SOLUTION BRIEF

Communications Service Providers
Service Assurance

Intracom Telecom Automates NFV Performance Improvement

Tests with colocated packet-processing network functions show the NFV Resource Intelligence platform using Intel® Resource Director Technology optimizes cache allocation for deterministic VNF performance and lower latency.[1]

Service assurance is a paradigm designed to improve deterministic performance in complex network systems and meet service level agreements (SLAs). Now, Intracom Telecom's solution featuring Intel® technology is being applied to virtual network functions (VNFs) to allow communications service providers (CommSPs) to deploy virtualized services that feature deterministic performance necessary for meeting SLAs.

Intracom Telecom is a Solution Partner in the Intel® Network Builders ecosystem. The company's NFV Resource Intelligence (NFV-RI) platform runs on Intel® architecture servers and leverages Intel® Resource Director Technology (Intel® RDT) to provide improved performance on shared cache and memory systems. Intracom Telecom's NFV-RI solution demonstrated reduced packet processing latency and significantly improved deterministic performance[1] in the tests on a forwarding VNF accepting traffic through a virtual switch.

The Challenge

NFV adds agility to network services by the use of non-dedicated, Intel architecture-based servers that can run multiple VNFs. NFV-RI platform allows orchestrated, remote service lifecycle management.

VNF performance is important, as well as its determinism in terms of predictable latency and throughput. In the legacy fixed-function network appliances, deterministic performance could be engineered into the system, at the cost of network agility. In a virtualized deployment, on the other hand, there are certain shared resources, including the CPU and last level cache (LLC), that have a direct effect on the packet processing performance.

The density of software in modern server platforms is increasing due to increasing core counts, the prevalence of containers as a low-footprint alternative to VMs, and other reasons. In this environment, the deployed applications tend to compete for the shared system resources, resulting in the performance being impacted in unpredictable ways as a result of contention for these resources. When an application exhibits "noisy" behavior against others (either temporarily or for a sustained period, and either running independently or as part of a larger service chain), its excessive resource consumption can cause evictions of packet processing code and data of other VNFs, which impacts their performance. In application-dense environments with multiple colocated VNFs, deciding the ideal distribution of resources is a highly complex task that goes beyond human expertise.

Using Intracom Telecom's NFV-RI platform, CommSPs can apply optimization policies and processes to their VNFs to meet service levels promised in their SLAs. In a virtualized environment, service assurance has the job of ensuring that applications get the CPU and LLC resources they need to deliver optimized performance. Once a VM has been given the CPU resources by the operating system or Enhanced Platform Awareness (EPA) software, NFV-RI can then inventory the CPU features of the compute host and exploit them to maximize VNF performance. This resource optimization is enabled using Intel RDT to provide programmatic partitioning and monitoring of the LLC.

NFV Resource Intelligence platform

NFV-RI employs artificial intelligence (AI)-powered exploration to search for resource allocations that optimize the performance of one or more VNFs, in a fully automated fashion, requiring zero domain expertise from the user. The platform considers CPU placement, LLC capacity, DRAM bandwidth, and CPU frequency for each of the colocated workloads. It operates in closed-loop form, taking feedback from performance probes provided by the VNFs to correlate the impact of its resource decisions on their performance.

Once the automated exploration is complete, the NFV-RI defines and assigns shared resources like CPU cores and LLC into chunks dedicated for private use by a VNF. By isolating the compute resources in the right amount according to each VNF's needs, the platform delivers the following:

- High network service performance
- Deterministic network service performance
- High infrastructure utilization
- Low energy consumption

Intel® Resource Director Technology (Intel® RDT)
This resource isolation is expanded using Intel Resource Director Technology (Intel RDT), which provides programmatic partitioning and monitoring of the LLC, DRAM, and CPU. Intel RDT is composed of the following elements:

- Cache Allocation Technology (CAT)
- Cache Monitoring Technology (CMT)
- Code and Data Prioritization (CDP)
- Memory Bandwidth Allocation (MBA)
- Memory Bandwidth Management (MBM)

Combined, these technologies provide the hardware framework to monitor and control the utilization of shared resources. In NFV applications, LLC and DRAM bandwidth are key resources to manage, and Intel RDT technologies can manage these workloads across shared resources.

