Flash Memory Summit

Non-Disruptive Firmware Upgrades for Enterprise SSDs

Radjendirane Codandaramane
Sr. Manager, Applications
Microsemi

Agenda

- Enterprise NVMe SSD
- Firmware Upgrades
- Problems
- Solutions

Flash Memory Summit

Enterprise NVMe SSDs

- NVMe SSD controller firmware upgrades are unavoidable
  - Bug fixes, enhancements, new features, and so on
- Enterprise and data center applications also require non-disruptive upgrades (NDU)

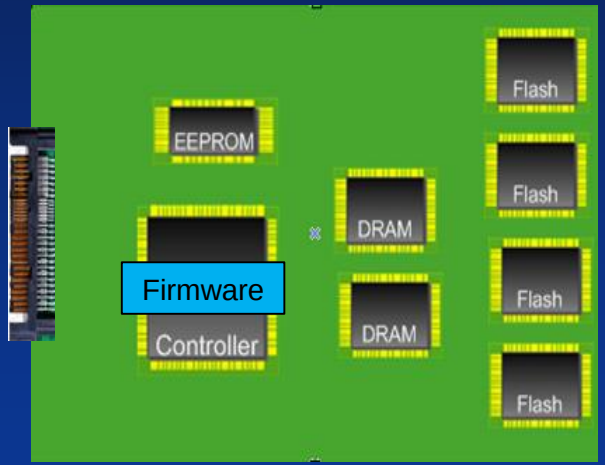

Controller FW Upgrade Challenges

- System Downtime
- Reboot
- Impact to Applications
- Data Loss
- Configuration Loss
- Connectivity Loss
- Lower Performance
- Security Compromise
- Impact to Other Resources
- Time Consuming
- Complex Procedures
- Errors

NVMe Standard Features

Name	Type	Host	Controller
Controller Capabilities	Register	Used to know the capabilities (such as timeout value)	
Controller Status	Register	Used to know when the FW activation is done	Sets CSTS.PP and sends async event to host
Identify	Admin Command	Used to know the controller features	
FW Image Download	Admin Command	Issues the FW image	
FW Commit	Admin Command	Activates the FW image	

NVMe Standard: Identify Command

Byte	Bit	Field	Description
71:64		FR	Firmware revision contains the active FW revision
95:92	9	OAES	Optional async events supported. If bit 9 is set to 1, controller can send FW activation notices to host (only after host allows through the set feature command).
257:256	2	OACS	Optional admin command support, bit 2. If set, controller supports FW commit and image download commands.
260	4:0	FRMW	FW updates: Bit 4, if set to 1—controller supports FW activation without a reset Bits 3:1—number of FW slots Bit 0, if set to 1—first FW slot is read only
271:270		MTFA	Max. time for FW activation: number of 100 ms units required for controller to activate the FW without a reset
319		FWUG	FW update granularity: indicates the minimum granularity and alignment of the FW image download data (0x1= 4 KB, 0x2= 8 KB, 0x00= no info, 0xFF= no restriction)

