Flash Memory Summit

Effective Resource Utilization with the Hardware Accelerators in SSD Controllers

Presenter: Ramyakanth Edupuganti, Applications Engineering

Microchip Technology Inc.

MICROCHIP

Agenda

- NVMe® SSD Resources
- Challenges of NVMe® SSD
- Data Processing Offload Engines

NVMe® SSD Resources

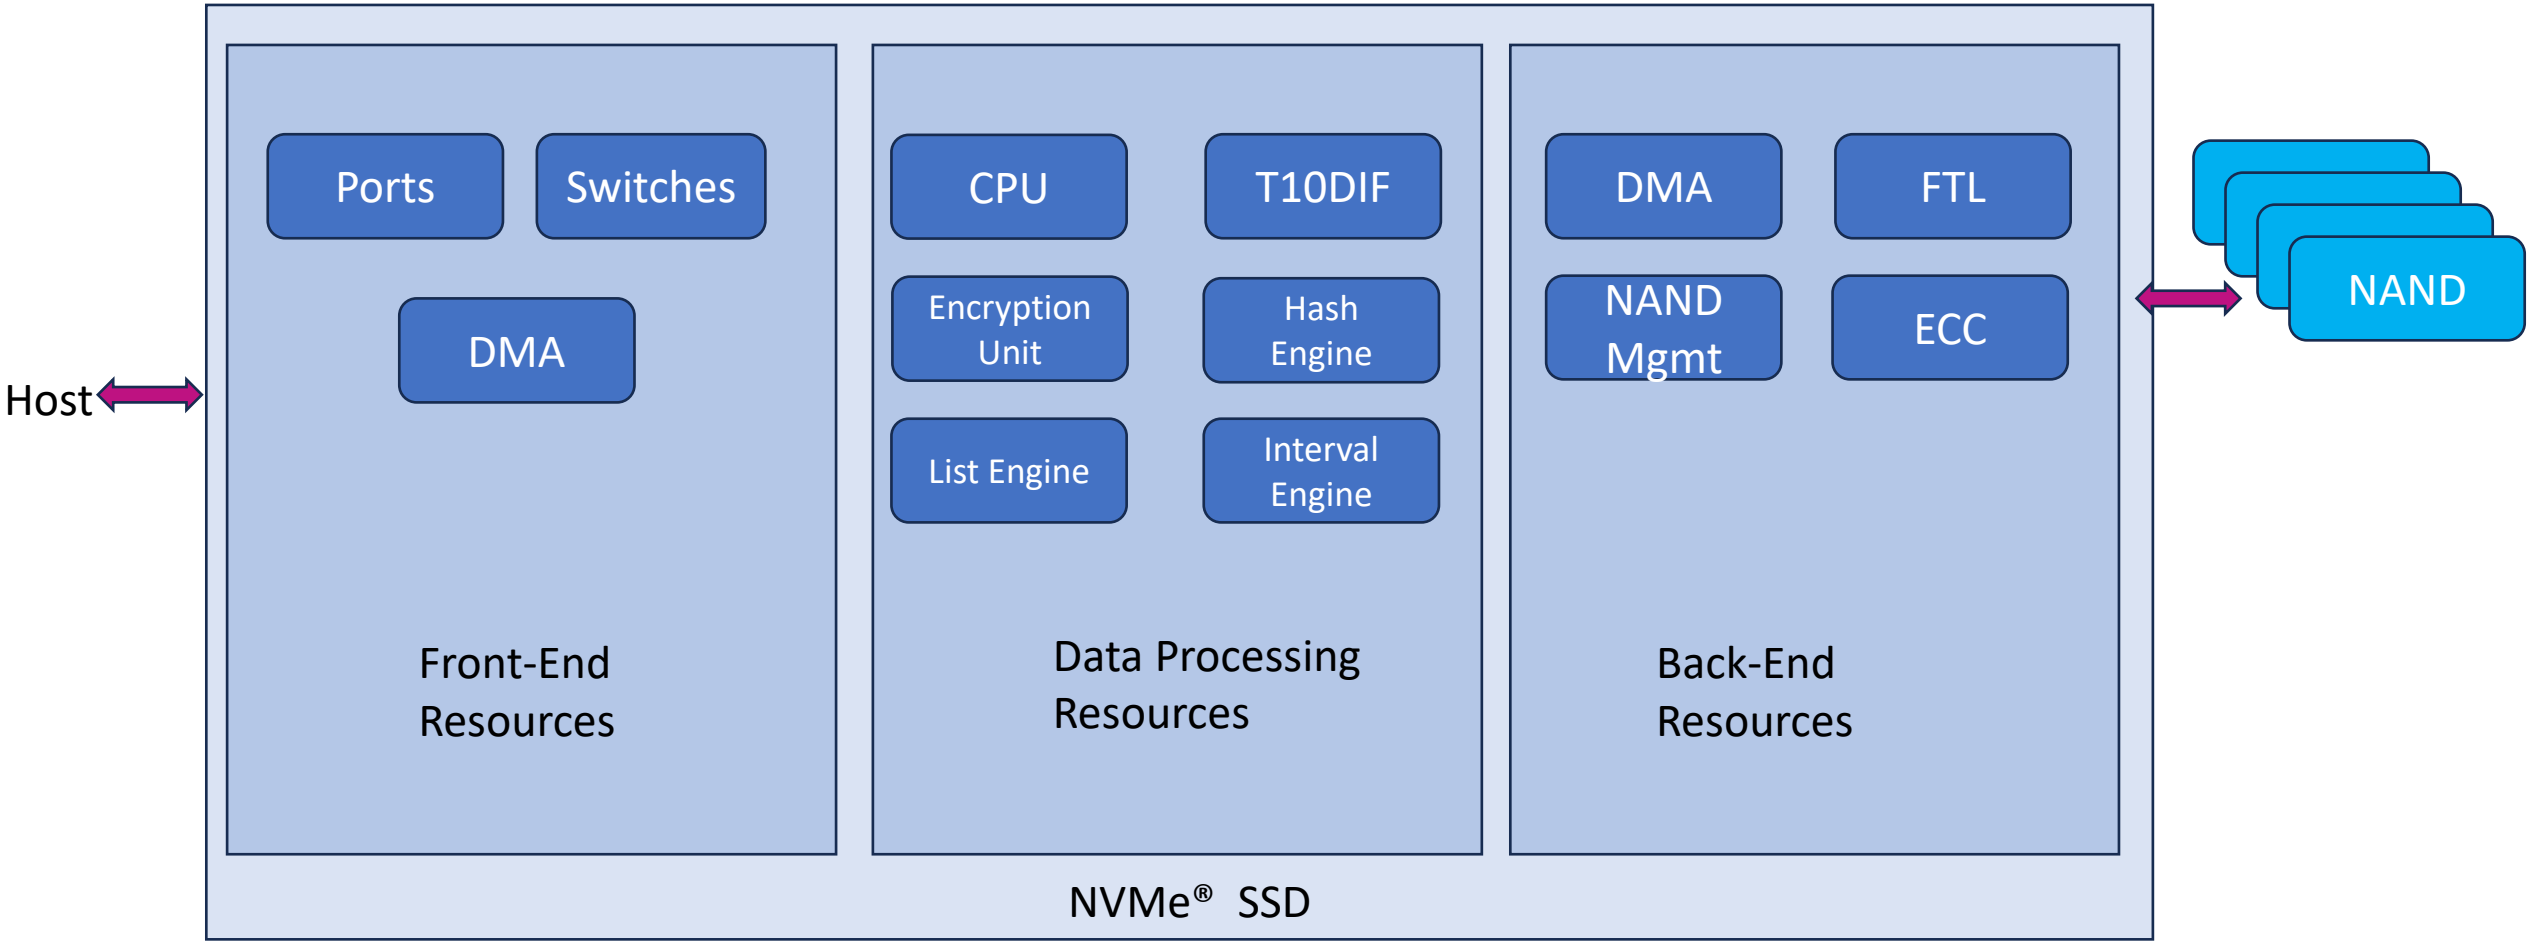

Challenges of NVMe® SSD

- Challenges of NVMe® SSD

  - Increase in density and performance
  - NAND devices becoming more complex, need for smart NAND management
    - Increasing number of layers, bits-per-cell
    - Change in technologies
  - Compute capability of NVMe® SSDs
  - Fault detection and prevention
  - Debug capabilities on failure

Computing resources in SSD controllers

- Effectively utilizing HW engines listed below can resolve some of the issues
  - HW accelerated machine learning engines
  - Offload engines such as XOR engines for data protection
  - CPU cores with additional capabilities like floating point
  - PCIe® expansion ports for connecting additional resources

MICROCHIP

Applications of ML in NVMe® SSD

Flash Memory Summit

- ML in Computational Storage
  - Firmware (FW) in SSD can repurpose ML engine in the NVMe® controller as a Computational Storage Engine (CSE), entail the SSD as a computational Storage Device (CSD)
  - Host CPU or peer devices can take advantage of this capability in the NVMe® SSD to free up their computing resources

