Computational Storage

Jamon Bowen

Product Planning & Storage Segment Director

XILINX

Bottlenecks Remain for Data Intensive Applications

Processor-centric architecture

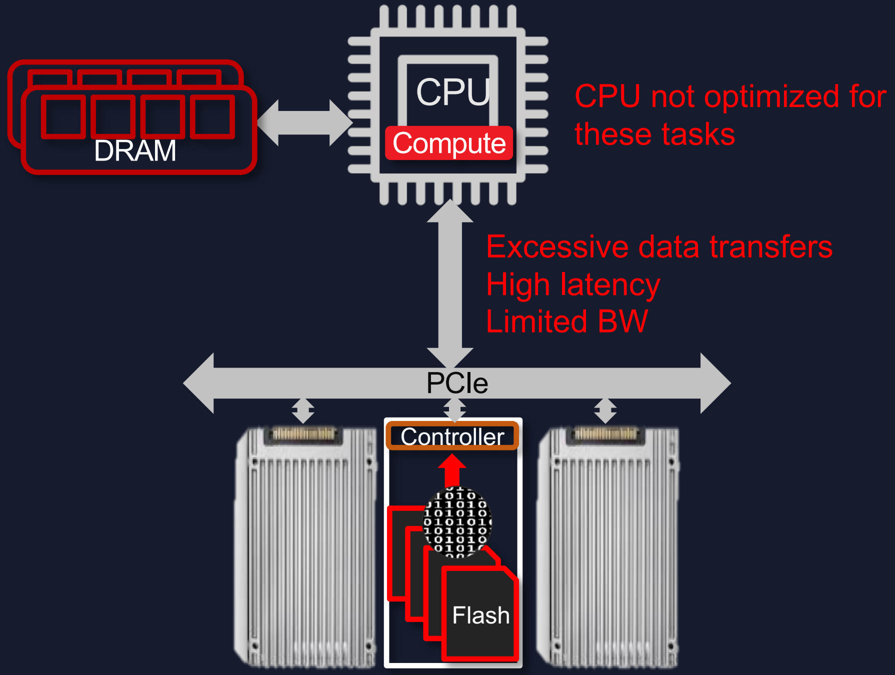

Emergence of Computational Storage as the Solution

Computational storage architecture

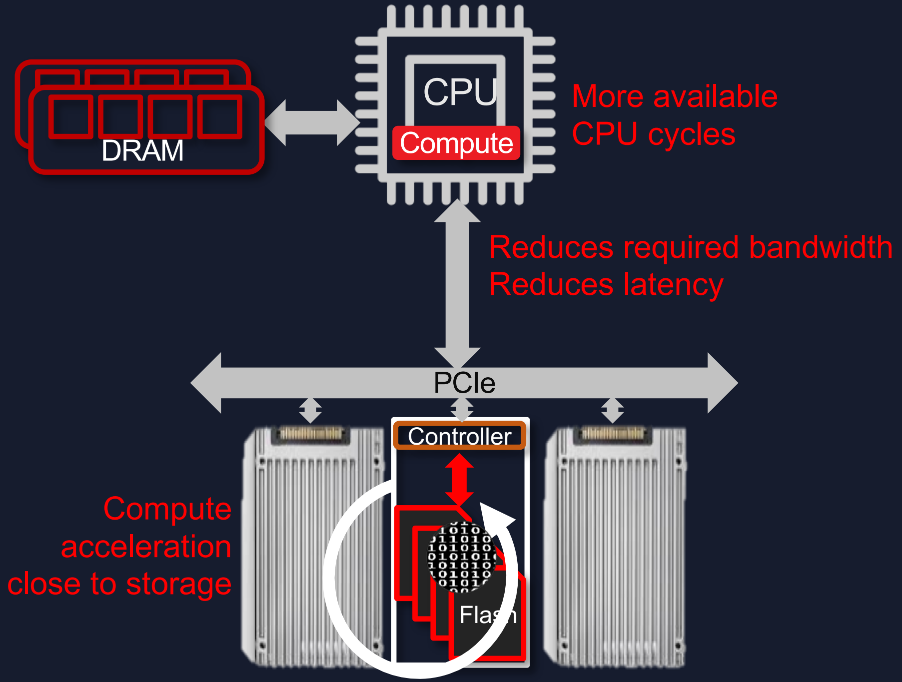

How FPGAs Address the Computational Storage Problem

FPGAs in Storage Today

- Flash controllers

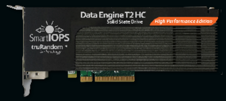

- Storage Systems
  - Cache-offload
  - Storage System & Switching connectivity
  - Data Reduction


FPGA Advantages for Computational Storage

- Flexible, fully customizable architecture adapts to specific applications
  - Massive parallelism, I/O and customizable data path
- Performance, power and latency of dedicated HW + reconfigurability of SW
- More economical than ASIC/ASSP for many applications

