Draft Document for Review April 29, 2024 11:02 am REDP-5725-00

Redbooks
ibm.com/redbooks

# IBM FlashCore Module (FCM) Product Guide

## Features the newly available FCM4 with AI-powered ransomware detection

Hartmut Lonzer

Jon Herd

Vasfi Gucer

IBM

Redpaper

# IBM FlashCore Module (FCM) Product Guide

This IBM® Redpaper Product Guide describes the IBM FlashCore Module (FCM) history, a general overview and then a deeper dive on the way IBM leads the field in the adoption of high speed, low latency storage.

The IBM FlashCore Module is used in the latest IBM FlashSystem® solutions, IBM Elastic Storage® System (ESS) and the IBM Storage Scale System (SSS). The IBM FlashCore Module (FCM) combines the performance of flash and a Non-Volatile Memory Express (NVMe) optimized architecture with the reliability and innovation of IBM FlashCore® technology and the rich feature set and high availability (HA) with IBM Storage Virtualize software. The IBM FlashCore Module (FCM) is also a computational storage device off loading tasks from the storage controller enabling overall storage functionality.

**Note:** Check out the recent article by *Sam Werner, VP, IBM Storage Product Management* on the launch of IBM's fourth-generation FlashCore Module technology.

## Terminology

Here are some explanations of the items commonly used when discussing *flash memory* in general.

### Memory cell states

The following considerations apply to memory cell states:

- Flash memory stores data in individual memory cells and each cell has N possible states,
- In SLC memory this N will be either a one or a zero, so one bit of data was stored in each cell.
- For other types of flash memory, each cell has N states depending on the type of flash**.**
- IBM uses the concept of voltage states as the way the data is stored. For SLC, it has only two voltage states, and MLC, there are 4 voltage states, and so on.
- SLC memory has the advantage of higher write speeds, lower power consumption and higher cell endurance. However, because SLC memory stores less data per cell than *multi-level cells*, or MLC memory, it costs more per megabyte of storage to manufacture. Due to higher transfer speeds and expected longer life, SLC flash technology is used in high-performance memory cards.
- Multi-level cells (MLC) are a memory cell capable of storing more than a single bit of information, compared to a single-level cell (SLC), which can store only one bit per

memory cell. However MLC is almost obsolete now, as density and 3D stacking has come. It therefore offers little advantage over TLC.

- *Triple-level cells* (TLC) and *quad-level cells* (QLC) can store three and four bits per cell respectively. The name *multi-level cell* is sometimes used specifically to refer to the *two-level cell*. Overall, the memories are named as follows:
  - Single-level cell or SLC (1 bit per cell).
  - Multi-level cell or MLC (2 bits per cell),
  - Triple-level cell or TLC (3 bits per cell).
  - Quad-level cell or QLC (4 bits per cell).

### NAND terminology

The following discusses the NAND terminology:

- The smallest piece of a NAND flash is a *cell*.
- Each cell is stored in a *page*. Each page can be written to, and they are the smallest piece of the NAND flash that can store data or be programmed.
- Groups of pages are called *blocks*.
- Blocks are manufactured as a batch called *dies*.
- A chip, or die, is what ends up being put on a *circuit board* in a finished product.

### What is Non-Volatile Memory Express (NVMe)?

NVM Express is an open collection of standards and information to fully expose the benefits of non-volatile memory in all types of computing environments from mobile to data center. NVMe is a protocol that is built on PCIe standards.

The original NVM Express Work Group was incorporated as NVM Express in 2014 and is the consortium responsible for the development of the NVM Express specification. The organization currently has over 100 member companies.

### What is NAND flash memory?

NAND flash memory is a type of non-volatile storage technology that does not require power to retain data. An important goal of NAND flash development has been to reduce the cost per bit and to increase maximum chip capacity so that flash memory can compete with magnetic storage devices, such as hard disks. NAND flash has found a market in devices to which large files are frequently uploaded and replaced. MP3 players, digital cameras and USB flash drives use NAND technology.

**Note:** Although it is true that it NAND memory doesn't require power, but it does require to be powered on at some point in the future, to keep the data correct. For example, if we leave a device off for more then 90 days at 40 degrees C, we start to see issues with data retention.

3D NAND flash stacks memory cells vertically in multiple layers, achieving a higher density than traditional NAND memory. It can be used in the same situations as traditional NAND thus taking advantage of the higher density for the same footprint. However going beyond QLC is very difficult with the current technology used.

## How Flash is evolving

Flash systems have been evolving:

- Lithography (more cells per die).
  - Lithography reduces the width of the traces and the size of the individual transistors. For a while NAND vendors just reduced this lithography, which reduced the size of the cell in a 2 dimensional structure. Then endurance became very difficult.
  - They then went 3 dimensional and did not rely as much on smaller lithography but on 3D stacking. Current stacking is 176 layers in FCM4.
  - We will see 300 and more in the next few years.
- More bits per cell.
  - Signal processing to discriminate between multiple bits in a single cell. 1 bit (SLC), 2 bits (MLC), 3 bits (TLC), and 4 bits (QLC) per cell.
  - Voltage detection is used to confirm what voltage state a cell is in, thus how the cell encoding level is determined
  - Analog reading of the cell voltage level is what determines the state. This is true for each of the cell types. QLC splits it into 16 voltage levels.
- More layers increases the density in 3D NAND.
- More dies per package.
  - More dies per package allows for increased density.
  - More dies also help to improve performance, as these dies can be accessed at the same time.
- Improved management.
  - Better management of NAND, decreases wear and improves performance.

Figure 1 on page 3 shows the different ways we can expand the NAND chips density.

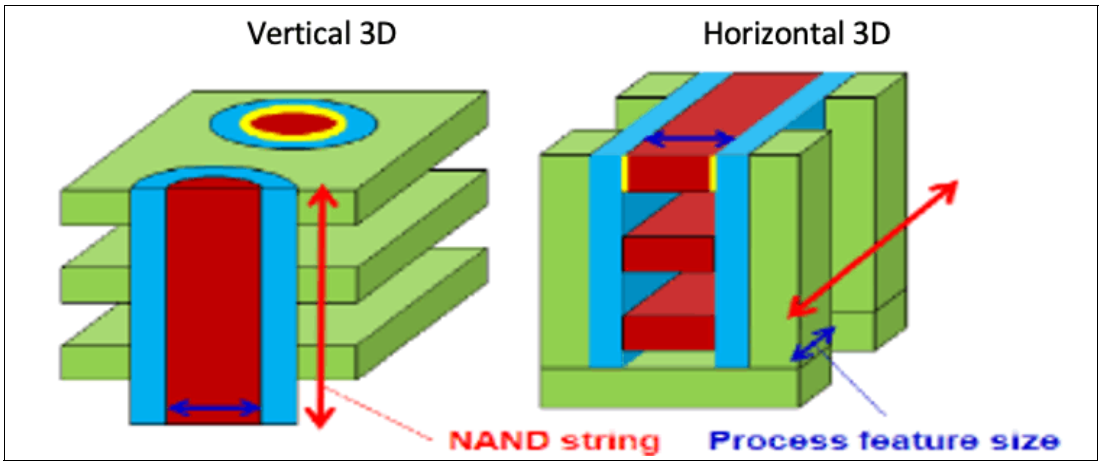


*Figure 1 NAND types*

With all NAND based flash systems there are many challenges that need to be overcome to ensure high speed access, low latency and durability. Table 1 explains each of these functions and how IBM is addressing these within the design and FCM technology.

*Table 1 Flash management challenges and solutions*

| Function | Description / Observations | Design Goals |
|---|---|---|
| Garbage Collection | Reclaim invalidated space due to out-of-place writes. Relocation of valid data leads to write amplification (WA). | - Smarter data placement using heat segregation reduces write amplification. |
| Wear Leveling | Traditional approaches equalize usage of Flash cells by balancing P/E cycles of blocks. Wear leveling moves further increase WA. As the drive fills up, the controller must move data and erase blocks before new data can be added. This process of erasing and writing (programming) data is referred to as the P/E cycle. | - Equalize block health instead of program / erase (P/E) cycles.<br>- Dynamic wear leveling: Smarter data placement using health binning.<br>- Static wear leveling: Reduce to strict minimum to ensure retention targets. |
| Health Management | Blocks that reach the error correction capability of the ECC must be retired. Retired blocks eat up over-provisioning and ultimately limit device endurance even if there are still many good blocks available. Reference link for heat binning and heat segregation here | - Continuously monitor block health and shift threshold voltages accordingly<br>- Actively narrow health distribution of all blocks with health binning. At end-of-life remaining good blocks only have little P/E cycles left. |
| Error Detection & Correction | Industry Standard SSDs perform read-retry and/or rely on ECC schemes using soft information. Read latency deteriorates with age of the device. | - Stronger ECC that does not require read-retry<br>- Variable Stripe RAID<br>- Array- level RAID |
| Data Reduction | Use NAND flash capacity efficiently, taking advantage of data that repeats. All NAND management preformed on data reduced sets | - Minimize impact to performance – Scaling<br>- Use standard methods - Dynamic Huffman, GZIP |

## Floating Gate versus Charge Trap NAND technology

FlashCore Module 1, 2 and 3 employed Floating Gate NAND flash. Each cell had its own *Floating Gate* to store the charge that represents either “1” or “0”. On top of that IBM added a lot of significant technology around it.

In Floating Gate NAND flash the electrons had a pretty small surface area to escape. It worked well but had some limitations that were becoming difficult to move past in order to continue delivering advanced features and even more performance.

As IBM was developing FlashCore Module FCM4s, there were changes beginning to take place in the flash industry which meant all vendors would need to prepare for a new type of NAND flash known as *Charge Trap flash*.

In Charge Trap flash, electrons get *UNTRAPPED* in every direction. What it means for the market is that flash manufacturers can create a more attractively priced solution where it is easier to add layers thereby creating better density.

The benefit to clients is that IBM's extensive flash knowledge and development skills, enabled them to very quickly adopt to this new approach. At some point all vendors will use Charge Trap flash, but by employing IBM's advanced FlashCore technology IBM has been able to bring greatly enhanced endurance and performance to this new implementation of flash and do so well in advance of other vendors.

