SOLUTION BRIEF

Communications Service Providers
Virtual Switch

Lagopus* Software Switch Helps CommSPs Meet SDN Requirements

Lagopus is a high-performance software OpenFlow* 1.3 switch that leverages the Data Plane Development Kit to realize almost 40 Gbps packet processing and 1 million flow rules.[1]

Summary

Communications service providers (CommSPs) face a number of challenges in today's market: fast growing data traffic, pressure from customers to expand their services, and revenue growth that is not keeping pace with traffic volume and network costs. Software-based network solutions such as software defined networking (SDN) and network functions virtualization (NFV) can help. The potential of these technologies to optimize capital and operational expenditures and help CommSPs realize their customers' needs in a timely manner is driving strong interest in high-performance software-based network solutions.

Lagopus* is a new open source SDN software switch. It leverages the Data Plane Development Kit (DPDK) and realizes almost 40 Gbps packet processing and 1 million flow rules.[1] Because it fully conforms to OpenFlow* switch specification 1.3, including the CommSP network protocol handling, multi tables, group tables, QoS functions, and more, its applications can be expanded to data center, CommSP, and enterprise networks.

Challenge

The evolution of virtualization technology and its maturity have changed the way data centers operate. Cloud computing has opened a way of creating new services, but it has then created demands for diversification of network services to meet each end user's requirements. SDN is a modern solution and a promising technology to provide the needed flexibility to highly integrated network services.

The increases in network complexity and the need for rapid service delivery put cost and time-to-market pressures on service providers. As network traffic continues to grow, communications service providers face strong pressure to optimize CAPEX and OPEX since revenues are only modestly increasing.[2] NFV was perceived as a direction to deal with these pressures through the use of standard, off-the-shelf servers. These servers give communications service providers the freedom to install software-based NFV applications, as opposed to dedicated hardware blocks, which both helps lower costs and deployment times, so new services can be deployed to meet customers' needs in a timely manner.

Since NFV applications inherently are required to handle traffic of millions of packets per second, there has been strong demand for a software-based solution to realize the performance. Furthermore, the advantages of SDN and NFV should be leveraged not only in the data centers, but also in CommSP networks. Networks require a high level of performance to keep up with increasing traffic. For example, the number of routes in the Internet that must be handled by CommSP routers has grown to around 0.6 million and is still growing.[3] Therefore, CommSPs are looking for the performance to handle 1 million flow rules.

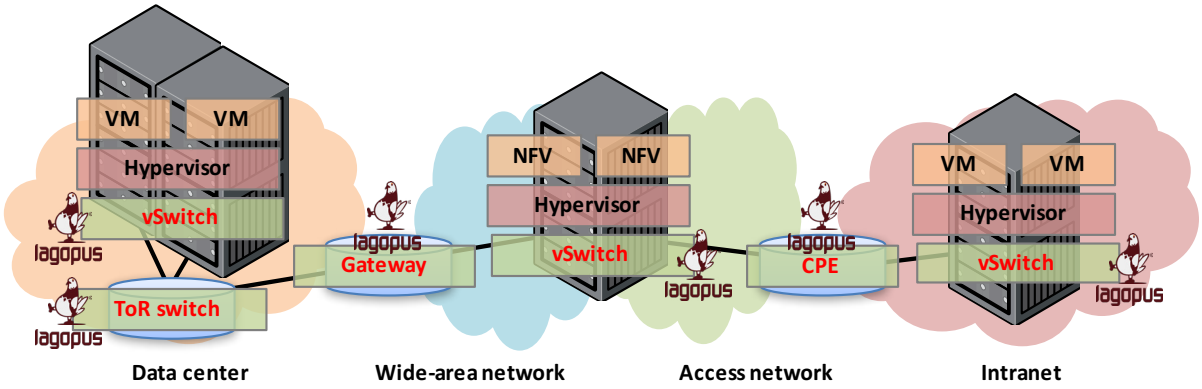

Figure 1. SDN applications in CommSP networks

To expand the SDN application to CommSP networks, not only performance, but also functionality, is also required. As communications service providers look to retrofit legacy network systems with SDN/NFV technology, they often are forced to take an evolutionary approach rather than a radical approach in order to avoid compromising quality. Therefore, many management interfaces and legacy network protocols are required for CommSP SDN to be integrated with legacy network infrastructure and existing OSS/BSS systems, such as MPLS, PBB, QinQ, as well as OAM functionalities. Software-only solutions on general purpose compute provide the flexibility to implement wide varieties of frequently changing standard and proprietary protocols.

