

NVMe PCIe SSD M.2 2280 Manual

NVMe PCIe SSD is a non-volatile, solid-state storage device delivering uncompromising performance, reliability and ruggedness for environmentally challenging applications.

Manual	8/19/2019
PSFNP5xxxx5xxx	Viking Technology
Revision C	Page 1 of 23

Revision History

Date	Revision	Description	Checked By
5/12/19	A	Initial Release from modified PSFN22xxxx5xxx_A revise PN table, mechanical and pin out(add 4 lanes) and supply voltage. Add IT and CT PNs, LBA, power, performance endurance.	
7/24/19	B	Add VPFNP5002T5C4WT3 VPFNP5002T5I4WT3 VPFNP5960G5IFWT3	
8/19/19	C	Add section for "Data Integrity Assurance After Unexpected Power Loss" and a reference to whitepaper AN00025	

Manual	8/19/2019
PSFNP5xxxx5xxx	Viking Technology
Revision C	Page 2 of 23

Legal Information

Legal Information

Copyright© 2019 Sanmina Corporation. All rights reserved. The information in this document is proprietary and confidential to Sanmina Corporation. No part of this document may be reproduced in any form or by any means or used to make any derivative work (such as translation, transformation, or adaptation) without written permission from Sanmina. Sanmina reserves the right to revise this documentation and to make changes in content from time to time without obligation on the part of Sanmina to provide notification of such revision or change.

Sanmina provides this documentation without warranty, term or condition of any kind, either expressed or implied, including, but not limited to, expressed and implied warranties of merchantability, fitness for a particular purpose, and non-infringement. While the information contained herein is believed to be accurate, such information is preliminary, and should not be relied upon for accuracy or completeness, and no representations or warranties of accuracy or completeness are made. In no event will Sanmina be liable for damages arising directly or indirectly from any use of or reliance upon the information contained in this document. Sanmina may make improvements or changes in the product(s) and/or the program(s) described in this documentation at any time.

Sanmina, Viking Technology, Viking Modular Solutions, and Element logo are trademarks of Sanmina Corporation. Other company, product or service names mentioned herein may be trademarks or service marks of their respective owners.

Manual	8/19/2019
PSFNP5xxxx5xxx	Viking Technology
Revision C	Page 3 of 23

Ordering Information: M.2 2280 PCIe SSD Solid-State Drive

Part Number	Interface	Application	User Capacity (GB)	Encryption	Temperature	NAND
VPFNP5240G5C5WT3	PCIe/NVMe	Enterprise	240	Pyrite/AES256 OPAL 2.0	(0to+70°C)	TSB BiCS TLC 3D
VPFNP5480G5CHWT3	PCIe/NVMe	Enterprise	480	Pyrite/AES256 OPAL 2.0	(0to+70°C)	TSB BiCS TLC 3D
VPFNP5960G5CFWT3	PCIe/NVMe	Enterprise	960	Pyrite/AES256 OPAL 2.0	(0to+70°C)	TSB BiCS TLC 3D
VPFNP5240G5I5WT3	PCIe/NVMe	Industrial	240	Pyrite/AES256 OPAL 2.0	(-40to+85°C)	TSB BiCS TLC 3D
VPFNP5480G5IHWT3	PCIe/NVMe	Industrial	480	Pyrite/AES256 OPAL 2.0	(-40to+85°C)	TSB BiCS TLC 3D
VPFNP5960G5IFWT3	PCIe/NVMe	Industrial	960	Pyrite/AES256 OPAL 2.0	(-40to+85°C)	TSB BiCS TLC 3D
VPFNP5002T5C4WT3	PCIe/NVMe	Enterprise	2048	Pyrite/AES256 OPAL 2.0	(0to+70°C)	TSB BiCS TLC 3D
VPFNP5002T5I4WT3	PCIe/NVMe	Industrial	2048	Pyrite/AES256 OPAL 2.0	(-40to+85°C)	TSB BiCS TLC 3D

Notes:

1. Usable capacity based on a level of over-provisioning applied to wear leveling, bad sectors, index tables etc.
2. SSD's ship unformatted from the factory unless otherwise requested.
3. 1 GB = 1,000,000,000 Byte
4. One Sector = 512 Byte.
5. Lowercase x is a wildcard character that represents the device code for Flash device capacity

