

RoHS Compliant PCI Express Flash Drive

PV910-M242 Product Specifications



May 26, 2021

Version 1.0



Apacer Technology Inc.

1F, No.32, Zhongcheng Rd., Tucheng Dist., New Taipei City, Taiwan, R.O.C

Tel: +886-2-2267-8000 Fax: +886-2-2267-2261

www.apacer.com

Specifications Overview:

- **PCIe Interface**
 - Compliant with NVMe 1.3
 - Compatible with PCIe Gen3 x2 interface
- **Capacity**
 - 60, 120, 240, 480 GB
- **Performance***
 - Interface burst read/write: 2 GB/sec
 - Sequential read: up to 1,775 MB/sec
 - Sequential write: up to 1,255 MB/sec
 - Random read (4K): up to 130,000 IOPS
 - Random write (4K): up to 131,000 IOPS
- **Flash Management**
 - Low-Density Parity-Check (LDPC) Code
 - Global Wear Leveling
 - Flash bad-block management
 - Flash Translation Layer: Page Mapping
 - Power Failure Management
 - S.M.A.R.T.
 - TRIM
 - Hyper Cache Technology
 - Over-Provisioning
 - SMART Read Refresh™
 - NVMe Secure Erase
- **NAND Flash Type: 3D TLC (BiCS3)**
- **MTBF: >3,000,000 hours**
- **Endurance (in drive writes per day: DWPD)**
 - 60 GB: 3.07 DWPD
 - 120 GB: 3 DWPD
 - 240 GB: 2.98 DWPD
 - 480 GB: 2.94 DWPD
- **Temperature Range**
 - Operating:
 - Standard: 0°C to 70°C
 - Wide: -40°C to 85°C
 - Storage: -40°C to 100°C
- **Supply Voltage**
 - 3.3 V ± 5%
- **Power Consumption***
 - Active mode: 690 mA
 - Idle mode: 140 mA
- **Connector Type**
 - 75-pin M.2 module pinout
- **Power Management**
 - Supports APST
 - Supports ASPM L1.2
- **NVMe Features****
 - Supports HMB (Host Memory Buffer)
- **Security**
 - AES 256-bit hardware encryption
 - Signed Firmware
- **Reliability**
 - Thermal Sensor
 - Thermal Throttling
 - End-to-End Data Protection
- **Form Factor**
 - M.2 2242-D5-B-M Key
 - Dimensions: 42.00 x 22.00 x 3.88, unit: mm
 - Net weight: 3.53g ± 5%
- **LED Indicators for Drive Behavior**
- **RoHS Compliant**

*Varies from capacities. The values for performances and power consumptions presented are typical and may vary depending on flash configurations or platform settings.

**Windows 10 (version 1703) onwards supports the HMB (Host Memory Buffer) function.

Table of Contents

1. General Descriptions	4
2. Functional Block	4
3. Pin Assignments.....	5
4. Product Specifications.....	7
4.1 Capacity.....	7
4.2 Performance	7
4.3 Environmental Specifications	7
4.4 Mean Time Between Failures (MTBF)	8
4.5 Certification and Compliance.....	8
4.6 Endurance	8
4.7 LED Indicator Behavior.....	9
5. Flash Management	10
5.1 Error Correction/Detection.....	10
5.2 Bad Block Management	10
5.3 Global Wear Leveling	10
5.4 Flash Translation Layer – Page Mapping.....	10
5.5 Power Failure Management	10
5.6 TRIM.....	11
5.7 Hyper Cache Technology	11
5.8 Over-Provisioning	11
5.9 SMART Read Refresh™	11
5.10 NVMe Secure Erase.....	11
6. NVMe Support Features	12
6.1 Host Memory Buffer.....	12
7. Security & Reliability Features	13
7.1 Thermal Sensor	13
7.2 Thermal Throttling	13
7.3 Advanced Encryption Standard.....	13
7.4 End-to-End Data Protection.....	13
7.5 Signed Firmware	13

8. Software Interface	14
8.1 Command Set.....	14
8.2 S.M.A.R.T.	15
9. Electrical Specifications.....	17
9.1 Operating Voltage.....	17
9.2 Power Consumption	17
10. Physical Characteristics.....	18
10.1 Dimensions	18
10.2 Net Weight.....	18
11. Product Ordering Information.....	19
11.1 Product Code Designations.....	19
11.2 Valid Combinations.....	20

1. General Descriptions

Apacer PV910-M242 (M.2 2242) is the next generation modularized Solid State Drive (SSD) with the shape of all new M.2 form factor, aimed to be the more suitable for mobile and compact computers with standard width at only 22.00 mm. PV910-M242 appears in M.2 2242 mechanical dimensions and is believed to be the leading add-in storage solution for future host computing systems.

