

APPLICATION NOTE

VSPerf RFC 2544 Throughput

Characterize vSwitch Performance for Telco NFV Use Cases

How to benchmark and characterize virtual switch throughput performance for telco Network Functionality Virtualization (NFV) use cases - under the project VSPerf - using Valkyrie2544.

Contents

Application Note	3
Virtual Switch Performance Benchmarking.....	4
Physical to Physical (Phy2Phy) Scenario	4
VM Loopback (PVP) Scenario	4
Two VM Loopback (PVVP) Scenario	5
Physical Test Setup	6
Three Ways to VSPerf RFC2544 Test with Xena	6
Using Xena from VSPerf Script Framework	7
Step 1: Download the Project.....	7
Step 2: Configuration.....	8
Step 3: Import Xena APIs to Your Project.....	9
Step 4: Execute Test or Build Your Own	9
Using Valkyrie2544 GUI	10
Step 1: Select module and ports.....	10
Step 2: Define Test Topology and Frame Content.....	11
Step 3: Define Test Execution Control.....	12
Step 4: Configure Test Types	13
Step 5: Assign Port Group and Port Peer.....	14
Step 6: Edit Segment Header Profile	15
Step 7: Configure Multi-Stream.....	16
Step 8: Execute Test and Generate Report	17
Using Valkyrie2544 Command	18
Step 1: Locate Your V2544 Configuration File	18
Step 2: Execute Test from Command Line	19
VSPerf Results from Redhat using Xena	20

APPLICATION NOTE

Network Function Virtualization (NFV) is dramatically changing the way we design, build and operate networks, and has gained tremendous attention in the industry with its aim for elastic scaling of capacity, capex predictability, rapid and flexible applications/services deployment. As one of the key components in NFV, virtual switch (vSwitch) connects virtual machines where applications are running and its performance has become essential to the success of NFV deployment. Xena Networks is working closely with Open Platform for NFV (OPNFV) community on the vSwitch Performance (VSPerf) project towards a generic and architecture agnostic vSwitch performance benchmarking and test framework. It will serve as a basis for validating the suitability of different vSwitch implementations in a telco NFV deployment environment.

This application note describes how to benchmark and characterize virtual switch throughput performance for telco network functionality virtualization (NFV) use cases under the project VSPerf¹ with Valkyrie2544.

¹ <https://wiki.opnfv.org/display/vsperf/VSPerf+Home>

VIRTUAL SWITCH PERFORMANCE BENCHMARKING

Valkyrie2544 is a free PC application that makes it easy to create, edit and execute test configuration in accordance with RFC 2544. Valkyrie2544 is automatically installed together with ValkyrieManager and can be used with one or more Xena test chassis. The purpose of using existing RFCs is to establish a fair comparison between the performance of virtual and physical switches.

RFC 2544 is an important standard test suite referenced in the VSPerf test framework. It outlines a benchmarking methodology for network interconnect devices. The methodology results in performance metrics such as throughput, latency, and frame loss ratio.

Three test scenarios are defined by OPNFV VSPerf Characterize vSwitch Performance for Telco NFV Use Cases project (<https://wiki.opnfv.org/display/vsperf/Vsperf+Home>):

- Physical to Physical (Phy2Phy)
- VM Loopback (PVP)
- Two VM Loopback (PVVP)

PHYSICAL TO PHYSICAL (PHY2PHY) SCENARIO

The connection is configured as Physical port → vSwitch → Physical port. The vSwitch runs on the host. The deployment is illustrated in Figure 1.

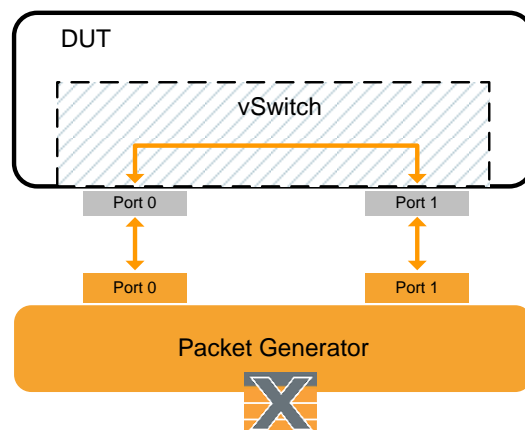


Figure 1. Phy2Phy scenario

VM Loopback (PVP) Scenario

The connection is configured as Physical port → vSwitch → VNF → vSwitch → Physical port. The vSwitch runs on the host. The VM runs DPDK testpmd/L2fwd/Linux Bridge to forward traffic. Traffic consists of uniform UDP packets. The deployment is illustrated in Figure 2.

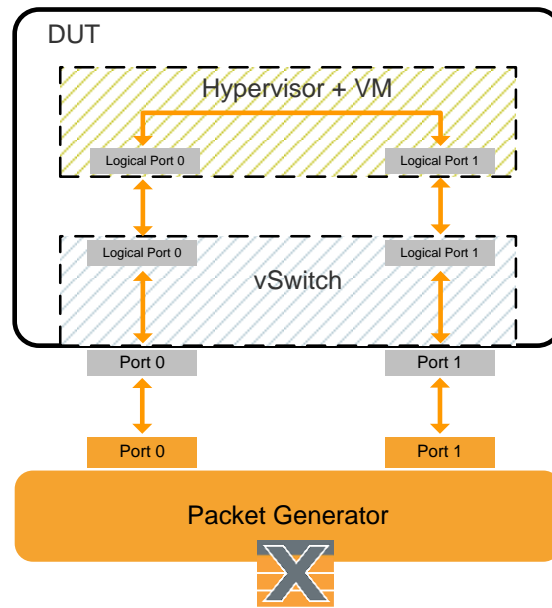


Figure 2. PVP scenario

TWO VM LOOPBACK (PVVP) SCENARIO

The connection is configured as Physical port → vSwitch → VNF → vSwitch → VNF → vSwitch → Physical port. The VM runs DPDK testpmd/L2fwd/Linux Bridge to forward traffic. Traffic consists of uniform UDP packets. The deployment is illustrated in Figure 3.




Figure 3. PVVP scenario

PHYSICAL TEST SETUP

Seen from a tester, the three test scenarios (Phy2Phy, PVP, PVVP) are only different internally from one another. Thus, they can be set up using the test setup shown in Figure 4 below. Two ports from Xena tester are used for testing the device under test (DUT). The DUT in this test is a host that runs vSwitch.




Figure 4. Physical test setup

If you have not installed a virtual switch on your DUT, you can go to <http://openvswitch.org/download/> and follow the instructions to install. It is out of the scope of this application note to demonstrate how to install and configure a virtual switch. You can also find more information about how to deploy Phy2Phy, PVP, and PVVP on <https://wiki.opnfv.org/display/vsperf/Vsperf+Home>.

THREE WAYS TO VSPERF RFC2544 TEST WITH XENA

Benchmarking methodologies on a virtual switch should not have significant difference from a physical switch. From a traffic generator point of view, the test scenarios should be the same when using RFC 2544, so that test results from a virtual switch are comparable with results from a physical one.

