Apacer

RoHS Compliant

PCI Express Flash Drive

PV140-25 Product Specifications

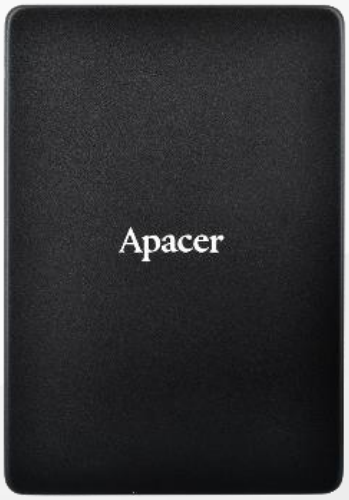

September 28, 2020

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

Apacer

Specifications Overview:

- **PCIe Interface**
  - Compliant with PCIe Express 3.1
  - Compliant with NVMe 1.3
  - Compatible with PCIe Gen3 x4 interface
- **Capacity**
  - 960, 1,920, 3,840 GB
- **Performance***
  - Interface burst read/write: 4 GB/sec
  - Sequential read: up to 3,340 MB/sec
  - Sequential write: up to 1,175 MB/sec
  - Random read (4K): up to 574,000 IOPS
  - Random write (4K): up to 266,000 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
  - Over-Provisioning
  - DataRAID™
  - NVMe Secure Erase
- **NAND Flash Type:** 3D TLC (BiCS3)
- **MTBF:** >3,000,000 hours
- **Endurance (in drive writes per day: DWPD)**
  - 960 GB: 2.2 DWPD
  - 1,920 GB: 2.3 DWPD
  - 3,840 GB: 3.3 DWPD
- **Temperature Range**
  - Operating: -40°C to 85°C
  - Storage: -40°C to 85°C
- **Supply Voltage**
  - 12 V ± 5%
- **Power Consumption***
  - Active mode: 705 mA
  - Idle mode: 180 mA
- **Connector Type**
  - U.2 (SFF-8639)
- **Form Factor**
  - 2.5 inch
  - Dimensions: 100.00 x 69.85 x 7.00, unit: mm
  - Net Weight: 71.5 ± 5%
- **Security**
  - AES 256-bit hardware encryption
- **Reliability**
  - Thermal Sensor
  - Thermal Throttling
  - End-to-End Data Protection
- **Power Management****
- **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.
**APST and ASPM power managements are not supported.

Table of Contents

1. General Descriptions

Apacer PV140-25 is the fastest Solid State Drive (SSD) designed with U.2 mechanical dimensions, providing full compliance with PCIe Gen3 x4 interface and NVMe 1.3 specifications. Built with a powerful PCIe controller, PV140-25 delivers outstanding performance in data transfer, reaching up to 574,000/266,000 and 3,340/1,175 MB/s in IOPS and sequential read/write. PV140-25 utilizes 3D NAND for higher capacity up to 3,840GB, making it the ideal choice for rugged, embedded applications.

Regarding reliability, PV140-25 is implemented with LDPC (Low Density Parity Check) ECC engine to extend SSD endurance and increase data reliability. For highly-intensive applications, End-to-End Data Protection ensures that data integrity can be assured at multiple points in the path to enable reliable delivery of data transfers, while Advanced Encryption Standard (AES) ensures data security and provides users with a peace of mind knowing their data is safeguarded at all times.

