

R3 Test Document of IEC Type 5: SmartNIC for Integrated Edge Cloud (IEC) Blueprint Family

- Introduction
- Akraino Test Group Information
 - Overall Test Architecture
 - Test Bed
 - Test Framework
 - Traffic Generator
- Test API description
- Test Dashboards
 - Functional Tests
 - Performance Tests
 - Single PF
 - Single PF, Single VF
 - Single PF, 4 VFs (only test with offloaded)
- Additional Testing
- Bottlenecks/Errata

Introduction

The R3 release will evaluation the throughput and packet forwarding performance of the Mellanox BlueField SmartNIC card.

A DPDK based [Open vSwitch](#) (OVS-DPDK) is used as the virtual switch, and the network traffic is virtualized with the VXLAN encapsulation.

Currently, the community version OVS-DPDK is considered experimental and not mature enough, it only supports "partial offload" which cannot utilize the full performance advantage of Mellanox NICs. Thus the OVS-DPDK we used is a fork of the community Open vSwitch. We develop our own offload code which enables the full hardware offload with DPDK rte_flow APIs.

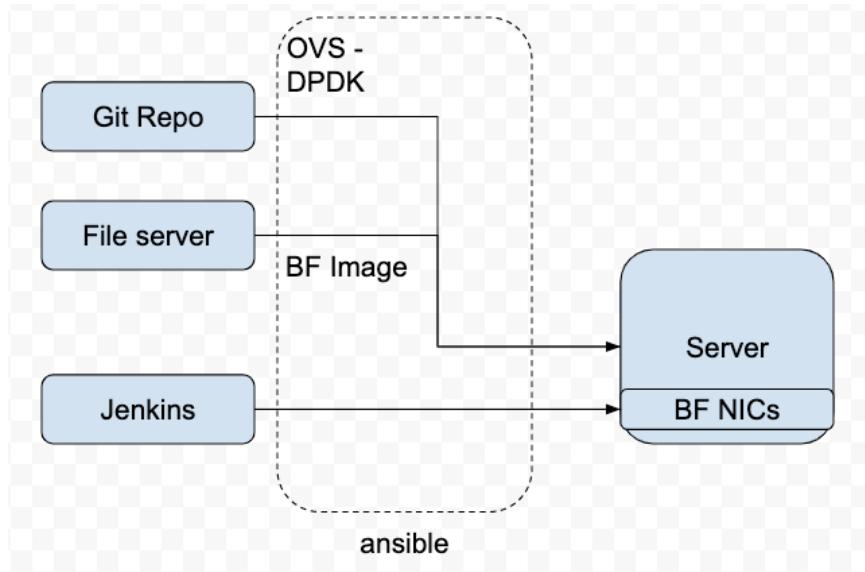
We have plans to open-source this OVS-DPDK. More details will be provided in future documentation.

Akraino Test Group Information

Overall Test Architecture

To deploy the Test architecture, we use a private Jenkins and an Intel server equipped with a BlueField v1 SmartNIC.

We use [Ansible](#) to automatically setup the filesystem image and install the OVS-DPDK in the SmartNICs.



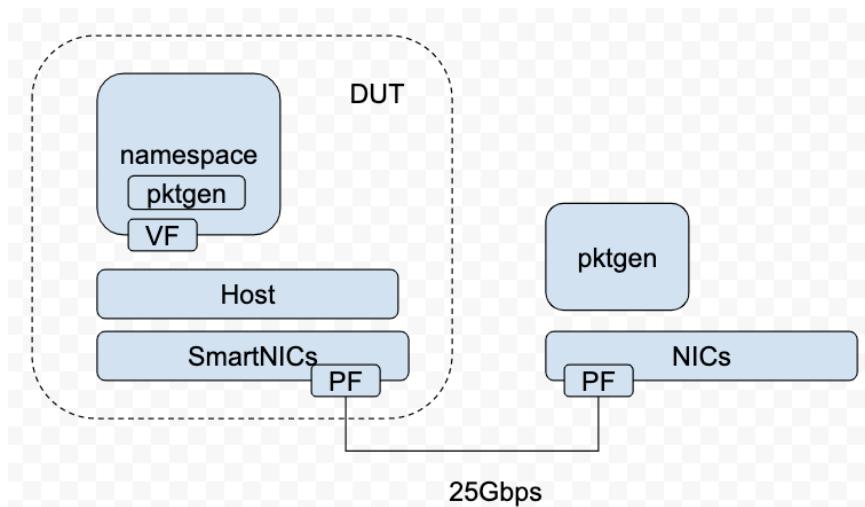
The File Server is a simple [Nginx](#) based web server where stores the BF drivers, FS image.

The Git repo is our own git repo where hosts OVS-DPDK and DPDK code.

The Jenkins will use ansible plugin to download BF drivers and FS image in the test server and setup the environment according to the ansible-playbook.

Image	download link
BlueField-2.5.1.11213.tar.xz	https://www.mellanox.com/products/software/bluefield
core-image-full-dev-BlueField-2.5.1.11213.2.5.3.tar.xz	https://www.mellanox.com/products/software/bluefield
mft-4.14.0-105-x86_64-deb.tgz	
MLNX_OFED_LINUX-5.0-2.1.8.0-debian8.11-x86_64.tgz	

Test Bed



The testbed setup is shown in the above diagram. DUT stands for Device Under Test

Test Framework

The software used and the OVS-DPDK setup is shown below.

