



VSPERF User Guide

Release draft (5853025)

OPNFV

March 12, 2016

1	vSwitchPerf test suites userguide	1
1.1	General	1
1.2	VSPERF Installation	1
1.3	Traffic Generator Setup	1
1.4	Cloning and building src dependencies	1
1.5	Configure the <code>./conf/10_custom.conf</code> file	2
1.6	Using a custom settings file	2
1.7	<code>vloop_vnf</code>	2
1.8	<code>l2fwd</code> Kernel Module	2
1.9	Executing tests	3
1.10	Executing Vanilla OVS tests	3
1.11	Executing PVP and PVVP tests	4
1.12	Executing PVP tests using Vanilla OVS	4
1.13	Using <code>vfio_pci</code> with DPDK	5
1.14	Selection of loopback application for PVP and PVVP tests	5
1.15	Executing Packet Forwarding tests	6
1.16	VSPERF modes of operation	6
1.17	Code change verification by <code>pylint</code>	7
1.18	GOTCHAs:	7
1.19	More information	8
2	Integration tests	9
2.1	Executing Tunnel encapsulation tests	9
2.2	Executing VXLAN decapsulation tests	10
2.3	Executing GRE decapsulation tests	11
2.4	Executing GENEVE decapsulation tests	11

VSWITCHPERF TEST SUITES USERGUIDE

1.1 General

VSPERF requires a traffic generators to run tests, automated traffic gen support in VSPERF includes:

- IXIA traffic generator (IxNetwork hardware) and a machine that runs the IXIA client software.
- Spirent traffic generator (TestCenter hardware chassis or TestCenter virtual in a VM) and a VM to run the Spirent Virtual Deployment Service image, formerly known as “Spirent LabServer”.

If you want to use another traffic generator, please select the Dummy generator option as shown in [Traffic generator instructions](#)

1.2 VSPERF Installation

To see the supported Operating Systems, vSwitches and system requirements, please follow the [installation instructions](#) to install.

1.3 Traffic Generator Setup

Follow the [Traffic generator instructions](#) to install and configure a suitable traffic generator.

1.4 Cloning and building src dependencies

In order to run VSPERF, you will need to download DPDK and OVS. You can do this manually and build them in a preferred location, OR you could use vswitchperf/src. The vswitchperf/src directory contains makefiles that will allow you to clone and build the libraries that VSPERF depends on, such as DPDK and OVS. To clone and build simply:

```
$ cd src
$ make
```

VSPERF can be used with stock OVS (without DPDK support). When build is finished, the libraries are stored in src_vanilla directory.

The ‘make’ builds all options in src:

- Vanilla OVS
- OVS with vhost_user as the guest access method (with DPDK support)

- OVS with vhost_cuse s the guest access method (with DPDK support)

The vhost_user build will reside in src/ovs/ The vhost_cuse build will reside in vswitchperf/src_cuse The Vanilla OVS build will reside in vswitchperf/src_vanilla

To delete a src subdirectory and its contents to allow you to re-clone simply use:

```
$ make clobber
```

1.5 Configure the ./conf/10_custom.conf file

The 10_custom.conf file is the configuration file that overrides default configurations in all the other configuration files in ./conf The supplied 10_custom.conf file **MUST** be modified, as it contains configuration items for which there are no reasonable default values.

The configuration items that can be added is not limited to the initial contents. Any configuration item mentioned in any .conf file in ./conf directory can be added and that item will be overridden by the custom configuration value.

