STORAGE DEVELOPER CONFERENCE
SDC 21
BY Developers FOR Developers

Virtual Conference
September 28-29, 2021

A SNIA Event

# Next Generation Cloud Data Centers

Presented by Dr. Jai Menon, Chief Scientist, Fungible

# Agenda

01

Current Cloud Data Center Challenges

02

How to Address These Challenges

03

Next Generation Data Center Building Blocks

 

# Current Data Center Challenges

 

# Five Current Data Center Challenges

## High Power/Footprint/TCO

- Underutilization of expensive SSDs, GPUs, networks
- General purpose server cores consumed by data-centric[1] tasks

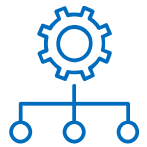

## Lack of Infrastructure Agility, Infrastructure Silos

## SKU Explosion

## Expensive or Weak Security, Lack of Reliability

## Inadequate Performance

[1]Data-centric = stateful, multiplexed processing of high b/w data streams

 

# Approach to Addressing the 5 Challenges

## 3 key elements to solution -- DPUs, penalty-free disaggregation of resources, and composability

- Improve power/footprint/TCO by
  - Penalty-free disaggregation and pooling of resources such as SSDs, GPUs
  - Judicious use of DPUs for data-centric tasks such as storage, networking, security
  - More efficient data-center networking
- Enhance agility by using composable infrastructure
- Reduce # of SKUs needed through composable disaggregated infrastructures
- DPU enhanced security and reliability
- Optimize performance by using CPUs, GPUs and DPUs in the data center

 

# Fungible’s DPU

 

# Fungible Offerings are Powered by the Fungible DPU™

## A New Class of Microprocessor Purpose-Built for the Data-Centric Era

The Fungible DPU is a new class of programmable microprocessor that:

Enables 10x more efficient execution of data-centric tasks

Implements an endpoint for a more efficient data center network

**CPU**

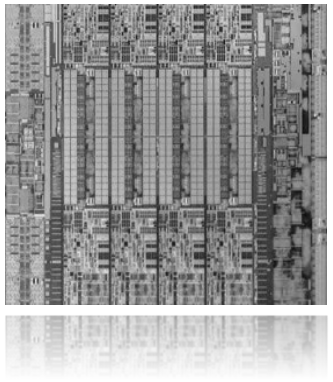

General-Purpose

**GPU**

Vector Floating Point

**Fungible DPU™**

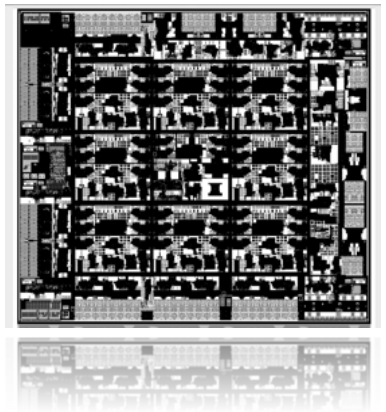

Data-Centric

- Standard external interfaces – PCIe, Ethernet
- Programmable in C
- Many multi-threaded cores & hardware accelerators for crypto, compression, etc.,

[1] Data-centric = stateful, multiplexed processing of high b/w data streams

 

# Fungible DPUs Have Unique Architectural Advantages

10X faster at data-centric tasks

## Fungible DPU

- Tightly Coupled CPUs & Accelerators
  - Very fast interactions
- Purpose-Built CPUs for data-centric tasks
- Purpose-Built Memory Systems & Fabrics
  - Heavily threaded
  - Latency-tolerant
- Specialized Accelerators
  - Multi-threaded
  - Fully integrated with memory systems

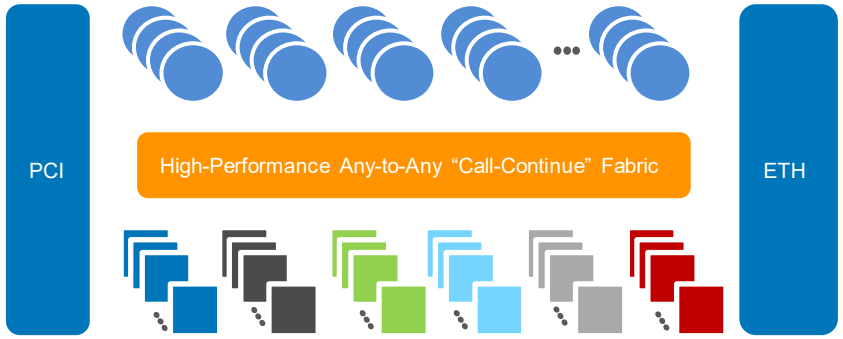


## Traditional DPUs

- Loose Coupling Between CPUs and Accelerators
  - Communication via interrupts/traps is slow
  - Separate memory slows communication
  - Switching to/from GP-CPUs is inefficient
- General-Purpose CPUs not good for data-centric tasks
  - Classical coherency only
  - Caches get destroyed by streaming data
  - Latency intolerant

