

From the Data Center to the Cloud to the Intelligent Edge, the Network of Business Runs on Intel

A common platform built on Intel® architecture delivers consistency, security, and scalability for a modernized, “cloud-optimized” network that prepares businesses for the future.

Customer Strategic Challenges

The digital economy has fundamentally changed the role of IT as an enabler of the business. Advancements in storage, memory, and compute technologies—as well as the rapid growth of connected devices and data-intensive applications—is changing how businesses compete. These changes combine with increased adoption of cloud computing as well as emerging security and compliance needs to make network infrastructure critical to business success.

New demands on the network are driving changes to application traffic patterns, causing bottlenecks and making it difficult to meet service-level agreements (SLAs), manage total cost of ownership (TCO), and maintain uniform security policies across private and public clouds. To enable the business, the network must connect all parts of the enterprise to the data and services needed across the organization, and it must do so in a reliable and secure manner. This is simply not possible with legacy, fixed-function network equipment, which cannot provide the necessary performance, agility, and consistent application experience across multicloud and edge environments. To compete today and be ready to adapt to the needs of tomorrow, enterprise IT must rethink how it architects and manages its network resources.

Optimize and Simplify from the Data Center to the Edge

A modern network typically extends from the core data center out to the edge and across multiple public, private, and hybrid clouds, as illustrated in Figure 1. This reality is changing how networks are architected, deployed, and managed. Modern, future-ready networks must bring the scale, reliability, and efficiency associated with the cloud into the enterprise network. This can be done through server-based network virtualization and orchestrated virtual network functions (VNFs) that allow network infrastructures to deliver innovative services with efficiency and ease.




Figure 1. Edge-generated data creates a wave of end-to-end implications. Network infrastructures must extend out to interconnect the edge, multiple clouds, and branch locations to enable the business to use data to make smarter decisions and engage customers in real time.

Software-defined networking (SDN) and network function virtualization (NFV) are attractive to many businesses due to the potential for cost savings, improved network performance, increased productivity, and enhanced security. A virtualized network increases the level of abstraction, making the network simpler and more efficient to manage. According to IDC, by 2021, 70 percent of enterprises will adopt dynamic software-defined branch and network solutions that deliver security and flexibility across cloud, data center, and edge interactions.¹

Virtualized networks can support extreme levels of scalability, programmability, and security across an ever-growing volume and variety of networking workloads—from the data center and network core to the edge. And when combined with the latest hardware and software acceleration technology, they can enable agility within the data center to effectively deliver valuable network services in a more efficient, flexible, and cost-effective way.

Growth of the Intelligent Edge

With the proliferation of internet of things (IoT) and the explosion of data being generated from smart devices, the strategic importance of the edge for business grows as well. Gartner predicts that by 2022, 50 percent of enterprise-generated data will be created and processed outside the data center or cloud.² This data is in close proximity to the customer or end user, and it holds the potential for faster, better insights for the business—but only if it can be managed and processed fast enough.

This trend of processing data where it is generated will require an extension of IT infrastructure out to the edge with more sophisticated compute, storage, and networking resources to enable analytics and AI scenarios and deliver real-time outcomes. IDC predicts this will grow in importance, and by 2022, 40 percent of enterprises will have doubled their IT asset spending in edge locations to better deliver digital services to local users and things.³

The growth in edge infrastructure use cases including those shown in Table 1 will help reduce application latency, improve service capabilities, and optimize TCO, as well as address data privacy requirements. It will also support the growing use cases that are emerging at the edge, such as smart retail experiences, supply chain optimization for faster delivery, digital security surveillance, intelligent manufacturing, and improved health care interactions with patients. With the right edge infrastructure, enterprises can run a greater volume and variety of data-intensive workloads at the edge with the bandwidth and latency required for a consistent customer experience.

Table 1. Edge computing example use cases.

