Solution Brief

Packet Processing

Quad-Core Intel® Xeon® Processor

Intel® QuickAssist Technology

Packet Processing with Intel® Multi-Core Processors

General-Purpose Processor Exceeds Network Processor Performance

> "In all likelihood, multiple processing cores in the Dual-Core Intel® Xeon® processors will enable a significant improvement in performance density for this type of intensive signaling plane processing."
>
> Jorgen Trank,
> Product Manager
> Signaling Solutions,
> TietoEnator

Executive Summary

In fewer than ten years from now, industry experts are forecasting network capacity demands will grow by a factor of one thousand.[1] Preparing for this data and bandwidth boom, service providers are looking for cost-effective equipment that is more flexible and scalable. They are seeking IP-based telecom and networking equipment based on general-purpose processors and open standards, which improves hardware and software interoperability and reuse, and ultimately reduces deployment costs. Open standards are the foundation for the next-generation wireless network, called 3GPP Long Term Evolution (LTE), which will increase IP traffic throughput and decrease the number of network elements by consolidating functionality.

Supporting a standards-based approach, embedded Intel® architecture components have a long history of serving wired, wireless and networking infrastructure. They perform a broad range of functions including *control plane* (radio network controller), *application server* (home location register) and *security* (firewalls and VPN). But with the latest Intel® multi-core processors, equipment manufacturers are dedicating some of the extra processor cores to data plane and packet processing and achieving throughputs over 1,260 megabits per second (Mbps). Traditionally, this level of data plane performance necessitated the use of dedicated hardware like network processors, ASICs and FPGAs. Today, developers are combining control and data plane functions on a single board and eliminating packet processing boards, which lowers cost and speeds up content processing.

This solution brief focuses on consolidating control and data processing on general-purpose multi-core processors. Benchmark data is provided for two prototypes, a radio network controller (RNC) and standard packet processing, running on a single board based on Quad-Core Intel® Xeon® processors. This brief also describes the future of Intel® QuickAssist Technology, an initiative led by Intel and other industry leaders to enable high-bandwidth and low-latency accelerators within a standards-based framework.

Radio Network Controller (RNC) Benchmark

This section describes radio access networks and compares RNC performance for systems using network processors and general-purpose Intel architecture processors. The RNC benchmarks show that general-purpose processors can exceed network processors in packet processing performance.

A 3G Radio Access Network consists of a variety of network elements, as shown in Figure 1. The RNC and Node B connect wireless subscribers to circuit-switched and packet-switched networks. This connection comprises separate paths for call data and signaling, identified as the data plane and control plane, respectively. In addition, the RNC provides radio resource management and call handover control.

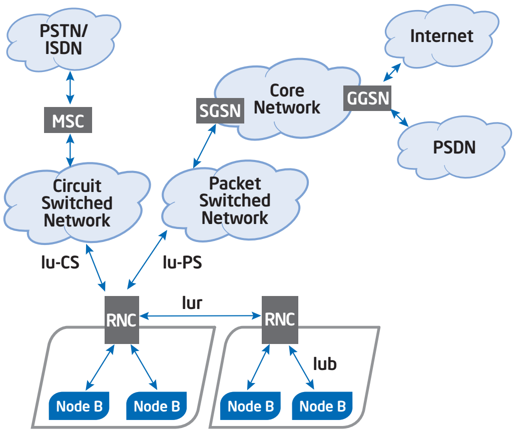

Figure 1. 3G Wireless Network Architecture

The RNC supports four primary interfaces that carry network traffic to other elements, as shown in Figure 1:

- Iu-CS to the circuit-switched network (uplink)
- Iu-PS to the packet-switched network (uplink)
- Iub to a Node B (downlink)
- Iur to another RNC (downlink)

The following discussion compares two implementations, based on Intel® network processors and Intel architecture multi-core processors, modeled to host over 500,000 subscribers. The prototype includes processing for the radio link control (RLC), media access control (MAC) and framing protocols. The benchmark testing exercises the most demanding workloads, the uplinks from circuit-switched and packet-switched networks.

Network Processor Implementation

An RNC implementation, using Intel NetStructure® IXB2800 3G boards, is illustrated in Figure 2. This configuration is based on a traffic model and performance validation results from IXB2800 3G board testing performed at Intel. These boards, based on Intel IXP2800 network processors, were designed to provide the performance and density required by next-generation WCDMA RNC data plane solutions.

