White Paper

NEXSEC, Intel Collaborate on uCPE 2.0 White Box Server Solution

The uCPE 2.0 described in this paper is a next-generation universal customer premise equipment (uCPE) white-box reference design based on Intel Atom® P5000 SoC and Intel Atom® C5000 SoC. It incorporates the advantages of pin-to-pin compatibility of the two SoC families, scaling from four cores to 24 cores on a single hardware form factor platform, to support service deployment demands across a wide range of networking performance and compute power levels. These include Intel's uCPE 2.0-enriched uCPE solution portfolio with continued platform optimization journey and innovated uCPE usages for software-defined wide area network (SD-WAN), edge, network security, and SRv6 gateway.

Authors

Frank Liu

Andy Ng

Dekun Li

Dong Wang

Yulong Pei

Intel Corporation

Hadwin Liu

NEXSEC

Introduction

A successful enterprise digitalization solution can be measured by four indicators:

- Is it a fully software programmable platform?
- What can it connect?
- Does it provide insightful data to drive business efficiency?
- Is it intelligent and automated?

An enterprise network is the foundational infrastructure of an enterprise's digitalization journey. It plays a vital role: connecting all digital enterprise assets, carrying enterprise data, empowering 24x7 networking services from branch data center to cloud. Enterprise network systems have become increasingly complex to address all these challenges.

The uCPE is an on-prem edge/branch office server platform that combines advanced network technologies and leading compute power that supports the edge applications possibly installed and running on the same platforms. These servers benefit from the performance delivered by Intel Atom SoCs in addition to highly integrated Intel network platforms. The uCPE is a software-defined edge server that is capable of hosting virtualized network functions (VNFs) or cloud-native network functions (CNFs) such as routing, SD-WAN, firewall, deep packet inspection (DPI) and enterprise edge applications in a single platform.

This simplifies the end-to-end connectivity between the enterprise network and the cloud, drives innovation in application development and improves quality of experience for the end users. A common CPU architecture featuring a consistent software environment and rich platform options delivered by Intel processors, application developers benefit from one-time development of an application that can run anywhere.

Intel and NEXSEC*, a leading Intel® Network Builders uCPE ecosystem member, have jointly launched uCPE 2.0, a white-box server platform based on the latest Intel Atom P5000 and Intel Atom C5000 processors that provides scalability for a variety of applications thanks to its wide range of core counts (from four to 24 depending on the SKU) in addition to advanced networking using 4G, 5G, and Wi-Fi 6.

Table of Contents

1. uCPE 2.0 White Box Reference Design

uCPE 2.0 is designed to host extensive workloads and deliver high-performance network connectivity. It features the Intel Atom P5000 / Intel Atom C5000 processor series and various network interfaces in a single compact system. In response to the growing demand for wirelessly connected, flexible, and space-saving edge appliances, uCPE 2.0 was developed to be widely deployable and to deliver excellent performance.

uCPE 2.0 devices are designed (see Table 1) to be deployed in 5G, edge compute, SD-WAN, or Secure Access Service Edge (SASE) applications. These devices also provide ample compute power to run virtualized apps at the edge, eliminating the need to route the data back to the cloud for processing.

FEATURE	DESCRIPTION
CPU	Intel Atom C5000 / Intel Atom P5000 processor, with 4-24 cores
Form Factor / Dimension	Desktop: 260x300x44mm Rackmount: 430x440x44mm
Memory	Support 2 DIMM sockets for DDR4 2933 ECC RDIMM, UDIMM Max. memory capacity, up to 32GB
Network Support	4x 10GbE SFP28 (integrated in CPU) 6x 1GbE RJ45 (4xGbE integrated in CPU+PHY) 2x 1GbE by Intel i211 Two pairs copper LAN bypass support
M.2 Module Support	2x M.2 sockets, supporting 5G/LTE/Wi-Fi 6/SSD/AI modules ▪ M.2 3580 B+M key socket (Low profile type) ∘ PCIe Gen3 x2 ∘ PCIe x1+ USB3.0 x1 ∘ SATA for SSD x1 ∘ SIM A Slot (default front access) ▪ M.2 5858 B key socket (high profile type) ∘ PCIe x1 + USB3.0 x1 ∘ IM slot (default internal access)
PCIE Slot Support	One PCIe x8 slot Riser card with configurable HSIO for 1x PCIe x8 or 2xPCIe x4, Support 2 LAN modules, up to 16GbE/4x10GbE Ports Support direct plug PCIe x8 FHHL card HSIO configurable, support up to 8x SATA ports Two OCulink CNN x2 for PCIe x16, when Intel Atom P5000 CPU is installed
Storage Capacity	2xSATAIII Ports On board 32GB eMMC
Front I/O Ports	1x DC-IN 1xPWR Button 2xUSB 2.0 Ports 1x console port (RJ45) 4x10/25G SFP28+ ports 6x1G RJ45 ports 1x SW programmable button 1x Micro SIM slot 6xLED ▪ PWR LED(Green) ▪ 5x programmable GPIO LED include below event ▪ System alert user define ▪ Cloud status, user define ▪ Wi-Fi Status ▪ LTE status (defect GPI O thru LTE module) ▪ LAN Bypass
PSU	Desktop: 120Watts PWR adapter Rackmount: 120watt PSU

Table 1. uCPE 2.0 Main Features

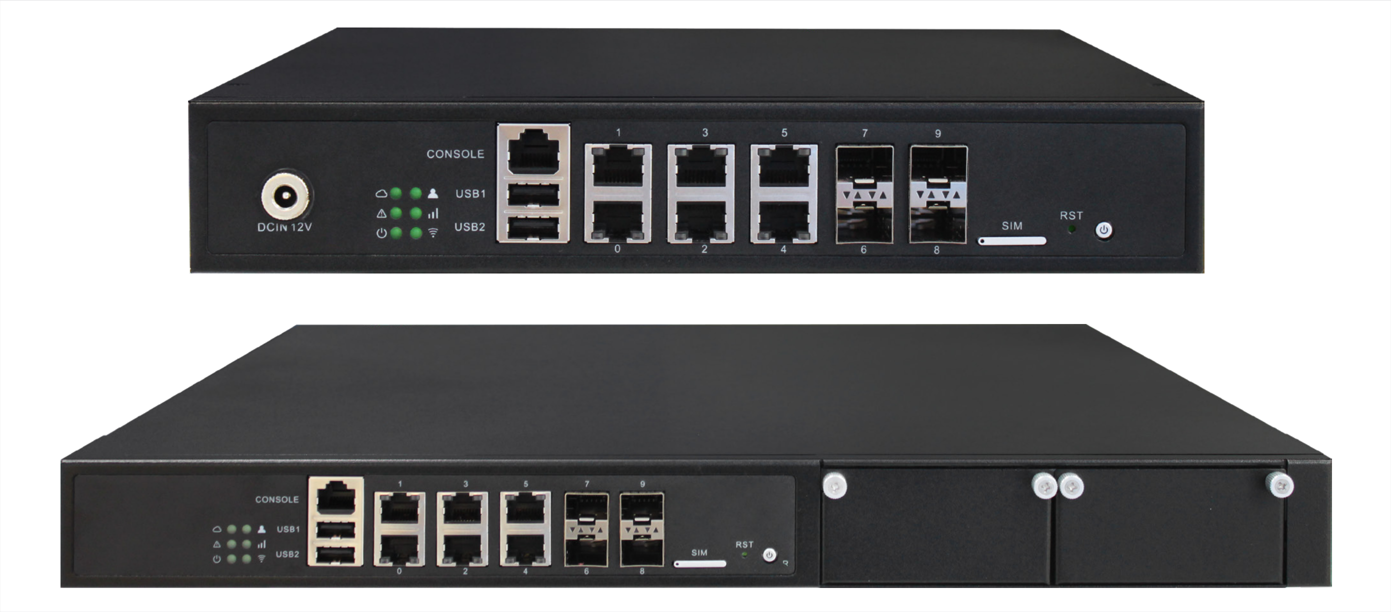

Figure 1. Two models of the uCPE 2.0 white box product; bottom version has two expansion bays.