The M.2 SSD is designed with PCIe-based connector pinouts, providing full compliance with the latest PCIe Gen3 x2 interface specifications. Aside from PCIe compliance, PV910-M242 delivers exceptional performance and power efficiency. On the other hand, the extreme thin and light form factor makes PV910-M242 the ideal choice for mobile computing systems, which appears to be the trend in near future.

Regarding reliability, PV910-M242 is built with a powerful PCIe controller that supports on-the-module ECC as well as efficient wear leveling scheme. In terms of power efficiency, PV910-M242 is compliant with PCIe Gen3 x2 interface standard so that it can operate on power management modes, which greatly save on power consumption.

2. Functional Block




Figure 2-1 Functional Block Diagram

3. Pin Assignments

This connector does not support hot plug capability. There are a total of 75 pins. 12 pin locations are used for mechanical key locations; this allows such a module to plug into both Key B and Key M connectors.




Table 3-1 Pin Assignments

Pin No.	Type	Description
1	GND	CONFIG_3 = GND
2	3.3V	3.3V source
3	GND	Ground
4	3.3V	3.3V source
5	N/C	No connect*
6	N/C	No connect*
7	N/C	No connect*
8	N/C	No connect*
9	GND	Ground
10	LED1#	Open drain, active low signal. These signals are used to allow the add-in card to provide status indicators via LED devices that will be provided by the system.
11	N/C	No connect*
12	Module Key B	Module Key
13	Module Key B	
14	Module Key B	
15	Module Key B	
16	Module Key B	
17	Module Key B	
18	Module Key B	
19	Module Key B	
20	N/C	No connect*
21	GND	Ground
22	N/C	No connect*
23	N/C	No connect*
24	N/C	No connect*
25	N/C	No connect*
26	N/C	No connect*
27	GND	Ground
28	N/C	No connect*
29	PETn1	PCIe TX Differential signal defined by the PCI Express M.2 spec

Table 3-1 Pin Assignments

Pin No.	Type	Description
30	N/C (WP)	Default: No connect*
31	PETp1	PCIe TX Differential signal defined by the PCI Express M.2 spec
32	N/C	No connect*
33	GND	Ground
34	N/C	No connect*
35	PERn1	PCIe RX Differential signal defined by the PCI Express M.2 spec
36	N/C	No connect*
37	PERp1	PCIe RX Differential signal defined by the PCI Express M.2 spec
38	N/C	No connect*
39	GND	Ground
40	N/C	No connect*
41	PETn0	PCIe TX Differential signal defined by the PCI Express M.2 spec
42	N/C	No connect*
43	PETp0	PCIe TX Differential signal defined by the PCI Express M.2 spec
44	N/C	No connect*
45	GND	Ground
46	N/C	No connect*
47	PERn0	PCIe RX Differential signal defined by the PCI Express M.2 spec
48	N/C	No connect*
49	PERp0	PCIe RX Differential signal defined by the PCI Express M.2 spec
50	PERST#(I/O)(0/3.3V)	PE-Reset is a functional reset to the card as defined by the PCIe Mini CEM specification.
51	GND	Ground
52	CLKREQ#(I/O)(0/3.3V)	Clock Request is a reference clock request signal as defined by the PCIe Mini CEM specification; Also used by L1 PM Sub-states.
53	REFCLKn	PCIe Reference Clock signals (100 MHz) defined by the PCI Express M.2 spec.
54	N/C	No connect*
55	REFCLKp	PCIe Reference Clock signals (100 MHz) defined by the PCI Express M.2 spec.
56	N/C	No connect*
57	GND	Ground
58	N/C	No connect*
59	Module Key M	Module Key
60	Module Key M	
61	Module Key M	
62	Module Key M	
63	Module Key M	
64	Module Key M	
65	Module Key M	
66	Module Key M	
67	N/C (Erase)	Default: No connect*
68	N/C	No connect*
69	NC	CONFIG_1 = No connect*
70	3.3V	3.3V source
71	GND	Ground
72	3.3V	3.3V source
73	GND	Ground
74	3.3V	3.3V source
75	GND	CONFIG_2 = Ground

*Reserved by Apacer, please do not connect on a host.

