CHARACTERIZE VSWITCH PERFORMANCE FOR TELCO NFV USE CASES LEVEL TEST REPORT

Release draft (18ff286)

OPNFV

February 14, 2016

CHAPTER

ONE

PERFORMANCE REPORT FOR OPEN VSWITCH WITH DPDK SUPPORT

1.1 Introduction

The objective of the OPNFV project titled **"Characterise vSwitch Performance for Telco NFV Use Cases"**, is to evaluate a virtual switch to identify its suitability for a Telco Network Function Virtualization (NFV) environment. As well as this, the project aims to identify any gaps or bottlenecks in order to drive architectural changes to improve virtual switch performance and determinism. The purpose of this document is to summarize the results of the tests carried out on the virtual switch in the Network Function Virtualization Infrastructure (NFVI) and, from these results, provide evaluations and recommendations for the virtual switch. Test results will be outlined in *details-of-LTR*, preceded by the *document-identifier* and the *scope* and *references*).

This document is currently in draft form.

1.1.1 Document identifier

The document id will be used to uniquely identify versions of the LTR. The format for the document id will be: OPNFV_vswitchperf_LTR_rel_STATUS, the status is one of: DRAFT, REVIEWED, CORRECTED or FINAL. The document id for this version of the LTR is: OPNFV_vswitchperf_LTR_Brahmaputra_DRAFT.

1.1.2 Scope

The scope of this report is to detail the results of the tests that have been performed on the virtual switch. This report will also evaluate the results of these tests and, based on these evaluations, provide recommendations on the suitability of the virtual switch for use in a Telco NFV environment.

1.1.3 References

OPNFV_vswitchperf_LTD_Brahmaputra_REVIEWED

1.2 Details of the Level Test Report

This section provides a *test-results-overview*. Also included are the *rationale* and the *conclusions*.

1.2.1 Test ID: PHY2PHY_TPUT

Test Environment

Below is the environment that the test was performed in:

- OS: centos 7.2.1511 Core
- Kernel Version: 3.10.0-229.14.1.el7.x86_64
- **NIC(s):**
  - Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01)
  - Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01)
- Board: Intel Corporation S2600JF [2 sockets]
- CPU: Intel(R) Xeon(R) CPU E5-2690 0 @ 2.90GHz
- CPU cores: 32
- Memory: 32727096 kB
- Virtual Switch Set-up: p2p
- vswitchperf: GIT tag: 18ff28682b0c2a000af61c628046749ab7a1434e
- Traffic Generator: IxNet, Version: 7.50.1009.20, GIT tag: None
- vSwitch: OvsDpdkVhost, Version: 2.5.90, GIT tag: 02ab4b1a6a173979a51cabd7000a34546d517e60
- DPDK Version: 2.2.0, GIT tag: a38e5ec15e3fe615b94f3cc5edca5974dab325ab

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
tx_rate_fps	Unknown
throughput_rx_fps	22626044.857
tx_rate_mbps	Unknown
throughput_rx_mbps	11584.535
tx_rate_percent	76.023
throughput_rx_percent	76.023
min_latency_ns	5880.000
max_latency_ns	342500.000
avg_latency_ns	7478.500
type	rfc2544
packet_size	64
traffic_type	udp

Statistics collected

The following system statistics were collected during testcase execution:

Process: ovs-vswitchd	
Statistic	Value
UID	0
PID	62307
%usr	200.05
%system	0.03
%guest	0.00
%CPU	200.07
CPU	•
minflt/s	46.50
majflt/s	0.00
VSZ	2668616
RSS	20015
%MEM	0.06
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

Process: ovsdb-server	
Statistic	Value
UID	0
PID	62306
%usr	0.00
%system	0.00
%guest	0.00
%CPU	0.00
CPU	•
minflt/s	0.00
majflt/s	0.00
VSZ	47632
RSS	3344
%MEM	0.01
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

Test results for packet size: 128

A detailed summary of the main results is outlined below.

Results/Metrics Collected

The following are the metrics obtained during this test:

Metric	Result
tx_rate_fps	Unknown
throughput_rx_fps	16238705.813
tx_rate_mbps	Unknown
throughput_rx_mbps	16628.435
tx_rate_percent	96.133
throughput_rx_percent	96.133
min_latency_ns	5340.000
max_latency_ns	140940.000
avg_latency_ns	12091.000
type	rfc2544
packet_size	128
traffic_type	udp

Statistics collected

The following system statistics were collected during testcase execution:

Process: ovs-vswitchd	
Statistic	Value
UID	0
PID	62307
%usr	200.05
%system	0.03
%guest	0.00
%CPU	200.07
CPU	•
minflt/s	46.50
majflt/s	0.00
VSZ	2668616
RSS	20015
%MEM	0.06
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

Process: ovsdb-server	
Statistic	Value
UID	0
PID	62306
%usr	0.00
%system	0.00
%guest	0.00
%CPU	0.00
CPU	•
minflt/s	0.00
majflt/s	0.00
VSZ	47632
RSS	3344
%MEM	0.01
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.

