How to Implement SMART Embedded for SATA & PCIe NVMe SSD?

This application note provides instructions to use SP SMART Embedded utility program to integrate with customer's program to get SMART information for SP Industrial SATA & PCIe NVMe SSD.

Support Environment

- OS : Windows 10 and Linux
- SP SMART Embedded utility program : smartctl 7.2
- Host : Intel x86 Platform

Support List for SP Industrial SSD

- SATA SSD & Cfast (MLC) : SSD700/500/300, MSA500/300, MDC500/300, CFX510/310
- SATA SSD & CFast (3D TLC) : SSD550/350/3K0, MSA550/350/3K0, MDC550/350, MDB550/350, MDA550/350/3K0 series, CFX550/350
- PCIe NVMe : MEC350, MEC3F0, MEC3K0 series

SMART Attribute

- SATA SSD & Cfast (MLC)

Attribute	SM2246EN	SM2246XT
	SSD700/500/300R/S series MSA500/300S MDC500/300 R/S series	CFX510/310
01	Read error rate CRC Error count	Read error rate CRC Error count
05	Reallocated sectors count	Reallocated sectors count
09	Power-on hours	Reserved
0C	Power cycle count	Power cycle count
A0	Uncorrectable sector count when read/Write	Uncorrectable sector count when read/Write
A1	Number of valid spare block	Number of valid spare block
A2		Number of valid spare block
A3	Number of initial invalid block	Number of initial invalid block
A4	Total erase count	Total erase count
A5	Maximum erase count	Maximum erase count
A6	Minimum erase count	Average erase count
A7	Max erase count of spec	
A8	Remain Life	

Attribute	SM2246EN SSD700/500/300R/S series MSA500/300S MDC500/300 R/S series	SM2246XT CFX510/310
A9	Remain Life	
AF	Program fail count in worst die	
B0	Erase fail count in worst die	
B1	Total wear level count	
B2	Runtime invalid block count	
B5	Total program fail count	
B6	Total erase fail count	
BB	Uncorrectable error count	
C0	Power-off retract count	Power-off retract count
C2	Controlled temperature	Controlled temperature
C3	Hardware ECC recovered	Hardware ECC recovered
C4	Reallocated event count	Reallocated event count
C6	Uncorrectable error count off-line	
C7	Ultra DMA CRC error count	Ultra DMA CRC error count
E1	Total LBAs written	
E8	Available reserved space	
F1	Write Sector Count Total LBAs Written (each write unit = 32MB)	Total LBAs written
F2	Read Sector Count Total LBAs Read (each read unit = 32MB)	Total LBAs read

• SATA SSD & Cfast (3D TLC)

Attribute	SM2258H	SM2258XT	RL5735
	SSD550/350 R/S series MSA550/350 S series MDC550/350 R/S series MDB550/350 S series MDA550/350 S series CFX550/350 S series	CFX550/350 series	SSD3K0E, MSA3K0E, MDA3K0E series
01	TRead error rate (CRC Error count)	TRead error rate (CRC Error count)	TRead error rate (CRC Error count)
05	Reallocated sectors count	Reallocated sectors count	Reallocated sectors count
09	Power-on hours	Power-On Hours Count	Power-On Hours Count
0C	Power cycle count	Power cycle count	Power cycle count
94	Total erase count (SLC) (pSLC model)		
95	Maximum erase count (SLC) (pSLC model)		
96	Minimum erase count (SLC) (pSLC model)		
97	Average erase count (SLC) (pSLC model)		
A0	Uncorrectable Sector Count On Line (Uncorrectable sector count when read/Write)	Online Uncorrect Sector Count (Uncorrectable sector count when read/Write)	
A1	Number of Pure Spare (Number of valid spare block)	Number of valid spare block	Grow defect number (Later bad block)
A2			Total erase count
A3	Number of initial invalid block	Number of initial invalid block	Max PE cycle Spec
A4	Total erase count (TLC)	Total Erase Count (TLC)	Average erase count
A5	Maximum erase count (TLC)	Maximum erase count (TLC)	
A6	Minimum erase count (TLC)	Minimum erase count (TLC)	Total bad block count
A7	Average erase count (TLC)	Average erase count (TLC)	SSD protect mode
A8	Max Erase Count in Spec (Max erase count of spec)	Max Erase Count in Spec	SATA Phy error count
A9	Remaining Life Percentage	Remaining Life Percentage	Remaining Life Percentage
AB			Program fail count
AC			Erase fail count
AE			Unexpected power loss count
AF			ECC fail count (host read fail)

