Solution Brief

Intel® NetSec Accelerator Reference Design

Reference Design Enables Deploying Critical Network and Security Functions Inline

The Intel® NetSec Accelerator Reference Design is a blueprint to commercialize an Intel® architecture-based PCIe add-in card for processor-intensive network security workloads. The card features all of the functionality of a server with the capability to support full orchestration and management capabilities. It is ideal for security workloads such as IPsec, SSL/TLS, firewall, SASE, analytics and inferencing for network security. This reference design can help improve performance, scale and efficiency for customers from edge to cloud.

With the ongoing transformation toward edge computing, the increase of devices and employees connecting from anywhere at any time is making enterprise environments more distributed than ever before. Traditional perimeter-focused security models and fixed deployment models no longer apply. Monolithic applications have been replaced by chains of containerized microservices that traverse on-premises and cloud infrastructure, decoupled from the underlying hardware; workloads need to be deployed where they are needed. These dynamic, software-defined environments require new approaches to apply security functions at the per-workload, per-user and per-device level.

The secure access service edge (SASE) model meets these new distributed security requirements by converging software-defined security and wide-area network (WAN) functions into a cloud-delivered set of services. The virtualized or containerized services enhance efficiency with centralized orchestration and reduce equipment costs by using cloud infrastructure based on commercial off the shelf (COTS) servers in place of legacy, single-purpose hardware.

Most SASE solutions are fully integrated stacks of network and security functions, with licensing models based on enabling specific components. SASE vendors make massive investments to evaluate, obtain and integrate them. SASE software functions require performance and stability in a server that shares multiple compute-intensive workloads and multiple tenants. Integrating a performant solution of software-defined WAN (SD-WAN) together with a security stack that includes NGFW, ZTNA, CASB, SWG, DLP and more is particularly challenging.

A successful SASE solution is comprised of geographically dispersed deployments of edges or points of presence, and such edges have unique challenges not always found in traditional on-premises data centers. Edge infrastructure may have stringent space, thermal and power constraints, and it may be deployed in isolated locations that require zero-touch management. The edge infrastructure may need to be augmented at speed to meet the growing demand for SASE services. Such requirements found at the edge for SASE may be common to many other workloads deployed at the edge.

The Intel® NetSec Accelerator Reference Design offers an alternative approach to SASE and the isolation of various edge workload functions that can dramatically reduce the infrastructure footprint for network and security workloads. It provides users the full functionality of a server on a PCIe card, including an Intel processor, Intel® Ethernet 800 Series Controller and substantial onboard high-speed memory.

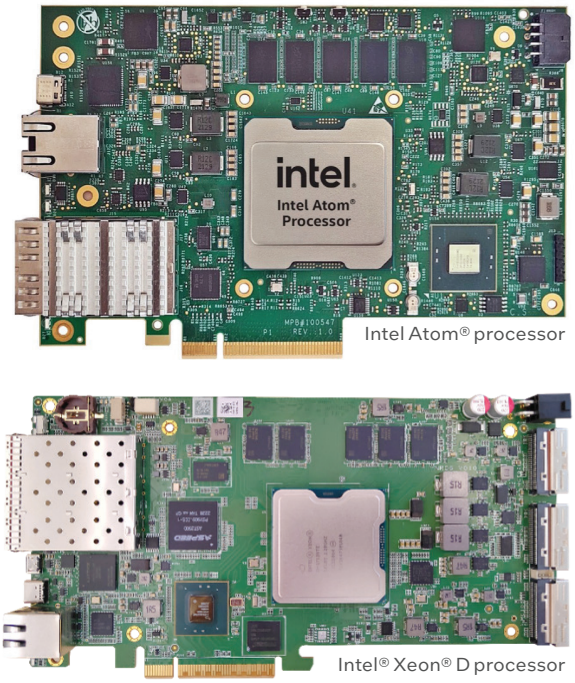

Figure 1. Intel® NetSec Accelerator Reference Design.

