
CSIT REPORT

Release rls1804

May 23, 2018

Contents

| | | |
|----------|---|------------|
| 1 | Introduction | 1 |
| 1.1 | Overview | 1 |
| 1.2 | General Notes | 2 |
| 2 | VPP Performance Tests | 4 |
| 2.1 | Overview | 4 |
| 2.2 | CSIT Release Notes | 8 |
| 2.3 | Test Methodology | 10 |
| 2.4 | Packet Throughput Graphs | 15 |
| 2.5 | Throughput Speedup Multi-Core | 63 |
| 2.6 | Packet Latency Graphs | 96 |
| 2.7 | VPP HTTP Server Performance Results | 120 |
| 2.8 | Experimental: MDR Search | 122 |
| 2.9 | Test Environment | 128 |
| 2.10 | Documentation | 165 |
| 3 | DPDK Performance Tests | 174 |
| 3.1 | Overview | 174 |
| 3.2 | CSIT Release Notes | 178 |
| 3.3 | Packet Throughput Graphs | 178 |
| 3.4 | Packet Latency Graphs | 186 |
| 3.5 | Test Environment | 190 |
| 3.6 | Documentation | 225 |
| 4 | VPP Functional Tests | 226 |
| 4.1 | Overview | 226 |
| 4.2 | CSIT Release Notes | 228 |
| 4.3 | Test Environment | 229 |
| 4.4 | Documentation | 239 |
| 5 | Honeycomb Functional Tests | 240 |
| 5.1 | Overview | 240 |
| 5.2 | CSIT Release Notes | 243 |
| 5.3 | Test Environment | 243 |
| 5.4 | Documentation | 243 |
| 6 | NSH_SFC Functional Tests | 244 |
| 6.1 | Overview | 244 |
| 6.2 | CSIT Release Notes | 246 |
| 6.3 | Test Environment | 246 |
| 6.4 | Documentation | 256 |

| | |
|--|------------|
| 7 CSIT Framework Documentation | 257 |
| 7.1 CSIT Design | 257 |
| 7.2 CSIT Test Naming | 261 |
| 7.3 Presentation and Analytics Layer | 263 |
| 7.4 CSIT TAGs Descriptions | 289 |
| Bibliography | 301 |

1.1 Overview

This is the **Fast Data I/O Project (FD.io) Continuous System Integration and Testing (CSIT)** project report for CSIT rls1804 system testing of VPP-18.04 release.

This is the full html version, there is also a reduced [PDF version of this report](#)¹.

The report describes CSIT functional and performance tests and their continuous execution delivered in CSIT rls1804. A high-level overview is provided for each CSIT test environment running in LF (Linux Foundation) FD.io Continuous Performance Labs. This is followed by summary of all executed tests against the VPP-18.04 release and associated FD.io projects and sub-systems (Honeycomb, DPDK, NSH_SFC), CSIT rls1804 release notes, result highlights and known issues discovered in CSIT. More detailed description of each environment, pointers to CSIT test code documentation and detailed test results with links to the source data files are also provided.

CSIT rls1804 report contains following main sections and sub-sections:

1. **Introduction** - general introduction to CSIT project; *Overview* - this section; *CSIT Test Naming* - CSIT general naming convention for test suites and test cases, important to recognize the high-level test content and interpret reported results; *General Notes* - general notes related to this report.
2. **VPP Performance Tests** - VPP performance tests executed in physical FD.io testbeds; *Overview* - tested topologies, test coverage and naming specifics, methodology for multi-core, packet throughput and latency, and KVM VM vhost tests; *CSIT Release Notes* - changes in CSIT rls1804, added tests, performance changes, environment or methodology changes, known CSIT issues; *Packet Throughput Graphs* and *Packet Latency Graphs* - plotted NDR, PDR throughput and latency results from multiple test job executions; *Throughput Speedup Multi-Core* - plotted core configuration speedup comparison; *Test Environment* - environment description; *VPP HTTP Server Performance Results* - plotted HTTP Server performance; *Documentation* - CSIT source code documentation for VPP performance tests.
3. **DPDK Performance Tests** - DPDK performance tests executed in physical FD.io testbeds; *Overview* - tested topologies, test coverage; *CSIT Release Notes* - changes in CSIT rls1804, any known CSIT issues; *Packet Throughput Graphs* and *Packet Latency Graphs* - plotted NDR, PDR throughput and latency results from multiple test job executions; *Test Environment* - environment description; *Documentation* - CSIT source code documentation for DPDK performance tests.

¹ https://docs.fd.io/csit/rls1804/report/_static/archive/csit_rls1804.pdf

4. **VPP Functional Tests** - VPP functional tests executed in virtual FD.io testbeds; *Overview* - tested virtual topologies, test coverage and naming specifics; *CSIT Release Notes* - changes in CSIT rls1804, added tests, environment or methodology changes, known CSIT issues, tests to be added; *Test Environment* - environment description ; *Documentation* - source code documentation for VPP functional tests.
5. **Honeycomb Functional Tests** - Honeycomb functional tests executed in virtual FD.io testbeds; *Overview* - tested virtual topologies, test coverage and naming specifics; *CSIT Release Notes* - changes in CSIT rls1804, added tests, environment or methodology changes, known CSIT issues; *Test Environment* - environment description ; *Documentation* - source code documentation for Honeycomb functional tests.
6. **NSH_SFC Functional Tests** - NSH_SFC functional tests executed in virtual FD.io testbeds; *Overview* - tested virtual topologies, test coverage and naming specifics; *CSIT Release Notes* - changes in CSIT rls1804, added tests, environment or methodology changes, known CSIT issues; *Test Environment* - environment description ; *Documentation* - source code documentation for NSH_SFC functional tests.
7. **Detailed Test Results** - auto-generated results from CSIT jobs executions using CSIT Robot Framework output files as source data; *VPP Performance Results*, *DPDK Performance Results*, *VPP Functional Results*, *Honeycomb Functional Results*, *VPPtest Functional Results*.
8. **Test Configuration** - auto-generated DUT configuration data from CSIT jobs executions using CSIT Robot Framework output files as source data; *VPP Performance Test Configs*, *VPP Functional Test Configs*.
9. **Test Operational Data** - auto-generated DUT operational data from CSIT jobs executions using CSIT Robot Framework output files as source data; *VPP Performance Operational Data*.
10. **CSIT Framework Documentation** - description of the overall CSIT framework design hierarchy, CSIT test naming convention, followed by description of Presentation and Analytics Layer (PAL) introduced in CSIT rls1804.

1.2 General Notes

All CSIT test results listed in this report are sourced and auto-generated from output.xml RF (Robot Framework) files resulting from LF FD.io Jenkins jobs execution against VPP-18.04 release artifacts. References are provided to the original LF FD.io Jenkins job results. Additional references are provided to the RF result files that got archived in FD.io nexus online storage system.

FD.io CSIT project currently covers multiple FD.io system and sub-system testing areas and this is reflected in this report, where each testing area is listed separately, as follows:

1. **VPP - Performance** - VPP benchmarking tests are executed in physical FD.io testbeds, focusing on VPP network data plane performance at this stage, both for Phy-to-Phy (NIC-to-NIC) and Phy-to-VM-to-Phy (NIC-to-VM-to-NIC) forwarding topologies. Tested across a range of NICs, 10GE and 40GE interfaces, range of multi-thread and multi-core configurations. VPP application runs in host user-mode. TRex is used as a traffic generator.
2. **LXC and Docker Containers VPP memif - Performance** - VPP memif virtual interface tests interconnect multiple VPP instances running in containers. VPP vswitch instance runs in bare-metal user-mode handling Intel x520 NIC 10GbE, Intel x710 NIC 10GbE, Intel xl710 NIC 40GbE interfaces and connecting over memif (Slave side) virtual interfaces to more instances of VPP running in LXC or in Docker Containers, both with memif virtual interfaces (Master side). Tested across a range of multi-thread and multi-core configurations. TRex is used as a traffic generator.
3. **Container Topologies Orchestrated by K8s - Performance** - CSIT Container topologies connected over the memif virtual interface (shared memory interface). For these tests VPP vswitch instance runs in a Docker Container handling Intel x520 NIC 10GbE, Intel x710 NIC 10GbE interfaces and connecting over memif virtual interfaces to more instances of VPP running in Docker Containers

with memif virtual interfaces. All containers are orchestrated by Kubernetes, with [Ligato](#)² for container networking. TRex is used as a traffic generator.

4. **DPDK Performance** - VPP is using DPDK code to control and drive the NICs and physical interfaces. Tests are used as a baseline to profile performance of the DPDK sub-system. DPDK tests are executed in physical FD.io testbeds, focusing on Testpmd/L3FWD data plane performance for Phy-to-Phy (NIC-to-NIC). Tests cover a range of NICs, 10GE and 40GE interfaces, range of multi-thread and multi-core configurations. Testpmd/L3FWD application runs in host user-mode. TRex is used as a traffic generator.
5. **VPP Functional** - VPP functional tests are executed in virtual FD.io testbeds focusing on VPP packet processing functionality, including network data plane and in-line control plane. Tests cover vNIC-to-vNIC vNIC-to-VM-to-vNIC forwarding topologies. Scapy is used as a traffic generator.
6. **Honeycomb Functional** - Honeycomb functional tests are executed in virtual FD.io testbeds, focusing on Honeycomb management and programming functionality of VPP. Tests cover a range of CRUD operations executed against VPP.
7. **NSH_SFC Functional** - NSH_SFC functional tests are executed in virtual FD.io testbeds focusing on NSH_SFC of VPP. Tests cover a range of CRUD operations executed against VPP.

In addition to above, CSIT rls1804 report does also include VPP unit test results. VPP unit tests are developed within the FD.io VPP project and as they complement CSIT system functional tests, they are provided mainly as a reference and to provide a more complete view of automated testing executed against VPP-18.04 release.

FD.io CSIT system is developed using two main coding platforms RF and Python2.7. CSIT rls1804 source code for the executed test suites is available in CSIT branch rls1804 in the directory `./tests/<name_of_the_test_suite>`. A local copy of CSIT source code can be obtained by cloning CSIT git repository - `git clone https://gerrit.fd.io/r/csit`. The CSIT testing virtual environment can be run on a local computer workstation (laptop, server) using Vagrant by following the instructions in [CSIT tutorials](#)³.


² <https://github.com/ligato>

³ <https://wiki.fd.io/view/CSIT#Tutorials>

2.1 Overview

2.1.1 Tested Physical Topologies

CSIT VPP performance tests are executed on physical baremetal servers hosted by LF FD.io project. Testbed physical topology is shown in the figure below.:



SUT1 and SUT2 are two System Under Test servers (Cisco UCS C240, each with two Intel XEON CPUs), TG is a Traffic Generator (TG, another Cisco UCS C240, with two Intel XEON CPUs). SUTs run VPP SW application in Linux user-mode as a Device Under Test (DUT). TG runs TRex SW application as a packet Traffic Generator. Physical connectivity between SUTs and to TG is provided using different NIC models that need to be tested for performance. Currently installed and tested NIC models include:

1. 2port10GE X520-DA2 Intel.
2. 2port10GE X710 Intel.
3. 2port10GE VIC1227 Cisco.
4. 2port40GE VIC1385 Cisco.

5. 2port40GE XL710 Intel.


From SUT and DUT perspective, all performance tests involve forwarding packets between two physical Ethernet ports (10GE or 40GE). Due to the number of listed NIC models tested and available PCI slot capacity in SUT servers, in all of the above cases both physical ports are located on the same NIC. In some test cases this results in measured packet throughput being limited not by VPP DUT but by either the physical interface or the NIC capacity.

Going forward CSIT project will be looking to add more hardware into FD.io performance labs to address larger scale multi-interface and multi-NIC performance testing scenarios.

For service chain topology test cases that require DUT (VPP) to communicate with VirtualMachines (VMs) or with Linux/Docker Containers (Ctrs) over vhost-user/memif interfaces, N of VM/Ctr instances are created on SUT1 and SUT2. Three types of service chain topologies are tested in CSIT rls1804:

1. "Parallel" topology with packets flowing from NIC via DUT (VPP) to VM/Container and back to VPP and NIC;
2. "Chained" topology (a.k.a. "Snake") with packets flowing via DUT (VPP) to VM/Container, back to DUT, then to the next VM/Container, back to DUT and so on until the last VM/Container in a chain, then back to DUT and NIC;
3. "Horizontal" topology with packets flowing via DUT (VPP) to Container, then via "horizontal" memif to the next Container, and so on until the last Container, then back to DUT and NIC. "Horizontal" topology is not supported for VMs;

For each of the above topologies, DUT (VPP) is tested in a range of L2 or IPv4/IPv6 configurations depending on the test suite. A sample DUT "Chained" service topology with N of VM/Ctr instances is shown in the figure below. Packet flow thru the DUTs and VMs/Ctrs is marked with ***:



In above "Chained" topology, packets are switched by DUT multiple times: twice for a single VM/Ctr, three times for two VMs/Ctrs, N+1 times for N VMs/Ctrs. Hence the external throughput rates measured by TG and listed in this report must be multiplied by (N+1) to represent the actual DUT aggregate packet forwarding rate.

For a "Parallel" and "Horizontal" service topologies packets are always switched by DUT twice per service chain.

Note that reported DUT (VPP) performance results are specific to the SUTs tested. Current LF FD.io SUTs are based on Intel XEON E5-2699v3 2.3GHz CPUs. SUTs with other CPUs are likely to yield different results. A good rule of thumb, that can be applied to estimate VPP packet throughput for Phy-to-Phy (NIC-to-NIC, PCI-to-PCI) topology, is to expect the forwarding performance to be proportional to CPU core frequency, assuming CPU is the only limiting factor and all other SUT parameters equivalent to FD.io CSIT

environment. The same rule of thumb can be also applied for Phy-to-VM/Ctr-to-Phy (NIC-to-VM/Ctr-to-NIC) topology, but due to much higher dependency on intensive memory operations and sensitivity to Linux kernel scheduler settings and behaviour, this estimation may not always yield good enough accuracy.

For detailed FD.io CSIT testbed specification and topology, as well as configuration and setup of SUTs and DUTs testbeds please refer to `test_environment`.

Similar SUT compute node can be arrived to in a standalone VPP setup by using a [vpp-config configuration tool](#)⁴ developed within the VPP project using CSIT recommended settings and scripts.

2.1.2 Performance Tests Coverage

Performance tests are split into two main categories:

- Throughput discovery - discovery of packet forwarding rate using binary search in accordance to [RFC 2544](#)⁵.
 - NDR - discovery of Non Drop Rate packet throughput, at zero packet loss; followed by one-way packet latency measurements at 10%, 50% and 100% of discovered NDR throughput.
 - PDR - discovery of Partial Drop Rate, with specified non-zero packet loss currently set to 0.5%; followed by one-way packet latency measurements at 100% of discovered PDR throughput.

CSIT rls1804 includes following performance test suites, listed per NIC type:

- 2port10GE X520-DA2 Intel
 - **L2XC** - L2 Cross-Connect switched-forwarding of untagged, dot1q, dot1ad VLAN tagged Ethernet frames.
 - **L2BD** - L2 Bridge-Domain switched-forwarding of untagged Ethernet frames with MAC learning; disabled MAC learning i.e. static MAC tests to be added.
 - **L2BD Scale** - L2 Bridge-Domain switched-forwarding of untagged Ethernet frames with MAC learning; disabled MAC learning i.e. static MAC tests to be added with 20k, 200k and 2M FIB entries.
 - **IPv4** - IPv4 routed-forwarding.
 - **IPv6** - IPv6 routed-forwarding.
 - **IPv4 Scale** - IPv4 routed-forwarding with 20k, 200k and 2M FIB entries.
 - **IPv6 Scale** - IPv6 routed-forwarding with 20k, 200k and 2M FIB entries.
 - **VMs with vhost-user** - virtual topologies with 1 VM and service chains of 2 VMs using vhost-user interfaces, with VPP forwarding modes incl. L2 Cross-Connect, L2 Bridge-Domain, VXLAN with L2BD, IPv4 routed-forwarding.
 - **COP** - IPv4 and IPv6 routed-forwarding with COP address security.
 - **ACL** - L2 Bridge-Domain switched-forwarding and IPv4 and IPv6 routed-forwarding with iACL and oACL IP address, MAC address and L4 port security.
 - **LISP** - LISP overlay tunneling for IPv4-over-IPv4, IPv6-over-IPv4, IPv6-over-IPv6, IPv4-over-IPv6 in IPv4 and IPv6 routed-forwarding modes.
 - **VXLAN** - VXLAN overlay tunnelling integration with L2XC and L2BD.
 - **QoS Policer** - ingress packet rate measuring, marking and limiting (IPv4).
 - **NAT** - (Source) Network Address Translation tests with varying number of users and ports per user.
 - **Container memif connections** - VPP memif virtual interface tests to interconnect VPP instances with L2XC and L2BD.

⁴ https://wiki.fd.io/view/VPP/Configuration_Tool

⁵ <https://tools.ietf.org/html/rfc2544.html>

- **Container K8s Orchestrated Topologies** - Container topologies connected over the memif virtual interface.
- **SRv6** - Segment Routing IPv6 tests.
- 2port40GE XL710 Intel
 - **L2XC** - L2 Cross-Connect switched-forwarding of untagged Ethernet frames.
 - **L2BD** - L2 Bridge-Domain switched-forwarding of untagged Ethernet frames with MAC learning.
 - **IPv4** - IPv4 routed-forwarding.
 - **IPv6** - IPv6 routed-forwarding.
 - **VMs with vhost-user** - virtual topologies with 1 VM and service chains of 2 VMs using vhost-user interfaces, with VPP forwarding modes incl. L2 Cross-Connect, L2 Bridge-Domain, VXLAN with L2BD, IPv4 routed-forwarding.
 - **IPSecSW** - IPSec encryption with AES-GCM, CBC-SHA1 ciphers, in combination with IPv4 routed-forwarding.
 - **IPSecHW** - IPSec encryption with AES-GCM, CBC-SHA1 ciphers, in combination with IPv4 routed-forwarding. Intel QAT HW acceleration.
 - **IPSec+LISP** - IPSec encryption with CBC-SHA1 ciphers, in combination with LISP-GPE overlay tunneling for IPv4-over-IPv4.
 - **VPP TCP/IP stack** - tests of VPP TCP/IP stack used with VPP built-in HTTP server.
 - **Container memif connections** - VPP memif virtual interface tests to interconnect VPP instances with L2XC and L2BD.
- 2port10GE X710 Intel
 - **L2BD** - L2 Bridge-Domain switched-forwarding of untagged Ethernet frames with MAC learning.
 - **VMs with vhost-user** - virtual topologies with 1 VM using vhost-user interfaces, with VPP forwarding modes incl. L2 Bridge-Domain.
 - **Container memif connections** - VPP memif virtual interface tests to interconnect VPP instances with L2XC and L2BD.
 - **Container K8s Orchestrated Topologies** - Container topologies connected over the memif virtual interface.
- 2port10GE VIC1227 Cisco
 - **L2BD** - L2 Bridge-Domain switched-forwarding of untagged Ethernet frames with MAC learning.
- 2port40GE VIC1385 Cisco
 - **L2BD - L2 Bridge-Domain switched-forwarding of untagged Ethernet frames** with MAC learning.

Execution of performance tests takes time, especially the throughput discovery tests. Due to limited HW testbed resources available within FD.io labs hosted by LF, the number of tests for NICs other than X520 (a.k.a. Niantic) has been limited to few baseline tests. CSIT team expect the HW testbed resources to grow over time, so that complete set of performance tests can be regularly and(or) continuously executed against all models of hardware present in FD.io labs.

2.1.3 Performance Tests Naming

CSIT rls1804 follows a common structured naming convention for all performance and system functional tests, introduced in CSIT rls1701.

The naming should be intuitive for majority of the tests. Complete description of CSIT test naming convention is provided on *CSIT Test Naming* (page 261).

2.2 CSIT Release Notes

2.2.1 Changes in CSIT rls1804

1. VPP performance tests

- *MRR tests* - New Maximum Receive Rate tests measure the packet forwarding rate under the maximum load offered by traffic generator over a set trial duration, regardless of packet loss. MRR tests are used for continuous performance trending and for comparison between releases.
- *Service Chaining with SRv6* - New SRv6 (Segment Routing IPv6) proxy tests measure performance of SRv6 Endpoint fronting SR-unaware appliance via masquerading (End.AM), dynamic proxy (End.AD) or static proxy (End.AS) SR functions.

2. Presentation and Analytics Layer

- *Performance trending* - Added continuous performance trending and analysis. New Performance Trending and Performance Analysis jobs executed regular throughput tests, with results being subsequently analysed and trend and anomalies summarized and presented in VPP Performance Dashboard and trendline graphs.

3. Test Framework Optimizations

- *Performance tests efficiency* - Qemu build/install optimizations, warmup phase handling, vpp restart handling. Resulted in improved stability and reduced total execution time by 30% for single pkt size e.g. 64B/78B.
- *General code housekeeping* - ongoing RF keywords optimizations, removal of redundant RF keywords.

2.2.2 Performance Changes

Relative performance changes in measured packet throughput in CSIT rls1804 are calculated against the results from CSIT rls1801 report. Listed mean and standard deviation values are computed based on a series of the same tests executed against respective VPP releases to verify test results repeatability, with percentage change calculated for mean values. Note that the standard deviation is quite high for a small number of packet throughput tests, what indicates poor test results repeatability and makes the relative change of mean throughput value not fully representative for these tests. The root causes behind poor results repeatability vary between the test cases.

NDR Changes

NDR small packet throughput changes between releases are available in a CSV and pretty ASCII formats:

- [csv format for 1t1c](#),
- [csv format for 2t2c](#),
- [pretty ASCII format for 1t1c](#),
- [pretty ASCII format for 2t2c](#).

Note: Test results have been generated by [FD.io test executor vpp performance jobs](#)⁶ with Robot Framework result files csit-vpp-perf-1804-*.zip [archived here](#).

PDR Changes

NDR small packet throughput changes between releases are available in a CSV and pretty ASCII formats:

- [csv format for 1t1c](#),
- [csv format for 2t2c](#),
- [pretty ASCII format for 1t1c](#),
- [pretty ASCII format for 2t2c](#).

Note: Test results have been generated by [FD.io test executor vpp performance jobs](#)⁷ with Robot Framework result files csit-vpp-perf-1804-*.zip [archived here](#).

MRR Changes

MRR small packet throughput changes between releases are available in a CSV and pretty ASCII formats:

- [csv format for 1t1c](#),
- [csv format for 2t2c](#),
- [csv format for 4t4c](#),
- [pretty ASCII format for 1t1c](#),
- [pretty ASCII format for 2t2c](#),
- [pretty ASCII format for 4t4c](#).

Note: Test results have been generated by [FD.io test executor vpp mrr jobs](#)⁸ with Robot Framework result files csit-vpp-perf-mrr-daily-master__*_output.xml.gz [archived here](#)⁹.

2.2.3 Throughput Trending

In addition to reporting throughput changes between VPP releases, CSIT provides continuous performance trending for VPP master branch:

1. [VPP Performance Dashboard](#)¹⁰ - per VPP test case throughput trend, trend compliance and summary of detected anomalies.
2. [Trending Methodology](#)¹¹ - throughput test metrics, trend calculations and anomaly classification (progression, regression, outlier).
3. [Trendline Graphs](#)¹² - per VPP build MRR throughput measurements against the trendline with anomaly highlights, with associated CSIT test jobs.

⁶ <https://jenkins.fd.io/view/csit/job/csit-vpp-perf-1804-all>

⁷ <https://jenkins.fd.io/view/csit/job/csit-vpp-perf-1804-all>

⁸ <https://jenkins.fd.io/view/csit/job/csit-vpp-perf-mrr-daily-master/>

⁹ https://docs.fd.io/csit/master/trending/_static/archive/

¹⁰ <https://docs.fd.io/csit/master/trending/introduction/index.html>

¹¹ <https://docs.fd.io/csit/master/trending/methodology/index.html>

¹² <https://docs.fd.io/csit/master/trending/trending/index.html>

2.2.4 Known Issues

List of known issues in CSIT rls1804 for VPP performance tests:

| # | Issue | Jira ID | Description |
|---|--|-----------|---|
| 1 | Sporadic (1 in 200) NDR discovery test failures on x520. | CSIT-570 | DPDK reporting rx-errors, indicating L1 issue. Suspected issue with HW combination of X710-X520 in LF testbeds. Not observed outside of LF testbeds. |
| 2 | Lower than expected NDR throughput of DPDK testpmd and VPP L2 path NDR throughput with xl710 and x710 NICs, compared to x520 NICs. | CSIT-571 | Suspected NIC firmware or DPDK driver issue affecting NDR and PDR throughput on XL710 and X710 NICs. |
| 3 | Tagged Ethernet dot1q and dot1ad L2 path throughput regression. | CSIT-1066 | Tagged Ethernet dot1q and dot1ad L2 path throughput regression: NDR -2%..-5%, PDR -2%..-6%, MRR. Affects I2xc and I2bd performance tests. |
| 4 | IPSec (software, no QAT HW) throughput regression. | CSIT-1064 | IPSec throughput regression: NDR -3%..-8%, PDR -2%..-8%, MRR -3%..-7%. Affects IPSec SW tests, QAT HW tests not affected. |
| 5 | High failure rate of creating working container topologies with K8s/Ligato orchestration. | CSIT-1065 | About 20% of orchestrated container topology tests failing data plane verification indicating configuration issue. Suspected issue with Ligato vpp-agent. |

2.3 Test Methodology

2.3.1 Multi-Core and Multi-Threading

Intel Hyper-Threading - CSIT rls1804 performance tests are executed with SUT servers' Intel XEON processors configured in Intel Hyper-Threading Disabled mode (BIOS setting). This is the simplest configuration used to establish baseline single-thread single-core application packet processing and forwarding performance. Subsequent releases of CSIT will add performance tests with Intel Hyper-Threading Enabled (requires BIOS settings change and hard reboot of server).

Multi-core Tests - CSIT rls1804 multi-core tests are executed in the following VPP thread and core configurations:

- 1t1c - 1 VPP worker thread on 1 CPU physical core.
- 2t2c - 2 VPP worker threads on 2 CPU physical cores.
- 4t4c - 4 VPP worker threads on 4 CPU physical cores.

VPP worker threads are the data plane threads. VPP control thread is running on a separate non-isolated core together with other Linux processes. Note that in quite a few test cases running VPP workers on 2 or 4 physical cores hits the I/O bandwidth or packets-per-second limit of tested NIC.

Section throughput_speedup_multi_core includes a set of graphs illustrating packet throughput speedup when running VPP on multiple cores.

2.3.2 Packet Throughput

Following values are measured and reported for packet throughput tests:

- NDR binary search per [RFC 2544](https://tools.ietf.org/html/rfc2544)¹³:

¹³ <https://tools.ietf.org/html/rfc2544.html>

- Packet rate: "RATE: <aggregate packet rate in packets-per-second> pps (2x <per direction packets-per-second>);"
- Aggregate bandwidth: "BANDWIDTH: <aggregate bandwidth in Gigabits per second> Gbps (untagged)";
- PDR binary search per **RFC 2544**¹⁴:
 - Packet rate: "RATE: <aggregate packet rate in packets-per-second> pps (2x <per direction packets-per-second>);"
 - Aggregate bandwidth: "BANDWIDTH: <aggregate bandwidth in Gigabits per second> Gbps (untagged)";
 - Packet loss tolerance: "LOSS_ACCEPTANCE <accepted percentage of packets lost at PDR rate>";
- NDR and PDR are measured for the following L2 frame sizes (untagged Ethernet):
 - IPv4 payload: 64B, IMIX_v4_1 (28x64B,16x570B,4x1518B), 1518B, 9000B;
 - IPv6 payload: 78B, 1518B, 9000B;
- NDR and PDR binary search resolution is determined by the final value of the rate change, referred to as the final step:
 - The final step is set to 50kpps for all NIC to NIC tests and all L2 frame sizes except 9000B (changed from 100kpps used in previous releases).
 - The final step is set to 10kpps for all remaining tests, including 9000B and all vhost VM and memif Container tests.

All rates are reported from external Traffic Generator perspective.

2.3.3 Maximum Receive Rate (MRR)

MRR tests measure the packet forwarding rate under the maximum load offered by traffic generator over a set trial duration, regardless of packet loss. Maximum load for specified Ethernet frame size is set to the bi-directional link rate.

Current parameters for MRR tests:

- Ethernet frame sizes: 64B (78B for IPv6 tests) for all tests, IMIX for selected tests (vhost, memif); all quoted sizes include frame CRC, but exclude per frame transmission overhead of 20B (preamble, inter frame gap).
- Maximum load offered: 10GE and 40GE link (sub-)rates depending on NIC tested, with the actual packet rate depending on frame size, transmission overhead and traffic generator NIC forwarding capacity.
 - For 10GE NICs the maximum packet rate load is 2* 14.88 Mpps for 64B, a 10GE bi-directional link rate.
 - For 40GE NICs the maximum packet rate load is 2* 18.75 Mpps for 64B, a 40GE bi-directional link sub-rate limited by TG 40GE NIC used, XL710.
- Trial duration: 10sec.

Similarly to NDR/PDR throughput tests, MRR test should be reporting bi- directional link rate (or NIC rate, if lower) if tested VPP configuration can handle the packet rate higher than bi-directional link rate, e.g. large packet tests and/or multi-core tests.

MRR tests are used for continuous performance trending and for comparison between releases.

¹⁴ <https://tools.ietf.org/html/rfc2544.html>

2.3.4 Packet Latency

TRex Traffic Generator (TG) is used for measuring latency of VPP DUTs. Reported latency values are measured using following methodology:

- Latency tests are performed at 10%, 50% of discovered NDR rate (non drop rate) for each NDR throughput test and packet size (except IMIX).
- TG sends dedicated latency streams, one per direction, each at the rate of 10kpps at the prescribed packet size; these are sent in addition to the main load streams.
- TG reports min/avg/max latency values per stream direction, hence two sets of latency values are reported per test case; future release of TRex is expected to report latency percentiles.
- Reported latency values are aggregate across two SUTs due to three node topology used for all performance tests; for per SUT latency, reported value should be divided by two.
- 1usec is the measurement accuracy advertised by TRex TG for the setup used in FD.io labs used by CSIT project.
- TRex setup introduces an always-on error of about 2*2usec per latency flow - additional Tx/Rx interface latency induced by TRex SW writing and reading packet timestamps on CPU cores without HW acceleration on NICs closer to the interface line.

2.3.5 vhostuser with KVM VMs

FD.io CSIT performance lab is testing VPP vhost with KVM VMs using following environment settings:

- Tests with varying Qemu virtio queue (a.k.a. vring) sizes: [vr256] default 256 descriptors, [vr1024] 1024 descriptors to optimize for packet throughput.
- Tests with varying Linux CFS (Completely Fair Scheduler) settings: [cfs] default settings, [cfsrr1] CFS RoundRobin(1) policy applied to all data plane threads handling test packet path including all VPP worker threads and all Qemu testpmd poll-mode threads.
- Resulting test cases are all combinations with [vr256,vr1024] and [cfs,cfsrr1] settings.
- Adjusted Linux kernel CFS scheduler policy for data plane threads used in CSIT is documented in [CSIT Performance Environment Tuning wiki](#)¹⁵. The purpose is to verify performance impact (MRR and NDR/PDR throughput) and same test measurements repeatability, by making VPP and VM data plane threads less susceptible to other Linux OS system tasks hijacking CPU cores running those data plane threads.

2.3.6 Memif with LXC and Docker Containers

CSIT rls1804 includes tests taking advantage of VPP memif virtual interface (shared memory interface) to interconnect VPP running in Containers. VPP vswitch instance runs in bare-metal user-mode handling NIC interfaces and connecting over memif (Slave side) to VPPs running in LINUX CONTAINER (LXC) or in Docker Container (DRC) configured with memif (Master side). LXC and DRCs run in a privileged mode with VPP data plane worker threads pinned to dedicated physical CPU cores per usual CSIT practice. All VPP instances run the same version of software. This test topology is equivalent to existing tests with vhost-user and VMs as described earlier in *Tested Physical Topologies* (page 4).

More information about CSIT LXC and DRC setup and control is available in *Container Orchestration in CSIT* (page 165).

¹⁵ <https://wiki.fd.io/view/CSIT/csit-perf-env-tuning-ubuntu1604>

2.3.7 Memif with K8s Pods/Containers

CSIT rls1804 includes tests of VPP topologies running in K8s orchestrated Pods/Containers and connected over memif virtual interfaces. In order to provide simple topology coding flexibility and extensibility container orchestration is done with [Kubernetes](#)¹⁶ using [Docker](#)¹⁷ images for all container applications including VPP. [Ligato](#)¹⁸ is used for the Pod/Container networking orchestration that is integrated with K8s, including memif support.

In these tests VPP vswitch runs in a K8s Pod with Docker Container (DRC) handling NIC interfaces and connecting over memif to more instances of VPP running in Pods/DRCs. All DRCs run in a privileged mode with VPP data plane worker threads pinned to dedicated physical CPU cores per usual CSIT practice. All VPP instances run the same version of software. This test topology is equivalent to existing tests with vhost-user and VMs as described earlier in *Tested Physical Topologies* (page 4).

Further documentation is available in *Container Orchestration in CSIT* (page 165).

2.3.8 IPsec with Intel QAT HW cards

VPP IPsec performance tests are using DPDK cryptodev device driver in combination with HW cryptodev devices - Intel QAT 8950 50G - present in LF FD.io physical testbeds. DPDK cryptodev can be used for all IPsec data plane functions supported by VPP.

Currently CSIT rls1804 implements following IPsec test cases:

- AES-GCM, CBC-SHA1 ciphers, in combination with IPv4 routed-forwarding with Intel xl710 NIC.
- CBC-SHA1 ciphers, in combination with LISP-GPE overlay tunneling for IPv4-over-IPv4 with Intel xl710 NIC.

2.3.9 TRex Traffic Generator Usage

[TRex traffic generator](#)¹⁹ is used for all CSIT performance tests. TRex stateless mode is used to measure NDR and PDR throughputs using binary search (NDR and PDR discovery tests) and for quick checks of DUT performance against the reference NDRs (NDR check tests) for specific configuration.

TRex is installed and run on the TG compute node. The typical procedure is:

- If the TRex is not already installed on TG, it is installed in the suite setup phase - see [TRex installation](#)²⁰.
- TRex configuration is set in its configuration file

```
/etc/trex_cfg.yaml
```

- TRex is started in the background mode

```
$ sh -c 'cd <t-rex-install-dir>/scripts/ && sudo nohup ./t-rex-64 -i -c 7 --iom 0 > /tmp/trex.  
↪ log 2>&1 &' > /dev/null
```

- There are traffic streams dynamically prepared for each test, based on traffic profiles. The traffic is sent and the statistics obtained using `trex_stl_lib.api.STLClient`.

Measuring packet loss

- Create an instance of STLClient
- Connect to the client

¹⁶ <https://github.com/kubernetes>

¹⁷ <https://github.com/docker>

¹⁸ <https://github.com/ligato>

¹⁹ <https://wiki.fd.io/view/TRex>

²⁰ https://git.fd.io/csit/tree/resources/tools/trex/trex_installer.sh?h=rls1804

- Add all streams
- Clear statistics
- Send the traffic for defined time
- Get the statistics

If there is a warm-up phase required, the traffic is sent also before test and the statistics are ignored.

Measuring latency

If measurement of latency is requested, two more packet streams are created (one for each direction) with TRex flow_stats parameter set to STLFlowLatencyStats. In that case, returned statistics will also include min/avg/max latency values.

2.3.10 TCP/IP tests with WRK tool

WRK HTTP benchmarking tool²¹ is used for experimental TCP/IP and HTTP tests of VPP TCP/IP stack and built-in static HTTP server. WRK has been chosen as it is capable of generating significant TCP/IP and HTTP loads by scaling number of threads across multi-core processors.

This in turn enables quite high scale benchmarking of the main TCP/IP and HTTP service including HTTP TCP/IP Connections-Per-Second (CPS), HTTP Requests-Per-Second and HTTP Bandwidth Throughput.

The initial tests are designed as follows:

- HTTP and TCP/IP Connections-Per-Second (CPS)
 - WRK configured to use 8 threads across 8 cores, 1 thread per core.
 - Maximum of 50 concurrent connections across all WRK threads.
 - Timeout for server responses set to 5 seconds.
 - Test duration is 30 seconds.
 - Expected HTTP test sequence:
 - * Single HTTP GET Request sent per open connection.
 - * Connection close after valid HTTP reply.
 - * Resulting flow sequence - 8 packets: >S,<S-A,>A,>Req,<Rep,>F,<F,> A.
- HTTP Requests-Per-Second
 - WRK configured to use 8 threads across 8 cores, 1 thread per core.
 - Maximum of 50 concurrent connections across all WRK threads.
 - Timeout for server responses set to 5 seconds.
 - Test duration is 30 seconds.
 - Expected HTTP test sequence:
 - * Multiple HTTP GET Requests sent in sequence per open connection.
 - * Connection close after set test duration time.
 - * Resulting flow sequence: >S,<S-A,>A,>Req[1],<Rep[1],...,>Req[n],<Rep[n],>F,<F,>A.

²¹ <https://github.com/wg/wrk>

2.4 Packet Throughput Graphs

Plotted results are generated by multiple executions of the same CSIT performance tests across three physical testbeds within LF FD.io labs. To provide a descriptive summary view, Box-and-Whisker plots are used to display variation in measured throughput values, without making any assumptions of the underlying statistical distribution.

For each plotted test case, Box-and-Whisker plots show the quartiles (Min, 1st quartile / 25th percentile, 2nd quartile / 50th percentile / mean, 3rd quartile / 75th percentile, Max) across collected data set (data set size stated in the note below). Outliers are plotted as individual points. Min and max values are plotted as bottom and top Whiskers respectively. 2nd and 3rd quartiles are plotted as bottom and top edge of the box. If multiple samples match only two values, and all samples fall between them, then no whiskers are plotted. If all samples have the same value, only a horizontal line is plotted.

Title of each graph is a regex (regular expression) matching all throughput test cases plotted on this graph, *X-axis labels* are indices of individual test suites executed by [FD.io test executor vpp performance jobs](#)²² jobs that created result output files used as data sources for the graph, *Y-axis labels* are measured Packets Per Second [pps] values, and the *Graph legend* lists the plotted test suites and their indices.

Note: Test results have been generated by [FD.io test executor vpp performance jobs](#)²³ with Robot Framework result files csit-vpp-perf-1804-*.zip [archived here](#). Plotted data set size per test case is equal to the number of job executions presented in this report version: **10**.

2.4.1 L2 Ethernet Switching

Following sections include summary graphs of VPP Phy-to-Phy performance with L2 Ethernet switching, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss). Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

NDR Throughput

VPP NDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below.

²² <https://jenkins.fd.io/view/csit/job/csit-vpp-perf-1804-all>

²³ <https://jenkins.fd.io/view/csit/job/csit-vpp-perf-1804-all>




Figure 1. VPP 1thread 1core - NDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)²⁴.

²⁴ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>




Figure 2. VPP 1thread 1core - NDR Throughput for Phy-to-Phy L2 Ethernet Switching (feature).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)²⁵.

VPP NDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below.

²⁵ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>




Figure 3. VPP 2threads 2cores - NDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)²⁶.

²⁶ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>




Figure 4. VPP 2threads 2cores - NDR Throughput for Phy-to-Phy L2 Ethernet Switching (feature).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)²⁷.

PDR Throughput

VPP PDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

²⁷ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>




Figure 5. VPP 1thread 1core - PDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)²⁸.

²⁸ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>




Figure 6. VPP 1thread 1core - PDR Throughput for Phy-to-Phy L2 Ethernet Switching (feature).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)²⁹.

VPP PDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

²⁹ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>




Figure 7. VPP 2thread 2core - PDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)³⁰.

³⁰ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>




Figure 8. VPP 2thread 2core - PDR Throughput for Phy-to-Phy L2 Ethernet Switching (feature).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)³¹.

2.4.2 IPv4 Routed-Forwarding

Following sections include summary graphs of VPP Phy-to-Phy performance with IPv4 Routed-Forwarding, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss). Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

NDR Throughput

VPP NDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below.

³¹ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>




Figure 1. VPP 1thread 1core - NDR Throughput for Phy-to-Phy IPv4 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)³².

VPP NDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below.

³² <https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls1804>




Figure 2. VPP 2threads 2cores - NDR Throughput for Phy-to-Phy IPv4 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)³³.

PDR Throughput

VPP PDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

³³ <https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls1804>




Figure 3. VPP 1thread 1core - PDR Throughput for Phy-to-Phy IPv4 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)³⁴.

VPP PDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

³⁴ <https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls1804>




Figure 4. VPP 2thread 2core - PDR Throughput for Phy-to-Phy IPv4 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)³⁵.

2.4.3 IPv6 Routed-Forwarding

Following sections include summary graphs of VPP Phy-to-Phy performance with IPv6 Routed-Forwarding, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss). Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

NDR Throughput

VPP NDR 78B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below.

³⁵ <https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls1804>




Figure 1. VPP 1thread 1core - NDR Throughput for Phy-to-Phy IPv6 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)³⁶.

VPP NDR 78B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below.

³⁶ <https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls1804>




Figure 2. VPP 2threads 2cores - NDR Throughput for Phy-to-Phy IPv6 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)³⁷.

PDR Throughput

VPP PDR 78B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

³⁷ <https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls1804>




Figure 3. VPP 1thread 1core - PDR Throughput for Phy-to-Phy IPv6 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](https://git.fd.io/csit)³⁸.

³⁸ <https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls1804>




Figure 4. VPP 2thread 2core - PDR Throughput for Phy-to-Phy IPv6 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)³⁹.

2.4.4 SRv6

Following sections include summary graphs of VPP Phy-to-Phy performance with SRv6, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss). Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

NDR Throughput

VPP NDR 78B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below.

³⁹ <https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls1804>




Figure 1. VPP 1thread 1core - NDR Throughput for Phy-to-Phy SRv6.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁴⁰.

VPP NDR 78B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below.

⁴⁰ <https://git.fd.io/csit/tree/tests/vpp/perf/srv6?h=rls1804>




Figure 2. VPP 2threads 2cores - NDR Throughput for Phy-to-Phy SRv6.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁴¹.

PDR Throughput

VPP PDR 78B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

⁴¹ <https://git.fd.io/csit/tree/tests/vpp/perf/srv6?h=rls1804>




Figure 3. VPP 1thread 1core - PDR Throughput for Phy-to-Phy SRv6.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁴².

VPP PDR 78B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

⁴² <https://git.fd.io/csit/tree/tests/vpp/perf/srv6?h=rls1804>




Figure 4. VPP 2thread 2core - PDR Throughput for Phy-to-Phy IPv6 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁴³.

2.4.5 IPv4 Overlay Tunnels

Following sections include summary graphs of VPP Phy-to-Phy performance with IPv4 Overlay Tunnels, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss). Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

NDR Throughput

VPP NDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below.

⁴³ <https://git.fd.io/csit/tree/tests/vpp/perf/srv6?h=rls1804>




Figure 1. VPP 1thread 1core - NDR Throughput for Phy-to-Phy IPv4 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁴⁴.

VPP NDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below.

⁴⁴ https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls1804




Figure 2. VPP 2threads 2cores - NDR Throughput for Phy-to-Phy IPv4 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁴⁵.

PDR Throughput

VPP PDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

⁴⁵ https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls1804




Figure 3. VPP 1thread 1core - PDR Throughput for Phy-to-Phy IPv4 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁴⁶.

VPP PDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

⁴⁶ https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls1804




Figure 4. VPP 2thread 2core - PDR Throughput for Phy-to-Phy IPv4 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁴⁷.

2.4.6 IPv6 Overlay Tunnels

Following sections include summary graphs of VPP Phy-to-Phy performance with IPv6 Overlay Tunnels, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss). Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

NDR Throughput

VPP NDR 78B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below.

⁴⁷ https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls1804




Figure 1. VPP 1thread 1core - NDR Throughput for Phy-to-Phy IPv6 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁴⁸.

VPP NDR 78B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below.

⁴⁸ https://git.fd.io/csit/tree/tests/vpp/perf/ip6_tunnels?h=rls1804




Figure 2. VPP 2threads 2cores - NDR Throughput for Phy-to-Phy IPv6 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁴⁹.

PDR Throughput

VPP PDR 78B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

⁴⁹ https://git.fd.io/csit/tree/tests/vpp/perf/ip6_tunnels?h=rls1804




Figure 3. VPP 1thread 1core - PDR Throughput for Phy-to-Phy IPv6 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁵⁰.

VPP PDR 78B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

⁵⁰ https://git.fd.io/csit/tree/tests/vpp/perf/ip6_tunnels?h=rls1804




Figure 4. VPP 2thread 2core - PDR Throughput for Phy-to-Phy IPv6 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁵¹.

2.4.7 VM vhost Connections

Following sections include summary graphs of VPP Phy-to-VM(s)-to-Phy performance with VM virtio and VPP vhost-user virtual interfaces, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss). Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

NDR Throughput

VPP NDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below.

⁵¹ https://git.fd.io/csit/tree/tests/vpp/perf/ip6_tunnels?h=rls1804




Figure 1a. VPP 1thread 1core - NDR Throughput for Phy-to-VM-to-Phy VM vhost-user selected TCs. CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁵².

⁵² https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 1b. VPP 1thread 1core - NDR Throughput for Phy-to-VM-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁵³.

VPP NDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below.

⁵³ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 2a. VPP 2threads 2cores - NDR Throughput for Phy-to-VM-to-Phy VM vhost-user selected TCs. CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁵⁴.

⁵⁴ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 2b. VPP 2threads 2cores - NDR Throughput for Phy-to-VM-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁵⁵.

PDR Throughput

VPP PDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

⁵⁵ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 3a. VPP 1thread 1core - PDR Throughput for Phy-to-VM-to-Phy VM vhost-user selected TCs. CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁵⁶.

⁵⁶ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 3b. VPP 1thread 1core - PDR Throughput for Phy-to-VM-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁵⁷.

VPP PDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

⁵⁷ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 4a. VPP 2thread 2core - PDR Throughput for Phy-to-VM-to-Phy VM vhost-user selected TCs. CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁵⁸.

⁵⁸ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 4b. VPP 2thread 2core - PDR Throughput for Phy-to-VM-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁵⁹.

2.4.8 Container memif Connections

Following sections include summary graphs of VPP Phy-to-Phy performance with Container memif Connections, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss). Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

NDR Throughput

VPP NDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below.

⁵⁹ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 1. VPP 1thread 1core - NDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁶⁰.

VPP NDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below.

⁶⁰ https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls1804




Figure 2. VPP 2threads 2cores - NDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁶¹.

PDR Throughput

VPP PDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

⁶¹ https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls1804




Figure 3. VPP 1thread 1core - PDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁶².

VPP PDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

⁶² https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls1804




Figure 4. VPP 2thread 2core - PDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁶³.

2.4.9 Container Orchestrated Topologies

Following sections include summary graphs of VPP Phy-to-Phy performance with Container Orchestrated Topologies, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss). Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

NDR Throughput

VPP NDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below.

⁶³ https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls1804




Figure 1. VPP 1thread 1core - NDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁶⁴.

VPP NDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below.

⁶⁴ https://git.fd.io/csit/tree/tests/kubernetes/perf/container_memif?h=rls1804




Figure 2. VPP 2threads 2cores - NDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁶⁵.

PDR Throughput

VPP PDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

⁶⁵ https://git.fd.io/csit/tree/tests/kubernetes/perf/container_memif?h=rls1804




Figure 3. VPP 1thread 1core - PDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁶⁶.

⁶⁶ https://git.fd.io/csit/tree/tests/kubernetes/perf/container_memif?h=rls1804




Figure 4. VPP 2thread 2core - PDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁶⁷.

2.4.10 IPsec Crypto HW: IP4 Routed-Forwarding

Following sections include summary graphs of VPP Phy-to-Phy performance with IPsec encryption used in combination with IPv4 routed-forwarding, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss). VPP IPsec encryption is accelerated using DPDK cryptodev library driving Intel Quick Assist (QAT) crypto PCIe hardware cards. Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

NDR Throughput

VPP NDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below.

⁶⁷ https://git.fd.io/csit/tree/tests/kubernetes/perf/container_memif?h=rls1804




Figure 1. VPP 1thread 1core - NDR Throughput for Phy-to-Phy IPSEC HW.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁶⁸.

VPP NDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below.

⁶⁸ <https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls1804>




Figure 2. VPP 2threads 2cores - NDR Throughput for Phy-to-Phy IPSEC HW.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁶⁹.

PDR Throughput

VPP PDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

⁶⁹ <https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls1804>




Figure 3. VPP 1thread 1core - PDR Throughput for Phy-to-Phy IPSEC HW.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](https://github.com/fredmoo/csit)⁷⁰.

VPP PDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.

⁷⁰ <https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls1804>




Figure 4. VPP 2thread 2core - PDR Throughput for Phy-to-Phy IPSEC HW.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁷¹.

2.5 Throughput Speedup Multi-Core

Throughput Speedup Multi-Core graphs are generated based on results from multiple executions of the same CSIT performance tests across the three physical testbeds within LF FD.io labs. Plotted grouped bars illustrate the 64B/78B packet throughput speedup ratio for 2- and 4-core multi-threaded VPP configurations relative to 1-core configurations.

Note: Test results have been generated by [FD.io test executor vpp performance jobs](#)⁷² with Robot Framework result files `csit-vpp-perf-*.zip` [archived here](#).

⁷¹ <https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls1804>

⁷² <https://jenkins.fd.io/view/csit/job/csit-vpp-perf-1804-all>

2.5.1 L2 Ethernet Switching

Following sections include Throughput Speedup Analysis for VPP multi-core multi-thread configurations with no Hyper-Threading, specifically for tested 2t2c (2threads, 2cores) and 4t4c scenarios. 1t1c throughput results are used as a reference for reported speedup ratio. Input data used for the graphs comes from Phy-to-Phy 64B performance tests with VPP L2 Ethernet switching, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss).

NDR Throughput

VPP NDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 and 40ge2p1x1710 network interface cards.

NIC 10ge2p1x520




Figure 1. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy L2 Ethernet Switching.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁷³.

NIC 40ge2p1x1710




Figure 2. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy L2 Ethernet Switching.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁷⁴.

PDR Throughput

VPP PDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 and 40ge2p1x1710 network interface cards. PDR measured for 0.5% packet loss ratio.

⁷³ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>

⁷⁴ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>

NIC 10ge2p1x520




Figure 3. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized PDR Throughput for Phy-to-Phy L2 Ethernet Switching.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁷⁵.

2.5.2 IPv4 Routed-Forwarding

Following sections include Throughput Speedup Analysis for VPP multi-core multi-thread configurations with no Hyper-Threading, specifically for tested 2t2c (2threads, 2cores) and 4t4c scenarios. 1t1c throughput results are used as a reference for reported speedup ratio. Input data used for the graphs comes from Phy-to-Phy 64B performance tests with VPP IPv4 Routed-Forwarding, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss).

⁷⁵ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>

NDR Throughput

VPP NDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 and 40ge2p1x1710 network interface cards.

NIC 10ge2p1x520




Figure 1. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy IPv4 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁷⁶.

⁷⁶ <https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls1804>

NIC 40ge2p1x1710




Figure 2. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy IPv4 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁷⁷.

PDR Throughput

VPP PDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 and 40ge2p1x1710 network interface cards. PDR measured for 0.5% packet loss ratio.

⁷⁷ <https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls1804>

NIC 10ge2p1x520




Figure 3. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized PDR Throughput for Phy-to-Phy IPv4 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁷⁸.

2.5.3 IPv6 Routed-Forwarding

Following sections include Throughput Speedup Analysis for VPP multi-core multi-thread configurations with no Hyper-Threading, specifically for tested 2t2c (2threads, 2cores) and 4t4c scenarios. 1t1c throughput results are used as a reference for reported speedup ratio. Input data used for the graphs comes from Phy-to-Phy 78B performance tests with VPP IPv6 Routed-Forwarding, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss).

⁷⁸ <https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls1804>

NDR Throughput

VPP NDR 78B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 and 40ge2p1x1710 network interface cards.

NIC 10ge2p1x520




Figure 1. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy IPv6 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁷⁹.

⁷⁹ <https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls1804>

NIC 40ge2p1x1710




Figure 2. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy IPv6 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁸⁰.

PDR Throughput

VPP PDR 78B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 and 40ge2p1x1710 network interface cards. PDR measured for 0.5% packet loss ratio.

⁸⁰ <https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls1804>

NIC 10ge2p1x520




Figure 3. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized PDR Throughput for Phy-to-Phy IPv6 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls1804)⁸¹.

2.5.4 SRv6

Following sections include Throughput Speedup Analysis for VPP multi-core multi-thread configurations with no Hyper-Threading, specifically for tested 2t2c (2threads, 2cores) and 4t4c scenarios. 1t1c throughput results are used as a reference for reported speedup ratio. Input data used for the graphs comes from Phy-to-Phy 78B performance tests with VPP SRv6, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss).

⁸¹ <https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls1804>

NDR Throughput

VPP NDR 78B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 network interface card.

NIC 10ge2p1x520




Figure 1. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy SRv6.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](https://git.fd.io/csit/tree/tests/vpp/perf/srv6?h=rls1804)⁸².

PDR Throughput

VPP PDR 78B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 network interface card. PDR measured for 0.5% packet loss ratio.

⁸² <https://git.fd.io/csit/tree/tests/vpp/perf/srv6?h=rls1804>

NIC 10ge2p1x520




Figure 3. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized PDR Throughput for Phy-to-Phy SRv6.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁸³.

2.5.5 IPv4 Overlay Tunnels

Following sections include Throughput Speedup Analysis for VPP multi-core multi-thread configurations with no Hyper-Threading, specifically for tested 2t2c (2threads, 2cores) and 4t4c scenarios. 1t1c throughput results are used as a reference for reported speedup ratio. Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

⁸³ <https://git.fd.io/csit/tree/tests/vpp/perf/srv6?h=rls1804>

NDR Throughput

VPP NDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 network interface card.

NIC 10ge2p1x520




Figure 1. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy IPv4 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls1804)⁸⁴.

PDR Throughput

VPP PDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 network interface card.

⁸⁴ https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls1804

NIC 10ge2p1x520




Figure 2. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized PDR Throughput for Phy-to-Phy IPv4 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁸⁵.

2.5.6 IPv6 Overlay Tunnels

Following sections include Throughput Speedup Analysis for VPP multi-core multi-thread configurations with no Hyper-Threading, specifically for tested 2t2c (2threads, 2cores) and 4t4c scenarios. 1t1c throughput results are used as a reference for reported speedup ratio. Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

⁸⁵ https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls1804

NDR Throughput

VPP NDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 network interface card.

NIC 10ge2p1x520




Figure 1. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy IPv6 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁸⁶.

PDR Throughput

VPP PDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 network interface card.

⁸⁶ https://git.fd.io/csit/tree/tests/vpp/perf/ip6_tunnels?h=rls1804

NIC 10ge2p1x520




Figure 2. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized PDR Throughput for Phy-to-Phy IPv6 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁸⁷.

2.5.7 VM vhost Connections

Following sections include Throughput Speedup Analysis for VPP multi-core multi-thread configurations with no Hyper-Threading, specifically for tested 2t2c (2threads, 2cores) and 4t4c scenarios. 1t1c throughput results are used as a reference for reported speedup ratio. Input data used for the graphs comes from Phy-to-Phy 64B performance tests with VM vhost-user, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss).

⁸⁷ https://git.fd.io/csit/tree/tests/vpp/perf/ip6_tunnels?h=rls1804

NDR Throughput

VPP NDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520, 10ge2p1x710 and 40ge2p1x1710 network interface cards.

NIC 10ge2p1x520




Figure 1a. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁸⁸.

⁸⁸ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 1b. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁸⁹.

⁸⁹ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804

NIC 10ge2p1x710




Figure 2. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁹⁰.

⁹⁰ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804

NIC 40ge2p1x1710




Figure 3a. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁹¹.

⁹¹ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 3b. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁹².

PDR Throughput

VPP PDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520, 10ge2p1x710 and 40ge2p1x1710 network interface cards.

⁹² https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804

NIC 10ge2p1x520




Figure 4a. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized PDR Throughput for Phy-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁹³.

⁹³ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 4b. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized PDR Throughput for Phy-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁹⁴.

⁹⁴ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804

NIC 10ge2p1x710




Figure 5. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized PDR Throughput for Phy-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁹⁵.

⁹⁵ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804

NIC 40ge2p1x1710




Figure 6a. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized PDR Throughput for Phy-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁹⁶.

⁹⁶ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 6b. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized PDR Throughput for Phy-to-Phy VM vhost-user selected TCs.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁹⁷.

2.5.8 Container memif Connections

Following sections include Throughput Speedup Analysis for VPP multi-core multi-thread configurations with no Hyper-Threading, specifically for tested 2t2c (2threads, 2cores) and 4t4c scenarios. 1t1c throughput results are used as a reference for reported speedup ratio. Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

NDR Throughput

VPP NDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 network interface card.

⁹⁷ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804

NIC 10ge2p1x520




Figure 1. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁹⁸.

PDR Throughput

VPP PDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 network interface card.

⁹⁸ https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls1804

NIC 10ge2p1x520




Figure 2. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized PDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)⁹⁹.

2.5.9 Container Orchestrated Topologies

Following sections include Throughput Speedup Analysis for VPP multi-core multi-thread configurations with no Hyper-Threading, specifically for tested 2t2c (2threads, 2cores) and 4t4c scenarios. 1t1c throughput results are used as a reference for reported speedup ratio. Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

⁹⁹ https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls1804

NDR Throughput

VPP NDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 and 10ge2p1x710 network interface cards.

NIC 10ge2p1x520




Figure 1. VPP 1thread 1core - NDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](https://git.fd.io/csit/tree/tests/kubernetes/perf/container_memif?h=rls1804)¹⁰⁰.

¹⁰⁰ https://git.fd.io/csit/tree/tests/kubernetes/perf/container_memif?h=rls1804

NIC 10ge2p1x710




Figure 2. VPP 1thread 1core - NDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹⁰¹.

PDR Throughput

VPP PDR 64B packet throughput speedup ratio is presented in the graphs below for 10ge2p1x520 and 10ge2p1x710 network interface cards.

¹⁰¹ https://git.fd.io/csit/tree/tests/kubernetes/perf/container_memif?h=rls1804

NIC 10ge2p1x520




Figure 3. VPP 1thread 1core - NDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹⁰².

¹⁰² https://git.fd.io/csit/tree/tests/kubernetes/perf/container_memif?h=rls1804

NIC 10ge2p1x710




Figure 4. VPP 1thread 1core - NDR Throughput for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹⁰³.

2.5.10 IPsec Crypto HW: IP4 Routed-Forwarding

Following sections include Throughput Speedup Analysis for VPP multi-core multi-thread configurations with no Hyper-Threading, specifically for tested 2t2c (2threads, 2cores) and 4t4c scenarios. 1t1c throughput results are used as a reference for reported speedup ratio. VPP IPsec encryption is accelerated using DPDK cryptodev library driving Intel Quick Assist (QAT) crypto PCIe hardware cards. Performance is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

¹⁰³ https://git.fd.io/csit/tree/tests/kubernetes/perf/container_memif?h=rls1804

NDR Throughput

VPP NDR 64B packet throughput speedup ratio is presented in the graphs below for 40ge2p1x1710 network interface card.

NIC 40ge2p1x1710




Figure 1. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized NDR Throughput for Phy-to-Phy IPSEC HW.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls1804)¹⁰⁴.

PDR Throughput

VPP PDR 64B packet throughput speedup ratio is presented in the graphs below for 40ge2p1x1710 network interface card.

¹⁰⁴ <https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls1804>

NIC 40ge2p1x1710

VPP PDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.




Figure 2. Throughput Speedup Analysis - Multi-Core Speedup Ratio - Normalized PDR Throughput for Phy-to-Phy IPSEC HW.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹⁰⁵.

2.6 Packet Latency Graphs

Plotted results are generated from a single execution of CSIT NDR discovery test. Box plots are used to show the Minimum, Median and Maximum packet latency per test.

Title of each graph is a regex (regular expression) matching all throughput test cases plotted on this graph, X-axis labels are indices of individual test suites executed by [FD.io test executor vpp performance jobs](#)¹⁰⁶

¹⁰⁵ <https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls1804>

¹⁰⁶ <https://jenkins.fd.io/view/csit/job/csit-vpp-perf-1804-all>

that created result output file used as data source for the graph, *Y-axis labels* are measured packet Latency [uSec] values, and the *Graph legend* lists the plotted test suites and their indices. Latency is reported for concurrent symmetric bi-directional flows, separately for each direction: i) West-to-East: TGint1-to-SUT1-to-SUT2-to-TGint2, and ii) East-to-West: TGint2-to-SUT2-to-SUT1-to-TGint1.

Note: Test results have been generated by [FD.io test executor vpp performance jobs](#)¹⁰⁷ with Robot Framework result files csit-vpp-perf-1804-*.zip [archived here](#).

2.6.1 L2 Ethernet Switching

This section includes summary graphs of VPP Phy-to-Phy packet latency with L2 Ethernet switching measured at 50% of discovered NDR throughput rate. Latency is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

VPP packet latency in 1t1c setup (1thread, 1core) is presented in the graph below.




Figure 1a. VPP 1thread 1core - packet latency for Phy-to-Phy L2 Ethernet Switching (base).

¹⁰⁷ <https://jenkins.fd.io/view/csit/job/csit-vpp-perf-1804-all>

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹⁰⁸.




Figure 1b. VPP 1thread 1core - packet latency for Phy-to-Phy L2 Ethernet Switching (feature).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹⁰⁹.

VPP packet latency in 2t2c setup (2thread, 2core) is presented in the graph below.

¹⁰⁸ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>

¹⁰⁹ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>




Figure 2a. VPP 2threads 2cores - packet latency for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804)¹¹⁰.

¹¹⁰ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>




Figure 2b. VPP 2threads 2cores - packet latency for Phy-to-Phy L2 Ethernet Switching (feature).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804)¹¹¹.

2.6.2 IPv4 Routed-Forwarding

This section includes summary graphs of VPP Phy-to-Phy packet latency with IPv4 Routed-Forwarding measured at 50% of discovered NDR throughput rate. Latency is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

VPP packet latency in 1t1c setup (1thread, 1core) is presented in the graph below.

¹¹¹ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1804>




Figure 1. VPP 1thread 1core - packet latency for Phy-to-Phy IPv4 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹¹².

VPP packet latency in 2t2c setup (2thread, 2core) is presented in the graph below.

¹¹² <https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls1804>




Figure 2. VPP 2threads 2cores - packet latency for Phy-to-Phy IPv4 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹¹³.

2.6.3 IPv6 Routed-Forwarding

This section includes summary graphs of VPP Phy-to-Phy packet latency with IPv6 Routed-Forwarding measured at 50% of discovered NDR throughput rate. Latency is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

VPP packet latency in 1t1c setup (1thread, 1core) is presented in the graph below.

