



**EPA Hands-On LAB #3 Dec 2018**

# **EPA Hands-On Lab #3**

Version 1.0

Date 30th Nov 2018



<b>1 Objective</b>	<b>3</b>
<b>2 Prerequisites</b>	<b>3</b>
<b>3 Network Diagram</b>	<b>4</b>
<b>4 DPDK Instance(s) Basic Details</b>	<b>5</b>
<b>5 DPDK Capable Instance creation from OpenStack Dashboard (Pre-configured to skip)</b>	<b>5</b>
5.1 Instance Creation with DPDK Network	5
5.2 Verification Post DPDK Instance Creation	8
<b>6 DPDK Capable Instance Login</b>	<b>10</b>
<b>7 Assigning IP Address to the DPDK Ports in the Ubuntu Instance</b>	<b>10</b>
<b>8 Latency Test</b>	<b>11</b>
8.1 East-West	11
4. Perform the steps #1 thru #3 between an instance with shared CPU and other instance.	13
8.2 Latency Comparison Graphs	13
8.2.1 Between instances with dedicated CPUs	13
<b>9 Running iperf</b>	<b>14</b>
9.1 East-West	14
9.2 iperf Comparison Graphs	17
9.2.1 East - West	17
<b>10 Appendix</b>	<b>18</b>
10.1 Configuring OVS with DPDK in Openstack Environment	18
10.1.1 Compute Node specific operations	18
10.1.2 Openstack Specific Settings	19
10.1.3 Openstack Dashboard Specific Changes	19
10.2 Reference	21



# 1 Objective

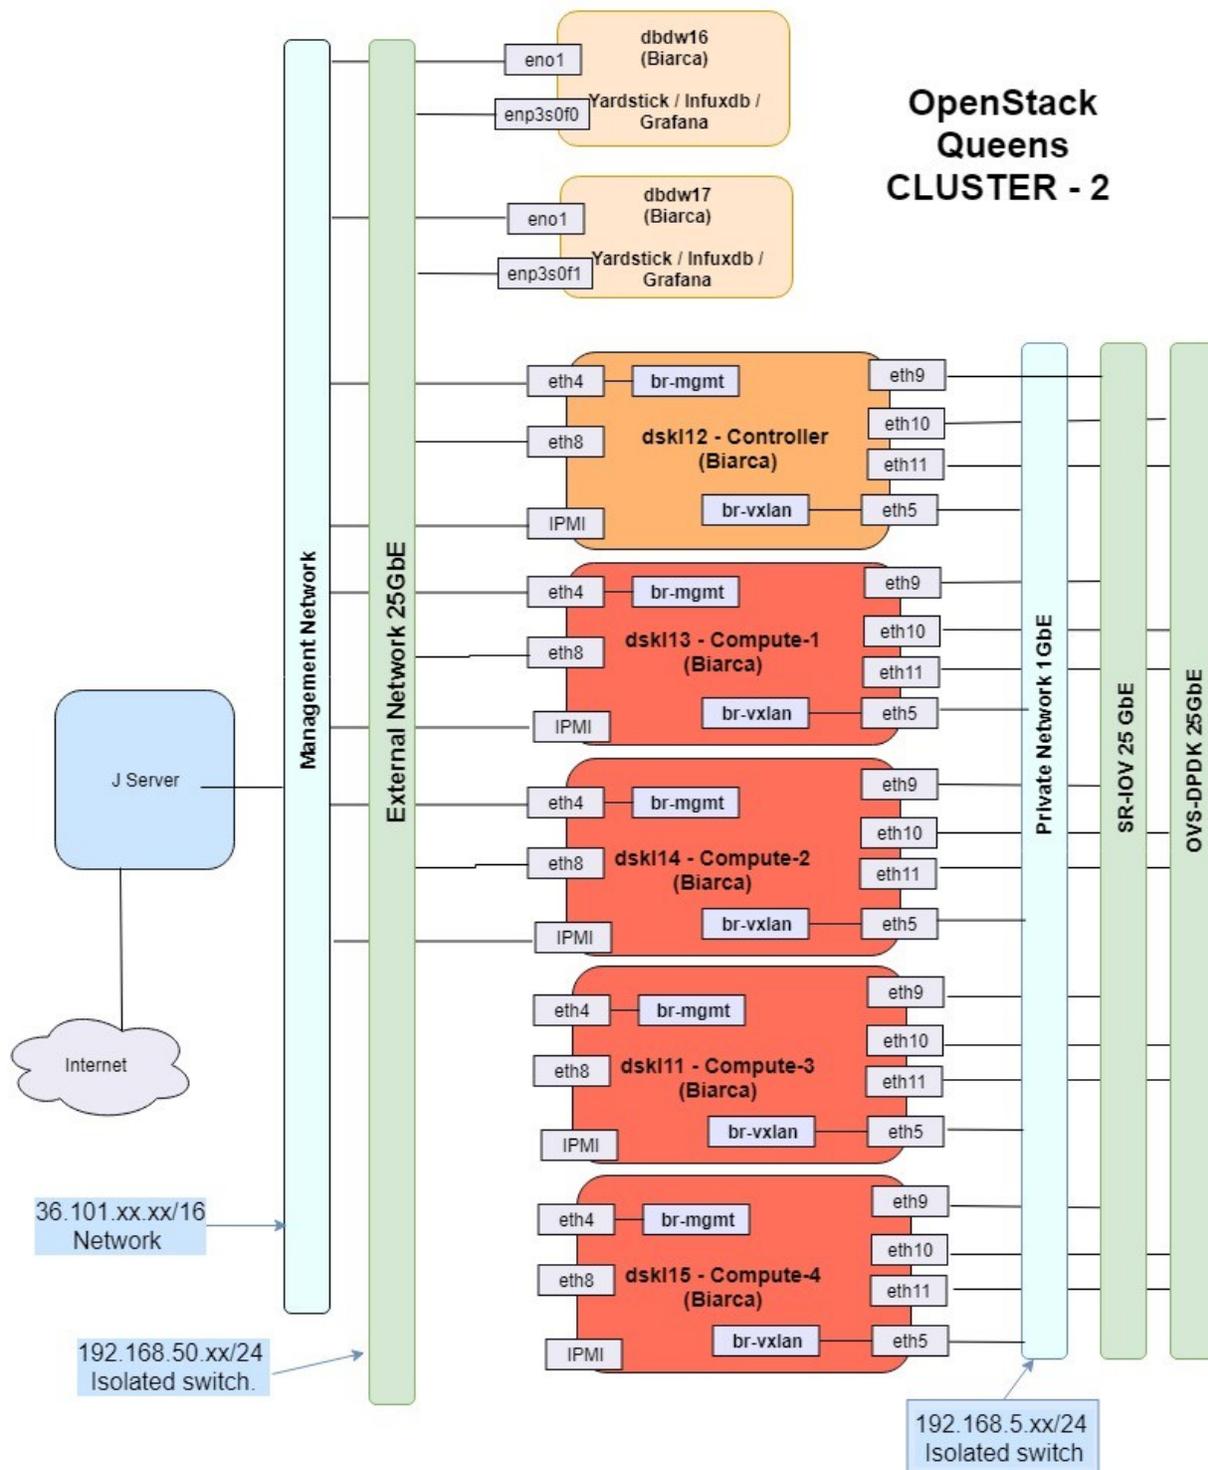
Using an OpenStack Queens based Multi-Node cluster, help get an understanding of Data Plane Development Kit (DPDK) as a simple option for network performance enhancement. Plan is to use iperf3 as a tool of choice running on Ubuntu 16.04 Instances using DPDK based network ports and validate East-West communication. Throughput will be measured from a CPUs pinned instance and an instance whose CPUs are shared.