There are three ways you can use Valkyrie2544 for testing:

- You can use XenaDriver and XenaTrafficGenerator from VSPerf framework, or
- You can configure and run your tests through Valkyrie2544 GUI, or
- You can run pre-configured tests through Valkyrie2544 command.

They are well illustrated in Figure 5. The following sections will guide you through VSPerf RFC 2544 Throughput test with Xena.

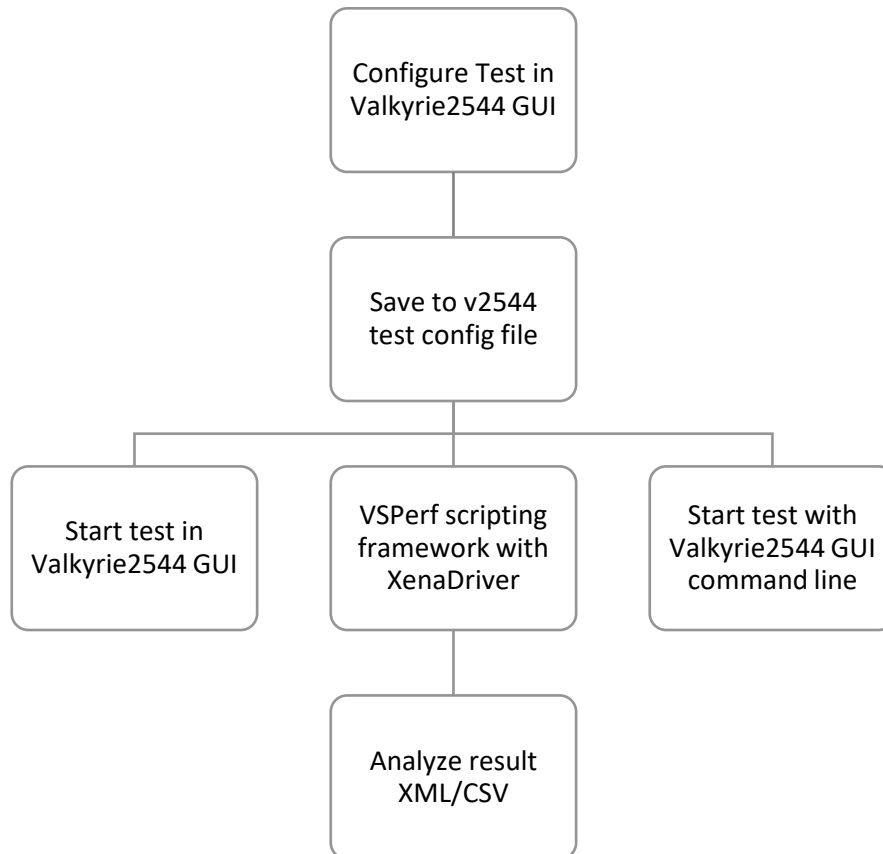


Figure 5. Three ways for VSPerf RFC2544 through test with Xena

USING XENA FROM VSPERF SCRIPT FRAMEWORK

Redhat and Xena are collaborating on Xena integration to VSPerf project for OPNVF. The open source project Xena-VSPerf can be found on:

<https://github.com/ctrautma/Xena-VSPerf>

XenaDriver is developed to provide a set of APIs for users to communicate with the Xena tester. Class Xena (traffic generator) is developed to provide APIs for traffic generating. Both are written in Python. VSPerf design document can be found on:

http://artifacts.opnfv.org/vswitchperf/docs/design/vswitchperf_design.html

Step 1: Download the Project

Go to <https://github.com/ctrautma/Xena-VSPerf> to download the project. Dependencies should also be downloaded and installed on your PC accordingly.

In the directory shown below, there are four important files for VSPerf testing:

1. Valkyrie2544.exe provides RFC 2544 test framework,
2. XenaDriver.py provides communication APIs to Xena testers,
3. xena.py is the traffic generator for starting/stopping test traffic on top of XenaDriver.py and Valkyrie2544.exe
4. xena_json.py converts Valkyrie2544 configuration file .v2544 to JSON format.

```
Xena-VSPerf/tools/pkt_gen/xena
-----profiles/
-----Valkyrie2544.exe
-----XenaDriver.py
-----xena.py
-----xena_json.py
```

In the directory `/profiles`, Valkyrie2544 test configuration files (.v2544) are stored. You can also store your v2544 files there.

An v2544 file is a JSON-formatted file that describes detailed configurations of a Valkyrie2544 test, e.g. which ports are included, topology, traffic pattern, packet header, which test should be carried out, output report format, etc.

An v2544 file is automatically generated by Valkyrie2544 when you save the test through the GUI. In the section *Test Steps – Using Valkyrie2544 GUI*, you can learn how to configure your tests through Valkyrie2544 GUI and generate your own v2544 files. Please note that if you decide to use your own .v2544 files, you need to modify xena.py accordingly.

Step 2: Configuration

A tester configuration file (*.conf) locates in the conf package `Xena-VSPerf/conf/`. You can define the IP address of the Xena tester, e.g. which ports and modules to use, and packet sizes as shown in the example below. You can also enable/disable the use of testers from other vendors in the file.

```
# Xena traffic generator connection info
TRAFFICGEN_XENA_IP = '10.19.15.19'
TRAFFICGEN_XENA_PORT1 = '0'
TRAFFICGEN_XENA_PORT2 = '1'
TRAFFICGEN_XENA_USER = 'vsperf'
TRAFFICGEN_XENA_PASSWORD = 'xena'
TRAFFICGEN_XENA_MODULE1 = '3'
TRAFFICGEN_XENA_MODULE2 = '3'

TEST_PARAMS = {'packet_sizes':'64'}
```

The configuration file also provides a settings object that exposes all of these settings. Settings are not passed from component to component. Rather they are available globally to all components once they import the conf package.

```
from conf import settings
...
log_file = settings.getValue('LOG_FILE_DEFAULT')
```


Step 3: Import Xena APIs to Your Project

To use Xena object (traffic generator), you need to import the modules:

```
from tools.pkt_gen.xena.xena import Xena
from tools.pkt_gen.trafficgen.trafficgenhelper import TRAFFIC_DEFAULTS
```

Xena object provides the following APIs:

```
1: [XENA.send_rfc2544_throughput],
2: [XENA.start_rfc2544_throughput,
    XENA.wait_rfc2544_throughput],
3: [XENA.send_burst_traffic],
4: [XENA.send_cont_traffic],
5: [XENA.start_cont_traffic, XENA_OBJ.stop_cont_traffic],
6: [XENA.send_rfc2544_back2back],
7: [XENA.start_rfc2544_back2back,
    XENA.wait_rfc2544_back2back],
```

Step 4: Execute Test or Build Your Own

You can now try to run the test by executing `xenatest.py` in the root directory. You can also write your own scripts using XenaDrive and Xena APIs.

USING VALKYRIE2544 GUI

The following sections show the steps of how to use Valkyrie2544 to perform VSPerf test with an example configuration. Before we start configuring Xena for RFC 2544 Throughput test, make sure you have installed the virtual switch for testing and have it correctly configured.

