day over 5 years, resulting in the following DW/D (Data Writes per Day) outlined in Table 6.

Table 6: Endurance and Capacity

Capacity	NAND	DW/D
800GB	eMLC	3.0
956GB	eMLC	0.8
1.6TB (1600GB)	eMLC	3.0
1.91TB (1910GB)	eMLC	0.8
3.2TB (3200GB)	eMLC	3.0
3.82TB (3820GB)	eMLC	0.8

4.7 Error Detection and Error Correction

The correction rate is 100 bits per 4kB for eMLC NAND. The EDC/ECC algorithm will maintain data integrity by allowing single or multiple bit corrections to the data stored in the flash array. If the data in the flash array is corrupted due to aging or the programming process, EDC/ECC will compensate for the errors to ensure the delivery of accurate data to the host system.

4.8 Hot-Plug Support

The SSD can be inserted or removed during operation. This capability is known as “hot plugging” or “surprise add/remove” capability. An active drive may be removed from a host system at any time, or powered off at any time, during a data transfer, while the flash is being programmed, etc. The error recovery procedures will recover from any errors introduced by hot plugging. The method used by the system to detect insertion or removal of the SSD is dependent upon the scenario. It is recommended that the user reference the *SSD Form Factor Working Group Enterprise SSD Form Factor Specification, Version 1.0, December 20, 2011 for details.*

4.9 Mean Time Between Failures (MTBF)

MTBF for the SSD is calculated at 2.0M Hours with an operating temperature between 5°C ~ 50°C.

4.10 Uncorrectable Bit Errors

Figure 1 illustrates the Uncorrectable Bit Error (UBER) calculation methodology. When all data correction mechanisms are enabled, the error rate will be sustained through all operating temperature ranges as specified in the upcoming sections.

$$\text{UBER} = \frac{\text{Uncorrectable Bit Errors}}{\text{Total Bits Read}} = 1E^{-17}$$

Figure 1: Uncorrectable Bit Error Methodology

4.11 Wear-Leveling

The wear-leveling algorithm guarantees that erase/write cycles are evenly distributed across all of the flash memory block locations to eliminate excessive writes to the same physical flash memory location.

5. Electrical Specifications

5.1 Overview

The SSD requires a +12 power supply; however, users should be aware of the additional 3.3Vaux power and ePERst0# signal specifications.

1. ***Connections to the drive should be made using a Safety Extra Low Voltage (SELV) circuit.***
2. ***The drive uses the +12V and +3.3Vaux supply input pins only. The +5V pins are not connected.***
3. ***The +3.3Vaux is implemented to support the SMBus.***

5.2 12-Volt Power Specifications

Table 7 lists the 12-Volt power specifications for all models of the SSD. Pins P13-P15 supply the+12V power for the SSD; however, the 5V pins P7-P8 and 3.3Vaux pins P1-P3 may be available for other drive types. *See 3.3Vaux Power Specifications*.

Table 7: 12V Power Specifications

Parameter	Value	Definition	Comment
P_{MAX}	25W	Maximum Power	Maximum value.
$12V_{tol}$	±15%	Voltage Tolerance (at pin).	Relaxed from PCIe CEM.
$12V_{amp}$	2.45A	Maximum Continuous Current.	Higher than PCI CEM. See Note 1.
$12V_{peak\text{-}amp}$	4.5A	Maximum Peak Current.	See Note 2.
$12V_{cap}$			See Note 3.
$12V_{drop}$	80mV	Voltage drop across connector.	
SMBus Delay	20ms (min) 1.0s (max)	Delay from 12V being within specification before SMBus access to any SM-Bus slave address other than VPD.	The Enterprise PCIe SSD has up to 1.0s before responding to SMBus transactions.
	20ms (min) 1.0s (max)	Delay from PRSNT# connector mating to an SM-Bus access to any SMBus slave address other than VPD.	

Notes:

1. The Maximum Continuous Current is defined as the highest averaged current value over any one (1) second period.
2. The maximum current to limit connector damage and limit instantaneous power.
3. The maximum capacitance presented by the SSD on the 12V power rail at the backplane connector.

5.3 3.3Vaux Power Specifications

Table 8 lists the +3.3Vaux power specifications for all models of the SSD. The +3.3Vaux is supplied for the SMBus. SMBus access is supported only if the +3.3Vaux is available. The option allows for the query of device information before power-on via VPD (Vital Product Data) over the SMBus.

Table 8: 3.3Vaux Power Specifications

Parameter	Value	Definition	Comment
$3.3VAux_{tol}$	±15%	Voltage Tolerance (at pin).	Relaxed from PCIe CEM.
$3.3VAux_{amp}$	20µA	Maximum Continuous Current; SMBus inactive.	See Note 1.
	1mA	Maximum Continuous Current; SMBus active.	See Note 1.
$3.3VAux_{cap}$	5µF	Maximum Capacitance Load.	See Note 2.
SMBus Delay	20ms (min) 1.0s (max)	Delay from 3.3V being within specification before SMBus access to VPD serial EEPROM may be performed.	The Enterprise PCIe SSD has up to 1.0s before responding to SMBus transactions.
	20ms (min) 1.0s (max)	Delay from PRSNT# and IDet# connector mating to when SMBus access to VPD serial EEPROM may be performed.	

Notes:

1. The Maximum Continuous Current is defined as the highest average current value over any one (1) second period.
2. The maximum capacitance presented by the PCIe SSD on the 12V power rail at the backplane connector.

5.4 ePERst0# Input Current

The ePERst0# (Enterprise Reset Signal) is identical as defined in the PCIe Specification (PCIe-SIG CEM, 2007), but designers should note the difference in the input current as outlined in Table 9.

Table 9: ePERst0# Input Current

Parameter	Value	Definition	Comment
I_{in}	1mA	Reset Pin Input Current	10µA in the PCIe Specification.

5.5 Power Management

The power management and hold-up circuits provide the various voltages, energy storage, and signaling required by the system. The interface has a 12-volt rail that is converted to 3.3V, 2.5V, 1.35V and 1.8V. The main voltage regulation is achieved using switching regulators set for a high efficiency (> 90%) buck configuration.

5.6 Minimum Off-Time

The drive must be powered off for at least 2.0 seconds before it is powered on again.

5.7 Power Backup Circuit

The power backup circuit uses power-good signals and power-failure indicators to determine the status of the host power, hold-up power, and internal system power conditions. The Power Backup circuit provides enough residual power to complete a write sequence, flush the cache, and save all the critical tables. This avoids potential data corruption and the need to rebuild the tables on a subsequent power cycle.

5.8 Start-Up Characteristics

Start-Up electrical characteristics include the start-up current limits, capacitor charge times and drive on-line time. The current, charge time and on-line values are dependent upon the charge state of the capacitors that comprise the power backup circuit. The measurements were made using a nominal 12-volt power supply at 25°C. All measurements reflect the test results of the 3.2TB (3200GB) unit.

5.9 Start-Up Current Limits

Figure 2 shows the 12V current limit with the capacitors in a charging state. The peak current is 1.9A. Table 11 lists the current limit values and capacitor charge-up time.

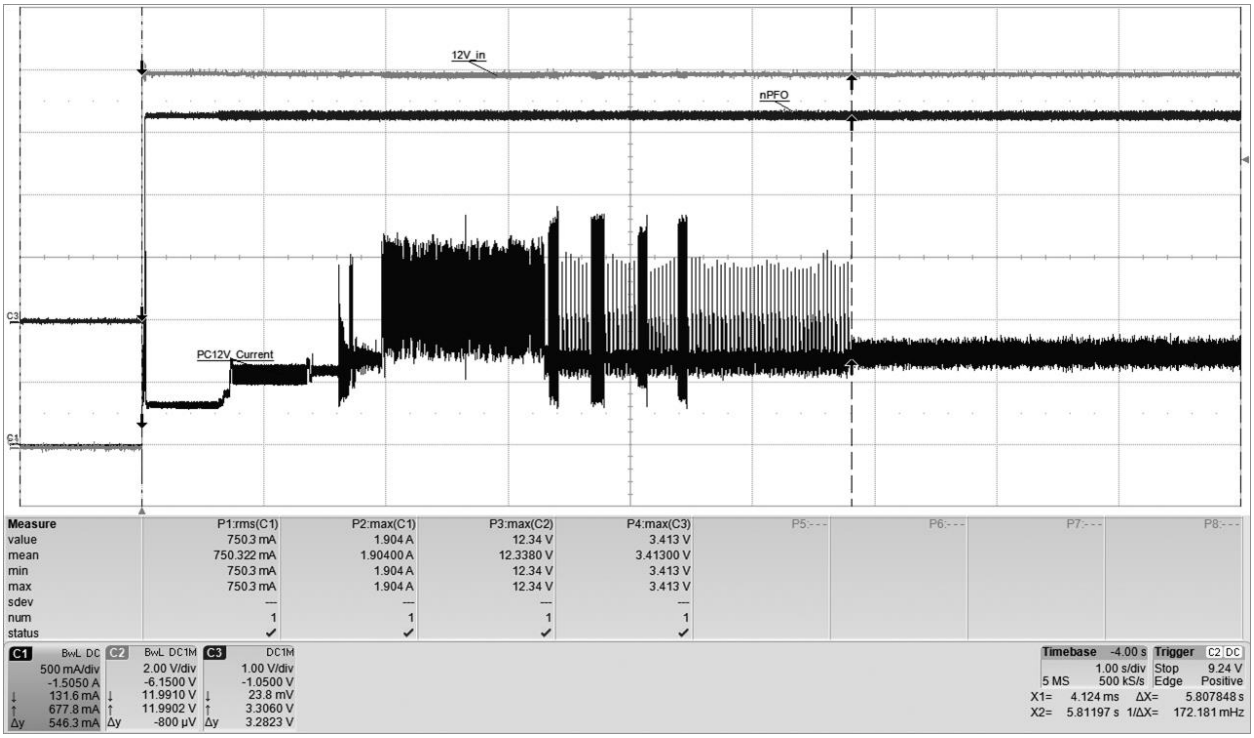

Figure 2: 12V Current Limit - Capacitors in Charging State

Table 10: Start-Up Current Limit and Drive On-Line Time

Maximum Start-Up Current (A)	Maximum Drive On-Line Time (Seconds)
1.90	5.81s

5.10 Backup Capacitor Charge-Up Time

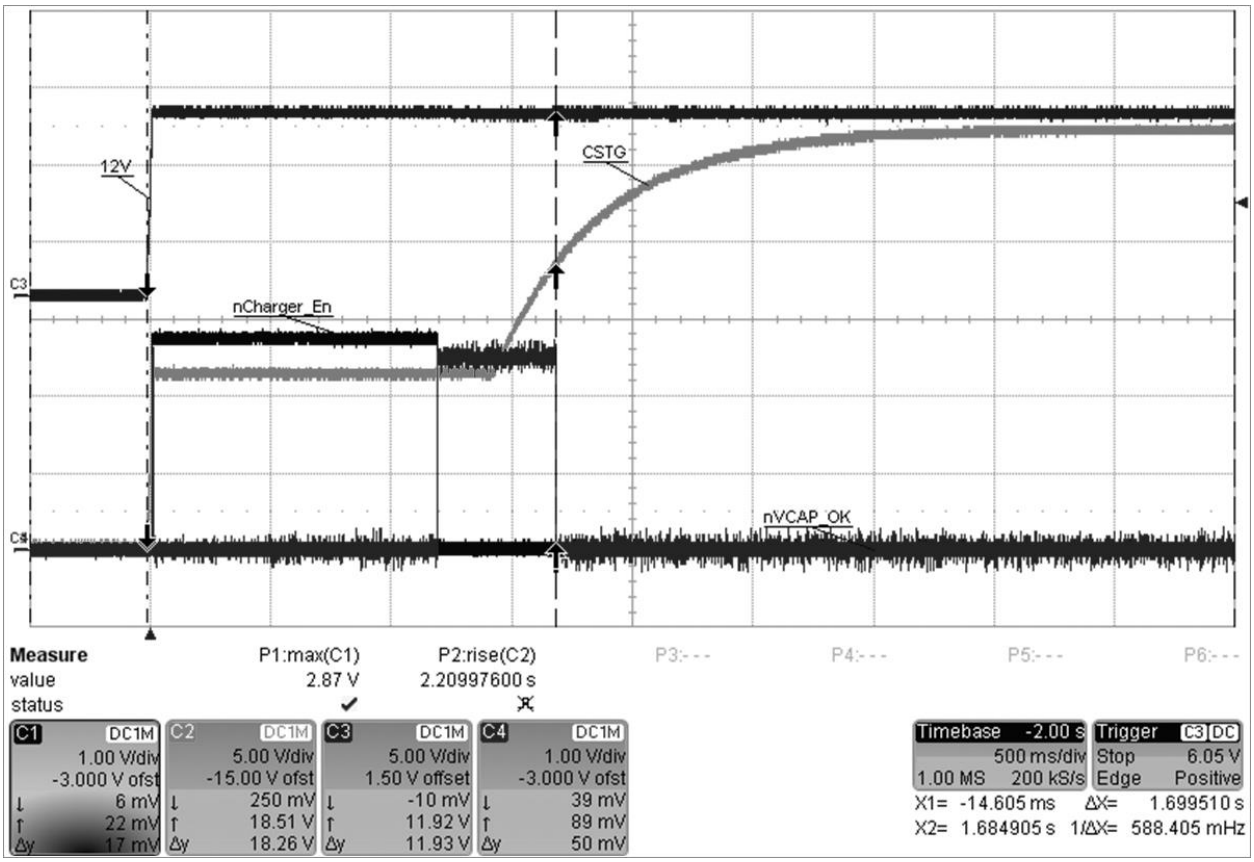

Figure 3: Backup Capacitor Charging State

Table 11: Backup Capacitor Charge-Up Time

Capacitor Charge State	Maximum Charge-Up Time (Seconds)
0V to 28VDC	1.70s

5.11 Backup Circuit Operation (Time, Voltage)

There must be sufficient hold-up time and margin to ensure that the cache is completely flushed to the NAND media during power loss. The resulting backup time is measured from the point when nPFO goes LOW to when POR goes LOW, between the super capacitor initial and final voltages. Figure 4 shows that the backup time was ∆62.50ms.

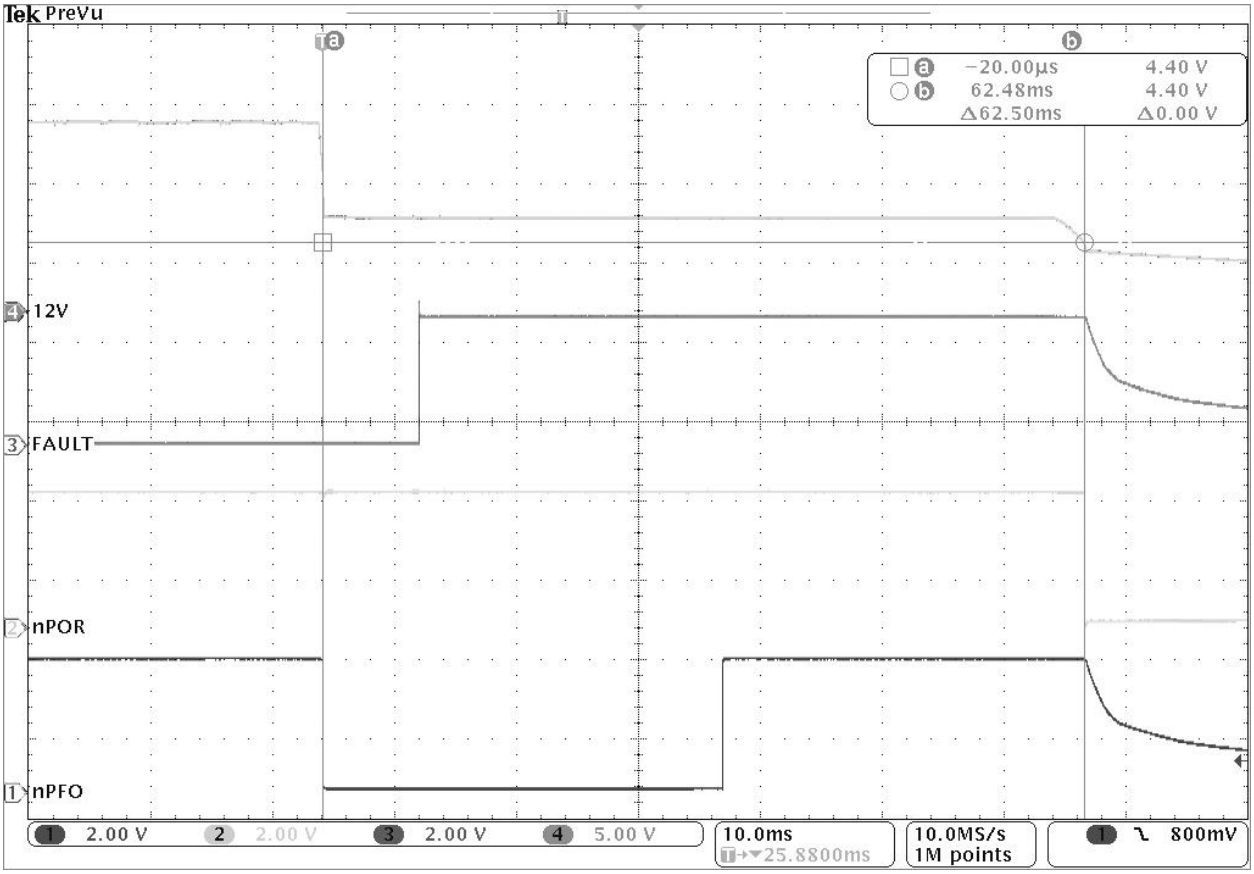

Figure 4: Backup Circuit Operation

5.12 Power Consumption

The amount of power consumed by the device is dependent on the workload profile (operation, transfer size, alignment, queue depth, threads, and temperature). The measurements presented here represent worst-case workload profiles using the nominal input voltage of 12V and a high temperature of 55°C. The power values are subject to change.

