
CSIT REPORT

Release rls1908_1

Sep 25, 2019

Contents

1	Introduction	1
1.1	Report History	1
1.2	Report Structure	1
1.3	Test Scenarios	3
1.4	Physical Testbeds	4
1.5	Test Methodology	10
2	VPP Performance	31
2.1	Overview	31
2.2	Release Notes	38
2.3	Packet Throughput	40
2.4	Speedup Multi-Core	152
2.5	Packet Latency	260
2.6	Reconfiguration Tests	385
2.7	Comparisons	390
2.8	Throughput Trending	396
2.9	Test Environment	397
2.10	Documentation	421
3	VPP Device	429
3.1	Overview	429
3.2	Release Notes	433
3.3	Integration Tests	433
3.4	Documentation	442
4	VPP Functional	443
4.1	Overview	443
4.2	Release Notes	445
4.3	Test Environment	446
4.4	Documentation	454
5	CSIT Framework	455
5.1	Design	455
5.2	Test Naming	459
5.3	Presentation and Analytics	461
5.4	CSIT RF Tags Descriptions	488
	Bibliography	503

1.1 Report History

FD.io CSIT-1908.1 Report history and per .[ww] revision changes are listed below.

.[ww] Revision	Changes
.39	Initial version

FD.io CSIT Reports follow CSIT-[yy][mm].[ww] numbering format, with version denoted by concatenation of two digit year [yy] and two digit month [mm], and maintenance revision identified by two digit calendar week number [ww].

1.2 Report Structure

FD.io CSIT-1908.1 report contains system performance and functional testing data of VPP-19.08.1 release. [PDF version of this report](#)¹ is available for download.

CSIT-1908.1 report is structured as follows:

1. INTRODUCTION: General introduction to FD.io CSIT-1908.1.
 - **Introduction:** This section.
 - **Test Scenarios Overview:** A brief overview of test scenarios covered in this report.
 - **Physical Testbeds:** Description of physical testbeds.
 - **Test Methodology:** Performance benchmarking and functional test methodologies.
2. VPP PERFORMANCE: VPP performance tests executed in physical FD.io testbeds.
 - **Overview:** Tested logical topologies, test coverage and naming specifics.
 - **Release Notes:** Changes in CSIT-1908.1, added tests, environment or methodology changes, known issues.
 - **Packet Throughput:** NDR, PDR throughput graphs based on results from repeated same test job executions to verify repeatability of measurements.

¹ https://docs.fd.io/csit/rls1908_1/report/_static/archive/csit_rls1908_1.39.pdf

- **Speedup Multi-Core:** NDR, PDR throughput multi-core speedup graphs based on results from test job executions.
 - **Packet Latency:** Latency graphs based on results from test job executions.
 - **Soak Tests:** Long duration soak tests are executed using PLRsearch algorithm.
 - **NFV Service Density:** Network Function Virtualization (NFV) service density tests focus on measuring total per server throughput at varied NFV service “packing” densities with vswitch providing host dataplane.
 - **Comparisons:** Performance comparisons between VPP releases and between different testbed types.
 - **Throughput Trending:** References to continuous VPP performance trending.
 - **Test Environment:** Performance test environment configuration.
 - **Documentation:** Pointers to CSIT source code documentation for VPP performance tests.
3. DPDK PERFORMANCE: DPDK performance tests executed in physical FD.io testbeds.
- **Overview:** Tested logical topologies, test coverage.
 - **Release Notes:** Changes in CSIT-1908.1, known issues.
 - **Packet Throughput:** NDR, PDR throughput graphs based on results from repeated same test job executions to verify repeatability of measurements.
 - **Packet Latency:** Latency graphs based on results from test job executions.
 - **Comparisons:** Performance comparisons between DPDK releases and between different testbed types.
 - **Throughput Trending:** References to regular DPDK performance trending.
 - **Test Environment:** Performance test environment configuration.
 - **Documentation:** Pointers to CSIT source code documentation for DPDK performance tests.
4. VPP DEVICE: VPP functional tests executed in physical FD.io testbeds using containers.
- **Overview:** Tested virtual topologies, test coverage and naming specifics;
 - **Release Notes:** Changes in CSIT-1908.1, added tests, environment or methodology changes, known issues.
 - **Integration Tests:** Functional test environment configuration.
 - **Documentation:** Pointers to CSIT source code documentation for VPP functional tests.
5. VPP FUNCTIONAL: VPP functional tests executed in virtual FD.io testbeds.
- **Overview:** Tested virtual topologies, test coverage and naming specifics;
 - **Release Notes:** Changes in CSIT-1908.1, added tests, environment or methodology changes, known issues.
 - **Test Environment:** Functional test environment configuration.
 - **Documentation:** Pointers to CSIT source code documentation for VPP functional tests.
6. HONEYCOMB FUNCTIONAL: Honeycomb functional tests executed in virtual FD.io testbeds.
- **Overview:** Tested virtual topologies, test coverage and naming specifics;
 - **Release Notes:** Changes in CSIT-1908.1, known issues.
 - **Test Environment:** Functional test environment configuration.
 - **Documentation:** Pointers to CSIT source code documentation for Honeycomb functional tests.
7. DETAILED RESULTS: Detailed result tables auto-generated from CSIT test job executions using RF (Robot Framework) output files as sources.

- **VPP Performance NDR/PDR:** VPP NDR/PDR throughput and latency.
 - **VPP Performance MRR:** VPP MRR throughput.
 - **DPDK Performance:** DPDK Testpmd and L3fwd NDR/PDR throughput and latency.
 - **VPP Functional:** Detailed VPP functional results.
 - **Honeycomb Functional:** Detailed HoneyComb functional results.
8. **TEST CONFIGURATION:** VPP DUT configuration data based on VPP API Test (VAT) Commands History auto-generated from CSIT test job executions using RF output files as sources.
 - **VPP Performance NDR/PDR:** Configuration data.
 - **VPP Performance MRR:** Configuration data.
 - **VPP Functional:** Configuration data.
 9. **TEST OPERATIONAL DATA:** VPP DUT operational data auto-generated from CSIT test job executions using RFoutput files as sources.
 - **VPP Performance NDR/PDR:** VPP *show run* outputs under test load.
 10. **CSIT FRAMEWORK DOCUMENTATION:** Description of the overall FD.io CSIT framework.
 - **Design:** Framework modular design hierarchy.
 - **Test naming:** Test naming convention.
 - **Presentation and Analytics Layer:** Description of PAL CSIT analytics module.
 - **CSIT RF Tags Descriptions:** CSIT RF Tags used for test suite and test case grouping and selection.

1.3 Test Scenarios

FD.io CSIT-1908.1 report includes multiple test scenarios of VPP centric applications, topologies and use cases. In addition it also covers baseline tests of DPDK sample applications. Tests are executed in physical (performance tests) and virtual environments (functional tests).

Brief overview of test scenarios covered in this report:

1. **VPP Performance:** VPP performance tests are executed in physical FD.io testbeds, focusing on VPP network data plane performance in NIC-to-NIC switching topologies. Tested across Intel Xeon Haswell and Skylake servers, ARM, Denverton, range of NICs (10GE, 25GE, 40GE) and multi-thread/multi-core configurations. VPP application runs in bare-metal host user-mode handling NICs. TRex is used as a traffic generator.
2. **VPP Vhostuser Performance with KVM VMs:** VPP VM service switching performance tests using vhostuser virtual interface for interconnecting multiple NF-in-VM instances. VPP vswitch instance runs in bare-metal user-mode handling NICs and connecting over vhost-user interfaces to VM instances each running VPP with virtio virtual interfaces. Similarly to VPP Performance, tests are run across a range of configurations. TRex is used as a traffic generator.
3. **VPP Memif Performance with LXC and Docker Containers:** VPP Container service switching performance tests using memif virtual interface for interconnecting multiple VPP-in-container instances. VPP vswitch instance runs in bare-metal user-mode handling NICs and connecting over memif (Slave side) interfaces to more instances of VPP running in LXC or in Docker Containers, both with memif interfaces (Master side). Similarly to VPP Performance, tests are run across a range of configurations. TRex is used as a traffic generator.
4. **DPDK Performance:** VPP uses DPDK to drive the NICs and physical interfaces. DPDK performance tests are used as a baseline to profile performance of the DPDK sub-system. Two DPDK applications are tested: Testpmd and L3fwd. DPDK tests are executed in the same testing environment as VPP

tests. DPDK Testpmd and L3fwd applications run 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 both network data plane and in-line control plane. Tests cover vNIC-to-vNIC vNIC-to-nestedVM-to-vNIC forwarding topologies. Scapy is used as a traffic generator.

All CSIT test data included in this report is auto-generated from RF (Robot Framework) output.xml files produced by LF (Linux Foundation) FD.io Jenkins jobs executed against VPP-19.08.1 release artifacts. References are provided to the original FD.io Jenkins job results and all archived source files.

FD.io CSIT system is developed using two main coding platforms: RF and Python2.7. CSIT-1908.1 source code for the executed test suites is available in CSIT branch rls1908_1 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`.

1.4 Physical Testbeds

All FD.io (Fast Data Input/Output) CSIT (Continuous System Integration and Testing) performance test results included in this report are executed on the physical testbeds hosted by LF FD.io project, unless otherwise noted.

Two physical server topology types are used:

- **2-Node Topology:** Consists of one server acting as a System Under Test (SUT) and one server acting as a Traffic Generator (TG), with both servers connected into a ring topology. Used for executing tests that require frame encapsulations supported by TG.
- **3-Node Topology:** Consists of two servers acting as Systems Under Test (SUTs) and one server acting as a Traffic Generator (TG), with all servers connected into a ring topology. Used for executing tests that require frame encapsulations not supported by TG e.g. certain overlay tunnel encapsulations and IPsec. Number of native Ethernet, IPv4 and IPv6 encapsulation tests are also executed on these testbeds, for comparison with 2-Node Topology.

Current FD.io production testbeds are built with SUT servers based on the following processor architectures:

- Intel Xeon: Skylake Platinum 8180 and Haswell-SP E5-2699v3.
- Intel Atom: Denverton C3858.
- ARM: TaiShan 2280, hip07-d05.


Server SUT performance depends on server and processor type, hence results for testbeds based on different servers must be reported separately, and compared if appropriate.

Complete technical specifications of compute servers used in CSIT physical testbeds are maintained in FD.io CSIT repository: https://git.fd.io/csit/tree/docs/lab/testbed_specifications.md.

Following is the description of existing production testbeds.

1.4.1 2-Node Xeon Skylake (2n-skx)

Four 2n-skx testbeds are in operation in FD.io labs. Each 2n-skx testbed is built with two SuperMicro SYS-7049GP-TRT servers, each in turn equipped with two Intel Xeon Skylake Platinum 8180 processors (38.5 MB Cache, 2.50 GHz, 28 cores). 2n-skx physical topology is shown below.



SUT servers are populated with the following NIC models:

1. NIC-1: x710-DA4 4p10GE Intel.
2. NIC-2: xxv710-DA2 2p25GE Intel.
3. NIC-3: mcx556a-edat ConnectX5 2p100GE Mellanox. (Not used yet.)
4. NIC-4: empty, future expansion.
5. NIC-5: empty, future expansion.

6. NIC-6: empty, future expansion.


TG servers run T-Rex application and are populated with the following NIC models:

1. NIC-1: x710-DA4 4p10GE Intel.
2. NIC-2: xxv710-DA2 2p25GE Intel.
3. NIC-3: mcx556a-edat ConnectX5 2p100GE Mellanox. (Not used yet.)
4. NIC-4: empty, future expansion.
5. NIC-5: empty, future expansion.
6. NIC-6: x710-DA4 4p10GE Intel. (For self-tests.)

All Intel Xeon Skylake servers run with Intel Hyper-Threading enabled, doubling the number of logical cores exposed to Linux, with 56 logical cores and 28 physical cores per processor socket.

1.4.2 3-Node Xeon Skylake (3n-skx)

Two 3n-skx testbeds are in operation in FD.io labs. Each 3n-skx testbed is built with three SuperMicro SYS-7049GP-TRT servers, each in turn equipped with two Intel Xeon Skylake Platinum 8180 processors (38.5 MB Cache, 2.50 GHz, 28 cores). 3n-skx physical topology is shown below.



SUT1 and SUT2 servers are populated with the following NIC models:

1. NIC-1: x710-DA4 4p10GE Intel.
2. NIC-2: xxv710-DA2 2p25GE Intel.
3. NIC-3: empty, future expansion.
4. NIC-4: empty, future expansion.

5. NIC-5: empty, future expansion.
6. NIC-6: empty, future expansion.


TG servers run T-Rex application and are populated with the following NIC models:

1. NIC-1: x710-DA4 4p10GE Intel.
2. NIC-2: xxv710-DA2 2p25GE Intel.
3. NIC-3: empty, future expansion.
4. NIC-4: empty, future expansion.
5. NIC-5: empty, future expansion.
6. NIC-6: x710-DA4 4p10GE Intel. (For self-tests.)

All Intel Xeon Skylake servers run with Intel Hyper-Threading enabled, doubling the number of logical cores exposed to Linux, with 56 logical cores and 28 physical cores per processor socket.

1.4.3 3-Node Xeon Haswell (3n-hsw)

Three 3n-hsw testbeds are in operation in FD.io labs. Each 3n-hsw testbed is built with three Cisco UCS-c240m3 servers, each in turn equipped with two Intel Xeon Haswell-SP E5-2699v3 processors (45 MB Cache, 2.3 GHz, 18 cores). 3n-hsw physical topology is shown below.



SUT1 and SUT2 servers are populated with the following NIC models:

1. NIC-1: VIC 1385 2p40GE Cisco.
2. NIC-2: NIC x520 2p10GE Intel.
3. NIC-3: empty.

4. NIC-4: NIC x710-QDA2 2p40GE Intel.
5. NIC-5: NIC x710-DA2 2p10GE Intel.
6. NIC-6: QAT 8950 50G (Walnut Hill) Intel.


TG servers run T-Rex application and are populated with the following NIC models:

1. NIC-1: NIC x710-QDA2 2p40GE Intel.
2. NIC-2: NIC x710-DA2 2p10GE Intel.
3. NIC-3: empty.
4. NIC-4: NIC x710-QDA2 2p40GE Intel.
5. NIC-5: NIC x710-DA2 2p10GE Intel.
6. NIC-6: NIC x710-DA2 2p10GE Intel. (For self-tests.)

All Intel Xeon Haswell servers run with Intel Hyper-Threading disabled, making the number of logical cores exposed to Linux match the number of 18 physical cores per processor socket.

1.4.4 2-Node Atom Denverton (2n-dnv)

2n-dnv testbed is built with: i) one Intel S2600WFT server acting as TG and equipped with two Intel Xeon Skylake Platinum 8180 processors (38.5 MB Cache, 2.50 GHz, 28 cores), and ii) one SuperMicro SYS-E300-9A server acting as SUT and equipped with one Intel Atom C3858 processor (12 MB Cache, 2.00 GHz, 12 cores). 2n-dnv physical topology is shown below.



SUT server have four internal 10G NIC port:

1. P-1: x553 copper port.
2. P-2: x553 copper port.
3. P-3: x553 fiber port.
4. P-4: x553 fiber port.

TG server run T-Rex software traffic generator and are populated with the following NIC models:

1. NIC-1: x550-T2 2p10GE Intel.
2. NIC-2: x550-T2 2p10GE Intel.
3. NIC-3: x520-DA2 2p10GE Intel.
4. NIC-4: x520-DA2 2p10GE Intel.

The 2n-dnv testbed is in operation in Intel SH labs.

1.4.5 3-Node Atom Denverton (3n-dnv)

One 3n-dnv testbed is built with: i) one SuperMicro SYS-7049GP-TRT server acting as TG and equipped with two Intel Xeon Skylake Platinum 8180 processors (38.5 MB Cache, 2.50 GHz, 28 cores), and ii) one SuperMicro SYS-E300-9A server acting as SUT and equipped with one Intel Atom C3858 processor (12 MB Cache, 2.00 GHz, 12 cores). 3n-dnv physical topology is shown below.

SUT1 and SUT2 servers are populated with the following NIC models:

1. NIC-1: x553 2p10GE fiber Intel.
2. NIC-2: x553 2p10GE copper Intel.

TG servers run T-Rex application and are populated with the following NIC models:

1. NIC-1: x710-DA4 4p10GE Intel.

1.4.6 3-Node ARM TaiShan (3n-tsh)

One 3n-tsh testbed is built with: i) one SuperMicro SYS-7049GP-TRT server acting as TG and equipped with two Intel Xeon Skylake Platinum 8180 processors (38.5 MB Cache, 2.50 GHz, 28 cores), and ii) one Huawei TaiShan 2280 server acting as SUT and equipped with one hip07-d05 processor (64* ARM Cortex-A72). 3n-tsh physical topology is shown below.

SUT1 and SUT2 servers are populated with the following NIC models:

1. NIC-1: connectx4 2p25GE Mellanox.
2. NIC-2: x520 2p10GE Intel.

TG servers run T-Rex application and are populated with the following NIC models:

1. NIC-1: x710-DA4 4p10GE Intel.
2. NIC-2: xxv710-DA2 2p25GE Intel.

1.5 Test Methodology

1.5.1 Terminology

- **Frame size:** size of an Ethernet Layer-2 frame on the wire, including any VLAN tags (dot1q, dot1ad) and Ethernet FCS, but excluding Ethernet preamble and inter-frame gap. Measured in Bytes.
- **Packet size:** same as frame size, both terms used interchangeably.
- **Inner L2 size:** for tunneled L2 frames only, size of an encapsulated Ethernet Layer-2 frame, preceded with tunnel header, and followed by tunnel trailer. Measured in Bytes.

- **Inner IP size:** for tunneled IP packets only, size of an encapsulated IPv4 or IPv6 packet, preceded with tunnel header, and followed by tunnel trailer. Measured in Bytes.
- **Device Under Test (DUT):** In software networking, “device” denotes a specific piece of software tasked with packet processing. Such device is surrounded with other software components (such as operating system kernel). It is not possible to run devices without also running the other components, and hardware resources are shared between both. For purposes of testing, the whole set of hardware and software components is called “System Under Test” (SUT). As SUT is the part of the whole test setup performance of which can be measured with [RFC 2544](https://tools.ietf.org/html/rfc2544)², using SUT instead of [RFC 2544](https://tools.ietf.org/html/rfc2544)³ DUT. Device under test (DUT) can be re-introduced when analyzing test results using whitebox techniques, but this document sticks to blackbox testing.
- **System Under Test (SUT):** System under test (SUT) is a part of the whole test setup whose performance is to be benchmarked. The complete methodology contains other parts, whose performance is either already established, or not affecting the benchmarking result.
- **Bi-directional throughput tests:** involve packets/frames flowing in both transmit and receive directions over every tested interface of SUT/DUT. Packet flow metrics are measured per direction, and can be reported as aggregate for both directions (i.e. throughput) and/or separately for each measured direction (i.e. latency). In most cases bi-directional tests use the same (symmetric) load in both directions.
- **Uni-directional throughput tests:** involve packets/frames flowing in only one direction, i.e. either transmit or receive direction, over every tested interface of SUT/DUT. Packet flow metrics are measured and are reported for measured direction.
- **Packet Loss Ratio (PLR):** ratio of packets received relative to packets transmitted over the test trial duration, calculated using formula: $PLR = (pkts_transmitted - pkts_received) / pkts_transmitted$. For bi-directional throughput tests aggregate PLR is calculated based on the aggregate number of packets transmitted and received.
- **Packet Throughput Rate:** maximum packet offered load DUT/SUT forwards within the specified Packet Loss Ratio (PLR). In many cases the rate depends on the frame size processed by DUT/SUT. Hence packet throughput rate MUST be quoted with specific frame size as received by DUT/SUT during the measurement. For bi-directional tests, packet throughput rate should be reported as aggregate for both directions. Measured in packets-per-second (pps) or frames-per-second (fps), equivalent metrics.
- **Bandwidth Throughput Rate:** a secondary metric calculated from packet throughput rate using formula: $bw_rate = pkt_rate - (frame_size + L1_overhead) - 8$, where L1_overhead for Ethernet includes preamble (8 Bytes) and inter-frame gap (12 Bytes). For bi-directional tests, bandwidth throughput rate should be reported as aggregate for both directions. Expressed in bits-per-second (bps).
- **Non Drop Rate (NDR):** maximum packet/bandwidth throughput rate sustained by DUT/SUT at PLR equal zero (zero packet loss) specific to tested frame size(s). MUST be quoted with specific packet size as received by DUT/SUT during the measurement. Packet NDR measured in packets-per-second (or fps), bandwidth NDR expressed in bits-per-second (bps).
- **Partial Drop Rate (PDR):** maximum packet/bandwidth throughput rate sustained by DUT/SUT at PLR greater than zero (non-zero packet loss) specific to tested frame size(s). MUST be quoted with specific packet size as received by DUT/SUT during the measurement. Packet PDR measured in packets-per-second (or fps), bandwidth PDR expressed in bits-per-second (bps).
- **Maximum Receive Rate (MRR):** packet/bandwidth rate regardless of PLR sustained by DUT/SUT under specified Maximum Transmit Rate (MTR) packet load offered by traffic generator. MUST be quoted with both specific packet size and MTR as received by DUT/SUT during the measurement. Packet MRR measured in packets-per-second (or fps), bandwidth MRR expressed in bits-per-second (bps).
- **Trial:** a single measurement step.

² <https://tools.ietf.org/html/rfc2544.html>

³ <https://tools.ietf.org/html/rfc2544.html>

- **Trial duration:** amount of time over which packets are transmitted and received in a single throughput measurement step.

1.5.2 VPP Forwarding Modes

VPP is tested in a number of L2 and IP packet lookup and forwarding modes. Within each mode baseline and scale tests are executed, the latter with varying number of lookup entries.

L2 Ethernet Switching

VPP is tested in three L2 forwarding modes:

- *l2patch*: L2 patch, the fastest point-to-point L2 path that loops packets between two interfaces without any Ethernet frame checks or lookups.
- *l2xc*: L2 cross-connect, point-to-point L2 path with all Ethernet frame checks, but no MAC learning and no MAC lookup.
- *l2bd*: L2 bridge-domain, multipoint-to-multipoint L2 path with all Ethernet frame checks, with MAC learning (unless static MACs are used) and MAC lookup.

l2bd tests are executed in baseline and scale configurations:

- *l2bdbase*: low number of L2 flows (254 per direction) is switched by VPP. They drive the content of MAC FIB size (508 total MAC entries). Both source and destination MAC addresses are incremented on a packet by packet basis.
- *l2bdscale*: high number of L2 flows is switched by VPP. Tested MAC FIB sizes include: i) 10k (5k unique flows per direction), ii) 100k (2x 50k flows) and iii) 1M (2x 500k). Both source and destination MAC addresses are incremented on a packet by packet basis, ensuring new entries are learned refreshed and looked up at every packet, making it the worst case scenario.

Ethernet wire encapsulations tested include: untagged, dot1q, dot1ad.

IPv4 Routing

IPv4 routing tests are executed in baseline and scale configurations:

- *ip4base*: low number of IPv4 flows (253 or 254 per direction) is routed by VPP. They drive the content of IPv4 FIB size (506 or 508 total /32 prefixes). Destination IPv4 addresses are incremented on a packet by packet basis.
- *ip4scale*: high number of IPv4 flows is routed by VPP. Tested IPv4 FIB sizes of /32 prefixes include: i) 20k (10k unique flows per direction), ii) 200k (2x 100k flows) and iii) 2M (2x 1M). Destination IPv4 addresses are incremented on a packet by packet basis, ensuring new FIB entries are looked up at every packet, making it the worst case scenario.

IPv6 Routing

IPv6 routing tests are executed in baseline and scale configurations:

- *ip6base*: low number of IPv6 flows (253 or 254 per direction) is routed by VPP. They drive the content of IPv6 FIB size (506 or 508 total /128 prefixes). Destination IPv6 addresses are incremented on a packet by packet basis.
- *ip6scale*: high number of IPv6 flows is routed by VPP. Tested IPv6 FIB sizes of /128 prefixes include: i) 20k (10k unique flows per direction), ii) 200k (2x 100k flows) and iii) 2M (2x 1M). Destination IPv6 addresses are incremented on a packet by packet basis, ensuring new FIB entries are looked up at every packet, making it the worst case scenario.

SRv6 Routing

SRv6 routing tests are executed in a number of baseline configurations, in each case SR policy and steering policy are configured for one direction and one (or two) SR behaviours (functions) in the other directions:

- *srv6enc1sid*: One SID (no SRH present), one SR function - End.
- *srv6enc2sids*: Two SIDs (SRH present), two SR functions - End and End.DX6.
- *srv6enc2sids-nodcaps*: Two SIDs (SRH present) without decapsulation, one SR function - End.
- *srv6proxy-dyn*: Dynamic SRv6 proxy, one SR function - End.AD.
- *srv6proxy-masq*: Masquerading SRv6 proxy, one SR function - End.AM.
- *srv6proxy-stat*: Static SRv6 proxy, one SR function - End.AS.

In all listed cases low number of IPv6 flows (253 per direction) is routed by VPP.

1.5.3 Tunnel Encapsulations

Tunnel encapsulations testing is grouped based on the type of outer header: IPv4 or IPv6.

IPv4 Tunnels

VPP is tested in the following IPv4 tunnel baseline configurations:

- *ip4vxlan-l2bdbase*: VXLAN over IPv4 tunnels with L2 bridge-domain MAC switching.
- *ip4vxlan-l2xcbase*: VXLAN over IPv4 tunnels with L2 cross-connect.
- *ip4lispip4-ip4base*: LISP over IPv4 tunnels with IPv4 routing.
- *ip4lispip6-ip6base*: LISP over IPv4 tunnels with IPv6 routing.

In all cases listed above low number of MAC, IPv4, IPv6 flows (254 or 253 per direction) is switched or routed by VPP.

In addition selected IPv4 tunnels are tested at scale:

- *dot1q-ip4vxlanscale-l2bd*: VXLAN over IPv4 tunnels with L2 bridge-domain MAC switching, with scaled up dot1q VLANs (10, 100, 1k), mapped to scaled up L2 bridge-domains (10, 100, 1k), that are in turn mapped to (10, 100, 1k) VXLAN tunnels. 64.5k flows are transmitted per direction.

IPv6 Tunnels

VPP is tested in the following IPv6 tunnel baseline configurations:

- *ip6lispip4-ip4base*: LISP over IPv4 tunnels with IPv4 routing.
- *ip6lispip6-ip6base*: LISP over IPv4 tunnels with IPv6 routing.

In all cases listed above low number of IPv4, IPv6 flows (253 per direction) is routed by VPP.

1.5.4 VPP Features

VPP is tested in a number of data plane feature configurations across different forwarding modes. Following sections list features tested.

ACL Security-Groups

Both stateless and stateful access control lists (ACL), also known as security-groups, are supported by VPP.

Following ACL configurations are tested for MAC switching with L2 bridge-domains:

- *l2bdbasemaclrn-iacl{E}sl-{F}flows*: Input stateless ACL, with {E} entries and {F} flows.
- *l2bdbasemaclrn-oacl{E}sl-{F}flows*: Output stateless ACL, with {E} entries and {F} flows.
- *l2bdbasemaclrn-iacl{E}sf-{F}flows*: Input stateful ACL, with {E} entries and {F} flows.
- *l2bdbasemaclrn-oacl{E}sf-{F}flows*: Output stateful ACL, with {E} entries and {F} flows.

Following ACL configurations are tested with IPv4 routing:

- *ip4base-iacl{E}sl-{F}flows*: Input stateless ACL, with {E} entries and {F} flows.
- *ip4base-oacl{E}sl-{F}flows*: Output stateless ACL, with {E} entries and {F} flows.
- *ip4base-iacl{E}sf-{F}flows*: Input stateful ACL, with {E} entries and {F} flows.
- *ip4base-oacl{E}sf-{F}flows*: Output stateful ACL, with {E} entries and {F} flows.

ACL tests are executed with the following combinations of ACL entries and number of flows:

- ACL entry definitions
 - flow non-matching deny entry: (src-ip4, dst-ip4, src-port, dst-port).
 - flow matching permit ACL entry: (src-ip4, dst-ip4).
- {E} - number of non-matching deny ACL entries, {E} = [1, 10, 50].
- {F} - number of UDP flows with different tuple (src-ip4, dst-ip4, src-port, dst-port), {F} = [100, 10k, 100k].
- All {E}x{F} combinations are tested per ACL type, total of 9.

ACL MAC-IP

MAC-IP binding ACLs are tested for MAC switching with L2 bridge-domains:

- *l2bdbasemaclrn-macip-iacl{E}sl-{F}flows*: Input stateless ACL, with {E} entries and {F} flows.

MAC-IP ACL tests are executed with the following combinations of ACL entries and number of flows:

- ACL entry definitions
 - flow non-matching deny entry: (dst-ip4, dst-mac, bit-mask)
 - flow matching permit ACL entry: (dst-ip4, dst-mac, bit-mask)
- {E} - number of non-matching deny ACL entries, {E} = [1, 10, 50]
- {F} - number of UDP flows with different tuple (dst-ip4, dst-mac), {F} = [100, 10k, 100k]
- All {E}x{F} combinations are tested per ACL type, total of 9.

NAT44

NAT44 is tested in baseline and scale configurations with IPv4 routing:

- *ip4base-nat44*: baseline test with single NAT entry (addr, port), single UDP flow.
- *ip4base-udpsrcscale{U}-nat44*: baseline test with {U} NAT entries (addr, {U}ports), {U}=15.
- *ip4scale{R}-udpsrcscale{U}-nat44*: scale tests with {R}*{U} NAT entries ({R}addr, {U}ports), {R}=[100, 1k, 2k, 4k], {U}=15.

1.5.5 Data Plane Throughput

Data Plane Throughput Tests

Network data plane throughput is measured using multiple test methods in order to obtain representative and repeatable results across the large set of performance test cases implemented and executed within CSIT.

Following throughput test methods are used:

- MLRsearch - Multiple Loss Ratio search
- MRR - Maximum Receive Rate
- PLRsearch - Probabilistic Loss Ratio search

Description of each test method is followed by generic test properties shared by all methods.

MLRsearch Tests

Description

Multiple Loss Ratio search (MLRsearch) tests discover multiple packet throughput rates in a single search, reducing the overall test execution time compared to a binary search. Each rate is associated with a distinct Packet Loss Ratio (PLR) criteria. In FD.io CSIT two throughput rates are discovered: Non-Drop Rate (NDR, with zero packet loss, PLR=0) and Partial Drop Rate (PDR, with PLR<0.5%). MLRsearch is compliant with [RFC 2544](#)⁴.

Usage

MLRsearch tests are run to discover NDR and PDR rates for each VPP and DPDK release covered by CSIT report. Results for small frame sizes (64b/78B, IMIX) are presented in packet throughput graphs (Box-and-Whisker Plots) with NDR and PDR rates plotted against the test cases covering popular VPP packet paths.

Each test is executed at least 10 times to verify measurements repeatability and results are compared between releases and test environments. NDR and PDR packet and bandwidth throughput results for all frame sizes and for all tests are presented in detailed results tables.

Details

See *MLRsearch Tests* (page 17) section for more detail. MLRsearch is being standardized in IETF in [draft-vpolak-mkonstan-mlrsearch](#)⁵.

MRR Tests

Description

Maximum Receive Rate (MRR) tests are complementary to MLRsearch tests, as they provide a maximum “raw” throughput benchmark for development and testing community.

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

⁴ <https://tools.ietf.org/html/rfc2544.html>

⁵ <https://tools.ietf.org/html/draft-vpolak-mkonstan-bmwg-mlrsearch>

Usage

MRR tests are much faster than MLRsearch as they rely on a single trial or a small set of trials with very short duration. It is this property that makes them suitable for continuous execution in daily performance trending jobs enabling detection of performance anomalies (regressions, progressions) resulting from data plane code changes.

MRR tests are also used for VPP per patch performance jobs verifying patch performance vs. parent. CSIT reports include MRR throughput comparisons between releases and test environments. Small frame sizes only (64b/78B, IMIX).

Details

See *MRR Throughput* (page 17) section for more detail about MRR tests configuration.

FD.io CSIT performance dashboard includes complete description of [daily performance trending tests](#)⁶ and [VPP per patch tests](#)⁷.

PLRsearch Tests

Description

Probabilistic Loss Ratio search (PLRsearch) tests discovers a packet throughput rate associated with configured Packet Loss Ratio (PLR) criteria for tests run over an extended period of time a.k.a. soak testing. PLRsearch assumes that system under test is probabilistic in nature, and not deterministic.

Usage

PLRsearch are run to discover a sustained throughput for $PLR=10^{-7}$ (close to NDR) for VPP release covered by CSIT report. Results for small frame sizes (64b/78B) are presented in packet throughput graphs (Box Plots) for a small subset of baseline tests.

Each soak test lasts 2hrs and is executed at least twice. Results are compared against NDR and PDR rates discovered with MLRsearch.

Details

See *PLRsearch* (page 18) methodology section for more detail. PLRsearch is being standardized in IETF in [draft-vpolak-bmwg-plrsearch](#)⁸.

Generic Test Properties

All data plane throughput test methodologies share following generic properties:

- Tested L2 frame sizes (untagged Ethernet):
 - IPv4 payload: 64B, IMIX (28x64B, 16x570B, 4x1518B), 1518B, 9000B.
 - IPv6 payload: 78B, IMIX (28x78B, 16x570B, 4x1518B), 1518B, 9000B.
 - All quoted sizes include frame CRC, but exclude per frame transmission overhead of 20B (preamble, inter frame gap).

⁶ https://docs.fd.io/csit/master/trending/methodology/performance_tests.html

⁷ https://docs.fd.io/csit/master/trending/methodology/perpatch_performance_tests.html

⁸ <https://tools.ietf.org/html/draft-vpolak-bmwg-plrsearch>

- Offered packet load is always bi-directional and symmetric.
- All measured and reported packet and bandwidth rates are aggregate bi-directional rates reported from external Traffic Generator perspective.

MLRsearch Tests

Overview

Multiple Loss Rate search (MLRsearch) tests use new search algorithm implemented in FD.io CSIT project. MLRsearch discovers multiple packet throughput rates in a single search, with each rate associated with a different 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).

MLRsearch discovers NDR and PDR in a single pass reducing required time duration compared to separate binary searches for NDR and PDR. Overall search time is reduced 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 execution time when compared to standard NDR/PDR binary search, while guaranteeing similar results.

If needed, MLRsearch 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.

Search Implementation

Detailed description of the MLRsearch algorithm is included in the IETF draft [draft-vpolak-mkonstan-mlrsearch](#)⁹ that is in the process of being standardized in the IETF Benchmarking Methodology Working Group (BMWG).

MLRsearch is also available as a [PyPI \(Python Package Index\) library](#)¹⁰.

Implementation Deviations

FD.io CSIT implementation of MLRsearch so far is fully based on the -01 version of the [draft-vpolak-mkonstan-mlrsearch-01](#)¹¹.

MRR Throughput

Maximum Receive Rate (MRR) tests are complementary to MLRsearch tests, as they provide a maximum “raw” throughput benchmark for development and testing community. MRR 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 currently used for following test jobs:

- Report performance comparison: 64B, IMIX for vhost, memif.
- Daily performance trending: 64B, IMIX for vhost, memif.

⁹ <https://tools.ietf.org/html/draft-vpolak-mkonstan-bmwg-mlrsearch>

¹⁰ <https://pypi.org/project/MLRsearch/>

¹¹ <https://tools.ietf.org/html/draft-vpolak-mkonstan-bmwg-mlrsearch-01>

- Per-patch performance verification: 64B.
- PLRsearch soaking tests: 64B.

Maximum offered load for specific L2 Ethernet frame size is set to either the maximum bi-directional link rate or tested NIC model capacity, as follows:

- For 10GE NICs the maximum packet rate load is 2x14.88 Mpps for 64B, a 10GE bi-directional link rate.
- For 25GE NICs the maximum packet rate load is 2x18.75 Mpps for 64B, a 25GE bi-directional link sub-rate limited by 25GE NIC used on TRex TG, XXV710.
- For 40GE NICs the maximum packet rate load is 2x18.75 Mpps for 64B, a 40GE bi-directional link sub-rate limited by 40GE NIC used on TRex TG, XL710. Packet rate for other tested frame sizes is limited by PCIeGen3 x8 bandwidth limitation of ~50Gbps.

MRR test code implements multiple bursts of offered packet load and has two configurable burst parameters: individual trial duration and number of trials in a single burst. This enables more precise performance trending by providing more results data for analysis.

Burst parameter settings vary between different tests using MRR:

- MRR individual trial duration:
 - Report performance comparison: 1 sec.
 - Daily performance trending: 1 sec.
 - Per-patch performance verification: 10 sec.
 - PLRsearch soaking tests: 5.2 sec.
- Number of MRR trials per burst:
 - Report performance comparison: 10.
 - Daily performance trending: 10.
 - Per-patch performance verification: 5.
 - PLRsearch soaking tests: 1.