Step 1: Select module and ports.

Select the two ports that are connected to the DUT. You can either reserve them now or the application will relinquish other users from the selected ports and reserve them for you when the test starts. Figure 6 shows an example where two ports are selected for the test. It is always good to check that the two ports you connect to the DUT (vSwitch) is shown green. If not, you need to check the connectivity problem.

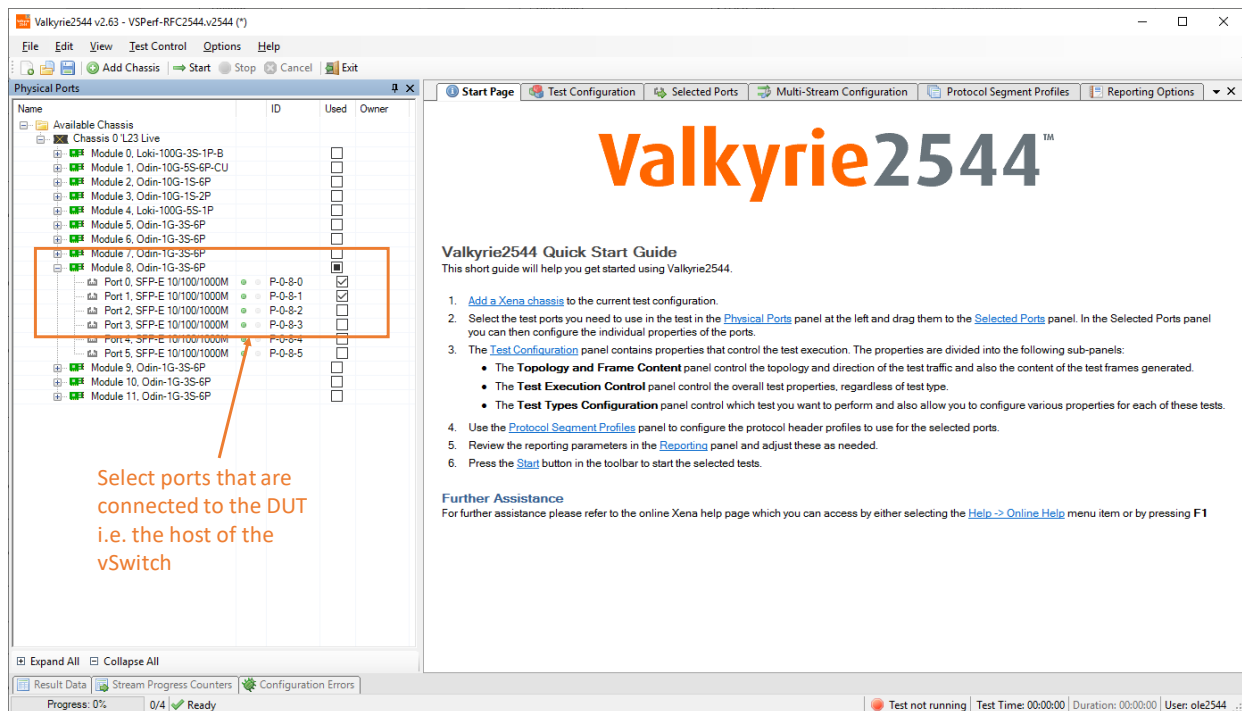


Figure 6. Select test module card and ports.

Step 2: Define Test Topology and Frame Content

As shown in Figure 7 below, on the right side of the window, you can select several panels. Go to panel *Test Configuration*.

The Test Configuration panel contains properties that control the test execution. The properties are divided into the following sub-panels:

- The *Topology and Frame Content* panel controls the topology and direction of the test traffic and the content of the test frames generated.
- The *Test Execution Control* panel controls the overall test properties, regardless of the test type.
- The *Test Types Configuration* panel controls which test you want to perform and allows you to configure various properties for each of these tests.

Select the *Topology and Frame Content* panel. In the *Overall Test Topology* section, select *Topology to Pairs*, and *Direction to Bidirectional* for a bidirectional traffic flow test. Unidirectional traffic flow should also be tested because some problem may not be visible with only testing with bidirectional flows. In addition to traffic directions, the vSwitch should be tested with different packet size because packet size can greatly affect the performance. Set frame size to 64 and 1500 bytes in order to examine the performance of the vSwitch under small and large packet sizes. Valkyrie2544 will execute two tests, one for each packet size.

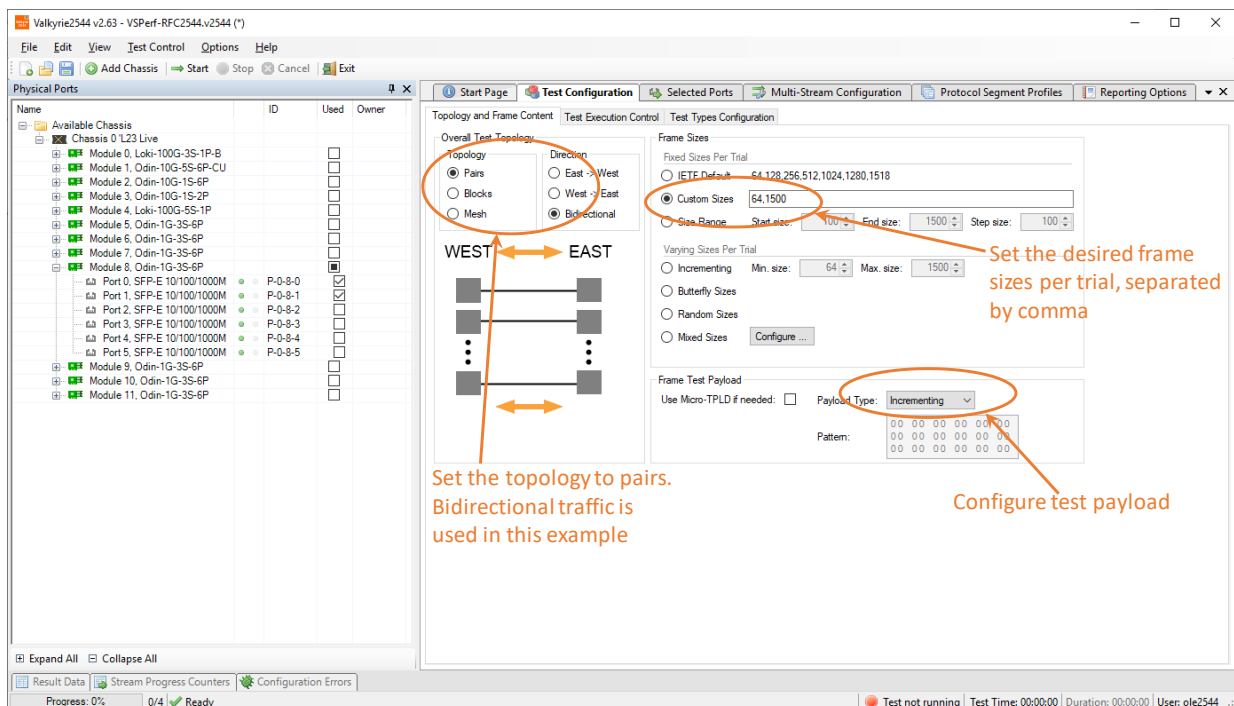


Figure 7. Define test topology and frame content.

