

3D NAND Flash with Advanced ECC Technology for Industrial Applications

A Guidebook for Designers- by Designers



What Are The Key Technologies of Current Storage Devices?

Learn about 3D NAND flash technology to help build a system of superior performance, ultra-scalability, and cost efficiency.

To achieve faster performance, ultra-scalability, and better cost efficiency, advanced 3D NAND flash technology was introduced onto the market in recent years.



The purpose of this playbook is to provide an overview of 3D NAND flash technology, its advanced Error Correcting Code (ECC) technology, and low-density parity-check (LDPC) code for endurance and reliability improvements.

What's NAND Flash Memory?

NAND flash memory is built up of many cells that hold bits, and those bits are either turned on or off through an electric charge. How those on/off cells are organized represents the data stored on the SSD.

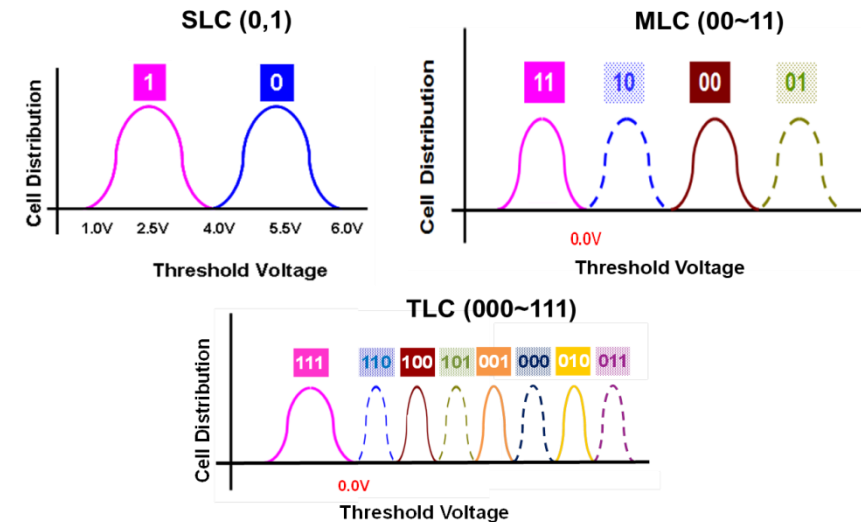
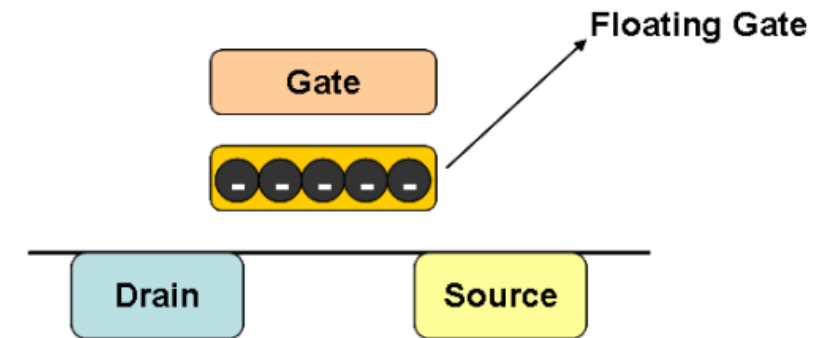
NAND flash memory is categorized into three types:

SLC (single-level cell)

MLC (multi-level cell)

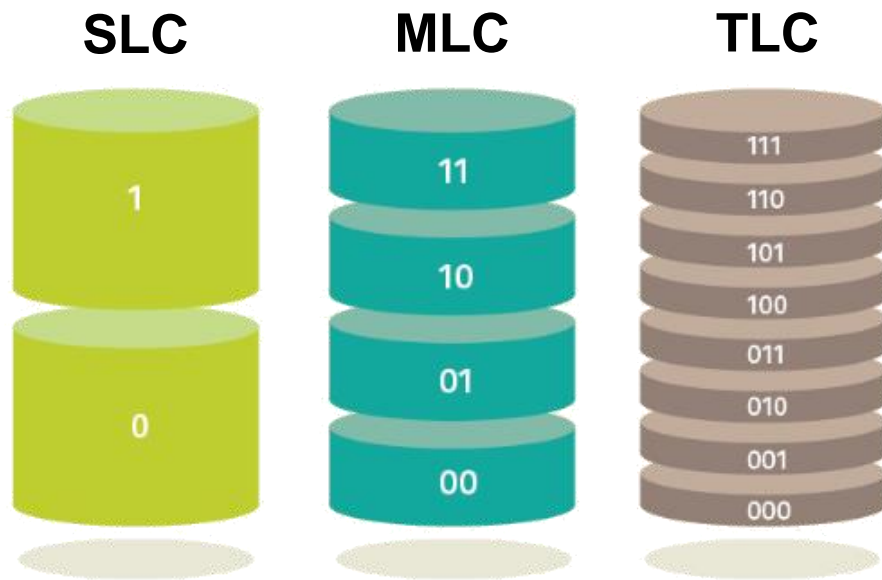
TLC (triple-level cell)

Basis Structure of a Memory Cell



NAND (SLC/MLC/TLC) Comparison

As the number of levels in a cell increases, more data can be stored on a single die for lower bit-cost. However, the trade-off for cost-saving is greater power consumption and lower endurance, due to more voltage levels required and technology limitations.



