Product Brief

Application Security and Delivery | Data Center Networking

Intel® Application Security and Delivery Acceleration Kit

The Intel® Application Security and Delivery Acceleration Kit, one of the Intel® Tofino™ Expandable Architecture Platform Acceleration Kits, provides all of the components required to implement a fully operational solution for more secure and reliable delivery of high-bandwidth application traffic in the data plane

Managing vast amounts of data flow in a secure network environment is becoming increasingly difficult as the volume of traffic grows rapidly to accommodate the expanding portfolio of new devices and applications like 4K video streaming, the Internet of Things (IoT) and work-from-home teleconferencing. Providing unlimited internet access for all users in this environment requires excessive heterogeneity that can make hostile traffic difficult to detect.

The Intel® Application Security and Delivery Acceleration Kit is a fast and easy-to-implement solution to efficiently deliver terabits of application traffic to the right servers while providing an effective means of security and protection against attacks. The Intel Application Security and Delivery Acceleration Kit leverages the power of Intel® Tofino™ Expandable Architecture and the acceleration capabilities of Intel® FPGAs and/or Intel® Infrastructure Processing Units (Intel® IPUs) to augment the Intel® Tofino™ Intelligent Fabric Processor (Intel® Tofino™ IFP) functionality for network solutions that are fast, more secure and reliable (see Figure 1).

Functionality

The Intel Application Security and Delivery Acceleration Kit provides all of the necessary software components for quick and easy implementation of reliable network traffic distribution solutions while helping to protect against attacks, reducing development complexity, accelerating time to deployment and lowering TCO. The Intel Application Security and Delivery Acceleration Kit is an extension of the Intel® P4 Studio Software Development Environment (Intel® P4 Studio SDE). It consists of reference implementations of three major functions that are combined for efficient acceleration of application security and delivery:

- The Intel Tofino IFP P4 reference profile.
- The control plane software for managing and controlling the execution of the Intel Tofino IFP and Intel FPGA data plane.
- The Intel FPGA implementation of the extra-large lookup table for an extremely large session table.

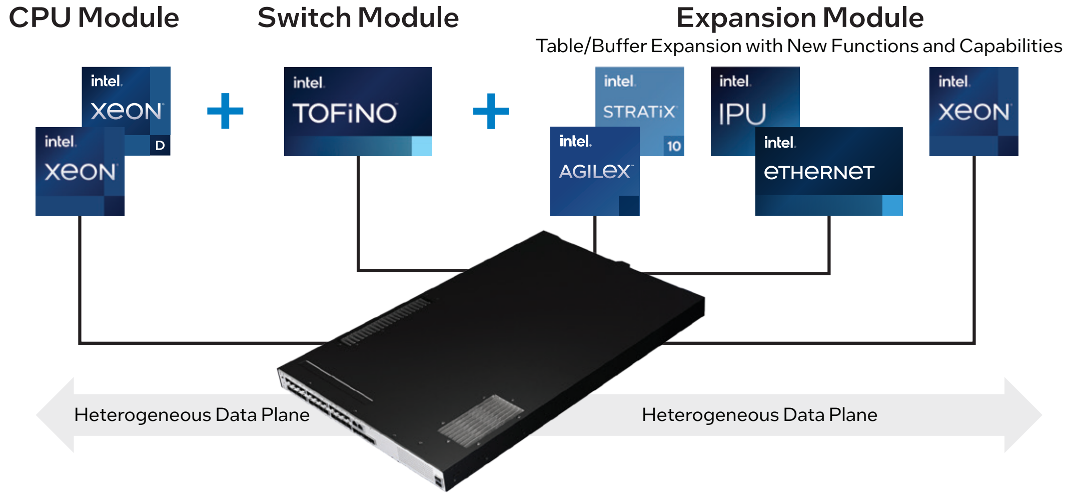

Figure 1. Intel® Tofino™ Expandable Architecture combines Intel® Xeon® processors, Intel® Tofino™ IFPs, and Intel® FPGAs and/or IPUs into a single system.

Intel Tofino IFP P4 Reference Profile

Hardened L2/L3 Switching and Routing

Data plane implementation in Intel Tofino IFPs includes customizable L2/L3 switching and routing functionality based on a hardened Switch P4 profile that is secure, reliable and switch abstraction interface (SAI)-controllable, making it possible to take advantage of broad community support and existing projects (e.g., SONiC OS) and tools.

Server Load Balancing

The L4 server load balancing directs the distribution of network traffic. It operates at a transport layer (e.g., fourth) of an ISO/OSI model and decides where every packet gets forwarded. It uses a flow cache to store a certain number of sessions including instructions on how to process the session packets. Packet operations are user-defined through the P4 code, allowing custom functionalities (e.g., different load balancing modes).

Tiered Cache Hierarchy

The Intel Application Security and Delivery Acceleration Kit implements a tiered cache hierarchy to benefit from the extremely fast lookup memories available in Intel Tofino IFPs. When there is a miss in the first-level cache due to limited capacity, the second-level cache seamlessly leverages the Intel® Tofino™ Expandable Architecture to perform lookup requests in larger memory resources.