## History of IBM FlashCore Modules

IBM has been delivering high-performance, highly-reliable customized flash modules for many years. It started with MicroLatency Modules for the high-performance FlashSystem family. These modules separated the control path and the data path so that nothing would slow down data. They added endurance features and implemented RAID within the modules. Numerous additional technologies and benefits were implemented.

Figure 2 shows the original IBM FlashCore MicroLatency Module used in the IBM FlashSystem 840 and 900 products on the left and the newer IBM FlashCore Modules on the right, as used in the current IBM FlashSystem product line.

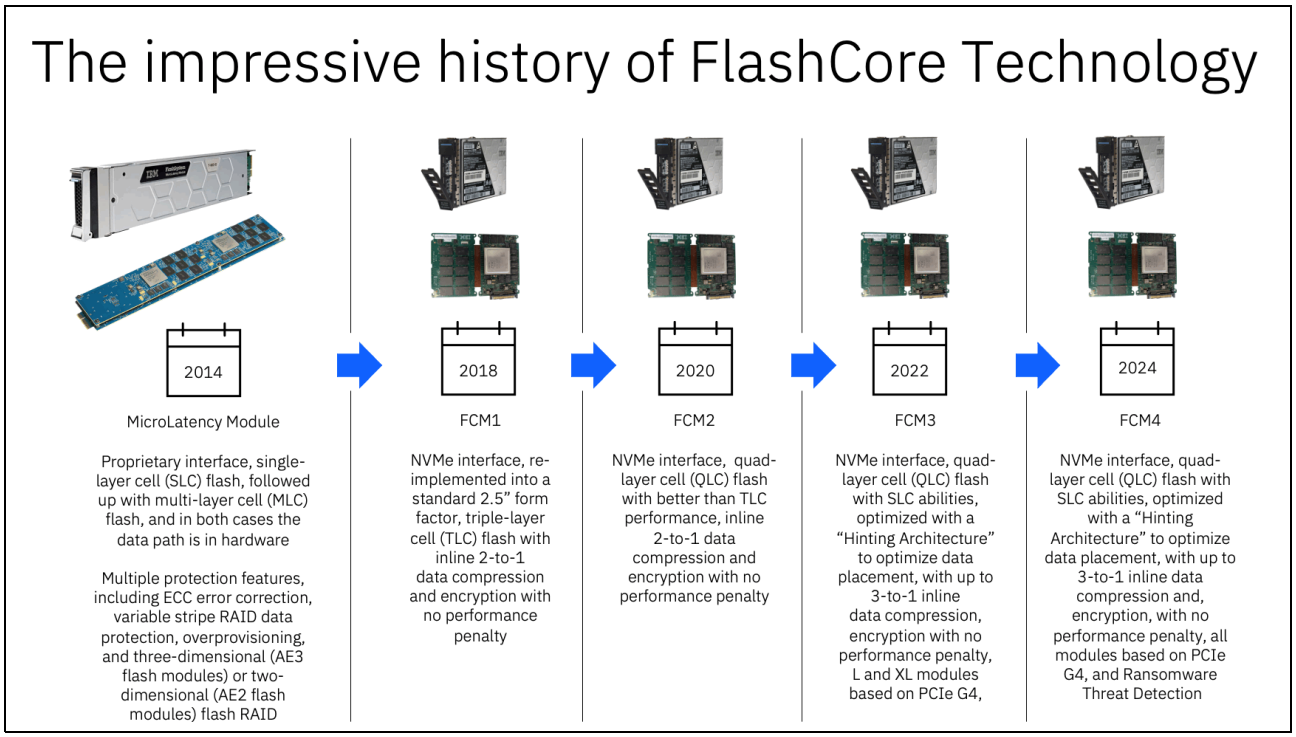


*Figure 2 Evolution of FlashCore*

As flash drives entered the mainstream storage market IBM noticed the majority of vendors were willing to use off-the-shelf commodity flash drives. IBM was confident that real value and efficiency would only be found in unique custom-built flash devices, rather than commodity off-the-shelf drives. On August 16, 2012, IBM announced that it had signed a definitive agreement to acquire Texas Memory Systems, who was a leading developer of high-performance flash memory solutions.

**2014**, IBM introduced MicroLatency Modules. These were custom built flash media that employed a proprietary interface and used single-layer cell (SLC) flash. In the following years IBM delivered MicroLatency Modules employing multi-layer cell (MLC) flash. To achieve maximum performance the data path was entirely in hardware. MicroLatency Modules had encryption and multiple protection features, including ECC error correction, variable stripe RAID data protection, over provisioning, and three-dimensional (AE3 flash modules) or two-dimensional (AE2 flash modules) flash RAID. AE3 modules also had compression and TLC level technology flash.

**2018**, IBM re-engineered MicroLatency Modules and among other things converted them to utilize a standard 2.5” form factor employing a Non-Volatile Memory express (NVMe) interface. The key item here was taking what had been done with the MicroLatency Module and moving it into a standard form factor and a standard interface.These new devices were known as FlashCore Modules (FCM1). They were based on triple-layer cell (TLC) flash and delivered 2-to-1 data compression and encryption with no performance penalty. Each FCM contained a large custom-designed flash controller that provided the penalty-free compression. Large numbers of FCMs could be added to a FlashSystem array and performance would remain consistent. These modules were unique in the storage industry.

**2020**, IBM introduce FCM2s. These second-generation devices continued to use an NVMe interface, and in a surprising move they employed quad-layer cell (QLC) flash. The rest of the storage industry largely avoided QLC because they felt it was too slow for primary storage. IBM shocked the industry by demonstrating that the QLC-based FCMs delivered superior performance to their TLC-based predecessors. With FCM2, IBM completely switched to QLC. No one else has been able to do this. IBM Research and Development worked together to figure out how to get the endurance and performance needed to use QLC everywhere.

FCM2 had the ability to compress up to 2.3 to 1 if the drive was full. We define effective capacity as the amount of capacity that can be addressed. With FCM2 you could compress more than 2.3 to 1 if you had less physical usage. For example, if you had 21 TB of data with 3 to 1 compressibility, you could store that in 7 TB on a 9.6 TB drive. However, you would only have 1 TB of addressing left because the FCM could go to 22 TB of effective addressing.

**2022**, IBM introduced FCM3s. These third-generation devices continued to use an NVMe interface and quad-layer cell (QLC) flash. IBM demonstrated their design genius by enabling these new modules to support high-performance single-layer cell (SLC) pages and high capacity QLC pages. IBM delivered improved performance and efficiency with a “Hinting Architecture” to optimize data placement on the SLC and QLC pages. In another improvement the FCM3s had built-in 3-to-1 data compression and encryption with no performance penalty, So, you could now write up to 28.8 TB to a 9.6 TB FCM. The compression algorithm stayed exactly the same.

**2024**, IBM introduced FCM4s. IBM once again improved the modules. IBM introduced another industry leading breakthrough called *Ransomware Threat Detection*, which is a process that identifies and responds to security threats before they can damage data or systems. The FCM4 collects detailed statistics on every I/O operation (IOP) for each virtual disk (Vdisk). This data is then intelligently summarized for efficient processing. The FCM4 transmits this summary to Storage Virtualize, which relays it to an AI-powered inference engine. This engine can identify unusual activity, like potential ransomware attacks, in under a minute. Upon detection, an immediate alert is sent to IBM Storage Insights Pro, allowing for swift action. Additionally, the information can be shared with IBM Storage Defender if available, further strengthening your security posture.

**Note:** The ransomware threat detection enablement requires FCM4 drives running FCM firmware 4.1 or higher, plus IBM Storage Virtualize code v8.6.3 or higher.

There is also the ability to upgrade up to IBM Storage Virtualize code v8.6.3 and FCM firmware 4.1 with an existing array, assuming you initially created your array with FCM4's and IBM Storage Virtualize code v8.6.2 or later.

IBM also advanced the data reduction abilities of the FCM4 by supporting the ability to have up to 3-to-1 data compression if the user has the right type of data. In a move to further optimize performance, all FCM4s are based on PCIe Gen4. We also moved to support the latest 176 layer flash from Micron with Charge Trap NAND and kept the endurance required.

IBM pioneered custom inline compression and encryption into their FlashCore Modules (FCM), which again has no impact on performance.

Figure 3 on page 7 shows a summary of the evolution of the FCM technology.

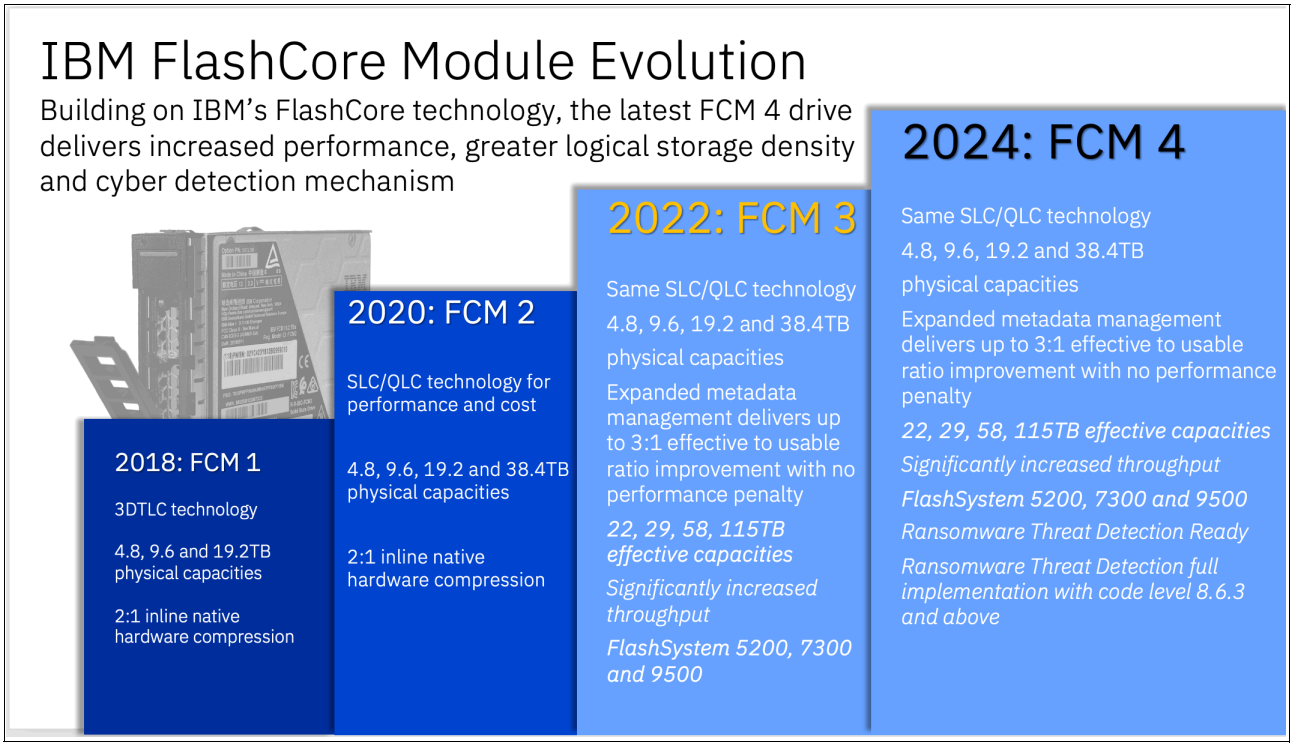


*Figure 3 IBM FCM evolution*

**Note:** All capacities are PCIe G4 with FCM4 vs FCM3, the 4.8 TB and 9.6 TB drives were PCIe G3.

The new FCM4 drives will be offered on IBM FlashSystem 7300 and IBM FlashSystem 9500, but they will also work on FlashSystem 5200.