PLRsearch

Motivation for PLRsearch

Network providers are interested in throughput a system can sustain.

[RFC 2544](#)¹² assumes loss ratio is given by a deterministic function of offered load. But NFV software systems are not deterministic enough. This makes deterministic algorithms (such as [binary search](#)¹³ per RFC 2544 and MLRsearch with single trial) to return results, which when repeated show relatively high standard deviation, thus making it harder to tell what “the throughput” actually is.

We need another algorithm, which takes this indeterminism into account.

Generic Algorithm

Detailed description of the PLRsearch algorithm is included in the IETF draft [draft-vpolak-bmwg-plrsearch-02](#)¹⁴ that is in the process of being standardized in the IETF Benchmarking Methodology Working Group (BMWG).

¹² <https://tools.ietf.org/html/rfc2544>

¹³ https://en.wikipedia.org/wiki/Binary_search_algorithm

¹⁴ <https://tools.ietf.org/html/draft-vpolak-bmwg-plrsearch-02>

Terms

The rest of this page assumes the reader is familiar with the following terms defined in the IETF draft:

- Trial Order Independent System
- Duration Independent System
- Target Loss Ratio
- Critical Load
- Offered Load regions
 - Zero Loss Region
 - Non-Deterministic Region
 - Guaranteed Loss Region
- Fitting Function
 - Stretch Function
 - Erf Function
- Bayesian Inference
 - Prior distribution
 - Posterior Distribution
- Numeric Integration
 - Monte Carlo
 - Importance Sampling

FD.io CSIT Implementation Specifics

The search receives `min_rate` and `max_rate` values, to avoid measurements at offered loads not supported by the traffic generator.

The implemented tests cases use bidirectional traffic. The algorithm stores each rate as bidirectional rate (internally, the algorithm is agnostic to flows and directions, it only cares about aggregate counts of packets sent and packets lost), but debug output from traffic generator lists unidirectional values.

In a sample implementation in FD.io CSIT project, there is roughly 0.5 second delay between trials due to restrictions imposed by packet traffic generator in use (T-Rex).

As measurements results come in, posterior distribution computation takes more time (per sample), although there is a considerable constant part (mostly for inverting the fitting functions).

Also, the integrator needs a fair amount of samples to reach the region the posterior distribution is concentrated at.

And of course, the speed of the integrator depends on computing power of the CPU the algorithm is able to use.

All those timing related effects are addressed by arithmetically increasing trial durations with configurable coefficients (currently 5.1 seconds for the first trial, each subsequent trial being 0.1 second longer).

In order to avoid them, the current implementation tracks natural logarithm (instead of the original quantity) for any quantity which is never negative. Logarithm of zero is minus infinity (not supported by Python), so special value "None" is used instead. Specific functions for frequent operations (such as "logarithm of sum of exponentials") are defined to handle None correctly.

Current implementation uses two fitting functions. In general, their estimates for critical rate differ, which adds a simple source of systematic error, on top of randomness error reported by integrator. Otherwise the reported stdev of critical rate estimate is unrealistically low.

Both functions are not only increasing, but also convex (meaning the rate of increase is also increasing).

Both fitting functions have several mathematically equivalent formulas, each can lead to an overflow or underflow in different sub-terms. Overflows can be eliminated by using different exact formulas for different argument ranges. Underflows can be avoided by using approximate formulas in affected argument ranges, such ranges have their own formulas to compute. At the end, both fitting function implementations contain multiple “if” branches, discontinuities are a possibility at range boundaries.

The numeric integrator expects all the parameters to be distributed (independently and) uniformly on an interval (-1, 1).

As both “mrr” and “spread” parameters are positive and not dimensionless, a transformation is needed. Dimensionality is inherited from max_rate value.

The “mrr” parameter follows a [Lomax distribution](#)¹⁵ with alpha equal to one, but shifted so that mrr is always greater than 1 packet per second.

The “stretch” parameter is generated simply as the “mrr” value raised to a random power between zero and one; thus it follows a [reciprocal distribution](#)¹⁶.

After few measurements, the posterior distribution of fitting function arguments gets quite concentrated into a small area. The integrator is using [Monte Carlo](#)¹⁷ with [importance sampling](#)¹⁸ where the biased distribution is [bivariate Gaussian](#)¹⁹ distribution, with deliberately larger variance. If the generated sample falls outside (-1, 1) interval, another sample is generated.

The center and the covariance matrix for the biased distribution is based on the first and second moments of samples seen so far (within the computation). The center is used directly, covariance matrix is scaled up by a heuristic constant (8.0 by default). The following additional features are applied designed to avoid hyper-focused distributions.

Each computation starts with the biased distribution inherited from the previous computation (zero point and unit covariance matrix is used in the first computation), but the overall weight of the data is set to the weight of the first sample of the computation. Also, the center is set to the first sample point. When additional samples come, their weight (including the importance correction) is compared to sum of the weights of data seen so far (within the iteration). If the new sample is more than one e-fold more impactful, both weight values (for data so far and for the new sample) are set to (geometric) average of the two weights.

This combination showed the best behavior, as the integrator usually follows two phases. First phase (where inherited biased distribution or single big sample are dominating) is mainly important for locating the new area the posterior distribution is concentrated at. The second phase (dominated by whole sample population) is actually relevant for the critical rate estimation.

First two measurements are hardcoded to happen at the middle of rate interval and at max_rate. Next two measurements follow MRR-like logic, offered load is decreased so that it would reach target loss ratio if offered load decrease lead to equal decrease of loss rate.

The rest of measurements start directly in between erf and stretch estimate average. There is one workaround implemented, aimed at reducing the number of consequent zero loss measurements (per fitting function). The workaround first stores every measurement result which loss ratio was the targeted loss ratio or higher. Sorted list (called lossy loads) of such results is maintained.

When a sequence of one or more zero loss measurement results is encountered, a smallest of lossy loads is drained from the list. If the estimate average is smaller than the drained value, a weighted average of this estimate and the drained value is used as the next offered load. The weight of the estimate decreases exponentially with the length of consecutive zero loss results.

¹⁵ https://en.wikipedia.org/wiki/Lomax_distribution

¹⁶ https://en.wikipedia.org/wiki/Reciprocal_distribution

¹⁷ https://en.wikipedia.org/wiki/Monte_Carlo_integration

¹⁸ https://en.wikipedia.org/wiki/Importance_sampling

¹⁹ https://en.wikipedia.org/wiki/Multivariate_normal_distribution

This behavior helps the algorithm with convergence speed, as it does not need so many zero loss result to get near critical region. Using the smallest (not drained yet) of lossy loads makes it sure the new offered load is unlikely to result in big loss region. Draining even if the estimate is large enough helps to discard early measurements when loss hapened at too low offered load. Current implementation adds 4 copies of lossy loads and drains 3 of them, which leads to fairly stable behavior even for somewhat inconsistent SUTs.

As high loss count measurements add many bits of information, they need a large amount of small loss count measurements to balance them, making the algorithm converge quite slowly. Typically, this happens when few initial measurements suggest spread way bigger then later measurements. The workaround in offered load selection helps, but more intelligent workarounds could get faster convergence still.

Some systems evidently do not follow the assumption of repeated measurements having the same average loss rate (when the offered load is the same). The idea of estimating the trend is not implemented at all, as the observed trends have varied characteristics.

Probably, using a more realistic fitting functions will give better estimates than trend analysis.

Bottom Line

The notion of Throughput is easy to grasp, but it is harder to measure with any accuracy for non-deterministic systems.

Even though the notion of critical rate is harder to grasp than the notion of throughput, it is easier to measure using probabilistic methods.

In testing, the difference between througput measurements and critical rate measurements is usually small, see soak vs ndr comparison.

In pactice, rules of thumb such as "send at max 95% of purported throughput" are common. The correct benchmarking analysis should ask "Which notion is 95% of throughput an approximation to?" before attempting to answer "Is 95% of critical rate safe enough?".

Algorithmic Analysis

While the estimation computation is based on hard probability science; the offered load selection part of PLRsearch logic is pure heuristics, motivated by what would a human do based on measurement and computation results.

The quality of any heuristic is not affected by soundness of its motivation, just by its ability to achieve the intended goals. In case of offered load selection, the goal is to help the search to converge to the long duration estimates sooner.

But even those long duration estimates could still be of poor quality. Even though the estimate computation is Bayesian (so it is the best it could be within the applied assumptions), it can still of poor quality when compared to what a human would estimate.

One possible source of poor quality is the randomness inherently present in Monte Carlo numeric integration, but that can be suppressed by tweaking the time related input parameters.

The most likely source of poor quality then are the assumptions. Most importantly, the number and the shape of fitting functions; but also others, such as trial order independence and duration independence.

The result can have poor quality in basically two ways. One way is related to location. Both upper and lower bounds can be overestimates or underestimates, meaning the entire estimated interval between lower bound and upper bound lays above or below (respectively) of human-estimated interval. The other way is related to the estimation interval width. The interval can be too wide or too narrow, compared to human estimation.

An estimate from a particular fitting function can be classified as an overestimate (or underestimate) just by looking at time evolution (without human examining measurement results). Overestimates decrease by time, underestimates increase by time (assuming the system performance stays constant).

Quality of the width of the estimation interval needs human evaluation, and is unrelated to both rate of narrowing (both good and bad estimate intervals get narrower at approximately the same relative rate) and relative width (depends heavily on the system being tested).

The following pictures show the upper (red) and lower (blue) bound, as well as average of Stretch (pink) and Erf (light green) estimate, and offered load chosen (grey), as computed by PLRsearch, after each trial measurement within the 30 minute duration of a test run.

Both graphs are focusing on later estimates. Estimates computed from few initial measurements are wildly off the y-axis range shown.

The following analysis will rely on frequency of zero loss measurements and magnitude of loss ratio if nonzero.

The offered load selection strategy used implies zero loss measurements can be gleamed from the graph by looking at offered load points. When the points move up farther from lower estimate, it means the previous measurement had zero loss. After non-zero loss, the offered load starts again right between (the previous values of) the estimate curves.

The very big loss ratio results are visible as noticeable jumps of both estimates downwards. Medium and small loss ratios are much harder to distinguish just by looking at the estimate curves, the analysis is based on raw loss ratio measurement results.


The following descriptions should explain why the graphs seem to signal low quality estimate at first sight, but a more detailed look reveals the quality is good (considering the measurement results).

L2 patch

Both fitting functions give similar estimates, the graph shows “stochasticity” of measurements (estimates increase and decrease within small time regions), and an overall trend of decreasing estimates.

On the first look, the final interval looks fairly narrow, especially compared to the region the estimates have travelled during the search. But the look at the frequency of zero loss results shows this is not a case of overestimation. Measurements at around the same offered load have higher probability of zero loss earlier (when performed farther from upper bound), but smaller probability later (when performed closer to upper bound). That means it is the performance of the system under test that decreases (slightly) over time.

With that in mind, the apparent narrowness of the interval is not a sign of low quality, just a consequence of PLRsearch assuming the performance stays constant.




Vhost

This test case shows what looks like a quite broad estimation interval, compared to other test cases with similarly looking zero loss frequencies. Notable features are infrequent high-loss measurement results causing big drops of estimates, and lack of long-term convergence.

Any convergence in medium-sized intervals (during zero loss results) is reverted by the big loss results, as they happen quite far from the critical load estimates, and the two fitting functions extrapolate differently.

In other words, human only seeing estimates from one fitting function would expect narrower end interval, but human seeing the measured loss ratios agrees that the interval should be wider than that.



Summary

The two graphs show the behavior of PLRsearch algorithm applied to soaking test when some of PLRsearch assumptions do not hold:

- L2 patch measurement results violate the assumption of performance not changing over time.
- Vhost measurement results violate the assumption of Poisson distribution matching the loss counts.

The reported upper and lower bounds can have distance larger or smaller than a first look by a human would expect, but a more closer look reveals the quality is good, considering the circumstances.

The usefulness of the critical load estimate is of questionable value when the assumptions are violated.

Some improvements can be made via more specific workarounds, for example long term limit of L2 patch performance could be estimated by some heuristic.

Other improvements can be achieved only by asking users whether loss patterns matter. Is it better to have single digit losses distributed fairly evenly over time (as Poisson distribution would suggest), or is it better to have short periods of medium losses mixed with long periods of zero losses (as happens in Vhost test) with the same overall loss ratio?

1.5.6 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 100% of discovered NDR and PDR rates for each throughput test and packet size (except IMIX).
- TG sends dedicated latency streams, one per direction, each at the rate of 9 kpps 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 if the three node topology is used for given performance test; 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 2×2 usec 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.

1.5.7 Multi-Core Speedup

All performance tests are executed with single processor core and with multiple cores scenarios.

Intel Hyper-Threading (HT)

Intel Xeon processors used in FD.io CSIT can operate either in HT Disabled mode (single logical core per each physical core) or in HT Enabled mode (two logical cores per each physical core). HT setting is applied in BIOS and requires server SUT reload for it to take effect, making it impractical for continuous changes of HT mode of operation.

CSIT-1908.1 performance tests are executed with server SUTs' Intel XEON processors configured with Intel Hyper-Threading Disabled for all Xeon Haswell testbeds (3n-hsw) and with Intel Hyper-Threading Enabled for all Xeon Skylake testbeds.

More information about physical testbeds is provided in *Physical Testbeds* (page 4).

Multi-core Tests

CSIT-1908.1 multi-core tests are executed in the following VPP worker thread and physical core configurations:

1. Intel Xeon Haswell testbeds (3n-hsw) with Intel HT disabled (1 logical CPU core per each physical core):
 1. 1t1c - 1 VPP worker thread on 1 physical core.
 2. 2t2c - 2 VPP worker threads on 2 physical cores.
 3. 4t4c - 4 VPP worker threads on 4 physical cores.
1. Intel Xeon Skylake testbeds (2n-skx, 3n-skx) with Intel HT enabled (2 logical CPU cores per each physical core):
 1. 2t1c - 2 VPP worker threads on 1 physical core.
 2. 4t2c - 4 VPP worker threads on 2 physical cores.
 3. 8t4c - 8 VPP worker threads on 4 physical cores.

VPP worker threads are the data plane threads running on isolated logical cores. With Intel HT enabled VPP workers are placed as sibling threads on each used physical core. VPP control threads (main, stats) are running on a separate non-isolated core together with other Linux processes.

In all CSIT tests care is taken to ensure that each VPP worker handles the same amount of received packet load and does the same amount of packet processing work. This is achieved by evenly distributing per interface type (e.g. physical, virtual) receive queues over VPP workers using default VPP round-robin mapping and by loading these queues with the same amount of packet flows.

If number of VPP workers is higher than number of physical or virtual interfaces, multiple receive queues are configured on each interface. NIC Receive Side Scaling (RSS) for physical interfaces and multi-queue for virtual interfaces are used for this purpose.

Section *Speedup Multi-Core* (page 152) includes a set of graphs illustrating packet throughput speedup when running VPP worker threads on multiple cores. 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.

1.5.8 Reconfiguration Tests

Overview

Reconf tests are designed to measure the impact of VPP re-configuration on data plane traffic. While VPP takes some measures against the traffic being entirely stopped for a prolonged time, the immediate forwarding rate varies during the re-configuration, as some configurations steps need the active dataplane worker threads to be stopped temporarily.

As the usual methods of measuring throughput need multiple trial measurements with somewhat long durations, and the re-configuration process can also be long, finding an offered load which would result in zero loss during the re-configuration process would be time-consuming.

Instead, reconf tests find a throughput value (lower bound for NDR) without re-configuration, and then maintain that offered load during re-configuration. The measured loss count is then assumed to be caused by the re-configuration process. The result published by reconf tests is the effective blocked time, that is the loss count divided by the offered load.

Current Implementation

Each reconf suite is based on a similar MLRsearch performance suite.

MLRsearch parameters are changed to speed up the throughput discovery. For example, PDR is not searched for, and final trial duration is shorter.

The MLRsearch suite has to contain a configuration parameter that can be scaled up, e.g. number of routes or number of service chains. Currently, only increasing the scale is supported as the re-configuration operation. In future, scale decrease or other operations can be implemented.

The traffic profile is not changed, so the traffic present is processed only by the smaller scale configuration. The added routes / chains are not targetted by the traffic.

For the re-configuration, the same Robot Framework and Python libraries are used, as were used in the initial configuration, with the exception of the final calls that do not interact with VPP (e.g. starting virtual machines) being skipped to reduce the test overall duration.

Discussion

Robot Framework introduces a certain overhead, which may affect timing of individual VPP API calls, which in turn may affect the number of packets lost.

The exact calls executed may contain unnecessary info dumps, repeated commands, or commands which change a value that do not need to be changed (e.g. MTU). Thus, implementation details are affecting the results, even if their effect on the corresponding MLRsearch suite is negligible.

The lower bound for NDR is the only value safe to be used when zero packets lost are expected without re-configuration. But different suites show different "jitter" in that value. For some suites, the lower bound

is not tight, allowing full NIC buffers to drain quickly between worker pauses. For other suites, lower bound for NDR still has quite a large probability of non-zero packet loss even without re-configuration.

But the results show very high effective blocked time, so the two objections related to NDR lower bound are negligible in comparison.

1.5.9 VPP Startup Settings

CSIT code manipulates a number of VPP settings in `startup.conf` for optimized performance. List of common settings applied to all tests and test dependent settings follows.

See [VPP startup.conf](#)²⁰ for a complete set and description of listed settings.

Common Settings

List of VPP `startup.conf` settings applied to all tests:

1. `heap-size <value>` - set separately for ip4, ip6, stats, main depending on scale tested.
2. `no-tx-checksum-offload` - disables UDP / TCP TX checksum offload in DPDK. Typically needed for use faster vector PMDs (together with `no-multi-seg`).
3. `buffers-per-numa <value>` - sets a number of memory buffers allocated to VPP per CPU socket. VPP default is 16384. Needs to be increased for scenarios with large number of interfaces and worker threads. To accommodate for scale tests, CSIT is setting it to the maximum possible value corresponding to the limit of DPDK memory mappings (currently 256). For Xeon Skylake platforms configured with 2MB hugepages and VPP data-size and buffer-size defaults (2048B and 2496B respectively), this results in value of 215040 ($256 * 840 = 215040$, $840 * 2496B$ buffers fit in 2MB hugepage). For Xeon Haswell nodes value of 107520 is used.

Per Test Settings

List of vpp `startup.conf` settings applied dynamically per test:

1. `corelist-workers <list_of_cores>` - list of logical cores to run VPP worker data plane threads. Depends on HyperThreading and core per test configuration.
2. `num-rx-queues <value>` - depends on a number of VPP threads and NIC interfaces.
3. `no-multi-seg` - disables multi-segment buffers in DPDK, improves packet throughput, but disables Jumbo MTU support. Disabled for all tests apart from the ones that require Jumbo 9000B frame support.
4. `UIO driver` - depends on topology file definition.
5. `QAT VFs` - depends on NRThreads, each thread = 1QAT VFs.

1.5.10 KVM VMs vhost-user

QEMU is used for KVM VM vhost-user testing environment. By default, standard QEMU version is used, preinstalled from OS repositories (qemu-2.11.1 for Ubuntu 18.04, qemu-2.5.0 for Ubuntu 16.04). The path to the QEMU binary can be adjusted in `Constants.py`.

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

CSIT supports two types of VMs:

- **Image-VM:** used for all functional, VPP_device, and regular performance tests except NFV density tests.

²⁰ https://git.fd.io/vpp/tree/src/vpp/conf/startup.conf?h=stable/1908_1&id=f4dcae4164f93dac80d4af19af0ee20e712ec673

- **Kernel-VM:** new VM type introduced for NFV density tests to provide greater in-VM application install flexibility and to further reduce test execution time by simpler VM lifecycle management.

Image-VM

CSIT can use a pre-created VM image. The path to the image can be adjusted in *Constants.py*. For convenience and full compatibility CSIT repository contains a set of scripts to prepare **Built-root**²¹ based embedded Linux image with all the dependencies needed to run DPDK Testpmd, DPDK L3Fwd, Linux bridge or Linux IPv4 forwarding.

Built-root was chosen for a VM image to make it lightweight and with fast booting time to limit impact on tests duration.

In order to execute CSIT tests, VM image must have following software installed: qemu-guest-agent, sshd, bridge-utils, VirtIO support and DPDK Testpmd/L3fwd applications. Username/password for the VM must be cisco/cisco and NOPASSWD sudo access. The interface naming is based on the driver (management interface type is Intel E1000), all E1000 interfaces will be named mgmt<n> and all VirtIO interfaces will be named virtio<n>. In VM /etc/init.d/qemu-guest-agent must be set to TRANSPORT=isa-serial:/dev/ttyS1 because ttyS0 is used by serial console and ttyS1 is dedicated for qemu-guest-agent in QEMU setup.

Kernel-VM

CSIT can use a kernel KVM image as a boot kernel, as an alternative to image VM. This option allows better configurability of what application is running in VM userspace. Using root9p filesystem allows mapping the host-OS filesystem as read only guest-OS filesystem.

Example of custom init script for the kernel-VM:

```
#!/bin/bash
mount -t sysfs -o "nodev,noexec,nosuid" sysfs /sys
mount -t proc -o "nodev,noexec,nosuid" proc /proc
mkdir /dev/pts
mkdir /dev/hugepages
mount -t devpts -o "rw,noexec,nosuid,gid=5,mode=0620" devpts /dev/pts || true
mount -t tmpfs -o "rw,noexec,nosuid,size=10%,mode=0755" tmpfs /run
mount -t tmpfs -o "rw,noexec,nosuid,size=10%,mode=0755" tmpfs /tmp
mount -t hugetlbfs -o "rw,relatime,pagesize=2M" hugetlbfs /dev/hugepages
echo 0000:00:06.0 > /sys/bus/pci/devices/0000:00:06.0/driver/unbind
echo 0000:00:07.0 > /sys/bus/pci/devices/0000:00:07.0/driver/unbind
echo vfio-pci > /sys/bus/pci/devices/0000:00:06.0/driver_override
echo vfio-pci > /sys/bus/pci/devices/0000:00:07.0/driver_override
echo 0000:00:06.0 > /sys/bus/pci/drivers/vfio-pci/bind
echo 0000:00:07.0 > /sys/bus/pci/drivers/vfio-pci/bind
$vnf_bin
poweroff -f
```

QemuUtils library during runtime replaces the \$vnf_bin variable by the path to NF binary and its parameters. This allows CSIT to run any application installed on host OS, for example the same version of VPP as running on the host-OS.

Kernel-VM image must be available in the host filesystem as a prerequisite. The path to kernel-VM image is defined in *Constants.py*.

1.5.11 LXC/DRC Container Memif

CSIT 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

²¹ <https://buildroot.org/>

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 *Logical Topologies* (page 31).

In addition to above vswitch tests, a single memif interface test is executed. It runs in a simple topology of two VPP container instances connected over memif interface in order to verify standalone memif interface performance.

More information about CSIT LXC and DRC setup and control is available in *Container Orchestration in CSIT* (page 421).

1.5.12 NFV Service Density

Network Function Virtualization (NFV) service density tests focus on measuring total per server throughput at varied NFV service “packing” densities with vswitch providing host dataplane. The goal is to compare and contrast performance of a shared vswitch for different network topologies and virtualization technologies, and their impact on vswitch performance and efficiency in a range of NFV service configurations.

Each NFV service instance consists of a set of Network Functions (NFs), running in VMs (VNFs) or in Containers (CNFs), that are connected into a virtual network topology using VPP vswitch running in Linux user-mode. Multiple service instances share the vswitch that in turn provides per service chain forwarding context(s). In order to provide a most complete picture, each network topology and service configuration is tested in different service density setups by varying two parameters:

- Number of service instances (e.g. 1,2,4..10).
- Number of NFs per service instance (e.g. 1,2,4..10).

Implementation of NFV service density tests in CSIT-1908.1 is using two NF applications:

- VNF: VPP of the same version as vswitch running in KVM VM, configured with /8 IPv4 prefix routing.
- CNF: VPP of the same version as vswitch running in Docker Container, configured with /8 IPv4 prefix routing.

Tests are designed such that in all tested cases VPP vswitch is the most stressed application, as for each flow vswitch is processing each packet multiple times, whereas VNFs and CNFs process each packets only once. To that end, all VNFs and CNFs are allocated enough resources to not become a bottleneck.

Service Configurations

Following NFV network topologies and configurations are tested:

- VNF Service Chains (VSC) with L2 vswitch
 - *Network Topology*: Sets of VNFs dual-homed to VPP vswitch over virtio-vhost links. Each set belongs to separate service instance.
 - *Network Configuration*: VPP L2 bridge-domain contexts form logical service chains of VNF sets and connect each chain to physical interfaces.
- CNF Service Chains (CSC) with L2 vswitch
 - *Network Topology*: Sets of CNFs dual-homed to VPP vswitch over memif links. Each set belongs to separate service instance.
 - *Network Configuration*: VPP L2 bridge-domain contexts form logical service chains of CNF sets and connect each chain to physical interfaces.
- CNF Service Pipelines (CSP) with L2 vswitch

- *Network Topology*: Sets of CNFs connected into pipelines over a series of memif links, with edge CNFs single-homed to VPP vswitch over memif links. Each set belongs to separate service instance.
- *Network Configuration*: VPP L2 bridge-domain contexts connect each CNF pipeline to physical interfaces.

Thread-to-Core Mapping

CSIT defines specific ratios for mapping software threads of vswitch and VNFs/CNFs to physical cores, with separate ratios defined for main control threads and data-plane threads.

In CSIT-1908.1 NFV service density tests run on Intel Xeon testbeds with Intel Hyper-Threading enabled, so each physical core is associated with a pair of sibling logical cores corresponding to the hyper-threads.

CSIT-1908.1 executes tests with the following software thread to physical core mapping ratios:

- vSwitch
 - Data-plane on single core
 - * (main:core) = (1:1) => 1mt1c - 1 main thread on 1 core.
 - * (data:core) = (1:1) => 2dt1c - 2 Data-plane Threads on 1 Core.
 - Data-plane on two cores
 - * (main:core) = (1:1) => 1mt1c - 1 Main Thread on 1 Core.
 - * (data:core) = (1:2) => 4dt2c - 4 Data-plane Threads on 2 Cores.
- VNF and CNF
 - Data-plane on single core
 - * (main:core) = (2:1) => 2mt1c - 2 Main Threads on 1 Core, 1 Thread per NF, core shared between two NFs.
 - * (data:core) = (1:1) => 2dt1c - 2 Data-plane Threads on 1 Core per NF.
 - Data-plane on single logical core (Two NFs per physical core)
 - * (main:core) = (2:1) => 2mt1c - 2 Main Threads on 1 Core, 1 Thread per NF, core shared between two NFs.
 - * (data:core) = (2:1) => 2dt1c - 2 Data-plane Threads on 1 Core, 1 Thread per NF, core shared between two NFs.

Maximum tested service densities are limited by a number of physical cores per NUMA. CSIT-1908.1 allocates cores within NUMA0. Support for multi NUMA tests is to be added in future release.

1.5.13 VPP_Device Functional

CSIT-1908.1 includes VPP_Device test environment for functional VPP device tests integrated into LFN CI/CD infrastructure. VPP_Device tests run on 1-Node testbeds (1n-skx, 1n-arm) and rely on Linux SRIOV Virtual Function (VF), dot1q VLAN tagging and external loopback cables to facilitate packet passing over external physical links. Initial focus is on few baseline tests. Existing CSIT Performance tests can be moved to VPP_Device framework. RF test definition code stays unchanged with the exception of traffic generator related L2 KWs.

1.5.14 IPsec on Intel QAT

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-1908.1 implements following IPsec test cases:

- AES-GCM, CBC-SHA1 ciphers, in combination with IPv4 routed-forwarding with Intel x1710 NIC.
- CBC-SHA1 ciphers, in combination with LISP-GPE overlay tunneling for IPv4-over-IPv4 with Intel x1710 NIC.

1.5.15 TRex Traffic Generator

Usage

TRex traffic generator²² is used for all CSIT performance tests. TRex stateless mode is used to measure NDR and PDR throughputs using MLRsearch and to measure maximum transfer rate in MRR tests.

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](#)²³.
- 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 > /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.api.STLClient`.

Measuring Packet Loss

Following sequence is followed to measure 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.

²² <https://wiki.fd.io/view/TRex>

²³ https://git.fd.io/csit/tree/resources/tools/trex/trex_installer.sh?h=rls1908_1

2.1 Overview

VPP performance test results are reported for all three physical testbed types present in FD.io labs: 3-Node Xeon Haswell (3n-hsw), 3-Node Xeon Skylake (3n-skx), 2-Node Xeon Skylake (2n-skx) and installed NIC models. For description of physical testbeds used for VPP performance tests please refer to *Physical Testbeds* (page 4).

2.1.1 Logical Topologies

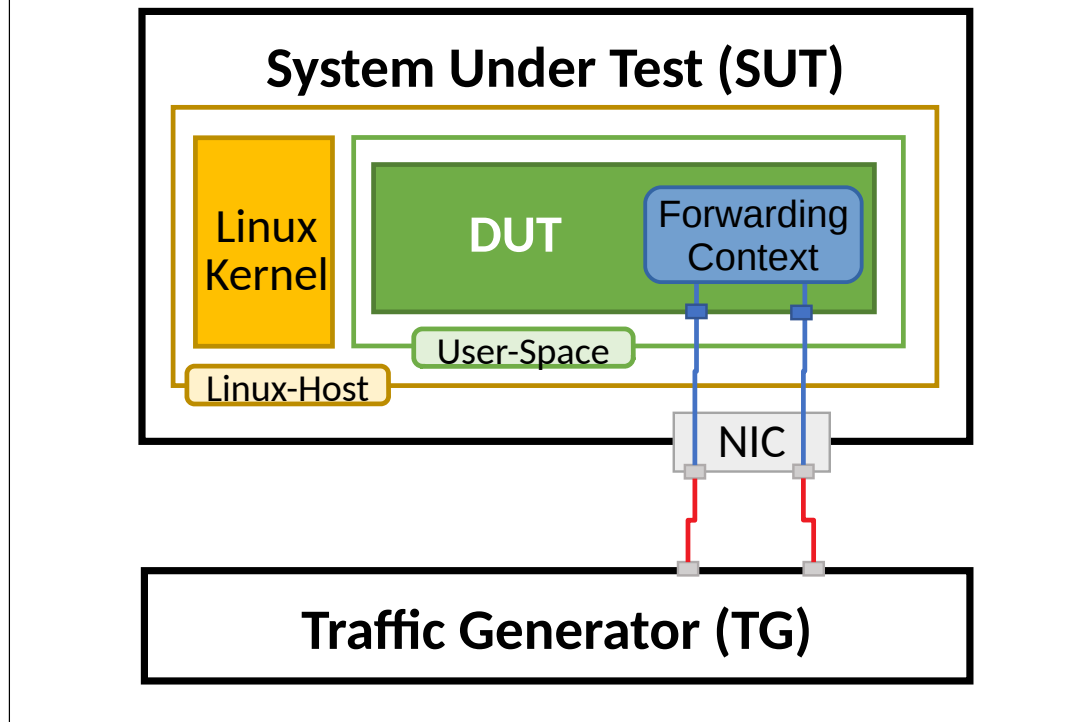
CSIT VPP performance tests are executed on physical testbeds described in *Physical Testbeds* (page 4). Based on the packet path thru server SUTs, three distinct logical topology types are used for VPP DUT data plane testing:

1. NIC-to-NIC switching topologies.
2. VM service switching topologies.
3. Container service switching topologies.


NIC-to-NIC Switching

The simplest logical topology for software data plane application like VPP is NIC-to-NIC switching. Tested topologies for 2-Node and 3-Node testbeds are shown in figures below.

2-Node Topology: NIC-to-NIC Switching



3-Node Topology: NIC-to-NIC Switching



Server Systems Under Test (SUT) run VPP application in Linux user-mode as a Device Under Test (DUT). Server Traffic Generator (TG) runs T-Rex application. Physical connectivity between SUTs and TG is provided using different drivers and NIC models that need to be tested for performance (packet/bandwidth throughput and latency).

From SUT and DUT perspectives, all performance tests involve forwarding packets between two (or more) physical Ethernet ports (10GE, 25GE, 40GE, 100GE). In most cases both physical ports on SUT are located on the same NIC. The only exceptions are link bonding and 100GE tests. In the latter case only one port per NIC can be driven at linerate due to PCIe Gen3 x16 slot bandwidth limitations. 100GE NICs are not supported in PCIe Gen3 x8 slots.

Note that reported VPP DUT performance results are specific to the SUTs tested. SUTs with other processors than the ones used in FD.io lab are likely to yield different results. A good rule of thumb, that can be applied to estimate VPP packet throughput for NIC-to-NIC switching topology, is to expect the forwarding performance to be proportional to processor core frequency for the same processor architecture, assuming processor is the only limiting factor and all other SUT parameters are equivalent to FD.io CSIT environment.


VM Service Switching


VM service switching topology test cases require VPP DUT to communicate with Virtual Machines (VMs) over vhost-user virtual interfaces.

Two types of VM service topologies are tested in CSIT-1908.1:

1. "Parallel" topology with packets flowing within SUT from NIC(s) via VPP DUT to VM, back to VPP DUT, then out thru NIC(s).
2. "Chained" topology (a.k.a. "Snake") with packets flowing within SUT from NIC(s) via VPP DUT to VM, back to VPP DUT, then to the next VM, back to VPP DUT and so on and so forth until the last VM in a chain, then back to VPP DUT and out thru NIC(s).

For each of the above topologies, VPP DUT is tested in a range of L2 or IPv4/IPv6 configurations depending on the test suite. Sample VPP DUT "Chained" VM service topologies for 2-Node and 3-Node testbeds with each SUT running N of VM instances is shown in the figures below.





In “Chained” VM topologies, packets are switched by VPP DUT multiple times: twice for a single VM, three times for two VMs, $N+1$ times for N VMs. Hence the external throughput rates measured by TG and listed in this report must be multiplied by $N+1$ to represent the actual VPP DUT aggregate packet forwarding rate.

For “Parallel” service topology packets are always switched twice by VPP DUT per service chain.

Note that reported VPP DUT performance results are specific to the SUTs tested. SUTs with other processor than the ones used in FD.io lab are likely to yield different results. Similarly to NIC-to-NIC switching topology, here one can also expect the forwarding performance to be proportional to processor core frequency for the same processor architecture, assuming processor is the only limiting factor. However due to much higher dependency on intensive memory operations in VM service chained topologies and sensitivity to Linux scheduler settings and behaviour, this estimation may not always yield good enough accuracy.