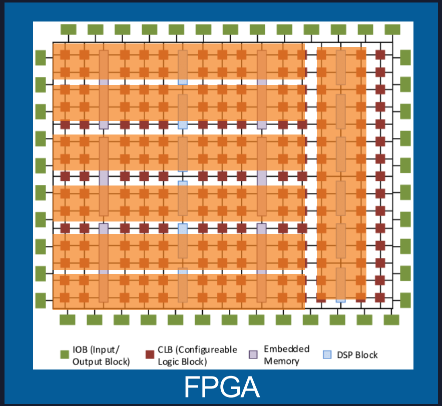

Encryption Accelerator

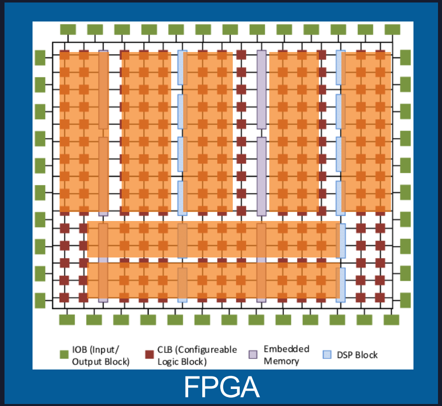

Decryption Accelerator

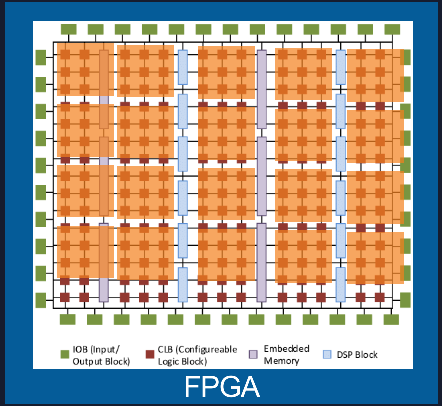

Analytics Accelerator

XILINX

FPGA Advantages for Changing Standards

Architecture easily adapts to latest compression algorithms

Gzip Accelerator

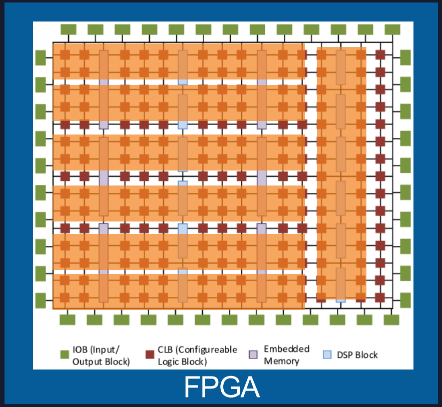

Brotli Accelerator

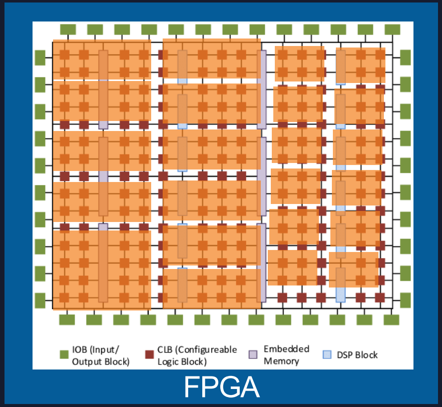

Zipline Accelerator

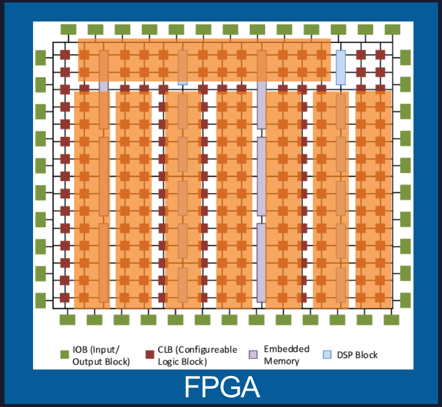

XILINX

Computational Storage Deployment Options

Computational Storage Drive (CSD)

› Integrated Accelerator and Flash

› Benefits:

  » Easy to implement- plug & play
  » Adding capacity adds accelerators + performance
  » Ability to optimize BW between accelerator and flash
  » Ability to customize FTL for specific workloads

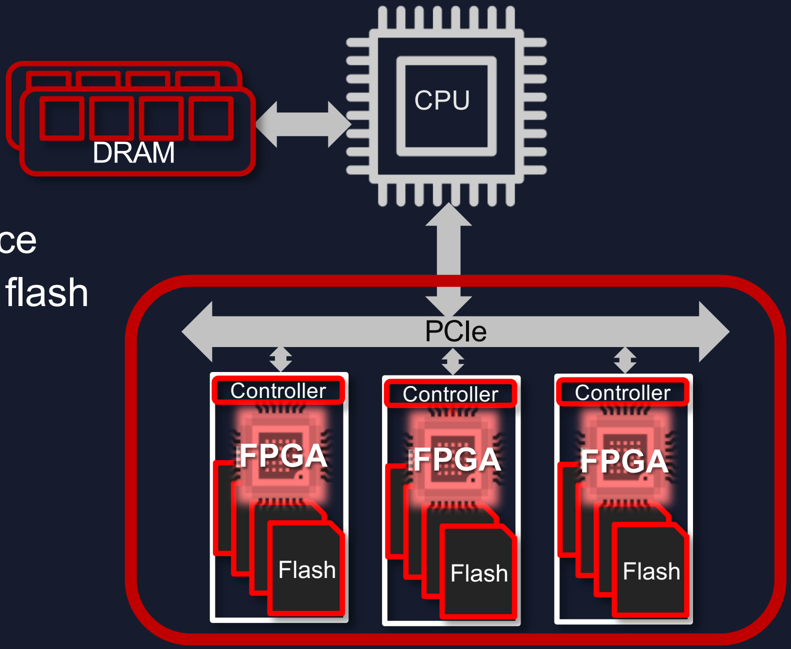

› Xilinx Partners:

  » Samsung
  » Scaleflux

Computational Storage Processor (CSP)

- Accelerator and Storage on same PCIe subsystem
- Benefits:
  - Independent SSD & acceleration scaling
  - Plugs into standard slot
  - PCIe Peer-to-peer transfers for high bandwidth and low latency
- Xilinx Partners:
  - Bittware
  - Eideticom
  - Xilinx

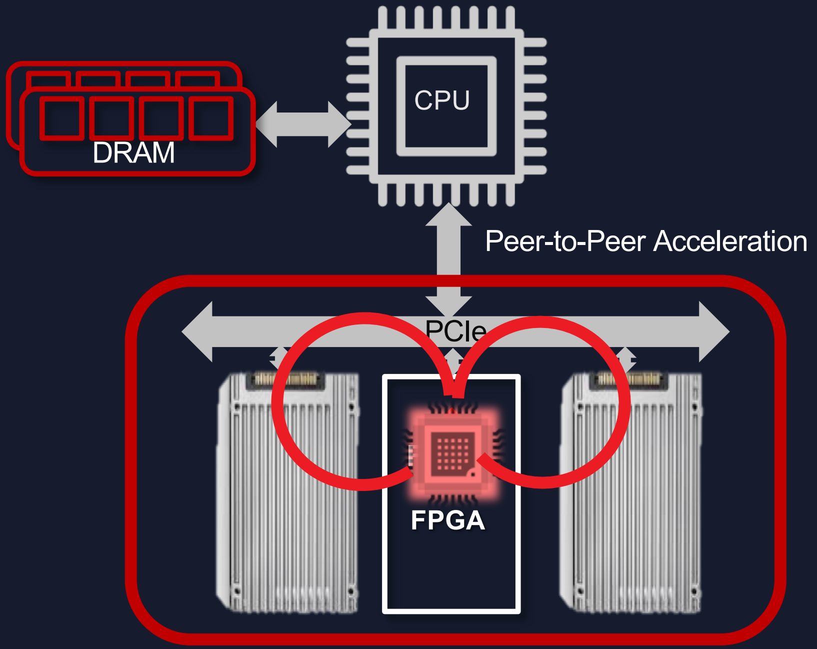

Computational Storage Array (CSA)

