

Week 1

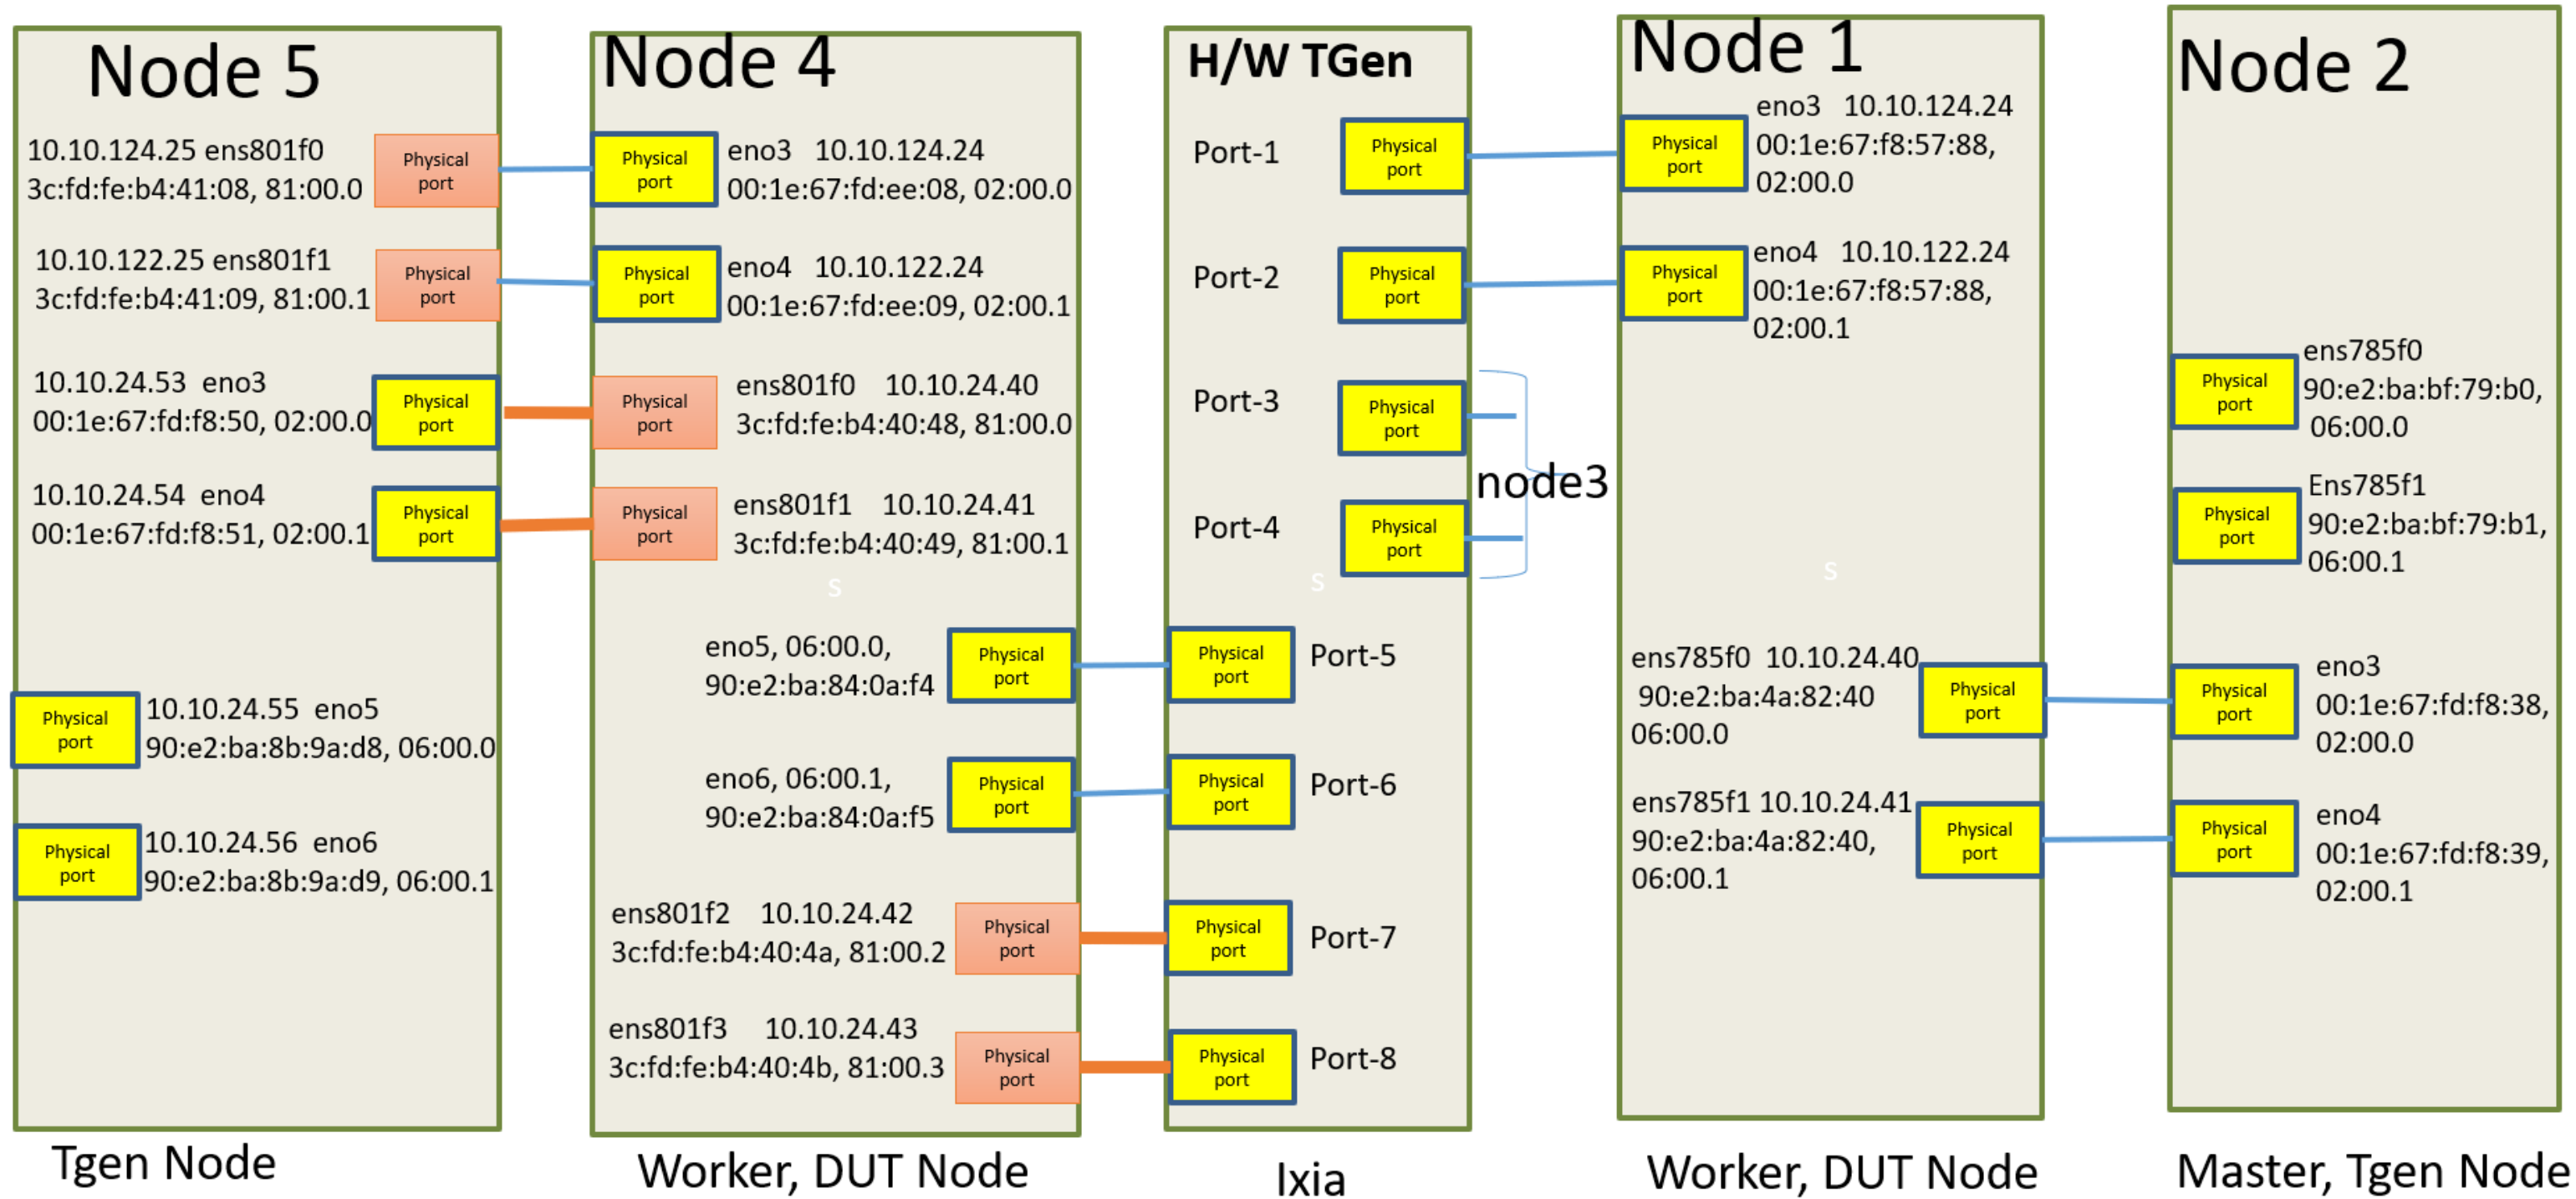
LF Networking

Project: XDP Performance studies for cloud-native NFV usecases

Tasks performed:

- Explored the architecture of OPNFV POD12 Cluster
- Installed VPP(Vector Packet Processing) on Node4
- Gone through RFC2544
- Added interfaces to VPP
- Used Trex(2.86) to generate traffic to benchmark VPP
- Results

Architecture of OPNFV POD12 Cluster



- Gone through RFC2544

Benchmarking Methodology for Network Interconnect Devices

Network Working Group
Request for Comments: 2544
Obsoletes: [1944](#)
Category: Informational

S. Bradner
Harvard University
J. McQuaid
NetScout Systems
March 1999

Benchmarking Methodology for Network Interconnect Devices

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

IESG Note

This document is a republication of [RFC 1944](#) correcting the values for the IP addresses which were assigned to be used as the default addresses for networking test equipment. (See section C.2.2). This RFC replaces and obsoletes [RFC 1944](#).

Abstract

This document discusses and defines a number of tests that may be used to describe the performance characteristics of a network interconnecting device. In addition to defining the tests this document also describes specific formats for reporting the results