- IXB28004XOC3 (eight line cards) – OC-3/STM-1 interfaces (1:1 sparing)
- IXB2800 (two RNL-c/sh boards) – paging and other data exchanges (1:1 sparing)
- IXB2800KAS (five RNL-d boards) – data plane for 504K subscribers (one spare)

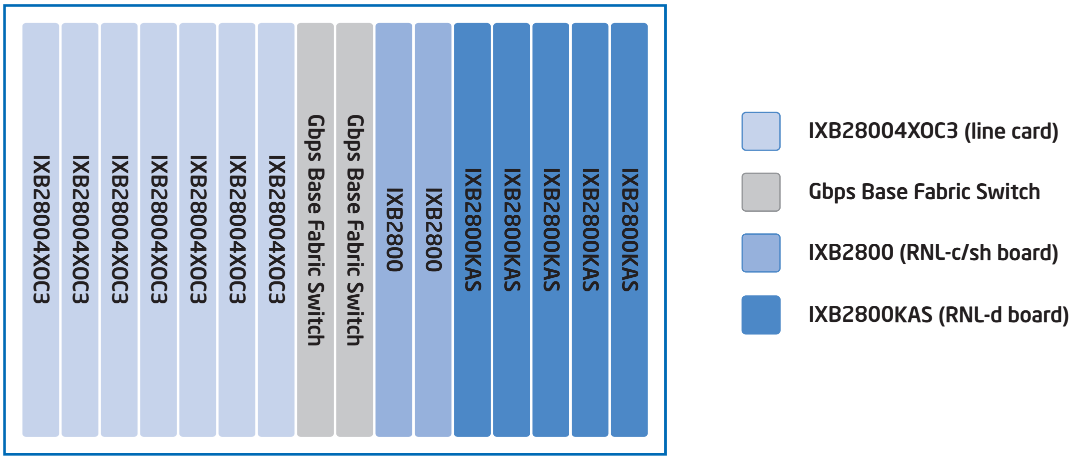

Figure 2. RNC Configuration – Network Processor-based Data Plane

Intel® Architecture Implementation

A second RNC implementation uses standard Intel architecture multi-core processors for data and control plane functions, as shown in Figure 3. This configuration is the same as the previous implementation, except the five IXB2800KAS boards are replaced with three Quad-Core Intel Xeon Processor 5300-based blades (1:1 sparing). As a result, two slots are now free and available for future expansion. Each blade has two quad-core processors – a total of eight processor cores – running at 2.13 GHz.

RNC Test Results

Both platforms supported the RNC traffic associated with a 504K subscriber network, approximately 1,260 megabits per second (Mbps), as shown in Table 1.[1,2] The network processor implementation employed four active IXB2800KAS blades, all of which operated at full capacity. The Intel architecture implementation has two active blades, where six processor cores were responsible for data plane traffic, allowing the remaining two processor cores to be used for other applications such as control plane processing. The six active cores operated at 50 percent utilization. At slightly higher processor core utilization, this platform handles even more network traffic, demonstrating that general-purpose processors can exceed network processors in packet processing performance.

The Intel architecture data plane implementation requires less space (three versus five blades), thereby freeing up two blade slots, compared to the network processor implementation.

Table 1. RNC Benchmark Results[2]

	RNC Data Plane	
	Network Processor Implementation	Intel® Architecture Implementation
Blades	IXB2800KAS (RNL-d board)	Quad-Core Intel® Xeon® Processor 5300 Server
Number of Blades: Active Spare	4 1	2 (6 of 8 CPU cores active) 1
Total Blades	5	3
Supported Throughput Incoming IuCS/PS traffic	1,260 Mbps	1,260 Mbps
Utilization	100%	6 CPU cores (50%)
Control plane function	Not supported by IXB2800KAS boards	2 CPU cores available

