/TRP1«/ca 1/CA 1»

CSIT REPORT

Release rls2206

Aug 03, 2022

CONTENTS

CHAPTER

ONE

INTRODUCTION

1.1 Report History

FD.io CSIT-2206 Report history and per .[ww] revision changes are listed below.

.[ww] Revision	Changes
.31	1. Added packet throughput, speedup multi-core and packet latency graphs for 2n-tx2 LXC/DRC Container Memif.
.30	1. Added VPP iterative data from 2n-clx, 2n-icx, 2n-skx - 100B-ethip4tcp-ip4base-tput - 100B-ethip4tcp-nat44ed-tput - 100B-ethip4udp-ip4base-tput
.29	1. Added VPP iterative data: • 2n-zn2 • 3n-alt 2. Added VPP coverage data: • 2n-zn2 • 3n-alt 3. Added RCAs, see *Root Cause Analysis for Performance Changes* (page 83).
.28	Initial revision

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-2206 report contains system performance and functional testing data of VPP-22.06 release. PDF version of this report[1] is available for download.

CSIT-2206 report is structured as follows:

1. INTRODUCTION: General introduction to FD.io CSIT-2206.
    - **Introduction**: This section.
    - **Test Scenarios Overview**: A brief overview of test scenarios covered in this report.

[1] https://s3-docs.fd.io/csit/rls2206/report/_static/archive/csit_rls2206.31.pdf

- **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-2206, 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 repeatibility of measurements.
   - **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-2206, known issues.
   - **Packet Throughput**: NDR, PDR throughput graphs based on results from repeated same test job executions to verify repeatibility 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. TREX PERFORMANCE: TREX performance tests executed in physical FD.io testbeds.
   - **Overview**: Tested logical topologies, test coverage.
   - **Release Notes**: Changes in CSIT-2206, known issues.
   - **Packet Throughput**: NDR, PDR throughput graphs based on results from repeated same test job executions to verify repeatibility of measurements.
   - **Packet Latency**: Latency graphs based on results from test job executions.
   - **Throughput Trending**: References to regular TREX performance trending.
   - **Test Environment**: Performance test environment configuration.
5. 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-2206, 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.

6. 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.
7. 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.
8. 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.
9. 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-2206 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 Cascadelake and Skylake servers, ARM, Denverton, range of NICs (10GE, 25GE, 40GE, 100GE) 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. **T-Rex Performance**: T-Rex perfomance tests are executed in physical FD.io testbeds, focusing on T-Rex data plane performance in NIC-to-NIC loopback topologies. Tested across Intel Skylake servers, range of NICs (10GE) and selected traffic profiles. TRex is used as a traffic generator.
6. **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-22.06 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 Python. CSIT-2206 source code for the executed test suites is available in CSIT branch rls2206 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 Performance Physical Testbeds

All FD.io (Fast Data Input/Ouput) 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 a 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, Cascadelake 6252N, Icelake 8358.
- Intel Atom: Denverton C3858.
- Arm: TaiShan 2280, hip07-d05.
- AMD EPYC: Zen2 7532.

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.

1.4.1 Physical NICs and Drivers

SUT and TG servers are equipped with a number of different NIC models.

VPP is performance tested on SUTs with the following NICs and drivers:

1. 2p10GE: x550, x553 Intel (codename Niantic) - DPDK Poll Mode Driver (PMD).
2. 4p10GE: x710-DA4 Intel (codename Fortville, FVL) - DPDK PMD. - AVF in PMD mode. - AF_XDP in PMD mode.
3. 2p25GE: xxv710-DA2 Intel (codename Fortville, FVL) - DPDK PMD. - AVF in PMD mode. - AF_XDP in PMD mode.
4. 2p100GE: cx556a-edat Mellanox ConnectX5 - RDMA_core in PMD mode.
5. 2p100GE: E810-2CQDA2 Intel (codename Columbiaville, CVL) - DPDK PMD. - AVF in PMD mode.

DPDK applications, testpmd and l3fwd, are performance tested on the same SUTs exclusively with DPDK drivers for all NICs.

TRex running on TGs is using DPDK drivers for all NICs.

VPP hoststack tests utilize ab (Apache HTTP server benchmarking tool) running on TGs and using Linux drivers for all NICs.

For more information see *Test Environment* (page 1453) and *Test Environment* (page 1649).

1.4.2 2-Node AMD EPYC Zen2 (2n-zn2)

One 2n-zn2 testbed in in operation in FD.io labs. It is built based on two SuperMicro SuperMicro AS-1114S-WTRT servers, with SUT and TG servers equipped with one AMD EPYC Zen2 7532 processor each (256 MB Cache, 2.40 GHz, 32 cores). 2n-zn2 physical topology is shown below.

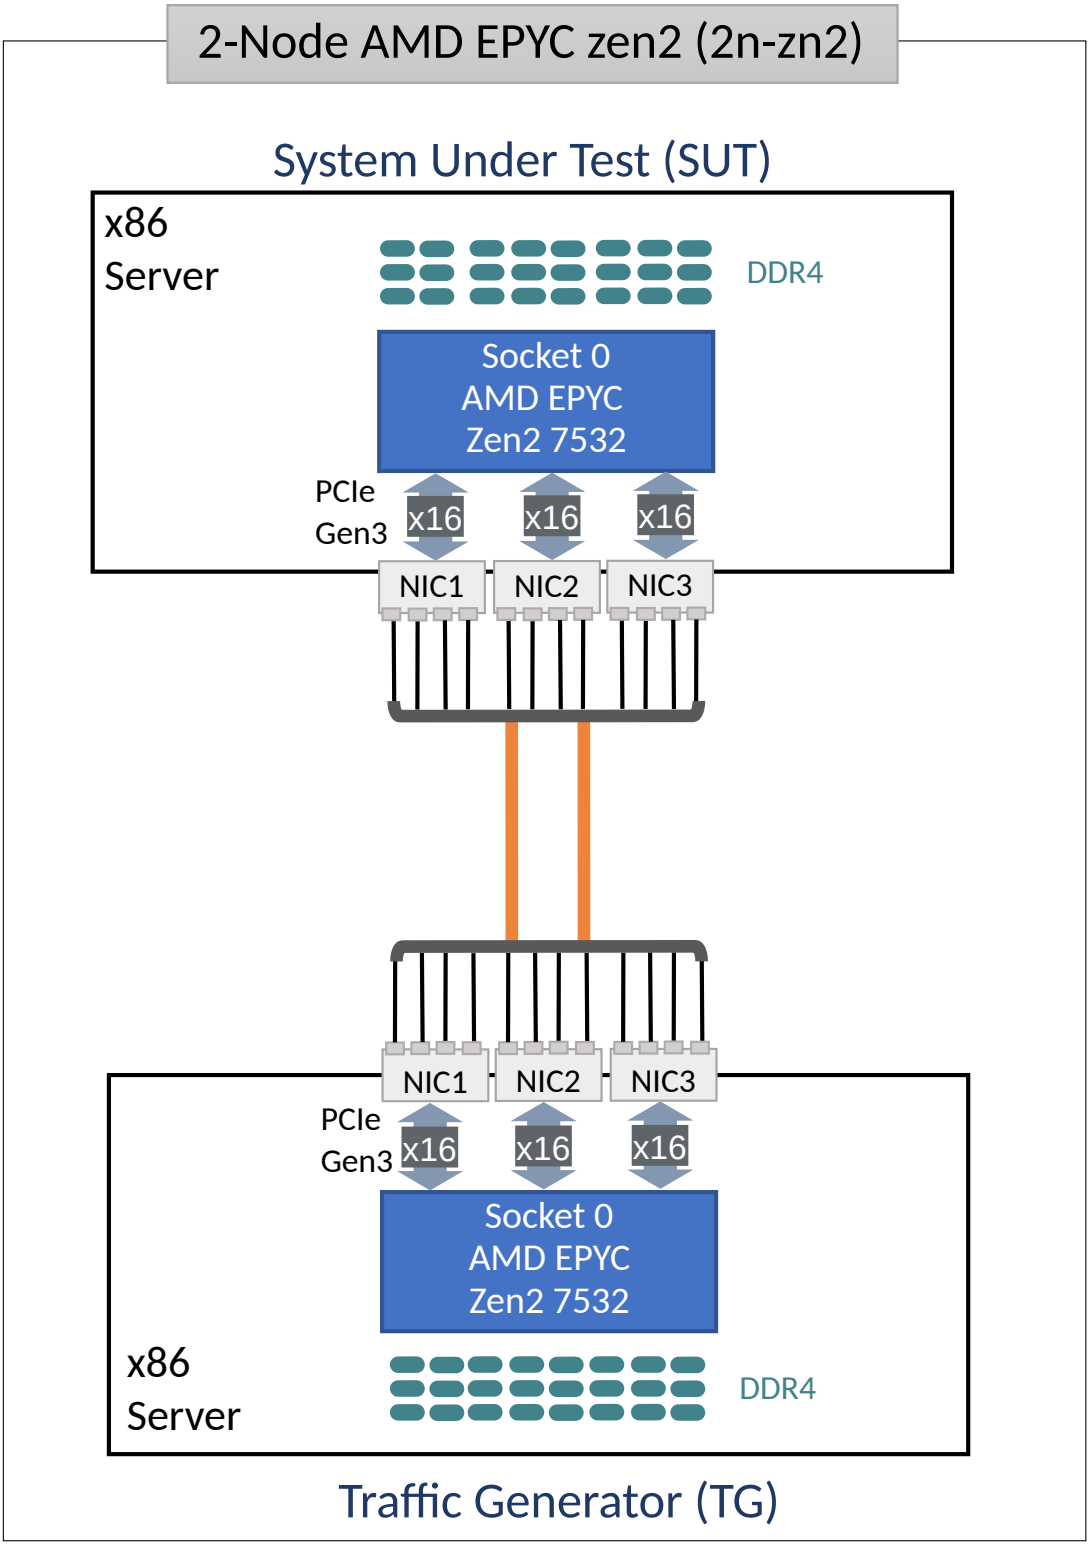

SUT NICs:

1. NIC-1: x710-DA4 4p10GE Intel.
2. NIC-2: xxv710-DA2 2p25GE Intel.
3. NIC-3: cx556a-edat ConnectX5 2p100GE Mellanox.

TG NICs:

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

2. NIC-2: xxv710-DA2 2p25GE Intel.
3. NIC-3: cx556a-edat ConnectX5 2p100GE Mellanox.

All AMD EPYC Zen2 7532 servers run with AMD SMT enabled, doubling the number of logical cores exposed to Linux.

1.4.3 2-Node Xeon Cascadelake (2n-clx)

Three 2n-clx testbeds are in operation in FD.io labs. Each 2n-clx testbed is built with two SuperMicro SYS-7049GP-TRT servers, SUTs are equipped with two Intel Xeon Gold 6252N processors (35.75 MB Cache, 2.30 GHz, 24 cores). TGs are equiped with Intel Xeon Cascade Lake Platinum 8280 processors (38.5 MB Cache, 2.70 GHz, 28 cores). 2n-clx physical topology is shown below.

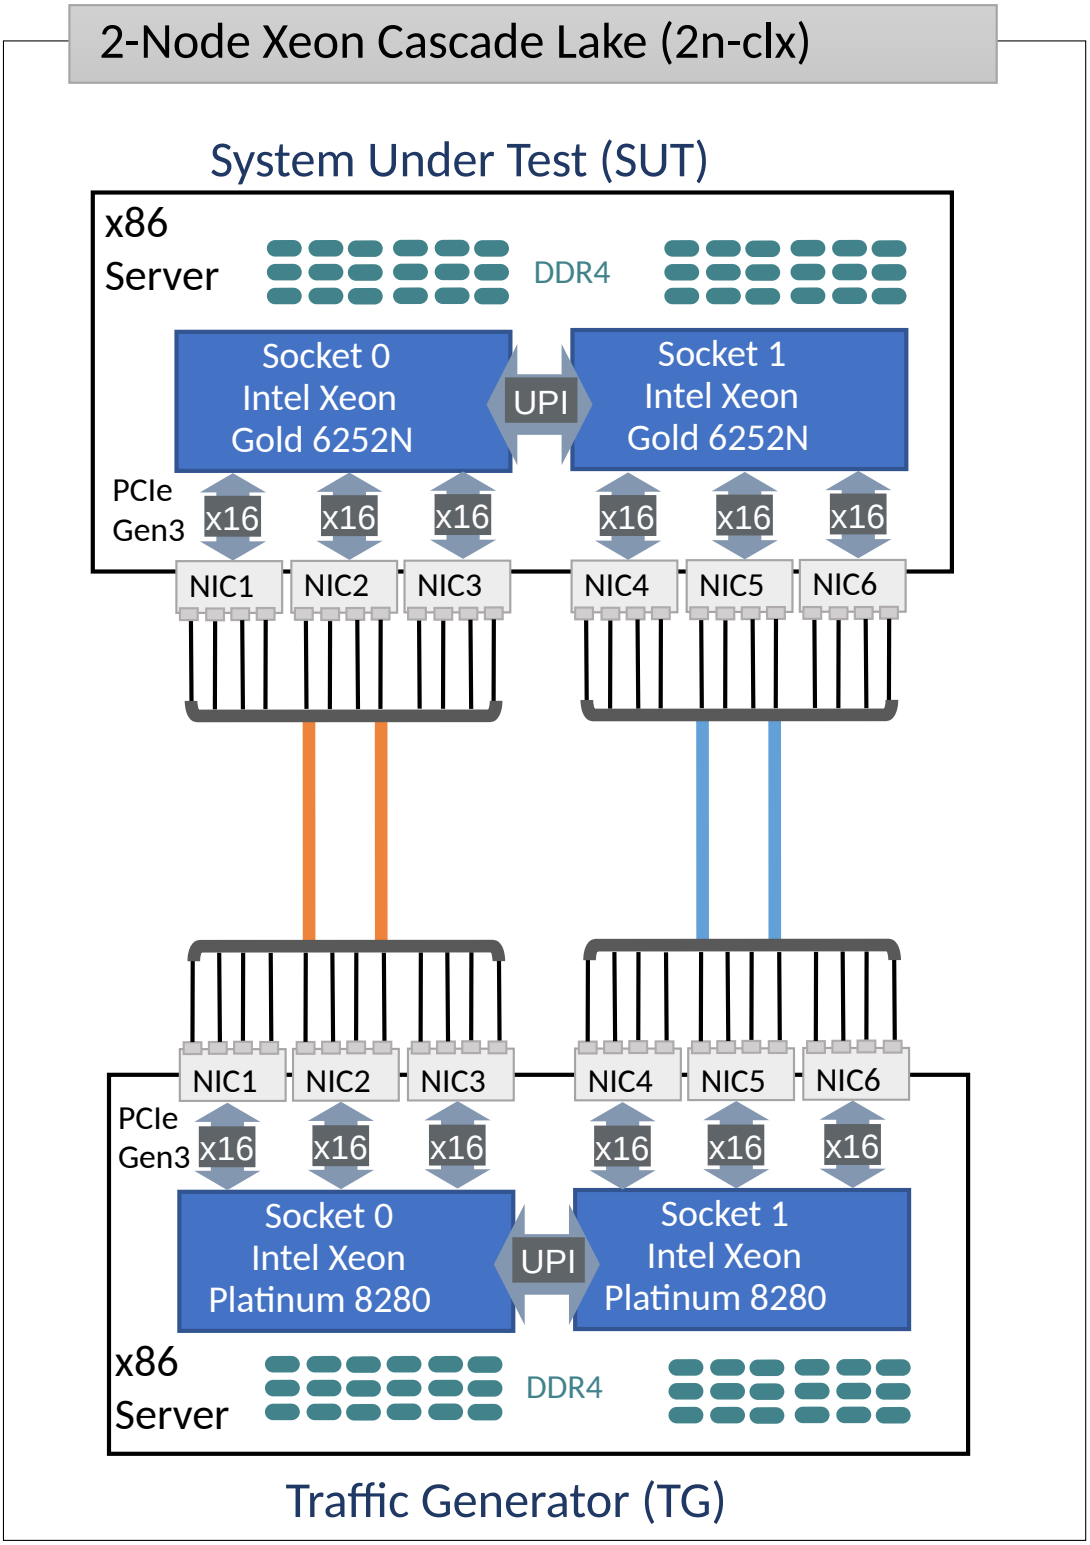

SUT NICs:

1. NIC-1: x710-DA4 4p10GE Intel.
2. NIC-2: xxv710-DA2 2p25GE Intel.
3. NIC-3: cx556a-edat ConnectX5 2p100GE Mellanox.
4. NIC-4: empty, future expansion.
5. NIC-5: empty, future expansion.

6. NIC-6: empty, future expansion.

TG NICs:

1. NIC-1: x710-DA4 4p10GE Intel.
2. NIC-2: xxv710-DA2 2p25GE Intel.
3. NIC-3: cx556a-edat ConnectX5 2p100GE Mellanox.
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 Cascadelake servers run with Intel Hyper-Threading enabled, doubling the number of logical cores exposed to Linux.

1.4.4 2-Node Xeon Icelake (2n-icx)

One 2n-icx testbed is in operation in FD.io labs. It is built with two SuperMicro SYS-740GP-TNRT servers, each in turn equipped with two Intel Xeon Platinum 8358 processors (48 MB Cache, 2.60 GHz, 32 cores).

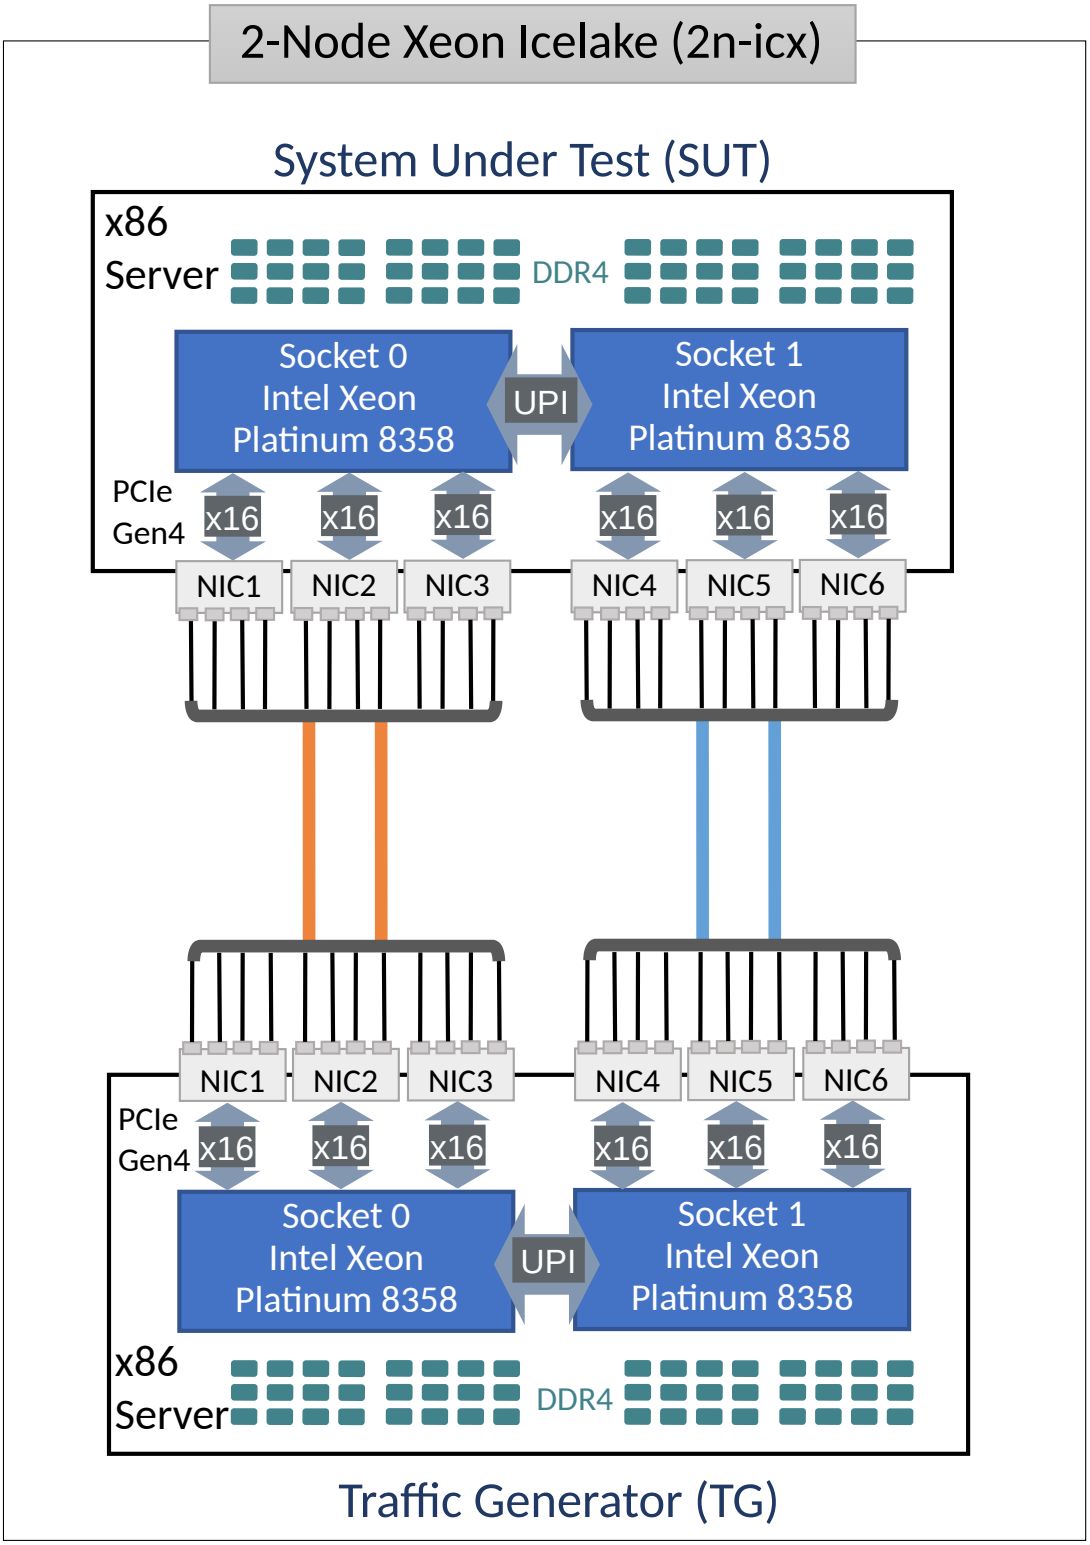

SUT and TG NICs:

1. NIC-1: xxv710-DA2 2p25GE Intel.
2. NIC-2: E810-2CQDA2 2p100GbE Intel (* to be added).
3. NIC-3: E810-CQDA4 4p100GbE Intel (* to be added).

All Intel Xeon Icelake servers run with Intel Hyper-Threading enabled, doubling the number of logical cores exposed to Linux.

1.4.5 3-Node Xeon Icelake (3n-icx)

One 3n-icx testbed is in operation in FD.io labs. It is built with three SuperMicro SYS-740GP-TNRT servers, each in turn equipped with two Intel Xeon Platinum 8358 processors (48 MB Cache, 2.60 GHz, 32 cores).

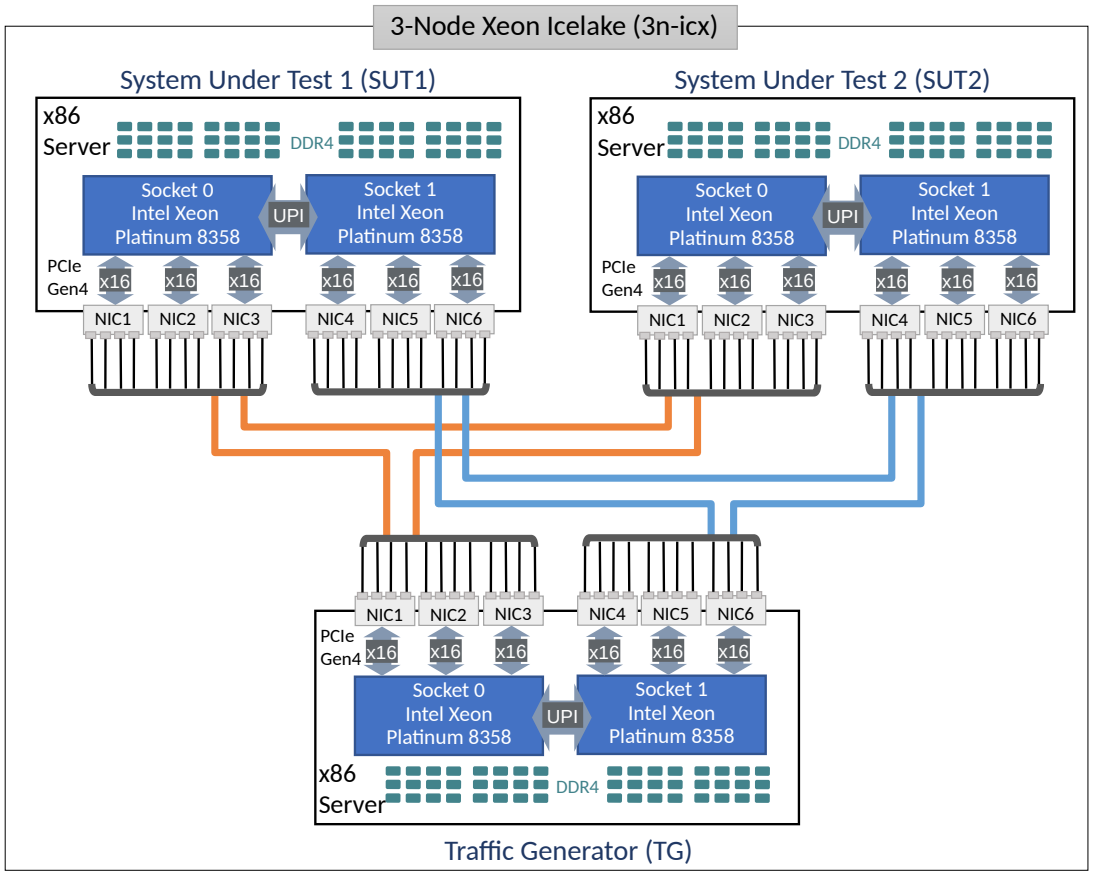

SUT and TG NICs:

1. NIC-1: xxv710-DA2 2p25GE Intel.
2. NIC-2: E810-2CQDA2 2p100GbE Intel (* to be added).
3. NIC-3: E810-CQDA4 4p100GbE Intel (* to be added).

All Intel Xeon Icelake servers run with Intel Hyper-Threading enabled, doubling the number of logical cores exposed to Linux.

1.4.6 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.

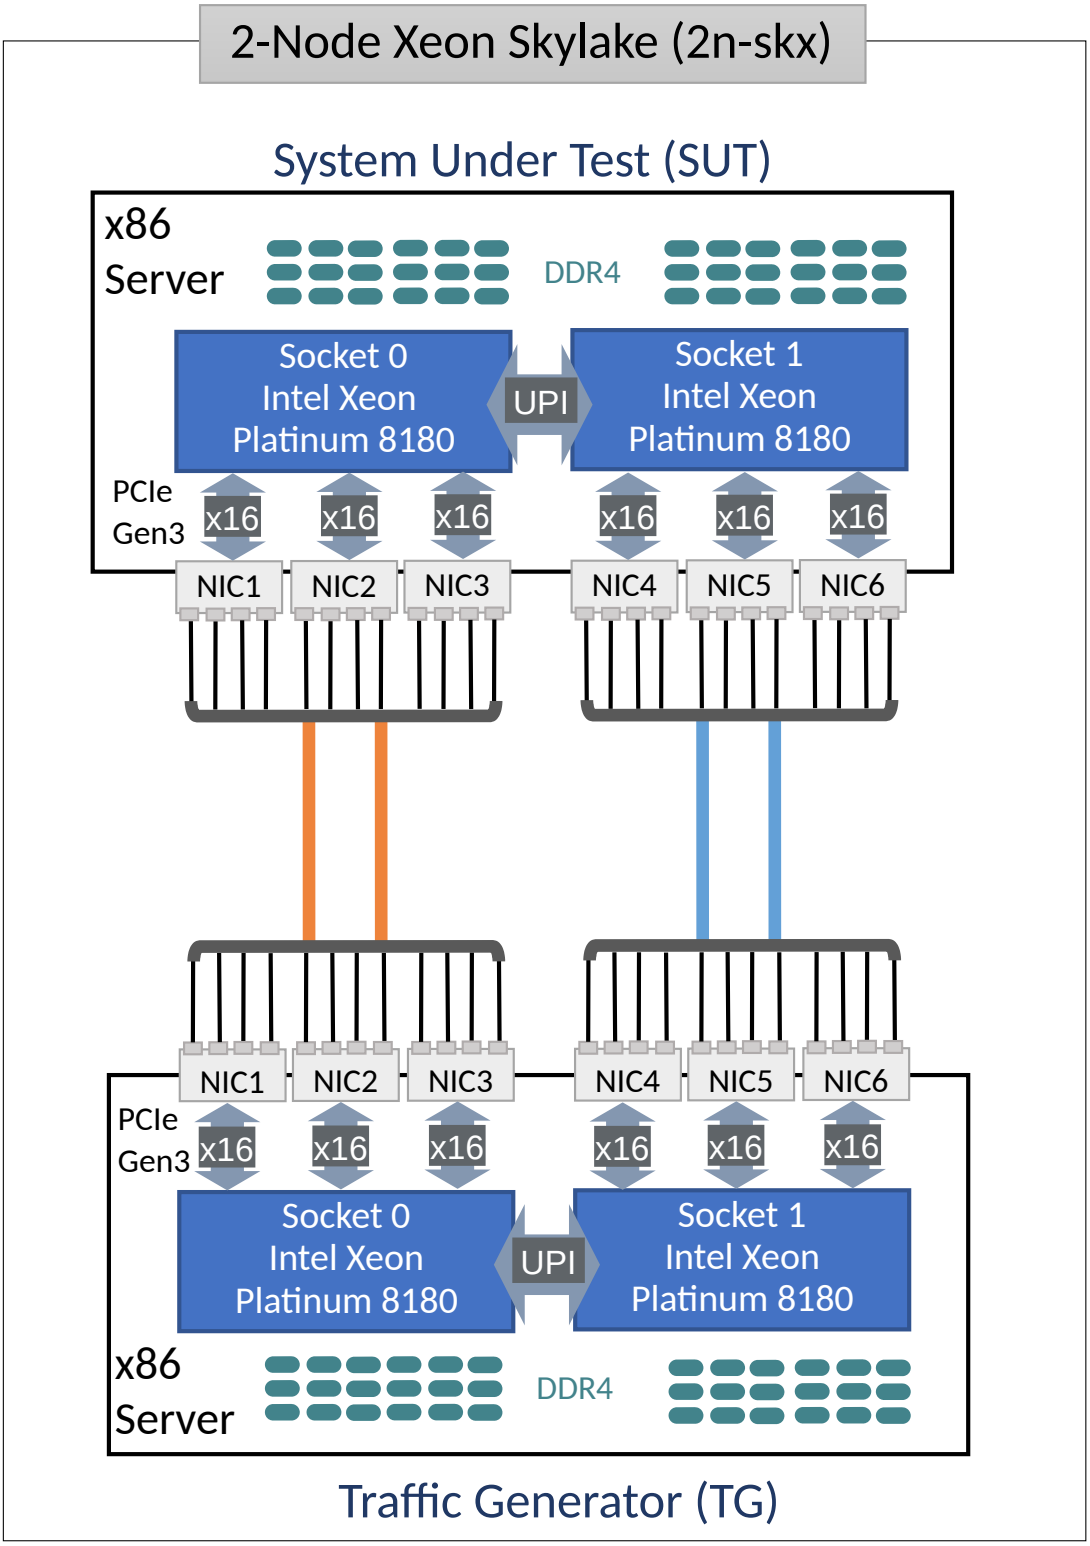

SUT NICs:

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 NICs:

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.

1.4.7 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.

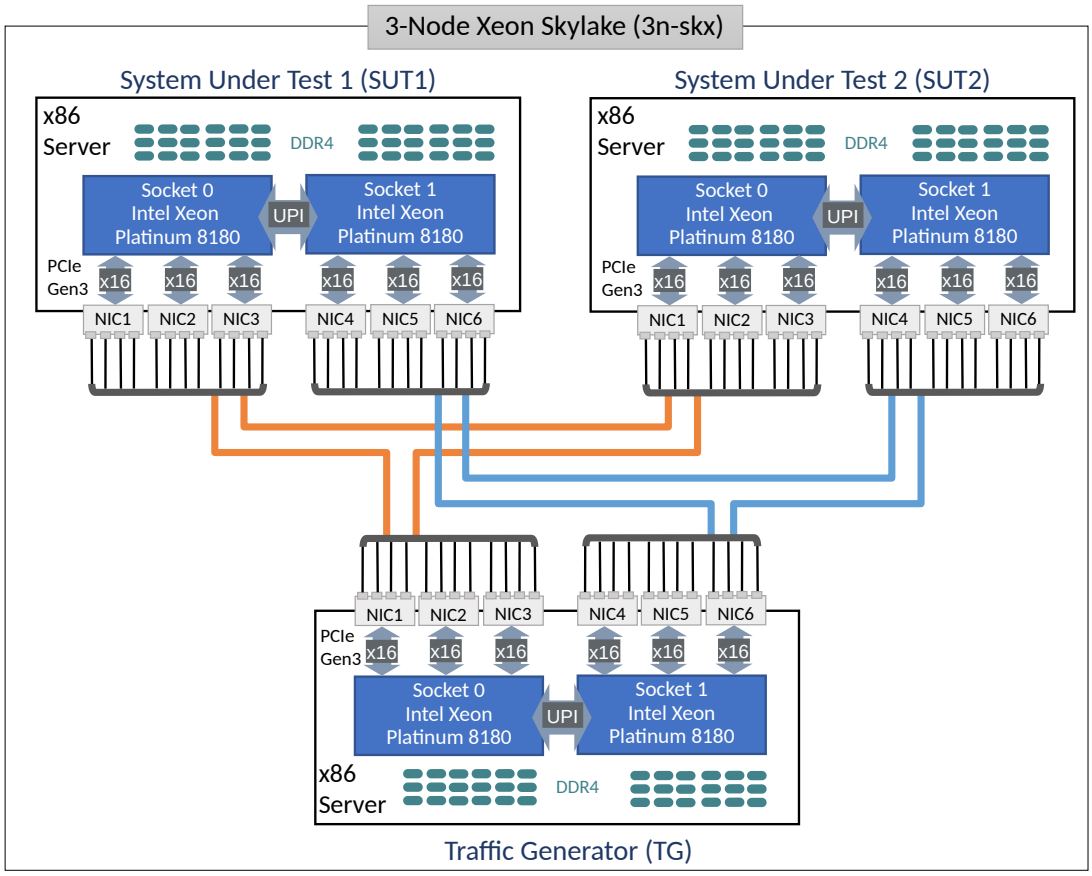

SUT1 and SUT2 NICs:

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 NICs:

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.

1.4.8 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.

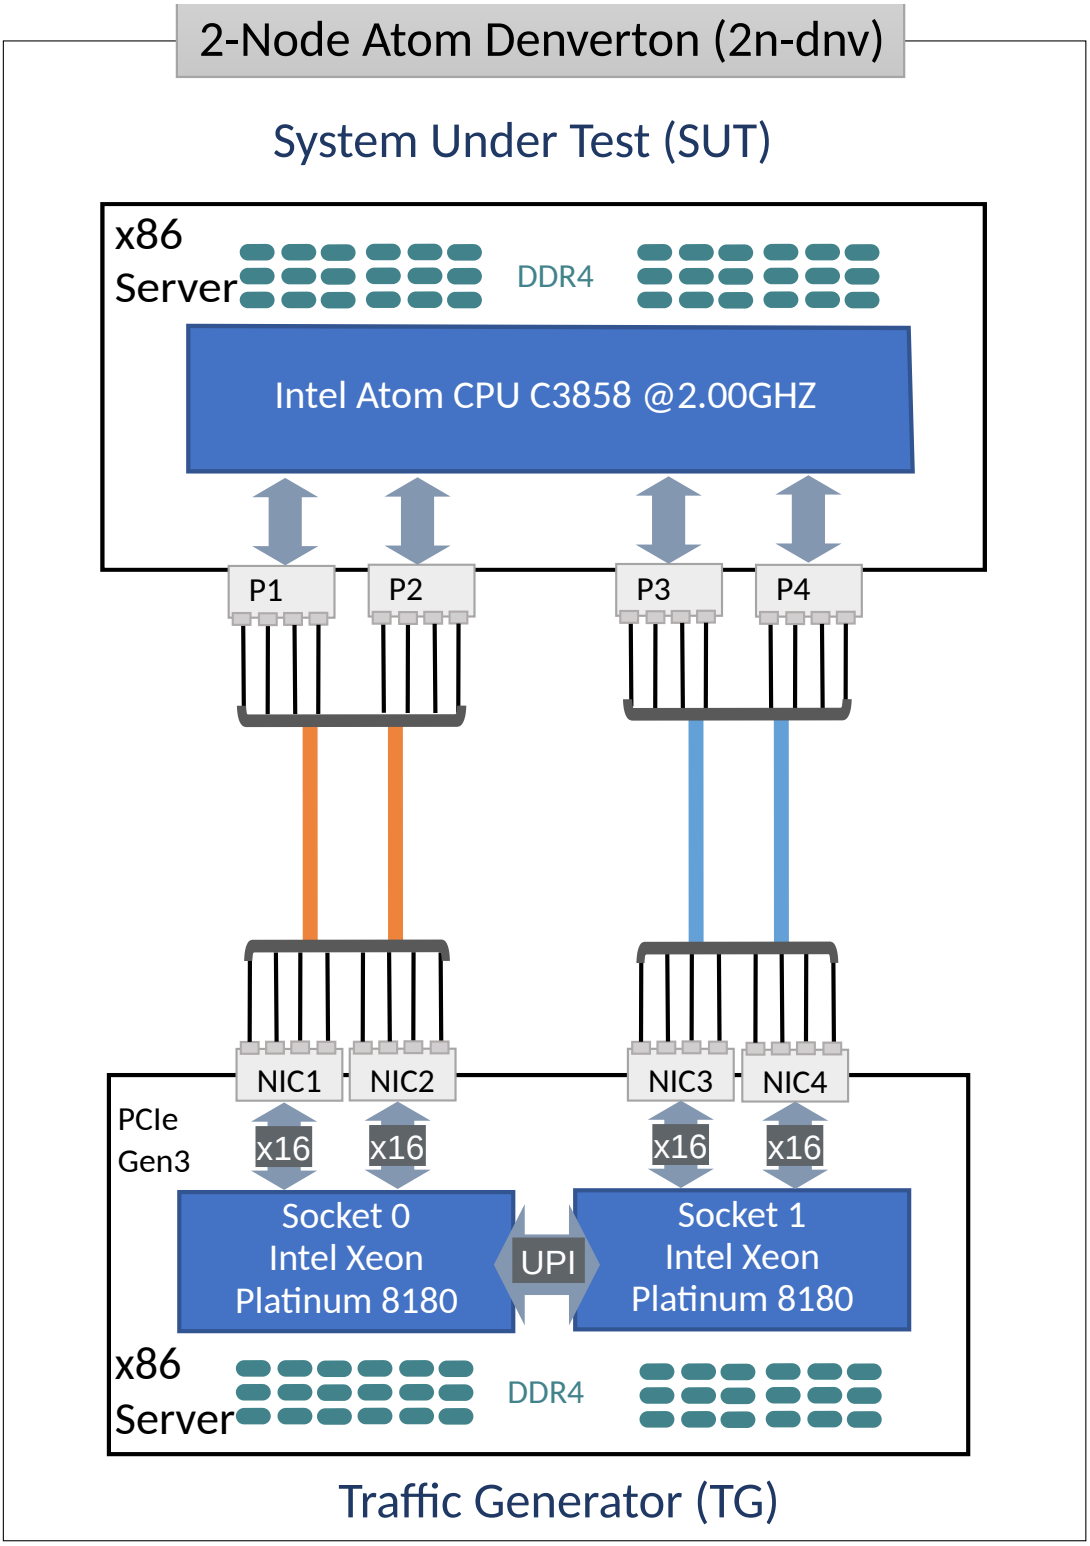

SUT 10GE NIC ports:

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 NICs:

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.

1.4.9 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.

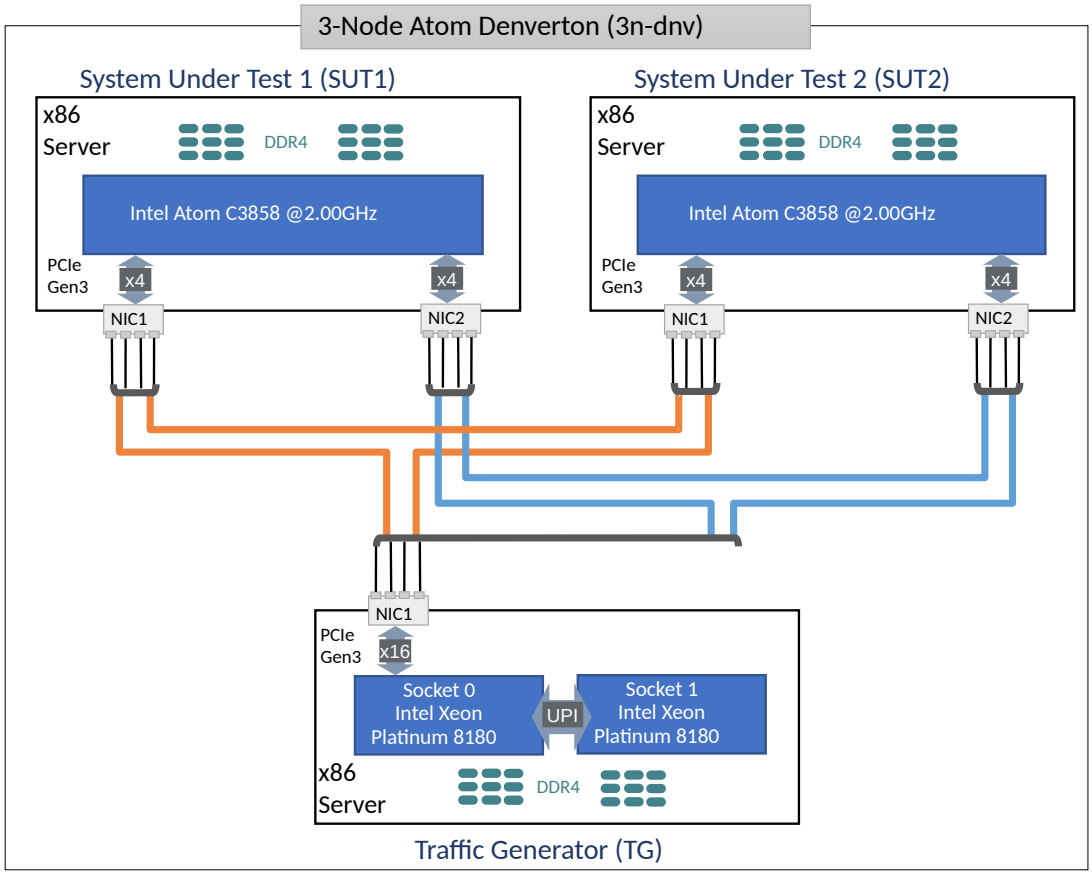

SUT1 and SUT2 NICs:

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

TG NICs:

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

1.4.10 3-Node ARM Altra (3n-alt)

One 3n-tsh testbed is built with: i) one SuperMicro SYS-740GP-TNRT server acting as TG and equipped with two Intel Xeon Icelake Platinum 8358 processors (80 MB Cache, 2.60 GHz, 32 cores), and ii) one Ampere Altra server acting as SUT and equipped with two Q80-30 processors (80* ARM Neoverse N1). 3n-alt physical topology is shown below.

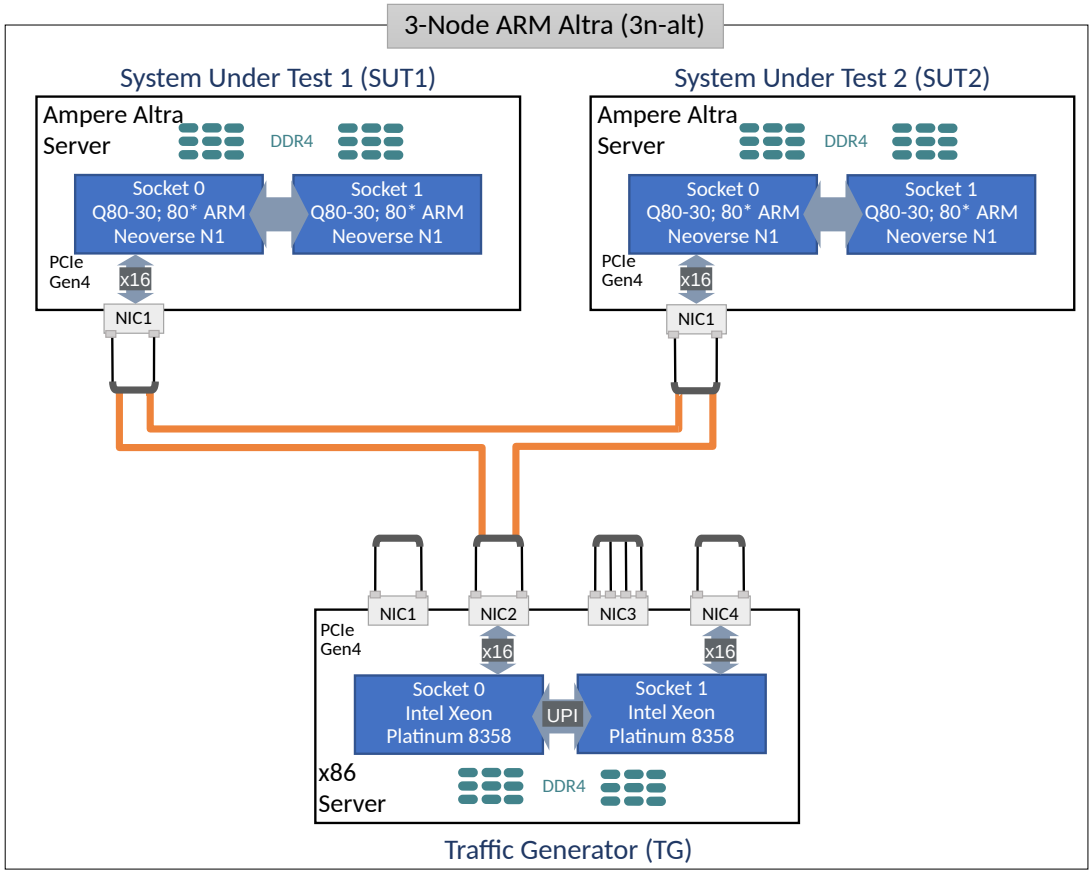

SUT1 and SUT2 NICs:

1. NIC-1: xl710-QDA2-2p40GE Intel.

TG NICs:

1. NIC-1: xxv710-DA2-2p25GE Intel.
2. NIC-2: xl710-QDA2-2p40GE Intel.
3. NIC-3: e810-XXVDA4-4p25GE Intel.
4. NIC-4: e810-2CQDA2-2p100GE Intel.

1.4.11 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.

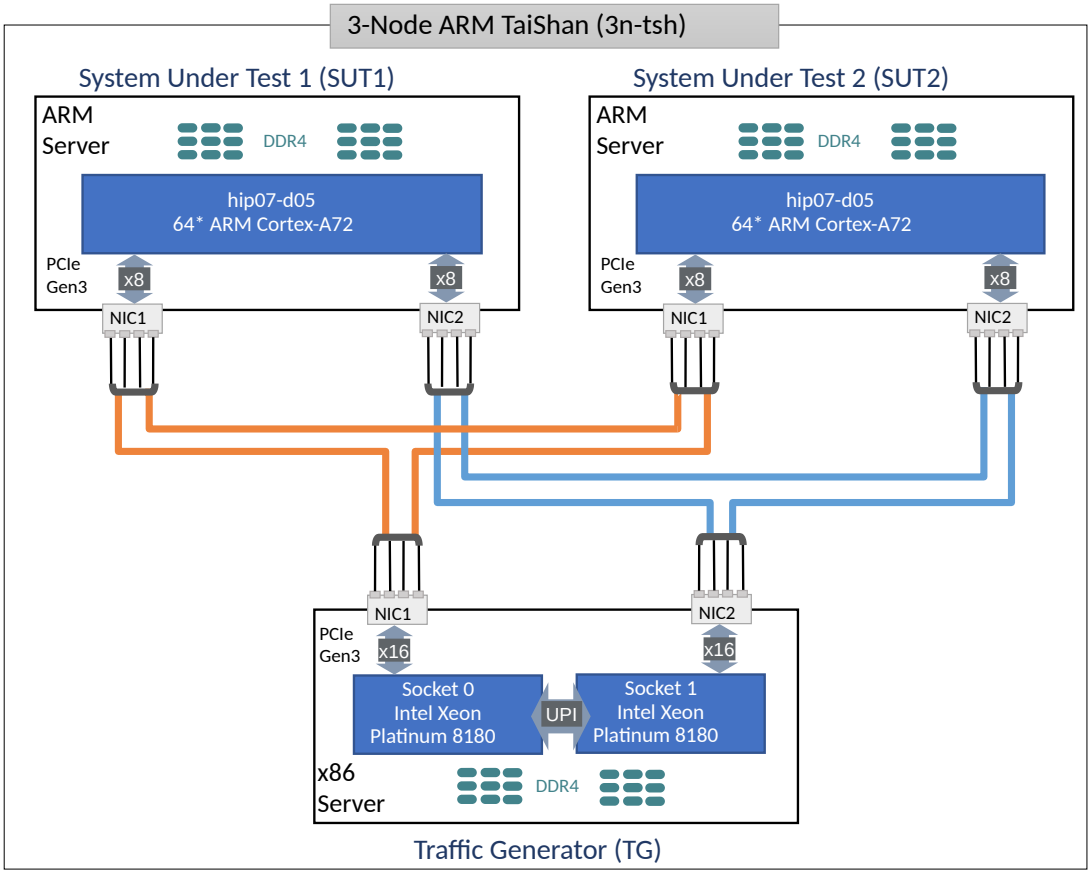

SUT1 and SUT2 NICs:

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

TG NICs:

1. NIC-1: x710-DA4 4p10GE Intel.
2. NIC-2: xxv710-DA2 2p25GE Intel.
3. NIC-3: xl710-QDA2 2p40GE Intel.

1.4.12 2-Node ARM ThunderX2 (2n-tx2)

One 2n-tx2 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 Marvell ThnderX2 9975 (28* ThunderX2) server acting as SUT and equipped with two ThunderX2 ARMv8 CN9975 processors. 2n-tx2 physical topology is shown below.

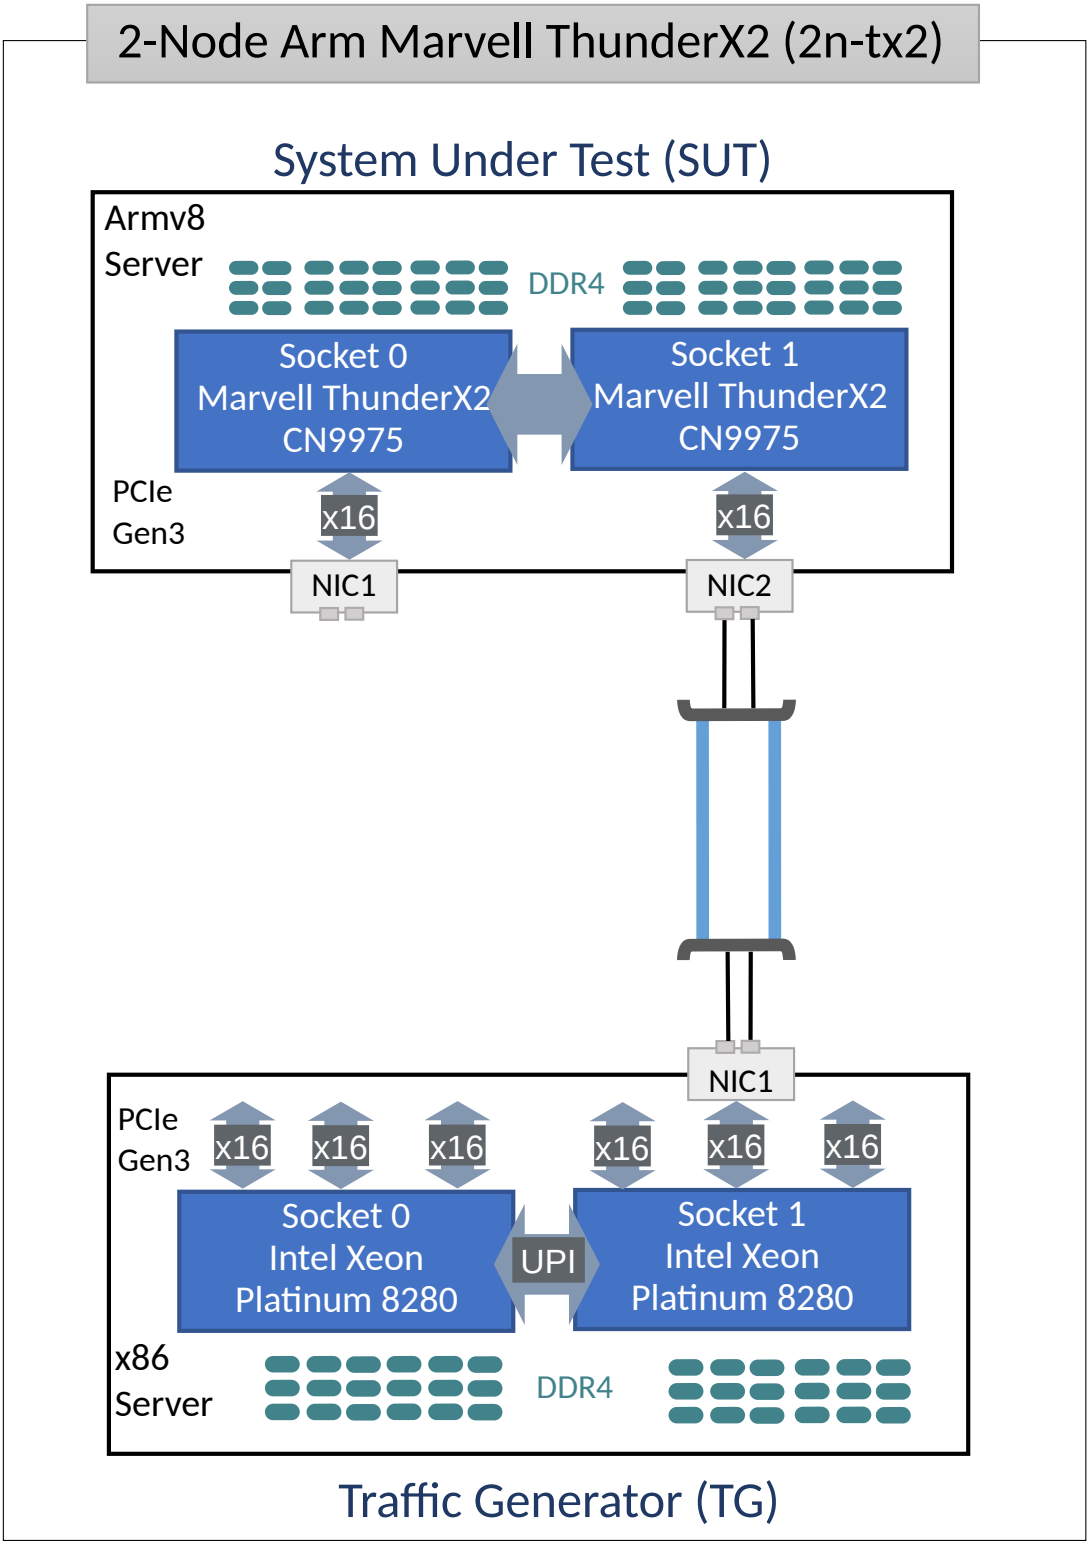

SUT NICs:

1. NIC-1: xl710-QDA2 2p40GE Intel (not connected).
2. NIC-2: xl710-QDA2 2p40GE Intel.

TG NICs:

1. NIC-1: xl710-QDA2 2p40GE 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[2], using SUT instead of RFC 2544[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 east-west and west-east 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 east-west or west-east 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

[2] https://tools.ietf.org/html/rfc2544.html
[3] https://tools.ietf.org/html/rfc2544.html

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

1.5.2 Per Thread Resources

CSIT test framework is managing mapping of the following resources per thread:

1. Cores, physical cores (pcores) allocated as pairs of sibling logical cores (lcores) if server in Hyper-Threading/SMT mode, or as single lcores if server not in HyperThreading/SMT mode. Note that if server's processors are running in HyperThreading/SMT mode sibling lcores are always used.
2. Receive Queues (RxQ), packet receive queues allocated on each physical and logical interface tested.
3. Transmit Queues(TxQ), packet transmit queues allocated on each physical and logical interface tested.

Approach to mapping per thread resources depends on the application/DUT tested (VPP or DPDK apps) and associated thread types, as follows:

1. Data-plane workers, used for data-plane packet processing, when no feature workers present.
    - Cores: data-plane workers are typically tested in 1, 2 and 4 pcore configurations, running on single lcore per pcore or on sibling lcores per pcore. Result is a set of {T}t{C}c thread-core configurations, where{T} stands for a total number of threads (lcores), and {C} for a total number of pcores. Tested configurations are encoded in CSIT test case names, e.g. "1c", "2c", "4c", and test tags "2T1C"(or "1T1C"), "4T2C" (or "2T2C"), "8T4C" (or "4T4C").
    - Interface Receive Queues (RxQ): as of CSIT-2106 release, number of RxQs used on each physical or virtual interface is equal to the number of data-plane workers. In other words each worker has a dedicated RxQ on each interface tested. This ensures packet processing load to be equal for each worker, subject to RSS flow load balancing efficacy. Note: Before CSIT-2106 total number of RxQs across all interfaces of specific type was equal to the number of data-plane workers.
    - Interface Transmit Queues (TxQ): number of TxQs used on each physical or virtual interface is equal to the number of data-plane workers. In other words each worker has a dedicated TxQ on each interface tested.
    - Applies to VPP and DPDK Testpmd and L3Fwd.
2. Data-plane and feature workers (e.g. IPsec async crypto workers), the latter dedicated to specific feature processing.

- Cores: data-plane and feature workers are tested in 2, 3 and 4 pcore configurations, running on single lcore per pcore or on sibling lcores per pcore. This results in a two sets of thread-core combinations separated by "-", {T}t{C}c-{T}t{C}c, with the leading set denoting total number of threads (lcores) and pcores used for data-plane workers, and the trailing set denoting total number of lcores and pcores used for feature workers. Accordingly, tested configurations are encoded in CSIT test case names, e.g. "1c-1c", "1c-2c", "1c-3c", and test tags "2T1C_2T1C" (or "1T1C_1T1C"), "2T1C_4T2C"(or "1T1C_2T2C"), "2T1C_6T3C" (or "1T1C_3T3C").
- RxQ and TxQ: no RxQs and no TxQs are used by feature workers.
- Applies to VPP only.

3. Management/main worker, control plane and management.
   - Cores: single lcore.
   - RxQ: not used (VPP default behaviour).
   - TxQ: single TxQ per interface, allocated but not used (VPP default behaviour).
   - Applies to VPP only.

VPP Thread Configuration

Mapping of cores and RxQs to VPP data-plane worker threads is done in the VPP startup.conf during test suite setup:

1. *corelist-workers <list_of_cores>*: List of logical cores to run VPP data-plane workers and feature workers. The actual lcores' allocations depends on HyperThreading/SMT server configuration and per test core configuration.
   - For tests without feature workers, by default, all CPU cores configured in startup.conf are used for data-plane workers.
   - For tests with feature workers, CSIT code distributes lcores across data-plane and feature workers.
2. *num-rx-queues <value>*: Number of Rx queues used per interface.

Mapping of TxQs to VPP data-plane worker threads uses the default VPP setting of one TxQ per interface per data-plane worker.

DPDK Thread Configuration

Mapping of cores and RxQs to DPDK Testpmd/L3Fwd data-plane worker threads is done in the startup CLI:

1. *-l <list_of_cores>* - List of logical cores to run DPDK application.
2. *nb-cores=<N>* - Number of forwarding cores.
3. *rxq=<N>* - Number of Rx queues used per interface.

1.5.3 VPP Forwarding Modes

VPP is tested in a number of L2, IPv4 and IPv6 packet lookup and forwarding modes. Within each mode baseline and scale tests are executed, the latter with varying number of FIB 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*: Two MAC FIB entries are learned by VPP to enable packet switching between two interfaces in two directions. VPP L2 switching is tested with 254 IPv4 unique flows per direction, varying IPv4 source address per flow in order to invoke RSS based packet distribution across VPP workers. The same source and destination MAC address is used for all flows per direction. IPv4 source address is incremented for every packet.
- *l2bdscale*: A high number of MAC FIB entries are learned by VPP to enable packet switching between two interfaces in two directions. Tested MAC FIB sizes include: i) 10k with 5k unique flows per direction, ii) 100k with 2 x 50k flows and iii) 1M with 2 x 500k flows. Unique flows are created by using distinct source and destination MAC addresses that are changed for every packet using incremental ordering, making VPP learn (or refresh) distinct src MAC entries and look up distinct dst MAC entries for every packet. For details, see *Packet Flow Ordering* (page 45).

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

IPv4 Routing

IPv4 routing tests are executed in baseline and scale configurations:

- *ip4base*: Two /32 IPv4 FIB entries are configured in VPP to enable packet routing between two interfaces in two directions. VPP routing is tested with 253 IPv4 unique flows per direction, varying IPv4 source address per flow in order to invoke RSS based packet distribution across VPP workers. IPv4 source address is incremented for every packet.
- *ip4scale*: A high number of /32 IPv4 FIB entries are configured in VPP. Tested IPv4 FIB sizes include: i) 20k with 10k unique flows per direction, ii) 200k with 2 * 100k flows and iii) 2M with 2 * 1M flows. Unique flows are created by using distinct IPv4 destination addresses that are changed for every packet, using incremental or random ordering. For details, see *Packet Flow Ordering* (page 45).

IPv6 Routing

Similarly to IPv4, IPv6 routing tests are executed in baseline and scale configurations:

- *ip6base*: Two /128 IPv4 FIB entries are configured in VPP to enable packet routing between two interfaces in two directions. VPP routing is tested with 253 IPv6 unique flows per direction, varying IPv6 source address per flow in order to invoke RSS based packet distribution across VPP workers. IPv6 source address is incremented for every packet.
- *ip4scale*: A high number of /128 IPv6 FIB entries are configured in VPP. Tested IPv6 FIB sizes include: i) 20k with 10k unique flows per direction, ii) 200k with 2 * 100k flows and iii) 2M with 2 * 1M flows. Unique flows are created by using distinct IPv6 destination addresses that are changed for every packet, using incremental or random ordering. For details, see *Packet Flow Ordering* (page 45).

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-nodecaps*: 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.

1.5.4 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.

[4] https://tools.ietf.org/html/rfc2544.html

Details

See *MLRsearch Tests* (page 26) section for more detail. MLRsearch is being standardized in IETF in draft-ietf-bmwg-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.

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 27) 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.

[5] https://datatracker.ietf.org/doc/html/draft-ietf-bmwg-mlrsearch-01
[6] https://s3-docs.fd.io/csit/master/trending/methodology/performance_tests.html
[7] https://s3-docs.fd.io/csit/master/trending/methodology/perpatch_performance_tests.html

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 30 minutes and is executed at least twice. Results are compared against NDR and PDR rates discovered with MLRsearch.

Details

See *PLRsearch* (page 28) 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).
- 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 any number of packet throughput rates in a single search, with each rate associated with a different Packet Loss Ratio (PLR) criteria.

Two throughput rates of interest 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, currently 0.5%).

MLRsearch discovers all the rates in a single pass, reducing required time duration compared to separate `binary search`_es for each rate. 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.

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.

---

[8] https://tools.ietf.org/html/draft-vpolak-bmwg-plrsearch

Search Implementation

Detailed description of the MLRsearch algorithm is included in the IETF draft draft-ietf-bmwg-mlrsearch-01[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 is currently fully based on the -01` version of the draft-ietf-bmwg-mlrsearch[11], the PyPI version is slightly older.

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.
- Per-patch performance verification: 64B.
- Initial iterations of MLRsearch and PLRsearch: 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.
  - Initial iteration for MLRsearch: 1 sec.
  - Initial iteration for PLRsearch: 5.2 sec.
- Number of MRR trials per burst:

[9] https://datatracker.ietf.org/doc/html/draft-ietf-bmwg-mlrsearch-01
[10] https://pypi.org/project/MLRsearch/
[11] https://datatracker.ietf.org/doc/html/draft-ietf-bmwg-mlrsearch-01

  - Report performance comparison: 10.
  - Daily performance trending: 10.
  - Per-patch performance verification: 5.
  - Initial iteration for MLRsearch: 1.
  - Initial iteration for PLRsearch: 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).

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

[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

- 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 supporeted 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 implemenation in FD.io CSIT project, there is roughly 0.5 second delay between trials due to restrictons 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, called "stretch" and "erf". 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 arithmetic 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. Dimentionality 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

[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

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 heurictic 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 overal 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 targed 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.

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.

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

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 Tests vs NDR Tests* (page 1452).

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 randomnes inherently present in Monte Carlo numeric integration, but that can be supressed 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 relatative 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 gleaned 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.

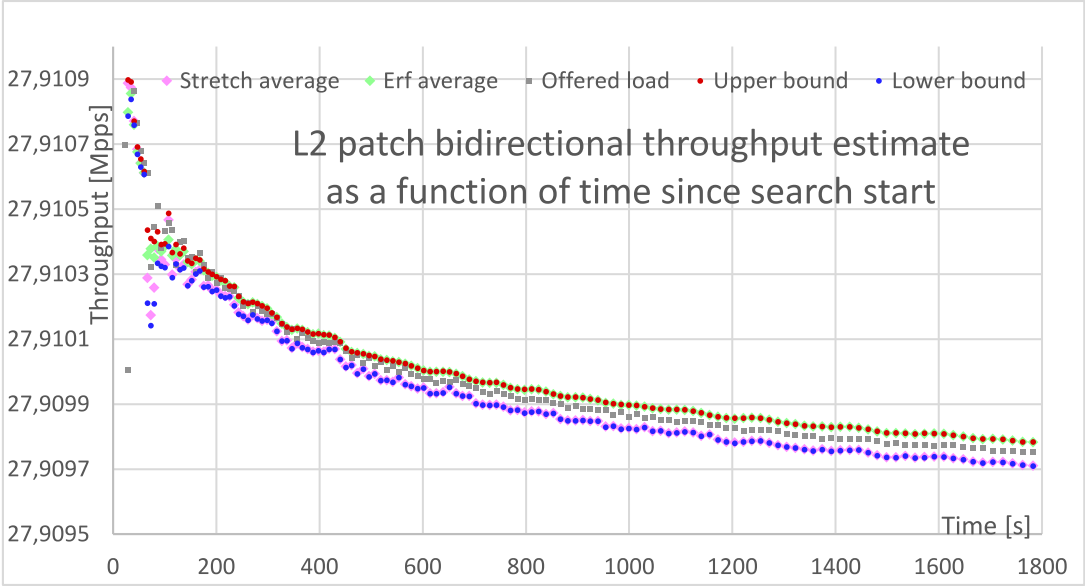

Vhost

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.

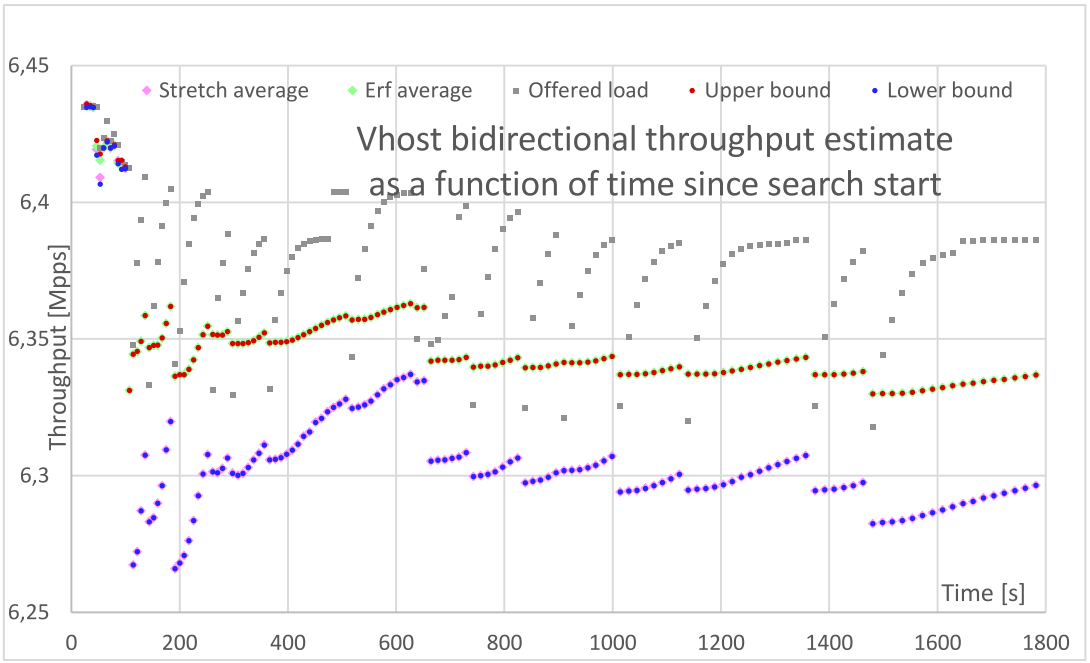

Summary

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 usefullness 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 estmated 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?

1.5.5 TRex Traffic Generator

Usage

TRex traffic generator[20] is used for majority of CSIT performance tests. TRex is used in multiple types of performance tests, see *Data Plane Throughput Tests* (page 24) for more detail.

TRex is installed and run on the TG compute node. Versioning, installation and startup is documented in *TG Settings - TRex* (page 1457).

[20] https://trex-tgn.cisco.com

Traffic modes

TRex is primarily used in two (mutually incompatible) modes.

Stateless mode

Sometimes abbreviated as STL. A mode with high performance, which is unable to react to incoming traffic. We use this mode whenever it is possible. Typical test where this mode is not applicable is NAT44ED, as DUT does not assign deterministic outside address+port combinations, so we are unable to create traffic that does not lose packets in out2in direction.

Measurement results are based on simple L2 counters (opackets, ipackets) for each traffic direction.

Stateful mode

A mode capable of reacting to incoming traffic. Contrary to the stateless mode, only UDP and TCP is supported (carried over IPv4 or IPv6 packets). Performance is limited, as TRex needs to do more CPU processing. TRex suports two subtypes of stateful traffic, CSIT uses ASTF (Advanced STateFul mode).

This mode is suitable for NAT44ED tests, as clients send packets from inside, and servers react to it, so they see the outside address and port to respond to. Also, they do not send traffic before NAT44ED has created the corresponding translation entry.

When possible, L2 counters (opackets, ipackets) are used. Some tests need L7 counters, which track protocol state (e.g. TCP), but those values are less than reliable on high loads.

Traffic Continuity

Generated traffic is either continuous, or limited (by number of transactions). Both modes support both continuities in principle.

Continuous traffic

Traffic is started without any data size goal. Traffic is ended based on time duration, as hinted by search algorithm. This is useful when DUT behavior does not depend on the traffic duration. The default for stateless mode.

Limited traffic

Traffic has defined data size goal (given as number of transactions), duration is computed based on this goal. Traffic is ended when the size goal is reached, or when the computed duration is reached. This is useful when DUT behavior depends on traffic size, e.g. target number of NAT translation entries, each to be hit exactly once per direction. This is used mainly for stateful mode.

Traffic synchronicity

Traffic can be generated synchronously (test waits for duration) or asynchronously (test operates during traffic and stops traffic explicitly).

Synchronous traffic

Trial measurement is driven by given (or precomputed) duration, no activity from test driver during the traffic. Used for most trials.

Asynchronous traffic

Traffic is started, but then the test driver is free to perform other actions, before stopping the traffic explicitly. This is used mainly by reconf tests, but also by some trials used for runtime telemetry.

Trafic profiles

TRex supports several ways to define the traffic. CSIT uses small Python modules based on Scapy as definitions. Details of traffic profiles depend on modes (STL or ASTF), but some are common for both modes.

Search algorithms are intentionally unaware of the traffic mode used, so CSIT defines some terms to use instead of mode-specific TRex terms.

Transactions

TRex traffic profile defines a small number of behaviors, in CSIT called transaction templates. Traffic profiles also instruct TRex how to create a large number of transactions based on the templates.

Continuous traffic loops over the generated transactions. Limited traffic usually executes each transaction once (typically as constant number of loops over source addresses, each loop with different source ports).

Currently, ASTF profiles define one transaction template each. Number of packets expected per one transaction varies based on profile details, as does the criterion for when a transaction is considered successful.

Stateless transactions are just one packet (sent from one TG port, successful if received on the other TG port). Thus unidirectional stateless profiles define one transaction template, bidirectional stateless profiles define two transaction templates.

TPS multiplier

TRex aims to open transaction specified by the profile at a steady rate. While TRex allows the transaction template to define its intended “cps” value, CSIT does not specify it, so the default value of 1 is applied, meaning TRex will open one transaction per second (and transaction template) by default. But CSIT invocation uses “multiplier” (mult) argument when starting the traffic, that multiplies the cps value, meaning it acts as TPS (transactions per second) input.

With a slight abuse of nomenclature, bidirectional stateless tests set “packets per transaction” value to 2, just to keep the TPS semantics as a unidirectional input value.

Duration stretching

TRex can be IO-bound, CPU-bound, or have any other reason why it is not able to generate the traffic at the requested TPS. Some conditions are detected, leading to TRex failure, for example when the bandwidth does not fit into the line capacity. But many reasons are not detected.

Unfortunately, TRex frequently reacts by not honoring the duration in synchronous mode, taking longer to send the traffic, leading to lower then requested load offered to DUT. This usualy breaks assumptions used in search algorithms, so it has to be avoided.

For stateless traffic, the behavior is quite deterministic, so the workaround is to apply a fictional TPS limit (max_rate) to search algorithms, usually depending only on the NIC used.

For stateful traffic the behavior is not deterministic enough, for example the limit for TCP traffic depends on DUT packet loss. In CSIT we decided to use logic similar to asynchronous traffic. The traffic driver sleeps for a time, then stops the traffic explicitly. The library that parses counters into measurement results than usually treats unsent packets/transactions as lost/failed.

We have added a IP4base tests for every NAT44ED test, so that users can compare results. If the results are very similar, it is probable TRex was the bottleneck.

Startup delay

By investigating TRex behavior, it was found that TRex does not start the traffic in ASTF mode immediately. There is a delay of zero traffic, after which the traffic rate ramps up to the defined TPS value.

It is possible to poll for counters during the traffic (fist nonzero means traffic has started), but that was found to influence the NDR results.

Thus “sleep and stop” stategy is used, which needs a correction to the computed duration so traffic is stopped after the intended duration of real traffic. Luckily, it turns out this correction is not dependend on traffic profile nor CPU used by TRex, so a fixed constant (0.112 seconds) works well. Unfortunately, the constant may depend on TRex version, or execution environment (e.g. TRex in AWS).

The result computations need a precise enough duration of the real traffic, luckily server side of TRex has precise enough counter for that.

It is unknown whether stateless traffic profiles also exhibit a startup delay. Unfortunately, stateless mode does not have similarly precise duration counter, so some results (mostly MRR) are affected by less precise duration measurement in Python part of CSIT code.

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 and encoded HDRHistogram data.

1.5.6 DUT state considerations

This page discusses considerations for Device Under Test (DUT) state. DUTs such as VPP require configuration, to be provided before the aplication starts (via config files) or just after it starts (via API or CLI access).

During operation DUTs gather various telemetry data, depending on configuration. This internal state handling is part of normal operation, so any performance impact is included in the test results. Accessing telemetry data is additional load on DUT, so we are not doing that in main trial measurements that affect results, but we include separate trials specifically for gathering runtime telemetry.

But there is one kind of state that needs specific handling. This kind of DUT state is dynamically created based on incoming traffic, it affects how DUT handles the traffic, and (unlike telemetry counters) it has

uneven impact on CPU load. Typical example is NAT, where detecting new sessions takes more CPU than forwarding packet on existing (open or recently closed) sessions. We call DUT configurations with this kind of state "stateful", and configurations without them "stateless". (Even though stateless configurations contain state described in previous paragraphs, and some configuration items may have "stateful" in their name, such as stateful ACLs.)

Stateful DUT configurations

Typically, the level of CPU impact of traffic depends on DUT state. The first packets causing DUT state to change have higher impact, subsequent packets matching that state have lower impact.

From performance point of view, this is similar to traffic phases for stateful protocols, see *NGFW draft <https://tools.ietf.org/html/draft-ietf-bmwg-ngfw-performance-05#section-4.3.4>*. In CSIT we borrow the terminology (even if it does not fit perfectly, see discussion below). Ramp-up traffic causes the state change, sustain traffic does not change the state.

As the performance is different, each test has to choose which traffic it wants to test, and manipulate the DUT state to achieve the intended impact.

Ramp-up trial

Tests aiming at sustain performance need to make sure DUT state is created. We achieve this via a ramp-up trial, specific purpose of which is to create the state.

Subsequent trials need no specific handling, as long as the state remains the same. But some state can time-out, so additional ramp-up trials are inserted whenever the code detects the state can time-out. Note that a trial with zero loss refreshes the state, so only the time since the last non-zero loss trial is tracked.

For the state to be set completely, it is important both DUT and TG do not lose any packets. We achieve this by setting the profile multiplier (TPS from now on) to low enough value.

It is also important each state-affecting packet is sent. For size-limited traffic profile it is guaranteed by the size limit. For continuous traffic, we set a long enough duration (based on TPS).

At the end of the ramp-up trial, we check DUT state to confirm it has been created as expected. Test fails if the state is not (completely) created.

State Reset

Tests aiming at ramp-up performance do not use ramp-up trial, and they need to reset the DUT state before each trial measurement. The way of resetting the state depends on test, usually an API call is used to partially de-configure the part that holds the state, and then re-configure it back.

In CSIT we control the DUT state behavior via a test variable "resetter". If it is not set, DUT state is not reset. If it is set, each search algorithm (including MRR) will invoke it before all trial measurements (both main and telemetry ones). Any configuration keyword enabling a feature with DUT state will check whether a test variable for ramp-up rate is present. If it is present, resetter is not set. If it is not present, the keyword sets the apropriate resetter value. This logic makes sure either ramp-up or state reset are used.

Notes: If both ramp-up and state reset were used, the DUT behavior would be identical to just reset, while test would take longer to execute. If neither were used, DUT will show different performance in subsequent trials, violating assumptions of search algorithms.

DUT versus protocol ramp-up

There are at least three different causes for bandwidth possibly increasing within a single measurement trial.

The first is DUT switching from state modification phase to constant phase, it is the primary focus of this document. Using ramp-up traffic before main trials eliminates this cause for tests wishing to measure the performance of the next phase. Using size-limited profiles eliminates the next phase for tests wishing to measure performance of this phase.

The second is protocol such as TCP ramping up their throughput to utilize the bandwidth available. This is the original meaning of "ramp up" in the NGFW draft (see above). In existing tests we are not using this meaning of TCP ramp-up. Instead we use only small transactions, and large enough initial window so TCP acts as ramped-up already.

The third is TCP increasing offered load due to retransmissions triggered by packet loss. In CSIT we again try to avoid this behavior by using small enough data to transfer, so overlap of multiple transactions (primary cause of packet loss) is unlikely. But in MRR tests, packet loss and non-constant offered load are still expected.

Stateless DUT configuratons

These are simple configurations, which do not set any resetter value (even if ramp-up duration is not configured). Majority of existing tests are of this type, using continuous traffic profiles.

In order to identify limits of Trex performance, we have added suites with stateless DUT configuration (VPP ip4base) subjected to size-limited ASTF traffic. The discovered rates serve as a basis of comparison for evaluating the results for stateful DUT configurations (VPP NAT44ed) subjected to the same traffic profiles.

DUT versus TG state

Traffic Generator profiles can be stateful (ASTF) or stateless (STL). DUT configuration can be stateful or stateless (with respect to packet traffic).

In CSIT we currently use all four possible configurations:

- Regular stateless VPP tests use stateless traffic profiles.
- Stateless VPP configuration with stateful profile is used as a base for comparison.
- Some stateful DUT configurations (NAT44DET, NAT44ED unidirectional) are tested using stateless traffic profiles and continuous traffic.
- The rest of stateful DUT configurations (NAT44ED bidirectional) are tested using stateful traffic profiles and size limited traffic.

1.5.7 Network Address Translation IPv4 to IPv4

NAT44 Prefix Bindings

NAT44 prefix bindings should be representative to target applications, where a number of private IPv4 addresses from the range defined by **RFC 1918**[21] is mapped to a smaller set of public IPv4 addresses from the public range.

Following quantities are used to describe inside to outside IP address and port bindings scenarios:

- Inside-addresses, number of inside source addresses (representing inside hosts).
- Ports-per-inside-address, number of TCP/UDP source ports per inside source address.

[21] https://tools.ietf.org/html/rfc1918.html

- Outside-addresses, number of outside (public) source addresses allocated to NAT44.
- Ports-per-outside-address, number of TCP/UDP source ports per outside source address. The maximal number of ports-per-outside-address usable for NAT is 64 512 (in non-reserved port range 1024-65535, RFC 4787[22]).
- Sharing-ratio, equal to inside-addresses divided by outside-addresses.

CSIT NAT44 tests are designed to take into account the maximum number of ports (sessions) required per inside host (inside-address) and at the same time to maximize the use of outside-address range by using all available outside ports. With this in mind, the following scheme of NAT44 sharing ratios has been devised for use in CSIT:

ports-per-inside-address	sharing-ratio
63	1024
126	512
252	256
504	128

Initial CSIT NAT44 tests, including associated TG/TRex traffic profiles, are based on ports-per-inside-address set to 63 and the sharing ratio of 1024. This approach is currently used for all NAT44 tests including NAT44det (NAT44 deterministic used for Carrier Grade NAT applications) and NAT44ed (Endpoint Dependent).

Private address ranges to be used in tests:

- 192.168.0.0 - 192.168.255.255 (192.168/16 prefix)
  - Total of 2^16 (65 536) of usable IPv4 addresses.
  - Used in tests for up to 65 536 inside addresses (inside hosts).
- 172.16.0.0 - 172.31.255.255 (172.16/12 prefix)
  - Total of 2^20 (1 048 576) of usable IPv4 addresses.
  - Used in tests for up to 1 048 576 inside addresses (inside hosts).

NAT44 Session Scale

NAT44 session scale tested is govern by the following logic:

- Number of inside-addresses(hosts) H[i] = (H[i-1] x 2^2) with H(0)=1 024, i = 1,2,3, ...
  - H[i] = 1 024, 4 096, 16 384, 65 536, 262 144, ...
- Number of sessions S[i] = H[i] * ports-per-inside-address
  - ports-per-inside-address = 63

i	hosts	sessions
0	1 024	64 512
1	4 096	258 048
2	16 384	1 032 192
3	65 536	4 128 768
4	262 144	16 515 072

[22] https://tools.ietf.org/html/rfc4787.html

NAT44 Deterministic

NAT44det performance tests are using TRex STL (Stateless) API and traffic profiles, similar to all other stateless packet forwarding tests like ip4, ip6 and l2, sending UDP packets in both directions inside-to-outside and outside-to-inside. See *Data Plane Throughput Tests* (page 24) for more detail.

The inside-to-outside traffic uses single destination address (20.0.0.0) and port (1024). The inside-to-outside traffic covers whole inside address and port range, the outside-to-inside traffic covers whole outside address and port range.

NAT44det translation entries are created during the ramp-up phase, followed by verification that all entries are present, before proceeding to the main measurements of the test. This ensures session setup does not impact the forwarding performance test.

Associated CSIT test cases use the following naming scheme to indicate NAT44det scenario tested:

- ethip4udp-nat44det-h{H}-p{P}-s{S}-[mrr|ndrpdr|soak]
  - {H}, number of inside hosts, H = 1024, 4096, 16384, 65536, 262144.
  - {P}, number of ports per inside host, P = 63.
  - {S}, number of sessions, S = 64512, 258048, 1032192, 4128768, 16515072.
  - [mrr|ndrpdr|soak], MRR, NDRPDR or SOAK test.

NAT44 Endpoint-Dependent

In order to excercise NAT44ed ability to translate based on both source and destination address and port, the inside-to-outside traffic varies also destination address and port. Destination port is the same as source port, destination address has the same offset as the source address, but applied to different subnet (starting with 20.0.0.0).

As the mapping is not deterministic (for security reasons), we cannot easily use stateless bidirectional traffic profiles. Inside address and port range is fully covered, but we do not know which outside-to-inside source address and port to use to hit an open session.

Therefore, NAT44ed is benchmarked using following methodologies:

- Unidirectional throughput using *stateless* traffic profile.
- Connections-per-second (CPS) using *stateful* traffic profile.
- Bidirectional throughput (TPUT, see below) using *stateful* traffic profile.

Unidirectional NAT44ed throughput tests are using TRex STL (Stateless) APIs and traffic profiles, but with packets sent only in inside-to-outside direction. Similarly to NAT44det, NAT44ed unidirectional throughput tests include a ramp-up phase to establish and verify the presence of required NAT44ed binding entries. As the sessions have finite duration, the test code keeps inserting ramp-up trials during the search, if it detects a risk of sessions timing out. Any zero loss trial visits all sessions, so it acts also as a ramp-up.

Stateful NAT44ed tests are using TRex ASTF (Advanced Stateful) APIs and traffic profiles, with packets sent in both directions. Tests are run with both UDP and TCP sessions. As NAT44ed CPS (connections-per-second) stateful tests measure (also) session opening performance, they use state reset instead of ramp-up trial. NAT44ed TPUT (bidirectional throughput) tests prepend ramp-up trials as in the unidirectional tests, so the test results describe performance without translation entry creation overhead.

Associated CSIT test cases use the following naming scheme to indicate NAT44det case tested:

- Stateless: ethip4udp-nat44ed-h{H}-p{P}-s{S}-udir-[mrr|ndrpdr|soak]
  - {H}, number of inside hosts, H = 1024, 4096, 16384, 65536, 262144.
  - {P}, number of ports per inside host, P = 63.
  - {S}, number of sessions, S = 64512, 258048, 1032192, 4128768, 16515072.

  - udir-[mrr|ndrpdr|soak], unidirectional stateless tests MRR, NDRPDR or SOAK.
- Stateful: ethip4[udp|tcp]-nat44ed-h{H}-p{P}-s{S}-[cps|tput]-[mrr|ndrpdr|soak]
  - [udp|tcp], UDP or TCP sessions
  - {H}, number of inside hosts, H = 1024, 4096, 16384, 65536, 262144.
  - {P}, number of ports per inside host, P = 63.
  - {S}, number of sessions, S = 64512, 258048, 1032192, 4128768, 16515072.
  - [cps|tput], connections-per-second session establishment rate or packets-per-second average rate, or packets-per-second rate without session establishment.
  - [mrr|ndrpdr|soak], bidirectional stateful tests MRR, NDRPDR, or SOAK.

Stateful traffic profiles

There are several important details which distinguish ASTF profiles from stateless profiles.

General considerations

Protocols

ASTF profiles are limited to either UDP or TCP protocol.

Programs

Each template in the profile defines two “programs”, one for the client side and one for the server side.

Each program specifies when that side has to wait until enough data is received (counted in packets for UDP and in bytes for TCP) and when to send additional data. Together, the two programs define a single transaction. Due to packet loss, transaction may take longer, use more packets (retransmission) or never finish in its entirety.

Instances

A client instance is created according to TPS parameter for the trial, and sends the first packet of the transaction (in some cases more packets). Each client instance uses a different source address (see sequencing below) and some source port. The destination address also comes from a range, but destination port has to be constant for a given program.

TRex uses an opaque way to chose source ports, but as session counting shows, next client with the same source address uses a different source port.

Server instance is created when the first packet arrives to the server side. Source address and port of the first packet are used as destination address and port for the server responses. This is the ability we need when outside surface is not predictable.

When a program reaches its end, the instance is deleted. This creates possible issues with server instances. If the server instance does not read all the data client has sent, late data packets can cause a second copy of server instance to be created, which breaks assumptions on how many packet a transaction should have.

The need for server instances to read all the data reduces the overall bandwidth TRex is able to create in ASTF mode.

Note that client instances are not created on packets, so it is safe to end client program without reading all server data (unless the definition of transaction success requires that).

Sequencing

ASTF profiles offer two modes for choosing source and destination IP addresses for client programs: seqential and pseudorandom. In current tests we are using sequential addressing only (if destination address varies at all).

For client destination UDP/TCP port, we use a single constant value. (TRex can support multiple program pairs in the same traffic profile, distinguished by the port number.)

Transaction overlap

If a transaction takes longer to finish, compared to period implied by TPS, TRex will have multiple client or server instances active at a time.

During calibration testing we have found this increases CPU utilization, and for high TPS it can lead to TRex's Rx or Tx buffers becoming full. This generally leads to duration stretching, and/or packet loss on TRex.

Currently used transactions were chosen to be short, so risk of bad behavior is decreased. But in MRR tests, where load is computed based on NIC ability, not TRex ability, anomalous behavior is still possible (e.g. MRR values being way lower than NDR).

Delays

TRex supports adding constant delays to ASTF programs. This can be useful, for example if we want to separate connection establishment from data transfer.

But as TRex tracks delayed instances as active, this still results in higher CPU utilization and reduced performance issues (as other overlaping transactions). So the current tests do not use any delays.

Keepalives

Both UDP and TCP protocol implementations in TRex programs support keepalive duration. That means there is a configurable period of keepalive time, and TRex sends keepalive packets automatically (outside the program) for the time the program is active (started, not ended yet) but not sending any packets.

For TCP this is generally not a big deal, as the other side usually retransmits faster. But for UDP it means a packet loss may leave the receiving program running.

In order to avoid keepalive packets, keepalive value is set to a high number. Here, "high number" means that even at maximum scale and minimum TPS, there are still no keepalive packets sent within the corresponding (computed) trial duration. This number is kept the same also for smaller scale traffic profiles, to simplify maintenance.

Transaction success

The transaction is considered successful at Layer-7 (L7) level when both program instances close. At this point, various L7 counters (unofficial name) are updated on TRex.

We found that proper close and L7 counter update can be CPU intensive, whereas lower-level counters (ipackets, opackets) called L2 counters can keep up with higher loads.

For some tests, we do not need to confirm the whole transaction was successful. CPS (connections per second) tests are a typical example. We care only for NAT44ed creating a session (needs one packet in inside-to-outside direction per session) and being able to use it (needs one packet in outside-to-inside direction).

Similarly in TPUT tests (packet throuput, counting both control and data packets), we care about NAT44ed ability to forward packets, we do not care whether aplications (TRex) can fully process them at that rate.

Therefore each type of tests has its own formula (usually just one counter already provided by TRex) to count "successful enough" transactions and attempted transactions. Currently, all tests relying on L7 counters use size-limited profiles, so they know what the count of attempted transactions should be, but due to duration stretching TRex might have been unable to send that many packets. For search purposes, unattempted transactions are treated the same as attempted but failed transactions.

Sometimes even the number of transactions as tracked by search algorithm does not match the transactions as defined by ASTF programs. See TCP TPUT profile below.

UDP CPS

This profile uses a minimalistic transaction to verify NAT44ed session has been created and it allows outside-to-inside traffic.

Client instance sends one packet and ends. Server instance sends one packet upon creation and ends.

In principle, packet size is configurable, but currently used tests apply only one value (100 bytes frame).

Transaction counts as attempted when opackets counter increases on client side. Transaction counts as successful when ipackets counter increases on client side.

TCP CPS

This profile uses a minimalistic transaction to verify NAT44ed session has been created and it allows outside-to-inside traffic.

Client initiates TCP connection. Client waits until connection is confirmed (by reading zero data bytes). Client ends. Server accepts the connection. Server waits for indirect confirmation from client (by waiting for client to initiate close). Server ends.

Without packet loss, the whole transaction takes 7 packets to finish (4 and 3 per direction). From NAT44ed point of view, only the first two are needed to verify the session got created.

Packet size is not configurable, but currently used tests report frame size as 64 bytes.

Transaction counts as attempted when tcps_connattempt counter increases on client side. Transaction counts as successful when tcps_connects counter increases on client side.

UDP TPUT

This profile uses a small transaction of "request-response" type, with several packets simulating data payload.

Client sends 5 packets and closes immediately. Server reads all 5 packets (needed to avoid late packets creating new server instances), then sends 5 packets and closes. The value 5 was chosen to mirror what TCP TPUT (see below) choses.

Packet size is configurable, currently we have tests for 100, 1518 and 9000 bytes frame (to match size of TCP TPUT data frames, see below).

As this is a packet oriented test, we do not track the whole 10 packet transaction. Similarly to stateless tests, we treat each packet as a "transaction" for search algorthm packet loss ratio purposes. Therefore a "transaction" is attempted when opacket counter on client or server side is increased. Transaction is successful if ipacket counter on client or server side is increased.

If one of 5 client packets is lost, server instance will get stuck in the reading phase. This probably decreases TRex performance, but it leads to more stable results then alternatives.

TCP TPUT

This profile uses a small transaction of "request-response" type, with some data amount to be transferred both ways.

In CSIT release 22.06, TRex behavior changed, so we needed to edit the traffic profile. Let us describe the pre-22.06 profile first.

Client connects, sends 5 data packets worth of data, receives 5 data packets worth of data and closes its side of the connection. Server accepts connection, reads 5 data packets worth of data, sends 5 data packets worth of data and closes its side of the connection. As usual in TCP, sending side waits for ACK from the receiving side before proceeding with next step of its program.

Server read is needed to avoid premature close and second server instance. Client read is not stricly needed, but ACKs allow TRex to close the server instance quickly, thus saving CPU and improving performance.

The number 5 of data packets was chosen so TRex is able to send them in a single burst, even with 9000 byte frame size (TRex has a hard limit on initial window size). That leads to 16 packets (9 of them in c2s direction) to be exchanged if no loss occurs. The size of data packets is controlled by the traffic profile setting the appropriate maximum segment size. Due to TRex restrictions, the minimal size for IPv4 data frame achievable by this method is 70 bytes, which is more than our usual minimum of 64 bytes. For that reason, the data frame sizes available for testing are 100 bytes (that allows room for eventually adding IPv6 ASTF tests), 1518 bytes and 9000 bytes. There is no control over control packet sizes.

Exactly as in UDP TPUT, ipackets and opackets counters are used for counting "transactions" (in fact packets).

If packet loss occurs, there can be large transaction overlap, even if most ASTF programs finish eventually. This can lead to big duration stretching and somehow uneven rate of packets sent. This makes it hard to interpret MRR results (frequently MRR is below NDR for this reason), but NDR and PDR results tend to be stable enough.

In 22.06, the "ACK from the receiving side" behavior changed, the receiving side started sending ACK sometimes also before receiving the full set of 5 data packets. If the previous profile is understood as a "single challenge, single response" where challenge (and also response) is sent as a burst of 5 data packets, the new profile uses "bursts" of 1 packet instead, but issues the challenge-response part 5 times sequentially (waiting for receiving the response before sending next challenge). This new profile happens to have the same overall packet count (when no re-transmissions are needed). Although it is possibly more taxing for TRex CPU, the results are comparable to the old traffic profile.

Ip4base tests

Contrary to stateless traffic profiles, we do not have a simple limit that would guarantee TRex is able to send traffic at specified load. For that reason, we have added tests where "nat44ed" is replaced by "ip4base". Instead of NAT44ed processing, the tests set minimalistic IPv4 routes, so that packets are forwarded in both inside-to-outside and outside-to-inside directions.

The packets arrive to server end of TRex with different source address&port than in NAT44ed tests (no translation to outside values is done with ip4base), but those are not specified in the stateful traffic profiles. The server end (as always) uses the received address&port as destination for outside-to-inside traffic. Therefore the same stateful traffic profile works for both NAT44ed and ip4base test (of the same scale).

The NAT44ed results are displayed together with corresponding ip4base results. If they are similar, TRex is probably the bottleneck. If NAT44ed result is visibly smaller, it describes the real VPP performance.

1.5.8 Packet Latency

TRex Traffic Generator (TG) is used for measuring one-way latency in 2-Node and 3-Node physical testbed topologies. TRex integrates High Dynamic Range Histogram (HDRH)[23] functionality and reports per packet latency distribution for latency streams sent in parallel to the main load packet streams.

Following methodology is used:

- Only NDRPDR test type measures latency and only after NDR and PDR values are determined. Other test types do not involve latency streams.
- Latency is measured at different background load packet rates:
  - No-Load: latency streams only.
  - Low-Load: at 10% PDR.
  - Mid-Load: at 50% PDR.
  - High-Load: at 90% PDR.
- Latency is measured for all tested packet sizes except IMIX due to TRex TG restriction.
- 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 and HDRH latency values distribution per stream direction, hence two sets of latency values are reported per test case (marked as E-W and W-E).
- +/- 1 usec is the measurement accuracy of TRex TG and the data in HDRH latency values distribution is rounded to microseconds.
- TRex TG introduces a (background) always-on Tx + Rx latency bias of 4 usec on average per direction resulting from TRex software writing and reading packet timestamps on CPU cores. Quoted values are based on TG back-to-back latency measurements.
- Latency graphs are not smoothed, each latency value has its own horizontal line across corresponding packet percentiles.
- Percentiles are shown on X-axis using a logarithmic scale, so the maximal latency value (ending at 100% percentile) would be in infinity. The graphs are cut at 99.9999% (hover information still lists 100%).

1.5.9 Packet Flow Ordering

TRex Traffic Generator (TG) supports two main ways how to cover address space (on allowed ranges) in scale tests.

In most cases only one field value (e.g. IPv4 destination address) is altered, in some cases two fields (e.g. IPv4 destination address and UDP destination port) are altered.

Incremental Ordering

This case is simpler to implement and offers greater control.

When changing two fields, they can be incremented synchronously, or one after another. In the latter case we can specify which one is incremented each iteration and which is incremented by "carrying over" only when the other "wraps around". This way also visits all combinations once before the "carry" field also wraps around.

It is possible to use increments other than 1.

[23] http://hdrhistogram.org/

Randomized Ordering

This case chooses each field value at random (from the allowed range). In case of two fields, they are treated independently. TRex allows to set random seed to get deterministic numbers. We use a different seed for each field and traffic direction. The seed has to be a non-zero number, we use 1, 2, 3, and so on.

The seeded random mode in TRex requires a "limit" value, which acts as a cycle length limit (after this many iterations, the seed resets to its initial value). We use the maximal allowed limit value (computed as 2^24 - 1).

Randomized profiles do not avoid duplicated values, and do not guarantee each possible value is visited, so it is not very useful for stateful tests.

1.5.10 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.
- *ip4gtpusw-ip4base*: GTPU over IPv4 tunnels with IPv4 routing.

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

1.5.11 Internet Protocol Security (IPsec)

VPP IPsec performance tests are executed for the following crypto plugins:

- *crypto_native*, used for software based crypto leveraging CPU platform optimizations e.g. Intel's AES-NI instruction set.
- *crypto_ipsecmb*, used for hardware based crypto with Intel QAT PCIe cards.

IPsec with VPP Native SW Crypto

Currently CSIT-2206 implements following IPsec test cases relying on VPP native crypto (*crypto_native* plugin):

VPP Crypto Engine	ESP Encryption	ESP Integrity	Scale Tested
crypto_native	AES[128|256]-GCM	GCM	1 to 60k tunnels
crypto_native	AES128-CBC	SHA[256|512]	1 to 60k tunnels

VPP IPsec with SW crypto are executed in both tunnel and policy modes, with tests running on 3-node testbeds: 3n-skx, 3n-tsh.

IPsec with Intel QAT HW

Currently CSIT-2206 implements following IPsec test cases relying on ipsecmb library (*crypto_ipsecmb* plugin) and Intel QAT 8950 (50G HW crypto card):

dpdk_cryptodev

VPP Crypto Engine	VPP Crypto Workers	ESP Encryption	ESP Integrity	Scale Tested
crypto_ipsecmb	sync/all workers	AES[128|256]-GCM	GCM	1, 1k tunnels
crypto_ipsecmb	sync/all workers	AES[128]-CBC	SHA[256|512]	1, 1k tunnels
crypto_ipsecmb	async/crypto worker	AES[128|256]-GCM	GCM	1, 4, 1k tunnels
crypto_ipsecmb	async/crypto worker	AES[128]-CBC	SHA[256|512]	1, 4, 1k tunnels

IPsec with Async Crypto Feature Workers

TODO Description to be added

IPsec Uni-Directional Tests with VPP Native SW Crypto

Currently CSIT-2206 implements following IPsec uni-directional test cases relying on VPP native crypto (*crypto_native* plugin) in tunnel mode:

VPP Crypto Engine	ESP Encryption	ESP Integrity	Scale Tested
crypto_native	AES[128|256]-GCM	GCM	4, 1k, 10k tunnels
crypto_native	AES128-CBC	SHA[512]	4, 1k, 10k tunnels

In policy mode:

VPP Crypto Engine	ESP Encryption	ESP Integrity	Scale Tested
crypto_native	AES[256]-GCM	GCM	1, 40, 1k tunnels

The tests are running on 2-node testbeds: 2n-tx2. The uni-directional tests are partially addressing a weakness in 2-node testbed setups with T-Rex as the traffic generator. With just one DUT node, we can either encrypt or decrypt traffic in each direction.

The testcases are only doing encryption - packets are encrypted on the DUT and then arrive at TG where no additional packet processing is needed (just counting packets).

Decryption would require that the traffic generator generated encrypted packets which the DUT then would decrypt. However, T-Rex does not have the capability to encrypt packets.

1.5.12 Access Control Lists

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.

1.5.13 Multi-Core Speedup

All performance tests are executed with single physical 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-2206 performance tests are executed with server SUTs' Intel XEON processors configured with Intel Hyper-Threading Enabled for all Xeon Skylake and Xeon Cascadelake testbeds.

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

Multi-core Tests

CSIT-2206 multi-core tests are executed in the following VPP worker thread and physical core configurations:

1. Intel Xeon Skylake and Cascadelake testbeds (2n-skx, 3n-skx, 2n-clx) 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 470) includes a set of graphs illustrating packet throughout 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.

1.5.14 Hoststack Testing

TCP/IP with iperf3

iperf3 goodput measurement tool[24] is used for measuring the maximum attainable goodput of the VPP Host Stack connection across two instances of VPP running on separate DUT nodes. iperf3 is a popular open source tool for active measurements of the maximum achievable goodput on IP networks.

Because iperf3 utilizes the POSIX socket interface APIs, the current test configuration utilizes the LD_PRELOAD mechanism in the linux kernel to connect iperf3 to the VPP Host Stack using the VPP Communications Library (VCL) LD_PRELOAD library (libvcl_ldpreload.so).

[24] https://github.com/esnet/iperf

In the future, a forked version of iperf3 which has been modified to directly use the VCL application APIs may be added to determine the difference in performance of 'VCL Native' applications versus utilizing LD_PRELOAD which inherently has more overhead and other limitations.

The test configuration is as follows:

```
       DUT1              Network              DUT2
[ iperf3-client -> VPP1 ]=======[ VPP2 -> iperf3-server]
```

where,

1. iperf3 server attaches to VPP2 and LISTENs on VPP2:TCP port 5201.
2. iperf3 client attaches to VPP1 and opens one or more stream connections to VPP2:TCP port 5201.
3. iperf3 client transmits a uni-directional stream as fast as the VPP Host Stack allows to the iperf3 server for the test duration.
4. At the end of the test the iperf3 client emits the goodput measurements for all streams and the sum of all streams.

Test cases include 1 and 10 Streams with a 20 second test duration with the VPP Host Stack configured to utilize the Cubic TCP congestion algorithm.

Note: iperf3 is single threaded, so it is expected that the 10 stream test shows little or no performance improvement due to multi-thread/multi-core execution.

There are also variations of these test cases which use the VPP Network Simulator (NSIM) plugin to test the VPP Hoststack goodput with 1 percent of the traffic being dropped at the output interface of VPP1 thereby simulating a lossy network. The NSIM tests are experimental and the test results are not currently representative of typical results in a lossy network.

UDP/IP with iperf3

iperf3 goodput measurement tool[25] is used for measuring the maximum attainable goodput of the VPP Host Stack connection across two instances of VPP running on separate DUT nodes. iperf3 is a popular open source tool for active measurements of the maximum achievable goodput on IP networks.

Because iperf3 utilizes the POSIX socket interface APIs, the current test configuration utilizes the LD_PRELOAD mechanism in the linux kernel to connect iperf3 to the VPP Host Stack using the VPP Communications Library (VCL) LD_PRELOAD library (libvcl_ldpreload.so).

In the future, a forked version of iperf3 which has been modified to directly use the VCL application APIs may be added to determine the difference in performance of 'VCL Native' applications versus utilizing LD_PRELOAD which inherently has more overhead and other limitations.

The test configuration is as follows:

```
       DUT1              Network              DUT2
[ iperf3-client -> VPP1 ]=======[ VPP2 -> iperf3-server]
```

where,

1. iperf3 server attaches to VPP2 and LISTENs on VPP2:UDP port 5201.
2. iperf3 client attaches to VPP1 and transmits one or more streams of packets to VPP2:UDP port 5201.
3. iperf3 client transmits a uni-directional stream as fast as the VPP Host Stack allows to the iperf3 server for the test duration.
4. At the end of the test the iperf3 client emits the goodput measurements for all streams and the sum of all streams.

[25] https://github.com/esnet/iperf

Test cases include 1 and 10 Streams with a 20 second test duration with the VPP Host Stack using the UDP transport layer..

Note: iperf3 is single threaded, so it is expected that the 10 stream test shows little or no performance improvement due to multi-thread/multi-core execution.

QUIC/UDP/IP with vpp_echo

vpp_echo performance testing tool[26] is a bespoke performance test application which utilizes the 'native HostStack APIs' to verify performance and correct handling of connection/stream events with uni-directional and bi-directional streams of data.

Because iperf3 does not support the QUIC transport protocol, vpp_echo is used for measuring the maximum attainable goodput of the VPP Host Stack connection utilizing the QUIC transport protocol across two instances of VPP running on separate DUT nodes. The QUIC transport protocol supports multiple streams per connection and test cases utilize different combinations of QUIC connections and number of streams per connection.

The test configuration is as follows:

```
        DUT1                Network                 DUT2
[ vpp_echo-client -> VPP1 ]=======[ VPP2 -> vpp_echo-server]
                      N-streams/connection
```

where,

1. vpp_echo server attaches to VPP2 and LISTENs on VPP2:TCP port 1234.
2. vpp_echo client creates one or more connections to VPP1 and opens one or more stream per connection to VPP2:TCP port 1234.
3. vpp_echo client transmits a uni-directional stream as fast as the VPP Host Stack allows to the vpp_echo server for the test duration.
4. At the end of the test the vpp_echo client emits the goodput measurements for all streams and the sum of all streams.

Test cases include

1. 1 QUIC Connection with 1 Stream
2. 1 QUIC connection with 10 Streams
3. 10 QUIC connetions with 1 Stream
4. 10 QUIC connections with 10 Streams

with stream sizes to provide reasonable test durations. The VPP Host Stack QUIC transport is configured to utilize the picotls encryption library. In the future, tests utilizing addtional encryption algorithms will be added.

VSAP ab with nginx

VSAP (VPP Stack Acceleration Project)[27] aims to establish an industry user space application ecosystem based on the VPP hoststack. As a pre-requisite to adapting open source applications using VPP Communications Library to accelerate performance, the VSAP team has introduced baseline tests utilizing the LD_PRELOAD mechanism to capture baseline performance data.

AB (Apache HTTP server benchmarking tool)[28] is used for measuring the maximum connections-per-second and requests-per-second.

[26] https://wiki.fd.io/view/VPP/HostStack#External_Echo_Server.2FClient_.28vpp_echo.29
[27] https://wiki.fd.io/view/VSAP
[28] https://httpd.apache.org/docs/2.4/programs/ab.html

NGINX[29] is a popular open source HTTP server application. Because NGINX utilizes the POSIX socket interface APIs, the test configuration uses the LD_PRELOAD mechanism to connect NGINX to the VPP Hoststack using the VPP Communications Library (VCL) LD_PRELOAD library (libvcl_ldpreload.so).

In the future, a version of NGINX which has been modified to directly use the VCL application APIs will be added to determine the difference in performance of 'VCL Native' applications versus utilizing LD_PRELOAD which inherently has more overhead and other limitations.

The test configuration is as follows:

```
  TG     Network         DUT
[ AB ]=============[ VPP -> nginx ]
```

where,

1. nginx attaches to VPP and listens on TCP port 80
2. ab runs CPS and RPS tests with packets flowing from the Test Generator node, across 100G NICs, through VPP hoststack to NGINX.
3. At the end of the tests, the results are reported by AB.

1.5.15 Generic Segmentation Offload Tests

Overview

Generic Segmentation Offload (GSO) reduces per-packet processing overhead by enabling applications to pass a multi-packet buffer to (v)NIC and process a smaller number of large packets (e.g. frame size of 64 KB), instead of processing higher numbers of small packets (e.g. frame size of 1500 B), thus reducing per-packet overhead.

CSIT-2206 introduced GSO tests for VPP vhostuser and tapv2 interfaces. All tests cases use iPerf3 client and server applications running TCP/IP as a traffic generator. For performance comparison the same tests are run without GSO enabled.

GSO Test Topologies

Two VPP GSO test topologies are implemented in CSIT-2206:

1. iPerfC_GSOvirtio_LinuxVM — GSOvhost_VPP_GSOvhost — iPerfS_GSOvirtio_LinuxVM
   - Tests VPP GSO on vhostuser interfaces and interaction with Linux virtio with GSO enabled.
1. iPerfC_GSOtap_LinuxNspace — GSOtapv2_VPP_GSOtapv2 — iPerfS_GSOtap_LinuxNspace
   - Tests VPP GSO on tapv2 interfaces and interaction with Linux tap with GSO enabled.

Common configuration:

- iPerfC (client) and iPerfS (server) run in TCP/IP mode without upper bandwidth limit.
- Trial duration is set to 30 sec.
- iPerfC, iPerfS and VPP run in the single SUT node.

[29] https://www.nginx.com/

VPP GSOtap Topology

VPP Configuration

VPP GSOtap tests in CSIT-2206 are executed without using hyperthreading. VPP worker runs on a single core. Multi-core tests are not executed. Each interface belongs to separate namespace. Following core pinning scheme is used:

- 1t1c (rxq=1, rx_qsz=4096, tx_qsz=4096)
  - system isolated: 0,28,56,84
  - vpp mt: 1
  - vpp wt: 2
  - vhost: 3-5
  - iperf-s: 6
  - iperf-c: 7

iPerf3 Server Configuration

iPerf3 version used 3.7

```
$ sudo -E -S ip netns exec tap1_namespace iperf3 \
    --server --daemon --pidfile /tmp/iperf3_server.pid --logfile /tmp/iperf3.log --port 5201 --
↪affinity <X>
```

For the full iPerf3 reference please see: iPerf3 docs[30].

iPerf3 Client Configuration

iPerf3 version used 3.7

```
$ sudo -E -S ip netns exec tap1_namespace iperf3 \
    --client 2.2.2.2 --bind 1.1.1.1 --port 5201 --parallel <Y> --time 30.0 --affinity <X> --zerocopy
```

For the full iPerf3 reference please see: iPerf3 docs[31].

VPP GSOvhost Topology

VPP Configuration

VPP GSOvhost tests in CSIT-2206 are executed without using hyperthreading. VPP worker runs on a single core. Multi-core tests are not executed. Following core pinning scheme is used:

- 1t1c (rxq=1, rx_qsz=1024, tx_qsz=1024) - system isolated: 0,28,56,84 - vpp mt: 1 - vpp wt: 2 - vm-iperf-s: 3,4,5,6,7 - vm-iperf-c: 8,9,10,11,12 - iperf-s: 1 - iperf-c: 1

[30] https://github.com/esnet/iperf/blob/master/docs/invoking.rst
[31] https://github.com/esnet/iperf/blob/master/docs/invoking.rst

iPerf3 Server Configuration

iPerf3 version used 3.7

```
$ sudo iperf3 \
    --server --daemon --pidfile /tmp/iperf3_server.pid --logfile /tmp/iperf3.log --port 5201 --
↪affinity X
```

For the full iPerf3 reference please see: iPerf3 docs[32].

iPerf3 Client Configuration

iPerf3 version used 3.7

```
$ sudo iperf3 \
    --client 2.2.2.2 --bind 1.1.1.1 --port 5201 --parallel <Y> --time 30.0 --affinity X --zerocopy
```

For the full iPerf3 reference please see: iPerf3 docs[33].

1.5.16 Reconfiguration Tests

Important: **DISCLAIMER**: Described reconf test methodology is experimental, and subject to change following consultation within csit-dev, vpp-dev and user communities. Current test results should be treated as indicative.

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 first find a througput value (lower bound for NDR) without re-configuration, and then maintain that ofered 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 the final trial duration is shorter.

The MLRsearch suite has to contain a configuration parameter that can be scaled up, e.g. number of tunnels 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 tunnels / chains are not targetted by the traffic.

[32] https://github.com/esnet/iperf/blob/master/docs/invoking.rst
[33] https://github.com/esnet/iperf/blob/master/docs/invoking.rst

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.

1.5.17 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[34] 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).

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.

[34] https://git.fd.io/vpp/tree/src/vpp/conf/startup.conf?h=stable/2206&id=0d352a97c5e3ad1f5f6eab18a978a14b0b9e06a8

1.5.18 KVM VMs vhost-user

QEMU is used for KVM VM vhost-user testing enviroment. By default, standard QEMU version is used, preinstalled from OS repositories (qemu-2.11.1 for Ubuntu 18.04). 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.
- **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[35] 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
```

[35] https://buildroot.org/

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*.

1.5.19 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], [lxcnamespace] and [stgraber].

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] and [stgraber].

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

More information in [lxcsecurity] and [lxcsecfeatures].

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/uio_pci_generic/unbind`**. This is not the case of unprivileged containers where `sysfs` is mounted as read-only.

Orchestrating Container Lifecycle Events

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

1. Acquire - acquiring/downloading existing container images via **`docker pull`** or **`lxc-create -t download`**.
2. Build - building a container image from scratch or another container image via **`docker build <dockerfile/composefile>`** or customizing LXC templates in GitHub[36].
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.

LXC

LXC is the well-known and heavily tested low-level Linux container runtime [lxcsource], 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 is 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 vpp startup.conf.

[36] 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]. Currently it is using *runc*, a CLI tool for spawning and running containers according to the OCI specification[37].

A Docker container image is a lightweight, stand-alone, executable package that includes everything needed to run the container: 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.

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

```
+-----------------------------------------------------------------------+
|              RF Keywords (high level lifecycle control)               |
+-----------------------------------------------------------------------+
| Construct VNF containers on all DUTs                                  |
| Acquire all '${group}' containers                                     |
| Create all '${group}' containers                                      |
| Install all '${group}' containers                                     |
| Configure all '${group}' containers                                   |
| Stop all '${group}' containers                                        |
| Destroy all '${group}' containers                                     |
+-----------------+-----------------------------------------------------+
                  |  1
                  |
                  |  1..N
+-----------------v-----------------+        +--------------------------+
|          ContainerManager         |        |  ContainerEngine         |
+-----------------------------------+        +--------------------------+
| __init()__                        |        | __init(node)__           |
| construct_container()             |        | acquire(force)           |
| construct_containers()            |        | create()                 |
| acquire_all_containers()          |        | stop()                   |
| create_all_containers()           | 1    1 | destroy()                |
| execute_on_container()            <>-------| info()                   |
| execute_on_all_containers()       |        | execute(command)         |
```

(continues on next page)

[37] https://www.opencontainers.org/

(continued from previous page)

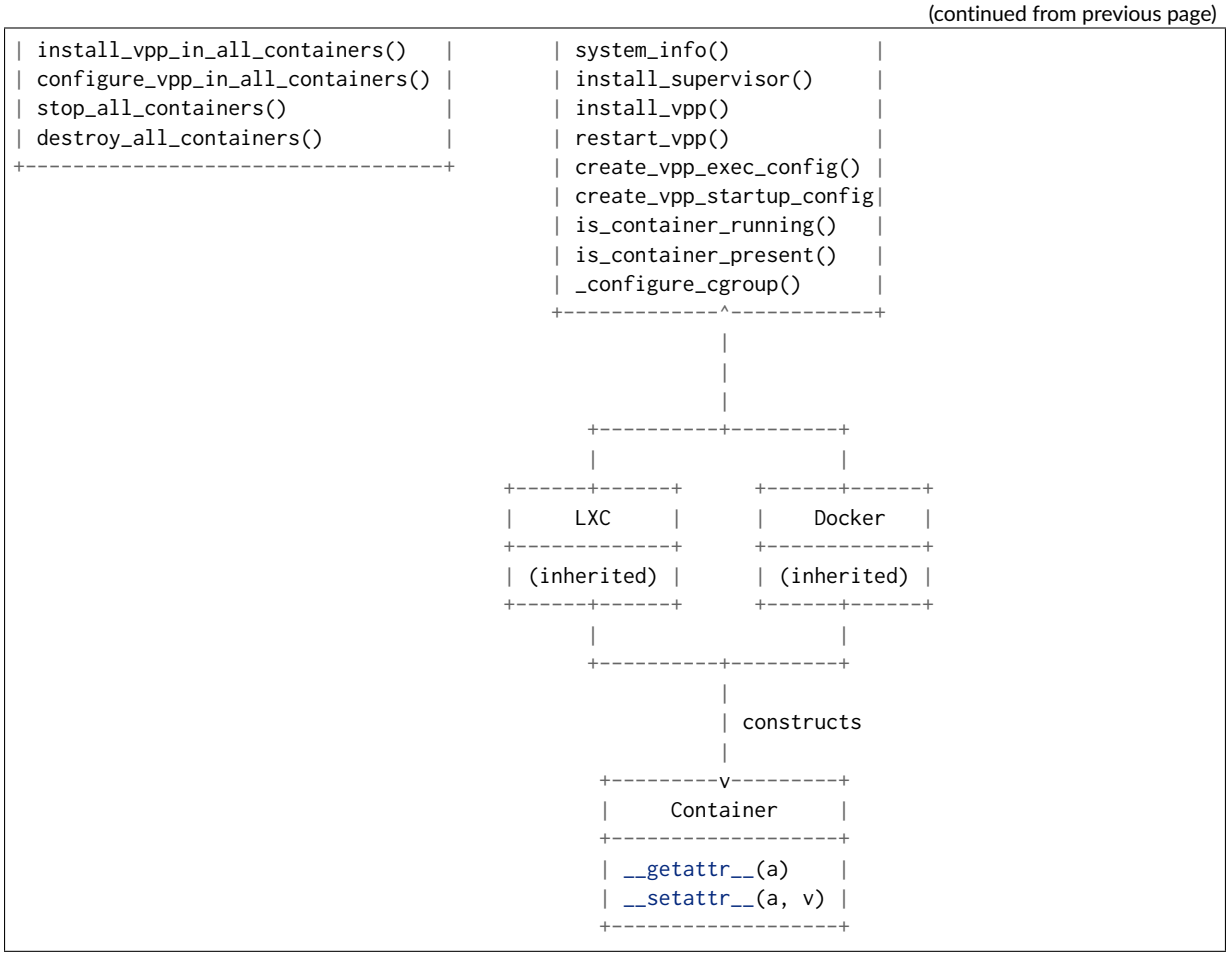

Sequentional diagram that illustrates the creation of a single container.

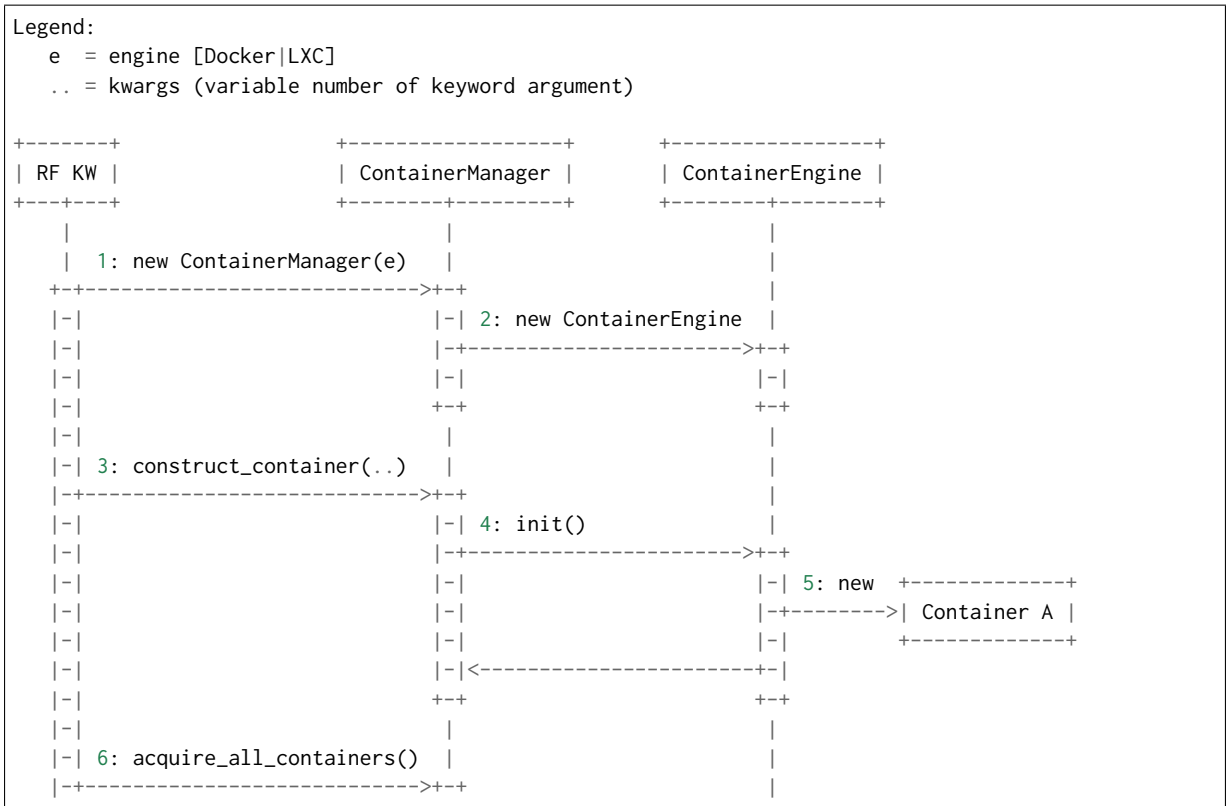

(continues on next page)

(continued from previous page)

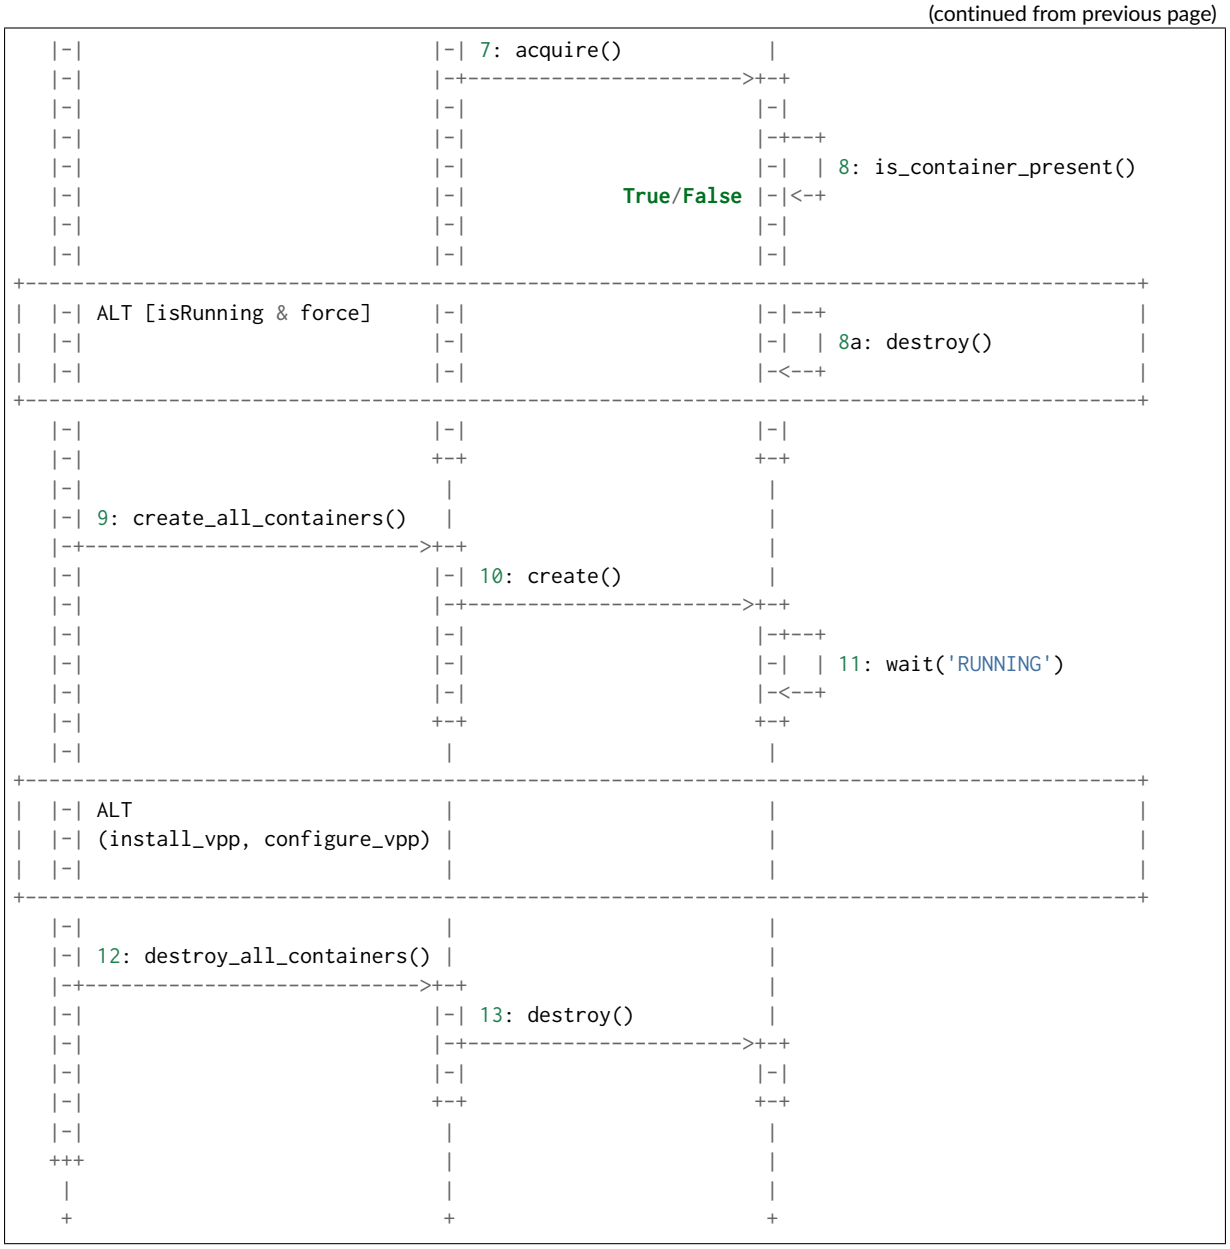

Container Data Structure

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

There is no parameters check functionality. Passing the correct arguments is a responsibility of the caller.

Examples

This section contains a high-level example of multiple initialization steps via ContainerManager; taken from an actual CSIT code, but with non-code lines (comments, Documentation) removed for brevity.

:

```
| Start containers for test
| | [Arguments] | ${dut}=${None} | ${nf_chains}=${1} | ${nf_nodes}=${1}
| | ... | ${auto_scale}=${True} | ${pinning}=${True}
| |
| | Set Test Variable | @{container_groups} | @{EMPTY}
| | Set Test Variable | ${container_group} | CNF
| | Set Test Variable | ${nf_nodes}
| | Import Library | resources.libraries.python.ContainerUtils.ContainerManager
| | ... | engine=${container_engine} | WITH NAME | ${container_group}
| | Construct chains of containers
| | ... | dut=${dut} | nf_chains=${nf_chains} | nf_nodes=${nf_nodes}
| | ... | auto_scale=${auto_scale} | pinning=${pinning}
| | Acquire all '${container_group}' containers
| | Create all '${container_group}' containers
| | Configure VPP in all '${container_group}' containers
| | Start VPP in all '${container_group}' containers
| | Append To List | ${container_groups} | ${container_group}
| | Save VPP PIDs
```

Kubernetes

For the future use, Kubernetes [k8sdoc] 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 `kubectl` installed on each of DUTs. One time initialization script, `resources/libraries/bash/k8s_setup.sh` does reset/init kubectl, 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.

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.

Making own python wrapper library of `kubectl` 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.

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-2206:

- LXC topologies:
  - eth-l2xcbase-eth-2memif-1lxc.
  - eth-l2bdbasemaclrn-eth-2memif-1lxc.

- Docker topologies:
  - eth-l2xcbase-eth-2memif-1docker.
  - eth-l2xcbase-eth-1memif-1docker

References

1.5.20 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 interfaces and connecting over memif (Slave side) to VPPs running in Linux Container (LXC) or in Docker Container (DRC) configured with memif (Master side). LXCs and DRCs run in a priviliged 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 73).

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 57).

1.5.21 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, 6, 8, 10).
- Number of NFs per service instance (e.g. 1, 2, 4, 6, 8, 10).

Implementation of NFV service density tests in CSIT-2206 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-2206 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-2206 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-2206 allocates cores within NUMA0. Support for multi NUMA tests is to be added in future release.

1.5.22 VPP_Device Functional

CSIT-2206 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. New device tests can be added by small edits to existing CSIT Performance (2-node) test. RF test definition code stays unchanged with the exception of traffic generator related L2 KWs.

1.5.23 Suite Generation

CSIT uses robot suite files to define tests. However, not all suite files available for Jenkins jobs (or manually started bootstrap scripts) are present in CSIT git repository. They are generated only when needed.

Autogen Library

There is a code generation layer implemented as Python library called “autogen”, called by various bash scripts.

It generates the full extent of CSIT suites, using the ones in git as templates.

Sources

The generated suites (and their contents) are affected by multiple information sources, listed below.

Git Suites

The suites present in git repository act as templates for generating suites. One of autogen design principles is that any template suite should also act as a full suite (no placeholders).

In practice, autogen always re-creates the template suite with exactly the same content, it is one of checks that autogen works correctly.

Regenerate Script

Not all suites present in CSIT git repository act as template for autogen. The distinction is on per-directory level. Directories with regenerate_testcases.py script usually consider all suites as templates (unless possibly not included by the glob patten in the script).

The script also specifies minimal frame size, indirectly, by specifying protocol (protocol “ip4” is the default, leading to 64B frame size).

Constants

Values in Constants.py are taken into consideration when generating suites. The values are mostly related to different NIC models and NIC drivers.

Python Code

Python code in resources/libraries/python/autogen contains several other information sources.

Testcase Templates

The test case part of template suite is ignored, test case lines are created according to text templates in Testcase.py file.

Testcase Argument Lists

Each testcase template has different number of “arguments”, e.g. values to put into various placeholders. Different test types need different lists of the argument values, the lists are in regenerate_glob method in Regenerator.py file.

Iteration Over Values

Python code detects the test type (usually by substrings of suite file name), then iterates over different quantities based on type. For example, only ndrpdr suite templates generate other types (mrr and soak).

Hardcoded Exclusions

Some combinations of values are known not to work, so they are excluded. Examples: Density tests for too much CPUs; IMIX for ASTF.

Non-Sources

Some information sources are available in CSIT repository, but do not affect the suites generated by autogen.

Testbeds

Overall, no information visible in topology yaml files is taken into account by autogen.

Testbed Architecture

Historically, suite files are agnostic to testbed architecture, e.g. ICX or ALT.

Testbed Size

Historically, 2-node and 3-node suites have diferent names, and while most of the code is common, the differences are not always simple enough. Autogen treat 2-node and 3-node suites as independent templates.

TRex suites are intended for a 1-node circuit of otherwise 2-node or 3-node testbeds, so they support all 3 robot tags. They are also detected and treated differently by autogen, mainly because they need different testcase arguments (no CPU count). Autogen does nothing specifically related to the fact they should run only in testbeds/NICs with TG-TG line available.

Other Topology Info

Some bonding tests need two (parallel) links between DUTs. Autogen does not care, as suites are agnostic. Robot tag marks the difference, but the link presence is not explicitly checked.

Job specs

Information in job spec files depend on generated suites (not the other way). Autogen should generate more suites, as job spec is limited by time budget. More suites should be available for manually triggered verify jobs, so autogen covers that.

Bootstrap Scripts

Historically, bootstrap scripts perform some logic, perhaps adding exclusion options to Robot invocation (e.g. skipping testbed+NIC combinations for tests that need parallel links).

Once again, the logic here relies on what autogen generates, autogen does not look into bootstrap scripts.

1.5.24 Testing in AWS EC2

AWS Performance Testbeds

CSIT implements two virtual machine topology types running in AWS EC2:

- **2-Node Topology**: Consists of one EC2 instance as a System Under Test (SUT) and one EC2 instance acting as a Traffic Generator (TG), with both instances connected into a ring topology. Used for executing tests that require frame encapsulations supported by TG.
- **3-Node Topology**: Consists of two EC2 instances acting as a Systems Under Test (SUTs) and one EC2 instance acting as a Traffic Generator (TG), with all instances 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.

AWS EC2 Instances

CSIT is using AWS EC2 C5n instances as System Under Test and TG virtual machines. C5n instances got selected to take advantage of high network throughput and packet rate performance. C5n instances offer up to 100 Gbps network bandwidth and increased memory over comparable C5 instances. For more information, see Instance types[52].

C5n features:

- 3.0 GHz Intel Xeon Platinum (Skylake) processors with Intel AVX-512 instructions.

[52] https://aws.amazon.com/ec2/instance-types/

- Sustained all core Turbo frequency of up to 3.4GHz, and single core turbo frequency of up to 3.5 GHz.
- Requires HVM AMIs (Amazon Machine Images) that include drivers for ENA and NVMe. See *CSIT Amazon Machine Images* (page 68) for more information.
- Network bandwidth to up to 100 Gbps.
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor.

Model	vCPU	Memory (GiB)	Instance Storage (GiB)	Network Bandwidth (Gbps)***	EBS Bandwidth (Mbps)
c5n.large	2	5.25	EBS-Only	Up to 25	Up to 4,750
c5n.xlarge	4	10.5	EBS-Only	Up to 25	Up to 4,750
c5n.2xlarge	8	21	EBS-Only	Up to 25	Up to 4,750
c5n.4xlarge	16	42	EBS-Only	Up to 25	4,750
c5n.9xlarge	36	96	EBS-Only	50	9,500
c5n.18xlarge	72	192	EBS-Only	100	19,000
c5n.metal	72	192	EBS-Only	100	19,000

CSIT is configured by default to use *c5n.4xlarge* in *eu-central-1* AWS region due to allocation stability issues with *c5n.9xlarge* in *eu-central-1* region.

AWS EC2 Networking

CSIT EC2 instances are equipped with AWS Elastic Network Adapter (ENA) supporting AWS enhanced networking. Enhanced networking uses single root I/O virtualization (SR-IOV) to provide high-performance networking capabilities. For more information, see Elastic Network Adapter[53].

For more information about the current advertised AWS ENA performance limits, see Computed optimized instances[54].

CSIT DUTs make use of AWS ENA DPDK driver supplied by AWS and specified in amzn drivers dpdk[55].

CSIT Amazon Machine Images

An Amazon Machine Image (AMI) provides the information required to launch an instance. CSIT is using Amazon Elastic Block Store (EBS) where the root device for an instance launched from the AMI is a volume created from an Amazon EBS snapshot.

As the TG and SUT instances have slightly different software requirements, we are defining two AMIs for TG and SUT separately. AMI details examples:

- TG:
  - AMI Name: csit_c5n_ubuntu_focal_tg
  - Platform details: Linux/UNIX
  - Architecture: x86_64
  - Usage operation: RunInstances
  - Image Type: machine
  - Virtualization type: hvm
  - Description: CSIT TG image based on Ubuntu Focal

[53] https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/enhanced-networking-ena.html
[54] https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/compute-optimized-instances.html
[55] https://github.com/amzn/amzn-drivers/tree/master/userspace/dpdk

  - Root Device Name: /dev/sda1
  - Root Device Type: ebs
- SUT:
  - AMI Name: csit_c5n_ubuntu_focal_sut
  - Platform details: Linux/UNIX
  - Architecture: x86_64
  - Usage operation: RunInstances
  - Image Type: machine
  - Virtualization type: hvm
  - Description: CSIT SUT image based on Ubuntu Focal
  - Root Device Name: /dev/sda1
  - Root Device Type: ebs

Both TG and SUT AMIs are created manually before launching topology and are not part of automated scripts.

Building AMIs requires Hashicorp Packer with Amazon plugin installed.

For more information, see Amazon Machine Images[56].

AWS Deployments

CSIT performance testbed deployments in AWS rely on Infrastructure-as-a-C (IaaC) Terraform AWS providers. Terraform providers specified in CSIT interact with resources provided by AWS to orchestrate virtual environment for running CSIT performance tests. For more information, see Terraform Registry aws[57].

Compatibility

Software	OSS Version
Terraform	1.0.3 or newer
Vault	1.8.4 or newer

Requirements

- Required Modules and Providers
  - Terraform Registry aws[58].
  - Terraform Registry null[59].
  - Terraform Registry tls[60].
  - Terraform Registry vault[61].
- Required software

[56] https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html
[57] https://registry.terraform.io/providers/hashicorp/aws/latest/
[58] https://registry.terraform.io/providers/hashicorp/aws/latest
[59] https://registry.terraform.io/providers/hashicorp/null/latest
[60] https://registry.terraform.io/providers/hashicorp/tls
[61] https://registry.terraform.io/providers/hashicorp/vault

- Vault[62] service available on specified ip/port.

Deployment Example

Following is an example of a Terraform deploy module[63] for a CSIT 2-Node testbed topology with AWS variables set to default values. A number of variables is also defined in a separate Terraform variable file[64].

```
module "deploy" {
  source = "./deploy"

  # Parameters starting with var. can be set using "TF_VAR_*" environment
  # variables or -var parameter when running "terraform apply", for default
  # values see ./variables.tf
  testbed_name          = var.testbed_name
  topology_name         = var.topology_name
  environment_name      = var.environment_name
  resources_name_prefix = var.resources_name_prefix

  # AWS general
  region        = var.region
  avail_zone    = var.avail_zone
  instance_type = var.instance_type
  ami_image_tg  = var.ami_image_tg
  ami_image_sut = var.ami_image_sut

  # AWS Network
  vpc_cidr_mgmt = "192.168.0.0/24"
  vpc_cidr_b    = "192.168.10.0/24"
  vpc_cidr_c    = "200.0.0.0/24"
  vpc_cidr_d    = "192.168.20.0/24"

  tg_mgmt_ip   = "192.168.0.10"
  dut1_mgmt_ip = "192.168.0.11"

  tg_if1_ip   = "192.168.10.254"
  tg_if2_ip   = "192.168.20.254"
  dut1_if1_ip = "192.168.10.11"
  dut1_if2_ip = "192.168.20.11"

  trex_dummy_cidr_port_0 = "10.0.0.0/24"
  trex_dummy_cidr_port_1 = "20.0.0.0/24"

  # Ansible
  ansible_python_executable = "/usr/bin/python3"
  ansible_file_path         = "../../fdio.infra.ansible/site.yaml"
  ansible_topology_path     = "../../fdio.infra.ansible/cloud_topology.yaml"
  ansible_provision_pwd     = "Csit1234"

  # First run
  first_run_commands = [
    "sudo sed -i 's/^PasswordAuthentication/#PasswordAuthentication/' /etc/ssh/sshd_config",
    "sudo systemctl restart sshd",
    "sudo useradd --create-home -s /bin/bash provisionuser",
    "echo 'provisionuser:Csit1234' | sudo chpasswd",
    "echo 'provisionuser ALL = (ALL) NOPASSWD: ALL' | sudo tee -a /etc/sudoers",
    "sudo useradd --create-home -s /bin/bash testuser",
    "echo 'testuser:Csit1234' | sudo chpasswd",
```

(continues on next page)

[62] https://releases.hashicorp.com/vault/
[63] https://git.fd.io/csit/tree/fdio.infra.terraform/2n_aws_c5n/main.tf
[64] https://git.fd.io/csit/tree/fdio.infra.terraform/2n_aws_c5n/variables.tf

(continued from previous page)

```
    "echo 'testuser ALL = (ALL) NOPASSWD: ALL' | sudo tee -a /etc/sudoers"
  ]
}
```

Secrets & Credentials

Set credentials manually

To set the credentials manually you first need to tell the module to not fetch credentials from Vault. To do that, set *provider "aws" access_key* and *secret_key* to custom value or use credentials file as a source.