¹¹³ <https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls1804>




Figure 1. VPP 1thread 1core - packet latency for Phy-to-Phy IPv6 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹¹⁴.

VPP packet latency in 2t2c setup (2thread, 2core) is presented in the graph below.

¹¹⁴ <https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls1804>




Figure 2. VPP 2threads 2cores - packet latency for Phy-to-Phy IPv6 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹¹⁵.

2.6.4 SRv6

This section includes summary graphs of VPP Phy-to-Phy packet latency with SRv6 measured at 50% of discovered NDR throughput rate. Latency is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

VPP packet latency in 1t1c setup (1thread, 1core) is presented in the graph below.

¹¹⁵ <https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls1804>




Figure 1. VPP 1thread 1core - packet latency for Phy-to-Phy SRv6.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹¹⁶.

VPP packet latency in 2t2c setup (2thread, 2core) is presented in the graph below.

¹¹⁶ <https://git.fd.io/csit/tree/tests/vpp/perf/srv6?h=rls1804>




Figure 2. VPP 2threads 2cores - packet latency for Phy-to-Phy SRv6.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹¹⁷.

2.6.5 IPv4 Overlay Tunnels

This section includes summary graphs of VPP Phy-to-Phy packet latency with IPv4 Overlay Tunnels measured at 50% of discovered NDR throughput rate. Latency is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

VPP packet latency in 1t1c setup (1thread, 1core) is presented in the graph below.

¹¹⁷ <https://git.fd.io/csit/tree/tests/vpp/perf/srv6?h=rls1804>




Figure 1. VPP 1thread 1core - packet latency for Phy-to-Phy IPv4 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹¹⁸.

VPP packet latency in 2t2c setup (2thread, 2core) is presented in the graph below.

¹¹⁸ https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls1804




Figure 2. VPP 2threads 2cores - packet latency for Phy-to-Phy IPv4 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹¹⁹.

2.6.6 IPv6 Overlay Tunnels

This section includes summary graphs of VPP Phy-to-Phy packet latency with IPv6 Overlay Tunnels measured at 50% of discovered NDR throughput rate. Latency is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

VPP packet latency in 1t1c setup (1thread, 1core) is presented in the graph below.

¹¹⁹ https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls1804




Figure 1. VPP 1thread 1core - packet latency for Phy-to-Phy IPv6 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹²⁰.

VPP packet latency in 2t2c setup (2thread, 2core) is presented in the graph below.

¹²⁰ https://git.fd.io/csit/tree/tests/vpp/perf/ip6_tunnels?h=rls1804




Figure 2. VPP 2threads 2cores - packet latency for Phy-to-Phy IPv6 Overlay Tunnels.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](https://git.fd.io/csit/tree/tests/vpp/perf/ip6_tunnels?h=rls1804)¹²¹.

2.6.7 VM vhost Connections

This section includes summary graphs of VPP Phy-to-VM(s)-to-Phy packet latency with with VM virtio and VPP vhost-user virtual interfaces measured at 50% of discovered NDR throughput rate. Latency is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

VPP packet latency in 1t1c setup (1thread, 1core) is presented in the graph below.

¹²¹ https://git.fd.io/csit/tree/tests/vpp/perf/ip6_tunnels?h=rls1804




Figure 1a. VPP 1thread 1core - packet latency for Phy-to-VM-to-Phy VM vhost-user selected TCs (ipv4).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹²².

¹²² https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 1b. VPP 1thread 1core - packet latency for Phy-to-VM-to-Phy VM vhost-user selected TCs (I2). CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹²³. VPP packet latency in 2t2c setup (2thread, 2core) is presented in the graph below.

¹²³ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 2a. VPP 2threads 2cores - packet latency for Phy-to-VM-to-Phy VM vhost-user selected TCs (ipv4). CSIT source code for the test cases used for above plots can be found in [CSIT git repository](https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804)¹²⁴.

¹²⁴ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 2b. VPP 2threads 2cores - packet latency for Phy-to-VM-to-Phy VM vhost-user selected TCs (12). CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹²⁵.

2.6.8 Container memif Connections

This section includes summary graphs of VPP Phy-to-Phy packet latency with Container memif Connections measured at 50% of discovered NDR throughput rate. Latency is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

VPP packet latency in 1t1c setup (1thread, 1core) is presented in the graph below.

¹²⁵ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1804




Figure 1. VPP 1thread 1core - packet latency for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹²⁶.

VPP packet latency in 2t2c setup (2thread, 2core) is presented in the graph below.

¹²⁶ https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls1804




Figure 2. VPP 2threads 2cores - packet latency for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹²⁷.

2.6.9 Container Orchestrated Topologies

This section includes summary graphs of VPP Phy-to-Phy packet latency with CContainer Orchestrated Topologies measured at 50% of discovered NDR throughput rate. Latency is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

VPP packet latency in 1t1c setup (1thread, 1core) is presented in the graph below.

¹²⁷ https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls1804




Figure 1. VPP 1thread 1core - packet latency for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹²⁸.

VPP packet latency in 2t2c setup (2thread, 2core) is presented in the graph below.

¹²⁸ https://git.fd.io/csit/tree/tests/kubernetes/perf/container_memif?h=rls1804




Figure 2. VPP 2threads 2cores - packet latency for Phy-to-Phy L2 Ethernet Switching (base).

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹²⁹.

2.6.10 IPsec Crypto HW: IP4 Routed-Forwarding

This section includes summary graphs of VPP Phy-to-Phy packet latency with IPsec encryption used in combination with IPv4 routed-forwarding, with latency measured at 50% of discovered NDR throughput rate. VPP IPsec encryption is accelerated using DPDK cryptodev library driving Intel Quick Assist (QAT) crypto PCIe hardware cards. Latency is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

VPP packet latency in 1t1c setup (1thread, 1core) is presented in the graph below.

¹²⁹ https://git.fd.io/csit/tree/tests/kubernetes/perf/container_memif?h=rls1804




Figure 1. VPP 1thread 1core - packet latency for Phy-to-Phy IPsec HW with IPv4 Routed-Forwarding.

CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹³⁰.

VPP packet latency in 2t2c setup (2thread, 2core) is presented in the graph below.

¹³⁰ <https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls1804>




Figure 2. VPP 2threads 2cores - packet latency for Phy-to-Phy IPSec HW with IPv4 Routed-Forwarding. CSIT source code for the test cases used for above plots can be found in [CSIT git repository](#)¹³¹.

2.7 VPP HTTP Server Performance Results

Plotted results are generated by multiple executions of the same CSIT performance tests across three physical testbeds within LF FD.io labs. To provide a descriptive summary view, Box-and-Whisker plots are used to display variation in measured performance values, without making any assumptions of the underlying statistical distribution.

For each plotted test case, Box-and-Whisker plots show the quartiles (Min, 1st quartile / 25th percentile, 2nd quartile / 50th percentile / mean, 3rd quartile / 75th percentile, Max) across collected data set (data set size stated in the note below). Outliers are plotted as individual points. Min and max values are plotted as bottom and top Whiskers respectively. 2nd and 3rd quartiles are plotted as bottom and top edge of the box. If multiple samples match only two values, and all samples fall between them, then no whiskers are plotted. If all samples have the same value, only a horizontal line is plotted.

¹³¹ <https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls1804>

Note: Data sources for reported test results: i) [FD.io test executor vpp performance jobs](#)¹³², ii) archived FD.io jobs test result [output files](#).

2.7.1 Connections per second




Figure 1. VPP HTTP Server Performance - Connections per Second

CSIT source code for the test cases used for above plots can be found in CSIT git repository:

```
$ cd tests/vpp/perf/tcp  
$ grep -HE '(1t1c|2t2c|4t4c)-ethip4tcphttp-httpserver-cps' *
```

¹³² <https://jenkins.fd.io/view/csit/job/csit-vpp-perf-1804-all>

2.7.2 Requests per second




Figure 2. VPP HTTP Server Performance - Requests per Second

CSIT source code for the test cases used for above plots can be found in CSIT git repository:

```
$ cd tests/vpp/perf/tcp  
$ grep -HE '(1t1c|2t2c|4t4c)-ethip4tcphttp-httpserver-rps' *
```

2.8 Experimental: MDR Search

Multiple Drop Rate (MDR) Search is a new search algorithm implemented in FD.io CSIT project. MDR discovers multiple packet throughput rates in a single search, with each rate associated with a distinct Packet Loss Ratio (PLR) criteria.

Two throughput measurements used in FD.io CSIT are Non-Drop Rate (NDR, with zero packet loss, PLR=0) and Partial Drop Rate (PDR, with packet loss rate not greater than the configured non-zero PLR). MDR search discovers NDR and PDR in a single pass reducing required execution time compared to separate binary searches for NDR and PDR. MDR reduces execution time even further by relying on shorter

trial durations of intermediate steps, with only the final measurements conducted at the specified final trial duration. This results in the shorter overall search execution time when compared to a standard NDR/PDR binary search, while guaranteeing the same or similar results.

If needed, MDR can be easily adopted to discover more throughput rates with different pre-defined PLRs.

Note: All throughput rates are *always* bi-directional aggregates of two equal (symmetric) uni-directional packet rates received and reported by an external traffic generator.

2.8.1 Overview

The main properties of MDR search:

- MDR is a duration aware multi-phase multi-rate search algorithm.
 - Initial phase determines promising starting interval for the search.
 - Intermediate phases progress towards defined final search criteria.
 - Final phase executes measurements according to the final search criteria.
- *Initial phase:*
 - Uses link rate as a starting transmit rate and discovers the Maximum Receive Rate (MRR) used as an input to the first intermediate phase.
- *Intermediate phases:*
 - Start with initial trial duration (in the first phase) and converge geometrically towards the final trial duration (in the final phase).
 - Track two values for NDR and two for PDR.
 - * The values are called (NDR or PDR) lower_bound and upper_bound.
 - * Each value comes from a specific trial measurement (most recent for that transmit rate), and as such the value is associated with that measurement's duration and loss.
 - * A bound can be invalid, for example if NDR lower_bound has been measured with nonzero loss.
 - * Invalid bounds are not real boundaries for the searched value, but are needed to track interval widths.
 - * Valid bounds are real boundaries for the searched value.
 - * Each non-initial phase ends with all bounds valid.
 - Start with a large (lower_bound, upper_bound) interval width and geometrically converge towards the width goal (measurement resolution) of the phase. Each phase halves the previous width goal.
 - Use internal and external searches:
 - * External search - measures at transmit rates outside the (lower_bound, upper_bound) interval. Activated when a bound is invalid, to search for a new valid bound by doubling the interval width. It is a variant of **exponential search**¹³³.
 - * Internal search - **binary search**¹³⁴, measures at transmit rates within the (lower_bound, upper_bound) valid interval, halving the interval width.
- *Final phase* is executed with the final test trial duration, and the final width goal that determines resolution of the overall search. Intermediate phases together with the final phase are called non-initial phases.

¹³³ https://en.wikipedia.org/wiki/Exponential_search

¹³⁴ https://en.wikipedia.org/wiki/Binary_search

The main benefits of MDR search vs. binary search include:

- In general MDR is likely to execute more search trials overall, but less trials at a set final duration.
- In well behaving cases it greatly reduces (>50%) the overall duration compared to a single PDR (or NDR) binary search duration, while finding multiple drop rates.
- In all cases MDR yields the same or similar results to binary search.
- Note: both binary search and MDR are susceptible to reporting non-repeatable results across multiple runs for very bad behaving cases.

Caveats:

- Worst case MDR can take longer than a binary search e.g. in case of drastic changes in behaviour for trials at varying durations.

2.8.2 Search Implementation

Following is a brief description of the current MDR search implementation in FD.io CSIT.

Input Parameters

1. *maximum_transmit_rate* - maximum packet transmit rate to be used by external traffic generator, limited by either the actual Ethernet link rate or traffic generator NIC model capabilities. Sample defaults: 2 * 14.88 Mpps for 64B 10GE link rate, 2 * 18.75 Mpps for 64B 40GE NIC maximum rate.
2. *minimum_transmit_rate* - minimum packet transmit rate to be used for measurements. MDR search fails if lower transmit rate needs to be used to meet search criteria. Default: 2 * 10 kpps (could be higher).
3. *final_trial_duration* - required trial duration for final rate measurements. Default: 30 sec.
4. *initial_trial_duration* - trial duration for initial MDR phase. Default: 1 sec.
5. *final_relative_width* - required measurement resolution expressed as (lower_bound, upper_bound) interval width relative to upper_bound. Default: 0.5%.
6. *packet_loss_ratio* - maximum acceptable PLR search criteria for PDR measurements. Default: 0.5%.
7. *number_of_intermediate_phases* - number of phases between the initial phase and the final phase. Impacts the overall MDR search duration. Less phases are required for well behaving cases, more phases may be needed to reduce the overall search duration for worse behaving cases. Default (2). (Value chosen based on limited experimentation to date. More experimentation needed to arrive to clearer guidelines.)

Initial phase

1. First trial measures at maximum rate and discovers MRR.
 - (a) *in*: trial_duration = initial_trial_duration.
 - (b) *in*: offered_transmit_rate = maximum_transmit_rate.
 - (c) *do*: single trial.
 - (d) *out*: measured loss ratio.
 - (e) *out*: mrr = measured receive rate.
2. Second trial measures at MRR and discovers MRR2.
 - (a) *in*: trial_duration = initial_trial_duration.
 - (b) *in*: offered_transmit_rate = MRR.

- (c) *do*: single trial.
 - (d) *out*: measured loss ratio.
 - (e) *out*: $mrr2$ = measured receive rate.
3. Third trial measures at MRR2.
- (a) *in*: $trial_duration = initial_trial_duration$.
 - (b) *in*: $offered_transmit_rate = MRR2$.
 - (c) *do*: single trial.
 - (d) *out*: measured loss ratio.

Non-initial phases

1. Main loop:
- (a) *in*: $trial_duration$ for the current phase. Set to $initial_trial_duration$ for the first intermediate phase; to $final_trial_duration$ for the final phase; or to the element of interpolating geometric sequence for other intermediate phases. For example with two intermediate phases, $trial_duration$ of the second intermediate phase is the geometric average of $initial_trial_duration$ and $final_trial_duration$.
 - (b) *in*: $relative_width_goal$ for the current phase. Set to $final_relative_width$ for the final phase; doubled for each preceding phase. For example with two intermediate phases, the first intermediate phase uses quadruple of $final_relative_width$ and the second intermediate phase uses double of $final_relative_width$.
 - (c) *in*: $ndr_interval$, $pdr_interval$ from the previous main loop iteration or the previous phase. If the previous phase is the initial phase, both intervals have $lower_bound = MRR2$, $upper_bound = MRR$. Note that the initial phase is likely to create intervals with invalid bounds.
 - (d) *do*: According to the procedure described in point 2, either exit the phase (by jumping to 1.g.), or prepare new transmit rate to measure with.
 - (e) *do*: Perform the trial measurement at the new transmit rate and $trial_duration$, compute its loss ratio.
 - (f) *do*: Update the bounds of both intervals, based on the new measurement. The actual update rules are numerous, as NDR external search can affect PDR interval and vice versa, but the result agrees with rules of both internal and external search. For example, any new measurement below an invalid $lower_bound$ becomes the new $lower_bound$, while the old measurement (previously acting as the invalid $lower_bound$) becomes a new and valid $upper_bound$. Go to next iteration (1.c.), taking the updated intervals as new input.
 - (g) *out*: current $ndr_interval$ and $pdr_interval$. In the final phase this is also considered to be the result of the whole search. For other phases, the next phase loop is started with the current results as an input.
2. New transmit rate (or exit) calculation (for 1.d.):
- If there is an invalid bound then prepare for external search:
 - If the most recent measurement at NDR $lower_bound$ transmit rate had the loss higher than zero, then the new transmit rate is NDR $lower_bound$ decreased by two NDR interval widths.
 - Else, if the most recent measurement at PDR $lower_bound$ transmit rate had the loss higher than PLR, then the new transmit rate is PDR $lower_bound$ decreased by two PDR interval widths.
 - Else, if the most recent measurement at NDR $upper_bound$ transmit rate had no loss, then the new transmit rate is NDR $upper_bound$ increased by two NDR interval widths.

- Else, if the most recent measurement at PDR upper_bound transmit rate had the loss lower or equal to PLR, then the new transmit rate is PDR upper_bound increased by two PDR interval widths.
- Else, if NDR (or PDR) interval does not meet the current phase width goal, prepare for internal search. The new transmit rate is $(\text{lower bound} + \text{upper bound}) / 2$. It does not matter much which interval is investigated first. The current implementation starts with NDR, unless PDR interval is wider (but always preferring NDR is slightly better).
- Else, if some bound has still only been measured at a lower duration, prepare to re-measure at the current duration (and the same transmit rate). The order of priorities is:
 - NDR lower_bound,
 - PDR lower_bound,
 - NDR upper_bound,
 - PDR upper_bound.
- Else, do not prepare any new rate, to exit the phase. This ensures that at the end of each non-initial phase all intervals are valid, narrow enough, and measured at current phase trial duration.

2.8.3 Implementation Deviations

This document so far has been describing a simplified version of MDR search algorithm. The full algorithm as implemented contains additional logic, which makes some of the details (but not general ideas) above incorrect. Here is a short description of the additional logic as a list of principles, explaining their main differences from (or additions to) the simplified description, but without detailing their mutual interaction.

1. *Logarithmic transmit rate.* In order to better fit the relative width goal, the interval doubling and halving is done differently. For example, the middle of 2 and 8 is 4, not 5.
2. *Optimistic maximum rate.* The increased rate is never higher than the maximum rate. Upper bound at that rate is always considered valid.
3. *Pessimistic minimum rate.* The decreased rate is never lower than the minimum rate. If a lower bound at that rate is invalid, a phase stops refining the interval further (until it gets re-measured).
4. *Conservative interval updates.* Measurements above current upper bound never update a valid upper bound, even if drop ratio is low. Measurements below current lower bound always update any lower bound if drop ratio is high.
5. *Ensure sufficient interval width.* Narrow intervals make external search take more time to find a valid bound. If the new transmit increased or decreased rate would result in width less than the current goal, increase/decrease more. This can happen if measurement for the other interval makes the current interval too narrow. Similarly, take care the measurements in the initial phase create wide enough interval.
6. *Timeout for bad cases.* The worst case for MDR search is when each phase converges to intervals way different than the results of the previous phase. Rather than suffer total search time several times larger than pure binary search, the implemented tests fail themselves when the search takes too long (given by argument *timeout*).

2.8.4 Test Effectiveness Comparison

Introduction

CSIT release 1804 contains two test suites that use the new MDR search to enable comparison against existing CSIT NDR and PDR binary searches. The suites got chosen based on the level of consistency of their historical NDR/PDR results:

1. *10Ge2P1X520-Ethip4-Ip4Base-Ndrpdr* - yielding very consistent binary search results.
2. *10Ge2P1X520-Eth-L2Bdbasemaclrn-Eth-2Vhostvr1024-1Vm-Ndrpdr* - yielding somewhat inconsistent results.

Here “inconsistent” means the values found differ between runs, even though the setup and the test are exactly the same.

The search part of CSIT binary search tests requires a single 5-second warmup and each trial measurement is set to 10 seconds.

New tests with MDR search do not have any warmup, as initial measurements are not critical to the final result.

Fairness of the following comparison has been achieved by setting MDR final relative width to values causing the width to match the binary NDR/PDR result. Each search algorithm has been run with three different (final) trial durations: 10s, 30s and 60s.

Tables below compares overall test duration between the search tests. For simplicity only data for single thread 64B packet tests is listed, as it takes the longest in all cases.

Data in tables is based on result of 6 runs.

Tables

Table 1: Table 1. Search part of test duration.

| Duration+-avgdev [s] | IP4 10s | IP4 30s | IP4 60s | Vhost 10s | Vhost 30s | Vhost 60s |
|----------------------|------------|-------------|-------------|------------|-------------|-------------|
| MDR (both intervals) | 50.8+-1.2 | 109.0+-10.0 | 202.8+-11.7 | 80.5+-9.0 | 201.9+-20.6 | 474.9+-58.2 |
| NDR binary | 98.9+-0.1 | 278.6+-0.1 | 548.8+-0.1 | 119.8+-0.1 | 339.3+-0.1 | 669.6+-0.2 |
| PDR binary | 98.9+-0.1 | 278.6+-0.1 | 548.8+-0.1 | 119.7+-0.1 | 339.3+-0.1 | 669.5+-0.1 |
| NDR+PDR sum | 197.8+-0.1 | 557.2+-0.2 | 1097.6+-0.1 | 239.5+-0.1 | 678.7+-0.1 | 1339.2+-0.1 |

Note: Here “avgdev” is the estimated difference between the average duration computed from the limited sample and a true average duration as its hypothetical limit for infinite samples. To get the usual “standard deviation” of duration, “avgdev” has to be multiplied by the square root of the number of samples. For the subtle details see [estimation of standard deviation](#)¹³⁵, we used zero ACF and $c4=1$.

Table 2: Table 2. MDR duration as percentage of NDR duration.

| Fraction+-stdev [%] | IP4 10s | IP4 30s | IP4 60s | Vhost 10s | Vhost 30s | Vhost 60s |
|--------------------------------------|-----------|-----------|-----------|-----------|-----------|-----------|
| MDR duration divided by NDR duration | 51.4+-1.2 | 39.1+-3.6 | 37.0+-2.1 | 67.2+-7.5 | 59.5+-6.1 | 70.9+-8.7 |

Note: Here “stdev” is standard deviation as computed by the [simplified error propagation formula](#)¹³⁶.

¹³⁵ https://en.wikipedia.org/wiki/Unbiased_estimation_of_standard_deviation

¹³⁶ https://en.wikipedia.org/wiki/Propagation_of_uncertainty#Simplification

Conclusions

In consistent tests, MDR is on average more than 50% faster than a single NDR binary search (even though MDR also detects PDR).

One exception is 10 second final trial duration, where MDR is (only) almost 50% faster than NDR binary search. Most probably presence of 2 intermediate phases (instead of just 1) hurts there.

In inconsistent tests MDR is still somewhat faster than NDR binary search, but it is not by 50%, and it is hard to quantify as MDR samples have wildly varying durations.

2.8.5 Search Time Graphs

The following graphs were created from the data gathered from comparison runs, for the vhost tests. The vertical axis has always the same values, zoomed into the interesting part of the search space. The faint blue vertical lines separate the phases of MDR search. The bound lines are sloped just to help locate the previous value, in reality the bounds are updated instantly at the end of the measurement.

The graphs do not directly show when a particular bound is invalid. However this can be gleaned indirectly by identifying that the measurement does not satisfy that bound's validity conditions (see point 2a). Also, the external search follows, and the measurement previously acting as an invalid upper or lower bound starts acting instead as a valid lower or upper bound, respectively.

The following three graphs are for MDR with 10 second final trial duration, showing different behavior in this inconsistent test, and different amount of "work" done by each phase. Also the horizontal axis has the same scaling here.

The next graph is for MDR with 60 second final trial duration, to showcase the final phase takes the most of the overall search time. The scaling of the horizontal axis is different.

Finally, here are two graphs showing NDR and PDR binary searches. The trial duration is again 60 seconds, but scaling of horizontal axis is once again different. This shows the binary search spends most time measuring outside the interesting rate region.

2.9 Test Environment

CSIT performance tests are executed on the three identical physical testbeds hosted by LF for FD.io project. Each testbed consists of two servers acting as Systems Under Test (SUT) and one server acting as Traffic Generator (TG).

2.9.1 Server Specification and Configuration

Complete specification and configuration of compute servers used in CSIT physical testbeds is maintained on wiki page [CSIT LF Testbeds](https://wiki.fd.io/view/CSIT/CSIT_LF_testbed)¹³⁷.

¹³⁷ https://wiki.fd.io/view/CSIT/CSIT_LF_testbed

2.9.2 SUT Configuration

Host configuration

All hosts are Cisco UCS C240-M4 (2x Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz, 18c, 512GB RAM)

```
$ lscpu
Architecture:          x86_64
CPU op-mode(s):       32-bit, 64-bit
Byte Order:           Little Endian
CPU(s):               36
On-line CPU(s) list:  0-35
Thread(s) per core:   1
Core(s) per socket:   18
Socket(s):            2
NUMA node(s):         2
Vendor ID:            GenuineIntel
CPU family:           6
Model:                63
Model name:           Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz
Stepping:             2
CPU MHz:              2294.249
BogoMIPS:             4589.82
Virtualization:       VT-x
L1d cache:           32K
L1i cache:           32K
L2 cache:            256K
L3 cache:            46080K
NUMA node0 CPU(s):   0-17
NUMA node1 CPU(s):   18-35
Flags:                fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36_
↳ clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx pdpe1gb rdtscp lm constant_tsc arch_
↳ perfmon pebs bts rep_good nopl xtopology nonstop_tsc aperfmperf eagerfpu pni pclmulqdq dtes64_
↳ monitor ds_cpl vmx smx est tm2 ssse3 sdbg fma cx16 xtpr pdcm pcid dca sse4_1 sse4_2 x2apic movbe_
↳ popcnt tsc_deadline_timer aes xsave avx f16c rdrand lahf_lm abm epb tpr_shadow vnmi flexpriority_
↳ ept vpid fsgsbase tsc_adjust bmi1 avx2 smep bmi2 erms invpcid cqm xsaveopt cqm_llc cqm_occup_llc_
↳ dtherm arat pln pts
```

BIOS settings

```
C240 /bios # show advanced detail
Set-up parameters:
  Intel(R) VT-d ATS Support: Enabled
  Adjacent Cache Line Prefetcher: Enabled
  All Onboard LOM Ports: Enabled
  Altitude: 300 M
  Bits per second: 115200
  Power Technology: Performance
  Channel Interleaving: Auto
  Intel(R) VT-d Coherency Support: Disabled
  Console Redirection: COM 0
  Number of Enabled Cores: All
  Energy Performance: Performance
  CPU Performance: Enterprise
  DCU IP Prefetcher: Enabled
  DCU Streamer Prefetch: Enabled
  Demand Scrub: Enabled
  Direct Cache Access Support: Auto
  Enhanced Intel Speedstep(R) Tec: Disabled
  Execute Disable: Enabled
  Flow Control: None
  Hardware Prefetcher: Enabled
  Intel(R) Hyper-Threading Techno: Disabled
```

(continues on next page)

```
Intel(R) Turbo Boost Technology: Disabled
Intel(R) VT: Enabled
Intel(R) VT-d: Enabled
Intel(R) Interrupt Remapping: Enabled
Legacy USB Support: Enabled
Extended APIC: XAPIC
LOM Port 1 OptionROM: Enabled
LOM Port 2 OptionROM: Enabled
MMIO above 4GB: Enabled
NUMA: Enabled
PCI ROM CLP: Disabled
Package C State Limit: C6 Retention
Intel(R) Pass Through DMA: Disabled
Patrol Scrub: Enabled
xHCI Mode: Disabled
All PCIe Slots OptionROM: Enabled
PCIe Slot:1 OptionROM: Disabled
PCIe Slot:2 OptionROM: Disabled
PCIe Slot:3 OptionROM: Disabled
PCIe Slot:4 OptionROM: Disabled
PCIe Slot:5 OptionROM: Disabled
PCIe Slot:6 OptionROM: Disabled
PCIe Slot:HBA Link Speed: GEN3
PCIe Slot:HBA OptionROM: Enabled
PCIe Slot:MLOM OptionROM: Enabled
PCIe Slot:N1 OptionROM: Enabled
PCIe Slot:N2 OptionROM: Enabled
Processor Power state C1 Enhanc: Disabled
Processor C3 Report: Disabled
Processor C6 Report: Disabled
P-STATE Coordination: HW ALL
Putty KeyPad: ESCN
Energy Performance Tuning: BIOS
QPI Link Frequency Select: Auto
QPI Snoop Mode: Home Snoop
Rank Interleaving: Auto
Redirection After BIOS POST: Always Enable
PCH SATA Mode: AHCI
Select Memory RAS: Maximum Performance
SR-IOV Support: Enabled
Terminal Type: VT100
Port 60/64 Emulation: Enabled
Workload Configuration: Balanced
CDN Support for VIC: Disabled
Out-of-Band Management: Disabled
```

NIC models and placement

In addition to CIMC and Management, each TG has 4x Intel X710 10GB NIC (=8 ports) and 2x Intel XL710 40GB NIC (=4 ports), whereas each SUT has:

- 1x Intel X520 NIC (10GB, 2 ports),
- 1x Cisco VIC 1385 (40GB, 2 ports),
- 1x Intel XL710 NIC (40GB, 2 ports),
- 1x Intel X710 NIC (10GB, 2 ports),
- 1x Cisco VIC 1227 (10GB, 2 ports).

This allows for a total of five ring topologies, each using ports on specific NIC model, enabling per NIC model benchmarking.

- 0a:00.0 Ethernet controller: Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01) Subsystem: Intel Corporation Ethernet Server Adapter X520-2
- 0a:00.1 Ethernet controller: Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01) Subsystem: Intel Corporation Ethernet Server Adapter X520-2
- 06:00.0 Ethernet controller: Cisco Systems Inc VIC Ethernet NIC (rev a2) Subsystem: Cisco Systems Inc VIC 1227 PCIe Ethernet NIC
- 07:00.0 Ethernet controller: Cisco Systems Inc VIC Ethernet NIC (rev a2) Subsystem: Cisco Systems Inc VIC 1227 PCIe Ethernet NIC
- 13:00.0 Ethernet controller: Cisco Systems Inc VIC Ethernet NIC (rev a2) Subsystem: Cisco Systems Inc VIC 1385 PCIe Ethernet NIC
- 15:00.0 Ethernet controller: Cisco Systems Inc VIC Ethernet NIC (rev a2) Subsystem: Cisco Systems Inc VIC 1385 PCIe Ethernet NIC
- 85:00.0 Ethernet controller: Intel Corporation Ethernet Controller XL710 for 40GbE QSFP+ (rev 01) Subsystem: Intel Corporation Ethernet Converged Network Adapter XL710-Q2
- 85:00.1 Ethernet controller: Intel Corporation Ethernet Controller XL710 for 40GbE QSFP+ (rev 01) Subsystem: Intel Corporation Ethernet Converged Network Adapter XL710-Q2
- 87:00.0 Ethernet controller: Intel Corporation Ethernet Controller X710 for 10GbE SFP+ (rev 01) Subsystem: Intel Corporation Ethernet Converged Network Adapter X710-2
- 87:00.1 Ethernet controller: Intel Corporation Ethernet Controller X710 for 10GbE SFP+ (rev 01) Subsystem: Intel Corporation Ethernet Converged Network Adapter X710-2

2.9.3 SUT Configuration - Host OS Linux

Software details (OS, configuration) of physical testbeds are maintained on wiki page [CSIT LF Testbeds](#)¹³⁸.

System provisioning is done by combination of PXE boot unattended install and [Ansible](#)¹³⁹ described in [CSIT Testbed Setup](#)¹⁴⁰.

Below a subset of the running configuration:

```
$ lsb_release -a
No LSB modules are available.
Distributor ID: Ubuntu
Description: Ubuntu 16.04.1 LTS
Release: 16.04
Codename: xenial
```

```
$ cat /sys/devices/system/node/node*/meminfo
Node 0 MemTotal: 264048168 kB
Node 0 MemFree: 257730716 kB
Node 0 MemUsed: 6317452 kB
Node 0 Active: 1079920 kB
Node 0 Inactive: 470064 kB
Node 0 Active(anon): 674772 kB
Node 0 Inactive(anon): 248572 kB
Node 0 Active(file): 405148 kB
Node 0 Inactive(file): 221492 kB
Node 0 Unevictable: 0 kB
Node 0 Mlocked: 0 kB
Node 0 Dirty: 12 kB
Node 0 Writeback: 0 kB
```

(continues on next page)

¹³⁸ https://wiki.fd.io/view/CSIT/CSIT_LF_testbed

¹³⁹ <https://www.ansible.com>

¹⁴⁰ <https://git.fd.io/csit/tree/resources/tools/testbed-setup/README.md?h=rls1804>

(continued from previous page)

```

Node 0 FilePages:      1270432 kB
Node 0 Mapped:         20116 kB
Node 0 AnonPages:     279548 kB
Node 0 Shmem:         643796 kB
Node 0 KernelStack:   3376 kB
Node 0 PageTables:    1316 kB
Node 0 NFS_Unstable:   0 kB
Node 0 Bounce:         0 kB
Node 0 WritebackTmp:  0 kB
Node 0 Slab:           80428 kB
Node 0 SReclaimable:  38288 kB
Node 0 SUnreclaim:   42140 kB
Node 0 AnonHugePages: 270336 kB
Node 0 HugePages_Total: 2048
Node 0 HugePages_Free: 2048
Node 0 HugePages_Surp: 0
Node 1 MemTotal:      264237596 kB
Node 1 MemFree:       256758976 kB
Node 1 MemUsed:       7478620 kB
Node 1 Active:        1746052 kB
Node 1 Inactive:      981104 kB
Node 1 Active(anon):  1272936 kB
Node 1 Inactive(anon): 849968 kB
Node 1 Active(file):  473116 kB
Node 1 Inactive(file): 131136 kB
Node 1 Unevictable:   0 kB
Node 1 Mlocked:       0 kB
Node 1 Dirty:         0 kB
Node 1 Writeback:     0 kB
Node 1 FilePages:    2715284 kB
Node 1 Mapped:        75928 kB
Node 1 AnonPages:    11920 kB
Node 1 Shmem:        2111036 kB
Node 1 KernelStack:  2576 kB
Node 1 PageTables:   1348 kB
Node 1 NFS_Unstable:  0 kB
Node 1 Bounce:        0 kB
Node 1 WritebackTmp:  0 kB
Node 1 Slab:          90604 kB
Node 1 SReclaimable: 55384 kB
Node 1 SUnreclaim:   35220 kB
Node 1 AnonHugePages: 6144 kB
Node 1 HugePages_Total: 2048
Node 1 HugePages_Free: 2048
Node 1 HugePages_Surp: 0

```

Kernel boot parameters used in CSIT performance testbeds

- **isolcpus=<cpu number>-<cpu number>** used for all cpu cores apart from first core of each socket used for running VPP worker threads and Qemu/LXC processes <https://www.kernel.org/doc/Documentation/admin-guide/kernel-parameters.txt>
- **intel_pstate=disable** - [X86] Do not enable intel_pstate as the default scaling driver for the supported processors. Intel P-State driver decide what P-state (CPU core power state) to use based on requesting policy from the cpufreq core. [X86 - Either 32-bit or 64-bit x86] <https://www.kernel.org/doc/Documentation/cpu-freq/intel-pstate.txt>
- **nohz_full=<cpu number>-<cpu number>** - [KNL,BOOT] In kernels built with CONFIG_NO_HZ_FULL=y, set the specified list of CPUs whose tick will be stopped whenever possible. The boot CPU will be forced outside the range to maintain the timekeeping. The CPUs in this range must also be included in the rcu_nocbs= set. Specifies the adaptive-ticks CPU cores, causing kernel to avoid sending scheduling-clock interrupts to listed cores as long as they have a

single runnable task. [KNL - Is a kernel start-up parameter, SMP - The kernel is an SMP kernel]. https://www.kernel.org/doc/Documentation/timers/NO_HZ.txt

- **rcu_nocbs** - [KNL] In kernels built with CONFIG_RCU_NOCB_CPU=y, set the specified list of CPUs to be no-callback CPUs, that never queue RCU callbacks (read-copy update). <https://www.kernel.org/doc/Documentation/admin-guide/kernel-parameters.txt>

Applied command line boot parameters:

```
$ cat /proc/cmdline
BOOT_IMAGE=/vmlinuz-4.4.0-72-generic root=UUID=35ea11e4-e44f-4f67-8cbe-12f09c49ed90 ro isolcpus=1-
↳17,19-35 nohz_full=1-17,19-35 rcu_nocbs=1-17,19-35 intel_pstate=disable console=tty0_
↳console=ttyS0,115200n8
```

Mount listing

```
$ cat /proc/mounts
sysfs /sys sysfs rw,nosuid,nodev,noexec,relatime 0 0
proc /proc proc rw,nosuid,nodev,noexec,relatime 0 0
udev /dev devtmpfs rw,nosuid,relatime,size=264125468k,nr_inodes=66031367,mode=755 0 0
devpts /dev/pts devpts rw,nosuid,noexec,relatime,gid=5,mode=620,ptmxmode=000 0 0
tmpfs /run tmpfs rw,nosuid,noexec,relatime,size=52828580k,mode=755 0 0
/dev/sda2 / ext4 rw,relatime,errors=remount-ro,data=ordered 0 0
securityfs /sys/kernel/security securityfs rw,nosuid,nodev,noexec,relatime 0 0
tmpfs /dev/shm tmpfs rw,nosuid,nodev 0 0
tmpfs /run/lock tmpfs rw,nosuid,nodev,noexec,relatime,size=5120k 0 0
tmpfs /sys/fs/cgroup tmpfs ro,nosuid,nodev,noexec,mode=755 0 0
cgroup /sys/fs/cgroup/systemd cgroup rw,nosuid,nodev,noexec,relatime,xattr,release_agent=/lib/
↳systemd/systemd-cgroups-agent,name=systemd 0 0
pstore /sys/fs/pstore pstore rw,nosuid,nodev,noexec,relatime 0 0
cgroup /sys/fs/cgroup/freezer cgroup rw,nosuid,nodev,noexec,relatime,freezer 0 0
cgroup /sys/fs/cgroup/net_cls,net_prio cgroup rw,nosuid,nodev,noexec,relatime,net_cls,net_prio 0 0
cgroup /sys/fs/cgroup/cpu,cpuacct cgroup rw,nosuid,nodev,noexec,relatime,cpu,cpuacct 0 0
cgroup /sys/fs/cgroup/memory cgroup rw,nosuid,nodev,noexec,relatime,memory 0 0
cgroup /sys/fs/cgroup/blkio cgroup rw,nosuid,nodev,noexec,relatime,blkio 0 0
cgroup /sys/fs/cgroup/perf_event cgroup rw,nosuid,nodev,noexec,relatime,perf_event 0 0
cgroup /sys/fs/cgroup/devices cgroup rw,nosuid,nodev,noexec,relatime,devices 0 0
cgroup /sys/fs/cgroup/cpuset cgroup rw,nosuid,nodev,noexec,relatime,cpuset,clone_children 0 0
cgroup /sys/fs/cgroup/hugetlb cgroup rw,nosuid,nodev,noexec,relatime,hugetlb 0 0
cgroup /sys/fs/cgroup/pids cgroup rw,nosuid,nodev,noexec,relatime,pids 0 0
systemd-1 /proc/sys/fs/binfmt_misc autofs rw,relatime,fd=26,pgrp=1,timeout=0,minproto=5,maxproto=5,
↳direct 0 0
hugetlbfs /dev/hugepages hugetlbfs rw,relatime 0 0
debugfs /sys/kernel/debug debugfs rw,relatime 0 0
mqueue /dev/mqueue mqueue rw,relatime 0 0
tracefs /sys/kernel/debug/tracing tracefs rw,relatime 0 0
fusectl /sys/fs/fuse/connections fusectl rw,relatime 0 0
/dev/sda1 /boot ext4 rw,relatime,data=ordered 0 0
none /mnt/huge hugetlbfs rw,relatime,pagesize=2048k 0 0
lxcfs /var/lib/lxcfs fuse.lxcfs rw,nosuid,nodev,relatime,user_id=0,group_id=0,allow_other 0 0
```

Package listing

```
$ dpkg -l | grep '^ii' | awk '{print $2 ": " $3}'
accountsservice: 0.6.40-2ubuntu11.1
acl: 2.2.52-3
adduser: 3.113+nmu3ubuntu4
apparmor: 2.10.95-0ubuntu2.6
apt: 1.2.12~ubuntu16.04.1
apt-transport-https: 1.2.24
apt-utils: 1.2.12~ubuntu16.04.1
aufs-tools: 1:3.2+20130722-1.1ubuntu1
autoconf: 2.69-9
```

(continues on next page)

(continued from previous page)

```
automake: 1:1.15-4ubuntu1
autotools-dev: 20150820.1
base-files: 9.4ubuntu4.2
base-passwd: 3.5.39
bash: 4.3-14ubuntu1.1
binutils: 2.26.1-1ubuntu1~16.04.3
bridge-utils: 1.5-9ubuntu1
bsdutils: 1:2.27.1-6ubuntu3.1
build-essential: 12.1ubuntu2
busybox-initramfs: 1:1.22.0-15ubuntu1
busybox-static: 1:1.22.0-15ubuntu1
bzip2: 1.0.6-8
busybox-static: 1:1.22.0-15ubuntu1
bzip2: 1.0.6-8
ca-certificates: 20160104ubuntu1
ca-certificates-java: 20160321
cgroup-bin: 0.41-7ubuntu1
cgroup-lite: 1.11
cgroup-tools: 0.41-7ubuntu1
cloud-image-utils: 0.27-0ubuntu24
console-setup: 1.108ubuntu15.2
console-setup-linux: 1.108ubuntu15.2
corekeeper: 1.6
coreutils: 8.25-2ubuntu2
cpio: 2.11+dfsg-5ubuntu1
cpp: 4:5.3.1-1ubuntu1
cpp-5: 5.4.0-6ubuntu1~16.04.2
cpu-checker: 0.7-0ubuntu7
cpufrequtils: 008-1
crda: 3.13-1
cron: 3.0pl1-128ubuntu2
crudini: 0.7-1
dash: 0.5.8-2.1ubuntu2
dbus: 1.10.6-1ubuntu3
debconf: 1.5.58ubuntu1
debconf-i18n: 1.5.58ubuntu1
debianutils: 4.7
debootstrap: 1.0.78+nmu1ubuntu1.3
dh-python: 2.20151103ubuntu1.1
diffutils: 1:3.3-3
distro-info: 0.14build1
distro-info-data: 0.28ubuntu0.1
dkms: 2.2.0.3-2ubuntu11.2
dmidecode: 3.0-2ubuntu0.1
dns-root-data: 2015052300+h+1
dnsmasq-base: 2.75-1ubuntu0.16.04.2
docker-ce: 17.09.0~ce-0~ubuntu
dpkg: 1.18.4ubuntu1.1
dpkg-dev: 1.18.4ubuntu1.1
e2fslibs:amd64: 1.42.13-1ubuntu1
e2fsprogs: 1.42.13-1ubuntu1
e2tables: 2.0.10.4-3.4ubuntu2
eject: 2.1.5+deb1+cvs20081104-13.1
ethtool: 1:4.5-1
expect: 5.45-7
fakeroot: 1.20.2-1ubuntu1
file: 1:5.25-2ubuntu1
findutils: 4.6.0+git+20160126-2
fontconfig: 2.11.94-0ubuntu1.1
fontconfig-config: 2.11.94-0ubuntu1.1
fonts-dejavu-core: 2.35-1
```

(continues on next page)

(continued from previous page)

```
fonts-dejavu-extra: 2.35-1
g++: 4:5.3.1-1ubuntu1
g++-5: 5.4.0-6ubuntu1~16.04.2
gcc: 4:5.3.1-1ubuntu1
gcc-5: 5.4.0-6ubuntu1~16.04.2
gcc-5-base:amd64: 5.4.0-6ubuntu1~16.04.2
gcc-6-base:amd64: 6.0.1-0ubuntu1
gdb: 7.11.1-0ubuntu1~16.5
gdbserver: 7.11.1-0ubuntu1~16.5
genisoimage: 9:1.1.11-3ubuntu1
gettext-base: 0.19.7-2ubuntu3
gir1.2-glib-2.0:amd64: 1.46.0-3ubuntu1
git: 1:2.7.4-0ubuntu1
git-man: 1:2.7.4-0ubuntu1
gnupg: 1.4.20-1ubuntu3.1
gpgv: 1.4.20-1ubuntu3.1
grep: 2.25-1~16.04.1
grub-common: 2.02~beta2-36ubuntu3.1
grub-gfxpayload-lists: 0.7
grub-pc: 2.02~beta2-36ubuntu3.1
grub-pc-bin: 2.02~beta2-36ubuntu3.1
grub2-common: 2.02~beta2-36ubuntu3.1
gzip: 1.6-4ubuntu1
hicolor-icon-theme: 0.15-0ubuntu1
hostname: 3.16ubuntu2
ifupdown: 0.8.10ubuntu1
init: 1.29ubuntu2
init-system-helpers: 1.29ubuntu2
initramfs-tools: 0.122ubuntu8.1
initramfs-tools-bin: 0.122ubuntu8.1
initramfs-tools-core: 0.122ubuntu8.1
initscripts: 2.88dsf-59.3ubuntu2
insserv: 1.14.0-5ubuntu3
installation-report: 2.60ubuntu1
iproute2: 4.3.0-1ubuntu3
iptables: 1.6.0-2ubuntu3
iputils-ping: 3:20121221-5ubuntu2
ipxe-qemu: 1.0.0+git-20150424.a25a16d-1ubuntu1
isc-dhcp-client: 4.3.3-5ubuntu12.1
isc-dhcp-common: 4.3.3-5ubuntu12.1
iso-codes: 3.65-1
iw: 3.17-1
java-common: 0.56ubuntu2
kbd: 1.15.5-1ubuntu4
keyboard-configuration: 1.108ubuntu15.2
klibc-utils: 2.0.4-8ubuntu1.16.04.1
kmod: 22-1ubuntu4
krb5-locales: 1.13.2+dfsg-5
kubeadm: 1.8.4-00
kubect1: 1.8.4-00
kubelet: 1.8.4-00
kubernetes-cni: 0.5.1-00
language-selector-common: 0.165.3
laptop-detect: 0.13.7ubuntu2
less: 481-2.1
libaccountsservice0:amd64: 0.6.40-2ubuntu11.1
libacl1:amd64: 2.2.52-3
libaio1:amd64: 0.3.110-2
libalgorithm-diff-perl: 1.19.03-1
libalgorithm-diff-xs-perl: 0.04-4build1
libalgorithm-merge-perl: 0.08-3
```

(continues on next page)

(continued from previous page)

```
libapparmor-perl: 2.10.95-0ubuntu2.6
libapparmor1:amd64: 2.10.95-0ubuntu2
libapr1:amd64: 1.5.2-3
libapt-inst2.0:amd64: 1.2.12~ubuntu16.04.1
libapt-pkg5.0:amd64: 1.2.12~ubuntu16.04.1
libasan2:amd64: 5.4.0-6ubuntu1~16.04.2
libasn1-8-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libasound2:amd64: 1.1.0-0ubuntu1
libasound2-data: 1.1.0-0ubuntu1
libasprintf0v5:amd64: 0.19.7-2ubuntu3
libasyncns0:amd64: 0.8-5build1
libatk1.0-0:amd64: 2.18.0-1
libatk1.0-data: 2.18.0-1
libatm1:amd64: 1:2.5.1-1.5
libatomic1:amd64: 5.4.0-6ubuntu1~16.04.2
libattr1:amd64: 1:2.4.47-2
libaudit-common: 1:2.4.5-1ubuntu2
libaudit1:amd64: 1:2.4.5-1ubuntu2
libavahi-client3:amd64: 0.6.32~rc+dfsg-1ubuntu2
libavahi-common-data:amd64: 0.6.32~rc+dfsg-1ubuntu2
libavahi-common3:amd64: 0.6.32~rc+dfsg-1ubuntu2
libbabeltrace-ctf1:amd64: 1.3.2-1
libbabeltrace1:amd64: 1.3.2-1
libblkid1:amd64: 2.27.1-6ubuntu3.1
libbluetooth3:amd64: 5.37-0ubuntu5
libboost-iostreams1.58.0:amd64: 1.58.0+dfsg-5ubuntu3.1
libboost-random1.58.0:amd64: 1.58.0+dfsg-5ubuntu3.1
libboost-system1.58.0:amd64: 1.58.0+dfsg-5ubuntu3.1
libboost-thread1.58.0:amd64: 1.58.0+dfsg-5ubuntu3.1
libbrlapi0.6:amd64: 5.3.1-2ubuntu2.1
libbsd0:amd64: 0.8.2-1
libbz2-1.0:amd64: 1.0.6-8
libc-bin: 2.23-0ubuntu3
libc-dev-bin: 2.23-0ubuntu9
libc6-dbg:amd64: 2.23-0ubuntu9
libc6-dev:amd64: 2.23-0ubuntu9
libcaca0:amd64: 0.99.beta19-2build2~gcc5.2
libcacard0:amd64: 1:2.5.0-2
libcairo2:amd64: 1.14.6-1
libcap-ng0:amd64: 0.7.7-1
libcap2:amd64: 1:2.24-12
libcap2-bin: 1:2.24-12
libcc1-0:amd64: 5.4.0-6ubuntu1~16.04.2
libcgroup1:amd64: 0.41-7ubuntu1
libcilkrts5:amd64: 5.4.0-6ubuntu1~16.04.2
libcomerr2:amd64: 1.42.13-1ubuntu1
libcpufreq0: 008-1
libcryptsetup4:amd64: 2:1.6.6-5ubuntu2
libcups2:amd64: 2.1.3-4
libcurl3-gnutls:amd64: 7.47.0-1ubuntu2.1
libdatrie1:amd64: 0.2.10-2
libdb5.3:amd64: 5.3.28-11
libdbus-1-3:amd64: 1.10.6-1ubuntu3
libdbus-glib-1-2:amd64: 0.106-1
libdebconfclient0:amd64: 0.198ubuntu1
libdevmapper1.02.1:amd64: 2:1.02.110-1ubuntu10
libdns-export162: 1:9.10.3.dfsg.P4-8ubuntu1.1
libdpkg-perl: 1.18.4ubuntu1.1
libdrm-amdgpu1:amd64: 2.4.67-1ubuntu0.16.04.2
libdrm-intel1:amd64: 2.4.67-1ubuntu0.16.04.2
libdrm-nouveau2:amd64: 2.4.67-1ubuntu0.16.04.2
```

(continues on next page)

(continued from previous page)

```
libdrm-radeon1:amd64: 2.4.67-1ubuntu0.16.04.2
libdrm2:amd64: 2.4.67-1ubuntu0.16.04.2
libedit2:amd64: 3.1-20150325-1ubuntu2
libelf1:amd64: 0.165-3ubuntu1
liberror-perl: 0.17-1.2
libestr0: 0.1.10-1
libexpat1:amd64: 2.1.0-7ubuntu0.16.04.2
libexpat1-dev:amd64: 2.1.0-7ubuntu0.16.04.2
libfakeroot:amd64: 1.20.2-1ubuntu1
libfdisk1:amd64: 2.27.1-6ubuntu3.1
libfdt1:amd64: 1.4.0+dfsg-2
libffi6:amd64: 3.2.1-4
libfile-fcntllock-perl: 0.22-3
libflac8:amd64: 1.3.1-4
libfontconfig1:amd64: 2.11.94-0ubuntu1.1
libfontenc1:amd64: 1:1.1.3-1
libfreetype6:amd64: 2.6.1-0.1ubuntu2
libfribidi0:amd64: 0.19.7-1
libfuse2:amd64: 2.9.4-1ubuntu3
libgcc-5-dev:amd64: 5.4.0-6ubuntu1~16.04.2
libgcc1:amd64: 1:6.0.1-0ubuntu1
libgcrypt20:amd64: 1.6.5-2ubuntu0.2
libgdbm3:amd64: 1.8.3-13.1
libgdk-pixbuf2.0-0:amd64: 2.32.2-1ubuntu1.2
libgdk-pixbuf2.0-common: 2.32.2-1ubuntu1.2
libgif7:amd64: 5.1.4-0.3~16.04
libgirepository-1.0-1:amd64: 1.46.0-3ubuntu1
libgl1-mesa-dri:amd64: 11.2.0-1ubuntu2.2
libgl1-mesa-glx:amd64: 11.2.0-1ubuntu2.2
libglapi-mesa:amd64: 11.2.0-1ubuntu2.2
libglb2.0-0:amd64: 2.48.1-1~ubuntu16.04.1
libglb2.0-bin: 2.48.1-1~ubuntu16.04.1
libglb2.0-data: 2.48.1-1~ubuntu16.04.1
libglb2.0-dev: 2.48.1-1~ubuntu16.04.1
libgmp10:amd64: 2:6.1.0+dfsg-2
libgnutls-openssl27:amd64: 3.4.10-4ubuntu1.1
libgnutls30:amd64: 3.4.10-4ubuntu1.1
libgomp1:amd64: 5.4.0-6ubuntu1~16.04.2
libgpg-error0:amd64: 1.21-2ubuntu1
libgraphite2-3:amd64: 1.3.6-1ubuntu1
libgssapi-krb5-2:amd64: 1.13.2+dfsg-5
libgssapi3-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libgtk2.0-0:amd64: 2.24.30-1ubuntu1.16.04.2
libgtk2.0-bin: 2.24.30-1ubuntu1.16.04.2
libgtk2.0-common: 2.24.30-1ubuntu1.16.04.2
libharfbuzz0b:amd64: 1.0.1-1ubuntu0.1
libhcrypto4-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libheimbase1-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libheimntlm0-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libhogweed4:amd64: 3.2-1
libhx509-5-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libice-dev:amd64: 2:1.0.9-1
libice6:amd64: 2:1.0.9-1
libicu55:amd64: 55.1-7
libidn11:amd64: 1.32-3ubuntu1.1
libisc-export160: 1:9.10.3.dfsg.P4-8ubuntu1.1
libiscsi2:amd64: 1.12.0-2
libisl15:amd64: 0.16.1-1
libitm1:amd64: 5.4.0-6ubuntu1~16.04.2
libjbig0:amd64: 2.1-3.1
libjpeg-turbo8:amd64: 1.4.2-0ubuntu3
```

(continues on next page)

(continued from previous page)

```
libjpeg8:amd64: 8c-2ubuntu8
libjson-c2:amd64: 0.11-4ubuntu2
libk5crypto3:amd64: 1.13.2+dfsg-5
libkeyutils1:amd64: 1.5.9-8ubuntu1
libklibc: 2.0.4-8ubuntu1.16.04.1
libkmod2:amd64: 22-1ubuntu4
libkrb5-26-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libkrb5-3:amd64: 1.13.2+dfsg-5
libkrb5support0:amd64: 1.13.2+dfsg-5
liblcms2-2:amd64: 2.6-3ubuntu2
libldap-2.4-2:amd64: 2.4.42+dfsg-2ubuntu3.1
libllvm3.8:amd64: 1:3.8-2ubuntu4
liblocale-gettext-perl: 1.07-1build1
liblsan0:amd64: 5.4.0-6ubuntu1~16.04.2
libltdl-dev:amd64: 2.4.6-0.1
libltdl7:amd64: 2.4.6-0.1
liblxc1: 2.0.7-0ubuntu1~16.04.2
liblz4-1:amd64: 0.0~r131-2ubuntu2
liblzma5:amd64: 5.1.1alpha+20120614-2ubuntu2
libmagic1:amd64: 1:5.25-2ubuntu1
libmnl0:amd64: 1.0.3-5
libmount1:amd64: 2.27.1-6ubuntu3.1
libmpc3:amd64: 1.0.3-1
libmpdec2:amd64: 2.4.2-1
libmpfr4:amd64: 3.1.4-1
libmpx0:amd64: 5.4.0-6ubuntu1~16.04.2
libncurses5:amd64: 6.0+20160213-1ubuntu1
libncursesw5:amd64: 6.0+20160213-1ubuntu1
libnetfilter-contrack3:amd64: 1.0.5-1
libnettle6:amd64: 3.2-1
libnewt0.52:amd64: 0.52.18-1ubuntu2
libnfnetwork0:amd64: 1.0.1-3
libnih-dbus1:amd64: 1.0.3-4.3ubuntu1
libnih1:amd64: 1.0.3-4.3ubuntu1
libnl-3-200:amd64: 3.2.27-1
libnl-genl-3-200:amd64: 3.2.27-1
libnspr4:amd64: 2:4.12-0ubuntu0.16.04.1
libnss3:amd64: 2:3.23-0ubuntu0.16.04.1
libnss3-nssdb: 2:3.23-0ubuntu0.16.04.1
libnuma-dev:amd64: 2.0.11-1ubuntu1
libnuma1:amd64: 2.0.11-1ubuntu1
libogg0:amd64: 1.3.2-1
libopus0:amd64: 1.1.2-1ubuntu1
libp11-kit0:amd64: 0.23.2-3
libpam-cgfs: 2.0.6-0ubuntu1~16.04.1
libpam-modules:amd64: 1.1.8-3.2ubuntu2
libpam-modules-bin: 1.1.8-3.2ubuntu2
libpam-runtime: 1.1.8-3.2ubuntu2
libpam0g:amd64: 1.1.8-3.2ubuntu2
libpango-1.0-0:amd64: 1.38.1-1
libpangocairo-1.0-0:amd64: 1.38.1-1
libpangoft2-1.0-0:amd64: 1.38.1-1
libpcap-dev: 1.7.4-2
libpcap0.8:amd64: 1.7.4-2
libpcap0.8-dev: 1.7.4-2
libpci3:amd64: 1:3.3.1-1.1ubuntu1
libpciaccess0:amd64: 0.13.4-1
libpcre16-3:amd64: 2:8.38-3.1
libpcre3:amd64: 2:8.38-3.1
libpcre3-dev:amd64: 2:8.38-3.1
libpcre32-3:amd64: 2:8.38-3.1
```

(continues on next page)

(continued from previous page)

```

libpcrecpp0v5:amd64: 2:8.38-3.1
libpcsclite1:amd64: 1.8.14-1ubuntu1.16.04.1
libper15.22:amd64: 5.22.1-9
libpixmap-1-0:amd64: 0.33.6-1
libplymouth4:amd64: 0.9.2-3ubuntu13.1
libpng12-0:amd64: 1.2.54-1ubuntu1
libpolkit-gobject-1-0:amd64: 0.105-14.1
libpopt0:amd64: 1.16-10
libprocps4:amd64: 2:3.3.10-4ubuntu2
libpthread-stubs0-dev:amd64: 0.3-4
libpulse0:amd64: 1:8.0-0ubuntu3
libpython-all-dev:amd64: 2.7.11-1
libpython-dev:amd64: 2.7.11-1
libpython-stdlib:amd64: 2.7.11-1
libpython2.7:amd64: 2.7.12-1~16.04
libpython2.7-dev:amd64: 2.7.12-1~16.04
libpython2.7-minimal:amd64: 2.7.12-1~16.04
libpython2.7-stdlib:amd64: 2.7.12-1~16.04
libpython3-stdlib:amd64: 3.5.1-3
libpython3.5:amd64: 3.5.2-2ubuntu0~16.04.1
libpython3.5-minimal:amd64: 3.5.2-2ubuntu0~16.04.1
libpython3.5-stdlib:amd64: 3.5.2-2ubuntu0~16.04.1
libquadmath0:amd64: 5.4.0-6ubuntu1~16.04.2
librados2: 10.2.2-0ubuntu0.16.04.2
librbd1: 10.2.2-0ubuntu0.16.04.2
libreadline6:amd64: 6.3-8ubuntu2
libroken18-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
librtmp1:amd64: 2.4+20151223.gitfa8646d-1build1
libsasl2-2:amd64: 2.1.26.dfsg1-14build1
libsasl2-modules:amd64: 2.1.26.dfsg1-14build1
libsasl2-modules-db:amd64: 2.1.26.dfsg1-14build1
libsdl1.2debian:amd64: 1.2.15+dfsg1-3
libseccomp2:amd64: 2.2.3-3ubuntu3
libselinux1:amd64: 2.4-3build2
libsemanage-common: 2.3-1build3
libsemanage1:amd64: 2.3-1build3
libsepol1:amd64: 2.4-2
libsigsegv2:amd64: 2.10-4
libslang2:amd64: 2.3.0-2ubuntu1
libsm-dev:amd64: 2:1.2.2-1
libsm6:amd64: 2:1.2.2-1
libsmartcols1:amd64: 2.27.1-6ubuntu3.1
libsndfile1:amd64: 1.0.25-10
libspice-server1:amd64: 0.12.6-4ubuntu0.1
libsqlite3-0:amd64: 3.11.0-1ubuntu1
libss2:amd64: 1.42.13-1ubuntu1
libssl1.0.0:amd64: 1.0.2g-1ubuntu4.5
libstdc++-5-dev:amd64: 5.4.0-6ubuntu1~16.04.2
libstdc++6:amd64: 5.4.0-6ubuntu1~16.04.2
libsystemd0:amd64: 229-4ubuntu10
libtasn1-6:amd64: 4.7-3ubuntu0.16.04.1
libtcl8.6:amd64: 8.6.5+dfsg-2
libtext-charwidth-perl: 0.04-7build5
libtext-iconv-perl: 1.7-5build4
libtext-wrapi18n-perl: 0.06-7.1
libthai-data: 0.1.24-2
libthai0:amd64: 0.1.24-2
libtiff5:amd64: 4.0.6-1ubuntu0.2
libtinfo5:amd64: 6.0+20160213-1ubuntu1
libtk8.6:amd64: 8.6.5-1
libtool: 2.4.6-0.1

```

(continues on next page)

(continued from previous page)

```
libtsan0:amd64: 5.4.0-6ubuntu1~16.04.2
libtxc-dxtn-s2tc0:amd64: 0~git20131104-1.1
libubsan0:amd64: 5.4.0-6ubuntu1~16.04.2
libudev1:amd64: 229-4ubuntu10
libusb-0.1-4:amd64: 2:0.1.12-28
libusb-1.0-0:amd64: 2:1.0.20-1
libusbredirparser1:amd64: 0.7.1-1
libustr-1.0-1:amd64: 1.0.4-5
libutempter0:amd64: 1.1.6-3
libuuid1:amd64: 2.27.1-6ubuntu3.1
libvorbis0a:amd64: 1.3.5-3
libvorbisenc2:amd64: 1.3.5-3
libwind0-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libwrap0:amd64: 7.6.q-25
libx11-6:amd64: 2:1.6.3-1ubuntu2
libx11-data: 2:1.6.3-1ubuntu2
libx11-dev:amd64: 2:1.6.3-1ubuntu2
libx11-doc: 2:1.6.3-1ubuntu2
libx11-xcb1:amd64: 2:1.6.3-1ubuntu2
libxau-dev:amd64: 1:1.0.8-1
libxau6:amd64: 1:1.0.8-1
libxaw7:amd64: 2:1.0.13-1
libxcb-dri2-0:amd64: 1.11.1-1ubuntu1
libxcb-dri3-0:amd64: 1.11.1-1ubuntu1
libxcb-glx0:amd64: 1.11.1-1ubuntu1
libxcb-present0:amd64: 1.11.1-1ubuntu1
libxcb-render0:amd64: 1.11.1-1ubuntu1
libxcb-shape0:amd64: 1.11.1-1ubuntu1
libxcb-shm0:amd64: 1.11.1-1ubuntu1
libxcb-sync1:amd64: 1.11.1-1ubuntu1
libxcb1:amd64: 1.11.1-1ubuntu1
libxcb1-dev:amd64: 1.11.1-1ubuntu1
libxcomposite1:amd64: 1:0.4.4-1
libxcursor1:amd64: 1:1.1.14-1
libxdamage1:amd64: 1:1.1.4-2
libxdmcp-dev:amd64: 1:1.1.2-1.1
libxdmcp6:amd64: 1:1.1.2-1.1
libxen-4.6:amd64: 4.6.0-1ubuntu4.1
libxenstore3.0:amd64: 4.6.0-1ubuntu4.1
libxext6:amd64: 2:1.3.3-1
libxfixes3:amd64: 1:5.0.1-2
libxft2:amd64: 2.3.2-1
libxi6:amd64: 2:1.7.6-1
libxinerama1:amd64: 2:1.1.3-1
libxml2:amd64: 2.9.3+dfsg1-1ubuntu0.1
libxmu6:amd64: 2:1.1.2-2
libxmuu1:amd64: 2:1.1.2-2
libxpm4:amd64: 1:3.5.11-1
libxrandr2:amd64: 2:1.5.0-1
libxrender1:amd64: 1:0.9.9-0ubuntu1
libxshmfence1:amd64: 1.2-1
libxss1:amd64: 1:1.2.2-1
libxt-dev:amd64: 1:1.1.5-0ubuntu1
libxt6:amd64: 1:1.1.5-0ubuntu1
libxtables11:amd64: 1.6.0-2ubuntu3
libxtst6:amd64: 2:1.2.2-1
libxv1:amd64: 2:1.0.10-1
libxxf86dga1:amd64: 2:1.1.4-1
libxxf86vm1:amd64: 1:1.1.4-1
libyajl2:amd64: 2.1.0-2
linux-base: 4.0ubuntu1
```

