Intel® Open Network Platform Release 2.1 Reference Architecture Guide

SDN/NFV Solutions with Intel® Open Network Platform

Document Revision 1.0

March 2016

Revision History

Date	Revision	Comments
March 31, 2016	1.0	Initial release of Intel® Open Network Platform Release 2.1

Contents

Figures

Tables

1.0 Audience and Purpose

Intel® Open Network Platform (Intel® ONP) is a Reference Architecture that provides engineering guidance and ecosystem-enablement support to encourage widespread adoption of Software Defined Networking (SDN) and Network Functions Virtualization (NFV) solutions in Telco, Enterprise, and Cloud.

The primary audiences for this document are architects and engineers developing or testing SDN/NFV solutions and looking to use the Intel® Open Network Platform Reference Architecture for reference. This document provides step by step instructions on how to build, configure and operate the system using open-source software. Ingredients include the following:

- OpenStack*
- OpenDaylight*
- Data Plane Development Kit (DPDK)*
- Open vSwitch* (OvS)
- Fedora 23*
- CentOS-7.2*

This document provides a guide for integration of these software elements on the Intel® Architecture (IA) platform for the Intel® ONP reference architecture. The content includes high-level architecture, setup, configuration and provisioning procedures. This information is intended to help architects and engineers evaluate Network Functions Virtualization (NFV) and Software Defined Networking (SDN) solutions with Intel® ONP.

The purpose of documenting configurations is not to imply any preferred methods. Providing a baseline configuration of well-tested procedures, however, it can help achieve optimal system performance on an IA platform when developing an NFV/SDN solutions.

Please note that Intel® offers a scripts available on 01.org to facilitate the installation of Intel® ONP reference software stack.

2.0 Summary

The Intel® ONP uses open-source software to help accelerate SDN and NFV commercialization with the latest Intel® Architecture Communications Platform. This document describes how to set up and configure the Controller and Compute Nodes for evaluating and developing NFV/SDN solutions using Intel® Open Network Platform ingredients.

Supported processor families include:

- Intel® Xeon® processor E5-2600 v4 product family
- Intel® Xeon® processor E5-2600 v3 product family
- Intel® Xeon® processor D-1500 family

Supported networking adapters include:

- Intel® Ethernet Server Adapter X520 Series
- Intel® Ethernet Converged Network Adapter X540-T2
- Intel® Ethernet Converged Network Adapters XL710-QDA2 and X710-DA4.

Supported host operating systems include:

- Fedora 23 with QEMU-KVM virtualization technology
- CentOS-7.2 with QEMU-KVM virtualization technology.

Additional software ingredients include Intel® DPDK, Open vSwitch, Open vSwitch with DPDK, OpenStack, and OpenDaylight. Figure 2-1 shows the corresponding version information for the components involved. For the list of new features and improvements, see *Intel® ONP Release 2.1 Release Notes*, available on 01.org.

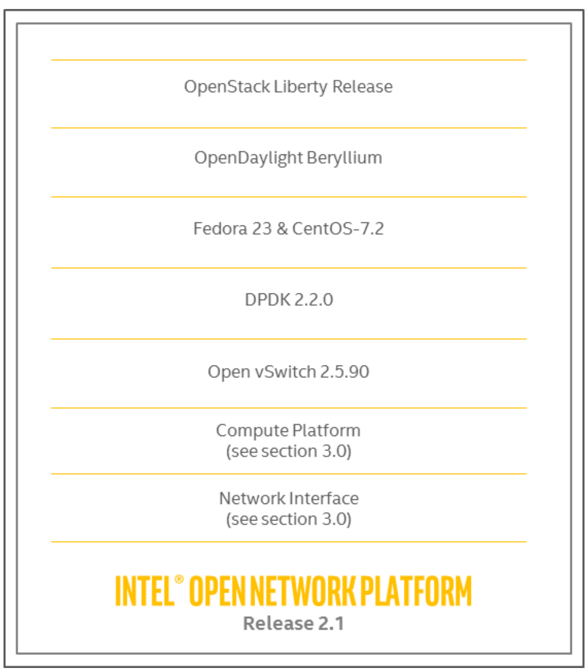

Figure 2-1 Intel® ONP 2.1 — Hardware and software ingredients

Figure 2-2 shows a generic SDN/NFV setup. In this configuration, the Orchestrator and Controller (Management and Control Planes) run on one server, and the two Compute Nodes (Data Plane) run on two individual server nodes.

The differences in the network configuration to enable this setup are shown with the management and data ports. Note that many variations of this setup can be deployed.