# 2 Prerequisites

- Networking fundamentals
- Basic knowledge of DPDK and familiarity with OpenStack environment
- Hands-On EPA Lab - Introduction
- Overview of current Infrastructure
- Host to be used for connecting to DPDK capable OpenStack Instances
- Infra-access details available (**Did you receive a note/slip? - If NO! - please STOP here and ASK!!!**)
  - Credentials to login into jumpserver
  - Host Access details
  - OpenStack Dashboard Credentials
  - DPDK Capable Instances details for East-West
    - External/Floating IP
    - Guest Net IP
    - DPDK Network IP - East - West
    - SSH Key pair Host Info
    - User Name: 'ubuntu'
  - DPDK Capable Compute Nodes Availability Zone
- **Commands in BLUE**
- **Dark Orange 2** text needs to be **updated before you attempt the command**
- **Dark Magenta** - "Cambria" font - **sample/console output - DO NOT USE for COPY/PASTE**
- **Highlighted information has yellow background**



### 3 Network Diagram



## 4 DPDK Instance(s) Basic Details

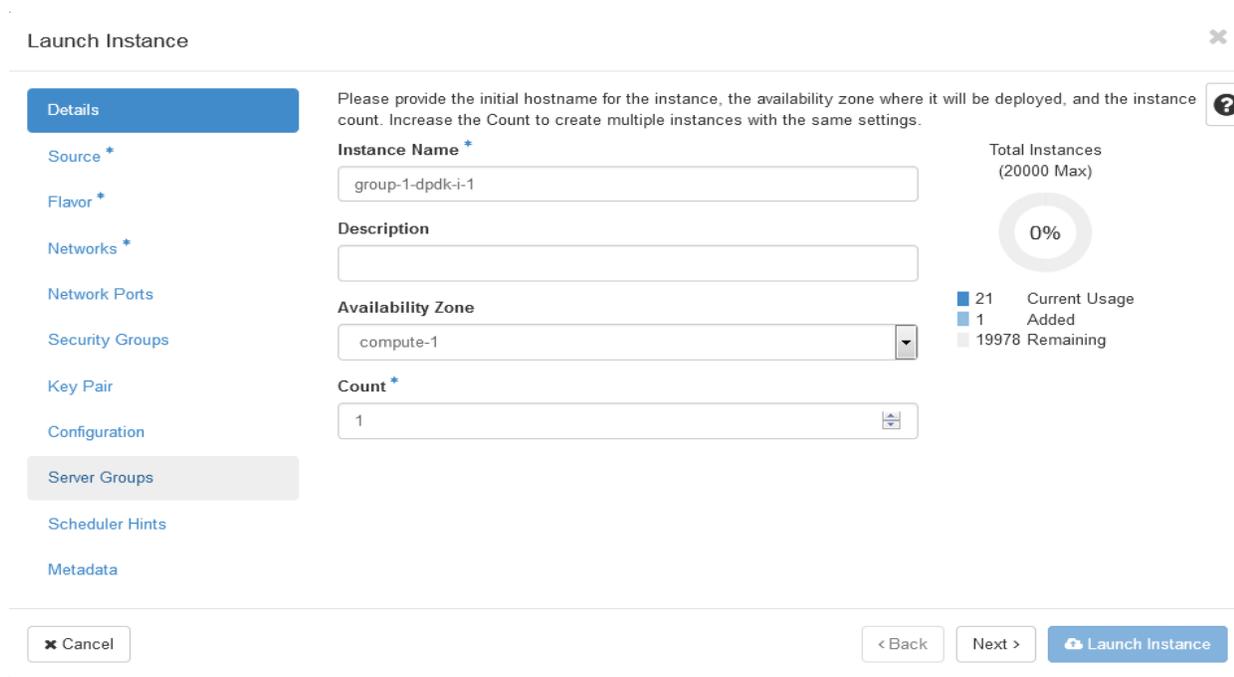
- Per Instance info
  - 2vCPUs
  - 4 GB RAM
  - 10 GB Disk
  - Ubuntu 16.04
  - Availability zones used “compute-1” and “compute-2” - Note Availability zones will be specific to user
  - Key pair as needed for SSH connection from host
- 2-3 Instances per user/group
- Additional details on DPDK are captured [DPDK Information](#)

## 5 DPDK Capable Instance creation from OpenStack Dashboard (Pre-configured to skip)

1. Login to OpenStack Dashboard using the link and access details provided to you

### 5.1 Instance Creation with DPDK Network

1. Go the Project → compute → instances. Click on new instance. Enter the name and select the availability zone.



Launch Instance

Please provide the initial hostname for the instance, the availability zone where it will be deployed, and the instance count. Increase the Count to create multiple instances with the same settings.

