



Flash Memory Summit

# Enabling Flexibility in NVMe<sup>®</sup> Solid State Drive (SSD) controllers for the Cloud Scale Architecture


Sanketh Srinivas

Manager, HW Applications Engineering

Microchip Technology Inc

# Converging Enterprise and Cloud SSD Needs

## PCIe® SSD Controller Units Doubling by 2026



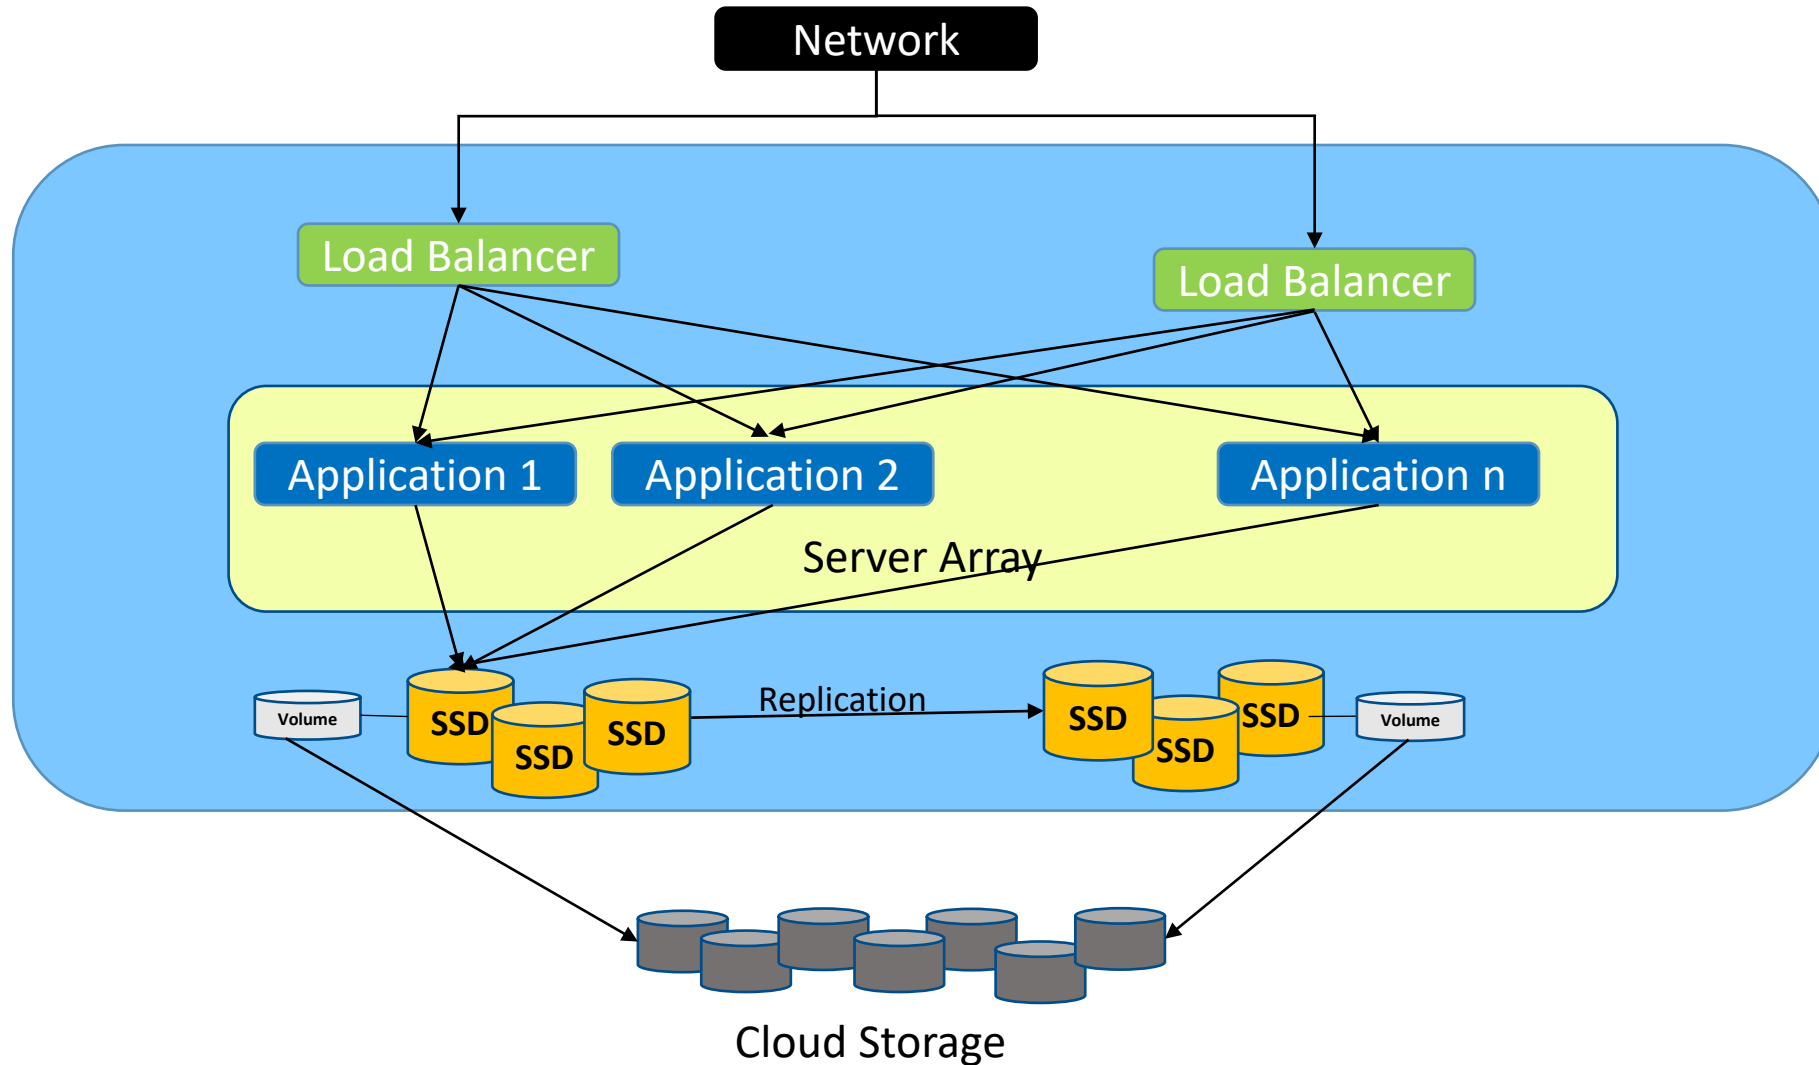
Source: Microchip Internal Forecasting

PCIe® Performance → Enterprise Market  
Performance Compute → Cloud/Hyper Scalars

- Accelerated transition into cloud and the nature of workloads and use cases that extensively use **Hot Data** are necessitating Cloud infrastructure to increasingly adopt SSDs

- Analysis extraction
  - Medical outcomes
  - Self-driving vehicles
- Filtered and sorting
  - Data analytics
  - Cat photos
- Time critical data
  - Stock market decisions
  - Task Prioritization
- Augmented reality
- Advertising
- Catastrophe avoidance

# Typical Cloud Infrastructure



# Requirements for SSDs in Cloud architectures

## Flexibility & Scalability

- Linear scaling for Performance with Capacity
- Support for SR-IOV and Multiple Physical Functions
- Interoperability

## Bandwidth Management

- High-performance storage for every computing instance

## Performance and Power

- Performance per \$\$\$
- Performance per Watt
- Performance in smaller form factors

