Computational SSDs

Rakesh Cheerla
Sr. Product Manager, Data Center Storage, Xilinx Inc.

Jan 2019

XILINX

Datacenter First …

Hyperscale Public Cloud | Enterprise Private Cloud | Telco Cloud / Edge | High Performance Computing

Compute

Storage

Network

Financial Analytics

Security

Genomics

Machine Learning

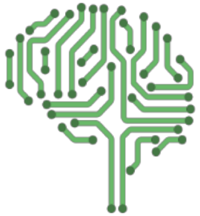

Video Transcoding

Big Data Analytics

Customer Use Cases – From Cloud to Enterprise

Contents

Datacenter Trends

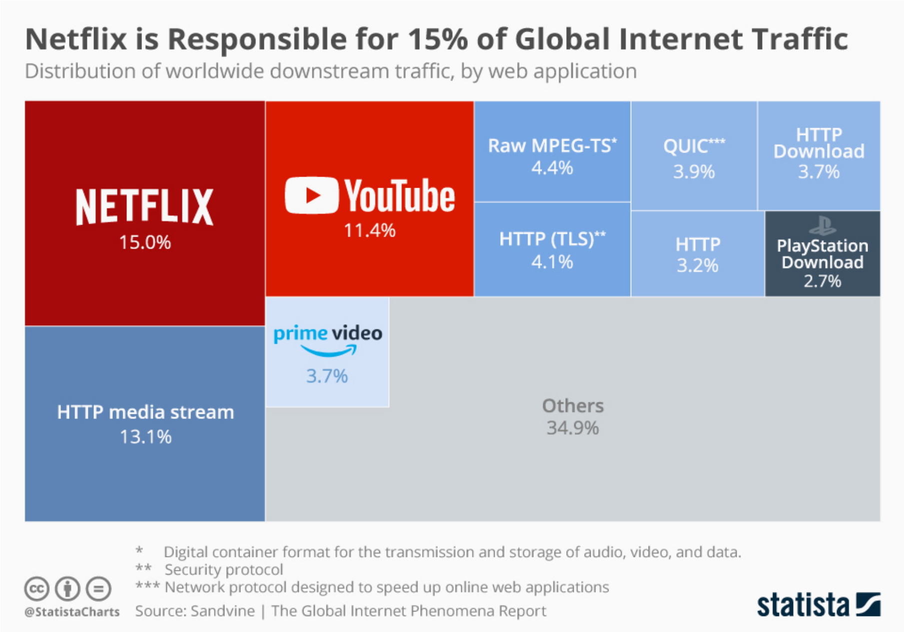

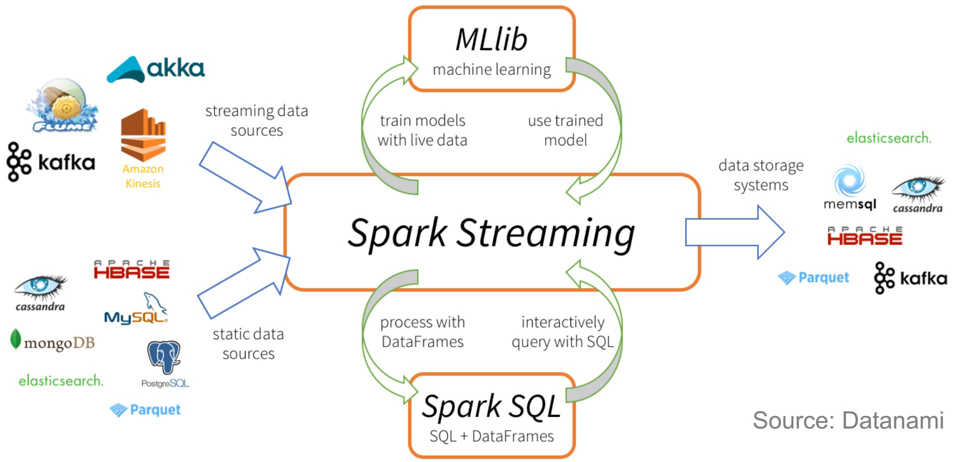

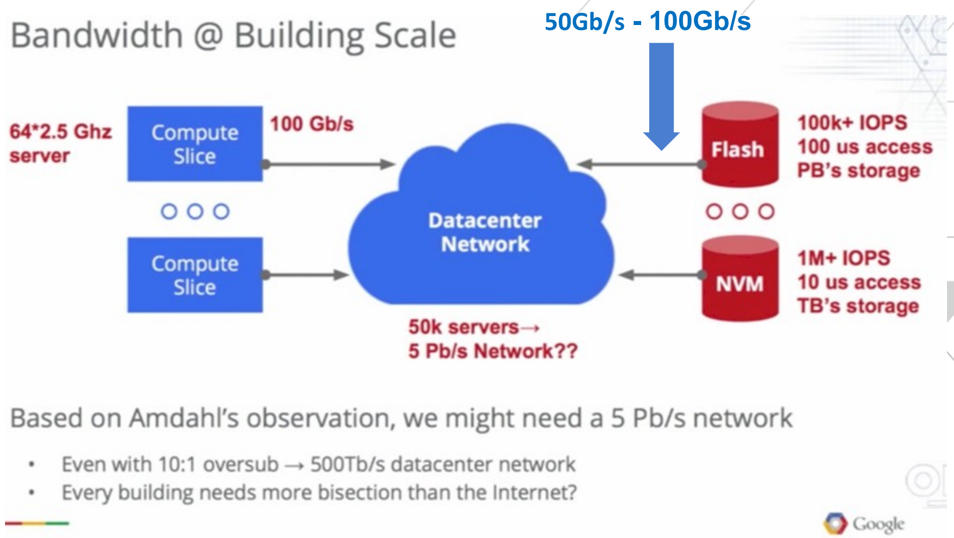

aws

Azure

Google Cloud Platform

Exponential Scale

Massively Parallel

Real Time

Huge Datasets

XILINX

The Decomposition of the Server

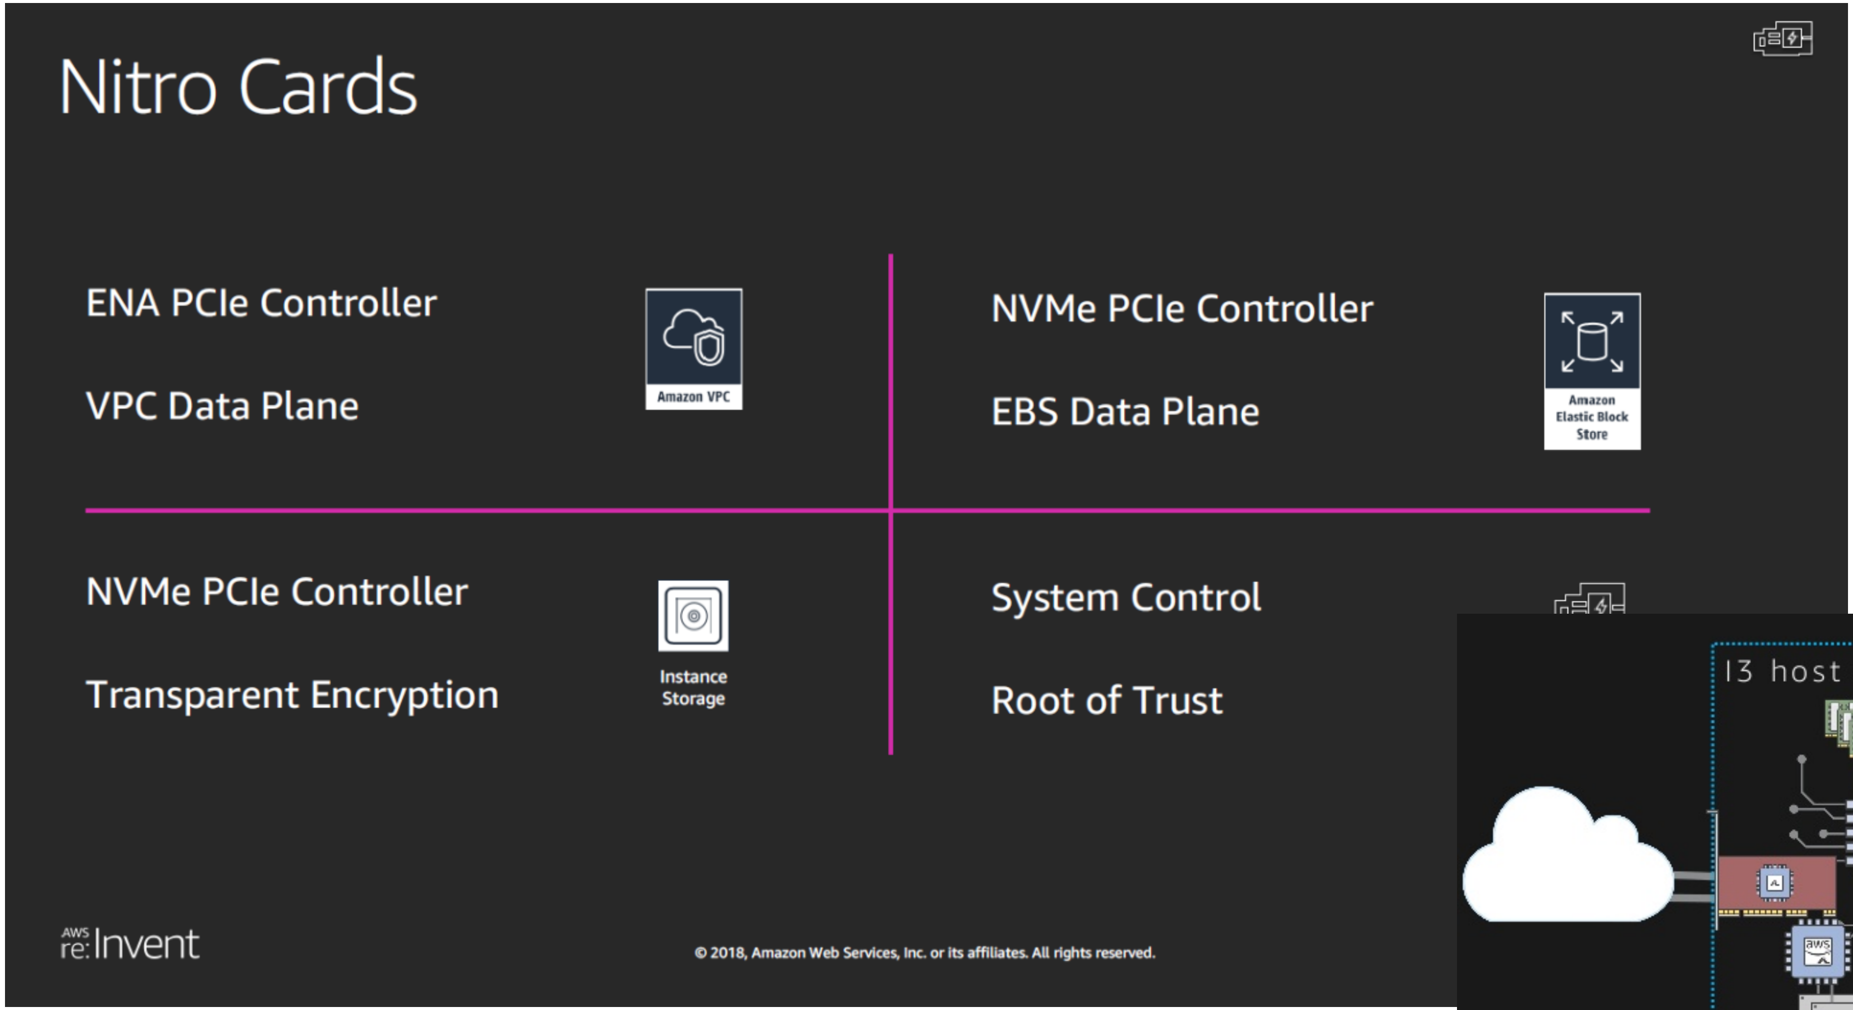