**Instance Name \***  
group-1-dpdk-i-1

**Description**

**Availability Zone**  
compute-1

**Count \***  
1

Total Instances (20000 Max)  
0%

21 Current Usage  
1 Added  
19978 Remaining

Cancel < Back Next > Launch Instance

2. Select **yardstick-samplevnfs** image for your instance.

### Launch Instance

Instance source is the template used to create an instance. You can use an image, a snapshot of an instance (image snapshot), a volume or a volume snapshot (if enabled). You can also choose to use persistent storage by creating a new volume.

**Source**

Select Boot Source:

**Allocated**

Name	Updated	Size	Type	Visibility
> Centos-7-netperf	11/6/18 7:03 PM	872.75 MB	qcow2	Public

**Available 4** Select one

Click here for filters.

Name	Updated	Size	Type	Visibility
> cirros-0.4.0	11/1/18 9:57 PM	12.13 MB	qcow2	Public
> Ubuntu-16.04	11/1/18 9:45 PM	283.06 MB	qcow2	Public
> yardstick-image	11/1/18 9:46 PM	607.06 MB	qcow2	Public
> yardstick-samplevifs	11/1/18 9:47 PM	3.55 GB	qcow2	Public

3. Select the **hpages** flavor



### Launch Instance

Details

Source

**Flavour**

Networks \*

Network Ports

Security Groups

Key Pair

Configuration

Server Groups

Scheduler Hints

Metadata

Flavours manage the sizing for the compute, memory and storage capacity of the instance.

**Allocated**

Name	VCPUS	RAM	Total Disk	Root Disk	Ephemeral Disk	Public
> small	2	2 GB	10 GB	10 GB	0 GB	Yes

**Available** 3 Select one

Click here for filters.

Name	VCPUS	RAM	Total Disk	Root Disk	Ephemeral Disk	Public
> tiny	1	1 GB	6 GB	6 GB	0 GB	Yes
> yardstick-flavor	1	1 GB	10 GB	10 GB	0 GB	Yes
> hpages	2	4 GB	10 GB	10 GB	0 GB	Yes

4. Select the **guest-net** as first network and **dpdk-nw-1**.

### Launch Instance

Details

Source \*

Flavor \*

**Networks \***

Network Ports

Security Groups

Key Pair

Configuration

Server Groups

Scheduler Hints

Metadata

Networks provide the communication channels for instances in the cloud.

**Allocated** Select networks from those listed below.

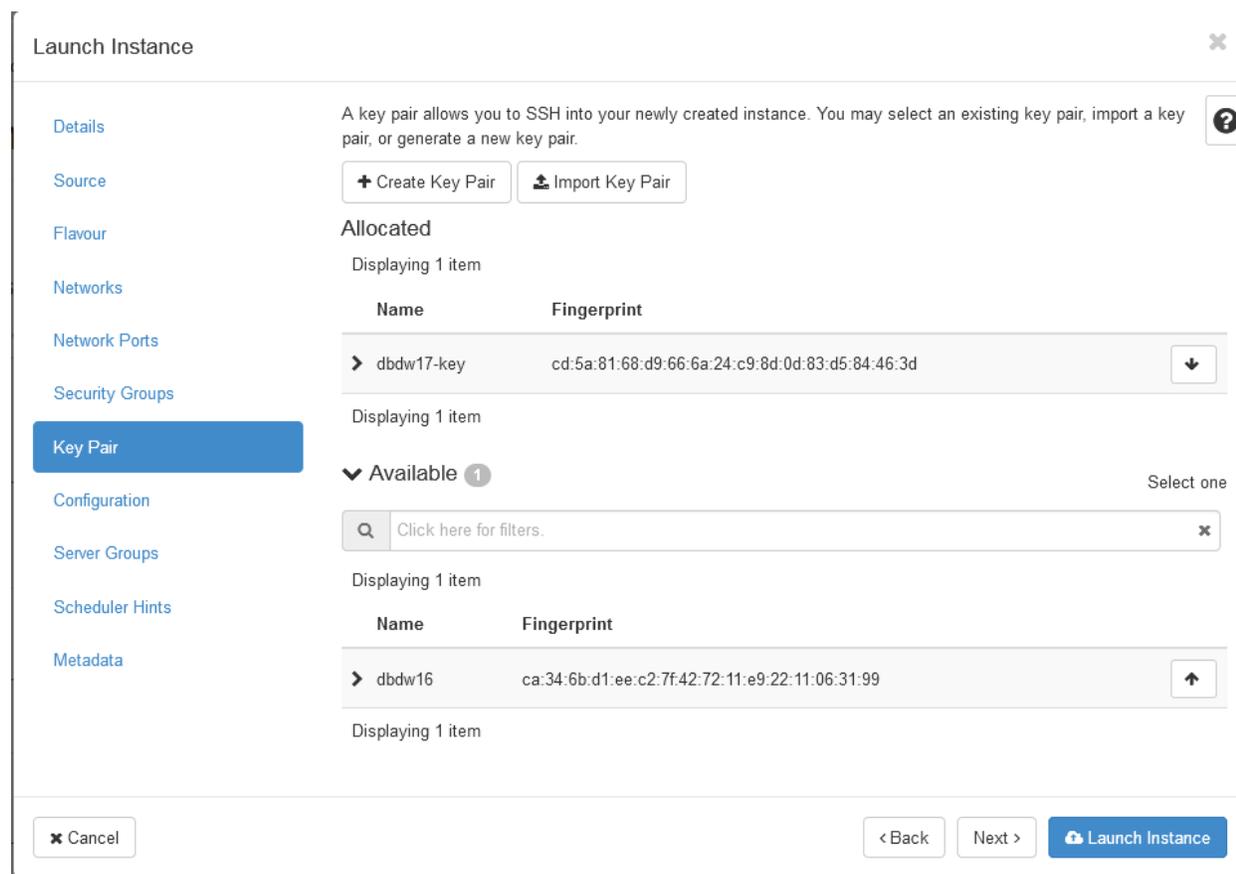
Network	Subnets Associated	Shared	Admin State	Status
Select an item from Available items below				

