SOLUTION BRIEF

Communications Service Providers
Service Assurance

Intracom Telecom* Improves VNF Placement for Network Services

The placement of VNFs on servers can have a significant impact on performance. Intracom Telecom's NFV Resource Intelligence platform discovers ideal server placements for each VNF and then monitors performance and energy consumption, utilizing data provided by Intel® Xeon® processors.

The Challenge

Network functions virtualization (NFV) is transforming the telecommunications industry by changing the way networks are built and operated. Standalone network functions like soft switching, edge routing, and firewalling, as well as composite network services like evolved packet core and IP multimedia subsystem are increasingly being migrated to software-based implementations.

Performance is a key factor for NFV adoption: first, virtual network functions (VNFs) must offer performance that is competitive with legacy, fixed-function alternatives. This means packet throughput that achieves carrier-grade key performance indicators (KPIs) including wireline data rates, low latency, low packet loss, and others.

In addition, VNFs must feature predictable performance so that the communications service providers (CommSPs) can confidently plan on certain service level objectives (SLOs). The challenge of predictability is that the VNFs might be assigned to run on spacious and controlled execution environments, or on highly consolidated, uncontrolled environments with no control on resource use of other VNFs (noisy neighbors). This is particularly important in emerging 5G networks, where latency-critical VNFs with tight latency SLOs need to run alongside other, diverse services on a single converged infrastructure.

Conventional NFV infrastructure (NFVI) approaches don't automatically place VNFs for very high performance, whether that's due to a lack of the suitable technology for optimal infrastructure utilization, or because the NFVI is configured too conservatively in terms of its ability to optimize for performance.

For example, latency-critical VNFs are usually deployed in isolation, perhaps being assigned an entire CPU socket or even a server, in order to eliminate contention with noisy neighbors for shared resources that would increase their latency. Similarly, throughput-sensitive workloads might be also deployed in isolation just to have enough resources reserved for demand peaks, even if their virtual CPUs sit almost idle and could be used for another VNF. Approaches like these ultimately guarantee deterministic performance under all conditions, but they lead to inefficient infrastructure usage due to low resource utilization, negatively impacting energy efficiency and capital and operating costs.

Achieving high throughput and predictable performance, and at the same time maintaining high levels of infrastructure and energy efficiency, is a tricky tradeoff.

The biggest challenge for automated or manual deployment processes (e.g., orchestrators, administrators) is to optimally schedule hardware resources for a virtual network service (NS) so that its SLOs are met, while also maximizing hardware utilization, which positively affects costs and energy usage.

This challenge has two aspects: intra-NS and inter-NS resource allocation. For intra-NS resources, the goal is to optimally allocate resources for every VNF that makes up the NS so their individual KPIs collectively satisfy a committed SLO. This problem becomes harder as NFV matures from monolithic, appliance-type network services with few VNFs, to microservice-type services with many VNFs. Inter-NS resource allocation adds a further complexity to this problem, as multiple, independent network services with possibly conflicting resource requirements need to colocate on the same machine without performance losses, as architectures becomes wider and network services more modular.

To help CommSPs solve this challenge, Intel® Network Builders ecosystem partner Intracom Telecom* has developed its own solution that adds intelligence for optimized execution on top of existing NFVI platforms.

Intracom Telecom's NFV Resource Intelligence platform

Intracom Telecom addresses the above challenges via its NFV Resource Intelligence platform, a resource intelligence solution for optimized execution of virtualized network services and optimized operation of the infrastructure where they run. The platform delivers:

- **High NS performance:** by discovering the right resource allocations for network service VNFs
- **Deterministic NS performance:** by isolating and protecting VNFs from noisy neighbors consuming resources like CPU processing cycles and the CPU's last level cache (LLC)[1]
- **High infrastructure utilization:** by enabling dense VNF consolidations within a single CPU socket, through fine-grain allocations of LLC capacity to VNFs
- **Low energy consumption:** by tuning resources and CPU frequencies towards energy-efficient configurations without compromising committed SLOs

To meet these goals, the NFV Resource Intelligence platform slices shared resources like CPU cores and LLC into chunks dedicated for private use by a VNF. This is enabled using Intel® Resource Director Technology (Intel® RDT) to provide programmatic partitioning and monitoring of the LLC.