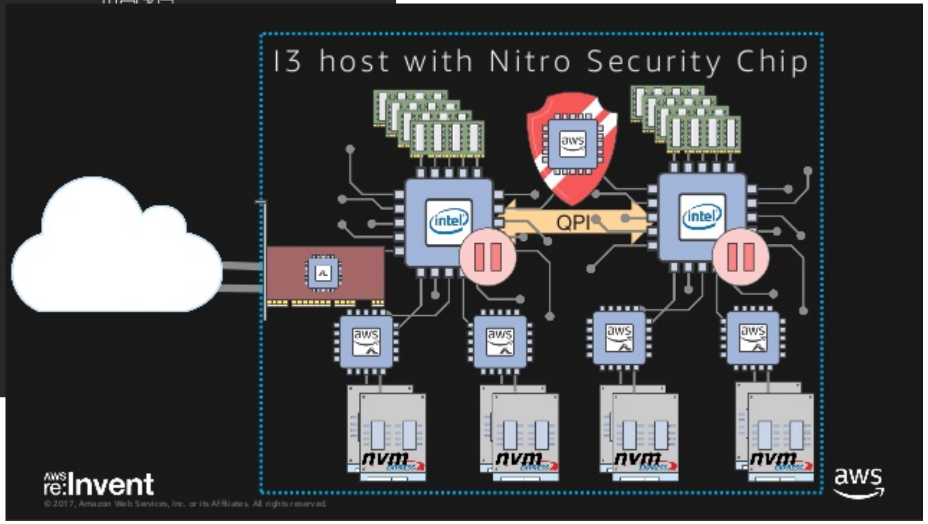

XILINX

Composition of the New Datacenter

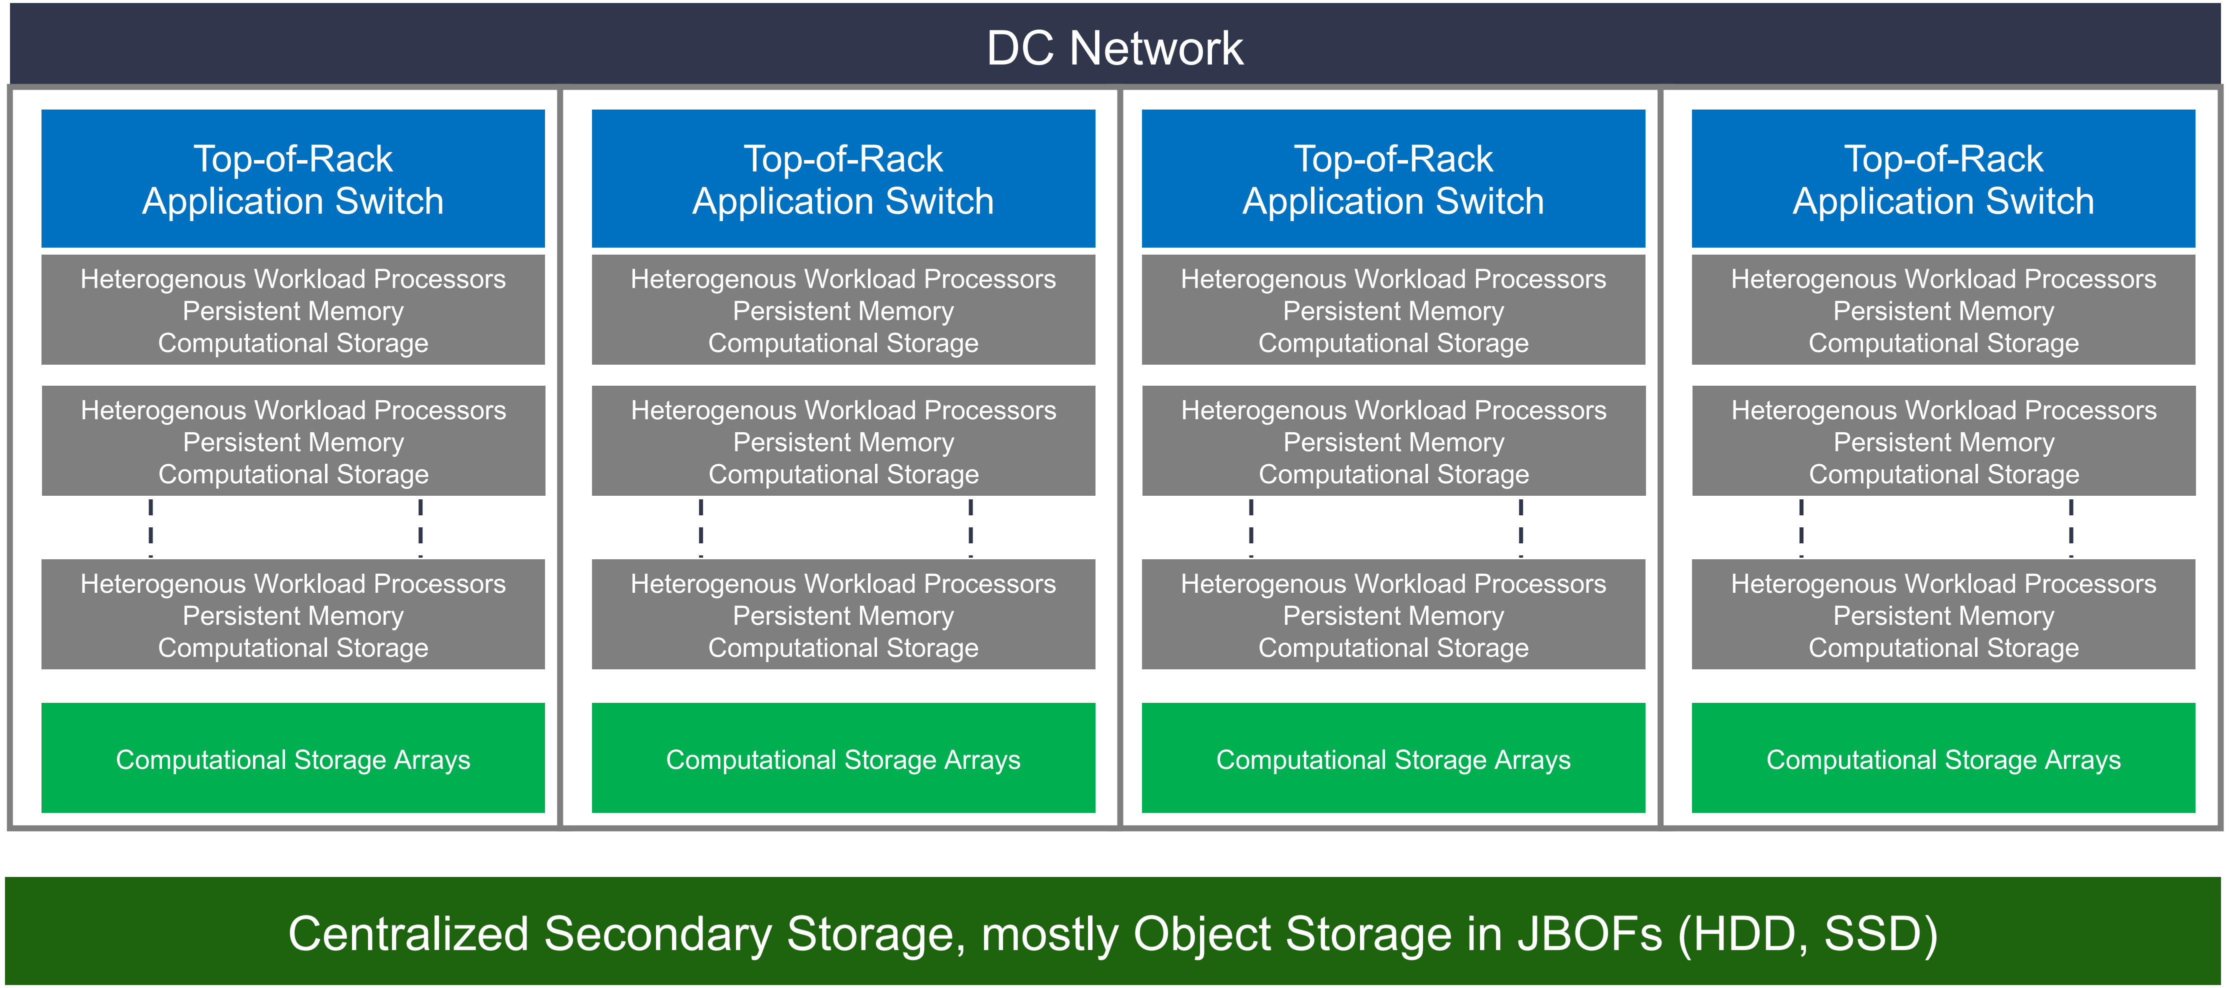

SNIA Computational Storage: Participating Companies

SNIA Computational

- EIDETICOM
- NGD systems — Bringing Intelligence to Storage
- ScaleFlux™
- arm
- Calypso Systems
- GIGAIO
- inspur
- KALRAY
- Lenovo
- MARVELL®
- Micron®
- NetApp®
- NETINT
- NYRIAD®
- SK hynix
- SAMSUNG
- TOSHIBA
- WD Western Digital®
- XILINX®

SNIA Computational Storage: Participating Companies

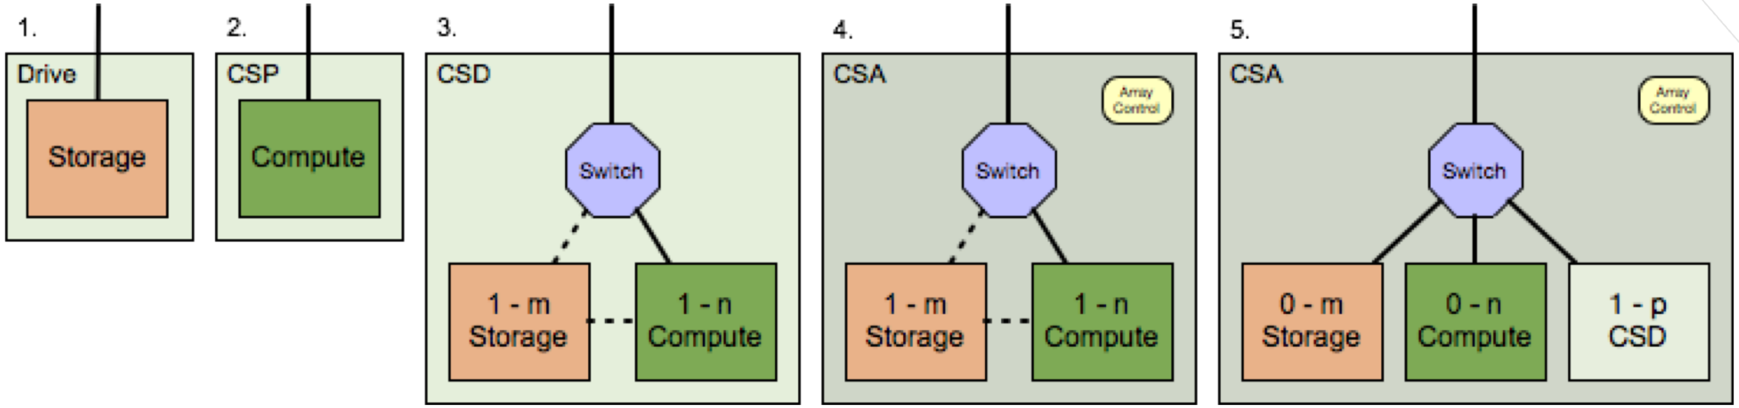

Computational Storage – Architectures that provide compute coupled to storage and/or reduce data movement

These architectures enable improvements in application performance and/or infrastructure efficiency through the integration of compute resources (outside of the traditional compute & memory architecture) either directly with storage or between the host and the storage. The goal of these architectures is to enable parallel computation and/or to alleviate constraints on existing compute, memory, storage, and I/O.

Fixed Purpose Computational Storage – Provides a well-defined computational storage service (for example: compression, RAID, erasure coding, regular expression, encryption). This service may be configurable.

General Purpose Computational Storage – Provides a programmable computational storage service (for example: this service may host an operating system, container, Berkeley packet filter, OpenCL, FPGA).

Moving Compute to Data ..

XILINX

Moving Compute to Data ….