of the tests. [Appendix A](#) lists the tests and conditions that we believe should be included for specific cases and gives additional information about testing practices. [Appendix B](#) is a reference listing of maximum frame rates to be used with specific frame sizes on various media and [Appendix C](#) gives some examples of frame formats to be used in testing.

Installing VPP(Vector Packet Processing) on Node4

```

Installed:
  vpp.x86_64 0:20.09-release

Dependency Installed:
  vpp-lib.x86_64 0:20.09-release          vpp-selinux-policy.x86_64 0:20.09-release

Complete!

```

```

[root@pod12-node4 opnfv]# sudo systemctl status vpp.service
● vpp.service - Vector Packet Processing Process
   Loaded: loaded (/usr/lib/systemd/system/vpp.service; disabled; vendor preset: disabled)
   Active: inactive (dead)

```

```

[root@pod12-node4 opnfv]# sudo systemctl start vpp.service
[root@pod12-node4 opnfv]# sudo systemctl status vpp.service
● vpp.service - Vector Packet Processing Process
   Loaded: loaded (/usr/lib/systemd/system/vpp.service; disabled; vendor preset: disabled)
   Active: active (running) since Thu 2022-06-09 09:00:09 PDT; 3s ago
     Process: 82782 ExecStartPre=/sbin/modprobe uio_pci_generic (code=exited, status=0/SUCCESS)