2. Functional Block

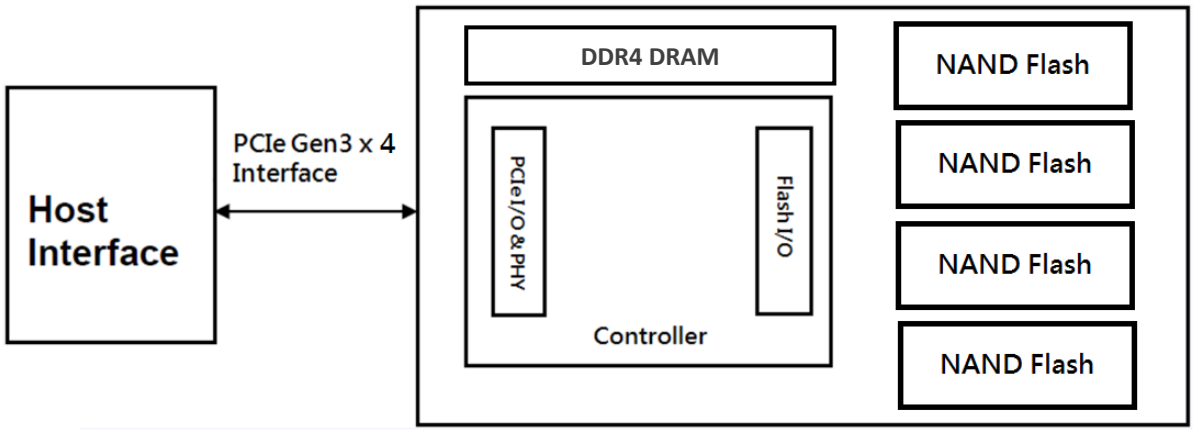

Figure 2-1 Functional Block Diagram

3. Pin Assignments

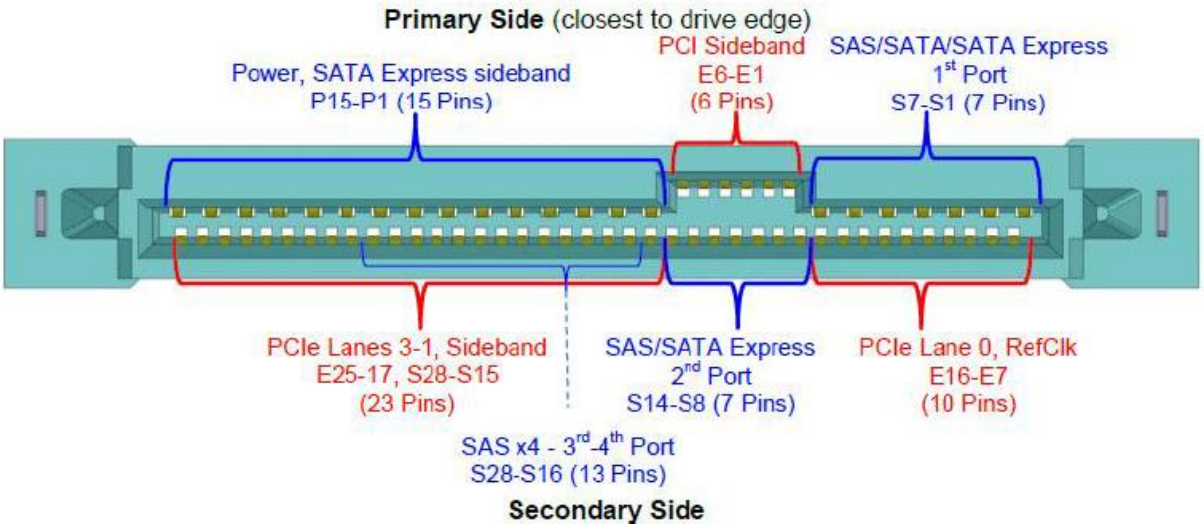

Table 3-1 Pin Assignments

Pin Number	Name	Type	Description
P1	WAKE#	Input	Signal for Link reactivation
P2	-	-	Outside scope of this specification
P3	CLKREQ#	Bi-Dir	Clock request
P4	IfDet#	Input	Interface Type Detect
P5	Ground	Ground	Ground
P6	Ground	Ground	Ground
P7	-	-	Outside scope of this specification
P8	-	-	Outside scope of this specification
P9	-	-	Outside scope of this specification
P10	PRSNT#	Input	Presence detect
P11	Activity	Input	-
P12	Ground	Ground	Ground
P13	+12V Precharge	Power	+12V Precharge power for SFF-8639 module
P14	+12V	Power	+12V power for SFF-8639 module
P15	+12V	Power	+12V power for SFF-8639 module
S1	Ground	Ground	Ground
S2	-	-	Outside scope of this specification
S3	-	-	Outside scope of this specification
S4	Ground	Ground	Ground
S5	-	-	Outside scope of this specification
S6	-	-	Outside scope of this specification
S7	Ground	Ground	Ground
S8	Ground	Ground	Ground

