

U.2 SSD

4TS2-P Series

Customer:

Customer

Part

Number:

Innodisk

Part

Number:

Innodisk

Model Name:

Date:

Innodisk Approver	Customer Approver

Features:

- PCIe Gen 4 x 4, NVMe SSD
- Kioxia 3D TLC NAND
- U.2 SSD 4TS2-P
- Standard/Wide-temperature
- iPowerguard
- iDataguard
- Dynamic Thermal Management
- Hybrid Write

Performance:

- Sequential Read up to 7,150 MB/s
- Sequential Write up to 6,100 MB/s

Power Requirements:

Input Voltage:	+12 DC +- 5%
Max Operating Wattage (R/W):	12.6W
Idle Wattage:	3.2W

Reliability:

Capacity	TBW (Enterprise)	DWPD
400GB	664	1
800GB	1328	1
1.6TB	2656	1
3.2TB	5312	1
6.4TB	7437	0.7

Data Retention	1 Year
Warranty	5 Years

DWPD based on Enterprise workload with 5 years limited warranty

For warranty details, please refer to:

https://www.innodisk.com/en/support_and_service/warranty

Table of contents

1. PRODUCT OVERVIEW	8
1.1 INTRODUCTION OF INNODISK U.2 SSD 4TS2-P	8
1.2 PRODUCT VIEW AND MODELS	8
1.3 PCIe INTERFACE	8
2. PRODUCT SPECIFICATIONS	9
2.1 CAPACITY AND DEVICE PARAMETERS.....	9
2.2 PERFORMANCE	9
2.3 ELECTRICAL SPECIFICATIONS	11
2.3.1 Power Requirement.....	11
2.3.2 Power Consumption	11
2.4 ENVIRONMENTAL SPECIFICATIONS	11
2.4.1 Temperature Ranges	11
2.4.2 Humidity.....	11
2.4.3 Shock and Vibration	11
2.4.4 Mean Time between Failures (MTBF)	12
2.5 CE AND FCC COMPATIBILITY	12
2.6 RoHS COMPLIANCE	12
2.7 RELIABILITY	12
2.8 TRANSFER MODE	13
2.9 PIN ASSIGNMENT	14
2.10 MECHANICAL DIMENSIONS.....	16
2.11 ASSEMBLY WEIGHT	17
2.12 SEEK TIME	17
2.13 NAND FLASH MEMORY.....	17
3. THEORY OF OPERATION	18
3.1 OVERVIEW.....	18
3.2 PCIe GEN 4x4 CONTROLLER	18
3.3 ERROR DETECTION AND CORRECTION.....	19
3.4 WEAR-LEVELING	19
3.5 BAD BLOCKS MANAGEMENT	19
3.6 iDATA GUARD	19
3.7 GARBAGE COLLECTION/TRIM	19
3.8 THERMAL MANAGEMENT	20
3.9 THERMAL MANAGEMENT	20
3.10 THERMAL THROTTLING.....	20
3.11 iPOWER GUARD	20

3.12 DIE RAID20

3.13 SLC CACHE.....21

4. INSTALLATION REQUIREMENTS22

4.1 U.2 SSD 4TS2-P PIN DIRECTIONS22

4.2 ELECTRICAL CONNECTIONS FOR U.2 SSD 4TS2-P23

4.3 DEVICE DRIVE23

5. SMART / HEALTH INFORMATION.....24

5.1 GET LOG PAGE(LOG IDENTIFIER 02H)24

REVISION HISTORY

Revision	Description	Date
V1.0	First release	Jul., 2022
V1.1	Update performance	Dec., 2022
V1.2	Revise Pin assignment	Feb., 2023
V1.3	Revise enterprise workload	Aug., 2023
V1.4	Revise enterprise workload Add Thermal Throttling Description	Oct., 2023
V1.5	Revise Performance Noted Add Thermal Throttling Description	Nov., 2023
V1.6	Update Performance	Jan., 2024
V1.7	Update Data Retention	Sep., 2024
V1.8	Update Performance Noted	Nov., 2024