MICROCHIP

Applications of ML in NVMe® SSD

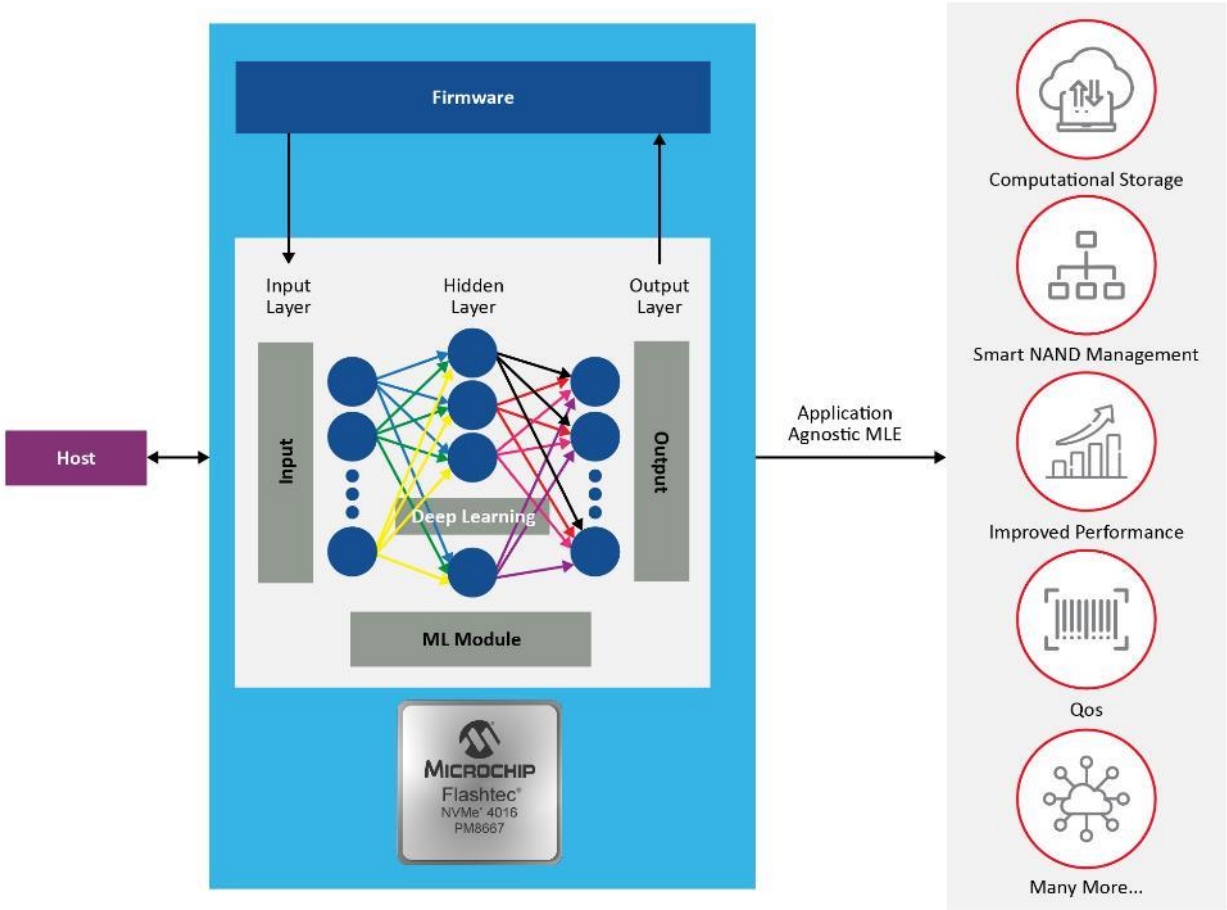

Applications of ML in NVMe® SSD

- ML for Smart NAND Management

As NAND technology advances, effective NAND management is becoming challenging:

- Layer count increasing
- New NAND technologies
- Number of bits/cell increasing

Need to improve error correction and develop techniques to reduce errors to take advantage of newer, denser and cost-efficient NAND