**Available** 5 Select at least one network

Click here for filters.

Network	Subnets Associated	Shared	Admin State	Status
> dpdk-nw-1	dpdk-subn-1	Yes	Up	Active
> sriov-net	sriov-subnet	Yes	Up	Active
> ext-net	ext-subnet	Yes	Up	Active
> dpdk-nw-2	dnp2-sn	Yes	Up	Active
> guest-net	guest-subnet	Yes	Up	Active

5. Select the key pair, **this very important to ssh login to the instance using floating ip.** Just click on launch instance once done.



Launch Instance

Details

Source

Flavour

Networks

Network Ports

Security Groups

**Key Pair**

Configuration

Server Groups

Scheduler Hints

Metadata

A key pair allows you to SSH into your newly created instance. You may select an existing key pair, import a key pair, or generate a new key pair.

+ Create Key Pair Import Key Pair

Allocated

Displaying 1 item

Name	Fingerprint
dbdw17-key	cd:5a:81:68:d9:66:6a:24:c9:8d:0d:83:d5:84:46:3d

Displaying 1 item

▼ Available 1 Select one

Click here for filters.

Displaying 1 item

Name	Fingerprint
dbdw16	ca:34:6b:d1:ee:c2:7f:42:72:11:e9:22:11:06:31:99

Displaying 1 item

Cancel < Back Next > Launch Instance

## 5.2 Verification Post DPDK Instance Creation

Follow the below steps to check if a DPDK instance is properly created.

1. ssh to the Openstack compute host on which DPDK instance is created.
2. Use **virsh list** to list all the instances on that host. Below is a sample.

```
root@dskl13:/home/ubuntu
# virsh list
 Id Name
 State
-----
-----
 1 instance-000001ee
 running
```



```
2 instance-000001ef
running
3 instance-000001f3
running
4 instance-000001f4
running
```

3. Instance details can be identified using the command **virsh dumpxml <instance id>**
4. To the newly created DPDK instance, the below should be available.

```
<memoryBacking>
  <hugepages>
    <page size='1048576' unit='KiB' nodeset='0'/>
  </hugepages>
</memoryBacking>

<vcpu placement='static'>2</vcpu>
<cputune>
  <shares>2048</shares>
  <vcpupin vcpu='0' cpuset='83'/>
  <vcpupin vcpu='1' cpuset='39'/>
  <emulatorpin cpuset='39,83'/>
</cputune>
<numatune>
  <memory mode='strict' nodeset='1'/>
  <memnode cellid='0' mode='strict' nodeset='1'/>
</numatune>

<interface type='vhostuser'>
  <mac address='fa:16:3e:65:e9:00'/>
  <source type='unix' path='/var/run/openvswitch/vhu34b907e3-cd'
mode='server'/>
  <target dev='vhu34b907e3-cd'/>
  <model type='virtio'/>
  <alias name='net1'/>
  <address type='pci' domain='0x0000' bus='0x00' slot='0x04'
function='0x0'/>
</interface>
```



- hugepages - indicate that large page size memory is assigned to the instance
- vcpupin - If dedicated CPUs are pinned to the instance, then as highlighted in the sample output above, single CPU is listed. If shared CPUs are configured, then it should look like `<vcpupin vcpu='0' cpuset='26-43,66-67,70-87'/>`
- numatune - Indicates on which NUMA node the instance is created on
- interface type “**vhostuser**” - Indicates that DPDK vhost user port is created in OVS and attached to the instance

## 6 DPDK Capable Instance Login

1. In Openstack Dashboard -> Project -> Compute -> Instances, find the floating ip of the instance along with the key-pair of the access host from which the instance can be accessed.
2. Have 3 terminal sessions and use each of the sessions to connect to the jump server and then to the access host in the previous step
3. From one session, “ssh” using the floating ip / ext-net IP to the instance from the access host. **Please make to execute “sudo su” before connecting to instance.**

```
root@<host-with-keypair>:# ssh ubuntu@<Instance external / floating IP>
```

```
root@<dbdw16>:# ssh ubuntu@192.168.50.94
```

**Note: If you see WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED! during login to the instance, run below command with the instance ip address.**

```
root@<hostname>:# ssh-keygen -f "/root/.ssh/known_hosts" -R <Instance external IP>
```

3. Need to repeat the same steps for 2 other instances for 2 other terminal sessions



# 7 Assigning IP Address to the DPDK Ports in the Ubuntu Instance

1. Update the interfaces file with the below lines. Make sure to provide your dpdk network IP at the **address** field.

```
# vi /etc/network/interfaces
auto ens4
iface ens4 inet static
address <DPDK Network IP Address provided>
netmask 255.255.255.0
```

2. Restart the network service or instance.

```
#systemctl restart networking
```

3. Repeat the same steps for the other instance also.

## 8 Latency Test

### 8.1 East-West

Of the 3 Instances on the same compute node, 2 instances have dedicated CPUs and 1 instance has shared CPUs. First ping between the instances with dedicated CPUs.