1.6 Using a custom settings file

If your 10_custom.conf doesn't reside in the ./conf directory of if you want to use an alternative configuration file, the file can be passed to vsperf via the --conf-file argument.

```
$ ./vsperf --conf-file <path_to_custom_conf> ...
```

Note that configuration passed in via the environment (--load-env) or via another command line argument will override both the default and your custom configuration files. This “priority hierarchy” can be described like so (1 = max priority):

1. Command line arguments
2. Environment variables
3. Configuration file(s)

1.7 vloop_vnf

vsperf uses a VM called vloop_vnf for looping traffic in the PVP and PVVP deployment scenarios. The image can be downloaded from <http://artifacts.opnfv.org/>.

```
$ wget http://artifacts.opnfv.org/vswitchperf/vloop-vnf-ubuntu-14.04_20151216.qcow2
```

vloop_vnf forwards traffic through a VM using one of: * DPDK testpmd * Linux Bridge * l2fwd kernel Module.

Alternatively you can use your own QEMU image.

1.8 l2fwd Kernel Module

A Kernel Module that provides OSI Layer 2 Ipv4 termination or forwarding with support for Destination Network Address Translation (DNAT) for both the MAC and IP addresses. l2fwd can be found in <vswitchperf_dir>/src/l2fwd

1.9 Executing tests

Before running any tests make sure you have root permissions by adding the following line to /etc/sudoers:

```
username ALL=(ALL) NOPASSWD: ALL
```

username in the example above should be replaced with a real username.

To list the available tests:

```
$ ./vsperf --list
```

To run a single test:

```
$ ./vsperf $TESTNAME
```

Where \$TESTNAME is the name of the vsperf test you would like to run.

To run a group of tests, for example all tests with a name containing 'RFC2544':

```
$ ./vsperf --conf-file=<path_to_custom_conf>/10_custom.conf --tests="RFC2544"
```

To run all tests:

```
$ ./vsperf --conf-file=<path_to_custom_conf>/10_custom.conf
```

Some tests allow for configurable parameters, including test duration (in seconds) as well as packet sizes (in bytes).

```
$ ./vsperf --conf-file user_settings.py
  --tests RFC2544Tput
  --test-param "duration=10;pkt_sizes=128"
```

For all available options, check out the help dialog:

```
$ ./vsperf --help
```

1.10 Executing Vanilla OVS tests

1. If needed, recompile src for all OVS variants

```
$ cd src
$ make distclean
$ make
```

2. Update your "10_custom.conf" file to use the appropriate variables for Vanilla OVS:

```
VSWITCH = 'OvsVanilla'
VSWITCH_VANILLA_PHY_PORT_NAMES = ['$PORT1', '$PORT2']
```

Where \$PORT1 and \$PORT2 are the Linux interfaces you'd like to bind to the vswitch.

3. Run test:

```
$ ./vsperf --conf-file=<path_to_custom_conf>
```

Please note if you don't want to configure Vanilla OVS through the configuration file, you can pass it as a CLI argument; BUT you must set the ports.

```
$ ./vsperf --vswitch OvsVanilla
```

1.11 Executing PVP and PVVP tests

To run tests using vhost-user as guest access method:

1. Set VHOST_METHOD and VNF of your settings file to:

```
VHOST_METHOD='user'  
VNF = 'QemuDpdkVhost'
```

2. If needed, recompile src for all OVS variants

```
$ cd src  
$ make distclean  
$ make
```

3. Run test:

```
$ ./vsperf --conf-file=<path_to_custom_conf>/10_custom.conf
```

To run tests using vhost-cuse as guest access method:

1. Set VHOST_METHOD and VNF of your settings file to:

```
VHOST_METHOD='cuse'  
VNF = 'QemuDpdkVhostCuse'
```

2. If needed, recompile src for all OVS variants

```
$ cd src  
$ make distclean  
$ make
```

3. Run test:

```
$ ./vsperf --conf-file=<path_to_custom_conf>/10_custom.conf
```

1.12 Executing PVP tests using Vanilla OVS

To run tests using Vanilla OVS:

1. Set the following variables:

```
VSWITCH = 'OvsVanilla'  
VNF = 'QemuVirtioNet'  
  
VANILLA_TGEN_PORT1_IP = n.n.n.n  
VANILLA_TGEN_PORT1_MAC = nn:nn:nn:nn:nn:nn  
  
VANILLA_TGEN_PORT2_IP = n.n.n.n  
VANILLA_TGEN_PORT2_MAC = nn:nn:nn:nn:nn:nn  
  
VANILLA_BRIDGE_IP = n.n.n.n  
  
or use --test-param
```



```
$ ./vsperf --conf-file=<path_to_custom_conf>/10_custom.conf
  --test-param "vanilla_tgen_tx_ip=n.n.n.n;
               vanilla_tgen_tx_mac=nn:nn:nn:nn:nn:nn"
```

2. If needed, recompile src for all OVS variants

```
$ cd src
$ make distclean
$ make
```

3. Run test:

```
$ ./vsperf --conf-file<path_to_custom_conf>/10_custom.conf
```

1.13 Using vfio_pci with DPDK

To use vfio with DPDK instead of igb_uio edit 'conf/02_vswitch.conf' with the following parameters:

```
DPDK_MODULES = [
  ('vfio-pci'),
]
SYS_MODULES = ['cuse']
```

NOTE: Please ensure that Intel VT-d is enabled in BIOS.

NOTE: Please ensure your boot/grub parameters include the following:

```
iommu=pt intel_iommu=on
```

To check that IOMMU is enabled on your platform:

```
$ dmesg | grep IOMMU
[ 0.000000] Intel-IOMMU: enabled
[ 0.139882] dmar: IOMMU 0: reg_base_addr fbffe000 ver 1:0 cap d2078c106f0466 ecap f020de
[ 0.139888] dmar: IOMMU 1: reg_base_addr ebffc000 ver 1:0 cap d2078c106f0466 ecap f020de
[ 0.139893] IOAPIC id 2 under DRHD base 0xfbffe000 IOMMU 0
[ 0.139894] IOAPIC id 0 under DRHD base 0xebffc000 IOMMU 1
[ 0.139895] IOAPIC id 1 under DRHD base 0xebffc000 IOMMU 1
[ 3.335744] IOMMU: dmar0 using Queued invalidation
[ 3.335746] IOMMU: dmar1 using Queued invalidation
....
```

1.14 Selection of loopback application for PVP and PVVP tests

To select loopback application, which will perform traffic forwarding inside VM, following configuration parameter should be configured:

```
GUEST_LOOPBACK = ['testpmd', 'testpmd']
```

or use `--test-param`

```
$ ./vsperf --conf-file=<path_to_custom_conf>/10_custom.conf
  --test-param "guest_loopback=testpmd"
```

Supported loopback applications are:

```
'testpmd'      - testpmd from dpdk will be built and used
'l2fwd'        - l2fwd module provided by Huawei will be built and used
'linux_bridge' - linux bridge will be configured
'buildin'      - nothing will be configured by vsperf; VM image must
                  ensure traffic forwarding between its interfaces
```

Guest loopback application must be configured, otherwise traffic will not be forwarded by VM and testcases with PVP and PVVP deployments will fail. Guest loopback application is set to 'testpmd' by default.

1.15 Executing Packet Forwarding tests

To select application, which will perform packet forwarding, following configuration parameter should be configured:

```
VSWITCH = 'none'
PKTFWD = 'TestPMD'

or use --vswitch and --fwdapp

$ ./vsperf --conf-file user_settings.py
    --vswitch none
    --fwdapp TestPMD
```

Supported Packet Forwarding applications are:

```
'testpmd'      - testpmd from dpdk
```

1. Update your "10_custom.conf" file to use the appropriate variables for selected Packet Forwarder:

```
# testpmd configuration
TESTPMD_ARGS = []
# packet forwarding mode: io|mac|mac_retry|macswap|flowgen|rxonly|txonly|csum|icmpecho
TESTPMD_FWD_MODE = 'csum'
# checksum calculation layer: ip|udp|tcp|sctp|outer-ip
TESTPMD_CSUM_LAYER = 'ip'
# checksum calculation place: hw (hardware) | sw (software)
TESTPMD_CSUM_CALC = 'sw'
# recognize tunnel headers: on|off
TESTPMD_CSUM_PARSE_TUNNEL = 'off'
```

