innodisk

Total Solution For Industrial Flash Storage

M.2 (P80)

4IE2 Series

Ultra iSLC

Customer: _____

Customer Part Number: _____

Innodisk Part Number: _____

Innodisk Model Name: _____

Date: _____

Innodisk Approver	Customer Approver

Features:

- PCIe Gen 4 x4, NVMe SSD
- Kioxia 3D TLC NAND
- M.2 Type 2280
- Standard/Wide-temperature
- Ultra iSLC technology
- iPower Guard
- iData Guard
- Support PLP (iCell) function (optional)
- 256-bit AES hardware-based encryption
- Support TCG OPAL function (optional)
- Support PCIe PLA_S3 & PLN (Optional)

Performance:

- Sequential Read up to 3,550 MB/s
- Sequential Write up to 3,300 MB/s

Power Requirements:

Input Voltage:	3.3V± 5%
Max Operating Wattage (R/W):	5.1W
Idle Wattage:	1.7W

Reliability:

Capacity	TBW	DWPD
40GB	2003	30.2
80GB	4006	30.2
160GB	7813	29.4
320GB	9766	18.4
640GB	20161	19.0

Data Retention	1 Year
Warranty	5 Years

For warranty details, please refer to:

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

Table of contents

Revision History

Revision	Description	Date
V1.0	First Release	Jul., 2024
V1.1	Remove Sustained Performance Update Pin Assignment Add PLA_S3 & PLN	Aug., 2024

List of Tables

List of Figures

1. Product Overview

1.1 Introduction of Innodisk M.2 (P80) 4IE2

The Innodisk M.2 (P80) 4IE2 is a DRAM-less NVMe SSD designed with a PCIe interface and industrial-grade 3D TLC NAND Flash. It supports PCIe Gen 4 x4 and complies with NVMe 1.4, offering exceptional top and sustained performance. With advanced error detection and correction (ECC) functions, the module ensures comprehensive End-to-End Data Path Protection, safeguarding data transmission between the host system and the NAND Flash.

Additionally, the Innodisk M.2 (P80) 4IE2 features an integrated AES engine within its controller. Upon receiving a data package from the host, the AES engine encrypts the data and stores it in the NAND Flash. This encryption mechanism ensures that unauthorized individuals cannot access or decrypt the data stored in the NAND Flash.

1.2 Product View and Models

Innodisk M.2 (P80) 4IE2 is available in follow capacities with industrial 3D TLC flash ICs.

M.2 (P80) 4IE2 40GB
M.2 (P80) 4IE2 80GB
M.2 (P80) 4IE2 160GB
M.2 (P80) 4IE2 320GB
M.2 (P80) 4IE2 640GB

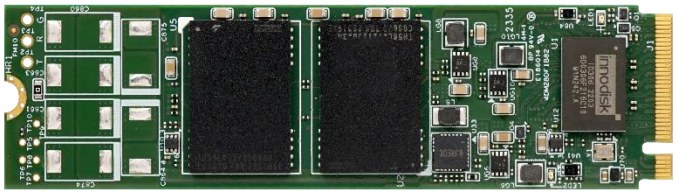

Figure 1: Innodisk M.2 (P80) 4IE2 (Standard)

Figure 2: Innodisk M.2 (P80) 4IE2 (Wide-temperature)

1.3 PCIe Interface

Innodisk M.2 (P80) 4IE2 supports PCIe Gen 4 interface and compliant with NVMe 1.4. M.2 (P80) 4IE2 can work under PCIe Gen 1, Gen 2 and Gen 3.

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) 4IE2 device parameters are shown in Table 1.