1. From the terminal session of the selected instance, ping external IP of other instance

```
[ubuntu@<instance1-name>~]$ping <external ip of other instance>
```

```
root@dpgk-gr4-i2:/home/ubuntu# ping -c 10 192.168.50.56
PING 192.168.50.56 (192.168.50.56) 56(84) bytes of data.
64 bytes from 192.168.50.56: icmp_seq=1 ttl=64 time=0.340 ms
64 bytes from 192.168.50.56: icmp_seq=2 ttl=64 time=0.218 ms
64 bytes from 192.168.50.56: icmp_seq=3 ttl=64 time=0.218 ms
64 bytes from 192.168.50.56: icmp_seq=4 ttl=64 time=0.273 ms
```



```
64 bytes from 192.168.50.56: icmp_seq=5 ttl=64 time=0.273 ms
64 bytes from 192.168.50.56: icmp_seq=6 ttl=64 time=0.277 ms
64 bytes from 192.168.50.56: icmp_seq=7 ttl=64 time=0.293 ms
64 bytes from 192.168.50.56: icmp_seq=8 ttl=64 time=0.276 ms
64 bytes from 192.168.50.56: icmp_seq=9 ttl=64 time=0.369 ms
64 bytes from 192.168.50.56: icmp_seq=10 ttl=64 time=0.300 ms
```

```
--- 192.168.50.56 ping statistics ---
10 packets transmitted, 10 received, 0% packet loss, time 9000ms
rtt min/avg/max/mdev = 0.218/0.283/0.369/0.048 ms
```

2. From the terminal session of the selected instance, ping guest-net IP of other instance.

```
[ubuntu@<instance1-name>~]$ping <guest-net ip of other instance>
```

```
root@dppk-gr4-i2:/home/ubuntu# ping -c 10 10.5.0.38
PING 10.5.0.38 (10.5.0.38) 56(84) bytes of data.
64 bytes from 10.5.0.38: icmp_seq=1 ttl=64 time=0.364 ms
64 bytes from 10.5.0.38: icmp_seq=2 ttl=64 time=0.322 ms
64 bytes from 10.5.0.38: icmp_seq=3 ttl=64 time=0.284 ms
64 bytes from 10.5.0.38: icmp_seq=4 ttl=64 time=0.234 ms
64 bytes from 10.5.0.38: icmp_seq=5 ttl=64 time=0.297 ms
64 bytes from 10.5.0.38: icmp_seq=6 ttl=64 time=0.233 ms
64 bytes from 10.5.0.38: icmp_seq=7 ttl=64 time=0.277 ms
64 bytes from 10.5.0.38: icmp_seq=8 ttl=64 time=0.279 ms
64 bytes from 10.5.0.38: icmp_seq=9 ttl=64 time=0.311 ms
64 bytes from 10.5.0.38: icmp_seq=10 ttl=64 time=0.265 ms
```

```
--- 10.5.0.38 ping statistics ---
10 packets transmitted, 10 received, 0% packet loss, time 8998ms
rtt min/avg/max/mdev = 0.233/0.286/0.364/0.042 ms
```

3. From the terminal session of the selected instance, ping DPDK IP of other instance

```
[ubuntu@<instance1-name>~]$ping <DPDK ip of other instance>
```

```
root@dppk-gr4-i2:/home/ubuntu# ping -c 10 12.0.0.13
PING 12.0.0.13 (12.0.0.13) 56(84) bytes of data.
64 bytes from 12.0.0.13: icmp_seq=1 ttl=64 time=0.452 ms
64 bytes from 12.0.0.13: icmp_seq=2 ttl=64 time=0.283 ms
64 bytes from 12.0.0.13: icmp_seq=3 ttl=64 time=0.245 ms
64 bytes from 12.0.0.13: icmp_seq=4 ttl=64 time=0.185 ms
64 bytes from 12.0.0.13: icmp_seq=5 ttl=64 time=0.170 ms
```



```
64 bytes from 12.0.0.13: icmp_seq=6 ttl=64 time=0.182 ms
64 bytes from 12.0.0.13: icmp_seq=7 ttl=64 time=0.125 ms
64 bytes from 12.0.0.13: icmp_seq=8 ttl=64 time=0.167 ms
64 bytes from 12.0.0.13: icmp_seq=9 ttl=64 time=0.173 ms
64 bytes from 12.0.0.13: icmp_seq=10 ttl=64 time=0.138 ms
```

```
--- 12.0.0.13 ping statistics ---
```

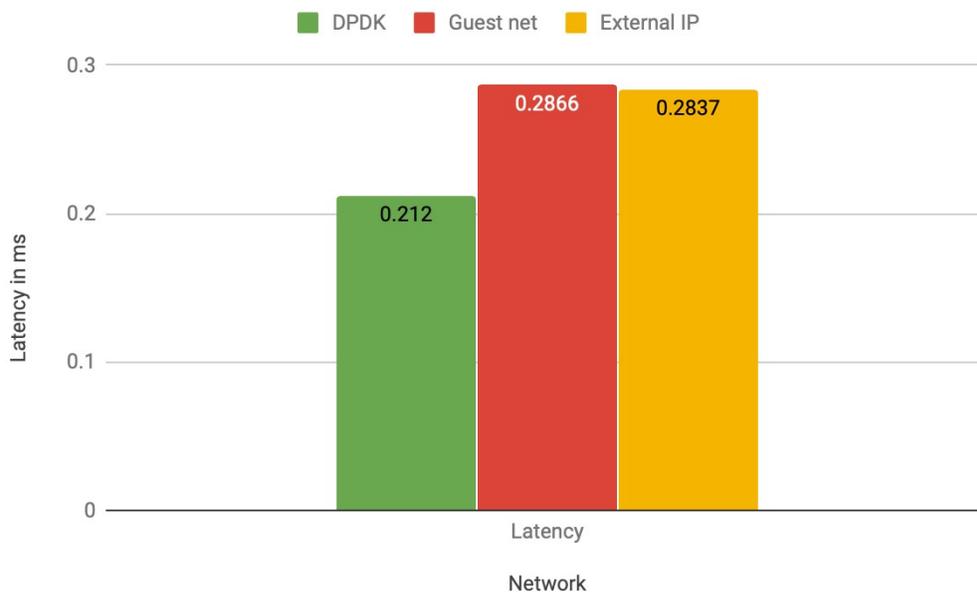
```
10 packets transmitted, 10 received, 0% packet loss, time 8999ms
```

```
rtt min/avg/max/mdev = 0.125/0.212/0.452/0.091 ms
```