(continues on next page)

(continued from previous page)

```
linux-firmware: 1.157.2
linux-generic: 4.4.0.72.78
linux-headers-4.4.0-72: 4.4.0-72.93
linux-headers-4.4.0-72-generic: 4.4.0-72.93
linux-headers-generic: 4.4.0.72.78
linux-image-4.4.0-72-generic: 4.4.0-72.93
linux-image-extra-4.4.0-72-generic: 4.4.0-72.93
linux-image-generic: 4.4.0.72.78
linux-libc-dev:amd64: 4.4.0-72.93
locales: 2.23-0ubuntu3
login: 1:4.2-3.1ubuntu5
logrotate: 3.8.7-2ubuntu2
lsb-base: 9.20160110ubuntu0.2
lsb-release: 9.20160110ubuntu0.2
lxc: 2.0.7-0ubuntu1~16.04.2
lxc-common: 2.0.7-0ubuntu1~16.04.2
lxc-templates: 2.0.7-0ubuntu1~16.04.2
lxc1: 2.0.7-0ubuntu1~16.04.2
lxcfs: 2.0.6-0ubuntu1~16.04.1
m4: 1.4.17-5
make: 4.1-6
makedev: 2.3.1-93ubuntu1
manpages: 4.04-2
manpages-dev: 4.04-2
mawk: 1.3.3-17ubuntu2
mime-support: 3.59ubuntu1
mount: 2.27.1-6ubuntu3.1
mountall: 2.54ubuntu1
msr-tools: 1.3-2
multiarch-support: 2.23-0ubuntu3
ncurses-base: 6.0+20160213-1ubuntu1
ncurses-bin: 6.0+20160213-1ubuntu1
ncurses-term: 6.0+20160213-1ubuntu1
net-tools: 1.60-26ubuntu1
netbase: 5.3
netcat-openbsd: 1.105-7ubuntu1
openjdk-8-jdk:amd64: 8u131-b11-2ubuntu1.16.04.3
openjdk-8-jdk-headless:amd64: 8u131-b11-2ubuntu1.16.04.3
openjdk-8-jre:amd64: 8u131-b11-2ubuntu1.16.04.3
openjdk-8-jre-headless:amd64: 8u131-b11-2ubuntu1.16.04.3
openssh-client: 1:7.2p2-4ubuntu2.1
openssh-server: 1:7.2p2-4ubuntu2.1
openssh-sftp-server: 1:7.2p2-4ubuntu2.1
openssl: 1.0.2g-1ubuntu4.5
os-prober: 1.70ubuntu3
passwd: 1:4.2-3.1ubuntu5
patch: 2.7.5-1
pciutils: 1:3.3.1-1.1ubuntu1
perl: 5.22.1-9
perl-base: 5.22.1-9
perl-modules-5.22: 5.22.1-9
pkg-config: 0.29.1-0ubuntu1
plymouth: 0.9.2-3ubuntu13.1
plymouth-theme-ubuntu-text: 0.9.2-3ubuntu13.1
procps: 2:3.3.10-4ubuntu2
python: 2.7.11-1
python-all: 2.7.11-1
python-all-dev: 2.7.11-1
python-apt: 1.1.0~beta1build1
python-apt-common: 1.1.0~beta1build1
python-dev: 2.7.11-1
```

(continues on next page)

(continued from previous page)

```
python-iniparse: 0.4-2.2
python-minimal: 2.7.11-1
python-pip: 8.1.1-2ubuntu0.2
python-pip-whl: 8.1.1-2ubuntu0.2
python-pkg-resources: 20.7.0-1
python-setuptools: 20.7.0-1
python-six: 1.10.0-3
python-virtualenv: 15.0.1+ds-3
python-wheel: 0.29.0-1
python2.7: 2.7.12-1~16.04
python2.7-dev: 2.7.12-1~16.04
python2.7-minimal: 2.7.12-1~16.04
python3: 3.5.1-3
python3-apt: 1.1.0~beta1build1
python3-chardet: 2.3.0-2
python3-dbus: 1.2.0-3
python3-gi: 3.20.0-0ubuntu1
python3-lxc: 2.0.7-0ubuntu1~16.04.2
python3-minimal: 3.5.1-3
python3-pkg-resources: 20.7.0-1
python3-requests: 2.9.1-3
python3-six: 1.10.0-3
python3-urllib3: 1.13.1-2ubuntu0.16.04.1
python3-virtualenv: 15.0.1+ds-3
python3.5: 3.5.2-2ubuntu0~16.04.1
python3.5-minimal: 3.5.2-2ubuntu0~16.04.1
qemu-block-extra:amd64: 1:2.5+dfsg-5ubuntu10.5
qemu-system-common: 1:2.5+dfsg-5ubuntu10.5
qemu-system-x86: 1:2.5+dfsg-5ubuntu10.5
qemu-utils: 1:2.5+dfsg-5ubuntu10.5
readline-common: 6.3-8ubuntu2
rename: 0.20-4
resolvconf: 1.78ubuntu2
rsync: 3.1.1-3ubuntu1
rsyslog: 8.16.0-1ubuntu3
screen: 4.3.1-2build1
seabios: 1.8.2-1ubuntu1
sed: 4.2.2-7
sensible-utils: 0.0.9
sgml-base: 1.26+nmu4ubuntu1
shared-mime-info: 1.5-2ubuntu0.1
sharutils: 1:4.15.2-1
socat: 1.7.3.1-1
ssh-import-id: 5.5-0ubuntu1
sudo: 1.8.16-0ubuntu1.1
systemd: 229-4ubuntu10
systemd-sysv: 229-4ubuntu10
sysv-rc: 2.88dsf-59.3ubuntu2
sysvinit-utils: 2.88dsf-59.3ubuntu2
tar: 1.28-2.1
tasksel: 3.34ubuntu3
tasksel-data: 3.34ubuntu3
tcl-expect:amd64: 5.45-7
tcl8.6: 8.6.5+dfsg-2
tcpd: 7.6.q-25
telnet: 0.17-40
tk8.6: 8.6.5-1
tzdata: 2016f-0ubuntu0.16.04
ubuntu-keyring: 2012.05.19
ubuntu-minimal: 1.361
ucf: 3.0036
```

(continues on next page)

(continued from previous page)

```

udev: 229-4ubuntu10
uidmap: 1:4.2-3.1ubuntu5.3
ureadahead: 0.100.0-19
usbutils: 1:007-4
util-linux: 2.27.1-6ubuntu3.1
uuid-runtime: 2.27.1-6ubuntu3.2
vim-common: 2:7.4.1689-3ubuntu1.1
vim-tiny: 2:7.4.1689-3ubuntu1.1
virtualenv: 15.0.1+ds-3
wamerican: 7.1-1
wget: 1.17.1-1ubuntu1.1
whiptail: 0.52.18-1ubuntu2
wireless-regdb: 2015.07.20-1ubuntu1
x11-common: 1:7.7+13ubuntu3
x11-utils: 7.7+3
x11proto-core-dev: 7.0.31-1~ubuntu16.04.1
x11proto-input-dev: 2.3.1-1
x11proto-kb-dev: 1.0.7-0ubuntu1
xauth: 1:1.0.9-1ubuntu2
xbitmaps: 1.1.1-2
xdg-user-dirs: 0.15-2ubuntu6
xkb-data: 2.16-1ubuntu1
xml-core: 0.13+nmu2
xorg-sgml-doctools: 1:1.11-1
xterm: 322-1ubuntu1
xtrans-dev: 1.3.5-1
xz-utils: 5.1.1alpha+20120614-2ubuntu2
zlib1g:amd64: 1:1.2.8.dfsg-2ubuntu4
zlib1g-dev:amd64: 1:1.2.8.dfsg-2ubuntu4

```

Kernel module listing

```

$ lsmod | sort
8250_fintek          16384  0
ablk_helper         16384  1 aesni_intel
acpi_pad            24576  0
acpi_power_meter    20480  0
aesni_intel         167936  0
aes_x86_64          20480  1 aesni_intel
ahci                 36864  0
authenc             16384  1 intel_qat
autofs4             40960  2
bridge              126976  0
coretemp            16384  0
crc32_pclmul        16384  0
crct10dif_pclmul    16384  0
cryptd              20480  3 ghash_clmulni_intel,aesni_intel,ablk_helper
dca                 16384  2 igb,ixgbe
edac_core           53248  1 sb_edac
enclosure           16384  1 ses
enic                81920  0
fjes                28672  0
fnic                106496  0
gf128mul            16384  1 lrw
ghash_clmulni_intel 16384  0
glue_helper         16384  1 aesni_intel
hid                 118784  2 hid_generic,usbhid
hid_generic         16384  0
i2c_algo_bit        16384  1 igb
i40e                286720  0
igb                 196608  0

```

(continues on next page)

(continued from previous page)

| | | | |
|------------------------|--------|---------|---|
| igb_uio | 16384 | 0 | |
| input_leds | 16384 | 0 | |
| intel_powerclamp | 16384 | 0 | |
| intel_qat | 110592 | 2 | qat_dh895xccvf, qat_dh895xcc |
| intel_rapl | 20480 | 0 | |
| ip6_udp_tunnel | 16384 | 1 | vxlan |
| ipmi_msghandler | 49152 | 2 | ipmi_ssif, ipmi_si |
| ipmi_si | 57344 | 0 | |
| ipmi_ssif | 24576 | 0 | |
| iptables_filter | 16384 | 1 | |
| iptables_mangle | 16384 | 1 | |
| iptables_nat | 16384 | 1 | |
| ip_tables | 24576 | 3 | iptables_filter, iptables_mangle, iptables_nat |
| ipt_MASQUERADE | 16384 | 1 | |
| irqbypass | 16384 | 1 | kvm |
| ixgbe | 290816 | 0 | |
| joydev | 20480 | 0 | |
| kvm | 544768 | 1 | kvm_intel |
| kvm_intel | 172032 | 0 | |
| libahci | 32768 | 1 | ahci |
| libfc | 114688 | 2 | fnic, libfcoc |
| libfcoc | 65536 | 1 | fnic |
| llc | 16384 | 2 | stp, bridge |
| lpc_ich | 24576 | 0 | |
| lrw | 16384 | 1 | aesni_intel |
| mac_hid | 16384 | 0 | |
| mdio | 16384 | 1 | ixgbe |
| megaraid_sas | 135168 | 3 | |
| mei | 98304 | 1 | mei_me |
| mei_me | 36864 | 0 | |
| Module | Size | Used by | |
| nf_conntrack | 106496 | 4 | nf_nat, nf_nat_ipv4, nf_nat_masquerade_ipv4, nf_conntrack_ipv4 |
| nf_conntrack_ipv4 | 16384 | 1 | |
| nf_defrag_ipv4 | 16384 | 1 | nf_conntrack_ipv4 |
| nf_nat | 24576 | 2 | nf_nat_ipv4, nf_nat_masquerade_ipv4 |
| nf_nat_ipv4 | 16384 | 1 | iptables_nat |
| nf_nat_masquerade_ipv4 | 16384 | 1 | ipt_MASQUERADE |
| pps_core | 20480 | 1 | ptp |
| ptp | 20480 | 3 | igb, i40e, ixgbe |
| qat_dh895xcc | 20480 | 0 | |
| qat_dh895xccvf | 20480 | 0 | |
| sb_edac | 32768 | 0 | |
| scsi_transport_fc | 61440 | 2 | fnic, libfc |
| ses | 20480 | 0 | |
| shpchp | 36864 | 0 | |
| stp | 16384 | 1 | bridge |
| udp_tunnel | 16384 | 1 | vxlan |
| uio | 20480 | 2 | uio_pci_generic, igb_uio |
| uio_pci_generic | 16384 | 0 | |
| usbhid | 49152 | 0 | |
| veth | 16384 | 0 | |
| vxlan | 49152 | 2 | i40e, ixgbe |
| wmi | 20480 | 0 | |
| x86_pkg_temp_thermal | 16384 | 0 | |
| x_tables | 36864 | 6 | xt_CHECKSUM, ip_tables, xt_tcpudp, ipt_MASQUERADE, iptables_filter, iptables_mangle |
| xt_CHECKSUM | 16384 | 1 | |
| xt_tcpudp | 16384 | 5 | |

Sysctl listing

```

$ sysctl -a
abi.vsyscall32 = 1
debug.exception-trace = 1
debug.kprobes-optimization = 1
dev.cdrom.autoclose = 1
dev.cdrom.autoeject = 0
dev.cdrom.check_media = 0
dev.cdrom.debug = 0
dev.cdrom.info = CD-ROM information, Id: cdrom.c 3.20 2003/12/17
dev.cdrom.info =
dev.cdrom.info = drive name:
dev.cdrom.info = drive speed:
dev.cdrom.info = drive # of slots:
dev.cdrom.info = Can close tray:
dev.cdrom.info = Can open tray:
dev.cdrom.info = Can lock tray:
dev.cdrom.info = Can change speed:
dev.cdrom.info = Can select disk:
dev.cdrom.info = Can read multisession:
dev.cdrom.info = Can read MCN:
dev.cdrom.info = Reports media changed:
dev.cdrom.info = Can play audio:
dev.cdrom.info = Can write CD-R:
dev.cdrom.info = Can write CD-RW:
dev.cdrom.info = Can read DVD:
dev.cdrom.info = Can write DVD-R:
dev.cdrom.info = Can write DVD-RAM:
dev.cdrom.info = Can read MRW:
dev.cdrom.info = Can write MRW:
dev.cdrom.info = Can write RAM:
dev.cdrom.info =
dev.cdrom.info =
dev.cdrom.lock = 0
dev.hpet.max-user-freq = 64
dev.mac_hid.mouse_button2_keycode = 97
dev.mac_hid.mouse_button3_keycode = 100
dev.mac_hid.mouse_button_emulation = 0
dev.raid.speed_limit_max = 200000
dev.raid.speed_limit_min = 1000
dev.scsi.logging_level = 0
fs.aio-max-nr = 65536
fs.aio-nr = 0
fs.binfmt_misc.status = enabled
fs.dentry-state = 69970      58326  45      0      0      0
fs.dir-notify-enable = 1
fs.epoll.max_user_watches = 108185784
fs.file-max = 52706330
fs.file-nr = 1224      0      52706330
fs.inode-nr = 42965 369
fs.inode-state = 42965      369      0      0      0      0
fs.inotify.max_queued_events = 16384
fs.inotify.max_user_instances = 128
fs.inotify.max_user_watches = 8192
fs.lease-break-time = 45
fs.leases-enable = 1
fs.mount-max = 100000
fs.mqueue.msg_default = 10
fs.mqueue.msg_max = 10
fs.mqueue.msgsize_default = 8192
fs.mqueue.msgsize_max = 8192
fs.mqueue.queues_max = 256
fs.nr_open = 1048576

```

(continues on next page)

(continued from previous page)

```
fs.overflowgid = 65534
fs.overflowuid = 65534
fs.pipe-max-size = 1048576
fs.pipe-user-pages-hard = 0
fs.pipe-user-pages-soft = 16384
fs.protected_hardlinks = 1
fs.protected_symlinks = 1
fs.quota.allocated_dquots = 0
fs.quota.cache_hits = 0
fs.quota.drops = 0
fs.quota.free_dquots = 0
fs.quota.lookups = 0
fs.quota.reads = 0
fs.quota.syncs = 0
fs.quota.writes = 0
fs.suid_dumpable = 0
kernel.acct = 4      2      30
kernel.acpi_video_flags = 0
kernel.auto_msgmni = 0
kernel.bootloader_type = 114
kernel.bootloader_version = 2
kernel.cad_pid = 1
kernel.cap_last_cap = 37
kernel.compat-log = 1
kernel.core_pattern = core
kernel.core_pipe_limit = 0
kernel.core_uses_pid = 0
kernel.ctrl-alt-del = 0
kernel.dmesg_restrict = 0
kernel.domainname = (none)
kernel.ftrace_dump_on_oops = 0
kernel.ftrace_enabled = 1
kernel.hardlockup_all_cpu_backtrace = 0
kernel.hardlockup_panic = 0
kernel.hostname = t2-sut1
kernel.hotplug =
kernel.hung_task_check_count = 4194304
kernel.hung_task_panic = 0
kernel.hung_task_timeout_secs = 120
kernel.hung_task_warnings = 10
kernel.io_delay_type = 1
kernel.kexec_load_disabled = 0
kernel.keys.gc_delay = 300
kernel.keys.maxbytes = 20000
kernel.keys.maxkeys = 200
kernel.keys.persistent_keyring_expiry = 259200
kernel.keys.root_maxbytes = 25000000
kernel.keys.root_maxkeys = 1000000
kernel.kptr_restrict = 1
kernel.kstack_depth_to_print = 12
kernel.max_lock_depth = 1024
kernel.modprobe = /sbin/modprobe
kernel.modules_disabled = 0
kernel.moksbstate_disabled = 0
kernel.msg_next_id = -1
kernel.msgmax = 8192
kernel.msgmnb = 16384
kernel.msgmni = 32000
kernel.ngroups_max = 65536
kernel.nmi_watchdog = 1
kernel.ns_last_pid = 11764
```

(continues on next page)

(continued from previous page)

```
kernel.numa_balancing = 1
kernel.numa_balancing_scan_delay_ms = 1000
kernel.numa_balancing_scan_period_max_ms = 60000
kernel.numa_balancing_scan_period_min_ms = 1000
kernel.numa_balancing_scan_size_mb = 256
kernel.osrelease = 4.4.0-72-generic
kernel.ostype = Linux
kernel.overflowgid = 65534
kernel.overflowuid = 65534
kernel.panic = 0
kernel.panic_on_io_nmi = 0
kernel.panic_on_oops = 0
kernel.panic_on_unrecovered_nmi = 0
kernel.panic_on_warn = 0
kernel.perf_cpu_time_max_percent = 25
kernel.perf_event_max_sample_rate = 12500
kernel.perf_event_mlock_kb = 516
kernel.perf_event_paranoid = 1
kernel.pid_max = 36864
kernel.poweroff_cmd = /sbin/poweroff
kernel.print_fatal_signals = 0
kernel.printk = 4 4 1 7
kernel.printk_delay = 0
kernel.printk_ratelimit = 5
kernel.printk_ratelimit_burst = 10
kernel.pty.max = 4096
kernel.pty.nr = 1
kernel.pty.reserve = 1024
kernel.random.boot_id = f683c836-6fc6-492a-a23b-62ab21895040
kernel.random.entropy_avail = 200
kernel.random.poolsize = 4096
kernel.random.read_wakeup_threshold = 64
kernel.random.urandom_min_reseed_secs = 60
kernel.random.uuid = 144ff2ba-1bc7-4836-8fb7-6aaa0ab7e65f
kernel.random.write_wakeup_threshold = 896
kernel.randomize_va_space = 0
kernel.real-root-dev = 0
kernel.sched_autogroup_enabled = 1
kernel.sched_cfs_bandwidth_slice_us = 5000
kernel.sched_child_runs_first = 0
kernel.sched_domain.cpu0.domain0.busy_factor = 32
kernel.sched_domain.cpu0.domain0.busy_idx = 3
kernel.sched_domain.cpu0.domain0.cache_nice_tries = 2
kernel.sched_domain.cpu0.domain0.flags = 25647
kernel.sched_domain.cpu0.domain0.forkexec_idx = 0
kernel.sched_domain.cpu0.domain0.idle_idx = 2
kernel.sched_domain.cpu0.domain0.imbalance_pct = 125
kernel.sched_domain.cpu0.domain0.max_interval = 72
kernel.sched_domain.cpu0.domain0.max_newidle_lb_cost = 1309
kernel.sched_domain.cpu0.domain0.min_interval = 36
kernel.sched_domain.cpu0.domain0.name = NUMA
kernel.sched_domain.cpu0.domain0.newidle_idx = 0
kernel.sched_domain.cpu0.domain0.wake_idx = 0
kernel.sched_domain.cpu18.domain0.busy_factor = 32
kernel.sched_domain.cpu18.domain0.busy_idx = 3
kernel.sched_domain.cpu18.domain0.cache_nice_tries = 2
kernel.sched_domain.cpu18.domain0.flags = 25647
kernel.sched_domain.cpu18.domain0.forkexec_idx = 0
kernel.sched_domain.cpu18.domain0.idle_idx = 2
kernel.sched_domain.cpu18.domain0.imbalance_pct = 125
kernel.sched_domain.cpu18.domain0.max_interval = 72
```

(continues on next page)

(continued from previous page)

```
kernel.sched_domain.cpu18.domain0.max_newidle_lb_cost = 2026
kernel.sched_domain.cpu18.domain0.min_interval = 36
kernel.sched_domain.cpu18.domain0.name = NUMA
kernel.sched_domain.cpu18.domain0.newidle_idx = 0
kernel.sched_domain.cpu18.domain0.wake_idx = 0
kernel.sched_latency_ns = 24000000
kernel.sched_migration_cost_ns = 500000
kernel.sched_min_granularity_ns = 3000000
kernel.sched_nr_migrate = 32
kernel.sched_rr_timeslice_ms = 25
kernel.sched_rt_period_us = 1000000
kernel.sched_rt_runtime_us = 950000
kernel.sched_shares_window_ns = 10000000
kernel.sched_time_avg_ms = 1000
kernel.sched_tunable_scaling = 1
kernel.sched_wakeup_granularity_ns = 4000000
kernel.secure_boot = 0
kernel.sem = 32000 1024000000 500 32000
kernel.sem_next_id = -1
kernel.sg-big-buff = 32768
kernel.shm_next_id = -1
kernel.shm_rmid_forced = 0
kernel.shmall = 18446744073692774399
kernel.shmmax = 8589934592
kernel.shmmni = 4096
kernel.soft_watchdog = 1
kernel.softlockup_all_cpu_backtrace = 0
kernel.softlockup_panic = 0
kernel.stack_tracer_enabled = 0
kernel.sysctl_writes_strict = 0
kernel.sysrq = 176
kernel.tainted = 12288
kernel.threads-max = 4126960
kernel.timer_migration = 1
kernel.traceoff_on_warning = 0
kernel.tracepoint_printk = 0
kernel.unknown_nmi_panic = 0
kernel.unprivileged_bpf_disabled = 0
kernel.unprivileged_userns_apparmor_policy = 1
kernel.unprivileged_userns_clone = 1
kernel.usermodehelper.bset = 4294967295 63
kernel.usermodehelper.inheritable = 4294967295 63
kernel.version = #93-Ubuntu SMP Fri Mar 31 14:07:41 UTC 2017
kernel.watchdog = 1
kernel.watchdog_cpumask = 0,18
kernel.watchdog_thresh = 10
kernel.yama.ptrace_scope = 1
net.core.bpf_jit_enable = 0
net.core.busy_poll = 0
net.core.busy_read = 0
net.core.default_qdisc = pfifo_fast
net.core.dev_weight = 64
net.core.flow_limit_cpu_bitmap = 0,00000000
net.core.flow_limit_table_len = 4096
net.core.max_skb_frags = 17
net.core.message_burst = 10
net.core.message_cost = 5
net.core.netdev_budget = 300
net.core.netdev_max_backlog = 1000
net.core.netdev_rss_key = _
↪29:61:61:e6:4e:d5:d0:a2:dc:81:6a:c8:44:1b:e2:8d:c8:6f:6a:2b:64:62:98:08:bb:63:48:8e:96:d1:6a:15:32:ca:da:8d:3c:0a:e
```

(continues on next page)

(continued from previous page)

```
net.core.netdev_tstamp_prequeue = 1
net.core.optmem_max = 20480
net.core.rmem_default = 212992
net.core.rmem_max = 212992
net.core.rps_sock_flow_entries = 0
net.core.somaxconn = 128
net.core.tstamp_allow_data = 1
net.core.warnings = 0
net.core.wmem_default = 212992
net.core.wmem_max = 212992
net.core.xfrm_acq_expires = 30
net.core.xfrm_aevent_etime = 10
net.core.xfrm_aevent_rseqth = 2
net.core.xfrm_larval_drop = 1
net.fan.vxlan = 4
net.ipv4.cipso_cache_bucket_size = 10
net.ipv4.cipso_cache_enable = 1
net.ipv4.cipso_rbm_optfmt = 0
net.ipv4.cipso_rbm_strictvalid = 1
net.ipv4.conf.all.accept_local = 0
net.ipv4.conf.all.accept_redirects = 0
net.ipv4.conf.all.accept_source_route = 0
net.ipv4.conf.all.arp_accept = 0
net.ipv4.conf.all.arp_announce = 0
net.ipv4.conf.all.arp_filter = 0
net.ipv4.conf.all.arp_ignore = 0
net.ipv4.conf.all.arp_notify = 0
net.ipv4.conf.all.bootp_relay = 0
net.ipv4.conf.all.disable_policy = 0
net.ipv4.conf.all.disable_xfrm = 0
net.ipv4.conf.all.force_igmp_version = 0
net.ipv4.conf.all.forwarding = 1
net.ipv4.conf.all.igmpv2_unsolicited_report_interval = 10000
net.ipv4.conf.all.igmpv3_unsolicited_report_interval = 1000
net.ipv4.conf.all.ignore_routes_with_linkdown = 0
net.ipv4.conf.all.log_martians = 0
net.ipv4.conf.all.mc_forwarding = 0
net.ipv4.conf.all.medium_id = 0
net.ipv4.conf.all.promote_secondaries = 0
net.ipv4.conf.all.proxy_arp = 0
net.ipv4.conf.all.proxy_arp_pvlan = 0
net.ipv4.conf.all.route_localnet = 0
net.ipv4.conf.all.rp_filter = 1
net.ipv4.conf.all.secure_redirects = 1
net.ipv4.conf.all.send_redirects = 1
net.ipv4.conf.all.shared_media = 1
net.ipv4.conf.all.src_valid_mark = 0
net.ipv4.conf.all.tag = 0
net.ipv4.conf.default.accept_local = 0
net.ipv4.conf.default.accept_redirects = 1
net.ipv4.conf.default.accept_source_route = 1
net.ipv4.conf.default.arp_accept = 0
net.ipv4.conf.default.arp_announce = 0
net.ipv4.conf.default.arp_filter = 0
net.ipv4.conf.default.arp_ignore = 0
net.ipv4.conf.default.arp_notify = 0
net.ipv4.conf.default.bootp_relay = 0
net.ipv4.conf.default.disable_policy = 0
net.ipv4.conf.default.disable_xfrm = 0
net.ipv4.conf.default.force_igmp_version = 0
net.ipv4.conf.default.forwarding = 1
```

(continues on next page)

(continued from previous page)

```
net.ipv4.conf.default.igmpv2_unsolicited_report_interval = 10000
net.ipv4.conf.default.igmpv3_unsolicited_report_interval = 1000
net.ipv4.conf.default.ignore_routes_with_linkdown = 0
net.ipv4.conf.default.log_martians = 0
net.ipv4.conf.default.mc_forwarding = 0
net.ipv4.conf.default.medium_id = 0
net.ipv4.conf.default.promote_secondaries = 0
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.default.proxy_arp_pvlan = 0
net.ipv4.conf.default.route_localnet = 0
net.ipv4.conf.default.rp_filter = 1
net.ipv4.conf.default.secure_redirects = 1
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.default.shared_media = 1
net.ipv4.conf.default.src_valid_mark = 0
net.ipv4.conf.default.tag = 0
net.ipv4.conf.enp25s0f0.accept_local = 0
net.ipv4.conf.enp25s0f0.accept_redirects = 1
net.ipv4.conf.enp25s0f0.accept_source_route = 1
net.ipv4.conf.enp25s0f0.arp_accept = 0
net.ipv4.conf.enp25s0f0.arp_announce = 0
net.ipv4.conf.enp25s0f0.arp_filter = 0
net.ipv4.conf.enp25s0f0.arp_ignore = 0
net.ipv4.conf.enp25s0f0.arp_notify = 0
net.ipv4.conf.enp25s0f0.bootp_relay = 0
net.ipv4.conf.enp25s0f0.disable_policy = 0
net.ipv4.conf.enp25s0f0.disable_xfrm = 0
net.ipv4.conf.enp25s0f0.force_igmp_version = 0
net.ipv4.conf.enp25s0f0.forwarding = 1
net.ipv4.conf.enp25s0f0.igmpv2_unsolicited_report_interval = 10000
net.ipv4.conf.enp25s0f0.igmpv3_unsolicited_report_interval = 1000
net.ipv4.conf.enp25s0f0.ignore_routes_with_linkdown = 0
net.ipv4.conf.enp25s0f0.log_martians = 0
net.ipv4.conf.enp25s0f0.mc_forwarding = 0
net.ipv4.conf.enp25s0f0.medium_id = 0
net.ipv4.conf.enp25s0f0.promote_secondaries = 0
net.ipv4.conf.enp25s0f0.proxy_arp = 0
net.ipv4.conf.enp25s0f0.proxy_arp_pvlan = 0
net.ipv4.conf.enp25s0f0.route_localnet = 0
net.ipv4.conf.enp25s0f0.rp_filter = 1
net.ipv4.conf.enp25s0f0.secure_redirects = 1
net.ipv4.conf.enp25s0f0.send_redirects = 1
net.ipv4.conf.enp25s0f0.shared_media = 1
net.ipv4.conf.enp25s0f0.src_valid_mark = 0
net.ipv4.conf.enp25s0f0.tag = 0
net.ipv4.conf.lo.accept_local = 0
net.ipv4.conf.lo.accept_redirects = 1
net.ipv4.conf.lo.accept_source_route = 1
net.ipv4.conf.lo.arp_accept = 0
net.ipv4.conf.lo.arp_announce = 0
net.ipv4.conf.lo.arp_filter = 0
net.ipv4.conf.lo.arp_ignore = 0
net.ipv4.conf.lo.arp_notify = 0
net.ipv4.conf.lo.bootp_relay = 0
net.ipv4.conf.lo.disable_policy = 1
net.ipv4.conf.lo.disable_xfrm = 1
net.ipv4.conf.lo.force_igmp_version = 0
net.ipv4.conf.lo.forwarding = 1
net.ipv4.conf.lo.igmpv2_unsolicited_report_interval = 10000
net.ipv4.conf.lo.igmpv3_unsolicited_report_interval = 1000
net.ipv4.conf.lo.ignore_routes_with_linkdown = 0
```

(continues on next page)

(continued from previous page)

```
net.ipv4.conf.lo.log_martians = 0
net.ipv4.conf.lo.mc_forwarding = 0
net.ipv4.conf.lo.medium_id = 0
net.ipv4.conf.lo.promote_secondaries = 0
net.ipv4.conf.lo.proxy_arp = 0
net.ipv4.conf.lo.proxy_arp_pvlan = 0
net.ipv4.conf.lo.route_localnet = 0
net.ipv4.conf.lo.rp_filter = 0
net.ipv4.conf.lo.secure_redirects = 1
net.ipv4.conf.lo.send_redirects = 1
net.ipv4.conf.lo.shared_media = 1
net.ipv4.conf.lo.src_valid_mark = 0
net.ipv4.conf.lo.tag = 0
net.ipv4.conf.lxcbr0.accept_local = 0
net.ipv4.conf.lxcbr0.accept_redirects = 1
net.ipv4.conf.lxcbr0.accept_source_route = 1
net.ipv4.conf.lxcbr0.arp_accept = 0
net.ipv4.conf.lxcbr0.arp_announce = 0
net.ipv4.conf.lxcbr0.arp_filter = 0
net.ipv4.conf.lxcbr0.arp_ignore = 0
net.ipv4.conf.lxcbr0.arp_notify = 0
net.ipv4.conf.lxcbr0.bootp_relay = 0
net.ipv4.conf.lxcbr0.disable_policy = 0
net.ipv4.conf.lxcbr0.disable_xfrm = 0
net.ipv4.conf.lxcbr0.force_igmp_version = 0
net.ipv4.conf.lxcbr0.forwarding = 1
net.ipv4.conf.lxcbr0.igmpv2_unsolicited_report_interval = 10000
net.ipv4.conf.lxcbr0.igmpv3_unsolicited_report_interval = 1000
net.ipv4.conf.lxcbr0.ignore_routes_with_linkdown = 0
net.ipv4.conf.lxcbr0.log_martians = 0
net.ipv4.conf.lxcbr0.mc_forwarding = 0
net.ipv4.conf.lxcbr0.medium_id = 0
net.ipv4.conf.lxcbr0.promote_secondaries = 0
net.ipv4.conf.lxcbr0.proxy_arp = 0
net.ipv4.conf.lxcbr0.proxy_arp_pvlan = 0
net.ipv4.conf.lxcbr0.route_localnet = 0
net.ipv4.conf.lxcbr0.rp_filter = 1
net.ipv4.conf.lxcbr0.secure_redirects = 1
net.ipv4.conf.lxcbr0.send_redirects = 1
net.ipv4.conf.lxcbr0.shared_media = 1
net.ipv4.conf.lxcbr0.src_valid_mark = 0
net.ipv4.conf.lxcbr0.tag = 0
net.ipv4.fwmark_reflect = 0
net.ipv4.icmp_echo_ignore_all = 0
net.ipv4.icmp_echo_ignore_broadcasts = 1
net.ipv4.icmp_errors_use_inbound_ifaddr = 0
net.ipv4.icmp_ignore_bogus_error_responses = 1
net.ipv4.icmp_msgs_burst = 50
net.ipv4.icmp_msgs_per_sec = 1000
net.ipv4.icmp_ratelimit = 1000
net.ipv4.icmp_ratemask = 6168
net.ipv4.igmp_link_local_mcast_reports = 1
net.ipv4.igmp_max_memberships = 20
net.ipv4.igmp_max_msf = 10
net.ipv4.igmp_qrv = 2
net.ipv4.inet_peer_maxttl = 600
net.ipv4.inet_peer_minttl = 120
net.ipv4.inet_peer_threshold = 65664
net.ipv4.ip_default_ttl = 64
net.ipv4.ip_dynaddr = 0
net.ipv4.ip_early_demux = 1
```

(continues on next page)

(continued from previous page)

```
net.ipv4.ip_forward = 1
net.ipv4.ip_forward_use_pmtu = 0
net.ipv4.ip_local_port_range = 32768      60999
net.ipv4.ip_local_reserved_ports =
net.ipv4.ip_no_pmtu_disc = 0
net.ipv4.ip_nonlocal_bind = 0
net.ipv4.ipfrag_high_thresh = 4194304
net.ipv4.ipfrag_low_thresh = 3145728
net.ipv4.ipfrag_max_dist = 64
net.ipv4.ipfrag_secret_interval = 0
net.ipv4.ipfrag_time = 30
net.ipv4.neigh.default.anycast_delay = 100
net.ipv4.neigh.default.app_solicit = 0
net.ipv4.neigh.default.base_reachable_time_ms = 30000
net.ipv4.neigh.default.delay_first_probe_time = 5
net.ipv4.neigh.default.gc_interval = 30
net.ipv4.neigh.default.gc_stale_time = 60
net.ipv4.neigh.default.gc_thresh1 = 128
net.ipv4.neigh.default.gc_thresh2 = 512
net.ipv4.neigh.default.gc_thresh3 = 1024
net.ipv4.neigh.default.locktime = 100
net.ipv4.neigh.default.mcast_resolicit = 0
net.ipv4.neigh.default.mcast_solicit = 3
net.ipv4.neigh.default.proxy_delay = 80
net.ipv4.neigh.default.proxy_qlen = 64
net.ipv4.neigh.default.retrans_time_ms = 1000
net.ipv4.neigh.default.ucast_solicit = 3
net.ipv4.neigh.default.unres_qlen = 31
net.ipv4.neigh.default.unres_qlen_bytes = 65536
net.ipv4.neigh.enp25s0f0.anycast_delay = 100
net.ipv4.neigh.enp25s0f0.app_solicit = 0
net.ipv4.neigh.enp25s0f0.base_reachable_time_ms = 30000
net.ipv4.neigh.enp25s0f0.delay_first_probe_time = 5
net.ipv4.neigh.enp25s0f0.gc_stale_time = 60
net.ipv4.neigh.enp25s0f0.locktime = 100
net.ipv4.neigh.enp25s0f0.mcast_resolicit = 0
net.ipv4.neigh.enp25s0f0.mcast_solicit = 3
net.ipv4.neigh.enp25s0f0.proxy_delay = 80
net.ipv4.neigh.enp25s0f0.proxy_qlen = 64
net.ipv4.neigh.enp25s0f0.retrans_time_ms = 1000
net.ipv4.neigh.enp25s0f0.ucast_solicit = 3
net.ipv4.neigh.enp25s0f0.unres_qlen = 31
net.ipv4.neigh.enp25s0f0.unres_qlen_bytes = 65536
net.ipv4.neigh.lo.anycast_delay = 100
net.ipv4.neigh.lo.app_solicit = 0
net.ipv4.neigh.lo.base_reachable_time_ms = 30000
net.ipv4.neigh.lo.delay_first_probe_time = 5
net.ipv4.neigh.lo.gc_stale_time = 60
net.ipv4.neigh.lo.locktime = 100
net.ipv4.neigh.lo.mcast_resolicit = 0
net.ipv4.neigh.lo.mcast_solicit = 3
net.ipv4.neigh.lo.proxy_delay = 80
net.ipv4.neigh.lo.proxy_qlen = 64
net.ipv4.neigh.lo.retrans_time_ms = 1000
net.ipv4.neigh.lo.ucast_solicit = 3
net.ipv4.neigh.lo.unres_qlen = 31
net.ipv4.neigh.lo.unres_qlen_bytes = 65536
net.ipv4.neigh.lxcbr0.anycast_delay = 100
net.ipv4.neigh.lxcbr0.app_solicit = 0
net.ipv4.neigh.lxcbr0.base_reachable_time_ms = 30000
net.ipv4.neigh.lxcbr0.delay_first_probe_time = 5
```

(continues on next page)

(continued from previous page)

```
net.ipv4.neigh.lxcbr0.gc_stale_time = 60
net.ipv4.neigh.lxcbr0.locktime = 100
net.ipv4.neigh.lxcbr0.mcast_resolicit = 0
net.ipv4.neigh.lxcbr0.mcast_solicit = 3
net.ipv4.neigh.lxcbr0.proxy_delay = 80
net.ipv4.neigh.lxcbr0.proxy_qlen = 64
net.ipv4.neigh.lxcbr0.retrans_time_ms = 1000
net.ipv4.neigh.lxcbr0.ucast_solicit = 3
net.ipv4.neigh.lxcbr0.unres_qlen = 31
net.ipv4.neigh.lxcbr0.unres_qlen_bytes = 65536
net.ipv4.ping_group_range = 1 0
net.ipv4.route.error_burst = 1250
net.ipv4.route.error_cost = 250
net.ipv4.route.gc_elasticity = 8
net.ipv4.route.gc_interval = 60
net.ipv4.route.gc_min_interval = 0
net.ipv4.route.gc_min_interval_ms = 500
net.ipv4.route.gc_thresh = -1
net.ipv4.route.gc_timeout = 300
net.ipv4.route.max_size = 2147483647
net.ipv4.route.min_adv_mss = 256
net.ipv4.route.min_pmtu = 552
net.ipv4.route.mtu_expires = 600
net.ipv4.route.redirect_load = 5
net.ipv4.route.redirect_number = 9
net.ipv4.route.redirect_silence = 5120
net.ipv4.tcp_abort_on_overflow = 0
net.ipv4.tcp_adv_win_scale = 1
net.ipv4.tcp_allowed_congestion_control = cubic reno
net.ipv4.tcp_app_win = 31
net.ipv4.tcp_autocorking = 1
net.ipv4.tcp_available_congestion_control = cubic reno
net.ipv4.tcp_base_mss = 1024
net.ipv4.tcp_challenge_ack_limit = 1000
net.ipv4.tcp_congestion_control = cubic
net.ipv4.tcp_dsack = 1
net.ipv4.tcp_early_retrans = 3
net.ipv4.tcp_ecn = 2
net.ipv4.tcp_ecn_fallback = 1
net.ipv4.tcp_fack = 1
net.ipv4.tcp_fastopen = 1
net.ipv4.tcp_fastopen_key = 00000000-00000000-00000000-00000000
net.ipv4.tcp_fin_timeout = 60
net.ipv4.tcp_frto = 2
net.ipv4.tcp_fwmark_accept = 0
net.ipv4.tcp_invalid_ratelimit = 500
net.ipv4.tcp_keepalive_intvl = 75
net.ipv4.tcp_keepalive_probes = 9
net.ipv4.tcp_keepalive_time = 7200
net.ipv4.tcp_limit_output_bytes = 262144
net.ipv4.tcp_low_latency = 0
net.ipv4.tcp_max_orphans = 262144
net.ipv4.tcp_max_reordering = 300
net.ipv4.tcp_max_syn_backlog = 2048
net.ipv4.tcp_max_tw_buckets = 262144
net.ipv4.tcp_mem = 6188856 8251809 12377712
net.ipv4.tcp_min_rtt_wlen = 300
net.ipv4.tcp_min_tso_segs = 2
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_mtu_probing = 0
net.ipv4.tcp_no_metrics_save = 0
```

(continues on next page)

(continued from previous page)

```
net.ipv4.tcp_notsent_lowat = -1
net.ipv4.tcp_orphan_retries = 0
net.ipv4.tcp_pacing_ca_ratio = 120
net.ipv4.tcp_pacing_ss_ratio = 200
net.ipv4.tcp_probe_interval = 600
net.ipv4.tcp_probe_threshold = 8
net.ipv4.tcp_recovery = 1
net.ipv4.tcp_reordering = 3
net.ipv4.tcp_retrans_collapse = 1
net.ipv4.tcp_retries1 = 3
net.ipv4.tcp_retries2 = 15
net.ipv4.tcp_rfc1337 = 0
net.ipv4.tcp_rmem = 4096 87380 6291456
net.ipv4.tcp_sack = 1
net.ipv4.tcp_slow_start_after_idle = 1
net.ipv4.tcp_stdurg = 0
net.ipv4.tcp_syn_retries = 6
net.ipv4.tcp_synack_retries = 5
net.ipv4.tcp_syncookies = 1
net.ipv4.tcp_thin_dupack = 0
net.ipv4.tcp_thin_linear_timeouts = 0
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_tso_win_divisor = 3
net.ipv4.tcp_tw_recycle = 0
net.ipv4.tcp_tw_reuse = 0
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_wmem = 4096 16384 4194304
net.ipv4.tcp_workaround_signed_windows = 0
net.ipv4.udp_mem = 12377712 16503618 24755424
net.ipv4.udp_rmem_min = 4096
net.ipv4.udp_wmem_min = 4096
net.ipv4.xfrm4_gc_thresh = 2147483647
net.ipv6.anycast_src_echo_reply = 0
net.ipv6.auto_flowlabels = 1
net.ipv6.bindv6only = 0
net.ipv6.conf.all.accept_dad = 1
net.ipv6.conf.all.accept_ra = 1
net.ipv6.conf.all.accept_ra_defrtr = 1
net.ipv6.conf.all.accept_ra_from_local = 0
net.ipv6.conf.all.accept_ra_min_hop_limit = 1
net.ipv6.conf.all.accept_ra_mtu = 1
net.ipv6.conf.all.accept_ra_pinfo = 1
net.ipv6.conf.all.accept_ra_rt_info_max_plen = 0
net.ipv6.conf.all.accept_ra_rtr_pref = 1
net.ipv6.conf.all.accept_redirects = 1
net.ipv6.conf.all.accept_source_route = 0
net.ipv6.conf.all.autoconf = 1
net.ipv6.conf.all.dad_transmits = 1
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.all.force_mld_version = 0
net.ipv6.conf.all.force_tllao = 0
net.ipv6.conf.all.forwarding = 0
net.ipv6.conf.all.hop_limit = 64
net.ipv6.conf.all.ignore_routes_with_linkdown = 0
net.ipv6.conf.all.max_addresses = 16
net.ipv6.conf.all.max_desync_factor = 600
net.ipv6.conf.all.mc_forwarding = 0
net.ipv6.conf.all.mldv1_unsolicited_report_interval = 10000
net.ipv6.conf.all.mldv2_unsolicited_report_interval = 1000
net.ipv6.conf.all.mtu = 1280
net.ipv6.conf.all.ndisc_notify = 0
```

(continues on next page)

(continued from previous page)

```
net.ipv6.conf.all.proxy_ndp = 0
net.ipv6.conf.all.regen_max_retry = 3
net.ipv6.conf.all.router_probe_interval = 60
net.ipv6.conf.all.router_solicitation_delay = 1
net.ipv6.conf.all.router_solicitation_interval = 4
net.ipv6.conf.all.router_solicitations = 3
sysctl: reading key "net.ipv6.conf.all.stable_secret"
net.ipv6.conf.all.suppress_frag_ndisc = 1
net.ipv6.conf.all.temp_prefered_lft = 86400
net.ipv6.conf.all.temp_valid_lft = 604800
net.ipv6.conf.all.use_oif_addrs_only = 0
net.ipv6.conf.all.use_tempaddr = 2
net.ipv6.conf.default.accept_dad = 1
net.ipv6.conf.default.accept_ra = 1
net.ipv6.conf.default.accept_ra_defrtr = 1
net.ipv6.conf.default.accept_ra_from_local = 0
net.ipv6.conf.default.accept_ra_min_hop_limit = 1
net.ipv6.conf.default.accept_ra_mtu = 1
net.ipv6.conf.default.accept_ra_pinfo = 1
net.ipv6.conf.default.accept_ra_rt_info_max_plen = 0
net.ipv6.conf.default.accept_ra_rtr_pref = 1
net.ipv6.conf.default.accept_redirects = 1
net.ipv6.conf.default.accept_source_route = 0
net.ipv6.conf.default.autoconf = 1
net.ipv6.conf.default.dad_transmits = 1
net.ipv6.conf.default.disable_ipv6 = 0
net.ipv6.conf.default.force_mld_version = 0
net.ipv6.conf.default.force_tllao = 0
net.ipv6.conf.default.forwarding = 0
net.ipv6.conf.default.hop_limit = 64
net.ipv6.conf.default.ignore_routes_with_linkdown = 0
net.ipv6.conf.default.max_addresses = 16
net.ipv6.conf.default.max_desync_factor = 600
net.ipv6.conf.default.mc_forwarding = 0
net.ipv6.conf.default.mldv1_unsolicited_report_interval = 10000
net.ipv6.conf.default.mldv2_unsolicited_report_interval = 1000
net.ipv6.conf.default.mtu = 1280
net.ipv6.conf.default.ndisc_notify = 0
net.ipv6.conf.default.proxy_ndp = 0
net.ipv6.conf.default.regen_max_retry = 3
net.ipv6.conf.default.router_probe_interval = 60
net.ipv6.conf.default.router_solicitation_delay = 1
net.ipv6.conf.default.router_solicitation_interval = 4
net.ipv6.conf.default.router_solicitations = 3
sysctl: reading key "net.ipv6.conf.default.stable_secret"
net.ipv6.conf.default.suppress_frag_ndisc = 1
net.ipv6.conf.default.temp_prefered_lft = 86400
net.ipv6.conf.default.temp_valid_lft = 604800
net.ipv6.conf.default.use_oif_addrs_only = 0
net.ipv6.conf.default.use_tempaddr = 2
net.ipv6.conf.enp25s0f0.accept_dad = 1
net.ipv6.conf.enp25s0f0.accept_ra = 1
net.ipv6.conf.enp25s0f0.accept_ra_defrtr = 1
net.ipv6.conf.enp25s0f0.accept_ra_from_local = 0
net.ipv6.conf.enp25s0f0.accept_ra_min_hop_limit = 1
net.ipv6.conf.enp25s0f0.accept_ra_mtu = 1
net.ipv6.conf.enp25s0f0.accept_ra_pinfo = 1
net.ipv6.conf.enp25s0f0.accept_ra_rt_info_max_plen = 0
net.ipv6.conf.enp25s0f0.accept_ra_rtr_pref = 1
net.ipv6.conf.enp25s0f0.accept_redirects = 1
net.ipv6.conf.enp25s0f0.accept_source_route = 0
```

(continues on next page)

(continued from previous page)

```
net.ipv6.conf.enp25s0f0.autoconf = 1
net.ipv6.conf.enp25s0f0.dad_transmits = 1
net.ipv6.conf.enp25s0f0.disable_ipv6 = 0
net.ipv6.conf.enp25s0f0.force_mld_version = 0
net.ipv6.conf.enp25s0f0.force_tllao = 0
net.ipv6.conf.enp25s0f0.forwarding = 0
net.ipv6.conf.enp25s0f0.hop_limit = 64
net.ipv6.conf.enp25s0f0.ignore_routes_with_linkdown = 0
net.ipv6.conf.enp25s0f0.max_addresses = 16
net.ipv6.conf.enp25s0f0.max_desync_factor = 600
net.ipv6.conf.enp25s0f0.mc_forwarding = 0
net.ipv6.conf.enp25s0f0.mldv1_unsolicited_report_interval = 10000
net.ipv6.conf.enp25s0f0.mldv2_unsolicited_report_interval = 1000
net.ipv6.conf.enp25s0f0.mtu = 1500
net.ipv6.conf.enp25s0f0.ndisc_notify = 0
net.ipv6.conf.enp25s0f0.proxy_ndp = 0
net.ipv6.conf.enp25s0f0.regen_max_retry = 3
net.ipv6.conf.enp25s0f0.router_probe_interval = 60
net.ipv6.conf.enp25s0f0.router_solicitation_delay = 1
net.ipv6.conf.enp25s0f0.router_solicitation_interval = 4
net.ipv6.conf.enp25s0f0.router_solicitations = 3
sysctl: reading key "net.ipv6.conf.enp25s0f0.stable_secret"
net.ipv6.conf.enp25s0f0.suppress_frag_ndisc = 1
net.ipv6.conf.enp25s0f0.temp_prefered_lft = 86400
net.ipv6.conf.enp25s0f0.temp_valid_lft = 604800
net.ipv6.conf.enp25s0f0.use_oif_addrs_only = 0
net.ipv6.conf.enp25s0f0.use_tempaddr = 0
net.ipv6.conf.lo.accept_dad = -1
net.ipv6.conf.lo.accept_ra = 1
net.ipv6.conf.lo.accept_ra_defrtr = 1
net.ipv6.conf.lo.accept_ra_from_local = 0
net.ipv6.conf.lo.accept_ra_min_hop_limit = 1
net.ipv6.conf.lo.accept_ra_mtu = 1
net.ipv6.conf.lo.accept_ra_rtr_pref = 1
net.ipv6.conf.lo.accept_ra_rt_info_max_plen = 0
net.ipv6.conf.lo.accept_ra_rtr_pref = 1
net.ipv6.conf.lo.accept_redirects = 1
net.ipv6.conf.lo.accept_source_route = 0
net.ipv6.conf.lo.autoconf = 1
net.ipv6.conf.lo.dad_transmits = 1
net.ipv6.conf.lo.disable_ipv6 = 0
net.ipv6.conf.lo.force_mld_version = 0
net.ipv6.conf.lo.force_tllao = 0
net.ipv6.conf.lo.forwarding = 0
net.ipv6.conf.lo.hop_limit = 64
net.ipv6.conf.lo.ignore_routes_with_linkdown = 0
net.ipv6.conf.lo.max_addresses = 16
net.ipv6.conf.lo.max_desync_factor = 600
net.ipv6.conf.lo.mc_forwarding = 0
net.ipv6.conf.lo.mldv1_unsolicited_report_interval = 10000
net.ipv6.conf.lo.mldv2_unsolicited_report_interval = 1000
net.ipv6.conf.lo.mtu = 65536
net.ipv6.conf.lo.ndisc_notify = 0
net.ipv6.conf.lo.proxy_ndp = 0
net.ipv6.conf.lo.regen_max_retry = 3
net.ipv6.conf.lo.router_probe_interval = 60
net.ipv6.conf.lo.router_solicitation_delay = 1
net.ipv6.conf.lo.router_solicitation_interval = 4
net.ipv6.conf.lo.router_solicitations = 3
sysctl: reading key "net.ipv6.conf.lo.stable_secret"
net.ipv6.conf.lo.suppress_frag_ndisc = 1
```

(continues on next page)

(continued from previous page)

```
net.ipv6.conf.lo.temp_prefered_lft = 86400
net.ipv6.conf.lo.temp_valid_lft = 604800
net.ipv6.conf.lo.use_oif_addrs_only = 0
net.ipv6.conf.lo.use_tempaddr = -1
net.ipv6.conf.lxcbr0.accept_dad = 0
net.ipv6.conf.lxcbr0.accept_ra = 1
net.ipv6.conf.lxcbr0.accept_ra_defrtr = 1
net.ipv6.conf.lxcbr0.accept_ra_from_local = 0
net.ipv6.conf.lxcbr0.accept_ra_min_hop_limit = 1
net.ipv6.conf.lxcbr0.accept_ra_mtu = 1
net.ipv6.conf.lxcbr0.accept_ra_pinfo = 1
net.ipv6.conf.lxcbr0.accept_ra_rt_info_max_plen = 0
net.ipv6.conf.lxcbr0.accept_ra_rtr_pref = 1
net.ipv6.conf.lxcbr0.accept_redirects = 1
net.ipv6.conf.lxcbr0.accept_source_route = 0
net.ipv6.conf.lxcbr0.autoconf = 1
net.ipv6.conf.lxcbr0.dad_transmits = 1
net.ipv6.conf.lxcbr0.disable_ipv6 = 0
net.ipv6.conf.lxcbr0.force_mld_version = 0
net.ipv6.conf.lxcbr0.force_tllao = 0
net.ipv6.conf.lxcbr0.forwarding = 0
net.ipv6.conf.lxcbr0.hop_limit = 64
net.ipv6.conf.lxcbr0.ignore_routes_with_linkdown = 0
net.ipv6.conf.lxcbr0.max_addresses = 16
net.ipv6.conf.lxcbr0.max_desync_factor = 600
net.ipv6.conf.lxcbr0.mc_forwarding = 0
net.ipv6.conf.lxcbr0.mldv1_unsolicited_report_interval = 10000
net.ipv6.conf.lxcbr0.mldv2_unsolicited_report_interval = 1000
net.ipv6.conf.lxcbr0.mtu = 1500
net.ipv6.conf.lxcbr0.ndisc_notify = 0
net.ipv6.conf.lxcbr0.proxy_ndp = 0
net.ipv6.conf.lxcbr0.regen_max_retry = 3
net.ipv6.conf.lxcbr0.router_probe_interval = 60
net.ipv6.conf.lxcbr0.router_solicitation_delay = 1
net.ipv6.conf.lxcbr0.router_solicitation_interval = 4
net.ipv6.conf.lxcbr0.router_solicitations = 3
sysctl: reading key "net.ipv6.conf.lxcbr0.stable_secret"
net.ipv6.conf.lxcbr0.suppress_frag_ndisc = 1
net.ipv6.conf.lxcbr0.temp_prefered_lft = 86400
net.ipv6.conf.lxcbr0.temp_valid_lft = 604800
net.ipv6.conf.lxcbr0.use_oif_addrs_only = 0
net.ipv6.conf.lxcbr0.use_tempaddr = 2
net.ipv6.flowlabel_consistency = 1
net.ipv6.flowlabel_state_ranges = 0
net.ipv6.fwmark_reflect = 0
net.ipv6.icmp.ratelimit = 1000
net.ipv6.idgen_delay = 1
net.ipv6.idgen_retries = 3
net.ipv6.ip6frag_high_thresh = 4194304
net.ipv6.ip6frag_low_thresh = 3145728
net.ipv6.ip6frag_secret_interval = 0
net.ipv6.ip6frag_time = 60
net.ipv6.ip_nonlocal_bind = 0
net.ipv6.mld_max_msf = 64
net.ipv6.mld_qrv = 2
net.ipv6.neigh.default.anycast_delay = 100
net.ipv6.neigh.default.app_solicit = 0
net.ipv6.neigh.default.base_reachable_time_ms = 30000
net.ipv6.neigh.default.delay_first_probe_time = 5
net.ipv6.neigh.default.gc_interval = 30
net.ipv6.neigh.default.gc_stale_time = 60
```

(continues on next page)

(continued from previous page)

```
net.ipv6.neigh.default.gc_thresh1 = 128
net.ipv6.neigh.default.gc_thresh2 = 512
net.ipv6.neigh.default.gc_thresh3 = 1024
net.ipv6.neigh.default.locktime = 0
net.ipv6.neigh.default.mcast_resolicit = 0
net.ipv6.neigh.default.mcast_solicit = 3
net.ipv6.neigh.default.proxy_delay = 80
net.ipv6.neigh.default.proxy_qlen = 64
net.ipv6.neigh.default.retrans_time_ms = 1000
net.ipv6.neigh.default.ucast_solicit = 3
net.ipv6.neigh.default.unres_qlen = 31
net.ipv6.neigh.default.unres_qlen_bytes = 65536
net.ipv6.neigh.enp25s0f0.anycast_delay = 100
net.ipv6.neigh.enp25s0f0.app_solicit = 0
net.ipv6.neigh.enp25s0f0.base_reachable_time_ms = 30000
net.ipv6.neigh.enp25s0f0.delay_first_probe_time = 5
net.ipv6.neigh.enp25s0f0.gc_stale_time = 60
net.ipv6.neigh.enp25s0f0.locktime = 0
net.ipv6.neigh.enp25s0f0.mcast_resolicit = 0
net.ipv6.neigh.enp25s0f0.mcast_solicit = 3
net.ipv6.neigh.enp25s0f0.proxy_delay = 80
net.ipv6.neigh.enp25s0f0.proxy_qlen = 64
net.ipv6.neigh.enp25s0f0.retrans_time_ms = 1000
net.ipv6.neigh.enp25s0f0.ucast_solicit = 3
net.ipv6.neigh.enp25s0f0.unres_qlen = 31
net.ipv6.neigh.enp25s0f0.unres_qlen_bytes = 65536
net.ipv6.neigh.lo.anycast_delay = 100
net.ipv6.neigh.lo.app_solicit = 0
net.ipv6.neigh.lo.base_reachable_time_ms = 30000
net.ipv6.neigh.lo.delay_first_probe_time = 5
net.ipv6.neigh.lo.gc_stale_time = 60
net.ipv6.neigh.lo.locktime = 0
net.ipv6.neigh.lo.mcast_resolicit = 0
net.ipv6.neigh.lo.mcast_solicit = 3
net.ipv6.neigh.lo.proxy_delay = 80
net.ipv6.neigh.lo.proxy_qlen = 64
net.ipv6.neigh.lo.retrans_time_ms = 1000
net.ipv6.neigh.lo.ucast_solicit = 3
net.ipv6.neigh.lo.unres_qlen = 31
net.ipv6.neigh.lo.unres_qlen_bytes = 65536
net.ipv6.neigh.lxcbr0.anycast_delay = 100
net.ipv6.neigh.lxcbr0.app_solicit = 0
net.ipv6.neigh.lxcbr0.base_reachable_time_ms = 30000
net.ipv6.neigh.lxcbr0.delay_first_probe_time = 5
net.ipv6.neigh.lxcbr0.gc_stale_time = 60
net.ipv6.neigh.lxcbr0.locktime = 0
net.ipv6.neigh.lxcbr0.mcast_resolicit = 0
net.ipv6.neigh.lxcbr0.mcast_solicit = 3
net.ipv6.neigh.lxcbr0.proxy_delay = 80
net.ipv6.neigh.lxcbr0.proxy_qlen = 64
net.ipv6.neigh.lxcbr0.retrans_time_ms = 1000
net.ipv6.neigh.lxcbr0.ucast_solicit = 3
net.ipv6.neigh.lxcbr0.unres_qlen = 31
net.ipv6.neigh.lxcbr0.unres_qlen_bytes = 65536
net.ipv6.route.gc_elasticity = 9
net.ipv6.route.gc_interval = 30
net.ipv6.route.gc_min_interval = 0
net.ipv6.route.gc_min_interval_ms = 500
net.ipv6.route.gc_thresh = 1024
net.ipv6.route.gc_timeout = 60
net.ipv6.route.max_size = 4096
```

(continues on next page)

(continued from previous page)

```
net.ipv6.route.min_adv_mss = 1220
net.ipv6.route.mtu_expires = 600
net.ipv6.xfrm6_gc_thresh = 2147483647
net.netfilter.nf_conntrack_acct = 0
net.netfilter.nf_conntrack_buckets = 65536
net.netfilter.nf_conntrack_checksum = 1
net.netfilter.nf_conntrack_count = 2
net.netfilter.nf_conntrack_events = 1
net.netfilter.nf_conntrack_expect_max = 1024
net.netfilter.nf_conntrack_generic_timeout = 600
net.netfilter.nf_conntrack_helper = 1
net.netfilter.nf_conntrack_icmp_timeout = 30
net.netfilter.nf_conntrack_log_invalid = 0
net.netfilter.nf_conntrack_max = 262144
net.netfilter.nf_conntrack_tcp_be_liberal = 0
net.netfilter.nf_conntrack_tcp_loose = 1
net.netfilter.nf_conntrack_tcp_max_retrans = 3
net.netfilter.nf_conntrack_tcp_timeout_close = 10
net.netfilter.nf_conntrack_tcp_timeout_close_wait = 60
net.netfilter.nf_conntrack_tcp_timeout_established = 432000
net.netfilter.nf_conntrack_tcp_timeout_fin_wait = 120
net.netfilter.nf_conntrack_tcp_timeout_last_ack = 30
net.netfilter.nf_conntrack_tcp_timeout_max_retrans = 300
net.netfilter.nf_conntrack_tcp_timeout_syn_recv = 60
net.netfilter.nf_conntrack_tcp_timeout_syn_sent = 120
net.netfilter.nf_conntrack_tcp_timeout_time_wait = 120
net.netfilter.nf_conntrack_tcp_timeout_unacknowledged = 300
net.netfilter.nf_conntrack_timestamp = 0
net.netfilter.nf_conntrack_udp_timeout = 30
net.netfilter.nf_conntrack_udp_timeout_stream = 180
net.netfilter.nf_log.0 = NONE
net.netfilter.nf_log.1 = NONE
net.netfilter.nf_log.10 = NONE
net.netfilter.nf_log.11 = NONE
net.netfilter.nf_log.12 = NONE
net.netfilter.nf_log.2 = NONE
net.netfilter.nf_log.3 = NONE
net.netfilter.nf_log.4 = NONE
net.netfilter.nf_log.5 = NONE
net.netfilter.nf_log.6 = NONE
net.netfilter.nf_log.7 = NONE
net.netfilter.nf_log.8 = NONE
net.netfilter.nf_log.9 = NONE
net.nf_conntrack_max = 262144
net.unix.max_dgram_qlen = 512
vm.admin_reserve_kbytes = 8192
vm.block_dump = 0
vm.compact_unevictable_allowed = 1
vm.dirty_background_bytes = 0
vm.dirty_background_ratio = 10
vm.dirty_bytes = 0
vm.dirty_expire_centisecs = 3000
vm.dirty_ratio = 20
vm.dirty_writeback_centisecs = 500
vm.dirtytime_expire_seconds = 43200
vm.drop_caches = 0
vm.extfrag_threshold = 500
vm.hugepages_treat_as_movable = 0
vm.hugetlb_shm_group = 0
vm.laptop_mode = 0
vm.legacy_va_layout = 0
```

(continues on next page)

(continued from previous page)

```

vm.lowmem_reserve_ratio = 256      256      32      1
vm.max_map_count = 200000
vm.memory_failure_early_kill = 0
vm.memory_failure_recovery = 1
vm.min_free_kbytes = 90112
vm.min_slab_ratio = 5
vm.min_unmapped_ratio = 1
vm.mmap_min_addr = 65536
vm.nr_hugepages = 4096
vm.nr_hugepages_mempolicy = 4096
vm.nr_overcommit_hugepages = 0
vm.nr_pdflush_threads = 0
vm.numa_zonelist_order = default
vm.oom_dump_tasks = 1
vm.oom_kill_allocating_task = 0
vm.overcommit_kbytes = 0
vm.overcommit_memory = 0
vm.overcommit_ratio = 50
vm.page-cluster = 3
vm.panic_on_oom = 0
vm.percpu_pagelist_fraction = 0
vm.stat_interval = 1
vm.swappiness = 0
vm.user_reserve_kbytes = 131072
vm.vfs_cache_pressure = 100
vm.zone_reclaim_mode = 0

```

Services listing

```

$ service --status-all
[ + ] apparmor
[ - ] bootmisc.sh
[ - ] checkfs.sh
[ - ] checkroot-bootclean.sh
[ - ] checkroot.sh
[ + ] console-setup
[ + ] cpufrequtils
[ + ] cron
[ + ] dbus
[ + ] docker
[ + ] ebttables
[ + ] grub-common
[ - ] hostname.sh
[ - ] hwclock.sh
[ + ] keyboard-setup
[ - ] killprocs
[ + ] kmod
[ + ] loadcpufreq
[ + ] lxcfs
[ - ] mountall-bootclean.sh
[ - ] mountall.sh
[ - ] mountdevsubfs.sh
[ - ] mountkernfs.sh
[ - ] mountnfs-bootclean.sh
[ - ] mountnfs.sh
[ + ] networking
[ - ] ondemand
[ - ] plymouth
[ - ] plymouth-log
[ + ] procps
[ + ] qemu-kvm

```

(continues on next page)

(continued from previous page)

```

[ - ] rc.local
[ + ] resolvconf
[ - ] rsync
[ + ] rsyslog
[ - ] screen-cleanup
[ - ] sendsigs
[ + ] ssh
[ + ] udev
[ - ] umountfs
[ - ] umountnfs.sh
[ - ] umountroot
[ + ] urandom
[ - ] uuid
[ - ] x11-common

```

Host CFS optimizations (QEMU+VPP)

Applying CFS scheduler tuning on all Qemu vcpu worker threads (those are handling testpmd - pmd threads) and VPP PMD worker threads. List of VPP PMD threads can be obtained e.g. from:

```

$ for psid in $(pgrep vpp)
$ do
$   for tid in $(ps -Lo tid --pid $psid | grep -v TID)
$   do
$     echo $tid
$   done
$ done

```

Or:

```
$ cat /proc/`pidof vpp`/task/*/stat | awk '{print $1" "$2" "$39}'
```

Applying Round-robin scheduling with highest priority

```

$ for psid in $(pgrep vpp)
$ do
$   for tid in $(ps -Lo tid --pid $psid | grep -v TID)
$   do
$     chrt -r -p 1 $tid
$   done
$ done

```

More information about Linux CFS can be found in: [Sched manual pages](#)¹⁴¹.