Expanded processing horsepower for security workloads

The Intel NetSec Accelerator Reference Design easily handles the compute capacity of network and security appliances, providing one or more separate physical execution environments. The accelerator supplements the server's primary processor with dedicated acceleration hardware for network and security functions.

Instruction set compatibility and a shared driver architecture between the host-server CPU and the CPU on the Intel NetSec Accelerator Reference Design help make the overall solution seamless using standards-based Intel architecture. The consistency of programming models enables platform commonality between the Intel processor on the accelerator and Intel® Xeon® Scalable processors or Intel Xeon D processors in the host machine.

The Intel® processor at the heart of the Intel® NetSec Accelerator Reference Design enables inline IPsec.

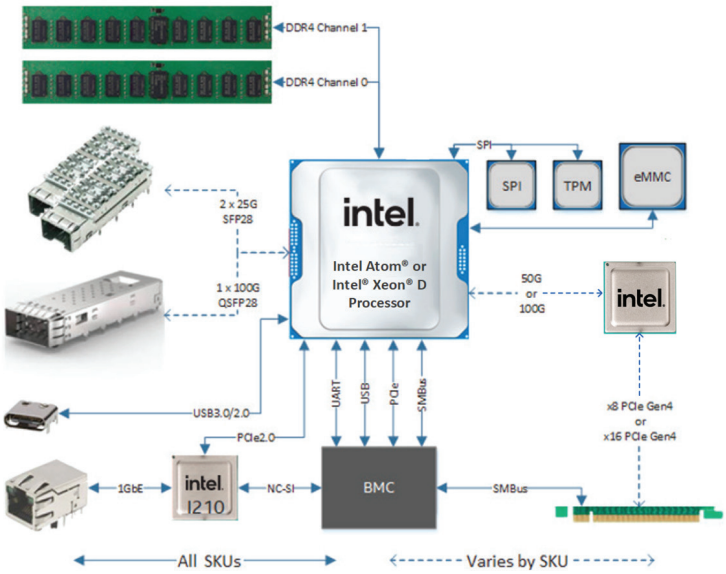

Figure 2. Expanded component view.

The autonomous compute resource isolates data and operations from the rest of the system, helping architects overcome incompatibilities among software components from multiple vendors, while providing better system-level security. The addition of hardware resources based on the reference design enables dramatic gains in solution capability and density. It also provides flexible capacity on demand for Day-2 upgradeability of in-place solutions.

Original equipment manufacturers (OEMs) and original design manufacturers (ODMs), working with network and security solutions providers, can take advantage of the reference design to bring network security accelerators to market more quickly. Intel is working with a number of partners to develop products, making a choice of technology vendors available to systems vendors, solutions integrators and end customers.

Reference Architecture	Intel Atom® Processor 8-Core Reference Design	Intel Atom Processor 16-Core Reference Design	Intel® Xeon® D Processor 4/8-Core Reference Design	Intel Xeon D Processor 10-Core Reference Design
CPU	Intel Atom P5721 processor	Intel Atom P5742 processor	Intel Xeon D-1713NT/1733NT processor	Intel Xeon D-1743NTE processor
Form Factor	Full height, half length		Full height, three-quarter length	
External Ports	2x 25GbE SFP28	1x 100GbE QSFP28	2x 10/25GbE SFP28	
Power Consumption	~50 to 90 watts	70 to 115 watts	90 watts to 145 watts	
Memory Capacity	Up to 32 GB @ 2933 MT/s		Up to 32 GB @ 2933 MT/s	
Host Interface	x8 PCIe Gen4	x16 PCIe Gen4	x16 PCIe Gen4	
Storage Capacity	Up to 256 GB eMMC		Up to 256 GB eMMC	
Throughput Target (Bi-directional Acceleration)	25 Gbps	50 Gbps	25 Gbps	50 Gbps
Throughput Target (Uni-directional Acceleration)	50 Gbps	100 Gbps	50 Gbps	100 Gbps

Reference design hardware specifications

The reference design includes variations, which are differentiated by processor (and core count), as well as I/O and networking resources.