Step 3: Define Test Execution Control

Test execution control can be configured in the panel *Test Execution Control* shown in Figure 8 below. Default configuration is used in this example with MAC learning mode on for every trial. Stream-based flow creation allows multiple streams to be generated from one physical port, while modifier-based allows only one. Port sync allows ports to generate traffic into the vSwitch simultaneously. Otherwise, ports will generate traffic individually at its own pace.

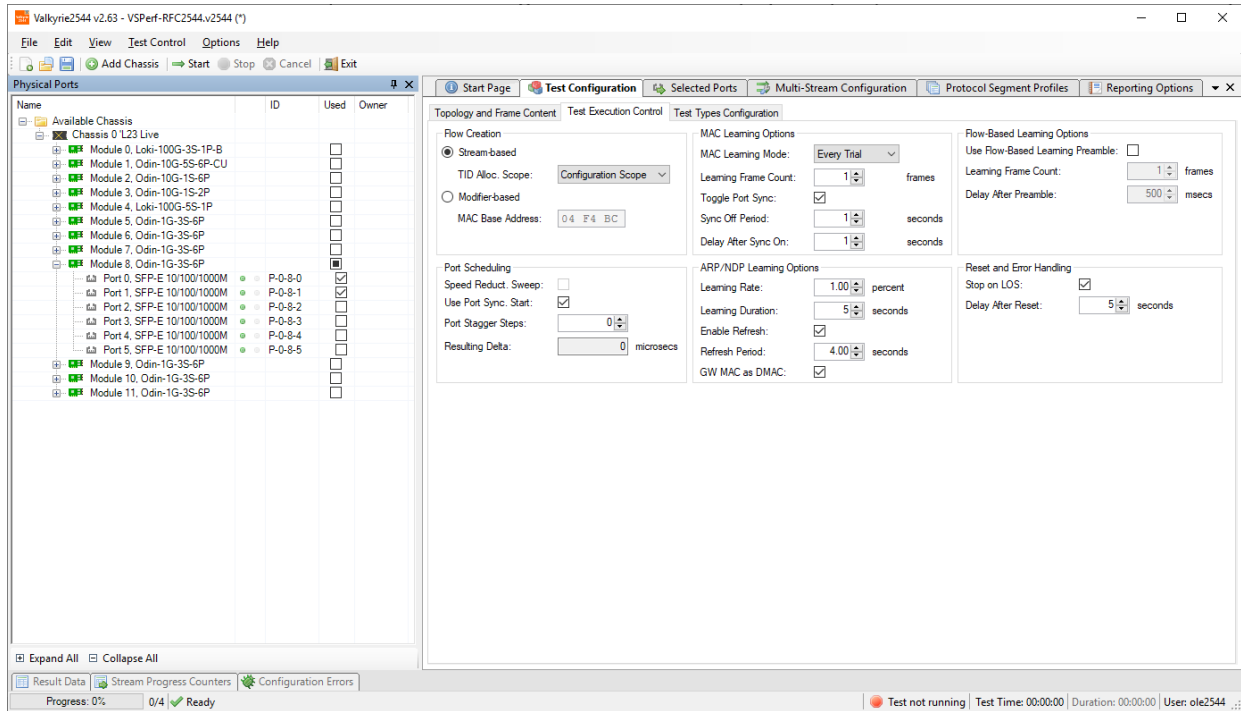


Figure 8. Define test execution control

Step 4: Configure Test Types

RFC 2544 test suite include four tests: throughput, latency and jitter, frame loss rate, and back-to-back. Select the *Throughput Test* for the test. Throughput test measures the maximum rate at which none of the offered frames is dropped by the DUT. The test may be executed using various frame sizes in Step 2.

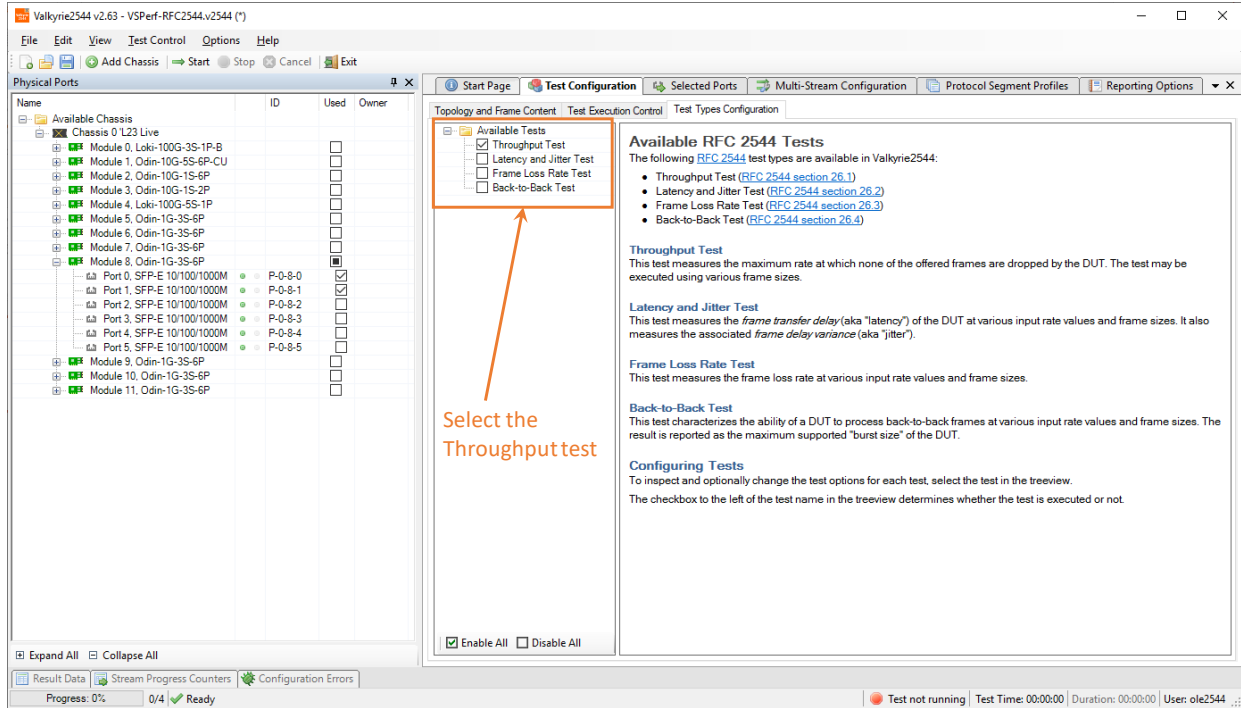


Figure 9. Configure test types to Throughput Test

Step 5: Assign Port Group and Port Peer

As shown in Figure 10, go to *Selected Ports* panel. Assign the first port to *East* group and the other *West*. Then, assign a peer port to each port. In the example shown in Figure 10 below, port 0 is in Group West and its peer, port 1 in group East. After that, choose the *Protocol Segment Profile*. Ethernet header profile is used in the example.