Host IRQ affinity

Changing the default pinning of every IRQ to core 0. (Same does apply on both guest VM and host OS)

```
$ for l in `ls /proc/irq`; do echo 1 | sudo tee /proc/irq/$l/smp_affinity; done
```

Host RCU affinity

Changing the default pinning of RCU to core 0. (Same does apply on both guest VM and host OS)

```
$ for i in `pgrep rcu[^c]` ; do sudo taskset -pc 0 $i ; done
```

Host Writeback affinity

Changing the default pinning of writebacks to core 0. (Same does apply on both guest VM and host OS)

¹⁴¹ <http://man7.org/linux/man-pages/man7/sched.7.html>

```
$ echo 1 | sudo tee /sys/bus/workqueue/devices/writeback/cpumask
```

2.9.4 DUT Configuration - VPP

VPP Version

VPP-18.04 release

VPP Compile Parameters

[FD.io VPP compile job](#)¹⁴²

VPP Install Parameters

```
$ dpkg -i --force-all vpp*
```

VPP Startup Configuration

VPP startup configuration changes per test case with different settings for CPU cores, rx-queues and no-multi-seg parameter. Startup config is aligned with applied test case tag:

Tagged by **1T1C**

```
ip
{
  heap-size 4G
}
unix
{
  cli-listen localhost:5002
  log /tmp/vpe.log
  nodaemon
}
ip6
{
  heap-size 4G
  hash-buckets 2000000
}
heapsize 4G
plugins
{
  plugin default
  {
    disable
  }
  plugin dpdk_plugin.so
  {
    enable
  }
}
cpu
{
  corelist-workers 2
  main-core 1
}
dpdk
{
  dev 0000:0a:00.0
  dev 0000:0a:00.1
  no-multi-seg
```

(continues on next page)

¹⁴² <https://jenkins.fd.io/view/vpp/job/vpp-merge-1804-ubuntu1604/>

(continued from previous page)

```

uio-driver uio_pci_generic
log-level debug
dev default
{
  num-rx-queues 1
}
socket-mem 1024,1024
no-tx-checksum-offload
}

```

Tagged by 2T2C

```

ip
{
  heap-size 4G
}
unix
{
  cli-listen localhost:5002
  log /tmp/vpe.log
  nodaemon
}
ip6
{
  heap-size 4G
  hash-buckets 2000000
}
heapsize 4G
plugins
{
  plugin default
  {
    disable
  }
  plugin dpdk_plugin.so
  {
    enable
  }
}
cpu
{
  corelist-workers 2,3
  main-core 1
}
dpdk
{
  dev 0000:0a:00.0
  dev 0000:0a:00.1
  no-multi-seg
  uio-driver uio_pci_generic
  log-level debug
  dev default
  {
    num-rx-queues 1
  }
  socket-mem 1024,1024
  no-tx-checksum-offload
}

```

Tagged by 4T4C

```
ip
{
  heap-size 4G
}
unix
{
  cli-listen localhost:5002
  log /tmp/vpe.log
  nodaemon
}
ip6
{
  heap-size 4G
  hash-buckets 2000000
}
heapsize 4G
plugins
{
  plugin default
  {
    disable
  }
  plugin dpdk_plugin.so
  {
    enable
  }
}
cpu
{
  corelist-workers 2,3,4,5
  main-core 1
}
dpdk
{
  dev 0000:0a:00.0
  dev 0000:0a:00.1
  no-multi-seg
  uio-driver uio_pci_generic
  log-level debug
  dev default
  {
    num-rx-queues 1
  }
  socket-mem 1024,1024
  no-tx-checksum-offload
}
```

2.9.5 TG Configuration - TRex

TG Version

TRex v2.35

DPDK version

DPDK v17.11

TG Build Script used

[TRex intallation](#)¹⁴³

TG Startup Configuration

¹⁴³ https://git.fd.io/csit/tree/resources/tools/trex/trex_installer.sh?h=rls1804

```
$ cat /etc/trex_cfg.yaml
- port_limit      : 2
  version         : 2
  interfaces      : ["0000:0d:00.0", "0000:0d:00.1"]
  port_info       :
    - dest_mac    : [0x3c,0xfd,0xfe,0x9c,0xee,0xf5]
      src_mac     : [0x3c,0xfd,0xfe,0x9c,0xee,0xf4]
    - dest_mac    : [0x3c,0xfd,0xfe,0x9c,0xee,0xf4]
      src_mac     : [0x3c,0xfd,0xfe,0x9c,0xee,0xf5]
```

TG Startup Command

```
$ sh -c 'cd <t-rex-install-dir>/scripts/ && sudo nohup ./t-rex-64 -i -c 7 --iom 0 > /tmp/trex.log 2>
->&1 &'> /dev/null
```

TG common API - pointer to driver

TRex driver¹⁴⁴

2.10 Documentation

2.10.1 Container Orchestration in CSIT

Overview

Linux Containers

Linux Containers is an OS-level virtualization method for running multiple isolated Linux systems (containers) on a compute host using a single Linux kernel. Containers rely on Linux kernel cgroups functionality for controlling usage of shared system resources (i.e. CPU, memory, block I/O, network) and for namespace isolation. The latter enables complete isolation of applications' view of operating environment, including process trees, networking, user IDs and mounted file systems.

LXC (Linux Containers) combine kernel's cgroups and support for isolated namespaces to provide an isolated environment for applications. Docker does use LXC as one of its execution drivers, enabling image management and providing deployment services. More information in *[lxc]* (page 301), *[lxc-namespace]* (page 301) and *[stgraber]* (page 301).

Linux containers can be of two kinds: privileged containers and unprivileged containers.

Unprivileged Containers

Running unprivileged containers is the safest way to run containers in a production environment. From LXC 1.0 one can start a full system container entirely as a user, allowing to map a range of UIDs on the host into a namespace inside of which a user with UID 0 can exist again. In other words an unprivileged container does mask the userid from the host, making it impossible to gain a root access on the host even if a user gets root in a container. With unprivileged containers, non-root users can create containers and will appear in the container as the root, but will appear as userid <non-zero> on the host. Unprivileged containers are also better suited to supporting multi-tenancy operating environments. More information in *[lxc-security]* (page 301) and *[stgraber]* (page 301).

¹⁴⁴ https://git.fd.io/csit/tree/resources/tools/trex/trex_stateless_profile.py?h=rls1804

Privileged Containers

Privileged containers do not mask UIDs, and container UID 0 is mapped to the host UID 0. Security and isolation is controlled by a good configuration of cgroup access, extensive AppArmor profile preventing the known attacks as well as container capabilities and SELinux. Here a list of applicable security control mechanisms:

- Capabilities - keep (whitelist) or drop (blacklist) Linux capabilities, *[capabilities]* (page 301).
- Control groups - cgroups, resource bean counting, resource quotas, access restrictions, *[cgroup1]* (page 301), *[cgroup2]* (page 301).
- AppArmor - apparmor profiles aim to prevent any of the known ways of escaping a container or cause harm to the host, *[apparmor]* (page 301).
- SELinux - Security Enhanced Linux is a Linux kernel security module that provides similar function to AppArmor, supporting access control security policies including United States Department of Defense-style mandatory access controls. Mandatory access controls allow an administrator of a system to define how applications and users can access different resources such as files, devices, networks and inter- process communication, *[selinux]* (page 301).
- Seccomp - secure computing mode, enables filtering of system calls, *[seccomp]* (page 301).

More information in *[lxc-security]* (page 301) and *[lxc-sec-features]* (page 301).

Linux Containers in CSIT

CSIT is using Privileged Containers as the `sysfs` is mounted with RW access. `sysfs` is required to be mounted as RW due to VPP accessing `/sys/bus/pci/drivers/uis_pci_generic/unbind`. This is not the case of unprivileged containers where `sysfs` is mounted as read-only.

Orchestrating Container Lifecycle Events

Following Linux container lifecycle events need to be addressed by an orchestration system:

1. Acquire - acquiring/downloading existing container images via `docker pull` or `lxc-create -t download`.
2. Build - building a container image from scratch or another container image via `docker build <dockerfile/composefile>` or customizing LXC templates in '<https://github.com/lxc/lxc/tree/master/templates>'
3. (Re-)Create - creating a running instance of a container application from anew, or re-creating one that failed. A.k.a. (re-)deploy via `docker run` or `lxc-start`
4. Execute - execute system operations within the container by attaching to running container. THIS is done by `lxc-attach` or `docker exec`
5. Distribute - distributing pre-built container images to the compute nodes. Currently not implemented in CSIT.

Container Orchestration Systems Used in CSIT

Current CSIT testing framework integrates following Linux container orchestration mechanisms:

- LXC/Docker for complete VPP container lifecycle control.
- Combination of Kubernetes (container orchestration), Docker (container images) and Ligato (container networking).

LXC

LXC is the well-known and heavily tested low-level Linux container runtime [*lxc-source*] (page 301), that provides a userspace interface for the Linux kernel containment features. With a powerful API and simple tools, LXC enables Linux users to easily create and manage system or application containers. LXC uses following kernel features to contain processes:

- Kernel namespaces: ipc, uts, mount, pid, network and user.
- AppArmor and SELinux security profiles.
- Seccomp policies.
- Chroot.
- Cgroups.

CSIT uses LXC runtime and LXC usertools to test VPP data plane performance in a range of virtual networking topologies.

Known Issues

- Current CSIT restriction: only single instance of lxc runtime due to the cgroup policies used in CSIT. There is plan to add the capability into code to create cgroups per container instance to address this issue. This sort of functionality is better supported in LXC 2.1 but can be done in current version as well.
- CSIT code is currently using cgroup to control the range of CPU cores the LXC container runs on. VPP thread pinning is defined in vpp startup.conf.

Docker

Docker builds on top of Linux kernel containment features, and offers a high-level tool for wrapping the processes, maintaining and executing them in containers [*docker*] (page 301). Currently it uses *runc* a CLI tool for spawning and running containers according to the [OCI specification](#)¹⁴⁵

A Docker container image is a lightweight, stand-alone, executable package of a piece of software that includes everything needed to run it: code, runtime, system tools, system libraries, settings.

CSIT uses Docker to manage the maintenance and execution of containerized applications used in CSIT performance tests.

- Data plane thread pinning to CPU cores - Docker CLI and/or Docker configuration file controls the range of CPU cores the Docker image must run on. VPP thread pinning defined in vpp startup.conf.

Kubernetes

Kubernetes [*k8s-doc*] (page 301), or K8s, is a production-grade container orchestration platform for automating the deployment, scaling and operating application containers. Kubernetes groups containers that make up an application into logical units, pods, for easy management and discovery. K8s pod definitions including compute resource allocation is provided in .yaml files.

CSIT uses K8s and its infrastructure components like etcd to control all phases of container based virtualized network topologies.

Ligato

Ligato [*ligato*] (page 301) is an open-source project developing a set of cloud-native tools for orchestrating container networking. Ligato integrates with FD.io VPP using goVPP [*govpp*] (page 301) and vpp-agent [*vpp-agent*] (page 302).

¹⁴⁵ <https://www.opencontainers.org/>

Known Issues

- Currently using a separate LF Jenkins job for building csit-centric prod_vpp_agent docker images vs. dockerhub/ligato ones.

Implementation

CSIT container orchestration is implemented in CSIT Level-1 keyword Python libraries following the Builder design pattern. Builder design pattern separates the construction of a complex object from its representation, so that the same construction process can create different representations e.g. LXC, Docker, other.

CSIT Robot Framework keywords are then responsible for higher level lifecycle control of of the named container groups. One can have multiple named groups, with 1..N containers in a group performing different role/functionality e.g. NFs, Switch, Kafka bus, ETCD datastore, etc. ContainerManager class acts as a Director and uses ContainerEngine class that encapsulate container control.

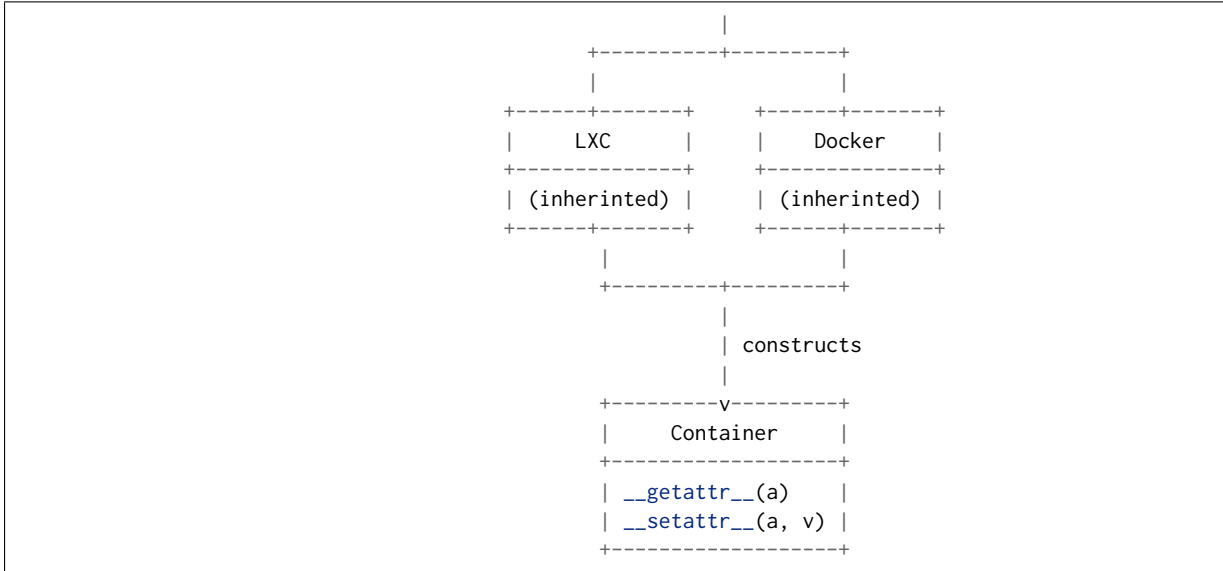
Current CSIT implementation is illustrated using UML Class diagram:

1. Acquire
2. Build
3. (Re-)Create
4. Execute

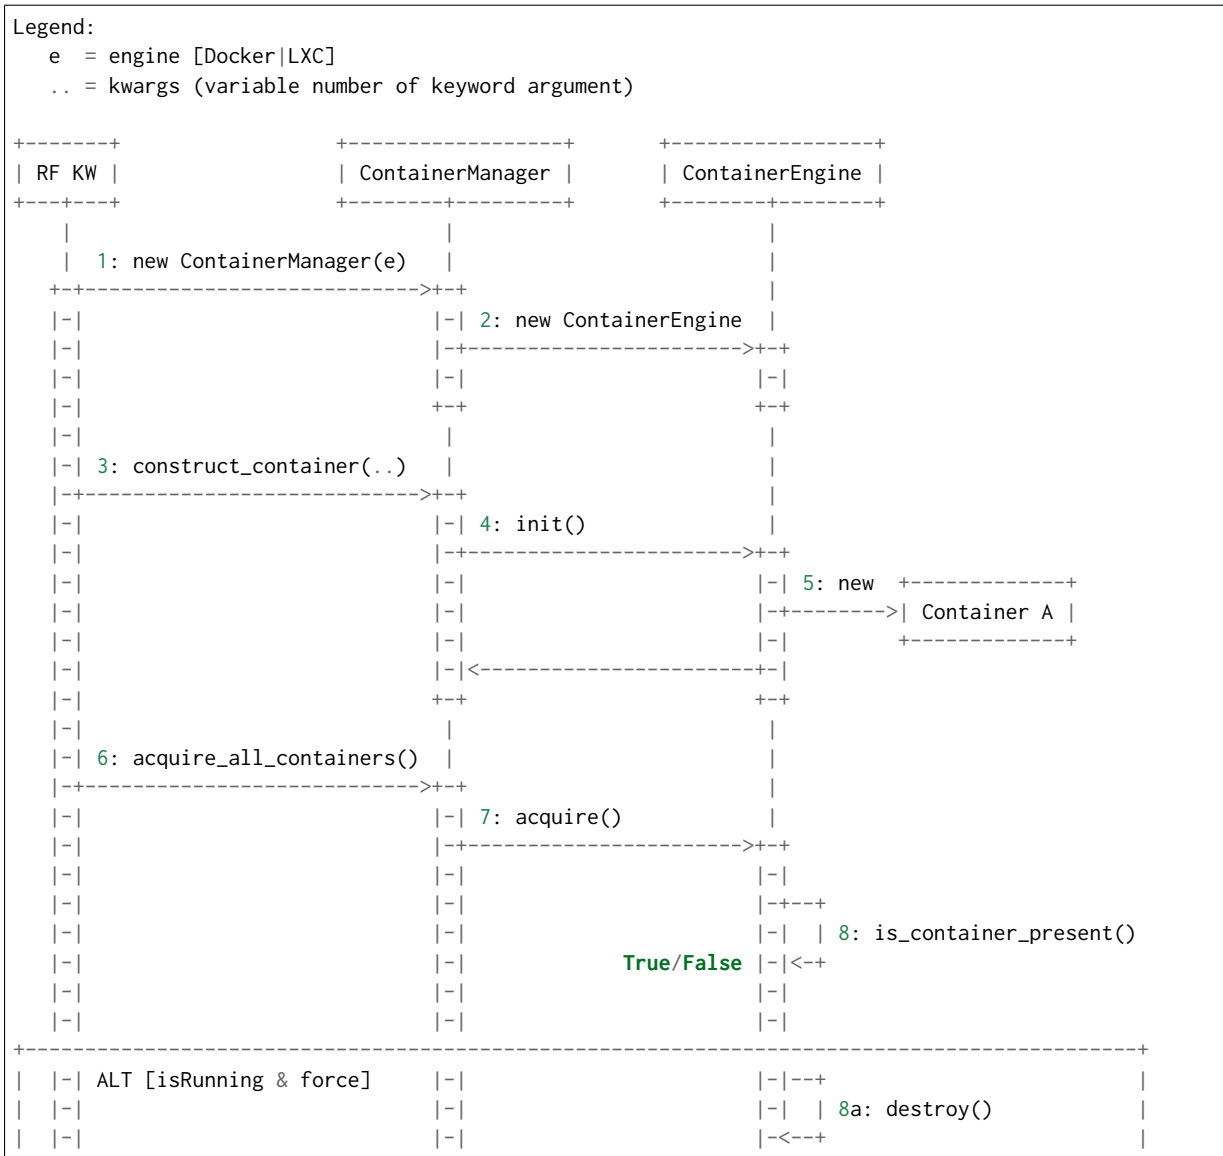


(continues on next page)

(continued from previous page)

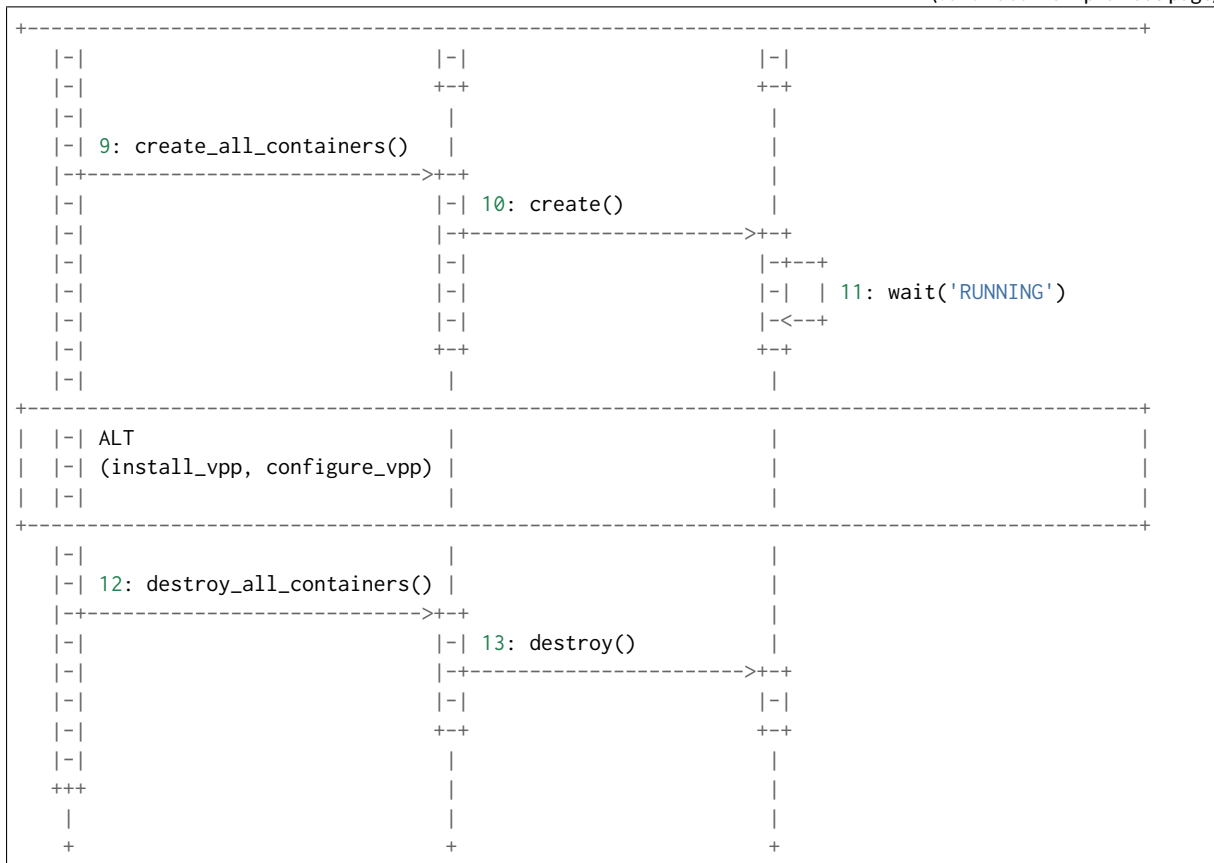


Sequential diagram that illustrates the creation of a single container.



(continues on next page)

(continued from previous page)



Container Data Structure

Container is represented in Python L1 library as a separate Class with instance variables and no methods except overridden `__getattr__` and `__setattr__`. Instance variables are assigned to container dynamically during the `construct_container(**kwargs)` call and are passed down from the RF keyword.

Usage example:

```

| Construct VNF containers on all DUTs
| | [Arguments] | ${technology} | ${image} | ${cpu_count}=${1} | ${count}=${1}
| | ...
| | ${group}= | Set Variable | VNF
| | ${guest_dir}= | Set Variable | /mnt/host
| | ${host_dir}= | Set Variable | /tmp
| | ${skip_cpus}= | Evaluate | ${vpp_cpus}+${system_cpus}
| | Import Library | resources.libraries.python.ContainerUtils.ContainerManager
| | ... | engine=${technology} | WITH NAME | ${group}
| | ${duts}= | Get Matches | ${nodes} | DUT*
| | :FOR | ${dut} | IN | @${duts}
| | {env}= | Create List | LC_ALL="en_US.UTF-8"
| | | ... | DEBIAN_FRONTEND=noninteractive | ETCDV3_ENDPOINTS=172.17.0.1:2379
| | | ${cpu_node}= | Get interfaces numa node | ${nodes['${dut}']}
| | | ... | ${dut1_if1} | ${dut1_if2}
| | | Run Keyword | ${group}.Construct containers
| | | ... | name=${dut}_${group}
| | | ... | node=${nodes['${dut}']}
| | | ... | host_dir=${host_dir}
| | | ... | guest_dir=${guest_dir}
| | | ... | image=${image}

```

(continues on next page)

(continued from previous page)

```

| | | ... | cpu_count=${cpu_count}
| | | ... | cpu_skip=${skip_cpus}
| | | ... | smt_used=${False}
| | | ... | cpuset_mems=${cpu_node}
| | | ... | cpu_shared=${False}
| | | ... | env=${env}

```

Mandatory parameters to create standalone container are: node, name, image [*image-var*] (page 302), cpu_count, cpu_skip, smt_used, cpuset_mems, cpu_shared.

There is no parameters check functionality. Passing required arguments is in coder responsibility. All the above parameters are required to calculate the correct cpu placement. See documentation for the full reference.

Kubernetes

Kubernetes is implemented as separate library `KubernetesUtils.py`, with a class with the same name. This utility provides an API for L2 Robot Keywords to control `kubect1` installed on each of DUTs. One time initialization script, `resources/libraries/bash/k8s_setup.sh` does reset/init `kubect1`, applies Calico v2.6.3 and initializes the `csit` namespace. CSIT namespace is required to not to interfere with existing setups and it further simplifies `apply/get/delete Pod/ConfigMap` operations on SUTs.

Kubernetes utility is based on YAML templates to avoid crafting the huge YAML configuration files, what would lower the readability of code and requires complicated algorithms. The templates can be found in `resources/templates/kubernetes` and can be leveraged in the future for other separate tasks.

Two types of YAML templates are defined:

- Static - do not change between deployments, that is infrastructure containers like Kafka, Calico, ETCD.
- Dynamic - per test suite/case topology YAML files e.g. `SFC_controller`, `VNF`, `VSWITCH`.

Making own python wrapper library of `kubect1` instead of using the official Python package allows to control and deploy environment over the SSH library without the need of using isolated driver running on each of DUTs.

Ligato

Ligato integration does require to compile the `vpp-agent` tool and build the bundled Docker image. Compilation of `vpp-agent` depends on specific VPP. In `ligato/vpp-agent` repository there are well prepared scripts for building the Docker image. Building docker image is possible via series of commands:

```

git clone https://github.com/ligato/vpp-agent
cd vpp_agent/docker/dev_vpp_agent
sudo docker build -t dev_vpp_agent --build-arg AGENT_COMMIT=<agent commit id>\
  --build-arg VPP_COMMIT=<vpp commit id> --no-cache .
sudo ./shrink.sh
cd ../prod_vpp_agent
sudo ./build.sh
sudo ./shrink.sh

```

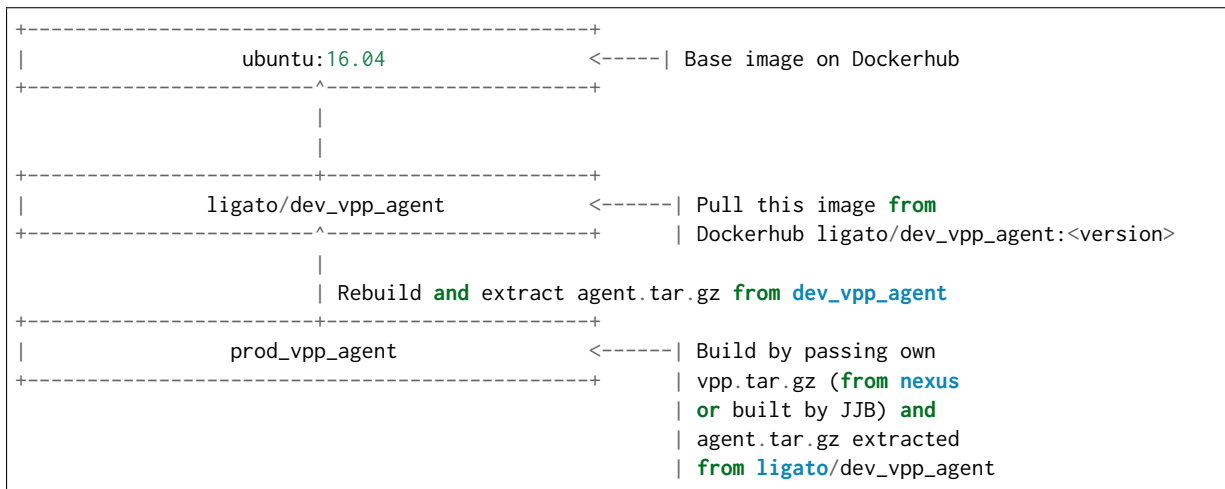
CSIT requires Docker image to include the desired VPP version (per patch testing, nightly testing, on demand testing).

The entire build process of building `dev_vpp_agent` image heavily depends on internet connectivity and also takes a significant amount of time (~1-1.5h based on internet bandwidth and allocated resources). The optimal solution would be to build the image on jenkins slave, transfer the Docker image to DUTs and execute separate suite of tests.

To address the amount of time required to build dev_vpp_agent image, we can pull existing specific version of `dev_vpp_agent` and extract the `vpp-agent` from it.

We created separate sets of Jenkins jobs, that will be executing following:

1. Clone latest CSIT and Ligato repositories.
2. Pull specific version of dev_vpp_agent image from Dockerhub.
3. Extract VPP API (from .deb package) and copy into dev_vpp_agent image
4. Rebuild vpp-agent and extract outside image.
5. Build prod_vpp_image Docker image from dev_vpp_agent image.
6. Transfer prod_vpp_agent image to DUTs.
7. Execute subset of performance tests designed for Ligato testing.



Approximate size of vnf-agent docker images:

| REPOSITORY | TAG | IMAGE ID | CREATED | SIZE |
|----------------|--------|--------------|-------------|--------|
| dev-vpp-agent | latest | 78c53bd57e2 | 6 weeks ago | 9.79GB |
| prod_vpp_agent | latest | f68af5afe601 | 5 weeks ago | 443MB |

In CSIT we need to create separate performance suite under tests/kubernetes/perf which contains modified Suite setup in comparison to standard perf tests. This is due to reason that VPP will act as vswitch in Docker image and not as standalone installed service.

Tested Topologies

Listed CSIT container networking test topologies are defined with DUT containerized VPP switch forwarding packets between NF containers. Each NF container runs their own instance of VPP in L2XC configuration.

Following container networking topologies are tested in CSIT rls1804:

- LXC topologies:
 - eth-l2xcbase-eth-2memif-1lxc.
 - eth-l2bdbasemaclrn-eth-2memif-1lxc.
- Docker topologies:
 - eth-l2xcbase-eth-2memif-1docker.
- Kubernetes/Ligato topologies:
 - eth-1drcl2bdbasemaclrn-eth-2memif-1drcl2xc-1paral

- eth-1drcl2bdbasemaclrn-eth-2memif-2drcl2xc-1horiz
- eth-1drcl2bdbasemaclrn-eth-2memif-4drcl2xc-1horiz
- eth-1drcl2bdbasemaclrn-eth-4memif-2drcl2xc-1chain
- eth-1drcl2bdbasemaclrn-eth-8memif-4drcl2xc-1chain
- eth-1drcl2xcbase-eth-2memif-1drcl2xc-1paral
- eth-1drcl2xcbase-eth-2memif-2drcl2xc-1horiz
- eth-1drcl2xcbase-eth-2memif-4drcl2xc-1horiz
- eth-1drcl2xcbase-eth-4memif-2drcl2xc-1chain
- eth-1drcl2xcbase-eth-8memif-4drcl2xc-1chain

References

2.10.2 VPP Performance Tests


CSIT VPP Performance Tests Documentation¹⁶² contains detailed functional description and input parameters for each test case.

¹⁶² <https://docs.fd.io/csit/rls1804/doc/tests.vpp.perf.html>

3.1 Overview

3.1.1 Tested Physical Topologies

CSIT DPDK performance tests are executed on physical baremetal servers hosted by LF FD.io project. Testbed physical topology is shown in the figure below.:



SUT1 and SUT2 are two System Under Test servers (Cisco UCS C240, each with two Intel XEON CPUs), TG is a Traffic Generator (TG, another Cisco UCS C240, with two Intel XEON CPUs). SUTs run Testpmd/L3FWD SW application in Linux user-mode as a Device Under Test (DUT). TG runs TRex SW application as a packet Traffic Generator. Physical connectivity between SUTs and to TG is provided using different NIC models that need to be tested for performance. Currently installed and tested NIC models include:

1. 2port10GE X520-DA2 Intel.
2. 2port10GE X710 Intel.
3. 2port10GE VIC1227 Cisco.
4. 2port40GE VIC1385 Cisco.

5. 2port40GE XL710 Intel.

From SUT and DUT perspective, all performance tests involve forwarding packets between two physical Ethernet ports (10GE or 40GE). Due to the number of listed NIC models tested and available PCI slot capacity in SUT servers, in all of the above cases both physical ports are located on the same NIC. In some test cases this results in measured packet throughput being limited not by VPP DUT but by either the physical interface or the NIC capacity.

Going forward CSIT project will be looking to add more hardware into FD.io performance labs to address larger scale multi-interface and multi-NIC performance testing scenarios.

Note that reported DUT (DPDK) performance results are specific to the SUTs tested. Current LF FD.io SUTs are based on Intel XEON E5-2699v3 2.3GHz CPUs. SUTs with other CPUs are likely to yield different results. A good rule of thumb, that can be applied to estimate DPDK packet throughput for Phy-to-Phy (NIC-to-NIC, PCI-to-PCI) topology, is to expect the forwarding performance to be proportional to CPU core frequency, assuming CPU is the only limiting factor and all other SUT parameters equivalent to FD.io CSIT environment. The same rule of thumb can be also applied for Phy-to-VM/LXC-to-Phy (NIC-to-VM/LXC-to-NIC) topology, but due to much higher dependency on intensive memory operations and sensitivity to Linux kernel scheduler settings and behaviour, this estimation may not always yield good enough accuracy.

For detailed LF FD.io test bed specification and physical topology please refer to [LF FD.io CSIT testbed wiki page](#)¹⁶³.

3.1.2 Performance Tests Coverage

Performance tests are split into two main categories:

- Throughput discovery - discovery of packet forwarding rate using binary search in accordance to [RFC 2544](#)¹⁶⁴.
 - NDR - discovery of Non Drop Rate packet throughput, at zero packet loss; followed by one-way packet latency measurements at 10%, 50% and 100% of discovered NDR throughput.
 - PDR - discovery of Partial Drop Rate, with specified non-zero packet loss currently set to 0.5%; followed by one-way packet latency measurements at 100% of discovered PDR throughput.
- Throughput verification - verification of packet forwarding rate against previously discovered throughput rate. These tests are currently done against 0.9 of reference NDR, with reference rates updated periodically.

CSIT rls1804 includes following performance test suites, listed per NIC type:

- 2port10GE X520-DA2 Intel
 - **L2IntLoop** - L2 Interface Loop forwarding any Ethernet frames between two Interfaces.
- 2port40GE XL710 Intel
 - **L2IntLoop** - L2 Interface Loop forwarding any Ethernet frames between two Interfaces.
- 2port10GE X520-DA2 Intel
 - **IPv4 Routed Forwarding** - L3 IP forwarding of Ethernet frames between two Interfaces.

Execution of performance tests takes time, especially the throughput discovery tests. Due to limited HW testbed resources available within FD.io labs hosted by Linux Foundation, the number of tests for NICs other than X520 (a.k.a. Niantic) has been limited to few baseline tests. Over time we expect the HW testbed resources to grow, and will be adding complete set of performance tests for all models of hardware to be executed regularly and(or) continuously.

¹⁶³ https://wiki.fd.io/view/CSIT/CSIT_LF_testbed

¹⁶⁴ <https://tools.ietf.org/html/rfc2544.html>

3.1.3 Performance Tests Naming

CSIT rls1804 follows a common structured naming convention for all performance and system functional tests, introduced in CSIT rls1701.

The naming should be intuitive for majority of the tests. Complete description of CSIT test naming convention is provided on *CSIT Test Naming* (page 261).

3.1.4 Methodology: Multi-Core and Multi-Threading

Intel Hyper-Threading - CSIT rls1804 performance tests are executed with SUT servers' Intel XEON processors configured in Intel Hyper-Threading Disabled mode (BIOS setting). This is the simplest configuration used to establish baseline single-thread single-core application packet processing and forwarding performance. Subsequent releases of CSIT will add performance tests with Intel Hyper-Threading Enabled (requires BIOS settings change and hard reboot of server).

Multi-core Tests - CSIT rls1804 multi-core tests are executed in the following VPP thread and core configurations:

1. 1t1c - 1 pmd worker thread on 1 CPU physical core.
2. 2t2c - 2 pmd worker threads on 2 CPU physical cores.

Note that in many tests running Testpmd/L3FWD reaches tested NIC I/O bandwidth or packets-per-second limit.

3.1.5 Methodology: Packet Throughput

Following values are measured and reported for packet throughput tests:

- NDR binary search per [RFC 2544](https://tools.ietf.org/html/rfc2544)¹⁶⁵:
 - Packet rate: "RATE: <aggregate packet rate in packets-per-second> pps (2x <per direction packets-per-second>)"
 - Aggregate bandwidth: "BANDWIDTH: <aggregate bandwidth in Gigabits per second> Gbps (untagged)"
- PDR binary search per [RFC 2544](https://tools.ietf.org/html/rfc2544)¹⁶⁶:
 - Packet rate: "RATE: <aggregate packet rate in packets-per-second> pps (2x <per direction packets-per-second>)"
 - Aggregate bandwidth: "BANDWIDTH: <aggregate bandwidth in Gigabits per second> Gbps (untagged)"
 - Packet loss tolerance: "LOSS_ACCEPTANCE <accepted percentage of packets lost at PDR rate>"
- NDR and PDR are measured for the following L2 frame sizes:
 - IPv4: 64B, 1518B, 9000B.

All rates are reported from external Traffic Generator perspective.

3.1.6 Methodology: Packet Latency

TREx Traffic Generator (TG) is used for measuring latency of Testpmd DUTs. Reported latency values are measured using following methodology:

¹⁶⁵ <https://tools.ietf.org/html/rfc2544.html>

¹⁶⁶ <https://tools.ietf.org/html/rfc2544.html>

- Latency tests are performed at 10%, 50% of discovered NDR rate (non drop rate) for each NDR throughput test and packet size (except IMIX).
- TG sends dedicated latency streams, one per direction, each at the rate of 10kpps at the prescribed packet size; these are sent in addition to the main load streams.
- TG reports min/avg/max latency values per stream direction, hence two sets of latency values are reported per test case; future release of TRex is expected to report latency percentiles.
- Reported latency values are aggregate across two SUTs due to three node topology used for all performance tests; for per SUT latency, reported value should be divided by two.
- 1usec is the measurement accuracy advertised by TRex TG for the setup used in FD.io labs used by CSIT project.
- TRex setup introduces an always-on error of about 2*2usec per latency flow - additional Tx/Rx interface latency induced by TRex SW writing and reading packet timestamps on CPU cores without HW acceleration on NICs closer to the interface line.

3.1.7 Methodology: TRex Traffic Generator Usage

The **TRex traffic generator**¹⁶⁷ is used for all CSIT performance tests. TRex stateless mode is used to measure NDR and PDR throughputs using binary search (NDR and PDR discovery tests) and for quick checks of DUT performance against the reference NDRs (NDR check tests) for specific configuration.

TRex is installed and run on the TG compute node. The typical procedure is:

- If the TRex is not already installed on TG, it is installed in the suite setup phase - see **TRex installation**¹⁶⁸.
- TRex configuration is set in its configuration file

```
/etc/trex_cfg.yaml
```

- TRex is started in the background mode

```
$ sh -c 'cd <t-rex-install-dir>/scripts/ && sudo nohup ./t-rex-64 -i -c 7 --iom 0 > /tmp/trex.  
↪log 2>&1 &' > /dev/null
```

- There are traffic streams dynamically prepared for each test, based on traffic profiles. The traffic is sent and the statistics obtained using `trex_stl_lib.api.STLClient`.

Measuring packet loss

- Create an instance of `STLClient`
- Connect to the client
- Add all streams
- Clear statistics
- Send the traffic for defined time
- Get the statistics

If there is a warm-up phase required, the traffic is sent also before test and the statistics are ignored.

Measuring latency

If measurement of latency is requested, two more packet streams are created (one for each direction) with TRex `flow_stats` parameter set to `STLFlowLatencyStats`. In that case, returned statistics will also include min/avg/max latency values.

¹⁶⁷ <https://wiki.fd.io/view/TRex>

¹⁶⁸ https://git.fd.io/csit/tree/resources/tools/trex/trex_installer.sh?h=rls1804

3.2 CSIT Release Notes

3.2.1 Changes in CSIT rls1804

No code changes apart from bug fixes.

3.2.2 Known Issues

Here is the list of known issues in CSIT rls1804 for Testpmd performance tests:

| # | Issue | Jira ID | Description |
|---|---|----------|--|
| 1 | Testpmd in 1t1c and 2t2c setups - large variation of discovered NDR throughput values across multiple test runs with xl710 and x710 NICs. | CSIT-569 | Suspected NIC firmware or DPDK driver issue affecting NDR throughput. Applies to XL710 and X710 NICs, no issues observed on x520 NICs. |
| 2 | Lower than expected NDR throughput with xl710 and x710 NICs, compared to x520 NICs. | CSIT-571 | Suspected NIC firmware or DPDK driver issue affecting NDR throughput. Applies to XL710 and X710 NICs. |

3.3 Packet Throughput Graphs

Plotted results are generated by multiple executions of the same CSIT performance tests across three physical testbeds within LF FD.io labs. To provide a descriptive summary view, Box-and-Whisker plots are used to display variation in measured throughput values, without making any assumptions of the underlying statistical distribution.

For each plotted test case, Box-and-Whisker plots show the quartiles (Min, 1st quartile / 25th percentile, 2nd quartile / 50th percentile / mean, 3rd quartile / 75th percentile, Max) across collected data set (data set size stated in the note below). Outliers are plotted as individual points. Min and max values are plotted as bottom and top Whiskers respectively. 2nd and 3rd quartiles are plotted as bottom and top edge of the box. If multiple samples match only two values, and all samples fall between them, then no whiskers are plotted. If all samples have the same value, only a horizontal line is plotted.

Title of each graph is a regex (regular expression) matching all throughput test cases plotted on this graph, *X-axis labels* are indices of individual test suites executed by [FD.io test executor dpdk performance jobs](#)¹⁶⁹ jobs that created result output files used as data sources for the graph, *Y-axis labels* are measured Packets Per Second [pps] values, and the *Graph legend* lists the plotted test suites and their indices.

Note: Test results have been generated by [FD.io test executor dpdk performance jobs](#)¹⁷⁰ with Robot Framework result files csit-dpdk-perf-*.zip [archived here](#). Plotted data set size per test case is equal to the number of job executions presented in this report version: **10**.

3.3.1 L2 Ethernet Interface Loop

Following sections include summary graphs of DPDK Testpmd Phy-to-Phy performance with L2 Ethernet Interface Loop, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss). Performance is reported for Testpmd running in multiple configurations of Testpmd pmd thread(s), a.k.a. Testpmd data plane thread(s), and their physical CPU core(s) placement.

¹⁶⁹ <https://jenkins.fd.io/view/csit/job/csit-dpdk-perf-1804-all>

¹⁷⁰ <https://jenkins.fd.io/view/csit/job/csit-dpdk-perf-1804-all>

NDR Throughput

Testpmd NDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below.




Figure 1. Testpmd 1thread 1core - NDR Throughput for Phy-to-Phy L2 Ethernet Looping.

CSIT source code for the test cases used for above plots can be found in CSIT git repository:

```
$ cd tests/dpdk/perf
$ grep -P '64B-1t1c-eth-l2xcbase-testpmd-ndrdisc' *
```

Testpmd NDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below.




Figure 2. Testpmd 2threads 2cores - NDR Throughput for Phy-to-Phy L2 Ethernet Looping.

CSIT source code for the test cases used for above plots can be found in CSIT git repository:

```
$ cd tests/dpdk/perf
$ grep -P '64B-2t2c-eth-l2xcbase-testpmd-ndrdisc' *
```

PDR Throughput

Testpmd PDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.




Figure 3. Testpmd 1thread 1core - PDR Throughput for Phy-to-Phy L2 Ethernet Looping.

CSIT source code for the test cases used for above plots can be found in CSIT git repository:

```
$ cd tests/dpdk/perf
$ grep -P '64B-1t1c-eth-l2xcbase-testpmd-pdrdisc' *
```

Testpmd PDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.




Figure 4. Testpmd 2thread 2core - PDR Throughput for Phy-to-Phy L2 Ethernet Looping.

CSIT source code for the test cases used for above plots can be found in CSIT git repository:

```
$ cd tests/dpdk/perf
$ grep -P '64B-2t2c-eth-l2xcbase-testpmd-pdrdisc' *
```

3.3.2 IPv4 Routed-Forwarding

Following sections include summary graphs of L3FWD Phy-to-Phy performance with packet routed forwarding, including NDR throughput (zero packet loss) and PDR throughput (<0.5% packet loss). Performance is reported for L3FWD running in multiple configurations of L3FWD pmd thread(s), a.k.a. L3FWD data plane thread(s), and their physical CPU core(s) placement.

NDR Throughput

Testpmd NDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below.




Figure 1. L3FWD 1thread 1core - NDR Throughput for Phy-to-Phy IPv4 Routed-Forwarding Looping.

CSIT source code for the test cases used for above plots can be found in CSIT git repository:

```
$ cd tests/dpdk/perf
$ grep -P '64B-1t1c-ethip4-ip4base-l3fwd-ndrdisc' *
```

Testpmd NDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below.




Figure 2. L3FWD 2threads 2cores - NDR Throughput for Phy-to-Phy IPv4 Routed-Forwarding Looping.
 CSIT source code for the test cases used for above plots can be found in CSIT git repository:

```
$ cd tests/dpdk/perf
$ grep -P '64B-2t2c-ethip4-ip4base-l3fwd-ndrdisc' *
```

PDR Throughput

L3FWD PDR 64B packet throughput in 1t1c setup (1thread, 1core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.




Figure 3. L3FWD 1thread 1core - PDR Throughput for Phy-to-Phy IPv4 Routed-Forwarding Looping.

CSIT source code for the test cases used for above plots can be found in CSIT git repository:

```
$ cd tests/dpdk/perf
$ grep -P '64B-1t1c-ethip4-ip4base-l3fwd-pdrdisc' *
```

L3FWD PDR 64B packet throughput in 2t2c setup (2thread, 2core) is presented in the graph below. PDR measured for 0.5% packet loss ratio.




Figure 4. L3FWD 2thread 2core - PDR Throughput for Phy-to-Phy IPv4 Routed-Forwarding Looping.

CSIT source code for the test cases used for above plots can be found in CSIT git repository:

```
$ cd tests/dpdk/perf
$ grep -P '64B-2t2c-ethip4-ip4base-l3fwd-pdrdisc' *
```

3.4 Packet Latency Graphs

Plotted results are generated from a single execution of CSIT NDR discovery test. Box plots are used to show the Minimum, Median and Maximum packet latency per test.

Title of each graph is a regex (regular expression) matching all throughput test cases plotted on this graph, *X-axis labels* are indices of individual test suites executed by [FD.io test executor dpdk performance jobs](https://jenkins.fd.io/view/csit/job/csit-dpdk-performance-jobs)¹⁷¹ that created result output file used as data source for the graph, *Y-axis labels* are measured packet Latency [uSec] values, and the *Graph legend* lists the plotted test suites and their indices. Latency is reported for concurrent symmetric bi-directional flows, separately for each direction: i) West-to-East: TGint1-to-SUT1-to-SUT2-to-TGint2, and ii) East-to-West: TGint2-to-SUT2-to-SUT1-to-TGint1.

¹⁷¹ <https://jenkins.fd.io/view/csit/job/csit-dpdk-perf-1804-all>

Note: Test results have been generated by [FD.io test executor dpdk performance jobs](#)¹⁷² with Robot Framework result files `csit-dpdk-perf-*.zip` [archived here](#).

3.4.1 L2 Ethernet Interface Loop

This section includes summary graphs of Testpmd Phy-to-Phy packet latency with L2 Ethernet Interface Loop measured at 50% of discovered NDR throughput rate. Latency is reported for Testpmd running in multiple configurations of Testpmd pmd thread(s), a.k.a. Testpmd data plane thread(s), and their physical CPU core(s) placement.

Testpmd packet latency - running in configuration of **one worker thread (1t) on one physical core (1c)** - is presented in the figure below.




Figure 1. Testpmd 1thread 1core - packet latency for Phy-to-Phy L2 Ethernet Looping.

CSIT test cases used to generate results presented above can be found in CSIT git repository by filtering with specified regex as follows:

¹⁷² <https://jenkins.fd.io/view/csit/job/csit-dpdk-perf-1804-all>

```
$ cd tests/dpdk/perf
$ grep -P '64B-1t1c-eth-l2xcbase-testpmd-ndrdisc' *
```

Testpmd packet latency - running in configuration of **two worker threads (2t) on two physical cores (2c)** - is presented in the figure below.




Figure 2. Testpmd 2thread 2core - packet latency for Phy-to-Phy L2 Ethernet Looping.

CSIT test cases used to generate results presented above can be found in CSIT git repository by filtering with specified regex as follows:

```
$ cd tests/dpdk/perf
$ grep -P '64B-2t2c-eth-l2xcbase-testpmd-ndrdisc' *
```

3.4.2 IPv4 Routed-Forwarding

This section includes summary graphs of L3FWD Phy-to-Phy performance with packet routed forwarding measured at 50% of discovered NDR throughput rate. Latency is reported for L3FWD running in multiple configurations of L3FWD pmd thread(s), a.k.a. L3FWD data plane thread(s), and their physical CPU core(s) placement.

L3FWD packet latency - running in configuration of **one worker thread (1t)** on **one physical core (1c)** - is presented in the figure below.




Figure 1. L3FWD 1thread 1core - packet latency for Phy-to-Phy IPv4 Routed-Forwarding.

CSIT test cases used to generate results presented above can be found in CSIT git repository by filtering with specified regex as follows:

```
$ cd tests/dpdk/perf
$ grep -E '64B-1t1c-ethip4-ip4base-l3fwd-ndrdisc' *
```

Testpmd packet latency - running in configuration of **two worker threads (2t)** on **two physical cores (2c)** - is presented in the figure below.




Figure 2. L3FWD 2thread 2core - packet latency for Phy-to-Phy IPv4 Routed-Forwarding.

CSIT test cases used to generate results presented above can be found in CSIT git repository by filtering with specified regex as follows:

```
$ cd tests/dpdk/perf
$ grep -E '64B-2t2c-ethip4-ip4base-l3fwd-ndrdisc' *
```

3.5 Test Environment

CSIT performance tests are executed on the three identical physical testbeds hosted by LF for FD.io project. Each testbed consists of two servers acting as Systems Under Test (SUT) and one server acting as Traffic Generator (TG).

3.5.1 Server Specification and Configuration

Complete specification and configuration of compute servers used in CSIT physical testbeds is maintained on wiki page [CSIT LF Testbeds](#)¹⁷³.

3.5.2 SUT Configuration

Host configuration

All hosts are Cisco UCS C240-M4 (2x Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz, 18c, 512GB RAM)

```
$ lscpu
Architecture:          x86_64
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Little Endian
CPU(s):                36
On-line CPU(s) list:  0-35
Thread(s) per core:    1
Core(s) per socket:    18
Socket(s):              2
NUMA node(s):          2
Vendor ID:              GenuineIntel
CPU family:             6
Model:                  63
Model name:             Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz
Stepping:               2
CPU MHz:                2294.249
BogoMIPS:               4589.82
Virtualization:        VT-x
L1d cache:              32K
L1i cache:              32K
L2 cache:               256K
L3 cache:               46080K
NUMA node0 CPU(s):     0-17
NUMA node1 CPU(s):     18-35
Flags:                  fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36_
↳ clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx pdpe1gb rdtscp lm constant_tsc arch_
↳ perfmon pebs bts rep_good nopl xtopology nonstop_tsc aperfmperf eagerfpu pni pclmulqdq dtes64_
↳ monitor ds_cpl vmx smx est tm2 ssse3 sdbg fma cx16 xtpr pdcm pcid dca sse4_1 sse4_2 x2apic movbe_
↳ popcnt tsc_deadline_timer aes xsave avx f16c rdrand lahf_lm abm epb tpr_shadow vnmi flexpriority_
↳ ept vpid fsgsbase tsc_adjust bmi1 avx2 smep bmi2 erms invpcid cqm xsaveopt cqm_llc cqm_occup_llc_
↳ dtherm arat pln pts
```

BIOS settings

```
C240 /bios # show advanced detail
Set-up parameters:
  Intel(R) VT-d ATS Support: Enabled
  Adjacent Cache Line Prefetcher: Enabled
  All Onboard LOM Ports: Enabled
  Altitude: 300 M
  Bits per second: 115200
  Power Technology: Performance
  Channel Interleaving: Auto
  Intel(R) VT-d Coherency Support: Disabled
  Console Redirection: COM 0
  Number of Enabled Cores: All
  Energy Performance: Performance
  CPU Performance: Enterprise
  DCU IP Prefetcher: Enabled
```

(continues on next page)

¹⁷³ https://wiki.fd.io/view/CSIT/CSIT_LF_testbed

```
DCU Streamer Prefetch: Enabled
Demand Scrub: Enabled
Direct Cache Access Support: Auto
Enhanced Intel Speedstep(R) Tec: Disabled
Execute Disable: Enabled
Flow Control: None
Hardware Prefetcher: Enabled
Intel(R) Hyper-Threading Techno: Disabled
Intel(R) Turbo Boost Technology: Disabled
Intel(R) VT: Enabled
Intel(R) VT-d: Enabled
Intel(R) Interrupt Remapping: Enabled
Legacy USB Support: Enabled
Extended APIC: XAPIC
LOM Port 1 OptionROM: Enabled
LOM Port 2 OptionROM: Enabled
MMIO above 4GB: Enabled
NUMA: Enabled
PCI ROM CLP: Disabled
Package C State Limit: C6 Retention
Intel(R) Pass Through DMA: Disabled
Patrol Scrub: Enabled
xHCI Mode: Disabled
All PCIe Slots OptionROM: Enabled
PCIe Slot:1 OptionROM: Disabled
PCIe Slot:2 OptionROM: Disabled
PCIe Slot:3 OptionROM: Disabled
PCIe Slot:4 OptionROM: Disabled
PCIe Slot:5 OptionROM: Disabled
PCIe Slot:6 OptionROM: Disabled
PCIe Slot:HBA Link Speed: GEN3
PCIe Slot:HBA OptionROM: Enabled
PCIe Slot:MLOM OptionROM: Enabled
PCIe Slot:N1 OptionROM: Enabled
PCIe Slot:N2 OptionROM: Enabled
Processor Power state C1 Enhanc: Disabled
Processor C3 Report: Disabled
Processor C6 Report: Disabled
P-STATE Coordination: HW ALL
Putty KeyPad: ESCN
Energy Performance Tuning: BIOS
QPI Link Frequency Select: Auto
QPI Snoop Mode: Home Snoop
Rank Interleaving: Auto
Redirection After BIOS POST: Always Enable
PCH SATA Mode: AHCI
Select Memory RAS: Maximum Performance
SR-IOV Support: Enabled
Terminal Type: VT100
Port 60/64 Emulation: Enabled
Workload Configuration: Balanced
CDN Support for VIC: Disabled
Out-of-Band Management: Disabled
```

NIC models and placement

In addition to CIMC and Management, each TG has 4x Intel X710 10GB NIC (=8 ports) and 2x Intel XL710 40GB NIC (=4 ports), whereas each SUT has:

- 1x Intel X520 NIC (10GB, 2 ports),
- 1x Cisco VIC 1385 (40GB, 2 ports),

- 1x Intel XL710 NIC (40GB, 2 ports),
- 1x Intel X710 NIC (10GB, 2 ports),
- 1x Cisco VIC 1227 (10GB, 2 ports).

This allows for a total of five ring topologies, each using ports on specific NIC model, enabling per NIC model benchmarking.

- 0a:00.0 Ethernet controller: Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01) Subsystem: Intel Corporation Ethernet Server Adapter X520-2
- 0a:00.1 Ethernet controller: Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01) Subsystem: Intel Corporation Ethernet Server Adapter X520-2
- 06:00.0 Ethernet controller: Cisco Systems Inc VIC Ethernet NIC (rev a2) Subsystem: Cisco Systems Inc VIC 1227 PCIe Ethernet NIC
- 07:00.0 Ethernet controller: Cisco Systems Inc VIC Ethernet NIC (rev a2) Subsystem: Cisco Systems Inc VIC 1227 PCIe Ethernet NIC
- 13:00.0 Ethernet controller: Cisco Systems Inc VIC Ethernet NIC (rev a2) Subsystem: Cisco Systems Inc VIC 1385 PCIe Ethernet NIC
- 15:00.0 Ethernet controller: Cisco Systems Inc VIC Ethernet NIC (rev a2) Subsystem: Cisco Systems Inc VIC 1385 PCIe Ethernet NIC
- 85:00.0 Ethernet controller: Intel Corporation Ethernet Controller XL710 for 40GbE QSFP+ (rev 01) Subsystem: Intel Corporation Ethernet Converged Network Adapter XL710-Q2
- 85:00.1 Ethernet controller: Intel Corporation Ethernet Controller XL710 for 40GbE QSFP+ (rev 01) Subsystem: Intel Corporation Ethernet Converged Network Adapter XL710-Q2
- 87:00.0 Ethernet controller: Intel Corporation Ethernet Controller X710 for 10GbE SFP+ (rev 01) Subsystem: Intel Corporation Ethernet Converged Network Adapter X710-2
- 87:00.1 Ethernet controller: Intel Corporation Ethernet Controller X710 for 10GbE SFP+ (rev 01) Subsystem: Intel Corporation Ethernet Converged Network Adapter X710-2

3.5.3 SUT Configuration - Host OS Linux

Software details (OS, configuration) of physical testbeds are maintained on wiki page [CSIT LF Testbeds](#)¹⁷⁴.

System provisioning is done by combination of PXE boot unattended install and [Ansible](#)¹⁷⁵ described in [CSIT Testbed Setup](#)¹⁷⁶.