The platform supports VNFs implemented as KVM* virtual machines, Docker* containers, Kubernetes* pods, or native Linux* applications. This broad virtualization support allows resource management for hybrid network services that are composed of service chains that include VMs, native applications, and container-based VNFs in a wide range of combinations. The NFV Resource Intelligence platform 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.

The NFV Resource Intelligence platform acts on VNFs after they have been deployed on servers in the customer's data center, regardless of whether they were deployed by an NFV orchestrator, cloud orchestrator, or even manually. For each deployed NS, the platform runs automated stress tests whose goal is to evaluate alternative resource allocations for each of the VNFs in terms of performance and energy. The resources considered by the platform include CPU placement (via several supported policies), LLC capacity, and CPU frequency. The platform uses the element managers provided by the VNFs to read the latest NS KPI values, as an indication of the impact of a certain resource allocation on performance. This is a key requirement for the platform to be able to drive the optimization process.

The platform tests many possible combinations of resource allocations for an NS in a policy-driven and consistent way. In the end, the platform exposes the performance and energy "footprint" of each tested VNF combination via a comprehensive *parallel coordinates* interactive visualization.

From this, the user is able to further explore the test results, narrowing them down to those meeting any desired combination of any number of NS KPI and platform energy constraints. With a click of a button, CommSPs can create an optimization profile out of a shortlisted combination, save it for permanent use, and apply it in real time. At any time, the user can switch between multiple defined optimization profiles in real time in order to drive the NS into different operating modes (e.g., best performance, maximum economy, etc.). The diagrams in Figure 1 show some of the variables for which optimization is available.

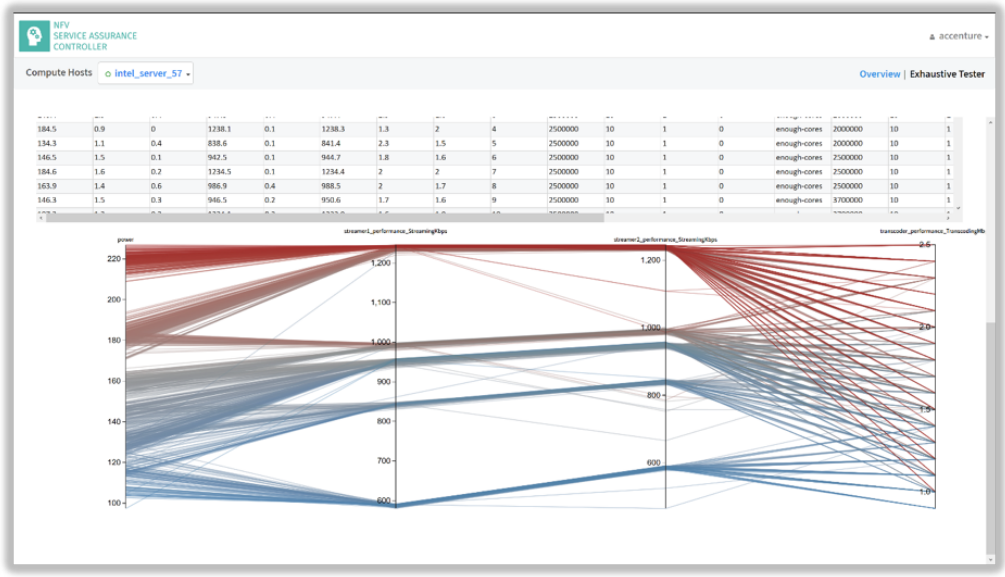

(a)

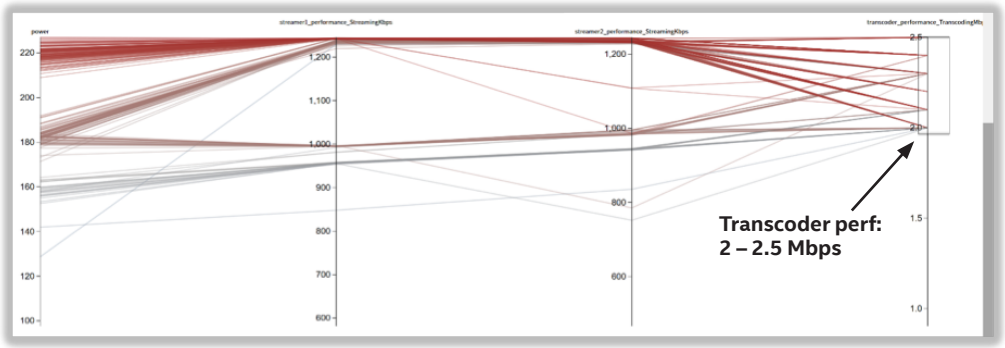

(b)

