SOLUTION BRIEF

Cloud Software
Highly Distributed Containerized Cloud Platform

Intel and Red Hat Are Developing Advanced NFVi for Carrier-grade Virtualized Infrastructure

Executive Summary

Intel and Red Hat are collaborating to build a fully integrated network functions virtualization infrastructure (NFVi) solution to help telecom service providers compete more effectively in an ever-changing business environment. This solution, featuring industry-standard high-volume servers (SHVS), can help reduce capital expenditure (CapEx) and operating expenditure (OpEx) and enables telecom service providers to deploy network services and applications at enterprise speed within an advanced microservices architecture. This solution brief describes the Intel and Red Hat NFVi solution and presents a basic set of performance data.

Introduction

Telecom service providers must increasingly compete with cloud service providers to provide not only cutting-edge applications, but even traditional network services. The approach that many are taking is to transform their networks from legacy network functions deployed in special-purpose hardware to virtualized network functions (VNFs), all deployed on SHVS.

This approach allows telecom service providers to scale services up and down in response to customer traffic without overprovisioning their networks with specialized legacy equipment. It also allows them to roll out new services quickly in response to changing consumer preferences and shifting business requirements.

Benefits

The Intel and Red Hat NFVi solution brings the strength of both companies to bear in serving the business needs of telecom service providers. Intel provides advanced server, networking, and storage technologies, in addition to experience developing virtualization technologies for enterprise and cloud customers; and Red Hat brings hardened, enterprise- and telecom-ready operating systems and virtualization technologies.

The NFVi solution integrates an advanced microservices architecture to enable dynamic provisioning of VNFs across a wide range of services by virtualizing the firewall, load balancing, router, deep packet inspection (DPI), and other workloads to deliver carrier-grade reliability.

Providing automated NFV orchestration and VNF management, the Intel and Red Hat NFVi solution delivers VNF health monitoring and lifecycle management to help ensure high availability and automated network-service delivery.

Built with the flexibility to accommodate components from a wide variety of commercial third-party vendors, the Intel and Red Hat NFVi solution provides the following benefits, while remaining fully serviceable and adaptable to future requirements:

- **Scale.** The solution can scale up or out to address a wide range of automated deployment scenarios and conditions.
- **Open infrastructure.** Built with commercially available, modular components using industry-standard high-volume servers, storage, and networking, the solution optimizes interoperability with third-party commercial components. This eases the cost and complexity of managing the network via a broad set of open interfaces with increased operational automation and simplified operational processes.
- **Open development environment**. Broadly available to the NFV developer community, the solution integrates foundational open source community software, such as the Data Plane Development Kit (DPDK), to improve the performance of VNFs.

Solution

Intel and Red Hat have joined forces with other third parties to create a solution for optimal interoperability and manageability through a broad set of open interfaces.

The Intel, Red Hat Enterprise Linux*, and OpenStack* Platform

As shown in Figure 1, the NFVi solution runs on Red Hat Enterprise Linux* (RHEL*) 7.3 atop a foundational layer of physical devices across compute, network, and storage components. This supports a virtual infrastructure manager (VIM) that uses the Red Hat OpenStack* platform (RHOSP*) version 10.0 to automate complex network overlay provisioning and management tasks. A complete European Telecommunications Standards Institute (ETSI) NFV system is created from this solution by adding management and orchestration (MANO) components and VNFs.

Installer scripts based on Ansible* enable dynamic modular provisioning of VNFs orchestrated via the VNF manager. This orchestration provides cataloging features and templating capabilities that simplify the deployment and management of VNFs to deliver an easy-to-use VNF library.

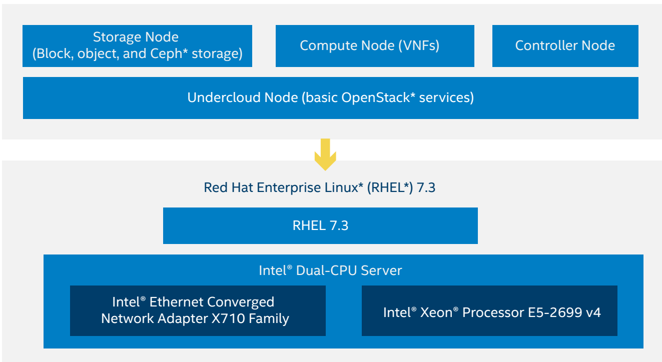

Figure 1. Intel and Red Hat NFVi solution components

NFVi Characterization Testing

NFVi characterization is the first step in problem solving for telecom-service-provider customers. Intel and Red Hat are committed to rigorously evaluating and testing performance across multiple configurations. This section provides an overview of the testing environment and performance results.