4. Product Specifications

4.1 Capacity

Capacity specifications of PV910-M242 are available as shown in Table 4-1. It lists the specific capacity and the default numbers of heads, sectors and cylinders for each product line.

Table 4-1 Capacity Specifications

Capacity	Total bytes*	Total LBA
60 GB	60,022,480,896	117,231,408
120 GB	120,034,123,776	234,441,648
240 GB	240,057,409,536	468,862,128
480 GB	480,103,981,056	937,703,088

*Display of total bytes varies from file systems, which means not all of the bytes can be used for storage.

**Notes: 1 GB = 1,000,000,000 bytes; 1 sector = 512 bytes.

LBA count addressed in the table above indicates total user storage capacity and will remain the same throughout the lifespan of the device. However, the total usable capacity of the SSD is most likely to be less than the total physical capacity because a small portion of the capacity is reserved for device maintenance usages.

4.2 Performance

Performance of PV910-M242 is listed below in Table 4-2.

Table 4-2 Performance Specifications

Capacity	60 GB	120 GB	240 GB	480 GB
Performance				
Sequential Read* (MB/s)	775	1,550	1,775	1,775
Sequential Write* (MB/s)	255	500	955	1,255
Random Read IOPS** (4K)	41,000	79,000	120,000	130,000
Random Write IOPS** (4K)	50,000	99,000	131,000	130,000

Note:

Measured with OS version: Win10 (64bit), version 1803 with HMB (Host Memory Buffer), performance may differ from various flash configurations or host system settings.

*Sequential performance is based on CrystalDiskMark 5.2.1 with file size 1,000MB.

**Random performance measured using IOMeter with Queue Depth 32.

4.3 Environmental Specifications

Environmental specifications of PV910-M242 are shown in Table 4-3.

Table 4-3 Environmental Specifications

Item	Specifications
Operating temp.	0°C to 70°C (Standard); -40°C to 85°C (Wide)
Non-operating temp.	-40°C to 100°C
Operating vibration	7.69 GRMS, 20~2000 Hz/random (compliant with MIL-STD-810G)
Non-operating vibration	4.02 GRMS, 15~2000 Hz/random (compliant with MIL-STD-810G)
Operating shock	50(G), 11(ms), half-sine wave
Non-operating shock	1,500(G), 0.5(ms), half-sine wave

Note: Shock and Vibration specifications are subject to change without notice.

4.4 Mean Time Between Failures (MTBF)

Mean Time Between Failures (MTBF) is predicted based on reliability data for the individual components in PV910-M242. The prediction result for PV910-M242 is more than 1,000,000 hours.

Note: The MTBF is predicated and calculated based on “Telcordia Technologies Special Report, SR-332, Issue 2” method.

4.5 Certification and Compliance

PV910-M242 complies with the following standards:

- FCC
- CE
- RoHS
- MIL-STD-810G

4.6 Endurance

The endurance of a storage device is predicted by Drive Writes Per Day based on several factors related to usage, such as the amount of data written into the drive, block management conditions, and daily workload for the drive. Thus, key factors, such as Write Amplifications and the number of P/E cycles, can influence the lifespan of the drive.

Table 4-4 Drive Writes Per Day

Capacity	Drive Writes Per Day
60 GB	3.07
120 GB	3
240 GB	2.98
480 GB	2.94

Note:

- This estimation complies with JEDEC JESD-219, enterprise endurance workload of random data with payload size distribution.
- Flash vendor guaranteed 3D NAND TLC P/E cycle: 3K
- WAF may vary from capacity, flash configurations and writing behavior on each platform.
- 1 Terabyte = 1,024GB
- DWPD (Drive Writes Per Day) is calculated based on the number of times that user overwrites the entire capacity of an SSD per day of its lifetime during the warranty period. (3D NAND TLC warranty: 2 years)

4.7 LED Indicator Behavior

The behavior of the PV910-M242 LED indicators is described in Table 4-5.

