

M.2 (P80)

4TS2-P Series

Customer: _____
Customer
Part
Number: _____
Innodisk
Part
Number: _____
Innodisk
Model Name: _____
Date: _____

Innodisk Approver	Customer Approver

**Total Solution For
Industrial Flash Storage**

Table of contents

1. PRODUCT OVERVIEW	8
1.1 INTRODUCTION OF INNODISK M.2 (P80) 4TS2-P	8
1.2 PRODUCT VIEW AND MODELS	8
PCIe INTERFACE	9
2. PRODUCT SPECIFICATIONS	10
2.1 CAPACITY AND DEVICE PARAMETERS	10
2.2 PERFORMANCE	10
2.3 ELECTRICAL SPECIFICATIONS	12
2.3.1 Power Requirement	12
2.3.2 Power Consumption	12
2.4 ENVIRONMENTAL SPECIFICATIONS	12
2.4.1 Temperature Ranges	12
2.4.2 Humidity	12
2.4.3 Shock and Vibration	12
2.4.4 Mean Time between Failures (MTBF)	13
2.5 CE AND FCC COMPATIBILITY	13
2.6 RoHS COMPLIANCE	13
2.7 RELIABILITY	14
2.8 TRANSFER MODE	14
2.9 PIN ASSIGNMENT	15
2.10 MECHANICAL DIMENSIONS	16
2.11 ASSEMBLY WEIGHT	18
2.12 SEEK TIME	18
2.13 NAND FLASH MEMORY	18
3. THEORY OF OPERATION	19
3.1 OVERVIEW	19
3.2 PCIe GEN. 4x4 CONTROLLER	19
3.3 ERROR DETECTION AND CORRECTION	20
3.4 WEAR-LEVELING	20
3.5 BAD BLOCKS MANAGEMENT	20
3.6 iDATAGUARD	20
3.7 GARBAGE COLLECTION/TRIM	20
3.8 THERMAL MANAGEMENT	21
3.9 THERMAL THROTTLING	21
3.10 iPOWERGUARD	21
3.11 DIE RAID	21
3.12 SLC CACHE	21

- 4. INSTALLATION REQUIREMENTS.....22**
 - 4.1 M.2 (P80) 4TS2-P PIN DIRECTIONS22**
 - 4.2 ELECTRICAL CONNECTIONS FOR M.2 (P80) 4TS2-P22**
 - 4.3 DEVICE DRIVE23**
- 5. SMART / HEALTH INFORMATION.....24**
 - 5.1 GET LOG PAGE(LOG IDENTIFIER 02H)24**
- 6. PART NUMBER RULE28**

Features:

- PCIe Gen. 4x 4, NVMe SSD
- Kioxia 3D TLC NAND
- M.2 2280-D5-M
- Standard & Wide temperature
- iPowerguard
- iDataguard
- Thermal Throttling Management
- 256-bit AES hardware-based encryption
- Hybrid Write Mode with SLC Cache Enable
- Support AES + TCG OPAL function

Performance:

- Sequential Read up to 6,950 MB/s
- Sequential Write up to 4,650 MB/s

Power Requirements:

Input Voltage:	3.3V±5%
Max Operating Wattage (R/W):	6.7W
Idle Wattage:	2.3W

Reliability:

Capacity	TBW (Enterprise)	DWPD
400GB	664	1
800GB	1,328	1
1.6TB	2,656	1
3.2TB	5,312	1

Data Retention	1 Year
Warranty	5 Years

DWPD based on Enterprise workload with 5 years limited warranty

1 year data retention is at NAND life end

For warranty details, please refer to:

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

REVISION HISTORY

Revision	Description	Date
V1.0	First release	Jul., 2024

List of Tables

TABLE 1: DEVICE PARAMETERS	10
TABLE 2: PERFORMANCE - 112 LAYERS 3D TLC	10
TABLE 3: LATENCY (QD1)	11
TABLE 4: QUALITY OF SERVICE (QOS)	11
TABLE 5: INNODISK M.2 (P80) 4TS2-P POWER REQUIREMENT	12
TABLE 6: TYPICAL POWER CONSUMPTION	12
TABLE 7: TEMPERATURE RANGE FOR M.2 (P80) 4TS2-P	12
TABLE 8: SHOCK/VIBRATION TESTING FOR M.2 (P80) 4TS2-P	12
TABLE 9: M.2 (P80) 4TS2-P MTBF	13
TABLE 10: M.2 (P80) 4TS2-P TBW	14
TABLE 11: INNODISK M.2 (P80) 4TS2-P PIN ASSIGNMENT	15
TABLE 12: INNODISK M.2 (P80) 4TS2-P LED INDICATOR	16
TABLE 13: M.2 (P80) 4TS2-P SLC CACHE	21
TABLE 14: GET LOG PAGE – SMART / HEALTH INFORMATION LOG	24