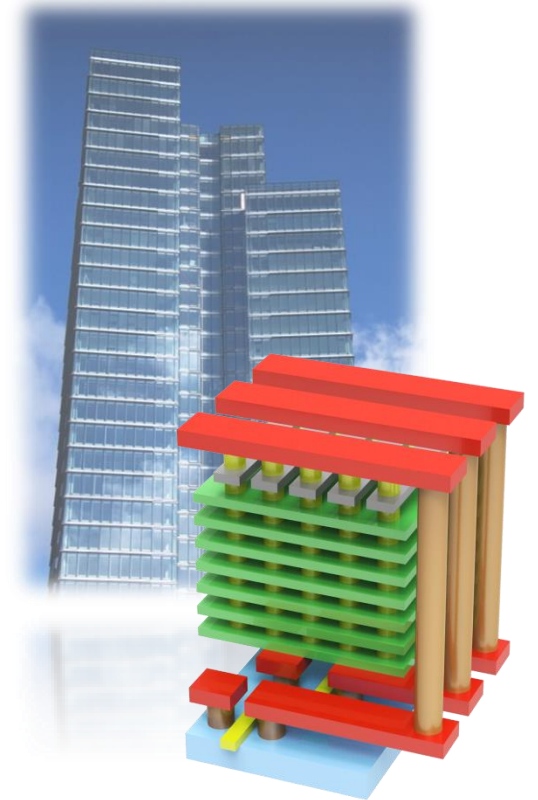
Flash Type	SLC	MLC	TLC
Storage	1 bits / cell	2 bits / cell	3 bits / cell
Program / Erase Cycle	100,000+	3,000+	500 ~ 1,000
Write Performance	Highest	High	Low
Cost-per-bit	Highest	High	Low
Power Consumption	Lowest	Middle	Highest
Application	Often Used in Industrial Grade Storage		Used More in Consumer Electronics

What is 3D NAND?

NAND Flash technology has been advancing to achieve higher storage density and lower bit-cost. However, in the planar process, individual cells interfere with each other in a space that is too narrow, so the 3D process was developed by scientists.

3D NAND is a precise process of vertically integrating NAND strings in a series. Memory transistors change from floating-gate types to trapped charge types.

3D NAND technology breaks the limitations of the 2D process by vertically stacking. By stacking more layers, 3D NAND can increase the unit density of storage



What Are the Benefits of 3D NAND?



Higher Capacity

By stacking more layers, 3D NAND can achieve higher storage density and ultra-scalability.



Reliable

P/E cycle count and data retention are higher than those of 2D TLC. With proper setting, the endurance level can reach 2D MLC grade.



Less Power

Less cell-to-cell program interference leads to lower power consumption.



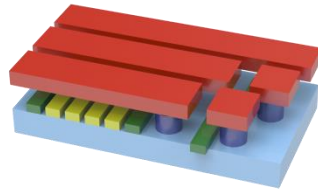
Saving TCO

Higher capacity and lower power consumption bring greater cost efficiency.

Differences Between 2D NAND and 3D NAND

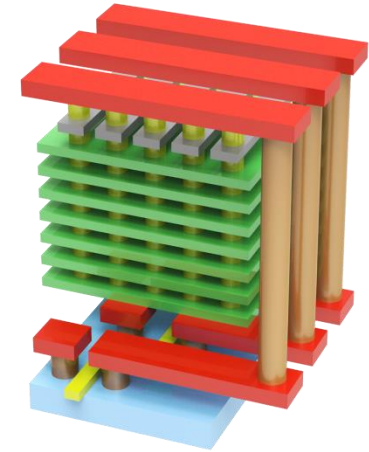
2D NAND Flash

15nm floating gate is approaching the physical limitation that pushes new lithography techniques to take over.



3D NAND Flash (BiCS)

Enlarged capacity within limited space provides good fab re-use by using existing lithography tools.



Flash Type	SLC	MLC	TLC	3D TLC (BiCS 3)
Storage	1 bits / cell	2 bits / cell	3 bits / cell	3 bits / cell
Program / Erase Cycle (ECC: BCH)	100,000+	3,000+	500 ~ 1,000	1,000 ~ 1,500
Write Performance	Highest	High	Low	Middle
Cost-per-bit	Highest	High	Low	Lowest
Power Consumption	Lowest	Middle	Highest	High

Better than 2D TLC but not good enough for industrial grade

How to Improve Endurance of 3D NAND Flash?