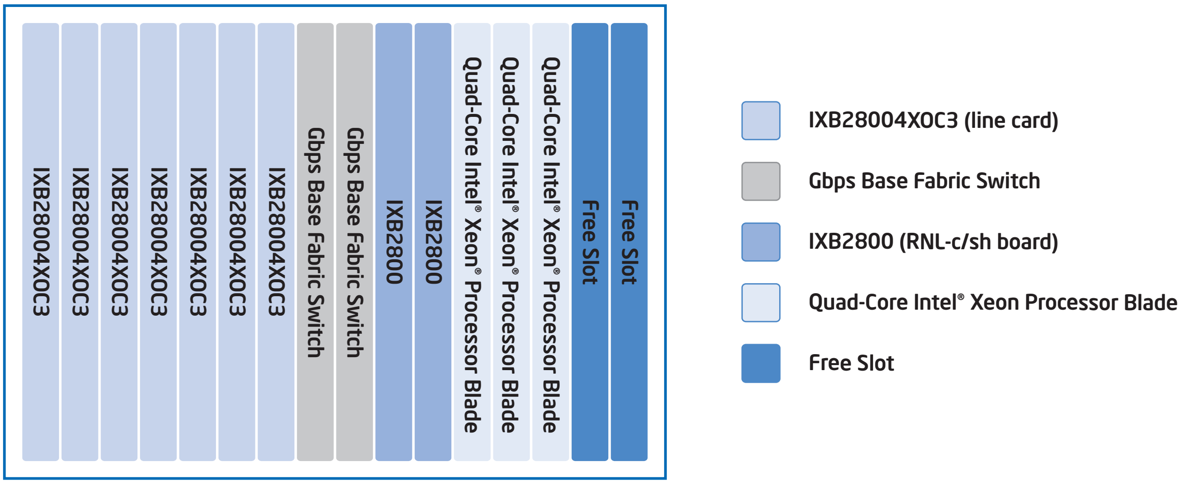

Figure 3. RNC Configuration – Intel® Architecture-based Data Plane

Flow Processing Benchmark

This section provides performance data for a flow processing prototype based on Intel® quad-core processors and shows the performance impact of executing control and data processing on a single system. The flow processing benchmarks demonstrate the ability of Intel general-purpose processors to perform both control and data processing on a single board at the same time.

Flow processing is a key function in various packet applications such as security processing, media processing and quality of service (QoS) processing. It classifies packets into flows based on a set of specified parameters and performs operations on those packets. The flow processing pipeline primarily consists of the blocks shown in Figure 4.

The flow processing pipeline is implemented as a Linux* kernel module, which bypasses the Linux network stack and receives packets directly from the network driver. The Hash Lookup block takes packets from the RX (receive) block and calculates a hash value that is used to assign packets to one of the CPU cores performing flow processing. This implementation employs the 'jhash' routines available in the Linux kernel, which operates on five fields in the packet: IP Source Address, IP Destination address, IP source port number, IP destination port number, and protocol ID.

The Policy Lookup block classifies the flows and determines which rules should be applied to the flows. The Action Handling routine block modifies packets and schedules them for transmission by the TX (transmit) block.

A Spirent Load Generator creates packet traffic for the receive and transmit blocks (RX and TX) by using sixteen Gigabit Ethernet links that are connected to the system under test (SUT), as shown in Figure 5.

Packet Forwarding Test Results

The SUT was equipped with two 3.0 GHz Quad-Core Intel Xeon Processors 5300, and the 16 Ethernet interfaces produced a maximum load up to 16 gigabytes per second (Gbps). The benchmark tests focused on packet forwarding to determine the SUT's maximum packet processing capability, although the flow processing application was capable of much more complex functions.

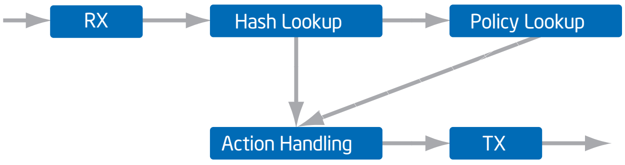

Figure 4. Flow Processing Pipeline

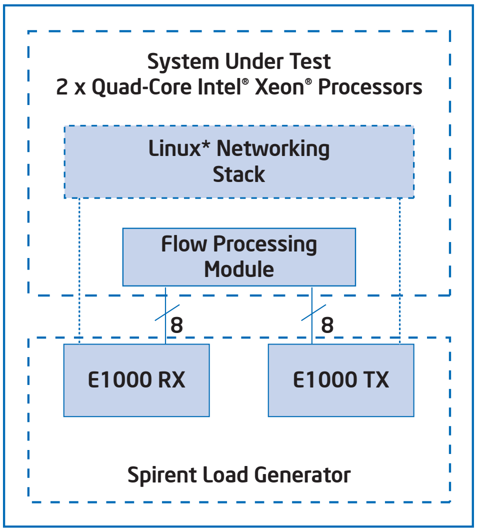

Figure 5. Implementation Model

The SUT performed duplex forwarding, where each of the eight cores handled two Gigabit Ethernet networking interfaces transmitting 64-byte Ethernet frames. The SUT processed 10.1 million packets per second (Mpps) without any dropped or lost, and the throughput increased as more cores were activated, as shown in Figure 6.[1,2]

Impact of Control Plane Processing

Another benchmark test assessed the impact of running flow processing on six cores while the other two cores executed a control plane module that read and wrote to a table. In this test, the size of the table varied between 0 and 8 megabytes, corresponding to the total amount of the L2 cache memory in the SUT. As the control plane table size increases, valuable cache memory resources are diverted from the data plane applications, which reduces flow processing performance. The resulting throughput decreases from 9.1 to 8.5 Mpps, as shown in Figure 7, which corresponds to a seven percent performance reduction, a relatively small amount.[2] This testing indicates that a single board based on embedded Intel multi-core processors provides flow processing performance comparable to network processors, while also supporting other applications like control plane processing.

