


FUNGIBLE F1 DATA PROCESSING UNIT

High-Performance, Fully Flexible Processor for Data-Centric Computing



KEY FEATURES

- Industry's first 800Gbps DPU
- Full offload of all infrastructure services from x86 processors
- High-level language programmable (i.e. C)
- Integrated 10x 100GE, 10x 40GE, 20x 50GE, 40x 25GE, 40x 10GE ports
- Separate 100M/1GE/10GE port for management Supports Fungible TrueFabric™ with non-drop and end-to-end QoS
- Support for NVMe-oF and NVMe over TCP
- Supports RDMA over TCP, TrueFabric™
- Integrated fully programmable L3/L4 IP router
- Supports overlay networks—NVGRE, VXLAN, GENEVE, MPLS, EVPN
- Latest generation MIPS64 cores—52 @ 1.6GHz
 - 200 hardware threads
 - Fully cache-coherent
- 64 PCIe Gen3/Gen4 lanes
 - 16 dual-mode controllers—each can be configured as endpoint or root complex
 - Support for PCIe SR-IOV with 64 PF/1024 VFs
- High-performance hardware accelerators
 - Programmable DMA engines—4Tbps
 - Crypto (AES-GCM/XTS)—1Tbps
 - Hash (SHA-3)—1Tbps
 - Compression/Decompression—512Gbps Full Duplex
 - Erasure coding/RAID—800Gbps
 - Regular expression (regex)—400Gbps
- Memory
 - Integrated 8GB HBM modules with ECC
 - Two-channel DDR4/NVDIMM with ECC
- Security
 - Secure boot and hardware Root of Trust
 - Secure Enclave and Key Vault
 - Public key authentication—signature 50K RSA 2K per second, 180K ECC per second
 - Physical unclonable function (PUF)
 - Line rate firewall/filtering
 - Deep packet inspection

BENEFITS

- TCO savings: Achieves 3x economic savings (network, compute, storage) across data center scales (mega scale to edge)
- Industry's highest performance processor for data-centric computing
 - I/O processor with 166MPPS
 - NVMe-oF controller @ 10MIOPS
 - L3-L7 security services @ 400Gbps
- Simplified server management with reduced server SKUs, enabled by disaggregation of compute and storage
- Ease-of-insertion—no changes to application software

OVERVIEW

Modern cloud-native applications are characterized by several key attributes: firstly, they are built as microservices that run on a distributed set of servers and secondly, they need to manage large datasets that extend beyond the capacity and performance of a single server. These attributes drive the need for high-performance, scale-out, hyperdisaggregated data centers, where servers interact as a unified pool of disaggregated compute and storage servers to serve the needs of these applications.

The Fungible DPU™ family of processors is purpose-built from the ground up to address the two biggest challenges in scale-out disaggregated data centers—inefficient execution of data-centric computations within server nodes and inefficient interchange of data among nodes. The Fungible DPU also strengthens the reliability, security and agility of these data centers. Note: Data-centric computations are computations performed in the network, storage, security and virtualization datapaths. Today, these computations represent more than 30% of computations in modern applications.

The Fungible F1 DPU is the flagship device of the Fungible DPU family of processors. Optimized for best-in-class performance, the F1 DPU delivers 800Gbps processing, executing data-centric computations an order of magnitude more efficiently than general-purpose CPUs. It fully implements the entire storage, networking, security and virtualization stack. The F1 DPU also facilitates highly efficient data interchange among server nodes through its TrueFabric™ technology. This enables disaggregation and pooling of all data center resources to be realized at massive scale. TrueFabric is a large-scale IP-over-Ethernet fabric protocol that provides full cross-sectional bandwidth with low average and tail latency, end-to-end QoS, and congestion-free connectivity. The TrueFabric protocol is fully standards-compliant and interoperable with TCP/IP over Ethernet. This ensures that the data center spine-leaf network can be built with standard off-the-shelf Ethernet switches.


F1 DPU ARCHITECTURE

The F1 DPU architecture leverages a unique hardware and software co-design that delivers maximum feature flexibility without compromising performance efficiencies for data-centric computing. This combination of attributes enables the F1 DPU to be designed for a multitude of use cases demanding high performance densities and low latencies such as storage, security, AI and analytics servers.

The F1's advanced SoC architecture integrates a large number of multi-core processors that are partitioned to run a separate control plane and data plane. The processors are interconnected through a fast network-on-chip (NoC) to a carefully selected collection of hardware accelerator blocks. The SoC interacts with external components through standard Ethernet network ports and PCIe Gen 3/Gen 4 controllers supporting Endpoint (EP) SR-IOV and Root Complex (RC) functionality.

APPLICATIONS

- Storage Target: Full storage stack offload optimized for NVMe-oF storage appliances
- Security Appliance: Full security stack (L3-L7) offload with built-in line rate firewall
- AI/ML: GPU disaggregation for scale-out AI/ML application
- Data Analytics: High-performance real-time big data analytics engines (OLAP, OLTP)
- Composable Disaggregated Infrastructure: Dynamic composability and resource pooling over the TrueFabric for CPUs, FPGAs, GPUs, SSDs, HDDs etc.