NVMe FW Update Process

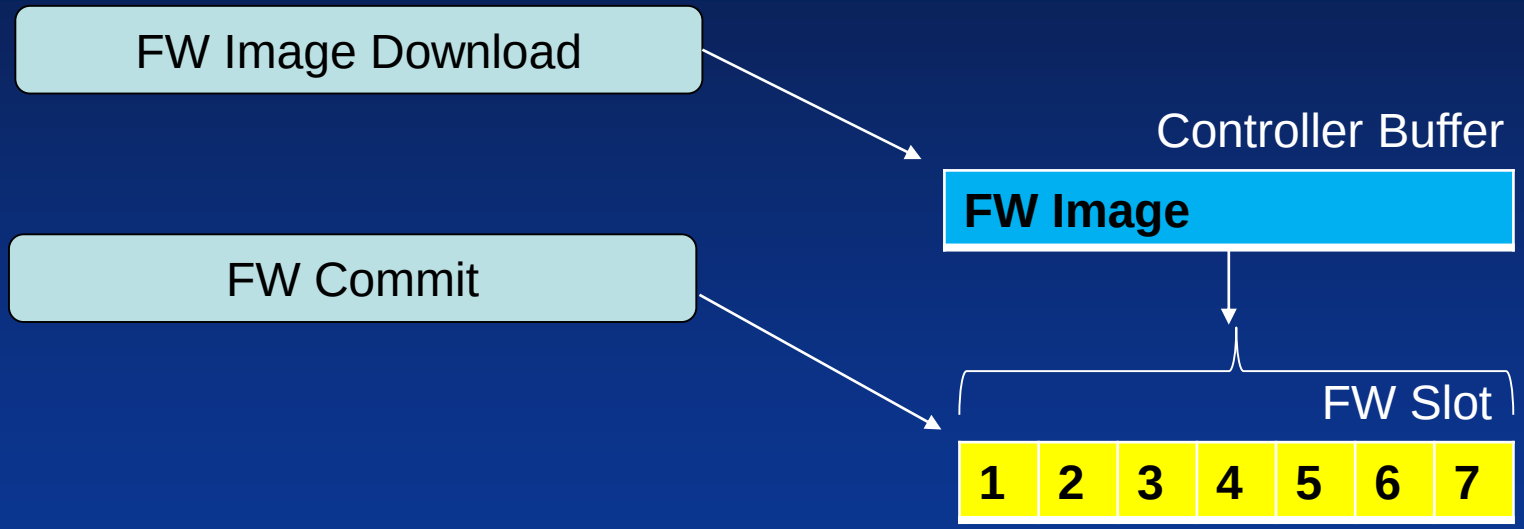

Flash Memory Summit

NVMe Controller Features

- Multi-Core Processors
- NVMe Queue Preservation
- Dual Port
- Security, Authentication, Integrity Checking
- Failsafe

Multi-Core Processors

- Distributes the workload
- Dedicates a core for front-end (PCIe, NVMe) processing
- Switch cores to ensure front end is alive to the host, while updating the FW to other cores

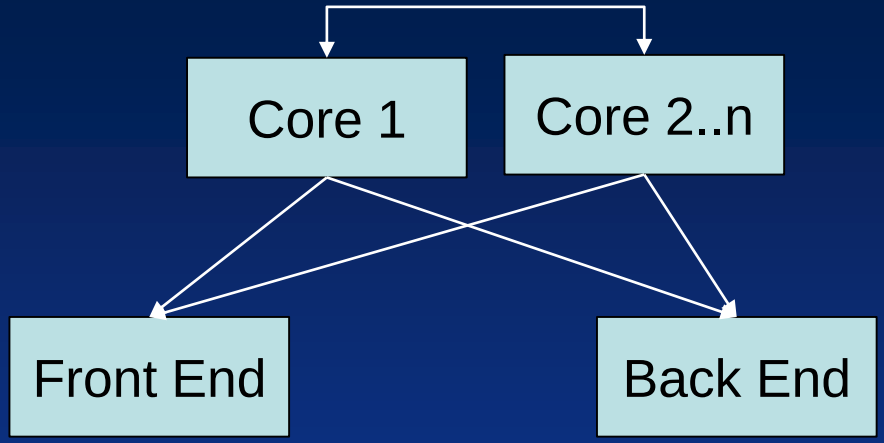

NVMe Queue Preservation

- Ensure HW mechanisms exist to save and restore NVMe queue states during the upgrade process
- Use controller status pause to notify the host