TCP SYN Flood DDoS Detection

The SYN proxy algorithm counts the number of TCP SYN packets that are received by the data plane during a certain time period and reports the count to the control plane software. SYN flood detection logic implemented in the control plane program maintains the exponentially weighted moving average (EWMA) and monitors whether the EWMA goes above a defined threshold. If it is above the threshold, data plane mitigation logic is enabled (see Figure 2).

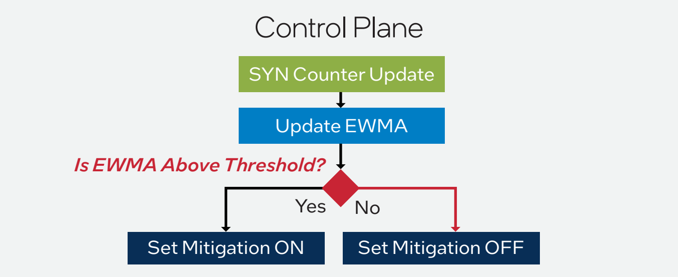

Figure 2. TCP SYN flood control plane logic.

TCP SYN Flood DDoS Mitigation

Once SYN flood mitigation logic is enabled upon detecting a SYN flood attack, the SYN proxy algorithm uses a SYN cookie challenge to authenticate the source of the communication. After receiving a correct acknowledge (ACK) response from the source of the communication, the source IP is whitelisted, and the TCP connection RST packet is sent (see Figure 3 for data plane algorithm depiction).

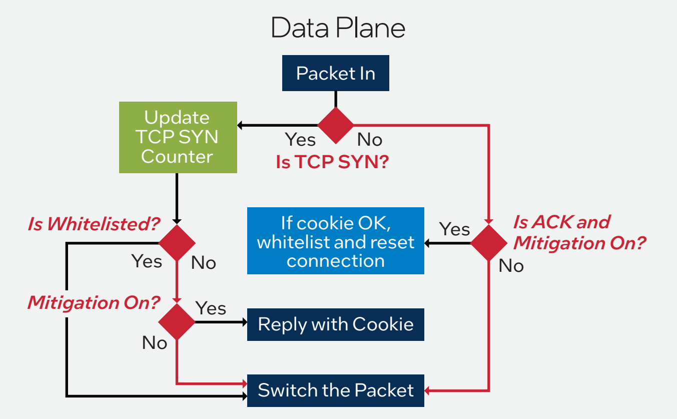

Figure 3. TCP SYN flood data plane logic.

Intel® FPGA Implementation of the Extra-Large Lookup Table

The FPGA provides an extra-large exact-match table implementation that uses tens of GBs of dynamic memory supporting up to hundreds of millions of sessions with speeds up to 4.8 billion lookups per second. Depending on the FPGA used, the dynamic memory is either:

- Internal to the FPGA package using high bandwidth memory (e.g., HBM2 in Intel® Stratix™ 10 FPGA), or
- External to the FPGA using DDR memory

The lookup key is provided from the Intel Tofino IFP to the Intel FPGA as part of the packet data, and the lookup result is then sent from the FPGA to the Intel Tofino IFP, again as part of the packet. The FPGA includes support for very high-speed table management through in-band data plane packets, taking advantage of the full speed of the dynamic memory without any overhead of additional software layers (see Table 1).

Table 1. Extra-Large Table Lookup Rates Based on DDR4 and HBM2 Configurations

Parameters	4x Intel® Stratix™ 10 GX FPGAs with 2x DDR4 per FPGA	4x Intel® Stratix™ 10 MX FPGAs with 8 GB of HBM2 per FPGA
Memory capacity	128 GB	32 GB
Table size[a]	256M session entries	128M session entries
Lookup rate[b]	Up to 600M lookups per second	Up to 4.8B lookups per second
Data path table update rate[c]	Up to 4M updates per second	Up to 32M updates per second
Cost[d]	$	$$

[a] Table size assumes 32 bytes per entry and an optimization that trades off between capacity and lookup performance that leads to per-entry overhead.

[b] Lookup rate assumes even distribution of flows across the universe of possible entries.

[c] Achievable update rate under no or minimal load. The actual update rate during standard operation depends on the number of lookup requests processed by the extra-large table.

[d] Indicative comparison. Actual pricing depends on customer volume commitments.

Learn More

You may find the following resources helpful:

- Intel® Application Security and Delivery Acceleration Kit Solution Brief
- OpenBNG Solutions from Intel
- Intel® Broadband Network Gateway Acceleration Kit (Intel® BNG Acceleration Kit) Product Brief
- Let Your Networks Soar with Intel® Tofino™ Expandable Architecture White Paper
- Intel® Tofino™ Intelligent Fabric Processors (Intel® Tofino™ IFPs)
- Intel® Xeon® Scalable processors
- Intel® FPGAs
- P4 open source programming language

For more information, contact your Intel representative and visit **intel.com/fabric**.

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. Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation. Intel, the Intel logo and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others. 1222/TKOE/KC/PDF 751816