



# Seagate® IronWolf® 525 SSD

## Product Manual

<b>User Capacity</b>	<b>Standard M.2 2280-D2-M Models</b>
500 GB	ZP500NM30002
1000 GB	ZP1000NM30002
2000 GB	ZP2000NM30002

201807300, Rev B  
September 2021

---

## Revision History

Version and Date	Description of Changes
Rev B, September 2021	Updated the following: <ul style="list-style-type: none"><li>■ <a href="#">Section 1, Introduction</a></li><li>■ <a href="#">Section 2, Specifications</a></li><li>■ <a href="#">Section 6, Feature Details</a></li></ul>
Rev A, July 2021	First document release.

© 2021, Seagate Technology LLC All rights reserved. Publication number: 201807300, Rev B, September 2021

Seagate Technology reserves the right to make changes to the product(s) or information disclosed herein at any time without notice.

Seagate, Seagate Technology and the Spiral logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. Nytro and SeaTools are either trademarks or registered trademarks of Seagate Technology LLC or one of its affiliated companies in the United States and/or other countries. All other trademarks or registered trademarks are the property of their respective owners.

No part of this publication may be reproduced in any form without written permission of Seagate Technology LLC. Call 877-PUB-TEK1(877-782-8351) to request permission.

The NVMe word mark and/or NVMeExpress design mark are trademarks of NVMeExpress, Inc. The PCIe word mark and/or PCIeExpress design mark are registered trademarks and/or service marks of PCI-SIG.

When referring to drive capacity, one gigabyte, or GB, equals one billion bytes and one terabyte, or TB, equals one trillion bytes. Your computer's operating system may use a different standard of measurement and report a lower capacity. In addition, some of the listed capacity is used for formatting and other functions, and thus will not be available for data storage. Actual quantities will vary based on various factors, including file size, file format, features and application software. Actual data rates may vary depending on operating environment and other factors. The export or re-export of hardware or software containing encryption may be regulated by the U.S. Department of Commerce, Bureau of Industry and Security (for more information, visit [www.bis.doc.gov](http://www.bis.doc.gov)), and controlled for import and use outside of the U.S. Seagate reserves the right to change, without notice, product offerings or specifications.

---

---

# Contents

<b>Support</b> .....	<b>4</b>
<b>1. Introduction</b> .....	<b>5</b>
1.1 Reference Documents .....	6
<b>2. Specifications</b> .....	<b>7</b>
2.1 Models and Capacity .....	7
2.2 Performance .....	7
2.3 Power Consumption .....	9
2.4 Environmental Conditions .....	10
2.5 Reliability/Endurance .....	11
<b>3. Mechanical Information</b> .....	<b>12</b>
<b>4. Pin and Signal Descriptions</b> .....	<b>14</b>
<b>5. SMART Support</b> .....	<b>17</b>
5.1 SMART Attributes .....	17
<b>6. Feature Details</b> .....	<b>18</b>
6.1 Flash Management .....	18
6.1.1 Error Correction Code (ECC) .....	18
6.1.2 Wear Leveling .....	18
6.1.3 Bad Block Management .....	18
6.1.4 TRIM .....	18
6.1.5 SMART .....	18
6.1.6 Over Provisioning .....	19
6.1.7 Firmware Upgrade .....	19
6.1.8 Thermal Throttling .....	19
6.2 Advanced Device Security Features .....	20
6.2.1 NVMe format .....	20
6.2.2 Physical Presence SID (PSID) .....	20
6.2.3 Manufacturer's Secure ID (MSID) .....	20
6.2.4 Sanitize Operation .....	20
6.3 SSD Lifetime Management .....	21
6.3.1 Total Bytes Written (TBW) .....	21
6.3.2 Media Wear Indicator .....	21
6.3.3 Read Only Mode (End of Life) .....	21
6.4 An Adaptive Approach to Performance Tuning .....	22
6.4.1 Throughput .....	22
6.4.2 Predict & Fetch .....	22
6.4.3 SLC Caching .....	22
<b>7. Safety, Standards, and Compliance</b> .....	<b>23</b>
7.1 Regulatory Model Numbers .....	23

---

## Support

For Internal SSD Support, visit: <https://www.seagate.com/support/products/>

For Firmware Download and Tools Download for Secure Erase, visit: <https://www.seagate.com/support/downloads/>

For information regarding online support and services, visit: <http://www.seagate.com/contacts/>

For information regarding Warranty Support, visit: <http://www.seagate.com/support/warranty-and-replacements/>

For information regarding data recovery services, visit:

<http://www.seagate.com/services-software/seagate-recovery-services/recover/>

For Seagate OEM and Distribution partner and Seagate reseller portal, visit: <http://www.seagate.com/partners>

# 1. Introduction

The Seagate® IronWolf® 525 SSD is built for NAS, and delivers high endurance, performance and caching speeds of up to 5000MB/s for NVMe-compatible systems. The Seagate IronWolf 525 SSD offers PCIe Gen 4 x4 interface with NVMe protocol support in an ultra-small M.2 form factor..