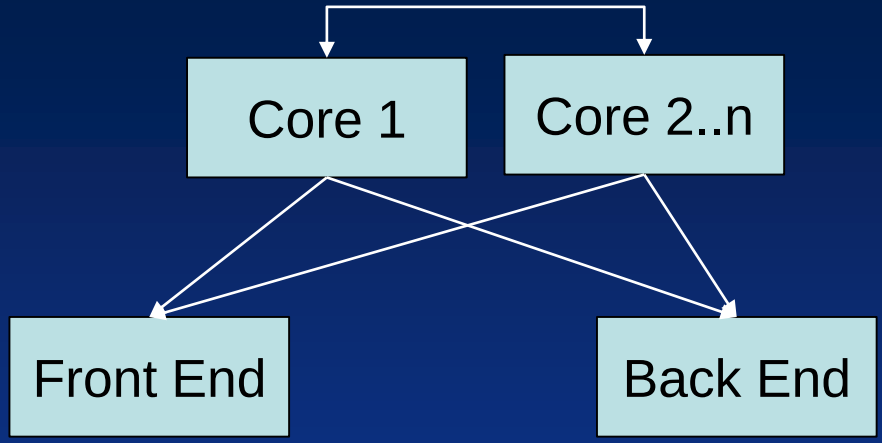

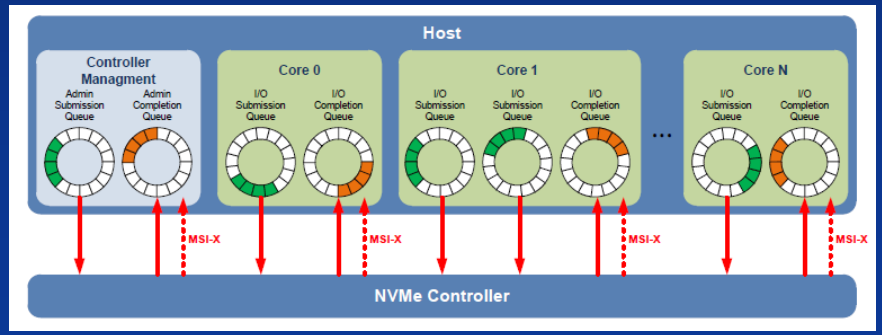

Dual Port

- Ensure the FW update command received in one port does not disrupt the other port through dedicated HW resources for dual-port operation
- Same applies to SR-IOV resources