1.2.2 Test ID: PHY2PHY_CONT

Test Environment

Below is the environment that the test was performed in:

- OS: centos 7.2.1511 Core
- Kernel Version: 3.10.0-229.14.1.el7.x86_64
- **NIC(s):**
  - Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01)
  - Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01)
- Board: Intel Corporation S2600JF [2 sockets]
- CPU: Intel(R) Xeon(R) CPU E5-2690 0 @ 2.90GHz
- CPU cores: 32
- Memory: 32727096 kB
- Virtual Switch Set-up: p2p
- vswitchperf: GIT tag: 18ff28682b0c2a000af61c628046749ab7a1434e
- Traffic Generator: IxNet, Version: 7.50.1009.20, GIT tag: None
- vSwitch: OvsDpdkVhost, Version: 2.5.90, GIT tag: 02ab4b1a6a173979a51cabd7000a34546d517e60
- DPDK Version: 2.2.0, GIT tag: a38e5ec15e3fe615b94f3cc5edca5974dab325ab

Below are test details:

- Test ID: phy2phy_cont
- Description: Phy2Phy Continuous Stream
- 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
tx_rate_fps	Unknown
throughput_rx_fps	23804916.941
tx_rate_mbps	Unknown
throughput_rx_mbps	12188.117
tx_rate_percent	100
throughput_rx_percent	79.985
min_latency_ns	21100.000
max_latency_ns	920080.000
avg_latency_ns	463522.500
type	rfc2544
packet_size	64
traffic_type	udp

Statistics collected

The following system statistics were collected during testcase execution:

Process: ovs-vswitchd	
Statistic	Value
UID	0
PID	63387
%usr	200.04
%system	0.03
%guest	0.00
%CPU	200.07
CPU	•
minflt/s	109.00
majflt/s	0.00
VSZ	2668616
RSS	22012
%MEM	0.07
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

Process: ovsdb-server	
Statistic	Value
UID	0
PID	63386
%usr	0.00
%system	0.00
%guest	0.00
%CPU	0.00
CPU	•
minflt/s	0.00
majflt/s	0.00
VSZ	47632
RSS	3344
%MEM	0.01
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

Test results for packet size: 128

A detailed summary of the main results is outlined below.

Results/Metrics Collected

The following are the metrics obtained during this test:

Metric	Result
tx_rate_fps	Unknown
throughput_rx_fps	16536984.557
tx_rate_mbps	Unknown
throughput_rx_mbps	16933.872
tx_rate_percent	100
throughput_rx_percent	97.899
min_latency_ns	20840.000
max_latency_ns	646360.000
avg_latency_ns	252419.500
type	rfc2544
packet_size	128
traffic_type	udp

Statistics collected

The following system statistics were collected during testcase execution:

Process: ovs-vswitchd	
Statistic	Value
UID	0
PID	63387
%usr	200.04
%system	0.03
%guest	0.00
%CPU	200.07
CPU	•
minflt/s	109.00
majflt/s	0.00
VSZ	2668616
RSS	22012
%MEM	0.07
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

Process: ovsdb-server	
Statistic	Value
UID	0
PID	63386
%usr	0.00
%system	0.00
%guest	0.00
%CPU	0.00
CPU	•
minflt/s	0.00
majflt/s	0.00
VSZ	47632
RSS	3344
%MEM	0.01
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.

1.3 Rationale for decisions

The tests conducted do not have pass/fail/conditional-pass criteria. The test is simply conducted and the results are reported.

1.4 Conclusions and recommendations

The test results are stable. The vsperf CI jobs that were used to obtain the results can be found at https://artifactory.opnfv.org/logs/vswitchperf/intel-pod3/2016-02-12_15-28-58.