Table 1: Device parameters

Capacity	Cylinders	Heads	Sectors	LBA	User Capacity(MB)
40GB	16383	16	63	78161328	38164
80GB				156301488	76319
160GB				312581808	152627
320GB				625142448	305245
640GB				78161328	38164

2.2 Performance

Burst Transfer Rate: 8 GB/s

Table 2: Performance- 112 Layers 3D TLC

Capacity	Unit	40GB	80GB	160GB	320GB	640GB
Sequential* Read (Q8T1)	MB/s	1,400	2,850	3,550	3,550	3,450
Sequential* Write (Q8T1)		530	1,050	2,000	2,950	3,300
4KB Random** Read (Q32T16)	IOPS	126,000	242,000	470,000	626,000	572,000
4KB Random** Write (Q32T16)		132,000	262,000	448,000	599,000	626,000

Note: * Performance results are 4IE2 with Kioxia BiCS5 NAND composition measured in Room Temperature with Out-of-Box devices and may vary depending on overall system setup. In addition, 4IE2 series adopt hybrid mode which enables SLC cache followed by TLC direct write to strike balance between burst performance and steady overall stability.

Note: ** Performance results are based on CrystalDiskMark 8.0.1 with file size 1000MB. Unit of 4KB item is IOPS.

2.3 Electrical Specifications

2.3.1 Power Requirement

Table 3: Innodisk M.2 (P80) 4IE2 Power Requirement

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

2.3.2 Power Consumption

Table 4: Power Consumption

Mode	Power Consumption (W)
Read	5.0
Write	5.1
Idle	1.7
Power on peak	5.2

Target: M.2 (P80) 4IE2 640GB

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 5: Temperature range for M.2 (P80) 4IE2

Temperature	Range
Operating	Standard Grade: 0°C to +70°C Industrial Grade: -40°C to +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 6: Shock/Vibration Testing for M.2 (P80) 4IE2

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 7 summarizes the MTBF prediction results for various M.2 (P80) 4IE2 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 7: M.2 (P80) 4IE2 MTBF

Product	Condition	MTBF (Hours)
Innodisk M.2 (P80) 4IE2	Telcordia SR-332 GB, 25°C	>3,000,000

2.5 CE and FCC Compatibility

M.2 (P80) 4IE2 conforms to CE and FCC requirements.

Table 8: M.2 (P42) 4IE2 ESD

Reliability	Reference standards
Electrostatic Discharge (ESD)	EC 61000-4-2 ESD

2.6 RoHS Compliance

M.2 (P80) 4IE2 is fully compliant with RoHS directive.

2.7 Reliability

Table 9: M.2 (P80) 4IE2 TBW

Parameter	Value	
Read Cycles	Unlimited Read Cycles	
Flash endurance	100,000 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**
40GB	3551	2003
80GB	7102	4006
160GB	14204	7813
320GB	28408	9766
640GB	56816	20161

* 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.
3. Based on out-of-box performance.

2.8 Transfer Mode

M.2 (P80) 4IE2 support following transfer mode:

PCIe Gen 4: 2 GB/s

PCIe Gen 3: 1 GB/s

PCIe Gen 2: 500 MB/s

PCIe Gen 1: 250 MB/s

2.9 Pin Assignment

Innodisk M.2 (P80) 4IE2 follows standard M.2 spec, socket 3 key M PCIe-based SSD pinout. See Table 10 for M.2 (P80) 4IE2 pin assignment.

Table 10: Innodisk M.2 (P80) 4IE2 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	57	GND
NC	56	55	REFCLKp
NC	54	53	REFCLKn
CLKREQ# (I/O) (0V/1.8V/3.3V)	52	51	GND
PERST# (I) (0V/1.8V/3.3V)	50	49	PERp0
NC	48	47	PERn0
NC	46	45	GND
ALERT# (I) (0/1.8V)	44	43	PETp0
SMB_DATA (I/O) (0/1.8V)	42	41	PETn0
SMB_CLK (I/O) (0/1.8V)	40	39	GND
NC	38	37	PERp1
NC	36	35	PERn1
NC	34	33	GND
NC	32	31	PETp1
PLA_S3# (I)(0/1.8/3.3V) or NC	30	29	PETn1
NC	28	27	GND
NC	26	25	PERp2
NC	24	23	PERn2
NC	22	21	GND
NC	20	19	PETp2
3.3V	18	17	PETn2
3.3V	16	15	GND
3.3V	14	13	PERp3

