SQFlash Intelligent SSD Self-Management

Author: Ares.Cheng

Keywords: SATA, SSD, S.M.A.R.T.

To meet the increasing demand for higher data volumes, the unit storage density of NAND flash has been continuously improved through substantial advancements in process technology that range from SLC to 3D NAND TLC. Yet, the downside of this improvement is reduced P/E cycles for high-density NAND flash and increased risk of data corruption. To facilitate monitoring of storage device status, the JEDEC Association established relevant Self-Monitoring, Analysis, and Reporting Technology (S.M.A.R.T.) specifications for SATA storage devices. However, because such instructions typically necessitate system technology-related knowledge, most users experience difficulty obtaining and utilizing such information, as well as performing adequate risk assessments and remote monitoring in accordance with S.M.A.R.T. instructions.

To reduce the risk of data corruption/loss and ensure data access, Advantech developed its SQFlash Utility and WISE-PaaS/DeviceOn platform to help customers implement S.M.A.R.T. specifications, predict product lifespans, and remotely monitor system status in real-time using a simple graphical user interface. This paper provides a detailed description of the intelligent SSD self-management system developed based on Advantech SQFlash.

Table of Content

Intelligent SSD Self-Management

SSD Information and Endurance Check

S.M.A.R.T. information in the SATA specification contains up to 256 information items. Some information is specific to traditional HDDs or SSD-related information, while some is considered manufacturer-defined or undefined information. To help users quickly understand the SSD status and information, Advantech's SQFlash Utility provides data for the 10 key parameters listed in the Information column, as shown in Figure 1. The data is obtained by decoding and calculating S.M.A.R.T. attributes. The remaining SSD endurance is presented in percentage using a pie chart. With this information, users can quickly understand the SSD status and establish contingency measures accordingly.

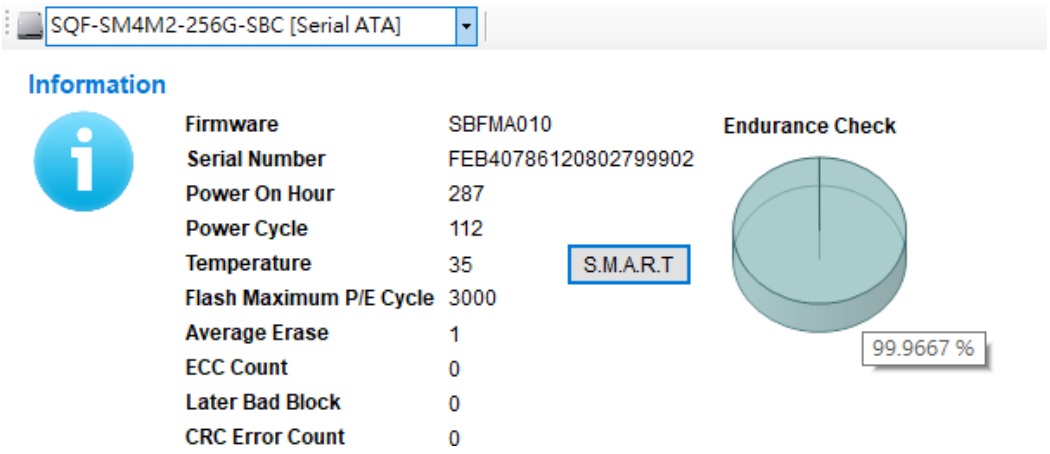

Figure 1. SSD Information and Endurance Check

Additionally, the provision of easy access to S.M.A.R.T data via the SQFlash Utility allows users to obtain more advanced information regarding the SSD status. As shown in Figure 2, many self-defined items are included in the SQFlash S.M.A.R.T. data, such as 10h, 11h, 64h, A8h, AAh, ADh, AFh, DAh, E7h, EAh, and EBh.

S.M.A.R.T

ID	Name	Raw	Value	Worst	Threshold
01	Raw_Read_Error_Rate	000000000000	100	100	50
09	Power_On_Hours	00000000011F	100	100	0
0C	Power_Cycle_Count	000000000070	100	100	0
0E	Device Capacity	00001DCF32B0	100	100	0
0F	User Capacity	00001DCF32B0	100	100	0
10	Initial Spare Blocks Available	000000000082	100	100	0
11	Spare Blocks Remaining	000000000073	100	100	0
64	Total Erase Count	000000001F44	100	100	0
A8	SATA PHY Error Count	000000000000	100	100	0
AA	Bad Block Count	0000000000E0	89	89	10
AD	Erase count	000000010003	100	100	0
AE	Unexpected Power Loss Count	000000000063	100	100	0
C0	Unexpected Power Loss Count	000000000063	100	100	0
AF	Power Failure Protection Status	000000000000	100	100	10
C2	Temperature	002600180023	65	62	0
CA	Percentage of Spares Remaining	000000000001	100	100	50
E7	SSD Life Remaining	000000000063	100	100	0
DA	CRC error	000000000000	100	100	50
EA	Total NAND Read	000044960A20	100	100	0
EB	Total NAND Written	000036BE6CE0	100	100	0
F1	Total Host Write	000071687C1B	100	100	0
F2	Total Host Read	000052613D91	100	100	0

