leidos

HIGH PERFORMANCE WORKLOAD SOLUTIONS

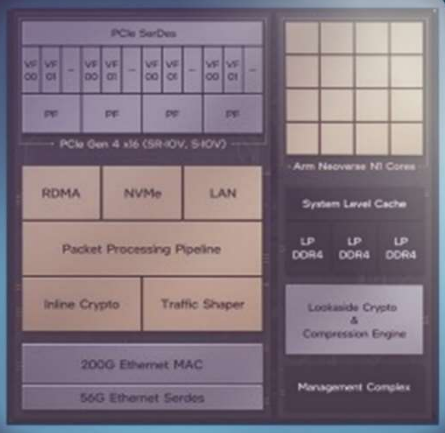

Deep Packet Inspection with the Intel® Infrastructure Processing Unit

SECURE HIGH-PERFORMANCE WORKLOAD SOLUTIONS FOR DEEP PACKET INSPECTION

In today's heavily contested cyber environment, incessant threats demand continuous cybersecurity innovation. Cyberattacks involving information breaches, policy infractions, network surveillance, malware and other dangers are ever-present. Faced with rising frequency and complexity of these attacks, tech leaders and network administrators depend increasingly upon deep packet inspection (DPI) to fortify network security, protect users, and safeguard cloud systems.

DPI is Data-Intensive

Because DPI is so data-intensive, special hardware is needed to accelerate workloads, provide higher performance, and deliver a lower total cost of ownership compared to traditional CPU-based processing. This specialized hardware must (1) identify data of interest without the overhead of post-processing; (2) parse through data faster than a CPU alone can; and (3) offload relevant data for post-processing.

The Intel® Infrastructure Processing Unit (Intel® IPU) Adapter E2100 is a programmable 200GbE, PCIe networking device that serves as an infrastructure acceleration platform optimized for power, performance, scale, and total cost of ownership. The adapter manages system-level infrastructure, allowing cloud operators to shift to a fully virtualized storage and network architecture while maintaining high performance, high predictability, and a high degree of control.

Full utilization of the adapter enables the Flexible Packet Processor (FPP) to efficiently identify packet importance for better workload benchmarks, optimized and accelerated network infrastructure, and maximized CPU utilization efficiency.

Deep Packet Inspection Use Cases

DPI workloads are vital to network security—they examine the data and header of the packet, determine its data path, and decide what actions the network should take on the packet. Many organizations prefer DPI workloads over conventional network packet filtering techniques, as conventional techniques do not provide the level of granularity that that is needed to eradicate advanced security threats. DPI workloads enable and enhance a range of network services, including firewalls, load balancing, quality-of-service (QoS), packet encapsulation/ decapsulation termination, and application monitoring.

- **Firewalls:** DPI provides deeper visibility into network traffic to support threat detection and prevention, enforce granular security policies, and optimize network performance.
- **Load Balancing:** Deeper visibility into network traffic enables more intelligent traffic distribution decisions, optimized resource utilization, improved application performance, and enhanced security posture.

- **Packet Encapsulation/Decapsulation Termination:** DPI can enforce and enhance security measures on the inner packet, mitigate risks associated with malicious activities, and ensure integrity of data transmission.
- **Application Monitoring:** DPI can provide deep insight into application performance, security, compliance and behavior, enabling optimization of application performance and enhancement of security.

FULL SYSTEM INTEGRATION WITH LEIDOS

Leidos creates high-performance workload solutions using technology available through our partnership with Intel. To accelerate compute-intensive DPI workloads we developed a fully integrated solution that leverages the 200 GbE Intel IPU Adapter and a complete software stack, all configured on the Dell PowerEdge R760 Server with Intel® Xeon® processors. Additionally, the integration of Red Hat OpenShift with Intel IPUs enables organizations to maximize performance, scalability, and efficiency by providing a layered approach to container and Kubernetes security. This holistic solution improves network performance and load balancing with QoS, enhances security features and performance, and adds versatility for the most demanding applications.

Example Use Case

ETLT (extract-transform-load-transform) is a term for ingest pipelines where data and enrichment and transformation are performed upon ingestion (extraction) and also after the data are loaded into the pipeline's datastore for further processing. In the example and illustrations below, two hosts ingesting sustained high data rate traffic flows are compared and contrasted:

Current State (Host A). Host A ingests the traffic through normal high speed network interfaces. The packet processing is performed on the host operating system and competes for host system/operating system resources (CPU, Memory, System DataBus between components, etc.). This processing/resource competition can result in disruption of stability and throughput; starvation of the critical processes of required resources can result in dropped packets and compromise of mission.

