Solution Brief

Network Acceleration
2nd Generation Intel® Xeon® Scalable Processors
Intel® Tofino™ P4-Programmable Switch ASICs
Intel® Stratix® 10 FPGAs

CASwell uses Intel Technology for Next-Generation Web Scale Server Switch

CASwell CAR-5056 server switch helps communications and cloud service providers manage high-volume cloud services by accelerating "middle box" functions like load balancing, application delivery control, and broadband network gateways using Intel® Xeon® Scalable processors, Intel® Tofino™ P4-programmable Switch ASICs, and Intel® Stratix® 10 FPGAs.

More Devices and Services = More Data

More connected devices, more cloud services, and more growth in consumer services such as streaming video, online games, and online AR/VR are driving internet traffic levels to new heights with no end in sight. Popular websites need to serve up millions of concurrent requests from clients and return the correct content quickly and reliably.

Impact on "Middle Box" Network Functions

This is impacting the performance of so called "middle box" network functions, which sit between the edge router and the application server providing specialized data packet processing, load balancing content for fast response, enforcing policies or providing security (see Figure 1). Middle box functions can accelerate more complex functionality, freeing up the host processors for revenue-generating workloads.

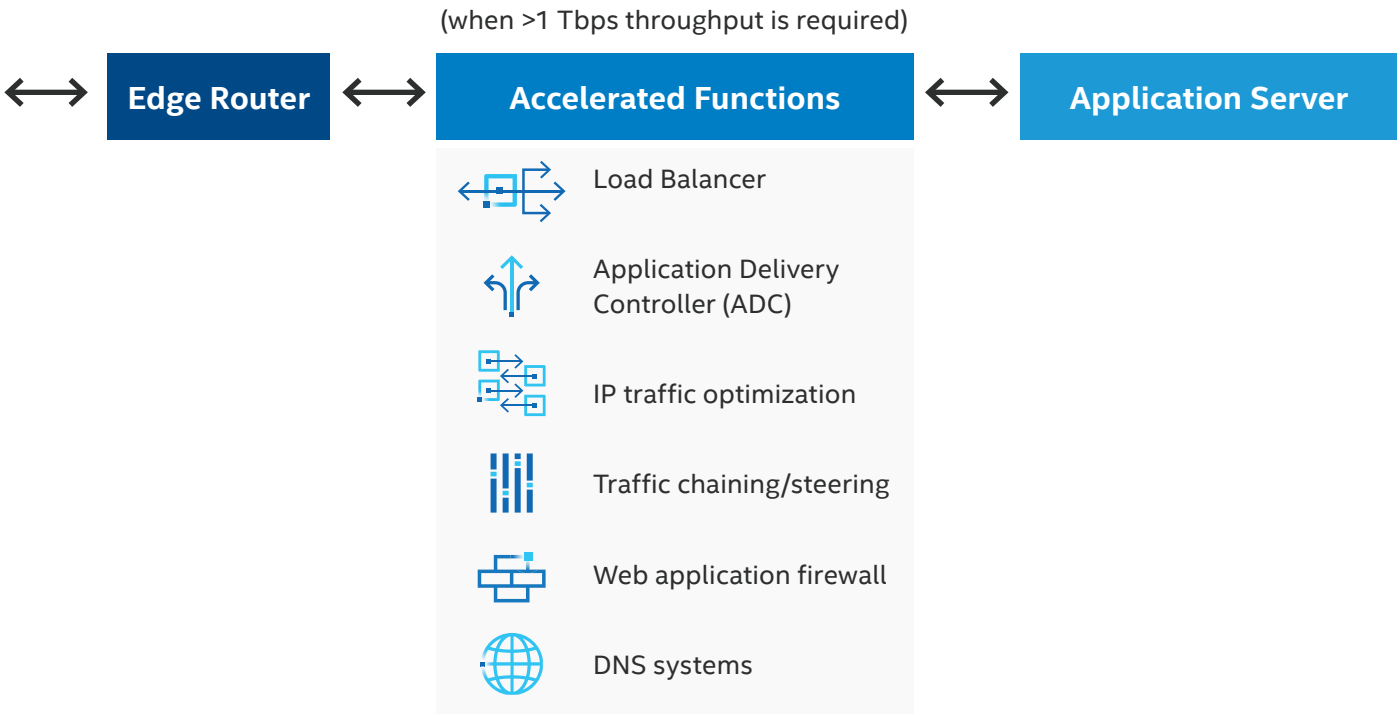

Figure 1. "Middle Box" network functions conceptual diagram with examples that can benefit from acceleration