List of Tables

TABLE 1: DEVICE PARAMETERS.....9

TABLE 2: PERFORMANCE - 112 LAYERS 3D TLC9

TABLE 4: LATENCY (QD1)10

TABLE 5: QUALITY OF SERVICE (QoS).....10

TABLE 6: INNODISK U.2 SSD 4TS2-P POWER REQUIREMENT 11

TABLE 7: POWER CONSUMPTION 11

TABLE 8: TEMPERATURE RANGE FOR U.2 SSD 4TS2-P..... 11

TABLE 9: SHOCK/VIBRATION TESTING FOR U.2 SSD 4TS2-P 11

TABLE 10: U.2 SSD 4TS2-P MTBF12

TABLE 11: U.2 SSD 4TS2-P TBW12

TABLE 12: INNODISK U.2 SSD 4TS2-P PIN ASSIGNMENT14

TABLE 13: INNODISK U.2 SSD 4TS2-P LED INDICATOR16

TABLE 14: U.2 SSD 4TS2-P SLC CACHE21

TABLE 15: GET LOG PAGE – SMART / HEALTH INFORMATION LOG24

List of Figures

FIGURE 1: INNODISK U.2 SSD 4TS2-P8

FIGURE 2: INNODISK U.2 SSD 4TS2-P16

FIGURE 3: INNODISK U.2 SSD 4TS2-P BLOCK DIAGRAM.....18

FIGURE 4: DEVICE SIGNAL SEGMENT POWER SEGMENT.....22

FIGURE 5: SIGNAL SEGMENT AND POWER SEGMENT(HOST/CABLE SIDE)23

1. Product Overview

1.1 Introduction of Innodisk U.2 SSD 4TS2-P

Innodisk U.2 SSD 4TS2-P is an NVM Express SSD designed as PCIe SFF-8639 module with PCIe interface and 3D TLC NAND Flash. U.2 SSD 4TS2-P supports PCIe Gen 4x4, and it is compliant with NVMe 1.4 providing excellent performance. With sophisticated error detection and correction (ECC) functions, the module can ensure full End-to-end Data Path Protection that secures the data transmission between host system and NAND Flash.

Innodisk U.2 SSD 4TS2-P provides ultra-speed and high IOPS and offers maximum capacity up to 6.4TB, making the SSD optimal for server and heavy data workload applications.

CAUTION *TRIM must be enabled.*

TRIM enables SSD's controller to skip invalid data instead of moving. It can free up significant amount of resources, extends the lifespan of SSD by reducing erase, and write cycles on the SSD. Innodisk's handling of garbage collection along with TRIM command improves write performance on SSDs.

1.2 Product View and Models

Innodisk U.2 SSD 4TS2-P is available in follow capacities within 3D TLC flash ICs.

U.2 SSD 4TS2-P 400GB

U.2 SSD 4TS2-P 800GB

U.2 SSD 4TS2-P 1.6TB

U.2 SSD 4TS2-P 3.2TB

U.2 SSD 4TS2-P 6.4TB




Figure 1: Innodisk U.2 SSD 4TS2-P

1.3 PCIe Interface

Innodisk U.2 SSD 4TS2-P supports PCIe Gen IV interface and compliant with NVMe 1.4. U.2 SSD 4TS2-P can work under PCIe Gen 1, Gen 2, Gen 3, & Gen 4.

Most of operating system includes NVMe in-box driver now. For more information about the driver support in each OS, please visit <https://nvmeexpress.org/drivers/>.

2. Product Specifications

2.1 Capacity and Device Parameters

U.2 SSD 4TS2-P device parameters are shown in Table 1.

Table 1: Device parameters

Capacity	LBA	User Capacity(MB)
400GB	781422768	381554
800GB	1562824368	763097
1.6TB	3125627568	1526185
3.2TB	6251233968	3052360
6.4TB	12502446768	6104710

2.2 Performance

Burst Transfer Rate: 8 GB/s

Table 2: Performance - 112 Layers 3D TLC

Capacity	Unit	400GB	800GB	1.6TB	3.2TB	6.4TB
Sequential** Read (Q8T1)	MB/s	5,800	6,950	6,650	6,650	6,650
Sequential** Write (Q8T1)		2,250	4,100	4,950	5,000	4,850
Sustained Sequential Read (Avg.)***		1,850	2,550	2,450	2,400	2,100
Sustained Sequential Write (Avg.)***		460	870	1,600	1,750	1,450
4KB Random** Read (Q32T16)	IOPS	494,000	818,000	819,000	814,000	818,000
4KB Random** Write (Q32T16)		593,000	700,000	714,000	719,000	711,000

Note: * Performance results are measured in Room Temperature with Out-of-Box devices and may vary depending on overall system setup. In addition, 4TG2-P series adopt hybrid mode which enables SLC Cache up to 3% of full disk capacity followed by TLC direct write to strike balance between burst performance and steady overall stability.