## FCM facts and features

Back in the early days of flash, a lot of people were discussing comparisons of serial-attached SCSI (SAS) hard disk drives (HDDs) and SAS solid-state drives (SSDs). The performance was much better with flash but there were concerns about longevity.

The conversation around solid-state storage transitioned several years ago, with a focus on comparing traditional SAS SSDs to the then-new NVMe technology. NVMe SSDs offered a significant leap in performance. While reliability and longevity of the underlying flash memory remained comparable between the two technologies, NVMe's performance edge became the dominant factor.

However as flash ages on competitor’s and commodity solid-state drives (SSDs), errors begin occurring. This is often due to improper gate activity. SSDS actually do not have moving parts, and we should not consider the gate as a moving part. It is more about forcing electrons through an insulator. Every time that the electrons are forced through this insulator, in causes damage to the insolation layer, which results in leakage. This makes the cells worse in being able to hold a specific voltage level, which is how the wearing of a cell occurs.

Cells can be retired if the are not able to hold a voltage. Pages and blocks can also be retired, and this is where over provisioning helps. Eventually enough of the blocks wear out, and this results in the SSD not being able to support the full capacity that was originally designed for, resulting in the SSD becoming inoperative.

To address this issue, IBM has implemented Variable Voltage technology into its FlashCore Modules (FCMs).

IBM performs dynamic read level shifting over the life of the flash blocks in the FCMs. This ensures that all flash cells are monitored and automatically have the necessary voltage applied in order to deliver the longest possible lifespan.

Because IBM monitors all flash cells, it continuously implements predictive techniques to adjust internal flash settings in advance, thereby minimizing the probability of uncorrectable errors.

Using IBM’s advanced characterization lab, IBM Flash developers have determined the best voltage levels to set for a block as it ages proactively. This is only possible because IBM develops their own and means that IBM can keep older flash fully productive.

The business benefits of FCMs is that Variable Voltage is the “Fountain of Youth.” It brings a long reliable life to the flash in the FCM. IBM is the only flash vendor who provides this highly unique and desirable capability.

More recently the conversation has centered around comparisons of NVMe SSDs and IBM’s unique FlashCore Modules (FCM). Both devices employ the NVMe communication protocol, and both are very fast. However, IBM’s unique FCMs deliver longevity protection, and they typically prove to be significantly less costly on a per-terabyte (TB) basis due to the built-in hardware compression, which comes without a performance penalty.

There are some competitors who can address one or two things that FCMs do, but none of them can provide the long list of technology found in FCMs or the value that technology delivers to clients.

One of the most important things that clients should understand regarding IBM’s unique FlashCore Modules is that the numerous unique capabilities we provide are only possible because IBM has the ability to monitor every single flash cell in our devices. This enables IBM

to perform custom operations to optimize numerous aspects of the flash. Clients obtain more robust flash which helps them maintain continuous operations.

The innovative design of these custom-designed modules delivers numerous real-world business benefits as follows:

### *Extremely low latency*

Clients can experience read cache latency as low as 50 microseconds which helps remove bottlenecks in their workloads. FCM's enables low system latency by off loading data reduction and freeing the storage software from doing metadata management and garbage collection. The FCM has a significant amount of IP to minimize the effects of internal garbage collection and storing frequently accessed data in SLC.

### *Enhanced endurance*

A huge, albeit less obvious benefit of the FCMs is greatly enhanced flash endurance. There are many things that help deliver up to 7 times greater flash endurance than an industry-standard, commodity solid-state-drive (SSD). That translates to fewer issues for clients. It means that time doesn't have to be spent dealing with failing SSDs and drive rebuilds.

### *No impact compression technology*

Implemented completely in hardware with no processor intervention means that these added benefits do not slow down the rest of the storage system.

The compression technology originated with the IBM Mainframe group and has been adapted to work in IBM FlashCore flash controllers. It is performed as the first step in the inbound data path, and decompression is the last step in the outbound data path. This minimizes the amount of data written to flash which in turn helps extend the longevity of the flash. With the multi-level techniques for wear leveling, and the new ransomware detection, the IBM FCMs are truly a computational storage device.

**Note:** Compression is always active in the FCMs. It cannot be switched off.

## Data write protection

Data write protection, also known as *Error Correction Code (ECC)*, is implemented on top of compressed data and therefore across more data.This allows the delivery of even better performance.The error correction code (ECC) IBM employs utilizes a hard-decision decoding approach. This method offers several advantages:

- High Correction Strength: It boasts a very high capability to correct errors within the data.
- Reduced Read Latency: By using hard decisions, it eliminates the need for re-reading data, resulting in faster operation.

Compression and decompression are completely transparent above the FlashCore Module except for management of space.

## SLC and Smart Data Placement

SLC is used as a both a staging area before moving data to QLC and as permanent storage, varying based on physical capacity utilization.

At ~20% physical capacity utilization and below, logical data is kept almost completely in SLC. While the endurance and performance numbers are not based on SLC storage, the life of the

drive may be extended for low capacity utilizations. As capacity utilization increases beyond 20%, the FCM will move data stored in SLC blocks to QLC blocks and will transition these SLC blocks to QLC. The controller aims to transition blocks to QLC such that it does not significantly affect the performance of the user data being written from the host. If the FCM detects that it can't keep up with the host pace, such as high bandwidth sequential writes, it will lower the write amplification by bypassing the SLC and directing incoming host data directly to QLC to allow for the best possible write bandwidth on the FCM. At ~80% capacity utilization and above, user data will be almost completely stored in QLC.

Once data is moved from SLC to QLC read latency may be increased as QLC pages on average have a higher read latency then SLC. However, in real world applications, such as database transactions, workloads are largely skewed. This means that some percentage of the logical space of the drive sees more accesses than the remainder of the drive.

The FCM2, 3 and 4 use QLC NAND devices that have tiered read latency based on page type. The tiered latency of the different page types allows the FCM to place data that is read often in the fastest pages of NAND to match SLC speeds. Alternatively, it will place cold read data in the long latency QLC pages. This function is called *Smart Data Placement*.

Smart Data Placement is designed to allow SLC performance in QLC with real world applications.

IBM FlashCore Modules utilize the following features and functions:

- The FlashSystem data compression/decompression algorithm is a *Modified Dynamic GZIP* algorithm.
- Takes advantage of already existing LSA mapping.
- There is less data to transfer in back-end making up for small added latency.
- Decompression done in line with minimal latency addition.
- Data is decompressed alongside compression and checked bit by bit. Never any corruption!
- Data protection (ECC) is implemented on top of compressed data.
  - Allows garbage collection and other background data transactions to operate on compressed data.
- Compression and decompression completely transparent above the Flash module except for management of space.

## Variable Stripe RAID

IBM has implemented a technology called *Variable Stripe RAID (VSR)* on each FCM. Each individual module performs chip-level RAID on the Flash within itself. If one die fails in a chip stripe, then only the failed die is bypassed. The data is restriped across the remaining chips and no system rebuild is required. This is significant because VSR helps reduce maintenance intervals caused by flash failures. Even more importantly, VSR helps avoid performance robbing system-level intervention most of the time.

There are many business benefits associated with this technology. VSR non-disruptively protects data from a chip level failure. By dynamically re-striping data at a sub-chip level, IBM can ensure continuous business operations. *The bottom line is that VSR helps preserve the life of a company's flash while also providing data protection and performance*.

So, who else has Variable Stripe RAID? No one! IBM is the only vendor to deliver Variable Stripe RAID for multiple dimensions of RAID protection while maintaining peak performance.

The multiple dimensions comes from also factoring in system-level RAID protection. The good news is that many of the things that would normally require intervention by system-level RAID are not a problem for IBM solutions because they are dealt with at the module-level.

**Note:** RAID stands for redundant array of inexpensive (or independent) disks.

*Note that we no longer implement RAID 5 due to RAID technology higher up in the IBM Storage Virtualize layer. However, we utilize RAID 0 at the FCM level, and with the variable stripe aspect, this is a key part of the architecture.* This ability to have variable stripe widths even though RAID 0 is extremely important. When we do a Stripe across the die, the firmware can have any number of die in that stripe. So, if one block is retired or busy for some reason we can make a stripe with 19 or 18 or 17.

## IBM differentiation – FlashCore Modules

IBM has developed how to get about the same endurance out of *Quad-level Cell (QLC)* as out of *Triple-Level Cell (TLC)* and thus has better performance out of the QLC version than our TLC version.

- Compression accelerator done in the SSD – off loads an expensive software task.
  - *The only SSD on the market that has this!*
  - It is a computational storage platform using *Field Programmable Gate Array (FPGA).*
- The storage controller gets compression completely transparently. The data path runs exactly the same as without compression.
- There is no remapping of data and need for metadata for a log structured array.
- It is QLC only.
- The controller (aka FPGA) has no complex tasks resulting in improved sustainability.
- Each write in each FCM goes to a new unused area of the drive. This is not done at the Storage Virtualize layer. Garbage collection runs in the background to reclaim these flagged areas.

