Computational Storage:
Deduplication aided by Compression

Presenter: Tim Fisher → FCM Lead Architect, STSM, Master Inventor

linkedin.com/in/timothy-fisher-ibm

the Future of Memory and Storage

Enterprise Storage

What is an “Enterprise” Storage Appliance

- Cost
- Reliability
- Performance
- Security
- Storage
  - Data Reduction

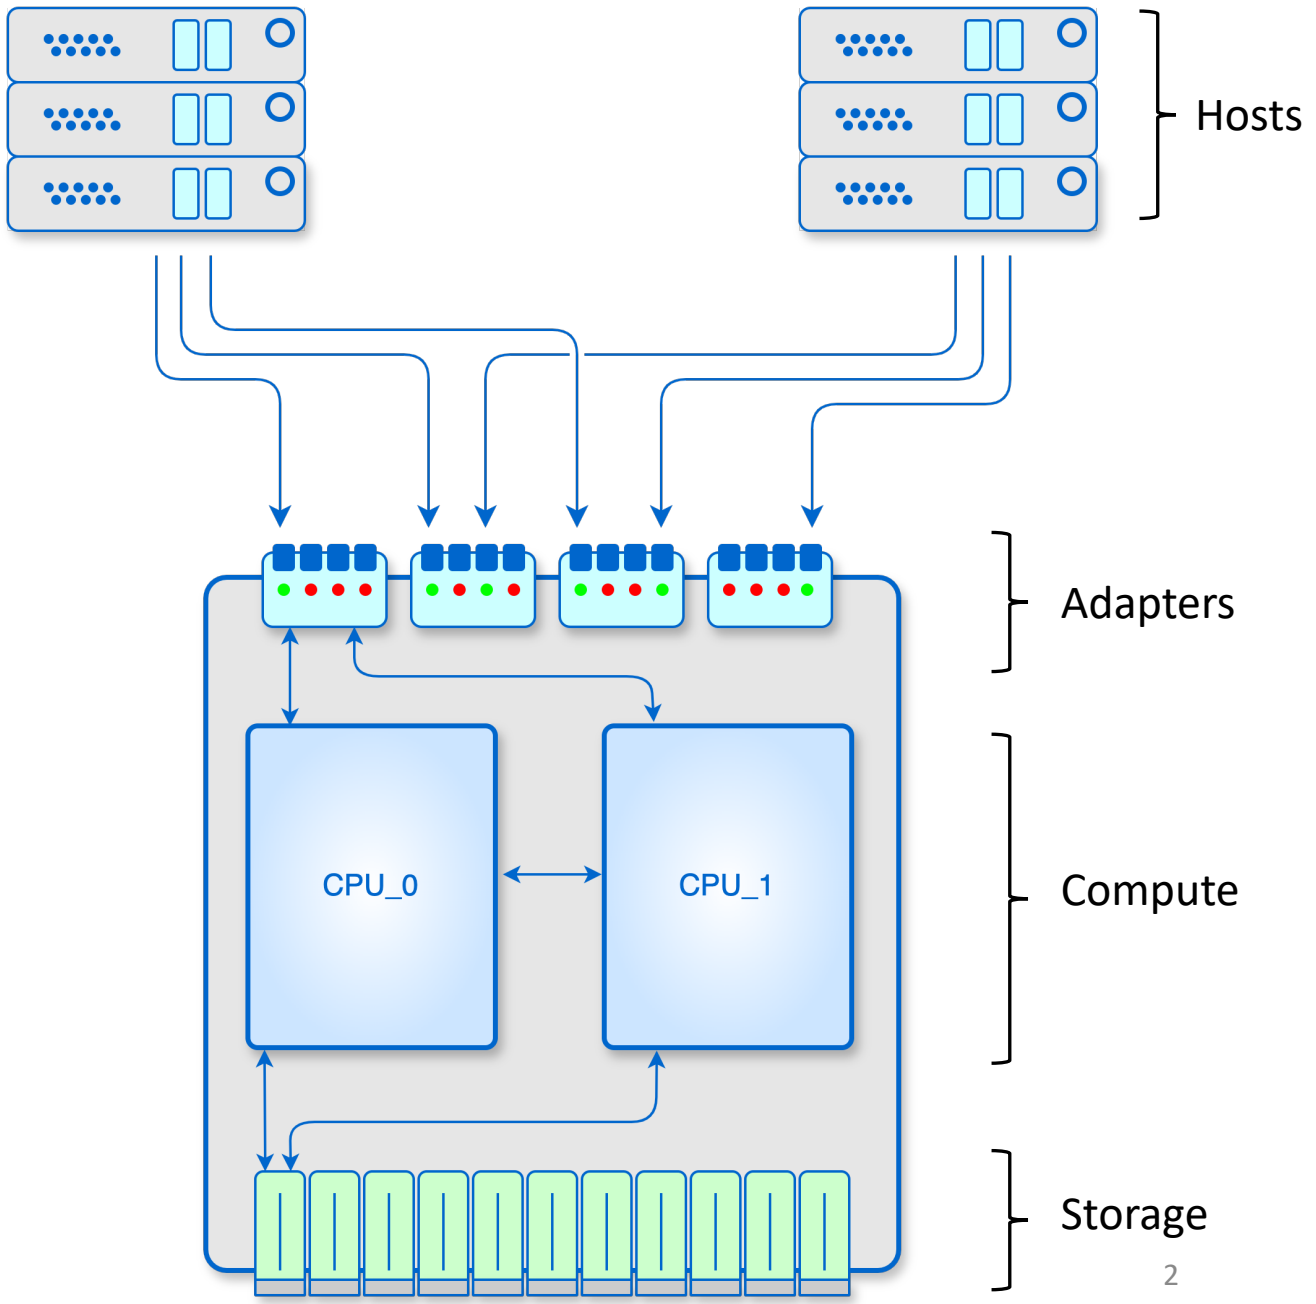

Data Reduction

- Compression
  - FCM CS Hardware Compression
  - Shipping in FlashSystems Storage products since 2017!
- Deduplication
  - FCM CS Deduplication Assists (Computational Storage Case Study)

FCM Compression

Benefits

- Compression in the FCM offers tremendous performance benefits
- The storage appliance uses ZERO CPU cycles to compress data
- The storage appliance uses ZERO DRAM for metadata to track compressed pages