Manual	8/19/2019
PSFNP5xxxx5xxx	Viking Technology
Revision C	Page 4 of 23

Table of Contents

1	INTRODUCTION	8
1.1	Features	8
1.2	PCIE Interface	9
2	PRODUCT SPECIFICATIONS	10
2.1	Capacity and LBA count	10
2.2	Performance	10
2.1	Power Consumption	11
2.1.1	Throughput	11
2.1.2	Predict & Fetch	11
2.2	Electrical Characteristics	11
2.2.1	Absolute Maximum Ratings	11
2.2.2	Supply Voltage	12
2.3	Environmental Conditions	12
2.3.1	Temperature and Altitude	12
2.3.2	Shock and Vibration	12
2.3.3	Electromagnetic Immunity	13
2.4	Reliability	13
2.5	Data Security	13
2.5.1	Secure Erase	13
2.5.2	Write Protect	13
2.5.3	Encryption OPTIONAL	14
2.5.4	Data Integrity Assurance After Unexpected Power Loss*	14
2.6	Flash Management	15
2.6.1	Error Correction Code (ECC)	15
2.6.2	Wear Leveling	16
2.6.3	Bad Block Management	16
2.6.4	TRIM	16

Manual	8/19/2019
PSFNP5xxxx5xxx	Viking Technology
Revision C	Page 5 of 23

2.6.5	SMART	17
2.6.6	Over-Provision	17
2.6.7	Firmware Upgrade	17

3	MECHANICAL INFORMATION	18
----------	-------------------------------	-----------

3.1	Signal and Power Description Tables	21
-----	-------------------------------------	----

Table of Tables

<i>Table 2-1: Maximum Sustained Read and Write Bandwidth</i> _____	10
<i>Table 2-2: Power Consumption</i> _____	11
<i>Table 2-3: Operating Voltage</i> _____	12
<i>Table 2-4: Temperature and Altitude Related Specifications</i> _____	12
<i>Table 2-5: Shock and Vibration Specifications</i> _____	12
<i>Table 2-6: Reliability Specifications</i> _____	13
<i>Table 3-1: M.2 PCIE Connector Pinouts</i> _____	21

Table of Figures

<i>Figure 3-1: Dimension Details for M.2 2280</i> _____	18
---	----

1 Introduction

This document describes the specification of Viking SSD which uses PCIe interface. The Viking SSD is fully consist of semiconductor device and using NAND Flash Memory which has a high reliability and a high technology in a small form factor for using a SSD and supporting Peripheral Component Interconnect Express (PCIe) 3.0 interface standard up to 4 lanes shows much faster performance than previous SATA SSDs It could also provide rugged features with an extreme environment with a high MTBF.

1.1 Features

The SSD delivers the following features:

- Native-PCIe SSD for enterprise application
- PCI Express Gen3: Single port X4 lanes
- Compliant with PCI Express Base Specification Rev. 3.1
- Compliant with NVM Express Specification Rev.1.3
- Static and Dynamic Wear Leveling and Bad Block Management
- RoHS / Halogen-Free Compliant
- Support up to queue depth 64K
- Support Power Management: ASPM/PCI-PM L0s, L1, L1.1 and L1.2
- Support SMART and TRIM commands
- Support 48-bit addressing mode
- Firmware update
- Firmware support for encryption
- Advanced Flash Management
- Advanced Wear Leveling
- Bad Block Management
- TRIM
- SMART
- Over-Provision
- Firmware Update
- Power Management
- Support APST
- Support ASPM
- Support L1.2
- Power Consumption²
 - Idle < 910 mW
 - L1.2 < 2 mW
- Temperature Range³
 - Industrial temperature: -40°C ~ 85°C
 - Storage: -40°C ~ 85°C

Manual	8/19/2019
PSFNP5xxx5xxx	Viking Technology
Revision C	Page 8 of 23

RoHS compliant
Features Support List⁴:
End to end data path protection
Thermal throttling
SmartECCTM
SmartRefreshTM
Drive log
Support of TCG OPAL⁴
Support TCG Pyrite⁴

Notes:

1. Refer to Chapter 2 for more details
2. Refer to section on Power Consumption for more details.
3. Operational temperature is measured by device temperature sensor.
4. Supported by a separate firmware version. Further information available upon request.

1.2 PCIE Interface

- PCI Express Gen3: Single port X4 lanes
- Compliant with PCI Express Base Specification Rev. 3.1
- Compliant with NVM Express Specification Rev.1.3

For a list of supported commands and other specifics, refer to PCI and NVME specifications.

Manual	8/19/2019
PSFNP5xxxx5xxx	Viking Technology
Revision C	Page 9 of 23

2 Product Specifications

2.1 Capacity and LBA count

Raw Capacity (GB)	User Capacity (GB)	LBA Count
256	240	468,862,128
512	480	937,703,088
1024	960	1,875,385,008
2048	2048	4,000,797,360

Notes:

- Per www.idema.org, LBA1-03 spec,
LBA counts = (97,696,368) + (1,953,504 * (Advertised Capacity in GBytes – 50))

