Seagate® FireCuda® X1070 SSD

Product Manual

User Capacity	Standard Models
1000GB	ZP1000GS30001
2000GB	ZP2000GS30001
4000GB	ZP4000GS30001

210892300, Rev A
February 2026

Revision History

Version and Date	Description of Changes
February 2026	First document release.

© 2026, Seagate Technology LLC All rights reserved. Publication number: 210892300, Rev A, February 2026

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. FireCuda 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 NVMExpress design mark are trademarks of NVMExpress, Inc. The PCIe word mark and/or PCIExpress 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), and controlled for import and use outside of the U.S. Seagate reserves the right to change, without notice, product offerings or specifications.

Contents

Support

Product support: www.seagate.com/support/products

Compliance data: www.seagate.com/support

Firmware and tools downloads: www.seagate.com/support/downloads

Online support and services: www.seagate.com/contacts

Warranty support: www.seagate.com/warranty

Data recovery services: www.seagate.com/rescue

Seagate OEM and distribution partners: www.seagate.com/partners

1. Introduction

Seagate® FireCuda® X1070 SSD is a mainstream PCIe Gen4 x4 SSD delivering high-speed performance for gamers.

Table 1 FireCuda X1070 SSD features

Feature	Description	
Capacity (User)	■ 1000GB, 2000GB, 4000GB	
Certifications, Eco-Compliance	■ BSMI, CE, FCC, KCC, Microsoft WHQL, RoHS, UL, VCCI	
Dimensions	1000GB ■ Length (max): 80.15 mm ■ Width (max): 22.15 mm ■ Height (max): 2.20 mm 2000GB, 4000GB ■ Length (max): 80.15 mm ■ Width (max): 22.15 mm ■ Height (max): 2.25 mm	
Form Factor	■ M.2 2280-S-M	
Weight	■ 1000GB, 2000GB: 6.0 g ■ 4000GB: 7.1 g	
Endurance	Total Bytes Written ■ 1000GB: 600TB ■ 2000GB: 1200TB ■ 4000GB: 2400TB	See Section 2.5, *Reliability and endurance*.
Interface Compliance	■ NVMe 1.4 ■ PCI Express Base 4.0, PCIe Gen4 x4 lane, and backward compatible to PCIe Gen3, Gen2, and Gen1	
NAND	■ QLC	
Operating Systems	■ Windows 11 ■ Ubuntu 20.04, 22.04 ■ CentOS 7, 8	
Performance Random	■ Read: Up to 900,000 IOPS ■ Write: Up to 1,000,000 IOPS	Actual performance may vary depending on use conditions and environment. See Section 2.2, *Performance*.
Performance Sequential	■ Read: Up to 7200MB/s ■ Write: Up to 6500MB/s	Actual performance may vary depending on the capacity, use conditions and environment. See Section 2.2, *Performance*.
Power Consumption	■ Active Power, Average: < 5.7 W ■ Idle Power PS3, Average: < 20 mW ■ Low Power L1.2 mode: < 2.5 mW	See Section 2.3, *Power consumption*.
Power Management	■ Supports ActiveStatePower Management (ASPM) ■ Supports Autonomous Power StateTransition (APST) ■ Supports L1.2	

Table 1 FireCuda X1070 SSD features *(continued)*

Feature	Description	
Reliability	■ End-to-end data path protection ■ MTBF: 1.8 million hours ■ UBER: 1 error in 10^{16} bits read	
Shock and Vibration	Shock ■ Non-Operating: 1,500 *g* at 0.5 ms	See Section 2.4, *Environmental conditions*.
	Vibration ■ Non-Operating: 1.52 GRMS (20 to 80 Hz Frequency)	
Temperature Range	■ Operating: 0°C to 70°C ■ Non-operating: -40°C to 85°C	
Voltage	■ Min = 3.14 V ±5% ■ Max = 3.47 V ±5%	
Warranty	Five years, or when the device reaches Host TBW, whichever happens first. Endurance rating valid for SSD Life Remaining > 1%.	