**Table 1 The IronWolf 525 SSD Features**

Feature	Description	
<b>Capacity (User)</b>	<ul style="list-style-type: none"> <li>500 GB, 1000 GB, 2000 GB</li> </ul>	
<b>Certifications, Eco-Compliance</b>	<ul style="list-style-type: none"> <li>CE, UL, FCC, BSMI, KCC, Microsoft WHQL, VCCI, CB</li> <li>RoHS</li> </ul>	
<b>Dimensions</b>	<ul style="list-style-type: none"> <li>Length, Max (mm/in): 80.15 mm/3.156 in</li> <li>Width, Max (mm/in): 22.15 mm/0.872 in</li> <li>Height, Max (mm/in): 3.58 mm/0.141 in</li> </ul>	
<b>Form Factor</b>	<ul style="list-style-type: none"> <li>M.2 2280-D2-M</li> </ul>	
<b>Weight</b>	<ul style="list-style-type: none"> <li>500 GB 8.0g, 1000 GB 8.5g, 2000 GB: 8.7g</li> </ul>	
<b>Endurance</b>	Total Bytes Written <ul style="list-style-type: none"> <li>500 GB: 700 TB</li> <li>1000 GB: 1400 TB</li> <li>2000 GB: 2800 TB</li> </ul>	See <a href="#">Section 2.5, Reliability/Endurance</a> .
<b>Interface Compliance</b>	<ul style="list-style-type: none"> <li>NVMe 1.3</li> <li>PCI Express Base 4.0, PCIe Gen 4 x 4 lane, and backward compatible to PCIe Gen3, Gen 2, and Gen 1</li> </ul>	
<b>NAND</b>	<ul style="list-style-type: none"> <li>3D TLC</li> </ul>	
<b>Operating Systems</b>	<ul style="list-style-type: none"> <li>Windows 10 (64 bit)</li> <li>Ubuntu 16.04, 18.04</li> <li>CentOS 6, 7</li> </ul>	
<b>Performance Random</b>	<ul style="list-style-type: none"> <li>Read: Up to 760,000 IOPS</li> <li>Write: Up to 700,000 IOPS</li> </ul>	Actual performance might vary depending on use conditions and environment. See <a href="#">Section 2.2, Performance</a> .
<b>Performance Sequential</b>	<ul style="list-style-type: none"> <li>Read: Up to 5000MB/s</li> <li>Write: Up to 4400MB/s</li> </ul>	Actual performance might vary depending on the capacity, use conditions and environment. See <a href="#">Section 2.2, Performance</a> .
<b>Power Consumption</b>	<ul style="list-style-type: none"> <li>Active Power, Average: &lt;7.6 W</li> <li>Idle Power PS3, Average: &lt;30 mW</li> <li>Low Power L1.2 mode: &lt; 2 mW</li> </ul>	See <a href="#">Section 2.3, Power Consumption</a> .
<b>Power Management</b>	<ul style="list-style-type: none"> <li>Supports ActiveStatePower Management (ASPM)</li> <li>Supports Autonomous Power StateTransition (APST)</li> <li>Supports L1.2</li> </ul>	
<b>Security</b>	<ul style="list-style-type: none"> <li>TCG Pyrite</li> </ul>	
<b>Reliability</b>	<ul style="list-style-type: none"> <li>End-to-end data path protection</li> <li>MTBF: 1.8 million hours</li> <li>UBER: 1 error in 10<sup>16</sup> bits read</li> </ul>	

