Flash Memory Summit

UNIVERSITY OF PATRAS
ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΩΝ

An Advanced Flash Emulator for Designing Today's High-Capacity Controllers

Th. Antonakopoulos, **N. Toulgaridis**, M. Varsamou,
E. Bougioukou and T. Petropoulos

University of Patras

Greece

Overview

- Why NAND Flash emulation at system level is needed?
- What are the challenges of emulation at system level?
- The architecture of NAND Flash Emulator
  - System Architecture
  - Memory Organization
  - Low-latency memory access
  - Experimental results
- Beyond the current NAND Flash Emulator

NAND Flash emulation at the system level

- Storage devices
  - Multiple memory ICs organized in channels
  - Multiple channels operating in parallel
  - Large memory capacity per channel (a few 10s or 100s GBs) – Huge capacity at system level (xTBs)
  - High IO rates and fast response time, especially when a page is read
  - Complicated functions (i.e. wear leveling, workload balancing) in the storage controller
- Full system prototyping and testing before the actual memory chips are available, based only on their specifications.
- Evaluate under different loading conditions the performance of the implemented algorithms
- Reduce time-to-market for the storage device when new memory chips become available

There is a need for an NAND Flash Channel Emulator that can emulate the whole memory capacity of a device and respond in real-time according to the NAND Flash specs

The main challenges of a system Flash emulator

UNIVERSITY OF PATRAS
ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΩΝ

- **Emulate the whole system capacity**
  - Single board emulators have limited fast memory capacity (x10GBs)
  - Storage systems have multiple channels with multiple dies/channel and their total capacity ranges from x100GBs up to xTBs.

  Solution: Exploit the DRAM capacity of server motherboards
    - Directly accessed by the host processor, indirectly accessed by devices attached to PCIe slots
    - Access is affected by the used host processor (number of DRAM controllers, internal data paths)

- **Respond in real-time according the Flash IC specs**
  - Data access time in NAND Flash: 30 to 50 usecs and the page transfer time depends on the NAND Flash interface supported and the Flash page size.
  - Multiple channels operating in parallel generate asynchronous access requests
  - The latency introduced by the Operating System has to be avoided

  Solution: Use a fast PCIe-based FPGA board where the DUT is attached
    - Custom logic has to be developed for direct access to the host's DRAM.
    - Modular design for supporting different Flash interfaces

The characteristics of the system level Flash emulator

- Uses a low-cost commercially available motherboard that support the maximum possible DRAM (xTB)
- Uses a PCIe card with a high-speed SoC FPGA that acts as the digital front-end to the DUT.
- Modular re-usable system design
- Split the design into two FPGAs boards, if needed.
  - FPGA boards are interconnected using a High-Speed Digital Link (HSDL), i.e. xSFP+.

Advantages:

- Supports minimum latency, high capacity, various I/O NAND Flash interfaces and provides flexibility on the mechanical attachment of the DUT.
- Minimum additional development effort when new memory devices have to be supported.