     Process: 82778 ExecStartPre=/bin/rm -f /dev/shm/db /dev/shm/global_vm /dev/shm/vpe-api (code=exited, status=0/SUCCESS)
   Main PID: 82789 (vpp_main)
     Tasks: 1
    Memory: 29.2M
    CGroup: /system.slice/vpp.service
           └─82789 /usr/bin/vpp -c /etc/vpp/startup.conf

```

Listing available devices:

```
[root@pod12-node4 opnfv]# sudo lshw -class network -businfo
```

Bus info	Device	Class	Description
pci@0000:02:00.0	eno3	network	82599ES 10-Gigabit SFI/SFP+ Network Connection
pci@0000:02:00.1	eno4	network	82599ES 10-Gigabit SFI/SFP+ Network Connection
pci@0000:04:00.0	eno1	network	I350 Gigabit Network Connection
pci@0000:04:00.3	eno2	network	I350 Gigabit Network Connection
pci@0000:06:00.0	eno5	network	82599ES 10-Gigabit SFI/SFP+ Network Connection
pci@0000:06:00.1	eno6	network	82599ES 10-Gigabit SFI/SFP+ Network Connection
pci@0000:81:00.0	ens801f0	network	Ethernet Controller X710 for 10GbE SFP+
pci@0000:81:00.1	ens801f1	network	Ethernet Controller X710 for 10GbE SFP+
pci@0000:81:00.2	ens801f2	network	Ethernet Controller X710 for 10GbE SFP+
pci@0000:81:00.3	ens801f3	network	Ethernet Controller X710 for 10GbE SFP+
	br0	network	Ethernet interface
	virbr0	network	Ethernet interface
	flannel.1	network	Ethernet interface
	virbr0-nic	network	Ethernet interface
	docker0	network	Ethernet interface