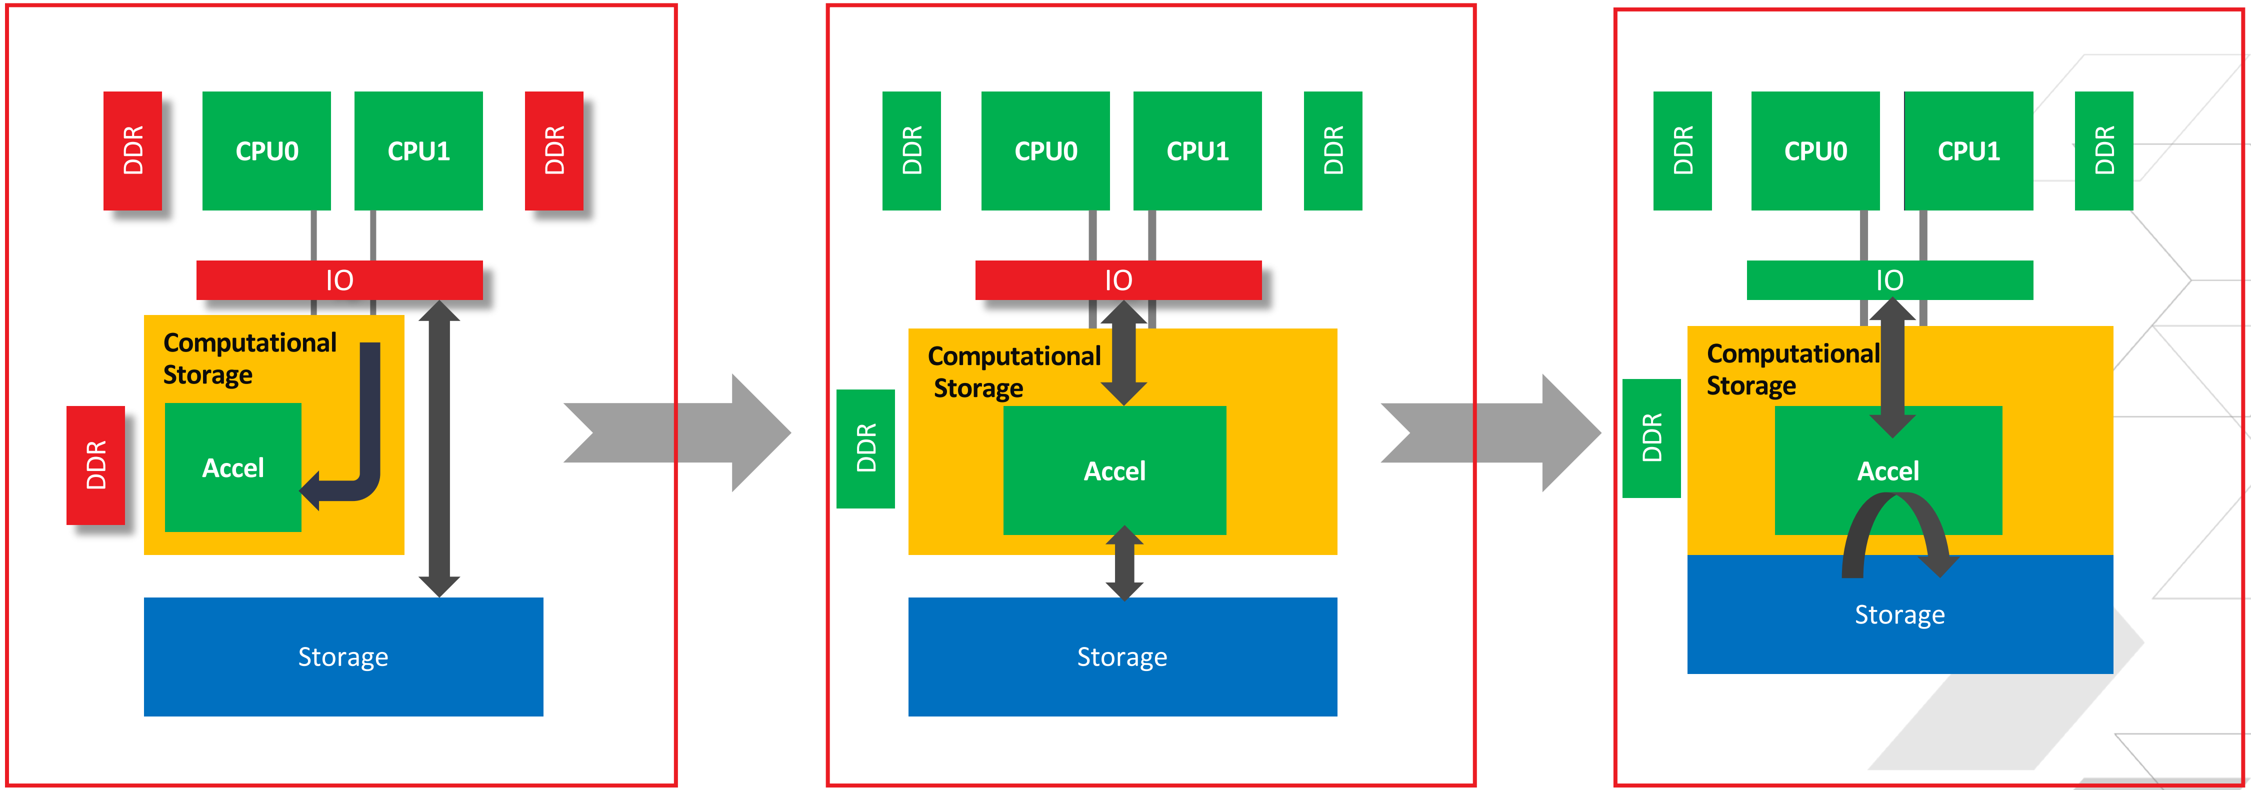

- "Offload" storage workloads
- Memory bottleneck remains
- "Inline" storage services
- Aggregation or End-Point
- Mostly Computational SSDs
- For Data-Intensive Workloads

XILINX

Moving Compute to Data ....

Adaptive Storage Acceleration

- Encryption
- Compression
- Data Dedupe
- RAID & Erasure codes
- Key-Value Offloads
- Database ETL & Query Offloads
- Spark-SQL / Map-Reduce
- Video / Image Transcoding, Processing and Delivery
- Search - Text, Image, Video etc.
- Stats / Counters
- Machine Learning

