

Transitioning to NFV with Radisys Networking Solutions Running on Intel® Processors

Designed to run consolidated, virtualized network functions, Radisys network appliances and ATCA compute modules are versatile building blocks for infrastructure transformation

Overview

Network operators and communications service providers (CSPs) are calling for more infrastructure solutions based on the principles outlined by network functions virtualization (NFV)¹ in order to reduce infrastructure expenditures (CapEx/OpEx) and accelerate the deployment of new services. At the heart of this architectural approach is the decoupling of network functions from hardware through abstraction and virtualization. The end result is software workloads that can be run on a common pool of hardware resources, allowing them to be deployed dynamically throughout the network as needed.

The benefits offered by this paradigm shift are amplified if the hardware resource pool has homogenous technology attributes, as in the case of standard Intel® x86 CPU-based servers. As proven by Web and enterprise data center deployments, the consolidation of network functions using a homogenous resource pool can reduce time to market and infrastructure costs by orders of magnitude.

This white paper describes how recent advances in Intel® x86 performance and software libraries are enabling solutions that address more network workloads, thus taking the telecom industry closer to the NFV design objectives. Examples include Radisys network appliances and ATCA compute modules based on Intel® processors, which achieve the performance, reliability and availability critical to address network infrastructure requirements as it transitions to an NFV model. The solutions support various networking functions, such as policy enforcement, the GTP-u protocol acceleration, media optimization, deep packet inspection (DPI), load balancing, switching and security, among others.

CONTENTS

Packet Processing on Intel® Processors	<i>pg. 2</i>
Enhancing Open Source for NFV	<i>pg. 2</i>
Benefits from Consolidation	<i>pg. 3</i>
Packet Processing Software Enhancements	<i>pg. 3</i>
Radisys Networking Solutions	<i>pg. 3</i>
Extensive Technical Support from Radisys	<i>pg. 6</i>
Workload Acceleration	<i>pg. 6</i>
Solutions for SDN/NFV	<i>pg. 6</i>
References	<i>pg. 7</i>

Packet Processing on Intel® Processors

The growing complexity of today's infrastructure is driving widespread interest in NFV among network operators and service providers. A major factor is the increasing variety of proprietary, fixed-function appliances that are unable to deliver the agility and economics needed to address constantly changing market requirements. This is because optimization for performance has mandated the use of specialized technologies, limiting flexibility, and the development and deployment of new services. The impact can be seen when a service provider wants to create a new service, but finds it requires a software change that only the OEM can make because the equipment is essentially a closed system.

One reason traditional network elements are rather complex is the use of different hardware architectures to run diverse workloads, like packet processing on network processors (NPU) and Application Specific ICs (ASICs), and control and application processing on general-purpose CPUs. As a significant step towards changing this paradigm, the latest Intel® Xeon® processors combined with the Intel® Data Plane Development Kit (Intel® DPDK) can provide a significant boost to networking/packet processing workload performance. This is possible because multi-core architecture with integrated memory controllers is delivering exceptional performance gains and significantly lower memory latencies (as much as a 40-50% reduction), allowing all these workloads to run on a single architecture. L3 forwarding throughput, for example, has sharply improved and is now on par with many NPUs, as detailed later in the paper.

Enhancing Open Source for NFV

Communications service providers have stringent timing constraints for their mission-critical applications and services, such as voice, video and charging. In many cases, open-source software components must be enhanced in order to satisfy the associated real-time requirements. Consequently, Intel, Radisys and Wind River have been working to improve the performance of network functions running in NFV environments.

An example is the Intel® DPDK Accelerated Open vSwitch—a fork of the open source Open vSwitch multilayer virtual switch found at <http://openvswitch.org/>. This software-based function re-created the kernel forwarding module (data plane) by building the switching logic on top of the Intel DPDK library to significantly boost packet switching throughput. The forwarding module runs in Linux User Space with BSD license rights.

The solution integrates the Wind River® Open Virtualization Profile (OVP), which is an add-on to Wind River Linux 5 that provides performance enhancements, management extensions and application services through open components. The Wind River OVP integrates a range of technologies to deliver adaptive-performance, streamlined interrupts, system partitioning and security management. It supports a heterogeneous collection of guests and hosts with options ranging from KVM guests and hosts with Wind River Linux only, through to KVM guests and hosts with mixed Linux guests, and finally to combinations of Linux and non-Linux guests. Additionally, Wind River OVP produces a set of packages that can be used on non-Wind River Linux distributions, allowing integration with third party or Wind River Linux orchestrated networks.