The devices we're going to use and bind with dpdk

Now, lets unbind the devices from kernel and attach them to dpdk so as to use them with vpp

```

Network devices using DPDK-compatible driver
=====
<none>

Network devices using kernel driver
=====
0000:02:00.0 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' if=eno3 drv=ixgbe unused=uiopci_generic
0000:02:00.1 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' if=eno4 drv=ixgbe unused=uiopci_generic
0000:04:00.0 'I350 Gigabit Network Connection 1521' if=eno1 drv=igb unused=uiopci_generic
0000:04:00.3 'I350 Gigabit Network Connection 1521' if=eno2 drv=igb unused=uiopci_generic *Active*
0000:06:00.0 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' if=eno5 drv=ixgbe unused=uiopci_generic
0000:06:00.1 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' if=eno6 drv=ixgbe unused=uiopci_generic
0000:81:00.0 'Ethernet Controller X710 for 10GbE SFP+ 1572' if=ens801f0 drv=i40e unused=uiopci_generic *Active*
0000:81:00.1 'Ethernet Controller X710 for 10GbE SFP+ 1572' if=ens801f1 drv=i40e unused=uiopci_generic *Active*
0000:81:00.2 'Ethernet Controller X710 for 10GbE SFP+ 1572' if=ens801f2 drv=i40e unused=uiopci_generic
0000:81:00.3 'Ethernet Controller X710 for 10GbE SFP+ 1572' if=ens801f3 drv=i40e unused=uiopci_generic

Other Network devices
=====
<none>

```

```
sudo ./dpdk-devbind.py -u 02:00.0 02:00.1
```

```
Network devices using DPDK-compatible driver
=====
<none>

Network devices using kernel driver
=====
0000:04:00.0 'I350 Gigabit Network Connection 1521' if=eno1 drv=igb unused=vfio-pci,uio_pci_generic
0000:04:00.3 'I350 Gigabit Network Connection 1521' if=eno2 drv=igb unused=vfio-pci,uio_pci_generic *Active*
0000:06:00.0 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' if=eno5 drv=ixgbe unused=vfio-pci,uio_pci_generic
0000:06:00.1 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' if=eno6 drv=ixgbe unused=vfio-pci,uio_pci_generic
0000:81:00.0 'Ethernet Controller X710 for 10GbE SFP+ 1572' if=ens801f0 drv=i40e unused=vfio-pci,uio_pci_generic *Active*
0000:81:00.1 'Ethernet Controller X710 for 10GbE SFP+ 1572' if=ens801f1 drv=i40e unused=vfio-pci,uio_pci_generic *Active*
0000:81:00.2 'Ethernet Controller X710 for 10GbE SFP+ 1572' if=ens801f2 drv=i40e unused=vfio-pci,uio_pci_generic
0000:81:00.3 'Ethernet Controller X710 for 10GbE SFP+ 1572' if=ens801f3 drv=i40e unused=vfio-pci,uio_pci_generic

Other Network devices
=====
0000:02:00.0 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' unused=ixgbe,vfio-pci,uio_pci_generic
0000:02:00.1 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' unused=ixgbe,vfio-pci,uio_pci_generic
```

```
sudo ./dpdk-devbind.py --bind=vfio-pci 02:00.0 02:00.1
```

```
Network devices using DPDK-compatible driver
=====
0000:02:00.0 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' drv=vfio-pci unused=ixgbe,uio_pci_generic
0000:02:00.1 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' drv=vfio-pci unused=ixgbe,uio_pci_generic

Network devices using kernel driver
=====
0000:04:00.0 'I350 Gigabit Network Connection 1521' if=eno1 drv=igb unused=vfio-pci,uio_pci_generic
0000:04:00.3 'I350 Gigabit Network Connection 1521' if=eno2 drv=igb unused=vfio-pci,uio_pci_generic *Active*
0000:06:00.0 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' if=eno5 drv=ixgbe unused=vfio-pci,uio_pci_generic
0000:06:00.1 '82599ES 10-Gigabit SFI/SFP+ Network Connection 10fb' if=eno6 drv=ixgbe unused=vfio-pci,uio_pci_generic
0000:81:00.0 'Ethernet Controller X710 for 10GbE SFP+ 1572' if=ens801f0 drv=i40e unused=vfio-pci,uio_pci_generic *Active*
0000:81:00.1 'Ethernet Controller X710 for 10GbE SFP+ 1572' if=ens801f1 drv=i40e unused=vfio-pci,uio_pci_generic *Active*
0000:81:00.2 'Ethernet Controller X710 for 10GbE SFP+ 1572' if=ens801f2 drv=i40e unused=vfio-pci,uio_pci_generic
0000:81:00.3 'Ethernet Controller X710 for 10GbE SFP+ 1572' if=ens801f3 drv=i40e unused=vfio-pci,uio_pci_generic

Other Network devices
=====
<none>
```


