



Flash Memory Summit

Computational Storage for the Masses

Mats Öberg

Associate Vice President, Marvell Technology

Compute where the data is



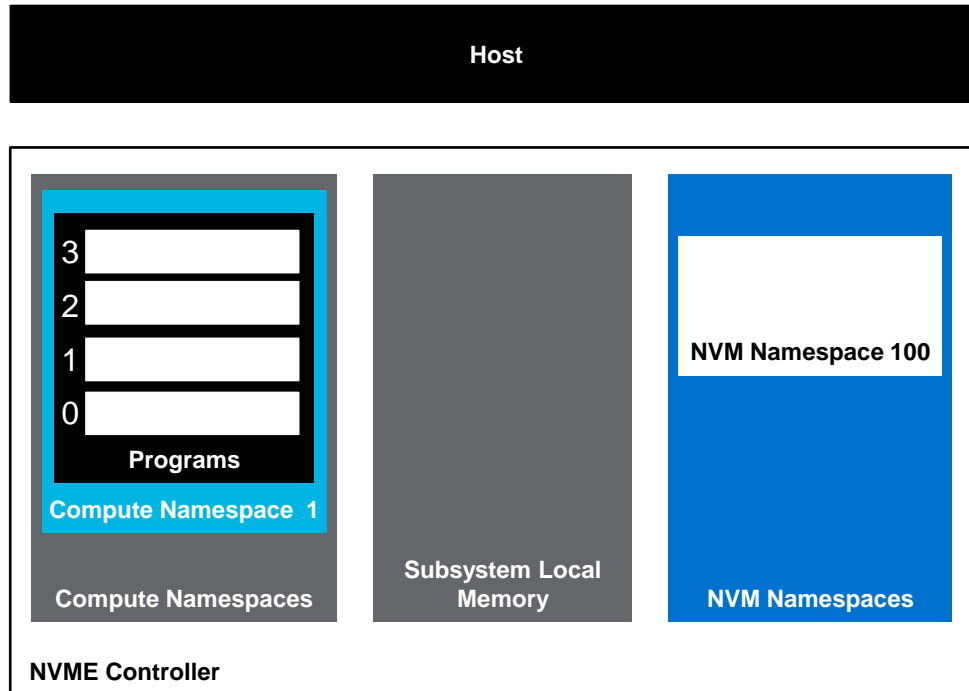
Flash Memory Summit

- Increase value of data
- Avoid major data movements
- Efficiency
 - Power
 - Performance
 - Cost