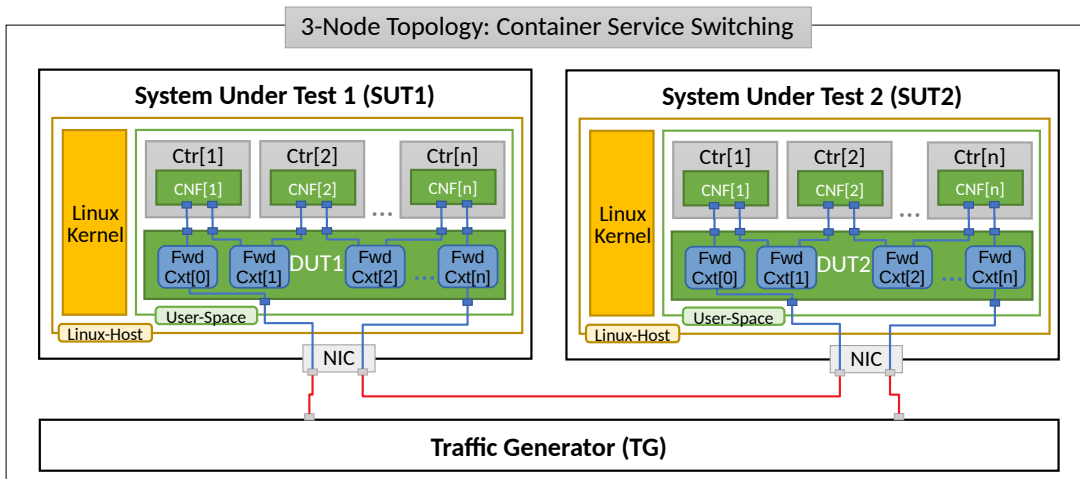
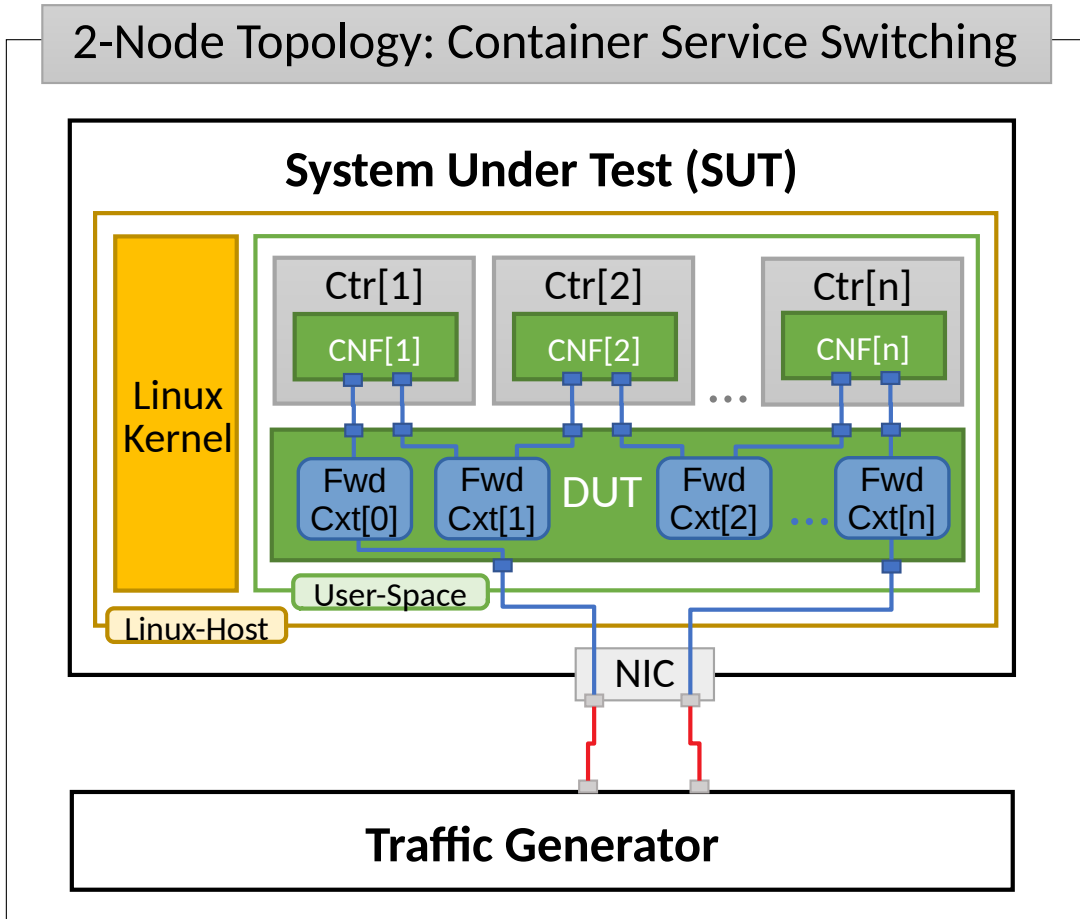
Container Service Switching

Container service switching topology test cases require VPP DUT to communicate with Containers (Ctrs) over memif virtual interfaces.

Three types of VM service topologies are tested in CSIT-1908.1:

1. “Parallel” topology with packets flowing within SUT from NIC(s) via VPP DUT to Container, back to VPP DUT, then out thru NIC(s).
2. “Chained” topology (a.k.a. “Snake”) with packets flowing within SUT from NIC(s) via VPP DUT to Container, back to VPP DUT, then to the next Container, back to VPP DUT and so on and so forth until the last Container in a chain, then back to VPP DUT and out thru NIC(s).
3. “Horizontal” topology with packets flowing within SUT from NIC(s) via VPP DUT to Container, then via “horizontal” memif to the next Container, and so on and so forth until the last Container, then back to VPP DUT and out thru NIC(s).

For each of the above topologies, VPP DUT is tested in a range of L2 or IPv4/IPv6 configurations depending on the test suite. Sample VPP DUT “Chained” Container service topologies for 2-Node and 3-Node testbeds with each SUT running N of Container instances is shown in the figures below.



In “Chained” Container topologies, packets are switched by VPP DUT multiple times: twice for a single Container, three times for two Containers, $N+1$ times for N Containers. Hence the external throughput rates measured by TG and listed in this report must be multiplied by $N+1$ to represent the actual VPP DUT aggregate packet forwarding rate.

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

Note that reported VPP DUT performance results are specific to the SUTs tested. SUTs with other processor than the ones used in FD.io lab are likely to yield different results. Similarly to NIC-to-NIC switching topology, here one can also expect the forwarding performance to be proportional to processor core frequency for the same processor architecture, assuming processor is the only limiting factor. However due

to much higher dependency on intensive memory operations in Container service chained topologies and sensitivity to Linux scheduler settings and behaviour, this estimation may not always yield good enough accuracy.

2.1.2 Performance Tests Coverage

Performance tests measure following metrics for tested VPP DUT topologies and configurations:

- Packet Throughput: measured in accordance with **RFC 2544**²⁴, using FD.io CSIT Multiple Loss Ratio search (MLRsearch), an optimized binary search algorithm, producing throughput at different Packet Loss Ratio (PLR) values:
 - Non Drop Rate (NDR): packet throughput at PLR=0%.
 - Partial Drop Rate (PDR): packet throughput at PLR=0.5%.
- One-Way Packet Latency: measured at different offered packet loads:
 - 100% of discovered NDR throughput.
 - 100% of discovered PDR throughput.
- Maximum Receive Rate (MRR): measure 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.

CSIT-1908.1 includes following VPP data plane functionality performance tested across a range of NIC drivers and NIC models:

²⁴ <https://tools.ietf.org/html/rfc2544.html>

Functionality	Description
ACL	L2 Bridge-Domain switching and IPv4 and IPv6 routing with iACL and oACL IP address, MAC address and L4 port security.
COP	IPv4 and IPv6 routing with COP address security.
IPv4	IPv4 routing.
IPv6	IPv6 routing.
IPv4 Scale	IPv4 routing with 20k, 200k and 2M FIB entries.
IPv6 Scale	IPv6 routing with 20k, 200k and 2M FIB entries.
IPSecHW	IPSec encryption with AES-GCM, CBC-SHA-256 ciphers, in combination with IPv4 routing. Intel QAT HW acceleration.
IPSec+LISP	IPSec encryption with CBC-SHA1 ciphers, in combination with LISP-GPE overlay tunneling for IPv4-over-IPv4.
IPSecSW	IPSec encryption with AES-GCM, CBC-SHA-256 ciphers, in combination with IPv4 routing.
K8s Containers Memif	K8s orchestrated container VPP service chain topologies connected over the memif virtual interface.
KVM VMs vhost-user	Virtual topologies with service chains of 1 VM using vhost-user interfaces, with different VPP forwarding modes incl. L2XC, L2BD, VXLAN with L2BD, IPv4 routing.
L2BD	L2 Bridge-Domain switching of untagged Ethernet frames with MAC learning; disabled MAC learning i.e. static MAC tests to be added.
L2BD Scale	L2 Bridge-Domain switching 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.
L2XC	L2 Cross-Connect switching of untagged, dot1q, dot1ad VLAN tagged Ethernet frames.
LISP	LISP overlay tunneling for IPv4-over-IPv4, IPv6-over-IPv4, IPv6-over-IPv6, IPv4-over-IPv6 in IPv4 and IPv6 routing modes.
LXC/DRC Containers Memif	Container VPP memif virtual interface tests with different VPP forwarding modes incl. L2XC, L2BD.
NAT	(Source) Network Address Translation tests with varying number of users and ports per user.
QoS Policer	Ingress packet rate measuring, marking and limiting (IPv4).
SRv6 Routing	Segment Routing IPv6 tests.
VPP TCP/IP stack	Tests of VPP TCP/IP stack used with VPP built-in HTTP server.
VTS	Virtual Topology System use case tests combining VXLAN overlay tunneling with L2BD, ACL and KVM VM vhost-user features.
VXLAN	VXLAN overlay tunnelling integration with L2XC and L2BD.

Execution of performance tests takes time, especially the throughput tests. Due to limited HW testbed resources available within FD.io labs hosted by LF, the number of tests for some NIC models has been limited to few baseline tests.

2.1.3 Performance Tests Naming

FD.io CSIT-1908.1 follows a common structured naming convention for all performance and system functional tests, introduced in CSIT-17.01.

The naming should be intuitive for majority of the tests. Complete description of FD.io CSIT test naming convention is provided on *Test Naming* (page 459).

2.2 Release Notes

2.2.1 Changes in CSIT-1908.1

1. VPP PERFORMANCE TESTS

- **Service density 2n-skx tests:** Added higher NF density tests with 802.1q (vlan) and VXLAN encapsulation from Traffic Generator.
- **GBP tests:** Added GBP (Group Based Policy) routing test cases with 802.1q (vlan) external traffic.
- **AVF IPv4 scale tests:** Increased coverage of AVF IPv4 base and scale test cases (Fortville NICs only).
- **2n-skx tests:** Increased coverage of selected (COP, iACL, Policer) test cases.
- **IPsec scale tests:** Added IPsec interface mode scale tests with 1, 40, 400, 1000, 5000, 10000, 20000, 40000, 60000 IPsec tunnels. Removed DPDK backend dependency. Major IPsec test code refactoring.
- **Hoststack TCP/IP tests:** Major refactor of Hoststack TCP performance tests using WRK generator talking to the VPP HTTP static server plugin measuring connections per second and requests per second.
- **Changed methodology of dot1q tests in 2-Node testbeds:** dot1q encapsulation is now used on both links of SUT. Previously dot1q was used only on a single link with the other link carrying untagged Ethernet frames. This change results in slightly lower throughput in CSIT-1908 for all dot1q tests in all 2-Node testbeds.
- **KVM VM vhost-user tests:** completed move to Kernel-VM for all tests. In addition to running DPDK Testpmd as VM workload, new tests created with VPP as VM workload. VPP in VM is the same version as the DUT VPP (acting as vSwitch) and its configuration depends on the test type. For all L2 Ethernet Switching tests it's vpp-l2xc (L2 cross-connect), for all IPv4 Routing tests it's vpp-ip4 (VPP IPv4 routing).

2. TEST FRAMEWORK

- **CSIT PAPI Support:** Finished conversion of CSIT VAT L1 keywords to PAPI L1 KWs in CSIT using VPP Python bindings (VPP PAPI). Redesign of key components of PAPI Socket Executor and PAPI history. Due to issues with PAPI performance, VAT is still used in CSIT for all VPP scale tests. See known issues below.
- **General Code Housekeeping:** Ongoing RF keywords optimizations, removal of redundant RF keywords and aligning of suite/test setup/teardowns.

3. PRESENTATION AND ANALYTICS LAYER

- **Graphs Layout Improvements:** Improved performance graphs layout for better readability and maintenance: test grouping, axis labels, descriptions, other informative decoration.

2.2.2 Known Issues

List of known issues in CSIT-1908.1 for VPP performance tests:

#	Ji- ralD	Issue Description
1	CSIT-570 ²⁵	Sporadic (1 in 200) NDR discovery test failures on x520. 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	VPP-662 ²⁶	9000B packets not supported by NICs VIC1227 and VIC1387.
3	CSIT-1503 ²⁷	[TRex-519 ²⁸] XL710/XXV710 with FW 6.0.1 will have Rx drop rate of 27MPPS.
4	CSIT-1498 ²⁹	Memif tests are sporadically failing on initialization of memif connection.
5	CSIT-1499 ³⁰	AVF tests are sporadically failing on initialization of AVF interface.
6	VPP-1676 ³¹	9000B ip4 memif errors - ip4-input: ip4 length > l2 length. IP4 jumbo frames (9000B) are dropped in case of tests with memif.
7	VPP-1677 ³²	9000B ip4 nat44: VPP crash + core dump. VPP crashes very often in case that NAT44 is configured and it has to process IP4 jumbo frames (9000B).
8	CSIT-1591 ³³ VPP-1763 ³⁴	All CSIT scale tests can not use PAPI due to much slower performance compared to VAT/CLI (it takes much longer to program VPP). This needs to be addressed on the PAPI side.
9	CSIT-1592 ³⁵ VPP-1764 ³⁶	VPP memif API does not enable memif zero-copy, resulting in different memif configuration vs. previously tested VAT/CLI where memif zero-copy was enabled by default. Needs to be fixed in VPP.
10	VPP-1675 ³⁷	IPv4 IPSEC 9000B packet tests are failing as no packet is forwarded. Reason: chained buffers are not supported.
11	CSIT-1593 ³⁸	IPv4 AVF 9000B packet tests are failing on 3n-skx while passing on 2n-skx.

²⁵ <https://jira.fd.io/browse/CSIT-570>

²⁶ <https://jira.fd.io/browse/VPP-662>

²⁷ <https://jira.fd.io/browse/CSIT-1503>

²⁸ <https://trex-tgn.cisco.com/youtrack/issue/trex-519>

²⁹ <https://jira.fd.io/browse/CSIT-1498>

³⁰ <https://jira.fd.io/browse/CSIT-1499>

³¹ <https://jira.fd.io/browse/VPP-1676>

³² <https://jira.fd.io/browse/VPP-1677>

³³ <https://jira.fd.io/browse/CSIT-1499>

³⁴ <https://jira.fd.io/browse/VPP-1763>

³⁵ <https://jira.fd.io/browse/CSIT-1592>

³⁶ <https://jira.fd.io/browse/VPP-1764>

³⁷ <https://jira.fd.io/browse/VPP-1675>

³⁸ <https://jira.fd.io/browse/CSIT-1593>

2.3 Packet Throughput

Throughput graphs are generated based on the results data obtained from the CSIT-1908.1 test jobs. In order to verify benchmark results repeatability selected, CSIT performance tests are executed multiple times (target: 10 times) on each physical testbed type. Box-and-Whisker plots are used to display variations in measured throughput values.

Lists of tests selected for multiple execution and graphing are captured per testbed type in `test_select_list_{testbed_type}.md`³⁹ files.

Graphs are split into sections as follows:

1. Header 1: VPP packet path and lookup types

- **L2 Ethernet Switching:** L2 bridge-domain, L2 cross-connect and L2 patch
- **IPv4 Routing:** IPv4 routing with /32 prefixes
- **IPv6 Routing:** IPv6 routing with /128 prefixes
- **SRv6 Routing:** SRv6 with IPv6 routing
- **IPv4 Tunnels:** IPv4 overlay tunnels
- **KVM VMs vhost-user:** KVM VMs connected over virtio and vhost-user interfaces
- **LXC/DRC Container Memif:** Linux containers and Docker containers connected over Memif interfaces
- **IPsec IPv4 Routing:** IPsec encryption/decryption with IPv4 routing
- **Virtual Topology System:** VXLAN configurations with L2 bridge-domains

2. Header 2: testbeds and NIC models

- section name format:
 - {testbed_type}-{nic_model}
- **testbed_type:**
 - 2n-skx: 2-node Xeon Skylake
 - 3n-skx: 3-node Xeon Skylake
 - 3n-hsw: 3-node Xeon Haswell
 - 3n-tsh: 3-node Arm TaiShan
 - 2n-dnv: 2-node Atom Denverton
 - 3n-dnv: 3-node Atom Denverton
- **nic_model:**
 - xxv710: xxv710 2p25GE Intel (Fortville)
 - x710: x710 4p10GE Intel (Fortville)
 - xl710: xl710 2p40GE Intel (Fortville)
 - x520: x520 2p10GE Intel (Niantic)
 - x553: x553 2p10GE Intel (Niantic)

3. Header 3: test group names

- section name format:
 - {frame_size}-{worker_thread_core_cfg}-{vpp_functionality}-{vpp_lookup_type}-{baseline_scale}-{nic_driver}

³⁹ https://git.fd.io/csit/tree/docs/job_specs

- **frame_size:**
 - 64b: 64 byte frames, smallest frame size for untagged IPv4 packets
 - 78b: 78 byte frames, smallest frame size for untagged IPv6 packets
 - 114b: VXLAN encapsulated L2 frames
 - imix: a sequence of (7x64B, 4x570, 1x1518) byte frames
- **worker_thread_core_cfg:**
 - 1t1c: 1 worker thread on 1 core, hyper-threading not used
 - 2t1c: 2 worker threads on 1 core, hyper-threading used
- **vpp_functionality** (optional):
 - features: including input-acl, output-acl, macip-iacl, nat44
 - srv6: srv6 encap/decap, proxy
 - link-bonding: L2 link aggregation with 1 or 2 bonded links
 - ipsec: IPsec encryption/decryption with different ciphers
 - vts: Virtual Topology System specific tests
- **vpp_lookup_type:**
 - l2switching, ip4routing, ip6routing, ip4tunnel, vhost, memif
- **baseline_scale:**
 - base: baseline tests with less than 10 forwarding entries
 - scale: scale tests with up to 2 million forwarding entries
 - base-scale: both baseline and scale tests grouped together
- **nic_driver:**
 - avf: VPP native avf driver for Intel Fortville NICs
 - i40e: dpdk poll mode driver for Intel Fortville NICs
 - ixgbe: dpdk poll mode driver for Intel Niantic NICs

For each 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. Outliers are plotted as individual points.

Additional information about graph data:

1. **Graph Title:** describes tested packet path, testbed topology, processor model, NIC model, packet size, number of cores and threads used by data plane workers and indication of VPP DUT configuration.
2. **X-axis Labels:** indices of individual test suites as listed in Graph Legend.
3. **Y-axis Labels:** measured Packets Per Second [pps] throughput values.
4. **Graph Legend:** lists X-axis indices with associated CSIT test suites executed to generate graphed test results.
5. **Hover Information:** lists minimum, first quartile, median, third quartile, and maximum. If either type of outlier is present the whisker on the appropriate side is taken to $1.5 \times \text{IQR}$ from the quartile (the "inner fence") rather than the max or min, and individual outlying data points are displayed as unfilled circles (for suspected outliers) or filled circles (for outliers). (The "outer fence" is $3 \times \text{IQR}$ from the quartile.)

Note: Test results have been generated by [FD.io test executor vpp performance job 2n-skx⁴⁰](#), [FD.io test executor vpp performance job 3n-skx⁴¹](#), [FD.io test executor vpp performance job 3n-hsw⁴²](#), [FD.io test executor vpp performance job 3n-tsh⁴³](#), [FD.io test executor vpp performance job 2n-dnv⁴⁴](#) and [FD.io test executor vpp performance job 3n-dnv⁴⁵](#) with RF result files csit-vpp-perf-1908_1-*.zip [archived here](#). Required per test case data set size is **10**, but for VPP tests the actual size varies per test case and is ≤ 10 .

⁴⁰ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-2n-skx

⁴¹ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-skx

⁴² https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-hsw

⁴³ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-tsh

⁴⁴ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-2n-dnv

⁴⁵ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-dnv

2.3.1 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.


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

⁴⁶ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1908>


2n-skx-xxv710


64b-2t1c-l2switching-base-avf







64b-2t1c-l2switching-base-i40e






64b-2t1c-l2switching-base-scale-i40e






2n-skx-x710


64b-2t1c-l2switching-base-scale-avf-i40e







3n-skx-xxv710


64b-2t1c-l2switching-base-avf







64b-2t1c-l2switching-base-i40e






64b-2t1c-l2switching-base-scale-i40e






3n-skx-x710

64b-2t1c-l2switching-base-scale-avf-i40e







3n-hsw-xl710

64b-1t1c-l2switching-base-scale-i40e

2n-dnv-x553


64b-1t1c-l2switching-base-ixgbe






3n-dnv-x553

64b-1t1c-l2switching-base-ixgbe





2.3.2 IPv4 Routing


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.


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

⁴⁷ <https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls1908>


2n-skx-xxv710


64b-2t1c-ip4routing-base-scale-avf







64b-2t1c-ip4routing-base-scale-i40e






64b-2t1c-features-ip4routing-base-i40e






2n-skx-x710


64b-2t1c-ip4routing-base-scale-avf-i40e







3n-skx-xxv710


64b-2t1c-ip4routing-base-scale-avf







64b-2t1c-ip4routing-base-scale-i40e





64b-2t1c-features-ip4routing-base-i40e







3n-skx-x710

64b-2t1c-ip4routing-base-scale-avf-i40e

3n-hsw-xl710


64b-1t1c-ip4routing-base-scale-i40e






2n-dnv-x553


64b-1t1c-ip4routing-base-ixgbe






3n-dnv-x553

64b-1t1c-ip4routing-base-ixgbe





2.3.3 IPv6 Routing


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.


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

⁴⁸ <https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls1908>

2n-skx-xxv710


78b-2t1c-ip6routing-base-scale-i40e






2n-skx-x710


78b-2t1c-ip6routing-base-scale-i40e






3n-skx-xxv710


78b-2t1c-ip6routing-base-scale-i40e






3n-skx-x710

78b-2t1c-ip6routing-base-scale-i40e







3n-hsw-xl710

78b-1t1c-ip6routing-base-scale-i40e

2n-dnv-x553


78b-1t1c-ip6routing-base-ixgbe






3n-dnv-x553

78b-1t1c-ip6routing-base-ixgbe





2.3.4 IPv4 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.


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

⁴⁹ https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls1908

3n-skx-xxv710


64b-2t1c-ip4tunnel-base-scale-i40e






3n-hsw-xl710


64b-1t1c-ip4tunnel-base-scale-i40e






3n-dnv-x553

64b-1t1c-ip4tunnel-base-ixgbe





2.3.5 KVM VMs vhost-user


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.


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

⁵⁰ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1908


2n-skx-xxv710


64b-2t1c-vhost-base-i40e-testpmd






64b-2t1c-vhost-base-i40e-vpp







3n-skx-xxv710


64b-2t1c-vhost-base-i40e-testpmd






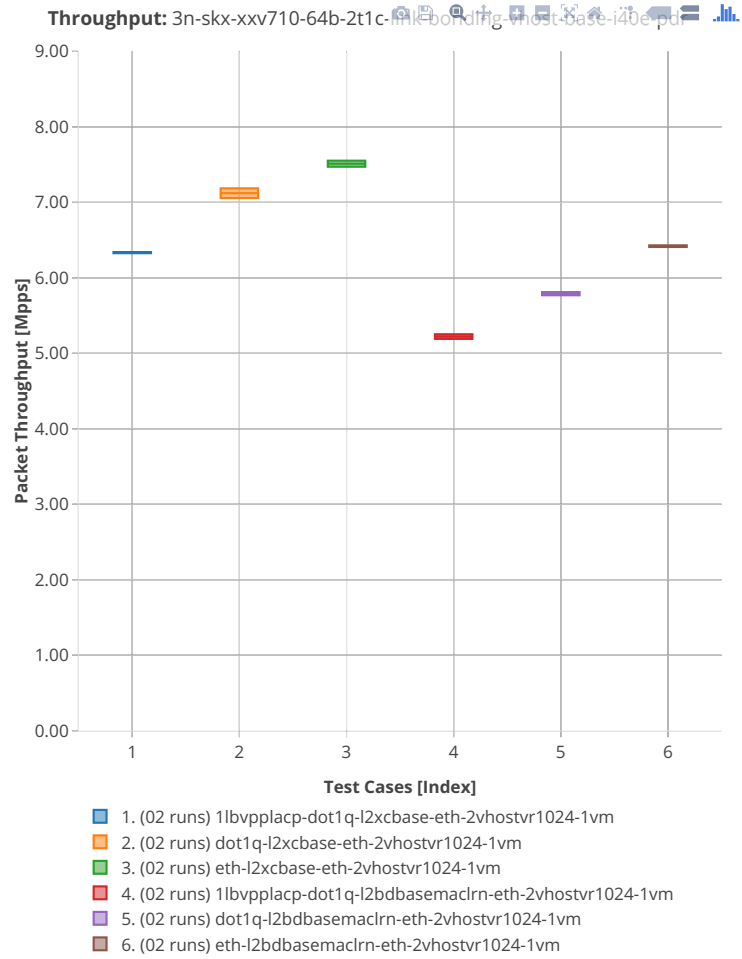
64b-2t1c-vhost-base-i40e-vpp



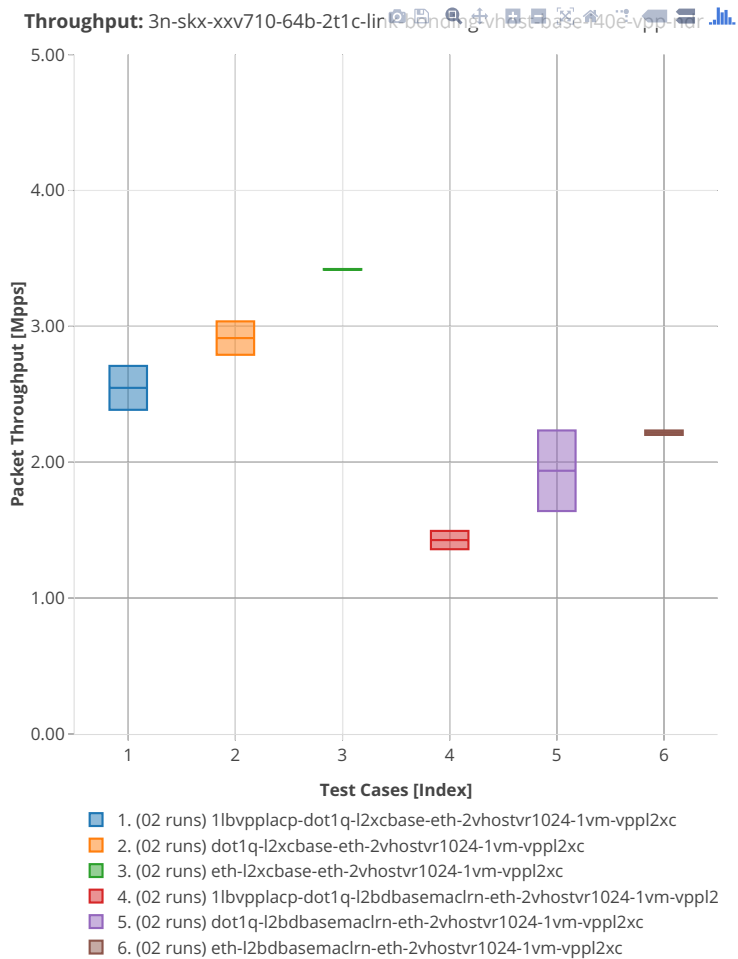



64b-2t1c-link-bonding-vhost-base-i40e-testpmd






64b-2t1c-link-bonding-vhost-base-i40e-vpp







3n-skx-x710


64b-2t1c-link-bonding-vhost-base-i40e-testpmd






64b-2t1c-link-bonding-vhost-base-i40e-vpp







3n-hsw-xl710


64b-1t1c-vhost-base-i40e-testpmd







64b-1t1c-vhost-base-i40e-vpp







64b-1t1c-link-bonding-vhost-base-i40e-testpmd





64b-1t1c-link-bonding-vhost-base-i40e-vpp





2.3.6 LXC/DRC Container Memif


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.


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

⁵¹ https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls1908

2n-skx-xxv710


64b-2t1c-memif-base-i40e






3n-skx-xxv710


64b-2t1c-memif-base-i40e






3n-hsw-xl710

64b-1t1c-memif-base-i40e





2.3.7 IPsec IPv4 Routing


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.


CSIT source code for the test cases used for plots can be found in [CSIT git repository](#)⁵².

⁵² <https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls1908>

3n-skx-xxv710


imix-2t1c-ipsec-ip4routing-base-scale-i40e







3n-hsw-xl710


imix-1t1c-ipsec-ip4routing-base-scale-sw-i40e






imix-1t1c-ipsec-ip4routing-base-scale-hw-i40e






3n-dnv-x553

imix-1t1c-ipsec-ip4routing-base-scale-sw-ixgbe





2.4 Speedup Multi-Core

Speedup Multi-Core throughput graphs are generated by multiple executions of the same performance tests across physical testbeds hosted LF FD.io labs: 3n-hsw, 2n-skx, 2n-skx. Grouped bars illustrate the 64B/78B packet throughput speedup ratio for 2- and 4-core multi- threaded VPP configurations relative to 1-core configurations.

Additional information about graph data:

1. **Graph Title:** describes tested packet path, testbed topology, processor model, NIC model, packet size used by data plane workers and indication of VPP DUT configuration.
2. **X-axis Labels:** number of cores.
3. **Y-axis Labels:** measured Packets Per Second [pps] throughput values.
4. **Graph Legend:** lists CSIT test suites executed to generate graphed test results.
5. **Hover Information:** lists number of runs executed, specific test substring, mean value of the measured packet throughput, calculated perfect throughput value, difference between measured and perfect values and relative speedup value.

Note: Test results have been generated by [FD.io test executor vpp performance job 2n-skx⁵³](#), [FD.io test executor vpp performance job 3n-skx⁵⁴](#), [FD.io test executor vpp performance job 3n-hsw⁵⁵](#), [FD.io test executor vpp performance job 3n-tsh⁵⁶](#), [FD.io test executor vpp performance job 2n-dnv⁵⁷](#) and [FD.io test executor vpp performance job 3n-dnv⁵⁸](#) with RF result files csit-vpp-perf-1908_1-*.zip [archived here](#). Required per test case data set size is **10**, but for VPP tests the actual size varies per test case and is ≤ 10 .

⁵³ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-2n-skx

⁵⁴ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-skx

⁵⁵ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-hsw

⁵⁶ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-tsh

⁵⁷ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-2n-dnv

⁵⁸ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-dnv

2.4.1 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).


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

⁵⁹ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1908>


2n-skx-xxv710


64b-l2switching-base-avf







64b-l2switching-base-i40e






64b-l2switching-base-scale-i40e






2n-skx-x710


64b-l2switching-base-scale-avf-i40e






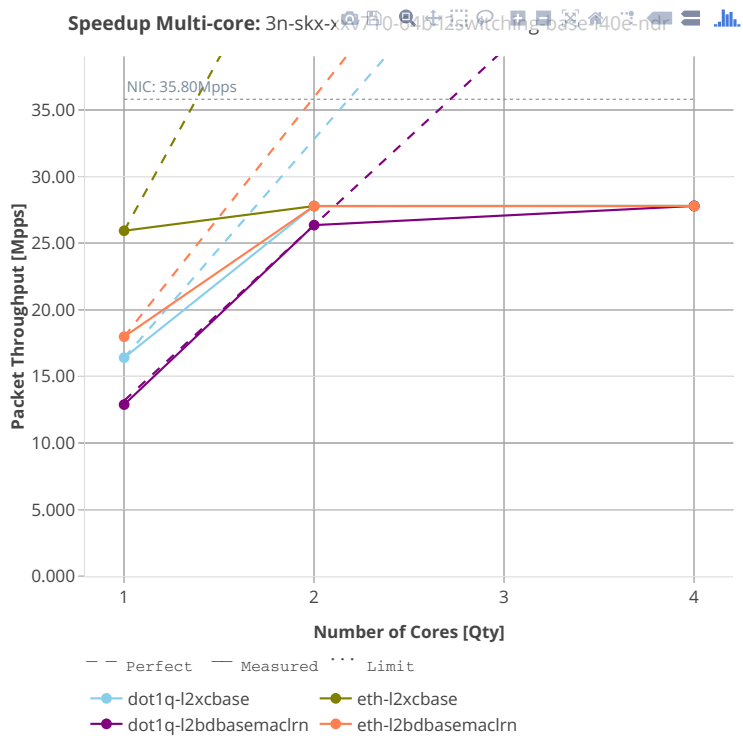
3n-skx-xxv710


64b-l2switching-base-avf







64b-l2switching-base-i40e






64b-l2switching-base-scale-i40e






3n-skx-x710

64b-l2switching-base-scale-avf-i40e







3n-hsw-xl710

64b-l2switching-base-scale-i40e

2n-dnv-x553


64b-l2switching-base-ixgbe






3n-dnv-x553

64b-l2switching-base-ixgbe





2.4.2 IPv4 Routing


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).


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

⁶⁰ <https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls1908>


2n-skx-xxv710


64b-ip4routing-base-scale-avf







64b-ip4routing-base-scale-i40e






64b-features-ip4routing-base-i40e






2n-skx-x710


64b-ip4routing-base-scale-avf-i40e







3n-skx-xxv710


64b-ip4routing-base-scale-avf







64b-ip4routing-base-scale-i40e





64b-features-ip4routing-base-i40e






3n-skx-x710

64b-ip4routing-base-scale-avf-i40e


3n-hsw-xl710


64b-ip4routing-base-scale-i40e



2n-dnv-x553


64b-ip4routing-base-ixgbe






3n-dnv-x553

64b-ip4routing-base-ixgbe





2.4.3 IPv6 Routing


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).


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

⁶¹ <https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls1908>

2n-skx-xxv710


78b-ip6routing-base-scale-i40e






2n-skx-x710


78b-ip6routing-base-scale-i40e






3n-skx-xxv710


78b-ip6routing-base-scale-i40e






3n-skx-x710

78b-ip6routing-base-scale-i40e







3n-hsw-xl710

78b-ip6routing-base-scale-i40e

2n-dnv-x553


78b-ip6routing-base-ixgbe






3n-dnv-x553

78b-ip6routing-base-ixgbe





2.4.4 IPv4 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.

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

⁶² https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls1908

3n-skx-xxv710


64b-ip4tunnel-base-scale-i40e


3n-hsw-xl710

64b-ip4tunnel-base-scale-i40e

3n-dnv-x553

64b-ip4tunnel-base-ixgbe





2.4.5 KVM VMs vhost-user


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).


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

⁶³ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1908


2n-skx-xxv710


64b-vhost-base-i40e-testpmd





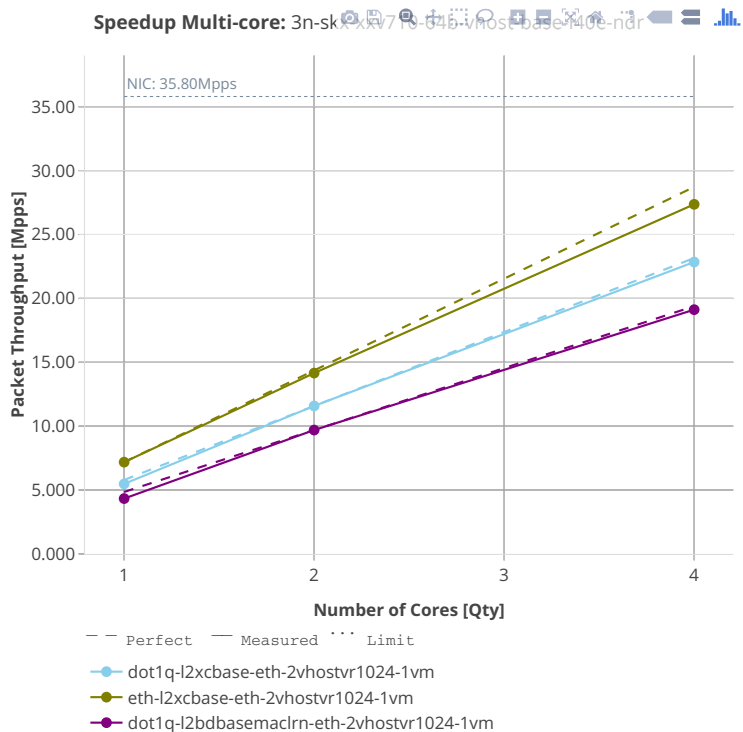
64b-vhost-base-i40e-vpp







3n-skx-xxv710


64b-vhost-base-i40e-testpmd







64b-vhost-base-i40e-vpp







64b-link-bonding-vhost-base-i40e-testpmd





64b-link-bonding-vhost-base-i40e-vpp






3n-skx-x710


64b-link-bonding-vhost-base-i40e-testpmd

64b-link-bonding-vhost-base-i40e-vpp

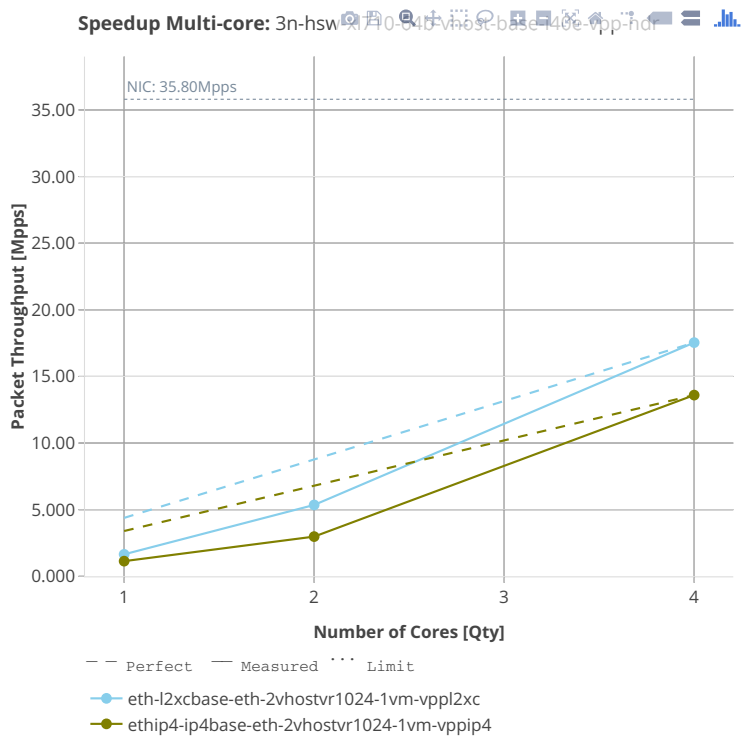
3n-hsw-xl710


64b-vhost-base-i40e-testpmd







64b-vhost-base-i40e-vpp





64b-link-bonding-vhost-base-i40e-vpp





2.4.6 LXC/DRC Container Memif

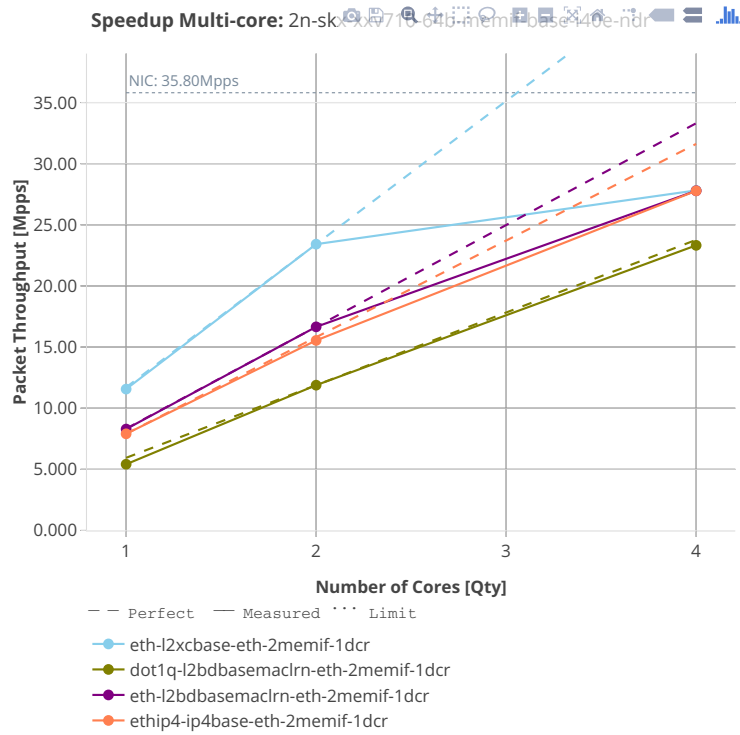
Following sections include Throughput Speedup Analysis for VPP multi-core multi-thread configurations with no Hyper-Threading, specifically for tested 2t2c (2threads, 2cores) and 4t4c scenarios. 1t1c throughput results are used as a reference for reported speedup ratio. Performance is reported for VPP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


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

⁶⁴ https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls1908

2n-skx-xxv710


64b-memif-base-i40e






3n-skx-xxv710


64b-memif-base-i40e






3n-hsw-xl710

64b-memif-base-i40e





2.4.7 IPsec IPv4 Routing


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.