5.12.1 800GB Power Consumption

Table 12: 800GB Power Consumption

Operation	Transfer Size	Align	Queue Depth	Threads	AvgMAX (W)
Idle	N/A	N/A	N/A	N/A	7.54
100% Random Read	4K	4K	256	4	16.17
100% Random Write	4K	4K	256	1	14.77
100% Sequential Read	128K	128K	32	1	16.52
100% Sequential Write	128K	128K	32	1	16.89

5.12.2 1.6TB (1600GB) Power Consumption

Table 13: 1.6TB (1600GB) Power Consumption

Operation	Transfer Size	Align	Queue Depth	Threads	AvgMAX (W)
Idle	N/A	N/A	N/A	N/A	7.46
100% Random Read	4K	4K	256	4	17.46
100% Random Write	4K	4K	256	1	19.36
100% Sequential Read	128K	128K	32	1	18.02
100% Sequential Write	128K	128K	32	1	19.61

5.12.3 3.2TB (3200GB) Power Consumption

Table 14: 3.2TB (3200GB) Power Consumption

Operation	Transfer Size	Align	Queue Depth	Threads	AvgMAX (W)
Idle	N/A	N/A	N/A	N/A	7.55
100% Random Read	4K	4K	256	4	17.65
100% Random Write	4K	4K	256	1	19.44
100% Sequential Read	128K	128K	32	1	17.75
100% Sequential Write	128K	128K	32	1	19.22

6. Interface Specifications

6.1 Overview

The SSD is comprised of the following functional blocks: the PCIe x4 Lane interface, an ASIC with an integrated processor, NAND flash memory and DRAM cache. Read/Write data transfer requests are initiated by the host via the PCIe bus interface and are processed by the ASIC. The controller then interfaces with the NAND flash devices and sequences the data flow between the DRAM and flash.

6.2 ASIC

The ASIC is responsible for initiating and controlling all activity within the controller, including bad block mapping and executing the wear-leveling algorithms. The controller decodes incoming host commands and will configure the appropriate interrupts and hardware engines for execution. For read and write transfer commands, there are hardware functions that minimize firmware overhead.

6.3 Read and Write Commands

Read and Write commands have dedicated hardware functions that do not require firmware support. Some commands may require the host controller to use external circuitry that does not involve the flash memory controller. When a Read or Write operation is requested, the integrated controllers sequence the data to and from the integrated DRAM cache.

6.4 User Data Cache

The SSD is configured with a DRAM cache, of which a portion is reserved for the user data cache; the cache is dynamic and will increase in size if space is available. The remainder of the DRAM is reserved for system data.

6.5 Connector Specifications

The connector supports the isolated use of SATA, SATA Express, SAS, and PCIe x4 Lane systems; however, the connector is only populated for the PCIe signaling interface. The connector is divided into P Series (Power) and E Series (ePCI Express or Enterprise PCIe) pin groups. The S Series (SATA/SATA Express x4 Lane/SAS-2) pin group is not discussed. The connector is designed to blind mate, has staggered contacts to facilitate “hot-plugging”, and conforms to the *SFF-8639, Revision 1.1 Standard*. Figure 5 provides a point of reference for the locations of the pin groups.

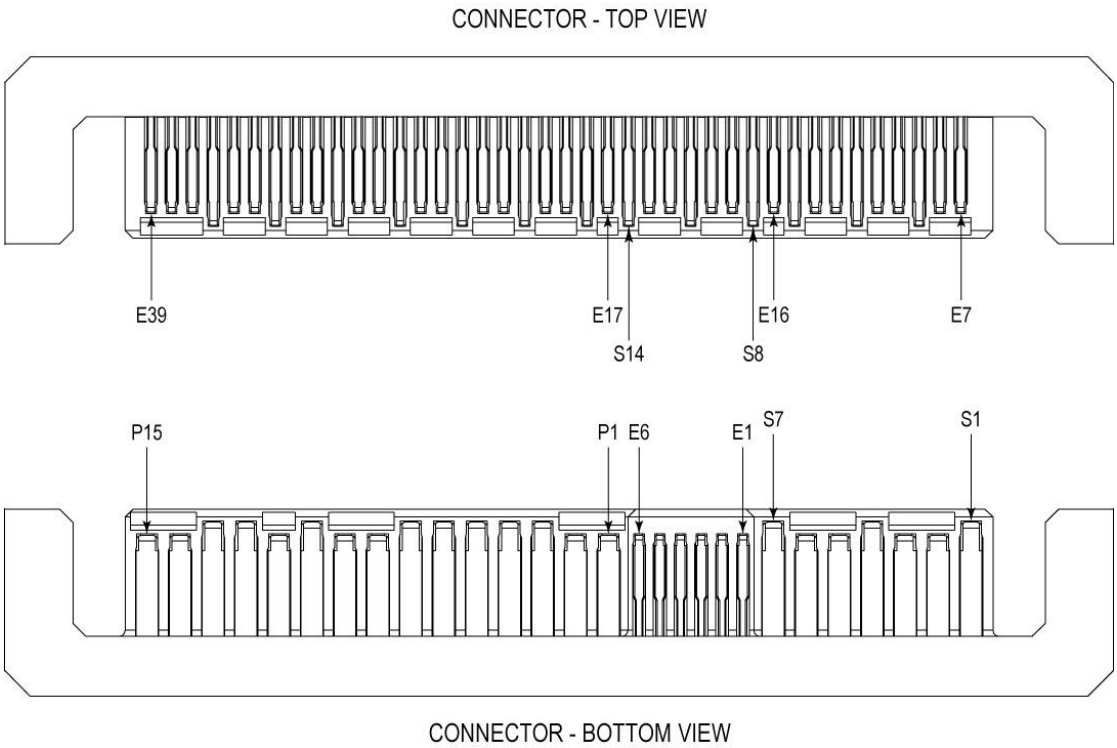

Figure 5: Connector Pin Series Groups

The maximum cable length has not yet been specified for the signal and power segments; seven (7) meter lengths are currently available. HGST, Inc. cannot at this time specify a cabled application; the SSD is designed to blindmate with a SFF-8639 compatible backplane receptacle.

6.6 Connector Dimensions

Figure 6 shows the dimensions of the drive connector. All measurements are in millimeters.

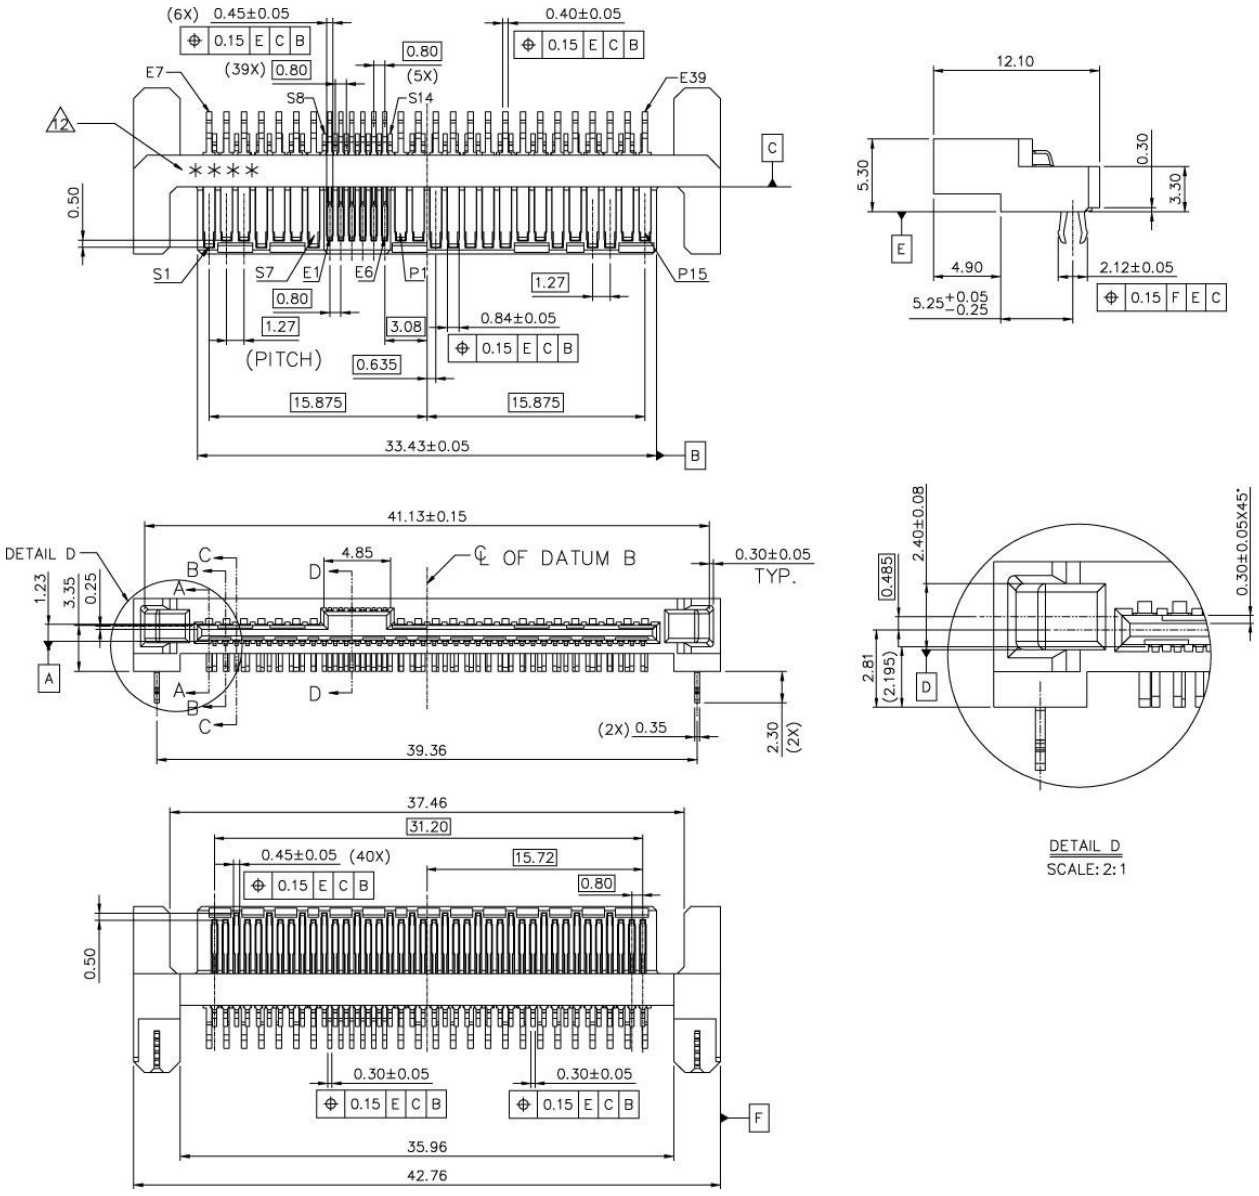

Figure 6: Drive Connector Dimensions

6.7 Signal Assignments

6.7.1 P Series Pin Assignments

Table 15 lists the P Series signal assignments and mate order for the SSD. It is recommended that the user reference the *SFF-8639, Revision 1.1 Standard*.

Table 15: P Series Signals

Pin No.	Mate	Signal Name	Type	Description	Notes
P1	3rd	Wake# (RSVD)	Not Used	No Connect	3
P2	3rd	sPCIeRst	Not Used	No Connect	3
P3	2nd	CLKREQ# (RSVD)	Not Used	No Connect	3
P4	1st	IDet#	Output	Interface-Detect	1, 3
P5	2nd	GND	Power	Ground	3
P6	2nd	GND	Power	Ground	3
P7	2nd	+5V Precharge	Not Used	No Connect	3
P8	3rd	+5V	Not Used	No Connect	3
P9	3rd	+5V	Not Used	No Connect	3
P10	2nd	PRSNT#	Output	Device Present	3
P11	3rd	Activity	Output	I/O Activity Output	2, 3
P12	1st	GND	Power	Ground	3
P13	2nd	+12V Precharge	Power	+12V Precharge	3
P14	3rd	+12V	Power	+12V	3
P15	3rd	+12V	Power	+12V	3

Notes:

1. P4; IDet# or Interface-Detect, is redefined from the SATA and SAS specifications that originally defined P4 as a GND (Ground). It is now defined as an input to host, and is used to identify the interface type of the SSD to the host; pulled to GND within the drive to select the Enterprise PCIe interface.
2. P11; Activity is an open drain output (optional) used for indicating I/O activity. If used, the host must provide an external pull-up to 3.3V.
3. The “Type” is with respect to the drive, not the host.

6.7.2 E Series Pin Assignments

Table 16 lists the E Series signal assignments and mate order for the SSD. It is recommended that the user reference the *SFF-8639, Revision 1.1 Standard*.

Table 16: E Series Signals

Pin No.	Mate	Signal Name	Type	Description
E1	3rd	RefCLK1+	Differential-Pair	PCIe Reference Clock+ (Port B), N/C
E2	3rd	RefCLK1-	Differential-Pair	PCIe Reference Clock- (Port B), N/C
E3	3rd	+3.3Vaux	Power	+3.3V for SMBus
E4	3rd	ePERst1#	Output	PCIe Reset (Port B), N/C
E5	3rd	ePERst0#	Output	PCIe Reset (Port A)
E6	3rd	RSVD1	-	Reserved 1, N/C
E7	3rd	RefCLK0+	Differential-Pair	PCIe Reference Clock+ (Port A)
E8	3rd	RefCLK0-	Differential-Pair	PCIe Reference Clock- (Port A)
E9	2nd	GND	Ground	Ground
E10	3rd	ePCIe0T+	Differential-Pair	PCIe 0 Tx+
E11	3rd	ePCIe0T-	Differential-Pair	PCIe 0 Tx-
E12	2nd	GND	Ground	Ground
E13	3rd	ePCIe0R-	Differential-Pair	PCIe 0 Rx-
E14	3rd	ePCIe0R+	Differential-Pair	PCIe 0 Rx+
E15	2nd	GND	Ground	Ground
E16	3rd	RSVD2	-	Reserved 2, N/C
E17	3rd	RSVD3	-	Reserved 3, N/C
E18	2nd	GND	Ground	Ground
E19	3rd	ePCIe1T+	Differential-Pair	ePCIe 1Tx+
E20	3rd	ePCIe1T-	Differential-Pair	ePCIe 1Tx-
E21	2nd	GND	Ground	Ground
E22	3rd	ePCIe1R-	Differential-Pair	ePCIe 2 Rx-
E23	3rd	ePCIe1R+	Differential-Pair	ePCIe 2 Rx+

Pin No.	Mate	Signal Name	Type	Description
E24	2nd	GND	Ground	Ground
E25	3rd	ePCIe2T+	Differential-Pair	ePCIe 2 Tx+
E26	3rd	ePCIe2T-	Differential-Pair	ePCIe 2 Tx-
E27	2nd	GND	Ground	Ground
E28	3rd	ePCIe2R-	Differential-Pair	ePCIe 2 Rx-
E29	3rd	ePCIe2R+	Differential-Pair	ePCIe 2 Rx+
E30	2nd	GND	Ground	Ground
E31	3rd	ePCIe3T+	Differential-Pair	ePCIe 3 Tx+
E32	3rd	ePCIe3T-	Differential-Pair	ePCIe 3 Tx-
E33	2nd	GND	Ground	Ground
E34	3rd	ePCIe3R-	Differential-Pair	ePCIe 3 Rx-
E35	3rd	ePCIe3R+	Differential-Pair	ePCIe 3 Rx+
E36	2nd	GND	Ground	Ground
E37	3rd	SBMClk	Bi-Directional	SMBus Clock
E38	3rd	SMBData	Bi-Directional	SMBus Data
E39	3rd	DualPortEn#	Output	ePCIe Dual Port Enable

6.7.3 Signal Definitions

Enterprise PCIe is a serial interface that uses LVDS (Low-Voltage Differential Signaling) pairs (one transmit pair and one receive pair). The following table lists the signal definitions.

Table 17: PCIe Signal Definitions

Name	Description
RefCLK0- and RefCLK1+	These are two differential low-voltage swing signals operating at 100MHz (x250MHz). It is recommended that the user reference the *PCI-SIG PCI Express Card Electromechanical Specification, Revision 3.0, Version 0.9, May 23, 2011* for more information.
ePERst0#	The ePERst0# or the Enterprise PCIe Reset signal uses the same logic as defined in the PCIe Specification, and allows the SSD to support hot removal and insertion (hot plugging).
ePCIe*x*T+ and ePCIe*x*T-	Transmitter Pairs. These are the equivalent to PETp*x* and PETn*x*, or the PCIe Lane Tx_*p* and PCIe Lane Tx_*n* differential transmitter pairs of the PCIe bus as defined by the PCIe specification.
ePCIe*x*R+ and ePCIe*x*R-	Receiver Pairs. These are the equivalent to PERp*x* and PERn*x*, or the PCIe Lane Rx_*p* and PCIe Lane Rx_*n* differential transmitter pairs of the PCIe bus as defined by the PCIe specification.
SMBClk	SMBus Clock. It is recommended that the user reference the *System Management Bus (SMBus) Specification 2.0, August 3, 2000*.
SMBData	SMBus Data. It is recommended that the user reference the *System Management Bus (SMBus) Specification 2.0, August 3, 2000*.
DualPortEn#	The DualPortEn# signal allows for dual port configuration. If pulled HIGH, the SSD is configured as a single x4 lane port. If pulled LOW (driven low or grounded to the backplane), then the SSD is configured as a dual x2 port device.