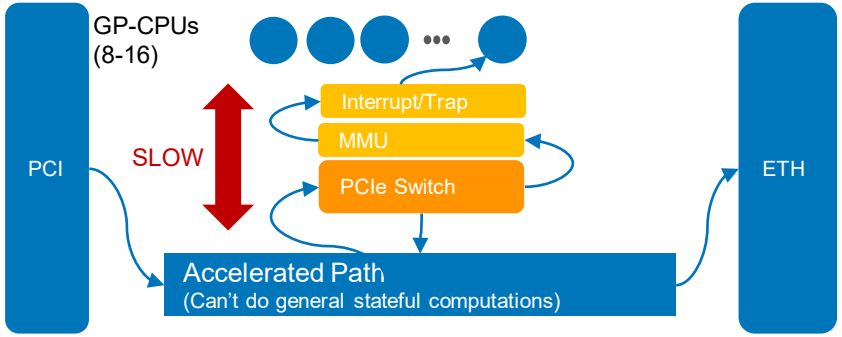


 

# Fungible DPU Converts a Standard IPoE Network Into a TrueFabric™

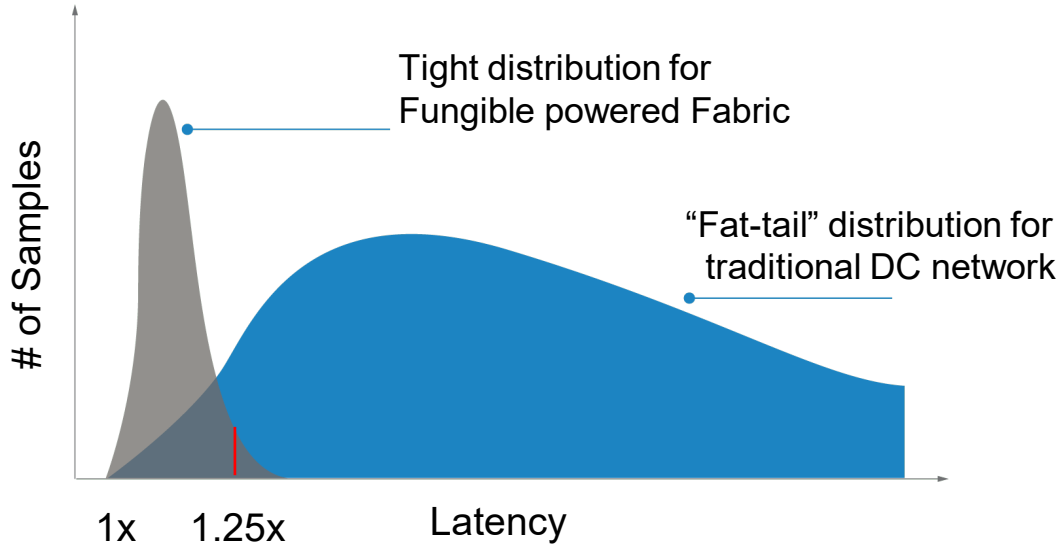


## Key Benefits of TrueFabric™

- Scalable (4 - 4000 racks)
- Full cross-sectional bandwidth
- Low latency and low jitter
- Fairness
- Congestion avoidance (even @ high loads)
- Fault tolerance (built-in detection/recovery)
- End-to-end security
- Open Standards (IPoE)

 

# Next Gen Data Center and its Building Blocks

 

# Building Blocks for the Next Generation Data Center

## Improve TCO, Performance, Security, Reliability and Agility of Cloud Data Centers

**DPU based Scale-out storage**
(Centralized Storage)

Improve the Storage


**DPU-based Host Cards**
(Handle Data-Centric Tasks)

Improve the Compute

Composer S/W

**Composer Software**
(Dynamic creation of VDCs)

Improve the Data Center

3 key elements of the solution – DPU, penalty-free disaggregation, composability

 

# DPU-Based Scale-Out Storage

 

# Cloud Data Center Requirements on Block Storage

| Requirement | Benefit |
|---|---|
| Storage is pooled and shared across all servers | High storage utilization; Independent storage scaling |
| **Very high and consistent performance** | Penalty-free disaggregated storage |
| Scale out | Grow as you need; Pay as you grow |
| **Line-rate Compression** | TCO; high storage utilization |
| **Line-rate Encryption** | Security; workload consolidation |
| Multi-tenancy (per vol protection, encryption, QoS) | Workload consolidation |
| REST API to manage PBs of data | TCO |
| **Rack scale resiliency @ low overhead using networked EC** | Very high reliability @ low cost |
| Supports VMs, containers, bare-metal | Workload consolidation |

**Blue rows need DPU**

 

