

Specifications Overview:

- **PCIe Interface**
 - Compliant with PCIe Express 3.1
 - Compliant with NVMe 1.3
 - Compatible with PCIe Gen3 x4 interface
- **Capacity**
 - 240, 480, 960 GB
- **Performance***
 - Interface burst read/write: 4 GB/sec
 - Sequential read: up to 2,660 MB/sec
 - Sequential write: up to 2,260 MB/sec
 - Random read (4K): up to 345,000 IOPS
 - Random write (4K): up to 172,000 IOPS
- **Flash Management**
 - Low-Density Parity-Check (LDPC) Code
 - Global Wear Leveling
 - Flash bad-block management
 - Flash Translation Layer: Page Mapping
 - S.M.A.R.T.
 - Power Failure Management
 - TRIM
 - Hyper Cache Technology
 - Over-Provisioning
 - DataRAID™
- **DRAM Cache for Enhanced Random Performance**
- **NAND Flash Type:** 3D TLC (BiCS3)
- **MTBF:** >1,000,000 hours
- **Endurance (in drive writes per day : DWPD)**
 - 240 GB: 3.99 DWPD
 - 480 GB: 4.03 DWPD
 - 960 GB: 3.99 DWPD
- **Temperature Range**
 - Operating: 0°C to 70°C
 - Storage: -40°C to 85°C
- **Supply Voltage**
 - 3.3 V ± 5%
- **Power Consumption***
 - Active mode: 1,725 mA
 - Idle mode: 250 mA
- **Connector Type**
 - 75-pin M.2 module pinout
- **Power Management**
 - Supports APST
 - Supports ASPM L1.2
- **Form Factor**
 - Form Factor: M.2 2280-D5-M
 - Dimensions: 80.00 x 22.00 x 3.88, unit: mm
 - Net Weight: 7.92 g
- **Shock & Vibration****
 - Shock: 1,500 G
 - Vibration: 15 G
- **Security**
 - End-to-End Data Protection
- **Reliability**
 - Thermal Sensor
 - Thermal Management Technique
- **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.

**Non-operating

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 Power Failure Management	10
5.5 Flash Translation Layer – Page Mapping.....	11
5.6 TRIM.....	11
5.7 Hyper Cache Technology	11
5.8 Over-Provisioning	11
5.9 DataRAID™	11
6. Security & Reliability Features	12
6.1 Thermal Sensor	12
6.2 Thermal Management Technique	12
6.3 End-to-End Data Protection.....	12
7. Software Interface	13
7.1 Command Set.....	13
7.2 S.M.A.R.T.	14
8. Electrical Specifications	16
8.1 Operating Voltage.....	16

8.2 Power Consumption	16
9. Physical Characteristics.....	17
9.1 Dimensions	17
9.2 Net Weight.....	17
10. Product Ordering Information.....	18
10.1 Product Code Designations.....	18
10.2 Valid Combinations.....	19

1. General Descriptions

Apacer PV310-M280 is the fastest SSD designed as M.2 2280 mechanical dimensions, providing full compliance with PCIe Gen3 x4 interface and NVMe 1.3 specifications. Built with a powerful PCIe controller, PV310-M280 delivers outstanding performance in data transfer, reaching up to 345,000/172,000 and 2,660/2,260 MB/s in IOPS and sequential read/write. Even for highly intensive applications, with End-to-End Data Protection technology, it ensures integrity and correctness of data transmission between the host and the NAND storage area, and greatly improves data reliability.