6.8 128b/130b Encoding

The user should reference the *PCIe Express Base Specification 3.0, Version 1.0, November 10, 2010* for more information, specifically *Section 4 Physical Layer Specification, Subsection 4.2.2, Encoding for 8.0 GT/s and Higher Data Rates*. PCIe Express uses 8b/10b encoding when the data rate is 2.5 GT/s or 5.0 GT/s. If the data rates are equal to or greater than 8.0 GT/s, it uses a per-lane code along with a physical layer encapsulation. The discussion of 128b/130b Encoding is beyond the scope of this document.

6.9 LED Indicators

LED indicators are used to annunciate the current conditions of the device. There are three on-board LED indicators: Heartbeat (Green), Activity (Green) and Fault (Yellow). The Heartbeat LED is located on the connector side of the drive; the Activity LED and Fault LED are located opposite the connector side.

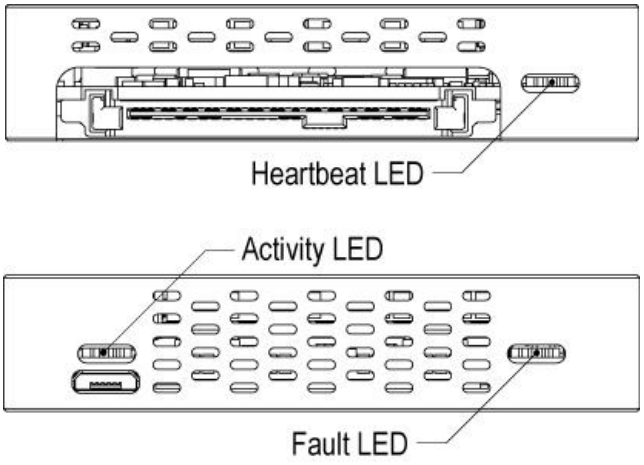

Figure 7: Heartbeat LED, Activity LED and Fault LED

Table 18: LED Descriptions

LED	Color	Description
Heartbeat	Green	The Heartbeat LED indicates that the device is functioning and ready to receive commands. The LED will blink at a steady rate of one (1) second ON, one (1) second OFF.
Activity	Green	The Activity LED indicates any current read/write processing or command activity. The blink rate of the LED is determined by the data rate or the command processing rate. The usual rate is 0.5 seconds ON, 0.5 seconds OFF.
Fault	Yellow	The Fault LED is only driven under operational fault conditions. *See Operational Fault Codes*.

6.9.1 Operational Fault Codes

The following codes indicate that a fault has been detected during operation. The Fault LED will strobe when power is applied, but will de-illuminate once normal operating conditions are established. If a fault occurs during operation, the Fault LED will strobe at a rate of 1Hz with a 50% duty cycle, followed by a 10 second delay. The number of strobes between each 10 second delay is used to indicate the different fault conditions. The unit will remain operational and accept READ commands regardless of the fault, but will reject any WRITE commands. The SSD should be replaced at the earliest possible opportunity.

The fault code should be noted, and if possible, the unit serial number. The user should contact the supplier or technical support.

Table 19: Operational Fault Codes

Fault Code	Action	Condition	Description
Solid Yellow	Replace unit.	Internal Failure	A solid yellow Fault LED indicates that an internal failure has occurred.
1 Flash	Replace unit.	Memory Tune Failed	Memory test failure.
2 Flashes	Replace unit.	Memory Size Failed	Memory test failure.
3 Flashes	Replace unit.	Memory Test Failed	Memory test failure.
4 Flashes	Replace unit.	SERDES Training Failed	SERDES training failure.
5 Flashes	Replace unit.	PCIe Initialization Failure	PCIe initialization failed. A failure occurred while the system tried to initialize the PCIe interface.
6 Flashes	Replace unit.	Power Backup Fault	Power backup circuit failure. The power backup circuit cannot charge the capacitors or the capacitors will not balance. The drive will save all data and prepare for power loss. The data in the cache buffer and the critical tables will be written to the media.
7 Flashes	Replace unit.	SMART Failed	SMART statistics and thresholds cannot be initialized.
8 Flashes	Replace unit.	NAND Write Protect Error	A write protect error was returned when the unit attempted to write to the flash media. If this error occurs continuously then access from the host system will be locked out.
9 Flashes	Replace unit.	Hardware Failure	A hardware failure occurred.
10 Flashes	Replace unit.	NAND Busy Error	A NAND Busy (timeout) error occurred.
11 Flashes	Replace unit.	Power Backup Fault	The power backup circuit has charged the capacitors, but an imbalance condition exists between the individual capacitors.
12 Flashes	Replace unit.	Power Backup Fault	The firmware test routine detected a failure; the capacitors cannot hold a minimum charge.
13 Flashes	Replace unit.	Power Backup Fault	The hold-up time is insufficient. The power backup circuit fails one (1) hour after the test due to the hold-up time being too short.
14 Flashes	Replace unit.	Power Backup Fault	The firmware test routine detected a failure; the capacitors will not hold a charge after multiple recharging attempts.
16 Flashes	Replace unit.	Drive Logging Fault	The drive crash log has not been cleared.
17 Flashes	Replace unit.	Drive Logging Fault	The drive crash log was not cleared in the last run and is still present.

6.9.2 Beacon State

The Beacon State provides the user with a method for locating a specific device within a large installation. The Activity LED will strobe with an “S-O-S” pattern after the host issues a Vendor-Unique Command to the drive; the drive will then enter the Beacon state to visually broadcast its location in the system.

6.9.3 BIST Operation

The Activity LED, Heartbeat LED and Fault LED will all strobe in unison during the execution of a BIST script. There are three states for the LEDs during the execution of the BIST script. The following table outlines the strobe patterns.

Table 20: BIST Script LED Patterns

LED Pattern	BIST State
1Hz	Currently executing BIST script.
10Hz	Completed with a failed status.
5Hz	Completed with a passed status.

6.10 USB Port

The USB Port is used for factory testing and debugging. The port conforms to the *Universal Serial Bus, Micro-USB Cables and Connectors Specification, Revision 1.01, April 4, 2007*. Figure 8 shows the approximate location of the 5-Pin, Type B Micro-USB port on the drive.

The USB port is reserved for use by authorized personnel only. Unauthorized use of the port will result in a violation of the terms of use and will void the warranty.

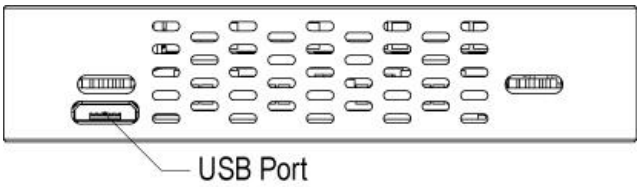

Figure 8: USB Port

7. Physical Characteristics

7.1 Overview

This section discusses the capacities, weights, and physical dimensions relative to the SSD.

7.2 Capacity

The following table lists the available capacities, logical block counts, and corresponding hexadecimal equivalents.

Table 21: Capacities

Block Sector Size	Unformatted Capacity	Logical Block Count	Hexadecimal
512-Byte	8.00166E+11	1,562,813,784	5D26 A558
	9.56166E+11	1,867,570,992	6F50 DF30
	1.60032E+12	3,125,627,568	BA4D 4AB0
	1.91032E+12	3,731,213,808	DE65 CDF0
	3.20063E+12	6,251,233,968	1 749A 9560
	3.82063E+12	7,462,406,448	1 BCCB 4930
4096 Alignment	8.00166E+11	195,351,723	BA4 D4AB
	9.56166E+11	233,446,374	DEA 1BE6
	1.60032E+12	390,703,446	1749 A956
	1.91032E+12	478,611,076	1C87 0684
	3.20063E+12	781,404,246	2E93 52AC
	3.82063E+12	932,800,806	3799 6926

7.3 Weight

The weight of the device varies according to the specific set of design characteristics. The capacity, materials, and form factor all determine the exact weight of the drive. Table 22 lists the approximate weights by capacity.

Table 22: Weights

Capacity	Pounds	Ounces	Grams
800GB	0.366	5.86	166.00
956GB	0.366	5.86	166.00
1.6TB (1600GB)	0.383	6.13	174.00
1.91TB (1910GB)	0.383	6.13	174.00
3.2TB (3200GB)	0.390	6.24	177.00
3.82TB (3820GB)	0.390	6.24	177.00

7.4 Form Factor Dimensions

The overall height for the 2.5-inch form factor is 15mm. Figure 9 shows the drive assembly and maximum case z-height. All measurements are in millimeters.

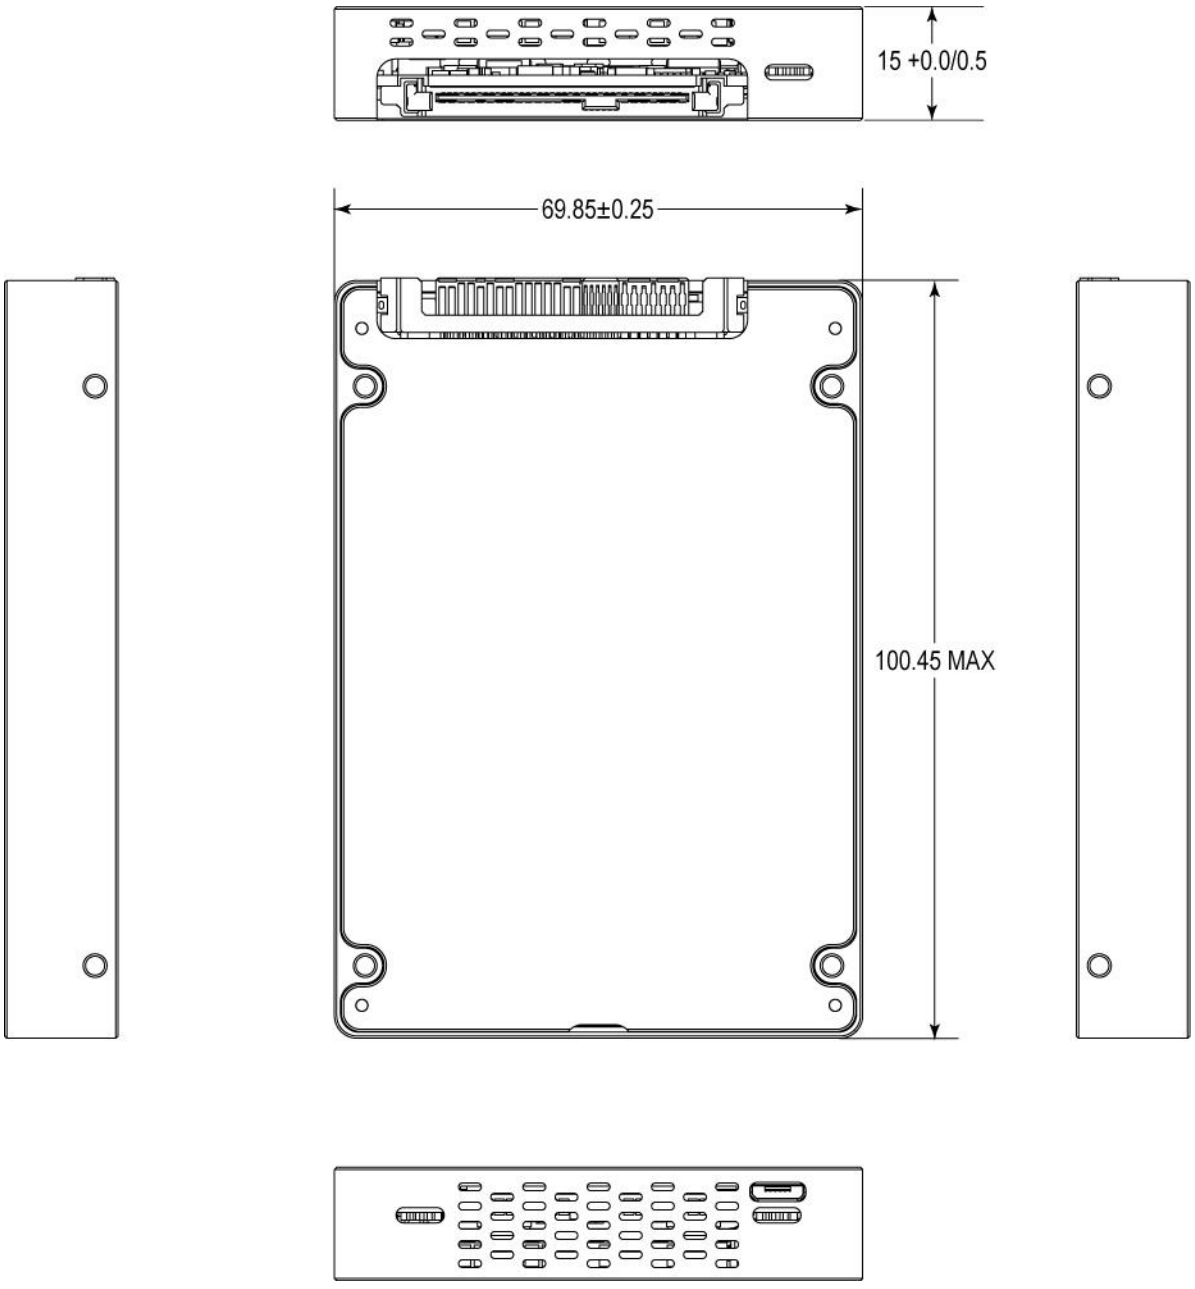

Figure 9: Assembly Dimensions

7.5 Connector Location-Bottom Mounting Screw

Figure 10 shows the relative location of the connector on the 2.5-inch form factor referenced to the bottom mounting screw. It is recommended that the user reference the *SFF-8223 Specification, Revision 2.5*. All measurements are in millimeters.

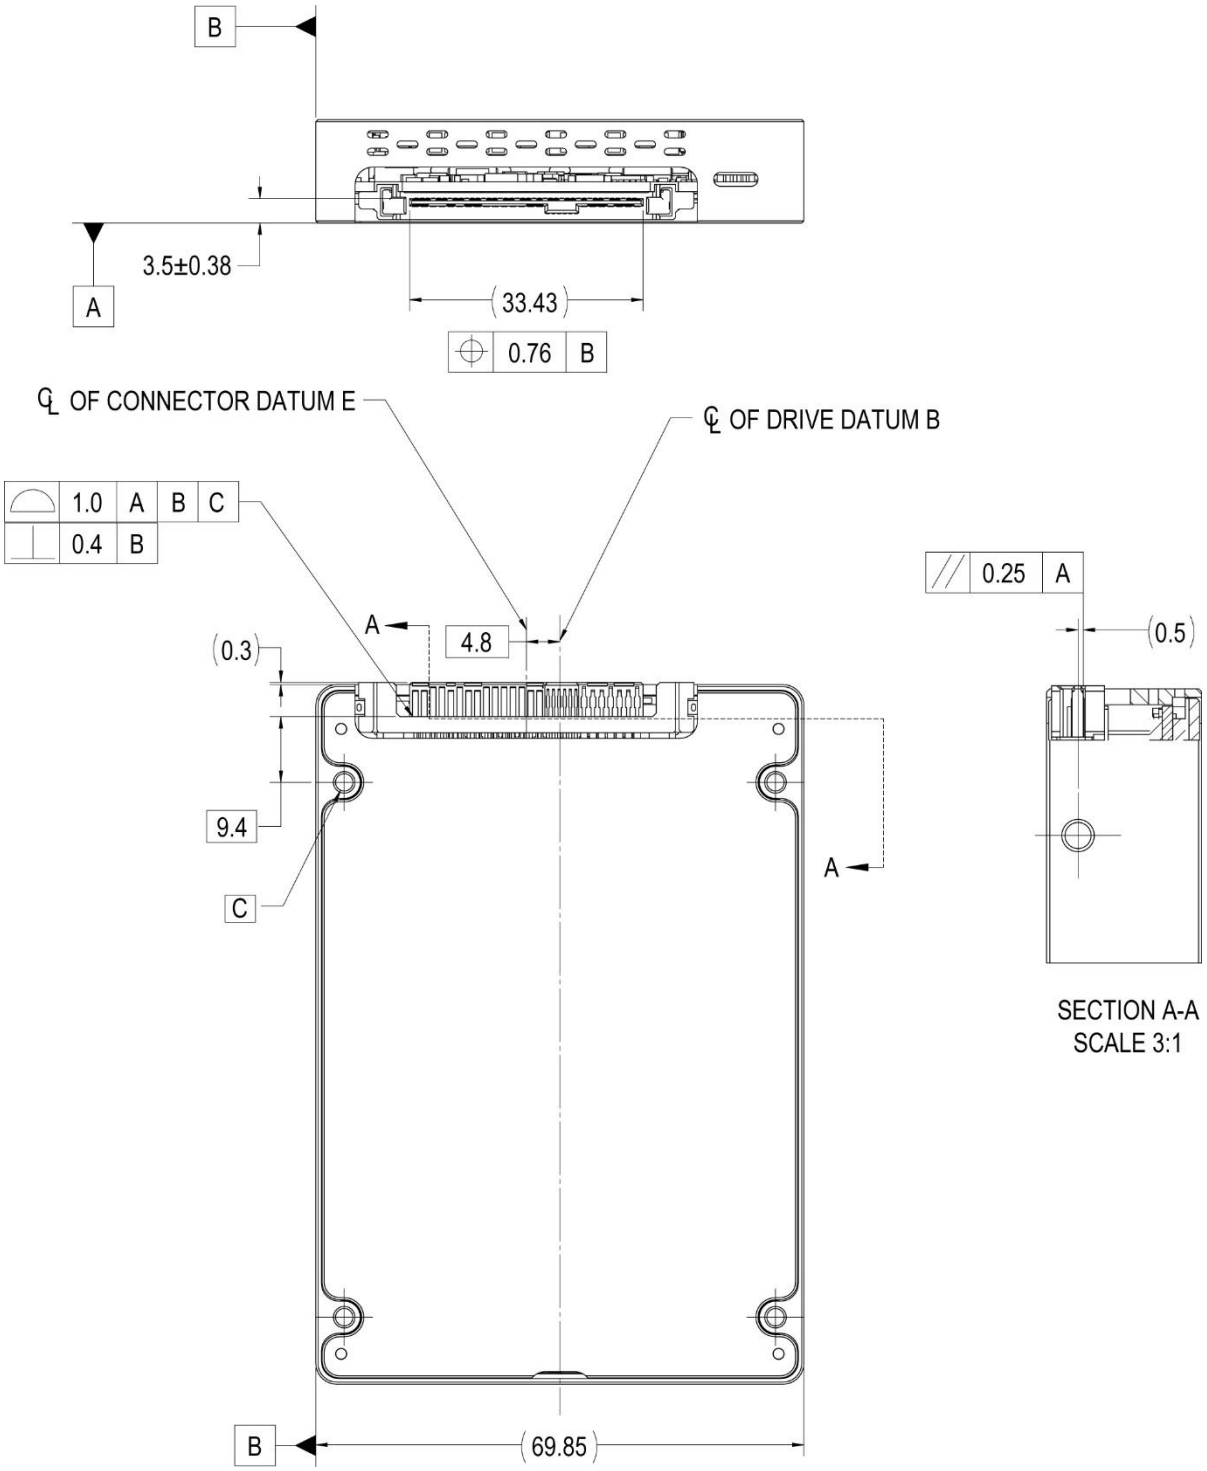

Figure 10: Connector Referenced to Bottom Mounting Screw

7.6 Connector Location-Side Mounting Screw

Figure 11 shows the relative location of the connector on the 2.5-inch form factor reference to the side mounting screw. It is recommended that the user reference the *SFF-8223 Specification, Revision 2.5*. All measurements are in millimeters.