- **Accelerator in-line with storage**
- **Benefits:**
  - SSD vendor independence
  - Independently scale accelerators and SSDs
  - Ability to optimize BW between accelerator and SSDs
- **Xilinx partners:**
  - Bittware

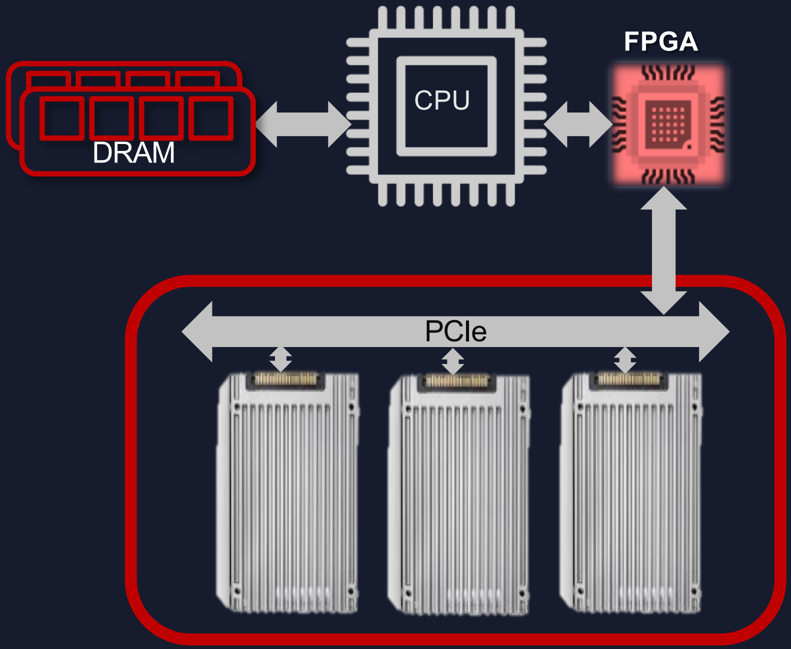

Computational Storage Example Application

Example of Analytics Acceleration

Q1: "Which cities originate the most flights with >10min delays?
Q2: "Which airport in the Bay Area has the worst record?

Airline traffic in the USA from 1970 to Present

Flight Data — 1.2B Entries
Airport Data — 500M Entries
Planes Data — 700M Entries