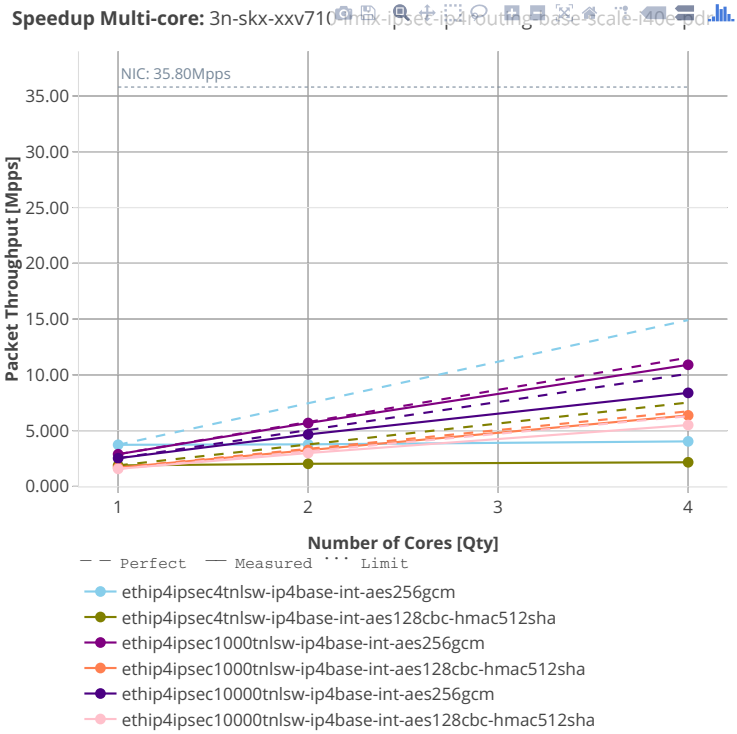
CSIT source code for the test cases used for plots can be found in [CSIT git repository](#)⁶⁵.

⁶⁵ <https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls1908>

3n-skx-xxv710


imix-ipsec-ip4routing-base-scale-i40e






3n-hsw-xl710


imix-ipsec-ip4routing-base-scale-sw-i40e

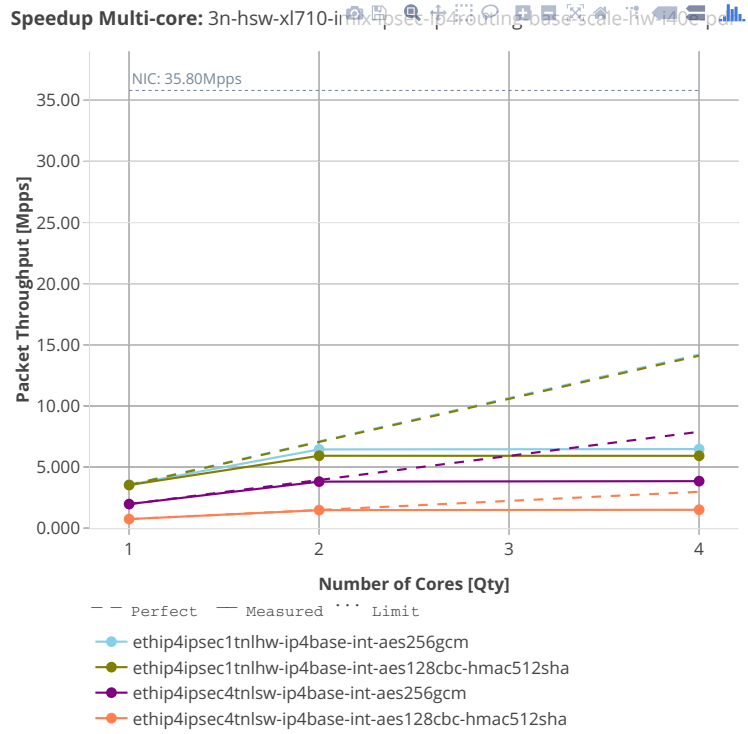


Speedup Multi-core: 3n-hsw-xl710-



imix-ipsec-ip4routing-base-scale-hw-i40e






3n-dnv-x553

imix-ipsec-ip4routing-base-scale-sw-ixgbe

Speedup Multi-core: 3n-dnv-x553-imix-ipsec-ip4routing-base-scale-sw-ixgbe



2.5 Packet Latency

Latency results are generated from a single execution of NDR discovery test across physical testbeds hosted LF FD.io labs: 3n-hsw, 2n-skx, 2n-skx. Box plots are used to show the Minimum, Median and Maximum packet latency per test.

Additional information about graph data:

1. **Graph Title:** describes tested packet path, testbed topology, processor model, NIC model, packet size, number of cores and threads used by data plane workers and indication of DUT configuration.
2. **X-axis Labels:** indices of individual test suites as listed in Graph Legend and direction of latency flow:
 - West-to-East: TGint1-to-SUT1-to-SUT2-to-TGint2.
 - East-to-West: TGint2-to-SUT2-to-SUT1-to-TGint1.
3. **Y-axis Labels:** measured packet latency values in [uSec].
4. **Graph Legend:** lists X-axis indices with associated CSIT test suites executed to generate graphed test results.
5. **Hover Information:** lists number of runs executed, specific test substring, packet flow direction, maximal, mean and minimal values of measured latencies.

Note: Test results have been generated by [FD.io test executor vpp performance job 2n-skx⁶⁶](#), [FD.io test executor vpp performance job 3n-skx⁶⁷](#), [FD.io test executor vpp performance job 3n-hsw⁶⁸](#) and [FD.io test executor vpp performance job 3n-tsh⁶⁹](#) with RF result files csit-vpp-perf-1908_1-*.zip [archived here](#).

⁶⁶ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-2n-skx

⁶⁷ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-skx

⁶⁸ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-hsw

⁶⁹ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-tsh

2.5.1 L2 Ethernet Switching


This section includes summary graphs of VPP Phy-to-Phy packet latency with L2 Ethernet switching measured at 100% 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.

CSIT source code for the test cases used for plots can be found in [CSIT git repository](#)⁷⁰.


⁷⁰ <https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls1908>

2n-skx-xxv710


64b-2t1c-l2switching-base-avf




64b-4t2c-l2switching-base-avf




64b-8t4c-l2switching-base-avf




64b-2t1c-l2switching-base-i40e




64b-4t2c-l2switching-base-i40e




64b-8t4c-l2switching-base-i40e



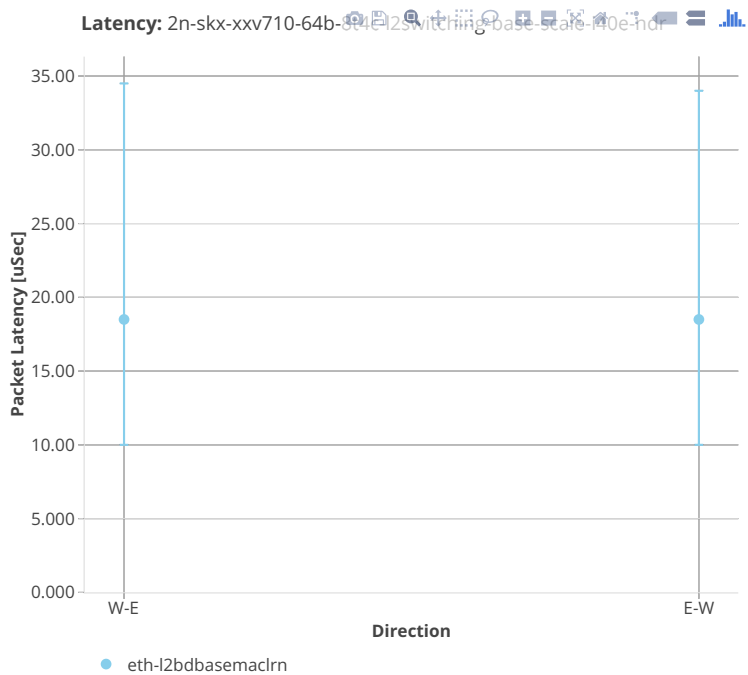
64b-2t1c-l2switching-base-scale-i40e



64b-4t2c-l2switching-base-scale-i40e




64b-8t4c-l2switching-base-scale-i40e




2n-skx-x710


64b-2t1c-l2switching-base-scale-avf-i40e



64b-4t2c-l2switching-base-scale-avf-i40e




64b-8t4c-l2switching-base-scale-avf-i40e




3n-skx-xxv710


64b-2t1c-l2switching-base-avf




64b-4t2c-l2switching-base-avf



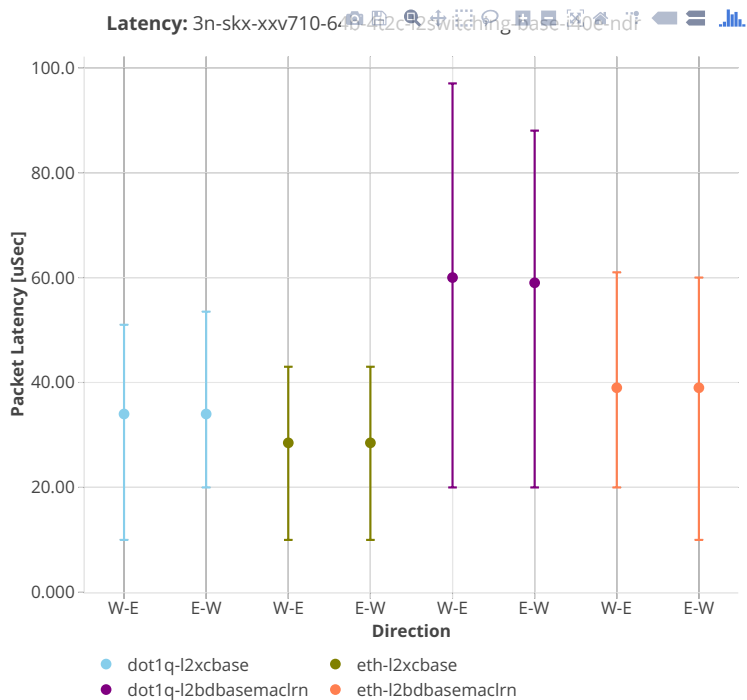
64b-8t4c-l2switching-base-avf



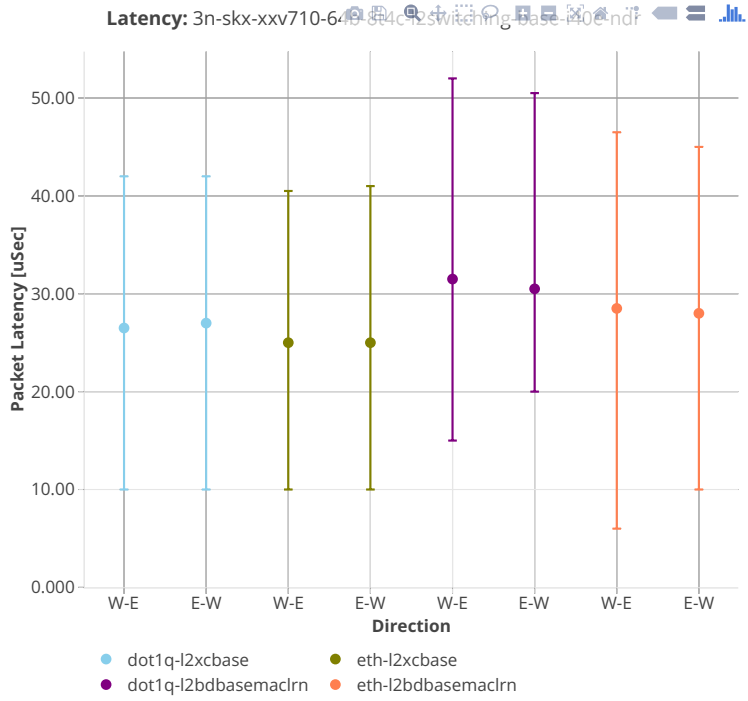
64b-2t1c-l2switching-base-i40e




64b-4t2c-l2switching-base-i40e




64b-8t4c-l2switching-base-i40e




64b-2t1c-l2switching-base-scale-i40e



64b-4t2c-l2switching-base-scale-i40e



64b-8t4c-l2switching-base-scale-i40e



64b-2t1c-features-l2switching-base-i40e

3n-skx-x710

64b-2t1c-l2switching-base-scale-avf-i40e

64b-4t2c-l2switching-base-scale-avf-i40e

64b-8t4c-l2switching-base-scale-avf-i40e

3n-hsw-xl710

64b-1t1c-l2switching-base-scale-i40e

64b-2t2c-l2switching-base-scale-i40e

64b-4t4c-l2switching-base-scale-i40e

2.5.2 IPv4 Routing


This section includes summary graphs of VPP Phy-to-Phy packet latency with IPv4 Routed-Forwarding measured at 100% 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.

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


⁷¹ <https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls1908>

2n-skx-xxv710


64b-2t1c-ip4routing-base-scale-avf




64b-4t2c-ip4routing-base-scale-avf




64b-8t4c-ip4routing-base-scale-avf




64b-2t1c-ip4routing-base-scale-i40e




64b-4t2c-ip4routing-base-scale-i40e




64b-8t4c-ip4routing-base-scale-i40e




64b-2t1c-features-ip4routing-base-i40e



64b-4t2c-features-ip4routing-base-i40e




64b-8t4c-features-ip4routing-base-i40e




2n-skx-x710


64b-2t1c-ip4routing-base-scale-avf-i40e



64b-4t2c-ip4routing-base-scale-avf-i40e




64b-8t4c-ip4routing-base-scale-avf-i40e




3n-skx-xxv710


64b-2t1c-ip4routing-base-scale-avf




64b-4t2c-ip4routing-base-scale-avf



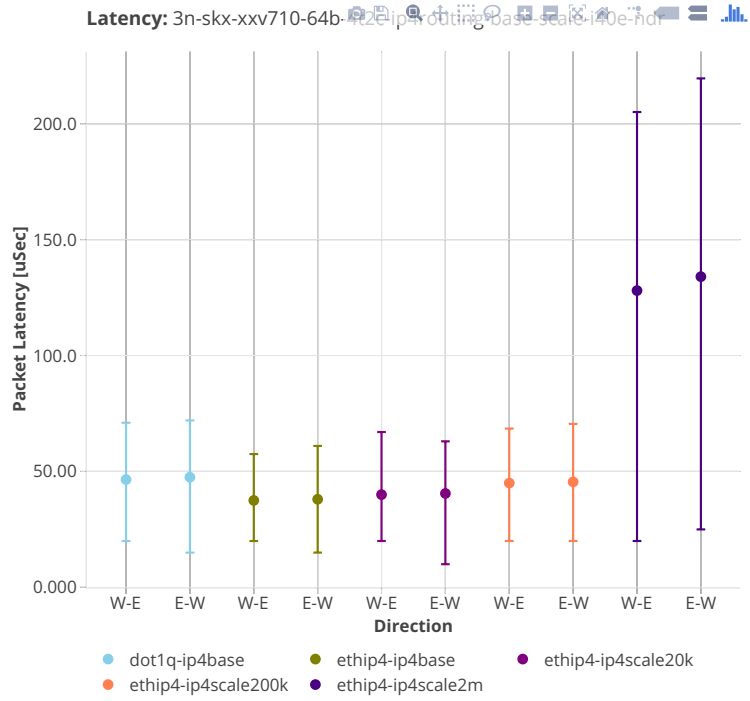
64b-8t4c-ip4routing-base-scale-avf




64b-2t1c-ip4routing-base-scale-i40e




64b-4t2c-ip4routing-base-scale-i40e




64b-8t4c-ip4routing-base-scale-i40e




64b-2t1c-features-ip4routing-base-i40e



64b-4t2c-features-ip4routing-base-i40e



64b-8t4c-features-ip4routing-base-i40e



3n-skx-x710


64b-2t1c-ip4routing-base-scale-avf-i40e

64b-4t2c-ip4routing-base-scale-avf-i40e


64b-8t4c-ip4routing-base-scale-avf-i40e

3n-hsw-xl710


64b-1t1c-ip4routing-base-scale-i40e



64b-2t2c-ip4routing-base-scale-i40e



64b-4t4c-ip4routing-base-scale-i40e



2.5.3 IPv6 Routing

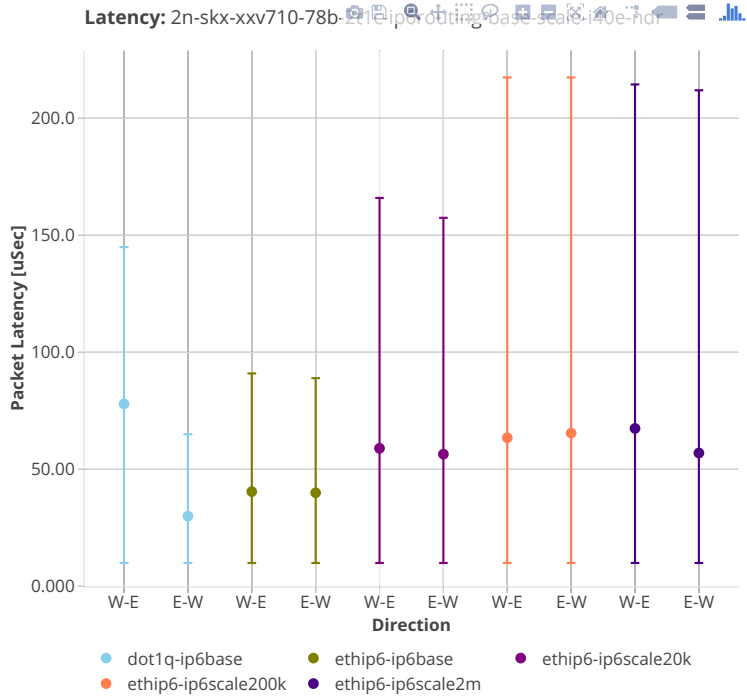
This section includes summary graphs of VPP Phy-to-Phy packet latency with IPv6 Routed-Forwarding measured at 100% of discovered NDR throughput rate. Latency is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

CSIT source code for the test cases used for plots can be found in [CSIT git repository](#)⁷².


⁷² <https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls1908>

2n-skx-xxv710


78b-2t1c-ip6routing-base-scale-i40e



78b-4t2c-ip6routing-base-scale-i40e

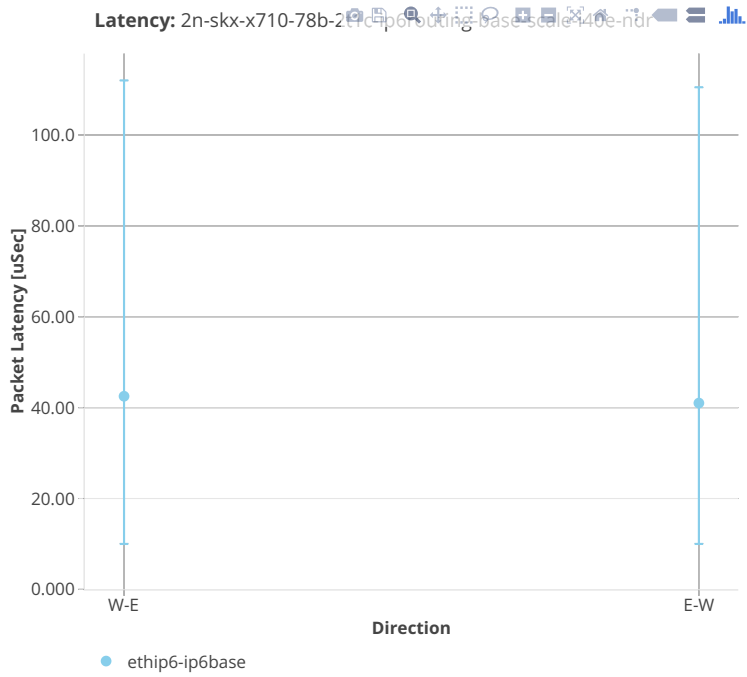


78b-8t4c-ip6routing-base-scale-i40e




2n-skx-x710


78b-2t1c-ip6routing-base-scale-i40e



78b-4t2c-ip6routing-base-scale-i40e

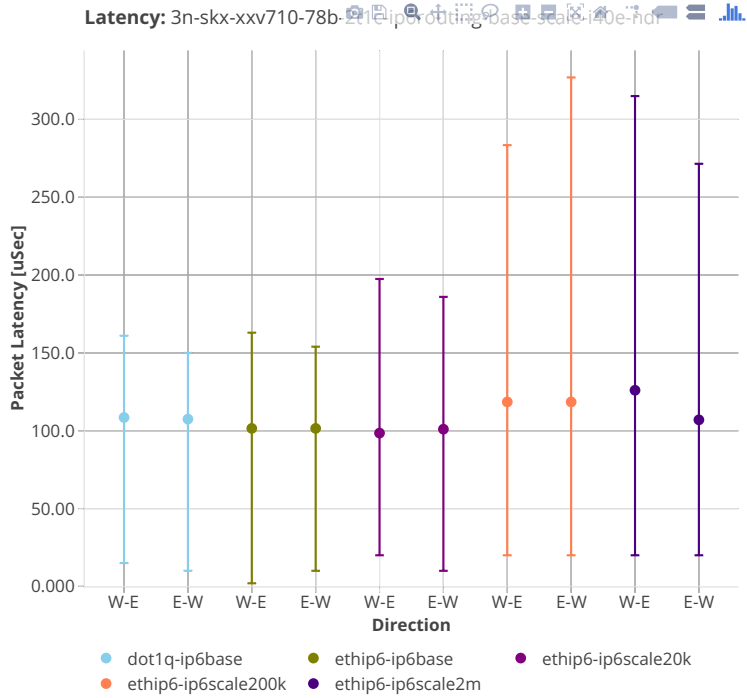


78b-8t4c-ip6routing-base-scale-i40e




3n-skx-xxv710


78b-2t1c-ip6routing-base-scale-i40e



78b-4t2c-ip6routing-base-scale-i40e

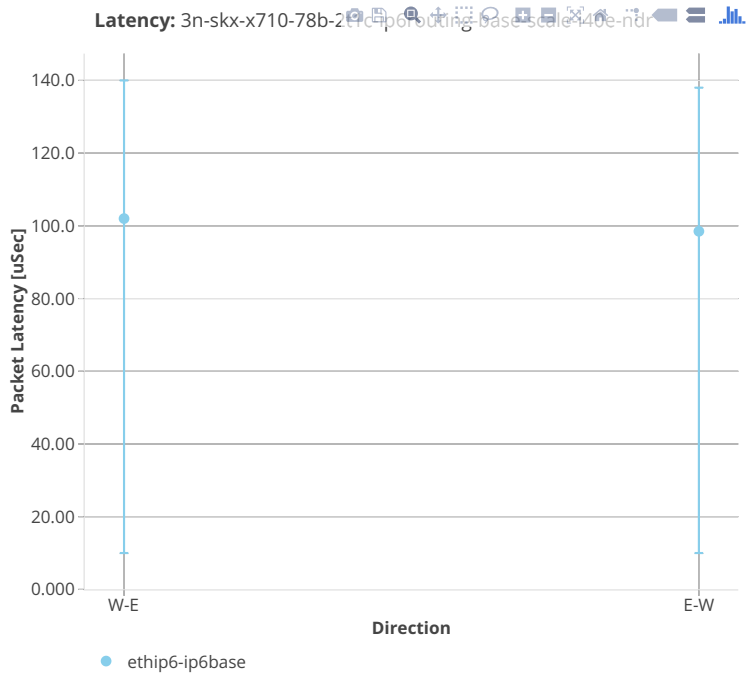


78b-8t4c-ip6routing-base-scale-i40e




3n-skx-x710

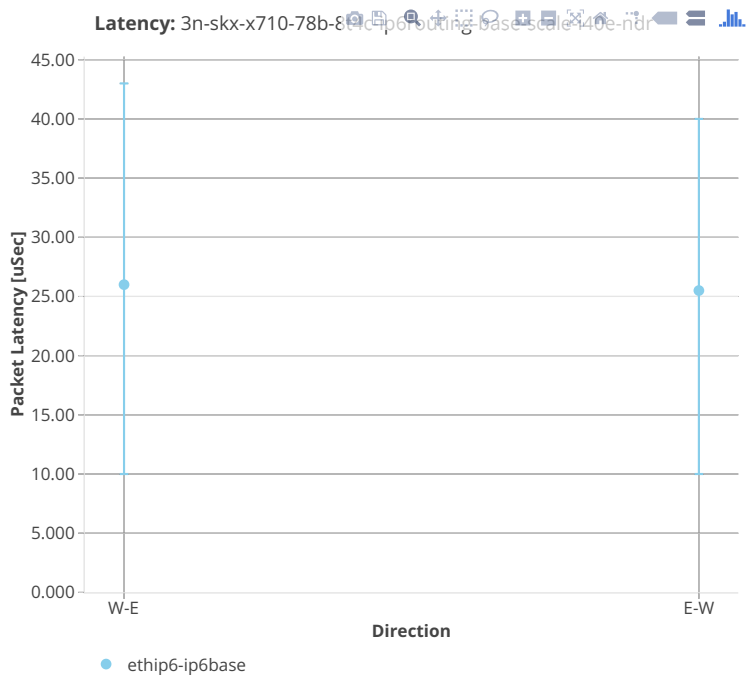
78b-2t1c-ip6routing-base-scale-i40e



78b-4t2c-ip6routing-base-scale-i40e



78b-8t4c-ip6routing-base-scale-i40e



3n-hsw-xl710

78b-1t1c-ip6routing-base-scale-i40e

78b-2t2c-ip6routing-base-scale-i40e

78b-4t4c-ip6routing-base-scale-i40e

2.5.4 IPv4 Tunnels

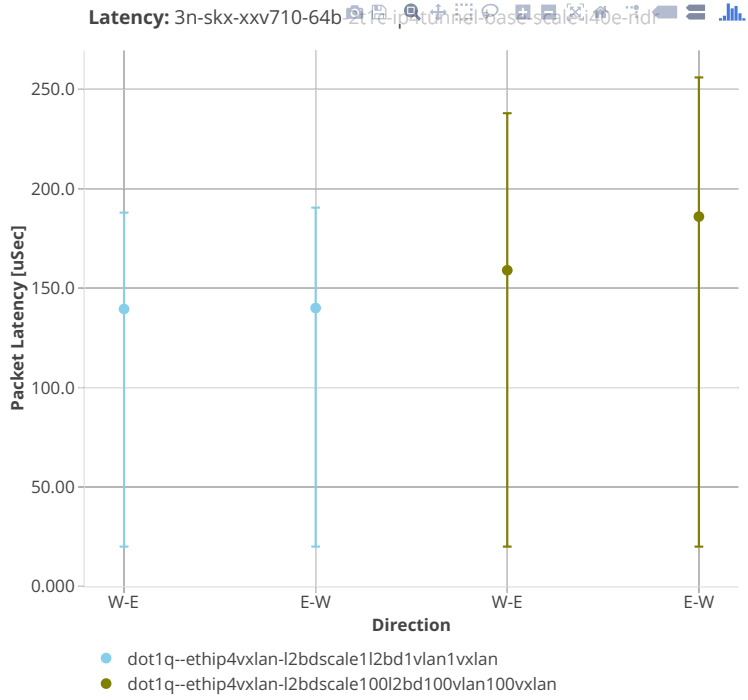
This section includes summary graphs of VPP Phy-to-Phy packet latency with IPv4 Overlay Tunnels measured at 100% of discovered NDR throughput rate. Latency is reported for VPP running in multiple configurations of VPP worker thread(s), a.k.a. VPP data plane thread(s), and their physical CPU core(s) placement.

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


⁷³ https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls1908

3n-skx-xxv710


64b-2t1c-ip4tunnel-base-scale-i40e



64b-4t2c-ip4tunnel-base-scale-i40e




64b-8t4c-ip4tunnel-base-scale-i40e




3n-hsw-xl710


64b-1t1c-ip4tunnel-base-scale-i40e



64b-2t2c-ip4tunnel-base-scale-i40e



64b-4t4c-ip4tunnel-base-scale-i40e



2.5.5 KVM VMs vhost-user


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 100% 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.

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


⁷⁴ https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls1908

2n-skx-xxv710


64b-2t1c-vhost-base-i40e-testpmd




64b-4t2c-vhost-base-i40e-testpmd




64b-8t4c-vhost-base-i40e-testpmd




64b-2t1c-vhost-base-i40e-vpp



64b-4t2c-vhost-base-i40e-vpp




64b-8t4c-vhost-base-i40e-vpp




3n-skx-xxv710


64b-2t1c-vhost-base-i40e-testpmd




64b-4t2c-vhost-base-i40e-testpmd




64b-8t4c-vhost-base-i40e-testpmd




64b-2t1c-vhost-base-i40e-vpp




64b-4t2c-vhost-base-i40e-vpp




64b-8t4c-vhost-base-i40e-vpp



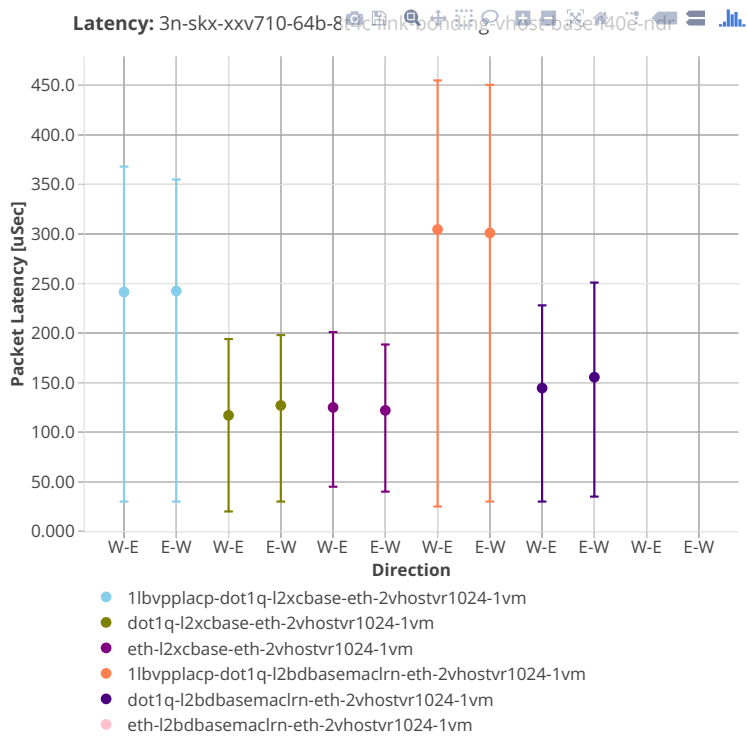
64b-2t1c-link-bonding-vhost-base-i40e-testpmd



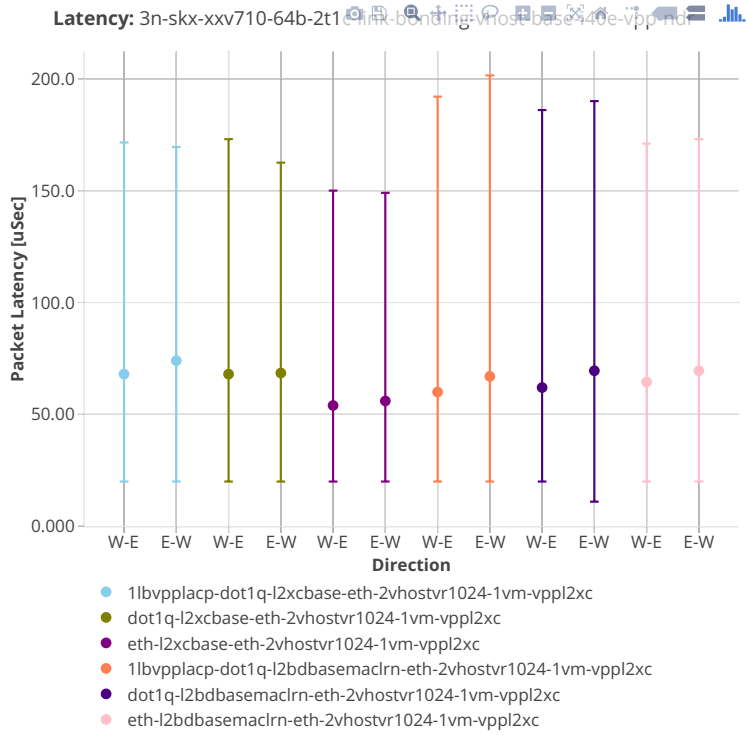
64b-4t2c-link-bonding-vhost-base-i40e-testpmd



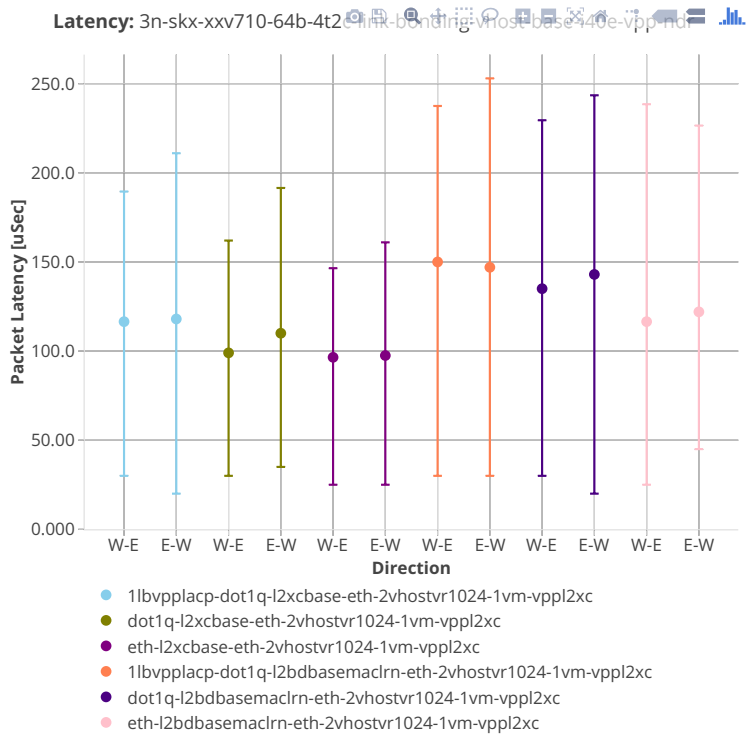
64b-8t4c-link-bonding-vhost-base-i40e-testpmd




64b-2t1c-link-bonding-vhost-base-i40e-vpp



64b-4t2c-link-bonding-vhost-base-i40e-vpp




64b-8t4c-link-bonding-vhost-base-i40e-vpp




3n-skx-x710


64b-2t1c-link-bonding-vhost-base-i40e-testpmd




64b-4t2c-link-bonding-vhost-base-i40e-testpmd




64b-8t4c-link-bonding-vhost-base-i40e-testpmd




64b-2t1c-link-bonding-vhost-base-i40e-vpp



64b-4t2c-link-bonding-vhost-base-i40e-vpp




64b-8t4c-link-bonding-vhost-base-i40e-vpp




3n-hsw-xl710


64b-1t1c-vhost-base-i40e-testpmd



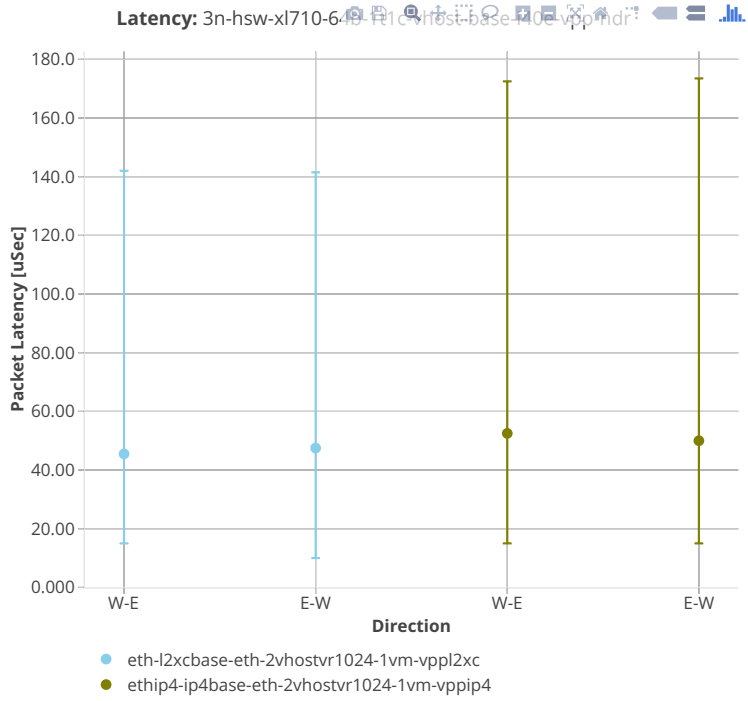
64b-2t2c-vhost-base-i40e-testpmd




64b-4t4c-vhost-base-i40e-testpmd




64b-1t1c-vhost-base-i40e-vpp




64b-2t2c-vhost-base-i40e-vpp




64b-4t4c-vhost-base-i40e-vpp




64b-1t1c-link-bonding-vhost-base-i40e-testpmd




64b-2t2c-link-bonding-vhost-base-i40e-testpmd




64b-4t4c-link-bonding-vhost-base-i40e-testpmd




64b-1t1c-link-bonding-vhost-base-i40e-vpp



64b-2t2c-link-bonding-vhost-base-i40e-vpp



64b-4t4c-link-bonding-vhost-base-i40e-vpp



2.5.6 LXC/DRC Container Memif


This section includes summary graphs of VPP Phy-to-Phy packet latency with Container memif Connections measured at 100% 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.