Pin Number	Name	Type	Description
S9	-	-	Outside scope of this specification
S10	-	-	Outside scope of this specification
S11	Ground	Ground	Ground
S12	-	-	Outside scope of this specification
S13	-	-	Outside scope of this specification
S14	Ground	Ground	Ground
S15	Reserved	-	Reserved
S16	Ground	Ground	Ground
S17	PETp1	Diff-Pair	Transmitter differential pair, Lane 1
S18	PETn1	Diff-Pair	Transmitter differential pair, Lane 1
S19	Ground	Ground	Ground
S20	PERn1	Diff-Pair	Receiver differential pair, Lane 1
S21	PERp1	Diff-Pair	Receiver differential pair, Lane 1
S22	Ground	Ground	Ground
S23	PETp2	Diff-Pair	Transmitter differential pair, Lane 2
S24	PETn2	Diff-Pair	Transmitter differential pair, Lane 2
S25	Ground	Ground	Ground
S26	PERn2	Diff-Pair	Receiver differential pair, Lane 2
S27	PERp2	Diff-Pair	Receiver differential pair, Lane 2
S28	Ground	Ground	Ground
E1	REFCLKB+	Diff-Pair	Reference clock (differential pair) for second X2 port
E2	REFCLKB-	Diff-Pair	Reference clock (differential pair) for second X2 port
E3	+3.3 Vaux	Power	3.3 V auxiliary power
E4	PERSTB#	Output	Fundamental reset for second X2 port
E5	PERST#	Output	Fundamental reset (if dual-port enabled, first X2 port)
E6	Reserved	-	Reserved
E7	REFCLK+	Diff-Pair	Reference clock (if dual-port enabled, first X2 port)
E8	REFCLK-	Diff-Pair	Reference clock (if dual-port enabled, first X2 port)
E9	Ground	Ground	Ground
E10	PETp0	Diff-Pair	Transmitter differential pair, Lane 0
E11	PETn0	Diff-Pair	Transmitter differential pair, Lane 0
E12	Ground	Ground	Ground
E13	PERn0	Diff-Pair	Receiver differential pair, Lane 0
E14	PERp0	Diff-Pair	Receiver differential pair, Lane 0
E15	Ground	Ground	Ground
E16	Reserved	-	Reserved
E17	PETp3	Diff-Pair	Transmitter differential pair, Lane 3
E18	PETn3	Diff-Pair	Transmitter differential pair, Lane 3
E19	Ground	Ground	Ground
E20	PERn3	Diff-Pair	Receiver differential pair, Lane 3
E21	PERp3	Diff-Pair	Receiver differential pair, Lane 3

Pin Number	Name	Type	Description
E22	Ground	Ground	Ground
E23	SMCLK	Bi-Dir	SMBus (System Management Bus) clock
E24	SMDAT	Bi-Dir	SMBus (System Management Bus) data
E25	-	-	-

4. Product Specifications

4.1 Capacity

Capacity specifications of PV140-25 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
960 GB	960,197,124,096	1,875,385,008
1,920 GB	1,920,383,410,176	3,750,748,848
3,840 GB	3,840,755,982,336	7,501,476,528

*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 PV140-25 is listed below in Table 4-2.

Table 4-2 Performance Specifications

Capacity / Performance	960 GB	1,920 GB	3,840 GB
Sequential Read* (MB/s)	3,340	3,225	2,920
Sequential Write* (MB/s)	1,165	1,175	1,120
Random Read IOPS (4K)**	396,000	574,000	455,000
Random Write IOPS (4K)**	266,000	264,000	250,000

Note:
Results may differ from various flash configurations or host system setting.
*Sequential performance is based on CrystalDiskMark 5.2.1 with file size 1,000MB.
**Random performance measured using IOMeter with Queue Depth 128.

4.3 Environmental Specifications

Environmental specifications of PV140-25 are shown in Table 4-3.

Table 4-3 Environmental Specifications

Item	Specifications
Operating temp.	-40°C to 85°C
Non-operating temp.	-40°C to 85°C
Non-operating vibration	20G, 20~2000 Hz/random
Non-operating shock	1,500(G), 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 PV140-25. The prediction result for PV140-25 is more than 3,000,000 hours.

Note: The MTBF is predicated and calculated based on “Telcordia Technologies Special Report, SR-332, Issue 3” method.

4.5 Certification and Compliance

PV140-25 complies with the following standards:

- CE
- FCC
- RoHS

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 Endurance Specifications

Capacity	Drive Writes Per Day
960 GB	2.2
1920 GB	2.3
3840 GB	3.3

Note:

- This estimation complies with JEDEC random enterprise workload.
- 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)

5. Flash Management

5.1 Error Correction/Detection

PV140-25 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

Current production technology is unable to guarantee total reliability of NAND flash memory array. When a flash memory device leaves factory, it comes with a minimal number of initial bad blocks during production or out-of-factory as there is no currently known technology that produce flash chips free of bad blocks. In addition, bad blocks may develop during program/erase cycles. Since bad blocks are inevitable, the solution is to keep them in control. Apacer flash devices are programmed with ECC, page mapping technique and S.M.A.R.T to reduce invalidity or error. Once bad blocks are detected, data in those blocks will be transferred to free blocks and error will be corrected by designated algorithms.

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 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.

5.8 DataRAID™

Apacer's DataRAID algorithm applies an additional level of protection and error-checking. Using this algorithm, a certain amount of space is given over to aggregating and resaving the existing parity data used for error checking. So, in the event that data becomes corrupted, the parity data can be compared to the existing uncorrupted data and the content of the corrupted data can be rebuilt.

5.9 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. Reliability Features

6.1 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.

6.2 Thermal Throttling

Thermal throttling can monitor the temperature of the SSD equipped with a built-in thermal sensor via S.M.A.R.T. commands. 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.