Attribute	SM2258H SSD550/350 R/S series MSA550/350 S series MDC550/350 R/S series MDB550/350 S series MDA550/350 S series CFX550/350 S series	SM2258XT CFX550/350 series	RL5735 SSD3K0E, MSA3K0E, MDA3K0E series
B1	Total wear level count	Wear leveling Count	
B2	Used Reserved Block Count (Runtime invalid block count)	Grown Bad Block Count	
B5	Total program fail count	Program Fail Count	Unaligned access count
B6	Total erase fail count	Erase Fail Count	
BB	Uncorrectable error count		Reported uncorrectable error
C0	Power-off retract count	Sudden Power Count (Power-off retract count)	
C2	Temperature_Celsius (Tjunction)	Enclosure Temperature (Tjunction)	Enclosure temperature (Tjunction)
C3	Hardware ECC recovered	Hardware ECC recovered	Cumulative corrected ecc
C4	Reallocated event count	Reallocated event count	Reallocation event count
C5	Current pending sector count:	Current Pending Sector Count	
C6	Uncorrectable error count off-line	Reported Uncorrectable Errors	
C7	UDMA CRC Error (Ultra DMA CRC error count)	CRC Error Count (Ultra DMA CRC error count)	Ultra DMA CRC error count
CE			Min. erase count
CF			Max erase count
E1	Host Writes (Total LBAs written)		
E8	Available reserved space	Max Erase Count in Spec	Available reserved space
E9	Total write to flash		Spare block
EA	Total Read from flash		
F1	Write Sector Count (Total Host Writes , each unit 32MB)	Host 32MB/unit Written (TLC)	Write life time
F2	Read Sector Count (Total Host Read , each unit 32MB)	Host 32MB/unit Read (TLC)	Read life time
F5	Flash Write count	NAND 32MB/unit Written (TLC)	Unexpected power loss count
F9			Total GB written to NAND (TLC)
FA			Total GB written to NAND (SLC)

- PCIe NVMe SSD (NVMe 1.3)

# of Bytes	Byte Index	Attributes	Description
1	0	Critical Warning: Bit Definition 00: If set to '1', then the available spare space has fallen below the threshold. 01: If set to '1', then a temperature is above an over temperature threshold or below an under temperature threshold. 02: If set to '1', then the NVM subsystem reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability. 03: If set to '1', then the media has been placed in read only mode. 04: If set to '1', then the volatile memory backup device has failed. This field is only valid if the controller has a volatile memory backup solution. 07:05: Reserved	This field indicates critical warnings for the state of the controller. Each bit corresponds to a critical warning type; multiple bits may be set. If a bit is cleared to '0', then that critical warning does not apply. Critical warnings may result in an asynchronous event notification to the host. Bits in this field represent the current associated state and are not persistent When the Available Spare falls below the threshold indicated in this field, an asynchronous event completion may occur. The value is indicated as a normalized percentage (0 to 100%).
2	2:1	Composite Temperature:	Contains a value corresponding to a temperature in degrees Kelvin that represents the current composite temperature of the controller and namespace(s) associated with that controller. The manner in which this value is computed is implementation specific and may not represent the actual temperature of any physical point in the NVM subsystem. The value of this field may be used to trigger an asynchronous event. Warning and critical overheating composite temperature threshold values are reported by the WCTEMP and CCTEMP fields in the Identify Controller data structure.
1	3	Available Spare:	Contains a normalized percentage (0 to 100%) of the remaining spare capacity available
1	4	Available Spare Threshold:	When the Available Spare falls below the threshold indicated in this field, an asynchronous event completion may occur. The value is indicated as a normalized percentage (0 to 100%).
1	5	Percentage Used:	Contains a vendor specific estimate of the percentage of NVM subsystem life used based on the actual usage and the manufacturer's prediction of NVM life. A value of 100 indicates that the estimated endurance of the NVM in the NVM subsystem has been consumed, but may not indicate an NVM subsystem 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). Refer to the JEDEC JESD218A standard for SSD device life and endurance measurement techniques
	31:6	Data Units Written:	
16	47:32	Data Units Read:	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. For the NVM command set, logical blocks read as part of Compare and Read operations shall be included in this value.

