



MPPA® DPU Coolidge™

A new class of processors specialized in intelligent data processing for data center infrastructure, compute, AI acceleration and edge applications.

High-performance computing

- Up to 1.5 TFLOPs (SP)/192 GFLOPs (DP)

Power efficiency

- As low as 30W

AI acceleration

- Up to 25 TFLOPs (16 bits)/50 TOPs (8 bits)

High-speed I/O & interfaces

- Up to 18GB/s
- 12M IOPS
- Low 30 µs latency
- X2 100GbE

Real-time data processing

- Massive parallel processing
- 80 cores
- 6-issue VLIW
- Ultra-low latency

Massive parallel multitask processing

- Scalable 80-core DPU processor

Fully programmable

- Open standards: C/C++, Linux, RTOS, POSIX
- Kalray SDK based on standard tools & APIs for the development of new and portability of existing applications.

Security/safety

- Hardware root of trust
- Secure boot
- Accelerated cryptographic functions (optional)

The Coolidge™ 1 and 2 **Data Processing Units (DPU)** are third-generation processors based on Kalray's **Massively Parallel Processor Array (MPPA®)** architecture. Kalray DPUs are natively capable of managing multiple workloads in parallel with no bottlenecks to enable smarter, more efficient, and energy-wise data-intensive applications.

The DPU presents a compelling alternative compared to GPU, ASIC, and FPGA, contributing distinct advantages and adaptability across numerous data-heavy applications, ranging from data centers to edge or in embedded systems.

Next-generation data center infrastructure & 5G

Data center infrastructure chip for seamless integration onto PCIe Gen4 cards for use cases including I/O controllers, storage initiator and target controllers, high-speed network processing offload

- x86 offload or stand-alone "CPU-free" applications
- Compatible with containerized, virtualized, and bare metal infrastructures
- Fully programmable with dynamic distribution of resources across data and control & management planes

Fully programmable acceleration of high-performance protocols, services & QoS

- Enhanced support for NVMe-oF, RoCE/RDMA, TCP/IP, NVMe
- Intelligent load-balancing, priority-based flow control, and stateless L1-L4 parsing
- High-speed data protection services for clustered and fully distributed applications
- Line rate erasure coding (Reed-Solomon) per cluster
- Line-rate encryption/decryption/hash (IPSEC, TLS, XTS, MACsec)
- Acceleration for open RAN L1*
- AI functionality for insightful analytics and adaptive configuration

Acceleration of compute-intensive applications

Acceleration of complex workloads

- Innovative, patented core and co-processor enhancement for machine learning inference
- Advanced computation for computer vision
- Signal processing (e.g., FFT), cryptography, mathematics

Development of autonomous intelligent embedded systems

- Compatibility with multiple operating systems (Linux, RTOS)
- Support of 'freedom from interference' for mixed criticality

Enable next-gen edge computing systems

- Real-time analytics for automation, prediction, and control
- Seamless incorporation into existing systems

	Coolidge™ 1	Coolidge™ 2
Processor architecture	64-bit / 32-bit	
Cores	80	
Clock speed	Up to 1GHz	Up to 1.2GHz
Instruction level parallelism (ILP)	6-issue VLIW	
L1 cache	16KB instruction cache/ 16KB data cache	32KB instruction cache/ 32KB data cache
Floating point unit (FPU) standard	IEEE 754-2008 Reciprocal square root operations in floating single precision 64-bit integer multiplication (asymmetric cryptography) 4 execution rings	
Load/store bandwidth	256-bits per cycle	
Co-processors	80; 1 per core	
Applications	Acceleration of INT8, INT16 or FP16 accuracy	Acceleration of INT8, FP16 accuracy
	High-performance acceleration for 32-bit floating-point transcendental functions	
	Optimized matrix operations for deep learning and artificial intelligence	
MAC operations	Up to 128/cycle	Up to 256/cycle
System-on-chip (SoC)	5 clusters (total of 80 application cores + 5 management cores)	
Compute performance	Up to 640 GFLOPs (SP)/160 GFLOPs (DP)	Up to 1.5 TFLOPs (SP)/192 GFLOPs (DP)
	Up to 2.5 TFLOPs (16 bits)/20 TOPs (8bits) for deep learning	Up to 25 TFLOPs (16 bits)/50 TOPs (8bits)
Cluster	16 application cores + 1 management/security core	
L2 cache/TCM	4 MB, 512GB/s	8 MB, 600GB/s
PCIe interface	16-lane PCIe GEN4 endpoint (EP) or root complex (RC) Bifurcation up to 8 downstream ports in RC mode SR-IOV up to 8 physical functions/248 virtual functions Support for hot pluggable Up to 512 DMAs for multi queues/kernel bypass drivers NVMe emulation-virtualization	
Memory interface	64-bit DDR4/ LPDDR4-3200 channels with sideband/inline ECC	64-bit DDR4 -3200 channels with sideband/ inline ECC

Up to two ranks per DDR4 channel
2 DDR channels (up to 32GB) with channel interleaving

MPPA® DPU Coolidge™ Processor Block Diagram

Coolidge™ is composed of 5 clusters with 16 cores dedicated to applications, and 1 core each for management and security. 16 nm Finfet technology.

Need more performance?

Connect several MPPA® DPU processors together to reach the level of performance you need.

	Coolidge™ 1	Coolidge™ 2
Network interface	2x100GbE Ethernet	
	2x1/2x10/2x25/2x40/2x50/2x100 GbE	8x1/8x10/8x25/4x40/4x50/2x100 GbE
	Jumbo frame support (9.6KB)	
		Support for PTP/IEEE 1588v2
	Priority flow control (PFC), IEEE 802.1Qbb Checksum offload header & payload Line rate packet classification/smart load balancing Hash & round-robin based dispatch policy	
		Accelerated receive flow steering
	RDMA over converged Ethernet (RoCE) v1 and v2	
Security	Hardware root of trust Secure boot with authentication & encryption True random number generators (TRNG) RSA, Diffie-Hellman, DSA, ECC, EC-DSA and EC-DH acceleration AES-128/192/256 (ECB/CBC/ICM/CTR/GCM/GMAC/CCM) AES-XTS for storage applications MD5/SHA-1, SHA-2, SHA-3 Kazumi/Snow 3G, ZUC	
Cryptography accelerators (optional)		
Secure execution	Mixed criticality support Lockable critical configuration Memory and cache partitioning for non-interference & time-predictable execution Configurable L1 cache coherency	
Management/control interfaces		1GbE management interface/RGMII
	GPIOs/UARTs/SPI/I2C/CAN/PWM SSI controller for serial NOR Flash with optional boot SDCARD UHS-I / eMMC 4.51 memory controller 2x USB 2.0 OTG ULPI JTAG IEEE 1149.1 16-bit parallel trace interface	
Compression/decompression acceleration		FC1950 (zlib), RFC1951(deflate) and RFC1952 (gzip) @100Gbps
Low-density parity check (LDPC) encoder and decoder accelerator		Compliant with 3GPP TS 38.212, 5G NR FEC UL, 5G NR FEC DL