NAND Flash Emulator Architecture

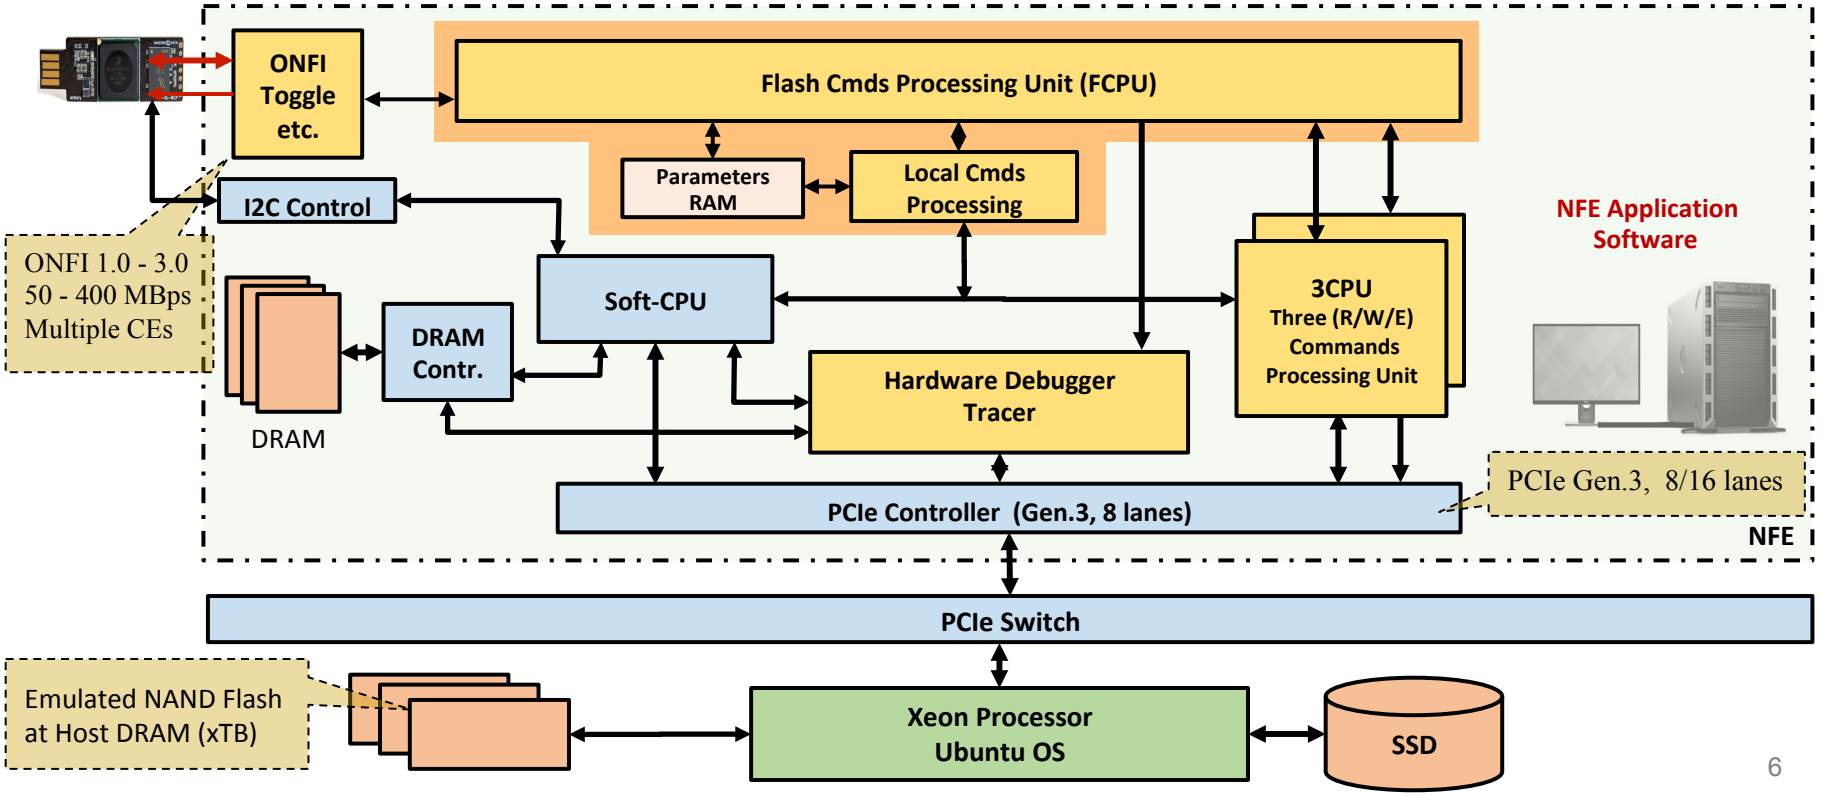

NAND Flash Emulator Architecture

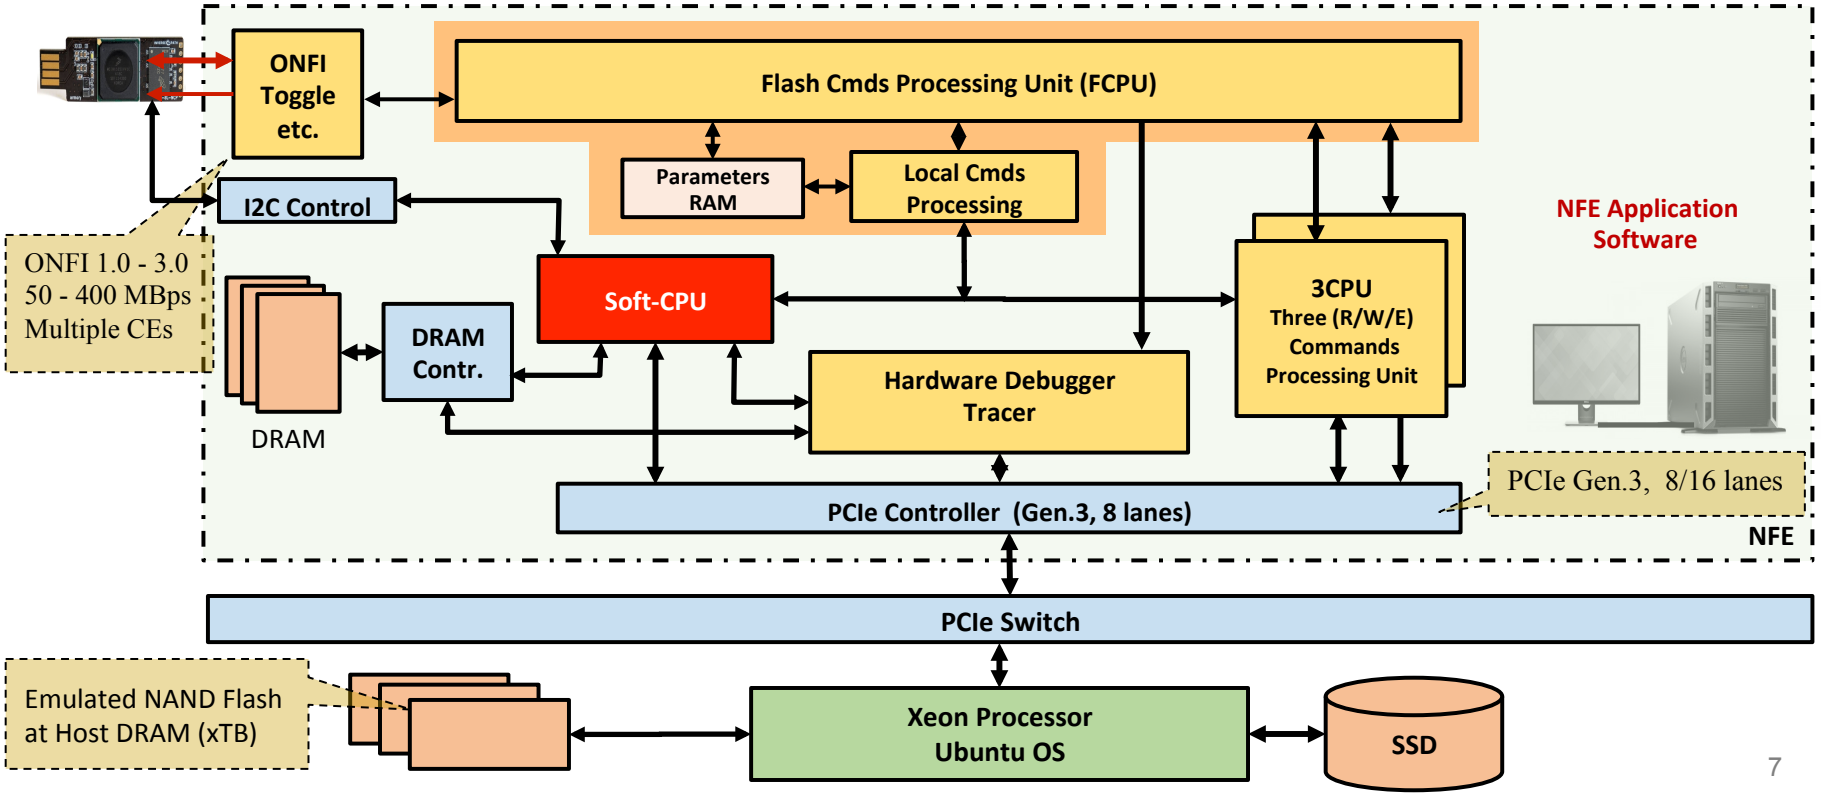

NAND Flash Emulator Architecture

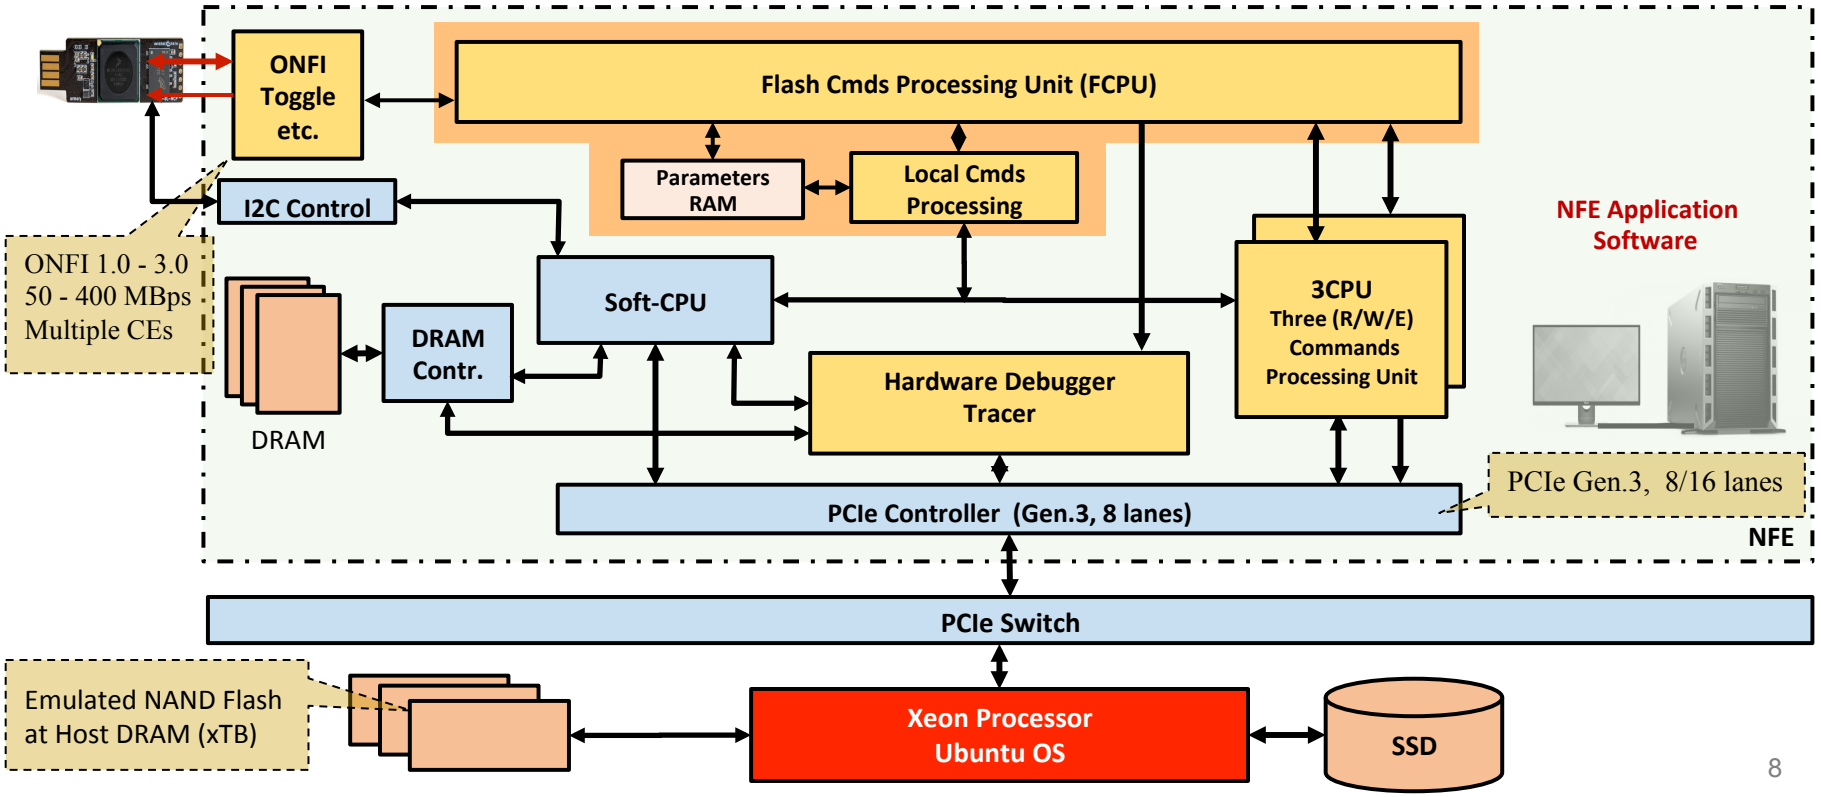

NAND Flash Emulator Architecture

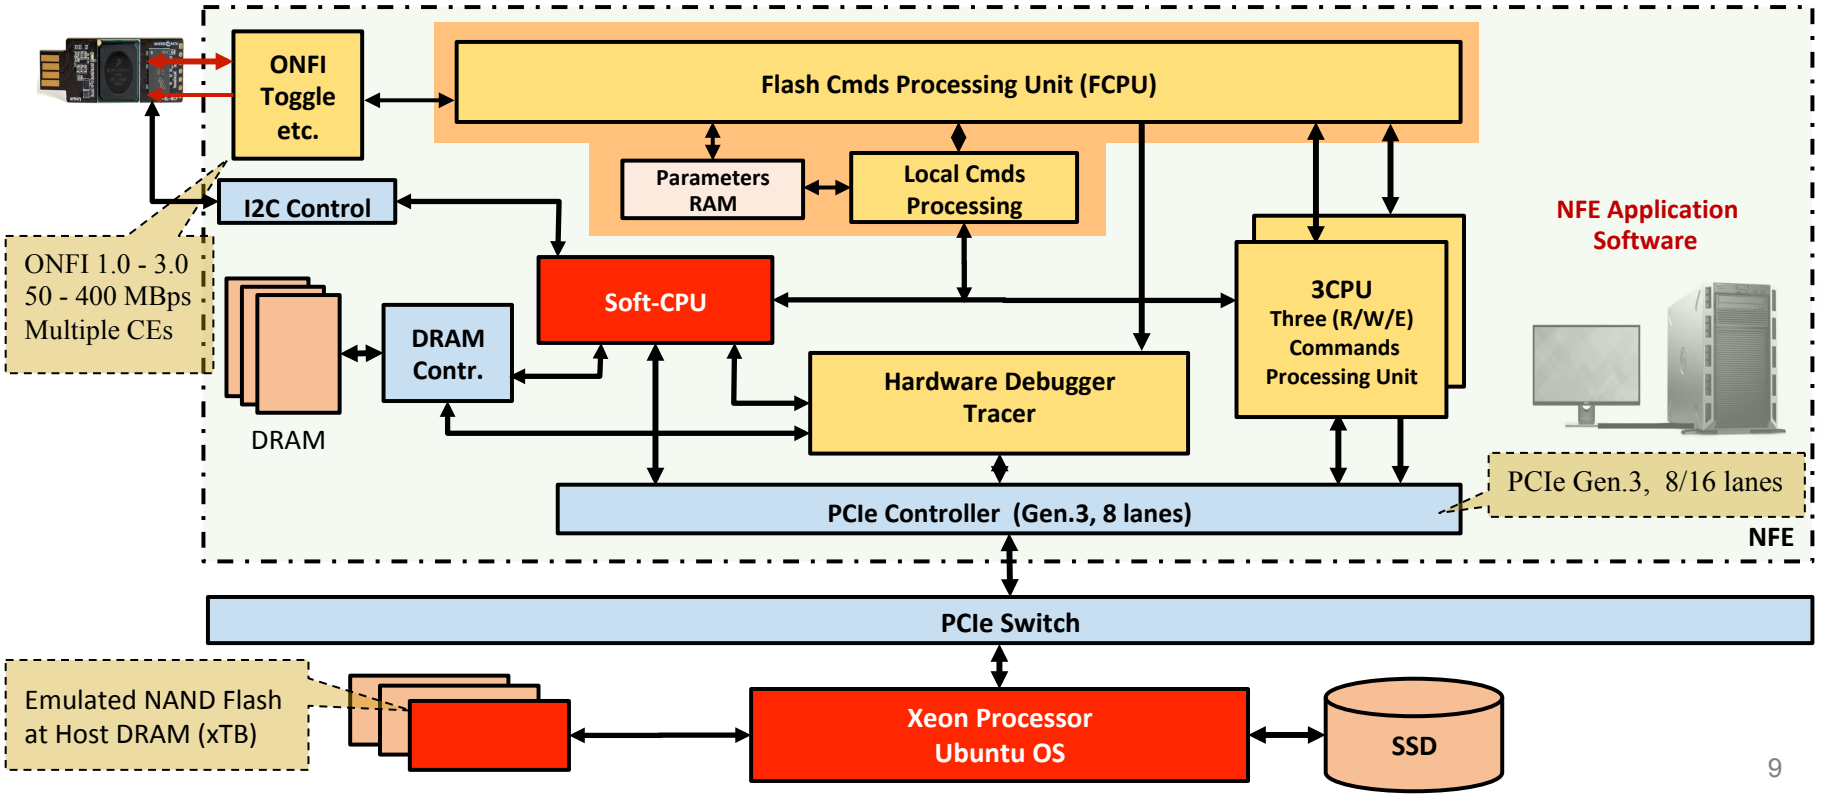

NAND Flash Emulator Architecture

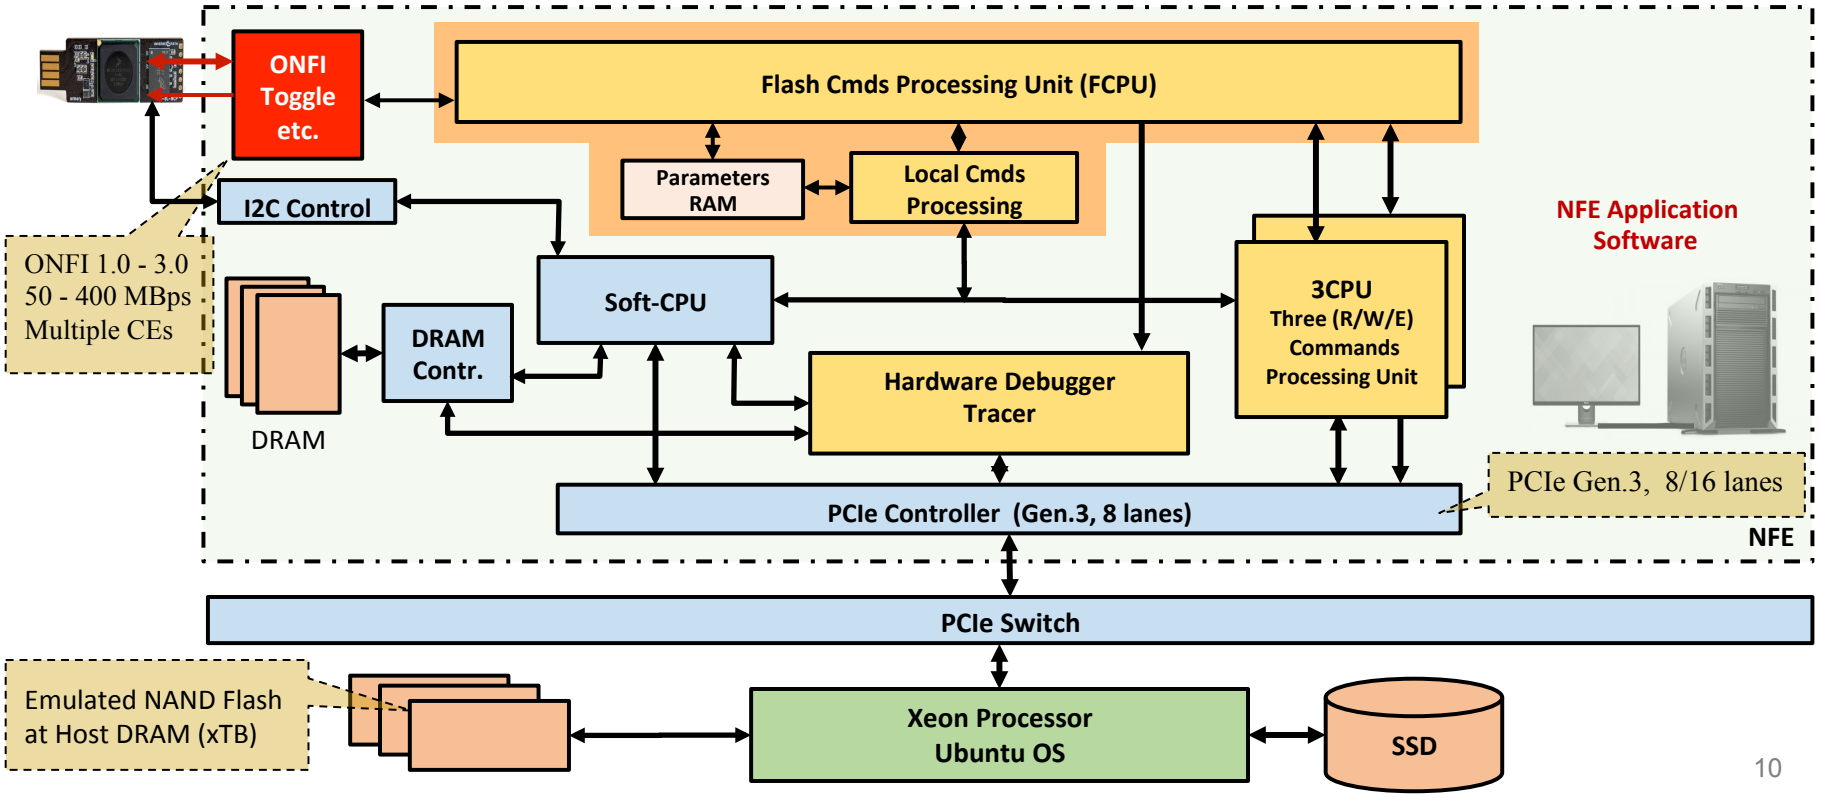

NAND Flash Emulator Architecture

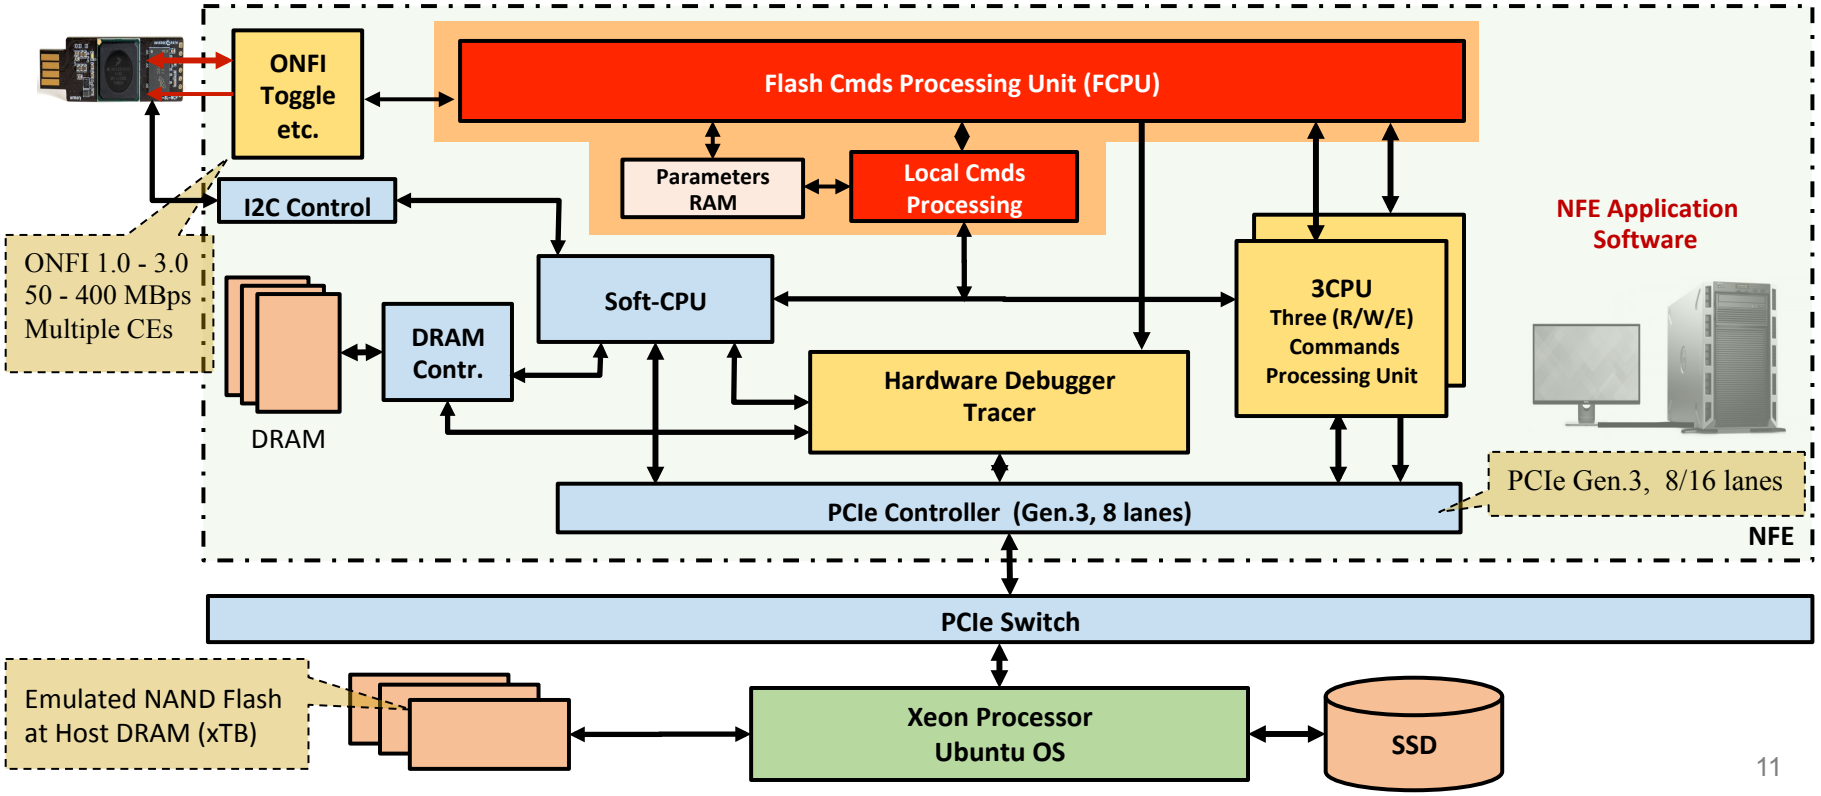

NAND Flash Emulator Architecture

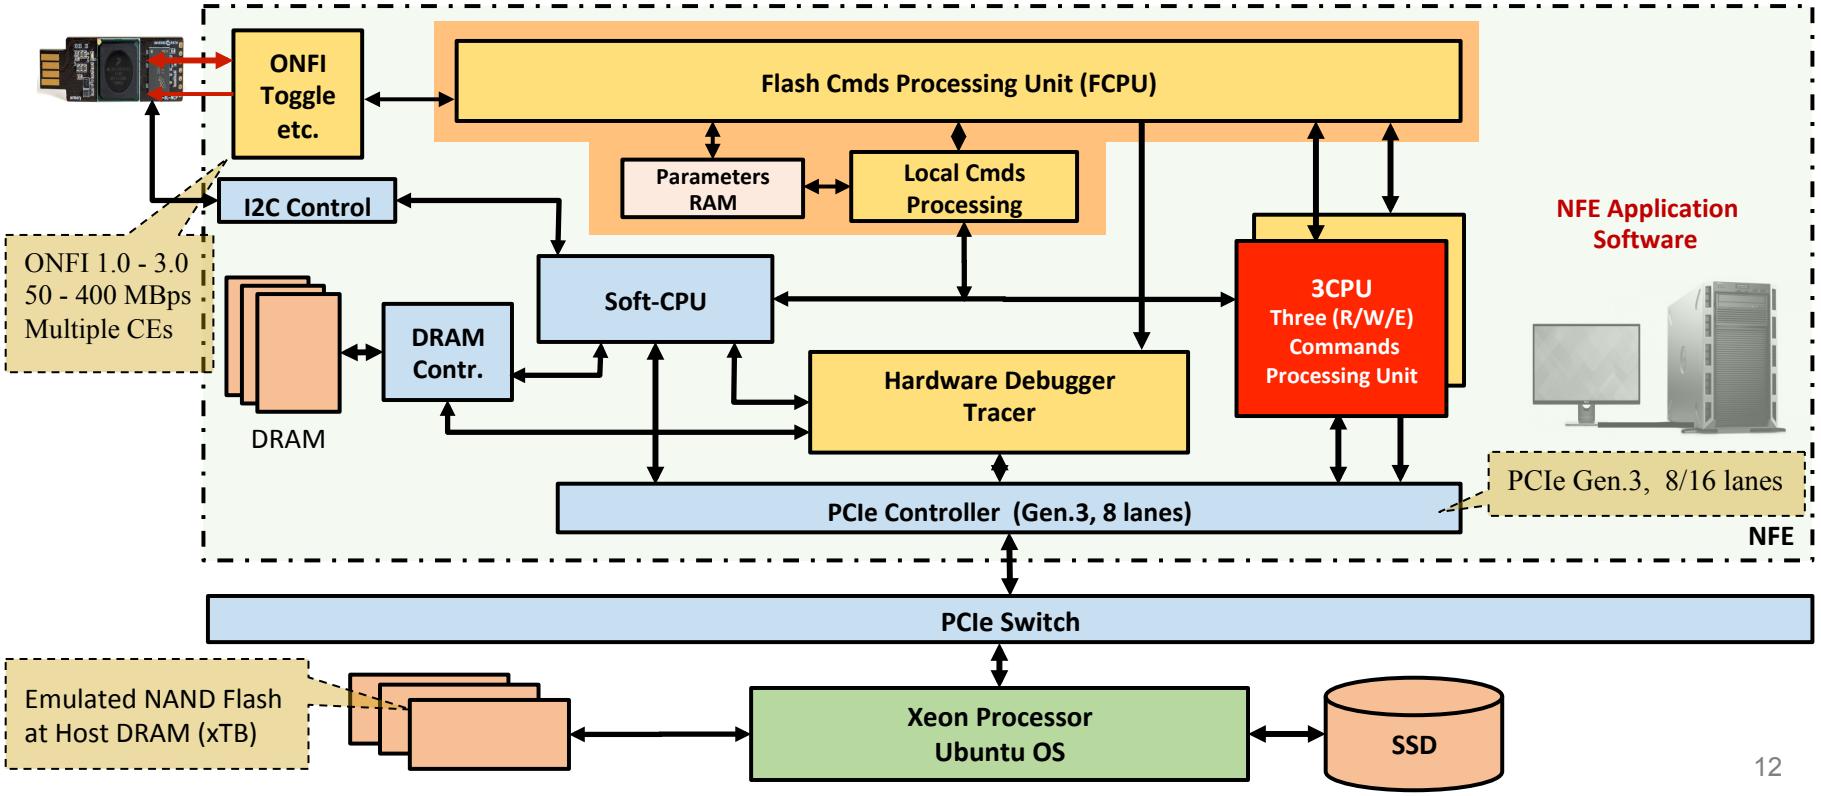