High Performance ECC Helps Data Accuracy and Improves Endurance

Error Correction Code (ECC) technology

Bose, Chaudhuri, and Hocquenghem (BCH)

Traditional algebraic coding method can correct up to a specified, fixed number of errors.

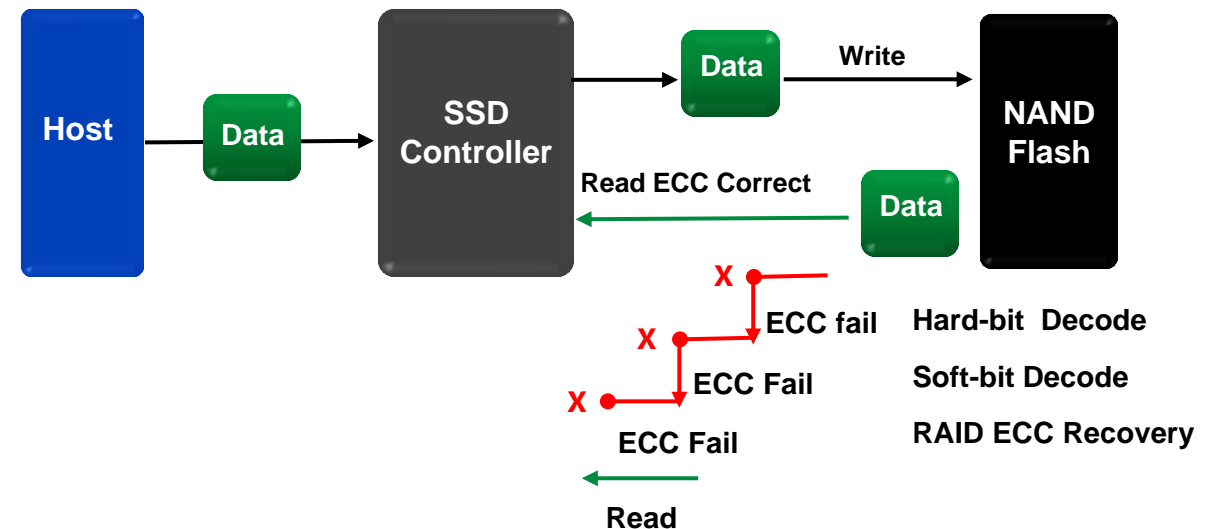
Low-Density Parity-Check (LDPC)

Using hard-bit (binary) and soft-bit (probabilities of bits) decode is 2~3x better performance over BCH.

RAID ECC

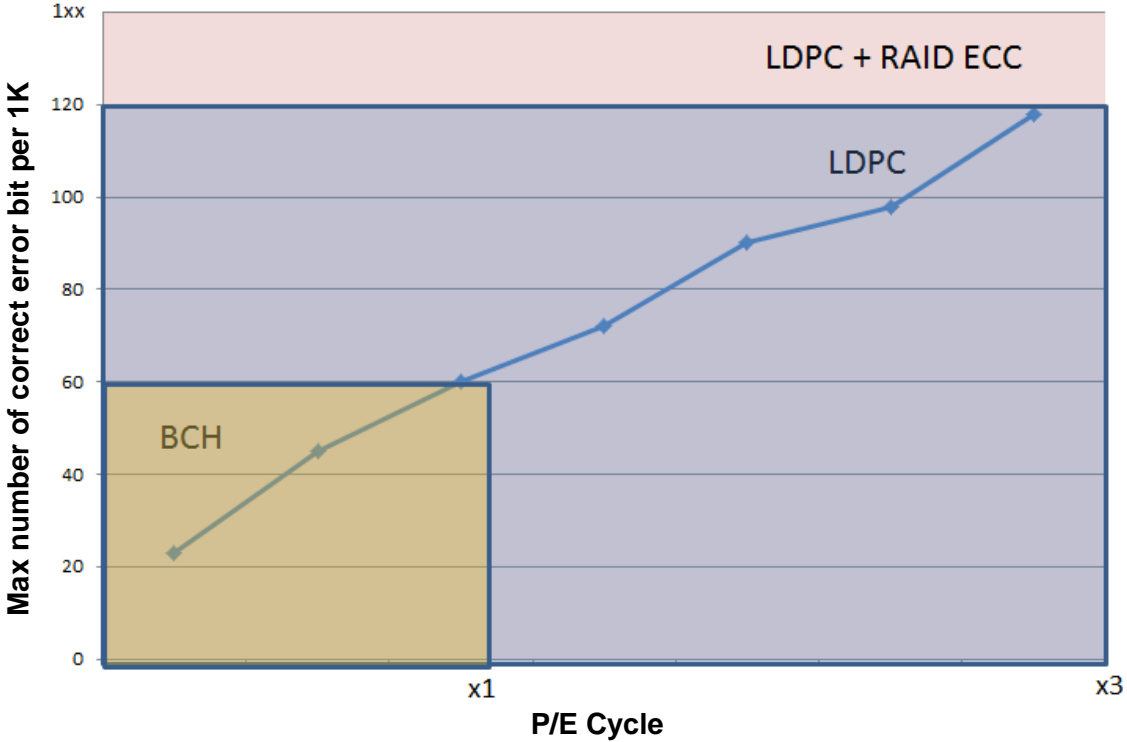
By identifying the RAID ECC parity between page and page, the performance can be enhanced more.