3.3V	12	11	PERn3
LED1# (O) (OD)	10	9	GND
PLN# (O)(0/1.8/3.3V) or NC	8	7	PETp3
NC	6	5	PETn3
3.3V	4	3	GND
3.3V	2	1	GND

2.10 Mechanical Dimensions

M.2 Type 2280

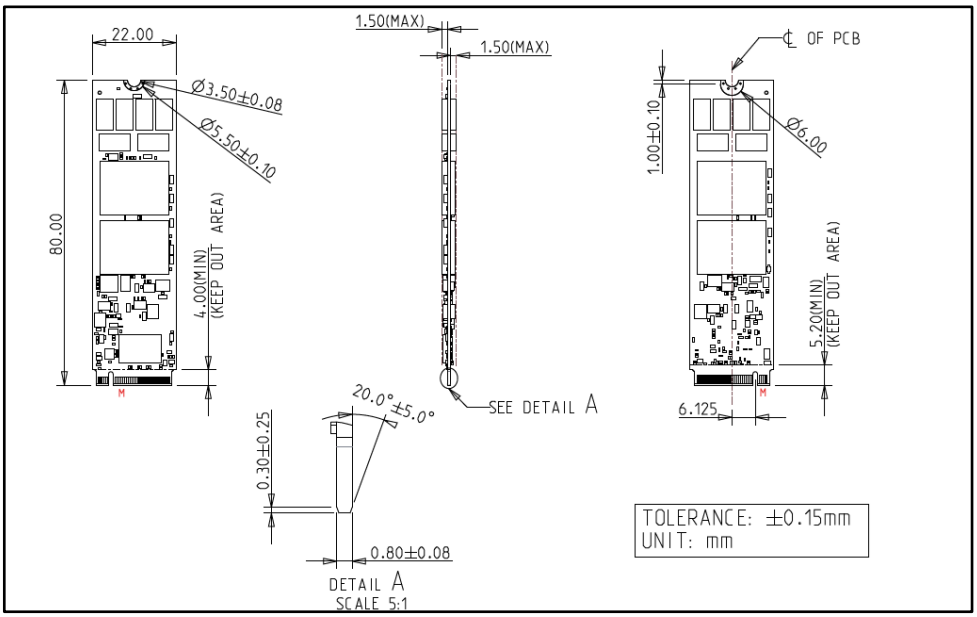

Figure 3: Innodisk M.2 (P80) 4IE2 mechanical drawing

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

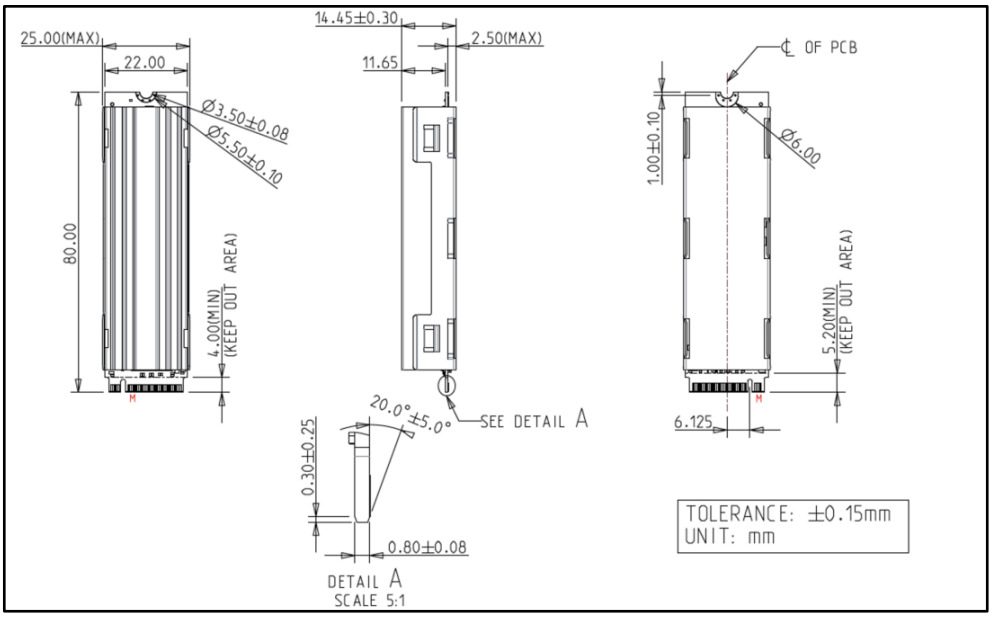

Figure 4: Innodisk M.2 (P80) 4IE2 mechanical drawing with heatsink diagram

2.11 Assembly Weight

An Innodisk M.2 (P80) 4IE2 within NAND flash ICs, 640GB's weight is 10 grams approximately.

2.12 Seek Time

Innodisk M.2 (P80) 4IE2 is not of magnetic rotating design. There is no seek or rotational latency.