NAND Flash Emulator Architecture

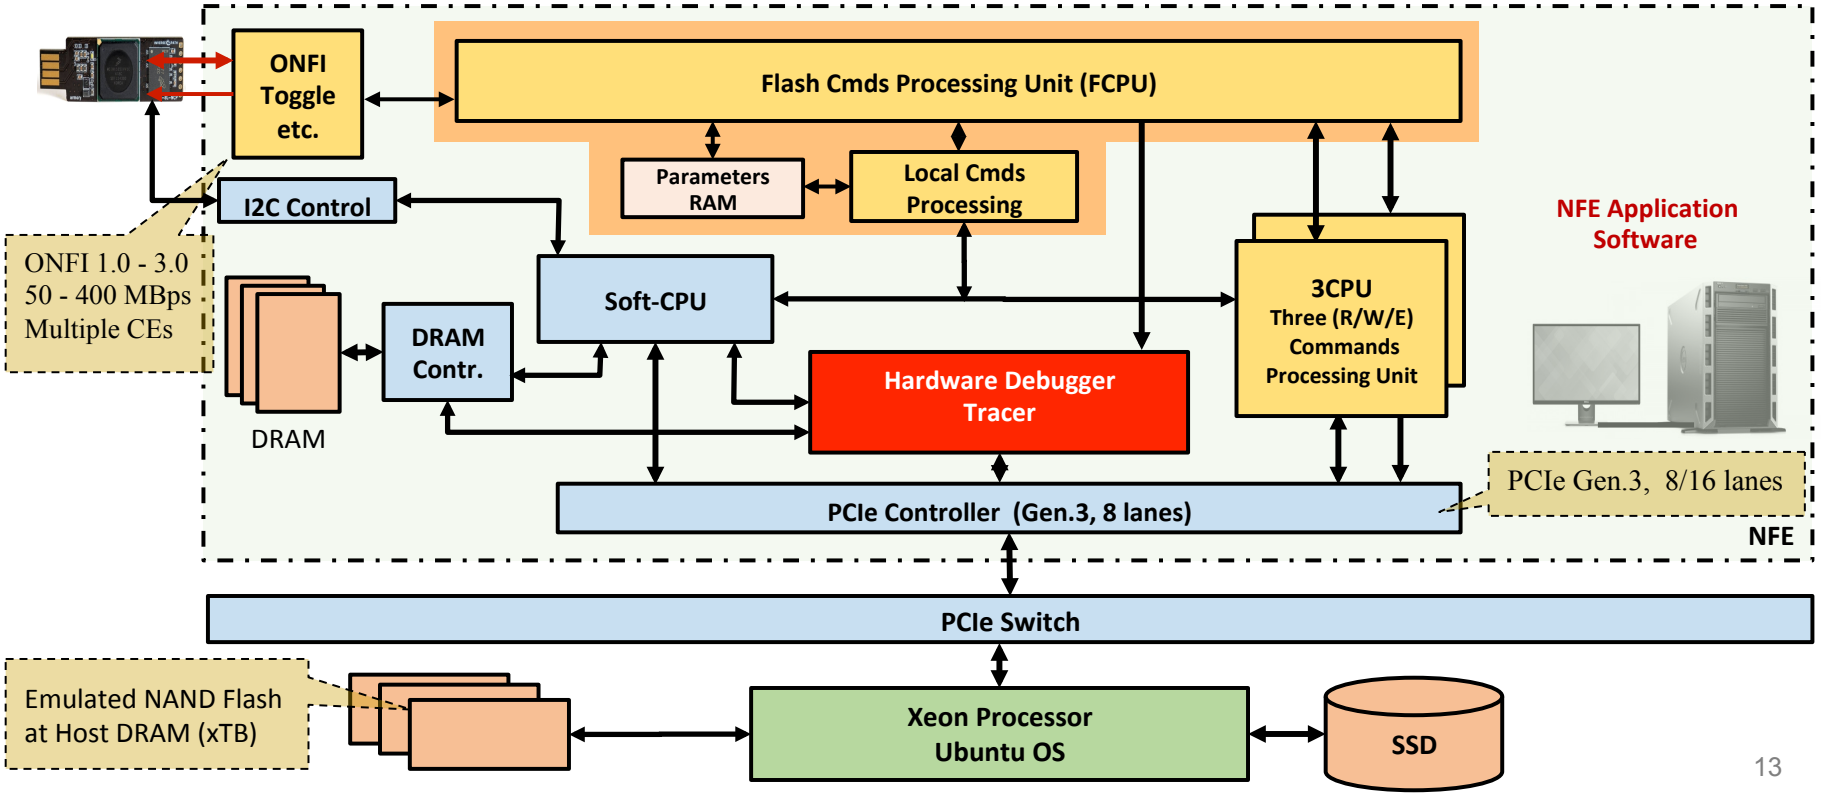

Data Flow in the NAND Flash Emulator

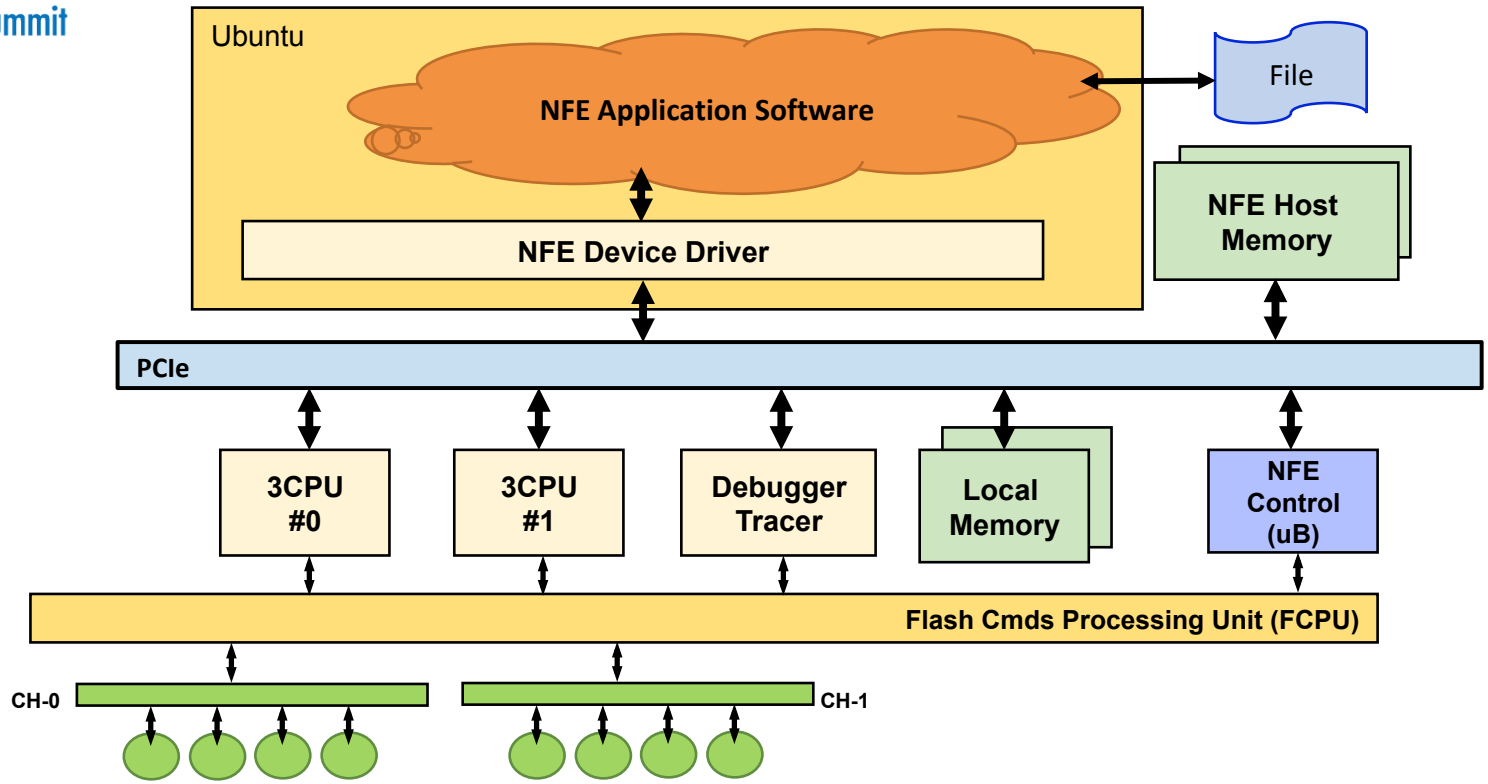

Data Flow in the NAND Flash Emulator

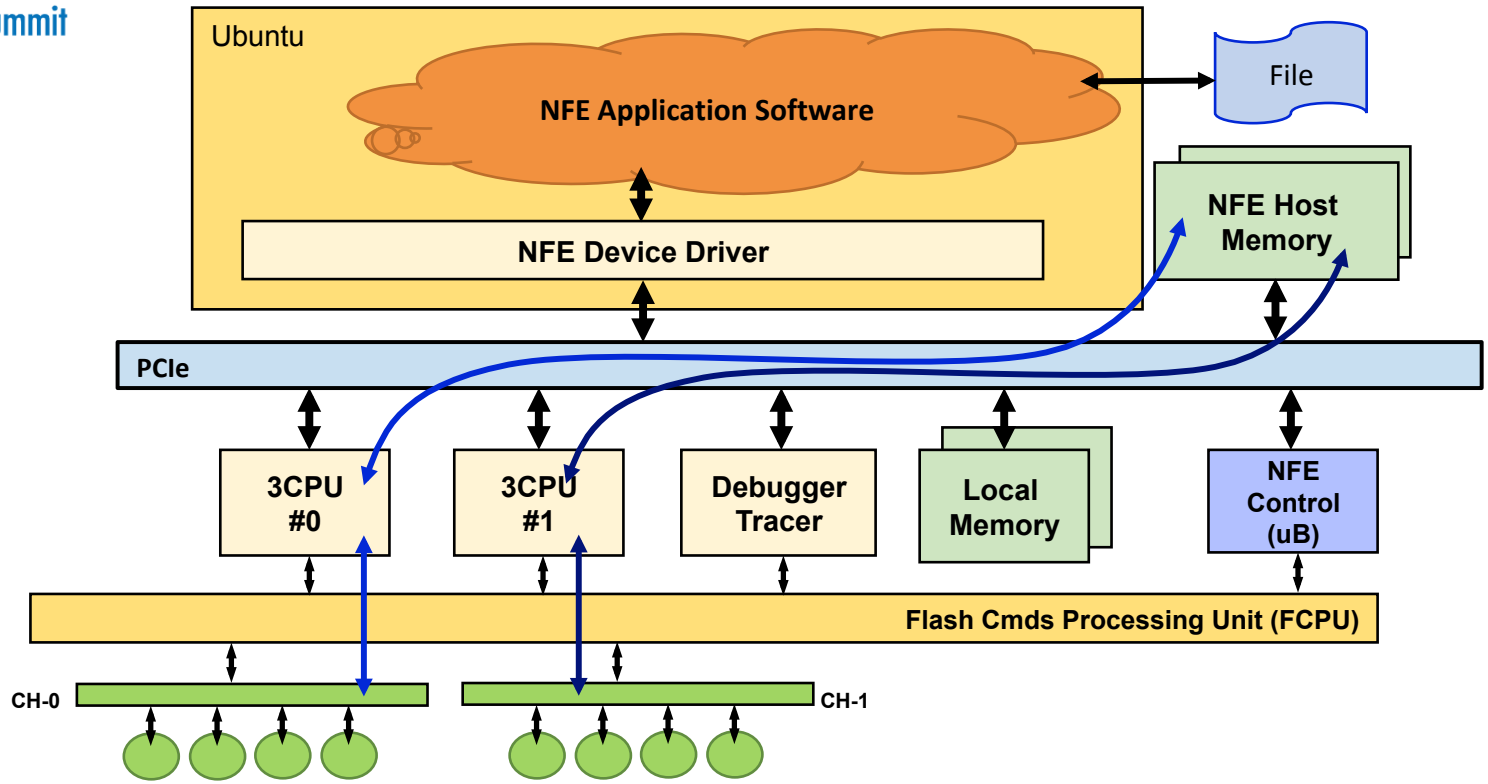

Data Flow in the NAND Flash Emulator

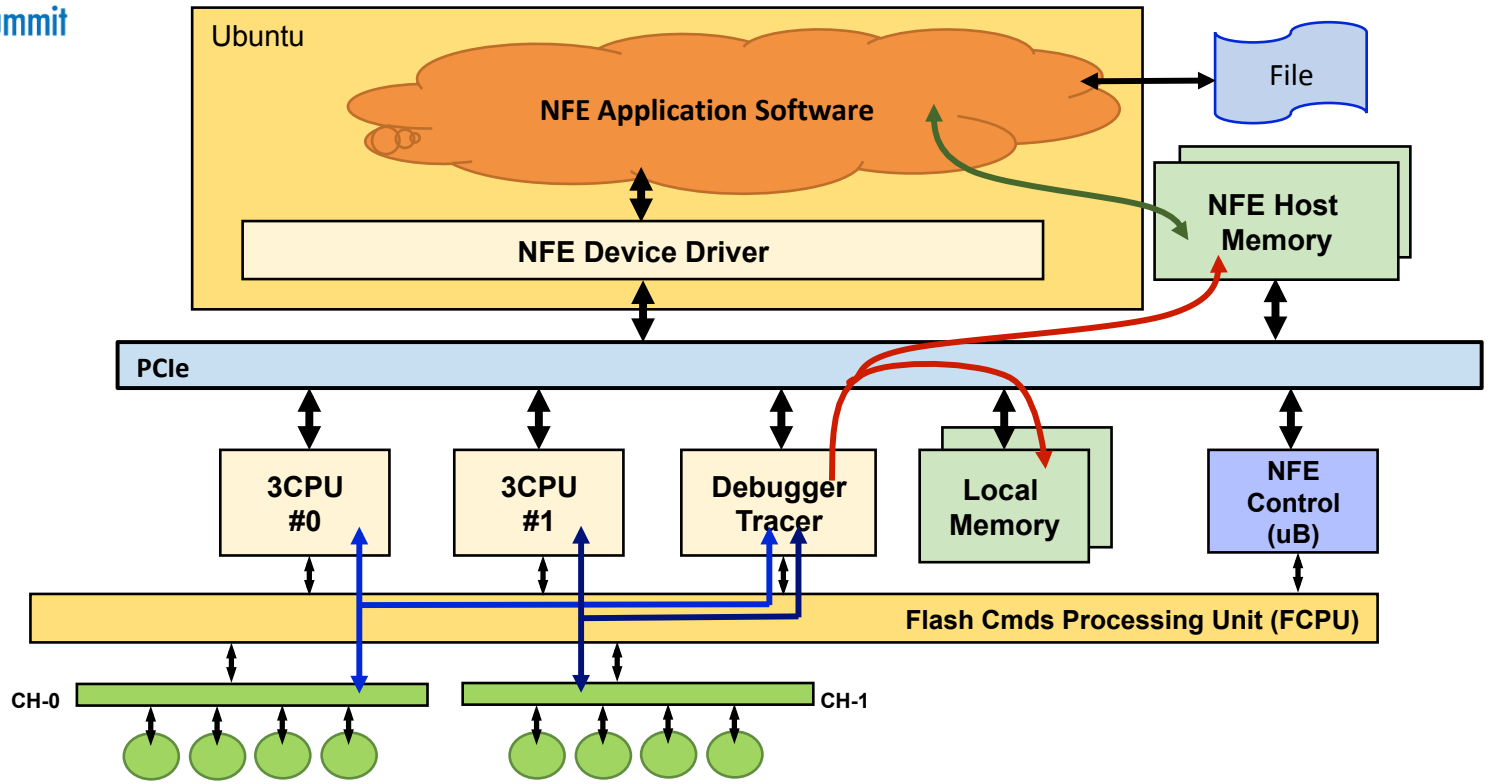

Data Flow in the NAND Flash Emulator

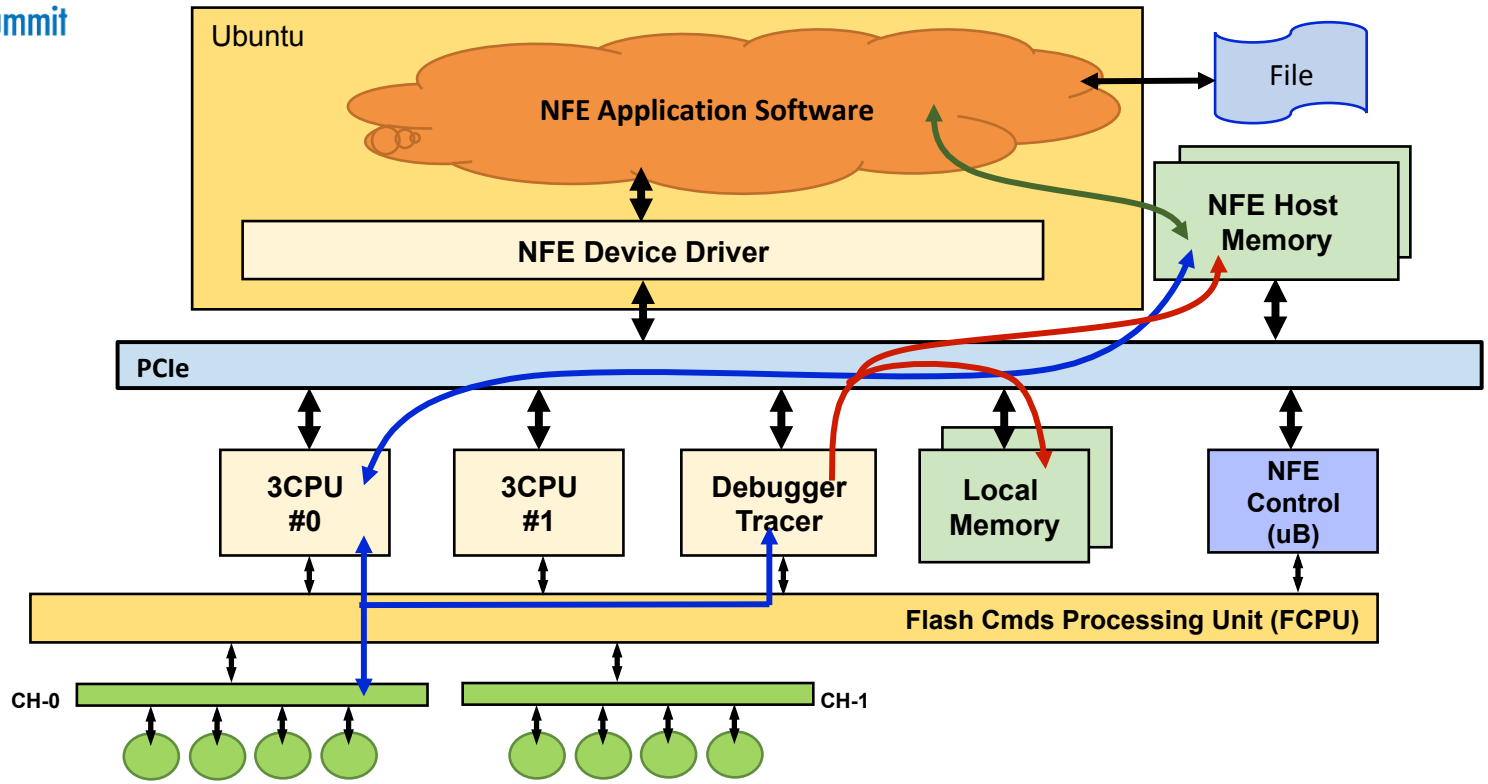

Memory Organization and Emulator Capabilities

UNIVERSITY OF PATRAS
ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΩΝ

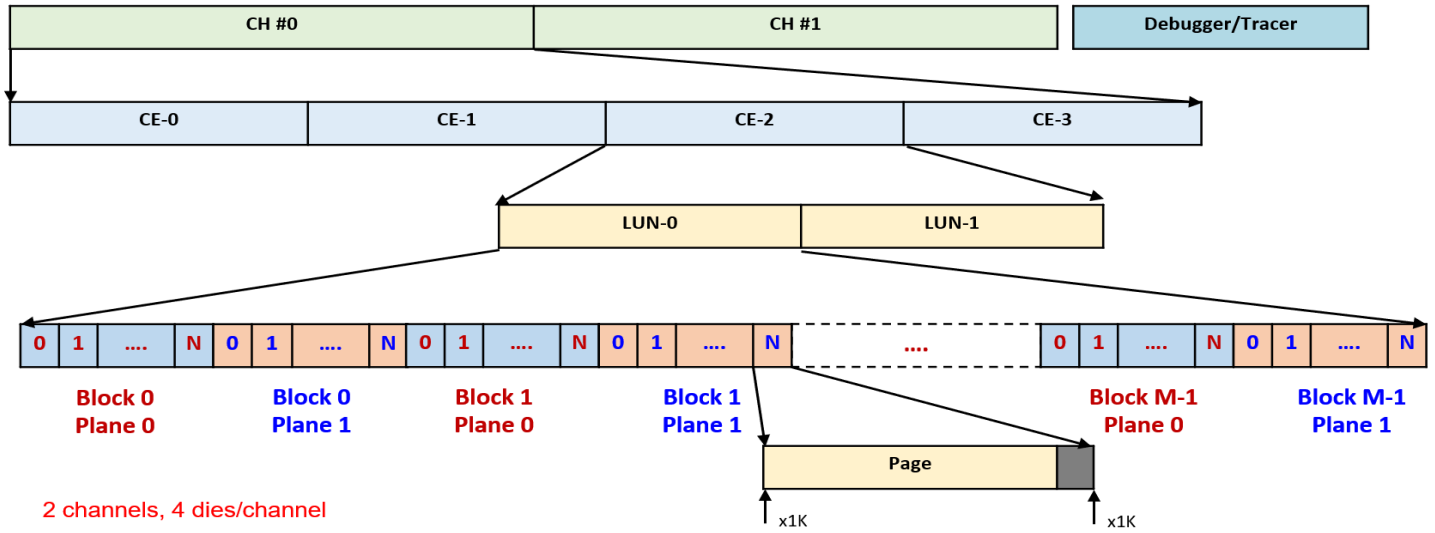

Memory Organization and Emulator Capabilities

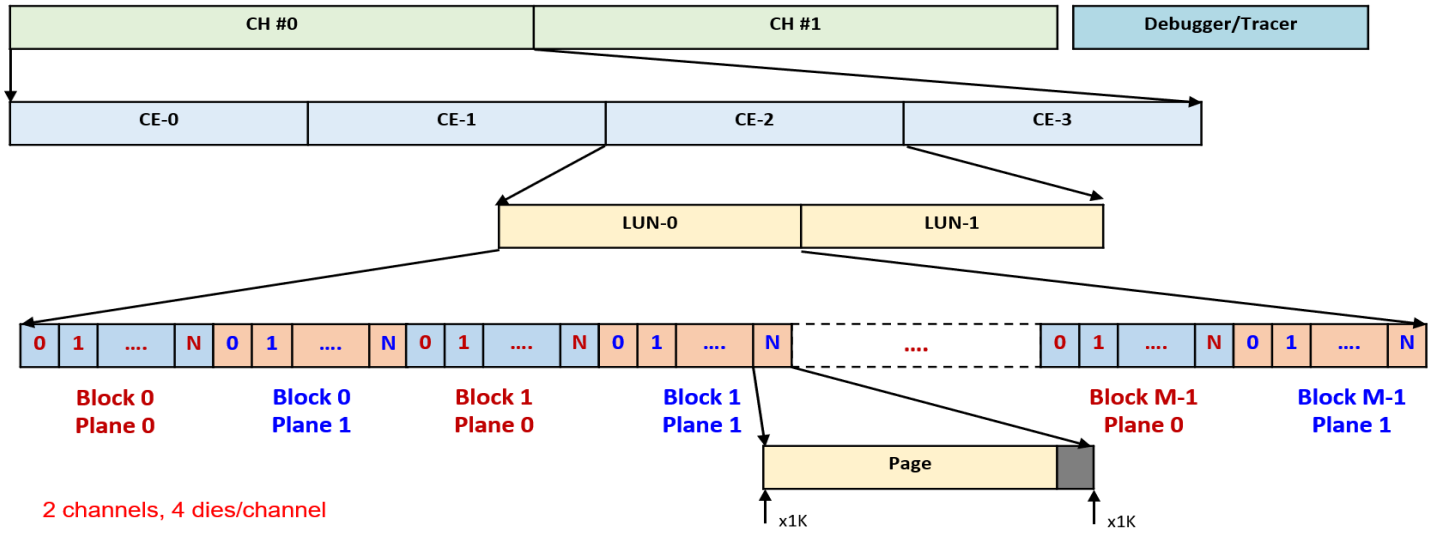

User Page [kB]	Flash Page [kB]	DRAM [x1K]	Total Number of emulated Pages	NAND Flash Capacity [Gbits]	Total Number of Dies	Supported Channels – Dies/channel
4.0	4.22	5.0	**107.3 M**	64	**64**	**16 – 4**
8.0	8.44	9.0	**59.6 M**	128	**32**	**8 – 4**
16.0	17.25	18.0	**29.8 M**	256	**16**	**4 - 4**
16.0	18.16	19.0	**28.2 M**	512	**8**	**4 - 2 or 2 - 4**

512 GB emulated NAND Flash (DRAM: 576 up to 640 GB)