# Fungible Storage Using DPU

## Improve TCO, Performance, Security, Reliability and Agility of Cloud Data Centers

HIGH PERFORMANCE

LOW TCO

SECURITY

AGILITY

SIMPLICITY

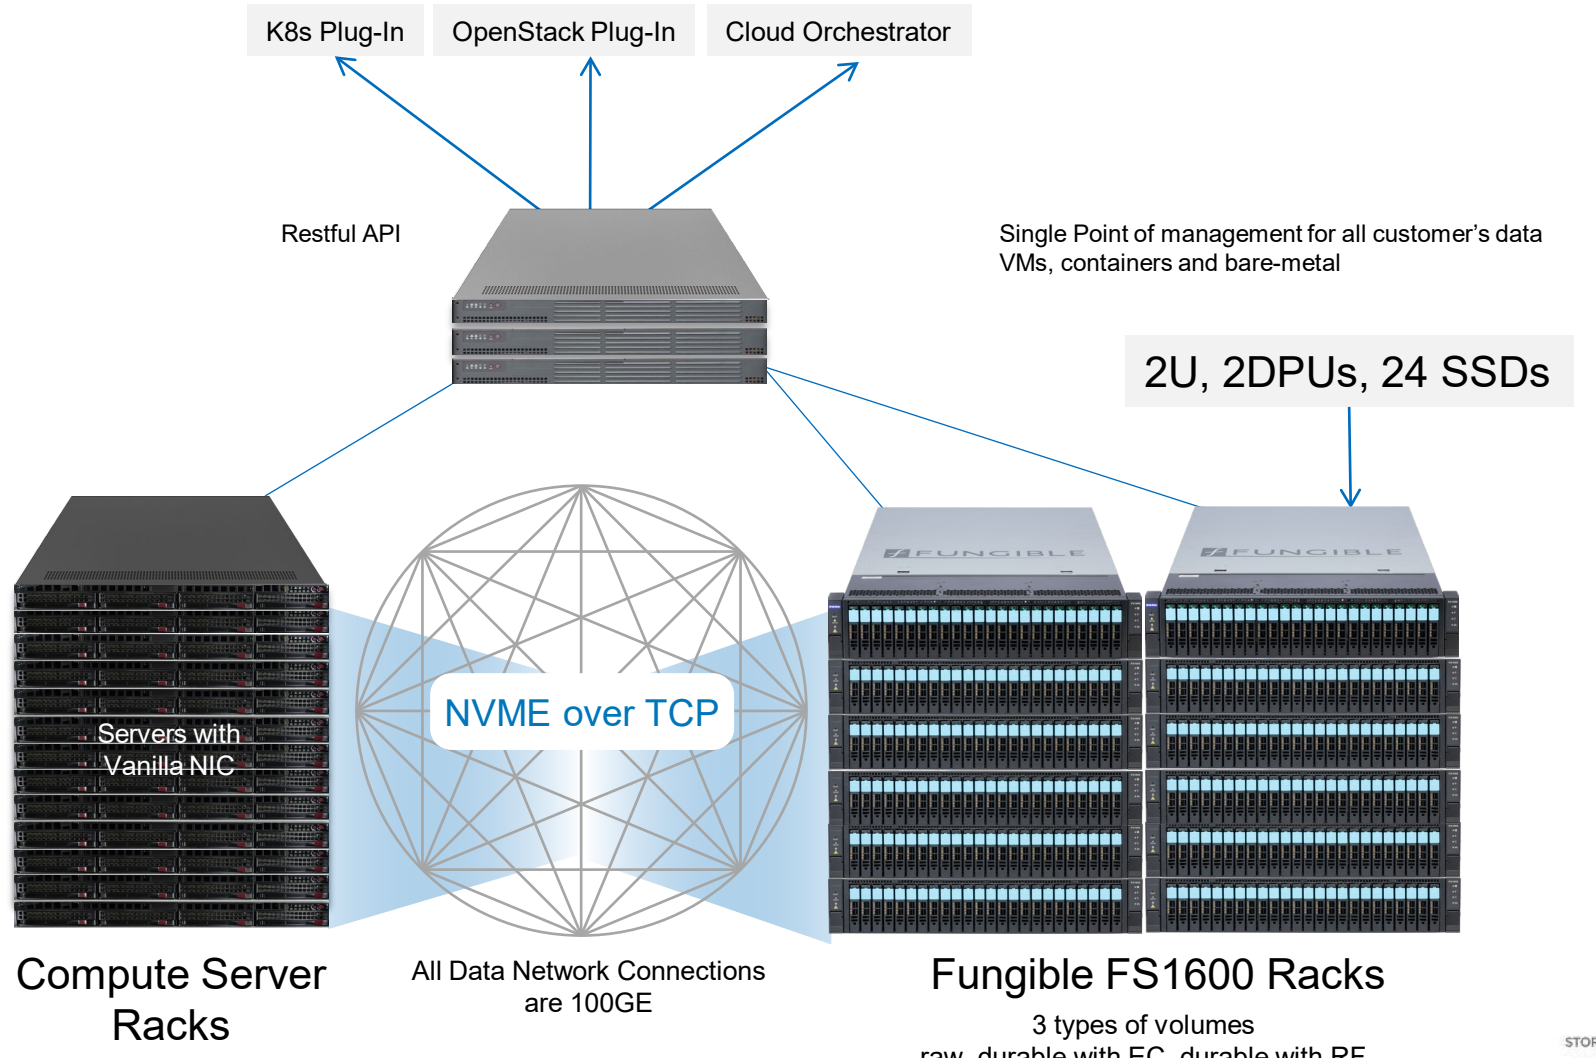


 