## Security

- Secure Boot with Root of Trust (ROT)
- Self Encryption

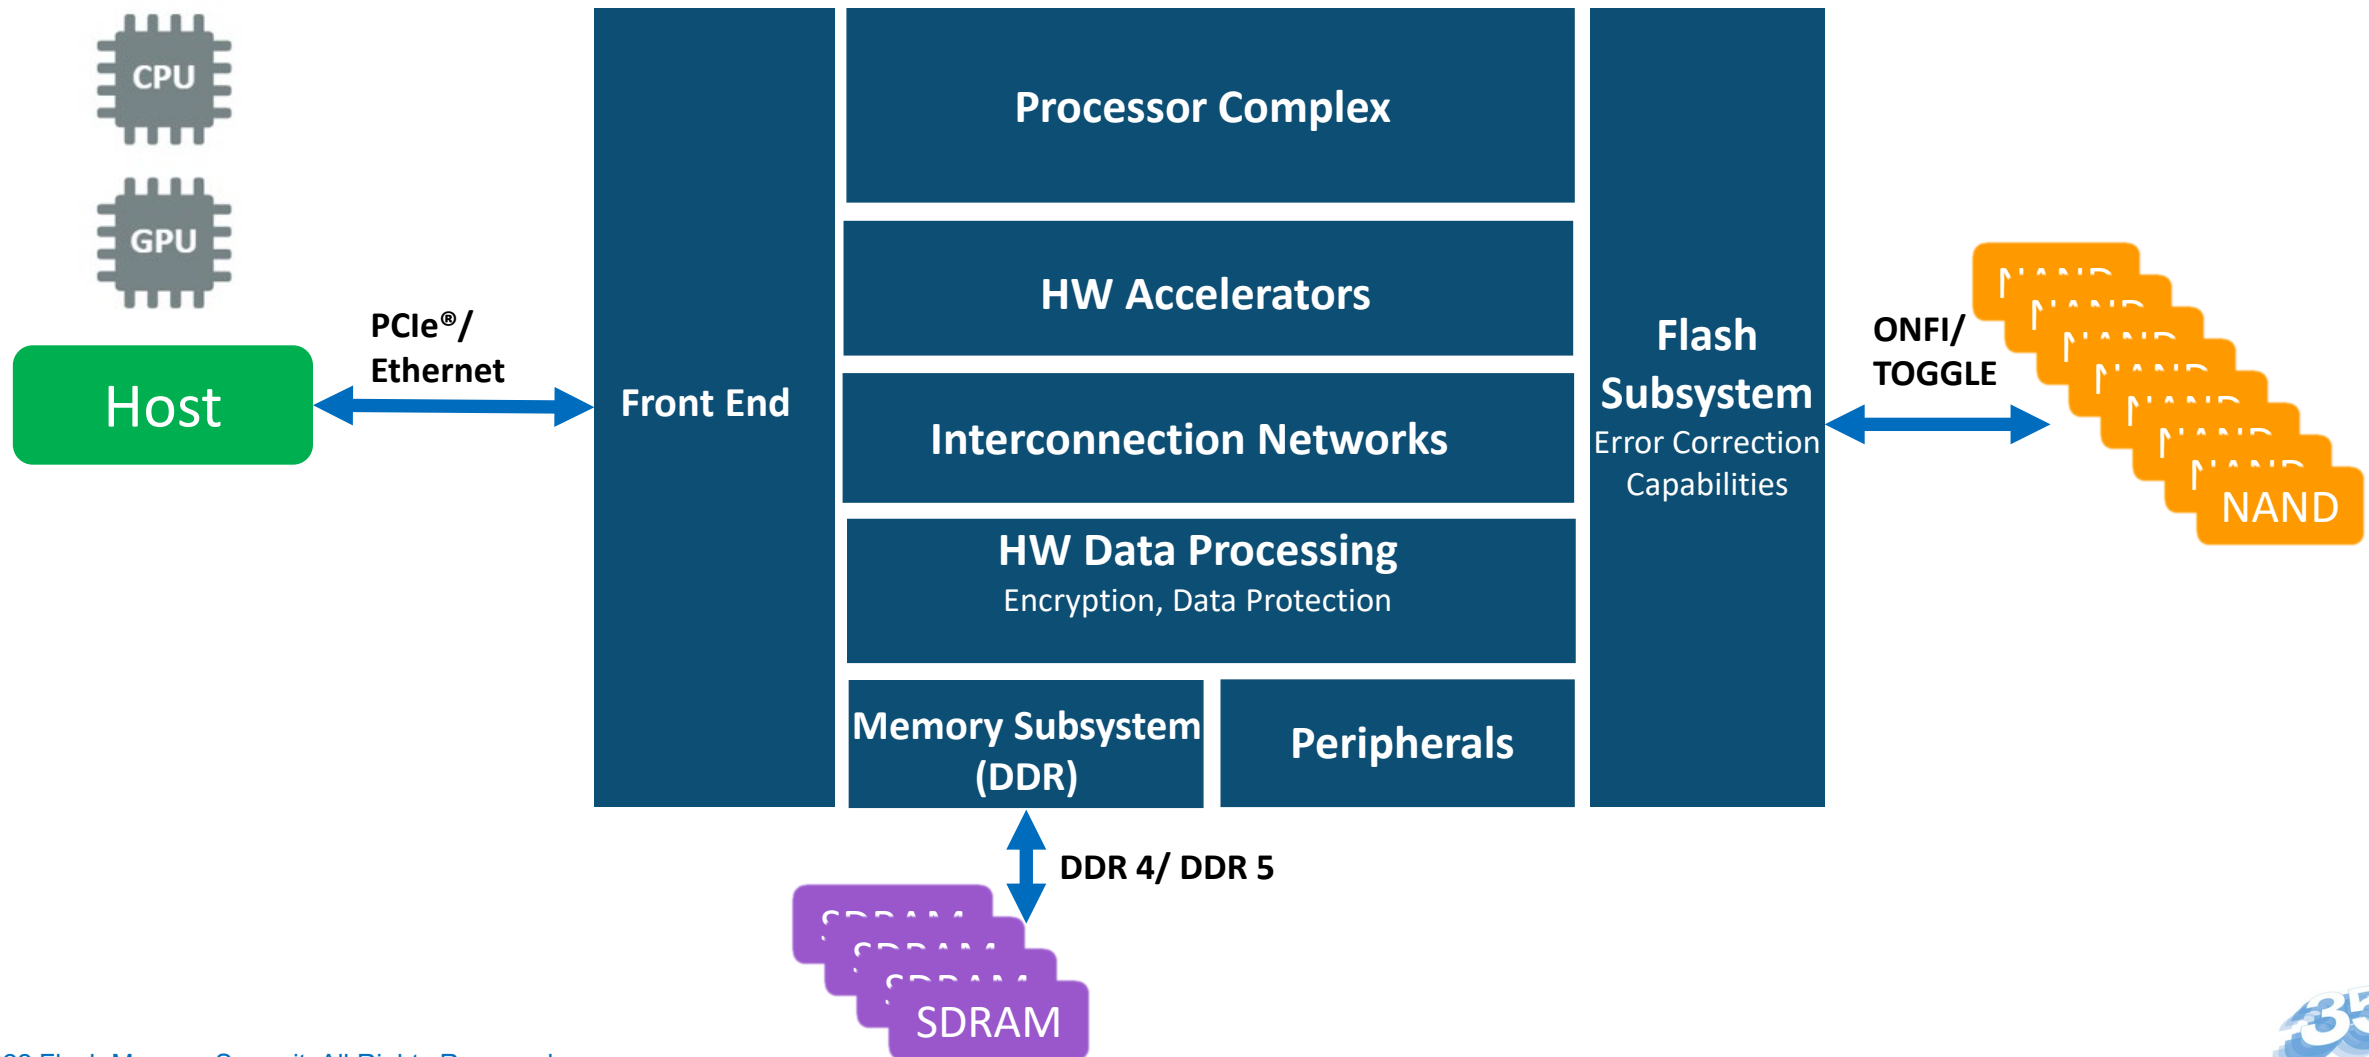
## Determinism

- Guaranteed minimum performance for each tenant
- Predictable QOS

## Efficiency

- Sustainability
- Resource Utilization
- Total Cost of Ownership (TCO)

# Typical SSD Controller Architectures



# SSD Controller Requirements for Cloud



## Flexibility for Cloud

Highly-programmable  
architecture with dynamic  
configurability

Security for Cloud

Virtualization for Cloud

Computational Storage

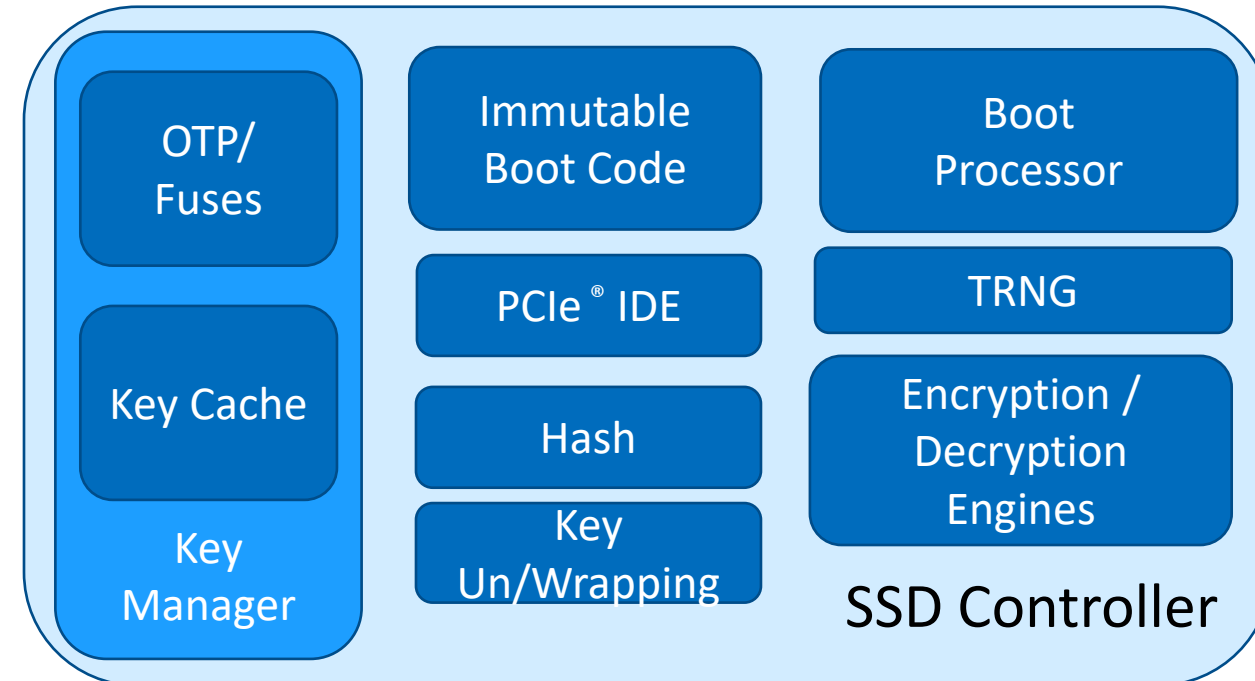
Bandwidth Manager

- SSD controllers need to have a flexible architecture to support the adoption in cloud infrastructure:
  - Efficient resource sharing and reconfigurability using **Virtualization**
  - Ability to **dynamically allocate** internal SSD HW resources for performance and power management as well as tenant management
  - Support for smart technologies like **Machine Learning and AI** for continuous housekeeping, error reduction and adaptation.
  - Ability to operate with **multiple generations of NAND** (ONFI and Toggle)
  - Advanced features and customization in Firmware fool proofing for forwards compatibility

# Security needs for Cloud SSD controllers

- Cloud Infrastructure requires flexible SSD security

- Secure boot with root of trust
- FW authentication and attestation
- PCIe<sup>®</sup> IDE (link encryption)
- Key Per IO
- Data encryption and decryption
- Key storage and wrapping
- Secure debugging
- FIPS 140-3 Level 2 Certifiable