In /etc/vpp/startup.conf

Adding the device PCI Address

```
## Whitelist specific interface by specifying PCI address
# dev 0000:06:00.0
# dev 0000:06:00.1
dev 0000:02:00.0
dev 0000:02:00.1
```

Now, restart vpp.service

```
● vpp.service - Vector Packet Processing Process
   Loaded: loaded (/usr/lib/systemd/system/vpp.service; disabled; vendor preset: disabled)
   Active: active (running) since Sat 2022-06-11 10:22:50 PDT; 17s ago
     Process: 15529 ExecStartPre=/sbin/modprobe uio_pci_generic (code=exited, status=0/SUCCESS)
     Process: 15526 ExecStartPre=/bin/rm -f /dev/shm/db /dev/shm/global_vm /dev/shm/vpe-api (code=exited, status=0/SUCCESS)
    Main PID: 15532 (vpp_main)
      Tasks: 11
     Memory: 85.1M
    CGroup: /system.slice/vpp.service
            └─15532 /usr/bin/vpp -c /etc/vpp/startup.conf
```

```
vpp# show interface
```

Name	Idx	State	MTU (L3/IP4/IP6/MPLS)	Counter	Count
TenGigabitEthernet2/0/0	1	down	9000/0/0/0		
TenGigabitEthernet2/0/1	2	down	9000/0/0/0		
local0	0	down	0/0/0/0		

Setting interfaces up

```
vpp# set interface state TenGigabitEthernet2/0/1 up
vpp# show interface
```

Name	Idx	State	MTU (L3/IP4/IP6/MPLS)	Counter	Count
TenGigabitEthernet2/0/0	1	up	9000/0/0/0	rx packets	4
				rx bytes	440
				drops	4
TenGigabitEthernet2/0/1	2	up	9000/0/0/0	rx packets	1
				rx bytes	110
				drops	1
local0	0	down	0/0/0/0		

```
vpp#
```

Creating a loopback to get the packets back to Trex

```
vpp#
vpp# test l2patch rx TenGigabitEthernet2/0/0 tx TenGigabitEthernet2/0/1
vpp# test l2patch rx TenGigabitEthernet2/0/1 tx TenGigabitEthernet2/0/0
vpp#
vpp#
vpp# show l2patch
    TenGigabitEthernet2/0/0 -> TenGigabitEthernet2/0/1
    TenGigabitEthernet2/0/1 -> TenGigabitEthernet2/0/0
vpp# █
```

Running Trex on Node 5

```

shivank@shivank-Legion-340-1... x  root@pod12-node5:~/trex_2.86 x op
-Per port stats table
  ports |           0 |           1
-----|-----|-----
  opackets |      20550381 |      20543764
   obytes |  25132609442 |  25125007666
  ipackets |      20543764 |      20550381
   ibytes |  25125007666 |  25132609442
   ierrors |           0 |           0
   oerrors |           0 |           0
   Tx Bw |  712.26 Kbps |  713.46 Kbps

-Global stats enabled
Cpu Utilization : 1.3 % 0.2 Gb/core
Platform_factor : 1.0
Total-Tx       :      1.43 Mbps
Total-Rx       :      1.43 Mbps
Total-PPS      :     117.40 pps
Total-CPS      :       0.00 cps

Expected-PPS   :       0.00 pps
Expected-CPS   :       0.00 cps
Expected-BPS   :       0.00 bps

Active-flows   :          0 Clients :          0 Socket-util : 0.0000 %
Open-flows    :          0 Servers :          0 Socket      : 0 Socket/Clients : -nan
Total_queue_full : 15863501
drop-rate     :       0.00 bps
current time  : 76.5 sec
test duration  : 0.0 sec

```

Thankyou