# Fungible Storage Performance

| | READ |
|---|---|
| **Raw**<br>**Single Node** | 15M 4KB IOPS<br>75 GBytes/sec |
| **Network Protected (RF=2)**<br>**Two Nodes**<br>(SSD and node failure protection) | 15M 4KB IOPS<br>120 GBytes/sec |
| **Network Protected 4+2 EC**<br>**6 nodes**<br>(SSD & node failure protection) | 20M KB IOPS<br>160 GBytes/sec |

**LINEAR PERFORMANCE SCALING MEASURED UP TO 16 NODES, EXPECT CONTINUED LINEAR SCALING BEYOND THIS**

 

# Fungible Storage Performance on MySQL Database

## DPU-based storage can be as fast as locally attached storage – penalty-free disaggregation

- Fungible FSC delivers better read/write latency than DAS even at high TPS
- Single instance MySQL 8.0
- XFS filesystem
- Innodb storage engine
- Innodb_buffer_pool_size= 16G
- DAS w/ MySQL table compression “zlib”
- FSC compression but no MySQL table compression
- Yahoo! Cloud Serving Benchmark (YCSB)
- 4KB record size
- 32,000,000 record count

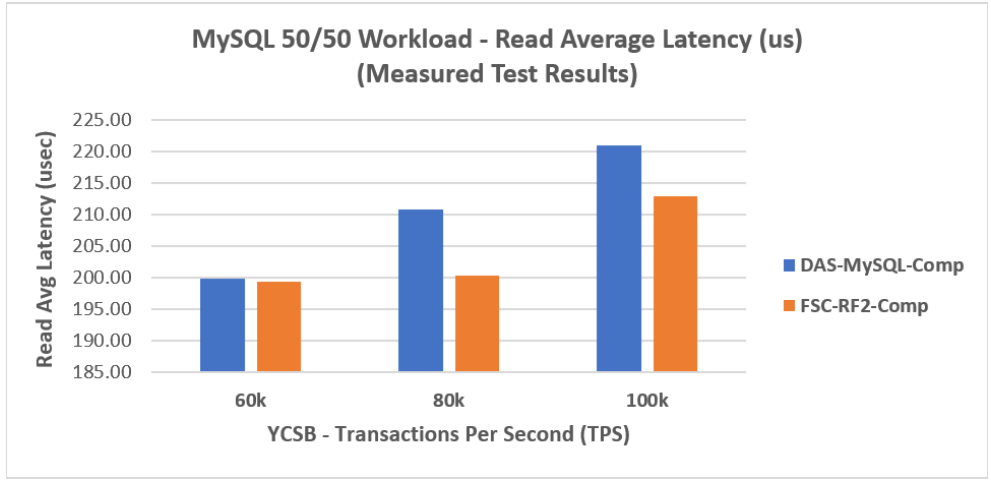


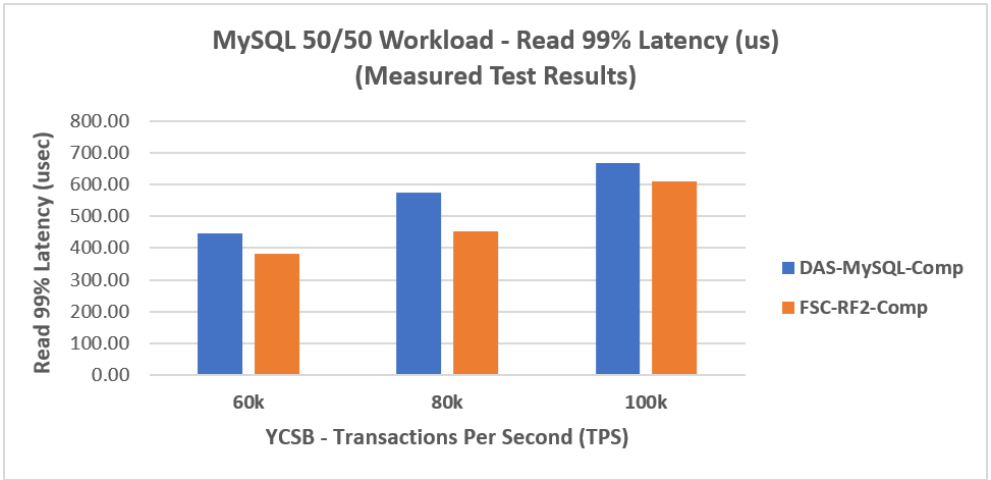


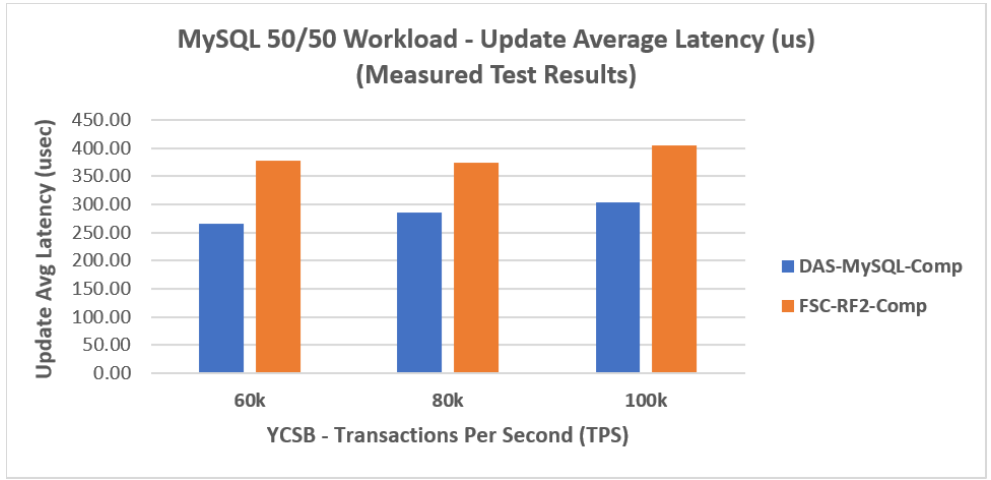


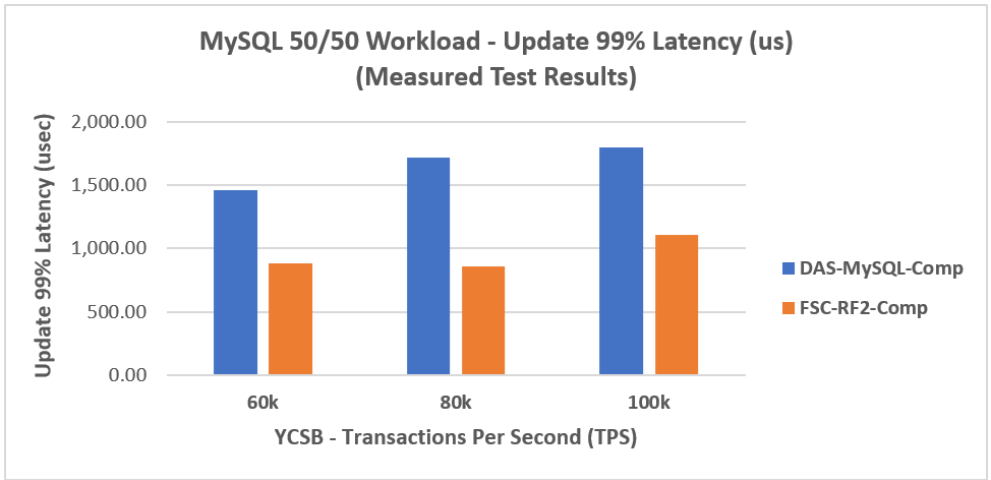


LOWER IS BETTER

 

# DPU-Based Host Cards

 

# Capabilities of DPU-Based Host cards

## Offload data-centric work from servers

| | Sample Features |
|---|---|
| Network | TCP, RoCE, TrueFabric™, Network overlays, switching/routing |
| Security | IPsec, SSL/TLS, kTLS, stateful firewall |
| Storage | **High speed NVMe/TCP offload**, **remote boot** |
| **Virtual PCIe Switch** | **GPU disaggregation** |


**Unique to Fungible DPU**

 

# Storage Initiator on DPU-Based Host Card

- Networked storage appears as local
- Server cores freed up for applications
- Applications in VMs have bare-metal performance
- End-to-end encryption and compression
- Diskless servers (remote boot)

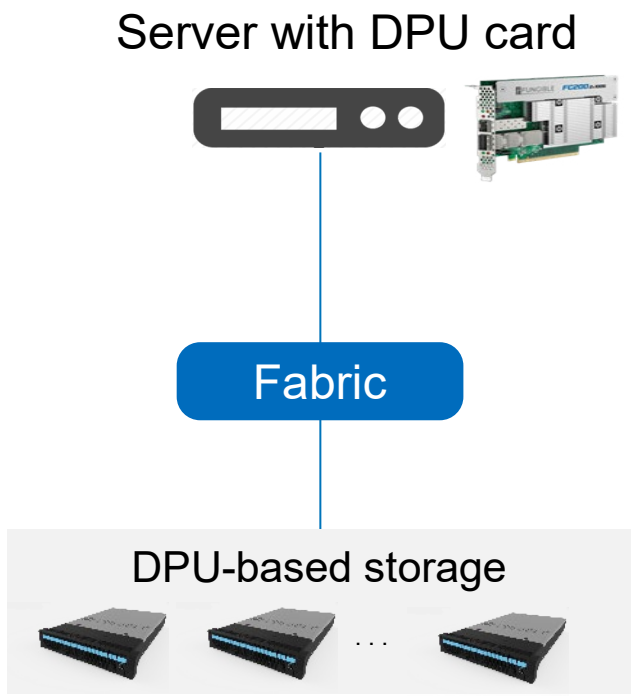


 