As the shift to SDN and NFV takes place, cost pressures make CommSPs sensitive to the cost of developing and testing these solutions. As software solutions target various networks (e.g., data center networks, CommSP networks, enterprise networks, etc.), multiple management interfaces must be developed and many interoperability tests must be performed for various environments. If software solutions are proprietary, the time and costs for development and testing could eliminate the advantages of SDN/NFV.

Solution

Lagopus, a new open source SDN software switch, leverages the Data Plane Development Kit (DPDK) and realizes almost 40 Gbps packet processing and 1 million flow rules.[1] It fully conforms to the OpenFlow switch specification 1.3, including the CommSP network protocol handling, multi tables, group tables, QoS functions, etc. to expand SDN applications to the various networks. It has received the top score from a major OpenFlow switch compliance test suite called Ryu Certification.*[4]

Lagopus is open source software and is available on the GitHub (http://lagopus.github.io/). Lagopus can be easily run on Intel® architecture–based servers, PCs, and also virtual machines. To leverage open innovation in the network industry, the Lagopus community site is open to users. Users can easily report issues to request bug fixes, and open source developers can easily commit their own code for bug fixes and new extensions. Open source development resolves the cost issue of the proprietary development and can enhance the advantages of SDN/NFV software solutions.

DPDK

DPDK is a set of highly tuned software libraries and drivers that unleash the performance of multicore Intel architecture.[5]

The software-based approach of DPDK is greatly beneficial to leading-edge customers who look to reuse software within their product portfolio from top to bottom. DPDK software libraries are optimized for Intel® Xeon® processors, Intel® Core™ processors, and Intel® Atom™ processors. Optimization on these processors helps alleviate customers' concerns about time to market because ease of development and testing on standard Intel architecture allows software developers to focus on only designing their applications. Intel's Tick-Tock model, which represents Intel's continuous

	ACTION (REQUIRED)	ACTION (OPTIONAL)	SET FIELD (OPTIONAL)	MATCH (REQUIRED)	MATCH (OPTIONAL)	GROUP (REQUIRED)	GROUP (OPTIONAL)	METER (OPTIONAL)	TOTAL
LAGOPUS	3 / 0	53 / 0	162 / 8	108 / 0	606 / 0	3 / 0	12 / 0	26 / 10	973 / 18
OPEN VSWITCH*	3 / 0	31 / 22	96 / 74	108 / 0	426 / 180	3 / 0	3 / 9	0 / 36	670 / 321
CPQD	3 / 0	47 / 6	159 / 7	102 / 6	598 / 8	3 / 0	12 / 0	31 / 5	955 / 32

Table 1. Results of Ryu Certification on OpenFlow 1.3[4]

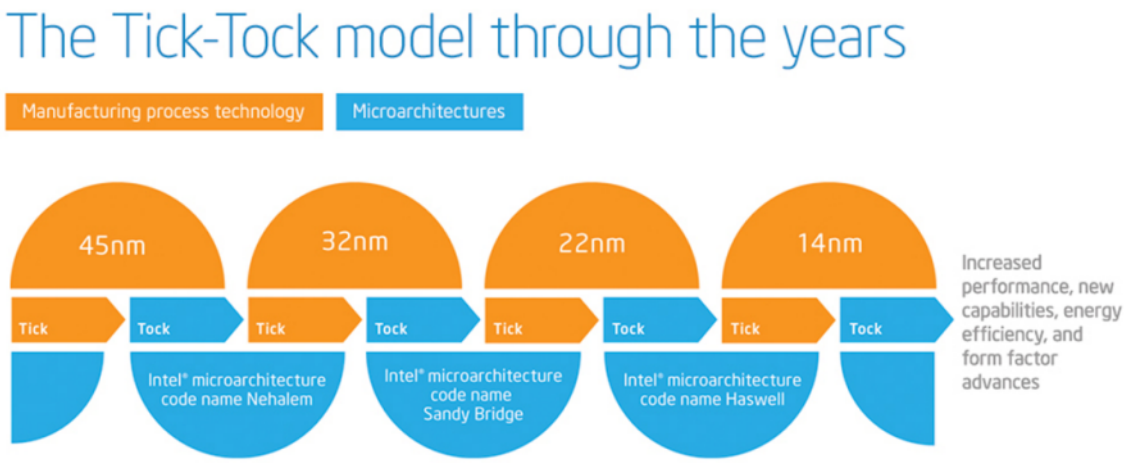

Figure 2. Intel's Tick-Tock model

innovations of leading edge manufacturing process technology and processor microarchitecture design, leverages the performance of software products and helps customers by cutting the cost of new hardware systems with less maintenance cost of software.