- The storage appliance uses ZERO system bandwidth for reading and modifying compressed blocks on overwrites
- All compression metadata is contained completely in the FCM using the existing LPT tables
- Garbage Collection in the FCM is done with compressed data

Challenges

- Storage capacity management

Without FCM Compression

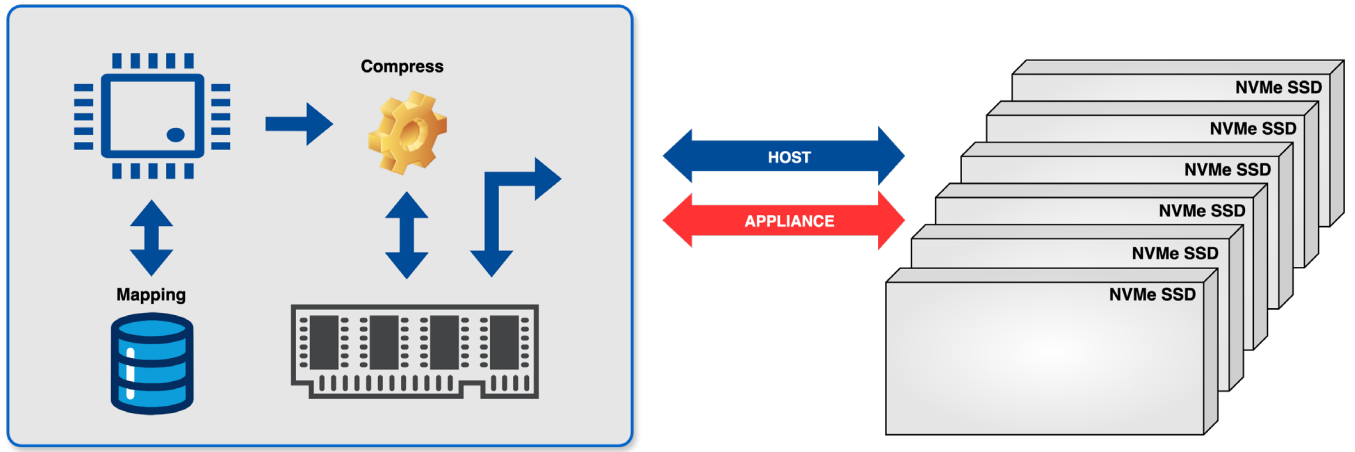

With FCM Compression

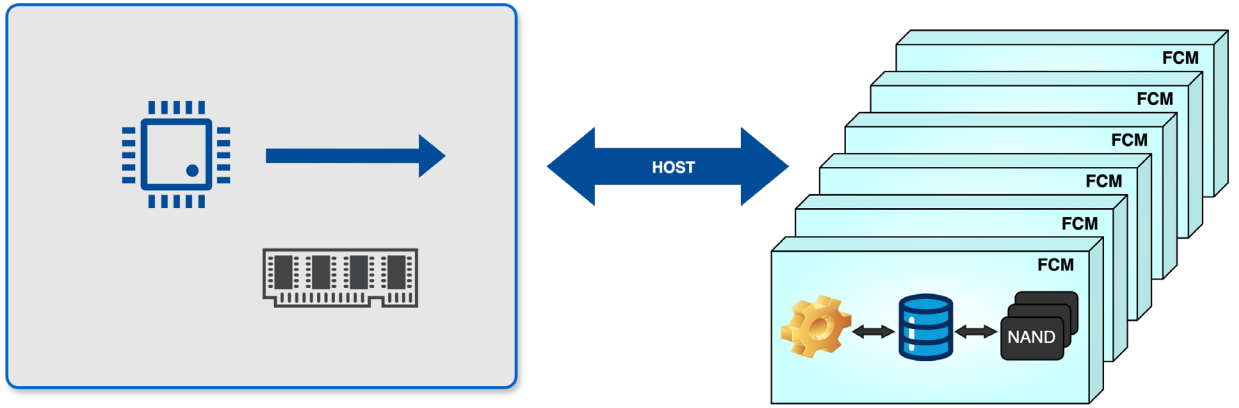

Deduplication

Benefits

- Enterprise storage appliances may have many volumes or partitions and may have multiple copies of data that can reduced to provide the user additional capacity
- Applications utilizing the storage appliance may write multiple copies of the same data which can also be reduced to provide additional capacity