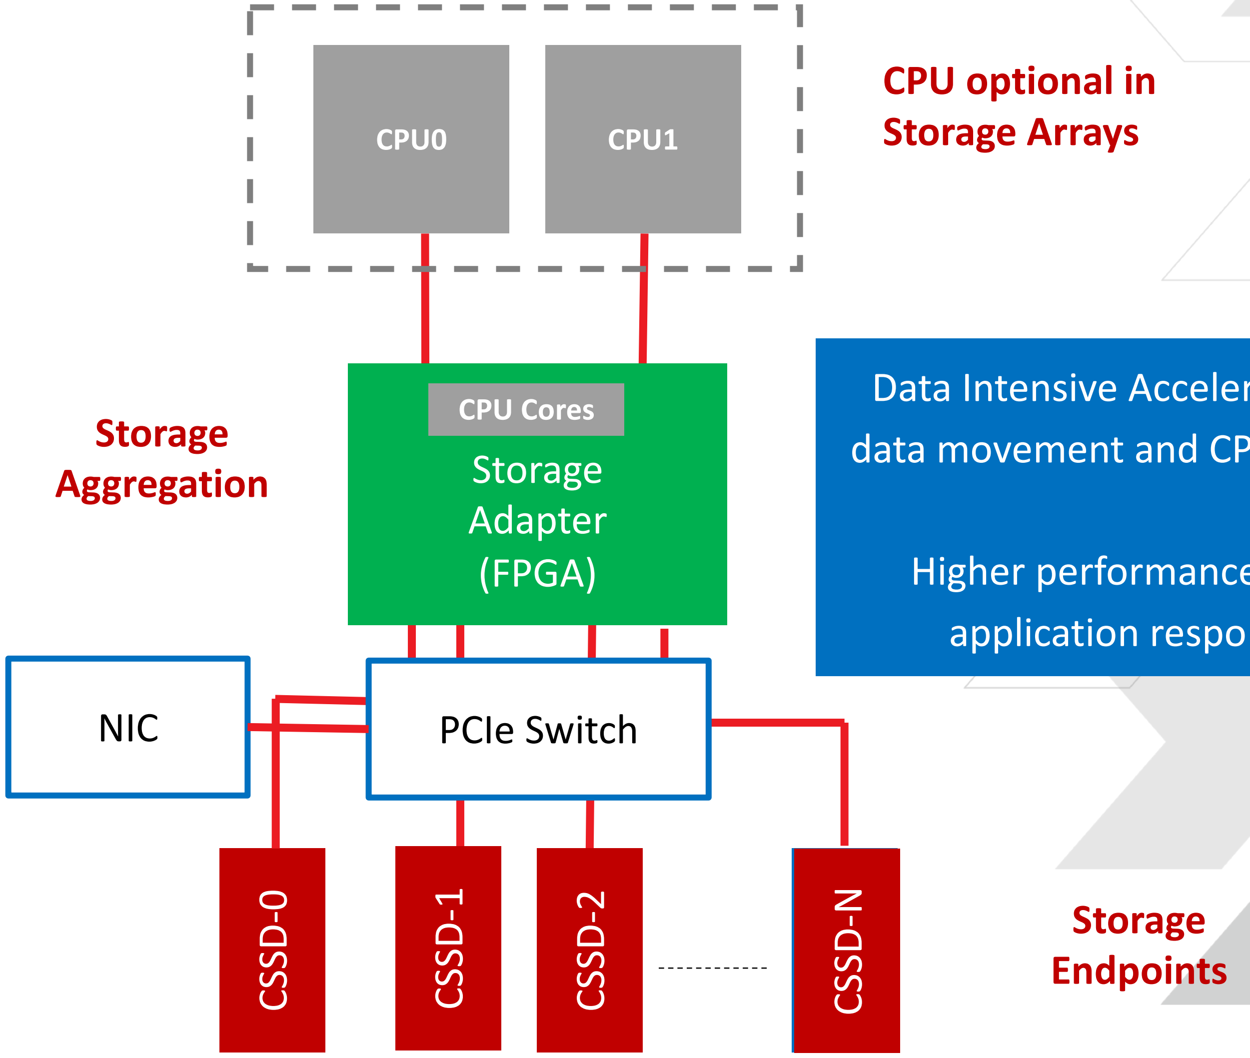

Data Intensive Acceleration reduces data movement and CPU interruption.

Higher performance and better application response times

XILINX

Computational SSDs

What is a Computational SSD ?

- Tightly coupled CSSD Sub-System
  - Embedded CPU Cores
  - Hardware Accelerators
  - Memory
  - NAND Flash Chips
- A purpose-built "internal" data paths
  - Any to any connectivity
  - 10X - 100X Internal Bandwidth
- Distributed, Scale-Out Model

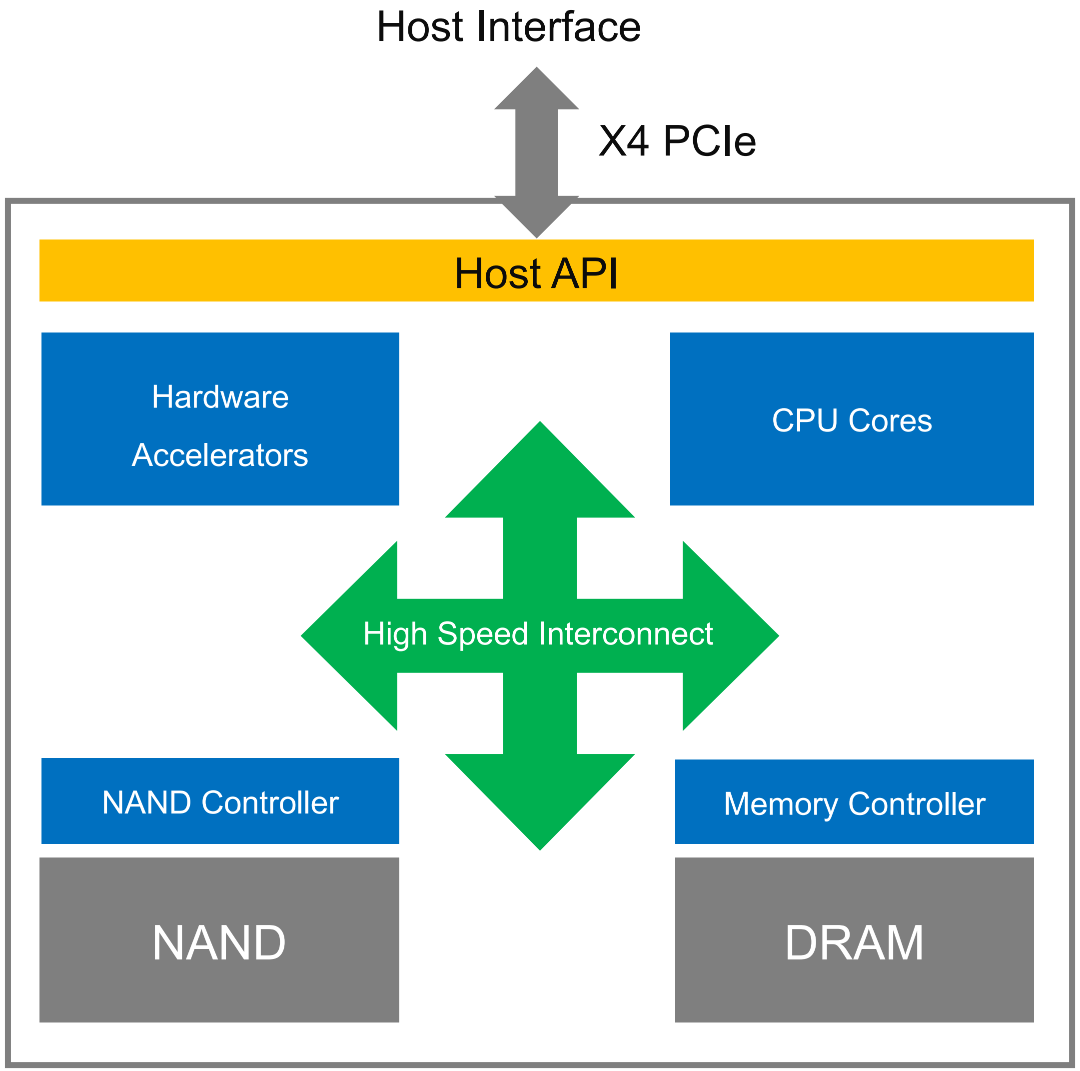

When do we use Computational SSDs ?

When to USE

- Large Data Transfers, PCIe Bandwidth is bottleneck
- Ingest/Data pre/post processing & analytics
- When data delivery can bypass the host e.g. video delivery
- Ability to move Software App to Storage

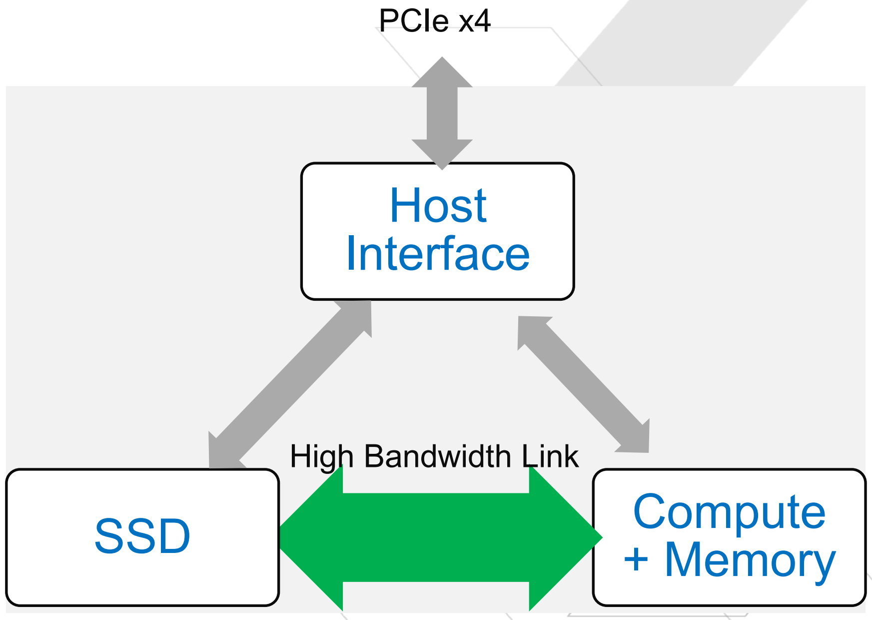

When to SKIP

- Compute heavy with very small data-transfer overhead
- Computation on small data e.g. in-memory compute
- Little to no parallelism. Algorithm is highly sequential