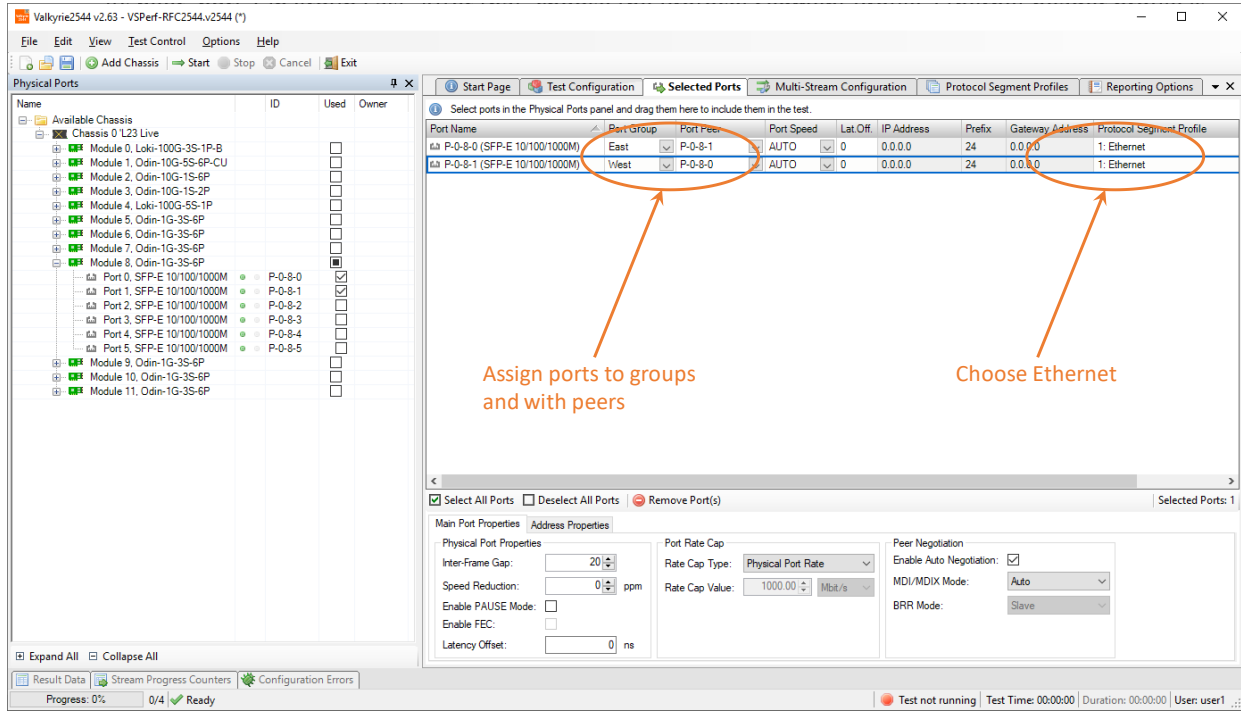


Figure 10. Assign Port Group and Port Peer

Step 6: Edit Segment Header Profile

As shown in Figure 11, go to panel *Protocol Segment Profiles* to check the protocol segment profile you have chosen in Step 5. User count suggests that there are two ports using that profile. You can also edit or create segment header profile.

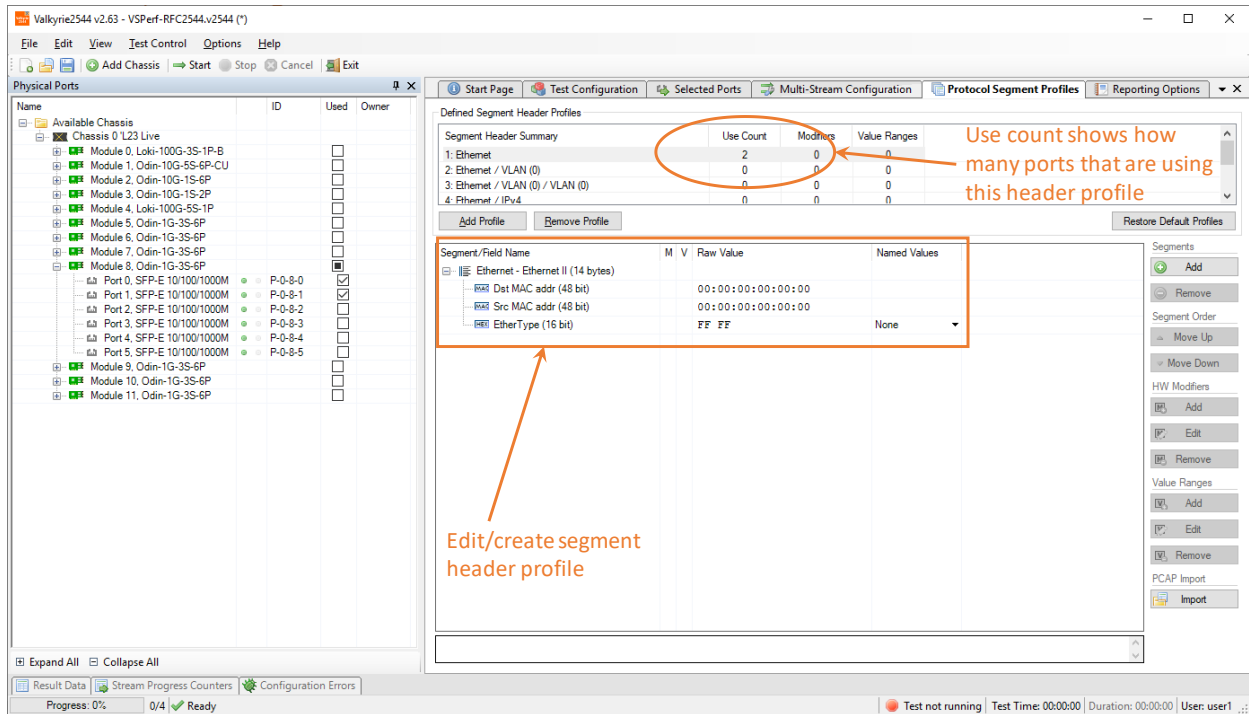


Figure 11. Edit/create segment header profile

Step 7: Configure Multi-Stream

As shown in Figure 12, go to panel *Multi-Stream Configuration*. Enable the Multi-Stream option. Create 32 streams per port. Example is shown in Figure 12. All 32 streams are generated from the same physical port with different source and destination MAC addresses and into the other port, vice versa. Since the traffic is bidirectional, the total amount of streams is 64.