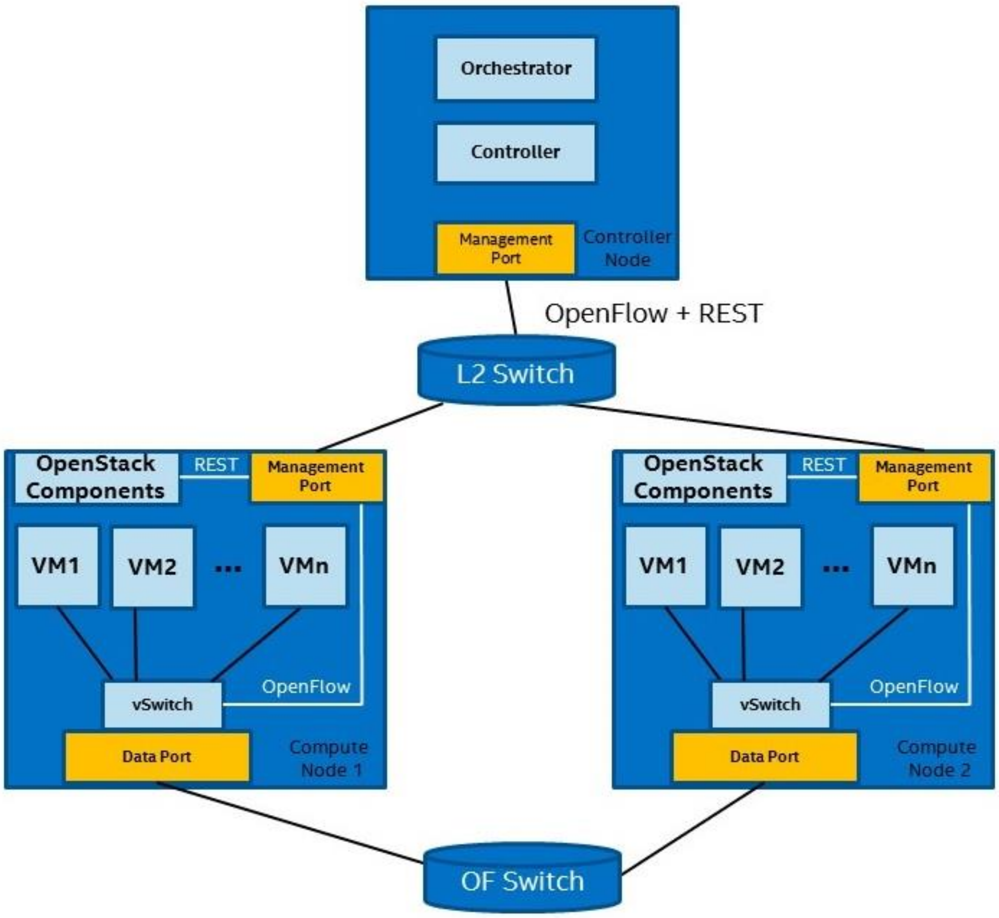

Figure 2-2 Generic setup with controller and two compute nodes

2.1 Network Services Examples

The network services presented in this section are included as examples that have been tested with Intel® ONP Reference Architecture. They are demonstrated as use cases running as virtualized instances deployed and controlled by OpenStack.

2.1.1 Suricata (Next Generation IDS/IPS Engine)

Suricata is a high-performance network Intrusion Detection System (IDS), Intrusion Prevention System (IPS), and network security monitoring engine developed by the Open Information Security Foundation, its supporting vendors, and the community. Refer to http://suricata-ids.org.

2.1.2 vBNG (Broadband Network Gateway)

A BNG, may also be known as a Broadband Remote Access Server, routes traffic to and from broadband remote access devices, such as Digital Subscriber Line Access Multiplexers. This network function is included as an example of a workload that can be virtualized on the Intel® ONP.

Intel® PROX (Packet pROcessing eXecution engine) tool is used to demonstrate workload functionality of a virtual BNG as a VNF. It is based on DPDK libraries to accelerate data plane packet processing. The PROX tool is configured to perform Layer 3 forwarding in a VNF acting as an intermediate node across two VNFs. With the options provided by the PROX tool the input and output characteristics of the traffic and packet processing across multiple cores with in the VNF can be configured among its different options.

Intel® PROX using DPDK can be found at https://01.org/intel-data-plane-performance-demonstrators/downloads/prox-application-v021.

Additional information on the performance characterization of this vBNG implementation can be found at https://networkbuilders.intel.com/docs/Network_Builders_RA_vBRAS_Final.pdf.

3.0 Hardware Components

Table 3-1 and Table 3-2 provide details of platform hardware components used for testing purposes. The Notes column describes some of the fine tunings enabled on the hardware.