HYPERDISAGGREGATE YOUR INFRASTRUCTURE TO IMPROVE UTILIZATION AND REDUCE FOOTPRINT

The F1 DPU transforms compute and storage resources into disaggregated, pooled network resources that can be shared among many remote servers over a secure, low-latency TrueFabric. The F1 also implements bare-metal virtualization that offloads the entire hypervisor data path from the x86 processor for applications demanding bare-metal performance and latency characteristics. Further, thanks to the high-performance security capabilities of the F1 DPU, it is also well-suited for security appliances.

SCALE-OUT STORAGE


The F1 DPU software implements a complete, industry-leading storage stack that guarantees the durability and security of user data, both in motion and at rest. A comprehensive set of features—including high-performance in-line erasure coding, text/image compression and encryption—delivers 5x savings over existing solutions.

Designed to compress and encrypt data in a single pass, the F1 DPU virtually eliminates latencies inherent in traditional networked storage compression and encryption schemes, which involve several data transfers between an x86 CPU and multiple devices. The F1’s advanced storage services include raw block, durable block, and key-value store for NVMe-over Fabrics (NVMe-oF). The F1 DPU enables scale-out all-flash array (AFA) storage systems with extremely high IOPS (Read (4KB): 15MIOPS—limited only by PCIe bandwidth), without the need for x86 CPUs.



SCALE-OUT COMPUTE

The F1 DPU offloads the entire network, storage and security stacks from attached x86 CPUs, freeing up more than 50% of the x86 CPUs’ cycles, making them available for additional user workloads. A single F1 DPU can connect to either four dual-socket servers through PCIe Gen 3 x16 or eight dual-socket servers through PCIe Gen 3 x8. The same architecture can be used to attach any PCIe resource such as GPUs, TPUs and FPGAs.



SCALE-OUT NETWORK

The F1 DPU incorporates a fully featured, high-performance and low-latency L3/L4 IP router built around a state-of-the-art, flexible forwarding pipeline that can be reconfigured to support future protocols. The F1 also fully offloads transport and overlay protocols including TCP and UDP and overlay technologies such as VXLAN, NVGRE, GENEVE. Furthermore, the F1 DPU implements TrueFabric, a large-scale IP over Ethernet fabric protocol that provides full cross-sectional bandwidth with low tail latency, along with end-to-end QoS, non-drop and congestion-free connectivity. TrueFabric allows the collapsing of spine and core switching layers into a single layer, achieving higher link utilization and delivering the highest economic savings compared to existing technologies.

SCALE-OUT SECURITY

The F1 DPU is designed to operate as a security gateway that provides uncompromising and comprehensive protection for all data in the data center. With its advanced security capabilities—deep packet inspection, regular expression parsing, line rate hashing, and encryption—the F1 is ideally equipped to protect all traffic flows, providing in-line firewall services for all east-west and north-south traffic at full line rate. In doing so, the F1 DPU minimizes the attack surface area and enables real-time threat detection and prevention.

The F1 DPU also incorporates a hardware Root of Trust that ensures all software is cryptographically authenticated before running on the DPU. It also incorporates a Secure Enclave, which implements a key vault and provides unprecedented asymmetric cryptographic performance levels with its integrated public key (RSA and ECC) accelerators.


The F1 DPU enables end-to-end encryption of up to 100 million flows.

SOFTWARE

The Fungible DPU family of processors share the exact same programming model.

The F1 DPU runs FunOS™ on its data plane. FunOS is an innovative, purpose-built operating system written in high-level programming languages (ANSI C) for the data plane. FunOS runs the following stacks and features; networking, storage, security, virtualization, analytics.

The control plane runs a standard OS (e.g. Linux) and contains agents that allow a cluster of F1 DPUs to be managed, controlled and monitored by a set of REST APIs. These REST APIs can be integrated into standard or third-party orchestration systems such as Kubernetes CSI plugins, OpenStack, OpenShift, etc.



TECHNICAL FEATURES

HYPER-THREADED DATA PROCESSING UNIT

- Clean-slate architecture optimized for data center infrastructure services
- Programmable high-performance and high-efficiency processors
- 52x MIPS64 R6 cores with hardware virtualization
- High associativity multi-level cache hierarchy
- 32MB total on-chip L2 memory with ECC
- Advanced scheduling extracts maximum efficiency from 200 independent hardware threads
- Global resource manager and work orchestrator
- Uniquely scalable cache coherent memory system

HIGH-PERFORMANCE CLUSTERED CORES

- MIPS64 Release 6 Instruction Set and Privileged Resource Architecture
- 64KB L1 I-cache, 80KB L1 D-cache
- Full 64b architecture
- Dual-issue pipeline
- Branch and jump prediction
- L1 data cache supports cache-to-cache data transfers
- Virtualization module support
- Out-of-order data return
- 48-bit virtual and physical addresses
- IEEE 754 - 2008 compliant floating-point unit with SIMD engine
- Fast MMU and multi-level on-core TLB
- ECC and parity support on data caches and L1 instruction
- Load/store bonding support
- Low-power features

