Solution Brief

Intel® Xeon® D-2700 Processors
Intel® Xeon® D-1700 Processors

High Performance and Density with Low Power at the Enterprise Edge

The Intel® Xeon® D processor is a workhorse engine for universal customer premise equipment (uCPE) that powers the enterprise edge's increasingly critical role in business computing. Optimized for density and delivered in a power-efficient system-on-chip (SoC), the processor provides the performance needed for demanding local workloads such as security and machine learning inference, as well as the convergence of information technology (IT), operations technology (OT), and communications technology (CT).

A growing proportion of enterprise data is generated on the network edge, far from on-premise (on-prem) infrastructure or public cloud data centers. The internet of things (IoT) exemplifies this trend, with massive, highly distributed data sets. At the same time, financial, technical, and privacy considerations may require the workloads based on that data to be executed close to the point where the data is generated. This is especially true for workloads that require low latency, or when the raw data is so large that transmitting it to a data center or public cloud for processing is cost-prohibitive.

For example, emergency shut-off of industrial equipment may require latencies on the order of one microsecond, and surveillance cameras may generate massive data sets that consist mainly of material that contains no events of interest. Both of these issues can be addressed by processing at the enterprise edge, as illustrated in Figure 1. The edge and associated services such as security may be hosted on-prem at the remote site, in the cloud, or a combination of both. The edge location itself may range from a home or mobile worker or retail shop to a regional data center.

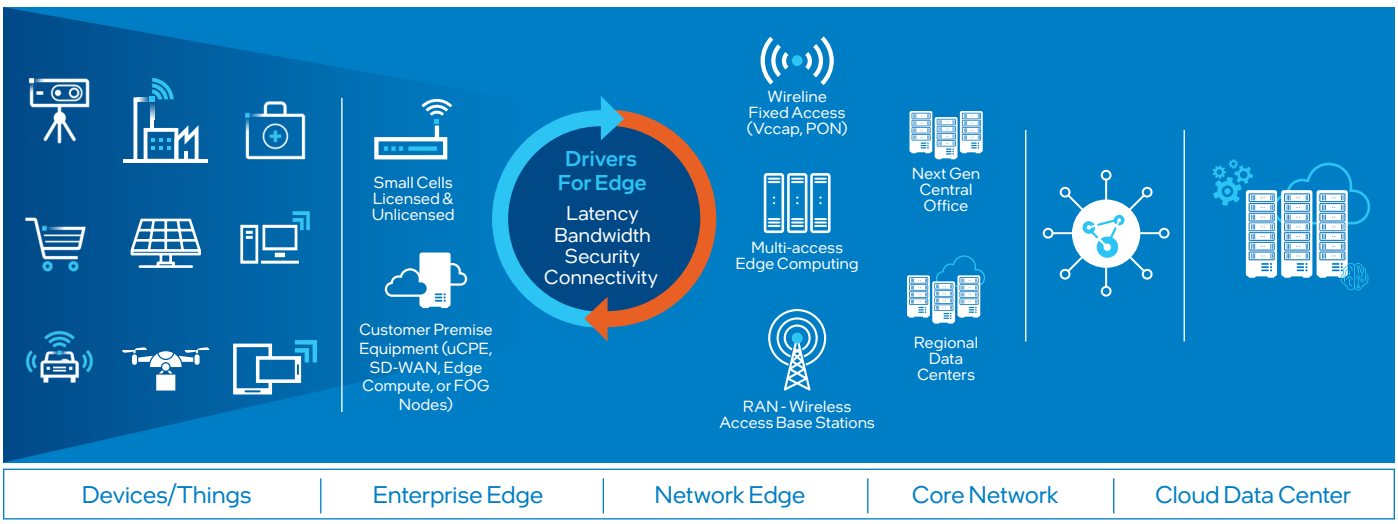

Figure 1. Flexibility to run workloads on-prem, at the network edge, or in the data center/cloud.