2. Functional Block

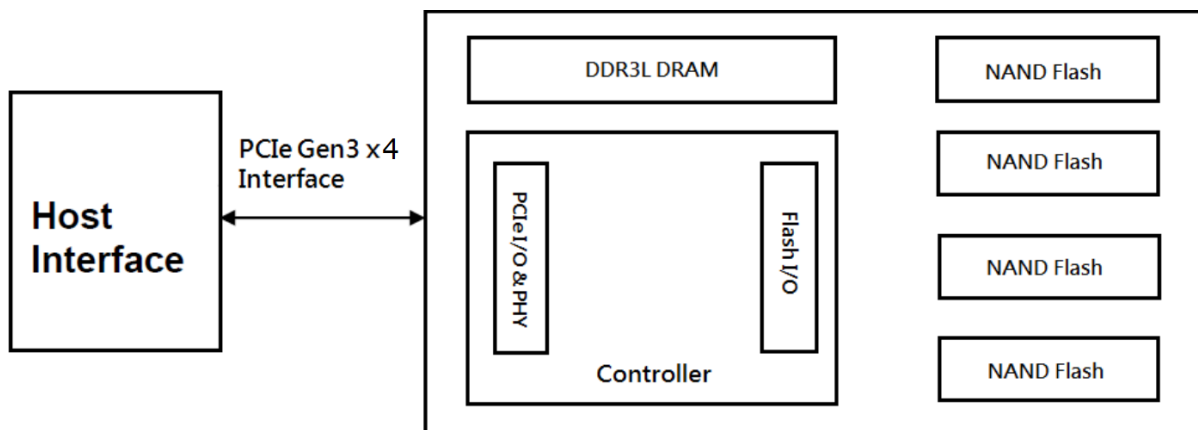


Figure 2-1 Functional Block Diagram

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 Key M connectors.




Table 3-1 Pin Assignments

Pin	Type	Description
1	GND	Ground
2	3.3V	3.3V source
3	GND	Ground
4	3.3V	3.3V source
5	PETn3	PCIe TX Differential signal defined by the PCI Express M.2 spec
6	N/C	No connect
7	PETp3	PCIe TX Differential signal defined by the PCI Express M.2 spec
8	N/C	No connect
9	GND	Ground
10	LED1#(O)	Status indicators via LED devices
11	PERn3	PCIe RX Differential signal defined by the PCI Express M.2 spec
12	3.3V	3.3V source
13	PERp3	PCIe RX Differential signals defined by the PCI Express M.2 spec
14	3.3V	3.3V source
15	GND	Ground
16	3.3V	3.3V source
17	PETn2	PCIe TX Differential signal defined by the PCI Express M.2 spec
18	3.3V	3.3V source
19	PETp2	PCIe TX Differential signal defined by the PCI Express M.2 spec
20	N/C	No connect
21	GND	Ground
22	N/C	No connect
23	PERn2	PCIe RX Differential signal defined by the PCI Express M.2 spec
24	N/C	No connect
25	PERp2	PCIe RX Differential signal defined by the PCI Express M.2 spec
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
30	N/C	No connect
31	PETp1	PCIe TX Differential signal defined by the PCI Express M.2 spec
32	N/C	No connect

Table 3-1 Pin Assignments

Pin	Type	Description
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	SMB_CLK	SMBus clock; Open Drain with pull up on platform
41	PETn0	PCIe TX Differential signal defined by the PCI Express M.2 spec
42	SMB_DATA	SMBus Data; Open Drain with pull up on platform
43	PETp0	PCIe TX Differential signal defined by the PCI Express M.2 spec
44	ALERT#	Alert notification to host system. Open Drain with pull up on platform, Active low Signals
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 specification. defined by the PCIe Mini CEM
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 Substates.
53	REFCLKn	PCIe Reference Clock signals (100 MHz) spec. defined by the PCI Express M.2
54	PEWAKE#(I/O)(0/3.3V)	Open Drain with pull up on platform; Active Low. PCIe PME Wake.
55	REFCLKp	PCIe Reference Clock signals (100 MHz) spec. defined by the PCI Express M.2
56	Reserved for MFG DATA	Manufacturing Data line. Used for SSD manufacturing only. Not used in normal operation. Pins should be left N/C in platform Socket.
57	GND	Ground
58	Reserved for MFG CLOCK	Manufacturing Clock line. Used for SSD manufacturing only. Not used in normal operation. Pins should be left N/C in platform Socket.
59	Module Key	Module Key
60	Module Key	Module Key
61	Module Key	Module Key
62	Module Key	Module Key
63	Module Key	Module Key
64	Module Key	Module Key
65	Module Key	Module Key
66	Module Key	Module Key
67	N/C	No connect
68	SUSCLK(32KHz) (I/O)(0/3.3V)	32.768 kHz clock supply input that is provided by the platform chipset to reduce power and cost for the module.
69	PEDET (NC-PCIe)	Host I/F Indication; No Connect for PCIe.
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	Ground