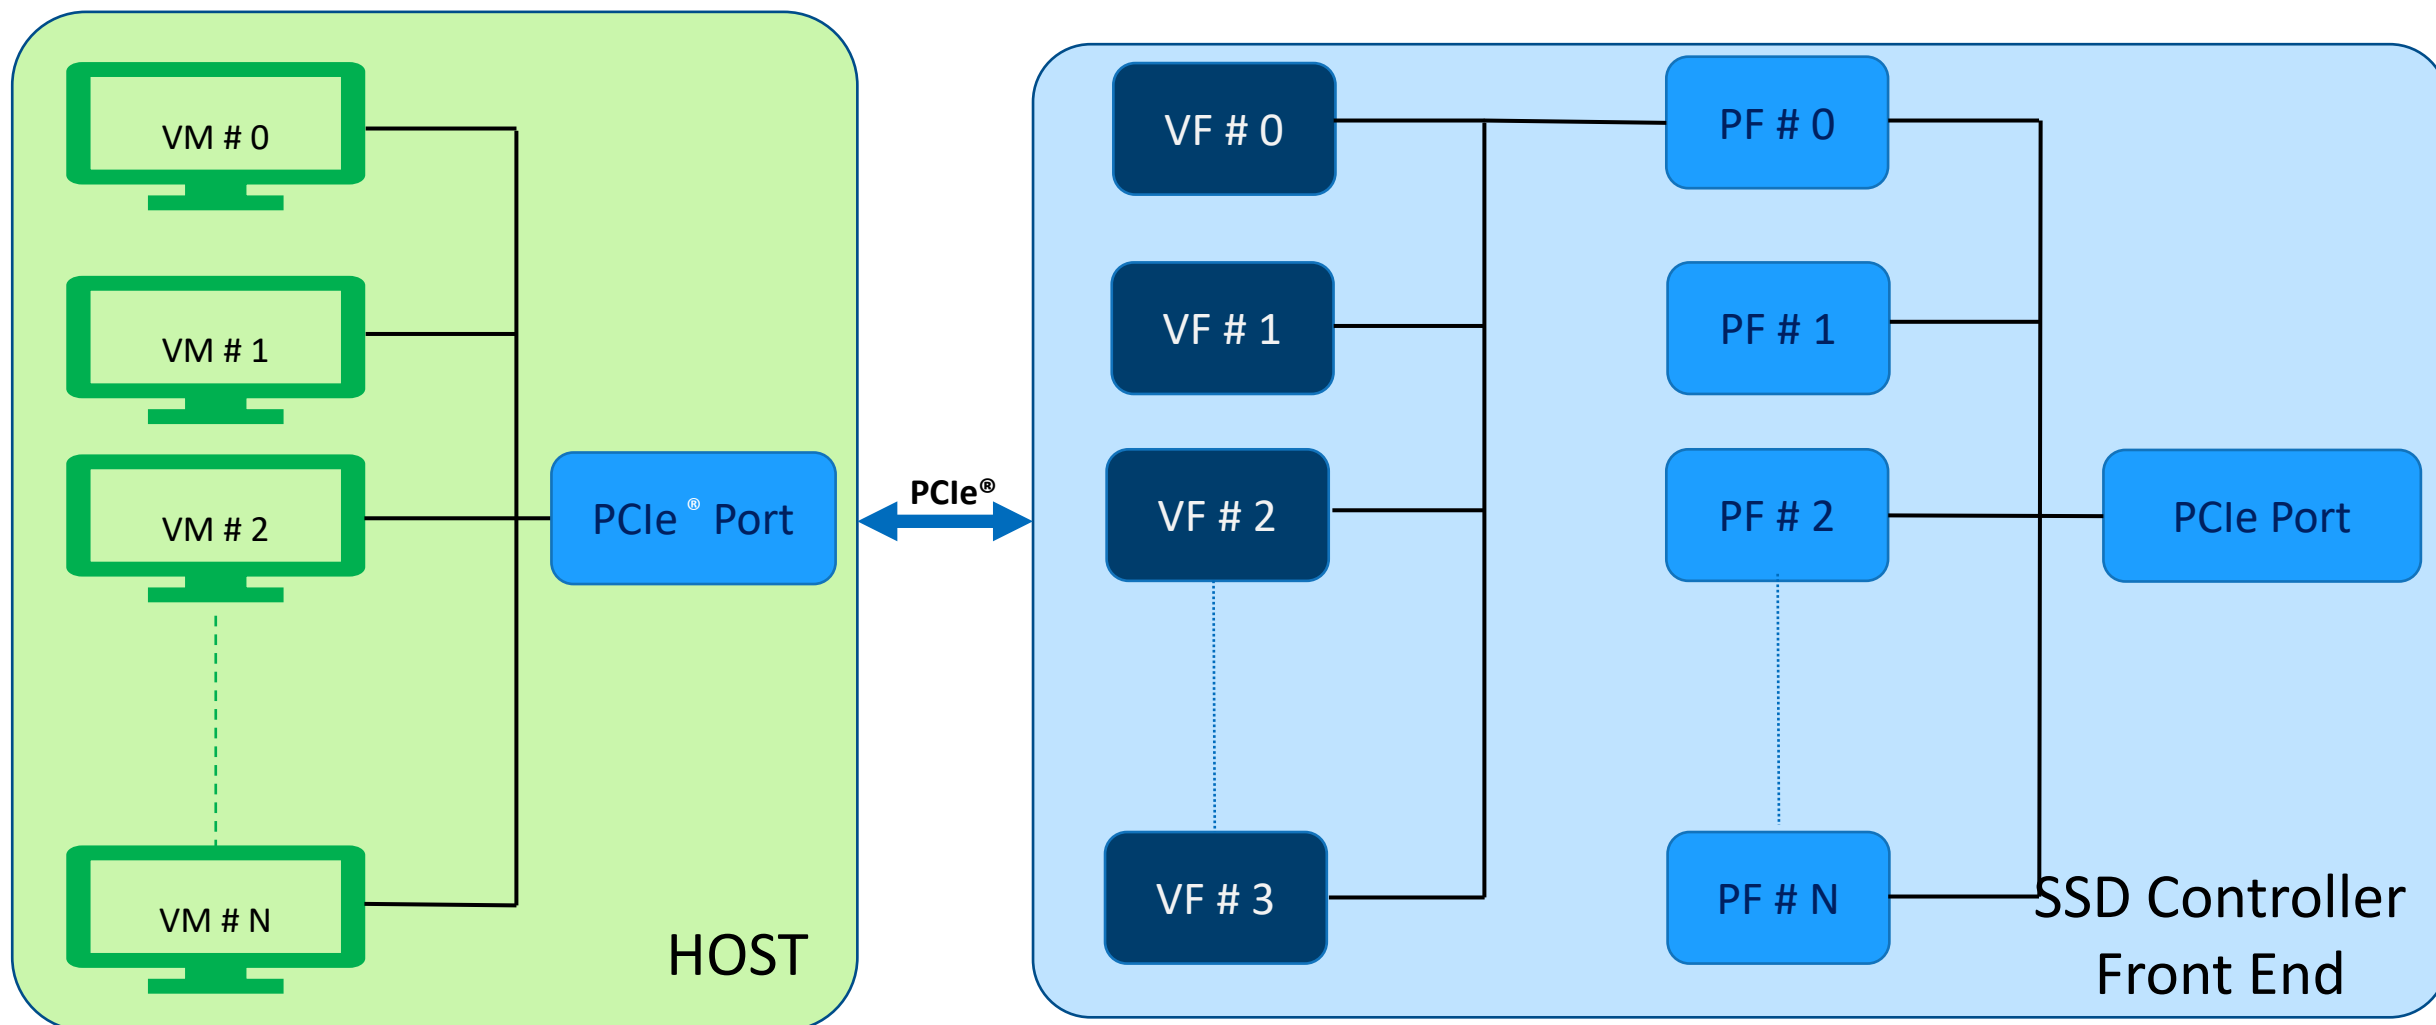




# Scalability using Virtualization for Cloud

- NVMe® SSD controllers should have the ability to scale and adapt to the dynamic workloads from the Cloud users
  - Industry standards such as SR-IOV and Multiple Physical Function (MPF) implementations help support scalability
    - VMs with their own dedicated share of an SSD that does not compete with other users
    - Data Isolation
    - Support for High User Density
    - Minimal Context switching overhead