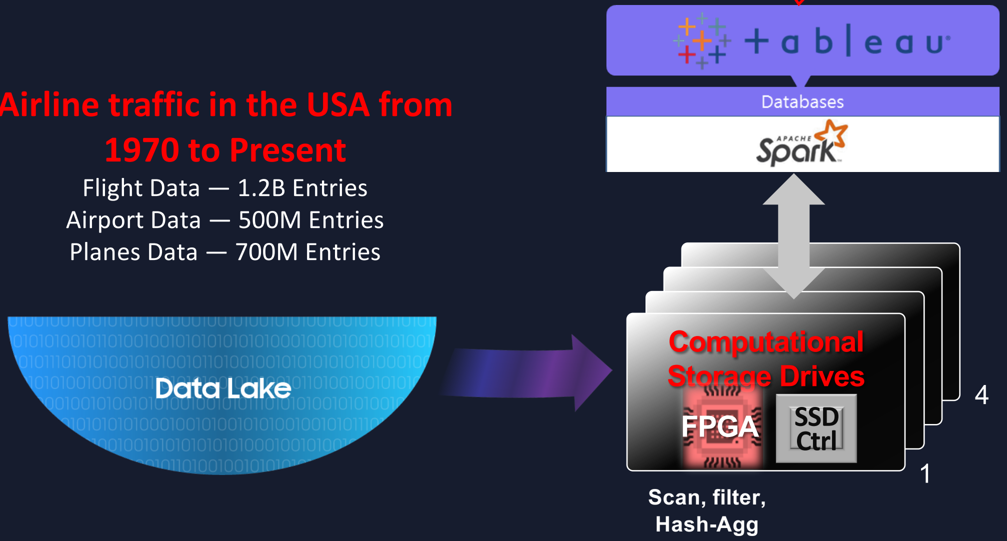

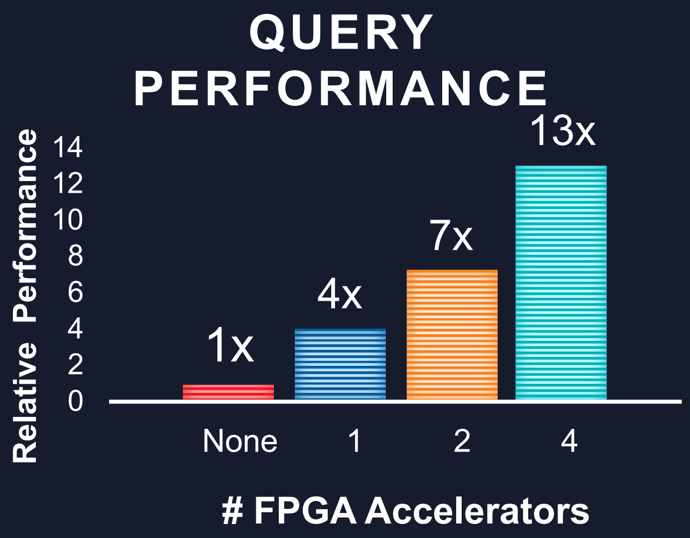

XILINX

Storage Developer Flow

Runtime Stack

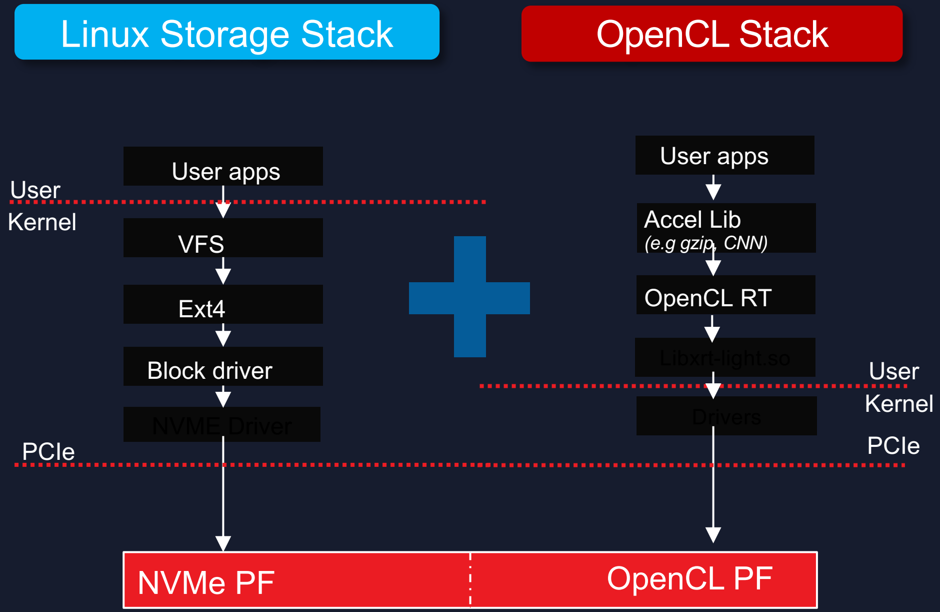

- Storage Accessed via NVMe Stack
- Computational Storage / Accelerator Discovered, Managed, Orchestrated by OpenCL Stack
- Shared Memory Space in the Compute Function Glues the Datapaths together

Database Acceleration Example

```
SELECT   sum(l_extendedprice * (1 - l_discount))
FROM     orders, lineitem
WHERE    l_orderkey = o_orderkey
         and o_orderdate >= date '1994-01-01' and o_orderdate < date '1995-01-01';
```

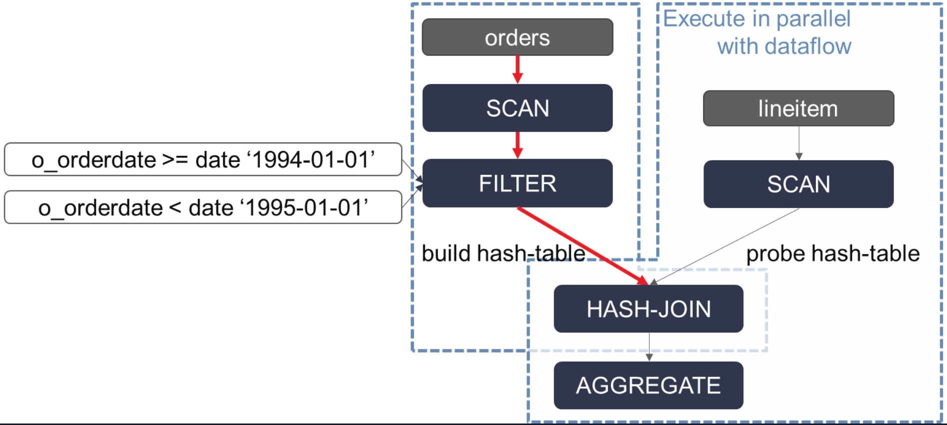

- **Problem: Need to parse through large amount of data to find the records of interest.**
- **Example:**
  - Analytics – Need the records for a time range for just one of many products included in the database.
- **Solution: Push down Scan, Filter, Aggregate to storage.**
- **Why?**
  - Higher Throughput
  - Lower Latency
  - CPU offload
  - Lower TCO