** Performance results are based on CrystalDiskMark 8.0.1 with file size 1000MB. Unit of 4KB items is I.O.P.S.

*** Performance results are based on AIDA 64 v5.98 with block size 1MB of Linear Read & Write Test

Performance may be different because ST and WT adopt different thermal solutions.

Table 3: Latency (QD1)

Capacity	Unit	400GB	800GB	1.6TB	3.2TB	6.4TB
Sequential Read	μs	43	42	41	37	41
Sequential Write		10	10	10	9	10
Random Read		63	63	64	64	87
Random Write		10	10	10	10	10

Note: Latency measured using 4KB(4,096 Bytes) transfer size with Queue Depth equal to 1 on a sequential and random workload

Table 4: Quality of Service (QoS)

Capacity	400GB	800GB	1.6TB	3.2TB	6.4TB
Quality of Service ^{1,2} (99.9%) (Unit: ms)					
Read Queue Depth 1	0.08	0.08	0.08	0.08	0.1
Write Queue Depth 1	0.02	0.02	0.02	0.02	0.02

Note:

¹Quality of Service measured using 4KB (4,096 bytes) transfer size on a random workload on a full Logical Block Address (LBA) span of the drive once the workload has reached steady state but including all background activities required for normal operation and data reliability.

²Based on Random 4KB QD=1 workloads, measured as the time taken for 99.9 percentile of commands to finish the round-trip from host to drive and back to host.

2.3 Electrical Specifications

2.3.1 Power Requirement

Table 5: Innodisk U.2 SSD 4TS2-P Power Requirement

Item	Symbol	Rating	Unit
Input voltage	V _{IN}	+12 DC +- 5%	V

2.3.2 Power Consumption

Table 6: Power Consumption

Mode	Power Consumption (W)
Read	10.3
Write	12.6
Idle	2.9
Power-on peak	8.7

* Target: U.2 SSD 4TS2-P 6.4TB

Note: Current results may vary depending on system components and power circuit design. Please refer to the test report for other capacities

2.4 Environmental Specifications

2.4.1 Temperature Ranges

Table 7: Temperature range for U.2 SSD 4TS2-P

Temperature	Range
Operating	Standard Grade: 0°C to +70°C Industrial Grade: -40°C ~ +85°C
Storage	-40°C ~ +85°C

2.4.2 Humidity

Relative Humidity: 10-95%, non-condensing

2.4.3 Shock and Vibration

Table 8: Shock/Vibration Testing for U.2 SSD 4TS2-P

Reliability	Test Conditions	Reference Standards
Vibration	7 Hz to 2K Hz, 20G, 3 axes	IEC 60068-2-6
Mechanical Shock	Duration: 0.5ms, 1500 G, 3 axes	IEC 60068-2-27

2.4.4 Mean Time between Failures (MTBF)

Table 10 summarizes the MTBF prediction results for various U.2 SSD 4TS2-P configurations. The analysis was performed using a RAM Commander™ failure rate prediction.

- **Failure Rate:** The total number of failures within an item population, divided by the total number of life units expended by that population, during a particular measurement interval under stated condition.
- **Mean Time between Failures (MTBF):** A basic measure of reliability for repairable items: The mean number of life units during which all parts of the item perform within their specified limits, during a particular measurement interval under stated conditions.

Table 9: U.2 SSD 4TS2-P MTBF

Product	Condition	MTBF (Hours)
Innodisk U.2 SSD 4TS2-P	Telcordia SR-332 GB, 25°C	>3,000,000

2.5 CE and FCC Compatibility

U.2 SSD 4TS2-P conforms to CE and FCC requirements.

2.6 RoHS Compliance

U.2 SSD 4TS2-P is fully compliant with RoHS directive.

2.7 Reliability

Table 10: U.2 SSD 4TS2-P TBW

Parameter	Value		
Flash endurance	3,000 P/E cycles		
Error Correct Code	Support (LDPC)		
Data Retention	Under 40°C: 1 Year at NAND Life End		
TBW* (Total Bytes Written) Unit: TB			
Capacity	Sequential workload	Client workload	Enterprise workload
400GB	1065	834	664
800GB	2130	1894	1328
1.6TB	4363	4145	2656
3.2TB	8727	8308	5312
6.4TB	17454	16752	7437
* Note: 1. Sequential: Mainly sequential write are estimated by PassMark Burnin Test 8.1 pro. 2. Client: Follow JESD218 Test method and JESD219A Workload, tested by ULINK. (The capacity lower than 64GB client workload is not specified in JEDEC219A, the values are estimated.) 3. Based on out-of-box performance.			