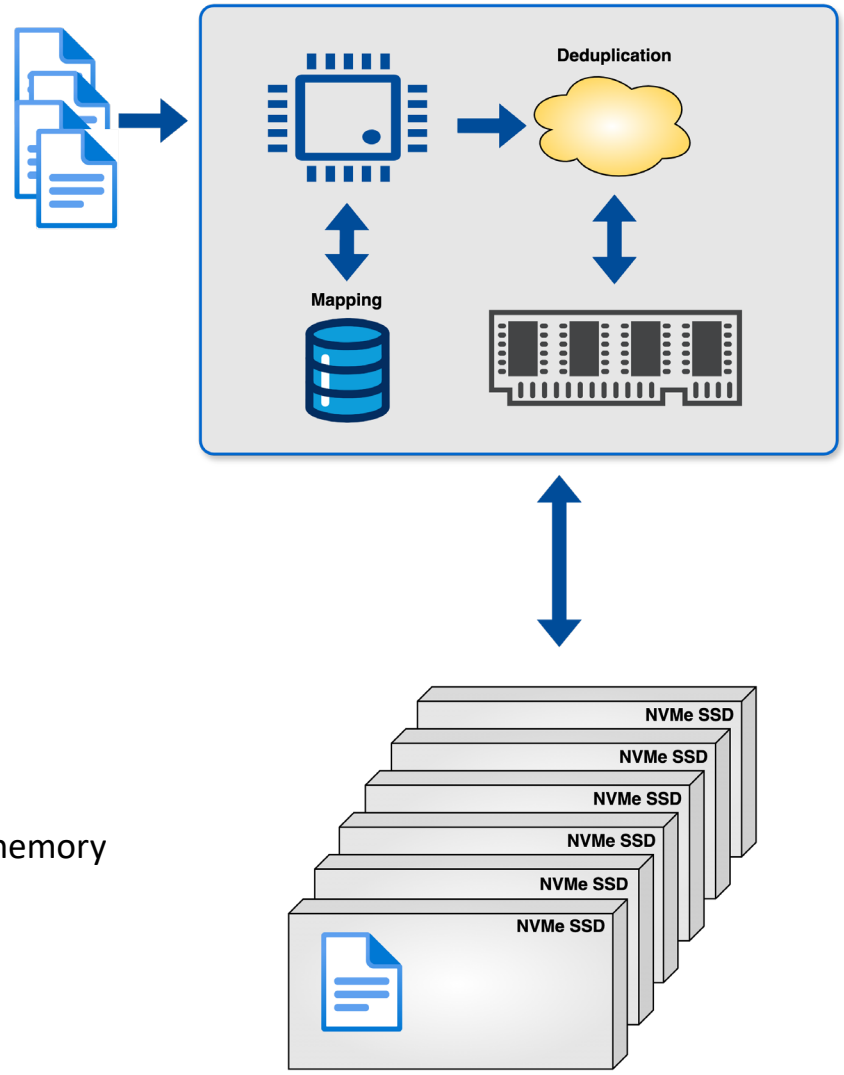

Challenges

- Enterprise storage appliances can have a LOT of capacity, and this will continue to grow
- Hash generation and storage may require a LOT of CPU cycles and memory
- Hash analysis may require a LOT of CPU cycles
  - Strong hash algorithms require 64+ rounds to produce a hash
- Deduplication metadata may require a LOT of memory
  - Example: 32B hash for every 4KB of user data, over 500TB array can require 1000s of GBs of memory
- As deduplication ratio increases, performance may degrade

How can the FCM help?

FCM SSD Overview

- FCM is an FPGA based controller design utilizing AMD/Xilinx Versal technology
- The FCM's primary function is an Enterprise QLC SSD that includes compression, encryption, high-performing FTL, and other features
- This computational storage use case is designed to not interfere with the storage performance or function

Computational Components

- While there are many other components that make up the hierarchy of our controller, we will focus on the roles of the APU(s), RPU(s), and programmable logic as they pertain to the computational storage infrastructure