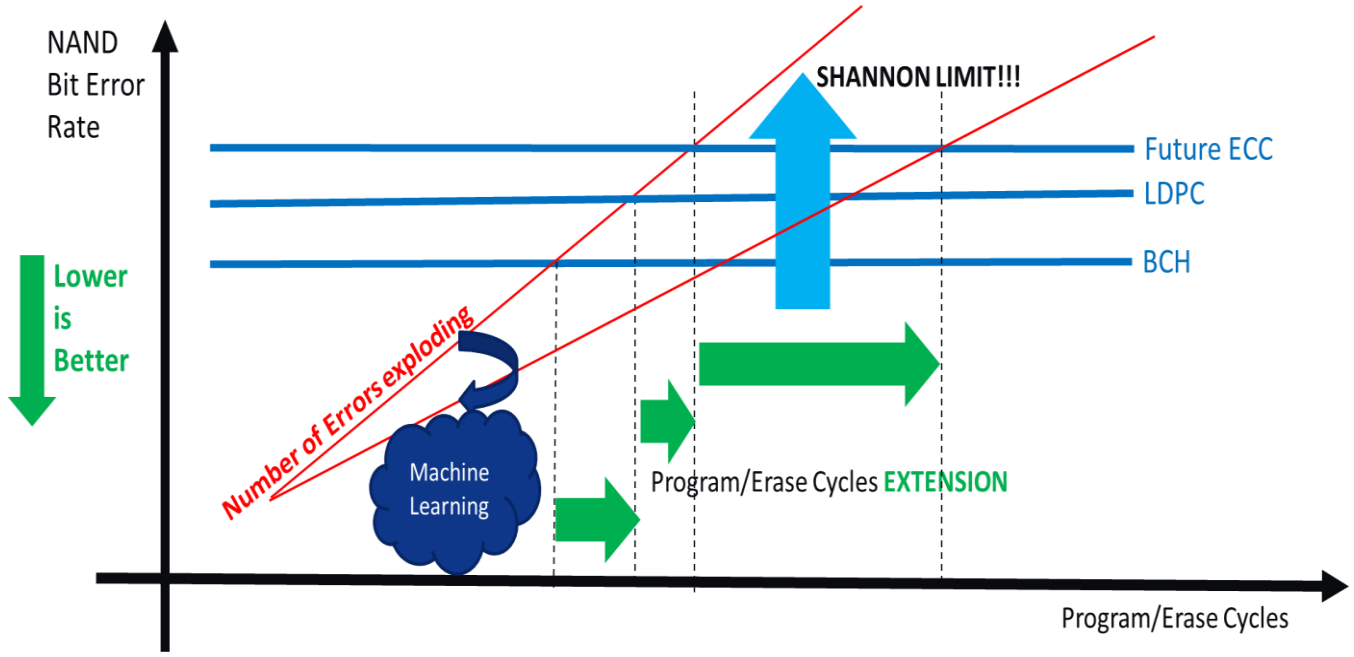

NAND BER After Applying Machine Learning

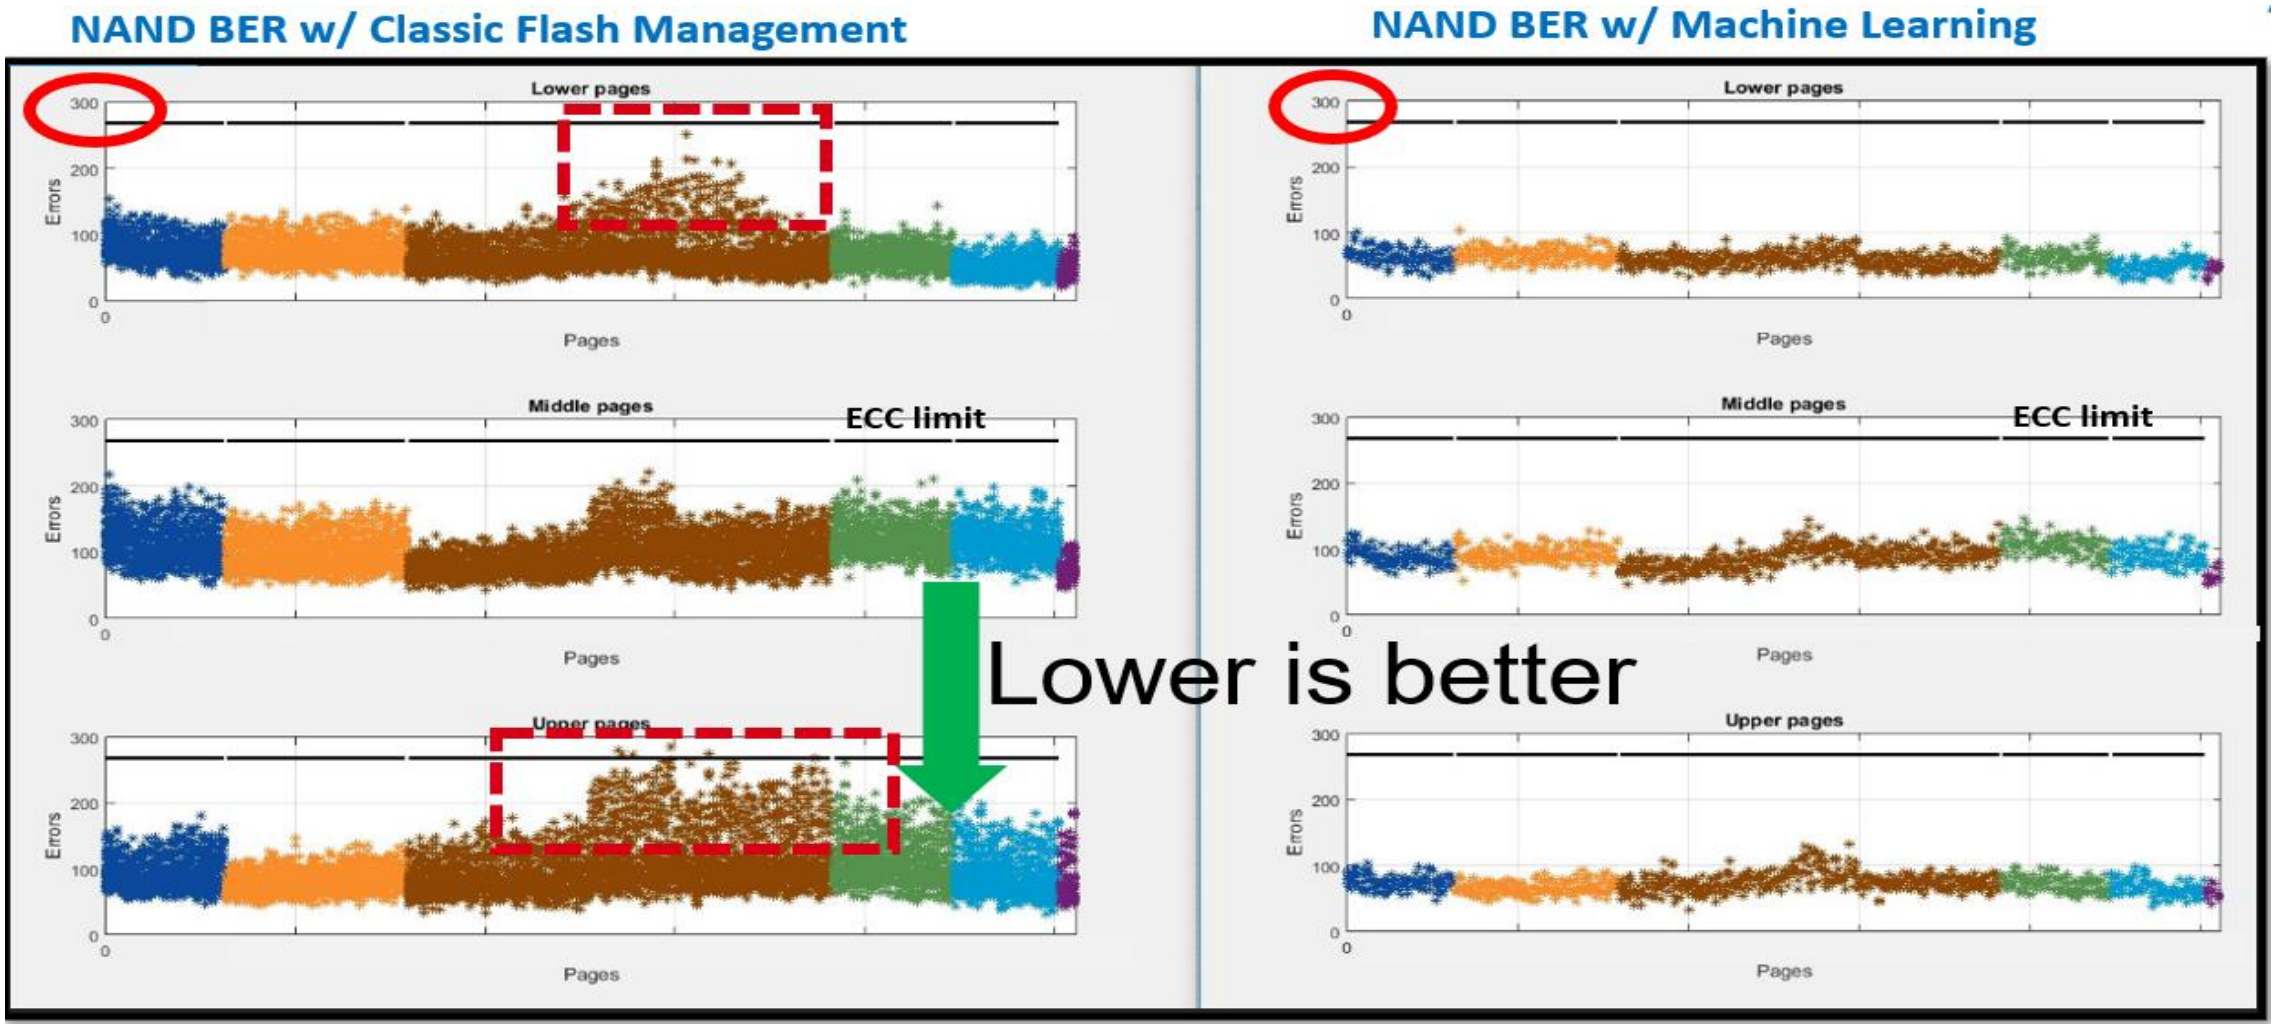

Applications of ML in NVMe® SSD

Flash Memory Summit

- ML for improved performance
  - Performance of an SSD in the context of I/O, endurance and power can be improved by ML engine
  - Some examples of AI/ML in the NVMe® SSD
    - Recognizing the I/O patterns
    - Allocating and adjusting the data path resources
    - Pre-fetching and keeping the data ready in the local memory
    - Using advanced cache management strategies
    - Applying dynamic power management strategies
    - Developing advanced Flash translation layer
    - Garbage collection
    - Wear-levelling methods

