



Intel® Solid State Drive Data Center for PCIe SMART Attributes

Application Note

August 2020



Revision History

Revision	Description	Date
001	Initial Release	February 2017
002	Revised for Minor Errors	August 2020

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

© Intel Corporation, Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.



Overview

SMART (Self-Monitoring, Analysis and Reporting Technology) is an open standard used by drives and hosts to monitor drive health and report potential problems. This document lists and describes the SMART attributes supported by Intel® Data Center Solid State Drives (SSDs) for PCIe.

Intel® Data Center Tool

Intel's Data Center Tool (DCT) can be used to read out the SMART attributes. The read out would include normalized, raw, status and threshold values. The following commands can be passed using the tool to read out the SMART attributes:

- Windows:


```
isdct.exe show -destination <filepath> -intelssd <device index> -smart
```
- Linux:


```
./issdct show -destination <filepath> -intelssd <device index> -smart
```

For more information, see the latest [Intel Data Center Tool User Guide](#).

SMART Attributes

The following table lists the SMART attributes supported by the Intel Data Center SSDs for PCIe.

SMART Attributes (LoByte)	# of Bytes	Attribute	Description
0	1	<p>Critical Warning: These bits if set, flag various warning sources.</p> <p>Bit 0: Available Spare is below Threshold</p> <p>Bit 1: Temperature has exceeded Threshold</p> <p>Bit 2: Reliability is degraded due to excessive media or internal errors</p> <p>Bit 3: Media is placed in Read- Only Mode</p> <p>Bit 4: Volatile Memory Backup System has failed (e.g., enhanced power loss capacitor test failure)</p> <p>Bits 5-7: Reserved</p>	Any of the critical warning can be tied to asynchronous event notification. Drive Health Indicator defined under bytes 3095-3076 of Identify Controller may still indicate "healthy" status when the critical warning flag is set.
1	2	<p>Temperature:</p> <p>Overall Device current temperature in Kelvin.</p>	<p>This reports media temperature.</p> <p>For AIC, it reports the NAND temperature, for 2.5" FF, it is the case temperature</p>
3	1	<p>Available Spare:</p> <p>Contains a normalized percentage (0 to 100%) of the remaining spare capacity available</p>	This field is not applicable to Intel SSDs' media and architecture. Available Spare will be set at 100% and decrements
4	1	Available Spare Threshold	This field is not applicable to Intel SSDs' media and architecture. Available Spare Threshold will be set at 0%



SMART Attributes (LoByte)	# of Bytes	Attribute	Description
5	1	Percentage Used Estimate (Value allowed to exceed 100%)	A value of 100 indicates that the estimated endurance of the device has been consumed but may not indicate a device failure. The value is allowed to exceed 100. Percentages greater than 254 shall be represented as 255. This value shall be updated once per power-on hour (when the controller is not in a sleep state). If the value reaches or exceeds 105, drive will enter a write protect mode with write bandwidth reaching <10MB/sec
32	16	Data Units Read (in LBAs)	Contains the number of 512 byte data units the host has read from the controller; this value does not include metadata. This value is reported in thousands (i.e., a value of 1 corresponds to 1000 units of 512 bytes read) and is rounded up. When the LBA size is a value other than 512 bytes, the controller shall convert the amount of data read to 512 byte units.
48	16	Data Units Write (in LBAs)	Contains the number of 512 byte data units the host has written to the controller; this value does not include metadata. This value is reported in thousands (i.e., a value of 1 corresponds to 1000 units of 512 bytes written) and is rounded up. When the LBA size is a value other than 512 bytes, the controller shall convert the amount of data written to 512 byte units. For the NVM command set, logical blocks written as part of Write operations shall be included in this value. Write Uncorrectable commands shall not impact this value
64	16	Host Read Commands	Contains the number of read commands issued to the controller.
80	16	Host Write Commands	Contains the number of write commands issued to the controller.
96	16	Controller Busy Time (in minutes)	Contains the amount of time the controller is busy with I/O commands. The controller is busy when there is a command outstanding to an I/O Queue (specifically, a command was issued by way of an I/O Submission Queue Tail doorbell write and the corresponding completion queue entry has not been posted yet to the associated I/O Completion Queue). This value is reported in minutes.
112	16	Power Cycles	Contains the number of power cycles.
128	16	Power On Hours	Contains the number of power-on hours. This does not include time that the controller was powered and in a low power state condition.



SMART Attributes (LoByte)	# of Bytes	Attribute	Description
144	16	Unsafe shutdowns	Contains the number of unsafe shutdowns. This count is incremented when a shutdown notification (CC.SHN) is not received prior to loss of power.
160	16	Media Errors	Contains the number of occurrences where the controller detected an unrecovered data integrity error. Errors such as uncorrectable ECC, CRC checksum failure, or LBA tag mismatch are included in this field.
176	16	Number of Error Information Log Entries	Contains the number of Error Information log entries over the life of the controller.
192	4	Warning Composite Temperature Time	Contains the amount of time in minutes that the controller is operational, and the Composite Temperature is greater than or equal to the Warning Composite Temperature Threshold (WCTEMP) field and less than the Critical Composite Temperature Threshold (CCTEMP) field in the Identify Controller data structure. (P3100)
196	4	Critical Composite Temperature Time	Contains the amount of time in minutes that the controller is operational, and the Composite Temperature is greater than the Critical Composite Temperature Threshold (CCTEMP) field in the identify Controller data structure.