2.8 Transfer Mode

U.2 SSD 4TS2-P support following transfer mode:

PCIe Gen 4: 8GB/s

PCIe Gen 3: 4GB/s

PCIe Gen 2: 2GB/s

PCIe Gen 1: 1GB/s

2.9 Pin Assignment

Innodisk U.2 SSD 4TS2-P follows standard SFF-8639 spec as below. Mechanical details are documented in SFF-8639: Multifunction 6X Unshielded Connector.

Table 11: Innodisk U.2 SSD 4TS2-P Pin Assignment


Pin	Mate	Name
P1	3rd	WAKE#
P2	3rd	-
P3	2nd	PWRDIS
P4	1st	IfDet#
P5	2nd	Ground
P6	2nd	Ground
P7	2nd	-
P8	3rd	-
P9	3rd	-
P10	2nd	-
P11	3rd	ACTIVITY#
P12	1st	Ground
P13	2nd	+12V Precharge
P14	3rd	+12V
P15	3rd	+12V
S1	2nd	Ground
S2	3rd	-
S3	3rd	-
S4	2nd	Ground
S5	3rd	-
S6	3rd	-
S7	2nd	Ground
S8	2nd	Ground
S9	3rd	-
S10	3rd	-
S11	2nd	Ground
S12	3rd	-
S13	3rd	-
S14	2nd	Ground
S15	3rd	-
S16	2nd	Ground
S17	3rd	PETp1

Pin	Mate	Name
S18	3rd	PETn1
S19	2nd	Ground
S20	3rd	PERn1
S21	3rd	PERp1
S22	2nd	Ground
S23	3rd	PETp2
S24	3rd	PETn2
S25	2nd	Ground
S26	3rd	PERn2
S27	3rd	PERp2
S28	2nd	Ground
E1	3rd	-
E2	3rd	-
E3	3rd	-
E4	3rd	CLKREQ#
E5	3rd	PERST#
E6	3rd	-
E7	3rd	RefClk0+
E8	3rd	RefClk0-
E9	2nd	Ground
E10	3rd	PETp0
E11	3rd	PETn0
E12	2nd	Ground
E13	3rd	PERn0
E14	3rd	PERp0
E15	2nd	Ground
E16	3rd	-
E17	3rd	PETp3
E18	3rd	PETn3
E19	2nd	Ground
E20	3rd	PERn3
E21	3rd	PERp3
E22	2nd	Ground
E23	3rd	-
E24	3rd	-
E25	3rd	-

Table 12: Innodisk U.2 SSD 4TS2-P LED indicator

LED Color	Function
Green	Power on
	Access

2.10 Mechanical Dimensions


Figure 2: Innodisk U.2 SSD 4TS2-P

2.11 Assembly Weight

An Innodisk U.2 SSD 4TS2-P within NAND flash ICs, 512GB's weight is 14 grams approximately.

2.12 Seek Time

Innodisk U.2 SSD 4TS2-P is not a magnetic rotating design. There is no seek or rotational latency required.

2.13 NAND Flash Memory

Innodisk U.2 SSD 4TS2-P uses 3D TLC NAND flash memory, which is non-volatility, high reliability and high speed memory storage.

3. Theory of Operation

3.1 Overview

Figure 3 shows the operation of Innodisk U.2 SSD 4TS2-P from the system level, including the major hardware blocks.




Figure 3: Innodisk U.2 SSD 4TS2-P Block Diagram

Innodisk U.2 SSD 4TS2-P integrates a PCIe Gen 4x4 controller and NAND flash memories. Communication with the host occurs through the host interface, using the standard NVM protocol. Communication with the flash device(s) occurs through the flash interface.

3.2 PCIe Gen 4x4 Controller

Innodisk U.2 SSD 4TS2-P is a PCIe Gen 4x4 controller is compliant with NVMe 1.4, up to 32.0Gbps transfer speed. Also it is compliant with PCIe Gen 1, Gen 2, Gen 3, and Gen 4 specification. The controller supports up to 8 channels for flash interface.