NVMe SSDs with flexible CPUs for data processing

NVMe TP4091 computational programs overview

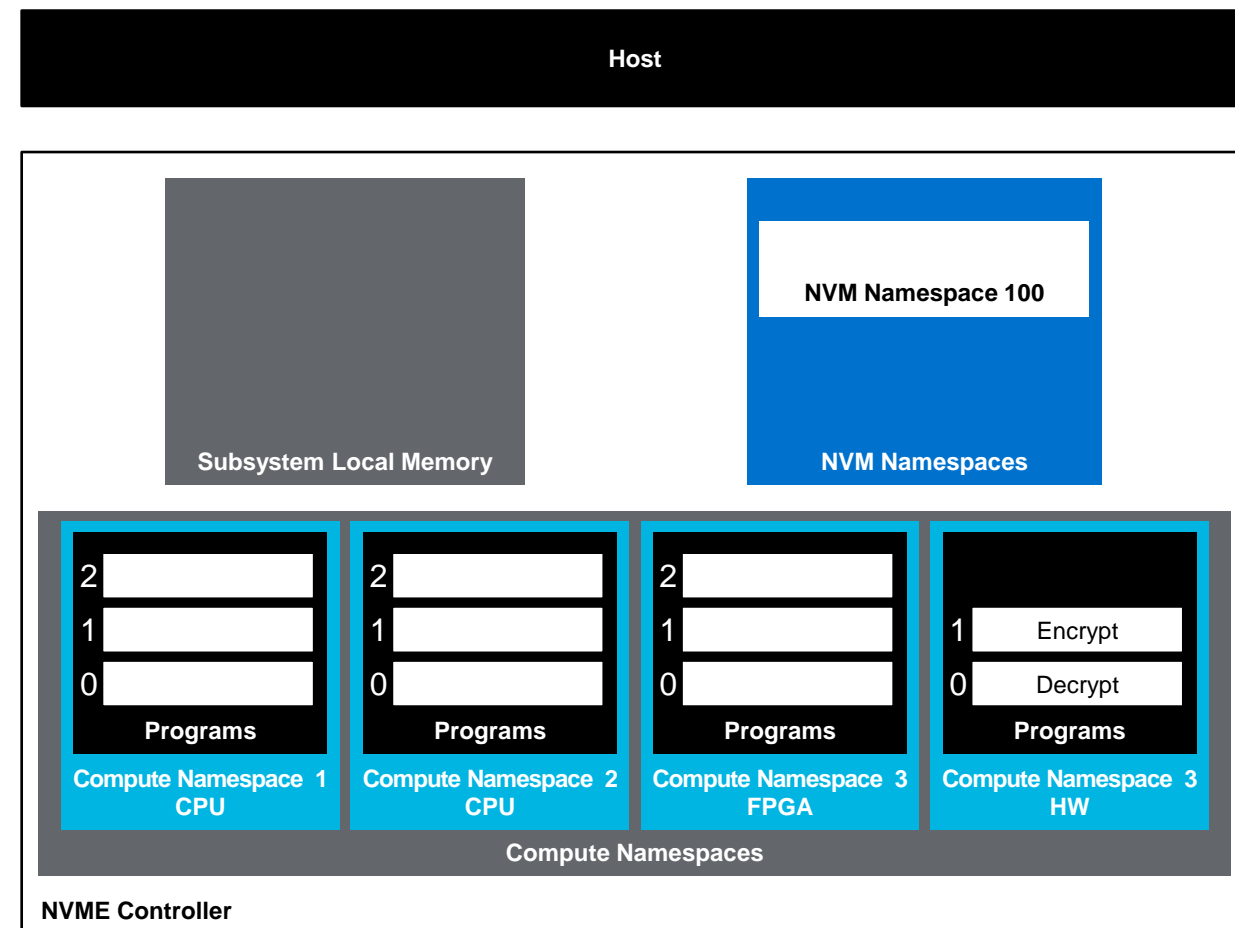


- Offloads program execution from host to NVM subsystem
- Programs compute on data in Subsystem Local Memory (SLM)
- New command set (TP4131) is being developed to copy data between
 - SLM and host
 - SLM and NVM namespaces