Additional SMART Attributes (Log Identifier CAh)

Byte	# of Bytes	Attribute	Description
0	1	AB (Program Fail Count)	These fields are not applicable to Intel SSDs' media and architecture and will not be populated. Raw value: shows total count of program fails. Normalized value: beginning at 100, shows the percent remaining of allowable program fails.
3	1	Normalized Value	
5	6	Current Raw Value	
12	1	AC (Erase Fail Count)	These fields are not applicable to Intel SSDs' media and architecture and will not be populated. Raw value: shows total count of erase fails. Normalized value: beginning at 100, shows the percent remaining of allowable erase fails.
15	1	Normalized Value	
17	6	Current Raw Value	
24	1	AD (Wear Leveling Count)	Raw value: Bytes 1-0: Min. erase cycle Bytes 3-2: Max. erase cycle Bytes 5-4: Avg. erase cycles Normalized value: decrements from 100 to 0.
27	1	Normalized Value	
29	6	Current Raw Value	
36	1	B8 (End to End Error Detection Count)	Raw value: reports number of End-to-End detected and corrected errors by hardware. Normalized value: always 100.
39	1	Normalized Value	
41	6	Current Raw Value	
48	1	C7 (CRC Error Count)	Raw value: total number of PCIe Interface CRC errors encountered, as specified in PCIe Link Performance Counter Parameter for "Bad TLP". Normalized value: always 100.
51	1	Normalized Value	
53	6	Current Raw Value	
60	1	E2 (Timed Workload, Media Wear)	These fields are not applicable to Intel SSDs' media and architecture and will not be populated. Raw value: measures the wear seen by the SSD (since reset of the workload timer, attribute E4h), as a percentage of the maximum rated cycles. Divide the raw value by 1024 to derive the percentage with 3 decimal points. Normalized value: always 100.
63	1	Normalized Value	
65	6	Current Raw Value	
72	1	E3 (Timed Workload, Host Reads %)	These fields are not applicable to Intel SSDs' media and architecture and will not be populated. Raw value: shows the percentage of I/O operations that are a read operation (since reset of the workload timer, attribute E4h). Reported as integer percentage from 0 to 100.
75	1	Normalized Value	
77	6	Current Raw Value	
84	1	E4 (Timed Workload, Timer)	These fields are not applicable to Intel SSDs' media and architecture and will not be populated. Raw value: measures the elapsed time (number of minutes since starting this workload timer). Normalized value: always 100.
87	1	Normalized Value	
89	6	Current Raw Value	



Byte	# of Bytes	Attribute	Description
96	1	EA (Thermal Throttle Status)	Raw value: reports Percent Throttle Status and Count of events
99	1	Normalized Value	
101	6	Current Raw Value	Byte 0: Throttle status reported as integer percentage. Bytes 1-4: Throttling event count. Number of times thermal throttle has activated. Preserved over power cycles. Byte 5: Reserved. Normalized value: always 100.
108	1	F0 (Retry Buffer Overflow Counter)	Raw Value: Counter to indicate the number of times Retry Buffer has overflown Normalized Value is always 100
111	1	Normalized Value	
113	6	Current Raw Value	
120	1	F3 (PLL Lock Loss Count)	Raw Value: Counter to indicate the number of times PCIe Refclock PLL has unlocked Normalized Value is always 100
123	1	Normalized Value	
125	6	Current Raw Value	
132	1	F4 (NAND Bytes Written)	Nand sectors written divided by 65536 (1 count = 32 MiB) Normalized value always 100
135	1	Normalized Value	
137	6	Current Value	
144	1	F5 (Host Bytes Written)	Host sectors written divided by 65536 (1 count = 32 MiB) Normalized value always 100
147	6	Normalized Value	
149	6	Current Value	
152	1	F6 (System Area Used)	Tracks the remaining available format changes and crypto erases. Raw value: Percentage from 0-100% starting at 0 and ending at 100% Normalized value always 100
155	1	Normalized Value	
162	6	Current Value	



Temperature Statistics

Byte	# of Bytes	Log Page Content
0-7	8	Current Internal Temperature in Celsius
0	1	Current Temperature (P4501/P4601)
0	8	Current Temperature
8-15	8	Overtemp shutdown Flag for last Drive Overheat
16-23	8	Overtemp Shutdown Flag for Temp Drive Overheat
24-31	8	Highest (Lifetime) Temperature in Celsius
32-39	8	Lowest (Lifetime) Temp in Celsius
40-79	40	Reserved
80-87	8	Specified PCB Maximum Operating Temperature in Celsius
88-95	8	Reserved
96-103	8	Specified PCB Minimum Operating Temperature in Celsius
104-111	8	Estimated Offset in Celsius
111-511	400	Reserved