```
provider "aws" {
  region     = var.region
  access_key = data.vault_aws_access_credentials.creds.access_key
  secret_key = data.vault_aws_access_credentials.creds.secret_key
}
```

1.6 Documentation

1.6.1 Test Code Documentation

CSIT Performance Tests Documentation[65] contains detailed functional description and input parameters for each test case.

[65] https://s3-docs.fd.io/csit/rls2206/docs/index.html

CHAPTER

TWO

VPP PERFORMANCE

2.1 Overview

VPP performance test results are reported for a range of processors. For description of physical testbeds used for VPP performance tests please refer to *Performance Physical Testbeds* (page 4).

2.1.1 Logical Topologies

CSIT VPP performance tests are executed on physical testbeds described in *Performance 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.

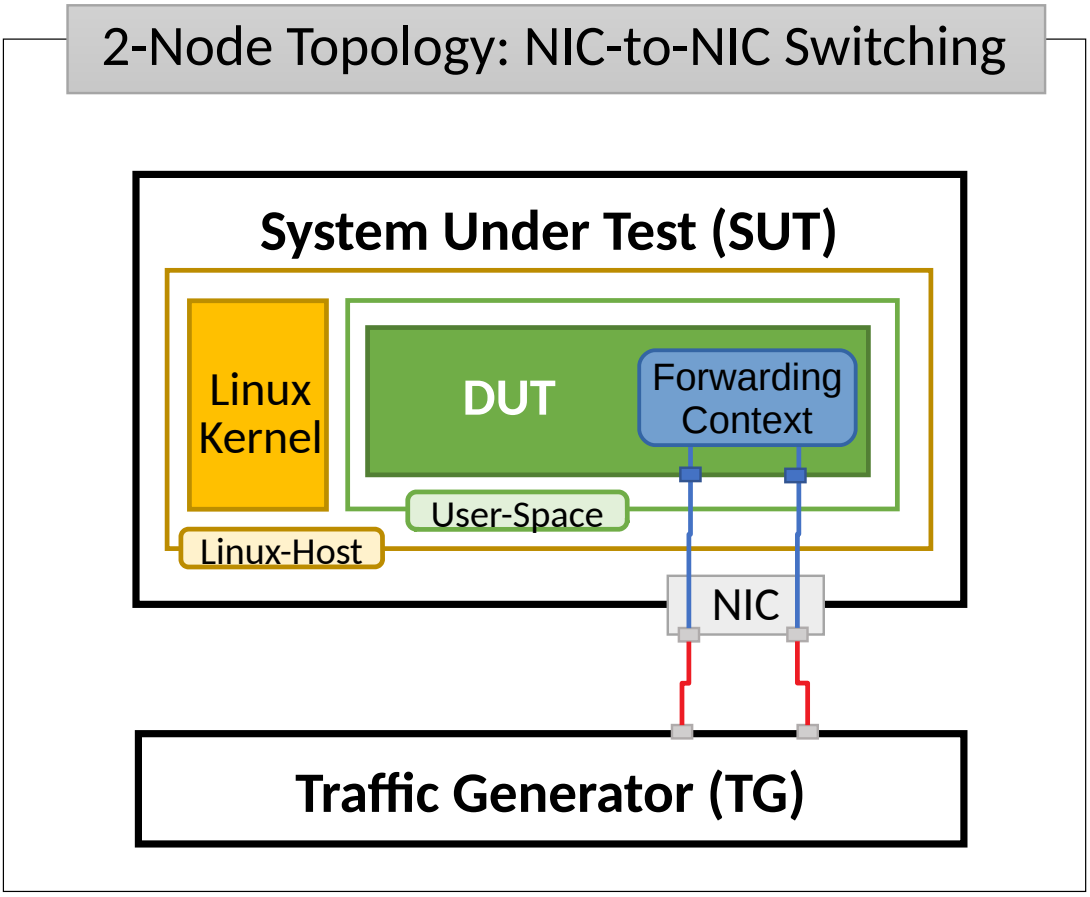

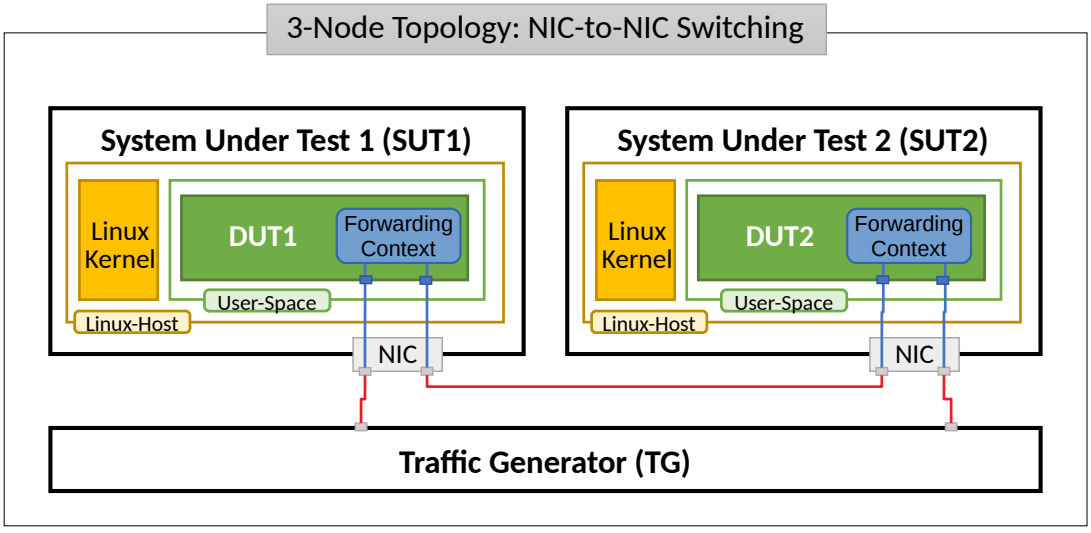

Server Systems Under Test (SUT) run VPP application in Linux user-mode as a Device Under Test (DUT). Server Traffic Generator (TG) runs T-Rex application. Physical connectivity between SUTs and TG is provided using different drivers and NIC models that need to be tested for performance (packet/bandwidth 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 limiations. 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 thoughput 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-2206:

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.

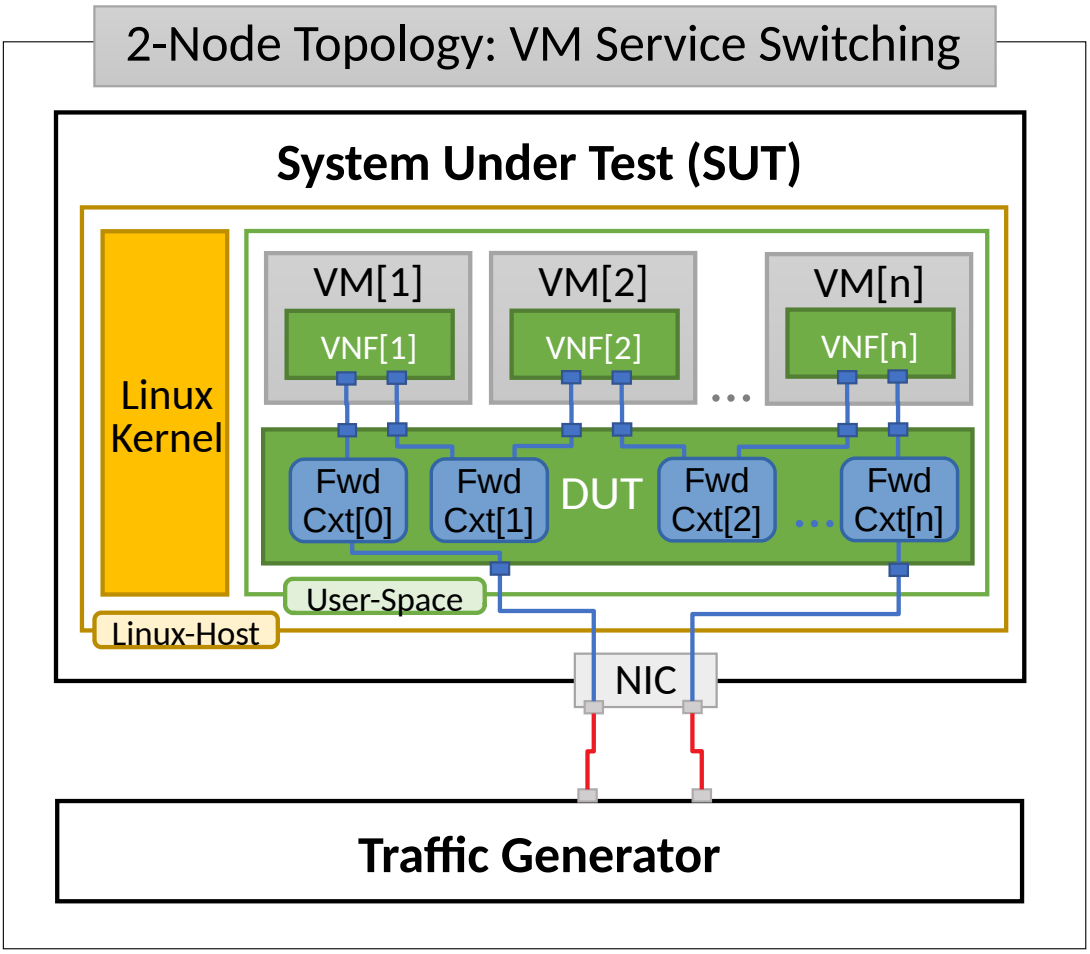

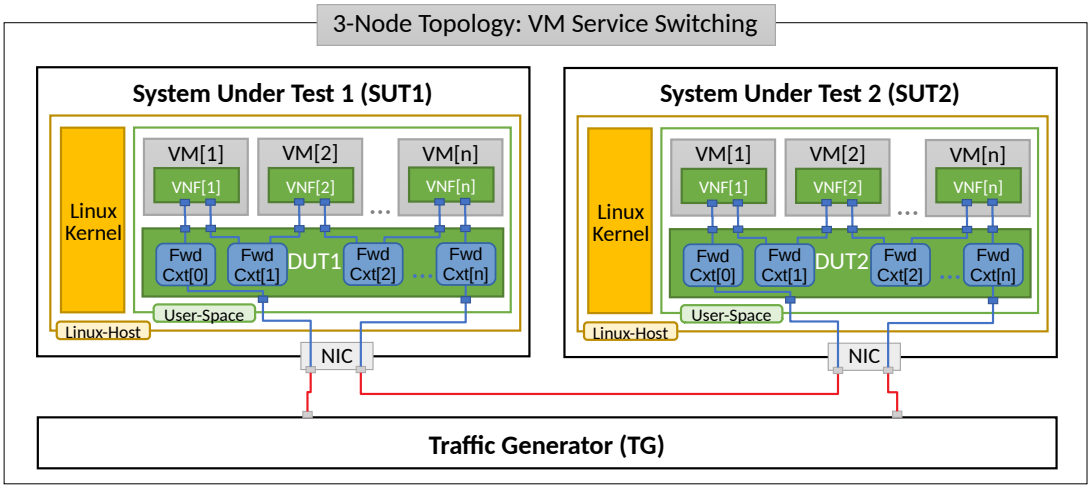

In “Chained” VM topologies, packets are switched by VPP DUT multiple times: twice for a single VM, three times for two VMs, N+1 times for N VMs. Hence the external throughput rates measured by TG and listed in this report must be multiplied by N+1 to represent the actual VPP DUT aggregate packet forwarding 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-2206:

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.

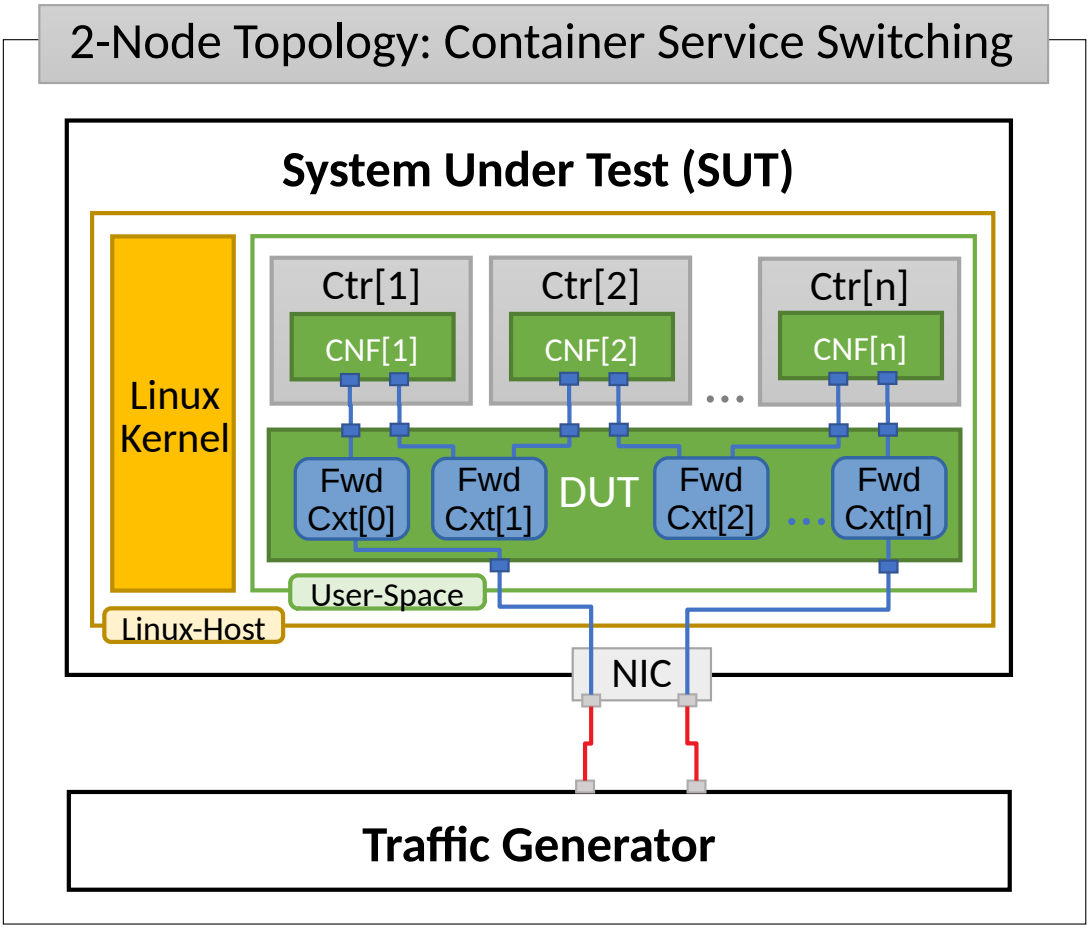

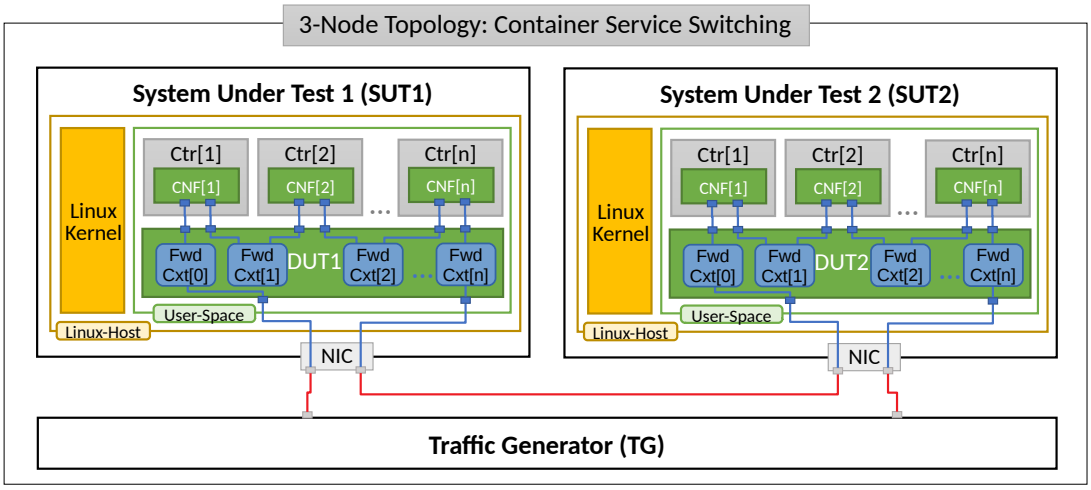

In “Chained” Container topologies, packets are switched by VPP DUT multiple times: twice for a single Container, three times for two Containers, N+1 times for N Containers. Hence the external throughput rates measured by TG and listed in this report must be multiplied by N+1 to represent the actual VPP 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[66], 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:
  - 90% of discovered PDR throughput.
  - 50% of discovered PDR throughput.
  - 10% of discovered PDR throughput.
  - Minimal offered load.
- 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, unless there is a known limitation preventing Traffic Generator from achieving the line rate.

CSIT-2206 includes following VPP data plane functionality performance tested across a range of NIC drivers and NIC models:

[66] https://tools.ietf.org/html/rfc2544.html

Functionality	Description
ACL	L2 Bridge-Domain switching and IPv4and IPv6 routing with iACL and oACL IP address, MAC address and L4 port security.
ADL	IPv4 and IPv6 routing with ADL address security.
GENEVE	GENEVE tunnels for IPv4 routing.
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.
IPSecAsyncHW	IPSec encryption with AES-GCM, CBC-SHA-256 ciphers in async mode, in combination with IPv4 routing. Intel QAT HW acceleration.
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.
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.
NAT44	(Source) Network Address Translation deterministic mode and endpoint-dependent mode tests with varying number of users and ports per user for IPv4.
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-2206 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 1743).

2.2 Release Notes

2.2.1 Changes in CSIT-2206

1. VPP PERFORMANCE TESTS
   - **Reduction of tests**: Removed certain test variations executed iteratively for the report (as well as in daily and weekly trending) due to physical testbeds overload.
2. TEST FRAMEWORK
   - **Removed ASTF PPS tests**: They provide no real benefit compared to TPUT tests. The ip4base variants renamed to TPUT.
   - **Changed TCP TPUT profiles**: The previous ones were found to be faulty. The new ones do not use bursts of packets to avoid CSIT-1830 and CSIT-1846.
   - **CSIT test environment** version has been updated to ver. 10, see *Environment Versioning* (page 1453).
   - **CSIT PAPI support**: Due to issues with PAPI performance, and deprecation of VAT, VPP CLI is used in CSIT for many VPP scale tests. See *Known Issues* (page 81).
   - **General Code Housekeeping**: Ongoing code optimizations and bug fixes.
3. PRESENTATION AND ANALYTICS LAYER
   - **Graphs improvements**: Updated Packet Latency graphs, see *Packet Latency* (page 45).

2.2.2 Known Issues

New

#	JiraID	Issue Description
1	CSIT-1827[67]	3n-icx, 3n-skx: all AVF crypto tests sporadically fail. 1518B with no traffic, IMIX with excessive packet loss.
2	CSIT-1830[68]	All testbeds: All TCP tput (and pps) tests are failing for small packets.
3	CSIT-1832[69]	3n-alt: NDR 1 packet lost on random tests.
4	CSIT-1834[70]	2n-icx, 2n-skx: sporadic AVF soak tests failing to find critical load with PLRsearch.
5	CSIT-1846[71]	2n-skx, 2n-clx, 2n-icx: ALL 1518B TCP tput tests failing with big packet loss.
6	CSIT-1847[72]	2n-skx: all 10vm-1t test failed with half of packets dropped.
7	CSIT-1849[73]	2n-skx: UDP 16m tput tests fail to create all sessions.

Previous

Issues reported in previous releases which still affect the current results.

[67] https://jira.fd.io/browse/CSIT-1827
[68] https://jira.fd.io/browse/CSIT-1830
[69] https://jira.fd.io/browse/CSIT-1832
[70] https://jira.fd.io/browse/CSIT-1834
[71] https://jira.fd.io/browse/CSIT-1846
[72] https://jira.fd.io/browse/CSIT-1847
[73] https://jira.fd.io/browse/CSIT-1849

#	JiraID	Issue Description
1	CSIT-1799[74]	All NAT44-ED 16M scale tests fail while setting NAT44 address range.
2	CSIT-1800[75]	All Geneve L3 mode scale tests (1024 tunnels) are failing.
3	CSIT-1802[76]	AF-XDP - NDR tests failing from time to time.
4	CSIT-1803[77]	3n-icx testbeds (Icelake): all IMIX aes128cbc-hmac512sha tests are failing due to excessive packet loss.
5	CSIT-1804[78]	3n-tsh, 3n-alt testbed (Taishan, Altra): NDR tests failing from time to time.
6	CSIT-1812[79]	All IMIX NAT44DET 4m 16m scale tests fail due to not creating required session count.
7	CSIT-1782[80]	Multicore AVF tests are failing when trying to create interface. Frequency is reduced by CSIT workaround, but occasional failures do still happen.
8	CSIT-1671[81] VPP-1763[82]	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. Currently, the time critical code uses VAT running large files with exec statements and CLI commands. Still, we needed to reduce the number of scale tests run to keep overall duration reasonable. More improvements needed to achieve sufficient configuration speed.
9	CSIT-1785[83] VPP-1972[84]	NAT44ED tests failing to establish all TCP sessions. At least for max scale, in allotted time (limited by session 500s timeout) due to worse slow path performance than previously measured and calibrated for. CSIT removed the max scale NAT tests to avoid this issue.
10	CSIT-1801[85]	9000B payload frames not forwarded over tunnels due to violating supported Max Frame Size (VxLAN, LISP, SRv6).
11	CSIT-1808[86]	All tests with 9000B payload frames not forwarded over memif interfaces.
12	CSIT-1809[87]	All tests with 9000B payload frames not forwarded over vhostuser interfaces.

[74] https://jira.fd.io/browse/CSIT-1799
[75] https://jira.fd.io/browse/CSIT-1800
[76] https://jira.fd.io/browse/CSIT-1802
[77] https://jira.fd.io/browse/CSIT-1803
[78] https://jira.fd.io/browse/CSIT-1804
[79] https://jira.fd.io/browse/CSIT-1812
[80] https://jira.fd.io/browse/CSIT-1782
[81] https://jira.fd.io/browse/CSIT-1671
[82] https://jira.fd.io/browse/VPP-1763
[83] https://jira.fd.io/browse/CSIT-1785
[84] https://jira.fd.io/browse/VPP-1972
[85] https://jira.fd.io/browse/CSIT-1801
[86] https://jira.fd.io/browse/CSIT-1808
[87] https://jira.fd.io/browse/CSIT-1809

Fixed

Issues reported in previous releases which were fixed in this release:

#	JiraID	Issue Description
1	CSIT-1810[88]	DPDK - performance regression with DPDK driver when Max Frame Size is set to less than 2023. Worse performance with DPDK driver when MTU is set to 2022 or less.
	VPP-1876[89]	
2	CSIT-1811[90]	All 9000B NAT44DET 64k 1m scale tests fail due to bps rate set to high on TRex.
3	CSIT-1791[91]	Performance regression in RDMA tests, due to CSIT environment changes. Two symptoms: 1. 10-20% regression across most tests. 2. DUT performance cap just below 38 Mpps.

2.2.3 Root Cause Analysis for Performance Changes

List of RCAs in CSIT-2206 for VPP performance changes:

#	JiraID	Issue Description
1	CSIT-1851[92]	trending regression: various icelake tests around 2202-04-15 Somewhat expected consequence of a VPP usability fix, the previous VPP compiler version was too new for the OS used.
2	VPP-2030[93]	regression: ip6base on ICX around 2022-03-23 “Loads blocked due to overlapping with a preceding store that cannot be forwarded.” started happening in ip6-lookup graph node.
3	CSIT-1852[94]	2n-zn2 mellanox performance cap Old issue, only now distinguished from CSIT-1751. This testbed+nic combination is capped below 28 Mpps, cause not identified yet.
4	CSIT-1853[95]	trending regression: nat44ed cps around 2202-04-01 VPP change added more computation to slow path (in order to support multiple VRFs). Not clear if the VPP implementation is optimized enough.

[88] https://jira.fd.io/browse/CSIT-1810
[89] https://jira.fd.io/browse/VPP-1876
[90] https://jira.fd.io/browse/CSIT-1811
[91] https://jira.fd.io/browse/CSIT-1791
[92] https://jira.fd.io/browse/CSIT-1851
[93] https://jira.fd.io/browse/VPP-2030
[94] https://jira.fd.io/browse/CSIT-1852
[95] https://jira.fd.io/browse/CSIT-1853

2.3 Packet Throughput

Throughput graphs are generated based on the results data obtained from the CSIT-2206 test jobs. In order to verify benchmark results repeatibility 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[96] files.

Graphs are split into sections as follows:

1. **Header 1**: VPP packet path and lookup types
    - **L2 Ethernet Switching**: L2 bridge-doman, 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
2. **Header 2**: testbeds and NIC models
    - section name format:
        - {**testbed_type**}-{**nic_model**}
    - **testbed_type**:
        - 2n-icx: 2-node Xeon Icelake
        - 3n-icx: 3-node Xeon Icelake
        - 2n-aws: 2-node AWS
        - 2n-skx: 2-node Xeon Skylake
        - 3n-skx: 3-node Xeon Skylake
        - 2n-clx: 2-node Xeon Cascade Lake
        - 2n-zn2: 2-node AMD Zen2
        - 3n-alt: 2-node Arm Altra
        - 3n-tsh: 3-node Arm TaiShan
        - 2n-tx2: 2-node Arm ThunderX2
        - 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)

[96] https://git.fd.io/csit/tree/docs/job_specs

  - x553: x553 2p10GE Intel (Niantic)
  - cx556a: cx556a-edat 2p100GE Mellanox ConnectX5
  - e810cq: E810-2CQDA2 2p100GE Intel Columbiaville

3. **Header 3**: test group names
   - section name format:
     - {**frame_size**}-{**worker_thread_core_cfg**}-{**vpp_functionality**}-{**vpp_lookup_type**}-{**baseline_scale**}-{**nic_driver**}
   - **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×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 are stored in build logs from FD.io vpp performance job 2n-icx[97], build logs from FD.io vpp performance job 3n-icx[98], build logs from FD.io vpp performance job 2n-aws[99], build logs from FD.io vpp performance job 2n-skx[100], build logs from FD.io vpp performance job 3n-skx[101], build logs from FD.io vpp performance job 2n-clx[102], build logs from FD.io vpp performance job 2n-zn2[103], build logs from FD.io vpp performance job 3n-alt[104], build logs from FD.io vpp performance job 3n-tsh[105], build logs from FD.io vpp performance job 2n-tx2[106], build logs from FD.io vpp performance job 2n-dnv[107] and build logs from FD.io vpp performance job 3n-dnv[108] with RF result files csit-vpp-perf-2206-*.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.

---

[97] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-icx
[98] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-icx
[99] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-aws
[100] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx
[101] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-skx
[102] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-clx
[103] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-zn2
[104] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-alt
[105] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-tsh
[106] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-tx2
[107] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-dnv
[108] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-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[109].

[109] https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls2206

2n-icx-xxv710

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

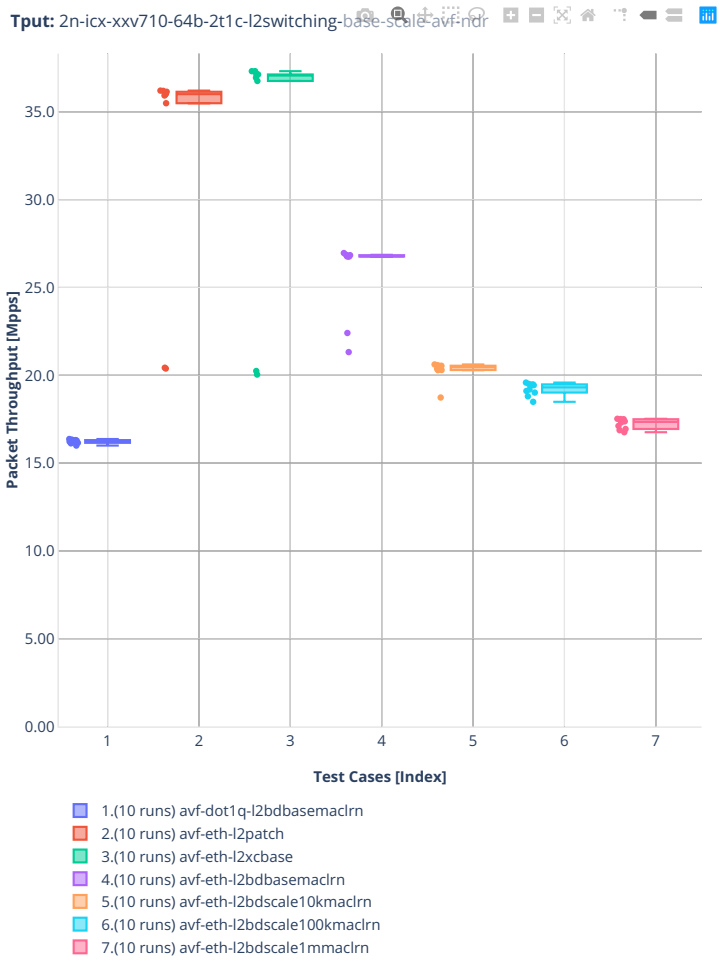

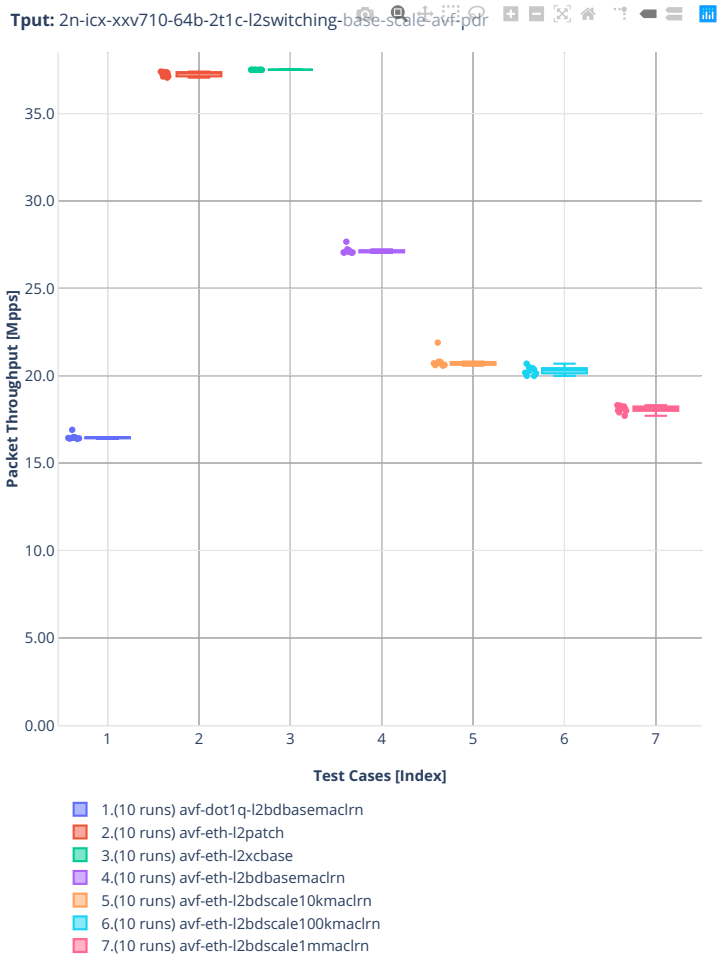
Tput: 2n-icx-xxv710-64b-2t1c-l2switching-base-scale-avf-pdr
Packet Throughput [Mpps]
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
1
2
3
4
5
6
7
Test Cases [Index]
1.(10 runs) avf-dot1q-l2bdbasemaclrn
2.(10 runs) avf-eth-l2patch
3.(10 runs) avf-eth-l2xcbase
4.(10 runs) avf-eth-l2bdbasemaclrn
5.(10 runs) avf-eth-l2bdscale10kmaclrn
6.(10 runs) avf-eth-l2bdscale100kmaclrn
7.(10 runs) avf-eth-l2bdscale1mmaclrn

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

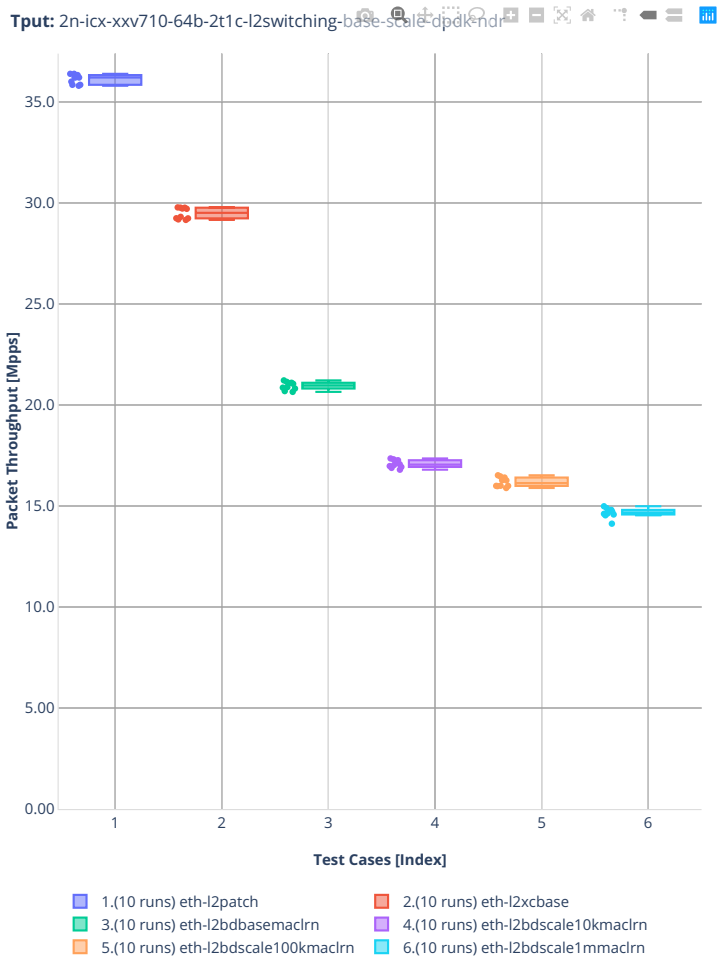

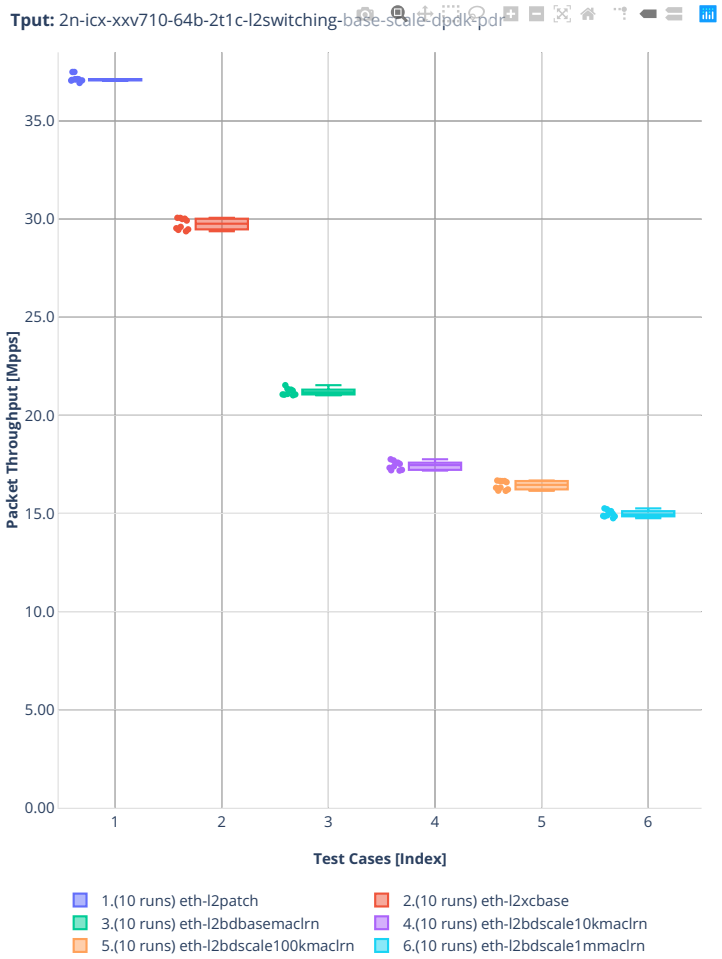
Tput: 2n-icx-xxv710-64b-2t1c-l2switching-base-scale-dpdk-pdr
Packet Throughput [Mpps]
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) eth-l2patch
2.(10 runs) eth-l2xcbase
3.(10 runs) eth-l2bdbasemaclrn
4.(10 runs) eth-l2bdscale10kmaclrn
5.(10 runs) eth-l2bdscale100kmaclrn
6.(10 runs) eth-l2bdscale1mmaclrn

3n-icx-xxv710

64b-2t1c-l2switching-base

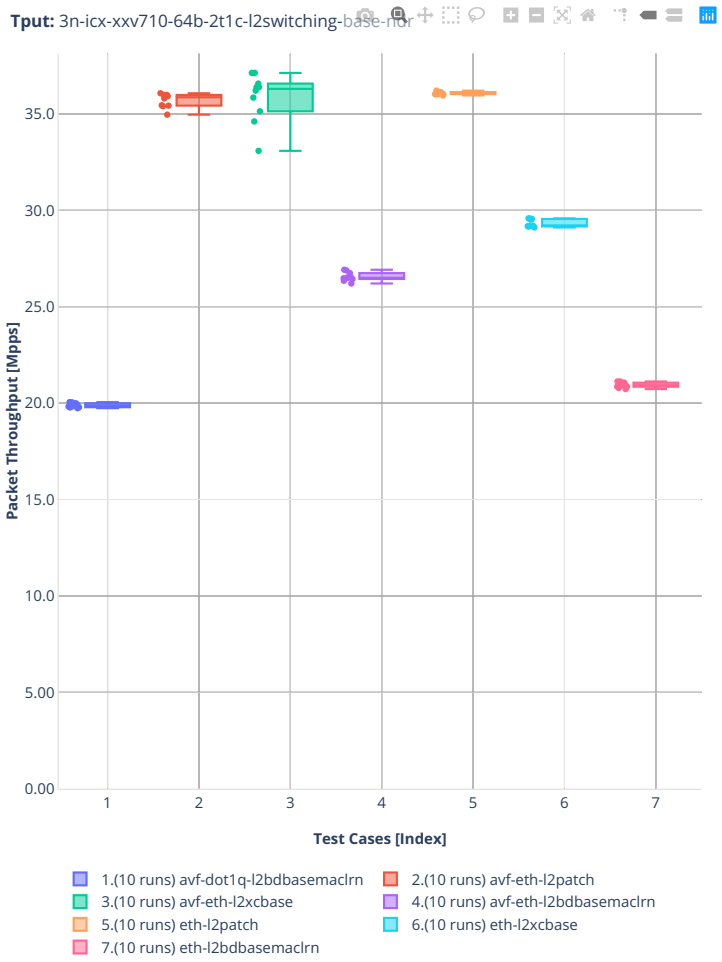

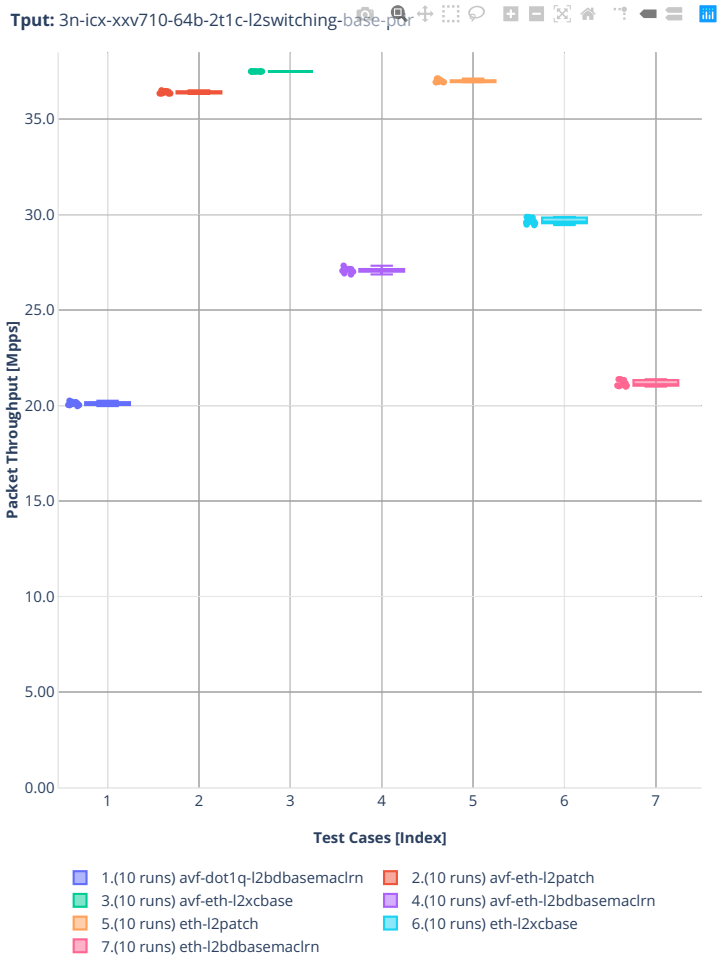
Tput: 3n-icx-xxv710-64b-2t1c-l2switching-base-pdr
Packet Throughput [Mpps]
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
1
2
3
4
5
6
7
Test Cases [Index]
1.(10 runs) avf-dot1q-l2bdbasemaclrn
2.(10 runs) avf-eth-l2patch
3.(10 runs) avf-eth-l2xcbase
4.(10 runs) avf-eth-l2bdbasemaclrn
5.(10 runs) eth-l2patch
6.(10 runs) eth-l2xcbase
7.(10 runs) eth-l2bdbasemaclrn

2n-skx-xxv710

64b-2t1c-l2switching-base-avf

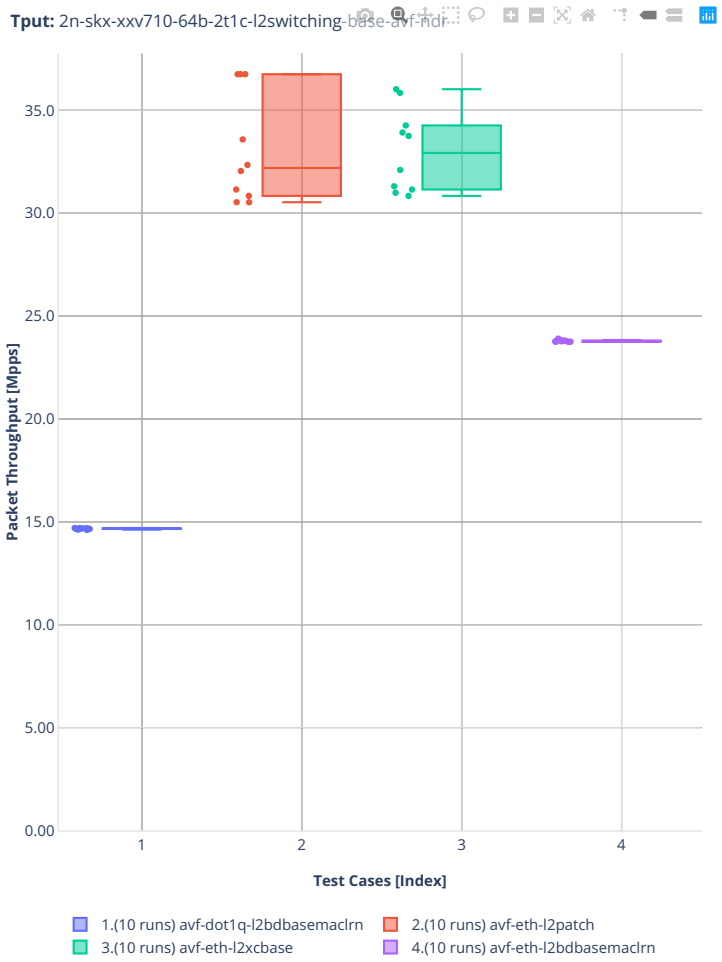

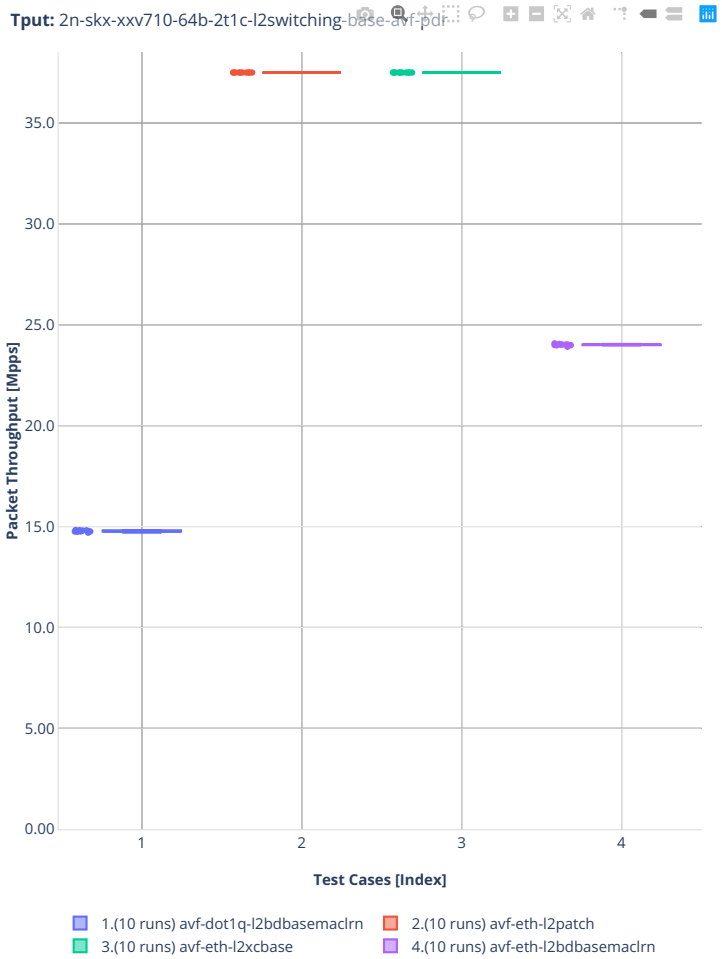
Tput: 2n-skx-xxv710-64b-2t1c-l2switching-base-avf-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) avf-dot1q-l2bdbasemaclrn
2.(10 runs) avf-eth-l2patch
3.(10 runs) avf-eth-l2xcbase
4.(10 runs) avf-eth-l2bdbasemaclrn

64b-2t1c-l2switching-base-dpdk

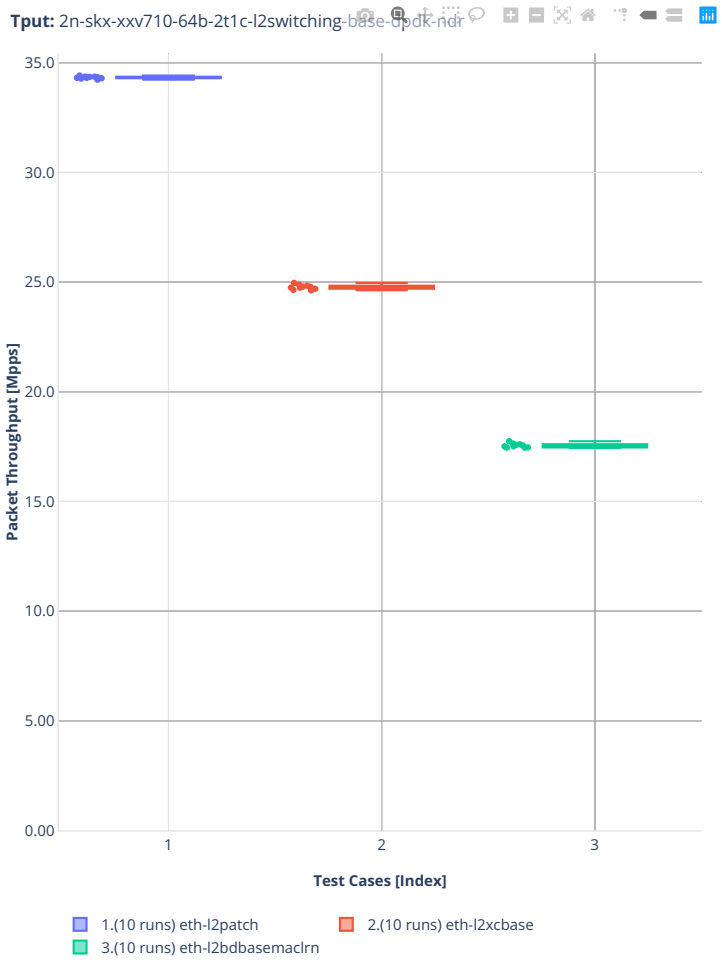

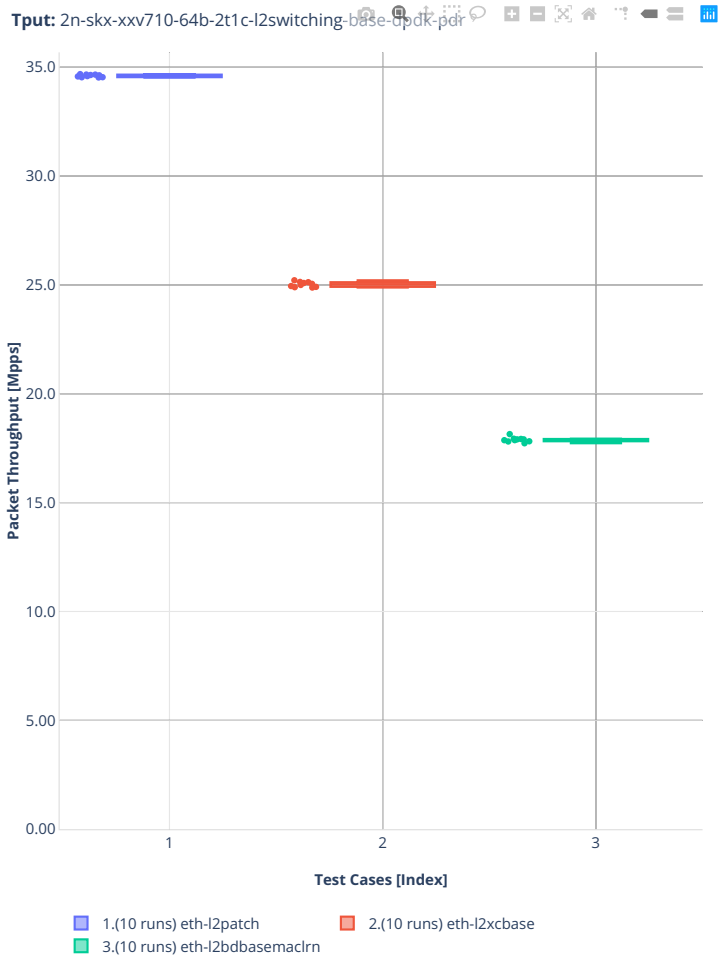
Tput: 2n-skx-xxv710-64b-2t1c-l2switching-base-dpdk-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) eth-l2patch
2.(10 runs) eth-l2xcbase
3.(10 runs) eth-l2bdbasemaclrn

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

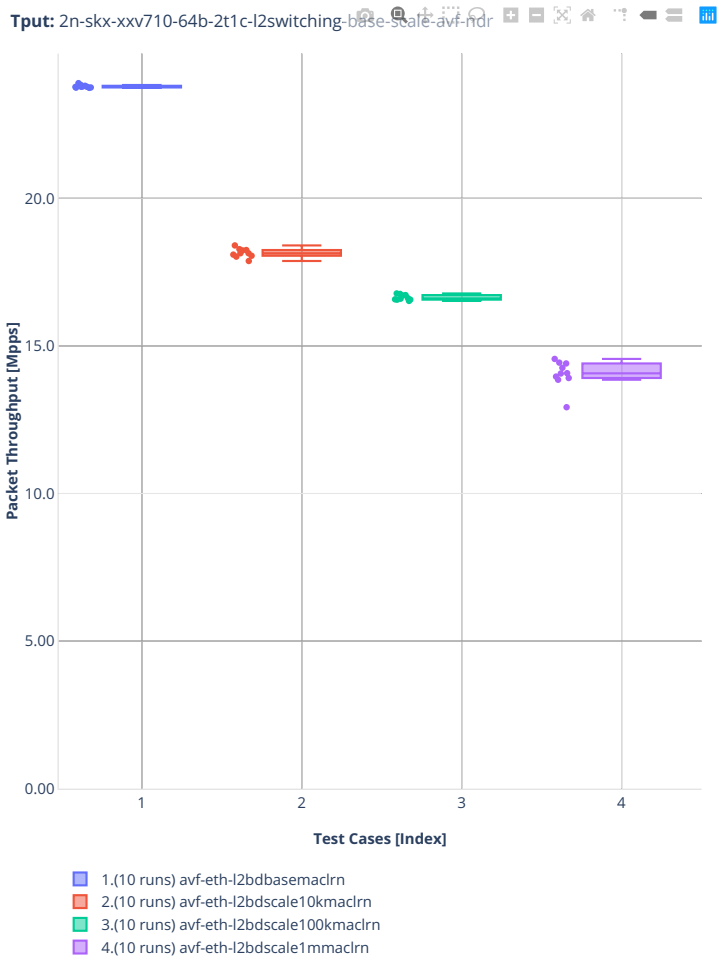

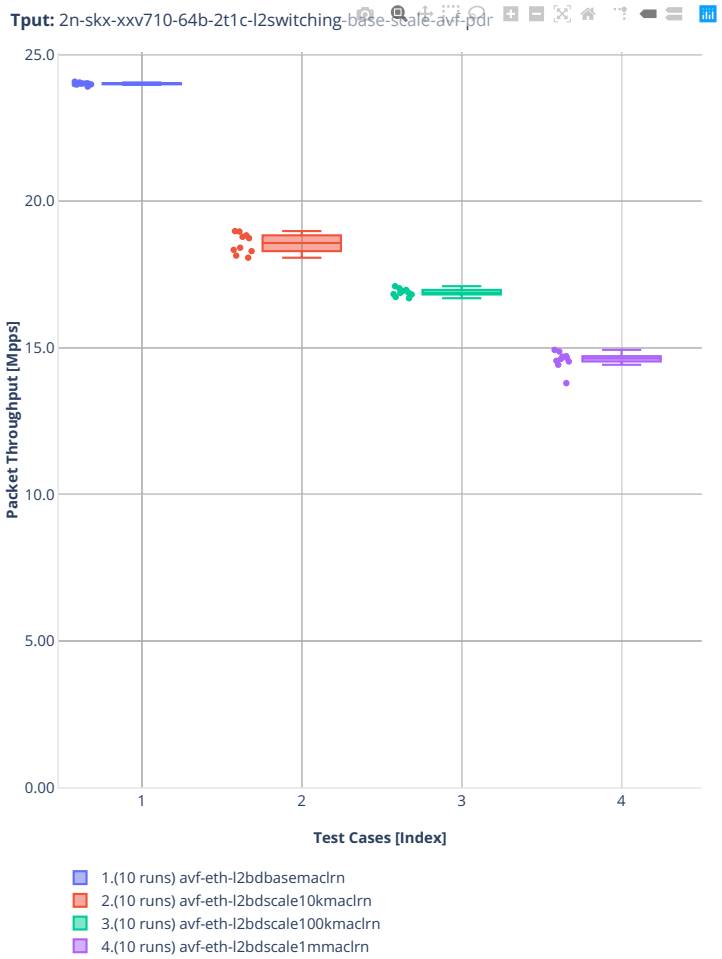
Tput: 2n-skx-xxv710-64b-2t1c-l2switching-base-scale-avf-pdr
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) avf-eth-l2bdbasemaclrn
2.(10 runs) avf-eth-l2bdscale10kmaclrn
3.(10 runs) avf-eth-l2bdscale100kmaclrn
4.(10 runs) avf-eth-l2bdscale1mmaclrn

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

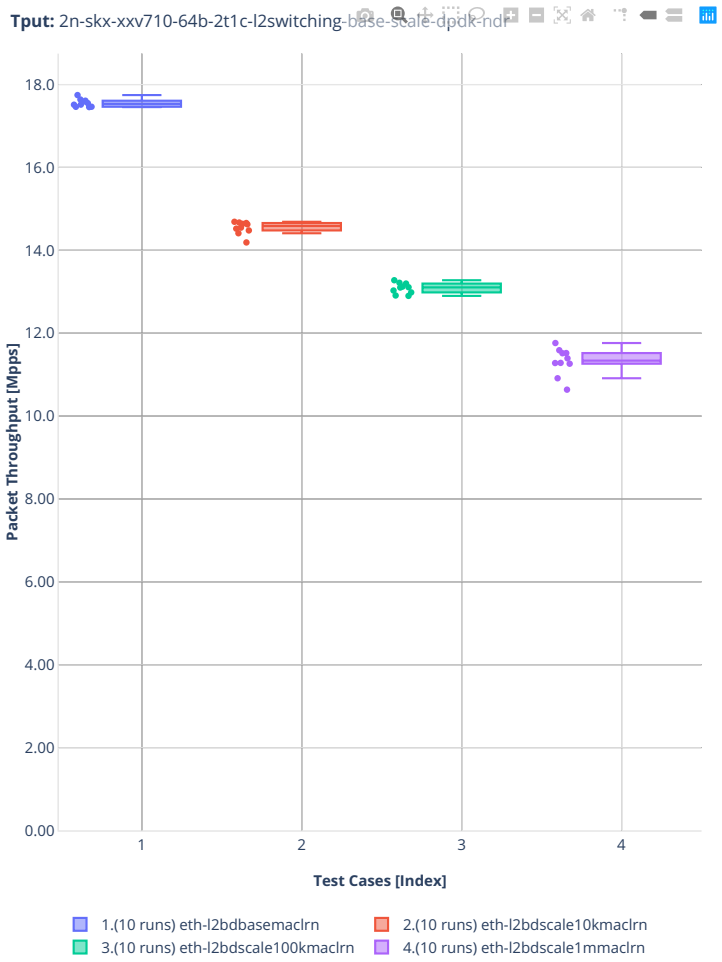

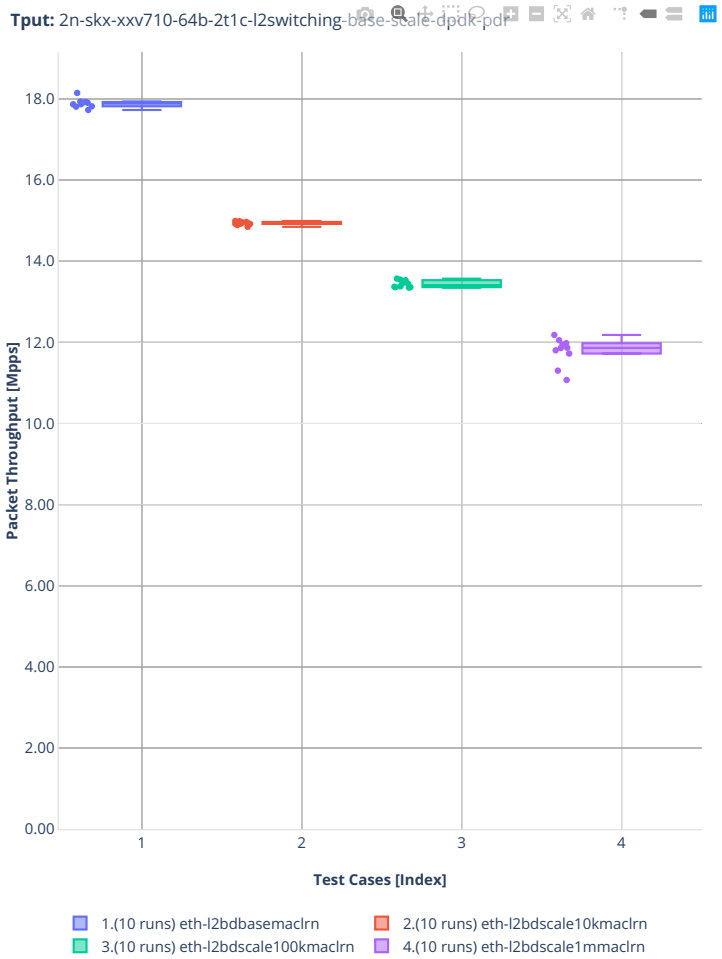
Tput: 2n-skx-xxv710-64b-2t1c-l2switching-base-scale-dpdk-pdr
Packet Throughput [Mpps]
18.0
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
Test Cases [Index]
1.(10 runs) eth-l2bdbasemaclrn
2.(10 runs) eth-l2bdscale10kmaclrn
3.(10 runs) eth-l2bdscale100kmaclrn
4.(10 runs) eth-l2bdscale1mmaclrn

2n-skx-x710

64b-2t1c-l2switching-base-scale-[avf,dpdk]

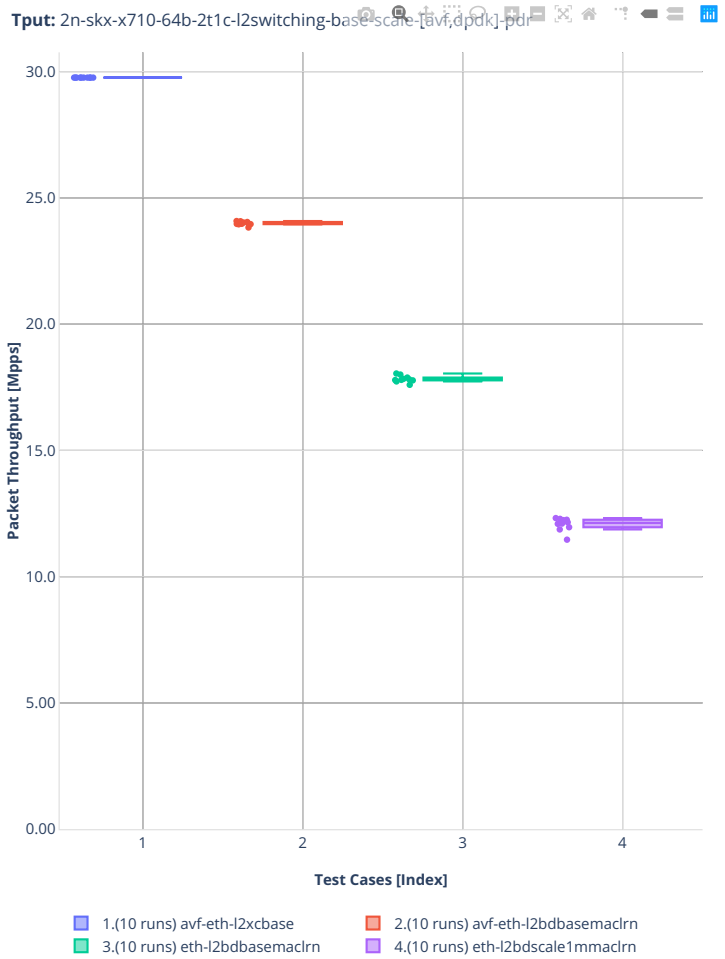
Tput: 2n-skx-x710-64b-2t1c-l2switching-base-scale-[avf,dpdk]-pdr
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) avf-eth-l2xcbase
2.(10 runs) avf-eth-l2bdbasemaclrn
3.(10 runs) eth-l2bdbasemaclrn
4.(10 runs) eth-l2bdscale1mmaclrn

3n-skx-xxv710

64b-2t1c-l2switching-base

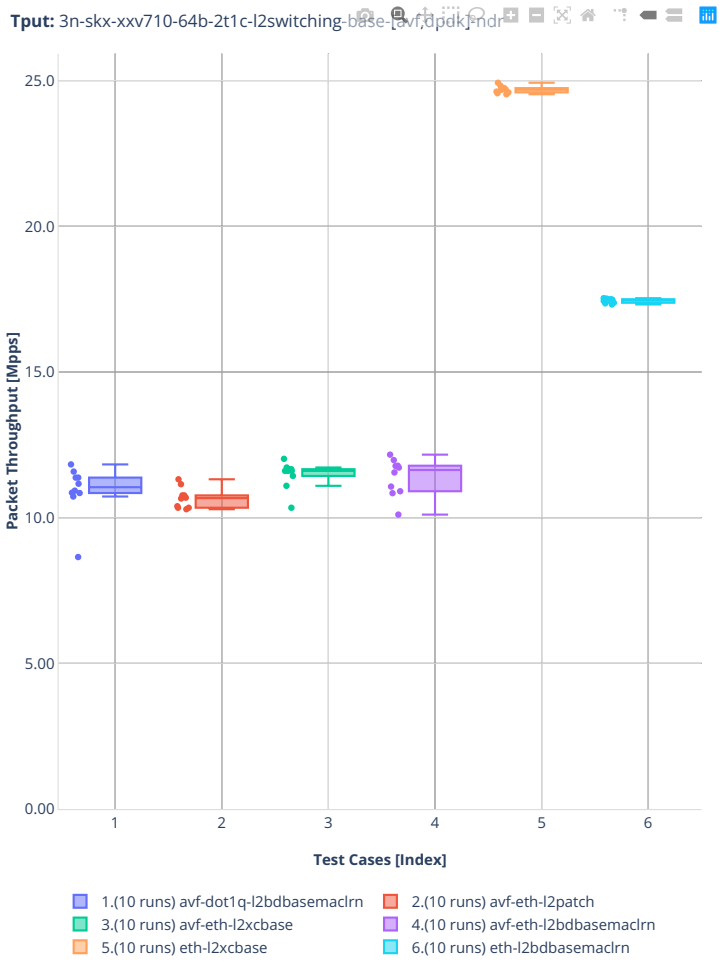

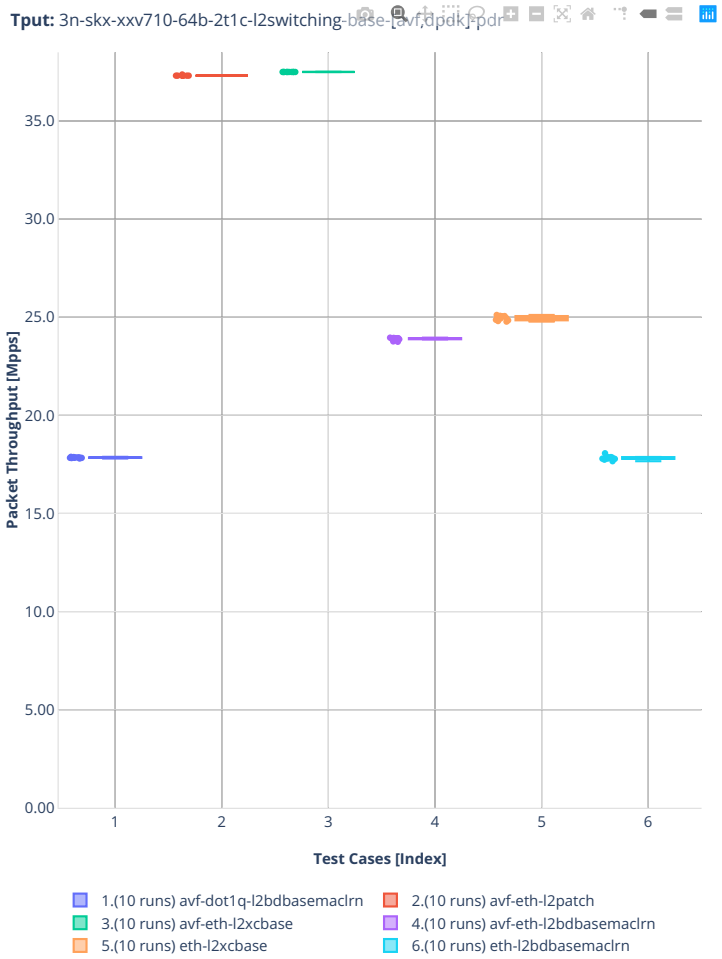
Tput: 3n-skx-xxv710-64b-2t1c-l2switching-base-[avf,dpdk]-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-dot1q-l2bdbasemaclrn
2.(10 runs) avf-eth-l2patch
3.(10 runs) avf-eth-l2xcbase
4.(10 runs) avf-eth-l2bdbasemaclrn
5.(10 runs) eth-l2xcbase
6.(10 runs) eth-l2bdbasemaclrn

3n-skx-x710

64b-2t1c-l2switching-base-avf

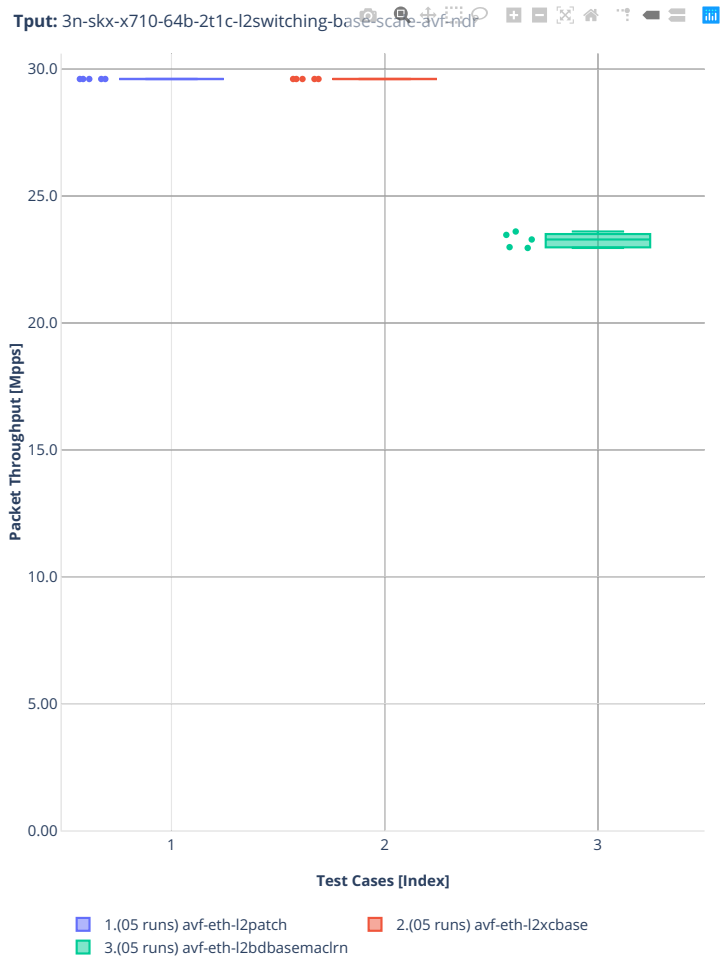

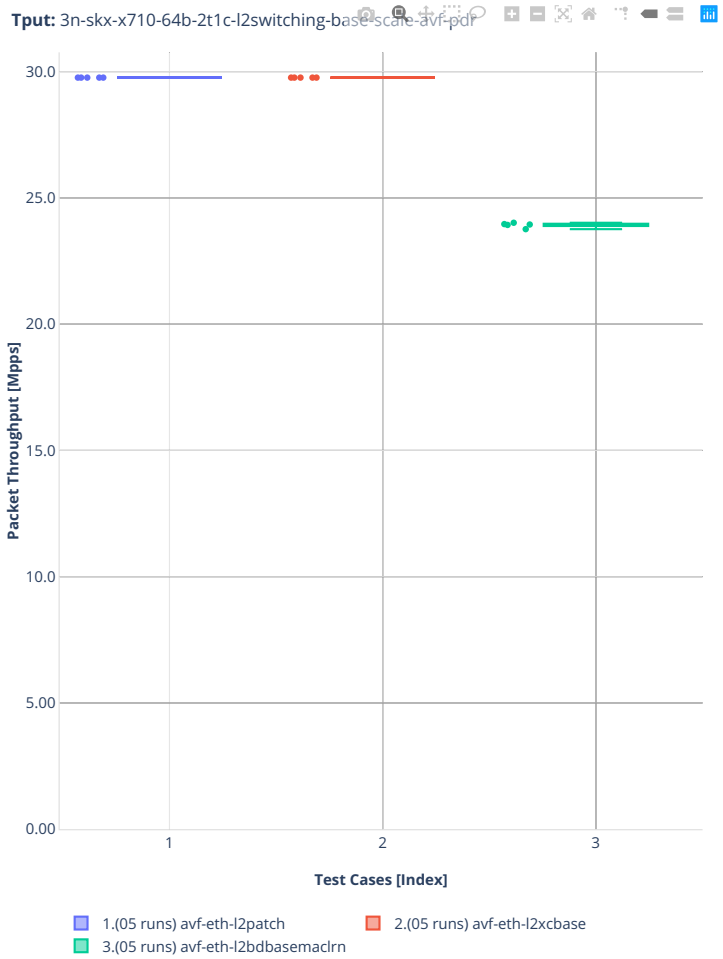
Tput: 3n-skx-x710-64b-2t1c-l2switching-base-scale-avf-pdr
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(05 runs) avf-eth-l2patch
2.(05 runs) avf-eth-l2xcbase
3.(05 runs) avf-eth-l2bdbasemaclrn

2n-clx-xxv710

64b-2t1c-l2switching-base-avf

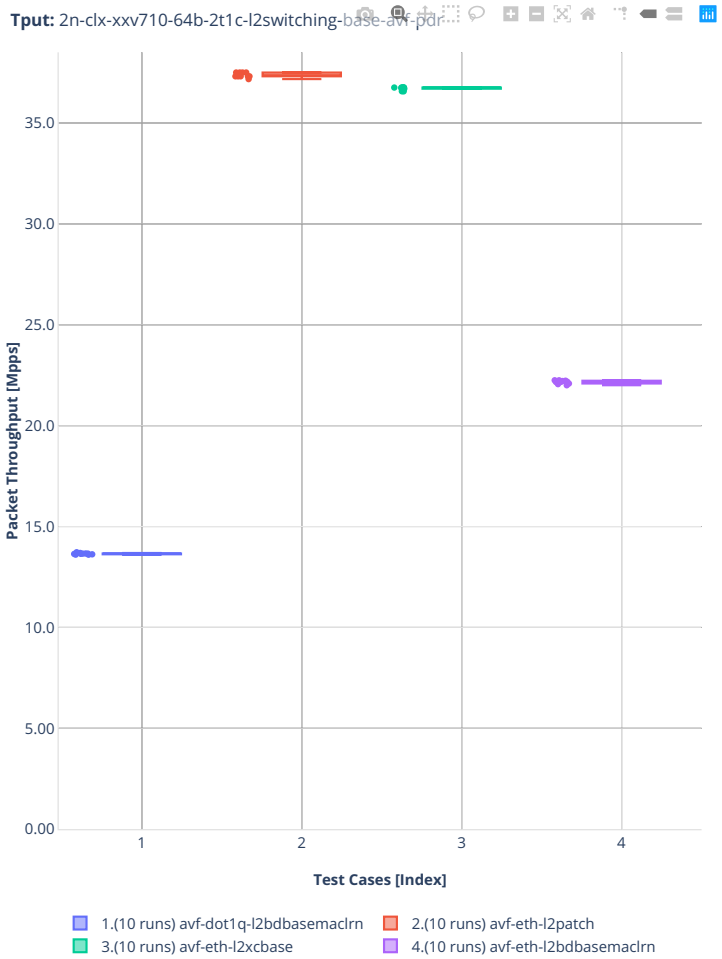
Tput: 2n-clx-xxv710-64b-2t1c-l2switching-base-avf-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) avf-dot1q-l2bdbasemaclrn
2.(10 runs) avf-eth-l2patch
3.(10 runs) avf-eth-l2xcbase
4.(10 runs) avf-eth-l2bdbasemaclrn

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

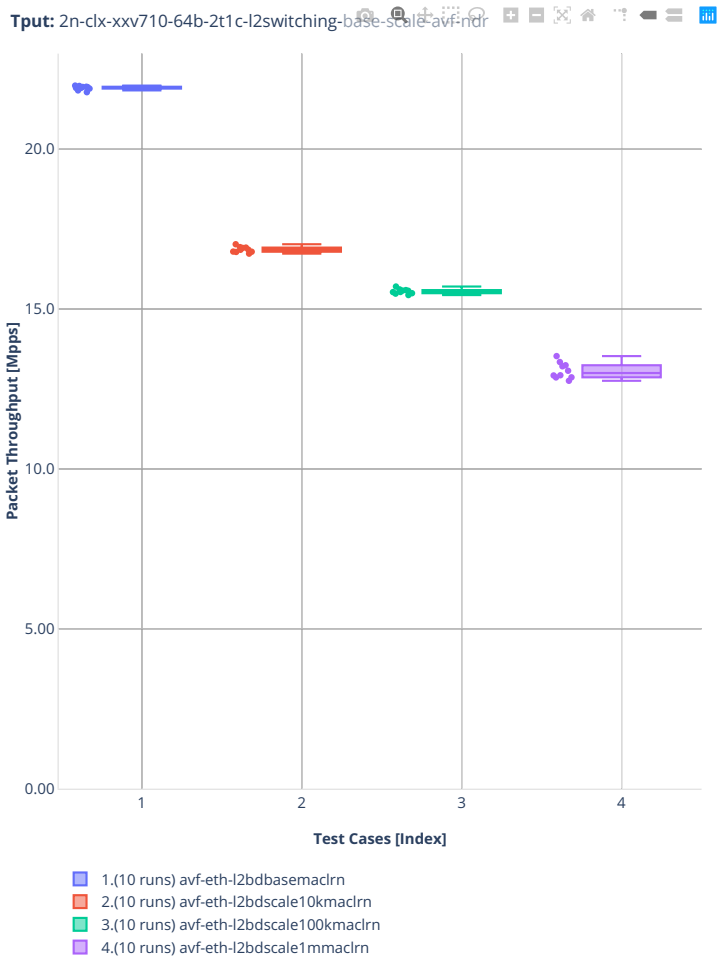

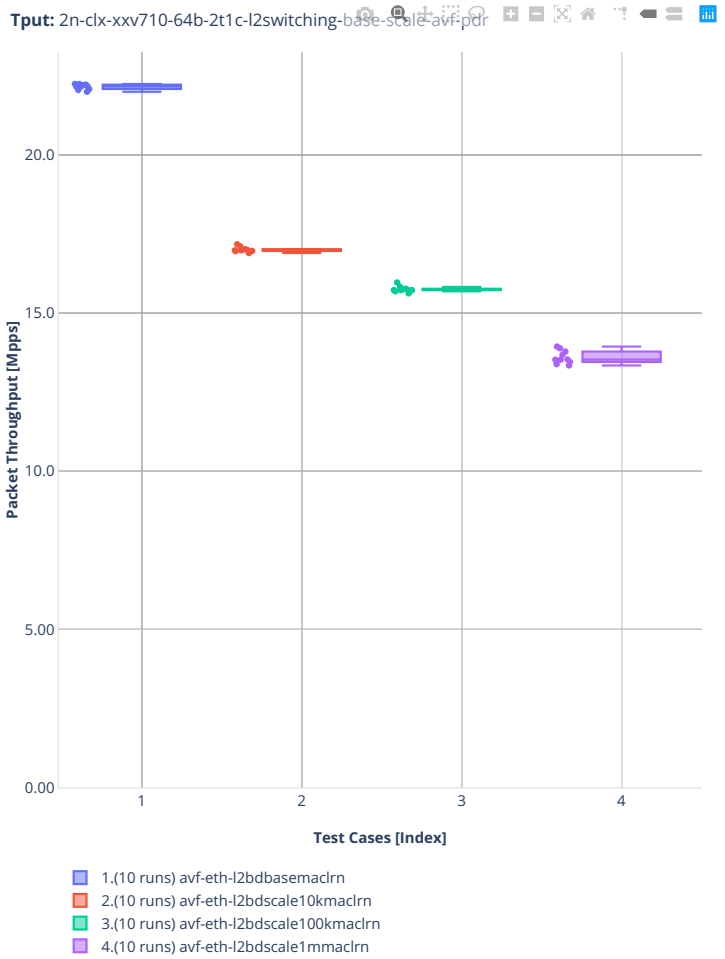
Tput: 2n-clx-xxv710-64b-2t1c-l2switching-base-scale-avf-pdr
Packet Throughput [Mpps]
20.0
15.0
10.0
5.00
0.00
1
2
3
4
Test Cases [Index]
1.(10 runs) avf-eth-l2bdbasemaclrn
2.(10 runs) avf-eth-l2bdscale10kmaclrn
3.(10 runs) avf-eth-l2bdscale100kmaclrn
4.(10 runs) avf-eth-l2bdscale1mmaclrn

64b-2t1c-l2switching-base-dpdk

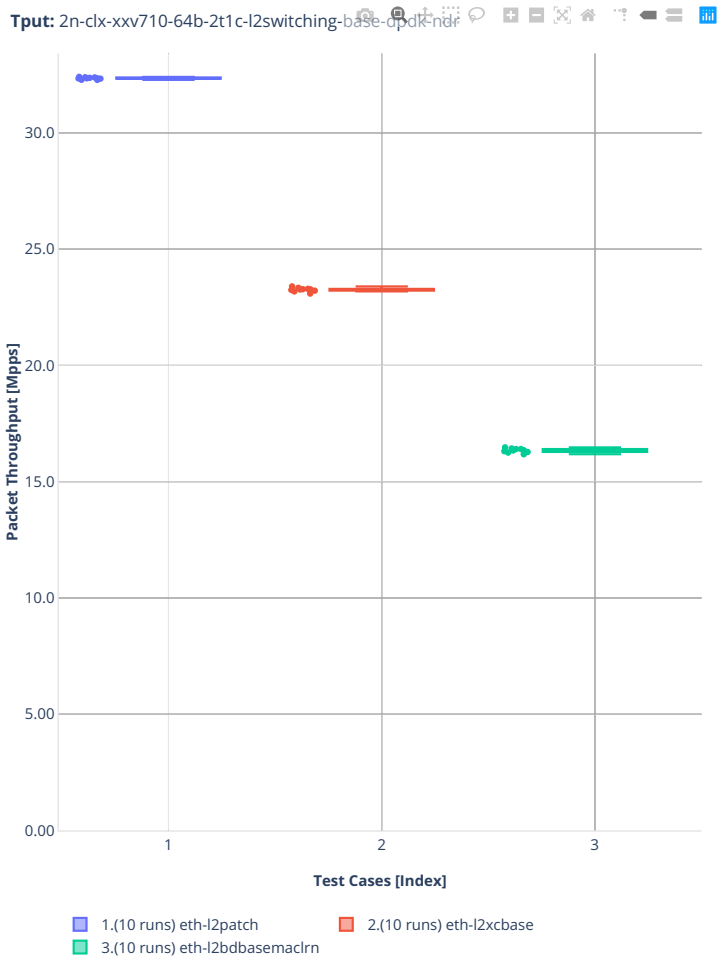

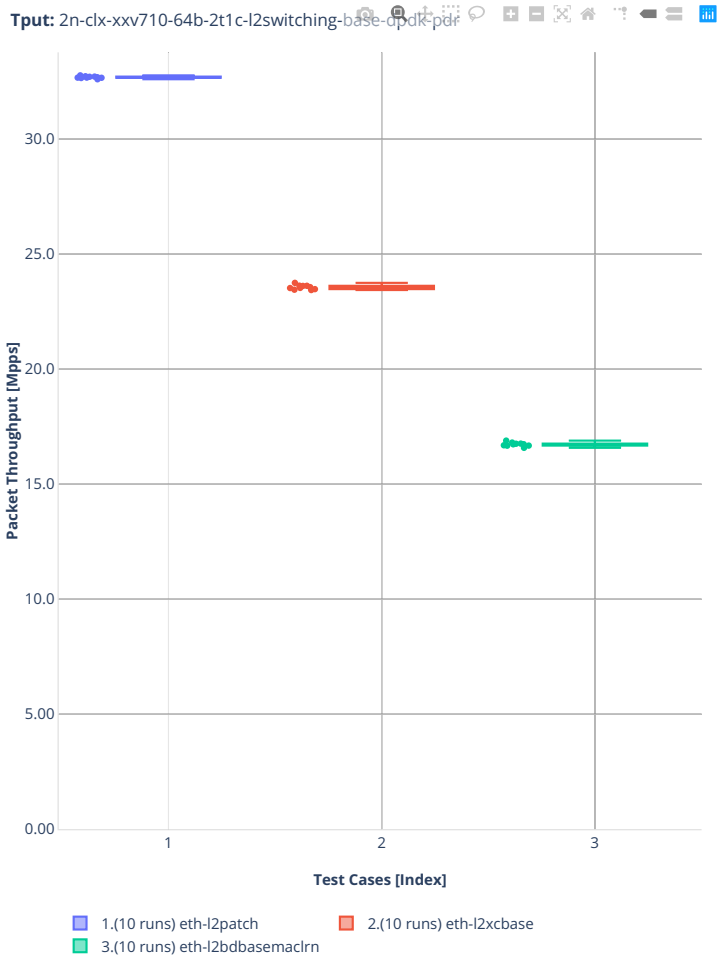
Tput: 2n-clx-xxv710-64b-2t1c-l2switching-base-dpdk-pdr
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) eth-l2patch
2.(10 runs) eth-l2xcbase
3.(10 runs) eth-l2bdbasemaclrn

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

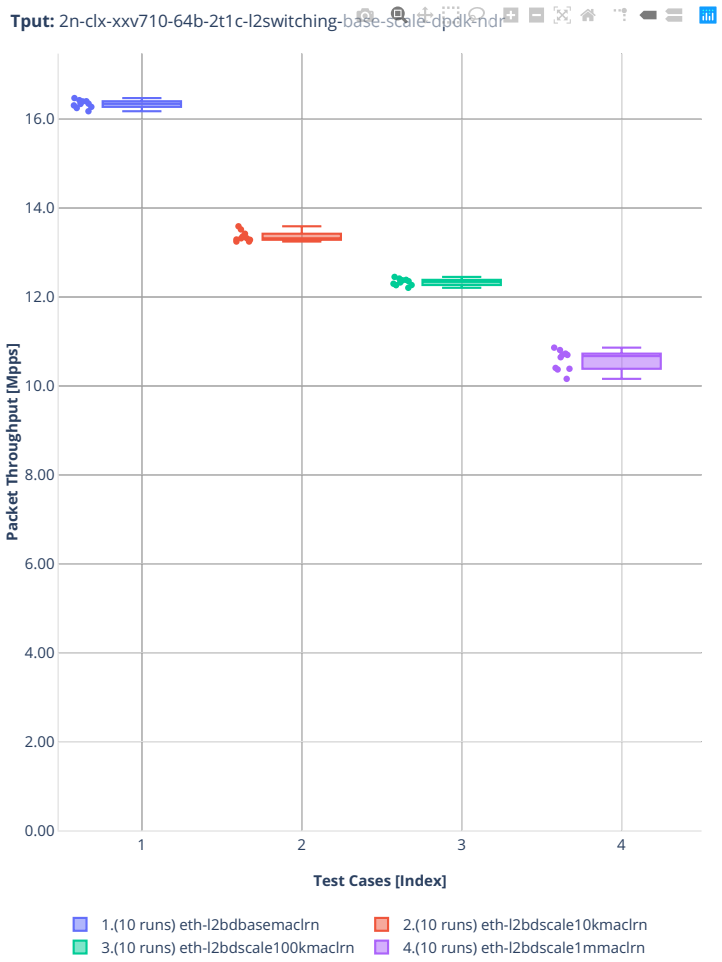

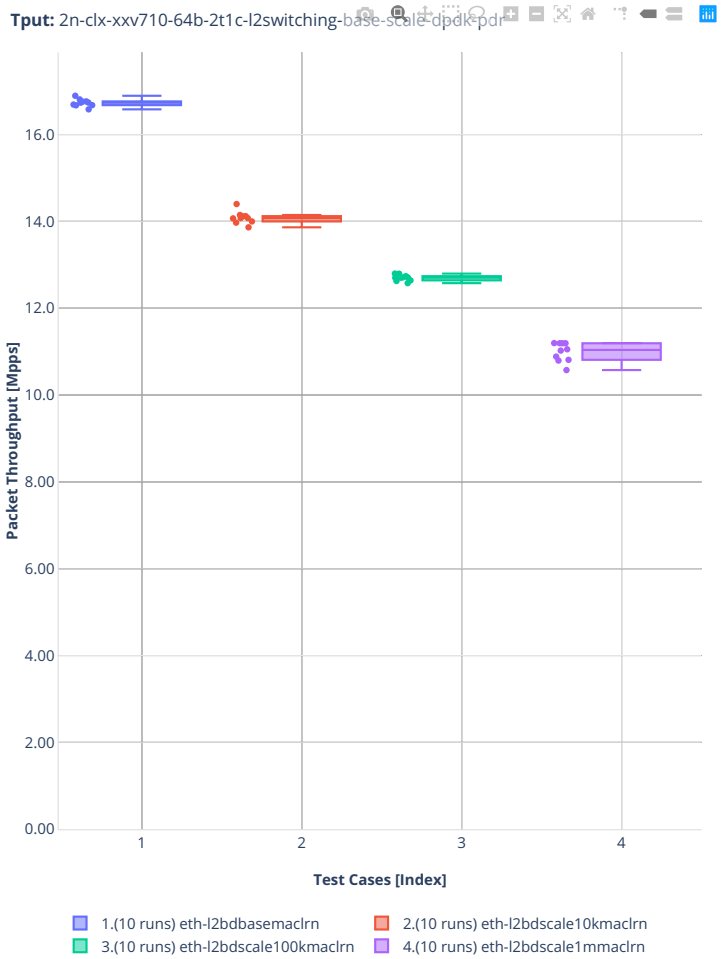
Tput: 2n-clx-xxv710-64b-2t1c-l2switching-base-scale-dpdk-pdr
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) eth-l2bdbasemaclrn
2.(10 runs) eth-l2bdscale10kmaclrn
3.(10 runs) eth-l2bdscale100kmaclrn
4.(10 runs) eth-l2bdscale1mmaclrn

2n-clx-x710

64b-2t1c-l2switching-base-avf

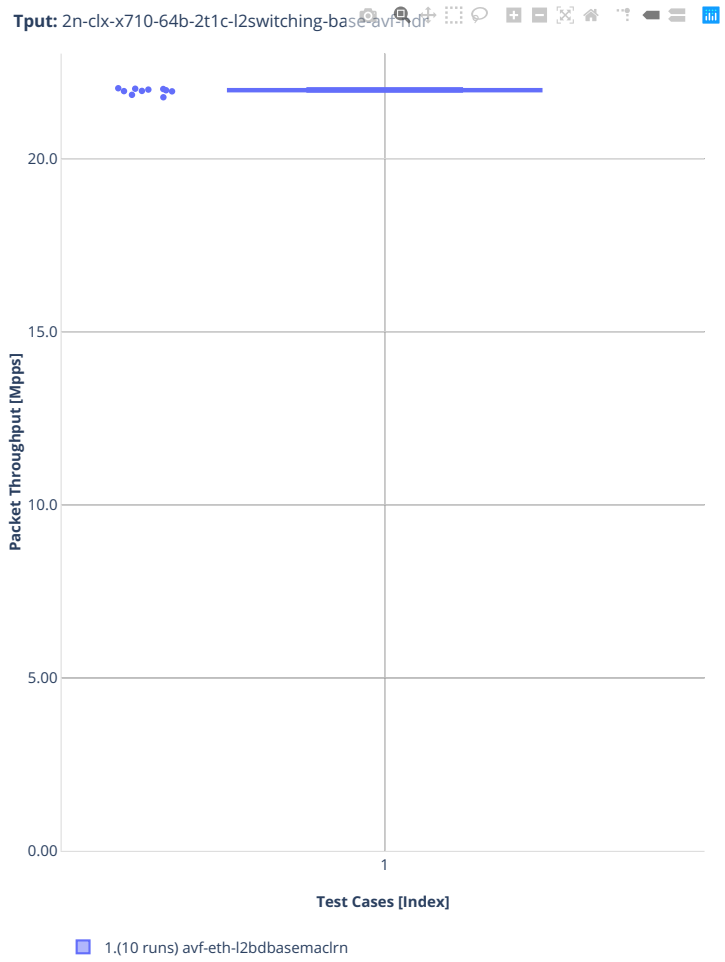

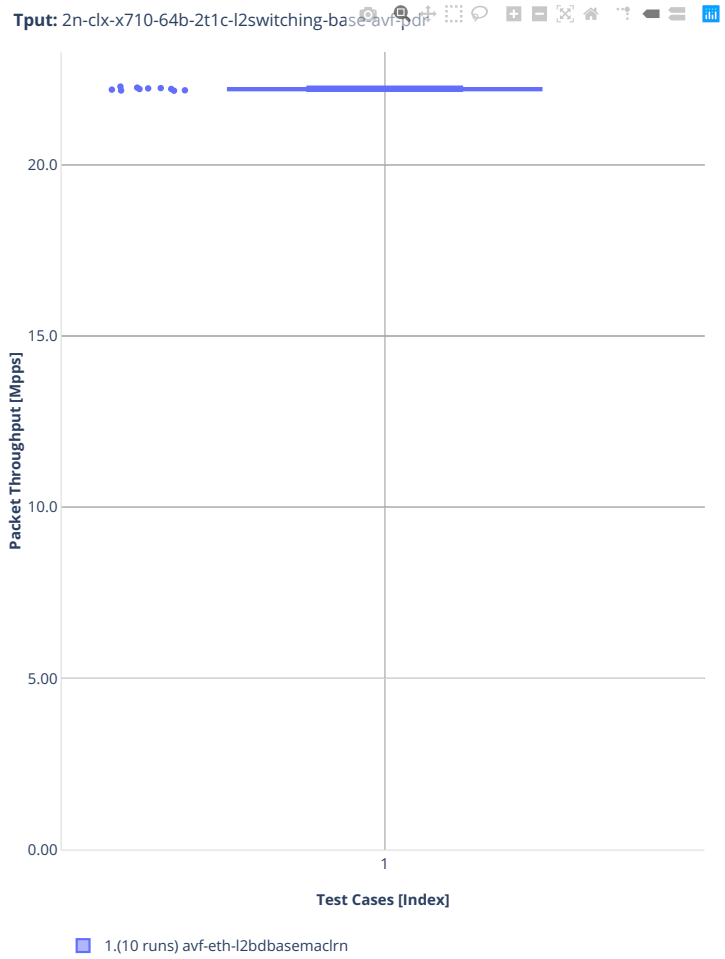
Tput: 2n-clx-x710-64b-2t1c-l2switching-base-avf-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
Test Cases [Index]
1.(10 runs) avf-eth-l2bdbasemaclrn

2n-clx-cx556a

64b-2t1c-l2switching-base-rdma-core

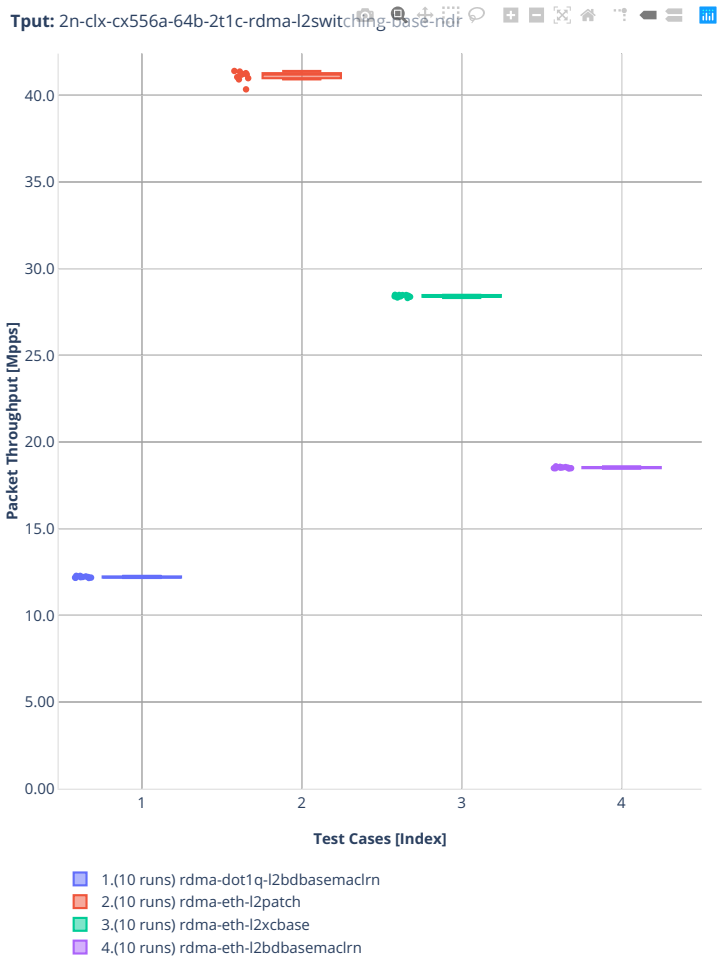

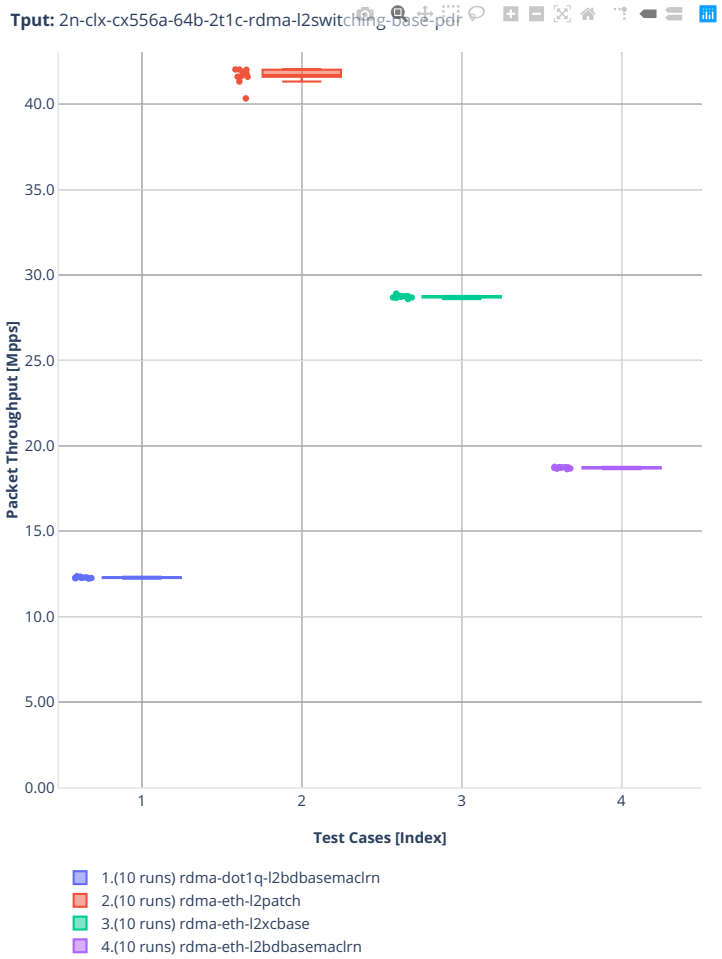
Tput: 2n-clx-cx556a-64b-2t1c-rdma-l2switching-base-pdr
Packet Throughput [Mpps]
40.0
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
1
2
3
4
Test Cases [Index]
1.(10 runs) rdma-dot1q-l2bdbasemaclrn
2.(10 runs) rdma-eth-l2patch
3.(10 runs) rdma-eth-l2xcbase
4.(10 runs) rdma-eth-l2bdbasemaclrn

64b-2t1c-l2switching-scale-rdma-core

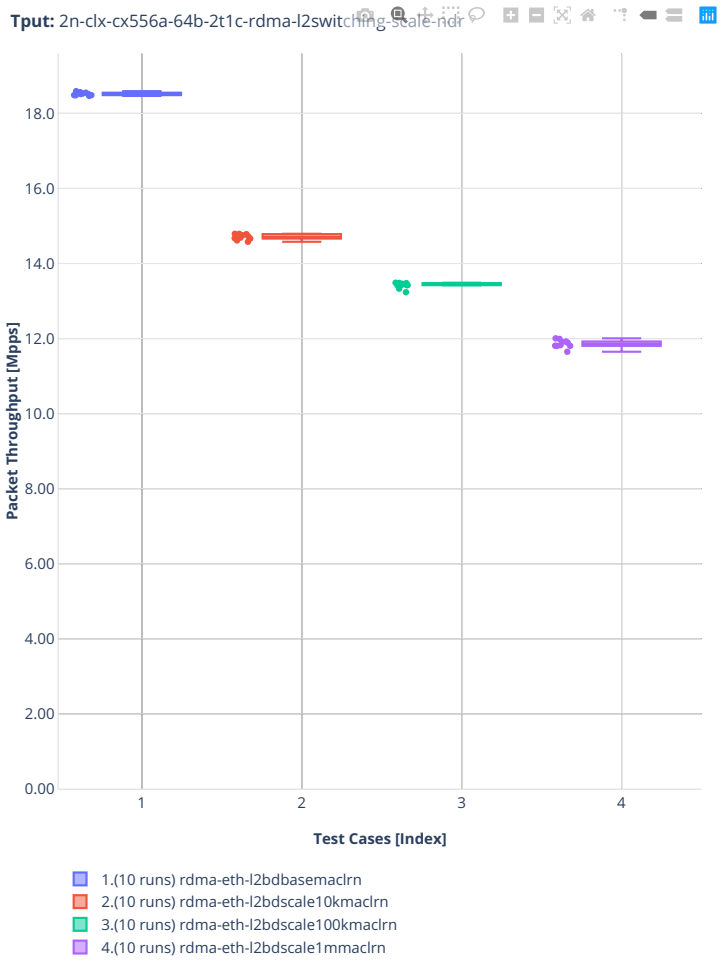

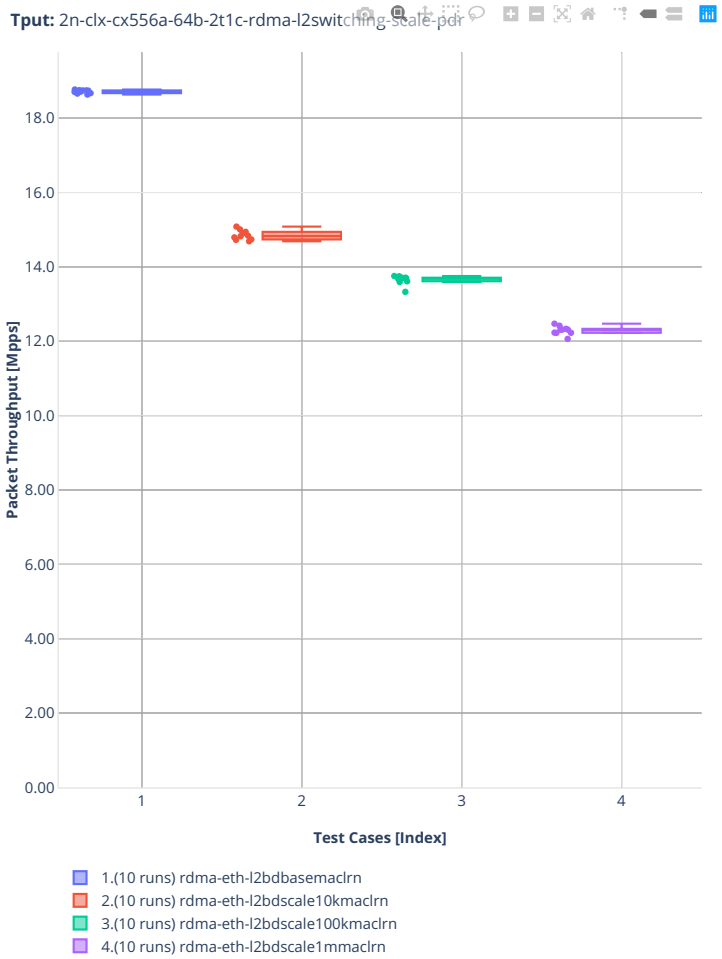
Tput: 2n-clx-cx556a-64b-2t1c-rdma-l2switching-scale-pdr
Packet Throughput [Mpps]
18.0
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
Test Cases [Index]
1.(10 runs) rdma-eth-l2bdbasemaclrn
2.(10 runs) rdma-eth-l2bdscale10kmaclrn
3.(10 runs) rdma-eth-l2bdscale100kmaclrn
4.(10 runs) rdma-eth-l2bdscale1mmaclrn

2n-clx-e810cq

64b-2t1c-l2switching-avf

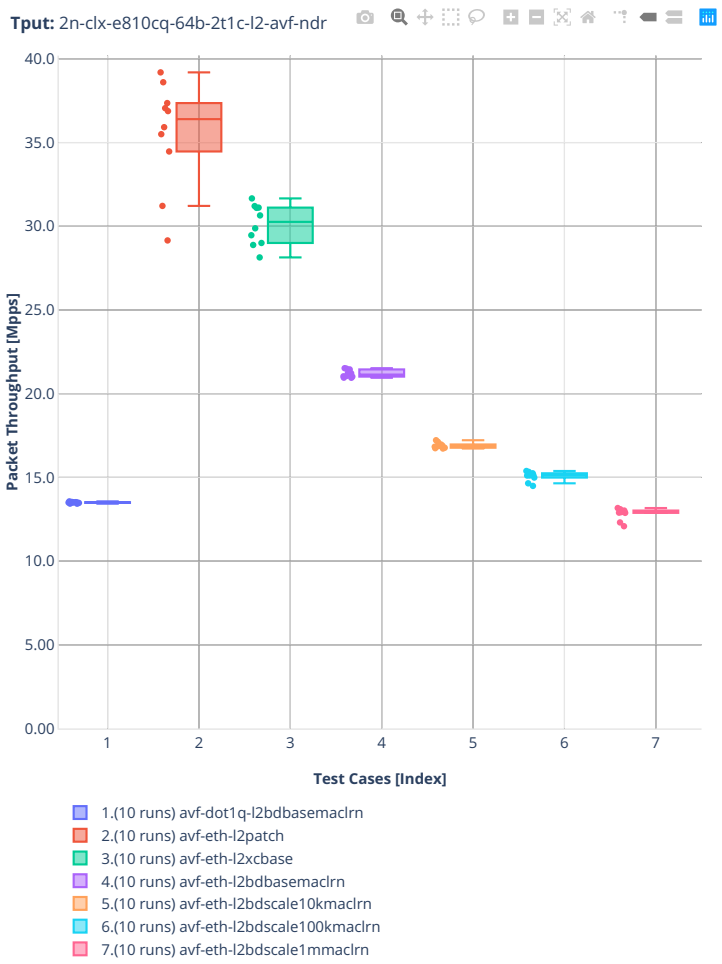

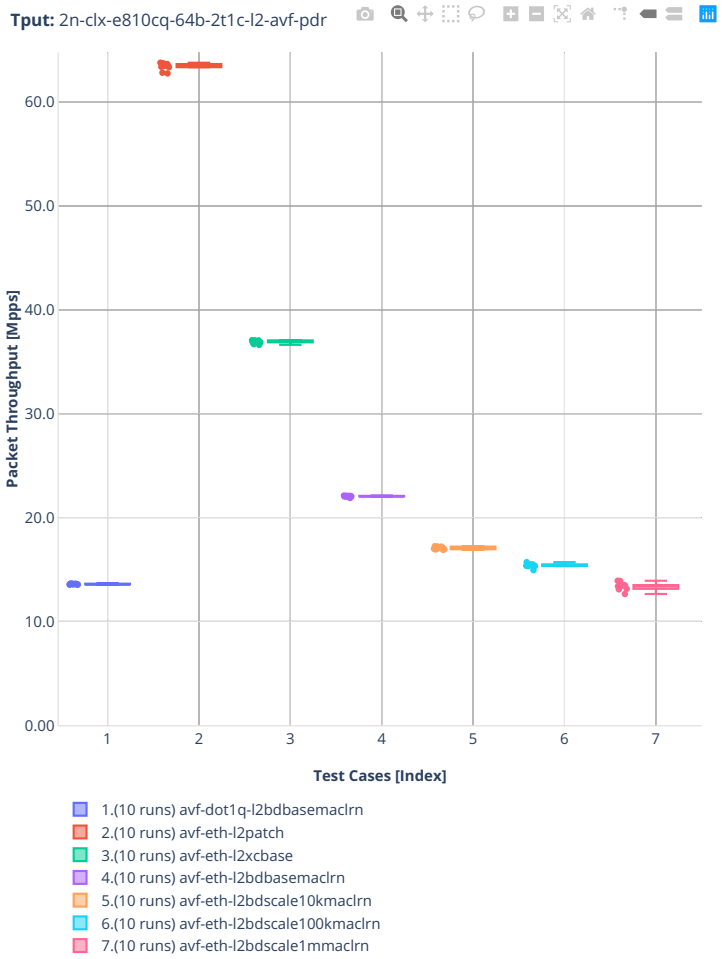
Tput: 2n-clx-e810cq-64b-2t1c-l2-avf-pdr
Packet Throughput [Mpps]
60.0
50.0
40.0
30.0
20.0
10.0
0.00
1
2
3
4
5
6
7
Test Cases [Index]
1.(10 runs) avf-dot1q-l2bdbasemaclrn
2.(10 runs) avf-eth-l2patch
3.(10 runs) avf-eth-l2xcbase
4.(10 runs) avf-eth-l2bdbasemaclrn
5.(10 runs) avf-eth-l2bdscale10kmaclrn
6.(10 runs) avf-eth-l2bdscale100kmaclrn
7.(10 runs) avf-eth-l2bdscale1mmaclrn

64b-2t1c-l2switching-dpdk

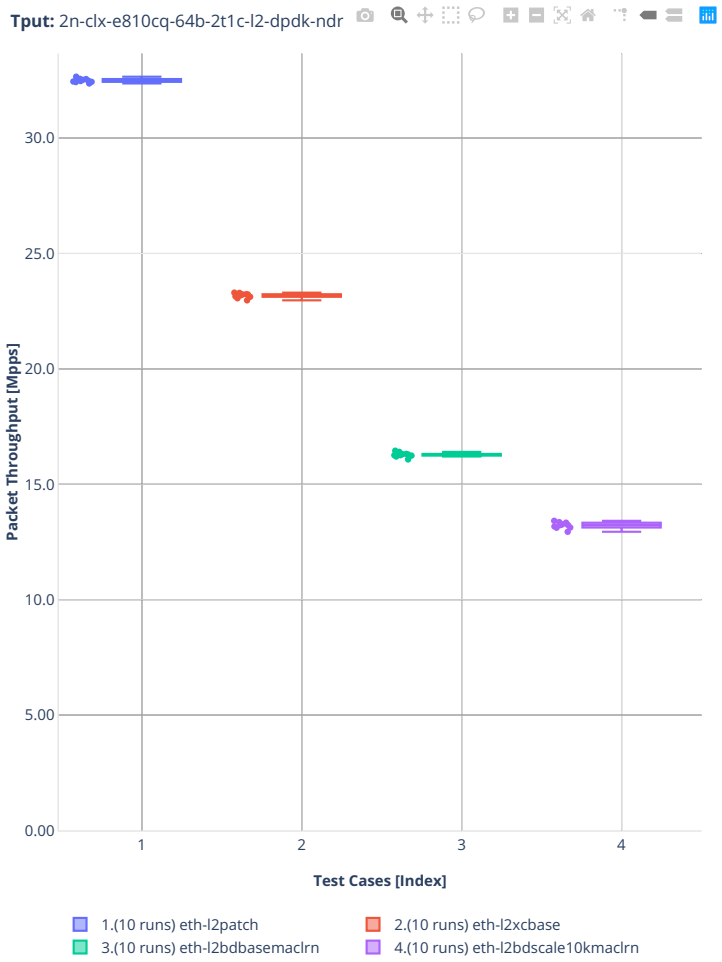


Tput: 2n-clx-e810cq-64b-2t1c-l2-dpdk-pdr
Packet Throughput [Mpps]
30.0
25.0
20.0
15.0
10.0
5.00
0.00
1
2
3
4
Test Cases [Index]
1.(10 runs) eth-l2patch
2.(10 runs) eth-l2xcbase
3.(10 runs) eth-l2bdbasemaclrn
4.(10 runs) eth-l2bdscale10kmaclrn

2n-zn2-xxv710

64b-2t1c-l2switching-base-avf

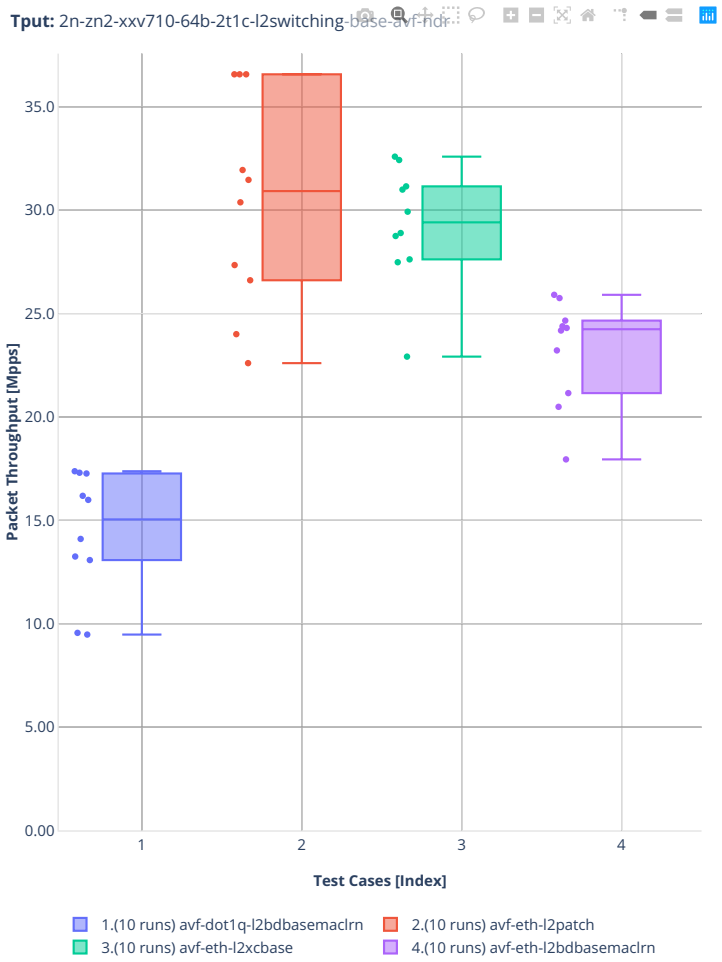

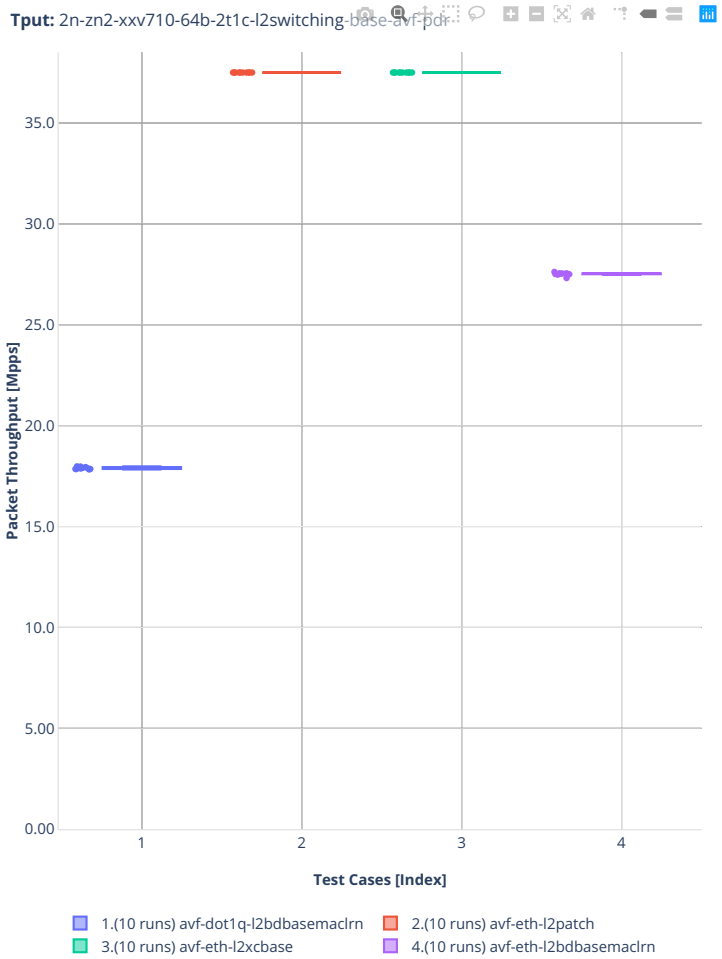
Tput: 2n-zn2-xxv710-64b-2t1c-l2switching-base-avf-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) avf-dot1q-l2bdbasemaclrn
2.(10 runs) avf-eth-l2patch
3.(10 runs) avf-eth-l2xcbase
4.(10 runs) avf-eth-l2bdbasemaclrn

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

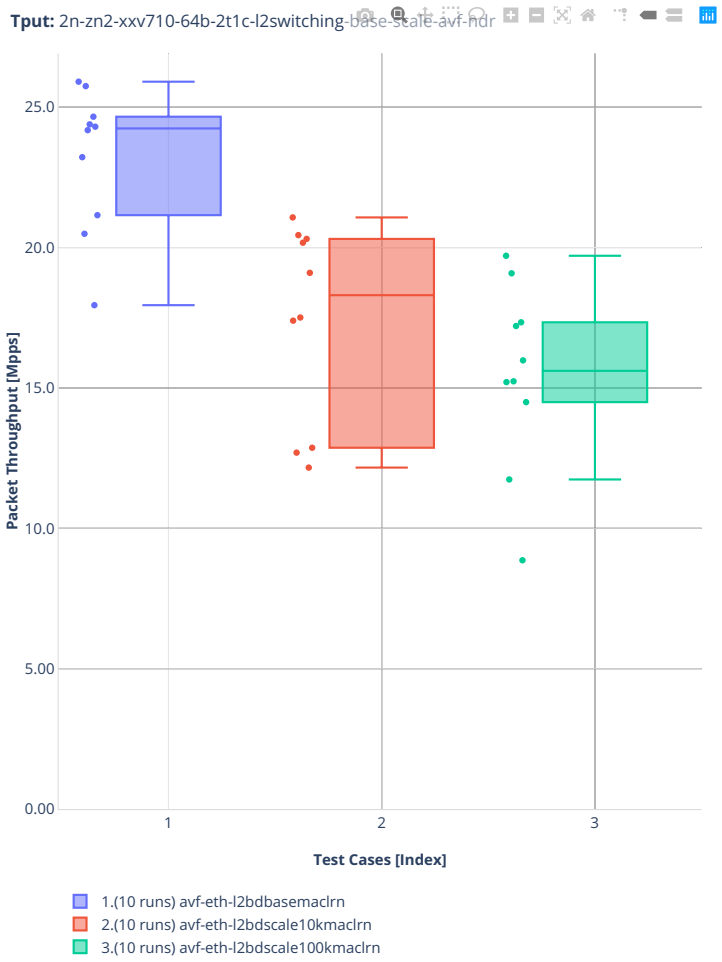

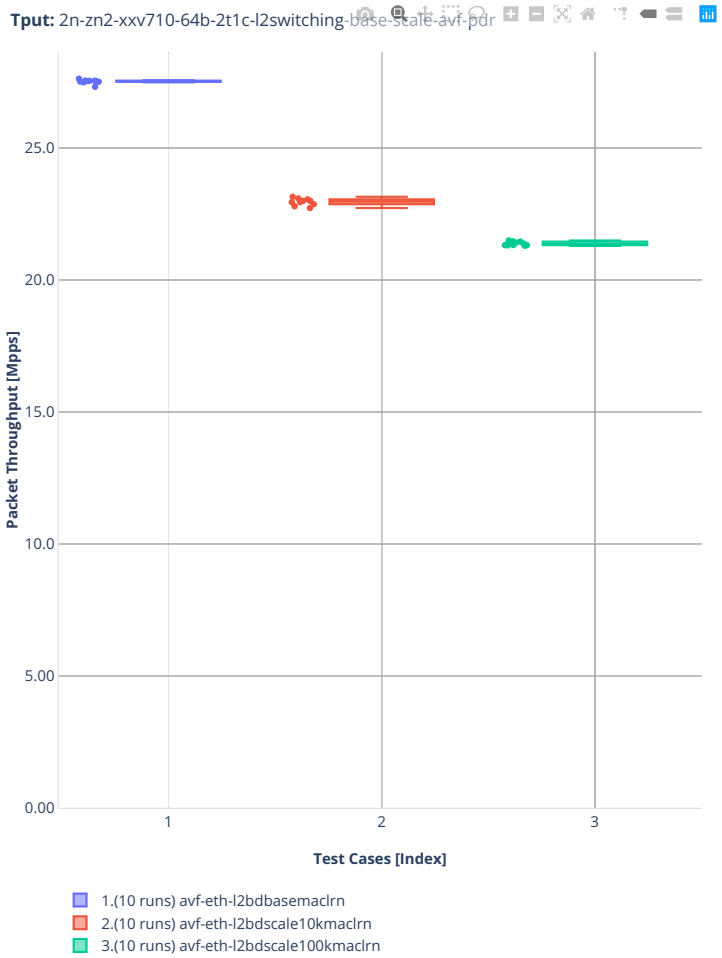
Tput: 2n-zn2-xxv710-64b-2t1c-l2switching-base-scale-avf-pdr
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-eth-l2bdbasemaclrn
2.(10 runs) avf-eth-l2bdscale10kmaclrn
3.(10 runs) avf-eth-l2bdscale100kmaclrn

64b-2t1c-l2switching-base-dpdk

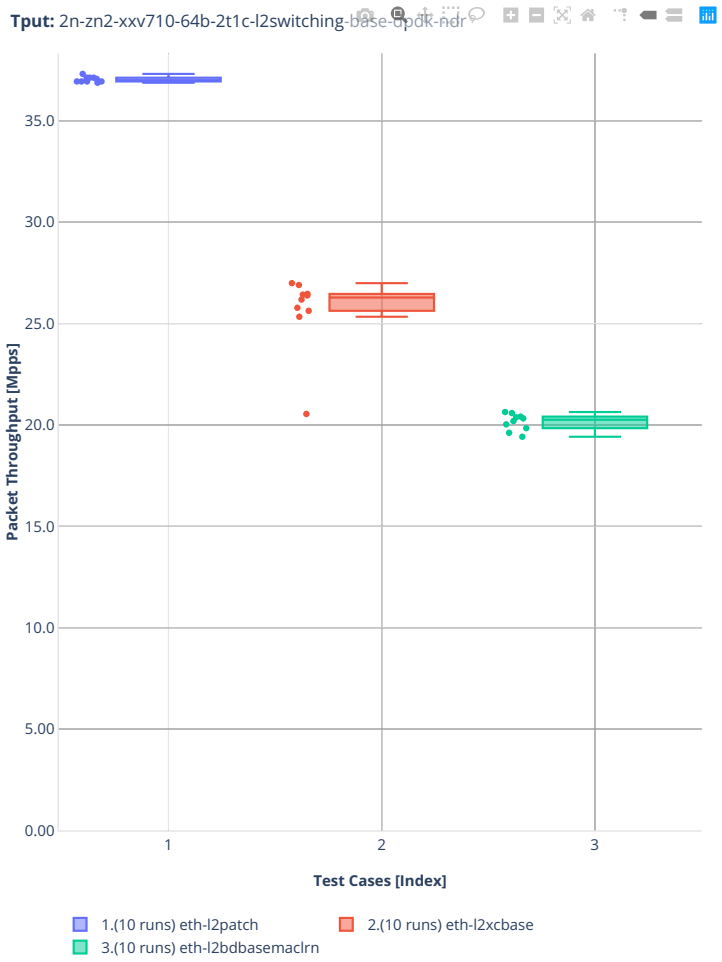

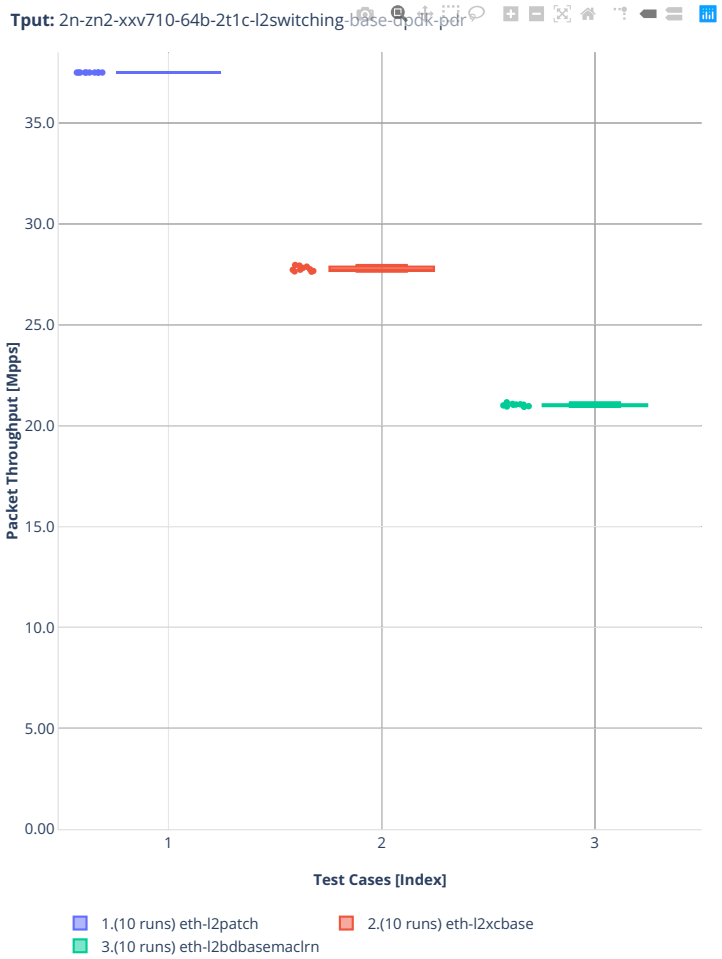
Tput: 2n-zn2-xxv710-64b-2t1c-l2switching-base-dpdk-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) eth-l2patch
2.(10 runs) eth-l2xcbase
3.(10 runs) eth-l2bdbasemaclrn

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

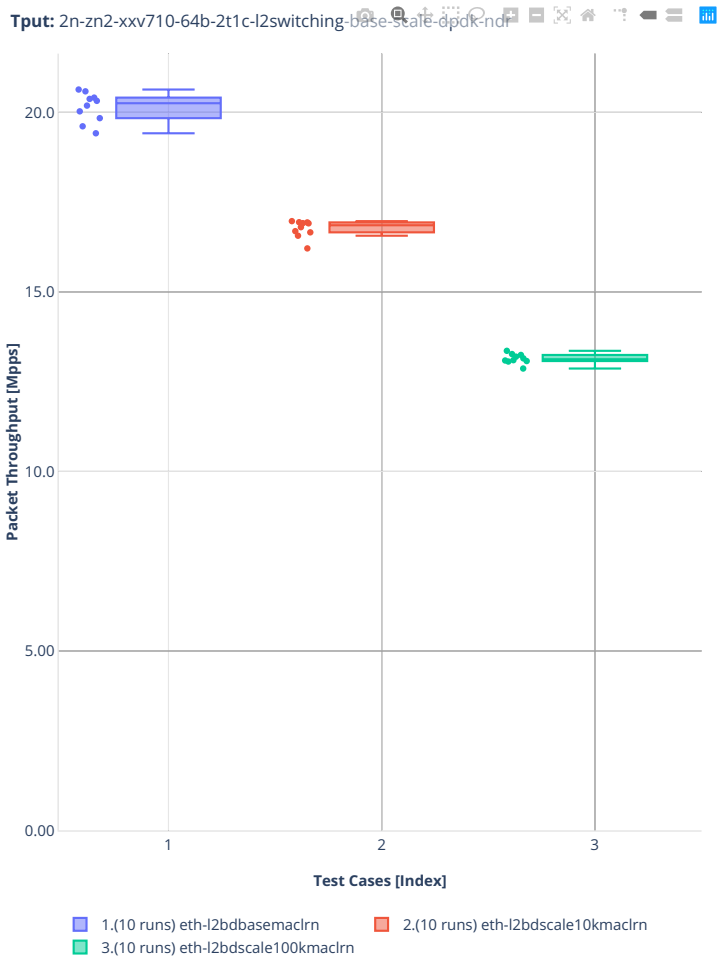

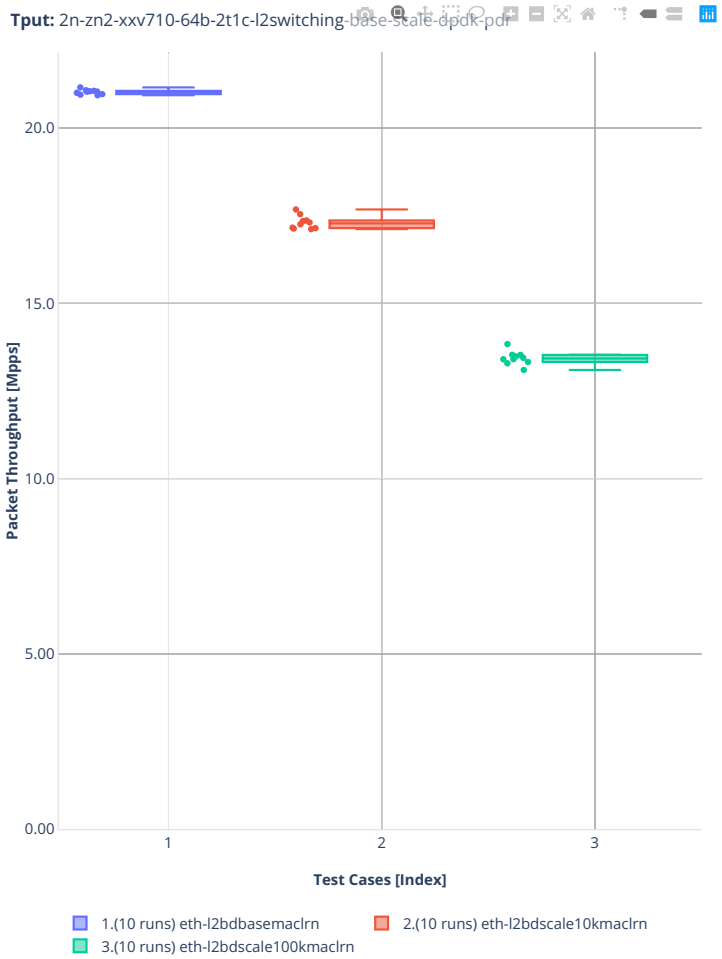
Tput: 2n-zn2-xxv710-64b-2t1c-l2switching-base-scale-dpdk-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) eth-l2bdbasemaclrn
2.(10 runs) eth-l2bdscale10kmaclrn
3.(10 runs) eth-l2bdscale100kmaclrn

2n-zn2-x710

64b-2t1c-l2switching-base

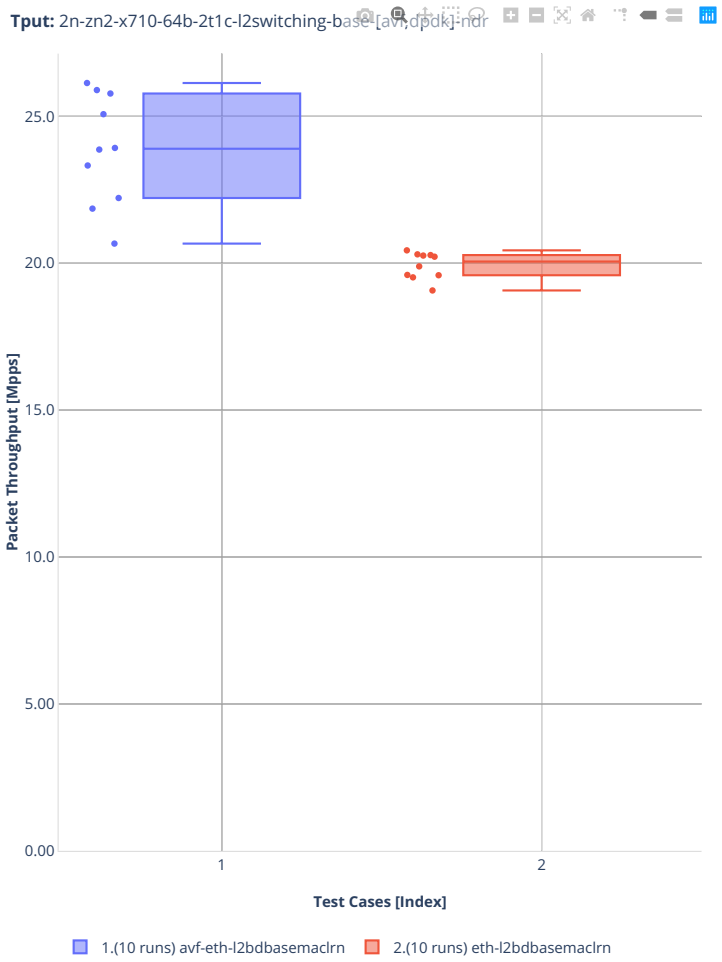

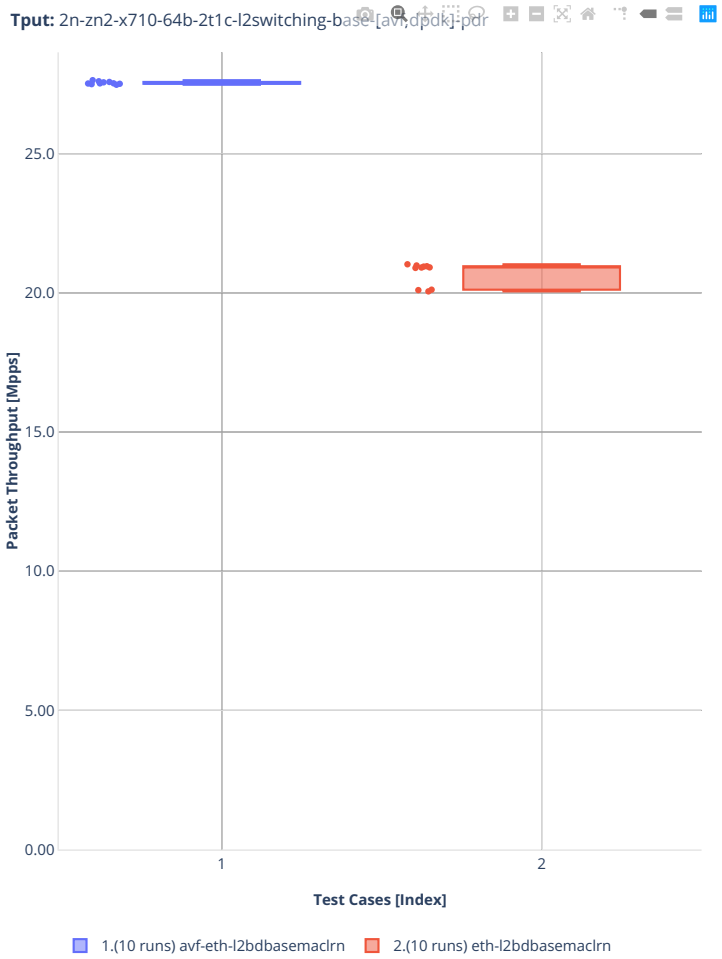
Tput: 2n-zn2-x710-64b-2t1c-l2switching-base-[avf,dpdk]-pdr
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) avf-eth-l2bdbasemaclrn
2.(10 runs) eth-l2bdbasemaclrn

2n-zn2-cx556a

64b-2t1c-l2switching-base-rdma-core

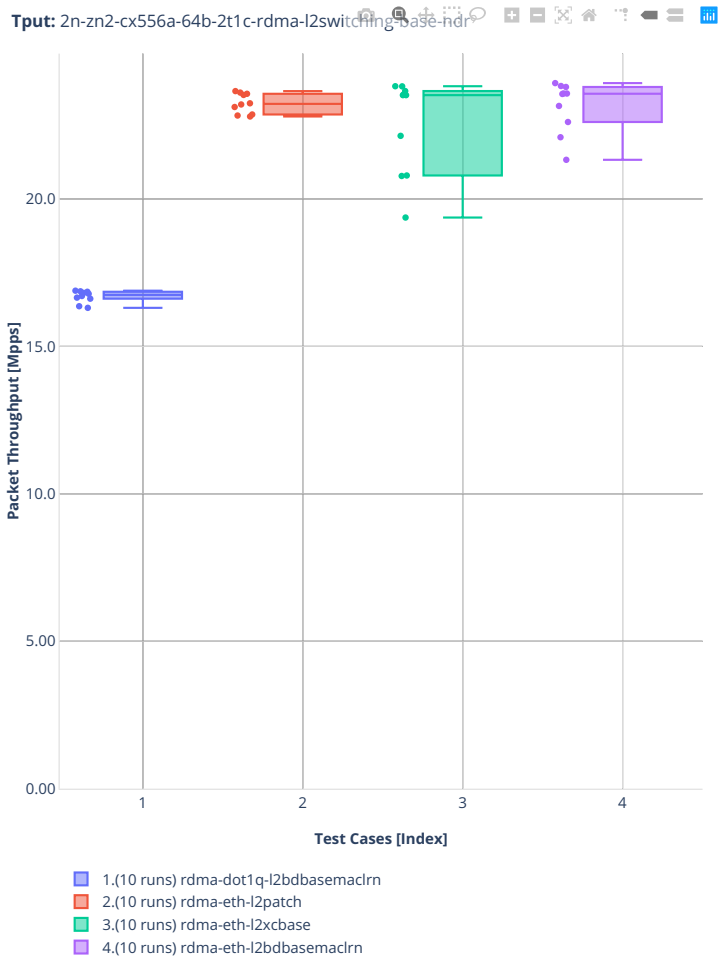

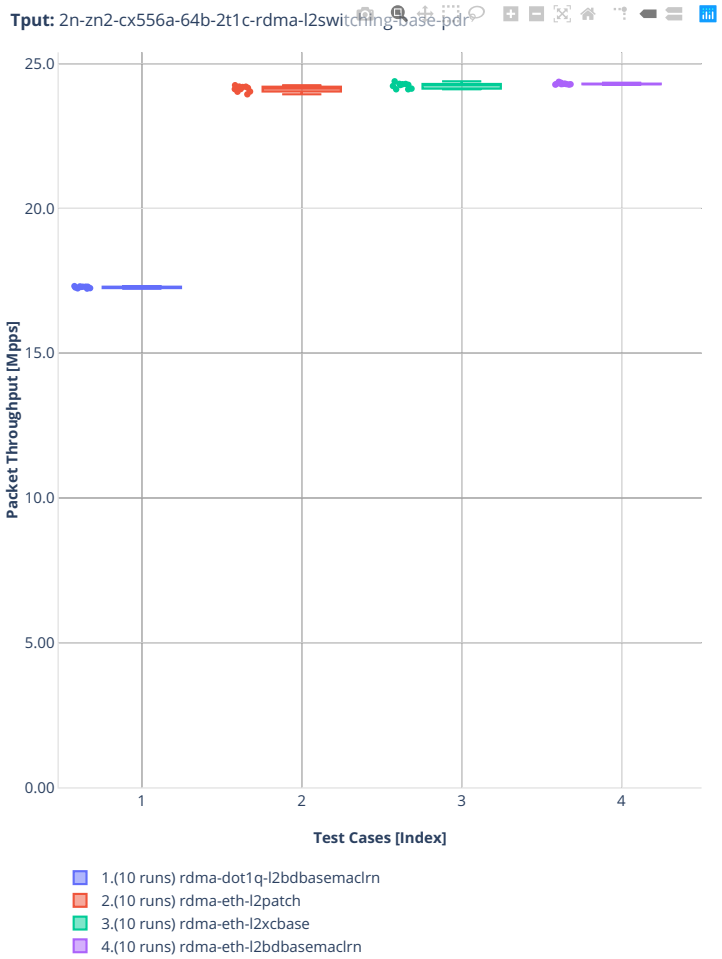
Tput: 2n-zn2-cx556a-64b-2t1c-rdma-l2switching-base-pdr
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) rdma-dot1q-l2bdbasemaclrn