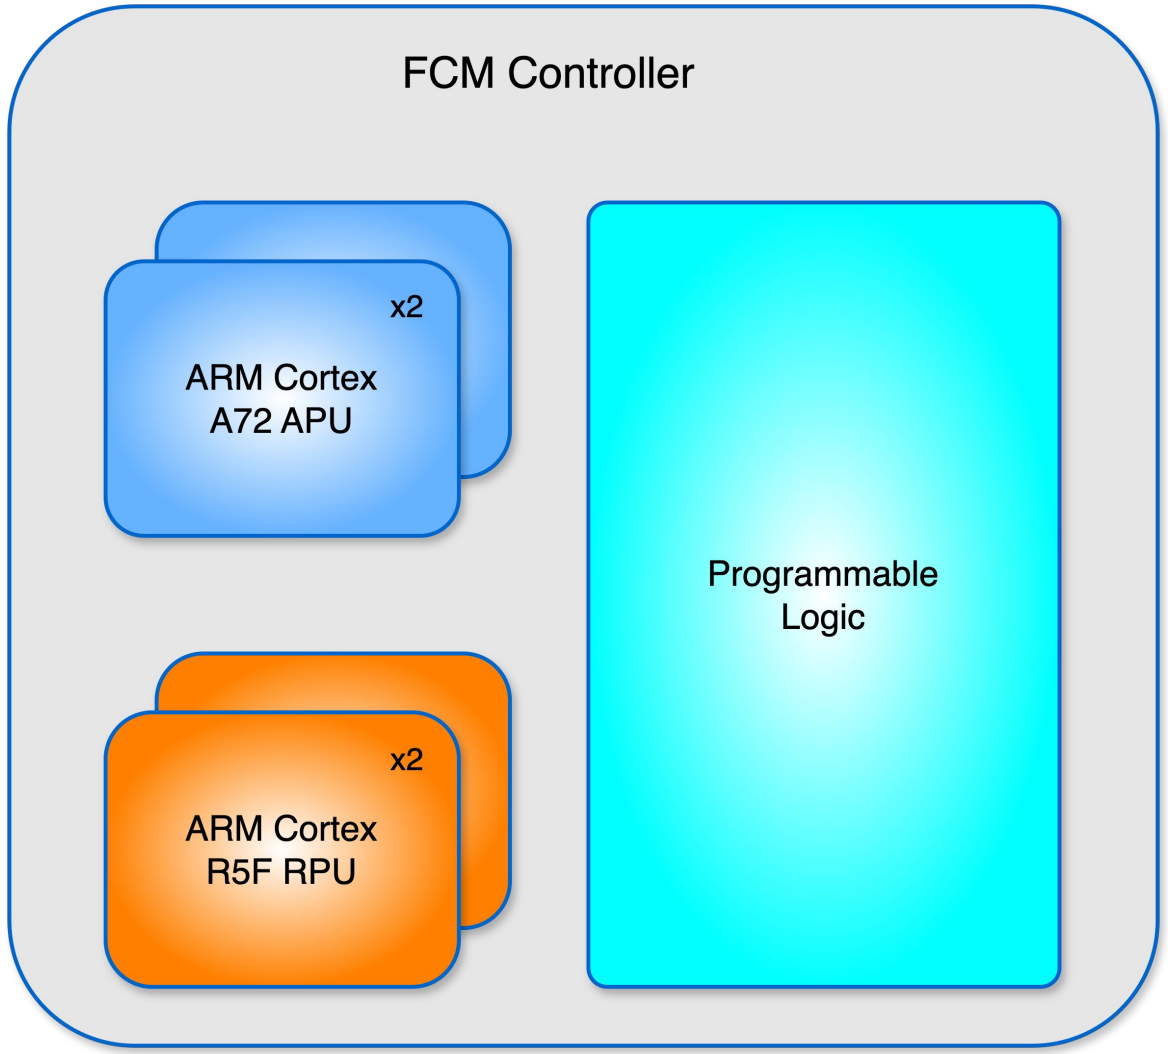

FCM CS Infrastructure

FCM has implemented a standardized process for computational storage, with an infrastructure built on NVMe® technology specifications:

NVM Express® Computational Programs
NVM Express® Subsystem Local Memory

While our use case is still primarily within our own appliances, where we control the software stack, this will allow the FCM to efficiently adapt and modify our computational programs for future use cases and platforms.