3.3 Error Detection and Correction

Innodisk U.2 SSD 4TS2-P is designed with hardware LDPC ECC engine with hard-decision and soft-decision decoding. Low-density parity-check (LDPC) codes have excellent error correcting performance close to the Shannon limit when decoded with the belief-propagation (BP) algorithm using soft-decision information.

3.4 Wear-Leveling

Flash memory can be erased within a limited number of times. This number is called the **erase cycle limit** or **write endurance limit** and is defined by the flash array vendor. The erase cycle limit applies to each individual erase block in the flash device.

Innodisk U.2 SSD 4TS2-P uses a combination of two types of wear leveling- dynamic and static wear leveling- to distribute write cycling across an SSD and balance erase count of each block, thereby extending flash lifetime.

3.5 Bad Blocks Management

Bad Blocks are blocks that contain one or more invalid bits whose reliability are not guaranteed. The Bad Blocks may be presented while the SSD is shipped, or may develop during the life time of the SSD. When the Bad Blocks is detected, it will be flagged, and not be used anymore. The SSD implement Bad Blocks management, Bad Blocks replacement, Error Correct Code to avoid data error occurred. The functions will be enabled automatically to transfer data from Bad Blocks to spare blocks, and correct error bit.

3.6 iData Guard

Innodisk's iData Guard is a comprehensive data protection mechanism that functions before and after a sudden power outage to the SSD. Low-power detection terminates data writing before an abnormal power-off, while table-remapping after power-on deletes corrupt data and maintains data integrity. Innodisk's iData Guard provides effective power cycling management, preventing data stored in flash from degrading with use.

3.7 Garbage Collection/TRIM

Garbage collection and TRIM technology is used to maintain data consistency and perform continual data cleansing on SSDs. It runs as a background process, freeing up valuable controller resources while sorting good data into available blocks, and deleting bad blocks. It also significantly reduces write operations to the drive, thereby increasing the SSD's speed and lifespan.

3.8 Thermal Management

U.2 SSD 4TS2-P has built-in thermal sensor which can detect environment temperature of SSD. In the meantime, firmware will monitor the thermal sensor to prevent any failure of overheating. During extreme temperature, firmware will adjust the data transfer behavior to maintain the SSD's reliable operation.

3.9 Thermal Management

Thermal throttling is a protective mechanism designed to safeguard components from potential damage caused by excessive temperatures. When an SSD approaches a critical temperature threshold, Innodisk firmware activates the thermal throttling mechanism to regulate the SSD's temperature. Thermal throttling is crucial for SSDs since it prevents drive damage, which could otherwise result in data loss. However, it's worth noting that when thermal throttling is activated, read and write tasks may experience a reduction in speed.

3.10 Thermal Throttling

Thermal throttling is a protective mechanism designed to safeguard components from potential damage caused by excessive temperatures. When an SSD approaches a critical temperature threshold, Innodisk firmware activates the thermal throttling mechanism to regulate the SSD's temperature. Thermal throttling is crucial for SSDs since it prevents drive damage, which could otherwise result in data loss. However, it's worth noting that when thermal throttling is activated, read and write tasks may experience a reduction in speed.

3.11 iPower Guard

iPower Guard technology is a set of preventive measures that protect the SSD in an unstable power supply environment. This comprehensive package comprises safeguards for startup and shutdown to maintain device performance and ensure data integrity.

3.12 Die RAID

Die RAID is a controller function which leveraged user capacity to back up the data in NAND flash. Die RAID supported can ensure the user data in the NAND Flash more consistent in certain scenario. Innodisk U.2 SSD 4TS2-P series is default enable the Die RAID function for the industrial application.

3.13 SLC Cache

4TS2-P series adopt hybrid mode which enables SLC Cache up to 3% of full disk capacity followed by TLC direct write to strike balance between burst performance and steady overall stability. The SLC Cache buffer size are defined as table below.