**Table 1 The IronWolf 525 SSD Features (continued)**

Feature	Description	
<b>Shock and Vibration</b>	Shock <ul style="list-style-type: none"> <li>■ Non-Operating: 1,500 G, at 0.5 ms</li> </ul>	<b>See Section 2.4, Environmental Conditions.</b>
	Vibration <ul style="list-style-type: none"> <li>■ Non-Operating: 1.52 G<sub>RMS</sub>, (20 to 80 Hz, Frequency)</li> </ul>	
<b>Temperature Range</b>	<ul style="list-style-type: none"> <li>■ Operating: 0°C to 70°C</li> <li>■ Non-operating: -40°C to 85°C</li> </ul>	
<b>Voltage</b>	<ul style="list-style-type: none"> <li>■ Min = 3.14V±5%</li> <li>■ Max = 3.47V±5%</li> </ul>	
<b>Warranty</b>	Five years, or when the device reaches Host TBW, whichever happens first. Endurance rating valid for SSD Life Remaining > 1%.	

## 1.1 Reference Documents

In case of conflict between this document and the following reference documents, this document takes precedence.

- PCIe Specifications
  - PCIe - PCI Express Electromechanical specification, revision 4.0
  - NVMe - Non Volatile Memory Express specification 1.3
  - PCIe CEM - PCI Express Card Electromechanical specification, revision 1.1
  - PCI Express M.2 Specification, revision 1.1
- Trusted Computing Group (TCG) Documents
  - Storage Work Group Security Subsystem Class: Opal, Version 2.00
- Seagate Documentation
  - SeaTools™ SSD GUI User Guide - Publication Number: 100837824
  - SeaChest for SSD User Guide - Publication Number: 100847684

## 2. Specifications

### 2.1 Models and Capacity

**Table 2 Models and Capacity**

User Capacity	Standard M.2 2280-D2-M Models
500 GB	ZP500NM30002
1000 GB	ZP1000NM30002
2000 GB	ZP2000NM30002

**NOTE** About capacity:

- Sector Size: 512 Bytes
- User-addressable LBA count = (97696368) + (1953504 x (Desired Capacity in Gb-50.0)) From International Disk Drive Equipment and Materials Association (IDEMA) (*LBA1-03\_standard.doc*)

### 2.2 Performance

**Table 3 FOB Random and Sequential Read and Write Performance (PCIe Gen4 x4)**

Parameter	500 GB	1000 GB	2000 GB
Sequential Read (Max, MB/s), 128KB	5000	5000	5000
Sequential Write (Max, MB/s), 128KB	2500	4400	4400
Random Read (Max, IOPS), 4KB QD32 T8	420,000	760,000	740,000
Random Write (Max, IOPS), 4KB QD32 T8	630,000	700,000	700,000

**NOTE** About Performance:

- Fresh out of box (FOB) performance obtained on newly formatted drive. Performance may vary based on the SSD's firmware version, system hardware, and configuration.
- Performance based on Windows 10 host with PCIe Gen4 motherboard.
- CrystalDiskMark 8.0.1, 1GB range, QD=32, Thread=1, 1GB range, QD=32, Thread=8 (random)
- Sequential Read/Write is measured while testing 1000 MB five times by CrystalDiskMark.

**Table 4 Sustained Random and Sequential Read and Write Performance (PCIe Gen4x4)**

Parameter	500 GB	1000 GB	2000 GB
Sequential Read (Max, MB/s), 128KB QD32	3300	4350	4300
Sequential Write (Max, MB/s), 128KB QD32	525	995	965
Random Read (Max, IOPS), 4KB QD256	230,000	445,000	425,000
Random Write (Max, IOPS), 4KB QD256	10,800	19,500	19,500

**NOTE**

About performance:

- Performance based on host with PCIe Gen4 motherboard and Ubuntu 16.04 operating system.
- Performance is measured with the following conditions
  - Sustained Sequential: FIO, full-range, 128K data size, QD=32
  - Sustained Random: FIO, full-range, 4K data size, QD256
- Performance may vary based on the SSD's firmware version, system hardware, and configuration.

**Table 5 FOB Random & Sequential, Read & Write Performance (PCIe Gen3 x4)**

Parameter	500 GB	1000 GB	2000 GB
Sequential Read (Max, MB/s), 128KB QD32	3400	3400	3400
Sequential Write (MB/s), 128KB QD32	2500	3200	3200
Random Read (IOPS), 4KB QD256	420,000	640,000	640,000
Random Write (IOPS), 4KB QD256	10,800	19,500	19,500

**NOTE**

About performance:

- Fresh out of box (FOB) performance obtained on newly formatted drive. Performance may vary based on SSD's firmware version, system hardware, and configuration.
- Performance based on Windows 10 host with PCIe Gen3 motherboard.
- CrystalDiskMark 8.0.1, 1GB range, QD=32, Thread=1, 1GB range, QD=32, Thread=8 (random)
- Sequential Read/Write is measured while testing 1000 MB five times by CrystalDiskMark.