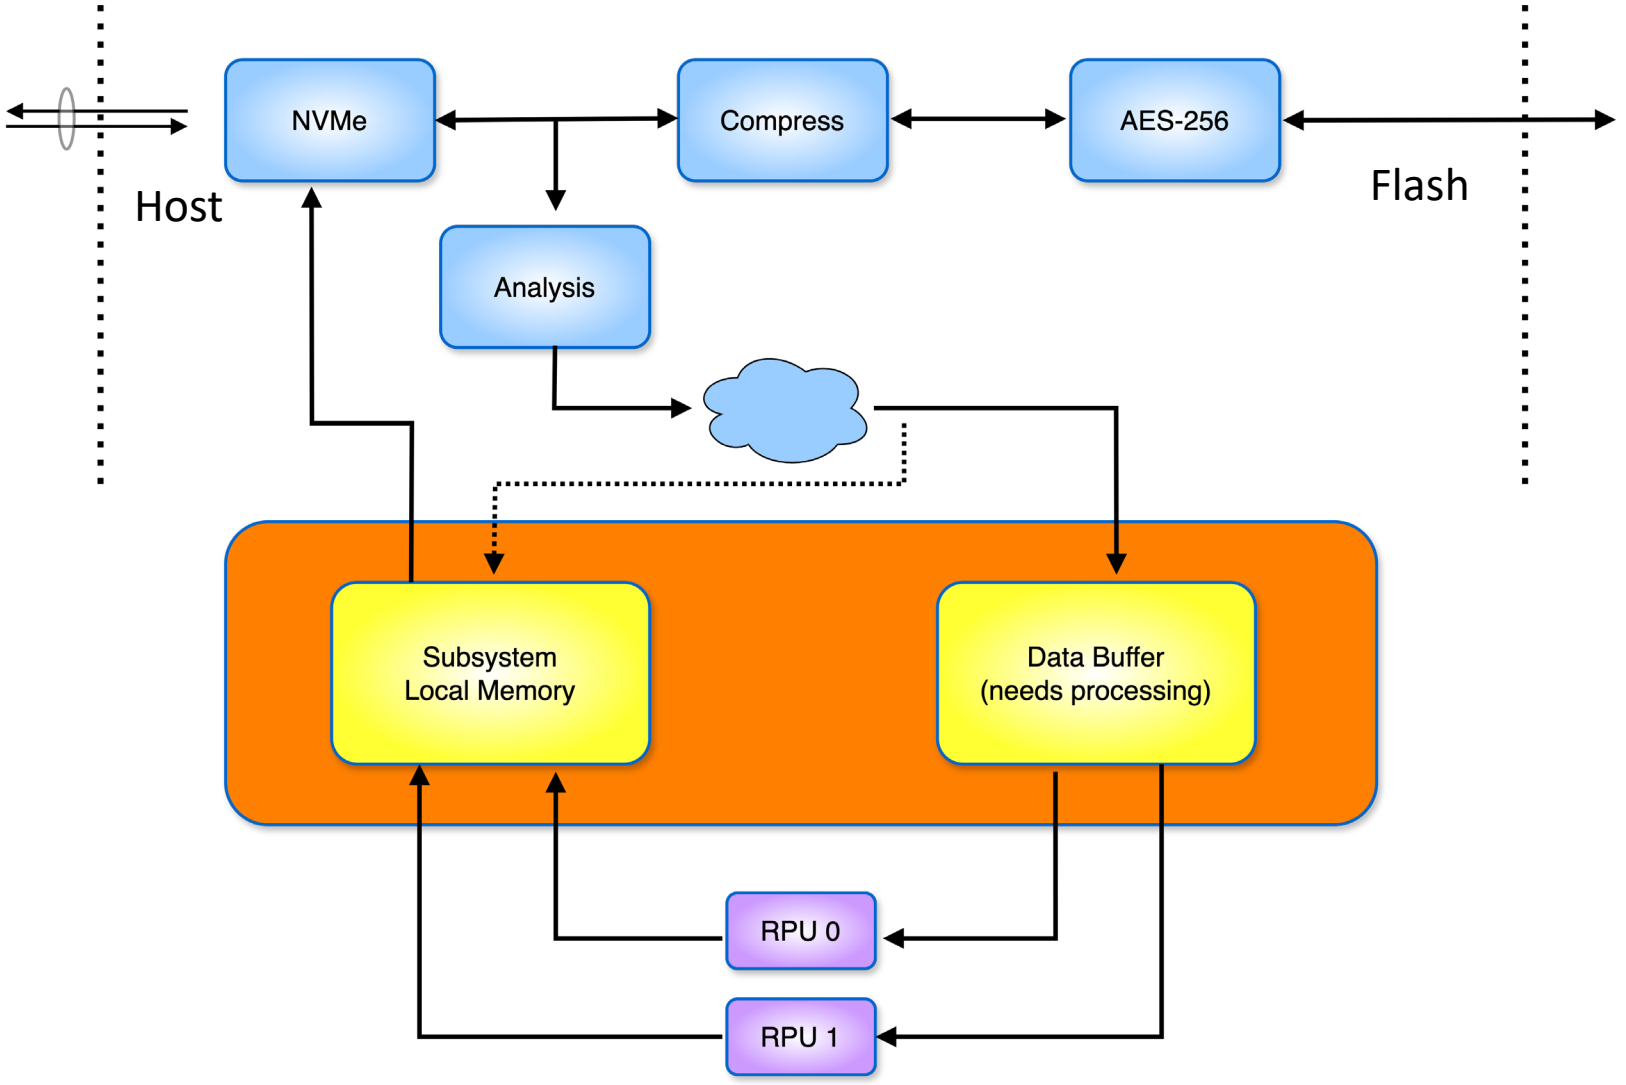

Use Case 1: FCM CS Deduplication Assist

Background Deduplication

- Inline deduplication is extremely hard, and in many cases, as IO rates increase, deduplication may be turned off
- Providing inline, hardware compression at the drive level provides a lot of benefit to performance for the system and host, as it requires no system performance, compute or memory overhead
- The immediate data reduction offered by drive compression serves allows the system to efficiently perform deduplication while still overing great value with
- Analysis done in the FCM, over time, can also allow the deduplication engine to prioritize logical regions for deduplication
- This makes a great case for performing Background Deduplication