Table 3-1 Intel® Xeon® processor E5-2600 v3 and Intel® Xeon® processor E5-2600 v4 product family-based platforms - hardware used in integration tests

Item	Description	Notes
Platform	Intel® Server Board S2600WTT	Formerly Wildcat Pass Intel® Xeon® processor-based DP server 2 x 10GbE integrated LAN ports based on Intel® Ethernet Controller X-540
	Intel® Server Board S2600WT2	Formerly Wildcat Pass Intel® Xeon® processor-based DP server 2 x 1 GbE integrated LAN ports based on Intel® Ethernet Controller I350-AM2
Processors	Dual Intel® Xeon® processor E5-2658 v3	Formerly Haswell 12 cores, 24 threads, 2.2 GHz, 105 W, 30 MB Intel® Smart cache per processor, 9.6 GT/s QPI, DDR4-1600/1866/2133, 24 hyper-threaded cores per CPU for 48 total cores. Supports CAT/CMT
	Dual Intel® Xeon® processor E5-2697 v3	Formerly Haswell 14 cores, 28 threads, 2.6 GHz, 145 W, 35 MB total cache per processor, 9.6 GT/s QPI, DDR4-1600/1866/2133, 28 hyper-threaded cores per CPU for 56 total cores.
	Dual Intel® Xeon® processor E5-2699 v3	Formerly Haswell 18 cores, 36 threads, 2.3 GHz, 145 W, 45 MB total cache per processor, 9.6 GT/s QPI, DDR4-1600/1866/2133, 36 hyper-threaded cores per CPU for 72 total cores.
	Dual Intel® Xeon® processor E5-2699 v4	Formerly Broadwell 22 cores, 44 threads, 2.2 GHz, 145 W, 55 MB total cache per processor, 9.6 GT/s QPI, DDR4-1600/1866/2133/2400, 44 hyper-threaded cores per CPU for 88 total cores. Supports CAT/CMT
Memory	64 GB total; Crucial CT8G4RFS4213	8x DDR4 RDIMM 2133 MHz, 8 GB
	64 GB total; Kingston KVR21R15S4/8	8x DDR4 RDIMM 2133 MHz, 8 GB
Intel® QuickAssist Technology	Intel® QuickAssist Adapter 8950	Formerly Walnut Hill Provides IPSec, SSL Acceleration and Compression services Support for SR-IOV PCIe Gen 3 (8GT/s)

Item	Description	Notes
NICs	Intel® Ethernet Converged Network Adapter X710-DA4	Formerly Fortville Intel® Ethernet Controller XL710-AM1 4 x 10 GbE ports Firmware version 4.53 Tested with Intel® FTLX8571D3BCV-IT and AFBR-703sDZ-IN2 transceivers
	Intel® Ethernet Converged Network Adapter XL710-QDA2	Formerly Fortville Intel® Ethernet Controller XL710-AM2 2 x 40 GbE ports Firmware version 4.53 Tested with Intel® E40QSFPSR transceiver
	Intel® Ethernet Converged Network Adapter X540-T2	Formerly Twinville Intel® Ethernet Controller X540-BT2 2 x 10 GbaseT ports
	Intel® Ethernet Converged Network Adapter X520-SR2	Formerly Niantic Intel® 82599ES 10 Gigabit Ethernet Controller 2 x 10 GbE ports Tested with Intel® FTLX8571D3BCV-IT transceiver
Local Storage	Intel® SSD DC S3500 Series	Formerly Wolfsville SSDSC2BB120G4 120 GB SSD 2.5in SATA 6GB/s
BIOS	Servers with Intel® Xeon® processor E5-2600 v3 product family: • SE5C610.86B.01.01.0009.060120151350 Release date: 06/01/2015 • SE5C610.86B.01.01.0011.081020151200 Release date: 08/10/2015 Servers with Intel® Xeon® processor E5-2600 v4 product family: • GRRFCRB1.86B.0267.R00.1509110656 RC revision 2.4.0 Release date: 09/11/2015	Hyper-Threading enabled Intel® Virtualization Technology (Intel® VT-x) enabled Intel® VT for Directed I/O (Intel® VT-d) enabled Turbo Boost enabled

Table 3-2 Intel® Xeon® processor D-1500 family-based SoC platforms - hardware used in integration tests

Item	Description	Notes
Platform	SuperMicro SuperServer 5018D-FN4T	Intel® Xeon® processor-based SOC server Motherboard: SuperMicro X10SDV-8C-TLN4F Dual LAN via onboard Intel® i350-AM2 Gigabit Ethernet Dual LAN via SoC 10GBase-T
Processors	Intel® Xeon® processor D-1540	Formerly Broadwell-DE 8 cores, 16 threads, 2.00 GHz, 12 MB cache Single Socket FCBGA 1667 supported CPU TDP 45W System-on-Chip
	Intel® Xeon® processor D-1520	Formerly Broadwell-DE 4 cores, 8 threads, 2.20 GHz, 6 MB cache Single Socket FCBGA 1667 CPU TDP 45W System-on-Chip
Memory	32 GB total; Kingston KVR21R15S4/8	4x DDR4 RDIMM 2133 MHz, 8 GB
Local Storage	Seagate Barracuda ST500DM002	500 GB HDD 3.5in SATA 6GB/s 7200RPM 16MB
BIOS	AMIBIOS Version: 1.0a Release Date: 05/27/2015	Hyper-Threading enabled Intel® Virtualization Technology (Intel® VT-x) enabled Intel® VT for Directed I/O (Intel® VT-d) enabled
Intel® QuickAssist Technology	Intel® QuickAssist Adapter 8950	Formerly Walnut Hill Provides IPSec, SSL Acceleration and Compression services Support for SR-IOV PCIe Gen 3 (8GT/s)

