Product Brief

Intel® Infrastructure Processing Unit SoC E2100

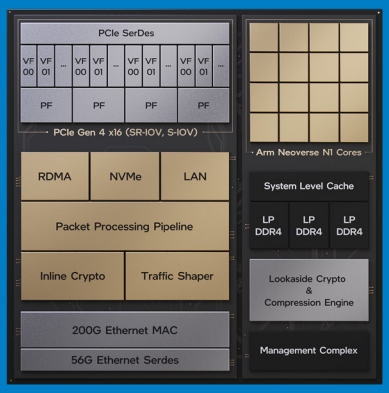

Key Features

Network Data Plane

- P4 programmable packet pipeline with QoS networking and telemetry up to 200Mpps
- Supports up to 4 host Intel® Xeon® processors with 200Gb/s full duplex
- NVMe device interface with inline AES-XTS and VM QoS
- RDMA with ROCEv2
- First hardware implementation of Falcon, a reliable low-latency hardware transport
- Inline crypto engine providing IPSec
- Lookaside Crypto and Compression Engine (LCE) for host or compute complex

Dedicated Compute Complex

- 16 ARM Neoverse N1 cores at up to 2.5 GHz
- Large system-level cache with 3 channels of LPDDR4x/DDR4 memory controllers
- Dedicated management processor

Overview

The Intel® Infrastructure Processing Unit (Intel® IPU) SoC E2100 is an infrastructure acceleration platform optimized for power, performance, and scale. It delivers feature-richness across various use cases including networking, storage and other infrastructure workloads.

The P4 programmable pipeline, and its paired P4 software suite, enable data plane acceleration for applications from virtualization to routing to ACLs. Advanced security accelerators, including inline and lookaside engines, enable wire-speed crypto and storage block encryption. With 16 high-performance Arm N1 Neoverse cores, a robust compute complex and integrated memory fabric, the IPU enables many application scenarios, from soft network pipelines to storage transport stacks.

The Intel® IPU SoC E2100 enables these key tenets of infrastructure acceleration:

- **Separation and isolation of infrastructure workloads:** Whether tenants in a cloud environment or application workloads in an edge or enterprise environment, IPUs optimize host CPU applications by removing the infrastructure overhead from traditional host-based network and storage infrastructure applications.
- **Infrastructure function acceleration:** Purpose-built accelerators for networking, storage, crypto, and telemetry provide performant hardware-based solutions, optimizing the infrastructure applications running on the IPU compute complex.
- **Virtualized storage:** IPUs enable the virtualization of storage resources, remote storage connectivity, and acceleration of storage-related computations. They provide data centers with greater flexibility to manage storage infrastructure, which is crucial for adapting to evolving data storage needs.

Uses Cases

Intel® Infrastructure Processing Unit Software Development Kit (Intel® IPU SDK) is a software stack that runs on the compute complex and the attached host. Developers can use the Intel® IPU SDK on the SoC to create targeted customer solutions:

- **Tenant hosting:** Separation of the compute complex from the host complex enhances security and leverages the inline crypto engine (ICE).
- **Accelerators as a Service:** Enables network-to-device memory data path and provides service abstraction for access to devices by implementing functions like QoS.
- **Appliance:** Performs packet processing either on a per-packet or per-flow basis. It can softly terminate packets and leverages a real-time packet processing engine to filter out suspicious packets.
- **Smart switch:** Multiple IPUs perform packet processing of select packets in the data center network.
- **Kubernetes acceleration:** Kubernetes facilitates containerized networking; the IPU SoC E2100 offers acceleration or offloading for the Kubernetes networking stack.

Packet Processing Pipeline

The flexible packet processor provides header processing service and enables data-plane use cases such as network virtualization, microservices, physical networking, and telemetry. Additionally, it supports legacy and advanced use cases for Cloud, Enterprise and Telco and consists of the packet processor and traffic shaper, which supports up to 200Mpps. The flexible packet processor supports P4 Programmable Pipeline with Inline IPsec, Hardware Connection Tracking, and Stateful ACLs, providing flexibility for defining and customizing the behavior of network data planes.

Compute Complex

The Intel IPU SoC E2100 compute complex has 16 Arm Neoverse N1 cores. These 16 high-frequency cores run up to 2.5 GHz and are backed by a large 32MB system-level cache and 3x dual-mode LPD-DR4x/DDR4 controllers for improved memory bandwidth. These give the IPU the bandwidth and horsepower to take on large infrastructure workloads.