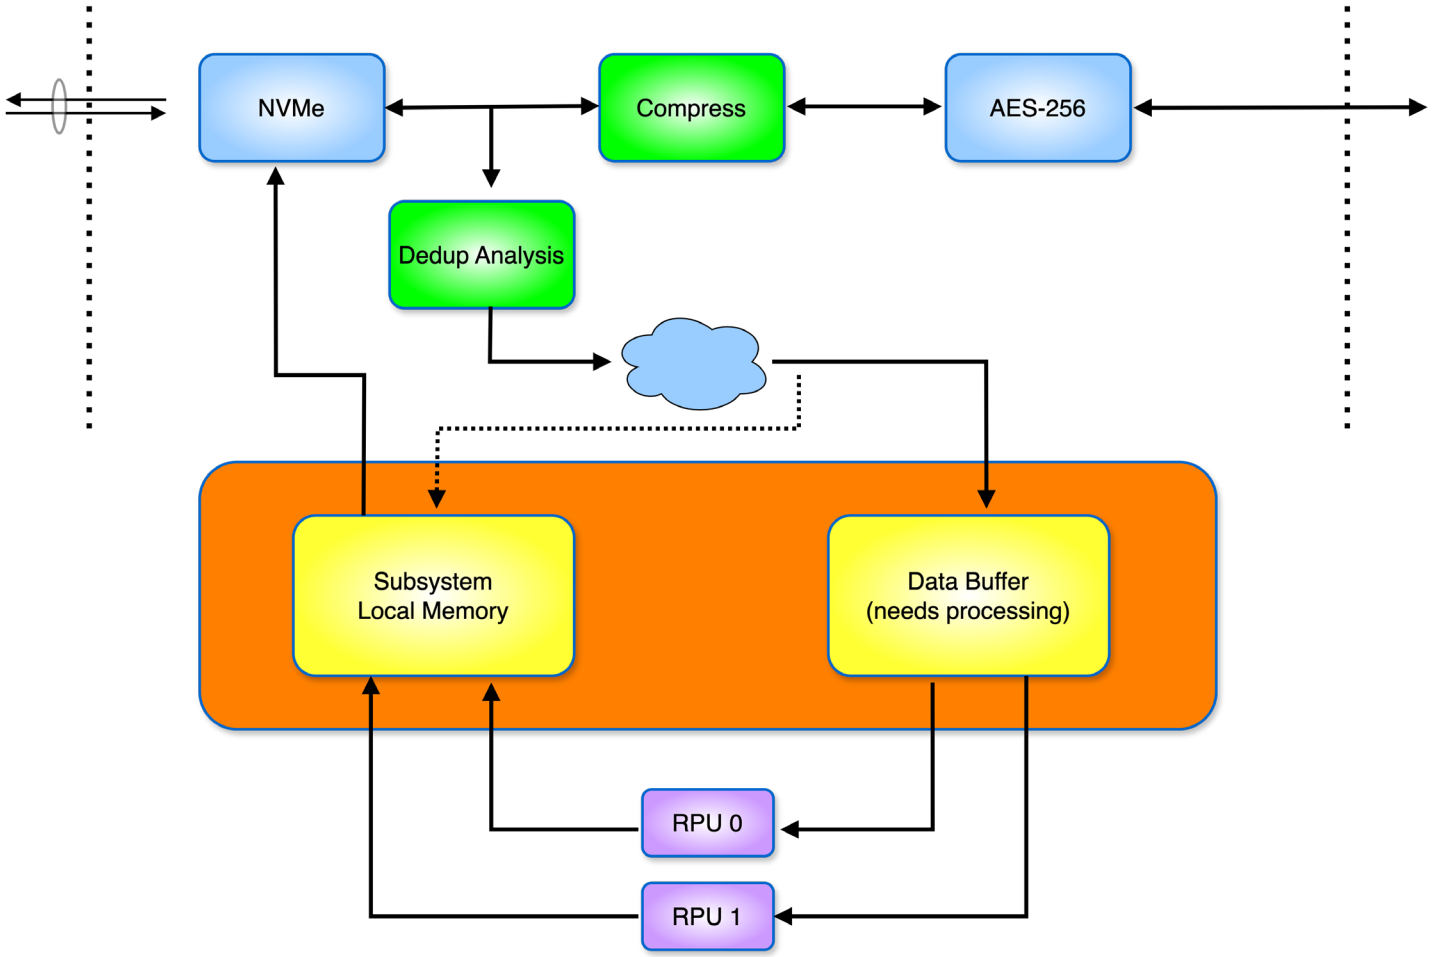

Use Case 2: FCM CS Deduplication Assist

Hash Generation

- Creating strong hashes can be a latency intensive operation as many rounds/iterations may be required to generate the hash
- Inline hashing at sizes similar to the block size can require many parallel hash engines and a lot of computation
- It makes a lot of sense to do hash generation on the FCM, in the background
- The FCM infrastructure can support many parallel hash engines, and there can be many FCMs in a storage array
- The FCM CS infrastructure can efficiently return hashes with little to no performance degradation
- Furthermore, the FCM can help guide deduplication…