4.0 Software Versions

Table 4-1 describes functions of the software ingredients along with their version or configuration. For open-source components, a specific commit ID set is used for this integration. Note that the commit IDs used are the latest working set at the time of this release.

Table 4-1 Software versions

Software Component	Function	Version/Configuration
Fedora 23	Host Operating System	Fedora 23 Server x86_64 Kernel version: 4.3.3-300.fc23.x86_64
CentOS-7.2	Host Operating System	CentOS-7.2 (1511) x86_64 DVD ISO Kernel version: 3.10.0-327.el7.x86_64
KVM4NFV Real-Time Kernel	Targeted towards low latency Telco environment	KVM4NFV Real-Time Kernel version: 4.1.10-rt10
QEMU-KVM	Virtualization technology	Fedora 23: • qemu-kvm version: 2.4.1-7.fc23.x86_64 • libvirt version: 1.2.18.2-2.fc23.x86_64 CentOS-7.2: • qemu-kvm version: 1.5.3-105.el7_2.3.x86_64 • libvirt version: 1.2.17-13.el7.x86_64
DPDK	Network stack bypass and libraries for packet processing; includes user space vhost drivers	DPDK 2.2.0
Open vSwitch	vSwitch	Open vSwitch 2.5.90 Commit ID 1589ee5ae97c3f71c50413db64ddd0546daeecc0 used for: • Open vSwitch (non-DPDK nodes) • Open vSwitch with DPDK
OpenStack	SDN orchestrator	OpenStack Liberty Release (see Table 4-3)
OpenDaylight	SDN controller	OpenDaylight beryllium-snapshot-0.4.1
Intel® Ethernet Drivers	Ethernet drivers	ixgbe-4.3.13 • Intel® Ethernet Server Adapter X520 Series • Intel® Ethernet Converged Network Adapter X540-T2 • Intel® Xeon® processor D-1500 family deployments i40e-1.4.25 • Intel® Ethernet Converged Network Adapters XL710-QDA2 and X710-DA4
Cache Allocation Technology (CAT) / Cache Monitoring Technology (CMT)	Intel® Resource Director Technology (RDT) components	Intel-cmt-cat Commit ID db381ce554528d96f394863aab6e985f5171cff9

4.1 Obtaining Software Ingredients

All of the open-source software ingredients involved are downloaded from the source repositories shown in the Table 4-2. Commit IDs for major OpenStack components are shown in the Table 4-3.

Table 4-2 Sources for software ingredients

Component	Location	Comments
Fedora 23	http://mirror.us.leaseweb.net/fedora/linux/releases/23/Server/x86_64/iso/	https://getfedora.org/en/server/ Fedora-Server-DVD-x86_64-23.iso
CentOS-7.2	http://isoredirect.centos.org/centos/7/isos/x86_64/CentOS-7-x86_64-DVD-1511.iso	https://www.CentOS.org/download/ CentOS-7-x86_64-DVD-1511.iso
KVM4NFV Real-Time Kernel	git clone https://gerrit.opnfv.org/gerrit/kvmfornfv git checkout WW-2015-52	v4.1.10-rt10
DPDK	git clone http://dpdk.org/git/dpdk checkout a38e5ec15e3fe615b94f3cc5edca5974dab325ab, v2.2 from master	Includes DPDK PMD, sample apps (bundled) DPDK download will be done through the DevStack script during installation.
Open vSwitch	git clone https://github.com/openvswitch.ovs.git checkout 1589ee5ae97c3f71c50413db64ddd0546daeecc0	OvS download will be done through the DevStack script during installation.
OpenStack	OpenStack Liberty Release	Deployed using DevStack (see next row). The commit IDs for the various OpenStack components are provided in the Table 4-3.
DevStack	git clone https://github.com/openstack-dev/devstack.git checkout stable/liberty 501bb07462ef4fbe81143f0a58364ada0da48fe2	
OpenDaylight	Beryllium checkout 9adb4907d8c542a31c4dd4a4f2219adb2b95ed7a	OpenDaylight download will be done through the DevStack script during installation.
CAT/CMT Technologies	git clone http://github.com/01org/intel-cmt-cat checkout db381ce554528d96f394863aab6e985f5171cff9	Intel® ONP Release 2.1 Platform QoS Application Note available on 01.org
Intel® ONP 2.1 Scripts	https://01.org/packet-processing/intel®-onp-servers	Includes helper scripts to set up Intel® ONP 2.1 using DevStack
Suricata	`# yum install suricata`	Suricata version 2.0.2-1.fc20.x86_64 was used in integration tests
vBNG using PROX	https://01.org/intel-data-plane-performance-demonstrators/downloads/prox-application-v021	
Intel® Ethernet Drivers	https://sourceforge.net/projects/e1000/files/ixgbe%20stable/4.3.13/ixgbe-4.3.13.tar.gz https://sourceforge.net/projects/e1000/files/i40e%20stable/1.4.25/i40e-1.4.25.tar.gz	ixgbe-4.3.13 • Intel® Ethernet Server Adapter X520 Series • Intel® Ethernet Converged Network Adapter X540-T2 • Intel® Xeon® processor D-1500 family i40e-1.4.25 • Intel® Ethernet Converged Network Adapters XL710-QDA2 and X710-DA4