# of Bytes	Byte Index	Attributes	Description
16	63:48	Data Units Written:	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.
16	79:64	Host Read Commands:	Contains the number of read commands completed by the controller. For the NVM command set, this is the number of Compare and Read commands.
16	95:80	Host Write Commands:	Contains the number of write commands completed by the controller. For the NVM command set, this is the number of Write commands.
16	111:96	Controller Busy Time:	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 via 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.
16	127:112	Power Cycles:Contains the number of power cycles.	
16	143:128	Power On Hours:	Contains the number of power-on hours. Power on hours is always logging, even when in low power mode.
16	159:144	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.
16	175:160	Media and Data Integrity 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.
16	191:176	Number of Error Information Log Entries:	Contains the number of Error Information log entries over the life of the controller.
4	195:192	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. If the value of the WCTEMP or CCTEMP field is 0h, then this field is always cleared to 0h regardless of the Composite Temperature value.
4	199:196	Critical Composite Temperature Time:	Contains the amount of time in minutes that the controller is operational and the Composite Temperature is greater the Critical Composite Temperature Threshold (CCTEMP) field in the Identify Controller data structure. If the value of the CCTEMP field is 0h, then this field is always cleared to 0h regardless of the Composite Temperature value.
2	201:200	Reserved	
2	203:202	Reserved	
2	205:204	Reserved	
2	207:206	Reserved	
2	209:208	Reserved	
2	211:210	Reserved	
2	213:212	Reserved	
2	215:214	Reserved	
296	511:216	Reserved	

Installation

- Please download the latest version of SMART Embedded utility program. (Download link by request)
- Unzip (In this case, unzip to E:\smartmontools-7.2.win32 folder)
- Run Command Prompt
- Run as Administrator
- C:\WINDOWS\system32> E:\smartmontools-7.2.win32\bin\smartctl.exe -h
- To get a usage summary

Command line tool to get SMART information (sdb : disk on PhysicalDrive 1)

- C:\WINDOWS\system32> E:\smartmontools-7.2.win32\bin\smartct.exe -a /dev/sdb
- Check the attached file SMART.TXT : https://www.silicon-power.com/support/lang/utf8/smart.txt

Output SMART information into JSON format. (sdb : disk on PhysicalDrive 1)

- C:\WINDOWS\system32> E:\smartmontools-7.2.win32\bin\smartctl.exe -a -j /dev/sdb
- Check the attached file JSON.TXT : https://www.silicon-power.com/support/lang/utf8/json.txt

Used Case 1: Remote monitoring SMART Dashboard via IBM Node-Red

- Install IBM Node Red, Node Red is a flow-based programming tool developed by IBM. We use Node Red to integrate SP SMART Embedded utility program to develop a remote monitoring tool " SP SMART Dashboard".
- Develop Script for Node Red and using " smartctl.exe"
- Script file as the attached SMARTDASHBOARD.TXT : https://www.silicon-power.com/support/lang/utf8/SMARTDASHBOARD.txt
- Open Browser, input "ip:1880/ui"
- ip is the IP address of machine which is running Node Red script. Defaul ip of local machine is 127.0.0.1

