Solution Brief

Life on the Edge: It's About Security and High-Performance

Take advantage of Edge Computing while staying cyber-secure

In the modern technology environment, yesterday's cutting-edge solution can quickly become today's legacy implementation. The ways in which businesses have had to respond to the pandemic over the past two years provides excellent examples of this. From adopting contactless purchasing practices to shifting to online ordering to supporting remote workers, companies that used technology to adapt and transform their business practices were able to thrive even during challenging times.

During the pandemic, cloud computing went truly mainstream. With massively scalable computing capabilities and a pay-as-you-go pricing model, companies trying to adapt to the pandemic restrictions eagerly embraced cloud solutions. As those deployments mature, additional technologies like IoT, AI, and video analytics are significantly increasing the amount of data being generated, resulting in higher bandwidth costs and increased latency as data travels to far-away data centers to be processed.

Edge computing works to improve the responsiveness of cloud solutions by moving data processing closer to where data is generated and used. This enables organizations to accelerate processing and reduce latency and bandwidth costs. The benefits of computing at the edge are so compelling that Gartner[1] predicts that by 2025, 75% of enterprise-generated data will be created and processed outside of the traditional data centers or clouds.

The benefits of Edge computing do not come without tradeoffs. Because these edge devices process an increasing amount of data critical to business operations, they are an enticing target for cybercriminals. Edge computing platforms must have a robust security foundation to secure them against attacks, along with tools to allow IT departments to properly configure and maintain them. Additionally, the explosive growth of hybrid work environments, with devices in many places and on different networks, gives cybercriminals more attack vectors than ever before. Organizations wanting to stay ahead of the criminals need to address both data processing at the edge and work-from-anywhere security.

Security Service Edge Protects from Endpoint to Edge to Cloud

Security Service Edge (SSE) describes an architectural framework to securely connect users, systems, and endpoint devices to applications and services that may be located anywhere. It is a combination of several technologies that had been developing independently. These technologies offer comprehensive capabilities with modern network security functions to support edge computing needs.

SSE allows enterprises to respond to security needs arising from today's network of remote workers and edge workloads connecting to corporate and cloud resources. Rather than forwarding all the remote worker and edge traffic to

Table of Contents

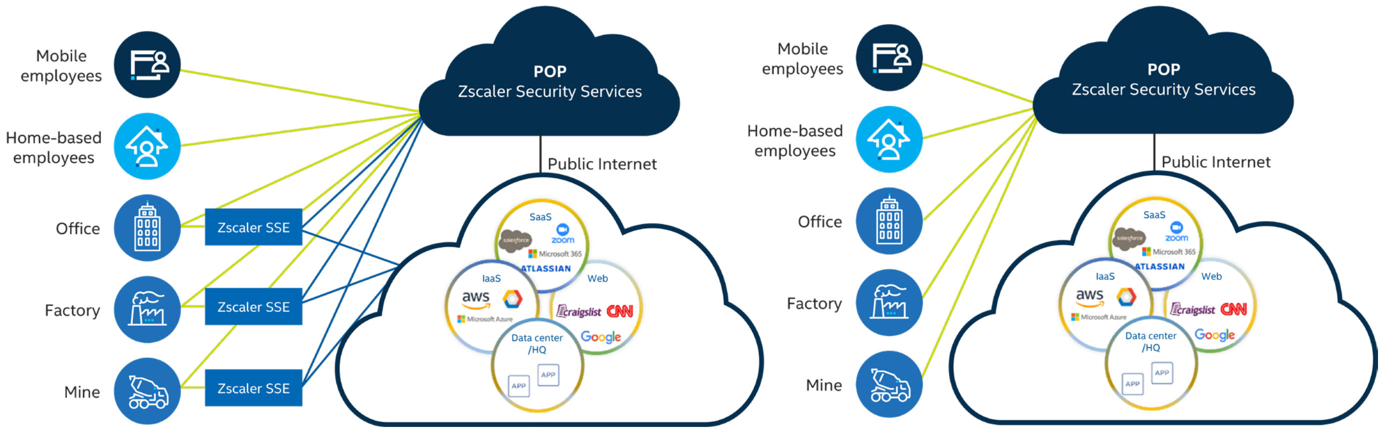

Figure 1. Zscaler Zero Trust secures real-time Edge compute (left diagram) in locations where massive amounts of Edge traffic are generated and only intermittent Internet connection is required or available. Zscaler Edge solutions interoperate with, and extend, the global Zscaler Zero Trust Exchange that connects any device on any network to any workload (right diagram).