# Virtualization for Cloud- VFs and MPFs

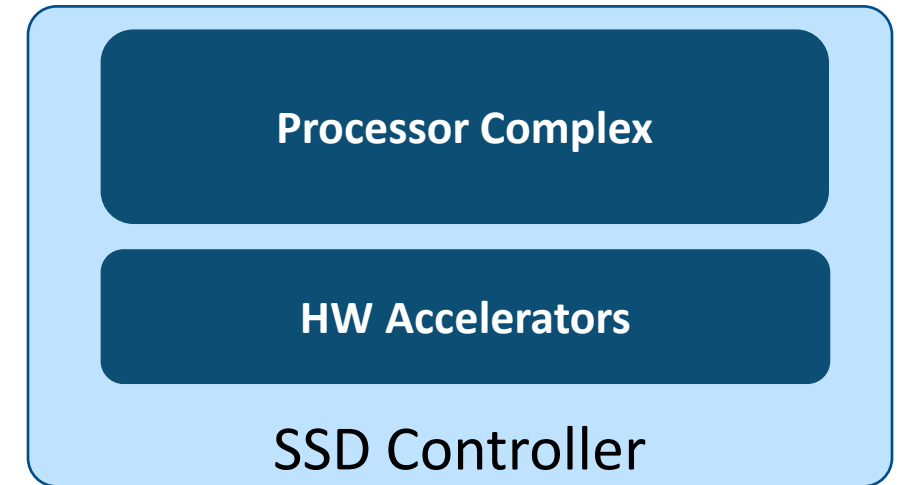


# Enabling Computational Storage for Cloud Infrastructure



Flash Memory Summit

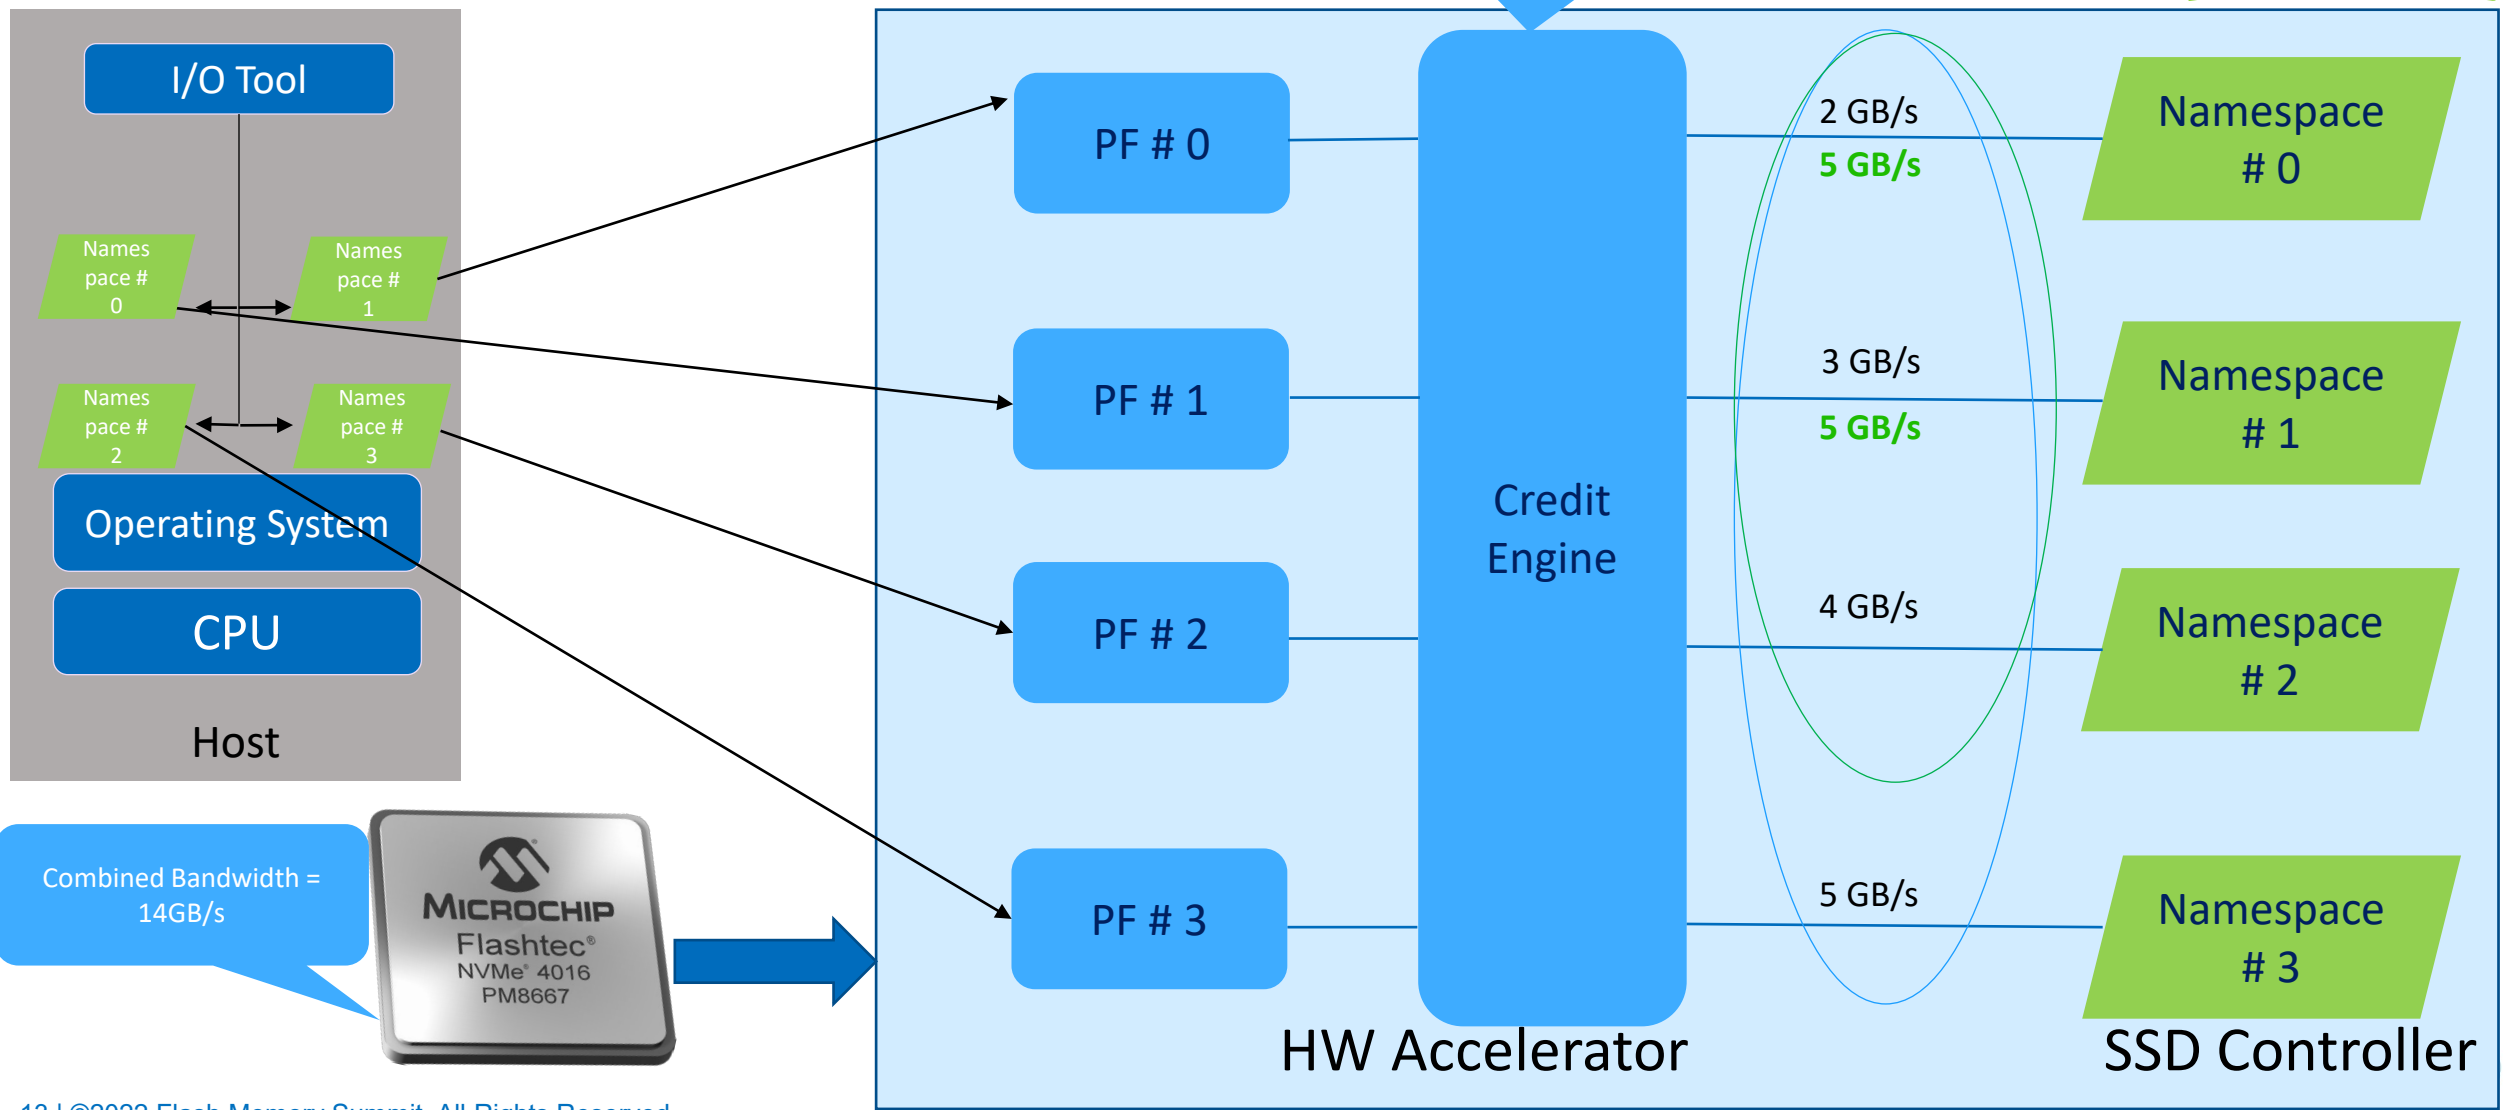
- SSD controllers should have computing HW to minimize the Host computational resource usage and efficiently use the SSD bandwidth
  - Processing Complex
    - General purpose CPU cores
  - HW Accelerators
    - Encryption
    - Data protection
    - Compression
    - De-Duplication
    - Artificial Intelligence/Machine Learning (AI/ML)
    - Field-Programmable Gate Arrays (FPGA)





- Determinism and Bandwidth management are some of the key requirements for a Multi-User environment such as Cloud infrastructure
  - Guaranteed QoS (Priority and Latency)
  - Guaranteed Performance (Bandwidth and IOPs)
  - Consistent Resource availability
  - Ability to dynamically allocate and deallocate resources to meet the demands

# Determinism and Bandwidth Management using Credit Engine



# Requirements for SSD Controllers in Cloud Infrastructure



Flash Memory Summit



## Cloud Enabling Features

### Optimizations for Cloud

Highly-programmable architecture, dynamically configurable power, performance and capacity

### Security for Cloud

Fully secured infrastructure with secure boot, double signing authentication, Integrity and Data Encryption (IDE), etc.

### Virtualization for Cloud

Advanced virtualization with large-scale physical and virtual functions per port

### Computational Storage

High-performance CPU cores, offload engines, Machine Learning and more

### Credit Engine

Dynamic allocation of hardware resources for on-demand cloud services



Thank you!  
Visit Microchip @ Booth #613