2. Technologies Implemented

The uCPE 2.0 White Box reference implementation (see Figure 1 and Figure 2) takes full advantage of the integrated features of the Intel Atom P5000/ Intel Atom C5000 SoCs, including the following features:

- Intel® QuickAssist™ Technology (Intel® QAT) accelerates encryption and compression; the platform can support up to 100 Gbps cryptography operation.
- Integrated Intel® Ethernet provides up to 100 Gbps throughput, provides up to 8 integrated Ethernet ports with link speed options from 1GbE to 100GbE.
- Intel® Dynamic Load Balancer™ (Intel® DLB) is a hardware managed system of queues and arbiters connecting producers and consumers. It enables pipelined packet processing models for load balancing and packet queueing. Intel DLB can be used to accelerate traffic distribution and bandwidth management in multi-core Intel architecture.

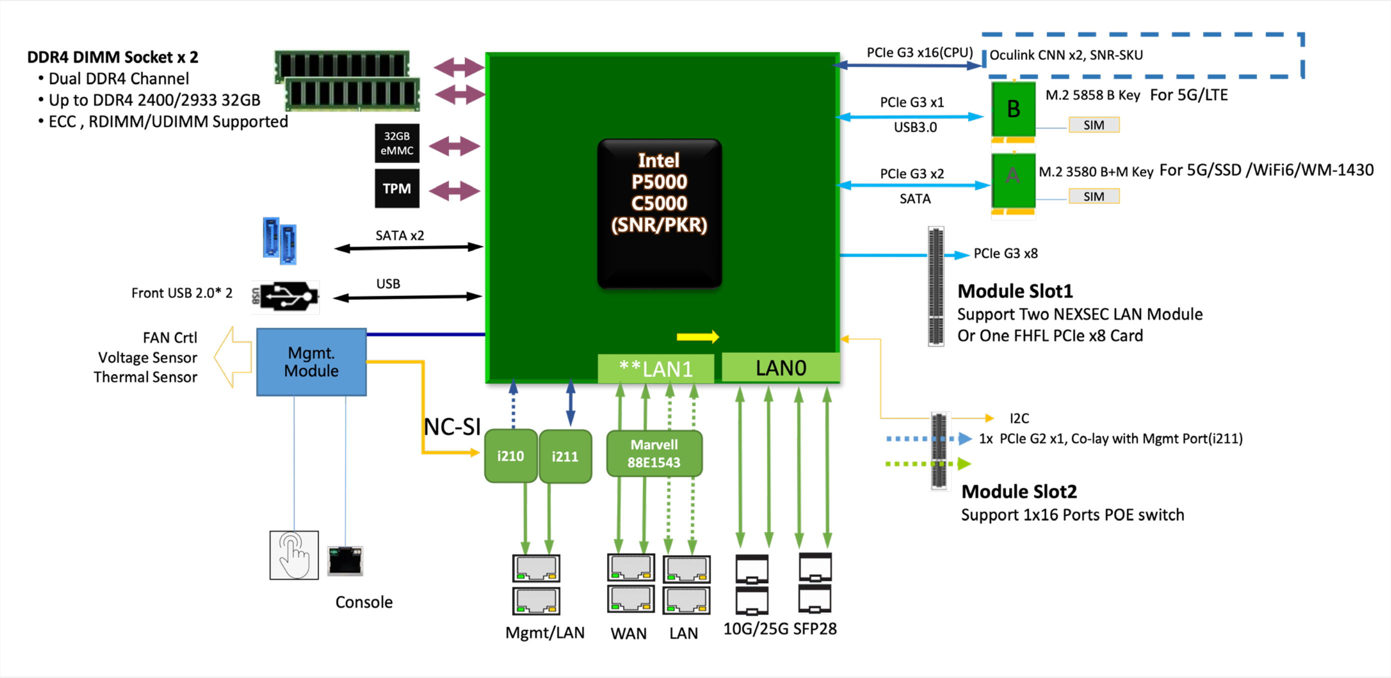

Figure 2. uCPE 2.0 Block Diagram

NEXSEC implemented numerous firmware features to enhance uCPE 2.0 White Box usability.

1. Firmware Level OS Fail-Over Design

Deployed uCPE systems experience frequent OS upgrade requests. Sometimes the system may be damaged by an unpredictable situation without any booting message. Then a manual inspection will be required, causing system downtime, and increasing operating costs. Therefore, optimizing the automation of operation and maintenance processes becomes an important feature.

The uCPE 2.0 provides a firmware-level OS recovery design. The system can autonomously carry out system failover and reset-to-factory-default functions without the intervention of OS. It also can roll back the OS to a previous workable system. As seen in Figure 3, the technical scheme is implemented by the firmware, BIOS and MCU / CPLD. To trigger the rollback, the user provides a boot policy and storage partition request. The failover mechanism then processes the change automatically.

2. Software Definition Lite-Manageable Hardware

uCPE 2.0 is a universal server but the field application may request new system status indicator, I/O control process or healthy monitor condition setting. It will be hard to implement these changes using a standard super I/O chip. To integrate a baseboard management controller (BMC) chip increases cost making the system uncompetitive in entry-level uCPE segment. NEXSEC has utilized the on-board microcontroller (MCU) function and capability to be a "lite-BMC" to offer super I/O function and programmable features to provide user-defined function include healthy monitor / control, firmware / OS fail-over, system status indicator and I/O control.

3. Lite Fail-Over WWAN 5G/LTE Support

uCPE2.0 provide two M.2 sockets with dual SIM slot that provides support for fail-over using the 5G/LTE wireless WAN (WWAN) module. But the M.2 socket may also be used for other modules like AI accelerator or a DPU module depending on the application. The NEXSEC uCPE2.0 lite fail over WWAN design supports one 5G/LTE module with intelligent SIM recovery mechanism. When the primary connection is lost, the system switches to the SIM signal automatically without any manual intervention. A fast switching dial-up mechanism is provided to support dual 5G/LTE link services cost effectively.