Below a subset of the running configuration:

```
$ lsb_release -a
No LSB modules are available.
Distributor ID: Ubuntu
Description: Ubuntu 16.04.1 LTS
Release: 16.04
Codename: xenial
```

```
$ cat /sys/devices/system/node/node*/meminfo
Node 0 MemTotal: 264048168 kB
Node 0 MemFree: 257730716 kB
Node 0 MemUsed: 6317452 kB
Node 0 Active: 1079920 kB
Node 0 Inactive: 470064 kB
```

(continues on next page)

¹⁷⁴ https://wiki.fd.io/view/CSIT/CSIT_LF_testbed

¹⁷⁵ <https://www.ansible.com>

¹⁷⁶ <https://git.fd.io/csit/tree/resources/tools/testbed-setup/README.md?h=rls1804>

(continued from previous page)

```

Node 0 Active(anon):      674772 kB
Node 0 Inactive(anon):   248572 kB
Node 0 Active(file):     405148 kB
Node 0 Inactive(file):   221492 kB
Node 0 Unevictable:      0 kB
Node 0 Mlocked:          0 kB
Node 0 Dirty:            12 kB
Node 0 Writeback:        0 kB
Node 0 FilePages:        1270432 kB
Node 0 Mapped:            20116 kB
Node 0 AnonPages:        279548 kB
Node 0 Shmem:            643796 kB
Node 0 KernelStack:     3376 kB
Node 0 PageTables:       1316 kB
Node 0 NFS_Unstable:     0 kB
Node 0 Bounce:           0 kB
Node 0 WritebackTmp:     0 kB
Node 0 Slab:              80428 kB
Node 0 SReclaimable:     38288 kB
Node 0 SUnreclaim:      42140 kB
Node 0 AnonHugePages:    270336 kB
Node 0 HugePages_Total:  2048
Node 0 HugePages_Free:   2048
Node 0 HugePages_Surp:   0
Node 1 MemTotal:         264237596 kB
Node 1 MemFree:          256758976 kB
Node 1 MemUsed:          7478620 kB
Node 1 Active:           1746052 kB
Node 1 Inactive:         981104 kB
Node 1 Active(anon):    1272936 kB
Node 1 Inactive(anon):  849968 kB
Node 1 Active(file):    473116 kB
Node 1 Inactive(file):  131136 kB
Node 1 Unevictable:     0 kB
Node 1 Mlocked:         0 kB
Node 1 Dirty:           0 kB
Node 1 Writeback:       0 kB
Node 1 FilePages:       2715284 kB
Node 1 Mapped:           75928 kB
Node 1 AnonPages:       11920 kB
Node 1 Shmem:           2111036 kB
Node 1 KernelStack:    2576 kB
Node 1 PageTables:     1348 kB
Node 1 NFS_Unstable:   0 kB
Node 1 Bounce:         0 kB
Node 1 WritebackTmp:   0 kB
Node 1 Slab:            90604 kB
Node 1 SReclaimable:   55384 kB
Node 1 SUnreclaim:    35220 kB
Node 1 AnonHugePages:  6144 kB
Node 1 HugePages_Total: 2048
Node 1 HugePages_Free: 2048
Node 1 HugePages_Surp: 0

```

Kernel boot parameters used in CSIT performance testbeds

- **isolcpus=<cpu number>-<cpu number>** used for all cpu cores apart from first core of each socket used for running VPP worker threads and Qemu/LXC processes <https://www.kernel.org/doc/Documentation/admin-guide/kernel-parameters.txt>
- **intel_pstate=disable** - [X86] Do not enable intel_pstate as the default scaling driver for the supported processors. Intel P-State driver decide what P-state (CPU core power state) to use based

on requesting policy from the cpufreq core. [X86 - Either 32-bit or 64-bit x86] <https://www.kernel.org/doc/Documentation/cpu-freq/intel-pstate.txt>

- **nohz_full=<cpu number>-<cpu number>** - [KNL,BOOT] In kernels built with CONFIG_NO_HZ_FULL=y, set the specified list of CPUs whose tick will be stopped whenever possible. The boot CPU will be forced outside the range to maintain the timekeeping. The CPUs in this range must also be included in the rcu_nocbs= set. Specifies the adaptive-ticks CPU cores, causing kernel to avoid sending scheduling-clock interrupts to listed cores as long as they have a single runnable task. [KNL - Is a kernel start-up parameter, SMP - The kernel is an SMP kernel]. https://www.kernel.org/doc/Documentation/timers/NO_HZ.txt
- **rcu_nocbs** - [KNL] In kernels built with CONFIG_RCU_NOCB_CPU=y, set the specified list of CPUs to be no-callback CPUs, that never queue RCU callbacks (read-copy update). <https://www.kernel.org/doc/Documentation/admin-guide/kernel-parameters.txt>

Applied command line boot parameters:

```
$ cat /proc/cmdline
BOOT_IMAGE=/vmlinuz-4.4.0-72-generic root=UUID=35ea11e4-e44f-4f67-8cbe-12f09c49ed90 ro isolcpus=1-
↪17,19-35 nohz_full=1-17,19-35 rcu_nocbs=1-17,19-35 intel_pstate=disable console=tty0_
↪console=ttyS0,115200n8
```

Mount listing

```
$ cat /proc/mounts
sysfs /sys sysfs rw,nosuid,nodev,noexec,relatime 0 0
proc /proc proc rw,nosuid,nodev,noexec,relatime 0 0
udev /dev devtmpfs rw,nosuid,relatime,size=264125468k,nr_inodes=66031367,mode=755 0 0
devpts /dev/pts devpts rw,nosuid,noexec,relatime,gid=5,mode=620,ptmxmode=000 0 0
tmpfs /run tmpfs rw,nosuid,noexec,relatime,size=52828580k,mode=755 0 0
/dev/sda2 / ext4 rw,relatime,errors=remount-ro,data=ordered 0 0
securityfs /sys/kernel/security securityfs rw,nosuid,nodev,noexec,relatime 0 0
tmpfs /dev/shm tmpfs rw,nosuid,nodev 0 0
tmpfs /run/lock tmpfs rw,nosuid,nodev,noexec,relatime,size=5120k 0 0
tmpfs /sys/fs/cgroup tmpfs ro,nosuid,nodev,noexec,mode=755 0 0
cgroup /sys/fs/cgroup/systemd cgroup rw,nosuid,nodev,noexec,relatime,xattr,release_agent=/lib/
↪systemd/systemd-cgroups-agent,name=systemd 0 0
pstore /sys/fs/pstore pstore rw,nosuid,nodev,noexec,relatime 0 0
cgroup /sys/fs/cgroup/freezer cgroup rw,nosuid,nodev,noexec,relatime,freezer 0 0
cgroup /sys/fs/cgroup/net_cls,net_prio cgroup rw,nosuid,nodev,noexec,relatime,net_cls,net_prio 0 0
cgroup /sys/fs/cgroup/cpu,cpuacct cgroup rw,nosuid,nodev,noexec,relatime,cpu,cpuacct 0 0
cgroup /sys/fs/cgroup/memory cgroup rw,nosuid,nodev,noexec,relatime,memory 0 0
cgroup /sys/fs/cgroup/blkio cgroup rw,nosuid,nodev,noexec,relatime,blkio 0 0
cgroup /sys/fs/cgroup/perf_event cgroup rw,nosuid,nodev,noexec,relatime,perf_event 0 0
cgroup /sys/fs/cgroup/devices cgroup rw,nosuid,nodev,noexec,relatime,devices 0 0
cgroup /sys/fs/cgroup/cpuset cgroup rw,nosuid,nodev,noexec,relatime,cpuset,clone_children 0 0
cgroup /sys/fs/cgroup/hugetlb cgroup rw,nosuid,nodev,noexec,relatime,hugetlb 0 0
cgroup /sys/fs/cgroup/pids cgroup rw,nosuid,nodev,noexec,relatime,pids 0 0
systemd-1 /proc/sys/fs/binfmt_misc autofs rw,relatime,fd=26,pgrp=1,timeout=0,minproto=5,maxproto=5,
↪direct 0 0
hugetlbfs /dev/hugepages hugetlbfs rw,relatime 0 0
debugfs /sys/kernel/debug debugfs rw,relatime 0 0
mqueue /dev/mqueue mqueue rw,relatime 0 0
tracefs /sys/kernel/debug/tracing tracefs rw,relatime 0 0
fusectl /sys/fs/fuse/connections fusectl rw,relatime 0 0
/dev/sda1 /boot ext4 rw,relatime,data=ordered 0 0
none /mnt/huge hugetlbfs rw,relatime,pagesize=2048k 0 0
lxcfs /var/lib/lxcfs fuse.lxcfs rw,nosuid,nodev,relatime,user_id=0,group_id=0,allow_other 0 0
```

Package listing

```
$ dpkg -l | grep '^ii' | awk '{print $2 ": " $3}'
accountsservice: 0.6.40-2ubuntu11.1
```

(continues on next page)

(continued from previous page)

```
ethtool: 1:4.5-1
expect: 5.45-7
fakeroot: 1.20.2-1ubuntu1
file: 1:5.25-2ubuntu1
findutils: 4.6.0+git+20160126-2
fontconfig: 2.11.94-0ubuntu1.1
fontconfig-config: 2.11.94-0ubuntu1.1
fonts-dejavu-core: 2.35-1
fonts-dejavu-extra: 2.35-1
g++: 4:5.3.1-1ubuntu1
g++-5: 5.4.0-6ubuntu1~16.04.2
gcc: 4:5.3.1-1ubuntu1
gcc-5: 5.4.0-6ubuntu1~16.04.2
gcc-5-base:amd64: 5.4.0-6ubuntu1~16.04.2
gcc-6-base:amd64: 6.0.1-0ubuntu1
gdb: 7.11.1-0ubuntu1~16.5
gdbserver: 7.11.1-0ubuntu1~16.5
genisoimage: 9:1.1.11-3ubuntu1
gettext-base: 0.19.7-2ubuntu3
gir1.2-glib-2.0:amd64: 1.46.0-3ubuntu1
git: 1:2.7.4-0ubuntu1
git-man: 1:2.7.4-0ubuntu1
gnupg: 1.4.20-1ubuntu3.1
gpgv: 1.4.20-1ubuntu3.1
grep: 2.25-1~16.04.1
grub-common: 2.02~beta2-36ubuntu3.1
grub-gfxpayload-lists: 0.7
grub-pc: 2.02~beta2-36ubuntu3.1
grub-pc-bin: 2.02~beta2-36ubuntu3.1
grub2-common: 2.02~beta2-36ubuntu3.1
gzip: 1.6-4ubuntu1
hicolor-icon-theme: 0.15-0ubuntu1
hostname: 3.16ubuntu2
ifupdown: 0.8.10ubuntu1
init: 1.29ubuntu2
init-system-helpers: 1.29ubuntu2
initramfs-tools: 0.122ubuntu8.1
initramfs-tools-bin: 0.122ubuntu8.1
initramfs-tools-core: 0.122ubuntu8.1
initscripts: 2.88dsf-59.3ubuntu2
insserv: 1.14.0-5ubuntu3
installation-report: 2.60ubuntu1
iproute2: 4.3.0-1ubuntu3
iptables: 1.6.0-2ubuntu3
iputils-ping: 3:20121221-5ubuntu2
ipxe-qemu: 1.0.0+git-20150424.a25a16d-1ubuntu1
isc-dhcp-client: 4.3.3-5ubuntu12.1
isc-dhcp-common: 4.3.3-5ubuntu12.1
iso-codes: 3.65-1
iw: 3.17-1
java-common: 0.56ubuntu2
kbd: 1.15.5-1ubuntu4
keyboard-configuration: 1.108ubuntu15.2
klibc-utils: 2.0.4-8ubuntu1.16.04.1
kmod: 22-1ubuntu4
krb5-locales: 1.13.2+dfsg-5
kubeadm: 1.8.4-00
kubect1: 1.8.4-00
kubelet: 1.8.4-00
kubernetes-cni: 0.5.1-00
language-selector-common: 0.165.3
```

(continues on next page)

(continued from previous page)

```
laptop-detect: 0.13.7ubuntu2
less: 481-2.1
libaccountsservice0:amd64: 0.6.40-2ubuntu11.1
libacl1:amd64: 2.2.52-3
libaio1:amd64: 0.3.110-2
libalgorithm-diff-perl: 1.19.03-1
libalgorithm-diff-xs-perl: 0.04-4build1
libalgorithm-merge-perl: 0.08-3
libapparmor-perl: 2.10.95-0ubuntu2.6
libapparmor1:amd64: 2.10.95-0ubuntu2
libapr1:amd64: 1.5.2-3
libapt-inst2.0:amd64: 1.2.12~ubuntu16.04.1
libapt-pkg5.0:amd64: 1.2.12~ubuntu16.04.1
libasan2:amd64: 5.4.0-6ubuntu1~16.04.2
libasan1-8-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libasound2:amd64: 1.1.0-0ubuntu1
libasound2-data: 1.1.0-0ubuntu1
libasprintf0v5:amd64: 0.19.7-2ubuntu3
libasyncns0:amd64: 0.8-5build1
libatk1.0-0:amd64: 2.18.0-1
libatk1.0-data: 2.18.0-1
libatm1:amd64: 1:2.5.1-1.5
libatomic1:amd64: 5.4.0-6ubuntu1~16.04.2
libattr1:amd64: 1:2.4.47-2
libaudit-common: 1:2.4.5-1ubuntu2
libaudit1:amd64: 1:2.4.5-1ubuntu2
libavahi-client3:amd64: 0.6.32~rc+dfsg-1ubuntu2
libavahi-common-data:amd64: 0.6.32~rc+dfsg-1ubuntu2
libavahi-common3:amd64: 0.6.32~rc+dfsg-1ubuntu2
libbabeltrace-ctf1:amd64: 1.3.2-1
libbabeltrace1:amd64: 1.3.2-1
libblkid1:amd64: 2.27.1-6ubuntu3.1
libbluetooth3:amd64: 5.37-0ubuntu5
libboost-iostreams1.58.0:amd64: 1.58.0+dfsg-5ubuntu3.1
libboost-random1.58.0:amd64: 1.58.0+dfsg-5ubuntu3.1
libboost-system1.58.0:amd64: 1.58.0+dfsg-5ubuntu3.1
libboost-thread1.58.0:amd64: 1.58.0+dfsg-5ubuntu3.1
libbrlapi0.6:amd64: 5.3.1-2ubuntu2.1
libbsd0:amd64: 0.8.2-1
libbz2-1.0:amd64: 1.0.6-8
libc-bin: 2.23-0ubuntu3
libc-dev-bin: 2.23-0ubuntu9
libc6-dbg:amd64: 2.23-0ubuntu9
libc6-dev:amd64: 2.23-0ubuntu9
libcaca0:amd64: 0.99.beta19-2build2~gcc5.2
libcacard0:amd64: 1:2.5.0-2
libcairo2:amd64: 1.14.6-1
libcap-ng0:amd64: 0.7.7-1
libcap2:amd64: 1:2.24-12
libcap2-bin: 1:2.24-12
libcc1-0:amd64: 5.4.0-6ubuntu1~16.04.2
libcgrouper1:amd64: 0.41-7ubuntu1
libcilkrts5:amd64: 5.4.0-6ubuntu1~16.04.2
libcomerr2:amd64: 1.42.13-1ubuntu1
libcpufreq0: 008-1
libcryptsetup4:amd64: 2:1.6.6-5ubuntu2
libcups2:amd64: 2.1.3-4
libcurl3-gnutls:amd64: 7.47.0-1ubuntu2.1
libdatrie1:amd64: 0.2.10-2
libdb5.3:amd64: 5.3.28-11
libdbus-1-3:amd64: 1.10.6-1ubuntu3
```

(continues on next page)

(continued from previous page)

```

libdbus-glib-1-2:amd64: 0.106-1
libdebconfclient0:amd64: 0.198ubuntu1
libdevmapper1.02.1:amd64: 2:1.02.110-1ubuntu10
libdns-export162: 1:9.10.3.dfsg.P4-8ubuntu1.1
libdpkg-perl: 1.18.4ubuntu1.1
libdrm-amdgpu1:amd64: 2.4.67-1ubuntu0.16.04.2
libdrm-intel1:amd64: 2.4.67-1ubuntu0.16.04.2
libdrm-nouveau2:amd64: 2.4.67-1ubuntu0.16.04.2
libdrm-radeon1:amd64: 2.4.67-1ubuntu0.16.04.2
libdrm2:amd64: 2.4.67-1ubuntu0.16.04.2
libedit2:amd64: 3.1-20150325-1ubuntu2
libelf1:amd64: 0.165-3ubuntu1
liberror-perl: 0.17-1.2
libestr0: 0.1.10-1
libexpat1:amd64: 2.1.0-7ubuntu0.16.04.2
libexpat1-dev:amd64: 2.1.0-7ubuntu0.16.04.2
libfakeroot:amd64: 1.20.2-1ubuntu1
libfdisk1:amd64: 2.27.1-6ubuntu3.1
libfdt1:amd64: 1.4.0+dfsg-2
libffi6:amd64: 3.2.1-4
libfile-fcntllock-perl: 0.22-3
libflac8:amd64: 1.3.1-4
libfontconfig1:amd64: 2.11.94-0ubuntu1.1
libfontenc1:amd64: 1:1.1.3-1
libfreetype6:amd64: 2.6.1-0.1ubuntu2
libfribidi0:amd64: 0.19.7-1
libfuse2:amd64: 2.9.4-1ubuntu3
libgcc-5-dev:amd64: 5.4.0-6ubuntu1~16.04.2
libgcc1:amd64: 1:6.0.1-0ubuntu1
libgcrypt20:amd64: 1.6.5-2ubuntu0.2
libgdbm3:amd64: 1.8.3-13.1
libgdk-pixbuf2.0-0:amd64: 2.32.2-1ubuntu1.2
libgdk-pixbuf2.0-common: 2.32.2-1ubuntu1.2
libgif7:amd64: 5.1.4-0.3~16.04
libgirepository-1.0-1:amd64: 1.46.0-3ubuntu1
libgl1-mesa-dri:amd64: 11.2.0-1ubuntu2.2
libgl1-mesa-glx:amd64: 11.2.0-1ubuntu2.2
libglapi-mesa:amd64: 11.2.0-1ubuntu2.2
libglib2.0-0:amd64: 2.48.1-1~ubuntu16.04.1
libglib2.0-bin: 2.48.1-1~ubuntu16.04.1
libglib2.0-data: 2.48.1-1~ubuntu16.04.1
libglib2.0-dev: 2.48.1-1~ubuntu16.04.1
libgmp10:amd64: 2:6.1.0+dfsg-2
libgnutls-openssl27:amd64: 3.4.10-4ubuntu1.1
libgnutls30:amd64: 3.4.10-4ubuntu1.1
libgomp1:amd64: 5.4.0-6ubuntu1~16.04.2
libgpg-error0:amd64: 1.21-2ubuntu1
libgraphite2-3:amd64: 1.3.6-1ubuntu1
libgssapi-krb5-2:amd64: 1.13.2+dfsg-5
libgssapi3-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libgtk2.0-0:amd64: 2.24.30-1ubuntu1.16.04.2
libgtk2.0-bin: 2.24.30-1ubuntu1.16.04.2
libgtk2.0-common: 2.24.30-1ubuntu1.16.04.2
libharfbuzz0b:amd64: 1.0.1-1ubuntu0.1
libhcrypto4-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libheimbase1-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libheimntlm0-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libhogweed4:amd64: 3.2-1
libhx509-5-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libice-dev:amd64: 2:1.0.9-1
libice6:amd64: 2:1.0.9-1

```

(continues on next page)

(continued from previous page)

```
libicu55:amd64: 55.1-7
libidn11:amd64: 1.32-3ubuntu1.1
libisc-export160: 1:9.10.3.dfsg.P4-8ubuntu1.1
libiscsi2:amd64: 1.12.0-2
libisl15:amd64: 0.16.1-1
libitm1:amd64: 5.4.0-6ubuntu1~16.04.2
libjbig0:amd64: 2.1-3.1
libjpeg-turbo8:amd64: 1.4.2-0ubuntu3
libjpeg8:amd64: 8c-2ubuntu8
libjson-c2:amd64: 0.11-4ubuntu2
libk5crypto3:amd64: 1.13.2+dfsg-5
libkeyutils1:amd64: 1.5.9-8ubuntu1
libklibc: 2.0.4-8ubuntu1.16.04.1
libkmod2:amd64: 22-1ubuntu4
libkrb5-26-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libkrb5-3:amd64: 1.13.2+dfsg-5
libkrb5support0:amd64: 1.13.2+dfsg-5
liblcms2-2:amd64: 2.6-3ubuntu2
libldap-2.4-2:amd64: 2.4.42+dfsg-2ubuntu3.1
libllvm3.8:amd64: 1:3.8-2ubuntu4
liblocale-gettext-perl: 1.07-1build1
liblsan0:amd64: 5.4.0-6ubuntu1~16.04.2
libltdl-dev:amd64: 2.4.6-0.1
libltdl7:amd64: 2.4.6-0.1
liblxc1: 2.0.7-0ubuntu1~16.04.2
liblz4-1:amd64: 0.0~r131-2ubuntu2
liblzma5:amd64: 5.1.1alpha+20120614-2ubuntu2
libmagic1:amd64: 1:5.25-2ubuntu1
libmnl0:amd64: 1.0.3-5
libmount1:amd64: 2.27.1-6ubuntu3.1
libmpc3:amd64: 1.0.3-1
libmpdec2:amd64: 2.4.2-1
libmpfr4:amd64: 3.1.4-1
libmpx0:amd64: 5.4.0-6ubuntu1~16.04.2
libncurses5:amd64: 6.0+20160213-1ubuntu1
libncursesw5:amd64: 6.0+20160213-1ubuntu1
libnetfilter-contrack3:amd64: 1.0.5-1
libnettle6:amd64: 3.2-1
libnewt0.52:amd64: 0.52.18-1ubuntu2
libnfnfnetlink0:amd64: 1.0.1-3
libnih-dbus1:amd64: 1.0.3-4.3ubuntu1
libnih1:amd64: 1.0.3-4.3ubuntu1
libnl-3-200:amd64: 3.2.27-1
libnl-genl-3-200:amd64: 3.2.27-1
libnspr4:amd64: 2:4.12-0ubuntu0.16.04.1
libnss3:amd64: 2:3.23-0ubuntu0.16.04.1
libnss3-nssdb: 2:3.23-0ubuntu0.16.04.1
libnuma-dev:amd64: 2.0.11-1ubuntu1
libnuma1:amd64: 2.0.11-1ubuntu1
libogg0:amd64: 1.3.2-1
libopus0:amd64: 1.1.2-1ubuntu1
libp11-kit0:amd64: 0.23.2-3
libpam-cgfs: 2.0.6-0ubuntu1~16.04.1
libpam-modules:amd64: 1.1.8-3.2ubuntu2
libpam-modules-bin: 1.1.8-3.2ubuntu2
libpam-runtime: 1.1.8-3.2ubuntu2
libpam0g:amd64: 1.1.8-3.2ubuntu2
libpango-1.0-0:amd64: 1.38.1-1
libpangocairo-1.0-0:amd64: 1.38.1-1
libpangoft2-1.0-0:amd64: 1.38.1-1
libpcap-dev: 1.7.4-2
```

(continues on next page)

(continued from previous page)

```

libpcap0.8:amd64: 1.7.4-2
libpcap0.8-dev: 1.7.4-2
libpci3:amd64: 1:3.3.1-1.1ubuntu1
libpciaccess0:amd64: 0.13.4-1
libpcre16-3:amd64: 2:8.38-3.1
libpcre3:amd64: 2:8.38-3.1
libpcre3-dev:amd64: 2:8.38-3.1
libpcre32-3:amd64: 2:8.38-3.1
libpcrecpp0v5:amd64: 2:8.38-3.1
libpcsc-lite1:amd64: 1.8.14-1ubuntu1.16.04.1
libperl5.22:amd64: 5.22.1-9
libpixman-1-0:amd64: 0.33.6-1
libplymouth4:amd64: 0.9.2-3ubuntu13.1
libpng12-0:amd64: 1.2.54-1ubuntu1
libpolkit-gobject-1-0:amd64: 0.105-14.1
libpopt0:amd64: 1.16-10
libprocps4:amd64: 2:3.3.10-4ubuntu2
libpthread-stubs0-dev:amd64: 0.3-4
libpulse0:amd64: 1:8.0-0ubuntu3
libpython-all-dev:amd64: 2.7.11-1
libpython-dev:amd64: 2.7.11-1
libpython-stdlib:amd64: 2.7.11-1
libpython2.7:amd64: 2.7.12-1~16.04
libpython2.7-dev:amd64: 2.7.12-1~16.04
libpython2.7-minimal:amd64: 2.7.12-1~16.04
libpython2.7-stdlib:amd64: 2.7.12-1~16.04
libpython3-stdlib:amd64: 3.5.1-3
libpython3.5:amd64: 3.5.2-2ubuntu0~16.04.1
libpython3.5-minimal:amd64: 3.5.2-2ubuntu0~16.04.1
libpython3.5-stdlib:amd64: 3.5.2-2ubuntu0~16.04.1
libquadmath0:amd64: 5.4.0-6ubuntu1~16.04.2
librados2: 10.2.2-0ubuntu0.16.04.2
librbd1: 10.2.2-0ubuntu0.16.04.2
libreadline6:amd64: 6.3-8ubuntu2
libroken18-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
librtmp1:amd64: 2.4+20151223.gitfa8646d-1build1
libsasl2-2:amd64: 2.1.26.dfsg1-14build1
libsasl2-modules:amd64: 2.1.26.dfsg1-14build1
libsasl2-modules-db:amd64: 2.1.26.dfsg1-14build1
libSDL1.2debian:amd64: 1.2.15+dfsg1-3
libseccomp2:amd64: 2.2.3-3ubuntu3
libselinux1:amd64: 2.4-3build2
libsemanage-common: 2.3-1build3
libsemanage1:amd64: 2.3-1build3
libsepol1:amd64: 2.4-2
libsigsegv2:amd64: 2.10-4
libslang2:amd64: 2.3.0-2ubuntu1
libsm-dev:amd64: 2:1.2.2-1
libsm6:amd64: 2:1.2.2-1
libsmartcols1:amd64: 2.27.1-6ubuntu3.1
libsndfile1:amd64: 1.0.25-10
libspice-server1:amd64: 0.12.6-4ubuntu0.1
libsqlite3-0:amd64: 3.11.0-1ubuntu1
libss2:amd64: 1.42.13-1ubuntu1
libssl1.0.0:amd64: 1.0.2g-1ubuntu4.5
libstdc++-5-dev:amd64: 5.4.0-6ubuntu1~16.04.2
libstdc++6:amd64: 5.4.0-6ubuntu1~16.04.2
libsystemd0:amd64: 229-4ubuntu10
libtasn1-6:amd64: 4.7-3ubuntu0.16.04.1
libtcl8.6:amd64: 8.6.5+dfsg-2
libtext-charwidth-perl: 0.04-7build5

```

(continues on next page)

(continued from previous page)

```
libtext-iconv-perl: 1.7-5build4
libtext-wrapi18n-perl: 0.06-7.1
libthai-data: 0.1.24-2
libthai0:amd64: 0.1.24-2
libtiff5:amd64: 4.0.6-1ubuntu0.2
libtinfo5:amd64: 6.0+20160213-1ubuntu1
libtk8.6:amd64: 8.6.5-1
libtool: 2.4.6-0.1
libtsan0:amd64: 5.4.0-6ubuntu1~16.04.2
libtxc-dxtn-s2tc0:amd64: 0~git20131104-1.1
libubsan0:amd64: 5.4.0-6ubuntu1~16.04.2
libudev1:amd64: 229-4ubuntu10
libusb-0.1-4:amd64: 2:0.1.12-28
libusb-1.0-0:amd64: 2:1.0.20-1
libusbredirparser1:amd64: 0.7.1-1
libustr-1.0-1:amd64: 1.0.4-5
libutempter0:amd64: 1.1.6-3
libuuid1:amd64: 2.27.1-6ubuntu3.1
libvorbis0a:amd64: 1.3.5-3
libvorbisenc2:amd64: 1.3.5-3
libwind0-heimdal:amd64: 1.7~git20150920+dfsg-4ubuntu1
libwrap0:amd64: 7.6.q-25
libx11-6:amd64: 2:1.6.3-1ubuntu2
libx11-data: 2:1.6.3-1ubuntu2
libx11-dev:amd64: 2:1.6.3-1ubuntu2
libx11-doc: 2:1.6.3-1ubuntu2
libx11-xcb1:amd64: 2:1.6.3-1ubuntu2
libxau-dev:amd64: 1:1.0.8-1
libxau6:amd64: 1:1.0.8-1
libxaw7:amd64: 2:1.0.13-1
libxcb-dri2-0:amd64: 1.11.1-1ubuntu1
libxcb-dri3-0:amd64: 1.11.1-1ubuntu1
libxcb-glx0:amd64: 1.11.1-1ubuntu1
libxcb-present0:amd64: 1.11.1-1ubuntu1
libxcb-render0:amd64: 1.11.1-1ubuntu1
libxcb-shape0:amd64: 1.11.1-1ubuntu1
libxcb-shm0:amd64: 1.11.1-1ubuntu1
libxcb-sync1:amd64: 1.11.1-1ubuntu1
libxcb1:amd64: 1.11.1-1ubuntu1
libxcb1-dev:amd64: 1.11.1-1ubuntu1
libxcomposite1:amd64: 1:0.4.4-1
libxcursor1:amd64: 1:1.1.14-1
libxdamage1:amd64: 1:1.1.4-2
libxdmcp-dev:amd64: 1:1.1.2-1.1
libxdmcp6:amd64: 1:1.1.2-1.1
libxen-4.6:amd64: 4.6.0-1ubuntu4.1
libxenstore3.0:amd64: 4.6.0-1ubuntu4.1
libxext6:amd64: 2:1.3.3-1
libxfixes3:amd64: 1:5.0.1-2
libxft2:amd64: 2.3.2-1
libxi6:amd64: 2:1.7.6-1
libxinerama1:amd64: 2:1.1.3-1
libxml2:amd64: 2.9.3+dfsg1-1ubuntu0.1
libxmu6:amd64: 2:1.1.2-2
libxmuu1:amd64: 2:1.1.2-2
libxpm4:amd64: 1:3.5.11-1
libxrandr2:amd64: 2:1.5.0-1
libxrender1:amd64: 1:0.9.9-0ubuntu1
libxshmfence1:amd64: 1.2-1
libxss1:amd64: 1:1.2.2-1
libxt-dev:amd64: 1:1.1.5-0ubuntu1
```

(continues on next page)

(continued from previous page)

```
libxt6:amd64: 1:1.1.5-0ubuntu1
libxtables11:amd64: 1.6.0-2ubuntu3
libxtst6:amd64: 2:1.2.2-1
libxv1:amd64: 2:1.0.10-1
libxxf86dga1:amd64: 2:1.1.4-1
libxxf86vm1:amd64: 1:1.1.4-1
libyajl2:amd64: 2.1.0-2
linux-base: 4.0ubuntu1
linux-firmware: 1.157.2
linux-generic: 4.4.0.72.78
linux-headers-4.4.0-72: 4.4.0-72.93
linux-headers-4.4.0-72-generic: 4.4.0-72.93
linux-headers-generic: 4.4.0.72.78
linux-image-4.4.0-72-generic: 4.4.0-72.93
linux-image-extra-4.4.0-72-generic: 4.4.0-72.93
linux-image-generic: 4.4.0.72.78
linux-libc-dev:amd64: 4.4.0-72.93
locales: 2.23-0ubuntu3
login: 1:4.2-3.1ubuntu5
logrotate: 3.8.7-2ubuntu2
lsb-base: 9.20160110ubuntu0.2
lsb-release: 9.20160110ubuntu0.2
lxc: 2.0.7-0ubuntu1~16.04.2
lxc-common: 2.0.7-0ubuntu1~16.04.2
lxc-templates: 2.0.7-0ubuntu1~16.04.2
lxc1: 2.0.7-0ubuntu1~16.04.2
lxcfs: 2.0.6-0ubuntu1~16.04.1
m4: 1.4.17-5
make: 4.1-6
makedev: 2.3.1-93ubuntu1
manpages: 4.04-2
manpages-dev: 4.04-2
mawk: 1.3.3-17ubuntu2
mime-support: 3.59ubuntu1
mount: 2.27.1-6ubuntu3.1
mountall: 2.54ubuntu1
msr-tools: 1.3-2
multiarch-support: 2.23-0ubuntu3
ncurses-base: 6.0+20160213-1ubuntu1
ncurses-bin: 6.0+20160213-1ubuntu1
ncurses-term: 6.0+20160213-1ubuntu1
net-tools: 1.60-26ubuntu1
netbase: 5.3
netcat-openbsd: 1.105-7ubuntu1
openjdk-8-jdk:amd64: 8u131-b11-2ubuntu1.16.04.3
openjdk-8-jdk-headless:amd64: 8u131-b11-2ubuntu1.16.04.3
openjdk-8-jre:amd64: 8u131-b11-2ubuntu1.16.04.3
openjdk-8-jre-headless:amd64: 8u131-b11-2ubuntu1.16.04.3
openssh-client: 1:7.2p2-4ubuntu2.1
openssh-server: 1:7.2p2-4ubuntu2.1
openssh-sftp-server: 1:7.2p2-4ubuntu2.1
openssl: 1.0.2g-1ubuntu4.5
os-prober: 1.70ubuntu3
passwd: 1:4.2-3.1ubuntu5
patch: 2.7.5-1
pciutils: 1:3.3.1-1.1ubuntu1
perl: 5.22.1-9
perl-base: 5.22.1-9
perl-modules-5.22: 5.22.1-9
pkg-config: 0.29.1-0ubuntu1
plymouth: 0.9.2-3ubuntu13.1
```

(continues on next page)

(continued from previous page)

```
plymouth-theme-ubuntu-text: 0.9.2-3ubuntu13.1
procps: 2:3.3.10-4ubuntu2
python: 2.7.11-1
python-all: 2.7.11-1
python-all-dev: 2.7.11-1
python-apt: 1.1.0~beta1build1
python-apt-common: 1.1.0~beta1build1
python-dev: 2.7.11-1
python-iniparse: 0.4-2.2
python-minimal: 2.7.11-1
python-pip: 8.1.1-2ubuntu0.2
python-pip-whl: 8.1.1-2ubuntu0.2
python-pkg-resources: 20.7.0-1
python-setuptools: 20.7.0-1
python-six: 1.10.0-3
python-virtualenv: 15.0.1+ds-3
python-wheel: 0.29.0-1
python2.7: 2.7.12-1~16.04
python2.7-dev: 2.7.12-1~16.04
python2.7-minimal: 2.7.12-1~16.04
python3: 3.5.1-3
python3-apt: 1.1.0~beta1build1
python3-chardet: 2.3.0-2
python3-dbus: 1.2.0-3
python3-gi: 3.20.0-0ubuntu1
python3-lxc: 2.0.7-0ubuntu1~16.04.2
python3-minimal: 3.5.1-3
python3-pkg-resources: 20.7.0-1
python3-requests: 2.9.1-3
python3-six: 1.10.0-3
python3-urllib3: 1.13.1-2ubuntu0.16.04.1
python3-virtualenv: 15.0.1+ds-3
python3.5: 3.5.2-2ubuntu0~16.04.1
python3.5-minimal: 3.5.2-2ubuntu0~16.04.1
qemu-block-extra:amd64: 1:2.5+dfsg-5ubuntu10.5
qemu-system-common: 1:2.5+dfsg-5ubuntu10.5
qemu-system-x86: 1:2.5+dfsg-5ubuntu10.5
qemu-utils: 1:2.5+dfsg-5ubuntu10.5
readline-common: 6.3-8ubuntu2
rename: 0.20-4
resolvconf: 1.78ubuntu2
rsync: 3.1.1-3ubuntu1
rsyslog: 8.16.0-1ubuntu3
screen: 4.3.1-2build1
seabios: 1.8.2-1ubuntu1
sed: 4.2.2-7
sensible-utils: 0.0.9
sgml-base: 1.26+nmu4ubuntu1
shared-mime-info: 1.5-2ubuntu0.1
sharutils: 1:4.15.2-1
socat: 1.7.3.1-1
ssh-import-id: 5.5-0ubuntu1
sudo: 1.8.16-0ubuntu1.1
systemd: 229-4ubuntu10
systemd-sysv: 229-4ubuntu10
sysv-rc: 2.88dsf-59.3ubuntu2
sysvinit-utils: 2.88dsf-59.3ubuntu2
tar: 1.28-2.1
tasksel: 3.34ubuntu3
tasksel-data: 3.34ubuntu3
tcl-expect:amd64: 5.45-7
```

(continues on next page)

(continued from previous page)

```

tc18.6: 8.6.5+dfsg-2
tcpd: 7.6.q-25
telnet: 0.17-40
tk8.6: 8.6.5-1
tzdata: 2016f-0ubuntu0.16.04
ubuntu-keyring: 2012.05.19
ubuntu-minimal: 1.361
ucf: 3.0036
udev: 229-4ubuntu10
uidmap: 1:4.2-3.1ubuntu5.3
ureadahead: 0.100.0-19
usbutils: 1:007-4
util-linux: 2.27.1-6ubuntu3.1
uuid-runtime: 2.27.1-6ubuntu3.2
vim-common: 2:7.4.1689-3ubuntu1.1
vim-tiny: 2:7.4.1689-3ubuntu1.1
virtualenv: 15.0.1+ds-3
wamerican: 7.1-1
wget: 1.17.1-1ubuntu1.1
whiptail: 0.52.18-1ubuntu2
wireless-regdb: 2015.07.20-1ubuntu1
x11-common: 1:7.7+13ubuntu3
x11-utils: 7.7+3
x11proto-core-dev: 7.0.31-1~ubuntu16.04.1
x11proto-input-dev: 2.3.1-1
x11proto-kb-dev: 1.0.7-0ubuntu1
xauth: 1:1.0.9-1ubuntu2
xbitmaps: 1.1.1-2
xdg-user-dirs: 0.15-2ubuntu6
xkb-data: 2.16-1ubuntu1
xml-core: 0.13+nmu2
xorg-sgml-doctools: 1:1.11-1
xterm: 322-1ubuntu1
xtrans-dev: 1.3.5-1
xz-utils: 5.1.1alpha+20120614-2ubuntu2
zlib1g:amd64: 1:1.2.8.dfsg-2ubuntu4
zlib1g-dev:amd64: 1:1.2.8.dfsg-2ubuntu4

```

Kernel module listing

```

$ lsmod | sort
8250_fintek          16384  0
ablk_helper         16384  1 aesni_intel
acpi_pad            24576  0
acpi_power_meter    20480  0
aesni_intel         167936  0
aes_x86_64          20480  1 aesni_intel
ahci                 36864  0
authenc             16384  1 intel_qat
autofs4             40960  2
bridge              126976  0
coretemp            16384  0
crc32_pclmul        16384  0
crct10dif_pclmul    16384  0
cryptd              20480  3 ghash_clmulni_intel,aesni_intel,ablk_helper
dca                 16384  2 igb,ixgbe
edac_core            53248  1 sb_edac
enclosure           16384  1 ses
enic                81920  0
fjes                28672  0
fnic                106496  0

```

(continues on next page)

(continued from previous page)

| | | | |
|------------------------|--------|---------|--|
| gf128mul | 16384 | 1 | lrw |
| ghash_clmulni_intel | 16384 | 0 | |
| glue_helper | 16384 | 1 | aesni_intel |
| hid | 118784 | 2 | hid_generic,usbhid |
| hid_generic | 16384 | 0 | |
| i2c_algo_bit | 16384 | 1 | igb |
| i40e | 286720 | 0 | |
| igb | 196608 | 0 | |
| igb_uio | 16384 | 0 | |
| input_leds | 16384 | 0 | |
| intel_powerclamp | 16384 | 0 | |
| intel_qat | 110592 | 2 | qat_dh895xccvf,qat_dh895xcc |
| intel_rapl | 20480 | 0 | |
| ip6_udp_tunnel | 16384 | 1 | vxlan |
| ipmi_msghandler | 49152 | 2 | ipmi_ssif,ipmi_si |
| ipmi_si | 57344 | 0 | |
| ipmi_ssif | 24576 | 0 | |
| iptable_filter | 16384 | 1 | |
| iptable_mangle | 16384 | 1 | |
| iptable_nat | 16384 | 1 | |
| ip_tables | 24576 | 3 | iptable_filter, iptable_mangle, iptable_nat |
| ipt_MASQUERADE | 16384 | 1 | |
| irqbypass | 16384 | 1 | kvm |
| ixgbe | 290816 | 0 | |
| joydev | 20480 | 0 | |
| kvm | 544768 | 1 | kvm_intel |
| kvm_intel | 172032 | 0 | |
| libahci | 32768 | 1 | ahci |
| libfc | 114688 | 2 | fnic,libfcoe |
| libfcoe | 65536 | 1 | fnic |
| llc | 16384 | 2 | stp,bridge |
| lpc_ich | 24576 | 0 | |
| lrw | 16384 | 1 | aesni_intel |
| mac_hid | 16384 | 0 | |
| mdio | 16384 | 1 | ixgbe |
| megaraid_sas | 135168 | 3 | |
| mei | 98304 | 1 | mei_me |
| mei_me | 36864 | 0 | |
| Module | Size | Used by | |
| nf_contrack | 106496 | 4 | nf_nat,nf_nat_ipv4,nf_nat_masquerade_ipv4,nf_contrack_ipv4 |
| nf_contrack_ipv4 | 16384 | 1 | |
| nf_defrag_ipv4 | 16384 | 1 | nf_contrack_ipv4 |
| nf_nat | 24576 | 2 | nf_nat_ipv4,nf_nat_masquerade_ipv4 |
| nf_nat_ipv4 | 16384 | 1 | iptable_nat |
| nf_nat_masquerade_ipv4 | 16384 | 1 | ipt_MASQUERADE |
| pps_core | 20480 | 1 | ptp |
| ptp | 20480 | 3 | igb,i40e,ixgbe |
| qat_dh895xcc | 20480 | 0 | |
| qat_dh895xccvf | 20480 | 0 | |
| sb_edac | 32768 | 0 | |
| scsi_transport_fc | 61440 | 2 | fnic,libfc |
| ses | 20480 | 0 | |
| shpchp | 36864 | 0 | |
| stp | 16384 | 1 | bridge |
| udp_tunnel | 16384 | 1 | vxlan |
| uio | 20480 | 2 | uio_pci_generic,igb_uio |
| uio_pci_generic | 16384 | 0 | |
| usbhid | 49152 | 0 | |
| veth | 16384 | 0 | |
| vxlan | 49152 | 2 | i40e,ixgbe |
| wmi | 20480 | 0 | |

(continues on next page)

(continued from previous page)

```
x86_pkg_temp_thermal    16384  0
x_tables                36864  6 xt_CHECKSUM,ip_tables,xt_tcpudp,ipt_MASQUERADE,iptable_filter,
↔ iptable_mangle
xt_CHECKSUM             16384  1
xt_tcpudp               16384  5
```

Sysctl listing

```
$ sysctl -a
abi.vsyscall32 = 1
debug.exception-trace = 1
debug.kprobes-optimization = 1
dev.cdrom.autoclose = 1
dev.cdrom.autoeject = 0
dev.cdrom.check_media = 0
dev.cdrom.debug = 0
dev.cdrom.info = CD-ROM information, Id: cdrom.c 3.20 2003/12/17
dev.cdrom.info =
dev.cdrom.info = drive name:
dev.cdrom.info = drive speed:
dev.cdrom.info = drive # of slots:
dev.cdrom.info = Can close tray:
dev.cdrom.info = Can open tray:
dev.cdrom.info = Can lock tray:
dev.cdrom.info = Can change speed:
dev.cdrom.info = Can select disk:
dev.cdrom.info = Can read multisession:
dev.cdrom.info = Can read MCN:
dev.cdrom.info = Reports media changed:
dev.cdrom.info = Can play audio:
dev.cdrom.info = Can write CD-R:
dev.cdrom.info = Can write CD-RW:
dev.cdrom.info = Can read DVD:
dev.cdrom.info = Can write DVD-R:
dev.cdrom.info = Can write DVD-RAM:
dev.cdrom.info = Can read MRW:
dev.cdrom.info = Can write MRW:
dev.cdrom.info = Can write RAM:
dev.cdrom.info =
dev.cdrom.info =
dev.cdrom.lock = 0
dev.hpet.max-user-freq = 64
dev.mac_hid.mouse_button2_keycode = 97
dev.mac_hid.mouse_button3_keycode = 100
dev.mac_hid.mouse_button_emulation = 0
dev.raid.speed_limit_max = 200000
dev.raid.speed_limit_min = 1000
dev.scsi.logging_level = 0
fs.aio-max-nr = 65536
fs.aio-nr = 0
fs.binfmt_misc.status = enabled
fs.dentry-state = 69970      58326  45      0      0      0
fs.dir-notify-enable = 1
fs.epoll.max_user_watches = 108185784
fs.file-max = 52706330
fs.file-nr = 1224  0      52706330
fs.inode-nr = 42965 369
fs.inode-state = 42965      369      0      0      0      0
fs.inotify.max_queued_events = 16384
fs.inotify.max_user_instances = 128
fs.inotify.max_user_watches = 8192
```

(continues on next page)

(continued from previous page)

```
fs.lease-break-time = 45
fs.leases-enable = 1
fs.mount-max = 100000
fs.mqueue.msg_default = 10
fs.mqueue.msg_max = 10
fs.mqueue.msgsize_default = 8192
fs.mqueue.msgsize_max = 8192
fs.mqueue.queues_max = 256
fs.nr_open = 1048576
fs.overflowgid = 65534
fs.overflowuid = 65534
fs.pipe-max-size = 1048576
fs.pipe-user-pages-hard = 0
fs.pipe-user-pages-soft = 16384
fs.protected_hardlinks = 1
fs.protected_symlinks = 1
fs.quota.allocated_dquots = 0
fs.quota.cache_hits = 0
fs.quota.drops = 0
fs.quota.free_dquots = 0
fs.quota.lookups = 0
fs.quota.reads = 0
fs.quota.syncs = 0
fs.quota.writes = 0
fs.suid_dumpable = 0
kernel.acct = 4      2      30
kernel.acpi_video_flags = 0
kernel.auto_msgmni = 0
kernel.bootloader_type = 114
kernel.bootloader_version = 2
kernel.cad_pid = 1
kernel.cap_last_cap = 37
kernel.compat-log = 1
kernel.core_pattern = core
kernel.core_pipe_limit = 0
kernel.core_uses_pid = 0
kernel.ctrl-alt-del = 0
kernel.dmesg_restrict = 0
kernel.domainname = (none)
kernel.ftrace_dump_on_oops = 0
kernel.ftrace_enabled = 1
kernel.hardlockup_all_cpu_backtrace = 0
kernel.hardlockup_panic = 0
kernel.hostname = t2-sut1
kernel.hotplug =
kernel.hung_task_check_count = 4194304
kernel.hung_task_panic = 0
kernel.hung_task_timeout_secs = 120
kernel.hung_task_warnings = 10
kernel.io_delay_type = 1
kernel.kexec_load_disabled = 0
kernel.keys.gc_delay = 300
kernel.keys.maxbytes = 20000
kernel.keys.maxkeys = 200
kernel.keys.persistent_keyring_expiry = 259200
kernel.keys.root_maxbytes = 25000000
kernel.keys.root_maxkeys = 1000000
kernel.kptr_restrict = 1
kernel.kstack_depth_to_print = 12
kernel.max_lock_depth = 1024
kernel.modprobe = /sbin/modprobe
```

(continues on next page)

(continued from previous page)

```
kernel.modules_disabled = 0
kernel.moksbstate_disabled = 0
kernel.msg_next_id = -1
kernel.msgmax = 8192
kernel.msgmnb = 16384
kernel.msgmni = 32000
kernel.ngroups_max = 65536
kernel.nmi_watchdog = 1
kernel.ns_last_pid = 11764
kernel.numa_balancing = 1
kernel.numa_balancing_scan_delay_ms = 1000
kernel.numa_balancing_scan_period_max_ms = 60000
kernel.numa_balancing_scan_period_min_ms = 1000
kernel.numa_balancing_scan_size_mb = 256
kernel.osrelease = 4.4.0-72-generic
kernel.ostype = Linux
kernel.overflowgid = 65534
kernel.overflowuid = 65534
kernel.panic = 0
kernel.panic_on_io_nmi = 0
kernel.panic_on_oops = 0
kernel.panic_on_unrecovered_nmi = 0
kernel.panic_on_warn = 0
kernel.perf_cpu_time_max_percent = 25
kernel.perf_event_max_sample_rate = 12500
kernel.perf_event_mlock_kb = 516
kernel.perf_event_paranoid = 1
kernel.pid_max = 36864
kernel.poweroff_cmd = /sbin/poweroff
kernel.print-fatal-signals = 0
kernel.printk = 4 4 1 7
kernel.printk_delay = 0
kernel.printk_ratelimit = 5
kernel.printk_ratelimit_burst = 10
kernel.ptty_max = 4096
kernel.ptty_nr = 1
kernel.ptty_reserve = 1024
kernel.random.boot_id = f683c836-6fc6-492a-a23b-62ab21895040
kernel.random.entropy_avail = 200
kernel.random.poolsize = 4096
kernel.random.read_wakeup_threshold = 64
kernel.random.urandom_min_reseed_secs = 60
kernel.random.uuid = 144ff2ba-1bc7-4836-8fb7-6aaa0ab7e65f
kernel.random.write_wakeup_threshold = 896
kernel.randomize_va_space = 0
kernel.real-root-dev = 0
kernel.sched_autogroup_enabled = 1
kernel.sched_cfs_bandwidth_slice_us = 5000
kernel.sched_child_runs_first = 0
kernel.sched_domain.cpu0.domain0.busy_factor = 32
kernel.sched_domain.cpu0.domain0.busy_idx = 3
kernel.sched_domain.cpu0.domain0.cache_nice_tries = 2
kernel.sched_domain.cpu0.domain0.flags = 25647
kernel.sched_domain.cpu0.domain0.forkexec_idx = 0
kernel.sched_domain.cpu0.domain0.idle_idx = 2
kernel.sched_domain.cpu0.domain0.imbalance_pct = 125
kernel.sched_domain.cpu0.domain0.max_interval = 72
kernel.sched_domain.cpu0.domain0.max_newidle_lb_cost = 1309
kernel.sched_domain.cpu0.domain0.min_interval = 36
kernel.sched_domain.cpu0.domain0.name = NUMA
kernel.sched_domain.cpu0.domain0.newidle_idx = 0
```

(continues on next page)

(continued from previous page)

```
kernel.sched_domain.cpu0.domain0.wake_idx = 0
kernel.sched_domain.cpu18.domain0.busy_factor = 32
kernel.sched_domain.cpu18.domain0.busy_idx = 3
kernel.sched_domain.cpu18.domain0.cache_nice_tries = 2
kernel.sched_domain.cpu18.domain0.flags = 25647
kernel.sched_domain.cpu18.domain0.forkexec_idx = 0
kernel.sched_domain.cpu18.domain0.idle_idx = 2
kernel.sched_domain.cpu18.domain0.imbalance_pct = 125
kernel.sched_domain.cpu18.domain0.max_interval = 72
kernel.sched_domain.cpu18.domain0.max_newidle_lb_cost = 2026
kernel.sched_domain.cpu18.domain0.min_interval = 36
kernel.sched_domain.cpu18.domain0.name = NUMA
kernel.sched_domain.cpu18.domain0.newidle_idx = 0
kernel.sched_domain.cpu18.domain0.wake_idx = 0
kernel.sched_latency_ns = 24000000
kernel.sched_migration_cost_ns = 500000
kernel.sched_min_granularity_ns = 3000000
kernel.sched_nr_migrate = 32
kernel.sched_rr_timeslice_ms = 25
kernel.sched_rt_period_us = 1000000
kernel.sched_rt_runtime_us = 950000
kernel.sched_shares_window_ns = 10000000
kernel.sched_time_avg_ms = 1000
kernel.sched_tunable_scaling = 1
kernel.sched_wakeup_granularity_ns = 4000000
kernel.secure_boot = 0
kernel.sem = 32000 1024000000 500 32000
kernel.sem_next_id = -1
kernel.sg-big-buff = 32768
kernel.shm_next_id = -1
kernel.shm_rmid_forced = 0
kernel.shmall = 18446744073692774399
kernel.shmmax = 8589934592
kernel.shmmni = 4096
kernel.soft_watchdog = 1
kernel.softlockup_all_cpu_backtrace = 0
kernel.softlockup_panic = 0
kernel.stack_tracer_enabled = 0
kernel.sysctl_writes_strict = 0
kernel.sysrq = 176
kernel.tainted = 12288
kernel.threads-max = 4126960
kernel.timer_migration = 1
kernel.traceoff_on_warning = 0
kernel.tracepoint_printk = 0
kernel.unknown_nmi_panic = 0
kernel.unprivileged_bpf_disabled = 0
kernel.unprivileged_usersns_apparmor_policy = 1
kernel.unprivileged_usersns_clone = 1
kernel.usermodehelper.bset = 4294967295 63
kernel.usermodehelper.inheritable = 4294967295 63
kernel.version = #93-Ubuntu SMP Fri Mar 31 14:07:41 UTC 2017
kernel.watchdog = 1
kernel.watchdog_cpumask = 0,18
kernel.watchdog_thresh = 10
kernel.yama.ptrace_scope = 1
net.core.bpf_jit_enable = 0
net.core.busy_poll = 0
net.core.busy_read = 0
net.core.default_qdisc = pfifo_fast
net.core.dev_weight = 64
```

(continues on next page)

(continued from previous page)

```
net.core.flow_limit_cpu_bitmap = 0,00000000
net.core.flow_limit_table_len = 4096
net.core.max_skb_frags = 17
net.core.message_burst = 10
net.core.message_cost = 5
net.core.netdev_budget = 300
net.core.netdev_max_backlog = 1000
net.core.netdev_rss_key =
↳29:61:61:e6:4e:d5:d0:a2:dc:81:6a:c8:44:1b:e2:8d:c8:6f:6a:2b:64:62:98:08:bb:63:48:8e:96:d1:6a:15:32:ca:da:8d:3c:0a:e
net.core.netdev_tstamp_prequeue = 1
net.core.optmem_max = 20480
net.core.rmem_default = 212992
net.core.rmem_max = 212992
net.core.rps_sock_flow_entries = 0
net.core.somaxconn = 128
net.core.tstamp_allow_data = 1
net.core.warnings = 0
net.core.wmem_default = 212992
net.core.wmem_max = 212992
net.core.xfrm_acq_expires = 30
net.core.xfrm_aevent_etime = 10
net.core.xfrm_aevent_rseqth = 2
net.core.xfrm_larval_drop = 1
net.fan.vxlan = 4
net.ipv4.cipso_cache_bucket_size = 10
net.ipv4.cipso_cache_enable = 1
net.ipv4.cipso_rbm_optfmt = 0
net.ipv4.cipso_rbm_strictvalid = 1
net.ipv4.conf.all.accept_local = 0
net.ipv4.conf.all.accept_redirects = 0
net.ipv4.conf.all.accept_source_route = 0
net.ipv4.conf.all.arp_accept = 0
net.ipv4.conf.all.arp_announce = 0
net.ipv4.conf.all.arp_filter = 0
net.ipv4.conf.all.arp_ignore = 0
net.ipv4.conf.all.arp_notify = 0
net.ipv4.conf.all.bootp_relay = 0
net.ipv4.conf.all.disable_policy = 0
net.ipv4.conf.all.disable_xfrm = 0
net.ipv4.conf.all.force_igmp_version = 0
net.ipv4.conf.all.forwarding = 1
net.ipv4.conf.all.igmpv2_unsolicited_report_interval = 10000
net.ipv4.conf.all.igmpv3_unsolicited_report_interval = 1000
net.ipv4.conf.all.ignore_routes_with_linkdown = 0
net.ipv4.conf.all.log_martians = 0
net.ipv4.conf.all.mc_forwarding = 0
net.ipv4.conf.all.medium_id = 0
net.ipv4.conf.all.promote_secondaries = 0
net.ipv4.conf.all.proxy_arp = 0
net.ipv4.conf.all.proxy_arp_pvlan = 0
net.ipv4.conf.all.route_localnet = 0
net.ipv4.conf.all.rp_filter = 1
net.ipv4.conf.all.secure_redirects = 1
net.ipv4.conf.all.send_redirects = 1
net.ipv4.conf.all.shared_media = 1
net.ipv4.conf.all.src_valid_mark = 0
net.ipv4.conf.all.tag = 0
net.ipv4.conf.default.accept_local = 0
net.ipv4.conf.default.accept_redirects = 1
net.ipv4.conf.default.accept_source_route = 1
net.ipv4.conf.default.arp_accept = 0
```

(continues on next page)

(continued from previous page)

```
net.ipv4.conf.default.arp_announce = 0
net.ipv4.conf.default.arp_filter = 0
net.ipv4.conf.default.arp_ignore = 0
net.ipv4.conf.default.arp_notify = 0
net.ipv4.conf.default.bootp_relay = 0
net.ipv4.conf.default.disable_policy = 0
net.ipv4.conf.default.disable_xfrm = 0
net.ipv4.conf.default.force_igmp_version = 0
net.ipv4.conf.default.forwarding = 1
net.ipv4.conf.default.igmpv2_unsolicited_report_interval = 10000
net.ipv4.conf.default.igmpv3_unsolicited_report_interval = 1000
net.ipv4.conf.default.ignore_routes_with_linkdown = 0
net.ipv4.conf.default.log_martians = 0
net.ipv4.conf.default.mc_forwarding = 0
net.ipv4.conf.default.medium_id = 0
net.ipv4.conf.default.promote_secondaries = 0
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.default.proxy_arp_pvlan = 0
net.ipv4.conf.default.route_localnet = 0
net.ipv4.conf.default.rp_filter = 1
net.ipv4.conf.default.secure_redirects = 1
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.default.shared_media = 1
net.ipv4.conf.default.src_valid_mark = 0
net.ipv4.conf.default.tag = 0
net.ipv4.conf.enp25s0f0.accept_local = 0
net.ipv4.conf.enp25s0f0.accept_redirects = 1
net.ipv4.conf.enp25s0f0.accept_source_route = 1
net.ipv4.conf.enp25s0f0.arp_accept = 0
net.ipv4.conf.enp25s0f0.arp_announce = 0
net.ipv4.conf.enp25s0f0.arp_filter = 0
net.ipv4.conf.enp25s0f0.arp_ignore = 0
net.ipv4.conf.enp25s0f0.arp_notify = 0
net.ipv4.conf.enp25s0f0.bootp_relay = 0
net.ipv4.conf.enp25s0f0.disable_policy = 0
net.ipv4.conf.enp25s0f0.disable_xfrm = 0
net.ipv4.conf.enp25s0f0.force_igmp_version = 0
net.ipv4.conf.enp25s0f0.forwarding = 1
net.ipv4.conf.enp25s0f0.igmpv2_unsolicited_report_interval = 10000
net.ipv4.conf.enp25s0f0.igmpv3_unsolicited_report_interval = 1000
net.ipv4.conf.enp25s0f0.ignore_routes_with_linkdown = 0
net.ipv4.conf.enp25s0f0.log_martians = 0
net.ipv4.conf.enp25s0f0.mc_forwarding = 0
net.ipv4.conf.enp25s0f0.medium_id = 0
net.ipv4.conf.enp25s0f0.promote_secondaries = 0
net.ipv4.conf.enp25s0f0.proxy_arp = 0
net.ipv4.conf.enp25s0f0.proxy_arp_pvlan = 0
net.ipv4.conf.enp25s0f0.route_localnet = 0
net.ipv4.conf.enp25s0f0.rp_filter = 1
net.ipv4.conf.enp25s0f0.secure_redirects = 1
net.ipv4.conf.enp25s0f0.send_redirects = 1
net.ipv4.conf.enp25s0f0.shared_media = 1
net.ipv4.conf.enp25s0f0.src_valid_mark = 0
net.ipv4.conf.enp25s0f0.tag = 0
net.ipv4.conf.lo.accept_local = 0
net.ipv4.conf.lo.accept_redirects = 1
net.ipv4.conf.lo.accept_source_route = 1
net.ipv4.conf.lo.arp_accept = 0
net.ipv4.conf.lo.arp_announce = 0
net.ipv4.conf.lo.arp_filter = 0
net.ipv4.conf.lo.arp_ignore = 0
```

(continues on next page)

(continued from previous page)

```
net.ipv4.conf.lo.arp_notify = 0
net.ipv4.conf.lo.bootp_relay = 0
net.ipv4.conf.lo.disable_policy = 1
net.ipv4.conf.lo.disable_xfrm = 1
net.ipv4.conf.lo.force_igmp_version = 0
net.ipv4.conf.lo.forwarding = 1
net.ipv4.conf.lo.igmpv2_unsolicited_report_interval = 10000
net.ipv4.conf.lo.igmpv3_unsolicited_report_interval = 1000
net.ipv4.conf.lo.ignore_routes_with_linkdown = 0
net.ipv4.conf.lo.log_martians = 0
net.ipv4.conf.lo.mc_forwarding = 0
net.ipv4.conf.lo.medium_id = 0
net.ipv4.conf.lo.promote_secondaries = 0
net.ipv4.conf.lo.proxy_arp = 0
net.ipv4.conf.lo.proxy_arp_pvlan = 0
net.ipv4.conf.lo.route_localnet = 0
net.ipv4.conf.lo.rp_filter = 0
net.ipv4.conf.lo.secure_redirects = 1
net.ipv4.conf.lo.send_redirects = 1
net.ipv4.conf.lo.shared_media = 1
net.ipv4.conf.lo.src_valid_mark = 0
net.ipv4.conf.lo.tag = 0
net.ipv4.conf.lxcbr0.accept_local = 0
net.ipv4.conf.lxcbr0.accept_redirects = 1
net.ipv4.conf.lxcbr0.accept_source_route = 1
net.ipv4.conf.lxcbr0.arp_accept = 0
net.ipv4.conf.lxcbr0.arp_announce = 0
net.ipv4.conf.lxcbr0.arp_filter = 0
net.ipv4.conf.lxcbr0.arp_ignore = 0
net.ipv4.conf.lxcbr0.arp_notify = 0
net.ipv4.conf.lxcbr0.bootp_relay = 0
net.ipv4.conf.lxcbr0.disable_policy = 0
net.ipv4.conf.lxcbr0.disable_xfrm = 0
net.ipv4.conf.lxcbr0.force_igmp_version = 0
net.ipv4.conf.lxcbr0.forwarding = 1
net.ipv4.conf.lxcbr0.igmpv2_unsolicited_report_interval = 10000
net.ipv4.conf.lxcbr0.igmpv3_unsolicited_report_interval = 1000
net.ipv4.conf.lxcbr0.ignore_routes_with_linkdown = 0
net.ipv4.conf.lxcbr0.log_martians = 0
net.ipv4.conf.lxcbr0.mc_forwarding = 0
net.ipv4.conf.lxcbr0.medium_id = 0
net.ipv4.conf.lxcbr0.promote_secondaries = 0
net.ipv4.conf.lxcbr0.proxy_arp = 0
net.ipv4.conf.lxcbr0.proxy_arp_pvlan = 0
net.ipv4.conf.lxcbr0.route_localnet = 0
net.ipv4.conf.lxcbr0.rp_filter = 1
net.ipv4.conf.lxcbr0.secure_redirects = 1
net.ipv4.conf.lxcbr0.send_redirects = 1
net.ipv4.conf.lxcbr0.shared_media = 1
net.ipv4.conf.lxcbr0.src_valid_mark = 0
net.ipv4.conf.lxcbr0.tag = 0
net.ipv4.fwmark_reflect = 0
net.ipv4.icmp_echo_ignore_all = 0
net.ipv4.icmp_echo_ignore_broadcasts = 1
net.ipv4.icmp_errors_use_inbound_ifaddr = 0
net.ipv4.icmp_ignore_bogus_error_responses = 1
net.ipv4.icmp_msgs_burst = 50
net.ipv4.icmp_msgs_per_sec = 1000
net.ipv4.icmp_ratelimit = 1000
net.ipv4.icmp_ratemask = 6168
net.ipv4.igmp_link_local_mcast_reports = 1
```