Characterizing an NFVi platform consists of measuring the performance of VNF workloads that represent "typical" VNFs that are executed on the platform in key test scenarios under different configurations. These are referred to as "test VNFs," because they implement core functions that might occur in production VNFs, such as layer-2 or layer-3 packet forwarding, as shown in Table 1.

Table 1. Test VNF workload functions

Test VNF Workload	Function
Layer-2 packet forwarding	Forwarding a layer-2 packet, such as an Ethernet frame, without consuming the packet and producing a new packet
Layer-2 packet forwarding	Forwarding a layer-2 packet while consuming a packet and producing a new packet
Multiprotocol Label Switching (MPLS) tagging	Adding or removing an MPLS tag to or from a packet
Load balancing/ quintuple lookup	Looking up fields from a packet in a flow-match table
Access control list	Matching fields in a packet to entries in an access control list
Buffering	Buffering packets in a flow for a finite time, such as 125 milliseconds before forwarding

NFVi Platform Configuration

Server platforms can be configured in multiple ways, including:

- Using bare metal (without virtualization)
- Deploying VNFs on virtual machines by using either unaccelerated or accelerated virtual switches (such as switches optimized with DPDK)
- Running another interconnect technology, such as single root input/output virtualization (SR-IOV) or Peripheral Component Interconnect Express (PCIe*) passthrough

Characterizing NFVi platform performance under multiple configurations gives important baseline data to telecom service providers needing to compare the performance of various orchestration policies methods under evaluation.

VNF Testing Method

The traffic received by each test VNF consists of packet flows of varying sizes from 64 bytes for voice traffic to 1,518 bytes for data packets in non-jumbo Ethernet frames. The traffic is transmitted at different line rates, beginning at a low rate and increasing until the NFVi begins to drop packets or until the physical line rate of the hardware is reached. The test measures the maximum rate that test VNFs can perform without dropping packets.

NFVi Characterization Results

Traffic was tested on the following four configurations:

- Host operating system in "bare-metal" mode
- Open vSwitch* to capture the performance of Open vSwitch independent of VNFs
- VNFs running in a virtual machine (VM)
- Simple Service Function Chain (SFC) consisting of two VNFs on VMs

Note: Traffic was tested on RHEL 7.3 only—pending future tests using both RHEL 7.3 and RHOSP 10.0.

All test-result details will be published upon completion.

NFVi Layer-2 Packet Forwarding Performance

Testing on VNFs for layer-2 packet forwarding was conducted using the same configurations described earlier.

Figure 2 shows the different configurations' results, with each platform configuration represented by a different line color. Each line reports the maximum data rate achieved without dropping packets for packets of different sizes. The data rate is reported as a percentage of the physical line rate, which is 10 gigabits per second (Gbps) for this server platform. As expected, the tests indicated the following:

- Performance is much lower for small packet sizes than for large packet sizes due to the overhead of processing packet headers as packets are forwarded through a system
- Using a virtual switch incurs significant overhead due to the longer processing pipelines for packet forwarding compared to other system configurations, such as SR-IOV or PCIe passthrough