Scan Query on standard NVMe

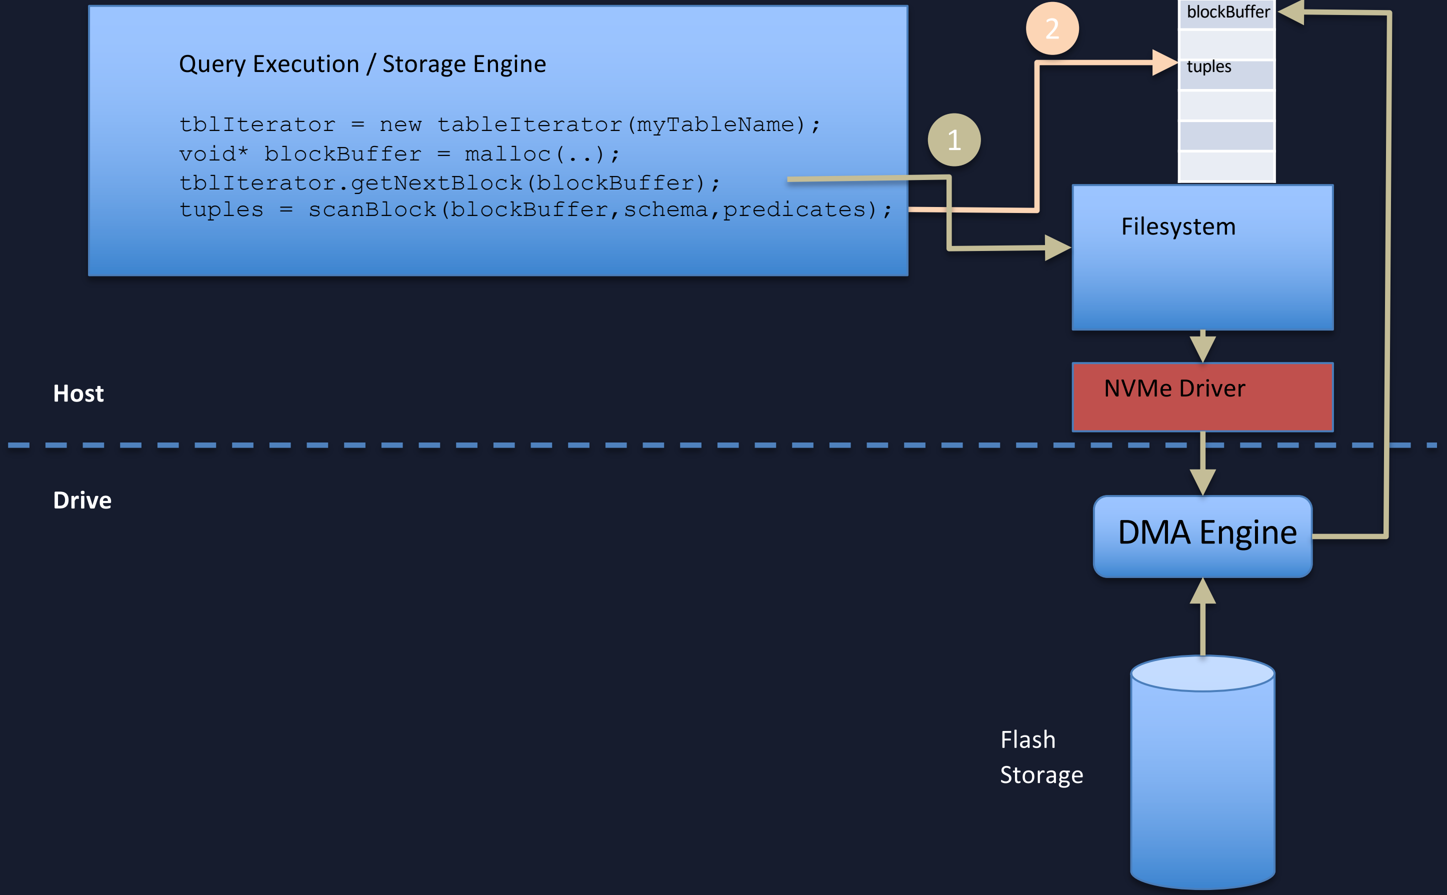

XILINX

Scan Query on computational NVMe

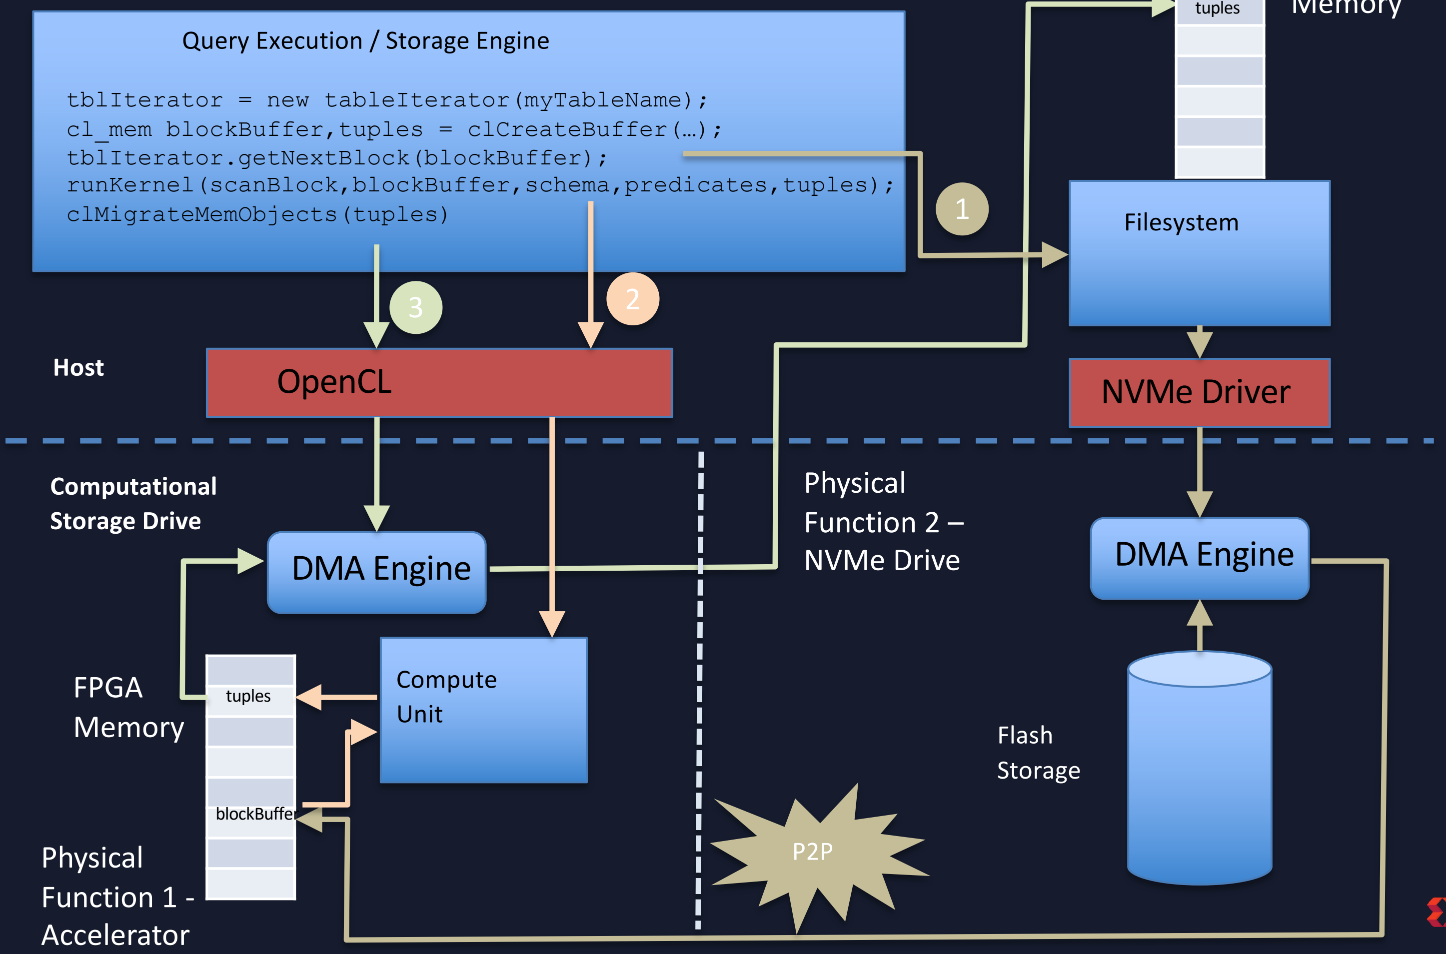

OpenCL CSD example

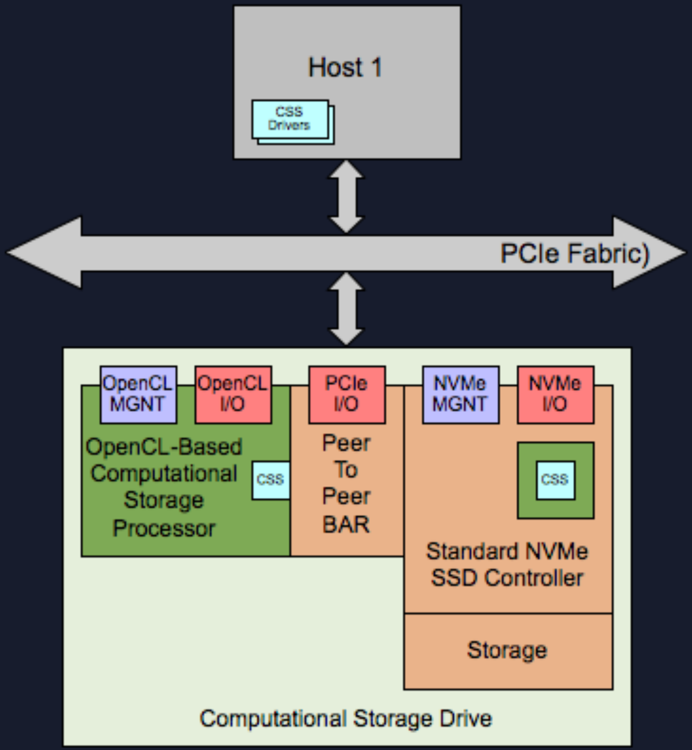

SNIA

XILINX

Why does P2P improve performance?