ECC Coding + RAID ECC



Can 3D NAND Flash (BiCS 3) Achieve Industrial Grade?

ECC Technology	BCH	LDPC	LDPC+RAID ECC
Decoding Algorithm	Algebraic Based	Probability Based	Probability Based
Performance of Hard-bit Decode	Fixed	1.3X BCH	1.3X BCH
Performance of Soft-bit Decode	-	2X~3X than BCH	3X+ than BCH <i>Winner</i>
Decoding Complexity	Low	Mid	High
Die Cost	Low	High	High



With LDPC + RAID ECC, BiCS 3 boosts P/E cycles to over **3,000** times (originally 1000~1500 times with BCH) into industrial grade.

NAND Flash Types in Industrial Application

SLC, MLC, and ultra MLC¹ SSD are still the mainstream for industrial applications due to their **high reliability, high endurance** and **stable performance** features.

Note1: Ultra MLC acts as SLC based on MLC structure

3D TLC (BiCS 3) SSD is new technology that is easier to adapt in commercial/industrial applications such as POS, Kiosks, and data centers due to its **cost-effective** and **large capacity** features.



Fan-less / Embedded System



Surveillance and Automation Machine / Write-Intensive Application



POS, Kiosk and Data Center / Read Intensive Application

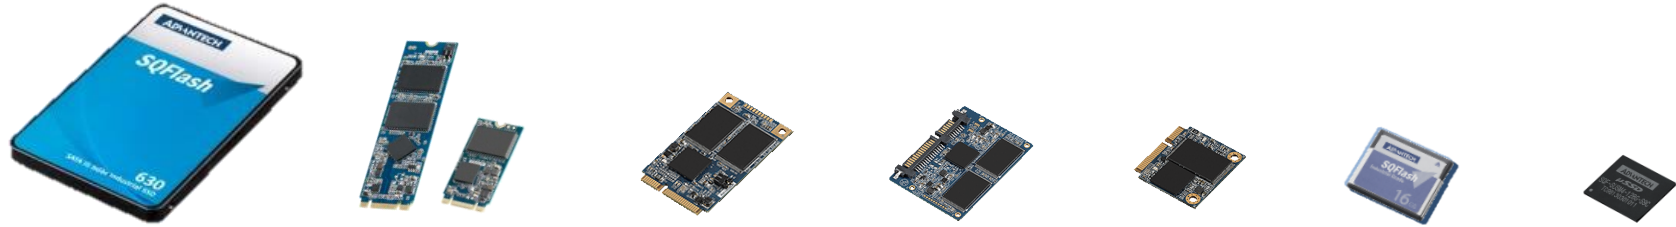
Conclusions



From 2D NAND flash to 3D NAND flash, higher storage density devices have emerged quickly in recent years. However, while eager to pursue cost-effectiveness, it is likely that endurance, reliability, and/or performance will be sacrificed.

LDPC ECC with RAID ECC technology could enhance the P/E cycle two- or three-fold, which enables 3D NAND flash (BiCS 3), to achieve MLC-level endurance that meets industrial-grade requirements for 3D NAND technology.

SATA 640 – Embedded Series



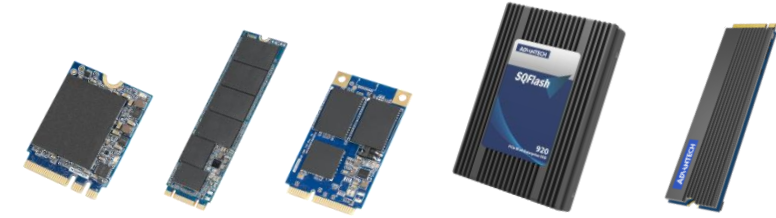
Product Series	Embedded 2.5" SSD	Embedded M.2 SSD	Embedded Modules				
Model Name	SQF-S25	SQF-SM8 / SQF-SM4	SQF-SMS	SQF-SLM	SQF-SHM	SQF-S10	SQF-SUS
Transfer Protocol	SATA 6Gb/s						
Connector	7 + 15 pin SATA	M.2 B + M Key	Mini PCIe with SATA pin-out	7 + 15 pin SATA	Mini PCIe with SATA pin-out	CFAST Type-I	Onboard
Flash Vendor	Toshiba						
Flash Type	Ultra MLC / MLC / 3D NAND						Ultra MLC / MLC
Capacity	16GB ~ 1TB (SM4SHM/S10 and SUS up to 256GB)						
ECC	LDPC with advanced ECC						
Add-on	Hardware write protect					N/A	
Op. Temp.	0 ~ 70° C / -40 ~ 85° C						
SQFlash Utility	GUI management tool & SW API pack						

