FUNGIBLE S1 DATA PROCESSING UNIT

Power and Form-Factor Optimized, High-Performance Processor for Data-Centric Computing

KEY FEATURES

- Industry's most flexible 200Gbps DPU
- Full offload of all infrastructure services from x86 processors
- High-level language programmable (i.e. C)
- Integrated 2x 100GE, 4x 50GE, 8x 25GE ports
- Separate 100M/1GE/10GE port for management
- Supports Fungible TrueFabric™ with non-drop and end-to-end QoS
- Supports NVMe, NVMe-oF and NVMe over TCP
- Supports RDMA over TCP, TrueFabric
- Supports up to 200G full TCP termination offload
- Supports overlay networks—NVGRE, VXLAN, GENEVE, MPLS, EVPN
- Supports 4x NVMe SSD local instance storage
- Thermal design power (TDP) of <35W for 100Gbps processing and <50W for 200Gbps processing
- Latest generation MIPS64 cores—16 @ 1.6GHz
  - 56 hardware threads
  - Fully cache coherent
- 32 PCIe Gen 3/Gen 4 lanes
  - 8 dual-mode controllers—each can be configured as endpoint or root complex
  - Supports PCIe SR-IOV with 32 PF/512 VFs
- High-performance hardware accelerators
  - Programmable DMA engines—2Tbps
  - Crypto (AES-GCM/CBC/XTS)—256Gbps
  - Hash (SHA-1/2/3)-256Gbps
  - Compression/Decompression—128Gbps Full Duplex
  - Erasure coding/RAID—200Gbps
  - Regular expression (regex)—100Gbps
- Security
  - Secure boot and hardware Root of Trust
  - Secure Enclave and Key Vault
  - Public key authentication
  - Physical unclonable function (PUF)
  - Line rate firewall/filtering
  - Deep packet inspection

BENEFITS

- TCO savings: Increased utilization from pooling of expensive resources and improved network utilization reduce need to over-provision.
- Bare metal performance for applications with high throughput and low latency requirements in virtualized environments
- High QoS enabled by TrueFabric, delivering low average latency and jitter, even at massive scale
- No-compromise end-to-end security
- Simplified server management with reduced server SKUs, enabled by disaggregation of compute and storage
- Ease-of-insertion—no changes to application software, support for standard APIs

OVERVIEW

The Fungible S1 DPU is the second device in the Fungible DPU™ family of purpose-built processors optimized for data-centric computing. Data-centric computing combines executing data-centric computations within server nodes and efficiently moving data among nodes. Data-centric computations are characterized by stateful processing of data streams at high rates. They are typified by networking, security and storage stacks, and today represent 30% of the modern server workloads.

The S1 DPU is the industry's most flexible 200Gbps DPU, capable of executing data-centric computations more efficiently than general-purpose CPUs. The S1 DPU fully offloads the entire storage, networking, security and virtualization stack from x86 cores in host servers, freeing up more than 50% of the x86 CPUs' cycles to run application workloads.

The S1 DPU also facilitates highly efficient data interchange among server nodes through its TrueFabric™ technology. TrueFabric is a large-scale IP-over-Ethernet fabric protocol that provides full network cross-sectional bandwidth with low average and tail latency, end-to-end QoS, congestion-free connectivity and security among server nodes. The TrueFabric™ protocol is fully standards-compliant and interoperable with TCP/IP over Ethernet, ensuring that the data center Spine-Leaf network can be built with standard off-the-shelf Ethernet switches.

By combining these key capabilities into a single solution, the Fungible DPU family of processors enables hyperdisaggregation and pooling of compute and storage resources - delivering a high-performance, massively scalable composable infrastructure for next generation data centers!

S1 DPU ARCHITECTURE

The Fungible DPU family of processors leverage the same hardware and software co-design and share the exact same programming model. However, while the F1 DPU is designed for high-performance standalone appliances such as storage, security, AI and analytics servers, the S1 DPU maximizes performance within the footprint and power envelope of a standard PCIe adapter.

The S1's advanced SoC architecture integrates clusters of multi-core processors that runs a cleanly separated control plane and data plane. These clusters 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.

The novel architecture of the Fungible DPU enables a unique combination of performance and flexibility - delivering a solution where the datapath is fully programmable, yet is fundamentally more efficient; achieving higher throughput, maintaining lower latency and consuming less power than existing solutions.