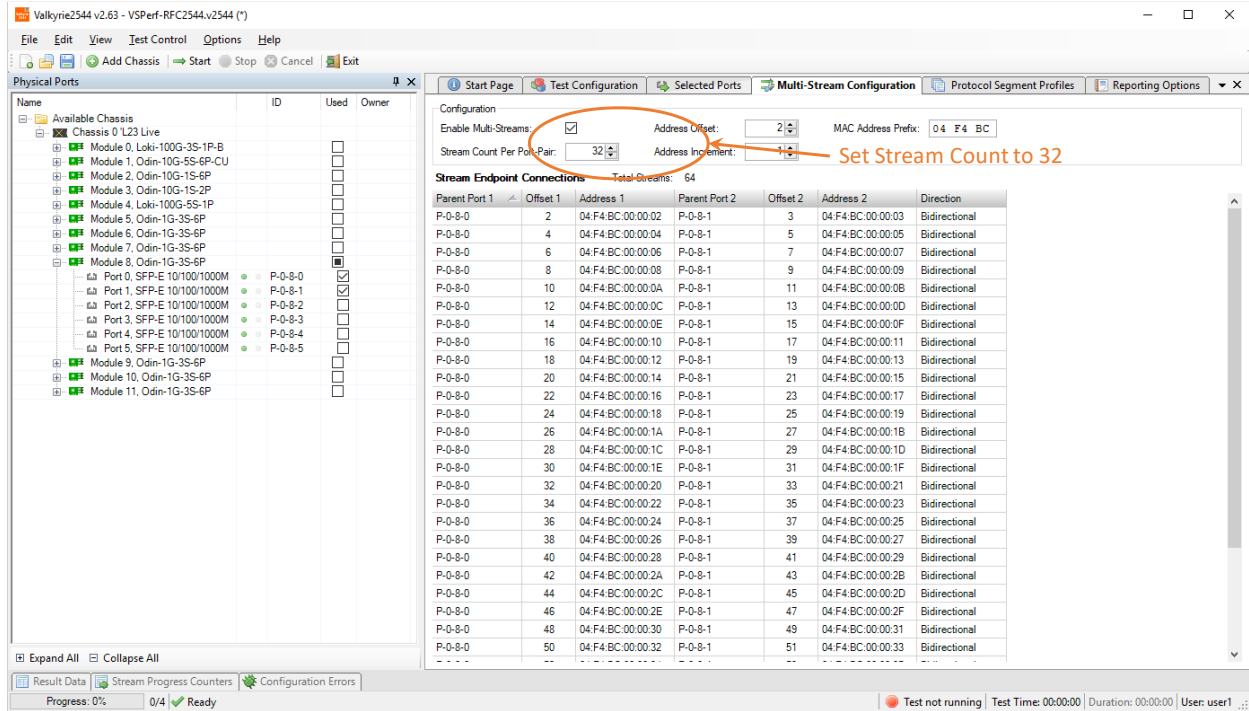


Figure 12. Configure multi-stream

Step 8: Execute Test and Generate Report

Save your configuration. The Valkyrie2544 configuration file has an extension name of v2544. This configuration file is very important if you want to execute your test in an automated fashion, which will be explained and shown in the next section.

As shown in Figure 13, you can select Valkyrie2544 to generate different types of reports, i.e. PDF, XML or CSV. It is very useful to select the XML or CSV formats if you want to automate some post-processes to the test results.

Click the *Start* button, and the test will begin to run. As mentioned previously, if you have not reserved the ports, Valkyrie2544 will reserve them for your test.

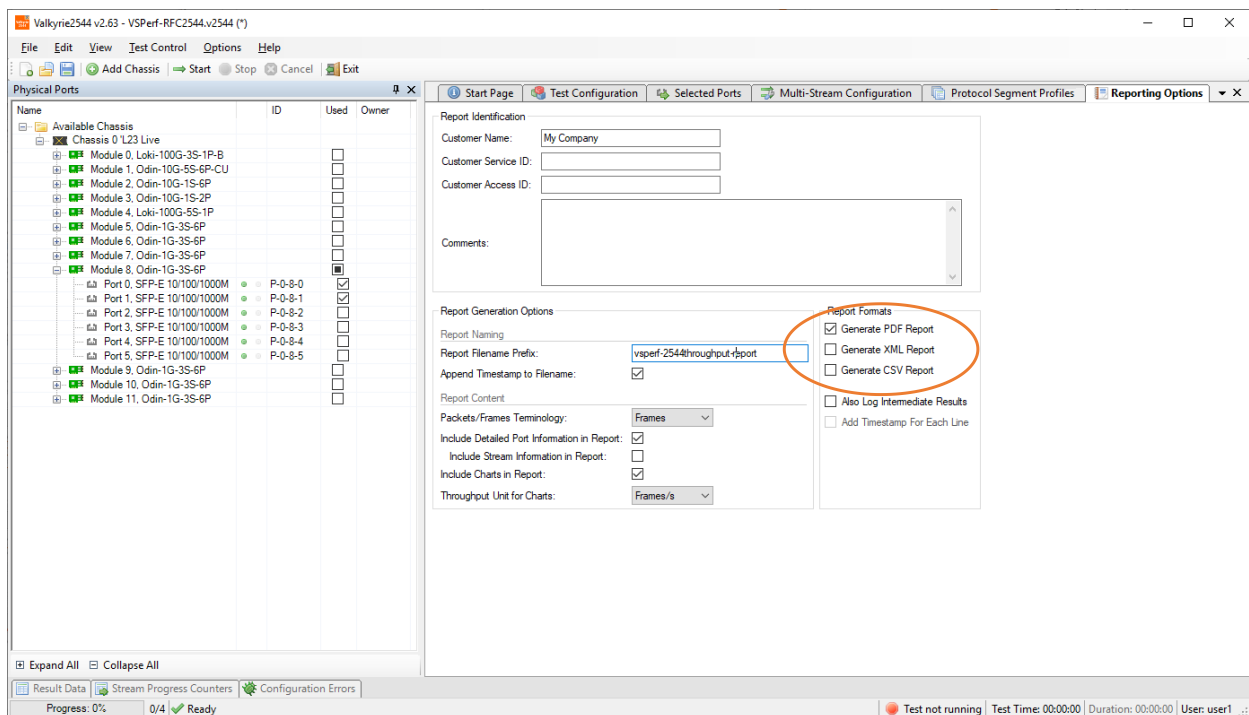


Figure 13. Select Report Formats

USING VALKYRIE2544 COMMAND

In addition to using Valkyrie2544 GUI to configure and run tests, Xena provides a great command line interface to execute pre-configured tests. With this feature, you can automate your test process with your own scripts.

Step 1: Locate Your V2544 Configuration File

Find the v2544 file that you previously saved. You can find it by clicking *Explore Xena Data Directory* as shown in Figure 14 below.

The file path in the example is:

```
C:\Users\me\Documents\Xena\Valkyrie2544-2G\VSPerf-RFC2544.v2544
```

and is referred to as `<CONFIG_FILE>` throughout the following sections.