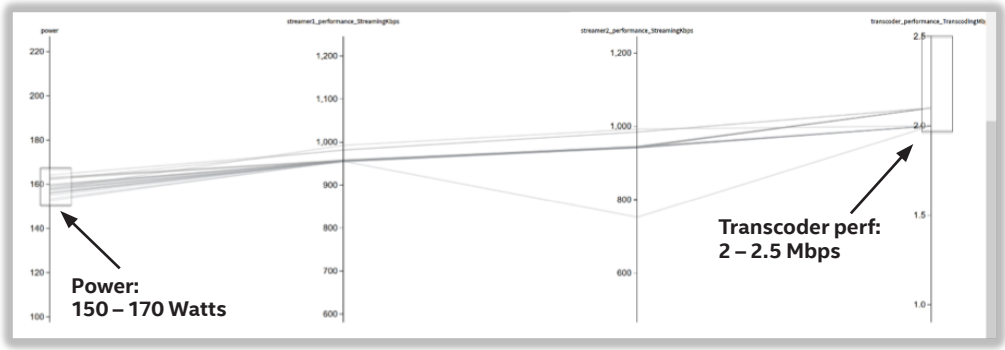

(c)

Figure 1. Inspection of stress tests results (a) and progressive shortlisting based on NS KPI and platform energy constraints (b), (c).[2]

The NFV Resource Intelligence platform can be offered either as a cloud service or as a self-hosted solution that can be deployed and operated in a private data center. The platform makes use of Smart Agents that are installed on every NS server that hosts a VNF. The Smart Agent communicates with a central component logic (controller) via a security-enabled channel to receive resource allocation directives, and send back telemetry data and other information (platform capabilities, VNF events, etc.).