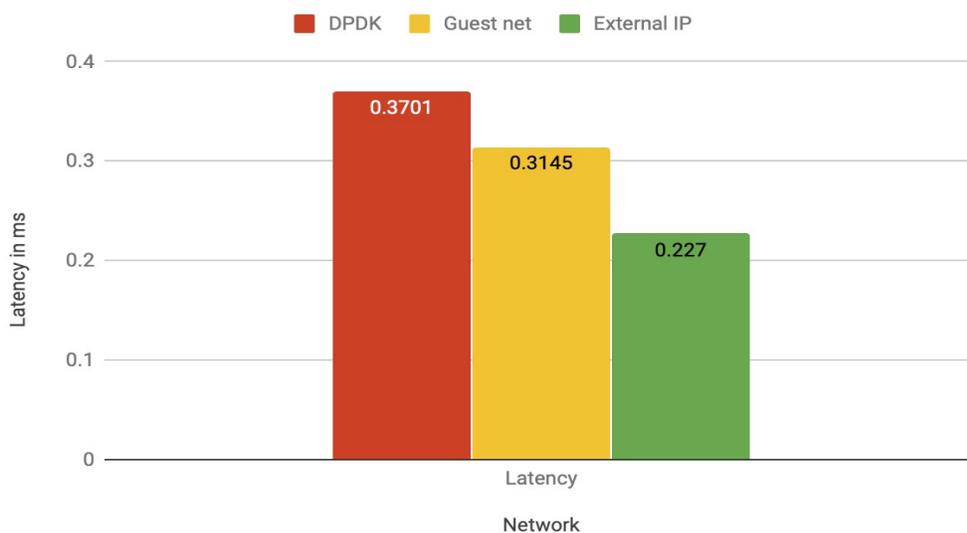
4. Perform the steps #1 thru #3 between an instance with shared CPU and other instance.

## 8.2 Latency Comparison Graphs

### 8.2.1 Between instances with dedicated CPUs



### 8.2.2 From an instance with shared CPUs





## 9 Running iperf

### 9.1 East-West

Of the 3 Instances on the same compute node, 2 instances have dedicated CPUs and 1 instance has shared CPUs. Pick an instance with dedicated CPUs to start iperf server and other instance with dedicated CPUs as iperf client.

1. From the terminal session of the selected instance, start “iperf server”

```
[ubuntu@<instance1-name>~]$ iperf3 -s
```

```
root@dppdk-gr4-i1:/home/ubuntu# iperf3 -s
```

```
-----  
Server listening on 5201  
-----
```

2. On the other instance, start iperf client on dpdk nic using the below command.

```
[root@<instance2-name>~]$ iperf3 -c <iperf server DPDK ip>
```

```
root@dppdk-gr4-i2:/home/ubuntu# iperf3 -c 12.0.0.13
```

```
Connecting to host 12.0.0.13, port 5201
```

```
[ 4] local 12.0.0.8 port 45804 connected to 12.0.0.13 port 5201
```

[ ID]	Interval		Transfer	Bandwidth	Retr	Cwnd	
[ 4]	0.00-1.00	sec	869 MBytes	7.29 Gbits/sec	22	683 KBytes	
[ 4]	1.00-2.00	sec	876 MBytes	7.35 Gbits/sec	0	683 KBytes	
[ 4]	2.00-3.00	sec	880 MBytes	7.38 Gbits/sec	0	683 KBytes	
[ 4]	3.00-4.00	sec	882 MBytes	7.40 Gbits/sec	0	683 KBytes	
[ 4]	4.00-5.00	sec	880 MBytes	7.38 Gbits/sec	0	684 KBytes	
[ 4]	5.00-6.00	sec	858 MBytes	7.19 Gbits/sec	0	684 KBytes	
[ 4]	6.00-7.00	sec	881 MBytes	7.39 Gbits/sec	32	615 KBytes	
[ 4]	7.00-8.00	sec	882 MBytes	7.40 Gbits/sec	0	615 KBytes	
[ 4]	8.00-9.00	sec	879 MBytes	7.37 Gbits/sec	0	626 KBytes	
[ 4]	9.00-10.00	sec	881 MBytes	7.39 Gbits/sec	0	656 KBytes	

```
-----  
[ ID] Interval          Transfer      Bandwidth      Retr  
[ 4] 0.00-10.00 sec 8.56 GBytes 7.35 Gbits/sec 54  
[ 4] 0.00-10.00 sec 8.56 GBytes 7.35 Gbits/sec  
sender  
receiver
```