SATA 840 – Enterprise Series



Product Series	Industrial 2.5" SSD	Industrial M.2 SSD
Model Name	SQF-S25	SQF-SM8
Transfer Protocol	SATA 6Gb/s	
Connector	7 + 15 pin SATA	M.2 B + M Key Mini PCIe with SATA pin-out
Flash Vendor	Toshiba	
Flash Type	3D NAND BiCS 3	
Capacity	256GB ~ 8TB	256GB ~ 2TB
ECC	LDPC, RAID ECC	
New Tech.	Global Fragment Writing, Flush Manager, Thermal Throttling	
Op. Temp.	0 ~ 70° C / -40 ~ 85° C	0 ~ 70° C / -40 ~ 85° C
SQFlash Utility	GUI management tool & SW API pack	

SQF 710 & 920 series – NVMe SSD

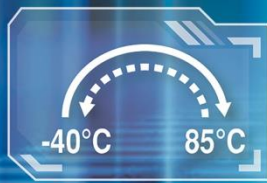


Embedded NVMe 710 series	Product Series	Enterprise NVMe 920 series
SQF-CM8 → M.2 2280 (B-M key) SQF-CM3 → M.2 2230 (A-E key) SQF-CMS → Full-size MiniPCIe	Form Factor	SQF-C25 → U.2 2.5" SSD (SFF-8639) SQF-CM8 → M.2 2280 (M key)
PCIe Gen3 x2, 2-lane	Transfer Protocol	PCIe Gen3 x4, 4-lane
Support ¹	AES256	Support
TCG-OPAL	SED feature	TCG-OPAL / TCG-Enterprise
128GB ~ 1TB	Capacity²	256GB ~ 8TB
Read: up to 1,600 MB/s Write: up to 1,100 MB/s Random 4K R/W: up to 240K/ 200K IOPS	Performance³	Read: up to 3,200 MB/s Write: up to 2,600 MB/s Random 4K R/W: up to 600K/ 600K IOPS
SmartECC	ECC mode	LDPC + RAID ECC
0 ~ 70° C / -40 ~ 85° C	Op. Temp.	0 ~ 70° C / -40 ~ 85° C

1. The SQF-CM8 of 710 series supports AES256.

2. The max. capacity of SQF-CM3 710 series upto 256GB. The SQF-CM8 920 series up to 4TB.

3. The speed result is max. capacity model.

**Wide Temp.****Thermal Solution****Design-in Service****Self-management**

Designing NVMe SSD for Industrial Applications

Author: Precyan Lee precyan.lee@advantech.com.tw

Keyword: NVMe SSD, PCIe SSD, Industrial SSD

Moving Forward –

Native PCIe Interface SSD for Industrial Applications

SSD performance is one of the biggest considerations when designing industrial applications. SSD technology is still evolving from PATA to SATA and the speed of its interface continues to increase. The PCIe interface appeared in the market almost at the same time as SATA 3.0 Gbps (Gen. 2) became popular. But due to the limitations of Flash IC technology and lack of host support, PCIe SSD wasn't successful. That is until today, Flash IC performance has now fully caught up. OS driver support for PCIe storage devices, and chipset native NVMe protocol support are ready to make the next big leap forward. As PCIe SSD has matured and become an important storage option for certain market segments, this paper discusses options for designing NVMe SSDs into industrial applications.

Performance Benefits

The throughput of the PCIe interface is the most crucial value for users or system integrators. The following table shows the general bus speeds and differences. The superiority of the PCIe product is plain to see.

	SATA Gen. 3	PCIe Gen. 3 x2	PCIe Gen. 3 x4
Interface bandwidth	6 Gb/s	16 Gb/s	32 Gb/s
Real product throughput	560 MB/s	1.97 GB/s	3.94 GB/s

PCIe SSD applications such as data centers, image / video processing, machine vision, and all applications that perform intensive data processing on storage media, all benefit from such performance levels.

In order to ensure the same high-performance levels with NVMe SSD media, which is at least 3 times faster than SATA-based SSD, the data bus signal quality needs to be taken into account. Poor signal quality not only compromises data throughput, but also causes major reliability concerns. The host PCIe bus design that compliments NVMe has to follow open device interface specifications for accessing non-volatile storage media attached via a PCI Express (PCIe) bus, and the system has to be verified by professional third-party signal integrity test labs. By doing so, compatibility between host (motherboard) and device (NVMe SSD) can be further assured.