2.13 NAND Flash Memory

Innodisk M.2 (P80) 4IE2 uses industrial 3D TLC NAND flash memory, which is non-volatility, high reliability and high speed memory storage.

3. Theory of Operation

3.1 Overview

Figure 5 shows the operation of Innodisk M.2 (P80) 4IE2 from the system level, including the major hardware blocks.

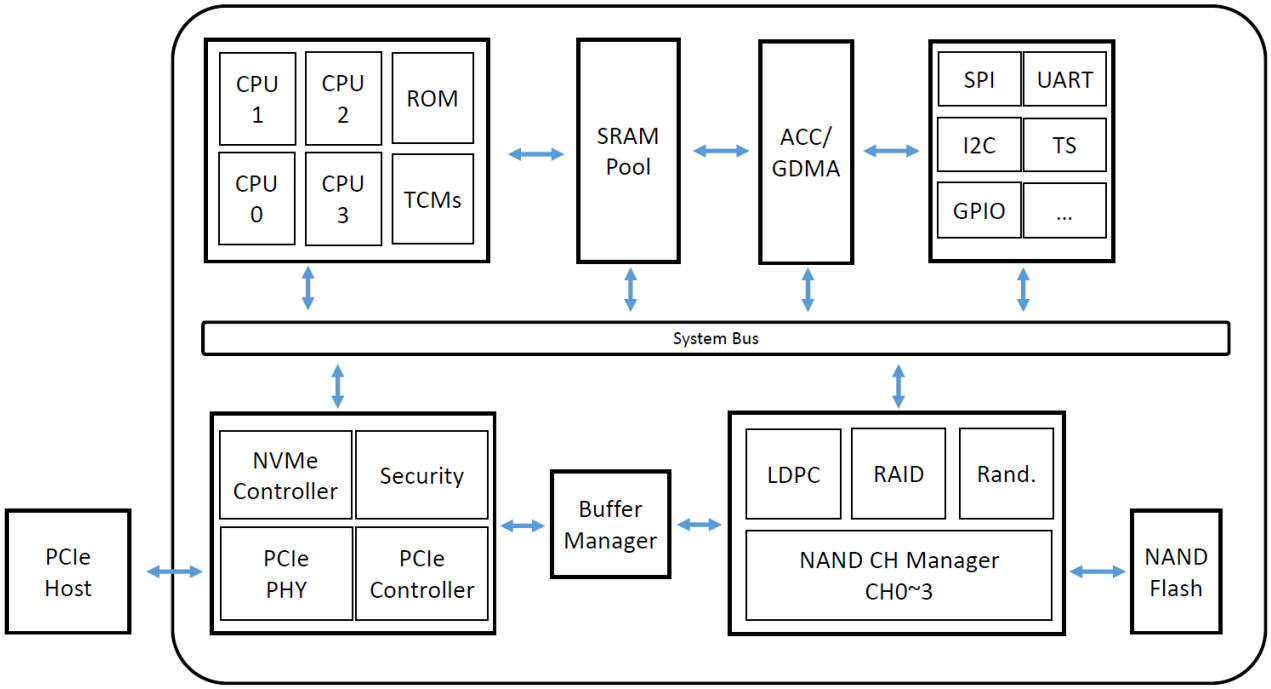

Figure 5: Innodisk M.2 (P80) 4IE2 Block Diagram

Innodisk M.2 (P80) 4IE2 integrates a PCIe Gen 4 x4 controller and NAND flash memories. Communication with the host occurs through the host interface, using the standard NVM Express protocol. Communication with the flash device(s) occurs through the flash interface.