Figure 4 on page 12 shows the internal layout of the IBM FlashCore Module.

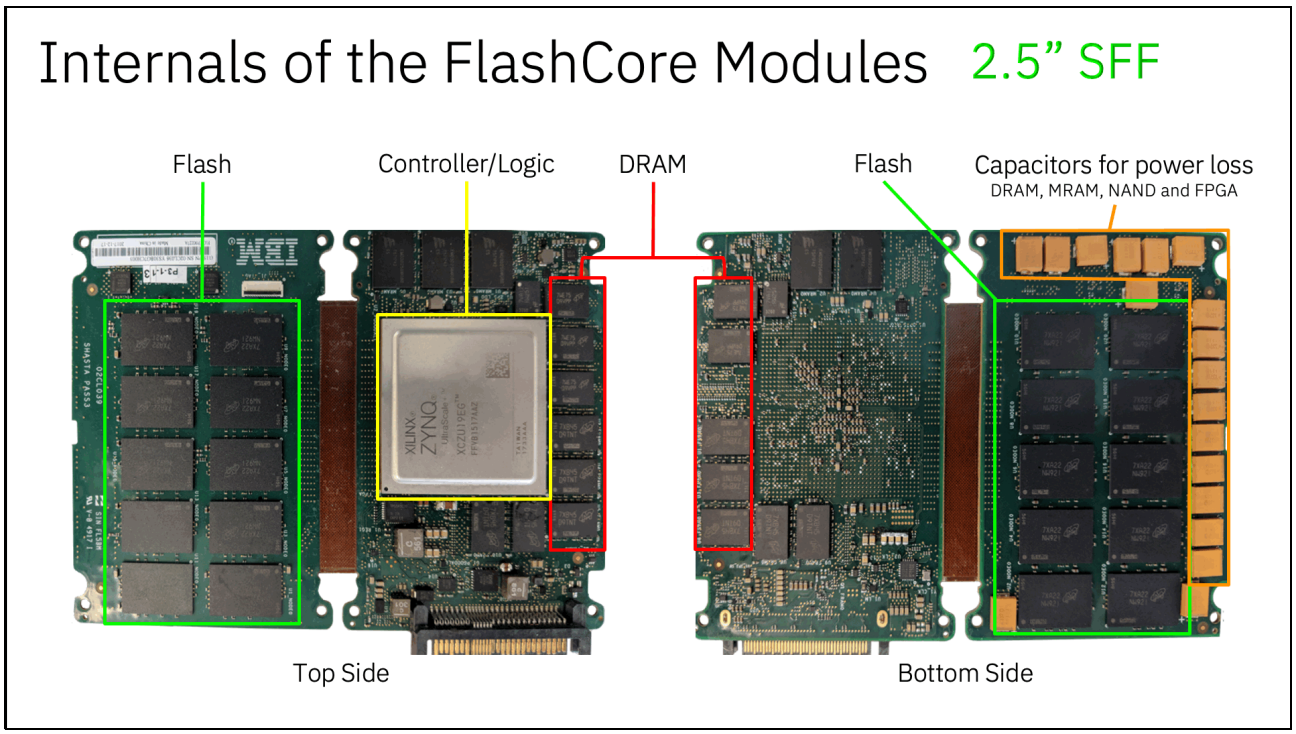


*Figure 4 Layout of the IBM FlashCore Module*

This is showing the interior of an IBM FlashCore Module (FCM). It is custom designed by IBM to meet the needs of a transforming storage market. The design employs IBM’s FlashCore Technology to solve problems that nearly all client's encounter.

Figure 5 on page 13 shows the internal components of the FCM and what each part does.

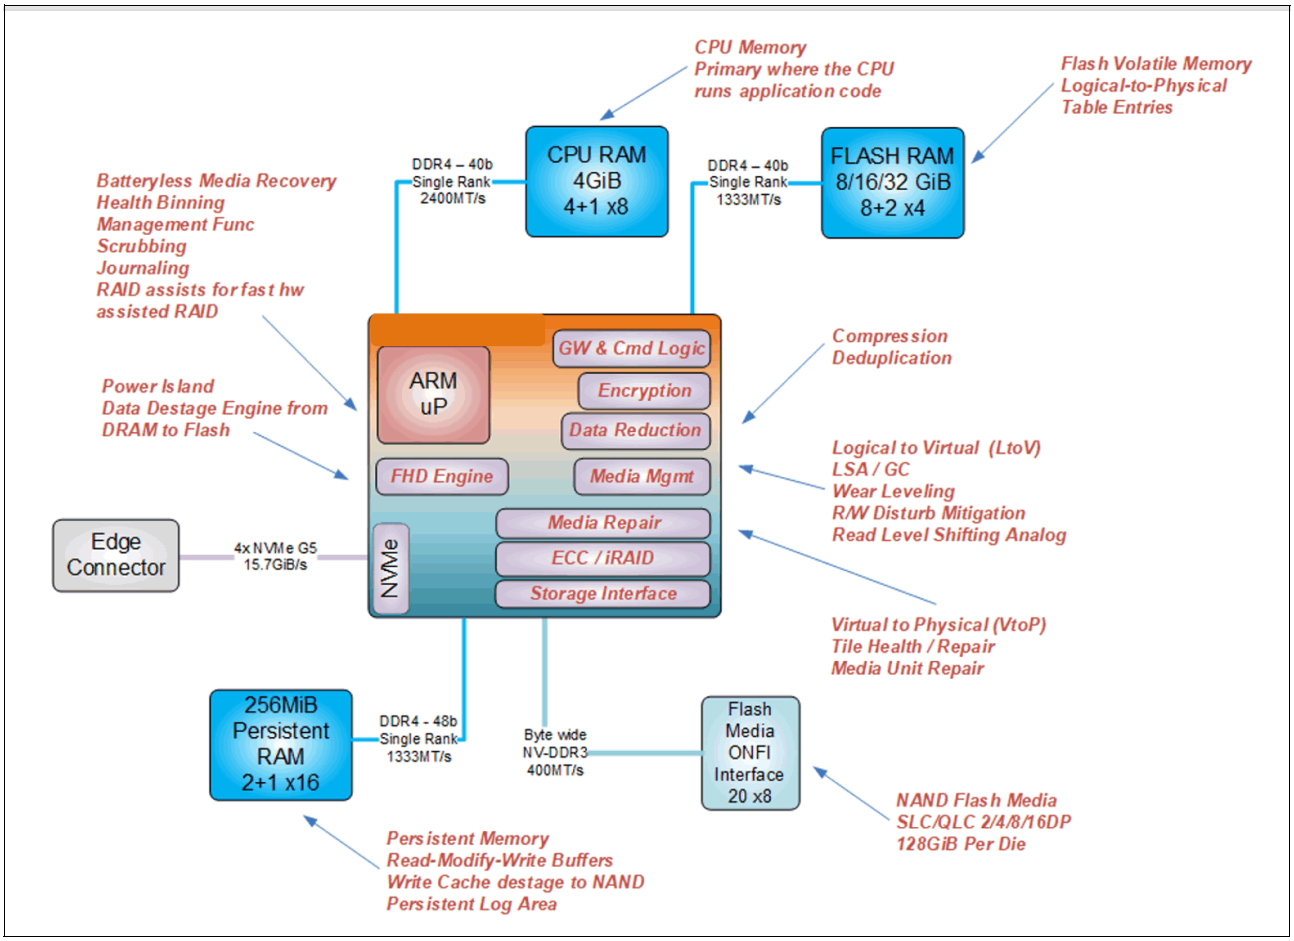


*Figure 5 FCM internal component description*

## FCM generations

This section describes the types of FCM drives that can be installed in the IBM FlashSystem control enclosures:

- FCM Generation 1
- FCM Generation 2
- FCM Generation 3
- FCM Generation 4

The different versions all have the same physical capacity, the difference is the effective capacity after the in-built hardware compression and encryption.

Each version of FCM drives have the following capacities (Physical size Terra Bytes usable (TBu)). See Table 2.

*Table 2 FCM type capacities*

| FCM module size | Physical size (TBu) |
|---|---|
| Small | 4.8 TBu |
| Medium | 9.6 TBu |
| Large | 19.2 TBu |

| FCM module size | Physical size (TBu) |
|---|---|
| XLarge | 38.4 TBu |

**Note:** The XL FCM drives require IBM Storage Virtualize V8.3.1 or later to be installed on the IBM FlashSystem control enclosure.

The following IBM FlashSystem products can support all three capacity versions of these drives:

- 9500**
- 9500R Rack Solution**
- 9200*
- 9200R Rack Solution*
- 7300**
- 7200*
- 5200**
- 5100*

* FCM3 was not supported in the 5100, 7200 and 9200.

** FCM4 is only supported in the IBM FlashSystem 9500, 7300 and 5200. *They are not supported in any of the expansion enclosures*.