- Interrupts
- Kernel / User mode transitions
- Copy time
- 1 us = 1000 (1 GHz) – 3000 (3 GHz cpu clock cycles

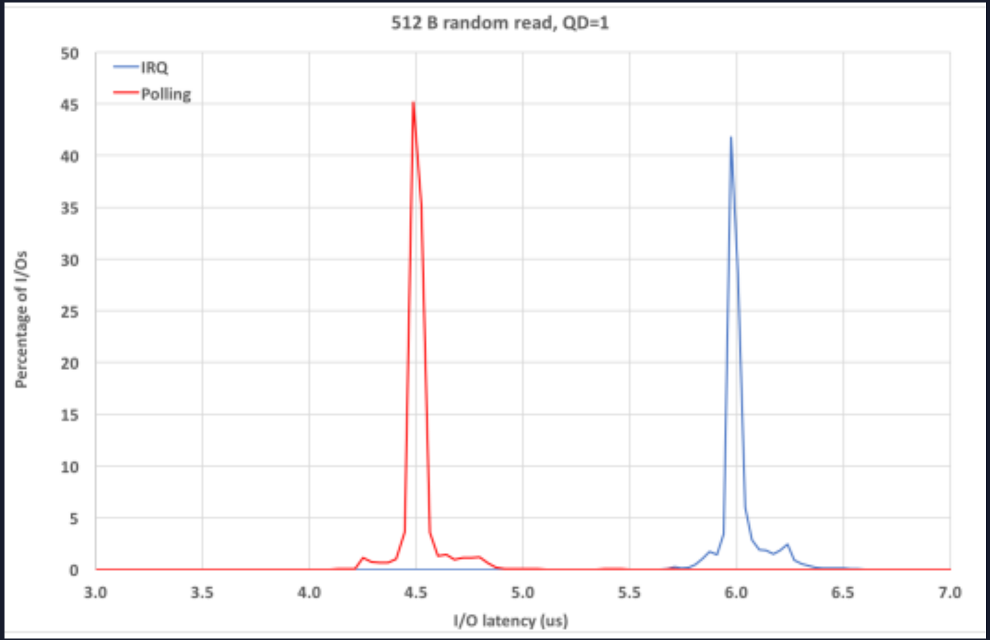

I/O Latency Optimization
with Polling, Damien *Le Moal*
Vault – Linux Storage and Filesystems Conference – 2017

Examples: SmartSSD® CSD and Ecosystem

Storage Services: Comp/Decomp, Encrypt/Decrypt, Metadata management, Erasure Coding,
Real-time Analytics & Biz Intelligence: DB Query (Spark, PostgreSQL, ..), Log analytics, genomics, physics
Rich Media and ML: transcoding, live streaming, object detection