Figure 2. SQFlash S.M.A.R.T. Data

ID	ATTRIBUTE_NAME	DATA ADDRESS (Byte)					
		10	9	8	7	6	5
01h	Raw_Read_Error_Rate	Uncorrectable ECC Count					
09h	Power_On_Hours	Power On Hours					
0Ch	Power_Cycle_Count	Power On/Off Counts					
0Eh	Device Capacity	0	0	Device Capacity			
0Fh	User Capacity	0	0	User Capacity			
10h	Initial Spare Blocks Available	0	0	Total Available Spare Block			
11h	Spare Blocks Remaining	0	0	Remaining Spare Block			
64h	Total Erase Count	0	0	Total Erase Count			
A8h	SATA PHY Error Count	SATA PHY Error Count					
AAh	Bad Block Count	Later Bad		0	0	Early Bad	
ADh	Erase Count	0	0	Avg. Erase		Max Erase	
AEh	Unexpected Power Loss Count	0	0	Unexpected Power Loss Count			
AFh	Power Failure Protection Status	Voltage Stabilizer Trigger Count				Guaranteed Flush	Drive Status (0x00 Normal)

ID	ATTRIBUTE_NAME	DATA ADDRESS (Byte)					
		10	9	8	7	6	5
						(0x01 Enable)	
C0h	Unexpected Power Loss Count	0	0	0	0	Unexpected Power Loss Count	
C2h	Temperature	Max. Temp.		Min. Temp.		Current Temp.	
CAh	Percentage of Spares Remaining	0	0	0	0	0	SSD Life Used
DAh	CRC Error	CRC Error Count					
E7h	SSD Life Remaining	0	0	0	0	0	SSD Life Left
EAh	Total NAND Read	Total NAND Read (Sector, 512B)					
EBh	Total NAND Written	Total NAND Written (Sector, 512B)					
F1h	Total Host Write	Host Write (Sector, 512B)					
F2h	Total Host Read	Host Read (Sector, 512B)					

Table 1. S.M.A.R.T. ID Definitions

ID	Description of the SQFlash-Defined S.M.A.R.T. IDs
10h	Initial Spare Blocks Available: The spare block count when the SSD is newly made.
11h	Spare Blocks Remaining: Currently available spare blocks.
64h	Total Erase Count: Sum of the erase count from all blocks.
A8h	SATA PHY Error Count: The PHY error count - such as data FIS CRC, code error, disparity error, command FIS CRC ... etc. This value resets to zero when powered off.
AAh	Bad Block Count: Block is a capacity unit of NAND flash. Bad blocks are damaged blocks that the SSD controller has marked as unusable. There are two types of bad blocks — early bad blocks and later bad blocks. Early bad blocks are the result of a manufacturing fault. Later bad blocks result from improper usage. If a block has too many ECCs, the SSD controller will mark the block as a later bad block. The erase count for the block will need to be checked. For SSDs nearing the end of their lifespan, later bad blocks are normal.
ADh	Erase Count: This value is directly related to the SSD lifespan. There are two types of erase counts — the average erase count and max. erase count. The average erase count is the average of the erase count for all blocks. The max. erase count is the highest erase count of all blocks. The SSD endurance is dependent on the NAND flash type, as shown in Table 2.