(continues on next page)

(continued from previous page)

```
net.ipv4.igmp_max_memberships = 20
net.ipv4.igmp_max_msf = 10
net.ipv4.igmp_qrv = 2
net.ipv4.inet_peer_maxttl = 600
net.ipv4.inet_peer_minttl = 120
net.ipv4.inet_peer_threshold = 65664
net.ipv4.ip_default_ttl = 64
net.ipv4.ip_dynaddr = 0
net.ipv4.ip_early_demux = 1
net.ipv4.ip_forward = 1
net.ipv4.ip_forward_use_pmtu = 0
net.ipv4.ip_local_port_range = 32768          60999
net.ipv4.ip_local_reserved_ports =
net.ipv4.ip_no_pmtu_disc = 0
net.ipv4.ip_nonlocal_bind = 0
net.ipv4.ipfrag_high_thresh = 4194304
net.ipv4.ipfrag_low_thresh = 3145728
net.ipv4.ipfrag_max_dist = 64
net.ipv4.ipfrag_secret_interval = 0
net.ipv4.ipfrag_time = 30
net.ipv4.neigh.default.anycast_delay = 100
net.ipv4.neigh.default.app_solicit = 0
net.ipv4.neigh.default.base_reachable_time_ms = 30000
net.ipv4.neigh.default.delay_first_probe_time = 5
net.ipv4.neigh.default.gc_interval = 30
net.ipv4.neigh.default.gc_stale_time = 60
net.ipv4.neigh.default.gc_thresh1 = 128
net.ipv4.neigh.default.gc_thresh2 = 512
net.ipv4.neigh.default.gc_thresh3 = 1024
net.ipv4.neigh.default.locktime = 100
net.ipv4.neigh.default.mcast_resolicit = 0
net.ipv4.neigh.default.mcast_solicit = 3
net.ipv4.neigh.default.proxy_delay = 80
net.ipv4.neigh.default.proxy_qlen = 64
net.ipv4.neigh.default.retrans_time_ms = 1000
net.ipv4.neigh.default.ucast_solicit = 3
net.ipv4.neigh.default.unres_qlen = 31
net.ipv4.neigh.default.unres_qlen_bytes = 65536
net.ipv4.neigh.enp25s0f0.anycast_delay = 100
net.ipv4.neigh.enp25s0f0.app_solicit = 0
net.ipv4.neigh.enp25s0f0.base_reachable_time_ms = 30000
net.ipv4.neigh.enp25s0f0.delay_first_probe_time = 5
net.ipv4.neigh.enp25s0f0.gc_stale_time = 60
net.ipv4.neigh.enp25s0f0.locktime = 100
net.ipv4.neigh.enp25s0f0.mcast_resolicit = 0
net.ipv4.neigh.enp25s0f0.mcast_solicit = 3
net.ipv4.neigh.enp25s0f0.proxy_delay = 80
net.ipv4.neigh.enp25s0f0.proxy_qlen = 64
net.ipv4.neigh.enp25s0f0.retrans_time_ms = 1000
net.ipv4.neigh.enp25s0f0.ucast_solicit = 3
net.ipv4.neigh.enp25s0f0.unres_qlen = 31
net.ipv4.neigh.enp25s0f0.unres_qlen_bytes = 65536
net.ipv4.neigh.lo.anycast_delay = 100
net.ipv4.neigh.lo.app_solicit = 0
net.ipv4.neigh.lo.base_reachable_time_ms = 30000
net.ipv4.neigh.lo.delay_first_probe_time = 5
net.ipv4.neigh.lo.gc_stale_time = 60
net.ipv4.neigh.lo.locktime = 100
net.ipv4.neigh.lo.mcast_resolicit = 0
net.ipv4.neigh.lo.mcast_solicit = 3
net.ipv4.neigh.lo.proxy_delay = 80
```

(continues on next page)

(continued from previous page)

```
net.ipv4.neigh.lo.proxy_qlen = 64
net.ipv4.neigh.lo.retrans_time_ms = 1000
net.ipv4.neigh.lo.ucast_solicit = 3
net.ipv4.neigh.lo.unres_qlen = 31
net.ipv4.neigh.lo.unres_qlen_bytes = 65536
net.ipv4.neigh.lxcbr0.anycast_delay = 100
net.ipv4.neigh.lxcbr0.app_solicit = 0
net.ipv4.neigh.lxcbr0.base_reachable_time_ms = 30000
net.ipv4.neigh.lxcbr0.delay_first_probe_time = 5
net.ipv4.neigh.lxcbr0.gc_stale_time = 60
net.ipv4.neigh.lxcbr0.locktime = 100
net.ipv4.neigh.lxcbr0.mcast_resolicit = 0
net.ipv4.neigh.lxcbr0.mcast_solicit = 3
net.ipv4.neigh.lxcbr0.proxy_delay = 80
net.ipv4.neigh.lxcbr0.proxy_qlen = 64
net.ipv4.neigh.lxcbr0.retrans_time_ms = 1000
net.ipv4.neigh.lxcbr0.ucast_solicit = 3
net.ipv4.neigh.lxcbr0.unres_qlen = 31
net.ipv4.neigh.lxcbr0.unres_qlen_bytes = 65536
net.ipv4.ping_group_range = 1 0
net.ipv4.route.error_burst = 1250
net.ipv4.route.error_cost = 250
net.ipv4.route.gc_elasticity = 8
net.ipv4.route.gc_interval = 60
net.ipv4.route.gc_min_interval = 0
net.ipv4.route.gc_min_interval_ms = 500
net.ipv4.route.gc_thresh = -1
net.ipv4.route.gc_timeout = 300
net.ipv4.route.max_size = 2147483647
net.ipv4.route.min_adv_mss = 256
net.ipv4.route.min_pmtu = 552
net.ipv4.route.mtu_expires = 600
net.ipv4.route.redirect_load = 5
net.ipv4.route.redirect_number = 9
net.ipv4.route.redirect_silence = 5120
net.ipv4.tcp_abort_on_overflow = 0
net.ipv4.tcp_adv_win_scale = 1
net.ipv4.tcp_allowed_congestion_control = cubic reno
net.ipv4.tcp_app_win = 31
net.ipv4.tcp_autocorking = 1
net.ipv4.tcp_available_congestion_control = cubic reno
net.ipv4.tcp_base_mss = 1024
net.ipv4.tcp_challenge_ack_limit = 1000
net.ipv4.tcp_congestion_control = cubic
net.ipv4.tcp_dsack = 1
net.ipv4.tcp_early_retrans = 3
net.ipv4.tcp_ecn = 2
net.ipv4.tcp_ecn_fallback = 1
net.ipv4.tcp_fack = 1
net.ipv4.tcp_fastopen = 1
net.ipv4.tcp_fastopen_key = 00000000-00000000-00000000-00000000
net.ipv4.tcp_fin_timeout = 60
net.ipv4.tcp_frto = 2
net.ipv4.tcp_fwmark_accept = 0
net.ipv4.tcp_invalid_ratelimit = 500
net.ipv4.tcp_keepalive_intvl = 75
net.ipv4.tcp_keepalive_probes = 9
net.ipv4.tcp_keepalive_time = 7200
net.ipv4.tcp_limit_output_bytes = 262144
net.ipv4.tcp_low_latency = 0
net.ipv4.tcp_max_orphans = 262144
```

(continues on next page)

(continued from previous page)

```
net.ipv4.tcp_max_reordering = 300
net.ipv4.tcp_max_syn_backlog = 2048
net.ipv4.tcp_max_tw_buckets = 262144
net.ipv4.tcp_mem = 6188856 8251809 12377712
net.ipv4.tcp_min_rtt_wlen = 300
net.ipv4.tcp_min_tso_segs = 2
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_mtu_probing = 0
net.ipv4.tcp_no_metrics_save = 0
net.ipv4.tcp_notsent_lowat = -1
net.ipv4.tcp_orphan_retries = 0
net.ipv4.tcp_pacing_ca_ratio = 120
net.ipv4.tcp_pacing_ss_ratio = 200
net.ipv4.tcp_probe_interval = 600
net.ipv4.tcp_probe_threshold = 8
net.ipv4.tcp_recovery = 1
net.ipv4.tcp_reordering = 3
net.ipv4.tcp_retrans_collapse = 1
net.ipv4.tcp_retries1 = 3
net.ipv4.tcp_retries2 = 15
net.ipv4.tcp_rfc1337 = 0
net.ipv4.tcp_rmem = 4096 87380 6291456
net.ipv4.tcp_sack = 1
net.ipv4.tcp_slow_start_after_idle = 1
net.ipv4.tcp_stdurg = 0
net.ipv4.tcp_syn_retries = 6
net.ipv4.tcp_synack_retries = 5
net.ipv4.tcp_syncookies = 1
net.ipv4.tcp_thin_dupack = 0
net.ipv4.tcp_thin_linear_timeouts = 0
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_tso_win_divisor = 3
net.ipv4.tcp_tw_recycle = 0
net.ipv4.tcp_tw_reuse = 0
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_wmem = 4096 16384 4194304
net.ipv4.tcp_workaround_signed_windows = 0
net.ipv4.udp_mem = 12377712 16503618 24755424
net.ipv4.udp_rmem_min = 4096
net.ipv4.udp_wmem_min = 4096
net.ipv4.xfrm4_gc_thresh = 2147483647
net.ipv6.anycast_src_echo_reply = 0
net.ipv6.auto_flowlabels = 1
net.ipv6.bindv6only = 0
net.ipv6.conf.all.accept_dad = 1
net.ipv6.conf.all.accept_ra = 1
net.ipv6.conf.all.accept_ra_defrtr = 1
net.ipv6.conf.all.accept_ra_from_local = 0
net.ipv6.conf.all.accept_ra_min_hop_limit = 1
net.ipv6.conf.all.accept_ra_mtu = 1
net.ipv6.conf.all.accept_ra_pininfo = 1
net.ipv6.conf.all.accept_ra_rt_info_max_plen = 0
net.ipv6.conf.all.accept_ra_rtr_pref = 1
net.ipv6.conf.all.accept_redirects = 1
net.ipv6.conf.all.accept_source_route = 0
net.ipv6.conf.all.autoconf = 1
net.ipv6.conf.all.dad_transmits = 1
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.all.force_mld_version = 0
net.ipv6.conf.all.force_tllao = 0
net.ipv6.conf.all.forwarding = 0
```

(continues on next page)

(continued from previous page)

```
net.ipv6.conf.all.hop_limit = 64
net.ipv6.conf.all.ignore_routes_with_linkdown = 0
net.ipv6.conf.all.max_addresses = 16
net.ipv6.conf.all.max_desync_factor = 600
net.ipv6.conf.all.mc_forwarding = 0
net.ipv6.conf.all.mldv1_unsolicited_report_interval = 10000
net.ipv6.conf.all.mldv2_unsolicited_report_interval = 1000
net.ipv6.conf.all.mtu = 1280
net.ipv6.conf.all.ndisc_notify = 0
net.ipv6.conf.all.proxy_ndp = 0
net.ipv6.conf.all.regen_max_retry = 3
net.ipv6.conf.all.router_probe_interval = 60
net.ipv6.conf.all.router_solicitation_delay = 1
net.ipv6.conf.all.router_solicitation_interval = 4
net.ipv6.conf.all.router_solicitations = 3
sysctl: reading key "net.ipv6.conf.all.stable_secret"
net.ipv6.conf.all.suppress_frag_ndisc = 1
net.ipv6.conf.all.temp_prefered_lft = 86400
net.ipv6.conf.all.temp_valid_lft = 604800
net.ipv6.conf.all.use_oif_addrs_only = 0
net.ipv6.conf.all.use_tempaddr = 2
net.ipv6.conf.default.accept_dad = 1
net.ipv6.conf.default.accept_ra = 1
net.ipv6.conf.default.accept_ra_defrtr = 1
net.ipv6.conf.default.accept_ra_from_local = 0
net.ipv6.conf.default.accept_ra_min_hop_limit = 1
net.ipv6.conf.default.accept_ra_mtu = 1
net.ipv6.conf.default.accept_ra_pinfo = 1
net.ipv6.conf.default.accept_ra_rt_info_max_plen = 0
net.ipv6.conf.default.accept_ra_rtr_pref = 1
net.ipv6.conf.default.accept_redirects = 1
net.ipv6.conf.default.accept_source_route = 0
net.ipv6.conf.default.autoconf = 1
net.ipv6.conf.default.dad_transmits = 1
net.ipv6.conf.default.disable_ipv6 = 0
net.ipv6.conf.default.force_mld_version = 0
net.ipv6.conf.default.force_tllao = 0
net.ipv6.conf.default.forwarding = 0
net.ipv6.conf.default.hop_limit = 64
net.ipv6.conf.default.ignore_routes_with_linkdown = 0
net.ipv6.conf.default.max_addresses = 16
net.ipv6.conf.default.max_desync_factor = 600
net.ipv6.conf.default.mc_forwarding = 0
net.ipv6.conf.default.mldv1_unsolicited_report_interval = 10000
net.ipv6.conf.default.mldv2_unsolicited_report_interval = 1000
net.ipv6.conf.default.mtu = 1280
net.ipv6.conf.default.ndisc_notify = 0
net.ipv6.conf.default.proxy_ndp = 0
net.ipv6.conf.default.regen_max_retry = 3
net.ipv6.conf.default.router_probe_interval = 60
net.ipv6.conf.default.router_solicitation_delay = 1
net.ipv6.conf.default.router_solicitation_interval = 4
net.ipv6.conf.default.router_solicitations = 3
sysctl: reading key "net.ipv6.conf.default.stable_secret"
net.ipv6.conf.default.suppress_frag_ndisc = 1
net.ipv6.conf.default.temp_prefered_lft = 86400
net.ipv6.conf.default.temp_valid_lft = 604800
net.ipv6.conf.default.use_oif_addrs_only = 0
net.ipv6.conf.default.use_tempaddr = 2
net.ipv6.conf.enp25s0f0.accept_dad = 1
net.ipv6.conf.enp25s0f0.accept_ra = 1
```

(continues on next page)

(continued from previous page)

```
net.ipv6.conf.enp25s0f0.accept_ra_defrtr = 1
net.ipv6.conf.enp25s0f0.accept_ra_from_local = 0
net.ipv6.conf.enp25s0f0.accept_ra_min_hop_limit = 1
net.ipv6.conf.enp25s0f0.accept_ra_mtu = 1
net.ipv6.conf.enp25s0f0.accept_ra_pinfo = 1
net.ipv6.conf.enp25s0f0.accept_ra_rt_info_max_plen = 0
net.ipv6.conf.enp25s0f0.accept_ra_rtr_pref = 1
net.ipv6.conf.enp25s0f0.accept_redirects = 1
net.ipv6.conf.enp25s0f0.accept_source_route = 0
net.ipv6.conf.enp25s0f0.autoconf = 1
net.ipv6.conf.enp25s0f0.dad_transmits = 1
net.ipv6.conf.enp25s0f0.disable_ipv6 = 0
net.ipv6.conf.enp25s0f0.force_mld_version = 0
net.ipv6.conf.enp25s0f0.force_tllao = 0
net.ipv6.conf.enp25s0f0.forwarding = 0
net.ipv6.conf.enp25s0f0.hop_limit = 64
net.ipv6.conf.enp25s0f0.ignore_routes_with_linkdown = 0
net.ipv6.conf.enp25s0f0.max_addresses = 16
net.ipv6.conf.enp25s0f0.max_desync_factor = 600
net.ipv6.conf.enp25s0f0.mc_forwarding = 0
net.ipv6.conf.enp25s0f0.mldv1_unsolicited_report_interval = 10000
net.ipv6.conf.enp25s0f0.mldv2_unsolicited_report_interval = 1000
net.ipv6.conf.enp25s0f0.mtu = 1500
net.ipv6.conf.enp25s0f0.ndisc_notify = 0
net.ipv6.conf.enp25s0f0.proxy_ndp = 0
net.ipv6.conf.enp25s0f0.regen_max_retry = 3
net.ipv6.conf.enp25s0f0.router_probe_interval = 60
net.ipv6.conf.enp25s0f0.router_solicitation_delay = 1
net.ipv6.conf.enp25s0f0.router_solicitation_interval = 4
net.ipv6.conf.enp25s0f0.router_solicitations = 3
sysctl: reading key "net.ipv6.conf.enp25s0f0.stable_secret"
net.ipv6.conf.enp25s0f0.suppress_frag_ndisc = 1
net.ipv6.conf.enp25s0f0.temp_prefered_lft = 86400
net.ipv6.conf.enp25s0f0.temp_valid_lft = 604800
net.ipv6.conf.enp25s0f0.use_oif_addrs_only = 0
net.ipv6.conf.enp25s0f0.use_tempaddr = 0
net.ipv6.conf.lo.accept_dad = -1
net.ipv6.conf.lo.accept_ra = 1
net.ipv6.conf.lo.accept_ra_defrtr = 1
net.ipv6.conf.lo.accept_ra_from_local = 0
net.ipv6.conf.lo.accept_ra_min_hop_limit = 1
net.ipv6.conf.lo.accept_ra_mtu = 1
net.ipv6.conf.lo.accept_ra_pinfo = 1
net.ipv6.conf.lo.accept_ra_rt_info_max_plen = 0
net.ipv6.conf.lo.accept_ra_rtr_pref = 1
net.ipv6.conf.lo.accept_redirects = 1
net.ipv6.conf.lo.accept_source_route = 0
net.ipv6.conf.lo.autoconf = 1
net.ipv6.conf.lo.dad_transmits = 1
net.ipv6.conf.lo.disable_ipv6 = 0
net.ipv6.conf.lo.force_mld_version = 0
net.ipv6.conf.lo.force_tllao = 0
net.ipv6.conf.lo.forwarding = 0
net.ipv6.conf.lo.hop_limit = 64
net.ipv6.conf.lo.ignore_routes_with_linkdown = 0
net.ipv6.conf.lo.max_addresses = 16
net.ipv6.conf.lo.max_desync_factor = 600
net.ipv6.conf.lo.mc_forwarding = 0
net.ipv6.conf.lo.mldv1_unsolicited_report_interval = 10000
net.ipv6.conf.lo.mldv2_unsolicited_report_interval = 1000
net.ipv6.conf.lo.mtu = 65536
```

(continues on next page)

(continued from previous page)

```
net.ipv6.conf.lo.ndisc_notify = 0
net.ipv6.conf.lo.proxy_ndp = 0
net.ipv6.conf.lo.regen_max_retry = 3
net.ipv6.conf.lo.router_probe_interval = 60
net.ipv6.conf.lo.router_solicitation_delay = 1
net.ipv6.conf.lo.router_solicitation_interval = 4
net.ipv6.conf.lo.router_solicitations = 3
sysctl: reading key "net.ipv6.conf.lo.stable_secret"
net.ipv6.conf.lo.suppress_frag_ndisc = 1
net.ipv6.conf.lo.temp_prefered_lft = 86400
net.ipv6.conf.lo.temp_valid_lft = 604800
net.ipv6.conf.lo.use_oif_addrs_only = 0
net.ipv6.conf.lo.use_tempaddr = -1
net.ipv6.conf.lxcbr0.accept_dad = 0
net.ipv6.conf.lxcbr0.accept_ra = 1
net.ipv6.conf.lxcbr0.accept_ra_defrtr = 1
net.ipv6.conf.lxcbr0.accept_ra_from_local = 0
net.ipv6.conf.lxcbr0.accept_ra_min_hop_limit = 1
net.ipv6.conf.lxcbr0.accept_ra_mtu = 1
net.ipv6.conf.lxcbr0.accept_ra_pininfo = 1
net.ipv6.conf.lxcbr0.accept_ra_rt_info_max_plen = 0
net.ipv6.conf.lxcbr0.accept_ra_rtr_pref = 1
net.ipv6.conf.lxcbr0.accept_redirects = 1
net.ipv6.conf.lxcbr0.accept_source_route = 0
net.ipv6.conf.lxcbr0.autoconf = 1
net.ipv6.conf.lxcbr0.dad_transmits = 1
net.ipv6.conf.lxcbr0.disable_ipv6 = 0
net.ipv6.conf.lxcbr0.force_mld_version = 0
net.ipv6.conf.lxcbr0.force_tllao = 0
net.ipv6.conf.lxcbr0.forwarding = 0
net.ipv6.conf.lxcbr0.hop_limit = 64
net.ipv6.conf.lxcbr0.ignore_routes_with_linkdown = 0
net.ipv6.conf.lxcbr0.max_addresses = 16
net.ipv6.conf.lxcbr0.max_desync_factor = 600
net.ipv6.conf.lxcbr0.mc_forwarding = 0
net.ipv6.conf.lxcbr0.mldv1_unsolicited_report_interval = 10000
net.ipv6.conf.lxcbr0.mldv2_unsolicited_report_interval = 1000
net.ipv6.conf.lxcbr0.mtu = 1500
net.ipv6.conf.lxcbr0.ndisc_notify = 0
net.ipv6.conf.lxcbr0.proxy_ndp = 0
net.ipv6.conf.lxcbr0.regen_max_retry = 3
net.ipv6.conf.lxcbr0.router_probe_interval = 60
net.ipv6.conf.lxcbr0.router_solicitation_delay = 1
net.ipv6.conf.lxcbr0.router_solicitation_interval = 4
net.ipv6.conf.lxcbr0.router_solicitations = 3
sysctl: reading key "net.ipv6.conf.lxcbr0.stable_secret"
net.ipv6.conf.lxcbr0.suppress_frag_ndisc = 1
net.ipv6.conf.lxcbr0.temp_prefered_lft = 86400
net.ipv6.conf.lxcbr0.temp_valid_lft = 604800
net.ipv6.conf.lxcbr0.use_oif_addrs_only = 0
net.ipv6.conf.lxcbr0.use_tempaddr = 2
net.ipv6.flowlabel_consistency = 1
net.ipv6.flowlabel_state_ranges = 0
net.ipv6.fwmark_reflect = 0
net.ipv6.icmp.ratelimit = 1000
net.ipv6.idgen_delay = 1
net.ipv6.idgen_retries = 3
net.ipv6.ip6frag_high_thresh = 4194304
net.ipv6.ip6frag_low_thresh = 3145728
net.ipv6.ip6frag_secret_interval = 0
net.ipv6.ip6frag_time = 60
```

(continues on next page)

(continued from previous page)

```
net.ipv6.ip_nonlocal_bind = 0
net.ipv6.mld_max_msf = 64
net.ipv6.mld_qrv = 2
net.ipv6.neigh.default.anycast_delay = 100
net.ipv6.neigh.default.app_solicit = 0
net.ipv6.neigh.default.base_reachable_time_ms = 30000
net.ipv6.neigh.default.delay_first_probe_time = 5
net.ipv6.neigh.default.gc_interval = 30
net.ipv6.neigh.default.gc_stale_time = 60
net.ipv6.neigh.default.gc_thresh1 = 128
net.ipv6.neigh.default.gc_thresh2 = 512
net.ipv6.neigh.default.gc_thresh3 = 1024
net.ipv6.neigh.default.locktime = 0
net.ipv6.neigh.default.mcast_resolicit = 0
net.ipv6.neigh.default.mcast_solicit = 3
net.ipv6.neigh.default.proxy_delay = 80
net.ipv6.neigh.default.proxy_qlen = 64
net.ipv6.neigh.default.retrans_time_ms = 1000
net.ipv6.neigh.default.ucast_solicit = 3
net.ipv6.neigh.default.unres_qlen = 31
net.ipv6.neigh.default.unres_qlen_bytes = 65536
net.ipv6.neigh.enp25s0f0.anycast_delay = 100
net.ipv6.neigh.enp25s0f0.app_solicit = 0
net.ipv6.neigh.enp25s0f0.base_reachable_time_ms = 30000
net.ipv6.neigh.enp25s0f0.delay_first_probe_time = 5
net.ipv6.neigh.enp25s0f0.gc_stale_time = 60
net.ipv6.neigh.enp25s0f0.locktime = 0
net.ipv6.neigh.enp25s0f0.mcast_resolicit = 0
net.ipv6.neigh.enp25s0f0.mcast_solicit = 3
net.ipv6.neigh.enp25s0f0.proxy_delay = 80
net.ipv6.neigh.enp25s0f0.proxy_qlen = 64
net.ipv6.neigh.enp25s0f0.retrans_time_ms = 1000
net.ipv6.neigh.enp25s0f0.ucast_solicit = 3
net.ipv6.neigh.enp25s0f0.unres_qlen = 31
net.ipv6.neigh.enp25s0f0.unres_qlen_bytes = 65536
net.ipv6.neigh.lo.anycast_delay = 100
net.ipv6.neigh.lo.app_solicit = 0
net.ipv6.neigh.lo.base_reachable_time_ms = 30000
net.ipv6.neigh.lo.delay_first_probe_time = 5
net.ipv6.neigh.lo.gc_stale_time = 60
net.ipv6.neigh.lo.locktime = 0
net.ipv6.neigh.lo.mcast_resolicit = 0
net.ipv6.neigh.lo.mcast_solicit = 3
net.ipv6.neigh.lo.proxy_delay = 80
net.ipv6.neigh.lo.proxy_qlen = 64
net.ipv6.neigh.lo.retrans_time_ms = 1000
net.ipv6.neigh.lo.ucast_solicit = 3
net.ipv6.neigh.lo.unres_qlen = 31
net.ipv6.neigh.lo.unres_qlen_bytes = 65536
net.ipv6.neigh.lxcbr0.anycast_delay = 100
net.ipv6.neigh.lxcbr0.app_solicit = 0
net.ipv6.neigh.lxcbr0.base_reachable_time_ms = 30000
net.ipv6.neigh.lxcbr0.delay_first_probe_time = 5
net.ipv6.neigh.lxcbr0.gc_stale_time = 60
net.ipv6.neigh.lxcbr0.locktime = 0
net.ipv6.neigh.lxcbr0.mcast_resolicit = 0
net.ipv6.neigh.lxcbr0.mcast_solicit = 3
net.ipv6.neigh.lxcbr0.proxy_delay = 80
net.ipv6.neigh.lxcbr0.proxy_qlen = 64
net.ipv6.neigh.lxcbr0.retrans_time_ms = 1000
net.ipv6.neigh.lxcbr0.ucast_solicit = 3
```

(continues on next page)

(continued from previous page)

```
net.ipv6.neigh.lxcbr0.unres_qlen = 31
net.ipv6.neigh.lxcbr0.unres_qlen_bytes = 65536
net.ipv6.route.gc_elasticity = 9
net.ipv6.route.gc_interval = 30
net.ipv6.route.gc_min_interval = 0
net.ipv6.route.gc_min_interval_ms = 500
net.ipv6.route.gc_thresh = 1024
net.ipv6.route.gc_timeout = 60
net.ipv6.route.max_size = 4096
net.ipv6.route.min_adv_mss = 1220
net.ipv6.route.mtu_expires = 600
net.ipv6.xfrm6_gc_thresh = 2147483647
net.netfilter.nf_contrack_acct = 0
net.netfilter.nf_contrack_buckets = 65536
net.netfilter.nf_contrack_checksum = 1
net.netfilter.nf_contrack_count = 2
net.netfilter.nf_contrack_events = 1
net.netfilter.nf_contrack_expect_max = 1024
net.netfilter.nf_contrack_generic_timeout = 600
net.netfilter.nf_contrack_helper = 1
net.netfilter.nf_contrack_icmp_timeout = 30
net.netfilter.nf_contrack_log_invalid = 0
net.netfilter.nf_contrack_max = 262144
net.netfilter.nf_contrack_tcp_be_liberal = 0
net.netfilter.nf_contrack_tcp_loose = 1
net.netfilter.nf_contrack_tcp_max_retrans = 3
net.netfilter.nf_contrack_tcp_timeout_close = 10
net.netfilter.nf_contrack_tcp_timeout_close_wait = 60
net.netfilter.nf_contrack_tcp_timeout_established = 432000
net.netfilter.nf_contrack_tcp_timeout_fin_wait = 120
net.netfilter.nf_contrack_tcp_timeout_last_ack = 30
net.netfilter.nf_contrack_tcp_timeout_max_retrans = 300
net.netfilter.nf_contrack_tcp_timeout_syn_rcv = 60
net.netfilter.nf_contrack_tcp_timeout_syn_sent = 120
net.netfilter.nf_contrack_tcp_timeout_time_wait = 120
net.netfilter.nf_contrack_tcp_timeout_unacknowledged = 300
net.netfilter.nf_contrack_timestamp = 0
net.netfilter.nf_contrack_udp_timeout = 30
net.netfilter.nf_contrack_udp_timeout_stream = 180
net.netfilter.nf_log.0 = NONE
net.netfilter.nf_log.1 = NONE
net.netfilter.nf_log.10 = NONE
net.netfilter.nf_log.11 = NONE
net.netfilter.nf_log.12 = NONE
net.netfilter.nf_log.2 = NONE
net.netfilter.nf_log.3 = NONE
net.netfilter.nf_log.4 = NONE
net.netfilter.nf_log.5 = NONE
net.netfilter.nf_log.6 = NONE
net.netfilter.nf_log.7 = NONE
net.netfilter.nf_log.8 = NONE
net.netfilter.nf_log.9 = NONE
net.nf_contrack_max = 262144
net.unix.max_dgram_qlen = 512
vm.admin_reserve_kbytes = 8192
vm.block_dump = 0
vm.compact_unevictable_allowed = 1
vm.dirty_background_bytes = 0
vm.dirty_background_ratio = 10
vm.dirty_bytes = 0
vm.dirty_expire_centisecs = 3000
```

(continues on next page)

(continued from previous page)

```

vm.dirty_ratio = 20
vm.dirty_writeback_centisecs = 500
vm.dirtytime_expire_seconds = 43200
vm.drop_caches = 0
vm.extfrag_threshold = 500
vm.hugepages_treat_as_movable = 0
vm.hugetlb_shm_group = 0
vm.laptop_mode = 0
vm.legacy_va_layout = 0
vm.lowmem_reserve_ratio = 256      256      32      1
vm.max_map_count = 200000
vm.memory_failure_early_kill = 0
vm.memory_failure_recovery = 1
vm.min_free_kbytes = 90112
vm.min_slab_ratio = 5
vm.min_unmapped_ratio = 1
vm.mmap_min_addr = 65536
vm.nr_hugepages = 4096
vm.nr_hugepages_mempolicy = 4096
vm.nr_overcommit_hugepages = 0
vm.nr_pdflush_threads = 0
vm.numa_zonelist_order = default
vm.oom_dump_tasks = 1
vm.oom_kill_allocating_task = 0
vm.overcommit_kbytes = 0
vm.overcommit_memory = 0
vm.overcommit_ratio = 50
vm.page-cluster = 3
vm.panic_on_oom = 0
vm.percpu_pagelist_fraction = 0
vm.stat_interval = 1
vm.swappiness = 0
vm.user_reserve_kbytes = 131072
vm.vfs_cache_pressure = 100
vm.zone_reclaim_mode = 0

```

Services listing

```

$ service --status-all
[ + ] apparmor
[ - ] bootmisc.sh
[ - ] checkfs.sh
[ - ] checkroot-bootclean.sh
[ - ] checkroot.sh
[ + ] console-setup
[ + ] cpufrequtils
[ + ] cron
[ + ] dbus
[ + ] docker
[ + ] ebttables
[ + ] grub-common
[ - ] hostname.sh
[ - ] hwclock.sh
[ + ] keyboard-setup
[ - ] killprocs
[ + ] kmod
[ + ] loadcpufreq
[ + ] lxcfs
[ - ] mountall-bootclean.sh
[ - ] mountall.sh
[ - ] mountdevsubfs.sh

```

(continues on next page)

(continued from previous page)

```

[ - ] mountkernfs.sh
[ - ] mountnfs-bootclean.sh
[ - ] mountnfs.sh
[ + ] networking
[ - ] ondemand
[ - ] plymouth
[ - ] plymouth-log
[ + ] procps
[ + ] qemu-kvm
[ - ] rc.local
[ + ] resolvconf
[ - ] rsync
[ + ] rsyslog
[ - ] screen-cleanup
[ - ] sendsigs
[ + ] ssh
[ + ] udev
[ - ] umountfs
[ - ] umountnfs.sh
[ - ] umountroot
[ + ] urandom
[ - ] uuidd
[ - ] x11-common

```

Host IRQ affinity

Changing the default pinning of every IRQ to core 0. (Same does apply on both guest VM and host OS)

```
$ for l in `ls /proc/irq`; do echo 1 | sudo tee /proc/irq/$l/smp_affinity; done
```

Host RCU affinity

Changing the default pinning of RCU to core 0. (Same does apply on both guest VM and host OS)

```
$ for i in `pgrep rcu[^c]`; do sudo taskset -pc 0 $i ; done
```

Host Writeback affinity

Changing the default pinning of writebacks to core 0. (Same does apply on both guest VM and host OS)

```
$ echo 1 | sudo tee /sys/bus/workqueue/devices/writeback/cpumask
```

3.5.4 DUT Configuration - DPDK**DPDK Version**

DPDK 18.02

DPDK Compile Parameters

```
make install T=x86_64-native-linuxapp-gcc -j
```

Testpmd Startup Configuration

Testpmd startup configuration changes per test case with different settings for CPU cores, rx-queues. Startup config is aligned with applied test case tag:

Tagged by **1T1C**

```
testpmd -c 0x3 -n 4 -- --numa --nb-ports=2 --portmask=0x3 --nb-cores=1 --max-pkt-len=9000 --
↪ txqflags=0 --forward-mode=io --rxq=1 --txq=1 --burst=64 --rxd=1024 --txd=1024 --disable-link-
↪ check --auto-start
```

Tagged by 2T2C

```
testpmd -c 0x403 -n 4 -- --numa --nb-ports=2 --portmask=0x3 --nb-cores=2 --max-pkt-len=9000 --
↳ txqflags=0 --forward-mode=io --rxq=1 --txq=1 --burst=64 --rxd=1024 --txd=1024 --disable-link-
↳ check --auto-start
```

Tagged by 4T4C

```
testpmd -c 0xc07 -n 4 -- --numa --nb-ports=2 --portmask=0x3 --nb-cores=4 --max-pkt-len=9000 --
↳ txqflags=0 --forward-mode=io --rxq=2 --txq=2 --burst=64 --rxd=1024 --txd=1024 --disable-link-
↳ check --auto-start
```

L3FWD Startup Configuration

L3FWD startup configuration changes per test case with different settings for CPU cores, rx-queues. Startup config is aligned with applied test case tag:

Tagged by 1T1C

```
l3fwd -l 1 -n 4 -- -P -L -p 0x3 --config='${port_config}' --enable-jumbo --max-pkt-len=9000 --eth-
↳ dest=0,${adj_mac0} --eth-dest=1,${adj_mac1} --parse-ptype
```

Tagged by 2T2C

```
l3fwd -l 1,2 -n 4 -- -P -L -p 0x3 --config='${port_config}' --enable-jumbo --max-pkt-len=9000 --eth-
↳ dest=0,${adj_mac0} --eth-dest=1,${adj_mac1} --parse-ptype
```

Tagged by 4T4C

```
l3fwd -l 1,2,3,4 -n 4 -- -P -L -p 0x3 --config='${port_config}' --enable-jumbo --max-pkt-len=9000 --
↳ eth-dest=0,${adj_mac0} --eth-dest=1,${adj_mac1} --parse-ptype
```

3.5.5 TG Configuration - TRex

TG Version

TRex v2.35

DPDK version

DPDK v17.11

TG Build Script used

TRex intallation¹⁷⁷

TG Startup Configuration

```
$ cat /etc/trex_cfg.yaml
- port_limit      : 2
  version        : 2
  interfaces     : ["0000:0d:00.0", "0000:0d:00.1"]
  port_info      :
    - dest_mac   : [0x3c,0xfd,0xfe,0x9c,0xee,0xf5]
      src_mac    : [0x3c,0xfd,0xfe,0x9c,0xee,0xf4]
    - dest_mac   : [0x3c,0xfd,0xfe,0x9c,0xee,0xf4]
      src_mac    : [0x3c,0xfd,0xfe,0x9c,0xee,0xf5]
```

TG Startup Command

```
$ sh -c 'cd <t-rex-install-dir>/scripts/ && sudo nohup ./t-rex-64 -i -c 7 --iom 0 > /tmp/trex.log 2>
↳ &1 &' > /dev/null
```

¹⁷⁷ https://git.fd.io/csit/tree/resources/tools/trex/trex_installer.sh?h=rls1804

TG common API - pointer to driver

Trex driver¹⁷⁸

3.6 Documentation

CSIT DPDK Performance Tests Documentation¹⁷⁹ contains detailed functional description and input parameters for each test case.

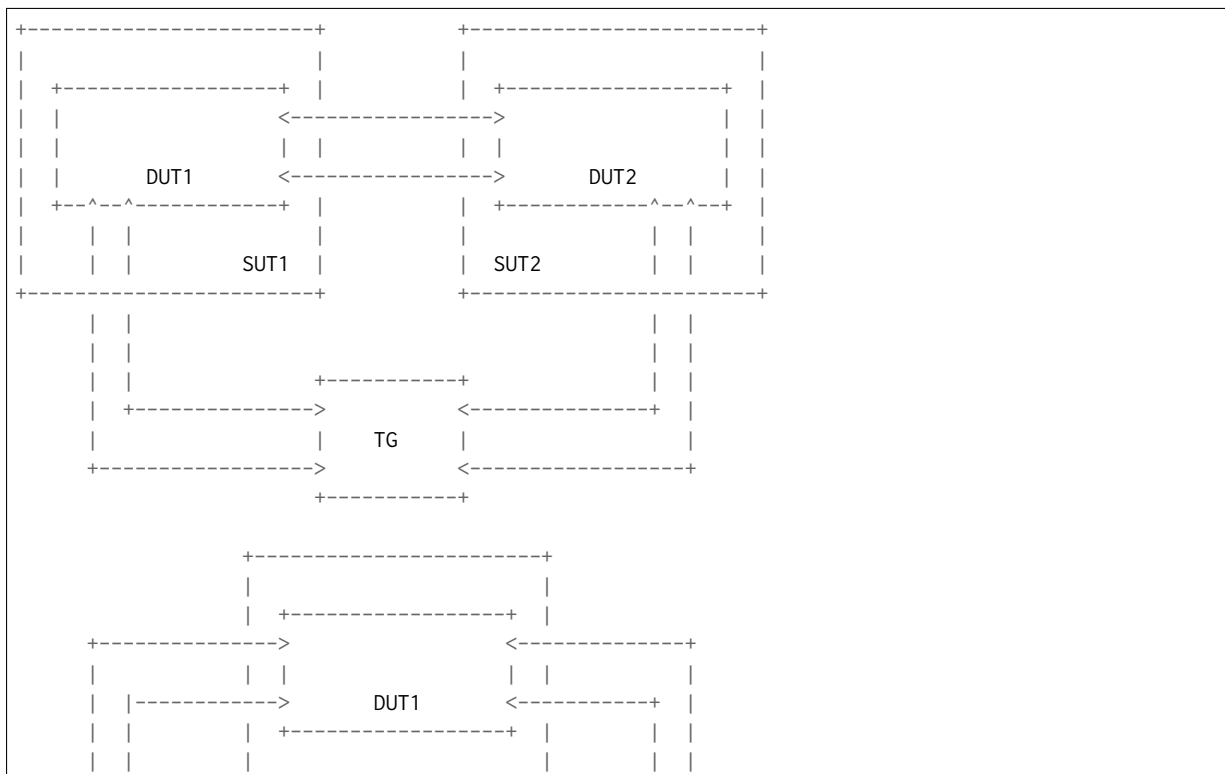
¹⁷⁸ https://git.fd.io/csit/tree/resources/tools/trex/trex_stateless_profile.py?h=rls1804

¹⁷⁹ <https://docs.fd.io/csit/rls1804/doc/tests.dpkg.perf.html>

4.1 Overview

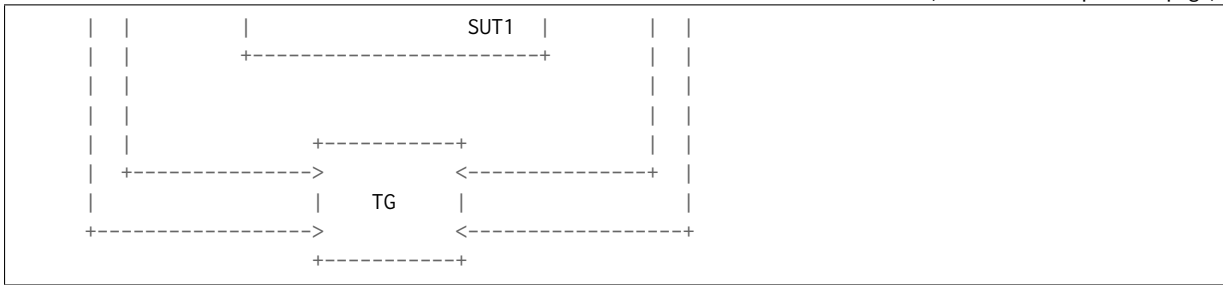
4.1.1 Tested Virtual Topologies

CSIT VPP functional tests are executed on virtualized topologies created using VIRT (Virtual Internet Routing Lab) simulation platform contributed by Cisco. VIRT runs on physical baremetal servers hosted by LF FD.io project. Majority of the tests are executed in the three node logical test topology - Traffic Generator (TG) node and two Systems Under Test (SUT) nodes connected in a loop. Some tests use two node logical test topology - TG node and SUT1 node. Both logical test topologies are shown in the figures below.:



(continues on next page)


(continued from previous page)



SUT1 and SUT2 are two VMs (Ubuntu or Centos, depending on the test suite), TG is a Traffic Generator (TG, another Ubuntu VM). SUTs run VPP SW application in Linux user-mode as a Device Under Test (DUT) within the VM. TG runs Scapy SW application as a packet Traffic Generator. Logical connectivity between SUTs and to TG is provided using virtual NICs using VMs' virtio driver.

Virtual testbeds are created on-demand whenever a verification job is started (e.g. triggered by the Gerrit patch submission) and destroyed upon completion of all functional tests. Each node is a Virtual Machine and each connection that is drawn on the diagram is available for use in any test case. During the test execution, all nodes are reachable thru the Management network connected to every node via dedicated virtual NICs and virtual links (not shown above for clarity).

For the test cases that require DUT (VPP) to communicate with VM over the vhost-user interfaces, a nested VM is created on SUT1 and/or SUT2 for the duration of these particular test cases only. DUT (VPP) test topology with VM is shown in the figure below including the applicable packet flow thru the VM (marked in the figure with ***):



4.1.2 Functional Tests Coverage

Following VPP functional test areas are covered in the CSIT rls1804 with results listed in this report:

- **DHCP - Client and Proxy** - Dynamic Host Control Protocol Client and Proxy for IPv4, IPv6.
- **GRE Overlay Tunnels** - Generic Routing Encapsulation for IPv4.
- **L2BD Ethernet Switching** - L2 Bridge-Domain switched-forwarding for untagged Ethernet, dot1q and dot1ad tagged.
- **L2XC Ethernet Switching** - L2 Cross-Connect switched-forwarding for untagged Ethernet, dot1q and dot1ad tagged.

- **LISP Overlay Tunnels** - Locator/ID Separation Protocol overlay tunnels and locator/id mapping control.
- **Softwire Tunnels** - IPv4-in-IPv6 softwire tunnels.
- **Cop Address Security** - address white-list and black-list filtering for IPv4, IPv6.
- **IPSec - Tunnels and Transport** - IPSec tunnel and transport modes.
- **IPv6 Routed-Forwarding** - IPv6 routed-forwarding, NS/ND, RA, ICMPv6.
- **uRPF Source Security** - unicast Reverse Path Forwarding security.
- **Tap Interface** - baseline Linux tap interface tests.
- **Telemetry - IPFIX and SPAN** - IPFIX netflow statistics and SPAN port mirroring.
- **VRF Routed-Forwarding** - multi-context IPVPN routed-forwarding for IPv4, IPv6.
- **iACL Security** - ingress Access Control List security for IPv4, IPv6, MAC.
- **IPv4 Routed-Forwarding** - IPv4 routed-forwarding, RPF, ARP, Proxy ARP, ICMPv4.
- **QoS Policer Metering** - ingress packet rate measuring and marking for IPv4, IPv6.
- **VLAN Tag Translation** - L2 VLAN tag translation 2to2, 2to1, 1to2, 1to1.
- **VXLAN Overlay Tunnels** - VXLAN tunneling for L2-over-IP, for IPv4, IPv6.

4.1.3 Functional Tests Naming

CSIT rls1804 follows a common structured naming convention for all performance and system functional tests, introduced in CSIT rls1701.

The naming should be intuitive for majority of the tests. Complete description of CSIT test naming convention is provided on *CSIT Test Naming* (page 261)..

Here few illustrative examples of the new naming usage for functional test suites:

1. **Physical port to physical port - a.k.a. NIC-to-NIC, Phy-to-Phy, P2P**
 - *eth2p-ethip4-ip4base-func.robot* => 2 ports of Ethernet, IPv4 baseline routed forwarding, functional tests.
2. **Physical port to VM (or VM chain) to physical port - a.k.a. NIC2VM2NIC, P2V2P, NIC2VMchain2NIC, P2V2V2P**
 - *eth2p-ethip4vxlan-l2bdbasemaclrn-eth-2vhost-1vm-func.robot* => 2 ports of Ethernet, IPv4 VXLAN Ethernet, L2 bridge-domain switching to/from two vhost interfaces and one VM, functional tests.

4.2 CSIT Release Notes

4.2.1 Changes in CSIT rls1804

1. VPP functional test framework changes:
 - functional tests are running on Centos image; corresponding CSIT Jenkins jobs are created.
2. Optimizations in CSIT functional framework:
 - *VPP install test* - VPP installation has been moved to a separate test in test suite setup phase to better identify any issues with VPP installation.
 - *VPP verify test* - test to check VPP health after installation.

4.2.2 Known Issues

List of known issues in CSIT rls1804 for VPP functional tests in VIRL:

| # | Issue | Jira ID | Description |
|---|--|------------------|--|
| 1 | DHCPv4 client: Client responses to DHCPv4 OFFER sent with different XID. | CSIT-129 VPP-99 | Client replies with DHCPv4 REQUEST message when received DHCPv4 OFFER message with different (wrong) XID. |
| 2 | Softwire - MAP-E: Incorrect calculation of IPv6 destination address when IPv4 prefix is 0. | CSIT-398 VPP-380 | IPv6 destination address is wrongly calculated in case that IPv4 prefix is equal to 0 and IPv6 prefix is less than 40. |
| 3 | Softwire - MAP-E: Map domain is created when incorrect parameters provided. | CSIT-399 VPP-435 | Map domain is created in case that the sum of suffix length of IPv4 prefix and PSID length is greater than EA bits length. IPv6 destination address contains bits written with PSID over the EA-bit length when IPv4 packet is sent. |
| 4 | IPv6 RA: Incorrect IPv6 destination address in response to ICMPv6 Router Solicitation. | CSIT-409 VPP-406 | Wrong IPv6 destination address (ff02::1) is used in ICMPv6 Router Advertisement packet sent as a response to received ICMPv6 Router Solicitation packet. |
| 5 | Vhost-user: QEMU reconnect does not work. | CSIT-565 | Used QEMU 2.5.0 does not support vhost-user reconnect. Requires upgrading CSIT VIRL environment to QEMU 2.7.0. |

4.3 Test Environment

CSIT functional tests are currently executed in FD.IO VIRL testbed. The physical VIRL testbed infrastructure consists of three VIRL hosts:

- All hosts are Cisco UCS C240-M4 (2x Intel(R) Xeon(R) CPU E5-2699 v3 @2.30GHz, 18c, 512GB RAM)
- tb4-vir1:
 - Status: Production
 - OS: Ubuntu 16.04.2

- STD server version 0.10.32.16
- UWM server version 0.10.32.16
- tb4-virl2:
 - Status: Production
 - OS: Ubuntu 16.04.2
 - STD server version 0.10.32.16
 - UWM server version 0.10.32.16
- tb4-virl3:
 - Status: Testing
 - OS: Ubuntu 16.04.2
 - STD server version 0.10.32.19
 - UWM server version 0.10.32.19

Whenever a patch is submitted to gerrit for review, parallel VIRT simulations are started to reduce the time of execution of all functional tests. The number of parallel VIRT simulations is equal to number of test groups defined by TEST_GROUPS variable in `csit/bootstrap.sh` file. The VIRT host to run VIRT simulation is selected based on least load algorithm per VIRT simulation.

Every VIRT simulation uses the same three-node - Traffic Generator (TG node) and two Systems Under Test (SUT1 and SUT2) - “double-ring” topology. The appropriate pre-built VPP packages built by Jenkins for the patch under review are then installed on the two SUTs, along with their `/etc/vpp/startup.conf` file, in all VIRT simulations.

4.3.1 SUT Configuration - VIRT Guest VM

Configurations of the SUT VMs is defined in [VIRT topologies directory](#)¹⁸⁰

- List of SUT VM interfaces::

```
<interface id="0" name="GigabitEthernet0/4/0"/>
<interface id="1" name="GigabitEthernet0/5/0"/>
<interface id="2" name="GigabitEthernet0/6/0"/>
<interface id="3" name="GigabitEthernet0/7/0"/>
```

- Number of 2MB hugepages: 1024
- Maximum number of memory map areas: 20000
- Kernel Shared Memory Max: 2147483648 ($vm.nr_hugepages * 2 * 1024 * 1024$)

4.3.2 SUT Configuration - VIRT Guest OS Linux

In CSIT terminology, the VM operating system for both SUTs that VPP-18.04 release has been tested with, is the following:

1. Ubuntu VIRT image

This image implies Ubuntu 16.04.1 LTS, current as of yyyy-mm-dd (that is, package versions are those that would have been installed by a `apt-get update`, `apt-get upgrade` on that day), produced by CSIT disk image build scripts.

The exact list of installed packages and their versions (including the Linux kernel package version) are included in [VIRT ubuntu images lists](#)¹⁸¹.

¹⁸⁰ <https://git.fd.io/csit/tree/resources/tools/virt/topologies/?h=rls1804>

¹⁸¹ <https://git.fd.io/csit/tree/resources/tools/disk-image-builder/ubuntu/lists/?h=rls1804>

A replica of this VM image can be built by running the `build.sh` script in CSIT repository.

2. CentOS VIRT image

This image implies Centos 7.4-1711, current as of yyyy-mm-dd (that is, package versions are those that would have been installed by a `yum update`, `yum upgrade` on that day), produced by CSIT disk image build scripts.

The exact list of installed packages and their versions (including the Linux kernel package version) are included in [VIRT centos images lists](#)¹⁸².

A replica of this VM image can be built by running the `build.sh` script in CSIT repository.

3. Nested VM image

In addition to the “main” VM image, tests which require VPP to communicate to a VM over a vhost-user interface, utilize a “nested” VM image.

This “nested” VM is dynamically created and destroyed as part of a test case, and therefore the “nested” VM image is optimized to be small, lightweight and have a short boot time. The “nested” VM image is not built around any established Linux distribution, but is based on [BuildRoot](#)¹⁸³, a tool for building embedded Linux systems. Just as for the “main” image, scripts to produce an identical replica of the “nested” image are included in CSIT GIT repository, and the image can be rebuilt using the “build.sh” script at [VIRT nested](#)¹⁸⁴.

4.3.3 DUT Configuration - VPP

Every System Under Test runs VPP SW application in Linux user-mode as a Device Under Test (DUT) node.

DUT port configuration

Port configuration of DUTs is defined in topology file that is generated per VIRT simulation based on the definition stored in [VIRT topologies directory](#)¹⁸⁵.

Example of DUT nodes configuration::

```
DUT1:
  type: DUT
  host: "10.30.51.157"
  port: 22
  username: cisco
  honeycomb:
    user: admin
    passwd: admin
    port: 8183
    netconf_port: 2831
  priv_key: |
    -----BEGIN RSA PRIVATE KEY-----
    MIIIEpgIBAACKAQEAwUD1TpzSHpwLQotZOFs4AgcPNEwCnP1AB2hWfMvI+8Kah/gb
    v8ruZU9RqhPs56tyKzxbhvNkY4VbH5F1Gi1HZu3mLqz4KfghMmaeMEj01T7BYyd
    vuBfTvIlu1jFQ2vAlnYrDwn+C1xJk81m0pDgvrLEX4qVvH2sGh7UEkYy5r82DNa2
    4VjzPB1J/c8a9zP8FoZUhYIzF4FLvRMjUADpbMXgJMsGpaZLmz95ap0Eot7vb1Cc
    1LvF97iyBCrtIOSKRKA50ZHLGjMKmOwnYU+cP5718tbproDVi6VJOo7zeuXyetMs
    8YB19kwb1WG9BqP9jctFvsmi5G7hXgq1Y8u+DwIDAQABAoIBAQC/W4E0DHjLMny7
    0bvW2YKzD0Zw3fttdB94tkm4PdZv5MybooPnsAvLaXVv0hEdfVi5kzSWN1/LY/tN
    EP1BgGphc2QgB59/PPxGwFIjDCvUz1sZpynBHe+B/qh5ExNQcVvsIQqWI7DX1XaN
    0i/khOzmJ6HncRRah1spKimYRsaUUDskyg7q3QqMwVaqBbbMvLs/w7ZWd/zoDqCU
    MY/pCI6hkB3QbRo0odiZLohphB12ShABTwjvVyyKL5UA4jAEneJrhH5gWVLXnfgD
    p62W5Col1KEYb1C8mUkPxpP7Qo277zw3xaq+oktIzhc5SUEUd7nJZtNqVAHqkItW
    79VmpKyxAoGBAPfU+kqNPATsvp+x1n5sn2SgipzDgtgi9QqNmC4cjtrQQaaqI57SG
```

(continues on next page)

¹⁸² <https://git.fd.io/csit/tree/resources/tools/disk-image-builder/centos/lists/?h=rls1804>

¹⁸³ <https://buildroot.org/>

¹⁸⁴ <https://git.fd.io/csit/tree/resources/tools/disk-image-builder/nested/?h=rls1804>

¹⁸⁵ <https://git.fd.io/csit/tree/resources/tools/virt/topologies/?h=rls1804>

(continued from previous page)

```

OHw1jX8i7L2G1WvVtkHg060n1EVo5n65fff0qeVBezLVJ7ghWI8U+oBiJJyQ4boD
GJVNs0oSUQ0rtuGd9eVwfdk3o19aCN0KK53oPfIYli29pyu4l095kg11AoGBAMef
bPEMBI/2XmCPshLswHGF1+dW8d+K1luj3CUQ/0vU1vma3dfBOYNsIwAgTP0iIUTg
8DYE6KBCdPtXAUEI0YAEAKB9ry1tKR2NQEIpfS1YytKErtwjAiqSi0heM6+zwEzu
f54Z4oBhMSL0jXoMnu+NZZec6EUdQeY40+jhjzAoGBAIOgC3dtjMPGKTP7+93u
UE/XIioI8fWg9fj3SMka4IMu+pVvRCRbAjRH7JrFLkjbUyuMqs3Arnk9K+gbdQt/
+m95Njtt6WoFXuPCwgbM3GidSmZwYT4454SfDzVBYScedcNm1FuR+8ov9bFLDtGT
D4gsngnGJj1MDFXTxZEn4nzZAoGBAKCg4WmpUPaCuXibyB+rZavxwTNSn2lJ83/
sYJGBhf/raiV/FLDUcM1vYg5dZnu37RsB/5/vqx0LZGyYd7x+Jo5HkQGPnKgNwhn
g8BkdZIRF8uEJqx0o0ycd0U7n/2093swIpKWo5LIiRPuqqzj+uZKnAL7vuVdxfaY
qVz2daMPAoGBALgaaKa3voU/HO1PYLWIhFrBThyJ+BQSQ80qrEzC8AnegWFxRAM8
EqrzZX17ACUuo1dH0Eipm41j2+BZWlQjiUgq5uj8+yzy+EU1ZRRyJc0KzbDACeuD
BpWWSXGBI5G4CppeYLjMUHZpJYeX1USULJQd2c4crLJKb76E8gz3Z9kN
-----END RSA PRIVATE KEY-----

```

interfaces:

```

port1:
  mac_address: "fa:16:3e:9b:89:52"
  pci_address: "0000:00:04.0"
  link: link1
port2:
  mac_address: "fa:16:3e:7a:33:60"
  pci_address: "0000:00:05.0"
  link: link4
port3:
  mac_address: "fa:16:3e:29:b7:ae"
  pci_address: "0000:00:06.0"
  link: link3
port4:
  mac_address: "fa:16:3e:76:8d:ff"
  pci_address: "0000:00:07.0"
  link: link6

```

DUT2:

```

type: DUT
host: "10.30.51.156"
port: 22
username: cisco
honeycomb:
  user: admin
  passwd: admin
  port: 8183
  netconf_port: 2831
priv_key: |

```

```

-----BEGIN RSA PRIVATE KEY-----
MIIePgIBAAKCAQEAWUDlTpzShpwLQotZ0FS4AgcPNEwCnP1AB2hWfMvI+8Kah/gb
v8ruZU9RqhPs56tyKzxbhvNkY4VbH5F1GillHzu3mLqzM4KfghMmaeMEj01T7BYyd
vuBftvIlu1jFQ2vAlnYrDwn+C1xJk81m0pDgvrLEX4qVvH2sGh7UEkYy5r82DNa2
4VjzPB1J/c8a9zP8FoZuHYIzF4FLvRMjUADpbMXgJMsGpaZLmz95ap0Eot7vb1Cc
1LvF97iyBCrtIOSKRKA50ZHLGjMKmOwnYU+cP5718tbproDVi6VJ0o7zeuXyEtMs
8YB19kwb1WG9BqP9jctFvsmi5G7hXgq1Y8u+DwIDAQABAoIBAQC/W4E0DHjLMny7
0bvW2YKzD0Zw3fttdB94tkm4PdZv5MybooPnsAvLaXVv0hEdfVi5kzSWN1/LY/tN
EP1BgGphc2QgB59/PPxGwFIjDCvUz1sZpynBHe+B/qh5ExNqCvVsIQWI7DX1XaN
0i/kh0zmJ6HncRRah1spKimYRsaUUDskyg7q3QqMwVaqBbbMvLs/w7Zwd/zoDqCU
MY/pCI6hkB3QbRo0diZLohphB12ShABTjvVyyKL5UA4jAeneJrhH5gWVWLnfgD
p62W5CollKEYb1C8mUkPxp7Qo277zw3xaq+oktIzhc5SUEud7nJZtNqVAHqkItW
79VmpKyxAoGBAPfU+kqNPaTSvp+x1n5sn2SgipzDgtgi9QqNmC4cjtrQaaqI57SG
OHw1jX8i7L2G1WvVtkHg060n1EVo5n65fff0qeVBezLVJ7ghWI8U+oBiJJyQ4boD
GJVNs0oSUQ0rtuGd9eVwfdk3o19aCN0KK53oPfIYli29pyu4l095kg11AoGBAMef
bPEMBI/2XmCPshLswHGF1+dW8d+K1luj3CUQ/0vU1vma3dfBOYNsIwAgTP0iIUTg
8DYE6KBCdPtXAUEI0YAEAKB9ry1tKR2NQEIpfS1YytKErtwjAiqSi0heM6+zwEzu
f54Z4oBhMSL0jXoMnu+NZZec6EUdQeY40+jhjzAoGBAIOgC3dtjMPGKTP7+93u

```

(continues on next page)

(continued from previous page)

```

UE/XIioI8fWg9fj3sMka4IMu+pVvRCRbAjRH7JrFLkjbUyuMqs3Arnk9K+gbdQt/
+m95Njtt6WoFXuPCwgbM3GidSmZwYT4454SfDzVBYScedCNm1FuR+8ov9bFLDtGT
D4gsngnGJj1MDFXTxZEn4nzZAoGBAKCg4WmpUPaCuXibyB+rZavxwsTNSn2lJ83/
sYJGBhf/raiV/FLDUcM1vYg5dZnu37RsB/5/vqxOLZGyYd7x+Jo5HkQGPnKgNwhn
g8BkdZIRF8uEJqx0o0ycd0U7n/2093swIpKwo5LIiRPuqzj+uZKnAL7vuVdxfaY
qVz2daMPAoGBALgaaKa3voU/H01PYLWIhFrBThyJ+BQSQ80qrEzC8AnegWFxRAM8
EqrzZXl7ACUuo1dH0Eipm41j2+BZWlQjiUgq5uj8+yzy+EU1ZRRyJcOKzbDACeuD
BpWWSXGBI5G4CppeYLjMUHZpJYeX1USULJQd2c4crLJKb76E8gz3Z9kN
-----END RSA PRIVATE KEY-----

```

interfaces:

```

port1:
  mac_address: "fa:16:3e:ad:6c:7d"
  pci_address: "0000:00:04.0"
  link: link2
port2:
  mac_address: "fa:16:3e:94:a4:99"
  pci_address: "0000:00:05.0"
  link: link5
port3:
  mac_address: "fa:16:3e:75:92:da"
  pci_address: "0000:00:06.0"
  link: link3
port4:
  mac_address: "fa:16:3e:2c:b1:2a"
  pci_address: "0000:00:07.0"
  link: link6

```

VPP Version

VPP-18.04 release

VPP Installed Packages - Ubuntu

```

$ dpkg -l vpp\*
Desired=Unknown/Install/Remove/Purge/Hold
| Status=Not/Inst/Conf-files/Unpacked/half-f-inst/trig-aWait/Trig-pend
|/ Err?=(none)/Reinst-required (Status,Err: uppercase=bad)
||/ Name          Version          Architecture Description
+++-----
ii vpp              18.04-release   amd64         Vector Packet Processing--executables
ii vpp-dbg         18.04-release   amd64         Vector Packet Processing--debug symbols
ii vpp-dev         18.04-release   amd64         Vector Packet Processing--development support
ii vpp-dpdk-dkms  18.02-vpp1      amd64         DPDK Development Package for VPP - Kernel Modules
ii vpp-lib         18.04-release   amd64         Vector Packet Processing--runtime libraries
ii vpp-plugins    18.04-release   amd64         Vector Packet Processing--runtime plugins

```

VPP Installed Packages - Centos

```

$ rpm -qai vpp*
Name       : vpp
Version    : 18.04
Release    : release
Architecture: x86_64
Install Date: Fri 04 May 2018 02:47:36 AM EDT
Group      : Unspecified
Size       : 2311324
License    : ASL 2.0
Signature  : (none)
Source RPM : vpp-18.04-release.src.rpm
Build Date : Wed 25 Apr 2018 01:00:24 PM EDT
Build Host : c29a51c12e88

```

(continues on next page)

(continued from previous page)

```
Relocations : (not relocatable)
Summary     : Vector Packet Processing
Description :
This package provides VPP executables: vpp, vpp_api_test, vpp_json_test
vpp - the vector packet engine
vpp_api_test - vector packet engine API test tool
vpp_json_test - vector packet engine JSON test tool
Name        : vpp-lib
Version     : 18.04
Release    : release
Architecture: x86_64
Install Date: Fri 04 May 2018 02:47:36 AM EDT
Group       : System Environment/Libraries
Size        : 24949034
License     : ASL 2.0
Signature   : (none)
Source RPM  : vpp-18.04-release.src.rpm
Build Date  : Wed 25 Apr 2018 01:00:24 PM EDT
Build Host  : c29a51c12e88
Relocations : (not relocatable)
Summary     : VPP libraries
Description :
This package contains the VPP shared libraries, including:
vppinfra - foundation library supporting vectors, hashes, bitmaps, pools, and string formatting.
svm - vm library
vlib - vector processing library
vlib-api - binary API library
vnet - network stack library
Name        : vpp-selinux-policy
Version     : 18.04
Release    : release
Architecture: x86_64
Install Date: Fri 04 May 2018 02:47:35 AM EDT
Group       : System Environment/Base
Size        : 86170
License     : ASL 2.0
Signature   : (none)
Source RPM  : vpp-18.04-release.src.rpm
Build Date  : Wed 25 Apr 2018 01:00:24 PM EDT
Build Host  : c29a51c12e88
Relocations : (not relocatable)
Summary     : VPP Security-Enhanced Linux (SELinux) policy
Description :
This package contains a tailored VPP SELinux policy
Name        : vpp-devel
Version     : 18.04
Release    : release
Architecture: x86_64
Install Date: Fri 04 May 2018 02:47:38 AM EDT
Group       : Development/Libraries
Size        : 10632218
License     : ASL 2.0
Signature   : (none)
Source RPM  : vpp-18.04-release.src.rpm
Build Date  : Wed 25 Apr 2018 01:00:24 PM EDT
Build Host  : c29a51c12e88
Relocations : (not relocatable)
Summary     : VPP header files, static libraries
Description :
This package contains the header files for VPP.
Install this package if you want to write a
```

(continues on next page)

(continued from previous page)

```

program for compilation and linking with vpp lib.
vlib
vlibmemory
vnet - devices, classify, dhcp, ethernet flow, gre, ip, etc.
vpp-api
vppinfra
Name       : vpp-plugins
Version    : 18.04
Release    : release
Architecture: x86_64
Install Date: Fri 04 May 2018 02:47:37 AM EDT
Group      : System Environment/Libraries
Size       : 44145968
License    : ASL 2.0
Signature  : (none)
Source RPM : vpp-18.04-release.src.rpm
Build Date : Wed 25 Apr 2018 01:00:24 PM EDT
Build Host : c29a51c12e88
Relocations : (not relocatable)
Summary    : Vector Packet Processing--runtime plugins
Description :
This package contains VPP plugins

```

VPP Startup Configuration

VPP startup configuration is common for all test cases except test cases related to SW Crypto device.