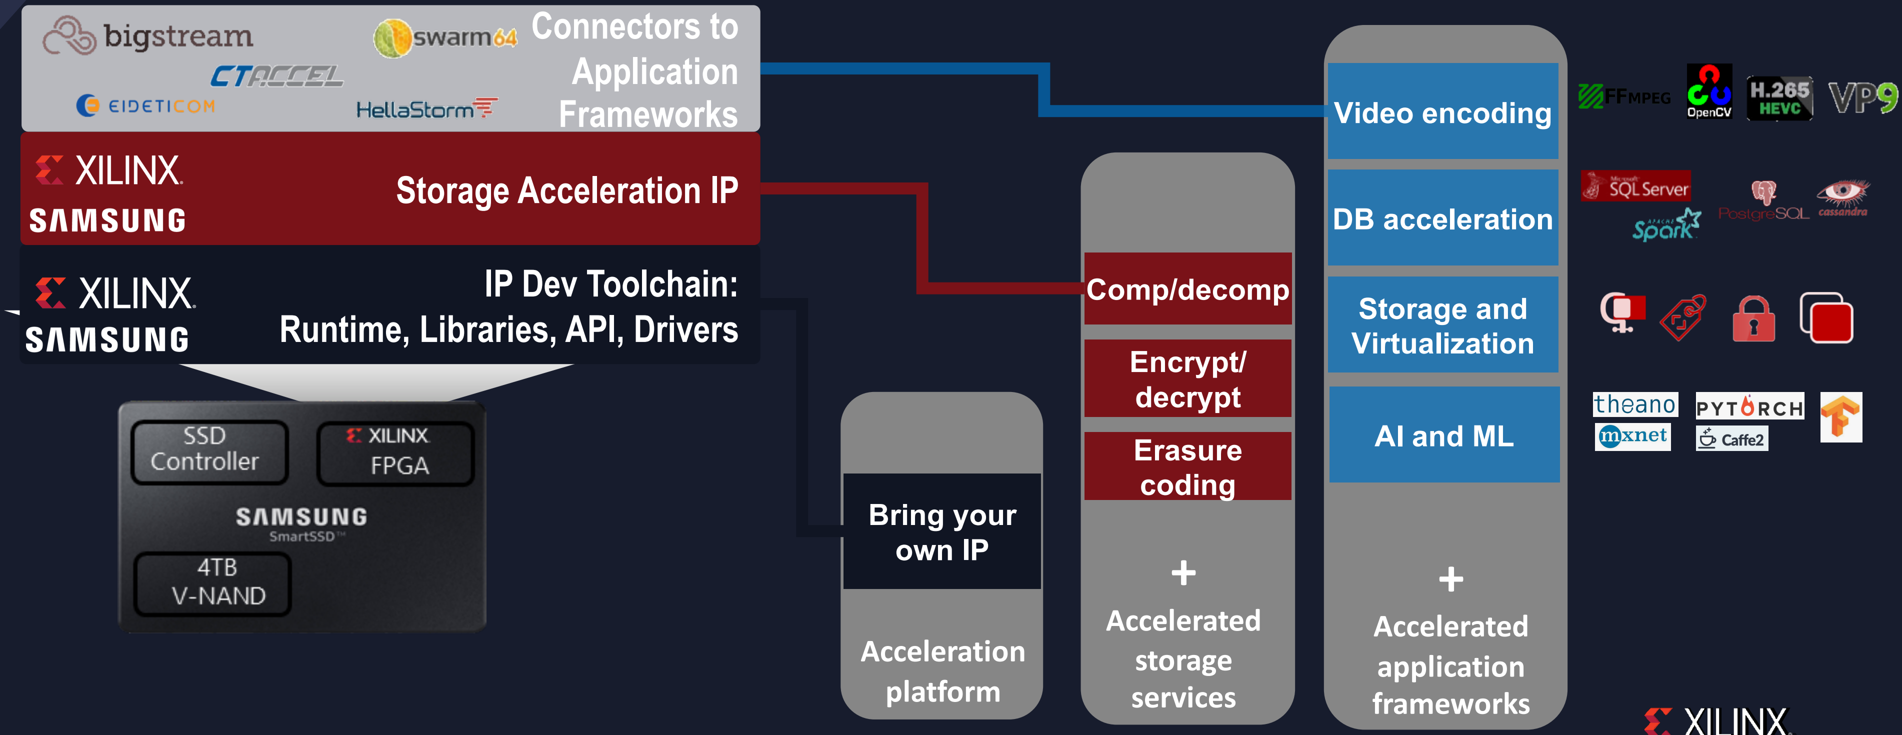

XILINX

Technology Preview SmartSSD® IP Development

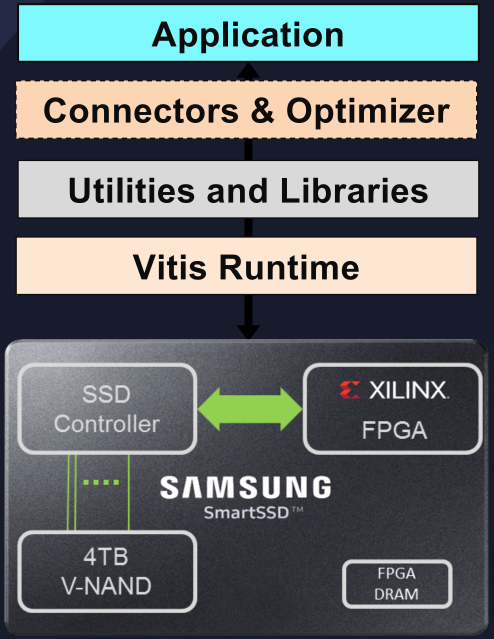

- Deploy off-the-shelf IP and solutions from our partners
- Use familiar Xilinx tools to develop new IP or redeploy existing accelerator IP from ASICs or FPGAs
- Use custom IP development services from Samsung and Xilinx partners

- Enterprise Class SSD Controller: NVMe1.3, CMB, AES256
- 4TB Capacity
- 330K LUTs total in dynamic region available for acceleration IP
- 4GB FPGA DDR
- External interface: PCIe Gen3x4

XILINX

Future Directions

Current Data Center Architecture:
Fixed Resources, Sub-optimal Utilization

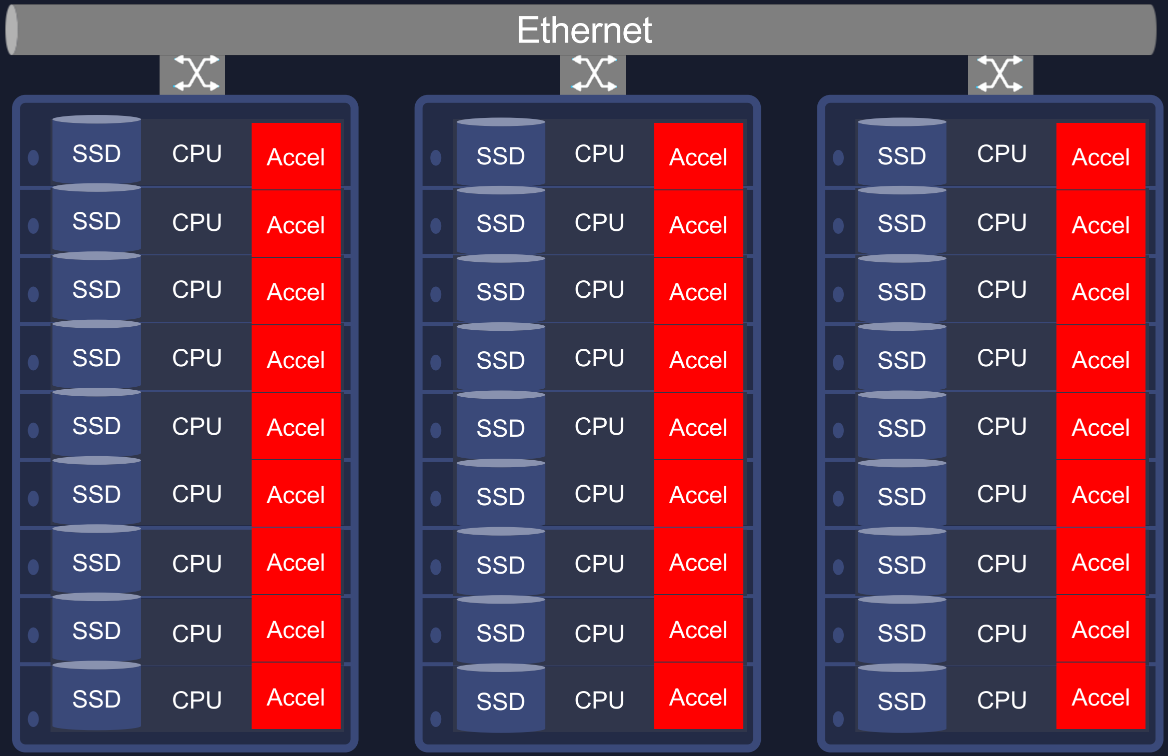

Future Data Center : Disaggregated and Composable

Challenge: Reduced Bandwidth and Increased Latency

XILINX

Computational Storage & Fabric

- Enables composability without significant performance penalty
- Benefits
  - Performance and latency benefits of computational storage
  - Scale compute / storage independently
  - Higher density per rack
  - Lowest TCO