Table 4-5 LED Behavior

Location	LED	Description
LED A	DAS	LED blinks when the drive is being accessed



5. Flash Management

5.1 Error Correction/Detection

PV910-M242 implements a hardware ECC scheme, based on the Low Density Parity Check (LDPC). LDPC is a class of linear block error correcting code which has apparent coding gain over BCH code because LDPC code includes both hard decoding and soft decoding algorithms. With the error rate decreasing, LDPC can extend SSD endurance and increase data reliability while reading raw data inside a flash chip.

5.2 Bad Block Management

Bad blocks are blocks that include one or more invalid bits, and their reliability is not guaranteed. Blocks that are identified and marked as bad by the manufacturer are referred to as “Initial Bad Blocks”. Bad blocks that are developed during the lifespan of the flash are named “Later Bad Blocks”. Apacer implements an efficient bad block management algorithm to detect the factory-produced bad blocks and manages any bad blocks that appear with use. This practice further prevents data being stored into bad blocks and improves the data reliability.

5.3 Global Wear Leveling

Flash memory devices differ from Hard Disk Drives (HDDs) in terms of how blocks are utilized. For HDDs, when a change is made to stored data, like erase or update, the controller mechanism on HDDs will perform overwrites on blocks. Unlike HDDs, flash blocks cannot be overwritten and each P/E cycle wears down the lifespan of blocks gradually. Repeatedly program/erase cycles performed on the same memory cells will eventually cause some blocks to age faster than others. This would bring flash storages to their end of service term sooner. Global wear leveling is an important mechanism that levels out the wearing of all blocks so that the wearing-down of all blocks can be almost evenly distributed. This will increase the lifespan of SSDs.

5.4 Flash Translation Layer – Page Mapping

Page mapping is an advanced flash management technology whose essence lies in the ability to gather data, distribute the data into flash pages automatically, and then schedule the data to be evenly written. Page-level mapping uses one page as the unit of mapping. The most important characteristic is that each logical page can be mapped to any physical page on the flash memory device. This mapping algorithm allows different sizes of data to be written to a block as if the data is written to a data pool and it does not need to take extra operations to process a write command. Thus, page mapping is adopted to increase random access speed and improve SSD lifespan, reduce block erase frequency, and achieve optimized performance and lifespan.

5.5 Power Failure Management

Power Failure Management plays a crucial role when power supply becomes unstable. Power disruption may occur when users are storing data into the SSD, leading to instability in the drive. However, with Power Failure Management, a firmware protection mechanism will be activated to scan pages and blocks once power is resumed. Valid data will be transferred to new blocks for merging and the mapping table will be rebuilt. Therefore, data reliability can be reinforced, preventing damage to data stored in the NAND Flash.

5.6 TRIM

TRIM is a feature which helps improve the read/write performance and speed of solid-state drives (SSD). Unlike hard disk drives (HDD), SSDs are not able to overwrite existing data, so the available space gradually becomes smaller with each use. With the TRIM command, the operating system can inform the SSD which blocks of data are no longer in use and can be removed permanently. Thus, the SSD will perform the erase action, which prevents unused data from occupying blocks all the time.

5.7 Hyper Cache Technology

Apacer proprietary Hyper Cache technology uses a portion of the available capacity as SLC (1bit-per-cell) NAND flash memory, called Hyper cache mode. When data is written to SSD, the firmware will direct the data to Hyper Cache mode, providing excellent performance to handle various scenarios in industrial use.

5.8 Over-Provisioning

Over-Provisioning (OP) is a certain portion of the SSD capacity exclusively for increasing Garbage Collection (GC) efficiency, especially when the SSD is filled to full capacity or performs a heavy mixed-random workload. OP has the advantages of providing extended life expectancy, reliable data integrity, and high sustained write performance.

5.9 SMART Read Refresh™

Apacer's SMART Read Refresh plays a proactive role in avoiding read disturb errors from occurring to ensure health status of all blocks of NAND flash. Developed for read-intensive applications in particular, SMART Read Refresh is employed to make sure that during read operations, when the read operation threshold is reached, the data is refreshed by re-writing it to a different block for subsequent use.