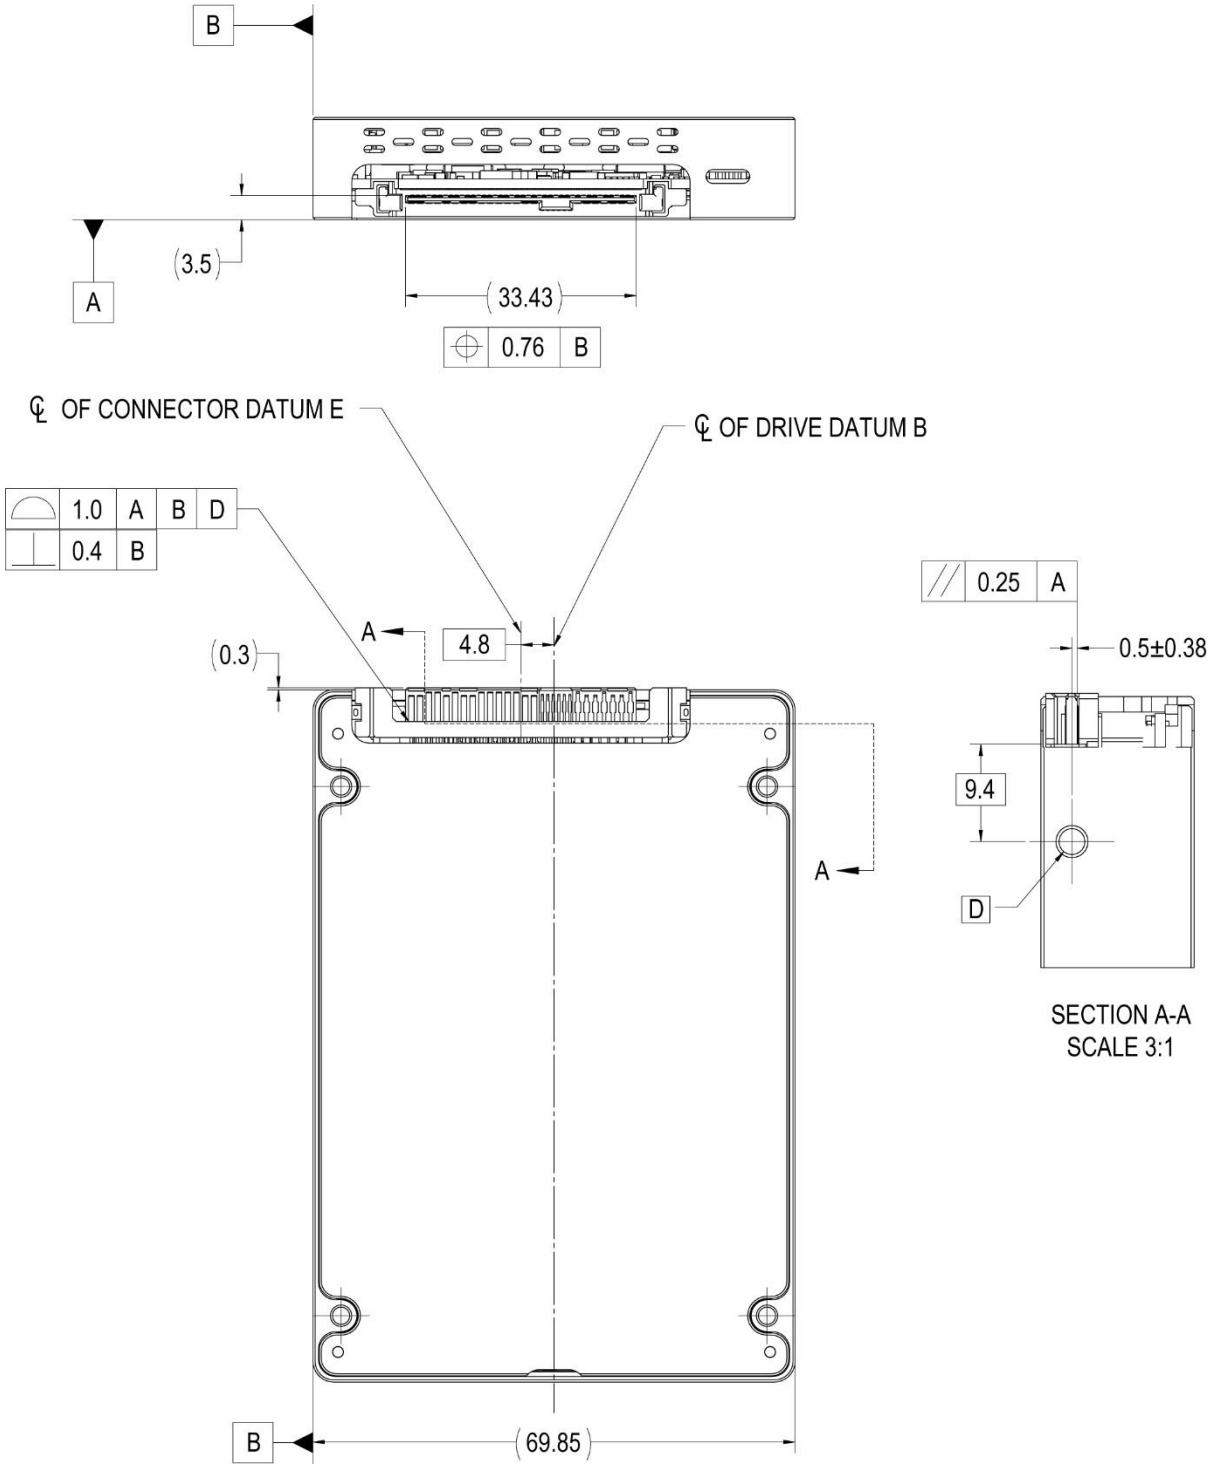

Figure 11: Connector Reference to Side Mounting Screw

8. Environmental Characteristics

8.1 Overview

The SSD is subjected to a series of environmental tests to validate it before shipment.

8.2 Temperature, Humidity and Altitude

The drive will operate to specification within the temperature, relative humidity and altitude conditions as outlined in Table 23.

Table 23: Temperature, Humidity and Altitude Conditions

Operating Conditions	
Ambient Temperature	0°C to 65°C
Relative Humidity	5% to 90%, non-condensing
Maximum Wet Bulb Temperature	29.4°C, non-condensing
Maximum Surface Temperature Gradient	20°C/Hour
Altitude	-305m to 3,048m
Shipping Conditions	
Ambient Temperature	55°C to 95°C
Relative Humidity	5% to 90%, non-condensing
Maximum Wet Bulb Temperature	35°C, non-condensing
Maximum Surface Temperature Gradient	30°C/Hour
Altitude	-305m to 3,048m
Storage Conditions	
Ambient Temperature	0°C to 60°C
Relative Humidity	5% to 90%, non-condensing
Maximum Wet Bulb Temperature	35°C, non-condensing
Altitude	-305m to 3,048m

8.3 Component Temperatures and Thermal Throttling

The SSD has multiple on-board temperature sensors to monitor the critical components. These sensors will trigger the thermal throttling system to prevent damage due to overheating. The combined output of the temperature sensors can be extracted to monitor the drive temperature. The thermal throttling mechanism should not activate when the drive is operating within the operating temperature specification and recommended airflow guidelines. *See Case Temperature*.

8.4 Storage Requirements

The drive or option kit is shipped within a sealed ESD bag by HGST.

8.5 Storage Time

Cumulative storage time in the sealed ESD bag before initial power-on must not exceed one (1) year. The drive must not remain inoperative for longer than six (6) months after it is unpackaged.

8.6 Case Temperature

The drive component temperatures must remain within the limits specified in Table 24. Maximum component temperature ratings must not be exceeded under any operating conditions. The drive may require forced air cooling to meet the specified maximum operating temperature.

See Cooling Requirements under Installation for the airflow guidelines.

Table 24: Maximum Allowable Surface Temperature

Module Name	Location	Maximum Allowable Surface Temperature
Frame Base	See Figure 12.	70°C

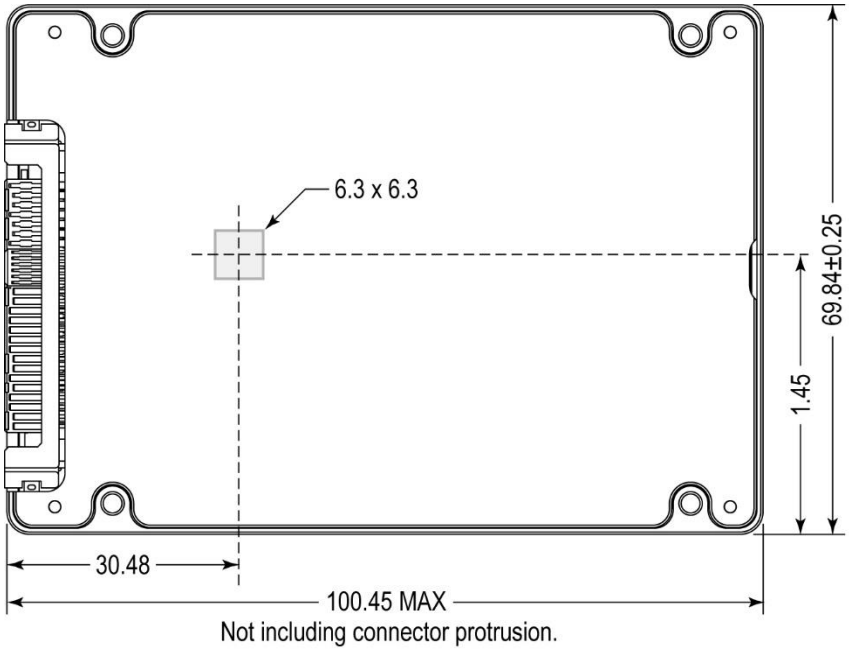

Figure 12: Location of Case Temperature Measurement

Notes:

1. The case temperature measurement device (i.e., thermocouple, thermistor, etc.) should be placed on the bottom of the enclosure nearest the ASIC (approximate location as noted by the 6.3 x 6.3 square in the above figure), allowing for the measurement of the hottest point on the enclosure.
2. The manufacturer cannot specify an ambient temperature and air flow, but only a maximum case temperature as indicated in the above figure.
3. Vendor-Unique Log Page C1h/Subpage Identifier 32h (Temperature History) can be used to verify whether the maximum temperature is being exceeded for a given application.
4. All dimensions are in millimeters (mm).
5. The image is not to scale.

8.7 Acoustics, Vibration and Shock

8.7.1 Acoustics

All SSD models have no acoustics (0dB).

8.7.2 Operating Vibration

8.7.2.1 Random Vibration

The drive is designed to operate without unrecoverable errors while being subjected to the vibration levels as defined in the following subsections. The assessments are determined during thirty (30) minutes of random vibration using the Power Specific Density (PSD) levels as follows:

No Errors: 2.17Grms, 5-700Hz; flat PSD profile for each of the three mutual proportional axes.

Note: The specified levels are measured at the mounting points.

8.7.3 Non-Operating Vibration

The drive will not sustain permanent damage or loss of recorded data after being subjected to the environments described in the following subsections.

8.7.3.1 Random Vibration

The test will consist of random vibration applied to each of the three (3) mutually perpendicular axes at a duration of ten (10) minutes per axis:

3.13Grms, 5-800Hz; flat PSD profile

8.7.4 Operating Shock

The drive will meet the following criteria while operating in the respective conditions as described in the following subsections. The shock pulses of each level are applied to the drive, ten (10) pulses for each direction and for all three (3) mutually perpendicular axes. There must be a minimum of thirty (30) seconds delay between shock pulses. The input level is applied to a base plate on which the drive is attached using four (4) mounting screws.

Shock	Duration	Pulse
20G	11ms	Half-sine wave shock pulse.
100G	0.5ms	Half-sine wave shock pulse.

8.7.4.1 Non-Operating Shock

The drive will not sustain permanent damage or loss of data after being subjected to the environments as described in the following subsection.

8.7.4.2 Half-Sine Wave Shock Pulse

The shocks are applied in each direction of the drive for the three (3) mutually perpendicular axes, one axis at a time. The input level is applied to the base plate on which the drive is attached using four (4) mounting screws.

Shock	Duration	Pulse
100G	11ms	Half-sine wave shock pulse.
150G	10ms	Half-sine wave shock pulse.
1000G	0.5ms	Half-sine wave shock pulse.

9. Installation

9.1 Overview

This section addresses issues regarding compatibility, system requirements, drive configuration, thermal dissipation, drive installation and operating system requirements.

There is a risk of electrocution! Use extreme caution when handling the device and while connecting it to a power source. Make sure to thoroughly read and understand this section before attempting to install the drive.

Observe all applicable electrical safety rules while installing the device.

This device can be damaged by Electrostatic Discharge (ESD). When handling the device, always wear a grounded wrist strap and use a static dissipative surface.

9.2 Compatibility

The SSD can be installed in any operating system environment that supports PCIe 2.x or greater devices. The drive can be connected to a host adapter, motherboard or installed in a PCIe-compliant server array that conforms to the SFF-8639 Specification.

9.3 System Requirements

Make sure the following items are available:

- Four (4) M3 machine screws.
- Optional Enterprise PCIe interface cable. (Maximum cable length not yet specified; seven (7) meter lengths currently available.)
- Optional Enterprise PCIe compatible power cable or adapter.
- Enterprise PCIe connector on the motherboard/backplane or host adapter, JBOD/RAID, flexible backplane or other disk array enclosure that supports the SFF-8639 Specification.
- Available +12V and optional +3.3Vaux DC power sources.
- Operating system or UEFI-aware operating system.

9.4 Drive Configuration

PCIe is a point-to-point, serial bus topology, capable of establishing full-duplex serial links with other PCIe devices via the root complex (host). Each device therefore requires a separate x4 Lane (minimum) PCIe interface on a system board, riser board, or backplane. While the SSD can be installed in any system with a backplane that adheres to the SFF-8639 specification, it is only configured for PCIe x4 Lane operation. The SSD will not function if inserted into a backplane configured for SATA, SATA Express or SAS connections.

9.5 Sector Size Configuration

The drive supports the use of native 512-byte or 4096K-aligned sector sizes with a metadata size of 0 bytes. The 520-byte or 4104K-byte sizes can be emulated by specifying a metadata size of 8 bytes should the user want to implement T10 DIF protection.

9.6 Diagnostic Software

The computer or system manufacturer is responsible for providing any diagnostic software or utilities.

9.7 Connector Requirements

Table 25 lists the drive side connector specifications:

Table 25: Drive Side Connector Requirements

Specification	Description
Insulation	
Material	High Temperature Thermoplastic
Flammability	UL 94V-0
Plating	
Underplating	50µ" Nickel
Wiping Area	30µ" Gold
Solder Tails	75µ" Tin over 50µ" Nickel
Electrical	
Contact Material	Copper Alloy
Low Level Contact Resistance	30 milliohms maximum for signal contacts.
Insulation Resistance	1000 Mega Ohms minimum
Dielectric Withstanding Voltage	No breakdown or flashover @500 VAC/1 minute.
Temperature Rise, Power Pins	Temperature Rise shall not exceed 30°C.
Environmental	
Operating Temperature	-40°C to +125°C
RoHS Compliance	Lead Free

9.8 Drive Orientation

The SSD can be installed in any number of orientations within the enclosure. The drive will operate and meet all the requirements as outlined in this specification regardless of mounting orientation.