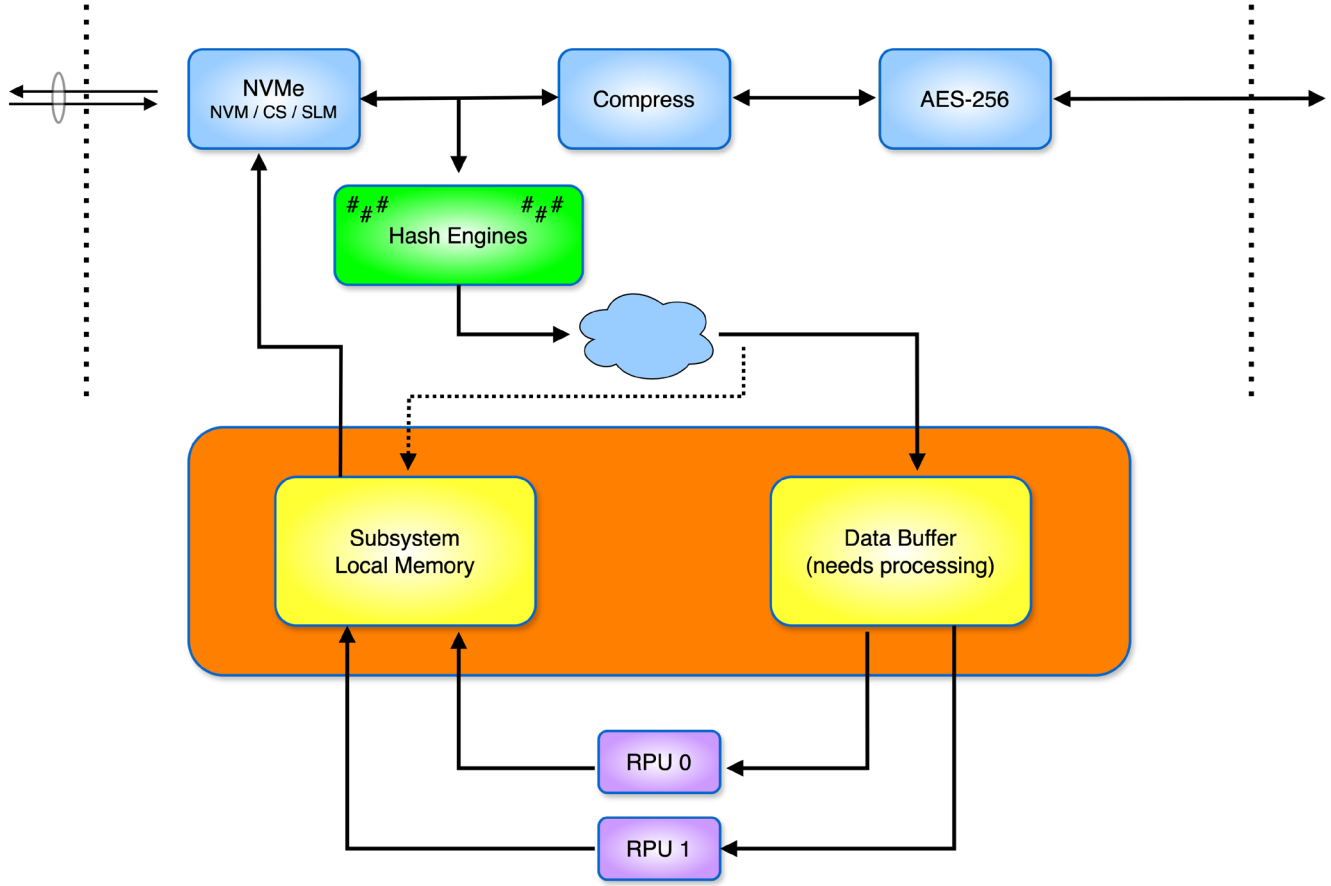

Use case 3: Deduplication aided by FCM compression
Volume Prioritization

- Typical Enterprise Storage appliances can have 1000s of volumes or partitions
- FCM metadata can be used to calculate compression ratio on each volume
- Deduplication layer in the storage appliance can choose to focus on volumes with a certain range of compression ratios to reduce system memory and CPU cycles
  - If compression for a given volume is good, the system can bypass the volume entirely or reduce priority
  - Alternatively, if there is no compression for a given volume, the appliance can try to determine if the data is encrypted and choose to bypass or lower priority