Attainable State (Host B). Host B uses Intel IPUs to offload the packet processing and related functions (flow processing, deep packet inspection, etc.). Offloading removes significant amounts of processing/resource competition from the system. Throughput, fault tolerance, and system stability increase, while the risk of dropped packets is reduced. The packet processing operations are isolated, with their own optimized resources separate from the host. Isolating the host from the data flow increases system capacity, reduces the systems needed to fulfill mission requirements, and allows the host to operate more efficiently and reliably.

Figure 1 on the next page illustrates (1) the collection of data from multiple sources in the form of telemetry; communications; Command, Control, Communications, Computers/Intelligence, Surveillance and Reconnaissance (C4ISR); operational technology (powerplant, electrical, etc.); and defensive cyber operations, (2) how that data would be managed in the current state (Host A) and the attainable state (Host B), and (3) the relative advantages of full-system integration using the 200 GbE Intel IPU Adapter.

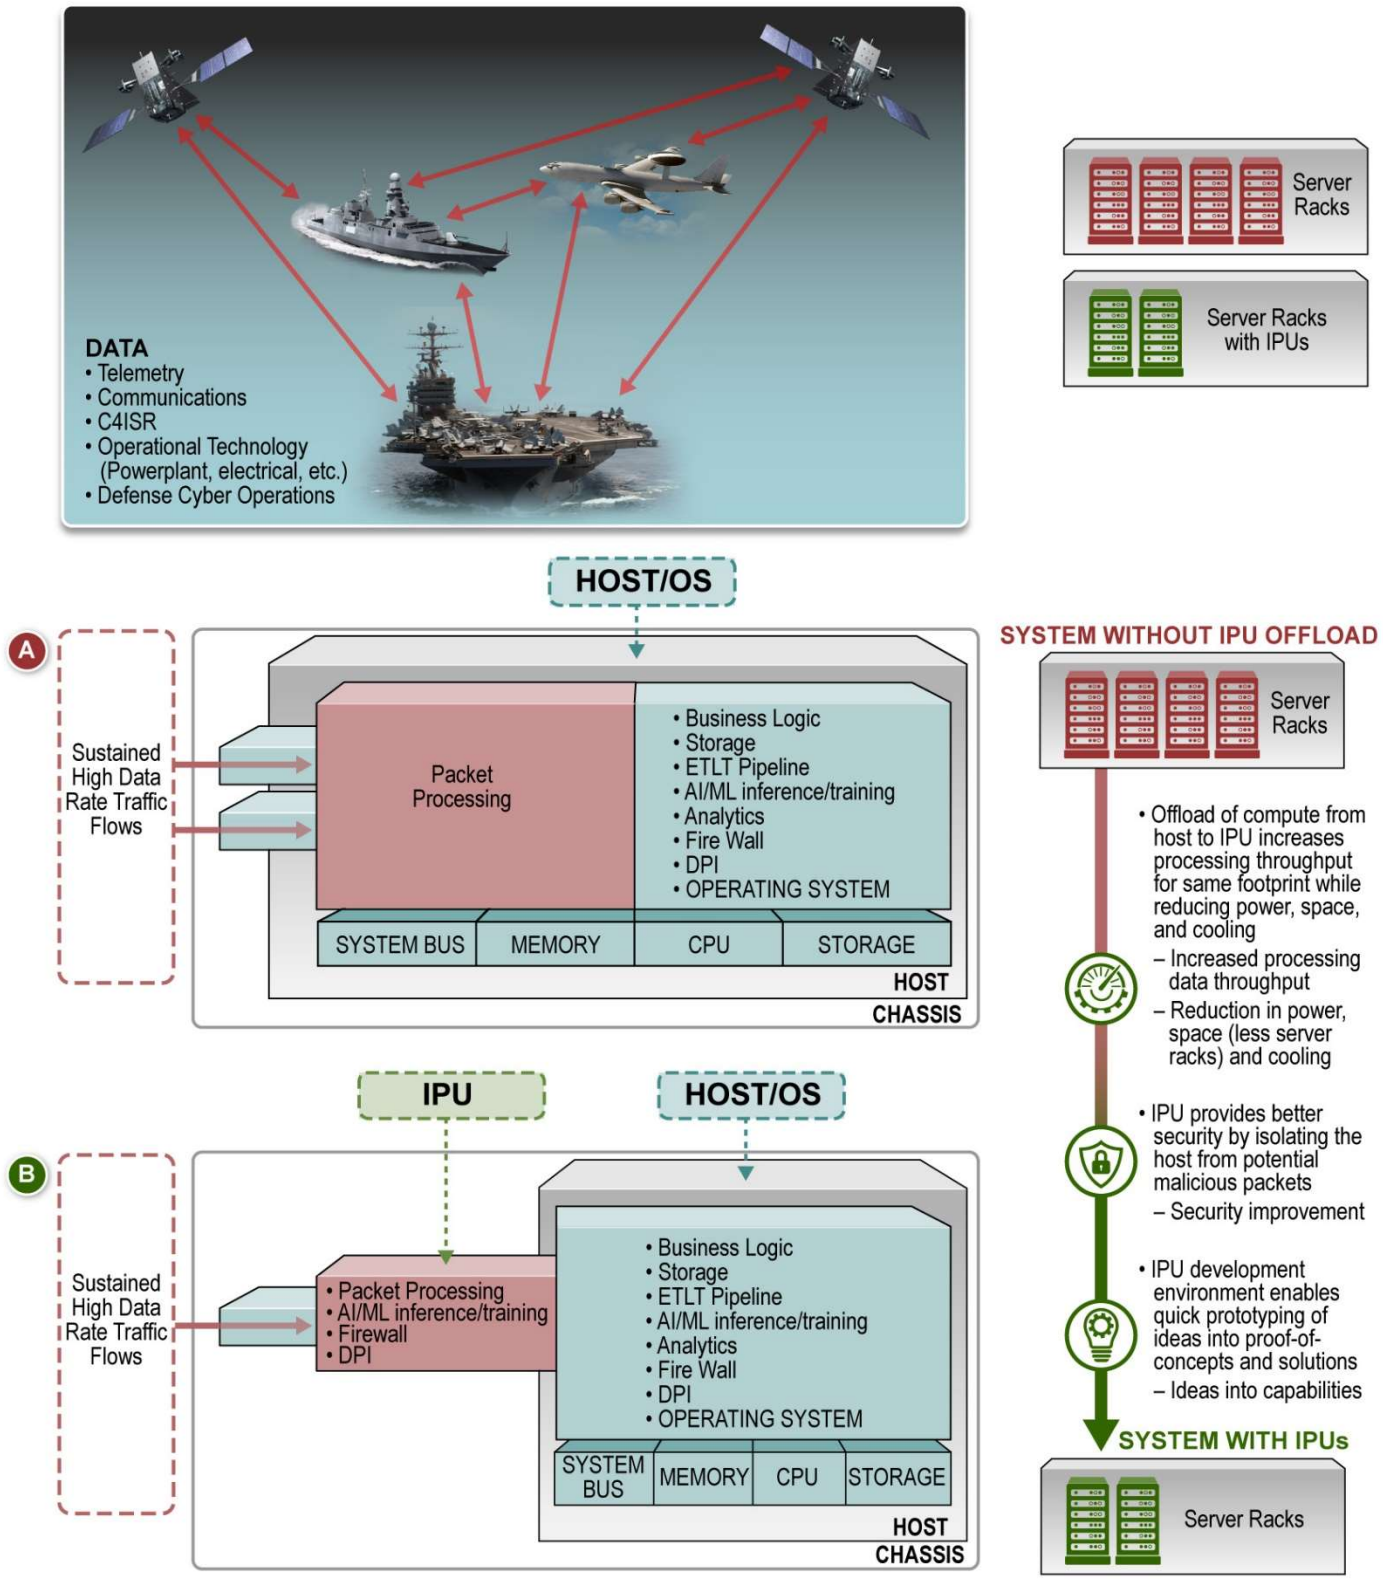

Figure 1: Example Use Case

The Intel IPU In Action

The Intel IPU Adapter features the Intel IPU SoC E2100 and was developed to address the needs of data-intensive applications requiring scalability, versatility and high-performance. The adapter enables infrastructure acceleration, efficient packet processing and improved data center security, as well as complete flexibility and programmability with the P4 programmable pipeline.