1.1 References

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 1.4
  - PCIe CEM - PCI Express Card Electromechanical specification, revision 1.1
  - PCI Express M.2 Specification, revision 1.1
- Seagate Downloads are available at Seagate Support: https://www.seagate.com/support/

2. Specifications

2.1 Models and capacity

Table 2 Models and capacity

User Capacity	Standard Models
1000GB	ZP1000GS30001
2000GB	ZP2000GS30001
4000GB	ZP4000GS30001

NOTE

Sector Size: 512 bytes (default)
User-addressable LBA counts = (97,696,368) + (1,953,504 x (Advertised Capacity in Gb-50.0)). From International Disk Drive Equipment and Materials Association (IDEMA) (*LBA1-03_standard.doc*)

Sector Size: 4096 bytes
User-addressable LBA count = (12,212,046) + (244,188 * (Desired Capacity in Gb-50.0)). From International Disk Drive Equipment and Materials Association (IDEMA) (*LBA1-03_standard.doc*)

2.2 Performance

Table 3 Random and sequential read and write performance

Capacity	CrystalDiskMark			
	Sequential Read (MB/s)	Sequential Write (MB/s)	Random Read (IOPS)	Random Write (IOPS)
1000GB	7200	6000	850,000	900,000
2000GB	7200	6500	900,000	1,000,000
4000GB	7200	6500	900,000	900,000

NOTE About performance:

- Fresh-ou- of-box (FOB) performance obtained on newly formatted drive. Performance may vary based on the SSD's firmware version, system hardware, and configuration.
- Performance is based on an AMD Gen4 X670E platform with an AMD Ryzen 9 7950X 16-core CPU and 16GB of DDR4-3200 memory.
- Performance is measured with the following conditions:
  - Sequential: CrystalDiskMark 8.0.4, 1GB range, QD=16, Thread=1 (Default Affinity

enabled)

- Random: CrystalDiskMark 8.0.4, 1GB range, QD=32, Thread=16 (Default Affinity enabled)

2.3 Power consumption

Table 4 Power consumption

Capacity	Power Consumption			
	Max Avg Read (W)	Max Avg Write (W)	Idle PS3 (mW)	L1.2 (mW)
1000GB	4.6	4.4	20	2.5
2000GB	4.7	4.6	20	2.5
4000GB	5.7	5.1	20	2.5

NOTE About power consumption:

- The average value of power consumption is based on 100% conversion efficiency.
- Power consumption measured during sequential read and sequential write while running CrystalDiskMark 8.0.4 x64 1GB Range, QD128, Thread=1.
- The measured power voltage is 3.3 V.
- Measured under ambient temperature.
- Power consumption can differ with configuration and platform.
- Power consumption during read and write operations is measured at the M.2 slot closest to the CPU on an AMD Gen4 X670E platform configured with an AMD Ryzen 9 7950X 16-core processor and 16GB of DDR4-3200 memory.

2.4 Environmental conditions

Table 5 Temperature, humidity, and shock

Specification	Value
Temperature Operating (case temperature at specific airflow) Non-operating	 0°C to 70°C -40°C to 85°C
Humidity Operating Non-operating (storage)	 90% 93%
Shock Non-operating	 1,500 *g*, duration 0.5 ms
Vibration Non-operating	 1.52 GRMS, (20Hz to 80Hz Frequency)

NOTE

Temperature is measured without condensation. For the purposes of this specification, the SSD's operating temperature refers to the composite temperature reported in the drive's SMART attributes.

Airflow is suggested. Airflow allows the device to be operated at the appropriate temperature for each component during heavy workload 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-axes.
- **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-axes.

2.5 Reliability and endurance

Table 6 Reliability and 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 ■ 1000GB: 600TB ■ 2000GB: 1200TB ■ 4000GB: 2400TB

NOTE The SSD achieves the specified MTBF with an operational temperature that complies with Table 5, *Temperature, humidity, and shock*, on page 10.

Endurance rating valid for SSD Life Remaining > 1%.

3. Mechanical dimensions

This section includes weights, dimensions, and mechanical drawings.