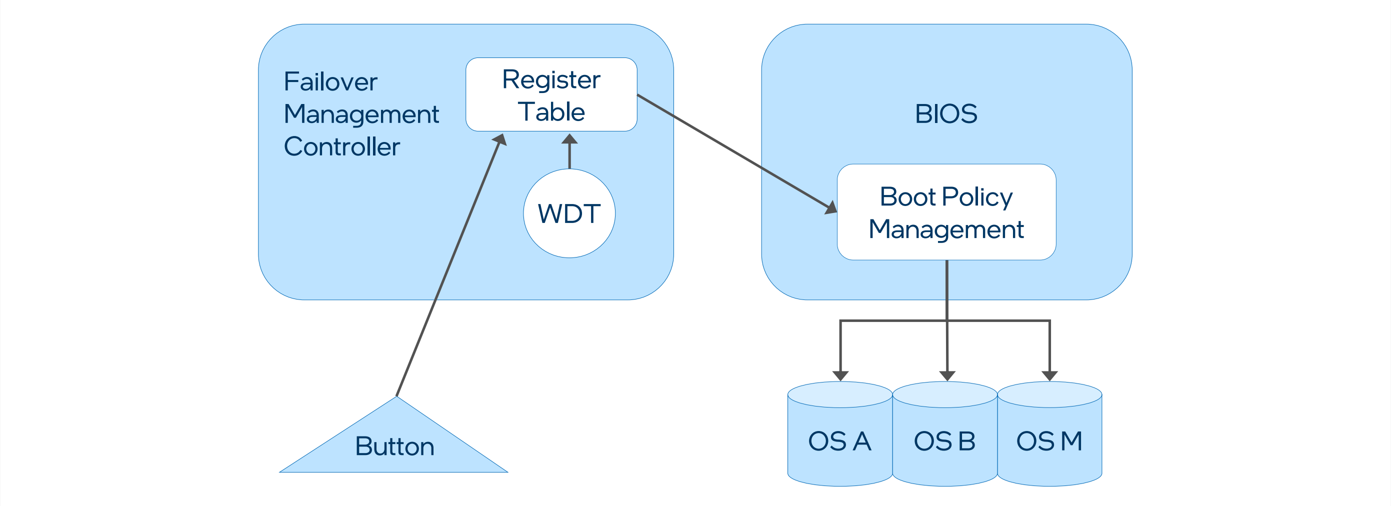

Figure 3. Fail Over Overview

3. uCPE 2.0 Use Cases

3.1 SD-WAN & SASE

3.1.1 SDWAN & SASE Overview

Enterprise applications are deployed in the cloud, edge and enterprise branch office to provide the lowest latency performance for employees. However, this network design provides more attack vectors for hackers. In addition, many companies have a bring-your-own-device (BYOD) policy that allows employees to use their own personal devices that may not be as secure.

SD-WAN addresses the challenges of branch offices connecting back to the corporate data centers and cloud by using software-defined networking to apply policies to data flows to ensure they use the right network. SD-WAN can simplify network management complexity through centralized and cloud-based management and enhance enterprise application experience through real-time traffic analytics and WAN optimization technology.

SD-WAN running on a uCPE 2.0 server accelerates enterprise applications through application recognition and bandwidth management and facilitates enterprise connectivity to the cloud through overlay network technology.

An SD-WAN uCPE has compute capacity to run other virtualized security software and the SD-WAN policies can direct the data flows through these security apps. SASE is a solution to help improve security for increasingly remote users to the enterprise networks, with built-in SD-WAN as the basic network infrastructure.

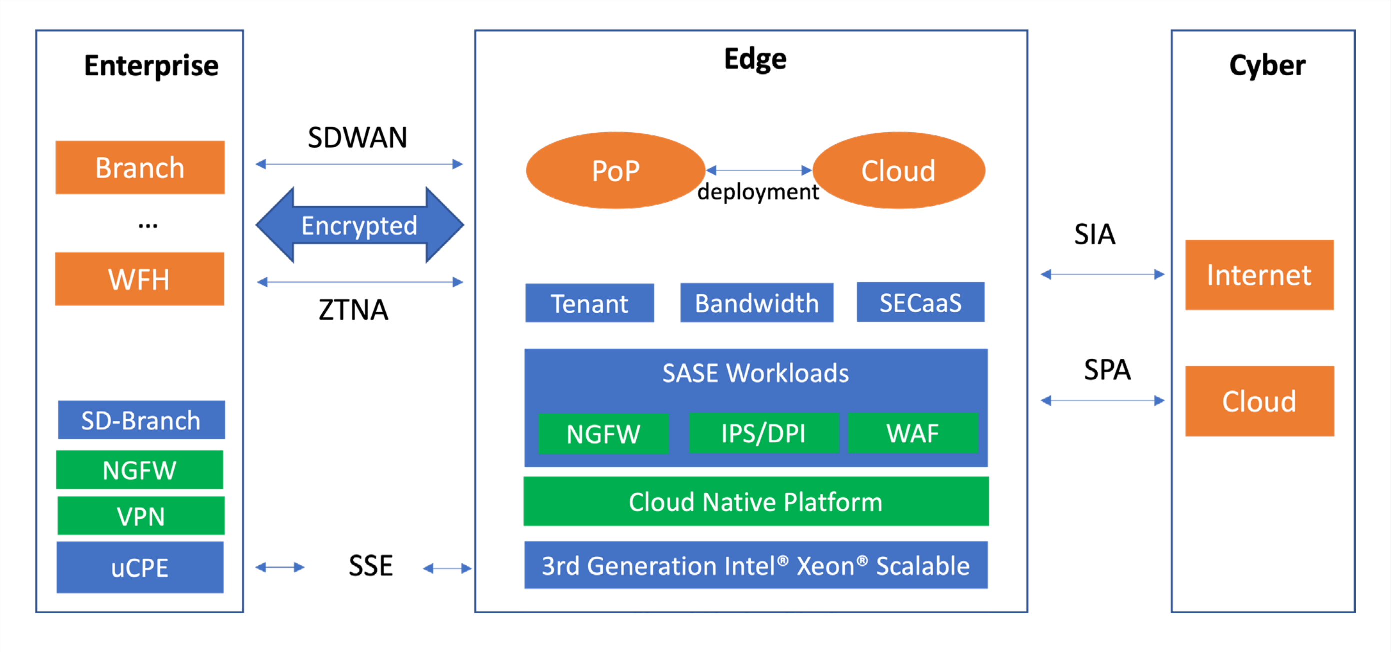

Figure 4. End-to-End SASE Reference Architecture

SASE transforms existing security architectures that are designed for protecting enterprise data and applications to work in hybrid cloud and edge computing environments. It combines security as a service or secure service edge (SSE) and network as a service (NaaS), providing high-performance tunnel connections and traffic scheduling in resource pools. Because multiple security functions run as micro services in the same Kubernetes cluster/node, a SASE platform can inspect the incoming traffic packet by packet through firewall, DLP or anti-virus functions without the need for multiple decryption and encryption passes. This improved security provides a platform for protecting remote users and corporate resources at the edge, corporate data centers, and in the cloud (see Figure 4).

The SASE reference architecture can be customized based on deployment, tenant management, security architecture, security service agility requirements. The SASE solution is deployed in a hybrid cloud environment, i.e., either a PoP, data center and/or public cloud, to enhance the access security for enterprise tenants who have branch offices distributed worldwide. Advanced routing such as the nearest network path is used to optimize network performance and to improve service efficiency and scalability.

New cloud native technologies are sought for SASE security architecture because they require much less resource consumption when hosting multiple workloads. Microservice and service function chains can improve service agility, service mesh can be used to scale and collaborate among multiple PoPs data centers and clouds. To support flexible office based or home-office work, SASE uses SD-WAN to provide security tunnels between enterprise branch and SASE edge locations. uCPE 2.0 SASE uses software-defined perimeter (SDP) with zero trust network access (ZTNA) to provide access services for work from home users.

Once traffic enters the SASE edge, series of security workloads such as next generation firewall (NGFW), deep packet inspection (DPI), intrusion prevention system (IPS) or web application firewall (WAF) are scheduled to scan network traffic based on tenant-subscribed services, traffic types, policies, and priorities. Finally, allowed traffic is redirected to either the Internet through secure internet access (SIA) service, or to public / private cloud if the access is from enterprise applications through secure private access (SPA)) service.

Security service edge (SSE) is the security part of comprehensive SASE solutions. NGFW is a typical SSE function deployed at the enterprise edge, providing software-defined branch (SD-Branch) baseline features, and access to more cloud-based security features such as DPI, WAF, data leakage protection (DLP), etc. Servers based on 3rd generation Intel® Xeon™ Scalable processors have the performance to support increased network and security workload generated by SSE platforms that host security workloads at the SASE edge. This whitepaper describes an example application where 1TB IPsec throughput was achieved: https://builders.intel.com/docs/networkbuilders/3rd-generation-intel-xeon-scalable-processor-achieving-1-tbps-ipsec-with-intel-advanced-vector-1617435344.pdf

3.1.2 SD-WAN with Intel® QAT and SMx acceleration

Tunneling is the technique used to implement SD-WAN. SD-WAN tunnels provide an overlay network on top of the transport underlay network. There are many tunnel implementation options such as VxLAN, IPsec, SSL, L2TP, GRE and combinations are used by SD-WAN, but IP security (IPsec) tunneling is the most important one. IPsec is defined by IETF RFC6072. In site-to-site tunnel mode, all user IP packets are encrypted and authenticated by IPsec and transported over an unsecured WAN such as a broadband connection to the Internet. There, the highly encrypted packets are transported to their destination. IPsec performance is one of the most important key performance indicators (KPI) to measure when calculating the efficiency of an SD-WAN solution.