MICROCHIP

Applications of ML in NVMe® SSD

- ML for QOS
  - NAND errors cause fallback to error correction and impact QOS
  - By using effective NAND management, the number of NAND errors can be reduced thus causing less impact to QOS
- Other uses of ML
  - AI/ML engine in NVMe® controller provides endless capabilities to the NVMe® device
  - Some additional scenarios are
    - Minimizing disruptions by early detecting and recovering or failover from faults, FW upgrades, security attacks
    - Optimizing power and performance by throttling the resources
    - Dynamically adjusting the information to be collected and stored for smart telemetry
    - Performing self-tests and corrections
    - Scaling with varying needs during the operation of the SSD without human intervention

XOR Engine for Data Protection

- Adding an XOR engine to the NVMe® controller can be used for data protection
  - Host can use this offload engine to free up host CPU utilization

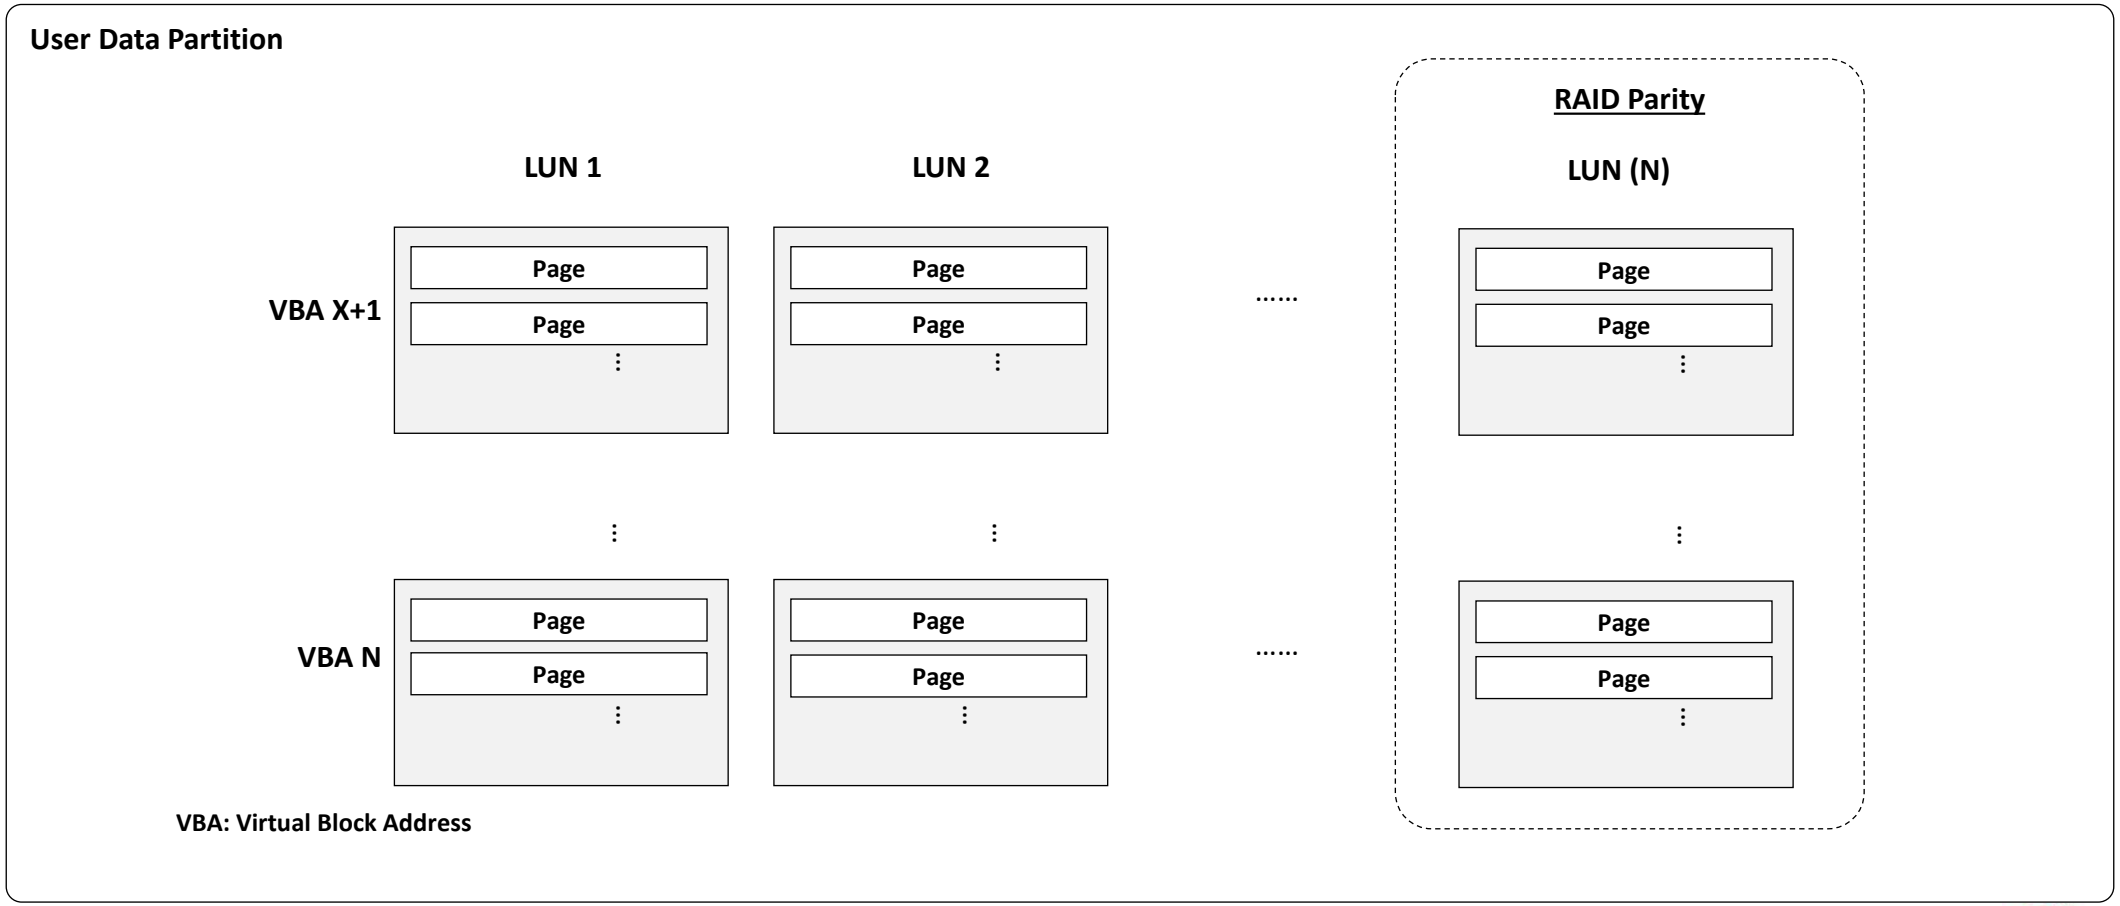

Effective utilization of CPU cores

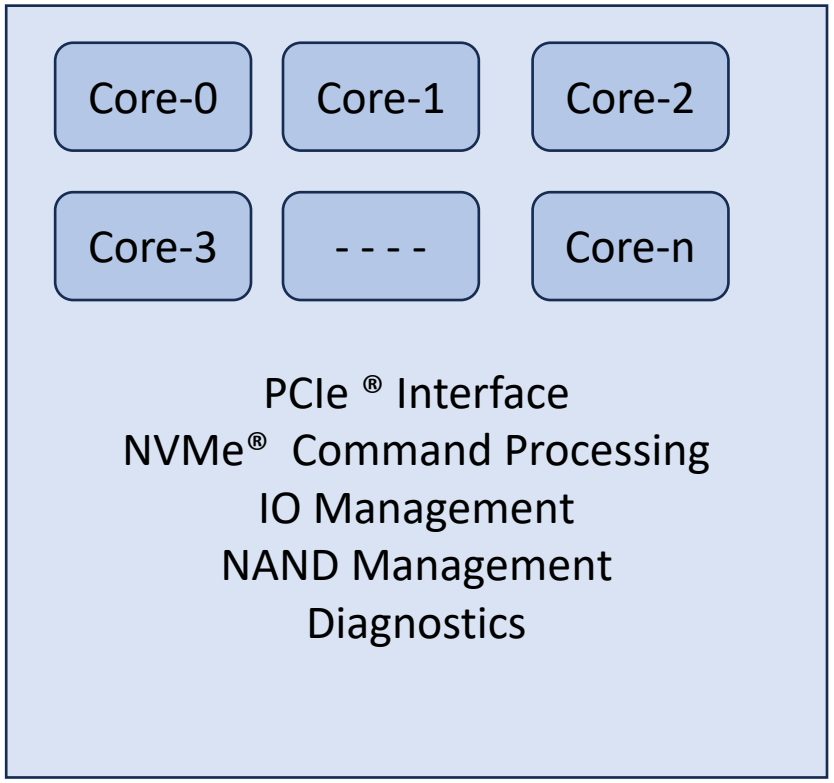

Typical SSD with n-cores

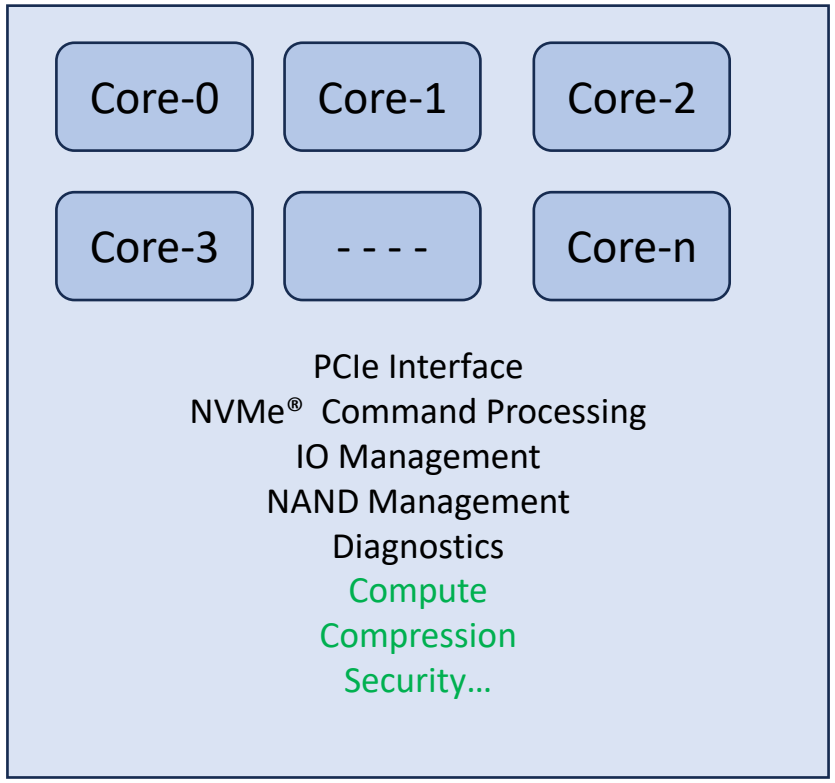

Effective utilization of SSD CPU cores