Networking	Security	Analytics	IoT
Routing	Firewall	Security Analytics	Video Surveillance
SD-WAN	IPS	Video Analytics	Smart Retail
WAN Optimization	IDS	AI Workloads	Smart Manufacturing

Network Architecture for the Intelligent Edge

Edge computing has quickly made current branch infrastructure outdated and ineffective. The traditional location-based router approach with WAN optimized appliances is no longer sufficient to meet the demands of enterprise users, applications, and services that reside in distributed locations. Fixed function appliances are not open and programmable. This limits the ability to adapt and scale for future needs. They also increase cost and complexity as each network function is tied to a fixed-function piece of hardware, making management and maintenance difficult.

uCPE and SD-WAN Enable Cloud-Centric Network

A flexible, cloud-centric network architecture replaces closed appliances with universal customer premises equipment (uCPE) platforms that are hosted on industry-standard servers to enable the dynamic deployment of a wide variety of software VNFs. Since uCPE platforms and software-defined wide-area networks (SD-WANs) are open and programmable, they provide a consistent platform for VNF deployment and help ensure the ability to scale for future needs.

This increased simplicity and flexibility gives IT more control by allowing it to intelligently route network traffic over standard, high performance servers, and dynamically change bandwidth requirements to avoid overprovisioning and lower the TCO of their WAN while maintaining high levels of service. And once VNFs are integrated, they can be easily deployed via software over a range of host platforms, from small branches to large data centers.

Architect with a Common Platform

Industry-standard hardware offers a range of platform options. As an enterprise begins the process of network modernization, it is important to consider the value of building a network architecture with a common platform that can extend across workloads and locations to allow for maximum flexibility and reduce complexity. This creates the ability to add capacity to scale performance without the need for architectural redesign. There are a variety of benefits to using the same processor architecture across an end-to-end network infrastructure, including the following:

- **Better control of security and trust.** With a common platform you can take advantage of built-in, silicon-based security technologies such as root of trust and cryptographic acceleration, across the environment. You achieve greater efficiencies and better manageability by standardizing your platform.

- **Performance scaling.** Predictable performance scaling is aided by architecture consistency that drives greater device platform synergy with respect to software, network interface card (NIC) features, virtualization technology, and accelerators. From Intel Atom® processors to Intel® Xeon® D processors to Intel® Xeon® Scalable processors, Intel architecture provides scalability across different edge use cases and from data center to cloud.
- **Application portability.** A common platform provides increased compatibility for application development, removing the edge device from the process and allowing for the reuse of application software. With multiplatform environments, customization is needed. An environment based entirely on Intel architecture enables an agile and portable application development environment.
- **Lower TCO:** A common platform reduces complexity. When deploying new features and functions at the data center, cloud, or edge, the financial cost benefit is likely to be higher. Likewise, the economics of developing software improves when it can be applied end-to-end.
- **Improved network response.** When different SD-WAN device types employ similar accelerators, they work better together. For example, if all devices can perform the same compute-intensive cryptography functions in hardware (instead of software), it becomes less likely that one of the devices will become a major bottleneck for network traffic.

A consistent and flexible architecture ensures a future-ready network that can better adapt in a dynamic environment without have to change or upgrade devices.

Across All Use Cases and Infrastructure Requirements—Intel Has You Covered

The Intel® processor portfolio is the basis for next-generation platforms to build virtualized, cloud-optimized networks. It offers a consistent architecture that scales and adapts with ease to support changing business and compute needs. Intel's range of processors can meet the variety of use cases coming to the enterprise edge.

- **Intel Xeon Scalable processors** deliver maximum performance and scalability for data center and edge environments with flexible operating space and power.
- **Intel Xeon D processors** are delivered in a system-on-a-chip (SOC) form factor optimized for edge environments that may be constrained in terms of space and power.
- **Intel Atom® C3000 processors** are purpose-built for very low power and optimized for edge environments constrained by physical size and extreme temperatures.

With this diverse processor portfolio, which is illustrated in Figure 2, IT can build a homogeneous network platform that supports a variety of use cases and has the flexibility to scale in or out without the added cost and complexity of re-architecting. This will be important for enterprises looking to deploy new, innovative services from the edge and increase the variety of AI workloads running at the edge.




Figure 2. A common Intel®-based platform provides scalability and consistency across different sizes of branch and edge locations, to enable flexibility and better TCO.