Table 7 Weight and dimensions

Capacity	Weight (g)	Length (Max)	Width (Max)	Height (Max)
1000GB	6.0	80.15 mm	22.15 mm	2.20 mm
2000GB	6.0			2.25 mm
4000GB	7.1			

Figure 1 FireCuda X1070 SSD M.2 2280-S-M top view

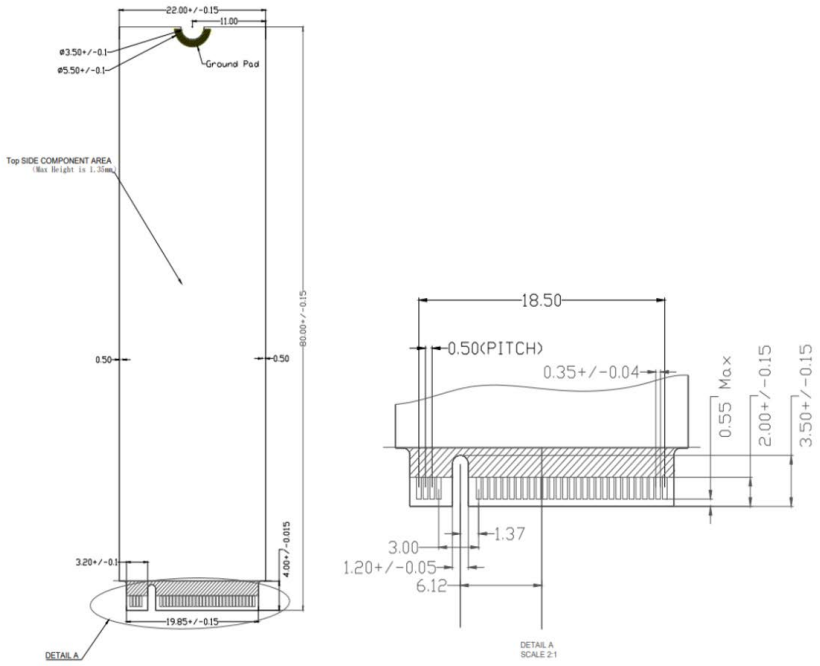

Figure 2 FireCuda X1070 SSD M.2 2280-S-M side view

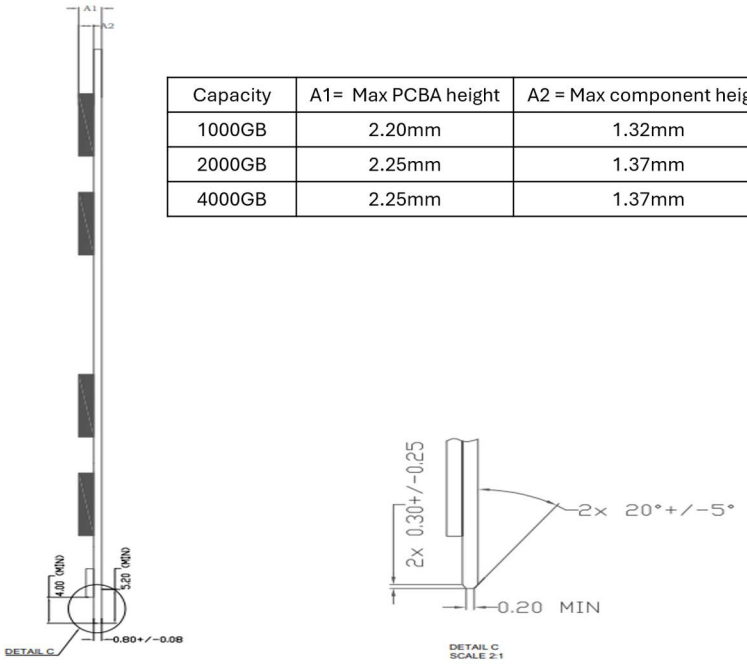

Capacity	A1= Max PCBA height	A2 = Max component height
1000GB	2.20mm	1.32mm
2000GB	2.25mm	1.37mm
4000GB	2.25mm	1.37mm