Figure 6 shows a comparison of all the versions of FCMs and their individual characteristics.

| | Introduced in 2018 | | | Introduced in 2020 | | | | Introduced in 2022 | | | | Introduced in 2024 | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | FlashCore Module 1 | | | FlashCore Module 2 | | | | FlashCore Module 3 | | | | FlashCore Module 4 | | | |
| Capacity Point | Small | Medium | Large | Small | Medium | Large | X-Large | Small | Medium | Large | X-Large | Small | Medium | Large | X-Large |
| NAND Dies | 2DP | 4DP | 8DP | 2DP | 4DP | 8DP | 16DP | 2DP | 4DP | 8DP | 16DP | 2DP | 4DP | 8DP | 16DP |
| Flash Technology | Triple-level cell (TLC) | Triple-level cell (TLC) | Triple-level cell (TLC) | Quad-level cell (QLC) | Quad-level cell (QLC) | Quad-level cell (QLC) | Quad-level cell (QLC) | SLC / QLC | SLC / QLC | SLC / QLC | SLC / QLC | SLC / QLC | SLC / QLC | SLC / QLC | SLC / QLC |
| Usable Capacity | 4.8 TB | 9.6 TB | 19.2 TB | 4.8 TB | 9.6 TB | 19.2 TB | 38.4 TB | 4.8 TB | 9.6TB | 19.2 TB | 38.4 TB | 4.8 TB | 9.6TB | 19.2 TB | 38.4 TB |
| Effective Capacity | ~9.6 TB | ~19.2 TB | ~38.4TB | ~22 TB | ~22 TB | ~44 TB | ~88 TB | 22 TB | 28.8 TB | 57.6 TB | 115.2 TB | 22 TB | 28.8 TB | 57.6 TB | 115.2 TB |
| PCIe | Gen 3 | Gen 3 | Gen 3 | Gen 3 | Gen 3 | Gen 3 | Gen 3 | Gen 3 | Gen 3 | Gen 4 | Gen 4 | Gen 4 | Gen 4 | Gen 4 | Gen 4 |
| Read Bandwidth | 2,000 MB/s | 2,000 MB/s | 2,000 MB/s | 2,200 MB/s | 2,200 MB/s | 2,200 MB/s | 2,200 MB/s | 2,300 MB/s | 2,300 MB/s | 3,500 MB/s | 3,500 MB/s | 2,200 MB/s | 2,200 MB/s | 3,500 MB/s | 3,500 MB/s |
| Write Bandwidth | 500 MB/s | 500 MB/s | 500 MB/s | 690 MB/s | 690 MB/s | 690 MB/s | 690 MB/s | 750 MB/s | 750 MB/s | 1,500 MB/s | 1,500 MB/s | 750 MB/s | 750 MB/s | 1,500 MB/s | 1,500 MB/s |
| Read IOPS (16K) | | | | 65K | 65K | 65K | 65K | 100k | 100k | 150k | 150k | 200k | 200k | 200k | 200k |
| Latency Read/Write | 100µs/20µs | 100µs/20µs | 100µs/20µs | 77µs/µs | 77µs/µs | 77µs/µs | 77µs/µs | | | | | | | | |
| Power Envelope | 25W | 25W | 25W | 25W | 25W | 25W | 25W | 25W to 30W | 25W to 30W | 25W to 30W | 25Wto 30W | 25W to 30W | 25W to 30W | 25W to 30W | 25W to 30W |
| Intrusion/Ransomware detection | No | No | No | No | No | No | No | No | No | No | No | Yes | Yes | Yes | Yes |
| Most likely FlashSystem targets | FlashSystem 5100, FlashSystem 7200, FlashSystem 9110/9150 | | | FlashSystem 5100, FlashSystem 5200, FlashSystem 7200, FlashSystem 9200 | | | | FlashSystem 5200, FlashSystem 7300, FlashSystem 9200/9500/9500R | | | | FlashSystem 5200, FlashSystem 7300, FlashSystem 9500/9500R | | | |
| Capacity, Benefits and Values | Flash Endurance w/ TLC | | | Flash Endurance w/ QLC | | | | Flash Endurance SLC/QLC Hinting | | | | Ransomware Threat Flash Endurance SLC/QLC Hinting | | | |
| | Physical | Effective | | Physical | Effective | | | Physical | Effective | | | Physical | Effective | | |

*Figure 6 FlashCore Module summary and comparisons*

**Note:** All capacities are PCIe G4 with FCM4. With FCM3, the 4.8 TB and 9.6 TB drives were PCIe G3 and the 19.2TB and 38.4TB drives were PCIe Gen 4. All FCM1 and FCM2 drives were PCIe Gen3.

Figure 7 shows an FCM (NVMe) with a capacity of 19.2 TB.

The first generation of FCM drives were built by using 64-layer Triple Level Cell (TLC) flash memory and an Everspin MRAM cache into a U.2 form factor.

Compared to previous generations, FCM4 takes a significant leap in performance by utilizing 176-layer Quad Level Cell (QLC) NAND flash memory. FCM generations 2 and 3 relied on 96-layer QLC NAND. These later FCM generations employ a technique called *pseudo-Single-Level Cell (pSLC) caching*. This innovative approach reserves a portion of the flash memory and operates it in a mode that mimics SLC flash. This delivers several benefits:

- **Improved performance**: pSLC caching significantly boosts performance by enabling faster read and write operations.
- **Reduced latency**: Data access becomes quicker due to the faster nature of pSLC mode.
- **Dynamic read cache on-device**: The pSLC cache acts as a dynamic on-device buffer, further accelerating data retrieval..

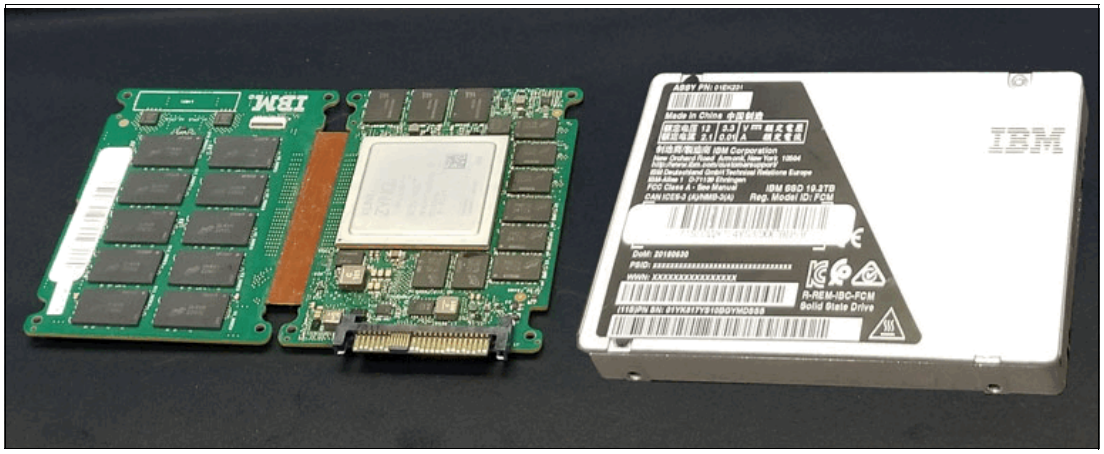


*Figure 7 IBM FlashCore Module (NVMe)*

Built for speed and efficiency, IBM FCM drives are the ideal solution for handling your most demanding workloads. Here's what sets them apart:

- **Parallel processing powerhouse:** Their high-parallelism architecture tackles complex tasks efficiently, delivering exceptional performance.
- **Optimized for 3D QLC NAND Flash**: Leveraging cutting-edge 3D QLC NAND technology, these drives offer significant performance gains and impressive storage capacity.
- **Advanced FPGA technology**: Updated FPGAs (Field-Programmable Gate Arrays) accelerate data processing, further boosting performance.
- **Minimized latency for compressed data**: The integrated read cache ensures smooth performance even when working with highly compressed data, minimizing lag.
- **Reduced write power consumption**: Efficient four-plane programming optimizes write operations, lowering power usage without sacrificing speed.

**Note:** FCM Generation 3 and 4 drives offer hardware-assisted compression up to 3:1.

At the time of writing, the FCM3 is currently undergoing FIPS 140-3 Level 2 certification, FCM4 will be undergoing the same certification shortly. Those certifications take 1-2 years to complete.They are currently both FIPS 140-3 Level 1complaint.

All FCM drives carry IBM Variable Stripe RAID (VSR) at the FCM level and use DRAID to protect data at the system level. VSR and DRAID together optimize RAID rebuilds by off-loading rebuilds to DRAID, and they offer protection against FCM failures.

## IBM FlashCore technology

At the heart of the IBM FlashSystem is IBM FlashCore technology (if the IBM FlashCore Module NVMe type drives are ordered), which consists of the following key elements:

- Hardware-accelerated architecture that is engineered for flash, with a hardware-only data path.
- The IBM FlashSystem data compression and decompression algorithm is a modified dynamic GZIP algorithm. Because it is implemented completely in hardware, no processor intervention is required.
- IBM FlashCore Modules, which are designed for low latency, density, and reliability.
- IBM Advanced Flash Management, which improves flash endurance over standard implementations without sacrificing latency.
- Depending the machine type and model of the control enclosure, the largest enclosure can contain up to 48 FCM drives.

Figure 8 on page 16 shows IBM FlashCore technology.

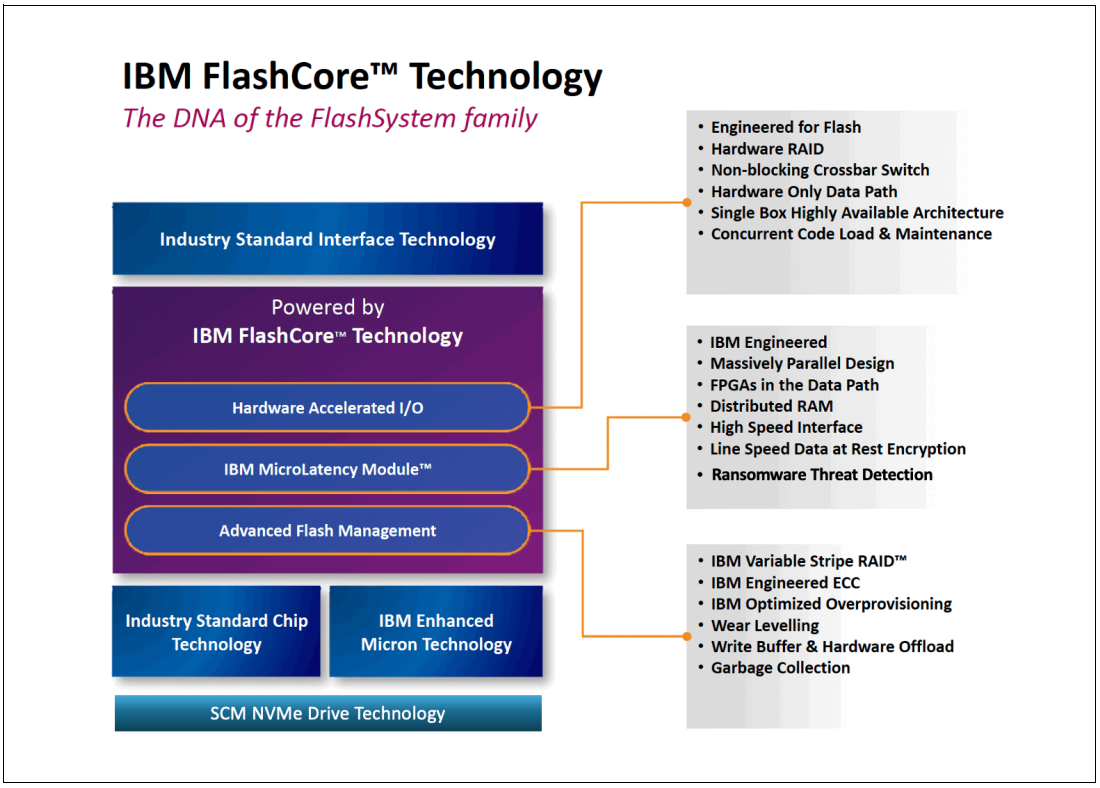


*Figure 8 IBM FlashCore technology*

For more information about IBM FlashCore technology, see this IBM Video.