3.2 PCIe Gen 4 x4 Controller

Innodisk M.2 (P80) 4IE2 is designed with a PCIe Gen 4 x4 controller which is compliant with NVMe 1.4, up to 64.0Gbps transfer speed. In addition, it is compliant with PCIe Gen 1, Gen 2 and Gen 3 specification. The controller supports up to four channels for flash interface.

3.3 Error Detection and Correction

Innodisk M.2 (P80) 4IE2 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 with a limited number of cycles. This number is called the ***erase cycle limit*** or ***write endurance limit*** and is defined by the flash NAND vendor. The erase cycle limit applies to each individual erase block in the flash device.

Innodisk M.2 (P80) 4IE2 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 device 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 lifetime of the SSD. When a Bad Block is detected, it will be flagged as unusable block by firmware. The SSD implement Bad Blocks management that consists of Bad Blocks replacement and Error Correcting 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 Garbage Collection/TRIM

Garbage collection and TRIM technology are 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.

3.7 End to End Data Path Protection

End-to-end Data Path Protection that secures the data transmission between host system and NAND Flash. In the transmission path, no matter in or out, all buffer and storage implement Error Code Correction that optimizes the data integrity in the whole transmission of SSD.

3.8 Thermal Management

M.2 (P80) 4IE2 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 iData Guard

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.

3.11 iCell Technology (Optional)

iCell circuit is designed with several capacitors to be able to provide power after host power off. The SSD controller can write all SRAM buffer data to flash, so that is why M.2 (P80) 4IE2 can ensure all data can be written to disk without any data loss.

3.12 TCG OPAL (Optional)

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 4IE2 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

3.12 PLA_S3# (Optional)

Power Loss Acknowledge. Open drain with pull-up on Platforms that support power loss notification. An Adapter that supports this function must drive the signal to reflect its current power loss processing complete state.

3.13 PLN (Optional)

Power Loss Notification. Open drain with a pull-up on Adapters that support power loss notification. When the Platform supports power loss notification, this signal is asserted to indicate a power loss event is expected to occur. When the Adapter supports this function and the signal is asserted then it must ready itself for power loss.