Figure 1 SMART Dashboard

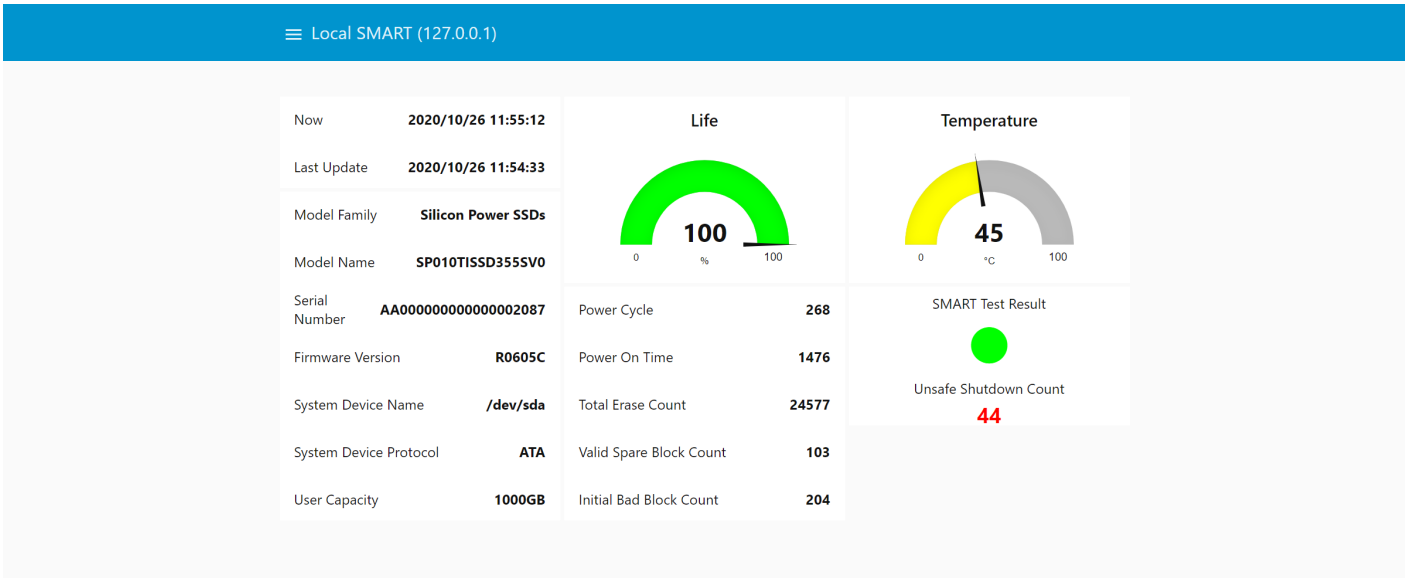

* Used case 2: Integration with Google Cloud Platform to manage SMART information of connected devices in the field

SP Industrial leverages Google Cloud Platform and SP SMART Embedded to develop a SMART IoT Sphere service platform. SP SMART IoT Sphere is a cloud-based service with alarm and maintenance notifications that monitors and analyzes the health and status of SP Industrial SSDs and Flash cards inside connected devices running Windows OS or Linux Ubuntu embedded OS.