Figure 3 FireCuda X1070 SSD M.2 2280-S-M bottom view

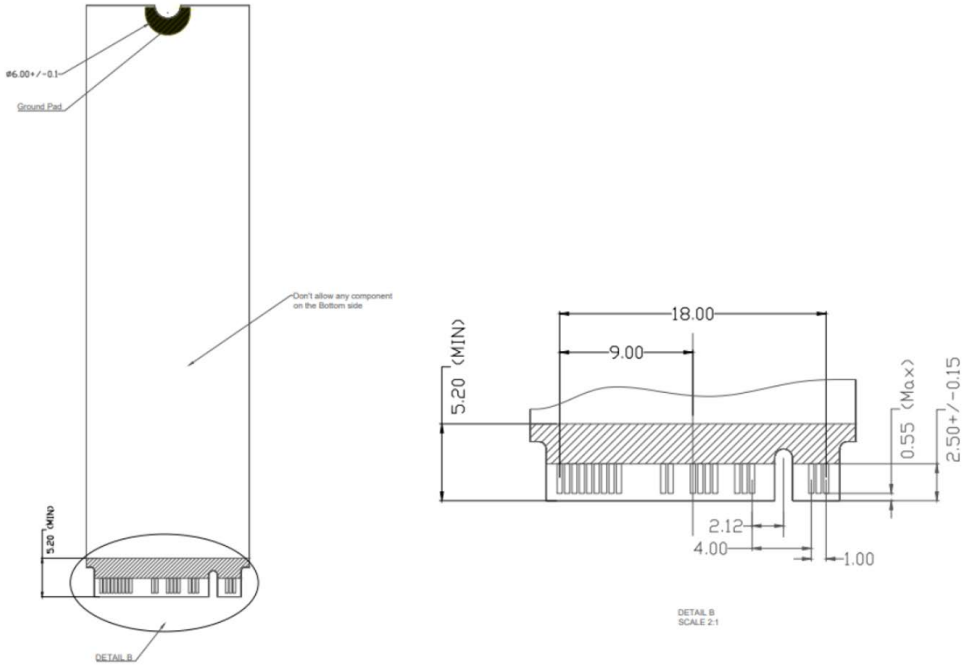

4. Pin and signal descriptions

Table 8 Pin and signal descriptions

Pin No.	Signal Name	Function	Pin No.	Signal Name	Function
1	GND	Ground	2	3.3V	3.3 V Source
3	GND	Ground	4	3.3V	3.3 V Source
5	PETn3	PCIe TX	6	PWRDIS	Power Disable
7	PETp3	PCIe TX	8	PLN#	Power Loss Notification
9	GND	Ground	10	LED_1#	Device Active Signal
11	PERn3	PCIe RX	12	3.3V	3.3 V Source
13	PERp3	PCIe RX	14	3.3V	3.3 V Source
15	GND	Ground	16	3.3V	3.3 V Source
17	PETn2	PCIe TX	18	3.3V	3.3 V Source
19	PETp2	PCIe TX	20	N/C	N/C
21	GND	Ground	22	VIO_1.8V	I/O Source
23	PERn2	PCIe RX	24	N/C	N/C
25	PERp2	PCIe RX	26	N/C	N/C
27	GND	Ground	28	N/C	N/C
29	PETn1	PCIe TX	30	PLA_S3#	Power Loss Acknowledge
31	PETp1	PCIe TX	32	N/C	N/C
33	GND	Ground	34	N/C	N/C
35	PERn1	PCIe RX	36	N/C	N/C
37	PERp1	PCIe RX	38	N/C	N/C
39	GND	Ground	40	SMB_CLK	SMBus Clock
41	PETn0	PCIe TX	42	SMB_DATA	SMBus Data
43	PETp0	PCIe TX	44	ALERT#	Alert Notification
45	GND	Ground	46	N/C	N/C
47	PERn0	PCIe RX	48	N/C	N/C
49	PERp0	PCIe RX	50	PERST#	PCIe Reset
51	GND	Ground	52	CLKREQ#	PCIe Device Clock Request
53	REFCLKN	PCIe Reference Clock	54	PEWAKE#	PCIe WAKE#
55	REFCLKP	PCIe Reference Clock	56	MFG_DATA	Reserved for MFG_DATA
57	GND	Ground	58	MFG_CLOCK	Reserved for MFG_CLOCK
59	N/C	Mechanical Notch	60	N/C	Mechanical Notch
61	N/C	Mechanical Notch	62	N/C	Mechanical Notch
63	N/C	Mechanical Notch	64	N/C	Mechanical Notch
65	N/C	Mechanical Notch	66	N/C	Mechanical Notch
67	N/C	N/C	68	SUSCLK	32.768 kHz Clock Supply
69	N/C	PEDET (NC-PCIe)	70	3.3V	3.3 V Source
71	GND	Ground	72	3.3V	3.3 V Source
73	VIO_CFG	IO Voltage configuration	74	3.3V	3.3 V Source
75	GND	Ground			