4. Once the test completes (10 seconds), note down the transfer and bandwidth displayed.

5. Re-start iperf client using the floating IP

```
# iperf3 -c <iperf server floating ip>
```

```
root@dpedk-gr4-i2:/home/ubuntu# iperf3 -c 10.5.0.56
iperf3: error - unable to connect to server: Connection refused
root@dpedk-gr4-i2:/home/ubuntu# iperf3 -c 192.168.50.56
Connecting to host 192.168.50.56, port 5201
[ 4] local 192.168.50.60 port 33816 connected to 192.168.50.56 port 5201
[ ID] Interval            Transfer      Bandwidth      Retr  Cwnd
[ 4]  0.00-1.00    sec    148 MBytes  1.24 Gbits/sec    0   1.38 MBytes
[ 4]  1.00-2.00    sec    154 MBytes  1.29 Gbits/sec    0   1.38 MBytes
[ 4]  2.00-3.00    sec    155 MBytes  1.30 Gbits/sec    0   1.45 MBytes
[ 4]  3.00-4.00    sec    152 MBytes  1.28 Gbits/sec    0   1.45 MBytes
[ 4]  4.00-5.00    sec    155 MBytes  1.30 Gbits/sec    0   1.45 MBytes
[ 4]  5.00-6.00    sec    151 MBytes  1.27 Gbits/sec    0   1.67 MBytes
[ 4]  6.00-7.00    sec    154 MBytes  1.29 Gbits/sec    0   1.67 MBytes
[ 4]  7.00-8.00    sec    152 MBytes  1.28 Gbits/sec    0   1.67 MBytes
[ 4]  8.00-9.00    sec    155 MBytes  1.30 Gbits/sec    0   1.67 MBytes
[ 4]  9.00-10.00   sec    152 MBytes  1.28 Gbits/sec    0   1.67 MBytes
-----
[ ID] Interval            Transfer      Bandwidth      Retr
[ 4]  0.00-10.00   sec    1.49 GBytes  1.28 Gbits/sec    0
[ 4]  0.00-10.00   sec    1.49 GBytes  1.28 Gbits/sec
                               sender
                               receiver
```

6. Once the test completes (10 seconds), note down the transfer and bandwidth displayed.

7. Re-start iperf client using the guest-net IP

```
# iperf -c <iperf server guest net ip>
```

```
root@dpedk-gr4-i2:/home/ubuntu# iperf3 -c 10.5.0.38
Connecting to host 10.5.0.38, port 5201
[ 4] local 10.5.0.22 port 48386 connected to 10.5.0.38 port 5201
[ ID] Interval            Transfer      Bandwidth      Retr  Cwnd
[ 4]  0.00-1.00    sec    157 MBytes  1.32 Gbits/sec    0   1.14 MBytes
```



[ 4]	1.00-2.00	sec	158 MBytes	1.32 Gbits/sec	0	1.21 MBytes
[ 4]	2.00-3.00	sec	156 MBytes	1.31 Gbits/sec	0	1.21 MBytes
[ 4]	3.00-4.00	sec	152 MBytes	1.28 Gbits/sec	0	1.21 MBytes
[ 4]	4.00-5.00	sec	158 MBytes	1.32 Gbits/sec	16	557 KBytes
[ 4]	5.00-6.00	sec	156 MBytes	1.31 Gbits/sec	0	1.10 MBytes
[ 4]	6.00-7.00	sec	159 MBytes	1.33 Gbits/sec	0	1.13 MBytes
[ 4]	7.00-8.00	sec	152 MBytes	1.28 Gbits/sec	0	1.26 MBytes
[ 4]	8.00-9.00	sec	158 MBytes	1.32 Gbits/sec	0	1.26 MBytes
[ 4]	9.00-10.00	sec	159 MBytes	1.33 Gbits/sec	0	1.26 MBytes
-----						
[ ID]	Interval		Transfer	Bandwidth	Retr	
[ 4]	0.00-10.00	sec	1.53 GBytes	1.31 Gbits/sec	16	sender
[ 4]	0.00-10.00	sec	1.53 GBytes	1.31 Gbits/sec		receiver

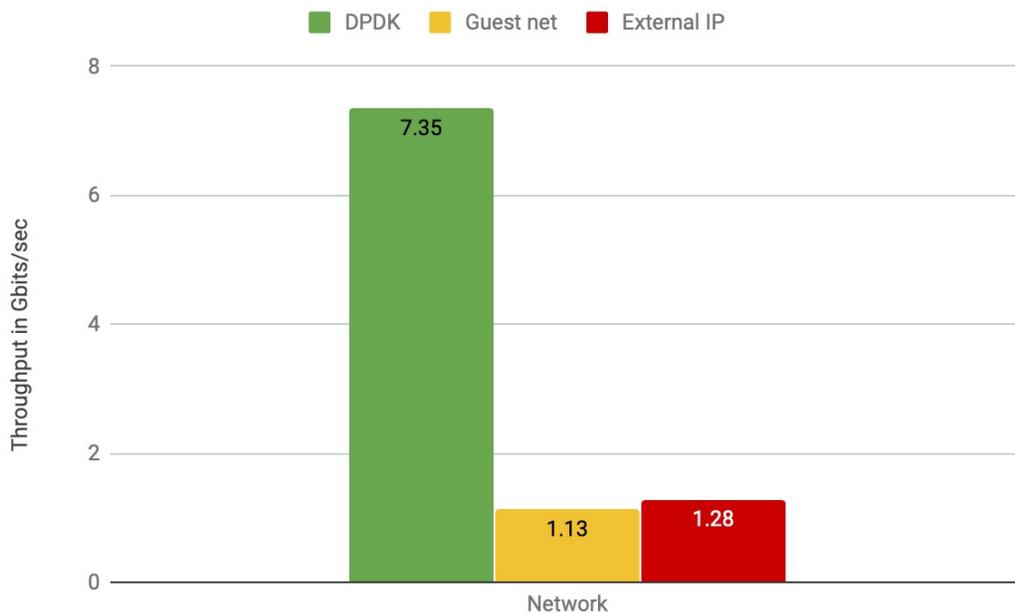
8. Once the test completes (10 seconds), note down the transfer and bandwidth displayed.

9. Compare the throughput between steps 4, 6 & 8. Step 4 throughput is in general higher than the throughput in steps 6 & 8. See the **highlighted** part in steps 4, 6 & 8 to observe the difference.