APPLICATIONS

- Bare metal virtualization, multi-tenancy with isolation: Offloading the entire hypervisor data path functionality into the S1 DPU
- Scale-out network: Offloading the full network stack - optimized for RDMA and RDMA over TrueFabric
- SDN/NFV: Offloading switching and routing, load balancer, security applications
- End-to-end network and application security: Offloading L3-L7 services
- Storage Initiator and Local Storage: Offloading network, storage and security services
- Data Analytics: Bringing analytics processing closer to data

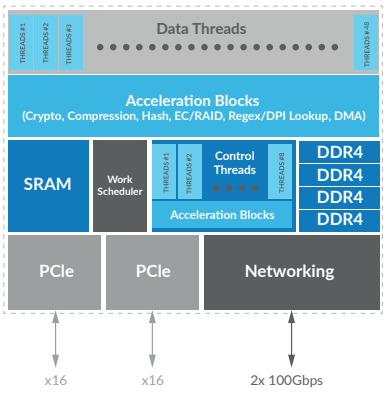

BUILD HIGHLY EFFICIENT, COST-OPTIMIZED, SECURE CLOUD, EDGE & ENTERPRISE DATA CENTERS

The broad capabilities of the Fungible S1 DPU enable various functions to be consolidated into a single PCIe adapter unifying network, security and storage services in an all-in-one platform!

BARE METAL VIRTUALIZATION

For virtualized applications requiring high performance and low latencies, direct access to underlying hardware resources is necessary. By offloading the hypervisor datapath into the S1 DPU, the hosts can maintain a lightweight hypervisor, thereby coming close to matching the bare metal performance characteristics of the hardware. This improved efficiency also allows more VMs or containers to be deployed on the same hardware.

The hypervisor datapath functions that can be offloaded into the S1 DPU include I/O virtualization, network services such as OVS, vRouter, network overlays, load balancer, storage services enabling local and remote storage, and security services.

SCALE-OUT NETWORK

The S1 DPU enables highly efficient data interchange among server nodes, scalable to thousands of server racks in the data center. It does so by fully offloading transport and overlay protocols, including TCP and UDP, as well as overlay technologies such as VXLAN, NVGRE and GENEVE.

The S1 DPU also implements standards-based RDMA, capable of leveraging Fungible's TrueFabric technology to deliver large-scale RDMA deployments. TrueFabric allows the collapsing of spine and core switching layers into a single layer, achieving higher link utilization (3x vs. ECMP) and delivering the highest economic savings compared to existing technologies.

SCALE-OUT NODE SECURITY

At the heart of the S1 DPU is a secure boot processor (SBP) that controls access permissions and implements a hardware-rooted chain of trust to ensure that all software running on the S1 DPU is cryptographically authenticated. The SBP thus safeguards against threats at all levels by making the S1 DPU a secure and trusted environment to run privileged infrastructure software. Such critical software is physically isolated from application software running on the host processor, thus shielding the system with another layer of security in hardware to thwart hardware-based attacks.

The S1 DPU also supports a plethora of L3-L7 network and security features to enable prevention, detection and response to potential threats in real-time. For example, the S1 DPU supports application-based policy management. This fine-grain access control allows the S1 DPU to minimize risks at the server by isolating and protecting each individual workload. Furthermore, the S1 DPU supports pervasive, end-to-end hardware-based authentication and encryption for network communications. It leverages cryptographic units for encryption and decrypting packets using industry-standard schemes including processing of IPsec, TLS payloads and associated Key Management System (KMS) functionality.

The S1 DPU also supports hardware-based deep packet inspection (DPI), facilitating dynamic monitoring and the logging of multiple layers of packets in packet streams. All these functions do not impose any host CPU processing tax.

SCALE-OUT STORAGE: STORAGE INITIATOR & LOCAL INSTANCE STORAGE

The S1 DPU can be integrated into host servers to complement the F1 DPU on the storage target. Serving as a bare metal NVMe-over-Fabrics storage initiator, the S1 DPU translates NVMe to NVMe-over-Fabrics, making NVMe remote storage appear local to the host. The S1 DPU can also execute data services such as end-to-end encryption for data security, and compression for more efficient network bandwidth utilization. This not only frees up expensive SSD space and precious CPU cores, but also boosts performance (IOPS) compared to what a server can achieve.

The S1 DPU also supports local instance storage and durable volume storage with up to 4x NVMe SSDs connected over PCIe. Local instance store volumes provide ephemeral data storage support, best suited for caching, scratchpad and certain storage applications.

DEPLOYMENT MODELS

The S1 DPU is designed with the flexibility to support various deployment models.