SASE acceleration use case

SASE service providers deploy point-of-presence (POP) solutions that are geographically dispersed to be in relative proximity to user endpoints and on-prem, edge and cloud services. Enterprise users access their resources via SASE POPs that act as access gateways to help meet service-level objectives for latency and throughput in a secure manner. Distributed delivery of cloud-native security services avoids the need to backhaul WAN traffic to centralized locations to apply security policies.

This novel topology provides substantial bandwidth cost savings while also improving the user experience by cutting out the transfer latency associated with backhaul. A POP server cluster hosts any or all of the primary SASE components in real time, for all user network traffic:

- **Next-Generation Firewall (NGFW)** combines traditional firewall functionality with complementary services such as deep packet inspection, intrusion protection and threat intelligence.
- **Software-Defined WAN (SD-WAN)** dynamically self-optimizes to connect users to applications, centrally directing traffic across any combination of transport services, such as MPLS, 4G/5G and cable broadband.
- **Zero Trust Network Access (ZTNA)** provides seamless remote access to resources and applications while granting the least privilege possible, regarding all entities as untrusted for all other purposes.
- **Secure Web Gateway (SWG)** filters user-initiated traffic to detect and remove malware and other unwanted software, helping enforce corporate security standards and maintain compliance.
- **Data Loss Prevention (DLP)** monitors outgoing user traffic to identify sensitive information and prevent unauthorized egress, whether malicious or otherwise.
- **Cloud Access Security Broker (CASB)** is an enforcement point between users and cloud services that applies policies such as authentication, encryption and logging.

Well-designed and constructed SASE POPs ensure delivery of these services with fidelity across locations and endpoint types, such as worker laptops, IoT sensors/actuators and mobile devices. Quality of service is dependent on the ability to provide adequate POP reach, including both the number of locations and the capacity for each to deliver services. SASE vendors optimize POP servers for the most cost/capacity efficiency.

SASE POP server provisioning

In addition to Intel Xeon Scalable processors, Intel Xeon D processors are designed specifically for provisioning dense compute at the network edge, making them a popular choice as the foundation for SASE POP servers. The platform provides energy-efficient performance, hardware-based security and acceleration technologies and advanced integrated Intel Ethernet connectivity.

Solution architects can expand the service capacity of Intel Xeon Scalable and Intel Xeon D processor-based POP servers with the addition of one or more accelerators based on the Intel NetSec Accelerator Reference Design. For example, a two-socket POP server based on 20-core Intel Xeon D processors would have a total of 40 cores. Deploying two 16-core accelerator cards in the system makes an additional 32 Intel Atom® processor cores available, for an 80% increase in core count without increasing the server footprint. Each accelerator can run a separate SASE service, with its own set of compute, memory and I/O resources, to provide further parallelization of the workload to improve deterministic performance and security.