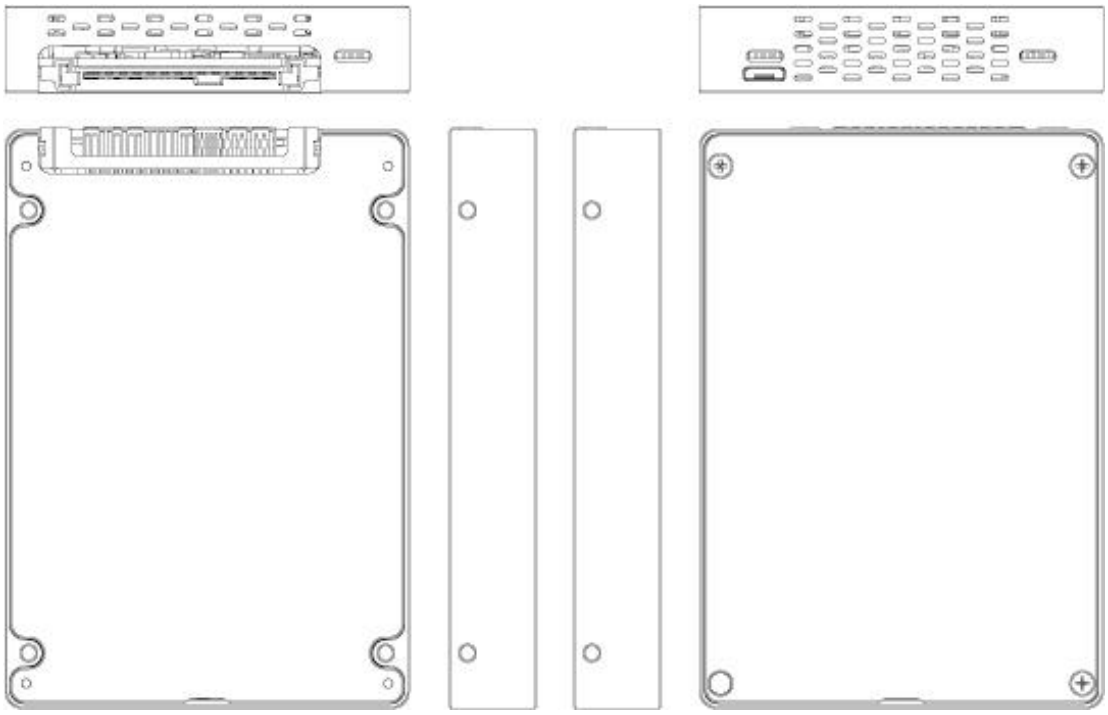

Figure 13: Possible Installation Orientations

9.9 Primary Heat Generation Area

Figure 14 indicates the approximate location of the primary heat generation area on the device. *See Cooling Requirements*.

Exercise caution when handling the drive after extended operation. The heat generated by the internal circuitry can be substantial.

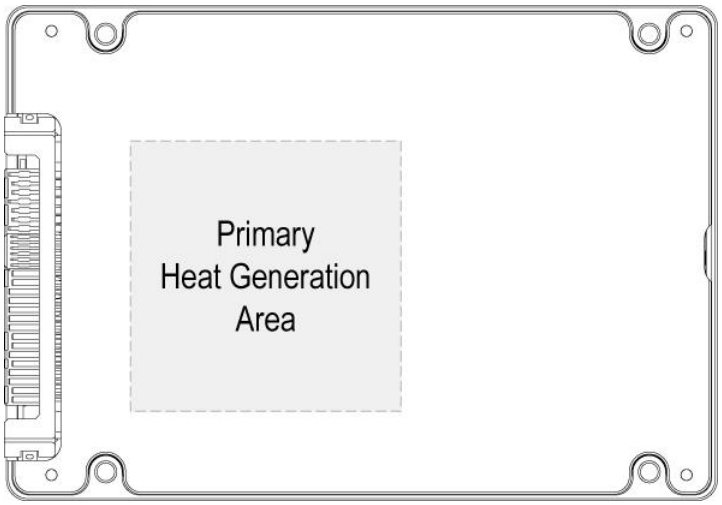

Figure 14: Primary Heat Generation Area

9.10 Cooling Requirements

The host enclosure may remove heat by conduction, convection, or other forced airflow to maintain the required operating temperature range. The optimal air flow should be from front-to-back or back-to-front, as illustrated in Figure 15.

√ *The user should be aware of the following conditions:*

1. *The internal drive temperature, as measured by the temperature sensor(s), should not exceed 75°C. The user should contact the manufacturer should the internal case temperature constantly spike above 75°C.*
2. *If forced air flow is used, an example of a recommended, sufficient flow is 3m/s, front-to-back, with an ambient temperature of 35°C.*
3. *The absolute case surface temperature is 70°C. See Case Temperature under Environmental Characteristics.*

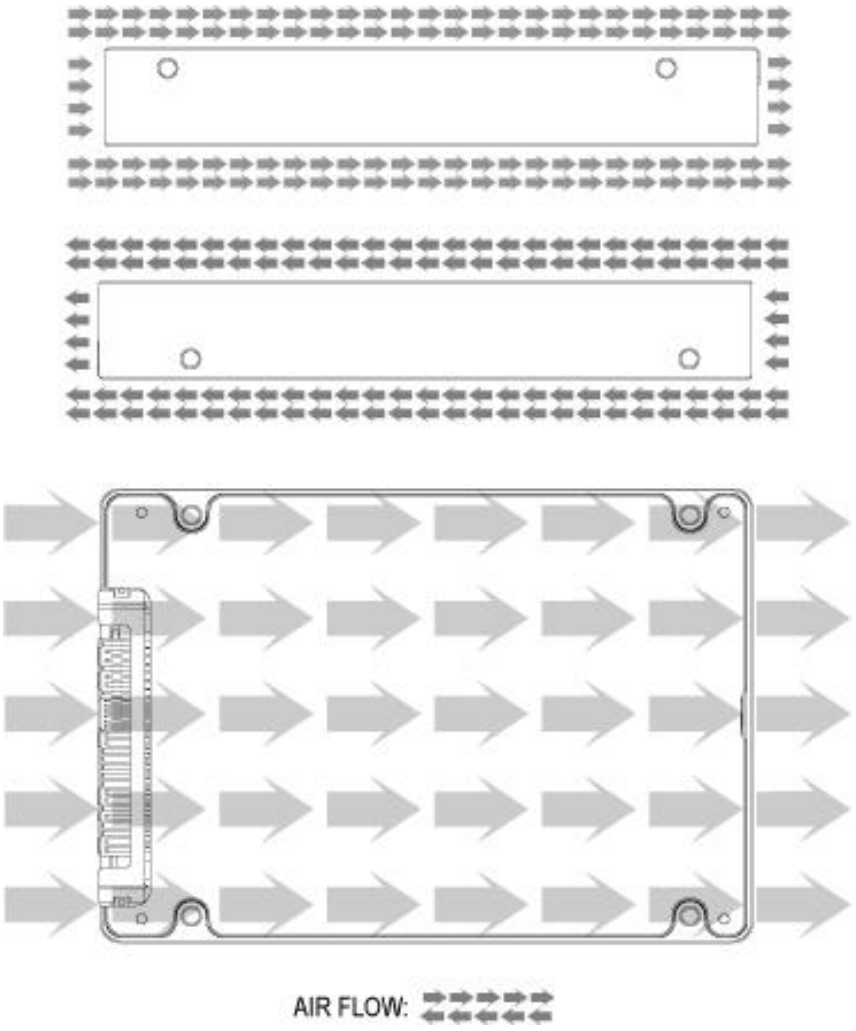

Figure 15: Optimal Air Flow Pattern for Cooling

9.11 Mounting Requirements

Figure 16 shows the relative locations of the mounting holes for the 2.5-inch form factor. Careful attention should be made to the length of the mounting screws and the recommended maximum torque to prevent damage to the enclosure. All measurements are in millimeters. *See* *Drive Installation*.

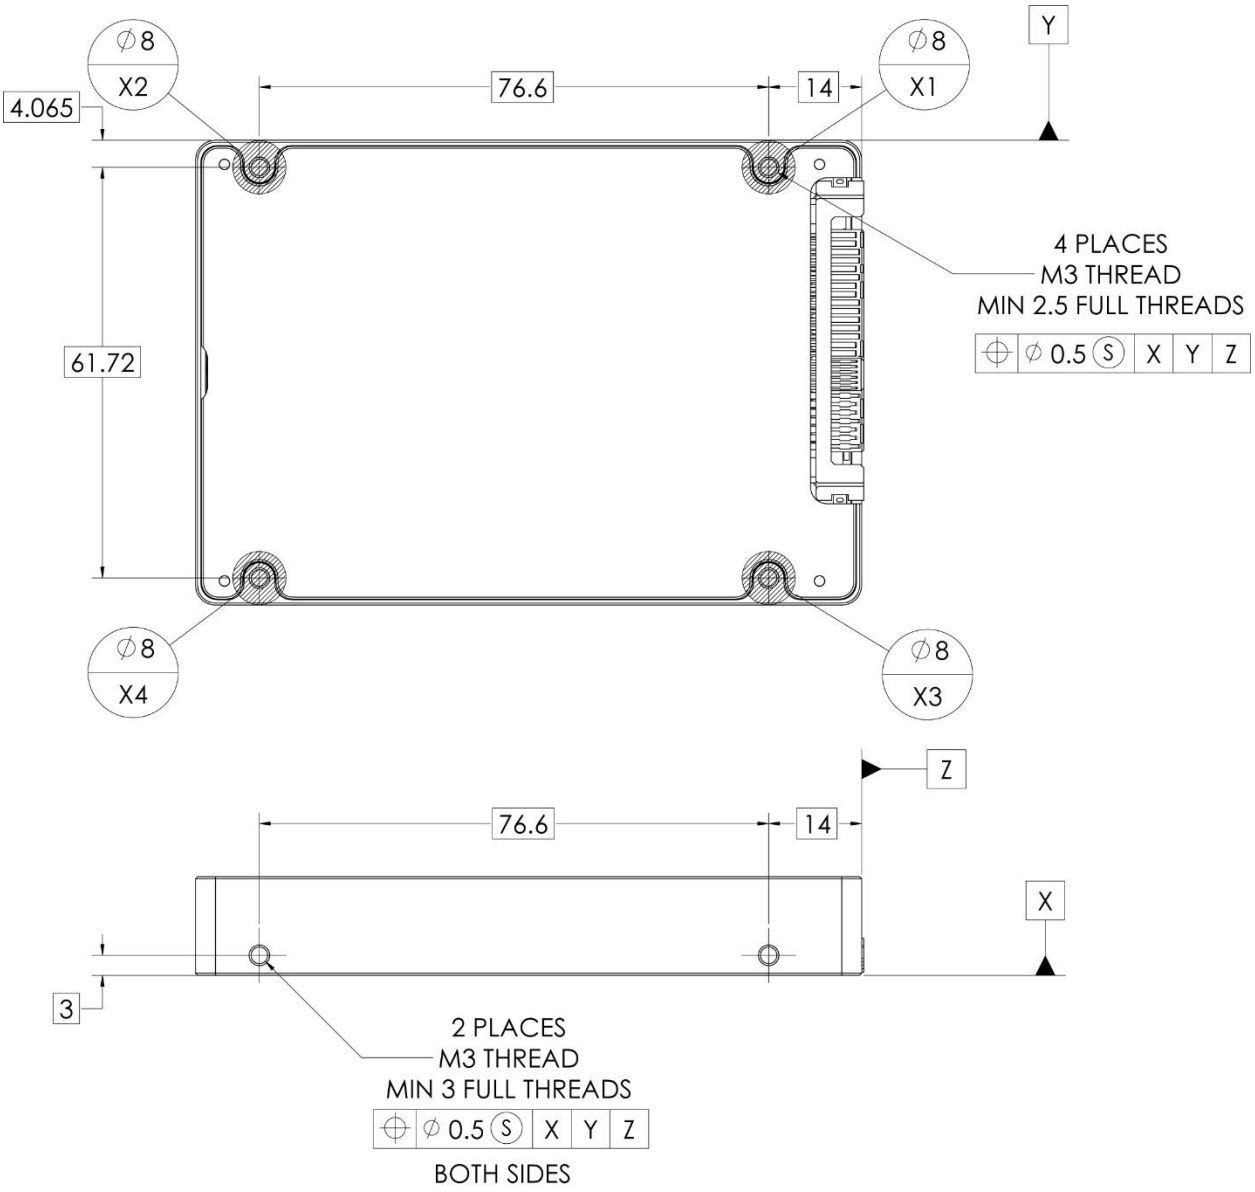

Figure 16: Mounting Specifications

Notes:

1. Screw Torque: 0.395±0.05 N-m (3.5±0.5 lbf-in) MAX screw torque.
2. OEM or user must determine best torque specification according to application.
3. OEM or user must determine best screw engagement according to application.

9.12 Drive Installation

There is a risk of electrocution! Use extreme caution when handling the device and while connecting it to a power source. Make sure to thoroughly read and understand this section before attempting to install the drive.

Observe all applicable electrical safety rules while installing the device.

Electrostatic Discharge or ESD can seriously damage the electronics of the host system, enclosure and device. It is very important that the user discharge any static electricity before starting the installation procedure. The user can touch an unpainted, grounded metallic surface to discharge any static charges that may be present on the body or clothing. As an alternative, the user can also wear an ESD protective wrist strap. The user can minimize the possibility of damage due to ESD by avoiding physical contact with the electronic components and interface connector.

The maximum cable length has not yet been specified for the signal and power segments; seven (7) meter lengths are currently available. The manufacturer cannot at this time specify a cabled application; the SSD is designed to blindmate with a SFF-8639 compatible backplane receptacle.

To install the SSD in a host system or alternate enclosure:

1. ***Power off the system***.
2. If necessary, remove the access cover.
3. ***Make sure the SSD is in the correct orientation before inserting it into the drive bay or attaching it to the mounting surface.***
4. ***The connector on the drive is keyed to ensure that the signal and power connections to the drive are correctly oriented***.
5. ***Do not force the SSD into the drive bay or into the backplane connector***. Doing so may damage the drive, the interface connector or the enclosure beyond repair.
    1. The SSD is designed to blindmate with a backplane receptacle that adheres to the *SFF-8639 Specification*.
    2. Position the SSD in an available drive bay or choose a suitable mounting location.
    3. Blindmate the SSD against the backplane connector or connect the SSD using a cable that complies with the *SFF-8639 Specification*.
6. ***Do not over tighten the mounting screws, but apply sufficient torque to ensure the drive is secure***. Secure the SSD in the drive bay or attach it to the mounting surface using the recommended machine screws.
7. If necessary, replace the access cover.
8. Power on the system.

10. Operating System Specifications

10.1 Overview

The SSD is compatible with Microsoft Windows, Linux Distributions, and VMware virtual environments. The device is low-level formatted at the factory (00h in all data sectors); however, it must be partitioned and high-level formatted. The device can be formatted as a boot or data storage drive using any standard disk partitioning and formatting utility.

10.2 UEFI Boot Configuration

UEFI boot can be implemented. UEFI is installed as a set of files (data tables, boot and runtime services) on a designated “boot partition” of the drive. These files are made available to the OS loader upon system power-on; the OS must be U(EFI)-aware. Usually, the default location for the OS loader is:

<EFI_SYSTEM_PARTITION>/BOOT/BOOT/<MACHINE_TYPE_SHORT_NAME>.EFI

Where <MACHINE_TYPE_SHORT_NAME> can be IA32, X64, IA64, ARM or AA64. The user should be aware that OS vendors may have proprietary boot loaders or modify the default boot location.

10.3 Microsoft Compatibility

The SSD is compatible with most Microsoft operating systems and has been verified on the following: using the native drivers supplied with the OS:

- Windows Server 2008 R2, 64-Bit
- Windows Server 2012 R2, 64-Bit

10.4 Linux Distributions

The device is compatible with Linux Distributions with PCIe support and has been verified on the following:

- CentOS 6.4, 6.5, 7.0
- Linux Distributions, Kernel 2.6.3 or higher.
- RHEL 6.4, 6.5, 7.0
- SLES 11 SP3
- SLES 12

11. Command Set Specification

11.1 Overview

NVM Express, or NVMe, is a register level interface. This approach is intended to simplify the management and scalability of NVMe devices installed in a system. The NVMe standard is comprised of various terms that are important to know when implementing NVMe devices in a system. This section will present the basic terms within context to help provide an intuitive understanding of the NVMe interface. The user should consult the cited NVME specification for details

11.2 NVMe Admin Command Set

Commands that can be submitted to the Admin Submission Queue belong to the Admin Command Set. As implied, the commands belonging to this set are used for device administration.

Table 26: NVMe Admin Command Set

Opcode	Command	Opcode (07)	Opcode (06:02)	Opcode (01:00)	Namespace Identifier Used
		Generic Command	Function	Data Transfer	
00h	Delete I/O Submission Queue	0b	000 00b	00b	No
01h	Create I/O Submission Queue	0b	000 00b	01b	No
02h	Get Log Page	0b	000 00b	10b	Yes
04h	Delete I/O Completion Queue	0b	000 01b	00b	No
05h	Create I/O Completion Queue	0b	000 01b	01b	No
06h	Identify	0b	000 01b	10b	Yes
08h	Abort	0b	000 10b	00b	No
09h	Set Features	0b	000 10b	01b	Yes
0Ah	Get Features	0b	000 10b	10b	Yes
10h	Firmware Activate	0b	001 00b	00b	No
11h	Firmware Image Download	0b	001 00b	01b	No
80h	Format NVM	1b	000 00b	00b	Yes
0Ch	Asynchronous Event Request	0b	000 11b	00b	No

11.3 NVMe I/O Command Set

The NVMe I/O Command Set is used for data manipulation. The commands read, write, compare or otherwise manage the actual data stored on the non-volatile media.