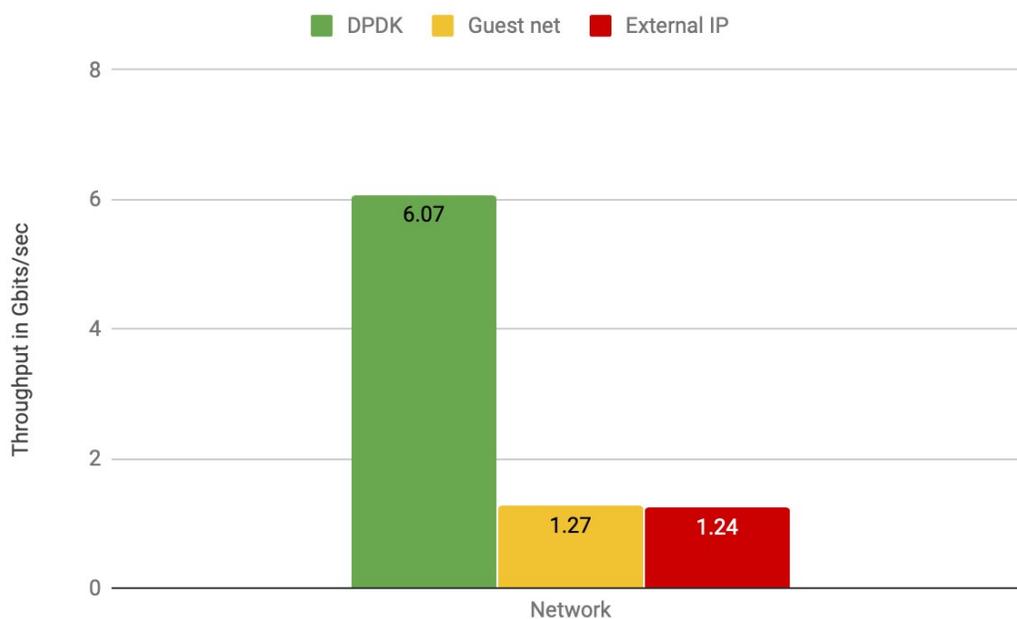


## 9.2 iperf Comparison Graphs

### 9.2.1 East - West



10. Perform the steps #1 thru #9 with iperf server on instance with dedicated CPUs and iperf client on instance with shared CPUs.





# 10 Appendix

## 10.1 Configuring OVS with DPDK in Openstack Environment

Below are a high level steps to configure OVS with DPDK in openstack environment.

### 10.1.1 Compute Node specific operations

- Configure Huge pages
- Install openvswitch-switch-dpdk package
- There are issues configuring OVS with DPDK that comes with the ubuntu package. So, using the link below, configure OVS with DPDK option.

<http://docs.openvswitch.org/en/latest/intro/install/dpdk/>

- Observed issues with Openstack Ansible, when OVS was configured from sources. To workaround the issue, we installed openvswitch-switch-dpdk package and then updating the files built from sources
- Once Openvswitch is configured with DPDK, set the below properties for OVS.
  - dpdk-extra="--pci-whitelist <interface PCI bus id>"
  - dpdk-init="true"
  - dpdk-lcore-mask=<specific cpu cores>
  - dpdk-socket-mem=<huge page memory>
  - pmd-cpu-mask=<cores dedicated for Poll mode drivers>
  - vhost-iommu-support="true"
- Below are the values configured for the lab environment.

```
{dpdk-extra="--pci-whitelist 0000:b1:00.0,0000:b1:00.1", dpdk-init="true", dpdk-lcore-mask="0xc00000", dpdk-socket-mem="0,4096", max-idle="50000", n-dpdk-rxqs="512", n-dpdk-txqs="512", pmd-cpu-mask="0x3000000000003000000", vhost-iommu-support="true"}
```
- For each NIC that you plan to use for DPDK, create a bridge in OVS. Below is the bridge created for Create Below are the values configured for the lab environment.

Bridge "br0"

Controller "tcp:127.0.0.1:6633"

is\_connected: true

fail\_mode: secure



```
Port "port0"  
  Interface "port0"  
    type: dpdk  
    options: {dpdk-devargs="0000:b1:00.0", n_rxq_desc="128",  
n_txq_desc="128"}
```

```
Port "br0"  
  Interface "br0"  
    type: internal
```

```
Port "phy-br0"  
  Interface "phy-br0"  
    type: patch  
    options: {peer="int-br0"}
```

### 10.1.2 Openstack Specific Settings

- On the controller and compute nodes, as part of ml2 configuration, add **openvswitch** and **I2population** to **mechanism\_drivers**
- On the compute node(s), in the neutron openvswitch agent, configure the below
  - Under the OVS section, add **integration\_bridge = br-int**
  - Under the OVS section, configure bridge\_mappings property to point to the dpdk bridge. In lab environment, the below is configured
    - bridge\_mappings = dpdk0:br0
  - Under the **Agent** section, **I2\_population** should be **True**
  - Under the **securitygroup** section, **firewall** driver should be **openvswitch**

### 10.1.3 Openstack Dashboard Specific Changes

1. Create a network for DPDK with **Provider Network Type** as **Flat** and for **Physical Network**, mention the name provided for device\_mappings in openvswitch\_agent.ini file.
2. A flavor needs to be configured with the below properties
  - a. aggregate\_instance\_extra\_specs:hpgs='true'
  - b. hw:cpu\_policy='dedicated'
  - c. hw:mem\_page\_size='1GB'
3. While creating instances, in the **Flavor** page, select a flavor with the properties mentioned in #2 above.
4. In the **Network** page, select the DPDK network
5. In the **Key Pair** page, provide a valid key-pair.





## 10.2 Reference

Add references of DPDK from Intel and OpenStack <TBD>

<https://01.org/packet-processing/blogs/nsundar/2018/nfv-i-host-configuration-low-latency>

<https://software.intel.com/sites/default/files/managed/4a/00/DPDK-Cookbook.pdf>

<https://software.intel.com/en-us/articles/dpdk-performance-optimization-guidelines-white-paper>

<https://software.intel.com/sites/default/files/managed/0e/b9/A-Low-Latency-NFV-Infrastructure-for-Performance-Critical-Applications.pdf>