Figure 2 Architecture of SMART IoT Sphere

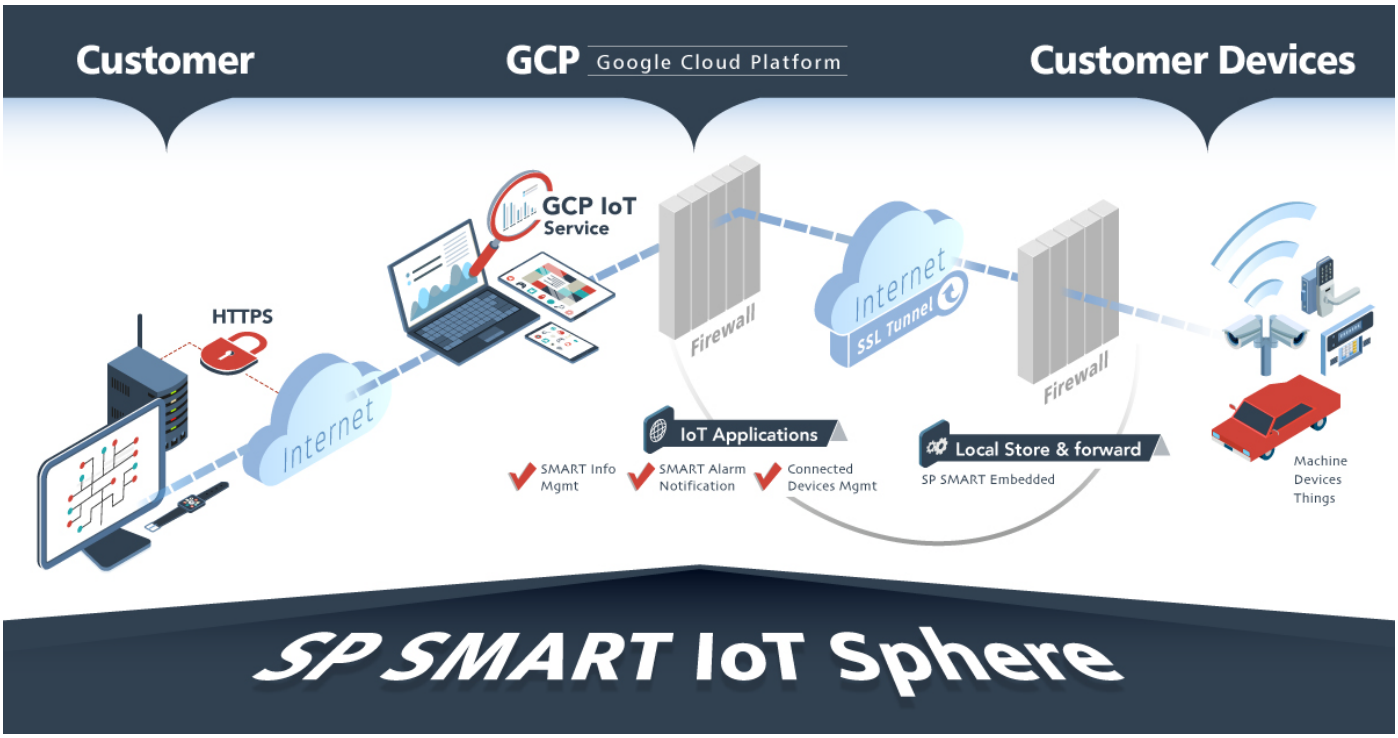

Figure 3 Multiple Devices management

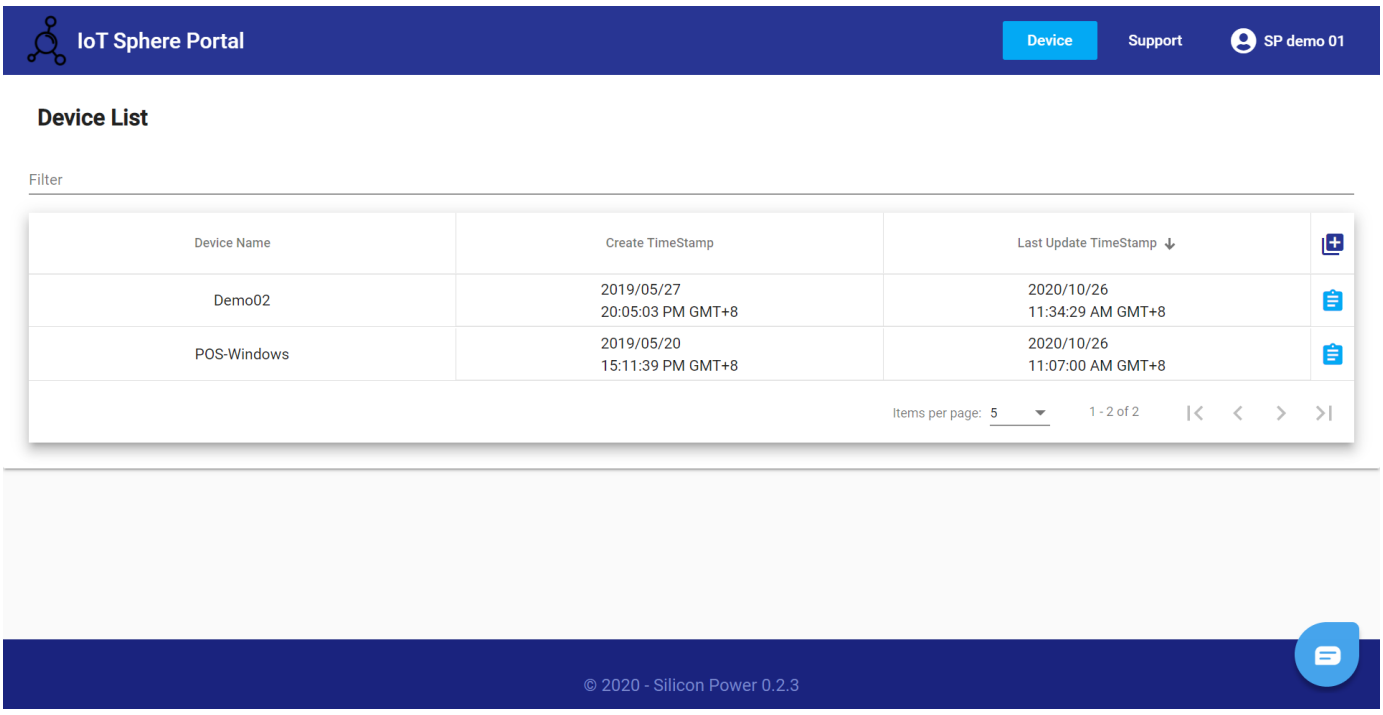