In an enterprise edge topology, low latency requirements are met and bandwidth costs are reduced by avoiding transmission back to the network core. Reduced levels of data transmission also reduce associated security exposures for potentially sensitive information, and the enterprise edge provides local connectivity for devices and IoT endpoints. In an effort to make the most efficient use of network resources possible, network architects are beginning to converge information technology (IT), operations technology (OT), and communications technology (CT) onto common shared computer hardware.

Enterprise Edge as a Driver for Network Transformation

In a traditional hub-and-spoke private network approach, multiple physical appliances installed at branch offices provide security and other services, with WAN connectivity from a local service provider. This arrangement requires onsite technical staff for networking and security, as well as multiprotocol label switching (MPLS) connectivity for each branch (including bandwidth headroom), which adds to OpEx. Cloud connectivity provided through the data center creates a potential bottleneck at the data center firewall that may compromise the user experience.

Software-defined WAN (SD-WAN) improves on the traditional model by providing branch offices with connectivity to both the data center and directly to the cloud over broadband and wireless networks, in addition to MPLS. SD-WAN provides a dynamic, flexible, cost-efficient network that increases agility, promotes user experience by relieving the bottleneck at the data center, and reduces the need for on-prem technical staff at branch offices. However, this approach increases the attack surface of the enterprise, requiring the deployment of a virtual firewall at each branch office to protect it from unwanted network traffic and to ward off outside intrusion into the private network.

An alternative approach is the secure access service edge (SASE), which extends the SD-WAN architecture by converging networking services with cloud-hosted security services such as firewall, secure web gateway (SWG), cloud access service broker (CASB), and zero-trust network access (ZTNA). The SASE architecture hosts these services from the cloud edge, from edge point-of-presence (POP) locations which consist of racks of commercial off-the-shelf (COTS) servers deployed at co-location facilities, or both. This approach supports enterprise IT strategic initiatives with a cloud-first orientation, intrinsic security, improved user experience, and operational simplicity. The global SASE market size is projected to grow from US$1.2 billion in 2021 to US$4.1 billion by 2026, for a CAGR of 26.4 percent.[1]

Edge Processing Requirements Based on Site Capacity and Functions

The level of processing needed at a given enterprise edge location depends both on the size of the location and the types of services that are provided locally. For example, the number of users at a site is a primary consideration, as is the degree of compute-intensive tasks performed locally, such as machine learning inference.

Processors chosen to meet those specific compute requirements may range from Intel Atom® processors for relatively low local compute requirements to Intel® Xeon® D processors as the midrange offering and Intel® Xeon® Scalable processors at the high end, as illustrated in Figure 2. The breadth of Intel's platform offerings enables flexibility of scale, functionality, and capacity in enterprise edge compute resources.