Table 13: U.2 SSD 4TS2-P SLC cache

Capacity	400GB	800GB	1.6TB	3.2TB	6.4TB
SLC cache (GB)	12	24	49.2	65.5	131.1
SLC cache (%)	3	3	3	2	2

4. Installation Requirements

4.1 U.2 SSD 4TS2-P Pin Directions




Figure 4: Device Signal Segment Power Segment

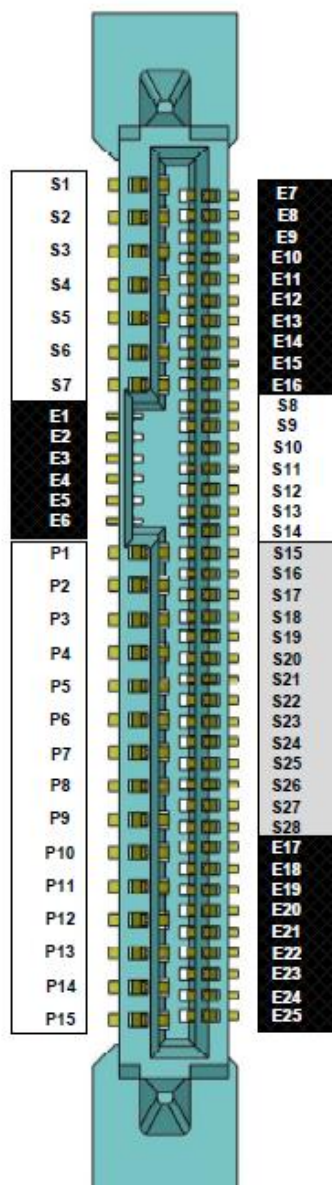


Figure 5: Signal Segment and Power Segment(Host/cable side)

4.2 Electrical Connections for U.2 SSD 4TS2-P

U.2 SSD 4TS2-P follows standard SFF-8639 spec, Mechanical details are documented in SFF-8639: Multifunction 6X Unshielded Connector, it is a total of 68 contacts, not all contacts may be utilized. U.2 SSD 4TS2-P only support PCIe interface, see more details in 2.9 *Pin Assignment*.

4.3 Device Drive

U.2 SSD 4TS2-P is compliant with NVMe 1.4. Both Operation System and BIOS should include NVMe driver to compatible with NVMe device. Nowadays, most of OS includes NVMe in-box driver now. For more information about the driver support in each OS, please visit the website <https://nvmexpress.org/drivers/>. For BIOS NVMe driver support please contact with motherboard manufacturers.

5. SMART / Health Information

This log page is used to provide SMART and general health information. The information provided is over the life of the controller and is retained across power cycles. More details about Set Features command; please refer to NVM Express 1.3

5.1 Get Log Page(Log Identifier 02h)

Innodisk 4TS2-P series SMART / Health Information Log are listed in following table.

Table 14: Get Log Page – SMART / Health Information Log

Bytes	Description																
0	<p>Critical Warning: This field indicates critical warnings for the state of the controller. Each bit corresponds to a critical warning type; multiple bits may be set. If a bit is cleared to '0', then that critical warning does not apply. Critical warnings may result in an asynchronous event notification to the host. Bits in this field represent the current associated state and are not persistent.</p> <table> <tr> <th>Bit</th><th>Definition</th></tr> <tr> <td>7:6</td><td>Reserved</td></tr> <tr> <td>5</td><td>If set to '1', then the Persistent Memory Region has become read-only or unreliable.</td></tr> <tr> <td>4</td><td>If set to '1', then the volatile memory backup device has failed. This field is only valid if the controller has a volatile memory backup solution.</td></tr> <tr> <td>3</td><td>If set to '1', then all of the media has been placed in read only mode. The controller shall not set this bit to '1' if the read-only condition on the media is a result of a change in the write protection state of a namespace.</td></tr> <tr> <td>2</td><td>If set to '1', then the NVM subsystem reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.</td></tr> <tr> <td>1</td><td>If set to '1', then a temperature is: a) greater than or equal to an over temperature threshold; or b) less than or equal to an under temperature threshold.</td></tr> <tr> <td>0</td><td>If set to '1', then the available spare capacity has fallen below the threshold.</td></tr> </table>	Bit	Definition	7:6	Reserved	5	If set to '1', then the Persistent Memory Region has become read-only or unreliable.	4	If set to '1', then the volatile memory backup device has failed. This field is only valid if the controller has a volatile memory backup solution.	3	If set to '1', then all of the media has been placed in read only mode. The controller shall not set this bit to '1' if the read-only condition on the media is a result of a change in the write protection state of a namespace.	2	If set to '1', then the NVM subsystem reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.	1	If set to '1', then a temperature is: a) greater than or equal to an over temperature threshold; or b) less than or equal to an under temperature threshold.	0	If set to '1', then the available spare capacity has fallen below the threshold.
Bit	Definition																
7:6	Reserved																
5	If set to '1', then the Persistent Memory Region has become read-only or unreliable.																
4	If set to '1', then the volatile memory backup device has failed. This field is only valid if the controller has a volatile memory backup solution.																
3	If set to '1', then all of the media has been placed in read only mode. The controller shall not set this bit to '1' if the read-only condition on the media is a result of a change in the write protection state of a namespace.																
2	If set to '1', then the NVM subsystem reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.																
1	If set to '1', then a temperature is: a) greater than or equal to an over temperature threshold; or b) less than or equal to an under temperature threshold.																
0	If set to '1', then the available spare capacity has fallen below the threshold.																
2:1	<p>Composite Temperature: Contains a value corresponding to a temperature in degrees Kelvin that represents the current composite temperature of the controller and namespace(s) associated with that controller. The manner in which this value is computed is implementation specific and may not represent the actual temperature of</p>																