2.(10 runs) rdma-eth-l2patch
3.(10 runs) rdma-eth-l2xcbase
4.(10 runs) rdma-eth-l2bdbasemaclrn

64b-2t1c-l2switching-scale-rdma-core

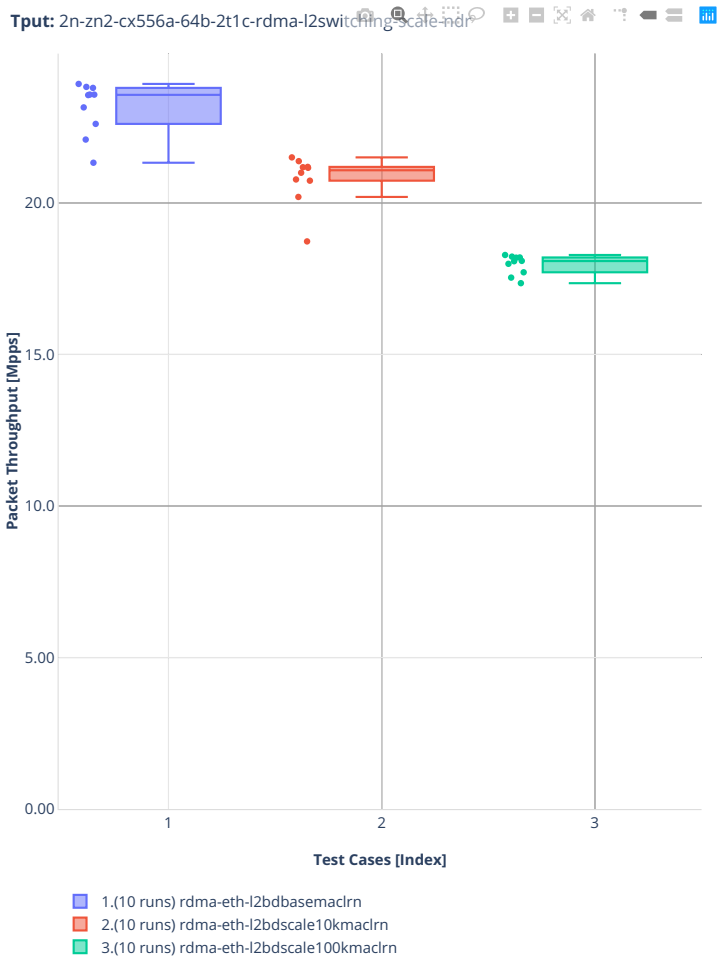

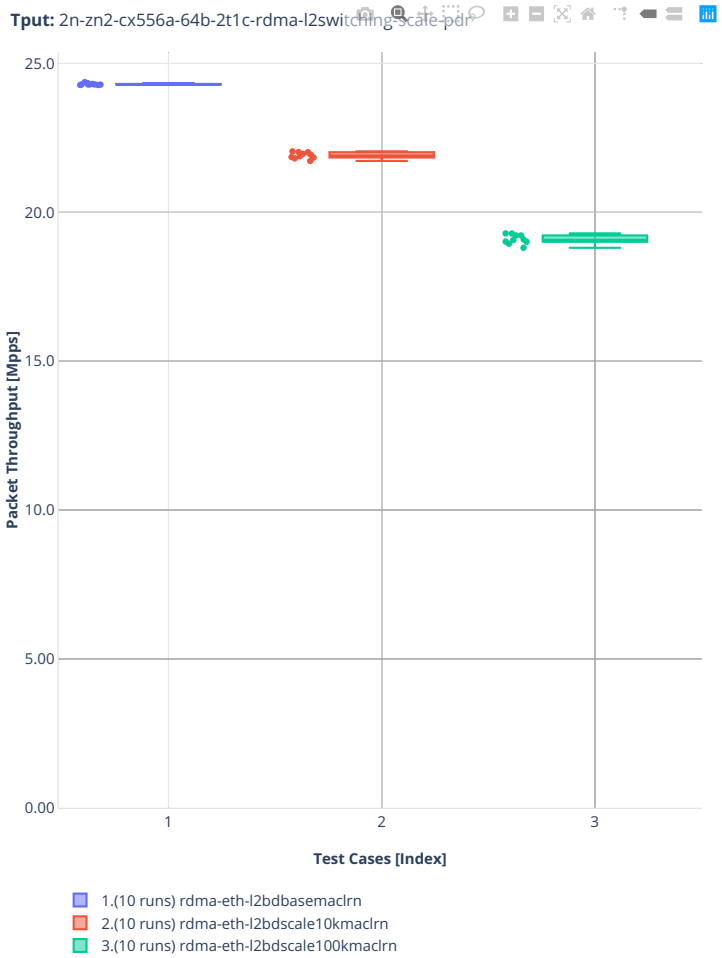
Tput: 2n-zn2-cx556a-64b-2t1c-rdma-l2switching-scale-pdr
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) rdma-eth-l2bdbasemaclrn
2.(10 runs) rdma-eth-l2bdscale10kmaclrn
3.(10 runs) rdma-eth-l2bdscale100kmaclrn

3n-alt-xl710

64b-1t1c-l2switching-base-scale

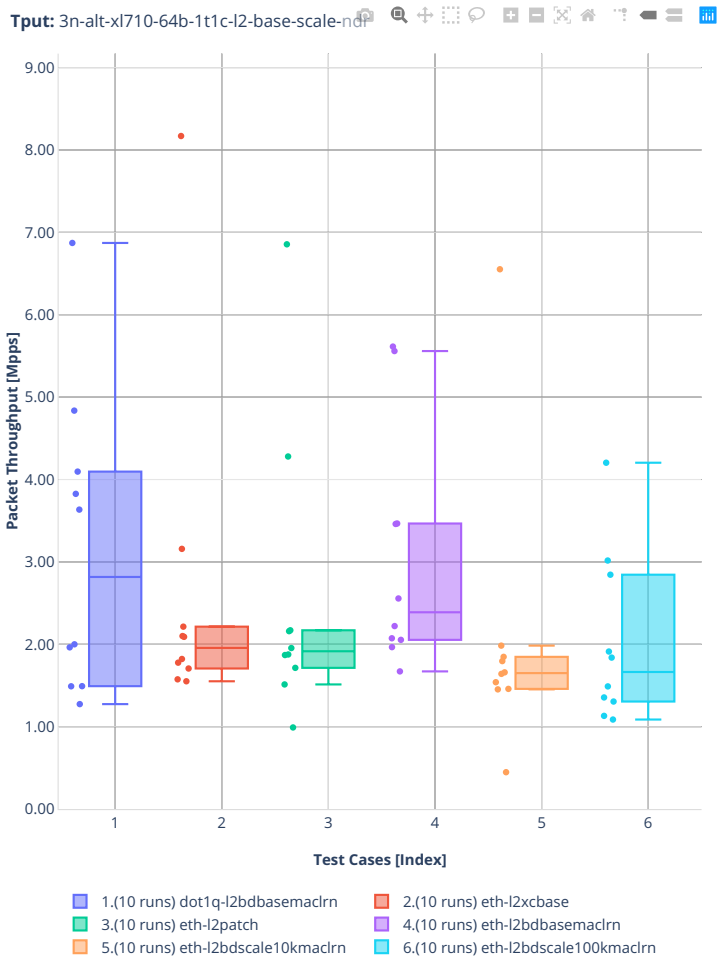

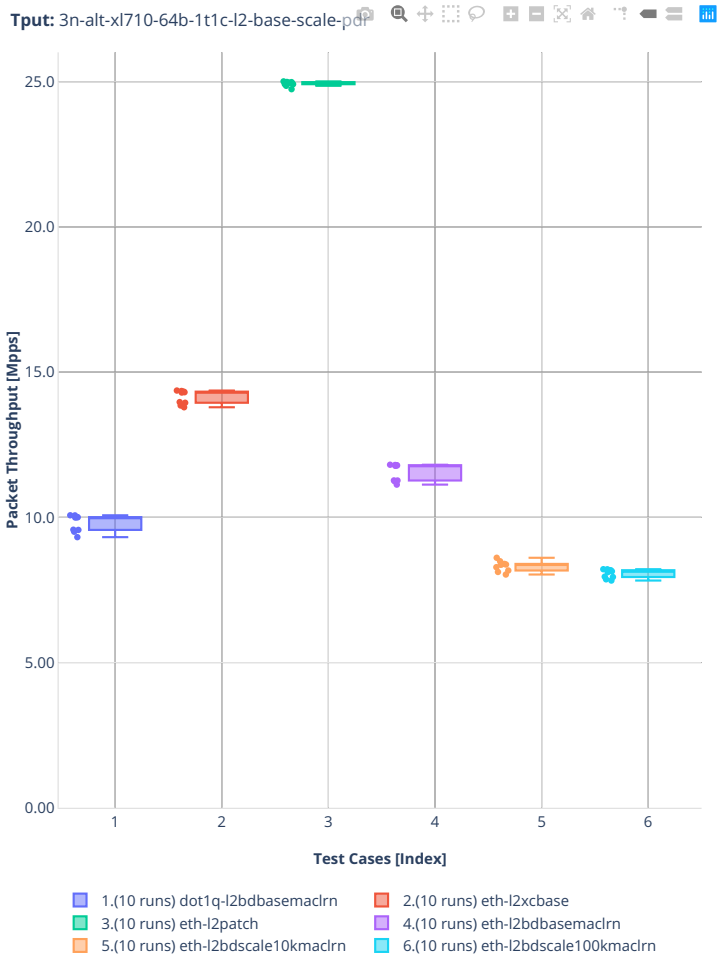
Tput: 3n-alt-xl710-64b-1t1c-l2-base-scale-pdr
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) dot1q-l2bdbasemaclrn
2.(10 runs) eth-l2xcbase
3.(10 runs) eth-l2patch
4.(10 runs) eth-l2bdbasemaclrn
5.(10 runs) eth-l2bdscale10kmaclrn
6.(10 runs) eth-l2bdscale100kmaclrn

64b-1t1c-l2switching-features

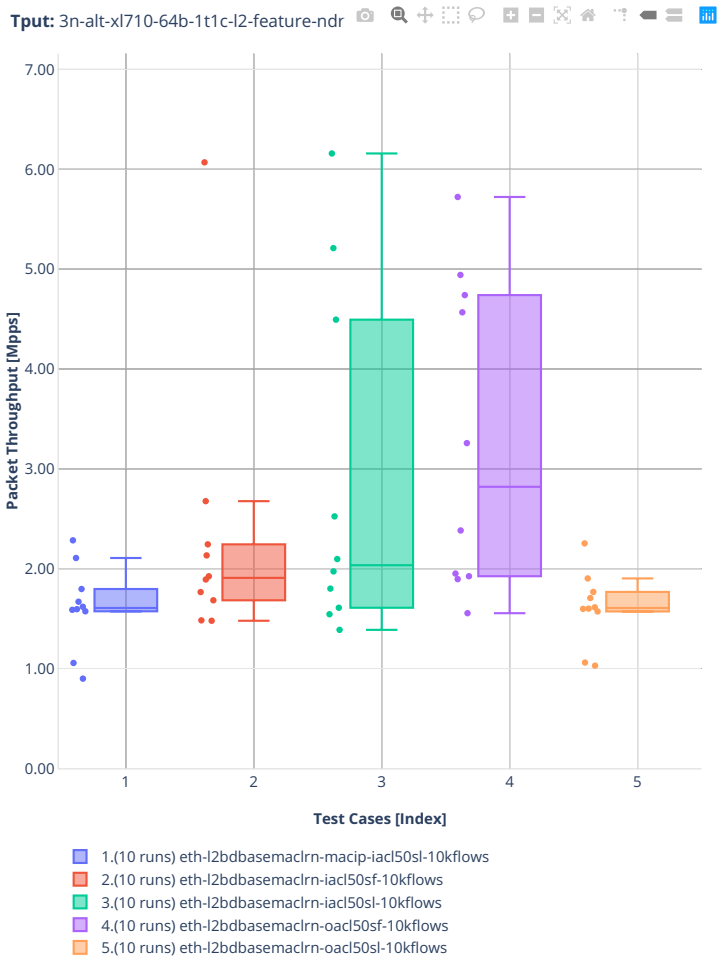

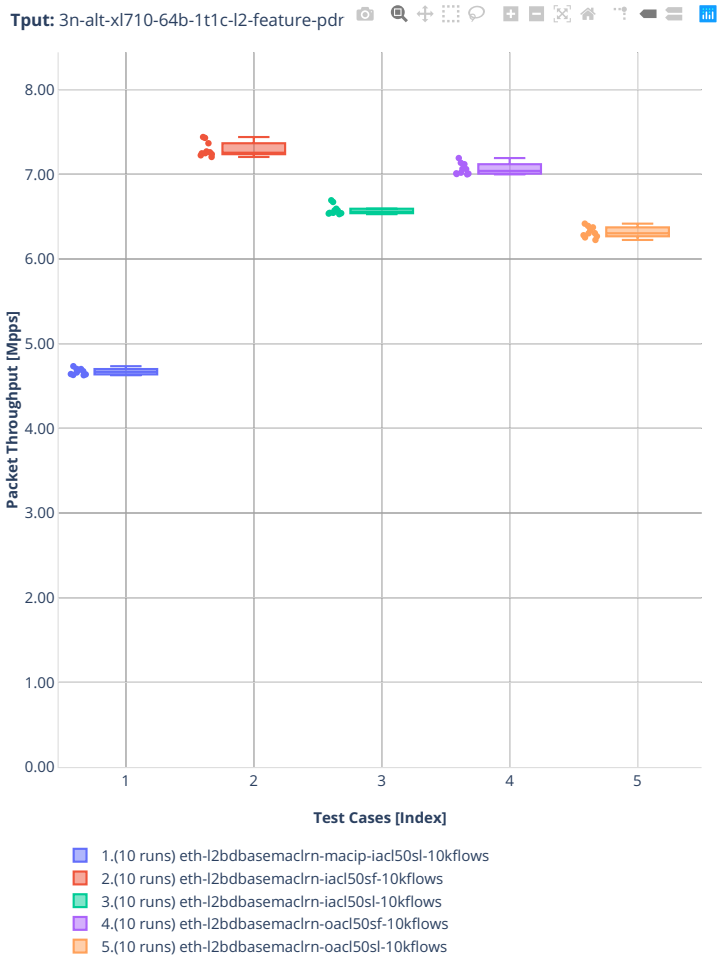
Tput: 3n-alt-xl710-64b-1t1c-l2-feature-pdr
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
Test Cases [Index]
1.(10 runs) eth-l2bdbasemaclrn-macip-iacl50sl-10kflows
2.(10 runs) eth-l2bdbasemaclrn-iacl50sf-10kflows
3.(10 runs) eth-l2bdbasemaclrn-iacl50sl-10kflows
4.(10 runs) eth-l2bdbasemaclrn-oacl50sf-10kflows
5.(10 runs) eth-l2bdbasemaclrn-oacl50sl-10kflows

3n-tsh-x520

64b-1t1c-l2switching-base-ixgbe

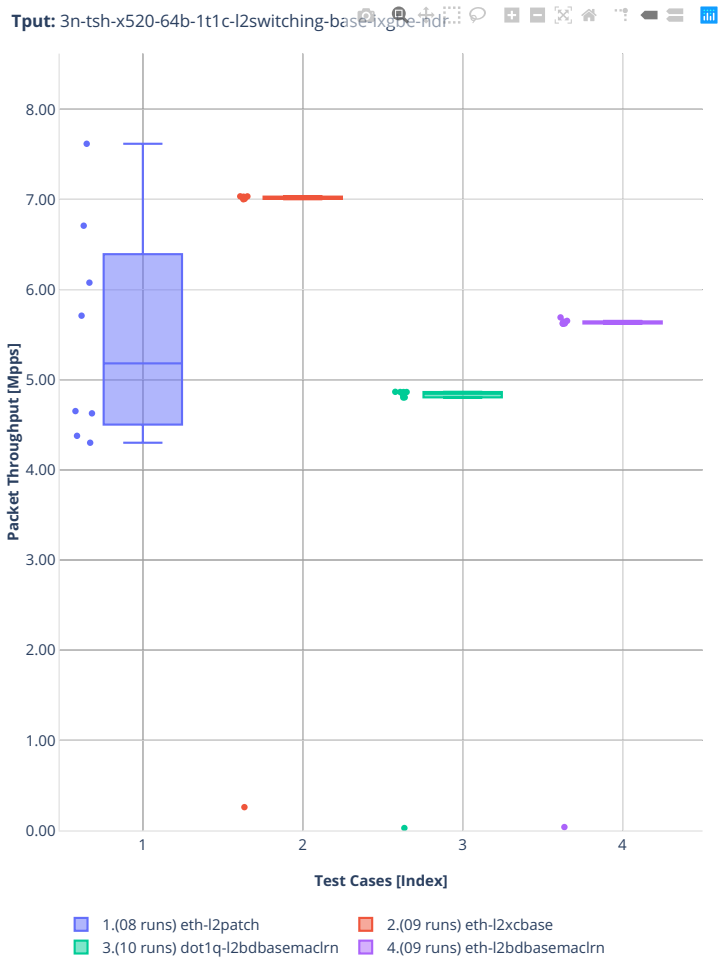

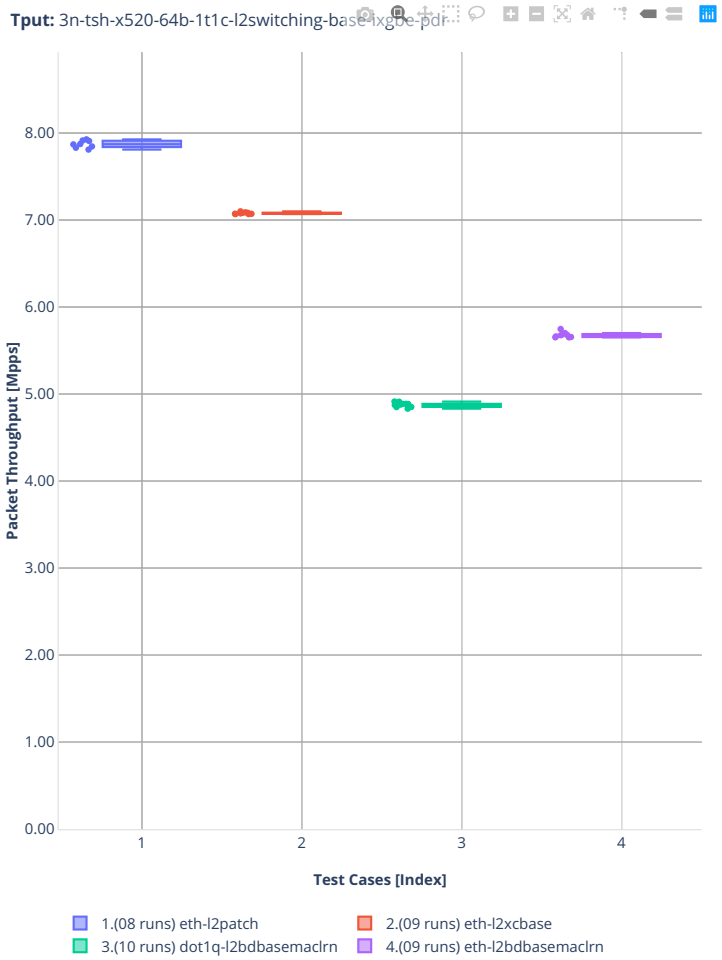
Tput: 3n-tsh-x520-64b-1t1c-l2switching-base-1xgbe-pdr
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(08 runs) eth-l2patch
2.(09 runs) eth-l2xcbase
3.(10 runs) dot1q-l2bdbasemaclrn
4.(09 runs) eth-l2bdbasemaclrn

64b-1t1c-features-l2switching-base-ixgbe

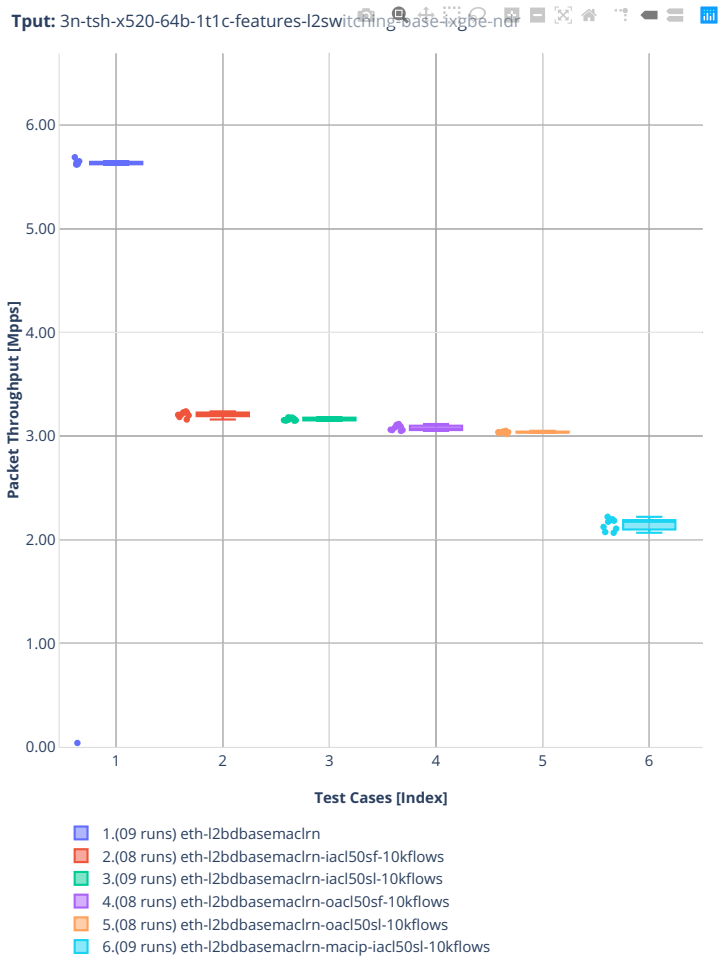

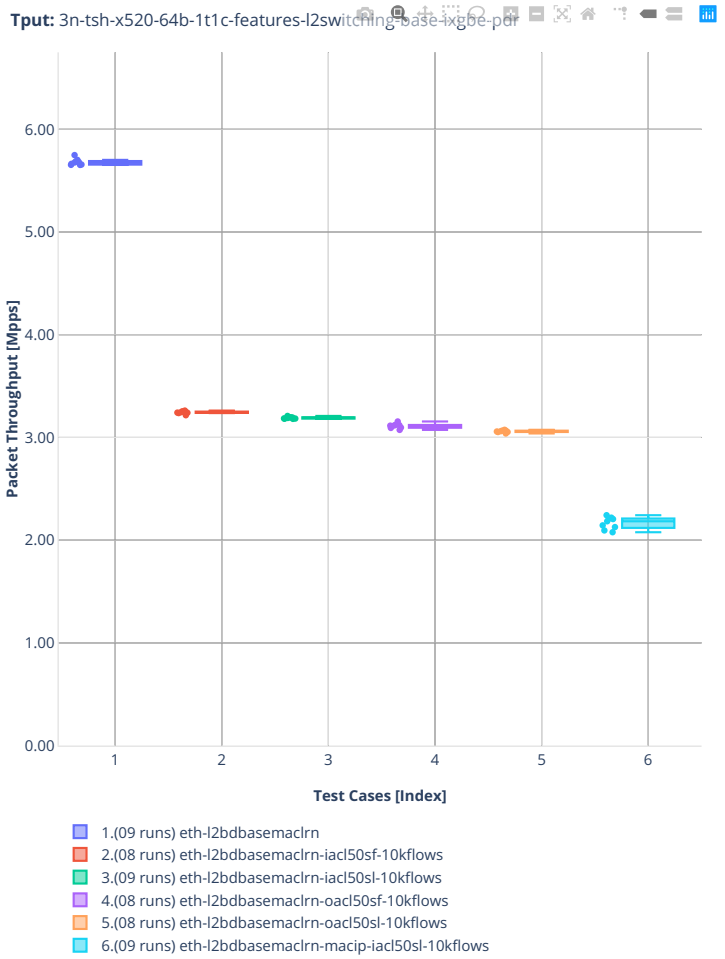
Tput: 3n-tsh-x520-64b-1t1c-features-l2switching-base-ixgbe-pdr
Packet Throughput [Mpps]
6.00
5.00
4.00
3.00
2.00
1.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(09 runs) eth-l2bdbasemaclrn
2.(08 runs) eth-l2bdbasemaclrn-iacl50sf-10kflows
3.(09 runs) eth-l2bdbasemaclrn-iacl50sl-10kflows
4.(08 runs) eth-l2bdbasemaclrn-oacl50sf-10kflows
5.(08 runs) eth-l2bdbasemaclrn-oacl50sl-10kflows
6.(09 runs) eth-l2bdbasemaclrn-macip-iacl50sl-10kflows

2n-tx2-xl710

64b-1t1c-l2switching-base-dpdk

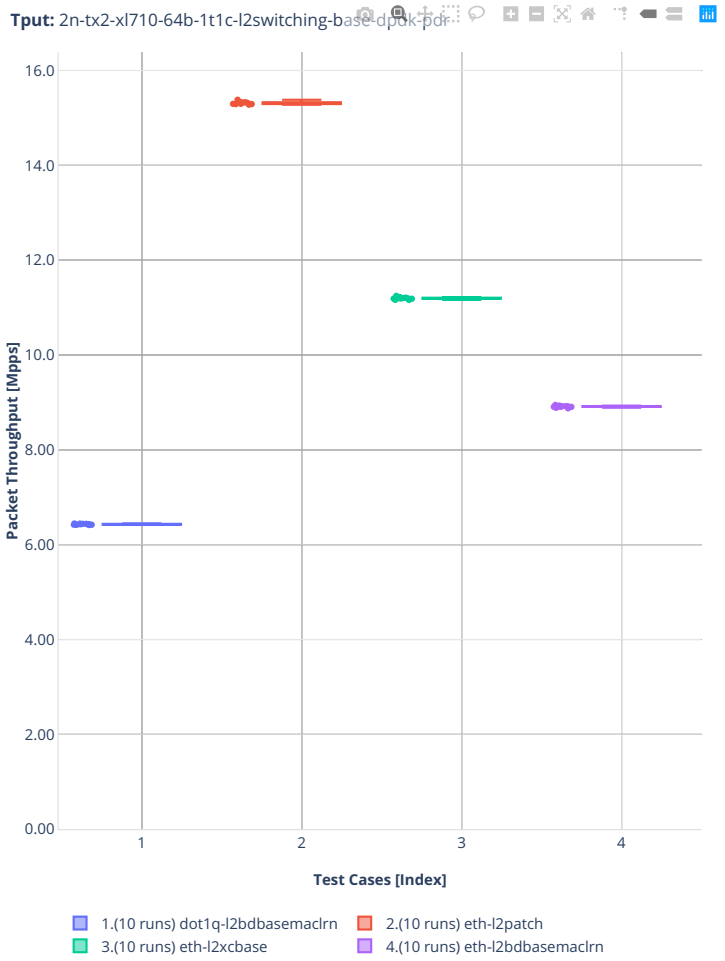
Tput: 2n-tx2-xl710-64b-1t1c-l2switching-base-dpdk-pdr
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) dot1q-l2bdbasemaclrn
2.(10 runs) eth-l2patch
3.(10 runs) eth-l2xcbase
4.(10 runs) eth-l2bdbasemaclrn

64b-1t1c-l2switching-scale-dpdk

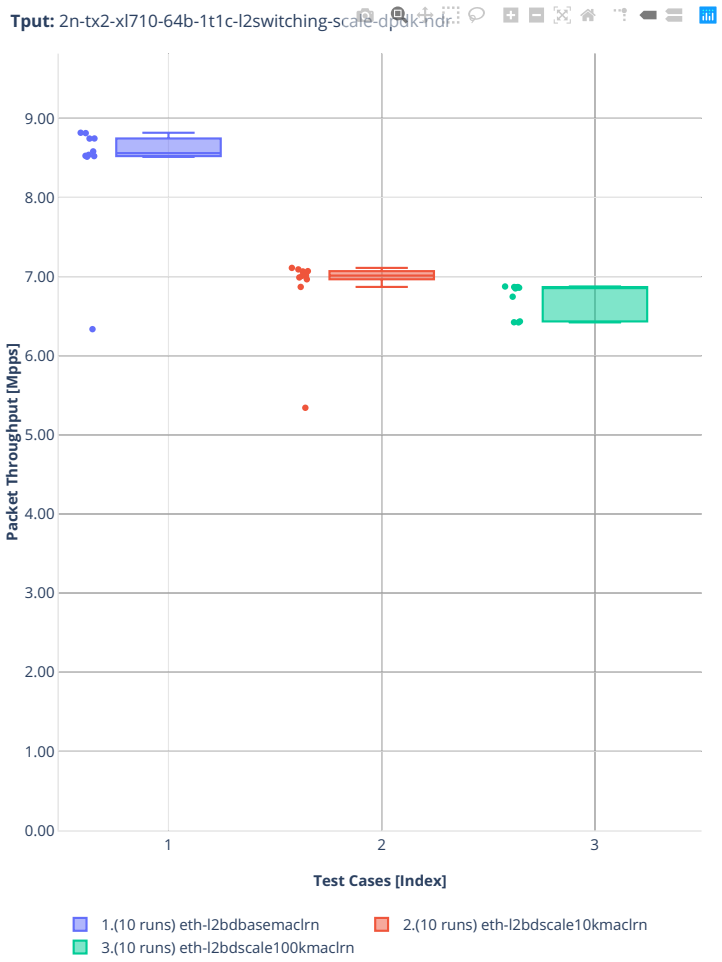

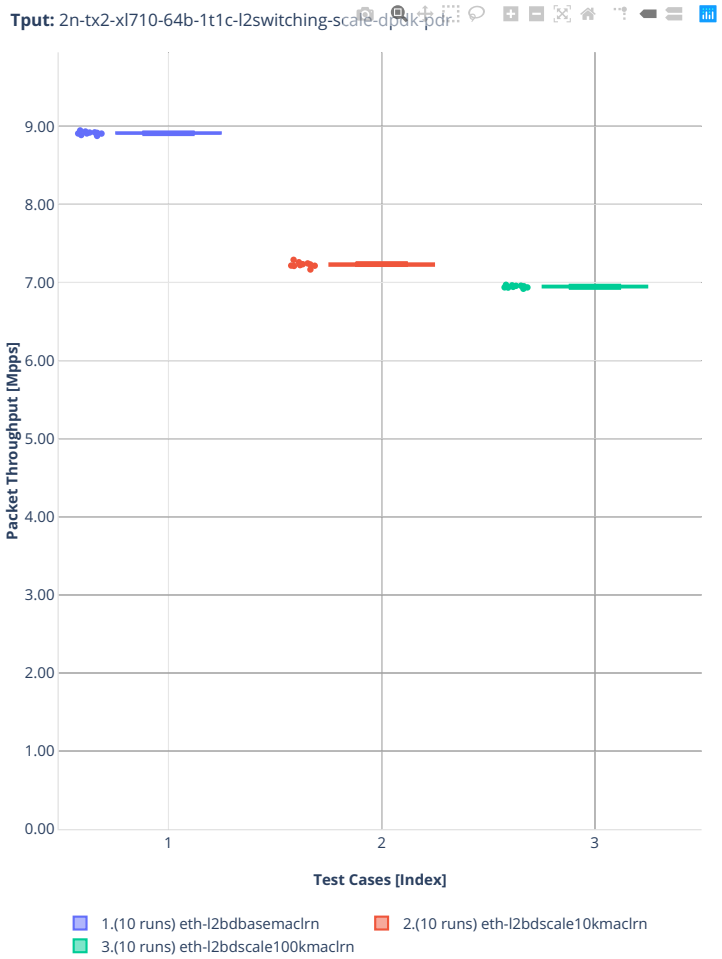
Tput: 2n-tx2-xl710-64b-1t1c-l2switching-scale-dpdk-pdr
9.00
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) eth-l2bdbasemaclrn
2.(10 runs) eth-l2bdscale10kmaclrn
3.(10 runs) eth-l2bdscale100kmaclrn

64b-1t1c-l2switching-scale-af-xdp

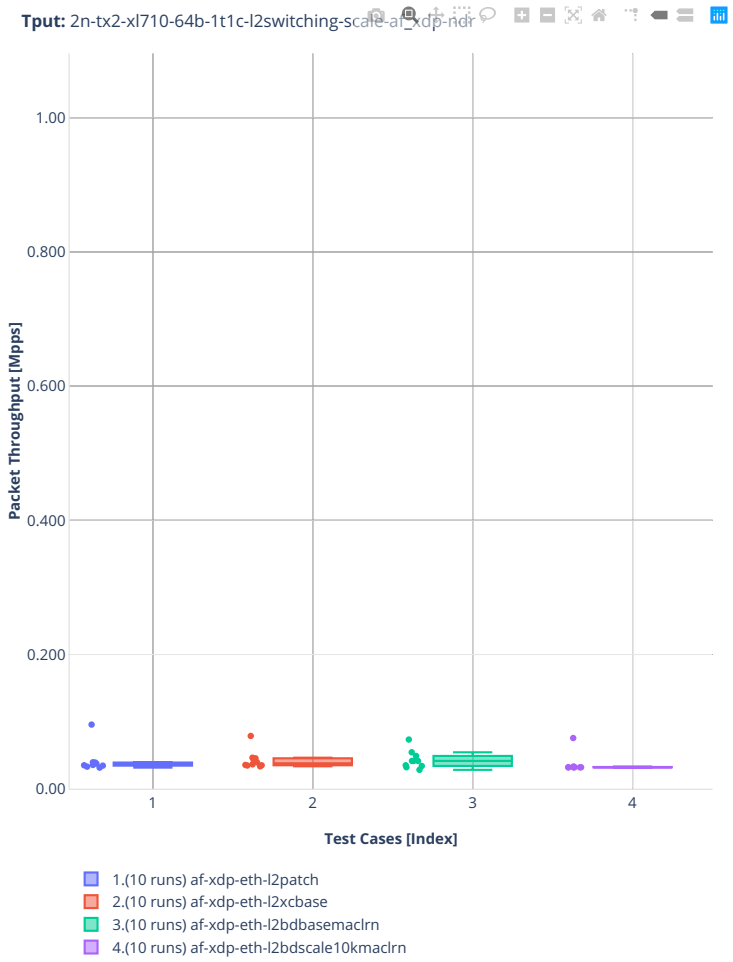

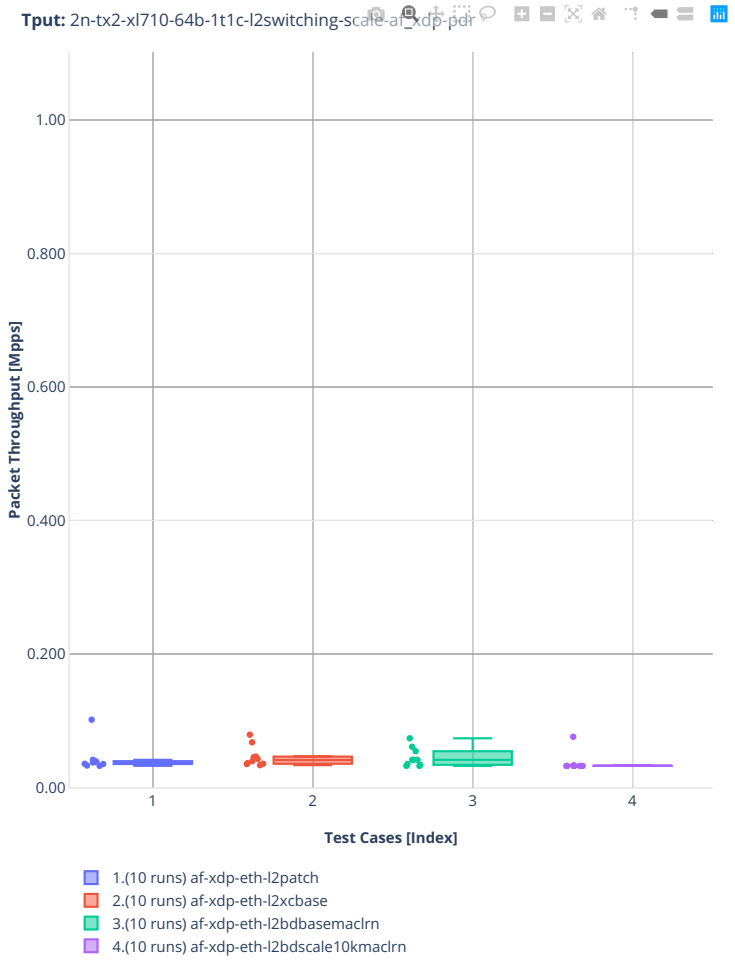
Tput: 2n-tx2-xl710-64b-1t1c-l2switching-scale-af_xdp-pdr
Packet Throughput [Mpps]
1.00
0.800
0.600
0.400
0.200
0.00
1
2
3
4
Test Cases [Index]
1.(10 runs) af-xdp-eth-l2patch
2.(10 runs) af-xdp-eth-l2xcbase
3.(10 runs) af-xdp-eth-l2bdbasemaclrn
4.(10 runs) af-xdp-eth-l2bdscale10kmaclrn

64b-1t1c-features-l2switching-base-dpdk

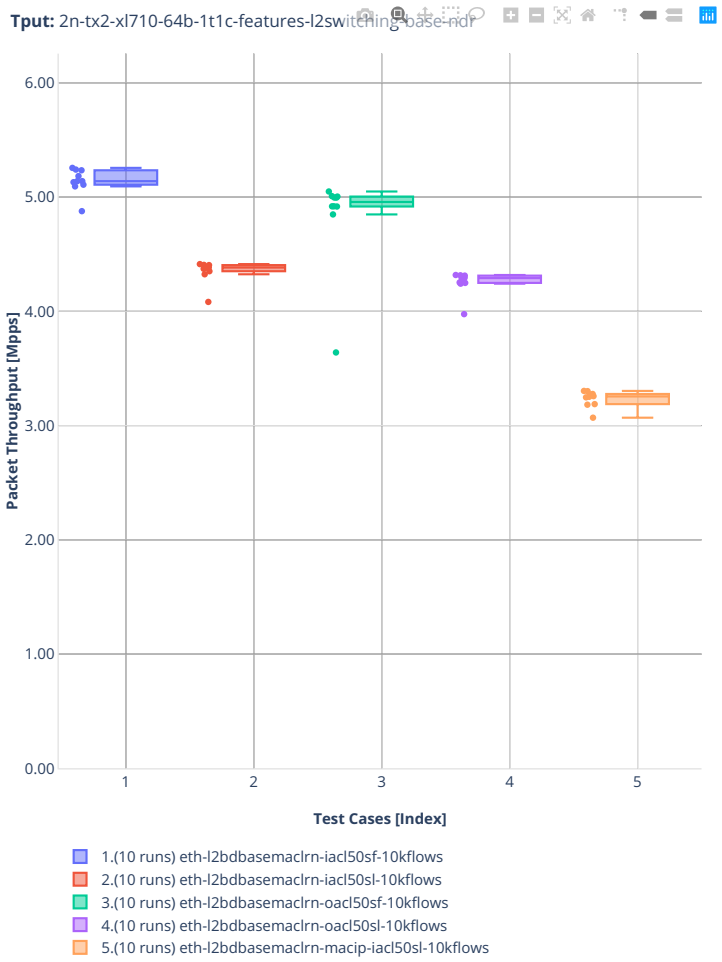

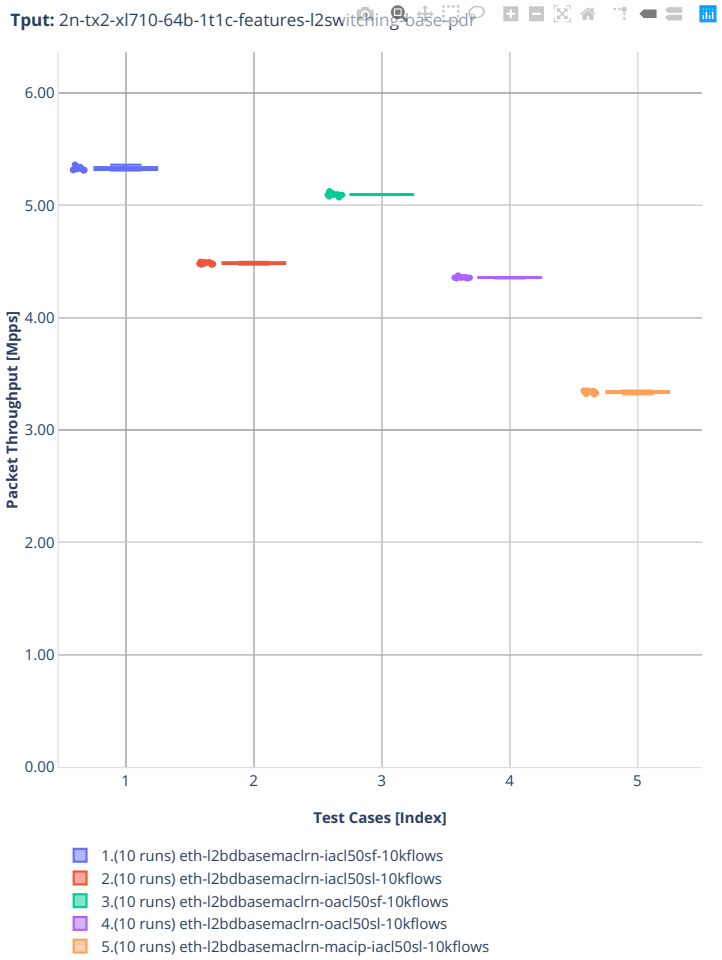
Tput: 2n-tx2-xl710-64b-1t1c-features-l2switching-base-pdr
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
Test Cases [Index]
1.(10 runs) eth-l2bdbasemaclrn-iacl50sf-10kflows
2.(10 runs) eth-l2bdbasemaclrn-iacl50sl-10kflows
3.(10 runs) eth-l2bdbasemaclrn-oacl50sf-10kflows
4.(10 runs) eth-l2bdbasemaclrn-oacl50sl-10kflows
5.(10 runs) eth-l2bdbasemaclrn-macip-iacl50sl-10kflows

2n-dnv-x553

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

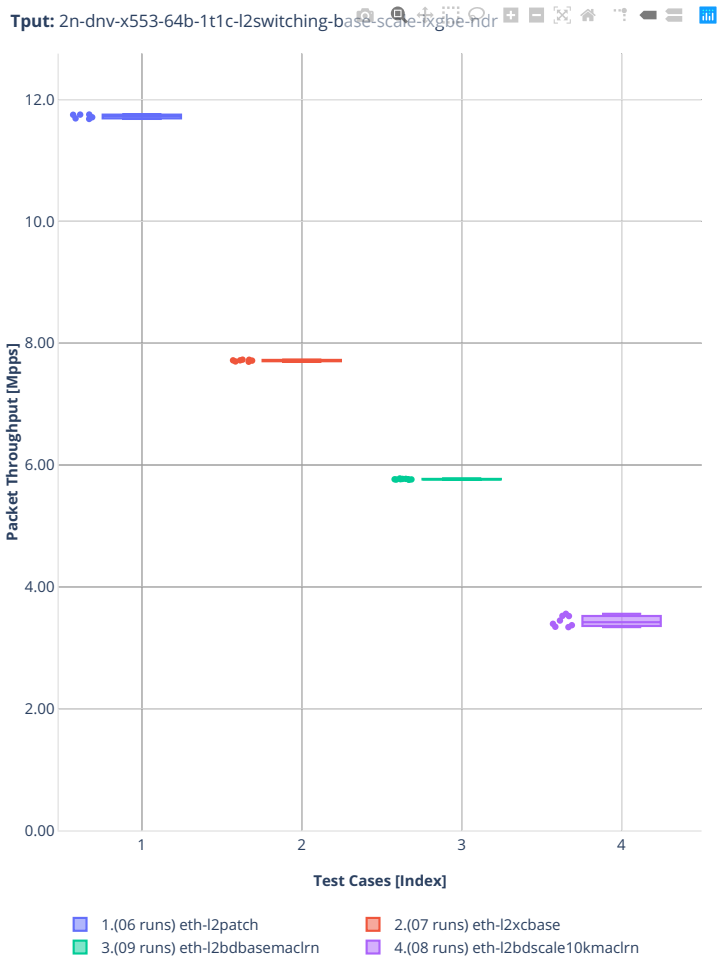

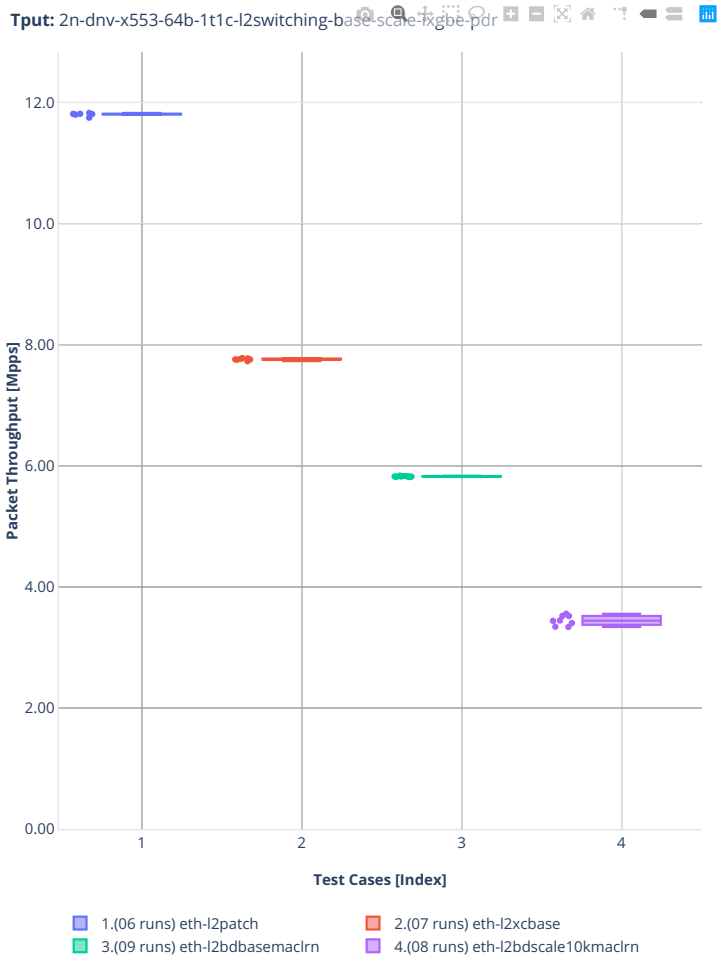
Tput: 2n-dnv-x553-64b-1t1c-l2switching-base-scale-ixgbe-pdr
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(06 runs) eth-l2patch
2.(07 runs) eth-l2xcbase
3.(09 runs) eth-l2bdbasemaclrn
4.(08 runs) eth-l2bdscale10kmaclrn

3n-dnv-x553

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

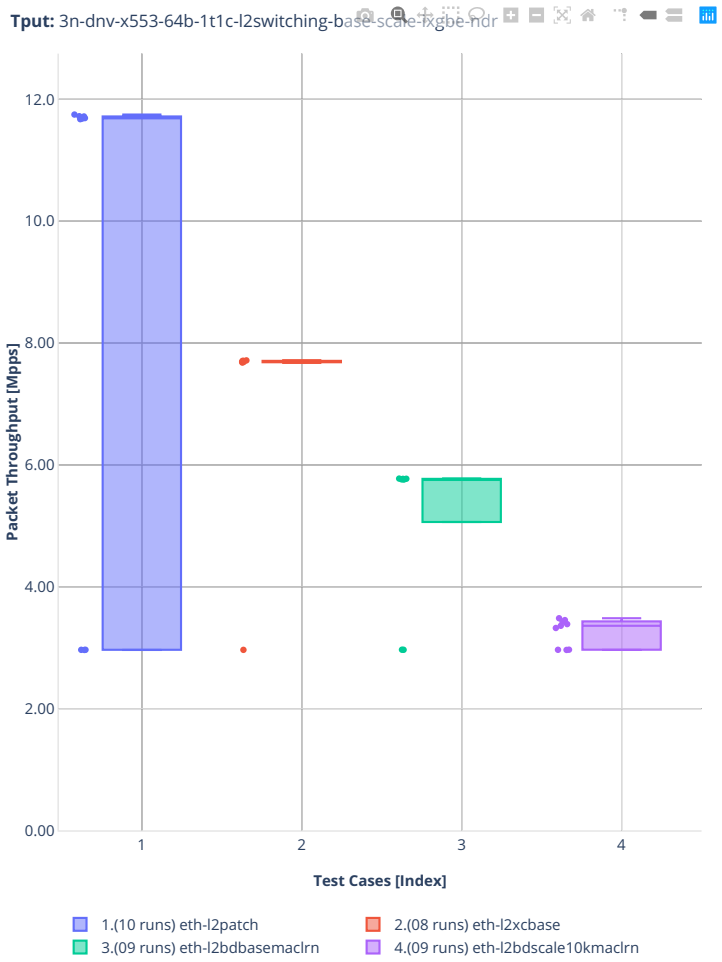

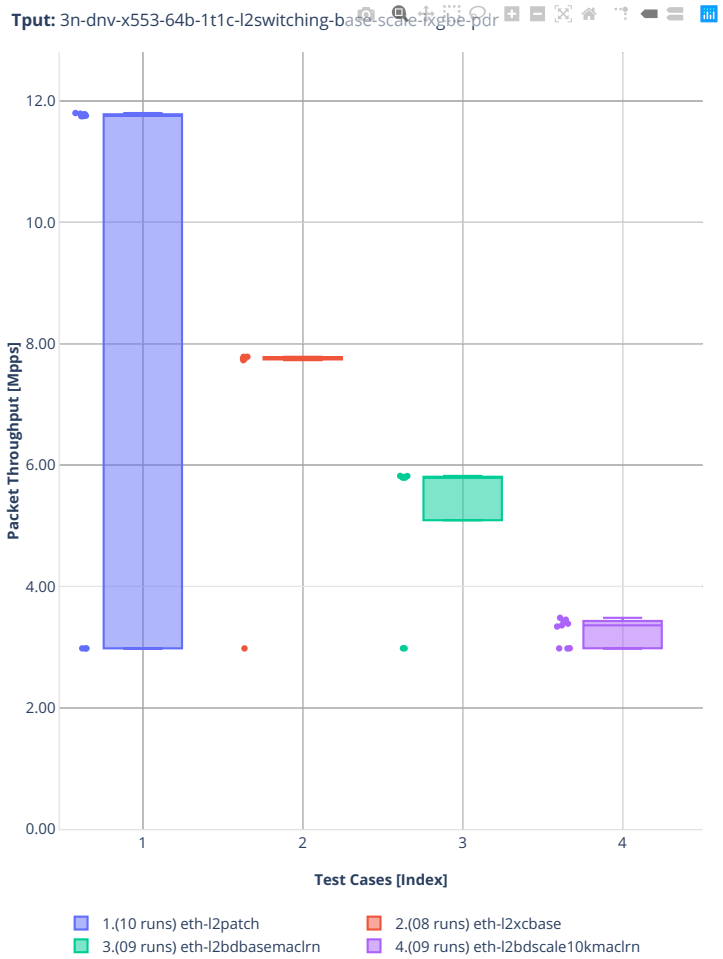
Tput: 3n-dnv-x553-64b-1t1c-l2switching-base-scale-ixgbe-pdr
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) eth-l2patch
2.(08 runs) eth-l2xcbase
3.(09 runs) eth-l2bdbasemaclrn
4.(09 runs) eth-l2bdscale10kmaclrn

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[110].

[110] https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls2206

2n-icx-xxv710

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

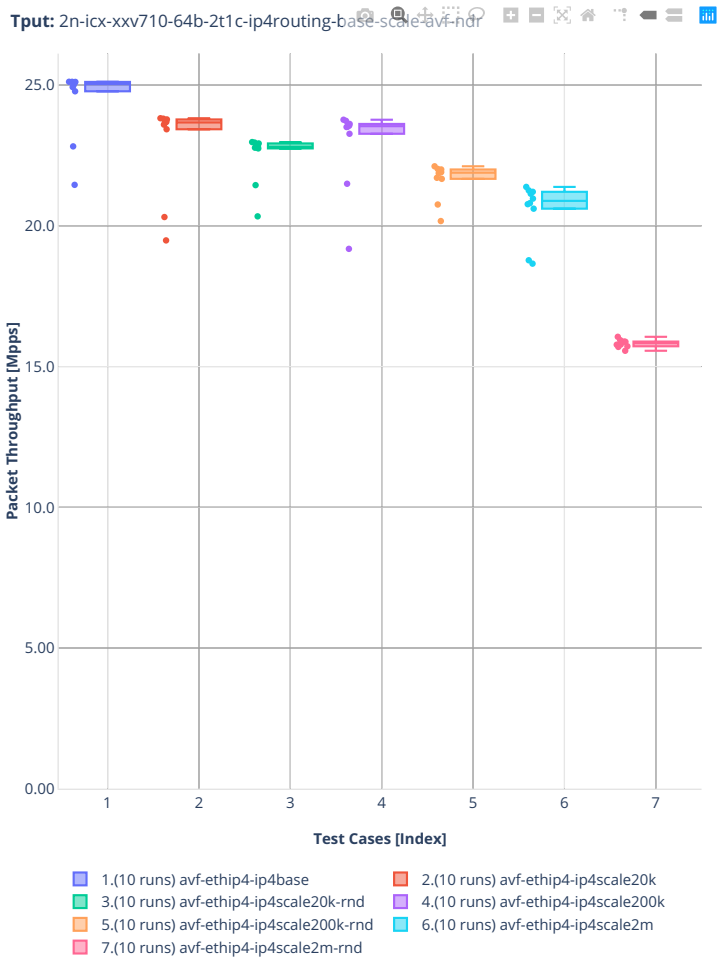

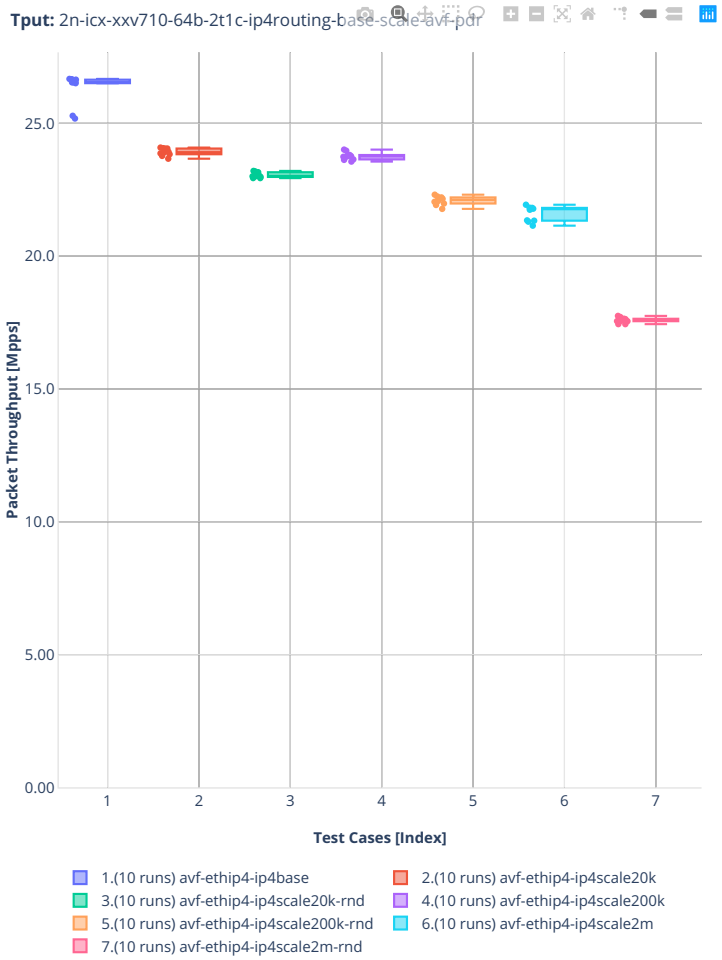
Tput: 2n-icx-xxv710-64b-2t1c-ip4routing-base-scale-avf-pdr
Packet Throughput [Mpps]
25.0
20.0
15.0
10.0
5.00
0.00
1
2
3
4
5
6
7
Test Cases [Index]
1.(10 runs) avf-ethip4-ip4base
2.(10 runs) avf-ethip4-ip4scale20k
3.(10 runs) avf-ethip4-ip4scale20k-rnd
4.(10 runs) avf-ethip4-ip4scale200k
5.(10 runs) avf-ethip4-ip4scale200k-rnd
6.(10 runs) avf-ethip4-ip4scale2m
7.(10 runs) avf-ethip4-ip4scale2m-rnd

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

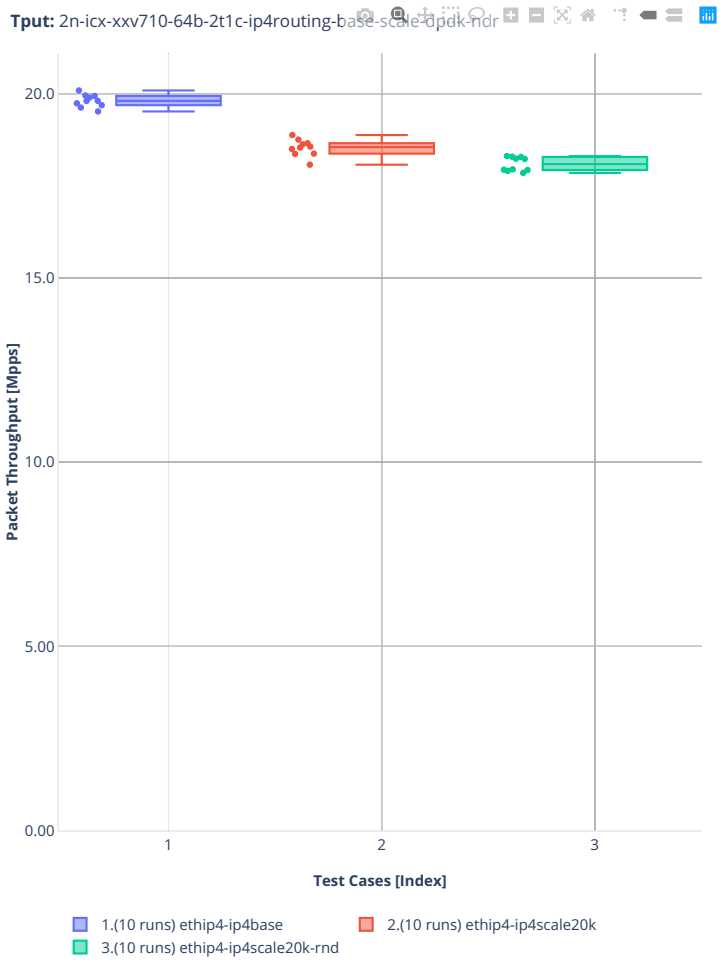

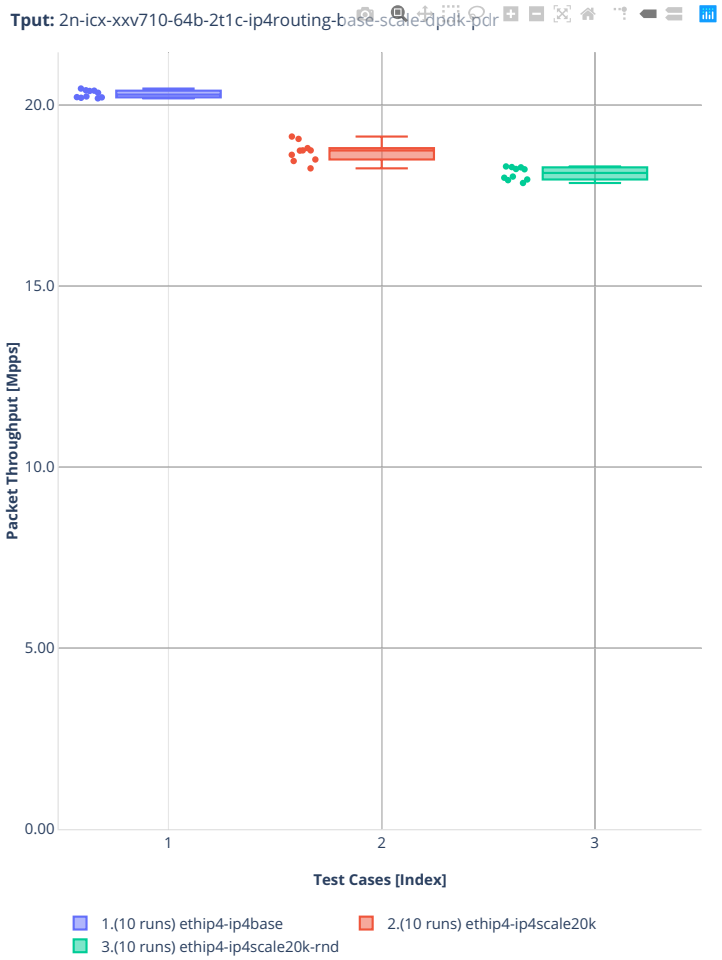
Tput: 2n-icx-xxv710-64b-2t1c-ip4routing-base-scale-dpdk-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip4-ip4base
2.(10 runs) ethip4-ip4scale20k
3.(10 runs) ethip4-ip4scale20k-rnd

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

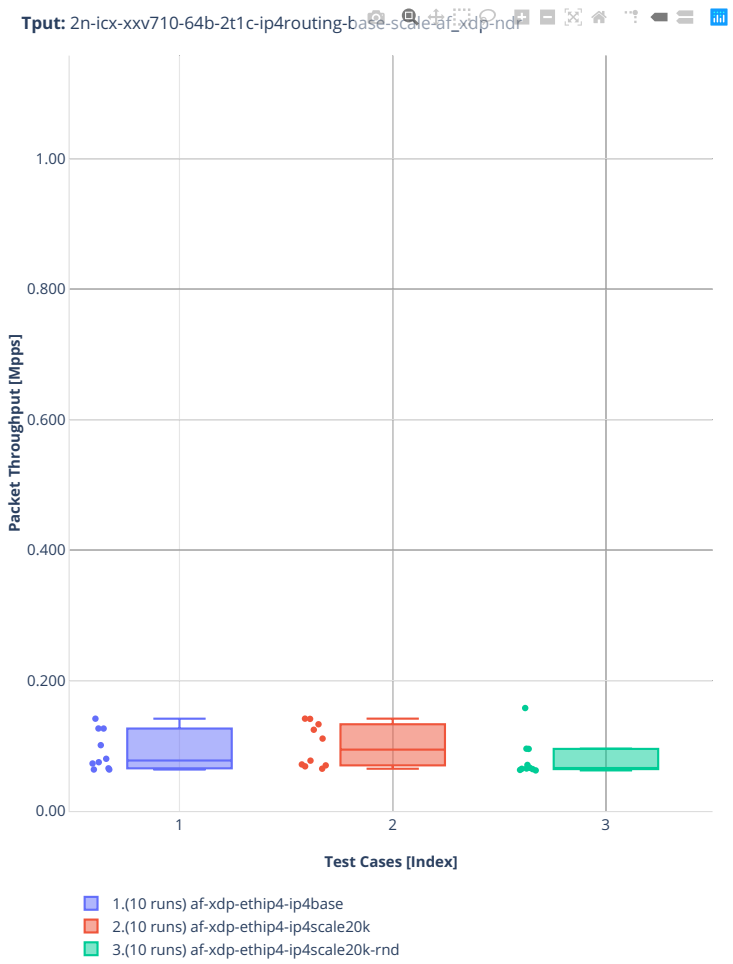

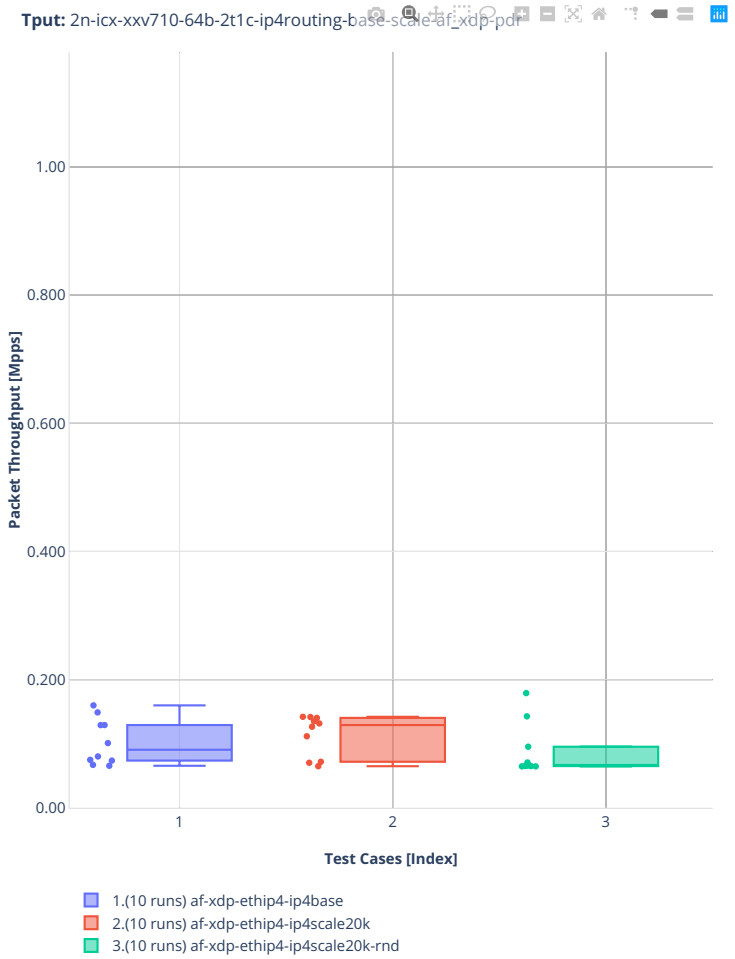
Tput: 2n-icx-xxv710-64b-2t1c-ip4routing-base-scale-af_xdp-pdr
Packet Throughput [Mpps]
1.00
0.800
0.600
0.400
0.200
0.00
1
2
3
Test Cases [Index]
1.(10 runs) af-xdp-ethip4-ip4base
2.(10 runs) af-xdp-ethip4-ip4scale20k
3.(10 runs) af-xdp-ethip4-ip4scale20k-rnd

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

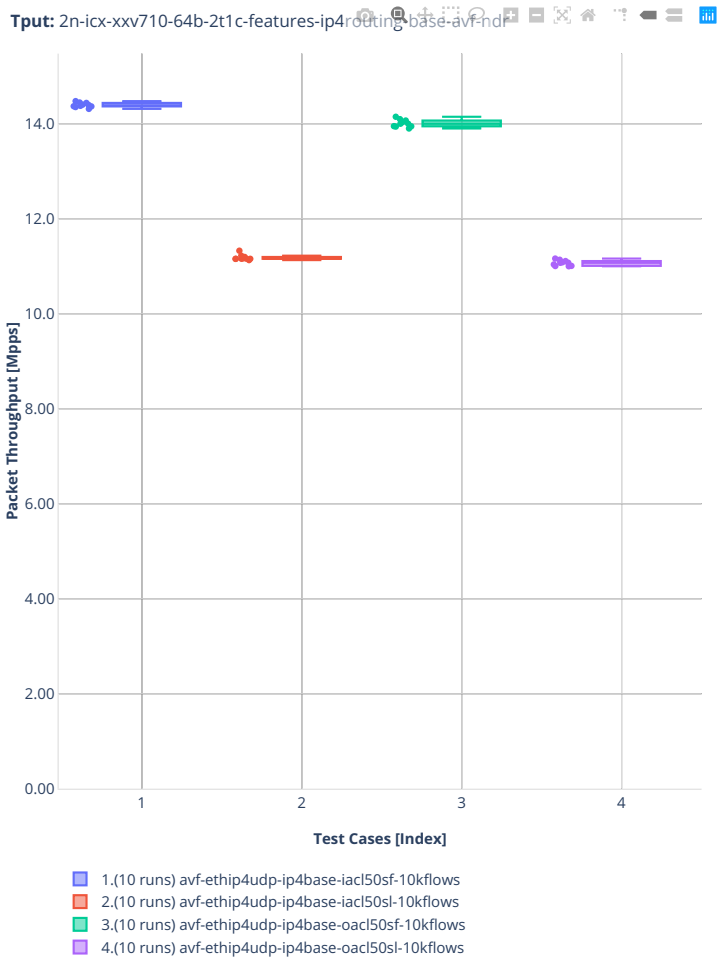

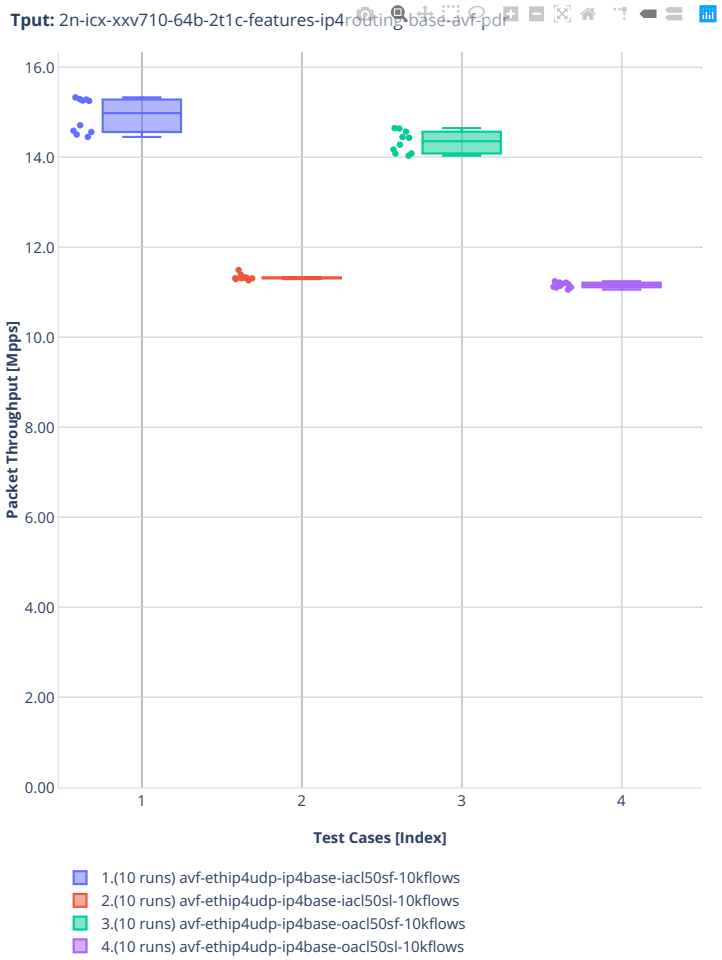
Tput: 2n-icx-xxv710-64b-2t1c-features-ip4routing-base-avf-pdr
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) avf-ethip4udp-ip4base-iacl50sf-10kflows
2.(10 runs) avf-ethip4udp-ip4base-iacl50sl-10kflows
3.(10 runs) avf-ethip4udp-ip4base-oacl50sf-10kflows
4.(10 runs) avf-ethip4udp-ip4base-oacl50sl-10kflows

3n-icx-xxv710

64b-2t1c-ip4routing-base

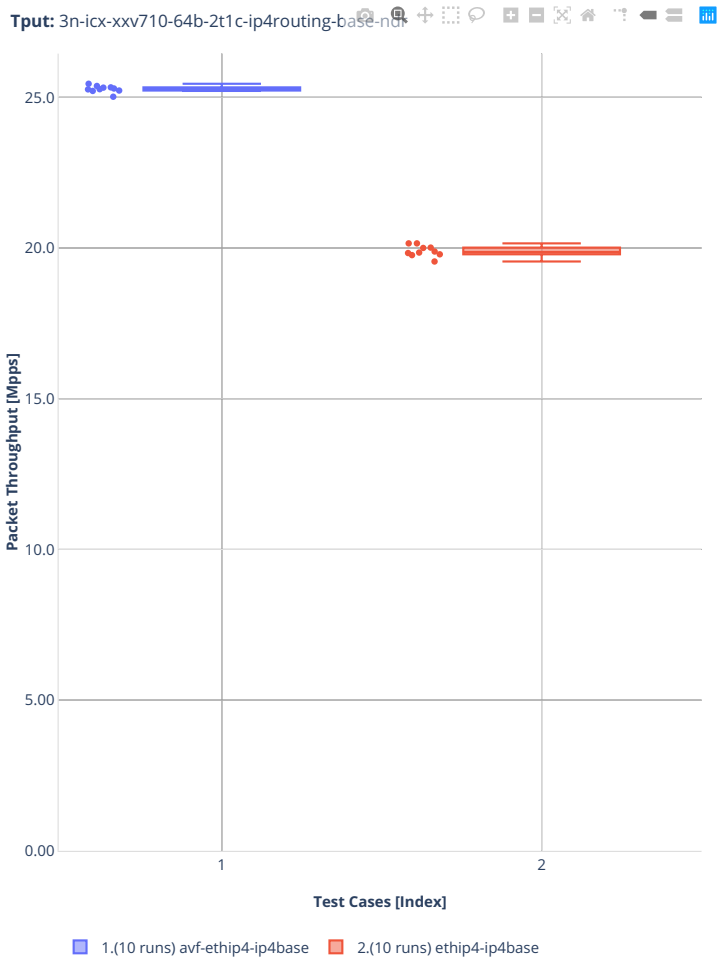

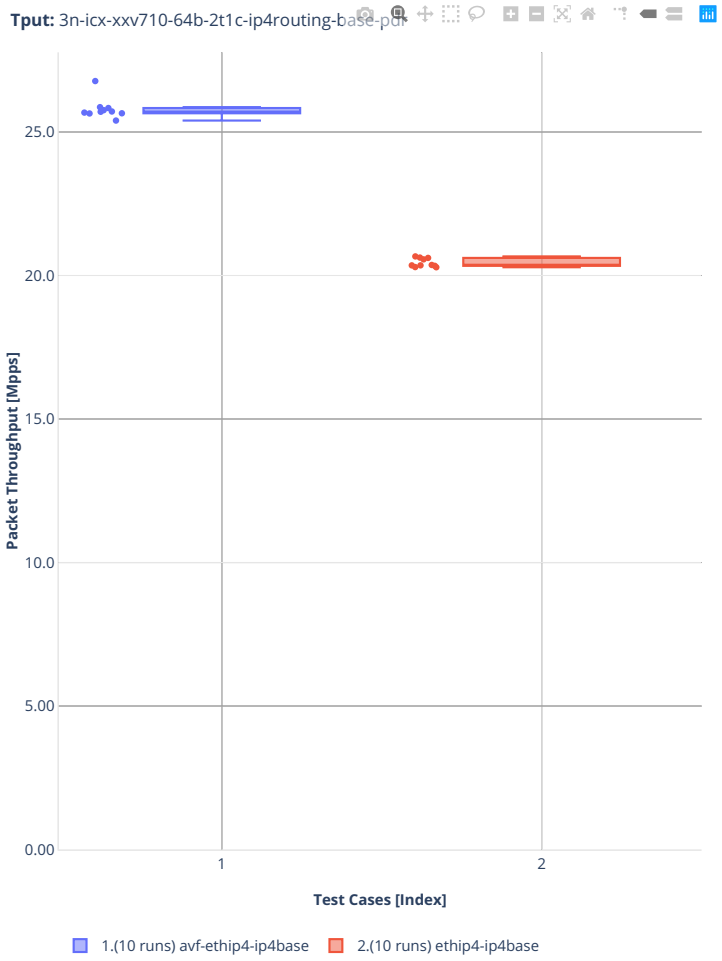
Tput: 3n-icx-xxv710-64b-2t1c-ip4routing-base-pdr
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) avf-ethip4-ip4base
2.(10 runs) ethip4-ip4base

2n-skx-xxv710

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

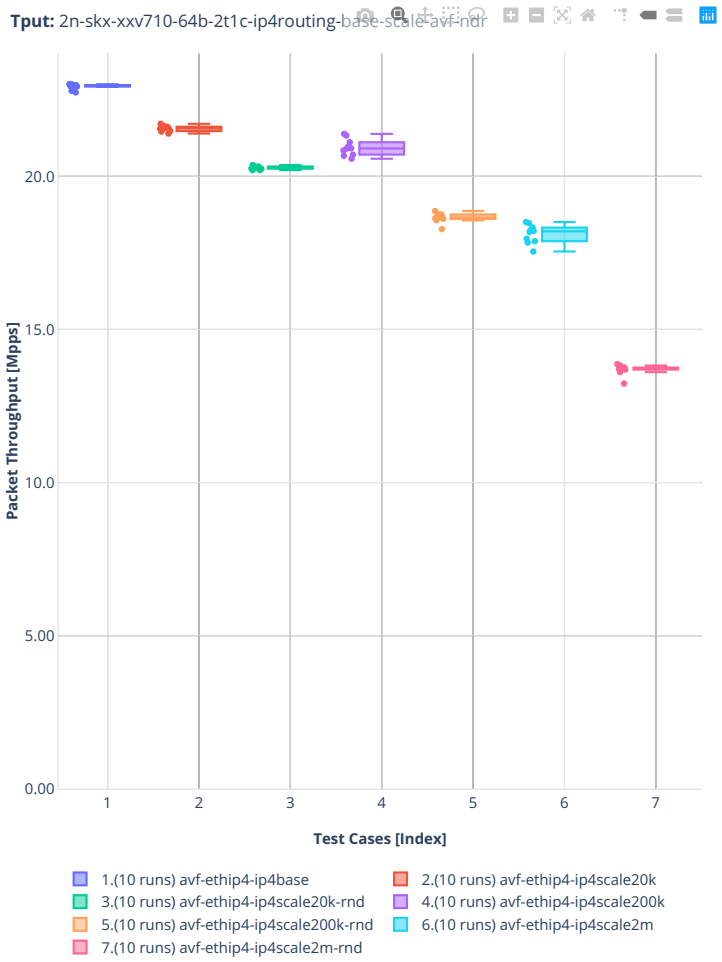

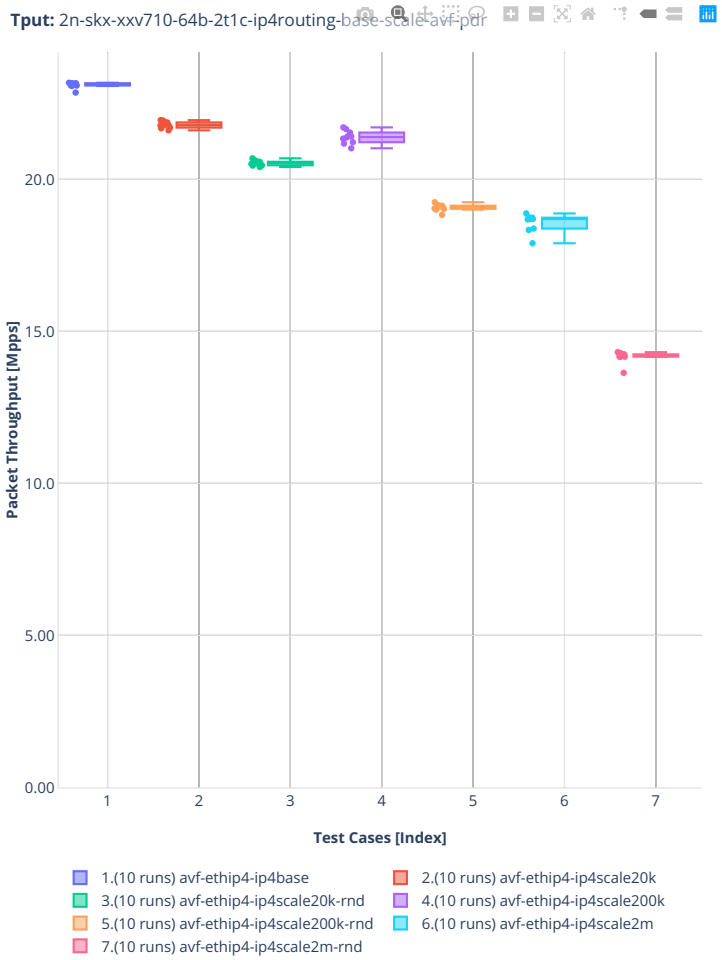
Tput: 2n-skx-xxv710-64b-2t1c-ip4routing-base-scale-avf-pdr
Packet Throughput [Mpps]
0.00
5.00
10.0
15.0
20.0
1
2
3
4
5
6
7
Test Cases [Index]
1.(10 runs) avf-ethip4-ip4base
2.(10 runs) avf-ethip4-ip4scale20k
3.(10 runs) avf-ethip4-ip4scale20k-rnd
4.(10 runs) avf-ethip4-ip4scale200k
5.(10 runs) avf-ethip4-ip4scale200k-rnd
6.(10 runs) avf-ethip4-ip4scale2m
7.(10 runs) avf-ethip4-ip4scale2m-rnd

64b-2t1c-ip4routing-base-scale-af-xdp

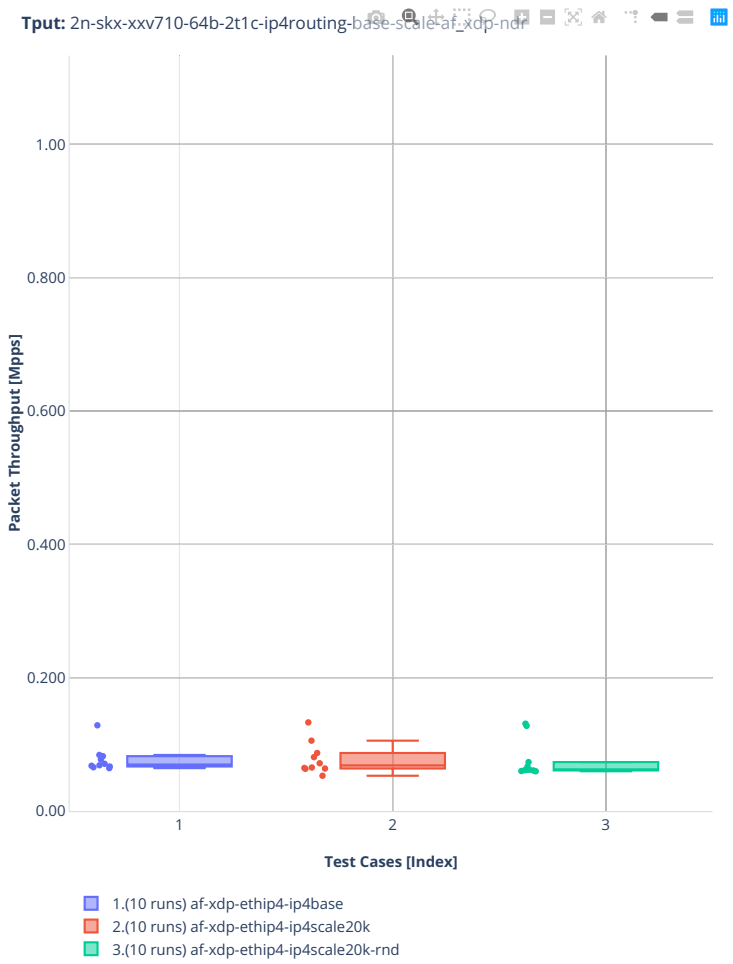

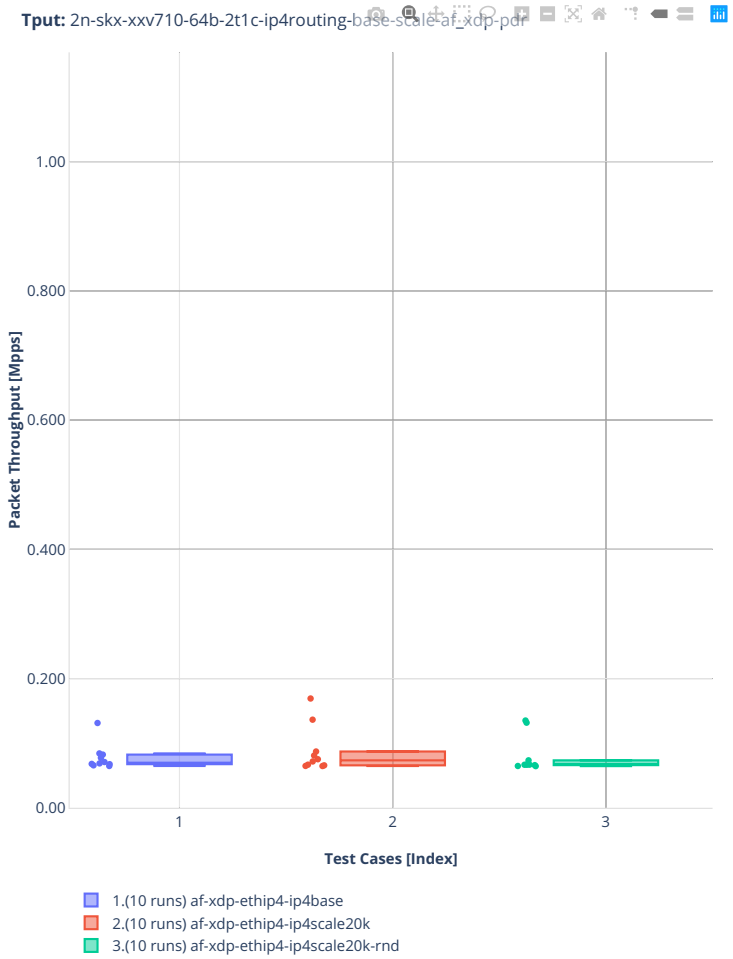
Tput: 2n-skx-xxv710-64b-2t1c-ip4routing-base-scale-af_xdp-pdr
Packet Throughput [Mpps]
1.00
0.800
0.600
0.400
0.200
0.00
1
2
3
Test Cases [Index]
1.(10 runs) af-xdp-ethip4-ip4base
2.(10 runs) af-xdp-ethip4-ip4scale20k
3.(10 runs) af-xdp-ethip4-ip4scale20k-rnd

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

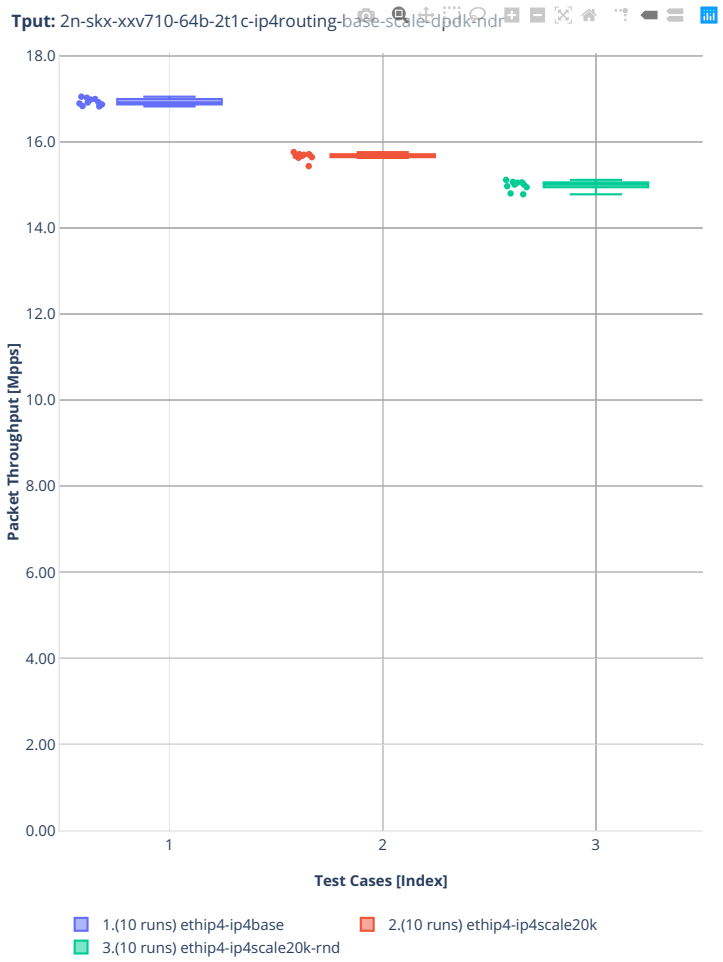

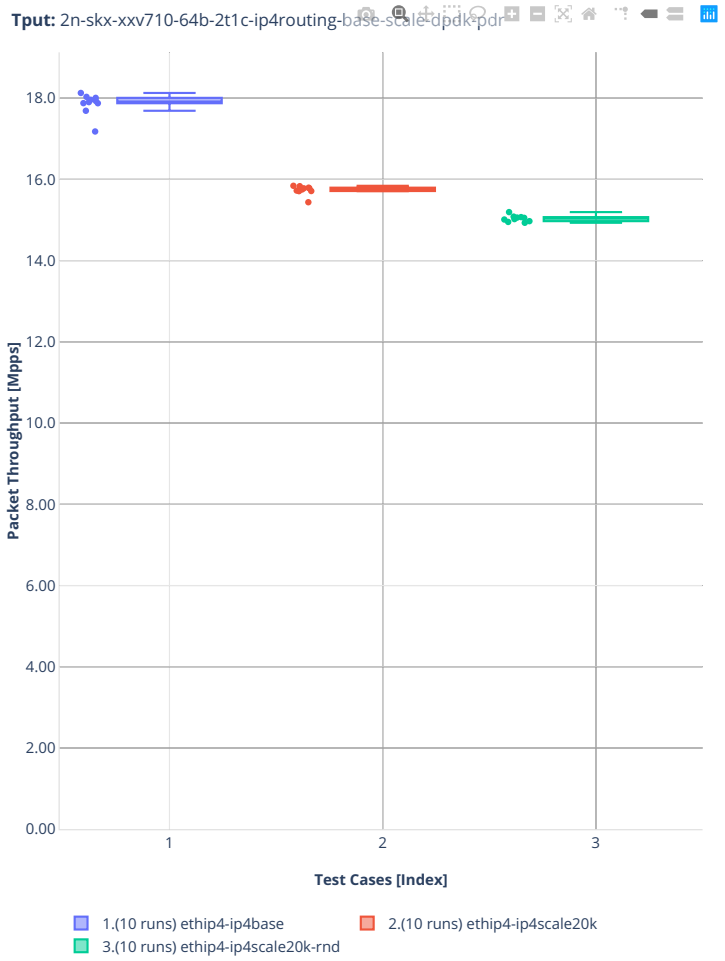
Tput: 2n-skx-xxv710-64b-2t1c-ip4routing-base-scale-dpdk-pdr
18.0
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip4-ip4base
2.(10 runs) ethip4-ip4scale20k
3.(10 runs) ethip4-ip4scale20k-rnd

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

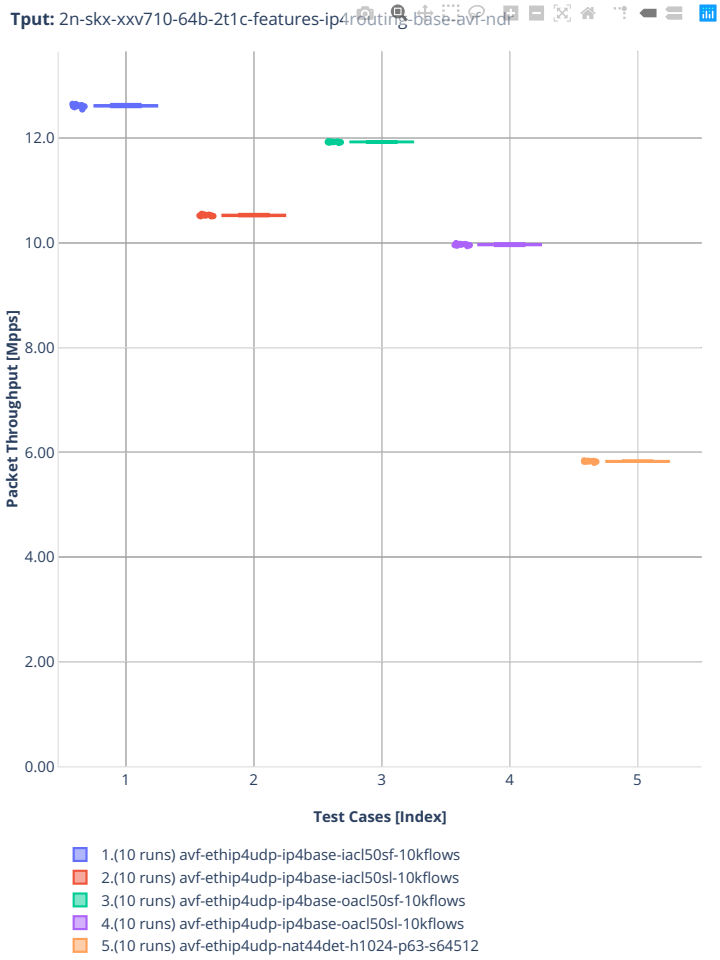

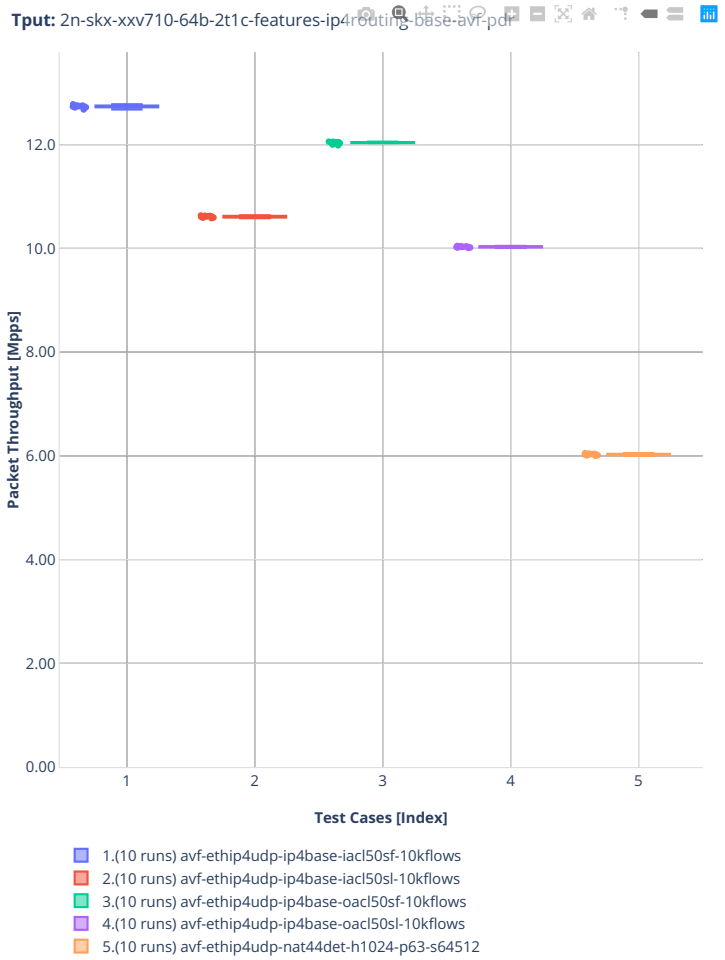
Tput: 2n-skx-xxv710-64b-2t1c-features-ip4routing-base-avf-pdr
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
Test Cases [Index]
1.(10 runs) avf-ethip4udp-ip4base-iacl50sf-10kflows
2.(10 runs) avf-ethip4udp-ip4base-iacl50sl-10kflows
3.(10 runs) avf-ethip4udp-ip4base-oacl50sf-10kflows
4.(10 runs) avf-ethip4udp-ip4base-oacl50sl-10kflows
5.(10 runs) avf-ethip4udp-nat44det-h1024-p63-s64512

2n-skx-x710

64b-2t1c-ip4routing-base-scale-[avf,dpdk]

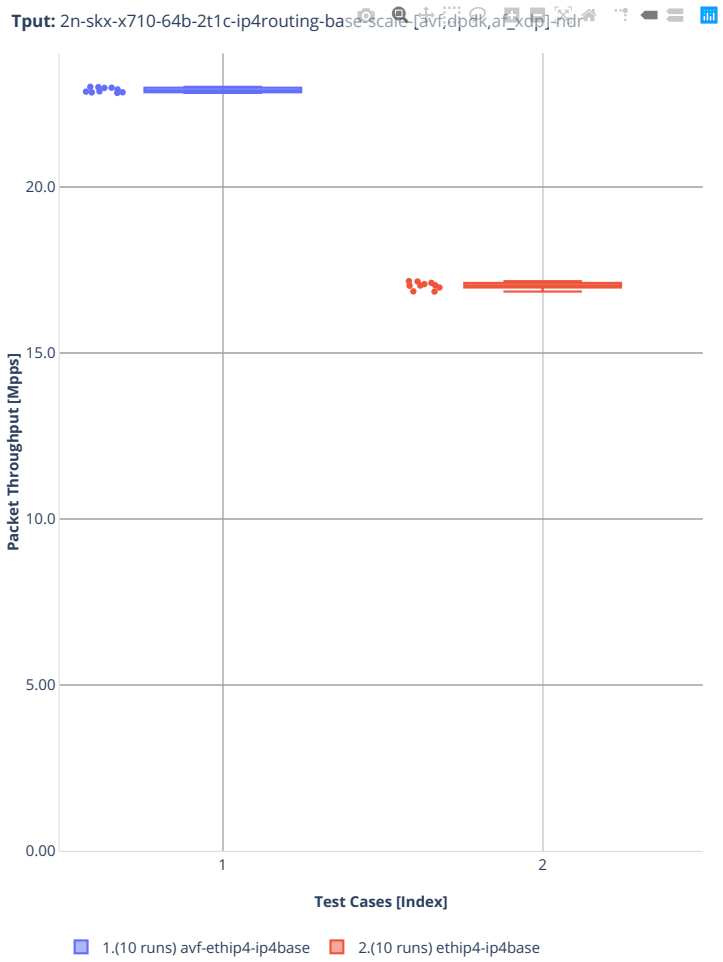

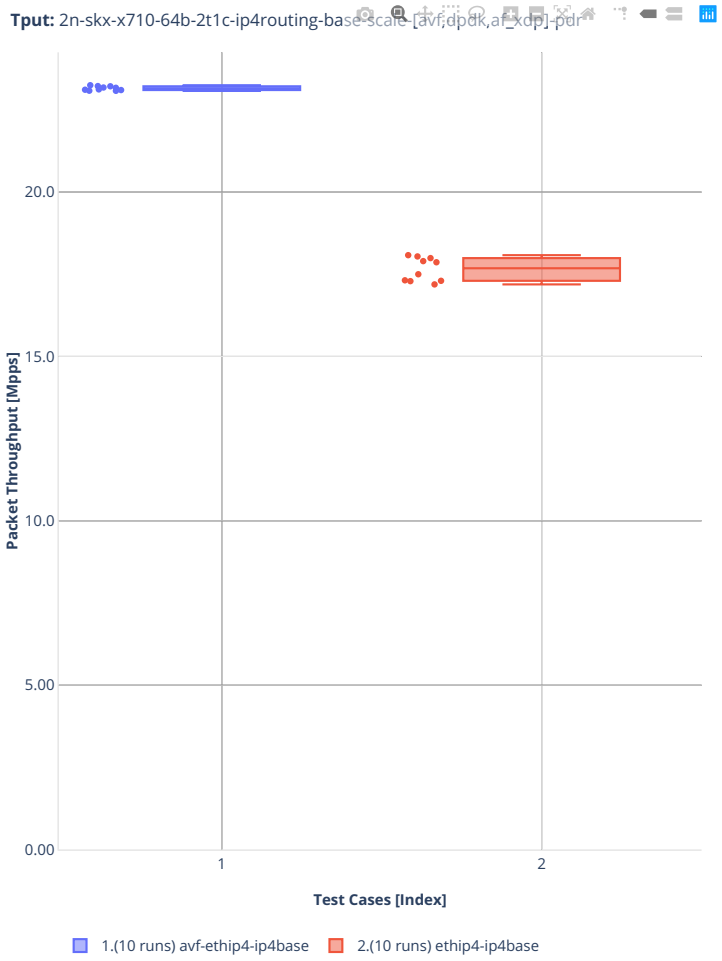
Tput: 2n-skx-x710-64b-2t1c-ip4routing-base-scale-[avf,dpdk,af_xdp]-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) avf-ethip4-ip4base
2.(10 runs) ethip4-ip4base

3n-skx-xxv710

64b-2t1c-ip4routing-base-[avf,dpdk]

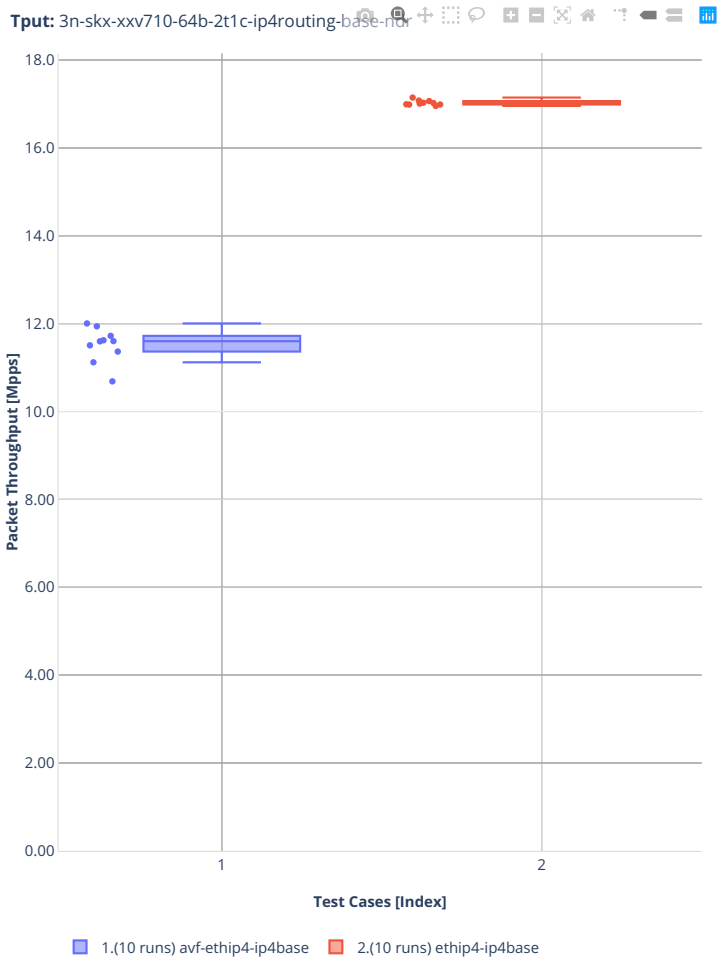

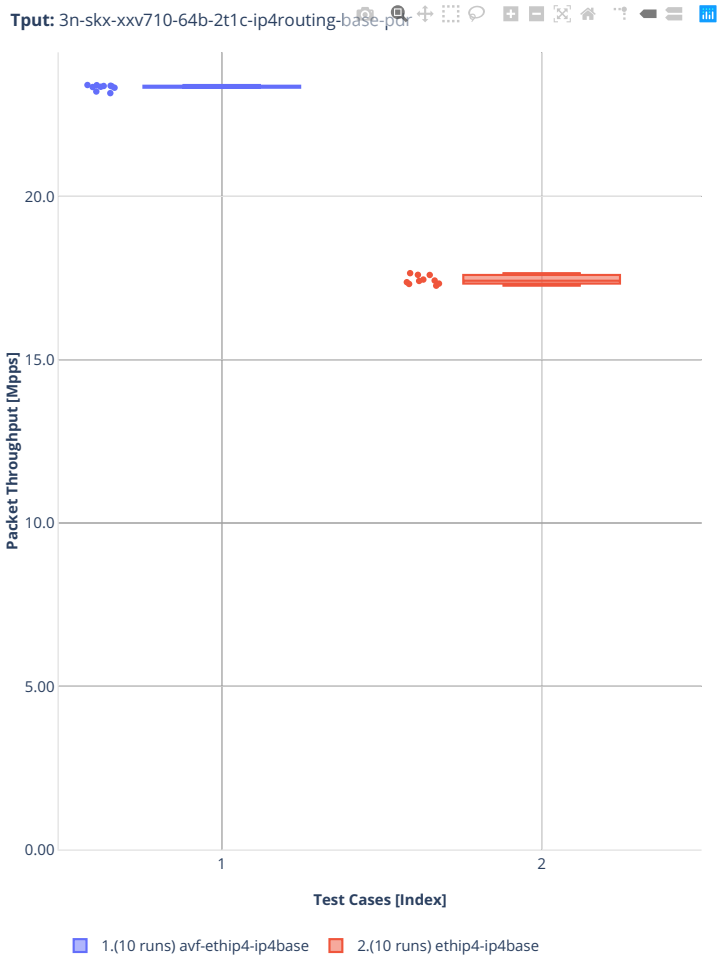
Tput: 3n-skx-xxv710-64b-2t1c-ip4routing-base-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) avf-ethip4-ip4base
2.(10 runs) ethip4-ip4base

3n-skx-x710

64b-2t1c-ip4routing-base-avf

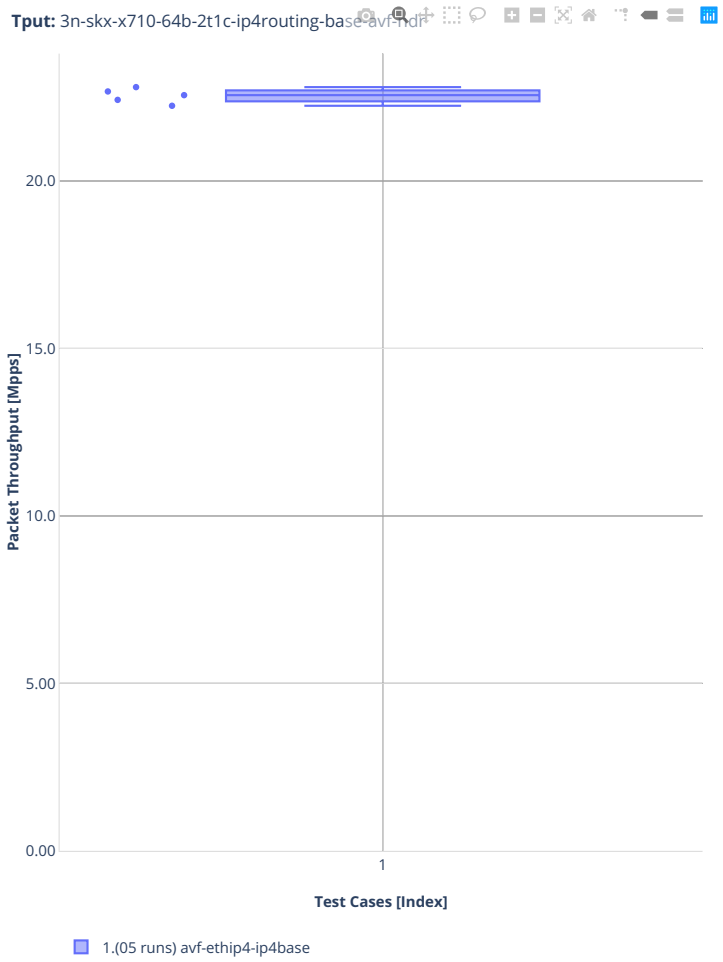

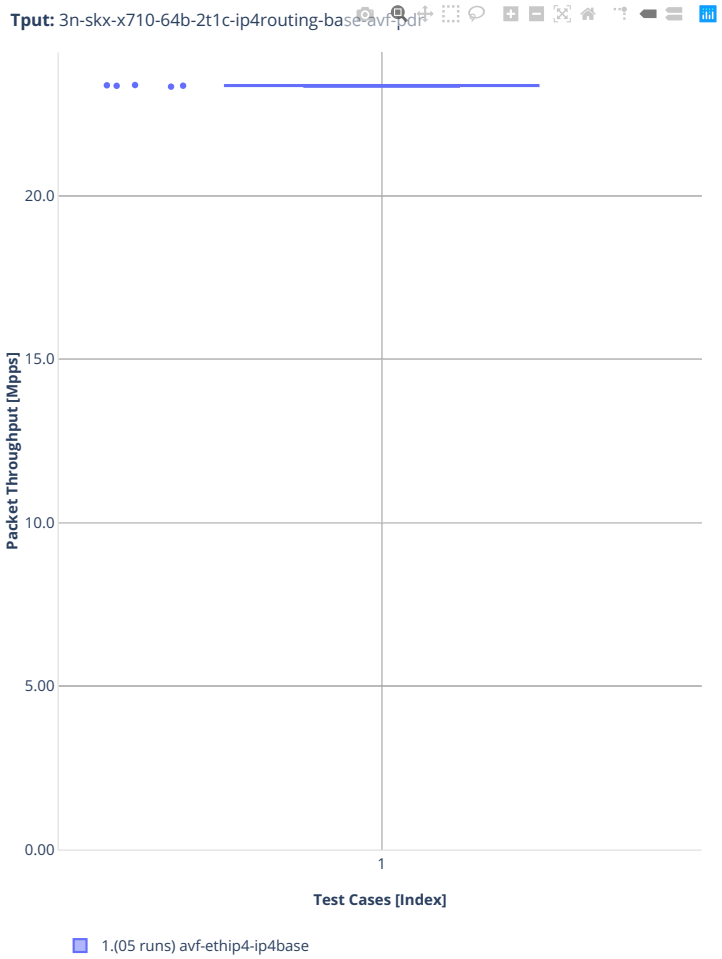
Tput: 3n-skx-x710-64b-2t1c-ip4routing-base-avf-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
Test Cases [Index]
1.(05 runs) avf-ethip4-ip4base

2n-clx-xxv710

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

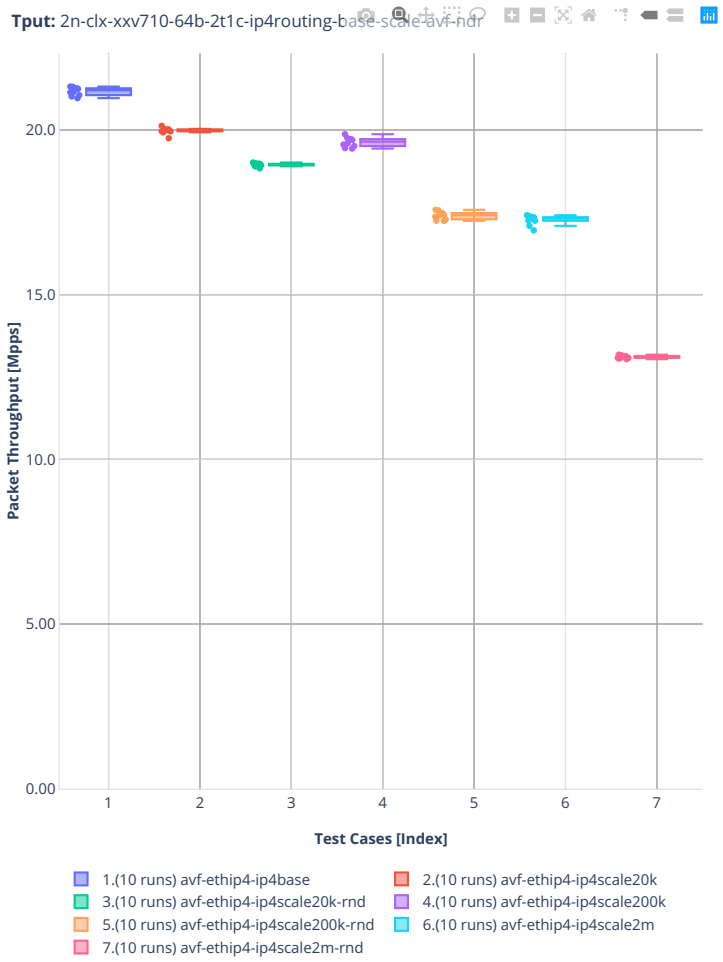

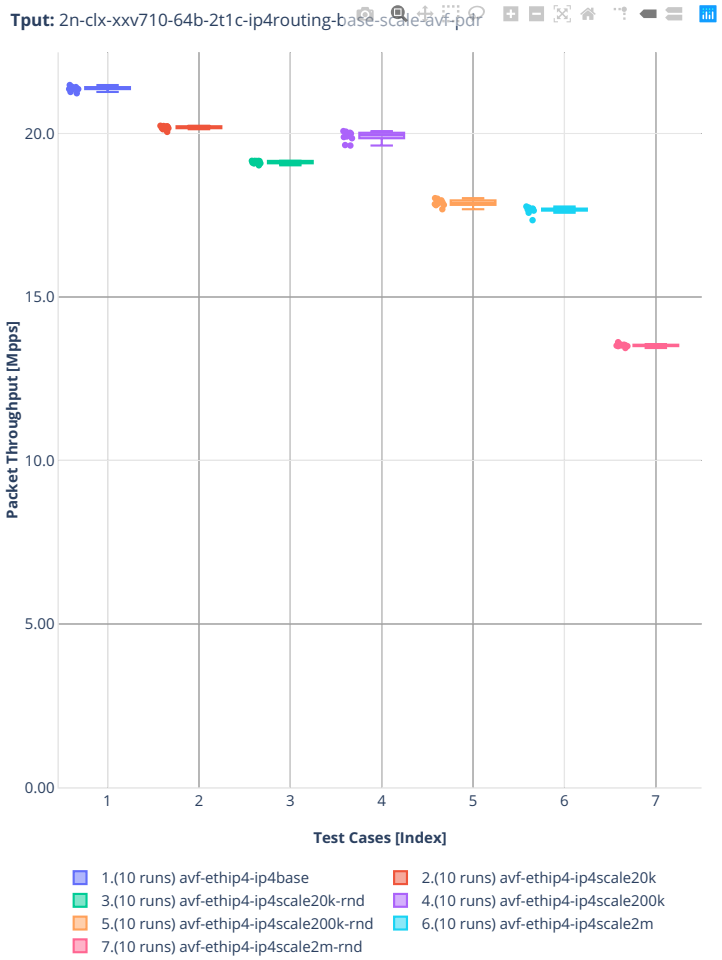
Tput: 2n-clx-xxv710-64b-2t1c-ip4routing-base-scale-avf-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
6
7
Test Cases [Index]
1.(10 runs) avf-ethip4-ip4base
2.(10 runs) avf-ethip4-ip4scale20k
3.(10 runs) avf-ethip4-ip4scale20k-rnd
4.(10 runs) avf-ethip4-ip4scale200k
5.(10 runs) avf-ethip4-ip4scale200k-rnd
6.(10 runs) avf-ethip4-ip4scale2m
7.(10 runs) avf-ethip4-ip4scale2m-rnd

64b-2t1c-ip4routing-base-scale-af-xdp

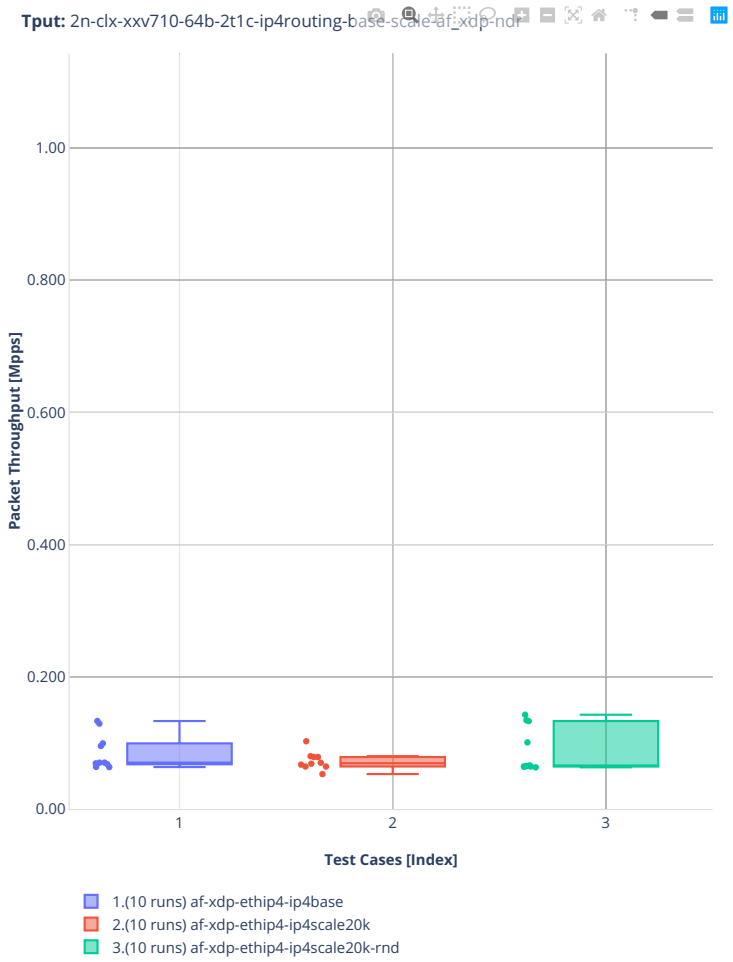

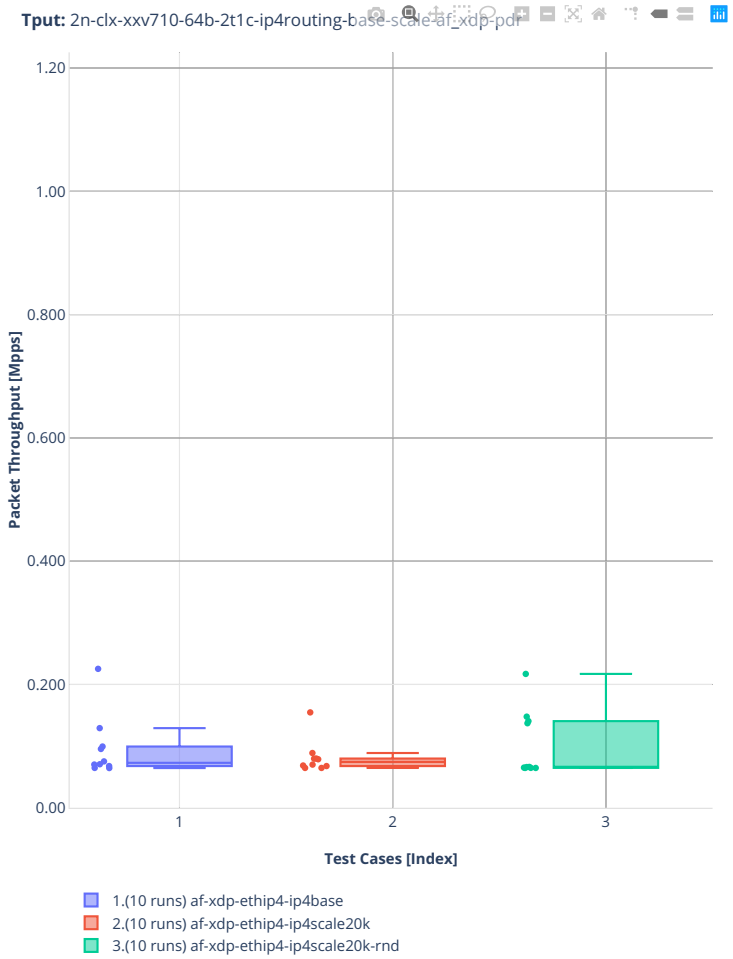
Tput: 2n-clx-xxv710-64b-2t1c-ip4routing-base-scale-af_xdp-pdr
1.20
1.00
0.800
0.600
0.400
0.200
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) af-xdp-ethip4-ip4base
2.(10 runs) af-xdp-ethip4-ip4scale20k
3.(10 runs) af-xdp-ethip4-ip4scale20k-rnd

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

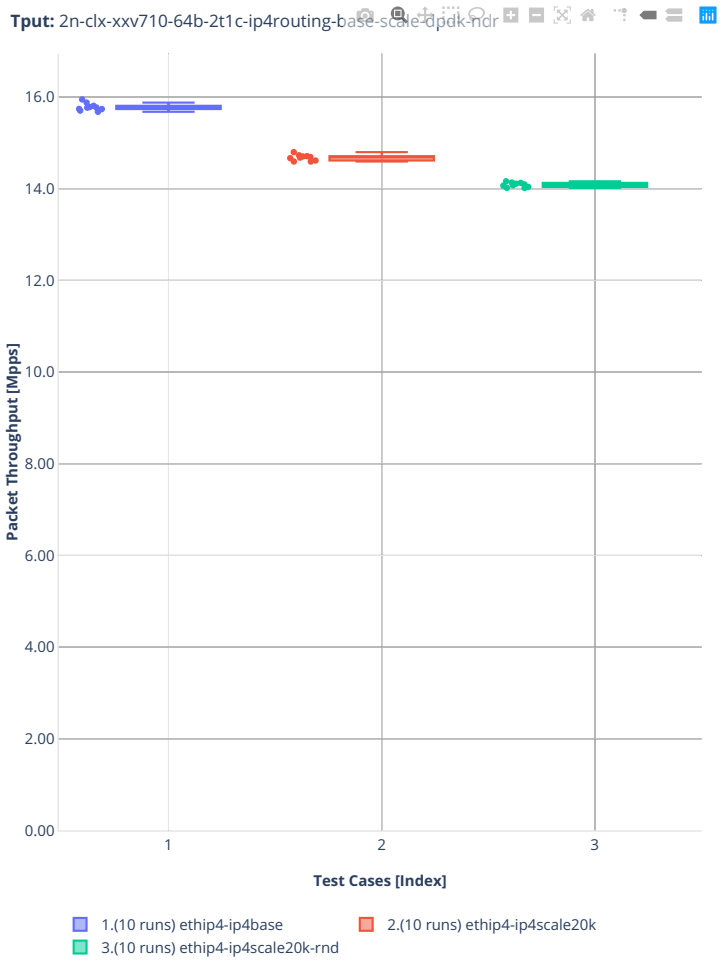

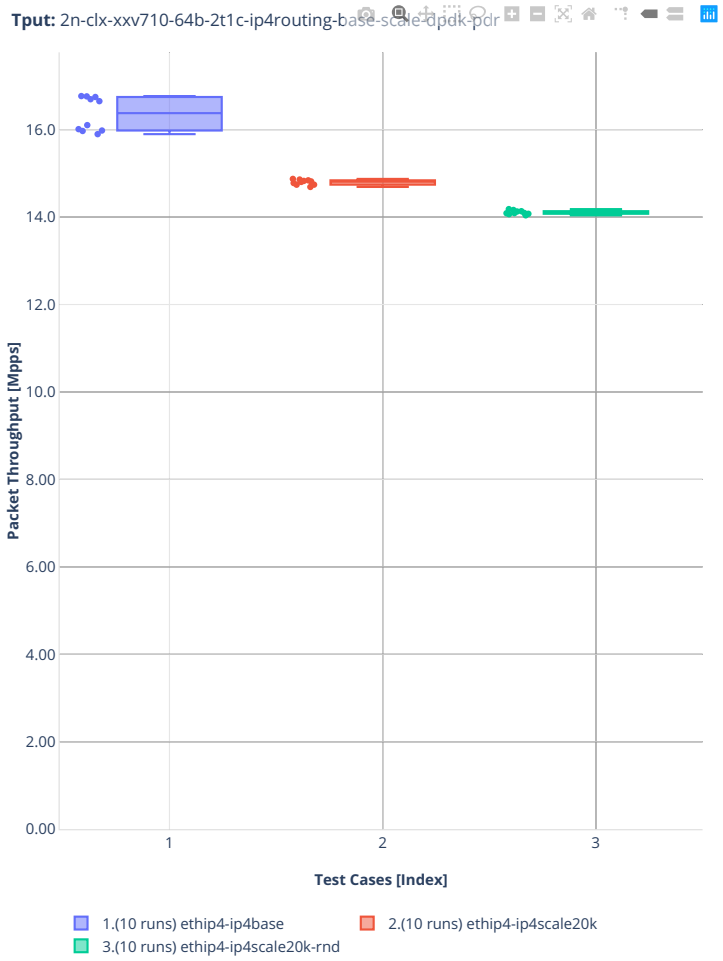
Tput: 2n-clx-xxv710-64b-2t1c-ip4routing-base-scale-dpdk-pdr
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip4-ip4base
2.(10 runs) ethip4-ip4scale20k
3.(10 runs) ethip4-ip4scale20k-rnd

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

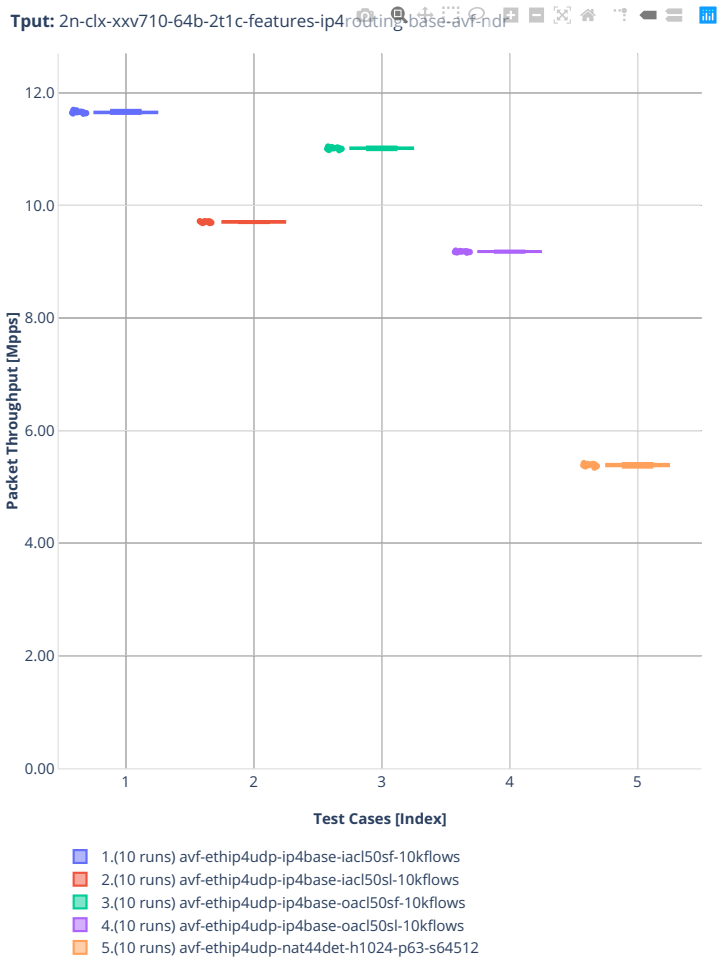

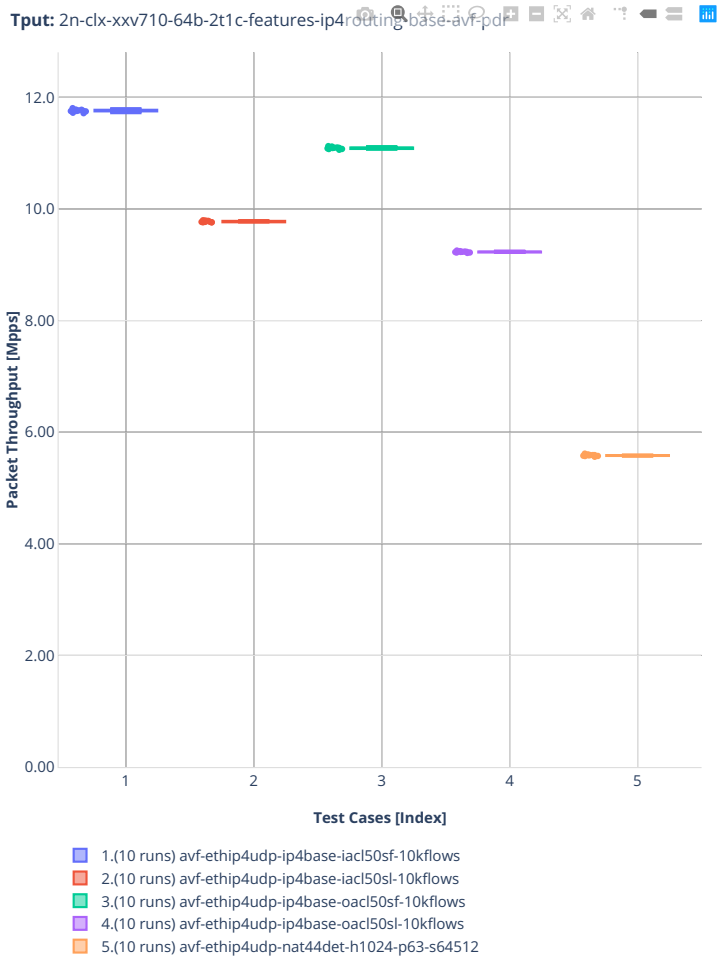
Tput: 2n-clx-xxv710-64b-2t1c-features-ip4routing-base-avf-pdr
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
Test Cases [Index]
1.(10 runs) avf-ethip4udp-ip4base-iacl50sf-10kflows
2.(10 runs) avf-ethip4udp-ip4base-iacl50sl-10kflows
3.(10 runs) avf-ethip4udp-ip4base-oacl50sf-10kflows
4.(10 runs) avf-ethip4udp-ip4base-oacl50sl-10kflows
5.(10 runs) avf-ethip4udp-nat44det-h1024-p63-s64512

2n-clx-x710

64b-2t1c-ip4routing-base-scale-[avf,dpdk]

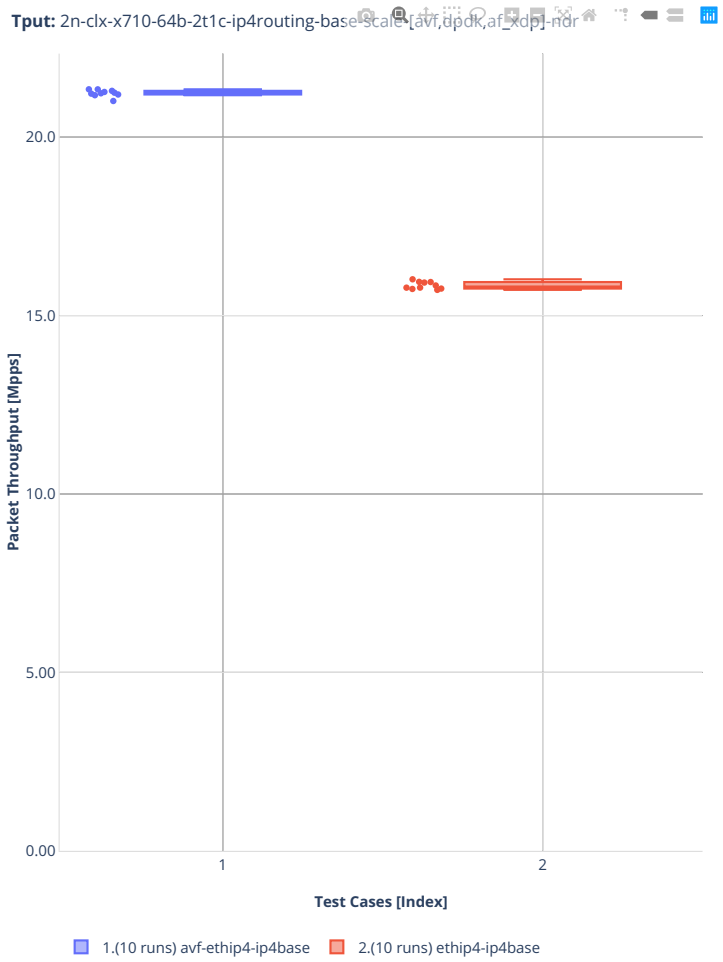

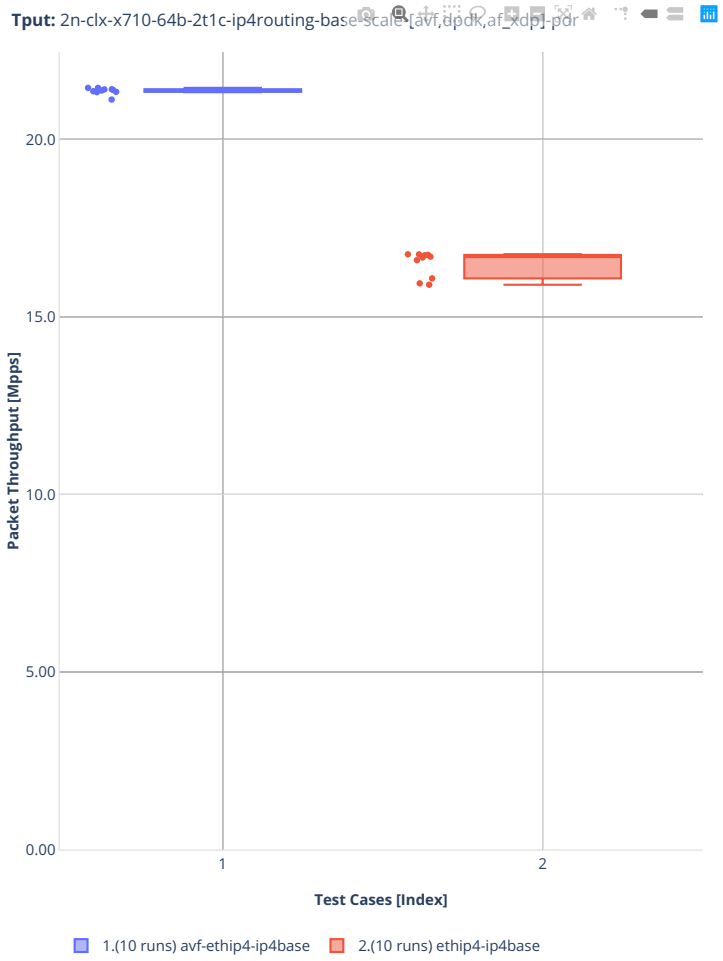
Tput: 2n-clx-x710-64b-2t1c-ip4routing-base-scale-[avf,dpdk,af_xdp]-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) avf-ethip4-ip4base
2.(10 runs) ethip4-ip4base

2n-clx-cx556a

64b-2t1c-ip4routing-base-scale-rdma-core

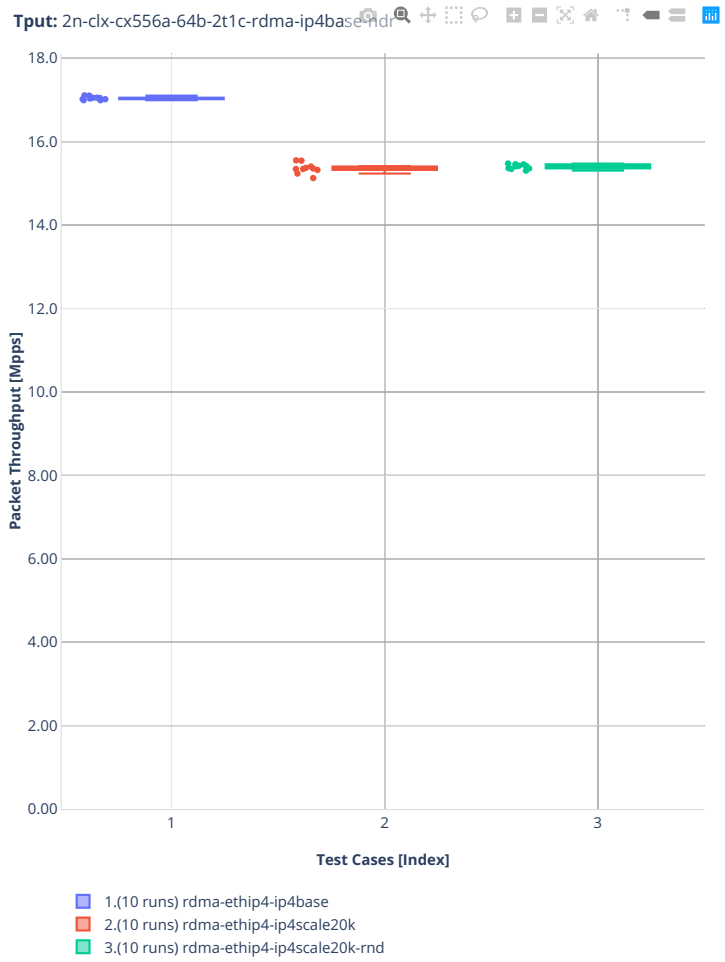

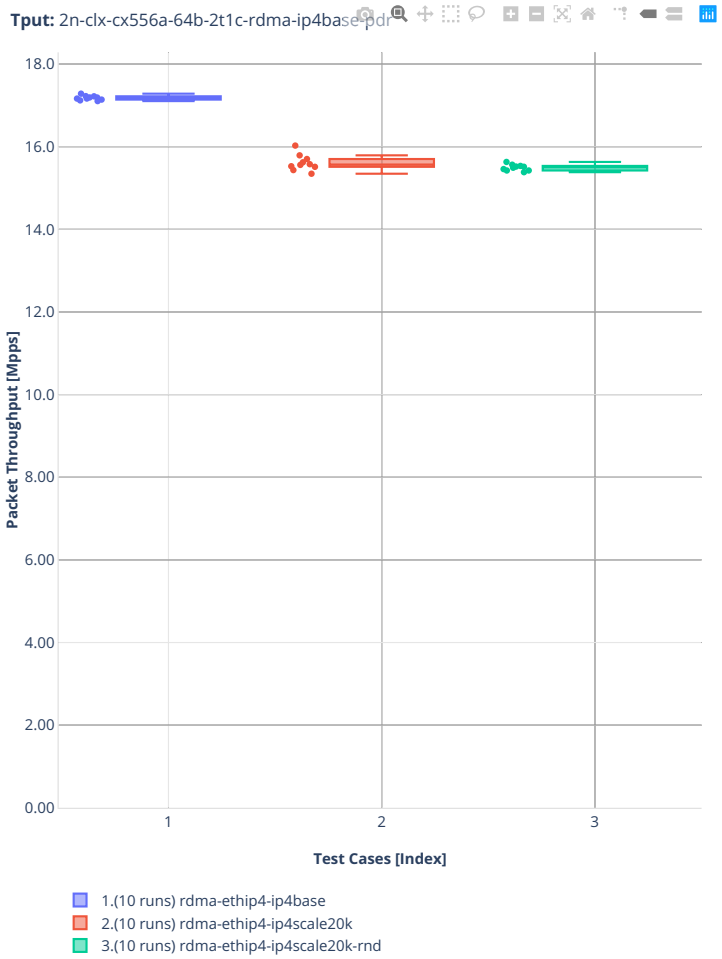
Tput: 2n-clx-cx556a-64b-2t1c-rdma-ip4base-pdr
18.0
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) rdma-ethip4-ip4base
2.(10 runs) rdma-ethip4-ip4scale20k
3.(10 runs) rdma-ethip4-ip4scale20k-rnd

64b-2t1c-ip4routing-features

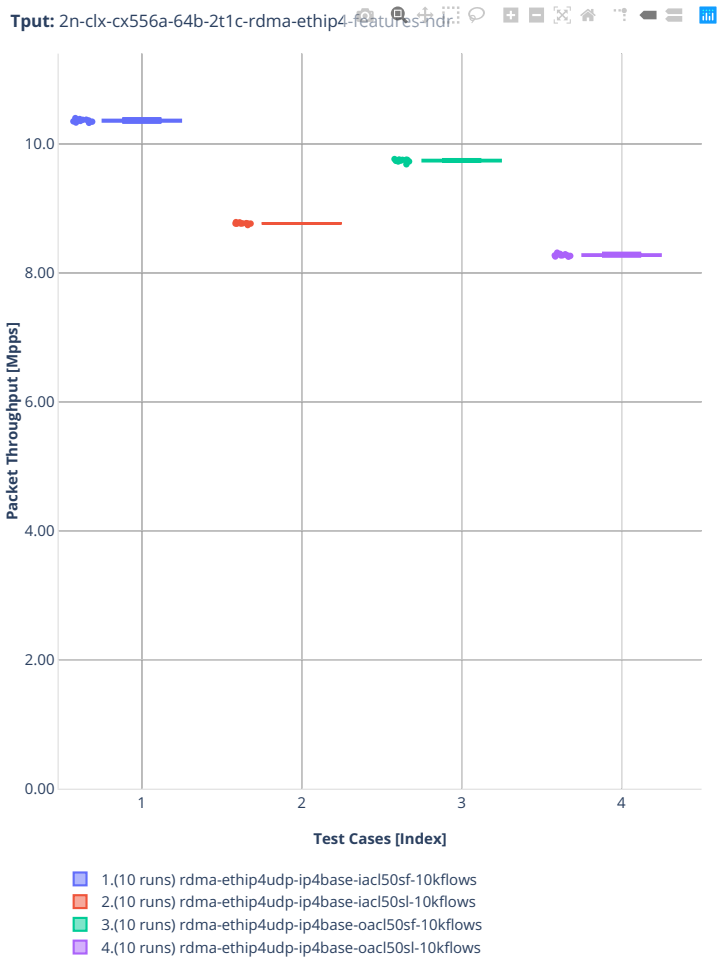

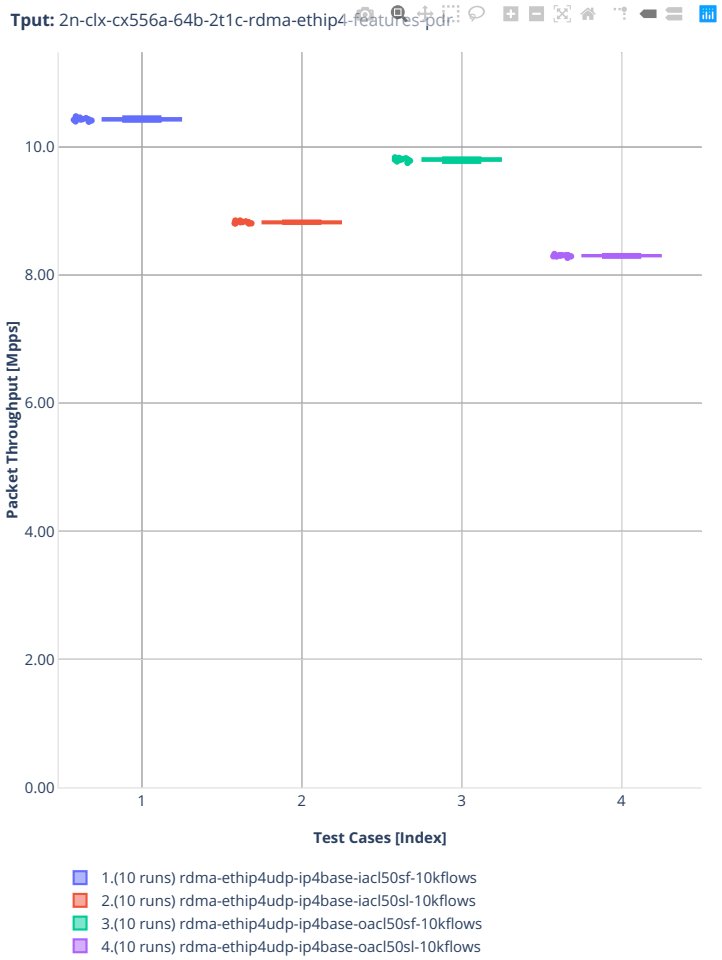
Tput: 2n-clx-cx556a-64b-2t1c-rdma-ethip4-features-pdr
Packet Throughput [Mpps]
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
Test Cases [Index]
1.(10 runs) rdma-ethip4udp-ip4base-iacl50sf-10kflows
2.(10 runs) rdma-ethip4udp-ip4base-iacl50sl-10kflows
3.(10 runs) rdma-ethip4udp-ip4base-oacl50sf-10kflows
4.(10 runs) rdma-ethip4udp-ip4base-oacl50sl-10kflows

2n-clx-e810cq

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

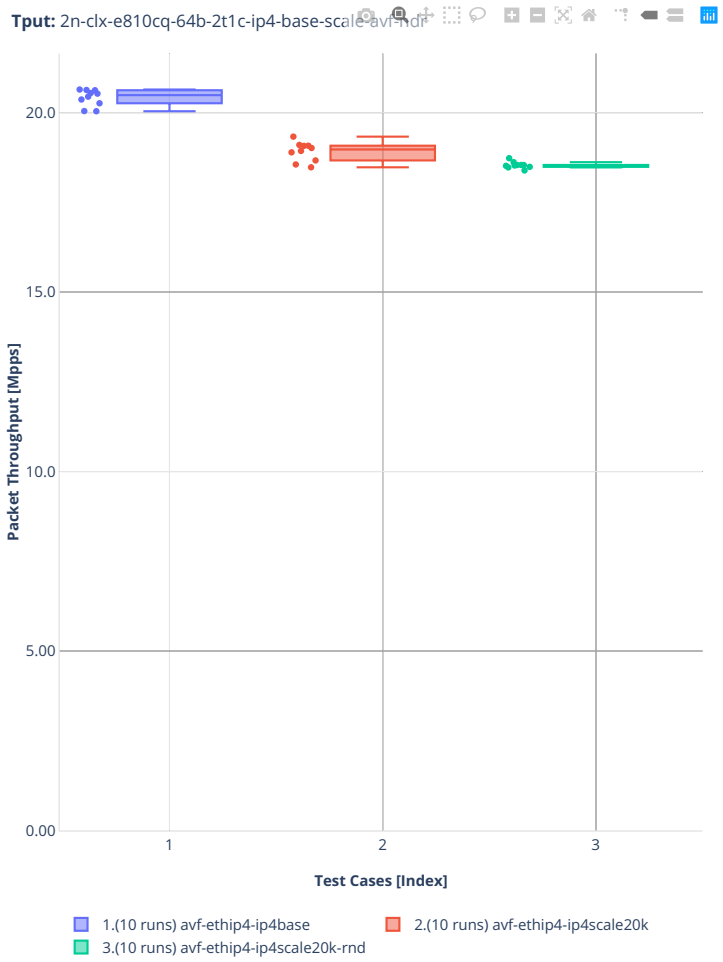

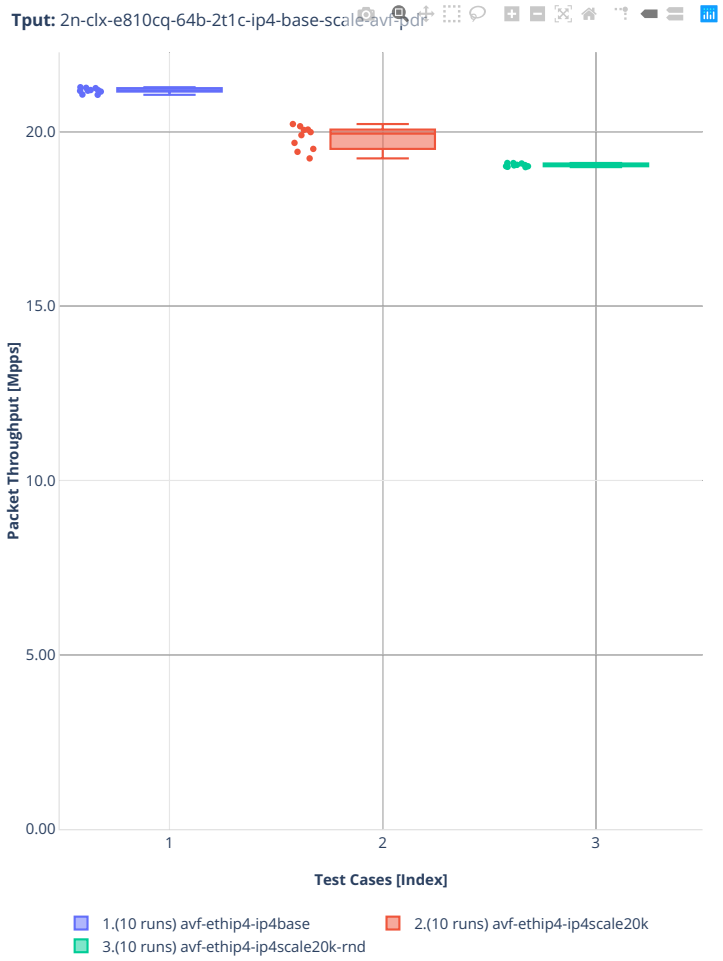
Tput: 2n-clx-e810cq-64b-2t1c-ip4-base-scale-avf-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-ethip4-ip4base
2.(10 runs) avf-ethip4-ip4scale20k
3.(10 runs) avf-ethip4-ip4scale20k-rnd

64b-2t1c-ip4routing-features-avf

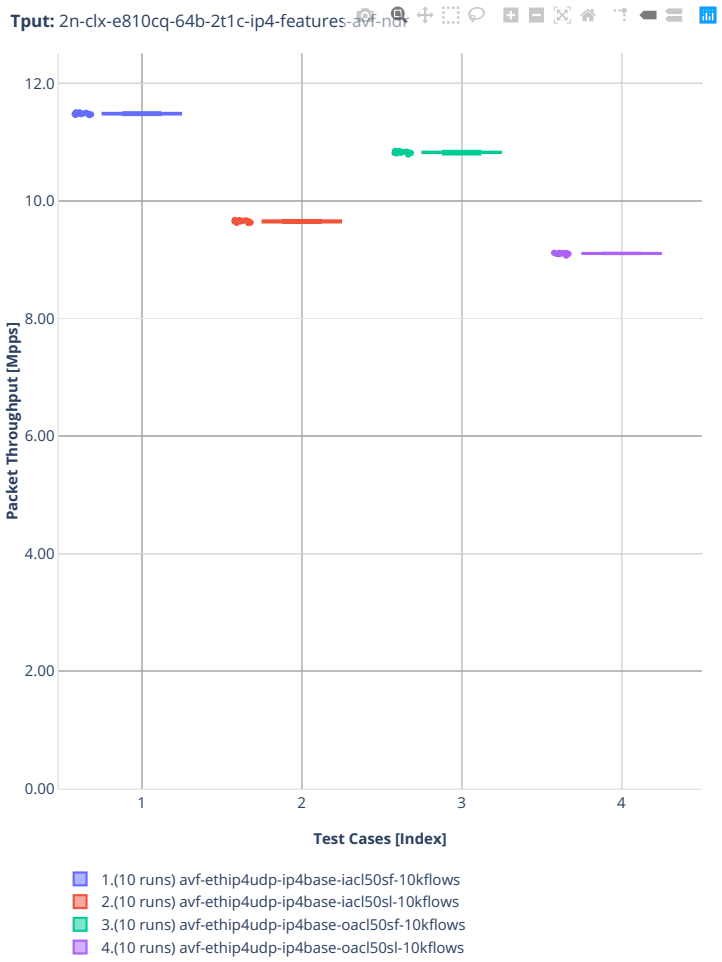

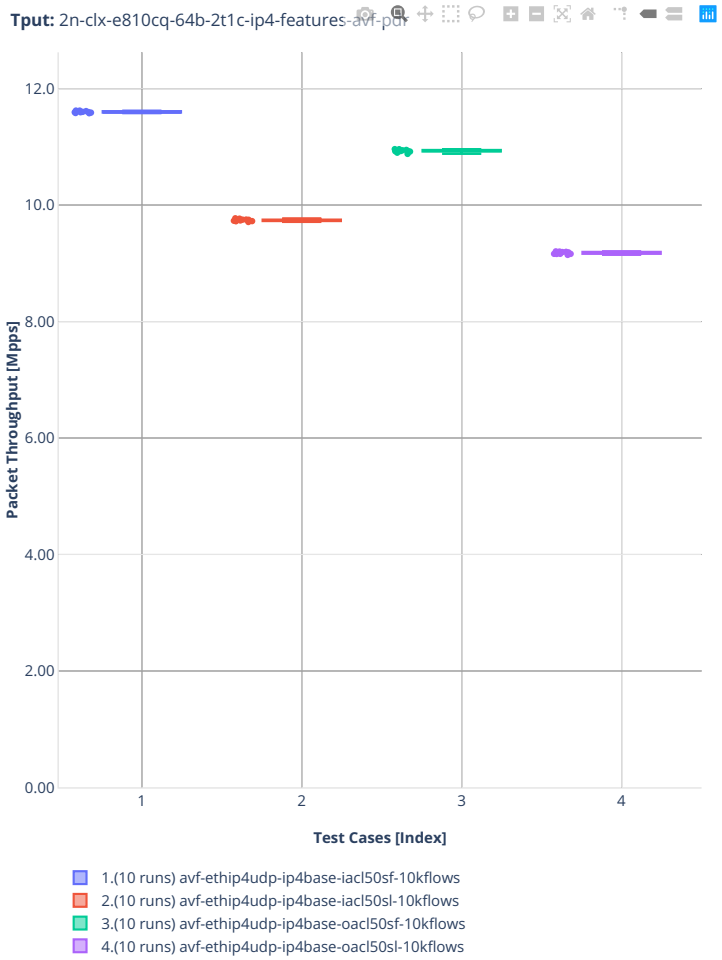
Tput: 2n-clx-e810cq-64b-2t1c-ip4-features-avf-pdr
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) avf-ethip4udp-ip4base-iacl50sf-10kflows
2.(10 runs) avf-ethip4udp-ip4base-iacl50sl-10kflows
3.(10 runs) avf-ethip4udp-ip4base-oacl50sf-10kflows
4.(10 runs) avf-ethip4udp-ip4base-oacl50sl-10kflows

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

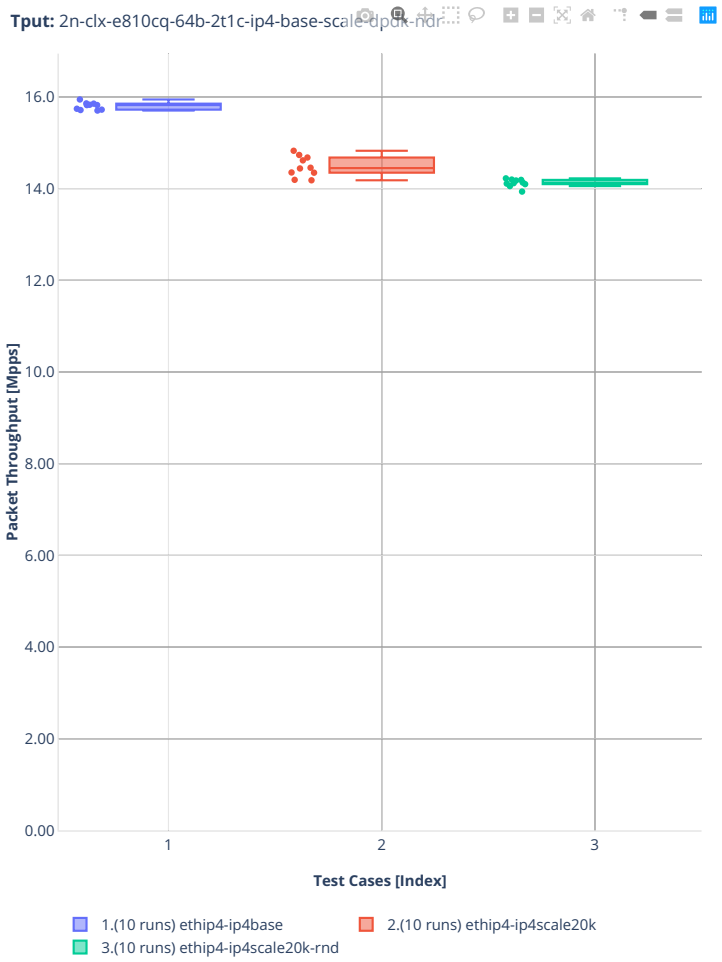

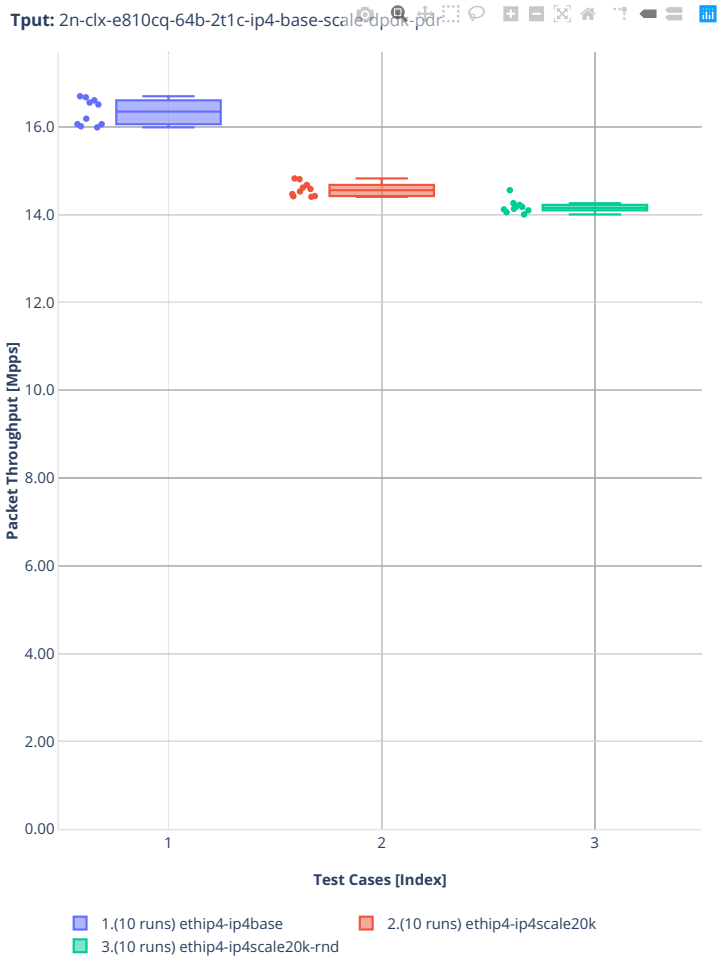
Tput: 2n-clx-e810cq-64b-2t1c-ip4-base-scale-dpdk-pdr
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip4-ip4base
2.(10 runs) ethip4-ip4scale20k
3.(10 runs) ethip4-ip4scale20k-rnd

2n-zn2-xxv710

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

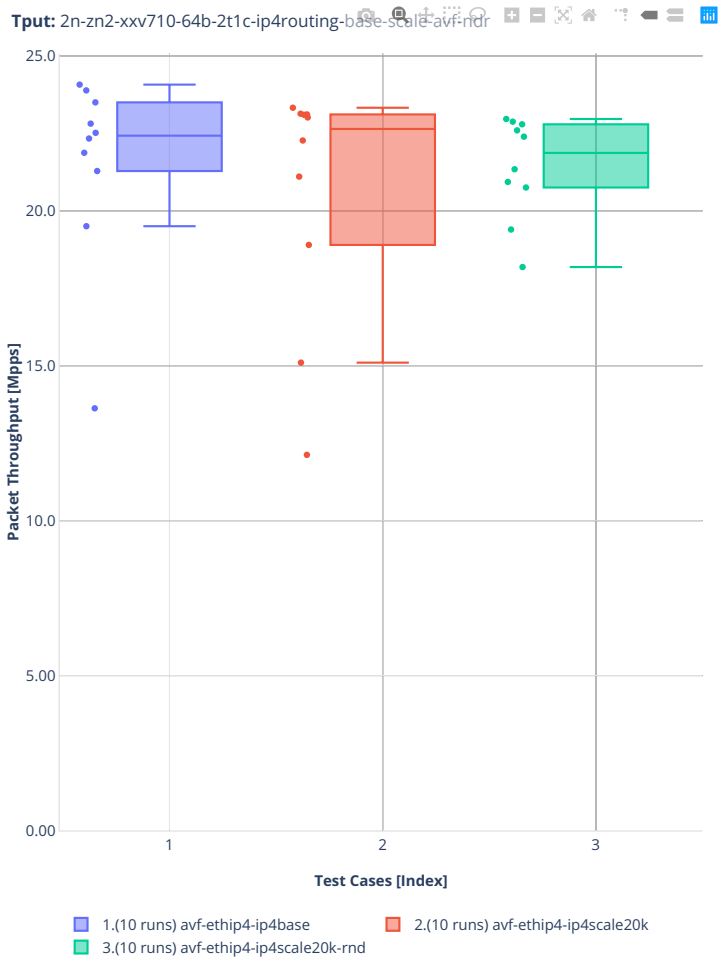

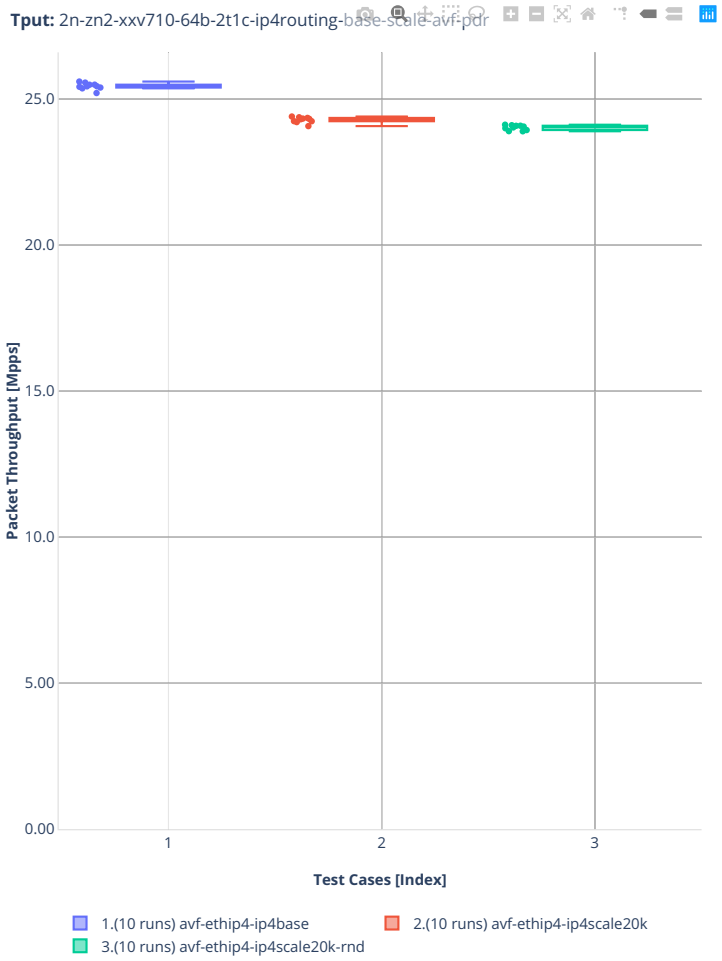
Tput: 2n-zn2-xxv710-64b-2t1c-ip4routing-base-scale-avf-pdr
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-ethip4-ip4base
2.(10 runs) avf-ethip4-ip4scale20k
3.(10 runs) avf-ethip4-ip4scale20k-rnd

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

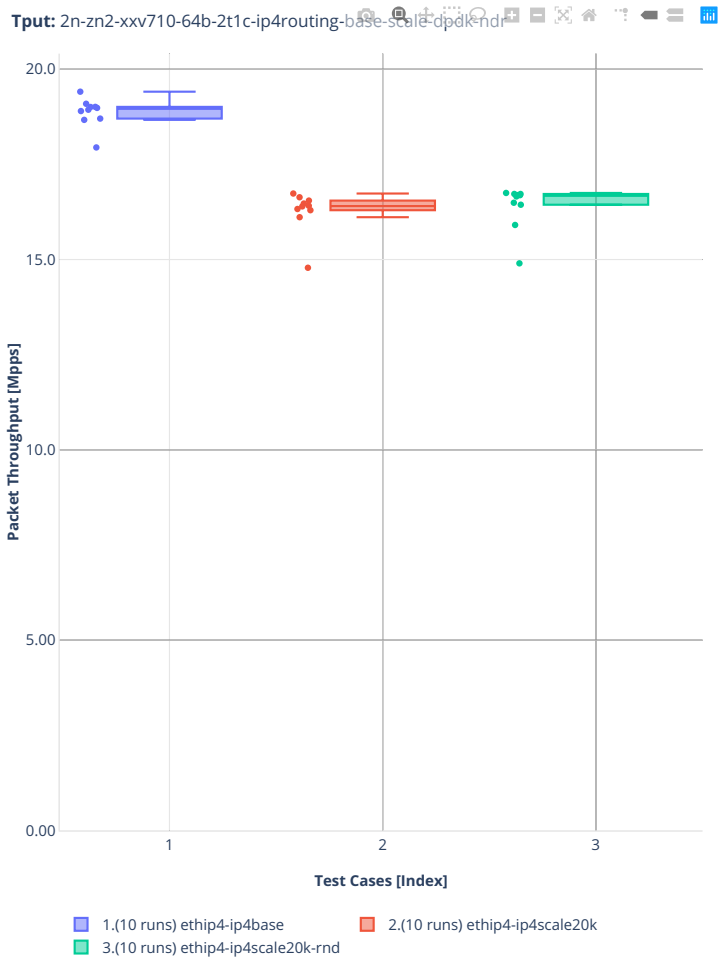

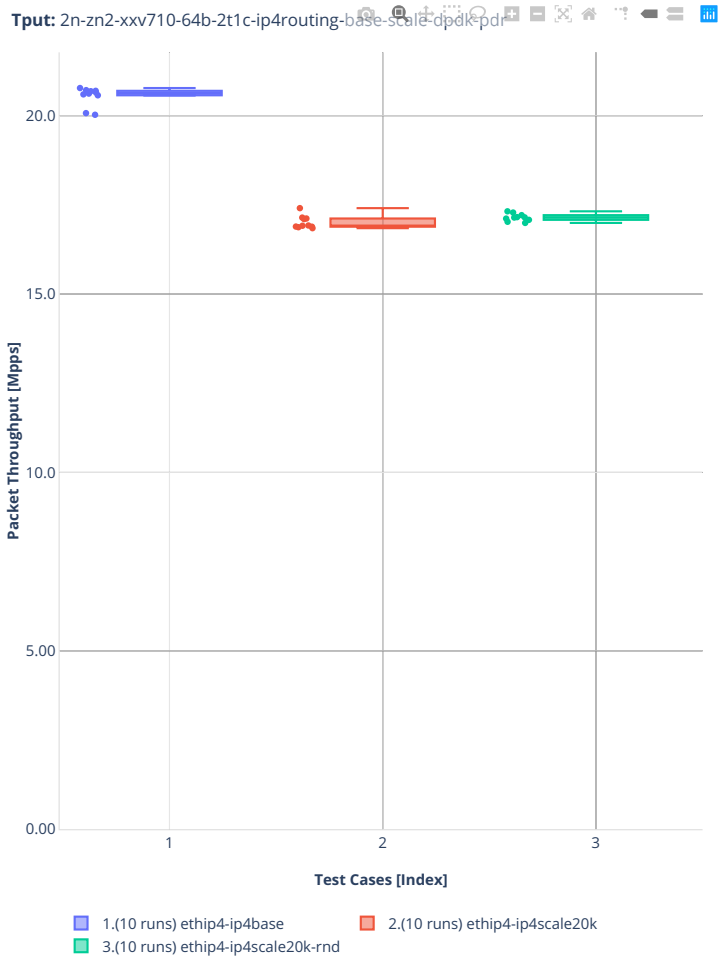
Tput: 2n-zn2-xxv710-64b-2t1c-ip4routing-base-scale-dpdk-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip4-ip4base
2.(10 runs) ethip4-ip4scale20k
3.(10 runs) ethip4-ip4scale20k-rnd

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

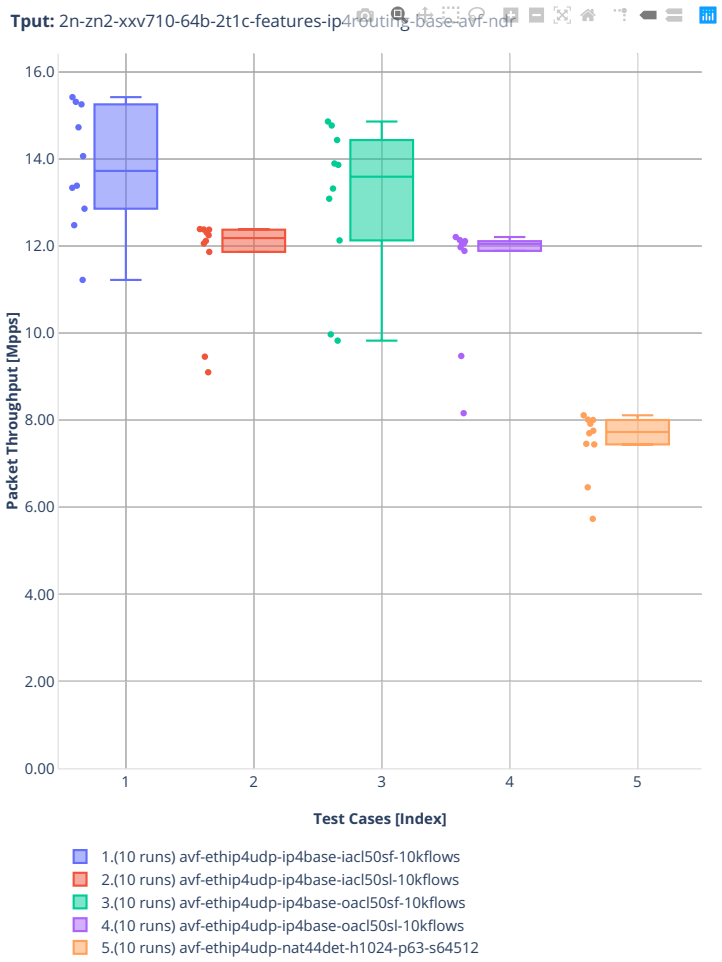

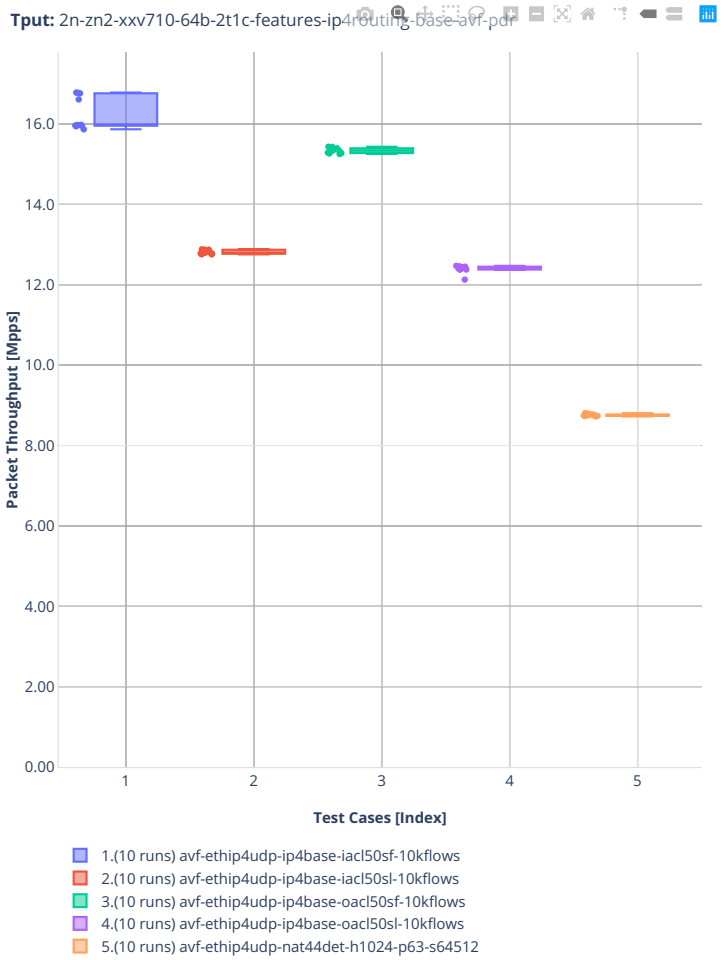
Tput: 2n-zn2-xxv710-64b-2t1c-features-ip4routing-base-avf-pdr
Packet Throughput [Mpps]
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
5
Test Cases [Index]
1.(10 runs) avf-ethip4udp-ip4base-iacl50sf-10kflows
2.(10 runs) avf-ethip4udp-ip4base-iacl50sl-10kflows
3.(10 runs) avf-ethip4udp-ip4base-oacl50sf-10kflows
4.(10 runs) avf-ethip4udp-ip4base-oacl50sl-10kflows
5.(10 runs) avf-ethip4udp-nat44det-h1024-p63-s64512

2n-zn2-x710

64b-2t1c-ip4routing-base-[avf,dpdk]

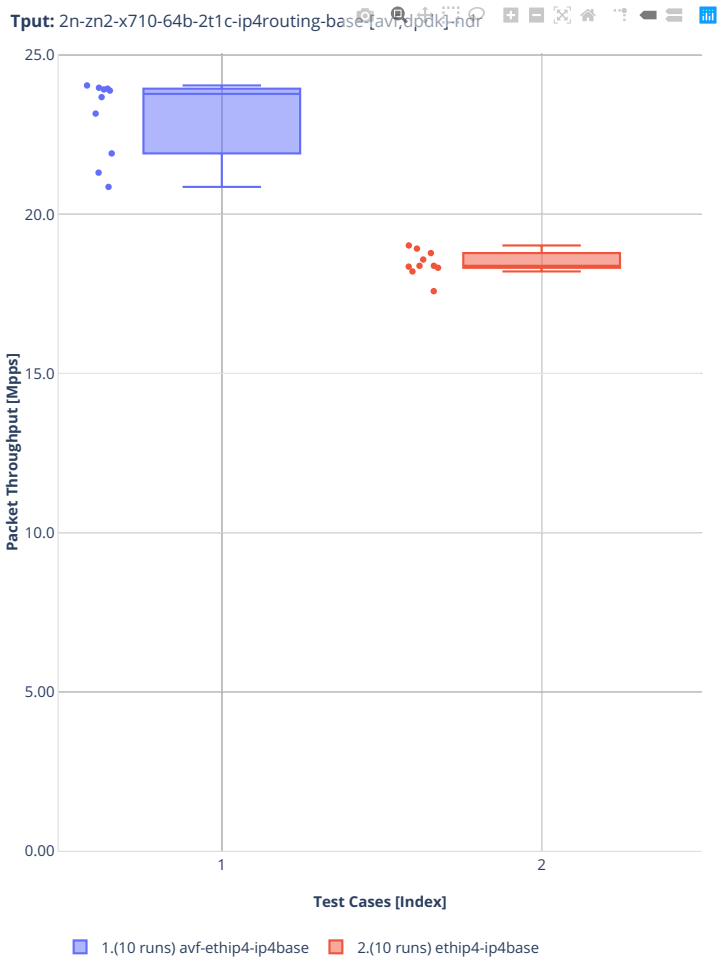

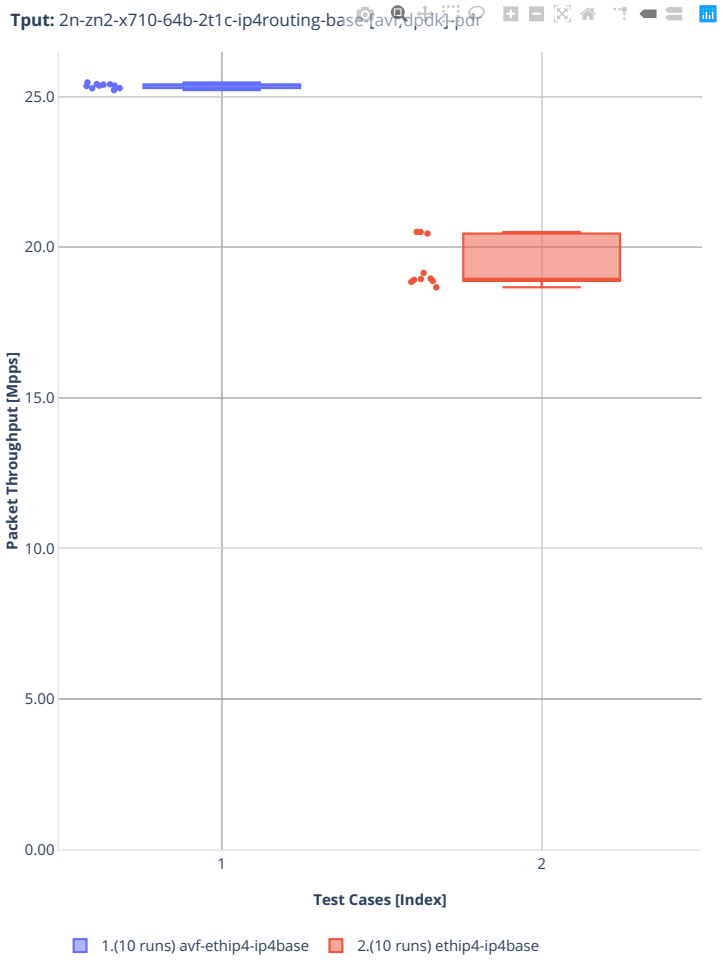
Tput: 2n-zn2-x710-64b-2t1c-ip4routing-base-[avf,dpdk]-pdr
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) avf-ethip4-ip4base
2.(10 runs) ethip4-ip4base

2n-zn2-cx556a

64b-2t1c-ip4routing-base-scale-rdma-core

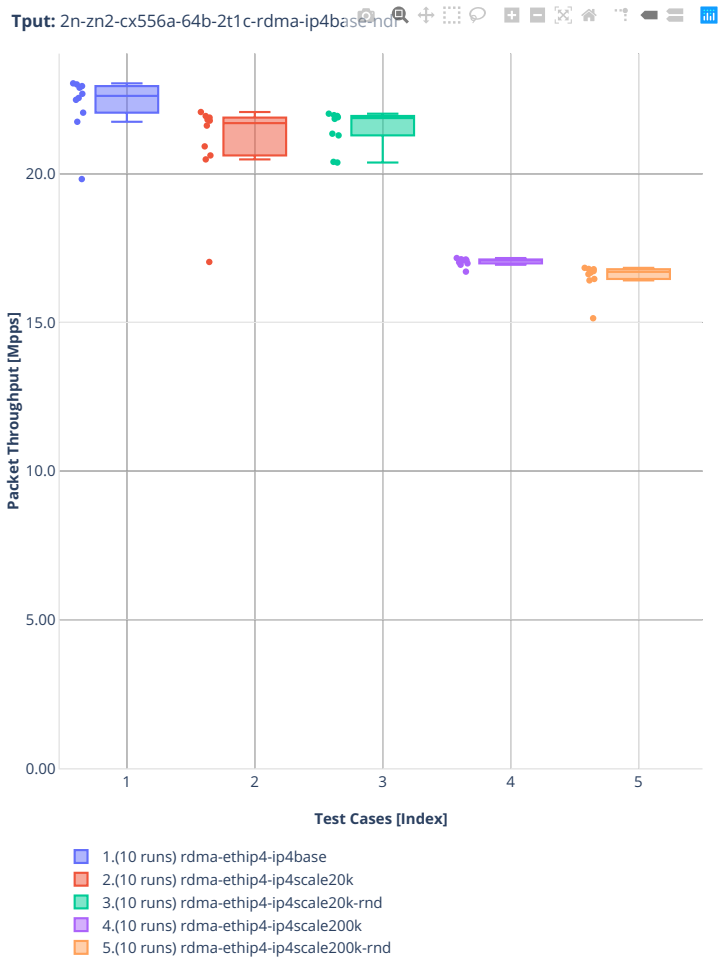

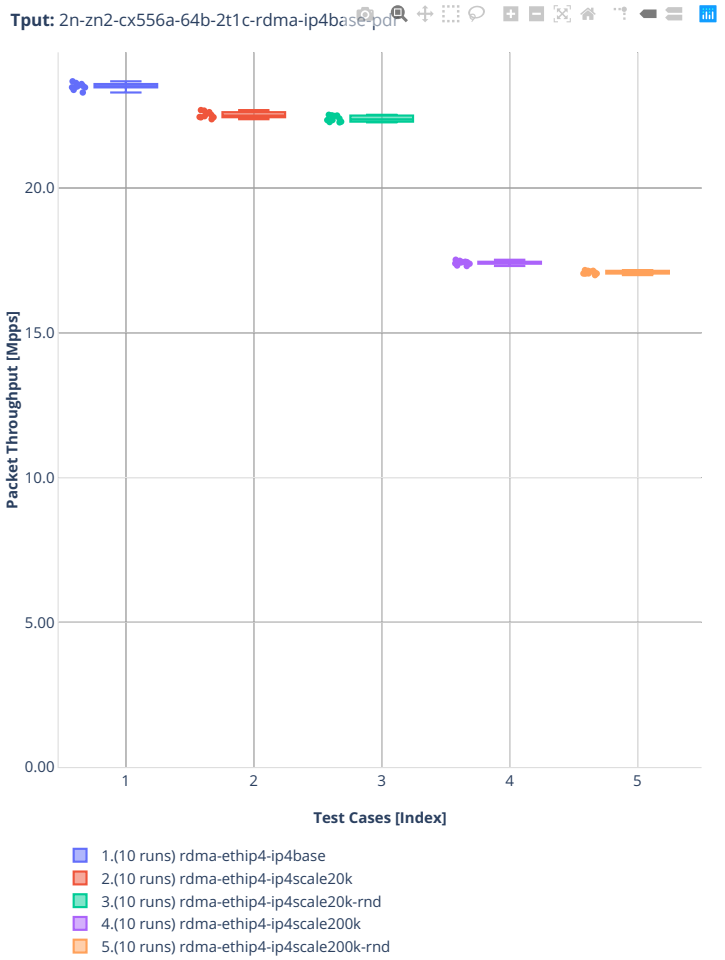
Tput: 2n-zn2-cx556a-64b-2t1c-rdma-ip4base-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
Test Cases [Index]
1.(10 runs) rdma-ethip4-ip4base
2.(10 runs) rdma-ethip4-ip4scale20k
3.(10 runs) rdma-ethip4-ip4scale20k-rnd
4.(10 runs) rdma-ethip4-ip4scale200k
5.(10 runs) rdma-ethip4-ip4scale200k-rnd

64b-2t1c-ip4routing-features

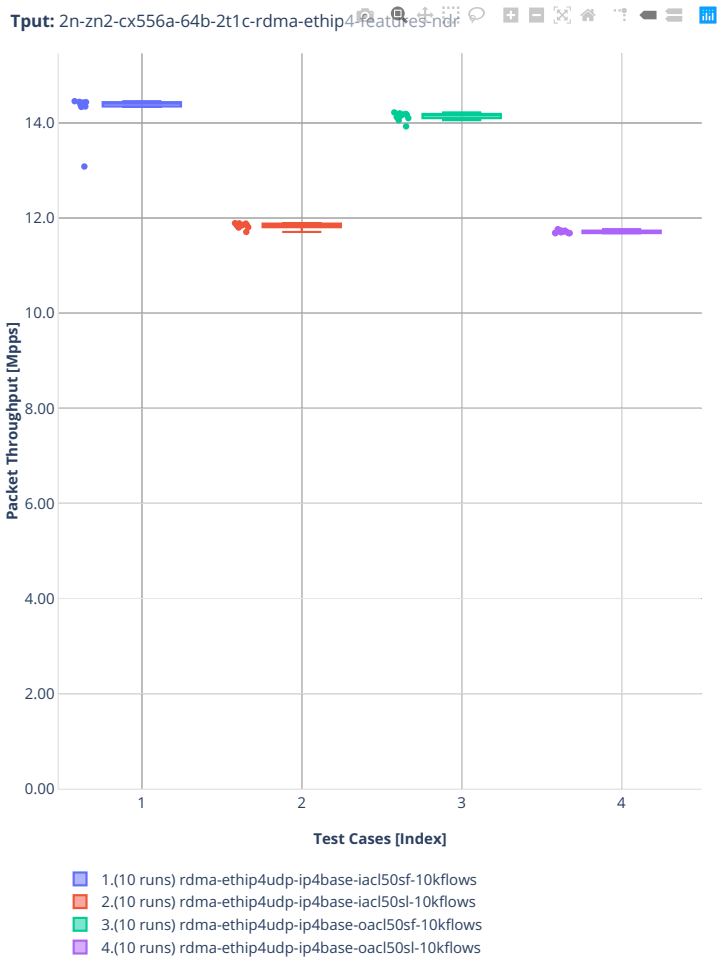

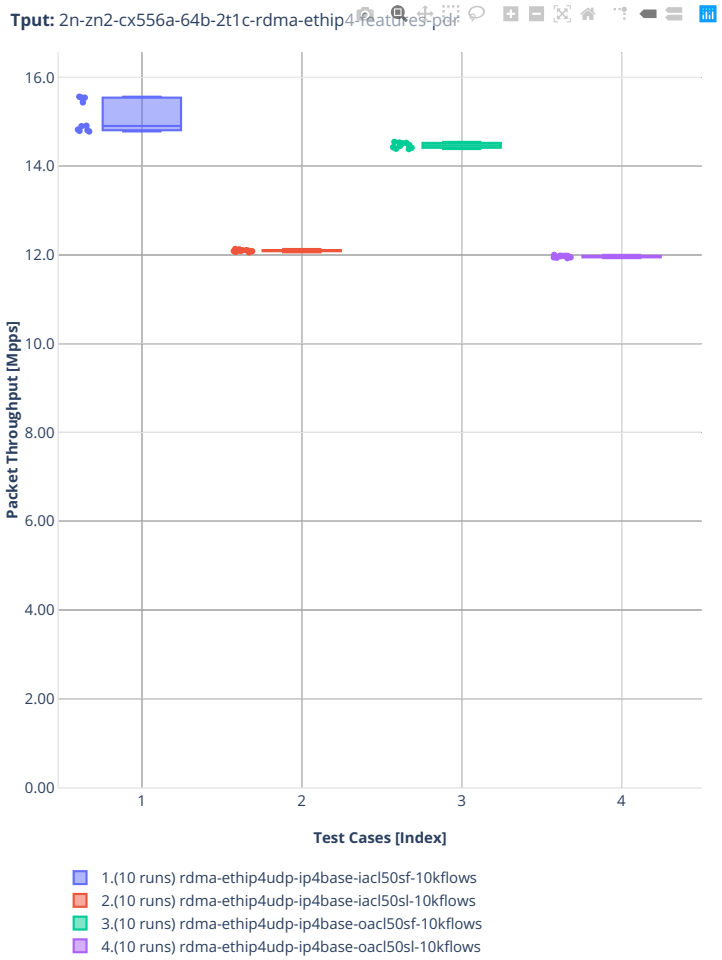
Tput: 2n-zn2-cx556a-64b-2t1c-rdma-ethip4-features-pdr
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) rdma-ethip4udp-ip4base-iacl50sf-10kflows
2.(10 runs) rdma-ethip4udp-ip4base-iacl50sl-10kflows
3.(10 runs) rdma-ethip4udp-ip4base-oacl50sf-10kflows
4.(10 runs) rdma-ethip4udp-ip4base-oacl50sl-10kflows

3n-alt-xl710

64b-1t1c-ip4routing-base-scale

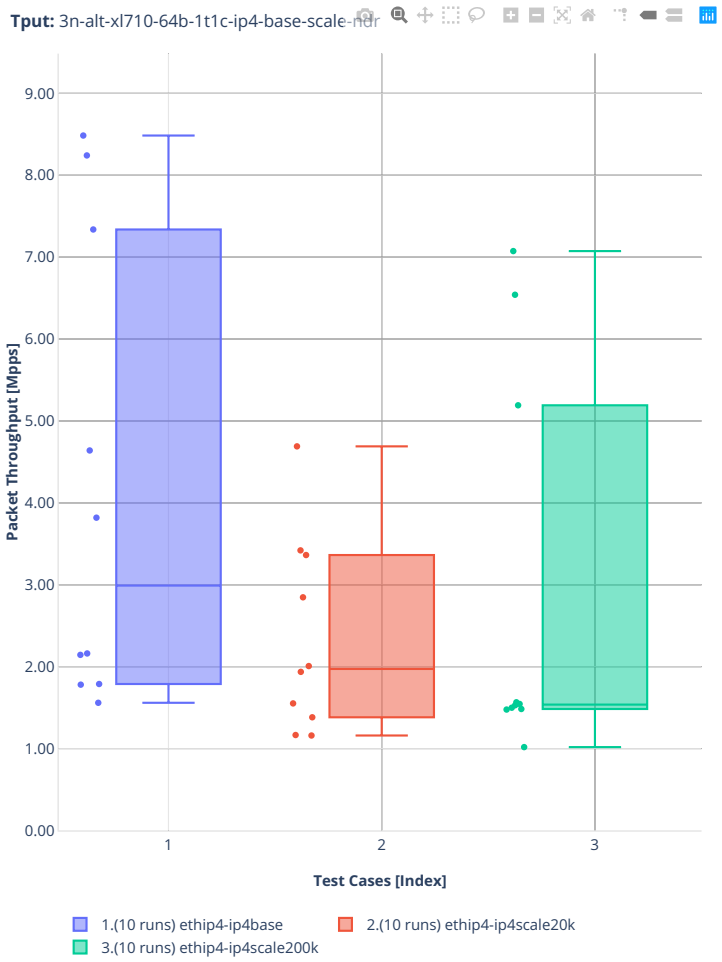


Tput: 3n-alt-xl710-64b-1t1c-ip4-base-scale-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip4-ip4base
2.(10 runs) ethip4-ip4scale20k
3.(10 runs) ethip4-ip4scale200k

64b-1t1c-ip4routing-features

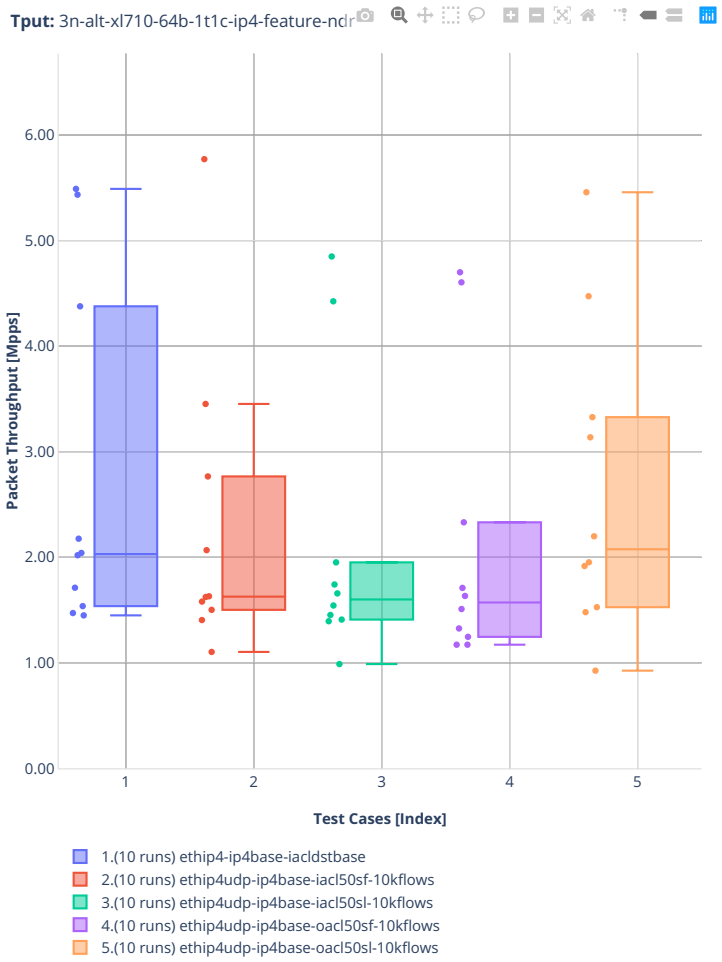

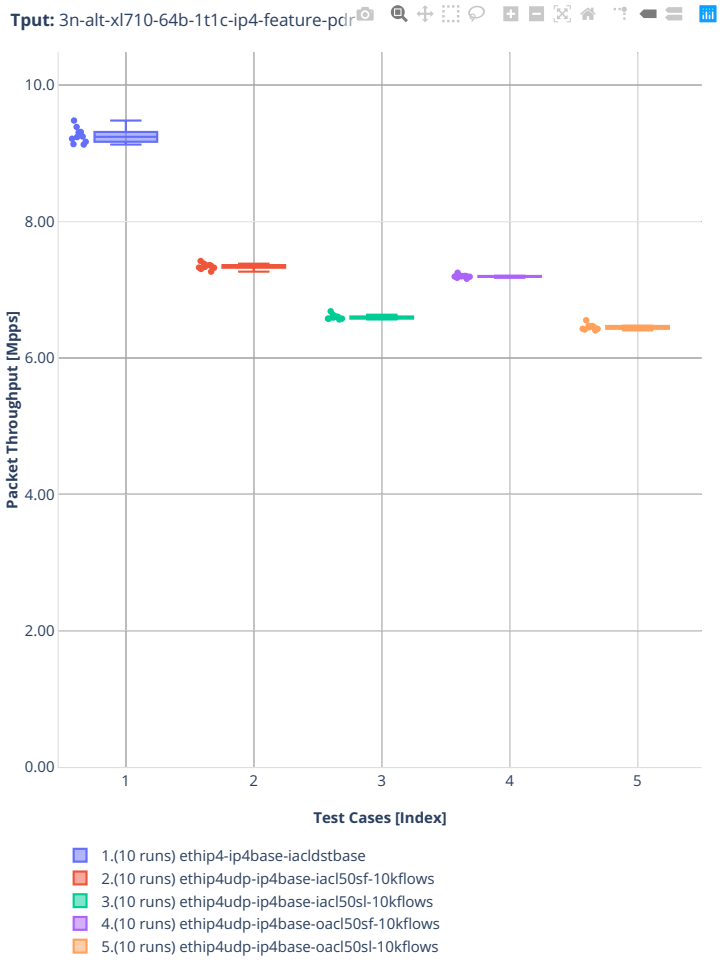
Tput: 3n-alt-xl710-64b-1t1c-ip4-feature-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
Test Cases [Index]
1.(10 runs) ethip4-ip4base-iacldstbase
2.(10 runs) ethip4udp-ip4base-iacl50sf-10kflows
3.(10 runs) ethip4udp-ip4base-iacl50sl-10kflows
4.(10 runs) ethip4udp-ip4base-oacl50sf-10kflows
5.(10 runs) ethip4udp-ip4base-oacl50sl-10kflows

3n-tsh-x520

64b-1t1c-features-ip4routing-base-ixgbe

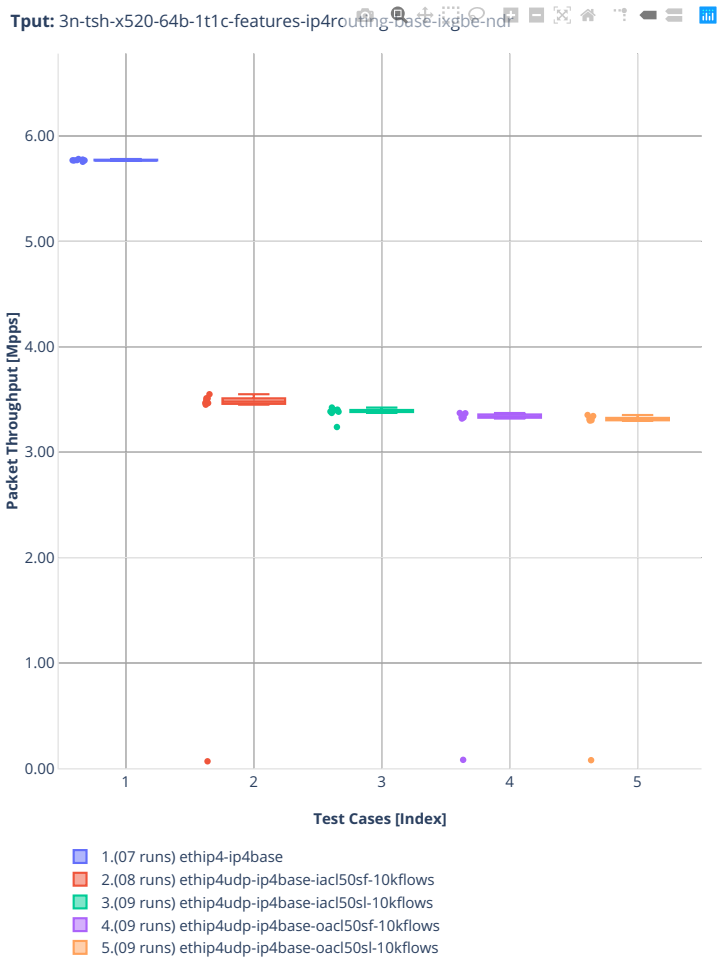

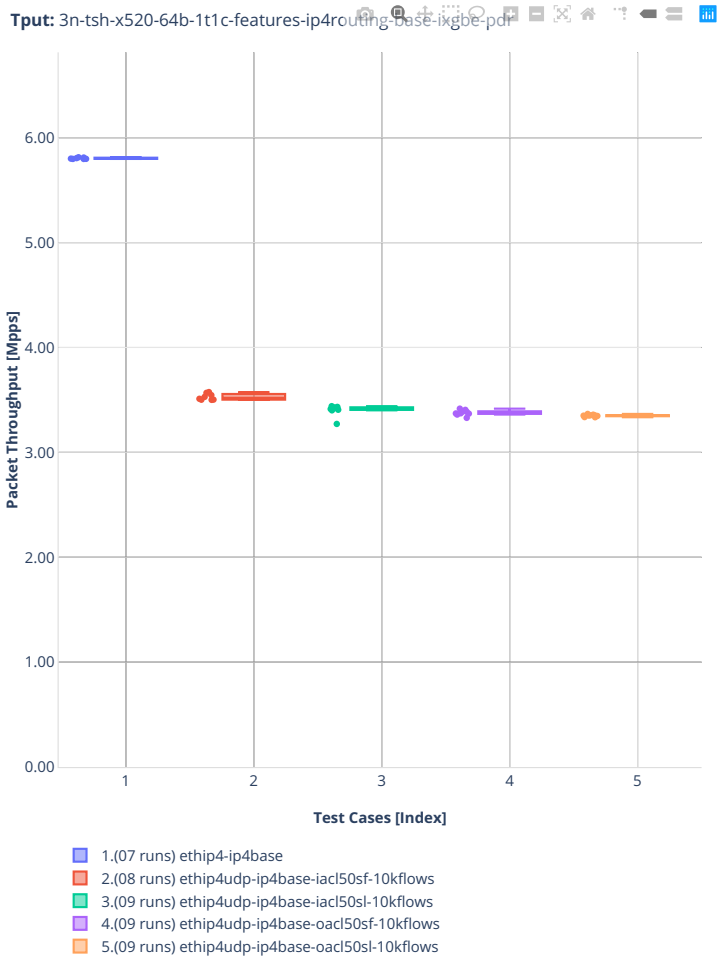
Tput: 3n-tsh-x520-64b-1t1c-features-ip4routing-base-ixgbe-pdr
Packet Throughput [Mpps]
6.00
5.00
4.00
3.00
2.00
1.00
0.00
1
2
3
4
5
Test Cases [Index]
1.(07 runs) ethip4-ip4base
2.(08 runs) ethip4udp-ip4base-iacl50sf-10kflows
3.(09 runs) ethip4udp-ip4base-iacl50sl-10kflows
4.(09 runs) ethip4udp-ip4base-oacl50sf-10kflows
5.(09 runs) ethip4udp-ip4base-oacl50sl-10kflows

2n-tx2-xl710

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

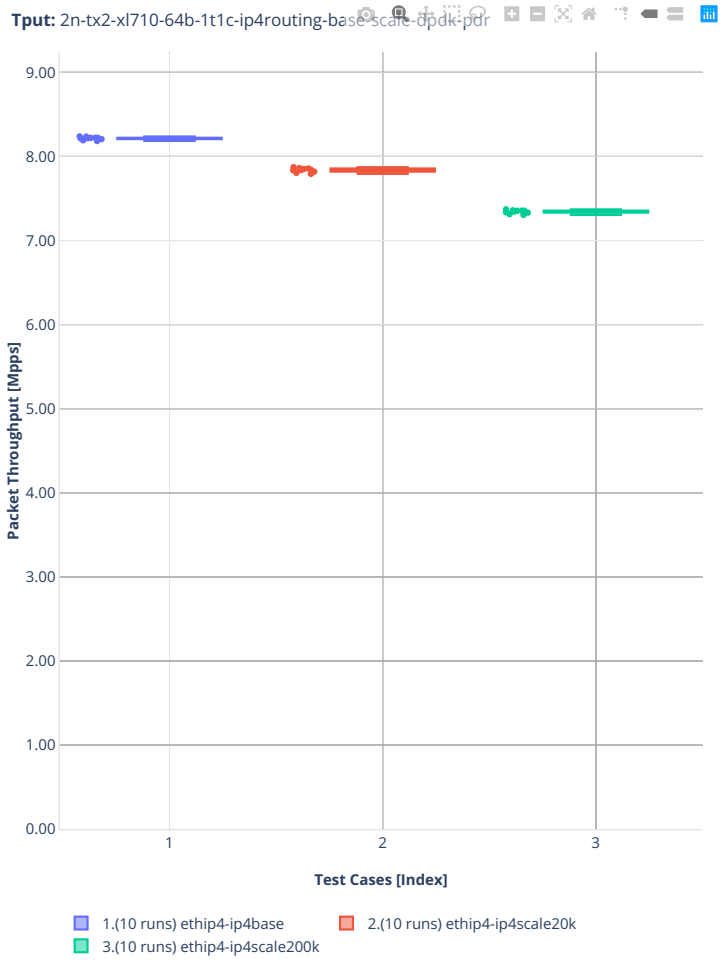
Tput: 2n-tx2-xl710-64b-1t1c-ip4routing-base-scale-dpdk-pdr
9.00
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip4-ip4base
2.(10 runs) ethip4-ip4scale20k
3.(10 runs) ethip4-ip4scale200k

64b-1t1c-ip4routing-base-scale-af-xdp

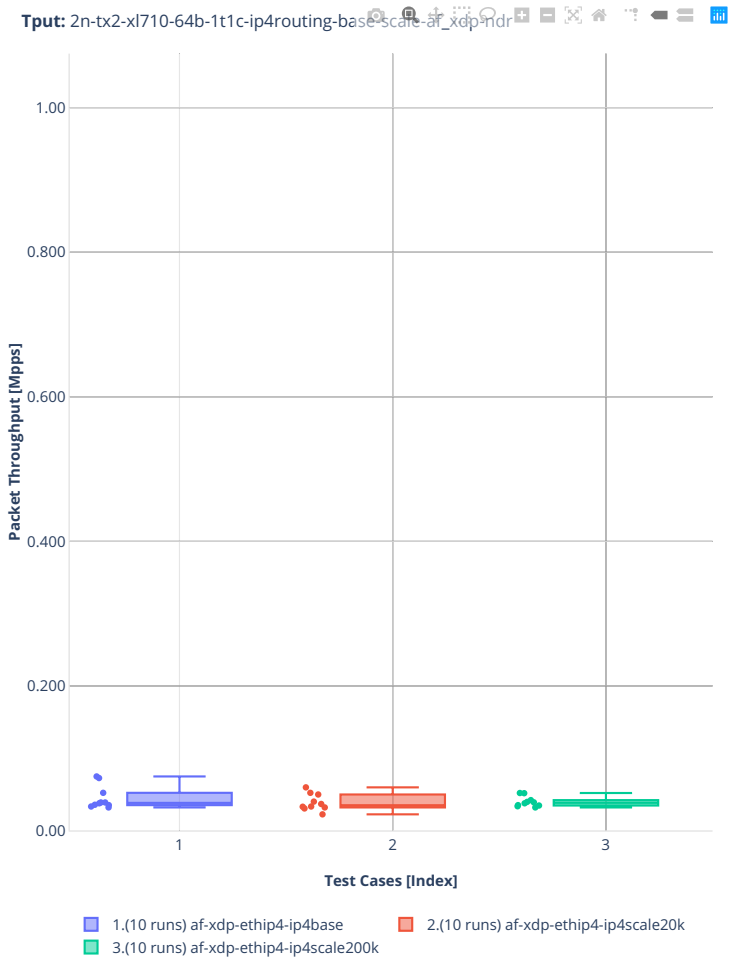

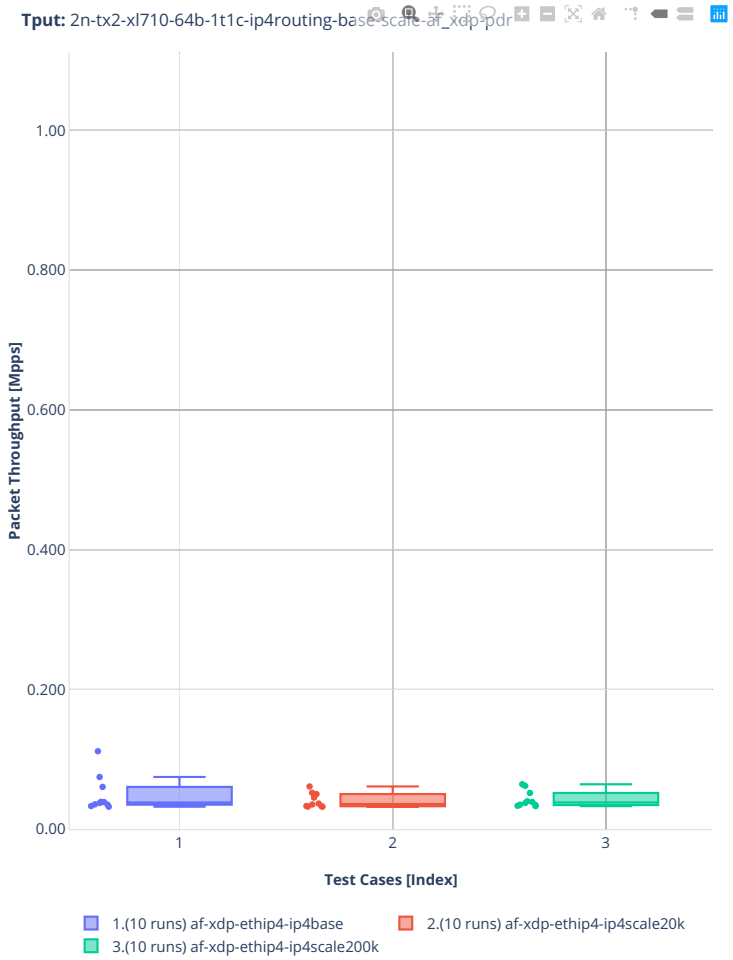
Tput: 2n-tx2-xl710-64b-1t1c-ip4routing-base-scale-af_xdp-pdr
1.00
0.800
0.600
0.400
0.200
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) af-xdp-ethip4-ip4base
2.(10 runs) af-xdp-ethip4-ip4scale20k
3.(10 runs) af-xdp-ethip4-ip4scale200k

64b-1t1c-features-ip4routing-base-dpdk

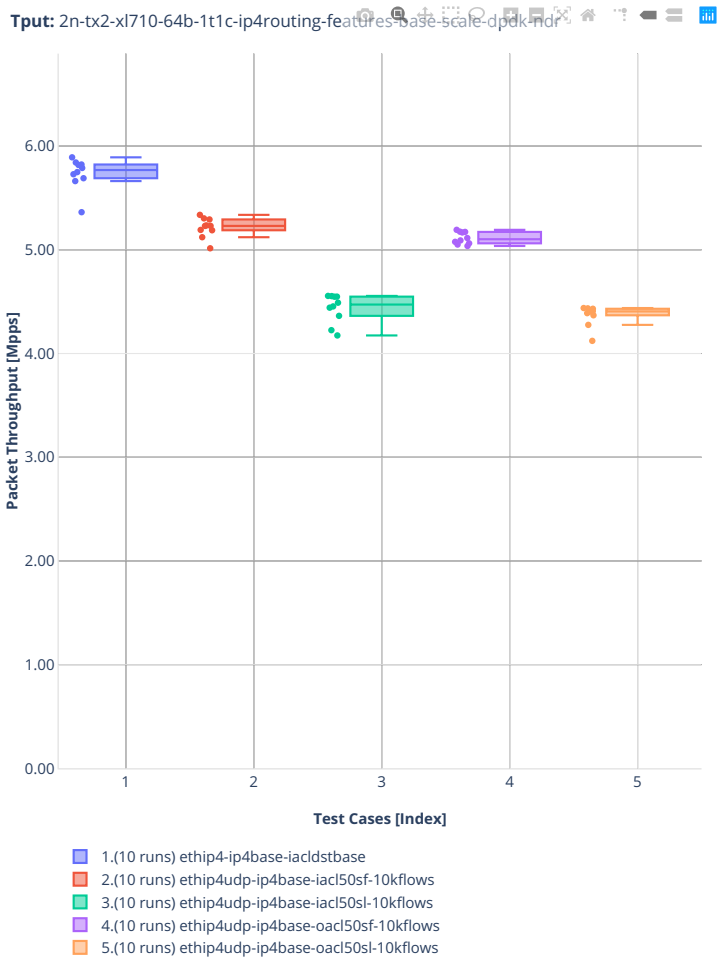

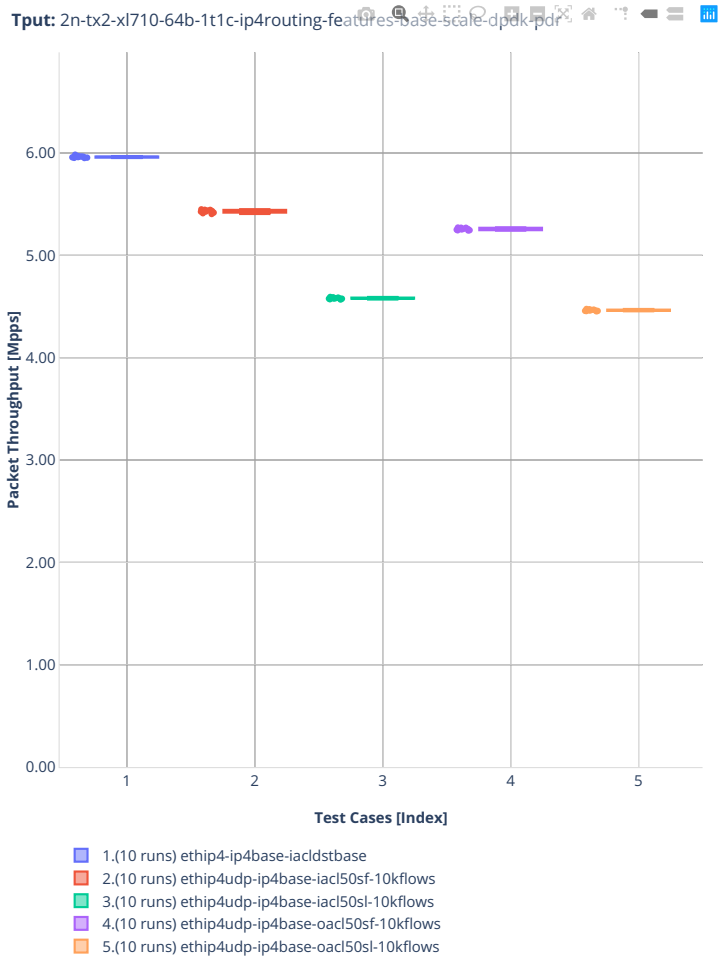
Tput: 2n-tx2-xl710-64b-1t1c-ip4routing-features-base-scale-dpdk-pdr
Packet Throughput [Mpps]
6.00
5.00
4.00
3.00
2.00
1.00
0.00
1
2
3
4
5
Test Cases [Index]
1.(10 runs) ethip4-ip4base-iacldstbase
2.(10 runs) ethip4udp-ip4base-iacl50sf-10kflows
3.(10 runs) ethip4udp-ip4base-iacl50sl-10kflows
4.(10 runs) ethip4udp-ip4base-oacl50sf-10kflows
5.(10 runs) ethip4udp-ip4base-oacl50sl-10kflows

2n-dnv-x553

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

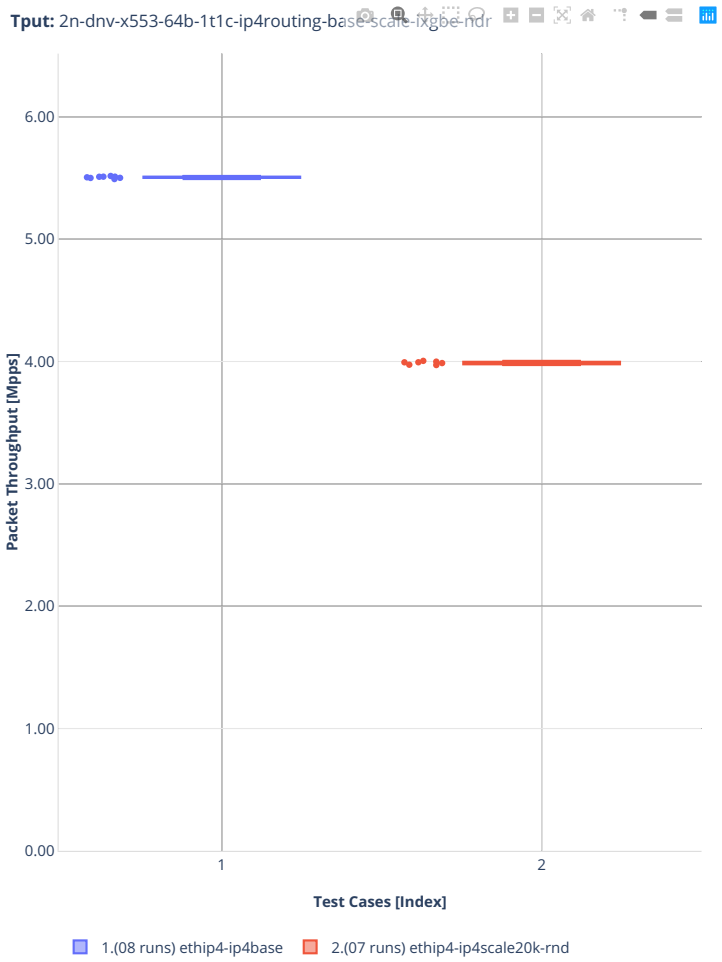

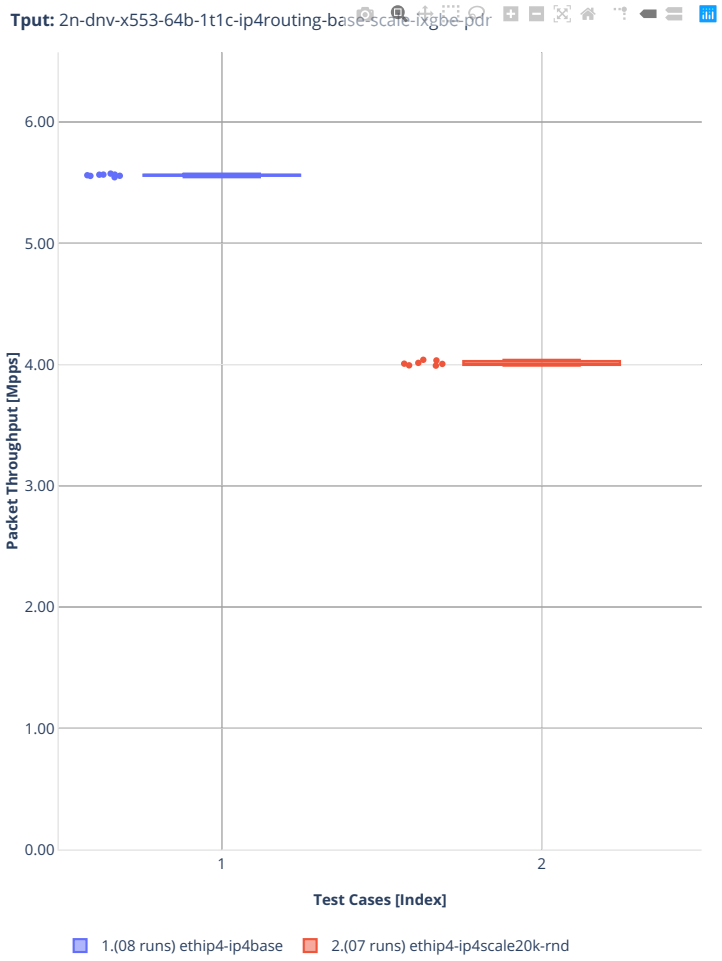
Tput: 2n-dnv-x553-64b-1t1c-ip4routing-base-scale-ixgbe-pdr
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(08 runs) ethip4-ip4base
2.(07 runs) ethip4-ip4scale20k-rnd

3n-dnv-x553

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

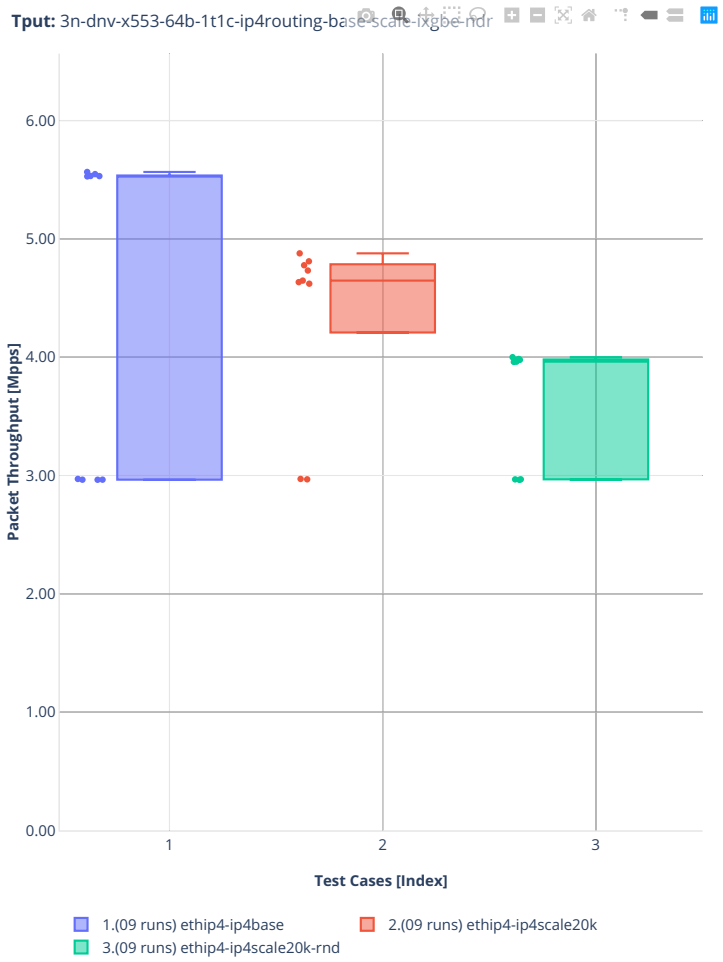

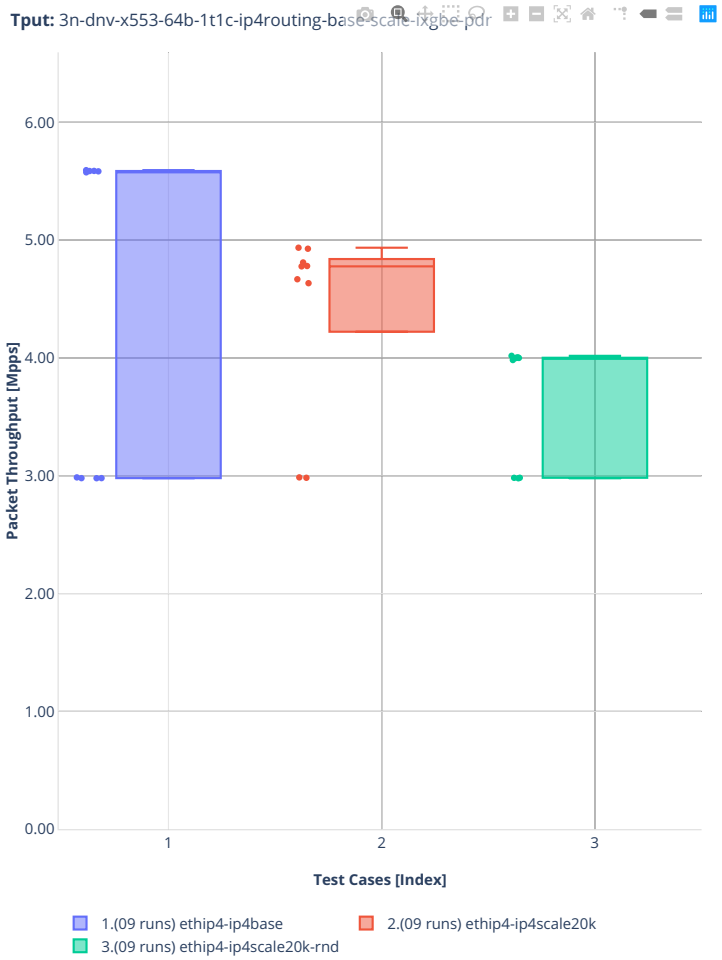
Tput: 3n-dnv-x553-64b-1t1c-ip4routing-base-scale-ixgbe-pdr
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(09 runs) ethip4-ip4base
2.(09 runs) ethip4-ip4scale20k
3.(09 runs) ethip4-ip4scale20k-rnd

2n-aws-nitro50g

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

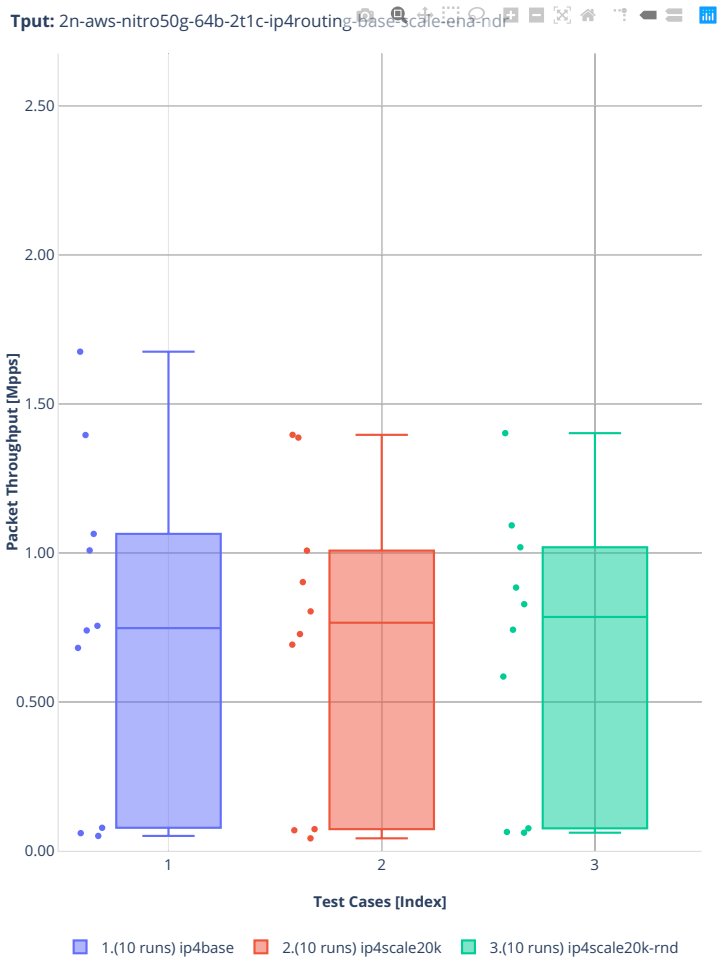

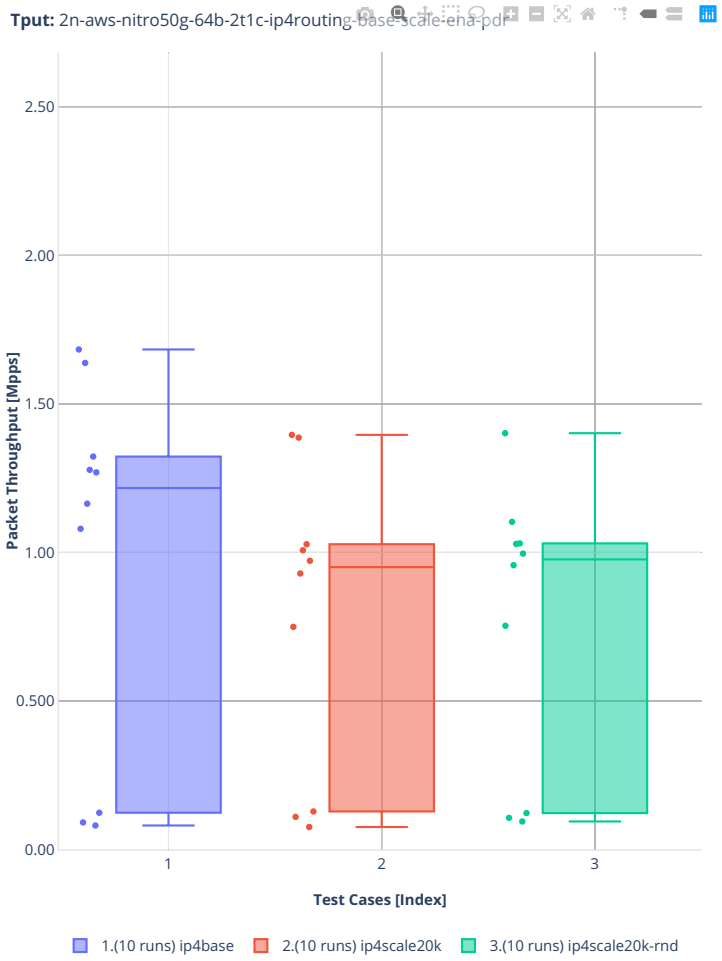
Tput: 2n-aws-nitro50g-64b-2t1c-ip4routing-base-scale-ena-pdr
2.50
2.00
1.50
1.00
0.500
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ip4base
2.(10 runs) ip4scale20k
3.(10 runs) ip4scale20k-rnd

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

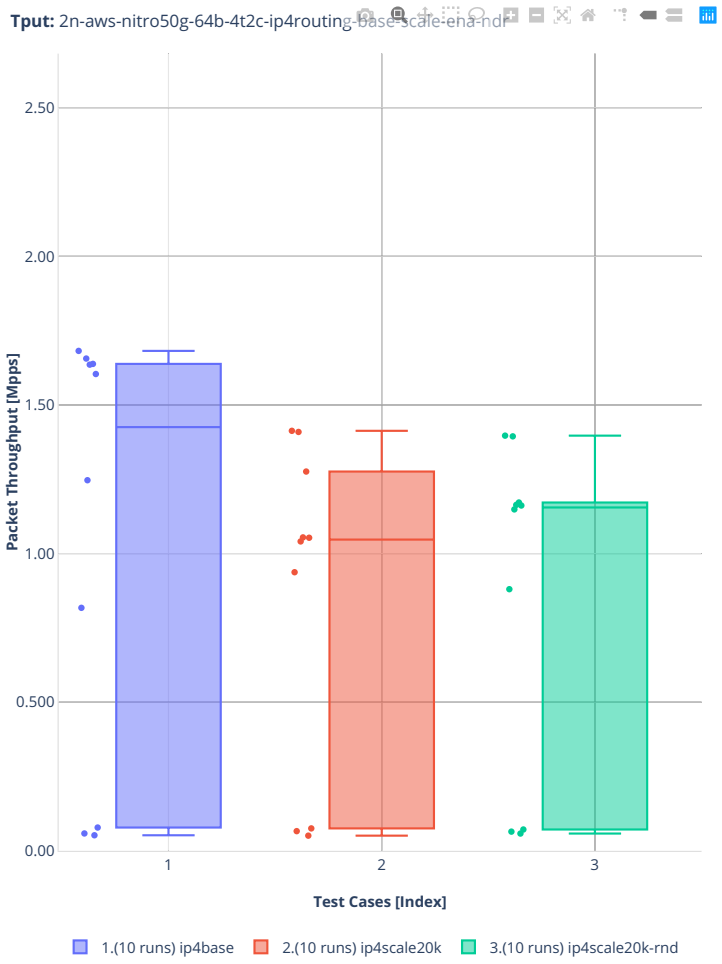

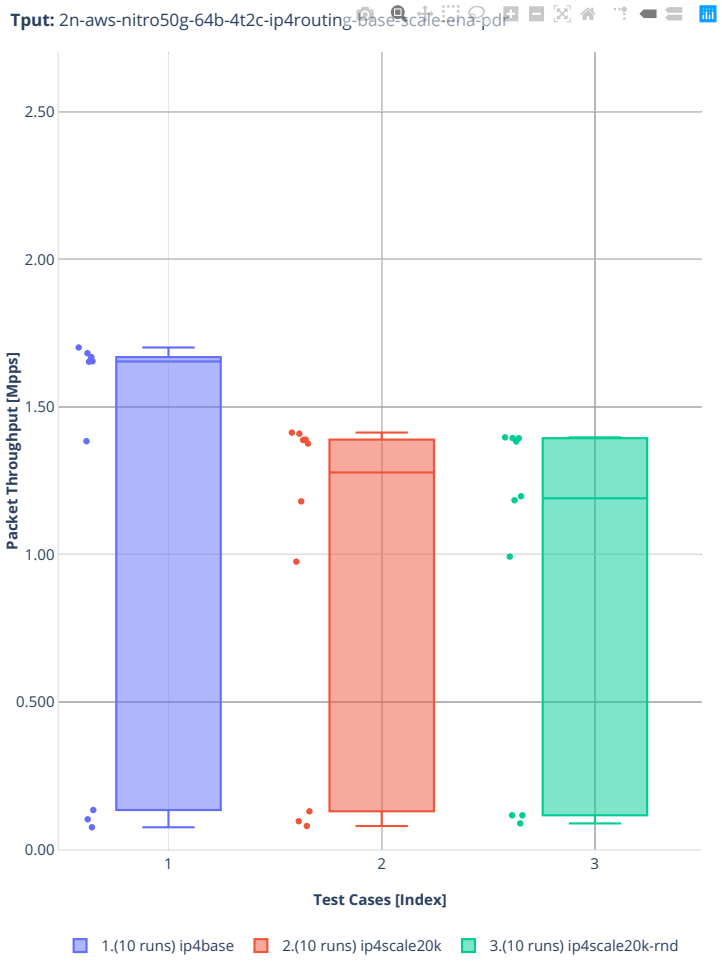
Tput: 2n-aws-nitro50g-64b-4t2c-ip4routing-base-scale-ena-pdr
2.50
2.00
1.50
1.00
0.500
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ip4base
2.(10 runs) ip4scale20k
3.(10 runs) ip4scale20k-rnd

1518b-2t1c-ip4routing-base-scale-ena

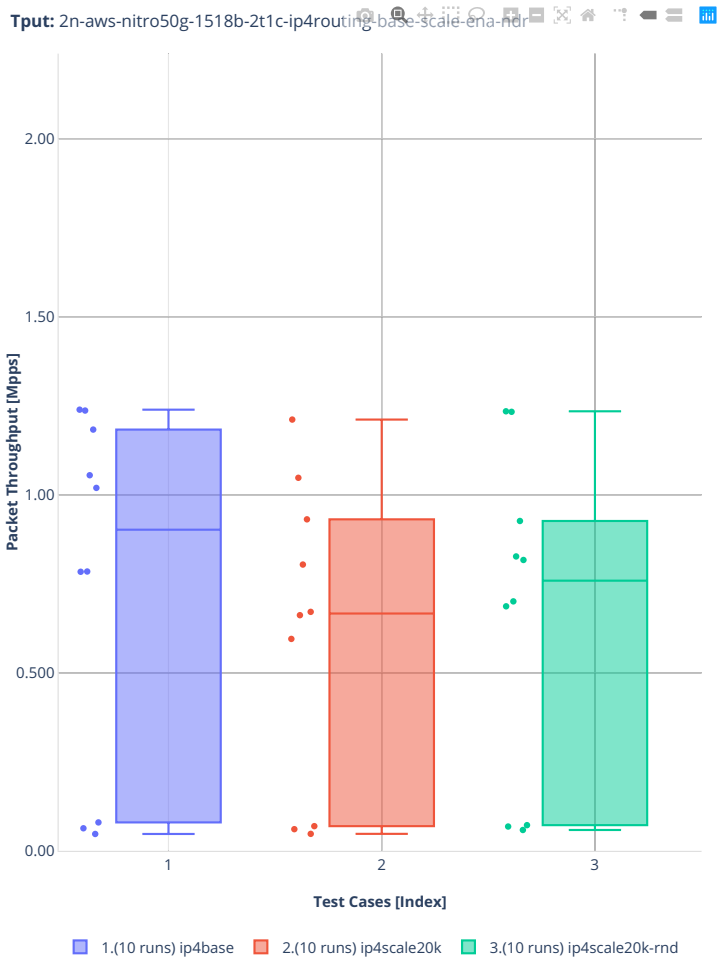

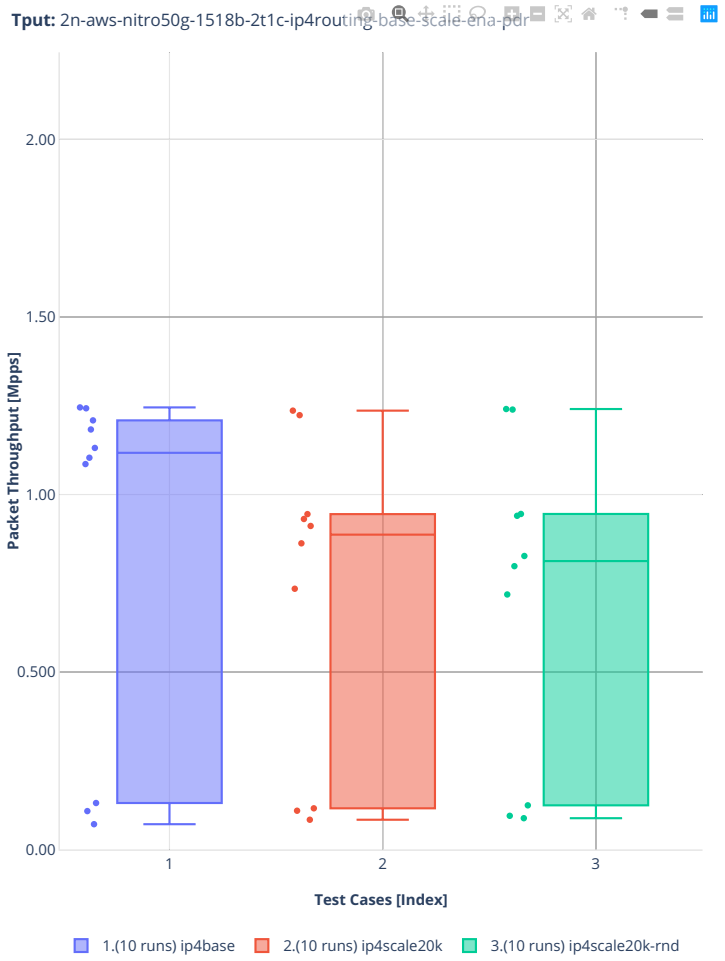
Tput: 2n-aws-nitro50g-1518b-2t1c-ip4routing-base-scale-ena-pdr
2.00
1.50
1.00
0.500
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ip4base
2.(10 runs) ip4scale20k
3.(10 runs) ip4scale20k-rnd

1518b-4t2c-ip4routing-base-scale-ena

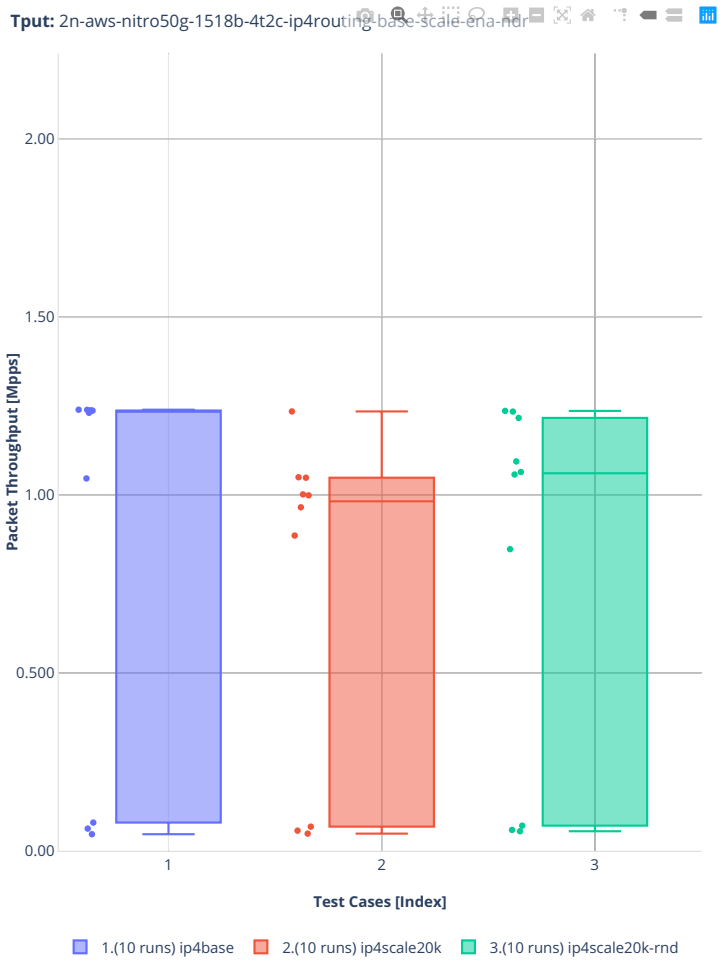

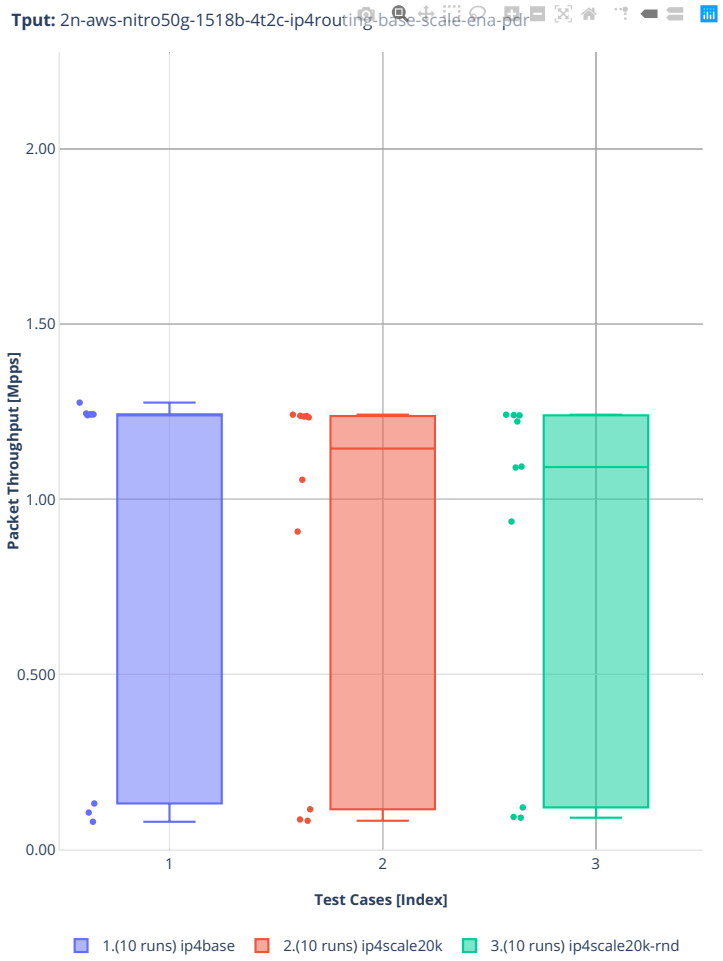
Tput: 2n-aws-nitro50g-1518b-4t2c-ip4routing-base-scale-ena-pdr
2.00
1.50
1.00
0.500
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ip4base
2.(10 runs) ip4scale20k
3.(10 runs) ip4scale20k-rnd

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[111].

[111] https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls2206

2n-icx-xxv710

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

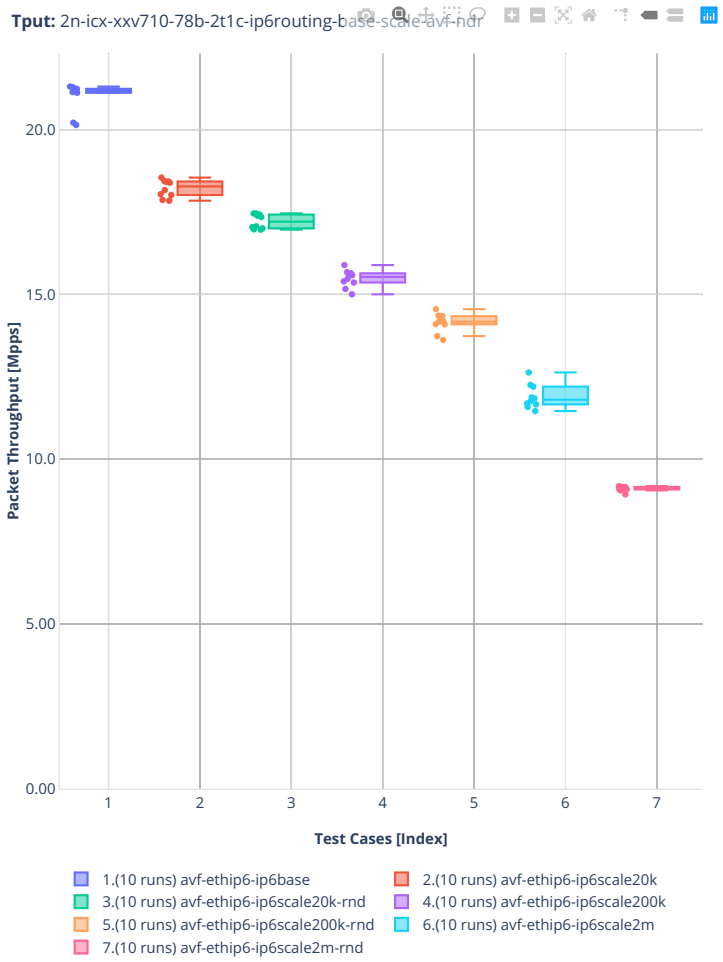

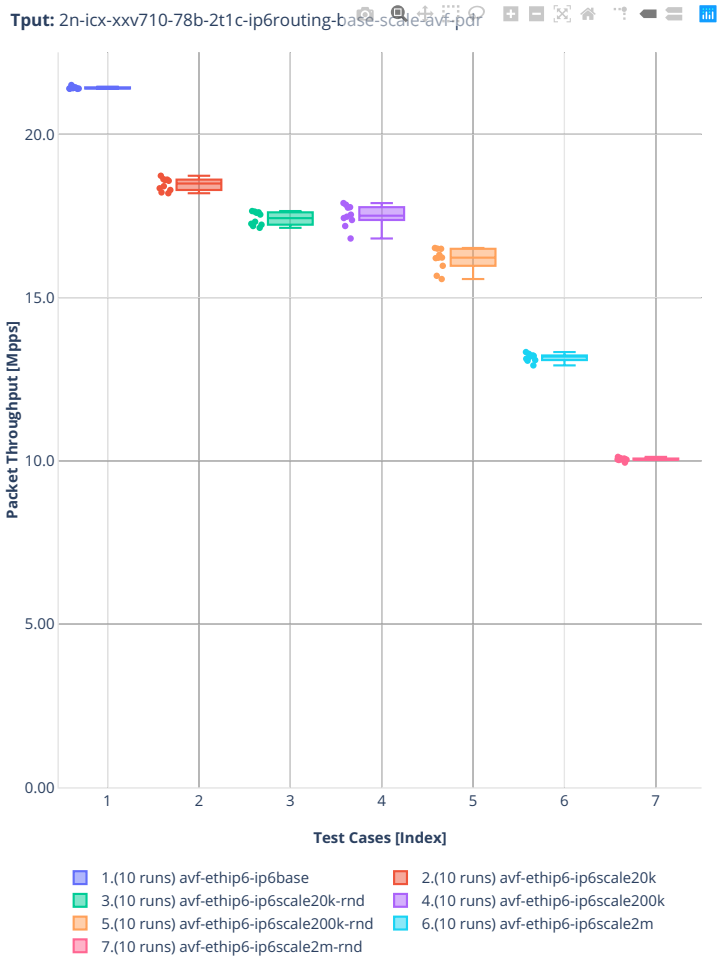
Tput: 2n-icx-xxv710-78b-2t1c-ip6routing-base-scale-avf-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
6
7
Test Cases [Index]
1.(10 runs) avf-ethip6-ip6base
2.(10 runs) avf-ethip6-ip6scale20k
3.(10 runs) avf-ethip6-ip6scale20k-rnd
4.(10 runs) avf-ethip6-ip6scale200k
5.(10 runs) avf-ethip6-ip6scale200k-rnd
6.(10 runs) avf-ethip6-ip6scale2m
7.(10 runs) avf-ethip6-ip6scale2m-rnd

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

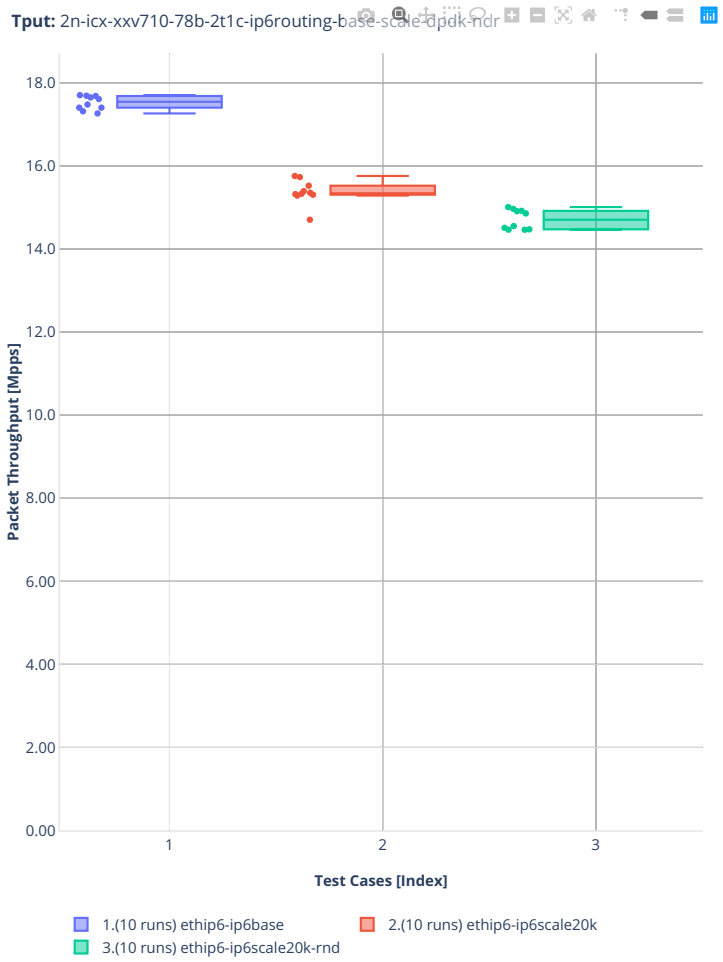

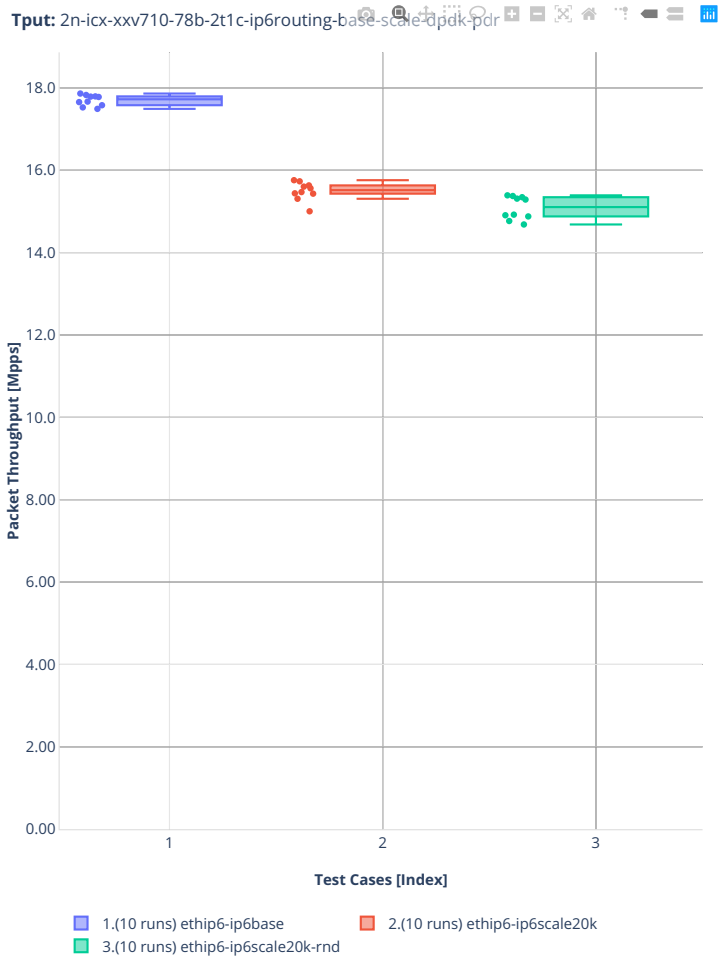
Tput: 2n-icx-xxv710-78b-2t1c-ip6routing-base-scale-dpdk-pdr
18.0
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip6-ip6base
2.(10 runs) ethip6-ip6scale20k
3.(10 runs) ethip6-ip6scale20k-rnd

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

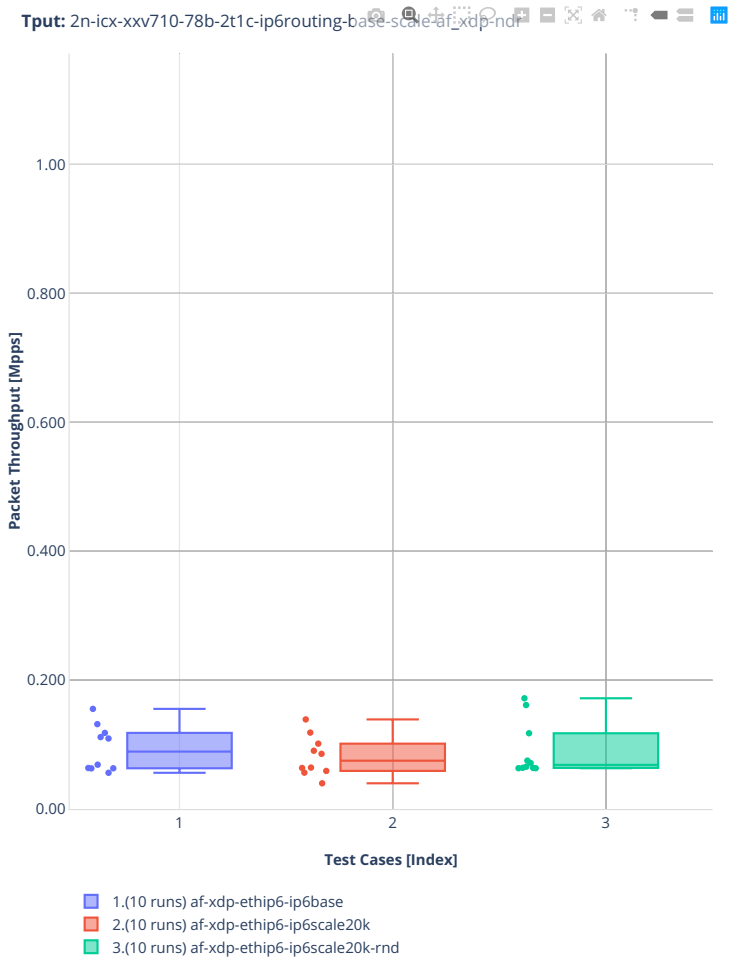

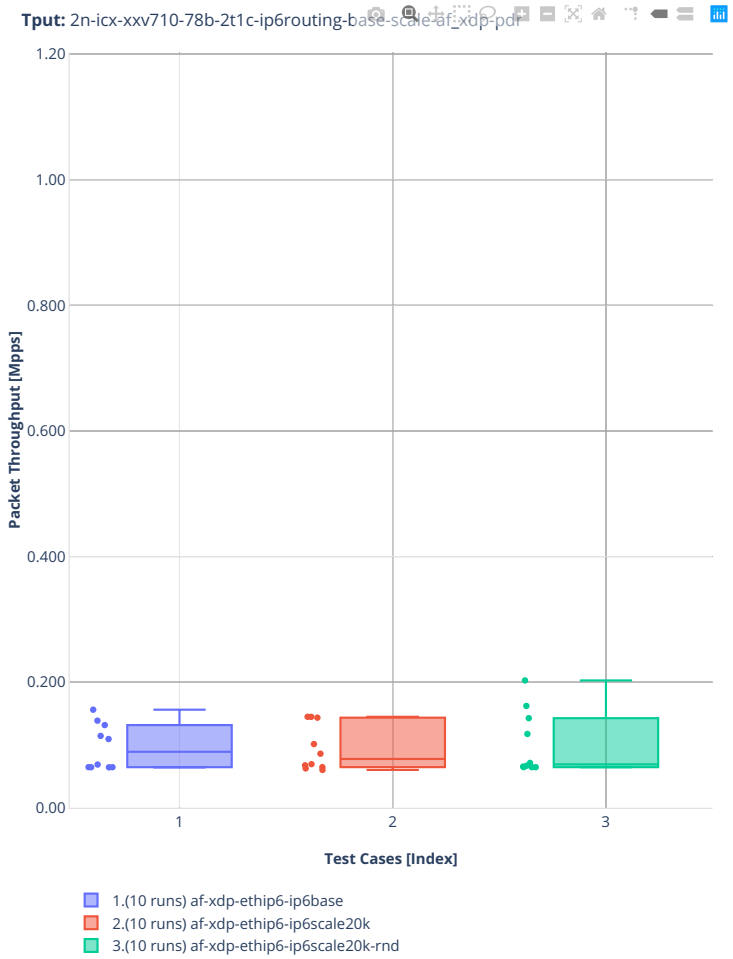
Tput: 2n-icx-xxv710-78b-2t1c-ip6routing-base-scale-af_xdp-pdr
1.20
1.00
0.800
0.600
0.400
0.200
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) af-xdp-ethip6-ip6base
2.(10 runs) af-xdp-ethip6-ip6scale20k
3.(10 runs) af-xdp-ethip6-ip6scale20k-rnd

3n-icx-xxv710

78b-2t1c-ip6routing-base

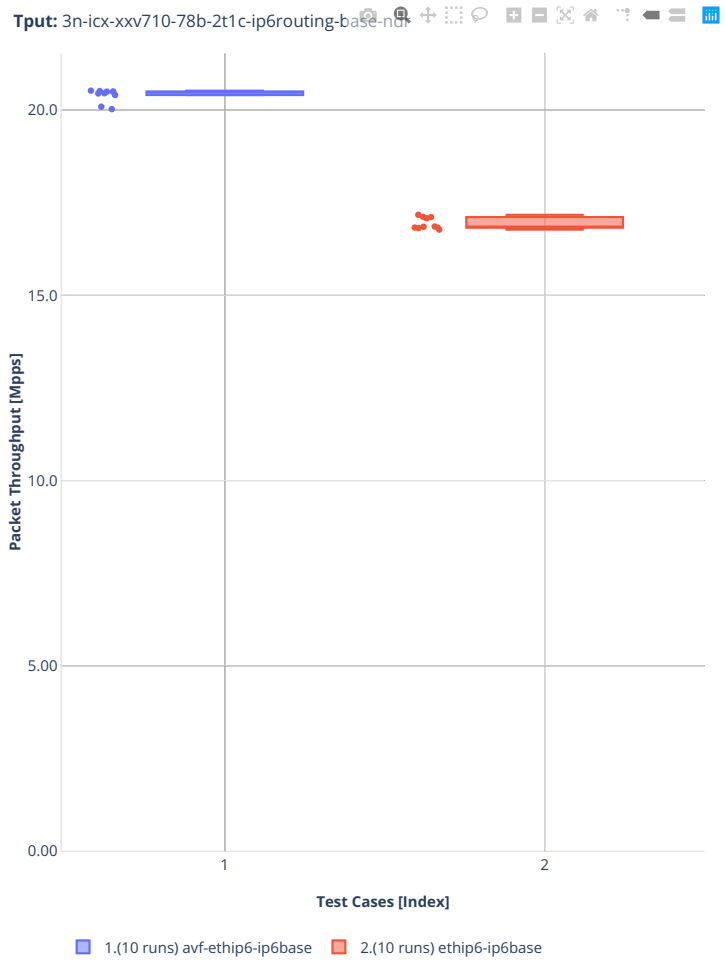

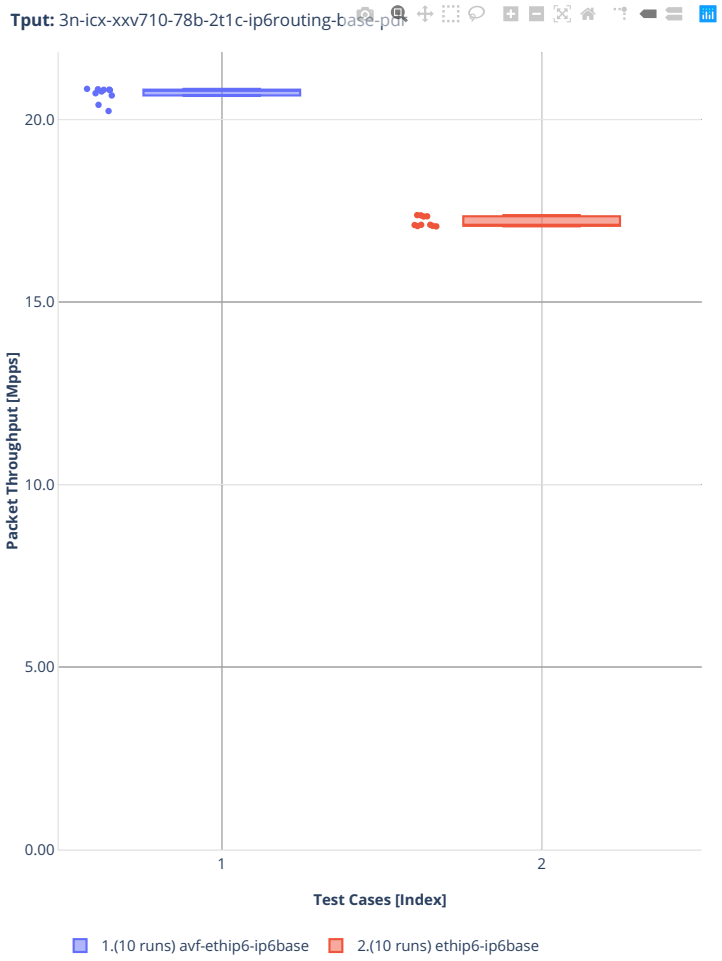
Tput: 3n-icx-xxv710-78b-2t1c-ip6routing-base-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) avf-ethip6-ip6base
2.(10 runs) ethip6-ip6base

2n-skx-xxv710

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

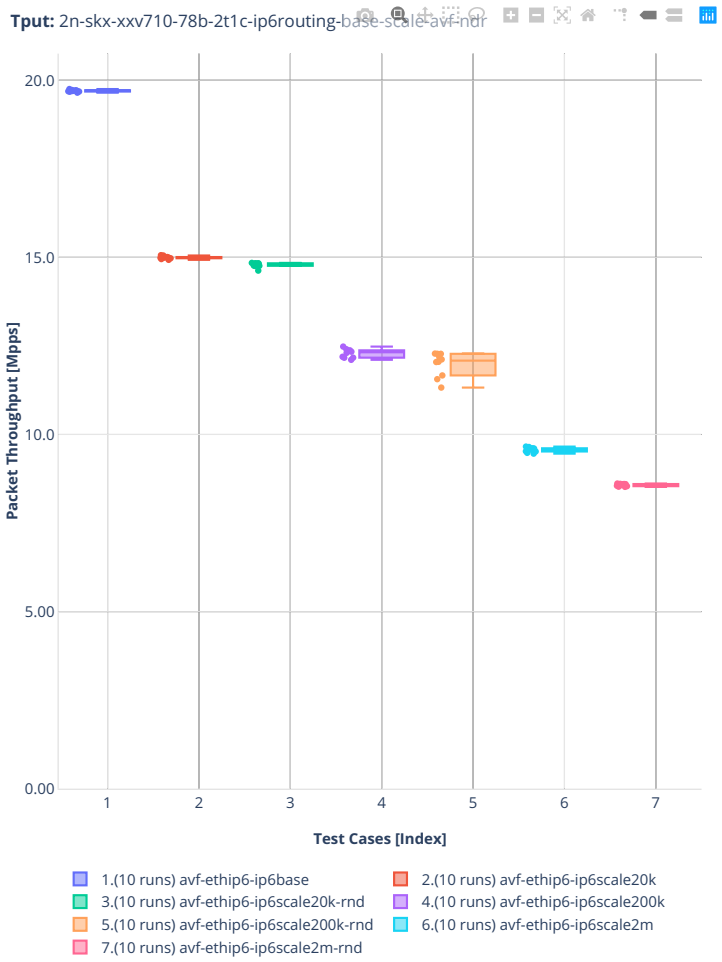

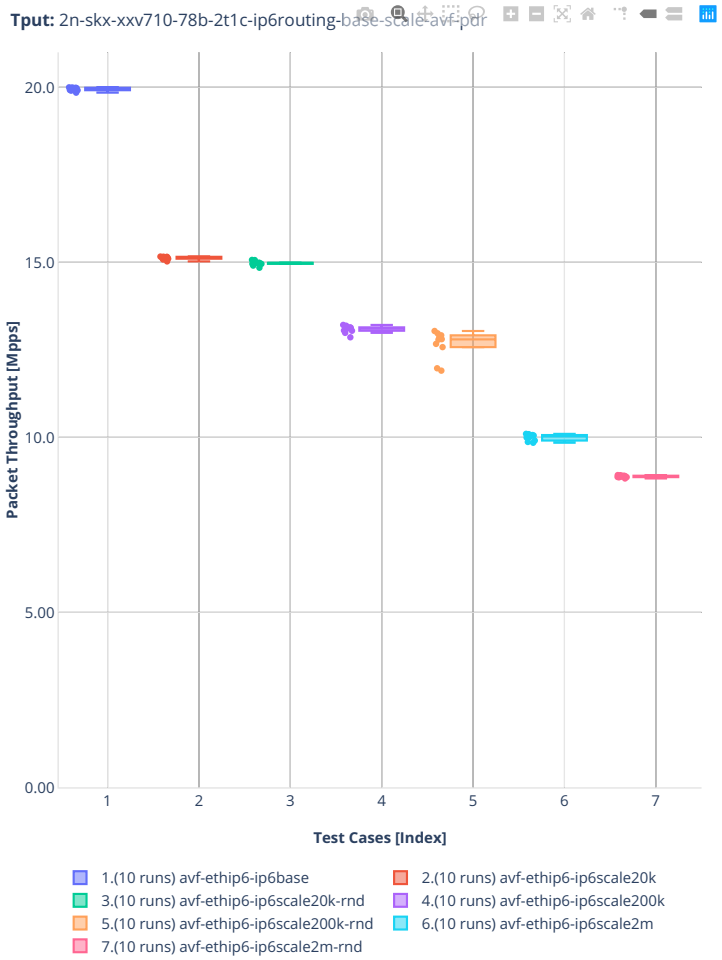
Tput: 2n-skx-xxv710-78b-2t1c-ip6routing-base-scale-avf-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
6
7
Test Cases [Index]
1.(10 runs) avf-ethip6-ip6base
2.(10 runs) avf-ethip6-ip6scale20k
3.(10 runs) avf-ethip6-ip6scale20k-rnd
4.(10 runs) avf-ethip6-ip6scale200k
5.(10 runs) avf-ethip6-ip6scale200k-rnd
6.(10 runs) avf-ethip6-ip6scale2m
7.(10 runs) avf-ethip6-ip6scale2m-rnd

78b-2t1c-ip6routing-base-scale-af-xdp