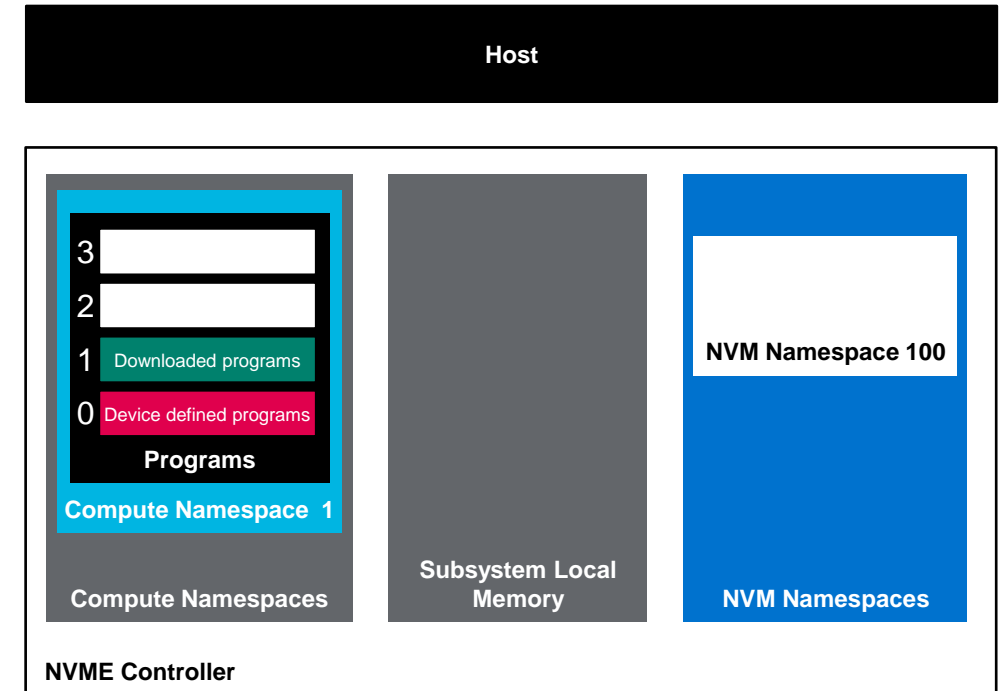
Computational programs

- **Standardized usage**
 - Called with parameters
 - Runs to completion
 - Orchestrated by host
- **Operates only on data in SLM**
- **Hardware or Software**
 - HW accelerator
 - FPGA
 - CPU



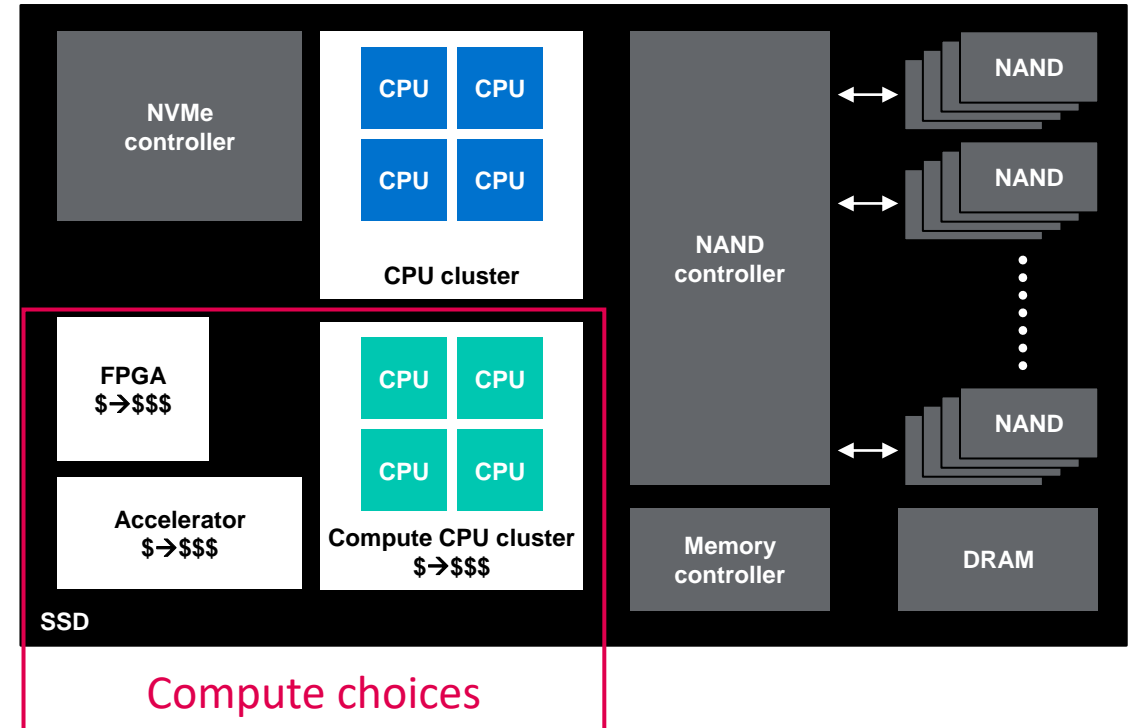
Device defined and downloadable programs

- **Device defined programs**
 - Functionality defined by manufacturer
 - Hardware or software
 - Common functions: Encryption, compression
- **Downloadable programs**
 - Loaded by the host
 - Software running on CPU
 - eBPF
 - FPGA programs