<table>
<tr><th>ID</th><th>Description of the SQFlash-Defined S.M.A.R.T. IDs</th></tr>
<tr><td></td><td>
<table>
<tr><th>NAND Flash Type</th><th>P/E Cycle (times)</th></tr>
<tr><td>SLC</td><td>100,000</td></tr>
<tr><td>UltraMLC</td><td>30,000</td></tr>
<tr><td>MLC</td><td>3,000</td></tr>
<tr><td>3D TLC (BiCS 3)</td><td>3,000</td></tr>
</table>
Table 2. P/E Cycle Times According to NAND Type
</td></tr>
<tr><td>AFh</td><td>Power Failure Protection Status: There are three power failure protection statuses.
Voltage Stabilizer Trigger Count: The number of times the voltage stabilizer circuit has been triggered.
Guaranteed Flush: This indicates the function is on/off status (0x01 Enable).
Drive Status: Error code of power failure protection functions (0x00 Normal).</td></tr>
<tr><td>DAh</td><td>CRC Error: Data error coding between the controller and host. If the CRC error count is not 0, the SATA signal is poor. Users should conduct a SATA signal trace from the host to the device (PCB layout, SATA cable, and SATA connector).</td></tr>
<tr><td>E7h</td><td>SSD Life Remaining, SSD Life Left: Presented as a percentage and calculated based on the average erase count and NAND reference erase count.
$$\textbf{SSD Life Left} = 1 - \left(\frac{\boldsymbol{Average\ Erase\ Count}}{\boldsymbol{NAND\ reference\ erase\ count}}\right)\%$$</td></tr>
<tr><td>EAh</td><td>Total NAND Read: The total data size that the SSD controller reads from NAND flash. Sector unit size is 512 bytes.</td></tr>
<tr><td>EBh</td><td>Total NAND Written: The total data size that the SSD controller writes to NAND flash. Sector unit size is 512 bytes.</td></tr>
</table>

Table 3. SQFlash-Defined S.M.A.R.T. IDs

SSD Lifespan Prediction

Although the Endurance Check pie chart can be referenced to determine the remaining product lifespan, the pie chart cannot predict the exact end-of-life (EOL). To help users predict the SSD lifespan both generally and under specific operating conditions, Advantech's SQFlash Utility software features General Prediction and Endurance Estimation functions.

General Prediction

The General Prediction function provides a general EOL prediction for the SSD. This prediction is based on the remaining SSD lifespan and calculated using the following formula:

$$\text{Estimated usable time} = \frac{(Flash\ P/E\ Cycle\ -\ Average\ Erase\ Cycle)}{Average\ Erase \div Power\ on\ hours} \times Safety\ Factor$$

The predicted EOL provided by the General Prediction function is presented as a specific date and time, as shown in Figure 3. This date is determined based on the estimated usable time and the current date and time. The Safety Factor acts as a buffer that allows users to prepare for SSD replacement. The Safety Factor value can be set as 0.8.

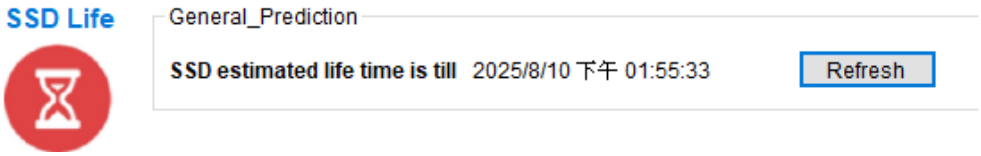

Figure 3. SSD Estimated Lifespan

Endurance Estimation

The Endurance Estimation function provides an exact EOL prediction for the SSD under specific operating scenarios. This function can assist users with risk assessments and spare parts inventory management. The predicted EOL is determined by using the “Start” and “Stop” buttons to measure the average erase cycle generated by writing data within a specified period of time. The Write Amplification and Terabyte Written (TBW) factors for the specific operating scenario can be calculated using the following formula:

$$\text{SSD Lifetime Estimation} = \frac{Flash\ Maximum\ P/E\ Cycle}{Average\ Erase\ cycles} * Test\ time$$

$$\text{Write Amplification} = \frac{Average\ Erase\ Cycles\ *\ SSD\ Capacity}{Host\ Write}$$

$$\text{TBW} = \frac{SSD\ Capacity * Flash\ \ P/E\ Cycle}{Write\ Amplification}$$

*Average Erase Cycles: The average erase cycles added since the Start button was last pressed.

*Host Write: The total amount of data written by the host to the SSD since the Start button was last pressed.

When the “Start” button is pressed, data for the specific operating scenario is collected in real time, as shown in Figure 4. When the task or operation is completed, press the “Stop” button to obtain the results. A pop-up window displaying the text results will then appear, as shown in Figure 5.