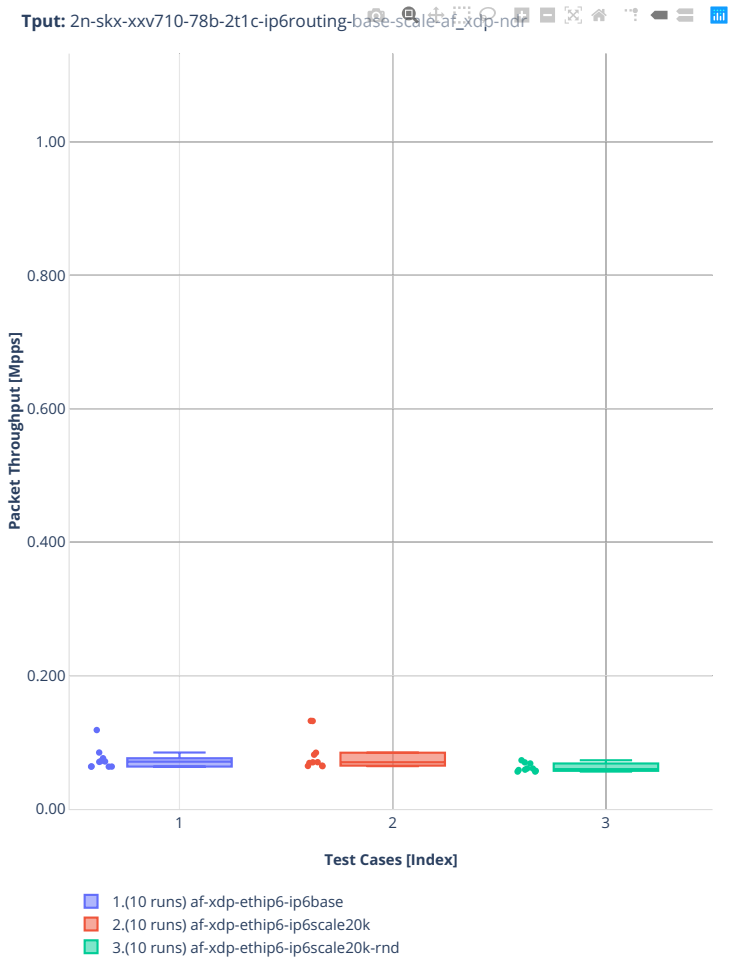

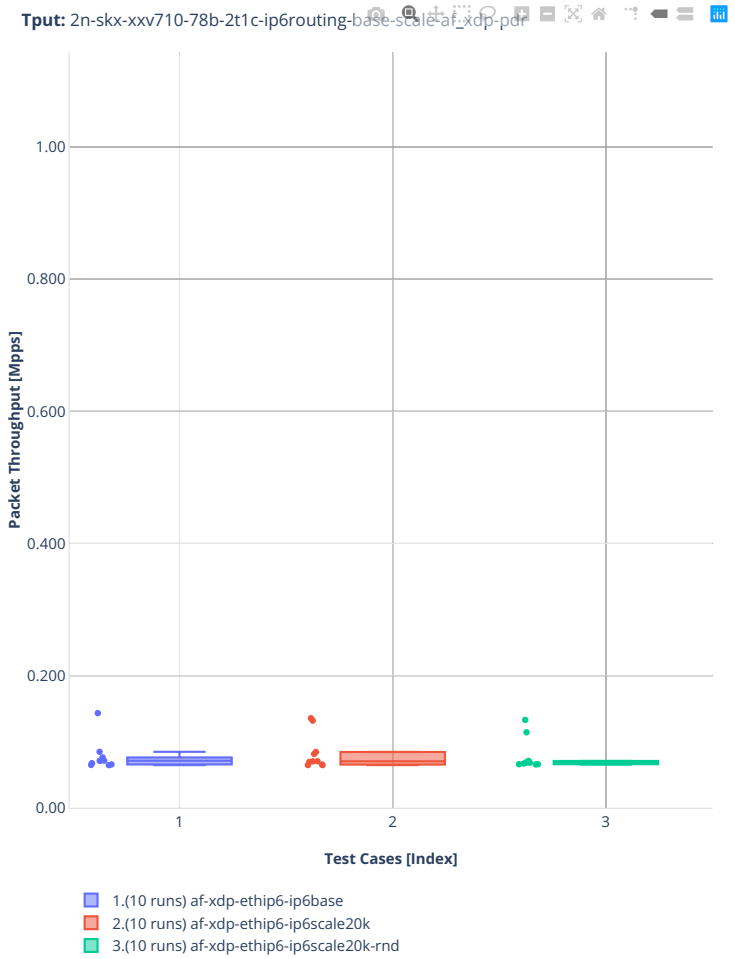
Tput: 2n-skx-xxv710-78b-2t1c-ip6routing-base-scale-af_xdp-pdr
Packet Throughput [Mpps]
1.00
0.800
0.600
0.400
0.200
0.00
1
2
3
Test Cases [Index]
1.(10 runs) af-xdp-ethip6-ip6base
2.(10 runs) af-xdp-ethip6-ip6scale20k
3.(10 runs) af-xdp-ethip6-ip6scale20k-rnd

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

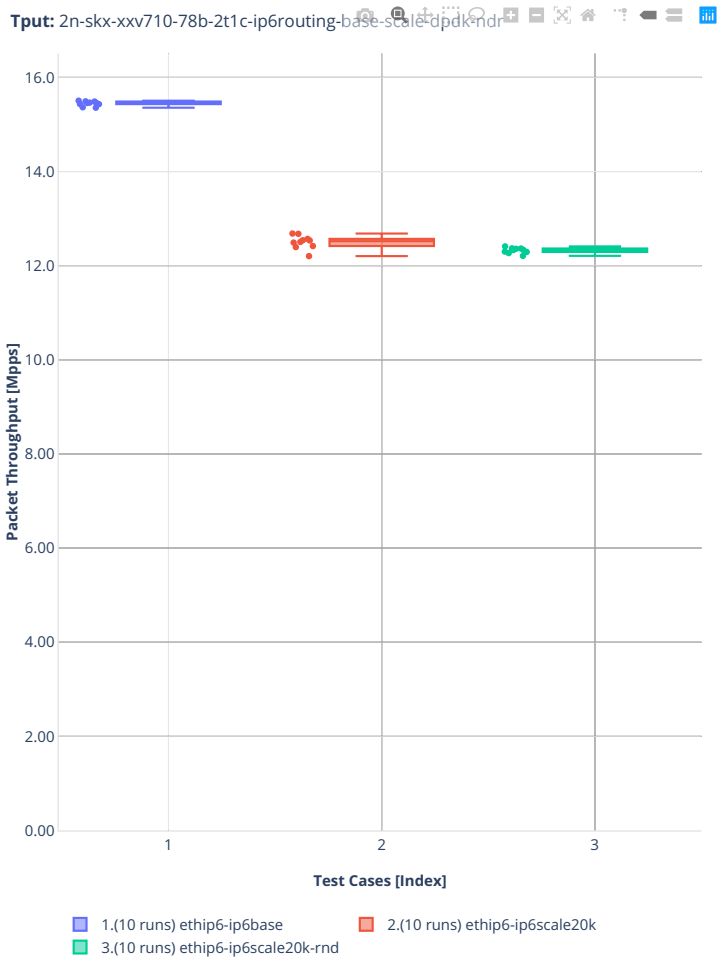

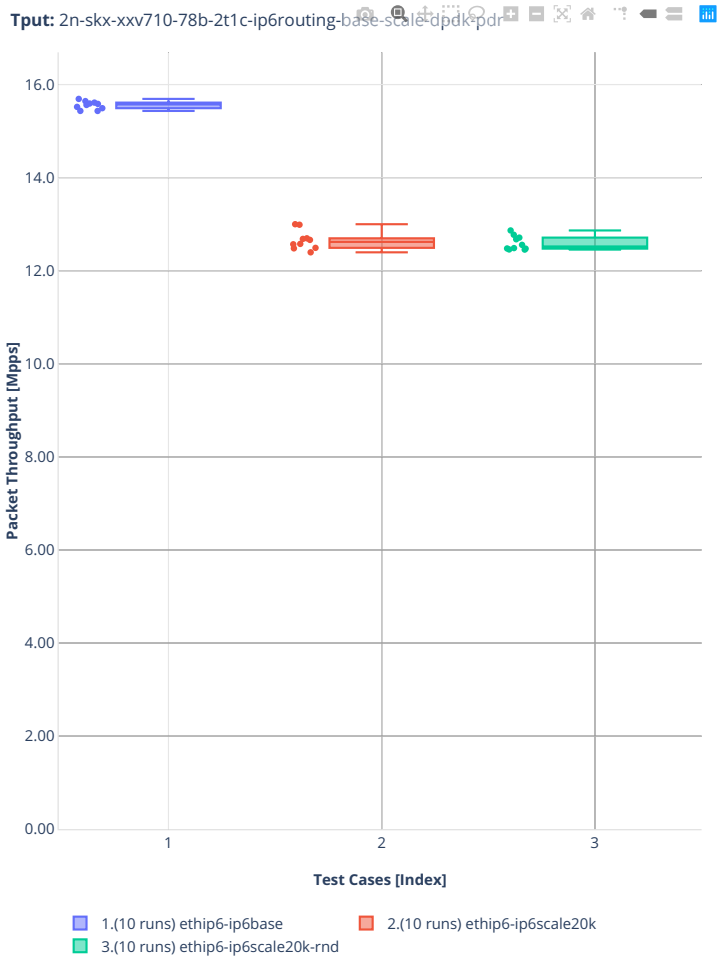
Tput: 2n-skx-xxv710-78b-2t1c-ip6routing-base-scale-dpdk-pdr
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip6-ip6base
2.(10 runs) ethip6-ip6scale20k
3.(10 runs) ethip6-ip6scale20k-rnd

2n-skx-x710

78b-2t1c-ip6routing-base-scale-[avf,dpdk]

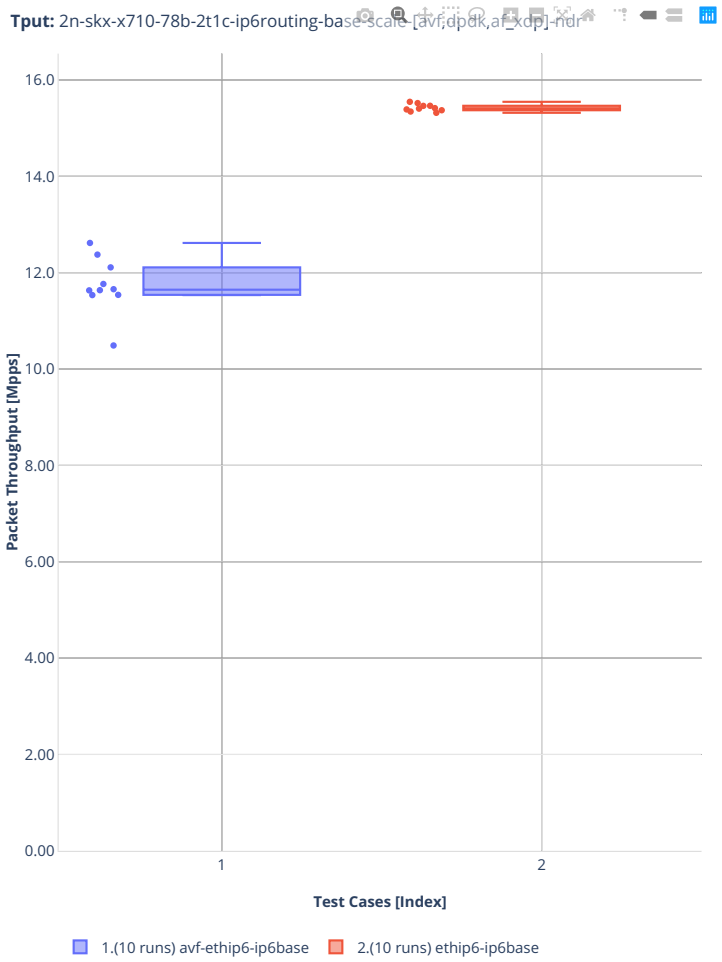

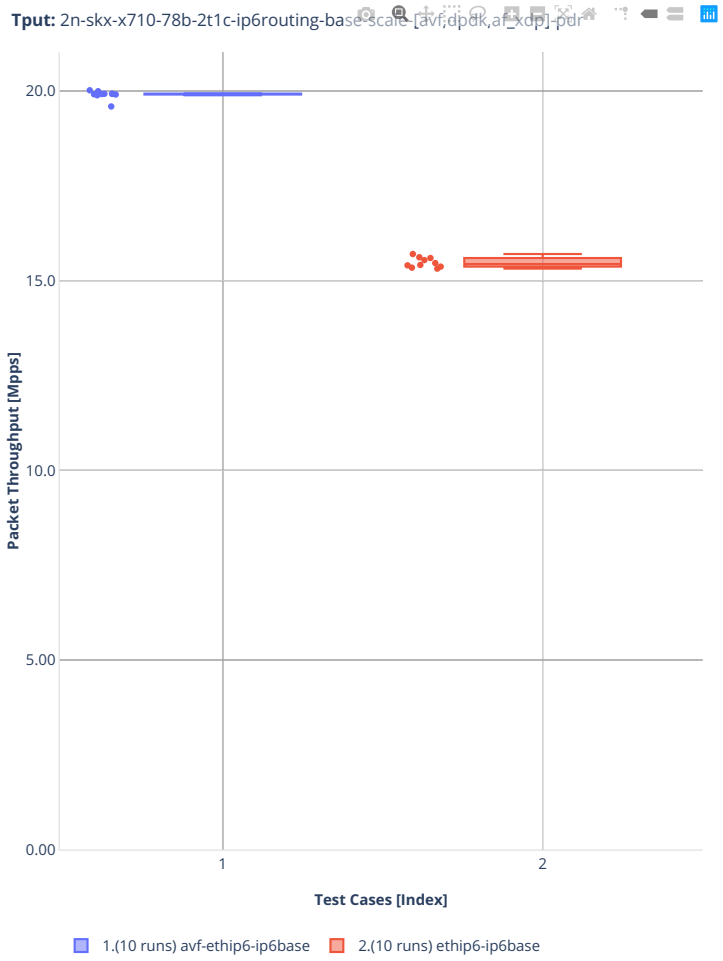
Tput: 2n-skx-x710-78b-2t1c-ip6routing-base-scale-[avf,dpdk,af_xdp]-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) avf-ethip6-ip6base
2.(10 runs) ethip6-ip6base

3n-skx-xxv710

78b-2t1c-ip6routing-base-[avf,dpdk]

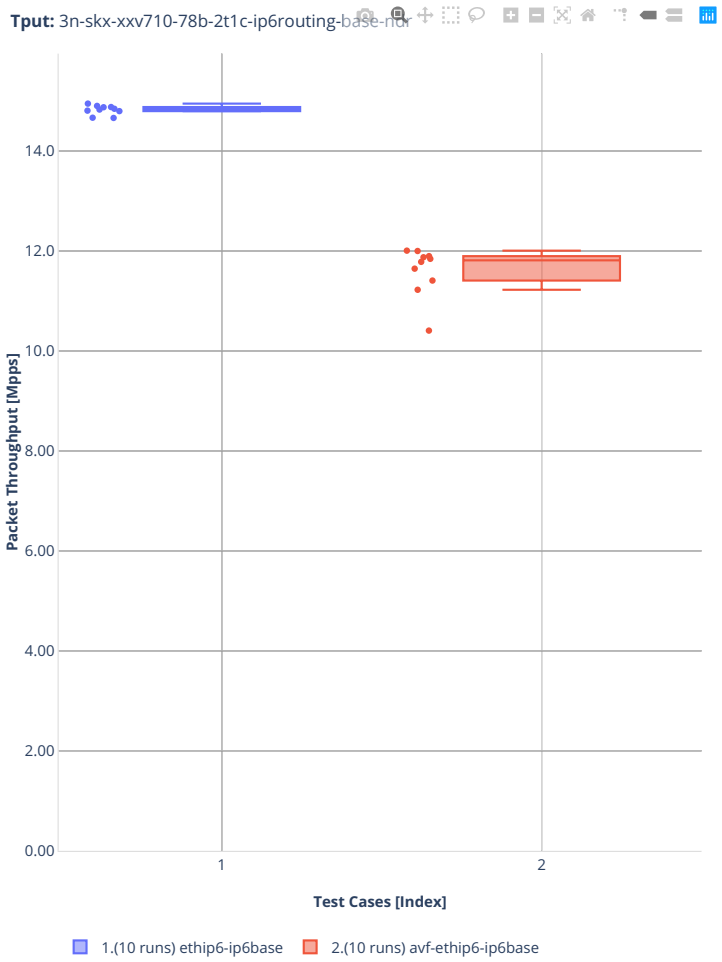

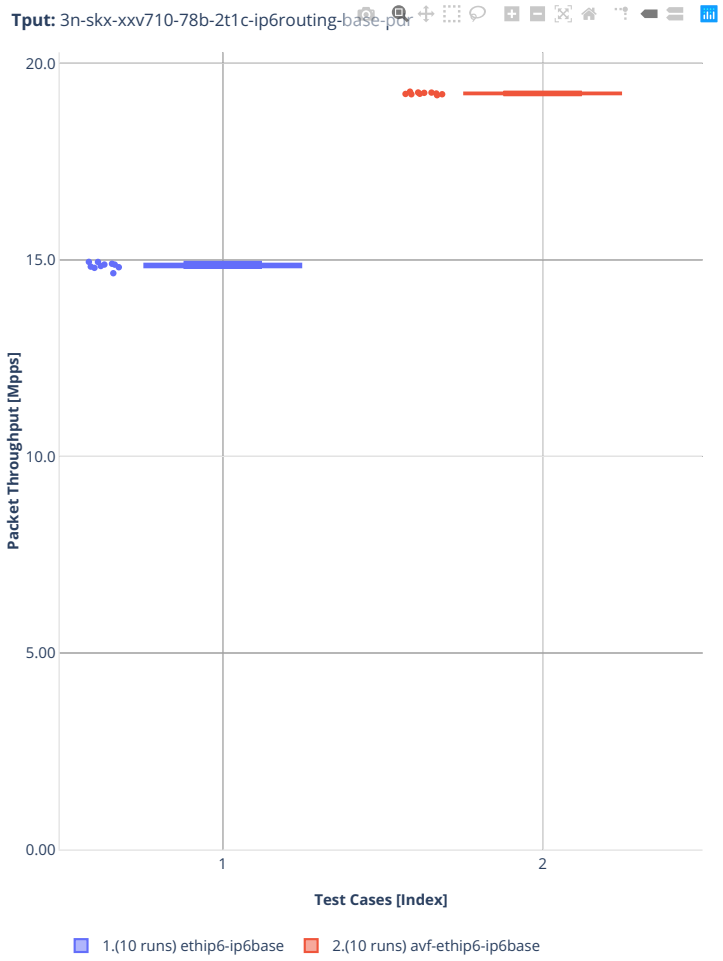
Tput: 3n-skx-xxv710-78b-2t1c-ip6routing-base-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) ethip6-ip6base
2.(10 runs) avf-ethip6-ip6base

3n-skx-x710

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

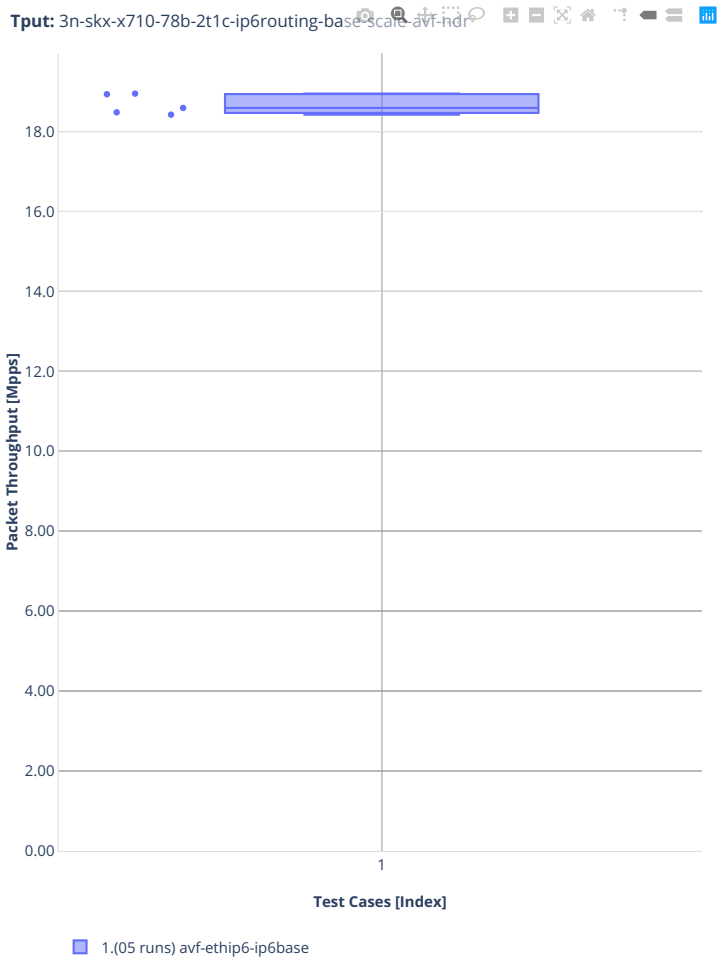

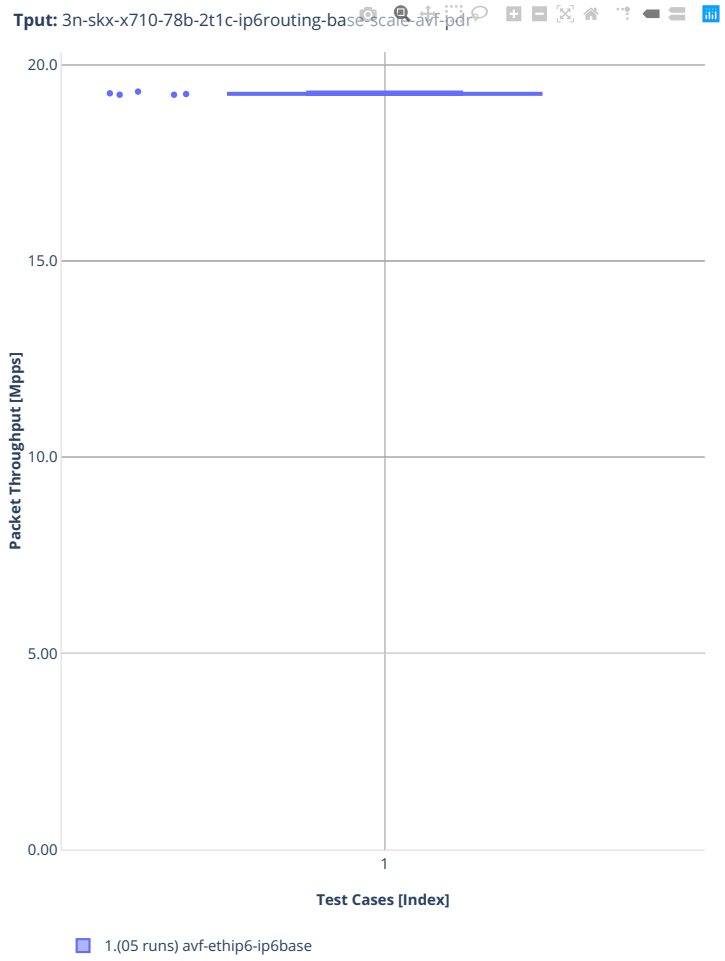
Tput: 3n-skx-x710-78b-2t1c-ip6routing-base-scale-avf-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
Test Cases [Index]
1.(05 runs) avf-ethip6-ip6base

2n-clx-xxv710

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

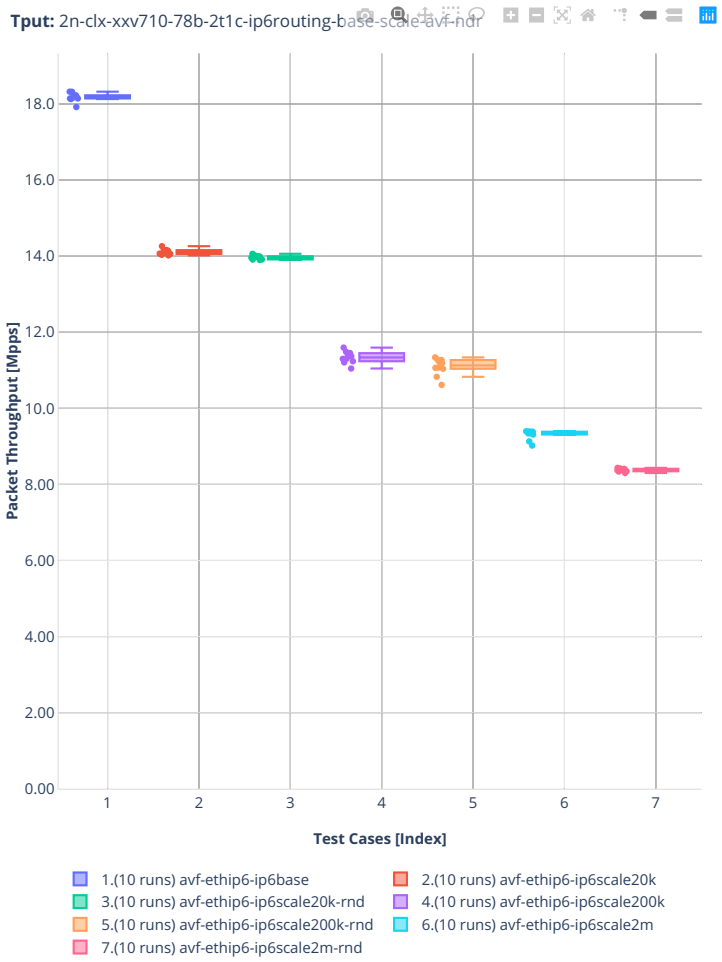

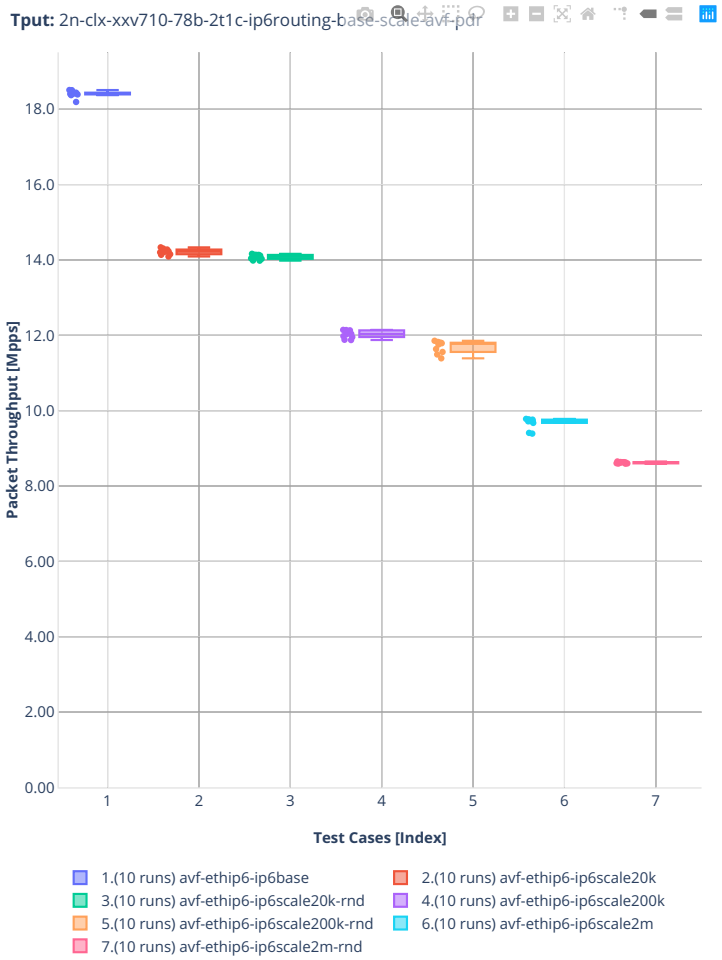
Tput: 2n-clx-xxv710-78b-2t1c-ip6routing-base-scale-avf-pdr
Packet Throughput [Mpps]
18.0
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
5
6
7
Test Cases [Index]
1.(10 runs) avf-ethip6-ip6base
2.(10 runs) avf-ethip6-ip6scale20k
3.(10 runs) avf-ethip6-ip6scale20k-rnd
4.(10 runs) avf-ethip6-ip6scale200k
5.(10 runs) avf-ethip6-ip6scale200k-rnd
6.(10 runs) avf-ethip6-ip6scale2m
7.(10 runs) avf-ethip6-ip6scale2m-rnd

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

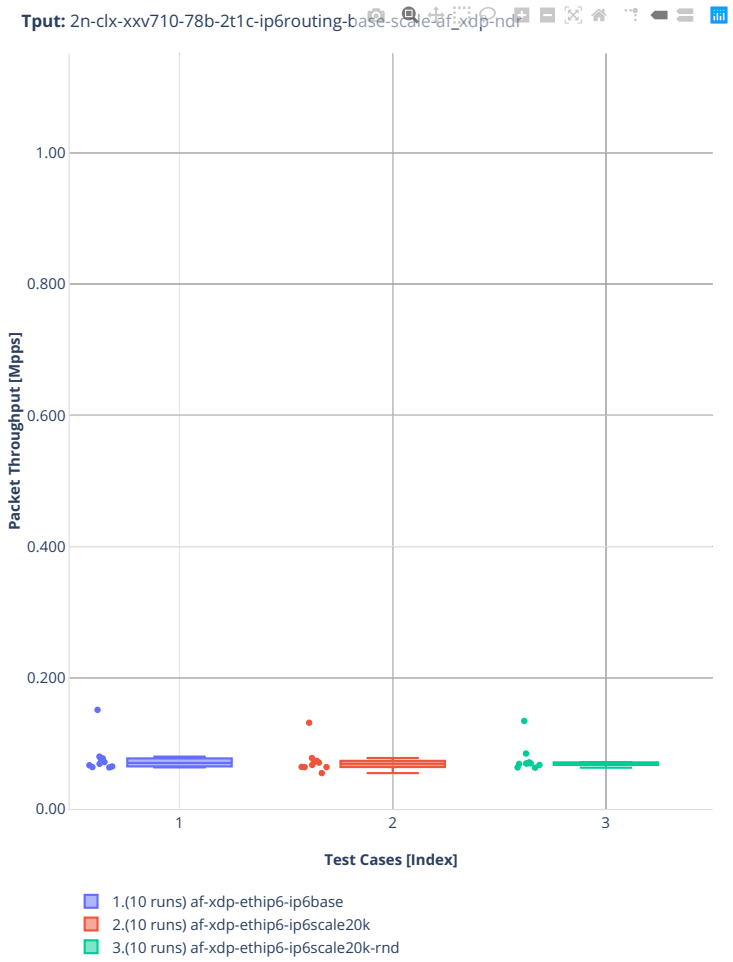

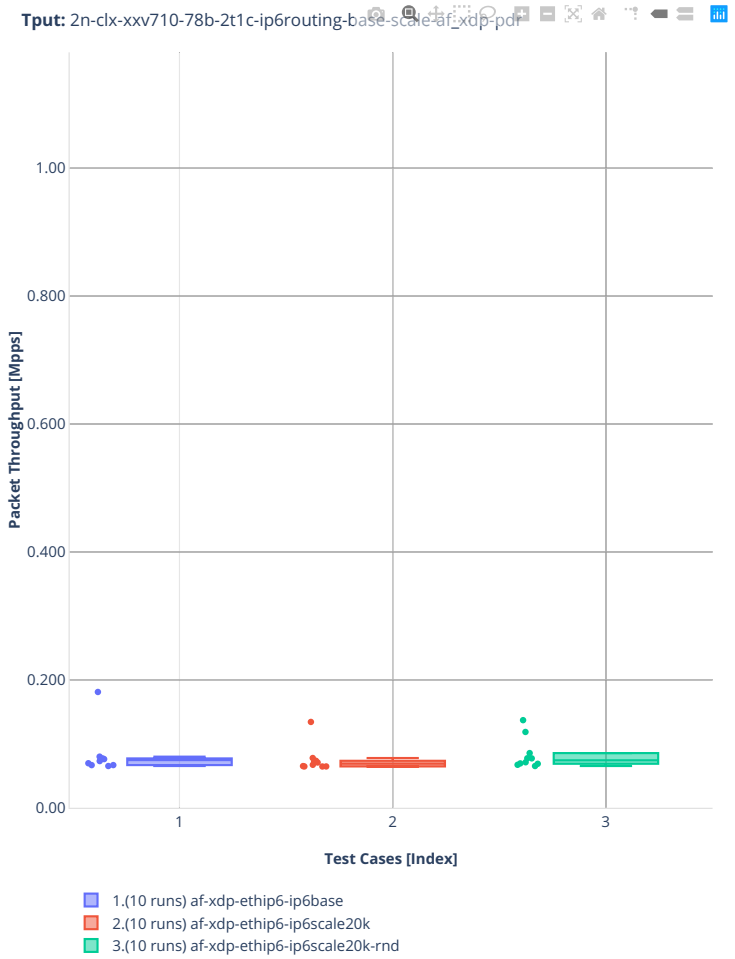
Tput: 2n-clx-xxv710-78b-2t1c-ip6routing-base-scale-af_xdp-pdr
Packet Throughput [Mpps]
1.00
0.800
0.600
0.400
0.200
0.00
1
2
3
Test Cases [Index]
1.(10 runs) af-xdp-ethip6-ip6base
2.(10 runs) af-xdp-ethip6-ip6scale20k
3.(10 runs) af-xdp-ethip6-ip6scale20k-rnd

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

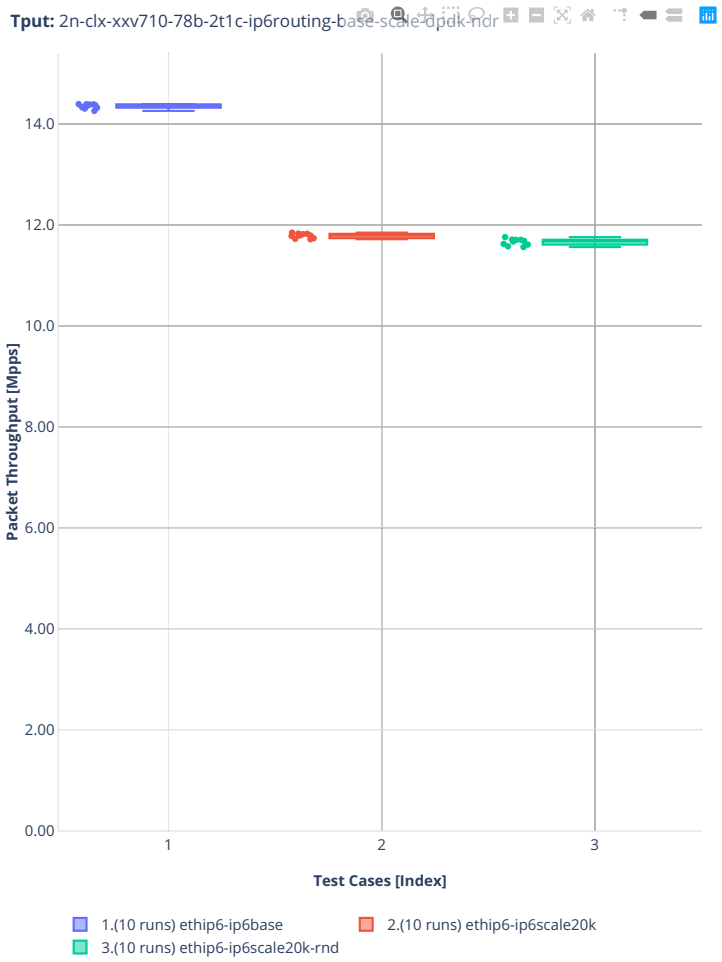

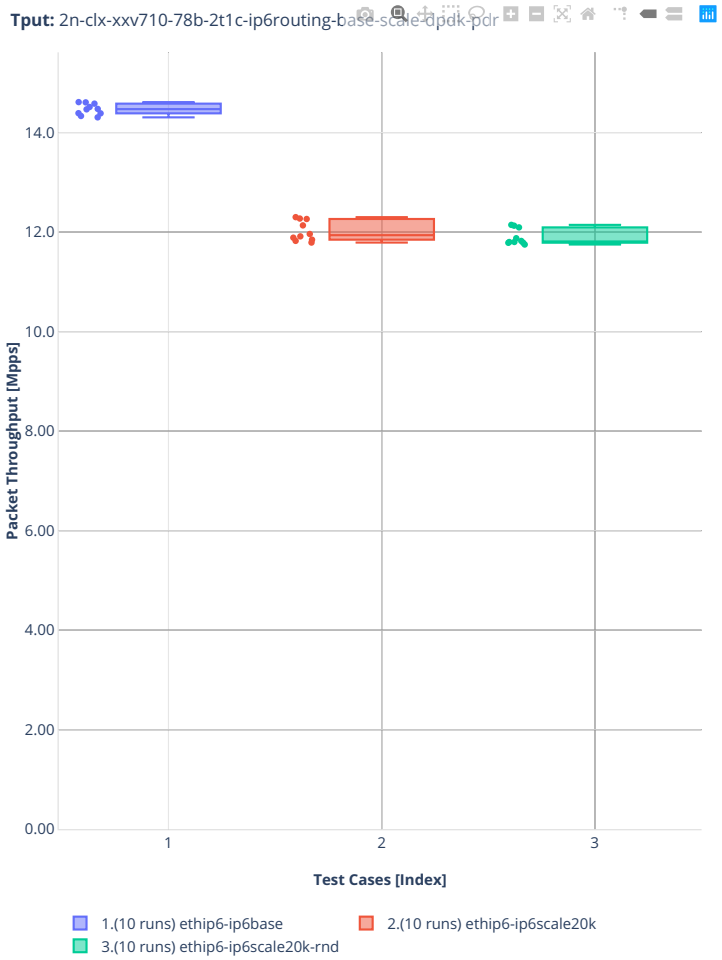
Tput: 2n-clx-xxv710-78b-2t1c-ip6routing-base-scale-dpdk-pdr
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip6-ip6base
2.(10 runs) ethip6-ip6scale20k
3.(10 runs) ethip6-ip6scale20k-rnd

2n-clx-x710

78b-2t1c-ip6routing-base-scale-[avf,dpdk]

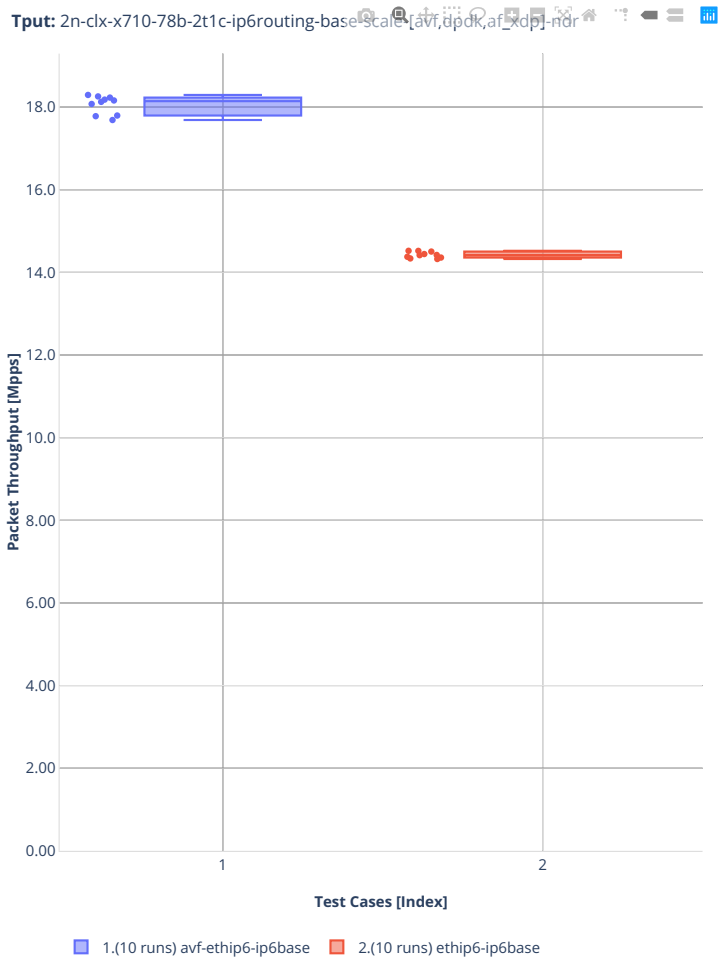

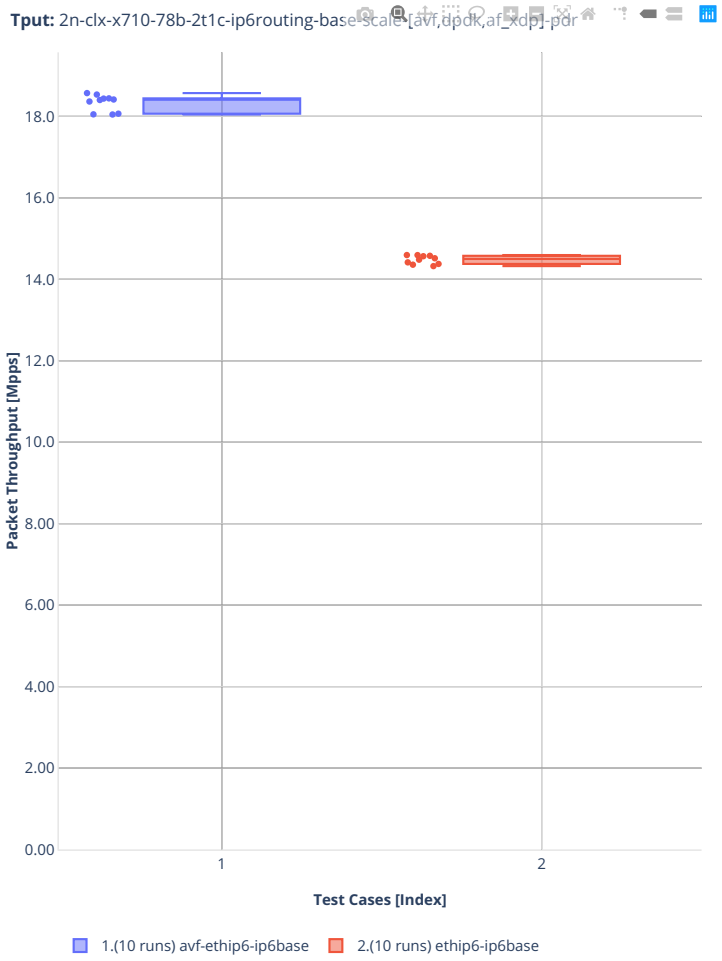
Tput: 2n-clx-x710-78b-2t1c-ip6routing-base-scale-[avf,dpdk,af_xdp]-pdr
Packet Throughput [Mpps]
18.0
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
1
2
Test Cases [Index]
1.(10 runs) avf-ethip6-ip6base
2.(10 runs) ethip6-ip6base

2n-clx-cx556a

78b-2t1c-ip6routing-base-scale-rdma-core

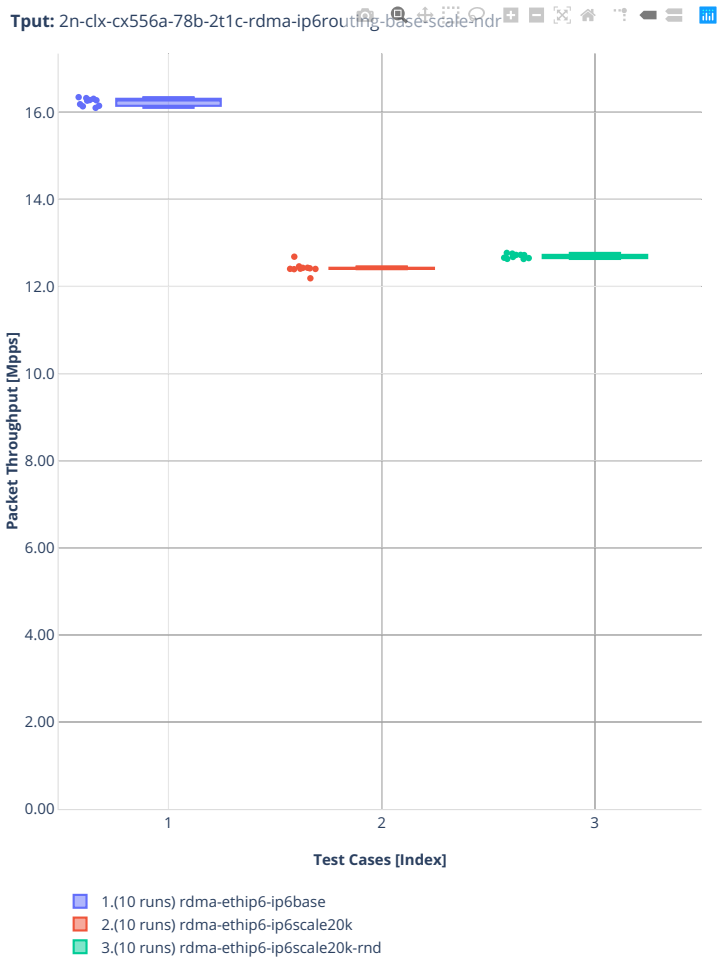

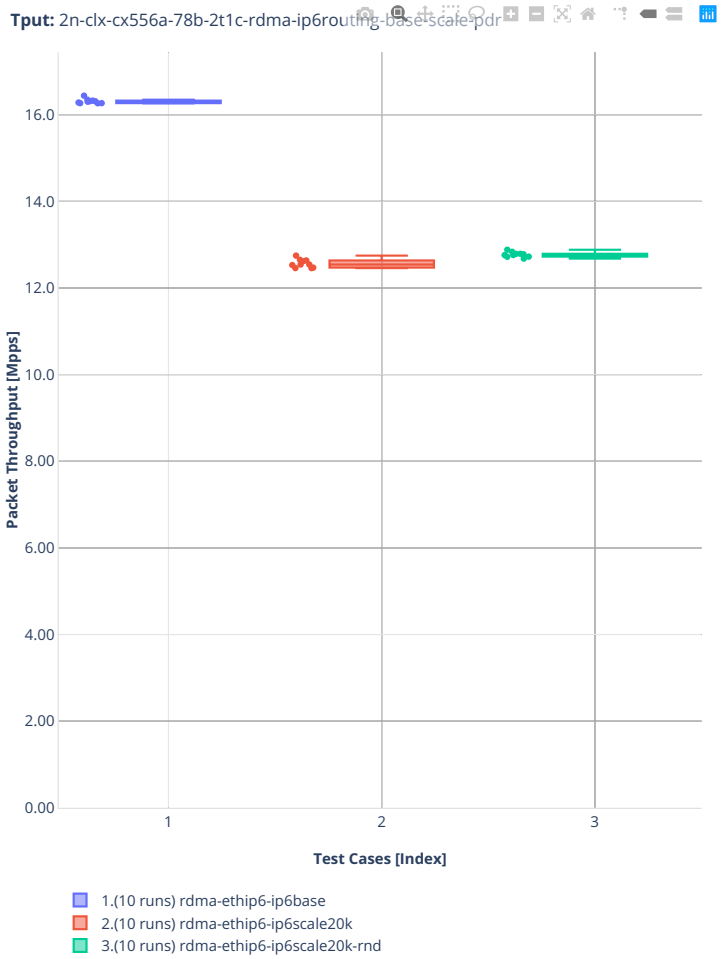
Tput: 2n-clx-cx556a-78b-2t1c-rdma-ip6routing-base-scale-pdr
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) rdma-ethip6-ip6base
2.(10 runs) rdma-ethip6-ip6scale20k
3.(10 runs) rdma-ethip6-ip6scale20k-rnd

2n-clx-e810cq

78b-2t1c-ip6routing-base-scale

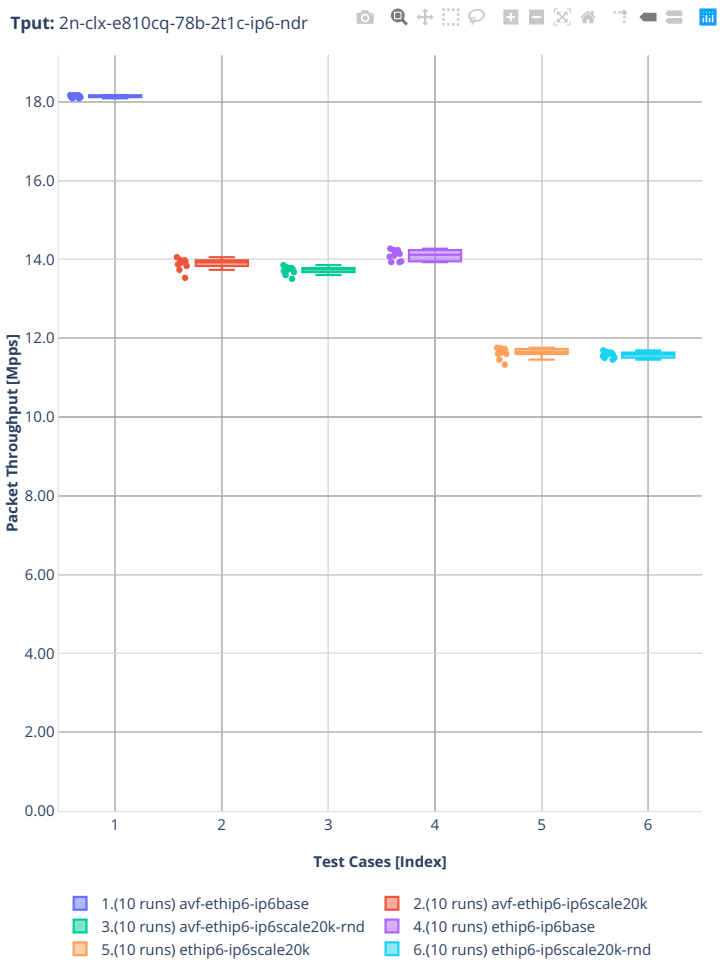

2n-zn2-xxv710

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

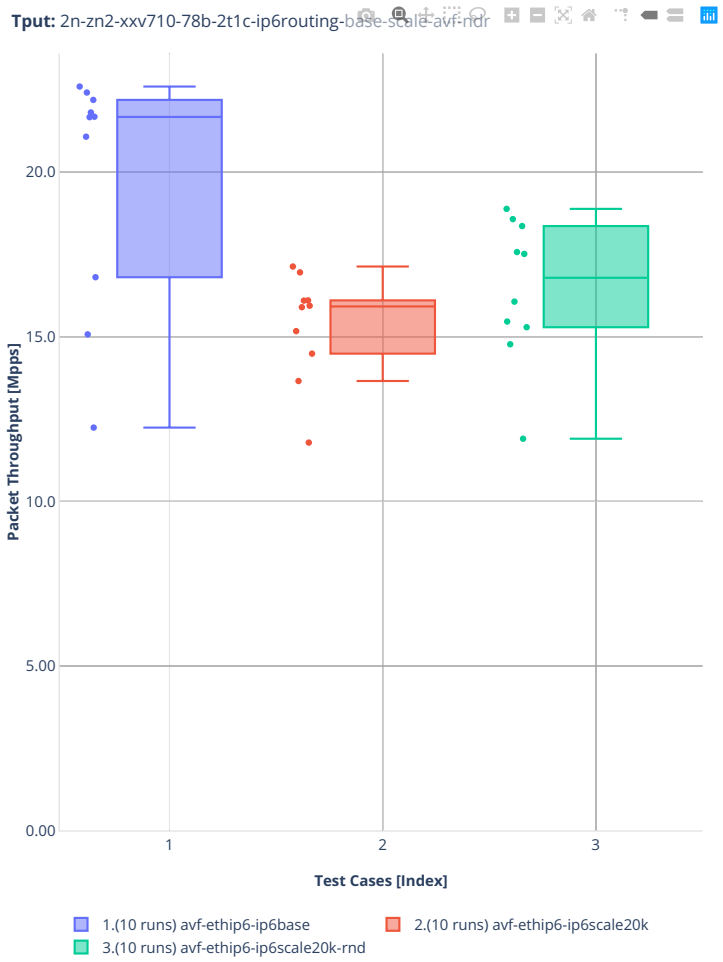

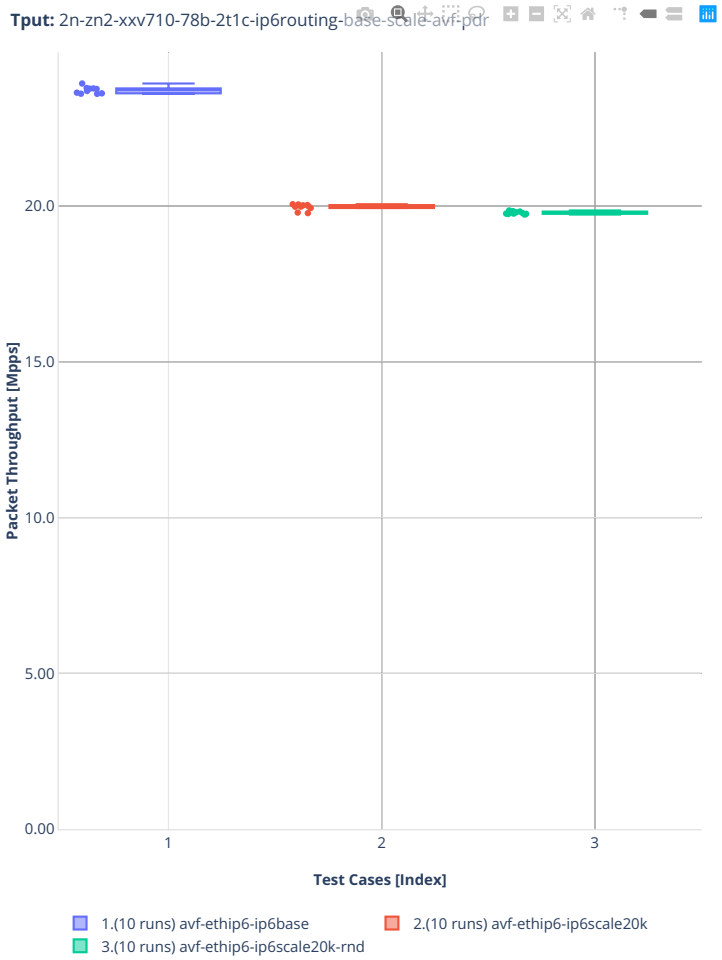
Tput: 2n-zn2-xxv710-78b-2t1c-ip6routing-base-scale-avf-pdr
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-ethip6-ip6base
2.(10 runs) avf-ethip6-ip6scale20k
3.(10 runs) avf-ethip6-ip6scale20k-rnd

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

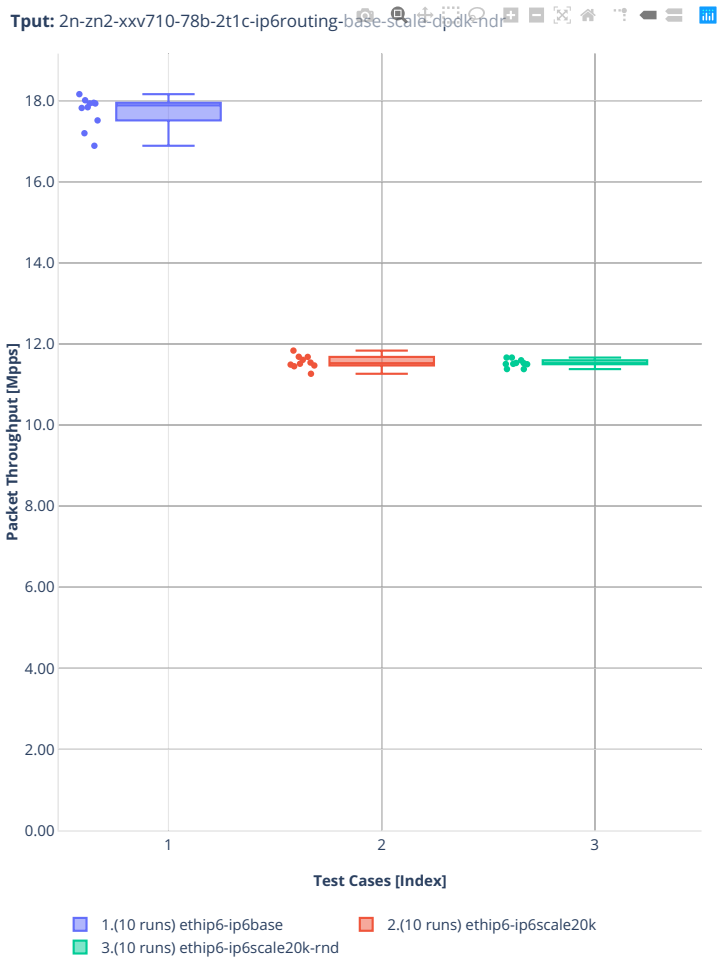

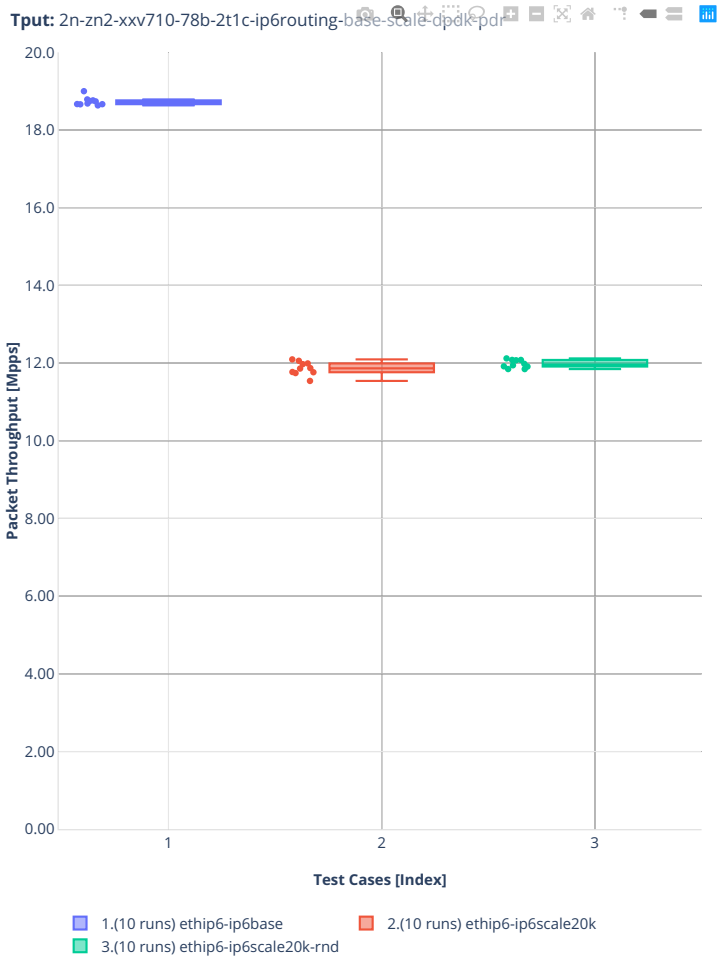
Tput: 2n-zn2-xxv710-78b-2t1c-ip6routing-base-scale-dpdk-pdr
20.0
18.0
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip6-ip6base
2.(10 runs) ethip6-ip6scale20k
3.(10 runs) ethip6-ip6scale20k-rnd

2n-zn2-x710

78b-2t1c-ip6routing-base-[avf,dpdk]

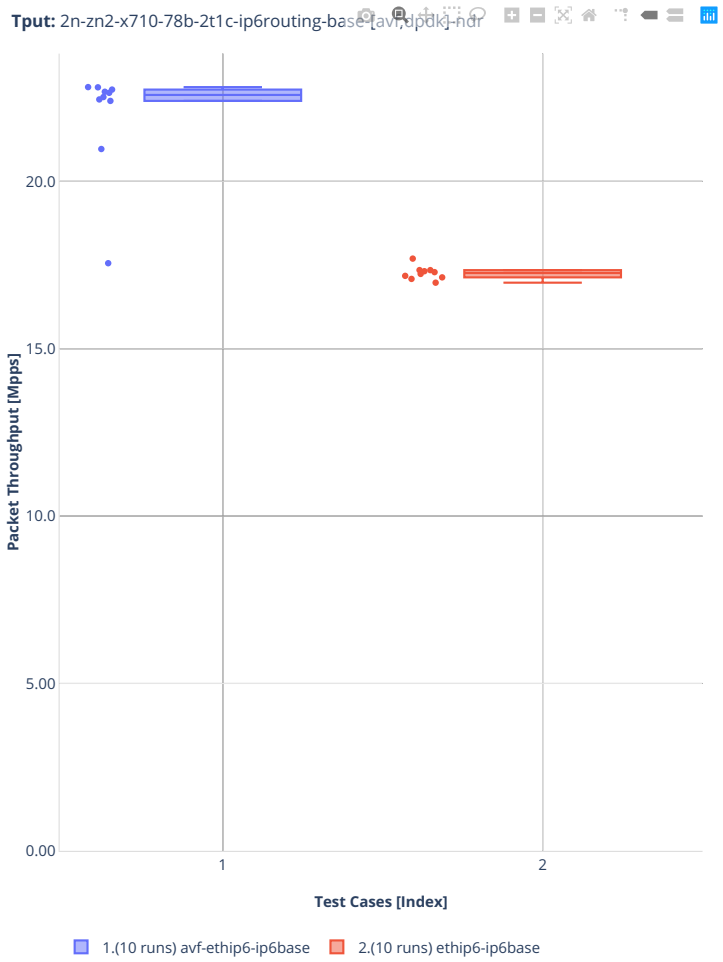

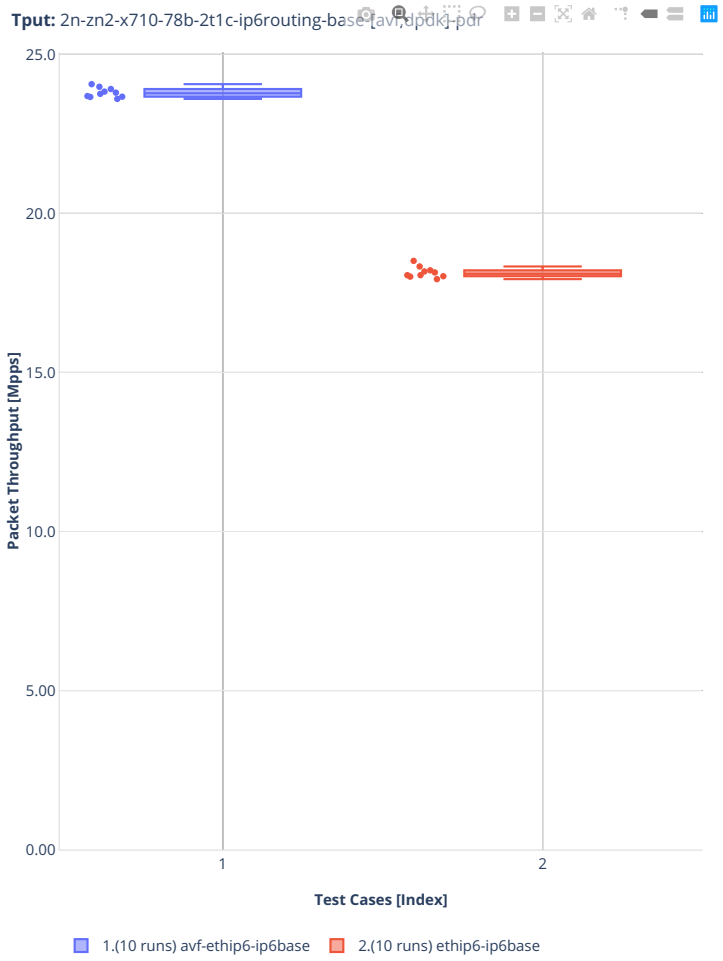
Tput: 2n-zn2-x710-78b-2t1c-ip6routing-base-[avf,dpdk]-pdr
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) avf-ethip6-ip6base
2.(10 runs) ethip6-ip6base

2n-zn2-cx556a

78b-2t1c-ip6routing-base-scale-rdma-core

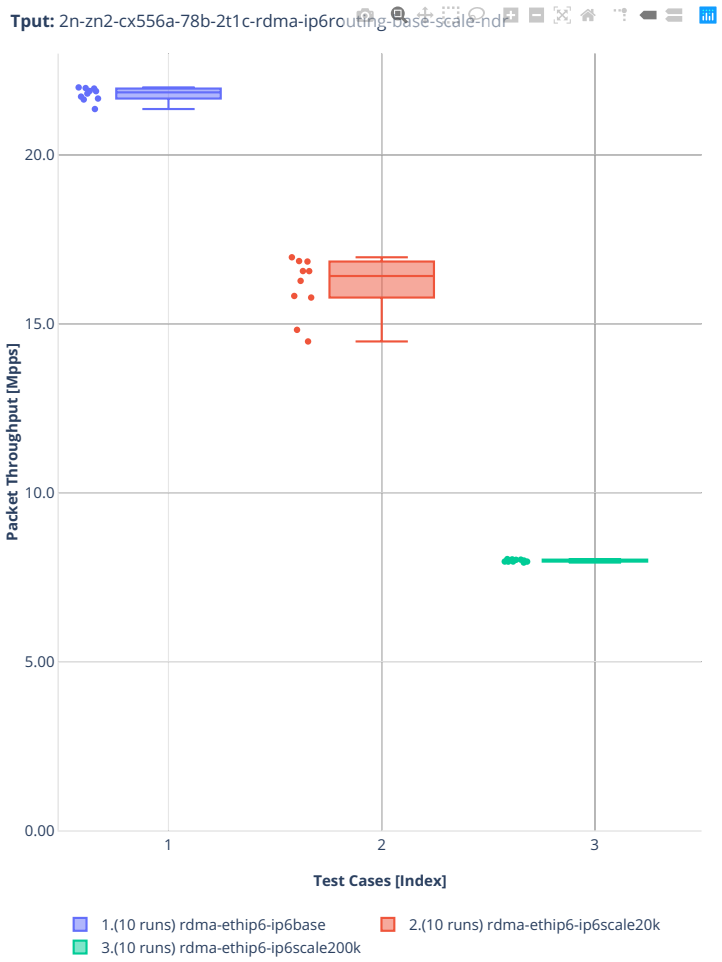

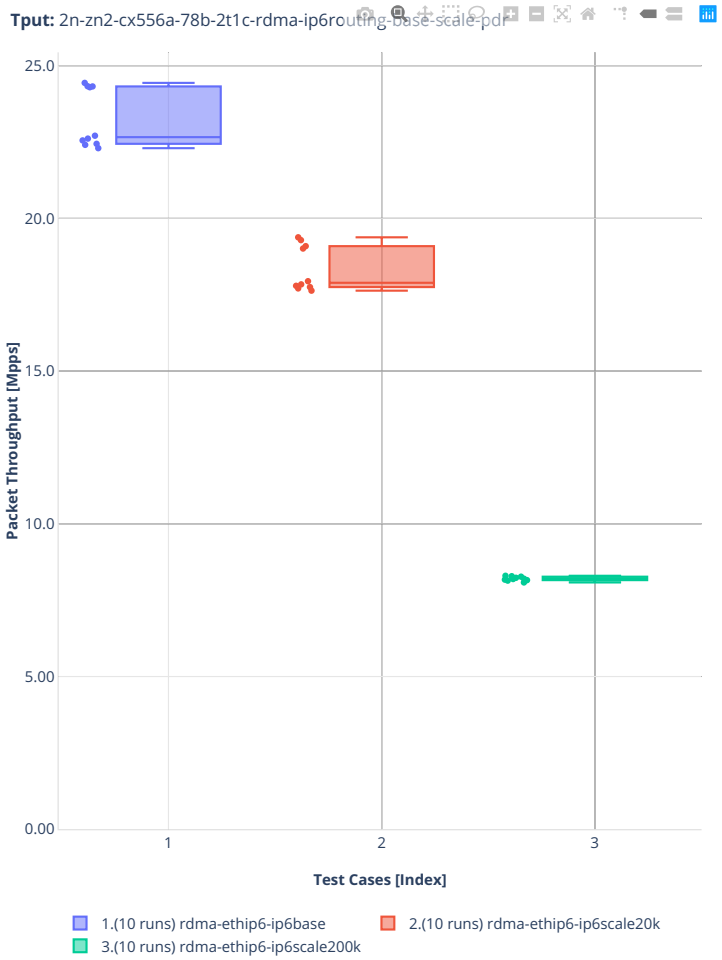
Tput: 2n-zn2-cx556a-78b-2t1c-rdma-ip6routing-base-scale-pdr
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) rdma-ethip6-ip6base
2.(10 runs) rdma-ethip6-ip6scale20k
3.(10 runs) rdma-ethip6-ip6scale200k

3n-alt-xl710

78b-1t1c-ip6routing-base-scale

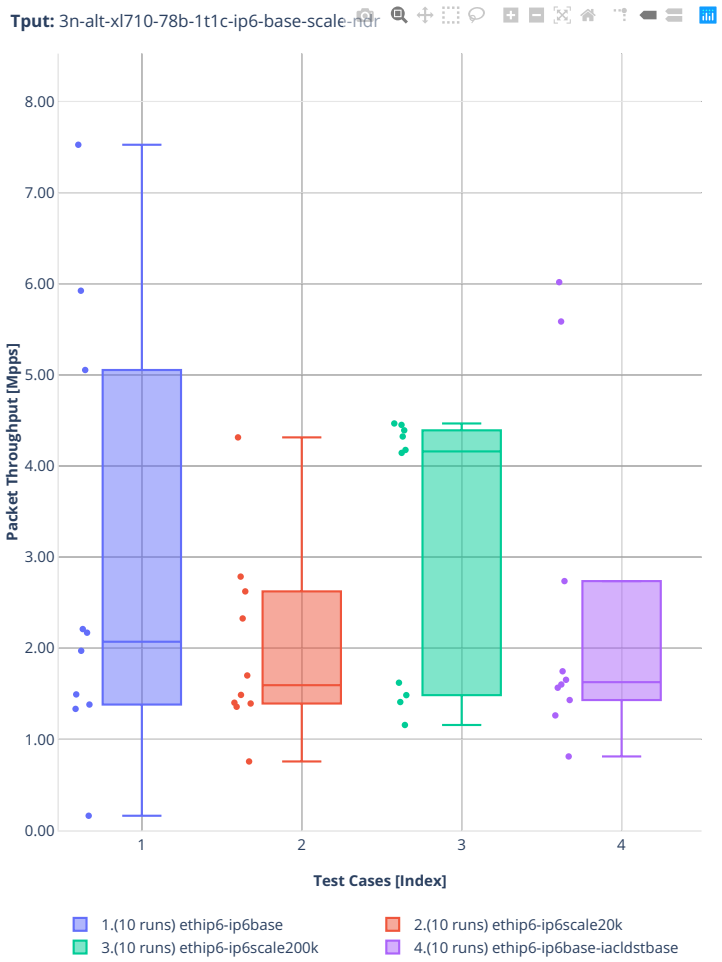

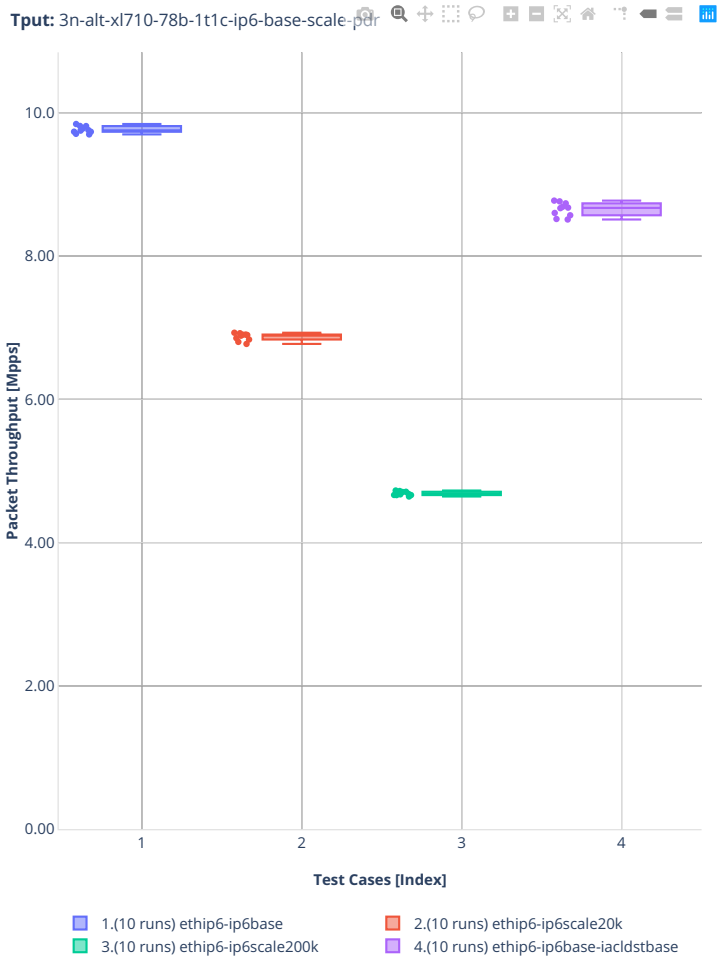
Tput: 3n-alt-xl710-78b-1t1c-ip6-base-scale-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) ethip6-ip6base
2.(10 runs) ethip6-ip6scale20k
3.(10 runs) ethip6-ip6scale200k
4.(10 runs) ethip6-ip6base-iacldstbase

3n-tsh-x520

78b-1t1c-ip6routing-base-ixgbe

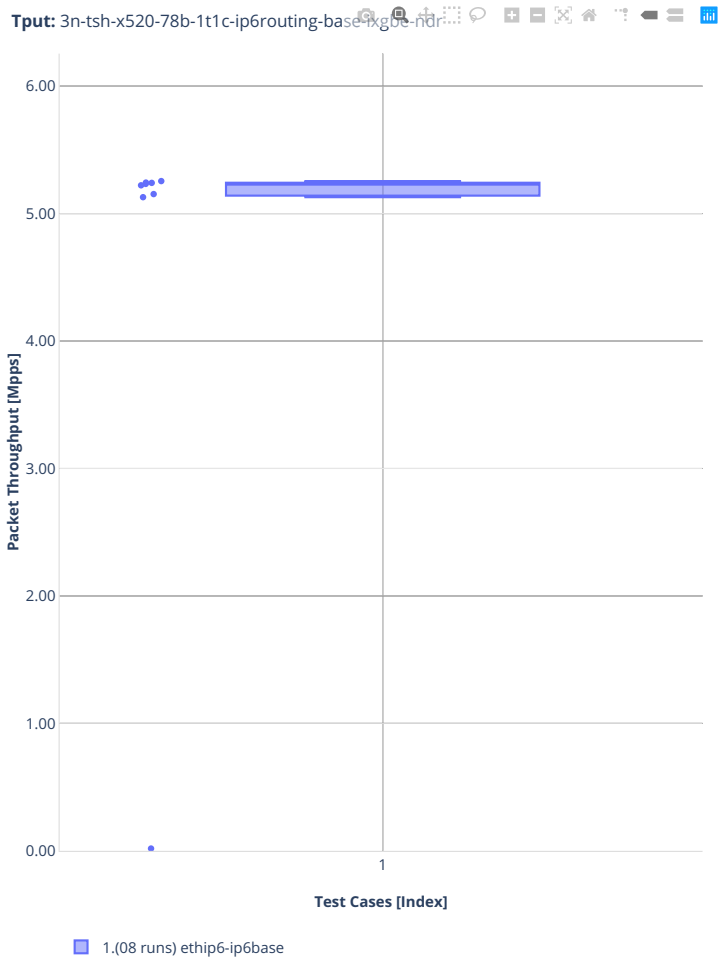

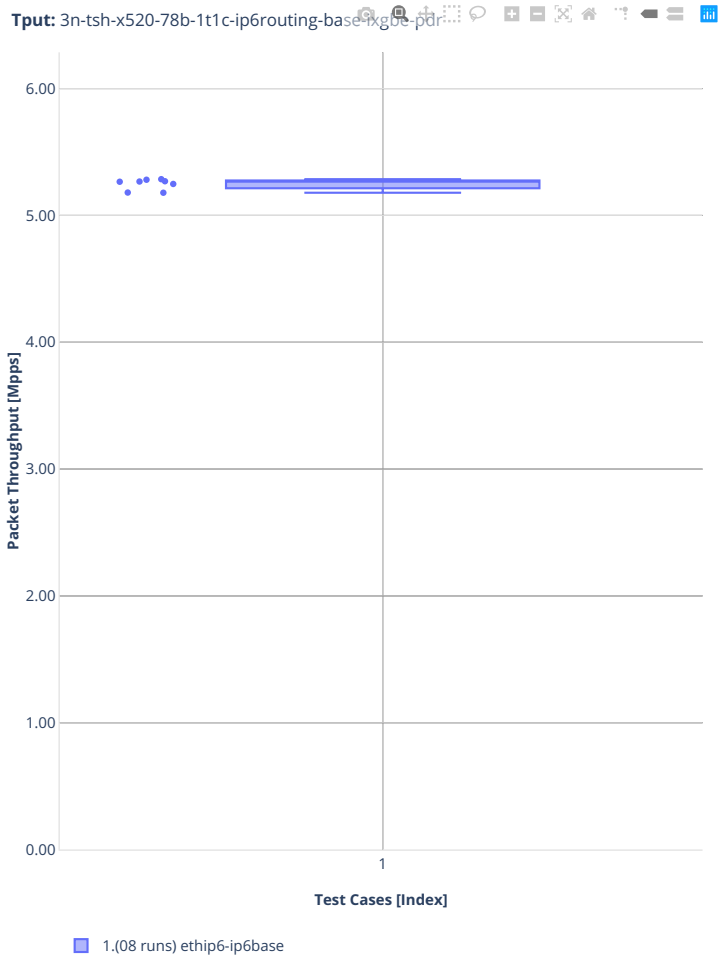
Tput: 3n-tsh-x520-78b-1t1c-ip6routing-base-ixgbe-pdr
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
Test Cases [Index]
1.(08 runs) ethip6-ip6base

2n-tx2-xl710

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

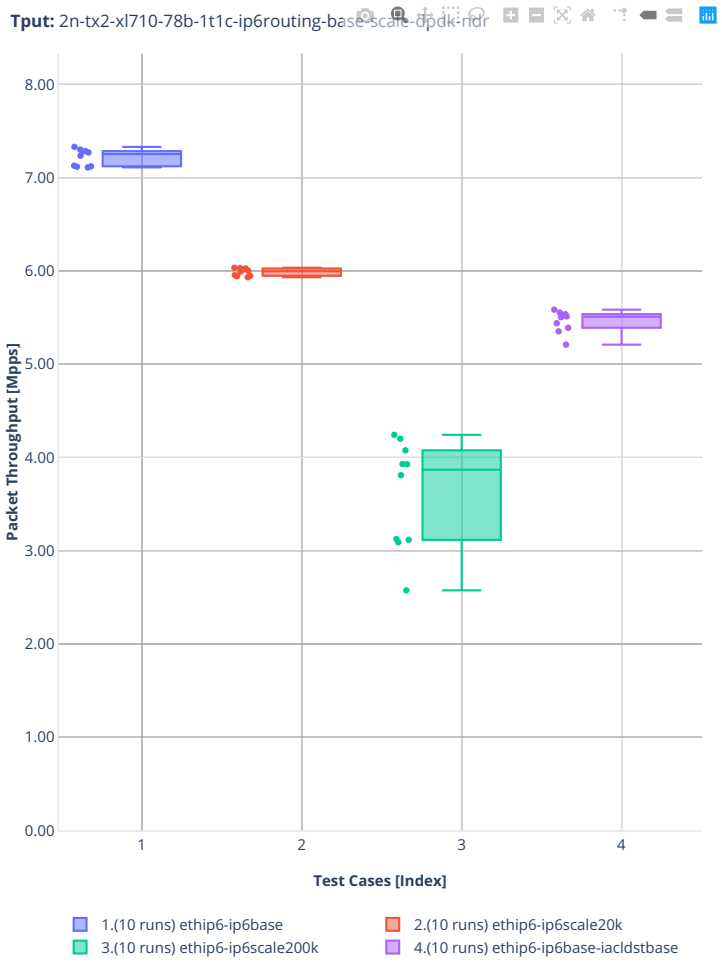

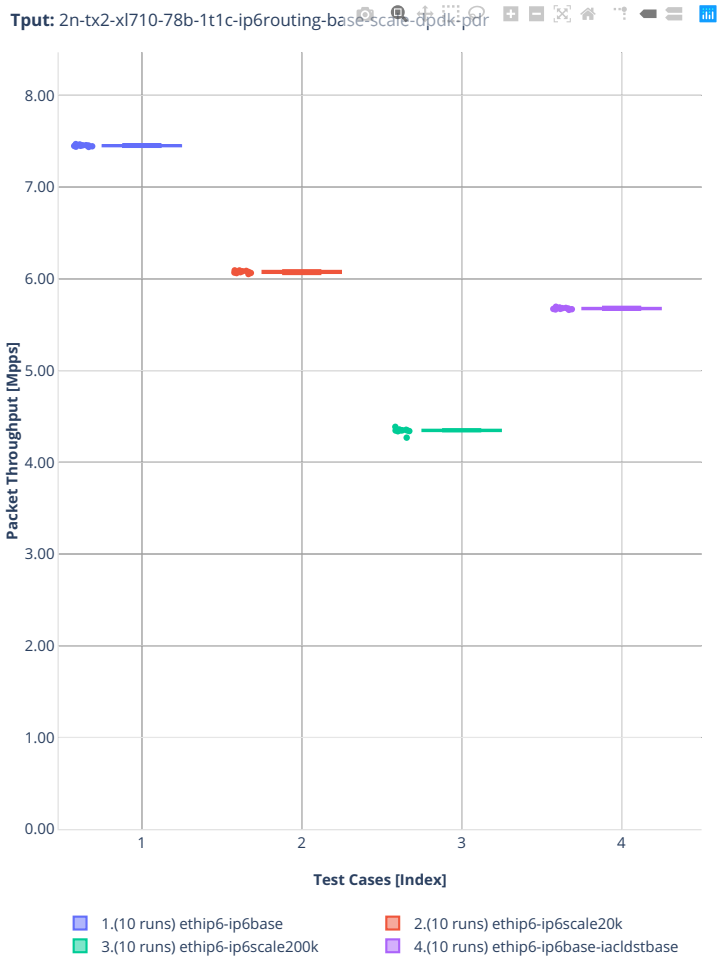
Tput: 2n-tx2-xl710-78b-1t1c-ip6routing-base-scale-dpdk-pdr
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) ethip6-ip6base
2.(10 runs) ethip6-ip6scale20k
3.(10 runs) ethip6-ip6scale200k
4.(10 runs) ethip6-ip6base-iacldstbase

78b-1t1c-ip6routing-base-scale-af-xdp

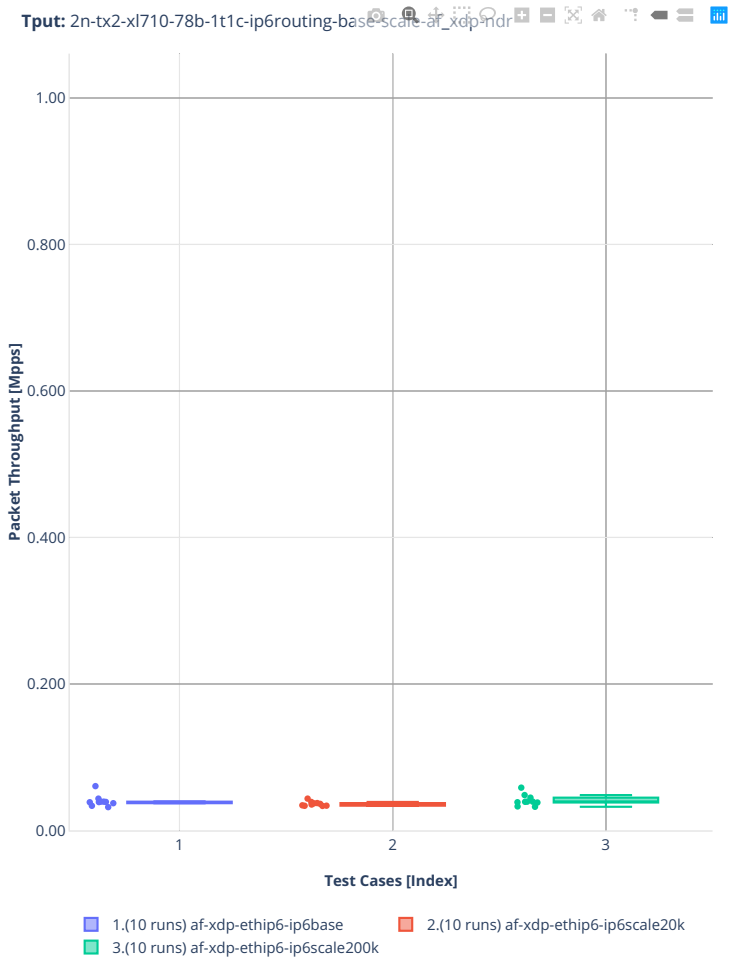

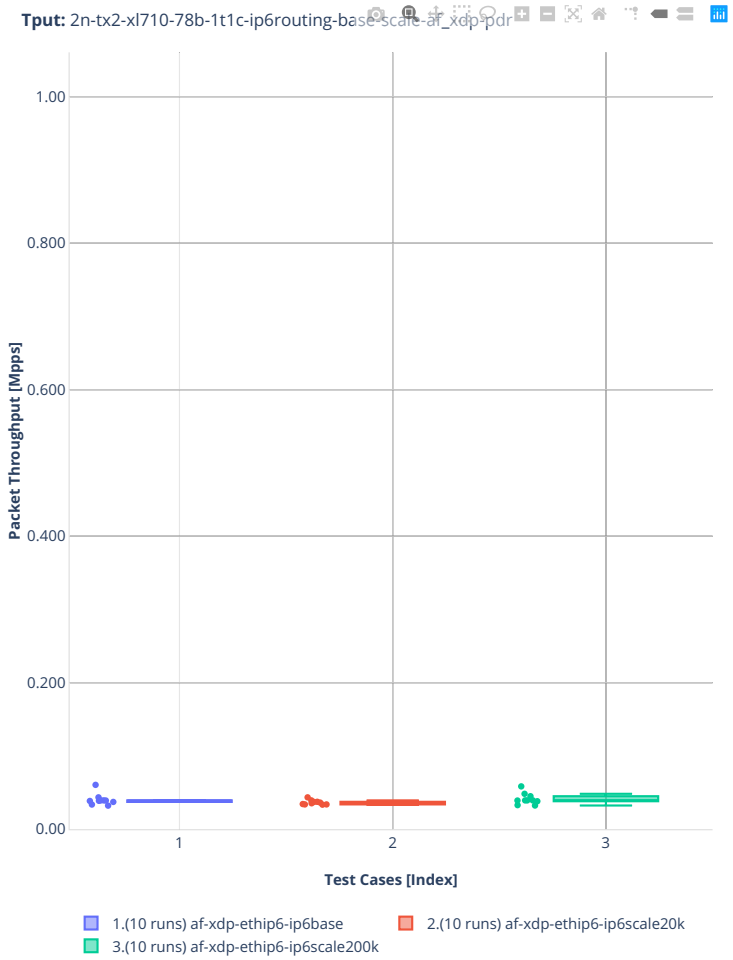
Tput: 2n-tx2-xl710-78b-1t1c-ip6routing-base-scale-af_xdp-pdr
1.00
0.800
0.600
0.400
0.200
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) af-xdp-ethip6-ip6base
2.(10 runs) af-xdp-ethip6-ip6scale20k
3.(10 runs) af-xdp-ethip6-ip6scale200k

2n-dnv-x553

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

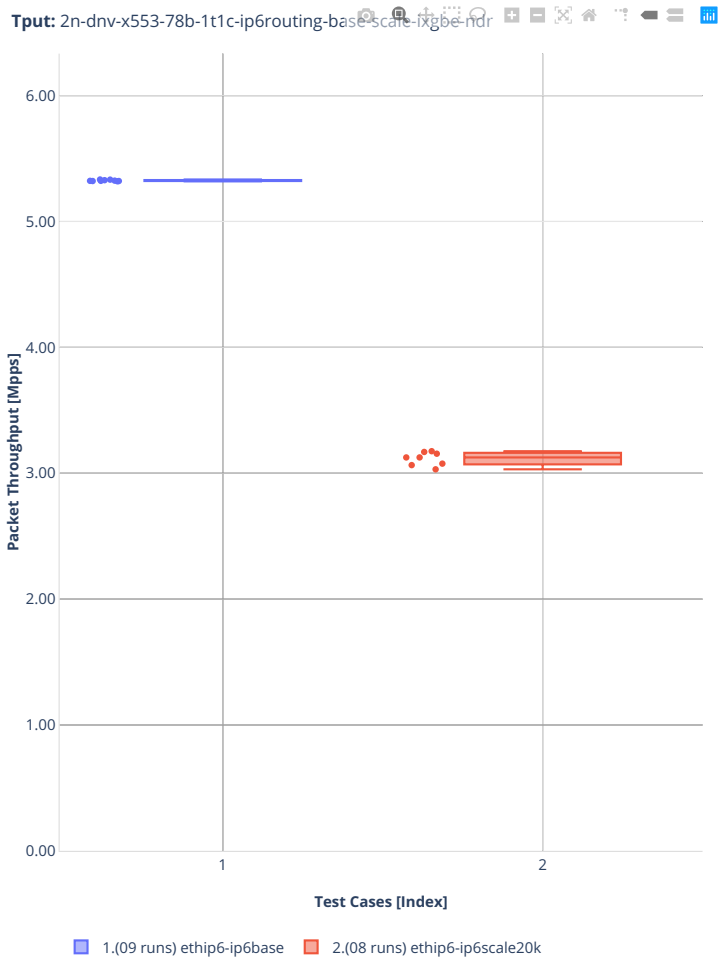

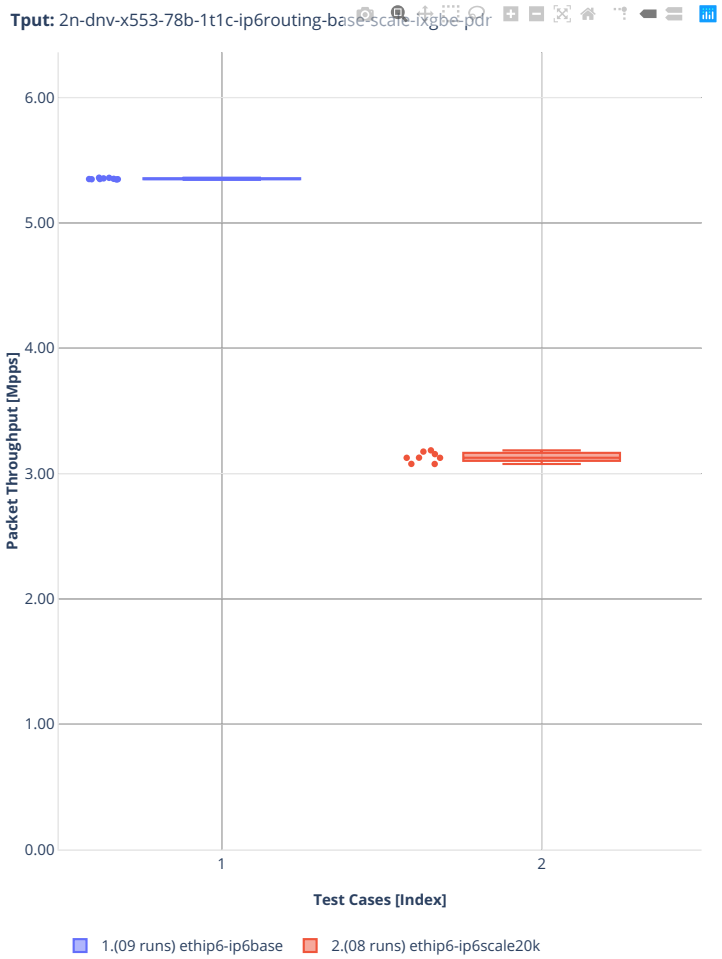
Tput: 2n-dnv-x553-78b-1t1c-ip6routing-base-scale-ixgbe-pdr
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(09 runs) ethip6-ip6base
2.(08 runs) ethip6-ip6scale20k

3n-dnv-x553

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

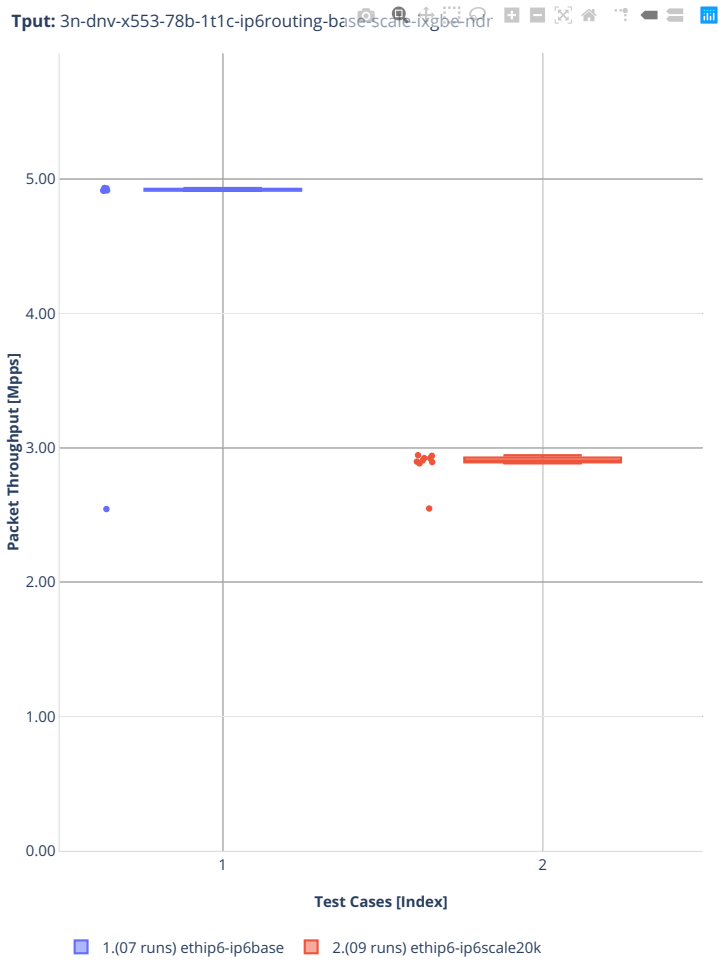

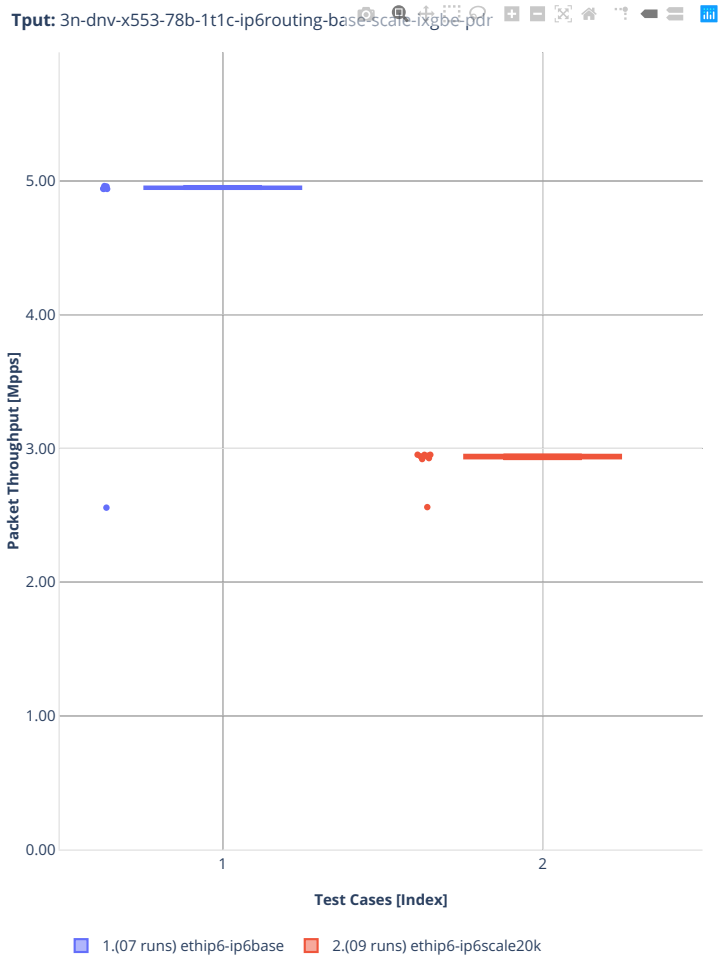
Tput: 3n-dnv-x553-78b-1t1c-ip6routing-base-scale-ixgbe-pdr
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(07 runs) ethip6-ip6base
2.(09 runs) ethip6-ip6scale20k

2n-aws-nitro50g

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

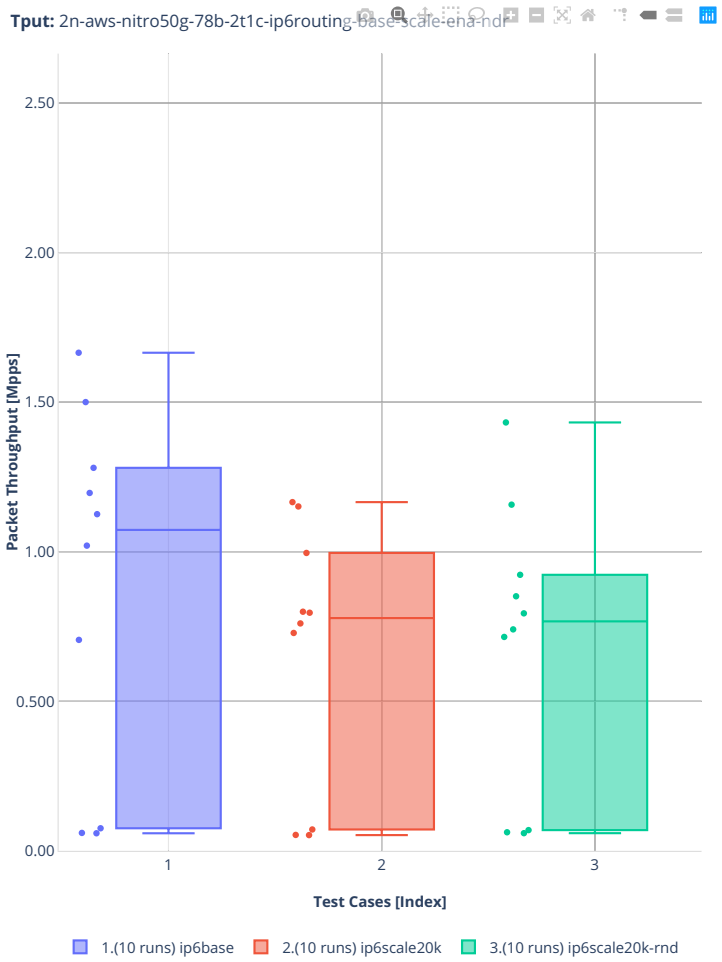

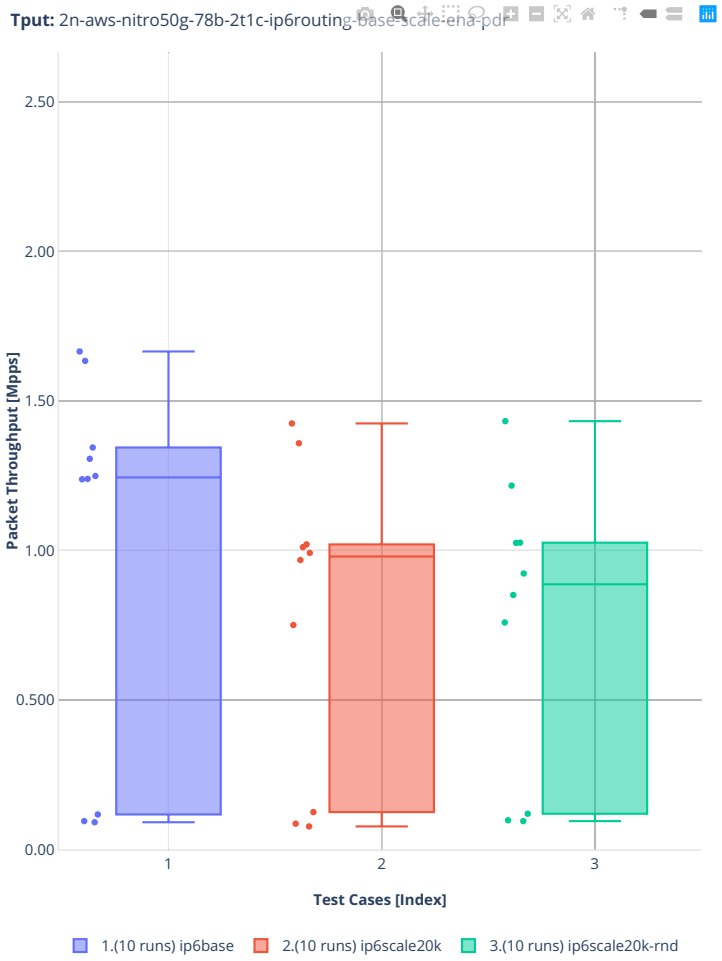
Tput: 2n-aws-nitro50g-78b-2t1c-ip6routing-base-scale-ena-pdr
2.50
2.00
1.50
1.00
0.500
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ip6base
2.(10 runs) ip6scale20k
3.(10 runs) ip6scale20k-rnd

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

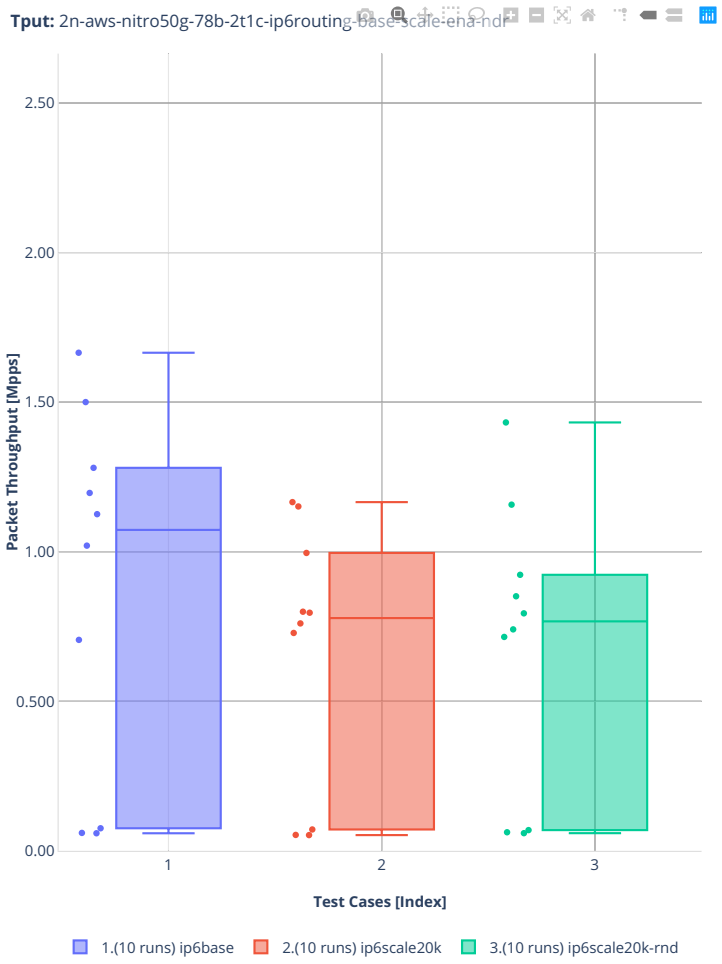

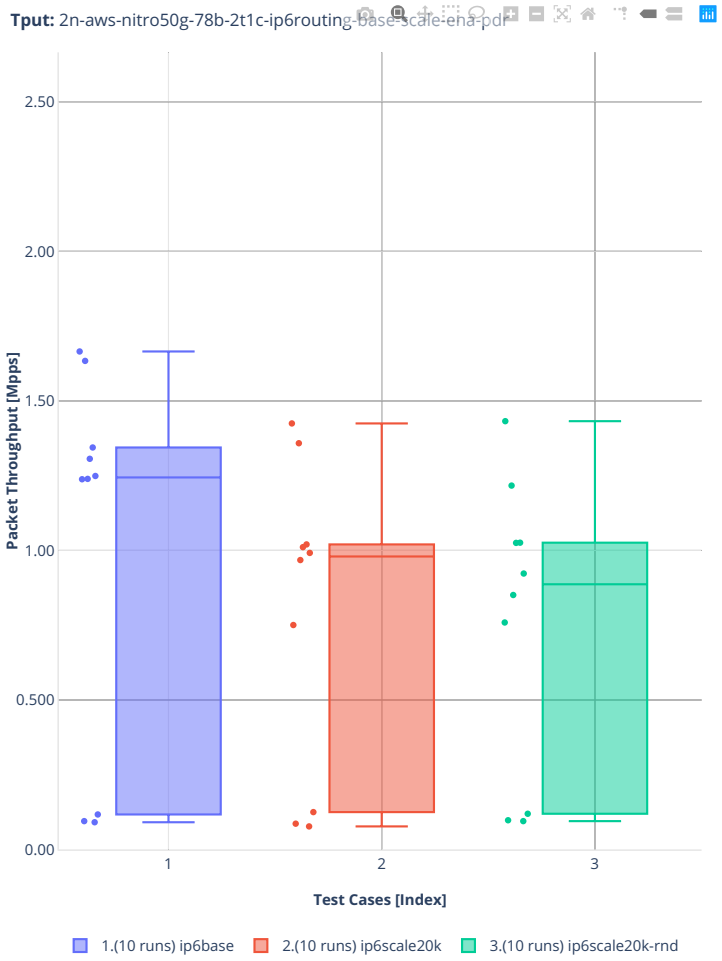
Tput: 2n-aws-nitro50g-78b-2t1c-ip6routing-base-scale-ena-pdr
2.50
2.00
1.50
1.00
0.500
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ip6base
2.(10 runs) ip6scale20k
3.(10 runs) ip6scale20k-rnd

1518b-2t1c-ip6routing-base-scale-ena

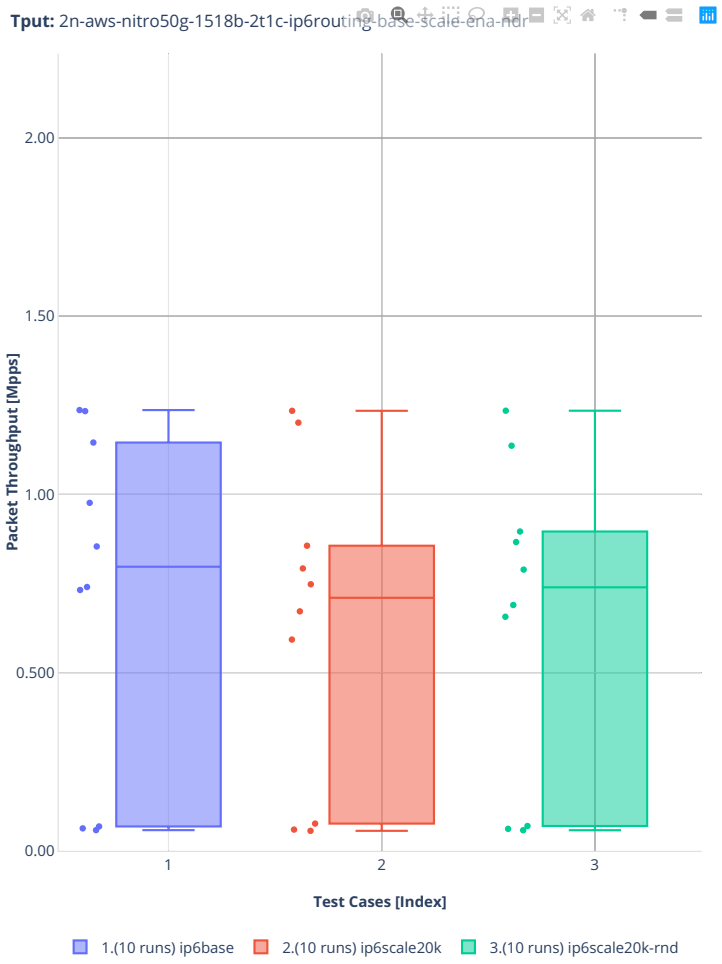

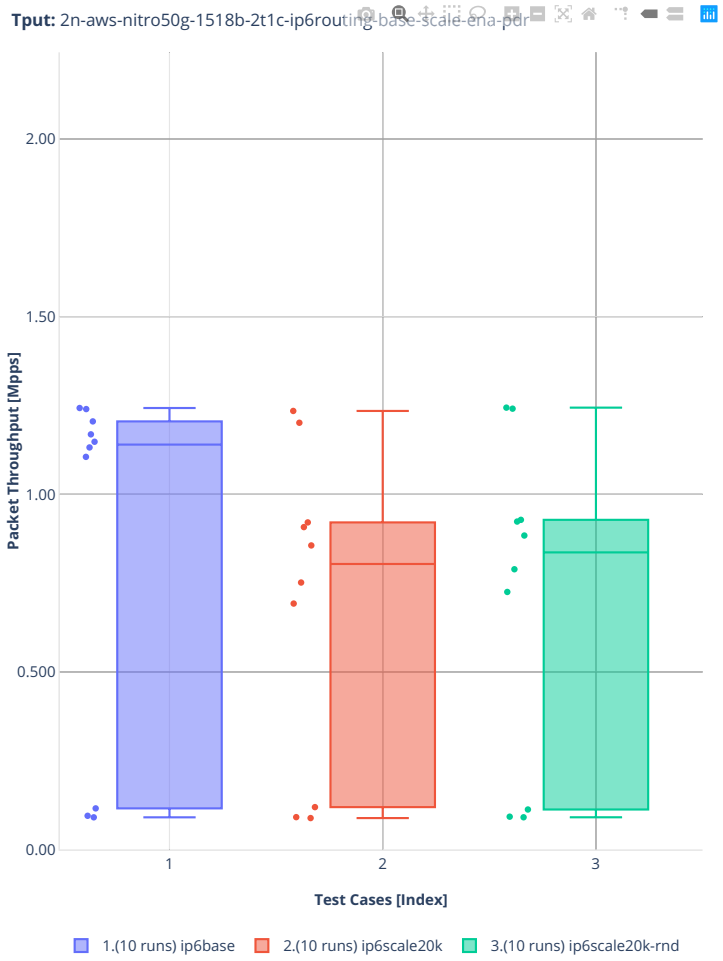
Tput: 2n-aws-nitro50g-1518b-2t1c-ip6routing-base-scale-ena-pdr
2.00
1.50
1.00
0.500
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ip6base
2.(10 runs) ip6scale20k
3.(10 runs) ip6scale20k-rnd

1518b-4t2c-ip6routing-base-scale-ena

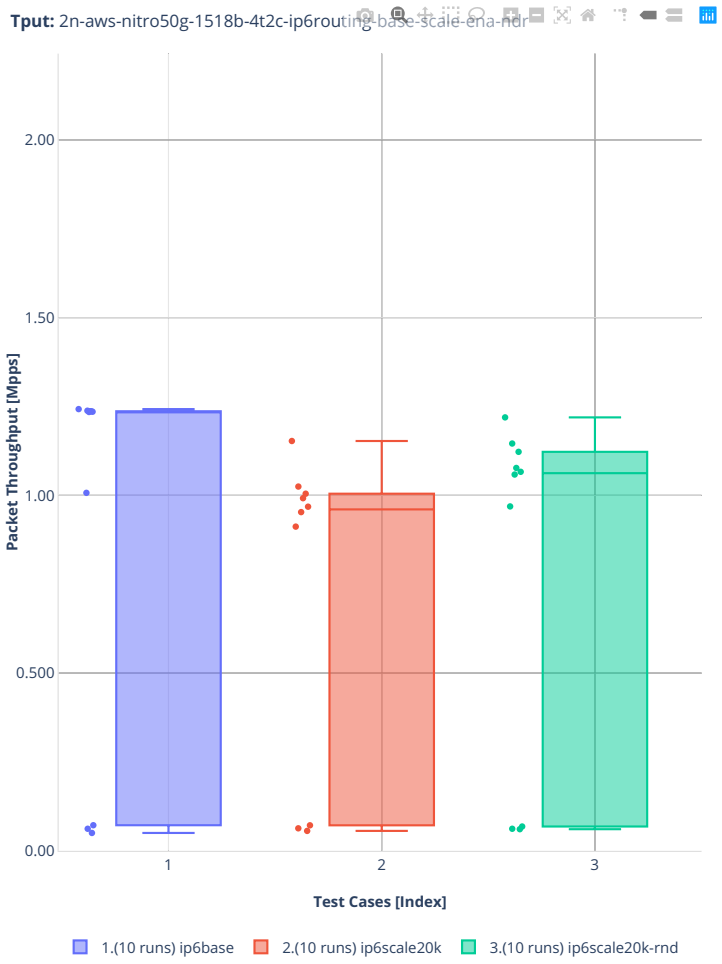

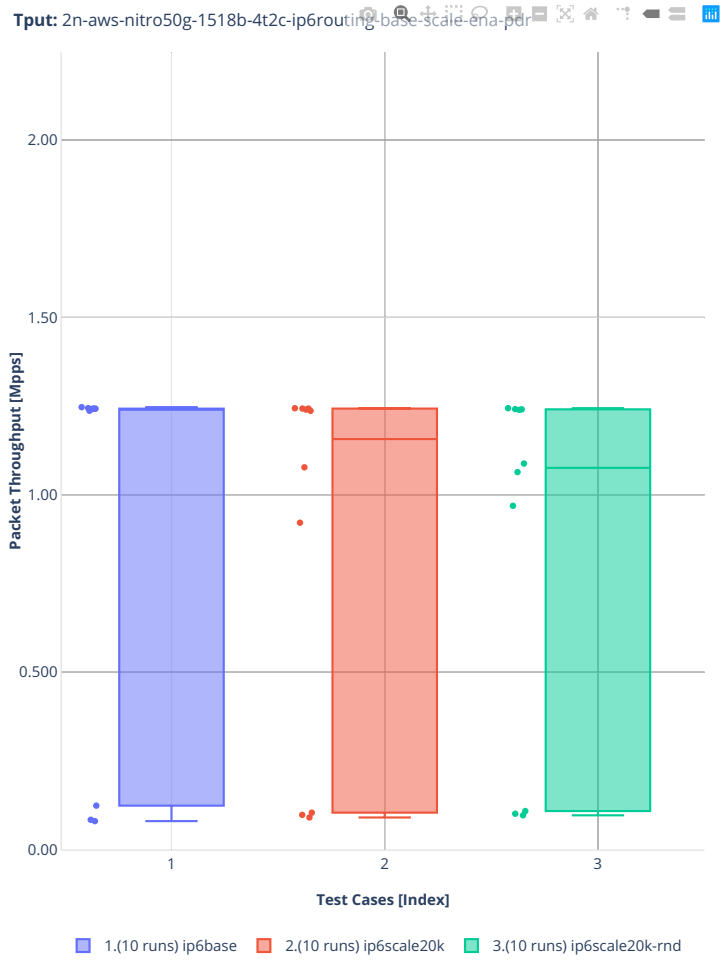
Tput: 2n-aws-nitro50g-1518b-4t2c-ip6routing-base-scale-ena-pdr
2.00
1.50
1.00
0.500
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ip6base
2.(10 runs) ip6scale20k
3.(10 runs) ip6scale20k-rnd

2.3.4 SRv6 Routing

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

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

[112] https://git.fd.io/csit/tree/tests/vpp/perf/srv6?h=rls2206

3n-icx-xxv710

78b-2t1c-srv6-ip6routing-base-avf

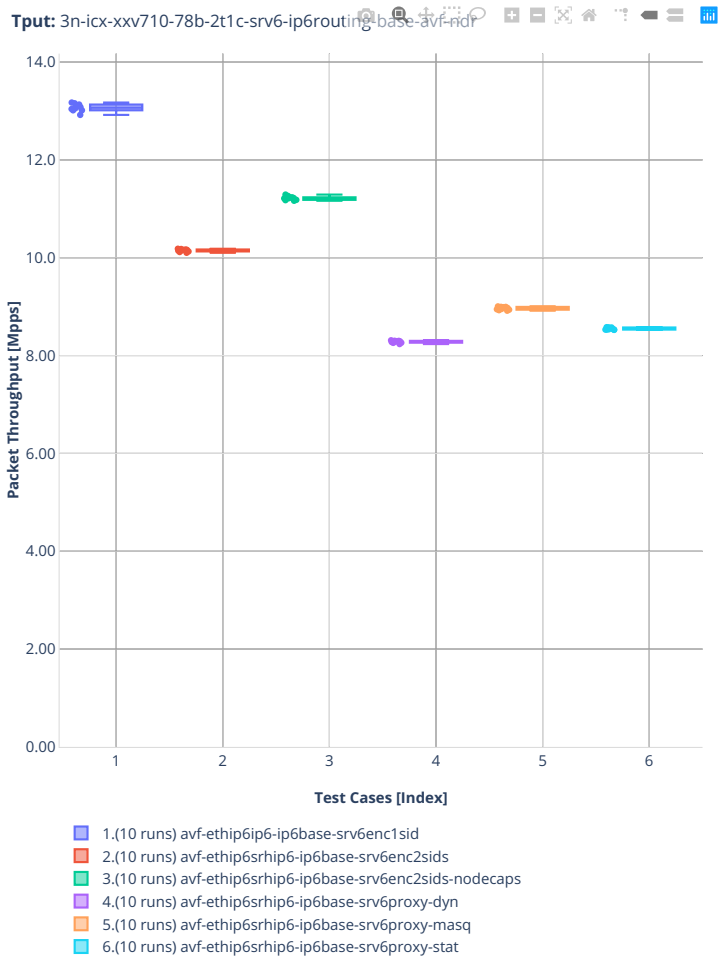

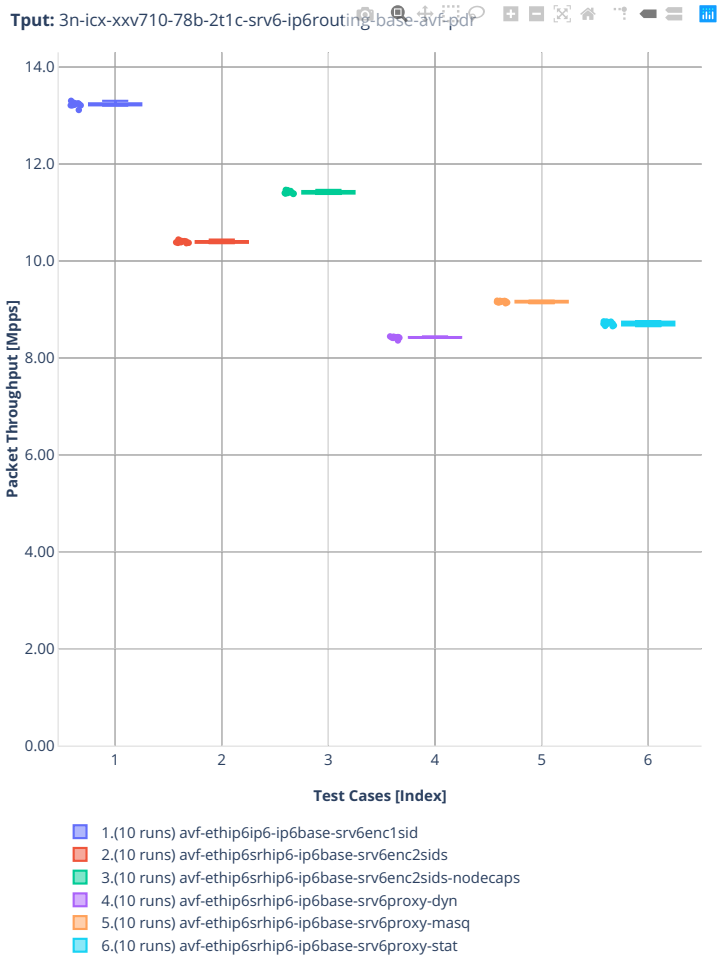
Tput: 3n-icx-xxv710-78b-2t1c-srv6-ip6routing-base-avf-pdr
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-ethip6ip6-ip6base-srv6enc1sid
2.(10 runs) avf-ethip6srhip6-ip6base-srv6enc2sids
3.(10 runs) avf-ethip6srhip6-ip6base-srv6enc2sids-nodecaps
4.(10 runs) avf-ethip6srhip6-ip6base-srv6proxy-dyn
5.(10 runs) avf-ethip6srhip6-ip6base-srv6proxy-masq
6.(10 runs) avf-ethip6srhip6-ip6base-srv6proxy-stat

3n-skx-xxv710

78b-2t1c-srv6-ip6routing-base-avf

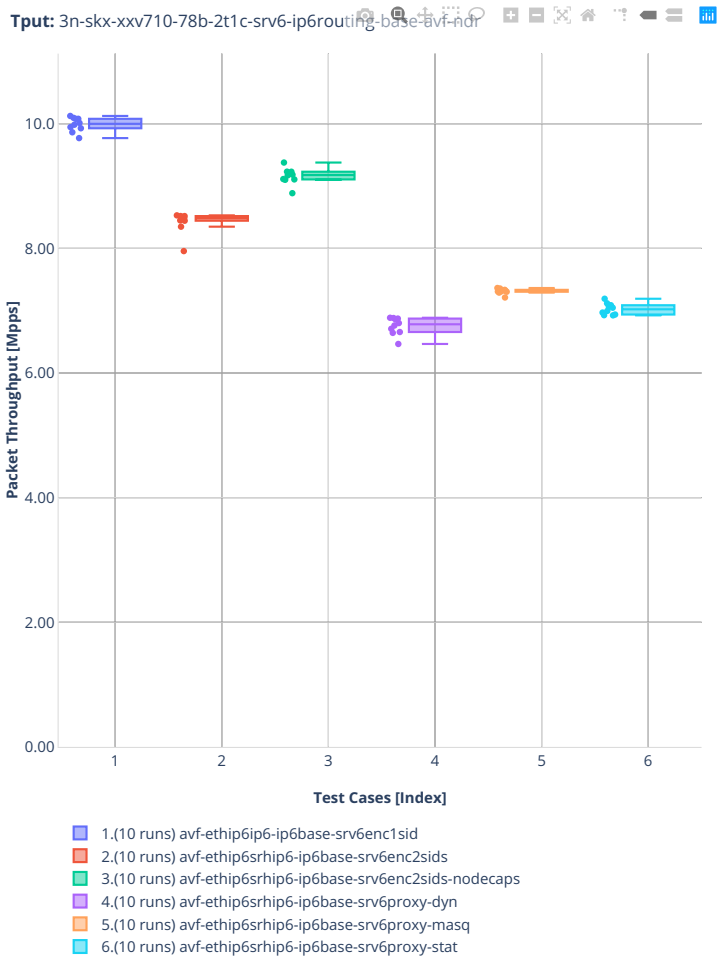

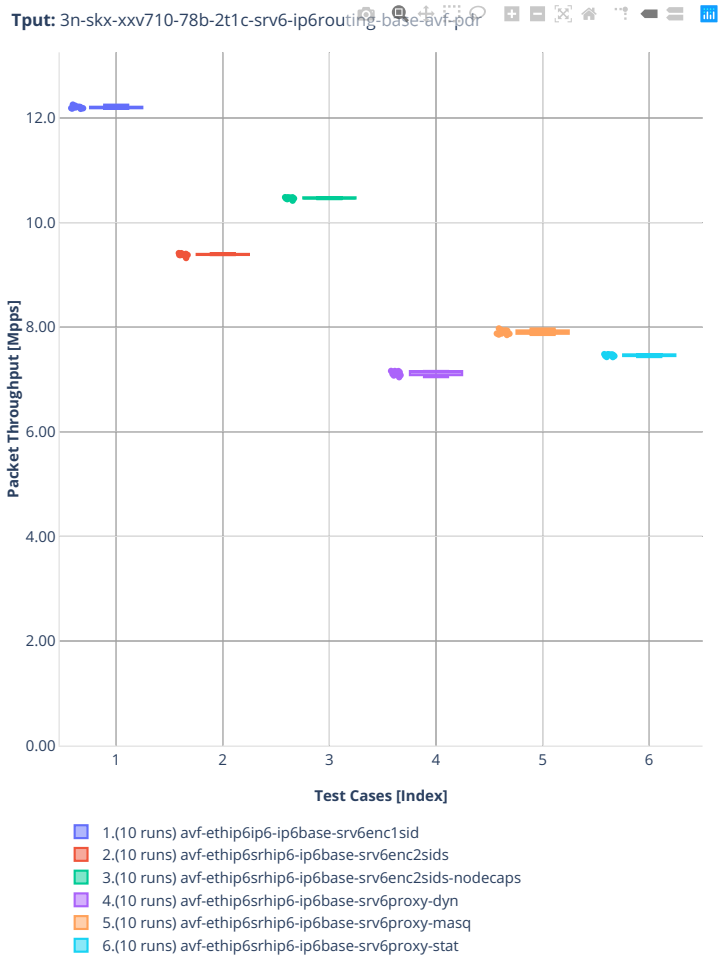
Tput: 3n-skx-xxv710-78b-2t1c-srv6-ip6routing-base-avf-pdr
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-ethip6ip6-ip6base-srv6enc1sid
2.(10 runs) avf-ethip6srhip6-ip6base-srv6enc2sids
3.(10 runs) avf-ethip6srhip6-ip6base-srv6enc2sids-nodecaps
4.(10 runs) avf-ethip6srhip6-ip6base-srv6proxy-dyn
5.(10 runs) avf-ethip6srhip6-ip6base-srv6proxy-masq
6.(10 runs) avf-ethip6srhip6-ip6base-srv6proxy-stat

3n-tsh-x520

78b-1t1c-srv6-ip6routing-base-ixgbe

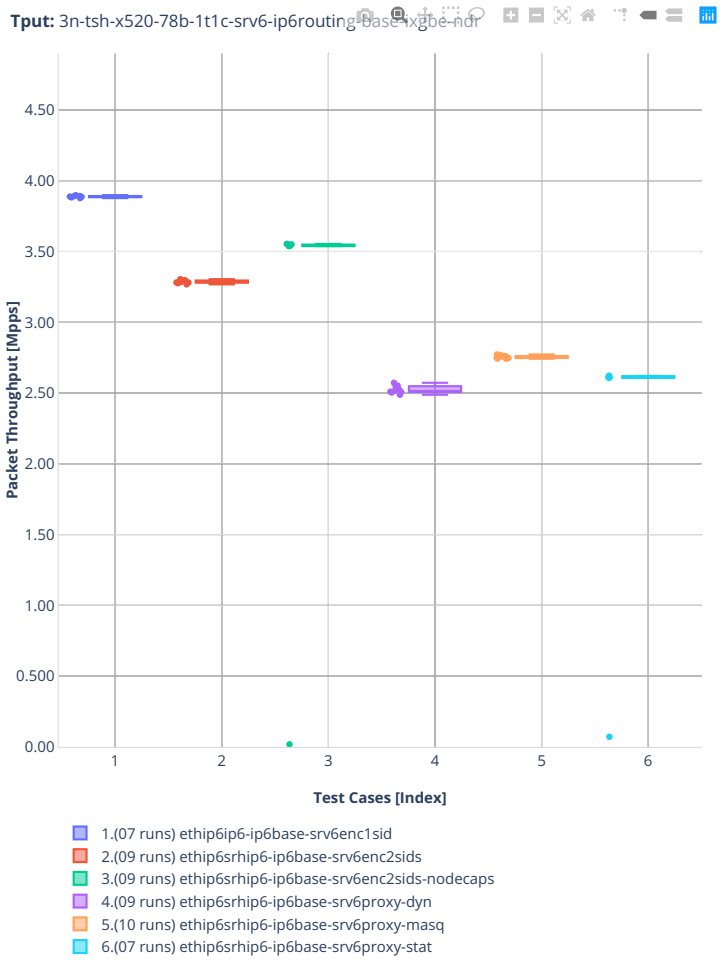

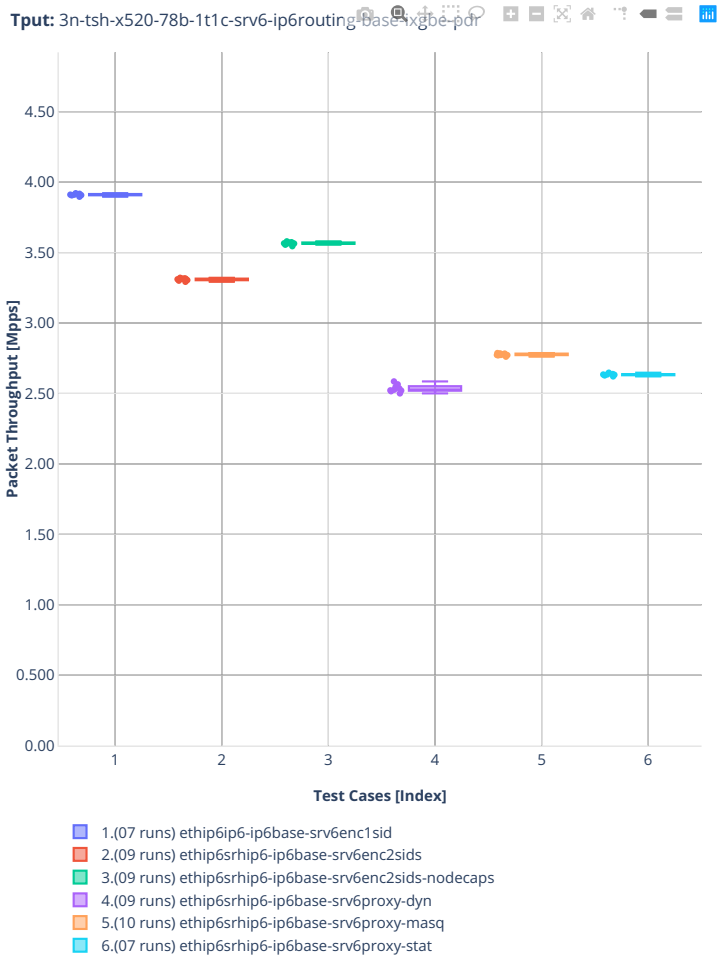
Tput: 3n-tsh-x520-78b-1t1c-srv6-ip6routing-base-ixgbe-pdr
Packet Throughput [Mpps]
4.50
4.00
3.50
3.00
2.50
2.00
1.50
1.00
0.500
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(07 runs) ethip6ip6-ip6base-srv6enc1sid
2.(09 runs) ethip6srhip6-ip6base-srv6enc2sids
3.(09 runs) ethip6srhip6-ip6base-srv6enc2sids-nodecaps
4.(09 runs) ethip6srhip6-ip6base-srv6proxy-dyn
5.(10 runs) ethip6srhip6-ip6base-srv6proxy-masq
6.(07 runs) ethip6srhip6-ip6base-srv6proxy-stat

2.3.5 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[113].

[113] https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls2206

2n-icx-xxv710

64b-2t1c-ethip4–ethip4udpgeneve-avf

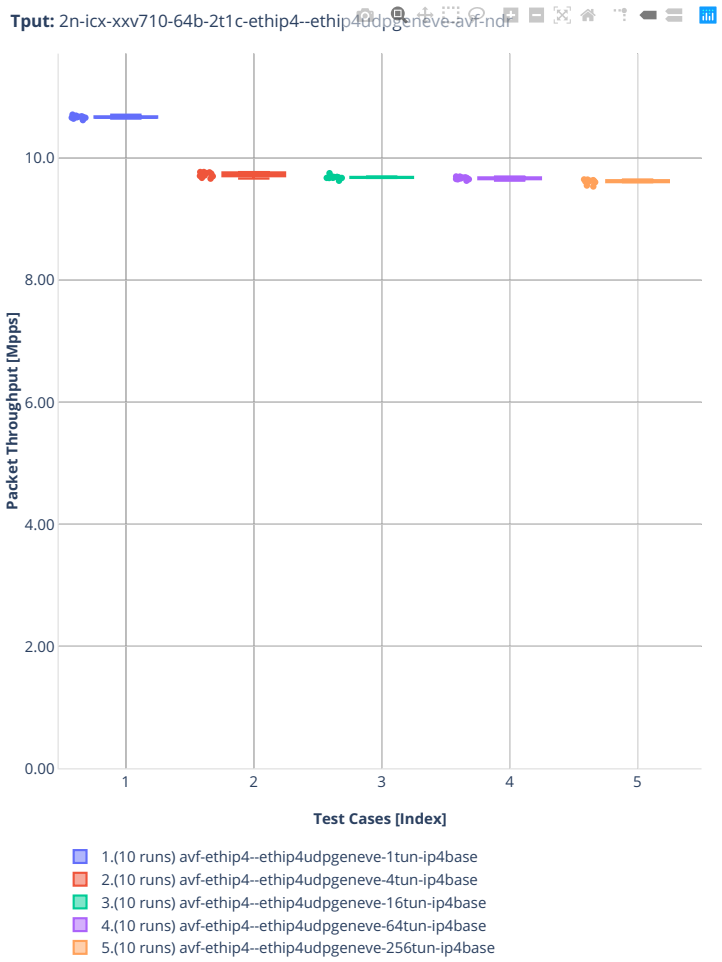

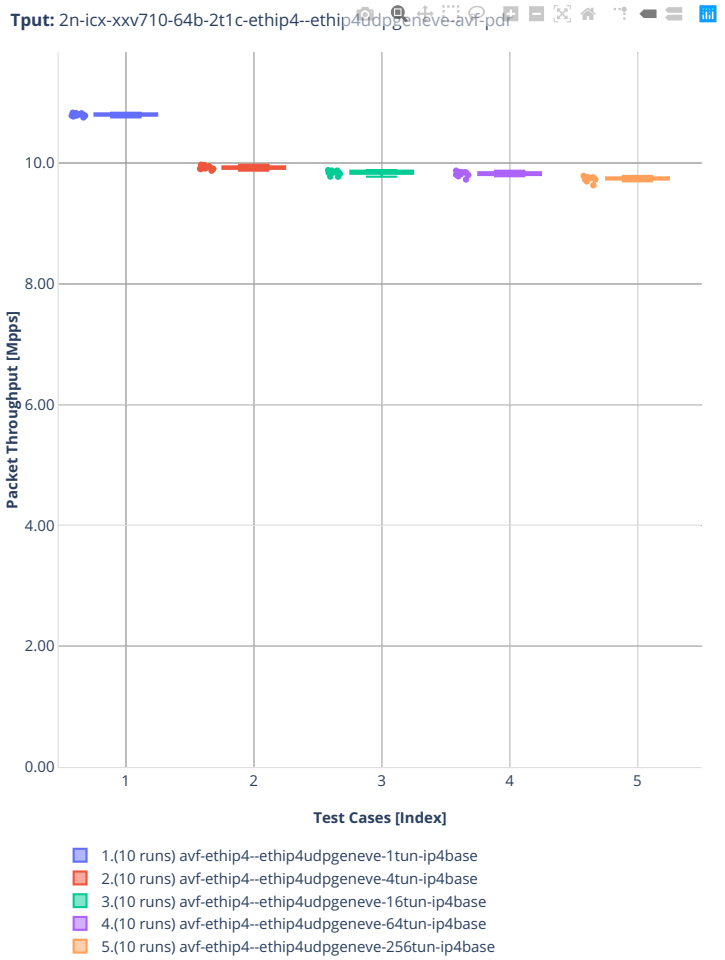
Tput: 2n-icx-xxv710-64b-2t1c-ethip4--ethip4udpgeneve-avf-pdr
Packet Throughput [Mpps]
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
5
Test Cases [Index]
1.(10 runs) avf-ethip4--ethip4udpgeneve-1tun-ip4base
2.(10 runs) avf-ethip4--ethip4udpgeneve-4tun-ip4base
3.(10 runs) avf-ethip4--ethip4udpgeneve-16tun-ip4base
4.(10 runs) avf-ethip4--ethip4udpgeneve-64tun-ip4base
5.(10 runs) avf-ethip4--ethip4udpgeneve-256tun-ip4base

3n-icx-xxv710

64b-2t1c-ip4tunnel-base

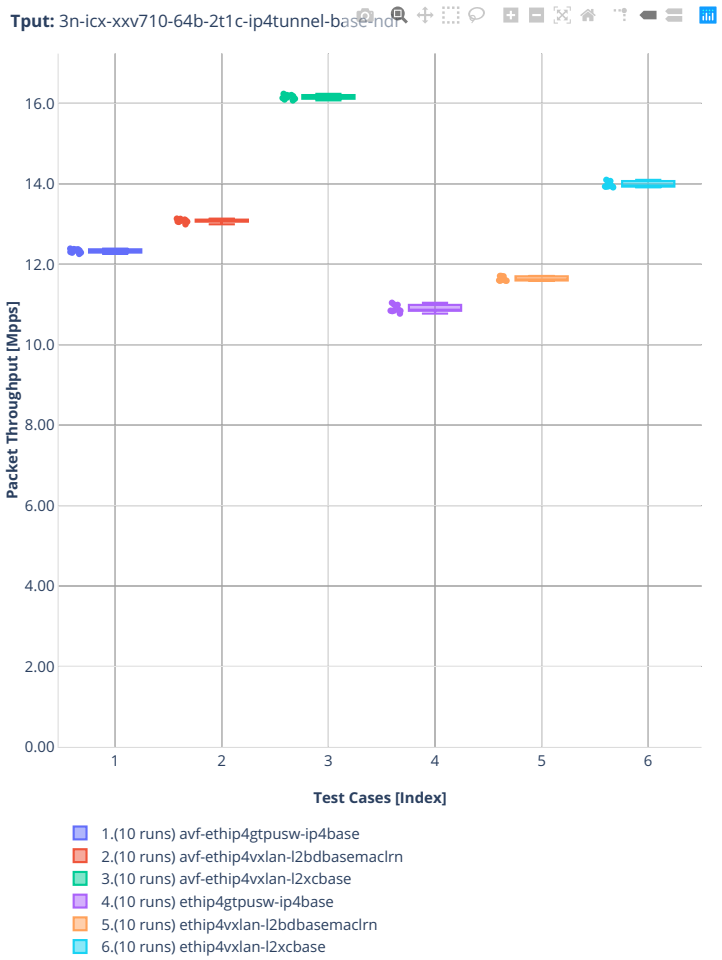

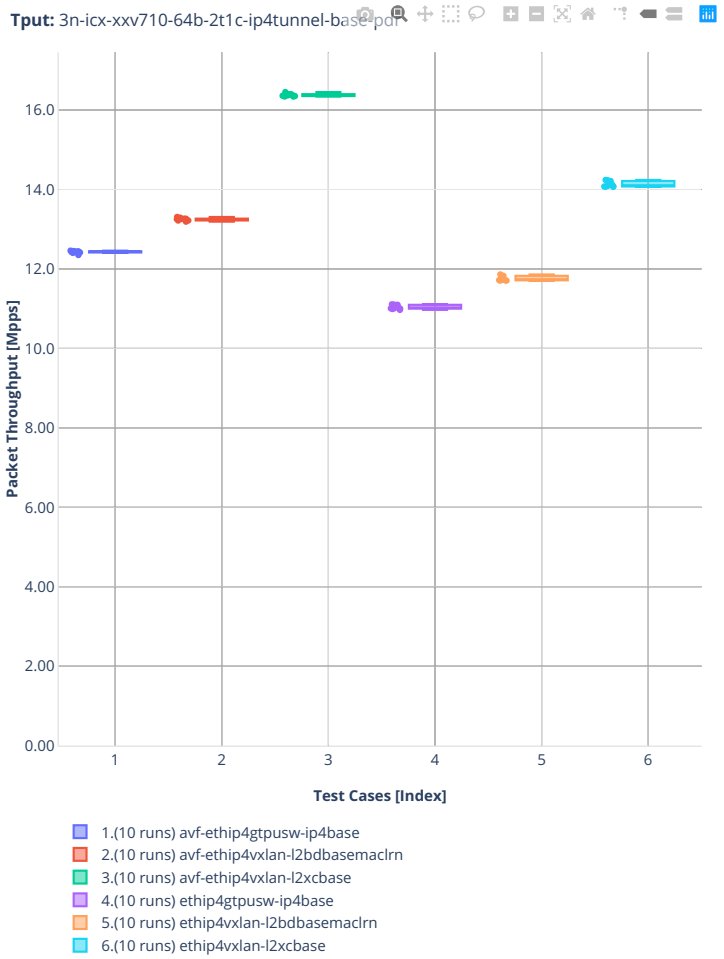
Tput: 3n-icx-xxv710-64b-2t1c-ip4tunnel-base.pdf
Packet Throughput [Mpps]
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-ethip4gtpusw-ip4base
2.(10 runs) avf-ethip4vxlan-l2bdbasemaclrn
3.(10 runs) avf-ethip4vxlan-l2xcbase
4.(10 runs) ethip4gtpusw-ip4base
5.(10 runs) ethip4vxlan-l2bdbasemaclrn
6.(10 runs) ethip4vxlan-l2xcbase

3n-icx-e810xxv

64b-2t1c-ip4tunnel-gtpu

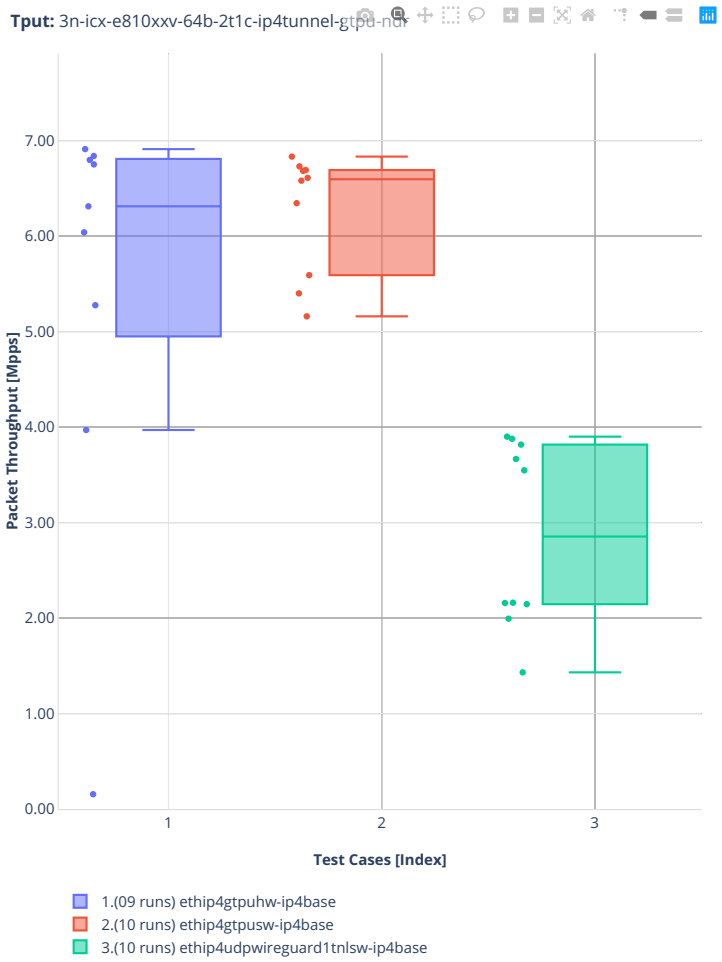

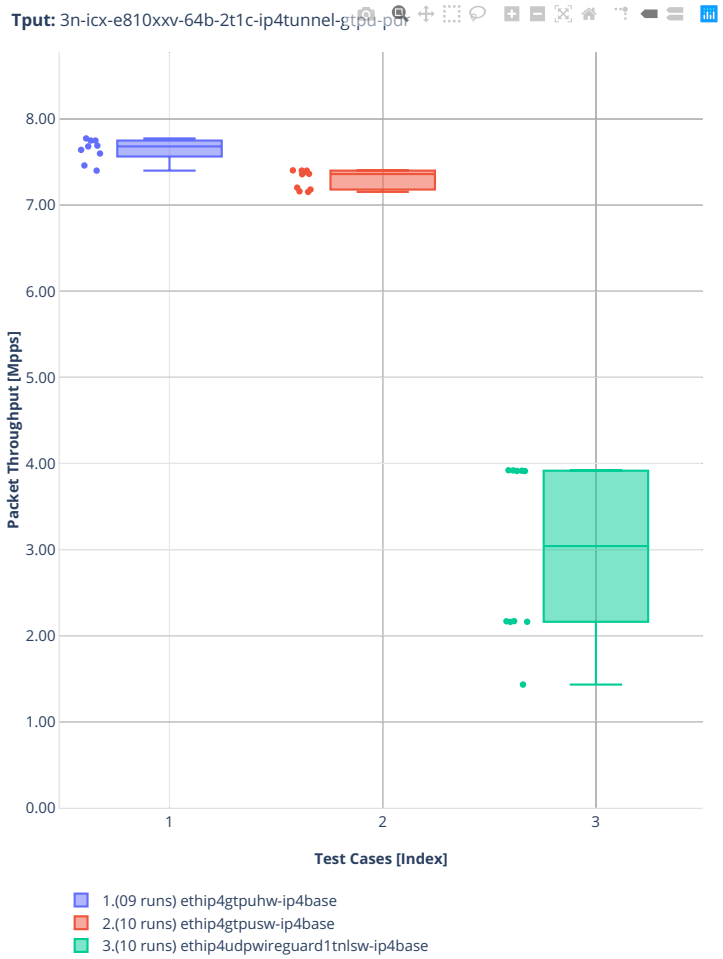
Tput: 3n-icx-e810xxv-64b-2t1c-ip4tunnel-gtpu-pdr
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(09 runs) ethip4gtpuhw-ip4base
2.(10 runs) ethip4gtpusw-ip4base
3.(10 runs) ethip4udpwireguard1tnlsw-ip4base

2n-skx-xxv710

64b-2t1c-ethip4–ethip4udpgeneve-avf

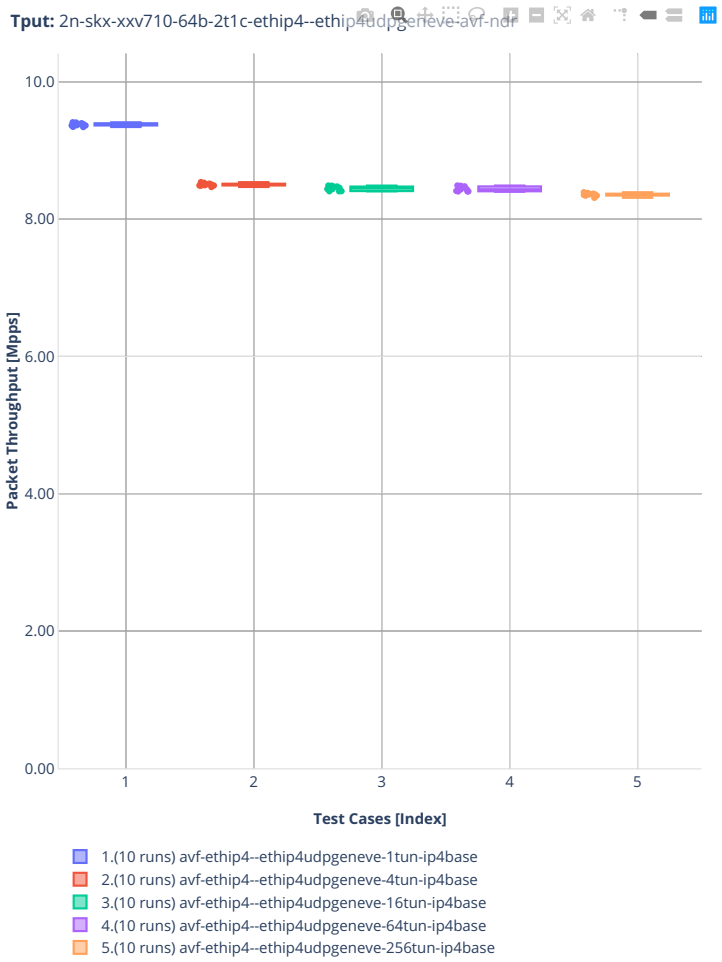

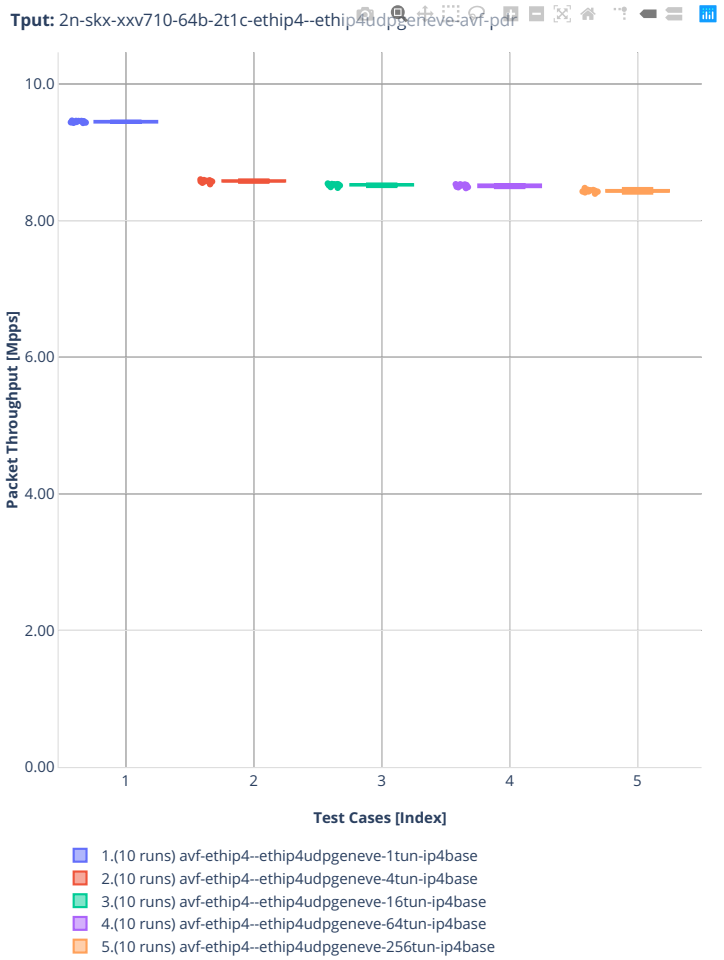
Tput: 2n-skx-xxv710-64b-2t1c-ethip4--ethip4udpgeneve-avf-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
Test Cases [Index]
1.(10 runs) avf-ethip4--ethip4udpgeneve-1tun-ip4base
2.(10 runs) avf-ethip4--ethip4udpgeneve-4tun-ip4base
3.(10 runs) avf-ethip4--ethip4udpgeneve-16tun-ip4base
4.(10 runs) avf-ethip4--ethip4udpgeneve-64tun-ip4base
5.(10 runs) avf-ethip4--ethip4udpgeneve-256tun-ip4base

2n-clx-xxv710

64b-2t1c-ethip4–ethip4udpgeneve-avf

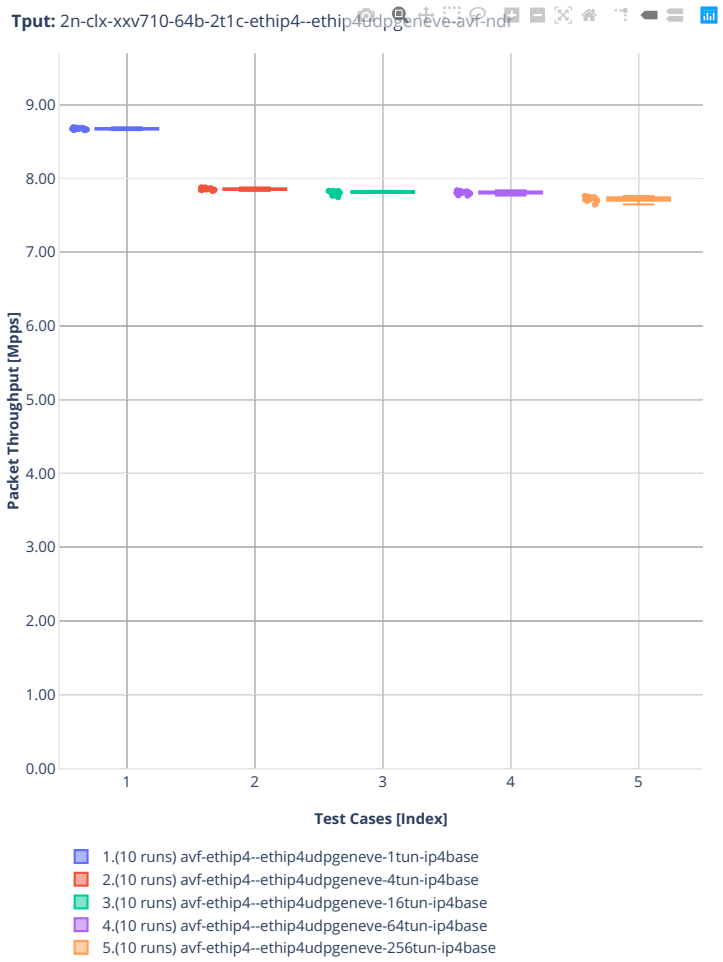

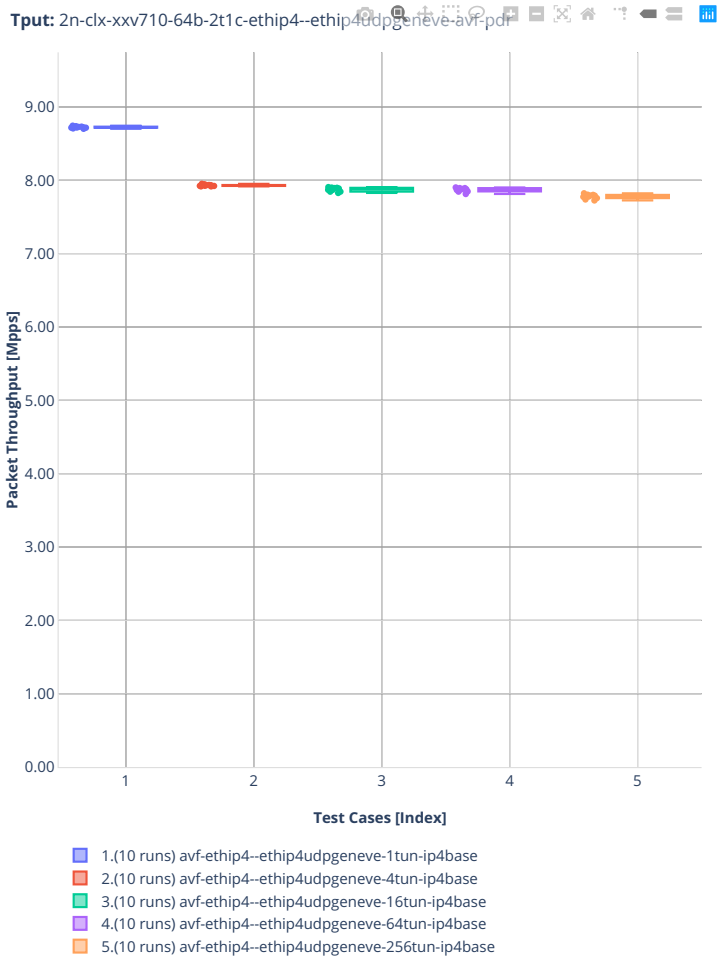
Tput: 2n-clx-xxv710-64b-2t1c-ethip4--ethip4udpgeneve-avf-pdr
9.00
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
Test Cases [Index]
1.(10 runs) avf-ethip4--ethip4udpgeneve-1tun-ip4base
2.(10 runs) avf-ethip4--ethip4udpgeneve-4tun-ip4base
3.(10 runs) avf-ethip4--ethip4udpgeneve-16tun-ip4base
4.(10 runs) avf-ethip4--ethip4udpgeneve-64tun-ip4base
5.(10 runs) avf-ethip4--ethip4udpgeneve-256tun-ip4base

2n-zn2-xxv710

64b-2t1c-ethip4–ethip4udpgeneve-avf

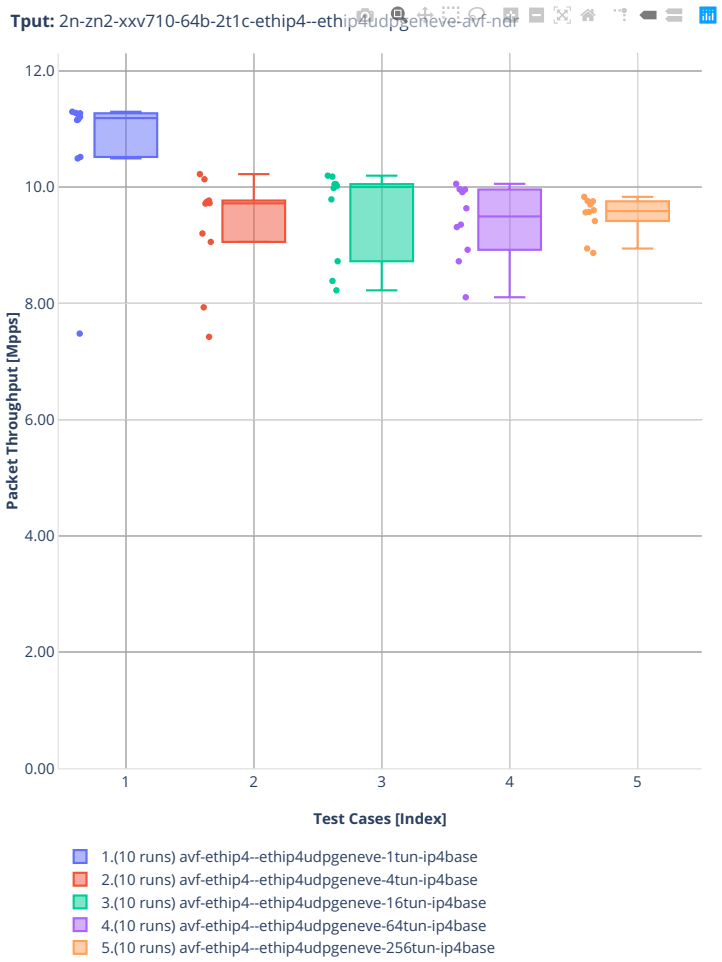

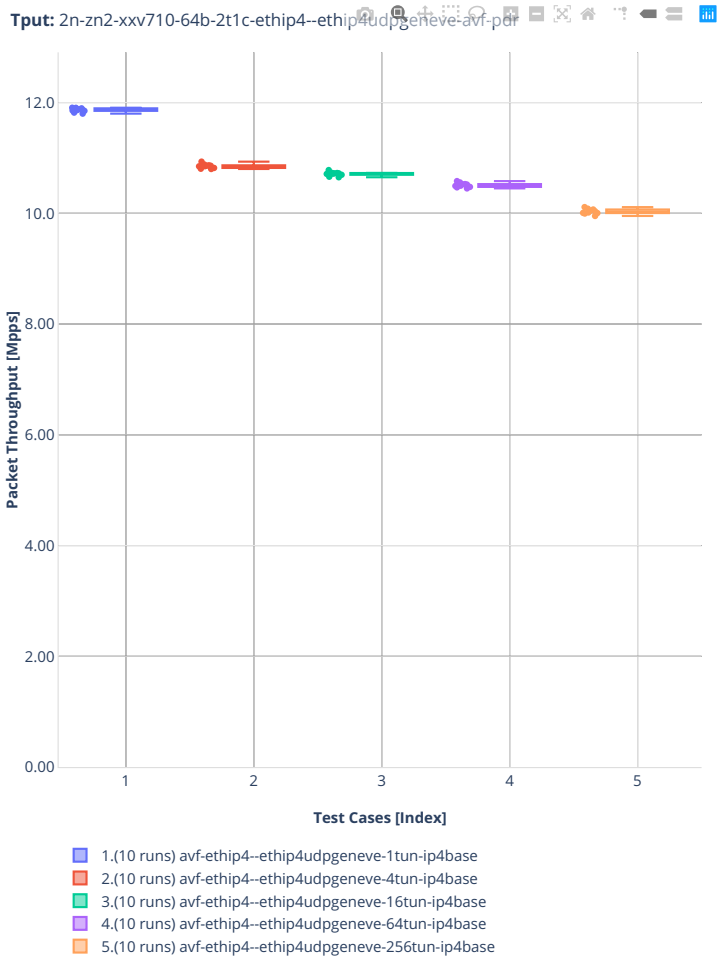
Tput: 2n-zn2-xxv710-64b-2t1c-ethip4--ethip4udpgeneve-avf-pdr
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
Test Cases [Index]
1.(10 runs) avf-ethip4--ethip4udpgeneve-1tun-ip4base
2.(10 runs) avf-ethip4--ethip4udpgeneve-4tun-ip4base
3.(10 runs) avf-ethip4--ethip4udpgeneve-16tun-ip4base
4.(10 runs) avf-ethip4--ethip4udpgeneve-64tun-ip4base
5.(10 runs) avf-ethip4--ethip4udpgeneve-256tun-ip4base

3n-skx-xxv710

64b-2t1c-ip4tunnel-base

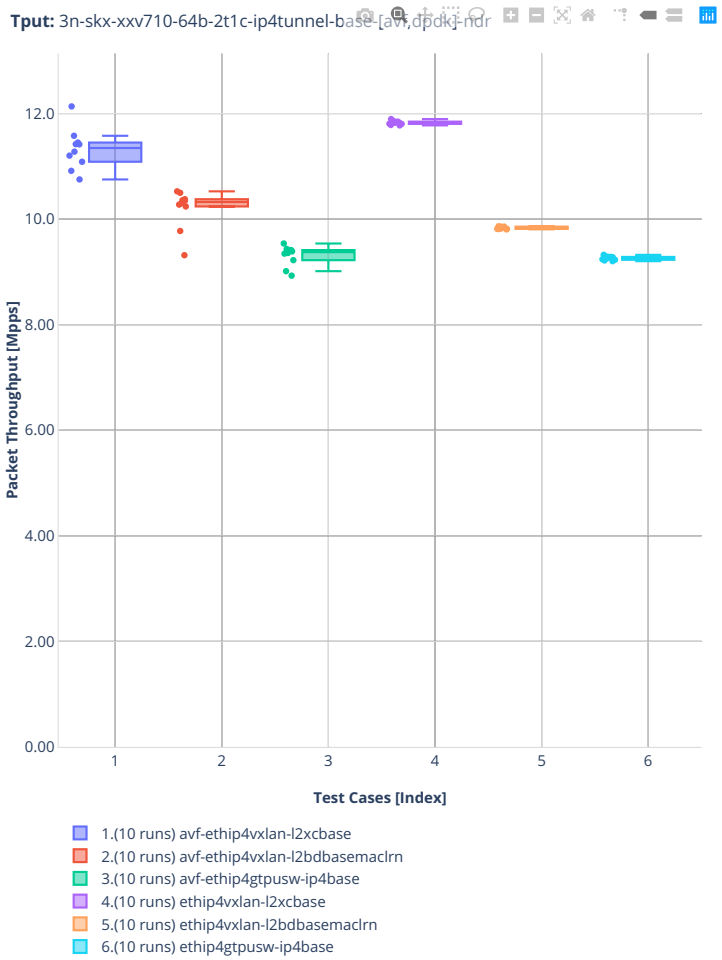

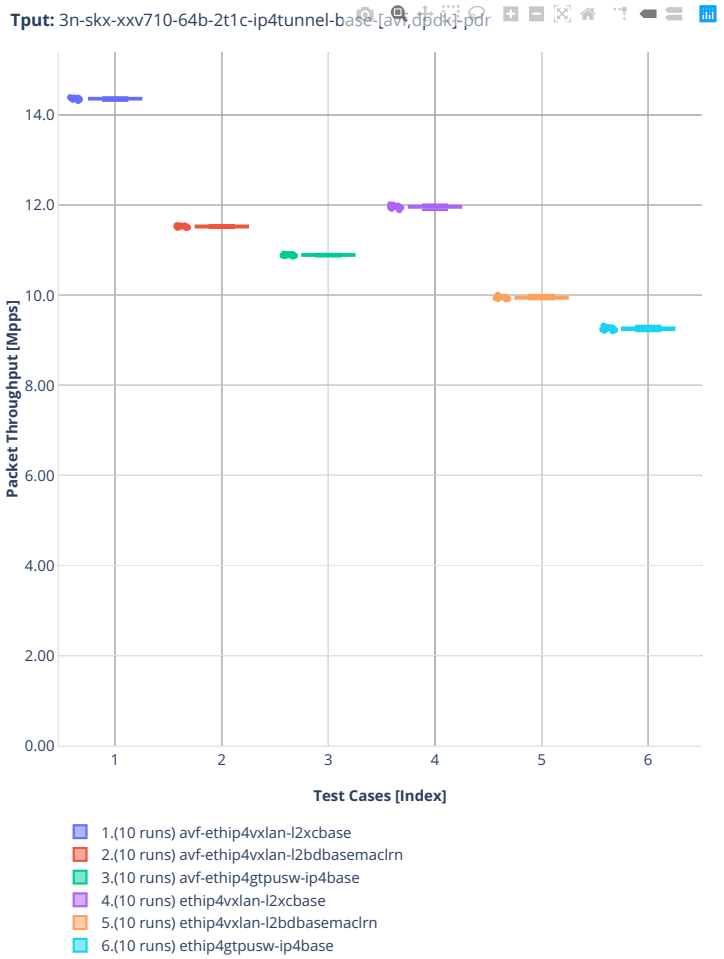
Tput: 3n-skx-xxv710-64b-2t1c-ip4tunnel-base-[avf,dpdk]-pdr
Packet Throughput [Mpps]
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-ethip4vxlan-l2xcbase
2.(10 runs) avf-ethip4vxlan-l2bdbasemaclrn
3.(10 runs) avf-ethip4gtpusw-ip4base
4.(10 runs) ethip4vxlan-l2xcbase
5.(10 runs) ethip4vxlan-l2bdbasemaclrn
6.(10 runs) ethip4gtpusw-ip4base

3n-alt-xl710

64b-1t1c-ip4tunnel-base

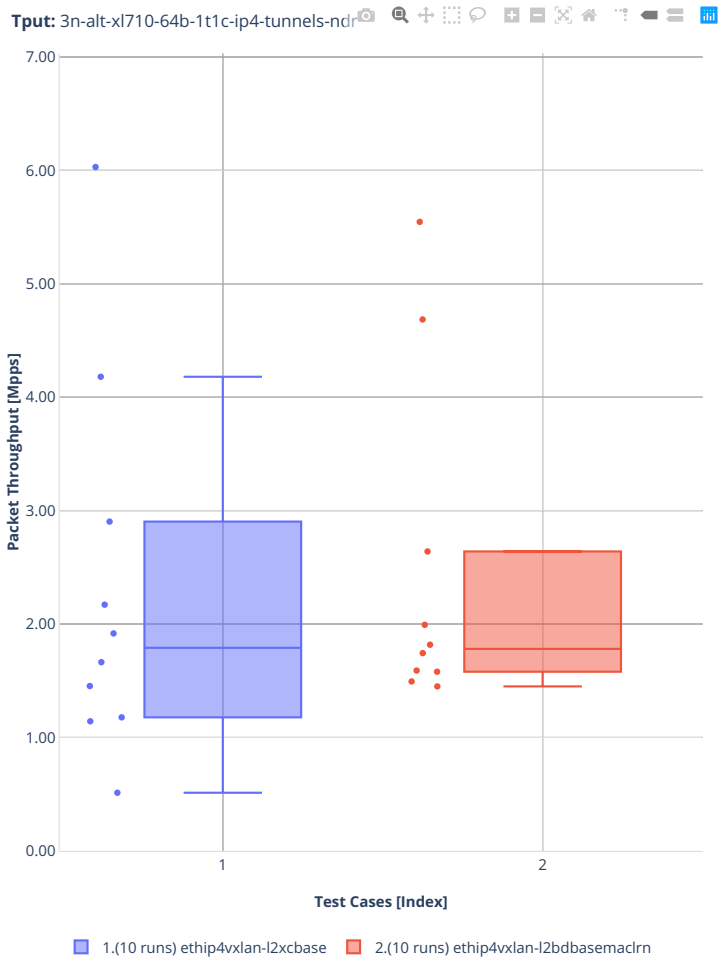

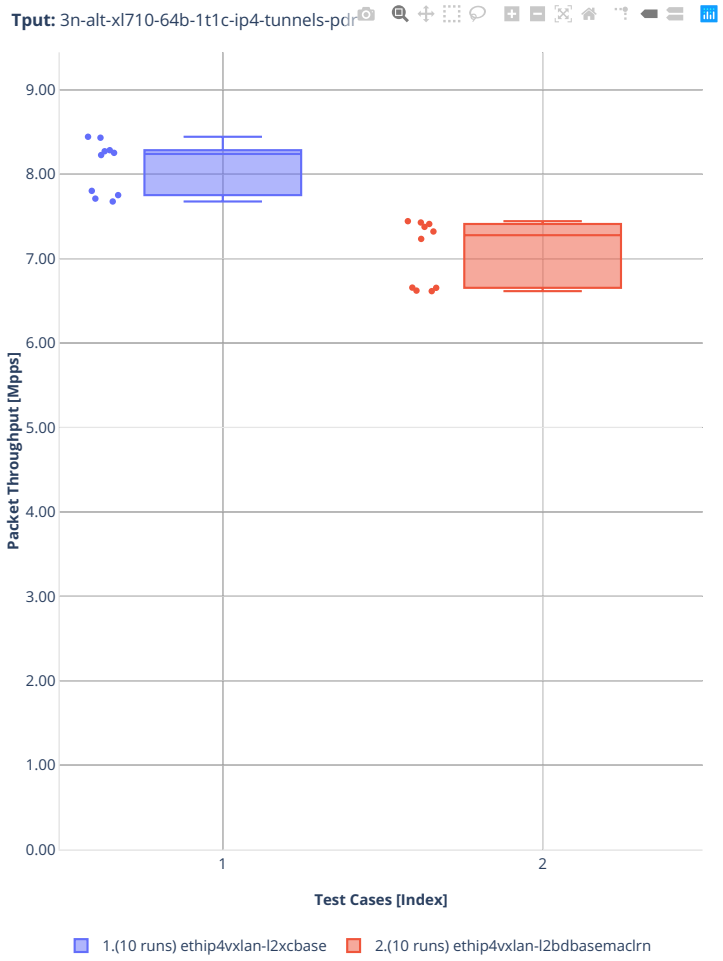
Tput: 3n-alt-xl710-64b-1t1c-ip4-tunnels-pdr
9.00
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) ethip4vxlan-l2xcbase
2.(10 runs) ethip4vxlan-l2bdbasemaclrn

3n-tsh-x520

64b-1t1c-ip4tunnel-base-ixgbe

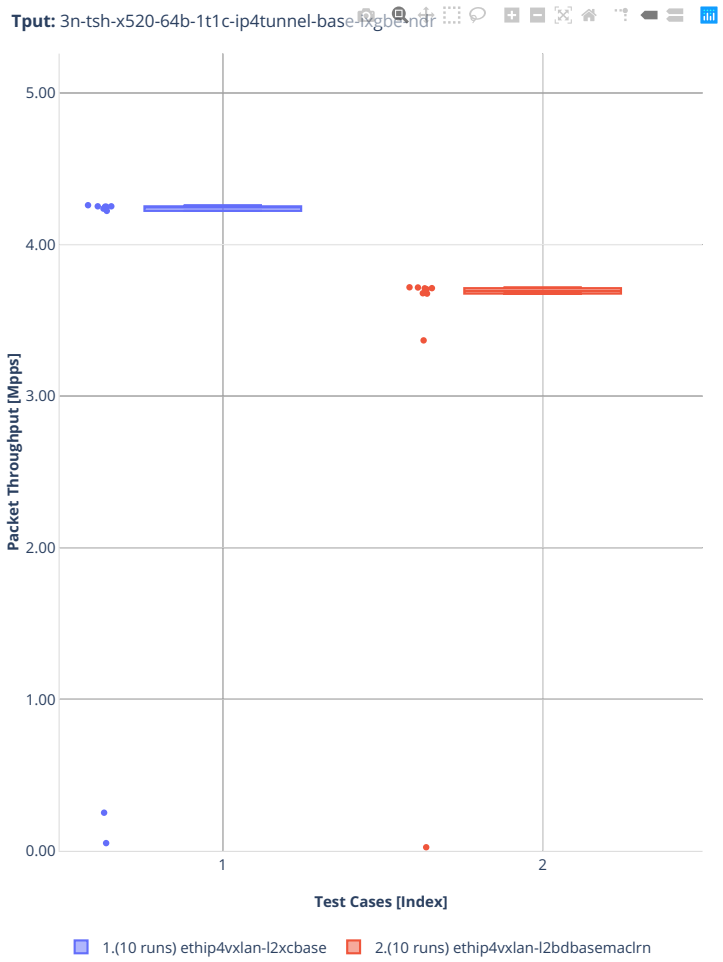

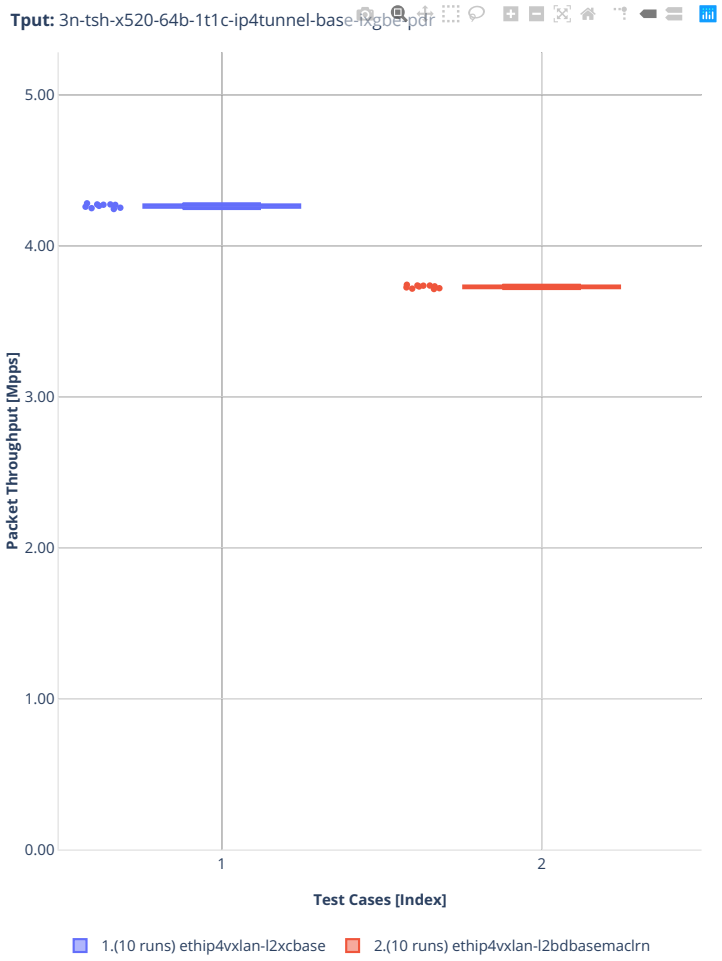
Tput: 3n-tsh-x520-64b-1t1c-ip4tunnel-base-1xgbe-pdr
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) ethip4vxlan-l2xcbase
2.(10 runs) ethip4vxlan-l2bdbasemaclrn

3n-dnv-x553

64b-1t1c-ip4tunnel-base-ixgbe

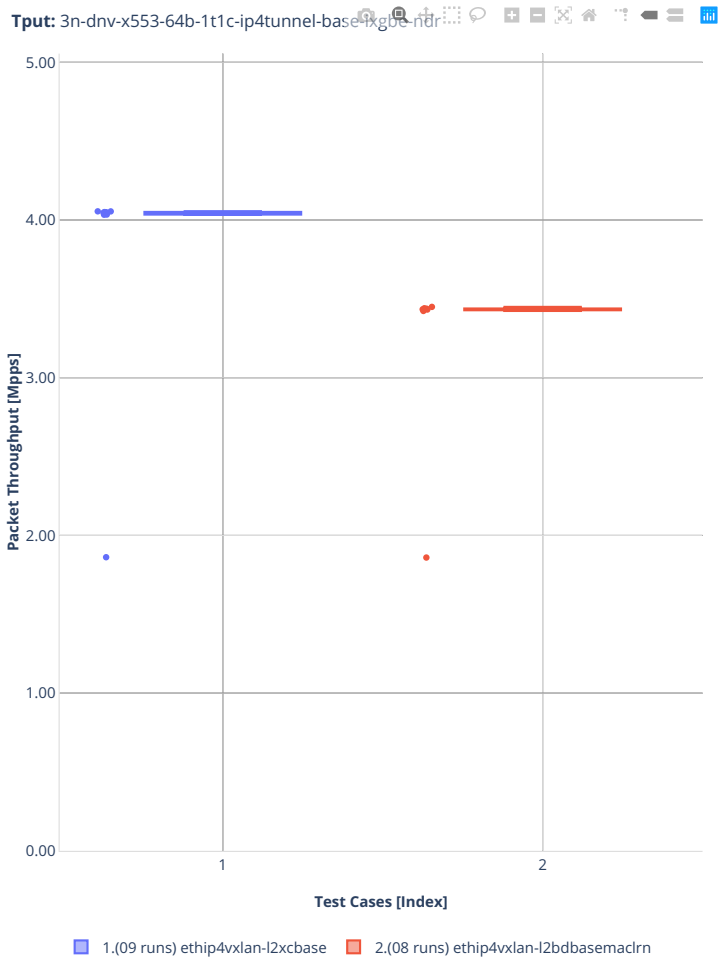

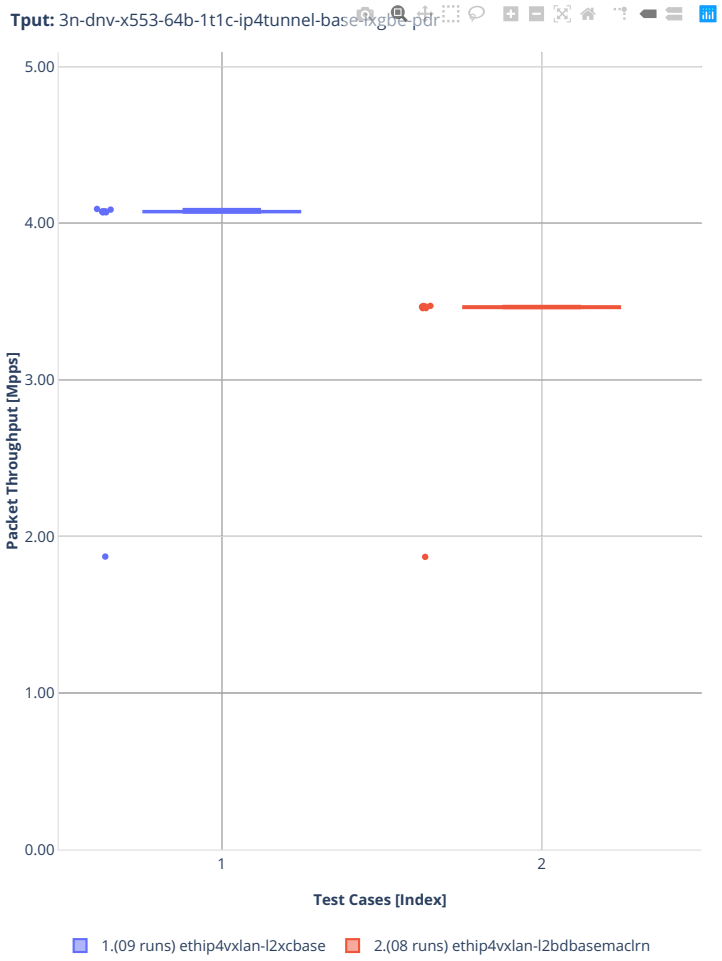
Tput: 3n-dnv-x553-64b-1t1c-ip4tunnel-base-ixgbe-pdr
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(09 runs) ethip4vxlan-l2xcbase
2.(08 runs) ethip4vxlan-l2bdbasemaclrn

2.3.6 NAT44 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[114].

[114] https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls2206

Det BiDir

2n-icx-xxv710

64b-nat44det-ip4routing-stl-bidir-avf

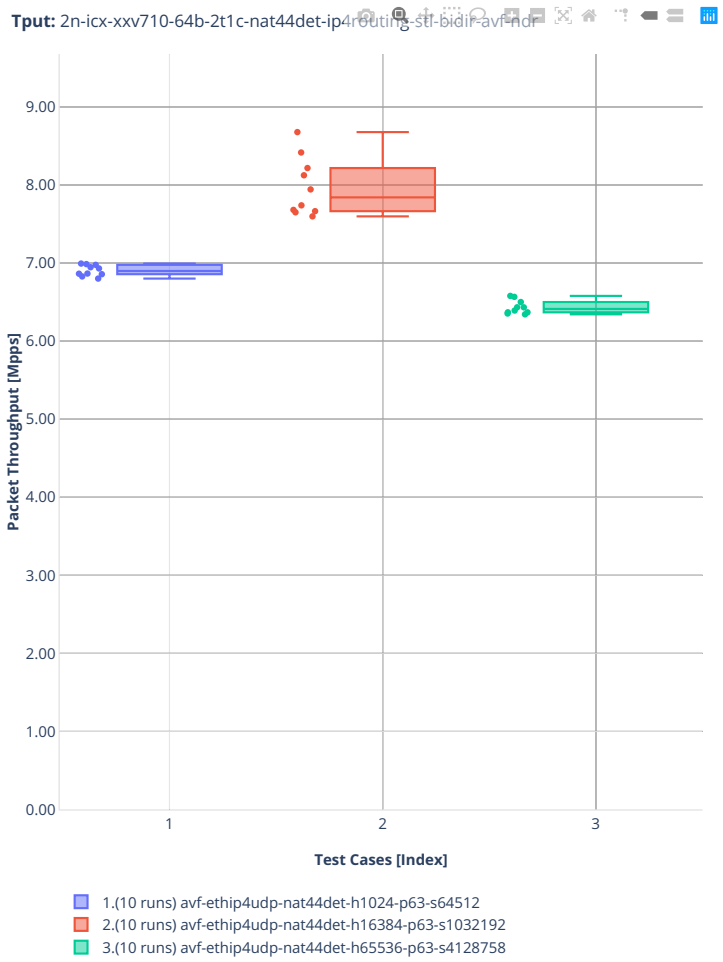

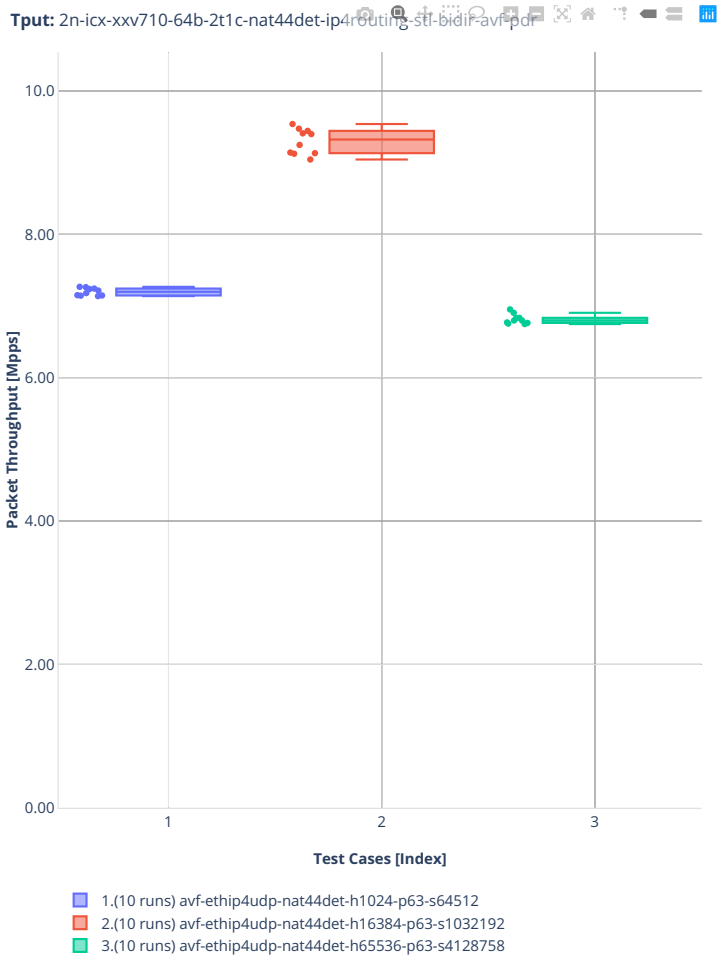
Tput: 2n-icx-xxv710-64b-2t1c-nat44det-ip4routing-stl-bidir-avf-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-ethip4udp-nat44det-h1024-p63-s64512
2.(10 runs) avf-ethip4udp-nat44det-h16384-p63-s1032192
3.(10 runs) avf-ethip4udp-nat44det-h65536-p63-s4128758

2n-clx-xxv710

64b-nat44det-ip4routing-stl-bidir-avf

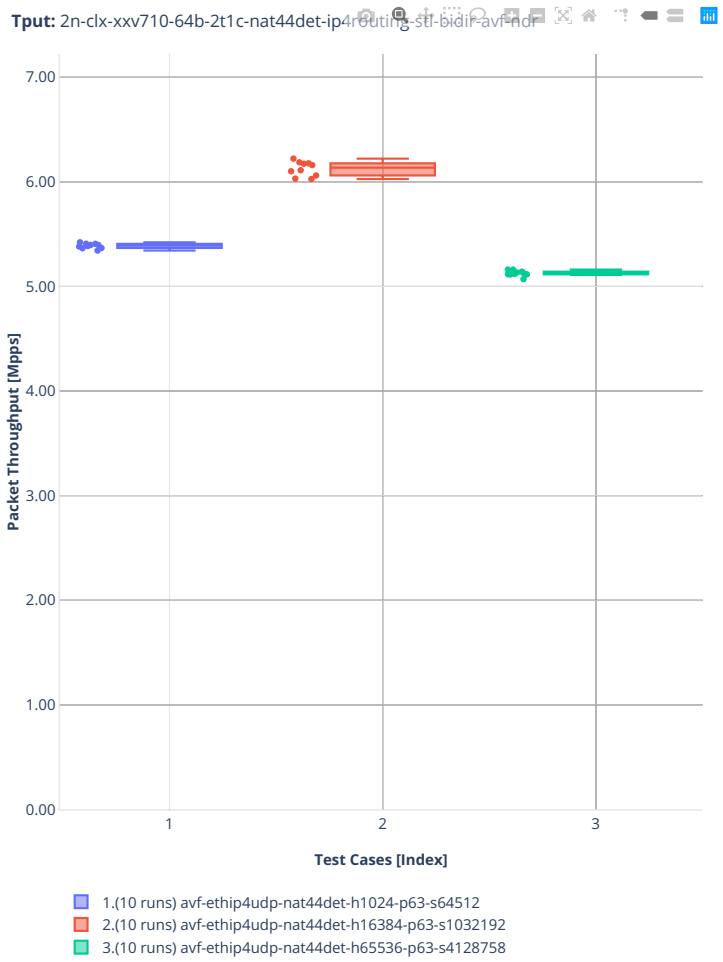

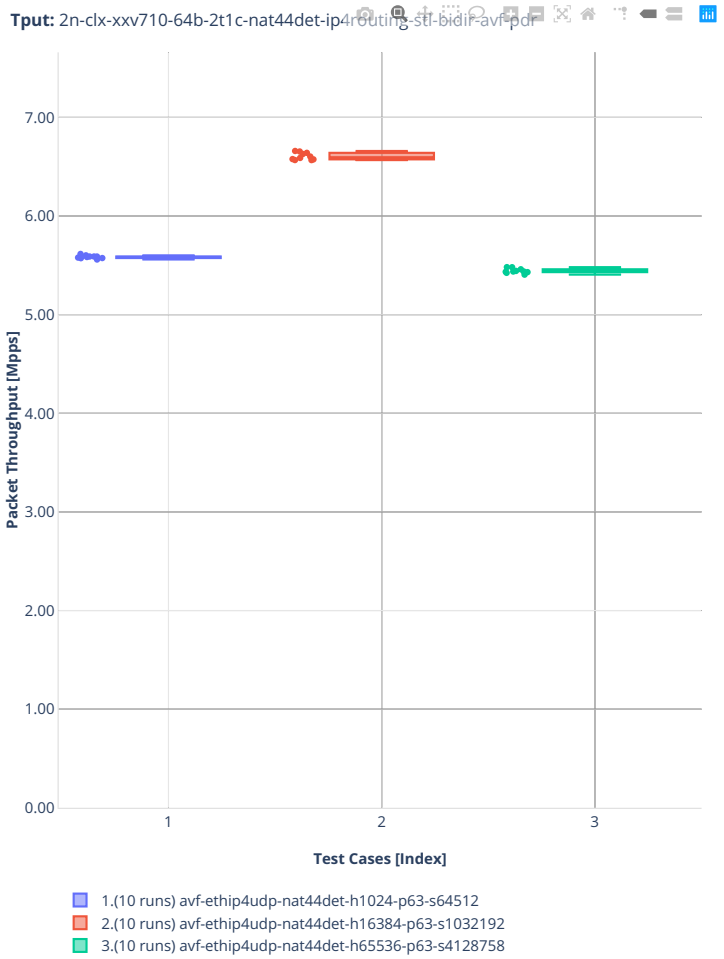
Tput: 2n-clx-xxv710-64b-2t1c-nat44det-ip4routing-stl-bidir-avf-pdr
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-ethip4udp-nat44det-h1024-p63-s64512
2.(10 runs) avf-ethip4udp-nat44det-h16384-p63-s1032192
3.(10 runs) avf-ethip4udp-nat44det-h65536-p63-s4128758

2n-skx-xxv710

64b-nat44det-ip4routing-stl-bidir-avf

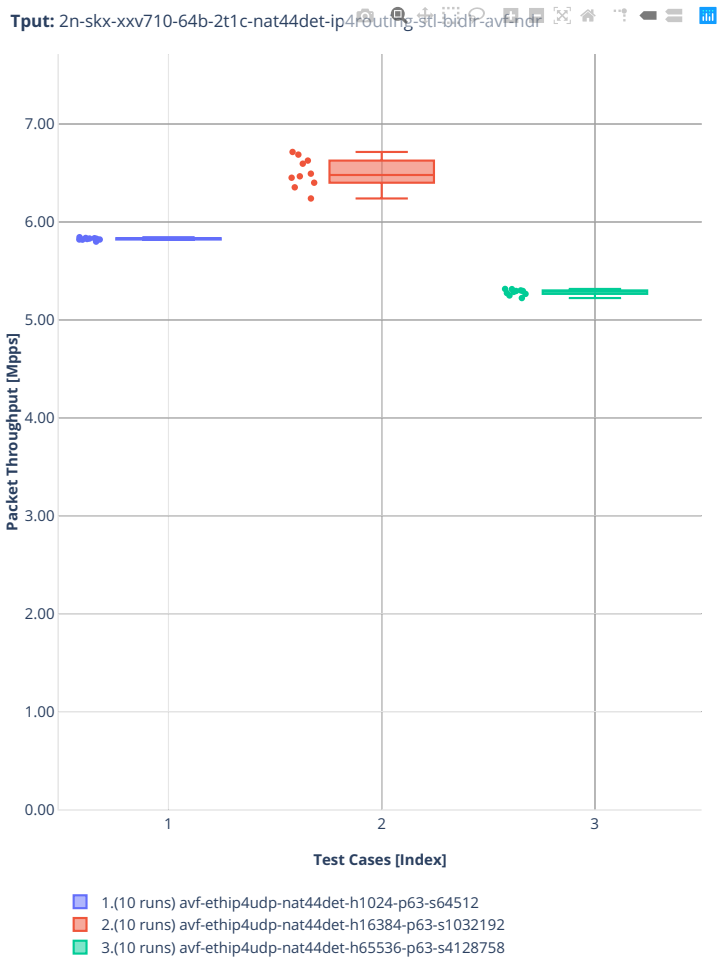

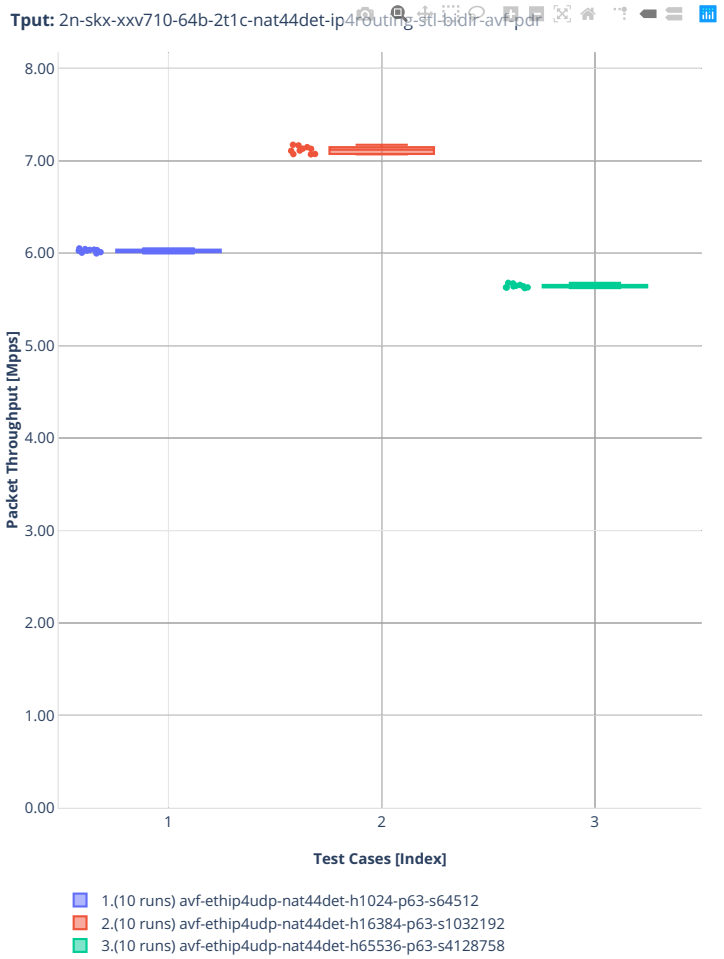
Tput: 2n-skx-xxv710-64b-2t1c-nat44det-ip4routing-stl-bidir-avf-pdr
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-ethip4udp-nat44det-h1024-p63-s64512
2.(10 runs) avf-ethip4udp-nat44det-h16384-p63-s1032192
3.(10 runs) avf-ethip4udp-nat44det-h65536-p63-s4128758

ED UniDir

2n-icx-xxv710

64b-nat44ed-ip4routing-stl-unidir-avf

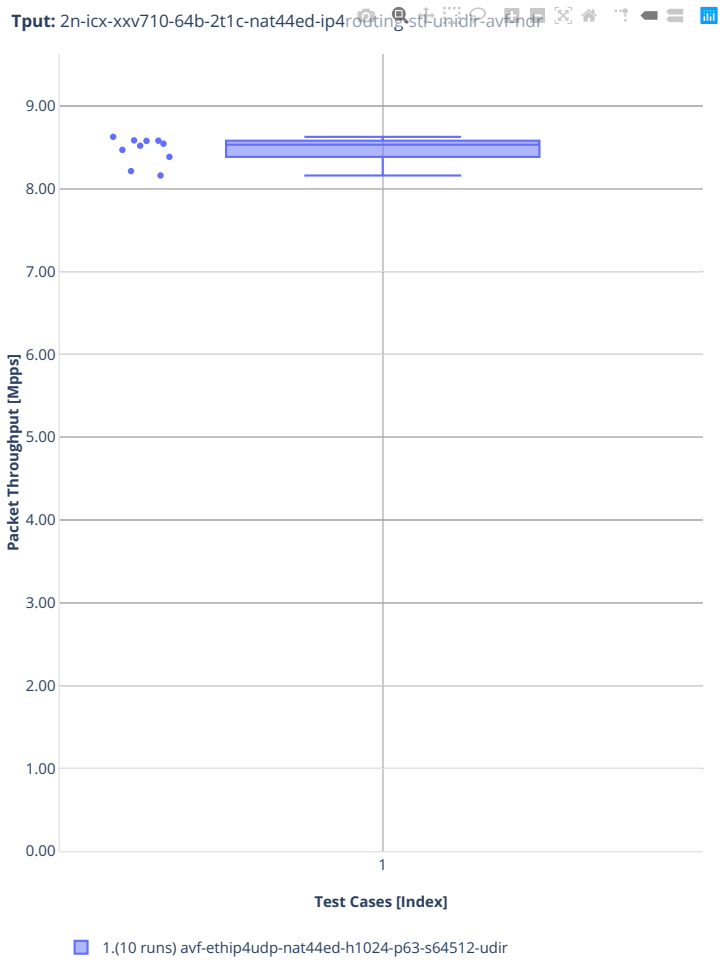

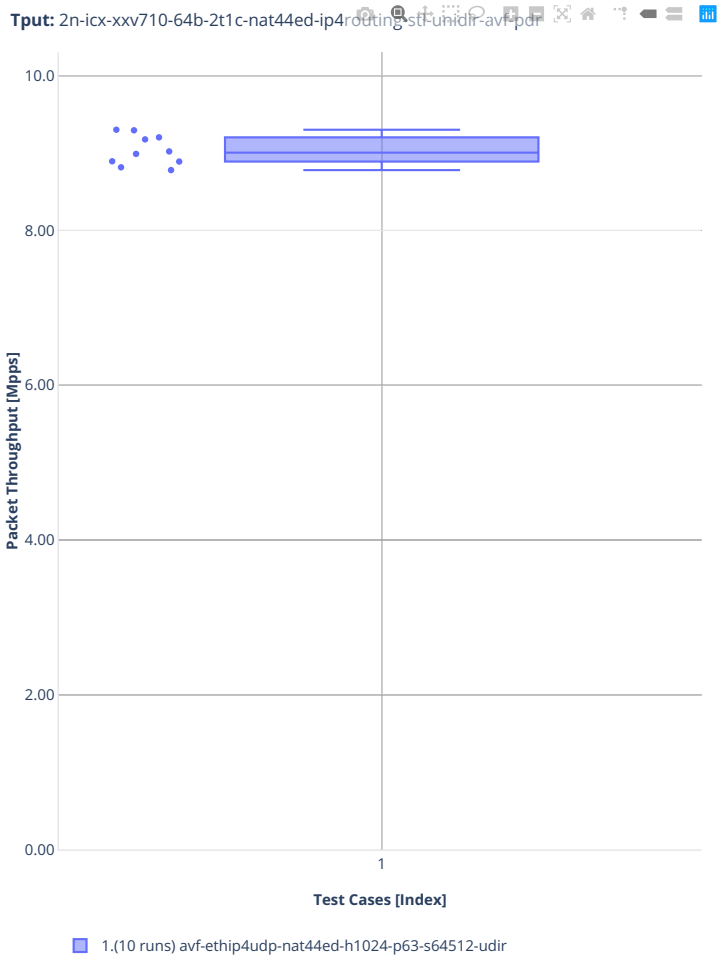
Tput: 2n-icx-xxv710-64b-2t1c-nat44ed-ip4routing-stl-unidir-avf-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
Test Cases [Index]
1.(10 runs) avf-ethip4udp-nat44ed-h1024-p63-s64512-udir

2n-clx-xxv710

64b-nat44ed-ip4routing-stl-unidir-avf

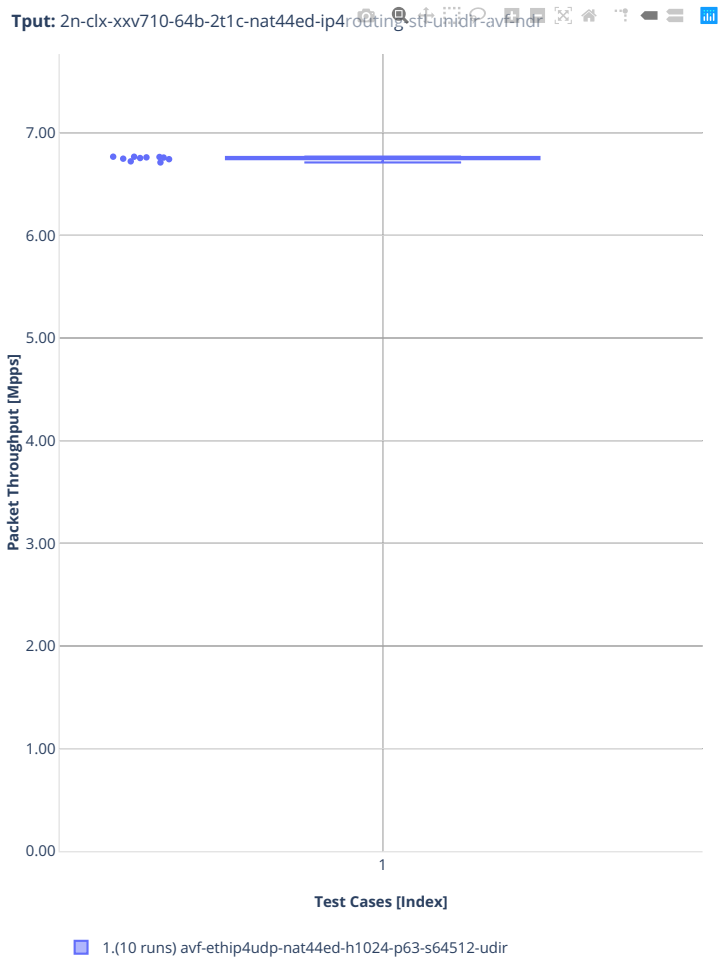

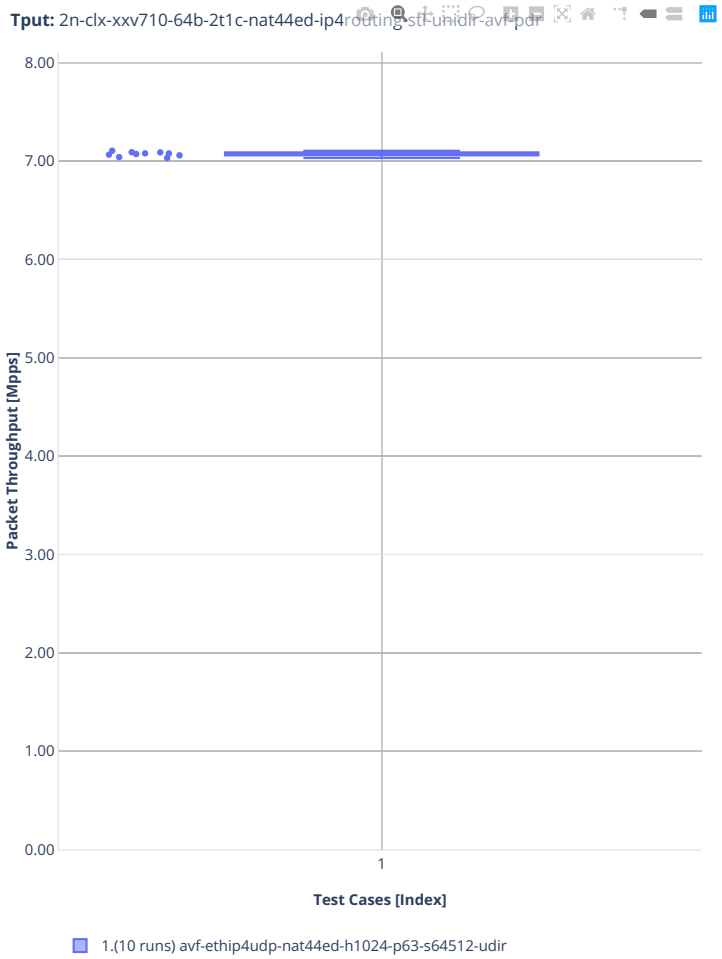
Tput: 2n-clx-xxv710-64b-2t1c-nat44ed-ip4routing-stl-unidir-avf-pdr
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
Test Cases [Index]
1.(10 runs) avf-ethip4udp-nat44ed-h1024-p63-s64512-udir

2n-skx-xxv710

64b-nat44ed-ip4routing-stl-unidir-avf

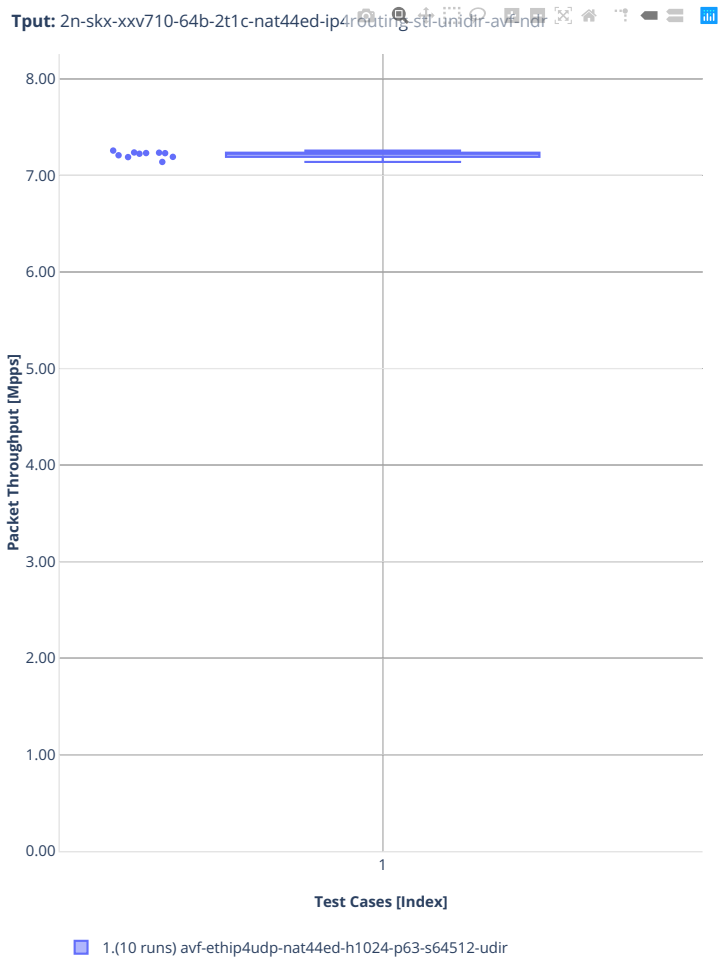

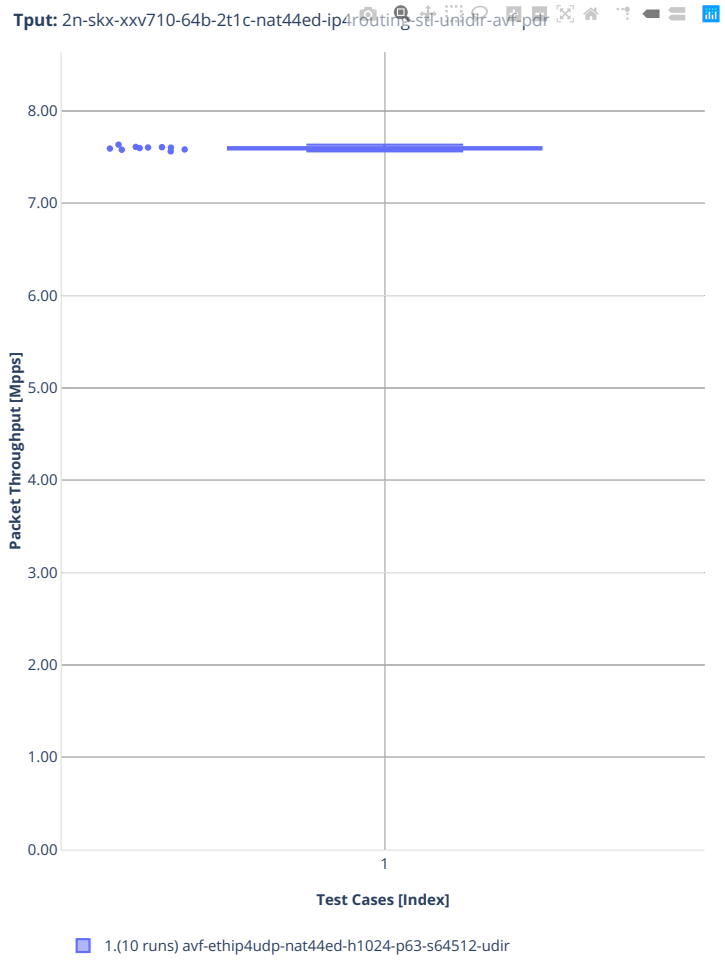
Tput: 2n-skx-xxv710-64b-2t1c-nat44ed-ip4routing-stl-unidir-avf-pdr
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
Test Cases [Index]
1.(10 runs) avf-ethip4udp-nat44ed-h1024-p63-s64512-udir

ED UDP CPS

2n-icx-xxv710

64b-nat44ed-ip4routing-udp-stf-cps-avf

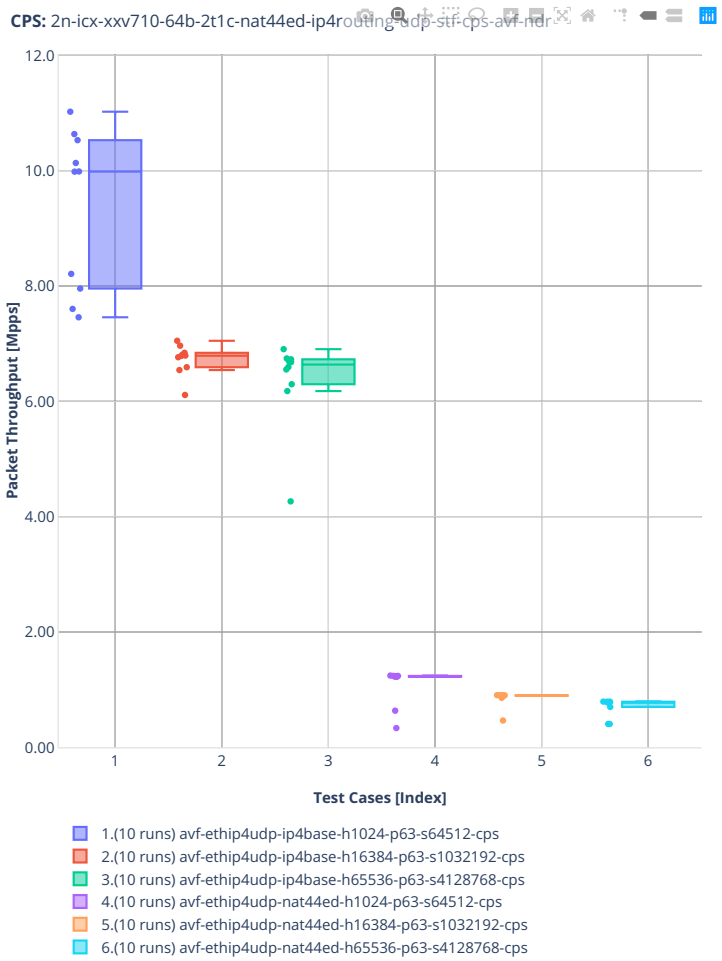

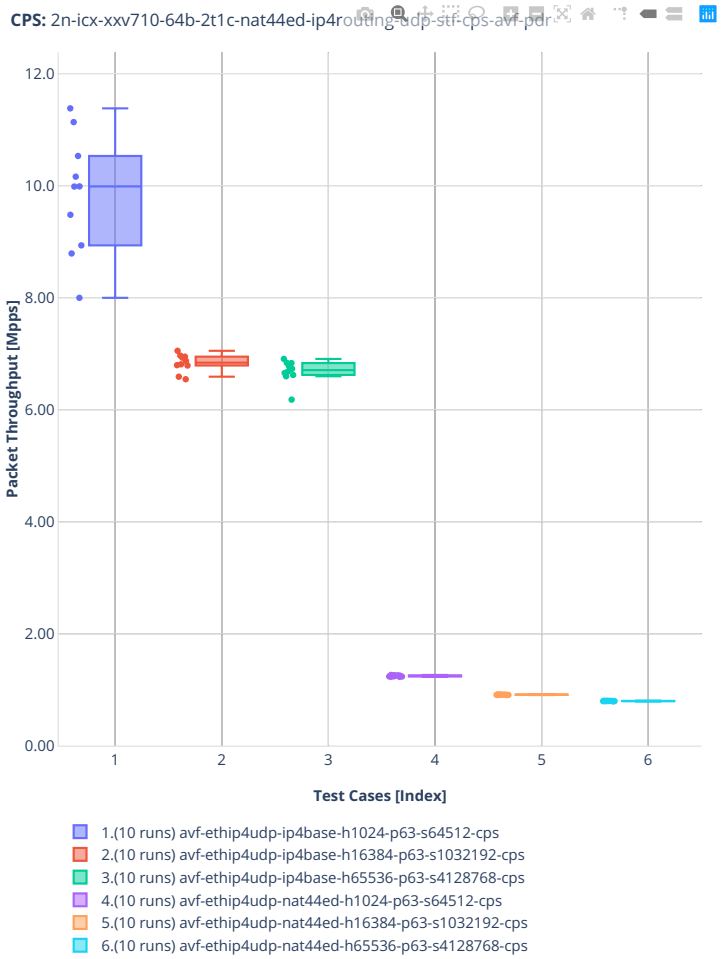
CPS: 2n-icx-xxv710-64b-2t1c-nat44ed-ip4routing-udp-stf-cps-avf-pdr
Packet Throughput [Mpps]
12.0
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-ethip4udp-ip4base-h1024-p63-s64512-cps
2.(10 runs) avf-ethip4udp-ip4base-h16384-p63-s1032192-cps
3.(10 runs) avf-ethip4udp-ip4base-h65536-p63-s4128768-cps
4.(10 runs) avf-ethip4udp-nat44ed-h1024-p63-s64512-cps
5.(10 runs) avf-ethip4udp-nat44ed-h16384-p63-s1032192-cps
6.(10 runs) avf-ethip4udp-nat44ed-h65536-p63-s4128768-cps

2n-clx-xxv710

64b-nat44ed-ip4routing-udp-stf-cps-avf

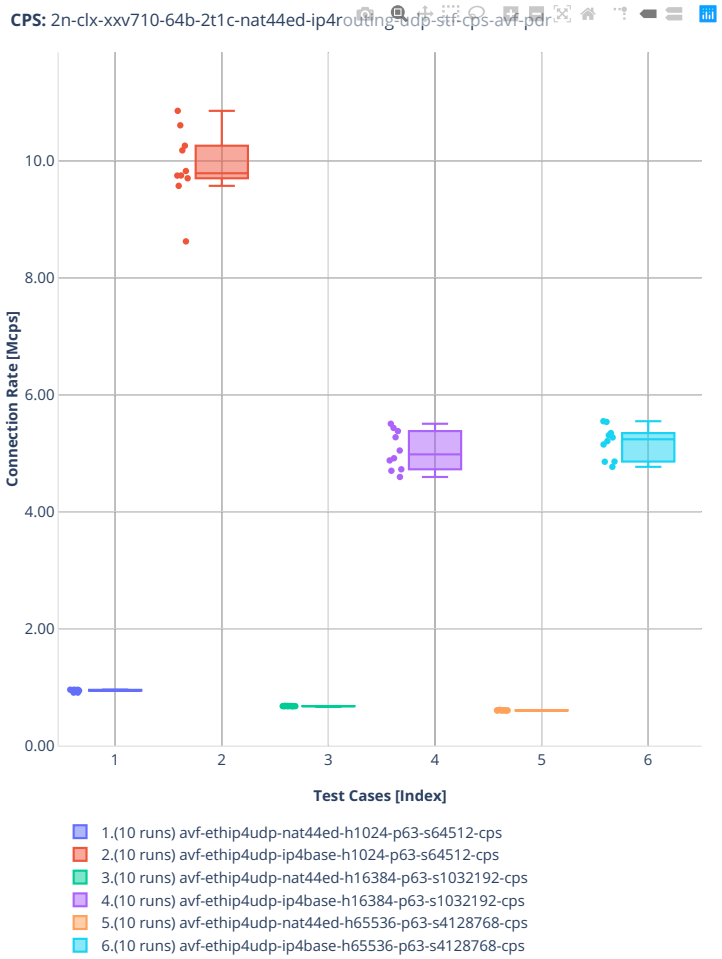
CPS: 2n-clx-xxv710-64b-2t1c-nat44ed-ip4routing-udp-stf-cps-avf-pdr
Connection Rate [Mcps]
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-ethip4udp-nat44ed-h1024-p63-s64512-cps
2.(10 runs) avf-ethip4udp-ip4base-h1024-p63-s64512-cps
3.(10 runs) avf-ethip4udp-nat44ed-h16384-p63-s1032192-cps
4.(10 runs) avf-ethip4udp-ip4base-h16384-p63-s1032192-cps
5.(10 runs) avf-ethip4udp-nat44ed-h65536-p63-s4128768-cps
6.(10 runs) avf-ethip4udp-ip4base-h65536-p63-s4128768-cps

2n-skx-xxv710

64b-nat44ed-ip4routing-udp-stf-cps-avf

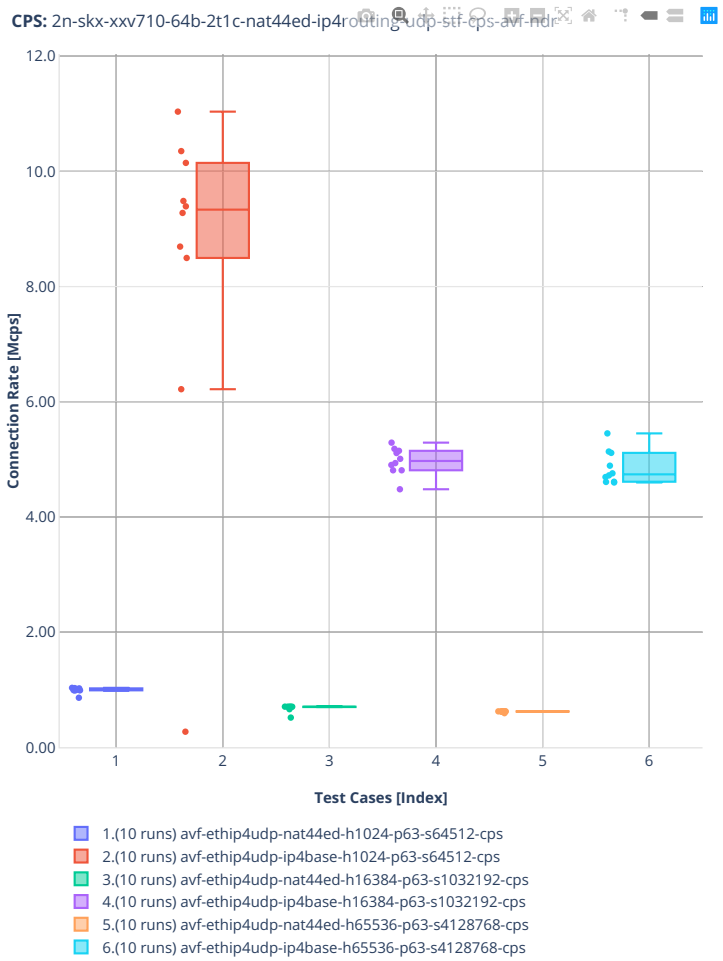

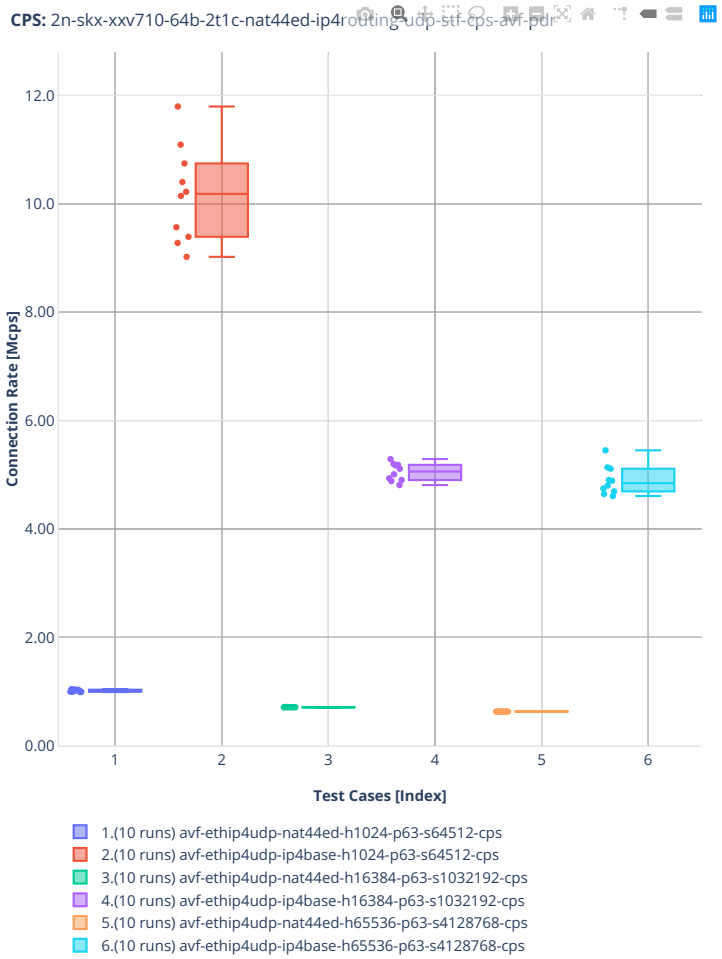
CPS: 2n-skx-xxv710-64b-2t1c-nat44ed-ip4routing-udp-stf-cps-avf-pdr
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Connection Rate [Mcps]
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-ethip4udp-nat44ed-h1024-p63-s64512-cps
2.(10 runs) avf-ethip4udp-ip4base-h1024-p63-s64512-cps
3.(10 runs) avf-ethip4udp-nat44ed-h16384-p63-s1032192-cps
4.(10 runs) avf-ethip4udp-ip4base-h16384-p63-s1032192-cps
5.(10 runs) avf-ethip4udp-nat44ed-h65536-p63-s4128768-cps
6.(10 runs) avf-ethip4udp-ip4base-h65536-p63-s4128768-cps

ED TCP CPS

2n-icx-xxv710

64b-nat44ed-ip4routing-tcp-stf-cps-avf

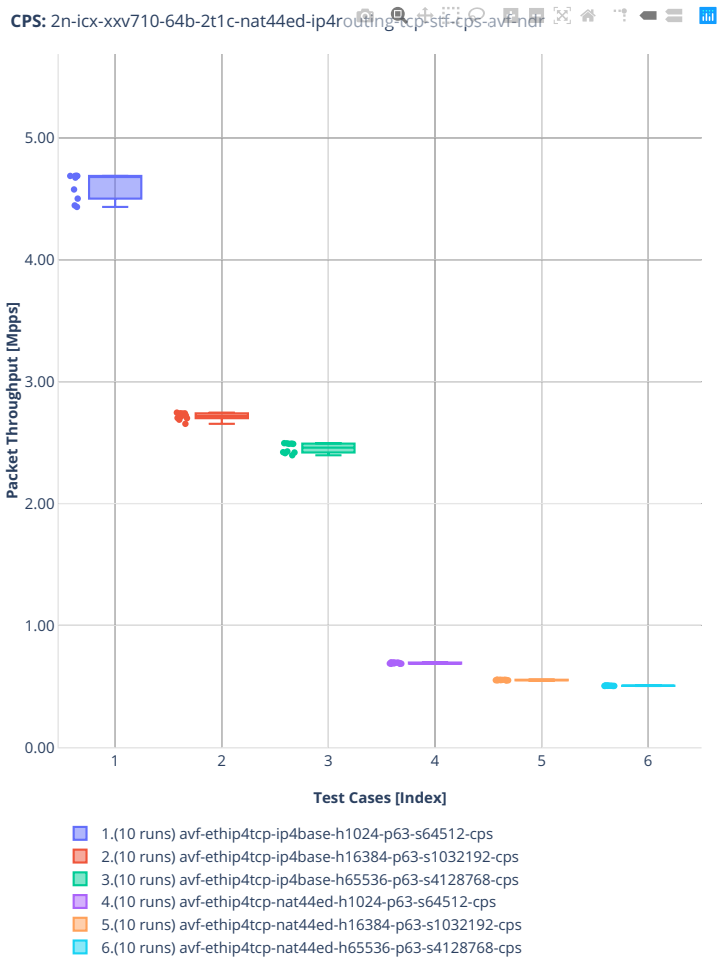

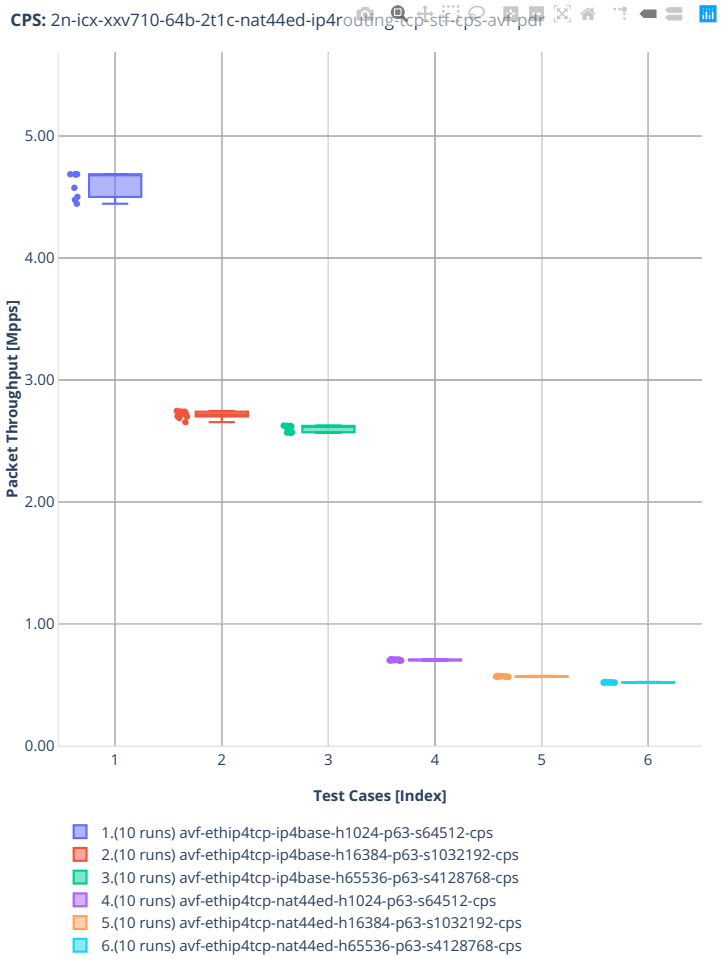
CPS: 2n-icx-xxv710-64b-2t1c-nat44ed-ip4routing-tcp-stf-cps-avf-pdr
Packet Throughput [Mpps]
5.00
4.00
3.00
2.00
1.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-ethip4tcp-ip4base-h1024-p63-s64512-cps
2.(10 runs) avf-ethip4tcp-ip4base-h16384-p63-s1032192-cps
3.(10 runs) avf-ethip4tcp-ip4base-h65536-p63-s4128768-cps
4.(10 runs) avf-ethip4tcp-nat44ed-h1024-p63-s64512-cps
5.(10 runs) avf-ethip4tcp-nat44ed-h16384-p63-s1032192-cps
6.(10 runs) avf-ethip4tcp-nat44ed-h65536-p63-s4128768-cps

2n-clx-xxv710

64b-nat44ed-ip4routing-tcp-stf-cps-avf

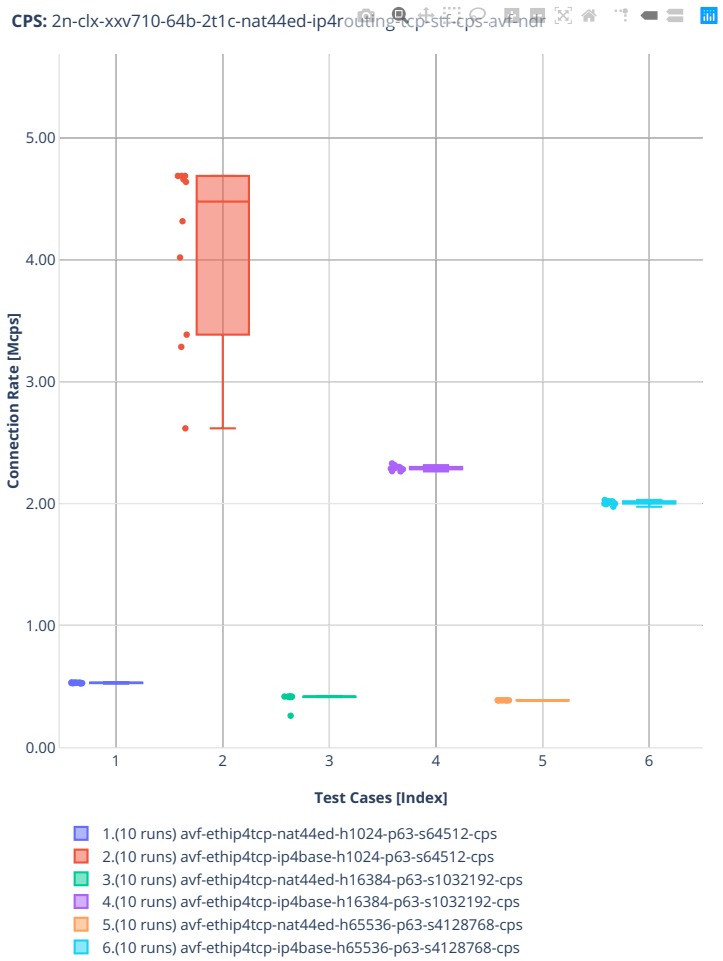

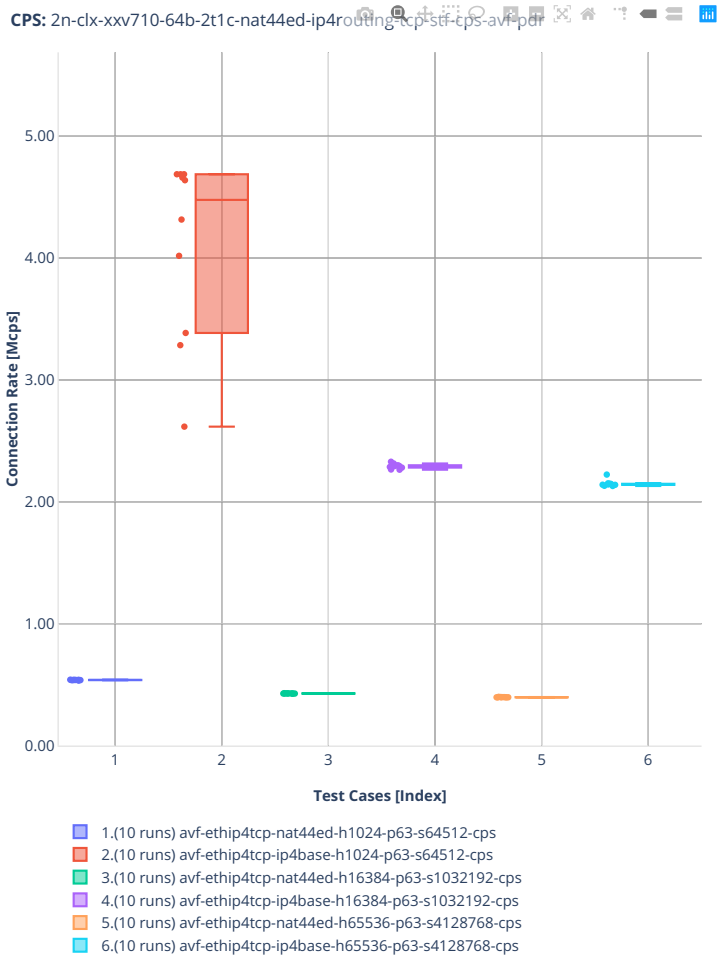
CPS: 2n-clx-xxv710-64b-2t1c-nat44ed-ip4routing-tcp-stf-cps-avf-pdr
Connection Rate [Mcps]
5.00
4.00
3.00
2.00
1.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-ethip4tcp-nat44ed-h1024-p63-s64512-cps
2.(10 runs) avf-ethip4tcp-ip4base-h1024-p63-s64512-cps
3.(10 runs) avf-ethip4tcp-nat44ed-h16384-p63-s1032192-cps
4.(10 runs) avf-ethip4tcp-ip4base-h16384-p63-s1032192-cps
5.(10 runs) avf-ethip4tcp-nat44ed-h65536-p63-s4128768-cps
6.(10 runs) avf-ethip4tcp-ip4base-h65536-p63-s4128768-cps

2n-skx-xxv710

64b-nat44ed-ip4routing-tcp-stf-cps-avf

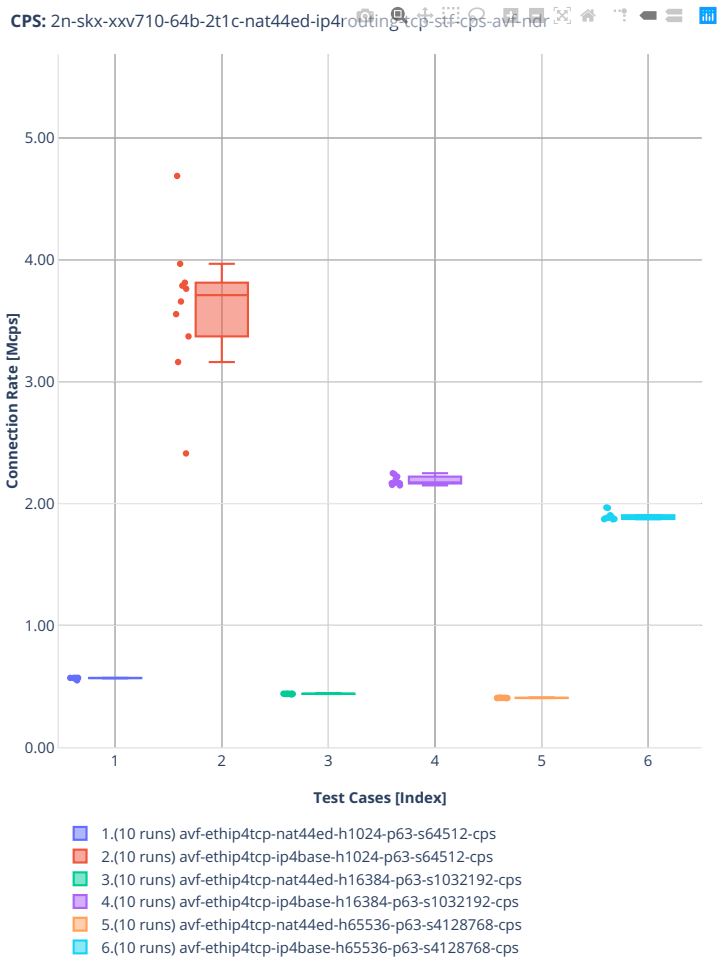

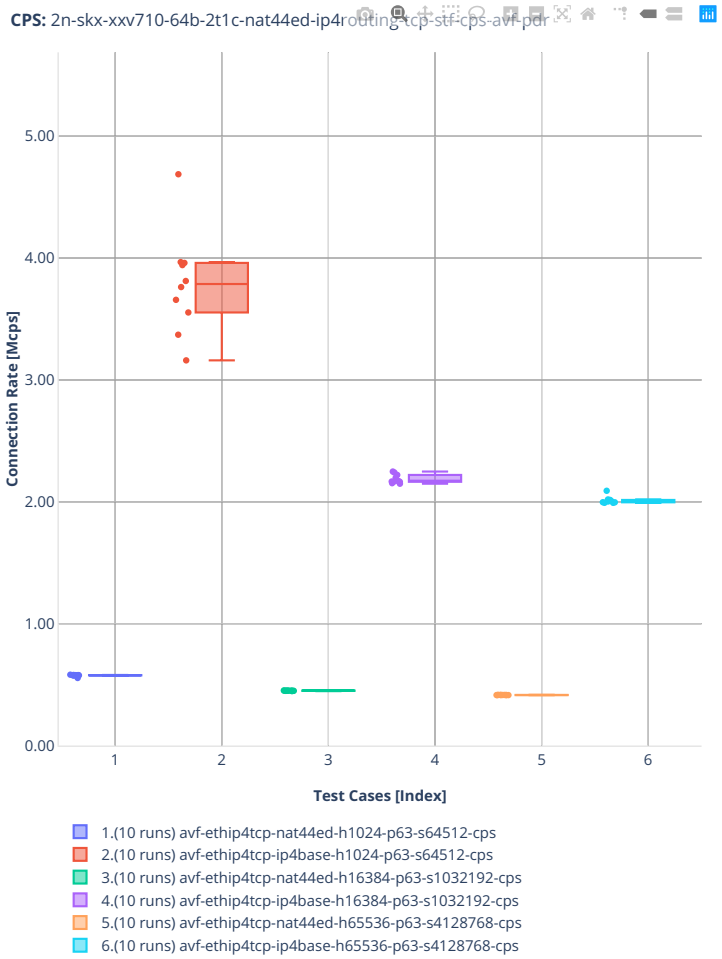
CPS: 2n-skx-xxv710-64b-2t1c-nat44ed-ip4routing-tcp-stf-cps-avf-pdr
Connection Rate [Mcps]
5.00
4.00
3.00
2.00
1.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-ethip4tcp-nat44ed-h1024-p63-s64512-cps
2.(10 runs) avf-ethip4tcp-ip4base-h1024-p63-s64512-cps
3.(10 runs) avf-ethip4tcp-nat44ed-h16384-p63-s1032192-cps
4.(10 runs) avf-ethip4tcp-ip4base-h16384-p63-s1032192-cps
5.(10 runs) avf-ethip4tcp-nat44ed-h65536-p63-s4128768-cps
6.(10 runs) avf-ethip4tcp-ip4base-h65536-p63-s4128768-cps

ED UDP TPUT

2n-icx-xxv710

100b-nat44ed-ip4routing-udp-tput-avf

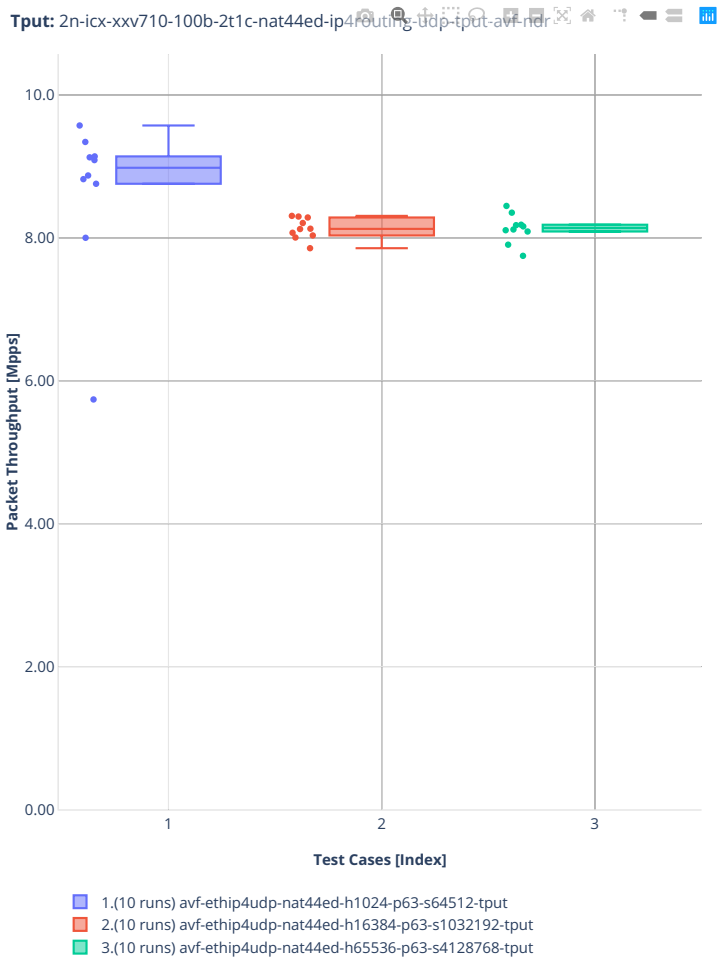

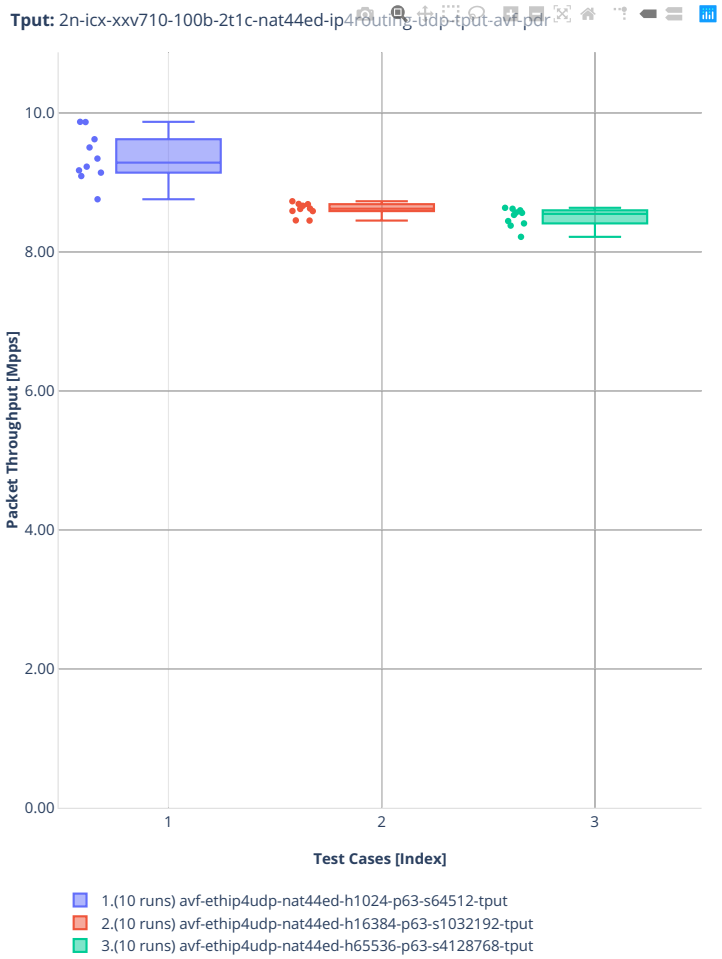
Tput: 2n-icx-xxv710-100b-2t1c-nat44ed-ip4routing-udp-tput-avf-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-ethip4udp-nat44ed-h1024-p63-s64512-tput
2.(10 runs) avf-ethip4udp-nat44ed-h16384-p63-s1032192-tput
3.(10 runs) avf-ethip4udp-nat44ed-h65536-p63-s4128768-tput

2n-clx-xxv710

100b-nat44ed-ip4routing-udp-tput-avf

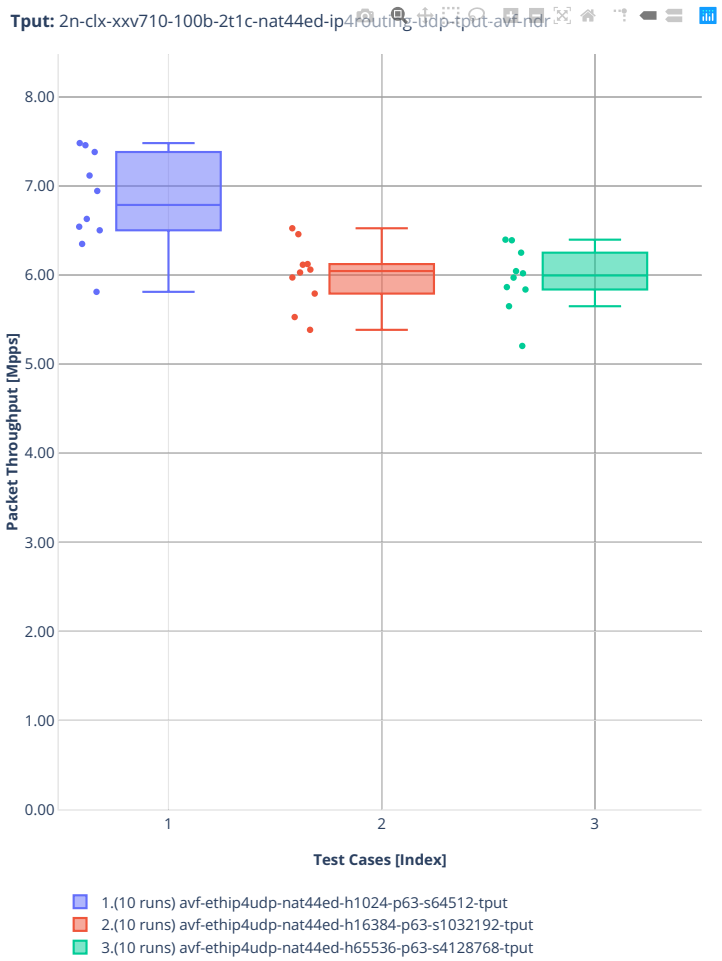

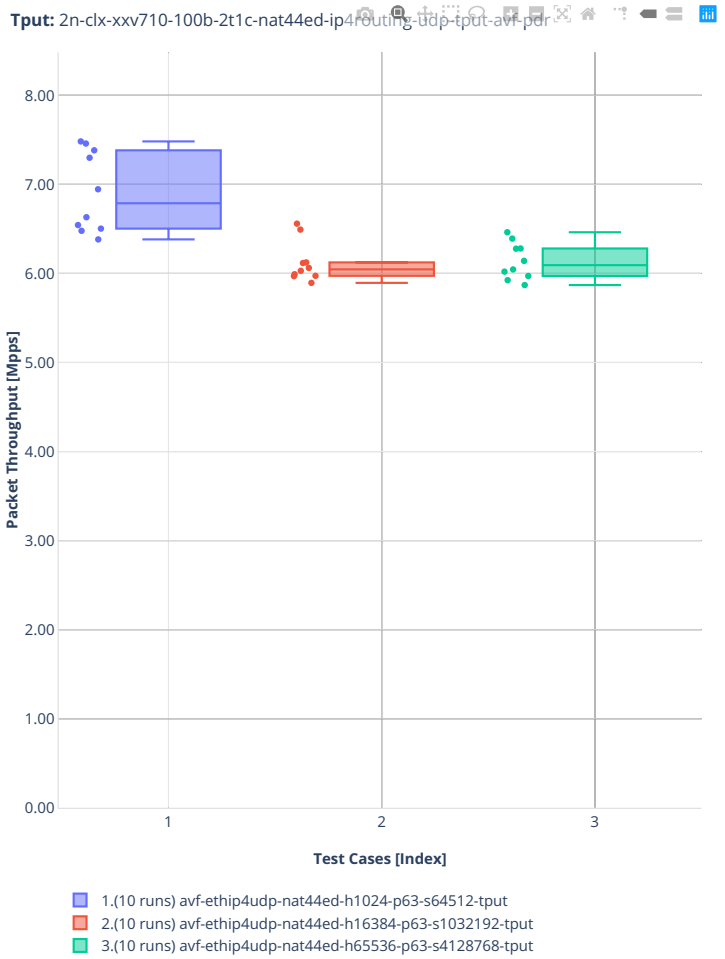
Tput: 2n-clx-xxv710-100b-2t1c-nat44ed-ip4routing-udp-tput-avf-pdr
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-ethip4udp-nat44ed-h1024-p63-s64512-tput
2.(10 runs) avf-ethip4udp-nat44ed-h16384-p63-s1032192-tput
3.(10 runs) avf-ethip4udp-nat44ed-h65536-p63-s4128768-tput

2n-skx-xxv710

100b-nat44ed-ip4routing-udp-tput-avf

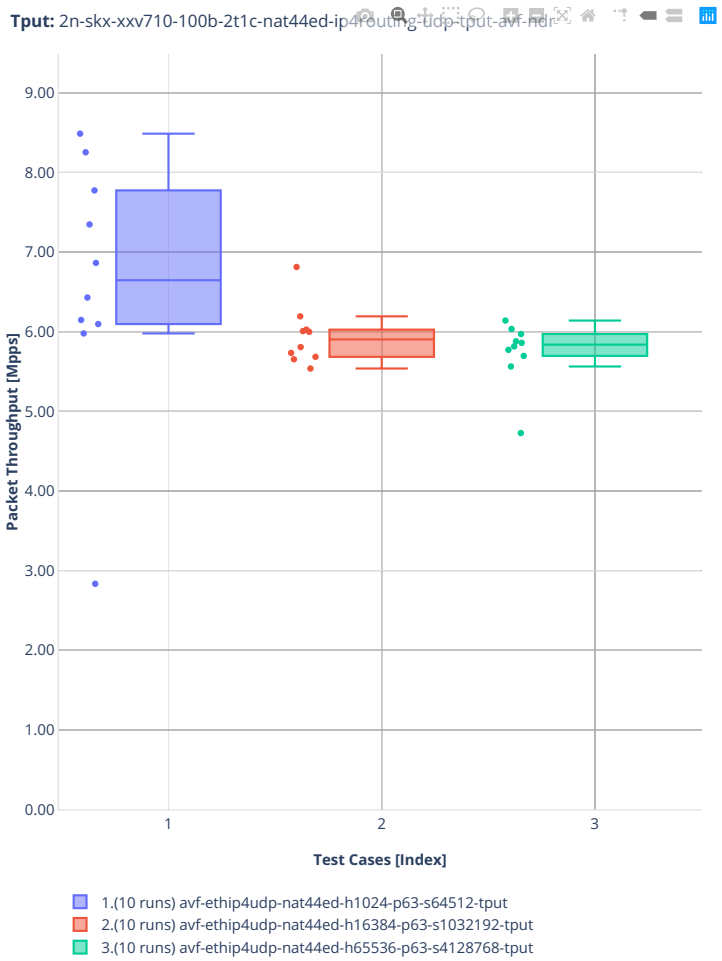

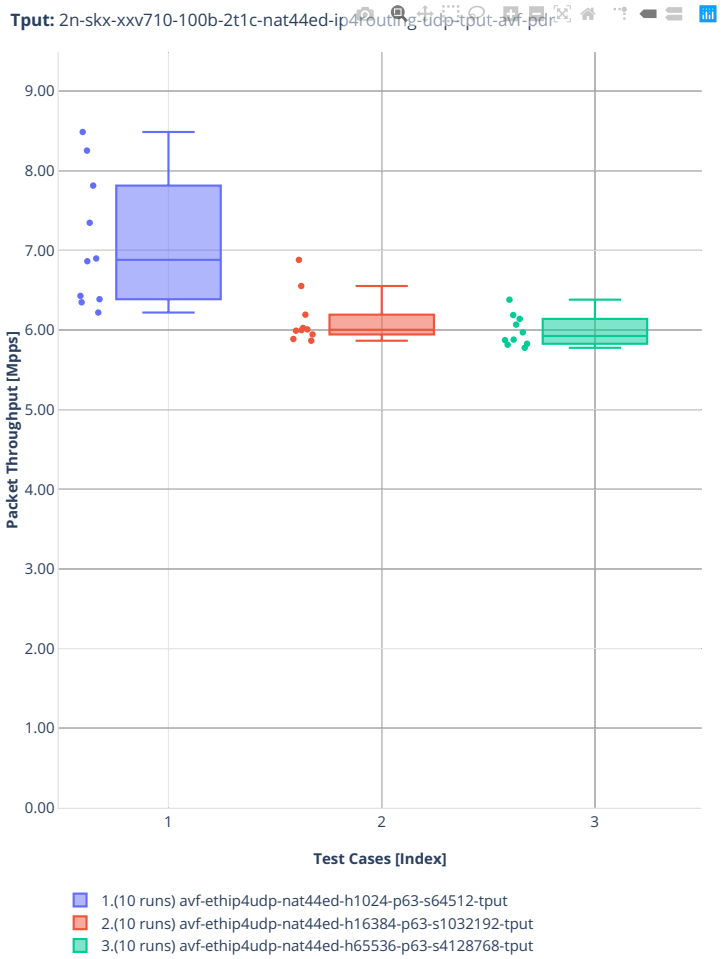
Tput: 2n-skx-xxv710-100b-2t1c-nat44ed-ip4routing-udp-tput-avf-pdr
9.00
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-ethip4udp-nat44ed-h1024-p63-s64512-tput
2.(10 runs) avf-ethip4udp-nat44ed-h16384-p63-s1032192-tput
3.(10 runs) avf-ethip4udp-nat44ed-h65536-p63-s4128768-tput

ED TCP TPUT

2n-icx-xxv710

100b-nat44ed-ip4routing-tcp-tput-avf

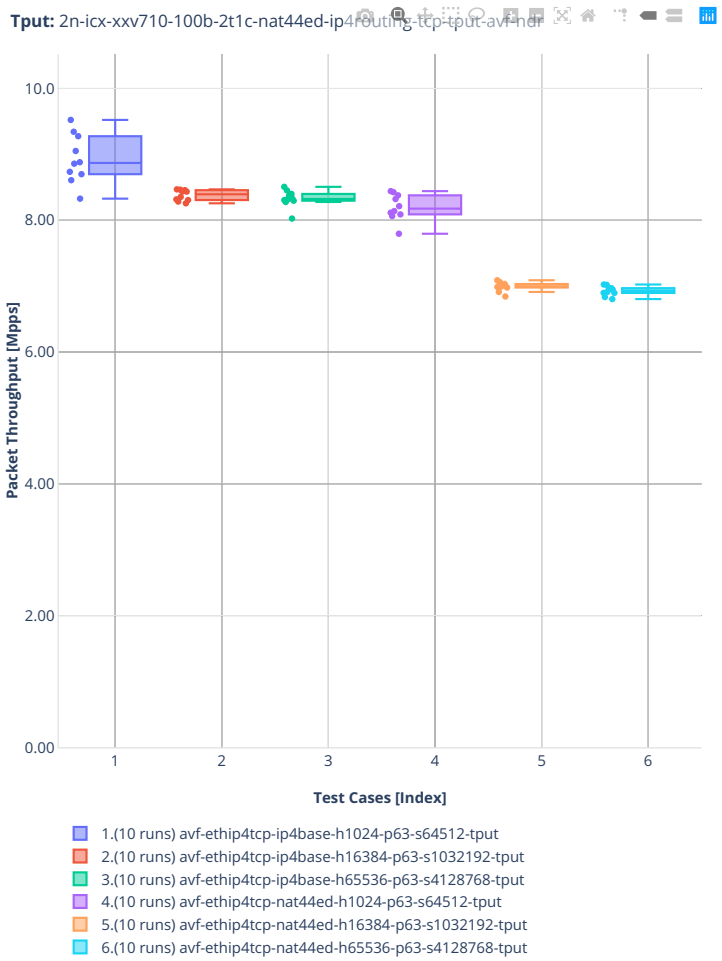

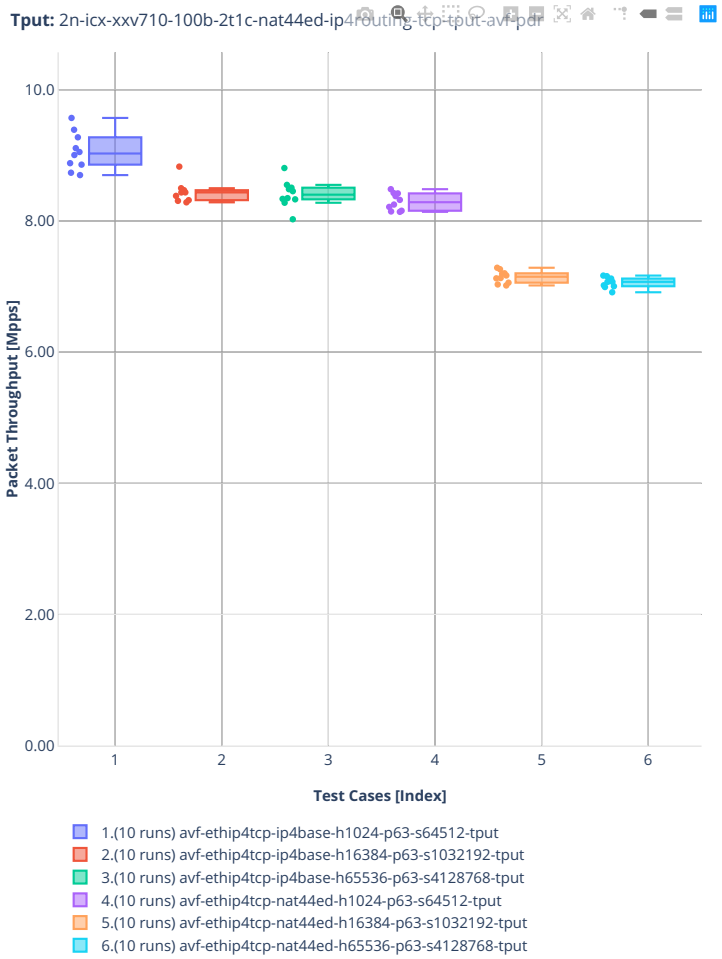
Tput: 2n-icx-xxv710-100b-2t1c-nat44ed-ip4routing-tcp-tput-avf-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-ethip4tcp-ip4base-h1024-p63-s64512-tput