4. Product Specifications

4.1 Capacity

Capacity specifications of PV310-M280 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*	Cylinders	Heads	Sectors	Max LBA
240 GB	240,056,795,136	16,383	16	63	468,862,128
480 GB	480,103,104,512	16,383	16	63	937,703,088
960 GB	960,196,771,840	16,383	16	63	1,875,385,008

*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 PV310-M280 is listed below in Table 4-2.

Table 4-2 Performance Specifications

Performance	Capacity	240 GB	480 GB	960 GB
	Sequential Read* (MB/s)		2,435	2,660
Sequential Write* (MB/s)		1,040	2,085	2,260
Random Read IOPS** (4K)		187,000	345,000	340,000
Random Write IOPS** (4K)		70,000	134,000	172,000

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 64.

4.3 Environmental Specifications

Environmental specifications of PV310-M280 are shown in Table 4-3.

Table 4-3 Environmental Specifications

Item	Specifications
Operating temp.	0°C to 70°C
Non-operating temp.	-40°C to 85°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 PV310-M280. The prediction result for PV310-M280 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

PV310-M280 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 Drive Writes Per Day

Capacity	Drive Writes Per Day
240 GB	3.99
480 GB	4.03
960 GB	3.99

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 PV310-M280 LED indicators is described in Table 4-5.

Table 4-5 LED Behavior

Location	LED	Color	Description
LED A	DAS	Red	LED blinks when the drive is being accessed
LED B	PHY	Green	LED blinks when PHY is connected



5. Flash Management

5.1 Error Correction/Detection

PV310-M280 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 Power Failure Management

Power Failure Management plays a crucial role when experiencing unstable power supply. Power disruption may occur when users are storing data into the SSD. In this urgent situation, the controller would run multiple write-to-flash cycles to store the metadata for later block rebuilding. This urgent operation requires about several milliseconds to get it done. At the next power up, the firmware will perform a status tracking to retrieve the mapping table and resume previously programmed NAND blocks to check if there is any incompleteness of transmission.

Note: The controller unit of this product model is designed with a DRAM as a write cache for improved performance and data efficiency. Though unlikely to happen in most cases, the data cached in the volatile DRAM might be potentially affected if a sudden power loss takes place before the cached data is flushed into non-volatile NAND flash memory.

5.5 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.

5.6 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 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.

6. Security & 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 Management Technique

Thermal management technique 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.

7. Software Interface

7.1 Command Set

Table 7-1 summarizes the commands supported by PV310-M280.

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

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
04h	Write Uncorrectable
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 PV310-M280.

Table 8-1 Operating Range

Item	Range
Supply Voltage	3.3V ± 5%

8.2 Power Consumption

Table 8-2 lists the power consumption for PV310-M280.

Table 8-2 Power Consumption

Mode \ Capacity	240 GB	480 GB	960 GB
Active (mA)	1,590	1,635	1,725
Idle (mA)	245	250	250