List of Figures

FIGURE 1: INNODISK M.2 (P80) 4TS2-P (STANDARD)	8
FIGURE 2: INNODISK M.2 (P80) 4TS2-P (WIDE-TEMPERATURE)	8
FIGURE 3: INNODISK M.2 (P80) 4TS2-P WITH HEAT-SPREADING COPPER LAYER DIAGRAM	16
FIGURE 4: INNODISK M.2 (P80) 4TS2-P WITH HEATSINK DIAGRAM	17
FIGURE 5: INNODISK M.2 (P80) 4TS2-P	17
FIGURE 6: INNODISK M.2 (P80) 4TS2-P BLOCK DIAGRAM	19
FIGURE 7: SIGNAL SEGMENT AND POWER SEGMENT	22

1. Product Overview

1.1 Introduction of Innodisk M.2 (P80) 4TS2-P

Innodisk M.2 (P80) 4TS2-P is an NVM Express SSD designed as the standard M.2 form factor with PCIe interface and 3D TLC NAND Flash. M.2 (P80) 4TS2-P supports PCIe Gen. 4x4, and it is compliant with NVMe 1.4 providing excellent performance. M.2 (P80) 4TS2-P with heat-spreading design dissipate heat generating from IC making SSD perform more steady. M.2 (P80) 4TS2-P has Die RAID protection to reduce bad blocks happening and optimize data integrity.

In addition, 4TS2-P series adopt hybrid mode which enables SLC Cache followed by TLC direct write to strike balance between burst performance and steady overall stability.

Innodisk M.2 (P80) 4TS2-P provides ultra-speed and high IOPS and offers maximum capacity up to 3.2TB, making the SSD optimal for server and heavy data workload applications.

Innodisk M.2 (P80) 4TG2-P is designed with AES engine, which is built-in the controller. When controller receives the data package from host, AES engine encrypts the data package and save the encrypted data into NAND flash. Thus, unauthorized personal has no access to decrypt the data in NAND flash.

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 M.2 (P80) 4TS2-P is available in follow capacities within 3D TLC flash ICs.

M.2 (P80) 4TS2-P 400GB M.2 (P80) 4TS2-P 1.6TB

M.2 (P80) 4TS2-P 800GB M.2 (P80) 4TS2-P 3.2TB



Figure 1: Innodisk M.2 (P80) 4TS2-P (Standard)



Figure 2: Innodisk M.2 (P80) 4TS2-P (Wide-temperature)

PCIe Interface

Innodisk M.2 (P80) 4TS2-P supports PCIe Gen. 4 interface and compliant with NVMe 1.4. M.2 (P80) 4TS2-P can work under PCIe Gen. 1, Gen. 2, Gen. 3 and 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://nvmexpress.org/drivers/>.

2. Product Specifications

2.1 Capacity and Device Parameters

M.2 (P80) 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

2.2 Performance

Burst Transfer Rate: 8 GB/s

Table 2: Performance - 112 Layers 3D TLC

Capacity	Unit	400GB	800GB	1.6TB	3.2TB
Sequential** Read (Q8T1)	MB/s	5,550	6,950	6,650	6,650
Sequential** Write (Q8T1)		2,250	4,050	4,650	4,500
Sustained Sequential Read (Avg.)***		1,950	2,450	2,400	2,150
Sustained Sequential Write (Avg.)***		460	910	1,300	1,000
4KB Random** Read (Q32T16)	IOPS	453,000	814,000	819,000	816,000
4KB Random** Write (Q32T16)		581,000	599,000	698,000	700,000

Note: * Performance results are measured in Room Temperature with Out-of-Box devices and may vary depending on overall system setup. In addition, 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.

** 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.

Performance is affected by thermal throttling if device temperatures is over 75C.