Remember Amdahl's Law
Balance compute vs. IO

Xilinx Adaptable Storage Platforms & Ecosystem

Xilinx Application Ecosystem Partners

LIFE SCIENCES
DRAGEN, edico genome, Falcon Accelerated Genomics

FINANCIAL
AON, MAXELER Technologies, ZoTech, CME Group, RYFT

VIDEO
NGCODEC, PATHPARTNER, skreens

MACHINE LEARNING
aws, InAccel, NIMBIX, Reconfigure.io, Atomic Rules, DeePhi, mle, PLUNIFY, XILINX Xilinx ML Suite

DATABASE
ALDEC, National Instruments, Titan, Falcon Computing, RYFT, XILINX Data Analytics Stack with Postgres

Framework, API, Python/Java/C++ Programmability

Xilinx Adaptable Storage Platforms

Layer				
Software Frameworks	SW Defined Storage	Spark	FFMPeg	TensorFLow
Libraries, Compilers, Middleware	Compression	Key Value	Video Encoder	Machine Learning
Firmware	**Storage Shell** (Firmware & Runtime)			
Integrated Development Environment & Hardware	SDAccel Environment			

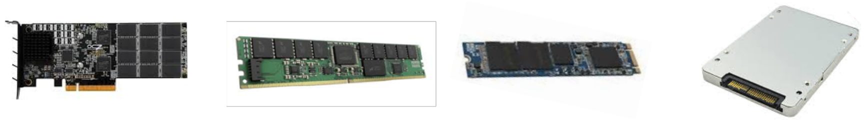

Open Sourced Xilinx Runtime Library
XILINX

XILINX

SDAccel Development Environment

Host code

- OpenCL APIs
- C / C++

Accelerator Kernel code

- OpenCL kernel code
- C/C++
- RTL IP
- 3rd party library code

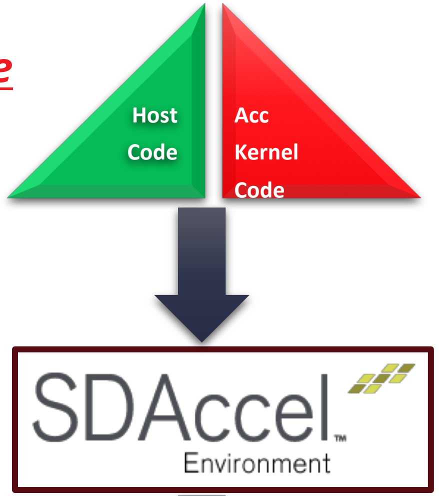

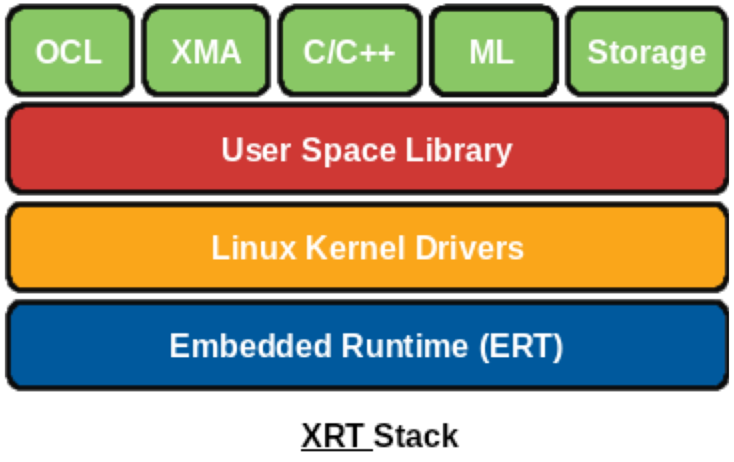

GitHub

Active Community Forums

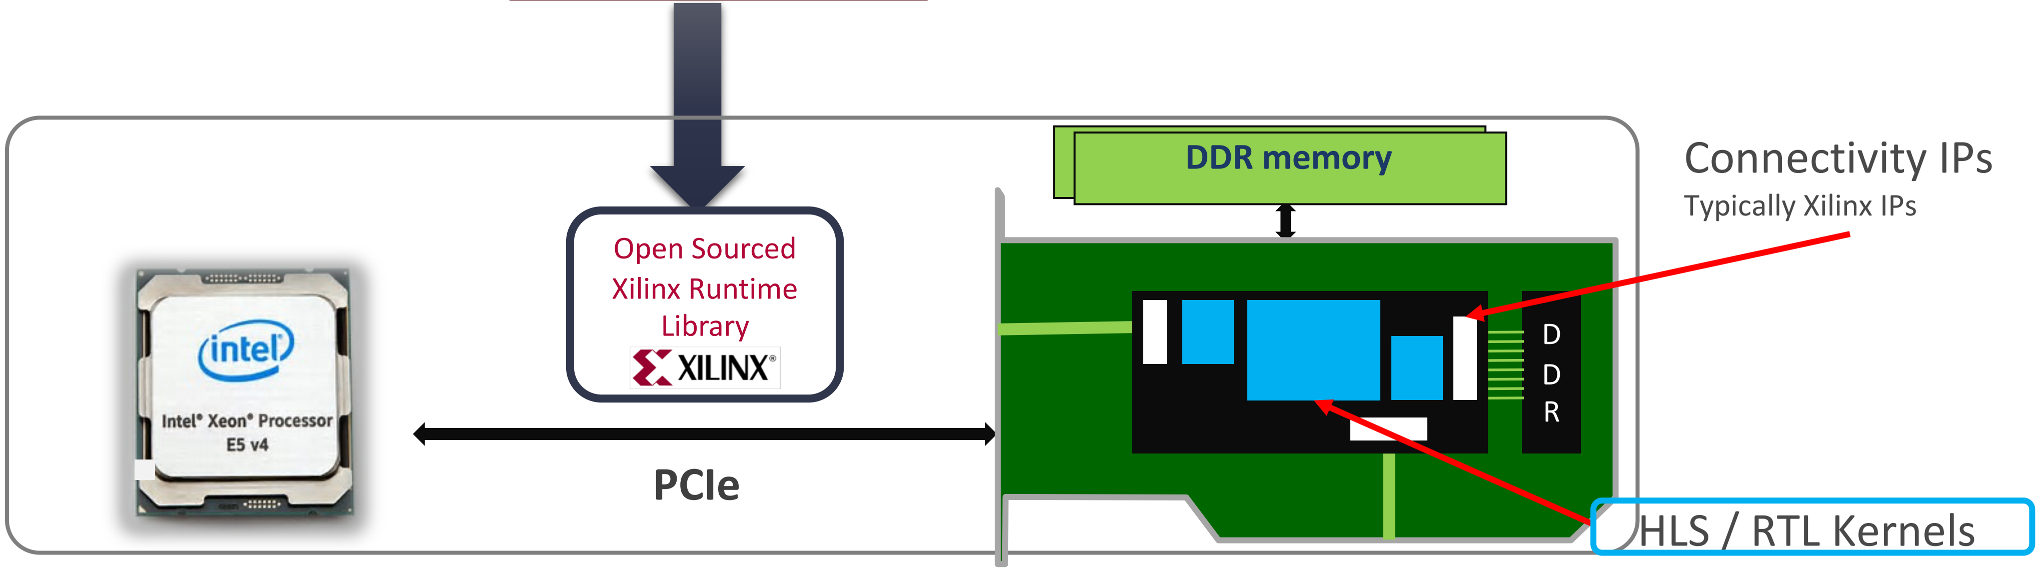

PCIe P2P, Multi-board / SSD support added. Implementing SW kernels / embedded runtime on MPSoC A53