Intel architecture has had a proven track record of high performance and innovations for decades, and DPDK has improved packet processing on Intel architecture.

High Performance SDN Switch

Based on DPDK, Lagopus leverages Intel architecture to improve OpenFlow packet processing. Figure 3 shows the packet processing model of Lagopus. Pipelining and parallelizing the packet processing allows Lagopus to take advantage of the Intel architecture multi-processor and multi-core performance and the performance can be scalable by increasing cores. Furthermore, Lagopus applies various optimizations to improve the performance. Packet batching and reduction of memory locks reduce parallel processing overhead. And leveraging memory pre-fetching for packet data and lookup tables improves cache memory efficiency and total throughput.

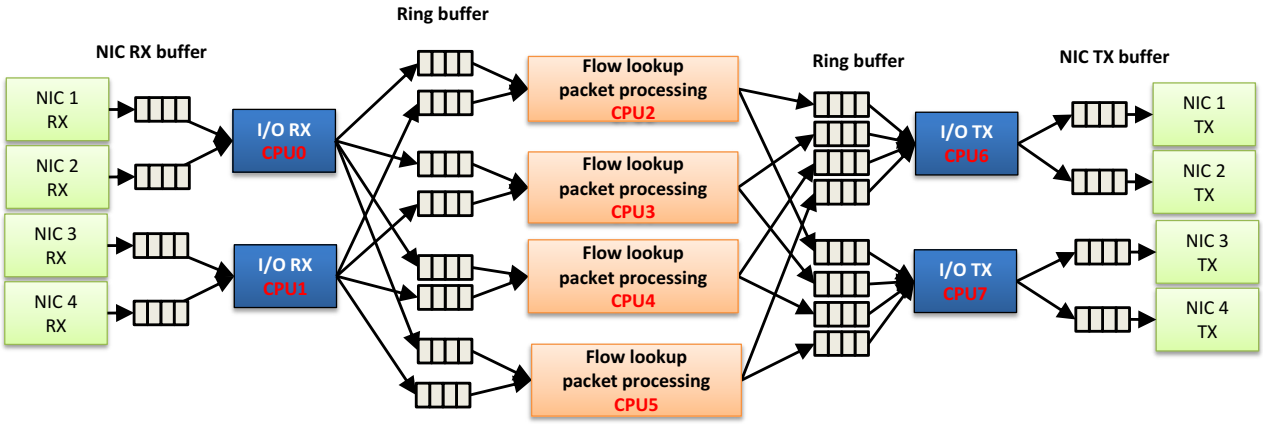

Figure 3. Packet processing using multi core CPU

Figure 4 shows the performance evaluation results. In the evaluation, Lagopus acted as a gateway between the data center and the CommSP network. Lagopus translated and forwarded the VLAN flows in the data center to MPLS VPN flows in the CommSP network. Lagopus was mapping flow rules between VLAN IDs and MPLS labels for each customer. Figure 4 shows throughput results with 1 million flow rules. Lagopus achieved almost 40 Gbps Ethernet wire rate for long packets.

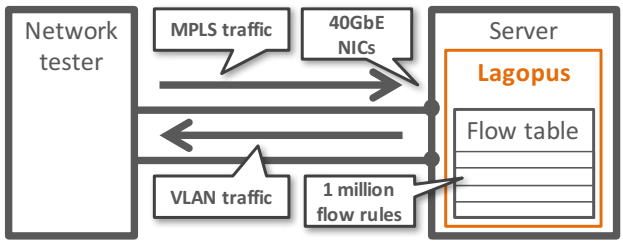

Server specification	
CPU	Intel® Xeon® processor E5-2667 v3 x2 (8 cores are assigned to Lagopus)
Memory	DDR4 64GB
NIC	Intel® Ethernet Controller XL710-QDA2 dual port 40GbE NIC
OS	Ubuntu* 14.04 LTS
DPDK	DPDK 2.2.0