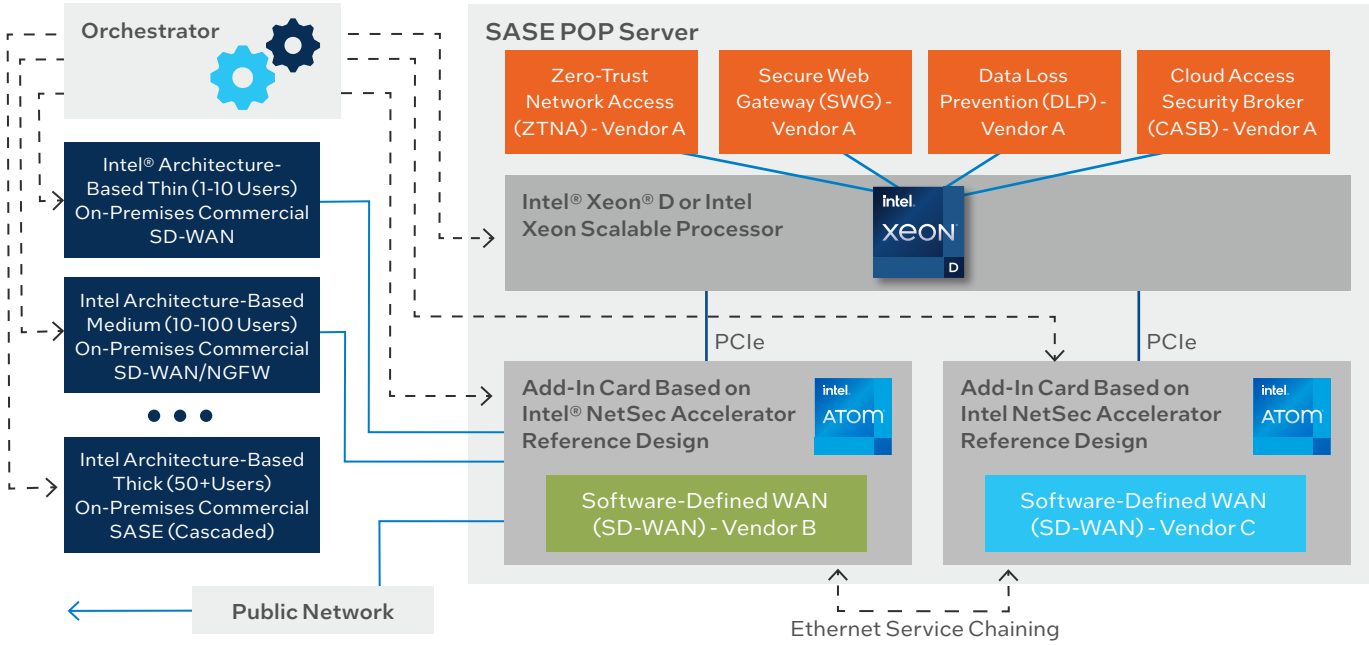

Figure 3. SASE Architecture enabled by Intel® NetSec Accelerator Reference Design.

Accelerator connection topology

An accelerator card based on the Intel NetSec Accelerator Reference Design may be connected directly to the external public network, enabling certain SASE functions to be accomplished independently of the main Intel Xeon Scalable processor or Intel Xeon D processor on the host server. This connectivity also makes it possible for the card to provide inline capabilities that independently push inbound data to the appropriate compute resource, for increased efficiency. Ethernet-based service chaining can interconnect services directly between accelerators, combining capacity or functionality, as in the example shown of delivering SD-WAN and security stack together as a single service from two distinct accelerators. The integrated capabilities of the Intel Xeon D or Atom processor facilitate resource sharing to improve efficiency, with load balancing between ports without involvement from the Intel cores running on the Intel NetSec Accelerator Reference Design solution.

For some implementations, an accelerator based on the Intel NetSec Accelerator Reference Design may not be connected to the outside world. For example, it may use service chaining to deliver services through another accelerator that is on the public network or provide functionality such as a sandboxing application for deep packet inspection that does not require external connectivity.

Potential for SASE acceleration in real world deployments

The SASE acceleration use case demonstrates some representative usage patterns for devices based on the reference design in SASE POP servers:

- **Increased density and infrastructure efficiency** based on deploying added compute capacity with one or more accelerators with the full range of server-on-a-card resources.
- **Multi-vendor integration** instantiated by the SASE POP server, providing unified services based on solutions that would otherwise be incompatible on a single system.
- **Advanced traffic control** using the integrated network switch in the Intel Atom processor to direct inbound data appropriately, independent of the main processor.
- **Service chaining and delivery of distributed SASE services** using multiple accelerators in a single SASE POP server.

By expanding the hardware and making it more capable, these factors help lower cost of operation by reducing the total number of servers required to accomplish a given performance goal.

Building blocks for accelerated networking and security

The reference design provides an independent, functional compute node that delivers server-class performance and reliability within a conservative power envelope. It provides inline cryptography for IPsec, as well as look-aside operation, which is appropriate for asynchronous bulk encryption workloads.