The Intel IPU SoC E2100

The Intel IPU SoC is an ASIC-based hardware accelerator designed to accelerate infrastructure workloads and optimize data center operations. It features dedicated programmable cores and Ethernet features, offering a host of benefits for both multi-tenant and bare metal cloud infrastructure as well as for Enterprise and Telco networks. Key tenets of infrastructure acceleration include:

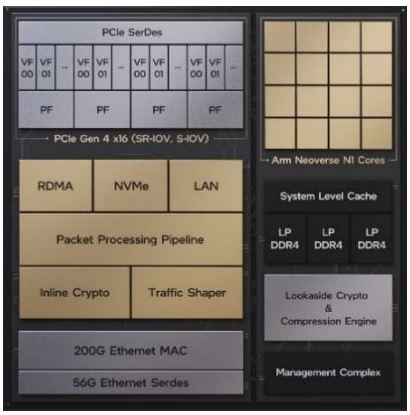

- **Separation and isolation of infrastructure workloads.** The Intel IPU SoC optimizes host CPU applications by removing the infrastructure overhead from traditional host-based network and storage infrastructure applications.
- **Infrastructure function acceleration.** Purpose-built accelerators for networking, storage, crypto, and telemetry provide performant hardware-based solutions, optimizing the infrastructure applications running on the Intel IPU SoC compute complex.
- **Virtualized storage.** The Intel IPU SoC enables the virtualization of storage resources; remote storage connectivity; and acceleration of storage-related computations. These capabilities provide data centers greater flexibility to manage storage infrastructure, which is crucial for adapting to evolving data storage needs.

The Packet Processor Complex

The packet processing complex is a highly performant subsystem designed to securely handle data-intensive workloads like DPI. The subsystem receives, transmits, and processes data packets through three engines: the FPP, the Inline Crypto Engine (ICE), and the traffic shaper. The FPP provides header processing services and enables data-plane use cases such as network virtualization, microservices, physical networking, and telemetry. Advanced use cases for Cloud and Enterprise are supported.

The FPP 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. It operates on a 256-byte header stream in both directions. The P4 Programmable Packet Pipeline supports QoS networking and telemetry with up to 200Mpps, critical for DPI workloads. The pipeline and its paired P4 software suite enable data plane acceleration for applications ranging from virtualization to routing and access control lists. Advanced security accelerators, including inline and lookaside engines, enable wire-speed crypto and storage block encryption.

The ICE processes packets bi-directionally (ingress and egress) for required infrastructure cryptography offload in the Intel IPU. The ICE is part of the network subsystem pipeline that, along with the packet processor, supports IPSec full offload in the adapter. Functions supported include:

- 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, and security protocol processing including packet modifications and error checking).

The Compute Complex

The Intel IPU SoC compute complex has 16 Arm Neoverse N1 cores. These high-frequency cores are backed by a large 32 MB system-level cache and 3x dual-mode LPDDR4x/DDR4 controllers for improved memory bandwidth. With 16 Arm N1 Neoverse cores providing a robust compute complex and integrated memory fabric, the Intel IPU SoC extends capabilities of data path processing in the FPP beyond 256-bytes and enables a plethora of application scenarios such as P4 extraction of traffic flow data, deploying machine learning algorithms and data analytics.

Complete Software Framework

Intel demonstrates a robust commitment to open-source infrastructure software by promoting collaboration and innovation within the development community. Intel IPUs offer a wealth of developer-friendly features within its comprehensive software stack:

- Complete programmable P4 stack includes toolchain, debuggers, analyzers, and application-facing libraries enabling all the Intel IPU 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.
- Seamless integration with Red Hat OpenShift as an inboxed solution, optimizing resource utilization, efficiency, and security.
- Supports IDPK – an open-source vendor-agnostic framework for infrastructure offload and management and enables network virtualization, storage virtualization, root-of-trust and other offload capabilities.

Finally, Intel's commitment to open source is underscored by its provision of open-source implementations for well-known networking and storage cases under the Open Programmable Infrastructure Project.

Dell PowerEdge R760 Server with Intel Xeon® Processors

The Dell PowerEdge R760 with Intel Xeon processors server provides performance and versatility for compute-intensive workloads. It is integrated with the Dell iDRAC manageability board management controller which among other things provides closed loop thermal control of the Intel® IPU.

TO LEARN MORE

DPI provides unparalleled insight into network traffic and enabling robust security measures against evolving threats. The compute-intensive nature of DPI necessitates specialized hardware to accelerate workloads, ensure high performance, and reduce TCO. Utilizing the Intel IPU Adapter E2100 presents a transformative solution for DPI, by offering the performance, flexibility, and scalability needed to fortify network defenses and ensure robust security measures.

Are you ready to unlock the power of Intel IPU Adapters for compute-intensive workloads like DPI? Contact your local Leidos representative for more details and visit Intel.com/IPU today to learn how the Intel IPU Adapter E2100 is the optimal choice for infrastructure processing.

LEIDOS REPRESENTATIVES

Colin Formulak	**(443) 367-7457**	**colin.formulak@leidos.com**
John Gorman	**(443) 367-7272**	**john.p.gorman@leidos.com**
Alan Carroll	**(410) 312-2000**	**alan.carroll@leidos.com**