Platform Compatibility

Form Factor Selection

I/O flexibility and compact design are often requested for industrial applications. As a result, embedded systems or embedded motherboards tend to incorporate different interfaces and sockets for different requirements. Even a small form factor board would still have a MiniPCIe socket, an M.2 socket, and even a PCIe connector for different application scenarios. Having an SSD selection would also enhance overall system design flexibility so several selections of NVMe based SSD for popular PCIe sockets are available for industrial motherboards

- MiniPCIe – one of the most popular sockets for industrial motherboards and available on most boards. By default, supports x1 interface, but with adjustments on both board and device x2 interface is possible.
- M.2 – is becoming popular not only for storage, but also wireless connection modules. The biggest problem with M.2 is their very diverse specifications, which can be categorized in different keys.

	Lanes	Common application
B key	x4	RF, cellular
M key	x4	Storage
A + E key	x2	WIFI

Hardware and Software Requirements

NVMe is new protocol for PCIe, it requests a certain level of chipset and OS driver support to work with its full functionality.

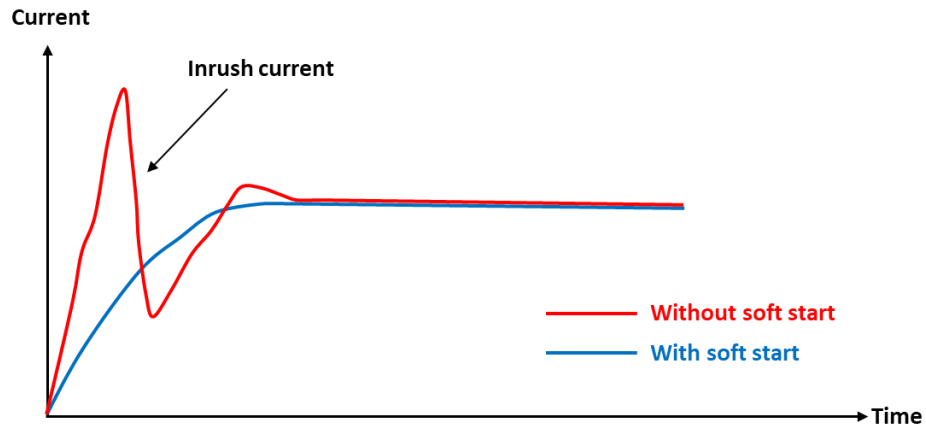
- Hardware – Intel 6th generation core processors (Skylake, Kaby Lake, Coffee Lake)
- OS – Windows 7 requests additional driver installation while Windows 8.1, Linux Kernel 4.4.16 and their newer versions already have the NVMe driver built-in

Sophisticated Power Management Design

In order to support stable performance and industrial level reliability, the stability of power and voltage inputs for the IC in NVMe SSDs is more important than the other kind of SSD. Several power management approaches should be considered.

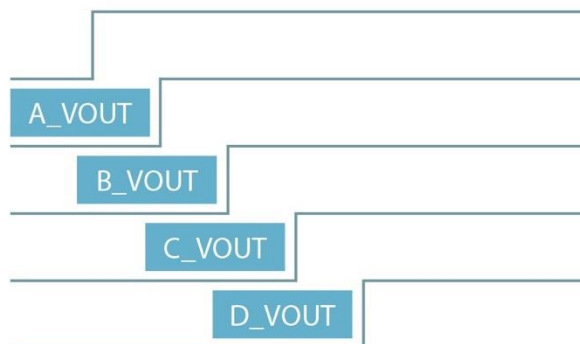
Automatic Internal Soft Start Operation

Reduce inrush current during instant power on, so the voltage level will be stable and not lead to data read / write errors.



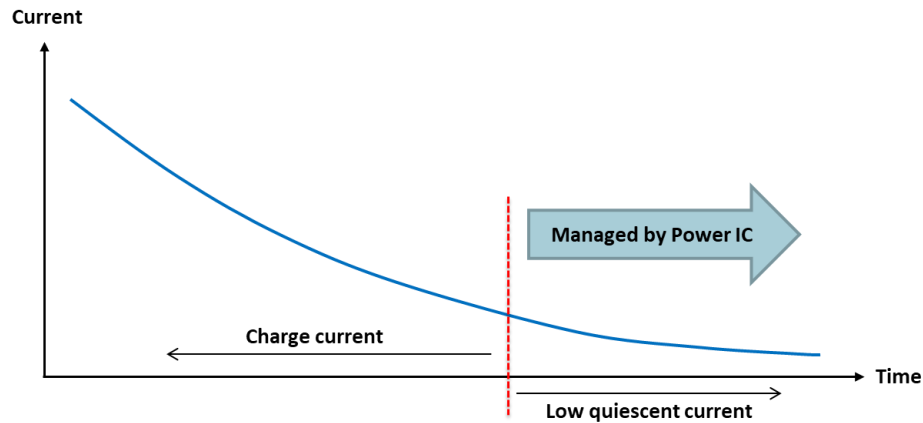
Power Sequence Control

There are multiple components (Flash IC / Controller IC / DDR IC) requesting separate power supply. A power sequence control function based on programmable on/off sequencing of multiple power supplies ensures the SSD internal power sequence will be well organized to achieve high working efficiency and stability.