XILINX

Introducing SmartSSD

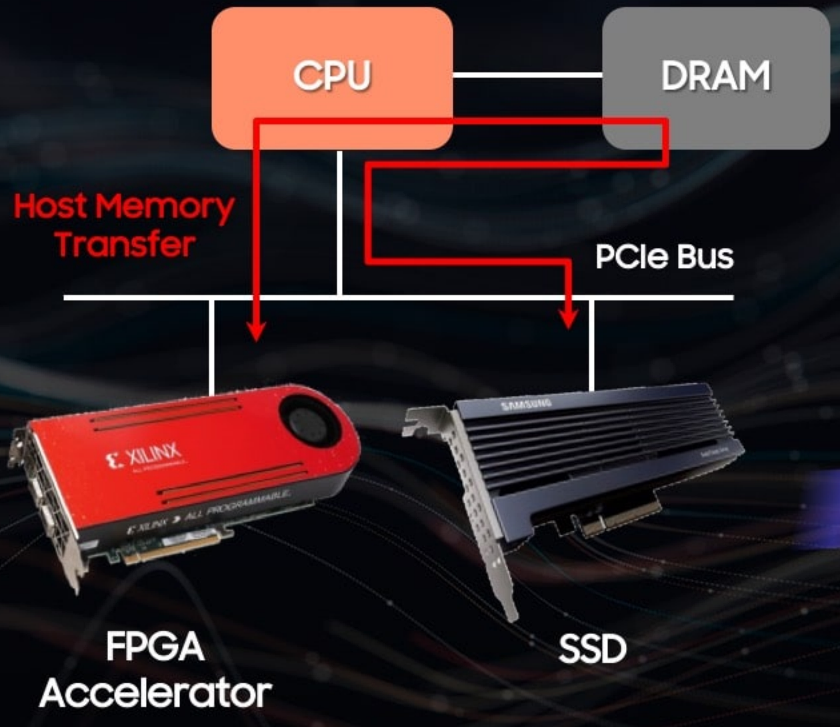

FPGA Accelerator
Data Movement through CPU and Memory

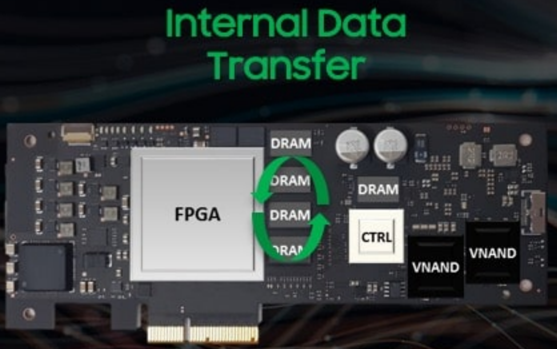

SmartSSD

Scalable and accelerated performance with SmartSSD

Next Gen SmartSSD

Greater integration

Continued...

SmartSSD

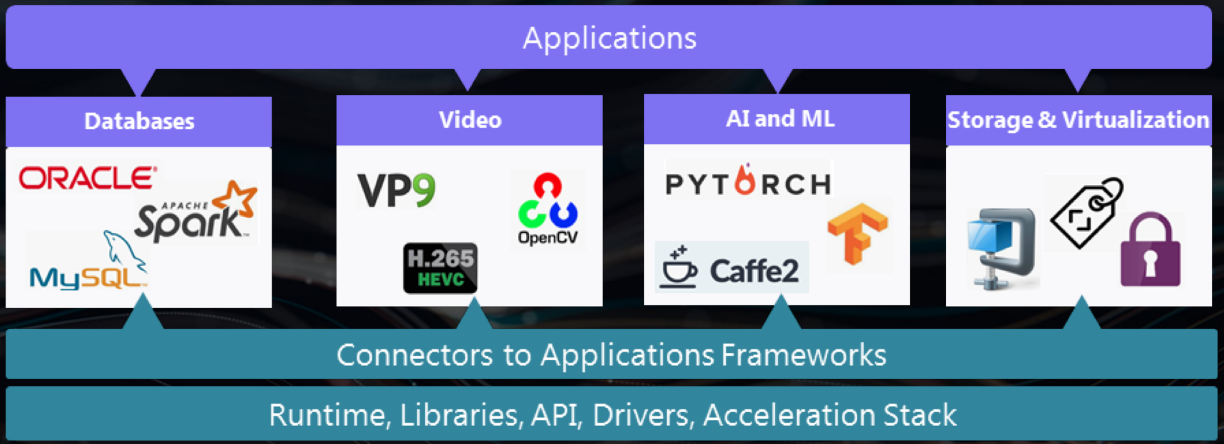

bigstream

XILINX
ALL PROGRAMMABLE.