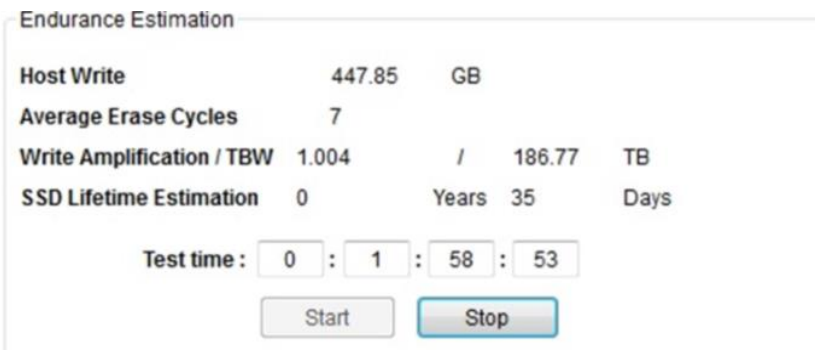

Figure 4. Endurance Estimation Real-Time Data

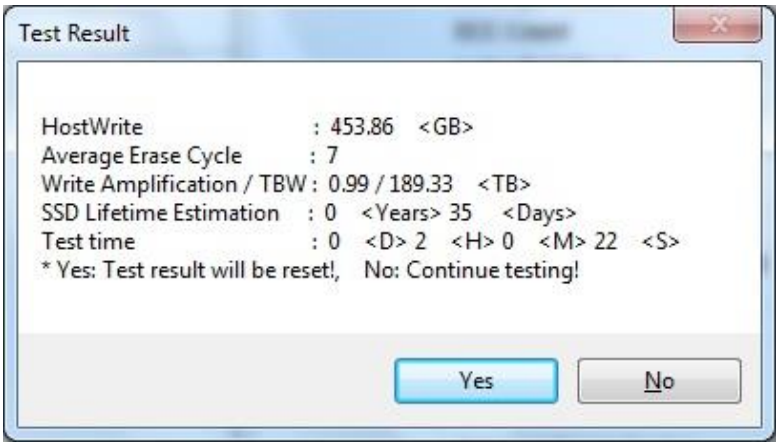

Figure 5. Test Results for a Specific Operating Scenario

Remote Management

For SSD status and lifespan monitoring, industrial application users may have one-to-many or many-to-one remote monitoring requirements. In order to meet diverse needs, Advantech's WISE-PaaS/DeviceOn remote management and monitoring platform was equipped with a Remote SSD Monitoring function. The WISE-PaaS/DeviceOn platform comprises two major components — a server-side system (the device manager) and a device management agent. The server-side system communicates with devices by using various protocols in order to provide both individual and bulk device control. Additionally, the server-side system facilitates remote management of the software and applications deployed on devices, and can be used to monitor/reset device applications or processes if necessary.