Table 3: Latency (QD1)

Capacity	Unit	400GB	800GB	1.6TB	3.2TB
Sequential Read	μs	46	44	45	51
Sequential Write		10	10	10	10
Random Read		68	68	68	86
Random Write		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
Quality of Service ^{1,2} (99.9%) (Unit: ms)				
Read Queue Depth 1	0.08	0.08	0.08	0.1
Write Queue Depth 1	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 M.2 (P80) 4TS2-P Power Requirement

Item	Symbol	Rating	Unit
Input voltage	V _{IN}	+3.3 DC +- 5%	V

2.3.2 Power Consumption

Table 6: Typical Power Consumption

Mode	Power Consumption (W)
Read	5.7
Write	6.7
Idle	2.3
Power on peak	8.3

Target: 3.2TB M.2 (P80) 4TS2-P

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 M.2 (P80) 4TS2-P

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

2.4.2 Humidity

Relative Humidity: 10-95%, non-condensing

2.4.3 Shock and Vibration

Table 8: Shock/Vibration Testing for M.2 (P80) 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 9 summarizes the MTBF prediction results for various M.2 (P80) 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: M.2 (P80) 4TS2-P MTBF

Product	Condition	MTBF (Hours)
Innodisk M.2 (P80) 4TS2-P	Telcordia SR-332 GB, 25°C	>3,000,000

2.5 CE and FCC Compatibility

M.2 (P80) 4TS2-P conforms to CE and FCC requirements.

2.6 RoHS Compliance

M.2 (P80) 4TS2-P is fully compliant with RoHS directive.

2.7 Reliability

Table 10: M.2 (P80) 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	1,065	985	664
800GB	2,130	2,065	1,328
1.6TB	4,363	4,413	2,656
3.2TB	8,727	9,136	5,312
* Note: 1. Sequential: Mainly sequential write are estimated by PassMark Burnin Test v8.1 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

M.2 (P80) 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 M.2 (P80) 4TS2-P follows standard M.2 spec, socket 3, key M PCIe-based SSD pinout. See Table 11 for M.2 (P80) 4TS2-P pin assignment.

Table 11: Innodisk M.2 (P80) 4TS2-P Pin Assignment

Signal Name	Pin #	Pin #	Signal Name
		75	GND
3.3V	74	73	GND
3.3V	72	71	GND
3.3V	70	69	NC
NC	68	67	NC
Notch	66	65	Notch
Notch	64	63	Notch
Notch	62	61	Notch
Notch	60	59	Notch
NC	58		
NC	56	57	GND
NC	54	55	REFCLKp
CLKREQ# (I/O)(0/3.3V)	52	53	REFCLKn
PERST# (I)(0/3.3V)	50	51	GND
NC	48	49	PERp0
NC	46	47	PERn0
ALERT	44	45	GND
NC(reserved for SMB_DATA)(I/O)(0/1.8V)	42	43	PETp0
NC(reserved for SMB_CLK)	40	41	PETn0
GND	38	39	GND
NC	36	37	PERp1
NC	34	35	PERn1
GND	32	33	GND
NC	30	31	PETp1
NC	28	29	PETn1
NC	26	27	GND
NC	24	25	PERp2
NC	22	23	PERn2
NC	20	21	GND
3.3V	18	19	PETp2
3.3V	16	17	PETn2
3.3V	14	15	GND
3.3V	12	13	PERp3
LED#(O)(OD)	10	11	PERn3
NC	8	9	GND
NC	6	7	PETp3
3.3V	4	5	PETn3
3.3V	2	3	GND
		1	GND

Table 12: Innodisk M.2 (P80) 4TS2-P LED indicator

LED Color	Function
Green	Power on
	Access

2.10 Mechanical Dimensions

M.2 Type 2280-D5-M with heat-spreading copper layer (Default accessory for ST)




Figure 3: Innodisk M.2 (P80) 4TS2-P with heat-spreading copper layer diagram

M.2 Type 2280-D5-M with heatsink (Default accessory for WT)




Figure 4: Innodisk M.2 (P80) 4TS2-P with heatsink diagram

M.2 Type 2280-D5-M




Figure 5: Innodisk M.2 (P80) 4TS2-P

2.11 Assembly Weight

Innodisk M.2 (P80) 4TS2-P within NAND flash ICs, 3.2TB's weight is 11 grams approximately.

2.12 Seek Time

Innodisk M.2 (P80) 4TS2-P is not a magnetic rotating design. There is no seek or rotational latency required.

2.13 NAND Flash Memory

Innodisk M.2 (P80) 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 6 shows the operation of Innodisk M.2 (P80) 4TS2-P from the system level, including the major hardware blocks.




Figure 6: Innodisk M.2 (P80) 4TS2-P Block Diagram

Innodisk M.2 (P80) 4TS2-P integrates a PCIe Gen III x4 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. The AES engine was built-in the DP2 controller. When M.2 (P80) 4TS2-P is initiated with Firmware, AES engine will generate a random number to be an AES key. Each SSD has a unique AES key when it leaves the factory.