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

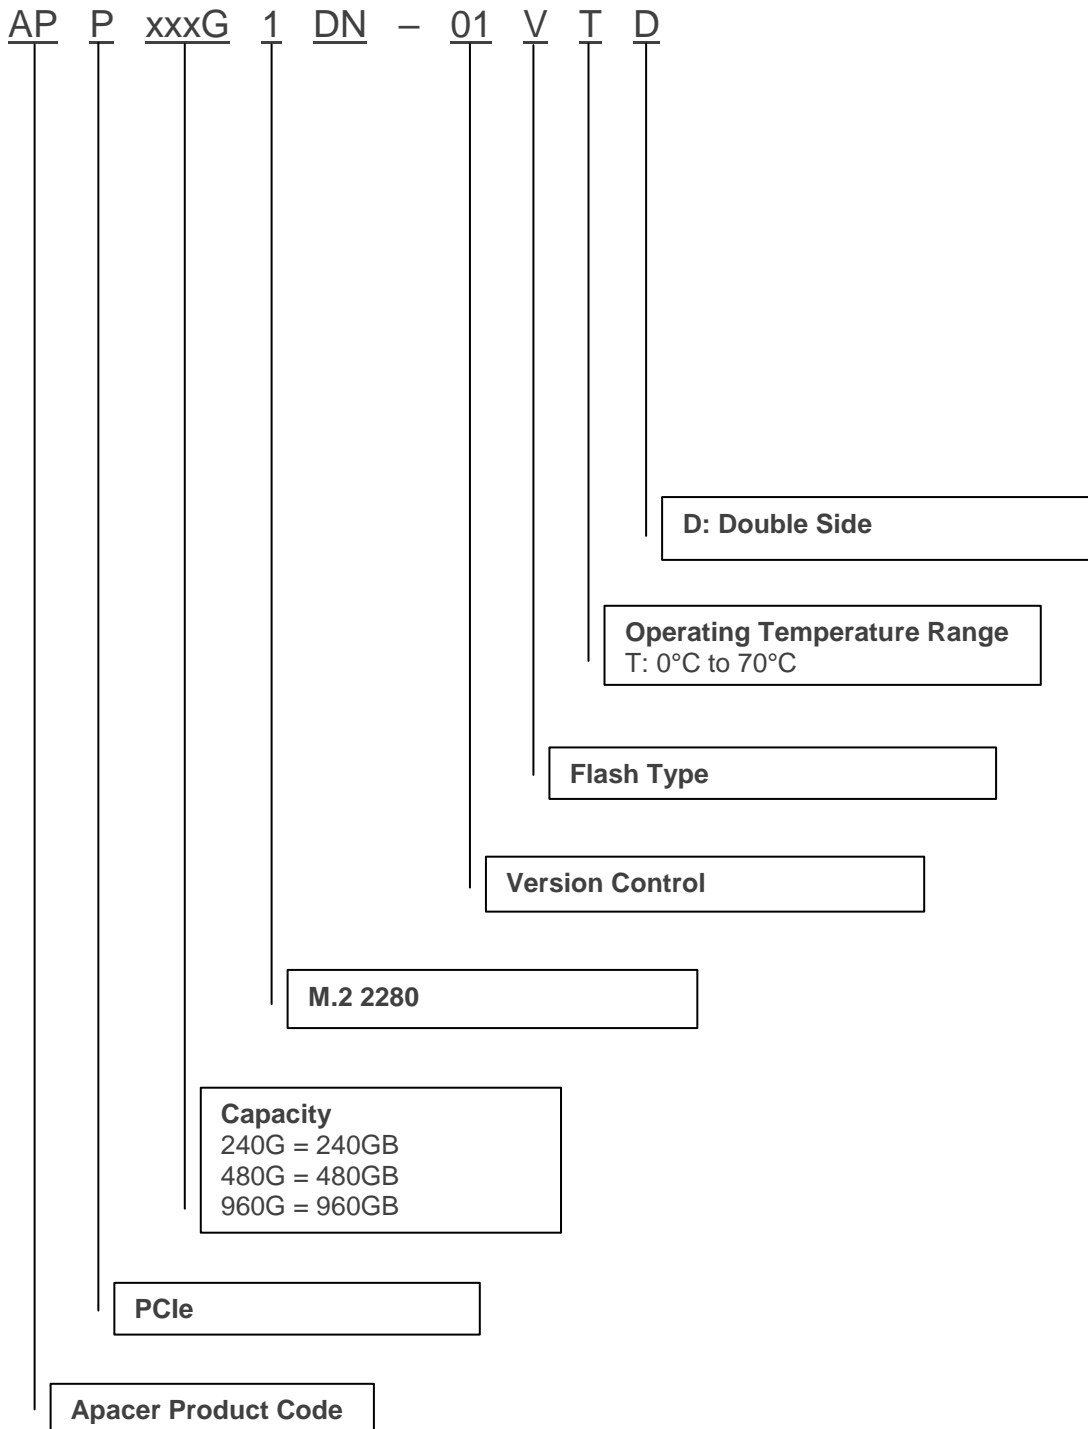


9.2 Net Weight

Capacity	Net Weight (g)
240GB	7.92
480GB	7.92
960GB	7.92

10. Product Ordering Information

10.1 Product Code Designations



10.2 Valid Combinations

Capacity	Part Number
240GB	APP240G1DN-01VTD
480GB	APP480G1DN-01VTD
960GB	APP960G1DN-01VTD

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	2/26/2019
0.2	<ul style="list-style-type: none"> - Updated product photo on the cover page - Added Endurance and LED Indicators for Drive Behavior to Specifications Overview - Added Over-Provisioning to Flash Management on Specifications Overview - Added 4.6 Endurance, 4.7 LED Indicator Behavior and 5.7 Over-Provisioning - Removed wide temperature and DataDefender support 	4/16/2019
0.3	<ul style="list-style-type: none"> - Updated Performance and Power Consumption on Specifications Overview page - Updated read/write performance for 1. General Descriptions - Updated 4.2 Performance and 8.2 Power Consumption - Updated Table 7-5 SMART (C0h) 	5/21/2019
1.0	<ul style="list-style-type: none"> - Completed endurance rating for Endurance on Specifications Overview page and 4.6 Endurance - Added Power Failure Management and DataRAID to Flash Management on Specifications Overview page - Added 5.4 Power Failure Management and 5.9 DataRAID 	6/12/2019
1.1	Removed 120GB support	6/20/2019