Some of the more popular middle box applications include:

Load Balancer: Busy websites often use many scaled-out servers to process high data volumes. A load balancer tracks the processing workload on each server and directs data flows to the server that has bandwidth to process it (see Figure 2). This "traffic cop" routing mechanism balances the traffic across all servers to maximize speed and capacity utilization.

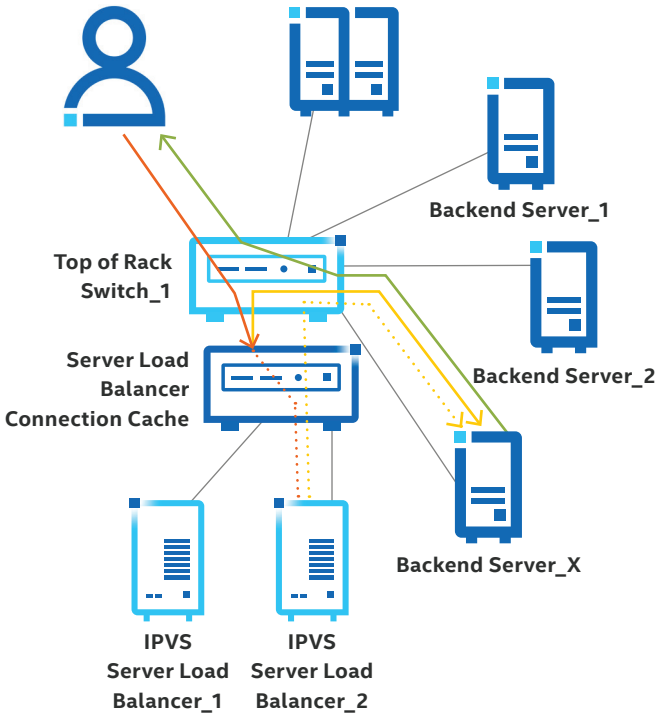

Figure 2. Example Load Balancer Architecture

Application Delivery Controller (ADC): This network function accelerates web traffic, provides rate shaping and SSL decryption, and serves as a Web application firewall. The ADC often also includes load balancing functionality.

Broadband Network Gateway (BNG): This network function sits at the edge of the wired IP telecommunications network and serves as the access point through which customers connect to the telecom broadband network. The BNG aggregates user sessions and is responsible for policy processing and for authentication, authorization, and accounting.

In addition to these popular network-adjacent applications, there are many other possible uses, including IP traffic optimization, traffic chaining/steering, web application firewall, DNS systems, and others.

In cases where extra high bandwidth (>1 Tbps) is required or in denser deployments, a server switch can boost the solution. Server switches offer both switching and compute in the same form factor. Server switches can benefit from Programming Protocol-independent Packet Processors (P4) programmability which enables data planes of the middle box applications to run in the switch fabric for very high throughput.

CASwell, an Intel partner, has developed the CAR-5056 server switch that is targeted at these applications.

A New Generation of Server Switch

The CASwell CAR-5056 is a 2RU-high server switch designed for cloud service providers (CSPs), communication service providers (CoSPs), and data centers providing web scale services (see Figure 3). The CASwell CAR-5056 is based on 2nd Generation Intel® Xeon® Scalable processors, Intel® Tofino™ P4- programmable switch ASICs, and Intel® Stratix® 10 field programmable gate arrays (FPGAs).

Figure 3. CASwell CAR-5056 server switch.