3.2 PCIe Gen. 4x4 Controller

Innodisk M.2 (P80) 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 M.2 (P80) 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 M.2 (P80) 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 iDataGuard

Innodisk's iDataGuard 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 iDataGuard 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

M.2 (P80) 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 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.

3.10 iPowerGuard

iPowerGuard 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.

3.11 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 M.2 (P80) 4TS2-P series is default enable the Die RAID function for the industrial application.

3.12 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: M.2 (P80) 4TS2-P SLC cache

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

3.13 TCG OPAL

OPAL is a set of specifications for features of data storage devices that enhance security. These specifications are published by the Trusted Computing Group’s Storage Work Group. Innodisk 4TS2-P is compliant with TCG OPAL 2.0(*1). The capability of TCG OPAL Security mode allows multiple users with independent access control to read/write/erase independent data areas (LBA ranges). Each locking range adjusts by authenticated authority. Note that by default there is a single “Global Range” that encompasses the whole user data area. In TCG Opal Security Mode, Revert, Revert SP and GenKey command can erase all of data including global range and locking range; in the meantime generate the new encrypted key.

*1. You need to install TCG OPAL software to implement OPAL function, which is supplied by TCG OPAL software developed company

4. Installation Requirements

4.1 M.2 (P80) 4TS2-P Pin Directions




Figure 7: Signal Segment and Power Segment

4.2 Electrical Connections for M.2 (P80) 4TS2-P

M.2 interconnect is based on a 75 position Edge Card connector. The 75 position connector is intended to be keyed so as to distinguish between families of host interfaces and the various Sockets used in general Platforms. M.2 (P80) 4TS2-P is compliant with M.2 Socket 3 key M. M.2

(P80) 4TS2-P is compatible with host connector H4.2.

4.3 Device Drive

M.2 (P80) 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 manufacture.

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 1.4

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 border="1"> <thead> <tr> <th>Bit</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>If set to '1', then the available spare space has fallen below the threshold.</td> </tr> <tr> <td>01</td> <td>If set to '1', then a temperature is above an over temperature threshold or below an under temperature threshold.</td> </tr> <tr> <td>02</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>03</td> <td>If set to '1', then the media has been placed in read only mode.</td> </tr> <tr> <td>04</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>07:05</td> <td>Reserved</td> </tr> </tbody> </table>	Bit	Definition	00	If set to '1', then the available spare space has fallen below the threshold.	01	If set to '1', then a temperature is above an over temperature threshold or below an under temperature threshold.	02	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.	03	If set to '1', then the media has been placed in read only mode.	04	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.	07:05	Reserved
Bit	Definition														
00	If set to '1', then the available spare space has fallen below the threshold.														
01	If set to '1', then a temperature is above an over temperature threshold or below an under temperature threshold.														
02	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.														
03	If set to '1', then the media has been placed in read only mode.														
04	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.														
07:05	Reserved														
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 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	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). Refer to the JEDEC JESD218A standard for SSD device life and endurance measurement techniques.
31:6	Reserved
47:32	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. For the NVM command set, logical blocks read as part of Compare and Read operations shall be included in this value.
63:48	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. 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.
79:64	Host Read Commands: Contains the number of read commands completed by the controller. For the NVM command set, this is the number of Compare and Read commands.
95:80	Host Write Commands: Contains the number of write commands completed by the controller. For the NVM command set, this is the number of Write commands.
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 M.2 (P80) 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).

6. 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	M	2	8	-	1	T	6	D	P	2	K	C	C	E	F	(H)	-	X	X
Definition																					
Code 1 st (Disk)											Code 14 th (Operation Temperature)										
D : Disk											C: Standard Grade (0°C~ +70°C)										
Code 2 nd (Feature set)											W: Industrial Grade (-40°C~ +85°C)										
S : Edge server series																					
Code 3 rd ~5 th (Form factor)											Code 15 th (Internal control)										
M28: M.2 Type 2280-D5-M											A~Z: BGA PCB version.										
Code 7 th ~9 th (Capacity)											Code 16 th (Channel of data transfer)										
400: 400GB			800: 800GB			1T6:1.6TB					E: Eight Channels										
3T2: 3.2TB																					
Code 10 th ~12 th (Controller)											Code 17 th (Flash Type)										
DP2: PCIe 4TG2-P with AES+TCG OPAL function											F: Kioxia 3D TLC										
Code 13 th (Flash mode)											Code 18 th (Optional function)										
K: 3D TLC 112 layers											H: with heatsink accessory (for WT)										
											Code 20 th ~ (Customize code)										