Table 27: NVM I/O Command Set

Opcode	Command	Opcode (06:02) Function	Opcode (01:00) Data Transfer
00h	Flush	000 00b	00b
01h	Write	000 00b	01b
02h	Read	000 00b	10b
04h	Write Uncorrectable	000 01b	00b
05h	Compare	000 01b	01b
08h	Write Zeros	000 10b	00b
09h	Dataset Management	000 10b	01b

11.4 Log Page Support

Table 28 lists the following mandatory log pages supported by the SSD as defined in the *NVMe 1.1a Specification*. *See the NVMe 1.1a Specification, Section 5.10.x.x*.

Table 28: NVMe Log Pages

Log Identifier	NVMe 1.1a Section	Description
00h	5.10.1	Reserved
01h	5.10.1.1	Error Information
02h	5.10.1.2	SMART / Health Information
03h	5.10.1.3	Firmware Slot Information
04h – 7Fh		Reserved
80h – BFh		I/O Command Set Specific.
C0h – FFh		Vendor-Specific.

11.4.1 Get Log Page (02h)

The Get Log Page (02h) command will return a data buffer of the requested log page. The command uses PRP Entry 1, PRP Entry 2, and the Command Dword 10 fields.

11.4.2 Log Identifier 01h – Error Information

The Error Information log page (Log Identifier 01h) may return the last *n* errors as compiled by the device. If the host specifies a data transfer of the size of *n* error logs, then the error logs for the last *n* errors are returned. The order for each entry is according to the time when the error occurred, with the most recent error being returned as the first log. The log page is a set of 64-byte entries and is global to the device.

11.4.3 Log Identifier 02h – SMART / Health Information

S.M.A.R.T. is an acronym for *Self-Monitoring, Analysis and Reporting Technology*. The specification has been adapted so that PCIe devices can report drive health, status and SMART statistics using Get Log Page (02h) / SMART Attributes (Log Identifier 02h). The information that is provided reflects the life of the controller and is retained across power cycles. The log page is supported on a global basis. The log page may be supported on a per namespace basis. A global log page may be requested by specifying a namespace of FFFFFFFh.

11.4.4 Log Identifier 03h – Firmware Slot Information

The Firmware Slot Information (Log Identifier 03h) log is used to return the firmware revision information stored in each supported firmware slot. The firmware is indicated as an ASCII string. The log page also returns the active slot number.

11.4.5 NVMe Vendor-Unique Log Pages

Log Page Identifiers 0xC0 through 0xFF are Vendor-Specific. The subpages of Log Page 0xC1 (C1h) contain vendor-unique information. The subpages are sequentially listed in 0xC1, and are associated with a Subpage Identifier. Table 29 lists the currently supported subpages and identifiers.

Table 29: Log Page 0xC1 and Subpages

Opcode DW0 [07:00]	CMD DW10 [07:00]	SUBCMC DW10 [27:16]	Namespace ID Used	Subpage Name	Subpage Identifier
0x02	0xC1	≤ Size of Page, 0s Based	No	Write Errors	0x02
0x02	0xC1	≤ Size of Page, 0s Based	No	Read Errors	0x03
0x02	0xC1	≤ Size of Page, 0s Based	No	Verify Errors	0x05
0x02	0xC1	≤ Size of Page, 0s Based	No	Self-Test Results	0x10
0x02	0xC1	≤ Size of Page, 0s Based	No	Background Scan	0x15
0x02	0xC1	≤ Size of Page, 0s Based	No	Erase Errors	0x30
0x02	0xC1	≤ Size of Page, 0s Based	No	Erase Counts	0x31
0x02	0xC1	≤ Size of Page, 0s Based	No	Temperature History	0x32
0x02	0xC1	≤ Size of Page, 0s Based	No	SSD Performance	0x37
0x02	0xC1	≤ Size of Page, 0s Based	No	Other Statistics	0x38

11.5 Get Features (0Ah) and Set Features (09h)

Table 30 lists the features supported by the SSD. The Get Features (0Ah) command can be used to extract and report features affecting the health and operation of the SSD to the host. The Set Features (09h) command can be used to define the settings. *See the NVMe 1.1a Specification, Section 5.12.*

Table 30: Feature Identifier Reference

Feature Identifier	NVMe 1.1a Section	Description
01h	5.12.1.1	Arbitration
02h	5.12.1.2	Power Management
03h	5.12.1.3	LBA Range Type
04h	5.12.1.4	Temperature Threshold
05h	5.12.1.5	Error Recovery
06h	5.12.1.6	Volatile Write Cache
07h	5.12.1.7	Number of Queues
08h	5.12.1.8	Interrupt Coalescing
09h	5.12.1.9	Interrupt Vector Configuration
80h	5.12.1.13	Software Progress Marker
0Ah	5.12.1.10	Atomic Write Unit Normal
0Bh	5.12.1.11	Asynchronous Event Configuration
0Ch	5.12.1.12	Autonomous Power State Transition

12. Label Specifications

12.1 Overview

The following certification marks may appear on the labels that are affixed to every drive shipped from the manufacturing location in accordance with the appropriate drive assembly drawing. These labels may be integrated with other labels.

12.2 Manufacturer Identification

A label with the “HGST, a Western Digital Company” logo, HGST model number and the statement “Made by HGST,” or HGST approved equivalent.

12.3 Product Identification

- A label having the drive model number, manufacturing date, formatted capacity, and country of origin or HGST approved equivalent.
- Certification Marks: BSMI, CE, CSA / cUL, EIP, FCC, KCC, RCM, TUV, UL, VCCI and WEEE.
- A bar code label symbolizing the drive serial number.
- As per agreement, a user designed label.
- Interface Definition Mark: PCIe

BSMI (Taiwan)

CE Mark (EU/EEA)

CSA, cUL, US

EIP (China RoHS)

FCC Mark (USA)

KCC (Korea)

RCM (Australia)

TUV

VCCI (Japan)

WEEE Directive

13. Electromagnetic Compatibility

13.1 Overview

The drive, when installed in a suitable enclosure and exercised with a random access routine at a maximum data rate will comply with the worldwide EMC requirements listed below.

The drive is designed for system integration and installation into a suitable enclosure for use. As such, the drive is supplied as a subassembly and is not subject to Subpart B of Part 15 of the FCC Rules and Regulations.

The design of the drive serves to minimize radiated emissions when installed in an enclosure that provides reasonable shielding. As such, the drive is capable of meeting FCC Class B limits; however, it is the responsibility of the user to ensure that the drive meets the appropriate EMC requirements in their system. The use of shielded I/O cables may be required if the enclosure does not provide adequate shielding, with the shields grounded to the enclosure and to the host computer.

13.2 Radiated and Conducted RF

Standard	Nation/Region
CISPR22:2009/Am1:2010	Australia, New Zealand
CNS 13438:2006	Taiwan
EN 55022:2010	EU
FCC Title 47 Part 15	USA
GB9254-2008	China
ICES-003, Issue 5, 2012	Canada
VCCI V-3/2013-04	Japan
KN 22:2013-3 (RRA Notice)	Korea
KN 22:2013-24 (RRA Notice)	Korea

13.3 ITE Immunity

Standard	Nation/Region
EN 55024:2010	EU
KN 24:2013-4 (RRA Notice)	Korea
KN 24:2013-25 (RRA Notice)	Korea

13.4 Power Line Harmonic Emissions

Standard	Nation/Region
EN61000-3-2:2006 Am1:2009, Am2:2009	EU
GB17625.1 2003	China

13.5 Voltage Fluctuations and Flicker

Standard	Nation/Region
EN 61000-3-3:2008	EU
GB 17625.2 1999	China

13.6 Immunity Specifications

Standard	Specification
KN 61000-4-2:2013-06	Electrostatic Discharge (ESD) Immunity
KN 61000-4-3:2011-10	Radiated RF Immunity
KN 61000-4-4:2011-10	Electrical Fast Transient/Burst (EFT/B) Immunity
KN 61000-4-5:2008-05	Surge Immunity
KN 61000-4-6:2013-06	Conducted RF Immunity
KN 61000-4-8:2013-06	Power Frequency Magnetic Field Immunity
KN 61000-4-11:2008-05	Voltage Dips and Interruptions Immunity

13.7 Class B Regulatory Notices

13.7.1 European Union

The product conforms with the protection requirements of EU Council Directive 2004/108/EC on the approximation of the laws of the Member States relating to electromagnetic compatibility. HGST cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-HGST option cards.

This product has been tested and found to comply with the limits for Class B Information Technology Equipment according to European Standard EN 55022. The limits for Class B equipment were derived for typical residential environments to provide reasonable protection against interference with licensed communication devices.

13.7.2 Canada

This Class B digital apparatus complies with Canadian ICES-003.

Cetappareilnumérique de la classe B est conforme à la norme NMB-003 du Canada.

13.7.3 Germany

Deutschsprachiger EU Hinweis:

HinweisfürGeräte der Klasse B EU-RichtliniezurElektromagnetischenVerträglichkeit Dieses Produktentspricht den Schutzanforderungen der EU-Richtlinie 2004/108/EC zurAngleichung der Rechtsvorschriftenüber die elektromagnetischeVerträglichkeit in den EU-Mitgliedsstaaten. undhält die Grenzwerte der EN 55022 Klasse B ein. Um dieses sicherzustellen, sind die Gerätewie in den Handbüchernbeschriebenzuinstallieren und zubetreiben. Des Weiterendürfenauchnur von der HGST empfohleneKabelangeschlossenwerden. HGST übernimmtkeineVerantwortungfür die Einhaltung der Schutzanforderungen, wenn das ProduktohneZustimmung der HGST verändertbzw. wennErweiterungskomponenten von FremdherstellernohneEmpfehlung der HGST gesteckt/eingebautwerden.

Deutschland: Einhaltung des Gesetzesüber die elektromagnetischeVerträglichkeit von Geräten

Dieses Produktentsprichtdem "Gesetzüber die elektromagnetischeVerträglichkeit von Geräten (EMVG)". Dies ist die Umsetzung der EU-Richtlinie 2004/108/EC in der Bundesrepublik Deutschland.

ZulassungsbescheinigunglautdemDeutschenGesetzüber die elektromagneti-scheVerträglichkeit von Geräten (EMVG) vom 20 July 2007 (bzw. der EMC EG Richtlinie 2004/108/EC) fürGeräte der Klasse B Dieses Gerätistberechtigt, in ÜbereinstimmungmitdemDeutschen EMVG das EG-Konformitätszeichen - CE - zuführen. Verantwortlichfür die Konformitätserklärungnach Paragraf 5 des EMVG ist die HGST, a Western Digital company, 3403 Yerba Buena Road, San Jose, California 95135. Informationen in Hinsicht EMVG Paragraf 4 Abs. (1) 4:

Das Geräterfüllt die Schutzanforderungennach EN 55024 und EN 55022 Klasse B

13.7.4 KCC (Korea)

이 기기는 가정용으로 전자파적합등록을 한 기기로서 주거 지역에서는 물론 모든 지역에서 사용할 수 있습니다.

13.7.5 BSMI (Taiwan)

The SSDs comply with the Chinese National Standard (CNS) 13438 and abides by the Electromagnetic Compatibility (EMC) Framework requirements of the Taiwanese Bureau of Standards, Metrology, and Inspection (BSMI).

警告使用者：

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾

在這種情況下，使用者會被要求採取其些適當的對策。

新加坡商日立環球儲存科技股份有限公司台灣分公司
台北市敦化北路 167 號 5 樓 (宏國大樓)

14. Standards

14.1 Overview

The following sections outline the safety standards for different countries.

14.2 UL and cUL Standard Conformity

The drive is certified under the following safety standards for use in Information Technology Equipment, including Electrical Business Equipment:

- EN 60950-1:2006 with A11:2009, A1:2010, A12:2011
- IEC 60950-1:2005, Second Edition; Am 1:2009
- UL 60950-1, Second Edition, 2011-12-19, USA
- CSA C22.2 No. 60950-1-07, Second Edition, 2011-12, Canada

The UL recognition, or the cUL certification, is maintained for the duration of the product manufacturing life cycle. The UL and cUL recognition marks appear on the drive label.

14.3 European Standards Compliance

This product is certified to the EN 60950-1:2006 with A11:2009, A1:2010, A12:2011 safety standard for Europe.

14.4 German Safety Mark

The product is certified by TUV to meet EN 60950-1:2006 with A11:2009, A1:2010, A12:2011 safety standard under the Bauart Mark.

14.5 Flammability

The printed circuit board (PCB), and connectors used in this drive meet or exceed the UL minimum flammability classifications listed in the table below. The flammability ratings are marked on the printed wiring boards and flex cables.

Table 31: Component Flammability Ratings

Component	Flammability Rating
Printed Circuit Board	Min. V-1
Connector	Min. V-2

14.6 References

The *Industry Standards* and *Manufacturing Location* sections list the formal standards, relevant reference specifications, and interface protocols that apply, in whole or in part, to the product.

14.6.1 Industry Standards

Organization	Document	Description
JTAG	1149.1-19901	Joint Test Action Group (JTAG), IEEE Standard 1149.1-1990, Test Access Port and Boundary Scan Architecture.
NVMe	NVMe 1.1a	NVM Express Working Group. NVM Express (NVMe) Specification, Revision 1.1a, September 23, 2013. *http://www.nvmexpress.org*.
NVMe	SMI 1.0	NVM Express Working Group. Technical Note: NVMe Simple Management Interface (SMI), Revision 1.0, February 24, 2015.
NVMe	Technical Note 1.0	NVMe Technical Note: NVMe Basic Management Command, Revision 1.0, February 24, 2015.
SFF Committee	SFF-8223	SFF Committee, SFF-8223 Specification for 2.5” Drive Form Factor with Serial Connector, Revision 2.5, May 25, 2006.
SFF Committee	SFF-8482	SFF Committee, SFF-8482 Specification for Unshielded Dual Port Serial Attachment Connector, Revision 2.2, February 1, 2006.
SFF Committee	SFF-8639	SFF Committee, SFF-8639 Specification for Multifunction 12Gb/s 6X Unshielded Connector, Revision 1.9, March 31, 2014.
SMBus	2.0	System Management Bus (SMBus) Specification 2.0, August 3, 2000.
SSD Work Group	Version 1.0	SSD Form Factor Working Group. Enterprise SSD Form Factor, Version 1.0, December 20, 2011.
UEFI	2.3, Errata B	Unified Extensible Firmware Interface (UEFI) Specification, Version 2.3, Errata B, February 25, 2010.

14.6.2 Manufacturing Location

Location	Certifications
Penang, Malaysia	ISO 9001 Certified ISO 14001 Certified

15. Ordering Information

15.1 Overview

Example: An HUSPM3232ADP301 would be an Ultrastar SN100 Series 2.5-Inch NVMe PCIe 3.0 x4 Lane SFF with a data rate of 5Gbps per lane and a capacity of 3.2TB. The drive is rated for the commercial temperature range and the memory subsystem would consist of eMLC NAND flash components.

15.2 Model Number Decoder

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
H	U	S	P	R	3	2	#	#	A	D	P	3	0	1

1	**Manufacturer**	
	H	HGST, Inc.
2	**Product Family**	
	U	Ultrastar
3	**Series**	
	S	Standard
4, 5	**Solid-State Media**	
	PR	PCIe Read Intensive (RI)
6, 7	**Capacity, Fully Populated**	
	32	3.2TB (3200GB)
8, 9	**Capacity, Actual**	
	80	800GB
	95	956GB
	16	1.6TB (1600GB)
	19	1.91TB (1910GB)
	32	3.2TB (3200GB)
	38	3.82TB (3820GB)
10	**Series Code**	
	A	A Series
11	**Form Factor**	
	D	2.5-Inch PCIe SFF
12, 13	**Interface Protocol / Data Rate**	
	P3	PCIe Express 3.0
14	**Feature Code**	
	0	Commercial Operating Temperature
15	**Special Features Code**	
	1	Encryption Support

16. Appendix A: Log Pages

16.1 Overview

This appendix documents the vendor-unique log pages supported by the device.

16.2 Vendor-Unique Logs Page 0xC1 (C1h)

The Vendor-Unique Logs Page 0xC1 (C1h) is a collection of vendor-unique logs stored as subpages. Each subpage has a Subpage Identifier or Subpage Code. Each subpage conforms to a generalized format consisting of a subpage header and a list of parameters. The Vendor-Unique Logs Page consists of a header describing the number of subpages and total length of the entire retrievable page.

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	Number of Subpages							
1	Reserved							
2-3	Total Length (Sum of all Subpage Lengths, in bytes.)							
4	Reserved		Subpage Code (First Subpage)					
5	Reserved							
6-7	Page Length							
8-9	Parameter Code (First Parameter)							
10-n	< Parameters >							
n+1	Parameter Code (Second Parameter)							
(n+2)-m	< Parameters >…							
m+1	Reserved		Subpage Code (Second Subpage)					
m+2	Reserved							
(m+3)-(m+4)	Page Length = 006Ch							
(m+5)-(m+6)	Parameter Code							
(m+7)-(m+…)	< Parameters >…							

16.3 Current Value Parameter Descriptions

There are current value parameters that are common to those subpages that collect statistical data. The field names and descriptions are presented here for reference purposes only.

Name	Description
DU	*Disable Update Bit*. If set to 0, the device shall update this parameter according to the user-specific levels. If set to 1, the device shall not update this parameter.
DS	*Disable Save*. If set to 1, the saving of the parameter is disabled. A value of 0 indicates that the value may be saved.
TSD	*Target Save Disable*. If set to 0, the parameter is saved by the device at vendor-unique intervals. If set to 1, implicit saving is disabled.
ETC	*Enable Threshold Comparison*. If set to 1, a comparison between the current and threshold value is made whenever the parameter is updated. Set to 0 if no comparison is made.
TMC	*Threshold Meet Criteria*. Defines the basis for which comparisons are made.
LBIN	*List Binary*. If set to 0 then the list is in ASCII format. If set to 1 then the list is in binary format.
LP	*List Parameters*. If set to 0 then the parameter is a data counter. If set to 1 then the parameter is a list parameter.