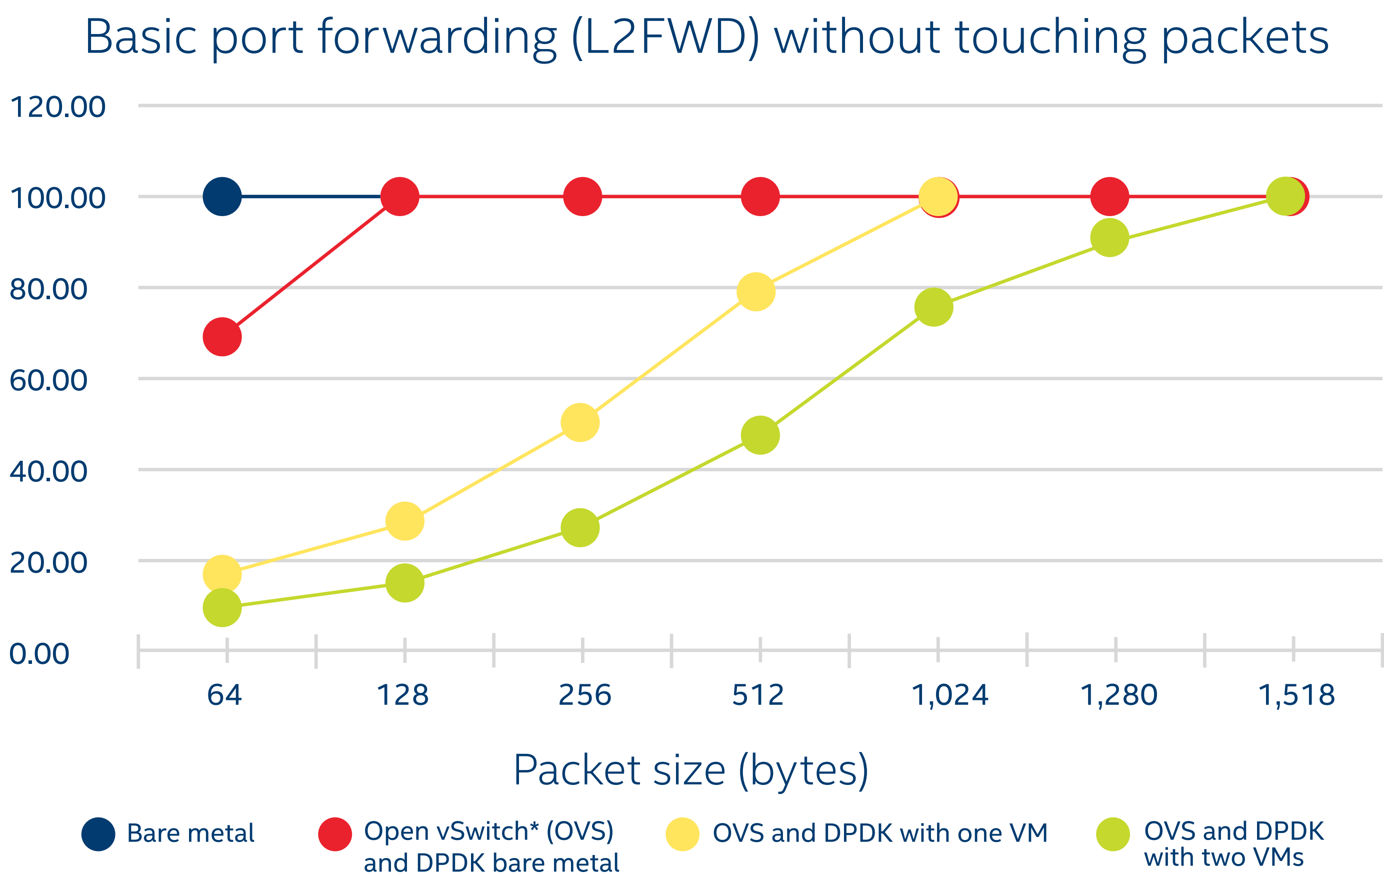

Figure 2. Layer-2 packet forwarding with different platform configurations

Conclusion

The Intel and Red Hat NFVi solution utilizes open source assets to enable a highly available, fully converged virtualized infrastructure that provides companies with the ability to dynamically respond to rapidly changing market conditions. A comprehensive solution with health monitoring for continual high availability, the Intel and Red Hat NFVi solution provides the carrier-grade reliability needed to compete in an increasingly crowded application-services market.

These performance results are the beginning of a thorough performance characterization of the Intel and Red Hat NFVi solution. Intel and Red Hat will extend this work to NFVi configurations including additional Red Hat products and use cases prioritized by their customers.

Next Steps

For more information, visit **networkbuilders.intel.com** and **redhat.com/en/technologies/industries/telecommunications**.

Printed in USA 0517/SH/PRW/PDF Printed in USA 335916-001US