The compute complex is tightly coupled with the network subsystem, allowing the network subsystem accelerators to use the system-level cache as a last-level cache. These high-bandwidth, low-latency connections enable a combination of hardware and software packet processing.

In addition, the management processor provides an interface to the server platform and an orchestration layer supporting robust system manageability and security configurations.

Accelerators

One of the pivotal movements for large Cloud Service Providers is the shift from legacy network standards like TCP to modern protocols that deliver data with better latency, predictability, and enhanced performance. The Intel IPU SoC E2100 supports multiple hardened accelerators that enable the high performance, low latency, and better efficiency required in the new generation of data centers.

With Intel's security-first mindset, the E2100 also provides features like root-of-trust, secure boot, and inline IPSec to secure every single packet being sent across the network and millions of connections.

NVMe Engine

Intel's high-performance NVMe engine enables storage offload with built-in support for encryption and additional options such as compression using the crypto and compression accelerator. Combined with a robust software stack, this engine can achieve up to 6M IOPs per direction. The NVMe engine exposes high-performance PCI Express (PCIe) NVMe devices to the host processor. It enables infrastructure providers to implement their storage protocol of choice using the cores on the IPU.

Falcon, a Reliable Low-Latency Hardware Transport

As modern data centers scale, providing scalable and predictable performance is a key challenge. The Intel IPU SoC E2100 features a next-generation reliable transport protocol, which addresses the long tail latency problems that exist with a lossy network which is critical for applications like storage, HPC, AI training clusters, and developing microservice deployment models.

Falcon is a hardware-assisted transport layer designed for high-performance, low-latency connections in Ethernet data center networks. Falcon enables deployments into existing networking infrastructure easily and is designed to support

backward compatibility with applications running on top of RDMA and NVMe interfaces. CPU utilization is minimized when NVMe/Falcon hardware transport is used to deliver high bandwidth storage with ultra-low latency.

Additionally, Falcon supports both in-order connection and out-order connection. It is a unified protocol that addresses the needs of new and demanding workloads that require high-burst bandwidth, high Op rate, and low latency. It supports usages including AI/ML, HPC, storage networking, and RDMA in cloud scale. Falcon achieves high performance, isolation, and efficiency at scale for high-bandwidth, low-latency workloads in today's Ethernet data center networks.

- High performance through low-tail latency: Hardware-assisted, delay-based congestion control, selective ACKs for fast loss recovery, and multipath-capable connections.
- Isolation and visibility in shared infrastructure: Includes µs-granularity per-flow traffic shaper and fine-grained stats for debuggability and software-defined network control.
- Efficiency and security: Implemented in hardware for low latency, high-op rate using industry-standard interfaces and PSP encryption.

The Falcon protocol consists of three layers. The upper layer protocol (ULP) mapping maps ULP operations such as messages and flow control into Falcon connections. The middle layer is the transaction layer. This layer exposes a request-response transaction interface to the ULPs. The bottom layer is the packet delivery layer that handles the network related functions, for example, congestion control, reliability, and multipathing.

RDMA with RoCEv2 Protocol
The RDMA with RoCEv2 protocol engine is a host-offload, kernel-bypass technology that enables direct memory-to-memory data communication between applications over a network. This engine supports up to 200Gb/s throughput, 150M messages per second, and up to 1M queue pairs (QP). It provides high-throughput and low-latency performance for high-speed Ethernet by eliminating three major sources of networking overhead: TCP/IP networking stack processing, memory copies, and application context switches. RoCEv2 operates on top of UDP/IP and is routable over IP networks.

Inline Crypto Engine
The Inline Crypto Engine (ICE) processes packets in both ingress and egress directions that require infrastructure cryptography offload in the Intel IPU SoC E2100. The ICE is part of the network subsystem pipeline along with the packet processor and supports IPSec full offload in the Intel IPU SoC E2100, supporting these scenarios:

- Security protocol processing for data in transit in the data center.
- Acceleration of encryption and decryption for select security protocols.
- Protocol processing security services:
  - Packet integrity authentication.
  - Anti-replay protection.
  - Security protocol processing that includes packet modifications and error checking.