Table 4-3 Commit IDs for major OpenStack components

OpenStack Component	Referenced Version (i.e. tag at release)	Commit ID
OpenStack Cinder	7.0.1	stable/liberty f51ffea673de5395aee6c789b07fb44d9e801b88
OpenStack Glance	11.0.1	stable/liberty 7296a5302b00bd066ddf6b14c7d5a9afb3b88e70
OpenStack Heat	5.0.1	stable/liberty dc8ccd8ee15bb336c4704ac1ff628bfb245a5593
OpenStack Horizon	8.0.1	stable/liberty fa47798f38b2a58514b93b6613129b0dfca18f36
OpenStack Ironic	4.2.2	stable/liberty 6eb970b71cb6ae629b733ced84917d9db5afc78a
OpenStack Keystone	8.1.0	stable/liberty c665080d4a700b6d92f29c40621d83bd7365de34
OpenStack Neutron	7.0.3	stable/liberty 197b188ea8bfdfd023b2da3b7572e9387568c500
OpenStack Nova	12.0.2	stable/liberty aa4edd349dde73739527ab793ff6209fe1907e2c
OpenStack Swift	2.5.0	stable/liberty 47eb6a37f86f29c355297b556c2ff898c98da9b2
OpenStack Requirements		stable/liberty f8579e166f45f6a580ef56cd3c0e734c03ae2f76
OpenStack Tempest		stable/liberty 271b3405729778a5bdb71004b8fa27484524295c
OpenStack noVNC		stable/liberty b403cb92fb8de82d04f305b4f14fa978003890d7
OpenStack networking-odl		stable/liberty 9adb4907d8c542a31c4dd4a4f2219adb2b95ed7a
OpenStack networking-ovs-dpdk		stable/liberty 3b800fea6255d2209565d0330c0ab73356f729d7

Note: See Intel® ONP Release 2.1 Scripts for commit IDs of minor components.

Note: Due to the number of ingredients involved, follow the instructions provided by the scripts and execute them in order to deploy and provision each of the software components.

Note: For the sake of simplicity, this document uses `yum` command for installing packages. As of Fedora 22, `yum` has been replaced by `dnf`.

5.0 Installation and Configuration Guide

This section describes the installation and configuration instructions to prepare the Controller and Compute Nodes.

5.1 BIOS Settings

For Intel® Xeon® processor based platforms, enter the BIOS menu and update the following configuration as described in the Table 5-1. These settings are common to both the Controller and Compute Nodes.

Table 5-1 BIOS Settings

Configuration	Controller Node Setting	Compute Node Setting
Intel® Virtualization Technology	Enabled	Enabled
Intel® Hyper-Threading Technology (HTT)	Enabled	Enabled
Intel® VT for Directed I/O (VT-d)	Enabled	Enabled

5.2 Operating System Installation and Configuration

Current Intel® ONP scripts support two operating systems – Fedora 23 and CentOS-7.2. Details below provide generic instructions for installing and configuring the operating system of choice. Other methods of installing the operating system, such as network installation, PXE boot installation, USB key installation, etc., are not described in this guide.

1. Download the image of Linux distribution supported by Intel® ONP from the source given in the Table 4-2.
2. Burn the ISO file to a DVD and create an installation disk.
3. Use the DVD to install the operating system. During the installation, click **Software selection**, then choose the following:
    - Development Tools
    - Virtualization
4. Create a user named **stack** and check the box **Make this user administrator** during the installation. The user **stack** is also used in the OpenStack installation. Reboot the system after completing the installation.

