

Storage Developer Conference September 22-23, 2020

Future Data Centers Will Be Hyperdisaggregated

Jai Menon Chief Scientist, Fungible September 22, 2020

Outline of Talk

SD₂₀

- Current data centers are inefficient for data-centric computations
- Future data centers will be hyperdisaggregated
- Fungible DPUTM and how it enables the next gen data center
- Fungible DPU-enabled storage fast, low cost, secure, reliable
- Summary

Data Center Inefficiencies Have Two Root Causes

SD₂₀

- 1. General-purpose processors are inefficient at data-centric processing
- 2. Data center networks are inadequate at scale



Data-centric = stateful, multiplexed processing of high b/w data streams; 25% today and growing Exponential growth in network traffic; amount of data; frequency & sophistication of cyber threats

Fungible Data Processing Unit[™] (DPU): Designed to Address the Two Root Cause Issues

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

- · Enables 10x more efficient execution of data-centric workloads
- Implements a scalable, low tail latency, congestion-free TrueFabric[™] endpoint



General-purpose application processing



Vector floating point operations, AI/ML workloads

Fungible DPU The third socket!

100 (1002) 11602	Thursday.			under ter geschapp m
				0 Ne 0 Ne
		-271 4(28) 		a and a constant of the second s
t onte Marya		- - - - - - - - - - - - - - - - - - -		
			an des Transforment	- Aler

Data-centric infrastructure workloads

- Standard external interfaces PCIe, Ethernet
- Programmable in C
- Tightly integrated multi-threaded cores & hardware accelerators for crypto, compression, EC



Fungible TrueFabric[™]



Complex, Multi-tier Network



- 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)

Simple, Flat Fabric



2020 Storage Developer Conference. © Fungible, Inc. All Rights Reserved.

SD@

TrueFabric Performance

SD²⁰

Traffic Pattern	1024 * (Node to Node)	1024 Node to 1024 Node	1023 Nodes to 1 Node
Fabric Utilization	90.7%	93%	90%
Latency Mean	1.84µs	2.10µs	1.71µs
Latency Variance	0.13µs	0.32µs	0.12µs
Latency P99	2.14µs	3.30µs	1.75µs

Network Configuration:

- 1,024 Nodes
- 200Gbps/Node
- Two-tier leaf-spine
- Leaf ZLL: 500ns
- Spine ZLL: 500ns
- iMix packet profile





- >2x better TCO, footprint and power
- Simpler and more reliable (fewer server types and network elements)
- More secure (pervasive encryption)



Fungible DPU

Fungible F1 DPU Architecture Targeted for Storage, Analytics, Security



8 Data Clusters

- 192 processor threads
- Full cache coherency
- Tightly integrated accelerators

SD₂₀

Control Cluster

- 8 processor threads
- Secure complex
 - HSM, Root-of-trust
- Work scheduler

Memory & I/O Interfaces

- 800Gbps Network Unit
- 4x16 G3/G4 PCIe Host Unit
- DDR4, HBM

On Chip Network

- High performance, low latency
- Any unit to any unit

Data Cluster

Runs the Data Plane on FunOS

6 Cores * 4 Threads

Multi-Threaded Accelerators

- Data movement
- Data lookup
- Data security (encryption)
- Data reduction (compression)
- Data protection (EC)
- Data analytics



SD@

Control Cluster

Runs the Control Plane on Linux

4 Cores * 2 Threads

Secure Enclave

- Secure boot
- Secure key vault
- Binary signing and authentication

Public Key Crypto Engines

- RSA
- Elliptic Curve

True Random Number Generator

Physical Unclonable Function



2020 Storage Developer Conference. © Fungible, Inc. All Rights Reserved.

SD@

CPUs, Caches, External Memory

CPUs

- MIPS-64, 9-stage, dual-issue, 4xSMT, FPU/SIMD unit
- · IPC on data-centric workload close to CPU-max
- Full hardware virtualization
- Large I+D L1\$, shared L2\$, full system-wide coherency

High Bandwidth HBM2 Memory

- 8GB, 4Tbits/sec
- Integrated in the package

High Capacity DDR4 Memory

- 2xDDR4 controllers, ECC enabled, up to 2666 MHz
- Up to 512GB
- Support of RDIMM, NVDIMM-N

• Fully general programmability

SD (20

- All code in ANSI-C
- Tight coupling with accelerators
- No performance compromises