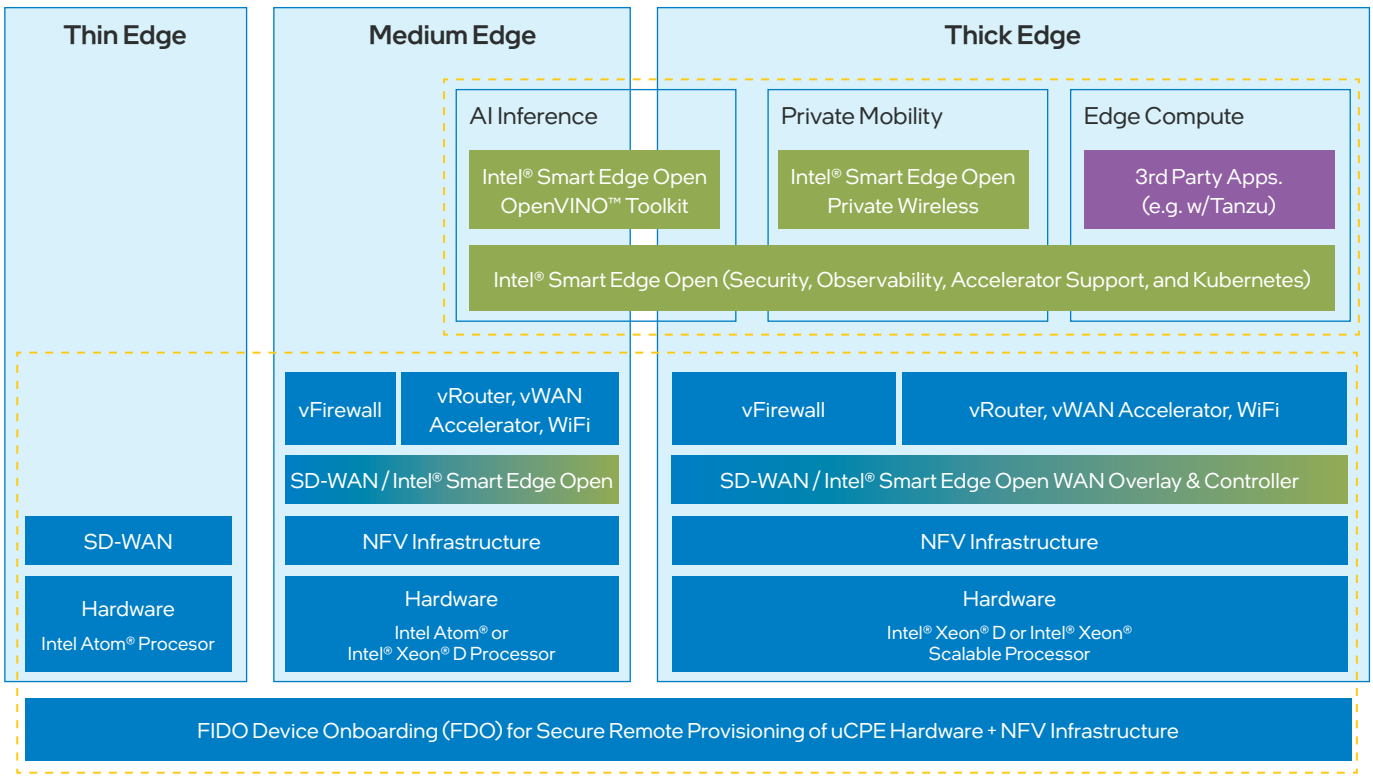

Figure 2. Enterprise edge deployment segmentation.

Enterprise edge deployments may be characterized as "thin edge," "medium edge," or "thick edge," based on the degree of functionality being provided by the edge compute platform:

- **Thin edge – small branch office, small retail store, work from home.** SASE and AI inference are provided as cloudhosted services. A typical thin edge device is a SD-WAN appliance based on a low core-count Intel Atom processor.
- **Medium edge – medium branch office, large retailer, bank, medical office.** SASE and AI inference are partly cloud-hosted and partly hosted on-prem at the edge. A typical medium edge device is uCPE based on an Intel Xeon D processor.
- **Thick edge – corporate headquarters, regional data center, factory, hospital.** SASE and significant local compute power are fully hosted at the enterprise edge. A typical thick edge device is uCPE based on an Intel Xeon D processor or Intel Xeon Scalable processor.

For many network edge deployments, the midrange Intel Xeon D processor provides the ideal balance between cost, processing power, and power consumption.

Intel® Xeon® D-2700 and D-1700 Processors for the Enterprise Edge

Built to deliver density- and power-optimized compute, Intel Xeon D-2700 and D-1700 processors feature advanced security and AI inference features, as well as per-core performance improvements over predecessors. The SoC package is highly integrated for simplicity and power efficiency, satisfying the requirements of compact device form factors for indoor, outdoor, and ruggedized implementations. The platform's built-in hardware acceleration for cryptography, compression, and machine learning inference are well suited to the needs of diverse and converged workloads at the enterprise edge. Key performance and security features of the platform include:

- **Intel® Deep Learning Boost (Intel® DL Boost)** accelerates machine learning workloads at the edge by eliminating unneeded precision in calculations so they can be completed more quickly.
- **Intel® AES New Instructions (Intel® AES-NI)** accelerate resource-intensive parts of the AES encryption algorithm in hardware.