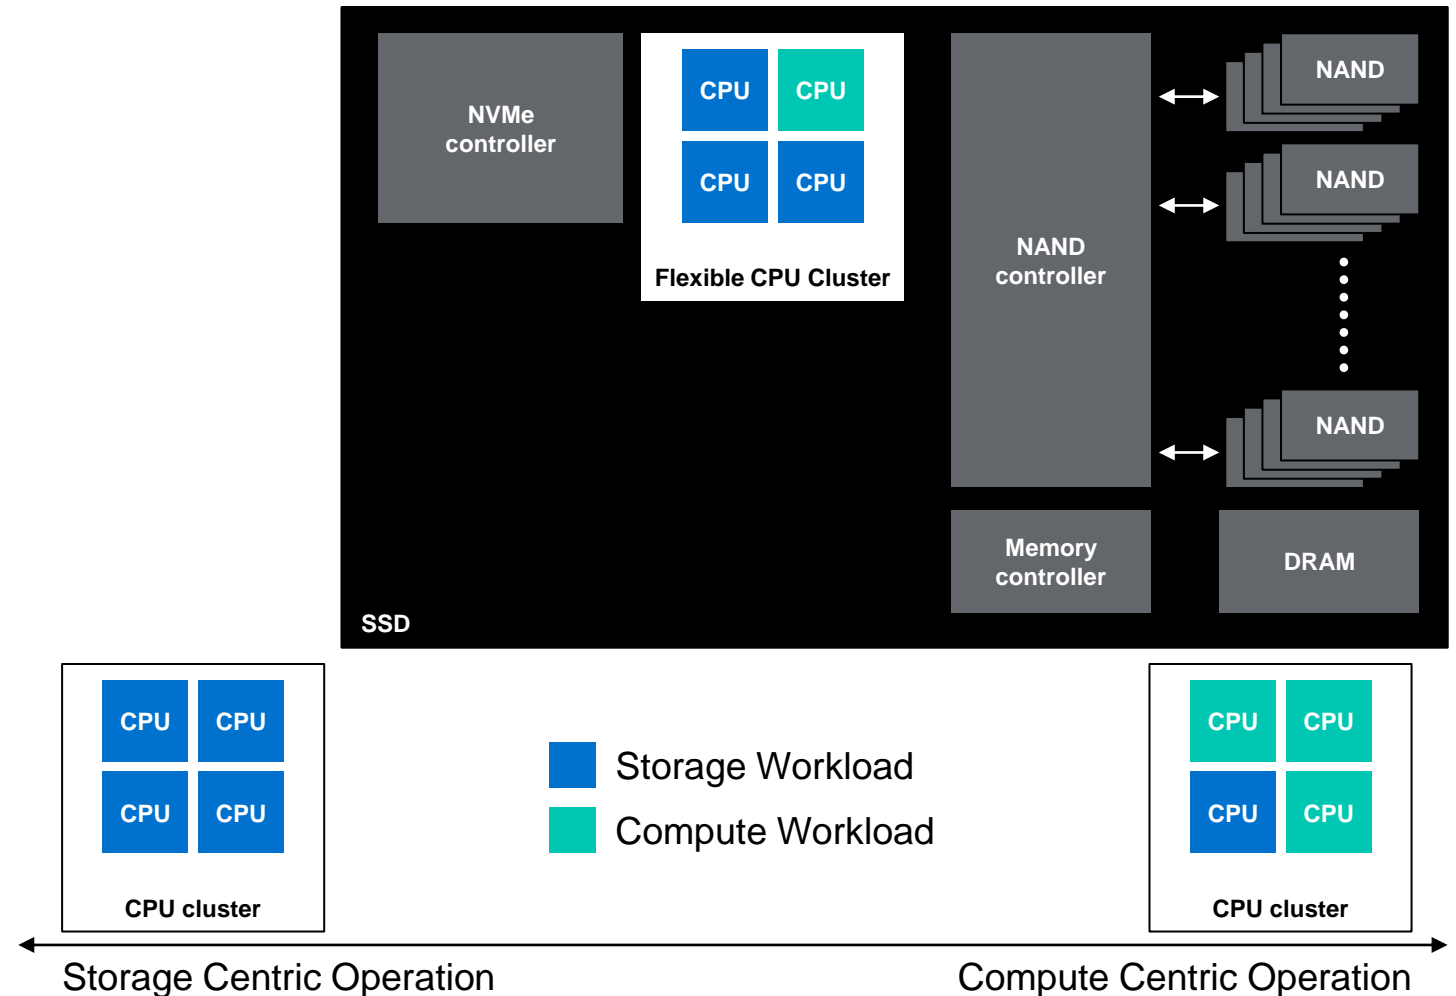
Compute inside SSD

- **Computational storage drive**
 - Compute power vs cost
 - FPGA
 - Hardware accelerator
 - CPU cores for compute
 - Memory
- **Niche applications vs General purpose**