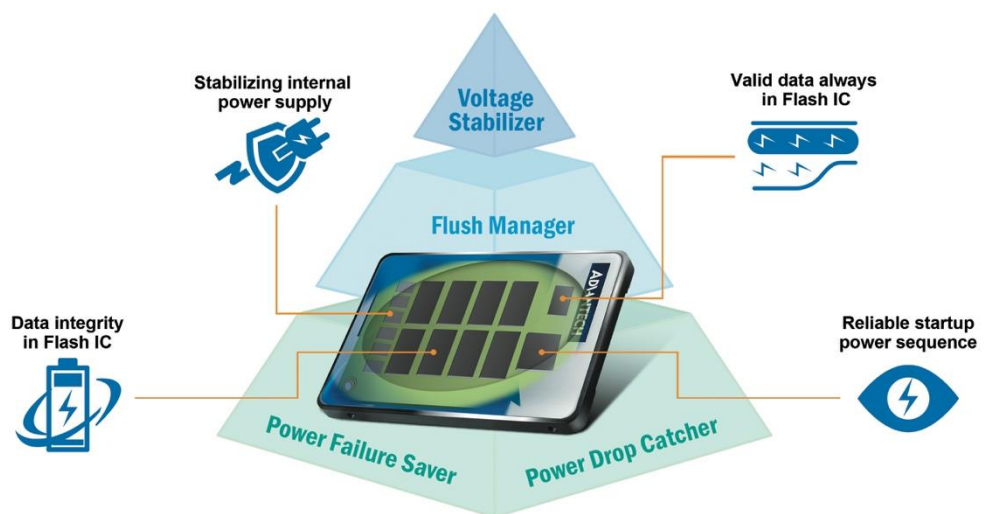
Low Power Mode Control

In low power mode setting, the SSD operates in light load mode. To keep SSD working efficiently it relies on stable power switching, driving loss protection, and quiescent current. Normally a power management IC will be needed to ensure stable 30 uA quiescent current.



Power Failure Protection

A proper power failure protection mechanism has to be implemented in order to maintain the highest level of data protection. This complete scheme combines four different functions – Power Failure Saver, Power Drop Catcher, Flush Manager, and Voltage Stabilizer. Details can be referred to another designated white paper ([link](#)).