2.2 Performance

Table 2-1: Maximum Sustained Read and Write Bandwidth

Capacity (GB)	Flash Structure	Performance			
		CrystalDiskMark		IOMeter	
		Read (MB/s)	Write (MB/s)	Read IOPS	Write IOPS
240	BiCS3, TLC, 8CE	3,100	1,040	187K	245K
480	BiCS3, TLC, 16CE	3,370	2,030	369K	470K
960	BiCS3, TLC, 32CE	3,470	3,000	600K	600K
2048	BiCS3, TLC, 32CE	TBD	TBD	TBD	TBD

Notes:

- Performance measured under the following conditions:
 - CrystalDiskMark 5.1.2, 1GB range, QD=32, Thread=1
 - IOMeter, 8GB range, 4K data size, QD=32 (3) ATTO, transfer Size 8192 KB
- Performance may vary from flash configuration and platform.
- Refer to Application Note AN0006 for Viking SSD Benchmarking Methodology.
- Data is based on SSD's using Toshiba TLC BiCS3
- Typical Power Consumption

2.1 Power Consumption

Table 2-2: Power Consumption

Capacity	Flash Configuration	Power Consumption ¹			
		Read	Write	PS3	PS4
		(mW)	(mW)	(mW)	(mW)
240GB	BiCS3 TLC, 8CE	6,400	3,900	16	2
480GB	BiCS3 TLC, 16CE	7,000	5,000	16	2
960GB	BiCS3 TLC, 32CE	7,200	6,100	16	2
2048GB	BiCS3 TLC, 32CE	TBD	TBD	16	2

NOTES:

1. Power consumption is measured during the sequential read and write operations performed by CrystalDiskMark 5.1.2, 1GB range, QD=32, Thread=1

2.1.1 Throughput

Based on the available space of the disk, the SSD will regulate the read/write speed and manage the performance of throughput. When there still remains a lot of space, the firmware will continuously perform read/write action. There is still no need to implement garbage collection to allocate and release memory, which will accelerate the read/write processing to improve the performance. Contrarily, when the space is going to be used up, the SSD will slow down the read/write processing, and implement garbage collection to release memory. Hence, read/write performance will become slower.

2.1.2 Predict & Fetch

Normally, when the Host tries to read data from a PCIe SSD, the PCIe SSD will only perform one read action after receiving one command. However, the Viking SSD applies Predict & Fetch to improve the read speed. When the host issues sequential read commands to the PCIe SSD, the PCIe SSD will automatically expect that the following will also be read commands. Thus, before receiving the next command, flash has already prepared the data. Accordingly, this accelerates the data processing time, and the host does not need to wait so long to receive data.

2.2 Electrical Characteristics

2.2.1 Absolute Maximum Ratings

Values shown are stress ratings only. Functional operation outside normal operating values is not implied. Extended exposure to absolute maximum ratings may affect reliability.

2.2.2 Supply Voltage

The operating voltage is 3.3V

Table 2-3: Operating Voltage

Description	Min	Max	Unit
Operating Voltage for 3.3 V (+/- 5%)	3.135	3.465	V

2.3 Environmental Conditions

2.3.1 Temperature and Altitude

Table 2-4: Temperature and Altitude Related Specifications

Conditions	Operating	Shipping	Storage
Commercial Temperature-Case ¹	0 to 70°C	-40 to 85°C	-40 to 85°C
Industrial Temperature-Case ¹	-40 to 85°C	-40 to 85°C	-40 to 85°C
Humidity (non-condensing)	90% under 40°C	93% under 40°C	93% under 40°C

Notes:

1. T_c is measured at the surface of NAND Flash package

2.3.2 Shock and Vibration

SSD products are tested in accordance with environmental specification for shock and vibration

Table 2-5: Shock and Vibration Specifications

Stimulus	Description
Shock(non-operating)	1500G (0.5ms duration x,y,z with 1/2 sine wave)
Vibration (non-operating)	(60min /axis on 3 axes) Displacement: 1.52mm (20 ~ 80 Hz) Acceleration: 20G (80 ~ 2,000 Hz)

2.3.3 Electromagnetic Immunity

M.2 is an embedded product for host systems and is designed not to impair with system functionality or hinder system EMI/FCC compliance.

2.4 Reliability

Table 2-6: Reliability Specifications

Parameter	Description	
ECC	Correct up to 120 bits error in 2K Byte data	
MTBF	2,000,000 hours	
Write Endurance (BiCS3 TLC)	Capacity	TBW
	120GB	170
	240GB	380
	480GB	800
	960GB	1665
	1920GB	TBD
	2048GB	TBD
	3845GB	TBD
Data Retention	> 90 days at NAND expiration	

Notes:

1. The reliability specification follows JEDEC standards JESD218A and JESD219A
2. Based on 0.8 DWPD based with 3,000 PE for TLC (TBW = DWPD * CAPACITY * 365 (days) * 3 years of Warranty/1000

2.5 Data Security

2.5.1 Secure Erase

Secure Erase is a standard ATA command and will write all “0xFF” to fully wipe all the data on hard drives and SSDs. When this command is issued, the SSD controller will empty its storage blocks and return to its factory default settings.