Prototype of the NAND Flash Emulator

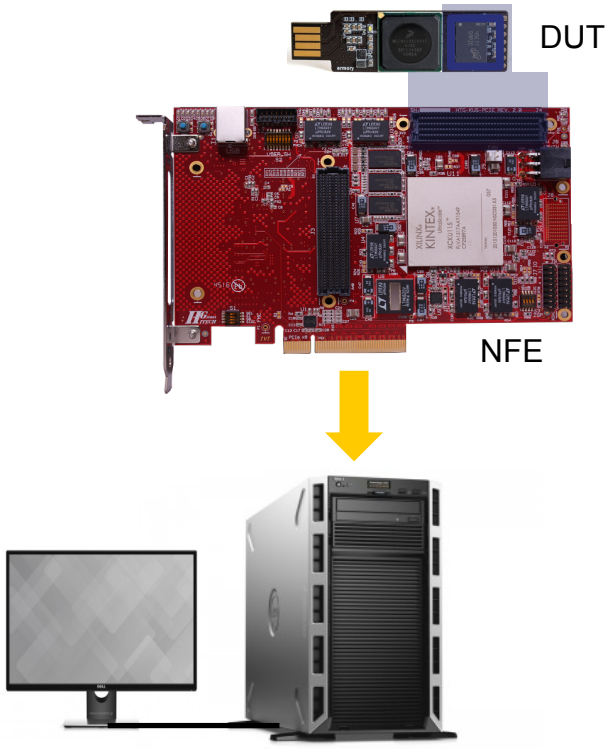

Supermicro X10SRi-F
Xeon E5-2650 v4

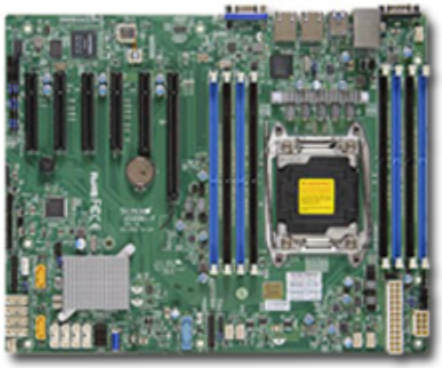

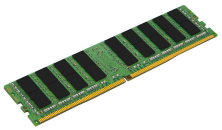

64 or 128 GB
DRAM SO-DIMMs

HTG-K800
Xilinx Kintex UltraScale KU085

Emulates up to

- 2 NAND Flash channels
- 4 - 8 CE per channel
- 2 GB L1 cache
- Specs
  - ONFI 1.0 - 3.0
  - Toggle 1.0 - 2.0

DIMM capacity	Total DRAM	DRAM for NFE	Emulated NAND Flash Memory
64 GB	512 GB	496 GB	480 GB
128 GB	**1 TB**	**768 GB**	**656 GB**