2.(10 runs) avf-ethip4tcp-ip4base-h16384-p63-s1032192-tput
3.(10 runs) avf-ethip4tcp-ip4base-h65536-p63-s4128768-tput
4.(10 runs) avf-ethip4tcp-nat44ed-h1024-p63-s64512-tput
5.(10 runs) avf-ethip4tcp-nat44ed-h16384-p63-s1032192-tput
6.(10 runs) avf-ethip4tcp-nat44ed-h65536-p63-s4128768-tput

2n-clx-xxv710

100b-nat44ed-ip4routing-tcp-tput-avf

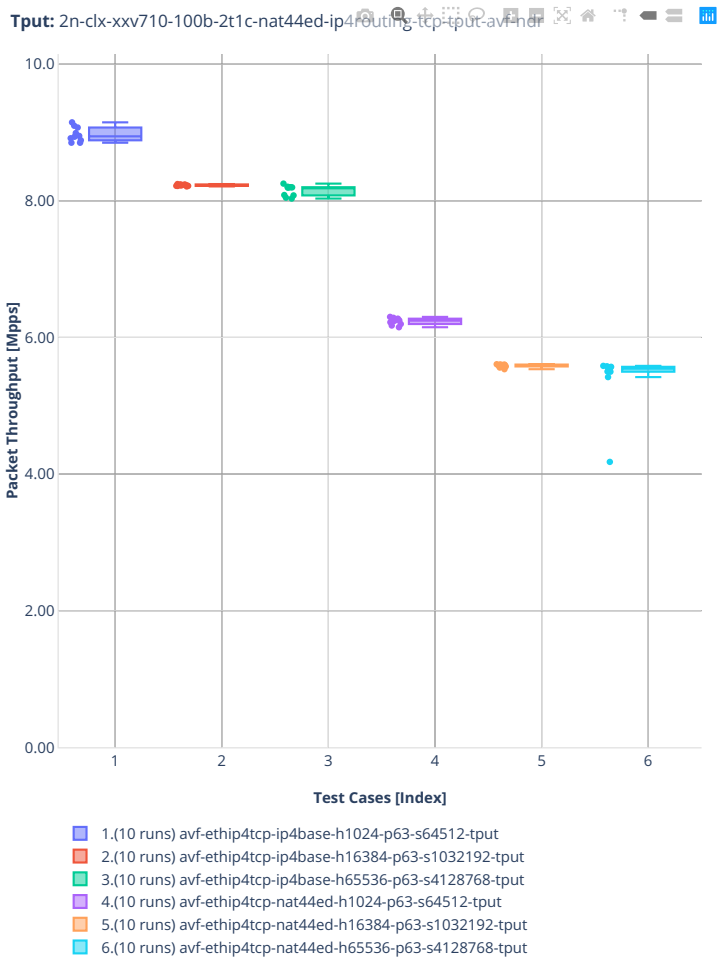

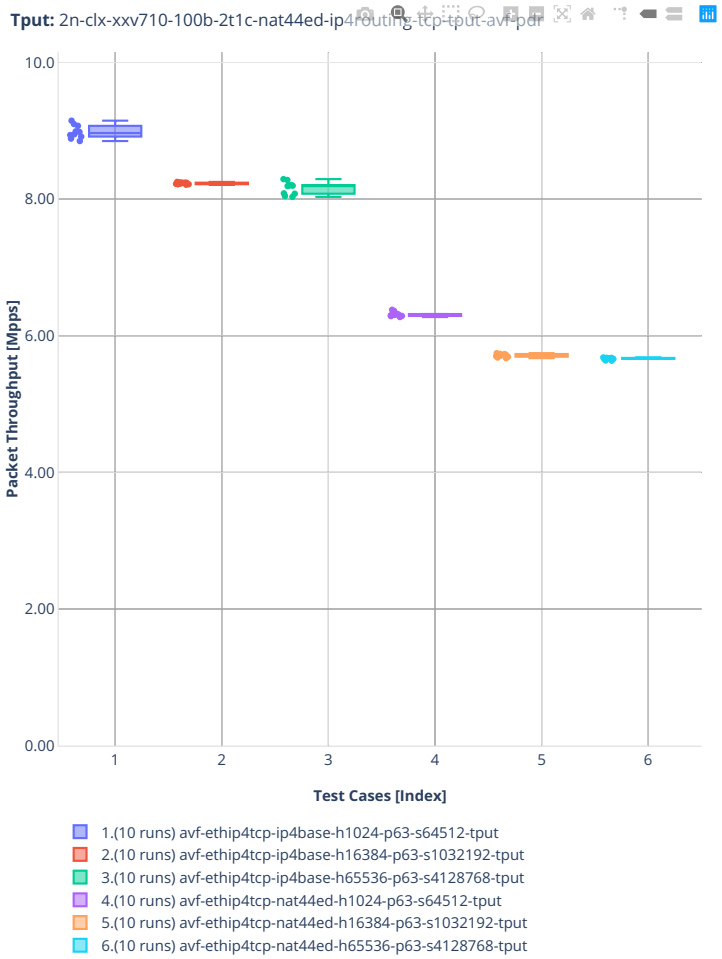
Tput: 2n-clx-xxv710-100b-2t1c-nat44ed-ip4routing-tcp-tput-avf-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-ethip4tcp-ip4base-h1024-p63-s64512-tput
2.(10 runs) avf-ethip4tcp-ip4base-h16384-p63-s1032192-tput
3.(10 runs) avf-ethip4tcp-ip4base-h65536-p63-s4128768-tput
4.(10 runs) avf-ethip4tcp-nat44ed-h1024-p63-s64512-tput
5.(10 runs) avf-ethip4tcp-nat44ed-h16384-p63-s1032192-tput
6.(10 runs) avf-ethip4tcp-nat44ed-h65536-p63-s4128768-tput

2n-skx-xxv710

100b-nat44ed-ip4routing-tcp-tput-avf

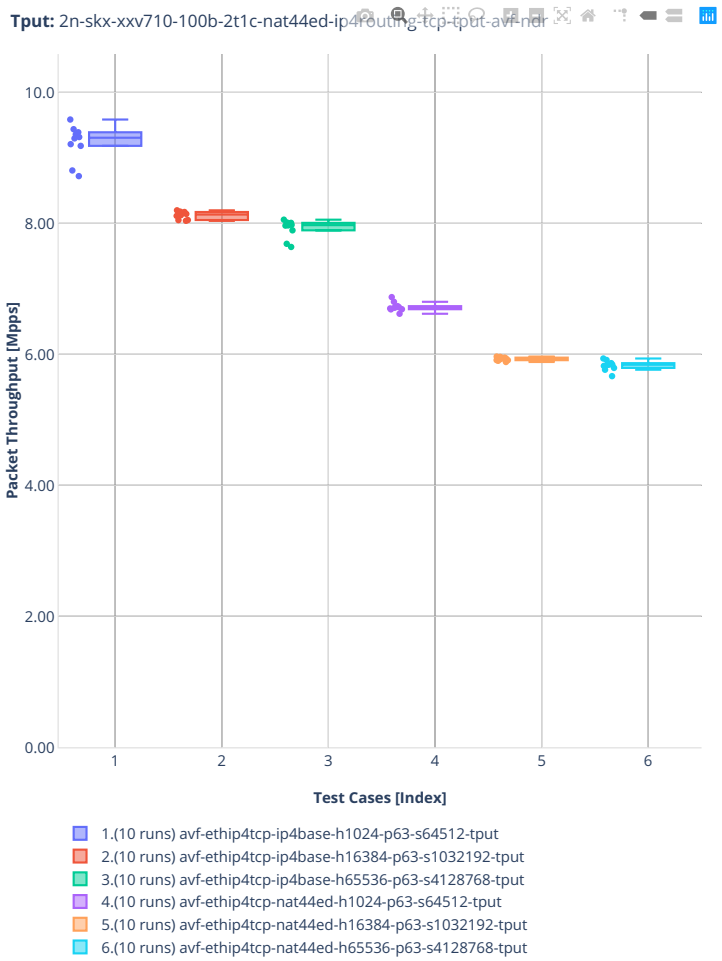

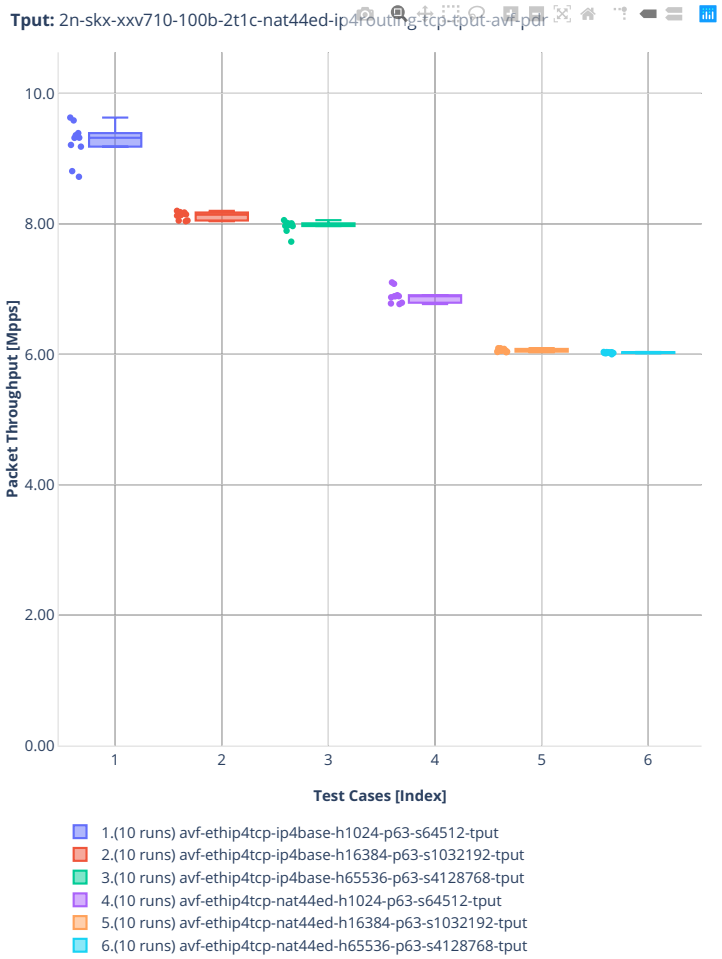
Tput: 2n-skx-xxv710-100b-2t1c-nat44ed-ip4routing-tcp-tput-avf-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-ethip4tcp-ip4base-h1024-p63-s64512-tput
2.(10 runs) avf-ethip4tcp-ip4base-h16384-p63-s1032192-tput
3.(10 runs) avf-ethip4tcp-ip4base-h65536-p63-s4128768-tput
4.(10 runs) avf-ethip4tcp-nat44ed-h1024-p63-s64512-tput
5.(10 runs) avf-ethip4tcp-nat44ed-h16384-p63-s1032192-tput
6.(10 runs) avf-ethip4tcp-nat44ed-h65536-p63-s4128768-tput

2.3.7 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[115].

[115] https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls2206

2n-icx-xxv710

64b-2t1c-vhost-base

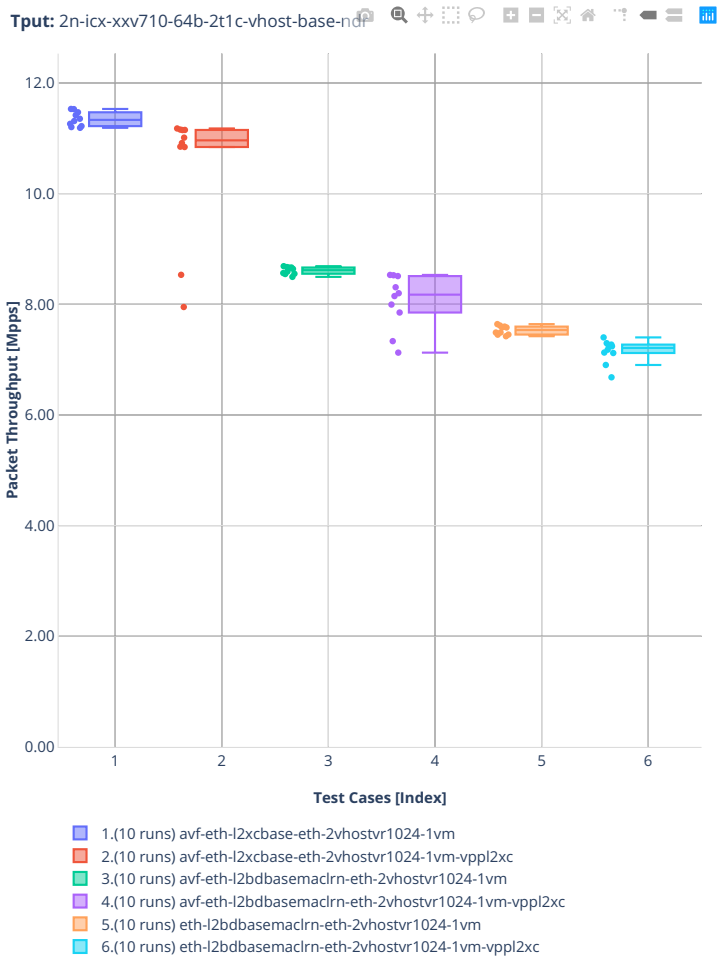

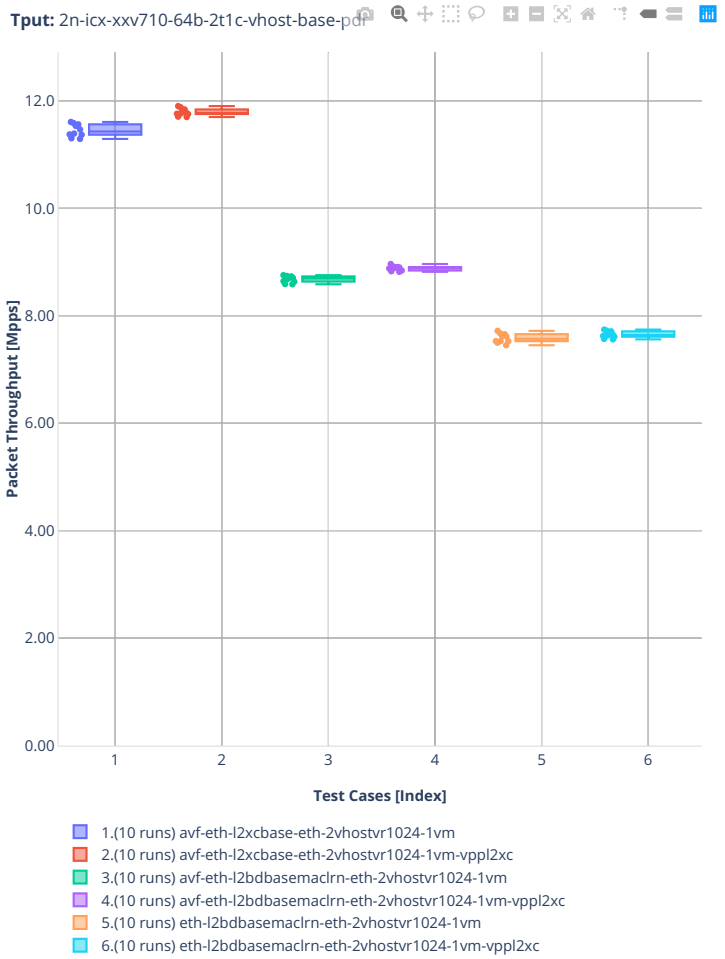
Tput: 2n-icx-xxv710-64b-2t1c-vhost-base-pdr
Packet Throughput [Mpps]
12.0
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-eth-l2xcbase-eth-2vhostvr1024-1vm
2.(10 runs) avf-eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
3.(10 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
4.(10 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
5.(10 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
6.(10 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2n-skx-xxv710

64b-2t1c-vhost-base-testpmd

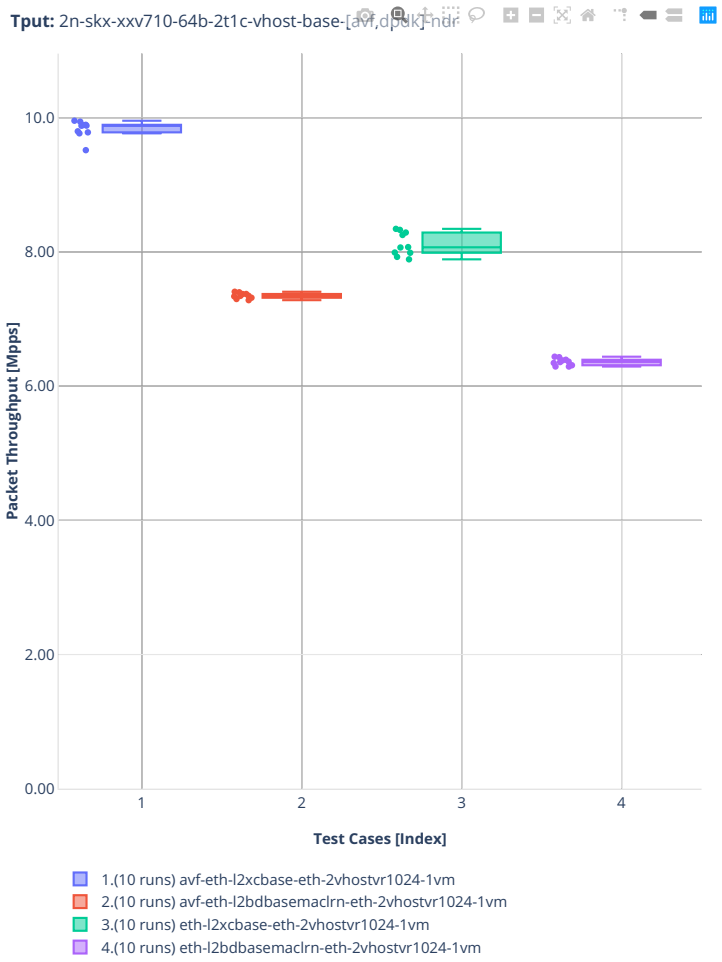

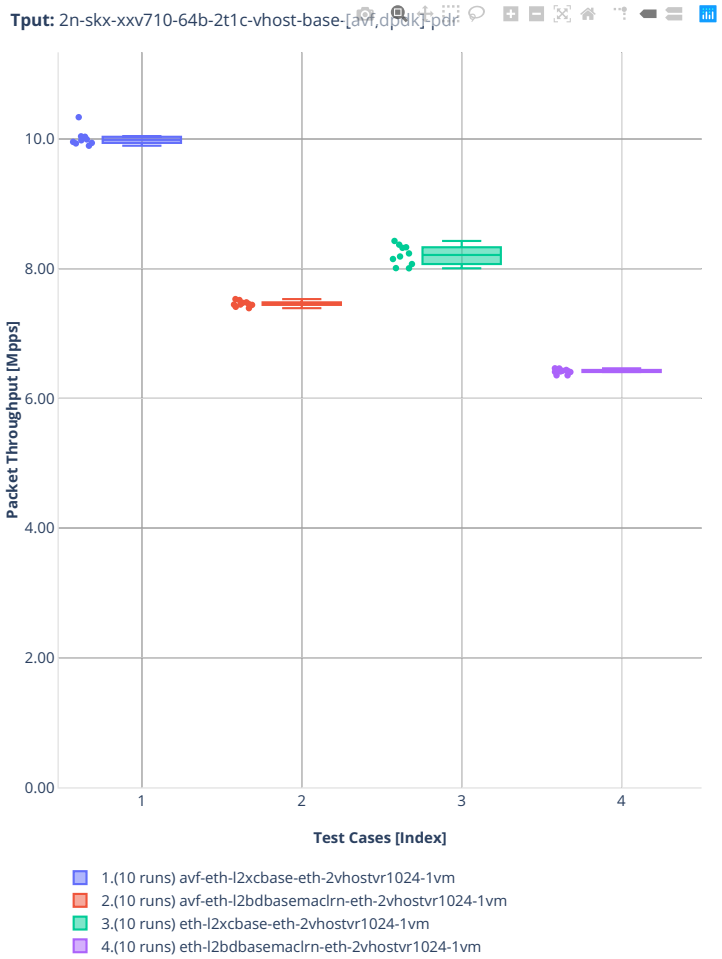
Tput: 2n-skx-xxv710-64b-2t1c-vhost-base-[avf,dpdk]-pdr
Packet Throughput [Mpps]
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
Test Cases [Index]
1.(10 runs) avf-eth-l2xcbase-eth-2vhostvr1024-1vm
2.(10 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
3.(10 runs) eth-l2xcbase-eth-2vhostvr1024-1vm
4.(10 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm

64b-2t1c-vhost-base-vpp

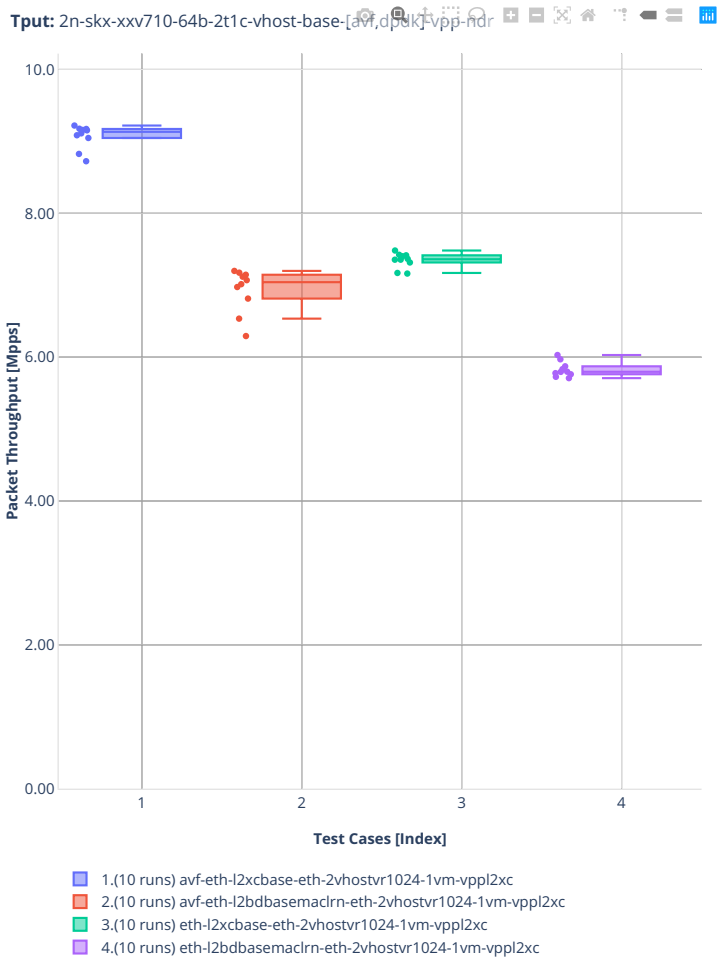

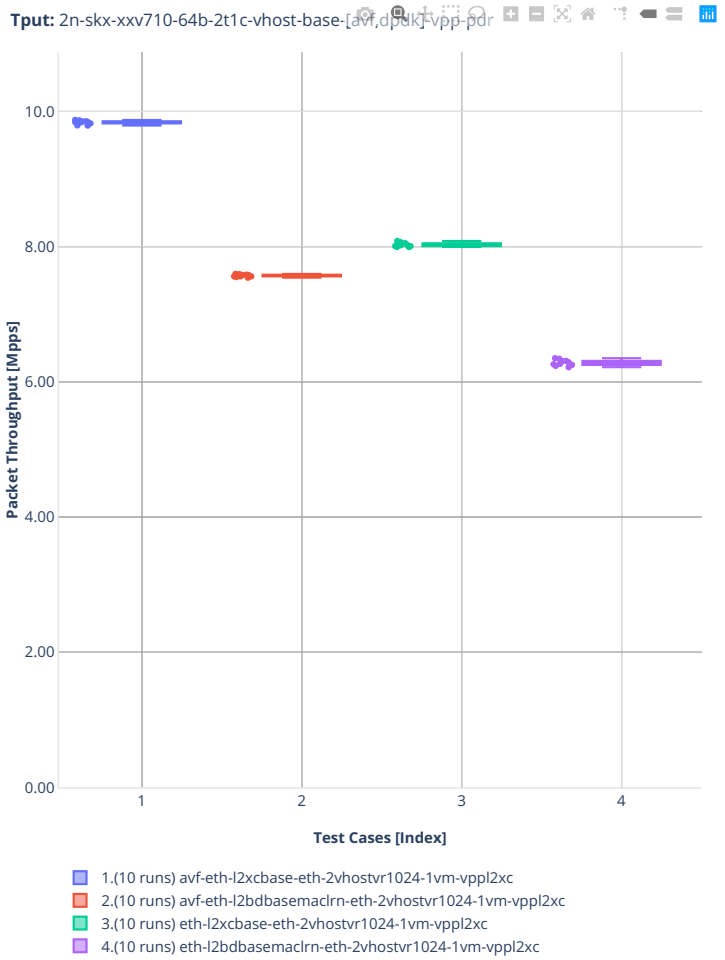
Tput: 2n-skx-xxv710-64b-2t1c-vhost-base-[avf,dpdk]-vpp-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) avf-eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
2.(10 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
3.(10 runs) eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
4.(10 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2n-clx-xxv710

64b-2t1c-vhost-base-testpmd

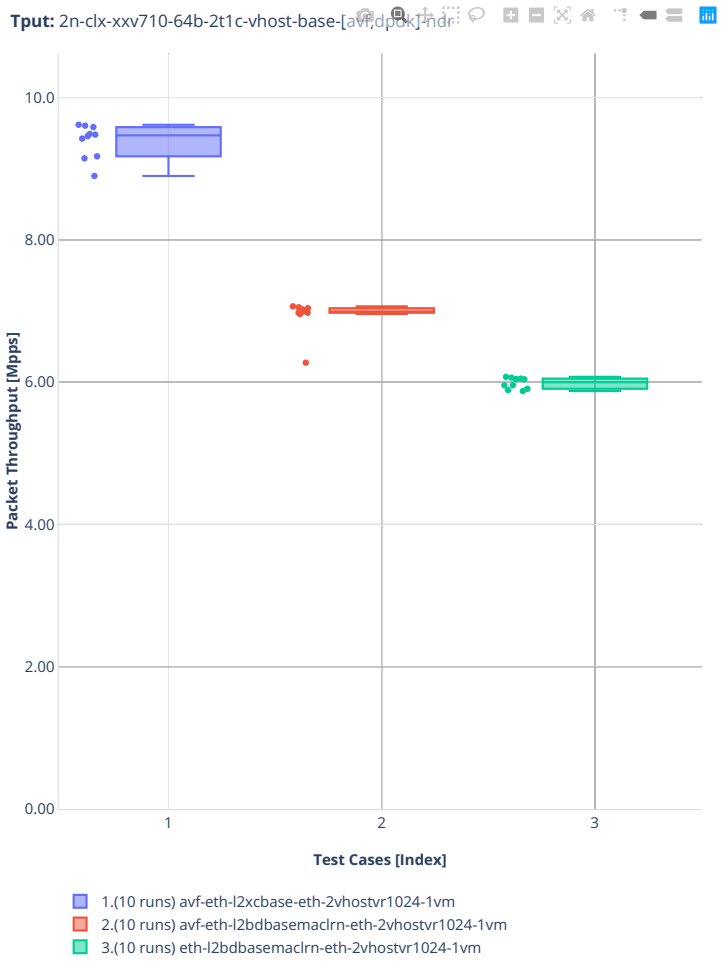

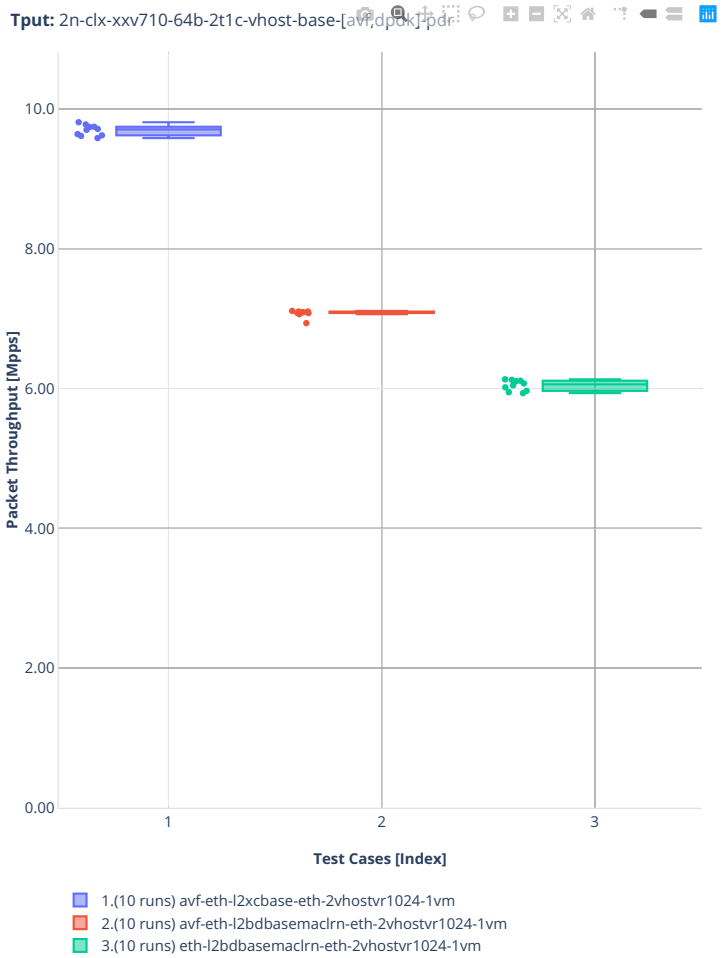
Tput: 2n-clx-xxv710-64b-2t1c-vhost-base-[avf,dpdk]-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-eth-l2xcbase-eth-2vhostvr1024-1vm
2.(10 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
3.(10 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm

64b-2t1c-vhost-base-vpp

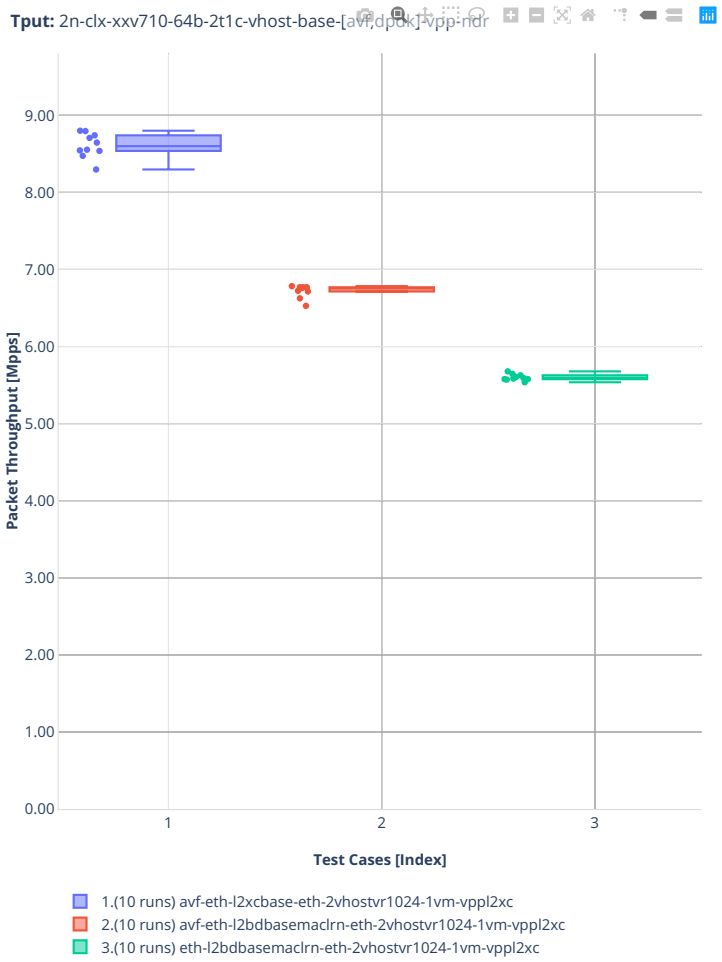

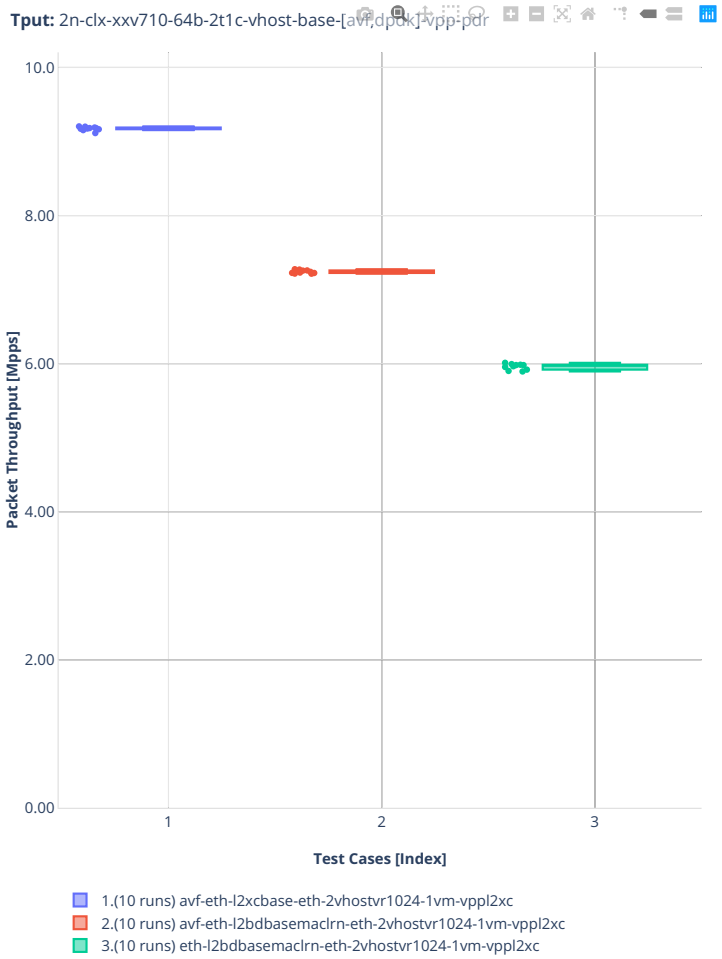
Tput: 2n-clx-xxv710-64b-2t1c-vhost-base-[avf,dpdk]-vpp-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
2.(10 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
3.(10 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2n-clx-cx556a

64b-2t1c-vhost-base-rdma-core

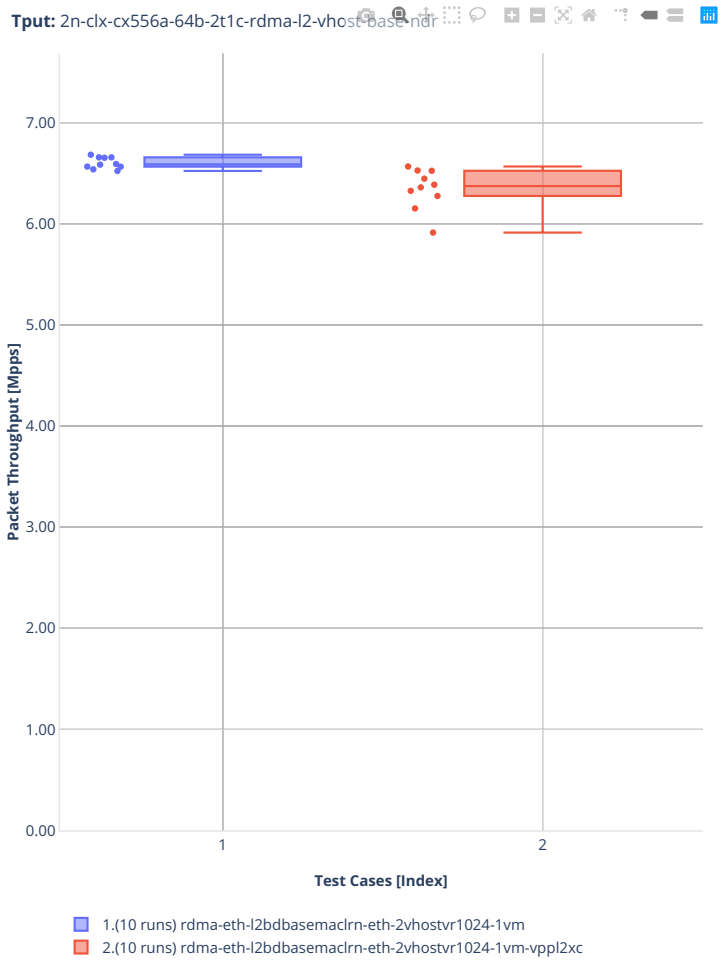

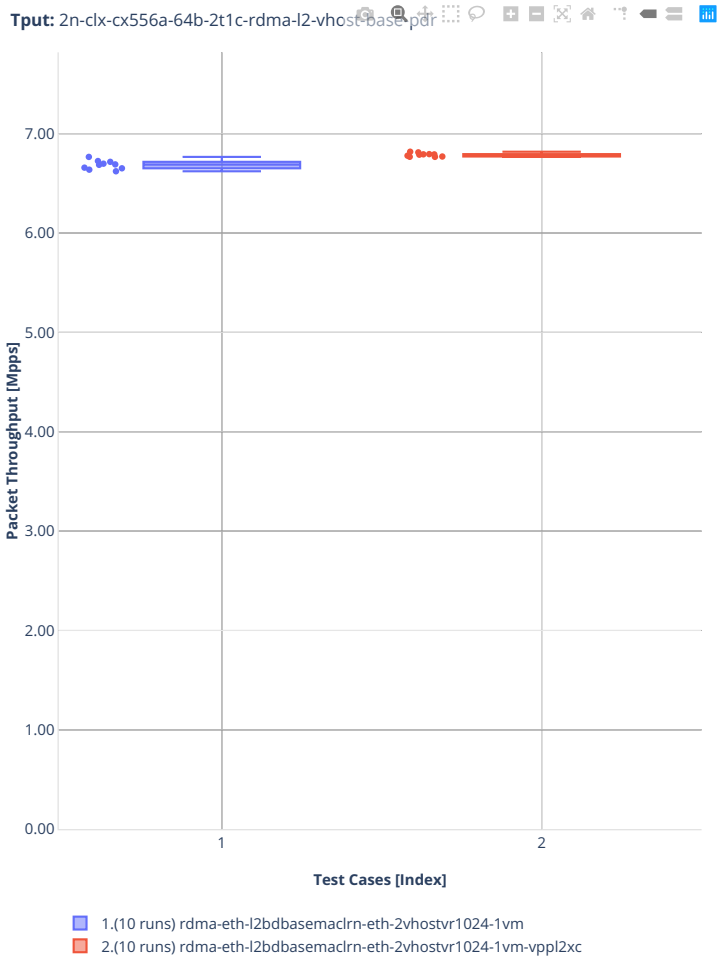
Tput: 2n-clx-cx556a-64b-2t1c-rdma-l2-vhost-base-pdr
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) rdma-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
2.(10 runs) rdma-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2n-clx-e810cq

64b-2t1c-vhost

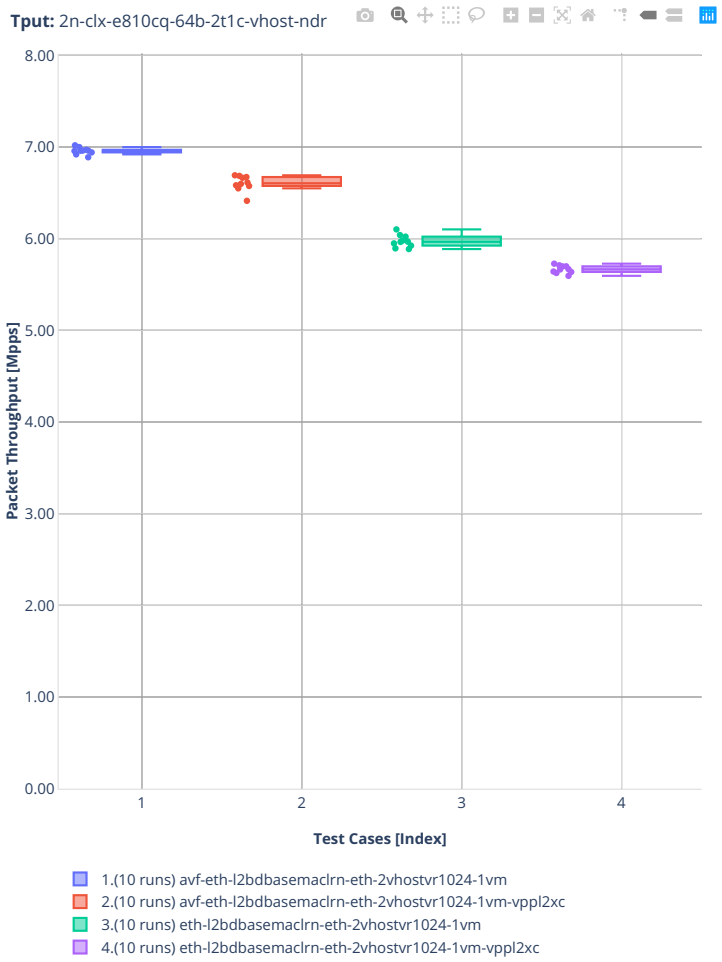

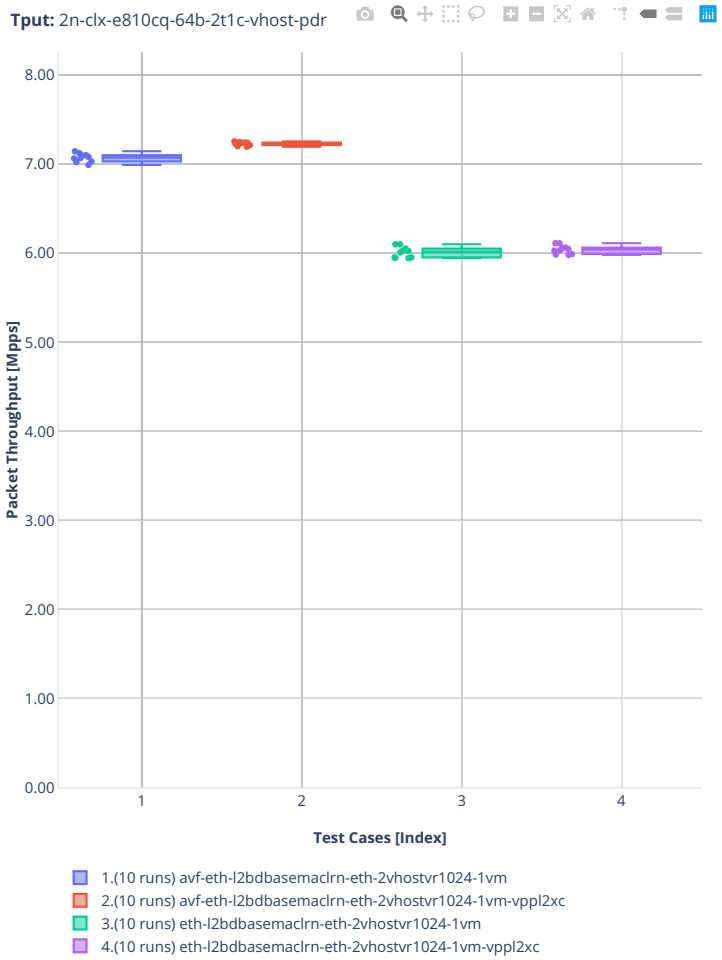
Tput: 2n-clx-e810cq-64b-2t1c-vhost-pdr
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
2.(10 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
3.(10 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
4.(10 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2n-zn2-xxv710

64b-2t1c-vhost-base-testpmd

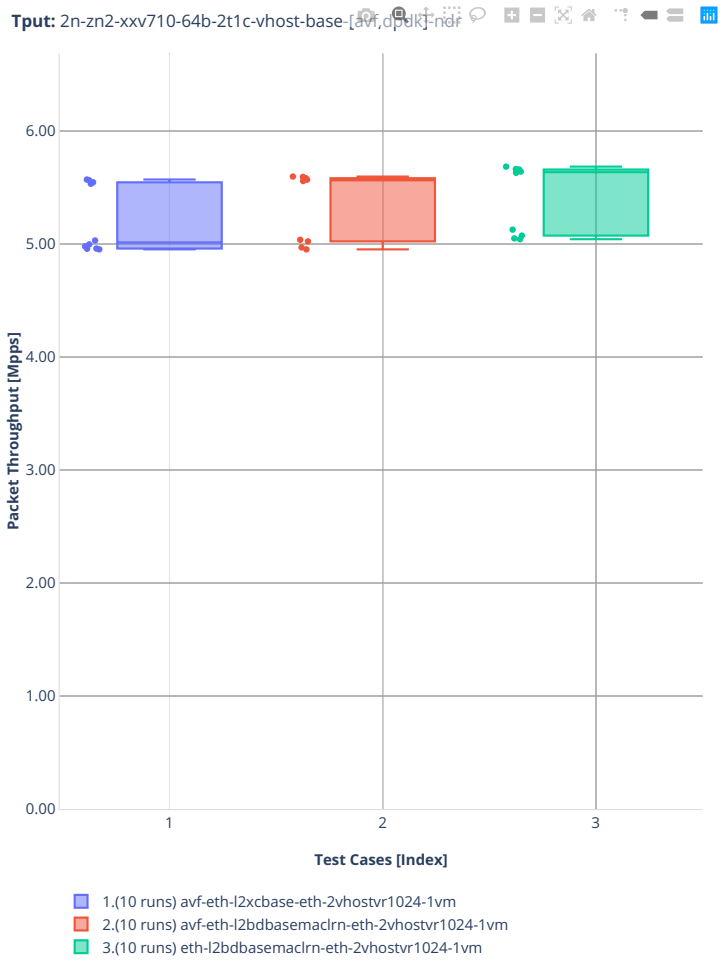

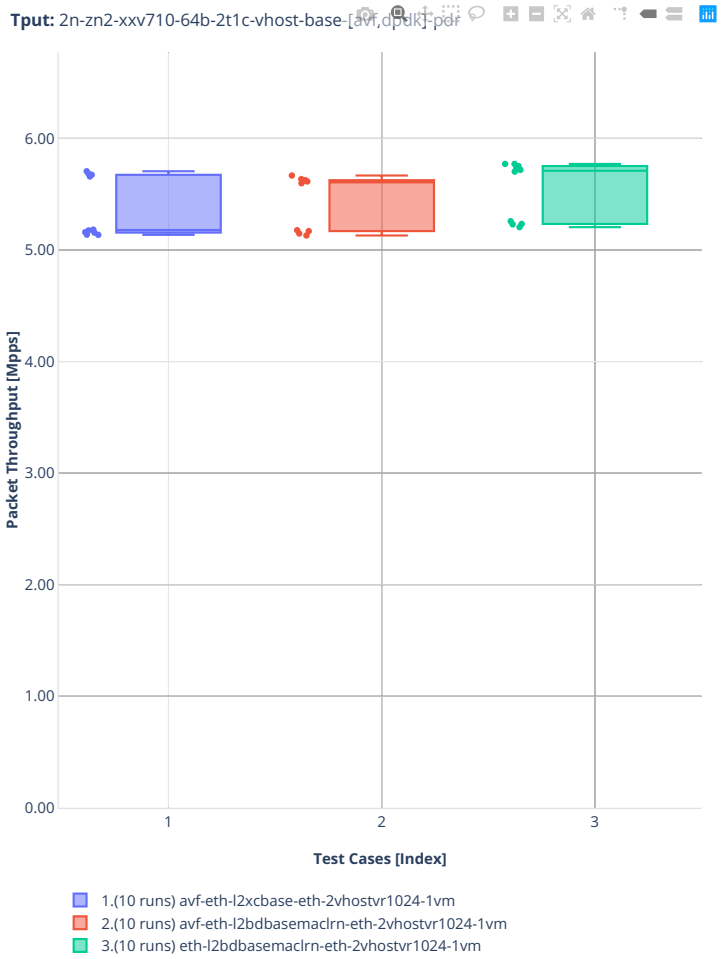
Tput: 2n-zn2-xxv710-64b-2t1c-vhost-base-[avf,dpdk]-pdr
Packet Throughput [Mpps]
6.00
5.00
4.00
3.00
2.00
1.00
0.00
1
2
3
Test Cases [Index]
1.(10 runs) avf-eth-l2xcbase-eth-2vhostvr1024-1vm
2.(10 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
3.(10 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm

64b-2t1c-vhost-base-vpp

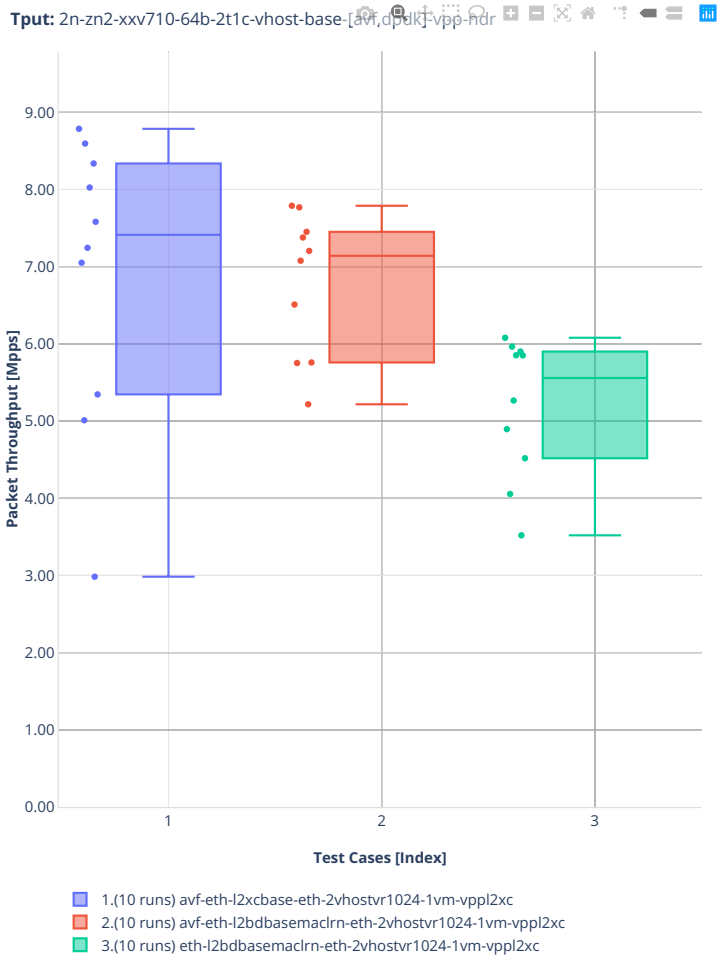

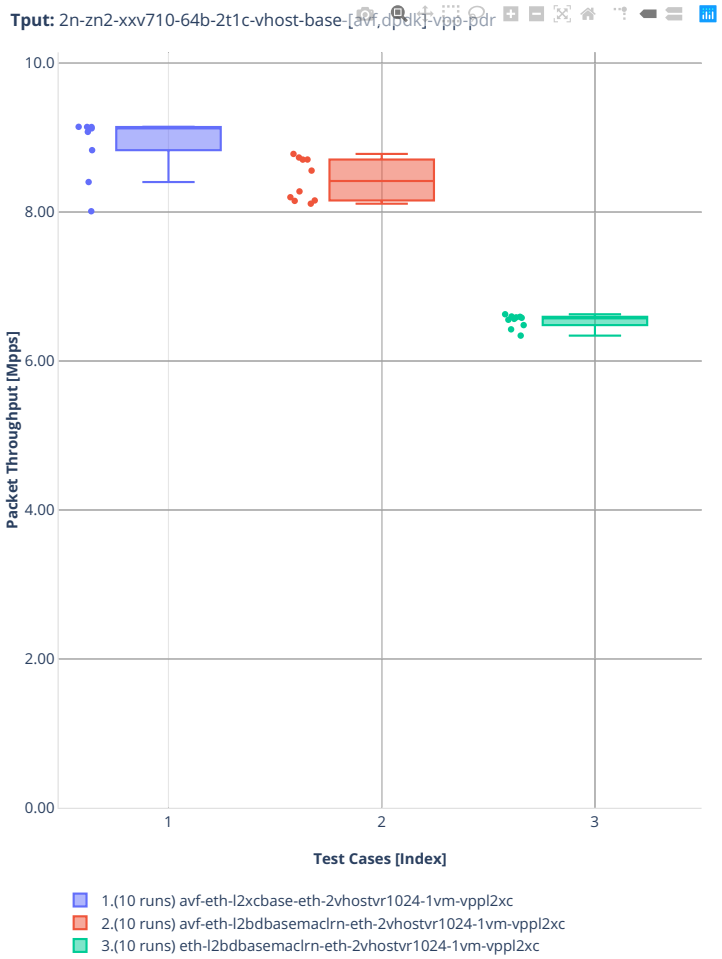
Tput: 2n-zn2-xxv710-64b-2t1c-vhost-base-[avf,dpdk]-vpp-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
2.(10 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
3.(10 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2n-zn2-cx556a

64b-2t1c-vhost-base-rdma-core

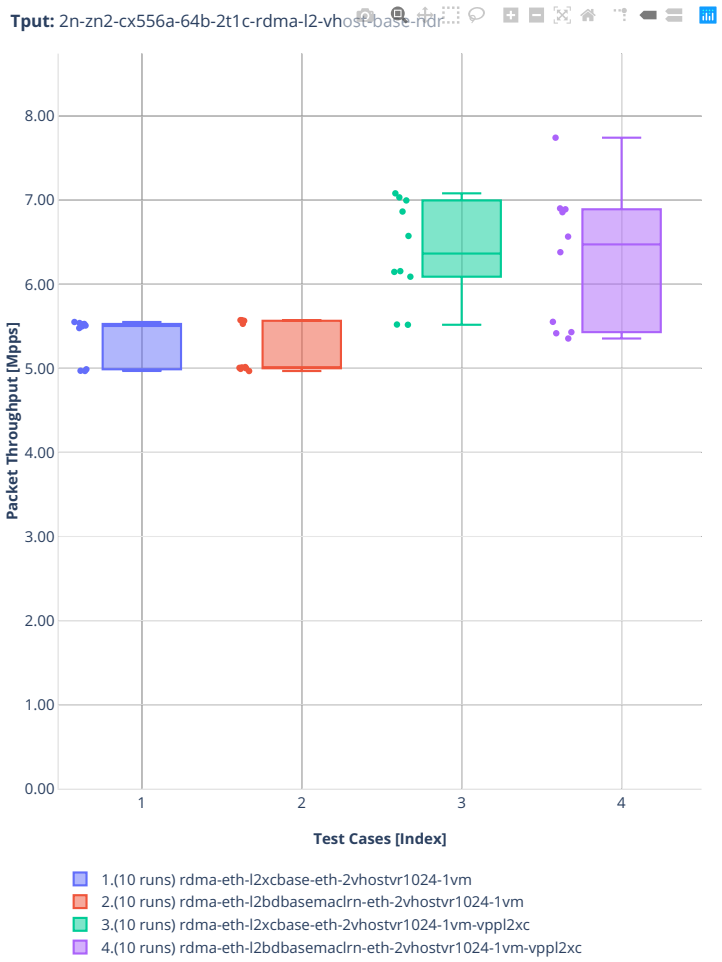

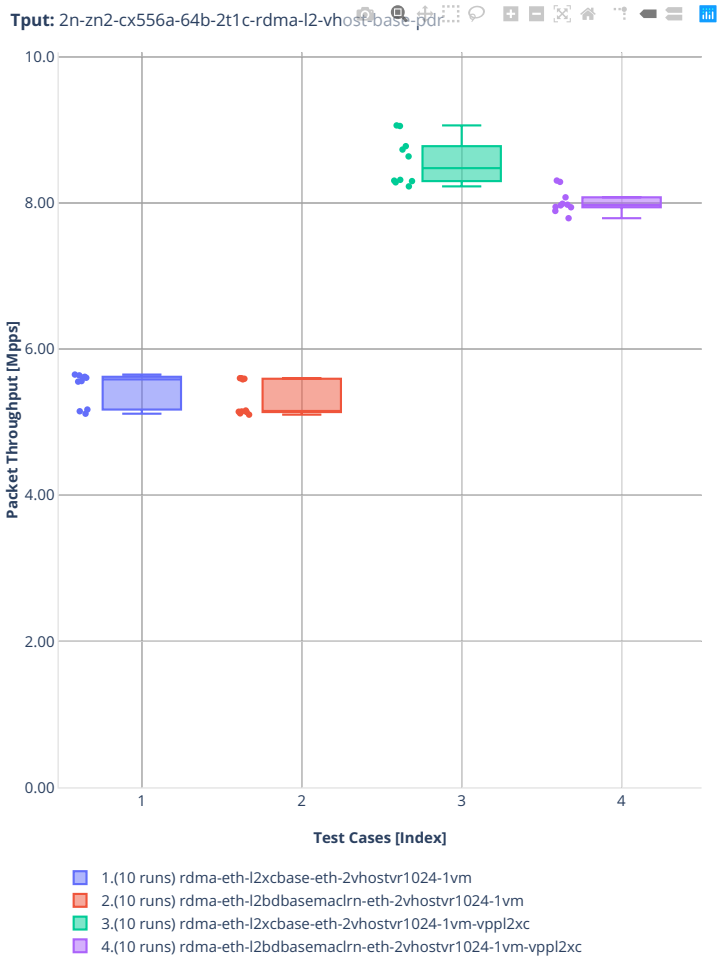
Tput: 2n-zn2-cx556a-64b-2t1c-rdma-l2-vhost-base-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
Test Cases [Index]
1.(10 runs) rdma-eth-l2xcbase-eth-2vhostvr1024-1vm
2.(10 runs) rdma-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
3.(10 runs) rdma-eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
4.(10 runs) rdma-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

3n-alt-xl710

64b-1t1c-vhost-base

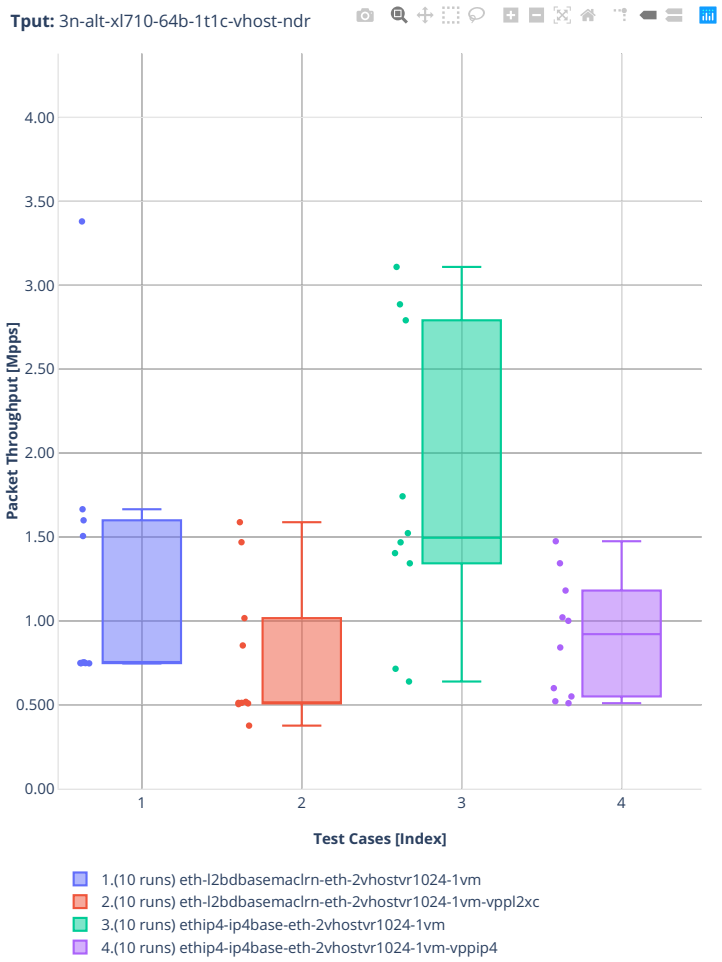

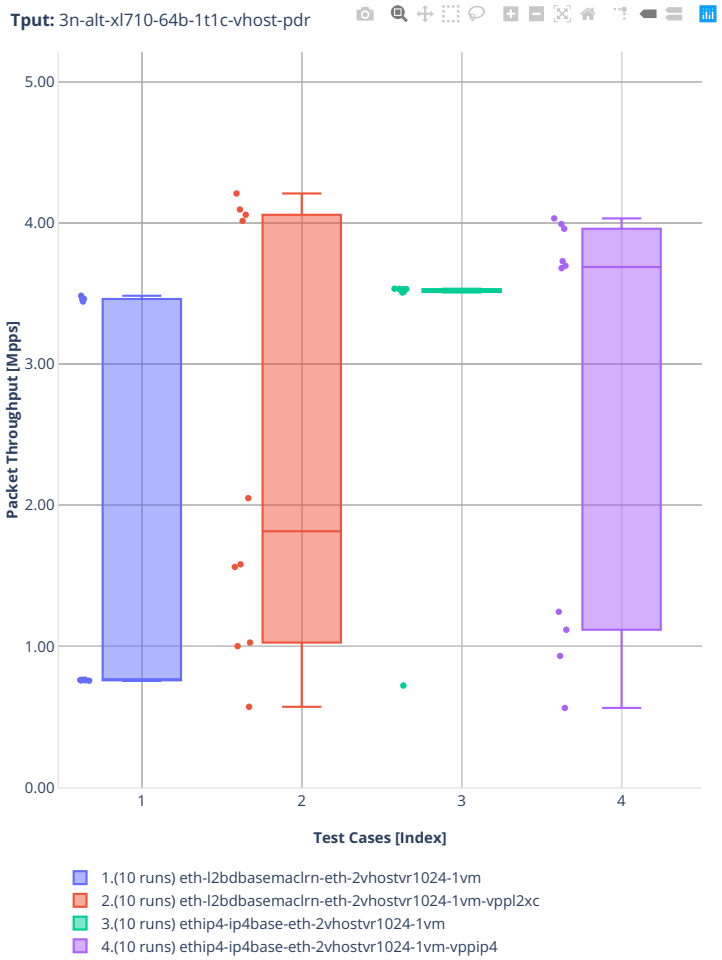
Tput: 3n-alt-xl710-64b-1t1c-vhost-pdr
Packet Throughput [Mpps]
5.00
4.00
3.00
2.00
1.00
0.00
1
2
3
4
Test Cases [Index]
1.(10 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
2.(10 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
3.(10 runs) ethip4-ip4base-eth-2vhostvr1024-1vm
4.(10 runs) ethip4-ip4base-eth-2vhostvr1024-1vm-vppip4

3n-tsh-x520

64b-1t1c-vhost-base-ixgbe

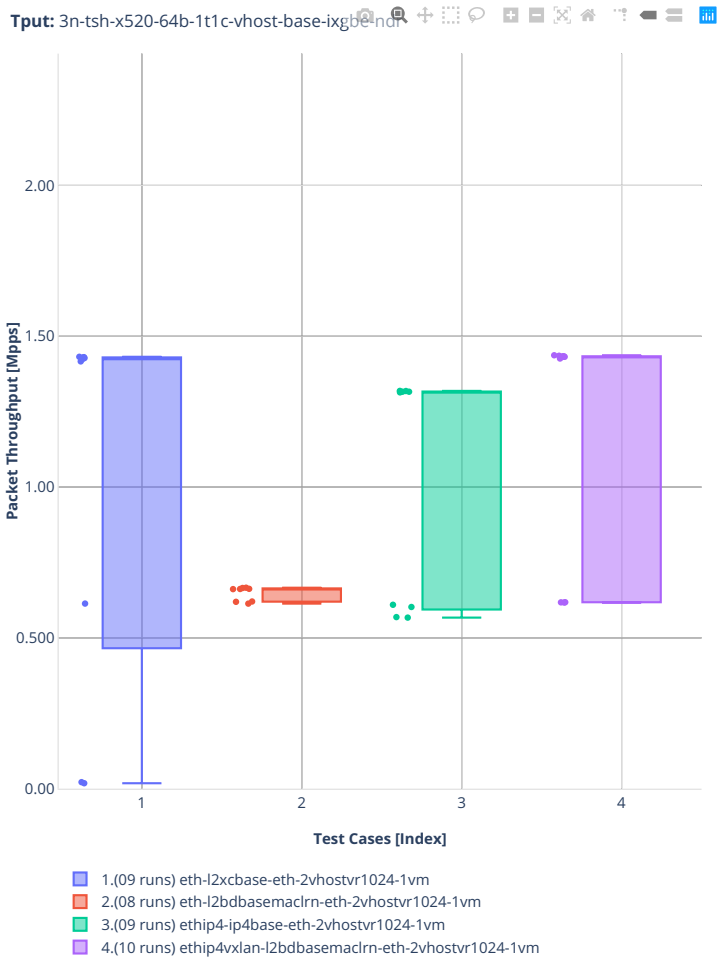

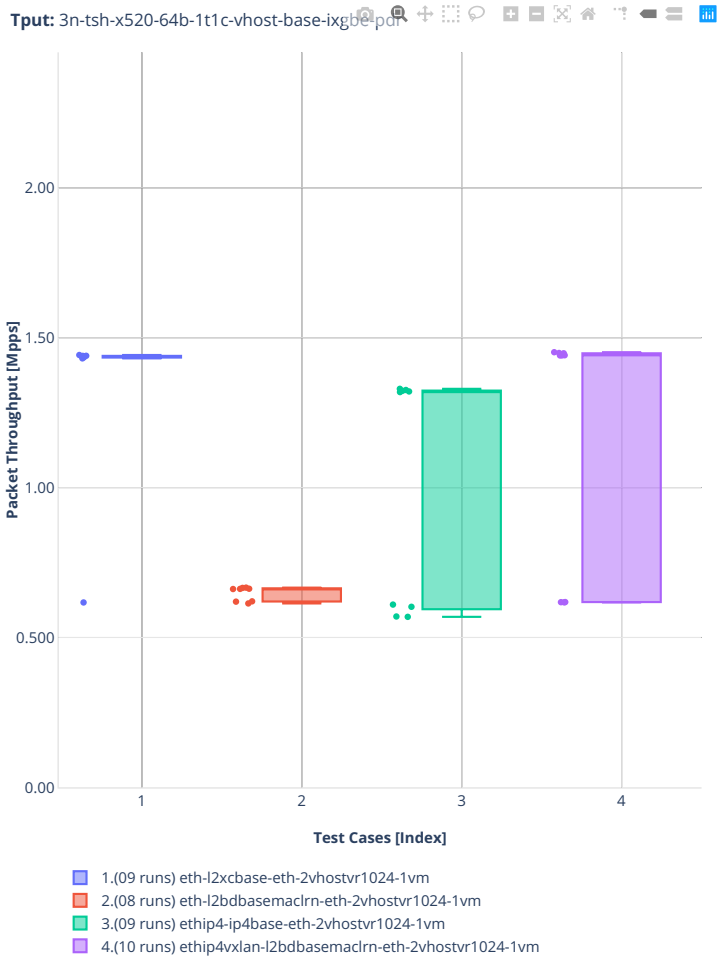
Tput: 3n-tsh-x520-64b-1t1c-vhost-base-ixgbe-pdr
Packet Throughput [Mpps]
2.00
1.50
1.00
0.500
0.00
1
2
3
4
Test Cases [Index]
1.(09 runs) eth-l2xcbase-eth-2vhostvr1024-1vm
2.(08 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
3.(09 runs) ethip4-ip4base-eth-2vhostvr1024-1vm
4.(10 runs) ethip4vxlan-l2bdbasemaclrn-eth-2vhostvr1024-1vm

64b-1t1c-vhost-base-ixgbe-vppl2xc

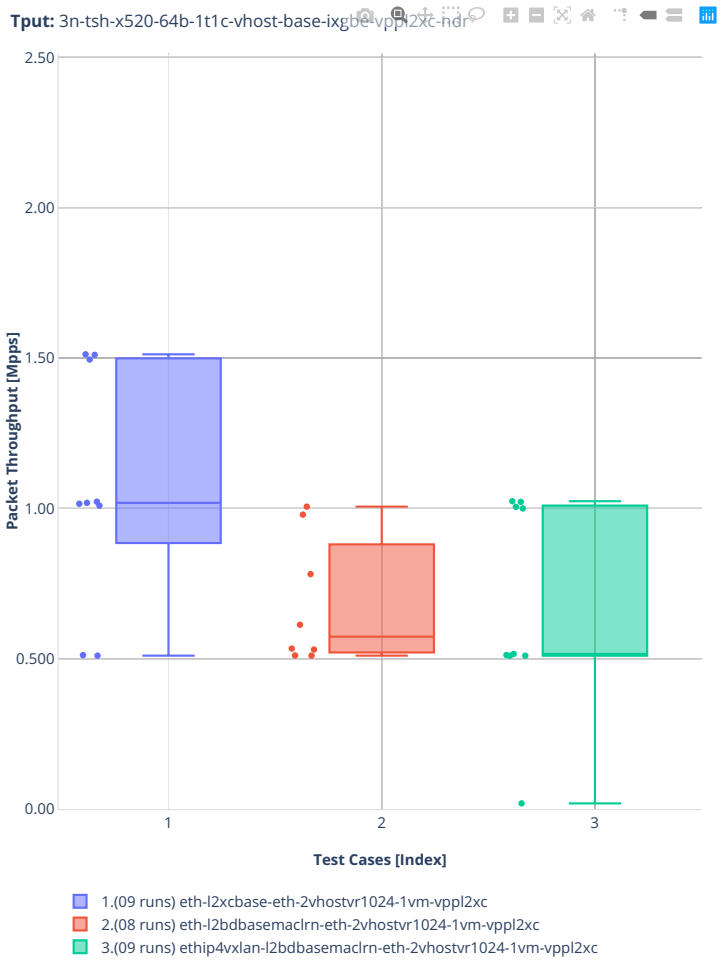

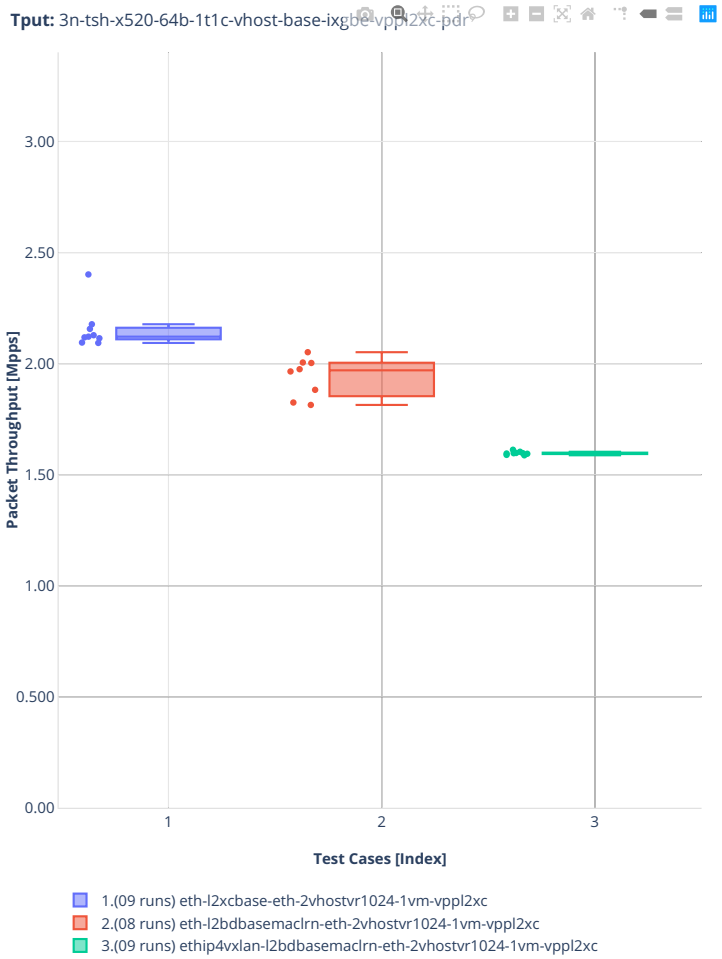
Tput: 3n-tsh-x520-64b-1t1c-vhost-base-ixgbe-vppl2xc-pdr
Packet Throughput [Mpps]
3.00
2.50
2.00
1.50
1.00
0.500
0.00
1
2
3
Test Cases [Index]
1.(09 runs) eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
2.(08 runs) eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
3.(09 runs) ethip4vxlan-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2.3.8 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[116].

[116] https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls2206

2n-icx-xxv710

64b-2t1c-memif-base

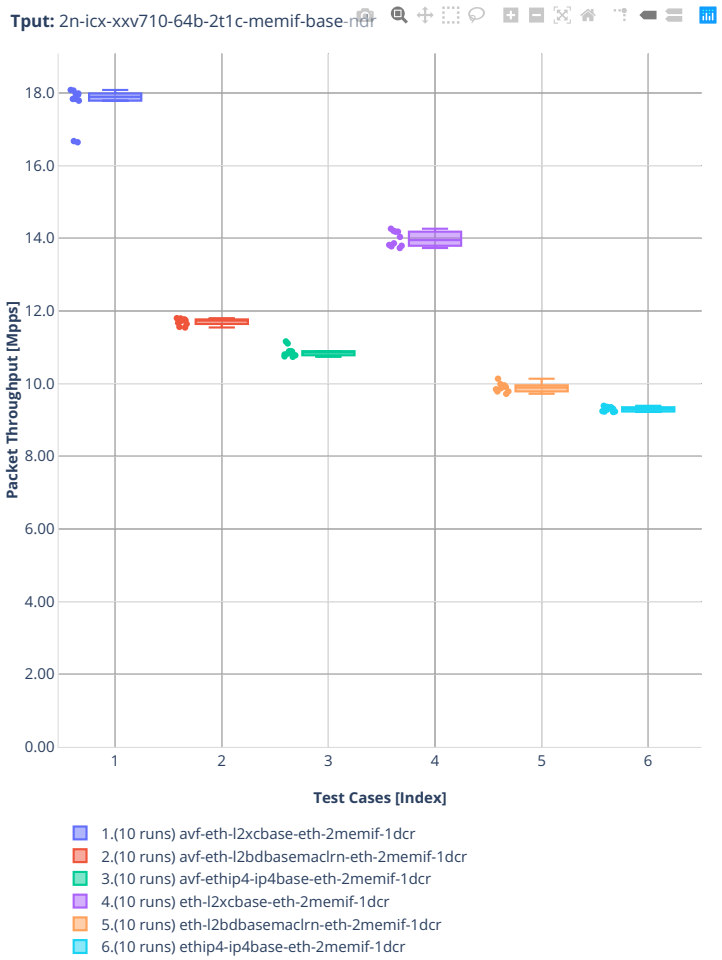

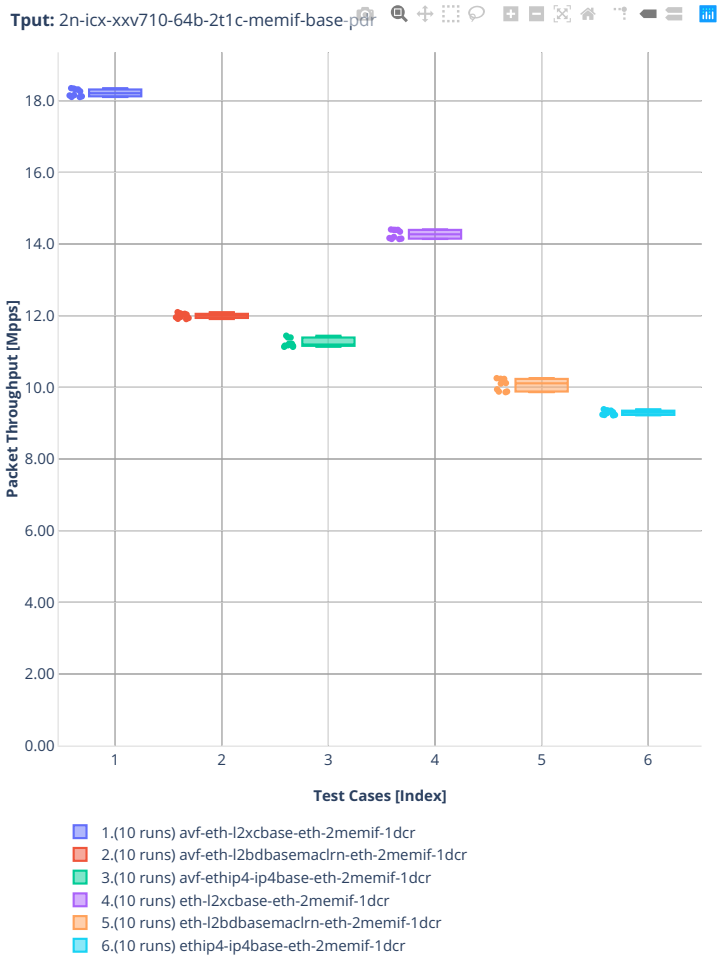
Tput: 2n-icx-xxv710-64b-2t1c-memif-base-pdr
Packet Throughput [Mpps]
18.0
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-eth-l2xcbase-eth-2memif-1dcr
2.(10 runs) avf-eth-l2bdbasemaclrn-eth-2memif-1dcr
3.(10 runs) avf-ethip4-ip4base-eth-2memif-1dcr
4.(10 runs) eth-l2xcbase-eth-2memif-1dcr
5.(10 runs) eth-l2bdbasemaclrn-eth-2memif-1dcr
6.(10 runs) ethip4-ip4base-eth-2memif-1dcr

2n-skx-xxv710

64b-2t1c-memif-base-avf

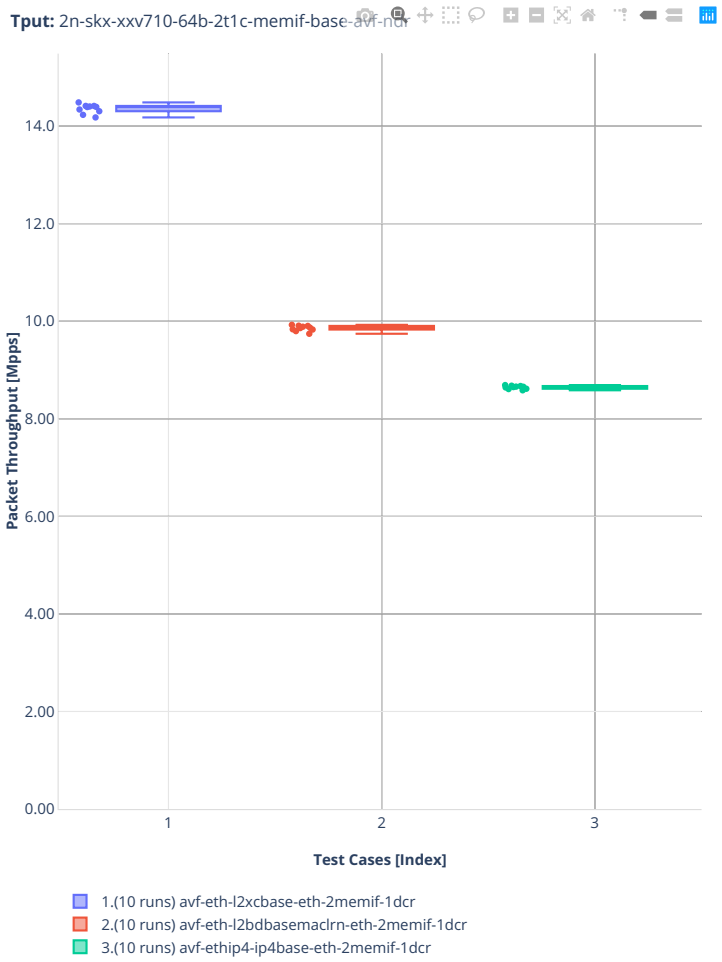

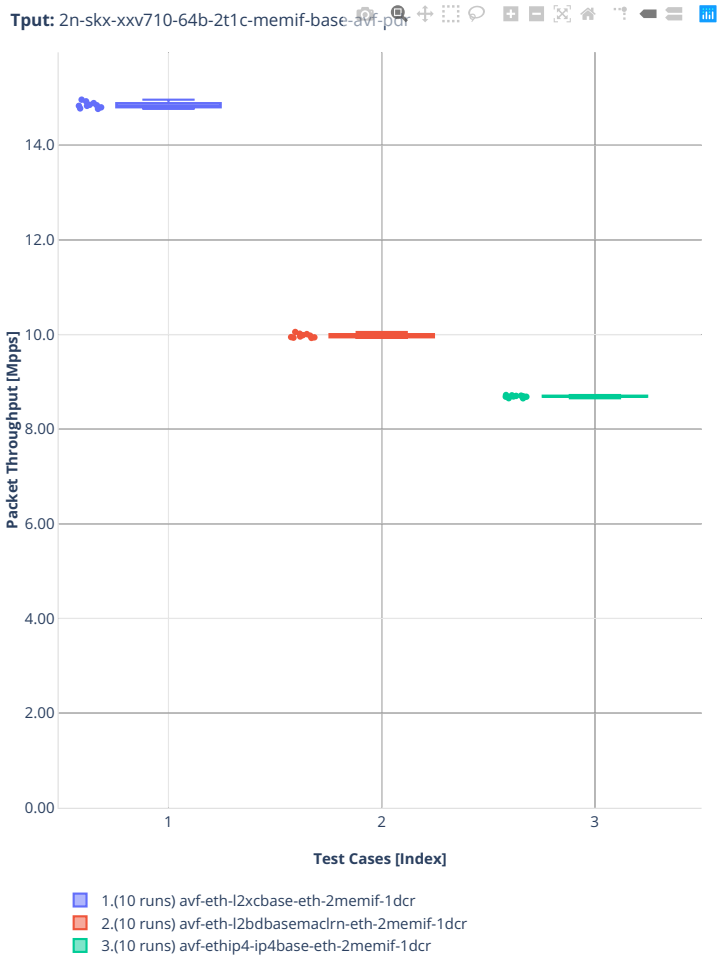
Tput: 2n-skx-xxv710-64b-2t1c-memif-base-avf-pdr
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-eth-l2xcbase-eth-2memif-1dcr
2.(10 runs) avf-eth-l2bdbasemaclrn-eth-2memif-1dcr
3.(10 runs) avf-ethip4-ip4base-eth-2memif-1dcr

64b-2t1c-memif-base-dpdk

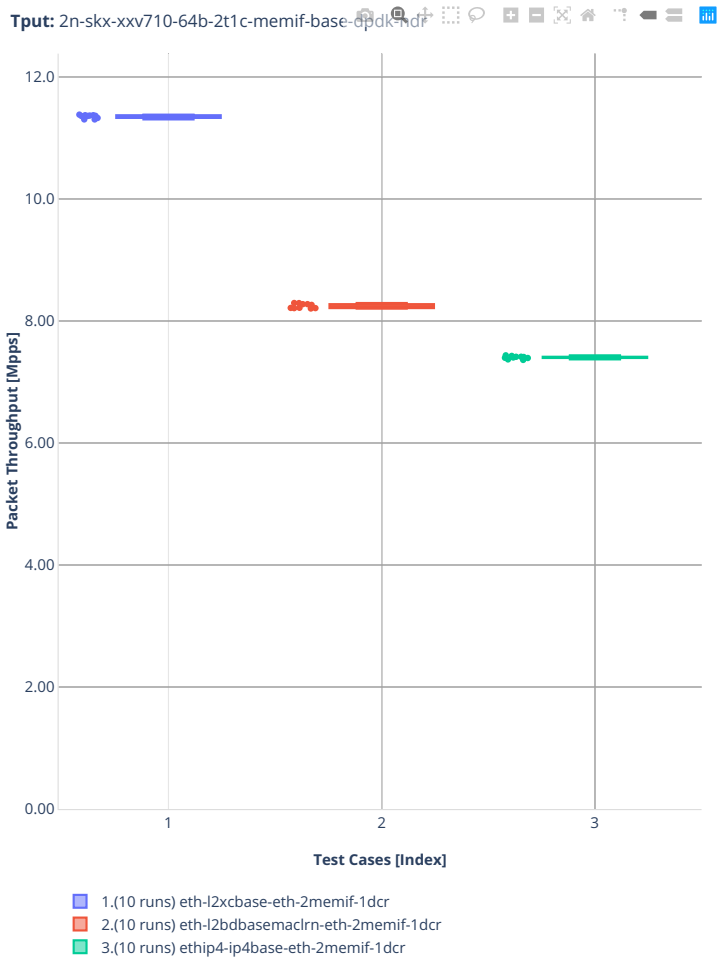

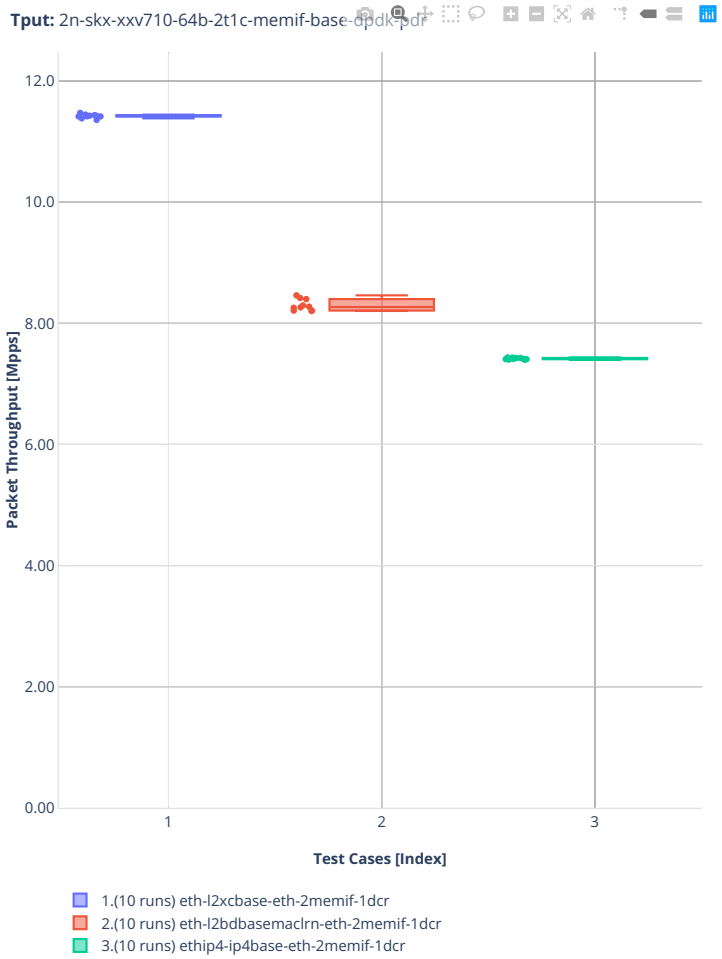
Tput: 2n-skx-xxv710-64b-2t1c-memif-base-dpdk-pdr
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) eth-l2xcbase-eth-2memif-1dcr
2.(10 runs) eth-l2bdbasemaclrn-eth-2memif-1dcr
3.(10 runs) ethip4-ip4base-eth-2memif-1dcr

2n-clx-xxv710

64b-2t1c-memif-base-avf

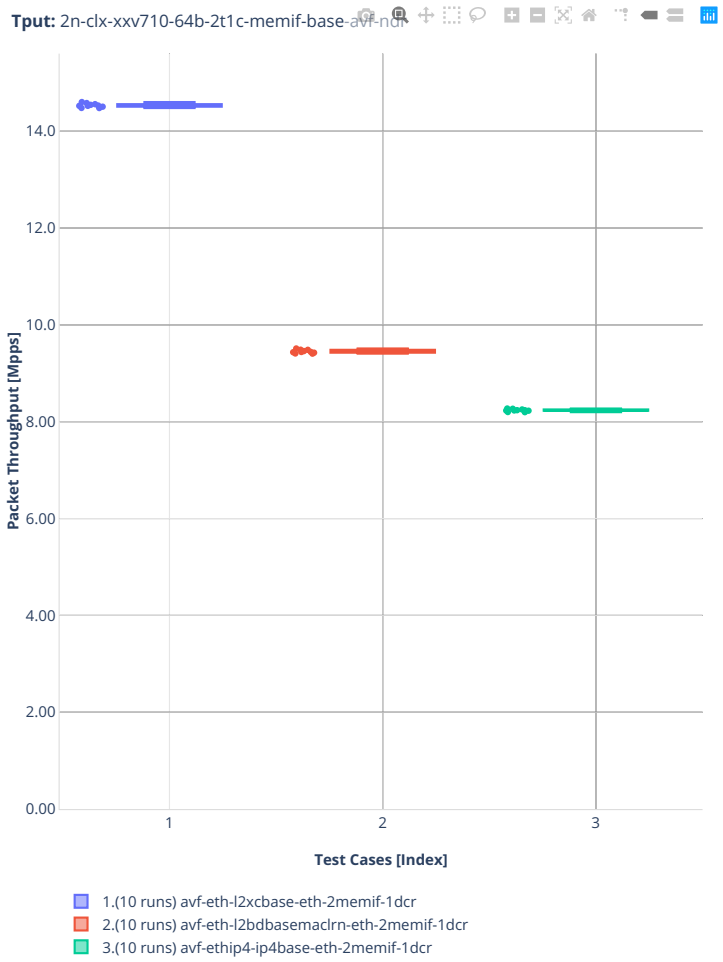

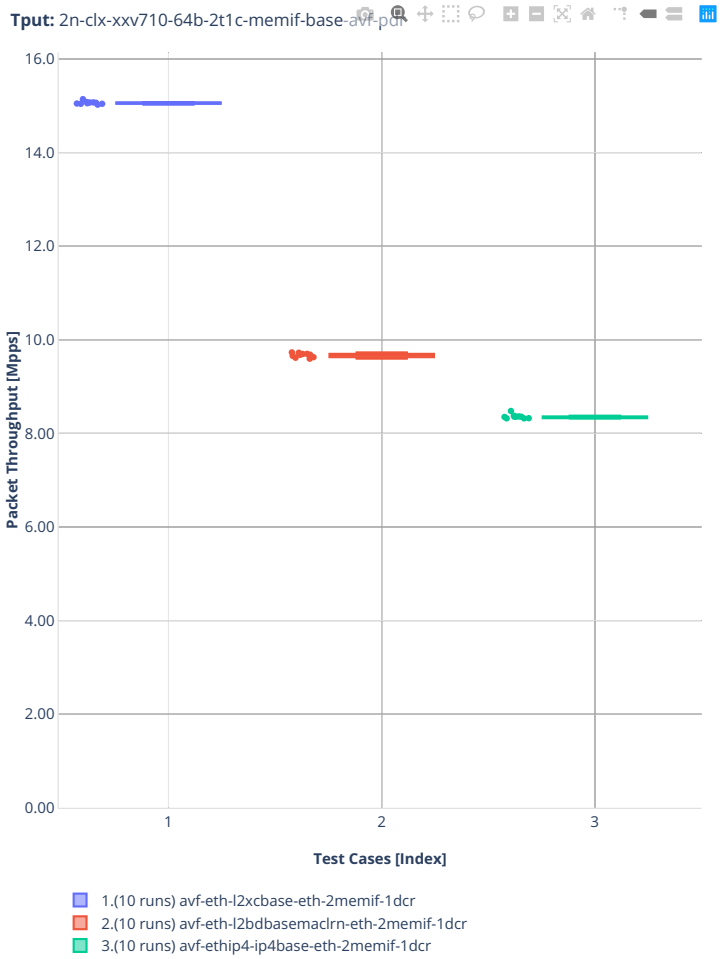
Tput: 2n-clx-xxv710-64b-2t1c-memif-base-avf-pdr
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-eth-l2xcbase-eth-2memif-1dcr
2.(10 runs) avf-eth-l2bdbasemaclrn-eth-2memif-1dcr
3.(10 runs) avf-ethip4-ip4base-eth-2memif-1dcr

64b-2t1c-memif-base-dpdk

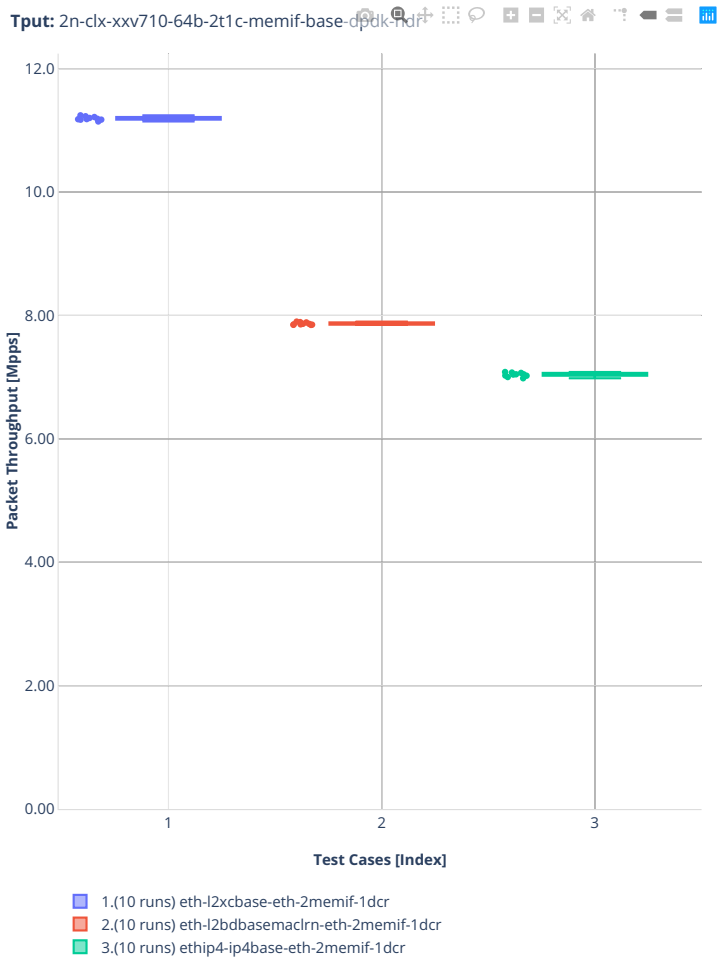

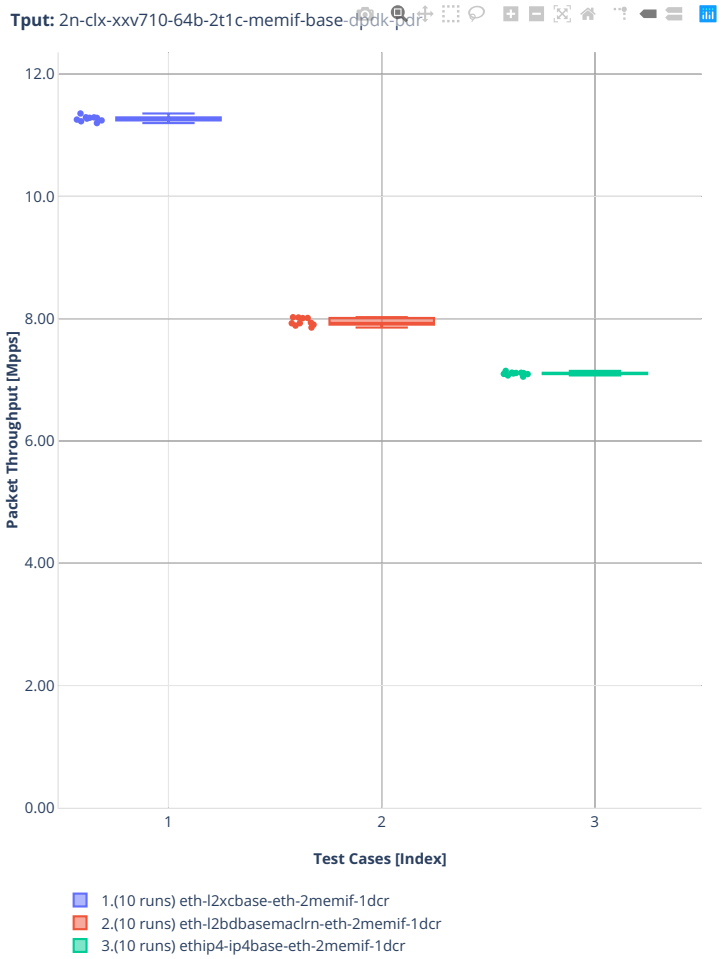
Tput: 2n-clx-xxv710-64b-2t1c-memif-base-dpdk-pdr
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) eth-l2xcbase-eth-2memif-1dcr
2.(10 runs) eth-l2bdbasemaclrn-eth-2memif-1dcr
3.(10 runs) ethip4-ip4base-eth-2memif-1dcr

2n-clx-cx556a

64b-2t1c-memif-base-rdma-core

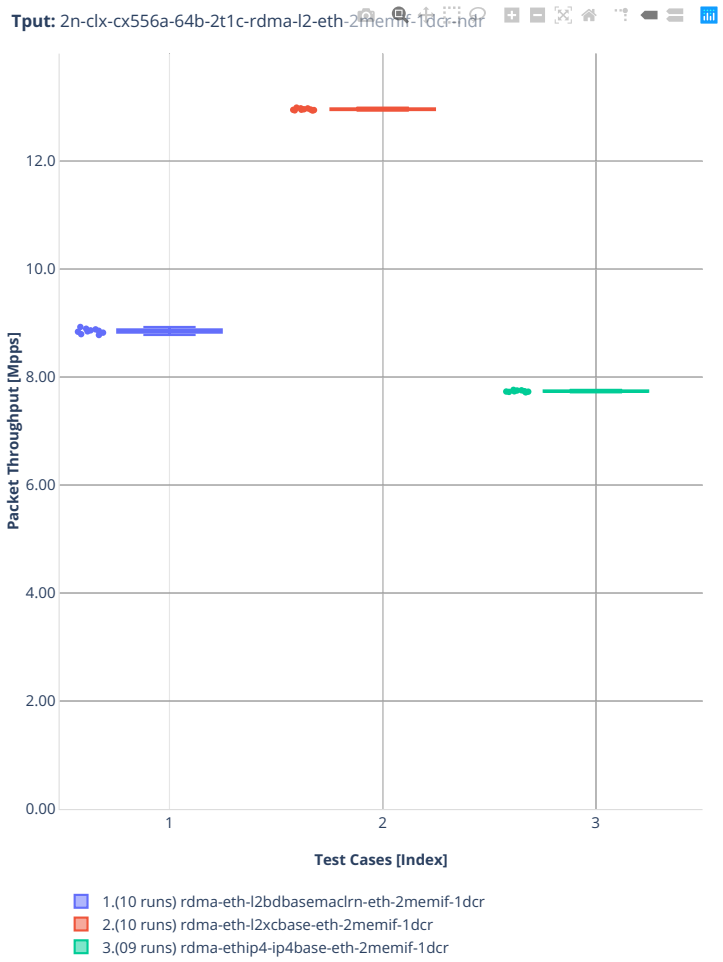

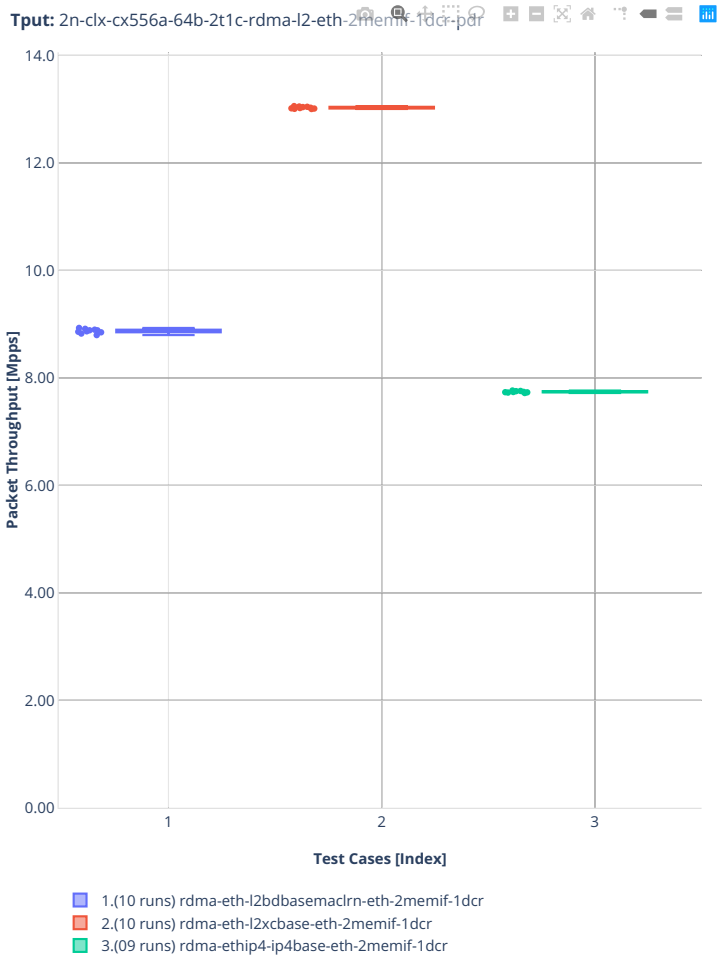
Tput: 2n-clx-cx556a-64b-2t1c-rdma-l2-eth-2memif-1dcr-pdr
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) rdma-eth-l2bdbasemaclrn-eth-2memif-1dcr
2.(10 runs) rdma-eth-l2xcbase-eth-2memif-1dcr
3.(09 runs) rdma-ethip4-ip4base-eth-2memif-1dcr

2n-clx-e810cq

64b-2t1c-memif-base

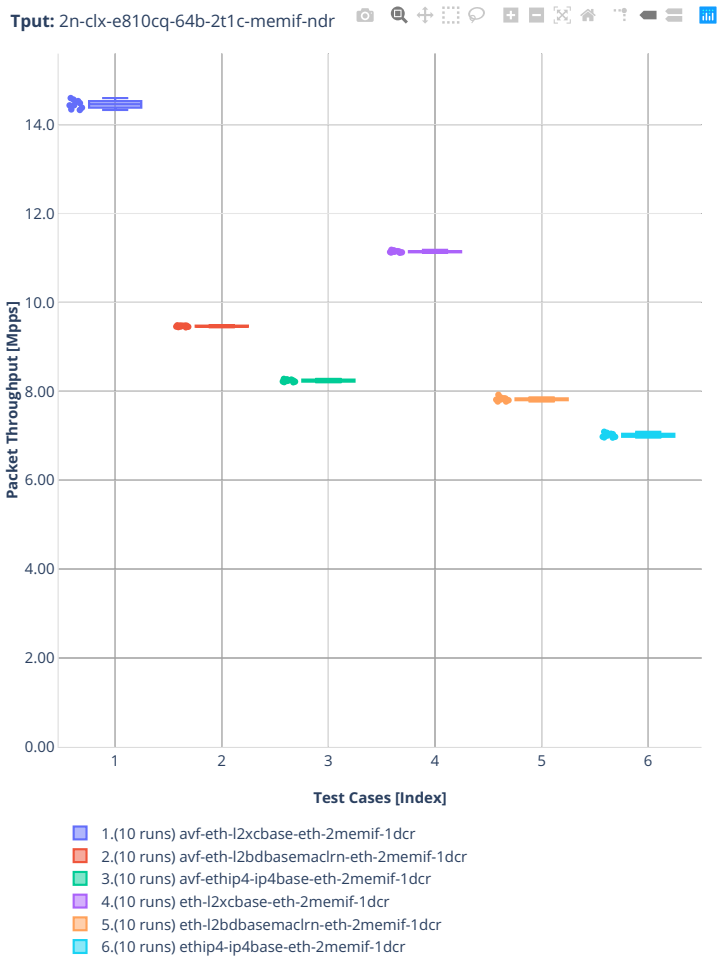

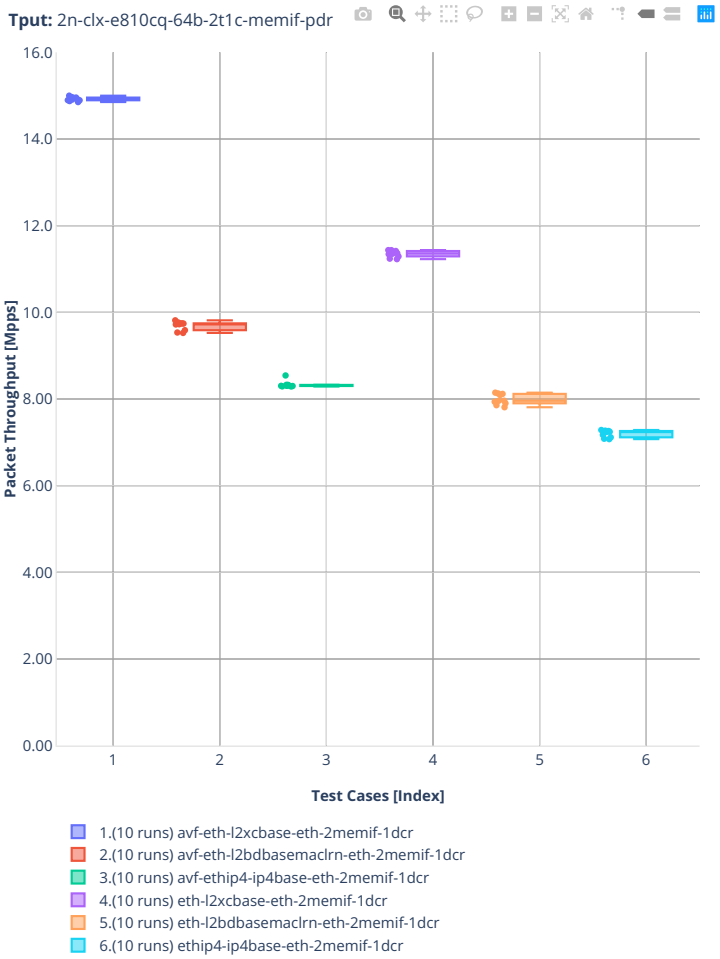
Tput: 2n-clx-e810cq-64b-2t1c-memif-pdr
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) avf-eth-l2xcbase-eth-2memif-1dcr
2.(10 runs) avf-eth-l2bdbasemaclrn-eth-2memif-1dcr
3.(10 runs) avf-ethip4-ip4base-eth-2memif-1dcr
4.(10 runs) eth-l2xcbase-eth-2memif-1dcr
5.(10 runs) eth-l2bdbasemaclrn-eth-2memif-1dcr
6.(10 runs) ethip4-ip4base-eth-2memif-1dcr

2n-tx2-xl710

64b-1t1c-memif-base-dpdk

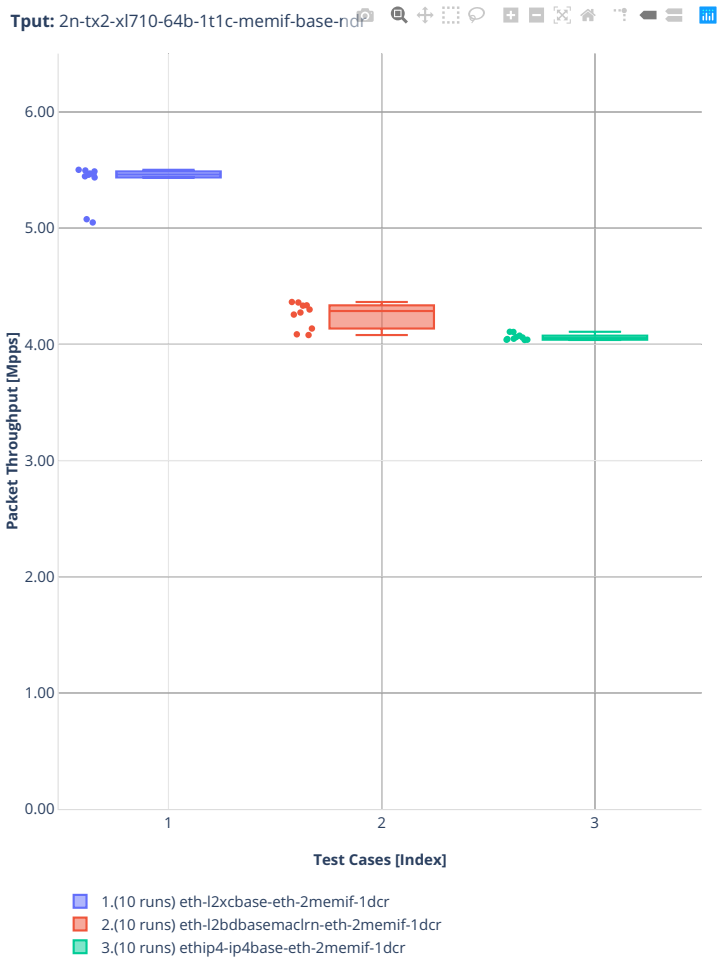

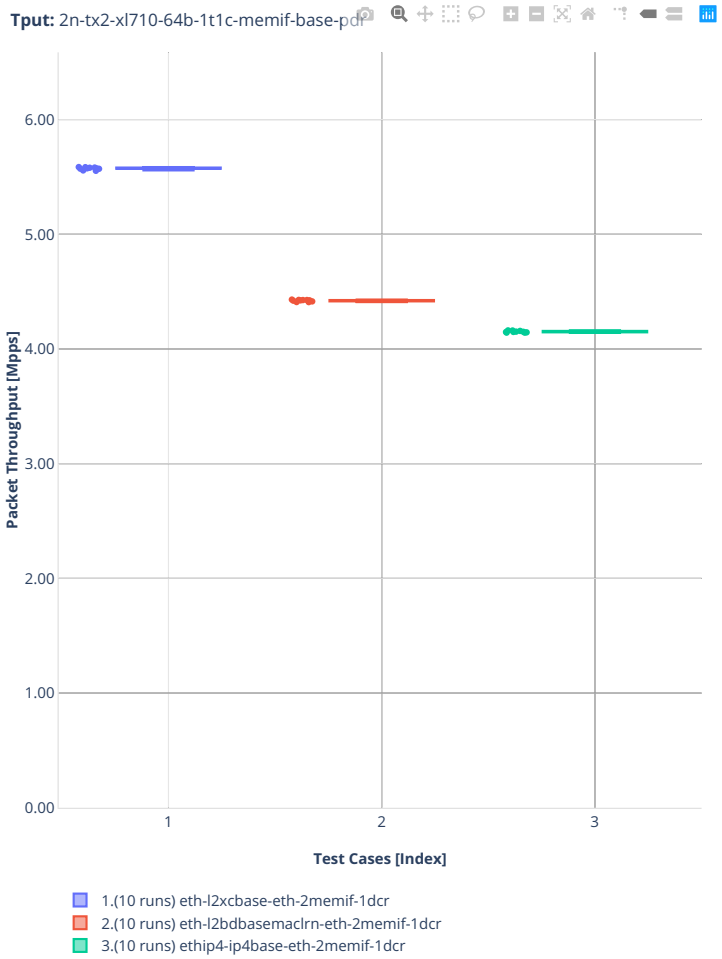
Tput: 2n-tx2-xl710-64b-1t1c-memif-base-pdr
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) eth-l2xcbase-eth-2memif-1dcr
2.(10 runs) eth-l2bdbasemaclrn-eth-2memif-1dcr
3.(10 runs) ethip4-ip4base-eth-2memif-1dcr

2n-zn2-xxv710

64b-2t1c-memif-base-avf

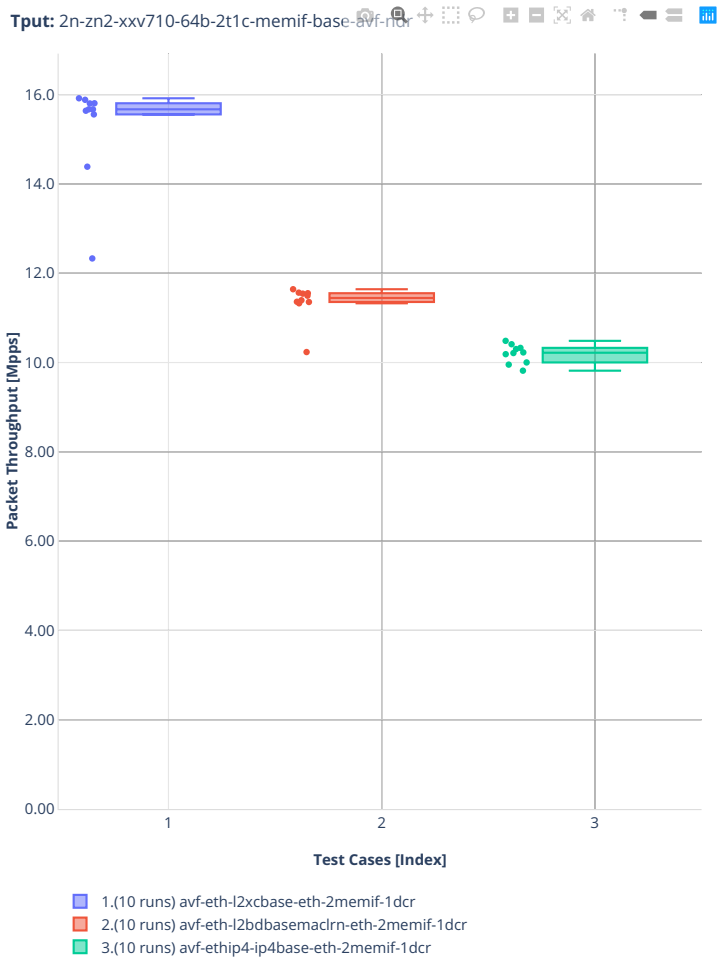

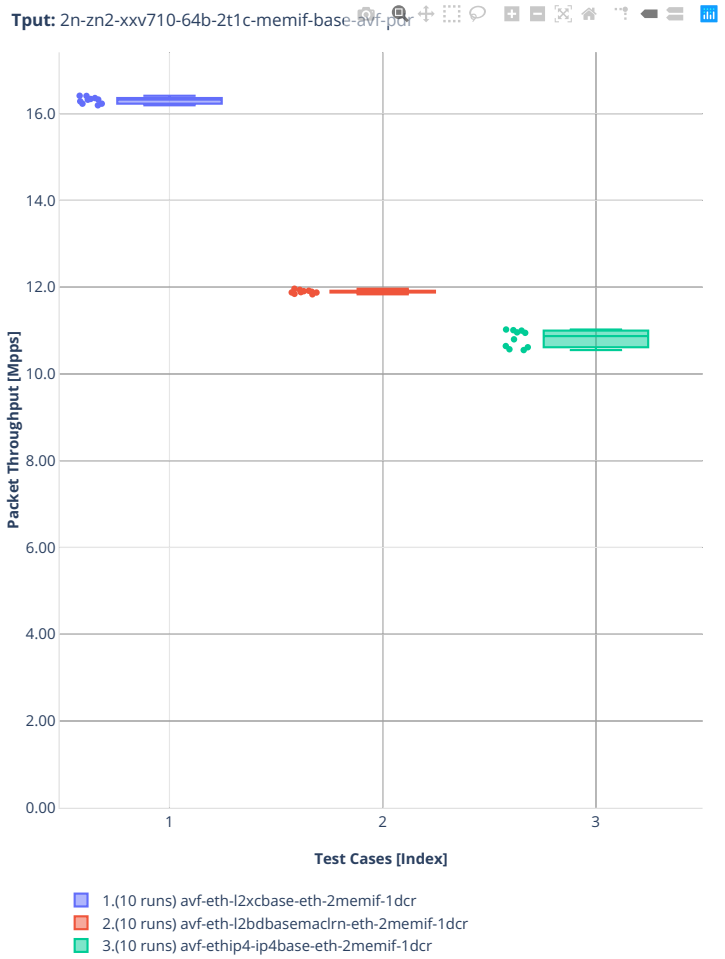
Tput: 2n-zn2-xxv710-64b-2t1c-memif-base-avf-pdr
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) avf-eth-l2xcbase-eth-2memif-1dcr
2.(10 runs) avf-eth-l2bdbasemaclrn-eth-2memif-1dcr
3.(10 runs) avf-ethip4-ip4base-eth-2memif-1dcr

64b-2t1c-memif-base-dpdk

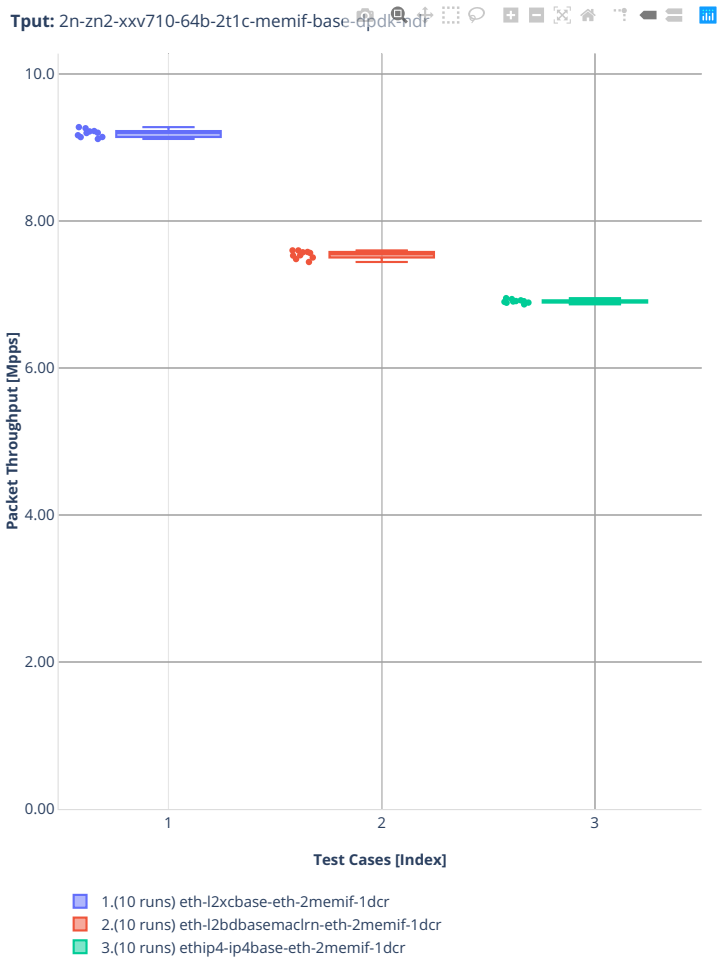

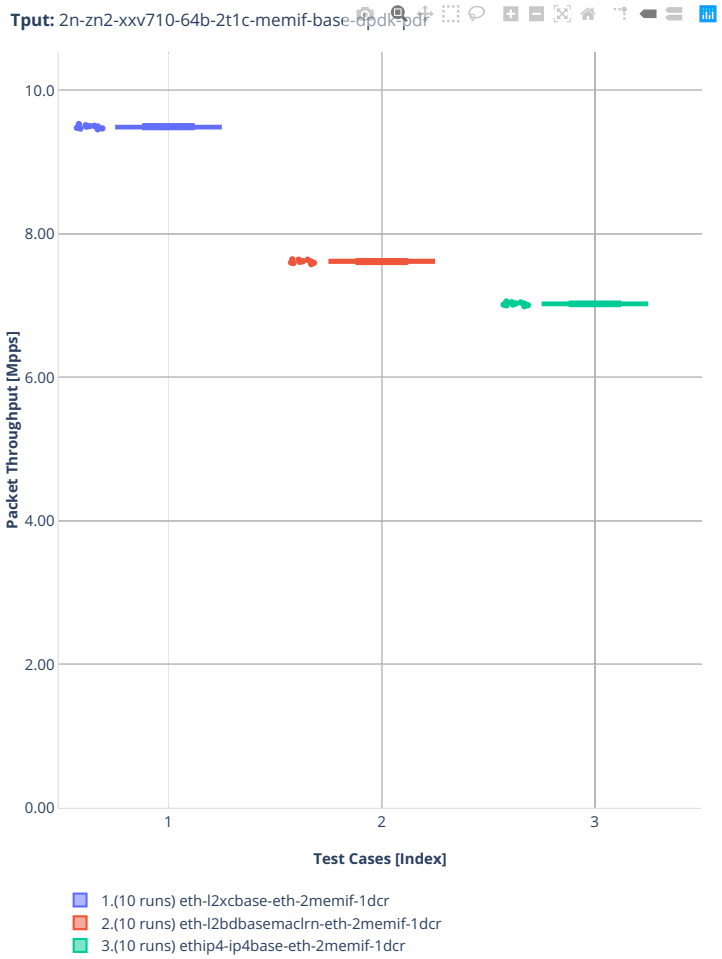
Tput: 2n-zn2-xxv710-64b-2t1c-memif-base-dpdk-pdr
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) eth-l2xcbase-eth-2memif-1dcr
2.(10 runs) eth-l2bdbasemaclrn-eth-2memif-1dcr
3.(10 runs) ethip4-ip4base-eth-2memif-1dcr

2n-zn2-cx556a

64b-2t1c-memif-base-rdma-core

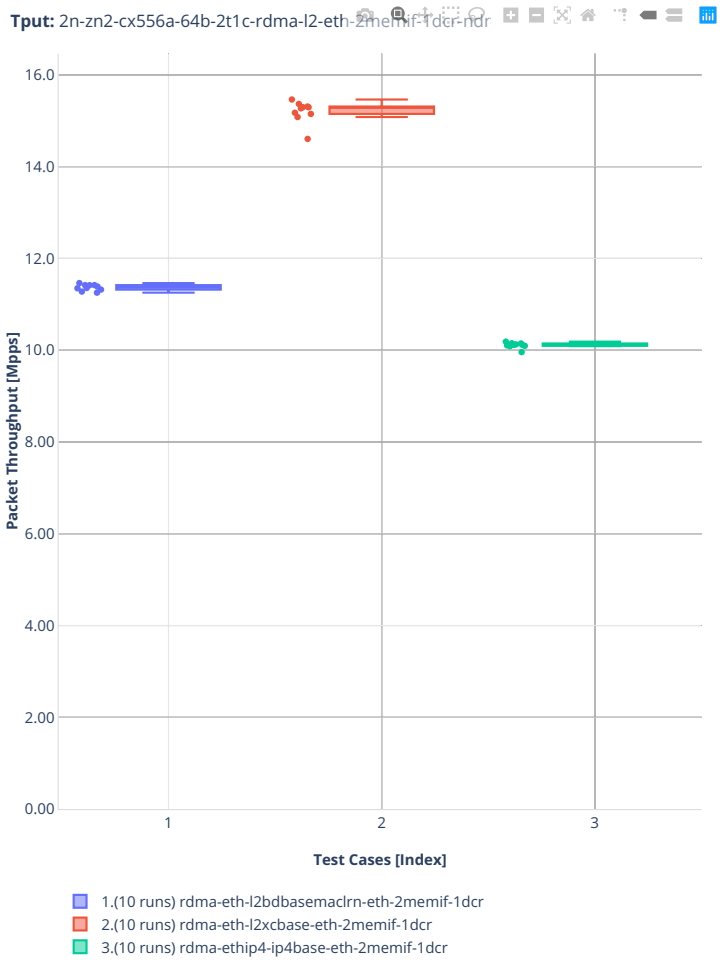

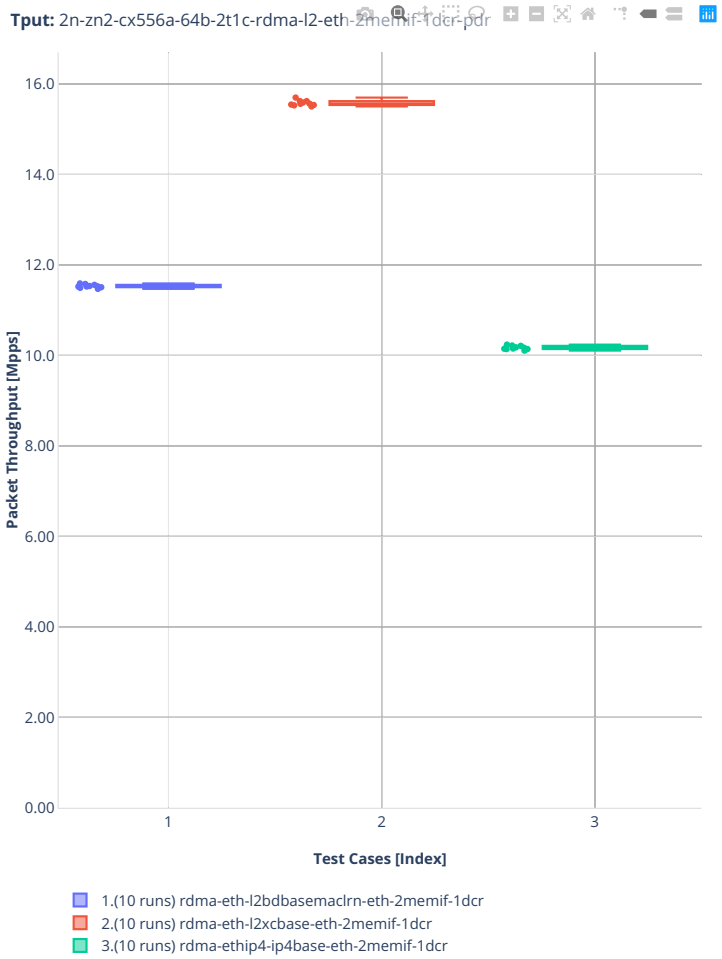
Tput: 2n-zn2-cx556a-64b-2t1c-rdma-l2-eth-2memif-1dcr-pdr
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) rdma-eth-l2bdbasemaclrn-eth-2memif-1dcr
2.(10 runs) rdma-eth-l2xcbase-eth-2memif-1dcr
3.(10 runs) rdma-ethip4-ip4base-eth-2memif-1dcr

2.3.9 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[117].

[117] https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls2206

3n-icx-xxv710

imix-2t1c-ipsec-aes256gcm-ip4routing-base-scale-sw-avf

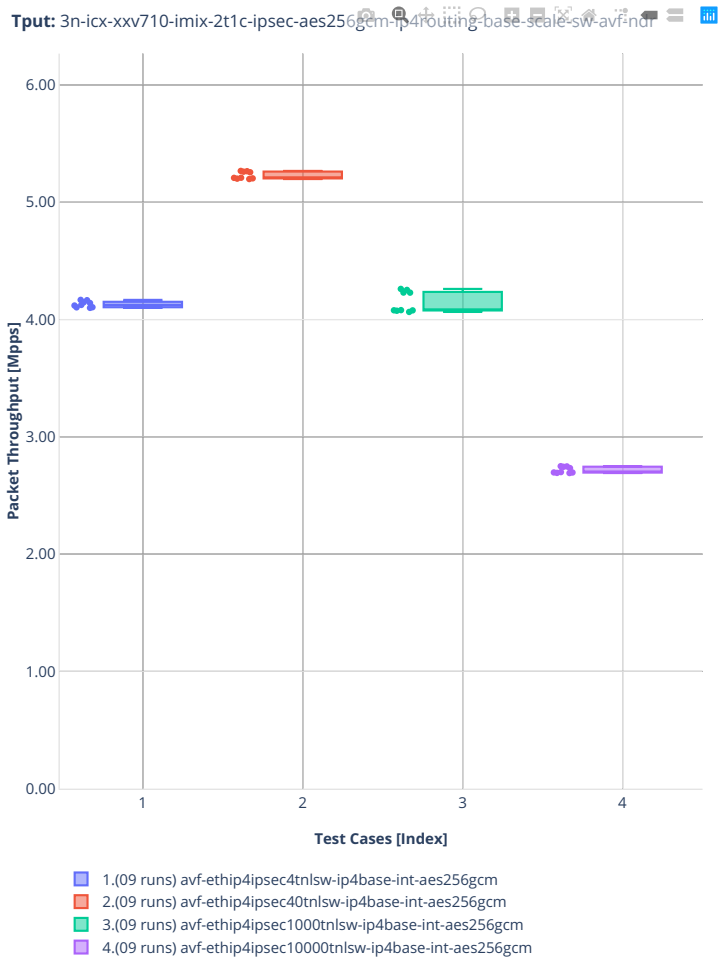

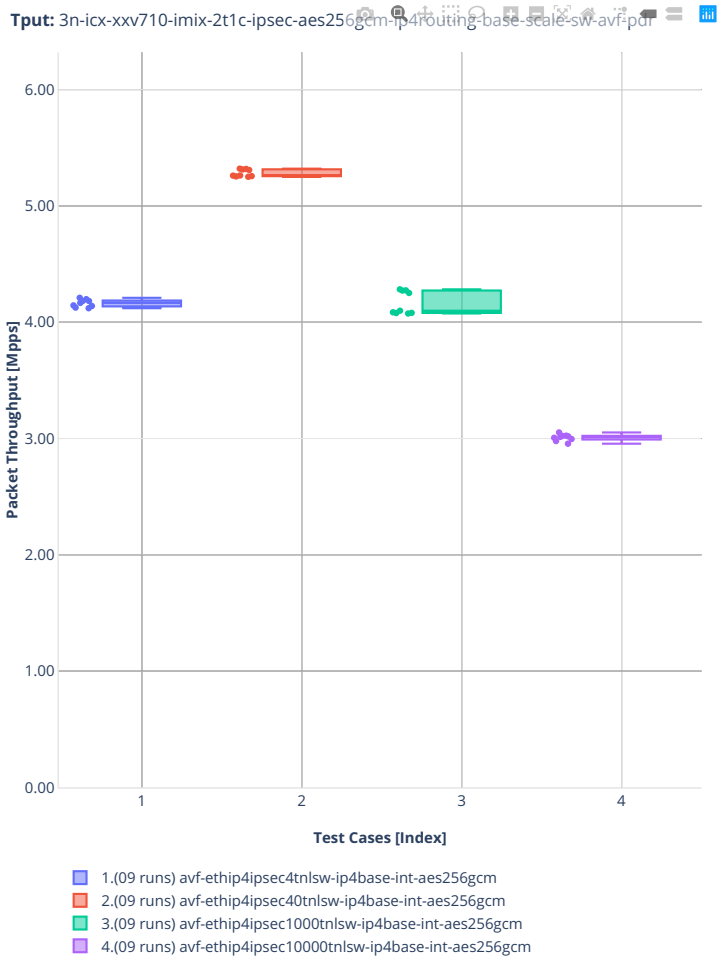
Tput: 3n-icx-xxv710-imix-2t1c-ipsec-aes256gcm-ip4routing-base-scale-sw-avf-pdr
Packet Throughput [Mpps]
6.00
5.00
4.00
3.00
2.00
1.00
0.00
1
2
3
4
Test Cases [Index]
1.(09 runs) avf-ethip4ipsec4tnlsw-ip4base-int-aes256gcm
2.(09 runs) avf-ethip4ipsec40tnlsw-ip4base-int-aes256gcm
3.(09 runs) avf-ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
4.(09 runs) avf-ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

imix-2t1c-ipsec-ip4routing-base-scale-sw-dpdk

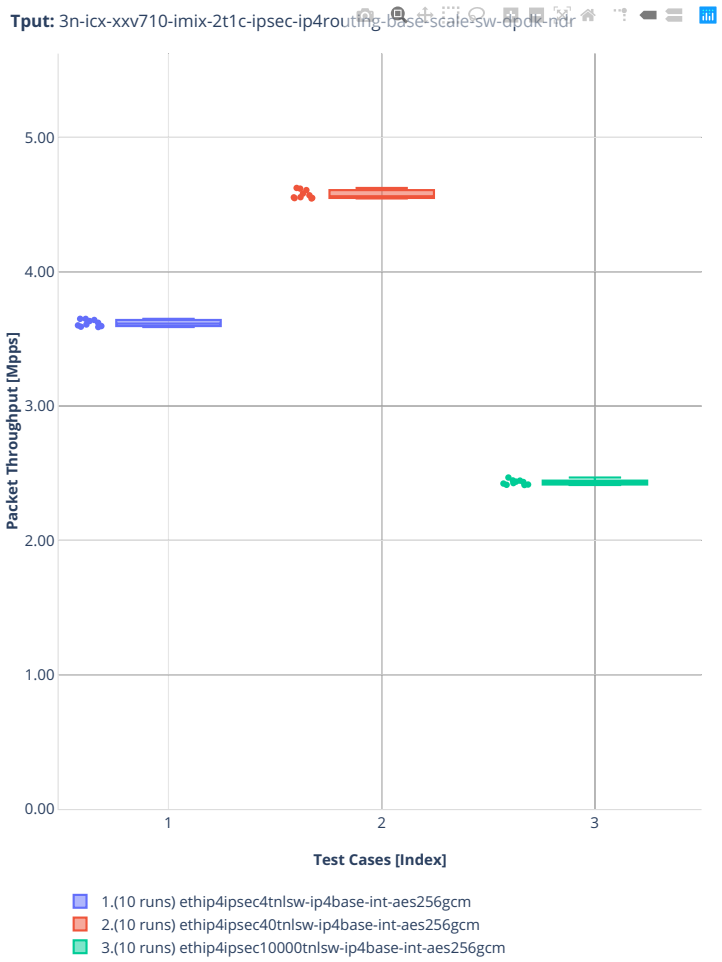

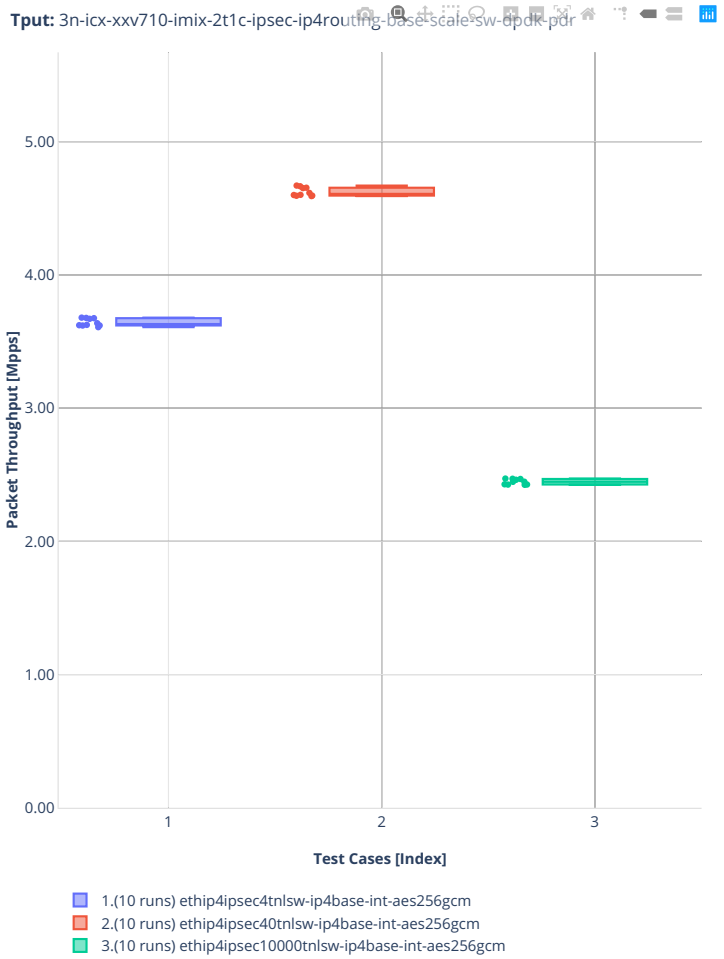
Tput: 3n-icx-xxv710-imix-2t1c-ipsec-ip4routing-base-scale-sw-dpdk-pdr
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes256gcm
2.(10 runs) ethip4ipsec40tnlsw-ip4base-int-aes256gcm
3.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

1518b-2t1c-ipsec-aes128cbc-ip4routing-base-scale-sw-avf

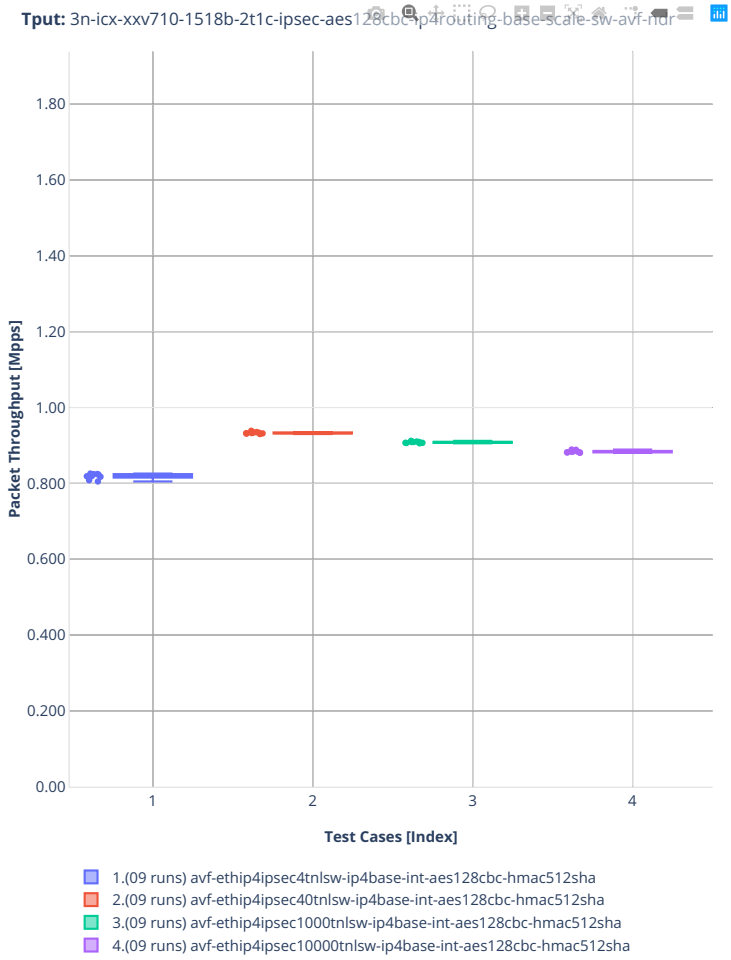

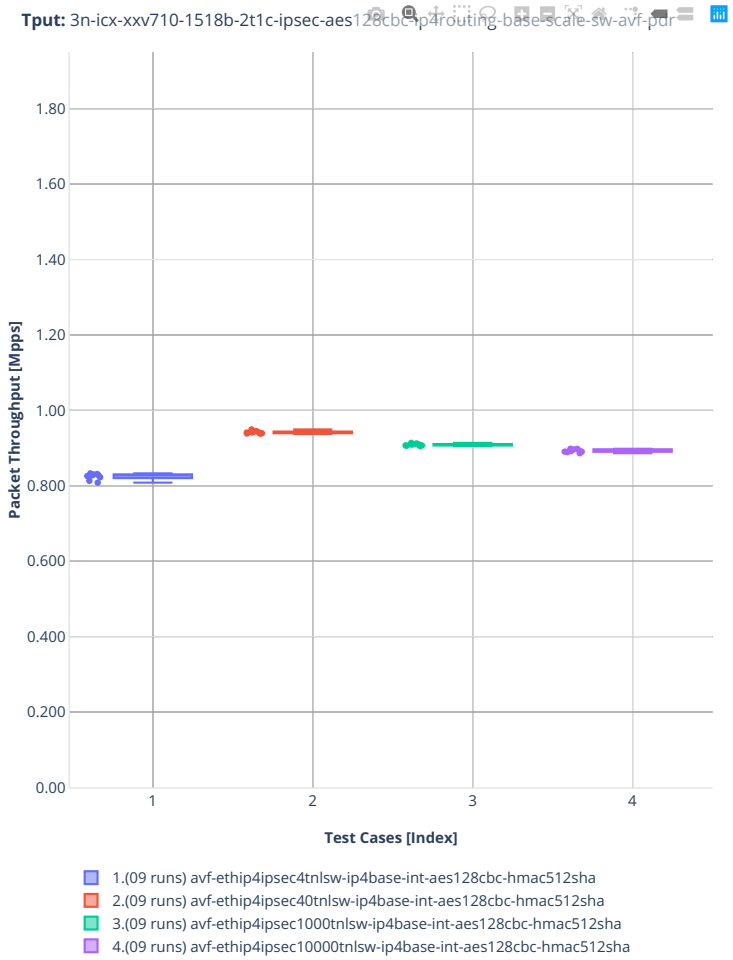
Tput: 3n-icx-xxv710-1518b-2t1c-ipsec-aes128cbc-ip4routing-base-scale-sw-avf-pdr
Packet Throughput [Mpps]
1.80
1.60
1.40
1.20
1.00
0.800
0.600
0.400
0.200
0.00
1
2
3
4
Test Cases [Index]
1.(09 runs) avf-ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
2.(09 runs) avf-ethip4ipsec40tnlsw-ip4base-int-aes128cbc-hmac512sha
3.(09 runs) avf-ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha
4.(09 runs) avf-ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha

1518b-2t1c-ipsec-aes256gcm-ip4routing-base-scale-sw-avf

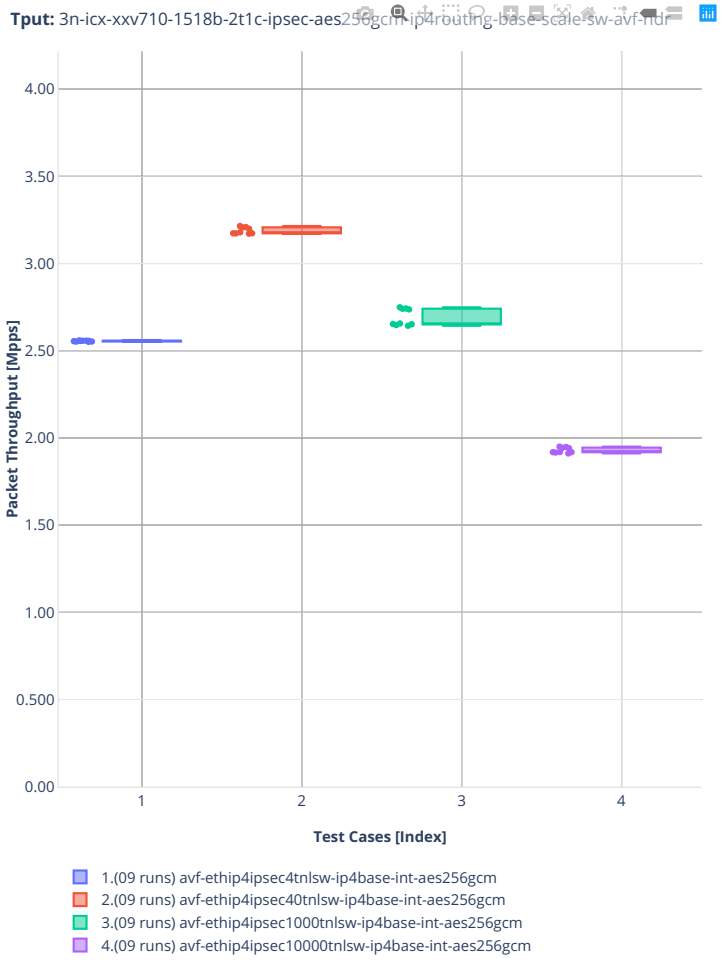

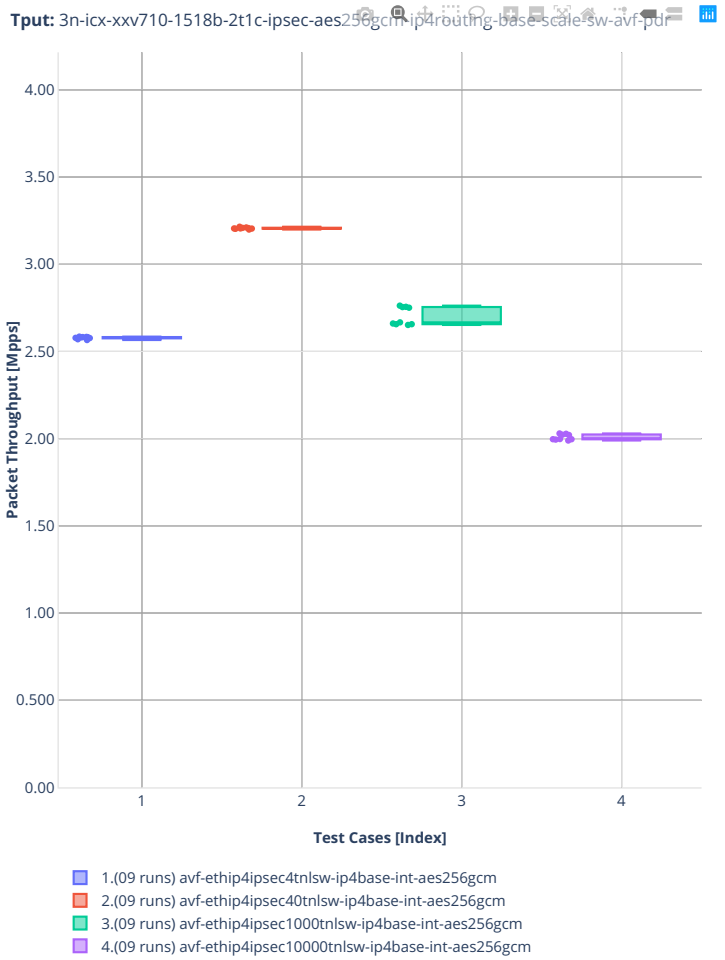
Tput: 3n-icx-xxv710-1518b-2t1c-ipsec-aes256gcm-ip4routing-base-scale-sw-avf-pdr
Packet Throughput [Mpps]
4.00
3.50
3.00
2.50
2.00
1.50
1.00
0.500
0.00
1
2
3
4
Test Cases [Index]
1.(09 runs) avf-ethip4ipsec4tnlsw-ip4base-int-aes256gcm
2.(09 runs) avf-ethip4ipsec40tnlsw-ip4base-int-aes256gcm
3.(09 runs) avf-ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
4.(09 runs) avf-ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

1518b-2t1c-ipsec-ip4routing-base-scale-sw-dpdk

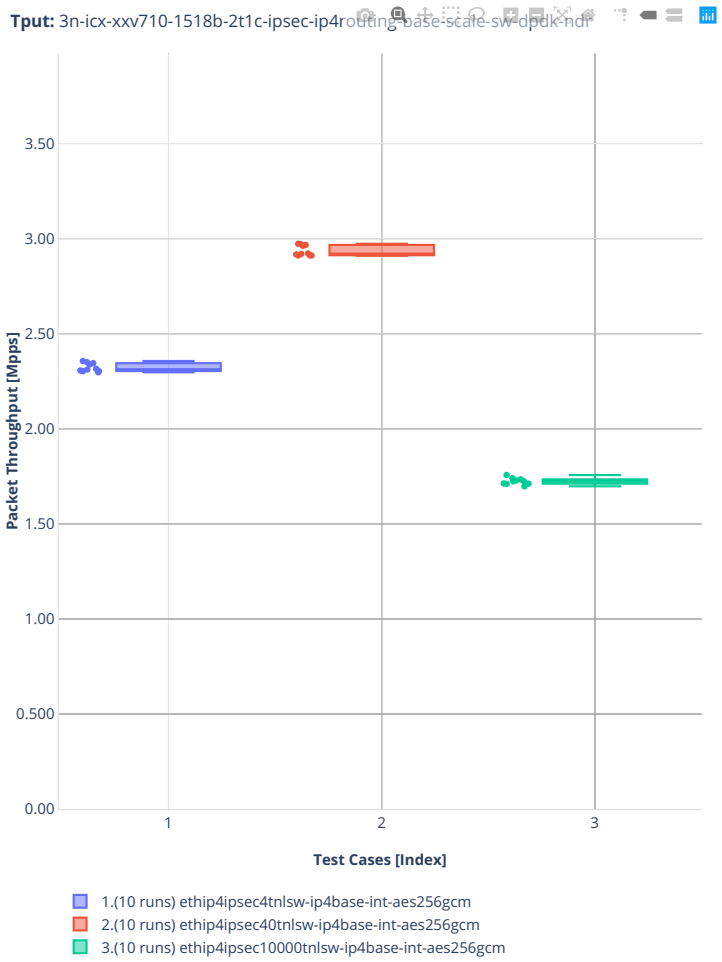

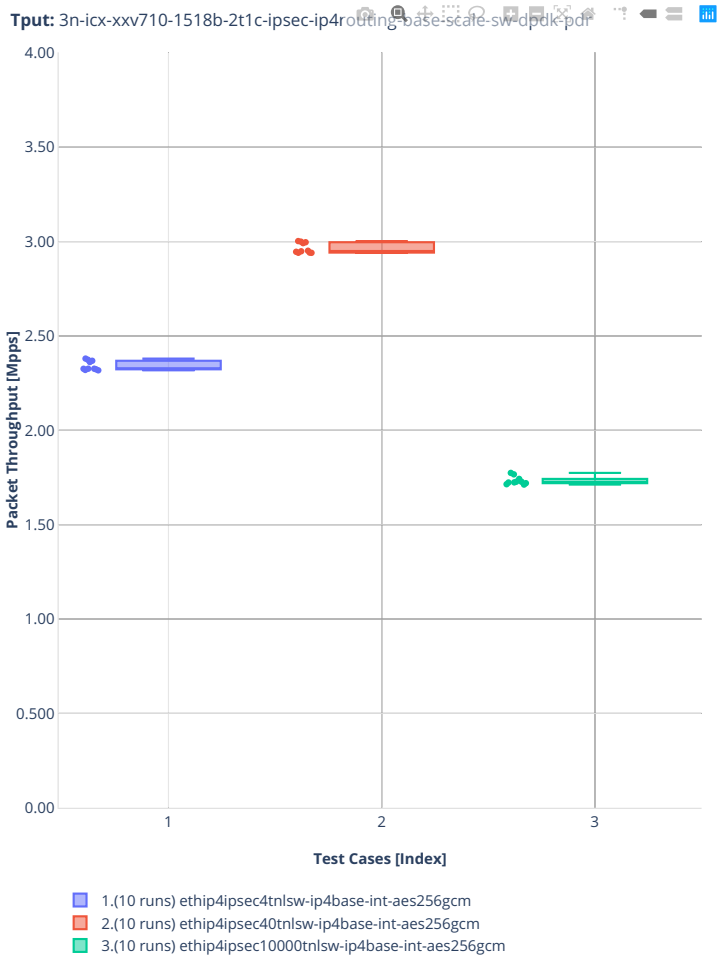
Tput: 3n-icx-xxv710-1518b-2t1c-ipsec-ip4routing-base-scale-sw-dpdk-pdr
4.00
3.50
3.00
2.50
2.00
1.50
1.00
0.500
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes256gcm
2.(10 runs) ethip4ipsec40tnlsw-ip4base-int-aes256gcm
3.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

3n-skx-xxv710

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

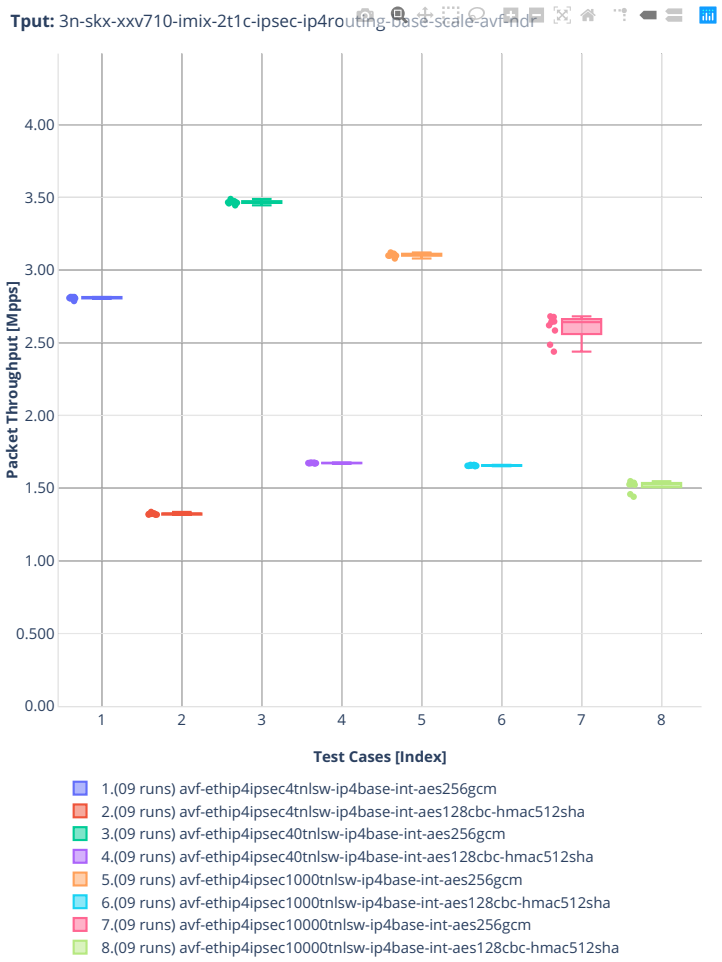

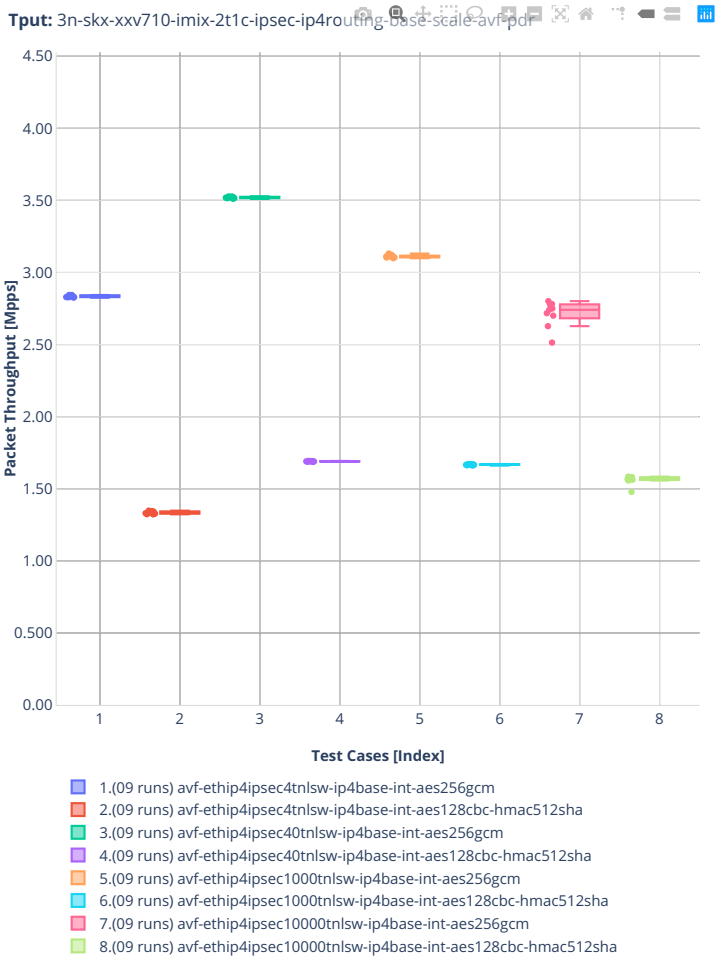
Tput: 3n-skx-xxv710-imix-2t1c-ipsec-ip4routing-base-scale-avf-pdr
Packet Throughput [Mpps]
4.50
4.00
3.50
3.00
2.50
2.00
1.50
1.00
0.500
0.00
1
2
3
4
5
6
7
8
Test Cases [Index]
1.(09 runs) avf-ethip4ipsec4tnlsw-ip4base-int-aes256gcm
2.(09 runs) avf-ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
3.(09 runs) avf-ethip4ipsec40tnlsw-ip4base-int-aes256gcm
4.(09 runs) avf-ethip4ipsec40tnlsw-ip4base-int-aes128cbc-hmac512sha
5.(09 runs) avf-ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
6.(09 runs) avf-ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha
7.(09 runs) avf-ethip4ipsec10000tnlsw-ip4base-int-aes256gcm
8.(09 runs) avf-ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha

1518b-2t1c-ipsec-ip4routing-base-scale-avf

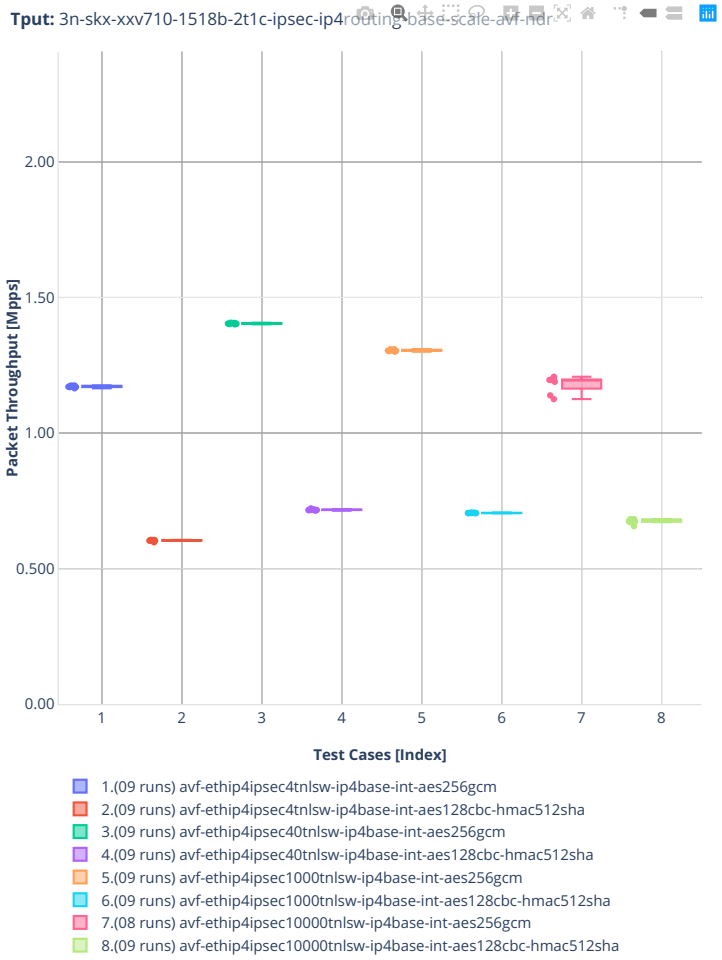

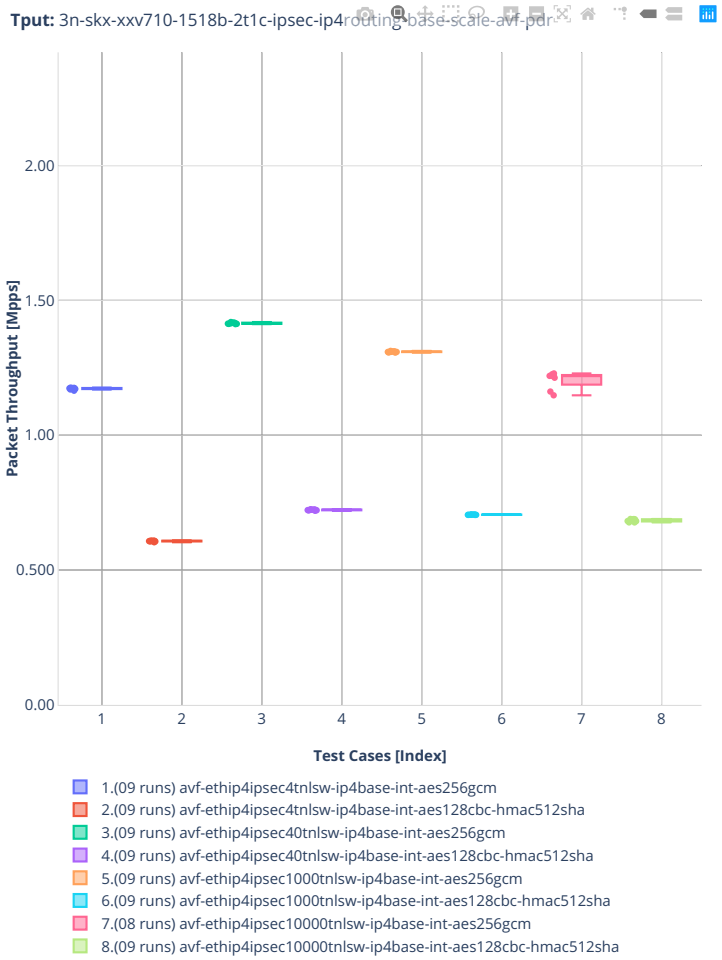
Tput: 3n-skx-xxv710-1518b-2t1c-ipsec-ip4routing-base-scale-avf-pdr
Packet Throughput [Mpps]
2.00
1.50
1.00
0.500
0.00
1
2
3
4
5
6
7
8
Test Cases [Index]
1.(09 runs) avf-ethip4ipsec4tnlsw-ip4base-int-aes256gcm
2.(09 runs) avf-ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
3.(09 runs) avf-ethip4ipsec40tnlsw-ip4base-int-aes256gcm
4.(09 runs) avf-ethip4ipsec40tnlsw-ip4base-int-aes128cbc-hmac512sha
5.(09 runs) avf-ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
6.(09 runs) avf-ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha
7.(08 runs) avf-ethip4ipsec10000tnlsw-ip4base-int-aes256gcm
8.(09 runs) avf-ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha

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

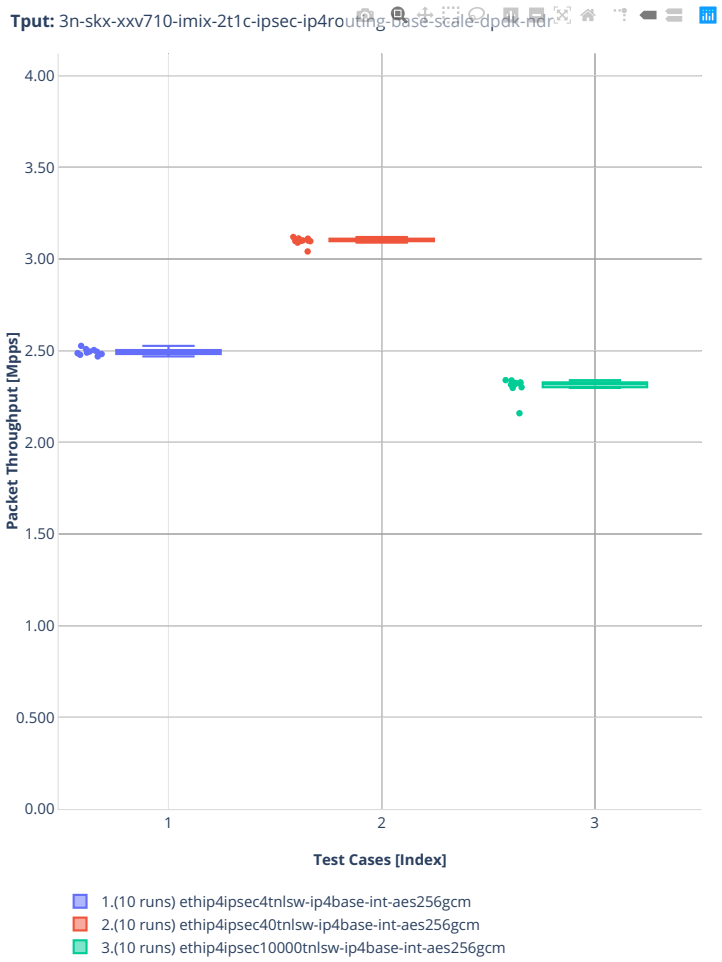

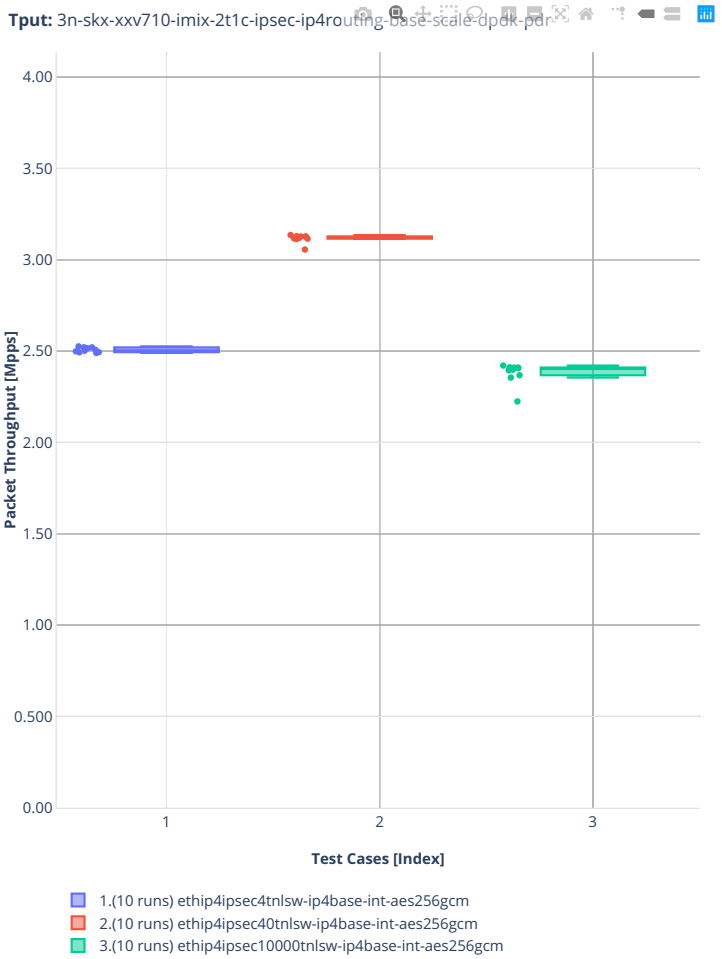
Tput: 3n-skx-xxv710-imix-2t1c-ipsec-ip4routing-base-scale-dpdk-pdr
4.00
3.50
3.00
2.50
2.00
1.50
1.00
0.500
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes256gcm
2.(10 runs) ethip4ipsec40tnlsw-ip4base-int-aes256gcm
3.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

1518b-2t1c-ipsec-ip4routing-base-scale-dpdk

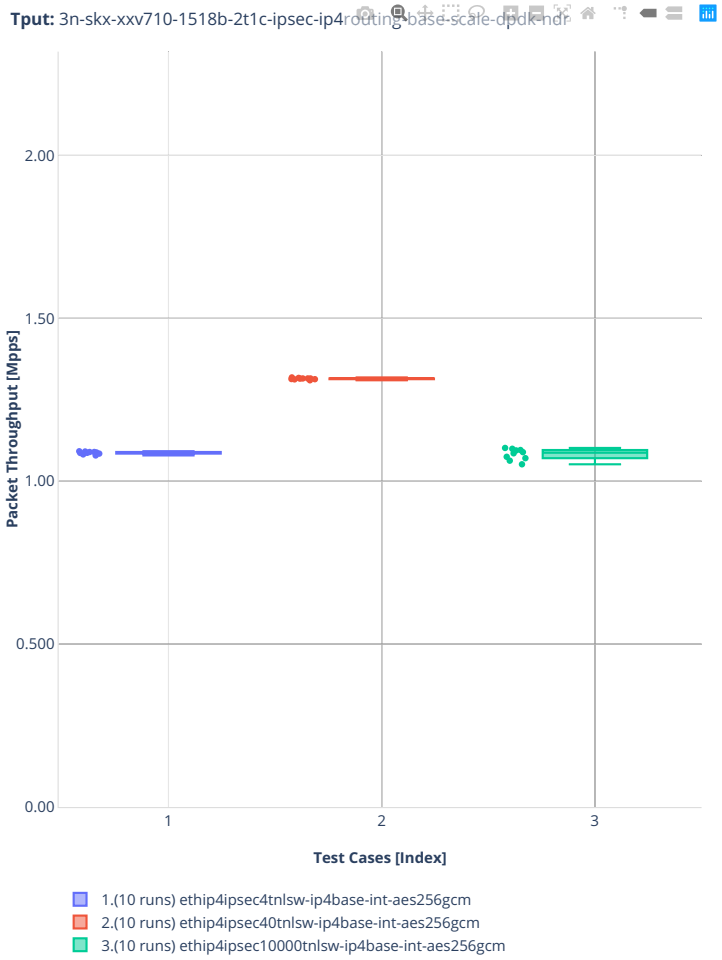

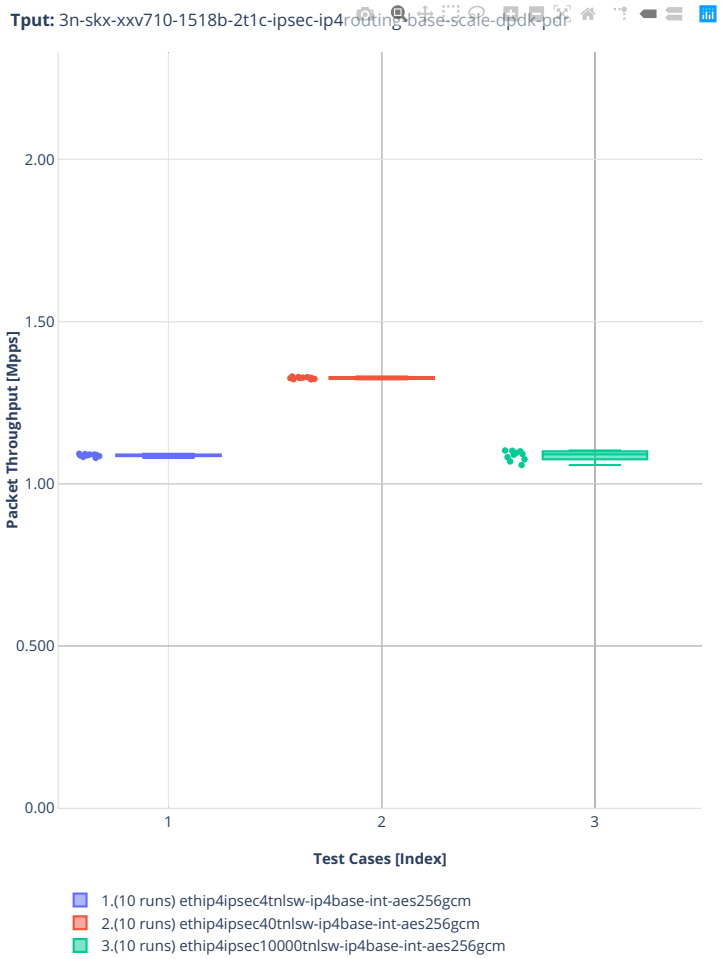
Tput: 3n-skx-xxv710-1518b-2t1c-ipsec-ip4routing-base-scale-dpdk-pdr
2.00
1.50
1.00
0.500
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes256gcm
2.(10 runs) ethip4ipsec40tnlsw-ip4base-int-aes256gcm
3.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

3n-alt-xl710

1518b-1t1c-ipsec-aes128cbc-ip4routing-scale-sw

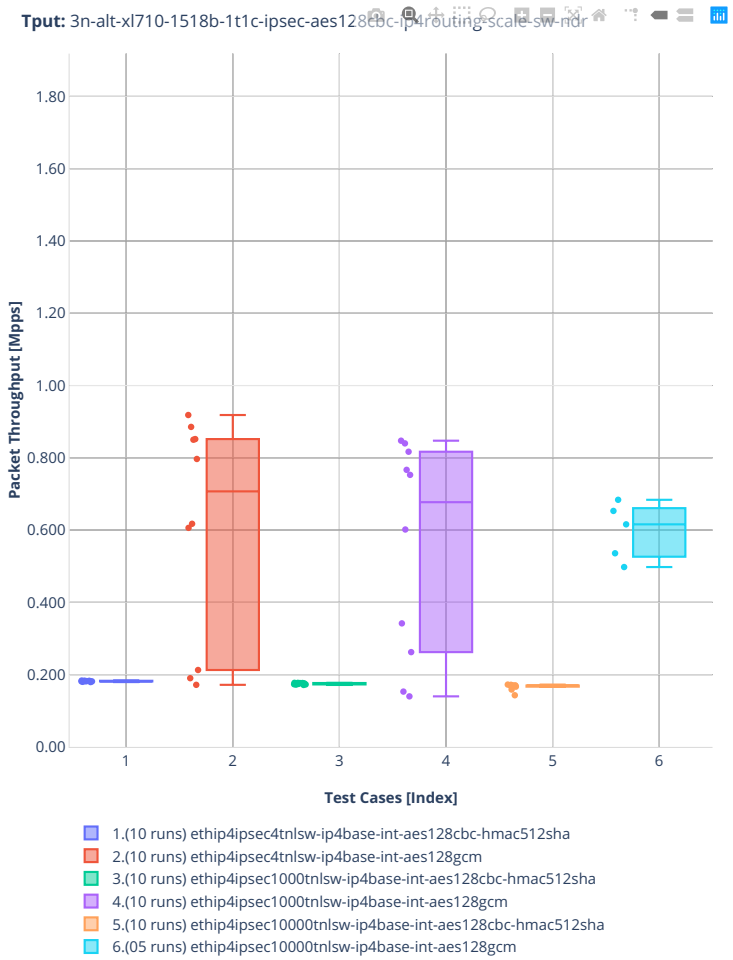

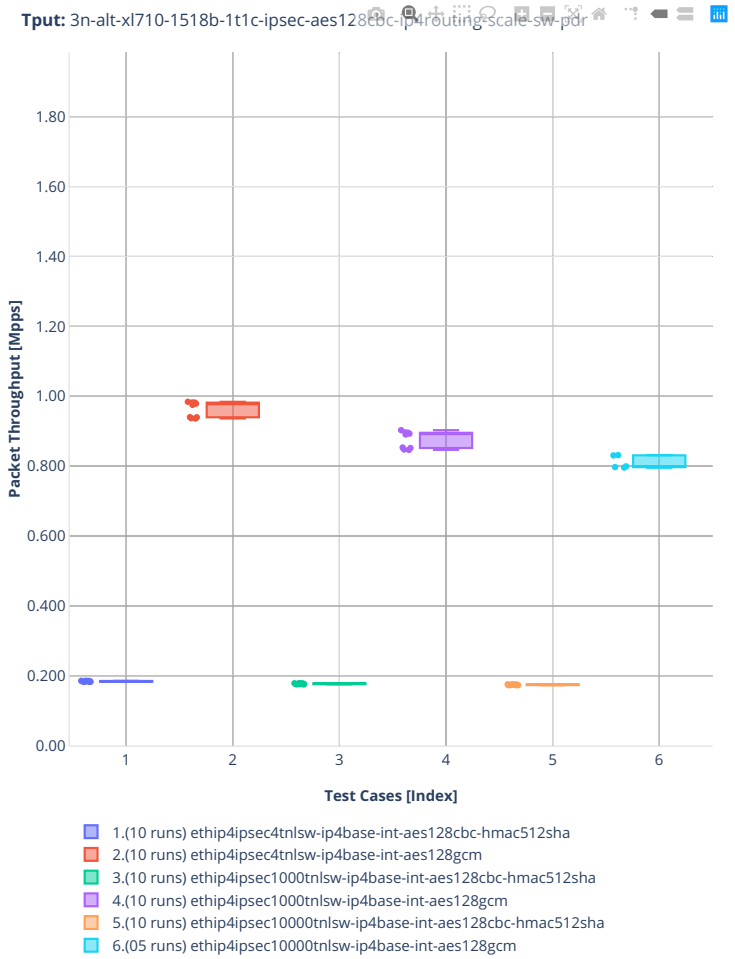
Tput: 3n-alt-xl710-1518b-1t1c-ipsec-aes128cbc-ip4routing-scale-sw-pdr
Packet Throughput [Mpps]
1.80
1.60
1.40
1.20
1.00
0.800
0.600
0.400
0.200
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
2.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes128gcm
3.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha
4.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes128gcm
5.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha
6.(05 runs) ethip4ipsec10000tnlsw-ip4base-int-aes128gcm

1518b-1t1c-ipsec-aes256gcm-ip4routing-scale-sw

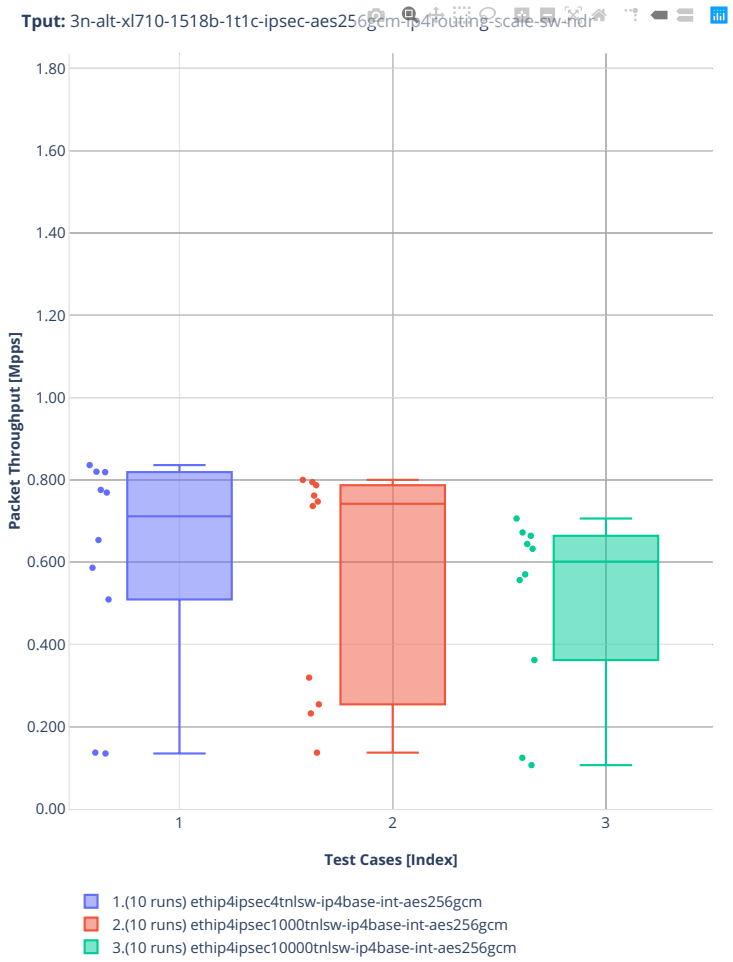

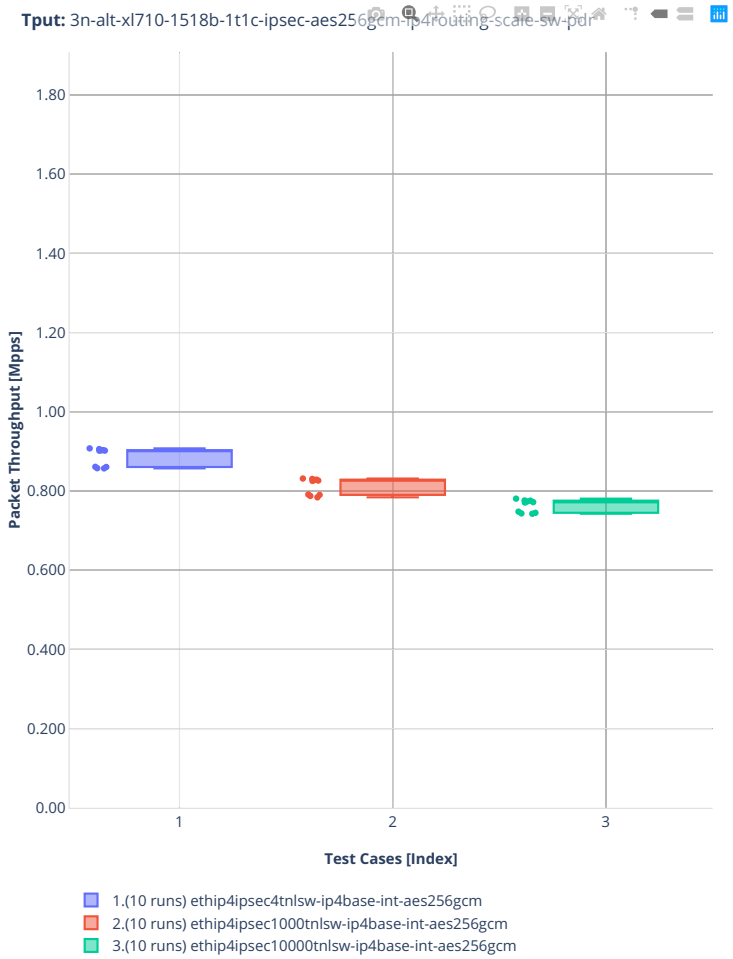
Tput: 3n-alt-xl710-1518b-1t1c-ipsec-aes256gcm-ip4routing-scale-sw-pdr
Packet Throughput [Mpps]
1.80
1.60
1.40
1.20
1.00
0.800
0.600
0.400
0.200
0.00
1
2
3
Test Cases [Index]
1.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes256gcm
2.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
3.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

imix-1t1c-ipsec-aes128cbc-ip4routing-scale-sw

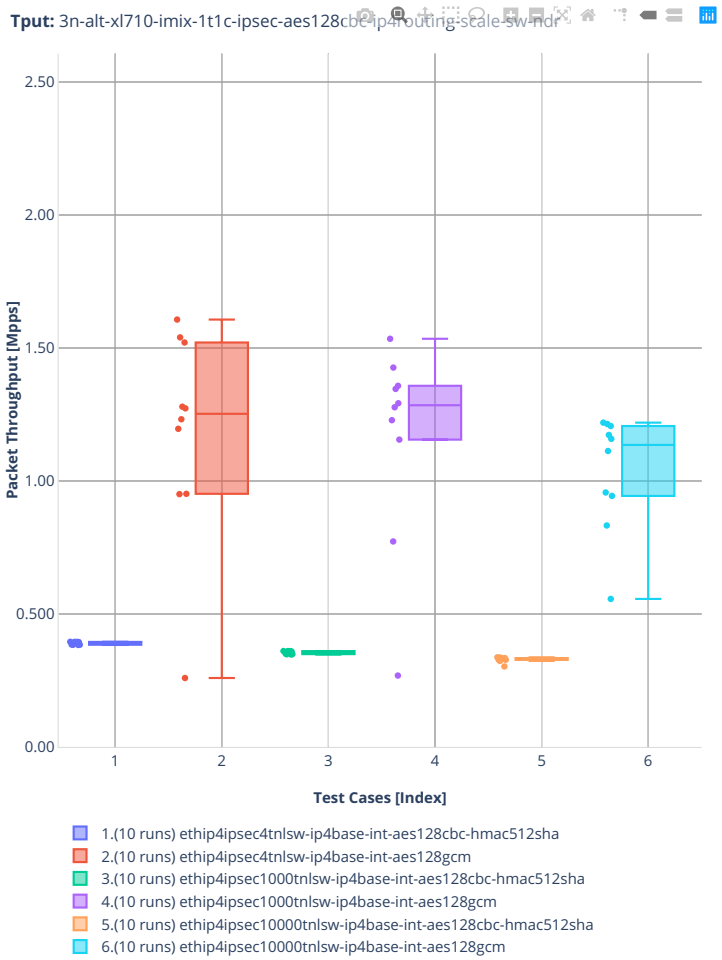

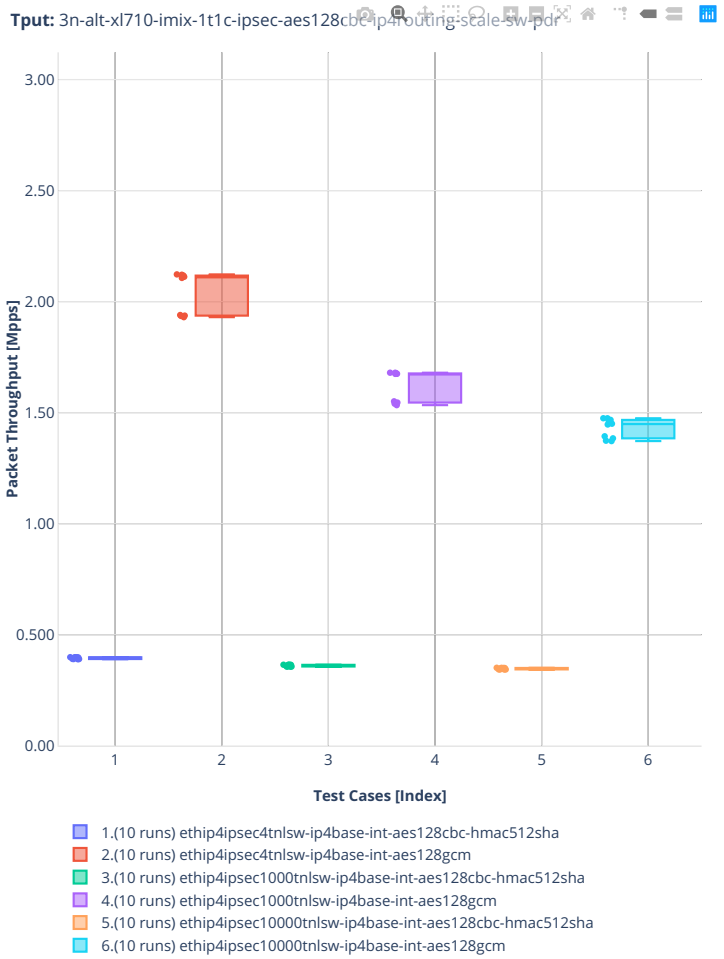
Tput: 3n-alt-xl710-imix-1t1c-ipsec-aes128cbc-ip4routing-scale-sw-pdr
Packet Throughput [Mpps]
3.00
2.50
2.00
1.50
1.00
0.500
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
2.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes128gcm
3.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha
4.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes128gcm
5.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha
6.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes128gcm

imix-1t1c-ipsec-aes256gcm-ip4routing-scale-sw

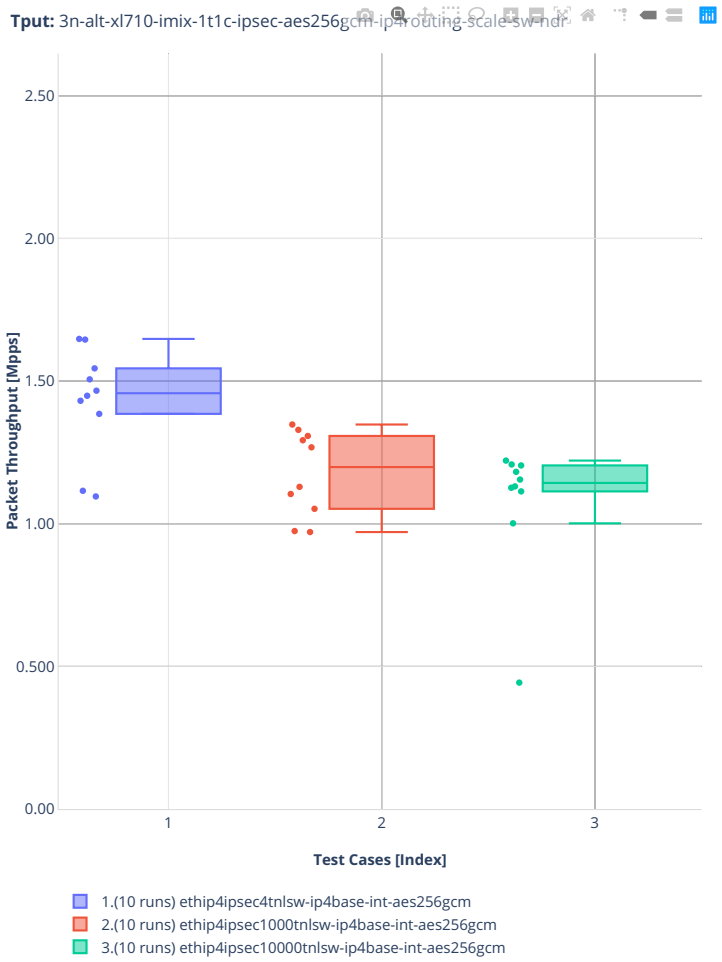

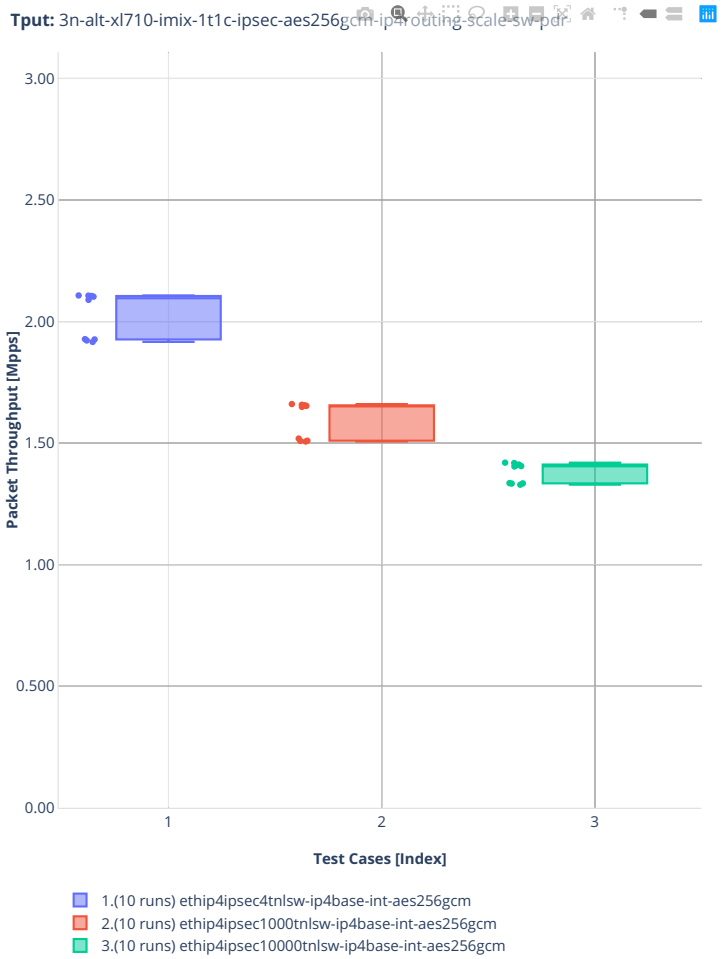
Tput: 3n-alt-xl710-imix-1t1c-ipsec-aes256gcm-ip4routing-scale-sw-pdr
3.00
2.50
2.00
1.50
1.00
0.500
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes256gcm
2.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
3.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

64b-1t1c-ipsec-ip4base-policy-outbound-nocrypto

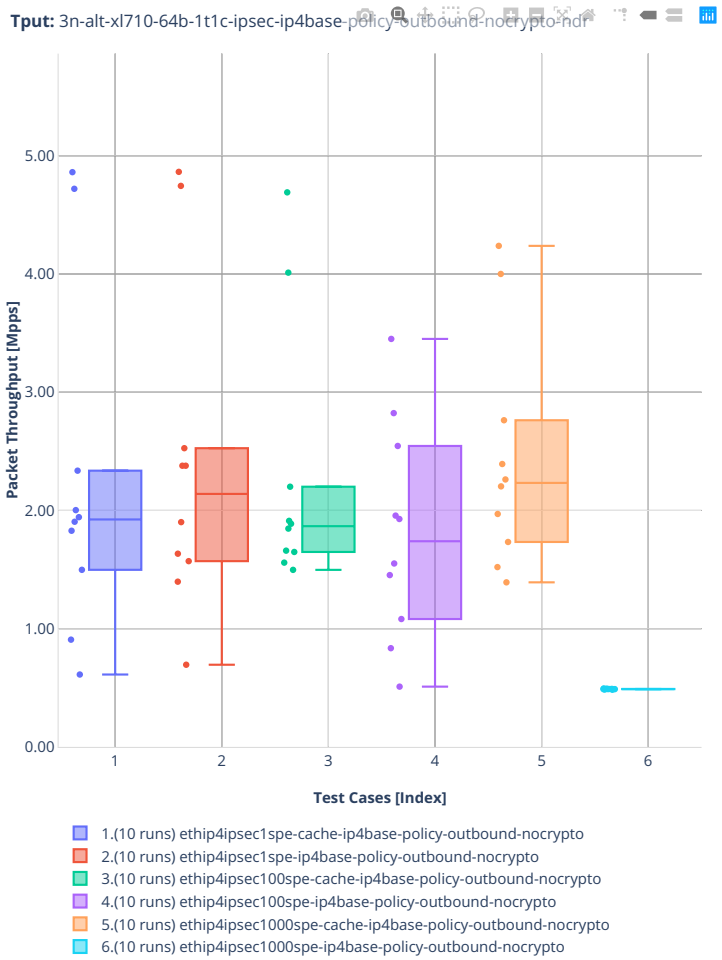

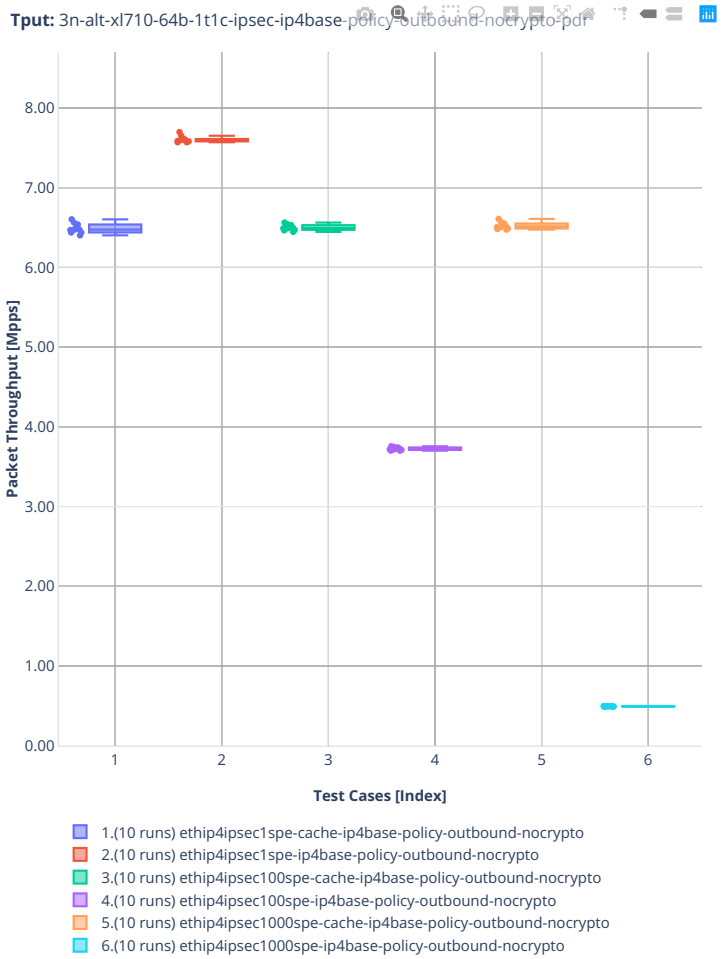
Tput: 3n-alt-xl710-64b-1t1c-ipsec-ip4base-policy-outbound-nocrypto-pdr
Packet Throughput [Mpps]
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) ethip4ipsec1spe-cache-ip4base-policy-outbound-nocrypto
2.(10 runs) ethip4ipsec1spe-ip4base-policy-outbound-nocrypto
3.(10 runs) ethip4ipsec100spe-cache-ip4base-policy-outbound-nocrypto
4.(10 runs) ethip4ipsec100spe-ip4base-policy-outbound-nocrypto
5.(10 runs) ethip4ipsec1000spe-cache-ip4base-policy-outbound-nocrypto
6.(10 runs) ethip4ipsec1000spe-ip4base-policy-outbound-nocrypto

64b-1t1c-ipsec-ip4base-policy-aes256gcm-sw

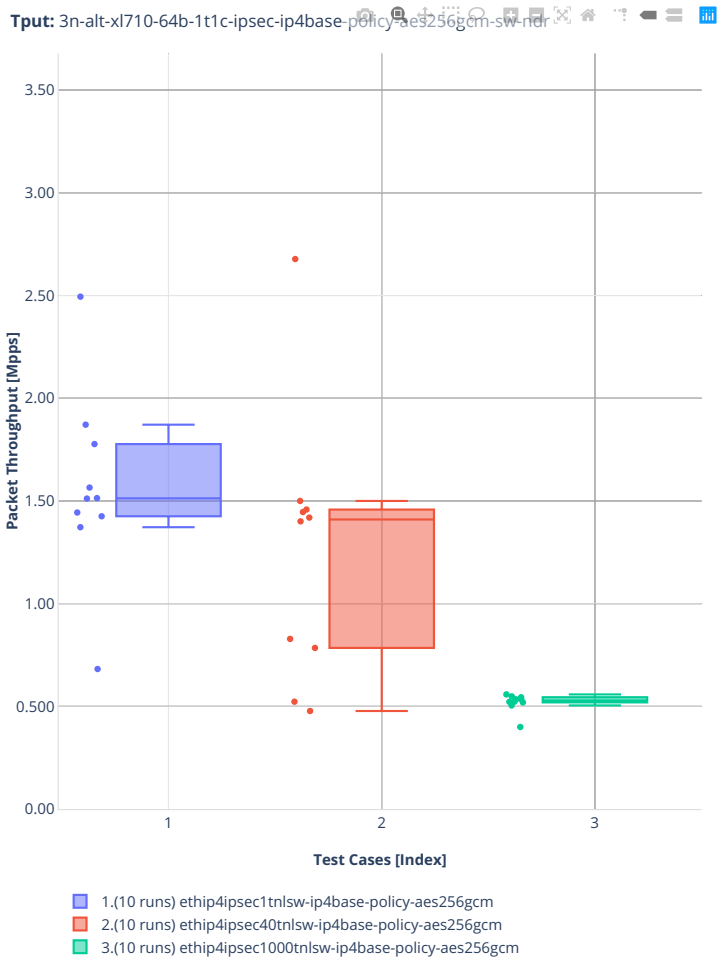

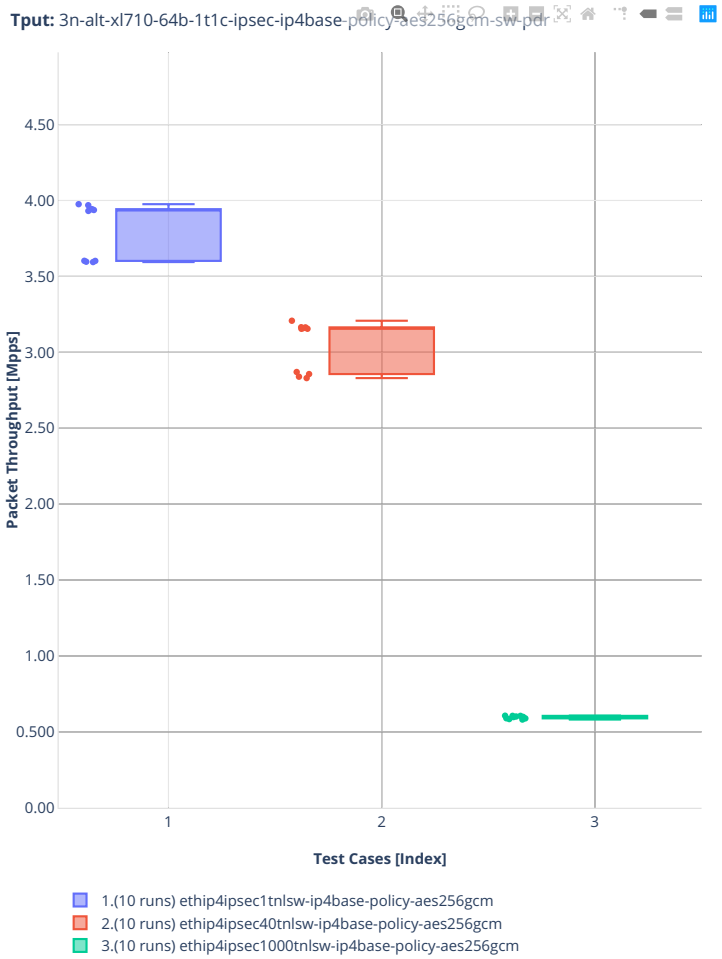
Tput: 3n-alt-xl710-64b-1t1c-ipsec-ip4base-policy-aes256gcm-sw-pdr
Packet Throughput [Mpps]
4.50
4.00
3.50
3.00
2.50
2.00
1.50
1.00
0.500
0.00
1
2
3
Test Cases [Index]
1.(10 runs) ethip4ipsec1tnlsw-ip4base-policy-aes256gcm
2.(10 runs) ethip4ipsec40tnlsw-ip4base-policy-aes256gcm
3.(10 runs) ethip4ipsec1000tnlsw-ip4base-policy-aes256gcm

3n-tsh-x520

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

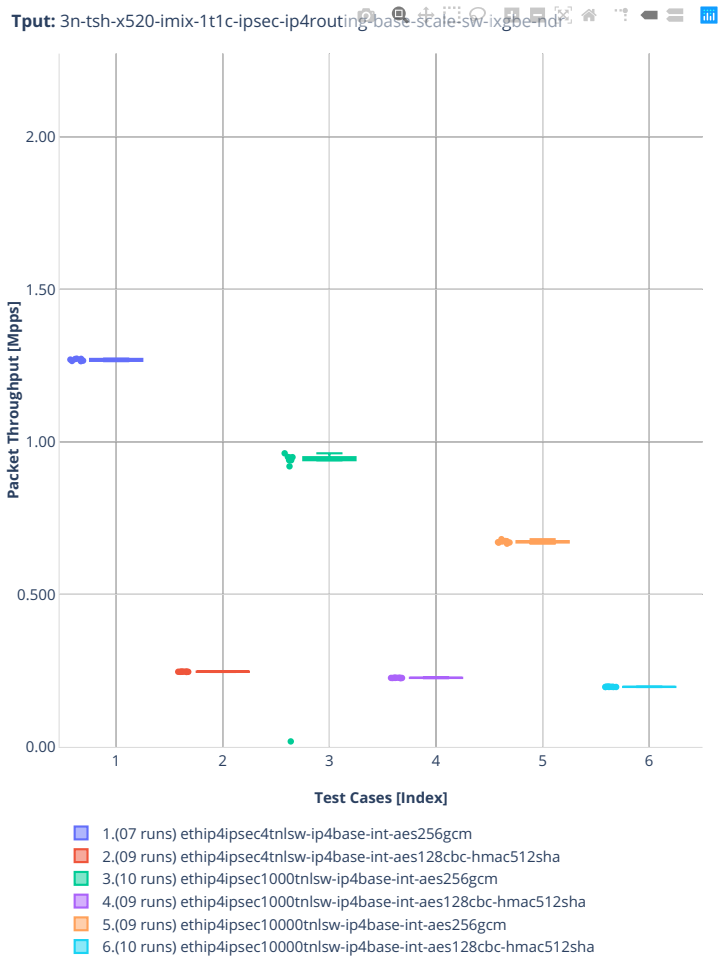

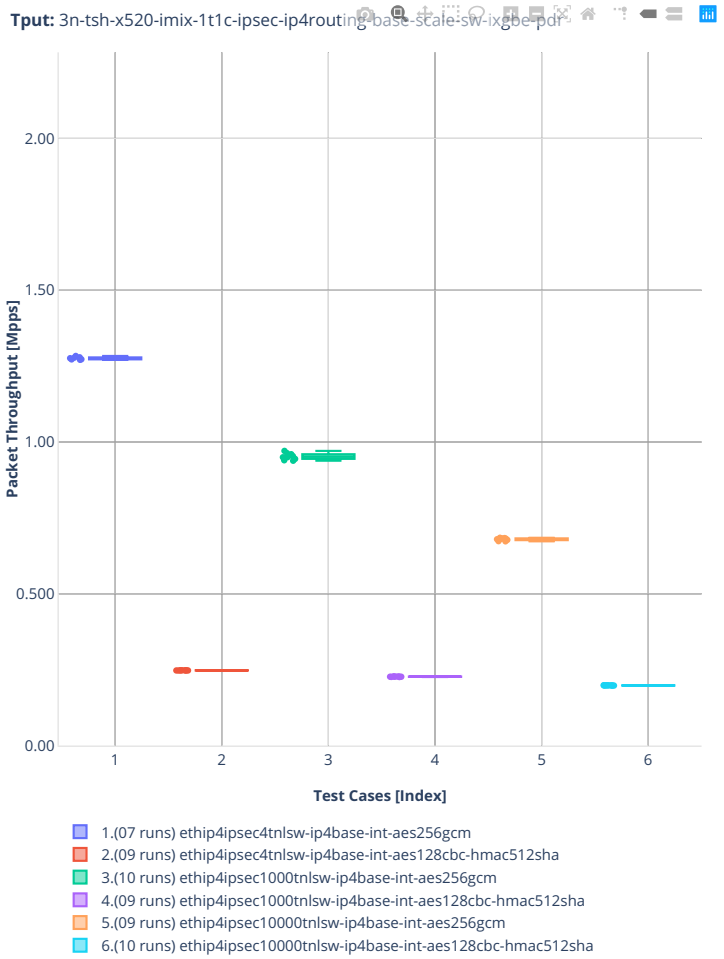
Tput: 3n-tsh-x520-imix-1t1c-ipsec-ip4routing-base-scale-sw-ixgbe-pdr
Packet Throughput [Mpps]
2.00
1.50
1.00
0.500
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(07 runs) ethip4ipsec4tnlsw-ip4base-int-aes256gcm
2.(09 runs) ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
3.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
4.(09 runs) ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha
5.(09 runs) ethip4ipsec10000tnlsw-ip4base-int-aes256gcm
6.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha

2n-tx2-xl710

64b-ipsec-spe-ip4routing-base-scale

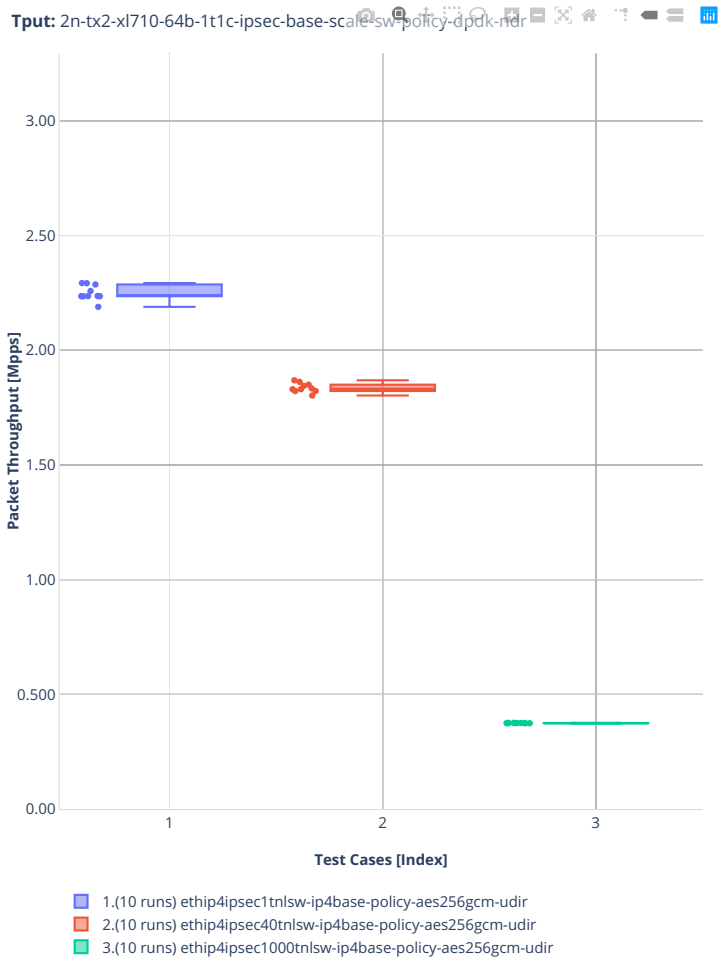

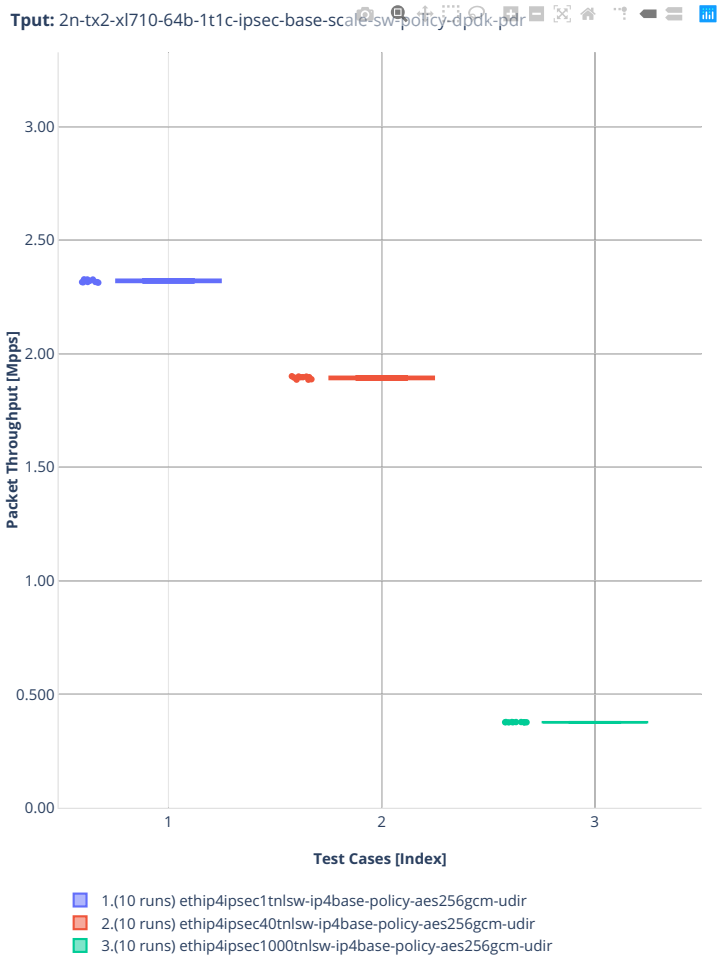
Tput: 2n-tx2-xl710-64b-1t1c-ipsec-base-scale-sw-policy-dpdk-pdr
Packet Throughput [Mpps]
3.00
2.50
2.00
1.50
1.00
0.500
0.00
1
2
3
Test Cases [Index]
1.(10 runs) ethip4ipsec1tnlsw-ip4base-policy-aes256gcm-udir
2.(10 runs) ethip4ipsec40tnlsw-ip4base-policy-aes256gcm-udir
3.(10 runs) ethip4ipsec1000tnlsw-ip4base-policy-aes256gcm-udir

64b-ipsec-ip4routing-base-scale-sw

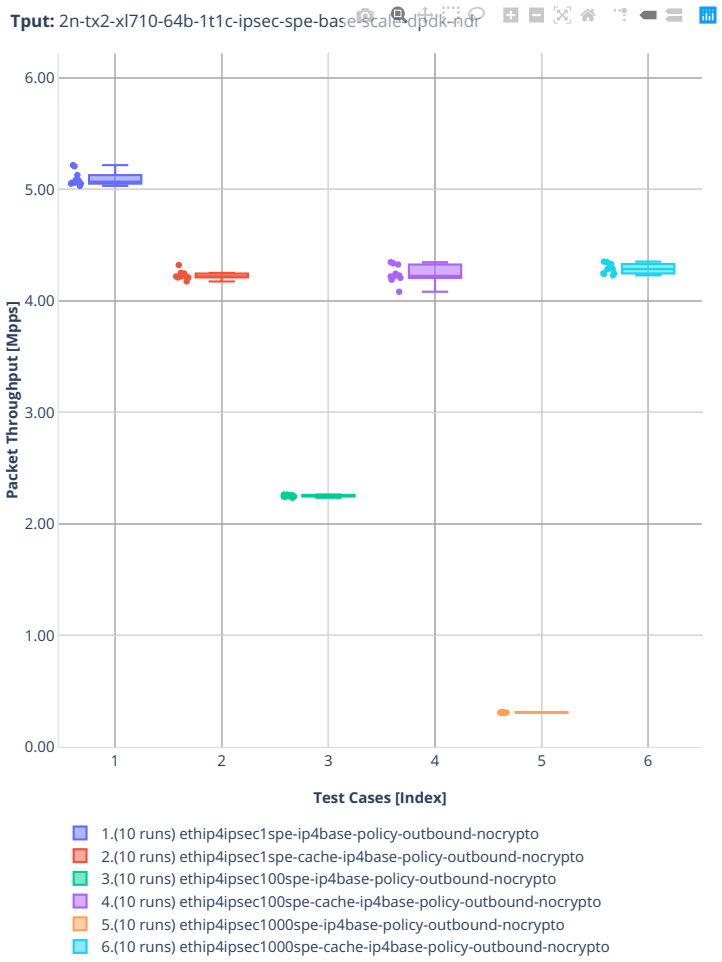

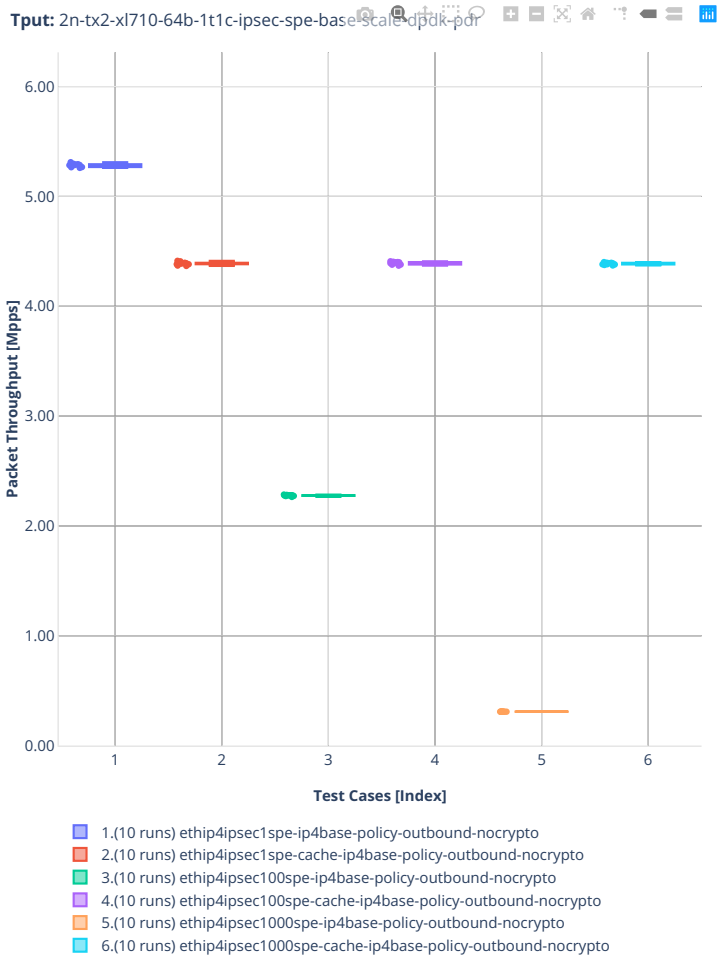
Tput: 2n-tx2-xl710-64b-1t1c-ipsec-spe-base-scale-dpdk-pdr
Packet Throughput [Mpps]
6.00
5.00
4.00
3.00
2.00
1.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) ethip4ipsec1spe-ip4base-policy-outbound-nocrypto
2.(10 runs) ethip4ipsec1spe-cache-ip4base-policy-outbound-nocrypto
3.(10 runs) ethip4ipsec100spe-ip4base-policy-outbound-nocrypto
4.(10 runs) ethip4ipsec100spe-cache-ip4base-policy-outbound-nocrypto
5.(10 runs) ethip4ipsec1000spe-ip4base-policy-outbound-nocrypto
6.(10 runs) ethip4ipsec1000spe-cache-ip4base-policy-outbound-nocrypto

1518b-ipsec-ip4routing-base-scale-sw-cbc

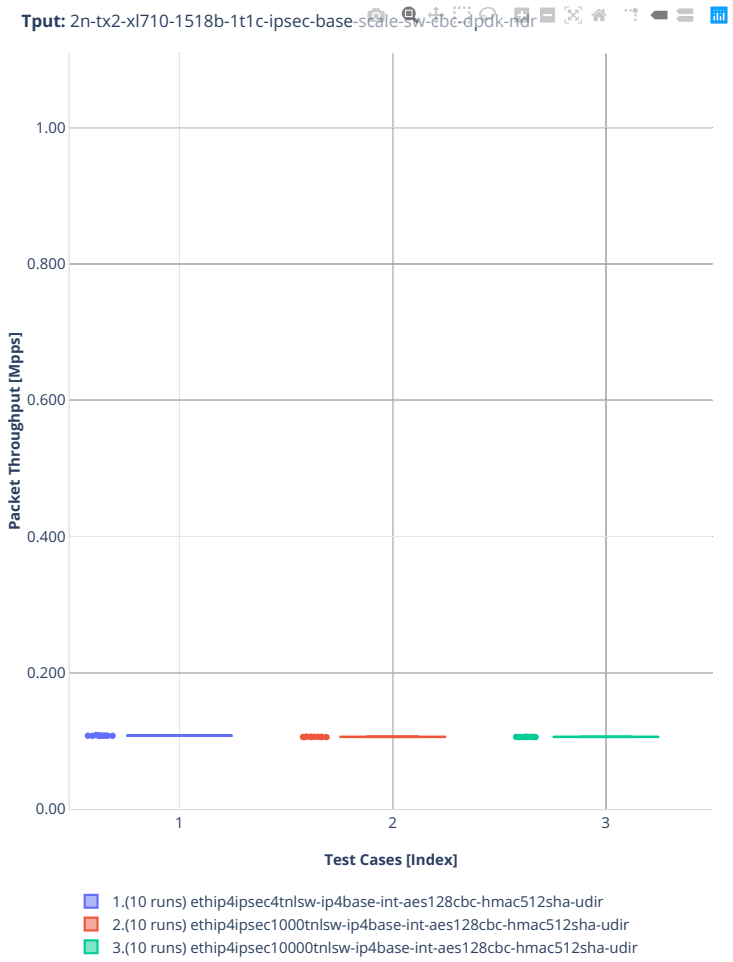

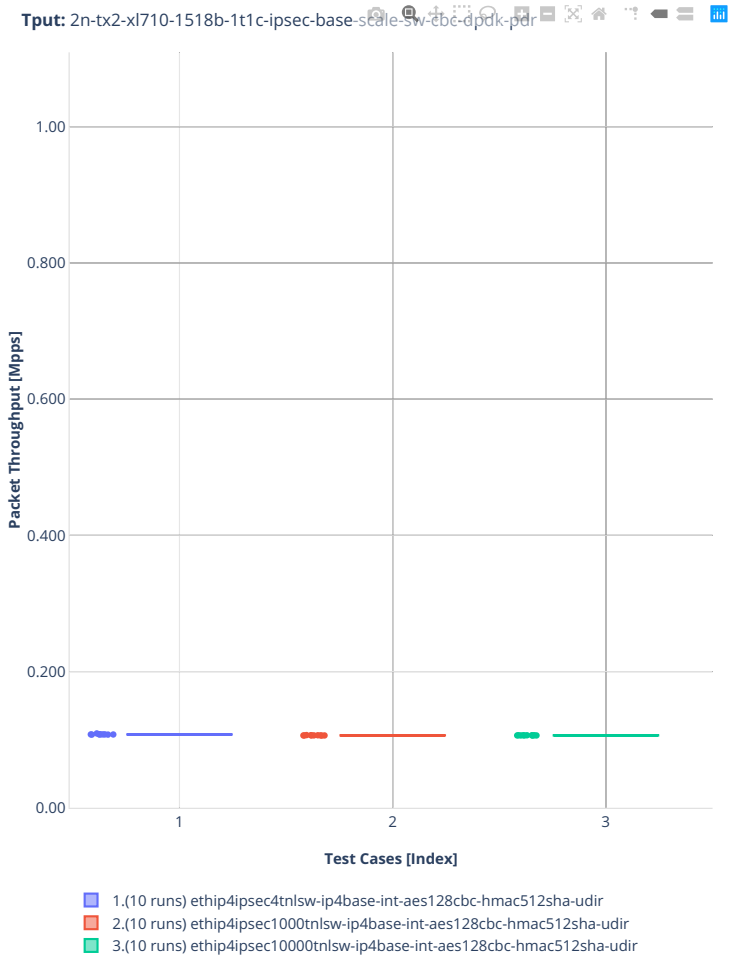
Tput: 2n-tx2-xl710-1518b-1t1c-ipsec-base-scale-sw-cbc-dpdk-pdr
Packet Throughput [Mpps]
1.00
0.800
0.600
0.400
0.200
0.00
1
2
3
Test Cases [Index]
1.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha-udir
2.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha-udir
3.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha-udir

imix-ipsec-ip4routing-base-scale-sw-cbc