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


⁷⁵ https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls1908

2n-skx-xxv710


64b-2t1c-memif-base-i40e



64b-4t2c-memif-base-i40e




64b-8t4c-memif-base-i40e




3n-skx-xxv710


64b-2t1c-memif-base-i40e



64b-4t2c-memif-base-i40e



64b-8t4c-memif-base-i40e



2.6 Reconfiguration Tests

See *Reconfiguration Tests* (page 25) for methodology description of this test type.

2.6.1 VNF Service Chains

In each test, a single service chain is added, the re-configuration contains all the steps the initial chains got, except the last step (starting VMs) is skipped.

Additional information about graph data:


1. **Graph Title:** describes tested VPP packet path. Format:
 - wire encapsulation dot1qip4v1xan,
 - VPP forwarding mode 12bd,
 - total number {Y} of service chains {Y}ch,
 - total number of chains being reconfigured 1ach,
 - total number of vhost-user interfaces forwarding packets on VPP with {Y} chains and {X} VMs per chain {2XY}vh (2 interfaces per {X} VMs per {Y} chains),
 - total number {XY} of VNF VMs forwarding packets {XY}vm and finally
 - VNF workload in VM testpmd.
2. **X-axis Labels:** indices of individual test suites as listed in Graph Legend.
3. **Y-axis Labels:** measured Implied time loss [s] values.
4. **Graph Legend:** lists X-axis indices with associated CSIT test suites executed to generate graphed test results and the average value and standard deviation of measured packet loss in millions of packets.
5. **Hover Information:** lists minimum, first quartile, median, third quartile, and maximum. If either type of outlier is present the whisker on the appropriate side is taken to 1.5×IQR from the quartile (the “inner fence”) rather than the max or min, and individual outlying data points are displayed as unfilled circles (for suspected outliers) or filled circles (for outliers). (The “outer fence” is 3×IQR from the quartile.)

Note: Test results have been generated by [FD.io test executor vpp performance job 2n-skx](#)⁷⁶ with RF result files csit-vpp-perf-1908_1-*.zip [archived here](#).


⁷⁶ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-2n-skx

2n-skx-xxv710


imix-2t1c-dot1qip4vxlan-l2bd



imix-4t2c-dot1qip4vxlan-l2bd



imix-8t4c-dot1qip4vxlan-l2bd



2.7 Comparisons

2.7.1 Current vs. Previous Release

Relative comparison of VPP packet throughput (NDR, PDR and MRR) between VPP-19.08.1 release and VPP-19.04 release (measured for CSIT-1908.1 and CSIT-1904 respectively) is calculated from results of tests running on 3-Node Intel Xeon Haswell testbeds (3n-hsw) in 1-core, 2-core and 4-core (MRR only) configurations.

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.

Note: Test results have been generated by

- [FD.io test executor vpp performance job 3n-hsw⁷⁷](#),
- [FD.io test executor vpp performance job 3n-skx⁷⁸](#),
- [FD.io test executor vpp performance job 2n-skx⁷⁹](#)

with RF result files csit-vpp-perf-1908_1-*.zip [archived here](#).

2n-skx

NDR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 2t1c NDR comparison](#)
- [ASCII 4t2c NDR comparison](#)
- [CSV 2t1c NDR comparison](#)
- [CSV 4t2c NDR comparison](#)

PDR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 2t1c PDR comparison](#)
- [ASCII 4t2c PDR comparison](#)
- [CSV 2t1c PDR comparison](#)
- [CSV 4t2c PDR comparison](#)

⁷⁷ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-hsw

⁷⁸ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-skx

⁷⁹ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-2n-skx

MRR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 2t1c MRR comparison](#)
- [ASCII 4t2c MRR comparison](#)
- [ASCII 8t4c MRR comparison](#)
- [CSV 2t1c MRR comparison](#)
- [CSV 4t2c MRR comparison](#)
- [CSV 8t4c MRR comparison](#)

3n-skx

NDR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 2t1c NDR comparison](#)
- [ASCII 4t2c NDR comparison](#)
- [CSV 2t1c NDR comparison](#)
- [CSV 4t2c NDR comparison](#)

PDR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 2t1c PDR comparison](#)
- [ASCII 4t2c PDR comparison](#)
- [CSV 2t1c PDR comparison](#)
- [CSV 4t2c PDR comparison](#)

MRR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 2t1c MRR comparison](#)
- [ASCII 4t2c MRR comparison](#)
- [ASCII 8t4c MRR comparison](#)
- [CSV 2t1c MRR comparison](#)
- [CSV 4t2c MRR comparison](#)
- [CSV 8t4c MRR comparison](#)

3n-hsw

NDR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 1t1c NDR comparison](#)
- [ASCII 2t2c NDR comparison](#)
- [CSV 1t1c NDR comparison](#)
- [CSV 2t2c NDR comparison](#)

PDR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 1t1c PDR comparison](#)
- [ASCII 2t2c PDR comparison](#)
- [CSV 1t1c PDR comparison](#)
- [CSV 2t2c PDR comparison](#)

MRR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 1t1c MRR comparison](#)
- [ASCII 2t2c MRR comparison](#)
- [ASCII 4t4c MRR comparison](#)
- [CSV 1t1c MRR comparison](#)
- [CSV 2t2c MRR comparison](#)
- [CSV 4t4c MRR comparison](#)

2.7.2 3n-Skx vs. 3n-Hsw Testbeds

Relative comparison of VPP-19.08.1 release packet throughput (NDR, PDR and MRR) is calculated for the same tests executed on 3-Node Skylake (3n- skx) and 3-Node Haswell (3n-hsw) physical testbed types, in 1-core, 2-core and 4-core configurations.

Note: Test results have been generated by [FD.io test executor vpp performance job 3n-hsw](#)⁸⁰ and [FD.io test executor vpp performance job 3n-skx](#)⁸¹ with RF result files csit-vpp-perf-1908_1-*.zip [archived here](#).

NDR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 1c NDR comparison](#)
- [ASCII 2c NDR comparison](#)
- [CSV 1c NDR comparison](#)
- [CSV 2c NDR comparison](#)

⁸⁰ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-hsw

⁸¹ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-skx

PDR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 1c PDR comparison](#)
- [ASCII 2c PDR comparison](#)
- [CSV 1c PDR comparison](#)
- [CSV 2c PDR comparison](#)

MRR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 1c MRR comparison](#)
- [ASCII 2c MRR comparison](#)
- [ASCII 4c MRR comparison](#)
- [CSV 1c MRR comparison](#)
- [CSV 2c MRR comparison](#)
- [CSV 4c MRR comparison](#)

2.7.3 3n-Skx vs. 2n-Skx Testbeds

Relative comparison of VPP-19.08.1 release packet throughput (NDR, PDR and MRR) is calculated for the same tests executed on 3-Node Skylake (3n- skx) and 2-Node Skylake (2n-skx) physical testbed types, in 1-core, 2-core and 4-core configurations.

Note: Test results have been generated by [FD.io test executor vpp performance job 3n-skx](#)⁸² and [FD.io test executor vpp performance job 2n-skx](#)⁸³ with RF result files csit-vpp-perf-1908_1-*.zip [archived here](#).

NDR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 1c NDR comparison](#)
- [ASCII 2c NDR comparison](#)
- [CSV 1c NDR comparison](#)
- [CSV 2c NDR comparison](#)

PDR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 1c PDR comparison](#)
- [ASCII 1c PDR comparison](#)
- [CSV 2c PDR comparison](#)
- [CSV 2c PDR comparison](#)

⁸² https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-skx

⁸³ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-2n-skx

MRR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 1c MRR comparison](#)
- [ASCII 2c MRR comparison](#)
- [ASCII 4c MRR comparison](#)
- [CSV 1c MRR comparison](#)
- [CSV 2c MRR comparison](#)
- [CSV 4c MRR comparison](#)

2.7.4 NICs Comparison

Relative comparison of VPP packet throughput (NDR, PDR and MRR) between NICs (measured for |csit-release) is calculated from results of tests running on 3-Node Intel Xeon Haswell testbeds (3n-hsw) in 1-core, 2-core and 4-core (MRR only) configurations.

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.

Note: Test results have been generated by

- [FD.io test executor vpp performance job 3n-skx](#)⁸⁴,
- [FD.io test executor vpp performance job 2n-skx](#)⁸⁵

with RF result files csit-vpp-perf-1908_1-*.zip [archived here](#).

3n-skx

NDR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 2t1c NDR Intel-x710 and Intel-xxv710 comparison](#)
- [ASCII 4t2c NDR Intel-x710 and Intel-xxv710 comparison](#)
- [CSV 2t1c NDR Intel-x710 and Intel-xxv710 comparison](#)
- [CSV 4t2c NDR Intel-x710 and Intel-xxv710 comparison](#)

PDR Comparison

Comparison tables in ASCII and CSV formats:

- [ASCII 2t1c PDR Intel-x710 and Intel-xxv710 comparison](#)
- [ASCII 4t2c PDR Intel-x710 and Intel-xxv710 comparison](#)
- [CSV 2t1c PDR Intel-x710 and Intel-xxv710 comparison](#)

⁸⁴ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-3n-skx

⁸⁵ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-2n-skx

- CSV 4t2c PDR Intel-x710 and Intel-xxv710 comparison

MRR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII 2t1c MRR Intel-x710 and Intel-xxv710 comparison
- ASCII 4t2c MRR Intel-x710 and Intel-xxv710 comparison
- ASCII 8t4c MRR Intel-x710 and Intel-xxv710 comparison
- CSV 2t1c MRR Intel-x710 and Intel-xxv710 comparison
- CSV 4t2c MRR Intel-x710 and Intel-xxv710 comparison
- CSV 8t4c MRR Intel-x710 and Intel-xxv710 comparison

2n-skx

NDR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII 2t1c NDR Intel-x710 and Intel-xxv710 comparison
- ASCII 4t2c NDR Intel-x710 and Intel-xxv710 comparison
- CSV 2t1c NDR Intel-x710 and Intel-xxv710 comparison
- CSV 4t2c NDR Intel-x710 and Intel-xxv710 comparison

PDR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII 2t1c PDR Intel-x710 and Intel-xxv710 comparison
- ASCII 4t2c PDR Intel-x710 and Intel-xxv710 comparison
- CSV 2t1c PDR Intel-x710 and Intel-xxv710 comparison
- CSV 4t2c PDR Intel-x710 and Intel-xxv710 comparison

MRR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII 2t1c MRR Intel-x710 and Intel-xxv710 comparison
- ASCII 4t2c MRR Intel-x710 and Intel-xxv710 comparison
- ASCII 8t4c MRR Intel-x710 and Intel-xxv710 comparison
- CSV 2t1c MRR Intel-x710 and Intel-xxv710 comparison
- CSV 4t2c MRR Intel-x710 and Intel-xxv710 comparison
- CSV 8t4c MRR Intel-x710 and Intel-xxv710 comparison

soak vs ndr comparison:

2.7.5 Soak Tests vs. NDR Tests

Relative comparison of VPP-19.08.1 release Soak PLRSearch vs. NDR packet throughput is calculated for the tests executed on 2-Node Skylake physical testbed types, in 1-core configurations.

Note: Test results have been generated by [FD.io test executor vpp performance job 2n-skx](#)⁸⁶ with RF result files csit-vpp-perf-1908_1-*.zip [archived here](#).

Comparison tables in ASCII and CSV formats:

- [ASCII Soak vs. NDR comparison](#)
- [CSV Soak vs. NDR comparison](#)

2.8 Throughput Trending

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

1. [Performance Dashboard](#)⁸⁷: per VPP test case throughput trend, trend compliance and summary of detected anomalies.
2. [Trending Methodology](#)⁸⁸: throughput test metrics, trend calculations and anomaly classification (progression, regression).
3. [VPP Trendline Graphs](#)⁸⁹: per VPP build MRR throughput measurements against the trendline with anomaly highlights and associated CSIT test jobs.

⁸⁶ https://jenkins.fd.io/view/csit/job/csit-vpp-perf-verify-1908_1-2n-skx

⁸⁷ <https://docs.fd.io/csit/master/trending/introduction/index.html>

⁸⁸ <https://docs.fd.io/csit/master/trending/methodology/index.html>

⁸⁹ <https://docs.fd.io/csit/master/trending/trending/index.html>

2.9 Test Environment

2.9.1 Physical Testbeds

FD.io CSIT performance tests are executed in physical testbeds hosted by LF for FD.io project. Two physical testbed topology types are used:

- **3-Node Topology:** Consisting of two servers acting as SUTs (Systems Under Test) and one server as TG (Traffic Generator), all connected in ring topology.
- **2-Node Topology:** Consisting of one server acting as SUTs and one server as TG both connected in ring topology.

Tested SUT servers are based on a range of processors including Intel Xeon Haswell-SP, Intel Xeon Skylake-SP, Arm, Intel Atom. More detailed description is provided in *Physical Testbeds* (page 4). Tested logical topologies are described in *Logical Topologies* (page 31).

2.9.2 Server Specifications

Complete technical specifications of compute servers used in CSIT physical testbeds are maintained in FD.io CSIT repository: [FD.io CSIT testbeds - Xeon Skylake, Arm, Atom](#)⁹⁰ and [FD.io CSIT Testbeds - Xeon Haswell](#)⁹¹.

2.9.3 Pre-Test Server Calibration

Number of SUT server sub-system runtime parameters have been identified as impacting data plane performance tests. Calibrating those parameters is part of FD.io CSIT pre-test activities, and includes measuring and reporting following:

1. System level core jitter – measure duration of core interrupts by Linux in clock cycles and how often interrupts happen. Using [CPU core jitter tool](#)⁹².
2. Memory bandwidth – measure bandwidth with [Intel MLC tool](#)⁹³.
3. Memory latency – measure memory latency with Intel MLC tool.
4. Cache latency at all levels (L1, L2, and Last Level Cache) – measure cache latency with Intel MLC tool.

Measured values of listed parameters are especially important for repeatable zero packet loss throughput measurements across multiple system instances. Generally they come useful as a background data for comparing data plane performance results across disparate servers.

Following sections include measured calibration data for Intel Xeon Haswell and Intel Xeon Skylake testbeds.

2.9.4 Calibration Data - Haswell

Following sections include sample calibration data measured on t1-sut1 server running in one of the Intel Xeon Haswell testbeds as specified in [FD.io CSIT Testbeds - Xeon Haswell](#)⁹⁴.

Calibration data obtained from all other servers in Haswell testbeds shows the same or similar values.

⁹⁰ https://git.fd.io/csit/tree/docs/lab/testbeds_sm_skx_hw_bios_cfg.md?h=rls1908_1

⁹¹ https://git.fd.io/csit/tree/docs/lab/testbeds_ucs_hsw_hw_bios_cfg.md?h=rls1908_1

⁹² https://git.fd.io/pma_tools/tree/jitter

⁹³ <https://software.intel.com/en-us/articles/intelr-memory-latency-checker>

⁹⁴ https://git.fd.io/csit/tree/docs/lab/testbeds_ucs_hsw_hw_bios_cfg.md?h=rls1908_1

Linux cmdline

```
$ cat /proc/cmdline
BOOT_IMAGE=/vmlinuz-4.15.0-36-generic root=UUID=5d2ecc97-245b-4e94-b0ae-c3548567de19 ro isolcpus=1-
↪17,19-35 nohz_full=1-17,19-35 rcu_nocbs=1-17,19-35 numa_balancing=disable intel_pstate=disable_
↪intel_iommu=on iommu=pt nmi_watchdog=0 audit=0 nosoftlockup processor.max_cstate=1 intel_idle.max_
↪cstate=1 hpet=disable tsc=reliable mce=off console=tty0 console=ttyS0,115200n8
```

Linux uname

```
$ uname -a
Linux t1-tg1 4.15.0-36-generic #39-Ubuntu SMP Mon Sep 24 16:19:09 UTC 2018 x86_64 x86_64 x86_64 GNU/
↪Linux
```

System-level Core Jitter

```
$ sudo taskset -c 3 /home/testuser/pma_tools/jitter/jitter -i 30
Linux Jitter testing program version 1.8
Iterations=30
The program will execute a dummy function 80000 times
Display is updated every 20000 displayUpdate intervals
Timings are in CPU Core cycles
Inst_Min:   Minimum Execution time during the display update interval(default is ~1 second)
Inst_Max:   Maximum Execution time during the display update interval(default is ~1 second)
Inst_jitter: Jitter in the Execution time during the display update interval. This is the value of _
↪interest
last_Exec:  The Execution time of last iteration just before the display update
Abs_Min:   Absolute Minimum Execution time since the program started or statistics were reset
Abs_Max:   Absolute Maximum Execution time since the program started or statistics were reset
tmp:       Cumulative value calculated by the dummy function
Interval:  Time interval between the display updates in Core Cycles
Sample No: Sample number
```

Inst_Min	Inst_Max	Inst_jitter	last_Exec	Abs_min	Abs_max	tmp	Interval	
↪Sample No								
160024	172636	12612	160028	160024	172636	1573060608	3205463144	↪
↪1								
160024	188236	28212	160028	160024	188236	958595072	3205500844	↪
↪2								
160024	185676	25652	160028	160024	188236	344129536	3205485976	↪
↪3								
160024	172608	12584	160024	160024	188236	4024631296	3205472740	↪
↪4								
160024	179260	19236	160028	160024	188236	3410165760	3205502164	↪
↪5								
160024	172432	12408	160024	160024	188236	2795700224	3205452036	↪
↪6								
160024	178820	18796	160024	160024	188236	2181234688	3205455408	↪
↪7								
160024	172512	12488	160028	160024	188236	1566769152	3205461528	↪
↪8								
160024	172636	12612	160028	160024	188236	952303616	3205478820	↪
↪9								
160024	173676	13652	160028	160024	188236	337838080	3205470412	↪
↪10								
160024	178776	18752	160028	160024	188236	4018339840	3205481472	↪
↪11								
160024	172788	12764	160028	160024	188236	3403874304	3205492336	↪
↪12								

(continues on next page)

(continued from previous page)

	160024	174616	14592	160028	160024	188236	2789408768	3205474904	┌
↔13									
	160024	174440	14416	160028	160024	188236	2174943232	3205479448	┌
↔14									
	160024	178748	18724	160024	160024	188236	1560477696	3205482668	┌
↔15									
	160024	172588	12564	169404	160024	188236	946012160	3205510496	┌
↔16									
	160024	172636	12612	160024	160024	188236	331546624	3205472204	┌
↔17									
	160024	172480	12456	160024	160024	188236	4012048384	3205455864	┌
↔18									
	160024	172740	12716	160028	160024	188236	3397582848	3205464932	┌
↔19									
	160024	179200	19176	160028	160024	188236	2783117312	3205476012	┌
↔20									
	160024	172480	12456	160028	160024	188236	2168651776	3205465632	┌
↔21									
	160024	172728	12704	160024	160024	188236	1554186240	3205497204	┌
↔22									
	160024	172620	12596	160028	160024	188236	939720704	3205466972	┌
↔23									
	160024	172640	12616	160028	160024	188236	325255168	3205471216	┌
↔24									
	160024	172484	12460	160028	160024	188236	4005756928	3205467388	┌
↔25									
	160024	172636	12612	160028	160024	188236	3391291392	3205482748	┌
↔26									
	160024	179056	19032	160024	160024	188236	2776825856	3205467152	┌
↔27									
	160024	172672	12648	160024	160024	188236	2162360320	3205483268	┌
↔28									
	160024	176932	16908	160024	160024	188236	1547894784	3205488536	┌
↔29									
	160024	172452	12428	160028	160024	188236	933429248	3205440636	┌
↔30									

Memory Bandwidth

```
$ sudo /home/testuser/mlc --bandwidth_matrix
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --bandwidth_matrix

Using buffer size of 100.000MB/thread for reads and an additional 100.000MB/thread for writes
Measuring Memory Bandwidths between nodes within system
Bandwidths are in MB/sec (1 MB/sec = 1,000,000 Bytes/sec)
Using all the threads from each core if Hyper-threading is enabled
Using Read-only traffic type
      Numa node
Numa node    0    1
0           57935.5  30265.2
1           30284.6  58409.9
```

```
$ sudo /home/testuser/mlc --peak_injection_bandwidth
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --peak_injection_bandwidth

Using buffer size of 100.000MB/thread for reads and an additional 100.000MB/thread for writes
```

(continues on next page)

(continued from previous page)

```

Measuring Peak Injection Memory Bandwidths for the system
Bandwidths are in MB/sec (1 MB/sec = 1,000,000 Bytes/sec)
Using all the threads from each core if Hyper-threading is enabled
Using traffic with the following read-write ratios
ALL Reads      : 115762.2
3:1 Reads-Writes : 106242.2
2:1 Reads-Writes : 103031.8
1:1 Reads-Writes : 87943.7
Stream-triad like: 100048.4

```

```

$ sudo /home/testuser/mlc --max_bandwidth
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --max_bandwidth

Using buffer size of 100.000MB/thread for reads and an additional 100.000MB/thread for writes

Measuring Maximum Memory Bandwidths for the system
Will take several minutes to complete as multiple injection rates will be tried to get the best_
↔_bandwidth
Bandwidths are in MB/sec (1 MB/sec = 1,000,000 Bytes/sec)
Using all the threads from each core if Hyper-threading is enabled
Using traffic with the following read-write ratios
ALL Reads      : 115782.41
3:1 Reads-Writes : 105965.78
2:1 Reads-Writes : 103162.38
1:1 Reads-Writes : 88255.82
Stream-triad like: 105608.10

```

Memory Latency

```

$ sudo /home/testuser/mlc --latency_matrix
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --latency_matrix

Using buffer size of 200.000MB
Measuring idle latencies (in ns)...

```

	Numa node	
Numa node	0	1
0	101.0	132.0
1	141.2	98.8

```

$ sudo /home/testuser/mlc --idle_latency
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --idle_latency

Using buffer size of 200.000MB
Each iteration took 227.2 core clocks ( 99.0 ns)

```

```

$ sudo /home/testuser/mlc --loaded_latency
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --loaded_latency

Using buffer size of 100.000MB/thread for reads and an additional 100.000MB/thread for writes

Measuring Loaded Latencies for the system
Using all the threads from each core if Hyper-threading is enabled
Using Read-only traffic type
Inject Latency Bandwidth

```

(continues on next page)

(continued from previous page)

Delay	(ns)	MB/sec
00000	294.08	115841.6
00002	294.27	115851.5
00008	293.67	115821.8
00015	278.92	115587.5
00050	246.80	113991.2
00100	206.86	104508.1
00200	123.72	72873.6
00300	113.35	52641.1
00400	108.89	41078.9
00500	108.11	33699.1
00700	106.19	24878.0
01000	104.75	17948.1
01300	103.72	14089.0
01700	102.95	11013.6
02500	102.25	7756.3
03500	101.81	5749.3
05000	101.46	4230.4
09000	101.05	2641.4
20000	100.77	1542.5

L1/L2/LLC Latency

```
$ sudo /home/testuser/mlc --c2c_latency
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --c2c_latency

Measuring cache-to-cache transfer latency (in ns)...
Local Socket L2->L2 HIT latency      42.1
Local Socket L2->L2 HITM latency     47.0
Remote Socket L2->L2 HITM latency (data address homed in writer socket)
      Reader Numa Node
Writer Numa Node   0      1
                  0      - 108.0
                  1    106.9   -
Remote Socket L2->L2 HITM latency (data address homed in reader socket)
      Reader Numa Node
Writer Numa Node   0      1
                  0      - 107.7
                  1    106.6   -
```

Spectre and Meltdown Checks

Following section displays the output of a running shell script to tell if system is vulnerable against the several “speculative execution” CVEs that were made public in 2018. Script is available on [Spectre & Meltdown Checker Github](#)⁹⁵.

```
Spectre and Meltdown mitigation detection tool v0.42

Checking for vulnerabilities on current system
Kernel is Linux 4.15.0-36-generic #39-Ubuntu SMP Mon Sep 24 16:19:09 UTC 2018 x86_64
CPU is Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz

Hardware check
* Hardware support (CPU microcode) for mitigation techniques
```

(continues on next page)

⁹⁵ <https://github.com/speed47/spectre-meltdown-checker>

(continued from previous page)

```

* Indirect Branch Restricted Speculation (IBRS)
  * SPEC_CTRL MSR is available: YES
  * CPU indicates IBRS capability: YES (SPEC_CTRL feature bit)
* Indirect Branch Prediction Barrier (IBPB)
  * PRED_CMD MSR is available: YES
  * CPU indicates IBPB capability: YES (SPEC_CTRL feature bit)
* Single Thread Indirect Branch Predictors (STIBP)
  * SPEC_CTRL MSR is available: YES
  * CPU indicates STIBP capability: YES (Intel STIBP feature bit)
* Speculative Store Bypass Disable (SSBD)
  * CPU indicates SSBD capability: YES (Intel SSBD)
* L1 data cache invalidation
  * FLUSH_CMD MSR is available: YES
  * CPU indicates L1D flush capability: YES (L1D flush feature bit)
* Microarchitecture Data Sampling
  * VERW instruction is available: NO
* Enhanced IBRS (IBRS_ALL)
  * CPU indicates ARCH_CAPABILITIES MSR availability: NO
  * ARCH_CAPABILITIES MSR advertises IBRS_ALL capability: NO
* CPU explicitly indicates not being vulnerable to Meltdown/L1TF (RDCL_NO): NO
* CPU explicitly indicates not being vulnerable to Variant 4 (SSB_NO): NO
* CPU/Hypervisor indicates L1D flushing is not necessary on this system: NO
* Hypervisor indicates host CPU might be vulnerable to RSB underflow (RSBA): NO
* CPU explicitly indicates not being vulnerable to Microarchitectural Data Sampling (MDS_NO): NO
* CPU supports Software Guard Extensions (SGX): NO
* CPU microcode is known to cause stability problems: NO (model 0x3f family 0x6 stepping 0x2_
↳ ucode 0x3d cpuid 0x306f2)
  * CPU microcode is the latest known available version: awk: cannot open bash (No such file or_
↳ directory)
UNKNOWN (latest microcode version for your CPU model is unknown)
* CPU vulnerability to the speculative execution attack variants
  * Vulnerable to CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Vulnerable to CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Vulnerable to CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): YES
  * Vulnerable to CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Vulnerable to CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Vulnerable to CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Vulnerable to CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
  * Vulnerable to CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
  * Vulnerable to CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)):_
↳ YES
  * Vulnerable to CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling_
↳ (MFBDS)): YES
  * Vulnerable to CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): YES
  * Vulnerable to CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory_
↳ (MDSUM)): YES

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: __user pointer sanitization)
* Kernel has array_index_mask_nospec: YES (1 occurrence(s) found of x86 64 bits array_index_mask_
↳ nospec())
* Kernel has the Red Hat/Ubuntu patch: NO
* Kernel has mask_nospec64 (arm64): NO
> STATUS: NOT VULNERABLE (Mitigation: __user pointer sanitization)

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: YES (Mitigation: Full generic retpoline, IBPB, IBRS_FW)
* Mitigation 1
  * Kernel is compiled with IBRS support: YES
    * IBRS enabled and active: YES (for firmware code only)
  * Kernel is compiled with IBPB support: YES

```

(continues on next page)

(continued from previous page)

```

* IBPB enabled and active: YES
* Mitigation 2
* Kernel has branch predictor hardening (arm): NO
* Kernel compiled with retpoline option: YES
  * Kernel compiled with a retpoline-aware compiler: YES (kernel reports full retpoline_
↳ compilation)
> STATUS: NOT VULNERABLE (Full retpoline + IBPB are mitigating the vulnerability)

CVE-2017-5754 aka Variant 3, Meltdown, rogue data cache load
* Mitigated according to the /sys interface: YES (Mitigation: PTI)
* Kernel supports Page Table Isolation (PTI): YES
  * PTI enabled and active: YES
  * Reduced performance impact of PTI: YES (CPU supports INVPCID, performance impact of PTI will be_
↳ greatly reduced)
* Running as a Xen PV DomU: NO
> STATUS: NOT VULNERABLE (Mitigation: PTI)

CVE-2018-3640 aka Variant 3a, rogue system register read
* CPU microcode mitigates the vulnerability: YES
> STATUS: NOT VULNERABLE (your CPU microcode mitigates the vulnerability)

CVE-2018-3639 aka Variant 4, speculative store bypass
* Mitigated according to the /sys interface: YES (Mitigation: Speculative Store Bypass disabled via_
↳ prctl and seccomp)
* Kernel supports disabling speculative store bypass (SSB): YES (found in /proc/self/status)
* SSB mitigation is enabled and active: YES (per-thread through prctl)
* SSB mitigation currently active for selected processes: YES (systemd-journald systemd-logind_
↳ systemd-networkd systemd-resolved systemd-timesyncd systemd-udev)
> STATUS: NOT VULNERABLE (Mitigation: Speculative Store Bypass disabled via prctl and seccomp)

CVE-2018-3615 aka Foreshadow (SGX), L1 terminal fault
* CPU microcode mitigates the vulnerability: N/A
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-3620 aka Foreshadow-NG (OS), L1 terminal fault
* Mitigated according to the /sys interface: YES (Mitigation: PTE Inversion; VMX: conditional cache_
↳ flushes, SMT disabled)
* Kernel supports PTE inversion: YES (found in kernel image)
* PTE inversion enabled and active: YES
> STATUS: NOT VULNERABLE (Mitigation: PTE Inversion; VMX: conditional cache flushes, SMT disabled)

CVE-2018-3646 aka Foreshadow-NG (VMM), L1 terminal fault
* Information from the /sys interface: Mitigation: PTE Inversion; VMX: conditional cache flushes,_
↳ SMT disabled
* This system is a host running a hypervisor: NO
* Mitigation 1 (KVM)
  * EPT is disabled: NO
* Mitigation 2
  * L1D flush is supported by kernel: YES (found flush_l1d in /proc/cpuinfo)
  * L1D flush enabled: YES (conditional flushes)
  * Hardware-backed L1D flush supported: YES (performance impact of the mitigation will be greatly_
↳ reduced)
  * Hyper-Threading (SMT) is enabled: NO
> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2018-12126 aka Fallout, microarchitectural store buffer data sampling (MSBDS)
* Kernel supports using MD_CLEAR mitigation: NO
> STATUS: VULNERABLE (Neither your kernel or your microcode support mitigation, upgrade both to_
↳ mitigate the vulnerability)

CVE-2018-12130 aka ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)

```

(continues on next page)

(continued from previous page)

```

* Kernel supports using MD_CLEAR mitigation: NO
> STATUS: VULNERABLE (Neither your kernel or your microcode support mitigation, upgrade both to_
↳mitigate the vulnerability)

CVE-2018-12127 aka RIDL, microarchitectural load port data sampling (MLPDS)
* Kernel supports using MD_CLEAR mitigation: NO
> STATUS: VULNERABLE (Neither your kernel or your microcode support mitigation, upgrade both to_
↳mitigate the vulnerability)

CVE-2019-11091 aka RIDL, microarchitectural data sampling uncacheable memory (MDSUM)
* Kernel supports using MD_CLEAR mitigation: NO
> STATUS: VULNERABLE (Neither your kernel or your microcode support mitigation, upgrade both to_
↳mitigate the vulnerability)

> SUMMARY: CVE-2017-5753:OK CVE-2017-5715:OK CVE-2017-5754:OK CVE-2018-3640:OK CVE-2018-3639:OK CVE-
↳2018-3615:OK CVE-2018-3620:OK CVE-2018-3646:OK CVE-2018-12126:KO CVE-2018-12130:KO CVE-2018-
↳12127:KO CVE-2019-11091:KO

```

2.9.5 Calibration Data - Skylake

Following sections include sample calibration data measured on s11-t31-sut1 server running in one of the Intel Xeon Skylake testbeds as specified in [FD.io CSIT testbeds - Xeon Skylake, Arm, Atom](#)⁹⁶.

Calibration data obtained from all other servers in Skylake testbeds shows the same or similar values.