In China, IPsec is replaced by the SM3 cryptographic hash function that is used in the Chinese National Standard. It's used as an authentication algorithm applied to complete IPsec encapsulating security payload (ESP) packets. The SM4 is block cipher symmetric algorithm used to replace Advanced Encryption Standard (AES) in the Chinese National Standard. It's used to protect the original packets in an IPsec stack. Both are mandatory for compliance with the Chinese Commercial Cryptography Scheme. The Intel Atom P5000 processor with the integrated Intel QAT supports both the SM3 and SM4 standards (see Fig 5).

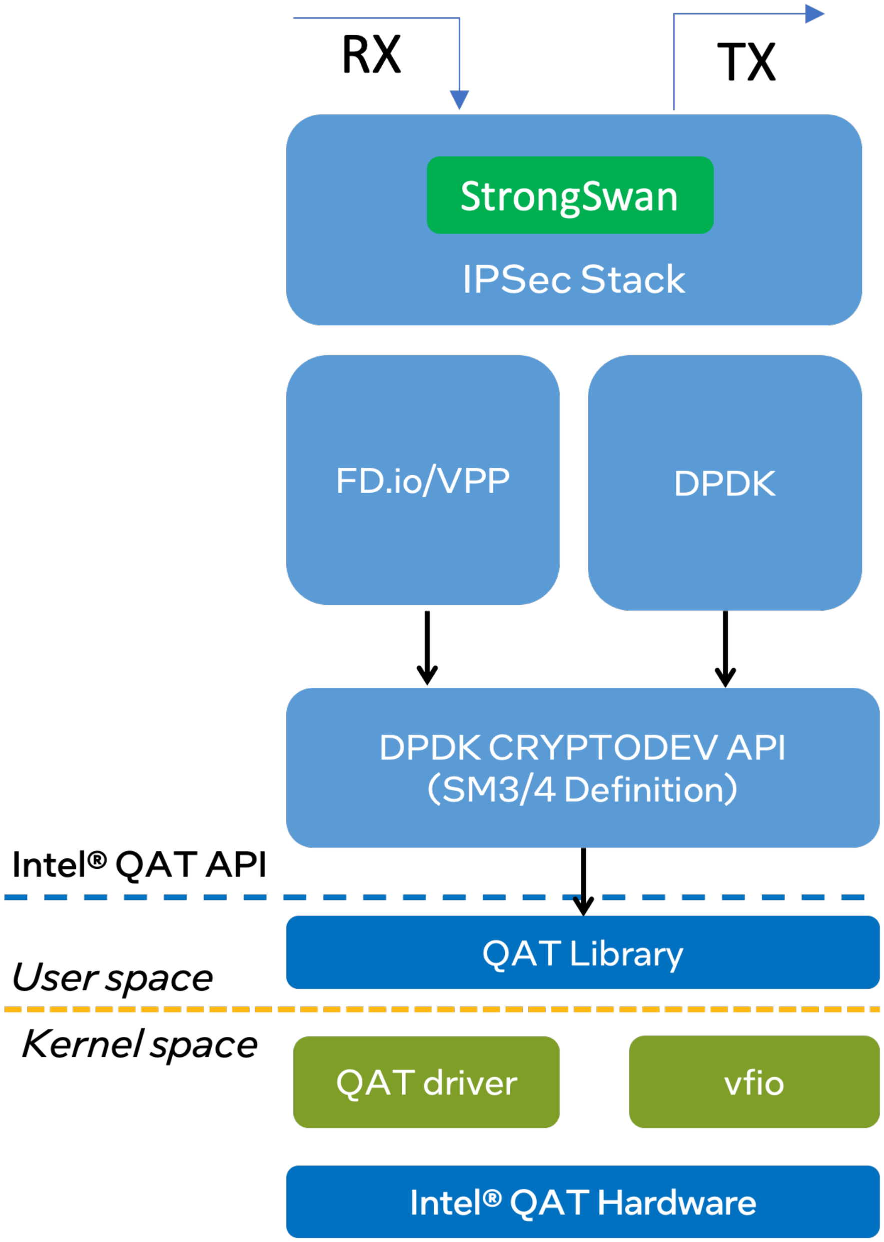

Figure 5. IPsec with SMx Acceleration with Intel QAT

3.2 Segment Routing IPv6 Gateway

3.2.1 Segment Routing IPv6 overview

Segment Routing IPv6 (SRv6) is used for large-scale networking and embeds a segment routing extension into an IPv6 header. The SRv6 application for uCPE 2.0 uses a 128-bit IPv6 address as the SR segment ID (SID) for network simplification, and path traceability.

It's designed to work with SDN technology to enable a programmable network that meets the requirements of emerging services such as 5G, cloud network collaboration and ubiquitous connectivity. G-SRv6 has been introduced by a leading Chinese mobile operator as carrier network standard through header compression of SRv6 packet as one of the key strategies of its future network.

3.2.2 SRv6 Packet Processing with DPDK and VPP

The Vector Packet Processor (VPP) data plane software has an up-to-date implementation of SRv6 network programming. The open source Data Plane Development Kit (DPDK) plays the role of user-space high performance network I/O. Network segment routing changes the way packets are forwarded on the network and enables network operators to better control the path of packets. The core of the SR protocol is to carry the routing information in the extension header of IPv6 packets, so that there is no need to maintain any state when it is forwarded in SRv6-supported devices. This extension header is the Segment Routing Header (SRH).

An SRH contains an ordered list of IPv6 addresses, and each IPv6 address indicates a next-hop segment address. The active segment is the IPv6 destination address of the packet. The segments left field of SRH indicates the next segment to be processed. If network packets are forwarded from SRv6-supporting devices, the segments left field of SRH is updated and the address is copied to the destination address of the IPv6 packet header.

On the current device node that supports SRv6, the available segment is called "local SID." The "local SID" is associated with a processing function on the local node, which includes an advance to the next SID in the SRH, or some other user-defined behaviors / functions (e.g., "END.DX6" that stands for "endpoint with decaps and IPv6 cross-connect"). This information is stored in the "My SID Table" entry. Each entry of the "My SID Table" indicates the function associated with the local SID and its parameters.

As mentioned in the previous section, SRv6 is widely used in many scenarios. As shown in Figure 6, the most typical application scenario is a layer three virtual private network (L3VPN).