**Table 6 Sustained Random and Sequential Read and Write Performance (PCIe Gen3 x4)**

Parameter	500 GB	1000 GB	2000 GB
Sequential Read (MB/s), Sustained, 128KB QD32	3250	3300	3300
Sequential Write (MB/s) Sustained, 128KB QD32	525	995	965
Random Read (IOPS) Sustained, 4KB QD256	230,000	445,000	425,000
Random Write (IOPS) Sustained, 4KB QD256	10,800	19,500	19,500

**NOTE**

About performance:

- Performance based on host with PCIe Gen3 motherboard and Ubuntu 16.04 operating system.
- Performance is measured with the following conditions:
  - Sustained Sequential: FIO, full-range, 128K data size, QD=32
  - Sustained Random: FIO, full-range, 4K data size, QD=256
- Performance may vary based on the SSD's firmware version, system hardware, and configuration.

## 2.3 Power Consumption

**Table 7 Power Consumption**

Parameter	500 GB	1000 GB	2000 GB
Max. Average Read (mW)	6200	7600	7600
Max. Average Write (mW)	5600	6500	6500
Idle PS3 (mW) (APST off)	20	20	30
L1.2 (mW)	2	2	2

**NOTE**

About power consumption:

- The average value of power consumption is based on 100% conversion efficiency.
- The measured power voltage is 3.3 V.
- Measured under ambient temperature.
- Power Consumption can differ according to flash configuration and platform.
- Power consumption is measured during the sequential read and write (128KB) operations performed by Linux FIO3.7.
- APST means Autonomous Power State Transitions in the NVMe specification.

## 2.4 Environmental Conditions

**Table 8 Temperature, Humidity, Shock**

Specification	Value
<b>Temperature</b> Operating (case temperature at specific airflow) Non-operating	0°C to 70°C -40°C to 85°C
<b>Humidity</b> Operating Non-operating (storage)	90% 93%
<b>Shock</b> Non-operating	1,500 G, duration 0.5 ms
<b>Vibration</b> Non-operating	1.52 G <sub>RMS</sub> , (20Hz to 80Hz, Frequency)

**NOTE**

**Temperature** is measured without condensation. Operating mode temperature is measured by temperature sensor, SMART Attribute.

**Airflow** is suggested. Airflow allows the device to be operated at the appropriate temperature for each component during heavy workloads environments.

**Shock and vibration** results assume that the SSD is mounted securely with the input vibration applied to the SSD mounting. These specifications do not cover connection issues that may result from testing at this level. The measured specification is in root mean square (RMS) form.

- **Non-operating Shock.** The limits of non-operating shock applies to all conditions of handling and transportation. This includes both isolated SSD and integrated SSDs. Shock may be applied in the X, Y, or Z-axis.
- **Non-Operating Vibration.** The limits of non-operating vibration shall apply to all conditions of handling and transportation. This includes both isolated SSD and integrated SSDs. Vibration may be applied in the X, Y, or Z-axis.

## 2.5 Reliability/Endurance

**Table 9 Reliability/Endurance**

Specification	Value
Mean time between failures (MTBF)	1.8 million hours
Bit Error Rate	1 error in $10^{16}$ bits read
Endurance	Total Bytes Written <ul style="list-style-type: none"><li>■ 500 GB: 700 TB</li><li>■ 1000 GB: 1400 TB</li><li>■ 2000 GB: 2800 TB</li></ul>

**NOTE**

About endurance:

- The SSD achieves the specified MTBF in an operational environment that complies with the operational temperature range specified in this manual. Operating temperatures are measured by temperature sensor.
- Endurance rating valid for SSD Life Remaining > 1%.
- Endurance is characterized while running Client JESD219A workload (per JESD218 specification).