Benefits Beyond the Processor

The value of an Intel®-based networking platform extends beyond performance, as illustrated in Figure 3. The choice of processor architecture also dictates different performance accelerators, security features, and software packages accompanying the CPU. These technologies help deliver higher performance and security consistency.




Figure 3. Intel® architecture benefits and advantages for edge to data center to cloud.

For example, the Data Plane Development Kit (DPDK) is a set of libraries and network interface controller drivers that increase fast packet processing on Intel architecture platforms by up to a factor of 10x.⁴ Packet processing software running on SD-WAN appliances—from SOHO to the data center—can benefit from DPDK optimizations for Intel architecture. In contrast, if a CPE appliance is built with another CPU architecture, it may not be able to reap the same gains from DPDK, possibly resulting in lower performance.

Additional features and capabilities available across select Intel architecture processors include the following:

- **Intel® QuickAssist Technology (Intel® QAT).** On-CPU hardware acceleration for security, authentication, and compression.
- **Hyperscan.** Open source regex matching library, optimized for Intel architecture and suitable for deep packet inspection (DPI), intrusion detection (IDS/IPS), and firewall applications.
- **Intel® Virtualization Technology (Intel® VT).** Hardware assist for virtualization software, eliminating performance overheads, improving security, and reducing software size, cost, and complexity.

- **Intel® Platform Trust Technology (Intel® PTT).** Integrated CPU capabilities that store keys, passwords, and digital certificates, and that support Microsoft* firmware requirements for a discrete Trusted Platform Module (TPM 2.0).

Accelerate with Optimized Solutions

The goal of network modernization is to reduce costs and complexity, improve security, and address latency and bandwidth challenges. That vision depends on an agile and secure network that can scale when and how it's needed to deliver a consistent and reliable user experience. Within a dynamic business environment, these needs will constantly be in flux as new applications are deployed, data volumes grow, and compute needs increase at the edge of the network.

As a global leader in data center technologies, Intel continues to drive platform innovation for next-generation network capabilities required in the data era. Intel also realizes the importance of standards and industry collaboration and the need for a seamless flow of data across the network, clouds, and devices. Along with leading ecosystem and solution partners, Intel is facilitating ecosystem alignment and collaboration for the delivery of world-class solutions to help enterprises deliver "cloud-optimized" networks.

For more information about how to get started on your network modernization journey, contact your Intel representative today:
www.intel.com/content/www/us/en/company-overview/contact-us.html



¹ IDC FutureScape: Worldwide IoT 2018 Predictions; <https://www.brighttalk.com/webcast/15593/336946/idc-futurescape-worldwide-internet-of-things-iot-2018-predictions>.

² Gartner, Living on the Edge: How Digital Business Will Push Infrastructure to the Edge, 2018; <https://www.equinox.com/resources/analyst-reports/edge-completes-cloud-gartner/>.

³ IDC FutureScape: Worldwide Data Center 2019 Predictions; <https://www.brighttalk.com/webcast/15909/334918/idc-futurescape-worldwide-datacenter-2019-predictions>.

⁴ Intel website, "Data Plane Development Kit (DPDK)," <https://www.intel.com/content/www/us/en/communications/data-plane-development-kit.html>.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit <http://www.intel.com/benchmarks>.

Performance results are based on testing as of December 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. This document may contain product features that are currently under development. This overview of new technology represents no commitment from VMware to deliver these features in any generally available product. Features are subject to change, and must not be included in contracts, purchase orders, or sales agreements of any kind. Technical feasibility and market demand will affect final delivery. Pricing and packaging for any new technologies or features discussed or presented have not been determined.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors known as errata which may cause deviations from published specifications. Current characterized errata are available on request.

Intel, the Intel logo, Intel Atom, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

Microsoft, Windows, and the Windows logo are trademarks, or registered trademarks of Microsoft Corporation in the United States and/or other countries.

*Other names and brands may be claimed as the property of others.

© 2019 Intel Corporation. All rights reserved. 0119/JL/MESH 338338-001US