Table 1: Benefits of platform improvements in Intel® Xeon® D-2700 and D-1700 processors.

New Technology		Benefit
Core Acceleration	Vector Neural Network Instructions (VNNI)	Acceleration of AI inferencing
	Vector Byte Manipulation Instructions (VBMI)	InCore compression/decompression acceleration for in-memory database workloads
	VPMADD52 instruction	Public key crypto generation – SSL front end web server acceleration
	New SHA extensions	Acceleration of hashing, SSL, TLS, IPsec, dedup, blockchain
	Vector AES	Database workloads acceleration
Performance Tuning and Management	Intel® Resource Director Technology (Intel® RDT)	Monitoring and control of memory and Last Level Cache usage
	Flexible All Core Turbo[2] and Prioritized Base Frequency[2]	Higher frequencies for a subset of cores while all cores are active to manage application-level performance
	Intel Speed Select[2]	Higher base frequency at lower core counts for dynamic SKUing
	New algorithms in Intel® QAT v1.8[2]	SHA3, SM3, SM4, ChachaPoly added to accelerate IPsec, TLS, and DTLS workloads
	Internal block grounding	Power optimized sub-segment SKUing
	Virtualization improvements	Enhancing performance of NFV workloads
	Asynchronous DRAM Refresh (ADR)	ADR with enhanced battery backup significantly decreases battery size requirements
Security	Intel® SGX-Trusted Environment Mode (Intel® SGX-TEM)	Fine grain data protection by means of application isolation in memory
	Intel® Total Memory Encryption – Multi-Tenant (Intel® TME-MT)	VM container isolation for multi-tenant platforms
	Intel® Platform Firmware Resilience (Intel® PFR)	Protect, detect, and correct security threats in transit, boot, and runtime

- **Intel® Advanced Vector Extensions 512 (Intel® AVX-512)** boosts performance for demanding workloads with ultra-wide 512-bit vector operations that work on more data per clock cycle than predecessor platforms without Intel AVX-512.
- **Intel® QuickAssist Technology (Intel® QAT)** accelerates encryption and compression; the platform is capable of supporting up to 100 Gbps crypto and 70 Gbps compression functionality simultaneously.
- **Integrated Intel® Etherne**t provides up to 100 Gbps of throughput, with link options from 1GbE to 100GbE, low power consumption, and advanced storage support with Remote Direct Memory Access (RDMA).
- **Intel® Software Guard Extensions (Intel® SGX)** protects data while in use by creating private, isolated areas of memory called secure enclaves where unencrypted data can be operated on, beyond the reach of software and users, regardless of their privilege level.
- **Memory Encryption** supports existing software without modification while cryptographically protecting memory against hardware attacks using the NIST AES XTS encryption standard with hardware-generated keys from a hardened random-number generator implemented in silicon.

The balanced platform delivers a broad-based set of technologies for workload acceleration, streamlined maintenance, and hardware-assisted security. Details about these technologies and the benefits they confer on enterprises are summarized in Table 1.

Intel® Select Solutions for Secure Access Service Edge (SASE):

The complexity of enterprise edge networks today requires the right mix of hardware and software components to build an infrastructure that meets each organization's requirements. The Intel® Select Solutions for SASE reference design eliminates guesswork with rigorously benchmark-tested and verified solutions optimized for real-world performance. The Intel® Select Solutions provide the latest tightly specified hardware and software components, purpose-built for SASE use cases, that can dramatically accelerate deployment and time to new services, while reducing implementation risk for enterprise customers.

Learn more:
https://networkbuilders.intel.com/solutionslibrary/intel-select-solutions-for-secure-access-service-edge-sase

Platform Features to Enable Convergence of IT, OT, and CT Workloads

Diverse workloads come together at the network edge, such as AI, network, and media. Converging IT, OT, and CT workloads onto a single enterprise platform based on the Intel Xeon D-2700 or D-1700 processor simplifies edge infrastructure, which can drive dramatic cost efficiencies and enable service innovation.