4. Installation Requirements

4.1 M.2 (P80) 4IE2 Pin Directions

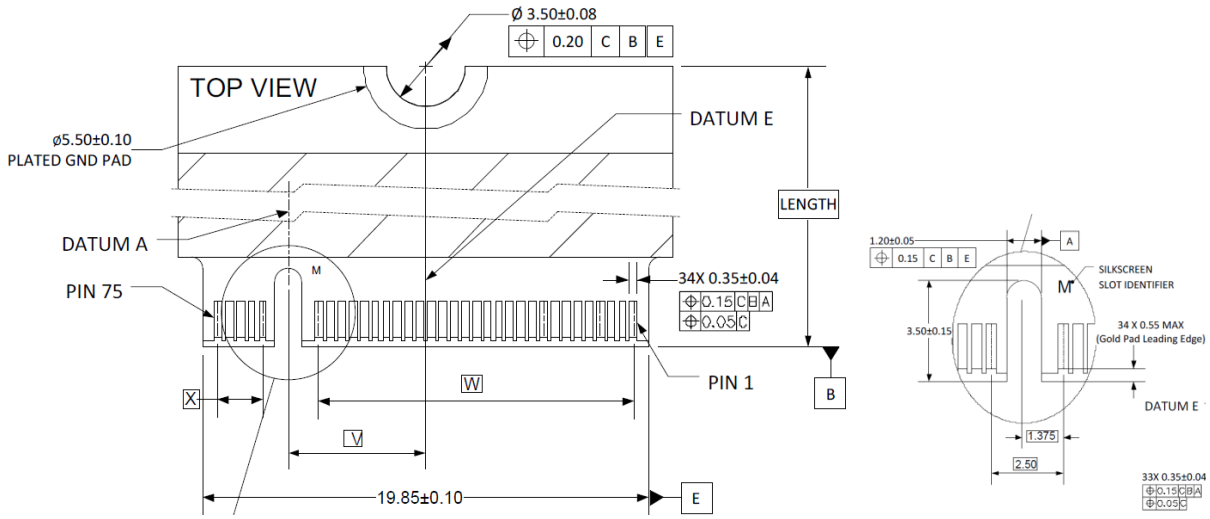

Figure 6: Signal Segment and Power Segment

4.2 Electrical Connections for M.2 (P80) 4IE2

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) 4IE2 is compliant with M.2 Socket 3 key M.

4.3 Device Drive

M.2 (P80) 4IE2 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 1.4

5.1 Get Log Page (Log Identifier 02h)

Innodisk 4IE2 series SMART / Health Information Log are listed in following table.

Table 11: Get Log Page – SMART / Health Information Log

<table>
<tr><th>Bytes</th><th>Description</th></tr>
<tr><td>0</td><td>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 to ‘1’. 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 state at the time the Get Log Page command is processed and may not reflect the state at the time a related asynchronous event notification, if any, occurs or occurred.
<table>
<tr><th>Bit</th><th>Definition</th></tr>
<tr><td>0</td><td>If set to ‘1’, then the available spare capacity has fallen below the threshold.</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.
b) less than or equal to an under temperature threshold.</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>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>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>5</td><td>If set to ‘1’, then the Persistent Memory Region has become read-only or unreliable.</td></tr>
<tr><td>7:6</td><td>Reserved</td></tr>
</table>
</td></tr>
</table>

1:2	**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. Warning and critical overheating composite temperature threshold values are reported by the WCTEMP and CCTEMP fields in the Identify Controller data structure.
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%). The values 101 to 255 are reserved.
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.
6	**Endurance Group Critical Warning Summary:** This field indicates critical warnings for the state of Endurance Groups. Each bit corresponds to a critical warning type, multiple bits may be set to '1'. If a bit is cleared to '0', then that critical warning does not apply to any Endurance Group. 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. If a bit is set to '1' in one or more Endurance Groups, then the corresponding bit shall be set to '1' in this field.

Bit	Definition
0	If set to '1', then the available spare capacity of one or more Endurance Groups has fallen below the threshold.
1	Reserved
2	If set to '1', then the reliability of one or more Endurance Groups has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.
3	If set to '1', then the namespaces in one or more Endurance Groups have been placed in read only mode not as a result of a change in the write protection state of a namespace.
7:4	Reserved

7:31	**Reserved**
32:47	**Data Units Read:** Contains the number of 512 byte data units the host has read from the controller as part of processing a SMART Data Units Read Command; this value does not include metadata. This value is reported in thousands (i.e., a value of 1 corresponds to 1,000 units of 512 bytes read) and is rounded up (e.g., one indicates that the number of 512 byte data units read is from 1 to 1,000, three indicates that the number of 512 byte data units read is from 2,001 to 3,000). Refer to the specific I/O Command Set specification for the list of SMART Data Units Read Commands that affect this field. A value of 0h in this field indicates that the number of SMART Data Units Read is not reported.
48:63	**Data Units Written:** Contains the number of 512 byte data units the host has written to the controller as part of processing a User Data Out Command; this value does not include metadata. This value is reported in thousands (i.e., a value of 1 corresponds to 1,000 units of 512 bytes written) and is rounded up (e.g., one indicates that the number of 512 byte data units written is from 1 to 1,000, three indicates that the number of 512 byte data units written is from 2,001 to 3,000). Refer to the specific I/O Command Set specification for the list of User Data Out Commands that affect this field. A value of 0h in this field indicates that the number of Data Units Written is not reported.
64:79	**Host Read Commands:** Contains the number of SMART Host Read Commands completed by the controller. Refer to the specific I/O Command Set specification for the list of SMART Host Read Commands that affect this field.
80:95	**Host Write Commands:** Contains the number of User Data Out Commands completed by the controller. Refer to the specific I/O Command Set specification for the list of User Data Out Commands that affect this field.
96:111	**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.
112:127	**Power Cycles:** Contains the number of power cycles.
128:143	**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.
144:159	**Unsafe Shutdowns:** Contains the number of unsafe shutdowns. This count is incremented when the controller does not report it is safe to power down prior to loss of main power.