5.3 Automated Installation Using Scripts

The Intel® ONP chooses DevStack for quick deployment of OpenStack in order to use its latest features, including, sometimes, its experimental ones. DevStack provides and maintains tools for the installation of OpenStack from upstream sources. Its main purpose is to support OpenStack development and testing of the components involved.

Due to its evolving nature, DevStack does not provide production-level stability. Since Devstack depends on the upstream sources in the internet to install the packages, there may be instability in the downloading process (for example SSL certificates updates, servers maintenance works etc.). If you see a failure in installation of DevStack due to difficulty in downloading the packages from internet, please wait for some time and try again.

In order to have predictable and validated setup using the scripts provided for the installation, it is optimal to proceed with the base set of Intel-recommended settings. Many of the manual operating system and OpenStack installations with DevStack are automated. The bulk of this procedure is condensed into a few steps that can be executed using these scripts. All the ingredients listed in Table 4-1 can be installed using this method and have been validated with hardware described in the section 3.0 Hardware Components.

Before continuing with the scripts, update the BIOS with the settings presented in Table 5-1. Installation scripts can be obtained from the download link in the section 4.1 Obtaining Software Ingredients.

Note: The automation scripts are only tested on Fedora 23 and CentOS-7.2 operating systems and will not necessarily work on other distributions of Linux.

5.3.1 Intel® ONP Scripts

Intel® ONP scripts provide a simplified procedure for installing and configuring OpenStack using DevStack. Based on the type of deployment preferred, the user will need to update a configuration file. The scripts will take care of preparing the system and deploying the software ingredients, including:

- Configuring, provisioning and updating the required operating system kernel, services and network settings.
- Installing OpenStack using DevStack installer.
- Configuring the server as either a Controller or Compute node, according to the configuration file.

The installation script tarball contains the following files to accelerate DevStack deployment:

README.md

This file provides instructions on how to update the `onps_config.yml` configuration file and run ansible. It is highly recommended to read this file before attempting installation.

onps_config.yml

A configuration file which determines the behavior of the installation scripts. The Intel® ONP deployment is based on the details provided in this file by user. Following are the current options provided to the user:

- the choice to provision the Controller and Compute Node with vanilla Open vSwitch or accelerated Open vSwitch using the DPDK.
- the choice of the overlay network type to configure: VXLAN or VLAN
- the choice to install the OpenDaylight SDN controller
- the hostname of the node
- the network interfaces to use for each of the OpenStack networks
- any proxy information
- the kernel to use.

Note: The selections in this file will determine the content of the `local.conf` configuration file required for DevStack installation.

prepare_system.yml

This is an ansible playbook that prepares the nodes to run DevStack. This ansible playbook does the following tasks:

- Configures network interfaces and services.
- Configures (and optionally updates) the kernel.

- Enables or disables system services.
- Pulls required updates for the operating systems.
- Creates the DevStack configuration file `local.conf`.
- Compiles and installs Linux Base Drivers for Intel® Ethernet Network Connections.

prepare_stack.yml

Once the system configuration tasks in `prepare_system.yml` are completed, ansible will run DevStack using the tasks in `prepare_stack_serial.yml`. This ansible playbook does the following:

- Executes DevStack installer script, `stack.sh`.
- Configures OvS and DPDK parameters based on the overlay network.
- Finalizes network settings.
- Capture log files from all the nodes.

onps_commit_ids.yml

Intel® ONP uses software components from multiple open-source repositories. In order to have a predictable installation and working setup specific commit IDs are used to ensure validated components are installed. This file contains list of frozen commit ids including, but not limited to OpenStack, OvS and DPDK components.

tests/*

DevStack uses a configuration file, `local.conf`, to setup and configure OpenStack services. The `local.conf` file is generated based on the user choices provided in the `onps_config.yml` file. As a point of reference, a set of sample `local.conf` configuration files are provided in the `tests`/ directory with various combinations of deployment. These files can be used as example configuration files for deployments of choice.

5.4 Controller and Compute Node Setup

The following procedure uses actual examples of an OpenStack (DevStack) installation performed in an Intel® test lab. It consists of one Controller Node (controller) and one Compute Node (compute). It is assumed that the user has successfully followed the hardware and software installation and configuration detailed in sections above.

Note: This procedure must be repeated on each node in the environment.

5.4.1 Network Configuration and Requirements

At least two networks are required to build the OpenStack infrastructure in a lab environment. One network is used to connect all nodes for OpenStack management (management network); the other is a private network exclusively for an OpenStack internal connection (private or tenant network) between instances (or VMs).

Some users might want to have Internet and/or external connectivity for OpenStack instances (VMs). In this case, an optional network (public network) can be used.

One additional network is required for Internet connectivity, because installing OpenStack requires pulling packages from various sources/repositories on the Internet.

The assumption is that the targeting OpenStack infrastructure contains multiple nodes: one is a Controller Node and one or more are Compute Nodes.