Integrated Intel® QuickAssist Technology (Intel® QAT) Gen3 accelerates symmetric and asymmetric encryption, driving up to 100 Gbps of throughput. The Intel QAT hardware communicates directly with the onboard Ethernet controller to decide which packets to process and which to pass to the processor. By shortening the data path, this capability enables inline IPsec.

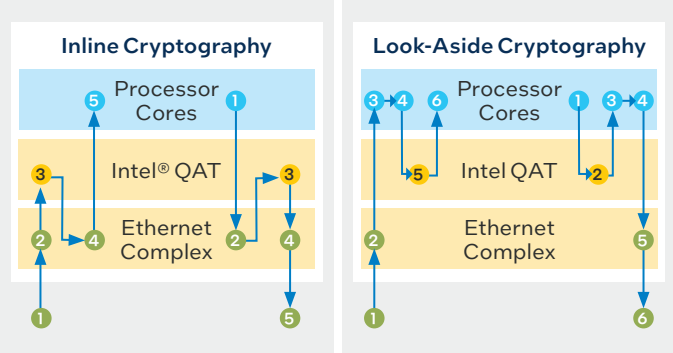

Figure 4. Enabling inline IPsec.

Intel Atom® processor — high performance per watt and inline IPsec

The foundation of the Intel NetSec Accelerator Reference Design is the Intel processor, which provides high throughput for security and networking workloads in an energy-efficient SoC form factor. The highly integrated Intel Atom processor incorporates Intel Ethernet and hardware accelerators into the SoC package, which provides low-latency operation and drives significant advantages in reduced equipment cost, space/server requirements and energy consumption.

Intel Xeon D processor

Intel Xeon D processors are designed for high computational throughput with low thermal design power (TDP). Packaged in a highly integrated system-on-chip design, the processor provides extensive capabilities for performance and security, including the following:

- **Intel® Software Guard Extensions (Intel® SGX)** protects data while in use by creating private, isolated areas of memory called secure execution enclaves where unencrypted data can be operated on, beyond the reach of software and users, regardless of their privilege levels.
- **Intel® AES New Instructions (Intel® AES-NI)** accelerate resource-intensive parts of the AES encryption algorithm in hardware.
- **Intel® Advanced Vector Extensions 512 (Intel® AVX-512)** boosts performance for demanding requirements such as AI and 5G workloads with ultra-wide 512-bit vector operations that work on more data per clock cycle than predecessor technologies.

Intel® Ethernet 800 Series Controller — advanced networking for security functions

The reference design includes Intel Ethernet 800 Series Controller with Dynamic Device Personalization (DDP) to perform packet processing and traffic shaping to improve performance. DDP enables network administrators to establish multiple profiles for traffic types, specifying packet-handling parameters and optimizations for each. Traffic prioritization based on DDP can be configured dynamically at runtime to enhance flexibility and agility.

Conclusion

The Intel NetSec Accelerator Reference Design is a highly efficient and easy-to-use blueprint for delivering all the functionality of an independent server in a PCIe form factor for effective slot-in integration into security appliances and edge platforms. It provides a separate physical compute environment within the server chassis.

The addition of resources including system memory in the reference design enables the accelerator device to run independently of the main server platform. The accelerator runs security services independently to expand server capacity, including support for multi-vendor software stacks that would otherwise not be compatible on a single system. This capability may help end customers reduce TCO by supporting more services per host.

The pre-validated reference design streamlines the development of new security appliances by OEMs and ODMs, working with network and security solutions providers. It enhances platform flexibility and reduces design-in requirements for those solution developers, helping them bring new security products to market more quickly and cost-effectively.

To learn more:

Intel® Xeon® D processors

Intel Atom® processors

Intel® Ethernet products

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Intel technologies may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

Code names are used by Intel to identify products, technologies or services that are in development and not publicly available. These are not "commercial" names and not intended to function as trademarks.

Any forecasts of goods and services needed for Intel's operations are provided for discussion purposes only. Intel will have no liability to make any purchase in connection with forecasts published in this document.

0224/DL/MESH/349364-002US