NAND Flash Emulator Software

Flash Memory Summit

UNIVERSITY OF PATRAS
ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΩΝ

DRAM initialization

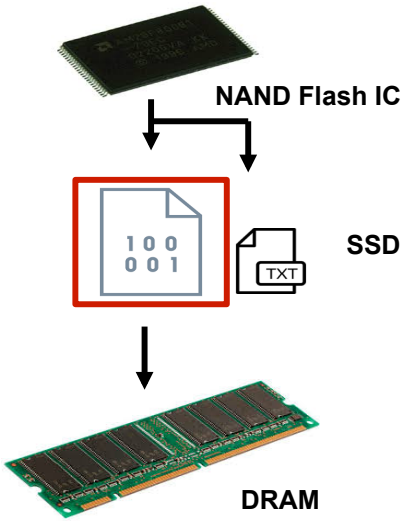

Operation

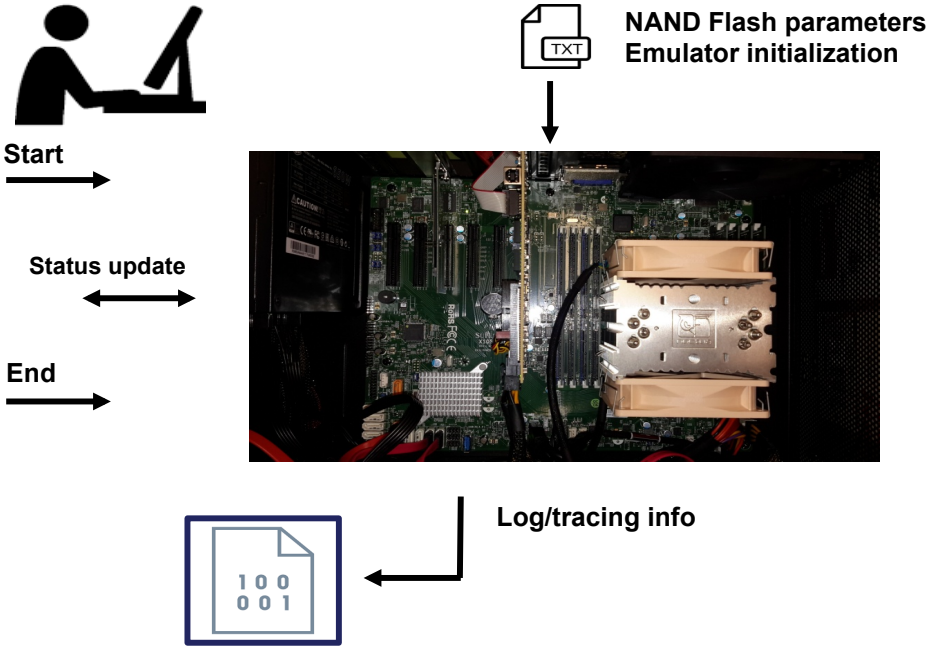

3CPU and Tracer Architecture

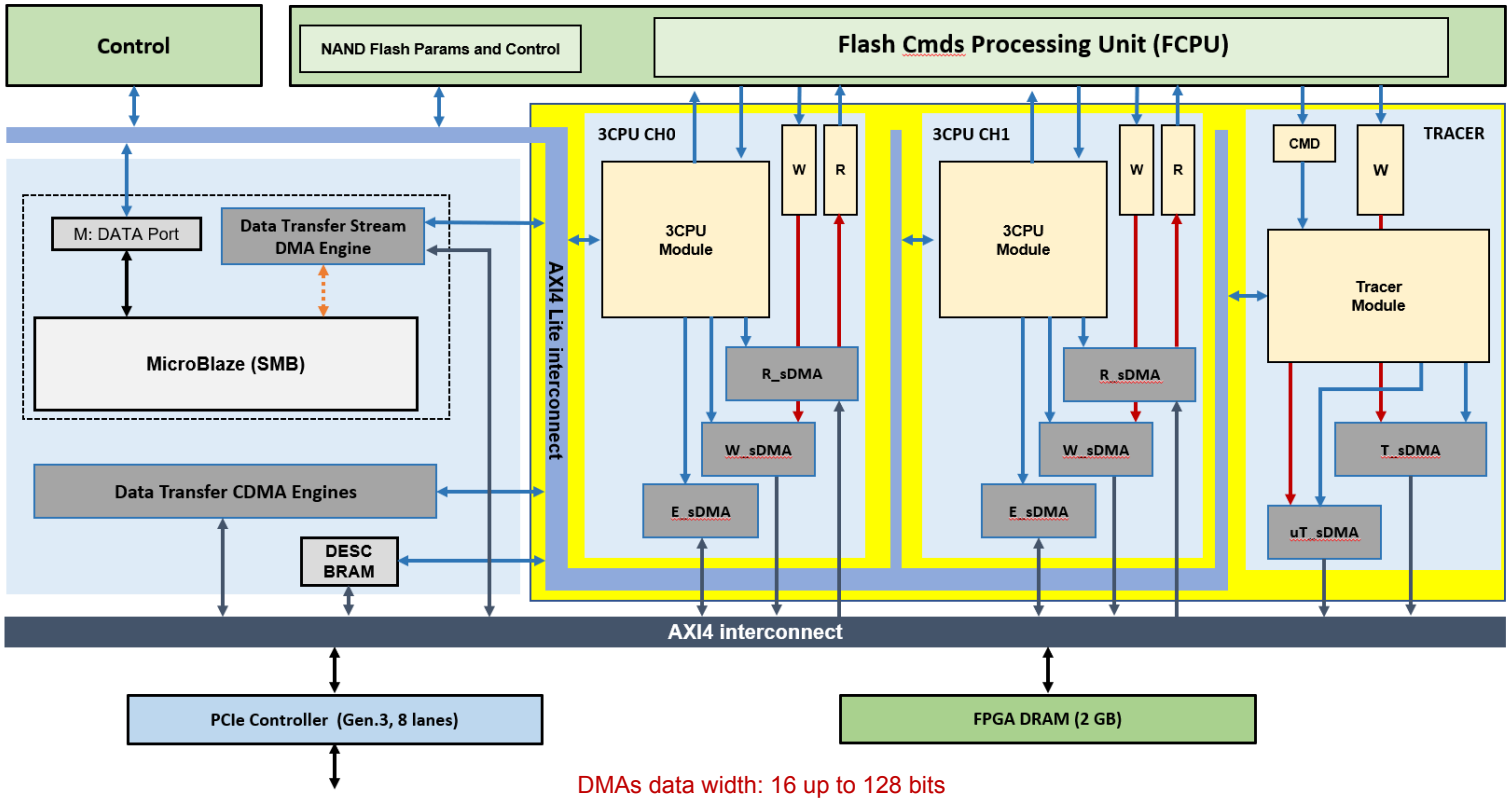

DMAs data width: 16 up to 128 bits

NAND Flash Emulator - Timing

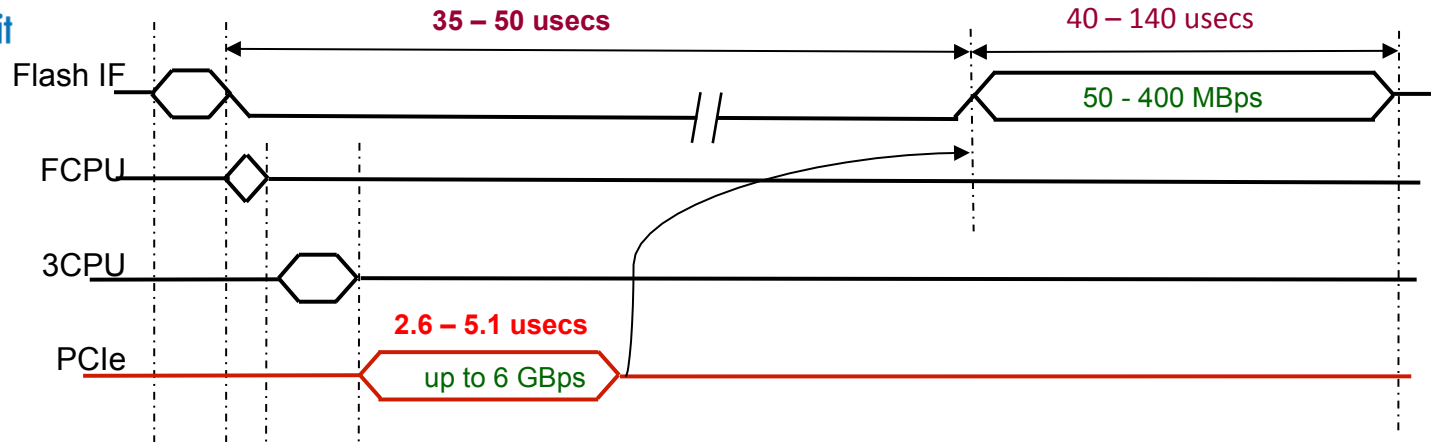

Page size [B]	Read Time [usecs]	Write Time [usecs]	ONFI	Transfer Rate [MBps]	Transfer Time [usecs]	Transfer Time over PCIe (8 lanes Gen3, 128 bits DMA) [usecs]	NAND Flash Channels supported
8,640	**50**	1,300	2.0	166	52.0	**2.6**	**18**
8,640	**35**	300	2.2	200	43.2	**2.6**	**12**
18,592	**50**	1,400	3.0	166	112.0	**5.1**	**8**
18,592	**50**	1,400	3.0	333	55.8	**5.1**	**8**

3CPU Experimental Results

Flash Memory Summit

UNIVERSITY OF PATRAS
ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΩΝ

17,408 B/page
128 pages/block

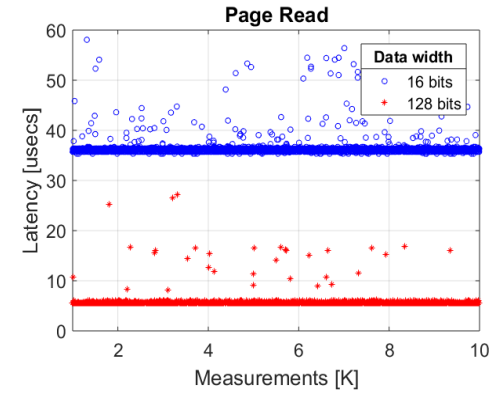

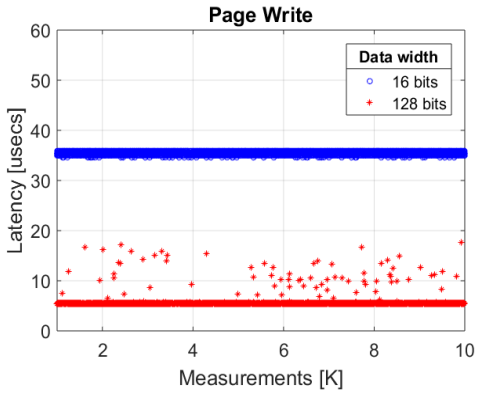

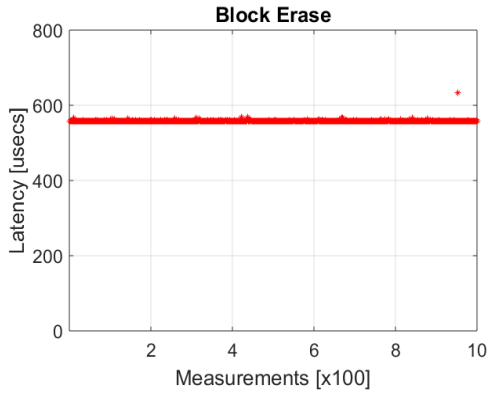

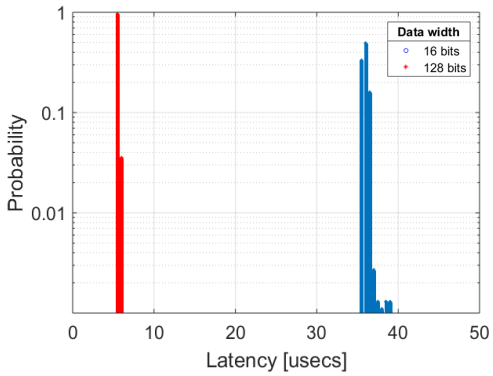

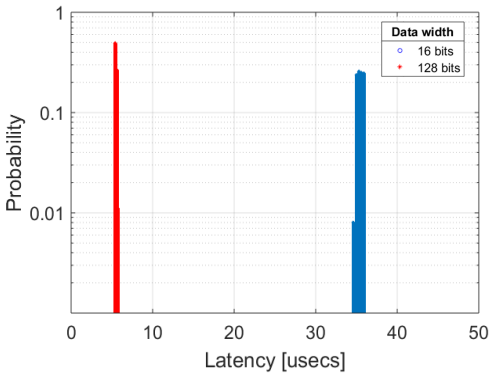

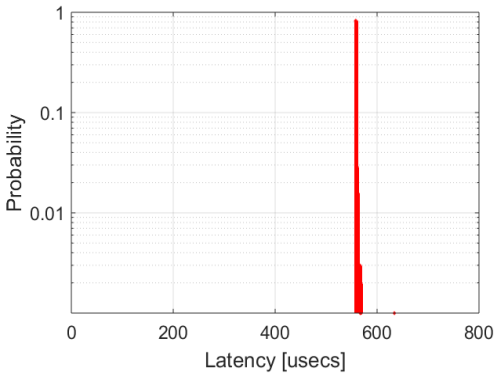

Tracer Performance

16 GB at
Host DRAM

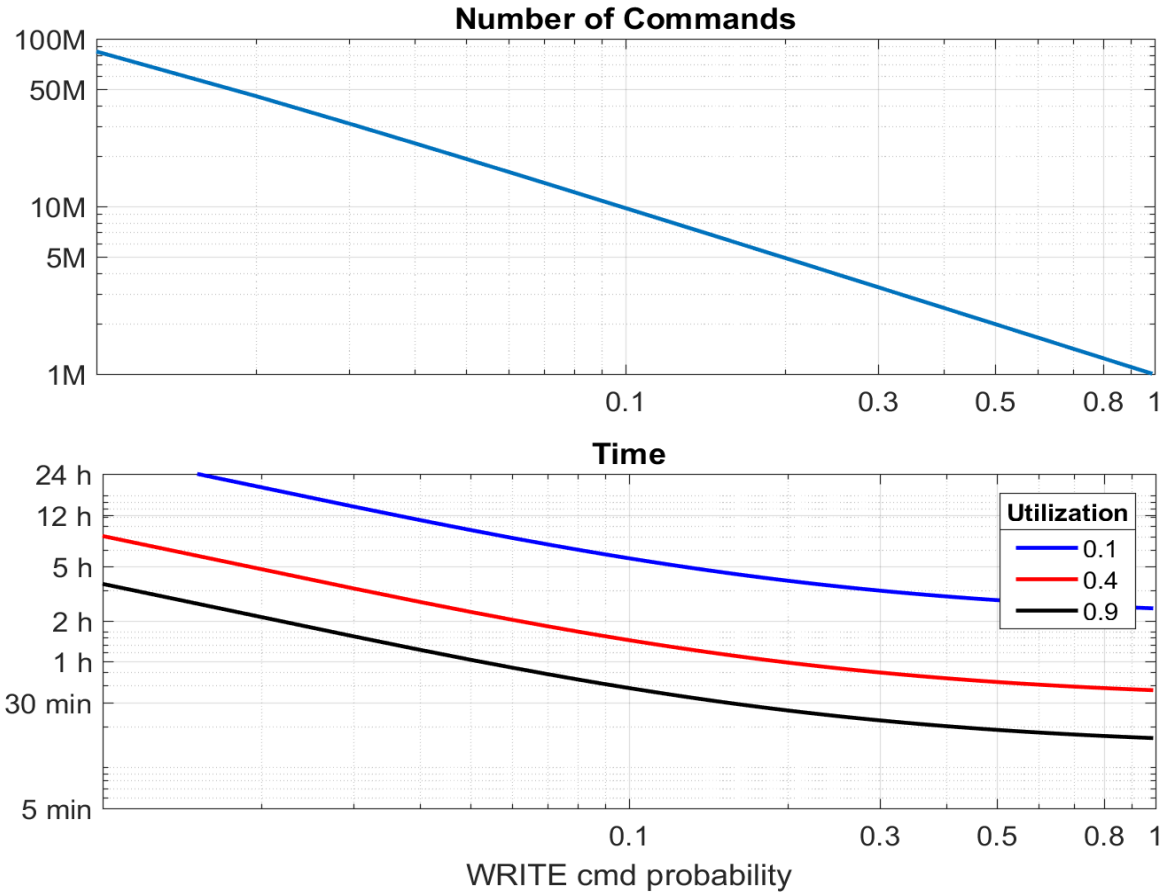

Conclusions

Multiple-channels NAND Flash emulator characteristics:

- Supports a large number of NAND Flash channels
- Practically supports unlimited NAND Flash Capacity
- Responds according to the NAND Flash chips
- Supports ONFI and Toggle interfaces
- Due to its modular design can be re-used for emulating other Non-volatile Memory (NVM) technologies and/or other IO interfaces (i.e. eMMC)
- Design of new algorithms based on data analytics (i.e. minimize read latency by predicting future read/write commands)

Thank you for your attention!