Intracom Telecom's NFV-RI supports VNFs implemented as KVM virtual machines, Docker containers, Kubernetes pods, or native Linux applications. With support for these compute modes, CommSPs can use the platform to manage resources for services that are composed of deployment units from multiple virtualization technologies. Service chains can be built that include VMs, native applications, and container-based VNFs in a wide range of combinations. The NFV-RI identifies and handles the unique needs of special classes of VNFs, like user-level soft switches, which need proper CPU placement to avoid resource conflicts with other VNFs.

Tests Demonstrate Automated Latency Optimization

Intracom Telecom tested the NFV-RI's latency optimization automation feature and demonstrated both its ability to reduce latency and to provide more deterministic performance.[1] The test scenario involved two Intracom Telecom-developed Data Plane Development Kit (DPDK) VNFs, one of which is a router and the other one a simple packet forwarder, accepting traffic through a DPDK-based Open vSwitch. The server has additional best-effort workloads on the same CPU socket that impact the latency performance of all three VNFs (i.e., router, forwarder, switch), which are highly sensitive to resource contention from colocated workloads.

The tests were designed to efficiently eliminate contention from colocated workloads. In many real-world applications, there is little or no contention, but there is always a possibility that the addition of a VNF would introduce some contention. In the cases where there is contention, although few, the performance of the VNFs can be significantly impacted in terms of latency and packet throughput. The tests were specifically designed to deliberately create contentious scenarios that could be improved by NFV-RI's AI-powered optimization engine.

Intel RDT was deployed on the server for LLC slicing (CAT) and memory bandwidth throttling using MBA functionality. As will be seen in the test results, these features improve the latency for the colocated VNFs when used in the default OS placement and when CPU pinning is deployed.

Test Setup
The tests were run on a server platform powered by 2.1 GHz Intel® Xeon® Gold 6252 CPUs. The vRouter used in the test was developed in-house, based on DPDK and featuring four packet processing states performing hash-based lookups each of which performs more elaborate packet processing. The vForwarder was also an in-house developed VNF that can switch packets using one packet processing stage. A DPDK-based Open vSwitch (OVS) was used to switch traffic between VNFs and the Ethernet NIC. For the "noisy neighbor" VNFs, two instances of an in-house developed streaming app (STR) were used to perform streaming-like memory accesses. Every workload ran in a QEMU virtual machine, except for the Open vSwitch, which ran on bare metal. Packet generation and measuring was done using two instances of the TRex traffic generator running on separate sockets, each feeding a separate VNF.

There were three deployment modes used in the tests: default placement by the operating system (OS) with no CPU pinning, and NFV-RI-optimized placement using both Enhanced Platform Awareness (EPA), and EPA combined with Intel RDT. EPA represents a methodology and a related suite of changes across multiple layers of the orchestration stack targeting intelligent platform capability, configuration, and capacity consumption. The key EPA feature used in these tests was CPU pinning, which assigns a VNF to a specific CPU.

The default OS placement was the baseline to show the improvement of both the EPA and the EPA/Intel RDT placements. In the CPU pinning tests all workloads under test were colocated on a single socket, and each workload got a dedicated slice of the CPU core. The CPU pinning plus Intel RDT tests differed in that the latency critical VNFs and the OVS got a larger portion of the LLC, while the STRs got a minimal LLC slice. The capacities of all slices were decided by NFV-RI's automated workflows.