### 3. Mechanical Information


**Weight:** 8.0 g (500 GB) 8.5 g (1000 GB), 8.7 g (2000 GB)

**Height:** 3.5 mm±0.08 mm

**Width:** 22 mm±0.15 mm

**Length:** 80 mm±0.15 mm

**Figure 1 IronWolf 525 SSD Top Side View**



**Figure 2 IronWolf 525 SSD Bottom Side View**





Figure 3 IronWolf 525 SSD Side View



Notes

Notes :

1.  = Max Component Height
2.  = No Component
3.  = No Component / Signal Vias / Signal Copper / Printing
4. General Tolerance  $\pm 0.15\text{mm}$
5.  are critical dimensions

## 4. Pin and Signal Descriptions

Table 10 Pin Descriptions

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
28	N/C	No connect
29	PETn1	PCIe TX Differential signal defined by the PCI Express M.2 spec
30	N/C	No connect

Table 10 Pin Descriptions (continued)

Pin No.	PCIe Pin	Description
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.
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.

Table 10 Pin Descriptions (continued)

Pin No.	PCIe Pin	Description
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



## 5. SMART Support

The IronWolf 525 SSD supports the SMART command set.

### 5.1 SMART Attributes

The following table lists SMART Attributes and Descriptions.

**Table 11 SMART Attributes (Log Identifier 02h)**

Bytes Index	Bytes	Description
[0]	1	Critical Warning
[2:1]	2	Composite Temperature
[3]	1	Available Spare
[4]	1	Available Spare Threshold
[5]	1	Percentage Used
[31:6]	26	Reserved
[47:32]	16	Data Units Read
[63:48]	16	Data Units Written
[79:64]	16	Host Read Commands
[95:80]	16	Host Write Commands
[111:96]	16	Controller Busy Time
[127:112]	16	Power Cycles
[143:128]	16	Power On Hours
[159:144]	16	Unsafe Shutdowns
[175:160]	16	Media and Data Integrity Errors
[191:176]	16	Number of Error Information Log Entries
[195:192]	4	Warning Composite Temperature Time
[199:196]	4	Critical Composite Temperature Time

## 6. Feature Details

### 6.1 Flash Management

#### 6.1.1 Error Correction Code (ECC)

Flash memory cells will deteriorate with use, which might generate random bit errors in the stored data. Thus, IronWolf 525 SSD applies the fourth generation LDPC (Low Density Parity Check) of ECC algorithm, which can detect and correct errors that occur during read process, ensure data has been read correctly, as well as protect data from corruption.

#### 6.1.2 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.

Seagate provides advanced Wear Leveling algorithm, which 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.

#### 6.1.3 Bad Block Management

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

#### 6.1.4 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. All trimmed blocks will instantly read as zeros regardless of their actual content.

#### 6.1.5 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.

### 6.1.6 Over Provisioning

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.

### 6.1.7 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.

### 6.1.8 Thermal Throttling

The purpose of thermal throttling is to prevent any components in a SSD from over-heating during read and write operations. The device is designed with an on-die and an on-board thermal sensor, and with its accuracy, firmware can apply different levels of throttling to achieve the purpose of protection efficiently and proactively via SMART reading.

**Table 12 Current version: Thermal Throttling 2.0**

Item	Content
Smart reporting temperature	Flash normalized case temperature
Reference of temp. reading	On-board thermal sensor, Controller on-die thermal sensor
tmt1 threshold	82°C per Smart reported
tmt2 threshold	85°C per Smart reported
Protect threshold	95°C per Smart reported
Protect controller threshold	110°C from on-die thermal sensor
Fatal threshold	120°C from on-die thermal sensor
Resume performance threshold	78°C per Smart reported
Temperature polling frequency	Every 1 sec
TMT1_state impact	-10% CE
TMT2_state impact	-20% CE

**NOTE**

For optimal performance:

- Provide sufficient airflow and cooling.
- Use a motherboard-integrated heat sink and cooling system.

## 6.2 Advanced Device Security Features

### 6.2.1 NVMe format

Secure Erase is a standard NVMe format command and it writes all "0xFF" to fully wipe all the data on the SSDs. When this command is issued, the SSD controller erases its storage blocks and returns the drive to its factory default settings.

### 6.2.2 Physical Presence SID (PSID)

The Physical Presence SID (PSID) is defined by TCG Pyrite as a 32-character string and its purpose is to revert the SSD back to its manufacturing setting when the drive is set via TCG Pyrite (non-SED). The PSID code is printed on the SSD's label. All user data will be erased when reverting the drive back to manufacturing settings.