Questions?

Back-up slides

NAND Flash Emulator - System Architecture

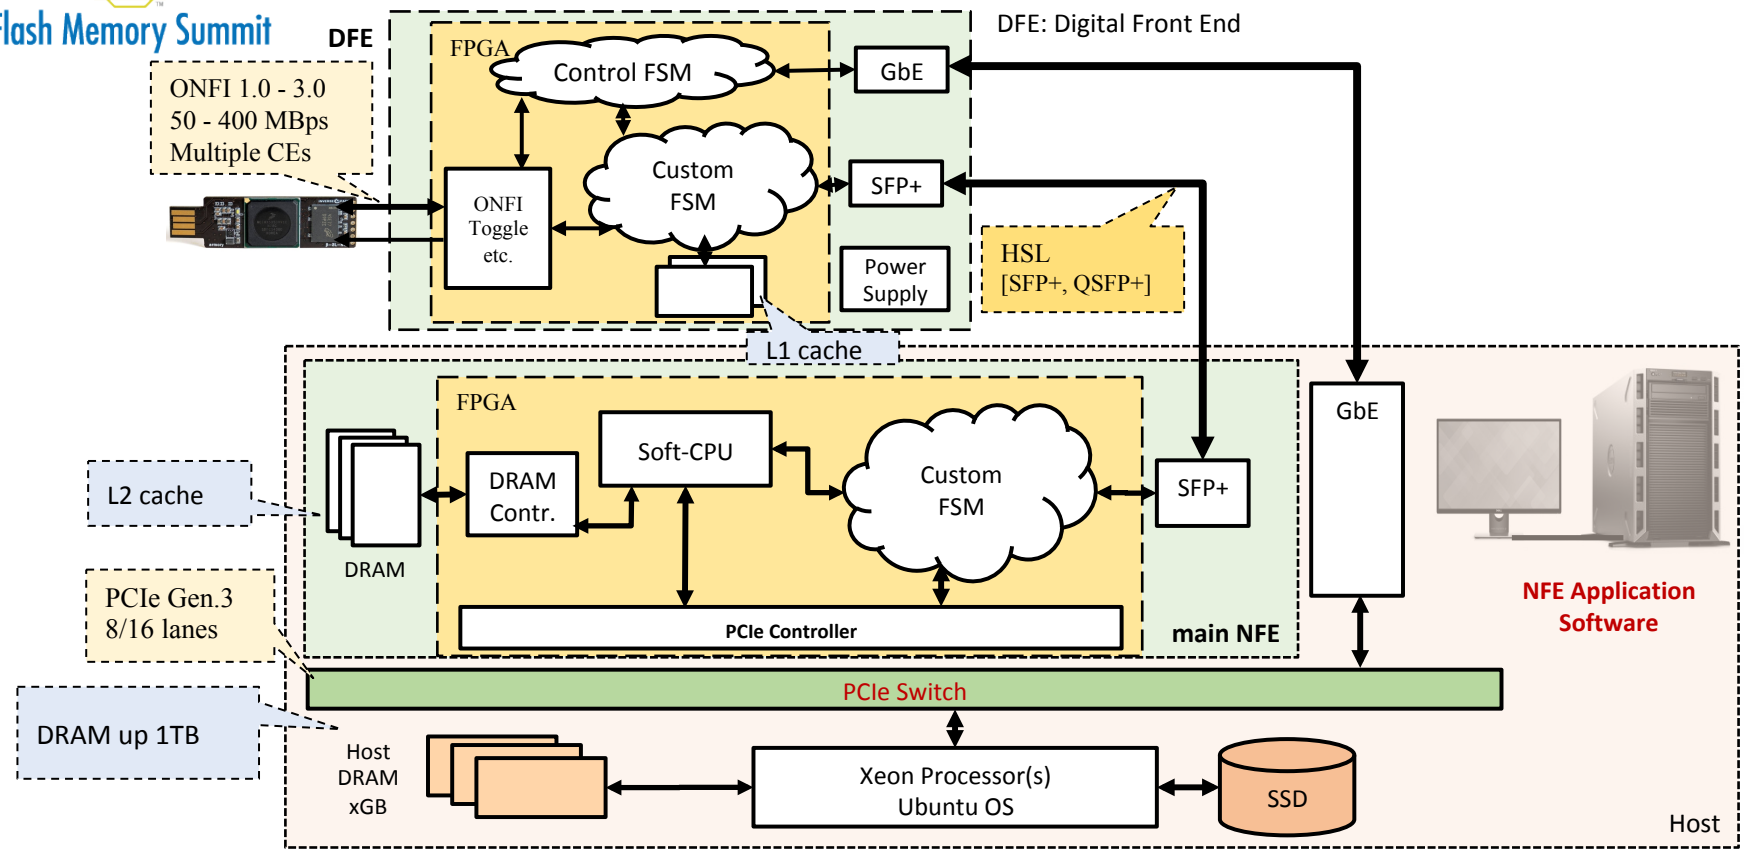

NAND Flash Emulator

Flash Memory Summit

UNIVERSITY OF PATRAS
ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΩΝ

Single board configuration

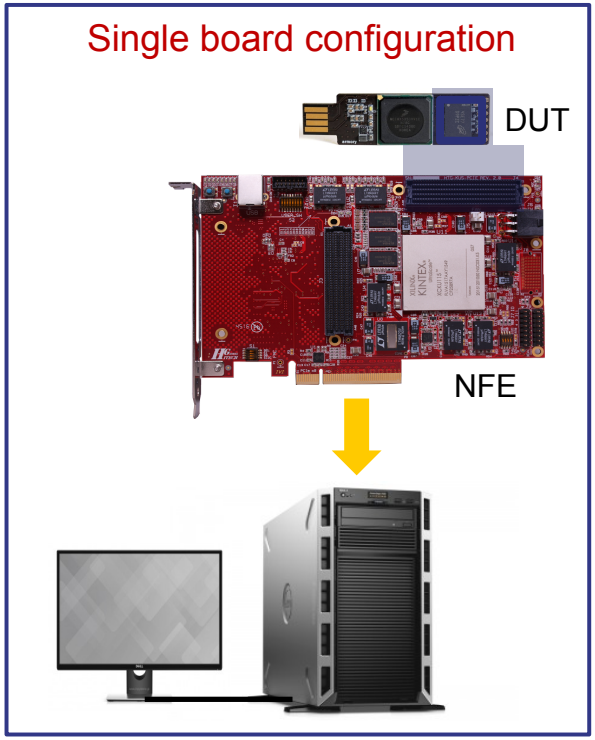

Two boards configuration

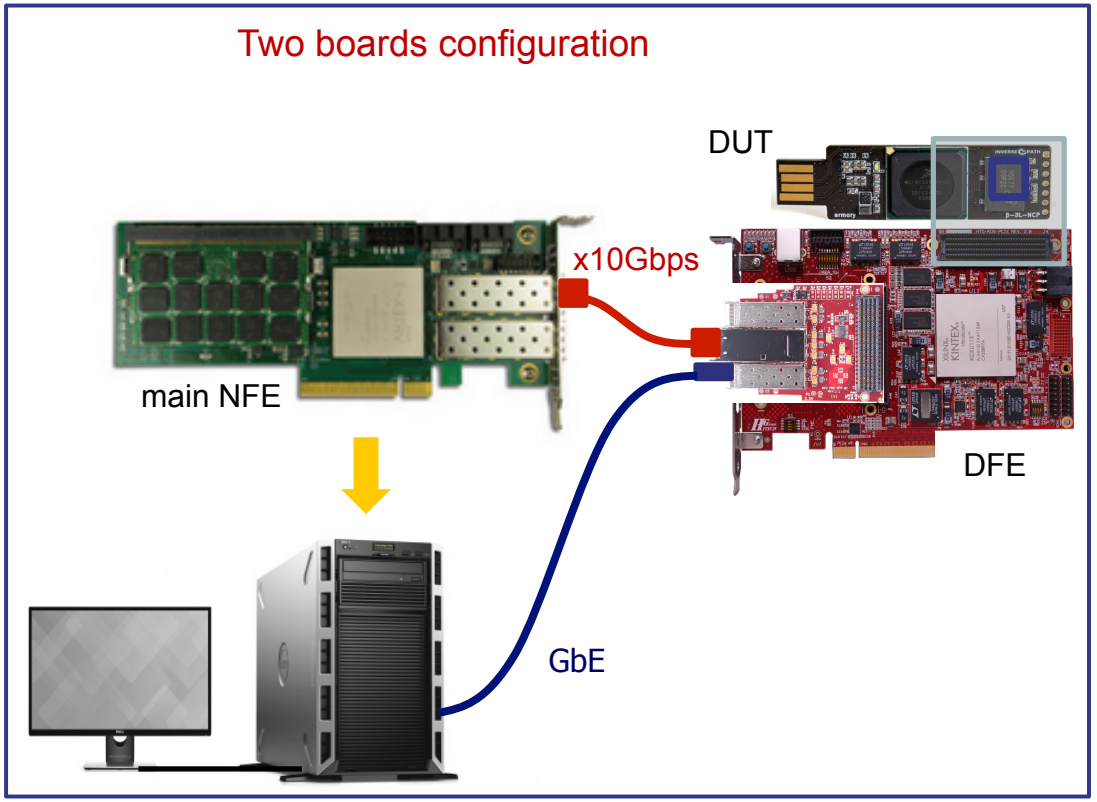

Emulation using a cluster of servers - case #1

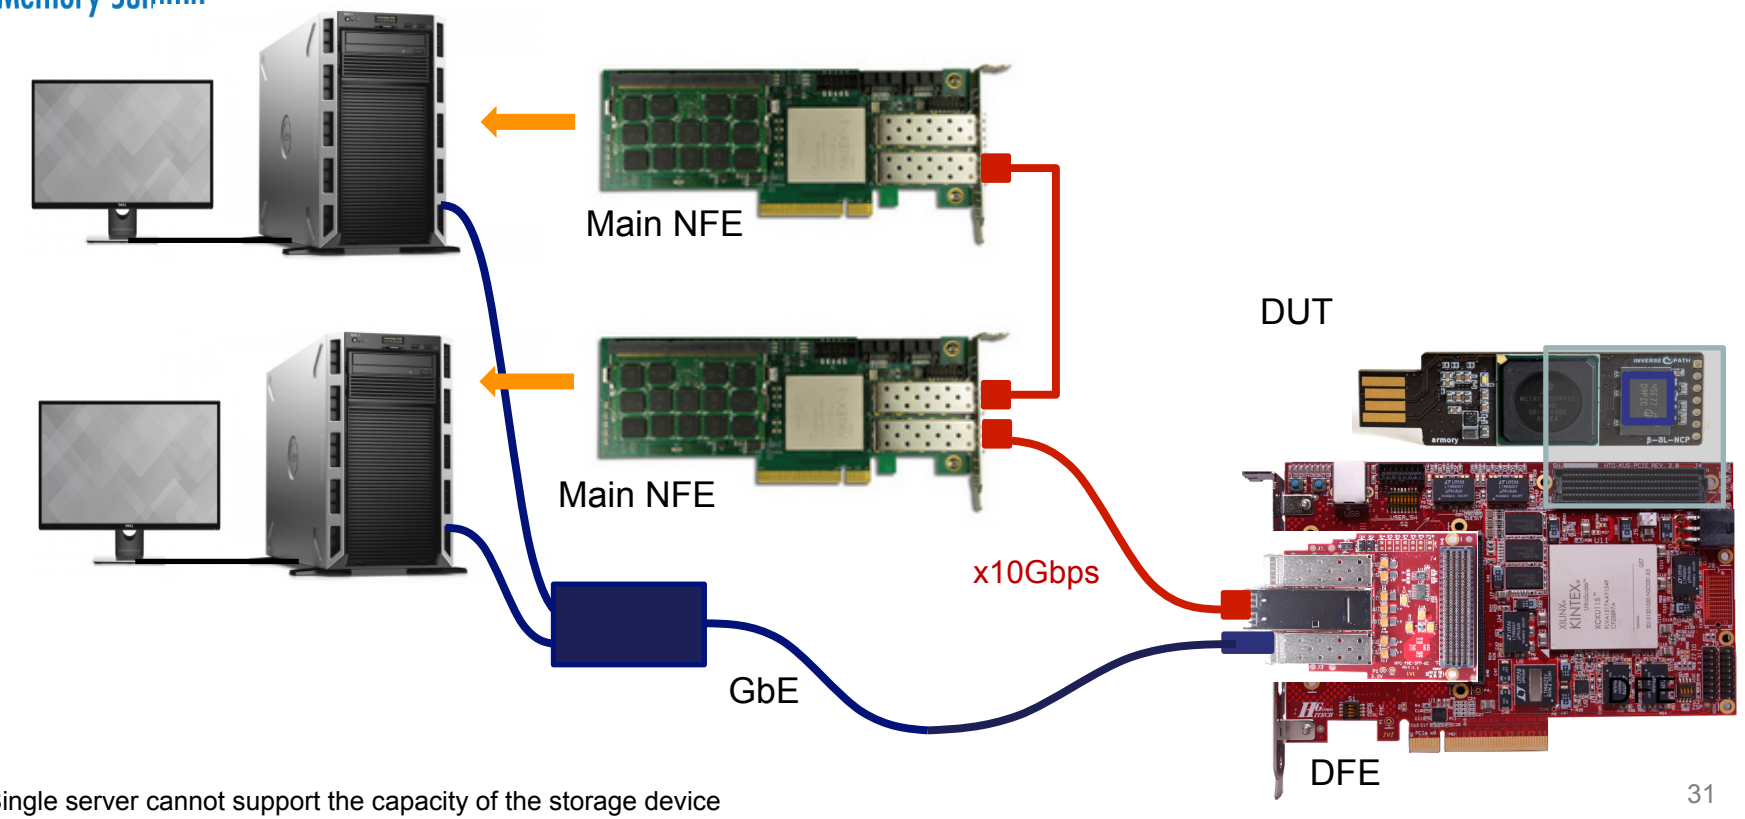

Single server cannot support the capacity of the storage device

Emulation using a cluster of boards

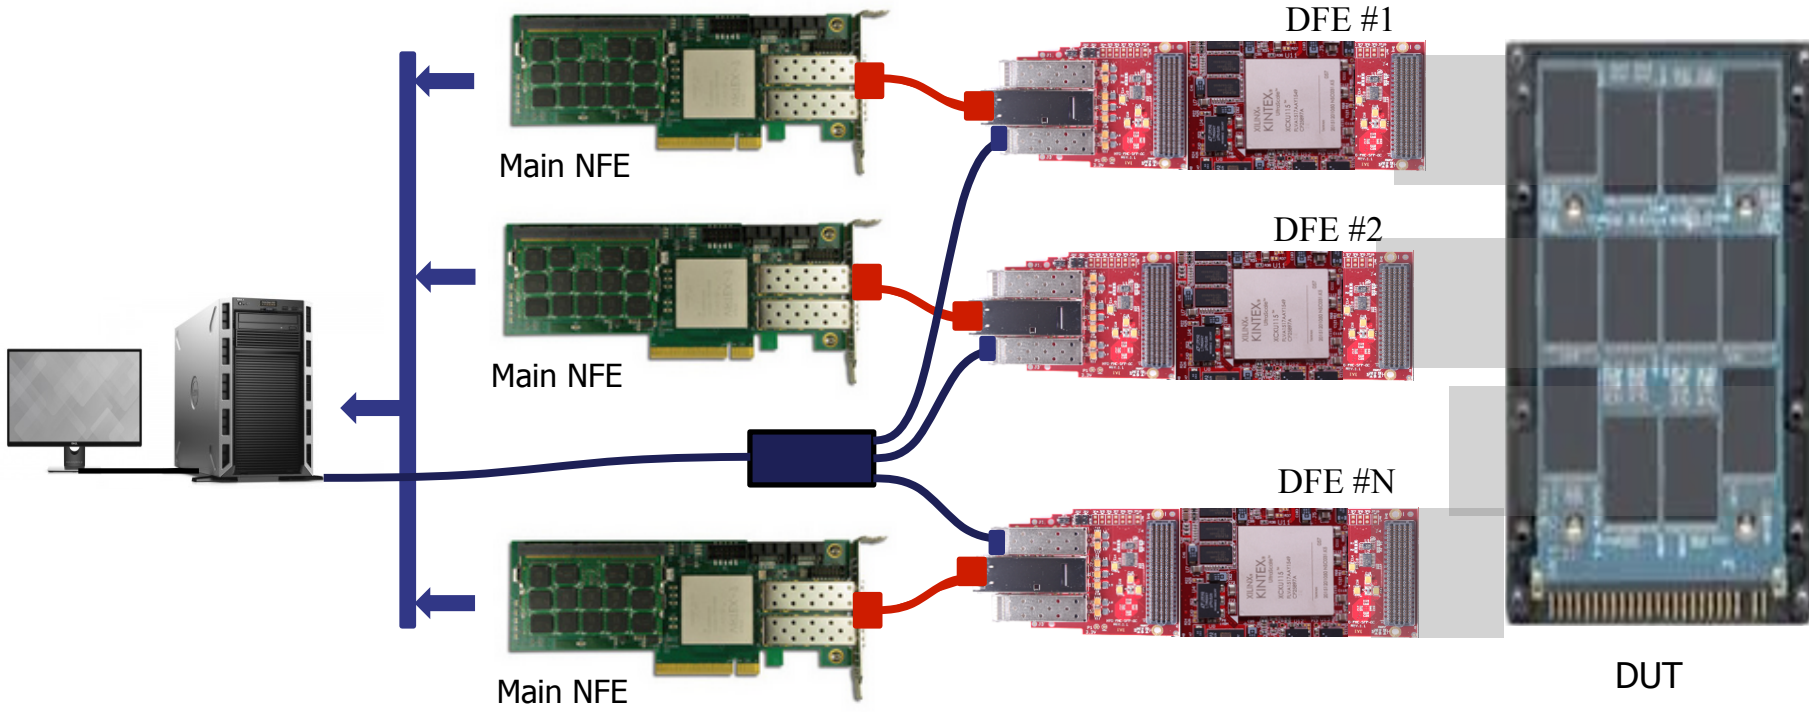

Single DFE cannot support the number of channels of the storage device

Emulation using a cluster of servers - case #2

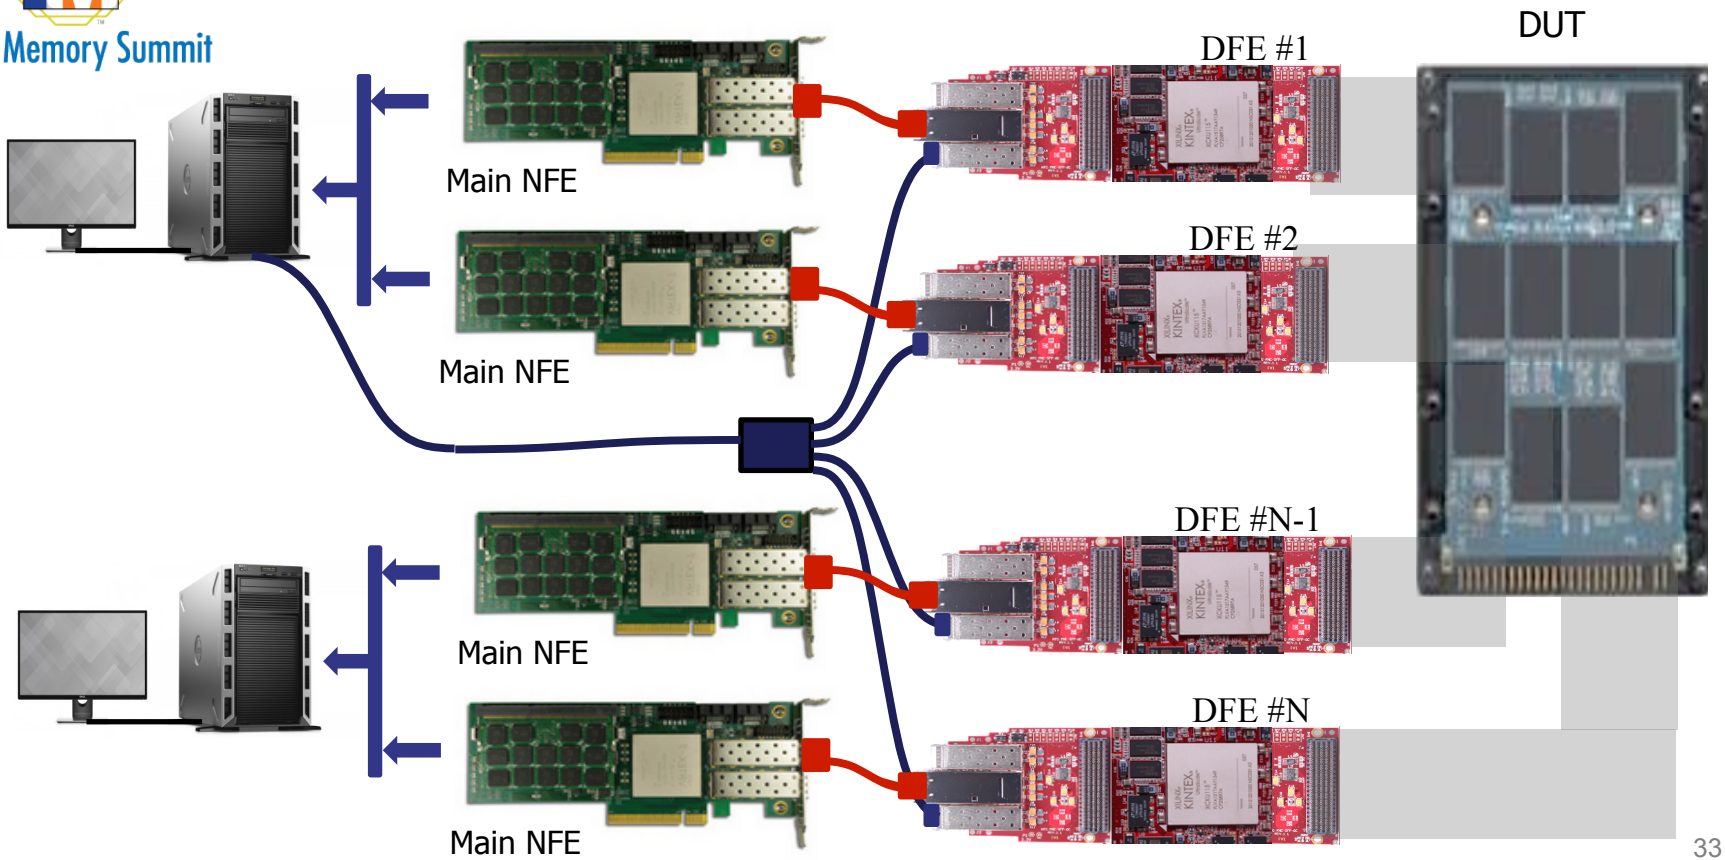

Prototype of the NAND Flash Emulator

Flash Memory Summit