Processor Performance and Security

For its server functionality, the CASwell CAR-5056 features dual 2nd Gen Intel Xeon Scalable processors. This processor family provides the foundation for powerful data center platforms that create an evolutionary leap in agility and scalability. Disruptive by design, these innovative processors set a new level of platform convergence and capabilities in compute, storage, memory, network, and security. CSPs and CoSPs can now drive forward their most ambitious digital initiatives with a feature-rich, highly versatile platform. CASwell CAR-5056 utilizes processors with up to 28 cores and also makes use of the on-chip Intel® QuickAssist Technology (Intel® QAT) to accelerate encryption and decryption performance in cloud, networking, big data, and storage applications.

Memory and Storage Options

Memory is important for the middle box functions needed for high-volume cloud services by applications, so the CASwell CAR-5056 supports up to 10 DIMMs for a maximum of 640GB main memory per CPU. The server switch also supports flexible storage options with one 2.5-inch SSD and one mSATA SSD.

Flexible Networking with P4 Programmability and Visibility

For networking, the CASwell CAR-5056 utilizes the Intel Tofino P4-programmable Ethernet switch ASIC. Intel Tofino provides the CASwell CAR-5056 with 32x100 GbE ports, two 10/25 GbE ports for server connectivity, and 1GbE port for management. Because the Intel Tofino switch ASIC is P4 programmable, the CASwell CAR-5056 enables service providers to update their packet protocols, header fields, or encapsulation to meet new IP standards or network performance goals. The CASwell CAR-5056 also supports in-band network telemetry (INT) to collect real-time network congestion data so it can redirect data flows to reduce congestion.

P4 Programming Language

By using the Intel® Tofino™ P4-programmable Ethernet switch ASIC, the CASwell CAR-5056 offers data plane programmability using the Programming Protocol-independent Packet Processors (P4) programming language to control packet-forwarding data planes in networking devices. P4 is a domain-specific programming language with constructs that are optimized for network data forwarding. P4 is target-independent and so can be deployed on switch ASICs, SmartNICs, FPGAs, host computers, and software switches. The language enables field reconfigurability so network engineers can change packet processing methodologies after the switches have been deployed. For more information visit **P4.org**

Extended Capabilities with FPGA Modules

With Intel Tofino and P4 programmability, the CASwell CAR-5056 is able to host P4-based network functions such as load balancers, data security, routing, etc., running within the network fabric for fast processing and minimal latency. This ability can be amplified by Intel Stratix 10 FPGA modules that bring flexibility and support for stateful processes and more complex processing.

The CASwell CAR-5056 supports up to four Intel Stratix 10 FPGAs modules, which deliver innovative advantages in performance, power efficiency, density, and system integration. The ICs feature the revolutionary Intel® Hyperflex™ FPGA Architecture and Intel's patented Embedded Multi-Die Interconnect Bridge (EMIB) technology, the Advanced Interface Bus (AIB), and a growing portfolio of chiplets.

The FPGA modules use PCI Express 3.0 to connect to both of the 2nd Gen Intel Xeon Scalable processors and use 20 embedded Ethernet connections to connect to the Intel Tofino P4-programmable switch ASIC to meet the performance demands of high-throughput systems.

CASwell CAR-5056 in a Load Balancing Application

Layer 4 load balancing (L4LB) is a popular service in cloud and telco data centers to provide efficient utilization of server and storage resources by distributing client requests and network load more efficiently across multiple servers.

L4LBs have evolved from dedicated appliances to virtualized applications running on servers. As separate appliances, an L4LB can increase the network management burden. The CASwell CAR-5056 (see Figure 4) can effectively use its programmable switch and FPGA modules for P4-based cloud-scale L4LB services to meet the data flow needs of scale-out applications while also ensuring the highest availability and reliability of the connected websites and web services.

A P4-based L4LB application is available for the CASwell CAR-5056 that builds on the foundations presented in the SilkRoad paper [**PDF download**].[1] The L4LB application features multiple redirection and encapsulation options, smart caching, data plane entry management for extra-large session table, and DDoS protection.