16.3.1 Subpage 0x02-Write Errors

Write Errors 0x02 (02h) enables the host to extract the number of flash WRITE commands that failed to complete successfully. A percentage is derived using WRITE ERRORS/WRITE commands to create a baseline for a threshold comparison. The parameters are persistent between power cycles and can be reset by specifying DW10.RESET=1.

	Bit							
Byte	7	6	5	4	3	2	1	0
0	Reserved		Page Code = 02h					
1	Reserved							
2-3	Page Length = 006Ch							
4-5	Parameter Code = 0000h (Errors Corrected without Delays)							
6	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
7	Parameter Length = 08h							
8-15	Errors Corrected without Delays = 0							
16-17	Parameter Code = 0001h (Errors Corrected with Possible Delays							
18	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
19	Parameter Length = 08h							
20-27	Errors Corrected with Possible Delays							
28-29	Parameter Code = 0002h (Total Re-Writes)							
30	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
31	Parameter Length = 08h							
32-39	Total Re-Writes							
40-41	Parameter Code = 0003h (Total Errors Corrected)							
42	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
43	Parameter Length = 08h							
44-51	Total Errors Corrected							

52-53	Parameter Code = 0004h (Total Times Correction Algorithm Processed)						
54	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
55	Parameter Length = 08h						
56-63	Total Times Correction Algorithm Processed						
64-65	Parameter Code = 0005h (Total Bytes Processed)						
66	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
67	Parameter Length = 08h						
68-75	Total Bytes Processed						
76-77	Parameter Code = 0006h (Total Uncorrected Errors)						
78	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
79	Parameter Length = 08h						
80-87	Total Uncorrected Errors = 0						
88-89	Parameter Code = 8000h (Vendor-Unique - Flash Write Commands)						
90	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
91	Parameter Length = 08h						
92-99	Total Flash Write Commands						
100-101	Parameter Code = 8001h (Vendor-Unique)						
102	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
103	Parameter Length = 08h						
104-111	Vendor-Unique = 0						

16.3.1.1 Current Parameter Code Descriptions

The Write Error Count is derived from the flash interface program error count (including user data and system data).

Code	Name	Description
0000h	Errors Corrected without Delays	0
0001h	Errors Corrected with Possible Delays	0
0002h	Total Re-Writes	Total Flash Page Programs that are the result of Reprogram Count x Dataframe Number per Page.
0003h	Total Errors Corrected	0
0004h	Total Times Correct Algorithm Processed	Obsolete
0005h	Total Bytes Processed	Flash Page Program x Dataframe Format per Page x 4320 (4320 is derived from the Dataframe Format using 100b4320 correction).
0006h	Total Uncorrected Errors	Total Flash Page Program Fail Count x Dataframe-Number per Page.
8000h	Vendor-Unique Flash Write Commands	Total Number of Flash Write Page commands.
8001h	Vendor-Unique	0

16.3.2 Subpage 0x03-Read Errors

Read Errors 0x03 (03h) enables the host to extract the number of flash READ commands that failed to complete successfully. A percentage is derived using READ ERRORS/READ commands to create a baseline for a threshold comparison. Should a READ command fail, the only course of action is to retry the READ. Errors corrected without delay are READ errors that completed the ECC correction. The parameters are persistent between power cycles and can be reset by specifying DW10.RESET=1.

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	Reserved		Page Code = 03h					
1	Reserved							
2-3	Page Length = 006Ch							
4-5	Parameter Code = 0000h (Errors Corrected without Delays)							
6	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
7	Parameter Length = 08h							
8-15	Errors Corrected without Delays							
16-17	Parameter Code = 0001h (Errors Corrected with Possible Delays)							
18	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
19	Parameter Length = 08h							
20-27	Errors Corrected with Possible Delays							
28-29	Parameter Code = 0002h (Total Re-Reads)							
30	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
31	Parameter Length = 08h							
32-39	Total Re-Reads							
40-41	Parameter Code = 0003h (Total Errors Corrected)							
42	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
43	Parameter Length = 08h							
44-51	Total Errors Corrected							
52-53	Parameter Code = 0004h (Total Errors Corrected)							

54	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
55	Parameter Length = 08h						
56-63	Total Times Correction Algorithm Processed						
64-65	Parameter Code = 0005h (Total Bytes Processed)						
66	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
67	Parameter Length = 08h						
68-75	Total Bytes Processed						
76-77	Parameter Code = 0006h (Total Uncorrected Errors)						
78	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
79	Parameter Length = 08h						
80-87	Total Uncorrected Errors						
88-89	Parameter Code = 8000h (Vendor-Unique - Flash Read Commands)						
90	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
91	Parameter Length = 08h						
92-99	Total Flash Read Commands						
100-101	Parameter Code = 8001h (Vendor-Unique)						
102	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
103	Parameter Length = 08h						
104-111	Total XOR Recovered						
112-113	Parameter Code = 8002h (Vendor-Unique)						
114	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
115	Parameter Length = 08h						
116-123	Total Corrected Bit Count						

16.3.2.1 Current Parameter Code Descriptions

The Read Error Count is derived from the flash interface read error count (including user data and system data).

Code	Name	Description
0000h	Errors Corrected without Delays	Total Flash Dataframe Reads corrected by the ECC engine.
0001h	Errors Corrected with Possible Delays	Total Re-Reads + Vendor-Unique – Total XOR Recovered.
0002h	Total Re-Reads	Total Flash Dataframe Reads that were recovered using Read-Retry.
0003h	Total Errors Corrected	Errors Corrected without Delays + Errors Corrected with Possible Delays.
0004h	Total Times Correct Algorithm Processed	Obsolete
0005h	Total Bytes Processed	Flash Dataframe Read x 4320 (4320 is derived from the Dataframe Format using 100b4320 correction).
0006h	Total Uncorrected Errors	Total Dataframes that failed all attempts to recover the data.
8000h	Vendor-Unique Flash Read Commands	Total Number of Flash Read Dataframes commands.
8001h	Vendor-Unique – Total XOR Recovered	Total Flash Dataframe Reads that were recovered using RAID.
8002h	Vendor-Unique – Total Corrected Bit Count	Total Corrected Bits of Corrected Dataframes.

Note: The host calculates the Total Bit Error Rate Counter:

*Total Corrected Bits of Corrected Dataframes / (Total Corrected Dataframes * Dataframe Size)*

16.3.3 Subpage 0x05-Verify Errors

Verify Errors 0x05 (05h) is an error counter, allowing the host to extract the number of flash VERIFY errors that occurred. The parameters are persistent between power cycles and can be reset by specifying DW10.RESET=1.

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	Reserved		Page Code = 05h					
1	Reserved							
2-3	Page Length = 0060h							
4-5	Parameter Code = 0000h (Errors Corrected without Delays)							
6	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
7	Parameter Length = 08h							
8-15	Errors Corrected without Delays							
16-17	Parameter Code = 0001h (Errors Corrected with Possible Delays)							
18	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
19	Parameter Length = 08h							
20-27	Errors Corrected with Possible Delays							
28-29	Parameter Code = 0002h (Total Re-Reads)							
30	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
31	Parameter Length = 08h							
32-39	Total Re-Reads							
40-41	Parameter Code = 0003h (Total Errors Corrected)							
42	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
43	Parameter Length = 08h							
44-51	Total Errors Corrected							
52-53	Parameter Code = 0004h (Total Times Correction Algorithm Processed)							
54	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
55	Parameter Length = 08h							

56-63	Total Times Correction Algorithm Processed						
64-65	Parameter Code = 0005h (Total Bytes Processed)						
66	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
67	Parameter Length = 08h						
68-75	Total Bytes Processed						
76-77	Parameter Code = 0006h (Total Uncorrected Errors)						
78	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
79	Parameter Length = 08h						
80-87	Total Uncorrected Errors						
88-89	Parameter Code = 8000h (Vendor-Unique - Total Commands Processed)						
90	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
91	Parameter Length = 08h						
92-99	Total Commands Processed						

16.3.3.1 Current Parameter Code Descriptions

Code	Name	Description
0000h	Errors Corrected without Delays	0
0001h	Errors Corrected with Possible Delays	0
0002h	Total Re-Reads	0
0003h	Total Errors Corrected	0
0004h	Total Times Correct Algorithm Processed	0
0005h	Total Bytes Processed	0
0006h	Total Uncorrected Errors	0
8000h	Total Commands Processed	0

16.3.4 Subpage 0x10-Self-Test Results

Self-Test Results 0x10 (10h) provides the results from the twenty (20) most recent self-tests. The results of the most recent self-test or the self-test that is currently in progress are reported in the first self-test log parameter; the results from the second most recent self-test are reported in the second parameter, etc. If there are fewer than twenty self-tests, the unused self-test log parameter entries are zero.

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	DS	SPF (0b)	Page Code = 10h					
1	Reserved							
2	(LSB)	Page Length (190h)						
3								(MSB)
Self-Test Results Log Parameters								
4	Self-Test Results Log Parameter (First)							
23								
	.							
	.							
	.							
384	Self-Test Results Log Parameter (Twentieth)							
403								

16.3.4.1 Self-Test Results 0x10 Log Parameter Format

The following table illustrates the format of a single Self-Test log parameter.

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	(LSB)	Parameter Code						
1	(0001h to 0014h)							(MSB)
2	DU	Obsolete	TSD	ECT	TMC		Format and Linking	
3	Parameter Length (10h)							
4	Self-Test Code			Reserved	Self-Test Results			
5	Self-Test Number							
6	(LSB)	Accumulated Power-On Hours						
7								(MSB)
8	(LSB	Address of First Failure						
15								(MSB)
16	Reserved				Sense Key			
17	Additional Sense Code							
18	Additional Sense Code Qualifier							
19	Vendor-Unique							

Notes:

1. The *Parameter Code* field will identify the log parameter being transferred. The results of the most recent self-test will contain 0001h; the second most recent test will contain 0002h, etc.
2. The *Format and Linking* field for each log parameter in the log page shall be set to 11b to indicate that the parameters are binary format list parameters.
3. The Parameter Length field shall contain 10h.
4. The Self-Test Code field contains the value in the Self-Test Code field of the Self-Test Command (0xC4 / 0x03) that started the self-test.

16.3.5 Subpage 0x15-Background Scan

Background Scan 0x15 (15h) will return information regarding the Background Media Scans (BMS) that are performed at regular intervals to ensure data integrity. The subpage will report the results of a background scan or return the current status of a background scan. Bytes 4 to 23 return the current state of the background scan. The remaining bytes (Byte 24 and afterward) will contain information about every block that was retired as a result of a background It also reports any errors in the logical blocks that were detected during a medium scan operation and whether the logical blocks were retired.

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	Page Code = 15h							
1	Reserved							
2	(LSB)	Page Length (n-3)						
3								(MSB)
Background Scan Results Log Parameters								
4		Background Scanning Status Parameter/						
19	Background Scan Parameter List							
20-43	Background Scan Results Log Parameter (First)							
	. . .							
n-23		Background Scan Results Log Parameter (Last)						
n								

16.3.5.1 Background Scan Results 0x15 Parameter Format

The following table illustrates the format of a single Background Scan Result log parameter.

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	(LSB)	Parameter Code						
1	(0000h_							(MSB)
2	DU	Obsolete	TSD	ETC	TMC		Format and Linking	
3	Parameter Length (10h)							
4-7	Accumulated Power-On Minutes							
8	Reserved							
9	Background Scanning Status 00h – Background Scan Idle 01h – Background Scan Active 08h – Background Scan Suspended							
10	(LSB)	Number of Scans Performed						
11	Total Number of BMS Completed up to the Current Time							(MSB)
12	(LSB)	Medium Scan Progress						
13	Current Percentage of the Entire Drive that has been Scanned in this Iteration of BMS							(MSB)
14-19	Reserved							

Note: The Background Scan will occur slowly over an extended time. When a scan begins, it is active for a section of the drive, then suspended until a timer triggers it to scan the next section. The BMS will switch between active and suspended for each section until the entire drive is scanned; it is at this point that the scan becomes idle.

16.3.5.2 Background Scan Results 0x15 Log Parameter

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	(LSB)	Parameter Code						
1	(0001h to 0800h)							(MSB)
2	DU = 0	Obsolete = 0	TSD = 0	ETC = 0	TMC = 0		Format and Linking = 0	
3	Parameter Length (14h)							
4-7	Accumulated Power On Minutes at Time of Background Scan Result Log							
8	Reassign Status = 04h				Sense Key = 01h			
9	Additional Sense Code = 0Bh							
10	Additional Sense Qualifier = 01h							
11-13	Vendor-Unique = 00h							
14-15	Vendor-Unique = 00h							
16-19	Reserved							
20	(LSB)	NAND Defect Detected by BGMS						
23	0x[C/T]BBBBPPP [31:27] = C (Channel) [26:24] = T (Target) [23:12] = B (Block Number) [11:00] = P (Page Numberl)							(MSB)

Notes:

1. The *Parameter Code* field will identify the log parameter being transferred. The results of the most recent Background Scan will contain 0001h; the second most recent scan will contain 0002h, etc.
2. The *Format and Linking* field for each log parameter in the log page shall be set to 00b to indicate that the parameters are data counters.
3. The *Parameter Length* field shall be 14h.
4. The BGMS results must be accumulated and the most recent 800h (2K) defects are to be reported via the log page. Each defect is reported in terms of defect discovery and location. The location (NAND Info) is reported as [C/T]BBBBPPP, similar to Read Defect Data.

16.3.6 Subpage 0x30-Erase Errors

Erase Errors 0x30 (30h) allows the host to extract the number of flash ERASE commands that failed to complete successfully. A percentage is calculated using ERASE ERRORS/ERASE commands to form the basis for a threshold comparison. The *Total Bytes Processed* field is calculated by multiplying the *Block Size* and the *Number of Commands*. The parameters are persistent between power cycles and can be reset by specifying DW10.RESET=1.

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	Reserved		Page Code = 30h					
1	Reserved							
2-3	Page Length = 0078h							
4-5	Parameter Code = 0000h (Errors Corrected without Delays)							
6	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
7	Parameter Length = 08h							
8-15	Errors Corrected without Delays							
16-17	Parameter Code = 0001h (Errors Corrected with Possible Delays)							
18	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
19	Parameter Length = 08h							
20-27	Errors Corrected with Possible Delays							
28-29	Parameter Code = 0002h (Total Re-Erase)							
30	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
31	Parameter Length = 08h							
32-39	Total Re-Erase							
40-41	Parameter Code = 0003h (Total Errors Corrected)							
42	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
43	Parameter Length = 08h							
44-51	Total Errors Corrected							
52-53	Parameter Code = 0004h (Total Times Correction Algorithm Processed)							
54	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
55	Parameter Length = 08h							
56-63	Total Times Correction Algorithm Processed							

64-65	Parameter Code = 0005h (Total Bytes Processed)						
66	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
67	Parameter Length = 08h						
68-75	Total Bytes Processed						
76-77	Parameter Code = 0006h (Total Uncorrected Errors)						
78	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
79	Parameter Length = 08h						
80-87	Total Uncorrected Errors						
88-89	Parameter Code = 8000h (Vendor-Unique - Flash Erase Commands)						
90	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
91	Parameter Length = 08h						
92-99	Flash Erase Commands						
100-101	Parameter Code = 8001h (Vendor-Unique - Manufacturers Defect Count)						
102	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
103	Parameter Length = 08h						
104-111	Manufacturers Defect Count						
112-113	Parameter Code = 8002h (Vendor-Unique - Grown Defect Count)						
114	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
115	Parameter Length = 08h						
116-123	Grown Defect Count						
124-125	Parameter Code = 8003h (Vendor-Unique – Max Erase Count of User Data)						
126	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
127	Parameter Length = 08h						
128-135	Max Erase Count of User Data						
136-137	Parameter Code = 8004h (Vendor-Unique – Maximum Erase Count of System Data)						
138	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	LBIN = 1	LP = 1
139	Parameter Length = 08h						
140-147	Max Erase Count of System Data						

16.3.6.1 Current Parameter Code Descriptions

The Erase Error Count is calculated using the flash interface erase error count (including user data and system data).

Code	Name	Description
0000h	Errors Corrected without Delays	0
0001h	Errors Corrected with Possible Delays	0
0002h	Total Re-Erase	0
0003h	Total Errors Corrected	0
0004h	Total Times Corrected Algorithm Processed	Obsolete
0005h	Total Bytes Processed	0
0006h	Total Uncorrected Errors	Total Blocks Erase Failed Count
8000h	Vendor-Unique – Flash Erase Commands	Total Number of Flash Erase Commands
8001h	Vendor-Unique – Manufacturers Defect Count	Total Number of MBB + BBB Defect Blocks
8002h	Vendor-Unique – Grown Defect Count	Total Number of GBB Defect Blocks
8003h	Vendor-Unique – Max Erase Count of User Data	The highest erase count of user data blocks.
8004h	Vendor-Unique – Max Erase Count of System Data	The highest erase count of system data blocks.

16.3.7 Subpage 0x31-Erase Counts

Erase Counts 0x31 (31h) will allow the host to count the number of flash ERASE commands that have occurred on a per channel basis. The lowest, highest and average erase counts for all the channels are recorded. The number of channels is a fixed value.