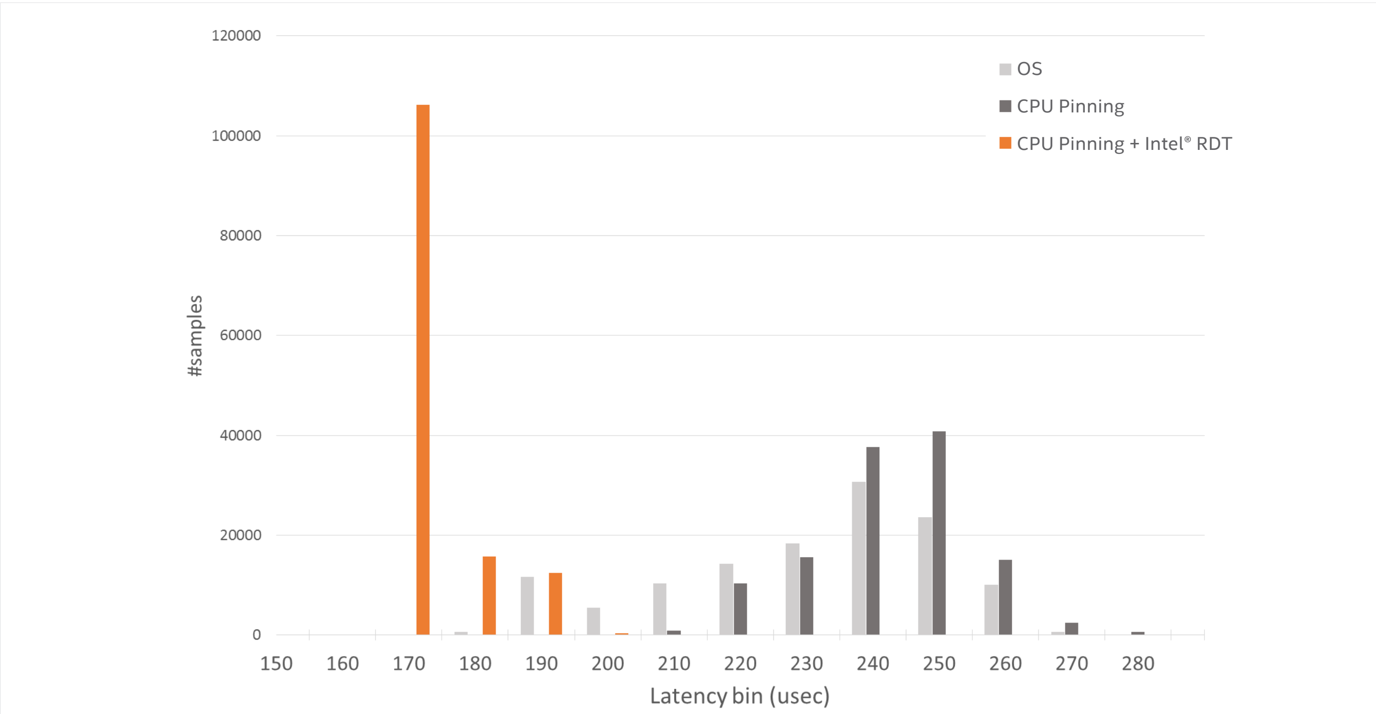

Figure 1. vRouter performance measured in microseconds (µs) under three test scenarios.[2]

Test Results

The test results in Figure 1 show that latency using just the OS default placement policy had a wide variation of results, ranging from 180 microseconds (µs) to 270 µs, with the highest number of samples in the 240 µs and 250 µs bins. The OS results featured the highest latency in the tests, and the wide variation reduced performance determinism.