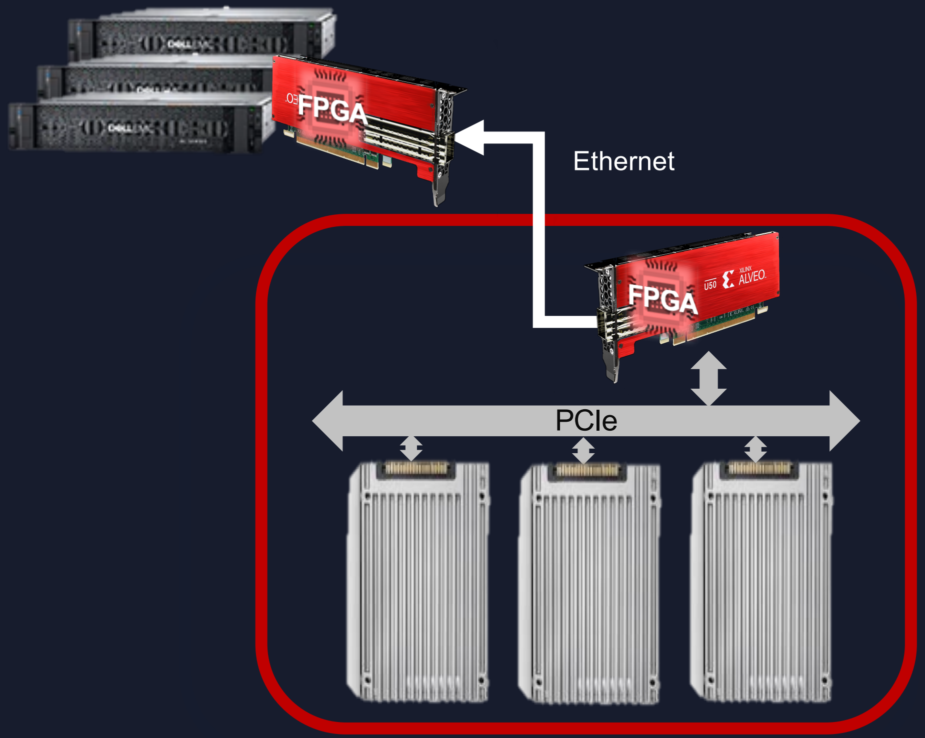

Future Data Center : Distributed Adaptive Acceleration

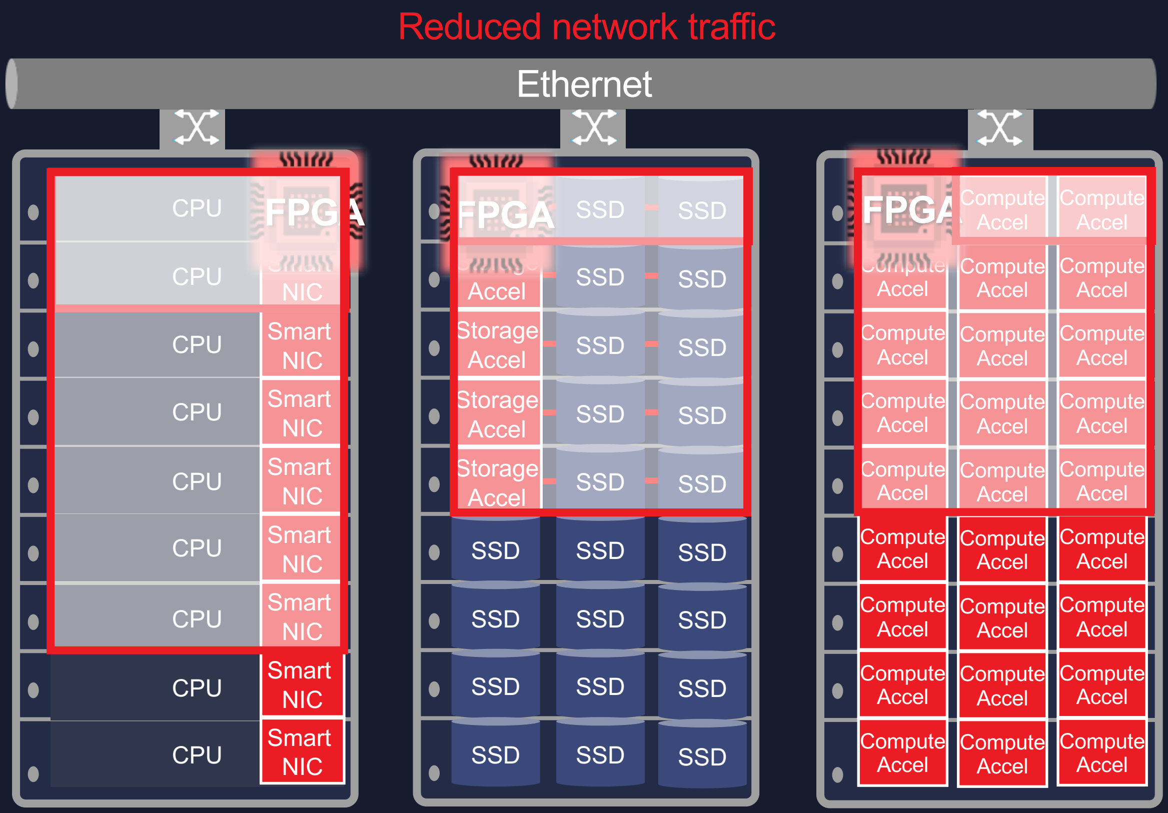

- > Composable accelerated storage, networking and compute
- > Optimized for each workload
- > Optimal infrastructure utilization

Computational Storage: Accelerating High-Speed Storage Systems

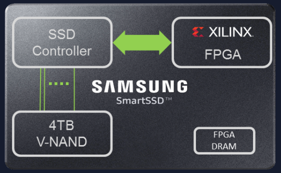

Computational storage addresses a broad range of application bottlenecks

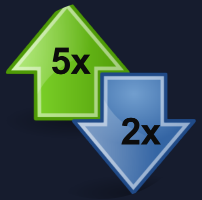

Offers data center operators >5x performance boost and up to 2x reduction of TCO

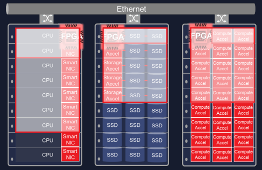

Xilinx is leading the way in distributed adaptive acceleration