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	Reserved		Page Code = 31h					
1	Reserved							
2-3	Page Length (n - 3) = 12 + 36 * Number of Channels							
4-5	Parameter Code = 8000h							
6	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
7	Parameter Length = 08h							
8-15	Highest Erase Count Across All Channels							
Channel Erase Count Descriptor List								
16	Channel Erase Count Descriptor (First Channel)							
51								
. . .								
n-35	Channel Erase Count Descriptor (Last Channel)							
n								

16.3.7.1 Channel Erase Count Descriptor

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0-1	Parameter Code = 8001h							
2	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
3	Parameter Length = 08h							
4-11	Highest Erase Count							
12-13	Parameter Code = 8002h							
14	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
15	Parameter Length = 08h							
16-23	Lowest Erase Count							
24-25	Parameter Code = 8003h							
26	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		LBIN = 1	LP = 1
27	Parameter Length = 08h							
28-35	Average Erase Count							

16.3.8 Subpage 0x32-Temperature History

Temperature History 0x32 (32h) allows the user to query the temperature data of a device. The user can also define the temperature parameters or reset the defaults. The parameters are persistent between power cycles and can be reset by specifying DW10.RESET=1.

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	Reserved = 0		Page Code = 32h					
1	Reserved = (00h)							
2-3	Page Length = 10h							
4	Current Temperature (°C)							
5	Reference Temperature (°C)							
6	Maximum Temperature (°C)							
7	Minimum Temperature (°C)							
8	(MSB)	Power-On Hour When Maximum Temperature Occurred (Minutes)						
11								(LSB)
12	(MSB)	Total Time Spent Over Reference Temperature (Minutes)						
15								(LSB)
16	(MSB)	Power-On Hour When Minimum Temperature Occurred (Minutes)						
19								(LSB)

16.3.8.1 Current Parameter Descriptions

Name	Description
Current Temperature (°C)	The current temperature reported by the internal temperature sensor(s).
Reference Temperature (°C)	The current reference temperature
Maximum Temperature (°C)	The highest temperature recorded by the internal temperature sensor(s).
Minimum Temperature (°C)	The lowest temperature recorded by the internal temperature sensor(s).
Power-On Hour When Maximum Temperature Occurred (Minutes)	A power-on stamp, in minutes, when the highest temperature was recorded. The default value for this field is the current power-on hour.
Total Time Spent Over Reference Temperature (Minutes)	The total time, in minutes, that the current temperature has been greater than the reference temperature. The default value for this field is 00h.
Power-On Hour When Minimum Temperature Occurred (Minutes)	The default value for this field is the current power-on hour.

16.3.9 Subpage 0x37-SSD Performance

SSD Performance 0x37 (37h) will return a set of statistical data in regards to performance. The user can specify a *Subpage Code* to return a statistical set. See *Statistical Set Descriptions* and *Subpage Code Sets*.

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	Reserved = 0		Page Code = 37h					
1	Subpage Code = [*xx*h] (*xx* from 0x01 to 0x0F)							
2-3	Page Length = 78h							
4-11	Host Read Commands							
12-19	Host Read Blocks							
20-27	Host Read Cache Hit Commands							
28-35	Host Read Cache Hit Blocks							
36-43	Host Read Commands Stalled							
44-51	Host Write Commands							
52-59	Host Write Blocks							
60-67	Host Write Odd Start Commands							
68-75	Host Write Odd End Commands							
76-83	Host Write Commands Stalled							
84-91	NAND Read Commands							
92-99	NAND Read Blocks							
100-107	NAND Write Commands							
108-115	NAND Write Blocks							
116-123	NAND Read Before Write							

16.3.9.1 Statistical Set Descriptions

Set	Description
0-13	These sets are (inclusive) of type Stats_Set.
14-15	These sets are the type Stats_Set_Long.

16.3.9.2 Subpage Code Sets

Set Name	Description
0	This set is currently updated by the drive code and is copied to Sets 1 through 12 in a cyclic manner every statistical period. Set 0 is then cleared and the process repeated.
1-12	These are the copies of the Set 0 statistics that are maintained over the last 12 statistical periods (5 minute duration). The 5 minute duration is fixed and cannot be changed.
13	This is the accumulated total of Sets 1 to 12 and will return the statistics collected over the previous hour by default.
14	These are the statistics accumulated since power-up and are updated when the current set (Set 0) is moved to Set 1, i.e., the total excluding Set 0.
15	These are the statistics accumulated during the drive life time and are updated when the current set (Set 0) is moved to Set 1, i.e., the total excluding Set 0.
16	Used as a work area.

16.3.9.3 SSD Performance 0x37 Definitions

Field	Description
Host Read Commands	Number of host read commands received during the reporting period.
Host Read Blocks	Number of 512-byte blocks requested during the reporting period. *Host Read Blocks / Host Read Commands = Average Read Size*
Host Read Cache Hit Commands	Number of host read commands that were serviced exclusively from the on-board read cache during the reporting period. No access to the NAND flash memory was required. This count is only updated if the entire command is serviced from the cache memory. *((Host Read Cache Hit Commands / Host Read Commands) * 100) will return a percentage of Host Read Commands satisfied from the cache.*
Host Read Cache Hit Blocks	Number of 512-byte blocks of data that have been returned for Host Read Cache Hit Commands during the reporting period. This count is only updated with the blocks returned for Host Read Commands that were serviced entirely from cache memory. *((Host Read Cache Hit Blocks / Host Read Blocks) * 100) will return a percentage of read data entirely from the cache.*
Host Read Commands Stalled	Number of Host Read Commands that were stalled due to lack of resources within the SSD during the reporting period (NAND flash command queue full, low cache page count, cache page contention, etc.). Commands are not considered stalled if the only reason for the delay was waiting for the data to be physically read from the NAND flash. It is normal to expect this count to equal zero on heavily utilized systems. *((Host Read Commands Stalled / Host Read Commands) * 100) will return a percentage of read commands that were stalled.* If the figure is consistently high, then consideration should be given to spreading the data across multiple SSDs.
Host Write Commands	Number of Host Write Commands received during the reporting period.
Host Write Blocks	Number of 512-byte blocks written during the reporting period. *Host Write Blocks / Host Write Commands = Average Write Size*
Host Write Odd Start Commands	Number of Host Write Commands that started on a non-aligned boundary during the reporting period. The size of the boundary alignment is normally 4K; therefore, this returns the number of commands that started on a non-4K aligned boundary. The SSD requires slightly more time to process non-aligned write commands than it does to process aligned write commands. *((Host Write Odd Start Commands / Host Write Commands) * 100) will return a percentage of Host Write Commands that started on a non-aligned boundary.* If this figure is equal to or approximately 100%, and the NAND Read Before Write value is also high, then the user should investigate the possibility of offsetting the file system. For Microsoft Windows systems, the user can use Diskpart. For Linux-based operating systems, there is normally a method whereby file system partitions can be placed where required.
Host Write Odd End Commands	Number of Host Write Commands that ended on a non-aligned boundary during the reporting period. This size of the boundary alignment is normally 4K; therefore, this returns the number of commands that ended on a non-4K aligned boundary. *((Host Write Odd End Commands / Host Write Commands) * 100) will return a percentage of Host Write Commands that ended on a non-aligned boundary.*

Field	Description
Host Write Commands Stalled	Number of Host Write Commands that were stalled due to lack of resources within the SSD during the reporting period. The most probable cause is that the write data was being received faster than it could be saved to the NAND flash memory. If there was a large volume of read commands being processed simultaneously, then other causes might include the NAND flash command queue being full, low cache page count, or cache page contention, etc. It is normal to expect this count to be zero on heavily utilized systems. *((Host Write Commands Stalled / Host Write Commands) * 100) will return a percentage of write commands that were stalled.* If the figure is consistently high, then consideration should be given to spreading the data across multiple SSDs.
NAND Read Commands	Number of read commands issued to the NAND devices during the reporting period. This value will normally be much higher than the Host Read Commands value, as the data needed to satisfy a single Host Read Command may be spread across several NAND flash devices.
NAND Read Blocks	Number of 512-byte blocks requested from the NAND flash devices during the reporting period. This value would normally be about the same as the Host Read Blocks value. *NAND Read Blocks / NAND Read Commands will return an average size of NAND Read Commands.* *((NAND Read Blocks / Host Read Blocks) * 100) will return a percentage of host read data that had to be physically read from the NAND flash devices.*
NAND Write Commands	Number of Write Commands issued to the NAND devices during the reporting period. There is no real correlation between the number of Host Write Commands issued and the number of NAND Write Commands.
NAND Write Blocks	Number of 512-byte blocks written to the NAND flash devices during the reporting period. This value would normally be about the same as the Host Write Blocks value. *NAND Write Blocks / NAND Write Commands will return the Average Size of NAND Write Commands. This value should never be greater than 128K, as this is the maximum size write that is every issued to a NAND device.* *((NAND Write Blocks / Host Write Blocks) * 100) will return a percentage of host write data that had to be physically written to the NAND devices. If this figure is less than 100%, it indicates that host blocks that were written to more than once were combined into a single write to the NAND flash device(s).*
NAND Read Before Write	The number of Read Before Write operations that were required to process non-aligned Host Write Commands during the reporting period. *See Host Write Odd Start Commands* and *Host Write Odd End Commands*. NAND Read Before Write operations have a detrimental effect on the overall performance of the device.

16.3.10 Subpage 0x38-Other Statistics

Other Statistics 0x38 (38h) will indicate the frequency of firmware downloads to the non-volatile memory. The parameters are persistent between power cycles and can be reset by specifying DW10.RESET=1 (Bit 31).

	Bit							
Byte	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	Reserved = 0		Page Code = 38h					
1	Reserved = (00h)							
2-3	Page Length = 08h							
4-7	FW Download Number							
8-11	Reserved = (00000000h)							

16.3.10.1 Current Parameter Descriptions

Name	Description
FW Download Number	This value will indicate the number of times the FW was downloaded to the non-volatile memory. The user can specify DW10.RESET=1 (Bit 31) to reset the value.

17. Appendix B: Inquiry VPD Pages

17.1 Overview

The Inquiry VPD (Vital Product Data) pages listed in this section can be extracted via the SMBus. The method implemented for the device conforms to the *NVMe Simple Management Interface Specification (NVM Express, Technical Note: NVMe Simple Management Interface, Revision 1.0, February 24, 2015)*.

The PCIe SMBus adheres to the *Enterprise SSD Form Factor 1.0 Specification* when the 3.3Vaux signal (3.3V Auxiliary Power) is implemented in the system; otherwise, if the 3.3Vaux signal is disabled (not connected), but the 12V rail is powered, then the system adheres to the *NVMe Management Interface (MI) Specification*.

17.2 Enterprise SSD Form Factor Vital Product Data (VPD) Table

The usual I^2C command syntax is used to read from the Enterprise SSD VPD table via the PCIe SMBus. The first two bytes are the address offset into the VPD table (Slave Address 0xA6), while the remainder of data is that read from the VPD table starting at the address offset. The address offset is in big-endian format. If data is read beyond the definition of the VPD Table, then the returned data is considered undetermined. Table 32 defines the VPD.

Table 32: Enterprise SSD Form Factor VPD Table

Register	Register Offset	Width (Bytes)	Type	Default Value	Description
Class Code	0	3	RO	0x01, 0x80, 0x00	Device type and programming interface.
PCI-Sig	3	2	RO	0x58, 0x1C	PCI-SIG Vendor ID
Serial # (1)	5	20	RO	“SerialNum” (remainder are zeros)	Serial Number (vendor/device unique)
Model # (1)	25	40	RO	“ModelNum” (remainder are zeros)	Model Number (ASCII string)
P0 Max Lnk Spd	65	1	RO	0x3	Port 0: Maximum Link Speed - Gen3
P0 Max Lnk Wd	66	1	RO	0x4	Port 0: Maximum Link Width - x4
P1 Max Lnk Spd	67	1	RO	0x0	Port 1: Maximum Link Speed - Not Supported.
P1 Max Lnk Wd	68	1	RO	0x0	Port 1: Maximum Link Width - Not Supported.
Pwr Rail Init	69	1	RO	0xA	12V Power Rail Initial Power Requirement - 10 Watts
Reserved	70	1	RO	0x0	Internal use only.
Reserved	71	1	RO	0x0	Internal use only.
Pwr Rail Max	72	1	RO	0x19	12V Power Rail Maximum Power Requirement - 25 Watts.
Reserved	73	1	RO	0x0	Internal use only.
Reserved	74	1	RO	0x0	Internal use only.
Cap List Offset	75	2	RO	0x4D, 0x00	16-bit address pointer to start of capability list.

Note: The “SerialNum” and “ModelNum” fields will be updated when the controller firmware sends the serial and model number data; the defaults are hard-coded.

17.3 Vendor-Unique Features

There are additional vendor-unique features that use the VPD standard method of extending the basic VPD definition. Figure 17 illustrates the definition. Table 33 lists the vendor-unique features that extend the definition.

Figure 17: SMBus Capability Definition

Notes:

1. The PCI-SIG Vendor ID for the manufacturer is **0x1C58**.
2. Capability List Headers are 4B in size.
3. Next Capability Address of zero indicates the last capability in the list.

Table 33: Vendor-Unique Extended Definitions

Register	Vendor-Specific ID	Description
Temperature Sensor	0x5A8C1C58	The Vendor-Specific and PCI-SIG Vendor ID values.

17.3.1 Temperature Sensor Access

The host should not make any assumptions about the offset of the Temperature Sensor(s), but should follow the linked list to scan for the vendor-unique command being sought. To read from or write to the Temperature Sensor(s), the usual I^2C command syntax is used. The first byte is the address offset pointing to the Temperature Sensor register table. The remainder of the data that is clocked is the data that is read from or written to starting at the address offset. If the data is read beyond the definition of the Temperature Sensor table, then the remaining data that is returned is undetermined.

Table 34: Temperature Sensor Vendor-Specific Command

Register	Register Offset	Width (Bytes)	Type	Default Value	Description
PCI-SIG (0xA5)	77	2	RO	0xA5, 0x00	**Start of Capability 0.**
Pointer Next Cap	79	2	RO	0x57, 0x00	Pointer to next capability.
PCI-SIG Vendor ID	81	2	RO	0x58, 0x1C	HGST Vendor ID.
Temp. Sensor ID Number	83	2	RO	0x8C, 0x5A	Manufacturer Vendor ID.
Value (Initial Software Default)	85	2	RO	0x00, 0x00	Temperature Value (Celsius)

17.4 NVMe Management Interface Table

The SMBus may switch between the Enterprise SSD Form Factor VPD table and the NVMe Management Interface (NVMe MI) table (Slave Address 0xD4). A microcontroller will switch to the NVMe MI table once it begins receiving NVMe MI data from the device controller. The device controller sends the NVMe MI data every five (5) seconds. The NVMe MI table adheres to the structure outlined in the specification and adds a Vendor-Unique Command Code at Byte 32.

The user should reference the NVMe Technical Note: NVMe Basic Management Command, Revision 1.0, February 24, 2015, for more information.

There are only a few fields, excluding the PEC calculation, within the NVMe MI table that are managed by the microcontroller. The microcontroller manages the “SMBus Arbitration” and “Drive Functional” flags that are within Status Flags (SFLGS) byte residing in Command Code 0. The “SMBus Arbitration” flag shall be set when a SMBus block read (32 bytes) is performed starting at Command Code 0. The “SMBus Arbitration” flag shall be cleared when the “SMBus Send Byte to Reset Arbitration Bit” sequence is performed (See the NVMe MI Specification for details.). The “Drive Functional” flag is set when the microcontroller does not receive a NVMe MI data transmission from the device controller within ten (10) seconds; otherwise, the “Drive Functional” flag is cleared.

17.4.1 Command Code 32 Structure

The Vendor-Unique Command Code 32 is one SMBus block read in length. Table 35 defines the Command Code 32 structure.

Table 35: NVMe MI Command Code 32

Byte	32	33	34	35	36	37	38	39	40	41	42	..	61	62	63
	Len	Status	Temp (Main)	Temp (Inlet)	Temp (DB1)	Temp (DB2)	Power (Sum)		Vendor ID		Serial Number			Reserved (Set to 0)	PEC

17.4.2 Command Code 32 Field Definitions

Table 36 lists the field definitions for Command Code 32.

Table 36: Command Code 32 Field Definitions

Field	Definition
Len	This field will indicate the number of bytes to read after the length before reaching PEC. The value for this field is 30.
Status	Status Flag. Only Bit 7 is used within Command Code 32; Bit 7 is the arbitration flag. When a SMBus block read of Command Code 32 is performed, the arbitration flag will be set to one (1). The arbitration flag will be cleared when the “SMBus Send Byte to Reset Arbitration Bit” sequence is performed. The remaining bits, 6-0, shall be zero.
Temp	The four temperature sensors for the Main, Inlet, Daughter Board 1 (DB1) and Daughter Board 2 (DB2). The temperatures should adhere to the format as described for the Composite Temperature in Command Code 0.
Power	The Power is the sum of the controller and flash channels. The value is represented as an unsigned integer, in little-endian format. The value is expressed in milliwatts.
Vendor ID	The two-byte Vendor ID assigned by the PCI SIG Working Group. It should match the VID in the Identify Controller command response.
Serial Number	The twenty (20) characters that match the serial number in the NVMe Identify Controller command response. The first character is transmitted first.
Reserved	Not used and shall be set to zero.
PEC	An eight (8) bit CRC calculated over the slave address, command code, second slave address, and the returned data.

18. Contact Information

18.1 General Information

Main Web Site: *http://www.hgst.com*

18.2 Technical Support

Software Support: *http://www.hgst.com/support/software-support*

Solid-State Drive Support: *http://www.hgst.com/support/solid-state-drive-support*

18.3 Email Support and Telephone Support

Email Support: *support@hgst.com*

Telephone Support: 1-855-778-2497. 24 x 7 Support. Please have the following information available when calling: Product Name, Model Number, Part Number and Operating System.

Index

HGST, a Western Digital company
3403 Yerba Buena Road
San Jose, CA 95135
Produced in the United States

United States: 800-801-4618 (Toll-Free)
International: 408-717-6000
FAX: 408-717-5000

March 22, 2016