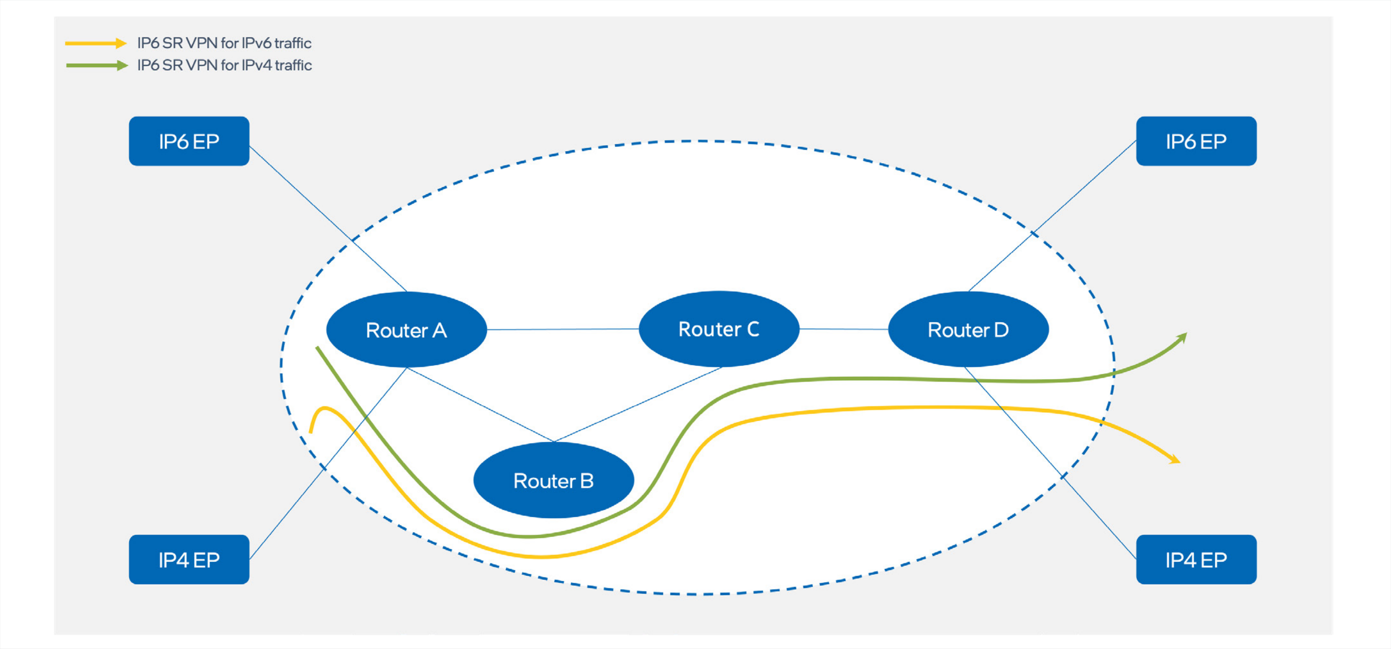

Figure 6. L3VPN use-case topology

3.2.3 Performance Analysis

As shown in the test results below, VPP-SRv6 has better performance on the uCPE2.0 device with Intel Atom P5000 / C5000 processors than when run on servers using older Intel Atom C3000 processors. For customer applications, the SRv6 protocol stack typically runs on top of user-space DPDK, with a VPP plug in and protocol stack. All SRv6 references mentioned in this paper refer to the L3VPN topology described in Figure 6, which is the simplest model in a real customer network.

3.2.3.1 Performance test setup

The tests were conducted using the FD.io CSIT test method (see Figure 7), in which two sets of uCPE 2.0 devices (with Intel Atom P5322 processors) act as DUT1 and DUT2, one server with an Intel® Xeon® 8280M processor @2.7GHz acts as traffic generator and performance collector. When compared to the previous test with Intel Atom C3000 series, the only difference is that the two DUTs are replaced with Intel Atom C3858 processors.

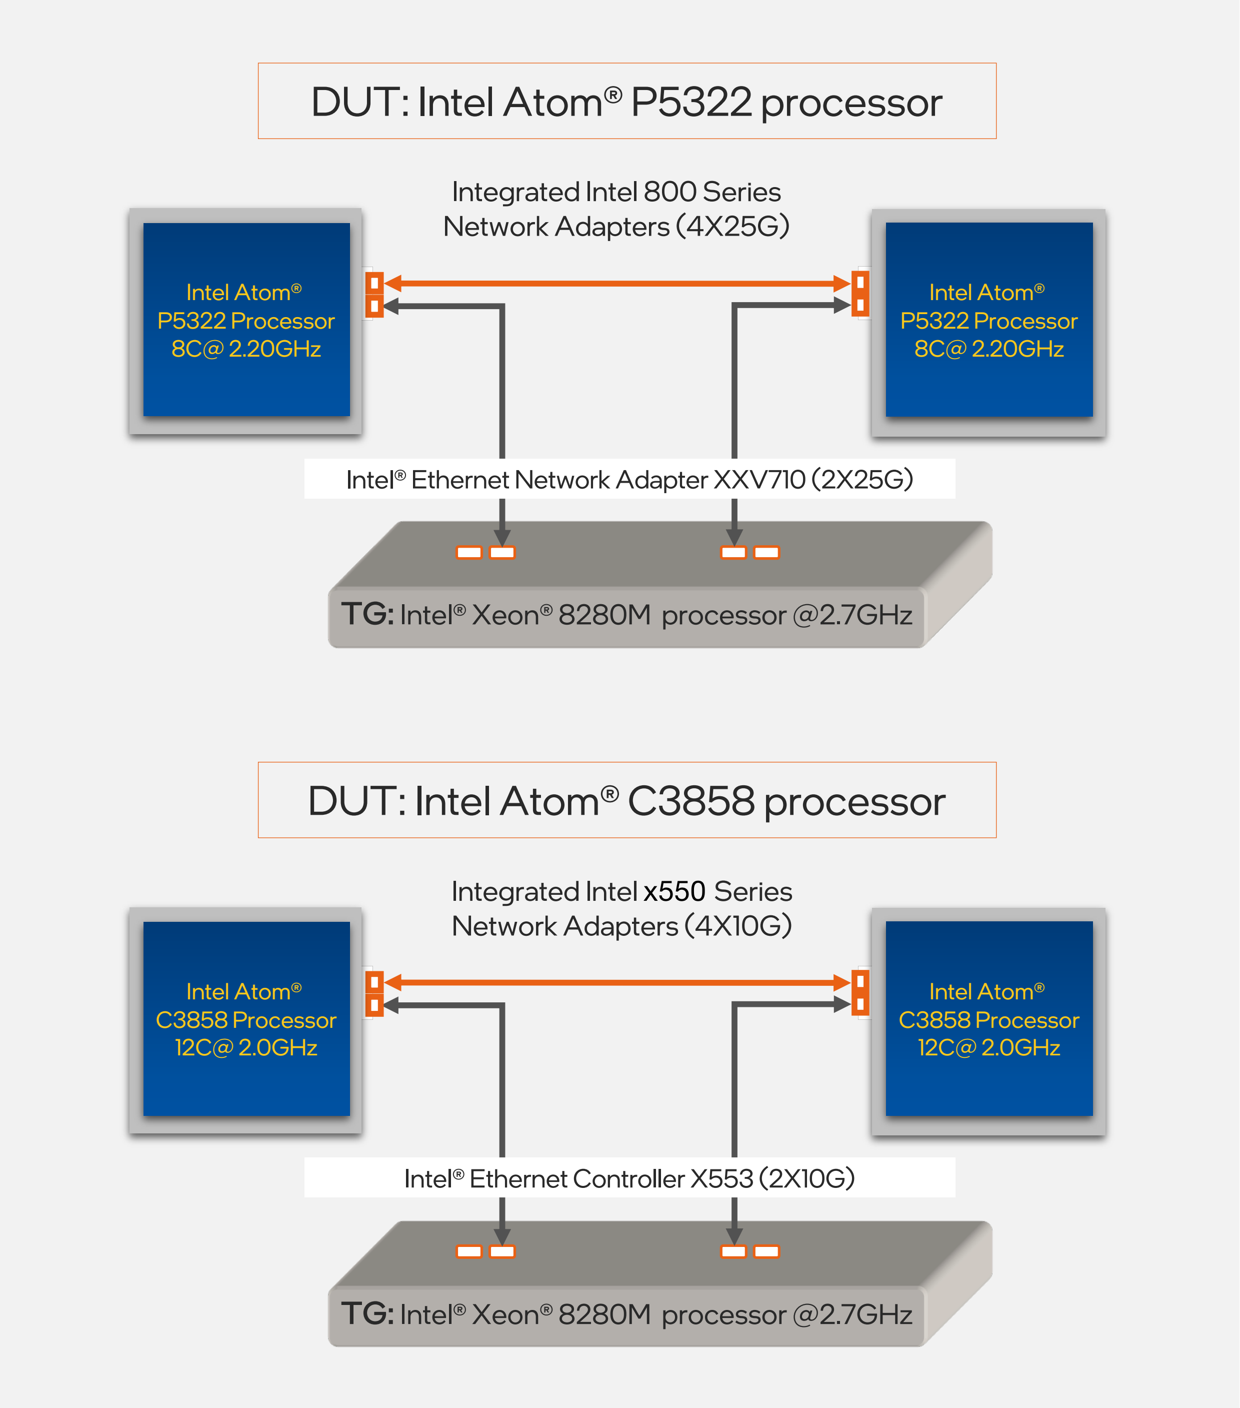
DUT: Intel Atom® P5322 processor
Integrated Intel 800 Series Network Adapters (4X25G)
Intel Atom® P5322 Processor 8C@ 2.20GHz
Intel Atom® P5322 Processor 8C@ 2.20GHz
Intel® Ethernet Network Adapter XXV710 (2X25G)
TG: Intel® Xeon® 8280M processor @2.7GHz
DUT: Intel Atom® C3858 processor
Integrated Intel x550 Series Network Adapters (4X10G)
Intel Atom® C3858 Processor 12C@ 2.0GHz
Intel Atom® C3858 Processor 12C@ 2.0GHz
Intel® Ethernet Controller X553 (2X10G)
TG: Intel® Xeon® 8280M processor @2.7GHz