UNIVERSITY OF PATRAS
ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΩΝ

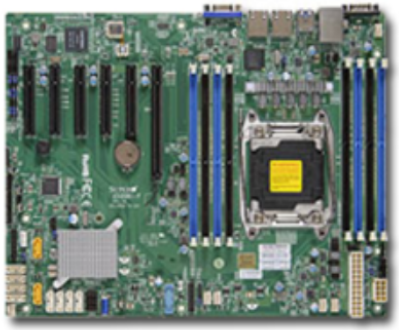

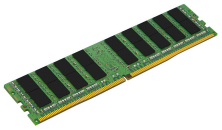

64 or 128 GB DRAM SO-DIMMs

DIMM capacity	Total DRAM	DRAM for NFE	Emulated NAND Flash Memory
16 GB	128 GB	112 GB	96 GB
32 GB	256 GB	250 GB	224 GB
64 GB	512 GB	496 GB	480 GB
128 GB	**1 TB**	**768 GB**	**656 GB**

Xeon E5-2650 v4

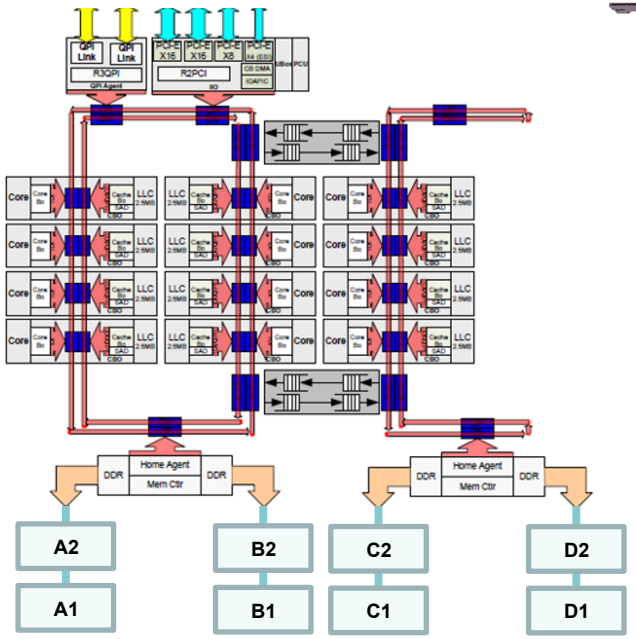

HTG-K800

Xilinx Kintex UltraScale KU085

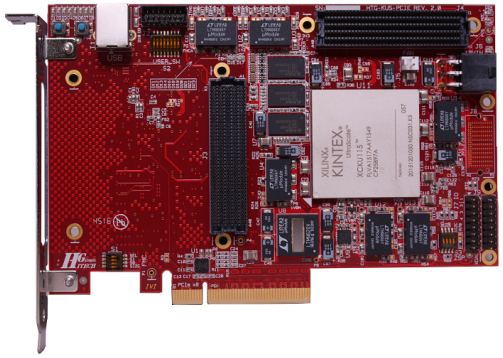

Emulates up to

- 2 NAND Flash channels
- 4 - 8 CE per channel
- 2 GB L1 cache
- Specs
  - ONFI 1.0 - 3.0
  - Toggle 1.0 – 2.0

Tracer Experimental Results

Flash Memory Summit

UNIVERSITY OF PATRAS
ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΩΝ

Page: 17,408 bytes

MBps	Transfer time [usecs]
66	264
200	87
400	44