## Data reduction tools

Compression and de-duplication are key features of IBM FlashSystem with FCM technology, significantly reducing storage requirements. To help customers estimate potential space savings, IBM offers the Comprestimator and Data Reduction Estimator Tool (DRET) tools. These tools can recognize the patterns of the client data, and estimate the compressibility of data per volume.

IBM FlashSystem models are supported by the IBM Comprestimator (which is available as a stand-alone tool and is found in the FlashSystem GUI) and also the stand-alone Data Reduction Estimator Tool (DRET). The DRET is a host-based application that the user uses to estimate the amount of compression and de-duplication on the IBM FlashSystem for specific workloads.

For more information about DRET, see Data Reduction Estimator Tool.

For more information about IBM FlashSystem Comprestimator, see IBM Comprestimator.

In order to determine the optimal storage requirements, select your preferred data reduction method and utilize the provided tools for storage usage estimation. Figure 9 shows how to start the Estimate Compression Saving option from the FS 9100 GUI.

From the main menu, click **Volumes** → **Volumes** and then, select one volume. Then, right-click the mouse to open the pop-up menu. From this menu, select **Capacity Savings** → **Estimate Compression Saving**.

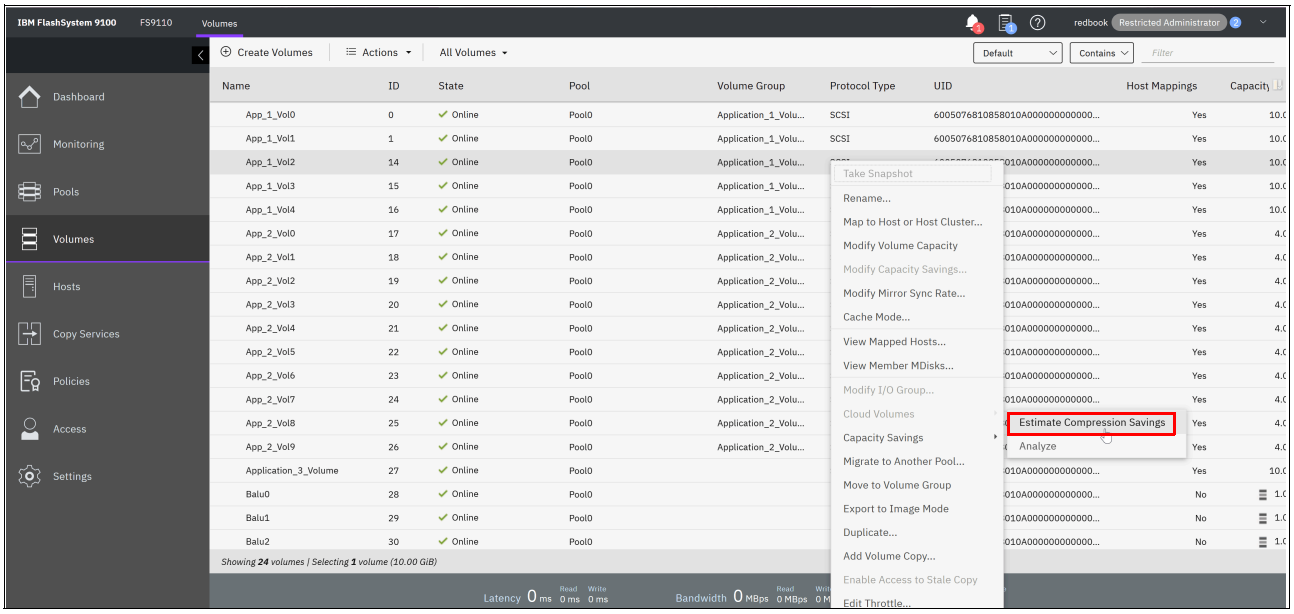


*Figure 9 Estimate Compression Savings from the GUI*

To estimate compression and savings on FCM and DRP approaches:

- IBM FlashCore Module - IBM FlashCore Module Compression:
  - Use the FCM option.
  - Do not use the Estimate Compression Saving option in the GUI to calculate the IBM FlashCore Module savings.

- Data Reduction Pool compression:
  - Use the DRP option.
  - Workloads that are on any IBM Storage Virtualize platforms can use the Estimate Compression Saving option in the GUI.
- Data Reduction Pool compression and deduplication:
  - IBM Comprestimator and Data Reduction Estimator Tool shows the savings for thin-provisioning, compression, and deduplication.
  - IBM Comprestimator and Data Reduction Estimator Tool reads entire volumes to identify de-duplicated data, so it takes longer to run.

For more information about data reduction pool compression and setup, see *Introduction and Implementation of Data Reduction Pools and Deduplication*, SG24-8430.

# Manageability and security

The following are the manageability and security features:

- Advanced security for data at rest with hardware-accelerated AES-XTS 256 encryption.
- IBM obtained FIPS 140-3 Level 1 certification for IBM FlashCore Modules generation 3 in the IBM FlashSystem products. This level of certification is fully supported by the FCM-3 modules in the system.
- FCM4 also supports RSA and CRYSTALS Kyber cryptography.
  - Secure Key Passing (SKP) data is encrypted twice, once by each cipher.
- CRYSTALS Kyber is a Quantum Safe Cryptography (QSC) algorithm.

## Encryption

Like its predecessors, IBM FlashSystem data encryption is based on the industry standard AES-XTS 256 encryption, as defined in the IEEE 1619-2007 standard and NIST Special Publication 800-38E as XTS-AES-256. The data encryption key is protected by a 256-bit AES key wrap of a key that is derived from the access key that is stored on the USB flash drive. The wrapped key is stored in the system in nonvolatile form.

**Note:** For more informatio about FlashSystem encryption you can refer to the following link.

### Self-encrypting drives

The IBM FlashCore Modules in the IBM FlashSystem control enclosure are self-encrypting drives (SEDs). With SEDs, you can encrypt the data on the drive within the hardware.

These types of flash drives include the following features:

- Encryption of data is done in the electrical circuit of the drive; therefore, it is not affected by performance issues from software encryption.
- Data Encryption Keys (DEK) never leave the confines of the SED, and are never loaded into CPU or memory.
- You can perform a fast cryptographic erasure of a SED by using a single CLI command to replace the DEK or revert the entire device to factory settings.

- Supports a security feature that is called *auto-lock*, which protects against someone plugging your drive into another system and accessing your data.
- Drives automatically lock themselves on power loss and require an access key at start time to unlock and allow I/O operations.
- If an SED drive is removed from a system with encryption and placed in another system, the drive data is not readable. The system posts an error message that indicates it is locked. The only way to use the drive is to format it. This formatting also performs a cryptographic erase by removing any encryption keys; therefore, all of the data on the drive is destroyed.

## Combining system encryption with self-encrypting drives

For control enclosures that support NVMe architecture, NVMe-attached drives are self-encrypting and self-compressing. With SEDs that use NVMe architecture, data encryption is completed in the drive. Data encryption keys remain on the drive without being stored in system memory.

In addition, the system supports automatic locks of encrypted drives when the system or drive is powered down. When the drive or system restarts, a master key is required to unlock the drive and continue I/O operations.

Because the encryption of data is done in the electrical circuit of the drive, it is not affected by any potential performance issues from software encryption.

**If SEDs are encrypting the data, why do you need to enable system encryption in IBM Storage Virtualize?** The answer is that you can use SEDs without enabling encryption on the system, but SEDs are unlocked by default at start time, unless configured with extra protection.

System-level encryption in IBM Storage Virtualize allows you to use USB flash drives or IBM Security® Guardium® Key Lifecycle Manager to manage access to encrypted objects on the system. This feature ensures that when a system is powered, this extra encryption key is required to read the data on the drives.

Consider the following points:

- SEDs are always encrypting, and you cannot stop them from being encrypted.
- You can use SEDs without enabling encryption on the system, but SEDs are unlocked by default unless they are configured with extra protection.
- With system encryption in IBM Storage Virtualize, you can use USB flash drives or IBM Security Guardium Key Lifecycle Manager to manage access to encrypted objects on the system.
- Software in the operating system is required to manage an access key that can be used to lock and unlock the SEDs and bring them online for I/O.

Therefore, the best solution is to use the SEDs with the Encryption Enablement Pack and USB or IBM Security Guardium Key Lifecycle Manager type encryption, or a mixture of both. This configuration ensures the maximum level of encrypting for your data that is on the system.

## Secure drive erasure process

The IBM FlashSystem system running IBM Storage Virtualize V8.5.0 or higher provides methods to securely erase data from a drive or boot drive when a control enclosure is decommissioned or before a drive is removed from the system during a repair activity.

Secure data deletion effectively erases or overwrites all traces of data from a data storage device. The original data on that device becomes inaccessible and cannot be reconstructed. You can securely delete data that is on individual drives and on a boot drive of a control enclosure. The methods and commands that are used to securely delete data enable the system to be used in compliance with European Regulation EU2019/424.

The following types of drives can be used for this process:

- Expansion enclosure SAS SSDs.
- NVMe drives (IBM FlashCore Modules and industry standard).
- Control enclosure node canister SSD boot drives.

The methods that the system uses to securely delete data from the drives varies according to the CLI commands that each type of drive can support. The completion time for the erase procedure also varies, depending on the amount of data and the method that is used to delete the data. In each case, when the operation completes, the result is that the data on the drive effectively becomes impossible to access.

Table 3 lists the types of erasure, the methods used, and the time taken.