Figure 7. Server setup for Intel Atom® P5322 and Intel Atom® C3858 series

For better comparison, it is important to ensure as much consistency between the two processor test environments as possible. Table 2 shows the detailed configuration information.

The performance benchmark test case is "ip6base-srv6enc2sids-End.DX6," which is already included in the FD.io CSIT project, The test case configured DUT1 and DUT2 with IPv6 routing, SR policy and steering policy, SR behavior is based on End.DX6. Segments left field in SRH header must be 0. The next header field is IPv6, decapsulate inner packet, forward on the interface associated with the Xconnect. This means the performance data for the comparison is all IPv6 based.

HARDWARE			
Testing as of		6/2022	6/2022
Motherboard		uCPE2.0 devices	Supermicro SYS-E300-9A
CPU	Product	Intel Atom® P5322 processor	Intel Atom® C3858 processor
	Speed (MHz)	2200	2000
	Number of Cores	8	12
	LLC Cache	7680K	2048K
Memory	Speed (MHz)	2400	2400
	Type	DDR4-2933 RDIMM	DDR4-2400 SODIMM
	Size	64GB	32GB
	Channels	1 DIMM/Channel, 2 Channels	1 DIMM/Channel, 2 Channels
BIOS	Vendor	Intel® Corporation	American Megatrends Inc.
	Version	JBVLCRB2.86B.0015.P03.2111300431	1.0b
	Microcode	0x4c000018	N/A
OS	Vendor	Ubuntu 20.04	Ubuntu 20.04
	Version	5.11.0-27-generic	5.11.0-27-generic
NIC	Product	4 x 25G ports from NAC NIC mode (NS)	4 x 10G ports from Intel Ethernet Controller X553
	Firmware	1.99 0x800094c0 1.3038.0	0x8000083f

SOFTWARE		
VPP version	v22.02	v22.02
gcc compiler	gcc 9.3.0	gcc 9.3.0

Table 2. Platform configuration for the Intel Atom P5322 processor and Intel Atom C3858 processor

Figure 8 shows the graph node process of VPP SRv6 implementation for the test case on two DUTs (both for Intel Atom P5000 processor and Intel Atom C3858 processor).