PCIe ® Expansion Ports

- NVMe® SSDs may have multiple ports including PCIe expansion ports to connect additional devices
- These additional devices could be another SSD, FPGA, etc., and can be used to add additional capabilities to the devices such as
  - Compute
  - Additional feature enhancement
  - Offload engines

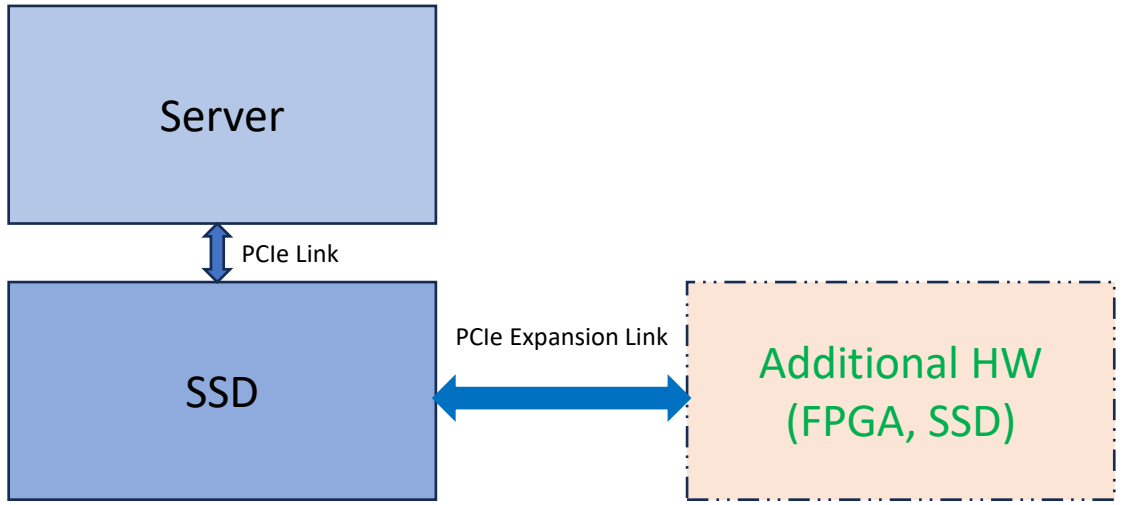

Thank you!

Visit Microchip Booth# 419

References

- https://www.microchip.com/en-us/about/media-center/blog/2022/how-ai-is-transforming-NVMe® -ssds
- https://www.snia.org/sites/default/files/computational/20190808_COMP-301A-1_Ram.pdf