2. Run test:

```
$ ./vsperf --conf-file <path_to_settings_py>
```

1.16 VSPERF modes of operation

VSPERF can be run in different modes. By default it will configure vSwitch, traffic generator and VNF. However it can be used just for configuration and execution of traffic generator. Another option is execution of all components except traffic generator itself.

Mode of operation is driven by configuration parameter -m or --mode

```
-m MODE, --mode MODE vsperf mode of operation;
Values:
    "normal" - execute vSwitch, VNF and traffic generator
    "trafficgen" - execute only traffic generator
```

```
"trafficgen-off" - execute vSwitch and VNF
"trafficgen-pause" - execute vSwitch and VNF but wait before traffic transmission
```

In case, that VSPERF is executed in “trafficgen” mode, then configuration of traffic generator should be configured through –test-param option. Supported CLI options useful for traffic generator configuration are:

```
'traffic_type' - One of the supported traffic types. E.g. rfc2544,
                 back2back or continuous
                 Default value is "rfc2544".
'bidirectional' - Specifies if generated traffic will be full-duplex (true)
                 or half-duplex (false)
                 Default value is "false".
'iloading' - Defines desired percentage of frame rate used during
             continuous stream tests.
             Default value is 100.
'multistream' - Defines number of flows simulated by traffic generator.
               Value 0 disables MultiStream feature
               Default value is 0.
'stream_type' - Stream Type is an extension of the "MultiStream" feature.
               If MultiStream is disabled, then Stream Type will be
               ignored. Stream Type defines ISO OSI network layer used
               for simulation of multiple streams.
               Default value is "L4".
```

Example of execution of VSPERF in “trafficgen” mode:

```
$ ./vsperf -m trafficgen --trafficgen IxNet --conf-file vsperf.conf
  --test-params "traffic_type=continuous;bidirectional=True;iloading=60"
```

1.17 Code change verification by pylint

Every developer participating in VSPERF project should run pylint before his python code is submitted for review. Project specific configuration for pylint is available at ‘pylint.rc’.

Example of manual pylint invocation:

```
$ pylint --rcfile ./pylintrc ./vsperf
```

1.18 GOTCHAs:

1.18.1 OVS with DPDK and QEMU

If you encounter the following error: “before (last 100 chars): ‘-path=/dev/hugepages,share=on: unable to map backing store for hugepages: Cannot allocate memory” with the PVP or PVVP deployment scenario, check the amount of hugepages on your system:

```
$ cat /proc/meminfo | grep HugePages
```

By default the vswitchd is launched with 1Gb of memory, to change this, modify –socket-mem parameter in conf/02_vswitch.conf to allocate an appropriate amount of memory:

```
VSWITCHD_DPDK_ARGS = ['-c', '0x4', '-n', '4', '--socket-mem 1024,0']
```

1.19 More information

For more information and details refer to the vSwitchPerf user guide at:
<http://artifacts.opnfv.org/vswitchperf/brahmaputra/userguide/index.html>

INTEGRATION TESTS

VSPERF includes a set of integration tests defined in `conf/integration`. These tests can be run by specifying `--run-integration` as a parameter to `vsperf`. Current tests in `conf/integration` are Overlay tests.

2.1 Executing Tunnel encapsulation tests

VSPERF supports VXLAN, GRE and GENEVE tunneling protocols. Testing of these protocols is limited to unidirectional traffic and P2P (Physical to Physical scenarios).

The VXLAN OVS DPDK encapsulation tests requires IPs, MAC addresses, bridge names and `WHITELIST_NICS` for DPDK.