5.10 NVMe Secure Erase

NVMe Secure Erase is an NVMe drive sanitize command currently embedded in most of the storage drives. Defined in NVMe specifications, NVMe Secure Erase is part of Format NVM command that allows storage drives to erase all user data areas. The erase process usually runs on the firmware level as most of the NVMe-based storage media currently in the market are built-in with this command. NVMe Secure Erase can securely wipe out the user data in the drive and protects it from malicious attack.

6. NVMe Support Features

6.1 Host Memory Buffer

Host Memory Buffer (HMB) allows HOST to allocate system memory for SSD's exclusive use in order to provide better performance and endurance, especially for DRAMless solutions.

7. Security & Reliability Features

7.1 Thermal Sensor

Apacer Thermal Sensor is a digital temperature sensor with serial interface. By using designated pins for transmission, storage device owners are able to read temperature data.

7.2 Thermal Throttling

Thermal throttling can monitor the temperature of the SSD equipped with a built-in thermal sensor. This method can ensure the temperature of the device stays within temperature limits by drive throttling, i.e. reducing the speed of the drive when the device temperature reaches the threshold level, so as to prevent overheating, guarantee data reliability, and prolong product lifespan. When the temperature exceeds the maximum threshold level, thermal throttling will be triggered to reduce performance step by step to prevent hardware components from being damaged. Performance is only permitted to drop to the extent necessary for recovering a stable temperature to cool down the device's temperature. Once the temperature decreases to the minimum threshold value, transfer speeds will rise back to its optimum performance level.

7.3 Advanced Encryption Standard

Advanced Encryption Standard (AES) is a specification for the encryption of electronic data. AES has been adopted by the U.S. government since 2001 to protect classified information and is now widely implemented in embedded computing applications. The AES algorithm used in software and hardware is symmetric so that encrypting/decrypting requires the same encryption key. Without the key, the encrypted data is inaccessible to ensure information security.

Notably in flash memory applications, AES 256-bit hardware encryption is the mainstream to protect sensitive or confidential data. The hardware encryption provides better performance, reliability, and security than software encryption. It uses a dedicated processor, which is built inside the controller, to process the encryption and decryption. This enormously shortens the processing time and makes it efficient.

7.4 End-to-End Data Protection

End-to-End Data Protection is a feature implemented in Apacer SSD products that extends error control to cover the entire path from the host computer to the drive and back, and that ensures data integrity at multiple points in the path to enable reliable delivery of data transfers. Unlike ECC which does not exhibit the ability to determine the occurrence of errors throughout the process of data transmission, End-to-End Data Protection allows SSD controller to identify an error created anywhere in the path and report the error to the host computer before it is written to the drive. This error-checking and error-reporting mechanism therefore guarantees the trustworthiness and reliability of the SSD.

7.5 Signed Firmware

Apacer's Signed Firmware technology is a secure way to update firmware. By including a digital signature, a firmware update will be authenticated by the Apacer SSD before a firmware update is performed. This extra layer of protection keeps drives secure.

8. Software Interface

8.1 Command Set

Table 8-1 summarizes the commands supported by PV910-M242.

Table 8-1 Admin Commands

Opcode	Command Description
00h	Delete I/O Submission Queue
01h	Create I/O Submission Queue
02h	Get Log Page
04h	Delete I/O Completion Queue
05h	Create I/O Completion Queue
06h	Identify
08h	Abort
09h	Set Features
0Ah	Get Features
0Ch	Asynchronous Event Request
10h	Firmware Activate
11h	Firmware Image Download
14h	Device Self-test

Table 8-2 Admin Commands – NVM Command Set Specific

Opcode	Command Description
80h	Format NVM

Table 8-3 NVM Commands

Opcode	Command Description
00h	Flush
01h	Write
02h	Read
09h	Dataset Management

8.2 S.M.A.R.T.

SMART, an acronym for Self-Monitoring, Analysis and Reporting Technology, is an open standard that allows a hard disk drive to automatically detect its health and report potential failures. When a failure is recorded by SMART, users can choose to replace the drive to prevent unexpected outage or data loss. Moreover, SMART can inform users of impending failures while there is still time to perform proactive actions, such as copy data to another device.

