Flash Memory Summit

Balancing Conflicting Design Constraints of GPUs, CPUs and SSDs

Presenter: Odie Killen, VP of Engineering, Viking Enterprise Solutions

Overview

- Typical System Layout
- Functions and Challenges of Each Device
- Preferred Solutions and Conflicts
- Currently Available Solutions
- Future Paradigm Shift

Typical System Layout

- Typical solutions consist to CPUs, GPUs and SSDs
- Mimic a traditional storage server implementation
  - Physical Size, Power and Cooling requirements of a GOU are challenging
- Future implementations will eliminate CPUs
  - Requires continued advancement on GPU and management SW
  - Various approaches are being evaluated

Typical System Layout

- Typical layout shown with CPU, SSDs and PCIe AICs

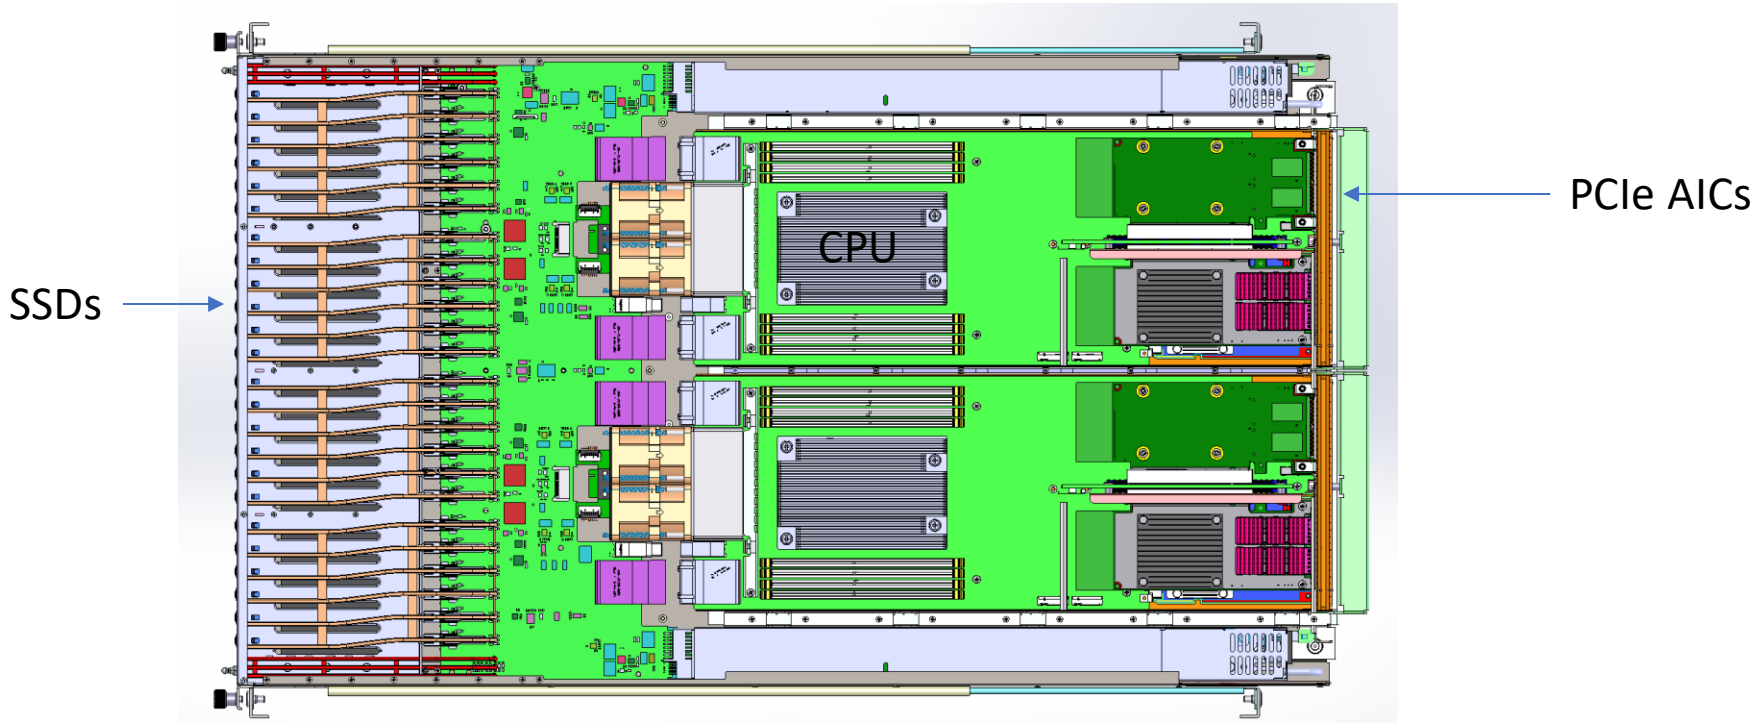

System Layout – GPU

- Controller/Server card shown with typical GPU

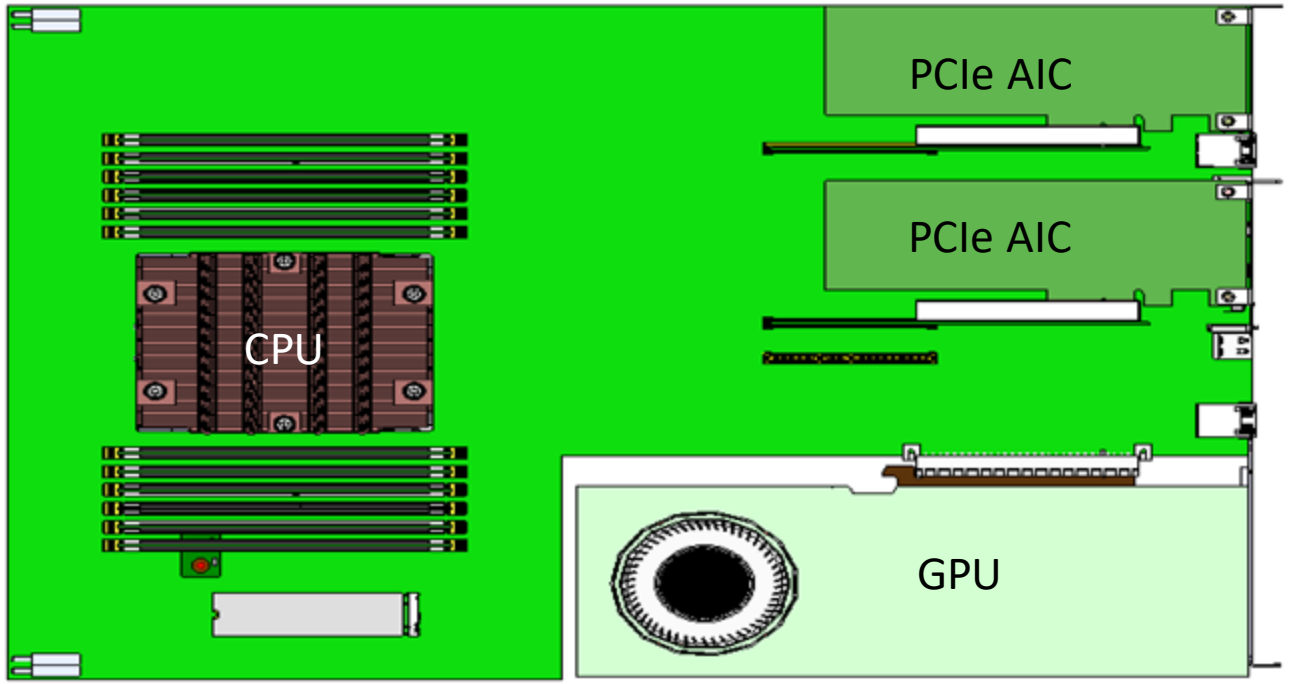

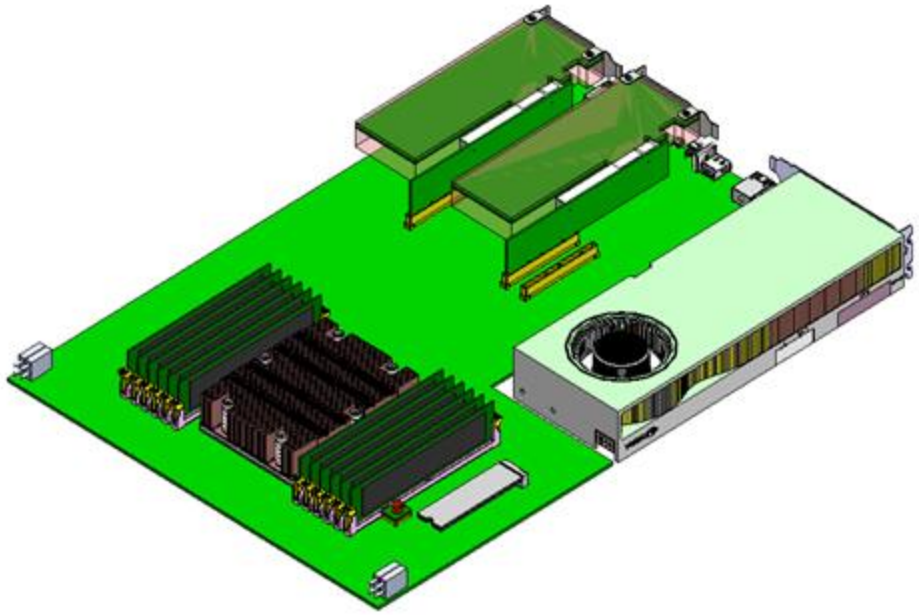

Device Functions

- SSDs
  - Provides very high throughput, low latency data storage using solid state memory
  - Operates at a higher temperature and power level than a traditional spinning disk
  - This is the high performance landing spot for any data stored during transactions