	<p>any physical point in the NVM subsystem. The value of this field may be used to trigger an asynchronous event.</p> <p>Warning and critical overheating composite temperature threshold values are reported by the WCTEMP and CCTEMP fields in the Identify Controller data structure.</p>
3	Available Spare: Contains a normalized percentage (0 to 100%) of the remaining spare capacity available.
4	Available Spare Threshold: When the Available Spare falls below the threshold indicated in this field, an asynchronous event completion may occur. The value is indicated as a normalized percentage (0 to 100%).
5	<p>Percentage Used: Contains a vendor specific estimate of the percentage of NVM subsystem life used based on the actual usage and the manufacturer's prediction of NVM life. A value of 100 indicates that the estimated endurance of the NVM in the NVM subsystem has been consumed, but may not indicate an NVM subsystem failure. The value is allowed to exceed 100. Percentages greater than 254 shall be represented as 255. This value shall be updated once per power-on hour (when the controller is not in a sleep state).</p> <p>Refer to the JEDEC JESD218A standard for SSD device life and endurance measurement techniques.</p>
31:6	Reserved
47:32	<p>Data Units Read: Contains the number of 512 byte data units the host has read from the controller; this value does not include metadata. This value is reported in thousands (i.e., a value of 1 corresponds to 1000 units of 512 bytes read) and is rounded up. When the LBA size is a value other than 512 bytes, the controller shall convert the amount of data read to 512 byte units.</p> <p>For the NVM command set, logical blocks read as part of Compare and Read operations shall be included in this value.</p>
63:48	<p>Data Units Written: Contains the number of 512 byte data units the host has written to the controller; this value does not include metadata. This value is reported in thousands (i.e., a value of 1 corresponds to 1000 units of 512 bytes written) and is rounded up. When the LBA size is a value other than 512 bytes, the controller shall convert the amount of data written to 512 byte units.</p> <p>For the NVM command set, logical blocks written as part of Write operations shall be included in this value. Write Uncorrectable commands shall not impact this value.</p>
79:64	<p>Host Read Commands: Contains the number of read commands completed by the controller.</p> <p>For the NVM command set, this is the number of Compare and Read commands.</p>
95:80	<p>Host Write Commands: Contains the number of write commands completed by the controller.</p> <p>For the NVM command set, this is the number of Write commands.</p>

111:96	Controller Busy Time: Contains the amount of time the controller is busy with I/O commands. The controller is busy when there is a command outstanding to an I/O Queue (specifically, a command was issued via an I/O Submission Queue Tail doorbell write and the corresponding completion queue entry has not been posted yet to the associated I/O Completion Queue). This value is reported in minutes.
127:112	Power Cycles: Contains the number of power cycles.
143:128	Power On Hours: Contains the number of power-on hours. This may not include time that the controller was powered and in a non-operational power state.
159:144	Unsafe Shutdowns: Contains the number of unsafe shutdowns. This count is incremented when a shutdown notification (CC.SHN) is not received prior to loss of power.
175:160	Media and Data Integrity Errors: Contains the number of occurrences where the controller detected an unrecovered data integrity error. Errors such as uncorrectable ECC, CRC checksum failure, or LBA tag mismatch are included in this field.
191:176	Number of Error Information Log Entries: Contains the number of Error Information log entries over the life of the controller.
195:192	Warning Composite Temperature Time: Contains the amount of time in minutes that the controller is operational and the Composite Temperature is greater than or equal to the Warning Composite Temperature Threshold (WCTEMP) field and less than the Critical Composite Temperature Threshold (CCTEMP) field in the Identify Controller data structure. If the value of the WCTEMP or CCTEMP field is 0h, then this field is always cleared to 0h regardless of the Composite Temperature value.
199:196	Critical Composite Temperature Time: Contains the amount of time in minutes that the controller is operational and the Composite Temperature is greater than the Critical Composite Temperature Threshold (CCTEMP) field in the Identify Controller data structure. If the value of the CCTEMP field is 0h, then this field is always cleared to 0h regardless of the Composite Temperature value.
201:200	Temperature Sensor 1: Controller's Tj temperature
203:202	Temperature Sensor 2: Flash package's Tj temperature (Channel #0 CE #0). This Flash package is located the closet to the controller IC on M.2 family.
205:204	Temperature Sensor 3: Flash package's Tj temperature (Channel #0 CE #0). This Flash package is located the closet to the controller IC on M.2 family.
207:206	Temperature Sensor 4: Flash package's Tj temperature (Channel #7 CE #0).
209:208	Temperature Sensor 5: Flash Tj max temperature from Channel #0 to Channel #3 Flash packages.
211:210	Temperature Sensor 6: Flash Tj max temperature from Channel #4 to Channel #7 Flash packages.