1.5 General

1.5.1 Glossary

- NFV - Network Function Virtualization
- Mbps - 1,000,000bps

1.5.2 Document change procedures and history

Document ID	Author	Date Modified
OPNFV_vswitchperf_LTR_ver_1.0_Jan_15_CN_DRAFT	Christopher Nolan	23/01/2015
OPNFV_vswitchperf_LTR_ver_1.1_Jan_15_CN_DRAFT	Christopher Nolan	28/01/2015

CHAPTER

TWO

PERFORMANCE REPORT FOR OPEN VSWITCH

2.1 Introduction

The objective of the OPNFV project titled **"Characterise vSwitch Performance for Telco NFV Use Cases"**, is to evaluate a virtual switch to identify its suitability for a Telco Network Function Virtualization (NFV) environment. As well as this, the project aims to identify any gaps or bottlenecks in order to drive architectural changes to improve virtual switch performance and determinism. The purpose of this document is to summarize the results of the tests carried out on the virtual switch in the Network Function Virtualization Infrastructure (NFVI) and, from these results, provide evaluations and recommendations for the virtual switch. Test results will be outlined in *details-of-LTR*, preceded by the *document-identifier* and the *scope* and *references*).

This document is currently in draft form.

2.1.1 Document identifier

The document id will be used to uniquely identify versions of the LTR. The format for the document id will be: OPNFV_vswitchperf_LTR_rel_STATUS, the status is one of: DRAFT, REVIEWED, CORRECTED or FINAL. The document id for this version of the LTR is: OPNFV_vswitchperf_LTR_Brahmaputra_DRAFT.

2.1.2 Scope

The scope of this report is to detail the results of the tests that have been performed on the virtual switch. This report will also evaluate the results of these tests and, based on these evaluations, provide recommendations on the suitability of the virtual switch for use in a Telco NFV environment.

2.1.3 References

OPNFV_vswitchperf_LTD_Brahmaputra_REVIEWED

2.2 Details of the Level Test Report

This section provides a *test-results-overview*. Also included are the *rationale* and the *conclusions*.

2.2.1 Test ID: PHY2PHY_TPUT

Test Environment

Below is the environment that the test was performed in:

- OS: centos 7.2.1511 Core
- Kernel Version: 3.10.0-229.14.1.el7.x86_64
- **NIC(s):**
  - Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01)
  - Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01)
- Board: Intel Corporation S2600JF [2 sockets]
- CPU: Intel(R) Xeon(R) CPU E5-2690 0 @ 2.90GHz
- CPU cores: 32
- Memory: 32727096 kB
- Virtual Switch Set-up: p2p
- vswitchperf: GIT tag: 18ff28682b0c2a000af61c628046749ab7a1434e
- Traffic Generator: IxNet, Version: 7.50.1009.20, GIT tag: None
- vSwitch: OvsVanilla, Version: 2.5.90, GIT tag: 02ab4b1a6a173979a51cabd7000a34546d517e60

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
tx_rate_fps	Unknown
throughput_rx_fps	527810.239
tx_rate_mbps	Unknown
throughput_rx_mbps	270.239
tx_rate_percent	1.773
throughput_rx_percent	1.773
min_latency_ns	3900.000
max_latency_ns	1198980.000
avg_latency_ns	76101.000
type	rfc2544
packet_size	64
traffic_type	udp

Statistics collected

The following system statistics were collected during testcase execution:

Process: ovs-vswitchd	
Statistic	Value
UID	0
PID	64022
%usr	0.17
%system	0.33
%guest	0.00
%CPU	0.49
CPU	•
minflt/s	0.14
majflt/s	0.00
VSZ	1300916
RSS	7587
%MEM	0.02
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

Process: ovsdb-server	
Statistic	Value
UID	0
PID	64021
%usr	0.00
%system	0.00
%guest	0.00
%CPU	0.01
CPU	•
minflt/s	0.01
majflt/s	0.00
VSZ	45760
RSS	2996
%MEM	0.01
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

Test results for packet size: 128

A detailed summary of the main results is outlined below.

Results/Metrics Collected

The following are the metrics obtained during this test:

Metric	Result
tx_rate_fps	Unknown
throughput_rx_fps	691514.294
tx_rate_mbps	Unknown
throughput_rx_mbps	708.111
tx_rate_percent	4.094
throughput_rx_percent	4.094
min_latency_ns	4020.000
max_latency_ns	9305120.000
avg_latency_ns	2748362.500
type	rfc2544
packet_size	128
traffic_type	udp

Statistics collected

The following system statistics were collected during testcase execution:

Process: ovs-vswitchd	
Statistic	Value
UID	0
PID	64022
%usr	0.17
%system	0.33
%guest	0.00
%CPU	0.49
CPU	•
minflt/s	0.14
majflt/s	0.00
VSZ	1300916
RSS	7587
%MEM	0.02
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

Process: ovsdb-server	
Statistic	Value
UID	0
PID	64021
%usr	0.00
%system	0.00
%guest	0.00
%CPU	0.01
CPU	•
minflt/s	0.01
majflt/s	0.00
VSZ	45760
RSS	2996
%MEM	0.01
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.