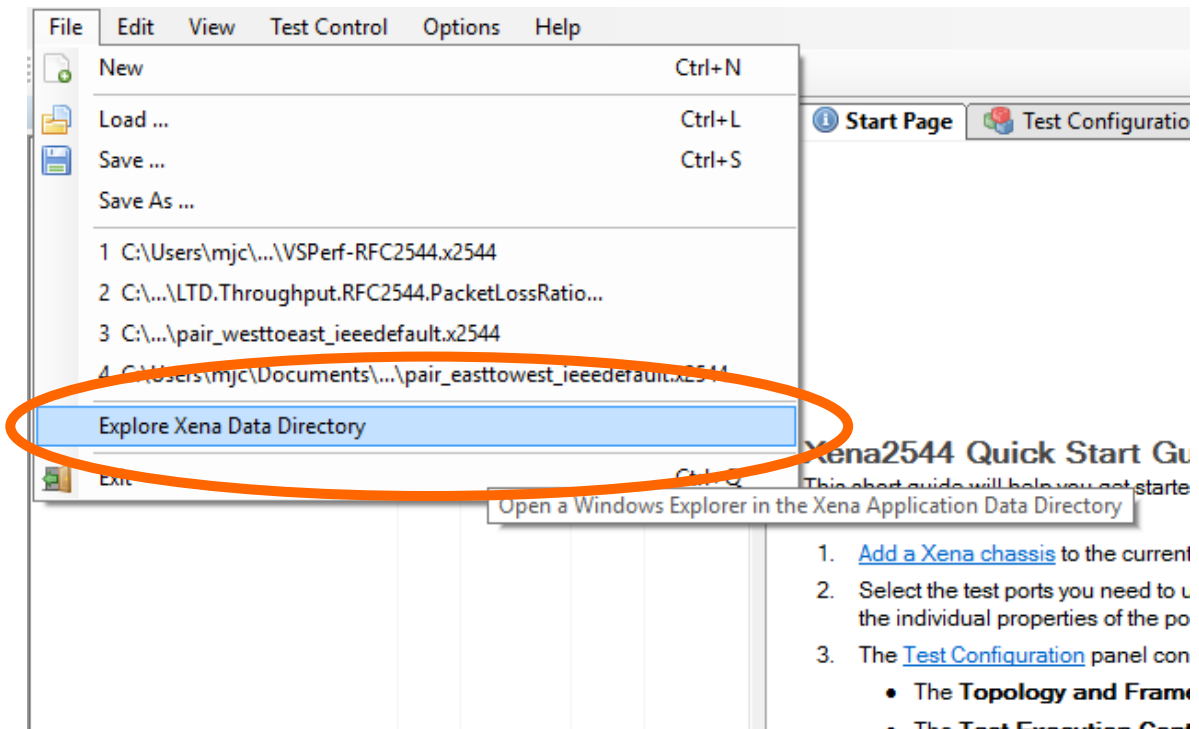
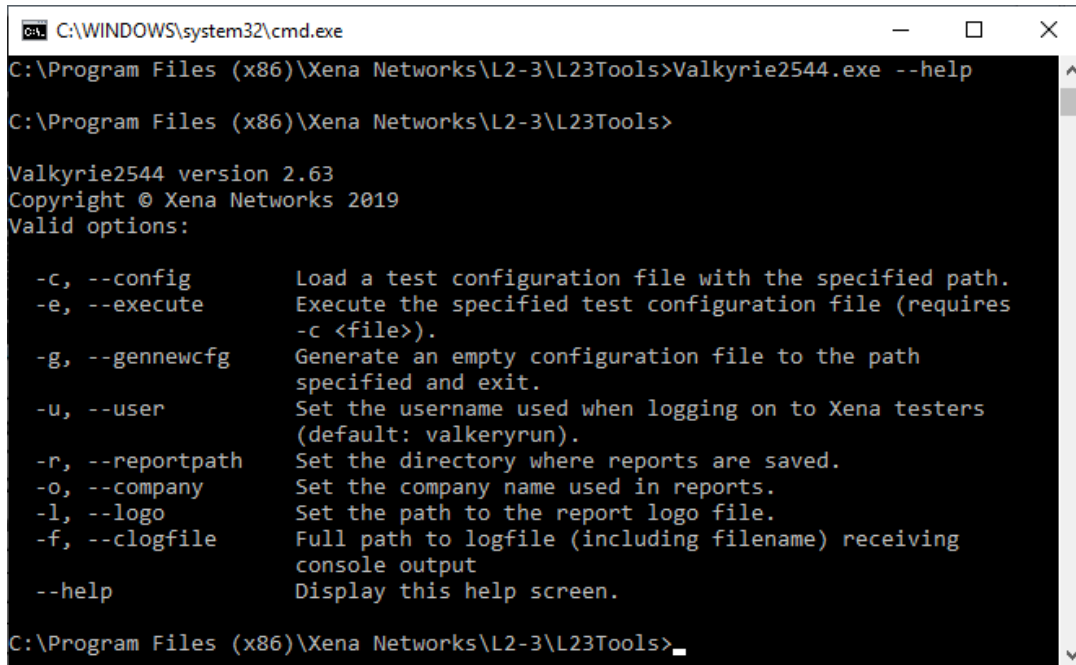


Figure 14. Find your v2544 configuration file

Step 2: Execute Test from Command Line

Go to command directory `C:\Program Files (x86)\Xena Networks\L2-3\L23Tools`, as shown in Figure 15. In this directory, the `Valkyrie2544.exe` is the command we will execute to execute the test configured in the `v2544` file. Several options are available as shown in Figure 15.



```
C:\WINDOWS\system32\cmd.exe
C:\Program Files (x86)\Xena Networks\L2-3\L23Tools>Valkyrie2544.exe --help
C:\Program Files (x86)\Xena Networks\L2-3\L23Tools>
Valkyrie2544 version 2.63
Copyright © Xena Networks 2019
Valid options:
-c, --config          Load a test configuration file with the specified path.
-e, --execute         Execute the specified test configuration file (requires
-c <file>).
-g, --gennewcfg       Generate an empty configuration file to the path
specified and exit.
-u, --user            Set the username used when logging on to Xena testers
(default: valkeryrun).
-r, --reportpath      Set the directory where reports are saved.
-o, --company         Set the company name used in reports.
-l, --logo            Set the path to the report logo file.
-f, --clogfile        Full path to logfile (including filename) receiving
console output
--help               Display this help screen.
C:\Program Files (x86)\Xena Networks\L2-3\L23Tools>
```