Figure 4 SP SMART Embedded supports both Windows 10 and Linux OS

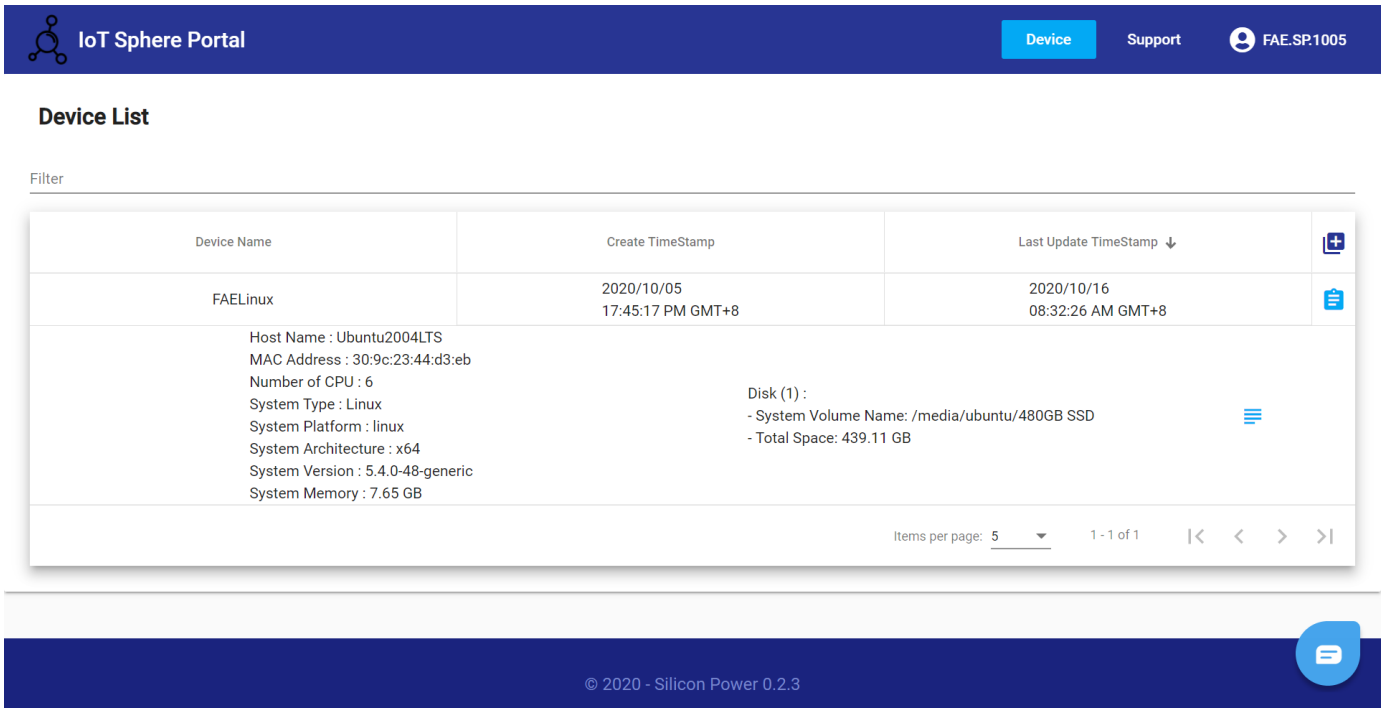

Figure 5 Realtime SMART Information display