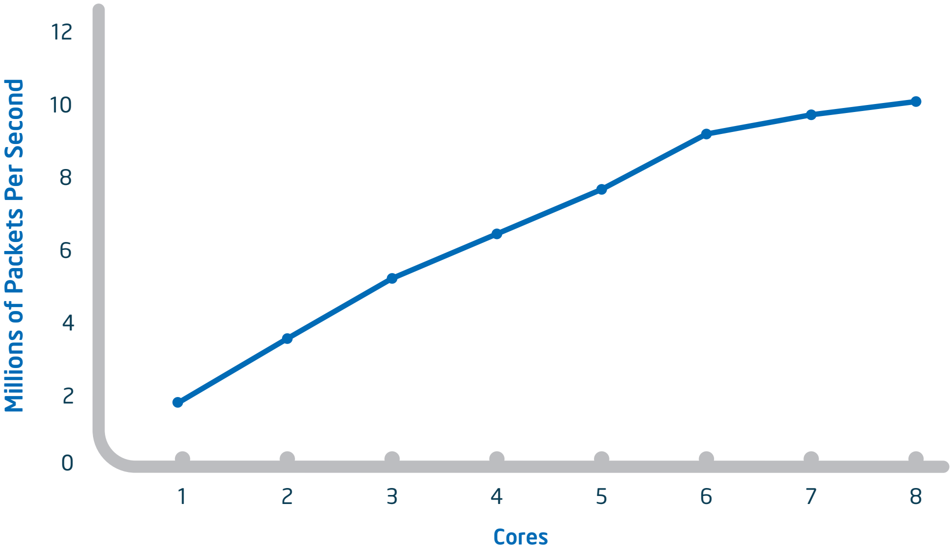

Figure 6. Duplex Forwarding Results

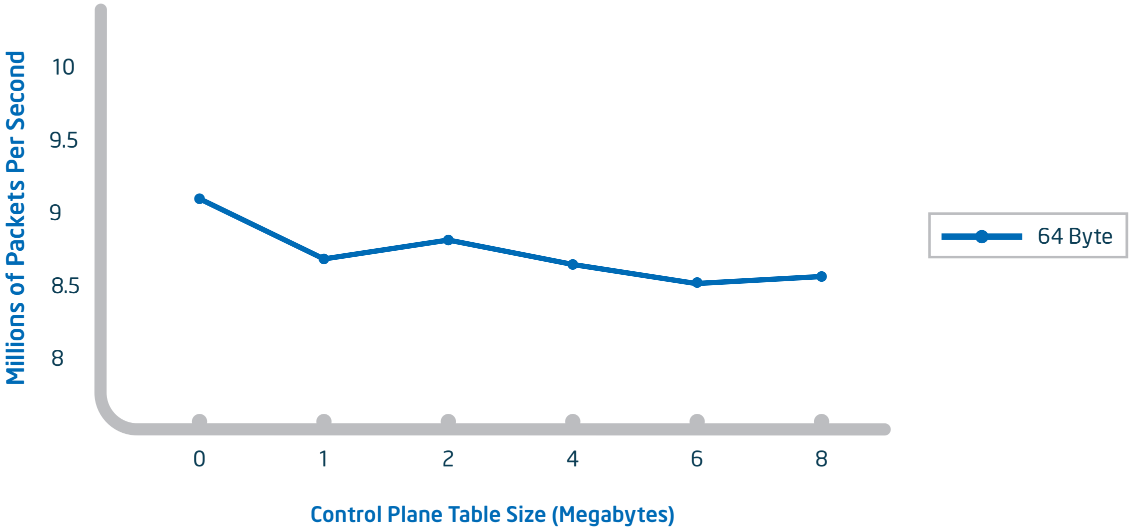

Figure 7. Data Plane Performance Versus Control Plane Table Size

Standards-Based Accelerator Technology

Large enterprise networks are transitioning to 10 Gigabit Ethernet and beyond, which requires security solutions capable of handling ever-increasing IP-based traffic. When network throughput exceeds the capability of software-only solutions, system developers make use of high-bandwidth and low-latency accelerators. Performance notwithstanding, networking and communication equipment vendors are demanding accelerators that are flexible, scalable and software-friendly.