213:212	Temperature Sensor 7: Flash Tj minimum temperature from Channel #0 to Channel #3 Flash packages.
215:214	Temperature Sensor 8: Flash Tj minimum temperature from Channel #4 to Channel #7 Flash packages.
219:216	Thermal Management Temperature 1 Transition Count: Contains the number of times the controller transitioned to lower power active power states or performed vendor specific thermal management actions while minimizing the impact on performance in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature (refer to section 8.4.5) (i.e., the Composite Temperature rose above the Thermal Management Temperature 1.) This counter shall not wrap once it reaches its maximum value. A value of zero, indicates that this transition has never occurred or this field is not implemented.
223:220	Thermal Management Temperature 2 Transition Count: Contains the number of times the controller transitioned to lower power active power states or performed vendor specific thermal management actions regardless of the impact on performance (e.g., heavy throttling) in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature (refer to section 8.4.5) (i.e., the Composite Temperature rose above the Thermal Management Temperature 2.) This counter shall not wrap once it reaches its maximum value. A value of zero, indicates that this transition has never occurred or this field is not implemented.
227:224	Total Time For Thermal Management Temperature 1: Contains the number of seconds that the controller had transitioned to lower power active power states or performed vendor specific thermal management actions while minimizing the impact on performance in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature (refer to section 8.4.5). This counter shall not wrap once it reaches its maximum value. A value of zero, indicates that this transition has never occurred or this field is not implemented.
231:228	Total Time For Thermal Management Temperature 2: Contains the number of seconds that the controller had transitioned to lower power active power states or performed vendor specific thermal management actions regardless of the impact on performance (e.g., heavy throttling) in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature (refer to section 8.4.5). This counter shall not wrap once it reaches its maximum value. A value of zero, indicates that this transition has never occurred or this field is not implemented.
511:232	Reserved

The innodisk U.2 SSD 4TS2-P series thermal sensor take ambient air temperature as a reference with any airflow condition, and the data can refer to iSMART.

Notes: More detailed health info has been defined by innodisk and will be shown on iSMART V5.3.21 (or later version).

7. Part Number Rule

CODE	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
	D	S	E	U	2	-	1	T	6	D	P	1	K	C	A	E	F	-	X	X	X
Definition																					
Code 1 st (Disk)											Code 14 th (Operation Temperature)										
D : Disk											C: Standard Grade (0℃~ +70℃)										
											W: Industrial Grade (-40℃~ +85℃)										
Code 2 nd (Feature set)											Code 15 th (Internal control)										
S : Edge server series											A~Z: BGA PCB version										
Code 3 rd ~5 th (Form factor)											Code 16 th (Channel of data transfer)										
EU2: U.2 SSD											E: Eight Channels										
Code 7 th ~9 th (Capacity)											Code 17 th (Flash Type)										
400: 400GB			800: 800GB			1T6:1.6TB					F: Kioxia 3D TLC										
3T2: 3.2TB			6T4: 6.4TB																		
Code 10 th ~12 th (Controller)											Code 18 th (Optional Function)										
DP1: PCIe 4TS2-P series																					
Code 13 th (Flash mode)											Code 19 th ~ 21 st (Customize code)										
K: 112 Layers 3D TLC																					