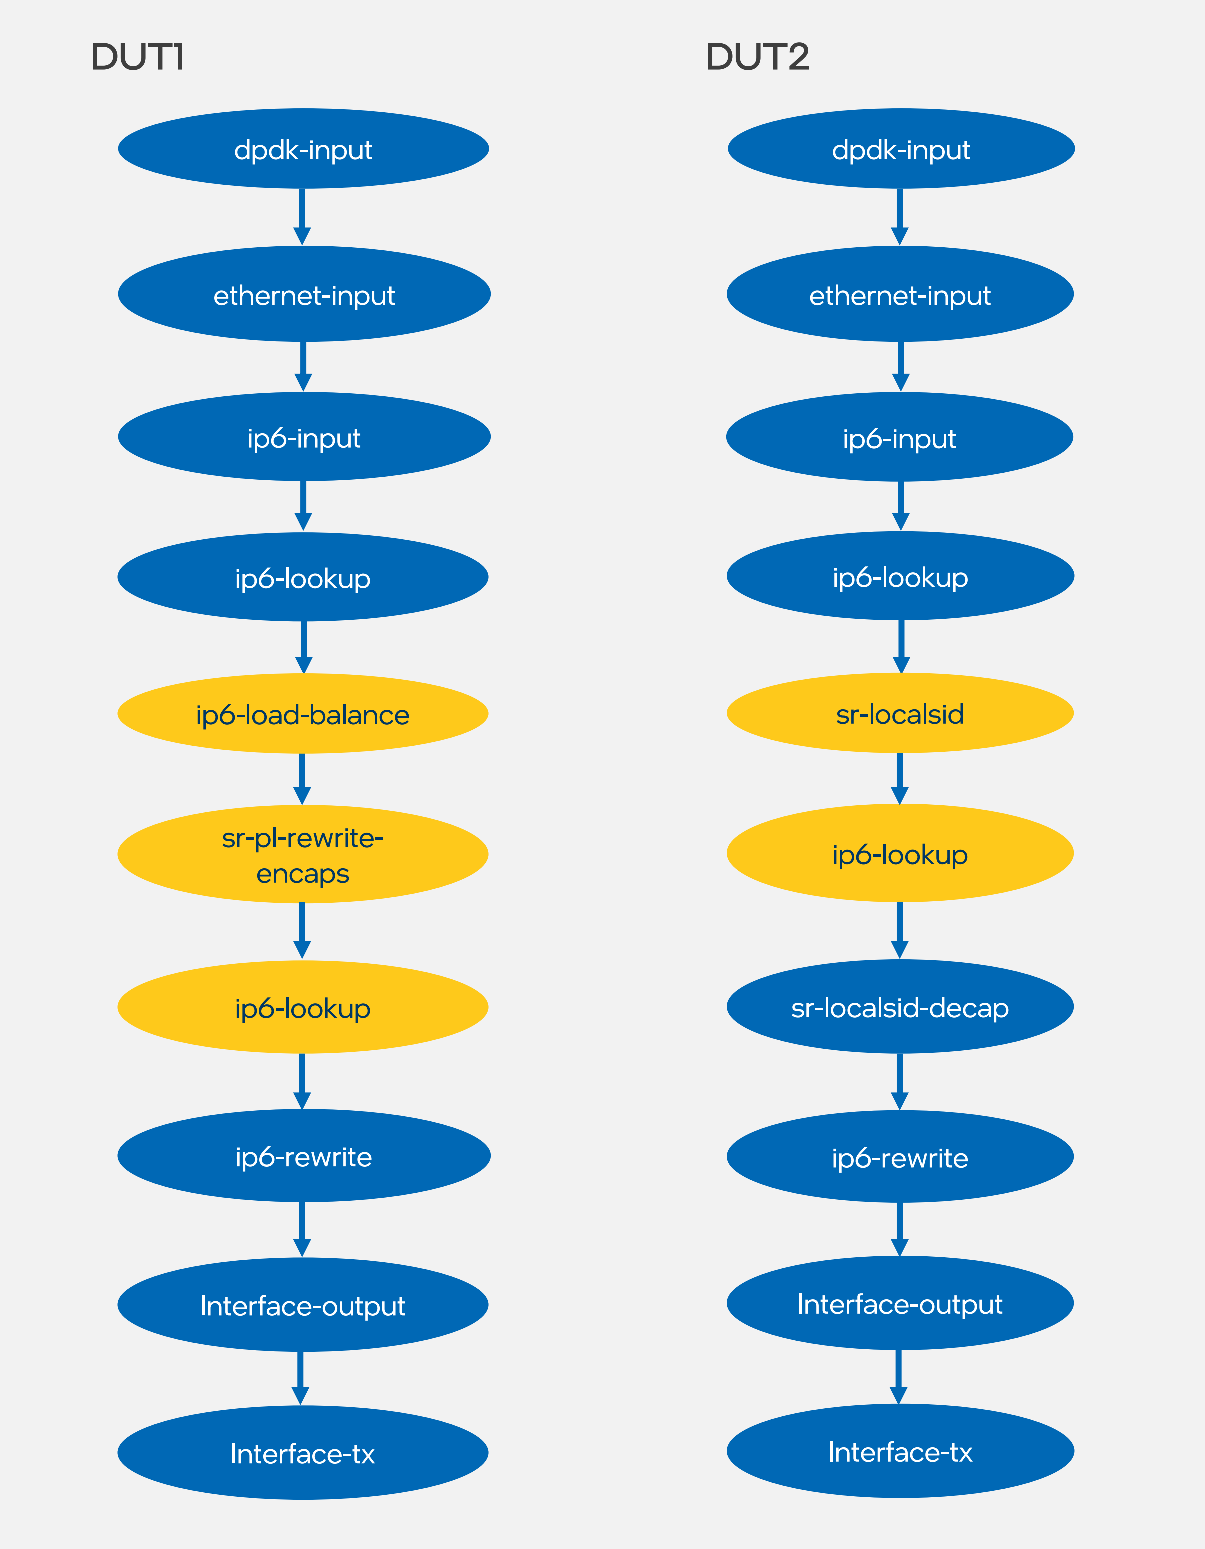

Figure 8. VPP SRv6 graph node process

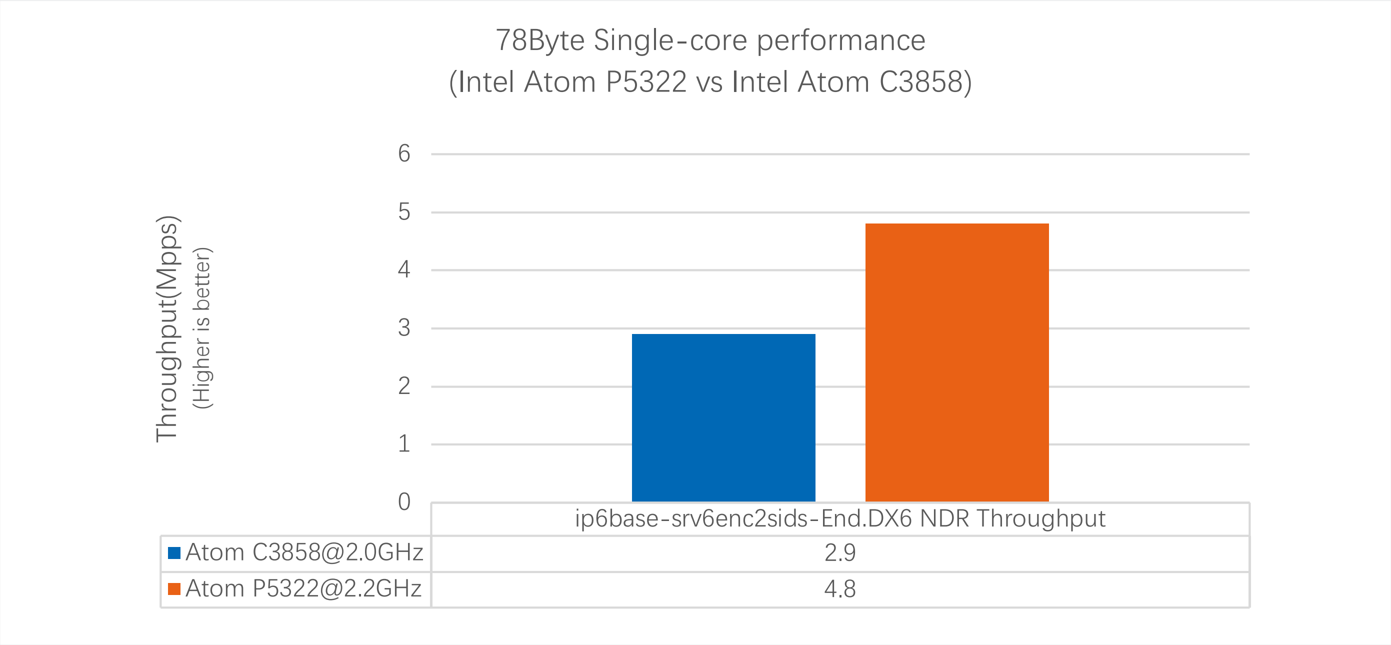

Figure 9. 78Byte single-core performance (Intel Atom P5322 processor vs Intel Atom C3858 processor)

3.2.3.2 Performance

Figure 9 shows the results of the tests. The uCPE2.0 device with Intel Atom P5000 processor delivers ~66% better performance improvement using one core / one 1 thread at 78-byte packet size compared with uCPE 1.0 (Supermicro SYS-E300-9A with Intel Atom C3858 processor). The performance data was collected for the NDR throughput method which shows maximum throughput at zero packet loss.

3.3 Intel® DLB-based Bandwidth Management

3.3.1 Intel DLB overview

Intel Dynamic Load Balancer (Intel DLB) is a hardware-managed system of queues and arbiters connecting producers and consumers.

Intel DLB interacts with producers and consumers running on Intel architecture processor cores. Figure 10 demonstrates the concept with four different traffic types.