To satisfy these requirements, Intel, ISVs and other industry leaders are setting the stage for accelerator innovation across security functions, markets and interconnect standards. This effort includes the Intel® QuickAssist Technology, which is a comprehensive initiative that enables the optimized use and deployment of accelerators on Intel® platforms, as illustrated in Figure 8.

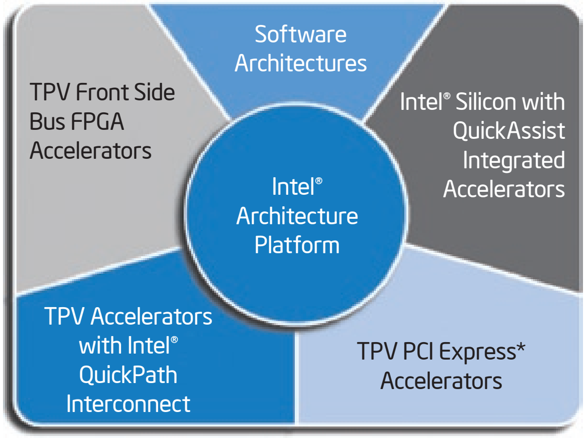

Figure 8. Intel® QuickAssist Technology Solutions

This initiative consists of a family of interrelated Intel and industry-standard technologies, including:

- **Accelerating packet processing for demanding applications** by attaching pre-processors to the Intel® processor Front Side Bus with Field Programmable Gate Arrays (FSB-FPGA) or through PCI Express.*
- **Easing the migration from one technology to another** by minimizing the impact to applications with the Intel QuickAssist Technology Accelerator Abstraction Layer (AAL).
- **Supporting small form factor accelerators** by combining numerous powerful networking and security technologies on a single chip, codenamed 'Tolapai.'
- **Increasing the speed and efficiency of standard interconnects** by working with the PCI-SIG to define PCI Express* 3.0 – initially proposed by Intel and IBM, codenamed 'Geneseo' – which will improve accelerator efficiency and deliver eight gigatransfers per second, double the bandwidth of PCI Express 2.0.

Today, many equipment makers deploy acceleration technology to boost performance. For example, the Nokia IP2450* security platform, running Check Point SecureXL* software, supports firewall throughput of approximately 9 Gbps, which can be doubled to 20 Gbps by upgrading the system with optional Accelerated Data Path* (ADP) add-in card technology. In the future, Intel QuickAssist Technology will enable the industry to develop accelerator add-in cards with faster communication links, like PCI Express 3.0 and processor busses, to further increase packet throughput.

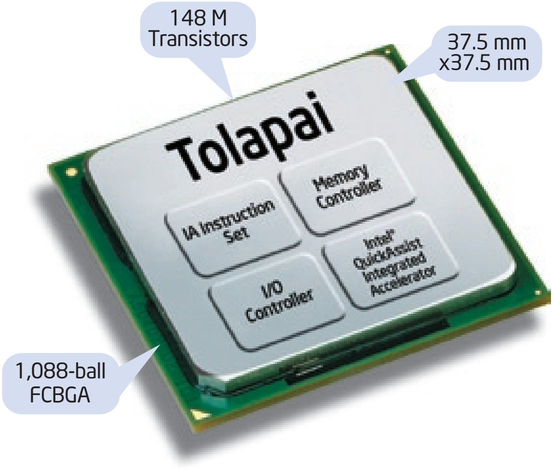

Figure 9. The Tolapai Processor

Ingredients for a Single Platform

Today, equipment vendors are under pressure to speed up development cycles and spin new products that address the latest emerging threats. They are looking for agile platforms that provide predictable performance and higher levels of scalability and flexibility. Combining standards-based technology, new accelerator initiatives and high-performance Intel multi-core processors, equipment vendors can create a single platform that readily scales network throughput and satisfies the needs of different market segments. Understanding this paradigm shift, developers are optimizing packet processing code for multi-core architecture and changing the game.

[i] Based on data from Nielsen/NetRatings, International Telecommunications Union, worldinternetstats.com and Intel analysis.

[1] Performance tests and ratings are measured using specific computer systems and/or components and reflect approximate performance of Intel® products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit http://www.intel.com/performance/resources/benchmark_limitations.htm

[2] For additional information about the benchmark results described in this paper, please contact your local Intel Sales Representative.

Printed in USA 0508/LDK/OCG/XX/PDF ♻ Please Recycle 319889-001US