Default

```

$ cat /etc/vpp/startup.conf
unix {
    nodaemon
    log /var/log/vpp/vpp.log
    full-coredump
    cli-listen /run/vpp/cli.sock
    gid vpp
}

api-trace {
    ## This stanza controls binary API tracing. Unless there is a very strong reason,
    ## please leave this feature enabled.
    on
    ## Additional parameters:
    ##
    ## To set the number of binary API trace records in the circular buffer, configure nitems
    ##
    ## nitems <nnn>
    ##
    ## To save the api message table decode tables, configure a filename. Results in /tmp/<filename>
    ## Very handy for understanding api message changes between versions, identifying missing
    ## plugins, and so forth.
    ##
    ## save-api-table <filename>
}

api-segment {
    gid vpp
}

cpu {
    ## In the VPP there is one main thread and optionally the user can create worker(s)
    ## The main thread and worker thread(s) can be pinned to CPU core(s) manually or automatically

```

(continues on next page)

(continued from previous page)

```
## Manual pinning of thread(s) to CPU core(s)

## Set logical CPU core where main thread runs
# main-core 1

## Set logical CPU core(s) where worker threads are running
# corelist-workers 2-3,18-19

## Automatic pinning of thread(s) to CPU core(s)

## Sets number of CPU core(s) to be skipped (1 ... N-1)
## Skipped CPU core(s) are not used for pinning main thread and working thread(s).
## The main thread is automatically pinned to the first available CPU core and worker(s)
## are pinned to next free CPU core(s) after core assigned to main thread
# skip-cores 4

## Specify a number of workers to be created
## Workers are pinned to N consecutive CPU cores while skipping "skip-cores" CPU core(s)
## and main thread's CPU core
# workers 2

## Set scheduling policy and priority of main and worker threads

## Scheduling policy options are: other (SCHED_OTHER), batch (SCHED_BATCH)
## idle (SCHED_IDLE), fifo (SCHED_FIFO), rr (SCHED_RR)
# scheduler-policy fifo

## Scheduling priority is used only for "real-time policies (fifo and rr),
## and has to be in the range of priorities supported for a particular policy
# scheduler-priority 50
}

# dpdk {
## Change default settings for all interfaces
# dev default {
## Number of receive queues, enables RSS
## Default is 1
# num-rx-queues 3

## Number of transmit queues, Default is equal
## to number of worker threads or 1 if no workers threads
# num-tx-queues 3

## Number of descriptors in transmit and receive rings
## increasing or reducing number can impact performance
## Default is 1024 for both rx and tx
# num-rx-desc 512
# num-tx-desc 512

## VLAN strip offload mode for interface
## Default is off
# vlan-strip-offload on
# }

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

## Whitelist specific interface by specifying PCI address and in
## addition specify custom parameters for this interface
# dev 0000:02:00.1 {
```

(continues on next page)

(continued from previous page)

```

#     num-rx-queues 2
# }

## Specify bonded interface and its slaves via PCI addresses
##
## Bonded interface in XOR load balance mode (mode 2) with L3 and L4 headers
# vdev eth_bond0,mode=2,slave=0000:02:00.0,slave=0000:03:00.0,xmit_policy=l34
# vdev eth_bond1,mode=2,slave=0000:02:00.1,slave=0000:03:00.1,xmit_policy=l34
##
## Bonded interface in Active-Back up mode (mode 1)
# vdev eth_bond0,mode=1,slave=0000:02:00.0,slave=0000:03:00.0
# vdev eth_bond1,mode=1,slave=0000:02:00.1,slave=0000:03:00.1

## Change UIO driver used by VPP, Options are: igb_uio, vfio-pci,
## uio_pci_generic or auto (default)
# uio-driver vfio-pci

## Disable mutli-segment buffers, improves performance but
## disables Jumbo MTU support
# no-multi-seg

## Increase number of buffers allocated, needed only in scenarios with
## large number of interfaces and worker threads. Value is per CPU socket.
## Default is 16384
# num-mbufs 128000

## Change hugepages allocation per-socket, needed only if there is need for
## larger number of mbufs. Default is 256M on each detected CPU socket
# socket-mem 2048,2048

## Disables UDP / TCP TX checksum offload. Typically needed for use
## faster vector PMDs (together with no-multi-seg)
# no-tx-checksum-offload
# }

# plugins {
## Adjusting the plugin path depending on where the VPP plugins are
#     path /home/bms/vpp/build-root/install-vpp-native/vpp/lib64/vpp_plugins

## Disable all plugins by default and then selectively enable specific plugins
# plugin default { disable }
# plugin dpdk_plugin.so { enable }
# plugin acl_plugin.so { enable }

## Enable all plugins by default and then selectively disable specific plugins
# plugin dpdk_plugin.so { disable }
# plugin acl_plugin.so { disable }
# }

## Alternate syntax to choose plugin path
# plugin_path /home/bms/vpp/build-root/install-vpp-native/vpp/lib64/vpp_plugins

```

SW Crypto Device

```

$ cat /etc/vpp/startup.conf
unix
{
  cli-listen /run/vpp/cli.sock
  gid vpp
  nodaemon

```

(continues on next page)

(continued from previous page)

```

full-coredump
log /tmp/vpp.log
}
api-segment
{
  gid vpp
}
dpdk
{
  vdev cryptodev_aesni_gcm_pmd, socket_id=0
  vdev cryptodev_aesni_mb_pmd, socket_id=0
}

```

4.3.4 TG Configuration

Traffic Generator node is VM running the same OS Linux as SUTs. Ports of this VM are used as source (Tx) and destination (Rx) ports for the traffic.

Traffic scripts of test cases are executed on this VM.

TG VM configuration

Configuration of the TG VMs is defined in [VIRL topologies directory](#)¹⁸⁶.

```
/csit/resources/tools/virl/topologies/double-ring-nested.xenial.virl
```

- List of TG VM interfaces::

```

<interface id="0" name="eth1"/>
<interface id="1" name="eth2"/>
<interface id="2" name="eth3"/>
<interface id="3" name="eth4"/>
<interface id="4" name="eth5"/>
<interface id="5" name="eth6"/>

```

TG node port configuration

Port configuration of TG is defined in topology file that is generated per VIRL simulation based on the definition stored in [VIRL topologies directory](#)¹⁸⁷.

Example of TG node configuration::

```

TG:
  type: TG
  host: "10.30.51.155"
  port: 22
  username: cisco
  priv_key: |
    -----BEGIN RSA PRIVATE KEY-----
    MIIIEPgIBAAKCAQEAwUD1TpzSHpwLQotZOFs4AgcPNEwCnP1AB2hWFmvI+8Kah/gb
    v8ruZU9RqhPs56tyKzxbhvNkY4VbH5F1Gi1HZu3mLqzM4KfghMmaeMEj01T7BYyd
    vuBfTvI1u1jfQ2vA1nYrDwn+C1xJk81m0pDgvrLEX4qVVh2sGh7UEkYy5r82DNa2
    4VjzPB1J/c8a9zP8FoZUhYIzF4FLvRMjUADpbMXgJMsGpaZLmz95ap0Eot7vb1Cc
    1LvF97iyBCrtIOSKRKA50ZHLGjMKmOwnYU+cP5718tbproDVi6VJ0o7zeuXyetMs
    8YB19kwb1WG9BqP9jctFvsmi5G7hXgq1Y8u+DwIDAQABAoIBAQC/W4E0DHjLMny7
    0bvW2YKzD0Zw3fttdB94tkm4PdZv5MybooPnsAvLaXVV0hEdfVi5kzSWN1/LY/tN
    EP1BgGphc2QgB59/PPxGwFIjDCvUz1sZpynBHe+B/qh5ExNQcVvsIQqWI7DX1XaN
    0i/kh0zmJ6HncRRah1spKimYRsaUUDskyg7q3QqMWVaqBbbMvLs/w7ZWd/zoDqCU
    MY/pCI6hkB3QbRo0diZLohphB12ShABTjvVyyKL5UA4jAENEJrhH5gWVLXnfgD

```

(continues on next page)

¹⁸⁶ <https://git.fd.io/csit/tree/resources/tools/virl/topologies/?h=rls1804>

¹⁸⁷ <https://git.fd.io/csit/tree/resources/tools/virl/topologies/?h=rls1804>

(continued from previous page)

```
p62W5Col1KEYb1C8mUkPxpP7Qo277zw3xaq+oktIZhc5SUEUd7nJZtNqVAHqkITW
79VmpKyxAoGBAPfU+kqNPATSvp+x1n5sn2SgipzDtgI9QqNmC4cjrQQaaqI57SG
OHw1jX8i7L2G1WvVtkHg060n1EVo5n65ffF0qeVBezLVJ7ghWI8U+oBiJJyQ4boD
GJVnso0SUQ0rtuGd9eVwFdk3o19aCN0KK53oPFIY1i29pyu4l095kg11AoGBAMef
bPEMBI/2XmCPshLSwhGF1+dW8d+K1luj3CUQ/0vU1vma3dfBOYNsIwAgTP0iIUTg
8DYE6KBCdPtxAUEI0YAEAKB9ry1tKR2NQEIPfs1YytKErtwjAiqSi0heM6+zwEzu
f54Z4oBhmsML0jXoMnu+NZzEc6EUdQeY40+jhJzAoGBAIOgC3dtjMPGKTP7+93u
UE/XIioI8fWg9fj3sMka4IMu+pVvRCRbAjRH7JrFLkjbUyuMqs3Arnk9K+gbdQt/
+m95Njtt6WoFXuPCwgbM3GidSmZwYT4454SfDzVBYScedCNm1FuR+8ov9bFLDtGT
D4gsngnGJj1MDFXTxZEn4nzZAoGBAKCg4WmpUPaCuXibyB+rZavxwsTNSn2lJ83/
sYJGBhf/raiV/FLDUcM1vYg5dZnu37RsB/5/vqxOLZGyYd7x+Jo5HKQGPnKgNwhn
g8BkdZIRF8uEJqx0o0ycd0U7n/2093swIpKWo5LIiRPuqzj+uZKnAL7vuVdxfaY
qVz2daMPAoGBALgaaKa3voU/H01PYLWIhFrBThyJ+BQSQ80qrEzC8AnegWFxRAM8
EqrzZX17ACUuo1dh0Eipm41j2+BZWlQjiUgq5uj8+zyz+EU1ZRRyJcOKzbDACeuD
BpWWSXGBI5G4CppeYLjMUHZpJYeX1USULJQd2c4crLJKb76E8gz3Z9kN
-----END RSA PRIVATE KEY-----
```

interfaces:

```
port3:
  mac_address: "fa:16:3e:b9:e1:27"
  pci_address: "0000:00:06.0"
  link: link1
  driver: virtio-pci
port4:
  mac_address: "fa:16:3e:e9:c8:68"
  pci_address: "0000:00:07.0"
  link: link4
  driver: virtio-pci
port5:
  mac_address: "fa:16:3e:e8:d3:47"
  pci_address: "0000:00:08.0"
  link: link2
  driver: virtio-pci
port6:
  mac_address: "fa:16:3e:cf:ca:58"
  pci_address: "0000:00:09.0"
  link: link5
  driver: virtio-pci
```

Traffic generator

Functional tests utilize Scapy as a traffic generator. There was used Scapy v2.3.1 for VPP-18.04 release tests.

4.4 Documentation

CSIT VPP Functional Tests Documentation¹⁸⁸ contains detailed functional description and input parameters for each test case.


¹⁸⁸ <https://docs.fd.io/csit/rls1804/doc/tests.vpp.func.html>

Honeycomb Functional Tests

5.1 Overview

5.1.1 Tested Virtual Topologies

CSIT Honeycomb functional tests are executed on virtualized topologies created using VIRT simulation platform contributed by Cisco. VIRT runs on physical baremetal servers hosted by LF FD.io project. All tests are executed in two node logical test topology - Traffic Generator (TG) node and Systems Under Test (SUT1) node connected in a loop. Logical test topology is shown in the figure below.:



SUT1 is a VM (Ubuntu or Centos, depending on the test suite), TG is a Traffic Generator (TG, another Ubuntu VM). SUTs run Honeycomb and VPP SW applications in Linux user-mode as a Device Under Test (DUT) within the VM. TG runs Scapy SW application as a packet Traffic Generator. Logical connectivity between SUTs and to TG is provided using virtual NICs using VMs' virtio driver.

Virtual testbeds are created on-demand whenever a verification job is started (e.g. triggered by the gerrit patch submission) and destroyed upon completion of all functional tests. Each node is a Virtual Machine and each connection that is drawn on the diagram is available for use in any test case. During the test execution, all nodes are reachable thru the Management network connected to every node via dedicated virtual NICs and virtual links (not shown above for clarity).

5.1.2 Functional Tests Coverage

The following Honeycomb functional test areas are included in the CSIT rls1804 with results listed in this report:

- **Basic interface management** - CRUD for interface state, - ipv4/ipv6 address, ipv4 neighbor, MTU value. - Test case count: 14
- **L2BD** - CRUD for L2 Bridge-Domain, interface assignment. - Create up to two bridge domains with all implemented functions turned on. - (flooding, unknown-unicast flooding, forwarding, learning, arp-termination) - Assign up to two physical interfaces to a single bridge domain. - Remove interface assignments, remove bridge domains. - Test case count: 5
- **L2FIB** - CRD for L2-FIB entries. - Create 4 FIB entries - (one of each for filter/forward, static/dynamic combinations). - Remove FIB entries. - Test case count: 7
- **VxLAN** - CRD for VxLAN tunnels. - Create VxLAN interface. - Disable VxLAN interface. - Re-create a disabled VxLAN interface. - Test case count: 6
- **VxLAN-GPE** - CRD for VxLAN GPE tunnels. - Create VxLAN GPE interface. - Disable VxLAN interface. - Re-create a disabled VxLAN interface. - Test case count: 7
- **Vhost-user** - CRUD for Vhost-user interfaces. - Create, modify and delete Vhost-user interface, as client and server. - Test case count: 8
- **TAP** - CRUD for Tap interface management. - Create, modify and delete TAP interface. - Test case count: 3
- **VLAN** - CRUD for VLAN sub-interface management. - Create VLAN sub-interface over a physical interface. - Toggle interface state separately for super-interface and sub-interface. - Configure IP address and bridge domain assignment on sub-interface. - Configure VLAN tag rewrite on sub-interface. - Test case count: 24
- **ACL** - CRD for low-level classifiers: table and session management, - interface assignment. - Configure up to 2 classify tables. - Configure up to 2 classify sessions on one table. - Assign classify session to a physical interface. - Remove tables, sessions, interface assignments. - Test case count: 9
- **PBB** - CRD for provider backbone bridge sub-interface. - Configure, modify and remove a PBB sub-interface over a physical interface. - Test case count: 8
- **NSH_SFC** - CRD for NSH maps and entries, using NSH_SFC plugin. - Configure up to 2 NSH entries. - Configure up to 2 NSH maps. - Modify and delete NSH maps and entries. - Test case count: 8
- **LISP** - CRD for Lisp: mapping, locator set, adjacency, map resolver. - Toggle Lisp feature status. - Configure and delete Lisp mapping as local and remote. - Configure and delete Lisp adjacency mapping - Configure and delete Lisp map resolver, proxy ITR. - Test case count: 18
- **LISP GPE** - CRUD for LISP GPE mappings. - Toggle Lisp GPE feature status. - Configure Lisp GPE mappings. - Traffic test verifying encapsulation. - Test case count: 12
- **NAT** - CRD for NAT entries, interface assignment. - Configure and delete up to two NAT entries. - Assign NAT entries to a physical interface. - Test case count: 6
- **Port mirroring** - CRD for SPAN port mirroring, interface assignment. - Configure SPAN port mirroring on a physical interface, mirroring - up to 2 interfaces. - Remove SPAN configuration from interfaces. - Test case count: 14
- **ACL-PLUGIN** - CRD for high-level classifier - MAC + IP address classification. - IPv4, IPv6 address classification. - TCP, UDP, ICMP, ICMPv6 protocol/next-header classification. - port number classification. - ICMP, ICMPv6 code and type classification. - Test case count: 15
- **ProxyARP** - CRD for proxyARP feature. - Configure proxyARP. - Assign to interface. - Test case count: 3
- **ProxyND6** - CRD for Neighbor Discovery Proxy. - Configure ProxyND6 feature on interface. - Test case count: 4

- **DHCP Relay** - CRD for DHCP relay feature. - Configure DHCP Relays. - IPv4 and IPv6 variants. - Test case count: 4
- **SLAAC** - CRD for Stateless Address AutoConfiguration. - Configure SLAAC feature on interfaces. - Test case count: 7
- **Routing** - CRD for routing. - Configure single-hop route. - Configure multi-hop routes. - Configure blackhole route. - IPv4 and IPv6 variants. - Test case count: 6
- **Policer** - CRD for traffic policing feature. - Configure Policing rules. - Assign to interface. - Test case count: 6
- **Border Gateway Protocol** - CRUD and functional tests for BGP. - Configure peers and routes - Check interactions with another BGP peer. - Test case count: 13
- **Honeycomb Infrastructure** - configuration persistence, - Netconf notifications for interface events, - Netconf negative tests aimed at specific issues - Netconf/Restconf northbound over IPv6 - Test case count: 12

Total 219 Honeycomb functional tests in the CSIT rls1804.

Operational data in Honeycomb should mirror configuration data at all times. Because of this, test cases follow this general pattern:

1. read operational data of the feature using restconf.
2. read status of the feature using VPP API dump.
3. modify configuration of the feature using restconf.
4. verify changes to operational data using restconf.
5. verify changes using VPP API dump, OR
6. send a packet to VPP node and observe behaviour to verify configuration.

Test cases involving network interfaces utilize the first two interfaces on the DUT node.

5.1.3 Functional Tests Naming

CSIT rls1804 introduced a common structured naming convention for all performance and functional tests. This change was driven by substantially growing number and type of CSIT test cases. Firstly, the original practice did not always follow any strict naming convention. Secondly test names did not always clearly capture tested packet encapsulations, and the actual type or content of the tests. Thirdly HW configurations in terms of NICs, ports and their locality were not captured either. These were but few reasons that drove the decision to change and define a new more complete and stricter test naming convention, and to apply this to all existing and new test cases.

The new naming should be intuitive for majority of the tests. The complete description of CSIT test naming convention is provided on [CSIT test naming page](#)¹⁸⁹.

Here few illustrative examples of the new naming usage for functional test suites:

1. **Physical port to physical port - a.k.a. NIC-to-NIC, Phy-to-Phy, P2P**
 - *eth2p-ethip4-ip4base-func.robot* => 2 ports of Ethernet, IPv4 baseline routed forwarding, functional tests.
2. **Physical port to VM (or VM chain) to physical port - a.k.a. NIC2VM2NIC, P2V2P, NIC2VMchain2NIC, P2V2V2P**
 - *eth2p-ethip4vxlان-l2bdbasemaclrn-eth-2vhost-1vm-func.robot* => 2 ports of Ethernet, IPv4 VXLAN Ethernet, L2 bridge-domain switching to/from two vhost interfaces and one VM, functional tests.

¹⁸⁹ <https://wiki.fd.io/view/CSIT/csit-test-naming>

5.2 CSIT Release Notes

5.2.1 Changes in CSIT rls1804

No changes.

5.2.2 Known Issues

Here is the list of known issues in CSIT rls1804 for Honeycomb functional tests in VIRL:

| # | Issue | Jira ID | Description |
|---|--|---------------|---|
| 1 | IPv6 BGP route configuration | HONEYCOMB-403 | Configuring Ipv6 route results in missing writer for IPv6Route and IPv6NextHop exception. |
| 2 | IP address subnet validation | VPP-649 | When configuring two IP addresses from the same subnet on an interface, VPP refuses the configuration but returns code 200:OK. This can cause desync between Honeycomb's config and operational data. |
| 3 | VxLAN GPE configuration crashes VPP | VPP-875 | Specific VxLAN GPE configurations cause VPP to crash and restart. |
| 4 | Operational data for IPv6 special routes | HC2VPP-254 | Special hop routes are misidentified as regular routes in operational data. |
| 5 | LISP PITR feature configuration | HC2VPP-263 | Locator set reference in operational data is incorrect. |

5.3 Test Environment

CSIT functional tests are currently executed in FD.IO VIRL testbed. See [VPP Functional Tests Environment](#) for more details.

5.4 Documentation

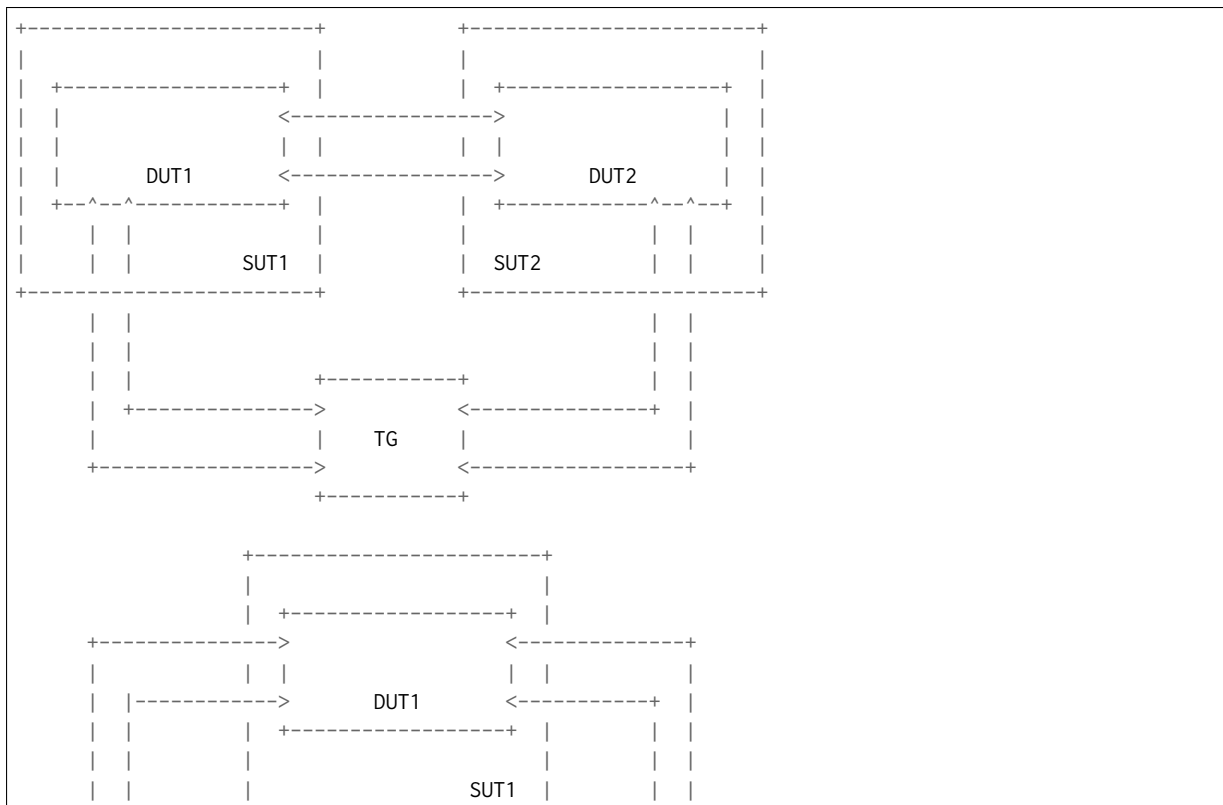
[CSIT Honeycomb Functional Tests Documentation](#)¹⁹⁰ contains detailed functional description and input parameters for each test case.

¹⁹⁰ <https://docs.fd.io/csit/rls1804/doc/tests.vpp.func.honeycomb.html>

6.1 Overview


6.1.1 Tested Virtual Topologies

CSIT NSH_SFC functional tests are executed on virtualized topologies created using VIRT simulation platform contributed by Cisco. VIRT runs on physical baremetal servers hosted by LF FD.io project. Majority of the tests are executed in the three node logical test topology - Traffic Generator (TG) node and two Systems Under Test (SUT) nodes connected in a loop. Some tests use two node logical test topology - TG node and SUT1 node. Both logical test topologies are shown in the figures below.:



(continues on next page)

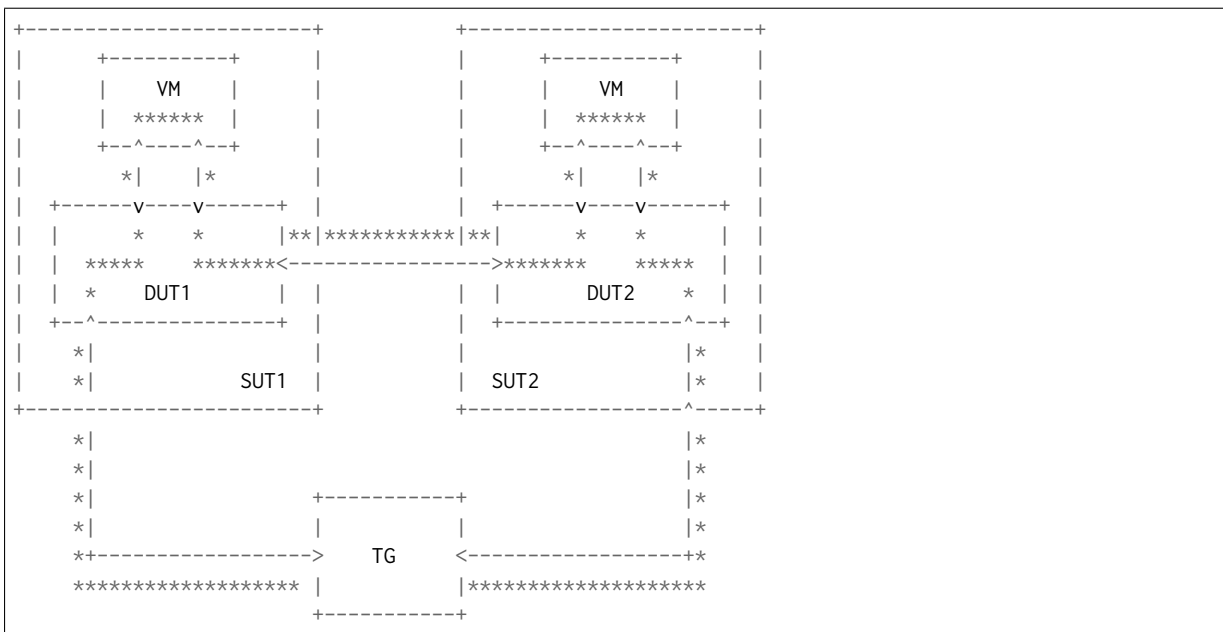
(continued from previous page)



SUT1 and SUT2 are two VMs (Ubuntu or Centos, depending on the test suite), TG is a Traffic Generator (TG, another Ubuntu VM). SUTs run VPP SW application in Linux user-mode as a Device Under Test (DUT) within the VM. TG runs Scapy SW application as a packet Traffic Generator. Logical connectivity between SUTs and to TG is provided using virtual NICs using VMs' virtio driver.

Virtual testbeds are created on-demand whenever a verification job is started (e.g. triggered by the gerrit patch submission) and destroyed upon completion of all functional tests. Each node is a Virtual Machine and each connection that is drawn on the diagram is available for use in any test case. During the test execution, all nodes are reachable thru the Management network connected to every node via dedicated virtual NICs and virtual links (not shown above for clarity).

For the test cases that require DUT (VPP) to communicate with VM over the vhost-user interfaces, a nested VM is created on SUT1 and/or SUT2 for the duration of these particular test cases only. DUT (VPP) test topology with VM is shown in the figure below including the applicable packet flow thru the VM (marked in the figure with ***):.



6.1.2 NSH_SFC Functional Tests Coverage

Following NSH_SFC functional test areas are covered in the CSIT rls1804 with results listed in this report:

- **NSH SFC Classifier** - TG sends some TCP packets to test NSH SFC Classifier functional. DUT1 will receive these packets from one NIC and loopback the VXLAN-GPE-NSH encapsulated packets to the TG from other NIC. - Test case count: 7
- **NSH SFC Proxy Inbound** - TG sends some VXLAN-GPE-NSH encapsulated packets to test NSH SFC Proxy Inbound functional. DUT1 will receive these packets from one NIC and loopback the VXLAN encapsulated packets to the TG from other NIC. - Test case count: 6
- **NSH SFC Proxy Outbound** - TG sends some VXLAN encapsulated packets to test NSH SFC Proxy

Outbound functional. DUT1 will receive these packets from one NIC and loopback the VXLAN-GPE-NSH encapsulated packets to the TG from other NIC. - Test case count: 6

- **NSH SFC Service Function Forward** - TG sends some VXLAN-GPE-NSH encapsulated packets to test NSH SFC Service Function Forward functional. DUT1 will receive these packets from one NIC and swap the VXLAN-GPE-NSH header, after that DUT1 loopback the VXLAN-GPE-NSH encapsulated packets to the TG from other NIC. - Test case count: 6

Total 25 NSH SFC functional tests in the CSIT rls1804.

6.2 CSIT Release Notes

6.2.1 Changes in CSIT rls1804

1. Added NSH SFC functional tests
 - None.

6.2.2 Known Issues

Here is the list of known issues in CSIT rls1804 for NSH_SFC functional tests in VIRT:

| # | Issue | Jira ID | Description |
|---|-------|---------|-------------|
| 1 | None | None | None |

6.3 Test Environment

CSIT functional tests are currently executed in FD.IO VIRT testbed. The physical VIRT testbed infrastructure consists of three VIRT hosts:

- All hosts are Cisco UCS C240-M4 (2x Intel(R) Xeon(R) CPU E5-2699 v3 @2.30GHz, 18c, 512GB RAM)
- tb4-virt1:
 - Status: Production
 - OS: Ubuntu 16.04.2
 - STD server version 0.10.32.16
 - UWM server version 0.10.32.16
- tb4-virt2:
 - Status: Production
 - OS: Ubuntu 16.04.2
 - STD server version 0.10.32.16
 - UWM server version 0.10.32.16
- tb4-virt3:
 - Status: Testing
 - OS: Ubuntu 16.04.2
 - STD server version 0.10.32.19
 - UWM server version 0.10.32.19

Whenever a patch is submitted to gerrit for review, parallel VIRT simulations are started to reduce the time of execution of all functional tests. The number of parallel VIRT simulations is equal to number of test groups defined by TEST_GROUPS variable in `csit/bootstrap.sh` file. The VIRT host to run VIRT simulation is selected based on least load algorithm per VIRT simulation.

Every VIRT simulation uses the same three-node - Traffic Generator (TG node) and two Systems Under Test (SUT1 and SUT2) - “double-ring” topology. The appropriate pre-built VPP packages built by Jenkins for the patch under review are then installed on the two SUTs, along with their `/etc/vpp/startup.conf` file, in all VIRT simulations.

6.3.1 SUT Configuration - VIRT Guest VM

Configurations of the SUT VMs is defined in [VIRT topologies directory](#)¹⁹¹

- List of SUT VM interfaces::

```
<interface id="0" name="GigabitEthernet0/4/0"/>
<interface id="1" name="GigabitEthernet0/5/0"/>
<interface id="2" name="GigabitEthernet0/6/0"/>
<interface id="3" name="GigabitEthernet0/7/0"/>
```

- Number of 2MB hugepages: 1024
- Maximum number of memory map areas: 20000
- Kernel Shared Memory Max: 2147483648 (`vm.nr_hugepages * 2 * 1024 * 1024`)

6.3.2 SUT Configuration - VIRT Guest OS Linux

In CSIT terminology, the VM operating system for both SUTs that VPP-18.04 release has been tested with, is the following:

1. Ubuntu VIRT image

This image implies Ubuntu 16.04.1 LTS, current as of yyyy-mm-dd (that is, package versions are those that would have been installed by a `apt-get update`, `apt-get upgrade` on that day), produced by CSIT disk image build scripts.

The exact list of installed packages and their versions (including the Linux kernel package version) are included in [VIRT ubuntu images lists](#)¹⁹².

A replica of this VM image can be built by running the `build.sh` script in CSIT repository.

2. CentOS VIRT image

This image implies Centos 7.4-1711, current as of yyyy-mm-dd (that is, package versions are those that would have been installed by a `yum update`, `yum upgrade` on that day), produced by CSIT disk image build scripts.

The exact list of installed packages and their versions (including the Linux kernel package version) are included in [VIRT centos images lists](#)¹⁹³.

A replica of this VM image can be built by running the `build.sh` script in CSIT repository.

3. Nested VM image

In addition to the “main” VM image, tests which require VPP to communicate to a VM over a vhost-user interface, utilize a “nested” VM image.

This “nested” VM is dynamically created and destroyed as part of a test case, and therefore the “nested” VM image is optimized to be small, lightweight and have a short boot time. The “nested”

¹⁹¹ <https://git.fd.io/csit/tree/resources/tools/virt/topologies/?h=rls1804>

¹⁹² <https://git.fd.io/csit/tree/resources/tools/disk-image-builder/ubuntu/lists/?h=rls1804>

¹⁹³ <https://git.fd.io/csit/tree/resources/tools/disk-image-builder/centos/lists/?h=rls1804>

VM image is not built around any established Linux distribution, but is based on [BuildRoot](https://buildroot.org/)¹⁹⁴, a tool for building embedded Linux systems. Just as for the “main” image, scripts to produce an identical replica of the “nested” image are included in CSIT GIT repository, and the image can be rebuilt using the “build.sh” script at [VIRL nested](https://git.fd.io/csit/tree/resources/tools/virl/topologies/?h=rls1804)¹⁹⁵.

6.3.3 DUT Configuration - VPP

Every System Under Test runs VPP SW application in Linux user-mode as a Device Under Test (DUT) node.

DUT port configuration

Port configuration of DUTs is defined in topology file that is generated per VIRL simulation based on the definition stored in [VIRL topologies directory](https://git.fd.io/csit/tree/resources/tools/virl/topologies/?h=rls1804)¹⁹⁶.

Example of DUT nodes configuration::

```
DUT1:
  type: DUT
  host: "10.30.51.157"
  port: 22
  username: cisco
  honeycomb:
    user: admin
    passwd: admin
    port: 8183
    netconf_port: 2831
  priv_key: |
    -----BEGIN RSA PRIVATE KEY-----
    MIIIEgIBAAKCAQEAWUD1TpzShpwLQotZOFS4AgcPNEWcNp1AB2hWfMvI+8Kah/gb
    v8ruZU9RqhPs56tyKzxbhvNkY4VbH5F1GihZu3mLqz4KfghMmaeMEj01T7BYyd
    vuBfTvIlu1jfQ2vA1nYrDwn+C1xJk81m0pDgvrLEX4qVvh2sGh7UEkYy5r82DNa2
    4VjzPB1J/c8a9zP8FoZuHYIzF4FLvRMjUADpbMXgJMsGpaZLmz95ap0Eot7vb1Cc
    1LvF97iyBCrtIOSKRKA50ZHLGjMKmOwnYU+cP5718tbproDVi6VJ0o7zeuXyetMs
    8YB19kwb1WG9BqP9jctFvsmi5G7hXgq1Y8u+DwIDAQABAoIBAQC/W4E0DHjLMny7
    0bvw2YkZD0Zw3fttdB94tkm4PdZv5MybooPnsAvLaXVV0hEdfVi5kzSWN1/LY/tN
    EP1BgGphc2QgB59/PPxGwFIjDCvUz1sZpynBHe+B/qh5ExNQCvvsIOqWI7DXLXaN
    0i/khOzmJ6HncRRah1spKimYRsaUUDskyg7q3QqMwVaqBbbMvLs/w7ZWd/zoDqCU
    MY/pCI6hkB3QbRo00diZLohphB12ShABTwjvVyyKL5UA4jAeneJrhH5gWVLXnfgD
    p62W5Col1KEYb1C8mUkPxpP7Qo277zw3xaq+oktIZhc5SUEUd7nJZtNqVAHqkItw
    79VmpKyxAoGBAPFUp+kNPaTSvp+x1n5sn2SgipzDgtgi9QqNmC4cjrQQaaqI57SG
    OHw1jX8i7L2G1WvVtkHg060n1EVo5n65fff0qeVBezLVJ7ghWI8U+oBiJJyQ4boD
    GJVns0OSUQ0rtuGd9eVwfdk3o19aCN0KK53oPfIYli29pyu41095kg11AoGBAMef
    bPEMBI/2XmCPshLSwhGF1+dW8d+K11uj3CUQ/0vU1vma3dfBOYNsIwAgTP0iIUTg
    8DYE6KBCDPtxAUEI0YAEAKB9ry1tKR2NQEIPfs1YytKErtwJaiQSi0heM6+zwEzu
    f54Z4oBhMSL0jXoMnu+NzZec6EUdQeY40+jhzAoGBAIogC3dtjMpgKTP7+93u
    UE/XIioI8fWg9fj3sMka4IMu+pVvRCRbAjRH7JrFLkjbUyuMqs3Arnk9K+gbdQt/
    +m95Njtt6WofXUcwgBm3GidSmZwYT4454SfDzVBYScedcNm1FuR+8ov9bFLDtGT
    D4gsngnGJj1MDFXTxZEn4nzZaOGBAKCg4WmpUPaCuXibyB+rZavxwsTNSn21J83/
    sYJGBhf/raiV/FLDUcM1vYg5dZnu37RsB/5/vqx0LZGyYd7x+Jo5HkQGpNkgNwhn
    g8BkdZIRF8uEJqx0o0ycd0U7n/2093swIpKWo5LIiRPuqqzj+uZKnAL7vuVdxfaY
    qVz2daMPAoGBALgaaKa3voU/H01PYLWIhFrBThyJ+BQSQ80qrEzC8AnegWFxRAM8
    EqrzZX17ACUuo1dh0Eipm41j2+BZW1QjiUgq5uj8+yzy+EU1ZRRyJcOKzbdACeuD
    BpWWSXGBI5G4CppeYLjMUHZpJYeX1USULJQd2c4crLJKb76E8gz3Z9kn
    -----END RSA PRIVATE KEY-----

  interfaces:
    port1:
      mac_address: "fa:16:3e:9b:89:52"
```

(continues on next page)

¹⁹⁴ <https://buildroot.org/>

¹⁹⁵ <https://git.fd.io/csit/tree/resources/tools/disk-image-builder/nested/?h=rls1804>

¹⁹⁶ <https://git.fd.io/csit/tree/resources/tools/virl/topologies/?h=rls1804>

(continued from previous page)

```

    pci_address: "0000:00:04.0"
    link: link1
port2:
    mac_address: "fa:16:3e:7a:33:60"
    pci_address: "0000:00:05.0"
    link: link4
port3:
    mac_address: "fa:16:3e:29:b7:ae"
    pci_address: "0000:00:06.0"
    link: link3
port4:
    mac_address: "fa:16:3e:76:8d:ff"
    pci_address: "0000:00:07.0"
    link: link6
DUT2:
type: DUT
host: "10.30.51.156"
port: 22
username: cisco
honeycomb:
  user: admin
  passwd: admin
  port: 8183
  netconf_port: 2831
priv_key: |
-----BEGIN RSA PRIVATE KEY-----
MIIePgIBAAKCAQEAWUD1TpzSHpwLQotZOFs4AgcPNEWcNp1AB2hWfMvI+8Kah/gb
v8ruZU9RqhPs56tyKzxbhvNkY4VbH5F1GiLHZu3mLqzM4KfghMmaeMEj0T7BYyd
vuBfTvIlu1jFQ2vAlnYrDwn+C1xJk81m0pDgvrLEX4qVvH2sGh7UEkYy5r82DNa2
4VjzPB1J/c8a9zP8FoZUhYIzF4FLvRMjUADpbMXgJMsGpaZLmz95ap0Eot7vb1Cc
1LvF97iyBCrtIOSKRKA50ZHLGjMKmOwNyu+cP5718tbproDVi6VJ0o7zeuXyetMs
8YB19kwb1WG9BqP9jctFvsmi5G7hXgq1Y8u+DwIDAQABAoIBAQC/W4E0DHjLMny7
0bvw2YKzD0Zw3fttdB94tkm4PdZv5MybooPnsAvLaXVV0hEdfVi5kzSWN1/LY/tN
EP1BgGphc2QgB59/PPxGwFIjDCvUz1sZpynBHe+B/qh5ExNQcVvsIOqWI7DXLXaN
0i/kh0zmJ6HncRRah1spKimYRsaUUDskyg7q3QqMwVaqBbbMvLs/w7Zwd/zoDqCU
MY/pCI6hkB3QbRo00diZLohphB12ShABTwjvVyyKL5UA4jAEneJrhH5gWVLXnfgD
p62W5Coll1KEYb1C8mUkPxpP7Qo277zw3xaq+oktIZhc5SUEud7nJZtNqVAHqkItW
79VmpKyxAoGBAPFuf+kqNPaTSvp+x1n5sn2SgipzDgtgi9QqNmC4cjtrQQAaqI57SG
OHw1jX8i7L2G1WvVtkHg060n1EVo5n65ffF0qeVBezLVJ7ghWI8U+oBiJjYQ4boD
GJVNs0S0UQ0rtuGd9eVwfdk3o19aCN0KK53oPFIY1i29pyu41095kg11AoGBAMef
bPEMBI/2XmCPshLShwGF1+dW8d+K11uj3CUQ/0vUlvma3dfBOYNsIwAgTP0iIUTg
8DYE6KBCdPtXAEI0YAEAKB9ry1tkR2NQEIpfslYytkErtwJaiqSi0heM6+zwEzu
f54Z4oBhMSL0jXoMnu+NZZEc6EUdQeY40+jhjzAoGBAIOgC3dtjMPGKTP7+93u
UE/XIioI8fWg9fj3sMka4IMu+pVvRCRBAjRH7JrFLkjbUyuMqs3Arnk9K+gbdQt/
+m95Njtt6WoFXuPCwgbM3GidSmZwYT4454SfDzVBYScedcNm1FuR+8ov9bFLDtGT
D4gsngnGJj1MDFXTxZEn4nzZAoGBAKCg4WmpUPaCuXibyB+rZavxwstNSn2lJ83/
sYJGBhf/raiV/FLDUcM1vYg5dZnu37RsB/5/vqx0LZGyYd7x+Jo5HkQGPnKgNwhn
g8BkdZIRF8uEJqx0o0ycd0U7n/2093swIpKWo5LIiRPuqqzj+uZKnAL7vuVdxfaY
qVz2daMPAoGBALgaaKa3voU/H01PYLWIhFrBThyJ+BQSQ80qrEzC8AnegWFxRAM8
EqrzZX17ACUuo1dh0Eipm41j2+BZW1QjiUgq5uj8+yzy+EU1ZRRyJcOKzbdACEud
BpWWSXGBI5G4CppeYLjMUHZpJYeX1USULJQd2c4crLJKb76E8gz3Z9kN
-----END RSA PRIVATE KEY-----

interfaces:
port1:
  mac_address: "fa:16:3e:ad:6c:7d"
  pci_address: "0000:00:04.0"
  link: link2
port2:
  mac_address: "fa:16:3e:94:a4:99"
  pci_address: "0000:00:05.0"

```

(continues on next page)

(continued from previous page)

```

link: link5
port3:
  mac_address: "fa:16:3e:75:92:da"
  pci_address: "0000:00:06.0"
link: link3
port4:
  mac_address: "fa:16:3e:2c:b1:2a"
  pci_address: "0000:00:07.0"
link: link6

```

VPP Version

VPP-18.04 release

VPP Installed Packages - Ubuntu

```

$ dpkg -l vpp\*
Desired=Unknown/Install/Remove/Purge/Hold
| Status=Not/Inst/Conf-files/Unpacked/half-conf/Half-inst/trig-aWait/Trig-pend
|/ Err?=(none)/Reinst-required (Status,Err: uppercase=bad)
||/ Name           Version           Architecture Description
+++-----+-----+-----+-----+
ii vpp              18.04-release    amd64         Vector Packet Processing--executables
ii vpp-dbg          18.04-release    amd64         Vector Packet Processing--debug symbols
ii vpp-dev          18.04-release    amd64         Vector Packet Processing--development support
ii vpp-dpdk-dkms    18.02-vpp1       amd64         DPDK Development Package for VPP - Kernel Modules
ii vpp-lib          18.04-release    amd64         Vector Packet Processing--runtime libraries
ii vpp-plugins     18.04-release    amd64         Vector Packet Processing--runtime plugins

```

VPP Installed Packages - Centos

```

$ rpm -qai vpp*
Name       : vpp
Version    : 18.04
Release    : release
Architecture: x86_64
Install Date: Fri 04 May 2018 02:47:36 AM EDT
Group      : Unspecified
Size       : 2311324
License    : ASL 2.0
Signature  : (none)
Source RPM : vpp-18.04-release.src.rpm
Build Date : Wed 25 Apr 2018 01:00:24 PM EDT
Build Host : c29a51c12e88
Relocations : (not relocatable)
Summary    : Vector Packet Processing
Description :
This package provides VPP executables: vpp, vpp_api_test, vpp_json_test
vpp - the vector packet engine
vpp_api_test - vector packet engine API test tool
vpp_json_test - vector packet engine JSON test tool
Name       : vpp-lib
Version    : 18.04
Release    : release
Architecture: x86_64
Install Date: Fri 04 May 2018 02:47:36 AM EDT
Group      : System Environment/Libraries
Size       : 24949034
License    : ASL 2.0
Signature  : (none)
Source RPM : vpp-18.04-release.src.rpm
Build Date : Wed 25 Apr 2018 01:00:24 PM EDT

```

(continues on next page)

(continued from previous page)

```

Build Host : c29a51c12e88
Relocations : (not relocatable)
Summary : VPP libraries
Description :
This package contains the VPP shared libraries, including:
vppinfra - foundation library supporting vectors, hashes, bitmaps, pools, and string formatting.
svm - vm library
vlib - vector processing library
vlib-api - binary API library
vnet - network stack library
Name : vpp-selinux-policy
Version : 18.04
Release : release
Architecture: x86_64
Install Date: Fri 04 May 2018 02:47:35 AM EDT
Group : System Environment/Base
Size : 86170
License : ASL 2.0
Signature : (none)
Source RPM : vpp-18.04-release.src.rpm
Build Date : Wed 25 Apr 2018 01:00:24 PM EDT
Build Host : c29a51c12e88
Relocations : (not relocatable)
Summary : VPP Security-Enhanced Linux (SELinux) policy
Description :
This package contains a tailored VPP SELinux policy
Name : vpp-devel
Version : 18.04
Release : release
Architecture: x86_64
Install Date: Fri 04 May 2018 02:47:38 AM EDT
Group : Development/Libraries
Size : 10632218
License : ASL 2.0
Signature : (none)
Source RPM : vpp-18.04-release.src.rpm
Build Date : Wed 25 Apr 2018 01:00:24 PM EDT
Build Host : c29a51c12e88
Relocations : (not relocatable)
Summary : VPP header files, static libraries
Description :
This package contains the header files for VPP.
Install this package if you want to write a
program for compilation and linking with vpp lib.
vlib
vlibmemory
vnet - devices, classify, dhcp, ethernet flow, gre, ip, etc.
vpp-api
vppinfra
Name : vpp-plugins
Version : 18.04
Release : release
Architecture: x86_64
Install Date: Fri 04 May 2018 02:47:37 AM EDT
Group : System Environment/Libraries
Size : 44145968
License : ASL 2.0
Signature : (none)
Source RPM : vpp-18.04-release.src.rpm
Build Date : Wed 25 Apr 2018 01:00:24 PM EDT
Build Host : c29a51c12e88

```

(continues on next page)

(continued from previous page)

```

Relocations : (not relocatable)
Summary      : Vector Packet Processing--runtime plugins
Description  :
This package contains VPP plugins

```

VPP Startup Configuration

VPP startup configuration is common for all test cases except test cases related to SW Crypto device.

Default

```

$ cat /etc/vpp/startup.conf
unix {
    nodaemon
    log /var/log/vpp/vpp.log
    full-coredump
    cli-listen /run/vpp/cli.sock
    gid vpp
}

api-trace {
## This stanza controls binary API tracing. Unless there is a very strong reason,
## please leave this feature enabled.
    on
## Additional parameters:
##
## To set the number of binary API trace records in the circular buffer, configure nitens
##
## nitens <nnn>
##
## To save the api message table decode tables, configure a filename. Results in /tmp/<filename>
## Very handy for understanding api message changes between versions, identifying missing
## plugins, and so forth.
##
## save-api-table <filename>
}

api-segment {
    gid vpp
}

cpu {
    ## In the VPP there is one main thread and optionally the user can create worker(s)
    ## The main thread and worker thread(s) can be pinned to CPU core(s) manually or automatically

    ## Manual pinning of thread(s) to CPU core(s)

    ## Set logical CPU core where main thread runs
    # main-core 1

    ## Set logical CPU core(s) where worker threads are running
    # corelist-workers 2-3,18-19

    ## Automatic pinning of thread(s) to CPU core(s)

    ## Sets number of CPU core(s) to be skipped (1 ... N-1)
    ## Skipped CPU core(s) are not used for pinning main thread and working thread(s).
    ## The main thread is automatically pinned to the first available CPU core and worker(s)
    ## are pinned to next free CPU core(s) after core assigned to main thread
    # skip-cores 4

    ## Specify a number of workers to be created

```

(continues on next page)

(continued from previous page)

```

## Workers are pinned to N consecutive CPU cores while skipping "skip-cores" CPU core(s)
## and main thread's CPU core
# workers 2

## Set scheduling policy and priority of main and worker threads

## Scheduling policy options are: other (SCHED_OTHER), batch (SCHED_BATCH)
## idle (SCHED_IDLE), fifo (SCHED_FIFO), rr (SCHED_RR)
# scheduler-policy fifo

## Scheduling priority is used only for "real-time policies (fifo and rr),
## and has to be in the range of priorities supported for a particular policy
# scheduler-priority 50
}

# dpdk {
## Change default settings for all interfaces
# dev default {
## Number of receive queues, enables RSS
## Default is 1
# num-rx-queues 3

## Number of transmit queues, Default is equal
## to number of worker threads or 1 if no workers threads
# num-tx-queues 3

## Number of descriptors in transmit and receive rings
## increasing or reducing number can impact performance
## Default is 1024 for both rx and tx
# num-rx-desc 512
# num-tx-desc 512

## VLAN strip offload mode for interface
## Default is off
# vlan-strip-offload on
# }

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

## Whitelist specific interface by specifying PCI address and in
## addition specify custom parameters for this interface
# dev 0000:02:00.1 {
#     num-rx-queues 2
# }

## Specify bonded interface and its slaves via PCI addresses
##
## Bonded interface in XOR load balance mode (mode 2) with L3 and L4 headers
# vdev eth_bond0,mode=2,slave=0000:02:00.0,slave=0000:03:00.0,xmit_policy=134
# vdev eth_bond1,mode=2,slave=0000:02:00.1,slave=0000:03:00.1,xmit_policy=134
##
## Bonded interface in Active-Back up mode (mode 1)
# vdev eth_bond0,mode=1,slave=0000:02:00.0,slave=0000:03:00.0
# vdev eth_bond1,mode=1,slave=0000:02:00.1,slave=0000:03:00.1

## Change UIO driver used by VPP, Options are: igb_uio, vfio-pci,
## uio_pci_generic or auto (default)
# uio-driver vfio-pci

## Disable multi-segment buffers, improves performance but

```

(continues on next page)

(continued from previous page)

```

## disables Jumbo MTU support
# no-multi-seg

## Increase number of buffers allocated, needed only in scenarios with
## large number of interfaces and worker threads. Value is per CPU socket.
## Default is 16384
# num-mbufs 128000

## Change hugepages allocation per-socket, needed only if there is need for
## larger number of mbufs. Default is 256M on each detected CPU socket
# socket-mem 2048,2048

## Disables UDP / TCP TX checksum offload. Typically needed for use
## faster vector PMDs (together with no-multi-seg)
# no-tx-checksum-offload
# }

# plugins {
## Adjusting the plugin path depending on where the VPP plugins are
#     path /home/bms/vpp/build-root/install-vpp-native/vpp/lib64/vpp_plugins

## Disable all plugins by default and then selectively enable specific plugins
# plugin default { disable }
# plugin dpdk_plugin.so { enable }
# plugin acl_plugin.so { enable }

## Enable all plugins by default and then selectively disable specific plugins
# plugin dpdk_plugin.so { disable }
# plugin acl_plugin.so { disable }
# }

## Alternate syntax to choose plugin path
# plugin_path /home/bms/vpp/build-root/install-vpp-native/vpp/lib64/vpp_plugins

```

SW Crypto Device

```

$ cat /etc/vpp/startup.conf
unix
{
  cli-listen /run/vpp/cli.sock
  gid vpp
  nodaemon
  full-coredump
  log /tmp/vpp.log
}
api-segment
{
  gid vpp
}
dpdk
{
  vdev cryptodev_aesni_gcm_pmd,socket_id=0
  vdev cryptodev_aesni_mb_pmd,socket_id=0
}

```

6.3.4 TG Configuration

Traffic Generator node is VM running the same OS Linux as SUTs. Ports of this VM are used as source (Tx) and destination (Rx) ports for the traffic.

Traffic scripts of test cases are executed on this VM.

TG VM configuration

Configuration of the TG VMs is defined in [VIRL topologies directory](#)¹⁹⁷.

```
/csit/resources/tools/virl/topologies/double-ring-nested.xenial.virl
```

- List of TG VM interfaces::

```
<interface id="0" name="eth1"/>
<interface id="1" name="eth2"/>
<interface id="2" name="eth3"/>
<interface id="3" name="eth4"/>
<interface id="4" name="eth5"/>
<interface id="5" name="eth6"/>
```

TG node port configuration

Port configuration of TG is defined in topology file that is generated per VIRL simulation based on the definition stored in [VIRL topologies directory](#)¹⁹⁸.

Example of TG node configuration::

```
TG:
  type: TG
  host: "10.30.51.155"
  port: 22
  username: cisco
  priv_key: |
    -----BEGIN RSA PRIVATE KEY-----
    MIIIEpgIBAACKAQEAwUDlTpzSHpwLQotZOFs4AgcPNEwCnP1AB2hWfMvI+8Kah/gb
    v8ruZU9RqhPs56tyKzxbhvNkY4VbH5F1Gi1HZu3mLqzM4KfghMmaeMEj01T7BYyd
    vuBfTvIlu1jFQ2vAlnYrDwn+C1xJk81m0pDgvrLEX4qVvH2sGh7UEkYy5r82DNa2
    4VjzPB1J/c8a9zP8FoZUHYIzF4FLvRMjUADpbMXgJMsGpaZLmz95ap0Eot7vb1Cc
    1LvF97iyBCrtIOSKRKA50ZHLGjMKmOwNyu+cP5718tbproDVi6VJ0o7zeuXyetMs
    8YB19kwb1WG9BqP9jctFvsmi5G7hXgq1Y8u+DwIDAQABAoIBAQC/W4E0DHjLMny7
    0bvW2YKzD0Zw3fttdB94tkm4PdZv5MybooPnsAvLaXVv0hEdfVi5kzSWN1/LY/tN
    EP1BgGphc2QgB59/PPxGwFIjDCvUz1sZpynBHe+B/qh5ExNqCvVsIOqWI7DX1XaN
    0i/khOzmJ6HncRRah1spKimYRsaUUDskyg7q3QqMwVaqBbbMvLs/w7Zwd/zoDqCU
    MY/pCI6hkB3QbRo0odiZLohphB12ShABTjvVyyKL5UA4jAeneJrhH5gWVLXnfgD
    p62W5CollKEYb1C8mUkPxpP7Qo277zw3xaq+oktIZhc5SUEUd7nJZtNqVAHqkItW
    79VmpKyxAoGBAPfU+kqNPaTSvp+x1n5sn2SgipzDgti9QqNmC4cjtrQqaqI57SG
    OHw1jX8i7L2G1WvVtkHg060n1EVo5n65fffQqeVBezLVJ7ghWI8U+oBiJJyQ4boD
    GJVnso0SUQ0rtuGd9eVwfDk3o19aCN0KK53oPFIYli29pyu4l095kg11AoGBAMef
    bPEMBI/2XmCPshLShwGF1+dW8d+K11uj3CUQ/0vU1vma3dfBOYNsIwAgTP0iIUTg
    8DYE6KBCdPtXAEUEI0YAEAKB9ry1tKR2NQEIPfs1YytkErtwjAiqSi0heM6+zwEzu
    f54Z4oBhMSL0jXoMnu+NZZec6EUdQeY40+jhjzAoGBAIOgc3dtjMPGKTP7+93u
    UE/XIioI8fWg9fj3sMka4IMu+pVvRCRbAjRH7JrFLkjbUyuMqs3Arnk9K+gbdQt/
    +m95Njtt6WoFXuPCwgbM3GidSmZwYT4454SfDzVBYScEDCNm1FuR+8ov9bFLDtGT
    D4gsngnGJj1MDFXTxZEn4nzZAoGBAKCg4WmpUPaCuXibyB+rZavxwsTNSn2lJ83/
    sYJGBhf/raiv/FLDUcM1vYg5dZnu37RsB/5/vqx0LZGyYd7x+Jo5HkQGPnKgNwhn
    g8BkdZIRF8uEJqx0o0ycd0U7n/2093swIpKWo5LIiRPuqqzj+uZKnAL7vuVdxfaY
    qVz2daMPAogBALgaak3voU/H01PYLWIhFrBThyJ+BQSQ80qrEzC8AnegWFxRAM8
    EqrzZX17ACUuo1dH0Eipm41j2+BZWLQjiUgq5uj8+zyzy+EU1ZRRyJc0KzbdACEud
    BpWWSXGBl5G4CppeYLjMUHZpJYeX1USULJQd2c4crLJKb76E8gz3Z9kN
    -----END RSA PRIVATE KEY-----

  interfaces:
    port3:
      mac_address: "fa:16:3e:b9:e1:27"
      pci_address: "0000:00:06.0"
```

(continues on next page)

¹⁹⁷ <https://git.fd.io/csit/tree/resources/tools/virl/topologies/?h=rls1804>

¹⁹⁸ <https://git.fd.io/csit/tree/resources/tools/virl/topologies/?h=rls1804>

(continued from previous page)

```
link: link1
driver: virtio-pci
port4:
  mac_address: "fa:16:3e:e9:c8:68"
  pci_address: "0000:00:07.0"
link: link4
driver: virtio-pci
port5:
  mac_address: "fa:16:3e:e8:d3:47"
  pci_address: "0000:00:08.0"
link: link2
driver: virtio-pci
port6:
  mac_address: "fa:16:3e:cf:ca:58"
  pci_address: "0000:00:09.0"
link: link5
driver: virtio-pci
```

Traffic generator

Functional tests utilize Scapy as a traffic generator. There was used Scapy v2.3.1 for VPP-18.04 release tests.

6.4 Documentation

[CSIT NSH_SFC Functional Tests Documentation](https://docs.fd.io/csit/rls1804/doc/tests.nsh_sfc.func.html)¹⁹⁹ contains detailed functional description and input parameters for each test case.

¹⁹⁹ https://docs.fd.io/csit/rls1804/doc/tests.nsh_sfc.func.html

7.1 CSIT Design

FD.io CSIT system design needs to meet continuously expanding requirements of FD.io projects including VPP, related sub-systems (e.g. plugin applications, DPDK drivers) and FD.io applications (e.g. DPDK applications), as well as growing number of compute platforms running those applications. With CSIT project scope and charter including both FD.io continuous testing AND performance trending/comparisons, those evolving requirements further amplify the need for CSIT framework modularity, flexibility and usability.