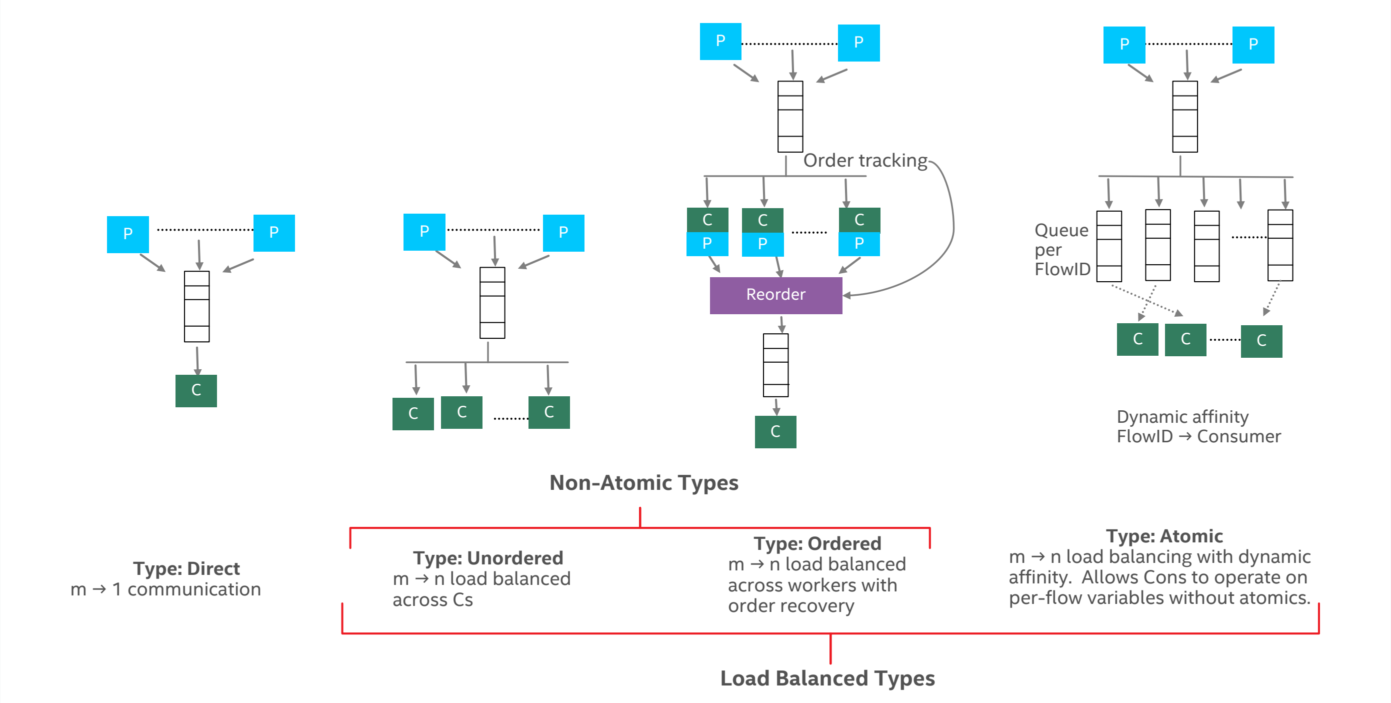

Figure 10. Intel DLB provides load balancing across four different traffic types

The four traffic types shown in Figure 10, include:

- **Directed:** 1-1 mapping between queue and consumer. Improves performance by enabling two independent HQM scheduling decisions in parallel.
- **Unordered:** Multiple queue entries are steered to multiple cores without restrictions on ordering. Improves performance with parallel execution on multiple cores.
- **Ordered:** Multiple queue entries are steered to multiple cores. Original order of QEs is restored by HQM on enqueue. Improves performance with parallel execution on multiple cores.
- **Atomic:** Queue entries are tagged with a flow ID. All QEs with the same flow ID are steered, one at a time, to the same core. This improves performance by removing critical sections from software.

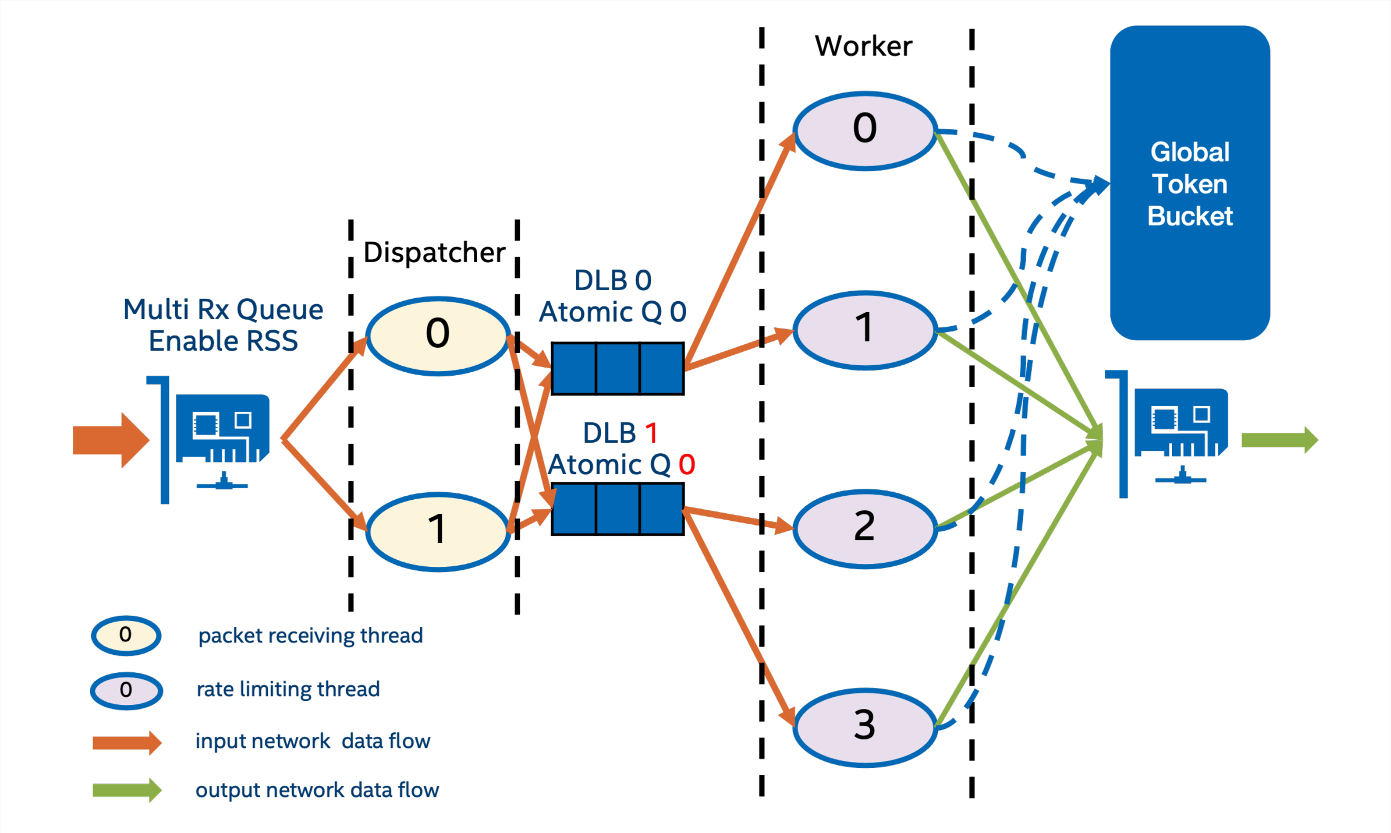

Figure 11. Lock-Less Rate Limiting Overview.

3.3.2 Intel DLB-based Rate Limiting Implementation

The lock-less rate limiting solution shown in Figure 11 utilizes the advantages of the Intel DLB Atomic Queues to evenly distribute packets across worker threads according to flow ID. Networking applications can get two benefits from this new solution:

- Good performance scaling with multi cores: A lock is used to protect data that might be concurrently accessed by multiple threads or processes; when lock conflict goes higher, performance will drop, and it will eventually become a packet processing performance bottleneck. The Intel DLB Atomic Queue allows packets with the same flow ID to be handled in same work thread, thus there is no need to add a lock to protect global token bucket which is shared among worker threads.
- Better performance at lower flow numbers: In an Intel DLB-based rate-limiting solution, unlike NIC RSS, the flow ID is not bound to a worker thread. Intel DLB Atomic Queue always chooses a suitable worker thread for a flow when it is not already handled by a worker. All worker thread loads are balanced, which lowers the possibility of dropping packets by an overloaded core.

4. Conclusion

uCPE 2.0 is the next generation uCPE white box server reference design. It enriches the Intel uCPE reference design portfolio with the enhanced features from Intel Atom P5000 and Intel Atom C5000 processors, enabling new use cases for SASE/SSE, SD-WAN, SRv6 and enterprise network appliances. By enabling the SM3/4 protocol from the integrated Intel QAT cryptography acceleration, high performance IPsec and SSL can be implemented and aid in compliance with the China National Security Standard* without additional hardware cost. It simplifies the hardware design and increases the network performance. SRv6 is a new underlay network technology that benefits from Intel's enhanced CPU microarchitecture allowing high performance SRv6 packet processing to be achieved on the uCPE 2.0 platform.

Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more on the Performance Index site.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. No product or component can be absolutely secure. Your costs and results may vary.

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

1222/TM/HO9/PDF ♻Please Recycle 353753-001US