Type	Description
SmartNICs	BlueField v1, 25Gbps
DPDK	version 19.11
vSwitch	OVS-DPDK 2.12 with VXLAN DECAP/ENCAP offload enabled. https://github.com/bytedance/ovs-dpdk

```

root@bluefield:/home/ovs-dpdk# ovs-vsctl show
2dccb148-526c-44a5-9351-67b04c5e2da4
  Bridge br-int
    datapath_type: netdev
    Port vxlan-vtp
      Interface vxlan-vtp
        type: vxlan
        options: {dst_port="4789", key=flow, local_ip="192.168.1.1", remote_ip=flow, tos=inherit}
  Port br-int
    Interface br-int
      type: internal
  Port pf1hpf
    Interface pf1hpf
      type: dpdk
      options: {dpdk-devargs="class=eth,mac=ae:d8:8a:c5:22:fb"}
Bridge br-ex
  datapath_type: netdev
  Port br-ex
    Interface br-ex
      type: internal
  Port p1
    Interface p1
      type: dpdk
      options: {dpdk-devargs="class=eth,mac=98:03:9b:af:7b:0b"}

```

```

root@bluefield:/home/ovs-dpdk# ovs-vsctl list open_vswitch
_uuid          : 2dccb148-526c-44a5-9351-67b04c5e2da4
bridges        : [22334686-733a-445e-9130-a42009a3586e, 38af610d-01f7-497d-878b-c6b6a44abf6a]
cur_cfg        : 10
datapath_types : [netdev, system]
datapaths      : {}
db_version     : []
dpdk_initialized: true
dpdk_version   : "DPDK 19.11.0"
external_ids   : {}
iface_types    : [dpdk, dpdkr, dpdkvhostuser, dpdkvhostuserclient, erspan, geneve, gre, internal,
ip6erspan, ip6gre, lisp, patch, stt, system, tap, vxlan]
manager_options: []
next_cfg       : 10
other_config   : {dpdk-extra="-w 03:00.1,representor=[0,65535] --legacy-mem ", dpdk-init="true", hw-
offload="true"}
ovs_version    : []
ssl            : []
statistics     : {cpu="16", file_systems="/,13521220,3918464 /data,243823,2064 /boot,357176,61104",
load_average="1.41,1.37,1.36", memory="16337652,5589928,1707640,0,0", process_ovs-vswitchd="5959388,256352,11515170,0,11484792,11484792,7", process_ovsdb-server="12620,6260,13190,0,68233832,68233832,6"}
system_type    : []
system_version: []

```

Traffic Generator

We will use DPDK pktgen as the Traffic Generator.

Test API description

The test is to evaluate the performance of SmartNIC offloading.

Thus we currently don't have any Test APIs provided.

Test Dashboards

Functional Tests

Open vSwitch itself contains a test suite for functional test, the link is <http://docs.openvswitch.org/en/latest/topics/testing/>

We have run the basic test suite according to the link.

By running,

```
make check TESTSUITEFLAGS=-j8
```

We got the below results.

Total Tests	Test Executed	Pass	Fail	In Progress
2225	2200	2198	2	0

The two failed cases are about sFlow sampling. We are investigating the internal reason.

25 cases are skipped due to the configuration.

Performance Tests

Single PF

OVS-DPDK	OF rules	Traffic	pktgen pps	received pps (hw)	received pps (no hw)	PMD idle cycles w/ hw offload	PMD idle cycles w/o hw offload
1 PMD(s)	Directly forwarding without the match "in_port=vxlan-vtp, actions=output: pf1hp1"	single TCP flow with VXLAN encapsulation	24.6Mpps	23.9Mpps	745156	99%	0%

Single PF, Single VF

OVS-DPDK	OF rules	Traffic	pktgen pps	received pps (hw)	received pps (no hw)	PMD idle cycles w/ hw offload	PMD idle cycles w/o hw offload
1 PMD(s)	Directly forwarding without the match "in_port=vxlan-vtp, actions=output: pf1vf0"	single TCP flow with VXLAN encapsulation	24.6Mpps	23.8Mpps	747097	99%	0%
1 PMD(s)	match and forwarding 100 rules "in_port=vxlan-vtp.ip,nw_dst=10.1.1.[1-100] actions=pf1vf0"	100 flows with nw_dst=10.1.1.[1-100]	24.6Mpps	21.6Mpps (100 megaflows offloaded)	624161	99%	0%
1 PMD(s)	match and forwarding 1000 rules "in_port=vxlan-vtp.ip,nw_dst=10.1.[1-10].[1-100] actions=pf1vf0"	1000 flows with nw_dst=10.1.[1-10].[1-100] 891 flows (due to pktgen limit, only 891 flows are loaded)	24.6Mpps	23.3Mpps (891 megaflows offloaded)	524144	99%	0%

Single PF, 4 VFs (only test with offloaded)

OVS-DPDK	OF rules	Traffic	pktgen pps	received pps (VF0)	received pps (VF1)	received pps (VF2)	received pps (VF3)	in total
----------	----------	---------	------------	--------------------	--------------------	--------------------	--------------------	----------

1 PMD(s)	match and forwarding 1000 rules "in_port=vxlan-vtp.ip,nw_dst=10.1.[1-10].[1-25] actions=pf1vf0" "in_port=vxlan-vtp.ip,nw_dst=10.1.[1-10].[26-50] actions=pf1vf1" "in_port=vxlan-vtp.ip,nw_dst=10.1.[1-10].[51-75] actions=pf1vf2" "in_port=vxlan-vtp.ip,nw_dst=10.1.[1-10].[75-100] actions=pf1vf3"	1000 flows with nw_dst=10.1.[1-10].[1-100] 891 flows (due to pktgen limit, only 891 flows are loaded)	24.6Mpps	5891903	5781310	5767825	5380159	232588 08 ~ 23.2 Mpps
----------	--	---	----------	---------	---------	---------	---------	------------------------------

Additional Testing

n/a

Bottlenecks/Errata

n/a