# Fungible Host Cards



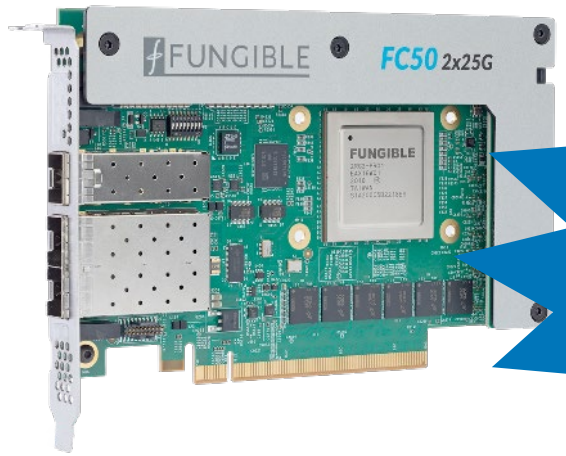


Up to 35+ Mpps
Up to ~2M IOPS
30W - 45W
FHHL Card

**ONE PLATFORM TO ADDRESS ALL INFRASTRUCTURE SERVICES**

| GEN 3.0 / 4.0 | High-Speed I/O | Virtualization | Local Storage | Accelerators | DMA |
|---|---|---|---|---|---|
| 16/32 Lanes<br>Gen 3 - 2x16, Gen3/4 - 4x8 | 2x 25G/50G/100G | SR-IOV and Virtio-Direct<br>32PFs, 512 VFs | *NVMe PCIe SSDs<br>Up to 4 NVMe SSDs | Crypto, Compression, Regex, EC, Lookup | Host Memory To DPU Memory |

 

# Composer Software

 

# Composer Software

## Data Center Composer Software

- Host & isolate multiple tenants using VPCs
- Composes bare-metal servers for tenant workloads
- Provision workloads
- Role based Access Control
- Application marketplace