In-line and Look-Aside Mode: While the S1 DPU is designed to provide excellent in-line performance, it is also a perfect candidate for look-aside acceleration in compute nodes. Designed to fit into a PCIe plug-in card, the S1 DPU can offload the CPUs at compelling cost-to-performance and performance-to-watt ratios on functions such as encryption, asymmetric cryptography, compression (text, JPG), deep packet inspection and filtering.

Composable Disaggregated Infrastructure and Hyperconverged Infrastructure: While the S1 DPU is designed for high performance, large-scale composable disaggregated infrastructure (CDI), it can also supercharge today's hyperconverged infrastructures (HCI). In HCI servers, the CPU sits in the critical path, bottlenecking data flow from the network to expensive resources such as SSDs, GPUs and FPGAs. The S1 DPU removes the CPU from the critical path and liberates the stranded resources to be first-class citizens on the network.

SOFTWARE

The S1 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 S1 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.

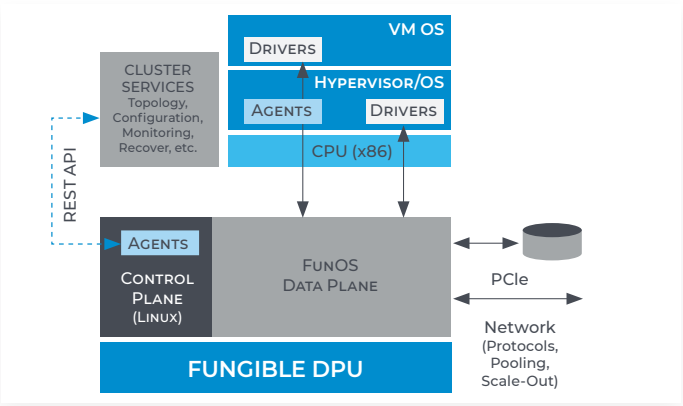

TECHNICAL FEATURES

Hyper-Threaded Data Processing Unit

- Clean-slate architecture optimized for data-centric infrastructure services
- Programmable high-performance and high efficiency processors
- 16x MIPS64 R6 cores with hardware virtualization
- High associativity multi-level cache hierarchy
- 52MB total on-chip SRAM memory with ECC
- Advanced scheduling extracts maximum efficiency from 56 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 RLB
- 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

- 32 PCIe Gen 3 or 16 Gen 4 lanes (bifurcatable)
- 2x Gen 3x16
- 4x Gen 3x8 or 4x Gen 4x8
- 8x Gen 3x4 or 8x Gen 4x4
- 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 32 PFs and 512 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 200Gbps
- 2x 100GE (IEEE802.3bj, IEEE802.3bm)
- 4x 40GE (IEEE802.3ba)
- 4x 50GE (Ethernet Consortium)
- 8x 25GE (Ethernet Consortium)
- 8x 10GE (802.3ae) IEEE 802.3ap based auto negotiation and backplane KR
- IEEE 802.1Qau (QCN)—congestion notification, IEEE802.Qaz (ETS) and IEEE802.Qbb (PRC)
- IEEE 802.1P VLAN tags and priority
- IEEE 1588v2
- 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—2Tbps
- Crypto (AES-GCM/XTS)—256Gbps
  - Various AES modes (CBC, XCBC, ECB, CTR, GCM, XTS) 128b, 192b and 256b
- Hash (SHA-3)—256Gbps
  - SHA-1/SHA-2/SHA-3 (160, 224, 256, 384, 512)
- Compression—128Gbps Full Duplex
- Erasure Coding/RAID—200Gbps
- Programmable CRC, T10-DIX—600Gbps
- Analytics/DPI—100Gbps

Memory Subsystem

- Four channels 32b DDR4 with ECC
- Up to 64GB

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
- 20K 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 timestamping with nanosecond granularity
  - PTP synchronized time reference distributed to all cores
  - Ultra-low jitter distribution within specialized internal fabric
- Two I2C interfaces (one master and one slave)
- One SMBus interface with MCTP support
- QSPI interface for Flash storage
- eMMC (5.1) interface support
- General Purpose I/O (GPIOs)
- Dual UART
- JTAG IEEE 1149.1 and IEEE 1149.6

Power

- 35W-50W 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, Fungible TrueFabric, routing/ network segmentation, TCP/IP offload, RDMA, RPC offload, SSL offload, firewall
- Storage—raw block, durable block and key-value store
- Storage—NVMe, NVMe-oF and NVMe over TCP
- Security—intrusion prevention, IPsec, 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