Figure 15. Find your Valkyrie2544 command directory

Execute the test by running the command in the current directory, shown in Figure 16.

```
Valkyrie2544.exe -c <CONFIG_FILE> -e or
Valkyrie2544.exe -config <CONFIG_FILE> -execute
```

```
INFO: Activating chassis: 192.168.1.170
INFO: Logon to 192.168.1.170:22606 successful
INFO: All chassis connected and read
INFO: Starting subtest 'Throughput Test'
INFO: Reserving ports
INFO: Resetting used ports
INFO: Port post-reset delay
INFO: Configuring ports
INFO: Toggling port sync state
INFO: InitTrial: 64 byte packets, iteration 1
INFO: Setting up streams
INFO: Sending L2 learning data stream preamble
INFO: Rate iteration: 100.000% traffic rate
INFO: Traffic on source port(s) started
INFO: Traffic on source port(s) stopped
INFO: Sending L2 learning data stream preamble
INFO: Rate iteration: 50.050% traffic rate
INFO: Traffic on source port(s) started
INFO: Traffic on source port(s) stopped
INFO: Sending L2 learning data stream preamble
INFO: Rate iteration: 25.075% traffic rate
INFO: Traffic on source port(s) started
INFO: Traffic on source port(s) stopped
INFO: Sending L2 learning data stream preamble
INFO: Rate iteration: 37.563% traffic rate
INFO: Traffic on source port(s) started
INFO: Traffic on source port(s) stopped
INFO: Sending L2 learning data stream preamble
INFO: Rate iteration: 43.806% traffic rate
INFO: Traffic on source port(s) started
INFO: Traffic on source port(s) stopped
INFO: Sending L2 learning data stream preamble
INFO: Rate iteration: 46.928% traffic rate
INFO: Traffic on source port(s) started
INFO: Traffic on source port(s) stopped
INFO: Sending L2 learning data stream preamble
INFO: Rate iteration: 48.489% traffic rate
INFO: Traffic on source port(s) started
INFO: Traffic on source port(s) stopped
INFO: Sending L2 learning data stream preamble
```

Figure 16. Executing test with Valkyrie2544 command

VSPERF RESULTS FROM REDHAT USING XENA

As one of the members of OPNFV, Redhat is using Xena for VSPerf. Below is a test report of PHY2PHY_TPUT. To understand different test IDs used by VSPerf, please read Xena's white paper Virtual Switch Performance Benchmarking.

```
Test ID: PHY2PHY_TPUT
-----

Test Environment
~~~~~

Below is the environment that the test was performed in:

* OS: redhat 7.3 Maipo
* Kernel Version: 3.10.0-470.el7.x86_64
* NIC(s):
  * Intel Corporation Ethernet 10G 2P X520 Adapter (rev 01)
  * Intel Corporation Ethernet 10G 2P X520 Adapter (rev 01)
* Board: Dell Inc. 0599V5 [2 sockets]
* CPU: Intel(R) Xeon(R) CPU E5-2637 v3 @ 3.50GHz
* CPU cores: 16
* Memory: 65759124 kB
* Virtual Switch Set-up: p2p
* vswitchperf: GIT tag: 07bab38905cf10cad6987bb9c5b302bbff2c6013
* Traffic Generator: Xena, Version: Valkyrie2544 v2.43, GIT tag: None
```

```
* vSwitch: OvsDpdkVhost, Version: None, GIT tag: None
* DPDK Version: 2.2.0, GIT tag: None
```

Below are test details:

```
* Test ID: phy2phy_tput
* Description: LTD.Throughput.RFC2544.PacketLossRatio
* Deployment: p2p
* Traffic type: rfc2544
* Bidirectional : True
```

Test results for packet size: 64

~~~~~  
A detailed summary of the main results is outlined below.

Results/Metrics Collected  
^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^

The following are the metrics obtained during this test:

| Metric                | Result             |
|-----------------------|--------------------|
| throughput_rx_fps     | 23115618.5733333   |
| throughput_rx_mbps    | 15533.695681279998 |
| throughput_rx_percent | 77.1484375         |
| tx_rate_fps           | 23115618.5733333   |
| tx_rate_mbps          | 15533.69568128     |
| tx_rate_percent       | 77.1484375         |
| min_latency_ns        | 4487.0             |
| max_latency_ns        | 51071.0            |
| avg_latency_ns        | 8352.0             |
| type                  | rfc2544            |
| packet_size           | 64                 |
| traffic_type          | udp                |
| test_execution_time   | 01:33:30           |

Statistics collected  
^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^

The following system statistics were collected during testcase execution:

```
Process: ovs-vswitchd
```

| Statistic | Value  |
|-----------|--------|
| UID       | 0      |
| PID       | 6870   |
| %usr      | 100.00 |
| %system   | 2.01   |
| %guest    | 0.00   |
| %CPU      | 100.00 |
| CPU       | -      |
| minflt/s  | 25.24  |

```

majflt/s                0.00
VSZ                    3292824
RSS                   30509
%MEM                  0.05
kB_rd/s              0.00
kB_wr/s              0.00
kB_ccwr/s            0.00
=====

```

Process: ovsdb-server

| Statistic | Value |
|-----------|-------|
| UID       | 0     |
| PID       | 6863  |
| %usr      | 0.01  |
| %system   | 0.00  |
| %guest    | 0.00  |
| %CPU      | 0.01  |
| CPU       | -     |
| minflt/s  | 0.00  |
| majflt/s  | 0.00  |
| VSZ       | 43728 |
| RSS       | 3260  |
| %MEM      | 0.00  |
| kB_rd/s   | 0.00  |
| kB_wr/s   | 0.00  |
| kB_ccwr/s | 0.00  |

Test results for packet size: 1500

~~~~~  
A detailed summary of the main results is outlined below.

Results/Metrics Collected
^^^^^^^^^^^^^^^^^^^^

The following are the metrics obtained during this test:

Metric	Result
throughput_rx_fps	1644717.9500000002
throughput_rx_mbps	19999.770272
throughput_rx_percent	99.99890255943811
tx_rate_fps	1644736
tx_rate_mbps	19999.98976
tx_rate_percent	100
min_latency_ns	8135.0
max_latency_ns	2584233.0
avg_latency_ns	2018627.0
type	rfc2544
packet_size	1500

```
traffic_type                udp
test_execution_time        01:33:30
=====
```

```
Statistics collected
^^^^^^^^^^^^^^^^^^^^
```

The following system statistics were collected during testcase execution:

```
=====
```

Process: ovs-vswitchd

```
-----
```

Statistic	Value
UID	0
PID	6870
%usr	100.00
%system	2.01
%guest	0.00
%CPU	100.00
CPU	-
minflt/s	25.24
majflt/s	0.00
VSZ	3292824
RSS	30509
%MEM	0.05
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

```
=====
```

```
=====
```

Process: ovssdb-server

```
-----
```

Statistic	Value
UID	0
PID	6863
%usr	0.01
%system	0.00
%guest	0.00
%CPU	0.01
CPU	-
minflt/s	0.00
majflt/s	0.00
VSZ	43728
RSS	3260
%MEM	0.00
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

```
=====
```

Anomalies

~~~~~

No anomalies were detected during the course of this test.

Testing Activities/Events

~~~~~

pidstat is used to collect the process statistics, as such some values such as %CPU and %USER maybe > 100% as the values are summed across multiple cores. For more info on pidstat please see: <http://linux.die.net/man/1/pidstat>.

Known issues: Some reported metrics have the value "unkown". These values are marked unknown as they are not values retrieved from the external tester (traffic generator). They were incorrectly derived in a way that made assumptions about packet sizes, as such they have been deprecated from vsperf and marked as unknown. They will be resolved in the next release.
.. There must be blank lines around to ensure correct formatting.