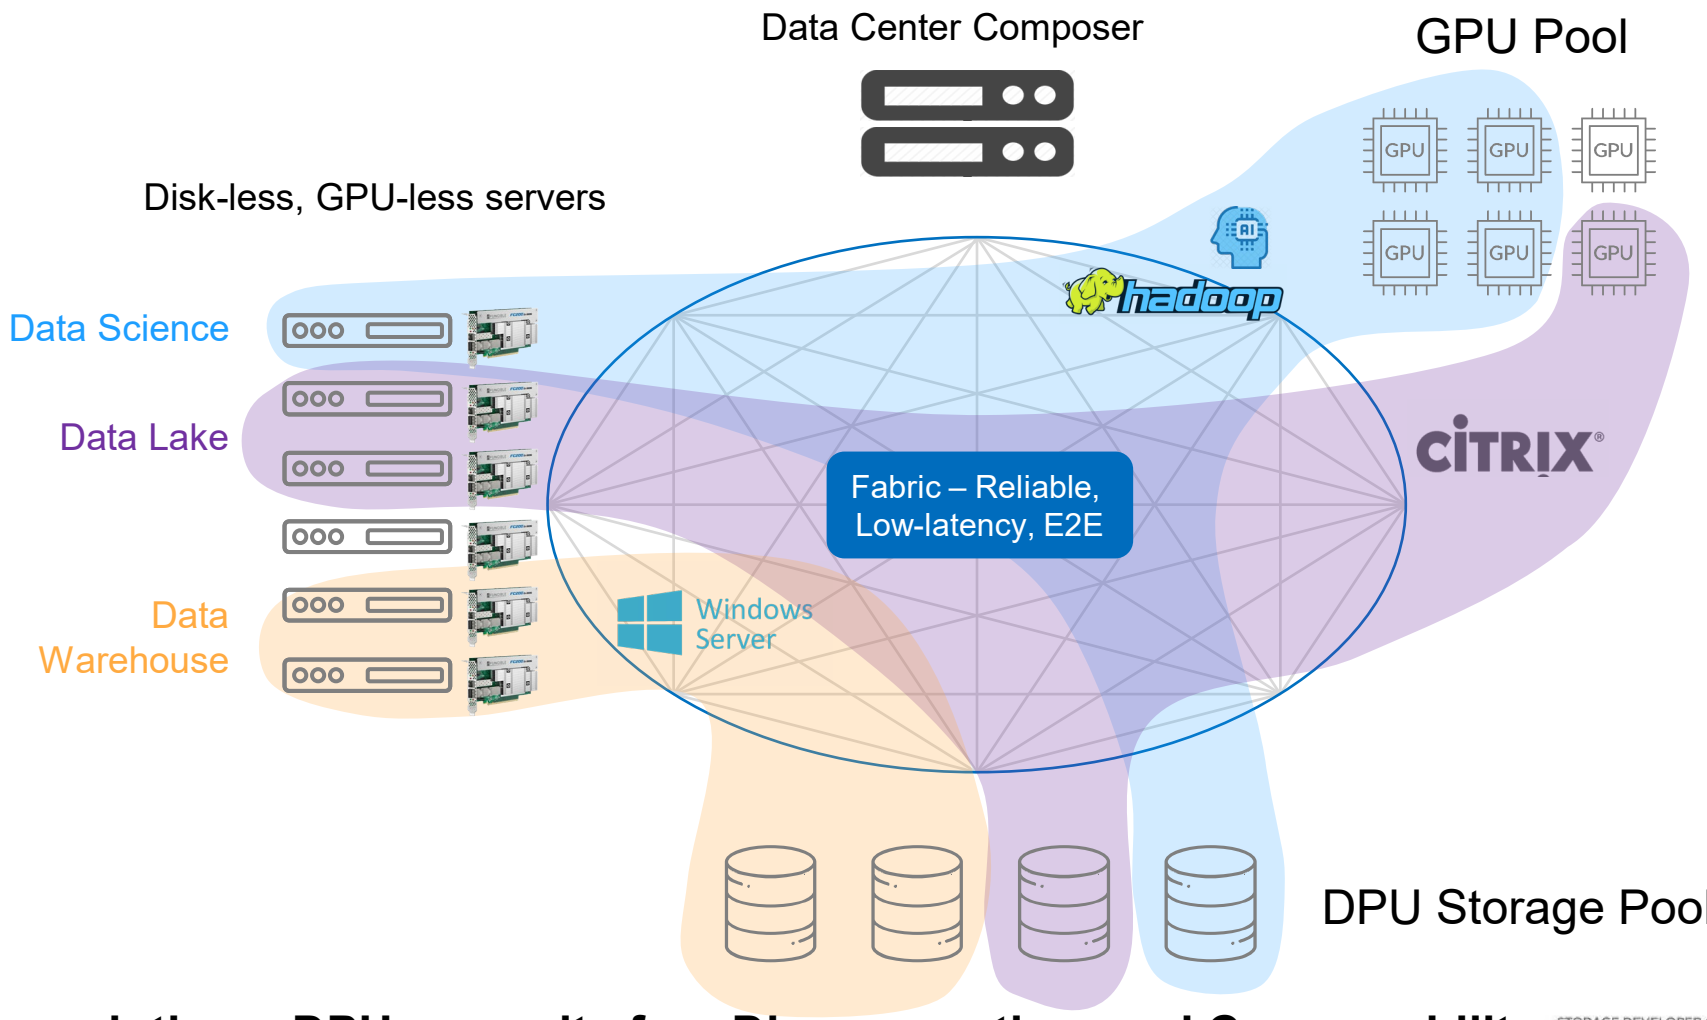


**3 Key elements of the solution – DPUs, penalty-free Disaggregation, and Composability**

 

# DPU-Based Next Gen Data Center

## Physical View

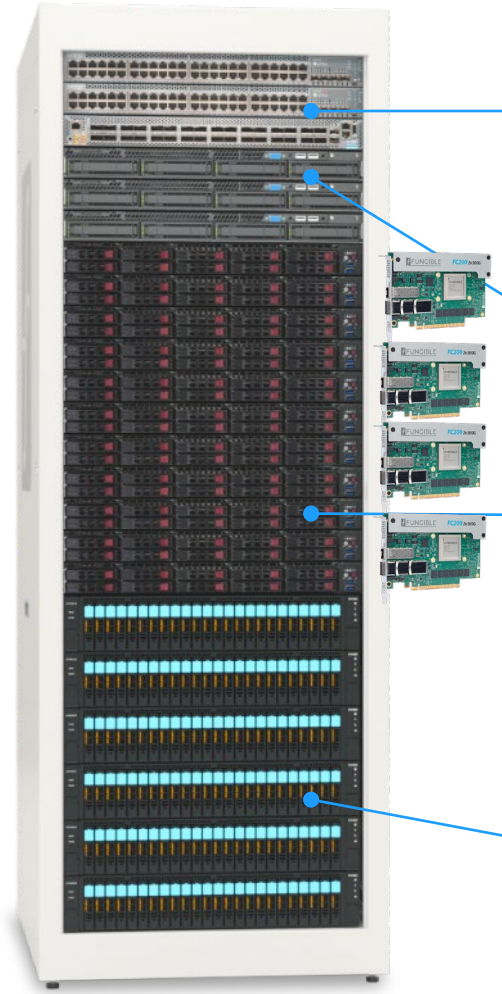

### STANDARD TOR SWITCHES

- Data, Management

### COMPOSER SOFTWARE

- Multi-tenant bare-metal virtualization solution
- API managed

### STANDARD SERVERS + DPU-based HOST CARDS

- Runs TrueFabric™
- Offload network, storage & security stacks
- Support GPU disaggregation

### DPU-Based STORAGE CLUSTER

- High performance, secure, agile, cost-efficient

**3 key elements of the solution – DPU, penalty-free Disaggregation, and Composability**

 