6.3 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.

6.4 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. Software Interface

7.1 Command Set

Table 7-1 summarizes the commands supported by PV140-25.

Table 7-1 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 7-2 Admin Commands – NVM Command Set Specific

Opcode	Command Description
80h	Format NVM
81h	Security Send
82h	Security Receive

Table 7-3 NVM Commands

Opcode	Command Description
00h	Flush
01h	Write
02h	Read
05h	Compare
08h	Write Zeroes
09h	Dataset Management

7.2 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 7-4 SMART (02h)

Byte	Length	Description
0	1	Critical Warning
1-2	2	Composite Temperature
3	1	Available Spare
4	1	Available Spare Threshold
5	1	Percentage Used
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
202-203	2	Temperature Sensor 2
204-205	2	Temperature Sensor 3
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 7-5 SMART (C0h)

Byte	Length	Description
0-255	256	Reserved
256-257	2	SSD Protect Mode
258-261	4	ECC Fail Count
262-273	12	Reserved
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
302-511	210	Reserved

8. Electrical Specifications

8.1 Operating Voltage

Table 8-1 lists the supply voltage for PV140-25.

Table 8-1 Operating Range

Item	Range
Supply Voltage	12V ± 5%

8.2 Power Consumption

Table 8-2 lists the power consumption for PV140-25.

Table 8-2 Power Consumption

Capacity / Mode	960 GB	1,920 GB	3,840 GB
Active (mA)	535	680	705
Idle (mA)	175	175	180

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.

9. Physical Characteristics

9.1 Dimensions

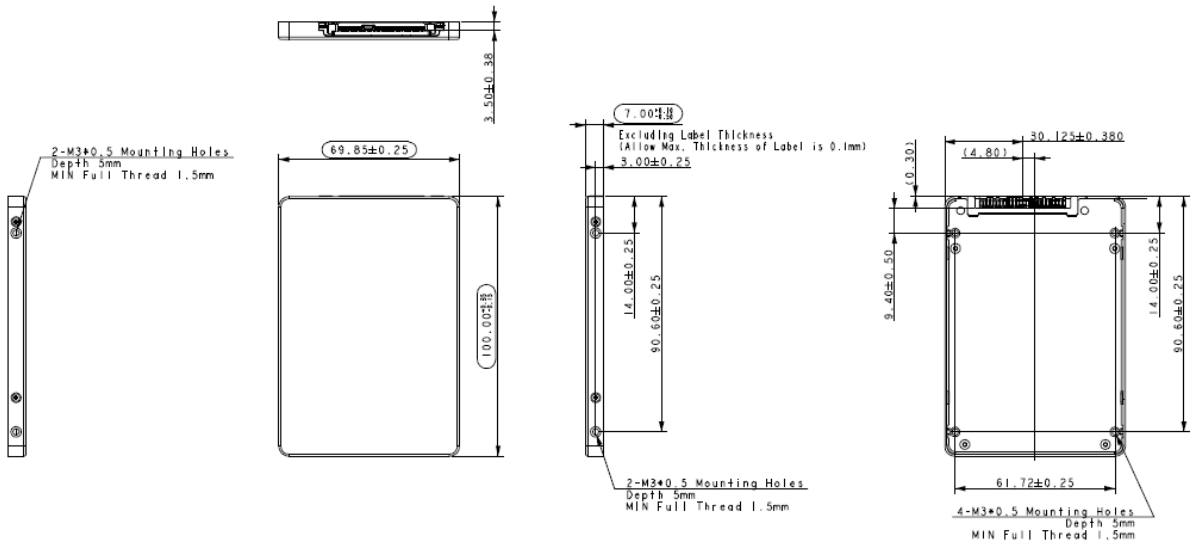

Figure 9-1 Dimensions

9.2 Net Weight

Table 9-1 Net Weight

Capacity	Net Weight (g ± 5%)
960GB	61.3
1,920GB	70.1
3,840GB	71.5

10. Product Ordering Information

10.1 Product Code Designations

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

Code 1-3 (Product Line & form factor)	PCIe U.2
Code 5-6 (Model/Solution)	PV140-25
Code 7-8 (Product Capacity)	5L: 960GB 5M: 1,920GB 5N: 3,840GB
Code 9 (Flash Type & Product Temp)	H: 3D TLC Wide temperature
Code 10 (Product Spec)	A: U.2 7mm housing
Code 12-14 (Version Number)	Random numbers generated by system
Code 15-16 (Firmware Version)	01: 960GB 02: 1,920/3,840GB

10.2 Valid Combinations

Capacity	Part Number
960GB	BB6.145LHA.00101
1,920GB	BB6.145MHA.00102
3,840GB	BB6.145NHA.00102

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
1.0	Initial release	9/28/2020