Table 8-4 SMART (02h)

Byte	Length	Description
0	1	Critical Warning
1-2	2	Composite Temperature (PCB Sensor)
3	1	Available Spare
4	1	Available Spare Threshold
5	1	Percentage Used (Average Erase Count / P/E Cycle Count)
6-31	26	Reserved
32-47	16	Data Units Read
48-63	16	Data Units Written
64-79	16	Host Read Commands
80-95	16	Host Write Commands
96-111	16	Controller Busy Time
112-127	16	Power Cycles
128-143	16	Power On Hours
144-159	16	Unsafe Shutdowns
160-175	16	Media and Data Integrity Errors
176-191	16	Number of Error Information Log Entries
192-195	4	Warning Composite Temperature Time
196-199	4	Critical Composite Temperature Time
200-201	2	Temperature Sensor 1: Controller Temperature
202-203	2	Temperature Sensor 2: PCB Temperature
204-205	2	Temperature Sensor 3: NAND Flash Temperature
206-207	2	Temperature Sensor 4
208-209	2	Temperature Sensor 5
210-211	2	Temperature Sensor 6
212-213	2	Temperature Sensor 7
214-215	2	Temperature Sensor 8
216-511	296	Reserved

Table 8-5 SMART (C0h)

Byte	Length	Description
0-255	256	Reserved
256-257	2	SSD Protect Mode
258-261	4	Host Read UNC Count
262-265	4	PHY Error Count
266-269	4	CRC Error Count
270-273	4	Total Early Bad Block Count
274-277	4	Total Later Bad Block Count
278-281	4	Max Erase Count
282-285	4	Average Erase Count
286-289	4	Program Fail Count
290-293	4	Erase Fail Count
294-301	8	Flash Write Sector
302-305	4	Total Spare Block
306-309	4	Current Spare Block
310-313	4	Read Retry Count
314-511	210	Reserved

9. Electrical Specifications

9.1 Operating Voltage

Table 9-1 lists the supply voltage for PV910-M242.

Table 9-1 Operating Range

Item	Range
Supply Voltage	3.3V \pm 5%

9.2 Power Consumption

Table 9-2 lists the power consumption for PV910-M242.

Table 9-2 Power Consumption

Mode \ Capacity	60 GB	120 GB	240 GB	480 GB
Active (mA)	340	545	600	690
Idle (mA)	140	140	140	140

Note:

*All values are typical and may vary depending on flash configurations or host system settings.

**Active power is an average power measurement performed using CrystalDiskMark with 128KB sequential read/write transfers.

10. Physical Characteristics

10.1 Dimensions




Figure 10-1 Dimensions

10.2 Net Weight

Table 10-1 Net Weight

Capacity	Net Weight (g \pm 5%)
60GB	3.01
120GB	3.39
240GB	3.39
480GB	3.53

11. Product Ordering Information

11.1 Product Code Designations

Code	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
	B	7	2	.	9	1	5	X	X	B	.	0	0	2	0	4

Code 1-3 (Product Line & form factor)	PCIe M.2 2242
Code 5-6 (Model/Solution)	PV910-M242
Code 7-8 (Product Capacity)	5G: 60GB 5H: 120GB 5J: 240GB 5K: 480GB
Code 9 (Flash Type & Product Temp)	G: 3D TLC Standard Temperature H: 3D TLC Wide Temperature
Code 10 (Product Spec)	B: Double side B+M key
Code 12-14 (Version Number)	Random numbers generated by system
Code 15-16 (Firmware Version)	04: Thermal Sensor + OP

11.2 Valid Combinations

Capacity	Standard Temperature	Wide Temperature
60GB	B72.915GGB.00204	B72.915GHB.00204
120GB	B72.915HGB.00204	B72.915HHB.00204
240GB	B72.915JGB.00204	B72.915JHB.00204
480GB	B72.915KGB.00204	B72.915KHB.00204

Note: Valid combinations are those products in mass production or will be in mass production. Consult your Apacer sales representative to confirm availability of valid combinations and to determine availability of new combinations.

Revision History

Revision	Description	Date
0.1	Preliminary release	5/24/2021
1.0	Official release	5/26/2021