5.4.2 Network Configuration Example

The following is an example on how to configure networks for the OpenStack infrastructure. The example uses four network interfaces. Note that the names of these network interfaces are only examples.

- **enp3s0f1:** For the Internet network — used to pull all necessary packages/patches from repositories on the Internet, configured to obtain a Dynamic Host Configuration Protocol (DHCP) address.
- **enp3s0f0:** For the management network — used to connect all nodes for OpenStack management, configured to use a private static address.
- **ens786f0:** For the tenant network — used for OpenStack internal connections for VMs. Configuration of the tenant network interface becomes more complicated with the introduction of the OpenDaylight SDN controller. Depending on whether OpenDaylight is used for network control, the configuration of this interface varies:
  - If OpenDaylight or VXLAN is used, configure this interface with an IP address. This address should be from a different network than the management network.
  - If OpenDaylight and VLXAN is not used, configure this interface with no IP address.
- **ens786f1:** For the optional external network — Used for VM Internet/external connectivity, configured with no IP address. This interface is only used in the controller node, if the external network is configured. For the compute node, this interface is not required.

Note: Among these interfaces, the interface for the tenant network (in this example, ens786f0) is used for DPDK and OvS with DPDK.

Note: Static IP address should be used for the interface of the management network.

Note: If VLAN is used, prior to configuring the network on the nodes, please make sure to configure specific VLAN tags on the switch ports assigned to corresponding nodes.

5.4.3 Controller and Compute Node Installation Procedure

Follow the steps below to configure the host to be an OpenStack Controller or Compute Node:

1. Plan ahead to decide what interfaces of your hosts will belong to management and/or the data plane network.
2. Manually edit the `onps_config.yml` configuration file on the deployment server for Controller and Compute Node, hostname, type of interfaces, type of vSwitch desired vanilla vs. DPDK based Open vSwitch, etc.
3. Execute `ansible-playbook –i inventory.ini multinode.yml on the deployment server.` The script will connect to the specified remotes systems to prepare the services and file system, install kernel and network drivers, reboot, install and deploy OpenStack controller and compute services.

Note: Intel® ONP scripts assume SOCKS proxy is available. Otherwise, git will not be able to clone repositories using git:// protocol. To change the git settings, use https:// protocol instead:

```
# git config --system url."https://".insteadOf git://
```

Note: DHCP is assumed on the internet interface. If DHCP is not present, add DNS1 & DNS2 to `/etc/hosts`. A reboot will be needed to fix `/etc/resolv.conf`.

Note: If the `prepare_stack.yml` playbook execution returns an error and/or fails during installation, follow the instructions below to clean the OpenStack node:

```
$ cd /home/<username>/devstack
$ ./unstack.sh
$ ./clean.sh
```

```
$ sudo rm -rf /opt/stack/
$ rm -rf /home/<username>/devstack
$ sudo reboot
```

Note: Once OpenStack installation is done, user is advised to check if all OpenStack services are up and running.

5.4.4 Real-Time Kernel Compute Node Enablement

Some use cases, such as Telco media applications, which are sensitive to low latency and jitter require a real time kernel.

During configuration of the `onps_config.yml` file, the user is presented with a multiple kernel options. To deploy a Compute Node using a real time kernel, specify:

- kernel_to_use=realtime
- version=<according to Table 4-1>
- kernelURL=<according to Table 4-2>

To enable the Real-Time Kernel you may need to ensure the `bc` is installed, which is the component required to generate the keys to sign the kernel modules during compilation. To install the `bc` run the following command.

```
# yum install bc
```

Appendix A: Special Considerations for DPDK enablement

A.1 OvS-DPDK Compute Node with a vhost-user

With the OvS with DPDK compute node with a vhost-user, a large memory page size should be configured for the OpenStack flavor for creating instances. This can be done in two steps: first create a flavor, and then modify it to allow a large memory page size.

The following commands create a flavor named `largepage-flavor` with an ID of 1002, 1024 Mb virtual memory, 4 Gb virtual disk space, 1 virtual CPU, and large memory page size:

```
$ nova flavor-create largepage-flavor 1002 1024 4 1
$ nova flavor-key 1002 set "hw:mem_page_size=large"
```

Use this flavor to create instances hosted by OvS with DPDK compute node with a vhost-user.