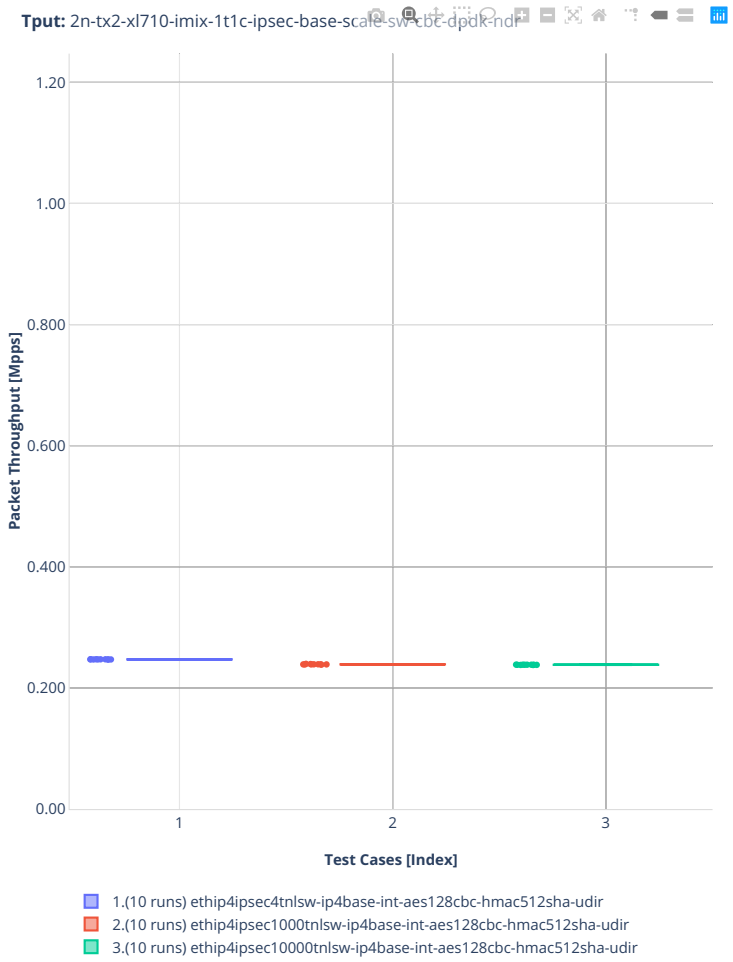

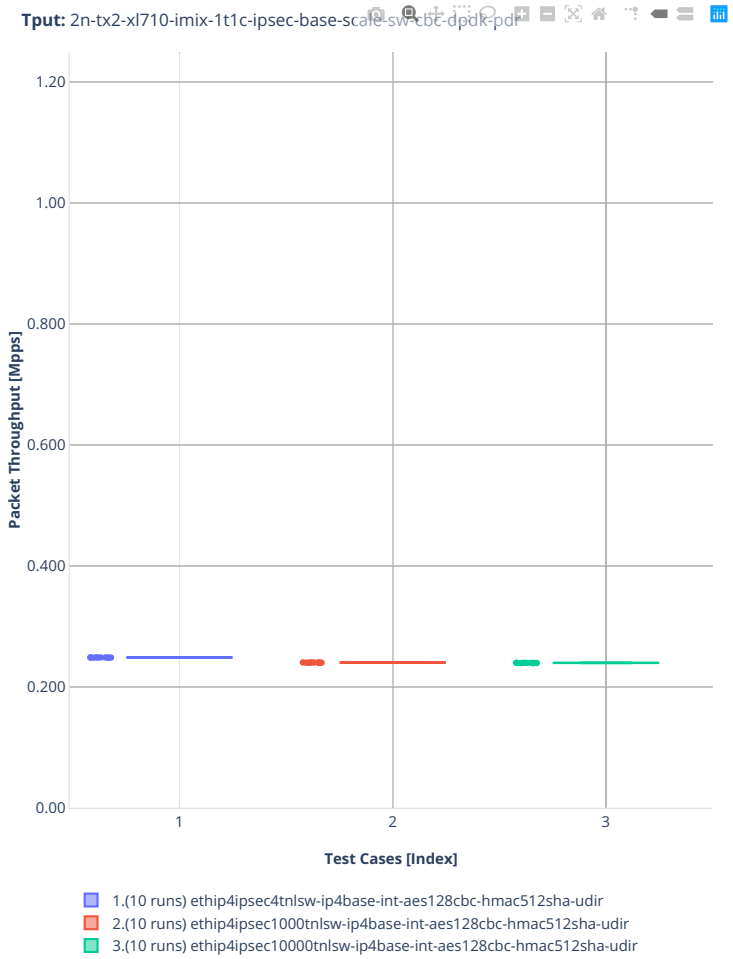
Tput: 2n-tx2-xl710-imix-1t1c-ipsec-base-scale-sw-cbc-dpdk-pdr
1.20
1.00
0.800
0.600
0.400
0.200
0.00
Packet Throughput [Mpps]
1
2
3
Test Cases [Index]
1.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha-udir
2.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha-udir
3.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha-udir

1518b-ipsec-ip4routing-base-scale-sw-gcm

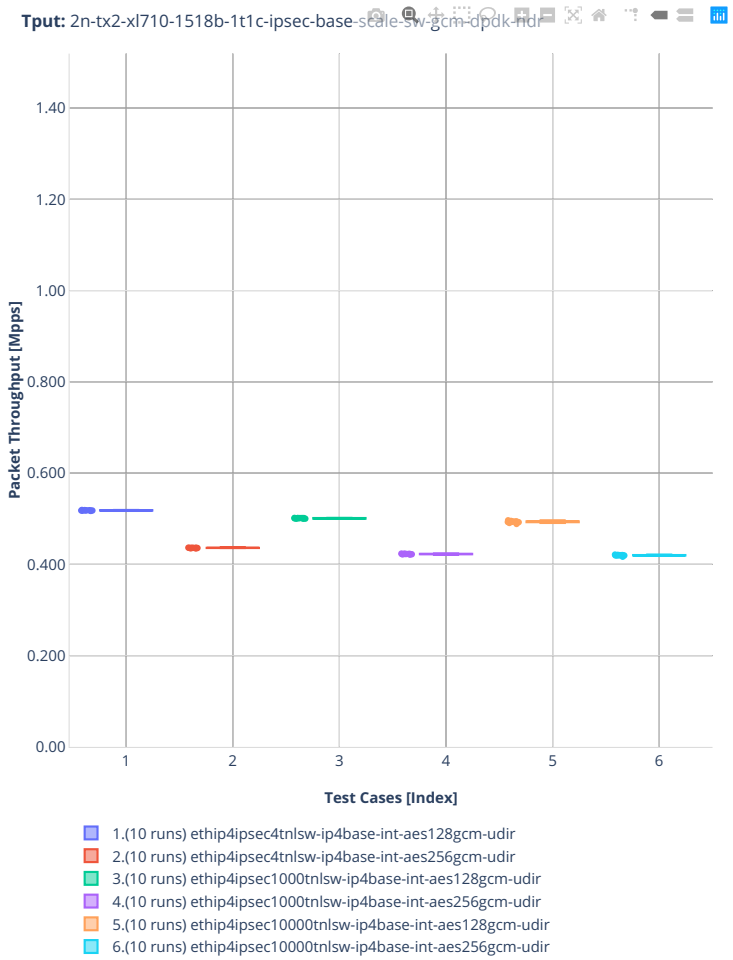

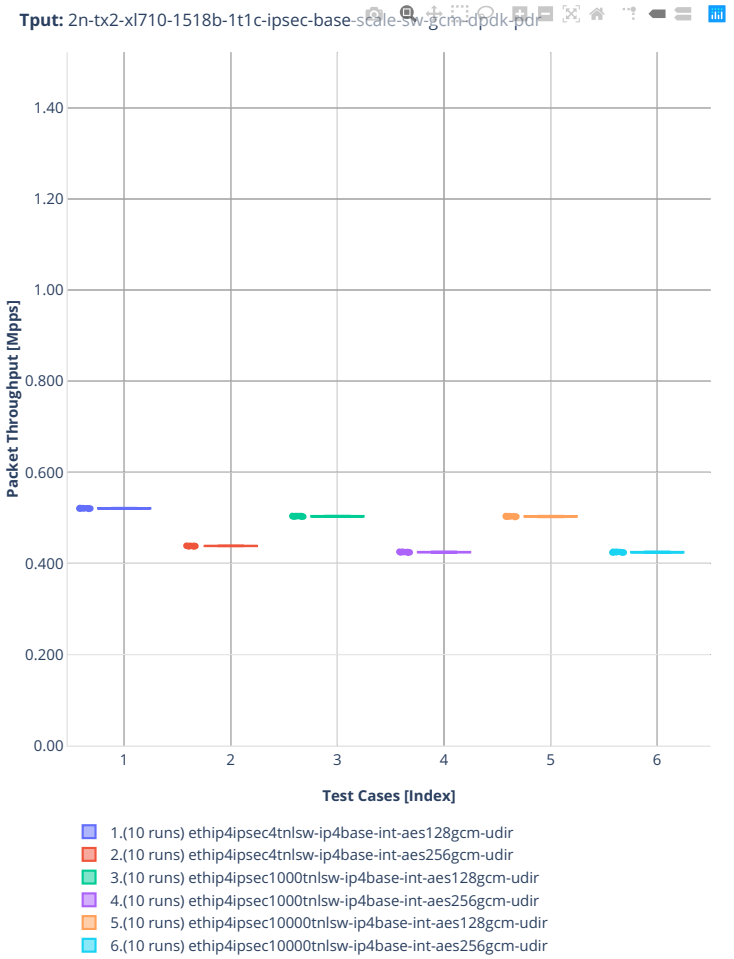
Tput: 2n-tx2-xl710-1518b-1t1c-ipsec-base-scale-sw-gcm-dpdk-pdr
Packet Throughput [Mpps]
1.40
1.20
1.00
0.800
0.600
0.400
0.200
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes128gcm-udir
2.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes256gcm-udir
3.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes128gcm-udir
4.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes256gcm-udir
5.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes128gcm-udir
6.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes256gcm-udir

imix-ipsec-ip4routing-base-scale-sw-gcm

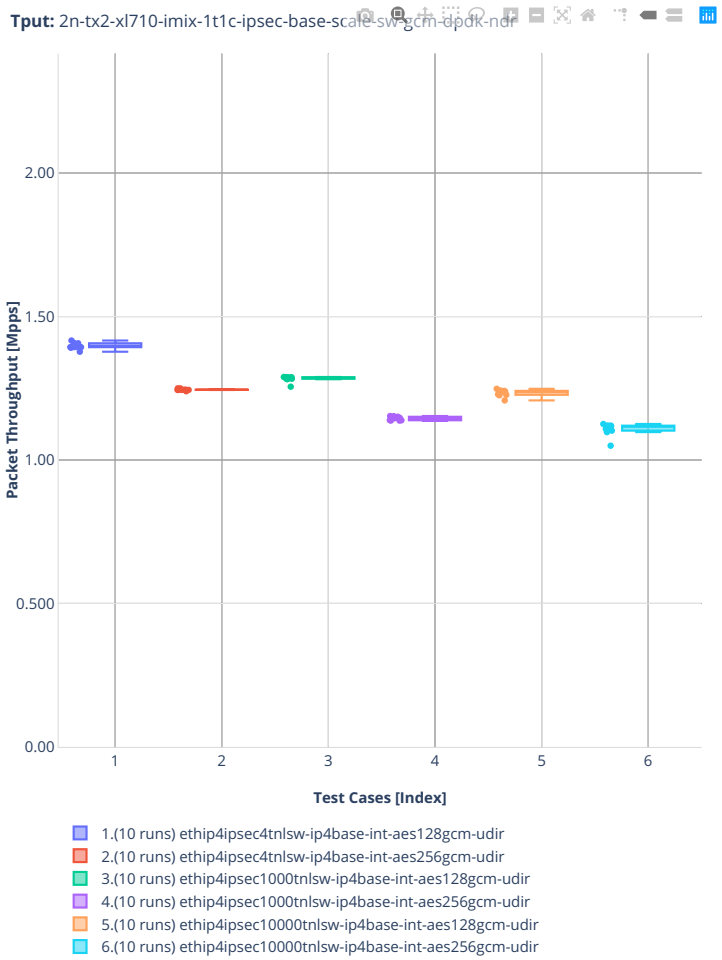

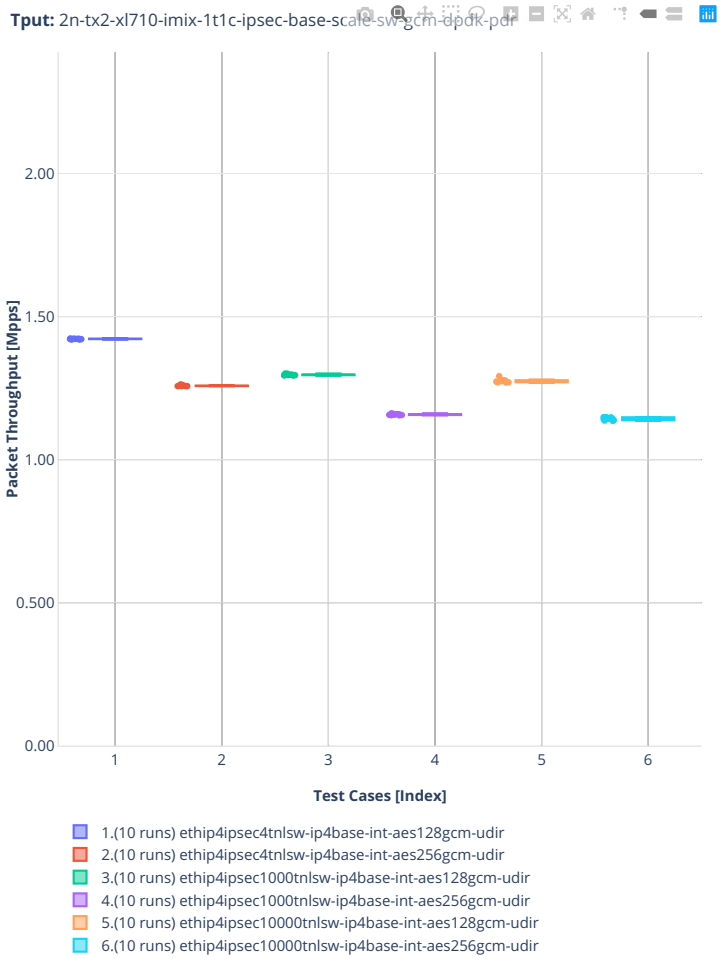
Tput: 2n-tx2-xl710-imix-1t1c-ipsec-base-scale-sw-gcm-dpdk-pdr
Packet Throughput [Mpps]
2.00
1.50
1.00
0.500
0.00
1
2
3
4
5
6
Test Cases [Index]
1.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes128gcm-udir
2.(10 runs) ethip4ipsec4tnlsw-ip4base-int-aes256gcm-udir
3.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes128gcm-udir
4.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes256gcm-udir
5.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes128gcm-udir
6.(10 runs) ethip4ipsec10000tnlsw-ip4base-int-aes256gcm-udir

3n-dnv-x553

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

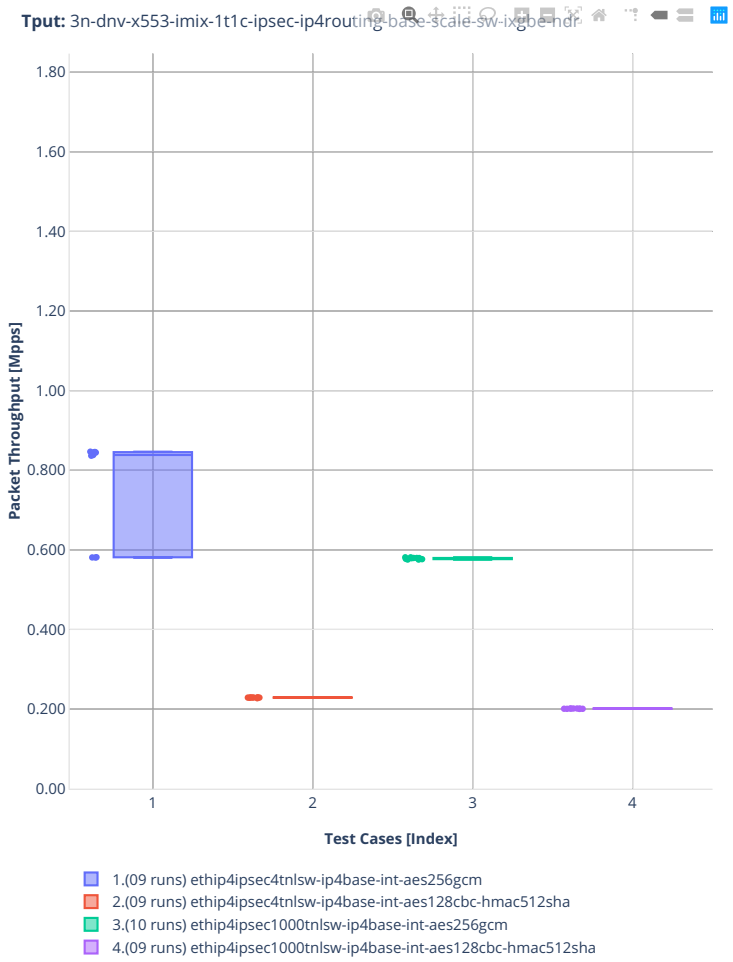

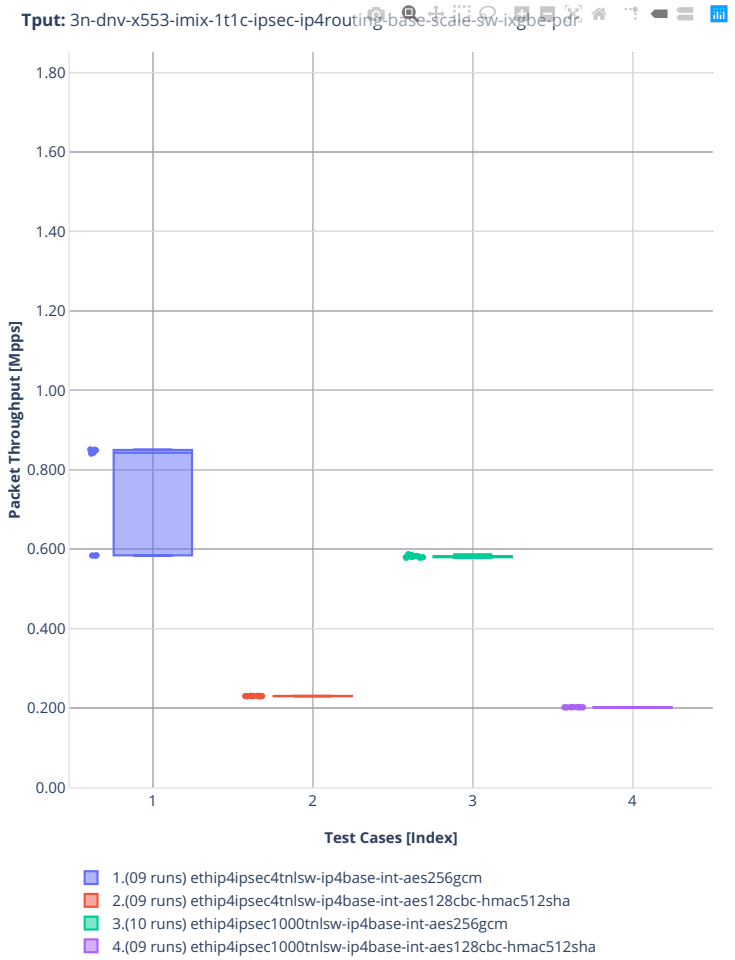
Tput: 3n-dnv-x553-imix-1t1c-ipsec-ip4routing-base-scale-sw-ixgbe-pdr
Packet Throughput [Mpps]
1.80
1.60
1.40
1.20
1.00
0.800
0.600
0.400
0.200
0.00
1
2
3
4
Test Cases [Index]
1.(09 runs) ethip4ipsec4tnlsw-ip4base-int-aes256gcm
2.(09 runs) ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
3.(10 runs) ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
4.(09 runs) ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha

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: 2n-icx, 3n-icx, 2n-aws, 2n-skx, 3n-skx, 2n-clx, 2n-zn2, 3n-alt, 3n-tsh, 2n-tx2, 2n-dnv, 3n-dnv. 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 are stored in build logs from FD.io vpp performance job 2n-icx[118], build logs from FD.io vpp performance job 3n-icx[119], build logs from FD.io vpp performance job 2n-aws[120], build logs from FD.io vpp performance job 2n-skx[121], build logs from FD.io vpp performance job 3n-skx[122], build logs from FD.io vpp performance job 2n-clx[123], build logs from FD.io vpp performance job 2n-zn2[124], build logs from FD.io vpp performance job 3n-alt[125], build logs from FD.io vpp performance job 3n-tsh[126], build logs from FD.io vpp performance job 2n-tx2[127], build logs from FD.io vpp performance job 2n-dnv[128] and build logs from FD.io vpp performance job 3n-dnv[129] with RF result files csit-vpp-perf-2206-*.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.

---

[118] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-icx
[119] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-icx
[120] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-aws
[121] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx
[122] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-skx
[123] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-clx
[124] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-zn2
[125] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-alt
[126] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-tsh
[127] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-tx2
[128] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-dnv
[129] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-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[130].

[130] https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls2206

2n-icx-xxv710

64b-l2switching-base-scale-avf

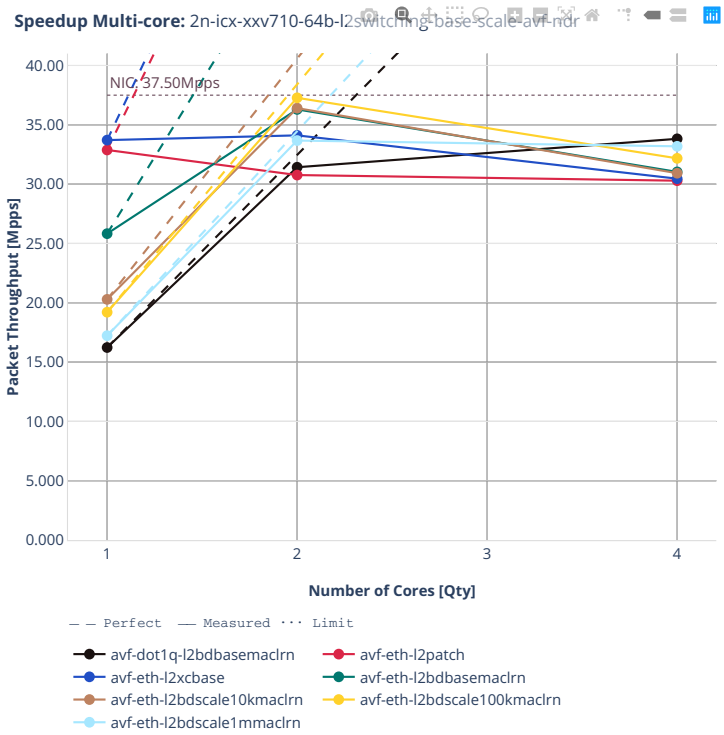

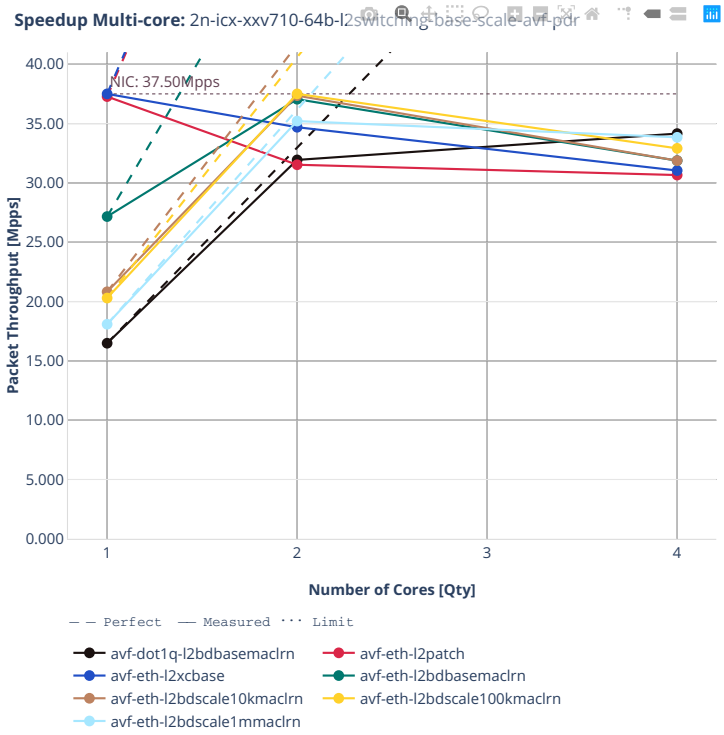
Speedup Multi-core: 2n-icx-xxv710-64b-l2switching-base-scale-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-dot1q-l2bdbasemaclrn
avf-eth-l2patch
avf-eth-l2xcbase
avf-eth-l2bdbasemaclrn
avf-eth-l2bdscale10kmaclrn
avf-eth-l2bdscale100kmaclrn
avf-eth-l2bdscale1mmaclrn

64b-l2switching-base-scale-dpdk

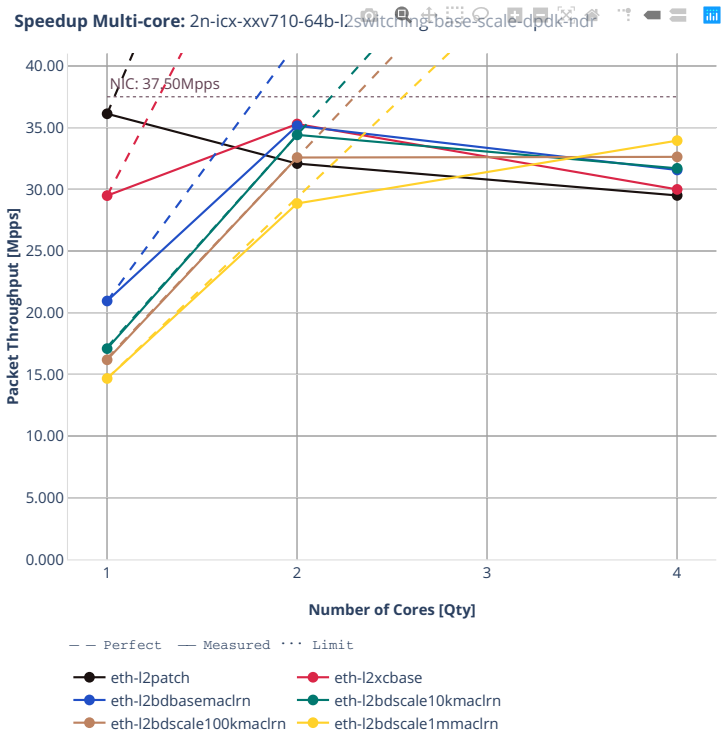

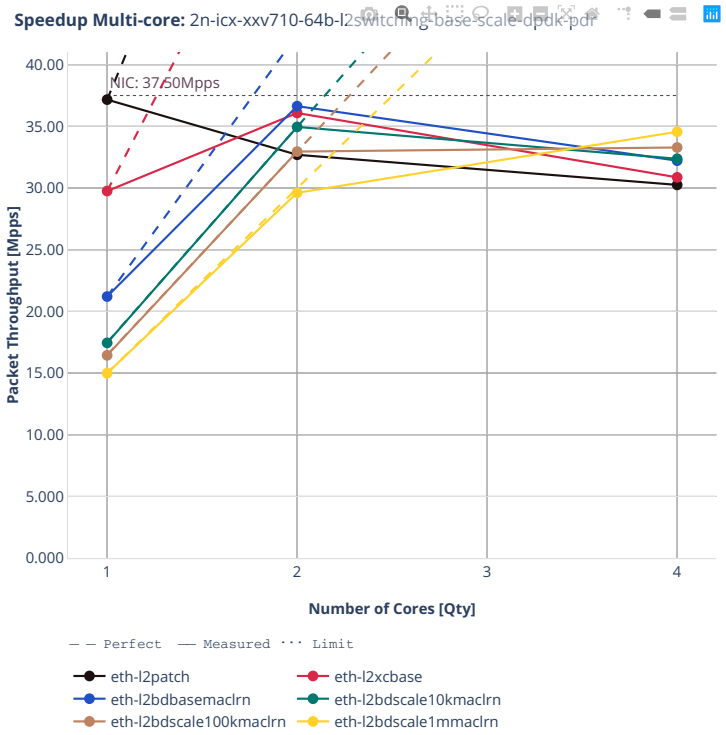
Speedup Multi-core: 2n-icx-xxv710-64b-l2switching-base-scale-dpdk-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2patch
eth-l2xcbase
eth-l2bdbasemaclrn
eth-l2bdscale10kmaclrn
eth-l2bdscale100kmaclrn
eth-l2bdscale1mmaclrn

3n-icx-xxv710

64b-l2switching-base

2n-skx-xxv710

64b-l2switching-base-avf

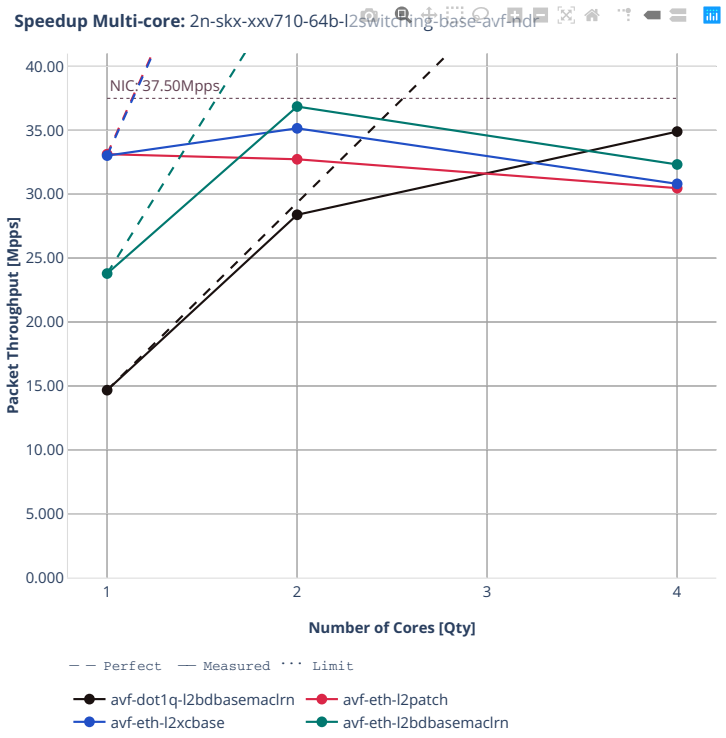

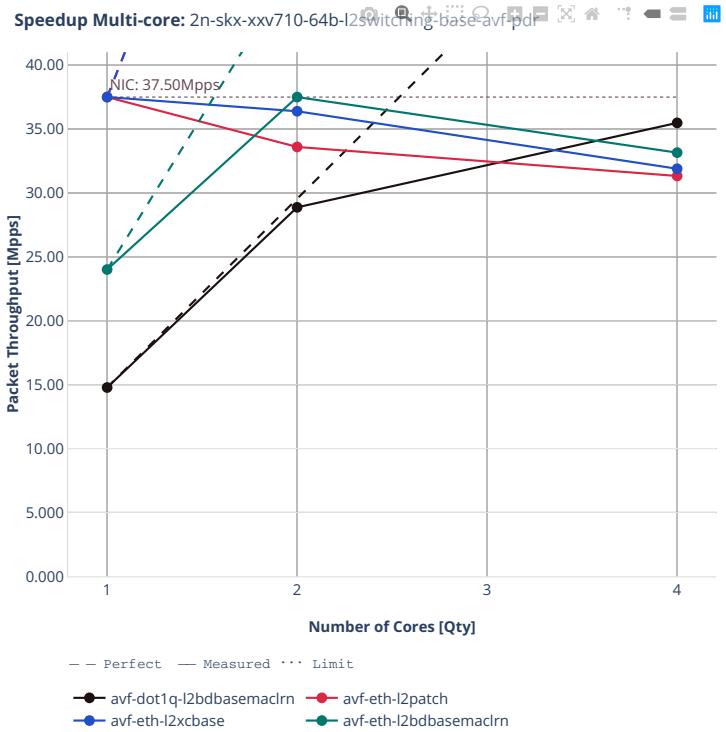
Speedup Multi-core: 2n-skx-xxv710-64b-l2switching-base-avf-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-dot1q-l2bdbasemaclrn
avf-eth-l2patch
avf-eth-l2xcbase
avf-eth-l2bdbasemaclrn

64b-l2switching-base-dpdk

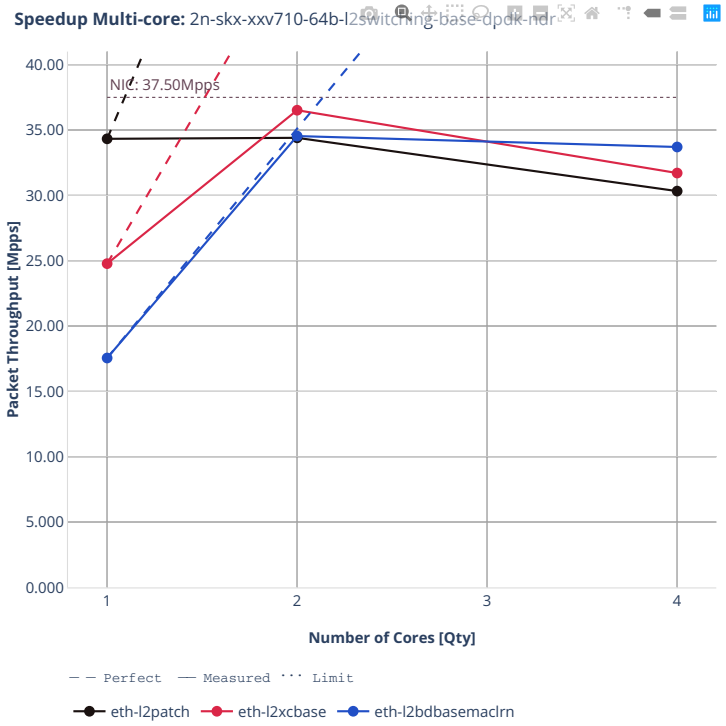

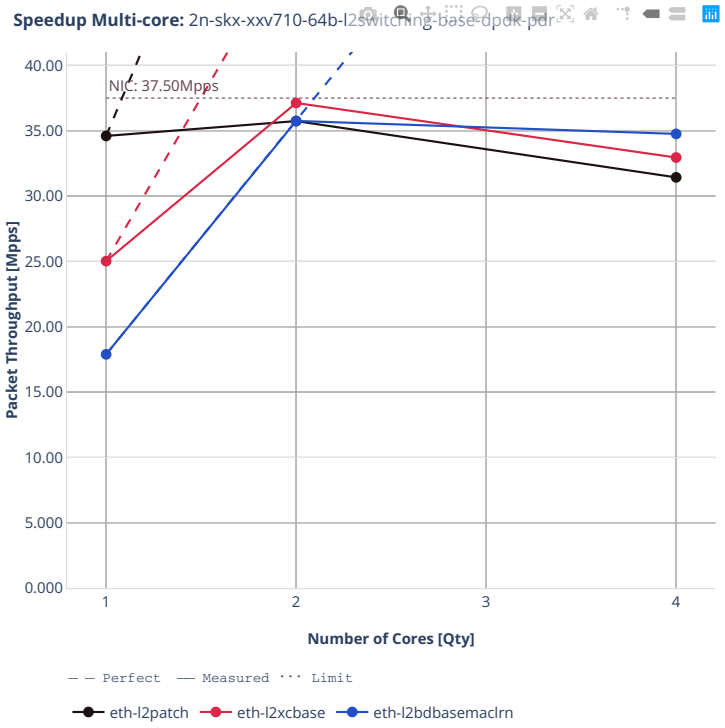
Speedup Multi-core: 2n-skx-xxv710-64b-l2switching-base-dpdk-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2patch
eth-l2xcbase
eth-l2bdbasemaclrn

64b-l2switching-base-scale-avf

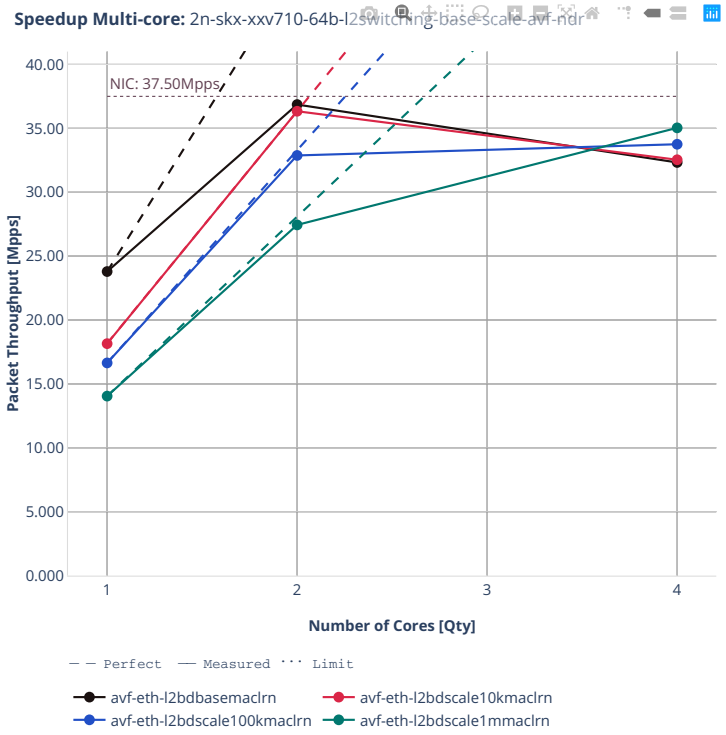

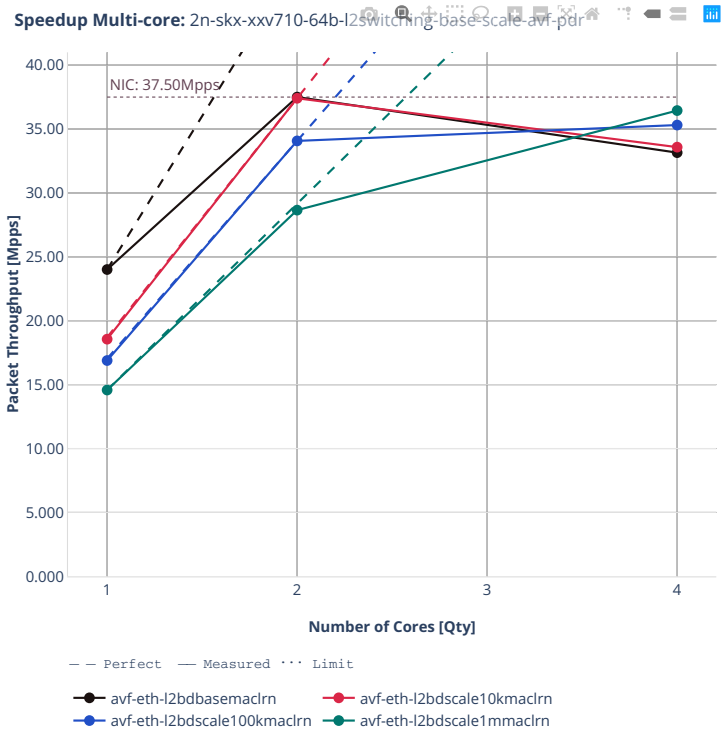
Speedup Multi-core: 2n-skx-xxv710-64b-l2switching-base-scale-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2bdbasemaclrn
avf-eth-l2bdscale10kmaclrn
avf-eth-l2bdscale100kmaclrn
avf-eth-l2bdscale1mmaclrn

64b-l2switching-base-scale-dpdk

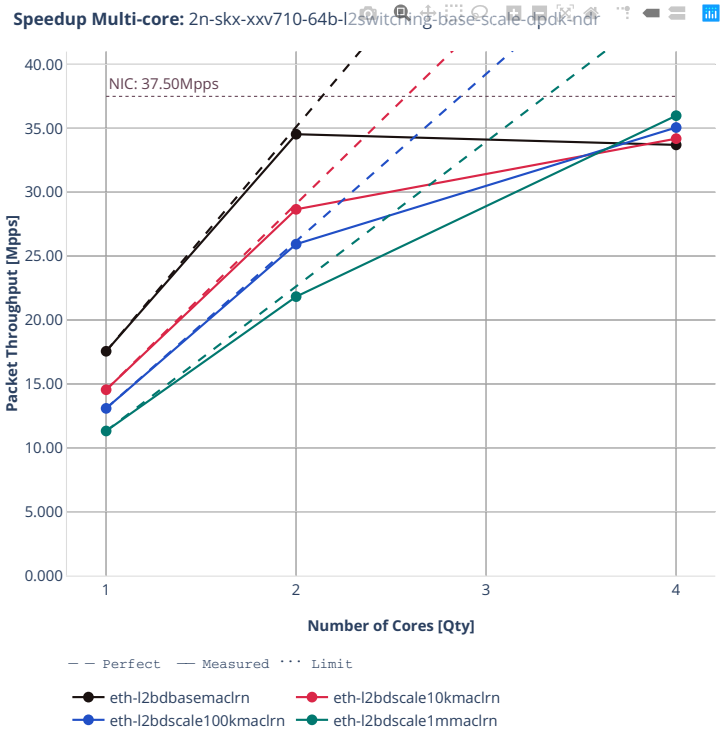

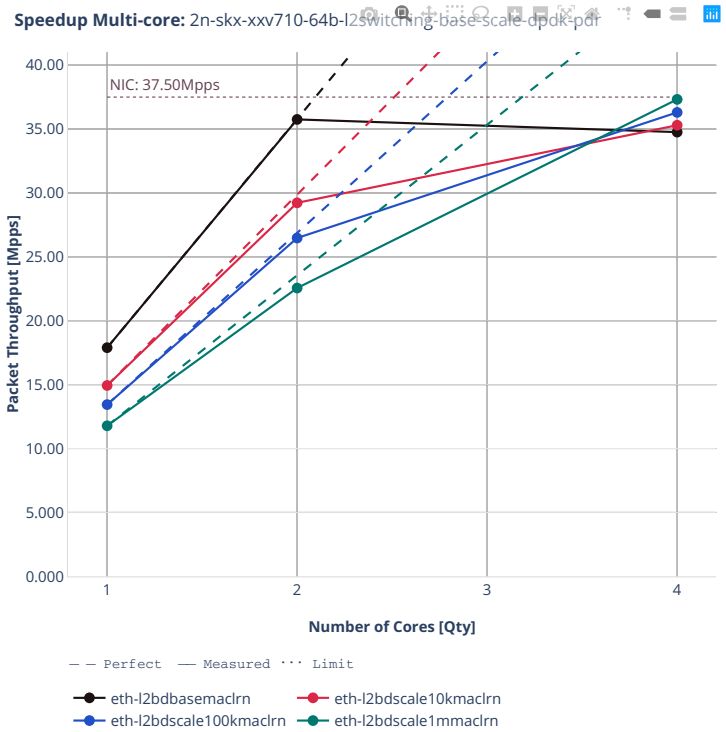
Speedup Multi-core: 2n-skx-xxv710-64b-l2switching-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2bdbasemaclrn
eth-l2bdscale10kmaclrn
eth-l2bdscale100kmaclrn
eth-l2bdscale1mmaclrn

2n-skx-x710

64b-l2switching-base-scale-[avf,dpdk]

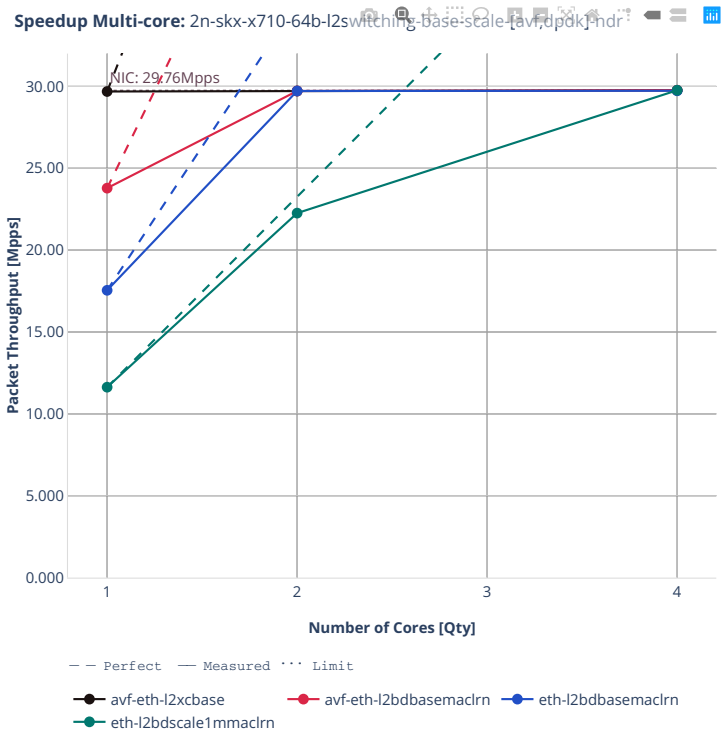

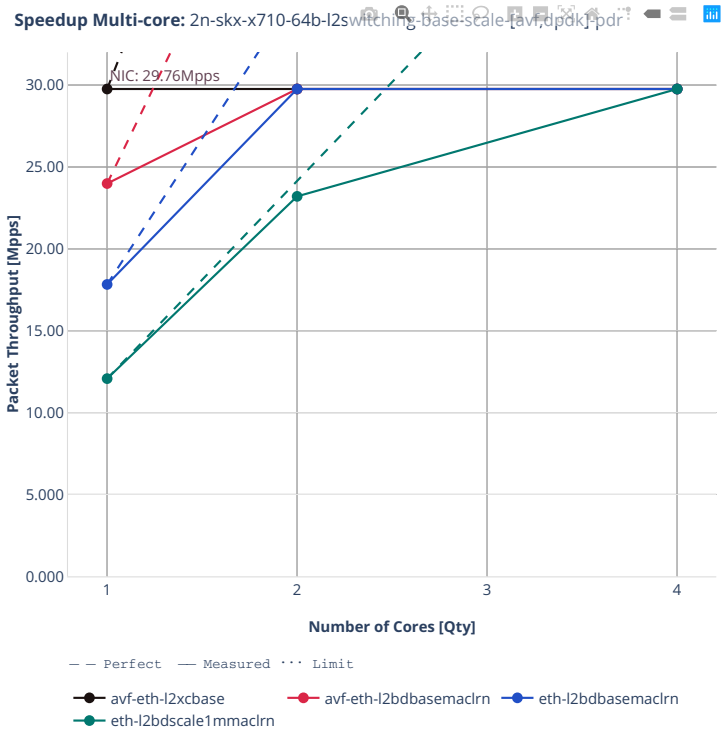
Speedup Multi-core: 2n-skx-x710-64b-l2switching-base-scale-[avf,dpdk]-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2xcbase
avf-eth-l2bdbasemaclrn
eth-l2bdbasemaclrn
eth-l2bdscale1mmaclrn

3n-skx-xxv710

64b-l2switching-base

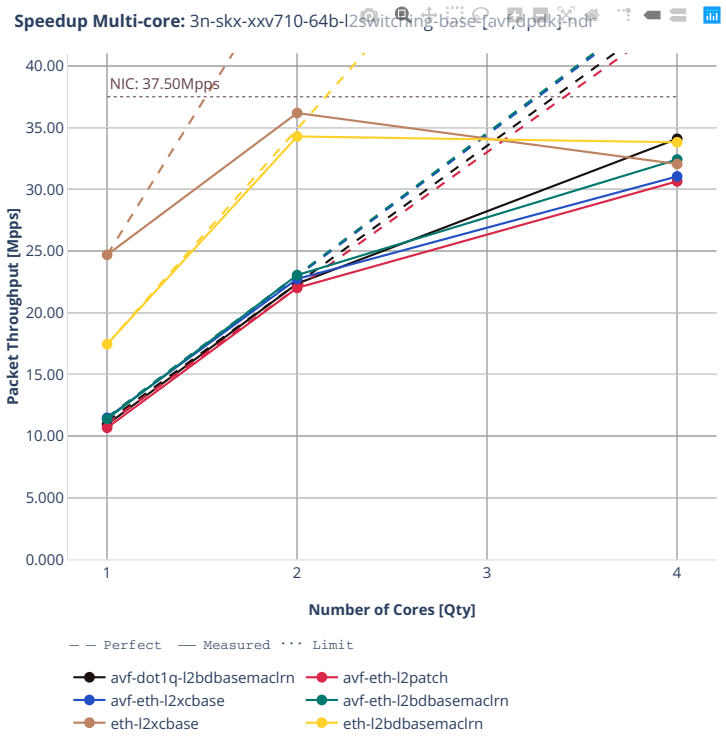

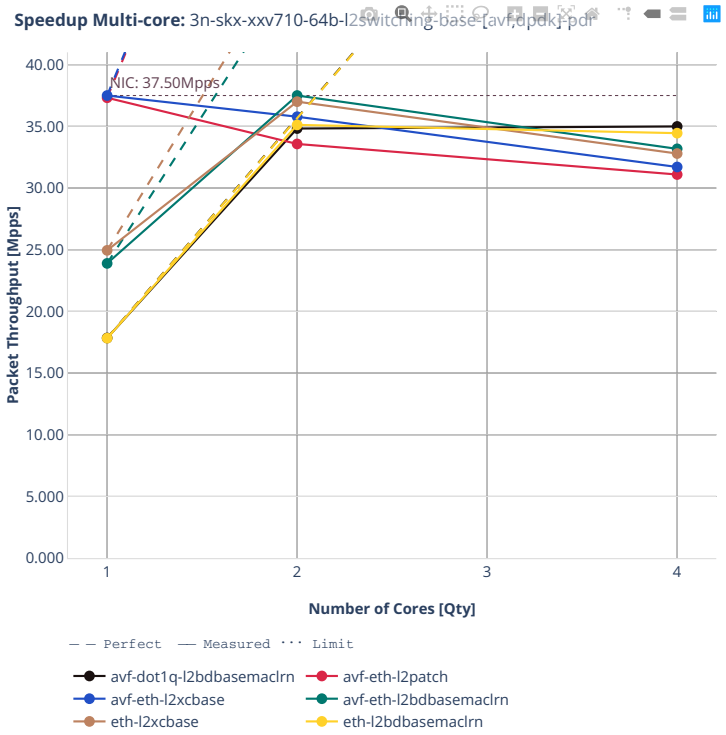
Speedup Multi-core: 3n-skx-xxv710-64b-l2switching-base-[avf,dpdk]-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-dot1q-l2bdbasemaclrn
avf-eth-l2patch
avf-eth-l2xcbase
avf-eth-l2bdbasemaclrn
eth-l2xcbase
eth-l2bdbasemaclrn

3n-skx-x710

64b-l2switching-base-avf

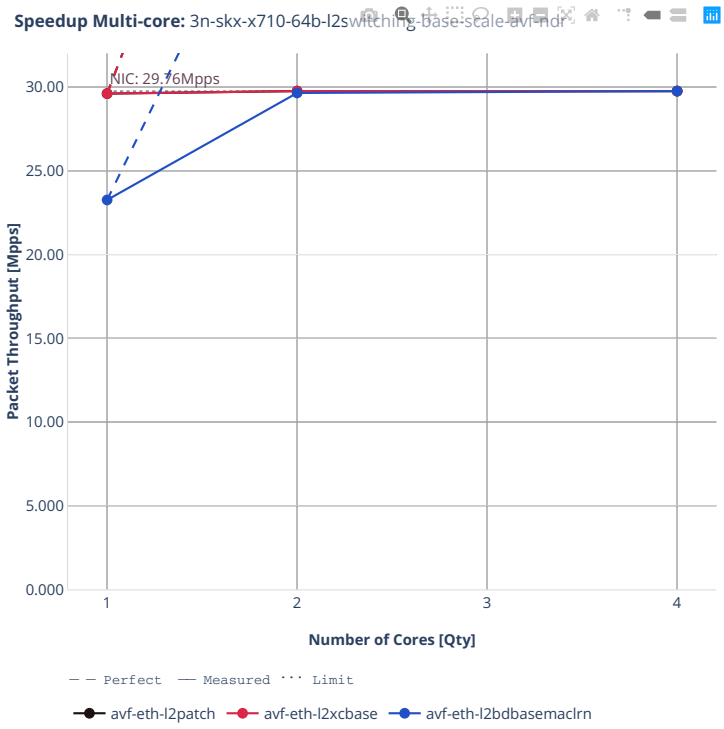

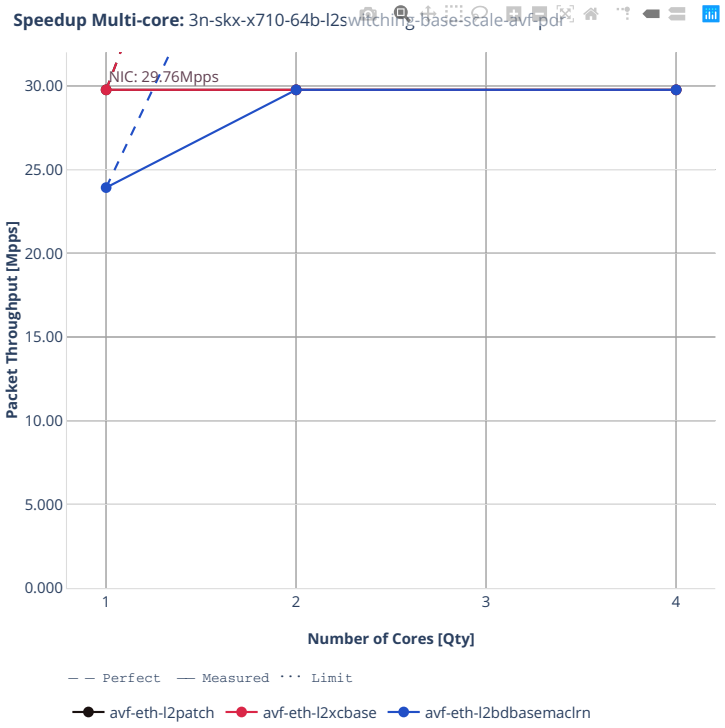
Speedup Multi-core: 3n-skx-x710-64b-l2switching-base-scale-avf-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2patch
avf-eth-l2xcbase
avf-eth-l2bdbasemaclrn

2n-clx-xxv710

64b-l2switching-base-avf

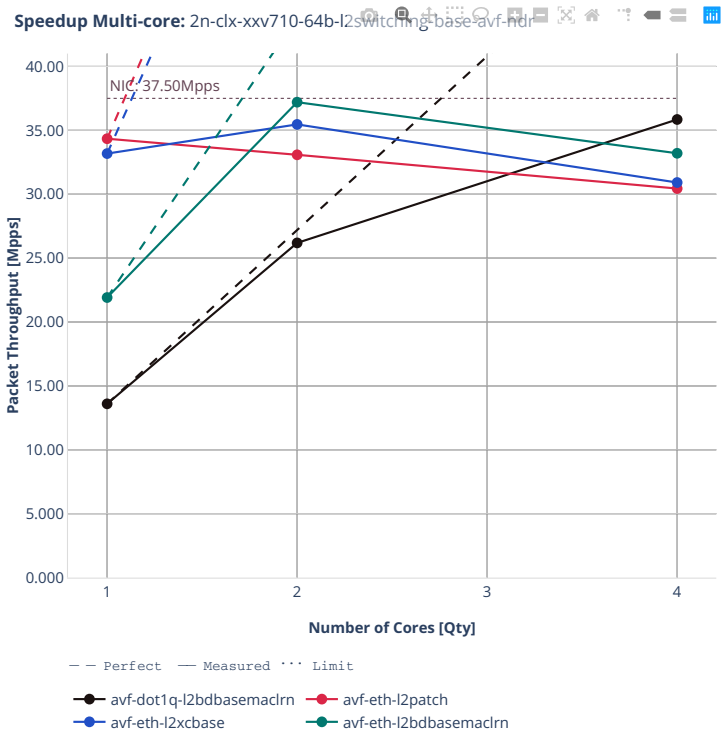

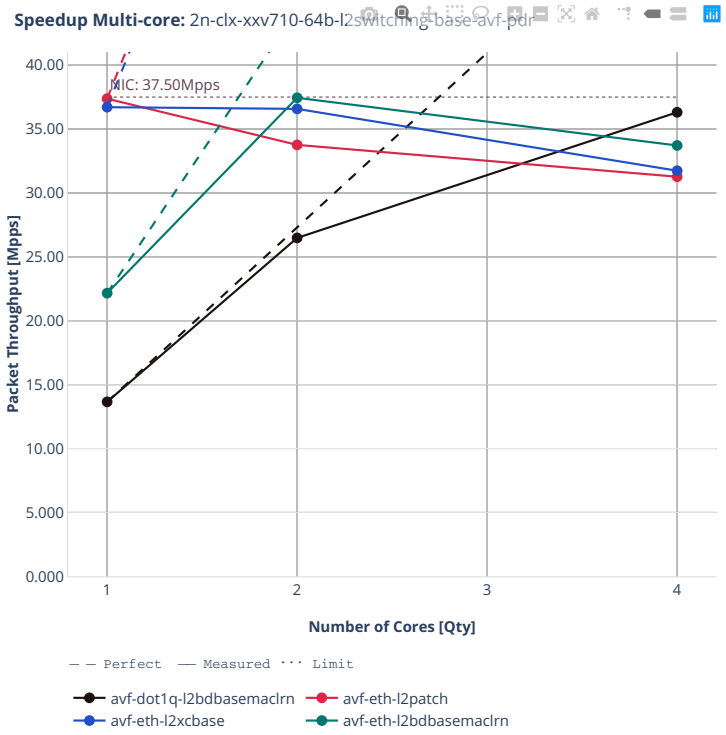
Speedup Multi-core: 2n-clx-xxv710-64b-l2switching-base-avf-pdr
MIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-dot1q-l2bdbasemaclrn
avf-eth-l2patch
avf-eth-l2xcbase
avf-eth-l2bdbasemaclrn

64b-l2switching-base-scale-avf

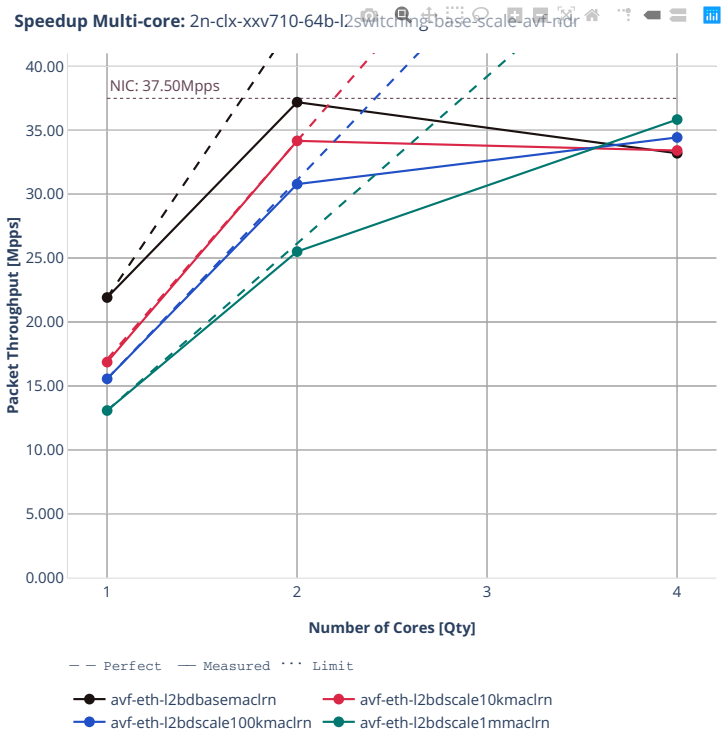

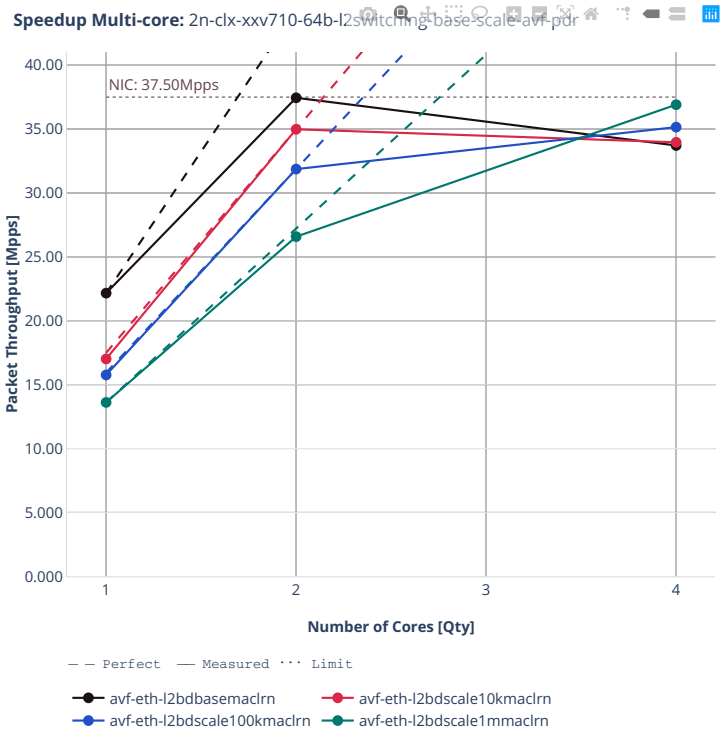
Speedup Multi-core: 2n-clx-xxv710-64b-l2switching-base-scale-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2bdbasemaclrn
avf-eth-l2bdscale10kmaclrn
avf-eth-l2bdscale100kmaclrn
avf-eth-l2bdscale1mmaclrn

64b-l2switching-base-dpdk

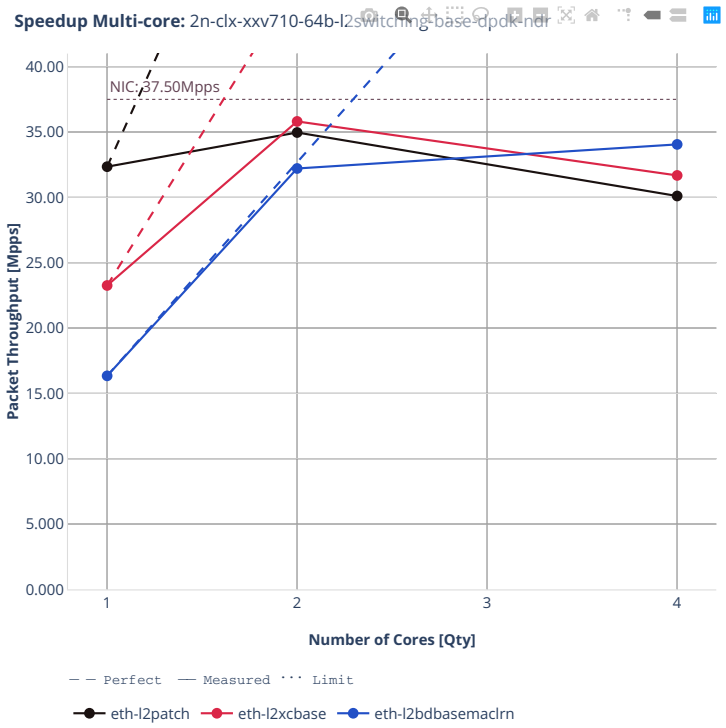

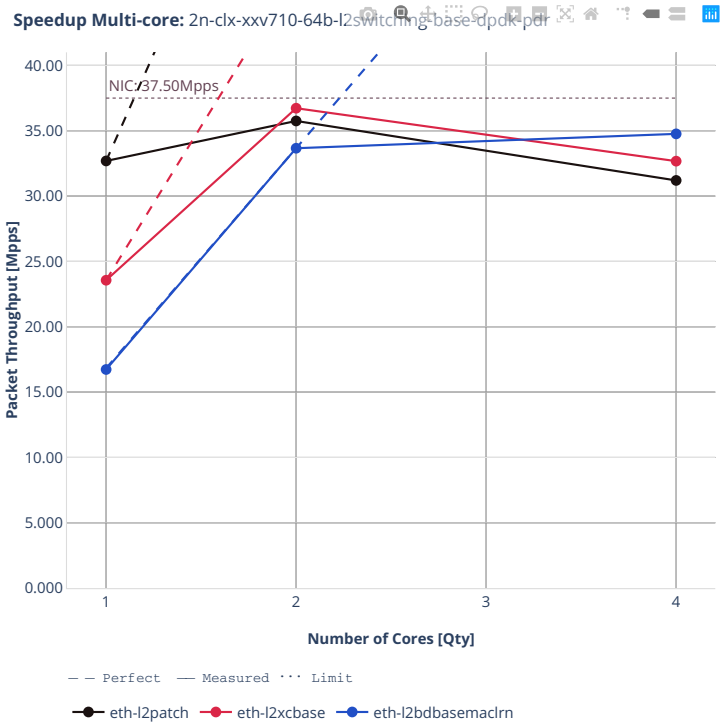
Speedup Multi-core: 2n-clx-xxv710-64b-l2switching-base-dpdk-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2patch
eth-l2xcbase
eth-l2bdbasemaclrn

64b-l2switching-base-scale-dpdk

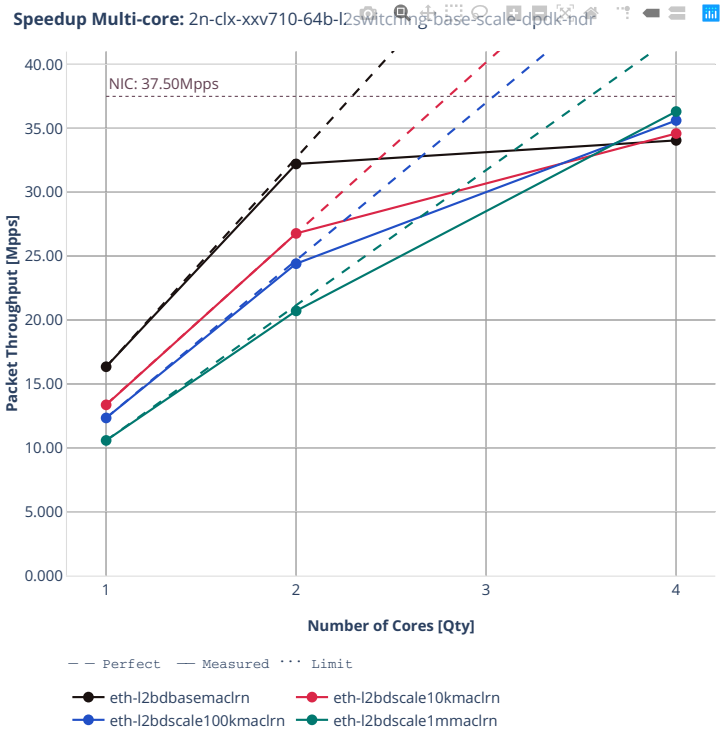

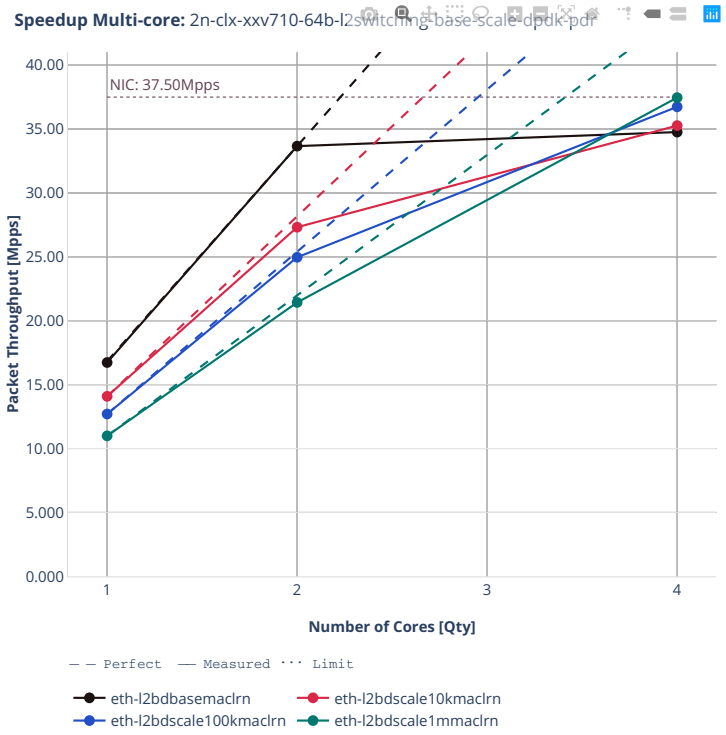
Speedup Multi-core: 2n-clx-xxv710-64b-l2switching-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2bdbasemaclrn
eth-l2bdscale10kmaclrn
eth-l2bdscale100kmaclrn
eth-l2bdscale1mmaclrn

2n-clx-x710

64b-l2switching-base-avf

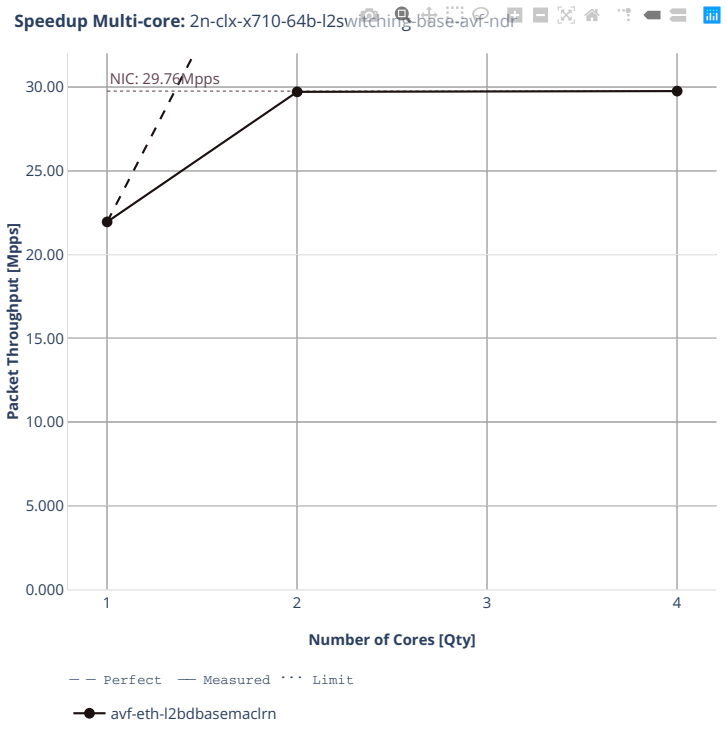

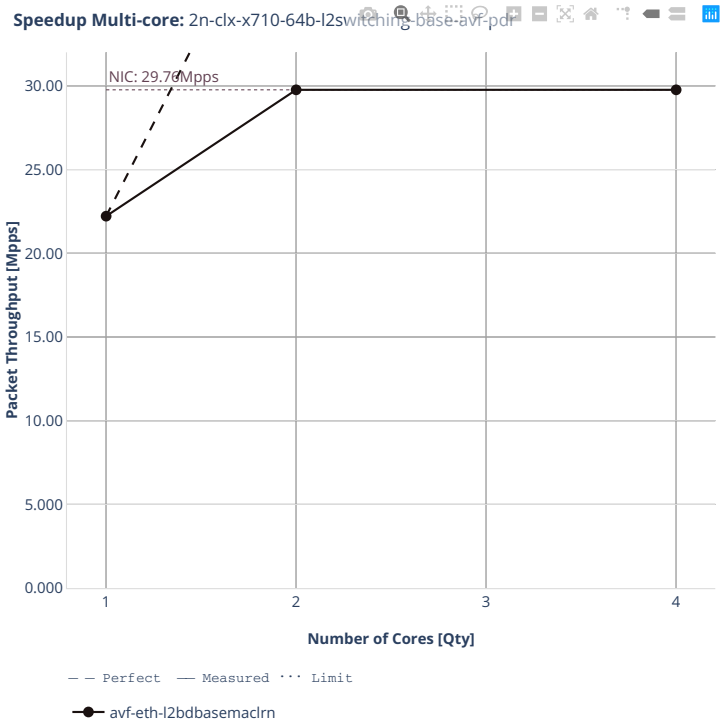
Speedup Multi-core: 2n-clx-x710-64b-l2switching-base-avf-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2bdbasemaclrn

2n-clx-cx556a

64b-l2switching-base-rdma-core

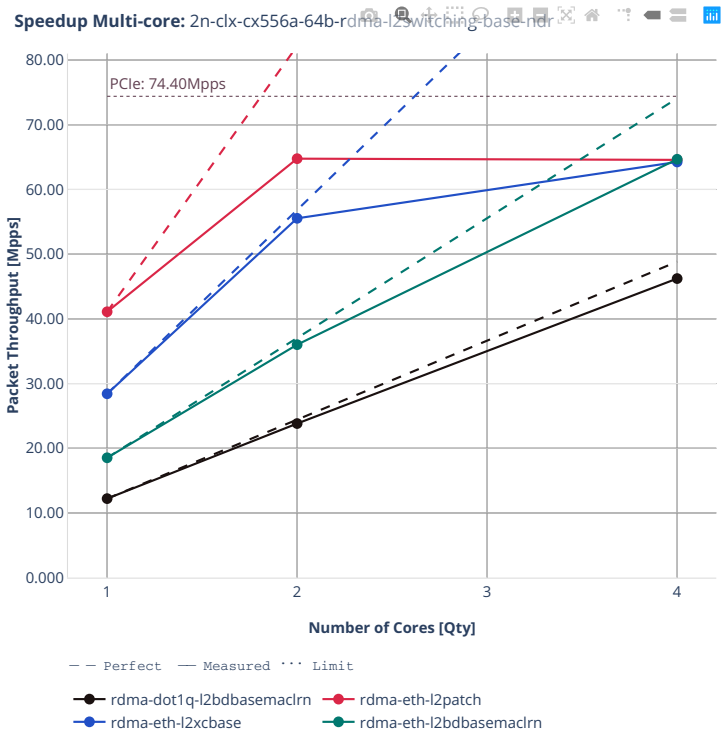

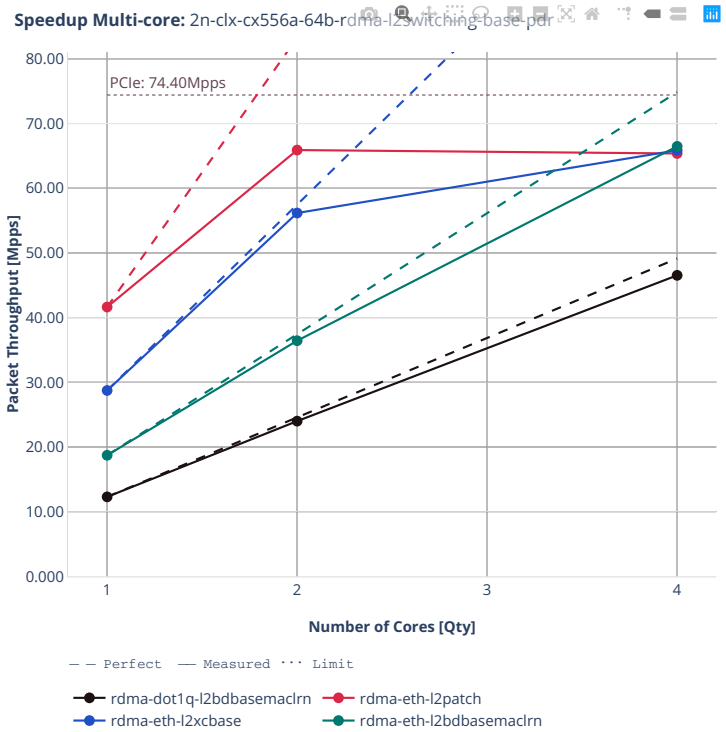
Speedup Multi-core: 2n-clx-cx556a-64b-rdma-l2switching-base-pdr
PCIe: 74.40Mpps
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-dot1q-l2bdbasemaclrn
rdma-eth-l2patch
rdma-eth-l2xcbase
rdma-eth-l2bdbasemaclrn

64b-l2switching-scale

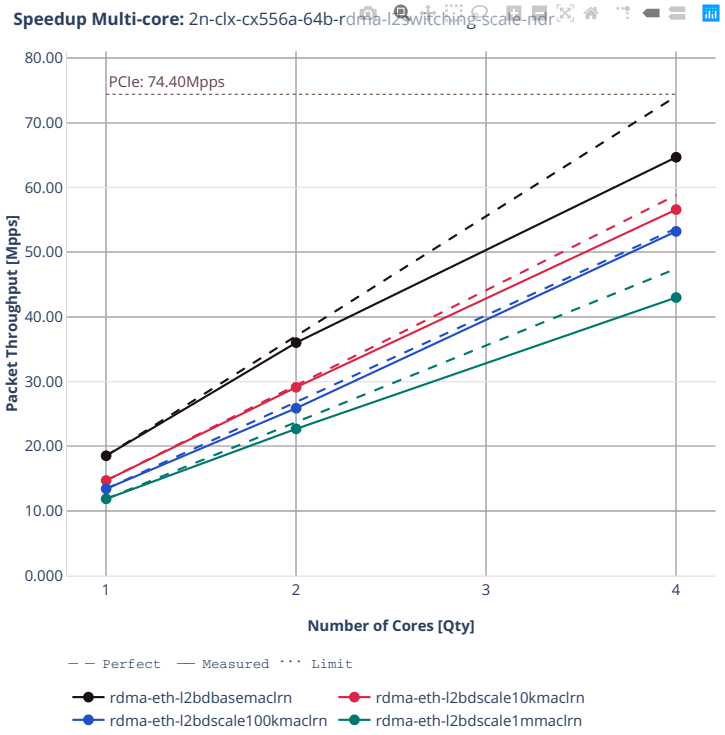

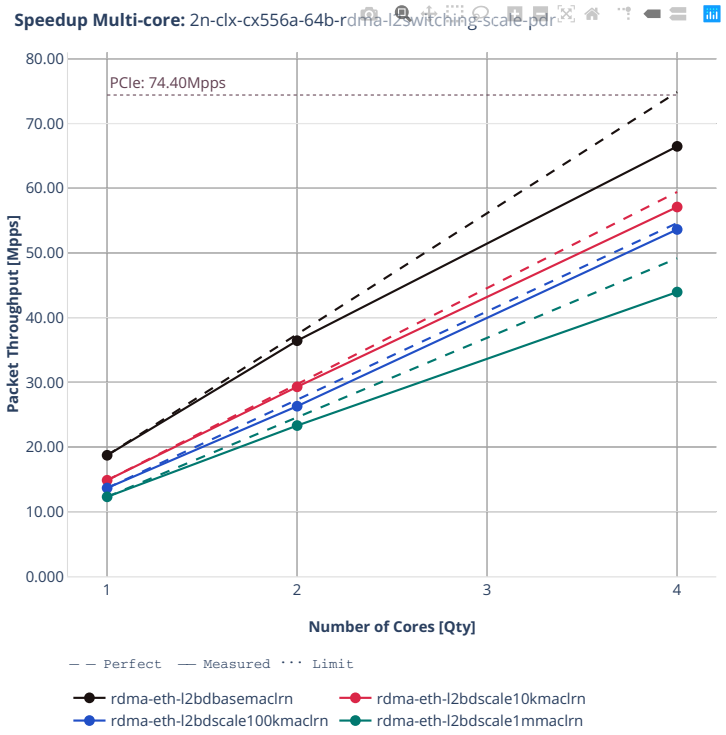
Speedup Multi-core: 2n-clx-cx556a-64b-rdma-l2switching-scale-pdr
PCIe: 74.40Mpps
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-eth-l2bdbasemaclrn
rdma-eth-l2bdscale10kmaclrn
rdma-eth-l2bdscale100kmaclrn
rdma-eth-l2bdscale1mmaclrn

2n-clx-e810cq

64b-l2switching-avf

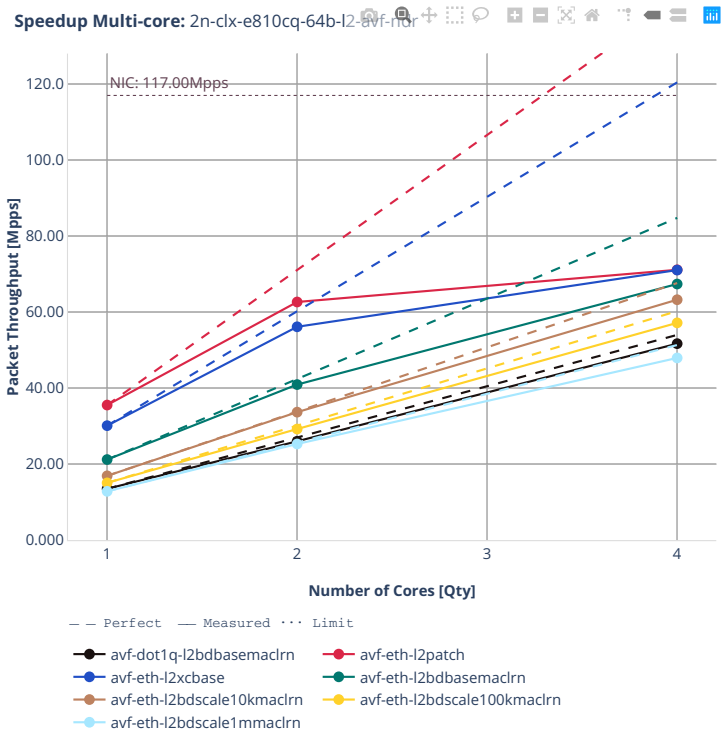

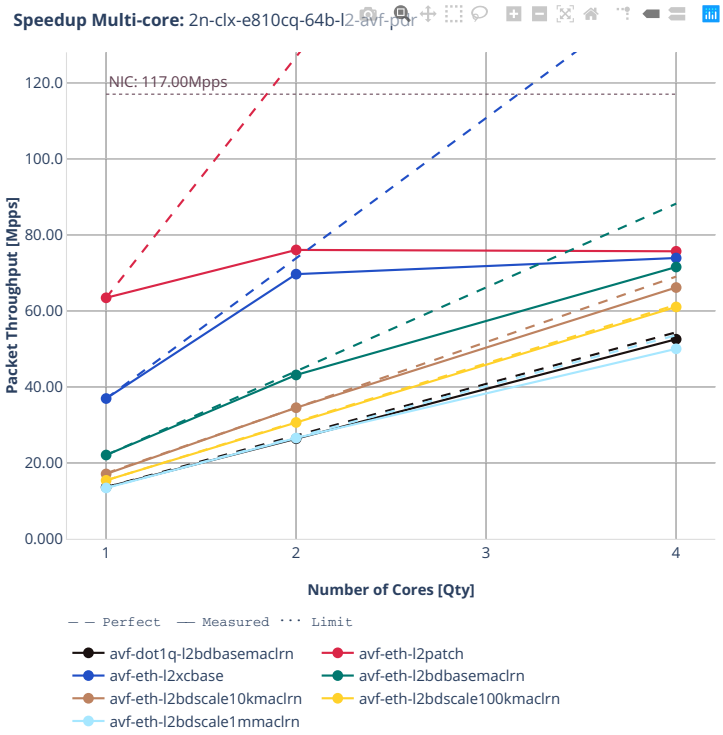
Speedup Multi-core: 2n-clx-e810cq-64b-l2-avf-pdr
NIC: 117.00Mpps
120.0
100.0
80.00
60.00
40.00
20.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-dot1q-l2bdbasemaclrn
avf-eth-l2patch
avf-eth-l2xcbase
avf-eth-l2bdbasemaclrn
avf-eth-l2bdscale10kmaclrn
avf-eth-l2bdscale100kmaclrn
avf-eth-l2bdscale1mmaclrn

64b-l2switching-dpdk

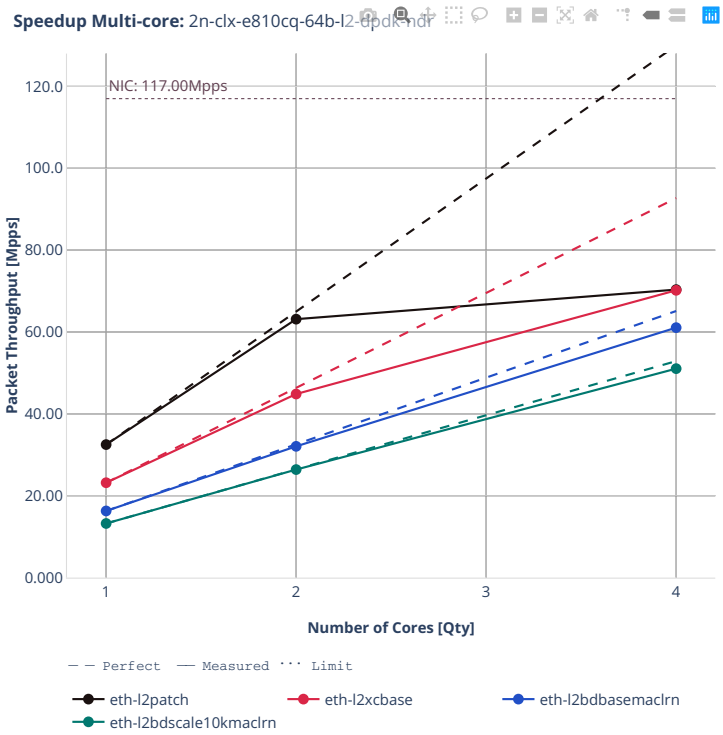

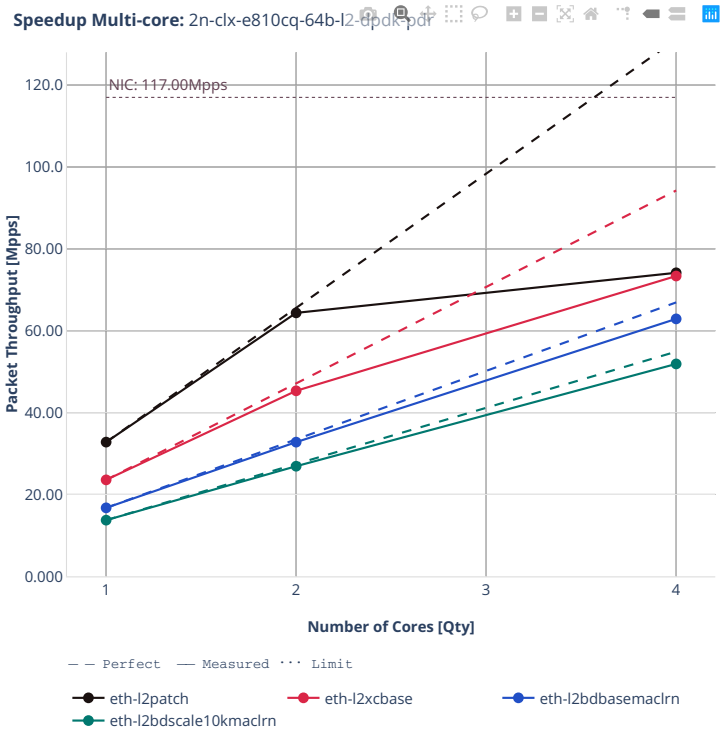
Speedup Multi-core: 2n-clx-e810cq-64b-l2-dpdk-pdr
NIC: 117.00Mpps
120.0
100.0
80.00
60.00
40.00
20.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2patch
eth-l2xcbase
eth-l2bdbasemaclrn
eth-l2bdscale10kmaclrn

2n-zn2-xxv710

64b-l2switching-base-avf

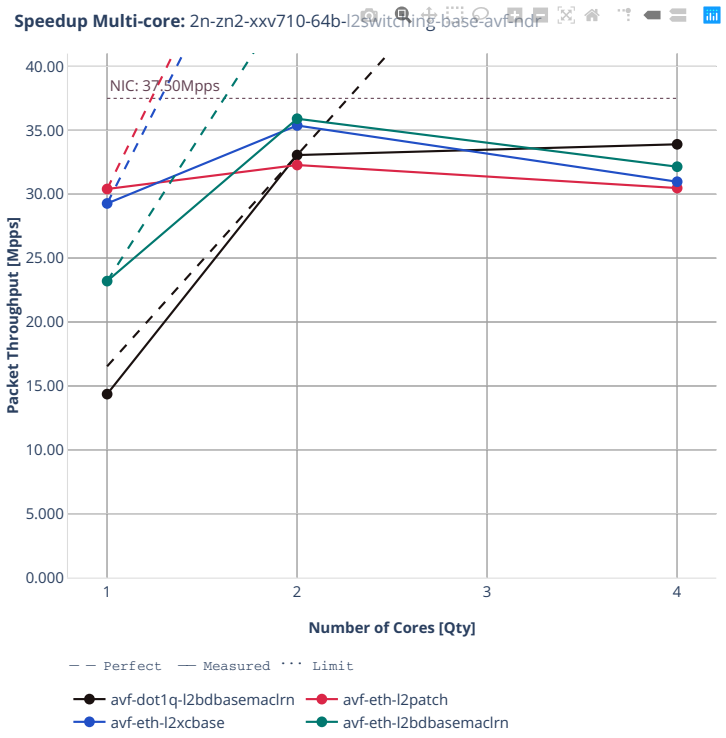

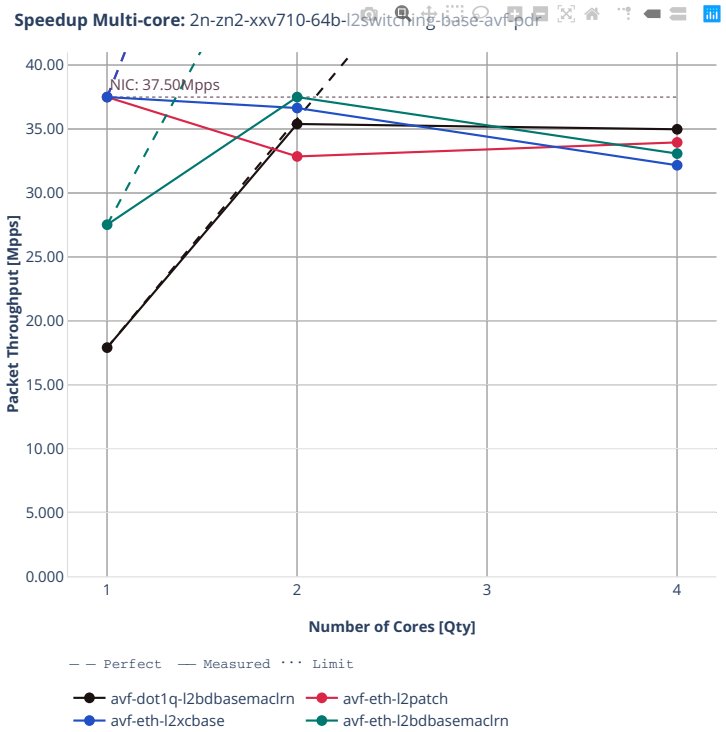
Speedup Multi-core: 2n-zn2-xxv710-64b-l2switching-base-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-dot1q-l2bdbasemaclrn
avf-eth-l2patch
avf-eth-l2xcbase
avf-eth-l2bdbasemaclrn

64b-l2switching-base-scale-avf

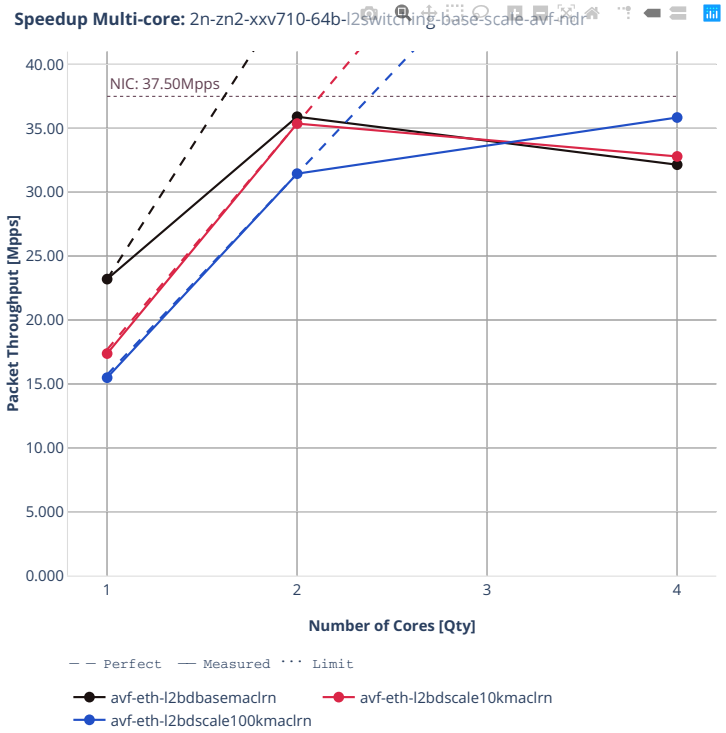

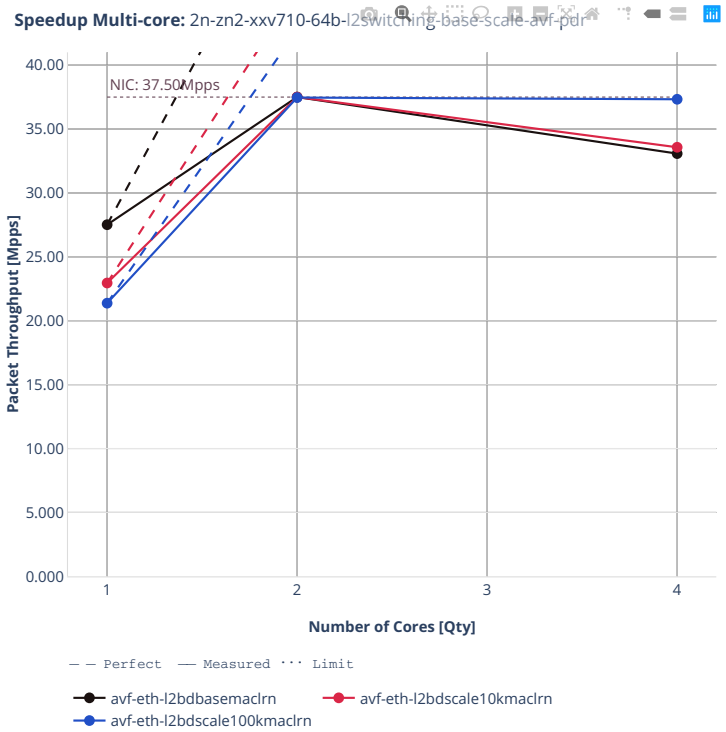
Speedup Multi-core: 2n-zn2-xxv710-64b-l2switching-base-scale-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2bdbasemaclrn
avf-eth-l2bdscale10kmaclrn
avf-eth-l2bdscale100kmaclrn

64b-l2switching-base-dpdk

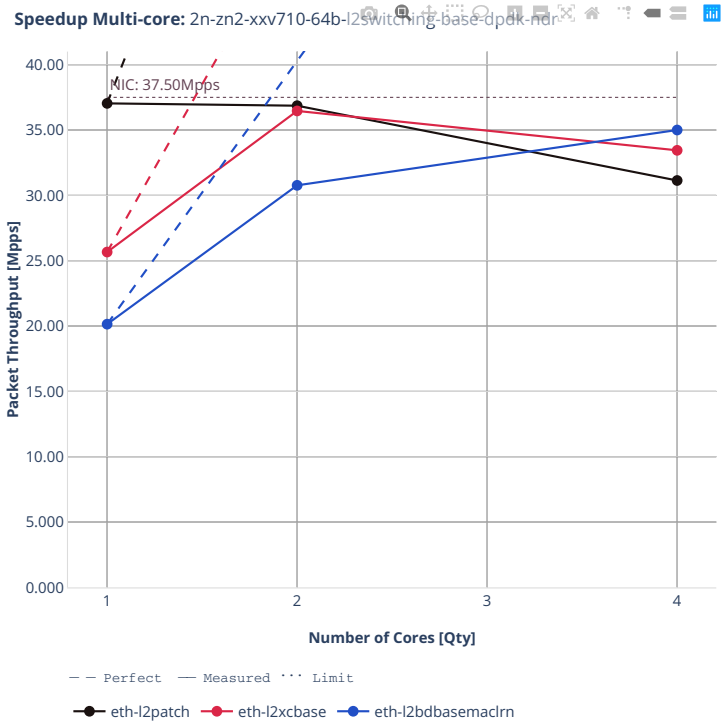

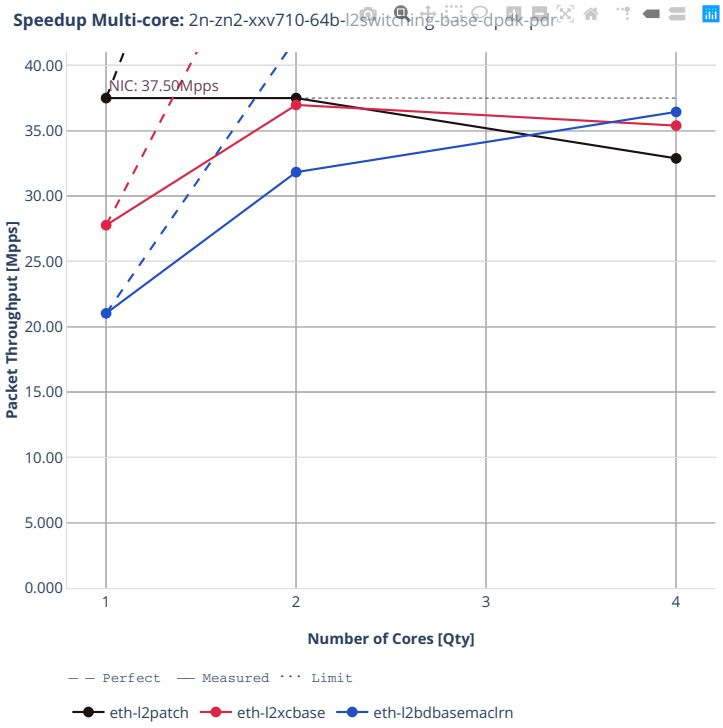
Speedup Multi-core: 2n-zn2-xxv710-64b-l2switching-base-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2patch
eth-l2xcbase
eth-l2bdbasemaclrn

64b-l2switching-base-scale-dpdk

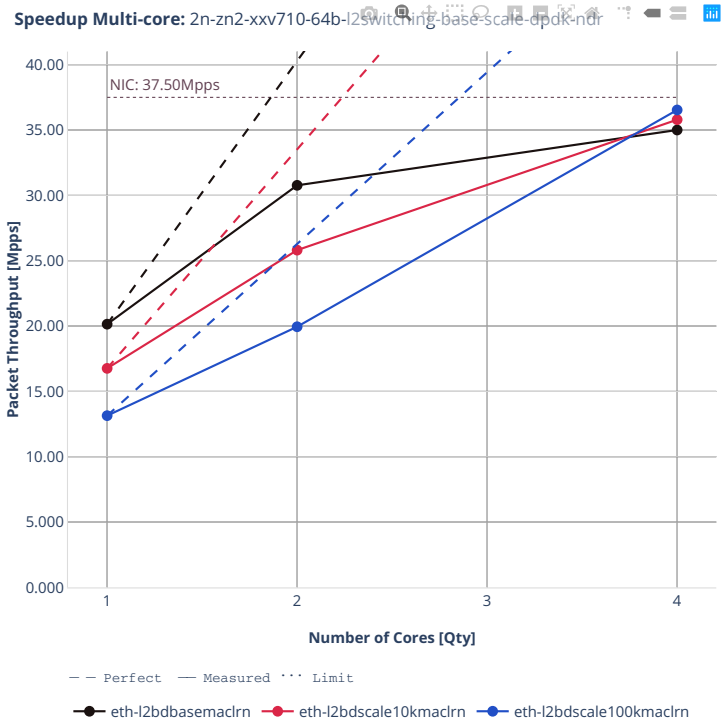

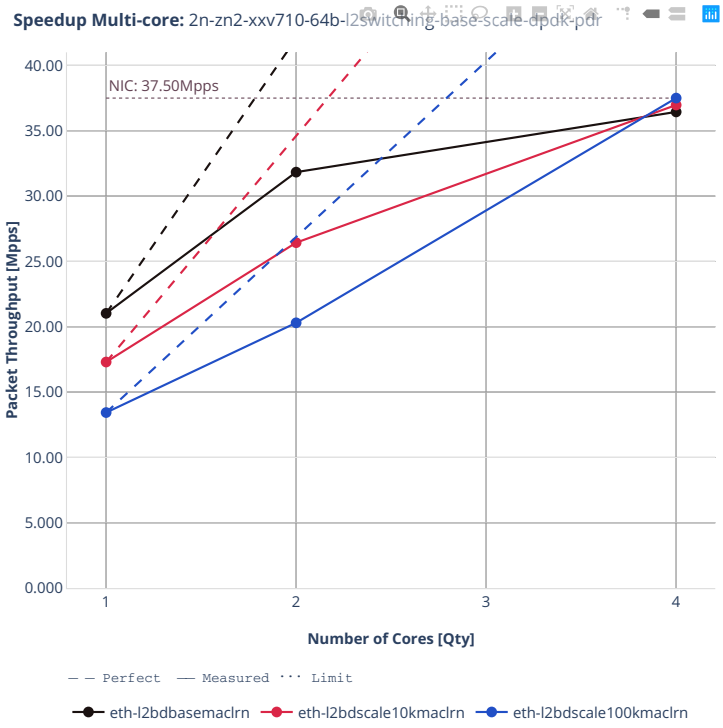
Speedup Multi-core: 2n-zn2-xxv710-64b-l2switching-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2bdbasemaclrn
eth-l2bdscale10kmaclrn
eth-l2bdscale100kmaclrn

2n-zn2-x710

64b-l2switching-base

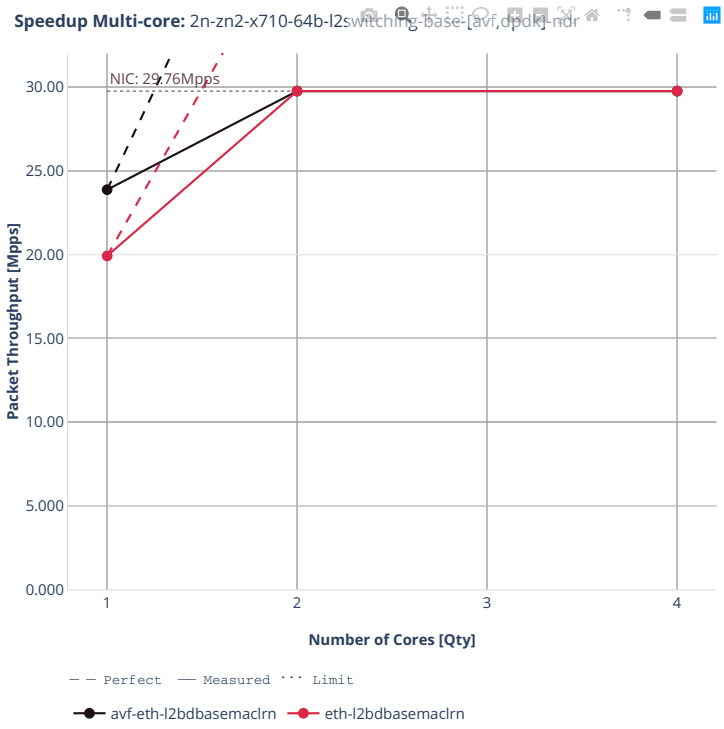

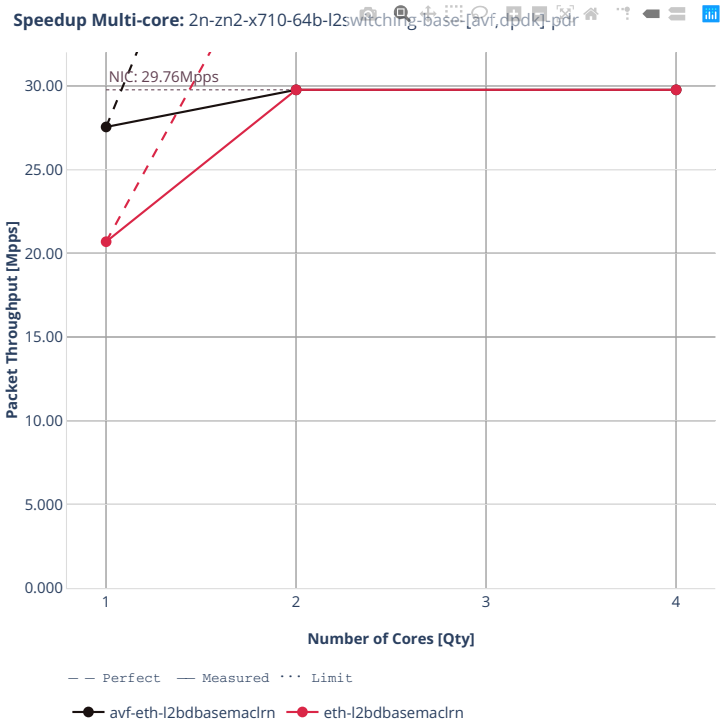
Speedup Multi-core: 2n-zn2-x710-64b-l2switching-base-[avf,dpdk]-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2bdbasemaclrn
eth-l2bdbasemaclrn

2n-zn2-cx556a

64b-l2switching-base-rdma-core

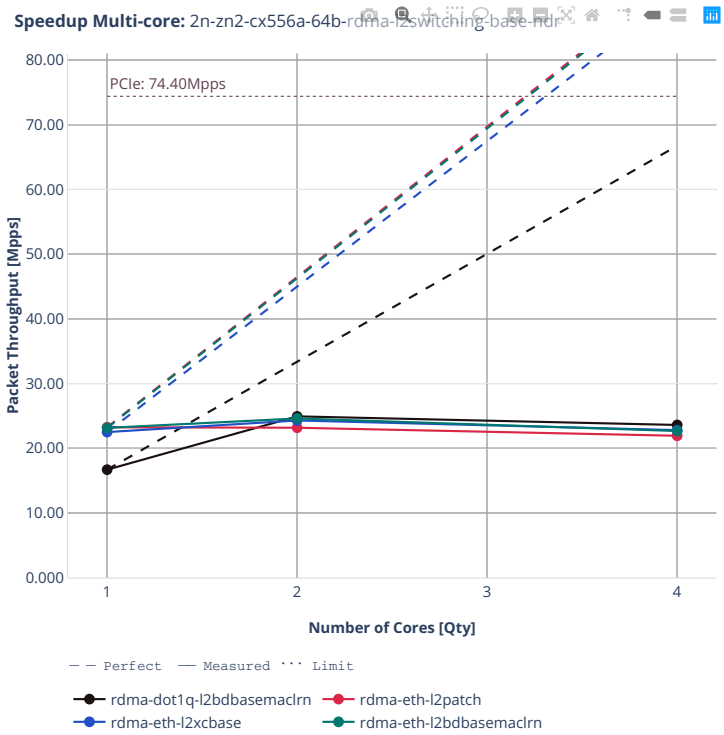

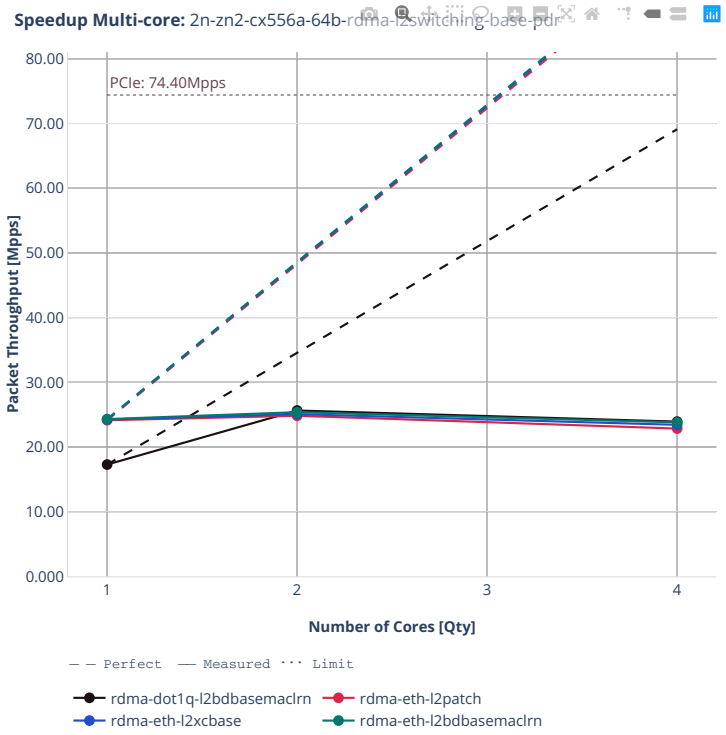
Speedup Multi-core: 2n-zn2-cx556a-64b-rdma-l2switching-base-pdr
PCIe: 74.40Mpps
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-dot1q-l2bdbasemaclrn
rdma-eth-l2patch
rdma-eth-l2xcbase
rdma-eth-l2bdbasemaclrn

64b-l2switching-scale

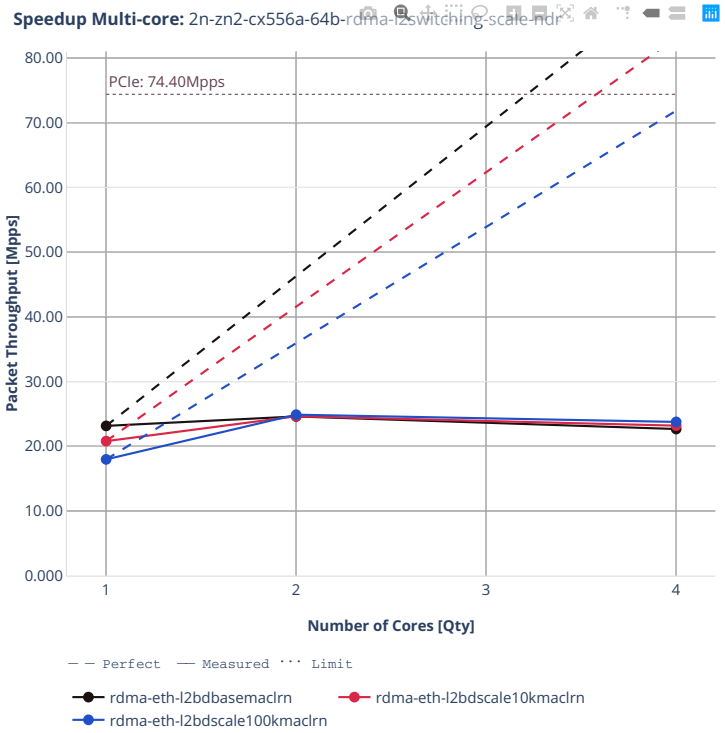

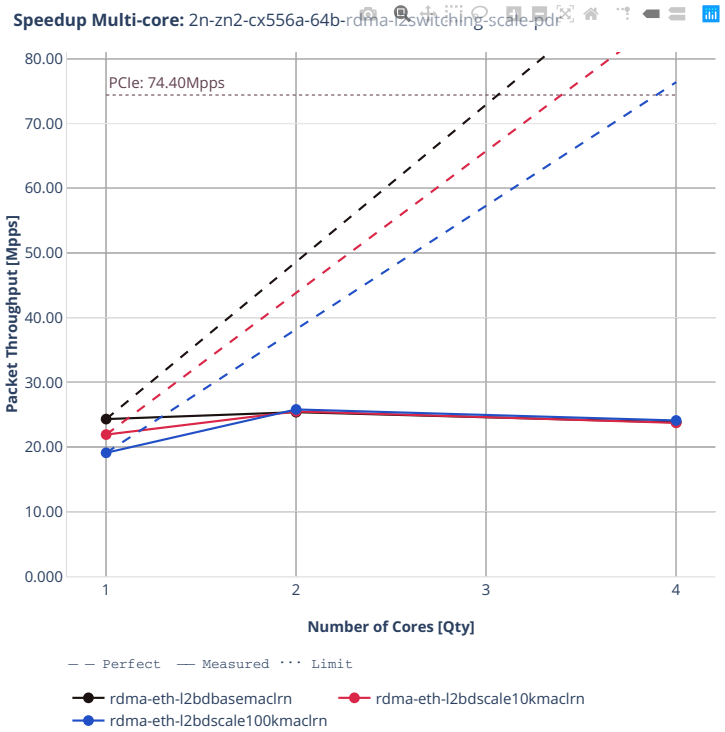
Speedup Multi-core: 2n-zn2-cx556a-64b-rdma-l2switching-scale-pdr
PCIe: 74.40Mpps
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-eth-l2bdbasemaclrn
rdma-eth-l2bdscale10kmaclrn
rdma-eth-l2bdscale100kmaclrn

3n-alt-xl710

64b-l2switching-base-scale

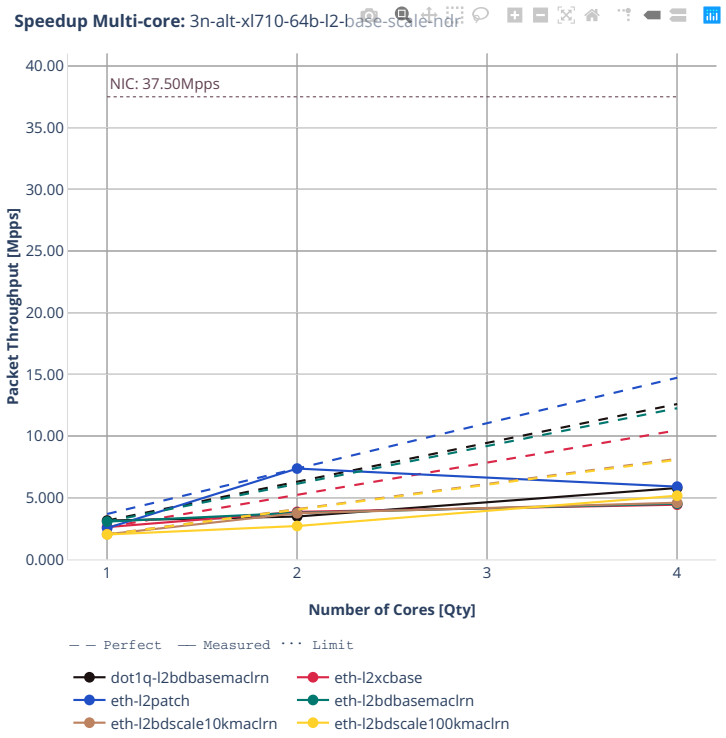

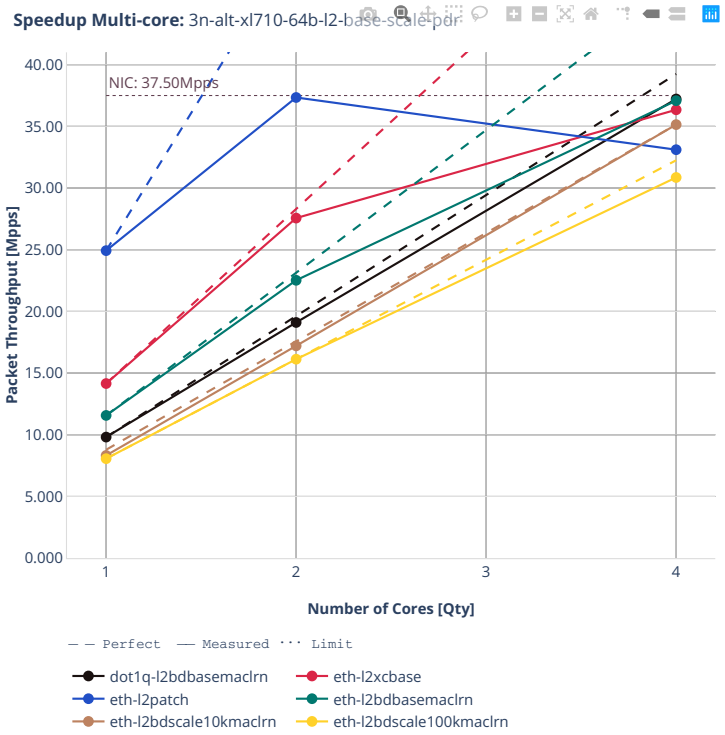
Speedup Multi-core: 3n-alt-xl710-64b-l2-base-scale-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
dot1q-l2bdbasemaclrn
eth-l2xcbase
eth-l2patch
eth-l2bdbasemaclrn
eth-l2bdscale10kmaclrn
eth-l2bdscale100kmaclrn

64b-l2switching-features

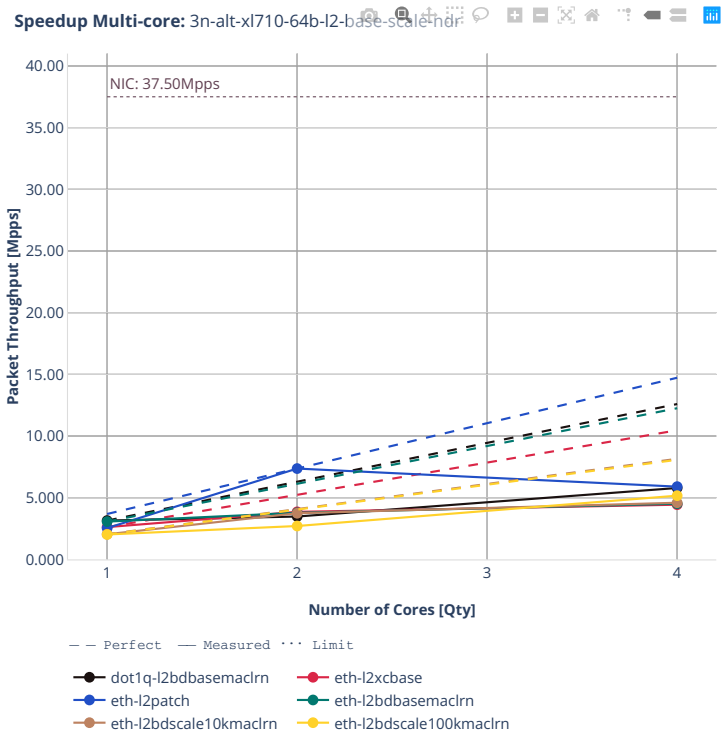

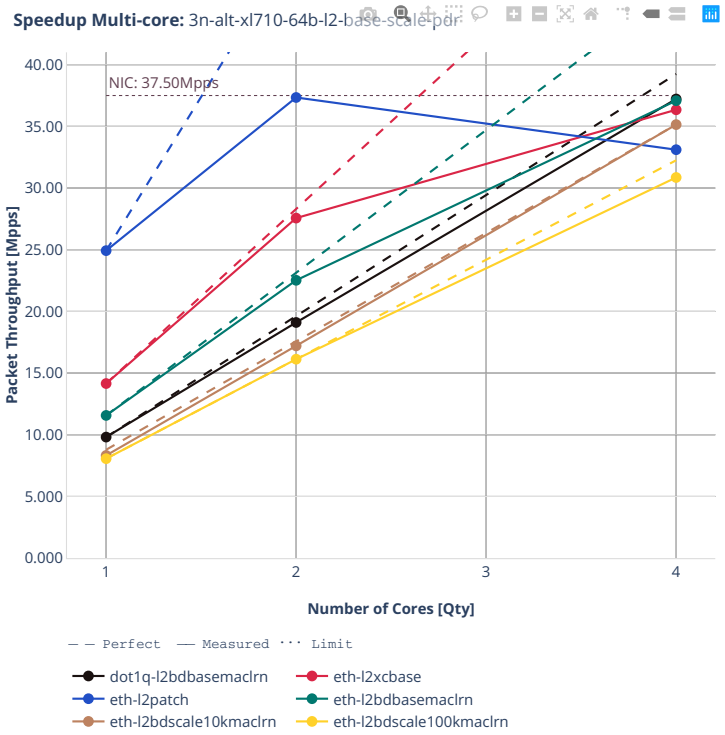
Speedup Multi-core: 3n-alt-xl710-64b-l2-base-scale-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
dot1q-l2bdbasemaclrn
eth-l2xcbase
eth-l2patch
eth-l2bdbasemaclrn
eth-l2bdscale10kmaclrn
eth-l2bdscale100kmaclrn

3n-tsh-x520

64b-l2switching-base-ixgbe

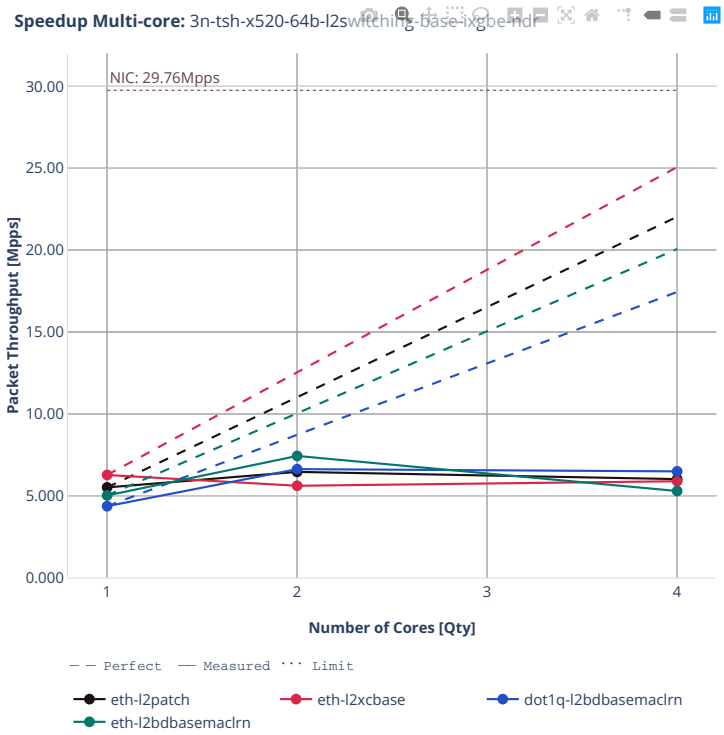

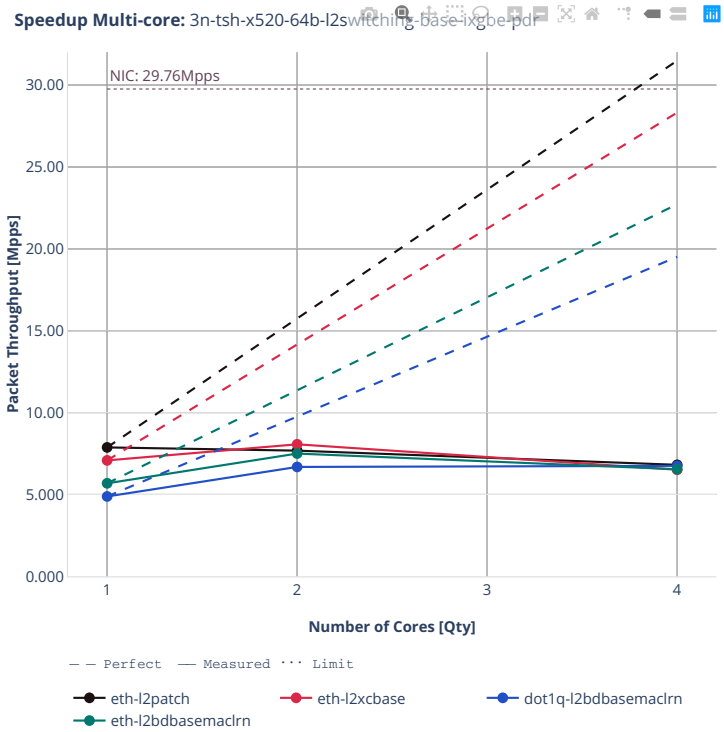
Speedup Multi-core: 3n-tsh-x520-64b-l2switching-base-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2patch
eth-l2xcbase
dot1q-l2bdbasemaclrn
eth-l2bdbasemaclrn

64b-features-l2switching-base-ixgbe

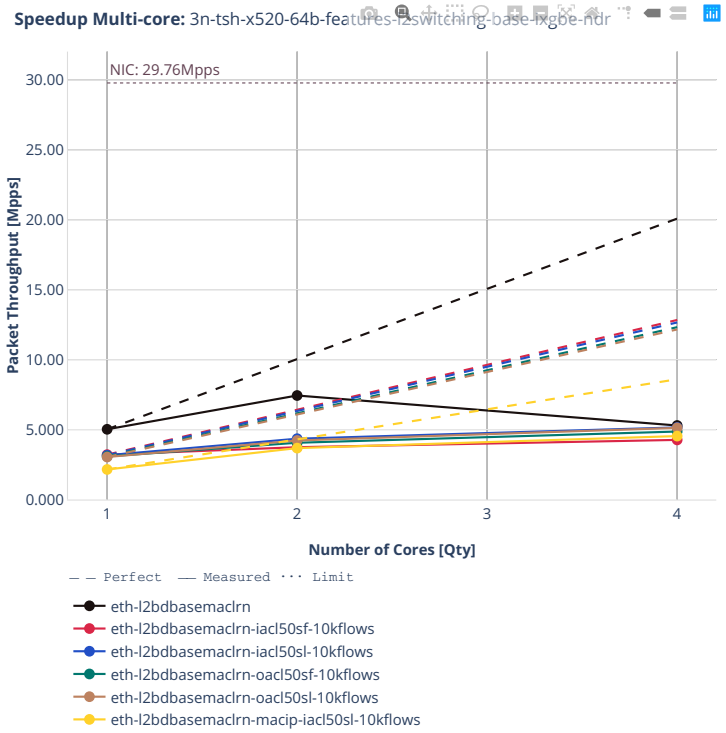

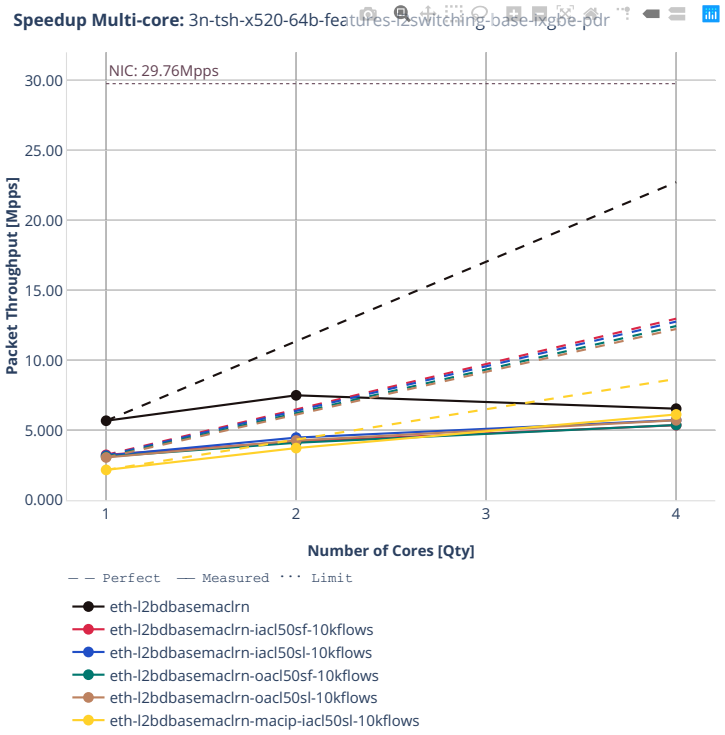
Speedup Multi-core: 3n-tsh-x520-64b-features-l2switching-base-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2bdbasemaclrn
eth-l2bdbasemaclrn-iacl50sf-10kflows
eth-l2bdbasemaclrn-iacl50sl-10kflows
eth-l2bdbasemaclrn-oacl50sf-10kflows
eth-l2bdbasemaclrn-oacl50sl-10kflows
eth-l2bdbasemaclrn-macip-iacl50sl-10kflows

2n-tx2-xl710

64b-l2switching-base-dpdk


Speedup Multi-core: 2n-tx2-xl710-64b-l2switching-base-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
dot1q-l2bdbasemaclrn
eth-l2patch
eth-l2xcbase
eth-l2bdbasemaclrn

64b-l2switching-scale-dpdk

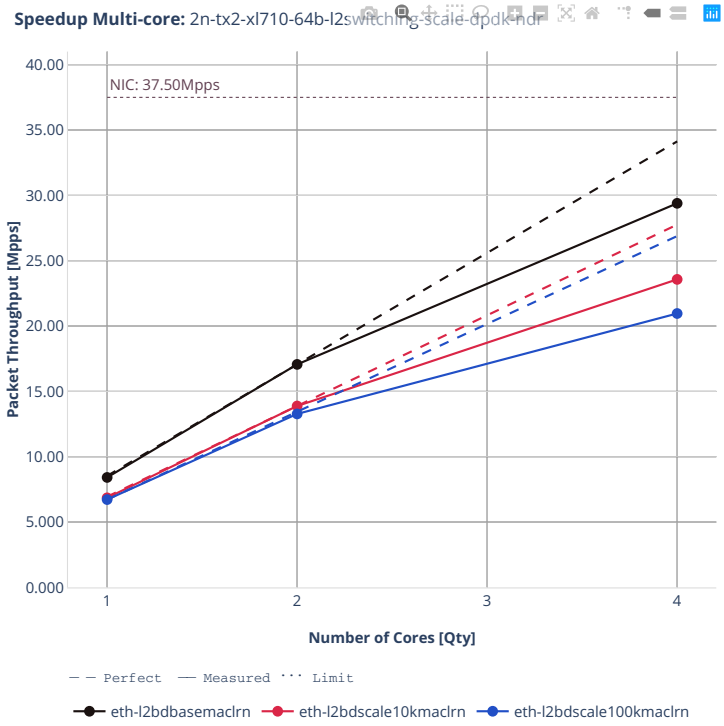

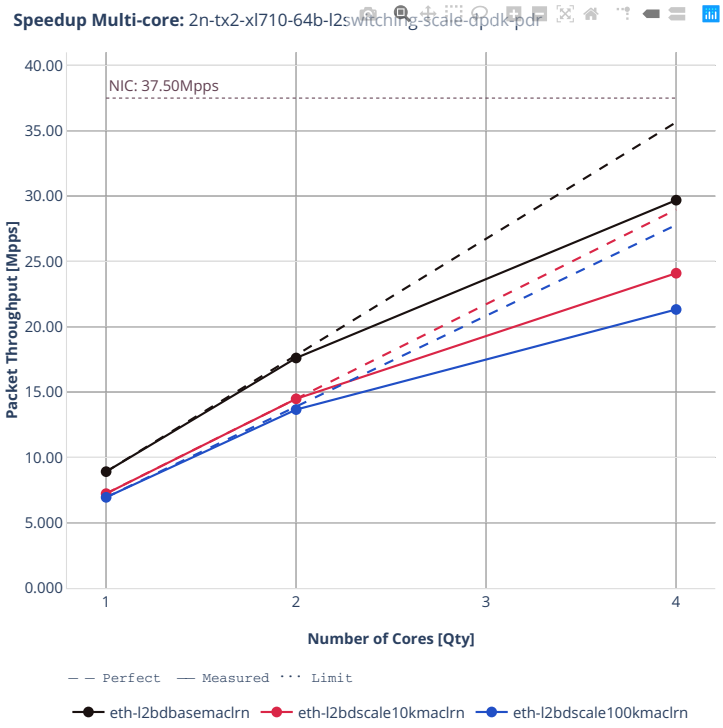
Speedup Multi-core: 2n-tx2-xl710-64b-l2switching-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2bdbasemaclrn
eth-l2bdscale10kmaclrn
eth-l2bdscale100kmaclrn

64b-l2switching-scale-af-xdp

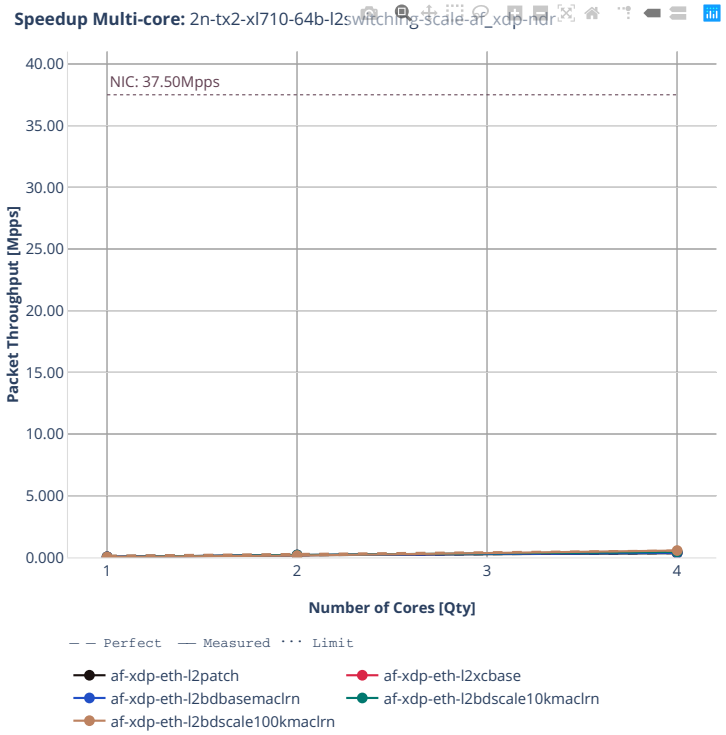

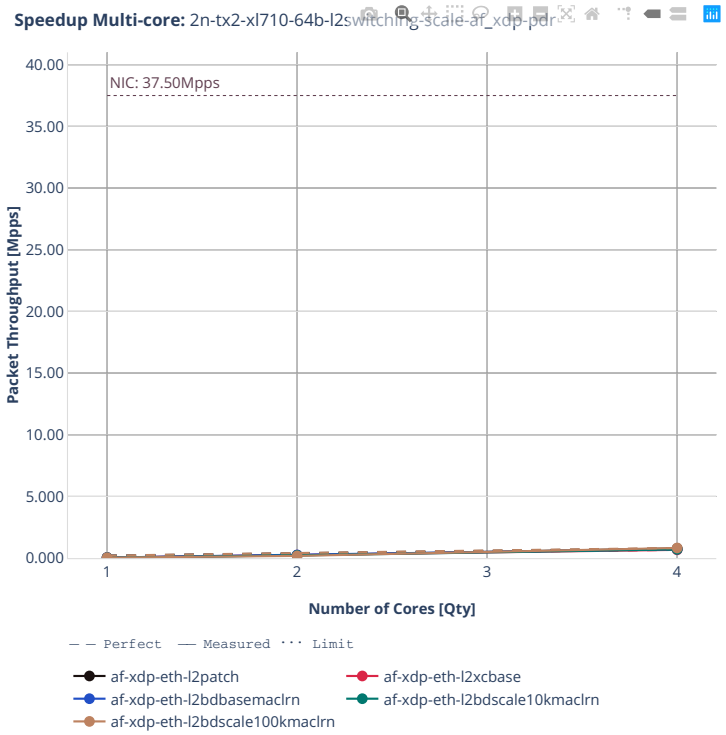
Speedup Multi-core: 2n-tx2-xl710-64b-l2switching-scale-af_xdp-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
af-xdp-eth-l2patch
af-xdp-eth-l2xcbase
af-xdp-eth-l2bdbasemaclrn
af-xdp-eth-l2bdscale10kmaclrn
af-xdp-eth-l2bdscale100kmaclrn

64b-features-l2switching-base-dpdk

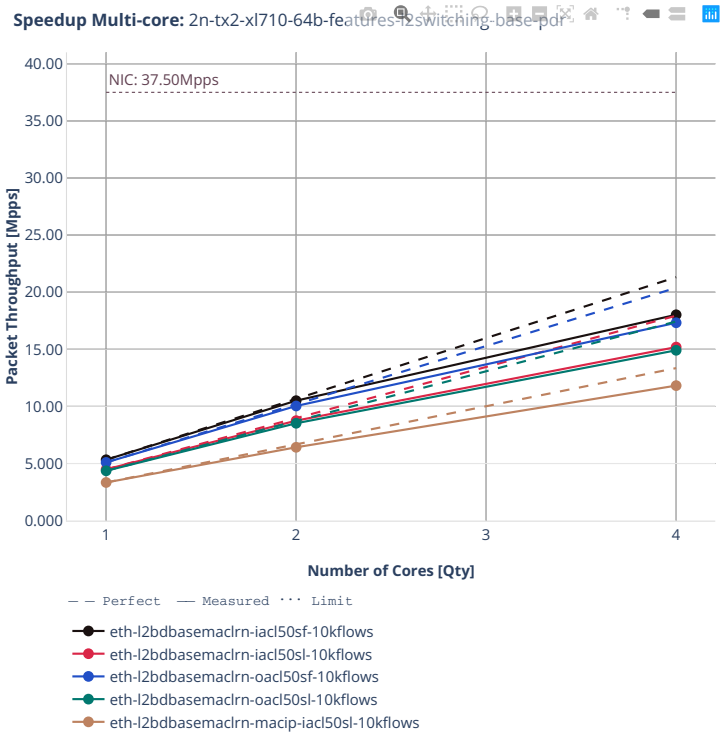
Speedup Multi-core: 2n-tx2-xl710-64b-features-l2switching-base-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2bdbasemaclrn-iacl50sf-10kflows
eth-l2bdbasemaclrn-iacl50sl-10kflows
eth-l2bdbasemaclrn-oacl50sf-10kflows
eth-l2bdbasemaclrn-oacl50sl-10kflows
eth-l2bdbasemaclrn-macip-iacl50sl-10kflows

2n-dnv-x553

64b-l2switching-base-scale-ixgbe

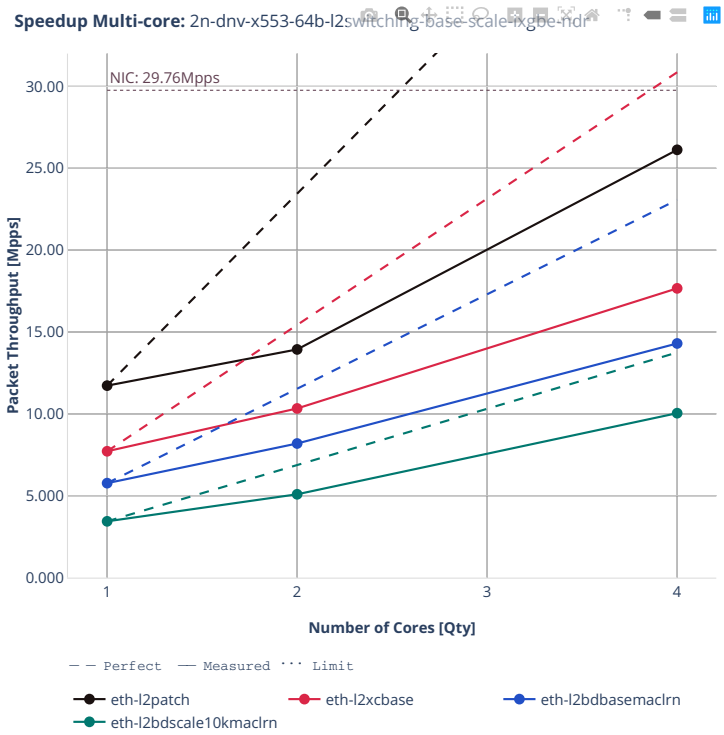

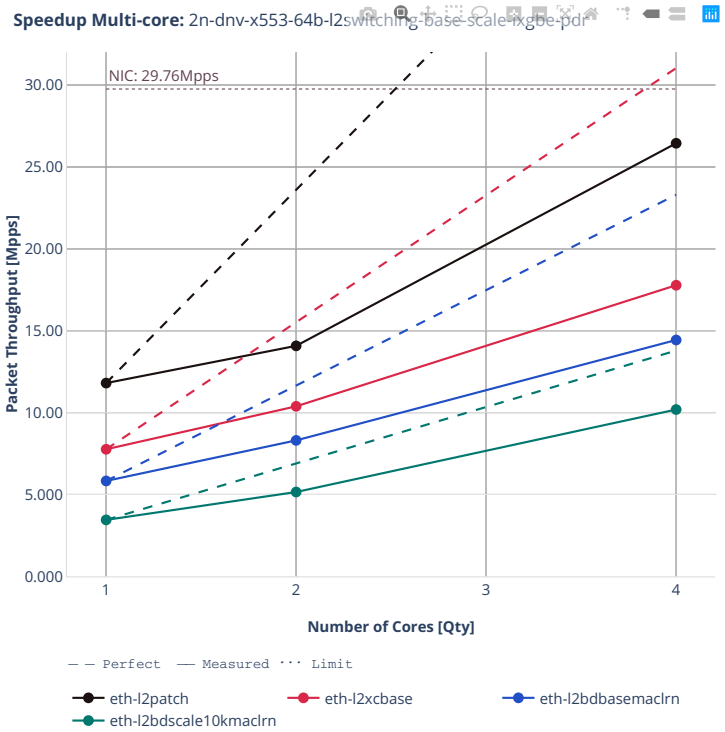
Speedup Multi-core: 2n-dnv-x553-64b-l2switching-base-scale-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2patch
eth-l2xcbase
eth-l2bdbasemaclrn
eth-l2bdscale10kmaclrn

3n-dnv-x553

64b-l2switching-base-scale-ixgbe

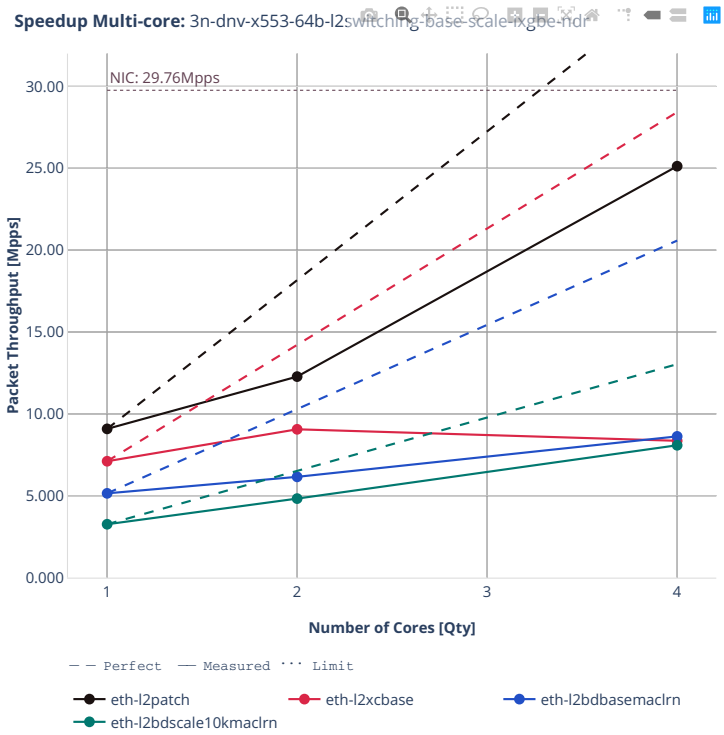

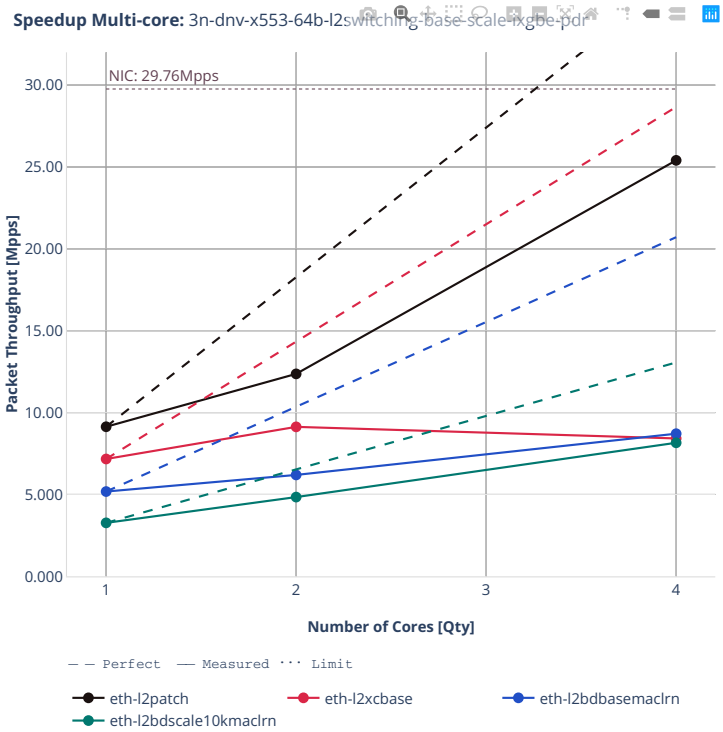
Speedup Multi-core: 3n-dnv-x553-64b-l2switching-base-scale-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2patch
eth-l2xcbase
eth-l2bdbasemaclrn
eth-l2bdscale10kmaclrn

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[131].

[131] https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls2206

2n-icx-xxv710

64b-ip4routing-base-scale-avf

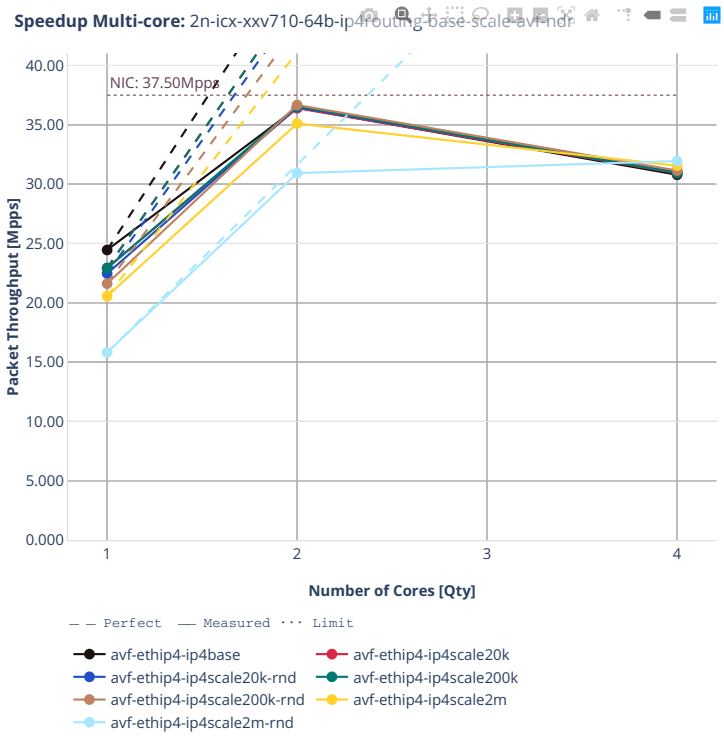

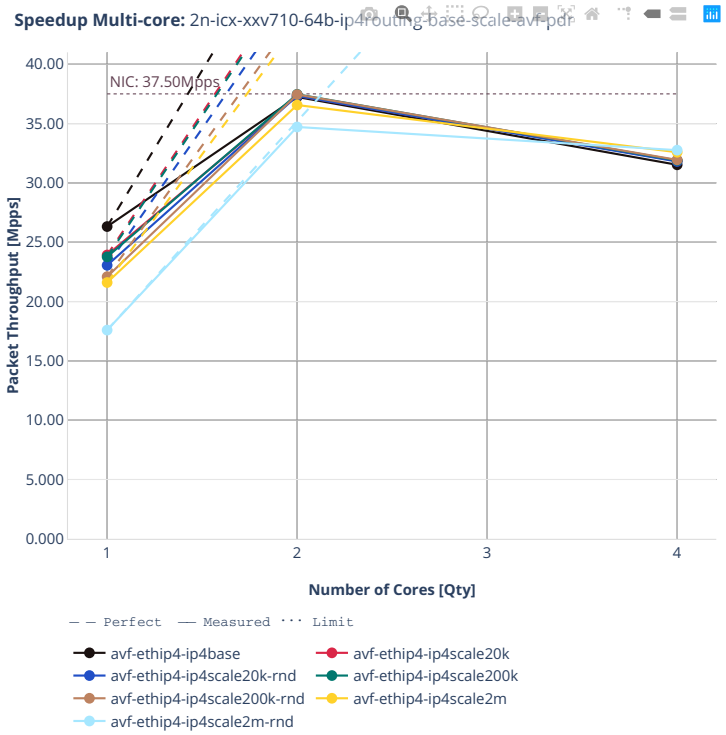
Speedup Multi-core: 2n-icx-xxv710-64b-ip4routing-base-scale-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4-ip4base
avf-ethip4-ip4scale20k
avf-ethip4-ip4scale20k-rnd
avf-ethip4-ip4scale200k
avf-ethip4-ip4scale200k-rnd
avf-ethip4-ip4scale2m
avf-ethip4-ip4scale2m-rnd

64b-ip4routing-base-scale-dpdk

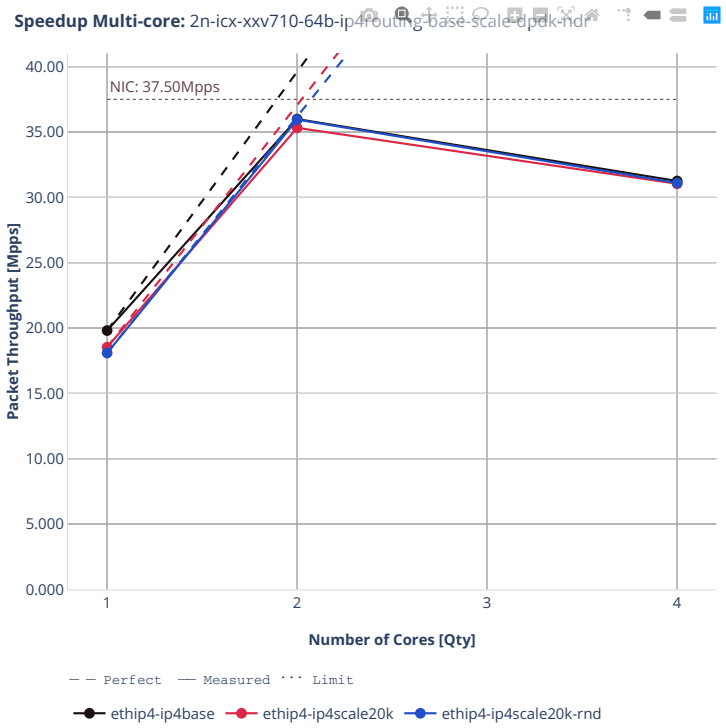

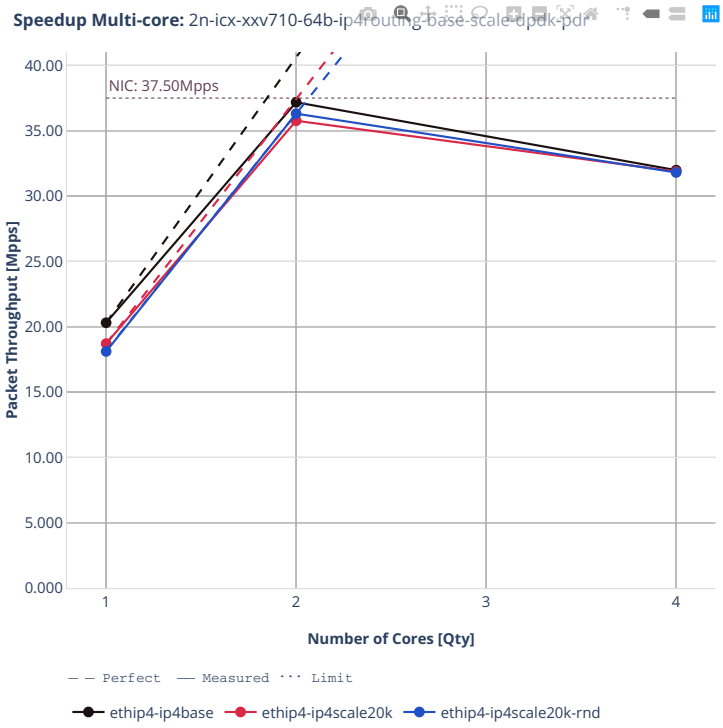
Speedup Multi-core: 2n-icx-xxv710-64b-ip4routing-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base
ethip4-ip4scale20k
ethip4-ip4scale20k-rnd

64b-ip4routing-base-scale-af_xdp

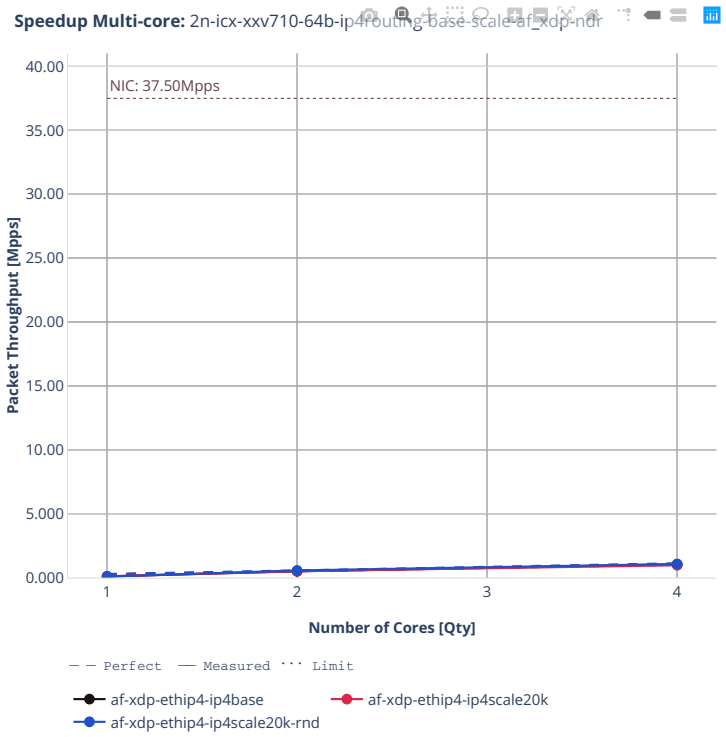

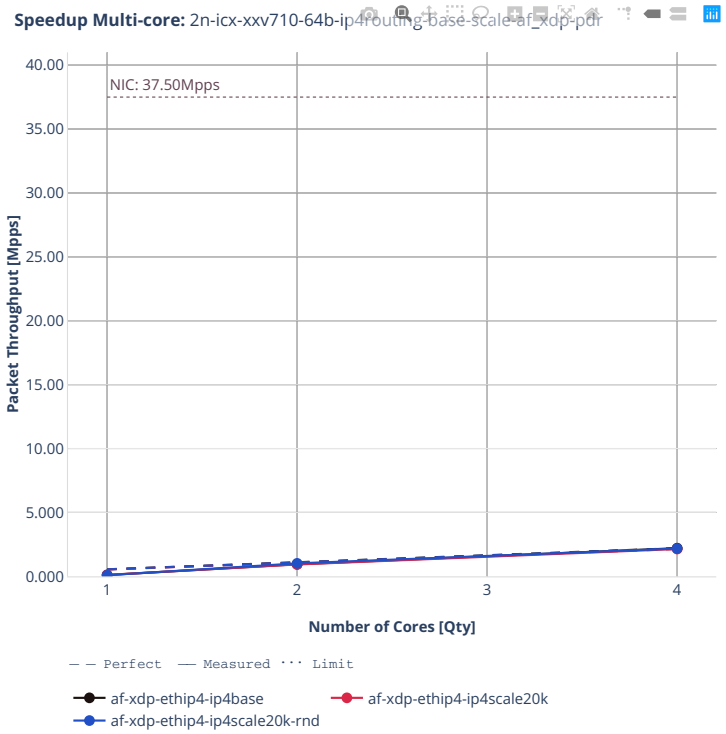
Speedup Multi-core: 2n-icx-xxv710-64b-ip4routing-base-scale-af_xdp-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
af-xdp-ethip4-ip4base
af-xdp-ethip4-ip4scale20k
af-xdp-ethip4-ip4scale20k-rnd

64b-features-ip4routing-base-avf

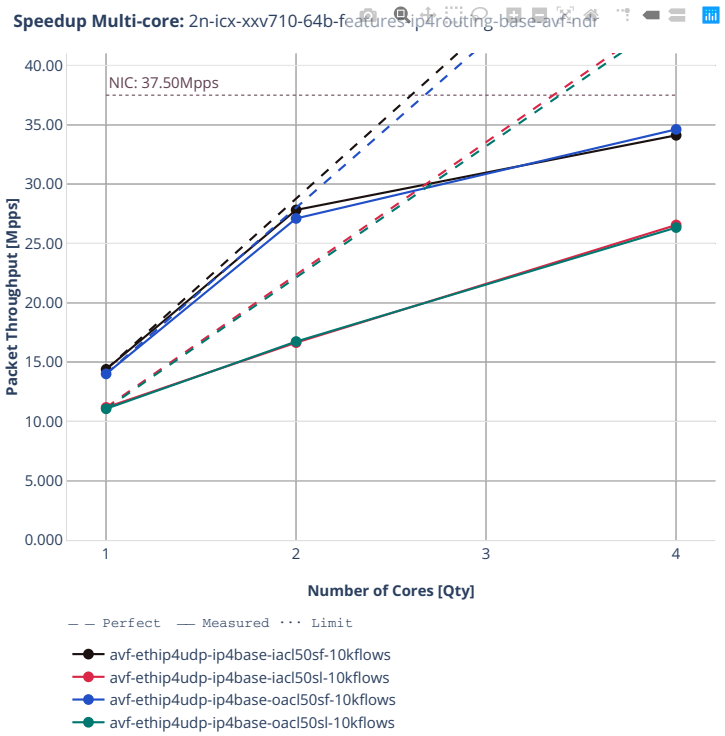

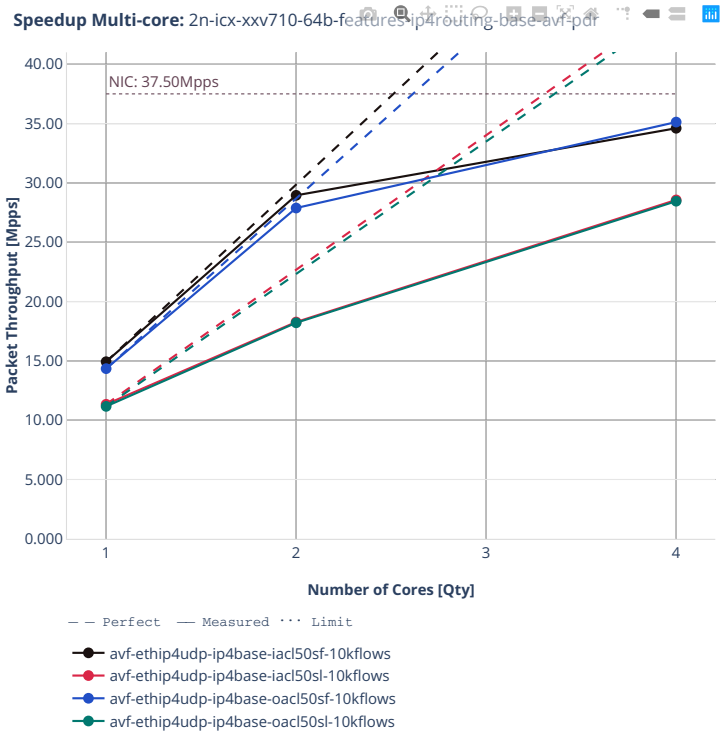
Speedup Multi-core: 2n-icx-xxv710-64b-features-ip4routing-base-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-ip4base-iacl50sf-10kflows
avf-ethip4udp-ip4base-iacl50sl-10kflows
avf-ethip4udp-ip4base-oacl50sf-10kflows
avf-ethip4udp-ip4base-oacl50sl-10kflows

3n-icx-xxv710

64b-ip4routing-base

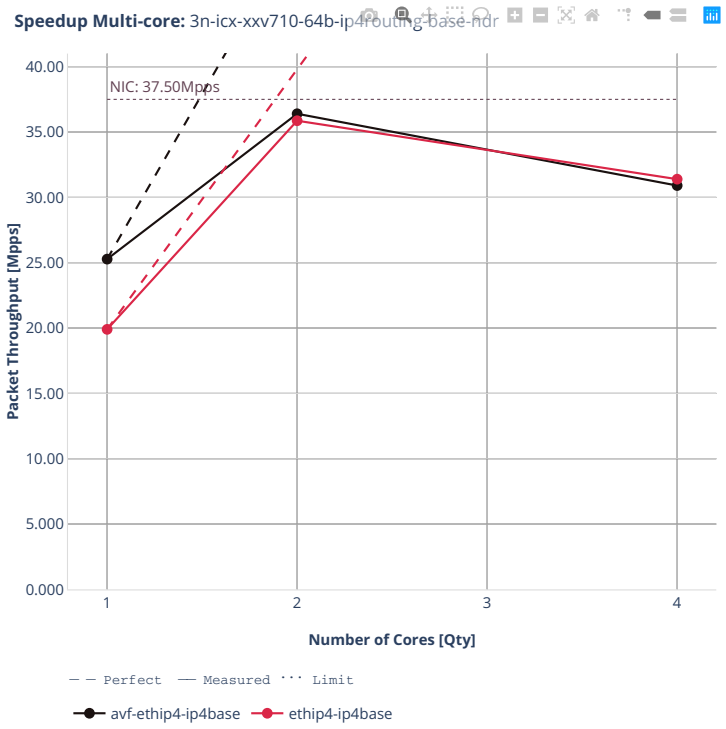

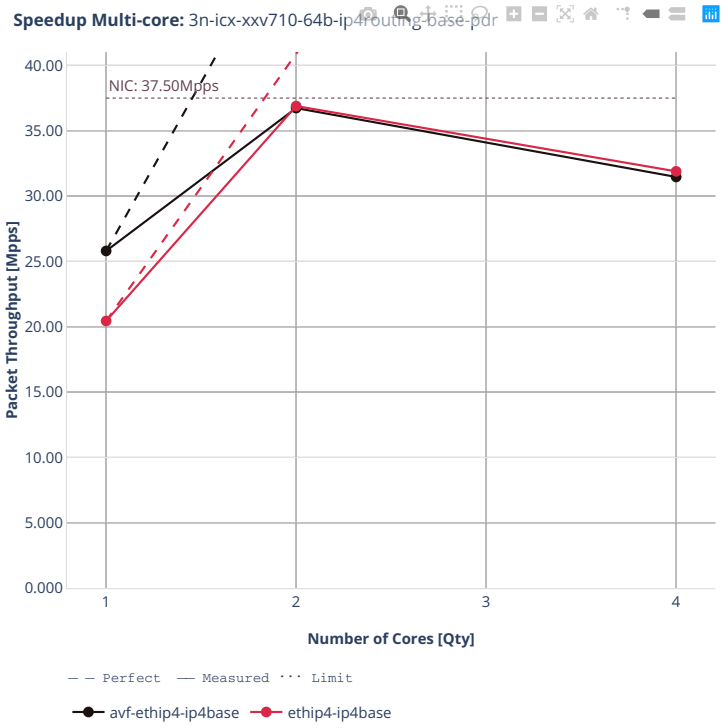
Speedup Multi-core: 3n-icx-xxv710-64b-ip4routing-base-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4-ip4base
ethip4-ip4base

2n-skx-xxv710

64b-ip4routing-base-scale-avf

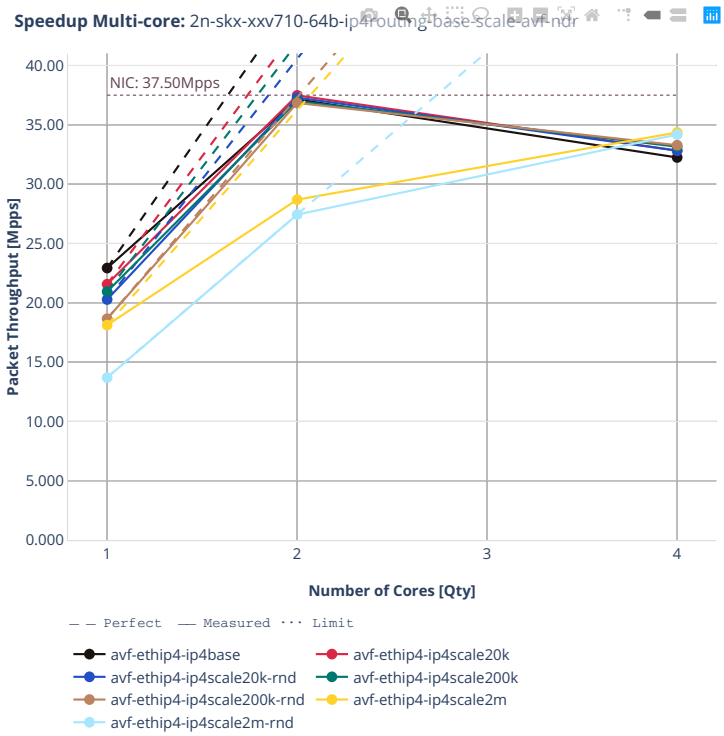

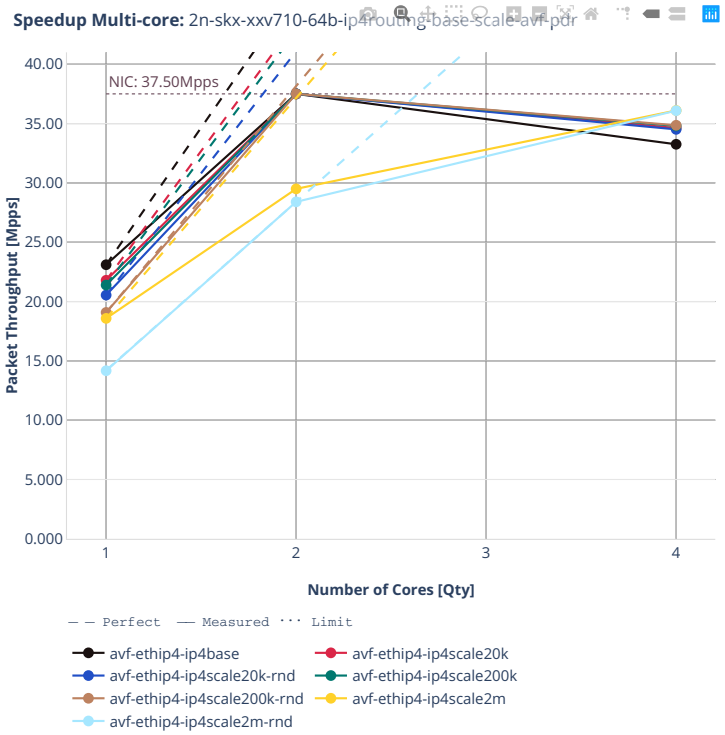
Speedup Multi-core: 2n-skx-xxv710-64b-ip4routing-base-scale-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4-ip4base
avf-ethip4-ip4scale20k
avf-ethip4-ip4scale20k-rnd
avf-ethip4-ip4scale200k
avf-ethip4-ip4scale200k-rnd
avf-ethip4-ip4scale2m
avf-ethip4-ip4scale2m-rnd

64b-ip4routing-base-scale-af-xdp

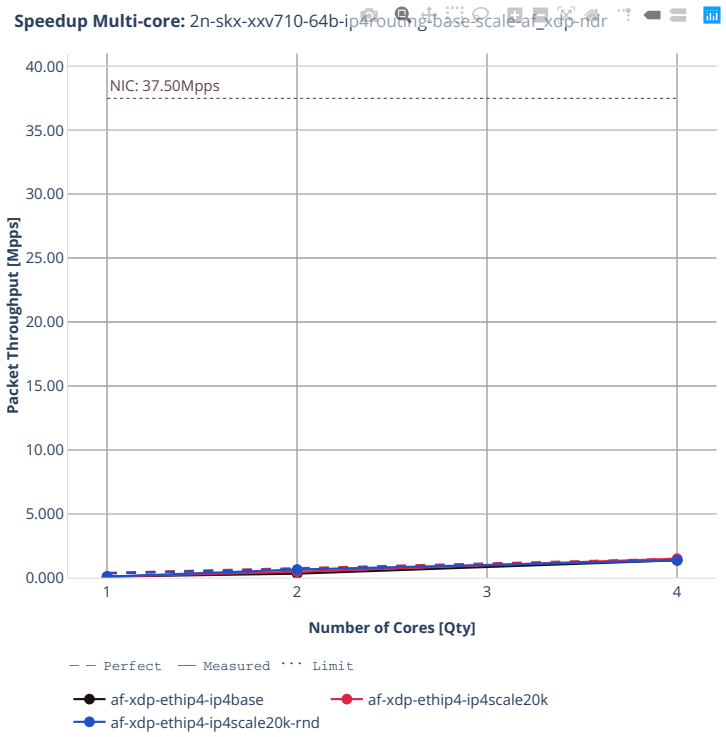

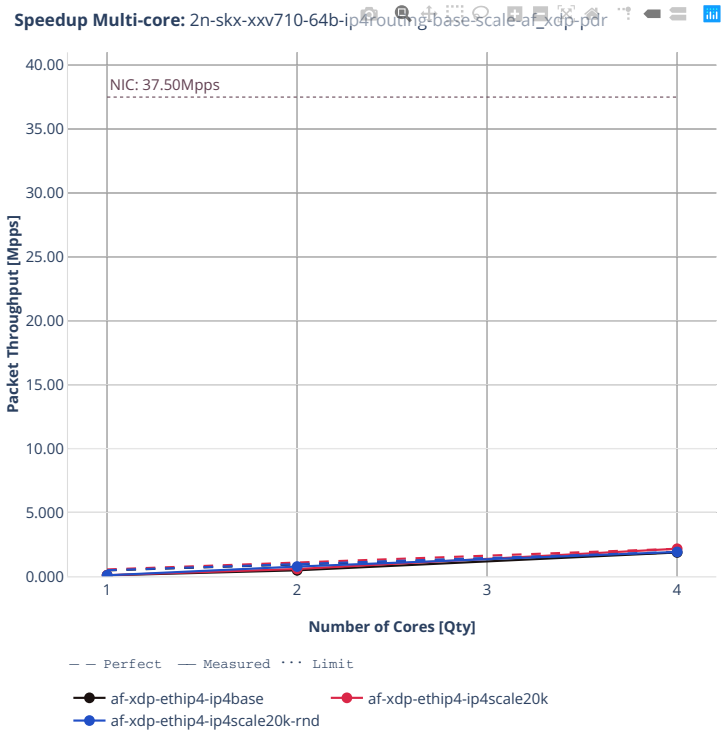
Speedup Multi-core: 2n-skx-xxv710-64b-ip4routing-base-scale-af_xdp-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
af-xdp-ethip4-ip4base
af-xdp-ethip4-ip4scale20k
af-xdp-ethip4-ip4scale20k-rnd

64b-ip4routing-base-scale-dpdk

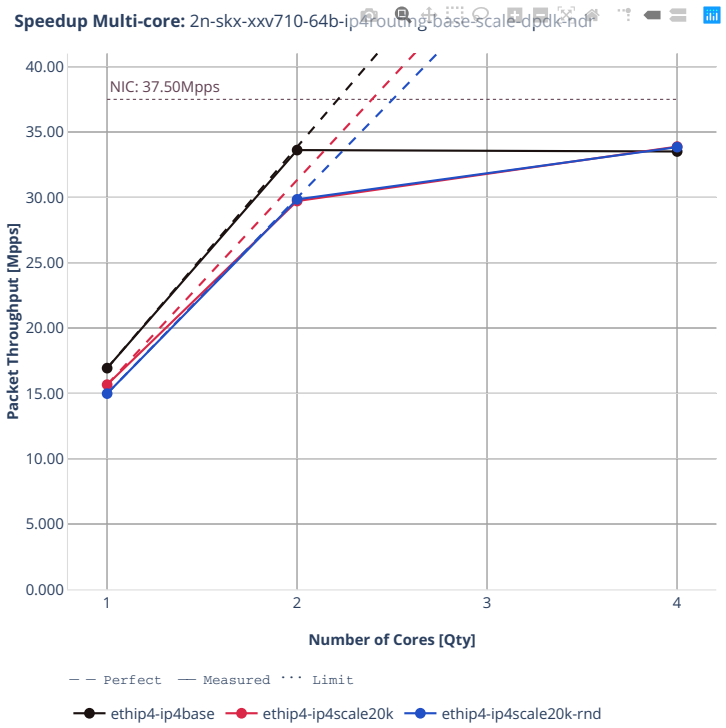

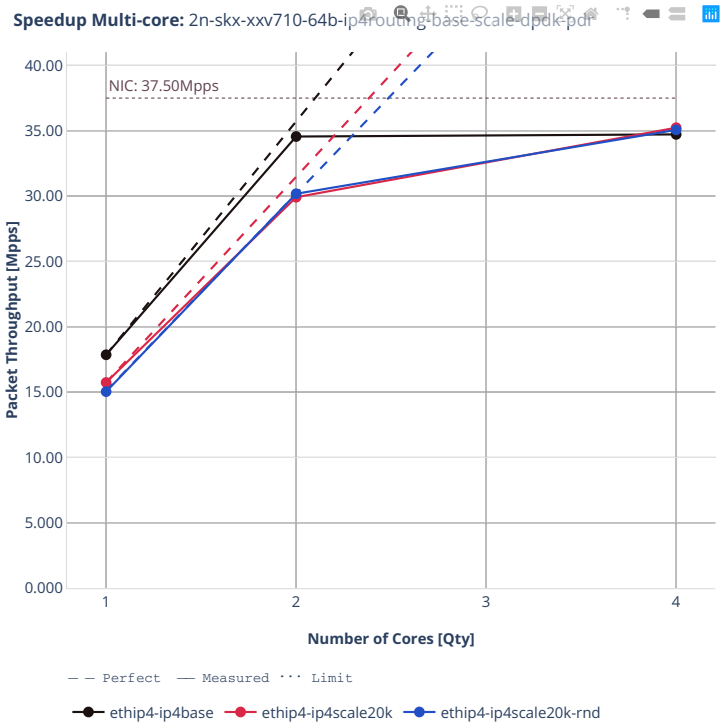
Speedup Multi-core: 2n-skx-xxv710-64b-ip4routing-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base
ethip4-ip4scale20k
ethip4-ip4scale20k-rnd

64b-features-ip4routing-base-avf

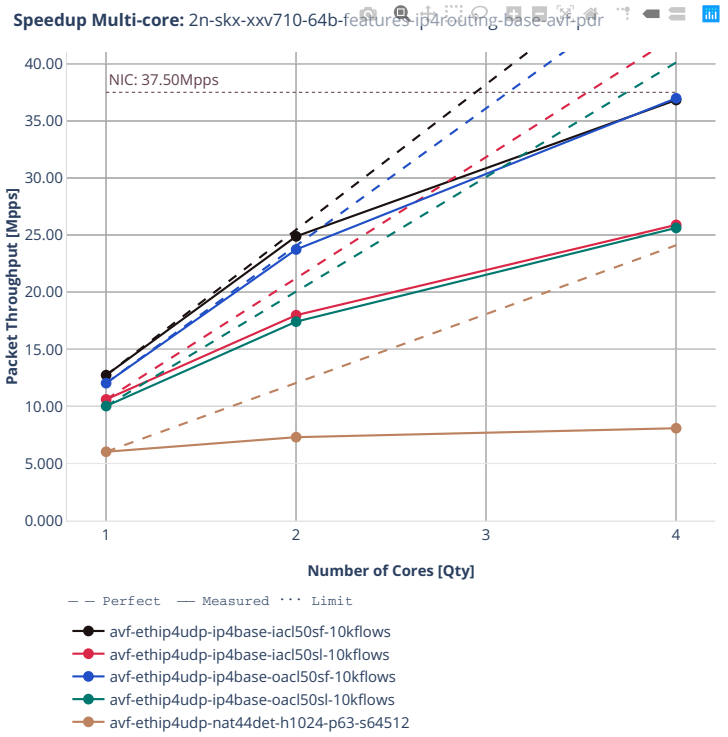
Speedup Multi-core: 2n-skx-xxv710-64b-features-ip4routing-base-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-ip4base-iacl50sf-10kflows
avf-ethip4udp-ip4base-iacl50sl-10kflows
avf-ethip4udp-ip4base-oacl50sf-10kflows
avf-ethip4udp-ip4base-oacl50sl-10kflows
avf-ethip4udp-nat44det-h1024-p63-s64512

2n-skx-x710

64b-ip4routing-base-scale-[avf,dpdk]

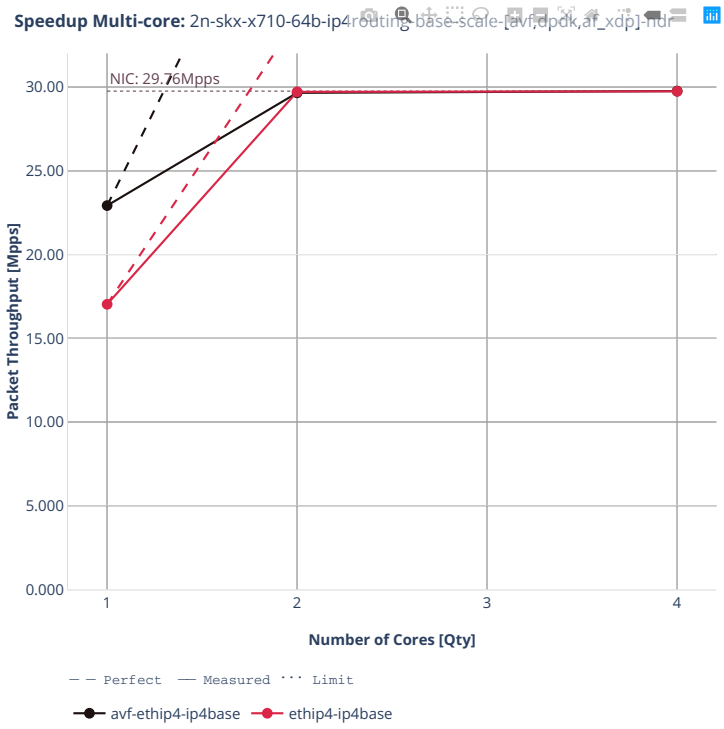

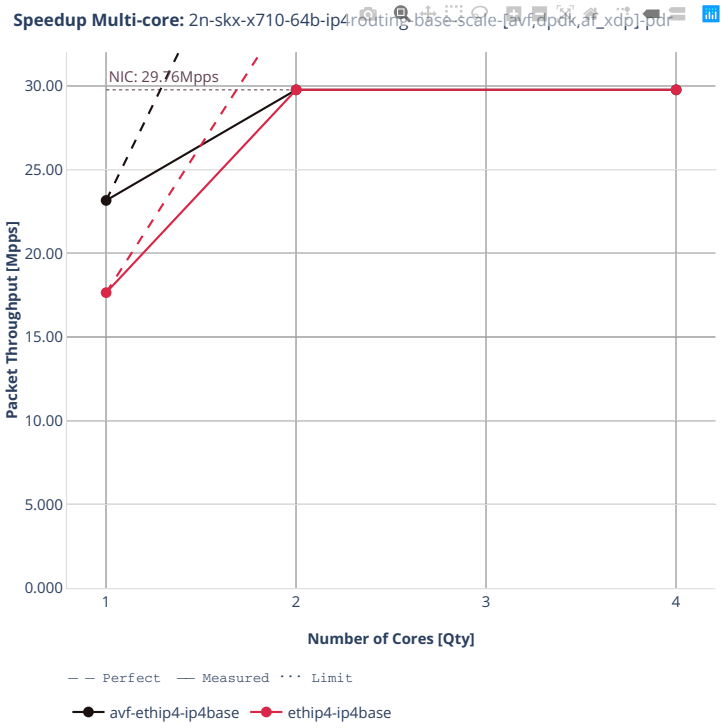
Speedup Multi-core: 2n-skx-x710-64b-ip4routing-base-scale-[avf,dpdk,af_xdp]-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4-ip4base
ethip4-ip4base

3n-skx-xxv710

64b-ip4routing-base-[avf,dpdk]

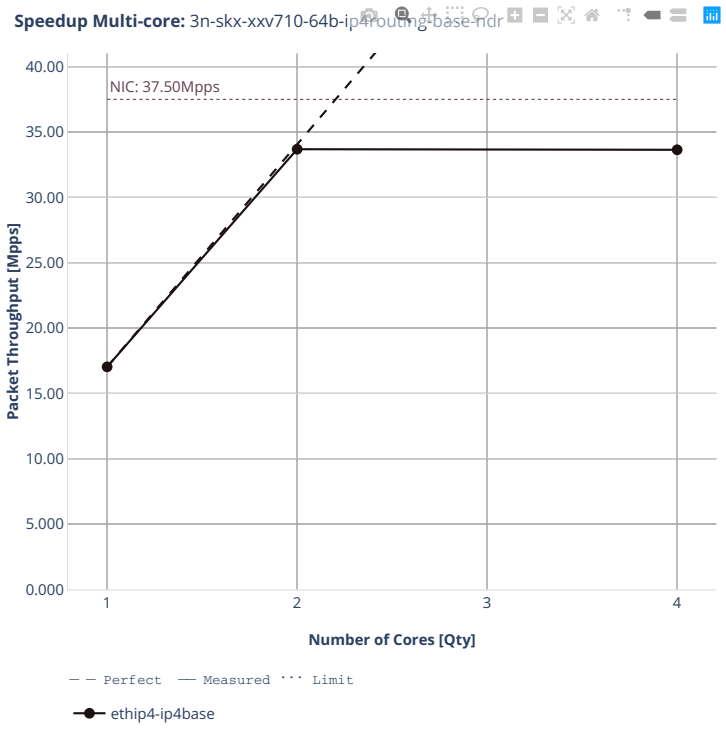

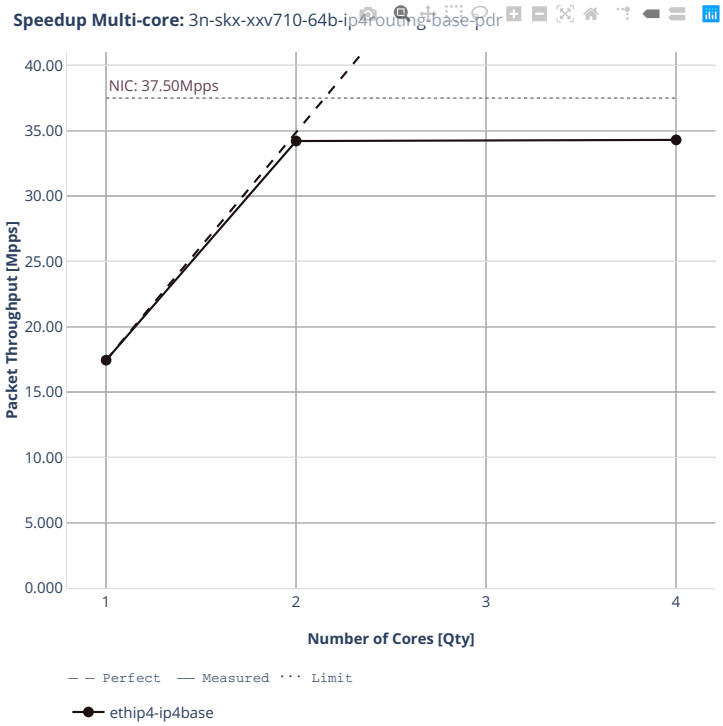
Speedup Multi-core: 3n-skx-xxv710-64b-ip4routing-base-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base

3n-skx-x710

64b-ip4routing-base-avf

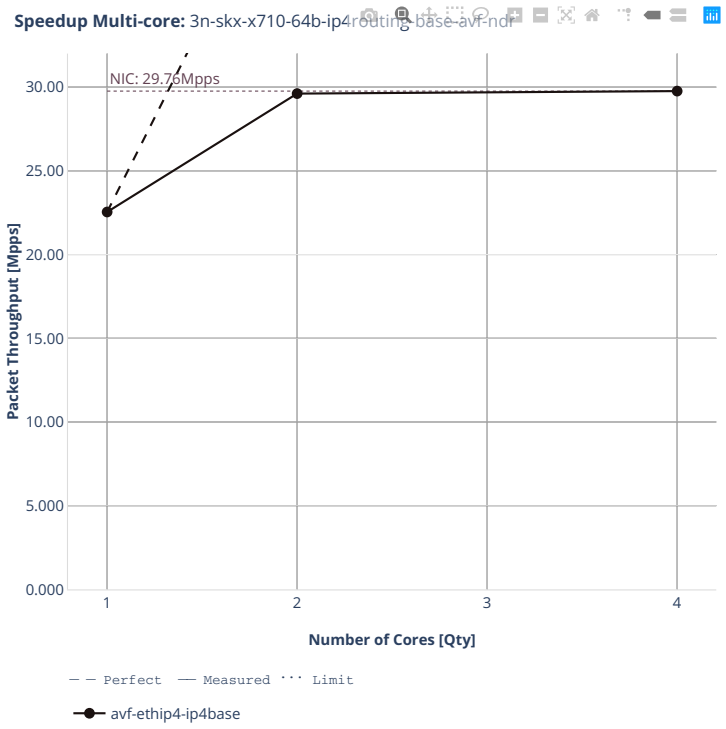

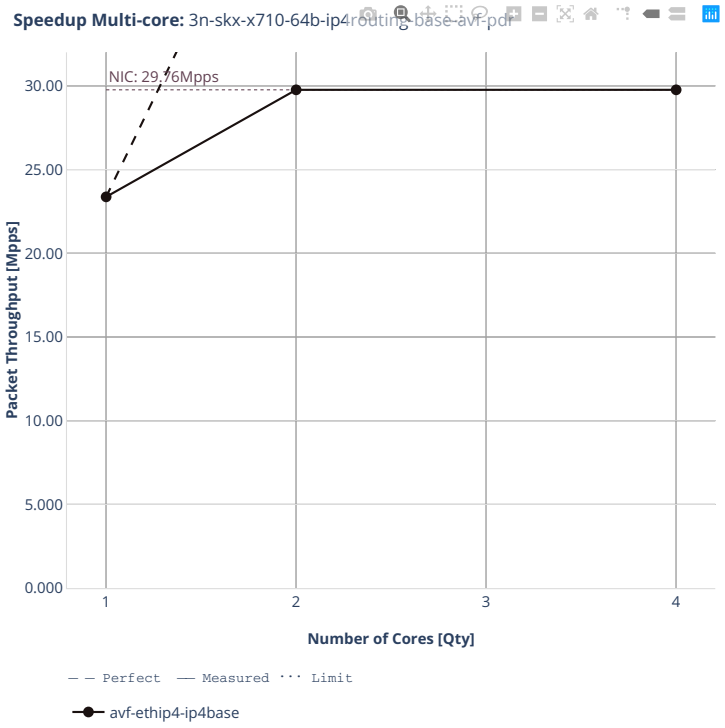
Speedup Multi-core: 3n-skx-x710-64b-ip4routing-base-avf-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4-ip4base

2n-clx-xxv710

64b-ip4routing-base-scale-avf

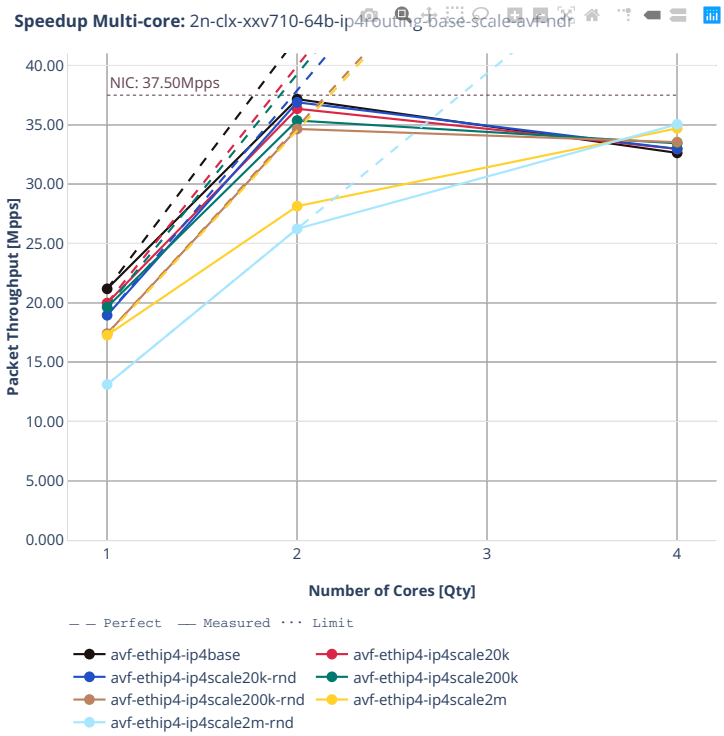

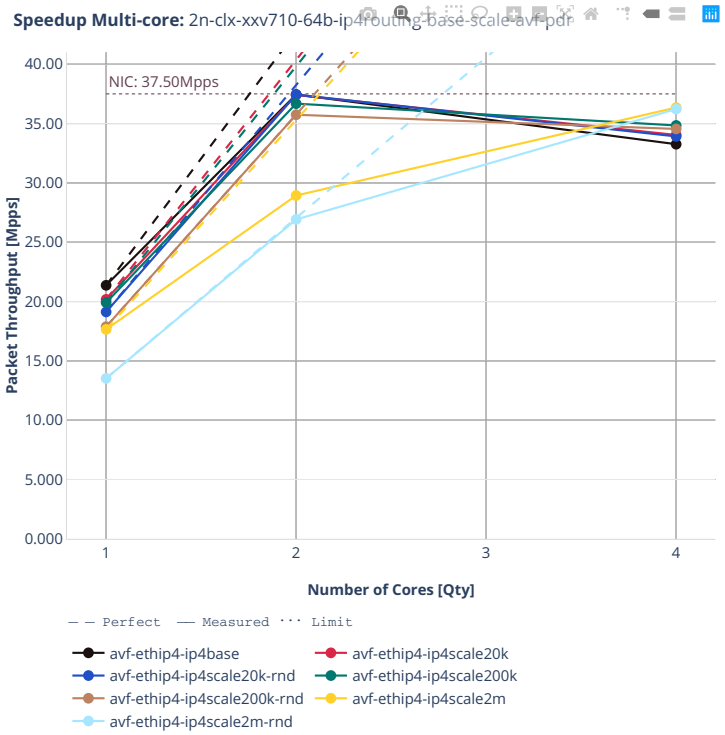
Speedup Multi-core: 2n-clx-xxv710-64b-ip4routing-base-scale-avf-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4-ip4base
avf-ethip4-ip4scale20k
avf-ethip4-ip4scale20k-rnd
avf-ethip4-ip4scale200k
avf-ethip4-ip4scale200k-rnd
avf-ethip4-ip4scale2m
avf-ethip4-ip4scale2m-rnd

64b-ip4routing-base-scale-af-xdp

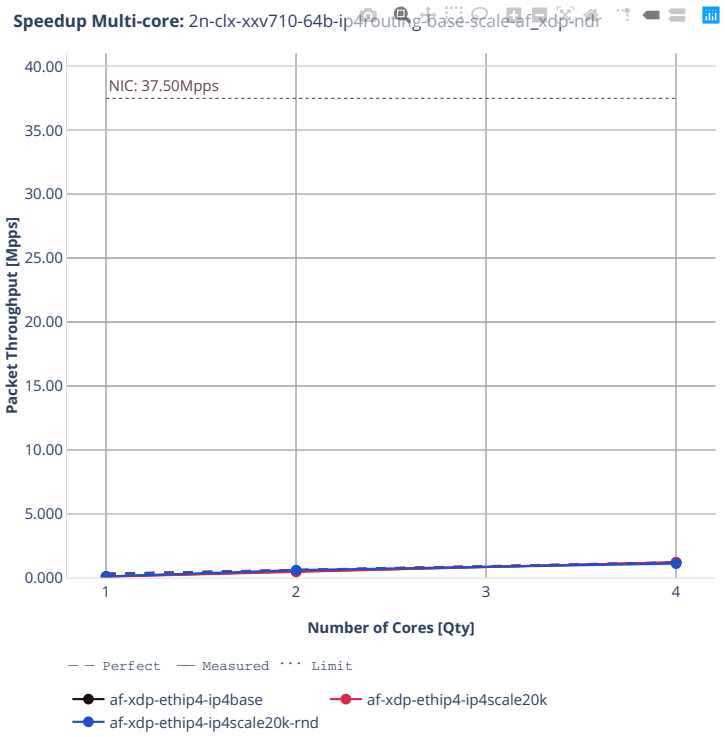

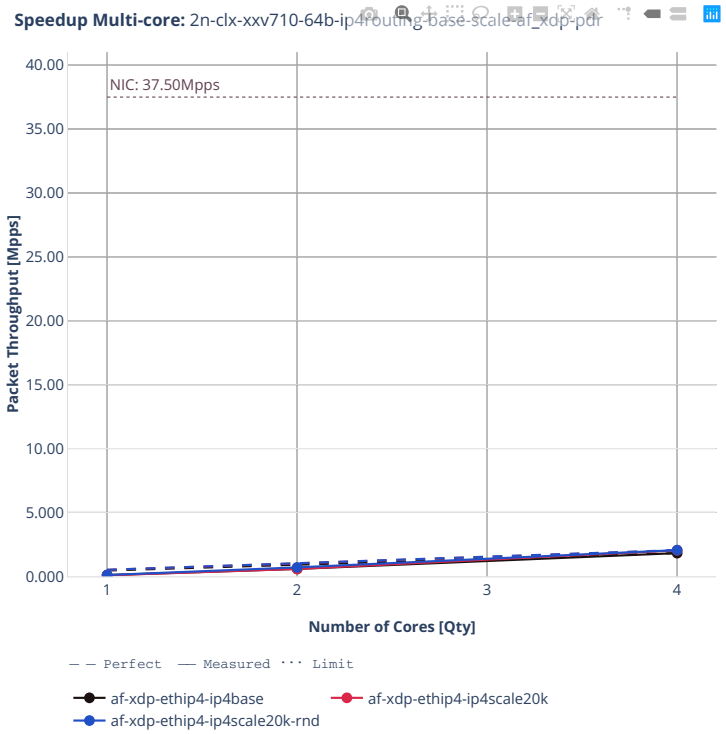
Speedup Multi-core: 2n-clx-xxv710-64b-ip4routing-base-scale-af_xdp-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
af-xdp-ethip4-ip4base
af-xdp-ethip4-ip4scale20k
af-xdp-ethip4-ip4scale20k-rnd

64b-ip4routing-base-scale-dpdk

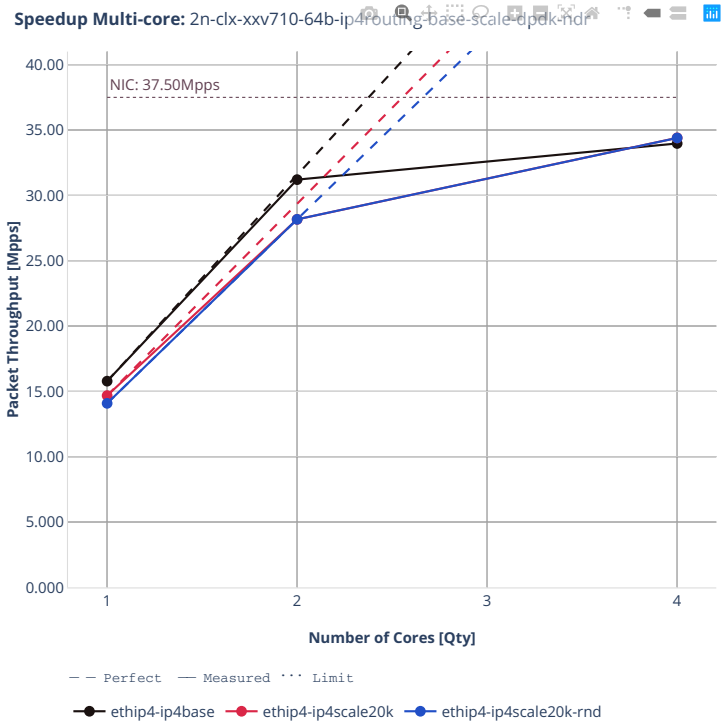

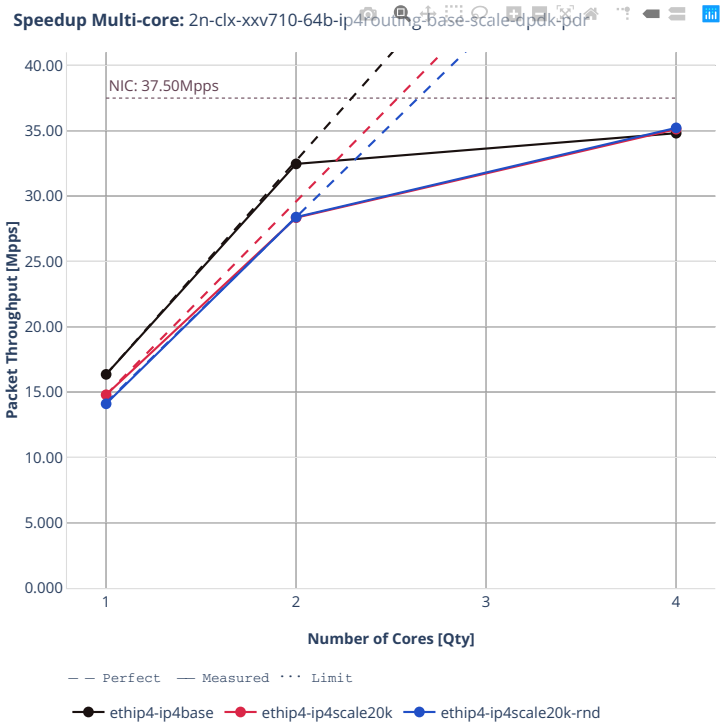
Speedup Multi-core: 2n-clx-xxv710-64b-ip4routing-base-scale-dpdk-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base
ethip4-ip4scale20k
ethip4-ip4scale20k-rnd

64b-features-ip4routing-base-avf

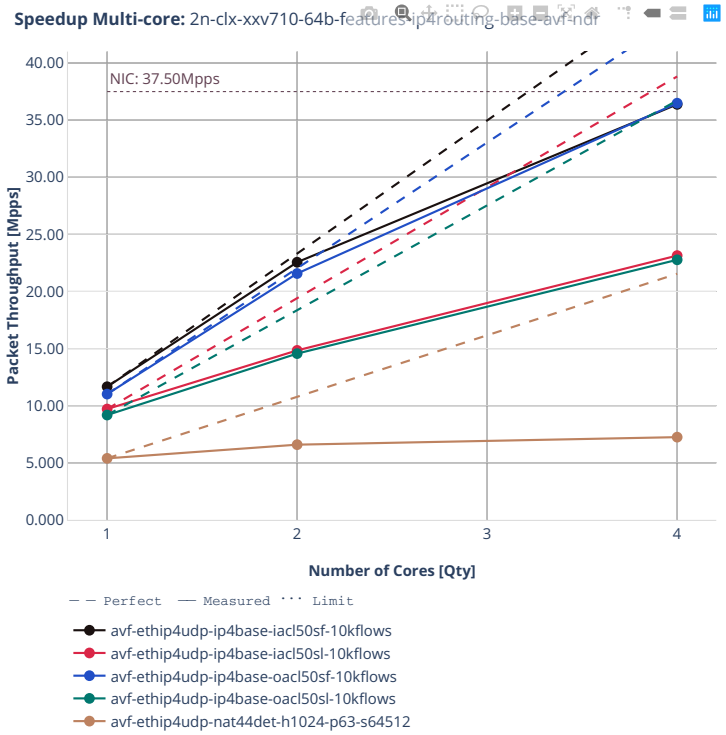

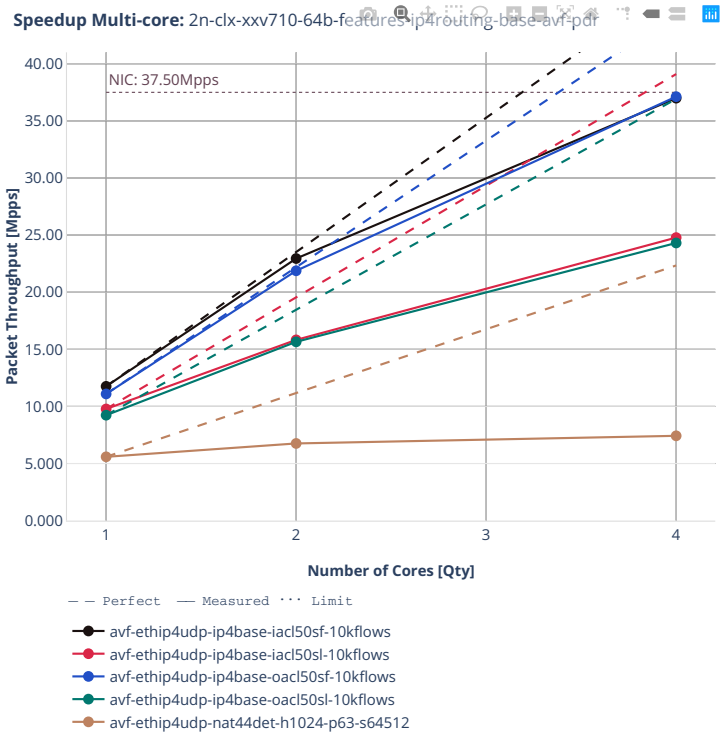
Speedup Multi-core: 2n-clx-xxv710-64b-features-ip4routing-base-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-ip4base-iacl50sf-10kflows
avf-ethip4udp-ip4base-iacl50sl-10kflows
avf-ethip4udp-ip4base-oacl50sf-10kflows
avf-ethip4udp-ip4base-oacl50sl-10kflows
avf-ethip4udp-nat44det-h1024-p63-s64512

2n-clx-x710

64b-ip4routing-base-scale-[avf,dpdk]

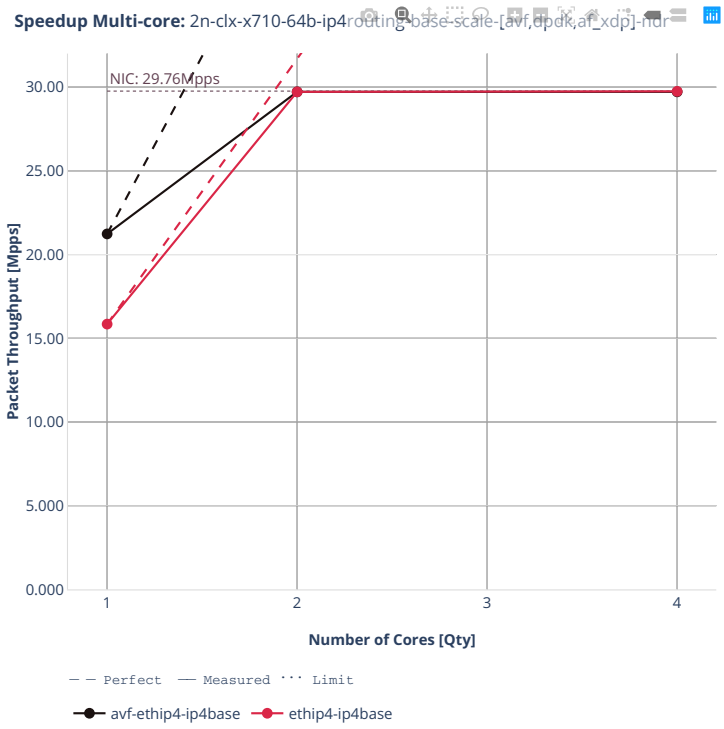

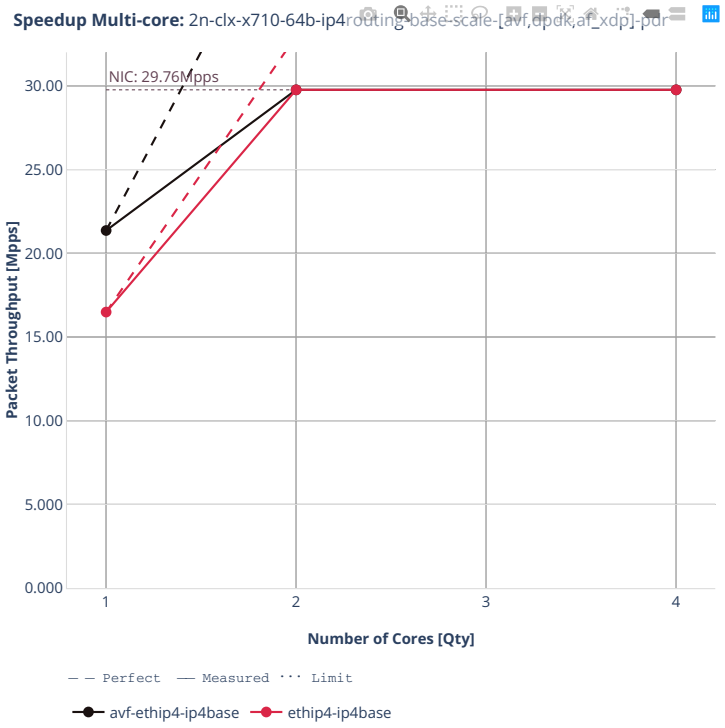
Speedup Multi-core: 2n-clx-x710-64b-ip4routing-base-scale-[avf,dpdk,af_xdp]-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4-ip4base
ethip4-ip4base

2n-clx-cx556a

64b-ip4routing-base-scale-rdma-core

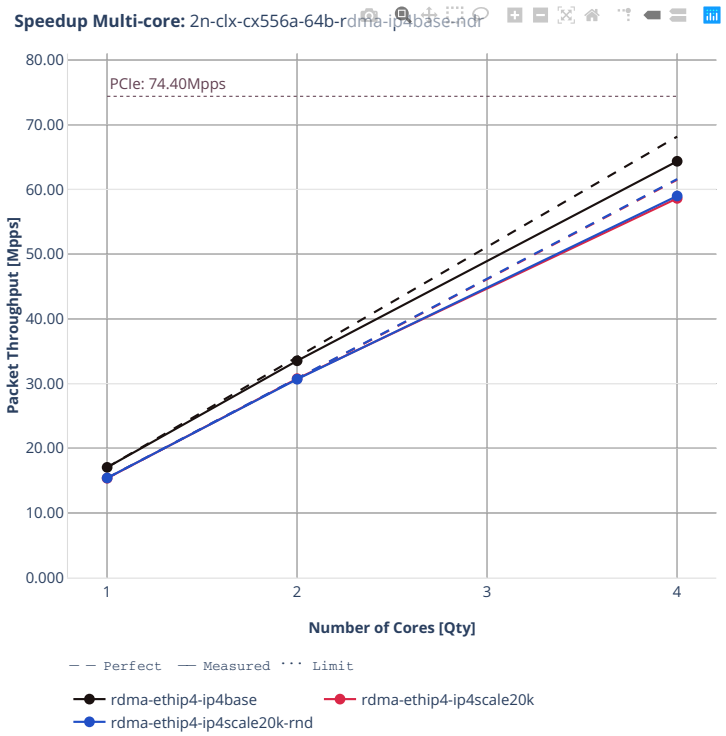

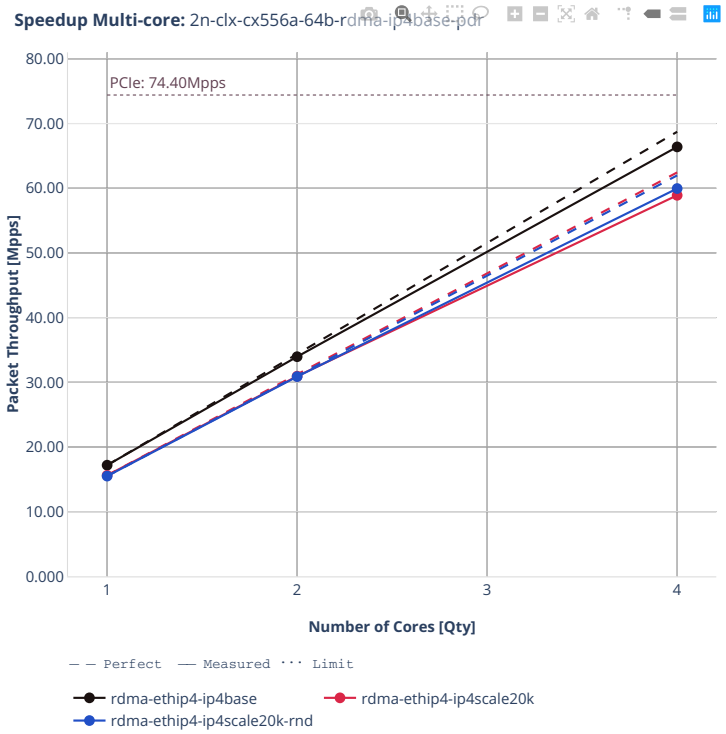
Speedup Multi-core: 2n-clx-cx556a-64b-rdma-ip4base-pdr
PCIe: 74.40Mpps
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-ethip4-ip4base
rdma-ethip4-ip4scale20k
rdma-ethip4-ip4scale20k-rnd

64b-ip4routing-features

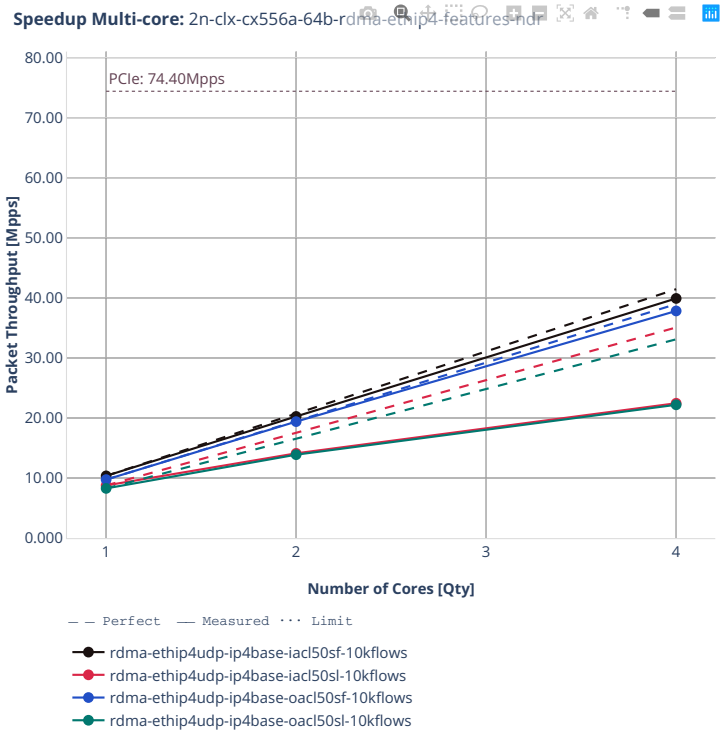

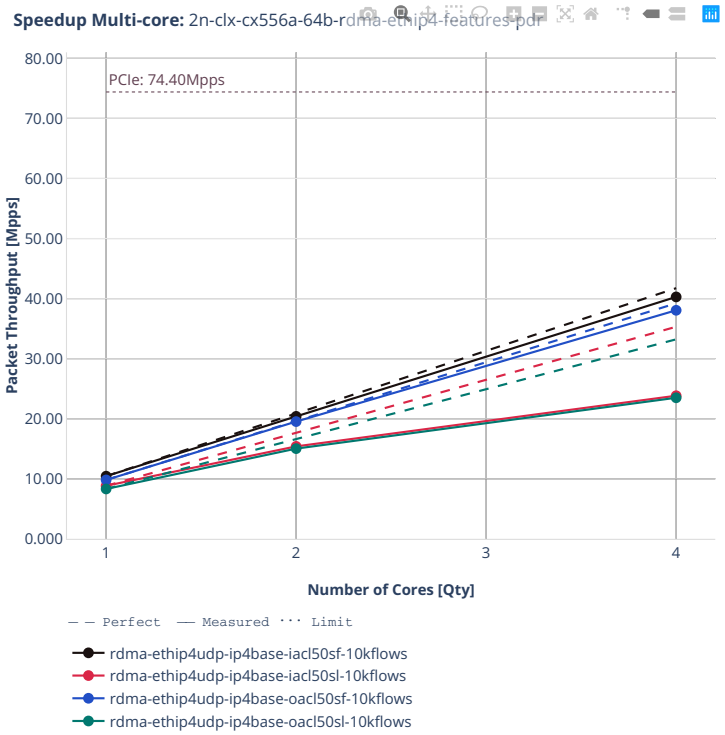
Speedup Multi-core: 2n-clx-cx556a-64b-rdma-ethip4-features-pdr
PCIe: 74.40Mpps
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-ethip4udp-ip4base-iacl50sf-10kflows
rdma-ethip4udp-ip4base-iacl50sl-10kflows
rdma-ethip4udp-ip4base-oacl50sf-10kflows
rdma-ethip4udp-ip4base-oacl50sl-10kflows

2n-clx-e810cq

64b-ip4routing-base-scale-avf

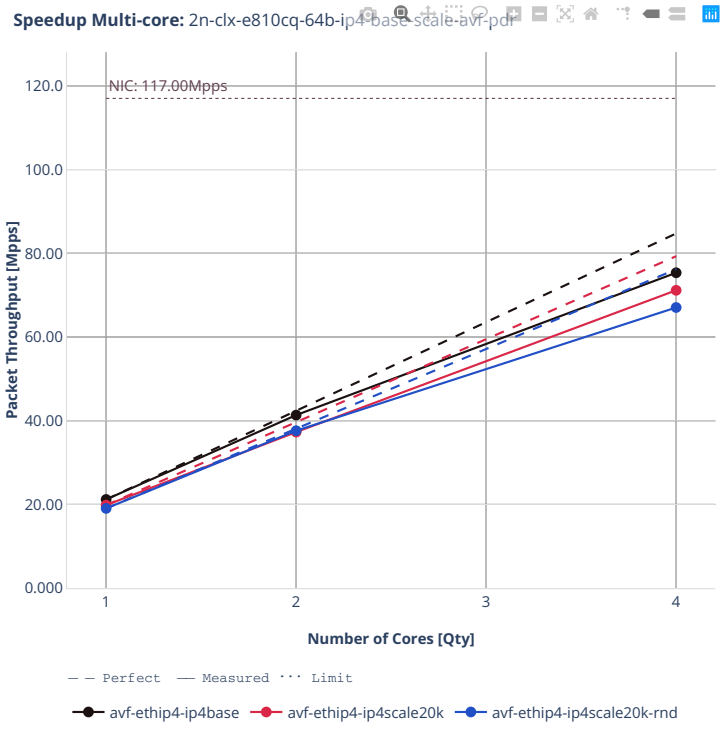
Speedup Multi-core: 2n-clx-e810cq-64b-ip4-base-scale-avf-pdr
NIC: 117.00Mpps
120.0
100.0
80.00
60.00
40.00
20.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4-ip4base
avf-ethip4-ip4scale20k
avf-ethip4-ip4scale20k-rnd

64b-ip4routing-features-avf

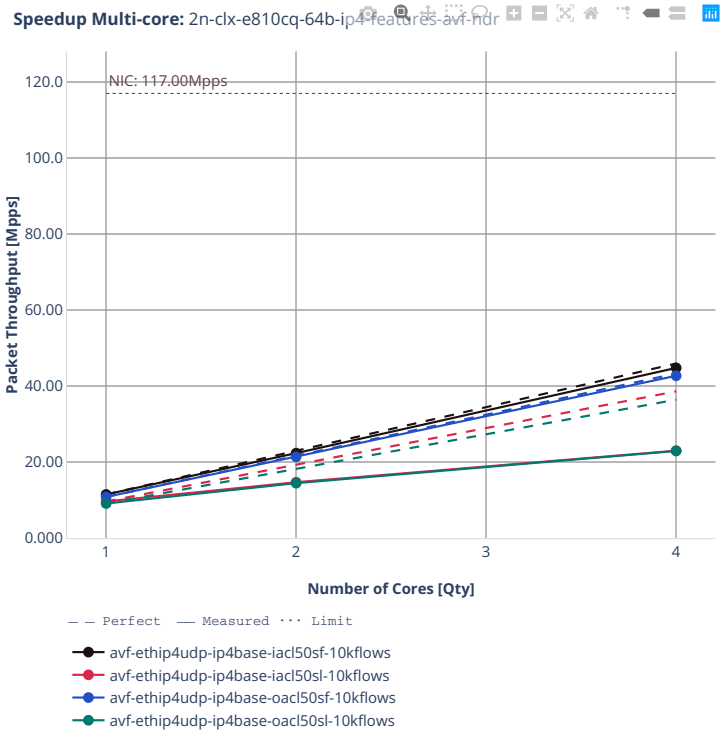

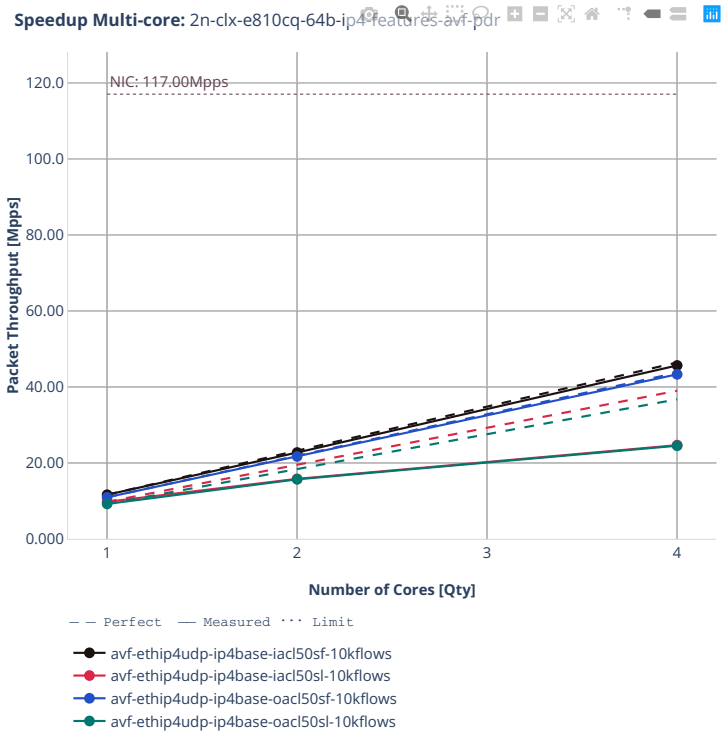
Speedup Multi-core: 2n-clx-e810cq-64b-ip4-features-avf-pdr
NIC: 117.00Mpps
120.0
100.0
80.00
60.00
40.00
20.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-ip4base-iacl50sf-10kflows
avf-ethip4udp-ip4base-iacl50sl-10kflows
avf-ethip4udp-ip4base-oacl50sf-10kflows
avf-ethip4udp-ip4base-oacl50sl-10kflows