To help realize the efficiencies of this converged architecture, the Intel Xeon D-2700 and D-1700 processors support enhancements to the Intel® Advanced Vector Extensions 512 (Intel® AVX-512) instruction set. The enhancements discussed in this section accelerate machine learning and security functions, optimizing converged IT, OT, and CT workloads for medium and thick edge applications.

Increased AI Inference Performance for Analytics with VNNI

Intel DL Boost Vector Neural Network Instructions (VNNI) deliver a significant performance improvement by combining three instructions into one, which helps optimize cache utilization and avoid potential bandwidth bottlenecks. VNNI accelerates deep learning inference for workloads such as image classification, object detection, speech recognition, language translation, and more.

The Intel Xeon D-2700 and D-1700 processors add VNNI support for TensorFlow 2.5 and Intel OneDNN, which increases inferencing performance for networking, security, image, and video analytics.

Improved IPsec and SSL Performance

New technologies for cryptographic performance increase IPsec throughput and the number of SSL connections established per second, as compared to previous generations.

Intel® AES New Instructions (Intel® AES-NI) accelerate resource-intensive parts of the AES encryption algorithm in hardware. The Intel Xeon D-2700 and D-1700 processors introduce Vector Intel AES instructions, which extend Intel AES-NI to support vector processing of up to four AES blocks (128 bits) at a time using 512-bit wide registers, improving efficiency.

Secure Hash Algorithm (SHA) is used for a wide variety of hashing functions in IPsec, SSL, and TLS connections, including to protect passwords, signatures, and certificates. The newly introduced Intel® SHA Extensions provide a significant improvement in SHA-256 performance with hardware-based acceleration.

Integrated Intel QAT v1.8 with in-line IPsec processing[2] enables even higher performance levels for uCPE, SD-WAN, and security appliances, including up to 100 Gbps crypto and up to 70 Gbps compression.

Accelerated Security Functions with Enhanced AVX-512 VBMI

New processor instructions increase Hyperscan pattern matching performance, which accelerates security network functions including next-generation firewalls (NGFW), intrusion protection systems (IPS), and intrusion detection systems (IDS).

Vector Bit-Manipulation Instructions (VBMI) increase efficiency by enabling data movement within registers at byte granularity, rather than only being able to move larger data constructs (i.e., words, doublewords, etc.).

The Intel Xeon D-2700 and D-1700 processors enhance Intel AVX-512 support for VBMI to include compression and decompression at the per-byte level. This increased granularity improves the efficiency of Hyperscan regular expression (RegEx) pattern matching, which enhances performance for application identification compared to predecessor platforms.

Conclusion

As networks transform to support the convergence of diverse workloads at the edge, Intel Xeon D-2700 and D-1700 processors are designed to accommodate widely variable and rugged compute environments by providing server-class processing in a highly compact, power-efficient SoC. For medium and thick edge deployments, the highly integrated SoC offers high performance and density at low TDP, with advanced hardware features to improve security and accelerate cryptography, compression, and machine learning inference.

More Information: www.intel.com/xeond

[1] MarketsandMarkets Research, August 2021. "Secure Access Service Edge (SASE) Market with COVID-19 Impact Analysis, by Offering (Network as a Service and Security as a Service), Organization Size (SMEs and Large Enterprises), Vertical, and Region - Global Forecast to 2026." https://www.marketsandmarkets.com/Market-Reports/secure-access-service-edge-market-220384224.html.

[2] Availability varies by SKU.

Features, SKUs, and frequencies are preliminary and subject to change.

Performance varies by use, configuration, and other factors. Learn more at https://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 configuration disclosure for configuration details. No product or component can be absolutely secure.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Your costs and results may vary.

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

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a nonexclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

0122/FS/MESH/346436-001US