*Table 3 Comparison of methods to securely delete data from drives*

| Priority | Deletion type | Method | Completion time |
|---|---|---|---|
| 1 | Cryptographic erase | Changes the encryption key and makes the data inaccessible. | Instant |
| 2 | Block erase | Quickly raises and lowers the voltage level of the storage element. Physical blocks are altered with a vendor-specific value. | Fast |
| 3 | Data Overwrite | Replaces the existing data with random data. | Slow |

The methods that are used to securely delete data vary according to manufacture, drive type, and drive firmware. For more information, see the documentation that is provided by the drive manufacturer.

If a drive supports more than one data deletion method, the system uses the highest-priority method.

For more information about the CLI commands that are used to run this secure erase function, see this IBM Documentation web page.

## Reliability, availability, and serviceability

If an IBM FlashCore Module failure occurs, critical customer applications can remain online while the defective module is replaced because a spare module is available. IBM Variable Stripe RAID is a patented IBM technology that provides an intra-module RAID stripe within each flash module. Variable Stripe RAID technology helps reduce downtime, and maintains performance and capacity during partial or full flash chip failures. Now with the new ransomware threat detection, it utilizes a patch apply process so no code load required

## IBM FlashCore Modules NVMe drive options

IBM FlashCore Modules combine IBM MicroLatency technology, advanced flash management, and reliability into a 2.5-inch SFF NVMe with built-in, performance-neutral hardware compression and encryption.

The following FCM NVMe flash drives are available (feature codes are machine specific - shown are for a IBM FlashSystem 9500):

- (#AHS9): 4.8 TB NVMe FlashCore Module 3.
- (#AHSA): 9.6 TB NVMe FlashCore Module 3.
- (#AHSB): 19.2 TB NVMe FlashCore Module 3.
- (#AHSC): 38.4 TB NVMe FlashCore Module 3.

For the FCM4 we have some extra feature codes as follows:

- (#AHSE) 4.8 TB NVMe FlashCore Module 4.
- (#AHSF) 9.6 TB NVMe FlashCore Module 4.
- (#AHSG) 19.2 TB NVMe FlashCore Module 4.
- (#AHSH) 38.4 TB NVMe FlashCore Module 4.

Consider the following points regarding limitations and drives:

- IBM FlashCore Modules:
  - Maximum 48 NVMe Drives. Normal ones: max 128.
  - Minimum-Maximum member drives per DRAID-6 array 6-128.
  - Minimum-Maximum member drives per DRAID-6 array (NVMe drives) 6-48.
  - Minimum-Maximum member drives per DRAID-1 array 2-16.
- IBM FlashCore Modules in the same DRAID array must be of the same capacity.

# New FCM developments

Throughout the development of the FCM technology we have gone through various versions of FCM which were designated FCM1, FCM2, FMC3 and now FCM4.

Each disassociation was a result of the FCM technology rather than the capacity as all of these variations had the 4 capacity sizes. The technology changes meant better access times and improved throughout.

### Inline data corruption detection (IDCD)

IDCD is a real-time data corruption mechanism that utilizes patterns seen, of data structure where something changes with the rates of data access both in reading and writing. The IDCD feature was made available in the IBM Storage Virtualize code v 8.6.0, where the detections of the potential ransomware attack was calculated in the storage virtualize layer.

In the majority of cyber-attacks they take the form of ransomware types, where the data on the storage array is encrypted by the attackers, thus making it unreadable by the legitimate owner. By looking at the read/encrypt/write actions on the arrays we can alert the customer to the fact that this type of activity is happening.

However, it is important to note that similar activity patterns could occur for legitimate reasons. For instance, a customer might choose to encrypt specific data stored on the arrays.

## Ransomware threat detection

The latest release of the IBM Storage Virtualize code v 8.6.3 and above, coupled with the new FCM4, provide even great real-time protection to ransomware detection The FCM4 his has some extra components included, towards assisting fighting cybercrime in the shape of a specialized set of hardware to support it. IBM has announced inline corruption detection that uses AI and ML to help detect ransomware attacks. This announcement and implementation is a combination of the new FCM4 technology and IBM Storage Virtualize code version 8.6.3 and above.

**Note:** The ransomware threat detection enablement requires FCM4 drives running FCM firmware 4.1 or higher, plus IBM Storage Virtualize code v 8.6.3 or higher.

There is also the ability to upgrade up to IBM Storage Virtualize code v8.6.3 and FCM firmware 4.1 with an existing array, assuming you initially created your array with FCM4's and IBM Storage Virtualize code v8.6.2 or later.

FCM4 is computational storage and adheres to the industry standards that are now specified, so we can easily adapt the functions on the FCMs in the field, through simple updates, and, as these are not in the data stream, this does not affect the performance of the FCM drives.

## New ransomware threat detection within Storage Virtualize

So, what is entropy?

- *File entropy* measures the randomness of the data in a file and is used to determine whether a file contains hidden data or suspicious scripts.
- File entropy is measured on a scale from 0 to 8, with 0 indicating no randomness (like a plain text file with repeated patterns) and 8 indicating maximum randomness (like a random noise file or a highly encrypted file).
- The more a unit can be compressed, the lower the entropy value; the less a unit can be compressed, the higher the entropy value.
- Entropy is used to detect highly random data, such as encrypted data written in by ransomware.

Figure 10 on page 23 shows the entropy patterns versus data throughput and the changes in the entropy count, that could indicate a ransomware attack.

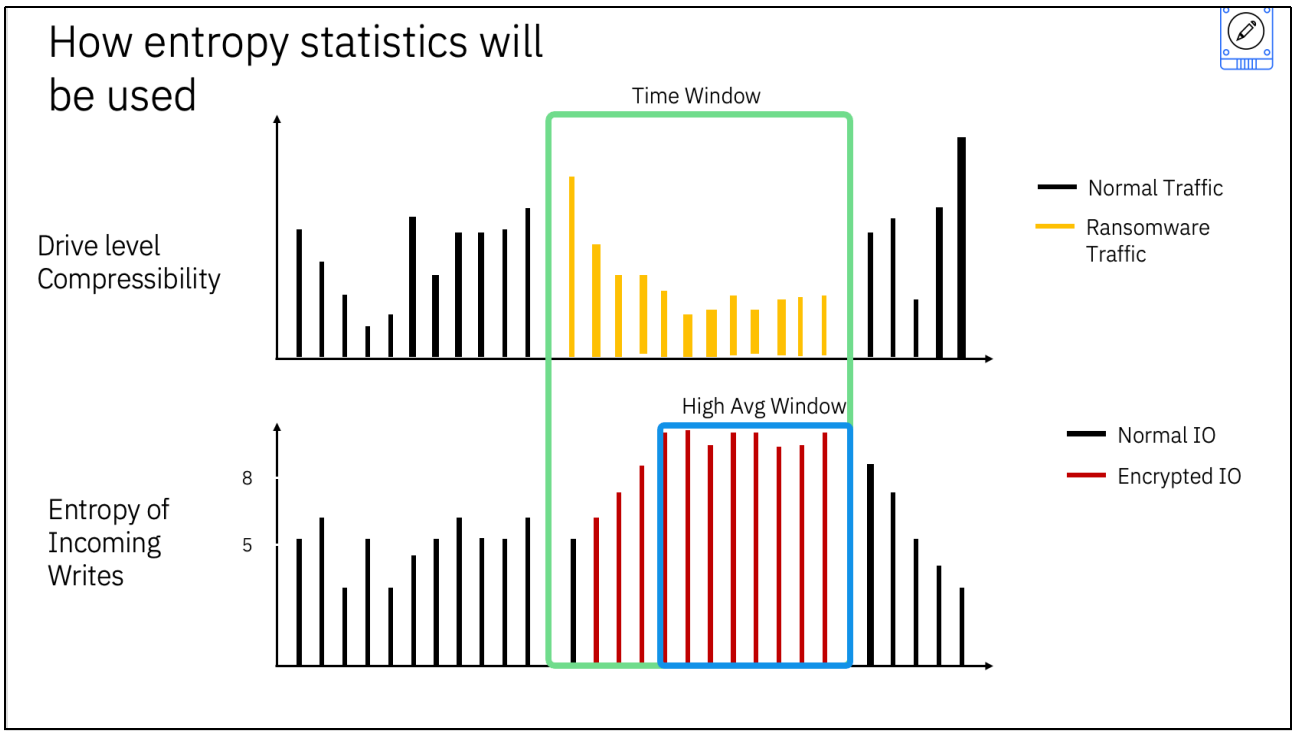


*Figure 10 Entropy count patterns*

In this example, the system has detected a decrease in compressibility and an increase in the entropy of incoming writes. This may be indicative of a cyber-attack.

IBM Storage Virtualize software will sample Entropy on every volume and send that back to Storage Insights Pro.

1. Detailed compressibility results are captured and sent back to IBM more frequently.
2. IBM Storage Insights Pro will analyze these statistics.
3. Alerts will be raised if a workload anomaly has been detected. It is like encryption turned on in an application.
4. Alerts will be tech previewed and clients can opt in.
5. The alerts will also go to support, if call home is enabled.

Figure 11 on page 24 shows the characteristics found in an I/O trace from ransomware attack and example of the entropy stages during an attack. This example shows the “Wannacry” virus patterns.