64b-ip4routing-base-scale-dpdk

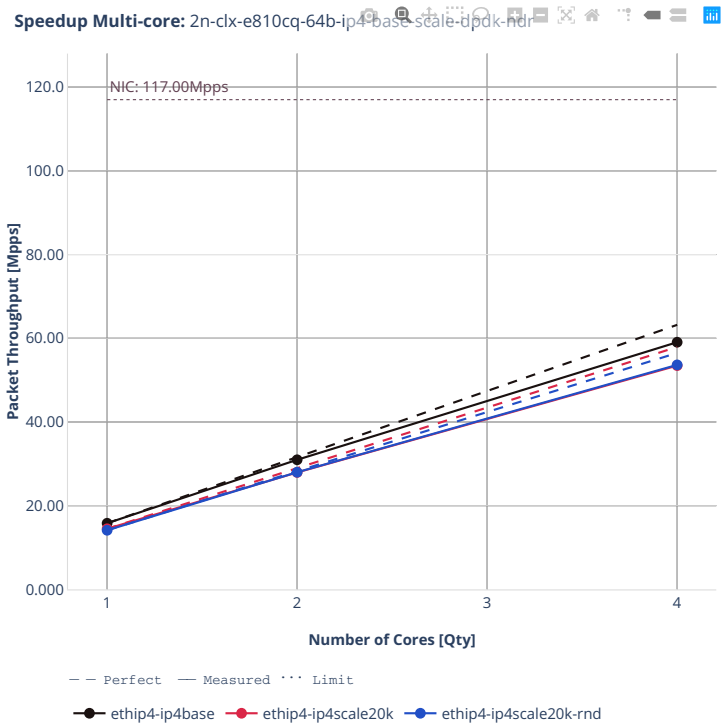

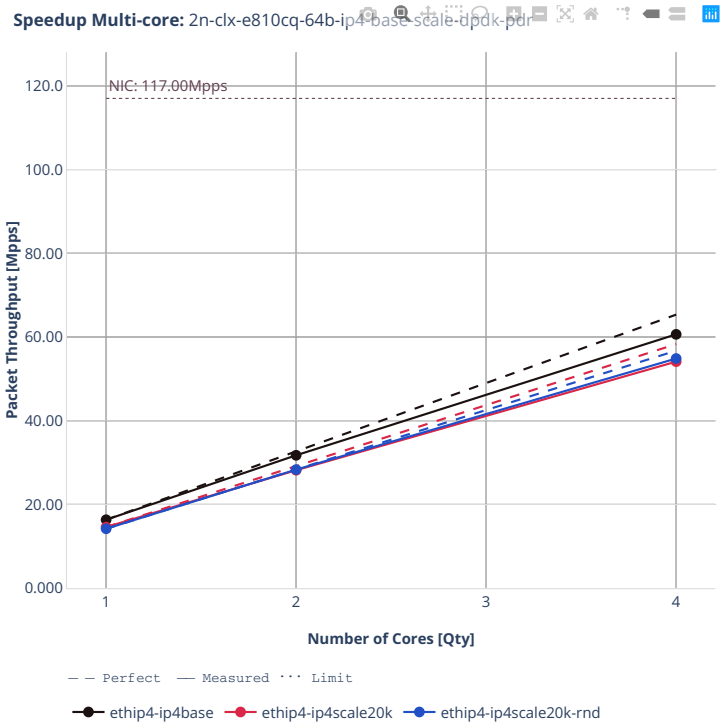
Speedup Multi-core: 2n-clx-e810cq-64b-ip4-base-scale-dpdk-pdr
NIC: 117.00Mpps
120.0
100.0
80.00
60.00
40.00
20.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base
ethip4-ip4scale20k
ethip4-ip4scale20k-rnd

2n-zn2-xxv710

64b-ip4routing-base-scale-avf

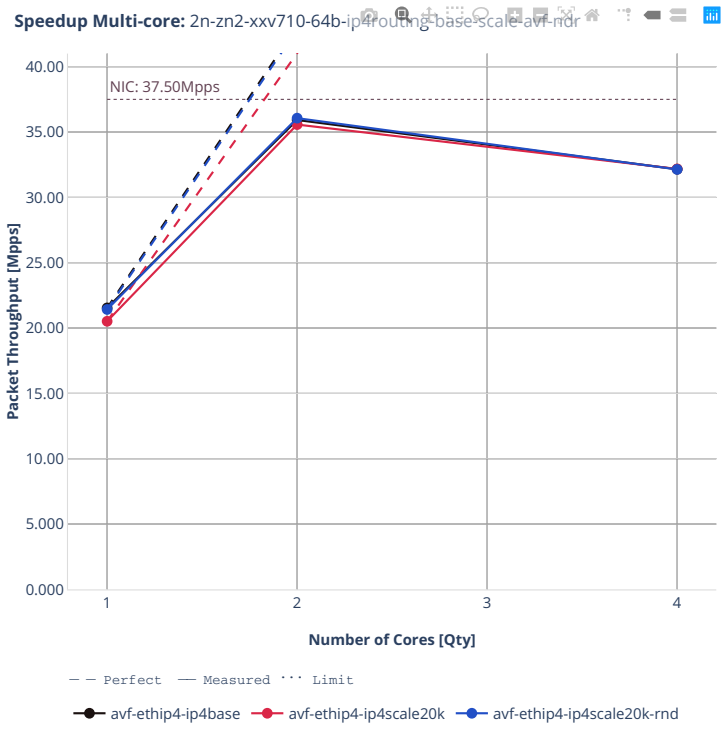

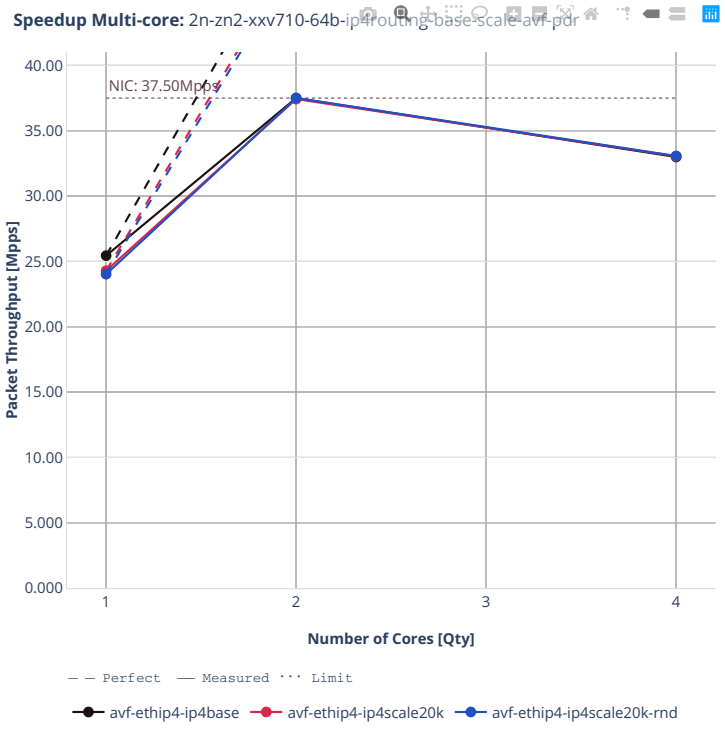
Speedup Multi-core: 2n-zn2-xxv710-64b-ip4routing-base-scale-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4-ip4base
avf-ethip4-ip4scale20k
avf-ethip4-ip4scale20k-rnd

64b-ip4routing-base-scale-dpdk

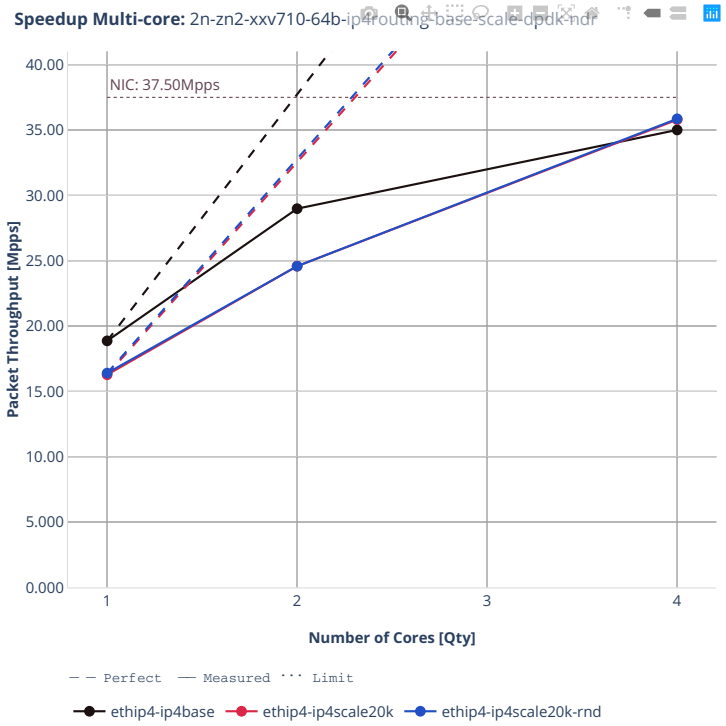

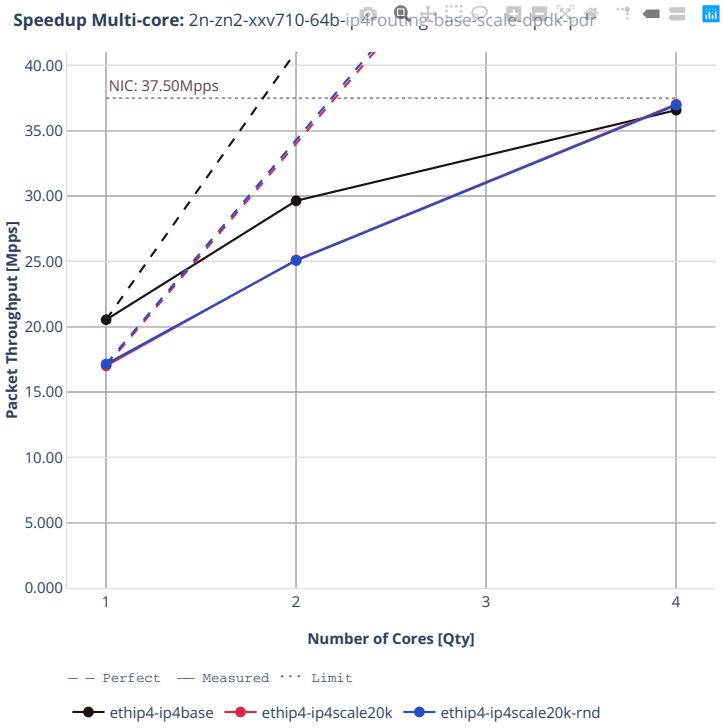
Speedup Multi-core: 2n-zn2-xxv710-64b-ip4routing-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base
ethip4-ip4scale20k
ethip4-ip4scale20k-rnd

64b-features-ip4routing-base-avf

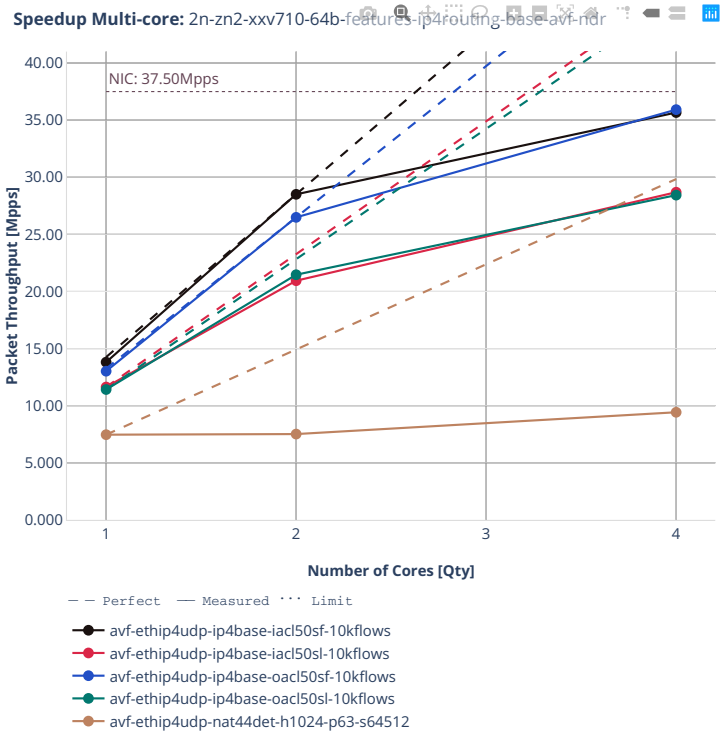

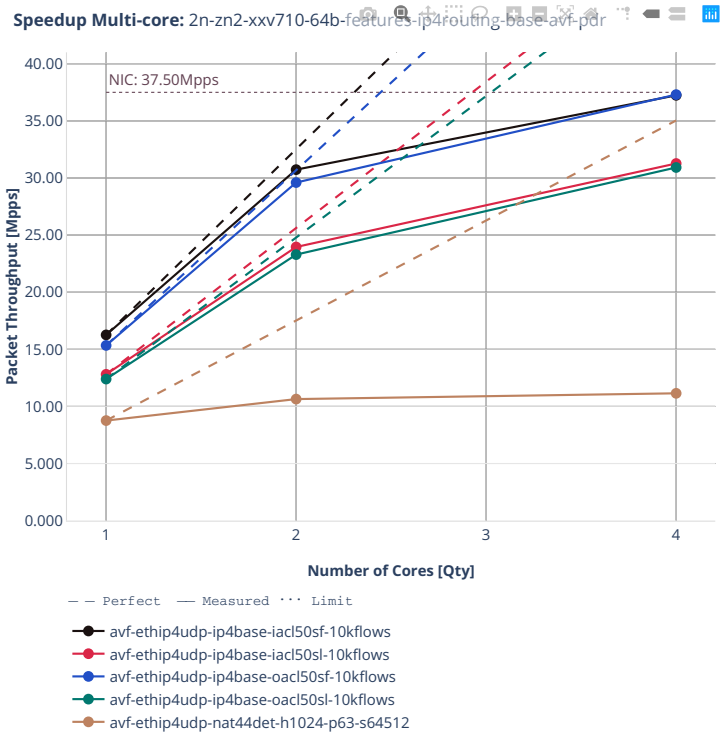
Speedup Multi-core: 2n-zn2-xxv710-64b-features-ip4routing-base-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-ip4base-iacl50sf-10kflows
avf-ethip4udp-ip4base-iacl50sl-10kflows
avf-ethip4udp-ip4base-oacl50sf-10kflows
avf-ethip4udp-ip4base-oacl50sl-10kflows
avf-ethip4udp-nat44det-h1024-p63-s64512

2n-zn2-x710

64b-ip4routing-base-[avf,dpdk]

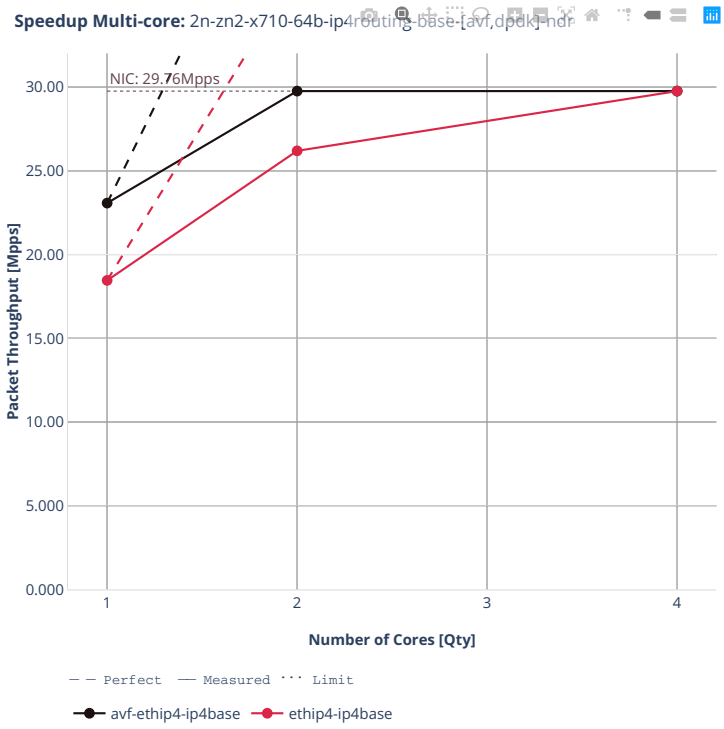

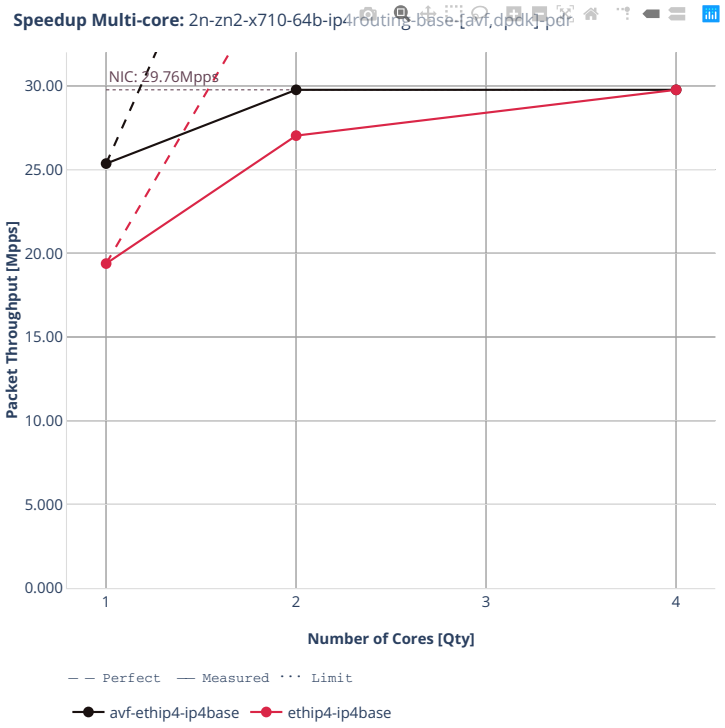
Speedup Multi-core: 2n-zn2-x710-64b-ip4routing-base-[avf,dpdk]-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4-ip4base
ethip4-ip4base

2n-zn2-cx556a

64b-ip4routing-base-scale-rdma-core

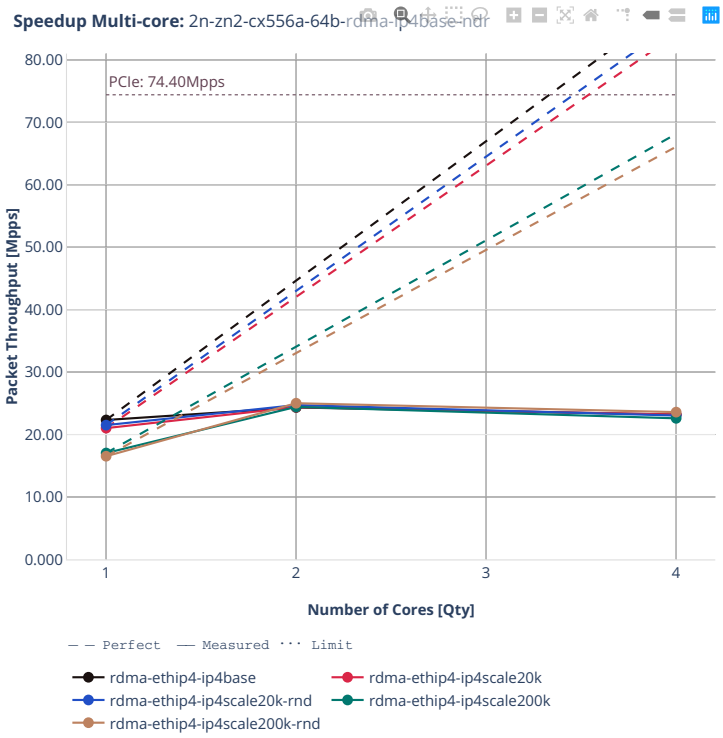

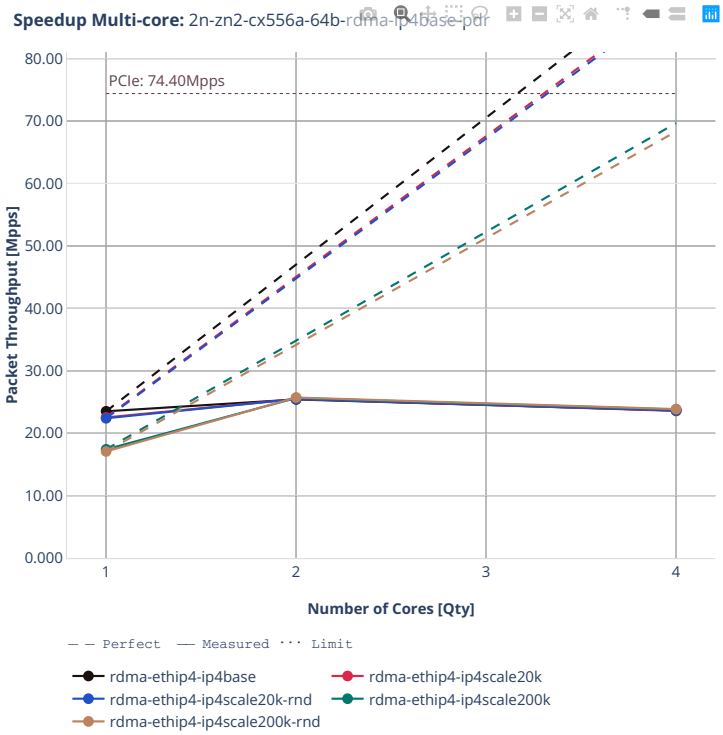
Speedup Multi-core: 2n-zn2-cx556a-64b-rdma-ip4base-pdr
PCIe: 74.40Mpps
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-ethip4-ip4base
rdma-ethip4-ip4scale20k
rdma-ethip4-ip4scale20k-rnd
rdma-ethip4-ip4scale200k
rdma-ethip4-ip4scale200k-rnd

64b-ip4routing-features

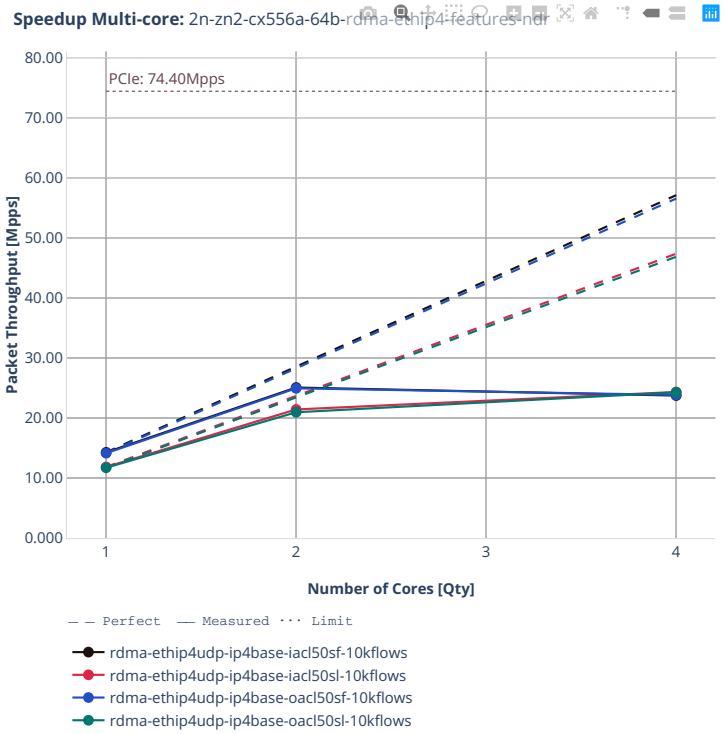

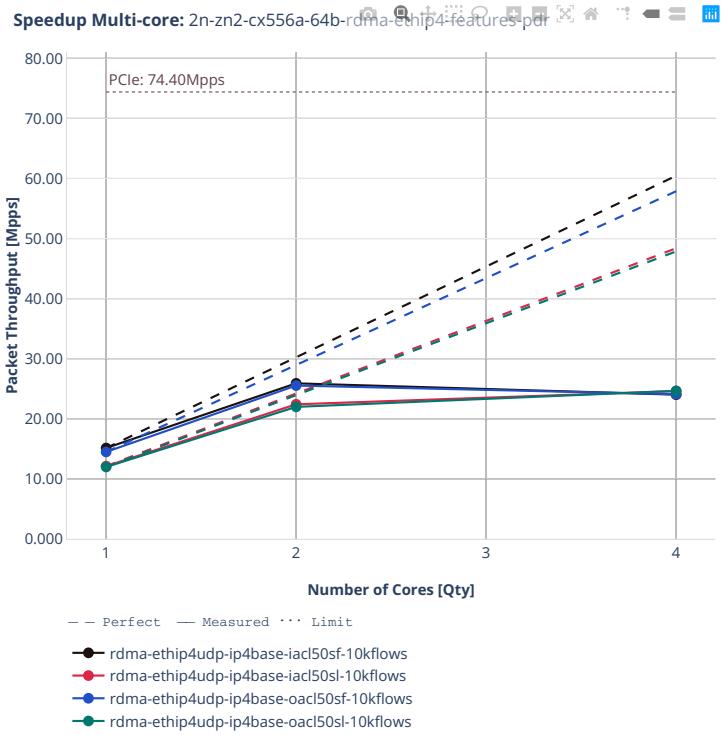
Speedup Multi-core: 2n-zn2-cx556a-64b-rdma-ethip4-features-pdr
PCIe: 74.40Mpps
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-ethip4udp-ip4base-iacl50sf-10kflows
rdma-ethip4udp-ip4base-iacl50sl-10kflows
rdma-ethip4udp-ip4base-oacl50sf-10kflows
rdma-ethip4udp-ip4base-oacl50sl-10kflows

3n-alt-xl710

64b-ip4routing-base-scale

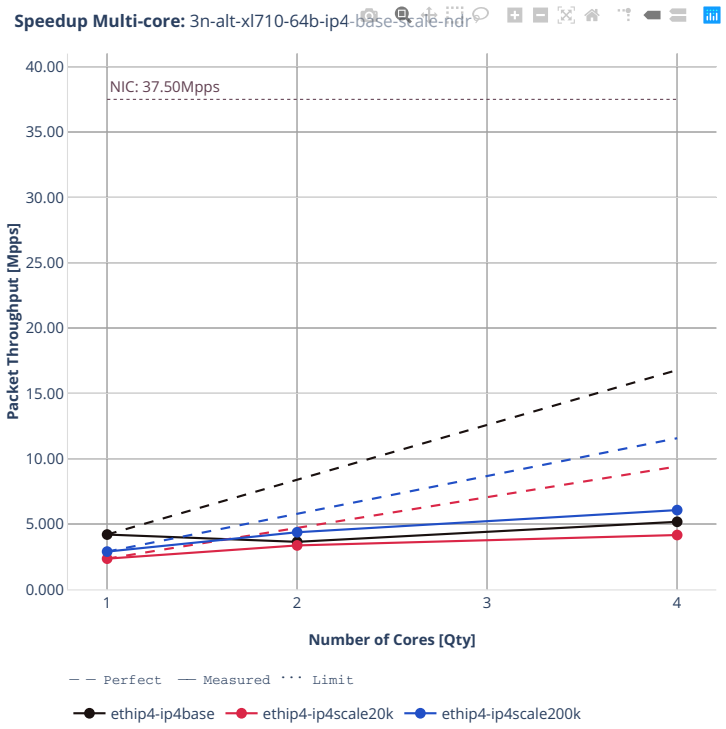

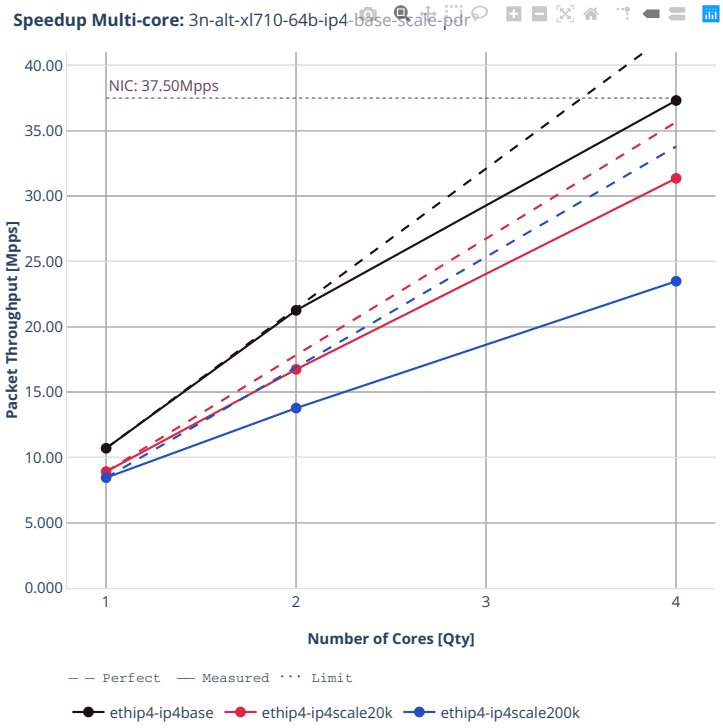
Speedup Multi-core: 3n-alt-xl710-64b-ip4-base-scale-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base
ethip4-ip4scale20k
ethip4-ip4scale200k

64b-ip4routing-features

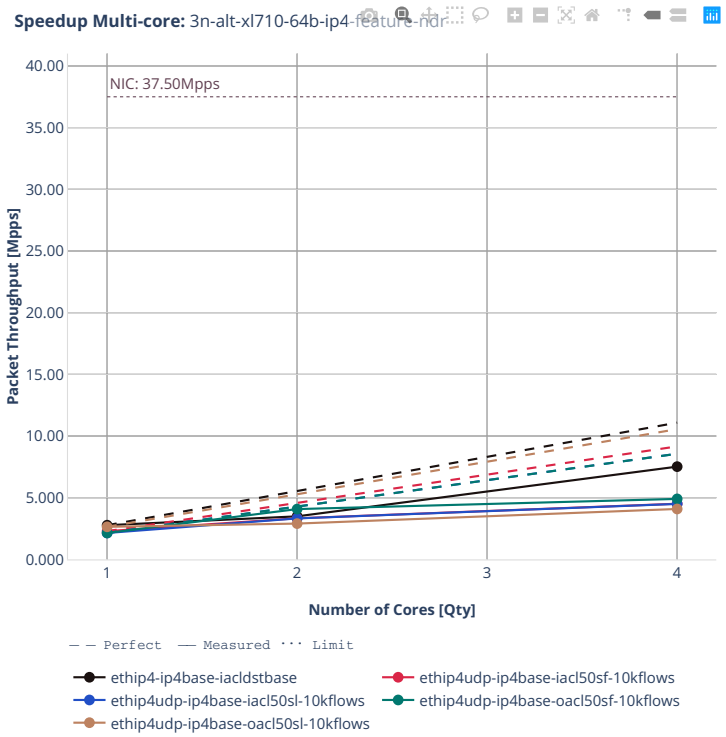

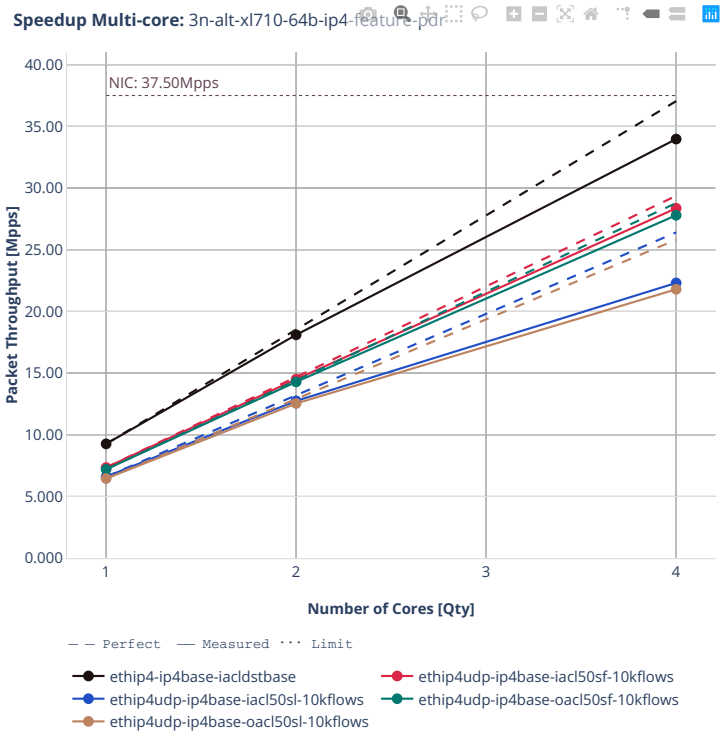
Speedup Multi-core: 3n-alt-xl710-64b-ip4-feature-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base-iacldstbase
ethip4udp-ip4base-iacl50sf-10kflows
ethip4udp-ip4base-iacl50sl-10kflows
ethip4udp-ip4base-oacl50sf-10kflows
ethip4udp-ip4base-oacl50sl-10kflows

3n-tsh-x520

64b-features-ip4routing-base-ixgbe

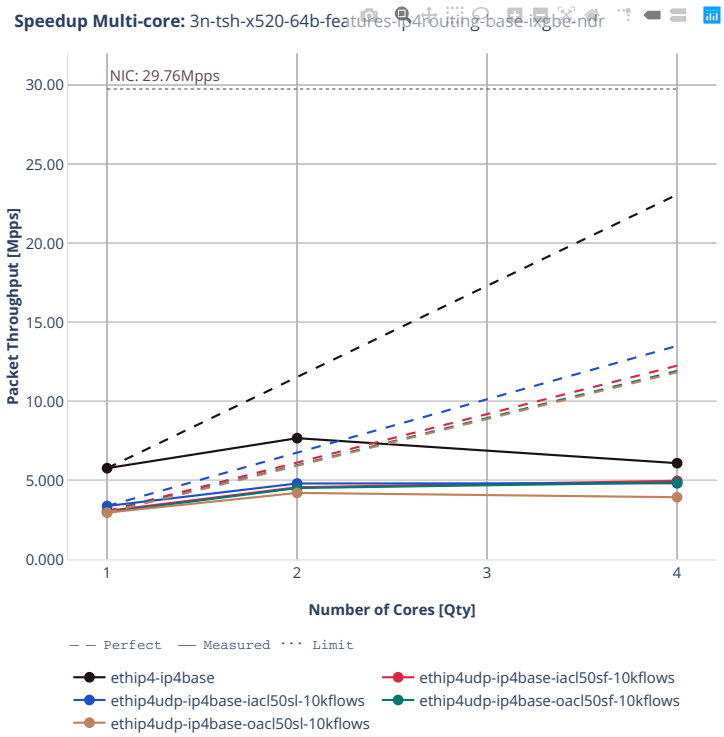

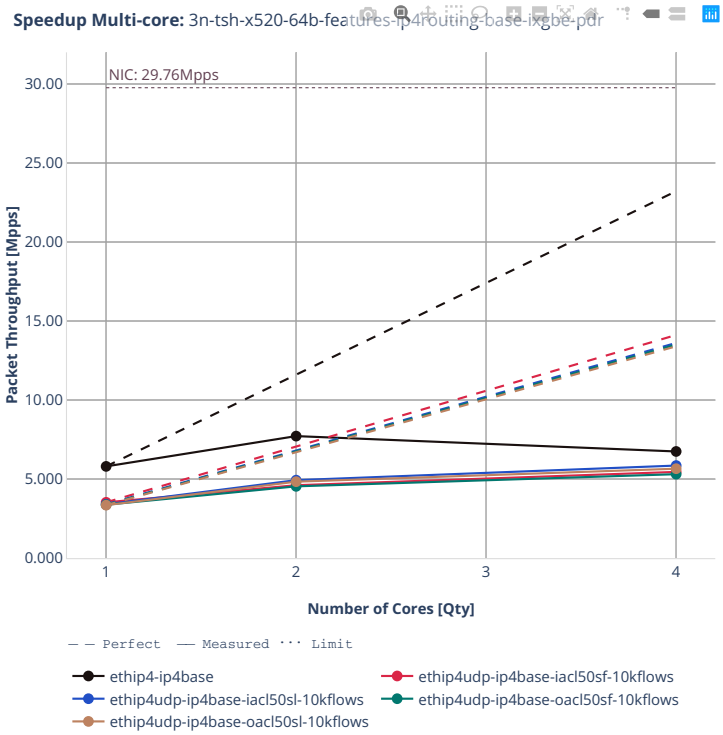
Speedup Multi-core: 3n-tsh-x520-64b-features-ip4routing-base-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base
ethip4udp-ip4base-iacl50sf-10kflows
ethip4udp-ip4base-iacl50sl-10kflows
ethip4udp-ip4base-oacl50sf-10kflows
ethip4udp-ip4base-oacl50sl-10kflows

2n-tx2-xl710

64b-ip4routing-base-scale-dpdk

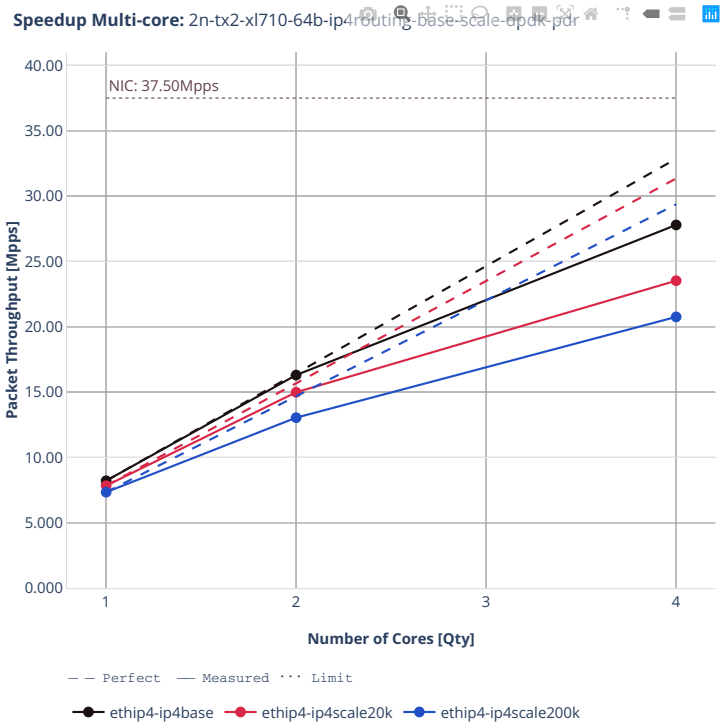
Speedup Multi-core: 2n-tx2-xl710-64b-ip4routing-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base
ethip4-ip4scale20k
ethip4-ip4scale200k

64b-ip4routing-base-scale-af-xdp

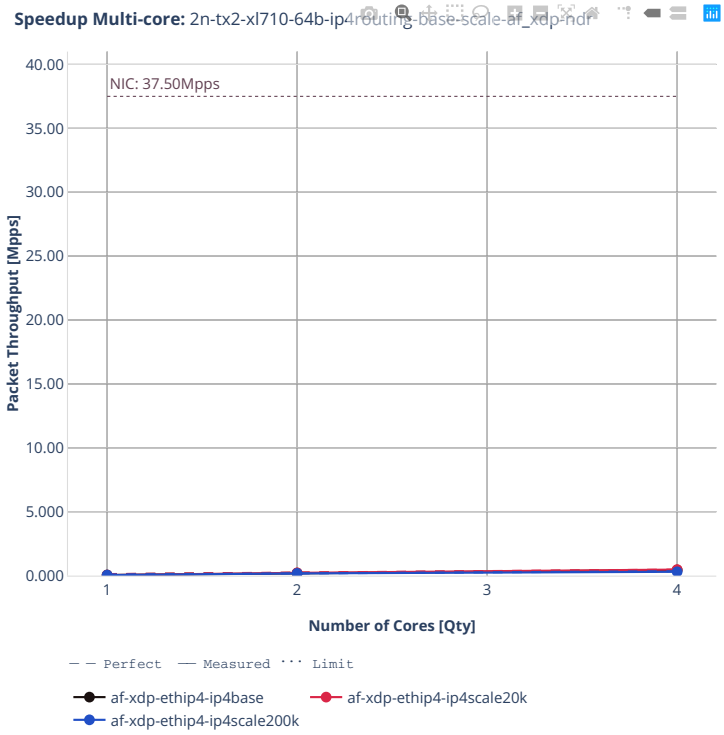

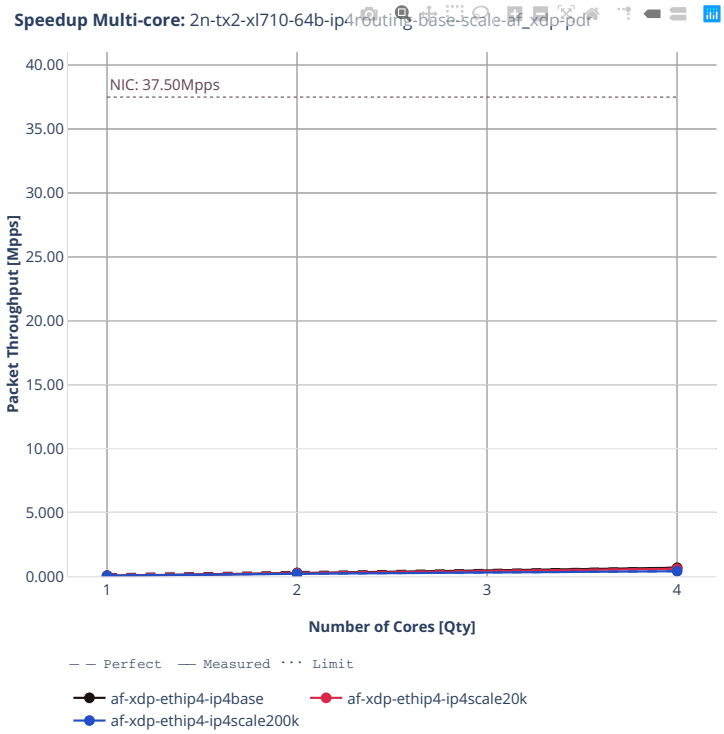
Speedup Multi-core: 2n-tx2-xl710-64b-ip4routing-base-scale-af_xdp-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
af-xdp-ethip4-ip4base
af-xdp-ethip4-ip4scale20k
af-xdp-ethip4-ip4scale200k

64b-features-ip4routing-base-dpdk

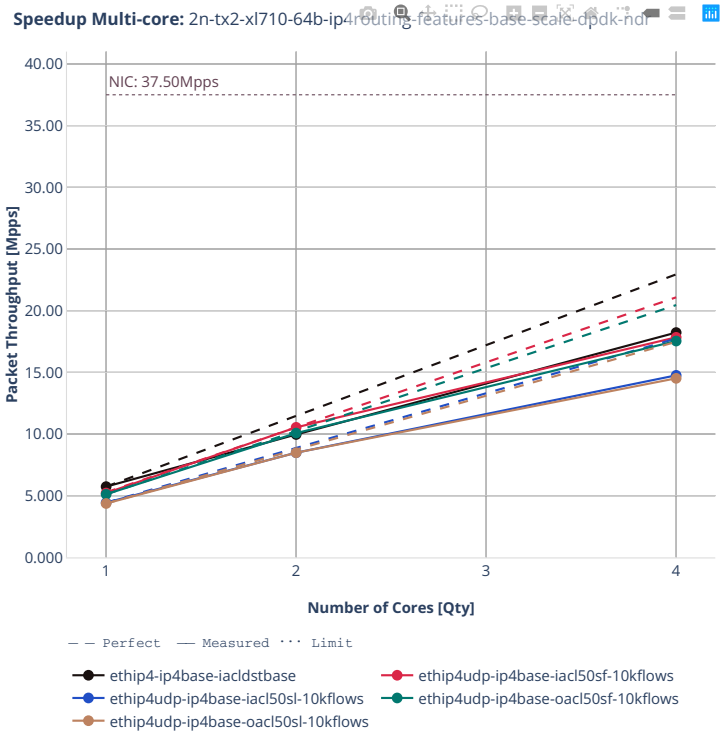

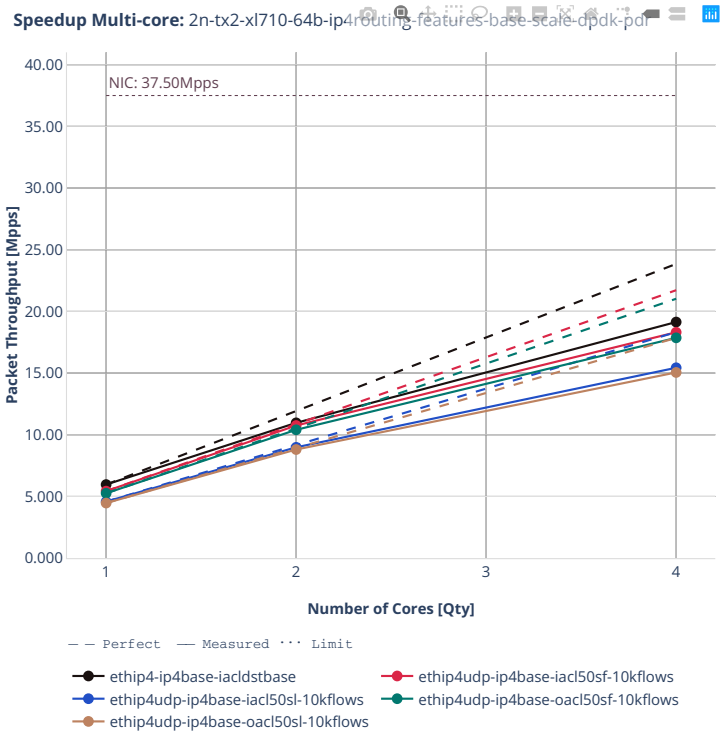
Speedup Multi-core: 2n-tx2-xl710-64b-ip4routing-features-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base-iacldstbase
ethip4udp-ip4base-iacl50sf-10kflows
ethip4udp-ip4base-iacl50sl-10kflows
ethip4udp-ip4base-oacl50sf-10kflows
ethip4udp-ip4base-oacl50sl-10kflows

2n-dnv-x553

64b-ip4routing-base-scale-ixgbe

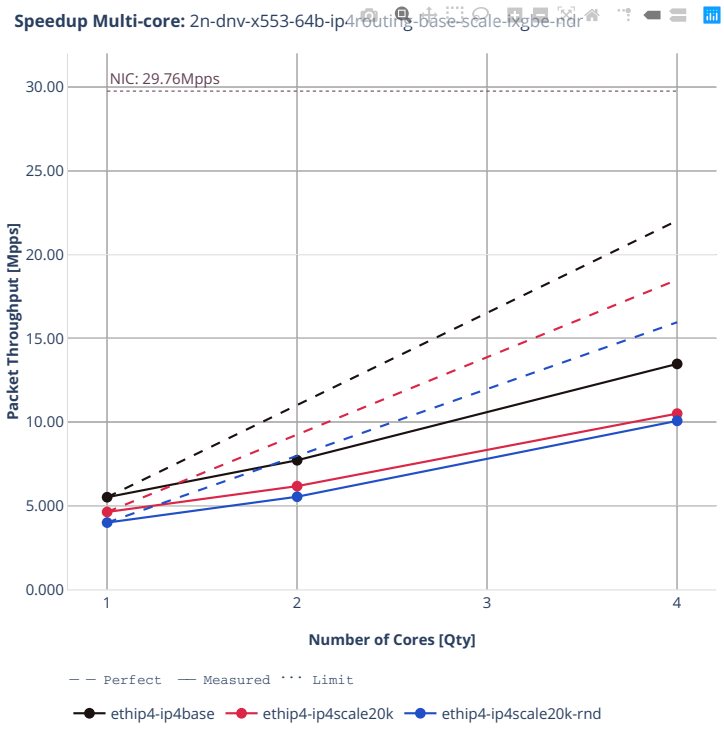

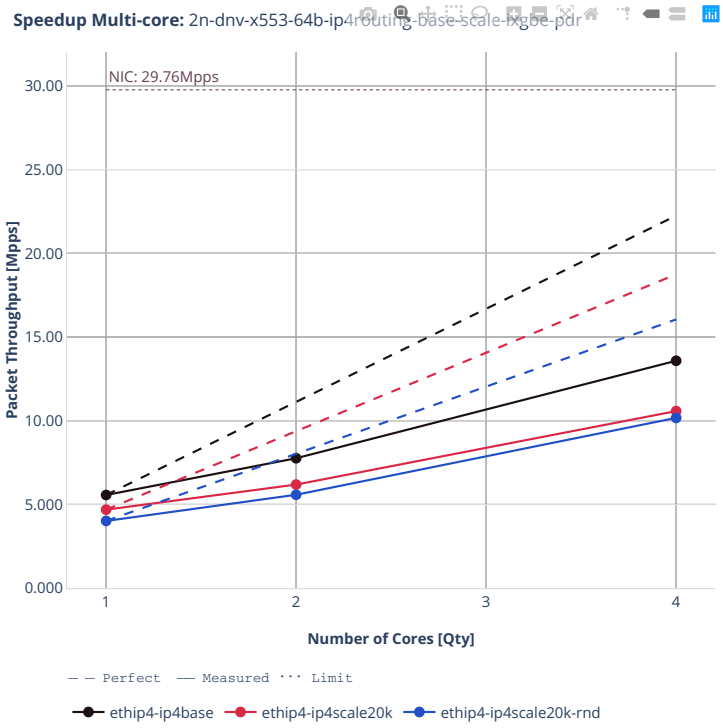
Speedup Multi-core: 2n-dnv-x553-64b-ip4routing-base-scale-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base
ethip4-ip4scale20k
ethip4-ip4scale20k-rnd

3n-dnv-x553

64b-ip4routing-base-scale-ixgbe

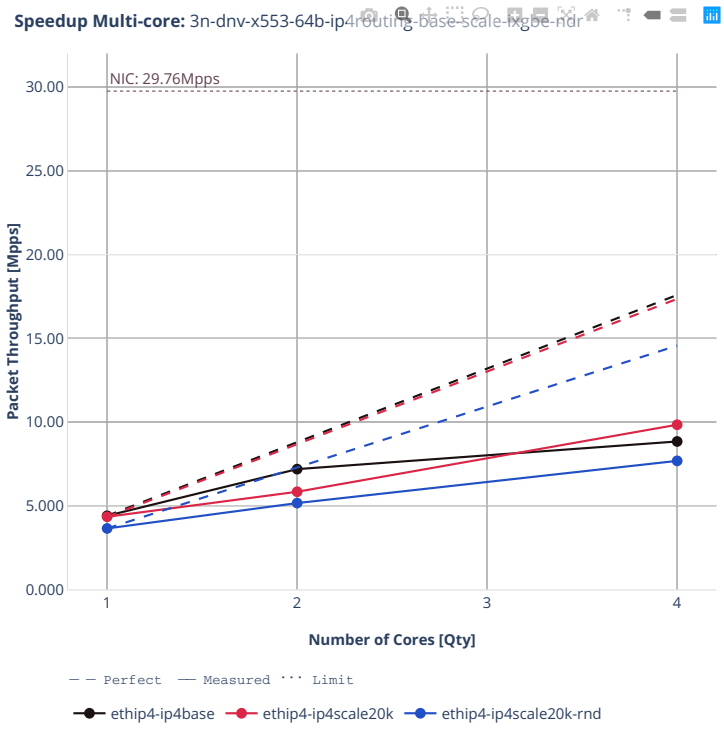

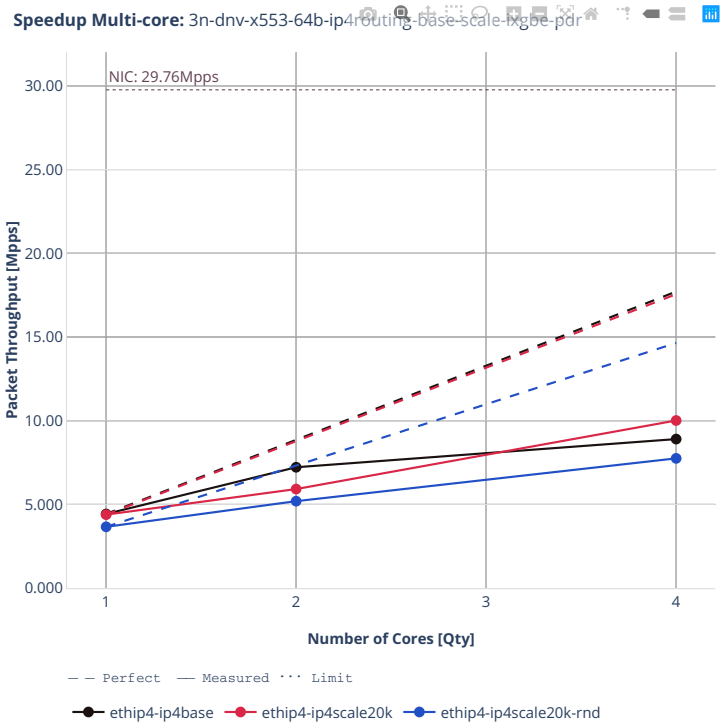
Speedup Multi-core: 3n-dnv-x553-64b-ip4routing-base-scale-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base
ethip4-ip4scale20k
ethip4-ip4scale20k-rnd

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[132].

[132] https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls2206

2n-icx-xxv710

78b-ip6routing-base-scale-avf

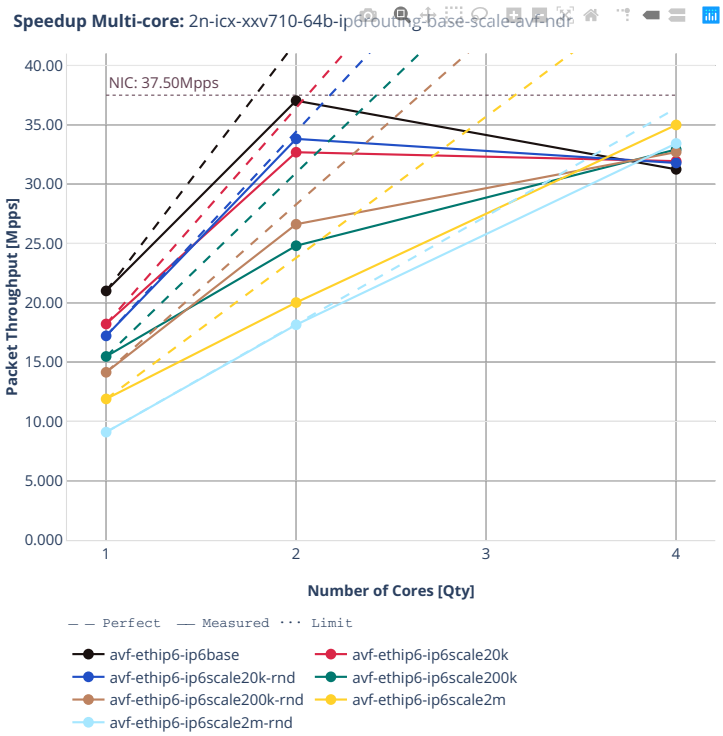

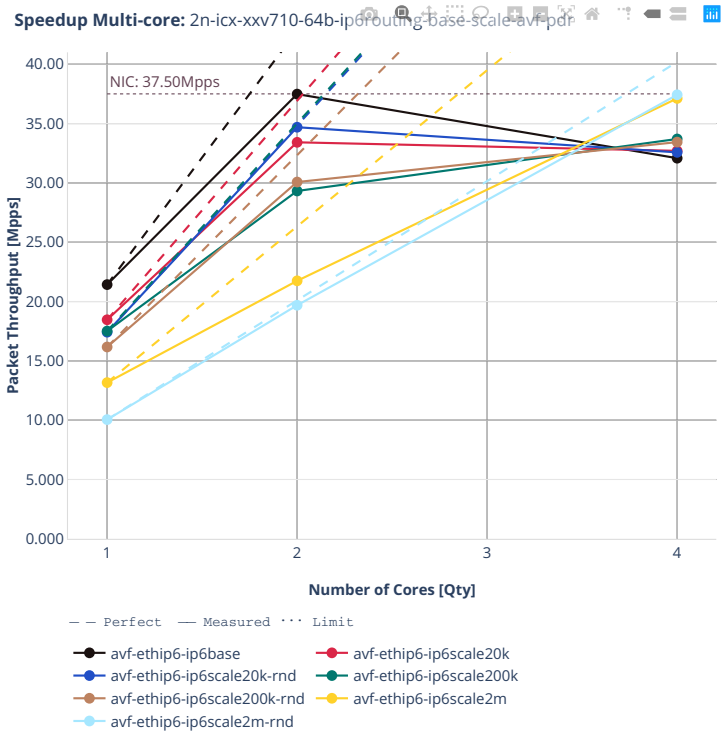
Speedup Multi-core: 2n-icx-xxv710-64b-ip6routing-base-scale-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip6-ip6base
avf-ethip6-ip6scale20k
avf-ethip6-ip6scale20k-rnd
avf-ethip6-ip6scale200k
avf-ethip6-ip6scale200k-rnd
avf-ethip6-ip6scale2m
avf-ethip6-ip6scale2m-rnd

78b-ip6routing-base-scale-dpdk

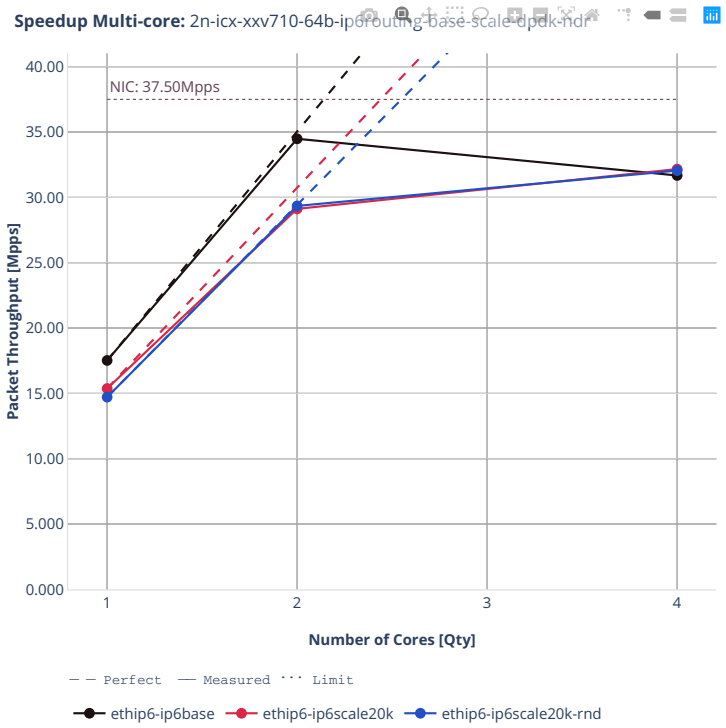

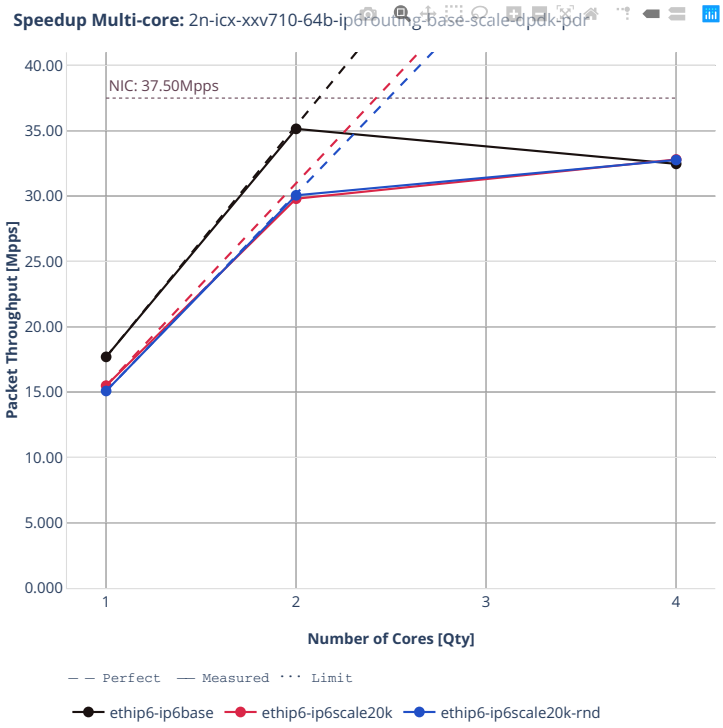
Speedup Multi-core: 2n-icx-xxv710-64b-ip6routing-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip6-ip6base
ethip6-ip6scale20k
ethip6-ip6scale20k-rnd

78b-ip6routing-base-scale-af_xdp

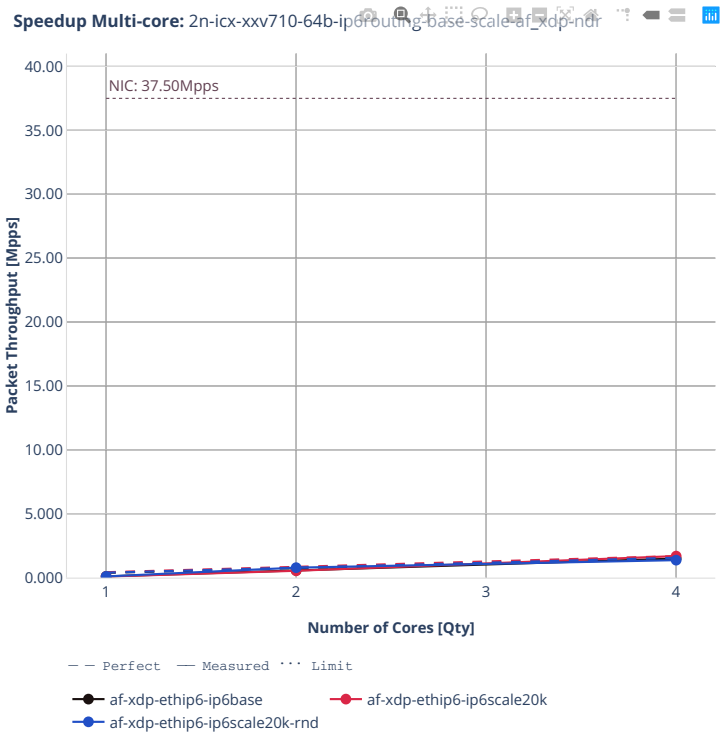

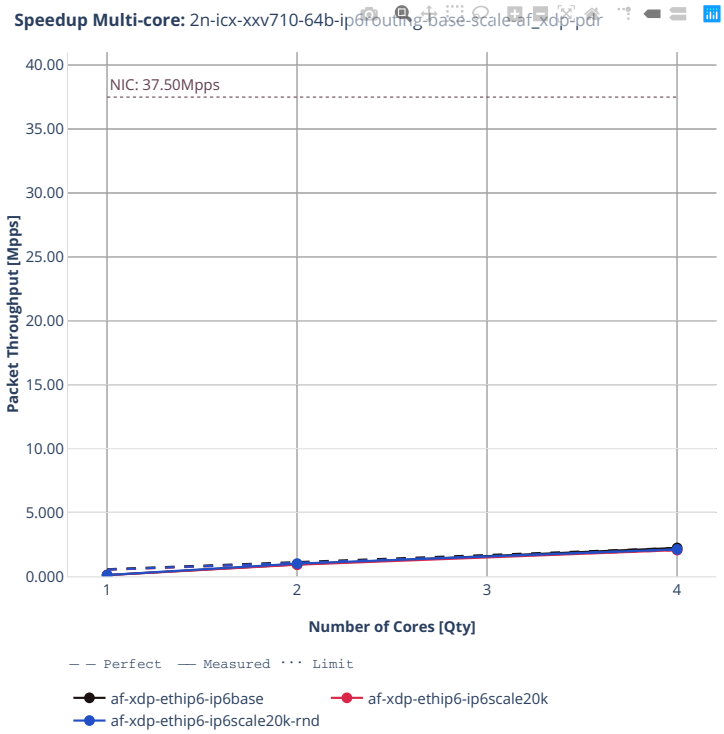
Speedup Multi-core: 2n-icx-xxv710-64b-ip6routing-base-scale-af_xdp-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
af-xdp-ethip6-ip6base
af-xdp-ethip6-ip6scale20k
af-xdp-ethip6-ip6scale20k-rnd

3n-icx-xxv710

78b-ip6routing-base

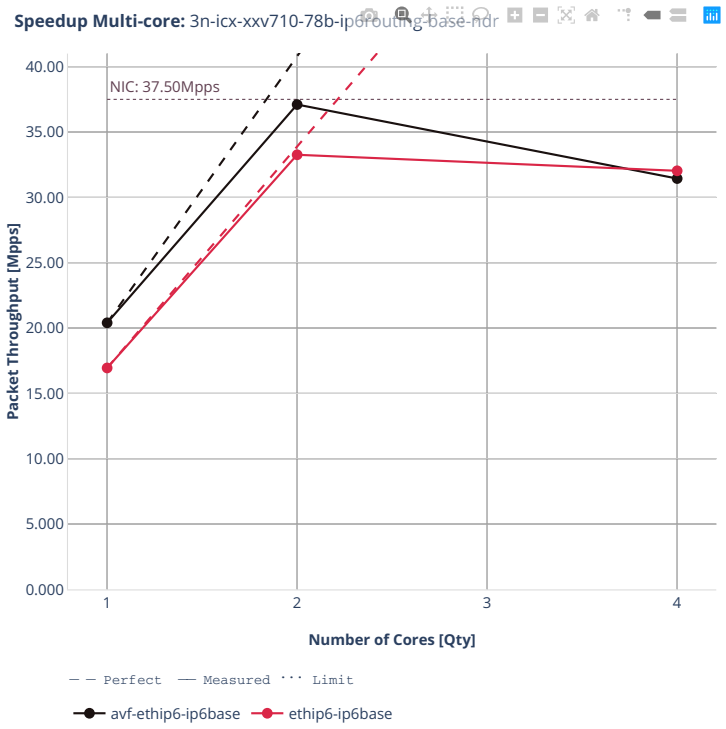

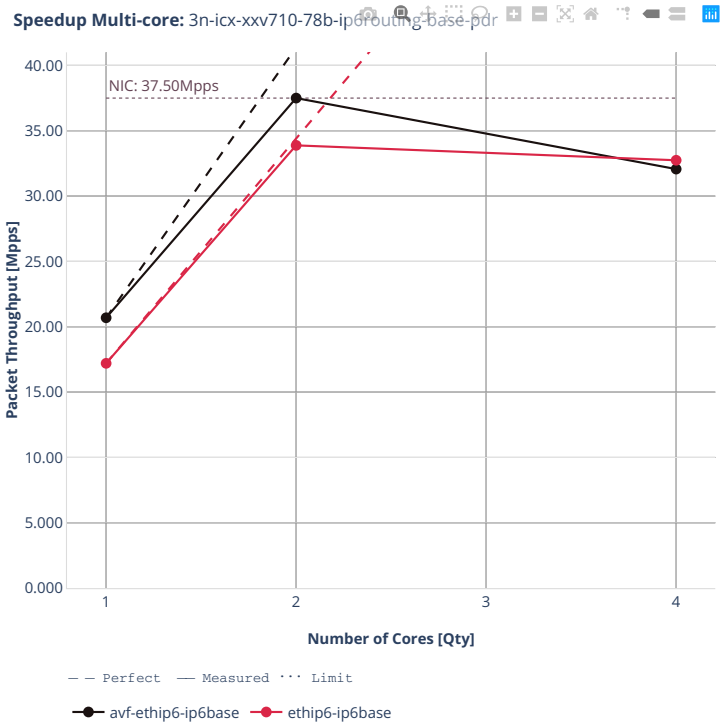
Speedup Multi-core: 3n-icx-xxv710-78b-ipe
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip6-ip6base
ethip6-ip6base

2n-skx-xxv710

78b-ip6routing-base-scale-avf

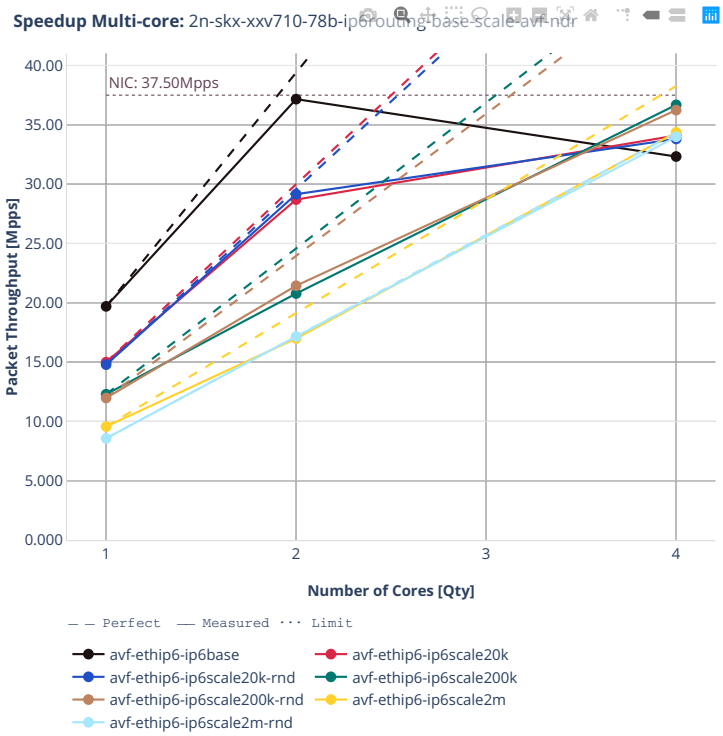

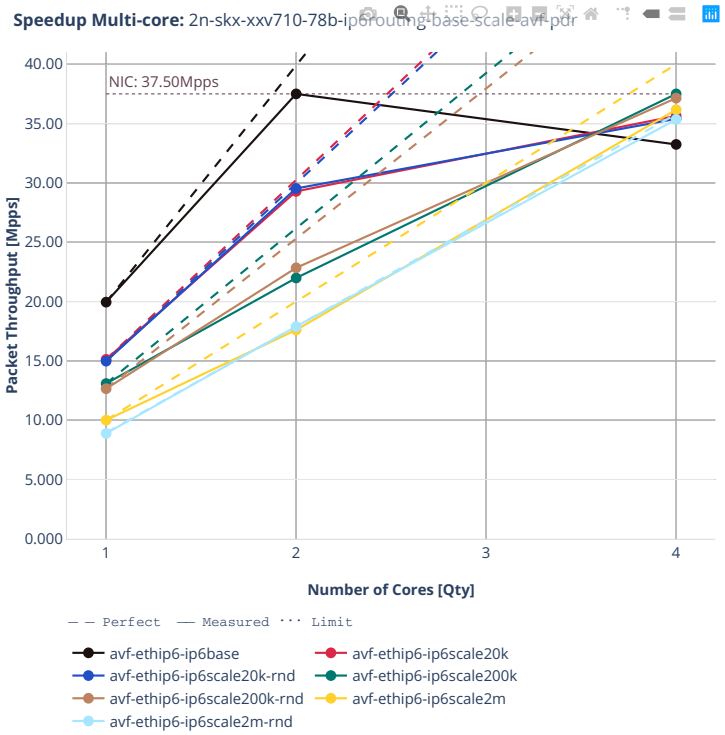
Speedup Multi-core: 2n-skx-xxv710-78b-ip6routing-base-scale-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip6-ip6base
avf-ethip6-ip6scale20k
avf-ethip6-ip6scale20k-rnd
avf-ethip6-ip6scale200k
avf-ethip6-ip6scale200k-rnd
avf-ethip6-ip6scale2m
avf-ethip6-ip6scale2m-rnd

78b-ip6routing-base-scale-af-xdp

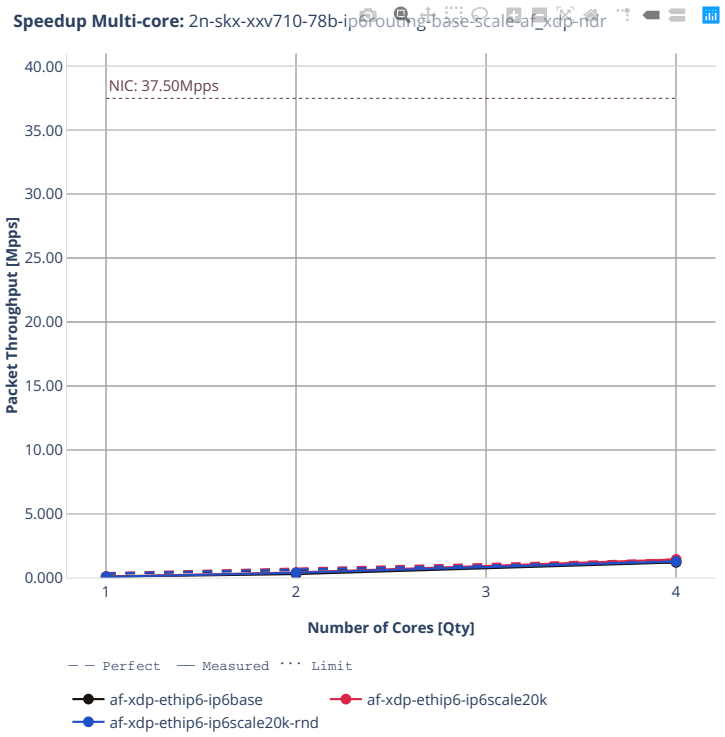

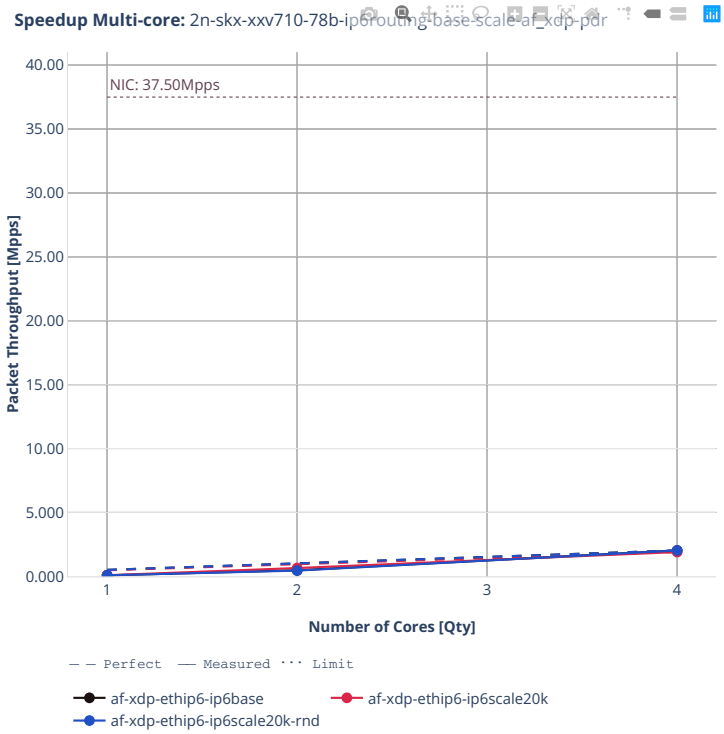
Speedup Multi-core: 2n-skx-xxv710-78b-ip6routing-base-scale-af_xdp-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
af-xdp-ethip6-ip6base
af-xdp-ethip6-ip6scale20k
af-xdp-ethip6-ip6scale20k-rnd

78b-ip6routing-base-scale-dpdk

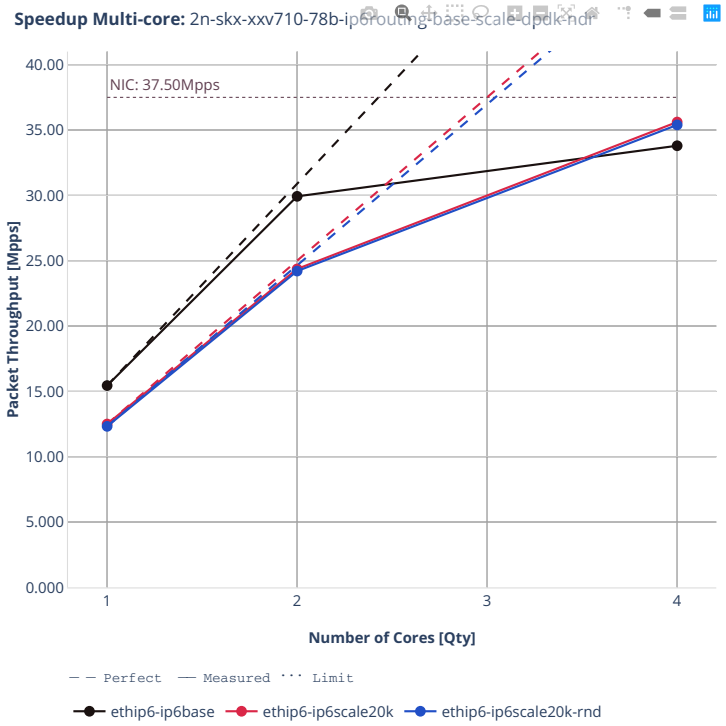

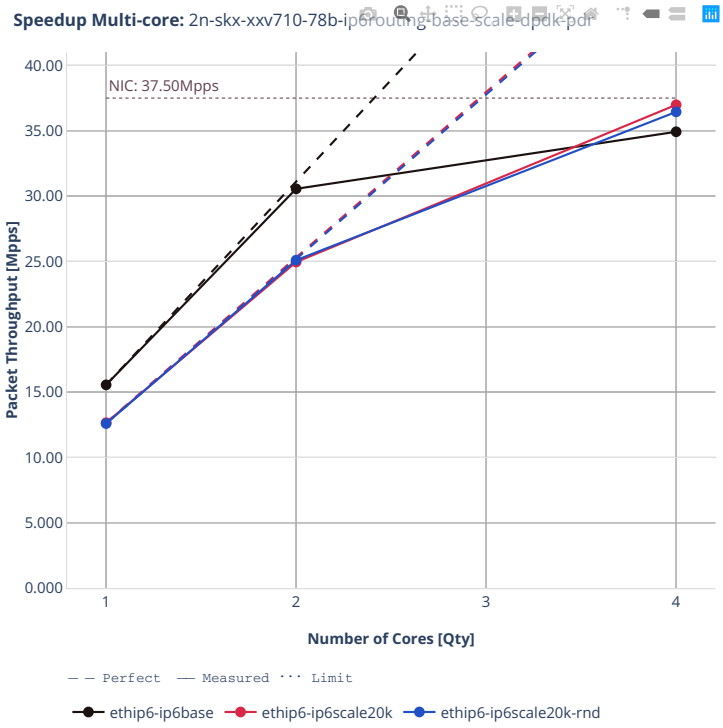
Speedup Multi-core: 2n-skx-xxv710-78b-ip6routing-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip6-ip6base
ethip6-ip6scale20k
ethip6-ip6scale20k-rnd

2n-skx-x710

78b-ip6routing-base-scale-[avf,dpdk]

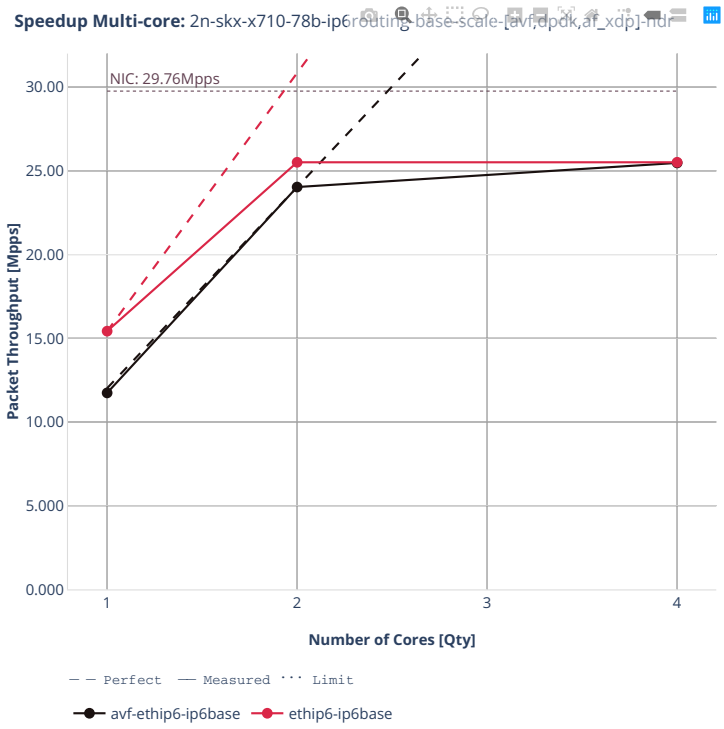

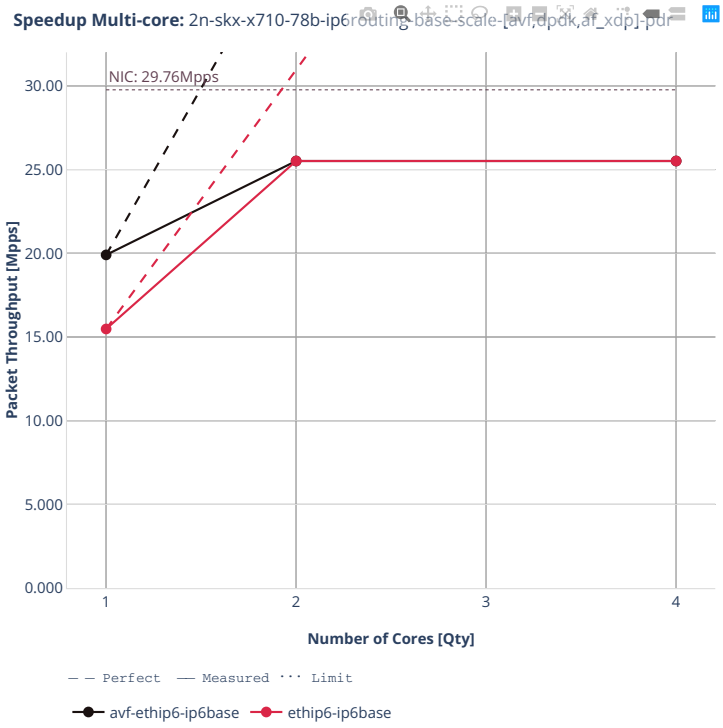
Speedup Multi-core: 2n-skx-x710-78b-ip6routing-base-scale-[avf,dpdk,af_xdp]-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip6-ip6base
ethip6-ip6base

3n-skx-xxv710

78b-ip6routing-base-[avf,dpdk]

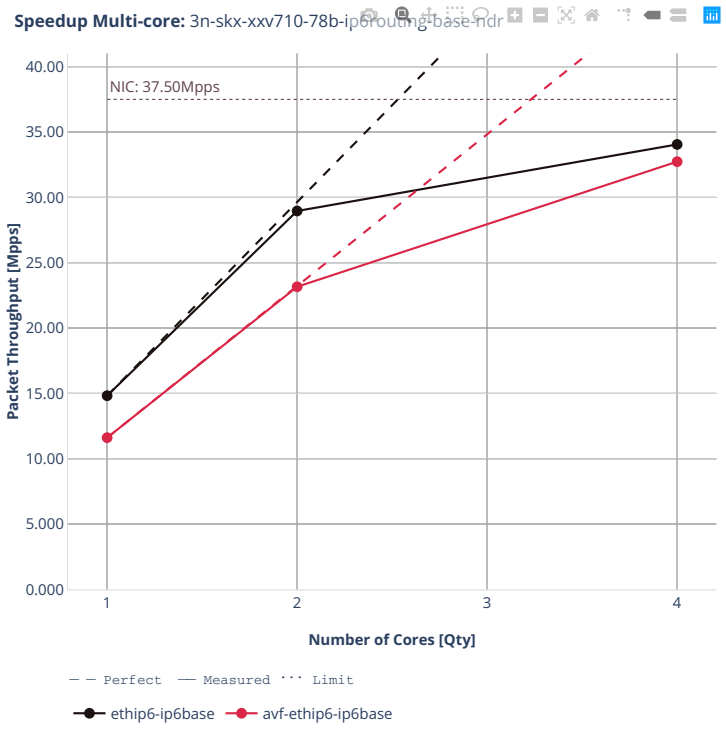

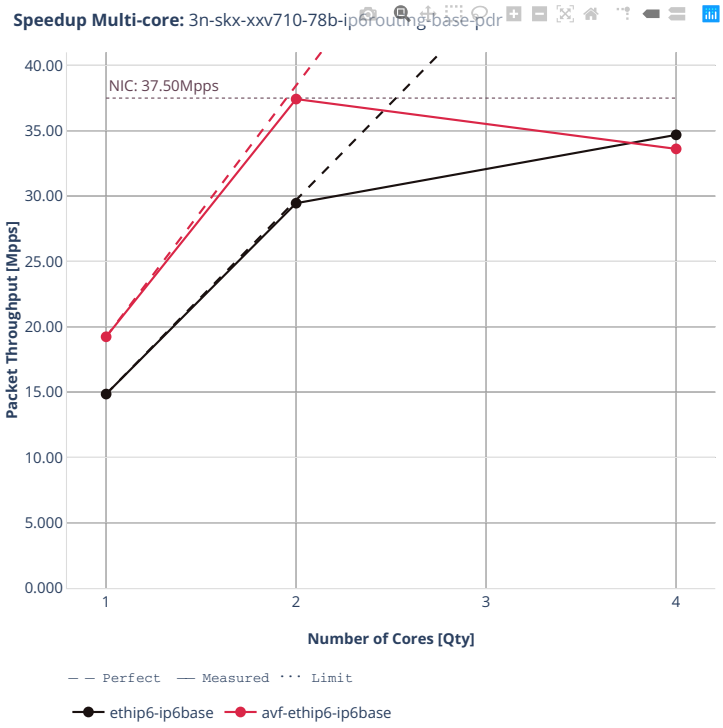
Speedup Multi-core: 3n-skx-xxv710-78b-ip6routing-base-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip6-ip6base
avf-ethip6-ip6base

3n-skx-x710

78b-ip6routing-base-scale-avf

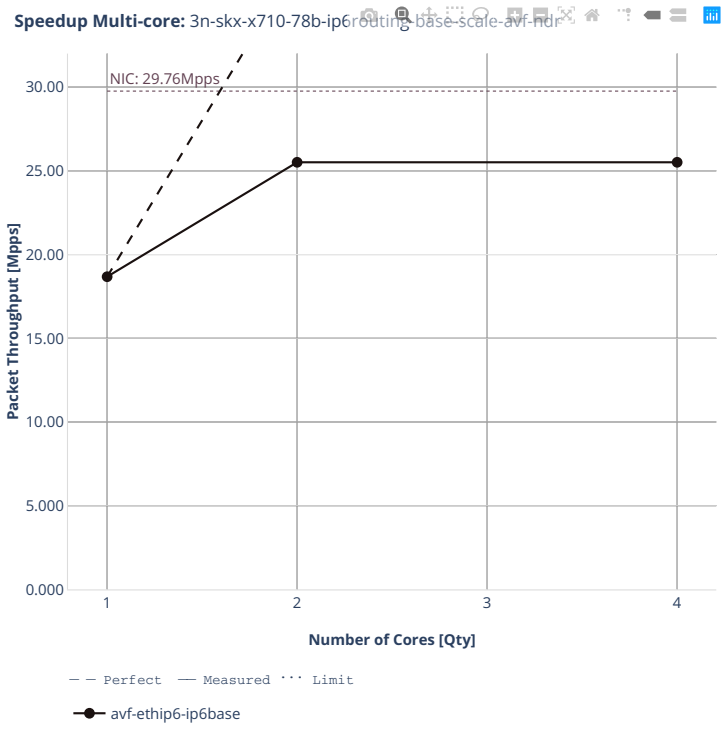

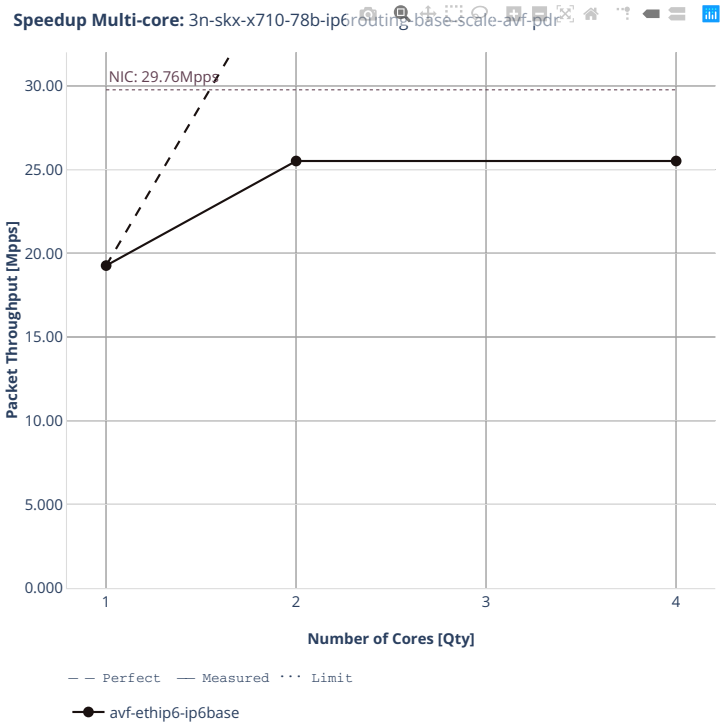
Speedup Multi-core: 3n-skx-x710-78b-ip6routing-base-scale-avf-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip6-ip6base

2n-clx-xxv710

78b-ip6routing-base-scale-avf

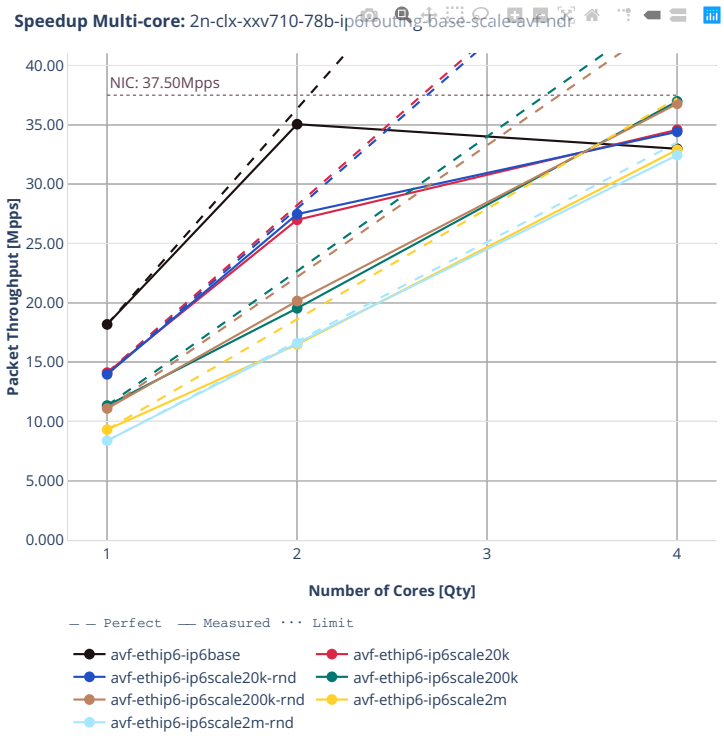

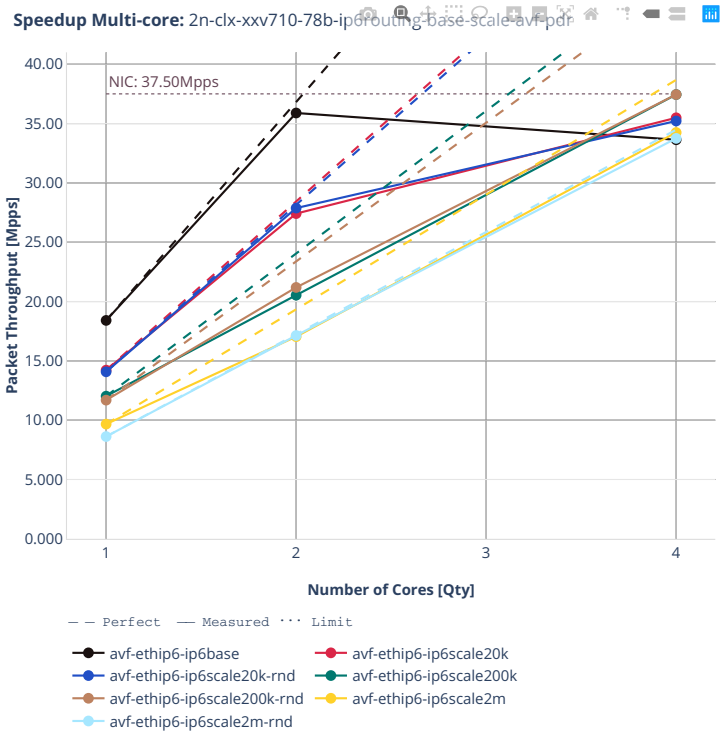
Speedup Multi-core: 2n-clx-xxv710-78b-ip6routing-base-scale-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip6-ip6base
avf-ethip6-ip6scale20k
avf-ethip6-ip6scale20k-rnd
avf-ethip6-ip6scale200k
avf-ethip6-ip6scale200k-rnd
avf-ethip6-ip6scale2m
avf-ethip6-ip6scale2m-rnd

78b-ip6routing-base-scale-af_xdp

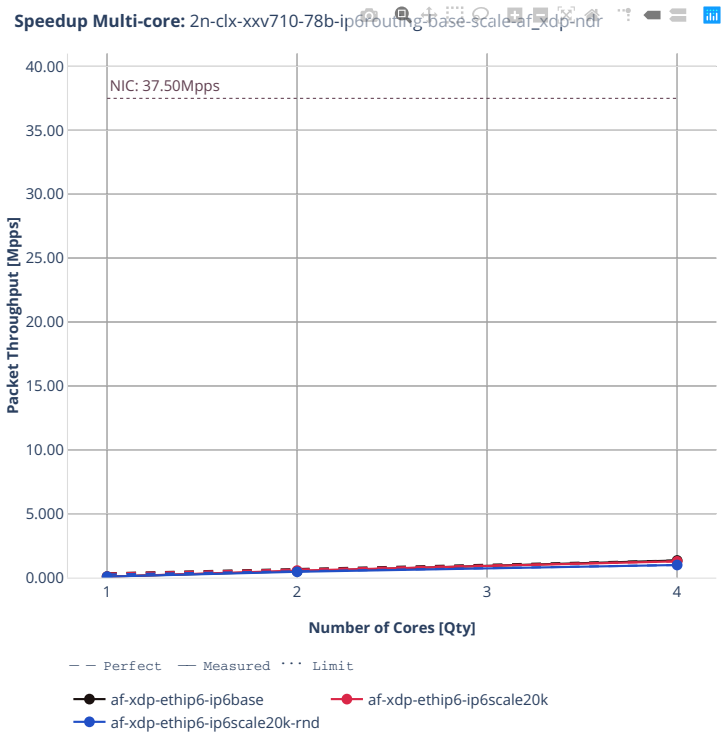

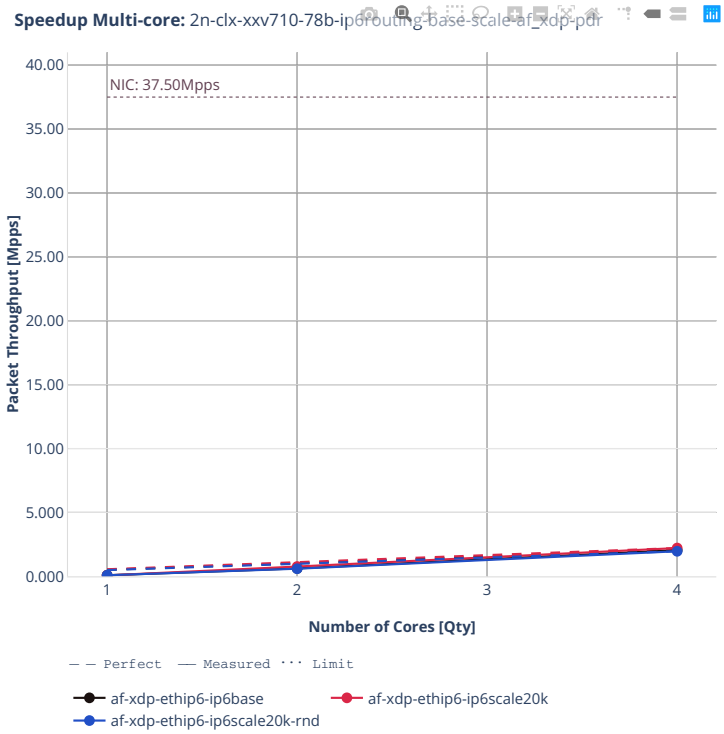
Speedup Multi-core: 2n-clx-xxv710-78b-ip6routing-base-scale-af_xdp-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
af-xdp-ethip6-ip6base
af-xdp-ethip6-ip6scale20k
af-xdp-ethip6-ip6scale20k-rnd

78b-ip6routing-base-scale-dpdk

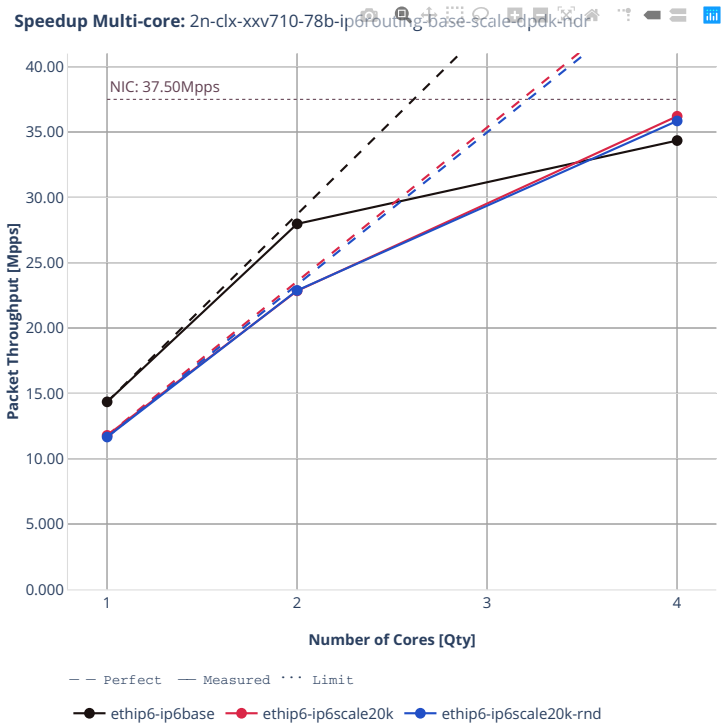

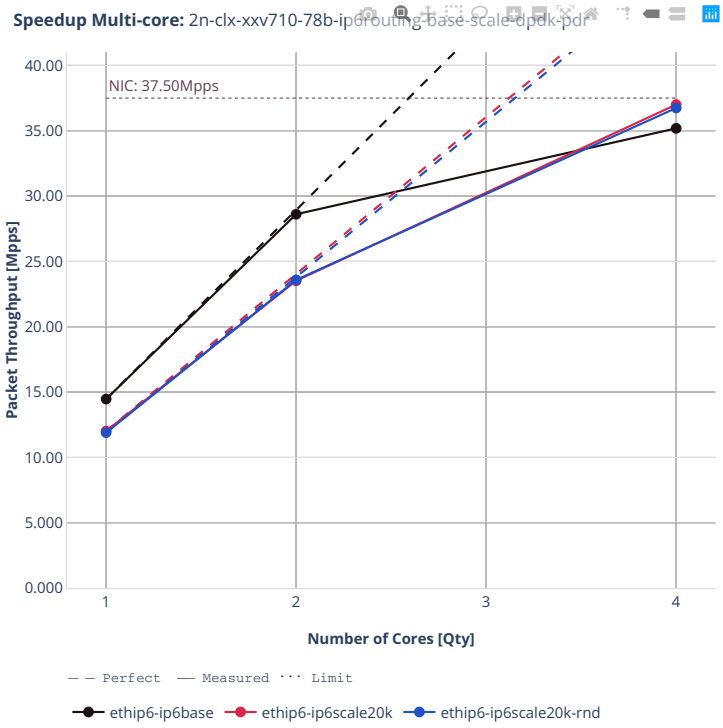
Speedup Multi-core: 2n-clx-xxv710-78b-ip6routing-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip6-ip6base
ethip6-ip6scale20k
ethip6-ip6scale20k-rnd

2n-clx-x710

78b-ip6routing-base-scale-[avf,dpdk,af_xdp]

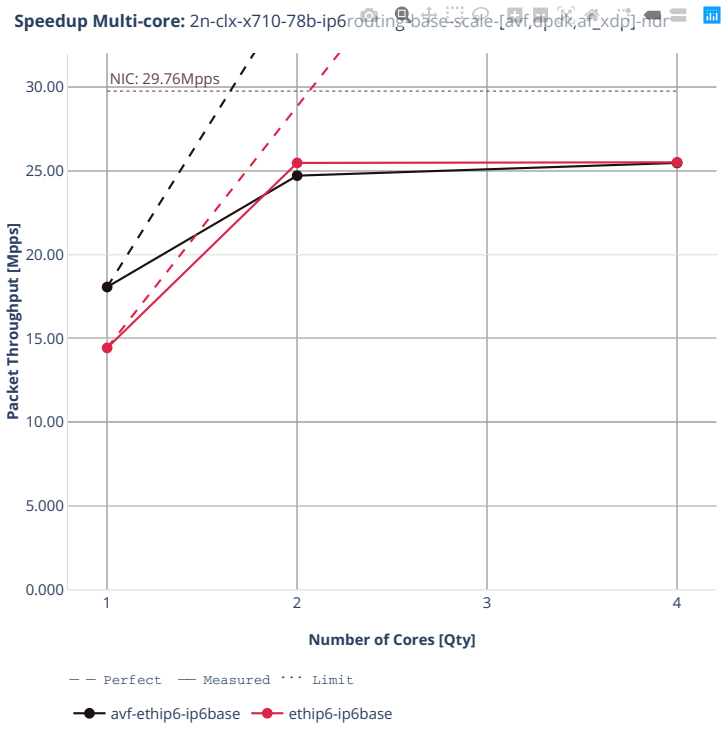

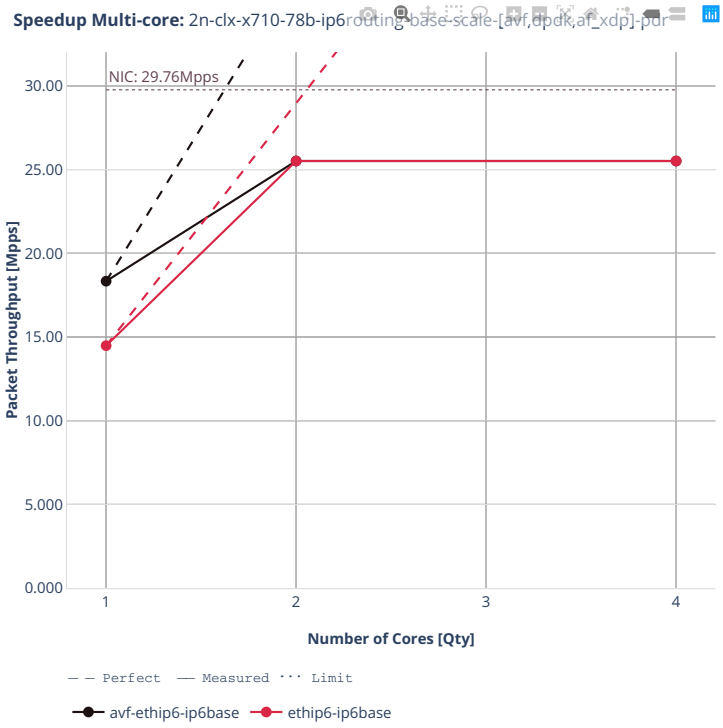
Speedup Multi-core: 2n-clx-x710-78b-ip6routing-base-scale-[avf,dpdk,af_xdp]-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip6-ip6base
ethip6-ip6base

2n-clx-cx556a

78b-ip6routing-base-scale-rdma-core

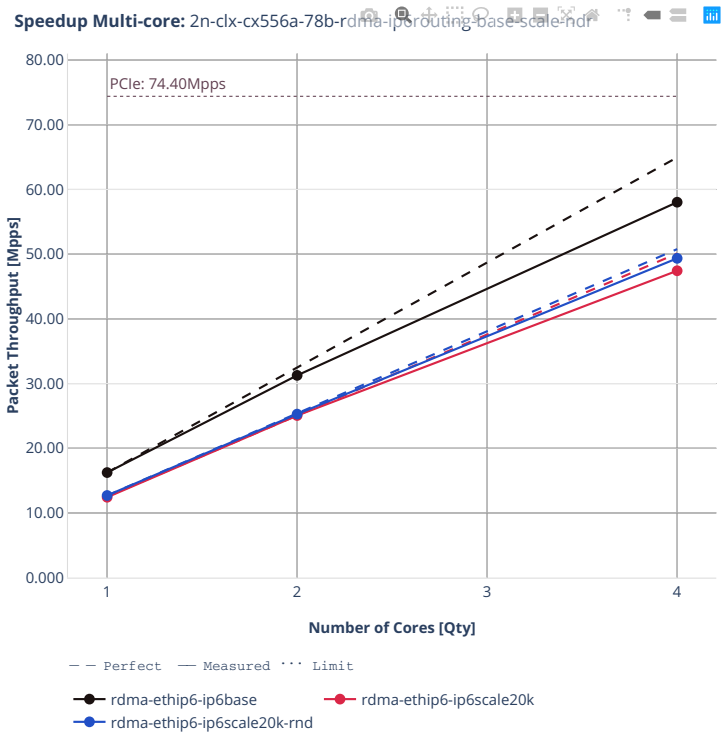

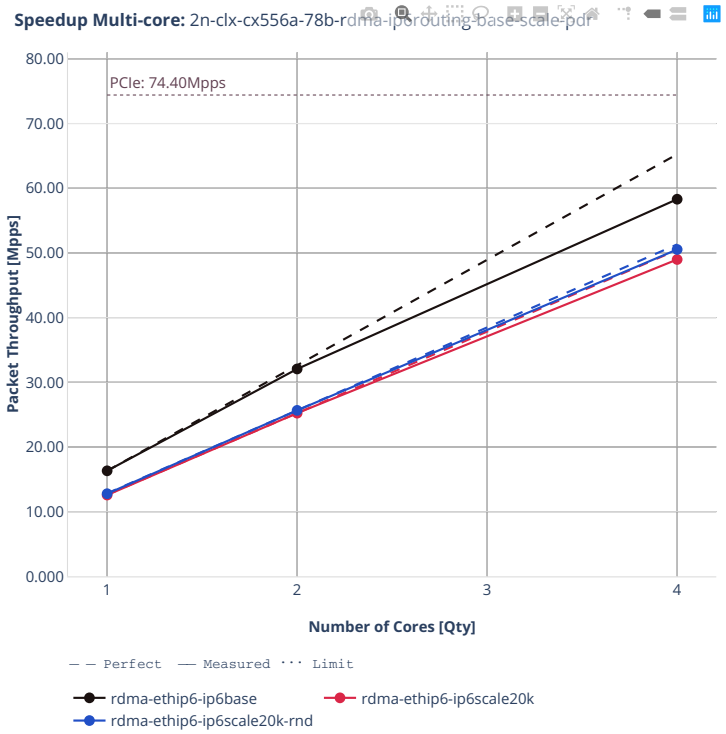
Speedup Multi-core: 2n-clx-cx556a-78b-rdma-ip6routing-base-scale-pdr
PCIe: 74.40Mpps
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-ethip6-ip6base
rdma-ethip6-ip6scale20k
rdma-ethip6-ip6scale20k-rnd

2n-clx-e810cq

78b-ip6routing-base-scale

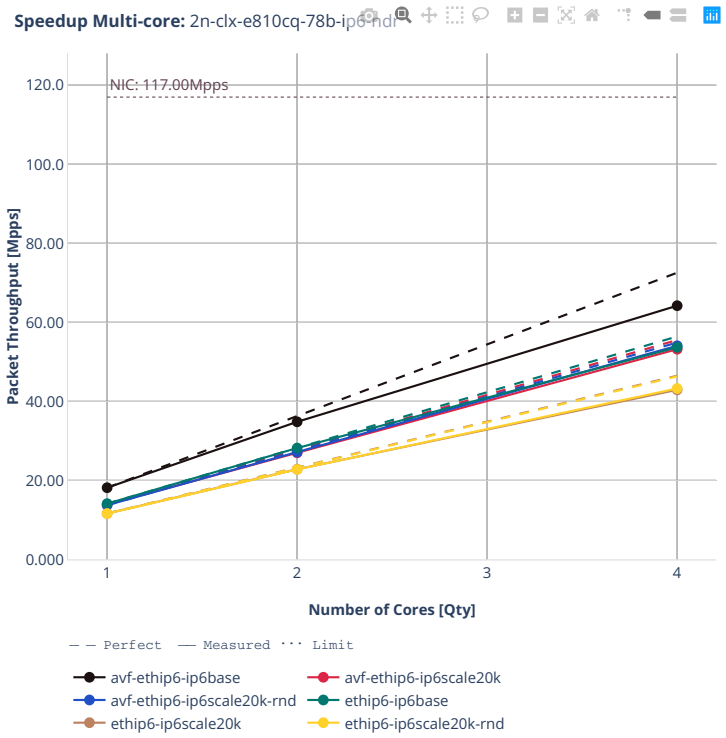

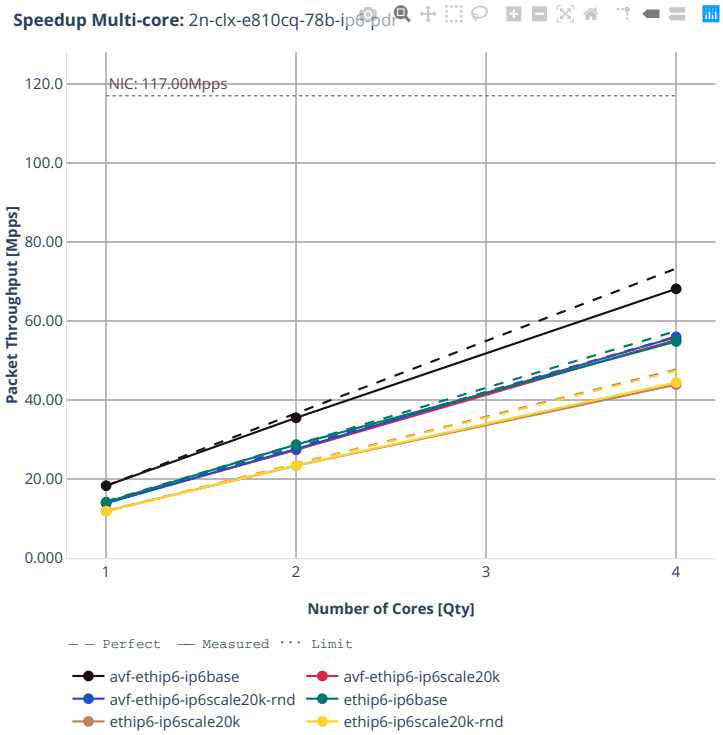
Speedup Multi-core: 2n-clx-e810cq-78b-ip6-pdr
NIC: 117.00Mpps
120.0
100.0
80.00
60.00
40.00
20.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip6-ip6base
avf-ethip6-ip6scale20k
avf-ethip6-ip6scale20k-rnd
ethip6-ip6base
ethip6-ip6scale20k
ethip6-ip6scale20k-rnd

2n-zn2-xxv710

78b-ip6routing-base-scale-avf

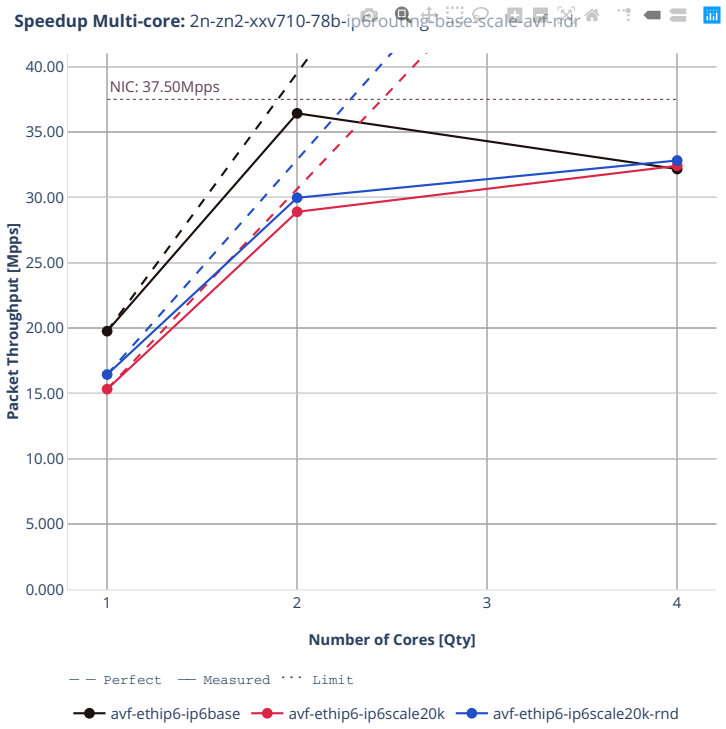

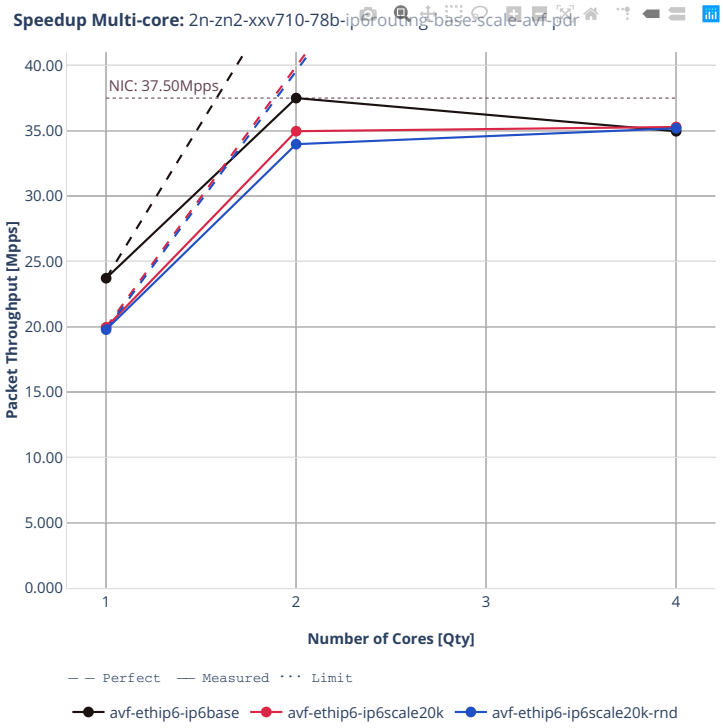
Speedup Multi-core: 2n-zn2-xxv710-78b-ip6routing-base-scale-avf-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip6-ip6base
avf-ethip6-ip6scale20k
avf-ethip6-ip6scale20k-rnd

78b-ip6routing-base-scale-dpdk


Speedup Multi-core: 2n-zn2-xxv710-78b-ip6routing-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip6-ip6base
ethip6-ip6scale20k
ethip6-ip6scale20k-rnd

2n-zn2-x710

78b-ip6routing-base-[avf,dpdk]

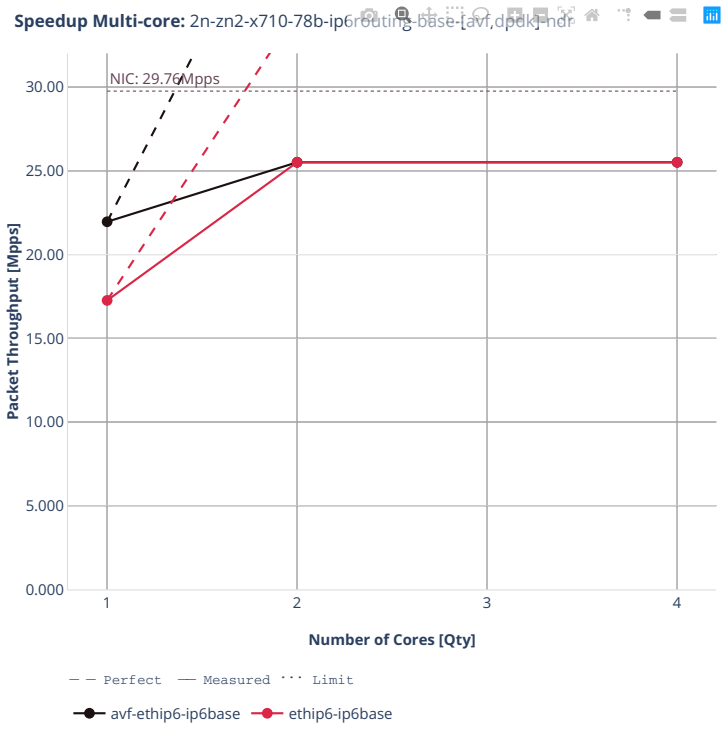

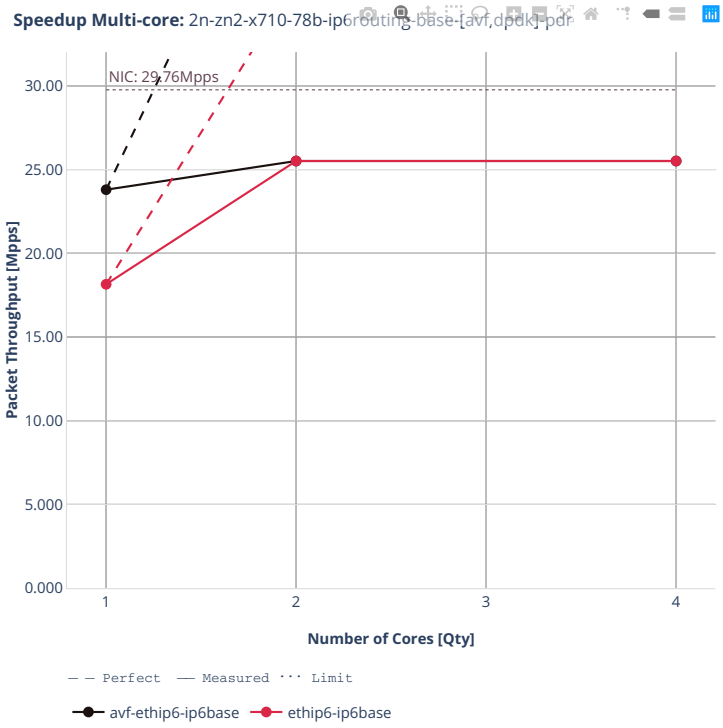
Speedup Multi-core: 2n-zn2-x710-78b-ip6routing-base-[avf,dpdk]-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip6-ip6base
ethip6-ip6base

2n-zn2-cx556a

78b-ip6routing-base-scale-rdma-core

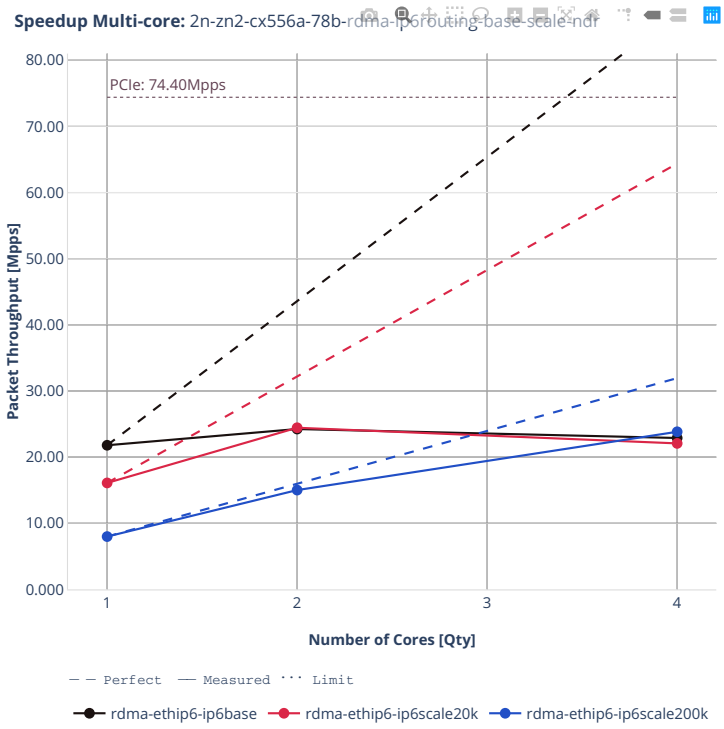

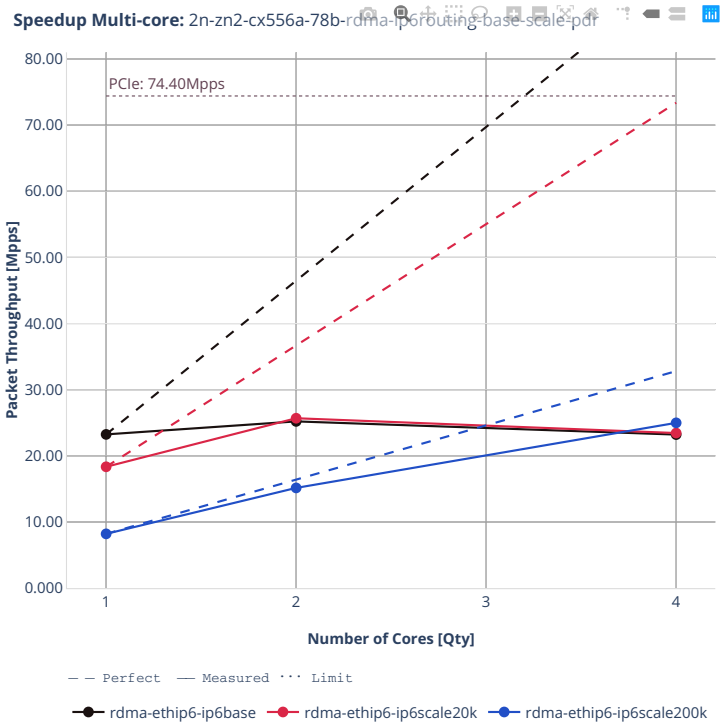
Speedup Multi-core: 2n-zn2-cx556a-78b-rdma-ip6routing-base-scale-pdr
PCIe: 74.40Mpps
Packet Throughput [Mpps]
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-ethip6-ip6base
rdma-ethip6-ip6scale20k
rdma-ethip6-ip6scale200k

3n-alt-xl710

78b-ip6routing-base-scale

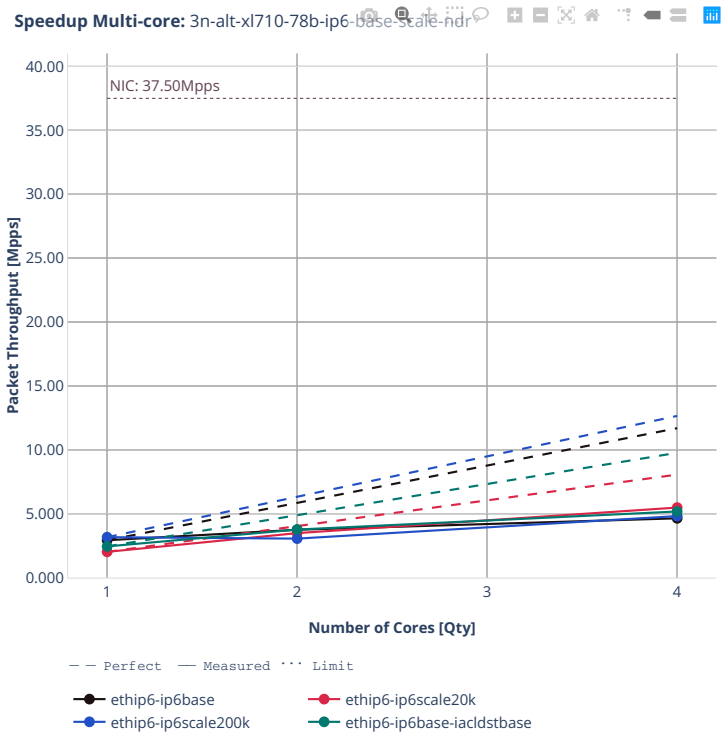

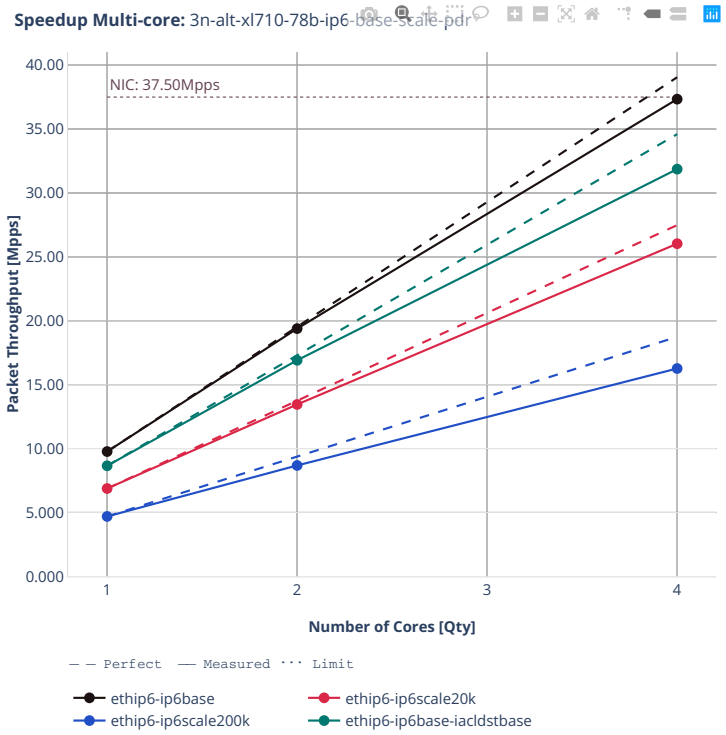
Speedup Multi-core: 3n-alt-xl710-78b-ip6-base-scale-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip6-ip6base
ethip6-ip6scale20k
ethip6-ip6scale200k
ethip6-ip6base-iacldstbase

3n-tsh-x520

78b-ip6routing-base-ixgbe

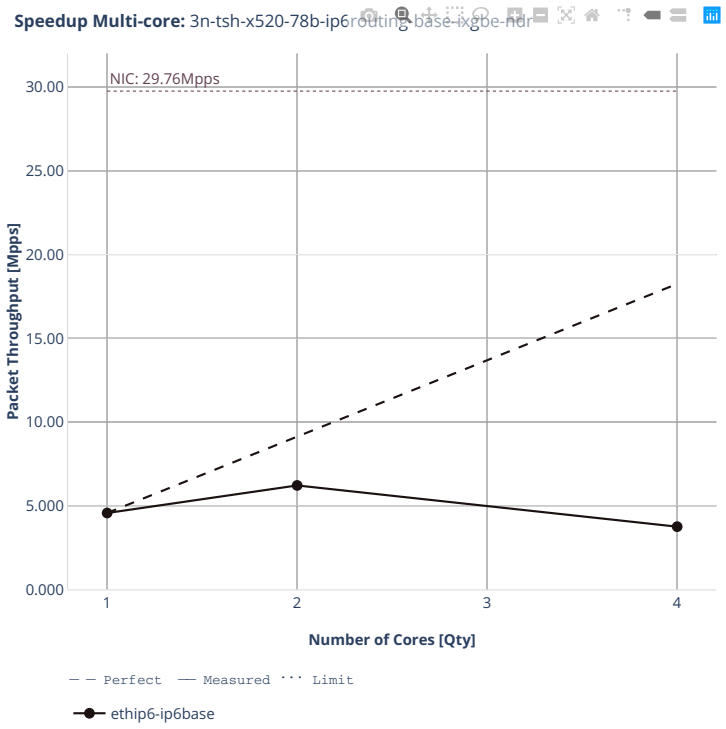

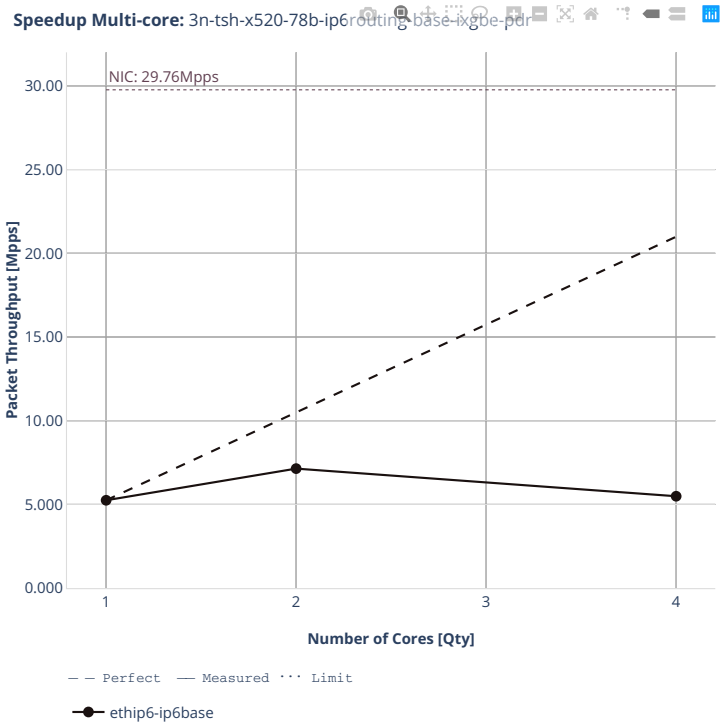
Speedup Multi-core: 3n-tsh-x520-78b-ip6routing-base-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip6-ip6base

2n-tx2-xl710

78b-ip6routing-base-scale-dpdk

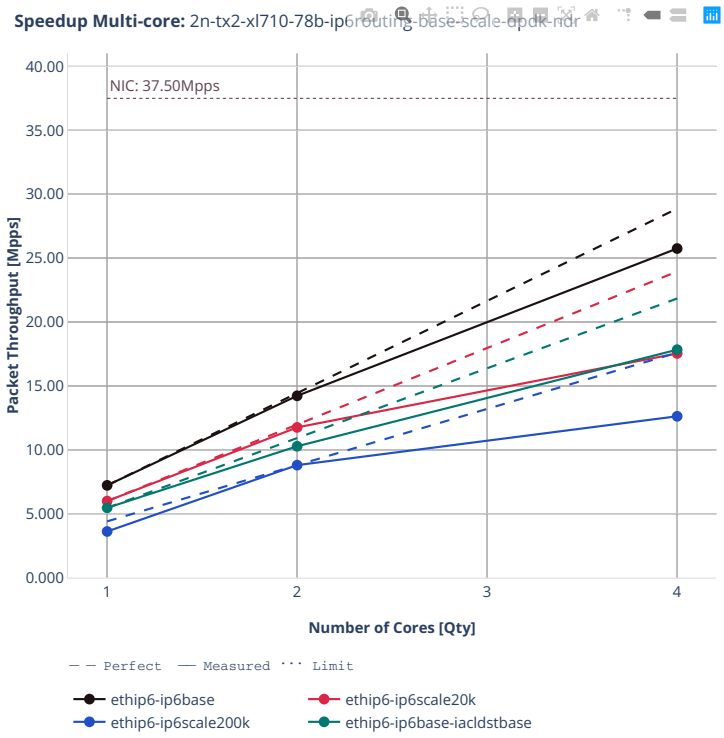

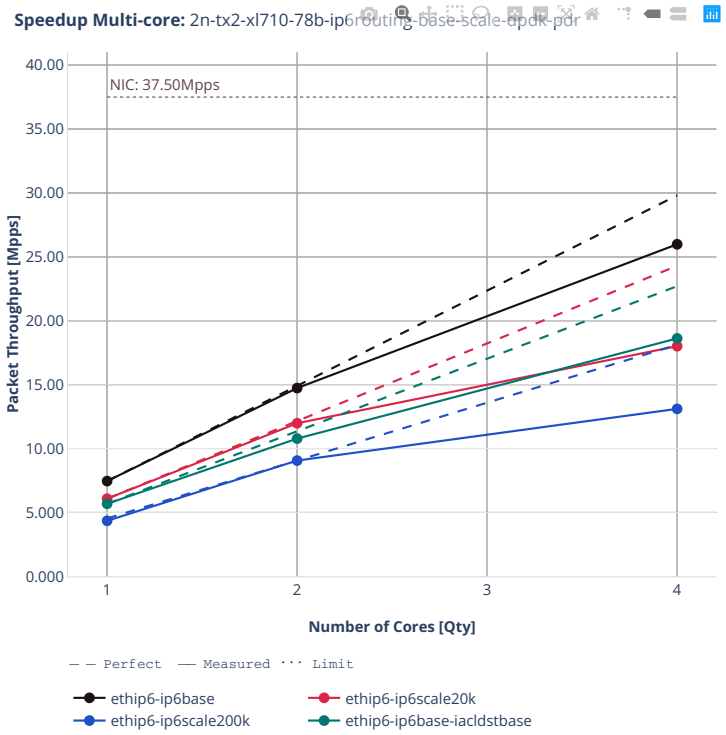
Speedup Multi-core: 2n-tx2-xl710-78b-ip6routing-base-scale-dpdk-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip6-ip6base
ethip6-ip6scale20k
ethip6-ip6scale200k
ethip6-ip6base-iacldstbase

78b-ip6routing-base-scale-af-xdp

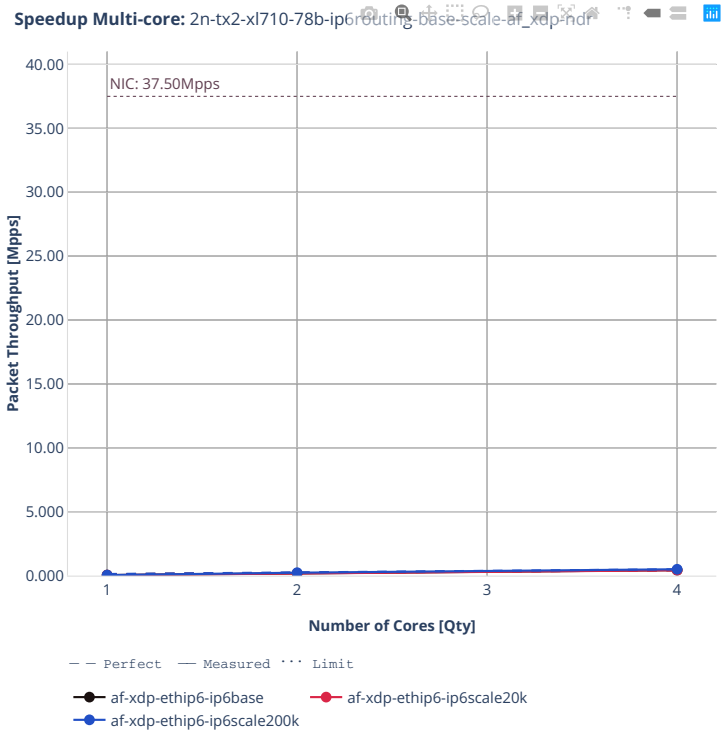

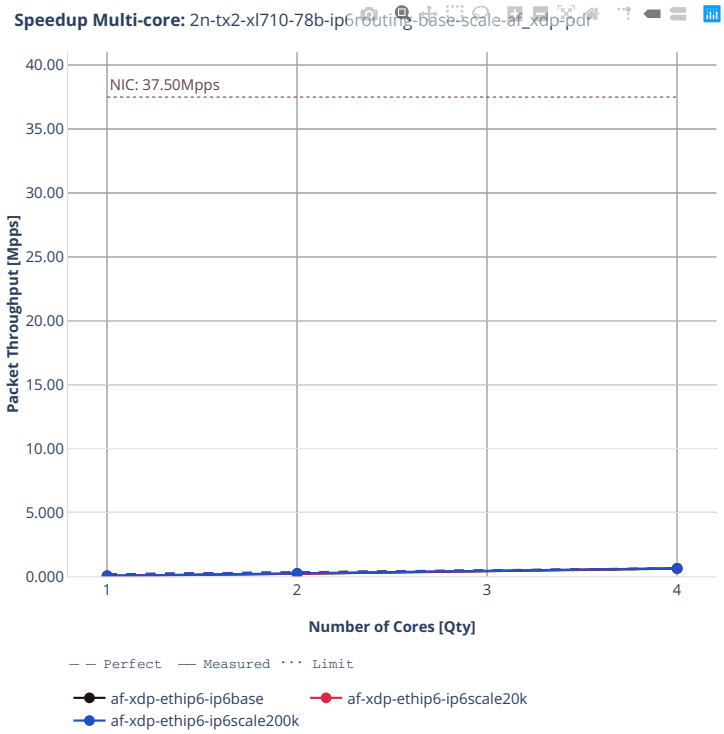
Speedup Multi-core: 2n-tx2-xl710-78b-ip6routing-base-scale-af_xdp-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
af-xdp-ethip6-ip6base
af-xdp-ethip6-ip6scale20k
af-xdp-ethip6-ip6scale200k

2n-dnv-x553

78b-ip6routing-base-scale-ixgbe

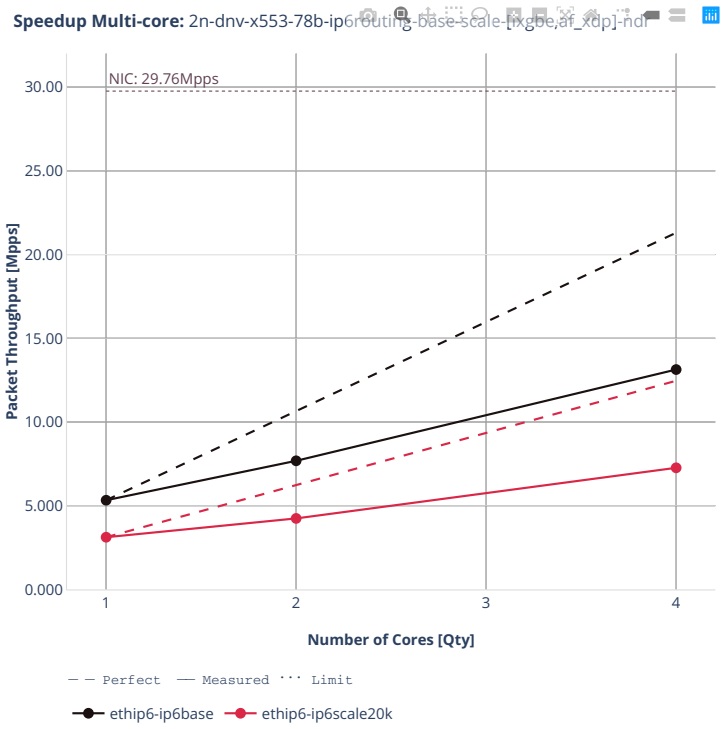

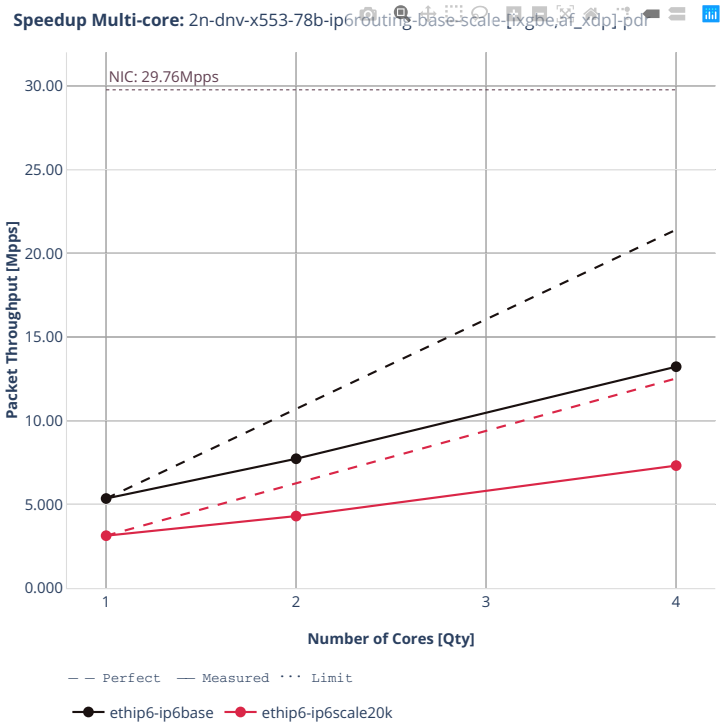
Speedup Multi-core: 2n-dnv-x553-78b-ip6routing-base-scale-[ixgbe,af_xdp]-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip6-ip6base
ethip6-ip6scale20k

3n-dnv-x553

78b-ip6routing-base-scale-ixgbe

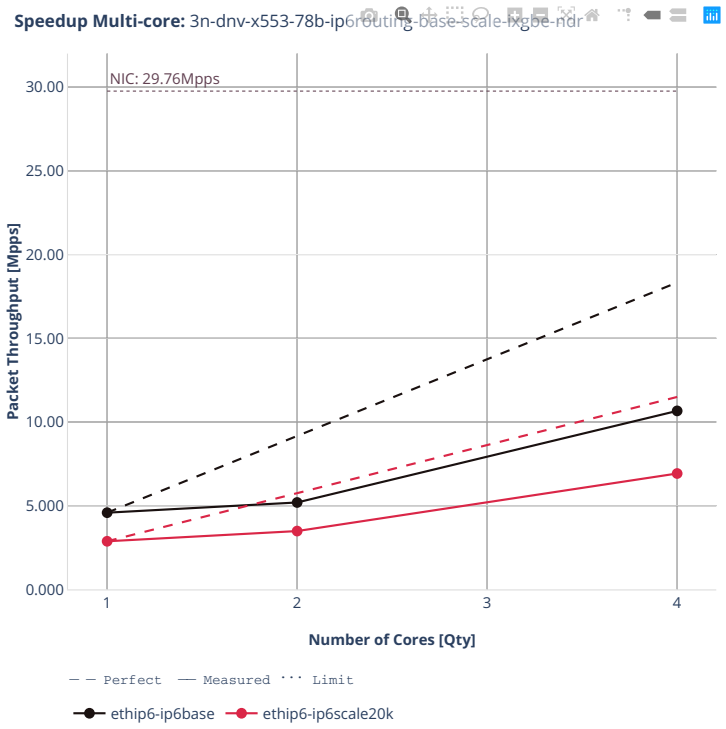

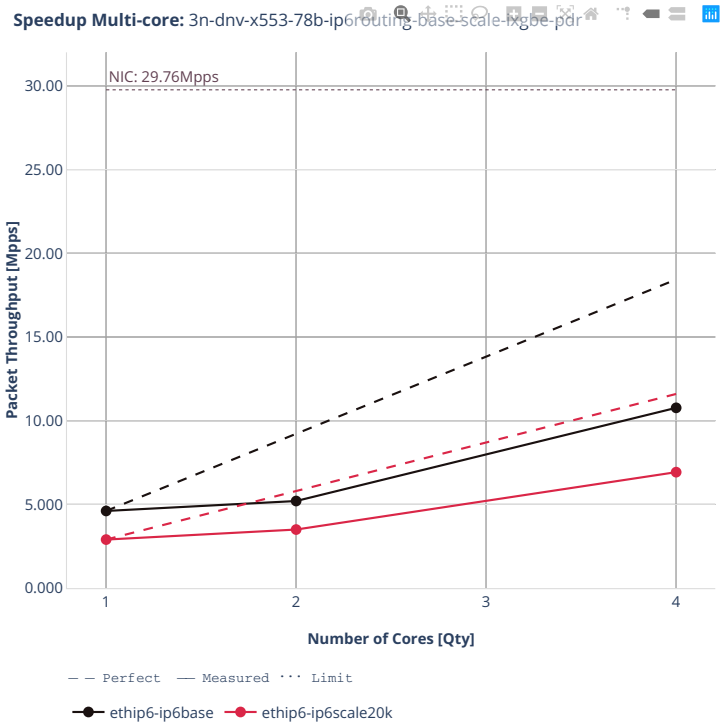
Speedup Multi-core: 3n-dnv-x553-78b-ip6routing-base-scale-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip6-ip6base
ethip6-ip6scale20k

2.4.4 SRv6 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 SRv6, 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[133].

[133] https://git.fd.io/csit/tree/tests/vpp/perf/srv6?h=rls2206

3n-icx-xxv710

78b-srv6-ip6routing-base-avf

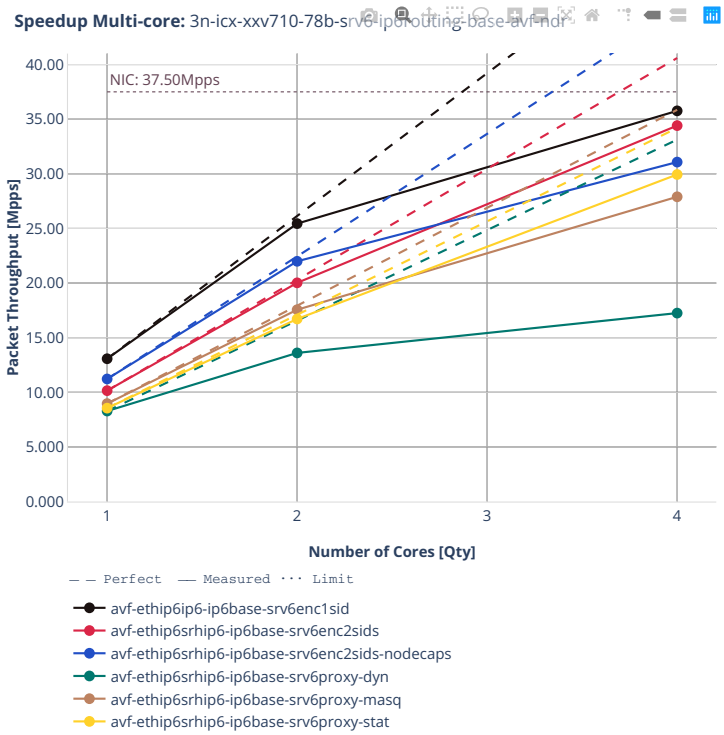

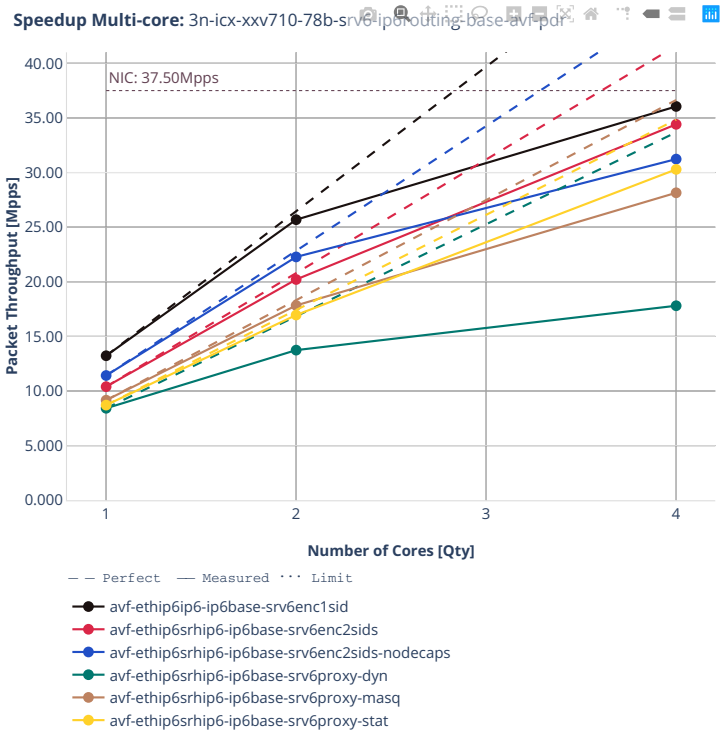
Speedup Multi-core: 3n-icx-xxv710-78b-srv6-ipv6routing-base-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip6ip6-ip6base-srv6enc1sid
avf-ethip6srhip6-ip6base-srv6enc2sids
avf-ethip6srhip6-ip6base-srv6enc2sids-nodecaps
avf-ethip6srhip6-ip6base-srv6proxy-dyn
avf-ethip6srhip6-ip6base-srv6proxy-masq
avf-ethip6srhip6-ip6base-srv6proxy-stat

3n-skx-xxv710

78b-srv6-ip6routing-base-avf

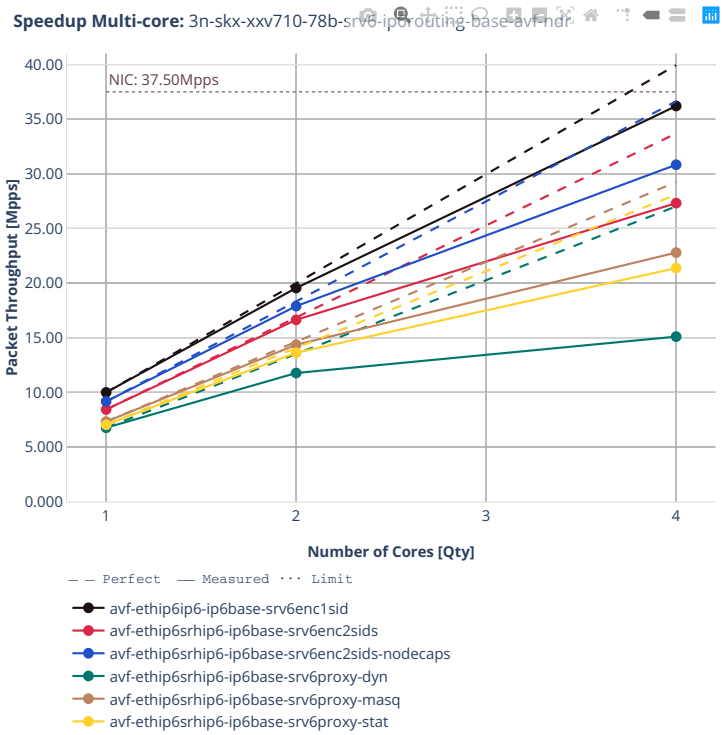

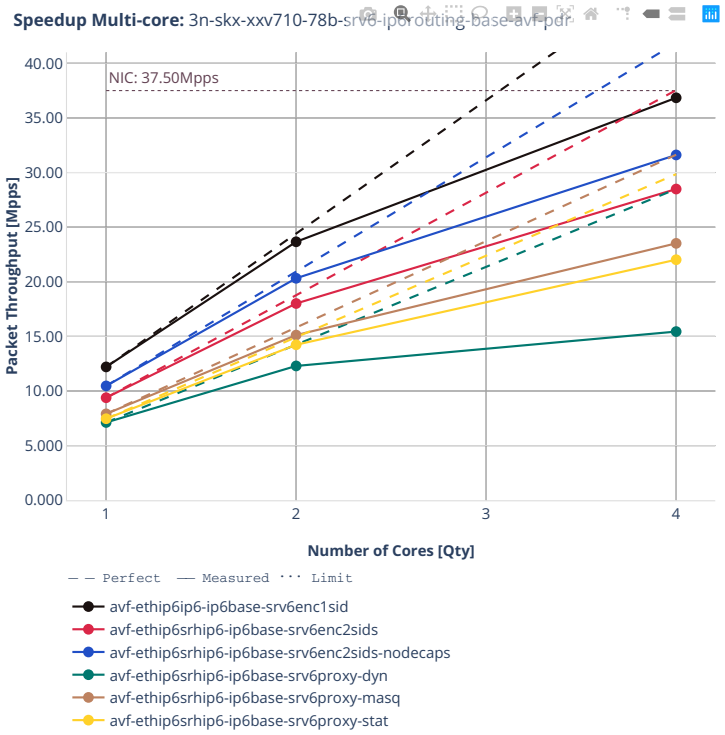
Speedup Multi-core: 3n-skx-xxv710-78b-srv6-ip6routing-base-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip6ip6-ip6base-srv6enc1sid
avf-ethip6srhip6-ip6base-srv6enc2sids
avf-ethip6srhip6-ip6base-srv6enc2sids-nodecaps
avf-ethip6srhip6-ip6base-srv6proxy-dyn
avf-ethip6srhip6-ip6base-srv6proxy-masq
avf-ethip6srhip6-ip6base-srv6proxy-stat

3n-tsh-x520

78b-srv6-ip6routing-base-ixgbe

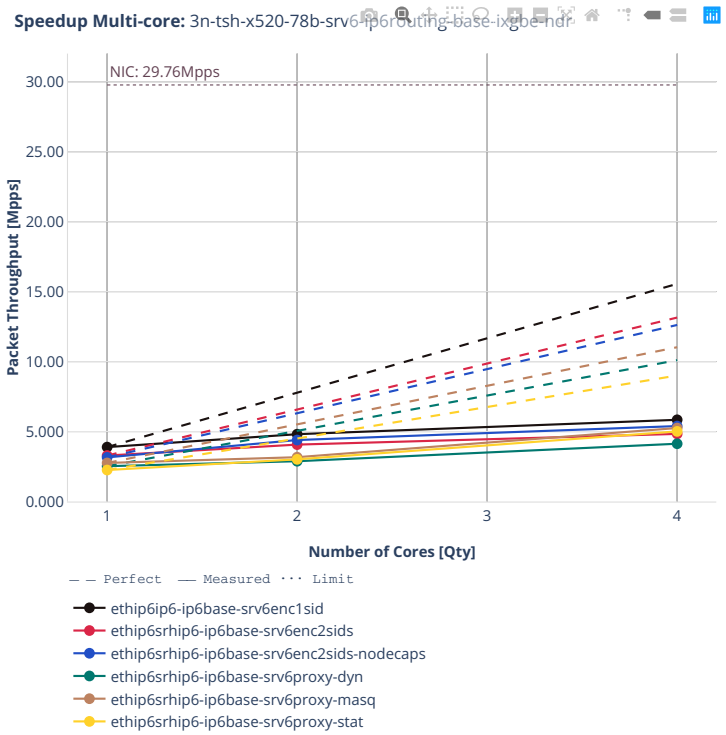

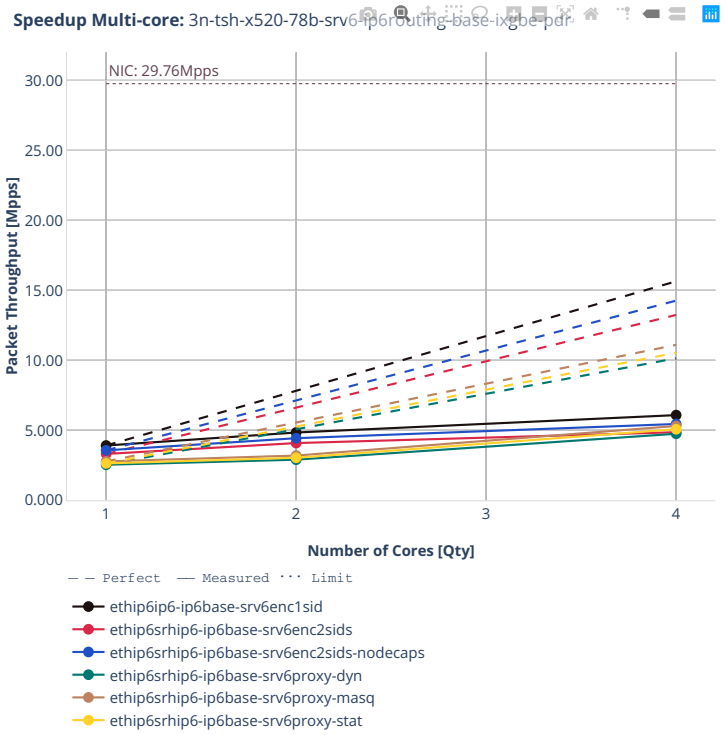
Speedup Multi-core: 3n-tsh-x520-78b-srv6-ip6routing-base-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip6ip6-ip6base-srv6enc1sid
ethip6srhip6-ip6base-srv6enc2sids
ethip6srhip6-ip6base-srv6enc2sids-nodecaps
ethip6srhip6-ip6base-srv6proxy-dyn
ethip6srhip6-ip6base-srv6proxy-masq
ethip6srhip6-ip6base-srv6proxy-stat

2.4.5 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[134].

[134] https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls2206

2n-icx-xxv710

64b-2t1c-ethip4–ethip4udpgeneve-avf

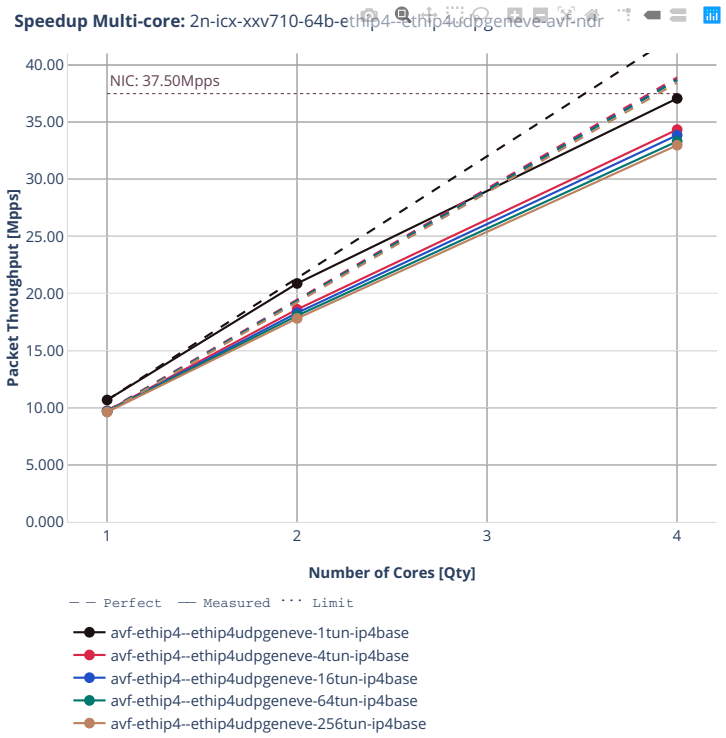

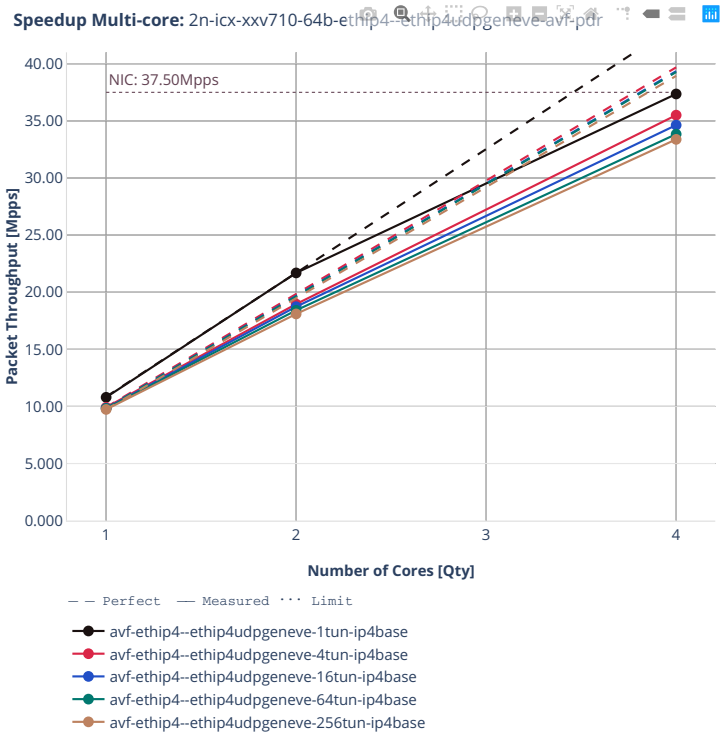
Speedup Multi-core: 2n-icx-xxv710-64b-ethip4-ethip4udpgeneve-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4--ethip4udpgeneve-1tun-ip4base
avf-ethip4--ethip4udpgeneve-4tun-ip4base
avf-ethip4--ethip4udpgeneve-16tun-ip4base
avf-ethip4--ethip4udpgeneve-64tun-ip4base
avf-ethip4--ethip4udpgeneve-256tun-ip4base

3n-icx-xxv710

64b-ip4tunnel-base

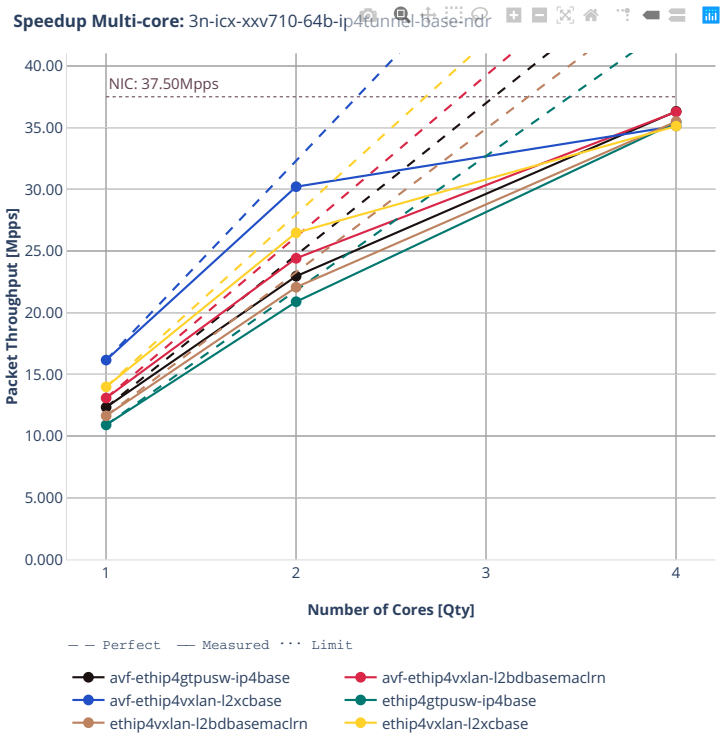

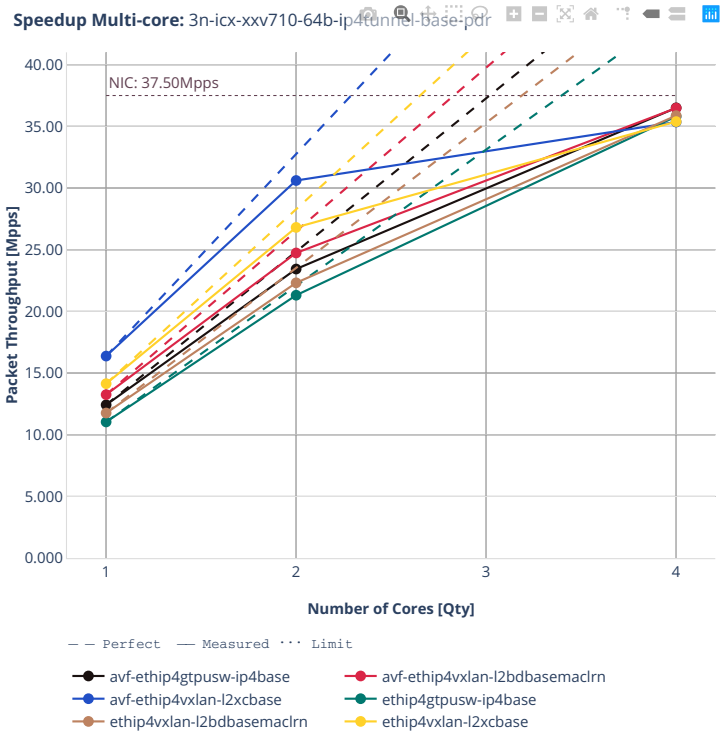
Speedup Multi-core: 3n-icx-xxv710-64b-ip4tunnel-base-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4gtpusw-ip4base
avf-ethip4vxlan-l2bdbasemaclrn
avf-ethip4vxlan-l2xcbase
ethip4gtpusw-ip4base
ethip4vxlan-l2bdbasemaclrn
ethip4vxlan-l2xcbase

3n-icx-e810xxv

64b-ip4tunnel-gtpu

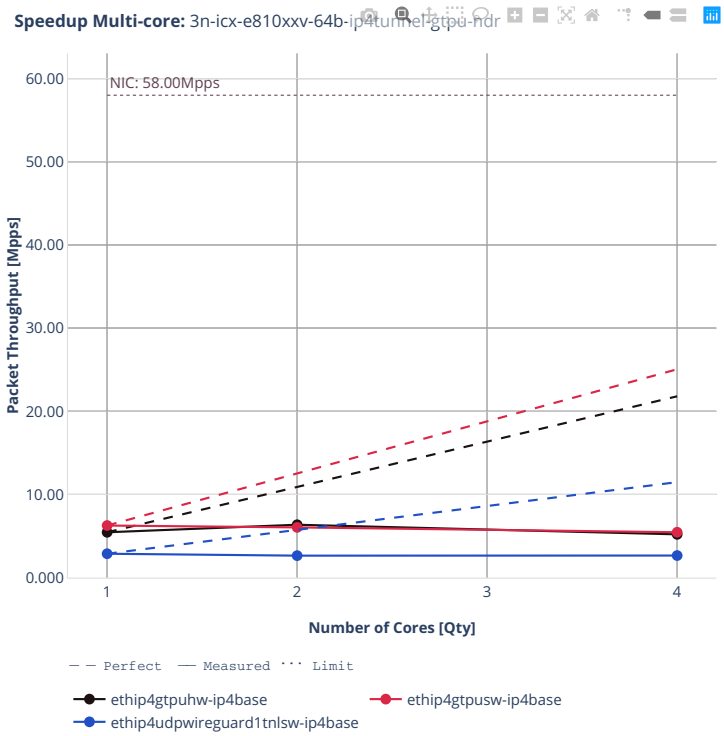

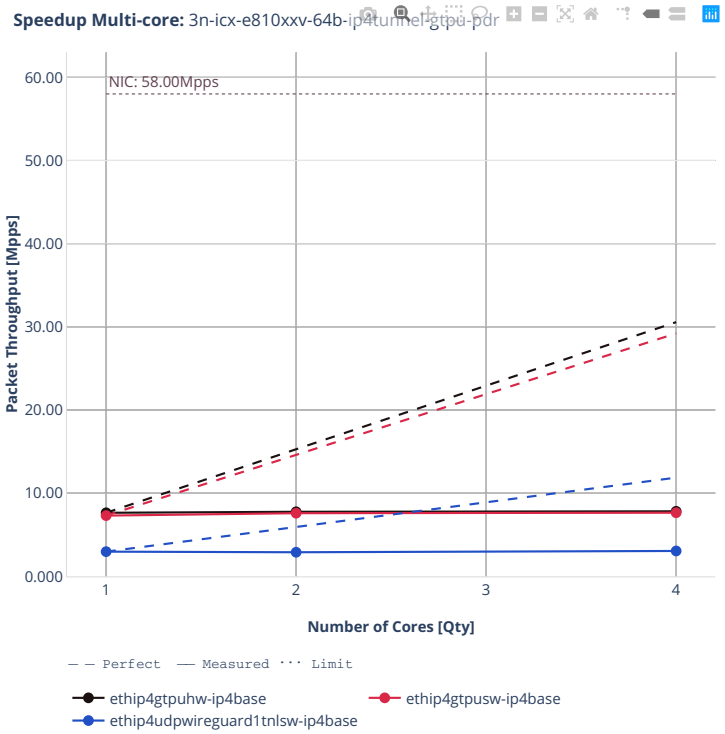
Speedup Multi-core: 3n-icx-e810xxv-64b-ip4tunnel-gtpu-pdr
NIC: 58.00Mpps
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4gtpuhw-ip4base
ethip4gtpusw-ip4base
ethip4udpwireguard1tnlsw-ip4base

2n-skx-xxv710

64b-2t1c-ethip4–ethip4udpgeneve-avf

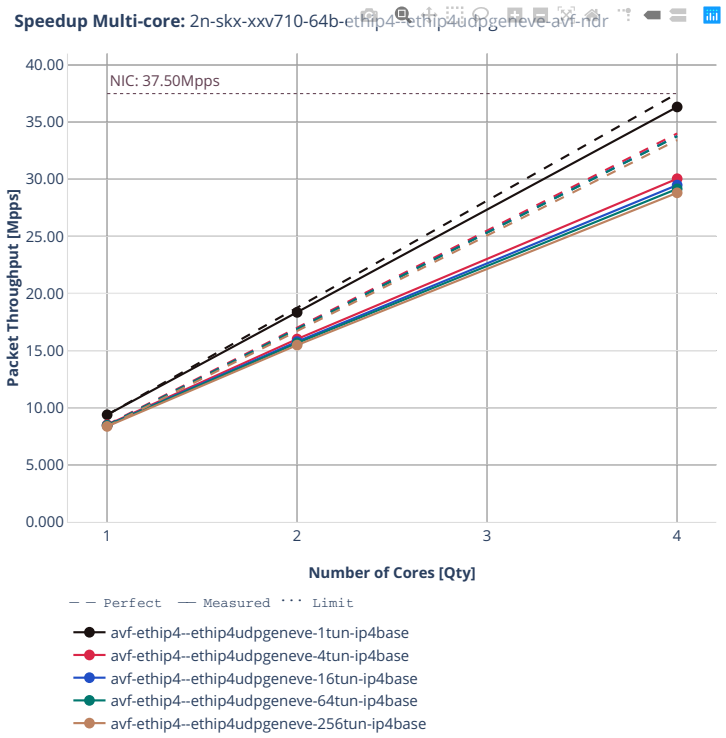

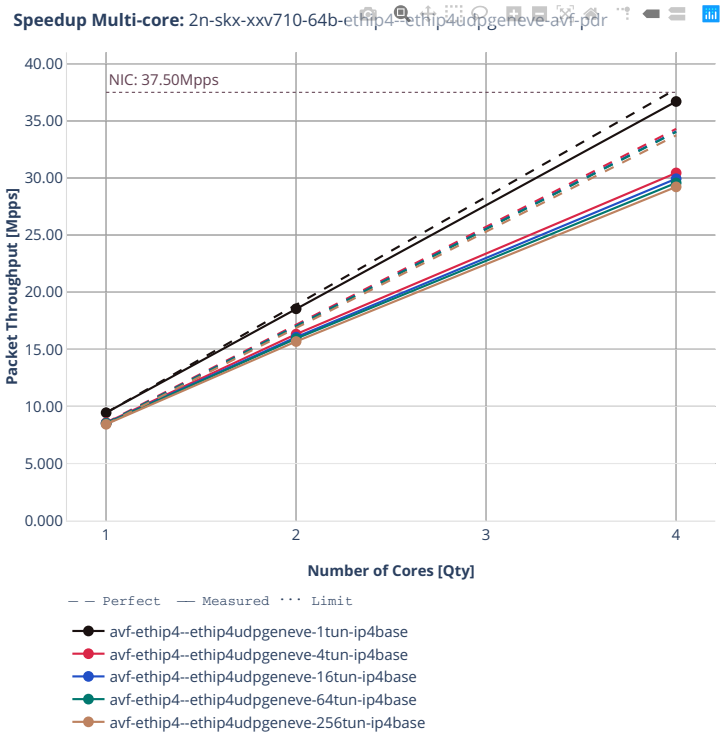
Speedup Multi-core: 2n-skx-xxv710-64b-ethip4-ethip4udpgeneve-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4--ethip4udpgeneve-1tun-ip4base
avf-ethip4--ethip4udpgeneve-4tun-ip4base
avf-ethip4--ethip4udpgeneve-16tun-ip4base
avf-ethip4--ethip4udpgeneve-64tun-ip4base
avf-ethip4--ethip4udpgeneve-256tun-ip4base

2n-clx-xxv710

64b-2t1c-ethip4–ethip4udpgeneve-avf

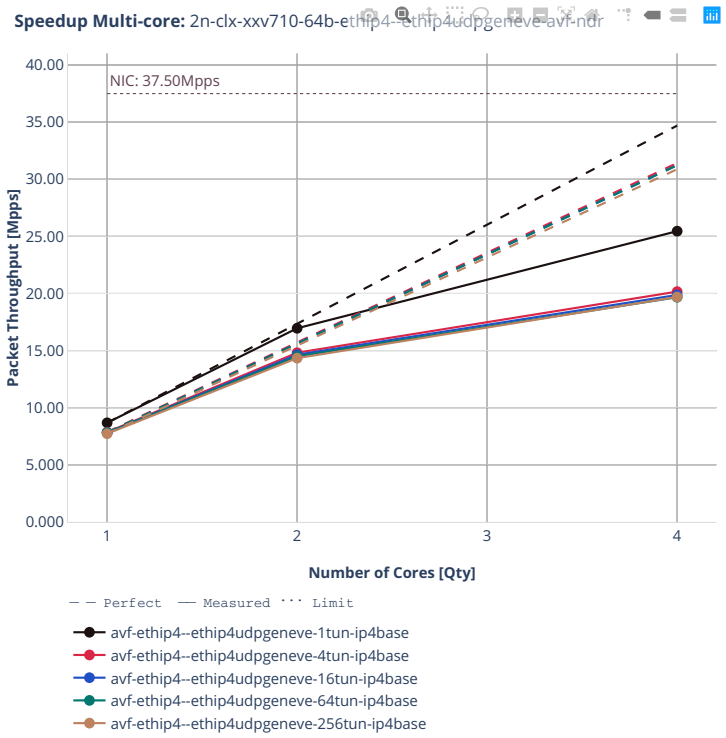

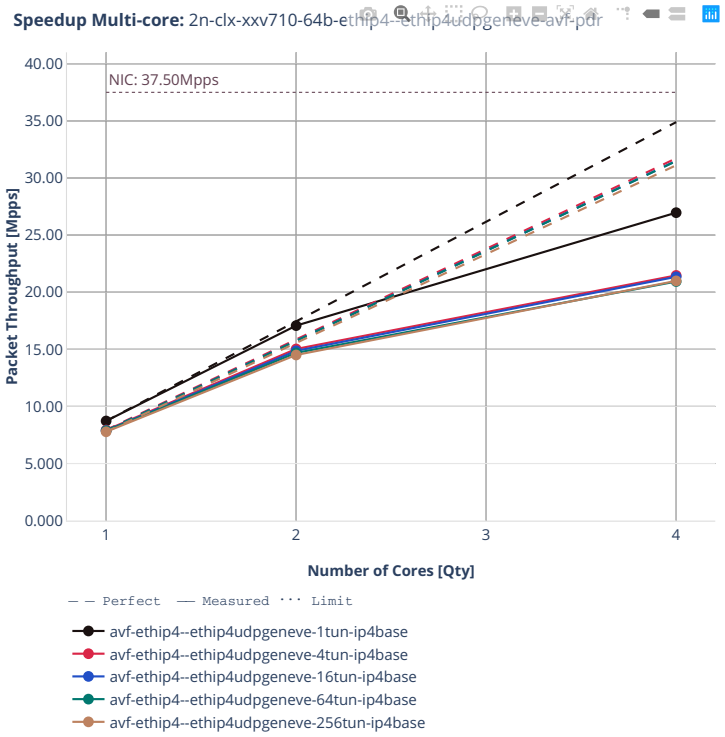
Speedup Multi-core: 2n-clx-xxv710-64b-ethip4-ethip4udpgeneve-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4--ethip4udpgeneve-1tun-ip4base
avf-ethip4--ethip4udpgeneve-4tun-ip4base
avf-ethip4--ethip4udpgeneve-16tun-ip4base
avf-ethip4--ethip4udpgeneve-64tun-ip4base
avf-ethip4--ethip4udpgeneve-256tun-ip4base

2n-zn2-xxv710

64b-2t1c-ethip4–ethip4udpgeneve-avf

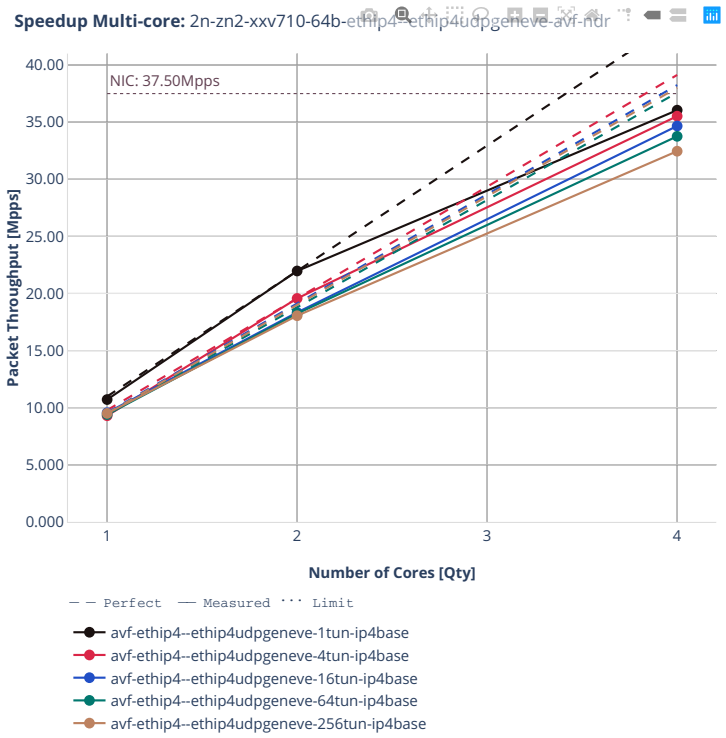

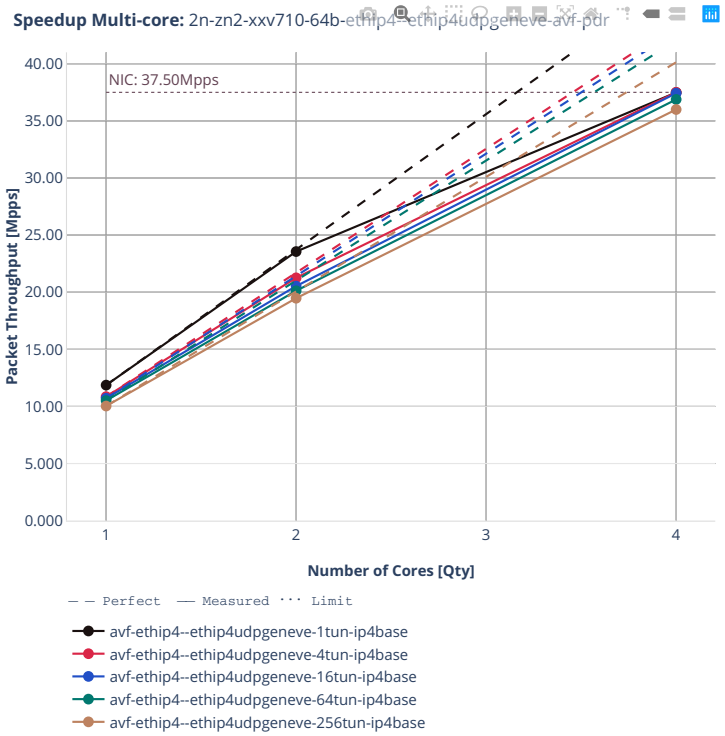
Speedup Multi-core: 2n-zn2-xxv710-64b-ethip4-ethip4udpgeneve-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4--ethip4udpgeneve-1tun-ip4base
avf-ethip4--ethip4udpgeneve-4tun-ip4base
avf-ethip4--ethip4udpgeneve-16tun-ip4base
avf-ethip4--ethip4udpgeneve-64tun-ip4base
avf-ethip4--ethip4udpgeneve-256tun-ip4base

3n-skx-xxv710

64b-ip4tunnel-base

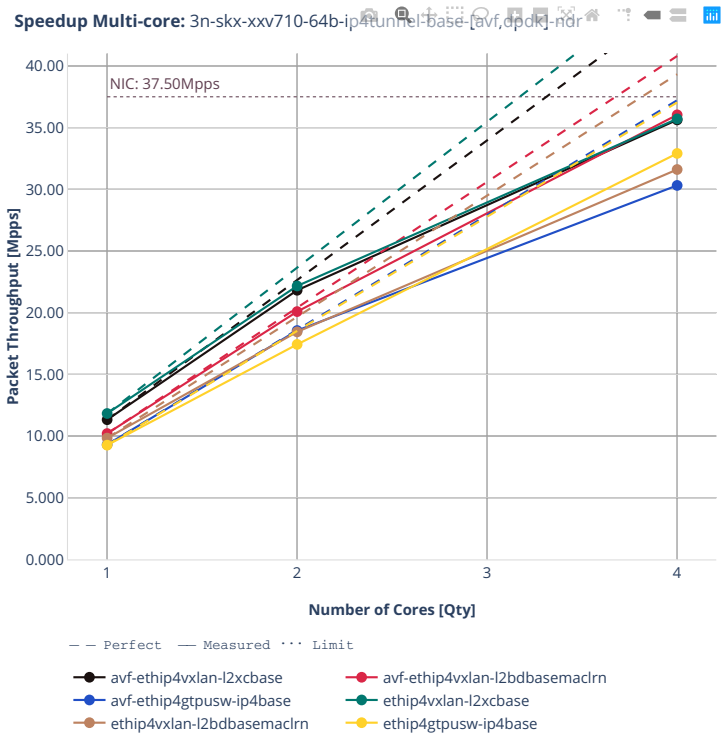

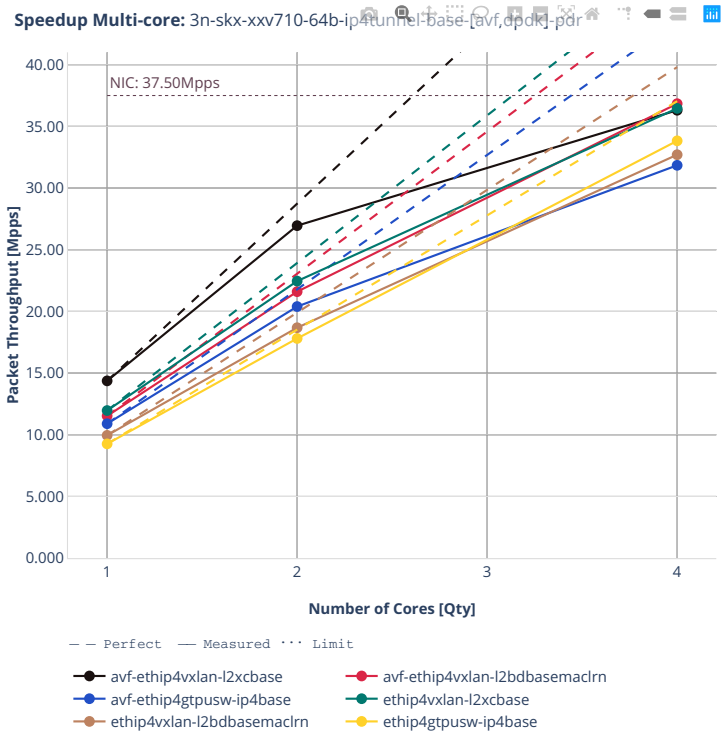
Speedup Multi-core: 3n-skx-xxv710-64b-ip4tunnel-base-[avf,dpdk]-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4vxlan-l2xcbase
avf-ethip4vxlan-l2bdbasemaclrn
avf-ethip4gtpusw-ip4base
ethip4vxlan-l2xcbase
ethip4vxlan-l2bdbasemaclrn
ethip4gtpusw-ip4base

3n-alt-xl710

64b-ip4tunnel-base

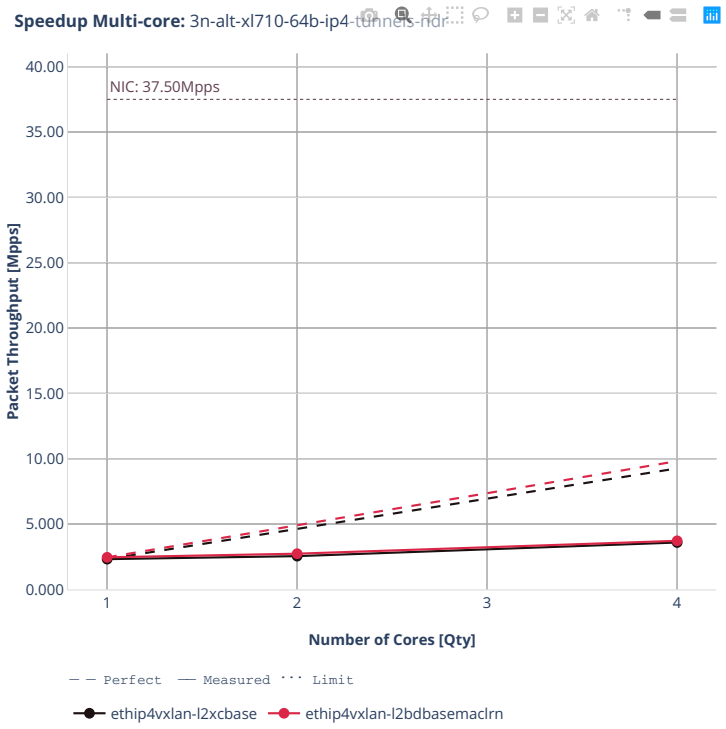

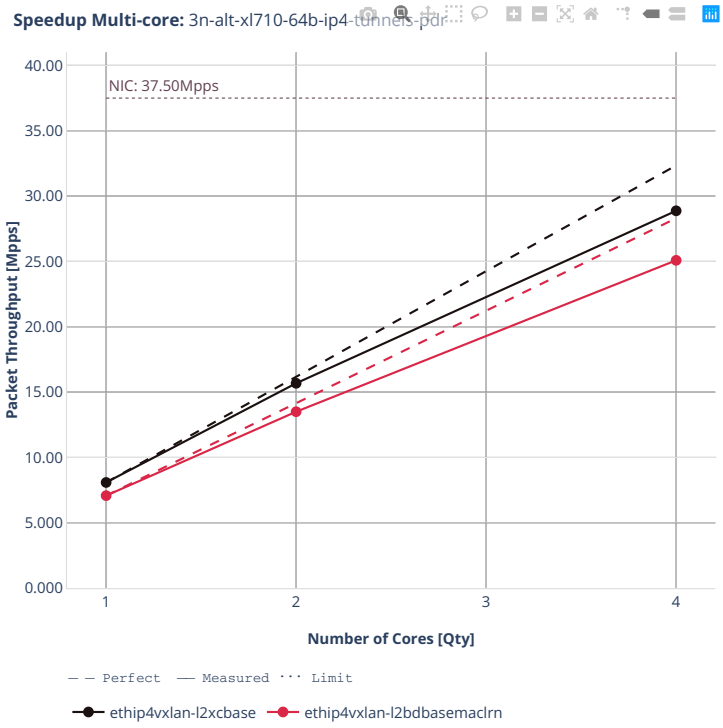
Speedup Multi-core: 3n-alt-xl710-64b-ip4-tunnels-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4vxlan-l2xcbase
ethip4vxlan-l2bdbasemaclrn

3n-tsh-x520

64b-ip4tunnel-base-ixgbe

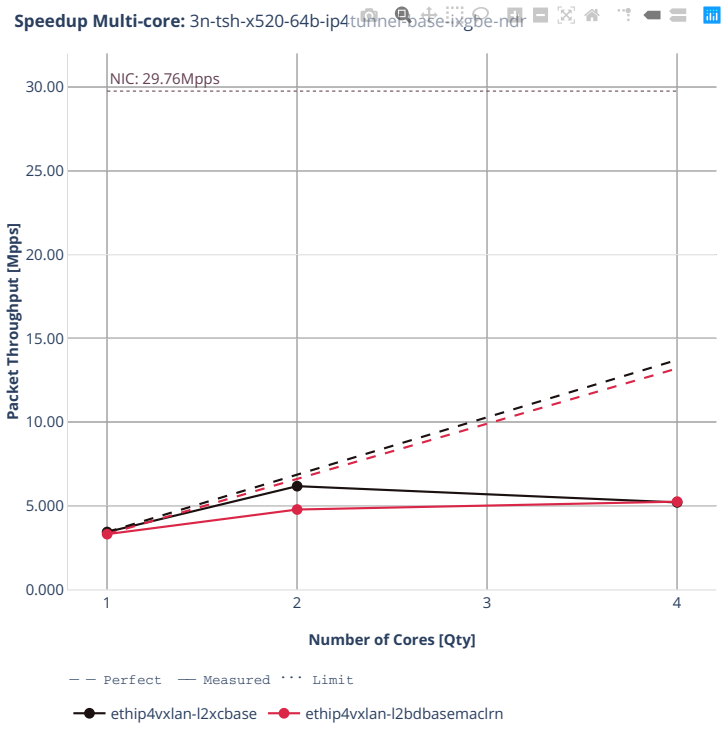

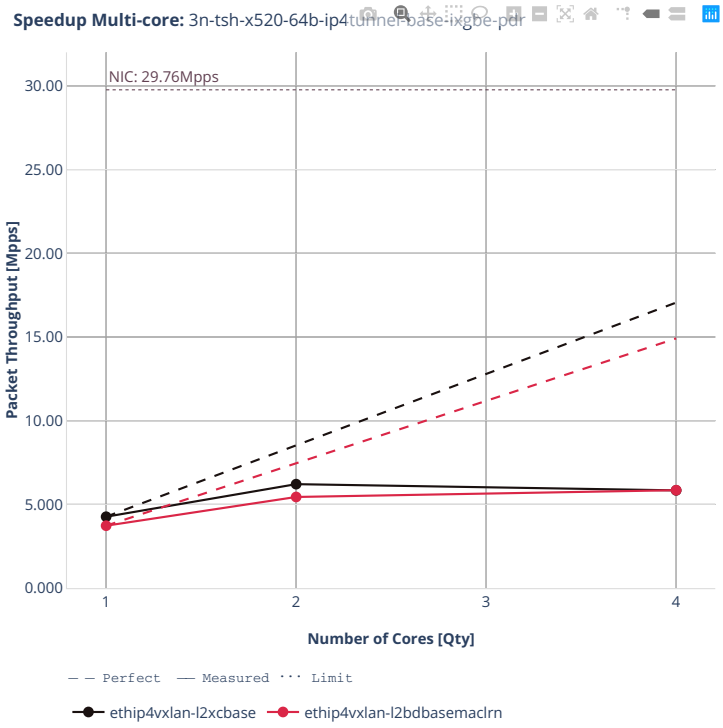
Speedup Multi-core: 3n-tsh-x520-64b-ip4tunnel-base-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4vxlan-l2xcbase
ethip4vxlan-l2bdbasemaclrn

3n-dnv-x553

64b-ip4tunnel-base-ixgbe

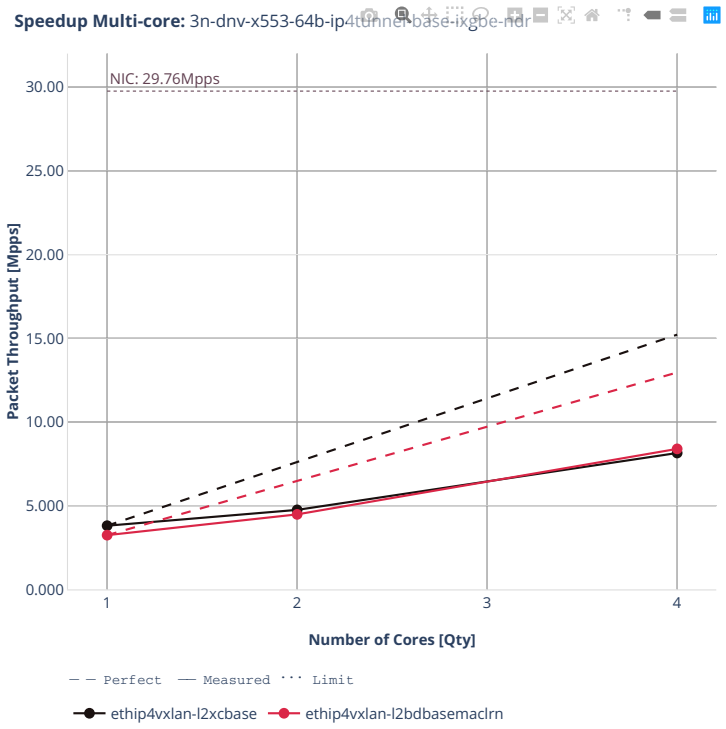

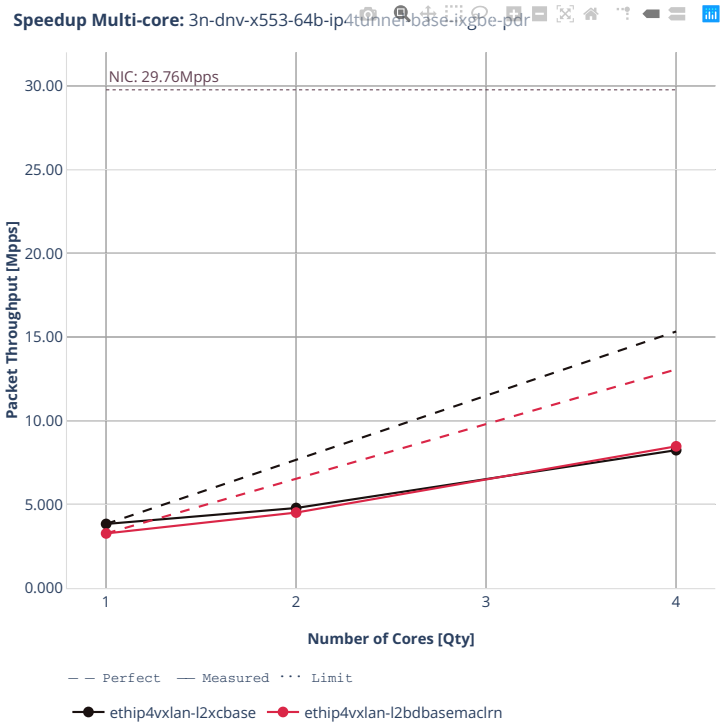
Speedup Multi-core: 3n-dnv-x553-64b-ip4tunnel-base-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4vxlan-l2xcbase
ethip4vxlan-l2bdbasemaclrn

2.4.6 NAT44 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[135].

[135] https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls2206

Det BiDir

2n-icx-xxv710

64b-nat44det-ip4routing-stl-bidir-avf

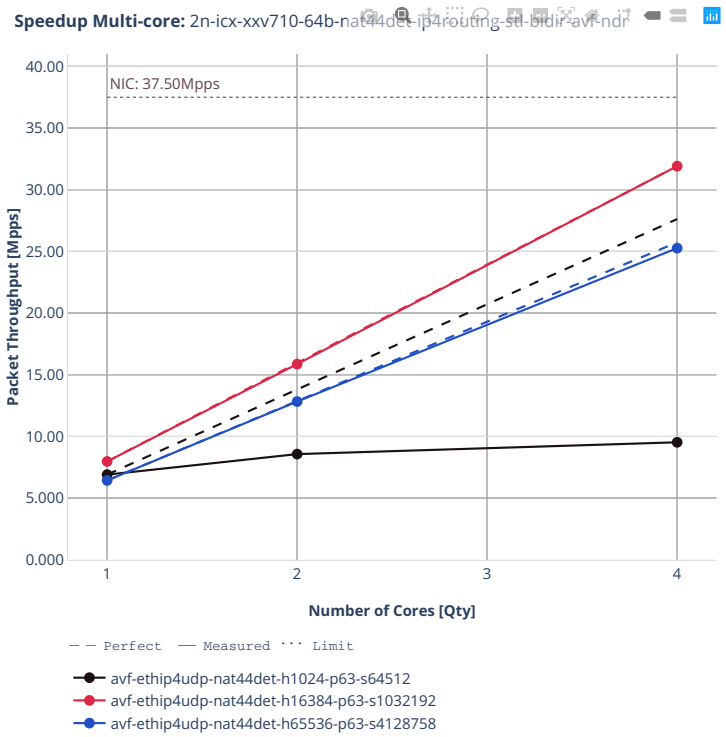

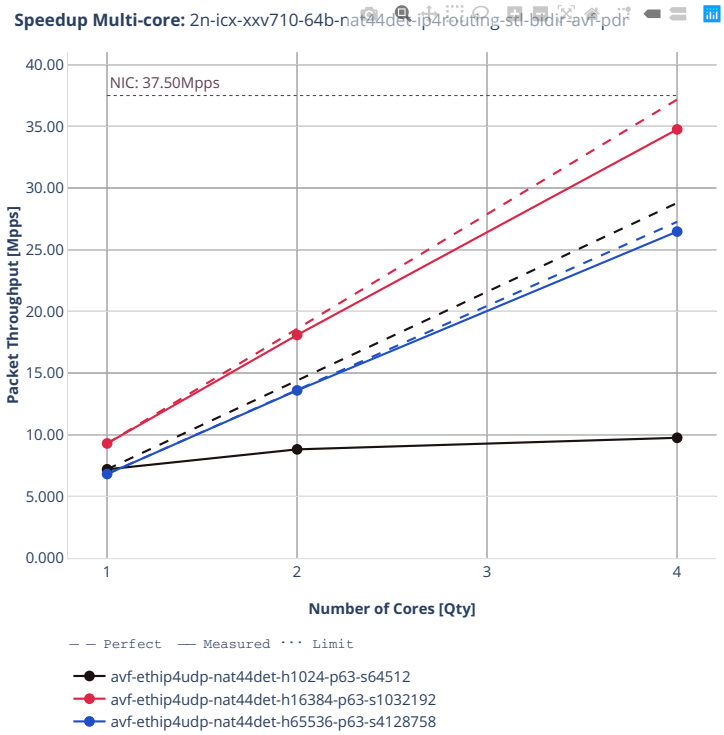
Speedup Multi-core: 2n-icx-xxv710-64b-nat44det-ip4routing-std-bidir-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-nat44det-h1024-p63-s64512
avf-ethip4udp-nat44det-h16384-p63-s1032192
avf-ethip4udp-nat44det-h65536-p63-s4128758

2n-clx-xxv710

64b-nat44det-ip4routing-stl-bidir-avf

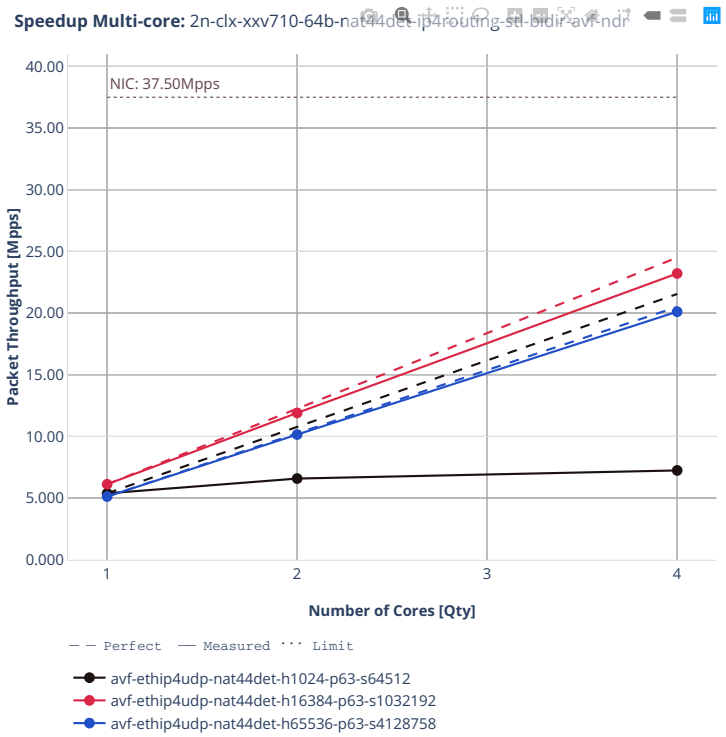

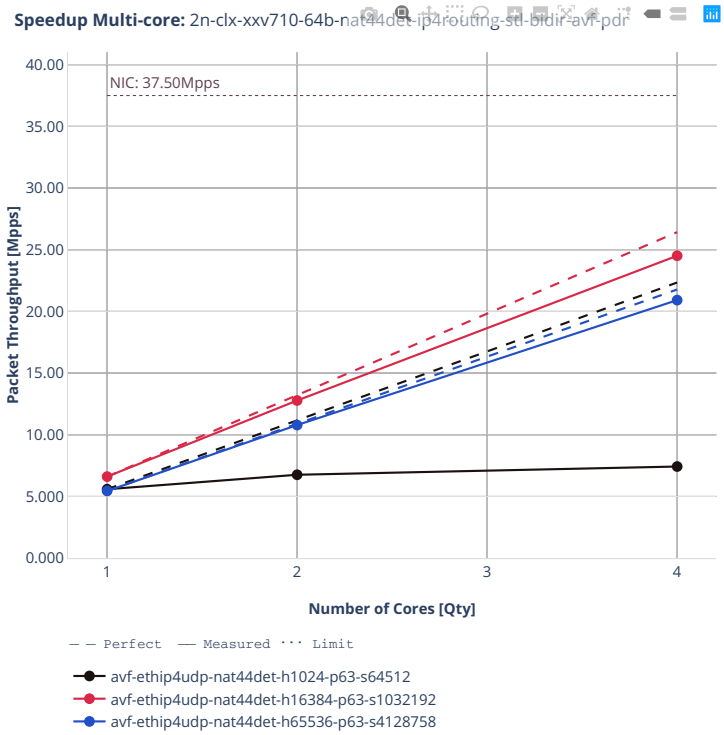
Speedup Multi-core: 2n-clx-xxv710-64b-nat44det-ip4routing-std-bidir-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-nat44det-h1024-p63-s64512
avf-ethip4udp-nat44det-h16384-p63-s1032192
avf-ethip4udp-nat44det-h65536-p63-s4128758

2n-skx-xxv710

64b-nat44det-ip4routing-stl-bidir-avf

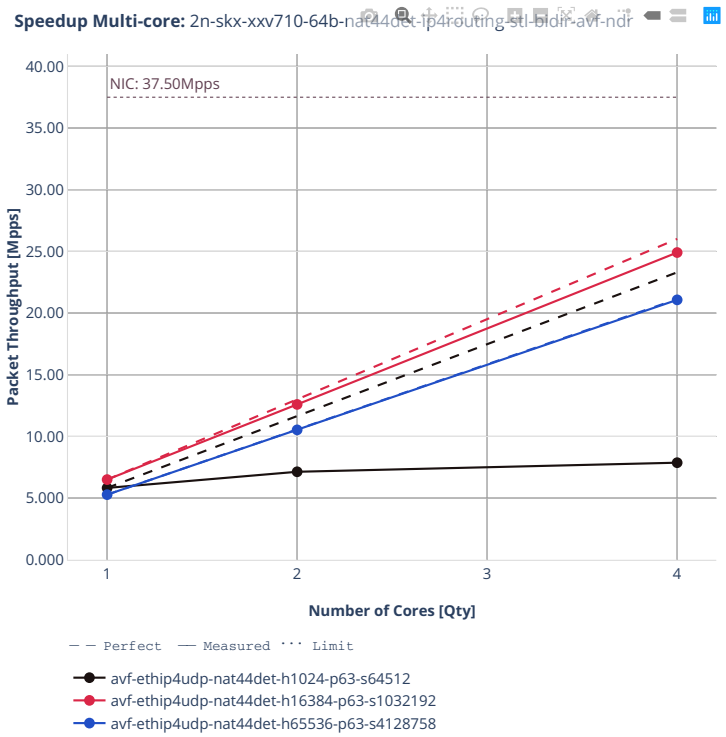

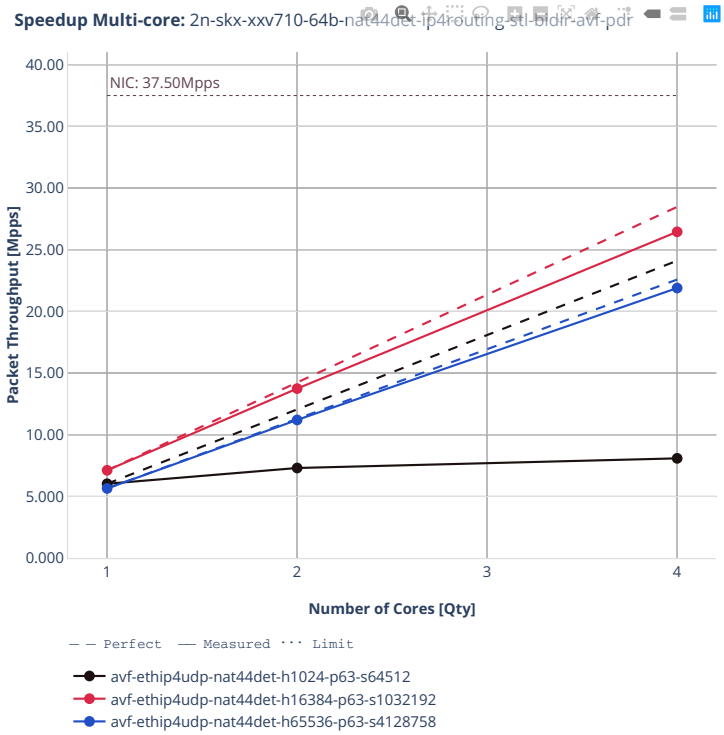
Speedup Multi-core: 2n-skx-xxv710-64b-nat44det-ip4routing-stl-bidir-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-nat44det-h1024-p63-s64512
avf-ethip4udp-nat44det-h16384-p63-s1032192
avf-ethip4udp-nat44det-h65536-p63-s4128758

ED UniDir

2n-icx-xxv710

64b-nat44ed-ip4routing-stl-unidir-avf

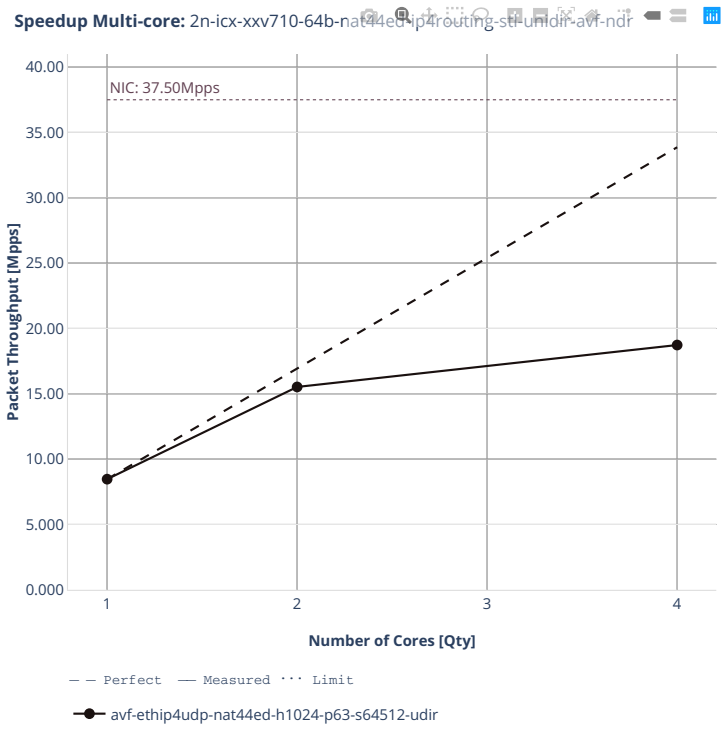

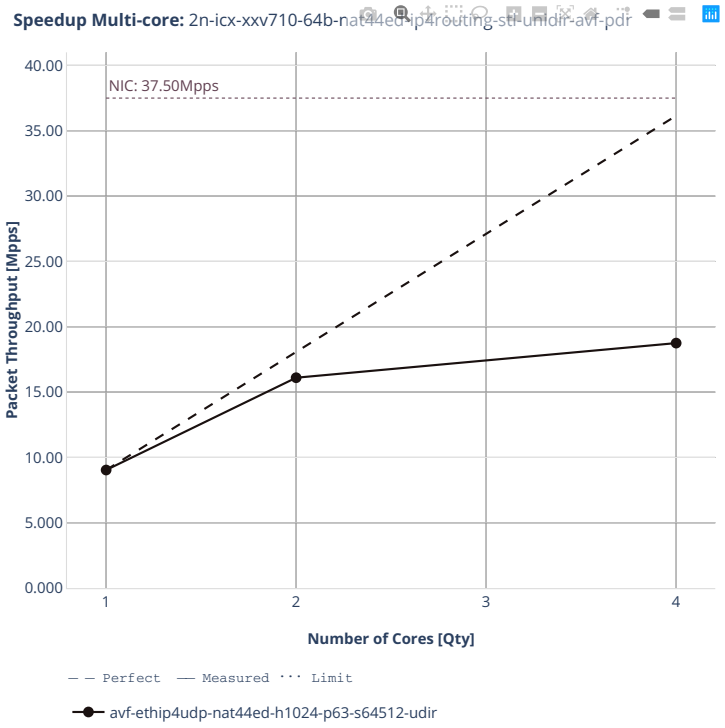
Speedup Multi-core: 2n-icx-xxv710-64b-nat44ed-ip4routing-std-udir-avf-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-nat44ed-h1024-p63-s64512-udir

2n-clx-xxv710

64b-nat44ed-ip4routing-stl-unidir-avf

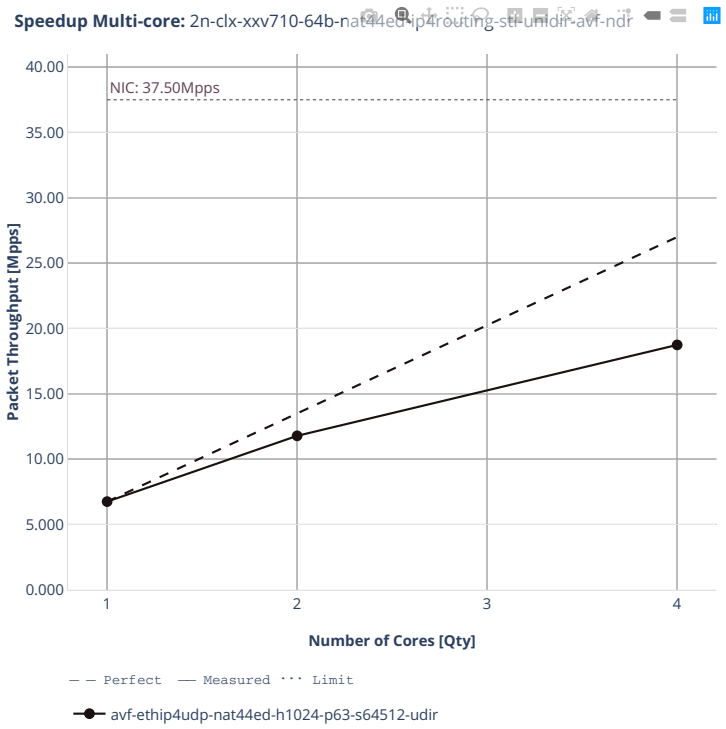

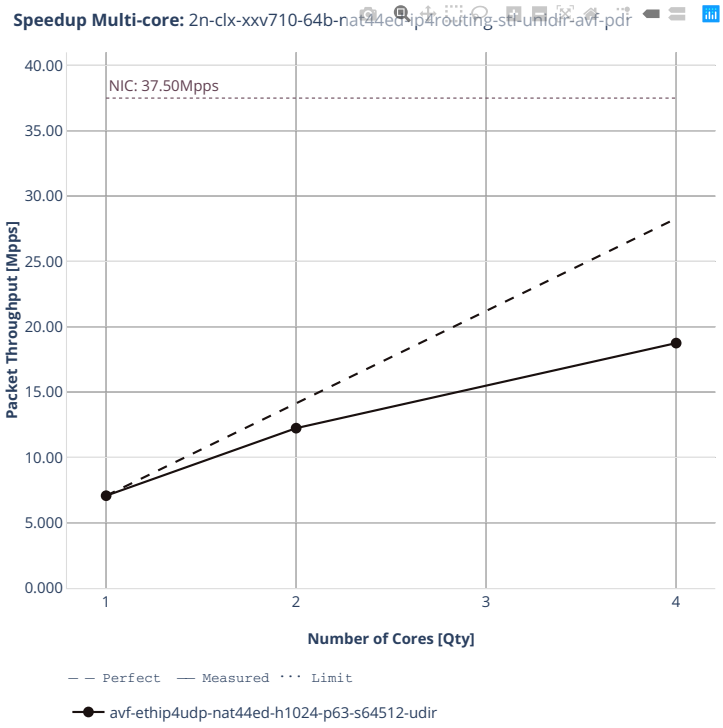
Speedup Multi-core: 2n-clx-xxv710-64b-nat44ed-ip4routing-std-undir-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-nat44ed-h1024-p63-s64512-udir

2n-skx-xxv710

64b-nat44ed-ip4routing-stl-unidir-avf

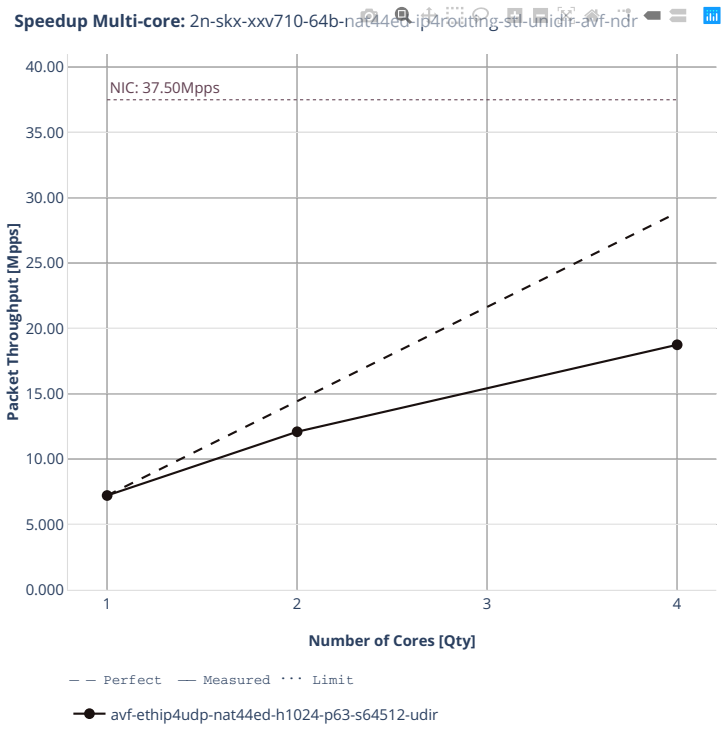

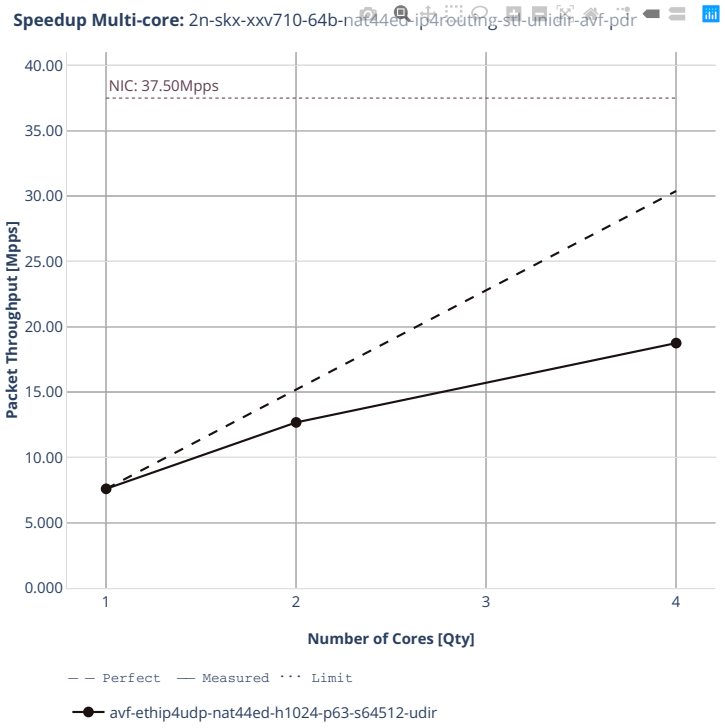
Speedup Multi-core: 2n-skx-xxv710-64b-nat44ed-ip4routing-std-unidir-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-nat44ed-h1024-p63-s64512-udir

ED UDP CPS

2n-icx-xxv710

64b-nat44ed-ip4routing-udp-stf-cps-avf

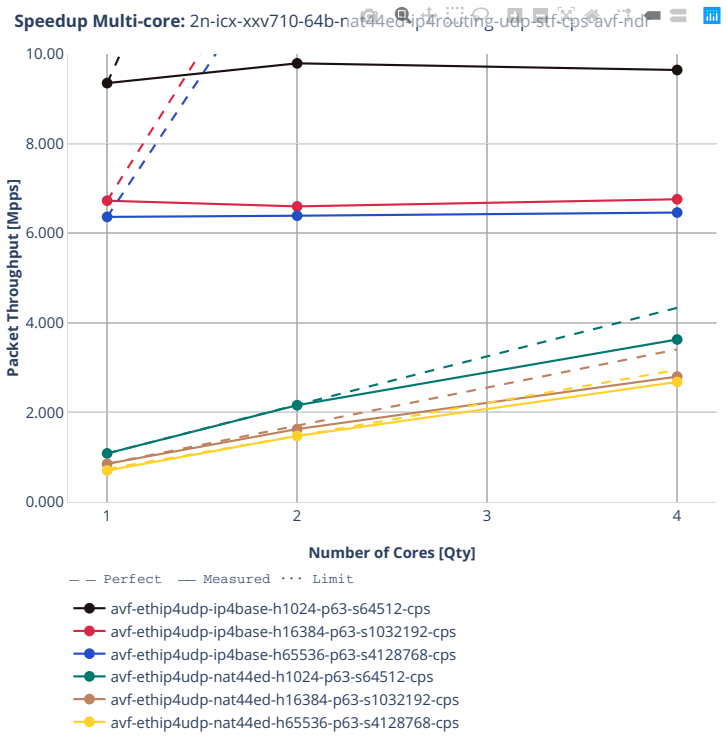

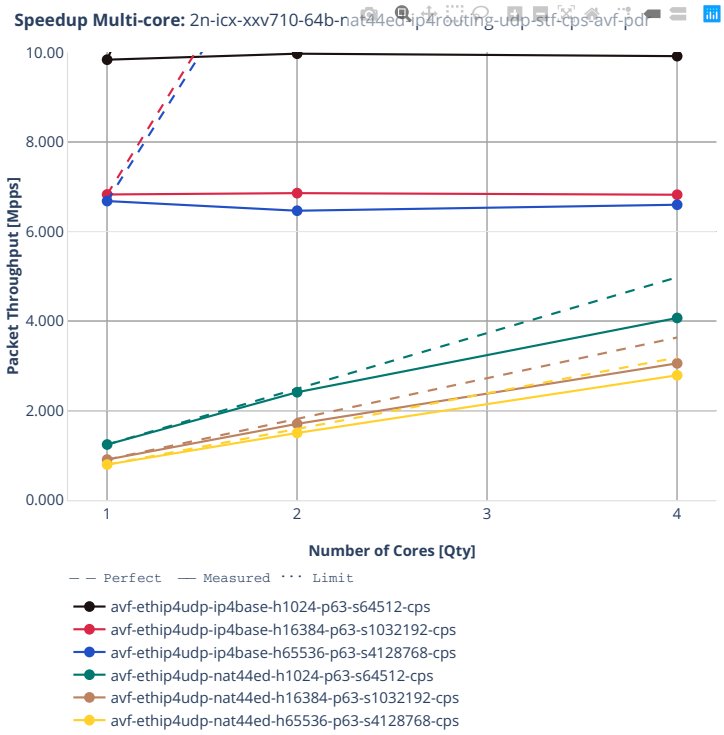
Speedup Multi-core: 2n-icx-xxv710-64b-nat44ed-ip4routing-udp-stf-cps-avf-pdr
Packet Throughput [Mpps]
10.00
8.000
6.000
4.000
2.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-ip4base-h1024-p63-s64512-cps
avf-ethip4udp-ip4base-h16384-p63-s1032192-cps
avf-ethip4udp-ip4base-h65536-p63-s4128768-cps
avf-ethip4udp-nat44ed-h1024-p63-s64512-cps
avf-ethip4udp-nat44ed-h16384-p63-s1032192-cps
avf-ethip4udp-nat44ed-h65536-p63-s4128768-cps

2n-clx-xxv710

64b-nat44ed-ip4routing-udp-stf-cps-avf

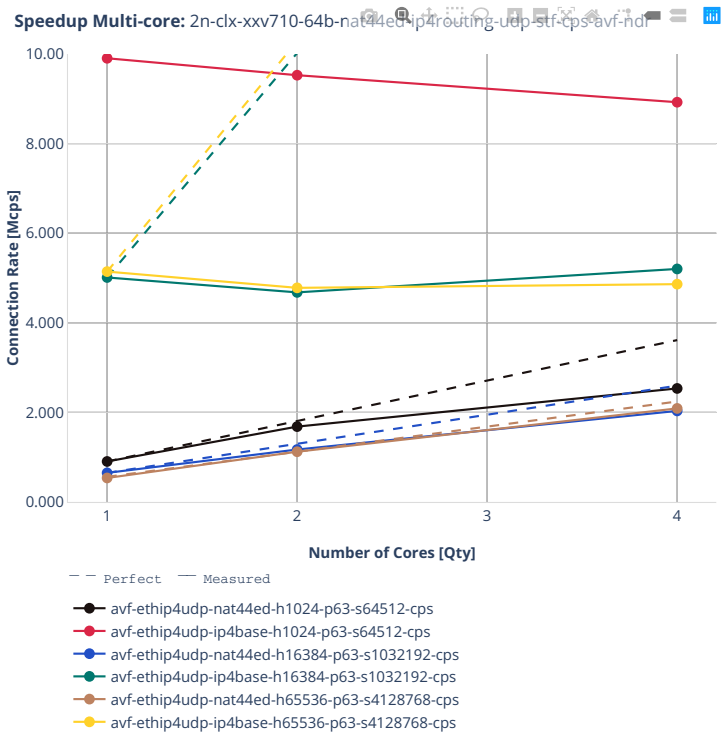

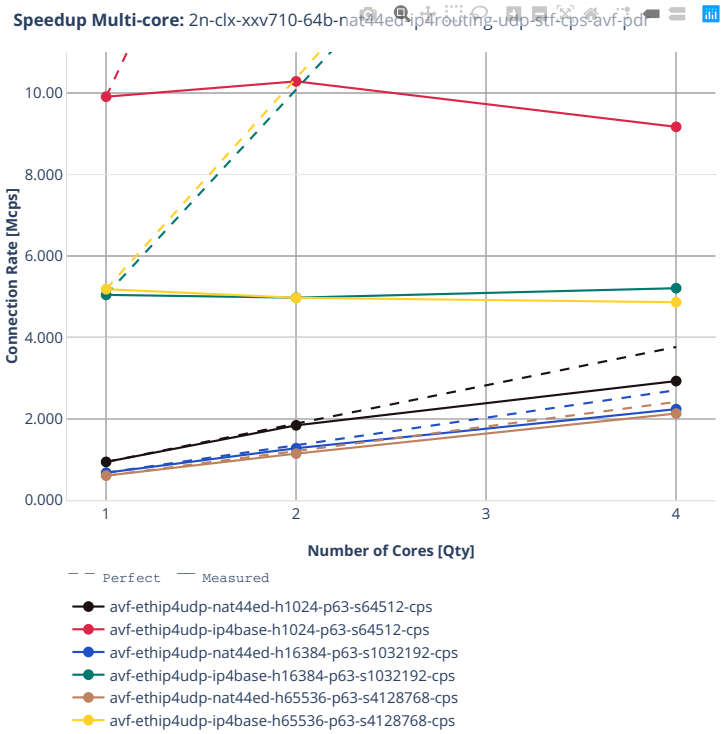
Speedup Multi-core: 2n-clx-xxv710-64b-nat44ed-ip4routing-udp-stf-cps-avf-pdr
Connection Rate [Mcps]
10.00
8.000
6.000
4.000
2.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
avf-ethip4udp-nat44ed-h1024-p63-s64512-cps
avf-ethip4udp-ip4base-h1024-p63-s64512-cps
avf-ethip4udp-nat44ed-h16384-p63-s1032192-cps
avf-ethip4udp-ip4base-h16384-p63-s1032192-cps
avf-ethip4udp-nat44ed-h65536-p63-s4128768-cps
avf-ethip4udp-ip4base-h65536-p63-s4128768-cps

2n-skx-xxv710

64b-nat44ed-ip4routing-udp-stf-cps-avf

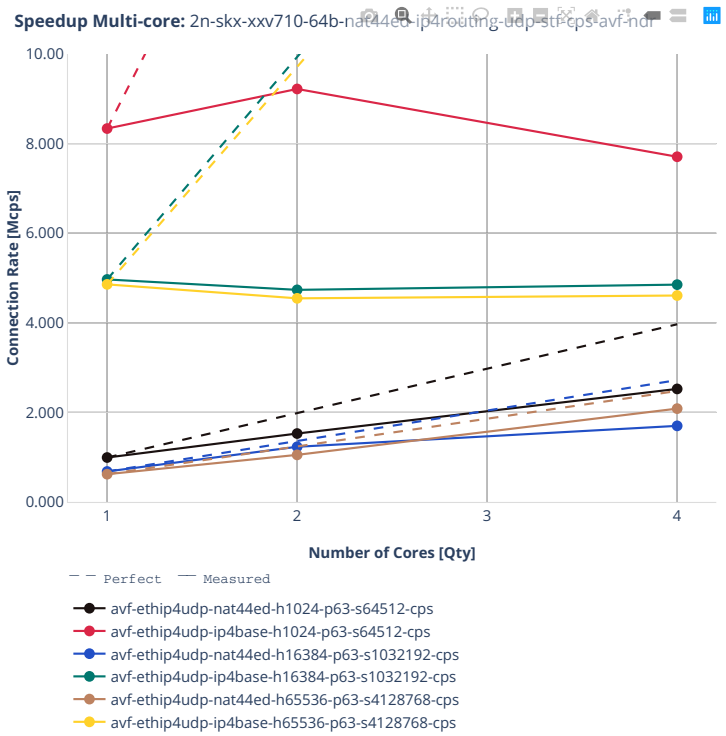

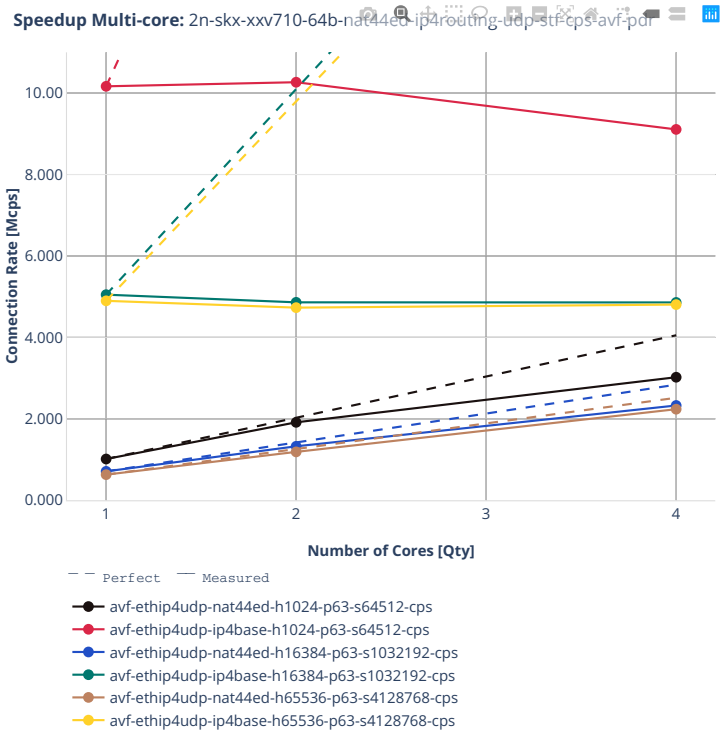
Speedup Multi-core: 2n-skx-xxv710-64b-nat44ed-ip4routing-udp-stf-cps-avf-pdr
Connection Rate [Mcps]
10.00
8.000
6.000
4.000
2.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
avf-ethip4udp-nat44ed-h1024-p63-s64512-cps
avf-ethip4udp-ip4base-h1024-p63-s64512-cps
avf-ethip4udp-nat44ed-h16384-p63-s1032192-cps
avf-ethip4udp-ip4base-h16384-p63-s1032192-cps
avf-ethip4udp-nat44ed-h65536-p63-s4128768-cps
avf-ethip4udp-ip4base-h65536-p63-s4128768-cps

ED TCP CPS

2n-icx-xxv710

64b-nat44ed-ip4routing-tcp-stf-cps-avf