corporate data centers, traffic is inspected through cloud deployed SSE before being sent directly to its destination. This lowers processing overhead in Multiprotocol Label Switching (MLPS) connections and network security appliances in the corporate data center.

While organizations know they need to increase agility and scalability while reducing complexity, finding solutions that deliver those capabilities can be a challenge. Intel, Zscaler, and Supermicro have partnered for years to create scalable, high-performance SSE implementations that keep up with the most demanding security needs of enterprise customers. These implementations can now secure high-throughput, low-latency workloads at the far edges of the network where sub-5 millisecond latencies are the new normal.

Zscaler: High-Reliability, High-Availability, and High-Scalability

Customers demand SSE solutions that deliver reliability, availability, and scalability 24/7. Zscaler has architected its solution using zero trust security principles: no user, application, or device is inherently trusted. Instead, the Zscaler Zero Trust Exchange (ZTE) cloud-based security establishes trust before any connection is made, determining appropriate levels of access and restrictions based on the user, device, application, location, and content to keep users and data safe (see Fig. 1). Zscaler delivers this zero-trust security between devices communicating over 5G/4G LTE, Wi-Fi 6, Low Earth Orbit, and any other high-speed wireless transport–in addition to traditional fiber and copper networking.

Zscaler ZTE delivers cloud and edge security services via 150 global points of presence. Zscaler's high-bandwidth, low-latency approach enables customers to carry out business securely while traversing networks whose level of trust varies from known, to uncertain, to unknowable. As customers move their workloads out to the far and deep edge, Zscaler cybersecurity moves with them.

Intel: Reference Architectures for Zscaler SSE

Zscaler's "Endpoint to Edge to Cloud" zero-trust solutions require hardware architectures that deliver high-performance computing to protect the most demanding enterprise workloads. Zscaler solutions deliver high-throughput packet processing with low latency and high determinism. Performance needs vary depending on the number of customer endpoints, the amount of traffic processed, and deployment locations. Servers that run virtualized and containerized SSE components have specific needs with respect to density, power delivery and cooling capabilities, and protections against physical tampering. Working with Zscaler, Intel developed the key hardware configurations necessary to implement a global SSE solution. This SSE reference architecture was based on Intel® Xeon® Scalable processors. The broad Intel Xeon processor portfolio enables Zscaler to deploy its SSE solution in a variety of form factors, from small, low-power platforms to high-capacity servers to address the specific workload requirements at each location. The broad industry support of Intel Xeon processors, from compilers and development tools to open-source libraries and research, simplified Zscaler's development efforts. For example, Intel® Software Guard Extensions (Intel® SGX) helps protect selected code and data by processing them in enclaves, allowing Zscaler developers to partition the application into trusted modules to help increase application security.

Standardizing on Intel architecture enabled Zscaler to take advantage of hardware acceleration features available with the Intel Xeon platform (see Fig. 2). Intel® QuickAssist Technology (Intel® QAT) accelerates authentication, encryption, digital signatures, and data compression while maintaining throughput and low-latency. The optimized Intel version of the Data Plane Development Kit (DPDK) is an open-source set of software libraries that accelerate packet processing by up to 10X–allowing Intel Xeon processors to be used for both packet processing and data processing, leading to simpler solutions that can scale with processor selection.