Default values are already provided. To customize for your environment, override the following variables in you `user_settings.py` file:

```
# Variables defined in conf/integration/02_vswitch.conf # Tunnel endpoint for Overlay P2P deployment scenario #
used for br0 VTEP_IP1 = '192.168.0.1/24'

# Used as remote_ip in adding OVS tunnel port and # to set ARP entry in OVS (e.g. tnl/arp/set br-ext 192.168.240.10
02:00:00:00:00:02 VTEP_IP2 = '192.168.240.10'

# Network to use when adding a route for inner frame data VTEP_IP2_SUBNET = '192.168.240.0/24'

# Bridge names TUNNEL_INTEGRATION_BRIDGE = 'br0' TUNNEL_EXTERNAL_BRIDGE = 'br-ext'

# IP of br-ext TUNNEL_EXTERNAL_BRIDGE_IP = '192.168.240.1/24'

# vxlan|gre|geneve TUNNEL_TYPE = 'vxlan'

# Variables defined conf/integration/03_traffic.conf # For OP2P deployment scenario TRAFFICGEN_PORT1_MAC
= '02:00:00:00:00:01' TRAFFICGEN_PORT2_MAC = '02:00:00:00:00:02' TRAFFICGEN_PORT1_IP = '1.1.1.1'
TRAFFICGEN_PORT2_IP = '192.168.240.10'
```

To run VXLAN encapsulation tests:

```
./vsperf --conf-file user_settings.py --run-integration --test-param 'tunnel_type=vxlan' overlay
```

To run GRE encapsulation tests:

```
./vsperf --conf-file user_settings.py --run-integration --test-param 'tunnel_type=gre' overlay_p
```

To run GENEVE encapsulation tests:

```
./vsperf --conf-file user_settings.py --run-integration --test-param 'tunnel_type=geneve' overla
```

To run OVS NATIVE tunnel tests (VXLAN/GRE/GENEVE):

1. Install the OVS kernel modules

```
cd src/ovs/ovs
sudo -E make modules_install
```

2. Set the following variables:

```
VSWITCH = 'OvsVanilla'
VSWITCH_VANILLA_PHY_PORT_NAMES = ['nic1name', 'nic2name']
# Specify vport_* kernel module to test.
VSWITCH_VANILLA_KERNEL_MODULES = ['vport_vxlan',
                                   'vport_gre',
                                   'vport_geneve',
                                   os.path.join(OVS_DIR_VANILLA, 'datapath/linux/openvswitch.ko')]
```

3. Run tests:

```
./vsperf --conf-file user_settings.py --run-integration --test-param 'tunnel_type=vxlan' overlay
```

2.2 Executing VXLAN decapsulation tests

To run VXLAN decapsulation tests:

1. Set the variables used in “Executing Tunnel encapsulation tests”
2. Set IXNET_TCL_SCRIPT, VXLAN_FRAME_L2, VLXAN_FRAME_L3 and DUT_NIC1_MAC of your settings file to:

```
IXNET_TCL_SCRIPT='ixnetrfc2544v2.tcl'

VXLAN_FRAME_L2 = {'srcmac':
                  '01:02:03:04:05:06',
                  'dstmac':
                  '<DUT's NIC1 MAC>',
                  }

VXLAN_FRAME_L3 = {'proto': 'udp',
                  'packetsize': 64,
                  'srcip': '1.1.1.1',
                  'dstip': '192.168.240.1',
                  }

VXLAN_FRAME_L4 = {'srcport': 4789,
                  'dstport': 4789,
                  'vni': VXLAN_VNI,
                  'inner_srcmac': '01:02:03:04:05:06',
                  'inner_dstmac': '06:05:04:03:02:01',
                  'inner_srcip': '192.168.0.10',
                  'inner_dstip': '192.168.240.9',
                  'inner_proto': 'udp',
                  'inner_srcport': 3000,
                  'inner_dstport': 3001,
                  }

# The receiving NIC of VXLAN traffic
DUT_NIC1_MAC = '<mac address>'
```