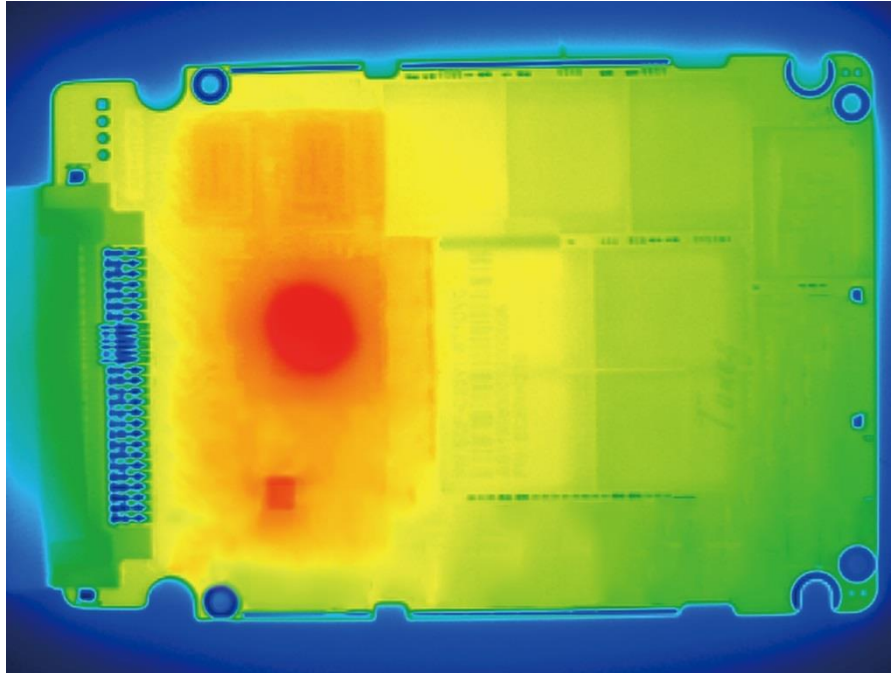


Thermal Solution Makes NVMe SSD Suitable for Industrial Applications

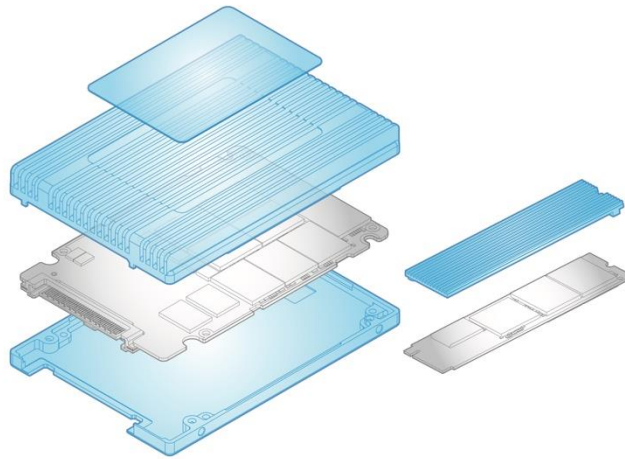
NVMe SSDs made with high frequency processors (controller) consume more power, which also generates more heat than traditional SSD. For reliable usage in industrial applications, even wide temperature scenarios, Advantech's SQFlash NVMe series

products are designed with several features that make them even more reliable.

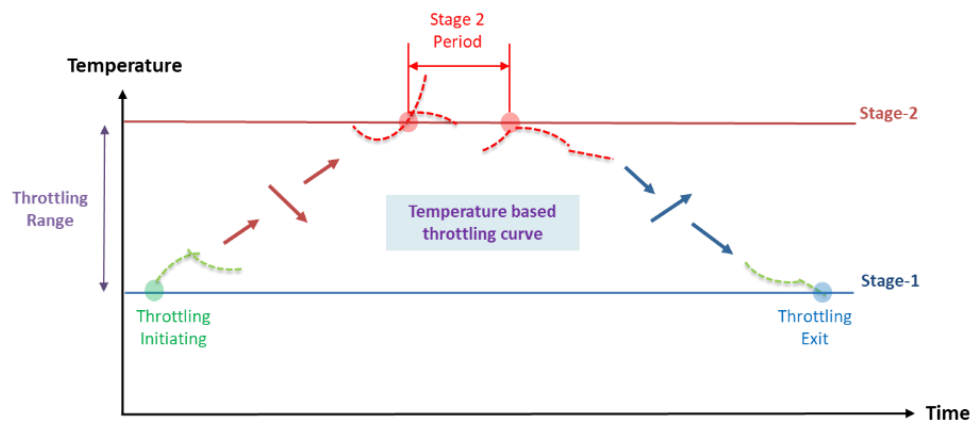
- Pre-design thermal simulation – thorough thermal simulation before and after PCB layout checks thermal weak points and deploys proper solutions to reduce massive heat buildup in real applications.



- Industrial heatsink design – for ruggedized applications that require high performance, NVMe SSD is the only option. Heat generation from SSDs is a big obstacle and prevents NVMe SSDs being implemented in many industrial application fields. Sophisticated thermal solutions with industrial grade heatsink designs can greatly improve their thermal performance and stable operation, enabling NVMe SSDs to be adopted in wide temperature / ruggedized environments.



- Thermal throttling management – when an SSD controller detects overheating by the internal sensor, a built-in intelligent firmware feature throttles the overall SSD performance to force the controller IC to cool down and prevent hang-ups or even physical damage to a device. The throttling mechanism is also divided into several stages to prevent sudden speed drops. After temperature returns to normal levels, the SSD automatically resumes full speed operation.



- Real-time temperature monitoring – built-in thermal sensor in the NVMe SSD can monitor SSD working temperatures and present SMART data. Thermal information can be accessed easily from the operation site or even be acquired through cloud ready solutions like Advantech WISE-PaaS platform with SSD PMQ (Predictive Maintenance Quality) function. Users can even set thresholds for the system to alarm or adjust workloads automatically.

A graphic for Advantech's SQFlash Management. It features a light blue cloud shape containing a QR code, the Advantech logo, the text 'WISE-PaaS Edge Intelligence Platform', and icons of a laptop, desktop monitor, and smartphone. To the right of the icons is a list of features: Remote Health Monitoring, Lifespan Prediction, and Online Predictive Maintenance. The main title 'Cloud-based SQFlash Management' is at the bottom in large blue font.

ADVANTECH
WISE-PaaS
Edge Intelligence Platform

Scan to Experience

- Remote Health Monitoring
- Lifespan Prediction
- Online Predictive Maintenance

Cloud-based
SQFlash Management



For more information about industrial storage modules SQFlash, please visit sqlflash.advantech.com