1.2	<ul style="list-style-type: none"> - Updated interface bust read/write from 2GB/sec to 4GB/sec at Performance on Specifications Overview page - Updated storage temperature range from -40°C to 100°C to -40°C to 85°C 	8/29/2019
1.3	Updated the note regarding random performance measuring method at 4.2 Performance	9/4/2019

Global Presence

Taiwan (Headquarters)

Apacer Technology Inc.

1F., No.32, Zhongcheng Rd., Tucheng Dist.,
New Taipei City 236, Taiwan R.O.C.

Tel: 886-2-2267-8000

Fax: 886-2-2267-2261

amtsales@apacer.com

U.S.A.

Apacer Memory America, Inc.

46732 Lakeview Blvd., Fremont, CA 94538

Tel: 1-408-518-8699

Fax: 1-510-249-9551

sa@apacerus.com

Japan

Apacer Technology Corp.

6F, Daiyontamachi Bldg., 2-17-12, Shibaura, Minato-Ku,
Tokyo, 108-0023, Japan

Tel: 81-3-5419-2668

Fax: 81-3-5419-0018

jpservices@apacer.com

Europe

Apacer Technology B.V.

Science Park Eindhoven 5051 5692 EB Son,
The Netherlands

Tel: 31-40-267-0000

Fax: 31-40-290-0686

sales@apacer.nl

China

Apacer Electronic (Shanghai) Co., Ltd

Room D, 22/FL, No.2, Lane 600, JieyunPlaza,
Tianshan RD, Shanghai, 200051, China

Tel: 86-21-6228-9939

Fax: 86-21-6228-9936

sales@apacer.com.cn

India

Apacer Technologies Pvt Ltd,

1874, South End C Cross, 9th Block Jayanagar,
Bangalore-560069, India

Tel: 91-80-4152-9061/62

Fax: 91-80-4170-0215

sales_india@apacer.com