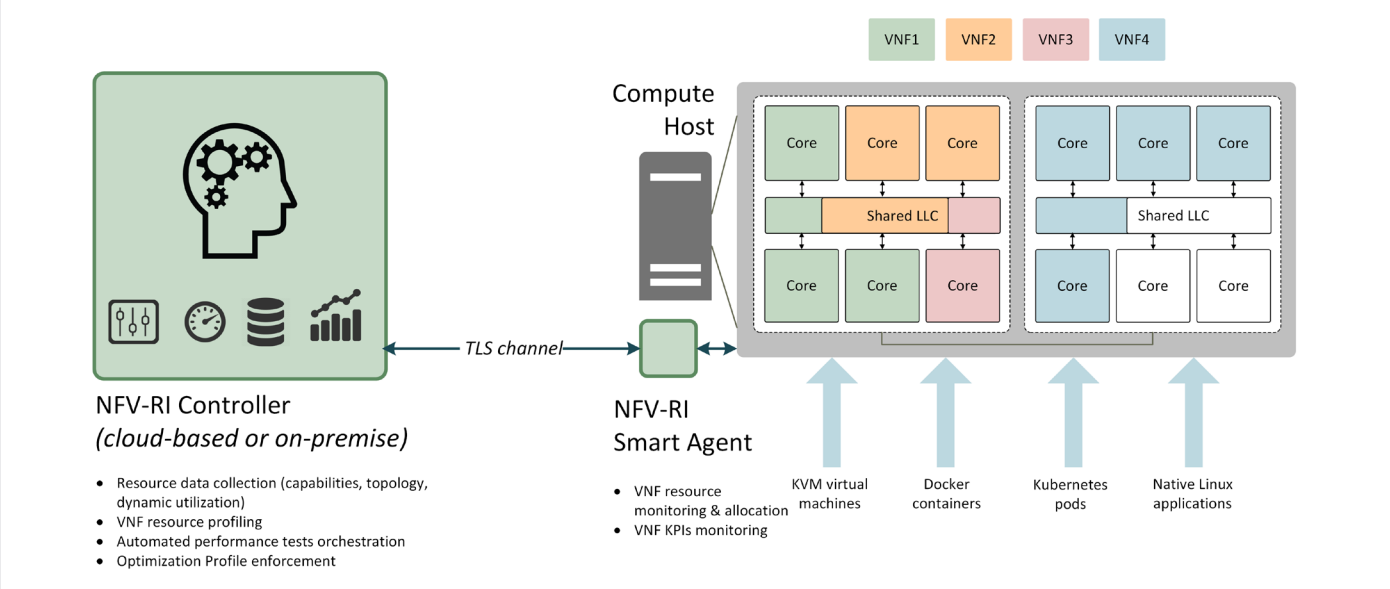

Figure 2. NFV Resource Intelligence platform architecture.

Intel® Components

Intracom Telecom specifies server platforms based on latest Intel® Xeon® processors for the NFV Resource Intelligence platform because the CPUs provide several key features for monitoring VNF optimization criteria. These include:

- Cache Allocation Technology (CAT) to allocate dedicated LLC slices to VNFs
- Cache Monitoring Technology (CMT) and Memory Bandwidth Monitoring (MBM) to monitor VNF LLC occupancy and memory bandwidth utilization, respectively
- Enhanced Intel SpeedStep® Technology to scale VNF frequencies
- Running Average Power Limit (RAPL) processor feature for CPU/memory power monitoring
- Intel® Node Manager for server power monitoring

The NFV Resource Intelligence platform is able to automatically detect the presence of each feature on a target machine during start-up, and automatically include it in the set of hardware features it will leverage to control and monitor resources on that machine.

Conclusion

Delivering high performance that is predictable is essential for NFV implementations to meet network KPIs. The NFV Resource Intelligence platform offers policy-based placement intelligence above and beyond NFVI capabilities that matches VNFs to the right server resources. Utilizing data from Intel Xeon processors, the NFV Resource Intelligence platform can deliver meaningful performance improvements for network services and can monitor changes in the server environment to maintain the right environment for these VNFs throughout the service lifecycle.

NFV Resource Intelligence platform Key Features

- Optimized NS performance by discovering ideal VNF resource allocations
- Deterministic NS performance by allocating hardware resources (cores, LLCs) for dedicated use by the VNFs
- Improved infrastructure utilization through dense VNF colocations
- Low energy consumption through elaborate power conservation techniques (per-VNF frequency scaling)
- Stress testing module for automated evaluation of thousands of possible VNF resource allocation combinations
- State-of-the-art UI for interactive exploration and shortlisting of stress test results, based on multiple NS KPIs and platform energy constraints
- Creation of Optimization Profiles out of shortlisted solutions, and real-time enforcement on the deployed network service
- Multiple VNF implementations supported (KVM VMs, Docker containers, Kubernetes pods, native Linux applications)
- Real-time collection and visualization of both application-specific KPIs and system-wide utilization/energy metrics
- Intuitive visualization of platform's resources slicing and utilization
- Cloud service and self-hosted offerings

About Intracom Telecom

Intracom Telecom is a global telecommunication systems and solutions vendor operating for over 40 years in the market. The company innovates in the wireless access and transmission field, offers a competitive telco software solutions portfolio and combines its offerings with a complete range of professional services. In the SDN/NFV area, Intracom telecom is committed to accelerating the adoption of SDN and NFV, creating solid foundation for NFV. As part of ongoing roadmap to innovation in the telco SDN/NFV domain, Intracom Telecom caries out advanced R&D activities in the SDN/NFV domain.
http://www.intracom-telecom.com/nfvRI

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] Memory bandwidth will be also addressed in upcoming releases.

[2] Figures provided courtesy of Intracom Telecom. The results presented in these figures were produced by Intracom Telecom on a compute host featuring dual Intel® Xeon® Platinum 8168 processors @ 2.7 GHz and 400GB DDR4 RAM @ 2.6 GHz. The operating system used was Ubuntu* 16.04.4 LTS.

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.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

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

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

1020/DO/H09/PDF ♻ Please Recycle 337865-002US