5. NVMe commands

In the tables below, O/M indicates whether the command is Mandatory (M) or Optional (O) according to the NVMe specification. Mandatory commands must be supported by all NVMe controllers, while optional commands may be implemented depending on the device's design or feature set.

Table 9 Admin commands

Identifier	O/M	Command Description	Supported
00h	M	Delete I/O Submission Queue	Supported
01h	M	Create I/O Submission Queue	Supported
02h	M	Get Log Page	Supported
04h	M	Delete I/O Completion Queue	Supported
05h	M	Create I/O Completion Queue	Supported
06h	M	Identify	Supported
08h	M	Abort	Supported
09h	M	Set Feature	Supported
0Ah	M	Get Feature	Supported
0Ch	M	Asynchronous Event Request	Supported
10h	O	Firmware Commit	Supported
11h	O	Firmware Image Download	Supported
14h	O	Device Self-test	Supported
80h	O	Format NVM	Supported
81h	O	Security Send	Supported
82h	O	Security Receive	Supported
84h	O	Sanitize	Supported

Table 10 I/O commands

Identifier	O/M	Command Description	Supported
00h	O	Flush	Supported
01h	O	Write	Supported
02h	O	Read	Supported
04h	O	Write Uncorrectable	Supported
05h	O	Compare	Not Supported
08h	O	Write Zeroes	Not Supported
09h	O	Dataset Management	Supported

Table 11 Set Feature commands

Identifier	O/M	Command Description	Supported
00h		Reserved	
01h	M	Arbitration	Supported

Table 11 Set Feature commands *(continued)*

02h	M	Power Management	Supported
03h	O	LBA Range Type	Not Supported
04h	M	Temperature Threshold	Supported
05h	M	Error Recovery	Supported
06h	O	Volatile Write Cache	Supported
07h	M	Number Of Queues	Supported
08h	M	Interrupt Coalescing	Supported
09h	M	Interrupt Vector Configuration	Supported
0Ah	M	Write Atomicity Normal	Supported
0Bh	M	Asynchronous Event Configuration	Supported
0Ch	O	Autonomous Power State Transition	Supported
0Dh	O	Host Memory Buffer	Supported
0Eh	O	Timestamp	Supported
10h	O	Host Controlled Thermal Management	Not Supported
11h	O	Non-Operational Power State Config	Supported
0Eh - 7Dh		Reserved	
80h	O	Software Progress Marker	Supported

Table 12 Get Log Page commands

Identifier	O/M	Command Description	Supported
00h		Reserved	
01h	M	Error Information	Supported
02h	M	SMART / Health Information	Supported
03h	M	Firmware Slot Information	Supported
04h	O	Changed Namespace List	Not Supported
06h	O	Device Self-Test	Supported
09h - 7Fh		Reserved	
81h	O	Sanitize Status	Supported
82h - FFh		Reserved	

6. SMART support

FireCuda X1070 SSD supports the SMART command set.

6.1 SMART attributes

The following table lists SMART Attributes and Descriptions.