This software, along with many other types of virtualized network functions, runs on Radisys network appliances and ATCA compute modules in support of NFV-based infrastructure.

Benefits from Consolidation

For many low to mid-density nodes in particular, consolidating multiple workloads onto Intel® architecture may allow NPUs to be removed from the design. This approach can lower development costs by creating more software reuse opportunities and simplifying the tool chain, which boosts efficiency, reduces training time, decreases license fees and enables programmers to work on any system function. Moreover, moving to a single architecture eliminates many integration and validation issues, saving time and effort.

If network equipment providers want to avoid hardware development altogether, they can use commercial, off-the-shelf (COTS) boards with carrier grade capabilities and long life support available from Radisys. These products take advantage of the Intel® processor roadmap, which continues to confirm Moore's Law through continuous investment in technology and manufacturing. On roughly an annual basis, Intel launches higher performance computing platforms used by equipment manufacturers to develop more capable products. Service providers can benefit from lower operating expenses (OpEx), because Intel processors optimize power consumption and lower the maintenance costs associated with managing complex multi-architecture systems.

Packet Processing Software Enhancements

Complementing its microarchitecture advancements, Intel® has made packet processing software enhancements provided by the Intel Data Plane Development Kit (Intel DPDK). The development kit reduces a significant amount of overhead when using an out-of-the-box, standard Linux operating system to host user plane stacks. Significant processing time is saved by using core affinity, disabling interrupts



Figure 1. Radisys R220 Network Appliance

generated by packet I/O, enforcing cache alignment, implementing huge pages to reduce cache misses, prefetching and many other concepts. The Intel DPDK runs in user space, thus removing the high overhead associated with kernel operations and with copying data between kernel and user memory space.

Radisys Networking Solutions

Radisys has a broad selection of x86 offerings, including several expressly designed for data plane applications, such as the R220 Network Appliance, ATCA compute modules and a GTP-u protocol stack, which are discussed in the following:

R220 Network Appliance

This 2U application-ready, DPI platform is well-suited to policy enforcement, network monitoring, lawful intercept, security/next-generation firewall and video optimization, among other applications. For these applications, operators need low- to high-density platforms spanning deployments across access, edge and core networks. The appliance features an innovative design (Figure 1) with modular front I/O and front storage modules that can be serviced/configured without de-racking. Radisys integrated the Intel DPDK, developing a comprehensive support package for it to enable impressive packet processing performance and ease of use.

Figure 2 shows the exceptional gains in packet processing performance for three generations of Intel Xeon processors, along with Intel DPDK integration by the 2012 generation. The configuration on the right is used by the R220 network appliance, providing nearly four times greater L3 forwarding throughput than the 2009 generation. When performing as a PCEF, a single appliance can deliver up to 40 gigabits per second (Gbps), and two to three clustered appliances can support around 100 Gbps.

ATCA-XE100 Compute Module

The ATCA-XE100 (Figure 3) is a ninth generation single board compute module from Radisys. It is a high-performance, single-slot, ATCA blade based on dual socket Intel® Xeon® processor E5-2600 with 16 very low profile (VLP) DIMM Sockets. The server-class compute module is ideal for DPI, packet processing, video optimization or any application that requires high-density I/O or the highest performance processing possible on ATCA. Individual modules can deliver up to 40 Gbps, and an ATCA platform with multiple blades can scale to over 300 Gbps processing throughput.

The ATCA-XE100 was designed to provide flexible, high-density I/O for many of today’s DPI and video optimization applications that need large amounts of I/O. For system architectures using a distributed I/O strategy, I/O comes directly into each processing blade. The ATCA-XE100 can easily support this approach with its 2 x 10 GbE ports on the front panel and an optional 8 x 10 GbE rear transition module (RTM). For centralized architectures, where all the I/O comes through the fabric, the ATCA-XE100 has dual 40 GbE NICs on the board, which provides up to 80 Gbps to the fabric in an active/active configuration. Radisys ported the NIC driver to the Intel DPDK and supports it directly.