Linux cmdline

```

$ cat /proc/cmdline
BOOT_IMAGE=/vmlinuz-4.15.0-23-generic root=UUID=759ad671-ad46-441b-a75b-9f54e81837bb ro isolcpus=1-
↳27,29-55,57-83,85-111 nohz_full=1-27,29-55,57-83,85-111 rcu_nocbs=1-27,29-55,57-83,85-111 numa_
↳balancing=disable intel_pstate=disable intel_iommu=on iommu=pt nmi_watchdog=0 audit=0_
↳nosoftlockup processor.max_cstate=1 intel_idle.max_cstate=1 hpet=disable tsc=reliable mce=off_
↳console=tty0 console=ttyS0,115200n8

```

Linux uname

```

$ uname -a
Linux s5-t22-sut1 4.15.0-23-generic #25-Ubuntu SMP Wed May 23 18:02:16 UTC 2018 x86_64 x86_64 x86_
↳64 GNU/Linux

```

System-level Core Jitter

```

$ sudo taskset -c 3 /home/testuser/pma_tools/jitter/jitter -i 20
Linux Jitter testing program version 1.8
Iterations=20
The program will execute a dummy function 80000 times
Display is updated every 20000 displayUpdate intervals
Timings are in CPU Core cycles
Inst_Min:   Minimum Execution time during the display update interval(default is ~1 second)
Inst_Max:   Maximum Execution time during the display update interval(default is ~1 second)
Inst_jitter: Jitter in the Execution time during the display update interval. This is the value of_
↳interest
last_Exec:  The Execution time of last iteration just before the display update

```

(continues on next page)

⁹⁶ https://git.fd.io/csit/tree/docs/lab/testbeds_sm_skx_hw_bios_cfg.md?h=rls1908_1

(continued from previous page)

Abs_Min:	Absolute Minimum Execution time since the program started or statistics were reset								
Abs_Max:	Absolute Maximum Execution time since the program started or statistics were reset								
tmp:	Cumulative value calculated by the dummy function								
Interval:	Time interval between the display updates in Core Cycles								
Sample No:	Sample number								
	Inst_Min	Inst_Max	Inst_jitter	last_Exec	Abs_min	Abs_max	tmp	Interval	
↔Sample No									
↔1	160022	171330	11308	160022	160022	171330	2538733568	3204142750	
↔2	160022	167294	7272	160026	160022	171330	328335360	3203873548	
↔3	160022	167560	7538	160026	160022	171330	2412904448	3203878736	
↔4	160022	169000	8978	160024	160022	171330	202506240	3203864588	
↔5	160022	166572	6550	160026	160022	171330	2287075328	3203866224	
↔6	160022	167460	7438	160026	160022	171330	76677120	3203854632	
↔7	160022	168134	8112	160024	160022	171330	2161246208	3203874674	
↔8	160022	169094	9072	160022	160022	171330	4245815296	3203878798	
↔9	160022	172460	12438	160024	160022	172460	2035417088	3204112010	
↔10	160022	167862	7840	160030	160022	172460	4119986176	3203856800	
↔11	160022	168398	8376	160024	160022	172460	1909587968	3203854192	
↔12	160022	167548	7526	160024	160022	172460	3994157056	3203847442	
↔13	160022	167562	7540	160026	160022	172460	1783758848	3203862936	
↔14	160022	167604	7582	160024	160022	172460	3868327936	3203859346	
↔15	160022	168262	8240	160024	160022	172460	1657929728	3203851120	
↔16	160022	169700	9678	160024	160022	172460	3742498816	3203877690	
↔17	160022	170476	10454	160026	160022	172460	1532100608	3204088480	
↔18	160022	167798	7776	160024	160022	172460	3616669696	3203862072	
↔19	160022	166540	6518	160024	160022	172460	1406271488	3203836904	
↔20	160022	167516	7494	160024	160022	172460	3490840576	3203848120	

Memory Bandwidth

```

$ sudo /home/testuser/mlc --bandwidth_matrix
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --bandwidth_matrix

Using buffer size of 100.000MB/thread for reads and an additional 100.000MB/thread for writes
Measuring Memory Bandwidths between nodes within system
Bandwidths are in MB/sec (1 MB/sec = 1,000,000 Bytes/sec)
Using all the threads from each core if Hyper-threading is enabled
    
```

(continues on next page)

(continued from previous page)

```
Using Read-only traffic type
```

		Numa node	
Numa node	0	1	
0	107947.7	50951.5	
1	50834.6	108183.4	

```
$ sudo /home/testuser/mlc --peak_injection_bandwidth
```

```
Intel(R) Memory Latency Checker - v3.5
```

```
Command line parameters: --peak_injection_bandwidth
```

```
Using buffer size of 100.000MB/thread for reads and an additional 100.000MB/thread for writes
```

```
Measuring Peak Injection Memory Bandwidths for the system
```

```
Bandwidths are in MB/sec (1 MB/sec = 1,000,000 Bytes/sec)
```

```
Using all the threads from each core if Hyper-threading is enabled
```

```
Using traffic with the following read-write ratios
```

```
ALL Reads      : 215733.9
```

```
3:1 Reads-Writes : 182141.9
```

```
2:1 Reads-Writes : 178615.7
```

```
1:1 Reads-Writes : 149911.3
```

```
Stream-triad like: 159533.6
```

```
$ sudo /home/testuser/mlc --max_bandwidth
```

```
Intel(R) Memory Latency Checker - v3.5
```

```
Command line parameters: --max_bandwidth
```

```
Using buffer size of 100.000MB/thread for reads and an additional 100.000MB/thread for writes
```

```
Measuring Maximum Memory Bandwidths for the system
```

```
Will take several minutes to complete as multiple injection rates will be tried to get the best_
↪bandwidth
```

```
Bandwidths are in MB/sec (1 MB/sec = 1,000,000 Bytes/sec)
```

```
Using all the threads from each core if Hyper-threading is enabled
```

```
Using traffic with the following read-write ratios
```

```
ALL Reads      : 216875.73
```

```
3:1 Reads-Writes : 182615.14
```

```
2:1 Reads-Writes : 178745.67
```

```
1:1 Reads-Writes : 149485.27
```

```
Stream-triad like: 180057.87
```

Memory Latency

```
$ sudo /home/testuser/mlc --latency_matrix
```

```
Intel(R) Memory Latency Checker - v3.5
```

```
Command line parameters: --latency_matrix
```

```
Using buffer size of 2000.000MB
```

```
Measuring idle latencies (in ns)...
```

		Numa node	
Numa node	0	1	
0	81.4	131.1	
1	131.1	81.3	

```
$ sudo /home/testuser/mlc --idle_latency
```

```
Intel(R) Memory Latency Checker - v3.5
```

```
Command line parameters: --idle_latency
```

```
Using buffer size of 2000.000MB
```

```
Each iteration took 202.0 core clocks ( 80.8 ns)
```



```

$ sudo /home/testuser/mlc --loaded_latency
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --loaded_latency

Using buffer size of 100.000MB/thread for reads and an additional 100.000MB/thread for writes

Measuring Loaded Latencies for the system
Using all the threads from each core if Hyper-threading is enabled
Using Read-only traffic type
Inject Latency Bandwidth
Delay (ns) MB/sec
=====
00000 282.66 215712.8
00002 282.14 215757.4
00008 280.21 215868.1
00015 279.20 216313.2
00050 275.25 216643.0
00100 227.05 215075.0
00200 121.92 160242.9
00300 101.21 111587.4
00400 95.48 85019.7
00500 94.46 68717.3
00700 92.27 49742.2
01000 91.03 35264.8
01300 90.11 27396.3
01700 89.34 21178.7
02500 90.15 14672.8
03500 89.00 10715.7
05000 82.00 7788.2
09000 81.46 4684.0
20000 81.40 2541.9

```

L1/L2/LLC Latency

```

$ sudo /home/testuser/mlc --c2c_latency
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --c2c_latency

Measuring cache-to-cache transfer latency (in ns)...
Local Socket L2->L2 HIT latency 53.7
Local Socket L2->L2 HITM latency 53.7
Remote Socket L2->L2 HITM latency (data address homed in writer socket)
      Reader Numa Node
Writer Numa Node 0 1
0 - 113.9
1 113.9 -
Remote Socket L2->L2 HITM latency (data address homed in reader socket)
      Reader Numa Node
Writer Numa Node 0 1
0 - 177.9
1 177.6 -

```

Spectre and Meltdown Checks

Following section displays the output of a running shell script to tell if system is vulnerable against the several “speculative execution” CVEs that were made public in 2018. Script is available on [Spectre & Meltdown Checker Github](#)⁹⁷.

⁹⁷ <https://github.com/speed47/spectre-meltdown-checker>

```

Spectre and Meltdown mitigation detection tool v0.42

Checking for vulnerabilities on current system
Kernel is Linux 4.15.0-46-generic #49-Ubuntu SMP Wed Feb 6 09:33:07 UTC 2019 x86_64
CPU is Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50GHz

Hardware check
* Hardware support (CPU microcode) for mitigation techniques
  * Indirect Branch Restricted Speculation (IBRS)
    * SPEC_CTRL MSR is available: YES
    * CPU indicates IBRS capability: YES (SPEC_CTRL feature bit)
  * Indirect Branch Prediction Barrier (IBPB)
    * PRED_CMD MSR is available: YES
    * CPU indicates IBPB capability: YES (SPEC_CTRL feature bit)
  * Single Thread Indirect Branch Predictors (STIBP)
    * SPEC_CTRL MSR is available: YES
    * CPU indicates STIBP capability: YES (Intel STIBP feature bit)
  * Speculative Store Bypass Disable (SSBD)
    * CPU indicates SSBD capability: YES (Intel SSBD)
  * L1 data cache invalidation
    * FLUSH_CMD MSR is available: YES
    * CPU indicates L1D flush capability: YES (L1D flush feature bit)
  * Microarchitecture Data Sampling
    * VERW instruction is available: NO
  * Enhanced IBRS (IBRS_ALL)
    * CPU indicates ARCH_CAPABILITIES MSR availability: NO
    * ARCH_CAPABILITIES MSR advertises IBRS_ALL capability: NO
  * CPU explicitly indicates not being vulnerable to Meltdown/L1TF (RDCL_NO): NO
  * CPU explicitly indicates not being vulnerable to Variant 4 (SSB_NO): NO
  * CPU/Hypervisor indicates L1D flushing is not necessary on this system: NO
  * Hypervisor indicates host CPU might be vulnerable to RSB underflow (RSBA): NO
  * CPU explicitly indicates not being vulnerable to Microarchitectural Data Sampling (MDS_NO): NO
  * CPU supports Software Guard Extensions (SGX): NO
  * CPU microcode is known to cause stability problems: NO (model 0x55 family 0x6 stepping 0x4_
↳ ucode 0x200004d cpuid 0x50654)
  * CPU microcode is the latest known available version: awk: cannot open bash (No such file or_
↳ directory)
UNKNOWN (latest microcode version for your CPU model is unknown)
* CPU vulnerability to the speculative execution attack variants
  * Vulnerable to CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Vulnerable to CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Vulnerable to CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): YES
  * Vulnerable to CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Vulnerable to CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Vulnerable to CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Vulnerable to CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
  * Vulnerable to CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
  * Vulnerable to CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)):_
↳ YES
  * Vulnerable to CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling_
↳ (MFBDS)): YES
  * Vulnerable to CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): YES
  * Vulnerable to CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory_
↳ (MDSUM)): YES

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: __user pointer sanitization)
* Kernel has array_index_mask_nospec: YES (1 occurrence(s) found of x86 64 bits array_index_mask_
↳ nospec())
* Kernel has the Red Hat/Ubuntu patch: NO
* Kernel has mask_nospec64 (arm64): NO
> STATUS: NOT VULNERABLE (Mitigation: __user pointer sanitization)

```

(continues on next page)

(continued from previous page)

```

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: YES (Mitigation: Full generic retpoline, IBPB, IBRS_FW)
* Mitigation 1
  * Kernel is compiled with IBRS support: YES
    * IBRS enabled and active: YES (for firmware code only)
  * Kernel is compiled with IBPB support: YES
    * IBPB enabled and active: YES
* Mitigation 2
  * Kernel has branch predictor hardening (arm): NO
  * Kernel compiled with retpoline option: YES
    * Kernel compiled with a retpoline-aware compiler: YES (kernel reports full retpoline_
↳ compilation)
  * Kernel supports RSB filling: YES
> STATUS: NOT VULNERABLE (Full retpoline + IBPB are mitigating the vulnerability)

CVE-2017-5754 aka Variant 3, Meltdown, rogue data cache load
* Mitigated according to the /sys interface: YES (Mitigation: PTI)
* Kernel supports Page Table Isolation (PTI): YES
  * PTI enabled and active: YES
  * Reduced performance impact of PTI: YES (CPU supports INVPCID, performance impact of PTI will be_
↳ greatly reduced)
* Running as a Xen PV DomU: NO
> STATUS: NOT VULNERABLE (Mitigation: PTI)

CVE-2018-3640 aka Variant 3a, rogue system register read
* CPU microcode mitigates the vulnerability: YES
> STATUS: NOT VULNERABLE (your CPU microcode mitigates the vulnerability)

CVE-2018-3639 aka Variant 4, speculative store bypass
* Mitigated according to the /sys interface: YES (Mitigation: Speculative Store Bypass disabled via_
↳ prctl and seccomp)
* Kernel supports disabling speculative store bypass (SSB): YES (found in /proc/self/status)
* SSB mitigation is enabled and active: YES (per-thread through prctl)
* SSB mitigation currently active for selected processes: YES (systemd-journald systemd-logind_
↳ systemd-networkd systemd-resolved systemd-timesyncd systemd-udev)
> STATUS: NOT VULNERABLE (Mitigation: Speculative Store Bypass disabled via prctl and seccomp)

CVE-2018-3615 aka Foreshadow (SGX), L1 terminal fault
* CPU microcode mitigates the vulnerability: N/A
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-3620 aka Foreshadow-NG (OS), L1 terminal fault
* Mitigated according to the /sys interface: YES (Mitigation: PTE Inversion; VMX: conditional cache_
↳ flushes, SMT vulnerable)
* Kernel supports PTE inversion: YES (found in kernel image)
* PTE inversion enabled and active: YES
> STATUS: NOT VULNERABLE (Mitigation: PTE Inversion; VMX: conditional cache flushes, SMT vulnerable)

CVE-2018-3646 aka Foreshadow-NG (VMM), L1 terminal fault
* Information from the /sys interface: Mitigation: PTE Inversion; VMX: conditional cache flushes,_
↳ SMT vulnerable
* This system is a host running a hypervisor: NO
* Mitigation 1 (KVM)
  * EPT is disabled: NO
* Mitigation 2
  * L1D flush is supported by kernel: YES (found flush_l1d in /proc/cpuinfo)
  * L1D flush enabled: YES (conditional flushes)
  * Hardware-backed L1D flush supported: YES (performance impact of the mitigation will be greatly_
↳ reduced)
  * Hyper-Threading (SMT) is enabled: YES

```

(continues on next page)

(continued from previous page)

```

> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2018-12126 aka Fallout, microarchitectural store buffer data sampling (MSBDS)
* Kernel supports using MD_CLEAR mitigation: NO
> STATUS: VULNERABLE (Neither your kernel or your microcode support mitigation, upgrade both to
↳mitigate the vulnerability)

CVE-2018-12130 aka ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)
* Kernel supports using MD_CLEAR mitigation: NO
> STATUS: VULNERABLE (Neither your kernel or your microcode support mitigation, upgrade both to
↳mitigate the vulnerability)

CVE-2018-12127 aka RIDL, microarchitectural load port data sampling (MLPDS)
* Kernel supports using MD_CLEAR mitigation: NO
> STATUS: VULNERABLE (Neither your kernel or your microcode support mitigation, upgrade both to
↳mitigate the vulnerability)

CVE-2019-11091 aka RIDL, microarchitectural data sampling uncacheable memory (MDSUM)
* Kernel supports using MD_CLEAR mitigation: NO
> STATUS: VULNERABLE (Neither your kernel or your microcode support mitigation, upgrade both to
↳mitigate the vulnerability)

> SUMMARY: CVE-2017-5753:OK CVE-2017-5715:OK CVE-2017-5754:OK CVE-2018-3640:OK CVE-2018-3639:OK CVE-
↳2018-3615:OK CVE-2018-3620:OK CVE-2018-3646:OK CVE-2018-12126:KO CVE-2018-12130:KO CVE-2018-
↳12127:KO CVE-2019-11091:KO

```

2.9.6 Calibration Data - Denverton

Following sections include sample calibration data measured on Denverton server at Intel SH labs.

And VPP-18.10 2-Node Atom Denverton testing took place at Intel Corporation carefully adhering to FD.io CSIT best practices.

Linux cmdline

```

$ cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-4.15.0-36-generic root=UUID=d3cfff0-1e77-423a-a53a-a117199b6025 ro intel_
↳iommu=on iommu=pt isolcpus=1-11 nohz_full=1-11 rcu_nocbs=1-11 default_hugepagesz=1G hugepagesz=1G
↳hugepages=8 intel_pstate=disable nmi_watchdog=0 numa_balancing=disable tsc=reliable nosoftlockup
↳quiet splash vt.handoff=7

```

Linux uname

```

$ uname -a
Linux 4.15.0-36-generic #39~16.04.1-Ubuntu SMP Tue Sep 25 08:59:23 UTC 2018 x86_64 x86_64 x86_64
↳GNU/Linux

```

System-level Core Jitter

```

$ sudo taskset -c 2 /home/testuser/pma_tools/jitter/jitter -c 2 -i 20
Linux Jitter testing program version 1.9
Iterations=20
The program will execute a dummy function 80000 times
Display is updated every 20000 displayUpdate intervals

```

(continues on next page)

(continued from previous page)

```
Thread affinity will be set to core_id:2
Timings are in CPU Core cycles
Inst_Min: Minimum Excution time during the display update interval(default is ~1 second)
Inst_Max: Maximum Excution time during the display update interval(default is ~1 second)
Inst_jitter: Jitter in the Excution time during rhe display update interval. This is the value of
↳interest
last_Exec: The Excution time of last iteration just before the display update
Abs_Min: Absolute Minimum Excution time since the program started or statistics were reset
Abs_Max: Absolute Maximum Excution time since the program started or statistics were reset
tmp: Cumulative value calcalted by the dummy function
Interval: Time interval between the display updates in Core Cycles
Sample No: Sample number
```

↳Sample No	Inst_Min	Inst_Max	Inst_jitter	last_Exec	Abs_min	Abs_max	tmp	Interval	↳
↳1	177530	196100	18570	177530	177530	196100	4156751872	3556820054	↳
↳2	177530	200784	23254	177530	177530	200784	3210600864	3556897644	↳
↳3	177530	196346	18816	177530	177530	200784	780337152	3556918674	↳
↳4	177530	195962	18432	177530	177530	200784	1239613440	3556847928	↳
↳5	177530	195960	18430	177530	177530	200784	1698889728	3556860214	↳
↳6	177530	198824	21294	177530	177530	200784	2158166016	3556854934	↳
↳7	177530	198522	20992	177530	177530	200784	2617442304	3556862410	↳
↳8	177530	196362	18832	177530	177530	200784	3076718592	3556851636	↳
↳9	177530	199114	21584	177530	177530	200784	3535994880	3556870846	↳
↳10	177530	197194	19664	177530	177530	200784	3995271168	3556933584	↳
↳11	177530	198272	20742	177536	177530	200784	159580160	3556869044	↳
↳12	177530	197586	20056	177530	177530	200784	618856448	3556903482	↳
↳13	177530	196072	18542	177530	177530	200784	1078132736	3556825540	↳
↳14	177530	196354	18824	177530	177530	200784	1537409024	3556881664	↳
↳15	177530	195906	18376	177530	177530	200784	1996685312	3556839924	↳
↳16	177530	199066	21536	177530	177530	200784	2455961600	3556860220	↳
↳17	177530	196968	19438	177530	177530	200784	2915237888	3556871890	↳
↳18	177530	195896	18366	177530	177530	200784	3374514176	3556855338	↳
↳19	177530	196020	18490	177530	177530	200784	3833790464	3556839820	↳
↳20	177530	196030	18500	177530	177530	200784	4293066752	3556889196	↳

Memory Bandwidth

```
$ sudo /home/testuser/mlc --bandwidth_matrix
```

(continues on next page)

(continued from previous page)

```

Intel(R) Memory Latency Checker - v3.5
Command line parameters: --bandwidth_matrix

Using buffer size of 100.000MB/thread for reads and an additional 100.000MB/thread for writes
Measuring Memory Bandwidths between nodes within system
Bandwidths are in MB/sec (1 MB/sec = 1,000,000 Bytes/sec)
Using all the threads from each core if Hyper-threading is enabled
Using Read-only traffic type
      Memory node
Socket      0
           0 28157.2

```

```

$ sudo /home/testuser/mlc --peak_injection_bandwidth
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --peak_injection_bandwidth

Using buffer size of 100.000MB/thread for reads and an additional 100.000MB/thread for writes

Measuring Peak Injection Memory Bandwidths for the system
Bandwidths are in MB/sec (1 MB/sec = 1,000,000 Bytes/sec)
Using all the threads from each core if Hyper-threading is enabled
Using traffic with the following read-write ratios
ALL Reads      :      28150.0
3:1 Reads-Writes :      27425.0
2:1 Reads-Writes :      27565.4
1:1 Reads-Writes :      27489.3
Stream-triad like:      26878.2

```

```

$ sudo /home/testuser/mlc --max_bandwidth
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --max_bandwidth

Using buffer size of 100.000MB/thread for reads and an additional 100.000MB/thread for writes

Measuring Maximum Memory Bandwidths for the system
Will take several minutes to complete as multiple injection rates will be tried to get the best_
↪bandwidth
Bandwidths are in MB/sec (1 MB/sec = 1,000,000 Bytes/sec)
Using all the threads from each core if Hyper-threading is enabled
Using traffic with the following read-write ratios
ALL Reads      :      30032.40
3:1 Reads-Writes :      27450.88
2:1 Reads-Writes :      27567.46
1:1 Reads-Writes :      27501.90
Stream-triad like:      27124.82

```

Memory Latency

```

$ sudo /home/testuser/mlc --latency_matrix
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --latency_matrix

Using buffer size of 2000.000MB
Intel(R) Memory Latency Checker - v3.5
Measuring idle latencies (in ns)...
      Memory node
Socket      0
           0  93.1

```

```
$ sudo /home/testuser/mlc --idle_latency
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --idle_latency

Using buffer size of 200.000MB
Each iteration took 186.7 core clocks ( 93.4 ns)
```

```
$ sudo /home/testuser/mlc --loaded_latency
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --loaded_latency

Using buffer size of 100.000MB/thread for reads and an additional 100.000MB/thread for writes

Measuring Loaded Latencies for the system
Using all the threads from each core if Hyper-threading is enabled
Using Read-only traffic type
Inject Latency Bandwidth
Delay (ns) MB/sec
=====
00000 135.35 27186.0
00002 135.47 27176.9
00008 134.97 27063.3
00015 134.41 26825.6
00050 139.83 28419.1
00100 124.28 22616.4
00200 109.40 14139.8
00300 104.56 10275.1
00400 102.02 8120.0
00500 100.38 6751.4
00700 98.30 5124.9
01000 96.56 3852.7
01300 95.65 3149.0
01700 95.06 2585.4
02500 94.43 1988.8
03500 94.16 1621.1
05000 93.95 1343.1
09000 93.65 1052.6
20000 93.43 851.7
```

L1/L2/LLC Latency

```
$ sudo /home/testuser/mlc --c2c_latency
Intel(R) Memory Latency Checker - v3.5
Command line parameters: --c2c_latency

Measuring cache-to-cache transfer latency (in ns)...
Local Socket L2->L2 HIT latency 8.8
Local Socket L2->L2 HITM latency 8.8
```

Spectre and Meltdown Checks

Following section displays the output of a running shell script to tell if system is vulnerable against the several “speculative execution” CVEs that were made public in 2018. Script is available on [Spectre & Meltdown Checker Github](#)⁹⁸.

⁹⁸ <https://github.com/speed47/spectre-meltdown-checker>

```

Spectre and Meltdown mitigation detection tool v0.42
Checking for vulnerabilities on current system
Kernel is Linux 4.15.0-51-generic #55-Ubuntu SMP Wed May 15 14:27:21 UTC 2019 x86_64
CPU is Intel(R) Atom(TM) CPU C3858 @ 2.00GHz

Hardware check
* Hardware support (CPU microcode) for mitigation techniques
  * Indirect Branch Restricted Speculation (IBRS)
    * SPEC_CTRL MSR is available: YES
    * CPU indicates IBRS capability: YES (SPEC_CTRL feature bit)
  * Indirect Branch Prediction Barrier (IBPB)
    * PRED_CMD MSR is available: YES
    * CPU indicates IBPB capability: YES (SPEC_CTRL feature bit)
  * Single Thread Indirect Branch Predictors (STIBP)
    * SPEC_CTRL MSR is available: YES
    * CPU indicates STIBP capability: YES (Intel STIBP feature bit)
  * Speculative Store Bypass Disable (SSBD)
    * CPU indicates SSBD capability: YES (Intel SSBD)
  * L1 data cache invalidation
    * FLUSH_CMD MSR is available: NO
    * CPU indicates L1D flush capability: NO
  * Microarchitecture Data Sampling
    * VERW instruction is available: YES (MD_CLEAR feature bit)
  * Enhanced IBRS (IBRS_ALL)
    * CPU indicates ARCH_CAPABILITIES MSR availability: YES
    * ARCH_CAPABILITIES MSR advertises IBRS_ALL capability: NO
  * CPU explicitly indicates not being vulnerable to Meltdown/L1TF (RDCL_NO): YES
  * CPU explicitly indicates not being vulnerable to Variant 4 (SSB_NO): NO
  * CPU/Hypervisor indicates L1D flushing is not necessary on this system: YES
  * Hypervisor indicates host CPU might be vulnerable to RSB underflow (RSBA): NO
  * CPU explicitly indicates not being vulnerable to Microarchitectural Data Sampling (MDS_NO): YES
  * CPU supports Software Guard Extensions (SGX): NO
  * CPU microcode is known to cause stability problems: NO (model 0x5f family 0x6 stepping 0x1_
↳ ucode 0x2e cpuid 0x506f1)
  * CPU microcode is the latest known available version: awk: fatal: cannot open file `bash for_
↳ reading (No such file or directory)
UNKNOWN (latest microcode version for your CPU model is unknown)
* CPU vulnerability to the speculative execution attack variants
  * Vulnerable to CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Vulnerable to CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Vulnerable to CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Vulnerable to CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Vulnerable to CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Vulnerable to CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Vulnerable to CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): NO
  * Vulnerable to CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): NO
  * Vulnerable to CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)):_
↳ NO
  * Vulnerable to CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling_
↳ (MFBDS)): NO
  * Vulnerable to CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Vulnerable to CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory_
↳ (MDSUM)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: __user pointer sanitization)
* Kernel has array_index_mask_nospec: YES (1 occurrence(s) found of x86 64 bits array_index_mask_
↳ nospec())
* Kernel has the Red Hat/Ubuntu patch: NO
* Kernel has mask_nospec64 (arm64): NO
> STATUS: NOT VULNERABLE (Mitigation: __user pointer sanitization)

```

(continues on next page)

(continued from previous page)

```

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: YES (Mitigation: Full generic retpoline, IBPB:
↳conditional, IBRS_FW, STIBP: disabled, RSB filling)
* Mitigation 1
  * Kernel is compiled with IBRS support: YES
    * IBRS enabled and active: YES (for firmware code only)
  * Kernel is compiled with IBPB support: YES
    * IBPB enabled and active: YES
* Mitigation 2
  * Kernel has branch predictor hardening (arm): NO
  * Kernel compiled with retpoline option: YES
    * Kernel compiled with a retpoline-aware compiler: YES (kernel reports full retpoline
↳compilation)
> STATUS: NOT VULNERABLE (Full retpoline + IBPB are mitigating the vulnerability)

CVE-2017-5754 aka Variant 3, Meltdown, rogue data cache load
* Mitigated according to the /sys interface: YES (Not affected)
* Kernel supports Page Table Isolation (PTI): YES
  * PTI enabled and active: UNKNOWN (dmesg truncated, please reboot and relaunch this script)
  * Reduced performance impact of PTI: NO (PCID/INVCID not supported, performance impact of PTI
↳will be significant)
* Running as a Xen PV DomU: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-3640 aka Variant 3a, rogue system register read
* CPU microcode mitigates the vulnerability: YES
> STATUS: NOT VULNERABLE (your CPU microcode mitigates the vulnerability)

CVE-2018-3639 aka Variant 4, speculative store bypass
* Mitigated according to the /sys interface: YES (Mitigation: Speculative Store Bypass disabled via
↳prctl and seccomp)
* Kernel supports disabling speculative store bypass (SSB): YES (found in /proc/self/status)
* SSB mitigation is enabled and active: YES (per-thread through prctl)
* SSB mitigation currently active for selected processes: YES (systemd-journald systemd-logind
↳systemd-networkd systemd-resolved systemd-timesyncd systemd-udev)
> STATUS: NOT VULNERABLE (Mitigation: Speculative Store Bypass disabled via prctl and seccomp)

CVE-2018-3615 aka Foreshadow (SGX), L1 terminal fault
* CPU microcode mitigates the vulnerability: N/A
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-3620 aka Foreshadow-NG (OS), L1 terminal fault
* Mitigated according to the /sys interface: YES (Not affected)
* Kernel supports PTE inversion: YES (found in kernel image)
* PTE inversion enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-3646 aka Foreshadow-NG (VMM), L1 terminal fault
* Information from the /sys interface: Not affected
* This system is a host running a hypervisor: NO
* Mitigation 1 (KVM)
  * EPT is disabled: NO
* Mitigation 2
  * L1D flush is supported by kernel: YES (found flush_l1d in kernel image)
  * L1D flush enabled: NO
  * Hardware-backed L1D flush supported: NO (flush will be done in software, this is slower)
  * Hyper-Threading (SMT) is enabled: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12126 aka Fallout, microarchitectural store buffer data sampling (MSBDS)
* Mitigated according to the /sys interface: YES (Not affected)

```

(continues on next page)

(continued from previous page)

```

* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: NO
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12130 aka ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: NO
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12127 aka RIDL, microarchitectural load port data sampling (MLPDS)
* Mitigated according to the /sys interface: YES (Not affected)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: NO
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2019-11091 aka RIDL, microarchitectural data sampling uncacheable memory (MDSUM)
* Mitigated according to the /sys interface: YES (Not affected)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: NO
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

> SUMMARY: CVE-2017-5753:OK CVE-2017-5715:OK CVE-2017-5754:OK CVE-2018-3640:OK CVE-2018-3639:OK CVE-
↳2018-3615:OK CVE-2018-3620:OK CVE-2018-3646:OK CVE-2018-12126:OK CVE-2018-12130:OK CVE-2018-
↳12127:OK CVE-2019-11091:OK

```

2.9.7 SUT Settings - Linux

System provisioning is done by combination of PXE boot unattended install and [Ansible](#)⁹⁹ described in [CSIT Testbed Setup](#)¹⁰⁰.

Below a subset of the running configuration:

1. Xeon Haswell - Ubuntu 18.04.1 LTS

```

$ lsb_release -a
No LSB modules are available.
Distributor ID: Ubuntu
Description:    Ubuntu 18.04.1 LTS
Release:        18.04
Codename:       bionic

```

2. Xeon Skylake - Ubuntu 18.04 LTS

```

$ lsb_release -a
No LSB modules are available.
Distributor ID: Ubuntu
Description:    Ubuntu 18.04 LTS
Release:        18.04
Codename:       bionic

```

⁹⁹ <https://www.ansible.com>

¹⁰⁰ https://git.fd.io/csit/tree/resources/tools/testbed-setup/README.md?h=rls1908_1

Linux Boot Parameters

- **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>
- **numa_balancing=disable** - [KNL,X86] Disable automatic NUMA balancing.
- **intel_iommu=enable** - [DMAR] Enable Intel IOMMU driver (DMAR) option.
- **iommu=on, iommu=pt** - [x86, IA-64] Disable IOMMU bypass, using IOMMU for PCI devices.
- **nmi_watchdog=0** - [KNL,BUGS=X86] Debugging features for SMP kernels. Turn hardlockup detector in nmi_watchdog off.
- **nosoftlockup** - [KNL] Disable the soft-lockup detector.
- **tsc=reliable** - Disable clocksource stability checks for TSC. [x86] reliable: mark tsc clocksource as reliable, this disables clocksource verification at runtime, as well as the stability checks done at bootup. Used to enable high-resolution timer mode on older hardware, and in virtualized environment.
- **hpet=disable** - [X86-32,HPET] Disable HPET and use PIT instead.

Hugepages Configuration

Huge pages are managed via sysctl configuration located in `/etc/sysctl.d/90-csit.conf` on each testbed. Default huge page size is 2M. The exact amount of huge pages depends on testbed. All the values are defined in *Ansible inventory - hosts* files.

Applied Boot Cmdline

1. Xeon Haswell - Ubuntu 18.04.1 LTS

```
$ cat /proc/cmdline
BOOT_IMAGE=/vmlinuz-4.15.0-36-generic root=UUID=5d2ecc97-245b-4e94-b0ae-c3548567de19 ro isolcpus=1-
↪17,19-35 nohz_full=1-17,19-35 rcu_nocbs=1-17,19-35 numa_balancing=disable intel_pstate=disable_
↪intel_iommu=on iommu=pt nmi_watchdog=0 audit=0 nosoftlockup processor.max_cstate=1 intel_idle.max_
↪cstate=1 hpet=disable tsc=reliable mce=off console=tty0 console=ttyS0,115200n8
```

2. Xeon Skylake - Ubuntu 18.04 LTS

```
$ cat /proc/cmdline
BOOT_IMAGE=/vmlinuz-4.15.0-23-generic root=UUID=3fa246fd-1b80-4361-bb90-f339a6bbed51 ro isolcpus=1-
↪27,29-55,57-83,85-111 nohz_full=1-27,29-55,57-83,85-111 rcu_nocbs=1-27,29-55,57-83,85-111 numa_
↪balancing=disable intel_pstate=disable intel_iommu=on iommu=pt nmi_watchdog=0 audit=0_
↪nosoftlockup processor.max_cstate=1 intel_idle.max_cstate=1 hpet=disable tsc=reliable mce=off_
↪console=tty0 console=ttyS0,115200n8
```

2.9. Test Environment

2.9.8 DUT Settings - VPP

VPP Version

VPP-19.08.1 release

VPP Compile Parameters

FD.io VPP compile job¹⁰¹

VPP Install Parameters

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

VPP Startup Configuration

VPP startup configuration vary per test case, with different settings for `$$CORELIST_WORKERS`, `$$NUM_RX_QUEUES`, `$$UIO_DRIVER`, `$$NUM-MBUFS` and `$$NO_MULTI_SEG` parameter. Default template is provided below:

```
ip
{
  heap-size 4G
}
statseg
{
  size 4G
}
unix
{
  cli-listen /run/vpp/cli.sock
  log /tmp/vpe.log
  nodaemon
}
socksvr {
  default
}
ip6
{
  heap-size 4G
  hash-buckets 2000000
}
heapsize 4G
plugins
{
  plugin default
  {
    disable
  }
  plugin dpdk_plugin.so
  {
```

(continues on next page)

¹⁰¹ https://jenkins.fd.io/view/vpp/job/vpp-merge-1908_1-ubuntu1604/

(continued from previous page)

```

    enable
  }
}
cpu
{
  corelist-workers $$CORELIST_WORKERS
  main-core 1
}
dpdk
{
  num-mbufs $$NUM-MBUFS
  uio-driver $$UIO_DRIVER
  $$NO_MULTI_SEG
  log-level debug
  dev default
  {
    num-rx-queues $$NUM_RX_QUEUES
  }
  no-tx-checksum-offload
  dev $$DEV_1
  dev $$DEV_2
}

```

Description of VPP startup settings used in CSIT is provided in *Test Methodology* (page 10).

2.9.9 TG Settings - TRex

TG Version

TRex v2.54

DPDK Version

DPDK v19.02

TG Build Script Used

TRex intallation¹⁰²

TG Startup Configuration

```

$ cat /etc/trex_cfg.yaml
- version      : 2
  interfaces   : ["0000:0d:00:00", "0000:0d:00:01"]
  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]

```

¹⁰² https://git.fd.io/csit/tree/resources/tools/trex/trex_installer.sh?h=rls1908_1

TG Startup Command

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

TG API Driver

Trex driver¹⁰³

¹⁰³ https://git.fd.io/csit/tree/resources/tools/trex/trex_stateless_profile.py?h=rls1908_1

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 503), *[lxcnamespace]* (page 503) and *[stgraber]* (page 503).

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 *[lxcsecurity]* (page 503) and *[stgraber]* (page 503).

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 503).
- Control groups - cgroups, resource bean counting, resource quotas, access restrictions, *[cgroup1]* (page 503), *[cgroup2]* (page 503).
- AppArmor - apparmor profiles aim to prevent any of the known ways of escaping a container or cause harm to the host, *[apparmor]* (page 503).
- 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 503).
- Seccomp - secure computing mode, enables filtering of system calls, *[seccomp]* (page 503).

More information in *[lxcsecurity]* (page 503) and *[lxcsecfeatures]* (page 503).

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/uiopci_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 [GitHub](#)¹⁰⁴.
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 [[lxcsource](#)] (page 503), 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`.

¹⁰⁴ <https://github.com/lxc/lxc/tree/master/templates>

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 503). Currently it using *runc* a CLI tool for spawning and running containers according to the [OCI specification](https://www.opencontainers.org/)¹⁰⁵

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 vpp startup.conf.

Kubernetes

Kubernetes [k8sdoc] (page 503), 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 503) 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 503) and vpp-agent [vppagent] (page 503).

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

¹⁰⁵ <https://www.opencontainers.org/>




Sequential diagram that illustrates the creation of a single container.