The latest versions of Intel Xeon Scalable processors have also added new artificial intelligence accelerators, collectively referred to as Intel® Deep Learning Boost (Intel® DL Boost). AI is rapidly becoming a key tool in security as the focus shifts

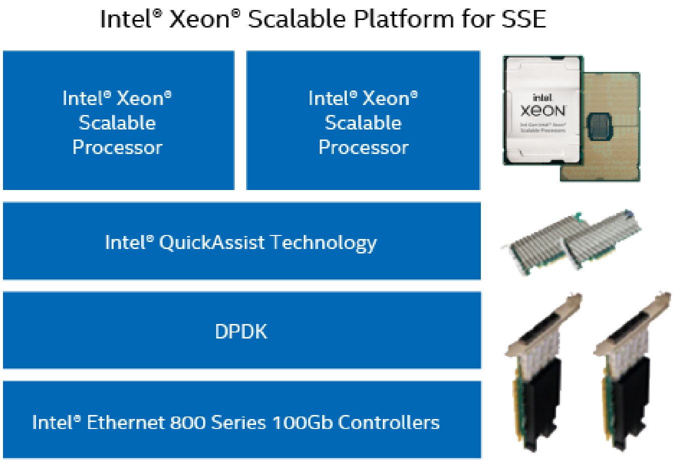
Intel® Xeon® Scalable Platform for SSE
Intel® Xeon® Scalable Processor
Intel® Xeon® Scalable Processor
Intel® QuickAssist Technology
DPDK
Intel® Ethernet 800 Series 100Gb Controllers
intel
XEON

Figure 2. Intel's SSE platform based on the Intel® Xeon® Scalable Processor, Intel® QuickAssist Technology for crypto and compression acceleration, DPDK for user space data plane acceleration and the Intel 100 GbE network interface card.

from looking for known vulnerabilities to looking for anomalous behavior that would indicate a new, yet-to-be-discovered attack. Zscaler is already evaluating these capabilities to develop new edge-based AI algorithms to enhance security for its customers.
Zscaler's collaboration with Intel also enables both companies to evaluate the impact of the yet-to-be-released next generation Intel Xeon Scalable processor family. The companies are evaluating how to ensure seamless migration to the new processor along with opportunities to incorporate new processor technologies to deliver superior value to customers.

Supermicro: SSE Hardware Solutions

The final stage in the SSE reference architecture was development of the physical platforms. For that, Intel and Zscaler turned to Supermicro, a global technology leader in hardware solutions for enterprise, cloud, AI, and telco/edge infrastructure. Supermicro developed two designs to support the wide range of processing tasks carried out by Zscaler SSE.

The standard design centers on a single-socket Intel Xeon processor, which offers up to 32 processor cores. The high-performance design supports dual-socket Intel Xeon processors, offering up to 64 processor cores for high-capacity, high-density implementations.

Both designs utilize the latest Intel® Ethernet 800 Series Controllers, supporting either 25 Gb or 100Gb networks. These high-performance Ethernet controllers include features to optimize performance based on workloads, from Application Device Queues that allow network traffic to be prioritized based on applications to support for Remote Direct Memory Access to transfer data with less overhead. As with other Intel acceleration technologies, the result is higher throughput while adding more security and intelligence to SSE implementations.

Better Together

The combination of Zscaler security software running on Supermicro hardware, optimized to take advantage of Intel Xeon Scalable processors, accelerators, and Ethernet controllers, results in a cutting-edge SSE solution delivered worldwide in support of securing enterprise business. This is especially important today as everyone braces for more cyberattacks by increasingly sophisticated and determined threat actors.

Better for You

This collaborative effort between Intel, Zscaler, and Supermicro delivers high-bandwidth, low-overhead, power-efficient SSE solutions to help you power and secure your unique Edge Compute needs.

For more information about Zscaler security solutions, please visit: www.zscaler.com

For more information about Supermicro hardware solutions, please visit: www.supermicro.com/en

For more information about Intel Xeon Scalable processors and other technologies, please visit: www.intel.com/xeon/scalable

Notices & Disclaimers

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.
Intel technologies may require enabled hardware, software or service activation.
No product or component can be absolutely secure.
Your costs and results may vary.
(1) Van der Meulen, Rob. "What Edge Computing Means for Infrastructure and Operations Leaders." Gartner, October 3, 2018, https://www.gartner.com/smarterwithgartner/what-edge-computing-means-for-infrastructure-and-operations-leaders

Printed in USA 351441-001EN Please Recycle 0922/DA/ICEL/PDF