Lookaside Crypto and Compression Engine
The LCE provides lookaside cryptography and compression services for data at rest and in transit. The LCE engine improves throughput on big data, file systems, and databases by reducing data over-head by compression. It concurrently supports up to 100Gbps compression plus 100Gbps decompression. When data is in transit, the LCE ensures that the data is intact using authenticity, integrity, and privacy protocols (for example, TLS, DTLS, QUIC, IPSec, PSP). When data is at rest, the LCE provides device-level encryption, storage-level encryption, and data protection on disk.

Complete Software Framework

The Intel IPU SoC E2100 promotes better collaboration, open-source compatibility, and developer-friendly features through a complete software stack. Notable features include:

- Complete P4 stack including toolchain, debuggers, analyzers and application-facing libraries enabling all of the E2100's packet-processing capabilities.
- Storage offload support through industry standard SPDK plugins, enabling NVMe-over-TCP and other storage transport protocols.
- Packet I/O support through both standard kernel interfaces and DPDK utilizing the Infrastructure Datapath Function standard driver.

Intel also offers open-source implementations of well-known networking and storage cases through the Infrastructure Programmer Development Kit (IPDK) under the Open Programmable Infrastructure Project, enabling solution development through validated, performance-tested recipes.

Host Interface

Intel Xeon processors with 200Gb/s full duplex

PCI-SIG PCIe Base Specification Revision 4.0, Version 1.0

16 PCIe lanes

PCIe Interface

PCIe 4.0 x16

MSI/MSI-X support

Ethernet Interface

Ethernet Port Logic: 4 ports up to 200Gb throughput

8 lanes of 56G SerDes, PAM4/ NRZ Support

1588 PTP timestamping (RX) and timestamp measurement (TX)

Up to 4 ports supported by single Ethernet Port Mac (EPM)

I/O Virtualization

12K Flexible Host Queues and Transmit Rate Limiters

Virtio-net, virtio-blk, vmxnet3, custom virtual devices

DRAM Interface

Dual-mode, three-channel interface supporting LPDDR4x memory using LPDDR4x-4267

LPDDR4x total bandwidth: 48GB

LAN Performance

Up to 200Gbps line rate bi-directional throughput

Up to 200Mpps bi-directional packet rate

Packet Processing Pipeline

High-performance Flexible Packet Processor operating on a 256-byte header stream in both directions

Programmable parser supports 1K parse graph nodes/256B parse depth

P4 Programmable Pipeline w/ Inline IPsec, Hardware Connection Tracking and Stateful ACLs

Programmable Parsing, Multi-stage Match-Action, Mirroring, Multicast, Modification and Recirculation

Parsing engine supports: Up to 1M LPM Routes, up to 16M Exact Match Entries, 1M Meters/ Policers/Shapers, TCAM and range tables

NVMe Performance

Up to 200Gbps line rate bi-directional throughput

Hardware paths support up to 6M 4KB read/write IOPS simultaneously

Reliable Transport

Falcon technology reliable transport

RDMA with RoCEv2 standard libraries

200Gbps line rate bi-directional throughput

Supports 6 physical functions (4x Intel Xeon processors, 1x Arm compute complex, and 1x integrated management complex)

Compute Complex

Up to 16 Neoverse N1 Arm cores at up to 2.5 GHz with 64KB L1 cache and 512KB L2 cache per core

32MB Distributed L3 Cache; up to 2.0 GHz

Coherent Mesh Network interconnect with 32MB System Level Cache (SLC)

Memory Subsystem consisting of LPDDR4x/DDR4 controllers for 3 channels

Time Synchronization

IEEE 1588 PTP (Tx and Rx time sync sampling flows) internal latency measurement (Egress/Ingress)

Traffic shaping

NVMe maximum command latency

Security and Crypto

Inline IPSec ESP, AES-GCM 128/256 engine

Lookaside Cryptography and Compression Engine (LCE)
- Support for chained operations
- 200Gb Bulk Crypto per direction, including TLS offload

Internal/external RoT, Secure Boot, Secure Debug TRNG via management complex

Meets security standard SP800-193

Product Information

For information about Intel® IPUs visit: intel.com/ipu

Learn more about IPDK and Open Programmable Infrastructure at ipdk.io and opiproject.org

0424/ED/Axiom 818147-002US