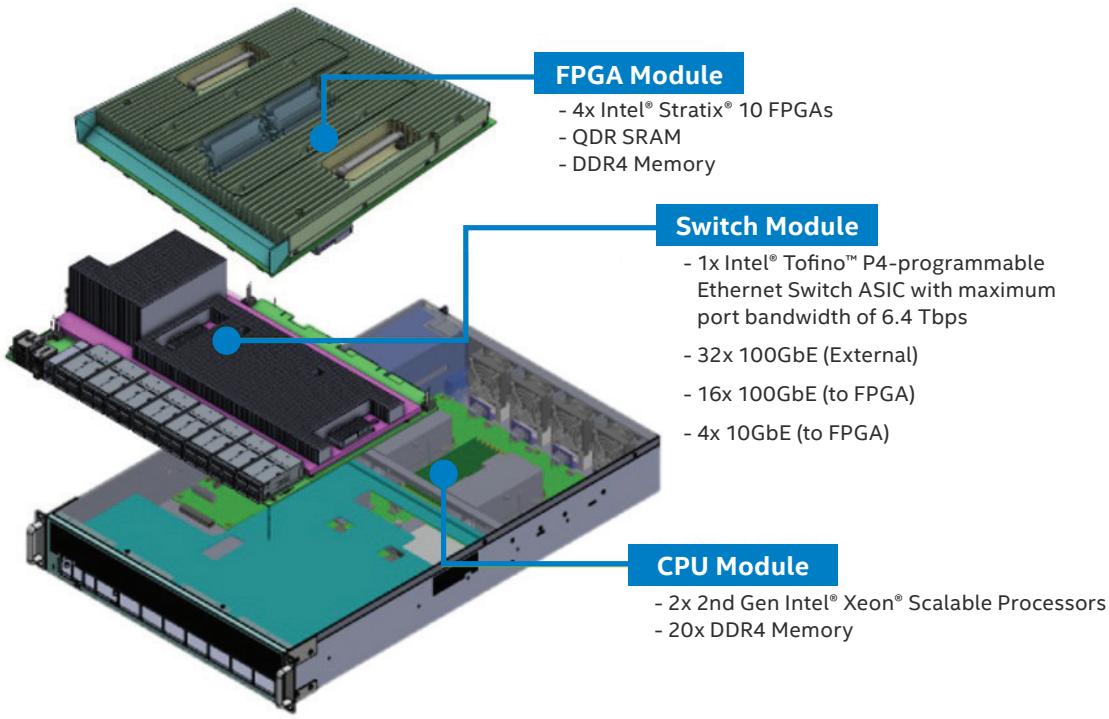

Figure 4. Interior of CASwell CAR-5056 with technology important to L4LB called out.

The Intel Tofino P4-programmable switch ASIC, capable of operating at up to 6.4 Tbps, provides the data plane processing for L4LB. Intel Stratix 10 FPGAs provide additional memory resources for tables and buffers and additional logic to optimize look up functions and implement use case-specific features. These FPGAs enable eXtra Large Tables (XLT) to support hundreds of millions of table entries, and eXtra Large Buffers (XLB) to provide very large buffers - up to tens of Gigabytes.

Conclusion

Middle box network functions are critical to effective web service performance at cloud scale, but they are also challenged by the high levels of network throughput that must be supported in a cloud data center. The CASwell CAR-5056 server switch is designed specifically to support P4-based middle box network functions. The Intel Tofino P4-programmable Ethernet switch ASIC allows the L4LB application to run in the network fabric with very high throughput. Modules using Intel Stratix 10 FPGAs provide expanded memory and table sizes to enable greater flexibility and support for more complex processing. With the CASwell CAR-5056, CASwell is delivering a next-generation switch server through innovation that adds significant value for CSPs and CoSPs.

Learn More

CASwell

CASwell CAR-5056

Intel 2nd Generation Xeon Scalable Processors

Intel Tofino P4-programmable Switch ASIC

Intel Stratix 10 FPGA

[1] https://rmiao.github.io/publications/silkroad-rui.pdf.

0821/BH/MESH/PDF 348068-001US