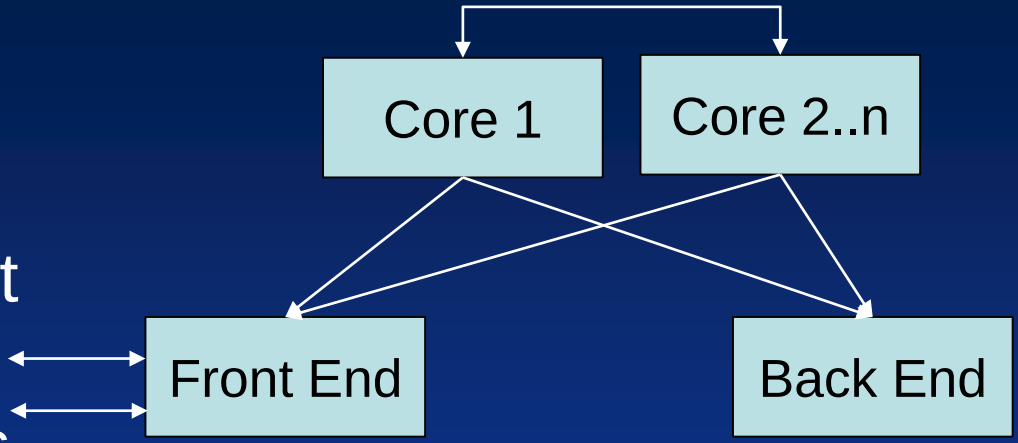

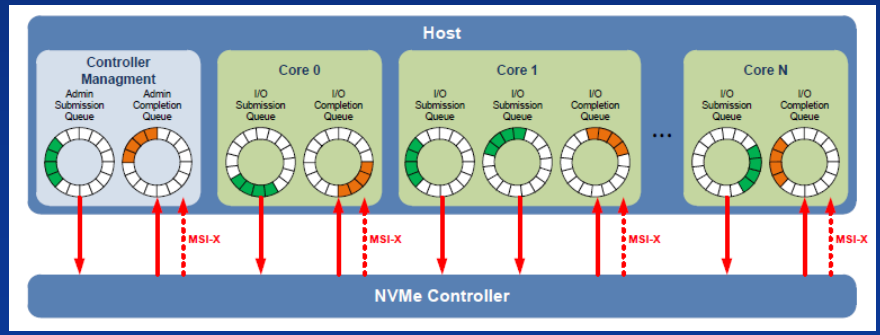

Security, Authentication, Integrity Checking

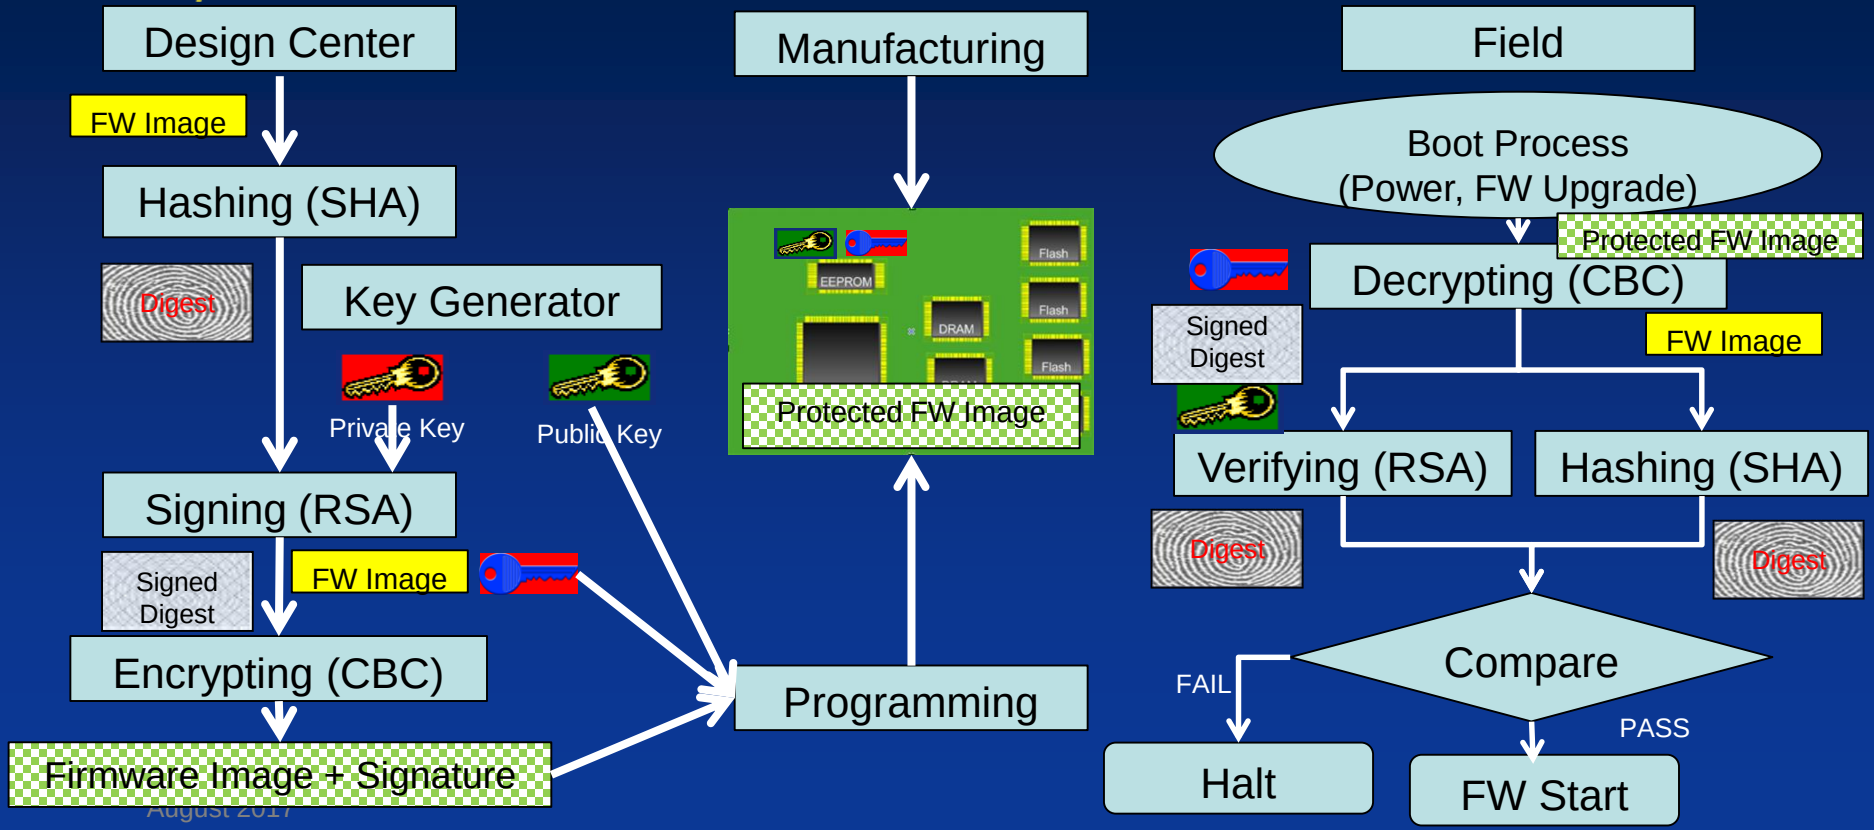

Fail Safe

- Protection mechanisms against power loss, user abort, connectivity loss, ASPM, security check fail, internal failures, and so on
- Use FW slots and have a baseline read-only FW image for recovery