- CPU Complex
  - Provides the computational engine for the computer. It processes input, performs operations on data, generates new data and exports results for storage
  - This is essentially the brain of the device

PCIe AIC Functions

- GPU/DPU
  - Provides the capability to process complex tasks on data without burdening the CPU engine.
  - Performs at a data rate much closer to line speed than is possible with a CPU
  - Extremely high performance and much high power/cooling requirements than other PCIe based devices (150-300W per device)
- General Purpose PCIe AICs
  - Provides an input path to the CPU.
  - Typically capable of performing some level of data manipulation and translation to the PCIe protocol
  - Typically much lower power than DPU/GPU with easier to achieve cooling requirements

Preferred Solution and Conflicts

- Preferred; more capabilities in a smaller space
  - More than 1 DPU/CPU per MOBO
  - 8x high-perf SSDs per GPU/DPU
  - Support for Full Size GPU/DPU
  - 2U, 1m deep packaging target
  - Prefer traditional, low cost cooling methods
  - Prefer the highest density solutions possible

- Conflict; existing technologies don't support desired capabilities
  - Power and cooling requirements exceed typical capabilities
  - Complicates cooling solution
  - Power. Packaging and cooling is much more complicated
  - Traditional air-cooled solutions are simply not adequate for more complex and dense solutions
  - Need to re-address the approach

Current Solutions

- More and more solutions are becoming available
  - Typical solutions are based on existing products and CPU based
  - 2U solutions with 1 GPU/DPU per controller, High Availability and 24 or less SSDs
  - 3U solutions with 2-3 GPU/DPU and reduced SSD counts, non-HA
  - Traditional cooling methods using existing rack infra-structure
  - Existing solutions don’t meet the need of DPU/GPU and SSD quantity
- CPU less solutions are starting to be evaluated

Future Paradigm Shifts

- Alternative cooling methods to include cold plates, liquid cooling and full immersion cooling
  - All require system to be designed with this in mind from the beginning, no more “shoe-horning” into existing platforms
- Purpose built solutions around GPU/DPU requirements
- Further advancement of dis-aggregated storage and compute
- Further drives a demarcation between traditional compute and GPU/DPU based solutions
- Increased instances of CPU-less designs

Thank You