Legend:
 e = engine [Docker|LXC]
 .. = kwargs (variable number of keyword argument)

(continues on next page)

(continued from previous page)



(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
| | ${skip_cpus}= | Evaluate | ${vpp_cpus}+${system_cpus}
| | Import Library | resources.libraries.python.ContainerUtils.ContainerManager
| | ... | engine=${container_engine} | WITH NAME | ${group}
| | ${duts}= | Get Matches | ${nodes} | DUT*
| | :FOR | ${dut} | IN | @${duts}
| | | ${env}= | Create List | DEBIAN_FRONTEND=noninteractive
| | | ${mnt}= | Create List | /tmp:/mnt/host | /dev:/dev
| | | ${cpu_node}= | Get interfaces numa node | ${nodes['${dut}']}
| | | ... | ${dut1_if1} | ${dut1_if2}
| | | Run Keyword | ${group}.Construct containers
| | | ... | name=${dut}_${group} | node=${nodes['${dut}']} | mnt=${mnt}
| | | ... | image=${container_image} | cpu_count=${container_cpus}
| | | ... | cpu_skip=${skip_cpus} | cpuset_mems=${cpu_node}
| | | ... | cpu_shared=${False} | env=${env} | count=${container_count}
| | | ... | install_dkms=${container_install_dkms}
| | Append To List | ${container_groups} | ${group}
```

Mandatory parameters to create standalone container are: node, name, image [*imagevar*] (page 504), cpu_count, cpu_skip, 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/k&s_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 adress 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-1908.1:

- LXC topologies:
 - eth-l2xcbase-eth-2memif-1lxc.
 - eth-l2bdbasemaclrn-eth-2memif-1lxc.
- Docker topologies:
 - eth-l2xcbase-eth-2memif-1docker.
 - eth-l2xcbase-eth-1memif-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 Test Code Documentation

[CSIT VPP Performance Tests Documentation](https://docs.fd.io/csit/rls1908_1/doc/tests.vpp.perf.html)¹²³ contains detailed functional description and input parameters for each test case.

¹²³ https://docs.fd.io/csit/rls1908_1/doc/tests.vpp.perf.html

3.1 Overview


3.1.1 Virtual Topologies

CSIT VPP Device tests are executed in Physical containerized topologies created on demand using set of scripts hosted and developed under CSIT repository. It runs on physical baremetal servers hosted by LF FD.io project. Based on the packet path thru SUT Containers, three distinct logical topology types are used for VPP DUT data plane testing:

1. vfNIC-to-vfNIC switching topologies.
2. vfNIC-to-vhost-user switching topologies.
3. vfNIC-to-memif switching topologies.

vfNIC-to-vfNIC Switching

The simplest physical topology for software data plane application like VPP is vfNIC-to-vfNIC switching. Tested virtual topologies for 2-Node testbeds are shown in figures below.



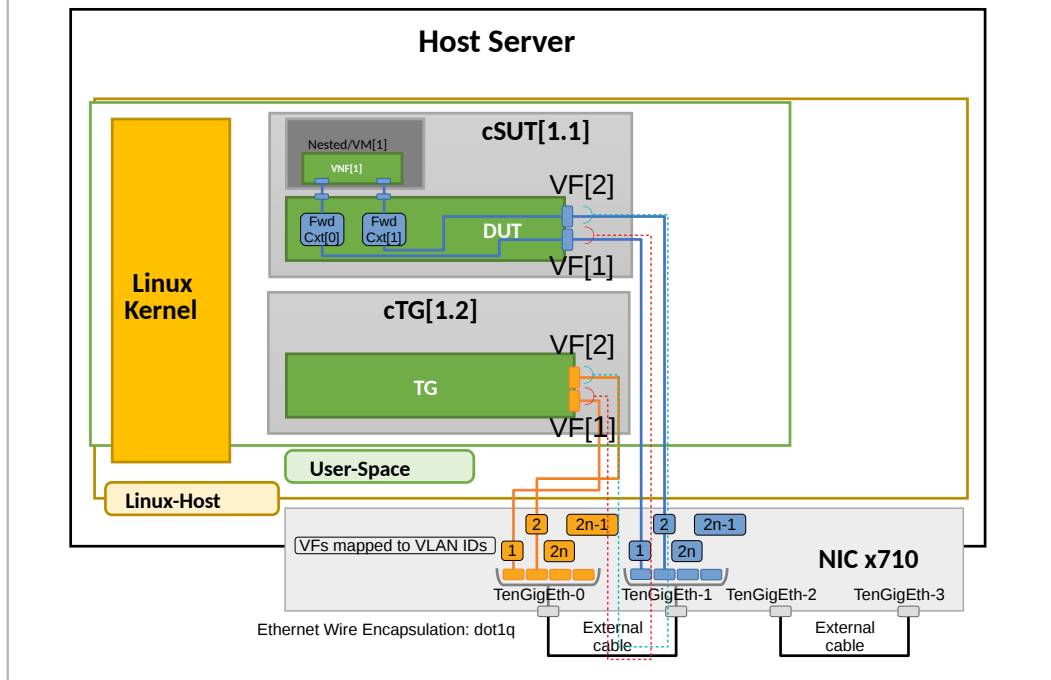
SUT1 is Docker Container (running Ubuntu, depending on the test suite), TG is a Traffic Generator (running Ubuntu Container). SUTs run VPP SW application in Linux user-mode as a Device Under Test (DUT) within the container. TG runs Scapy SW application as a packet Traffic Generator. Network connectivity between SUTs and to TG is provided using virtual function of physical NICs.

Virtual topologies 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 container running on physical server. During the test execution, all nodes are reachable thru the Management (not shown above for clarity).

vfNIC-to-vhost-user Switching


vfNIC-to-vhost-user switching topology test cases require VPP DUT to communicate with Virtual Machine (VM) over Vhost-user virtual interfaces. VM is created on SUT1 for the duration of these particular test cases only. Virtual test topology with VM is shown in the figure below.

2-Node Containerized Topologies: vfnic-to-vhost-user switching



vfnic-to-memif Switching

vfnic-to-memif switching topology test cases require VPP DUT to communicate with another Docker Container over memif interfaces. Container is created for the duration of these particular test cases only and it is running the same VPP version as running on DUT. Virtual test topology with Memif is shown in the figure below.



3.1.2 Functional Tests Coverage

CSIT-1908.1 includes following VPP functionality tested in VPP Device environment:

Functionality	Description
ACL	Ingress Access Control List security for L2 Bridge-Domain MAC switching, IPv4 routing, IPv6 routing.
COP	COP address white-list and black-list filtering for IPv4 and IPv6 routing.
IPSec	IPSec tunnel and transport modes.
IPv4	IPv4 routing, ICMPv4.
IPv6	IPv4 routing, ICMPv6.
L2BD	L2 Bridge-Domain switching for untagged Ethernet.
L2XC	L2 Cross-Connect switching for untagged Ethernet.
Memif Interface	Baseline VPP memif interface tests.
QoS Policer Metering	Ingress packet rate metering and marking for IPv4, IPv6.
Tap Interface	Baseline Linux tap interface tests.
VLAN Tag	L2 VLAN subinterfaces.
Vhost-user Interface	Baseline VPP vhost-user interface tests.
VXLAN	VXLAN overlay tunneling for L2-over-IPv4 and -over-IPv6.

3.1.3 Tests Naming

CSIT-1908.1 follows a common structured naming convention for all performance and system functional tests, introduced in CSIT-17.01.

The naming should be intuitive for majority of the tests. Complete description of CSIT test naming convention is provided on *Test Naming* (page 459).

3.2 Release Notes

3.2.1 Changes in CSIT-1908.1

1. TEST FRAMEWORK

- **Bug fixes.**
- **ARM platform compatibility.**

2. TEST COVERAGE

- Increased test coverage: **Dot1q, IPsec, 802.1ad VXLAN, COP whitelist, COP blacklist, QoS Policer Metering, iACL whitelist, AVF driver, TAP Interface.**
- Align vpp_device L2 Robot Keywords with performance L2 Robot Keywords.

3.2.2 Known Issues

List of known issues in CSIT-1908.1 for VPP functional tests in VPP Device:

#	JiraID	Issue Description
1		

3.3 Integration Tests

3.3.1 Abstract


FD.io VPP software data plane technology has become very popular across a wide range of VPP ecosystem use cases, putting higher pressure on continuous verification of VPP software quality.

This document describes a proposal for design and implementation of extended continuous VPP testing by extending existing test environments. Furthermore it describes and summarizes implementation details of Integration and System tests platform *1-Node VPP_Device*. It aims to provide a complete end-to-end view of *1-Node VPP_Device* environment in order to improve extendability and maintenance, under the guideline of VPP core team.

The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this document are to be interpreted as described in **RFC 8174**¹²⁴.

¹²⁴ <https://tools.ietf.org/html/rfc8174.html>

3.3.2 Overview



3.3.3 Physical Testbeds

All FD.io CSIT vpp-device tests are executed on physical testbeds built with bare-metal servers hosted by LF FD.io project. Two 1-node testbed topologies are used:


- **2-Container Topology:** Consisting of one Docker container acting as SUT (System Under Test) and one Docker container as TG (Traffic Generator), both connected in ring topology via physical NIC cross-connecting.

Current FD.io production testbeds are built with servers based on one processor generation of Intel Xeons: Skylake (Platinum 8180). Testbeds built with servers based on Arm processors are in the process of being added to FD.io production.

Following section describe existing production 1n-skx testbed.

1-Node Xeon Skylake (1n-skx)

1n-skx testbed is based on single SuperMicro SYS-7049GP-TRT server equipped with two Intel Xeon Skylake Platinum 8180 2.5 GHz 28 core processors. Physical testbed topology is depicted in a figure below.



Server is populated with the following NIC models:

1. NIC-1: x710-da4 4p10GE Intel.
2. NIC-2: x710-da4 4p10GE Intel.

All Intel Xeon Skylake servers run with Intel Hyper-Threading enabled, doubling the number of logical cores exposed to Linux, with 56 logical cores and 28 physical cores per processor socket.

NIC interfaces are shared using Linux vfiio_pci and VPP VF drivers:

- DPDK VF driver,
- Fortville AVF driver.

Provided Intel x710-da4 4p10GE NICs support 32 VFs per interface, 128 per NIC.

Complete 1n-skx testbeds specification is available on [CSIT LF Testbeds¹²⁵](#) wiki page.

Total of two 1n-skx testbeds are in operation in FD.io labs.

1-Node Virtualbox (1n-vbox)

1n-skx testbed can run in single VirtualBox VM machine. This solution replaces the previously used Vagrant environment based on 3 VMs.

VirtualBox VM MAY be created by Vagrant and MUST have additional 4 virtio NICs each pair attached to separate private networks to simulate back-to-back connections. It SHOULD be 82545EM device model (otherwise can be changed in bootstrap scripts). Example of Vagrant configuration:

```
Vagrant.configure(2) do |c|
  c.vm.network "private_network", type: "dhcp", auto_config: false,
```

(continues on next page)

¹²⁵ https://wiki.fd.io/view/CSIT/Testbeds:_Xeon_Skx,_Arm,_Atom.

(continued from previous page)

```
virtualbox__intnet: "port1", nic_type: "82545EM"
c.vm.network "private_network", type: "dhcp", auto_config: false,
  virtualbox__intnet: "port2", nic_type: "82545EM"

c.vm.provider :virtualbox do |v|
  v.customize ["modifyvm", :id, "--nicpromisc2", "allow-all"]
  v.customize ["modifyvm", :id, "--nicpromisc3", "allow-all"]
  v.customize ["modifyvm", :id, "--nicpromisc4", "allow-all"]
  v.customize ["modifyvm", :id, "--nicpromisc5", "allow-all"]
```

Vagrant VM is populated with the following NIC models:

1. NIC-1: 82545EM Intel.
2. NIC-2: 82545EM Intel.
3. NIC-3: 82545EM Intel.
4. NIC-4: 82545EM Intel.

3.3.4 Containers

It was agreed on TWS (Technical Work Stream) call to continue with Ubuntu 18.04 LTS as a baseline system with OPTIONAL extend to Centos 7 and SuSE per demand [TWSLink] (page 504).

All DCR (Docker container) images are REQUIRED to be hosted on Docker registry available from LF network, publicly available and trackable. For backup, tracking and contributing purposes all Dockerfiles (including files needed for building container) MUST be available and stored in [fdiocsitgerrit] (page 504) repository under appropriate folders. This allows the peer review process to be done for every change of infrastructure related to scope of this document. Currently only **csit-shim-dcr** and **csit-sut-dcr** containers will be stored and maintained under CSIT repository by CSIT contributors.

At the time of designing solution described in this document the interconnection between [dockerhub] (page 504) and [fdiocsitgerrit] (page 504) for automated build purposes and image hosting cannot be established with the trust and respectful to security of FD.io project. Unless addressed, DCR images will be placed in custom registry service [fdioregistry] (page 504). Automated Jenkins jobs will be created in align of long term solution for container lifecycle and ability to build new version of docker images.

In parallel, the effort is started to find the outsourced Docker registry service.

Versioning

As of initial version of vpp-device, we do have only single latest version of Docker image hosted on [dockerhub] (page 504). This will be addressed as further improvement with proper semantic versioning.

jenkins-slave-dcr

This DCR acts as the Jenkins slave (known also as jenkins minion). It can connect over SSH protocol to TCP port 6022 of **csit-shim-dcr** and executes non-interactive reservation script. Nomad is responsible for scheduling this container execution onto specific **1-Node VPP_Device** testbed. It executes CSIT environment including CSIT framework.

All software dependencies including VPP/DPDK that are not present in **csit-sut-dcr** container image and/or needs to be compiled prior running on **csit-sut-dcr** SHOULD be compiled in this container.

- *Container Image Location*: Docker image at snergster/vpp-ubuntu18.
- *Container Definition*: Docker file specified at [JenkinsSlaveDcrFile] (page 504).
- *Initializing*: Container is initialized from within *Consul by HashiCorp* and *Nomad by HashiCorp*.

csit-shim-dcr

This DCR acts as an intermediate layer running script responsible for orchestrating topologies under test and reservation. Responsible for managing VF resources and allocation to DUT (Device Under Test), TG (Traffic Generator) containers. This MUST to be done on **csit-shim-dcr**. This image also acts as the generic reservation mechanics arbiter to make sure that only Y number of simulations are spawned on any given HW node.

- *Container Image Location*: Docker image at snergster/csit-shim.
- *Container Definition*: Docker file specified at [CsitShimDcrFile] (page 504).
- *Initializing*: Container is initialized from within *Consul by HashiCorp* and *Nomad by HashiCorp*. Required docker parameters, to be able to run nested containers with VF reservation system are: privileged, net=host, pid=host.
- *Connectivity*: Over SSH only, using <host>:6022 format. Currently using root user account as primary. From the jenkins slave it will be able to connect via env variable, since the jenkins slave doesn't actually know what host its running on.

```
ssh -p 6022 root@10.30.51.node
```

csit-sut-dcr

This DCR acts as an SUT (System Under Test). Any DUT or TG application is installed there. It is RECOMMENDED to install DUT and all DUT dependencies via commands `rpm -ihv` on RedHat based OS or `dpkg -i` on Debian based OS.

Container is designed to be a very lightweight Docker image that only installs packages and execute binaries (previously built or downloaded on **jenkins-slave-dcr**) and contains libraries necessary to run CSIT framework including those required by DUT/TG.

- *Container Image Location*: Docker image at snergster/csit-sut.
- *Container Definition*: Docker file specified at [CsitSutDcrFile] (page 504).
- *Initializing*:

```
docker run
# Run the container in the background and print the new container ID.
--detach=true
# Give extended privileges to this container. A "privileged" container is
# given access to all devices and able to run nested containers.
--privileged
# Publish all exposed ports to random ports on the host interfaces.
--publish-all
# Automatically remove the container when it exits.
--rm
# Size of /dev/shm.
dcr_stc_params+="--shm-size 512M "
# Override access to PCI bus by attaching a filesystem mount to the
# container.
dcr_stc_params+="--mount type=tmpfs,destination=/sys/bus/pci/devices "
# Mount vfio to be able to bind to see binded interfaces. We cannot use
# --device=/dev/vfio as this does not see newly binded interfaces.
dcr_stc_params+="--volume /dev/vfio:/dev/vfio "
# Mount docker.sock to be able to use docker daemon of the host.
dcr_stc_params+="--volume /var/run/docker.sock:/var/run/docker.sock "
# Mount /opt/boot/ where VM kernel and initrd are located.
dcr_stc_params+="--volume /opt/boot:/opt/boot/ "
# Mount host hugepages for VMs.
dcr_stc_params+="--volume /dev/hugepages:/dev/hugepages/ "
```

Container name is catenated from **csit-** prefix and uuid generated uniquely for each container instance.

- **Connectivity:** Over SSH only, using <host>[:<port>] format. Currently using *root* user account as primary.

```
ssh -p <port> root@10.30.51.<node>
```

Container required to run as `--privileged` due to ability to create nested containers and have full read/write access to sysfs (for bind/unbind). Docker automatically pick free network port (`--publish-all`) for ability to connect over ssh. To be able to limit access to PCI bus, container is creating tmpfs mount type in PCI bus tree. CSIT reservation script is dynamically linking only PCI devices (NIC cards) that are reserved for particular container. This way it is not colliding with other containers. To make vfiowork, access to `/dev/vfio` must be granted.

3.3.5 Environment initialization

All 1-node servers are to be managed and provisioned via the [\[ansiblelink\]](#) (page 504) set of playbooks with *vpp-device* role. Full playbooks can be found under [\[fdiocsitansible\]](#) (page 504) directory. This way we are able to track all configuration changes of physical servers in gerrit (in structured yaml format) as well as we are able to extend *vpp-device* to additional servers with less effort or re-stage servers in case of failure.

SR-IOV VF initialization is done via `systemd` service during host system boot up. Service with name *csit-initialize-vfs.service* is created under `systemd` system context (`/etc/systemd/system/`). By default service is calling `/usr/local/bin/csit-initialize-vfs.sh` with single parameter:

- **start:** Creates maximum number of virtual functions (VFs) (detected from `sriov_totalvfs`) for each whitelisted PCI device.
- **stop:** Removes all VFs for all whitelisted PCI device.

Service is considered active even when all of its processes exited successfully. Stopping service will automatically remove VFs.

```
[Unit]
Description=CSIT Initialize SR-IOV VFs
After=network.target

[Service]
Type=one-shot
RemainAfterExit=True
ExecStart=/usr/local/bin/csit-initialize-vfs.sh start
ExecStop=/usr/local/bin/csit-initialize-vfs.sh stop

[Install]
WantedBy=default.target
```

Script is driven by two array variables `pci_blacklist`/`pci_whitelist`. They MUST store all PCI addresses in `<domain>:<bus>:<device>.<func>` format, where:

- **pci_blacklist:** PCI addresses to be skipped from VFs initialization (usefull for e.g. excluding management network interfaces).
- **pci_whitelist:** PCI addresses to be included for VFs initialization.

3.3.6 VF reservation

During topology initialization phase of script, mutex is used to avoid multiple instances of script to interact with each other during resources allocation. Mutual exclusion ensure that no two distinct instances of script will get same resource list.

Reservation function reads the list of all available virtual function network devices in system:

```
# Find the first ${device_count} number of available TG Linux network
# VF device names. Only allowed VF PCI IDs are filtered.
for netdev in ${tg_netdev[@]}
do
    for netdev_path in $(grep -l "${pci_id}" \
        /sys/class/net/${netdev}*/device/device \
        2> /dev/null)
    do
        if [[ ${#TG_NETDEVS[@]} -lt ${device_count} ]]; then
            tg_netdev_name=$(dirname ${netdev_path})
            tg_netdev_name=$(dirname ${tg_netdev_name})
            TG_NETDEVS+=($(basename ${tg_netdev_name}))
        else
            break
        fi
    done
    if [[ ${#TG_NETDEVS[@]} -eq ${device_count} ]]; then
        break
    fi
done
```

Where `${pci_id}` is ID of white-listed VF PCI ID. For more information please see [\[pciids\]](#) (page 504). This act as security constraint to prevent taking other unwanted interfaces. The output list of all VF network devices is split into two lists for TG and SUT side of connection. First two items from each TG or SUT network devices list are taken to expose directly to namespace of container. This can be done via commands:

```
$ ip link set ${netdev} netns ${DCR_CPIDS[tg]}
$ ip link set ${netdev} netns ${DCR_CPIDS[dut1]}
```

In this stage also symbolic links to PCI devices under `sysfs` bus directory tree are created in running containers. Once VF devices are assigned to container namespace and PCI devices are linked to running containers and mutex is exited. Selected VF network device automatically disappear from parent container namespace, so another instance of script will not find device under that namespace.

Once Docker container exits, network device is returned back into parent namespace and can be reused.

3.3.7 Network traffic isolation - Intel i40evf

In a virtualized environment, on Intel(R) Server Adapters that support SR-IOV, the virtual function (VF) may be subject to malicious behavior. Software-generated layer two frames, like IEEE 802.3x (link flow control), IEEE 802.1Qbb (priority based flow-control), and others of this type, are not expected and can throttle traffic between the host and the virtual switch, reducing performance. To resolve this issue, configure all SR-IOV enabled ports for VLAN tagging. This configuration allows unexpected, and potentially malicious, frames to be dropped. [\[inteli40e\]](#) (page 504)

To configure VLAN tagging for the ports on an SR-IOV enabled adapter, use the following command. The VLAN configuration SHOULD be done before the VF driver is loaded or the VM is booted. [\[inteli40e\]](#) (page 504)

```
$ ip link set dev <PF netdev id> vf <id> vlan <vlan id>
```

For example, the following instructions will configure PF eth0 and the first VF on VLAN 10.

```
$ ip link set dev eth0 vf 0 vlan 10
```

VLAN Tag Packet Steering allows to send all packets with a specific VLAN tag to a particular SR-IOV virtual function (VF). Further, this feature allows to designate a particular VF as trusted, and allows that trusted VF to request selective promiscuous mode on the Physical Function (PF). [\[inteli40e\]](#) (page 504)

To set a VF as trusted or untrusted, enter the following command in the Hypervisor:

```
$ ip link set dev eth0 vf 1 trust [on|off]
```

Once the VF is designated as trusted, use the following commands in the VM to set the VF to promiscuous mode. [intel40e] (page 504)

- For promiscuous all:

```
$ ip link set eth2 promisc on
```

- For promiscuous Multicast:

```
$ ip link set eth2 allmulti on
```

Note: By default, the ethtool priv-flag vf-true-promisc-support is set to *off*, meaning that promiscuous mode for the VF will be limited. To set the promiscuous mode for the VF to true promiscuous and allow the VF to see all ingress traffic, use the following command. `$ ethtool set-priv-flags p261p1 vf-true-promisc-support on` The vf-true-promisc-support priv-flag does not enable promiscuous mode; rather, it designates which type of promiscuous mode (limited or true) you will get when you enable promiscuous mode using the ip link commands above. Note that this is a global setting that affects the entire device. However, the vf-true-promisc-support priv-flag is only exposed to the first PF of the device. The PF remains in limited promiscuous mode (unless it is in MFP mode) regardless of the vf-true-promisc-support setting. [intel40e] (page 504)

Service described earlier `csit-initialize-vfs.service` is responsible for assigning 802.1Q vlan tagging to each virtual function via physical function from list of white-listed PCI addresses by following (simplified) code.

```
SCRIPT_DIR="$(dirname $(readlink -e "${BASH_SOURCE[0]}"))"
source "${SCRIPT_DIR}/csit-initialize-vfs-data.sh"

# Initilize whitelisted NICs with maximum number of VFs.
pci_idx=0
for pci_addr in ${PCI_WHITELIST[@]}; do
    if ! [[ ${PCI_BLACKLIST[*]} =~ "${pci_addr}" ]]; then
        pci_path="/sys/bus/pci/devices/${pci_addr}"
        # SR-IOV initialization
        case "${1:-start}" in
            "start" )
                sriov_totalvfs=$(< "${pci_path}"/sriov_totalvfs)
                ;;
            "stop" )
                sriov_totalvfs=0
                ;;
        esac
        echo ${sriov_totalvfs} > "${pci_path}"/sriov_numvfs
        # SR-IOV 802.1Q isolation
        case "${1:-start}" in
            "start" )
                pf=$(basename "${pci_path}"/net/*)
                for vf in $(seq "${sriov_totalvfs}"); do
                    # PCI address index in array (pairing siblings).
                    if [[ -n ${PF_INDICES[@]} ]]
                    then
                        vlan_pf_idx=${PF_INDICES[$pci_addr]}
                    else
                        vlan_pf_idx=$((pci_idx % (${#PCI_WHITELIST[@]}/2))
                    fi
                    # 802.1Q base offset.
                    vlan_bs_off=1100
                done
            ;;
        esac
    done
done
```

(continues on next page)

(continued from previous page)

```

# 802.1Q PF PCI address offset.
vlan_pf_off=$(( vlan_pf_idx * 100 + vlan_bs_off ))
# 802.1Q VF PCI address offset.
vlan_vf_off=$(( vlan_pf_off + vf - 1 ))
# VLAN string.
vlan_str="vlan ${vlan_vf_off}"
# MAC string.
mac5="$(printf '%x' ${pci_idx})"
mac6="$(printf '%x' $(( vf - 1 )))"
mac_str="mac ba:dc:0f:fe:${mac5}:${mac6}"
# Set 802.1Q VLAN id and MAC address
ip link set ${pf} vf $(( vf - 1 )) ${mac_str} ${vlan_str}
ip link set ${pf} vf $(( vf - 1 )) trust on
ip link set ${pf} vf $(( vf - 1 )) spoof off
done
pci_idx=$(( pci_idx + 1 ))
;;
esac
rmmod i40evf
modprobe i40evf
fi
done

```

Assignment starts at VLAN 1100 and incrementing by 1 for each VF and by 100 for each white-listed PCI address up to the middle of the PCI list. Second half of the lists is assumed to be directly (cable) paired siblings and assigned with same 802.1Q VLANs as its siblings.

3.3.8 Open tasks

Security

Note: Switch to non-privileged containers: As of now all three container flavors are using privileged containers to make it working. Explore options to switch containers to non-privileged with explicit rather implicit privileges.

Note: Switch to testuser account instead of root.

Maintainability

Note: Docker image distribution: Create Jenkins jobs with full pipeline of CI/CD for CSIT Docker images.

Stability

Note: Implement queueing mechanism: Currently there is no mechanics that would place starving jobs in queue in case of no resources available.

Note: Replace reservation script with Docker network plugin written in GOLANG/SH/Python - platform independent.

3.3.9 Links

3.4 Documentation

CSIT VPP Device Tests Documentation¹³⁶ contains detailed functional description and input parameters for each test case.

¹³⁶ https://docs.fd.io/csit/rls1908_1/doc/tests.vpp.device.html

4.1 Overview


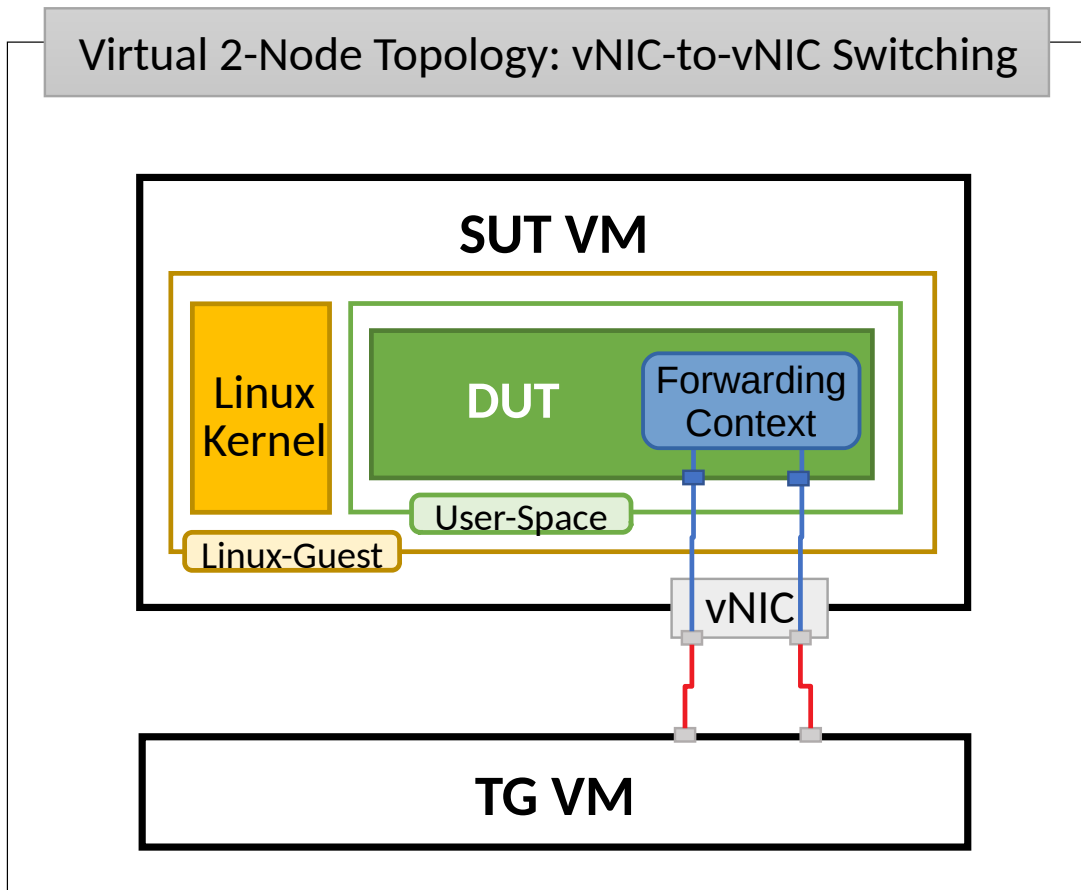
4.1.1 Virtual Topologies

CSIT VPP functional tests are executed in VM-based virtual topologies created on demand using VIRL (Virtual Internet Routing Lab) simulation platform contributed by Cisco. VIRL runs on physical baremetal servers hosted by LF FD.io project. Based on the packet path thru SUT VMs, two distinct logical topology types are used for VPP DUT data plane testing:

1. vNIC-to-vNIC switching topologies.
2. Nested-VM service switching topologies.

vNIC-to-vNIC Switching

The simplest virtual topology for software data plane application like VPP is vNIC-to-vNIC switching. Tested virtual topologies for 2-Node and 3-Node testbeds are shown in figures below.




SUT1 and SUT2 are two VMs (running Ubuntu or Centos, depending on the test suite), TG is a Traffic Generator (running 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. Network connectivity between SUTs and to TG is provided using virtual NICs and VMs' virtio drivers.

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).

Nested-VM Service Switching

Nested-VM (Virtual Machine) service switching topology test cases require VPP DUT to communicate with nested-VM(s) over vhost-user virtual interfaces. Nested-VM(s) is(are) created on SUT1 and/or SUT2 for the duration of these particular test cases only. Virtual test topology with nested-VM(s) is shown in the figure below.



4.1.2 Functional Tests Coverage

CSIT-1908.1 includes following VPP functionality tested in virtual VM environment:

Functionality	Description
LISP	Locator/ID Separation Protocol overlay tunnels and locator/id mapping control.

4.1.3 Functional Tests Naming

CSIT-1908.1 follows a common structured naming convention for all performance and system functional tests, introduced in CSIT-17.01.

The naming should be intuitive for majority of the tests. Complete description of CSIT test naming convention is provided on *Test Naming* (page 459).

4.2 Release Notes

4.2.1 Changes in CSIT-1908.1

1. TEST FRAMEWORK
 - **Bug fixes.**
2. CSIT TEST MIGRATION
 - **VPP_Device:** Continuing migration of the original FD.io CSIT VIRL tests to VPP-device tests for functional acceptance of VPP feature path(s) driven by performance tests.

4.2.2 Known Issues

List of known issues in CSIT-1908.1 for VPP functional tests in VIRL:

#	JiraID	Issue Description

4.3 Test Environment

CSIT VPP functional tests are executed in FD.io VIRL testbeds. The physical VIRL testbed infrastructure consists of three VIRL servers:

- tb4-virl1:
 - Status: Production
 - OS: Ubuntu 16.04.2
 - VIRL STD server version: 0.10.32.16
 - VIRL UWM server version: 0.10.32.16
- tb4-virl2:
 - Status: Production
 - OS: Ubuntu 16.04.2
 - VIRL STD server version: 0.10.32.16
 - VIRL UWM server version: 0.10.32.16
- tb4-virl3:
 - Status: Production
 - OS: Ubuntu 16.04.2
 - VIRL STD server version: 0.10.32.19
 - VIRL UWM server version: 0.10.32.19
- VIRL hosts: Cisco UCS C240-M4, each with 2x Intel Xeon E5-2699 v3 (2.30 GHz, 18c), 512GB RAM.

Whenever a patch is submitted to gerrit for review, parallel VIRL simulations are started to reduce the time of execution of all functional tests. The number of parallel VIRL simulations is equal to a number of test groups defined by TEST_GROUPS variable in `csit/bootstrap.sh` file. VIRL host to run VIRL simulation is selected based on least load algorithm per VIRL simulation.

Every VIRL simulation uses the same three-node logical ring topology - Traffic Generator (TG node) and two Systems Under Test (SUT1 and SUT2). 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 VIRL simulations.

4.3.1 SUT Settings - VIRL Guest VM

SUT VMs' settings are defined in [VIRL topologies directory](#)¹³⁷

- 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.

¹³⁷ https://git.fd.io/csit/tree/resources/tools/virl/topologies/?h=rls1908_1

- Kernel Shared Memory Max: 2147483648 (vm.nr_hugepages * 2 * 1024 * 1024).

4.3.2 SUT Settings - VIRL Guest OS Linux

In CSIT terminology, the VM operating system for both SUTs that VPP-19.08.1 release has been tested with, is the following:

1. Ubuntu VIRL 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 [VIRL ubuntu images lists](#)¹³⁸.

A replica of this VM image can be built by running the `build.sh` script in CSIT repository.

2. CentOS VIRL 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 [VIRL centos images lists](#)¹³⁹.

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](#)¹⁴⁰, 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](#)¹⁴¹.

4.3.3 DUT Settings - 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](#)¹⁴².

Example of DUT nodes configuration:

```
DUT1:
  type: DUT
  host: "10.30.51.157"
```

(continues on next page)

¹³⁸ https://git.fd.io/csit/tree/resources/tools/disk-image-builder/ubuntu/lists/?h=rls1908_1

¹³⁹ https://git.fd.io/csit/tree/resources/tools/disk-image-builder/centos/lists/?h=rls1908_1

¹⁴⁰ <https://buildroot.org/>

¹⁴¹ https://git.fd.io/csit/tree/resources/tools/disk-image-builder/nested/?h=rls1908_1

¹⁴² https://git.fd.io/csit/tree/resources/tools/virl/topologies/?h=rls1908_1

(continued from previous page)

```

arch: x86_64
port: 22
username: cisco
honeycomb:
  user: admin
  passwd: admin
  port: 8183
  netconf_port: 2831
priv_key: |
-----BEGIN RSA PRIVATE KEY-----
MIIePgIBAAKCAQEAwUDlTpzSHpwLQotZOFs4AgcPNEwCnP1AB2hWfMvI+8Kah/gb
v8ruZU9RqhPs56tyKzxbhvNkY4VbH5F1GiLHZu3mLqz4KfghMmaeMEj01T7BYyd
vuBfTvIlu1jfQ2vAlnYrDwn+C1xJk81m0pDgvrLEX4qVvH2sGh7UEkYy5r82DNa2
4VjzPB1J/c8a9zP8FoZUhYIzF4FLvRMjUADpbMXgJMsGpaZLmz95ap0Eot7vb1Cc
1LvF97iyBCrtIOSKRKA50ZHLGjMKmOwnYU+cP5718tbproDVi6VJ0o7zeuXyetMs
8YB19kWB1WG9BqP9jctFvsmi5G7hXgq1Y8u+DwIDAQABAoIBAQC/W4E0DHjLMny7
0bvW2YkZD0Zw3fttdB94tkm4PdZv5MybooPnsAvLaXVV0hEdfVi5kzSWN1/LY/tN
EP1BgGphc2QgB59/PPxGwFIjDCvUzlsZpynBHe+B/qh5ExNQcVvsIOqWI7DX1XaN
0i/kh0zmJ6HncRRah1spKimYRsaUUDskyg7q3QqMwVaqBbbMvLs/w7Zwd/zoDqCU
MY/pCI6hkB3QbRo0odiZLohphB12ShABTwjvVyyKL5UA4jAeneJrhH5gWVLXnfgD
p62W5CollKEYb1C8mUkPxp7Qo277zw3xaq+oktIZhc5SUEUd7nJZtNqVAHqkItW
79VmpKyxAoGBAPfU+kqNPaTSvp+x1n5sn2SgipzDgtgi9QqNmC4cjtrQaaqI57SG
OHw1jX8i7L2G1WvVtkHg060n1EVo5n65fff0qeVBezLVJ7ghWI8U+oBiJJyQ4boD
GJVnso0SUQ0rtuGd9eVwfDk3o19aCN0KK53oPfIYli29pyu4l095kg11AoGBAMef
bPEMBI/2XmCPshLSwhGF1+dW8d+K1luj3CUQ/0vU1vma3dfBOYNsIwAgTP0iIUTg
8DYE6KBCdPtXAEI0YAEAKB9ry1tKR2NQEIPfs1YytKerTjwAiqSi0heM6+zwEzu
f54Z4oBhmsML0jXoMnu+NZZec6EUdQeY40+jhjzAoGBAIogC3dtjMPGKTP7+93u
UE/XIioI8fWg9fj3sMka4IMu+pVvRCRBAjRH7JrFLkjbUyuMqs3Arnk9K+gbdQt/
+m95Njtt6WoFXuPCwgbM3GidSmZwYT4454SfDzVBYScEDCNm1FuR+8ov9bFLDtGT
D4gsngnGJj1MDFXTxZEn4nzZAoGBAKCg4WmpUPaCuXibyB+rZavxwstNSn2lJ83/
sYJGBhf/raiv/FLDUcM1vYg5dZnu37RsB/5/vqx0LZGyYd7x+Jo5HkQGPnKgNwhn
g8BkdZIRF8uEJqx0o0ycd0U7n/2093swIpKWo5LIiRPuqqzj+uZKnAL7vuVdxfaY
qVz2daMPAoGBALgaakA3voU/H01PYLWIhFrBThyJ+BQSQ80qrEzC8AnegWFxRAM8
EqrzZX17ACUuo1dH0Eipm41j2+BZWlQjiUgq5uj8+zyz+EU1ZRRyJc0KzbdACeud
BpWWSXGIBI5G4CppeYLjMUHZpJYeX1USULJQd2c4crLJKb76E8gz3Z9kN
-----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"
arch: x86_64
port: 22
username: cisco
honeycomb:

```

(continues on next page)

(continued from previous page)