3. Run test:

```
./vsperf --conf-file user_settings.py --run-integration overlay_p2p_decap_cont
```

2.3 Executing GRE decapsulation tests

To run GRE decapsulation tests:

1. Set the variables used in “Executing Tunnel encapsulation tests”
2. Set IXNET_TCL_SCRIPT, GRE_FRAME_L2, GRE_FRAME_L3 and DUT_NIC1_MAC of your settings file to:

```
IXNET_TCL_SCRIPT='ixnetrfc2544v2.tcl'

GRE_FRAME_L2 = {'srcmac':
                '01:02:03:04:05:06',
                'dstmac':
                '<DUT's NIC2 MAC>',
                }

GRE_FRAME_L3 = {'proto': 'gre',
                'packetize': 64,
                'srcip': '1.1.1.1',
                'dstip': '192.168.240.1',
                }

GRE_FRAME_L4 = {'srcport': 0,
                'dstport': 0,
                'inner_srcmac': '01:02:03:04:05:06',
                'inner_dstmac': '06:05:04:03:02:01',
                'inner_srcip': '192.168.0.10',
                'inner_dstip': '192.168.240.9',
                'inner_proto': 'udp',
                'inner_srcport': 3000,
                'inner_dstport': 3001,
                }

# The receiving NIC of GRE traffic
DUT_NIC1_MAC = '<mac address>'
```

3. Run test:

```
./vsperf --conf-file user_settings.py --test-param 'tunnel_type=gre' --run-integration overlay_p
```

2.4 Executing GENEVE decapsulation tests

IxNet 7.3X does not have native support of GENEVE protocol. The template, GeneveIxNetTemplate.xml_ClearText.xml, should be imported into IxNET for this testcase to work.

To import the template do: 1. Run the IxNetwork TCL Server 2. Click on the Traffic menu 3. Click on the Traffic actions and click Edit Packet Templates 4. On the Template editor window, click Import.

Select the template tools/pkt_gen/ixnet/GeneveIxNetTemplate.xml_ClearText.xml and click import.

To run GENEVE decapsulation tests:

1. Set the variables used in “Executing Tunnel encapsulation tests”

2. Set IXNET_TCL_SCRIPT, GENEVE_FRAME_L2, GENEVE_FRAME_L3 and DUT_NIC1_MAC of your settings file to:

```
IXNET_TCL_SCRIPT='ixnetrfc2544v2.tcl'

GENEVE_FRAME_L2 = {'srcmac':
                  '01:02:03:04:05:06',
                  'dstmac':
                  '<DUT's NIC2 MAC>',
                  }

GENEVE_FRAME_L3 = {'proto': 'udp',
                  'packetsize': 64,
                  'srcip': '1.1.1.1',
                  'dstip': '192.168.240.1',
                  'geneve_vni': 0,
                  'inner_srcmac': '01:02:03:04:05:06',
                  'inner_dstmac': '06:05:04:03:02:01',
                  'inner_srcip': '192.168.0.10',
                  'inner_dstip': '192.168.240.9',
                  'inner_proto': 'udp',
                  }

GENEVE_FRAME_L4 = {'srcport': 6081,
                  'dstport': 6081,
                  'geneve_vni': 0,
                  'inner_srcmac': '01:02:03:04:05:06',
                  'inner_dstmac': '06:05:04:03:02:01',
                  'inner_srcip': '192.168.0.10',
                  'inner_dstip': '192.168.240.9',
                  'inner_proto': 'udp',
                  'inner_srcport': 3000,
                  'inner_dstport': 3001,
                  }

# The receiving NIC of GENEVE traffic
DUT_NIC1_MAC = '<mac address>'
```

3. Run test:

```
./vsperf --conf-file user_settings.py --test-param 'tunnel_type=geneve' --run-integration overla
```