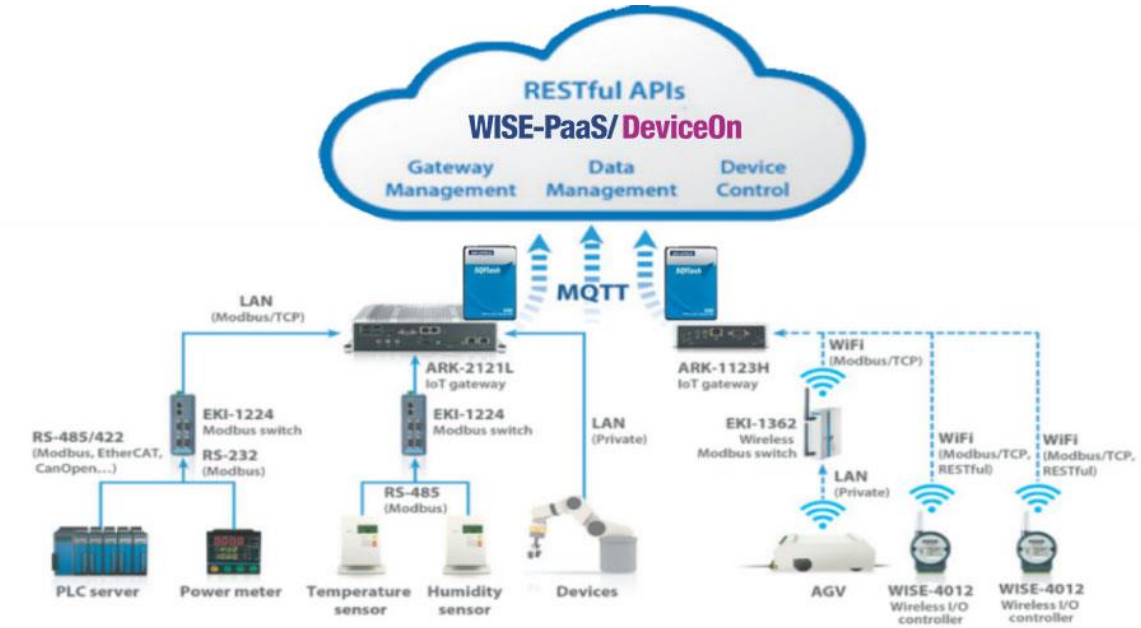

Figure 6. SQFlash Cloud-Based Management Diagram

Users can remotely monitor the SSD's current operating status and remaining lifespan in real time from a PC or cell phone via the WISE-PaaS/DeviceOn web page. The WISE-PaaS/DeviceOn web page provides critical information, such as the Endurance Check results, ECC count and good Block Rate, as shown in Figure 7.

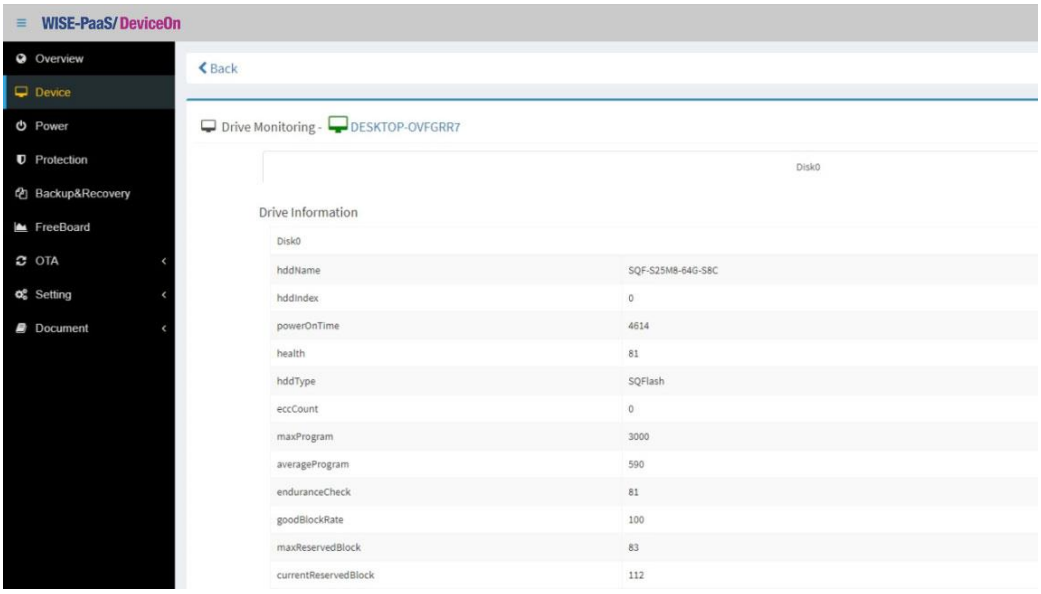

Figure 7. WISE-PaaS/DeviceOn Operation Interface

To allow users to understand the SSD status at a glance, the WISE-PaaS/DeviceOn platform also features a built-in dashboard. Users can customize the dashboard layout to emphasize or highlight specific data, as shown in Figure 8.

Figure 8. WISE-PaaS/DeviceOn SQFlash Dashboard

Predictive Maintenance Quality

To provide users with another means for quickly understanding the SSD status, the WISE-PaaS/DeviceOn platform is also equipped with a Predictive Maintenance Quality (PMQ) function. The PMQ value is calculated based on the predicted EOL provided by the General Prediction function (see Section 2.2.1), with historical RMA and DOA failure mode analysis used as the auxiliary criterion. The RMA and DOA failure mode analysis procedures are regularly updated to optimize machine learning algorithms. The PMQ user interface shows the current health status and latest event notifications. For events that require special attention, the PMQ interface will display a warning message or issue a user notification, as shown in Figure 9.

Figure 9. PMQ User Interface

Conclusion

The primary goal of Advantech's SQFlash development team is to assist SSD users with reducing data losses and securing assets. Advantech's SQFlash Utility provides an intuitive and convenient tool for obtaining the SSD health status. Using the SQFlash Utility's SSD Life Prediction function, users can also optimize risk assessments and spare parts inventory management. Regarding SSD management, Advantech's WISE-PaaS/DeviceOn platform can be used to conduct remote monitoring and management of SSDs in real time. Finally, Advantech's SQFlash Utility also supports customization to enable the achievement of intelligent SSD self-management according to specific usage requirements.

Download SQFlash Utility

For more information about industrial storage modules SQFlash, please visit http://sqflash.advantech.com