Table 13 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
[201:200]	2	Temperature Sensor 1 (Current Temperature)
[203:202]	2	Temperature Sensor 2 (N/A)
[205:204]	2	Temperature Sensor 3 (N/A)
[207:206]	2	Temperature Sensor 4 (N/A)
[209:208]	2	Temperature Sensor 5 (N/A)
[211:210]	2	Temperature Sensor 6 (N/A)
[213:212]	2	Temperature Sensor 7 (N/A)
[215:214]	2	Temperature Sensor 8 (N/A)
[511:216]	296	Reserved

NOTES

For Log Identifier 02h:

- "Critical Warning [Byte 0]"
- This field indicates critical warnings for the state of the controller.
  - Bit#0: Available spare is below threshold
  - Bit#1: Temperature is either above the over-temperature threshold or below the under-temperature threshold
  - Bit#2: Reliability is degraded due to excessive media or internal errors
  - Bit#3: Media is placed in read-only mode
  - BIt#4: Volatile memory backup device has failed
  - Bit#5 - Bit#7: Reserved
- "Available Spare [Byte 3]" Value (percentage) = 100* [(total reserved VB - consumed VB caused by early, later bad)/ total reserved VB]
- "Percentage Used [Byte 5]"
- Value (percentage) = 100* (total VB erase count/ PE cycle for total VB)

7. Feature details

7.1 Flash management

7.1.1 Error correction code (ECC)

Flash memory cells naturally wear out over time, which can lead to random bit errors in stored data. To maintain data integrity, FireCuda X1070 SSD uses a high efficiency 4KB-based LDPC (Low Density Parity Check) ECC algorithm, which can detect and correct errors that occur during the read process, ensuring accurate data retrieval and protecting against corruption.

7.1.2 Wear leveling

NAND flash devices can only sustain a limited number of program/erase cycles, and usage across the media is often uneven. When certain areas are updated more frequently than others, the device's lifespan can be significantly shortened. To prevent this, wear leveling is employed to distribute write and erase cycles evenly across the flash memory, extending its overall durability.

FireCuda X1070 SSD incorporates both dynamic and static wear leveling algorithms, which work together to significantly improve the drive's life expectancy.

7.1.3 Bad block management

Bad blocks are memory blocks that fail to function correctly or contain excessive invalid bits, making stored data unstable and unreliable. Blocks identified and marked as defective during manufacturing are known as "Early Bad Blocks", while those that develop over the device's lifespan are called "Later Bad Blocks". Seagate employs an advanced bad block management algorithm to detect factory-marked bad blocks and handle those that emerge during use. This approach ensures data is never written to defective blocks, enhancing overall reliability and data integrity.

7.1.4 TRIM

TRIM is a feature that improves the read/write performance of solid-state drives (SSDs). Unlike hard disk drives (HDD), SSDs cannot overwrite existing data directly, so available space gradually decreases as data is updated. The TRIM command allows the operating system to notify the SSD which blocks of data are no longer in use and can be removed permanently. By clearing unused blocks, the SSD prevents wasted storage space.

7.1.5 SMART

SMART (Self-Monitoring, Analysis, and Reporting Technology) is an open standard that enables hard drives to monitor their health and predict potential failures. When SMART detects a failure, users can proactively replace the drive to avoid unexpected downtime or data loss. Additionally, SMART provides early warnings of impending issues, giving users time to take preventive measures such as backing up data to another device.

7.1.6 Over-provisioning

Over-provisioning refers to the inclusion of extra NAND capacity in an 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.

Over-provisioning refers to allocating additional NAND capacity in an SSD that remains hidden and inaccessible to users. This extra space allows the controller to better manage program/erase (P/E) cycles, resulting in improved performance and increased IOPS (Input/Output Operations per Second). Over-provisioning enhances reliability and endurance by reducing wear on the flash memory. It also helps to lower write amplification, meaning the controller writes less data to the flash for each operation.

7.1.7 Firmware upgrade

Firmware is a set of instructions that governs how a device communicates with its host system. Upgrading the firmware ensures the device stays up to date by introducing new features, resolving compatibility issues, and improving read/write performance.

7.1.8 Thermal throttling

Thermal throttling is designed to prevent SSD components from overheating during read and write operations. The device incorporates an on-die thermal sensor to monitor temperature. Using this data, the firmware can proactively apply varying levels of throttling to maintain optimal protection and performance, leveraging SMART readings for accurate control.