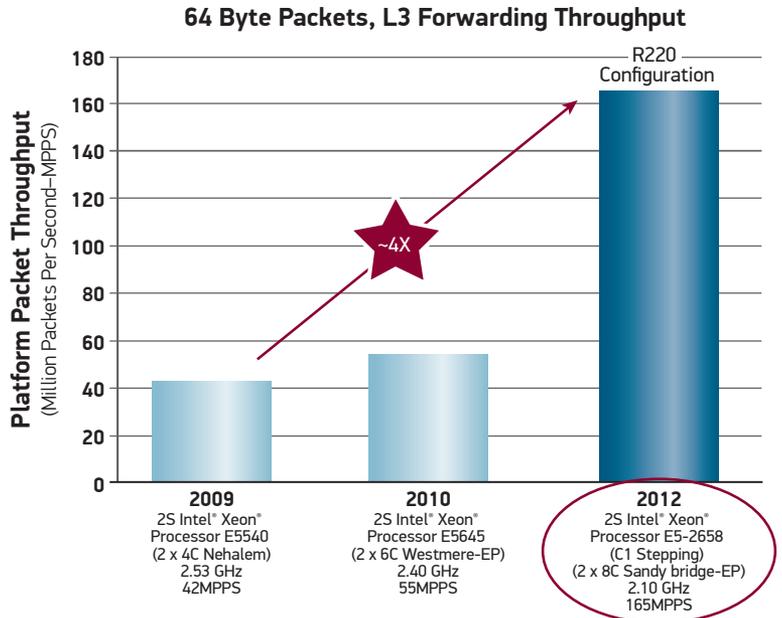


Figure 2. L3 Forwarding Performance Benchmark



Figure 3. Radisys ATCA-XE100 Compute Module

ATCA-4600 Compute Module

This module is ideal for control plane and server functions needed for LTE wireless infrastructure, IPTV and IP Multimedia Subsystem (IMS). It provides 40 Gbps fabric connectivity, a mezzanine slot, and several optional RTM offerings with additional storage and network interface choices. The dual, eight-core Intel® Xeon® processor E5-2448L processor provides high performance at a low cost per cycle and includes Intel® Hyper-Threading Technology (Intel® HT), which allows the 16 physical cores to process up to 32 threads simultaneously. The 12 DDR3 VLP RDIMMs allow up to 96GB of memory, while providing a cost-effective solution for applications with lower density memory requirements. The module also runs Intel DPDK software for packet processing applications.

GTP-u Fast Path Software

The explosive growth in mobile data traffic, forecasted to increase 18-fold between 2011 and 2016,² is putting unprecedented pressure on network operators. There were 175 million laptops on the mobile network in 2011, and each laptop generated 22 times for more traffic than the average smartphone.² Service providers are also introducing new services at a rapid pace, requiring network equipment providers (NEPs) to develop network elements that cost-effectively combine packet and application processing to address this dynamic landscape.

To better support these demanding requirements, Radisys optimized its Trillium GTP-u protocol software using the Intel DPDK running on the Intel Xeon processor. GTP-u, or the General Packet Radio Service (GPRS) Tunneling Protocol (GTP) user plane, refers to the protocol between:

- **3G:** Radio Network Controller (RNC) and the Serving GPRS Support Node (SGSN)
- **LTE:** eNodeB and the Serving Gateway (S-GW)

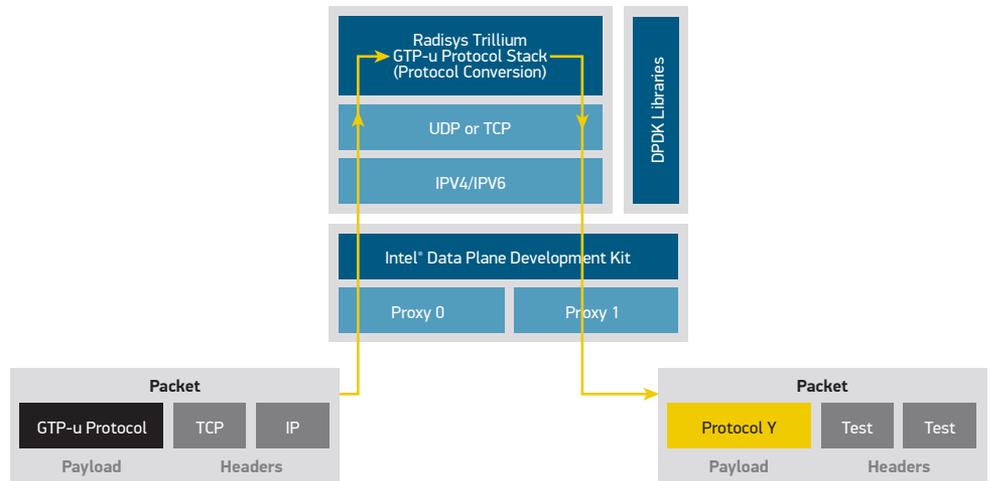


Figure 4. Trillium GTP-u Protocol Software Ported to the Intel DPDK

Since this cellular protocol is not associated with a programming API, developers often encapsulate the payload of the protocol into a standard TCP/IP socket, as shown on the left side of Figure 4.