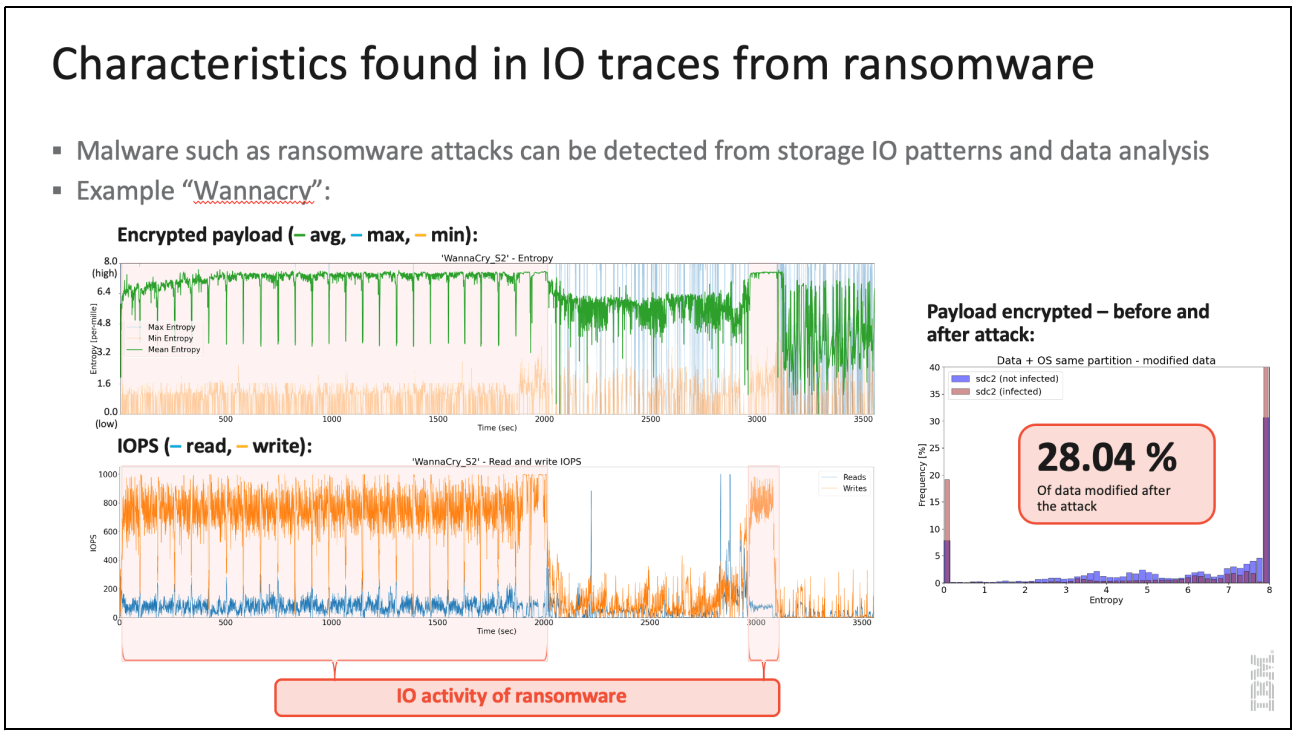


*Figure 11 Characteristics found in an IO trace from ransomware*

Figure 12 on page 24 shows the FlashSystem Ransomware Detection Logical Model and how the entropy detection is then handled, firstly by the IBM Storage Virtualize code and then passed to IBM Storage Insights Pro for alerting the customer and, if call home configured, direct to IBM support.

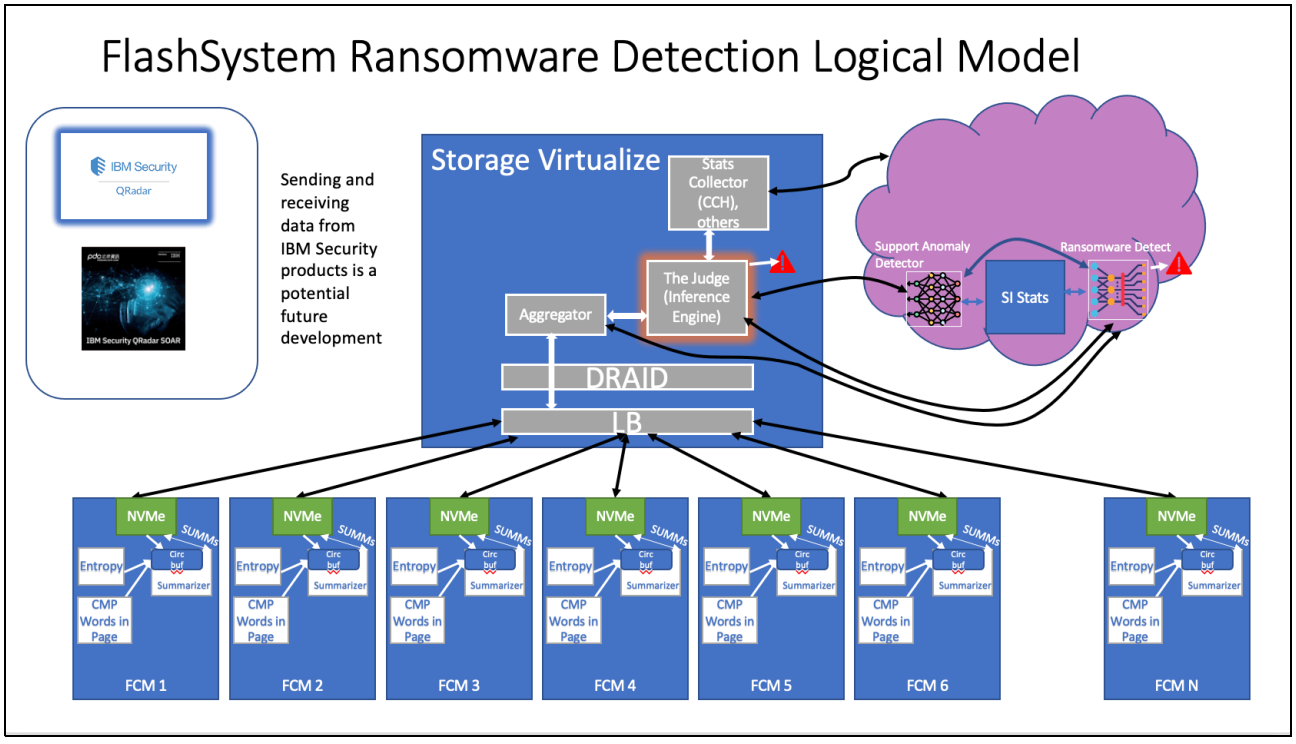


*Figure 12 FlashSystem Ransomware Detection conceptual model*

The FCM4s and the inference engine are informing us of vital signals to send back to Storage Insights Pro:

- These signals will be analyzed using AI Ops Machine Learning.
- This analysis will leverage idle network resources (bandwidth) while prioritizing critical system operations to avoid impacting host performance.
- FCM4 engine will take real time signals and summarize them.
- IBM Storage Virtualize will take all this information and collate it – and send summaries to IBM Storage Insights.
- There will be an inference engine in every FlashSystem NVMe system.
  - These engines will be fed information from Machine Learning models on anomalies and ransomware attacks.
  - They will learn what is normal for that system.
  - They will detect anomalous and dangerous behavior and take action.
  - This inference engine will be regularly updated.

## Futures of FCMs

IBM continues to improve the FCM family with the indirection on the FCM4 in 2024. This builds on the technology of the previous FCMs with the following enhancements:

- 2.5” dual ported U.2 NVMe form factor.
- Contains the Flash Translation Layer, Flash and accelerators for an external storage controller.
- 176 Layer Micron QLC.
- New Computational Storage Engines.

As technology improves and the integration of hardware and software within the FCM becomes tighter, IBM will strive to encompass these enhancements into the future generations of FCMs and implementation into future IBM FlashSystem products.

# Related information

For more information, see the following resources:

- *IBM FlashCore Module Cryptographic Erase*

  https://www.redbooks.ibm.com/abstracts/redp5529.html
- *Introduction and Implementation of Data Reduction Pools and Deduplication, SG24-8430*:

  http://www.redbooks.ibm.com/abstracts/sg248430.html
- IBM Documentation for IBM FlashSystem:

  https://www.ibm.com/support/knowledgecenter/en/search/flashsystem
- IBM FlashSystem 9500 product page:

  https://www.ibm.com/products/flashsystem-9500
- IBM Offering Information page (announcement letters and sales manuals):

  https://www.ibm.com/docs/en/announcements
- IBM Storage Virtualize FAQ

  Details about the IBM Storage Virtualize products, covering IBM FlashSystem family and SAN Volume Controller:

  https://www.ibm.com/downloads/cas/2DWAMWRB
- IBM FlashSystem Family Overview FAQ

  Overview of the IBM FlashSystem family with guidance on how to select the product that is right for you:

  https://www.ibm.com/downloads/cas/9OGKVW2R
- IBM FlashWatch FAQ

  Guidance for the IBM FlashWatch programs:

  https://www.ibm.com/downloads/cas/YVMYPEDE

# Authors

**Hartmut Lonzer** is Storage Advisory Partner Technical Specialist for DACH and SAN Offering Manager for DACH. His main focus is on the IBM FlashSystem Family and the IBM SAN Volume Controller. His experience with the IBM SAN Volume Controller and IBM FlashSystem products goes back to the beginning of these products. Hartmut has been with IBM in various technical and sales roles now for 46 years.

**Jon Herd** is an IBM Senior Executive Advocate working for the TLS EMEA Remote Technical Support and Client Care team in Germany, advising customers on a portfolio of IBM storage products, including IBM FlashSystem products. He also leads special projects for senior and executive management and is the SME lead for new product introduction in TLS EMEA. Jon has been with IBM for more than 49 years, and has held various technical roles, including Europe, Middle East, and Africa (EMEA) level support on mainframe servers and technical education development. He has written many IBM Redbooks® publications about IBM FlashSystem products and is an IBM Redbooks Platinum level author. He holds IBM certifications in Product Services at a Thought Leader L3 level, and Technical Specialist at an experienced L1 level. He is also a certified Chartered Member of the British Computer Society (MBCS - CITP), a Certified Member of the Institution of Engineering and Technology (MIET), and a Certified Technical Specialist of the Open Group (TOG).

**Vasfi Gucer** is works as the Storage Team Leader on the IBM Redbooks Team. He has more than 30 years of experience in the areas of systems management, networking hardware, and software. He writes extensively and teaches IBM classes worldwide about IBM products. His focus has been primarily on storage, cloud computing, and cloud storage technologies for the last 8 years. Vasfi is also an IBM Certified Senior IT Specialist, Project Management Professional (PMP), IT Infrastructure Library (ITIL) V2 Manager, and ITIL V3 Expert.

Thanks to the following people for their contributions to this project:

- Evelyn Perez
- James Whitaker
- Paul Edmonds
  **IBM UK**

- Luis F Lopez
- Edgar Daniel Gutierrez Haro
  **IBM Mexico**

- Brent Yardley
- Roger Kasten
- Andy Walls
- Brian Sherman
- Long Nguyen
  **IBM USA**

# Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

**ibm.com**/redbooks/residencies.html

# Stay connected to IBM Redbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

# Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
*IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

 

# Trademarks

REDP-5725-00

ISBN DocISBN

Printed in U.S.A.