Table 14 Thermal throttling

Item	Content
SMART reporting temperature	Flash normalized case temperature
Reference of temp. reading	Flash on-die thermal sensor and Controller on-die thermal sensor
tmt1 threshold	CTL Tj = 115°C or NAND = 80°C per SMART reported
tmt2 threshold	CTL Tj = 120°C or NAND = 85°C per SMART reported
Fatal threshold	CTL_Fatal = 128°C or NAND_Fatal = 92°C
Resume performance threshold	CTL Tj = 113°C or NAND = 78°C per SMART reported
Temperature polling frequency	About 0.8s
TMT1_state impact	APU add delay between both read and write
TMT2_state impact	APU add delay between both read and write

NOTE For optimal performance:

- Provide sufficient airflow and cooling

7.2 Advanced device security features

7.2.1 Secure Erase

Secure Erase is a standard NVMe format operation that writes "0xFF" across the entire SSD, effectively wiping all data. When executed, this command prompts the SSD controller to erase its storage blocks, restoring the drive to its factory default settings.

7.2.2 Sanitize operations

The Sanitize feature is an alternative to traditional secure erase methods via the Format NVM command. It provides robust data security by ensuring that user data is thoroughly wiped from the drive's media, caches, and Controller Memory Buffer. This is achieved through block erase operations, which overwrite or destroy the encryption key. The following table illustrates the types of Sanitize operations supported.

Table 15 Supported Sanitize operations

Drive Security Type	Sanitize Operation		
	Overwrite	Block Erase	Crypto Erase
Non-SED	No	Yes	No

7.3 SSD lifetime management

7.3.1 Media Scan

Media Scan is a technology designed to monitor the health status of NAND Flash, predict potential errors, and proactively replace weak blocks by transferring user data to healthy blocks while refreshing the weaker ones.

Media Scan operates through various modes, including background scan, read disturb scan, and power-up scan, each triggered by specific conditions. Despite their different triggers, all modes share the same purpose: ensuring the reliability and longevity of the NAND Flash. Media Scan systematically scans all blocks, evaluates their health status based on predefined rules and thresholds, and moves blocks meeting the criteria into a refresh queue for necessary maintenance.

7.3.2 Total Bytes Written (TBW)

Total Bytes Written (TBW) is a key metric used to estimate the lifespan of a Solid State Drive (SSD). It indicates the total amount of data that can be written to the SSD over its lifetime. To calculate TBW, the following equation is applied:

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

NAND Endurance: Refers to the P/E (Program/Erase) cycle of the NAND flash memory. It measures how many times the NAND cells can be written and erased before they wear out.

SSD Capacity: Total capacity of the SSD.

WAF: A ratio that compares the amount of data written by the SSD controller to the amount of data written by the host system. A lower WAF, ideally close to 1, indicates better efficiency, as it means less data is being written to the flash memory than is necessary, thereby enhancing the SSD's endurance.

7.3.3 Media Wear Indicator

The Media Wear Indicator tracks the drive's actual life through the SMART Attribute byte index [5]. It reports the percentage of usage, recommending the user to replace the drive once the indicator reaches 100%. This ensures optimal performance and prevents unexpected failures.

7.3.4 Read-Only Mode (End of Life)

As the drive ages due to accumulated program/erase cycles, media wear-out can lead to an increasing number of bad blocks. When the number of available spare blocks drops below the threshold (5%, as indicated by SMART attribute log ID 02h Byte4), the drive will notify the host system through an AER event and issue a Critical Warning. At this point, the drive enters Read-Only Mode to prevent further data corruption. Users are strongly advised to replace the drive immediately to maintain data integrity and system reliability.

7.4 An adaptive approach to performance tuning

7.4.1 Predict & Fetch

Predict & Fetch is a feature in PCIe SSDs that boosts read speeds and reduces latency by anticipating sequential read commands. Instead of responding to each command individually, the SSD predicts upcoming read operations and prepares data in advance, speeding up processing time.