```

user: admin
passwd: admin
port: 8183
netconf_port: 2831
priv_key: |
-----BEGIN RSA PRIVATE KEY-----
MIIIEPgIBAAKCAQEAWUD1TpzSHpwLQotZ0FS4AgcPNEWCnP1AB2hWFmvI+8Kah/gb
v8ruZU9RqhPs56tyKzxbhvNkY4VbH5F1GiLHZu3mLqzM4KfghMmaeMEj01T7BYyd
vuBfTvIluljFQ2vAlnYrDwn+C1xJk81m0pDgvrLEX4qVvh2sGh7UEkYy5r82DNa2
4VjzPB1J/c8a9zP8FoZuHYIzF4FLvRMjUADpbMXgJMsGpaZLmz95ap0Eot7vb1Cc
1LvF97iyBCrtIOSKRKA50ZHLGjMKmOwnYU+cP5718tbproDVi6VJ0o7zeuXyetMs
8YB19kWb1WG9BqP9jctFvsmi5G7hXgq1Y8u+DwIDAQABAoIBAQC/W4E0DHjLMny7
0bvW2YkZD0Zw3fttdB94tkm4PdZv5MybooPnsAvLaXVV0hEdFvI5kzSWN1/LY/tN
EP1BgGphc2QgB59/PPxGwFIjDCvUz1sZpynBHe+B/qh5ExNqCvvsIQqWI7DX1XaN
0i/khOzmJ6HncRRah1spKimYRsaUUDskyg7q3QqMwVaqBbbMvLs/w7ZWd/zoDqCU
MY/pC16hkB3QbRo00diZLohphB12ShABTwjvVyyKL5UA4jAEneJrhH5gWVLXnfgD
p62W5Col1KEYb1C8mUkPxpP7Qo277zw3xaq+oktIZhc5SUEUd7nJZtNqVAHqkItW
79VmpKyxAoGBAPFw+kqNPaTSvp+x1n5sn2SgipzDtgi9QqNmC4cjtrQqaaqI57SG
OHw1jX8i7L2G1WvVtkHg060n1EVo5n65ffF0qeVBezLVJ7ghWI8U+oBiJJyQ4boD
GJVnsoS0UQ0rtuGd9eVwfdk3o19aCN0KK53oPFIY1i29pyu4l095kg11AoGBAMef
bPEMBI/2XmCPshLShwGF1+dW8d+K11uj3CUQ/0vU1vma3dfBOYNsIwAgTP0iIUTg
8DYE6KBcDpTxAUEI0YAEAKB9ry1tKR2NQEIPfslYytKerTwiAiqSi0heM6+zwEzu
f54Z4oBhsMSL0jXoMnu+NZzEc6EUdQeY40+jhjzAoGBAIOgC3dtjMPGKTP7+93u
UE/XIioI8fWg9fj3sMka4IMu+pVvRCRbAjRH7JrFLkjbUyuMqs3Arnk9K+gbdQt/
+m95Njtt6WoF XuPCwgbM3GidSmZwYT4454SfDzVBYScedCNm1FuR+8ov9bFLDtGT
D4gsngnGJj1MDFXTxZEn4nzZAoGBAKCg4WmpUPaCuXibyB+rZavxwsTNSn21J83/
sYJGBhf/raiv/FLDUcM1vYg5dZnu37RsB/5/vqxOLZGyYd7x+Jo5HkQGPNkGnwhn
g8BkdZIRF8uEJqx0o0ycd0U7n/2093swIpKwo5LIiRPuqqzj+uZKnAL7vuVdxfaY
qVz2daMPAoGBALgaaKa3voU/H01PYLWIhFrBThyJ+BQSQ80qrEzC8AnegWFxRAM8
EqrzX17ACUuo1dH0Eipm41j2+BZW1QjiUgq5uj8+zyz+EU1ZRRyJc0KzbdACeuD
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-19.08.1 release

VPP Installed Packages - Ubuntu

```

$ dpkg -l | grep vpp
ii libvppinfra                    19.08-release          amd64          𐀀
↳Vector Packet Processing--runtime libraries
ii libvppinfra-dev                19.08-release          amd64          𐀀
↳Vector Packet Processing--runtime libraries
ii python3-vpp-api                19.08-release          amd64          𐀀
↳VPP Python3 API bindings
ii vpp                             19.08-release          amd64          𐀀
↳Vector Packet Processing--executables
ii vpp-api-python                  19.08-release          amd64          𐀀
↳VPP Python API bindings
ii vpp-dbg                         19.08-release          amd64          𐀀
↳Vector Packet Processing--debug symbols
ii vpp-dev                         19.08-release          amd64          𐀀
↳Vector Packet Processing--development support
ii vpp-plugin-core                 19.08-release          amd64          𐀀
↳Vector Packet Processing--runtime core plugins
ii vpp-plugin-dpdk                 19.08-release          amd64          𐀀
↳Vector Packet Processing--runtime dpdk plugin

```

VPP Installed Packages - Centos

```

$ rpm -qai *vpp*
vpp-api-python
Version      : 19.08
Release      : release
Architecture: x86_64
Install Date: Wed 28 Aug 2019 01:35:57 AM EDT
Group        : Development/Libraries
Size         : 195318
License      : ASL 2.0
Signature    : (none)
Source RPM   : vpp-19.08-release.src.rpm
Build Date   : Wed 21 Aug 2019 02:30:08 PM EDT
Build Host   : 05ab23f629ba
Relocations  : (not relocatable)
Summary      : VPP api python bindings
Description  :
This package contains the python bindings for the vpp api
Name         : vpp
Version      : 19.08
Release      : release
Architecture: x86_64
Install Date: Wed 28 Aug 2019 01:35:57 AM EDT
Group        : Unspecified
Size         : 2681317
License      : ASL 2.0
Signature    : (none)
Source RPM   : vpp-19.08-release.src.rpm
Build Date   : Wed 21 Aug 2019 02:30:08 PM EDT
Build Host   : 05ab23f629ba
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      : 19.08

```

(continues on next page)

(continued from previous page)

```

Release      : release
Architecture: x86_64
Install Date: Wed 28 Aug 2019 01:35:57 AM EDT
Group       : System Environment/Libraries
Size        : 43222411
License     : ASL 2.0
Signature   : (none)
Source RPM  : vpp-19.08-release.src.rpm
Build Date  : Wed 21 Aug 2019 02:30:08 PM EDT
Build Host  : 05ab23f629ba
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-devel
Version     : 19.08
Release     : release
Architecture: x86_64
Install Date: Wed 28 Aug 2019 01:35:59 AM EDT
Group       : Development/Libraries
Size        : 13906969
License     : ASL 2.0
Signature   : (none)
Source RPM  : vpp-19.08-release.src.rpm
Build Date  : Wed 21 Aug 2019 02:30:08 PM EDT
Build Host  : 05ab23f629ba
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-selinux-policy
Version     : 19.08
Release     : release
Architecture: x86_64
Install Date: Wed 28 Aug 2019 01:35:56 AM EDT
Group       : System Environment/Base
Size        : 102929
License     : ASL 2.0
Signature   : (none)
Source RPM  : vpp-19.08-release.src.rpm
Build Date  : Wed 21 Aug 2019 02:30:08 PM EDT
Build Host  : 05ab23f629ba
Relocations : (not relocatable)
Summary     : VPP Security-Enhanced Linux (SELinux) policy
Description :
This package contains a tailored VPP SELinux policy
Name        : vpp-plugins
Version     : 19.08
Release     : release

```

(continues on next page)

(continued from previous page)

```
Architecture: x86_64
Install Date: Wed 28 Aug 2019 01:35:59 AM EDT
Group        : System Environment/Libraries
Size         : 34150451
License      : ASL 2.0
Signature    : (none)
Source RPM   : vpp-19.08-release.src.rpm
Build Date   : Wed 21 Aug 2019 02:30:08 PM EDT
Build Host   : 05ab23f629ba
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.

Common Configuration

There is used the default startup configuration as defined in [VPP startup.conf](#)¹⁴³

4.3.4 TG Settings - Scapy

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](#)¹⁴⁴.

`/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 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](#)¹⁴⁵.

Example of TG node configuration::

¹⁴³ https://git.fd.io/vpp/tree/src/vpp/conf/startup.conf?h=stable/1908_1&id=f4dcae4164f93dac80d4af19af0ee20e712ec673

¹⁴⁴ https://git.fd.io/csit/tree/resources/tools/virl/topologies/?h=rls1908_1

¹⁴⁵ https://git.fd.io/csit/tree/resources/tools/virl/topologies/?h=rls1908_1

TG:

```

type: TG
host: "10.30.51.155"
arch: x86_64
port: 22
username: cisco
priv_key: |
-----BEGIN RSA PRIVATE KEY-----
MIIEpgIBAAKCAQEAwUDITpzShpwLQotZOFs4AgcPNEwCnP1AB2hWFmvI+8Kah/gb
v8ruZU9RqhPs56tyKzxbhvNkY4VbH5F1GiLHZu3mLqzM4KfghMmaeMEj01T7BYyd
vuBfTviIlu1jfQ2vAlnYrDwn+C1xJk81m0pDgvrLEX4qVVh2sGh7UEkYy5r82DNa2
4VjzPB1J/c8a9zP8FoZuHYIZF4FLvRMjUADpbMXgJMsGpaZLmz95ap0Eot7vb1Cc
1LvF97iyBCrtIOSKRKA50ZHLGjMKmOwnYU+cP5718tbproDVi6VJ0o7zeuXyetMs
8YBl9kWB1WG9BqP9jctFvsmi5G7hXgq1Y8u+DwIDAQABAoIBAQC/W4E0DHjLMny7
0bvW2YKzD0Zw3fttdB94tkm4PdZv5MybooPnsAvLaXVV0hEdFVi5kzSWN1/LY/tN
EP1BgGphc2QgB59/PPxGwFIjDcVuz1sZpynBHe+B/qh5ExNQCvVsIOqWI7DX1XaN
0i/kh0zmJ6HncRRah1spKimYRsAUUDskyg7q3QqMwVaqBbbMvLs/w7ZWd/zoDqCU
MY/pCI6hkB3QbRo0odiZLohphB12ShABTwjvVyyKL5UA4jAeneJrhH5gWVLXnfgD
p62W5CollKEYb1C8mUkPxpP7Qo277zw3xaq+oktIZhc5SUEUd7nJZtNqVAHqkItW
79VmpKyxAoGBAPfU+kqNPaTSvp+x1n5sn2SgipzDtgi9QqNmC4cjrQQaaqI57SG
OHw1jX8i7L2G1WvVtkHg060n1EVo5n65fff0QeVBezLVJ7ghWI8U+oBiJjyQ4boD
GJVns0oS0UQ0rtuGd9eVwfdk3o19aCN0KK53oPfIYli29pyu4l095kg11AoGBAMef
bPEMBI/2XmCPshLShwGF1+dW8d+K1luj3CUQ/0vU1vma3dfBOYNsIwAgTP0iIUTg
8DYE6KBCdPtXAEI0YAEAKB9ry1tkR2NQEIPfs1YytKErtwjAiqSi0heM6+zwEzu
f54Z4oBhsMSL0jXoMnu+NZZEc6EUdQeY40+jhjzAoGBAIogC3dtjMPGKTP7+93u
UE/XIioI8fWg9fj3sMka4IMu+pVvRCRbAjRH7JrFLkjbUyuMqs3Arn9K+gbdQt/
+m95Njtt6WoFxpUCwgbM3GidSmZwYT4454SfDzVBYScedCNm1FuR+8ov9bFLDtGT
D4gsngnGjJ1MDFXTxZEn4nzZAoGBAKCg4WmpUPaCuXibyB+rZavxwsTNSn2lJ83/
sYJGBhf/raiV/FLDUcM1vYg5dZnu37RsB/5/vqxOLZGyYd7x+Jo5HkQGPnKgNwhn
g8BkdZIRF8uEJqx0o0ycd0U7n/2093swIpKWo5LIiRPuqqzj+uZKnAL7vuVdxfaY
qVz2daMPAoGBALgaakA3voU/H01PYLWIhFrBThyJ+BQSQ80qrEzC8AnegWFxRAM8
EqrzZX17ACUuo1dH0Eipm41j2+BZW1QjiUgq5uj8+zyz+EU1ZRRyJc0KzbDACeud
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. Scapy v2.3.1 is used for VPP-19.08.1 release tests.

4.4 Documentation

CSIT VPP Functional Tests Documentation¹⁴⁶ contains detailed functional description and input parameters for each test case.

¹⁴⁶ https://docs.fd.io/csit/rls1908_1/doc/tests.vpp.func.html

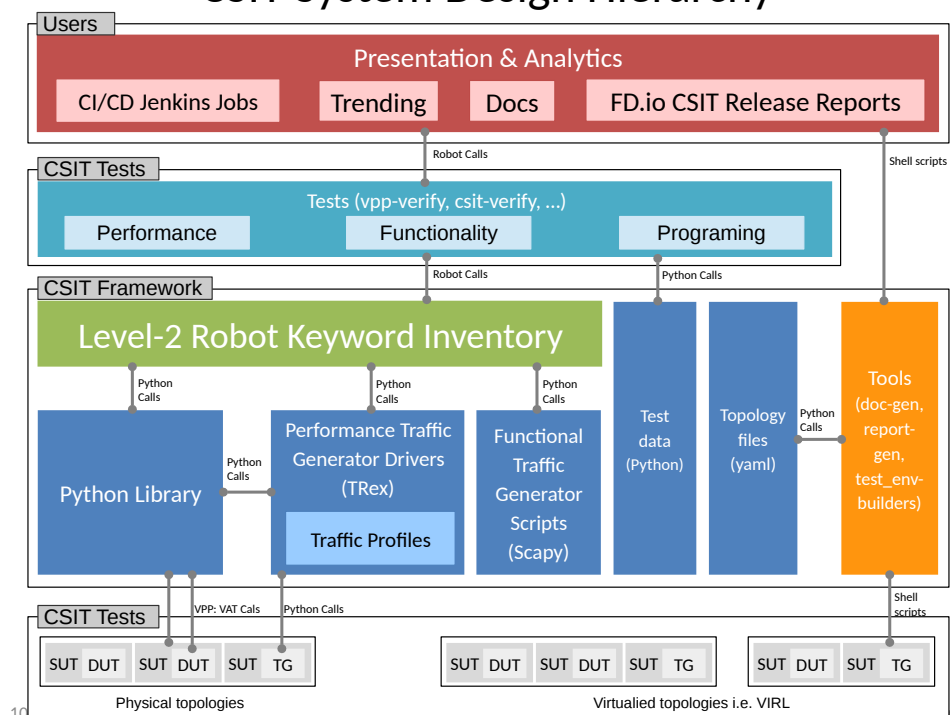
5.1 Design

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.

5.1.1 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.

CSIT System Design Hierarchy



A brief bottom-up description is provided here:

1. SUTs, DUTs, TGs

- SUTs - Systems Under Test;
- DUTs - Devices Under Test;
- TGs - Traffic Generators;

2. Level-1 libraries - Robot and Python

- Lowest level CSIT libraries abstracting underlying test environment, SUT, DUT and TG specifics;
- Used commonly across multiple L2 KWs;
- Performance and functional tests:
 - L1 KWs (KeyWords) are implemented as RF libraries and Python libraries;
- Performance TG L1 KWs:
 - All L1 KWs are implemented as Python libraries:
 - * Support for TRex only today;
 - * CSIT IXIA drivers in progress;
- Performance data plane traffic profiles:
 - TG-specific stream profiles provide full control of:
 - * Packet definition – layers, MACs, IPs, ports, combinations thereof e.g. IPs and UDP ports;
 - * Stream definitions - different streams can run together, delayed, one after each other;
 - * Stream profiles are independent of CSIT framework and can be used in any T-rex setup, can be sent anywhere to repeat tests with exactly the same setup;

- * 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;
 - Performance tests using physical testbed environment:
 - VPP;
 - DPDK-Testpmd;
 - DPDK-L3Fwd;
 - Tools:
 - Documentation generator;
 - Report generator;
 - Testbed environment setup ansible playbooks;
 - Operational debugging scripts;

5.1.2 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;

5.1.3 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;

5.1.4 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.

5.1.5 Coding Guidelines

Coding guidelines can be found on [Design optimizations wiki page](#)¹⁴⁷.

¹⁴⁷ https://wiki.fd.io/view/CSIT/Design_Optimizations

5.2 Test Naming

5.2.1 Background

CSIT-1908.1 follows a common structured naming convention for all performance and system functional tests, introduced in CSIT-1701.

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](#)¹⁴⁸. Below few illustrative examples of the naming usage for test suites across CSIT performance, functional and Honeycomb management test areas.

5.2.2 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.

5.2.3 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.*

¹⁴⁸ <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](#)¹⁴⁹.


¹⁴⁹ <https://wiki.fd.io/view/CSIT/csit-test-naming>

5.3 Presentation and Analytics

5.3.1 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`¹⁵⁰.
 - 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)



¹⁵⁰ <https://pandas.pydata.org/pandas-docs/stable/generated/pandas.Series.html>

5.3.2 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|l2bdbasemaclrn)-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

```

(continues on next page)

(continued from previous page)

```
- 15
- 16
- 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-vir1:
- lastSuccessfulBuild
hc2vpp-csit-perf-master-ubuntu1604:
- 8
- 9
hc2vpp-csit-integration-1707-ubuntu1604:
- lastSuccessfulBuild
csit-nsh_sfc-verify-func-1707-ubuntu1604-vir1:
- 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]"

```

(continues on next page)

(continued from previous page)

```

data: "template 3"
-
title: "VPP-17.01 mean [Mpps]"
data: "template 4"
-
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"
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

```

(continues on next page)

(continued from previous page)

```

- 14
- 15
- 16
- 17
- 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|l2bdbasemac|l2rn)-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"
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|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: ""
    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

```

(continues on next page)

(continued from previous page)

```

legend:
  orientation: "h"
  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.

5.3.3 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.

5.3.4 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.ly¹⁵¹ 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.

5.3.5 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.

¹⁵¹ <https://plot.ly/>

5.3.6 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:

General section - parameters common to all plots:

- type: "cpta";
- title: The title of this section;
- output-file-type: only ".html" is supported;
- output-file: path where the generated files will be stored.

Plots section:

- plot title;
- output file name;

- input data for plots;
 - 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;
- tests to be displayed in the plot defined by a filter;
- list of parameters to extract from the data;
- plot layout

Example:

```

-
type: "cpta"
title: "Continuous Performance Trending and Analysis"
output-file-type: ".html"
output-file: "{DIR[STATIC,VPP]}/cpta"
plots:
  - title: "VPP 1T1C L2 64B Packet Throughput - Trending"
    output-file-name: "l2-1t1c-x520"
    data: "plot-performance-trending-vpp"
    filter: "'NIC_Intel-X520-DA2' and 'MRR' and '64B' and ('BASE' or 'SCALE') and '1T1C' and (
↪ 'L2BDMACSTAT' or 'L2BDMACLRN' or 'L2XC fwd') and not 'VHOST' and not 'MEMIF'"
    parameters:
      - "result"
    layout: "plot-cpta-vpp"
  - title: "DPDK 4T4C IMIX MRR Trending"
    output-file-name: "dppk-imix-4t4c-xl710"
    data: "plot-performance-trending-dppk"
    filter: "'NIC_Intel-XL710' and 'IMIX' and 'MRR' and '4T4C' and 'DPDK'"
    parameters:
      - "result"
    layout: "plot-cpta-dppk"

```

The Dashboard

Performance dashboard tables provide the latest VPP throughput trend, trend compliance and detected anomalies, all on a per VPP test case basis. The Dashboard is generated as three tables for 1t1c, 2t2c and 4t4c MRR tests.

At first, the .csv tables are generated (only the table for 1t1c is shown):

```

-
type: "table"
title: "Performance trending dashboard"
algorithm: "table_performance_trending_dashboard"
output-file-ext: ".csv"
output-file: "{DIR[STATIC,VPP]}/performance-trending-dashboard-1t1c"
data: "plot-performance-trending-all"
filter: "'MRR' and '1T1C'"
parameters:
  - "name"
  - "parent"
  - "result"
ignore-list:
  - "tests.vpp.perf.l2.10ge2p1x520-eth-l2bdscale1mmaclrn-mrr.tc01-64b-1t1c-eth-l2bdscale1mmaclrn-
↪ ndrdisc"
outlier-const: 1.5

```

(continues on next page)

(continued from previous page)

```
window: 14
evaluated-window: 14
long-trend-window: 180
```

Then, html tables stored inside .rst files are generated:

```
-
type: "table"
title: "HTML performance trending dashboard 1t1c"
algorithm: "table_performance_trending_dashboard_html"
input-file: "{DIR[STATIC,VPP]}/performance-trending-dashboard-1t1c.csv"
output-file: "{DIR[STATIC,VPP]}/performance-trending-dashboard-1t1c.rst"
```

5.3.7 Root Cause Analysis

Root Cause Analysis (RCA) by analysing archived performance results – re-analyse available data for specified:

- range of jobs builds,
- set of specific tests and
- PASS/FAIL criteria to detect performance change.

In addition, PAL generates trending plots to show performance over the specified time interval.

Root Cause Analysis - Option 1: Analysing Archived VPP Results

It can be used to speed-up the process, or when the existing data is sufficient. In this case, PAL uses existing data saved in Nexus, searches for performance degradations and generates plots to show performance over the specified time interval for the selected tests.

Execution Sequence

1. Download and parse archived historical data and the new data.
2. Calculate trend metrics.
3. Find regression / progression.
4. Generate and publish results:
 - (a) Summary graphs to include measured values with Progression and Regression markers.
 - (b) List the DUT build(s) where the anomalies were detected.

CSIT PAL Specification

- What to test:
 - first build (Good); specified by the Jenkins job name and the build number
 - last build (Bad); specified by the Jenkins job name and the build number
 - step (1..n).
- Data:
 - tests of interest; list of tests (full name is used) which results are used

Example:

TODO

5.3.8 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
        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
```

(continues on next page)

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 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

PAL functional diagram



How to add an element

Element can be added by adding it's model to the specification file. If the element is to be generated by an existing algorithm, only it's parameters must be set.

If a brand new type of element needs to be added, also the algorithm must be implemented. Element generation algorithms are implemented in the files with names starting with "generator" prefix. The name of the function implementing the algorithm and the name of algorithm in the specification file have to be the same.

5.4 CSIT RF 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](#)¹⁵².

5.4.1 Testbed Topology Tags

2_NODE_DOUBLE_LINK_TOPO

2 nodes connected in a circular topology with two links interconnecting the devices.

2_NODE_SINGLE_LINK_TOPO

2 nodes connected in a circular topology with at least one link interconnecting devices.

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.

5.4.2 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).

5.4.3 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.

¹⁵² https://git.fd.io/csit/tree/docs/tag_documentation.rst?h=rls1908_1

5.4.4 NIC Model Tags

NIC_Intel-X520-DA2

Intel X520-DA2 NIC.

NIC_Intel-XL710

Intel XL710 NIC.

NIC_Intel-X710

Intel X710 NIC.

NIC_Intel-XXV710

Intel XXV710 NIC.

NIC_Cisco-VIC-1227

VIC-1227 by Cisco.

NIC_Cisco-VIC-1385

VIC-1385 by Cisco.

5.4.5 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

L2BD_1

Test with 1 L2 bridge domain.

L2BD_10

Test with 10 L2 bridge domains.

L2BD_100

Test with 100 L2 bridge domains.

L2BD_1K

Test with 1000 L2 bridge domains.

VLAN_1

Test with 1 VLAN sub-interface.

VLAN_10

Test with 10 VLAN sub-interfaces.

VLAN_100

Test with 100 VLAN sub-interfaces.

VLAN_1K

Test with 1000 VLAN sub-interfaces.

VXLAN_1

Test with 1 VXLAN tunnel.

VXLAN_10

Test with 10 VXLAN tunnels.

VXLAN_100

Test with 100 VXLAN tunnels.

VXLAN_1K

Test with 1000 VXLAN tunnels.

TNL_{t}

IPSec in tunnel mode - {t} tunnels.

SRC_USER_1

Traffic flow with 1 unique IP (users) in one direction.

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.

5.4.6 Test Category Tags

FUNCTEST

All functional test cases.

PERFTEST

All performance test cases.

5.4.7 Performance Type Tags

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. The results come from the final trials, which have duration of 30 seconds.

MRR

Performance tests where TG sends the traffic at maximum rate (line rate) and reports total sent/received packets over trial duration. The result is an average of 10 trials of 1 second duration.

SOAK

Performance tests using PLRsearch to find the critical load.

RECONF

Performance tests aimed to measure lost packets (time) when performing reconfiguration while full throughput offered load is applied.

5.4.8 Ethernet Frame Size Tags

These are describing the traffic offered by Traffic Generator, “primary” traffic in case of asymmetric load. For traffic between DUTs, or for “secondary” traffic, see $\{\text{overhead}\}$ value.

64B

64B frames used for test. Generic ethernet or IPv4.

78B

78B frames used for test. Ipv6.

114B

114B frames used for test. IPv4+vxlan.

118B

118B frames used for test. Dot1q+IPv4+vxlan.

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.

5.4.9 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.

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.

L2PATCH

L2PATCH 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.

DMM

Functional test cases for DMM.

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.

NF_DENSITY

Performance tests that measure throughput of multiple VNF and CNF service topologies at different service densities.

5.4.10 NF Service Density Tags

CHAIN

NF service density tests with VNF or CNF service chain topology(ies).

PIPE

NF service density tests with CNF service pipeline topology(ies).

NF_L3FWDIP4

NF service density tests with DPDK l3fwd IPv4 routing as NF workload.

NF_VPPIP4

NF service density tests with VPP IPv4 routing as NF workload.

{r}R{c}C

Service density matrix locator {r}R{c}C, {r}Row denoting number of service instances, {c}Column denoting number of NFs per service instance. {r}=(1,2,4,6,8,10), {c}=(1,2,4,6,8,10).

{n}VM{t}T

Service density {n}VM{t}T, {n}Number of NF Qemu VMs, {t}Number of threads per NF.

{n}DCR{t}T

Service density {n}DCR{t}T, {n}Number of NF Docker containers, {t}Number of threads per NF.

{n}_ADDED_CHAINS

{n}Number of chains (or pipelines) added (and/or removed) during RECONF test.

5.4.11 Forwarding Mode Tags

L2BDMACSTAT

VPP L2 bridge-domain, L2 MAC static.

L2BDMACLR

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.

5.4.12 Underlay Tags**IP4UNRLAY**

IPv4 underlay.

IP6UNRLAY

IPv6 underlay.

MPLSUNRLAY

MPLS underlay.

5.4.13 Overlay Tags**L2OVLAY**

L2 overlay.

IP4OVLAY

IPv4 overlay (IPv4 payload).

IP6OVLAY

IPv6 overlay (IPv6 payload).

5.4.14 Tagging Tags**DOT1Q**

All test cases with dot1q.

DOT1AD

All test cases with dot1ad.

5.4.15 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.

SRv6_1SID

All SRv6 test cases with single SID.

SRv6_2SID_DECAP

All SRv6 test cases with two SIDs and with decapsulation.

SRv6_2SID_NODECAP

All SRv6 test cases with two SIDs and without decapsulation.

5.4.16 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.

SINGLE_MEMIF

All test cases which uses only single Memif connection per DUT. One DUT instance is running in container having one physical interface exposed to container.

LBOND

All test cases which uses link bonding (BondEthernet interface).

LBOND_DPK

All test cases which uses DPK link bonding.

LBOND_VPP

All test cases which uses VPP link bonding.

LBOND_MODE_XOR

All test cases which uses link bonding with mode XOR.

LBOND_MODE_LACP

All test cases which uses link bonding with mode LACP.

LBOND_LB_L34

All test cases which uses link bonding with load-balance mode l34.

LBOND_1L

All test cases which uses one link for link bonding.

LBOND_2L

All test cases which uses two links for link bonding.

DRV_AVF

All test cases which uses Intel Adaptive Virtual Function (AVF) device plugin for VPP. This plugins provides native device support for Intel AVF. AVF is driver specification for current and future Intel Virtual Function devices. In essence, today this driver can be used only with Intel XL710 / X710 / XXV710 adapters.

5.4.17 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.

5.4.18 Encryption Tags**IPSECSW**

Crypto in software.

IPSECHW

Crypto in hardware.

IPSECTRAN

IPSec in transport mode.

IPSECTUN

IPSec in tunnel mode.

IPSECINT

IPSec in interface mode.

AES

IPSec using AES algorithms.

AES_128_CBC

IPSec using AES 128 CBC algorithms.

AES_128_GCM

IPSec using AES 128 GCM algorithms.

AES_256_GCM

IPSec using AES 256 GCM algorithms.

HMAC

IPSec using HMAC integrity algorithms.

HMAC_SHA_256

IPSec using HMAC SHA 256 integrity algorithms.

HMAC_SHA_512

IPSec using HMAC SHA 512 integrity algorithms.

5.4.19 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.

DRC

All test cases which use at least one Docker container.

DOCKER

All test cases which use Docker as container manager.

APP

All test cases with specific APP use.

5.4.20 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.

5.4.21 Multi-Threading Tags**STHREAD**

Dynamic tag. All test cases using single poll mode thread.

MTHREAD

Dynamic tag. 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.

1C

1 worker thread pinned to 1 dedicated physical core; or if HyperThreading is enabled, 2 worker threads each pinned to a separate logical core within 1 dedicated physical core. Main thread pinned to core 1.

2C

2 worker threads pinned to 2 dedicated physical cores; or if HyperThreading is enabled, 4 worker threads each pinned to a separate logical core within 2 dedicated physical cores. Main thread pinned to core 1.

4C

4 worker threads pinned to 4 dedicated physical cores; or if HyperThreading is enabled, 8 worker threads each pinned to a separate logical core within 4 dedicated physical cores. Main thread pinned to core 1.

1T1C

Dynamic tag. 1 worker thread pinned to 1 dedicated physical core. 1 receive queue per interface. Main thread pinned to core 1.

2T2C

Dynamic tag. 2 worker threads pinned to 2 dedicated physical cores. 1 receive queue per interface. Main thread pinned to core 1.

4T4C

Dynamic tag. 4 worker threads pinned to 4 dedicated physical cores. 2 receive queues per interface. Main thread pinned to core 1.

2T1C

Dynamic tag. 2 worker threads each pinned to a separate logical core within 1 dedicated physical core. 1 receive queue per interface. Main thread pinned to core 1.

4T2C

Dynamic tag. 4 worker threads each pinned to a separate logical core within 2 dedicated physical cores. 2 receive queues per interface. Main thread pinned to core 1.

8T4C

Dynamic tag. 8 worker threads each pinned to a separate logical core within 4 dedicated physical cores. 4 receive queues per interface. Main thread pinned to core 1.

5.4.22 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](#)¹⁰⁶
- [lxcnamespace] [Resource management: Linux kernel Namespaces and cgroups](#)¹⁰⁷.
- [stgraber] [LXC 1.0: Blog post series](#)¹⁰⁸.
- [lxcsecurity] [Linux Containers Security](#)¹⁰⁹.
- [capabilities] [Linux manual - capabilities - overview of Linux capabilities](#)¹¹⁰.
- [cgroup1] [Linux kernel documentation: cgroups](#)¹¹¹.
- [cgroup2] [Linux kernel documentation: Control Group v2](#)¹¹².
- [selinux] [SELinux Project Wiki](#)¹¹³.
- [lxcsecfeatures] [LXC 1.0: Security features](#)¹¹⁴.
- [lxcsource] [Linux Containers source](#)¹¹⁵.
- [apparmor] [Ubuntu AppArmor](#)¹¹⁶.
- [seccomp] [SECure COMputing with filters](#)¹¹⁷.
- [docker] [Docker](#)¹¹⁸.
- [k8sdoc] [Kubernetes documentation](#)¹¹⁹.
- [ligato] [Ligato](#)¹²⁰.
- [govpp] [FD.io goVPP project](#)¹²¹.
- [vppagent] [Ligato vpp-agent](#)¹²².

¹⁰⁶ <https://linuxcontainers.org/>

¹⁰⁷ <https://www.cs.ucsb.edu/~rich/class/cs293b-cloud/papers/lxc-namespace.pdf>

¹⁰⁸ <https://stgraber.org/2013/12/20/lxc-1-0-blog-post-series/>

¹⁰⁹ <https://linuxcontainers.org/lxc/security/>

¹¹⁰ <http://man7.org/linux/man-pages/man7/capabilities.7.html>

¹¹¹ <https://www.kernel.org/doc/Documentation/cgroup-v1/cgroups.txt>

¹¹² <https://www.kernel.org/doc/Documentation/cgroup-v2.txt>

¹¹³ http://selinuxproject.org/page/Main_Page

¹¹⁴ <https://stgraber.org/2014/01/01/lxc-1-0-security-features/>

¹¹⁵ <https://github.com/lxc/lxc>

¹¹⁶ <https://wiki.ubuntu.com/AppArmor>

¹¹⁷ https://www.kernel.org/doc/Documentation/prctl/seccomp_filter.txt

¹¹⁸ <https://www.docker.com/what-docker>

¹¹⁹ <https://kubernetes.io/docs/home/>

¹²⁰ <https://github.com/ligato>

¹²¹ <https://wiki.fd.io/view/GoVPP>

¹²² <https://github.com/ligato/vpp-agent>

[imagevar] Image parameter is required in initial commit version. There is plan to implement container build class to build Docker/LXC image.

[TWSLink] [TWS](#)¹²⁶

[dockerhub] [Docker hub](#)¹²⁷

[fdiocsitgerrit] [FD.io/CSIT gerrit](#)¹²⁸

[fdioregistry] [FD.io registry](#)

[JenkinsSlaveDcrFile] [jenkins-slave-dcr-file](#)¹²⁹

[CsitShimDcrFile] [csit-shim-dcr-file](#)¹³⁰

[CsitSutDcrFile] [csit-sut-dcr-file](#)¹³¹

[ansiblelink] [ansible](#)¹³²

[fdiocsitansible] [Fd.io/CSIT ansible](#)¹³³

[inteli40e] [Intel i40e](#)¹³⁴

[pciids] [pci ids](#)¹³⁵

¹²⁶ <https://wiki.fd.io/view/CSIT/TWS>

¹²⁷ <https://hub.docker.com/>

¹²⁸ <https://gerrit.fd.io/r/CSIT>

¹²⁹ <https://github.com/snergfdio/multivppcache/blob/master/ubuntu18/Dockerfile>

¹³⁰ <https://github.com/snergfdio/multivppcache/blob/master/csit-shim/Dockerfile>

¹³¹ <https://github.com/snergfdio/multivppcache/blob/master/csit-sut/Dockerfile>

¹³² <https://www.ansible.com/>

¹³³ <https://git.fd.io/csit/tree/resources/tools/testbed-setup/ansible>

¹³⁴ <https://downloadmirror.intel.com/26370/eng/readme.txt>

¹³⁵ <http://pci-ids.ucw.cz/v2.2/pci.ids>