# DPU-Based Next Gen Data Center

## Logical View

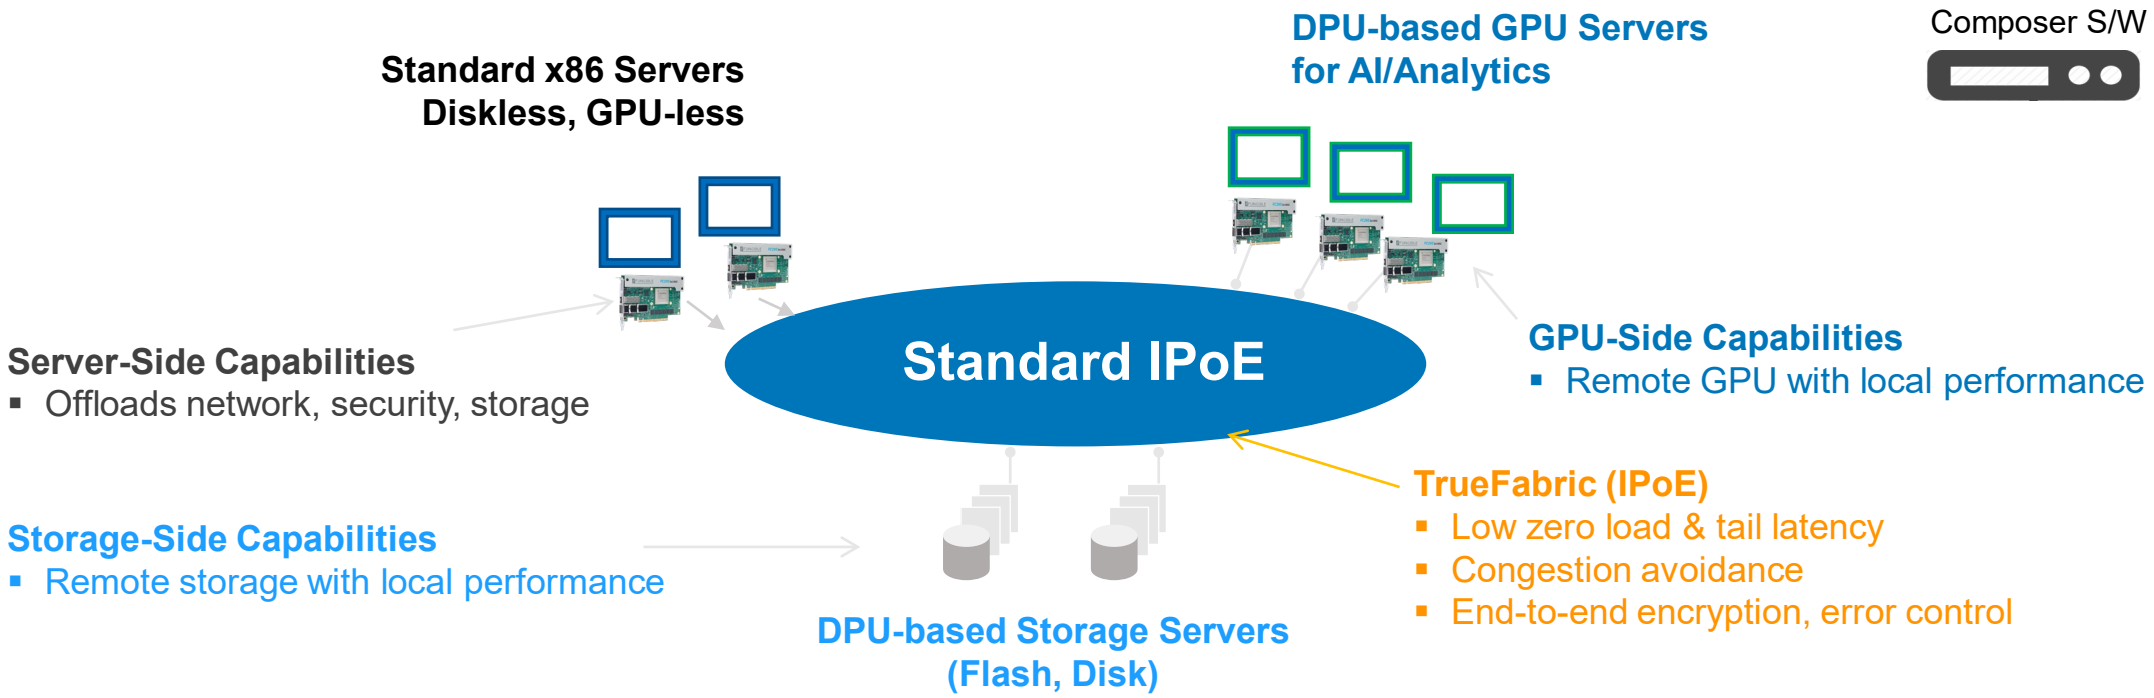


**3 key elements of the solution – DPU, penalty-free disaggregation, and composability**

 

# Fungible DPU-Based Next Generation Data Centers

## Addresses the 5 challenges of today's data centers with DPUs, penalty-free disaggregation and composability

LOW POWER
LOW SPACE

Disaggregation of storage, network and GPU resources enables higher utilization

Compression, low-cost EC based rack level resiliency

***Just-in-time composition*** of resources to exactly meet workload requirements with no wastage.

AGILITY

Redeploy resources across workloads ***in minutes*** to handle workload hot spots with composer software

SKUs

Reduce server SKUs to a minimal set through composability

SECURITY &
RELIABILITY

Independent ***hardware-accelerated security domains***, fine-grained segmentation, line rate encryption, rack-level resiliency

PERFORMANCE

High performance enabled by

(a) offloading data-centric tasks to the DPU-based cards

(b) high performance DPU-based storage

(c) high-perf networking with DPU-based TrueFabric™

 

# Summary

- Next Generation Data Centers will need
  - DPUs
    - To efficiently handle data-centric tasks
  - Penalty-free disaggregation
    - DPU based storage
    - DPU-based host cards
  - Composability software
    - To eliminate or reduce infrastructure silos

 

# THANK YOU

 

# Please take a moment to rate this session.

Your feedback is important to us.

 