7.4.2 SLC caching

The device's firmware uses dynamic caching to boost performance and endurance, enhancing the user experience. SLC caching can utilize up to 1/4 of the SSD's free capacity, optimizing storage efficiency.

7.4.3 Host Memory Buffer (HMB) support

FireCuda X1070 SSD supports Host Memory Buffer (HMB), which uses a small portion of system memory to improve performance, especially for random read operations over large data ranges. HMB is read-only and does not store user data.

8. Safety, standards, and compliance

Each SSD has a product label that includes certifications that apply to that specific drive. You can find up-to-date information on safety certifications, and component compliance requirements for Seagate devices on the Seagate Support page: www.seagate.com/support. Scroll down to the Compliance section to access all compliance resources.

8.1 Regulatory model numbers

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

- STA055

9. FireCuda X1070 SSD installation precautions

9.1 FireCuda X1070 SSD handling

An SSD device consists of numerous intricate components assembled together, requiring careful handling, especially if it includes WLCSP (Wafer Level Chip Scale Packaging) components such as PMIC, thermal sensors, or load switches. WLCSP is a packaging technology designed to achieve smaller footprints, making it ideal for compact designs. However, these ultra-small parts are highly sensitive to physical damage, such as bumps or scratches, which can compromise their functionality. It is essential to handle the drive with utmost care to ensure its reliability and performance. See Section 9.2, *FireCuda X1070 SSD installation*.

CAUTION! DO NOT DROP SSD

INSTALL SSD WITH CARE

STORE SSD IN A PROPER PACKAGE

During operations, FireCuda X1070 SSD may reach high temperatures.

CAUTION! Hot surface. Do not touch.

ATTENTION! Surface chaude. Ne pas toucher.

9.2 FireCuda X1070 SSD installation

FireCuda X1070 SSD features the PCIe connector with M-key, which is compatible only with the M-key socket. See Use Case 1 in Figure 4, *M-Key M.2 assembly precautions*. As shown in Use Case 2, misuse may cause severe damage to an SSD, including burnout.

Figure 4 M-Key M.2 assembly precautions

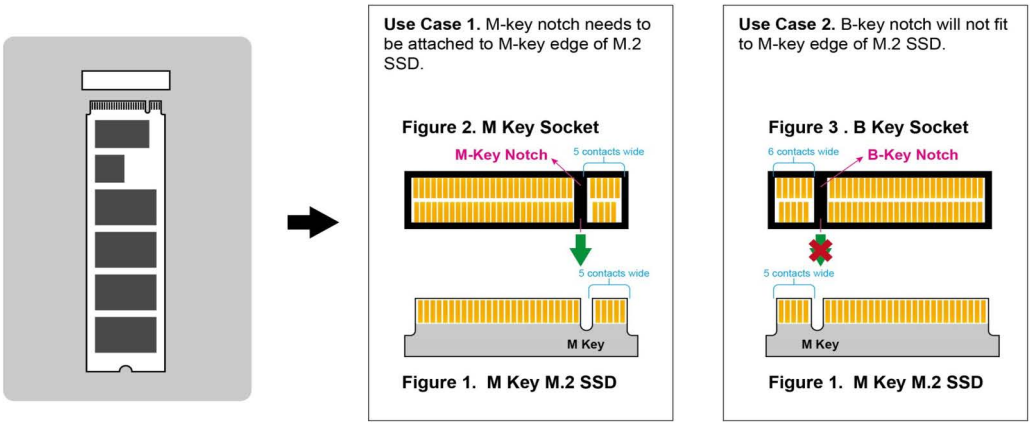

Seagate Technology LLC
AMERICAS Seagate Technology LLC 47488 Kato Road, Fremont, California 94538, United States, 510-661-1000
ASIA/PACIFIC Seagate Singapore International Headquarters Pte. Ltd. 90 Woodlands Avenue 7, Singapore 737911, 65-6412-3666
EUROPE, MIDDLE EAST AND AFRICA Seagate Technology (Netherlands) B.V. Koolhovenlaan 1, 1119 NB Schiphol-Rijk, Netherlands, 31-20-316-7300
Publication Number: 210892300, Rev A
February 2026