7.1.1 Design Hierarchy

CSIT follows a hierarchical system design with SUTs and DUTs at the bottom level of the hierarchy, presentation level at the top level and a number of functional layers in-between. The current CSIT system design including CSIT framework is depicted in the figure below.

- * Easily extensible – one can create a new stream profile that meets tests requirements;
- * Same stream profile can be used for different tests with the same traffic needs;
- Functional data plane traffic scripts:
 - Scapy specific traffic scripts;
- 3. Level-2 libraries - Robot resource files:
 - Higher level CSIT libraries abstracting required functions for executing tests;
 - L2 KWs are classified into the following functional categories:
 - Configuration, test, verification, state report;
 - Suite setup, suite teardown;
 - Test setup, test teardown;
- 4. Tests - Robot:
 - Test suites with test cases;
 - Functional tests using VIRT environment:
 - VPP;
 - Honeycomb;
 - NSH_SFC;
 - Performance tests using physical testbed environment:
 - VPP;
 - DPDK-Testpmd;
 - DPDK-L3Fwd;
 - Honeycomb;
 - VPP Container K8s orchestrated topologies;
 - Tools:
 - Documentation generator;
 - Report generator;
 - Testbed environment setup ansible playbooks;
 - Operational debugging scripts;

7.1.2 Test Lifecycle Abstraction

A well coded test must follow a disciplined abstraction of the test lifecycles that includes setup, configuration, test and verification. In addition to improve test execution efficiency, the common aspects of test setup and configuration shared across multiple test cases should be done only once. Translating these high-level guidelines into the Robot Framework one arrives to definition of a well coded RF tests for FD.io CSIT. Anatomy of Good Tests for CSIT:

1. Suite Setup - Suite startup Configuration common to all Test Cases in suite: uses Configuration KWs, Verification KWs, StateReport KWs;
2. Test Setup - Test startup Configuration common to multiple Test Cases: uses Configuration KWs, StateReport KWs;
3. Test Case - uses L2 KWs with RF Gherkin style:
 - prefixed with {Given} - Verification of Test setup, reading state: uses Configuration KWs, Verification KWs, StateReport KWs;

- prefixed with {When} - Test execution: Configuration KWs, Test KWs;
 - prefixed with {Then} - Verification of Test execution, reading state: uses Verification KWs, StateReport KWs;
4. Test Teardown - post Test teardown with Configuration cleanup and Verification common to multiple Test Cases - uses: Configuration KWs, Verification KWs, StateReport KWs;
 5. Suite Teardown - Suite post-test Configuration cleanup: uses Configuration KWs, Verification KWs, StateReport KWs;

7.1.3 RF Keywords Functional Classification

CSIT RF KWs are classified into the functional categories matching the test lifecycle events described earlier. All CSIT RF L2 and L1 KWs have been grouped into the following functional categories:

1. Configuration;
2. Test;
3. Verification;
4. StateReport;
5. SuiteSetup;
6. TestSetup;
7. SuiteTeardown;
8. TestTeardown;

7.1.4 RF Keywords Naming Guidelines

Readability counts: “..code is read much more often than it is written.” Hence following a good and consistent grammar practice is important when writing RF KeyWords and Tests. All CSIT test cases are coded using Gherkin style and include only L2 KWs references. L2 KWs are coded using simple style and include L2 KWs, L1 KWs, and L1 python references. To improve readability, the proposal is to use the same grammar for both RF KW styles, and to formalize the grammar of English sentences used for naming the RF KWs. RF KWs names are short sentences expressing functional description of the command. They must follow English sentence grammar in one of the following forms:

1. **Imperative** - verb-object(s): “*Do something*”, verb in base form.
2. **Declarative** - subject-verb-object(s): “*Subject does something*”, verb in a third-person singular present tense form.
3. **Affirmative** - modal_verb-verb-object(s): “*Subject should be something*”, “*Object should exist*”, verb in base form.
4. **Negative** - modal_verb-Not-verb-object(s): “*Subject should not be something*”, “*Object should not exist*”, verb in base form.

Passive form MUST NOT be used. However a usage of past participle as an adjective is okay. See usage examples provided in the Coding guidelines section below. Following sections list applicability of the above grammar forms to different RF KW categories. Usage examples are provided, both good and bad.

7.1.5 Coding guidelines

Coding guidelines can be found on [Design optimizations wiki page](https://wiki.fd.io/view/CSIT/Design_Optimizations)²⁰⁰.

²⁰⁰ https://wiki.fd.io/view/CSIT/Design_Optimizations

7.2 CSIT Test Naming

7.2.1 Background

CSIT rls1804 follows a common structured naming convention for all performance and system functional tests, introduced in CSIT rls1701.

The naming should be intuitive for majority of the tests. Complete description of CSIT test naming convention is provided on [CSIT test naming wiki page](#)²⁰¹. Below few illustrative examples of the naming usage for test suites across CSIT performance, functional and Honeycomb management test areas.

7.2.2 Naming Convention

The CSIT approach is to use tree naming convention and to encode following testing information into test suite and test case names:

1. packet network port configuration
 - port type, physical or virtual;
 - number of ports;
 - NIC model, if applicable;
 - port-NIC locality, if applicable;
2. packet encapsulations;
3. VPP packet processing
 - packet forwarding mode;
 - packet processing function(s);
4. packet forwarding path
 - if present, network functions (processes, containers, VMs) and their topology within the computer;
5. main measured variable, type of test.

Proposed convention is to encode ports and NICs on the left (underlay), followed by outer-most frame header, then other stacked headers up to the header processed by vSwitch-VPP, then VPP forwarding function, then encap on vhost interface, number of vhost interfaces, number of VMs. If chained VMs present, they get added on the right. Test topology is expected to be symmetric, in other words packets enter and leave SUT through ports specified on the left of the test name. Here some examples to illustrate the convention followed by the complete legend, and tables mapping the new test filenames to old ones.

7.2.3 Naming Examples

CSIT test suite naming examples (filename.robot) for common tested VPP topologies:

1. **Physical port to physical port - a.k.a. NIC-to-NIC, Phy-to-Phy, P2P**
 - *PortNICConfig-WireEncapsulation-PacketForwardingFunction- PacketProcessingFunction1-...-PacketProcessingFunctionN-TestType*
 - *10ge2p1x520-dot1q-l2bdbasemaclrn-ndrdisc.robot* => 2 ports of 10GE on Intel x520 NIC, dot1q tagged Ethernet, L2 bridge-domain baseline switching with MAC learning, NDR throughput discovery.

²⁰¹ <https://wiki.fd.io/view/CSIT/csit-test-naming>

- *10ge2p1x520-ethip4vxlan-l2bdbasemaclrn-ndrchk.robot* => 2 ports of 10GE on Intel x520 NIC, IPv4 VXLAN Ethernet, L2 bridge-domain baseline switching with MAC learning, NDR throughput discovery.
- *10ge2p1x520-ethip4-ip4base-ndrdisc.robot* => 2 ports of 10GE on Intel x520 NIC, IPv4 baseline routed forwarding, NDR throughput discovery.
- *10ge2p1x520-ethip6-ip6scale200k-ndrdisc.robot* => 2 ports of 10GE on Intel x520 NIC, IPv6 scaled up routed forwarding, NDR throughput discovery.
- *10ge2p1x520-ethip4-ip4base-iacldstbase-ndrdisc.robot* => 2 ports of 10GE on Intel x520 NIC, IPv4 baseline routed forwarding, ingress Access Control Lists baseline matching on destination, NDR throughput discovery.
- *40ge2p1vic1385-ethip4-ip4base-ndrdisc.robot* => 2 ports of 40GE on Cisco vic1385 NIC, IPv4 baseline routed forwarding, NDR throughput discovery.
- *eth2p-ethip4-ip4base-func.robot* => 2 ports of Ethernet, IPv4 baseline routed forwarding, functional tests.

2. Physical port to VM (or VM chain) to physical port - a.k.a. NIC2VM2NIC, P2V2P, NIC2VMchain2NIC, P2V2V2P

- *PortNICConfig-WireEncapsulation-PacketForwardingFunction- PacketProcessingFunction1-...- PacketProcessingFunctionN-VirtEncapsulation- VirtPortConfig-VMconfig-TestType*
- *10ge2p1x520-dot1q-l2bdbasemaclrn-eth-2vhost-1vm-ndrdisc.robot* => 2 ports of 10GE on Intel x520 NIC, dot1q tagged Ethernet, L2 bridge-domain switching to/from two vhost interfaces and one VM, NDR throughput discovery.
- *10ge2p1x520-ethip4vxlan-l2bdbasemaclrn-eth-2vhost-1vm-ndrdisc.robot* => 2 ports of 10GE on Intel x520 NIC, IPv4 VXLAN Ethernet, L2 bridge-domain switching to/from two vhost interfaces and one VM, NDR throughput discovery.
- *10ge2p1x520-ethip4vxlan-l2bdbasemaclrn-eth-4vhost-2vm-ndrdisc.robot* => 2 ports of 10GE on Intel x520 NIC, IPv4 VXLAN Ethernet, L2 bridge-domain switching to/from four vhost interfaces and two VMs, NDR throughput discovery.
- *eth2p-ethip4vxlan-l2bdbasemaclrn-eth-2vhost-1vm-func.robot* => 2 ports of Ethernet, IPv4 VXLAN Ethernet, L2 bridge-domain switching to/from two vhost interfaces and one VM, functional tests.

3. API CRUD tests - Create (Write), Read (Retrieve), Update (Modify), Delete (Destroy) operations for configuration and operational data

- *ManagementTestKeyword-ManagementOperation-ManagedFunction1-...- ManagedFunctionN- ManagementAPI1-ManagementAPIN-TestType*
- *mgmt-cfg-lisp-apivat-func* => configuration of LISP with VAT API calls, functional tests.
- *mgmt-cfg-l2bd-aphc-apivat-func* => configuration of L2 Bridge-Domain with Honeycomb API and VAT API calls, functional tests.
- *mgmt-oper-int-aphcnc-func* => reading status and operational data of interface with Honeycomb NetConf API calls, functional tests.
- *mgmt-cfg-int-tap-aphcnc-func* => configuration of tap interfaces with Honeycomb NetConf API calls, functional tests.
- *mgmt-notif-int-subint-aphcnc-func* => notifications of interface and sub-interface events with Honeycomb NetConf Notifications, functional tests.

For complete description of CSIT test naming convention please refer to [CSIT test naming wiki page](#)²⁰².

²⁰² <https://wiki.fd.io/view/CSIT/csit-test-naming>

7.3 Presentation and Analytics Layer

7.3.1 Overview

The presentation and analytics layer (PAL) is the fourth layer of CSIT hierarchy. The model of presentation and analytics layer consists of four sub-layers, bottom up:

- sL1 - Data - input data to be processed:
 - Static content - .rst text files, .svg static figures, and other files stored in the CSIT git repository.
 - Data to process - .xml files generated by Jenkins jobs executing tests, stored as robot results files (output.xml).
 - Specification - .yaml file with the models of report elements (tables, plots, layout, ...) generated by this tool. There is also the configuration of the tool and the specification of input data (jobs and builds).
- sL2 - Data processing
 - The data are read from the specified input files (.xml) and stored as multi-indexed `pandas.Series`²⁰³.
 - This layer provides also interface to input data and filtering of the input data.
- sL3 - Data presentation - This layer generates the elements specified in the specification file:
 - Tables: .csv files linked to static .rst files.
 - Plots: .html files generated using plot.ly linked to static .rst files.
- sL4 - Report generation - Sphinx generates required formats and versions:
 - formats: html, pdf
 - versions: minimal, full (TODO: define the names and scope of versions)

pal_layers.pdf

²⁰³ <https://pandas.pydata.org/pandas-docs/stable/generated/pandas.Series.html>

7.3.2 Data

Report Specification

The report specification file defines which data is used and which outputs are generated. It is human readable and structured. It is easy to add / remove / change items. The specification includes:

- Specification of the environment.
- Configuration of debug mode (optional).
- Specification of input data (jobs, builds, files, ...).
- Specification of the output.
- What and how is generated: - What: plots, tables. - How: specification of all properties and parameters.
- .yaml format.

Structure of the specification file

The specification file is organized as a list of dictionaries distinguished by the type:

```
-  
  type: "environment"  
-  
  type: "configuration"  
-  
  type: "debug"  
-  
  type: "static"  
-  
  type: "input"  
-  
  type: "output"  
-  
  type: "table"  
-  
  type: "plot"  
-  
  type: "file"
```

Each type represents a section. The sections "environment", "debug", "static", "input" and "output" are listed only once in the specification; "table", "file" and "plot" can be there multiple times.

Sections "debug", "table", "file" and "plot" are optional.

Table(s), files(s) and plot(s) are referred as "elements" in this text. It is possible to define and implement other elements if needed.

Section: Environment

This section has the following parts:

- type: "environment" - says that this is the section "environment".
- configuration - configuration of the PAL.
- paths - paths used by the PAL.
- urls - urls pointing to the data sources.
- make-dirs - a list of the directories to be created by the PAL while preparing the environment.

- remove-dirs - a list of the directories to be removed while cleaning the environment.
- build-dirs - a list of the directories where the results are stored.

The structure of the section “Environment” is as follows (example):

```

-
type: "environment"
configuration:
  # Debug mode:
  # - Skip:
  #   - Download of input data files
  # - Do:
  #   - Read data from given zip / xml files
  #   - Set the configuration as it is done in normal mode
  # If the section "type: debug" is missing, CFG[DEBUG] is set to 0.
  CFG[DEBUG]: 0

paths:
  # Top level directories:
  ## Working directory
  DIR[WORKING]: "_tmp"
  ## Build directories
  DIR[BUILD,HTML]: "_build"
  DIR[BUILD,LATEX]: "_build_latex"

  # Static .rst files
  DIR[RST]: "../.../docs/report"

  # Working directories
  ## Input data files (.zip, .xml)
  DIR[WORKING,DATA]: "${DIR[WORKING]}/data"
  ## Static source files from git
  DIR[WORKING,SRC]: "${DIR[WORKING]}/src"
  DIR[WORKING,SRC,STATIC]: "${DIR[WORKING,SRC]}/_static"

  # Static html content
  DIR[STATIC]: "${DIR[BUILD,HTML]}/_static"
  DIR[STATIC,VPP]: "${DIR[STATIC]}/vpp"
  DIR[STATIC,DPDK]: "${DIR[STATIC]}/dpdk"
  DIR[STATIC,ARCH]: "${DIR[STATIC]}/archive"

  # Detailed test results
  DIR[DTR]: "${DIR[WORKING,SRC]}/detailed_test_results"
  DIR[DTR,PERF,DPDK]: "${DIR[DTR]}/dpdk_performance_results"
  DIR[DTR,PERF,VPP]: "${DIR[DTR]}/vpp_performance_results"
  DIR[DTR,PERF,HC]: "${DIR[DTR]}/honeycomb_performance_results"
  DIR[DTR,FUNC,VPP]: "${DIR[DTR]}/vpp_functional_results"
  DIR[DTR,FUNC,HC]: "${DIR[DTR]}/honeycomb_functional_results"
  DIR[DTR,FUNC,NSHSFC]: "${DIR[DTR]}/nshsfc_functional_results"
  DIR[DTR,PERF,VPP,IMPRV]: "${DIR[WORKING,SRC]}/vpp_performance_tests/performance_improvements"

  # Detailed test configurations
  DIR[DTC]: "${DIR[WORKING,SRC]}/test_configuration"
  DIR[DTC,PERF,VPP]: "${DIR[DTC]}/vpp_performance_configuration"
  DIR[DTC,FUNC,VPP]: "${DIR[DTC]}/vpp_functional_configuration"

  # Detailed tests operational data
  DIR[DTO]: "${DIR[WORKING,SRC]}/test_operational_data"
  DIR[DTO,PERF,VPP]: "${DIR[DTO]}/vpp_performance_operational_data"

  # .css patch file to fix tables generated by Sphinx
  DIR[CSS_PATCH_FILE]: "${DIR[STATIC]}/theme_overrides.css"

```

(continues on next page)

(continued from previous page)

```

DIR[CSS_PATCH_FILE2]: "${DIR[WORKING, SRC, STATIC]}/theme_overrides.css"

urls:
  URL[JENKINS,CSIT]: "https://jenkins.fd.io/view/csit/job"
  URL[JENKINS,HC]: "https://jenkins.fd.io/view/hc2vpp/job"

make-dirs:
  # List the directories which are created while preparing the environment.
  # All directories MUST be defined in "paths" section.
  - "DIR[WORKING, DATA]"
  - "DIR[STATIC, VPP]"
  - "DIR[STATIC, DPDK]"
  - "DIR[STATIC, ARCH]"
  - "DIR[BUILD, LATEX]"
  - "DIR[WORKING, SRC]"
  - "DIR[WORKING, SRC, STATIC]"

remove-dirs:
  # List the directories which are deleted while cleaning the environment.
  # All directories MUST be defined in "paths" section.
  #- "DIR[BUILD, HTML]"

build-dirs:
  # List the directories where the results (build) is stored.
  # All directories MUST be defined in "paths" section.
  - "DIR[BUILD, HTML]"
  - "DIR[BUILD, LATEX]"

```

It is possible to use defined items in the definition of other items, e.g.:

```
DIR[WORKING, DATA]: "${DIR[WORKING]}/data"
```

will be automatically changed to

```
DIR[WORKING, DATA]: "_tmp/data"
```

Section: Configuration

This section specifies the groups of parameters which are repeatedly used in the elements defined later in the specification file. It has the following parts:

- data sets - Specification of data sets used later in element's specifications to define the input data.
- plot layouts - Specification of plot layouts used later in plots' specifications to define the plot layout.

The structure of the section "Configuration" is as follows (example):

```

-
  type: "configuration"
  data-sets:
    plot-vpp-throughput-latency:
      csit-vpp-perf-1710-all:
        - 11
        - 12
        - 13
        - 14
        - 15
        - 16
        - 17
        - 18

```

(continues on next page)

(continued from previous page)

```

- 19
- 20
vpp-perf-results:
  csit-vpp-perf-1710-all:
    - 20
    - 23
plot-layouts:
plot-throughput:
  xaxis:
    autorange: True
    autotick: False
    fixedrange: False
    gridcolor: "rgb(238, 238, 238)"
    linecolor: "rgb(238, 238, 238)"
    linewidth: 1
    showgrid: True
    showline: True
    showticklabels: True
    tickcolor: "rgb(238, 238, 238)"
    tickmode: "linear"
    title: "Indexed Test Cases"
    zeroline: False
  yaxis:
    gridcolor: "rgb(238, 238, 238)"
    hoverformat: ".4s"
    linecolor: "rgb(238, 238, 238)"
    linewidth: 1
    range: []
    showgrid: True
    showline: True
    showticklabels: True
    tickcolor: "rgb(238, 238, 238)"
    title: "Packets Per Second [pps]"
    zeroline: False
  boxmode: "group"
  boxgroupgap: 0.5
  autosize: False
  margin:
    t: 50
    b: 20
    l: 50
    r: 20
  showlegend: True
  legend:
    orientation: "h"
  width: 700
  height: 1000

```

The definitions from this sections are used in the elements, e.g.:

```

-
  type: "plot"
  title: "VPP Performance 64B-1t1c-(eth|dot1q|dot1ad)-(l2xcbase|l2bdbasemac|rn)-ndrdisc"
  algorithm: "plot_performance_box"
  output-file-type: ".html"
  output-file: "{DIR[STATIC,VPP]}/64B-1t1c-l2-se11-ndrdisc"
  data:
    "plot-vpp-throughput-latency"
  filter: "'64B' and ('BASE' or 'SCALE') and 'NDRDISC' and '1T1C' and ('L2BDMACSTAT' or 'L2BDMACLRN'
↪ or 'L2XCFWD') and not 'VHOST'"
  parameters:

```

(continues on next page)

(continued from previous page)

```

- "throughput"
- "parent"
traces:
  hoverinfo: "x+y"
  boxpoints: "outliers"
  whiskerwidth: 0
layout:
  title: "64B-1t1c-(eth|dot1q|dot1ad)-(l2xcbase|l2bdbasemaclrn)-ndrdisc"
  layout:
    "plot-throughput"

```

Section: Debug mode

This section is optional as it configures the debug mode. It is used if one does not want to download input data files and use local files instead.

If the debug mode is configured, the "input" section is ignored.

This section has the following parts:

- type: "debug" - says that this is the section "debug".
- general:
 - input-format - xml or zip.
 - extract - if "zip" is defined as the input format, this file is extracted from the zip file, otherwise this parameter is ignored.
- builds - list of builds from which the data is used. Must include a job name as a key and then a list of builds and their output files.

The structure of the section "Debug" is as follows (example):

```

-
type: "debug"
general:
  input-format: "zip" # zip or xml
  extract: "robot-plugin/output.xml" # Only for zip
builds:
  # The files must be in the directory DIR[WORKING,DATA]
  csit-dpdk-perf-1707-all:
  -
    build: 10
    file: "csit-dpdk-perf-1707-all__10.xml"
  -
    build: 9
    file: "csit-dpdk-perf-1707-all__9.xml"
  csit-nsh_sfc-verify-func-1707-ubuntu1604-v1r1:
  -
    build: 2
    file: "csit-nsh_sfc-verify-func-1707-ubuntu1604-v1r1-2.xml"
  csit-vpp-functional-1707-ubuntu1604-v1r1:
  -
    build: lastSuccessfulBuild
    file: "csit-vpp-functional-1707-ubuntu1604-v1r1-lastSuccessfulBuild.xml"
  hc2vpp-csit-integration-1707-ubuntu1604:
  -
    build: lastSuccessfulBuild
    file: "hc2vpp-csit-integration-1707-ubuntu1604-lastSuccessfulBuild.xml"
  csit-vpp-perf-1707-all:
  -

```

(continues on next page)

(continued from previous page)

```

build: 16
file: "csit-vpp-perf-1707-all__16__output.xml"
-
build: 17
file: "csit-vpp-perf-1707-all__17__output.xml"

```

Section: Static

This section defines the static content which is stored in git and will be used as a source to generate the report.

This section has these parts:

- type: "static" - says that this section is the "static".
- src-path - path to the static content.
- dst-path - destination path where the static content is copied and then processed.

::

- type: "static" src-path: "{DIR[RST]}" dst-path: "{DIR[WORKING,SRC]}"

Section: Input

This section defines the data used to generate elements. It is mandatory if the debug mode is not used.

This section has the following parts:

- type: "input" - says that this section is the "input".
- general - parameters common to all builds:
 - file-name: file to be downloaded.
 - file-format: format of the downloaded file, ".zip" or ".xml" are supported.
 - download-path: path to be added to url pointing to the file, e.g.: "{job}/{build}/robot/report/zip/{filename}"; {job}, {build} and {filename} are replaced by proper values defined in this section.
 - extract: file to be extracted from downloaded zip file, e.g.: "output.xml"; if xml file is downloaded, this parameter is ignored.
- builds - list of jobs (keys) and numbers of builds which output data will be downloaded.

The structure of the section "Input" is as follows (example from 17.07 report):

```

-
type: "input" # Ignored in debug mode
general:
  file-name: "robot-plugin.zip"
  file-format: ".zip"
  download-path: "{job}/{build}/robot/report/*zip*/{filename}"
  extract: "robot-plugin/output.xml"
builds:
  csit-vpp-perf-1707-all:
    - 9
    - 10
    - 13
    - 14
    - 15
    - 16

```

(continues on next page)

(continued from previous page)

```
- 17
- 18
- 19
- 21
- 22
csit-dpdk-perf-1707-all:
- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10
csit-vpp-functional-1707-ubuntu1604-v1r1:
- lastSuccessfulBuild
hc2vpp-csit-perf-master-ubuntu1604:
- 8
- 9
hc2vpp-csit-integration-1707-ubuntu1604:
- lastSuccessfulBuild
csit-nsh_sfc-verify-func-1707-ubuntu1604-v1r1:
- 2
```

Section: Output

This section specifies which format(s) will be generated (html, pdf) and which versions will be generated for each format.

This section has the following parts:

- type: “output” - says that this section is the “output”.
- format: html or pdf.
- version: defined for each format separately.

The structure of the section “Output” is as follows (example):

```
-
  type: "output"
  format:
    html:
      - full
    pdf:
      - full
      - minimal
```

TODO: define the names of versions

Content of “minimal” version

TODO: define the name and content of this version

Section: Table

This section defines a table to be generated. There can be 0 or more “table” sections.

This section has the following parts:

- type: “table” - says that this section defines a table.
- title: Title of the table.
- algorithm: Algorithm which is used to generate the table. The other parameters in this section must provide all information needed by the used algorithm.
- template: (optional) a .csv file used as a template while generating the table.
- output-file-ext: extension of the output file.
- output-file: file which the table will be written to.
- columns: specification of table columns:
 - title: The title used in the table header.
 - data: Specification of the data, it has two parts - command and arguments:
 - * command:
 - template - take the data from template, arguments:
 - number of column in the template.
 - data - take the data from the input data, arguments:
 - jobs and builds which data will be used.
 - operation - performs an operation with the data already in the table, arguments:
 - operation to be done, e.g.: mean, stdev, relative_change (compute the relative change between two columns) and display number of data samples ~= number of test jobs. The operations are implemented in the utils.py TODO: Move from utils.py to e.g. operations.py
 - numbers of columns which data will be used (optional).
- data: Specify the jobs and builds which data is used to generate the table.
- filter: filter based on tags applied on the input data, if “template” is used, filtering is based on the template.
- parameters: Only these parameters will be put to the output data structure.

The structure of the section “Table” is as follows (example of “table_performance_improvements”):

```

-
type: "table"
title: "Performance improvements"
algorithm: "table_performance_improvements"
template: "{DIR[DTR,PERF,VPP,IMPRV]}/tpl_performance_improvements.csv"
output-file-ext: ".csv"
output-file: "{DIR[DTR,PERF,VPP,IMPRV]}/performance_improvements"
columns:
-
  title: "VPP Functionality"
  data: "template 1"
-
  title: "Test Name"
  data: "template 2"
-
  title: "VPP-16.09 mean [Mpps]"
  data: "template 3"
-
  title: "VPP-17.01 mean [Mpps]"
  data: "template 4"
-

```

(continues on next page)

(continued from previous page)

```

title: "VPP-17.04 mean [Mpps]"
data: "template 5"
-
title: "VPP-17.07 mean [Mpps]"
data: "data csit-vpp-perf-1707-all mean"
-
title: "VPP-17.07 stdev [Mpps]"
data: "data csit-vpp-perf-1707-all stdev"
-
title: "17.04 to 17.07 change [%]"
data: "operation relative_change 5 4"
data:
  csit-vpp-perf-1707-all:
    - 9
    - 10
    - 13
    - 14
    - 15
    - 16
    - 17
    - 18
    - 19
    - 21
filter: "template"
parameters:
- "throughput"

```

Example of "table_details" which generates "Detailed Test Results - VPP Performance Results":

```

-
type: "table"
title: "Detailed Test Results - VPP Performance Results"
algorithm: "table_details"
output-file-ext: ".csv"
output-file: "${DIR[WORKING]}/vpp_performance_results"
columns:
-
  title: "Name"
  data: "data test_name"
-
  title: "Documentation"
  data: "data test_documentation"
-
  title: "Status"
  data: "data test_msg"
data:
  csit-vpp-perf-1707-all:
    - 17
filter: "all"
parameters:
- "parent"
- "doc"
- "msg"

```

Example of "table_details" which generates "Test configuration - VPP Performance Test Configs":

```

-
type: "table"
title: "Test configuration - VPP Performance Test Configs"
algorithm: "table_details"
output-file-ext: ".csv"

```

(continues on next page)

(continued from previous page)

```

output-file: "{DIR[WORKING]}/vpp_test_configuration"
columns:
-
  title: "Name"
  data: "data name"
-
  title: "VPP API Test (VAT) Commands History - Commands Used Per Test Case"
  data: "data show-run"
data:
  csit-vpp-perf-1707-all:
    - 17
filter: "all"
parameters:
- "parent"
- "name"
- "show-run"

```

Section: Plot

This section defines a plot to be generated. There can be 0 or more “plot” sections.

This section has these parts:

- type: “plot” - says that this section defines a plot.
- title: Plot title used in the logs. Title which is displayed is in the section “layout”.
- output-file-type: format of the output file.
- output-file: file which the plot will be written to.
- algorithm: Algorithm used to generate the plot. The other parameters in this section must provide all information needed by plot.ly to generate the plot. For example:
 - traces
 - layout
 - These parameters are transparently passed to plot.ly.
- data: Specify the jobs and numbers of builds which data is used to generate the plot.
- filter: filter applied on the input data.
- parameters: Only these parameters will be put to the output data structure.

The structure of the section “Plot” is as follows (example of a plot showing throughput in a chart box-with-whiskers):

```

-
type: "plot"
title: "VPP Performance 64B-1t1c-(eth|dot1q|dot1ad)-(l2xcbase|l2bdbasemac|rn)-ndrdisc"
algorithm: "plot_performance_box"
output-file-type: ".html"
output-file: "{DIR[STATIC,VPP]}/64B-1t1c-l2-se11-ndrdisc"
data:
  csit-vpp-perf-1707-all:
    - 9
    - 10
    - 13
    - 14
    - 15
    - 16
    - 17

```

(continues on next page)

(continued from previous page)

```

- 18
- 19
- 21
# Keep this formatting, the filter is enclosed with " (quotation mark) and
# each tag is enclosed with ' (apostrophe).
filter: "'64B' and 'BASE' and 'NDRDISC' and '1T1C' and ('L2BDMACSTAT' or 'L2BDMACLRN' or 'L2XCFWD
↵') and not 'VHOST'"
parameters:
- "throughput"
- "parent"
traces:
  hoverinfo: "x+y"
  boxpoints: "outliers"
  whiskerwidth: 0
layout:
  title: "64B-1t1c-(eth|dot1q|dot1ad)-(l2xcbase|l2bdbasemaclrn)-ndrdisc"
  xaxis:
    autorange: True
    autotick: False
    fixedrange: False
    gridcolor: "rgb(238, 238, 238)"
    linecolor: "rgb(238, 238, 238)"
    linewidth: 1
    showgrid: True
    showline: True
    showticklabels: True
    tickcolor: "rgb(238, 238, 238)"
    tickmode: "linear"
    title: "Indexed Test Cases"
    zeroline: False
  yaxis:
    gridcolor: "rgb(238, 238, 238)"
    hoverformat: ".4s"
    linecolor: "rgb(238, 238, 238)"
    linewidth: 1
    range: []
    showgrid: True
    showline: True
    showticklabels: True
    tickcolor: "rgb(238, 238, 238)"
    title: "Packets Per Second [pps]"
    zeroline: False
  boxmode: "group"
  boxgroupgap: 0.5
  autosize: False
  margin:
    t: 50
    b: 20
    l: 50
    r: 20
  showlegend: True
  legend:
    orientation: "h"
  width: 700
  height: 1000

```

The structure of the section "Plot" is as follows (example of a plot showing latency in a box chart):

```

-
type: "plot"
title: "VPP Latency 64B-1t1c-(eth|dot1q|dot1ad)-(l2xcbase|l2bdbasemaclrn)-ndrdisc"

```

(continues on next page)

(continued from previous page)

```

algorithm: "plot_latency_box"
output-file-type: ".html"
output-file: "{DIR[STATIC,VPP]}/64B-1t1c-l2-sel1-ndrdisc-lat50"
data:
  csit-vpp-perf-1707-all:
    - 9
    - 10
    - 13
    - 14
    - 15
    - 16
    - 17
    - 18
    - 19
    - 21
  filter: "'64B' and 'BASE' and 'NDRDISC' and '1T1C' and ('L2BDMACSTAT' or 'L2BDMACLRN' or 'L2XCFWD
↔') and not 'VHOST'"
parameters:
  - "latency"
  - "parent"
traces:
  boxmean: False
layout:
  title: "64B-1t1c-(eth|dot1q|dot1ad)-(l2xcbase|l2bdbasemac1rn)-ndrdisc"
  xaxis:
    autorange: True
    autotick: False
    fixedrange: False
    gridcolor: "rgb(238, 238, 238)"
    linecolor: "rgb(238, 238, 238)"
    linewidth: 1
    showgrid: True
    showline: True
    showticklabels: True
    tickcolor: "rgb(238, 238, 238)"
    tickmode: "linear"
    title: "Indexed Test Cases"
    zeroline: False
  yaxis:
    gridcolor: "rgb(238, 238, 238)"
    hoverformat: ""
    linecolor: "rgb(238, 238, 238)"
    linewidth: 1
    range: []
    showgrid: True
    showline: True
    showticklabels: True
    tickcolor: "rgb(238, 238, 238)"
    title: "Latency min/avg/max [uSec]"
    zeroline: False
  boxmode: "group"
  boxgroupgap: 0.5
  autosize: False
margin:
  t: 50
  b: 20
  l: 50
  r: 20
showlegend: True
legend:
  orientation: "h"

```

(continues on next page)

```
width: 700
height: 1000
```

The structure of the section “Plot” is as follows (example of a plot showing VPP HTTP server performance in a box chart with pre-defined data “plot-vpp-http-server-performance” set and plot layout “plot-cps”):

```
-
type: "plot"
title: "VPP HTTP Server Performance"
algorithm: "plot_http_server_performance_box"
output-file-type: ".html"
output-file: "{DIR[STATIC,VPP]}/http-server-performance-cps"
data:
  "plot-vpp-http-server-performance"
# Keep this formatting, the filter is enclosed with " (quotation mark) and
# each tag is enclosed with ' (apostrophe).
filter: "'HTTP' and 'TCP_CPS'"
parameters:
- "result"
- "name"
traces:
  hoverinfo: "x+y"
  boxpoints: "outliers"
  whiskerwidth: 0
layout:
  title: "VPP HTTP Server Performance"
  layout:
    "plot-cps"
```

Section: file

This section defines a file to be generated. There can be 0 or more “file” sections.

This section has the following parts:

- type: “file” - says that this section defines a file.
- title: Title of the table.
- algorithm: Algorithm which is used to generate the file. The other parameters in this section must provide all information needed by the used algorithm.
- output-file-ext: extension of the output file.
- output-file: file which the file will be written to.
- file-header: The header of the generated .rst file.
- dir-tables: The directory with the tables.
- data: Specify the jobs and builds which data is used to generate the table.
- filter: filter based on tags applied on the input data, if “all” is used, no filtering is done.
- parameters: Only these parameters will be put to the output data structure.
- chapters: the hierarchy of chapters in the generated file.
- start-level: the level of the the top-level chapter.

The structure of the section “file” is as follows (example):

```

-
type: "file"
title: "VPP Performance Results"
algorithm: "file_test_results"
output-file-ext: ".rst"
output-file: "{DIR[DTR,PERF,VPP]}/vpp_performance_results"
file-header: "\n.. |br| raw:: html\n\n <br />\n\n.. |prein| raw:: html\n\n <pre>\n\n\n..
↪|preout| raw:: html\n\n </pre>\n\n"
dir-tables: "{DIR[DTR,PERF,VPP]}"
data:
  csit-vpp-perf-1707-all:
    - 22
filter: "all"
parameters:
- "name"
- "doc"
- "level"
data-start-level: 2 # 0, 1, 2, ...
chapters-start-level: 2 # 0, 1, 2, ...

```

Static content

- Manually created / edited files.
- .rst files, static .csv files, static pictures (.svg), ...
- Stored in CSIT git repository.

No more details about the static content in this document.

Data to process

The PAL processes tests results and other information produced by Jenkins jobs. The data are now stored as robot results in Jenkins (TODO: store the data in nexus) either as .zip and / or .xml files.

7.3.3 Data processing

As the first step, the data are downloaded and stored locally (typically on a Jenkins slave). If .zip files are used, the given .xml files are extracted for further processing.

Parsing of the .xml files is performed by a class derived from “robot.api.ResultVisitor”, only necessary methods are overridden. All and only necessary data is extracted from .xml file and stored in a structured form.

The parsed data are stored as the multi-indexed pandas.Series data type. Its structure is as follows:

```

<job name>
  <build>
    <metadata>
    <suites>
    <tests>

```

“job name”, “build”, “metadata”, “suites”, “tests” are indexes to access the data. For example:

```

data =
job 1 name:
  build 1:
    metadata: metadata

```

(continues on next page)

(continued from previous page)

```

suites: suites
tests: tests
...
build N:
  metadata: metadata
  suites: suites
  tests: tests
...
job M name:
  build 1:
    metadata: metadata
    suites: suites
    tests: tests
    ...
  build N:
    metadata: metadata
    suites: suites
    tests: tests

```

Using indexes data["job 1 name"]["build 1"]["tests"] (e.g.: data["csit-vpp-perf-1704-all"]["17"]["tests"]) we get a list of all tests with all tests data.

Data will not be accessible directly using indexes, but using getters and filters.

Structure of metadata:

```

"metadata": {
  "version": "VPP version",
  "job": "Jenkins job name"
  "build": "Information about the build"
},

```

Structure of suites:

```

"suites": {
  "Suite name 1": {
    "doc": "Suite 1 documentation"
    "parent": "Suite 1 parent"
  }
  "Suite name N": {
    "doc": "Suite N documentation"
    "parent": "Suite N parent"
  }
}

```

Structure of tests:

Performance tests:

```

"tests": {
  "ID": {
    "name": "Test name",
    "parent": "Name of the parent of the test",
    "doc": "Test documentation"
    "msg": "Test message"
    "tags": ["tag 1", "tag 2", "tag n"],
    "type": "PDR" | "NDR",
    "throughput": {
      "value": int,
      "unit": "pps" | "bps" | "percentage"
    },
  },
  "latency": {
    "direction1": {

```

(continues on next page)

(continued from previous page)

```

    "100": {
      "min": int,
      "avg": int,
      "max": int
    },
    "50": { # Only for NDR
      "min": int,
      "avg": int,
      "max": int
    },
    "10": { # Only for NDR
      "min": int,
      "avg": int,
      "max": int
    }
  },
  "direction2": {
    "100": {
      "min": int,
      "avg": int,
      "max": int
    },
    "50": { # Only for NDR
      "min": int,
      "avg": int,
      "max": int
    },
    "10": { # Only for NDR
      "min": int,
      "avg": int,
      "max": int
    }
  }
},
"lossTolerance": "lossTolerance" # Only for PDR
"vat-history": "DUT1 and DUT2 VAT History"
},
"show-run": "Show Run"
},
"ID" {
  # next test
}

```

Functional tests:

```

"tests": {
  "ID": {
    "name": "Test name",
    "parent": "Name of the parent of the test",
    "doc": "Test documentation"
    "msg": "Test message"
    "tags": ["tag 1", "tag 2", "tag n"],
    "vat-history": "DUT1 and DUT2 VAT History"
    "show-run": "Show Run"
    "status": "PASS" | "FAIL"
  },
  "ID" {
    # next test
  }
}

```

Note: ID is the lowercase full path to the test.

Data filtering

The first step when generating an element is getting the data needed to construct the element. The data are filtered from the processed input data.

The data filtering is based on:

- job name(s).
- build number(s).
- tag(s).
- required data - only this data is included in the output.

WARNING: The filtering is based on tags, so be careful with tagging.

For example, the element which specification includes:

```
data:
  csit-vpp-perf-1707-all:
    - 9
    - 10
    - 13
    - 14
    - 15
    - 16
    - 17
    - 18
    - 19
    - 21
filter:
  - "'64B' and 'BASE' and 'NDRDISC' and '1T1C' and ('L2BDMACSTAT' or 'L2BDMACLRN' or 'L2XCFWD') and_
↳not 'VHOST'"
```

will be constructed using data from the job “csit-vpp-perf-1707-all”, for all listed builds and the tests with the list of tags matching the filter conditions.

The output data structure for filtered test data is:

```
- job 1
  - build 1
    - test 1
      - parameter 1
      - parameter 2
      ...
      - parameter n
      ...
    - test n
    ...
  - build n
  ...
- job n
```

Data analytics

Data analytics part implements:

- methods to compute statistical data from the filtered input data.
- trending.

Throughput Speedup Analysis - Multi-Core with Multi-Threading

Throughput Speedup Analysis (TSA) calculates throughput speedup ratios for tested 1-, 2- and 4-core multi-threaded VPP configurations using the following formula:

$$N_core_throughput_speedup = \frac{N_core_throughput}{1_core_throughput}$$

Multi-core throughput speedup ratios are plotted in grouped bar graphs for throughput tests with 64B/78B frame size, with number of cores on X-axis and speedup ratio on Y-axis.

For better comparison multiple test results' data sets are plotted per each graph:

- graph type: grouped bars;
- graph X-axis: (testcase index, number of cores);
- graph Y-axis: speedup factor.

Subset of existing performance tests is covered by TSA graphs.

Model for TSA:

```
-
type: "plot"
title: "TSA: 64B-*- (eth|dot1q|dot1ad)-(l2xcbase|l2bdbasemaclrn)-ndrdisc"
algorithm: "plot_throughput_speedup_analysis"
output-file-type: ".html"
output-file: "${DIR[STATIC,VPP]}/10ge2p1x520-64B-l2-tsa-ndrdisc"
data:
  "plot-throughput-speedup-analysis"
filter: "'NIC_Intel-X520-DA2' and '64B' and 'BASE' and 'NDRDISC' and ('L2BDMACSTAT' or 'L2BDMACLRN
↳ or 'L2XCFWD') and not 'VHOST'"
parameters:
- "throughput"
- "parent"
- "tags"
layout:
  title: "64B-*- (eth|dot1q|dot1ad)-(l2xcbase|l2bdbasemaclrn)-ndrdisc"
  layout:
    "plot-throughput-speedup-analysis"
```

Comparison of results from two sets of the same test executions

This algorithm enables comparison of results coming from two sets of the same test executions. It is used to quantify performance changes across all tests after test environment changes e.g. Operating System upgrades/patches, Hardware changes.

It is assumed that each set of test executions includes multiple runs of the same tests, 10 or more, to verify test results repeatability and to yield statistically meaningful results data.

Comparison results are presented in a table with a specified number of the best and the worst relative changes between the two sets. Following table columns are defined:

- name of the test;
- throughput mean values of the reference set;
- throughput standard deviation of the reference set;
- throughput mean values of the set to compare;
- throughput standard deviation of the set to compare;

- relative change of the mean values.

The model

The model specifies:

- type: "table" - means this section defines a table.
- title: Title of the table.
- algorithm: Algorithm which is used to generate the table. The other parameters in this section must provide all information needed by the used algorithm.
- output-file-ext: Extension of the output file.
- output-file: File which the table will be written to.
- reference - the builds which are used as the reference for comparison.
- compare - the builds which are compared to the reference.
- data: Specify the sources, jobs and builds, providing data for generating the table.
- filter: Filter based on tags applied on the input data, if "template" is used, filtering is based on the template.
- parameters: Only these parameters will be put to the output data structure.
- nr-of-tests-shown: Number of the best and the worst tests presented in the table. Use 0 (zero) to present all tests.

Example:

```
-
type: "table"
title: "Performance comparison"
algorithm: "table_performance_comparison"
output-file-ext: ".csv"
output-file: "{DIR[DTR,PERF,VPP,IMPRV]}/vpp_performance_comparison"
reference:
  title: "csit-vpp-perf-1801-all - 1"
  data:
    csit-vpp-perf-1801-all:
      - 1
      - 2
compare:
  title: "csit-vpp-perf-1801-all - 2"
  data:
    csit-vpp-perf-1801-all:
      - 1
      - 2
data:
  "vpp-perf-comparison"
filter: "all"
parameters:
- "name"
- "parent"
- "throughput"
nr-of-tests-shown: 20
```

Advanced data analytics

In the future advanced data analytics (ADA) will be added to analyze the telemetry data collected from SUT telemetry sources and correlate it to performance test results.

TODO

- describe the concept of ADA.
- add specification.

7.3.4 Data presentation

Generates the plots and tables according to the report models per specification file. The elements are generated using algorithms and data specified in their models.

Tables

- tables are generated by algorithms implemented in PAL, the model includes the algorithm and all necessary information.
- output format: csv
- generated tables are stored in specified directories and linked to .rst files.

Plots

- plot.ly²⁰⁴ is currently used to generate plots, the model includes the type of plot and all the necessary information to render it.
- output format: html.
- generated plots are stored in specified directories and linked to .rst files.

7.3.5 Report generation

Report is generated using Sphinx and Read_the_Docs template. PAL generates html and pdf formats. It is possible to define the content of the report by specifying the version (TODO: define the names and content of versions).

The process

1. Read the specification.
2. Read the input data.
3. Process the input data.
4. For element (plot, table, file) defined in specification:
 - (a) Get the data needed to construct the element using a filter.
 - (b) Generate the element.
 - (c) Store the element.
5. Generate the report.
6. Store the report (Nexus).

The process is model driven. The elements' models (tables, plots, files and report itself) are defined in the specification file. Script reads the elements' models from specification file and generates the elements.

It is easy to add elements to be generated in the report. If a new type of an element is required, only a new algorithm needs to be implemented and integrated.

²⁰⁴ <https://plot.ly/>

7.3.6 Continuous Performance Measurements and Trending

Performance analysis and trending execution sequence:

CSIT PA runs performance analysis, change detection and trending using specified trend analysis metrics over the rolling window of last <N> sets of historical measurement data. PA is defined as follows:

1. PA job triggers:
 - (a) By PT job at its completion.
 - (b) Manually from Jenkins UI.
2. Download and parse archived historical data and the new data:
 - (a) New data from latest PT job is evaluated against the rolling window of <N> sets of historical data.
 - (b) Download RF output.xml files and compressed archived data.
 - (c) Parse out the data filtering test cases listed in PA specification (part of CSIT PAL specification file).
3. Calculate trend metrics for the rolling window of <N> sets of historical data:
 - (a) Calculate quartiles Q1, Q2, Q3.
 - (b) Trim outliers using IQR.
 - (c) Calculate TMA and TMSD.
 - (d) Calculate normal trending range per test case based on TMA and TMSD.
4. Evaluate new test data against trend metrics:
 - (a) If within the range of $(TMA \pm 3 * TMSD) \Rightarrow$ Result = Pass, Reason = Normal.
 - (b) If below the range \Rightarrow Result = Fail, Reason = Regression.
 - (c) If above the range \Rightarrow Result = Pass, Reason = Progression.
5. Generate and publish results
 - (a) Relay evaluation result to job result.
 - (b) Generate a new set of trend analysis summary graphs and drill-down graphs.
 - i. Summary graphs to include measured values with Normal, Progression and Regression markers. MM shown in the background if possible.
 - ii. Drill-down graphs to include MM, TMA and TMSD.
 - (c) Publish trend analysis graphs in html format on <https://docs.fd.io/csit/master/trending/>.

Parameters to specify:

- job to be monitored - the Jenkins job which results are used as input data for this test;
- builds used for trending plot(s) - specified by a list of build numbers or by a range of builds defined by the first and the last build number;
- list plots to generate:
 - plot title;
 - output file name;
 - data for plots;
 - tests to be displayed in the plot defined by a filter;
 - list of parameters to extract from the data;

- periods (daily = 1, weekly = 5, monthly = 30);
- plot layout

Example:

```

-
type: "cpta"
title: "Continuous Performance Trending and Analysis"
algorithm: "cpta"
output-file-type: ".html"
output-file: "{DIR[STATIC,VPP]}/cpta"
data: "plot-performance-trending"
plots:
  - title: "VPP 1T1C L2 64B Packet Throughput - {period} Trending"
    output-file-name: "l2-1t1c-x520"
    data: "plot-performance-trending"
    filter: "'NIC_Intel-X520-DA2' and 'MRR' and '64B' and ('BASE' or 'SCALE') and '1T1C' and (
↪ 'L2BDMACSTAT' or 'L2BDMACLRN' or 'L2XCFWD') and not 'VHOST' and not 'MEMIF'"
    parameters:
      - "result"
#     - "name"
    periods:
      - 1
      - 5
      - 30
    layout: "plot-cpta"

  - title: "VPP 2T2C L2 64B Packet Throughput - {period} Trending"
    output-file-name: "l2-2t2c-x520"
    data: "plot-performance-trending"
    filter: "'NIC_Intel-X520-DA2' and 'MRR' and '64B' and ('BASE' or 'SCALE') and '2T2C' and (
↪ 'L2BDMACSTAT' or 'L2BDMACLRN' or 'L2XCFWD') and not 'VHOST' and not 'MEMIF'"
    parameters:
      - "result"
#     - "name"
    periods:
      - 1
      - 5
      - 30
    layout: "plot-cpta"

```

7.3.7 API

List of modules, classes, methods and functions

```

specification_parser.py

class Specification

    Methods:
        read_specification
        set_input_state
        set_input_file_name

    Getters:
        specification
        environment
        debug
        is_debug
        input

```

(continues on next page)

(continued from previous page)

```
        builds
        output
        tables
        plots
        files
        static

input_data_parser.py

class InputData

    Methods:
        read_data
        filter_data

    Getters:
        data
        metadata
        suites
        tests

environment.py

    Functions:
        clean_environment

class Environment

    Methods:
        set_environment

    Getters:
        environment

input_data_files.py

    Functions:
        download_data_files
        unzip_files

generator_tables.py

    Functions:
        generate_tables

    Functions implementing algorithms to generate particular types of
    tables (called by the function "generate_tables"):
        table_details
        table_performance_improvements

generator_plots.py

    Functions:
        generate_plots

    Functions implementing algorithms to generate particular types of
```

(continues on next page)

(continued from previous page)

```
plots (called by the function "generate_plots"):  
    plot_performance_box  
    plot_latency_box
```

```
generator_files.py
```

```
Functions:  
    generate_files
```

```
Functions implementing algorithms to generate particular types of  
files (called by the function "generate_files"):  
    file_test_results
```

```
report.py
```

```
Functions:  
    generate_report
```

```
Functions implementing algorithms to generate particular types of  
report (called by the function "generate_report"):  
    generate_html_report  
    generate_pdf_report
```

```
Other functions called by the function "generate_report":  
    archive_input_data  
    archive_report
```


7.4 CSIT TAGs Descriptions

All CSIT test cases are labelled with Robot Framework tags used to allow for easy test case type identification, test case grouping and selection for execution. Following sections list currently used CSIT TAGs and their documentation based on the content of [tag documentation rst file](#)²⁰⁵.

7.4.1 Topology TAGs

3_NODE_DOUBLE_LINK_TOPO

3 nodes connected in a circular topology with two links interconnecting the devices.

3_NODE_SINGLE_LINK_TOPO

3 nodes connected in a circular topology with at least one link interconnecting devices.

7.4.2 Objective TAGs

SKIP_PATCH

Test case(s) marked to not run in case of vpp-csit-verify (i.e. VPP patch) and csit-vpp-verify jobs (i.e. CSIT patch).

SKIP_VPP_PATCH

Test case(s) marked to not run in case of vpp-csit-verify (i.e. VPP patch).

7.4.3 Environment TAGs

HW_ENV

DUTs and TGs are running on bare metal.

VM_ENV

DUTs and TGs are running in virtual environment.

VPP_VM_ENV

DUTs with VPP and capable of running Virtual Machine.

7.4.4 NIC model tags

NIC_Intel-X520-DA2

Intel X520-DA2 NIC.

²⁰⁵ https://git.fd.io/csit/tree/docs/tag_documentation.rst?h=rls1804

NIC_Intel-XL710

Intel XL710 NIC.

NIC_Intel-X710

Intel X710 NIC.

NIC_Cisco-VIC-1227

VIC-1227 by Cisco.

NIC_Cisco-VIC-1385

VIC-1385 by Cisco.

7.4.5 Scaling TAGs

FIB_20K

2x10,000 entries in single fib table

FIB_200K

2x100,000 entries in single fib table

FIB_2M

2x1,000,000 entries in single fib table

TNL_1000

IPSec in tunnel mode - 1000 tunnels.

SRC_USER_10

Traffic flow with 10 unique IPs (users) in one direction.

SRC_USER_100

Traffic flow with 100 unique IPs (users) in one direction.

SRC_USER_1000

Traffic flow with 1000 unique IPs (users) in one direction.

SRC_USER_2000

Traffic flow with 2000 unique IPs (users) in one direction.

SRC_USER_4000

Traffic flow with 4000 unique IPs (users) in one direction.

100_FLOWS

Traffic stream with 100 unique flows (10 IPs/users x 10 UDP ports) in one direction.

10k_FLOWS

Traffic stream with 10 000 unique flows (10 IPs/users x 1000 UDP ports) in one direction.

100k_FLOWS

Traffic stream with 100 000 unique flows (100 IPs/users x 1000 UDP ports) in one direction.

7.4.6 Tags marking functional vs. performance of tests**FUNCTEST**

All functional test cases.

PERFTEST

All performance test cases.

7.4.7 Performance testing tags**PDRDISC**

Partial Drop Rate evaluation of single run result, with non-zero packet loss tolerance (LT) expressed in percentage of packets transmitted.

NDRDISC

Non Drop Rate evaluation of results. Loss acceptance of dropped packets is set to zero lost packets.

NDRPDR

Single test finding both No Drop Rate and Partial Drop Rate simultaneously. The search is done by optimized algorithm which performs multiple trial runs at different durations and transmit rates.

NDRCHK

Performance tests where TG verifies DUTs' throughput at ref-NDR (reference Non Drop Rate) with zero packet loss tolerance.

PDRCHK

Performance tests where TG verifies DUTs' throughput at ref-PDR (reference Partial Drop Rate) with 0.5% loss tolerance.

MRR

Performance tests where TG sends the traffic at maximum rate (line rate) and reports total sent/received packets over performance trial duration.

NDRPDRDISC

Find performance of DUT based on **RFC 2544**²⁰⁶ with linear / binary / combined search. (Previous LONG tests.)

7.4.8 Ethernet frame size tags for performance tests

64B

64B frames used for test.

78B

78B frames used for test.

IMIX

IMIX frame sequence (28x 64B, 16x 570B, 4x 1518B) used for test.

1460B

1460B frames used for test.

1480B

1480B frames used for test.

1514B

1514B frames used for test.

1518B

1518B frames used for test.

9000B

9000B frames used for test.

7.4.9 Test type tags

BASE

Baseline test cases, no encapsulation, no feature(s) configured in tests.

IP4BASE

IPv4 baseline test cases, no encapsulation, no feature(s) configured in tests.

IP6BASE

IPv6 baseline test cases, no encapsulation, no feature(s) configured in tests.

²⁰⁶ <https://tools.ietf.org/html/rfc2544.html>

L2XCBASE

L2XC baseline test cases, no encapsulation, no feature(s) configured in tests.

L2BDBASE

L2BD baseline test cases, no encapsulation, no feature(s) configured in tests.

SCALE

Scale test cases.

ENCAP

Test cases where encapsulation is used. Use also encapsulation tag(s).

FEATURE

At least one feature is configured in test cases. Use also feature tag(s).

TLDK

Functional test cases for TLDK.

TCP

Tests which use TCP.

TCP_CPS

Performance tests which measure connections per second using http requests.

TCP_RPS

Performance tests which measure requests per second using http requests.

HTTP

Tests which use HTTP.

7.4.10 Forwarding mode tags**L2BDMACSTAT**

VPP L2 bridge-domain, L2 MAC static.

L2BDMACLRNRN

VPP L2 bridge-domain, L2 MAC learning.

L2XCFWD

VPP L2 point-to-point cross-connect.

IP4FWD

VPP IPv4 routed forwarding.

IP6FWD

VPP IPv6 routed forwarding.

7.4.11 Underlay tags

IP4UNRLAY

IPv4 underlay.

IP6UNRLAY

IPv6 underlay.

MPLSUNRLAY

MPLS underlay.

7.4.12 Overlay tags

L2OVLAY

L2 overlay.

IP4OVLAY

IPv4 overlay (IPv4 payload).

IP6OVLAY

IPv6 overlay (IPv6 payload).

7.4.13 Tagging tags

DOT1Q

All test cases with dot1q.

DOT1AD

All test cases with dot1ad.

7.4.14 Encapsulation tags

ETH

All test cases with base Ethernet (no encapsulation).

LISP

All test cases with LISP.

LISPGPE

All test cases with LISP-GPE.

VXLAN

All test cases with Vxlan.

VXLANGPE

All test cases with VXLAN-GPE.

GRE

All test cases with GRE.

IPSEC

All test cases with IPSEC.

SRv6

All test cases with Segment routing over IPv6 dataplane.

7.4.15 Interface tags**PHY**

All test cases which use physical interface(s).

VHOST

All test cases which uses VHOST.

VHOST_256

All test cases which uses VHOST with qemu queue size set to 256.

VHOST_1024

All test cases which uses VHOST with qemu queue size set to 1024.

CFS_OPT

All test cases which uses VM with optimised scheduler policy.

TUNTAP

All test cases which uses TUN and TAP.

AFPKT

All test cases which uses AFPKT.

NETMAP

All test cases which uses Netmap.

MEMIF

All test cases which uses Memif.

7.4.16 Feature tags

IACLDST

iACL destination.

COPWHLIST

COP whitelist.

NAT44

NAT44 configured and tested.

NAT64

NAT44 configured and tested.

ACL

ACL plugin configured and tested.

IACL

ACL plugin configured and tested on input path.

OACL

ACL plugin configured and tested on output path.

ACL_STATELESS

ACL plugin configured and tested in stateless mode (permit action).

ACL_STATEFUL

ACL plugin configured and tested in stateful mode (permit+reflect action).

ACL1

ACL plugin configured and tested with 1 not-hitting ACE.

ACL10

ACL plugin configured and tested with 10 not-hitting ACEs.

ACL50

ACL plugin configured and tested with 50 not-hitting ACEs.

SRv6_PROXY

SRv6 endpoint to SR-unaware appliance via proxy.

SRv6_PROXY_STAT

SRv6 endpoint to SR-unaware appliance via static proxy.

SRv6_PROXY_DYN

SRv6 endpoint to SR-unaware appliance via dynamic proxy.

SRv6_PROXY_MASQ

SRv6 endpoint to SR-unaware appliance via masquerading proxy.

7.4.17 Encryption tags**IPSECSW**

Crypto in software.

IPSECHW

Crypto in hardware.

IPSECTRAN

IPSec in transport mode.

IPSECTUN

IPSec in tunnel mode.

7.4.18 Client-workload tags**VM**

All test cases which use at least one virtual machine.

LXC

All test cases which use Linux container and LXC utils.

DOCKER

All test cases which use Docker as container manager.

APP

All test cases with specific APP use.

7.4.19 Container orchestration tags

K8S

All test cases which use Kubernetes for orchestration.

SFC_CONTROLLER

All test cases which use ligato/sfc_controller for driving configuration of vpp inside container.

VPP_AGENT

All test cases which use Golang implementation of a control/management plane for VPP

1VSWITCH

VPP running in Docker container acting as VSWITCH.

1VNF

1 VPP running in Docker container acting as VNF work load.

2VNF

2 VPP running in 2 Docker containers acting as VNF work load.

4VNF

4 VPP running in 4 Docker containers acting as VNF work load.

7.4.20 Multi-threading tags

STHREAD

All test cases using single poll mode thread.

MTHREAD

All test cases using more than one poll mode driver thread.

1NUMA

All test cases with packet processing on single socket.

2NUMA

All test cases with packet processing on two sockets.

SMT

All test cases with symmetric Multi-Threading (HyperThreading) enabled.

NOSMT

All test cases with symmetric Multi-Threading (HyperThreading) disabled.

1T1C

1 worker thread pinned to 1 dedicated physical core. 1 receive queue per interface. Main thread pinned to core 0.

2T2C

2 worker threads pinned to 2 dedicated physical cores. 1 receive queue per interface. Main thread pinned to core 0.

4T4C

4 worker threads pinned to 4 dedicated physical cores. 2 receive queues per interface. Main thread pinned to core 0.

6T6C

6 worker threads pinned to 6 dedicated physical cores. 3 receive queues per interface. Main thread pinned to core 0.

8T8C

8 worker threads pinned to 8 dedicated physical cores. 4 receive queues per interface. Main thread pinned to core 0.

7.4.21 Honeycomb tags**HC_FUNC**

Honeycomb functional test cases.

HC_NSH

Honeycomb NSH test cases.

HC_PERSIST

Honeycomb persistence test cases.

HC_REST_ONLY

(Exclusion tag) Honeycomb test cases that cannot be run in Netconf mode using ODL client for Restfconf
-> Netconf translation.

Bibliography

- [lxc] [Linux Containers](#)¹⁴⁶
- [lxc-namespace] [Resource management: Linux kernel Namespaces and cgroups](#)¹⁴⁷.
- [stgraber] [LXC 1.0: Blog post series](#)¹⁴⁸.
- [lxc-security] [Linux Containers Security](#)¹⁴⁹.
- [capabilities] [‘Linux manual - capabilities - overview of Linux capabilities http://man7.org/linux/man-pages/man7/capabilities.7.html’](#).
- [cgroup1] [Linux kernel documentation: cgroups](#)¹⁵⁰.
- [cgroup2] [Linux kernel documentation: Control Group v2](#)¹⁵¹.
- [selinux] [SELinux Project Wiki](#)¹⁵².
- [lxc-sec-features] [LXC 1.0: Security features](#)¹⁵³.
- [lxc-source] [Linux Containers source](#)¹⁵⁴.
- [apparmor] [Ubuntu AppArmor](#)¹⁵⁵.
- [seccomp] [SECure COMputing with filters](#)¹⁵⁶.
- [docker] [Docker](#)¹⁵⁷.
- [k8s-doc] [Kubernetes documentation](#)¹⁵⁸.
- [ligato] [Ligato](#)¹⁵⁹.
- [govpp] [FD.io goVPP project](#)¹⁶⁰.

¹⁴⁶ <https://linuxcontainers.org/>

¹⁴⁷ <https://www.cs.ucsb.edu/~rich/class/cs293b-cloud/papers/lxc-namespace.pdf>

¹⁴⁸ <https://stgraber.org/2013/12/20/lxc-1-0-blog-post-series/>

¹⁴⁹ <https://linuxcontainers.org/lxc/security/>

¹⁵⁰ <https://www.kernel.org/doc/Documentation/cgroup-v1/cgroups.txt>

¹⁵¹ <https://www.kernel.org/doc/Documentation/cgroup-v2.txt>

¹⁵² http://selinuxproject.org/page/Main_Page

¹⁵³ <https://stgraber.org/2014/01/01/lxc-1-0-security-features/>

¹⁵⁴ <https://github.com/lxc/lxc>

¹⁵⁵ <https://wiki.ubuntu.com/AppArmor>

¹⁵⁶ https://www.kernel.org/doc/Documentation/prctl/seccomp_filter.txt

¹⁵⁷ <https://www.docker.com/what-docker>

¹⁵⁸ <https://kubernetes.io/docs/home/>

¹⁵⁹ <https://github.com/ligato>

¹⁶⁰ <https://wiki.fd.io/view/GoVPP>

[vpp-agent] [Ligato vpp-agent](https://github.com/ligato/vpp-agent)¹⁶¹.

[image-var] Image parameter is required in initial commit version. There is plan to implement container build class to build Docker/LXC image.

¹⁶¹ <https://github.com/ligato/vpp-agent>