160:175	**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. Errors introduced as a result of a Write Uncorrectable command (refer to the NVM Command Set Specification) may or may not be included in this field.
176:191	**Number of Error Information Log Entries:** Contains the number of Error Information log entries over the life of the controller.
192:195	**Warning Composite Temperature Time:** Contains the amount of time in minutes that the controller is operational and the Composite Temperature is greater than or equal to the Warning Composite Temperature Threshold (WCTEMP) field and less than the Critical Composite Temperature Threshold (CCTEMP) field in the Identify Controller data structure in Figure 275. 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.
196:199	**Critical 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 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.
200:201	**Temperature Sensor 1:** Contains the current temperature reported by the embedded thermal sensor in the controller.
202:203	**Temperature Sensor 2:** Contains the current temperature reported by the embedded thermal sensor in the NAND Flash (Channel #0 and CE #0).
204:205	**Temperature Sensor 3:** Contains the current temperature reported by the embedded thermal sensor in the NAND Flash (Channel #0 and CE #0).
206:207	**Temperature Sensor 4:** Contains the current temperature reported by the embedded thermal sensor in the NAND Flash (Last channel and CE #0).
208:209	**Temperature Sensor 5:** Contains the current temperature reported by temperature sensor 5.
210:211	**Temperature Sensor 6:** Contains the current temperature reported by temperature sensor 6.
212:213	**Temperature Sensor 7:** Contains the current temperature reported by temperature sensor 7.
214:215	**Temperature Sensor 8:** Contains the current temperature reported by temperature sensor 8.
216:219	**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.
220:223	**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.

224:227	**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.
228:231	**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.
232:337	**Reserved**
338:345	**Later Bad Count**
346:353	**Power-On hours Count**
354:361	**Drive Power Cycle Count**
362:369	**Total Bad Block Count**
370:377	**User Max Erase Count**
378:385	**User Avg Erase Count**
386:393	**Device Life**
394:401	**Spare Block Count**
402:409	**Program Fail Count**
410:417	**Erase Fail Count**
418:425	**Unexpected Power Loss Count**
426:433	**Temperature (Kelvin - K °K)**
434:441	**Flash ID**
442:449	**Later Bad Block Info (Read / Write / Erase)**
450:457	**Total LBAs Written (unit = 32MB)**
458:465	**Total LBAs Read (unit = 32MB)**

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	22
	D	**H**	**M**	**2**	**8**	**-**	**F**	**4**	**G**	**D**	**F**	**1**	**K**	**W**	**A**	**Q**	**F**	**P**	**(H)**	**-**	**X**	**X**

Definition	
Code 1st (Disk)	**Code 14th (Operation Temperature)**
D : Disk	C: Standard Grade (0℃~ +70℃)
Code 2nd (Feature set)	W: Industrial Grade (-40℃~ +85℃)
H : iSLC series	
Code 3rd ~5th (Form factor)	**Code 15th (Internal control)**
M28: M.2 Type 2280	A~Z: BGA PCB version.
Code 7th ~9th (Capacity)	**Code 16th (Channel of data transfer)**
40G: 40GB	D: Dual Channels
80G: 80GB	Q: Quad Channels
A60: 160GB	
D2G: 320GB	
F4G: 640GB	
Code 10th ~12th (Controller)	**Code 17th (Flash Type)**
DF1: PCIe 4IE2 series	F: Kioxia 3D TLC
DF2: PCIe 4IE2 series with AES & TCG OPAL functions	
Code 13th (Flash mode)	**Code 18th ~19th(Optional Function)**
K: 112 Layers 3D TLC	P: PLP (iCell) feature
	H: with heatsink accessory (WT default)
	PH: PLP (iCell) feature + heatsink accessory
	Code 21st~ (Customize code)