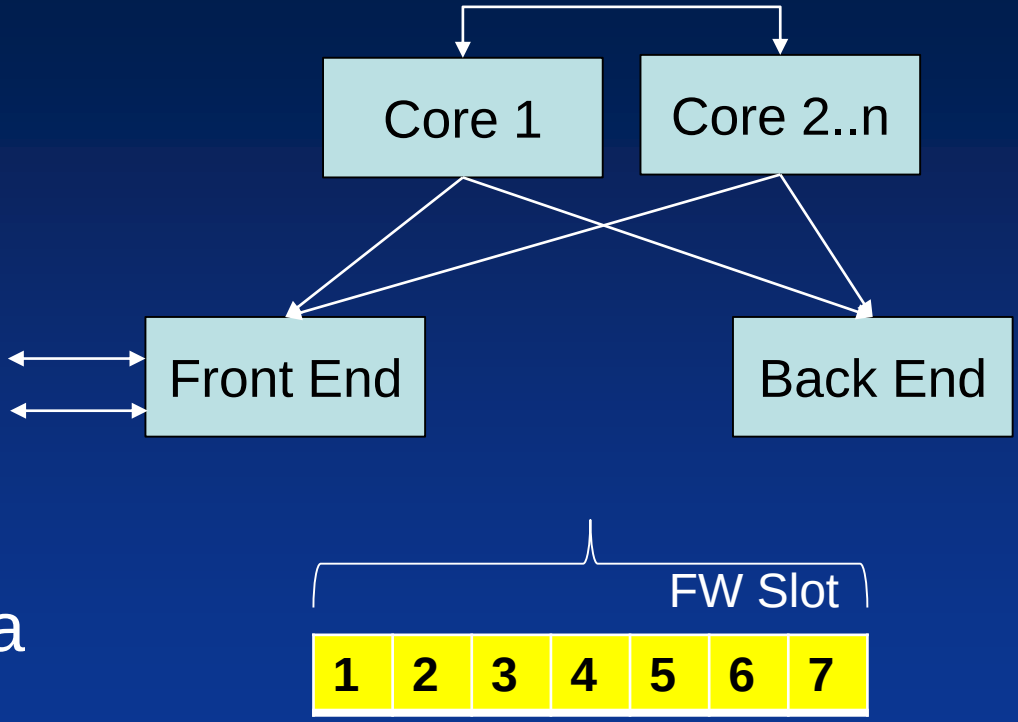

Summary

- Challenges of non-disruptive firmware upgrades in Enterprise SSDs
- NVMe mechanisms for NDU
- Controller mechanisms for NDU

Flash Memory Summit

Thanks!!

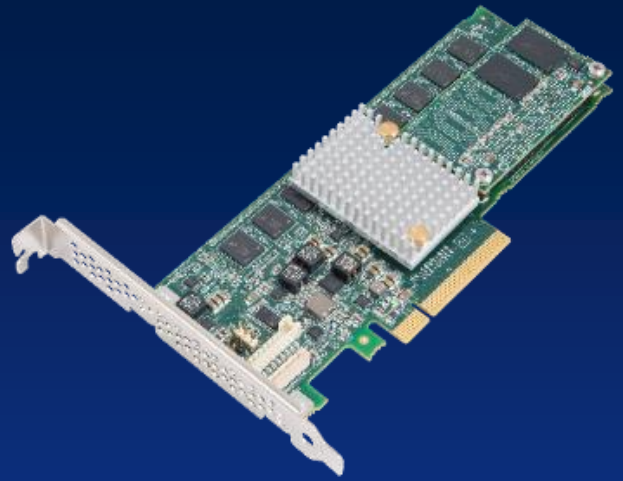

- Microsemi offers a range of products for enterprise storage and data center applications
- Come and visit us at booth #213
- www.microsemi.com