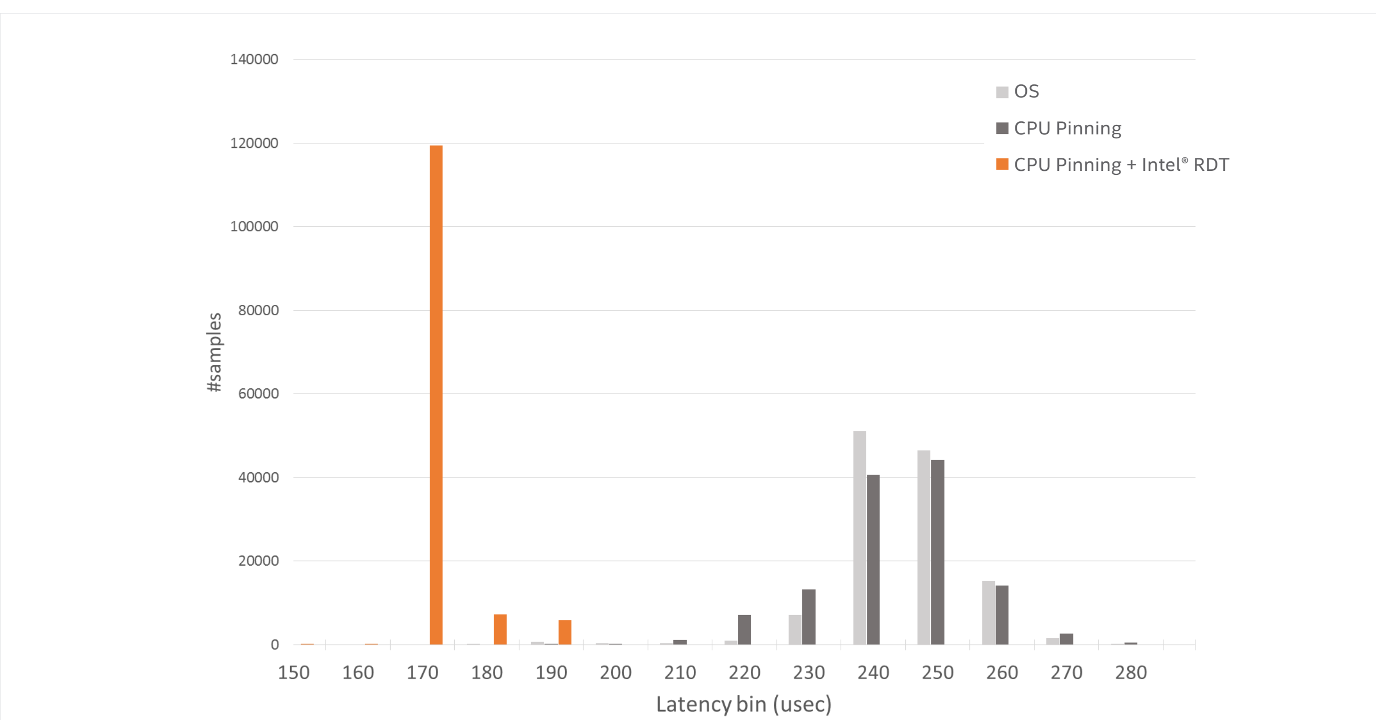

Figure 2. vForwarder performance results under three test scenarios.[2]

While CPU pinning improved the performance, these results show that deterministic latency was only slightly better than OS placement. In the CPU pinning results, the distribution of samples ranged from 210 to 280 µs, with a somewhat higher concentration of samples at 240 and 250 µs.

The EPA plus Intel RDT results show both the lowest latency and the most deterministic results with most of the results clustered on the 170 µs bin.

These results are for a vRouter, but Intracom Telecom simultaneously measured the impact on vForwarder traffic, which returned similar results (see Figure 2).

Conclusion

Using AI-based closed-loop optimization technology, Intracom Telecom's NFV-RI is able to optimize NFV performance for lower latency and in the same time offer deterministic performance.[1] The solution benefits the Intel RDT technology to provide critical DRAM and LLC information that offers performance optimization that is beyond the improvements that come from CPU pinning by itself. With NFV-RI, CommSPs can see NFV performance improvements and know that their services are optimized to meet customer needs.

About Intracom Telecom

Intracom Telecom is a global telecommunication systems and solutions vendor operating for over 40 years. The company has become the benchmark in fixed wireless access and it innovates in the 5G/4G wireless fronthaul, backhaul and small-cell SON backhaul international arena. Intracom Telecom offers a comprehensive portfolio of revenue-generating software solutions and a complete range of ICT services, focusing on IoT, SDN/NFV, big data analytics and data-driven intelligence, and Smart City solutions. The company also addresses the energy and utilities industries, emphasizing smart metering and end-to-end IT solutions. The company has a proven track record in the market, serving more than 100 renowned customers in over 70 countries. Its subsidiaries span across Europe, Russia/CIS, the Middle East and Africa, Asia and North America. More information is at http://www.intracom-telecom.com/index.htm

About Intel® Network Builders

Intel Network Builders is an ecosystem of infrastructure, software, and technology vendors coming together with communications service providers and end users to accelerate the adoption of solutions based on network functions virtualization (NFV) and software defined networking (SDN) in telecommunications and data center networks. The program offers technical support, matchmaking, and co-marketing opportunities to help facilitate joint collaboration through to the trial and deployment of NFV and SDN solutions. Learn more at http://networkbuilders.intel.com.

Notices & Disclaimers

[1] Tests conducted by Intracom Telecom in November 2019. Configurations: Server platform was based on 2.1 GHz Intel® Xeon® Gold 6252 CPUs (Microcode: 0x400001c). Installation was comprised of 96 logical threads in total, organized in two physical packages of 24 cores each with each core having hyper-threading enabled. The server featured 192 GB of DDR4, 2666 MT/s RAM, and Intel® Ethernet Controller XXV710 for 25 GbE SFP28 (rev 02). Workloads: The vRouter and vForwarder were internally developed and based on DPDK (version 17.11.3). Open vSwitch version: 2.11.0 used DPDK version: 18.11.1, OVS. Every workload ran in a QEMU version 3.1.0 VM, except for the Open vSwitch which ran on a bare metal server.

[2] See end note 1 for configurations. Figures provided courtesy of Intracom Telecom.

1020/DO/H09/PDF ♲ Please Recycle 342031-002US