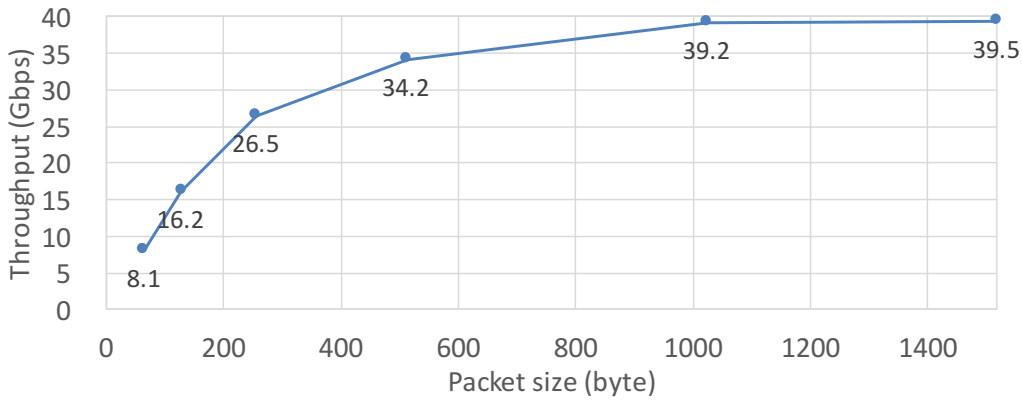

Figure 4. Gateway performance evaluation result[1]

Extensible Architecture

Figure 5 shows a brief design of Lagopus. Lagopus uses a modular architecture to realize software extensibility and improve applicability to the various networks.

Lagopus has two main components, the agent and the data plane. The agent consists of control and management functions. Since the agent fully conforms to the OpenFlow switch specification 1.3, it can be connected with various SDN controllers, e.g. OpenDaylight*[6] and Ryu.*[7] The data plane performs packet processing and forwarding.

All functions can be replaced and extended. For example, Lagopus has a unified switch configuration data store that prevents configuration conflicts by multiple management interfaces. Therefore developers can easily add management interfaces to improve interoperability with various operation systems.

Lagopus also has a switch hardware abstraction layer (HAL) between the agent and data plane; therefore, the whole data plane function can be replaced with different data plane implementations. Today Lagopus is implemented as a software-only data plane with DPDK. A data plane for advanced NICs that are equipped with offloading hardware is now being investigated.

Because the software data plane can be easily deployed on Intel Xeon processor–based servers, it can help CommSPs realize reductions in cost and time to market. But there are performance limitations for some functions, such as the QoS function, due to the software overhead. Advanced NICs have

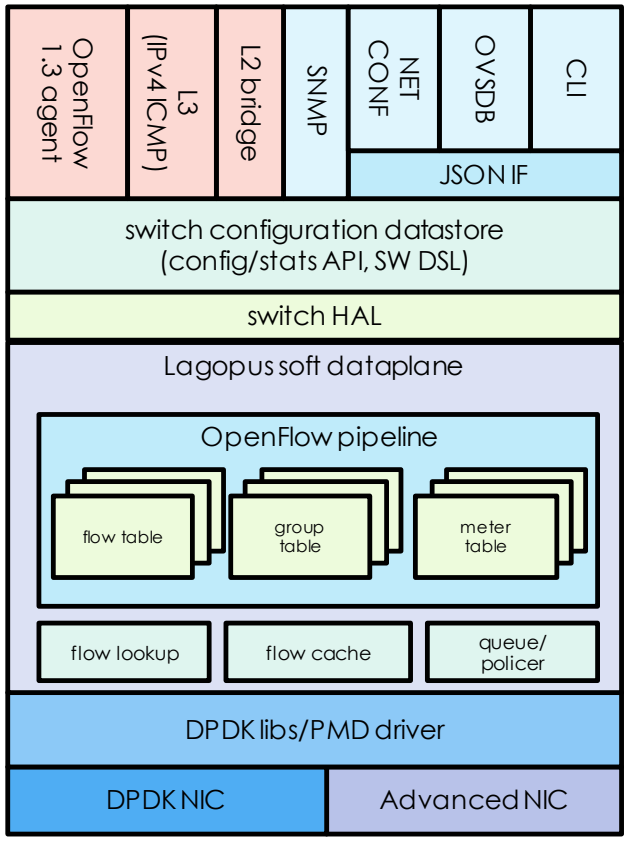

Figure 5. Modular architecture

Intel® Atom™ processor-based server

Intel® Xeon® processor-based server

Intel® Xeon® processor-based server plus advanced NIC

Figure 6. Lagopus platforms

the potential to improve the performance by offloading such functions to hardware. This also reduces the software processing load and provides the CPU resources to other processes, e.g., NFV.

Thus, the extensible architecture of Lagopus expands its platforms and improves its applicability.