Flexible controller

- Workload balancing
- Flexible CPU usage
 - Storage vs Compute
- Background processing







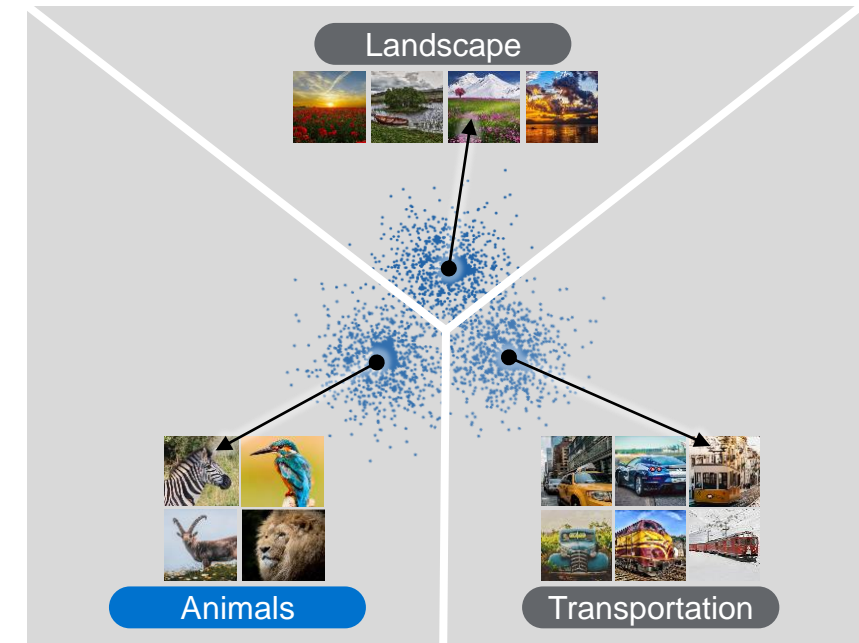
Types of offline compute operation

- Deduplication
- Filtering
- Aggregation
- Classification

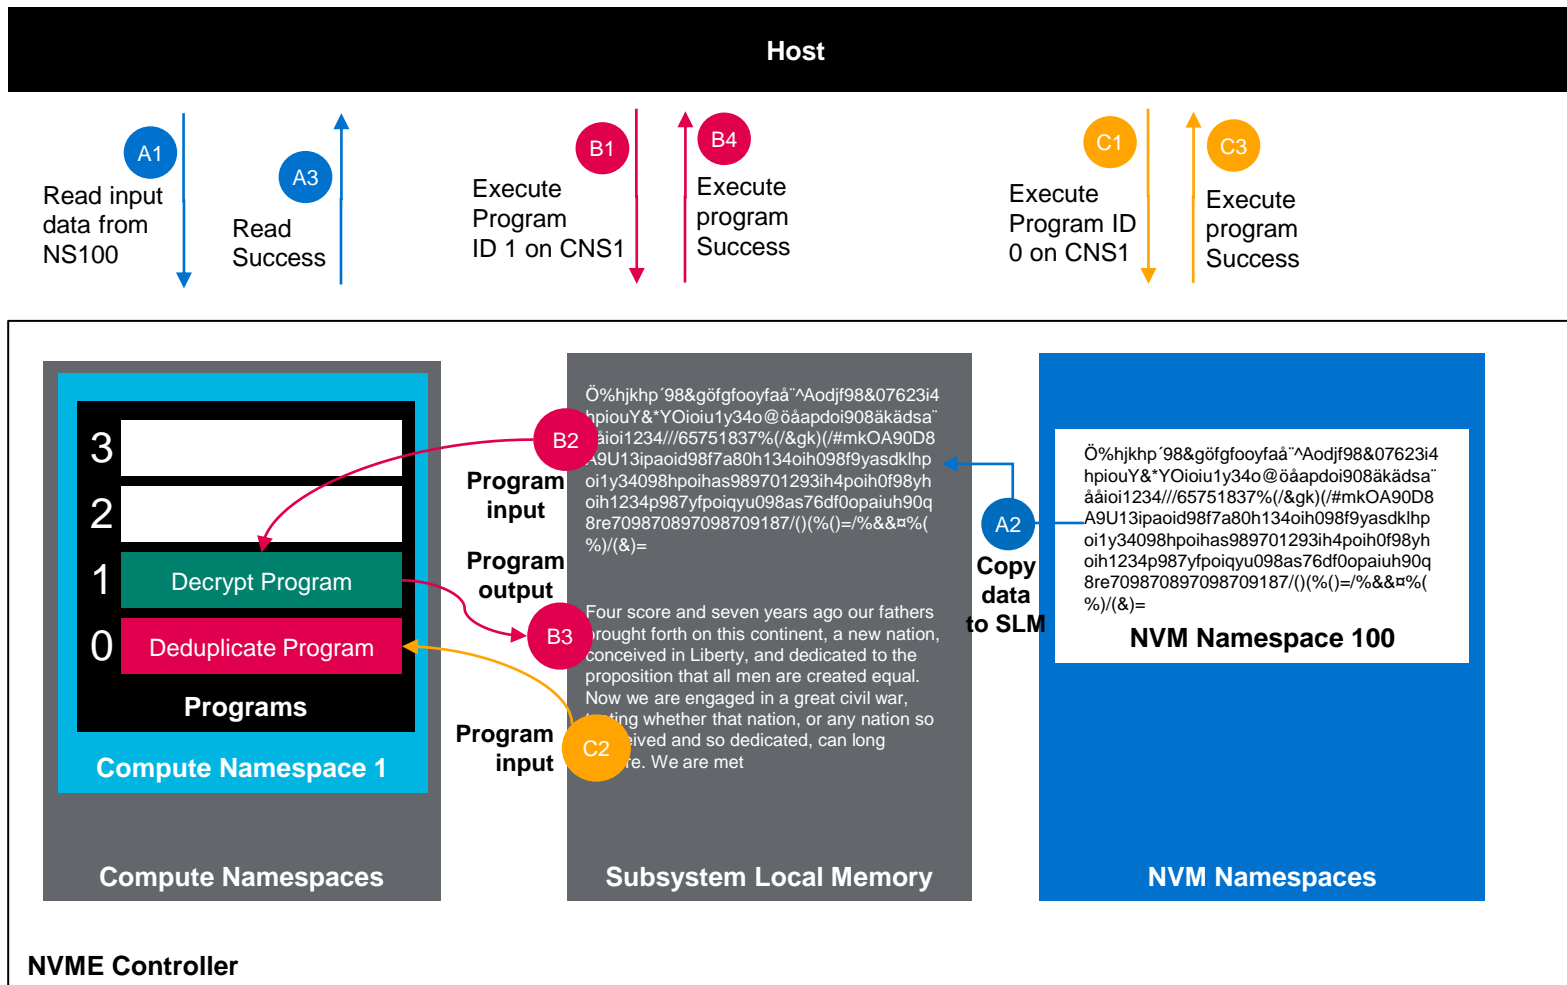
Name	Street	City	State	Email	Ph.
Andy Banner			NH		
Bruce Wayne			NY		
David Clapton			CA		
Jennie Adams			NY		
Tom Johnson			NY		
Tony Stark			TN		



Name	Street	City	State	Email	Ph.
Bruce Wayne			NY		
Jennie Adams			NY		
Tom Johnson			NY		



Simple offline deduplication example



Flow steps

- A** Read encrypted data into SLM
- B** Execute Program ID 1 on CNS1
- C** Execute Program ID 0 on CNS1

Repeat for next
block of data

Summary

1

NVMe TP4091 computational programs command set is under development

2

Standardization can propel adoption – CS for the masses

3

Flexible implementation with CPUs – storage centric vs compute centric

4

Heavy duty computational storage via dedicated CSPs