2.2.2 Test ID: PHY2PHY_CONT

Test Environment

Below is the environment that the test was performed in:

- OS: centos 7.2.1511 Core
- Kernel Version: 3.10.0-229.14.1.el7.x86_64
- **NIC(s):**
  - Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01)
  - Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01)
- Board: Intel Corporation S2600JF [2 sockets]
- CPU: Intel(R) Xeon(R) CPU E5-2690 0 @ 2.90GHz
- CPU cores: 32
- Memory: 32727096 kB
- Virtual Switch Set-up: p2p
- vswitchperf: GIT tag: 18ff28682b0c2a000af61c628046749ab7a1434e
- Traffic Generator: IxNet, Version: 7.50.1009.20, GIT tag: None
- vSwitch: OvsVanilla, Version: 2.5.90, GIT tag: 02ab4b1a6a173979a51cabd7000a34546d517e60

Below are test details:

- Test ID: phy2phy_cont
- Description: Phy2Phy Continuous Stream
- 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
tx_rate_fps	Unknown
throughput_rx_fps	1184202.403
tx_rate_mbps	Unknown
throughput_rx_mbps	606.312
tx_rate_percent	100
throughput_rx_percent	3.979
min_latency_ns	8806160.000
max_latency_ns	62921420.000
avg_latency_ns	8931913.500
type	rfc2544
packet_size	64
traffic_type	udp

Statistics collected

The following system statistics were collected during testcase execution:

Process: ovs-vswitchd	
Statistic	Value
UID	0
PID	64881
%usr	0.19
%system	0.35
%guest	0.00
%CPU	0.54
CPU	•
minflt/s	0.65
majflt/s	0.00
VSZ	1300912
RSS	6529
%MEM	0.02
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

Process: ovsdb-server	
Statistic	Value
UID	0
PID	64880
%usr	0.00
%system	0.00
%guest	0.00
%CPU	0.00
CPU	•
minflt/s	0.03
majflt/s	0.00
VSZ	45760
RSS	2996
%MEM	0.01
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

Test results for packet size: 128

A detailed summary of the main results is outlined below.

Results/Metrics Collected

The following are the metrics obtained during this test:

Metric	Result
tx_rate_fps	Unknown
throughput_rx_fps	1186633.007
tx_rate_mbps	Unknown
throughput_rx_mbps	1215.112
tx_rate_percent	100
throughput_rx_percent	7.025
min_latency_ns	5588200.000
max_latency_ns	62821820.000
avg_latency_ns	5691999.000
type	rfc2544
packet_size	128
traffic_type	udp

Statistics collected

The following system statistics were collected during testcase execution:

Process: ovs-vswitchd	
Statistic	Value
UID	0
PID	64881
%usr	0.19
%system	0.35
%guest	0.00
%CPU	0.54
CPU	•
minflt/s	0.65
majflt/s	0.00
VSZ	1300912
RSS	6529
%MEM	0.02
kB_rd/s	0.00
kB_wr/s	0.00
kB_ccwr/s	0.00

Process: ovsdb-server	
Statistic	Value
UID	0
PID	64880
%usr	0.00
%system	0.00
%guest	0.00
%CPU	0.00
CPU	•
minflt/s	0.03
majflt/s	0.00
VSZ	45760
RSS	2996
%MEM	0.01
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.

2.3 Rationale for decisions

The tests conducted do not have pass/fail/conditional-pass criteria. The test is simply conducted and the results are reported.

2.4 Conclusions and recommendations

The test results are stable. The vsperf CI jobs that were used to obtain the results can be found at https://artifactory.opnfv.org/logs/vswitchperf/intel-pod3/2016-02-12_15-28-58.

2.5 General

2.5.1 Glossary

- NFV - Network Function Virtualization
- Mbps - 1,000,000bps

2.5.2 Document change procedures and history

Document ID	Author	Date Modified
OPNFV_vswitchperf_LTR_ver_1.0_Jan_15_CN_DRAFT	Christopher Nolan	23/01/2015
OPNFV_vswitchperf_LTR_ver_1.1_Jan_15_CN_DRAFT	Christopher Nolan	28/01/2015