Ecosystem

The Lagopus community is now growing and is collaborating with various groups such as business users, universities, and research networks. Lagopus is released under Apache* License 2.0 and users can freely apply it to enhance their own solutions.

Expanding the community leads to improvements in software quality and enhancement of functions. The community involvement will also expand the SDN use cases. NTT is working to continue to develop Lagopus along with the community to meet the needs of a broad set of potential customers.

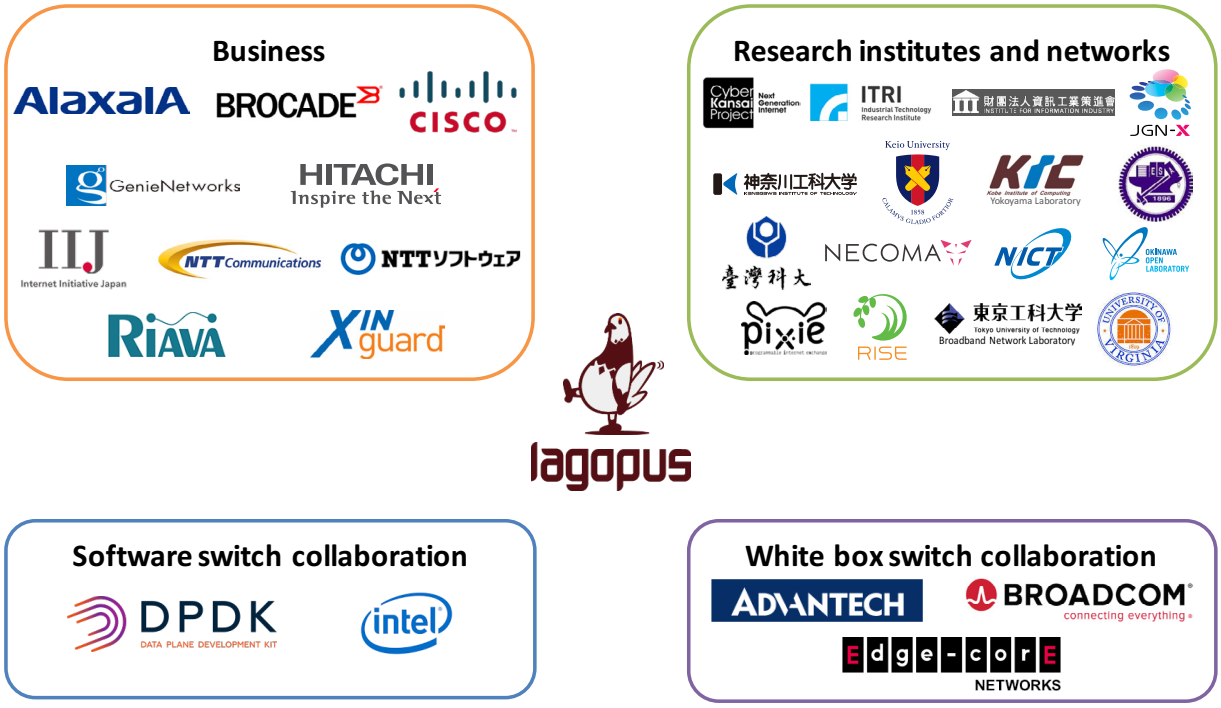

Figure 7. Lagopus community

[1] Test performed by NTT. Configurations: the performance is achieved by running Lagopus on a server that is equipped with an Intel® Xeon® processor E5-2667 v3, DDR4 64GB memory, and Intel® Ethernet Controller XL710-QDA2 dual port 40GbE NIC.

[2] Cisco 2012; Ovum 2012; ABI Research 2012, In-Stat 2012; Analyst Reports, Lit Search Bain Analysis

[3] CIDR REPORT; http://www.cidr-report.org/
[4] For details on the configurations and testing scenarios of Ryu Certification, visit http://osrg.github.io/ryu/certification.html.

[5] Results based on Intel internal analysis. Configurations: 1S Intel Xeon Processor E5-2600 (B0 stepping) 1x8Core, 2.0 GHz, PCIe* Gen2 Performance; More details about the performance of DPDK can be found at http://www.intel.com/content/www/us/en/intelligent-systems/intel-technology/packet-processing-is-enhanced-with-software-from-intel-dpdk.html

[6] OpenDaylight; http://www.opendaylight.org/
[7] Ryu SDN Framework; http://osrg.github.io/ryu/

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

0117/DO/PDF ♻ Please Recycle 335446-001US