A.2 Tenant Network with DPDK and/or VXLAN Tunneling

With the introduction of DPDK and/or VXLAN tunneling to tenant network, the user must consider the maximum transmission unit (MTU) associated with the network interface in order to have the tenant network function as expected. One method to do this is to decrease the MTU of all virtual network devices (i.e., network interfaces in all VMs) by 50 bytes to 1450 bytes. This is because the VXLAN adds an extra 50-byte header to the MTU of 1500 bytes for normal Ethernet packets. The example below changes the MTU eth0 to 1450 bytes:

```
# ip link set eth0 mtu 1450
```

Appendix B: Acronyms and Abbreviations

Abbreviation	Description
BIOS	Basic Input/Output System
BNG	Broadband (or Border) Network Gateway
CAT	Cache Allocation Technology
CMT	Cache Monitoring Technology
DHCP	Dynamic Host Configuration Protocol
DPDK	Data Plane Development Kit
HTT	Hyper-Threading Technology
IA	Intel® Architecture
IDS	Intrusion Detection System
IPS	Intrusion Prevention System
KVM	Kernel-based Virtual Machine
ML2	Mechanism Layer 2
NFV	Network Functions Virtualization
NIC	Network Interface Card
ONP	Open Network Platform
OvS	Open vSwitch
PROX	Packet pROcessing eXecution engine
RDT	Intel® Resource Director Technology
SDN	Software Defined Networking
SR-IOV	Single Root I/O Virtualization
vBNG	Virtual Broadband (or Border) Network Gateway
VM	Virtual Machine
VNF	Virtualized Network Function

Appendix C: References

Title	Source
01.org: Intel® Open Network Platform	https://01.org/packet-processing/intel%C2%AE-onp-servers
01.org: Intel® ONP 2.1 Release Notes	https://01.org/packet-processing/intel%C2%AE-onp-servers
DevStack	http://docs.openstack.org/developer/DevStack/
DPDK	http://www.intel.com/go/dpdk
Intel® Ethernet Converged Network Adapter X540-T2	http://ark.intel.com/products/58954/Intel-Ethernet-Converged-Network-Adapter-X540-T2
Intel® Converged Network Adapter X520-SR2	http://ark.intel.com/products/39774/Intel-Ethernet-Converged-Network-Adapter-X520-SR2
Intel® Ethernet Converged Network Adapter XL710-QDA2 2 x 40 GbE	http://ark.intel.com/products/83967/Intel-Ethernet-Converged-Network-Adapter-XL710-QDA2
Intel® Ethernet Converged Network Adapter X710-DA4 4 x 10 GbE	http://ark.intel.com/products/83965/Intel-Ethernet-Converged-Network-Adapter-X710-DA4
Intel® Server Board S2600WT2	http://ark.intel.com/products/82155/Intel-Server-Board-S2600WT2
Intel® Server Board S2600WTT	http://ark.intel.com/products/82156/Intel-Server-Board-S2600WTT
Intel® server product S2600WT family Product Brief	http://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/server-system-s2600wt-brief.pdf
Intel® Xeon® processor D-1500 family	http://ark.intel.com/products/series/87040/Intel-Xeon-Processor-D-1500-Family?q=Intel%20Xeon%20D-1500#@All
Intel® Xeon® processor D-1520	http://ark.intel.com/products/87038/Intel-Xeon-Processor-D-1520-6M-Cache-2_20-GHz
Intel® Xeon® processor D-1540	http://ark.intel.com/products/87039/Intel-Xeon-Processor-D-1540-12M-Cache-2_00-GHz
Intel® Xeon® processor E5-2600 v3 product family	http://ark.intel.com/products/series/81065/Intel-Xeon-Processor-E5-2600-v3-Product-Family#@Server
Intel® Xeon® processor E5-2699 v3	http://ark.intel.com/products/81061/Intel-Xeon-Processor-E5-2699-v3-45M-Cache-2_30-GHz
Intel® Xeon® processor E5-2697 v3	http://ark.intel.com/products/81059/Intel-Xeon-Processor-E5-2697-v3-35M-Cache-2_60-GHz
networking-ovs-dpdk ML2 Plugin	https://git.openstack.org/cgit/openstack/networking-ovs-dpdk https://github.com/openstack/networking-ovs-dpdk/blob/master/doc/source/usage.rst
OpenDaylight Beryllium	https://www.opendaylight.org/odlbe
OpenStack	https://www.openstack.org/
SuperMicro SuperServer 5018D-FN4T	http://www.supermicro.com/products/system/1u/5018/SYS-5018D-FN4T.cfm
SuperMicro X10SDV-8C-TLN4F Motherboard	http://www.supermicro.com/products/motherboard/Xeon/D/X10SDV-8C-TLN4F.cfm
Suricata	http://suricata-ids.org/

Legal Information

By using this document, in addition to any agreements you have with Intel, you accept the terms set forth below.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

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.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Intel processors of the same SKU may vary in frequency or power as a result of natural variability in the production process.

Intel technologies may require enabled hardware, specific software, or services activation. Check with your system manufacturer or retailer. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit http://www.intel.com/performance.

All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice. Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

No computer system can be absolutely secure. Intel does not assume any liability for lost or stolen data or systems or any damages resulting from such losses.

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

Intel Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.