### 6.2.3 Manufacturer's Secure ID (MSID)

The Manufacturer's Secure ID (MSID) is defined by TCG Pyrite as a 32-character string and is assigned during the manufacturing process, which is a password that cannot be changed by the host system. MSID can be obtained electronically from the drive across the interface. On acquiring the SSD, the user must set a new password. If a new password is not set the SSD might be potentially taken control by anyone who can reset the MSID. Such an attack on the drive is known as Denial of Service (DoS) since the rightful owner has been locked out.

### 6.2.4 Sanitize Operation

The Sanitize feature is an alternative to the existing secure erase capabilities through the Format NVM command and makes a robust data security by ensuring the user data from the drive's media, caches and the Controller Memory Buffer are all wiped by the block erase operations, overwriting or destroying the encryption key. The following table illustrates the types of Sanitize Operations supported.

**Table 13 Supported Sanitize Operations**

Drive Security Type	Sanitize Operation			TCG Commands	
	Overwrite	Block Erase	Crypto Erase	PSID Revert Process	Instant Security Erase
Non-SED (TCG Pyrite)	Yes	Yes	No	Yes	No

**NOTE**

Sanitize Overwrite command completion takes at least one hour per terabyte per pass. The number of passes is drive-selectable. The NVMe spec default is 16 passes. Contact Seagate Support for more detailed information.

## 6.3 SSD Lifetime Management

### 6.3.1 Total Bytes Written (TBW)

TBW (total bytes written) is a measurement of the SSDs' expected lifespan, which represents the amount of data written to the device. To calculate the TBW of a SSD, the following equation is applied:

$$TBW = [(NAND\ Endurance) \times (SSD\ Capacity)] / WAF$$

**NAND Endurance:** NAND endurance refers to the P/E (Program/Erase) cycle of a NAND flash. **SSD Capacity:** The SSD capacity is the specific capacity in total of a SSD.

**WAF:** Write Amplification Factor (WAF) is a numerical value representing the ratio between the amount of data that a SSD controller needs to write and the amount of data that the host's flash controller writes. A better WAF, which is near 1, guarantees better endurance and lower frequency of data written to flash memory.

### 6.3.2 Media Wear Indicator

Actual life indicator reported by SMART Attribute byte index [5], Percentage Used, recommends User to replace drive when reaching to 100%.

### 6.3.3 Read Only Mode (End of Life)

When drive is aged by cumulated program/erase cycles, media worn-out may cause increasing numbers of later bad block. When the number of available spare is less the threshold(5%, SMART attribute log ID 02h Byte4), the drive will notify Host through AER event and Critical Warning to enter Read Only Mode to prevent further data corruption. User should start to replace the drive with another one immediately.

## **6.4 An Adaptive Approach to Performance Tuning**

### **6.4.1 Throughput**

Based on the available space of the disk, the drive 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 being used up, the drive will slow down the read/write processing, and implement garbage collection to release memory. Hence, read/write performance will become slower.

### **6.4.2 Predict & Fetch**

Normally, when the Host tries to read data from the PCIe SSD, the PCIe SSD will only perform one read action after receiving one command. However, the drive 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.

### **6.4.3 SLC Caching**

The firmware design of the device currently adopts dynamic caching to deliver better performance for better endurance and consumer user experience. The SLC caching size is up to 1/3 of free capacity of the SSD.

## 7. Safety, Standards, and Compliance

Each Hard Drive and Solid State Drive ("device") has a product label that includes certifications that apply to that specific drive. The following information provides an overview of requirements that may apply to the drive.

**NOTE**

***The most up to date information on Safety, Standards, and Compliance for this product is available in the Seagate HDD and SSD Regulatory Compliance and Safety document. You can find this document on the Seagate Support page (Scroll down to the Compliance section.):***  
<https://www.seagate.com/support/products/>

### 7.1 Regulatory Model Numbers

The following regulatory model number represents all features and configurations in the series:

- STA020



**Seagate Technology LLC**

*AMERICAS Seagate Technology LLC 47488 Kato Road, Fremont, California 94538, United States, 510-661-1000*

*ASIA/PACIFIC Seagate Singapore International Headquarters Pte. Ltd. 90 Woodlands Avenue 7, Singapore 737911, 65-6412-3666*

*EUROPE, MIDDLE EAST AND AFRICA Seagate Technology (Netherlands) B.V. Koolhovenlaan 1, 1119 NB Schiphol-Rijk, Netherlands, 31-20-316-7300*

**Publication Number: 201807300, Rev B**

**September 2021**