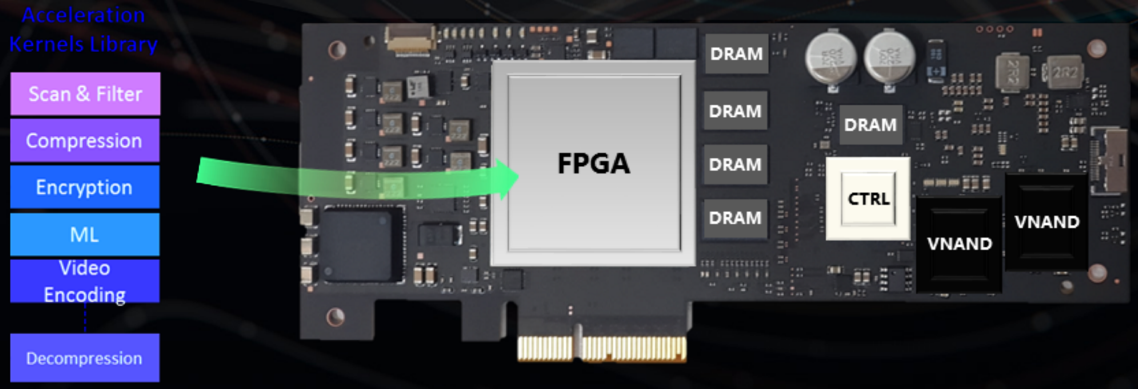

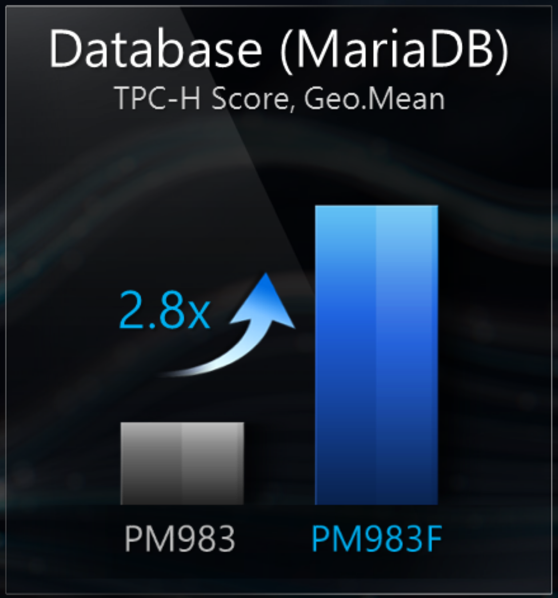

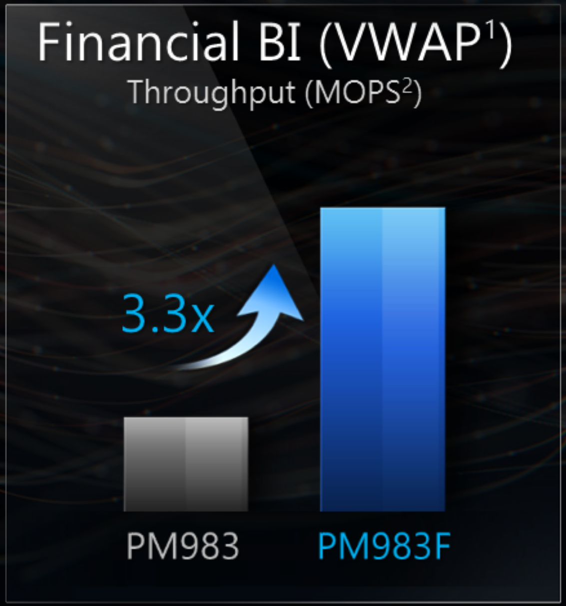

SAMSUNG

1. VWAP: Volume Weighted Average Price 2. MOPS: Million Operations per Second

XILINX

One Platform -> Many Services

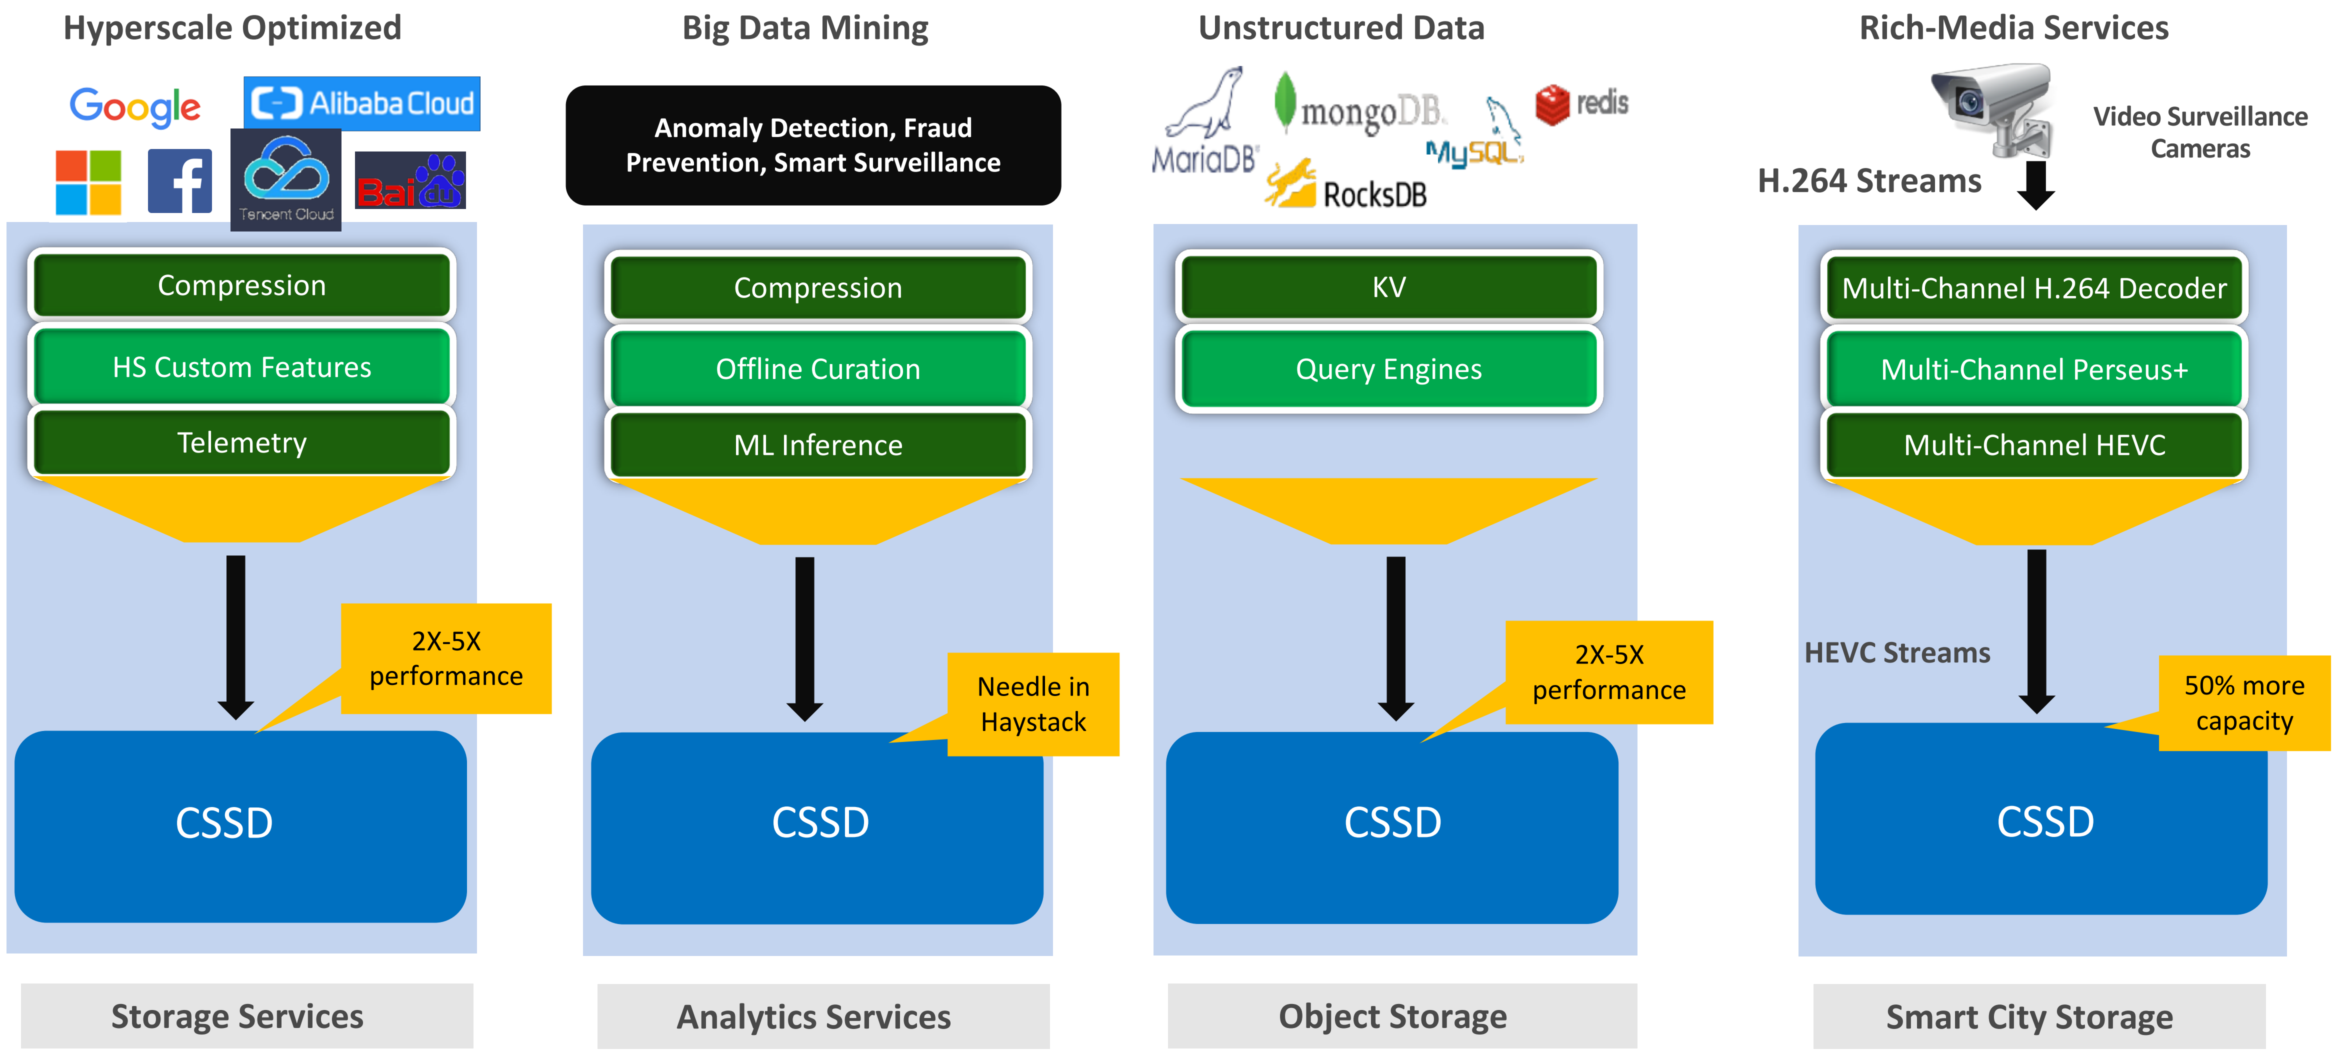

Multiple Service Personalities on one platform. Adaptable. Intelligent.

Computational Storage Summary

> Latency, throughput and power advantages

> System thinking key to accelerating applications

> Samsung SmartSSD platform for application development

> Join SNIA Computational Storage TWG: https://www.snia.org/computational

SNIA™ Computational

XILINX.

The best way to predict the future is to create it!

Peter Drucker

Adaptable.

Intelligent.

XILINX®