Datapath Programming Model

SD@



MIPS-64 Hardware Threads Execute Run-To-Completion C-Code



Heterogeneous Accelerator Threads

Fungible DPU Software Co-designed with the hardware



2020 Storage Developer Conference. © Fungible, Inc. All Rights Reserved.

SD@



Fungible DPU-Based Storage

Cloud Storage Requirements

Scale-out

• 100s of storage nodes; protection from multiple whole node failures

Low cost

· Low cost reliability, line rate compression, storage pooling

High performance (IOPS) & Low latency (avg. and tail)

Multi-tenancy

Security

Block, KV, File

2020 Storage Developer Conference. © Fungible, Inc. All Rights Reserved.



SD₂₀

Building Cloud Scale-out Storage With Fungible DPUs

SD@

Customer Objective	Existing Solutions	Fungible DPU-based Solution
Performance • IOPs and Bandwidth • Low Latency (avg and tail)	General purpose CPU is the bottleneck	Fungible DPU enables high performance TrueFabric enables low latency Remote = Local Performance
Cost Optimization	 Data Reduction limited by CPU Pooling limited by CPU and network bottlenecks 	 Data Reduction @ line rate (compression) High-speed EC vs. replication Efficient pooling with Fungible DPU and TrueFabric
Security	Limited by the CPU	At line rate for in motion and at rest
Scale	Limited by the network	TrueFabric provides scale
Multi-tenancy	Limited support for QoS, diversity	Per volume QoS, compression, encryption keys

Performance Efficiency Percentage (PEP)

SDI 20

New Performance Metric for Scale Out Storage Systems Intuitive, easily measured, applies to all workloads @ cloud scale, unaffected by SSD capacity





Fungible DPU-based Block Storage Performance

SD₂₀

Note: Most storage systems are built with 2 processors so below #s should be doubled

	4K Read Perf per DPU
Ephemeral	8M IOPS PEP=95%
RAID/EC Protected	5M IOPS ¹
(SSD failure protection)	PEP=60%
Network Protected (RF=2)	6M IOPS ¹
(SSD and node failure protection)	PEP=71%

Assumes 12 SSDs

Unprotected read performance is SSD limited; PEP = 95%

Essentially unchanged performance with compression and encryption turned on Assumes TCP not TrueFabric; Performance will be higher with TrueFabric

¹Estimated

Compression: Fungible DPU LZMA vs. Others 10,000 Fungible DPU LZMA н. . . 30x faster & higher C/R 100x faster & C/R 95% of Brotli:9 1,000 Throughput (Mbps) x gzip -6 100 Brotli:9 X I 2MA F1-LZMA 10 ▲ G7IP × BROTLI X ZSTD BZIP2 2.5 3 3.5 45 5 4 Compression Ratio (=Input/Compressed)

SD@

Fungible DPU Enables End-to-End Storage Benefits

- Improved performance
 - Fewer x86 cycles consumed with storage virtualization
 - improved tail latency with True Fabric
- Enables diskless servers (remote boot) & bare metal services
- End to end compression to save network bandwidth
- End to end encryption
- Pool local and remote drives





Block, KV, File, Computational SD@

Fungible DPU-Based Storage is Low Cost



 Superior compression at line rate almost as good as the best compression in the world (Google's Brotli) but 100x faster SD₂₀

- Low overhead durability (25% vs. 200% overhead) - e.g.12+3 EC vs 3way replication
- Encryption without self-encrypting drives which are more expensive
- Very high performance allows supporting the needs of a large # of customers in small footprint

Fungible DPU-Based Storage is Secure



• Encrypt at line rate (800Gbps)

SD₂₀

- Per tenant keys can be stored in secure enclave inside DPU
- Every customer can get a different key(s)
- Only signed code can execute in DPU using secure boot protocol

Fungible DPUs Enable Diskless Servers

SD@



Value to Cloud Provider

- Fewer server types lower TCO
- Create in minutes versus hours agility
- Fast live migration continuous ops

Value to End Customer

- More choices
- Less wait time

Summary

20

Fundamental new building blocks: Fungible DPU and TrueFabric

- <u>https://www.fungible.com/wp-content/uploads/2020/08/WP0033.01.02020820-TrueFabric-A-</u> <u>Fundamental-Advance-to-the-State-of-the-Art-in-Data-Center-Networks.pdf</u>
- <u>https://www.fungible.com/wp-content/uploads/2020/08/WP0027.00.02020818-The-Fungible-DPU-A-New-Category-of-Microprocessor.pdf</u>

Fungible DPU enables the hyperdisaggregated data center

>2x TCO reduction; simpler and more reliable; very secure

Fungible DPU can be used to build differentiated scale out storage

- · Great performance, cost efficiency, reliability, security
- Block, KV, File

Please take a moment to rate this session.

Your feedback matters to us.