Voice traffic is received on a single port, like proxy0, and then sent up to the network stack and applications, where it is manipulated and converted into another cellular protocol and forwarded on another socket. In this case, there are two sockets—one for terminating the GTP-u protocol and the other for sending the data out in a different protocol format.

The Trillium software also decrypts the packet and performs protocol-level deep packet inspection via a big table lookup. Exception packets are sent up to applications running on the same platform, exemplifying how user plane, control plane and applications can all run on a single platform.

What is unique about this solution is the addition of the Trillium GTP-u stack to the Intel DPDK and the ability to easily add applications and services. This solution is ideal for a cost-optimized signalling gateway or a high-performance eNodeB.

Extensive Technical Support from Radisys

With in-house technical resources and a global footprint, Radisys maintains the resources that enable developers to remain focused on meeting their customers' needs. Radisys offers remote diagnostics, which allows its technical support professionals to directly link into the live environment, and identify and resolve problems quickly, thus minimizing costly downtime. When configuring ATCA systems, OEMs and network operators can get guidance from Radisys technical experts about ways to maximize performance. Even if the hardware architecture includes products from multiple vendors, Radisys technical support can help ensure all hardware and software, across all platforms, is functioning properly.

Intel® DPDK Support

Radisys provides training, advice for performance optimization and issue resolution for the Intel DPDK running with Radisys hardware and software, so developers need not go to multiple ecosystem entities for help or questions. Radisys is the single front—with 24 x 7 technical support—for x86 data plane requirements: rackmount and bladed hardware, Intel DPDK software, support, consulting and professional services.

Workload Acceleration

In combination with GTP-u user plane processing, equipment developers can accelerate bulk encryption, data compression and other workloads using Intel® QuickAssist Technology: a set of software and hardware modules accessible via a unified set of industry-standard application programming interfaces (APIs) that provide consistent conventions and semantics across multiple accelerator implementations. For instance, the Intel® Platform for Communications Infrastructure with integrated Intel QuickAssist Technology is expected to deliver up to 100 Gbps crypto acceleration, which is comparable to today's solutions that typically employ multiple architectures to achieve this rate.

Intel promotes standard, open interfaces (i.e., APIs) through open source projects, such as OpenSSL and Openswan, that are also supported by Intel QuickAssist Technology. This capability enables developers to use built-in platform acceleration or add-in acceleration cards, with minimal software modification, to achieve different cost-performance objectives.

Solutions for SDN/NFV

Network equipment developers now have another option to consider when architecting solutions for SDN/NFV-based networks. It's possible to use a single architecture that consolidates all the workloads, thereby simplifying development and lowering hardware cost. Still, a good case can be made for designing with NPUs or ASICs when the highest throughput possible is required, as in a load balancer with low-packet touch. However, when low-to-medium packet forwarding performance is needed, along with services such as encryption, decryption, data compression, pattern matching or flow classification, an Intel Xeon processor-based platform may be ideal.

The scalability and virtualization capabilities of Intel processor-based platforms enables system engineers to develop a cost-competitive family of products based on a common code base. This ability significantly reduces software engineering effort, while the opportunity to use Radisys boards greatly minimizes hardware development cost. Radisys systems are supported by a broad ecosystem of independent software vendors (ISVs) and reduce the effort to bring SDN/NFV-based solutions to market.

References

- ¹ Source: “Network Functions Virtualization—Introductory White Paper,” published at the October 22-24, 2012 at the “SDN and OpenFlow World Congress,” Darmstadt-Germany, pg. 3, 4, 9, <http://portal.etsi.org/portal/server.pt/community/NFV/367>.
- ² Source: Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2011–2016. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf.



Radisys is an Associate member of the Intel® Intelligent Systems Alliance. From modular components to market-ready systems, Intel and the 200+ global member companies of the Alliance provide the performance, connectivity, manageability, and security developers need to create smart, connected systems. Learn more at: intel.com/go/intelligentsystems-alliance.

The Radisys logo, consisting of the word "radisys" in a lowercase, sans-serif font, with a registered trademark symbol (®) to the right. The logo is set against a dark red rectangular background.

Corporate Headquarters

5435 NE Dawson Creek Drive
Hillsboro, OR 97124 USA
503-615-1100 | Fax 503-615-1121
Toll-Free: 800-950-0044
www.radisys.com | info@radisys.com

©2013 Radisys Corporation.
Radisys and Trillium are registered trademarks of Radisys Corporation.
*All other trademarks are the properties of their respective owners.
September 2013