2.5.2 Write Protect

When a SSD contains too many bad blocks and data are continuously written in, then the SSD might not be usable anymore. Thus, Write Protect is a mechanism

to prevent data from being written in and protect the accuracy of data that are already stored in the SSD.

2.5.3 Encryption OPTIONAL

- Pyrite
- AES256
- OPAL 2.0

Note: Encryption is optional requiring a special firmware support

2.5.4 Data Integrity Assurance After Unexpected Power Loss*

Enterprise SSD – An Enterprise SSD contains optional PFAIL hardware and firmware that detect and manage power failures. This allows the drive to flush the controller cache and harden data to NAND flash. No data is lost or corrupted. Refer to Viking Application Note AN0025 for details.

Industrial/Client SSD's – Viking's Industrial/Client SSD contains sophisticated provisions to protect firmware and data from corruption due to unexpected power loss. Refer to Viking Application Note AN0025 for details.

***Refer to Viking Application Note AN0025 for details.**
Pfail function is optional for Enterprise SSDs only
Industrial/Client SSDs use special firmware algorithms

Manual	8/19/2019
PSFNP5xxx5xxx	Viking Technology
Revision C	Page 14 of 23

2.6 Flash Management

2.6.1 Error Correction Code (ECC)

Flash memory cells will deteriorate with use, which might generate random bit errors in the stored data. The SSD applies a BCH ECC algorithm, which can detect and correct errors occur during read process, ensure data been read correctly, as well as protect data from corruption.

Manual	8/19/2019
PSFNP5xxx5xxx	Viking Technology
Revision C	Page 15 of 23

2.6.2 Wear Leveling

NAND flash devices can only undergo a limited number of program/erase cycles, and in most cases, the flash media are not used evenly. If some areas get updated more frequently than others, the lifetime of the device would be reduced significantly. Thus, Wear Leveling is applied to extend the lifespan of NAND Flash by evenly distributing write and erase cycles across the media.

Advanced Wear Leveling algorithm, can efficiently spread out the flash usage through the whole flash media area. Moreover, by implementing both dynamic and static Wear Leveling algorithms, the life expectancy of the NAND flash is greatly improved.

2.6.3 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”. Viking 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.

2.6.4 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.

Manual	8/19/2019
PSFNP5xxx5xxx	Viking Technology
Revision C	Page 16 of 23

2.6.5 SMART

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.

2.6.6 Over-Provision

Over Provisioning refers to the inclusion of extra NAND capacity in a SSD, which is not visible and cannot be used by users. With Over Provisioning, the performance and IOPS (Input/Output Operations per Second) are improved by providing the controller additional space to manage P/E cycles, which enhances the reliability and endurance as well. Moreover, the write amplification of the SSD becomes lower when the controller writes data to the flash.


2.6.7 Firmware Upgrade

Firmware can be considered as a set of instructions on how the device communicates with the host. Firmware will be upgraded when new features are added, compatibility issues are fixed, or read/write performance gets improved.

Manual	8/19/2019
PSFNP5xxxx5xxx	Viking Technology
Revision C	Page 17 of 23

3 Mechanical Information

Figure 3-1: Dimension Details for M.2 2280




TOP VIEW

Notes:



1. All dimensions are in millimeter. General tolerance is ± 0.15 . PCB thickness 0.8 ± 0.08
2. Refer to Ordering Information table for the complete Viking part number that describes the "xxx".

Manual	8/19/2019
PSFNP5xxxx5xxx	Viking Technology
Revision C	Page 18 of 23



BOTTOM VIEW

Manual	8/19/2019
PSFNP5xxxx5xxx	Viking Technology
Revision C	Page 19 of 23



Manual	8/19/2019
PSFNP5xxxx5xxx	Viking Technology
Revision C	Page 20 of 23

3.1 Signal and Power Description Tables

Table 3-1: M.2 PCIe Connector Pinouts

Pin No.	PCIe Pin	Description
1	GND	CONFIG_3 = GND
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#	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	PERn3	PCIe RX Differential signal defined by the PCI Express M.2 spec
12	3.3V	3.3V source
13	PERp3	PCIe RX Differential signal 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

Pin No.	PCIe Pin	Description
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
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 (I/O)(0/1.8V)	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 (I/O)(0/1.8V)	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#(O) (0/1.8V)	Alert notification to master; Open Drain with pull-up on platform; Active low.
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)(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	PEWAKE#(I/O)(0/3.3V)	PCIe PME Wake. Open Drain with pull up on platform; Active Low.
55	REFCLKp	PCIe Reference Clock signals (100 MHz) defined by the PCI Express M.2 spec.

Pin No.	PCIe Pin	Description
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 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	No connect
68	SUSCLK(32KHz) (I)(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	N/C	PEDET (NC-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