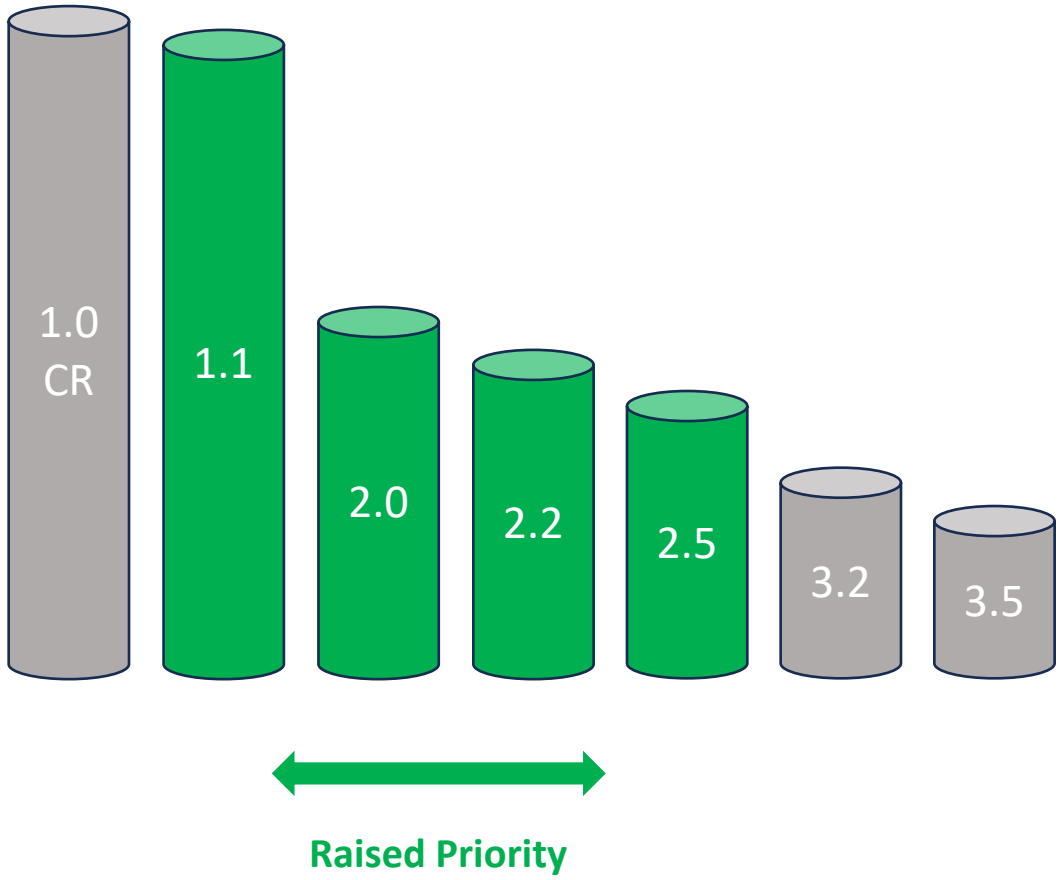

Use case 4: Deduplication aided by FCM compression

Volume grouping

- In an Enterprise Storage Appliance, it may be typical to create many different volumes, that due to the application or use case, may have similar data
- Furthermore, many backups and copies may commonly be stored on other volumes
- This leads to volumes that have very similar compression ratios
- By tracking the compression by volume and other statistics about the volume, deduplication can prioritize analyzing similar volumes

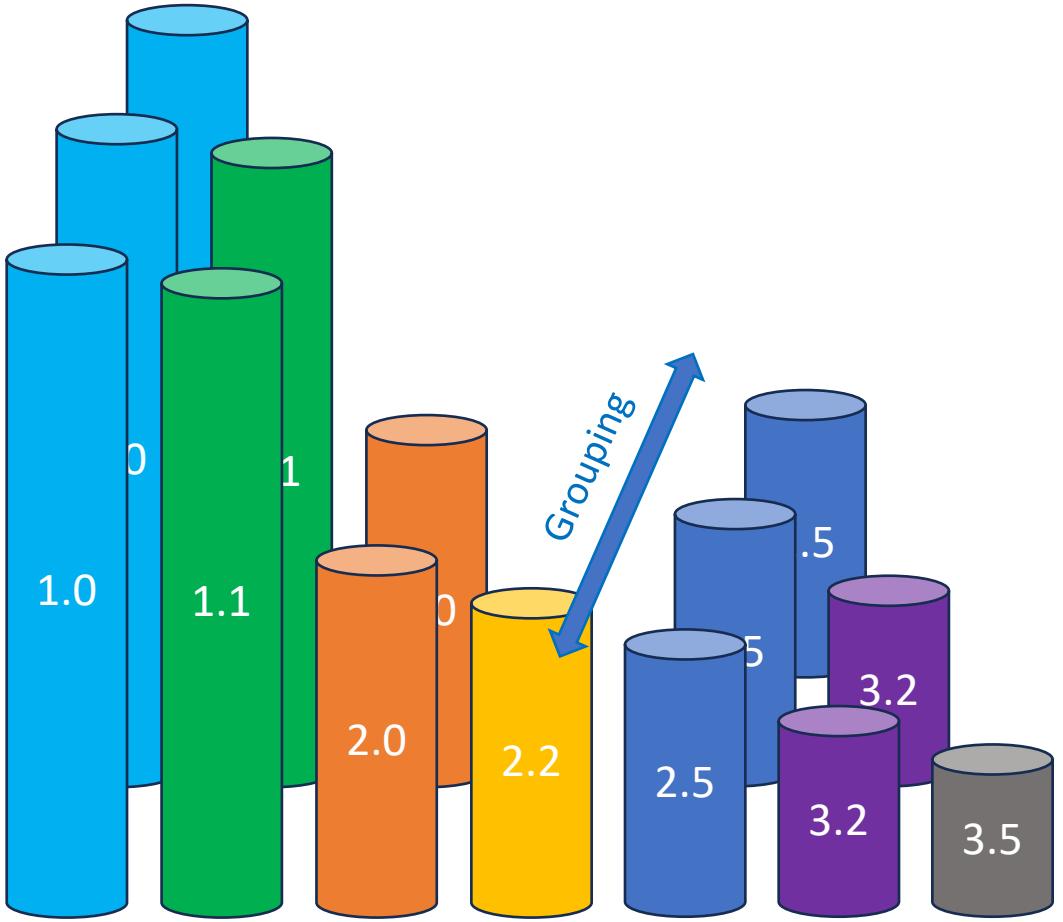

Deduplication aided by FCM compression
Summary

- Data Reduction is a necessary feature for an Enterprise Storage Appliance
- Compression done in an SSD/CSD can provide enormous benefits for both system performance and cost
- Deduplication is very “expensive” in large enterprise storage arrays
- Computational Storage and associated NVMe specifications can efficiently compute and deliver hash information
- Using compression information to prioritize or guide deduplication can provide a lower cost and higher performing Enterprise Storage appliance!