NETWORK-ON-CHIP

- Optimized high-bandwidth intelligent mesh
- Ultra-low latency messaging

PCIe INTERFACE

- 64 PCIe Gen 3 or 32 Gen 4 lanes (bifurcatable)
- 4x Gen 3x16 or 4x Gen 4x8
- 8x Gen 3x8 or 8x Gen 4x4
- 16x Gen 3x4 or 16x Gen 4x2
- 2.5G, 5G, 8G, 16G link rates
- 8 dual-mode controllers—each can be EP or RC
- Supports MSI/MSI-X
- Programmable QoS per VF
- SR-IOV support, with 64 PFs and 1024 VFs
 - 4 PFs per controller, up to 64 VFs per controller
 - Flexible resource allocation
 - Hardware QoS per virtual function
- Virtualized DMA support with QoS

NETWORKING

- Network bandwidth of 800Gbps
- 8x 100GE (IEEE802.3bj, IEEE 802.3bm)
- 8x 40GE (IEEE802.3ba)
- 8x 50GE (Ethernet Consortium)
- 32x 25GE (Ethernet Consortium)
- 32x 10GE (802.3ae)
- IEEE 802.3ap-based auto negotiation and backplane KR
- IEEE802.1Qau (QCN)—congestion notification, IEEE802.Qaz (ETS) and IEEE802.Qbb (PRC)
- IEEE1588v2
- Jumbo Frame Support (16KB)
- Integrated PHY supporting fiber and copper media
- Full hypervisor OVS offload
- P4-like programmable parser
- Fungible TrueFabric™
- All packets on wire can be encrypted (AES-GCM)

HIGH-PERFORMANCE ACCELERATION

- Programmable DMA engines—4Tbps
- Crypto (AES-GCM/XTS)—1Tbps
 - Various AES modes (CBC, XCBC, ECB, CTR, GCM, XTS) 128b, 192b, and 256b
- Hash (SHA-3)—1Tbps
 - SHA-1/SHA-2/SHA-3 (160, 224, 256, 384, 512)
- Compression—512Gbps
- Erasure Coding/RAID—800Gbps
- Programmable CRC, T10-DIX—1Tbps
- Analytics/DPI—400Gbps

MEMORY SUBSYSTEM

- Integrated HBM with a total of 8GB
- Two channels DDR4 w/ECC—up to 512GB each
- NVDIMM-N, RDIMM support

SECURE BOOT

- Dedicated secure boot processor—trusted key management and distribution, boots only validated certificates and firmware
- Root of trust for the server
- All DPU and x86 code is authenticated before execution
- Secure Enclave on device
- 50K RSA 2K signatures per second
- Physical unclonable function (PUF)

TRANSPORT & OVERLAY OFFLOAD

- RDMA over TCP, TrueFabric
- Complete TCP and UDP offload—stateless or stateful termination
- MPLS, VXLAN, NVGRE, GENEVE, EVPN

MANAGEMENT INTERFACES

- Orchestration interface for storage, compute, and networking
- 100M/1GE/10GE management interface
- IEEE 1588
 - Transparent clock, ordinary clock and slave support
 - One-step and two-step time-stamping with nanosecond granularity
 - PTP synchronized time reference distributed to all cores
 - Ultra-low jitter distribution within specialized internal fabric
- Four I2C interfaces (three masters and one slave)
- QSPI interface for flash
- eMMC (5.1) interface support
- General Purpose I/Os (GPIOs)
- Dual UART
- JTAG IEEE 1149.1 and IEEE 1149.6

POWER

- 120W–150W based on configuration

SOFTWARE DEVELOPMENT TOOLCHAIN

- Cross-compile GNU toolchain
- Data plane APIs—network, storage, security, data analytics

FUNOS DATA PLANE SOFTWARE

- Networking—underlay and overlay Ethernet, IPv4/6 routing, FungibleTrueFabric, routing/network segmentation, TCP/IP offload, RDMA, RPC offload, SSL offload, firewall
- Storage—raw block, durable block and key value store
- Storage—NVMe-oF and NVMe over TCP Security Stack—intrusion prevention, IPsec, transport layer security (TLS), URLF, logging, AV, ATP, DPI, regex
- Data Analytics—flexible pipeline for stream processing and columnar databases

LINUX CONTROL PLANE SOFTWARE

- Networking—underlay and overlay control plane
- Storage control plane
- Security control plane

LINUX SERVER SOFTWARE

- Linux netdev, RDMA Verbs and crypto device drivers
- eBPF offload
- OpenMPI support

ABOUT FUNGIBLE

Silicon Valley-based Fungible is reimagining the performance, economics, reliability and security of today's data centers.

CONTACT US

sales@fungible.com

Fungible, Inc.
3201 Scott Blvd | Santa Clara | CA 95054 | USA
669-292-5522

www.fungible.com | [in](#) [▶](#) [🐦](#) [✉](#)