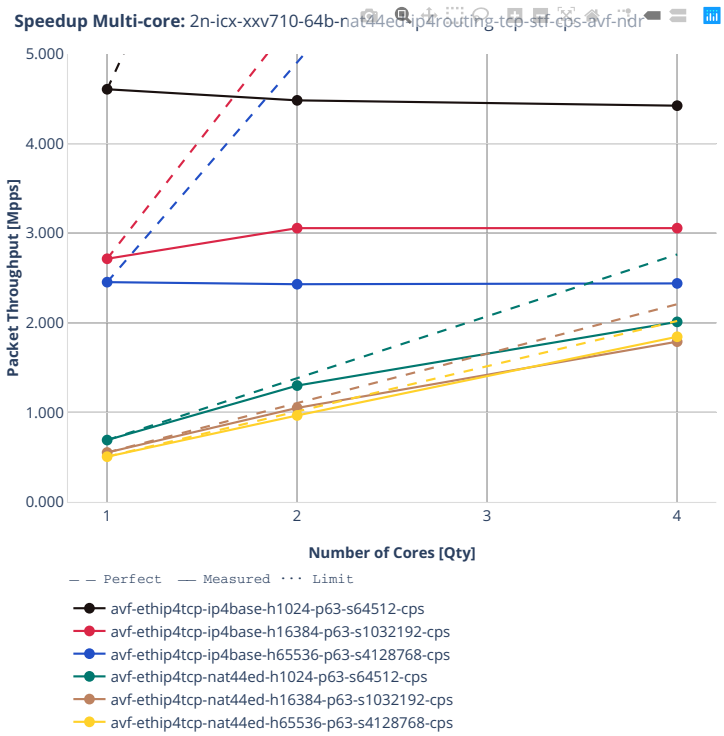

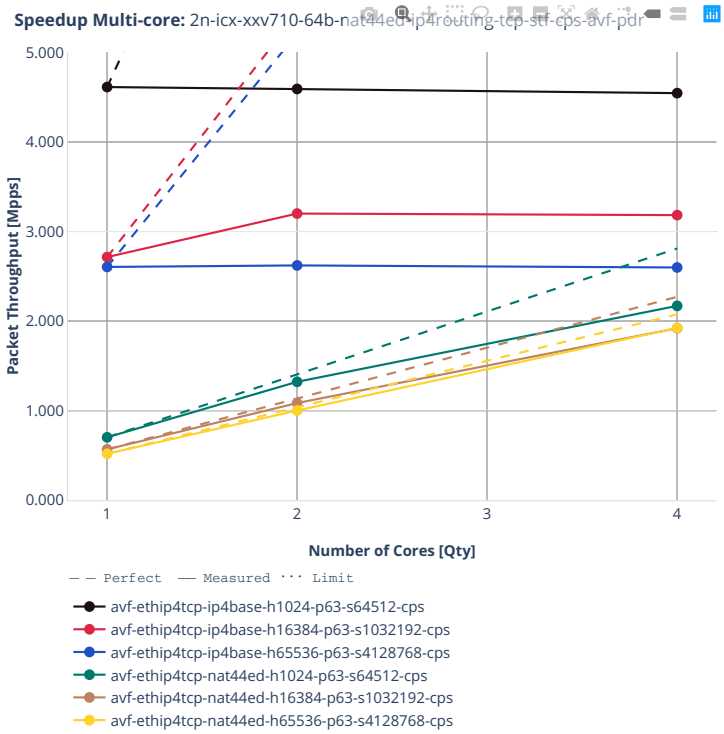
Speedup Multi-core: 2n-icx-xxv710-64b-nat44ed-ip4routing-tcp-stf-cps-avf-pdr
Packet Throughput [Mpps]
5.000
4.000
3.000
2.000
1.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4tcp-ip4base-h1024-p63-s64512-cps
avf-ethip4tcp-ip4base-h16384-p63-s1032192-cps
avf-ethip4tcp-ip4base-h65536-p63-s4128768-cps
avf-ethip4tcp-nat44ed-h1024-p63-s64512-cps
avf-ethip4tcp-nat44ed-h16384-p63-s1032192-cps
avf-ethip4tcp-nat44ed-h65536-p63-s4128768-cps

2n-clx-xxv710

64b-nat44ed-ip4routing-tcp-stf-cps-avf

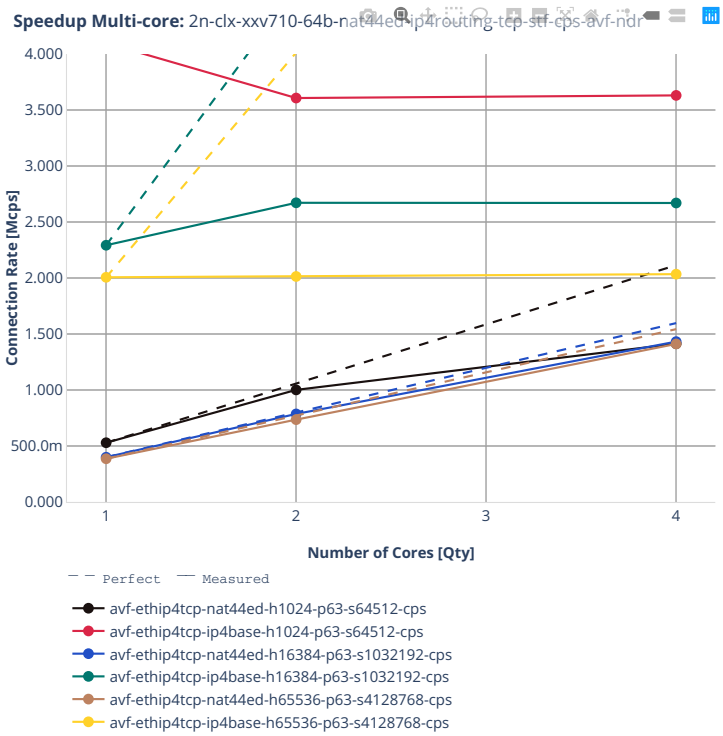

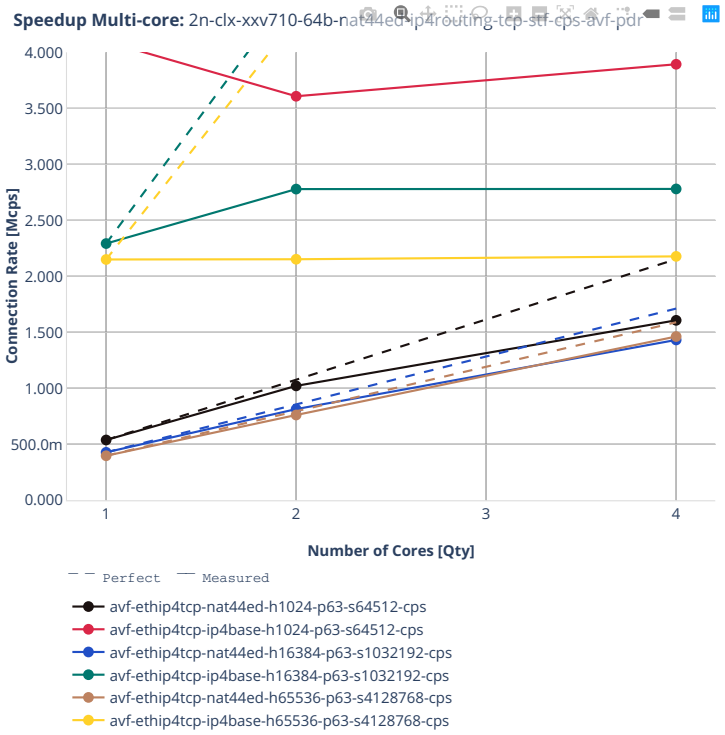
Speedup Multi-core: 2n-clx-xxv710-64b-nat44ed-ip4routing-tcp-stf-cps-avf-pdr
Connection Rate [Mcps]
4.000
3.500
3.000
2.500
2.000
1.500
1.000
500.0m
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
avf-ethip4tcp-nat44ed-h1024-p63-s64512-cps
avf-ethip4tcp-ip4base-h1024-p63-s64512-cps
avf-ethip4tcp-nat44ed-h16384-p63-s1032192-cps
avf-ethip4tcp-ip4base-h16384-p63-s1032192-cps
avf-ethip4tcp-nat44ed-h65536-p63-s4128768-cps
avf-ethip4tcp-ip4base-h65536-p63-s4128768-cps

2n-skx-xxv710

64b-nat44ed-ip4routing-tcp-stf-cps-avf

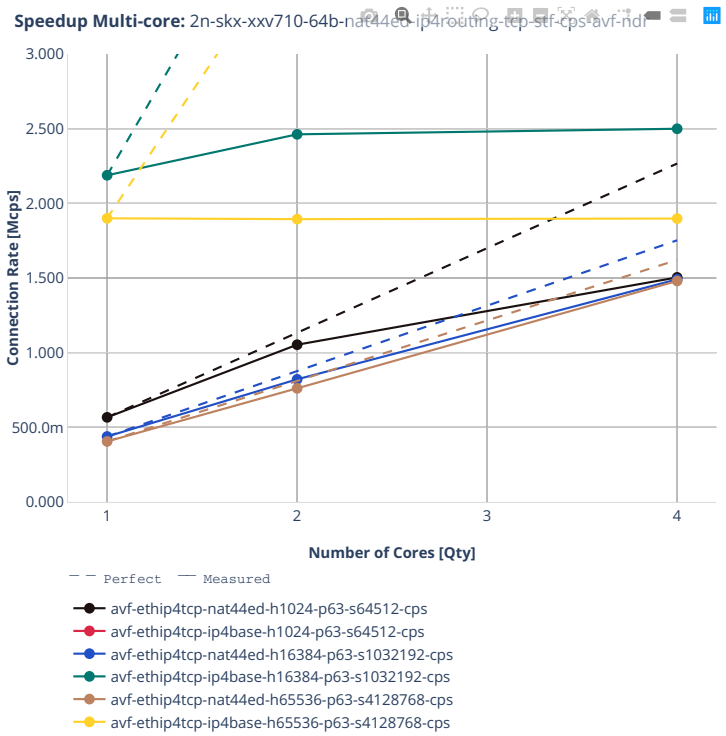

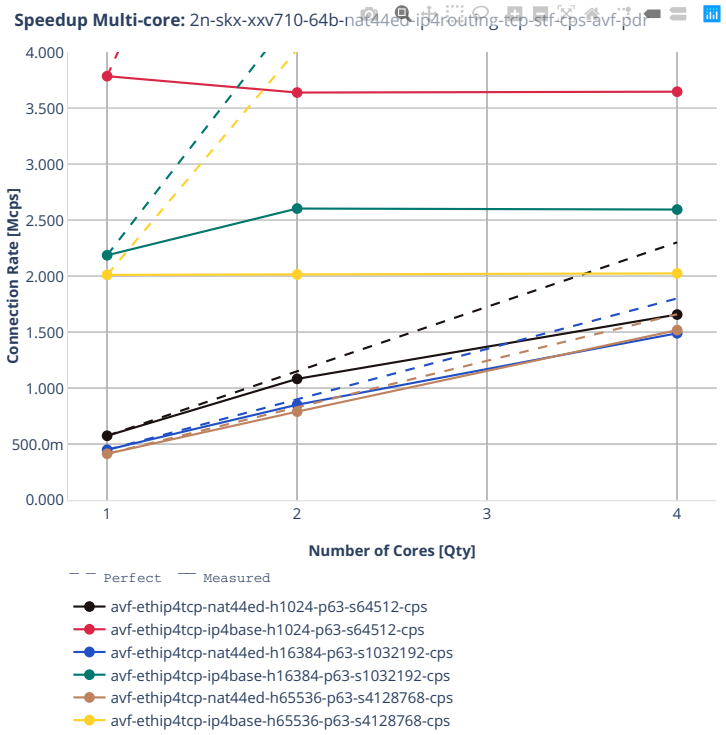
Speedup Multi-core: 2n-skx-xxv710-64b-nat44ed-ip4routing-tcp-stf-cps-avf-pdr
Connection Rate [Mcps]
4.000
3.500
3.000
2.500
2.000
1.500
1.000
500.0m
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
avf-ethip4tcp-nat44ed-h1024-p63-s64512-cps
avf-ethip4tcp-ip4base-h1024-p63-s64512-cps
avf-ethip4tcp-nat44ed-h16384-p63-s1032192-cps
avf-ethip4tcp-ip4base-h16384-p63-s1032192-cps
avf-ethip4tcp-nat44ed-h65536-p63-s4128768-cps
avf-ethip4tcp-ip4base-h65536-p63-s4128768-cps

ED UDP TPUT

2n-icx-xxv710

100b-nat44ed-ip4routing-udp-tput-avf

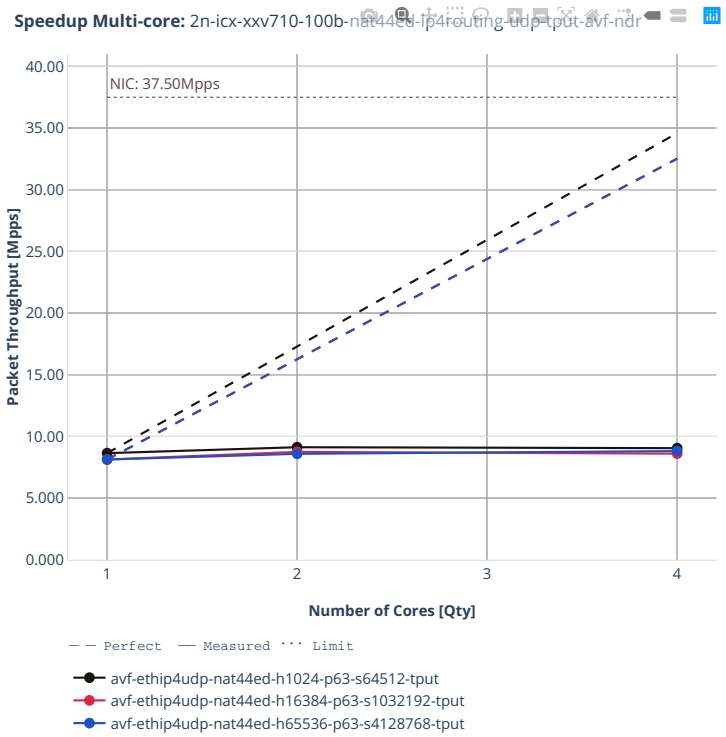

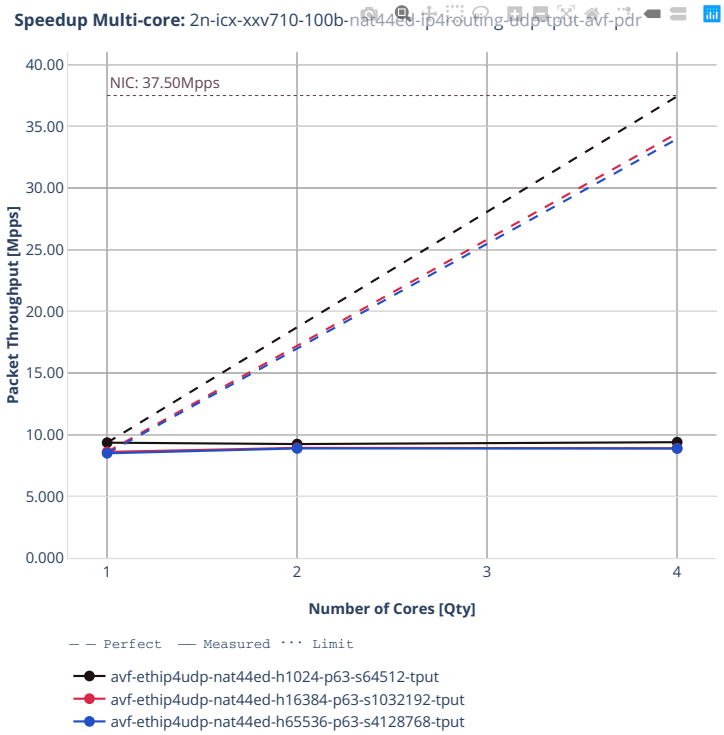
Speedup Multi-core: 2n-icx-xxv710-100b-nat44ed-ip4routing-udp-tput-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-nat44ed-h1024-p63-s64512-tput
avf-ethip4udp-nat44ed-h16384-p63-s1032192-tput
avf-ethip4udp-nat44ed-h65536-p63-s4128768-tput

2n-clx-xxv710

100b-nat44ed-ip4routing-udp-tput-avf

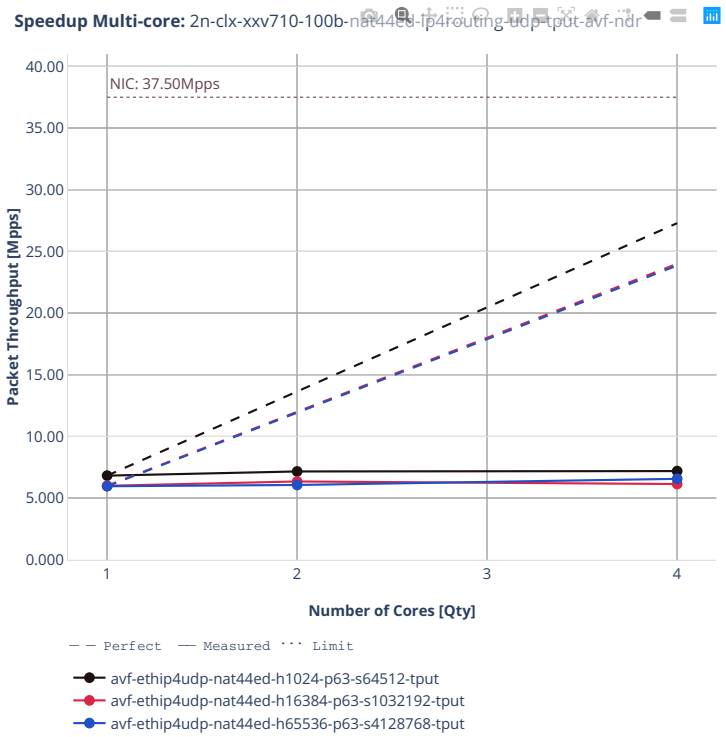

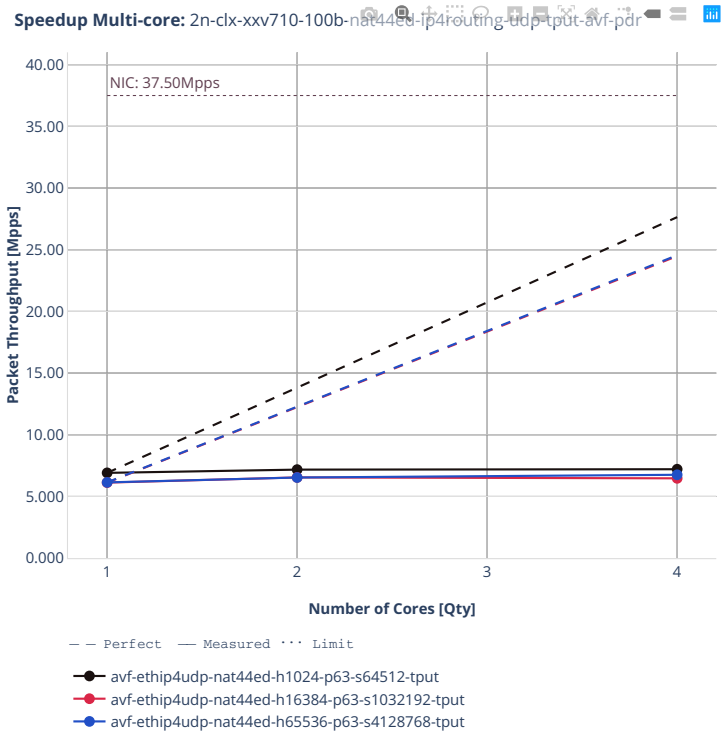
Speedup Multi-core: 2n-clx-xxv710-100b-nat44ed-ip4routing-udp-tput-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-nat44ed-h1024-p63-s64512-tput
avf-ethip4udp-nat44ed-h16384-p63-s1032192-tput
avf-ethip4udp-nat44ed-h65536-p63-s4128768-tput

2n-skx-xxv710

100b-nat44ed-ip4routing-udp-tput-avf

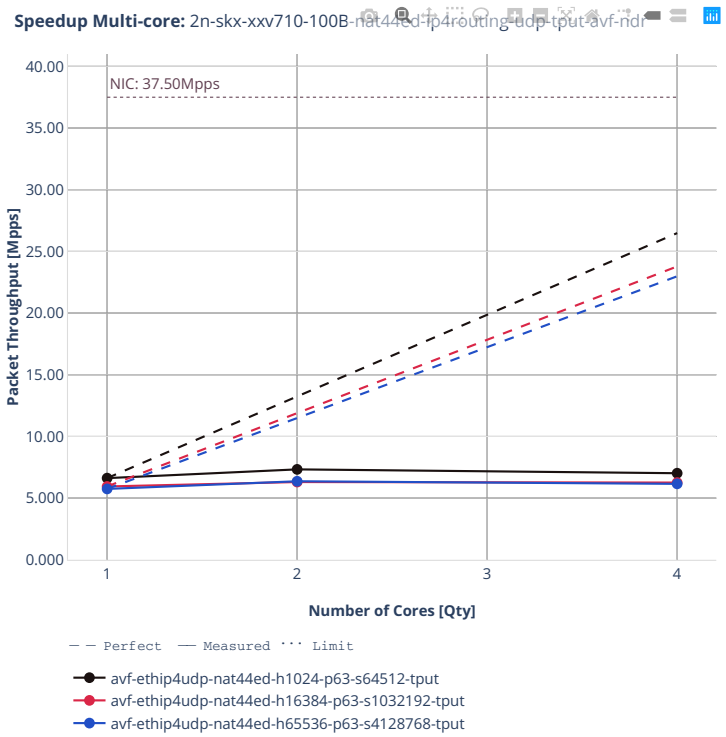

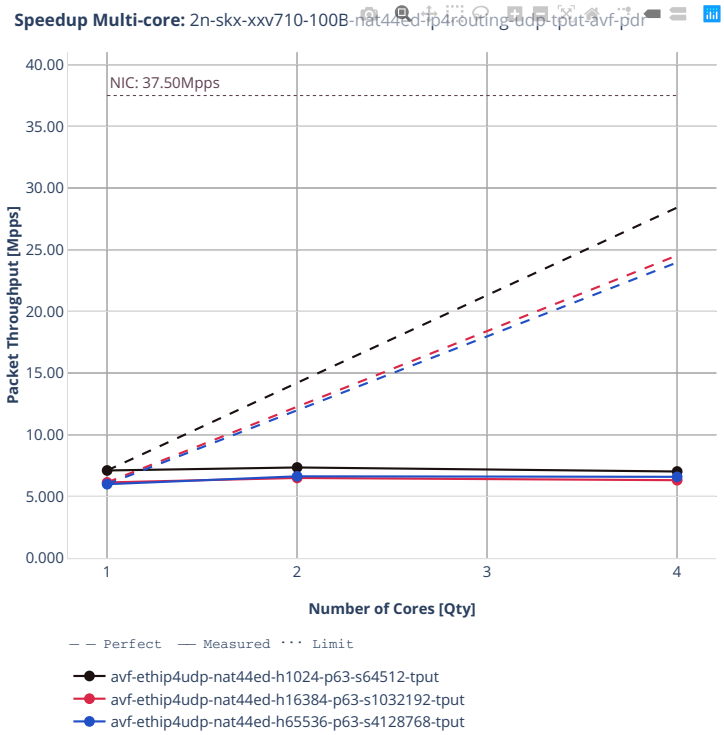
Speedup Multi-core: 2n-skx-xxv710-100B-nat44ed-ip4routing-udp-tput-avf-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4udp-nat44ed-h1024-p63-s64512-tput
avf-ethip4udp-nat44ed-h16384-p63-s1032192-tput
avf-ethip4udp-nat44ed-h65536-p63-s4128768-tput

ED TCP TPUT

2n-icx-xxv710

100b-nat44ed-ip4routing-tcp-tput-avf

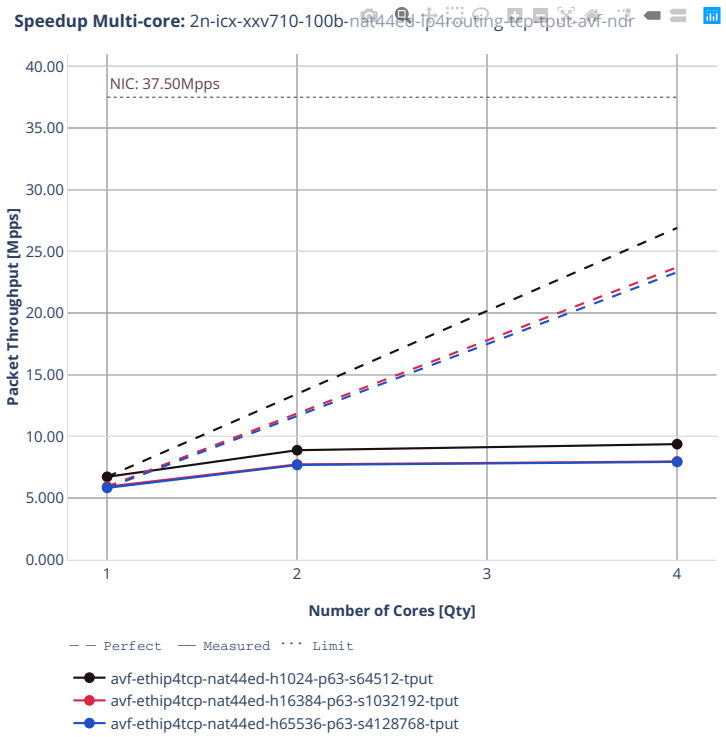

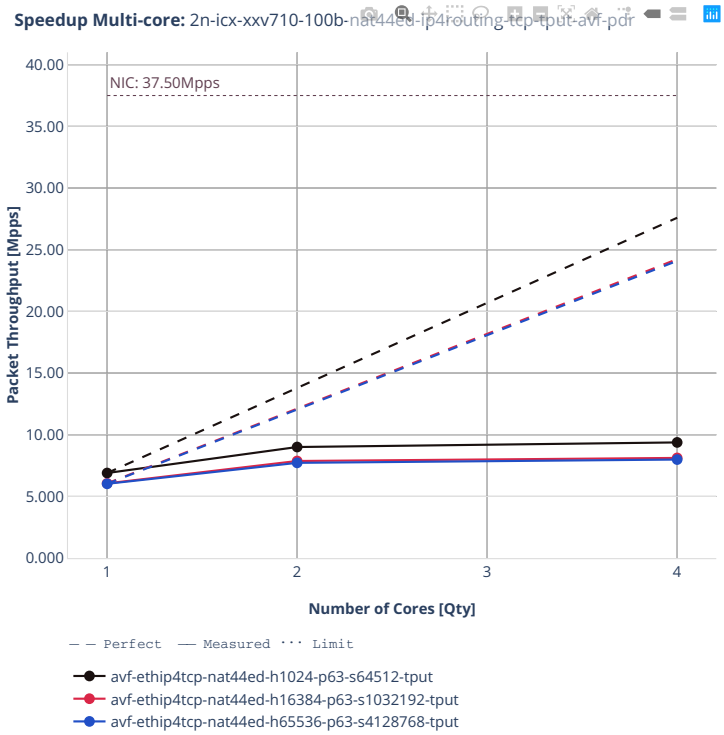
Speedup Multi-core: 2n-icx-xxv710-100b-nat44ed-ip4routing-tcp-tput-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4tcp-nat44ed-h1024-p63-s64512-tput
avf-ethip4tcp-nat44ed-h16384-p63-s1032192-tput
avf-ethip4tcp-nat44ed-h65536-p63-s4128768-tput

2n-clx-xxv710

100b-nat44ed-ip4routing-tcp-tput-avf

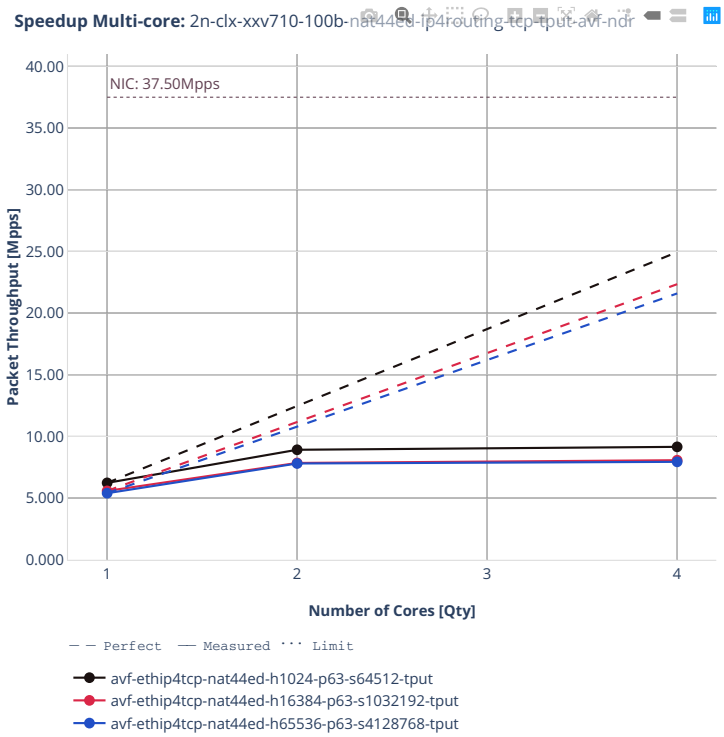

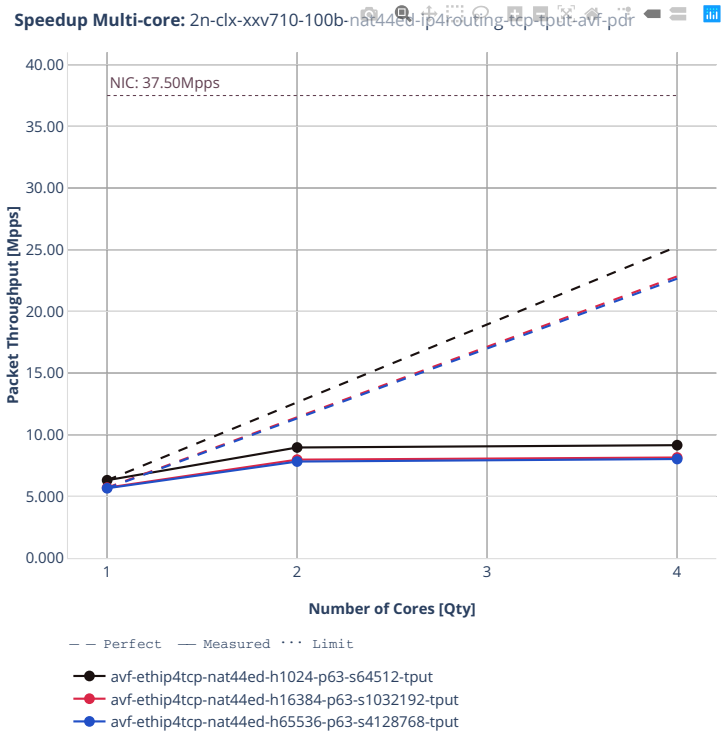
Speedup Multi-core: 2n-clx-xxv710-100b-nat44ed-ip4routing-tcp-tput-avf-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4tcp-nat44ed-h1024-p63-s64512-tput
avf-ethip4tcp-nat44ed-h16384-p63-s1032192-tput
avf-ethip4tcp-nat44ed-h65536-p63-s4128768-tput

2n-skx-xxv710

100b-nat44ed-ip4routing-tcp-tput-avf

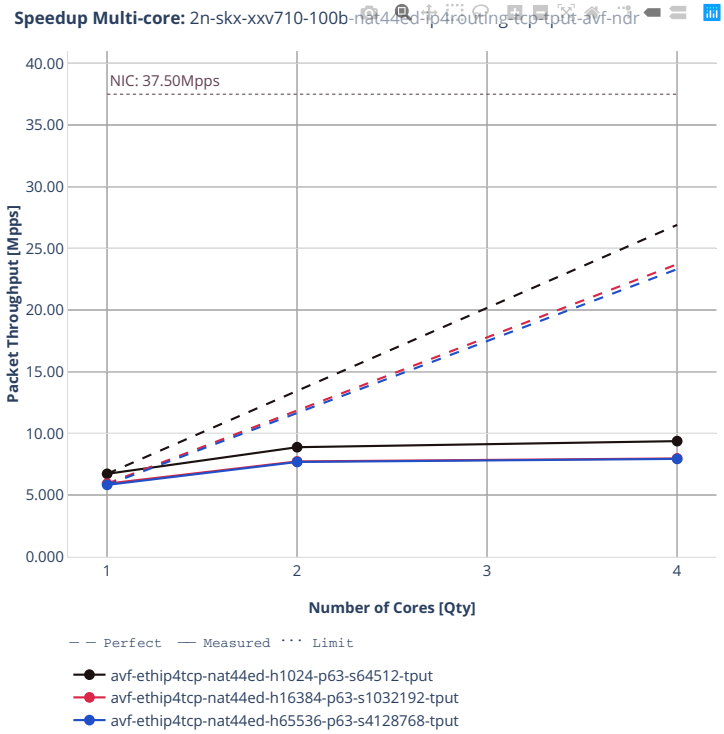

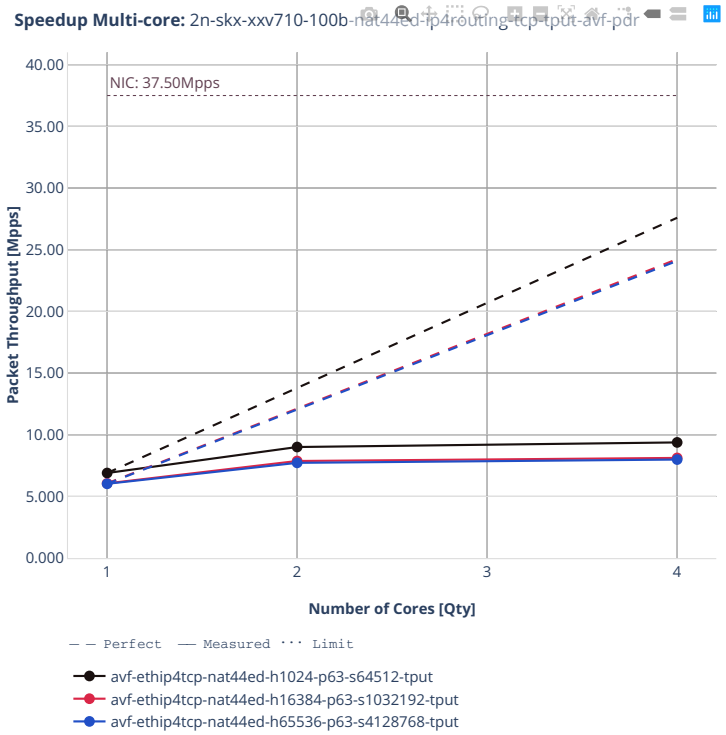
Speedup Multi-core: 2n-skx-xxv710-100b-nat44ed-ip4routing-tcp-tput-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4tcp-nat44ed-h1024-p63-s64512-tput
avf-ethip4tcp-nat44ed-h16384-p63-s1032192-tput
avf-ethip4tcp-nat44ed-h65536-p63-s4128768-tput

2.4.7 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[136].

[136] https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls2206

2n-icx-xxv710

64b-vhost-base

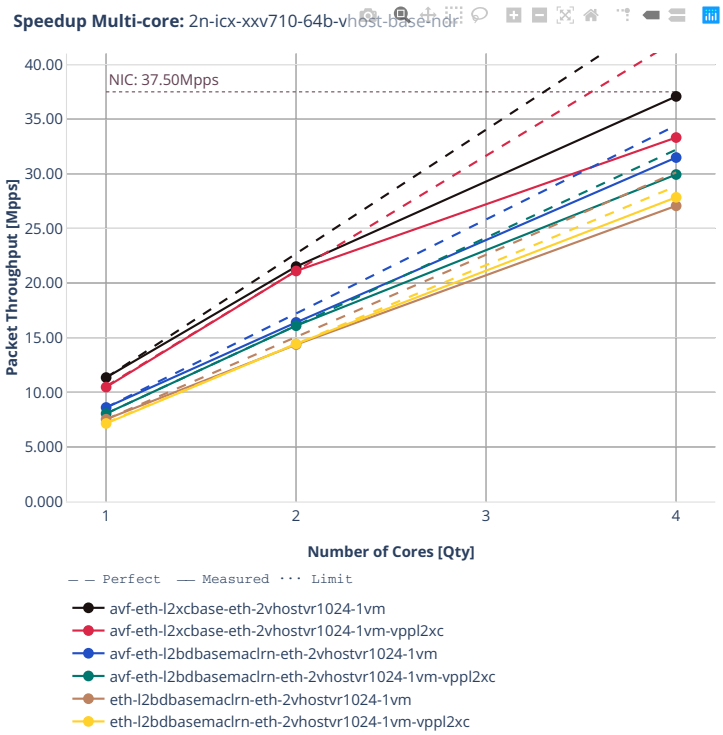

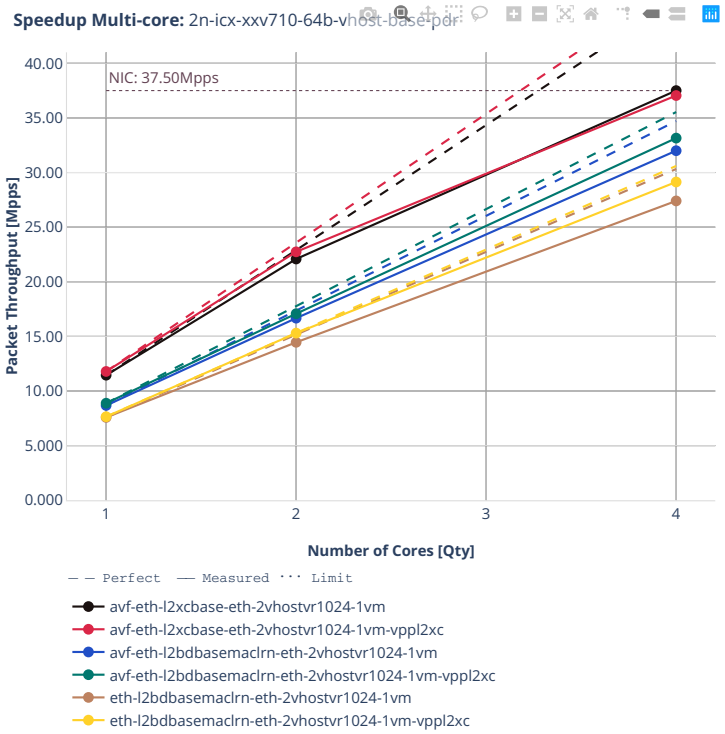
Speedup Multi-core: 2n-icx-xxv710-64b-vhost-base-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2xcbase-eth-2vhostvr1024-1vm
avf-eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2n-skx-xxv710

64b-vhost-base-testpmd

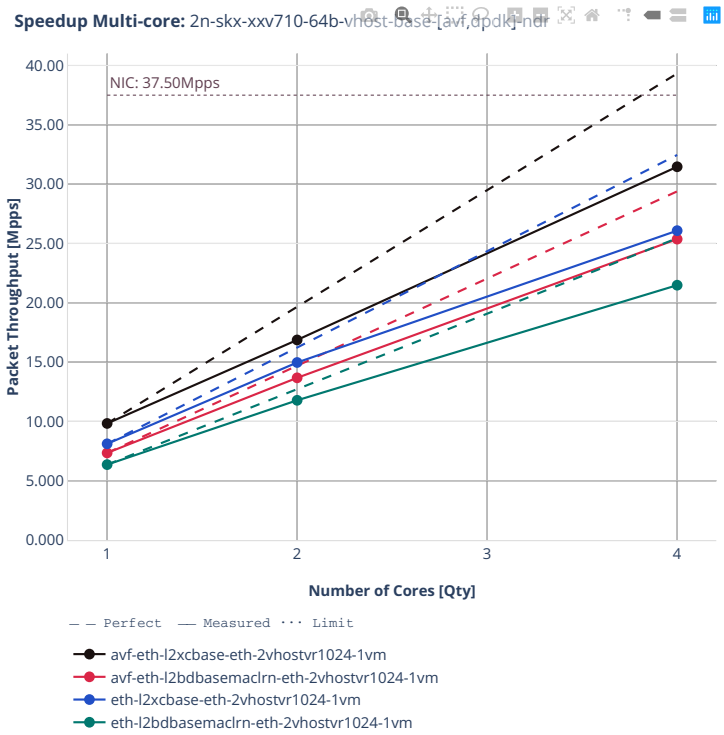

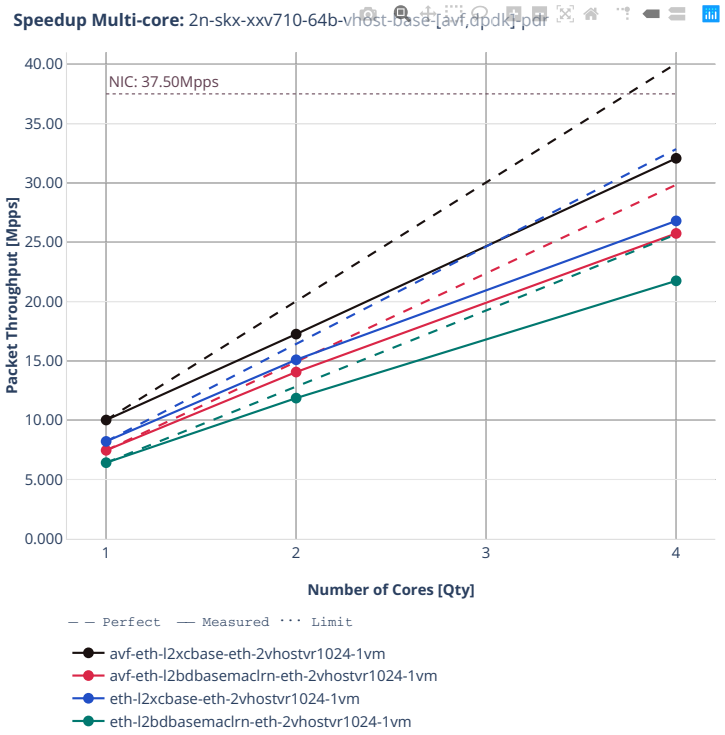
Speedup Multi-core: 2n-skx-xxv710-64b-vhost-base-[avf,dpdk]-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2xcbase-eth-2vhostvr1024-1vm
avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
eth-l2xcbase-eth-2vhostvr1024-1vm
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm

64b-vhost-base-vpp

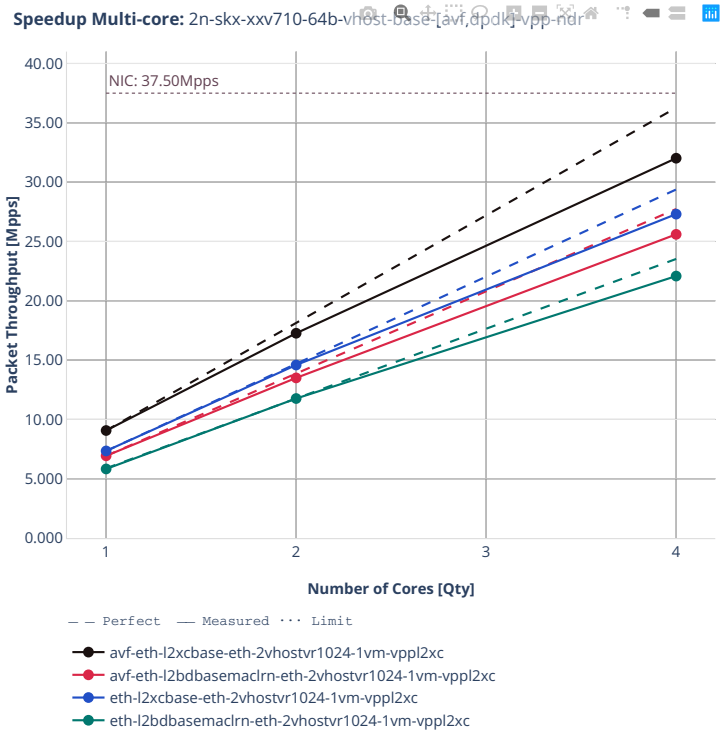

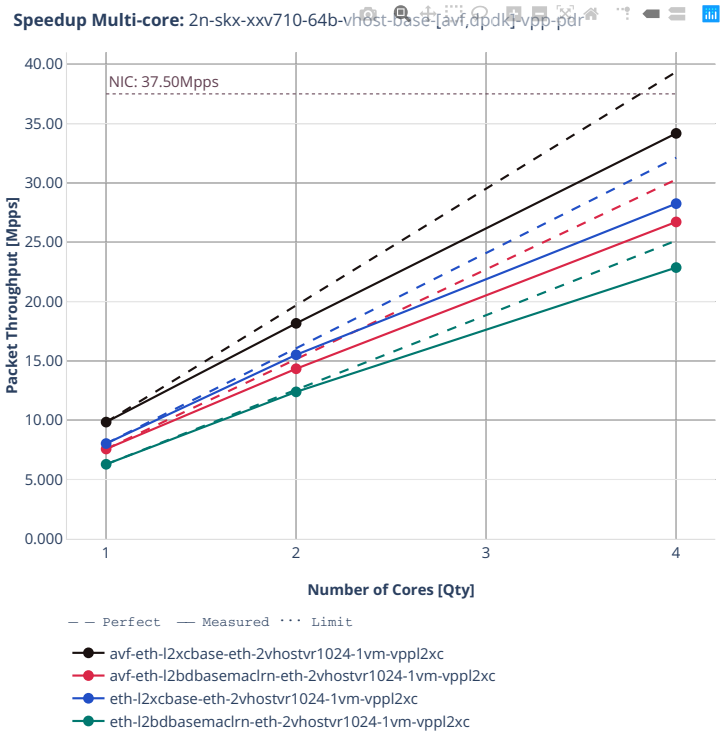
Speedup Multi-core: 2n-skx-xxv710-64b-vhost-base-[avf,dpdk]-vpp-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2n-clx-xxv710

64b-vhost-base-testpmd

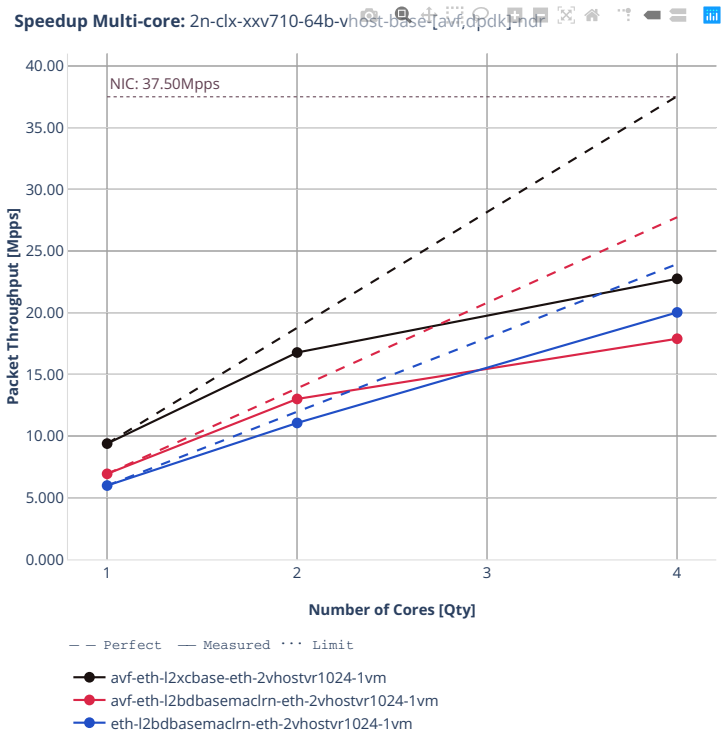

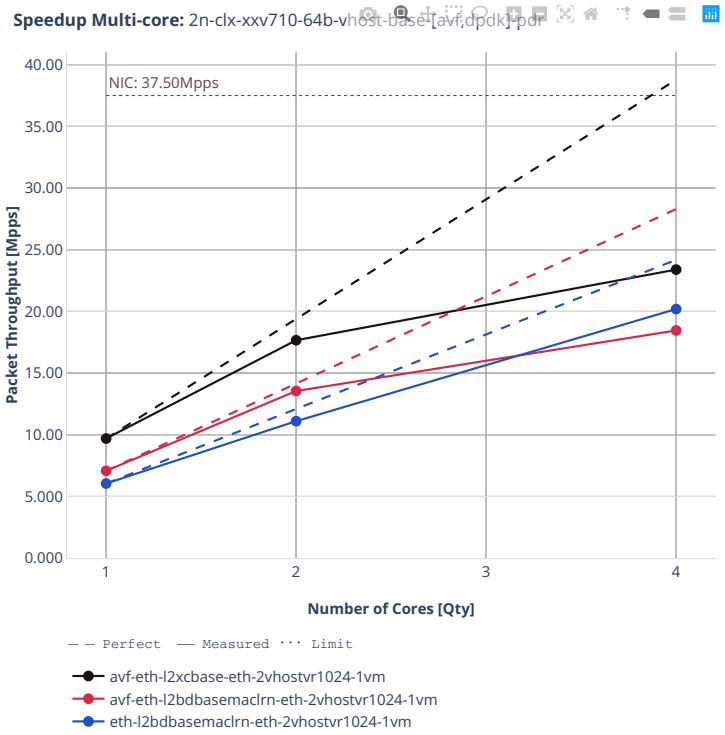
Speedup Multi-core: 2n-clx-xxv710-64b-vhost-base-[avf,dpdk]-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2xcbase-eth-2vhostvr1024-1vm
avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm

64b-vhost-base-vpp

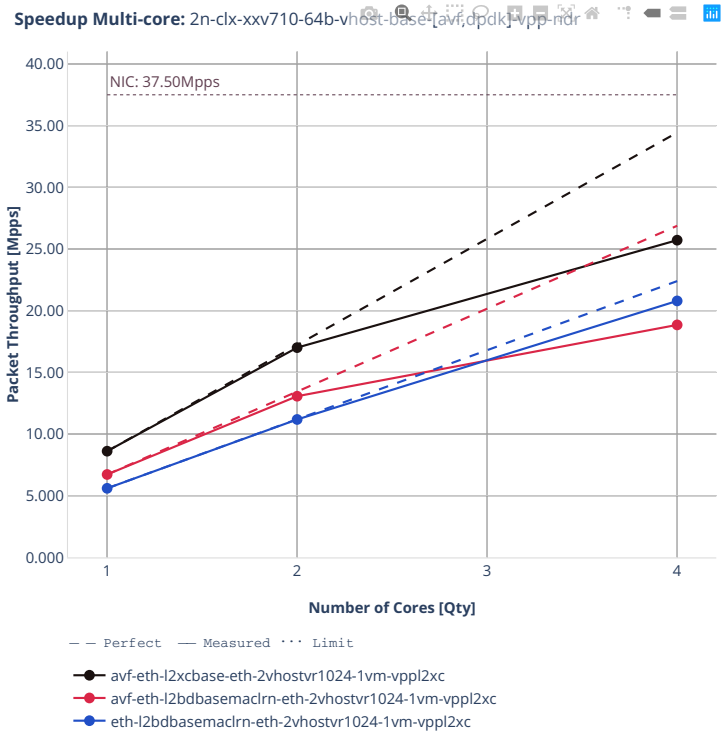

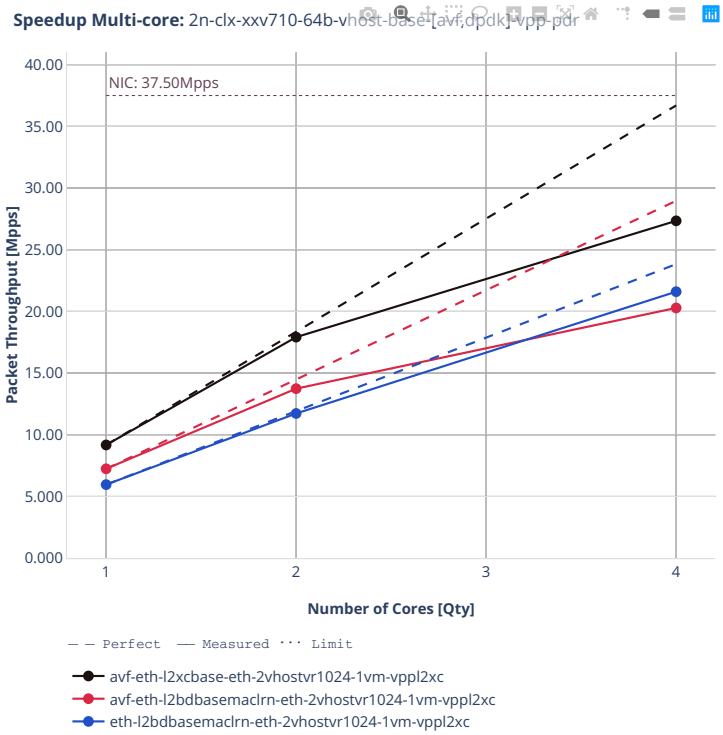
Speedup Multi-core: 2n-clx-xxv710-64b-vhost-base-[avf,dpdk]-vpp-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2n-clx-cx556a

64b-vhost-base-rdma-core

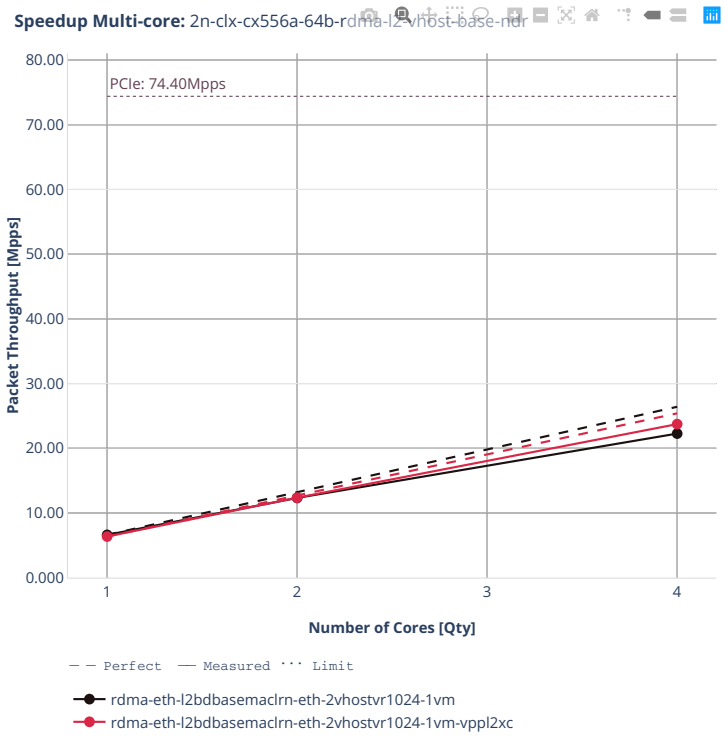

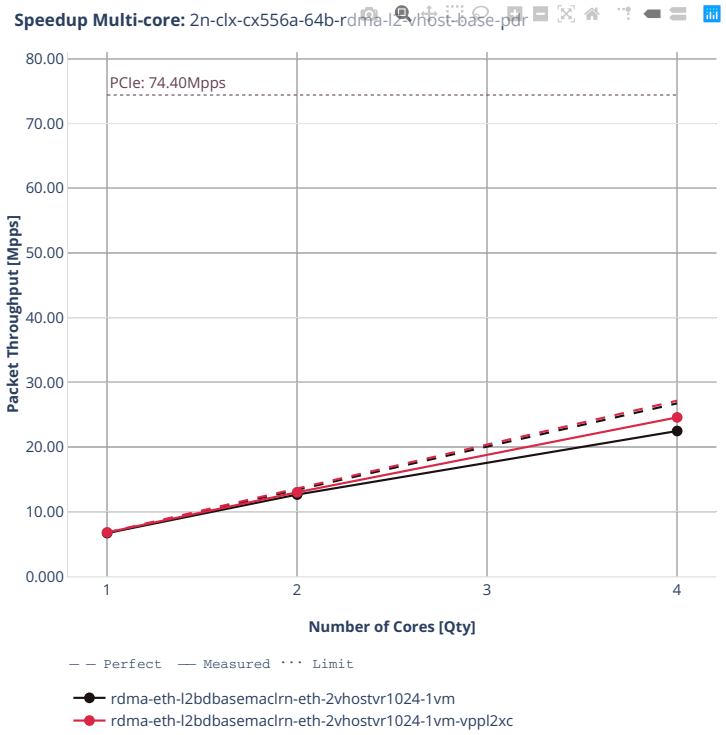
Speedup Multi-core: 2n-clx-cx556a-64b-rdma-l2-vhost-base-pdr
PCIe: 74.40Mpps
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
rdma-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2n-clx-e810cq

64b-vhost

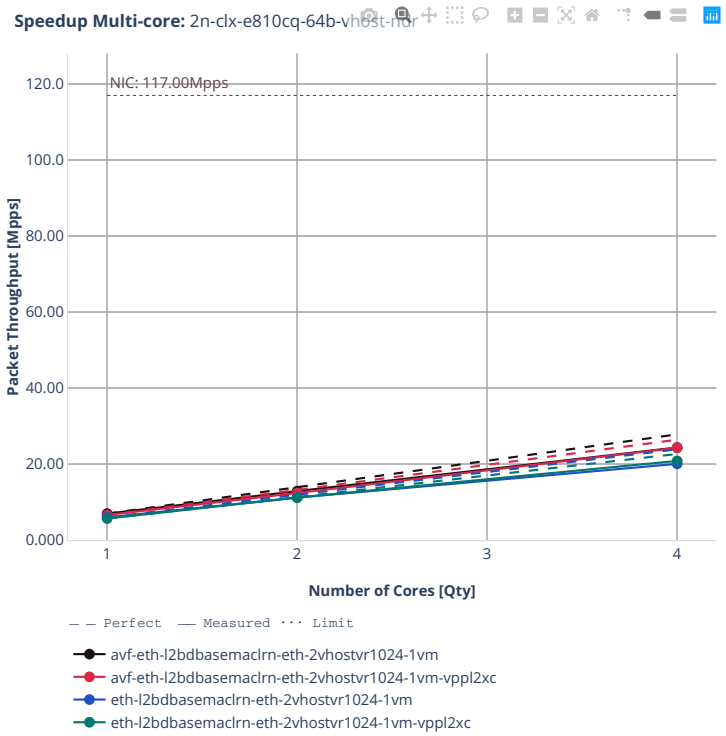

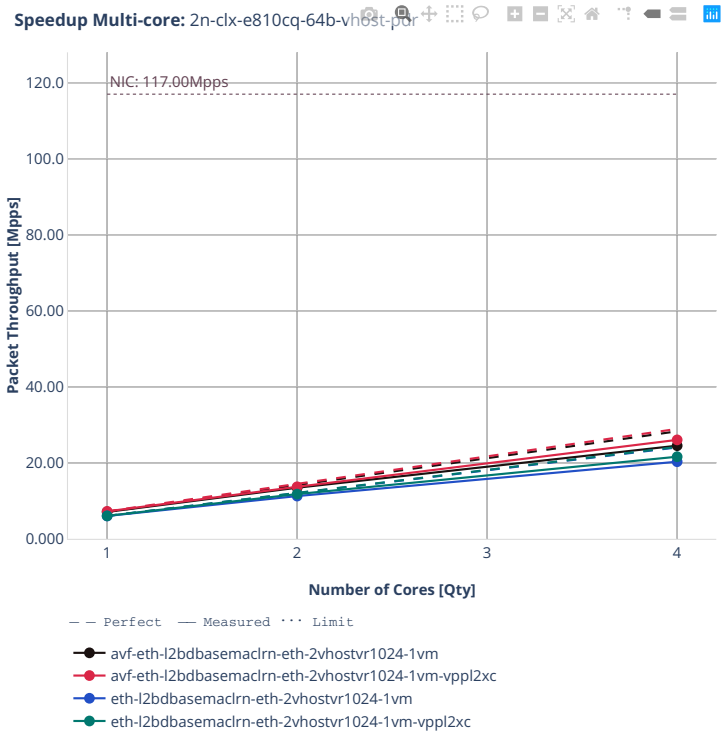
Speedup Multi-core: 2n-clx-e810cq-64b-vhost-pdr
NIC: 117.00Mpps
120.0
100.0
80.00
60.00
40.00
20.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2n-zn2-xxv710

64b-vhost-base-testpmd

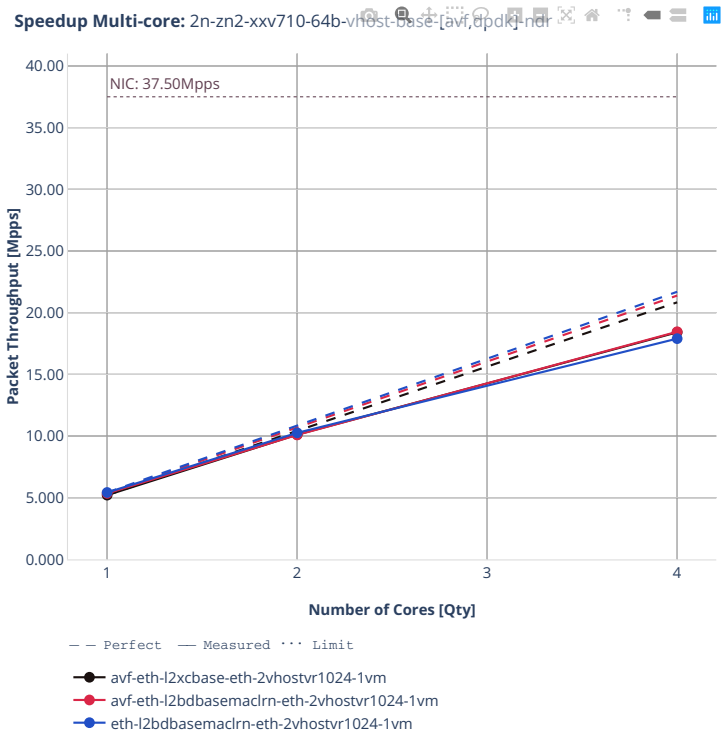

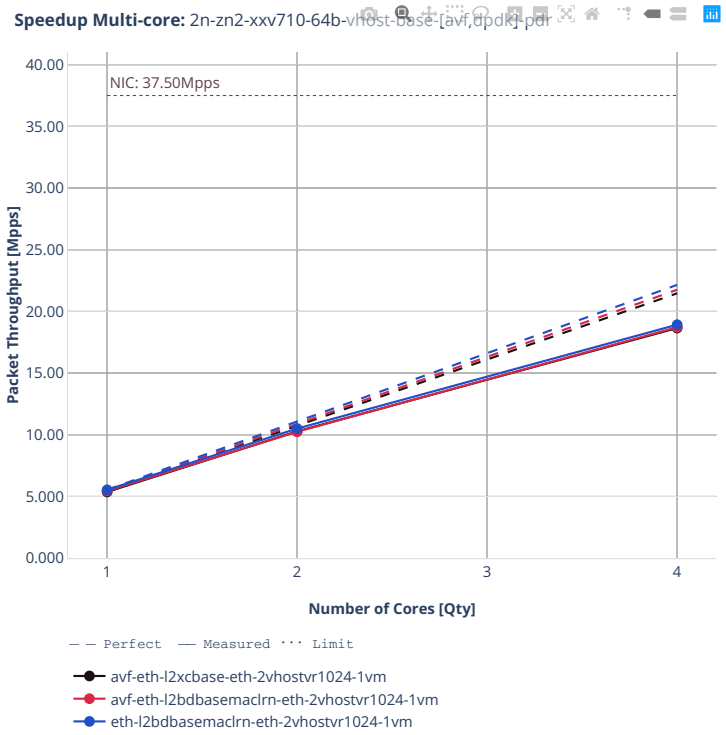
Speedup Multi-core: 2n-zn2-xxv710-64b-vhost-base-[avf,dpdk]-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2xcbase-eth-2vhostvr1024-1vm
avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm

64b-vhost-base-vpp

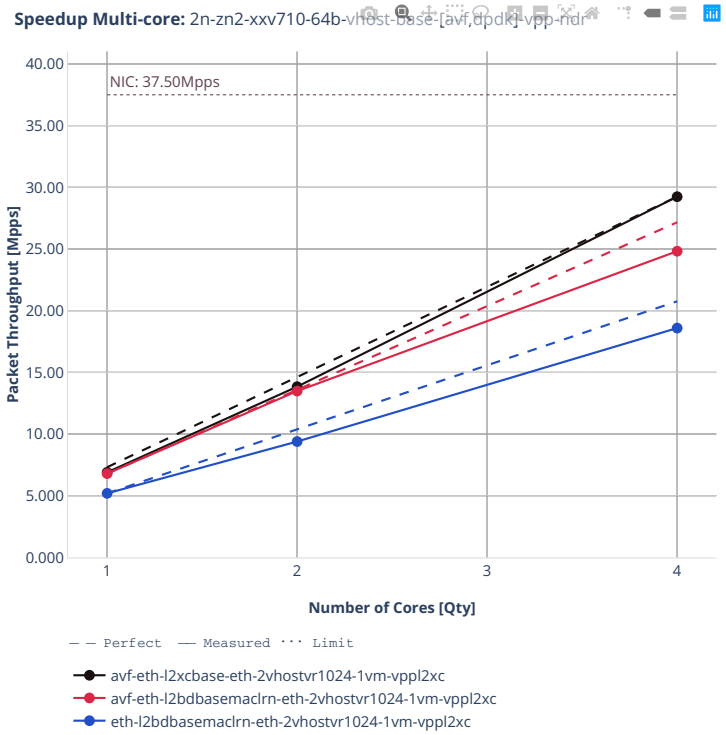

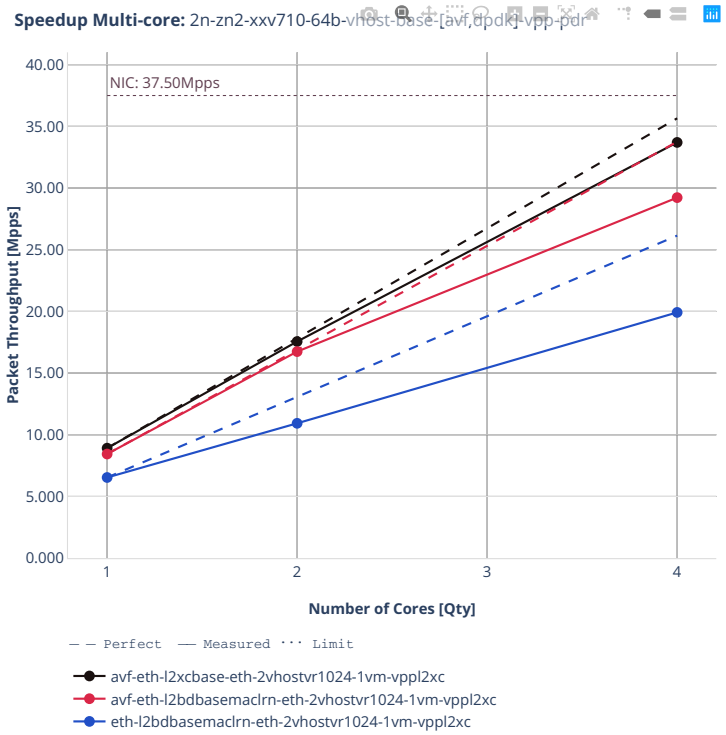
Speedup Multi-core: 2n-zn2-xxv710-64b-vhost-base-[avf,dpdk]-vpp-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2n-zn2-cx556a

64b-vhost-base-rdma-core

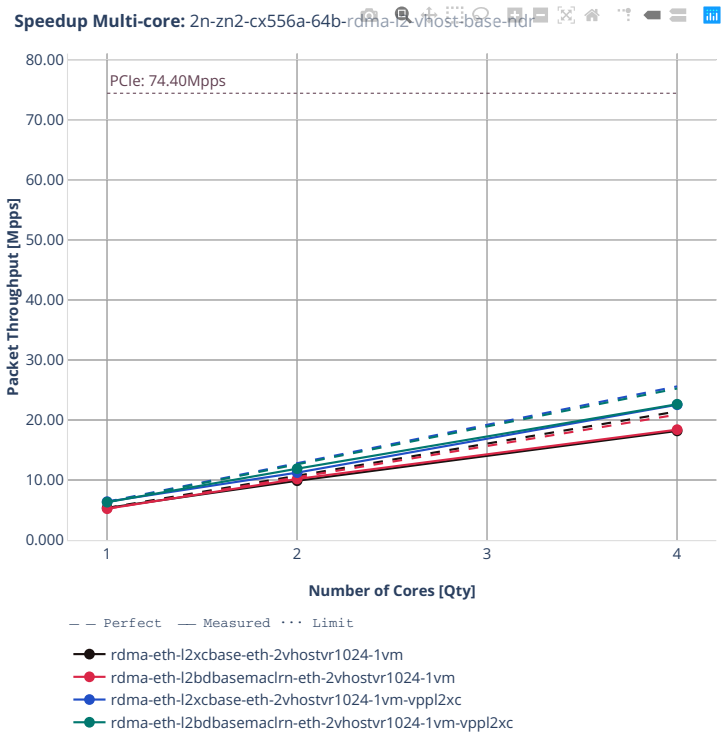

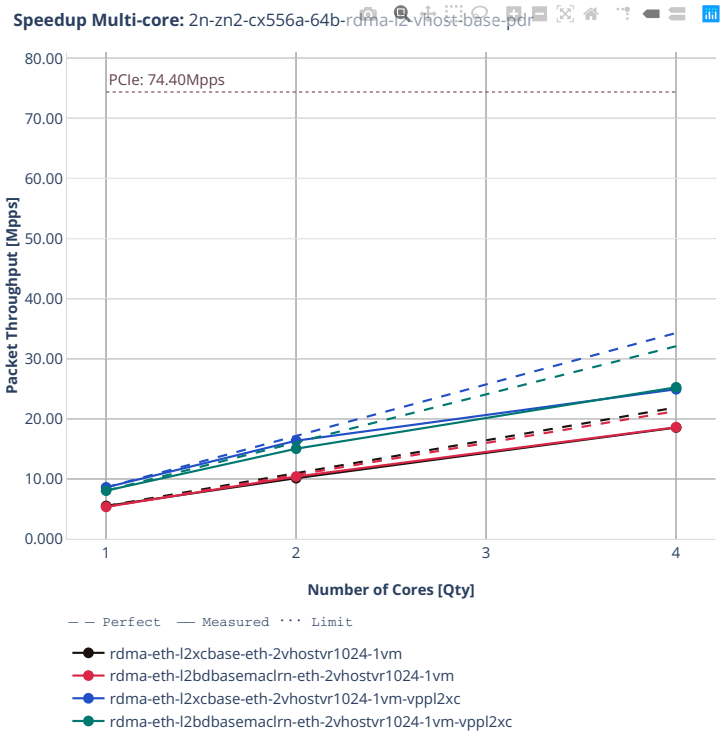
Speedup Multi-core: 2n-zn2-cx556a-64b-rdma-l2-vhost-base-pdr
PCIe: 74.40Mpps
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-eth-l2xcbase-eth-2vhostvr1024-1vm
rdma-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
rdma-eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
rdma-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

3n-alt-xl710

64b-vhost-base

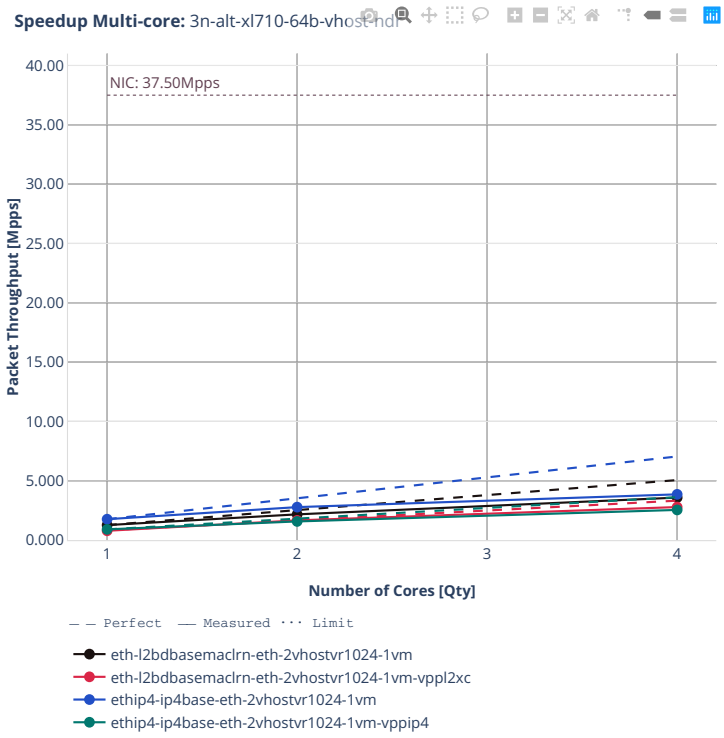

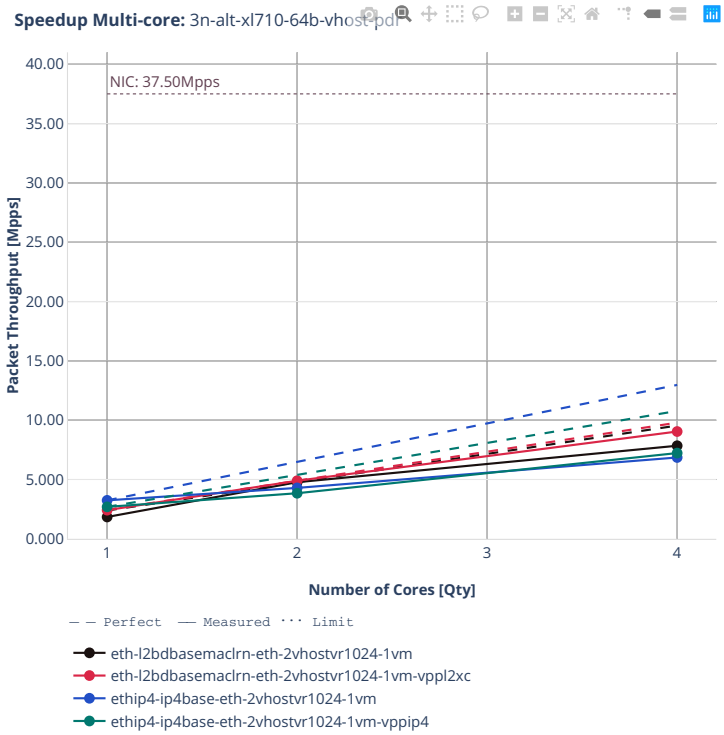
Speedup Multi-core: 3n-alt-xl710-64b-vhost-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
ethip4-ip4base-eth-2vhostvr1024-1vm
ethip4-ip4base-eth-2vhostvr1024-1vm-vppip4

3n-tsh-x520

64b-vhost-base-ixgbe

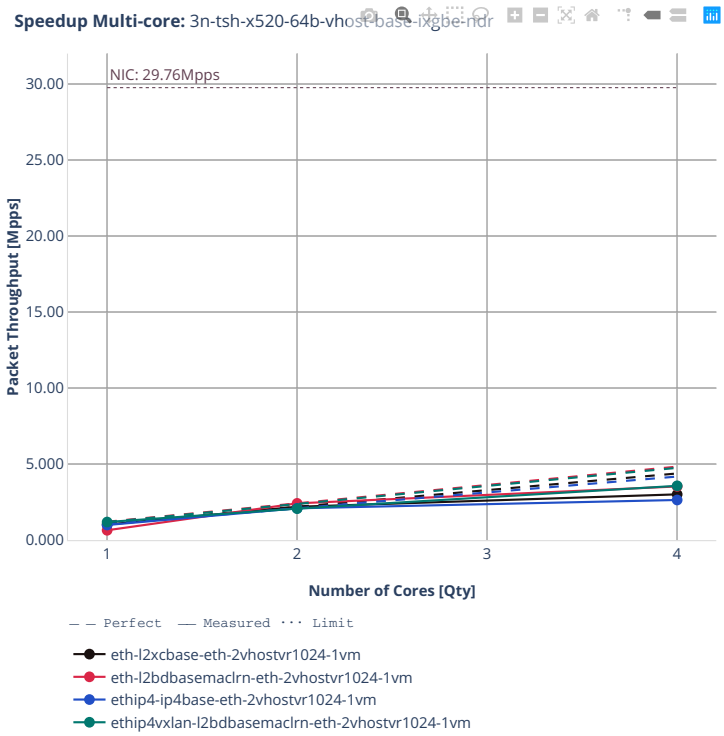

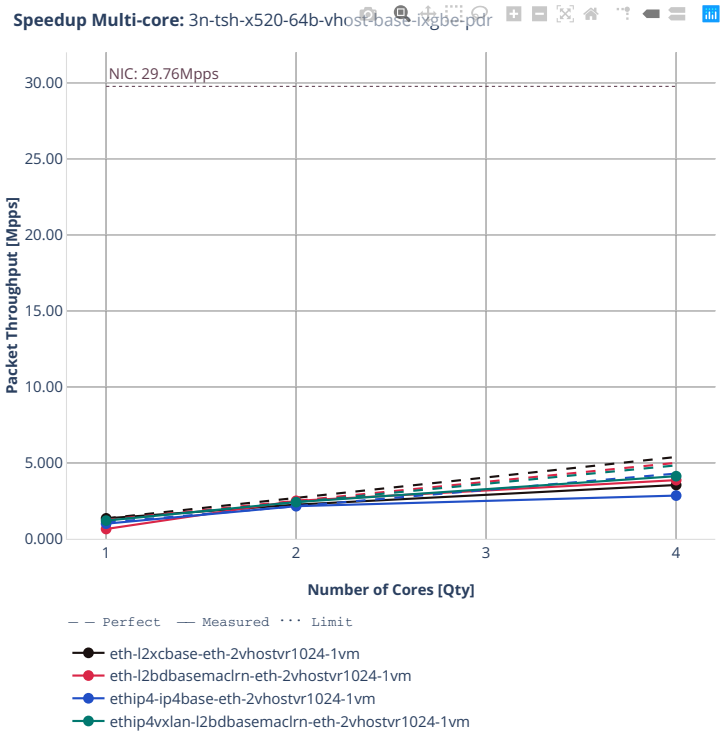
Speedup Multi-core: 3n-tsh-x520-64b-vhost-base-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-eth-2vhostvr1024-1vm
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
ethip4-ip4base-eth-2vhostvr1024-1vm
ethip4vxlan-l2bdbasemaclrn-eth-2vhostvr1024-1vm

64b-vhost-base-ixgbe-vppl2xc

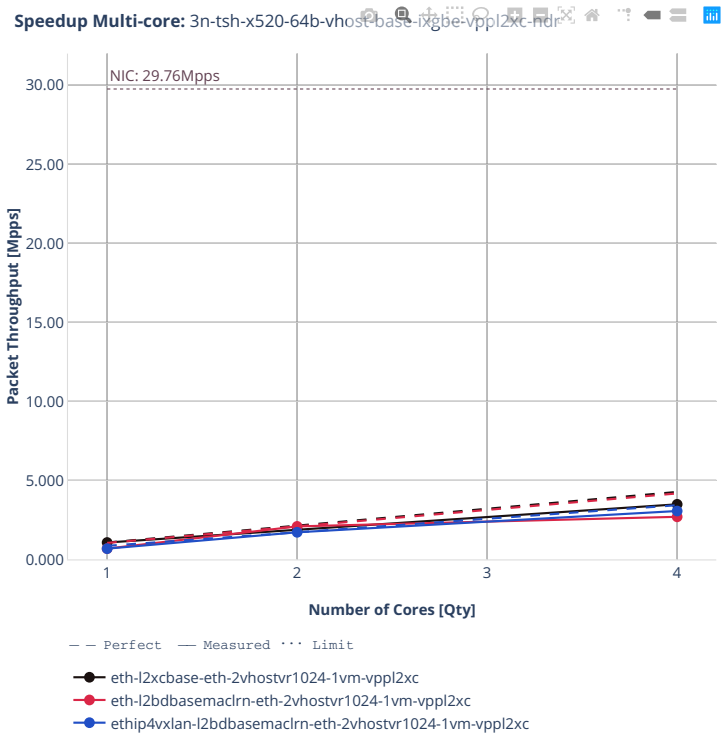

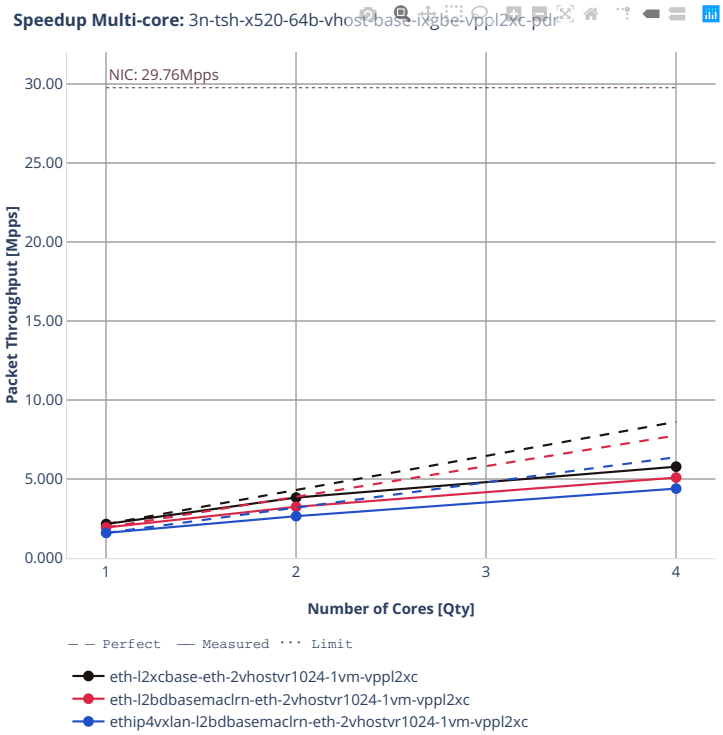
Speedup Multi-core: 3n-tsh-x520-64b-vhost-base-ixgbe-vppl2xc-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc
eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
ethip4vxlan-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

2.4.8 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[137].

[137] https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls2206

2n-icx-xxv710

64b-memif-base

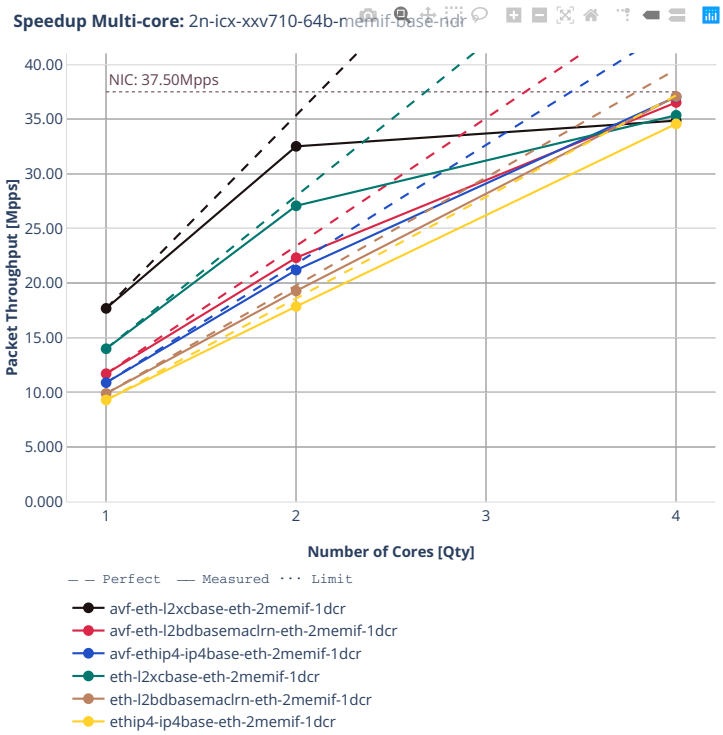

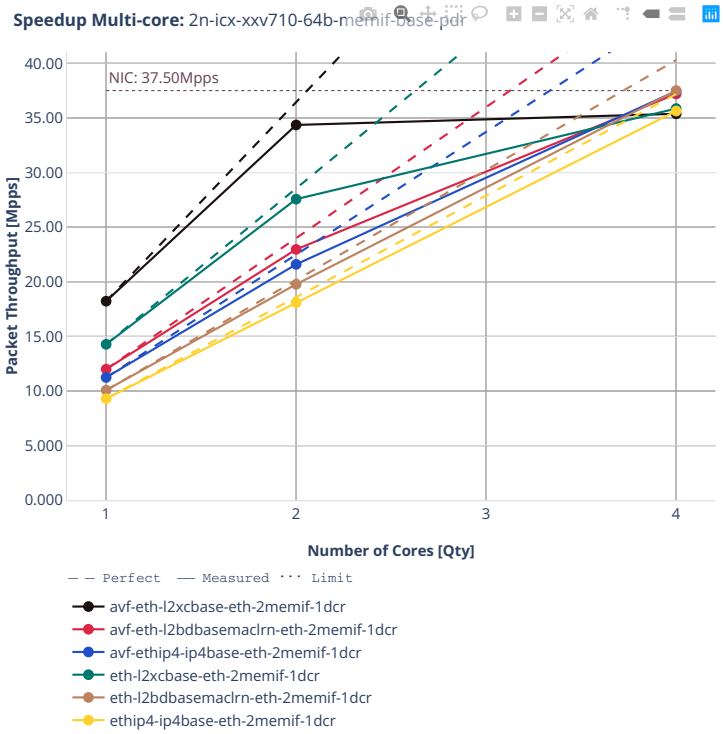
Speedup Multi-core: 2n-icx-xxv710-64b-memif-base-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2xcbase-eth-2memif-1dcr
avf-eth-l2bdbasemaclrn-eth-2memif-1dcr
avf-ethip4-ip4base-eth-2memif-1dcr
eth-l2xcbase-eth-2memif-1dcr
eth-l2bdbasemaclrn-eth-2memif-1dcr
ethip4-ip4base-eth-2memif-1dcr

2n-skx-xxv710

64b-memif-base-avf

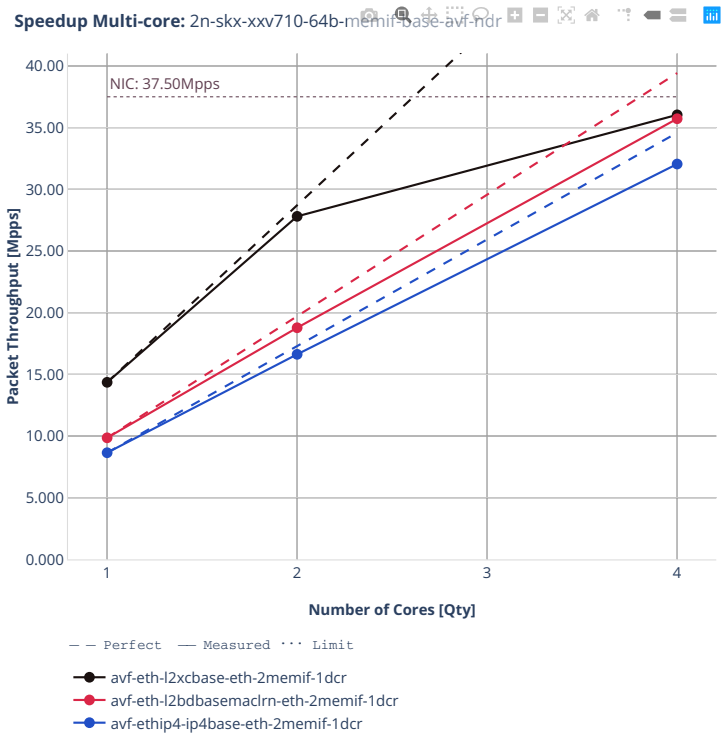

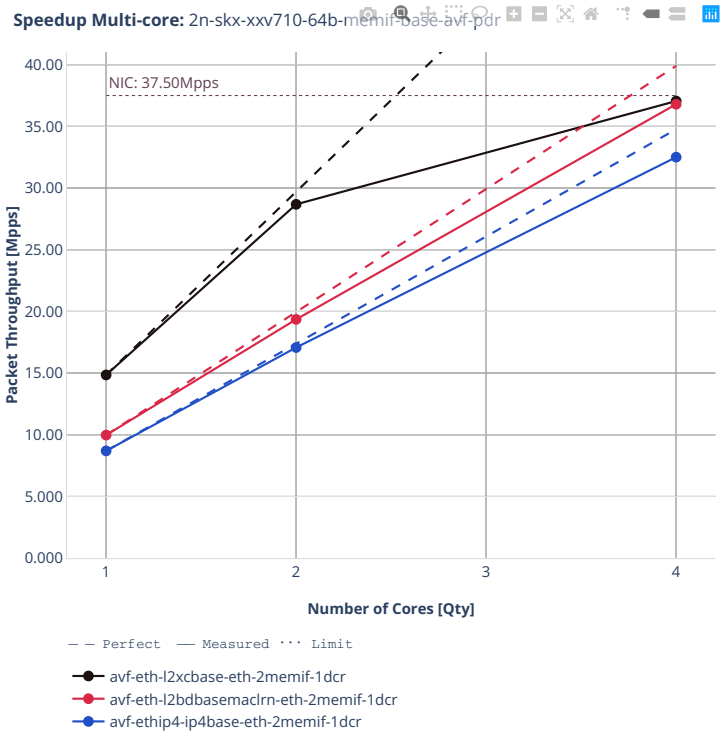
Speedup Multi-core: 2n-skx-xxv710-64b-memif-base-avf-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2xcbase-eth-2memif-1dcr
avf-eth-l2bdbasemaclrn-eth-2memif-1dcr
avf-ethip4-ip4base-eth-2memif-1dcr

64b-memif-base-dpdk

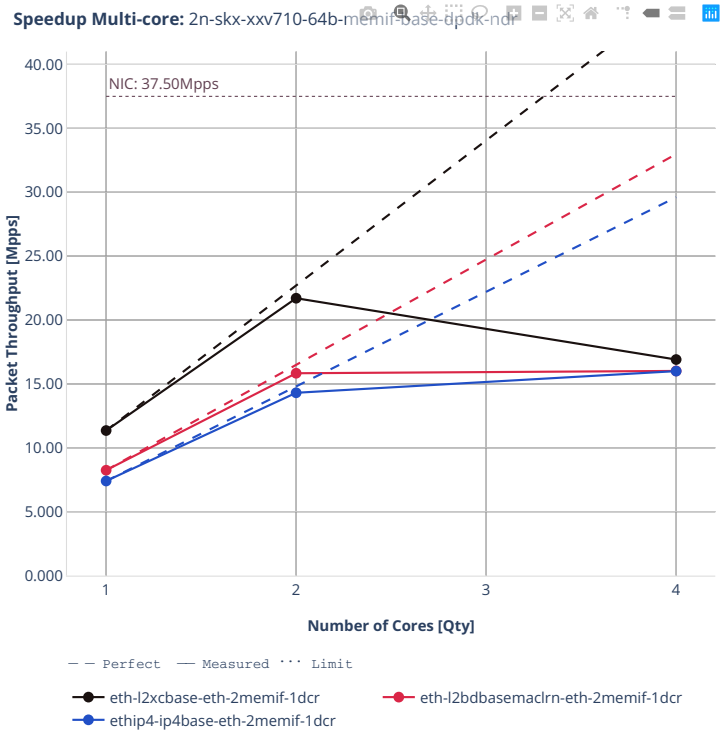

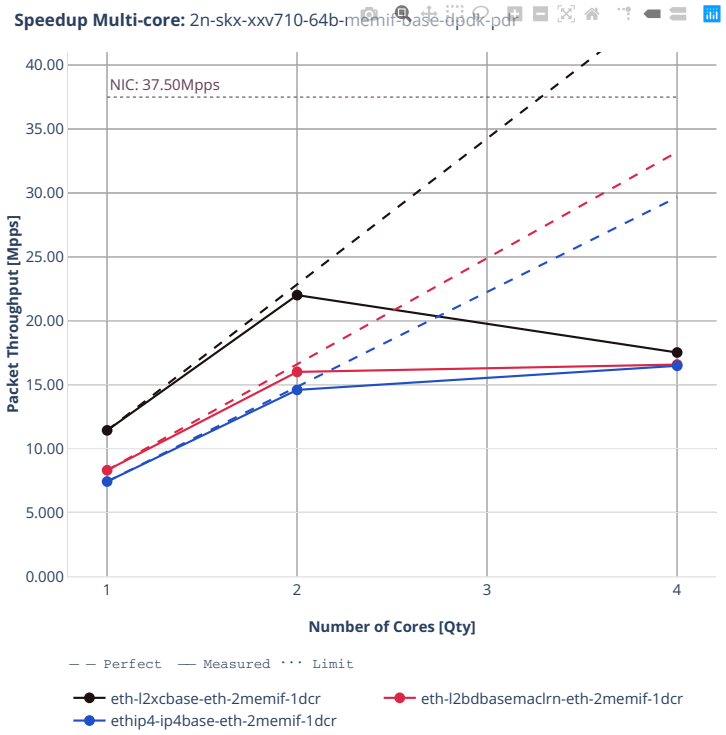
Speedup Multi-core: 2n-skx-xxv710-64b-memif-base-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-eth-2memif-1dcr
eth-l2bdbasemaclrn-eth-2memif-1dcr
ethip4-ip4base-eth-2memif-1dcr

2n-clx-xxv710

64b-memif-base-avf

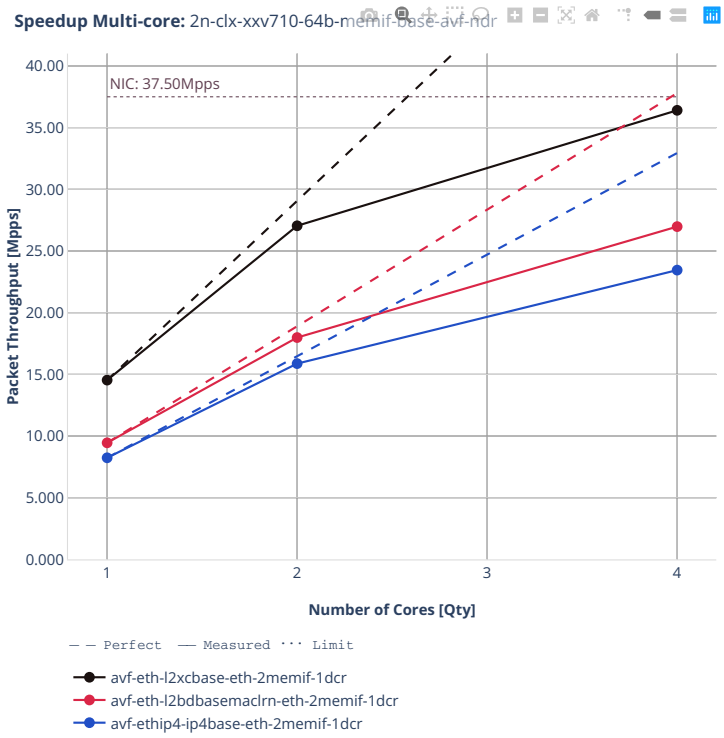

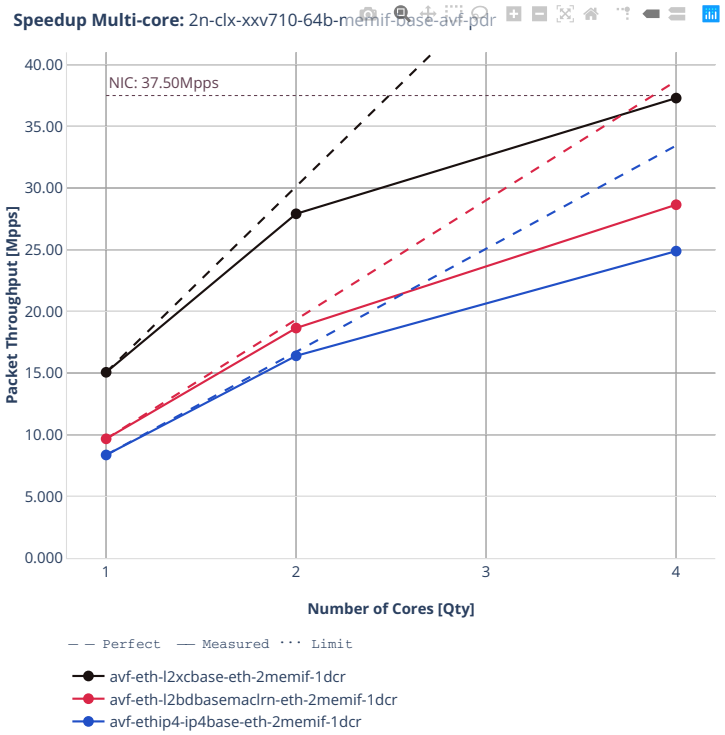
Speedup Multi-core: 2n-clx-xxv710-64b-memif-base-avf-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2xcbase-eth-2memif-1dcr
avf-eth-l2bdbasemaclrn-eth-2memif-1dcr
avf-ethip4-ip4base-eth-2memif-1dcr

64b-memif-base-dpdk

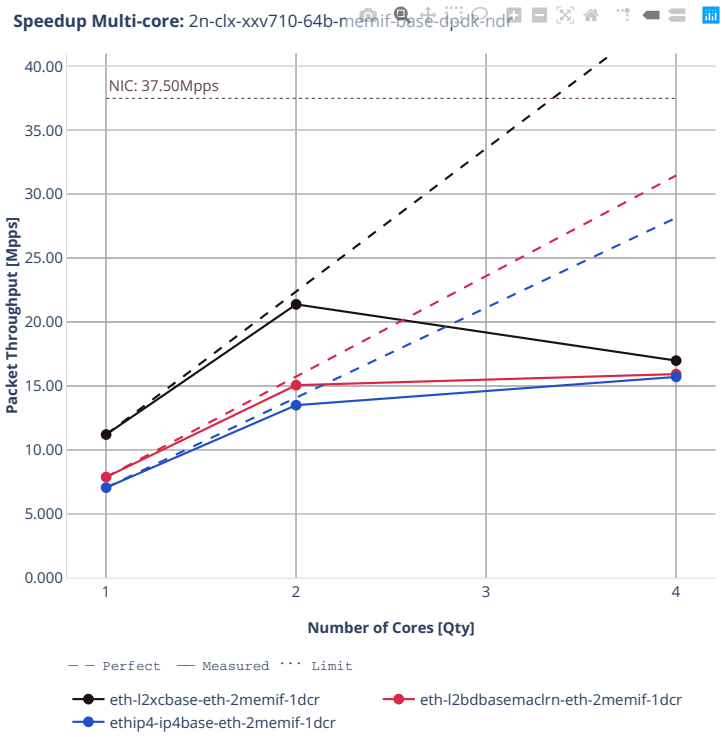

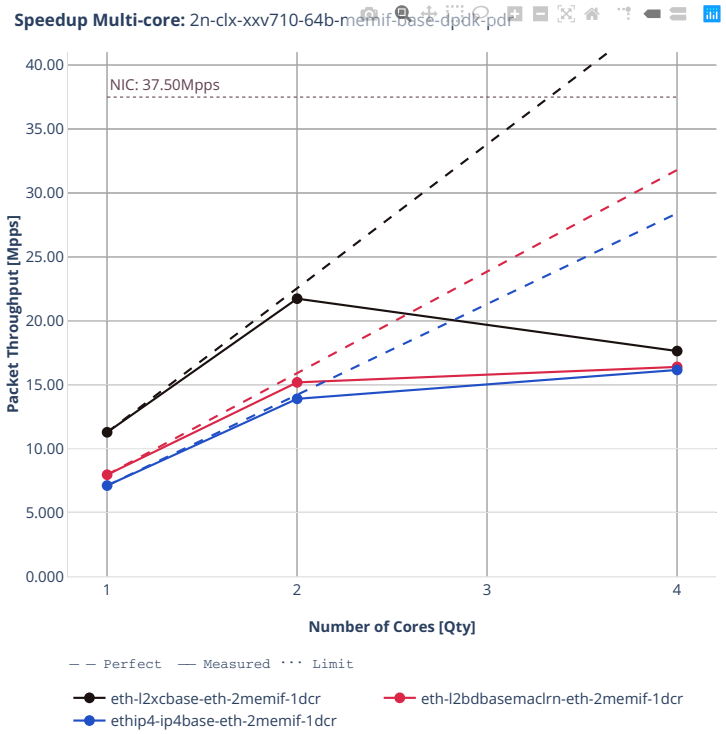
Speedup Multi-core: 2n-clx-xxv710-64b-memif-base-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-eth-2memif-1dcr
eth-l2bdbasemaclrn-eth-2memif-1dcr
ethip4-ip4base-eth-2memif-1dcr

2n-clx-cx556a

64b-memif-base-rdma-core

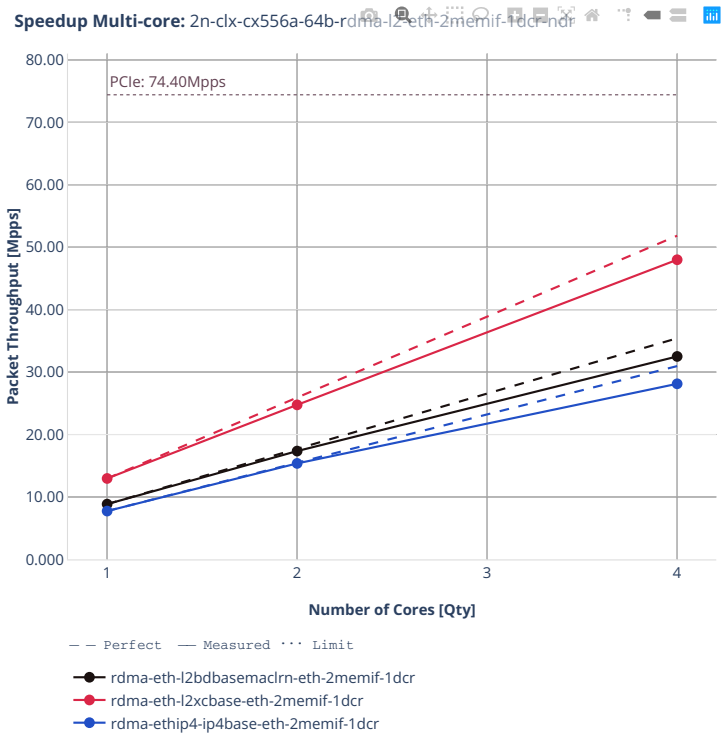

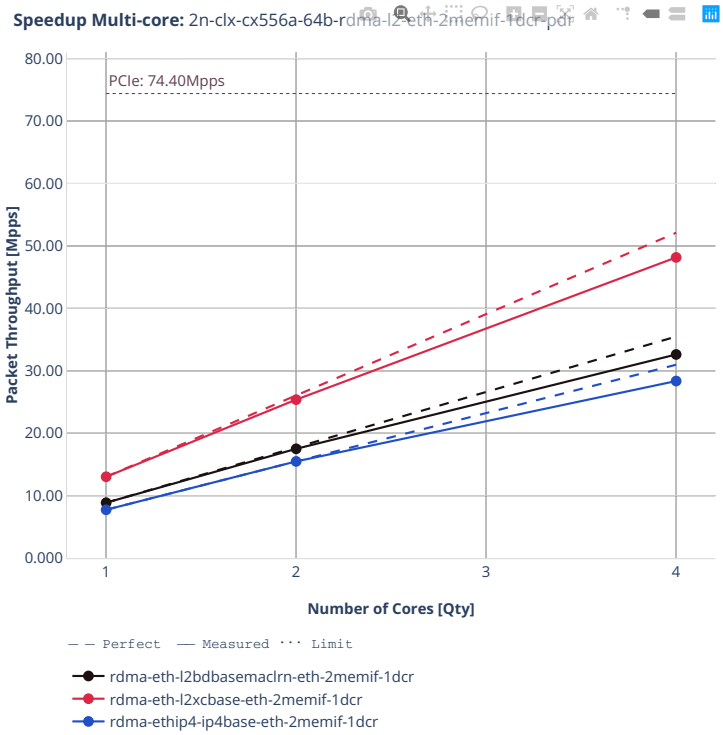
Speedup Multi-core: 2n-clx-cx556a-64b-rdma-l2-eth-2memif-1dcr-pdr
PCIe: 74.40Mpps
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-eth-l2bdbasemaclrn-eth-2memif-1dcr
rdma-eth-l2xcbase-eth-2memif-1dcr
rdma-ethip4-ip4base-eth-2memif-1dcr

2n-clx-e810cq

64b-memif-base

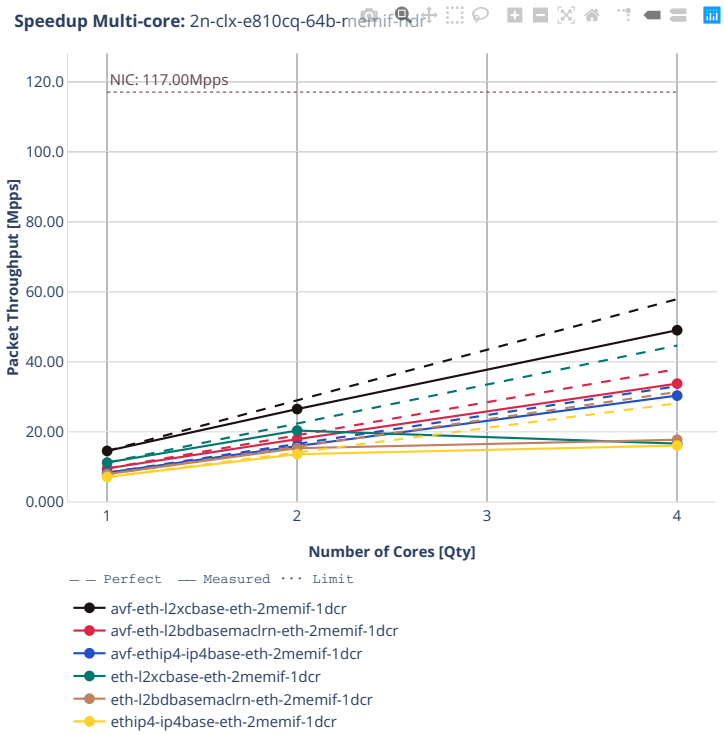

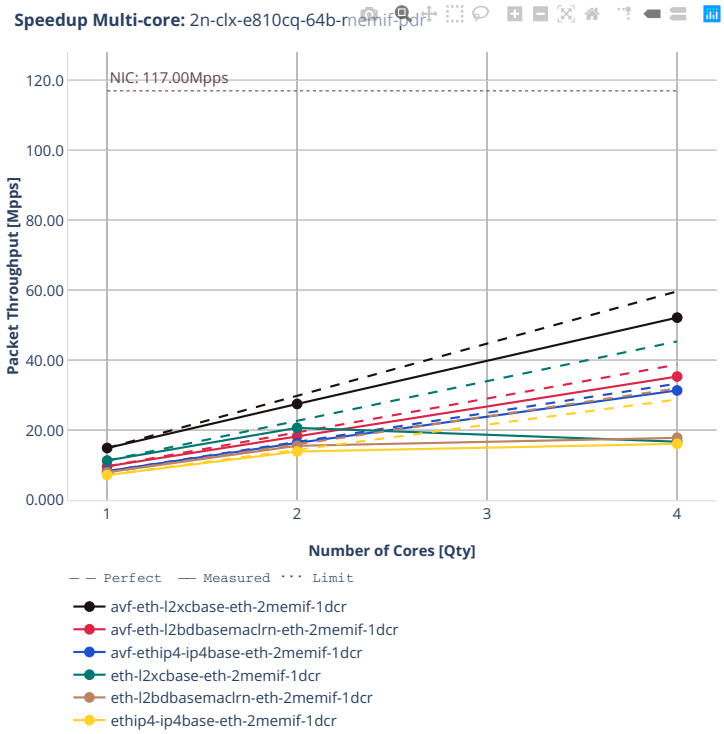
Speedup Multi-core: 2n-clx-e810cq-64b-memif-pdr
NIC: 117.00Mpps
120.0
100.0
80.00
60.00
40.00
20.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2xcbase-eth-2memif-1dcr
avf-eth-l2bdbasemaclrn-eth-2memif-1dcr
avf-ethip4-ip4base-eth-2memif-1dcr
eth-l2xcbase-eth-2memif-1dcr
eth-l2bdbasemaclrn-eth-2memif-1dcr
ethip4-ip4base-eth-2memif-1dcr

2n-tx2-xl710

64b-memif-base-dpdk

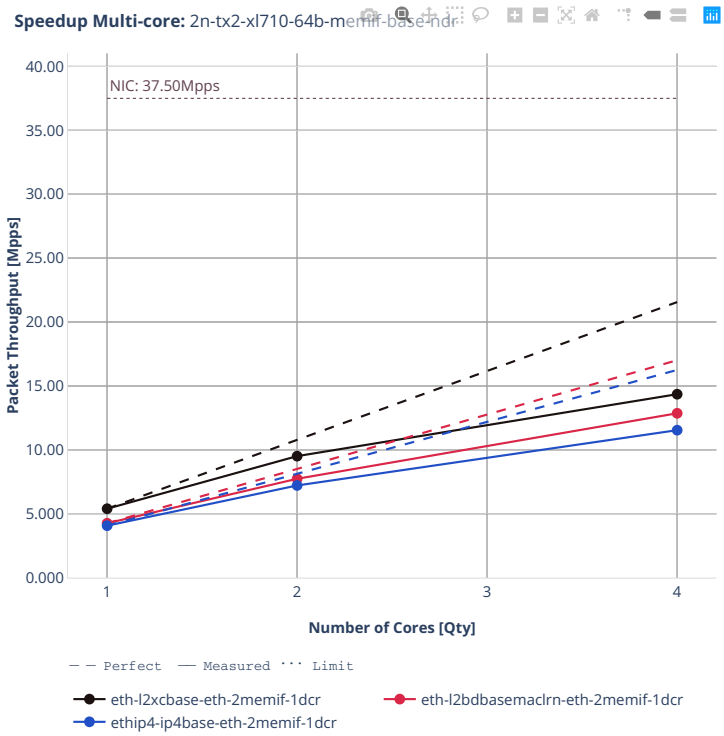

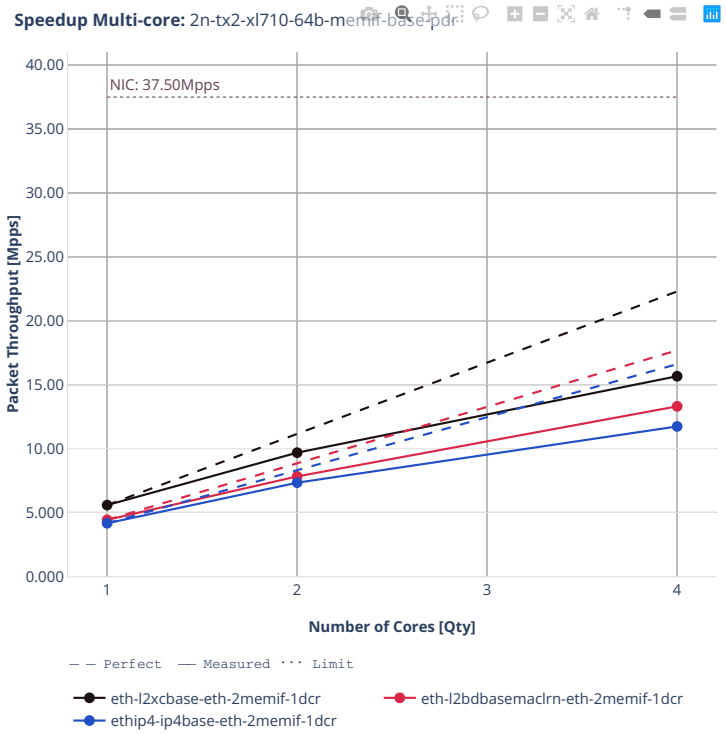
Speedup Multi-core: 2n-tx2-xl710-64b-memif-base-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-eth-2memif-1dcr
eth-l2bdbasemaclrn-eth-2memif-1dcr
ethip4-ip4base-eth-2memif-1dcr

2n-zn2-xxv710

64b-memif-base-avf

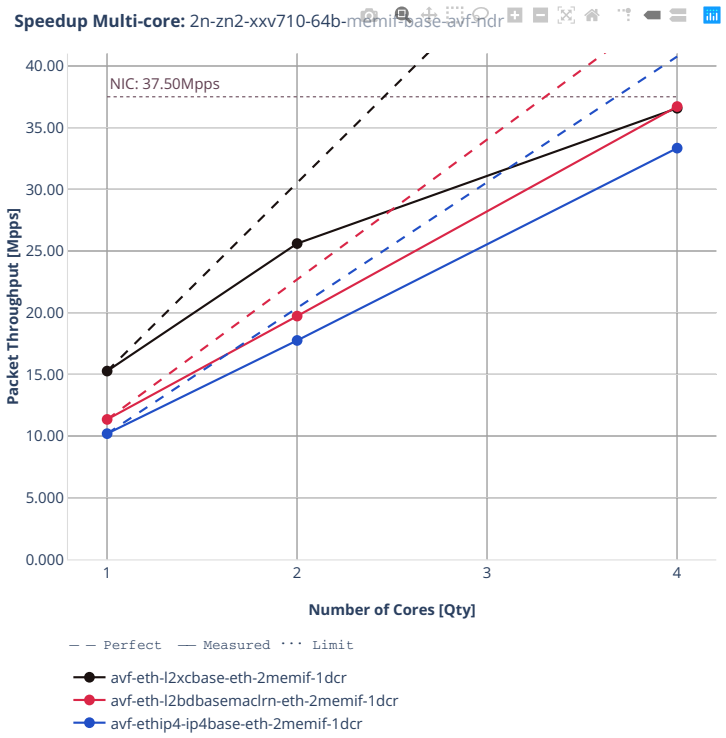

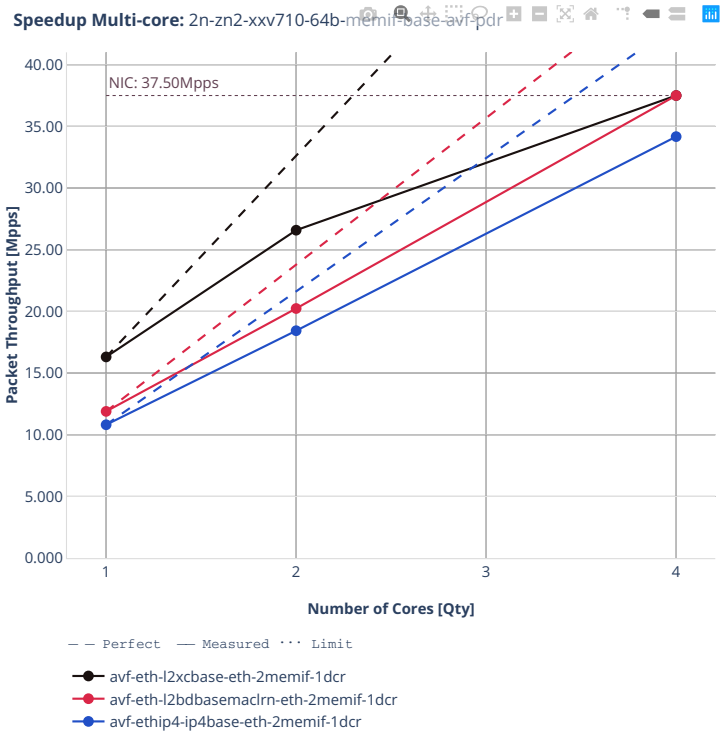
Speedup Multi-core: 2n-zn2-xxv710-64b-memif-base-avf-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-eth-l2xcbase-eth-2memif-1dcr
avf-eth-l2bdbasemaclrn-eth-2memif-1dcr
avf-ethip4-ip4base-eth-2memif-1dcr

64b-memif-base-dpdk

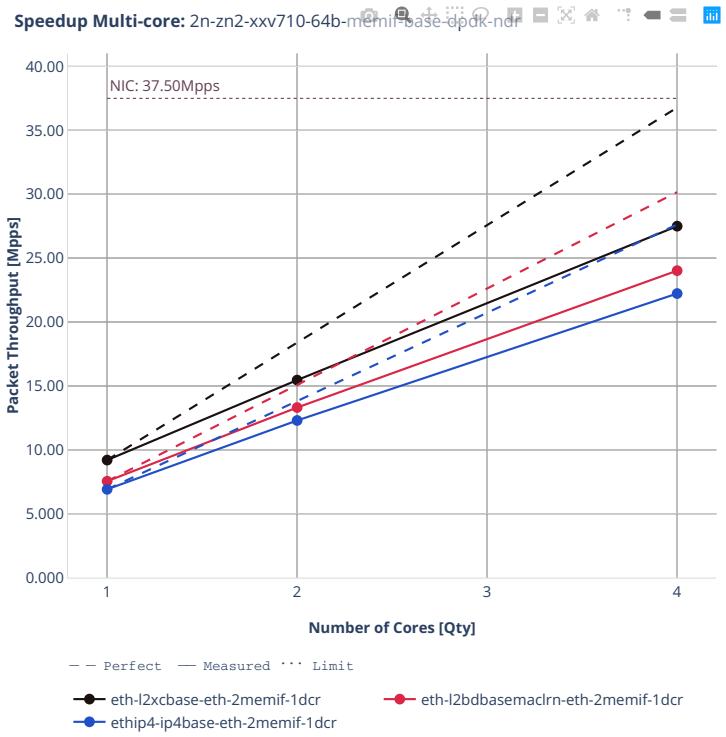

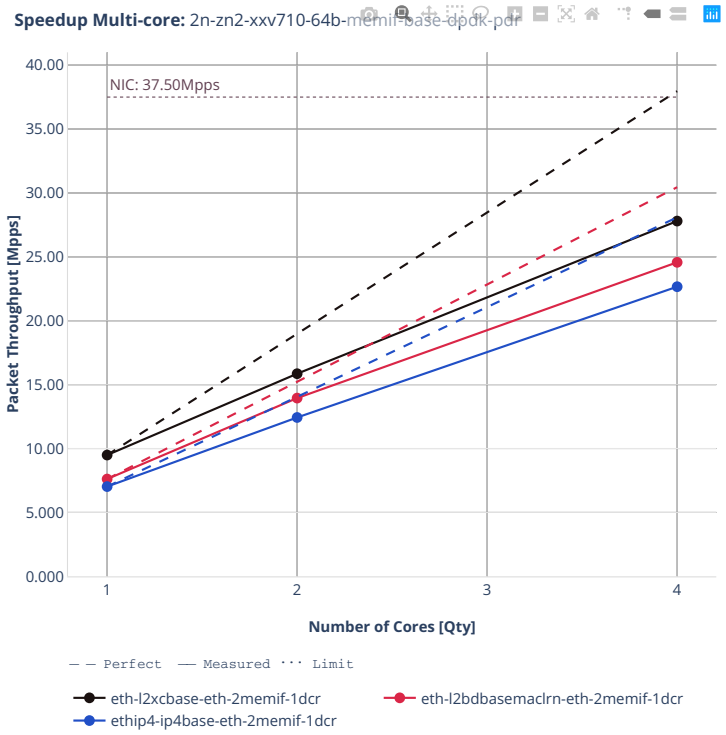
Speedup Multi-core: 2n-zn2-xxv710-64b-memif-base-dpdk-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-eth-2memif-1dcr
eth-l2bdbasemaclrn-eth-2memif-1dcr
ethip4-ip4base-eth-2memif-1dcr

2n-zn2-cx556a

64b-memif-base-rdma-core

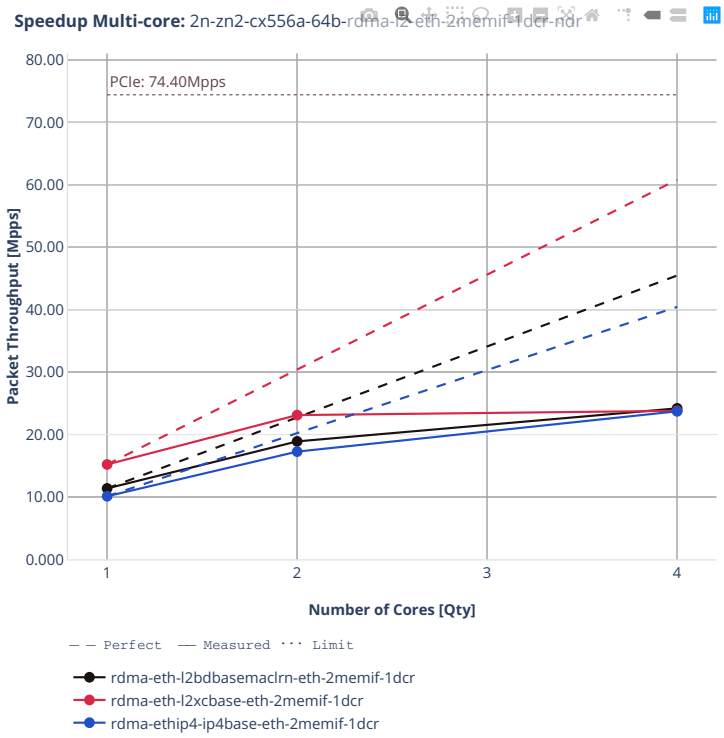

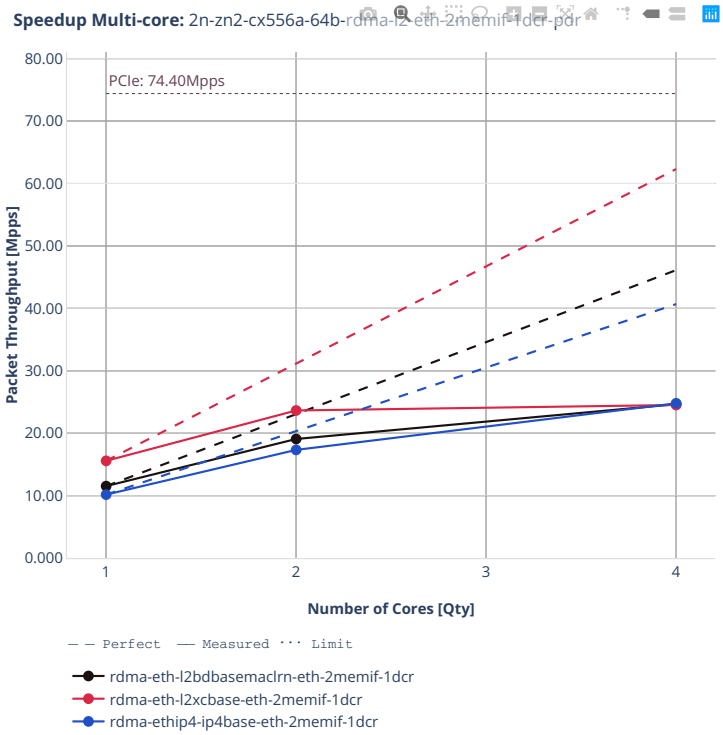
Speedup Multi-core: 2n-zn2-cx556a-64b-rdma-l2-eth-2memif-1dcr-pdr
PCIe: 74.40Mpps
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
rdma-eth-l2bdbasemaclrn-eth-2memif-1dcr
rdma-eth-l2xcbase-eth-2memif-1dcr
rdma-ethip4-ip4base-eth-2memif-1dcr

2.4.9 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[138].

[138] https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls2206

3n-icx-xxv710

1518b-2t1c-ipsec-aes128cbc-ip4routing-base-scale-sw-avf

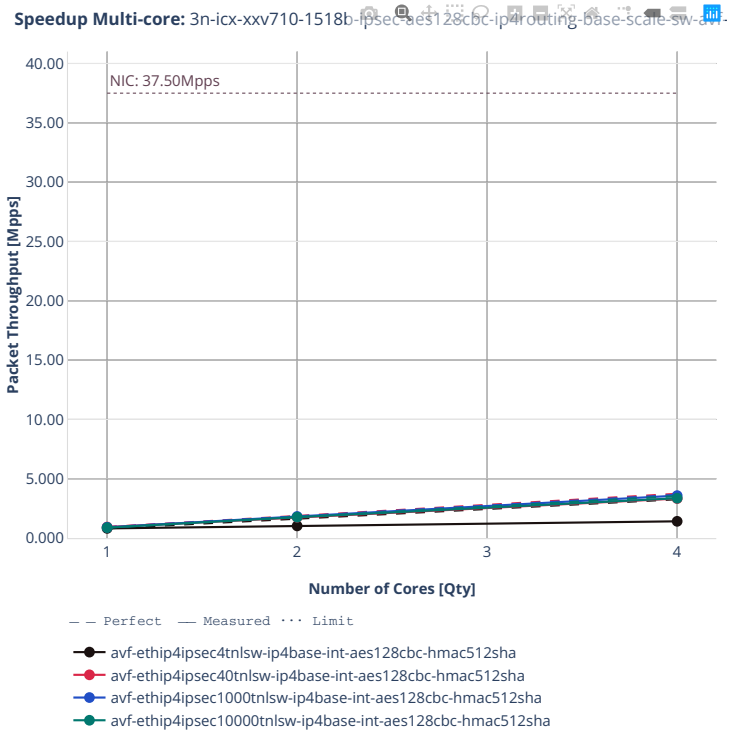

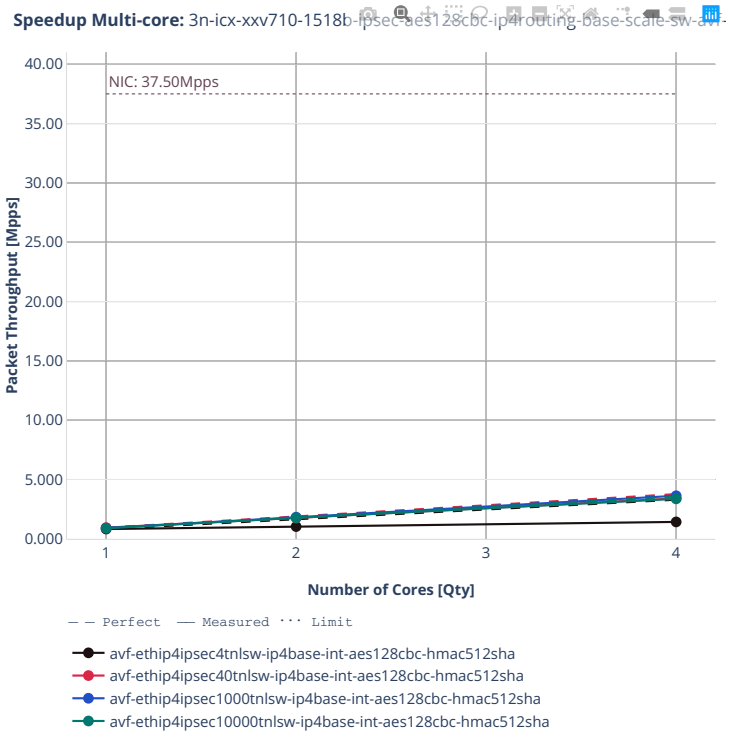
Speedup Multi-core: 3n-icx-xxv710-1518b-ipsec-aes128cbc-ip4routing-base-scale-sw-avf-
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
avf-ethip4ipsec40tnlsw-ip4base-int-aes128cbc-hmac512sha
avf-ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha
avf-ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha

1518b-2t1c-ipsec-aes256gcm-ip4routing-base-scale-sw-avf

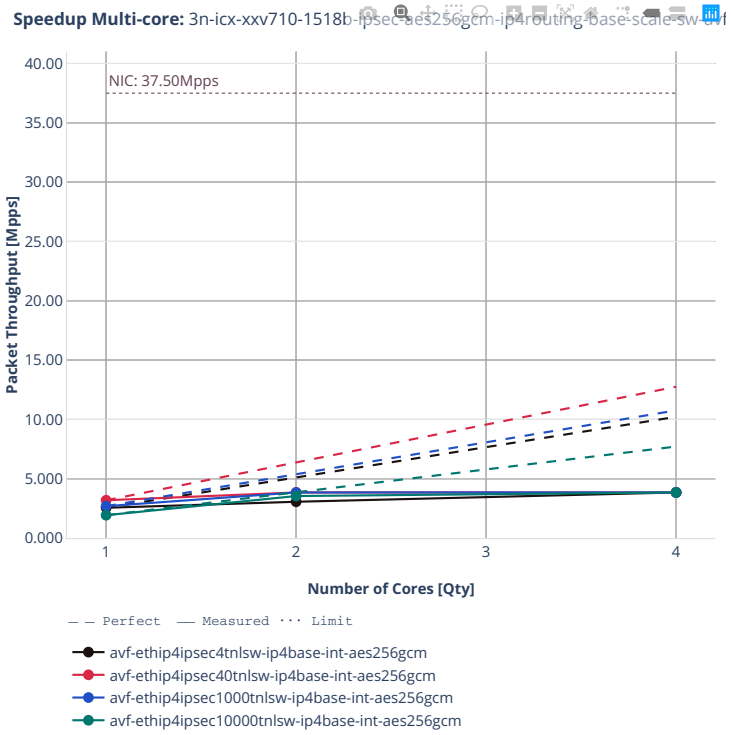

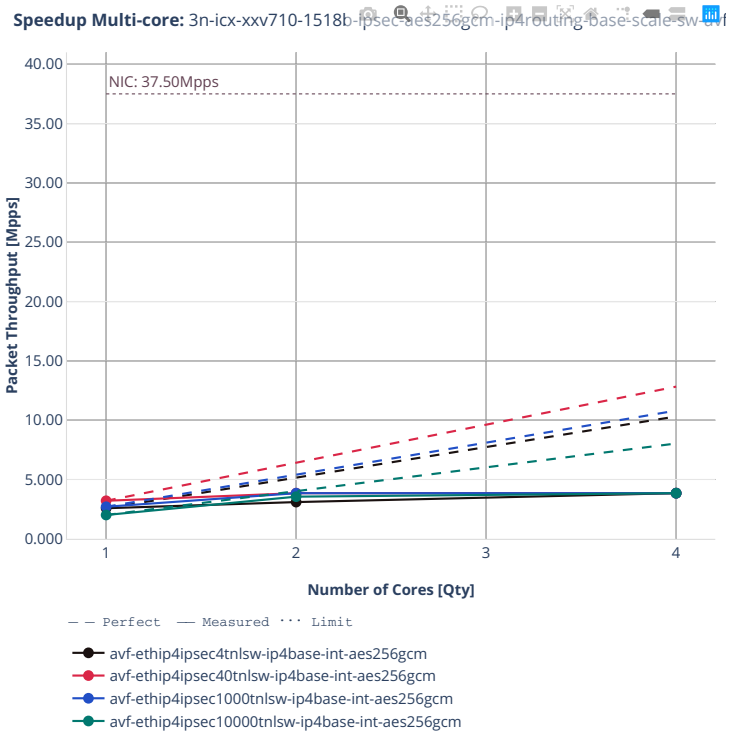
Speedup Multi-core: 3n-icx-xxv710-1518b-ipsec-aes256gcm-ip4routing-base-scale-sw-ull
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4ipsec4tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec40tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

1518b-2t1c-ipsec-ip4routing-base-scale-sw-dpdk

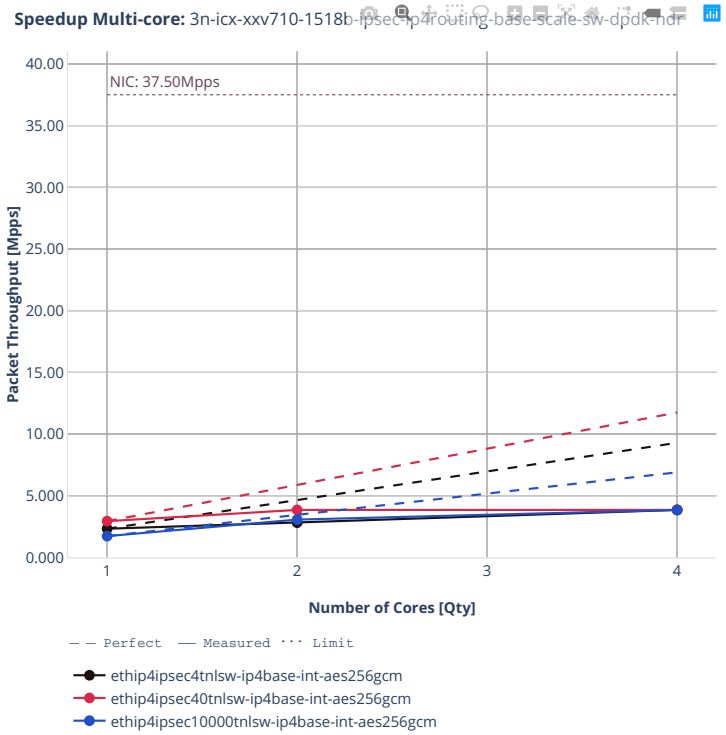

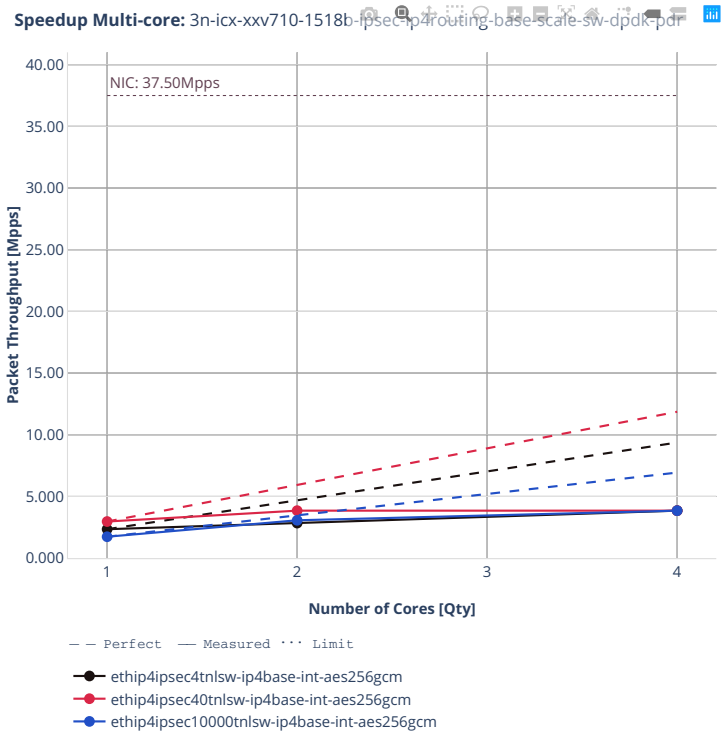
Speedup Multi-core: 3n-icx-xxv710-1518b-ipsec-ip4routing-base-scale-sw-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec4tnlsw-ip4base-int-aes256gcm
ethip4ipsec40tnlsw-ip4base-int-aes256gcm
ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

imix-2t1c-ipsec-aes256gcm-ip4routing-base-scale-sw-avf

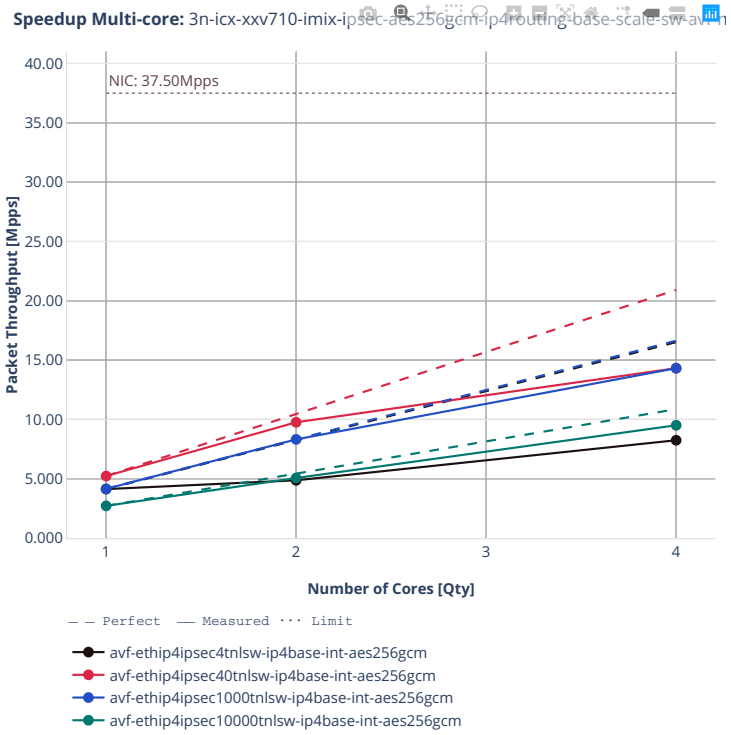

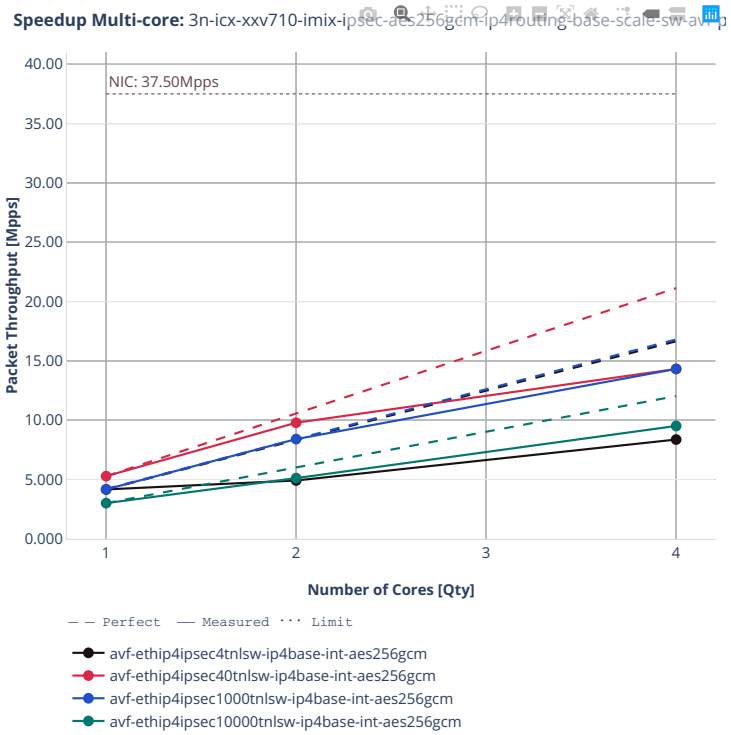
Speedup Multi-core: 3n-icx-xxv710-imix-ipsec-aes256gcm-ip4routing-base-scale-sw-avf
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4ipsec4tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec40tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

imix-2t1c-ipsec-ip4routing-base-scale-sw-dpdk

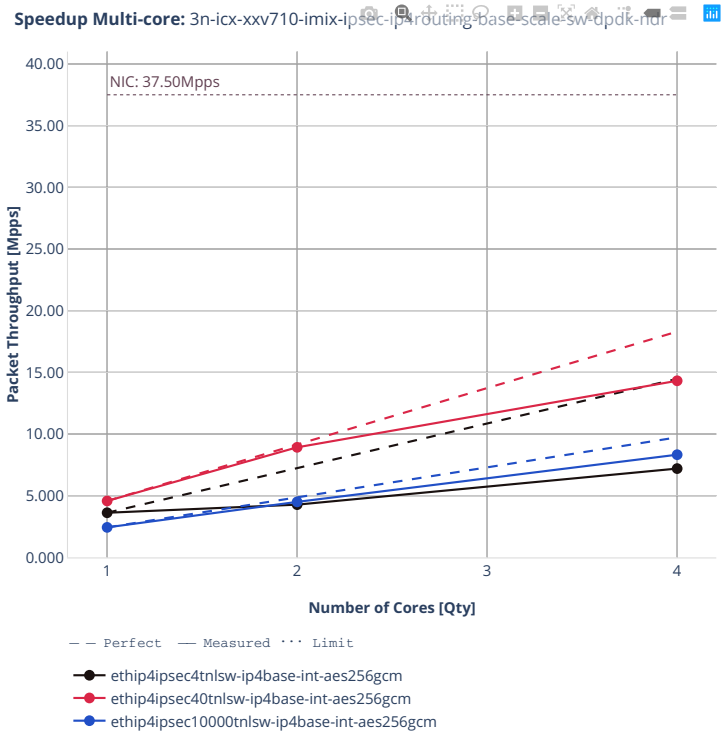

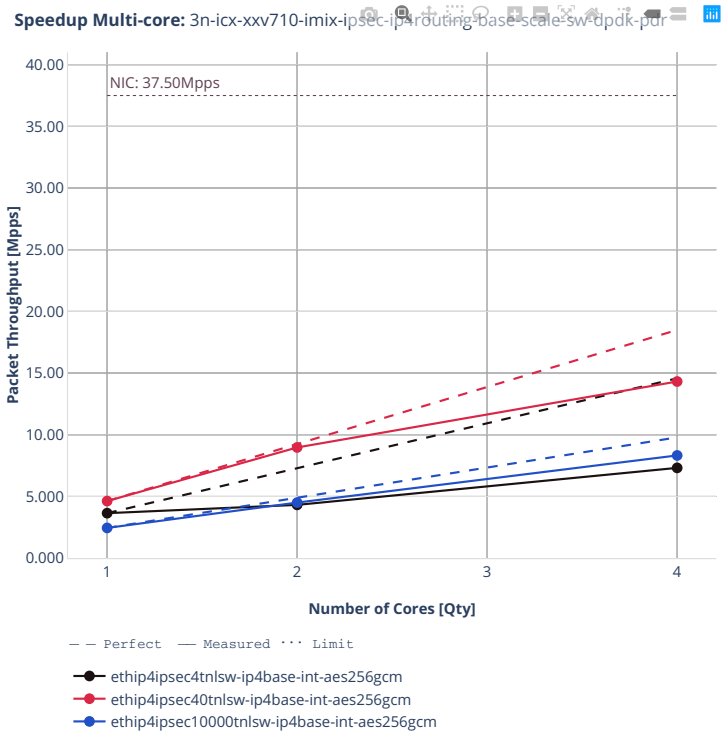
Speedup Multi-core: 3n-icx-xxv710-imix-ipsec-ip4routing-base-scale-sw-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec4tnlsw-ip4base-int-aes256gcm
ethip4ipsec40tnlsw-ip4base-int-aes256gcm
ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

3n-skx-xxv710

imix-ipsec-ip4routing-base-scale-avf

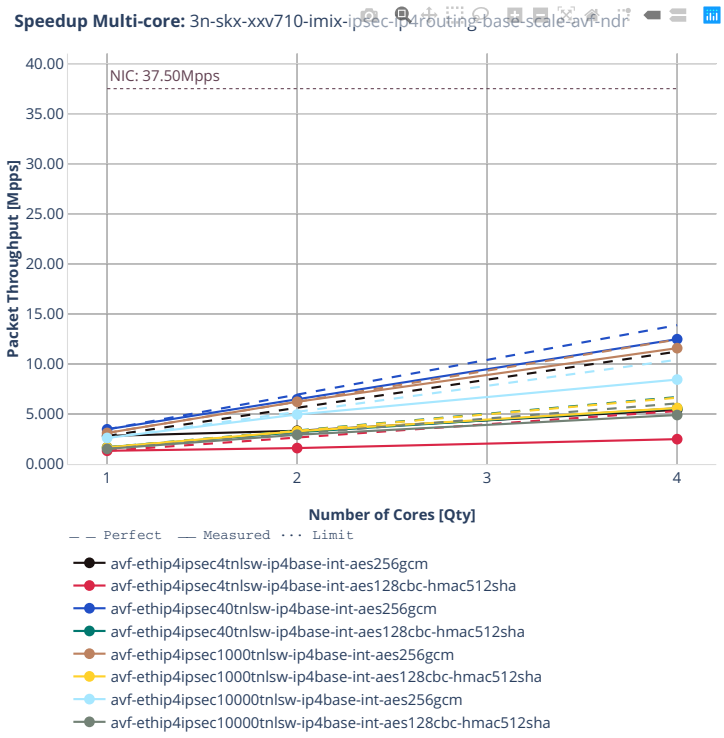

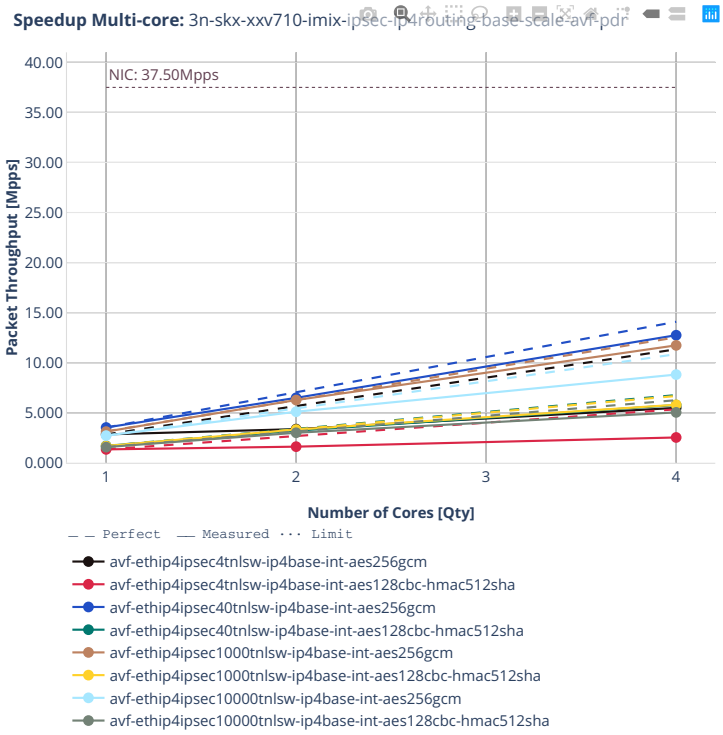
Speedup Multi-core: 3n-skx-xxv710-imix-ipsec-ip4routing-base-scale-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4ipsec4tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
avf-ethip4ipsec40tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec40tnlsw-ip4base-int-aes128cbc-hmac512sha
avf-ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha
avf-ethip4ipsec10000tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha

1518b-ipsec-ip4routing-base-scale-avf

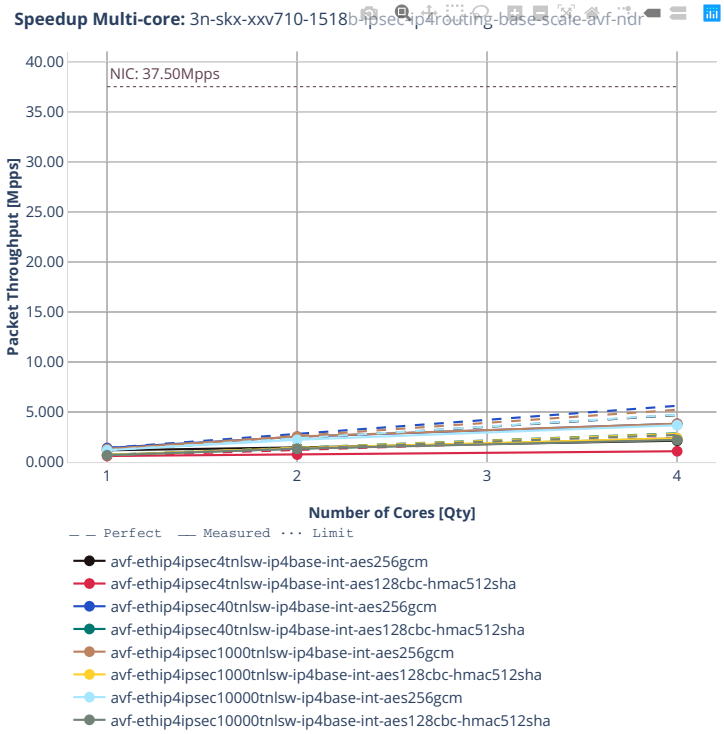

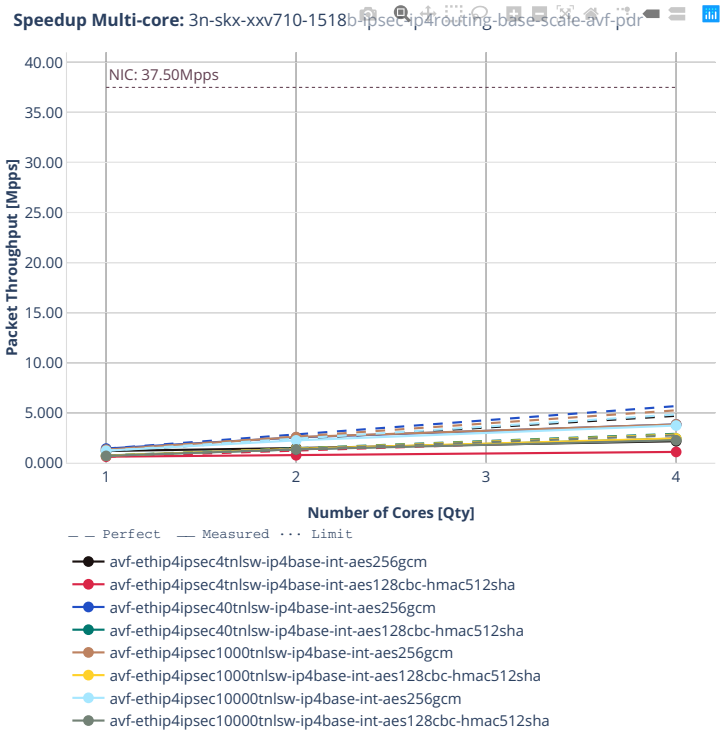
Speedup Multi-core: 3n-skx-xxv710-1518b-ipsec-ip4routing-base-scale-avf-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
avf-ethip4ipsec4tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
avf-ethip4ipsec40tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec40tnlsw-ip4base-int-aes128cbc-hmac512sha
avf-ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha
avf-ethip4ipsec10000tnlsw-ip4base-int-aes256gcm
avf-ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha

imix-ipsec-ip4routing-base-scale-dpdk

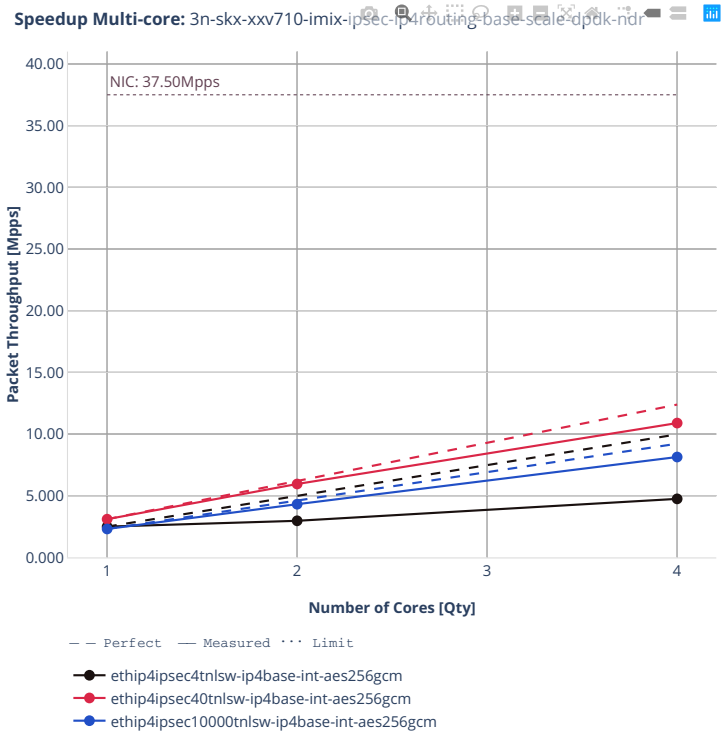

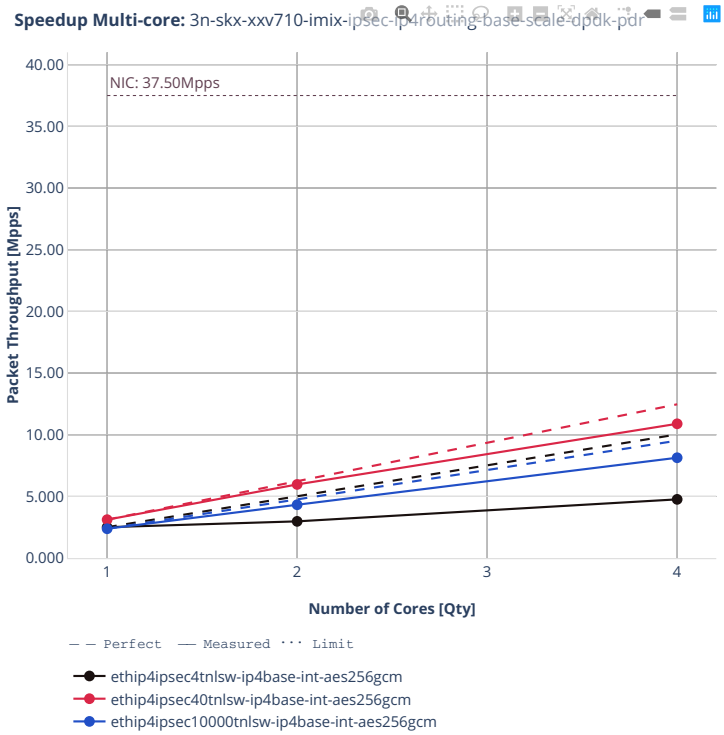
Speedup Multi-core: 3n-skx-xxv710-imix-ipsec-ip4routing-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec4tnlsw-ip4base-int-aes256gcm
ethip4ipsec40tnlsw-ip4base-int-aes256gcm
ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

3n-alt-xl710

1518b-ipsec-aes128cbc-ip4routing-scale-sw

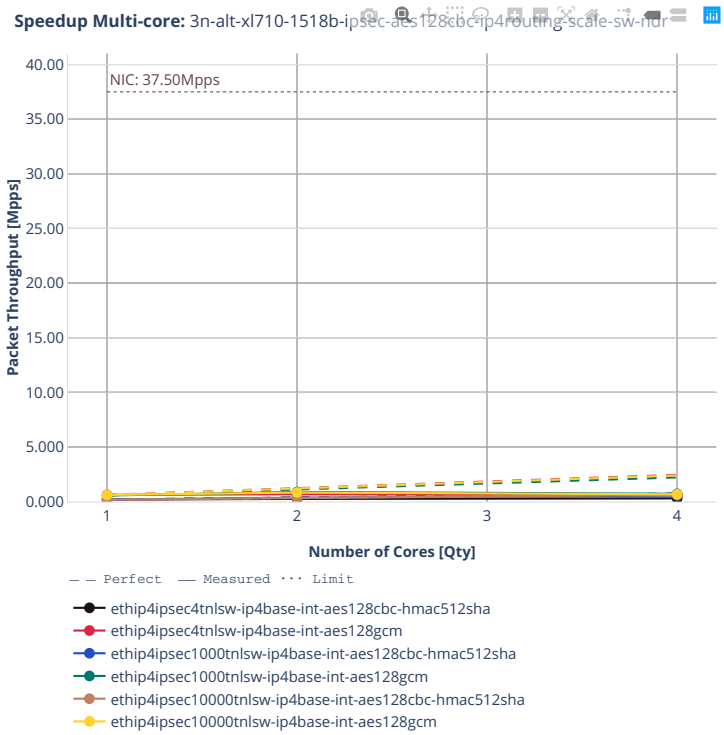

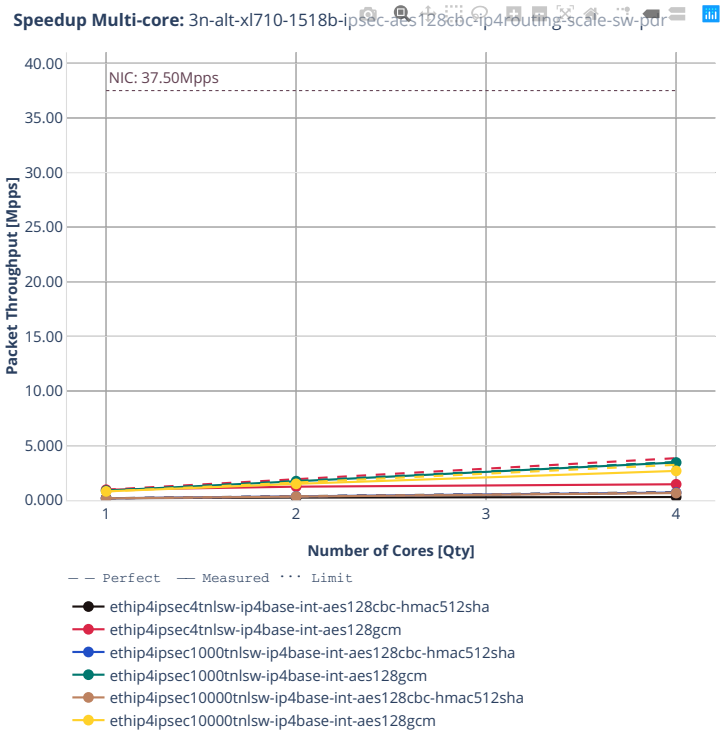
Speedup Multi-core: 3n-alt-xl710-1518b-ipsec-aes128cbc-ip4routing-scale-sw-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
ethip4ipsec4tnlsw-ip4base-int-aes128gcm
ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha
ethip4ipsec1000tnlsw-ip4base-int-aes128gcm
ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha
ethip4ipsec10000tnlsw-ip4base-int-aes128gcm

1518b-ipsec-aes256gcm-ip4routing-scale-sw

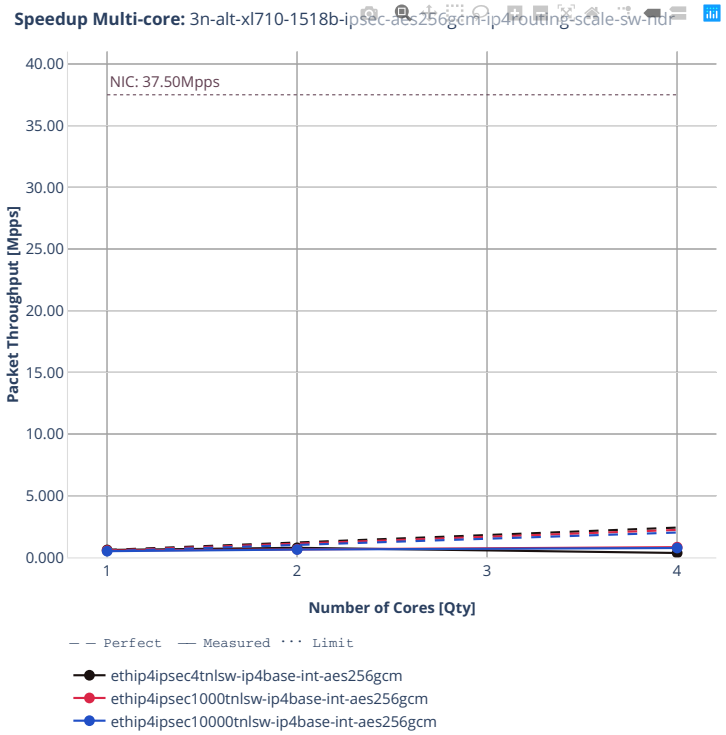

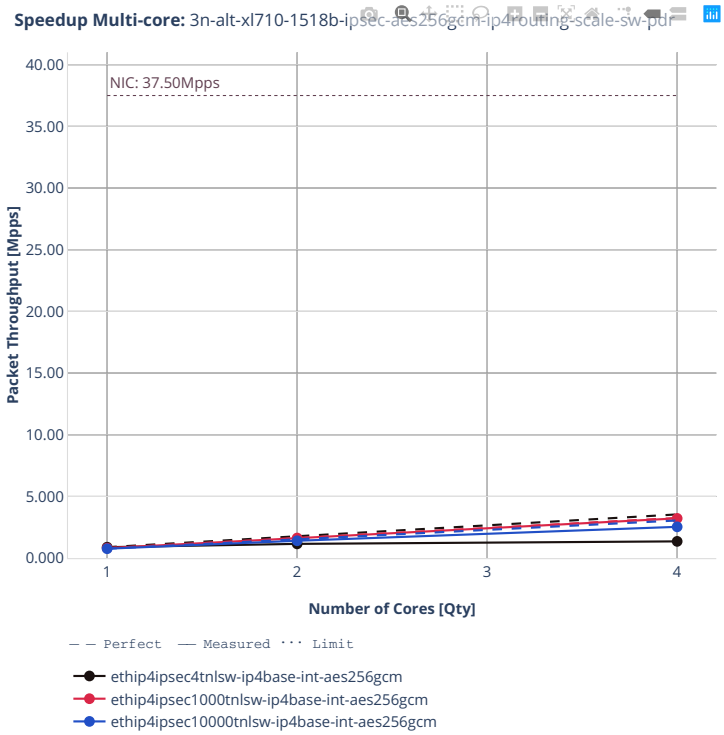
Speedup Multi-core: 3n-alt-xl710-1518b-ipsec-aes256gcm-ip4routing-scale-sw-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec4tnlsw-ip4base-int-aes256gcm
ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

imix-ipsec-aes128cbc-ip4routing-scale-sw

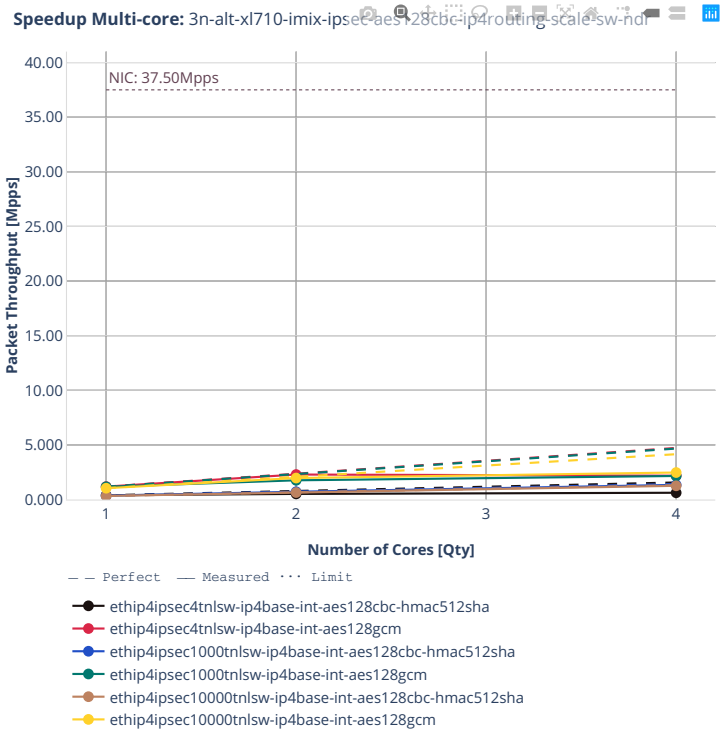

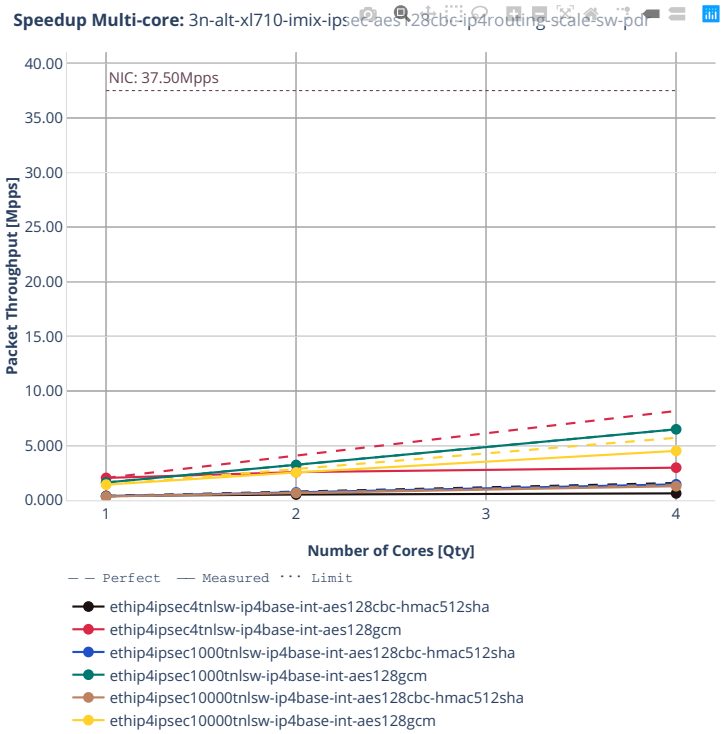
Speedup Multi-core: 3n-alt-xl710-imix-ipsec-aes128cbc-ip4routing-scale-sw-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
ethip4ipsec4tnlsw-ip4base-int-aes128gcm
ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha
ethip4ipsec1000tnlsw-ip4base-int-aes128gcm
ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha
ethip4ipsec10000tnlsw-ip4base-int-aes128gcm

imix-ipsec-aes256gcm-ip4routing-scale-sw

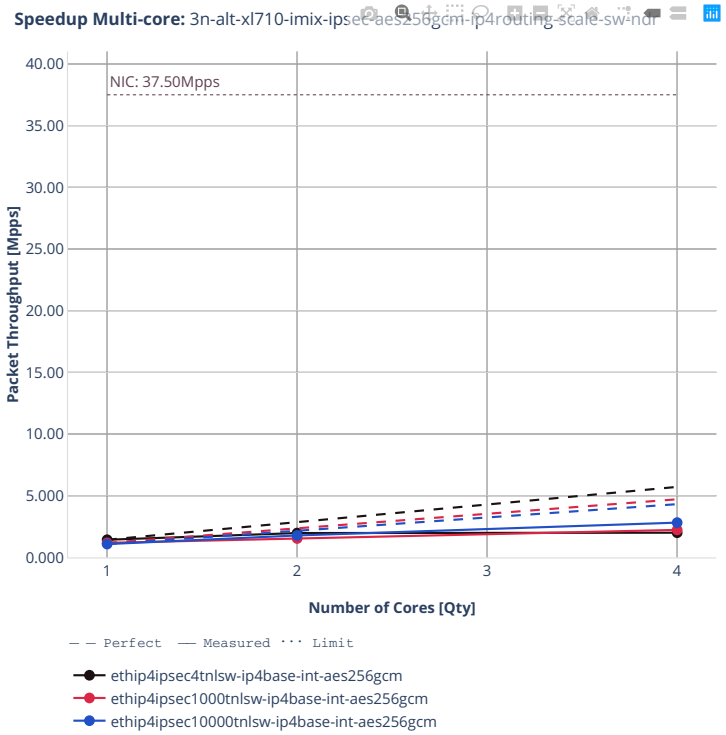

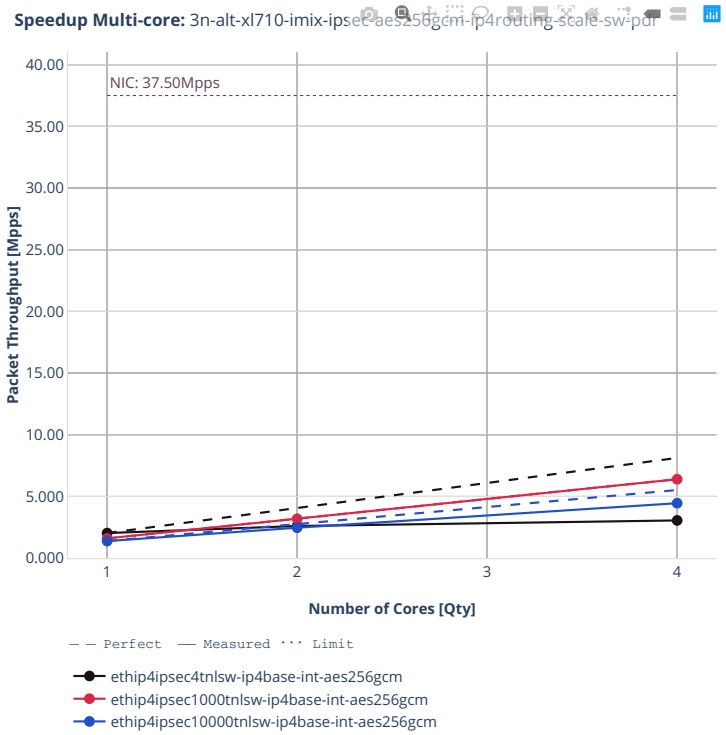
Speedup Multi-core: 3n-alt-xl710-imix-ipsec-aes256gcm-ip4routing-scale-sw-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec4tnlsw-ip4base-int-aes256gcm
ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
ethip4ipsec10000tnlsw-ip4base-int-aes256gcm

64b-ipsec-ip4base-policy-outbound-nocrypto

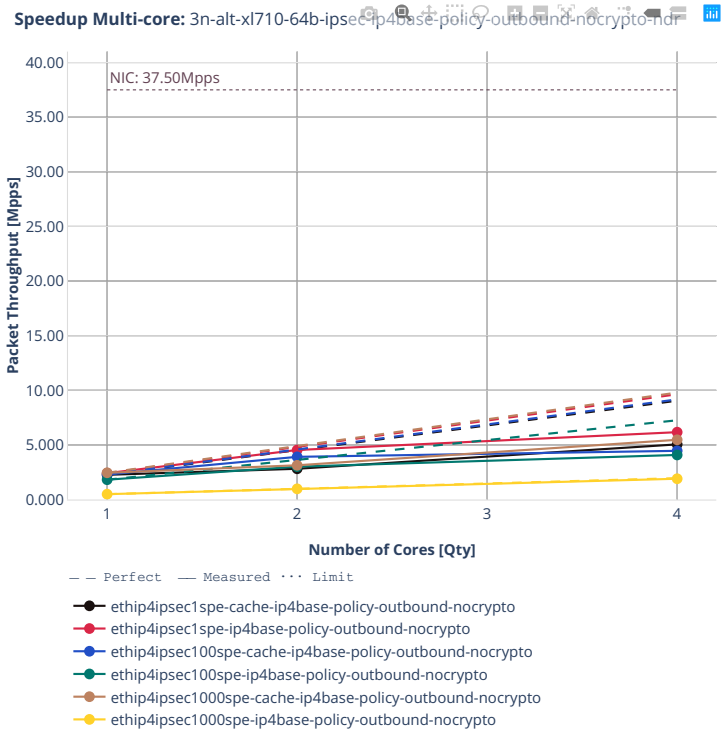

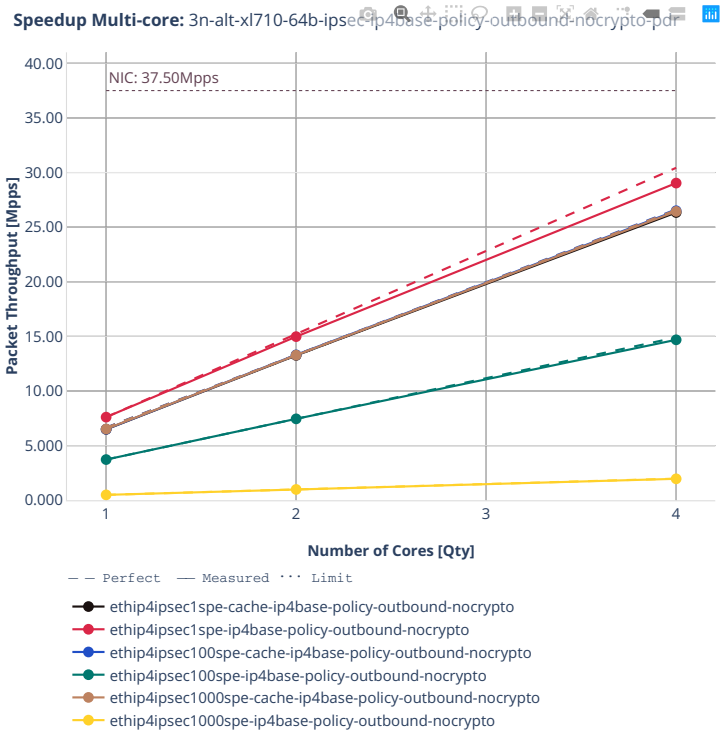
Speedup Multi-core: 3n-alt-xl710-64b-ipsec-ip4base-policy-outbound-nocrypto-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec1spe-cache-ip4base-policy-outbound-nocrypto
ethip4ipsec1spe-ip4base-policy-outbound-nocrypto
ethip4ipsec100spe-cache-ip4base-policy-outbound-nocrypto
ethip4ipsec100spe-ip4base-policy-outbound-nocrypto
ethip4ipsec1000spe-cache-ip4base-policy-outbound-nocrypto
ethip4ipsec1000spe-ip4base-policy-outbound-nocrypto

64b-ipsec-ip4base-policy-aes256gcm-sw

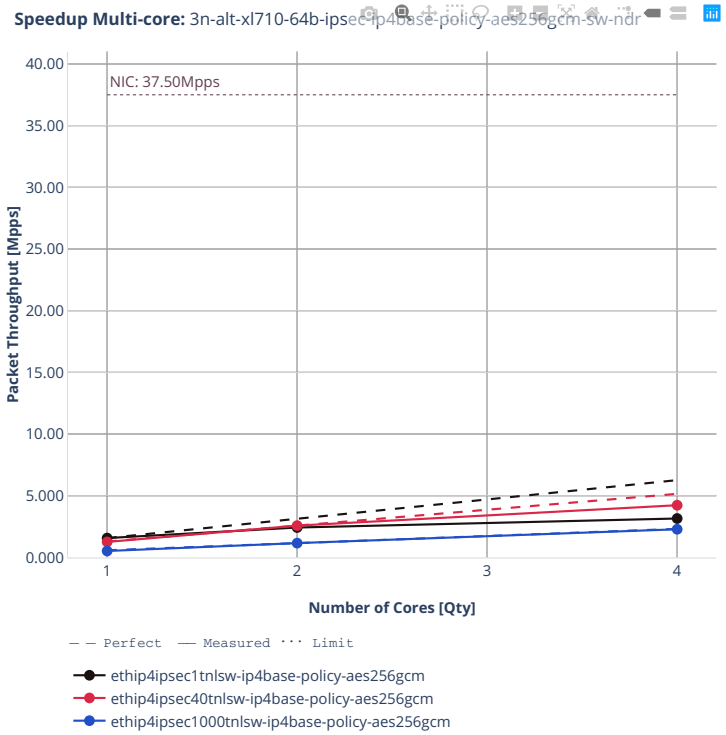

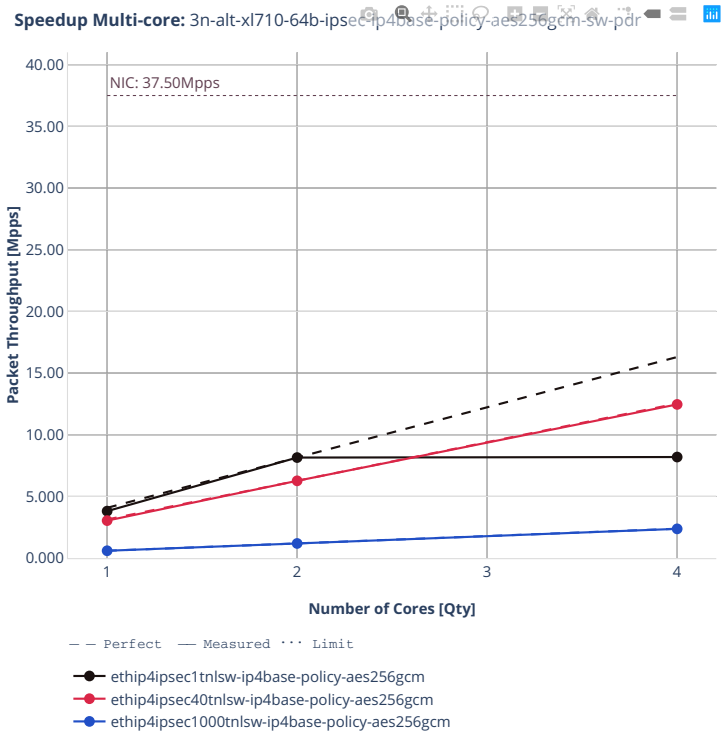
Speedup Multi-core: 3n-alt-xl710-64b-ipsec-ip4base-policy-aes256gcm-sw-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec1tnlsw-ip4base-policy-aes256gcm
ethip4ipsec40tnlsw-ip4base-policy-aes256gcm
ethip4ipsec1000tnlsw-ip4base-policy-aes256gcm

3n-tsh-x520

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

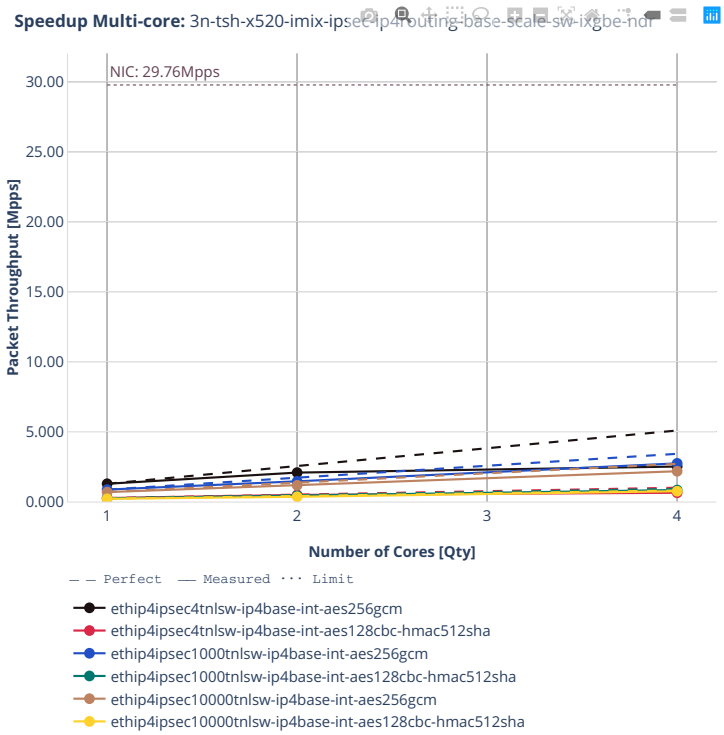

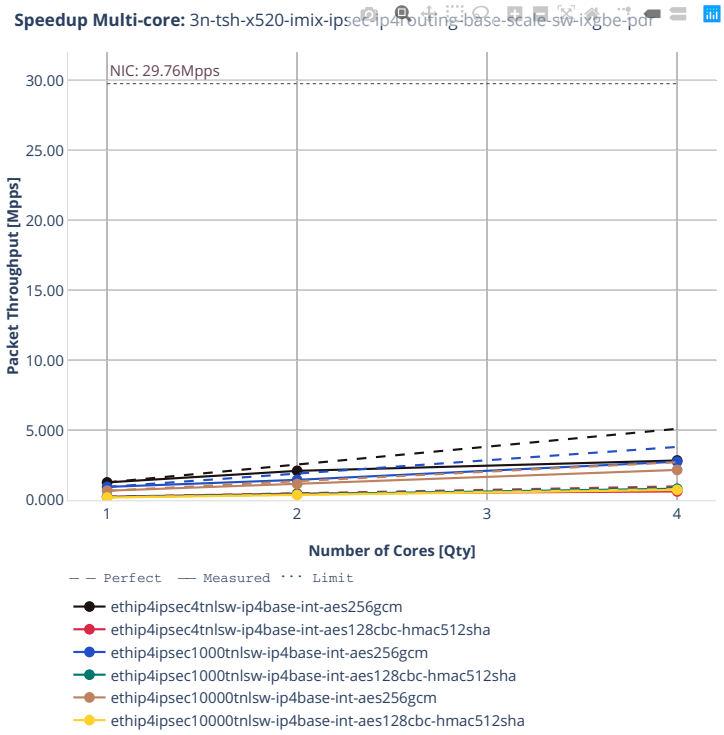
Speedup Multi-core: 3n-tsh-x520-imix-ipsec-ip4routing-base-scale-sw-ixgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec4tnlsw-ip4base-int-aes256gcm
ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha
ethip4ipsec10000tnlsw-ip4base-int-aes256gcm
ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha

2n-tx2-xl710

64b-ipsec-spe-ip4routing-base-scale

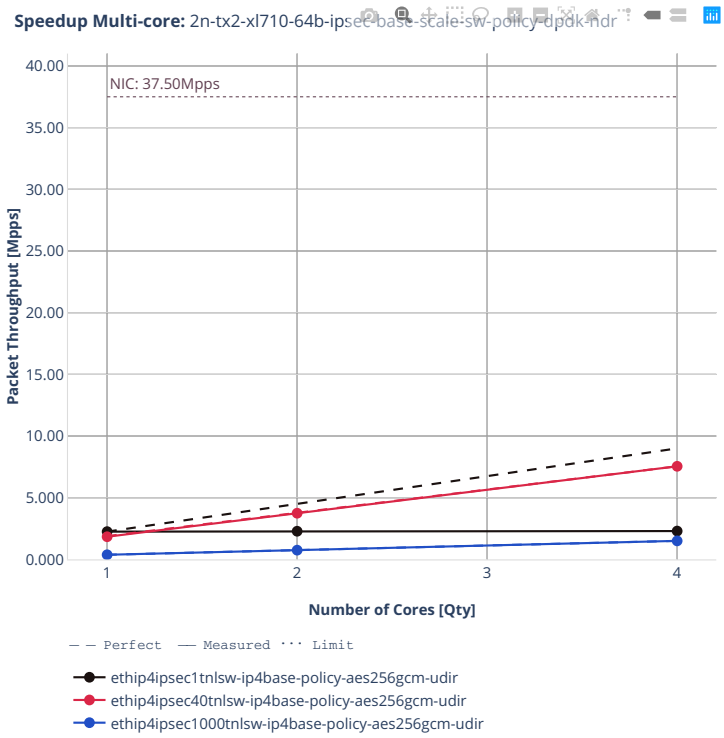

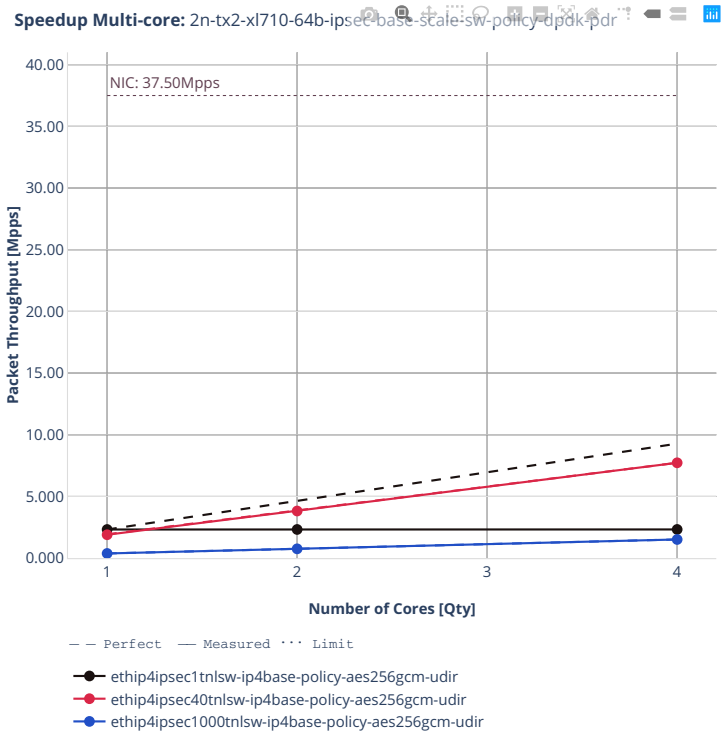
Speedup Multi-core: 2n-tx2-xl710-64b-ipsec-base-scale-sw-policy-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec1tnlsw-ip4base-policy-aes256gcm-udir
ethip4ipsec40tnlsw-ip4base-policy-aes256gcm-udir
ethip4ipsec1000tnlsw-ip4base-policy-aes256gcm-udir

64b-ipsec-ip4routing-base-scale-sw

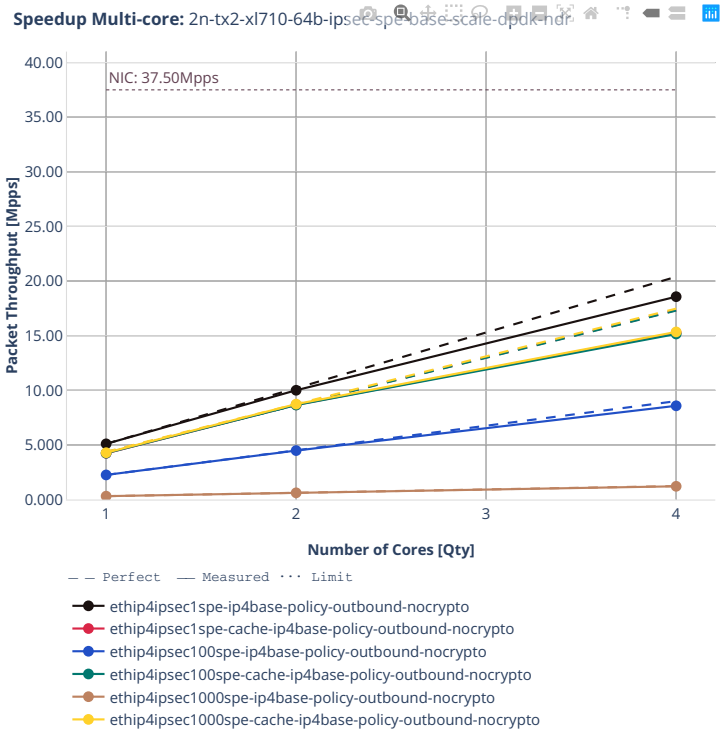

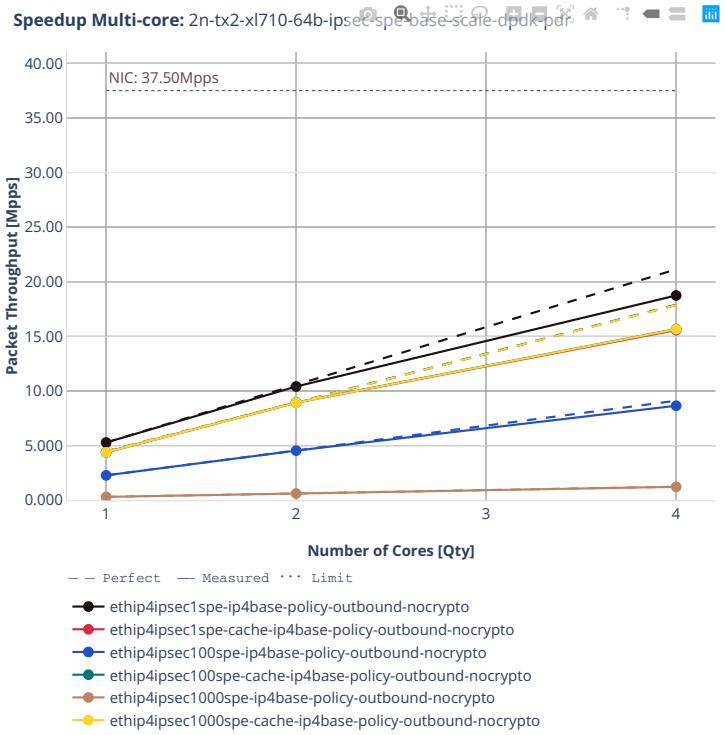
Speedup Multi-core: 2n-tx2-xl710-64b-ipsec-spe-base-scale-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec1spe-ip4base-policy-outbound-nocrypto
ethip4ipsec1spe-cache-ip4base-policy-outbound-nocrypto
ethip4ipsec100spe-ip4base-policy-outbound-nocrypto
ethip4ipsec100spe-cache-ip4base-policy-outbound-nocrypto
ethip4ipsec1000spe-ip4base-policy-outbound-nocrypto
ethip4ipsec1000spe-cache-ip4base-policy-outbound-nocrypto

1518b-ipsec-ip4routing-base-scale-sw-cbc

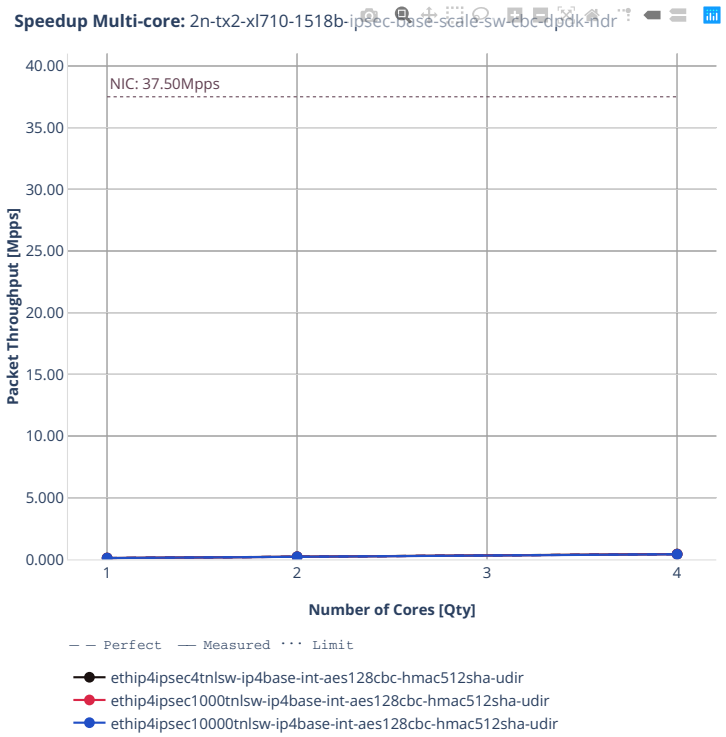

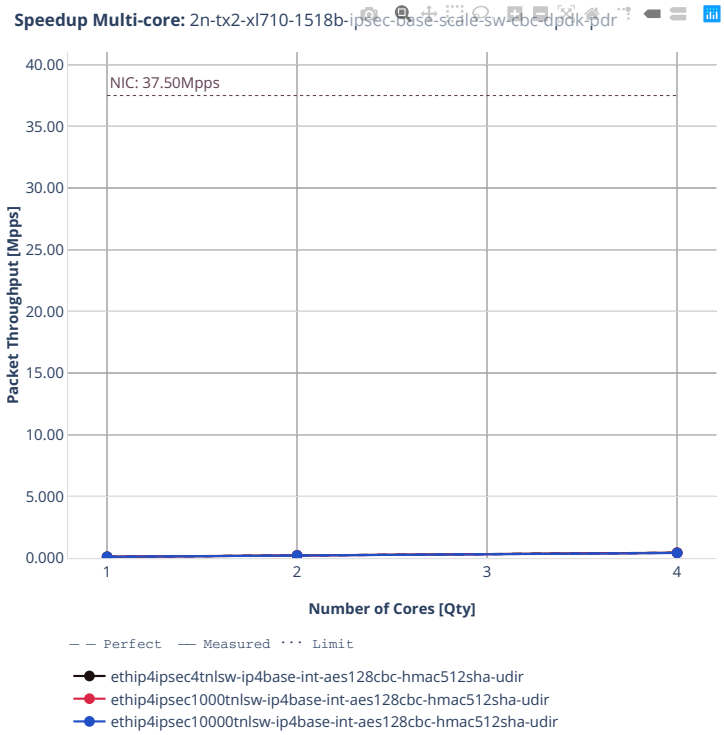
Speedup Multi-core: 2n-tx2-xl710-1518b-ipsec-base-scale-sw-cbc-dpdk-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha-udir
ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha-udir
ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha-udir

imix-ipsec-ip4routing-base-scale-sw-cbc

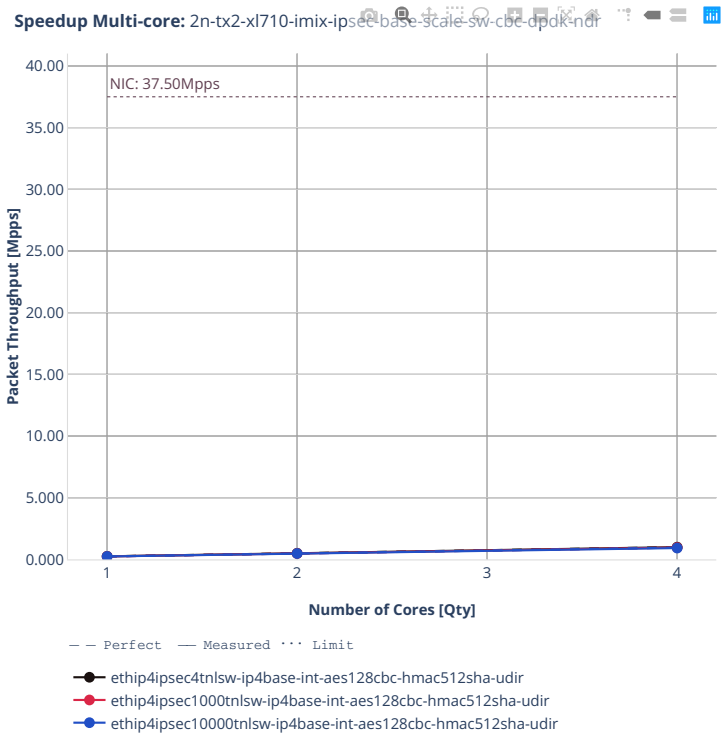

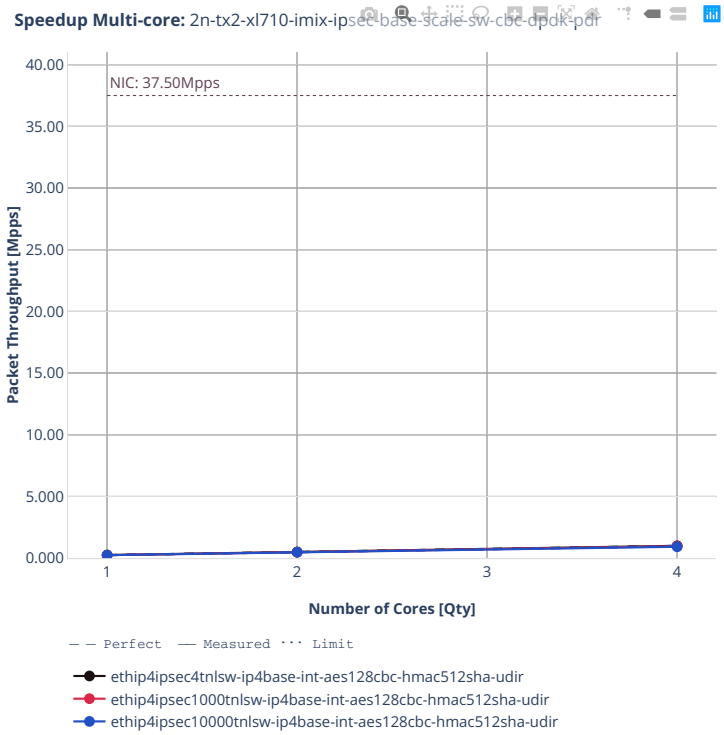
Speedup Multi-core: 2n-tx2-xl710-imix-ipsec-base-scale-sw-cbc-dpdk-pdr
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha-udir
ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha-udir
ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha-udir

1518b-ipsec-ip4routing-base-scale-sw-gcm

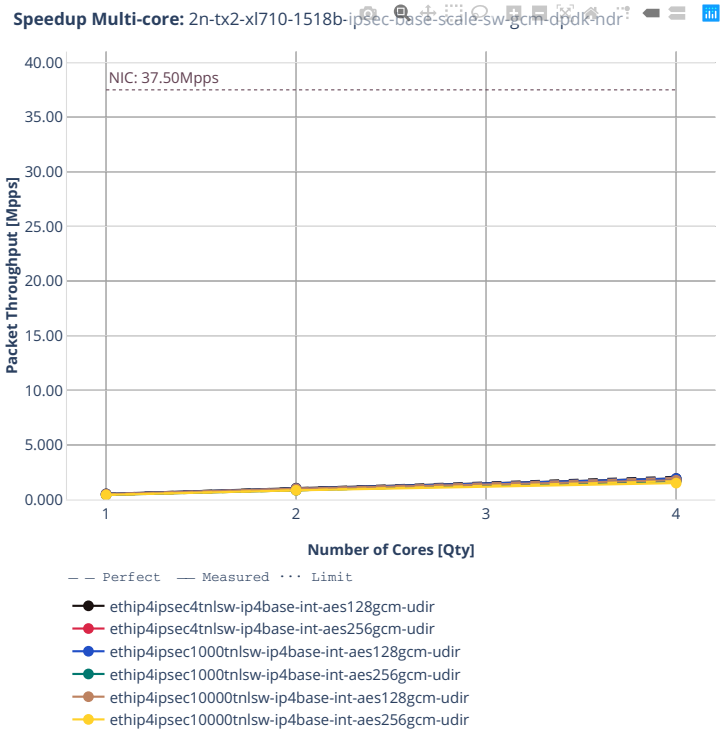

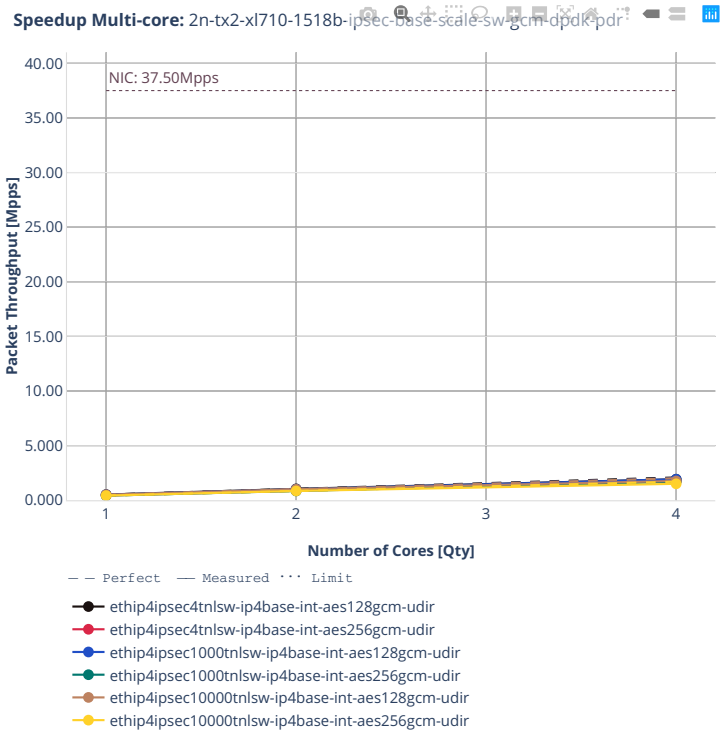
Speedup Multi-core: 2n-tx2-xl710-1518b-ipsec-base-scale-sw-gcm-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec4tnlsw-ip4base-int-aes128gcm-udir
ethip4ipsec4tnlsw-ip4base-int-aes256gcm-udir
ethip4ipsec1000tnlsw-ip4base-int-aes128gcm-udir
ethip4ipsec1000tnlsw-ip4base-int-aes256gcm-udir
ethip4ipsec10000tnlsw-ip4base-int-aes128gcm-udir
ethip4ipsec10000tnlsw-ip4base-int-aes256gcm-udir

imix-ipsec-ip4routing-base-scale-sw-gcm

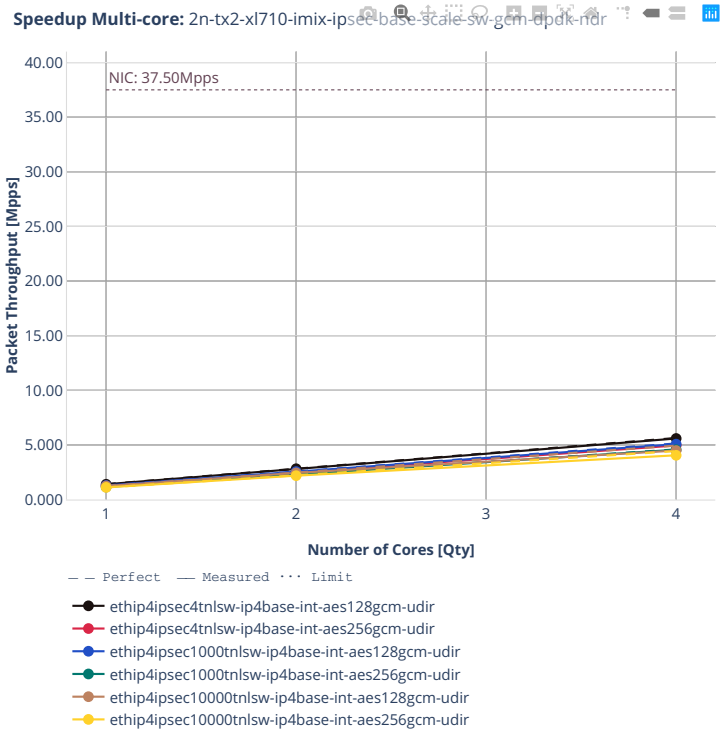

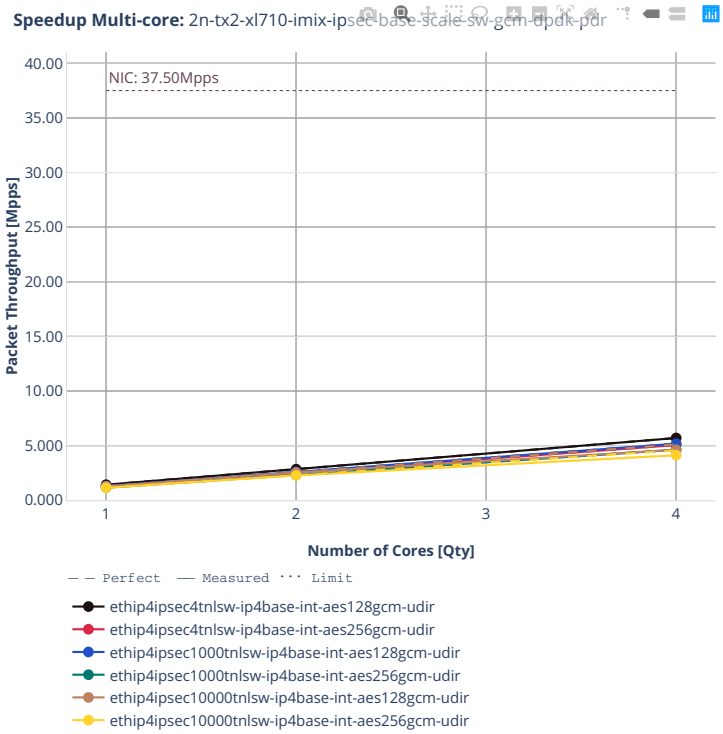
Speedup Multi-core: 2n-tx2-xl710-imix-ipsec-base-scale-sw-gcm-dpdk-pdr
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec4tnlsw-ip4base-int-aes128gcm-udir
ethip4ipsec4tnlsw-ip4base-int-aes256gcm-udir
ethip4ipsec1000tnlsw-ip4base-int-aes128gcm-udir
ethip4ipsec1000tnlsw-ip4base-int-aes256gcm-udir
ethip4ipsec10000tnlsw-ip4base-int-aes128gcm-udir
ethip4ipsec10000tnlsw-ip4base-int-aes256gcm-udir

3n-dnv-x553

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

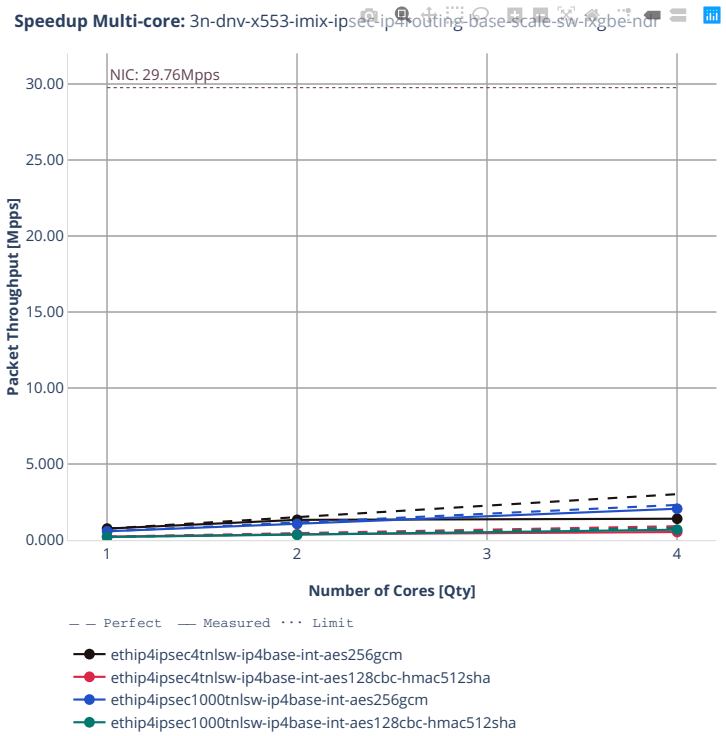

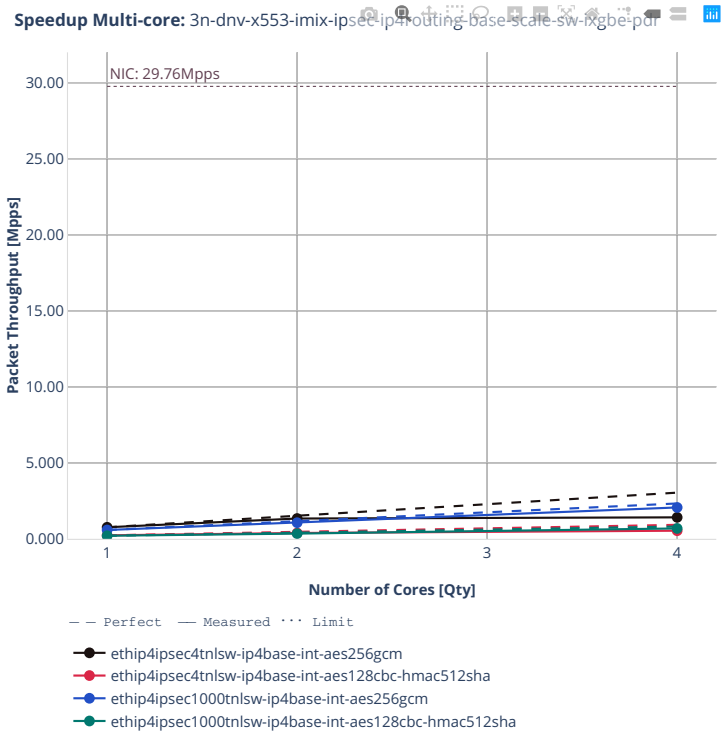
Speedup Multi-core: 3n-dnv-x553-imix-ipsec-ip4routing-base-scale-sw-nxgbe-pdr
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4ipsec4tnlsw-ip4base-int-aes256gcm
ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha

2.5 Packet Latency

VPP latency results are generated based on the test data obtained from CSIT-2206 NDR-PDR throughput tests executed across physical testbeds hosted in LF FD.io labs: 2n-icx, 3n-icx, 2n-aws, 2n-skx, 3n-skx, 2n-clx, 2n-zn2, 3n-alt, 3n-tsh, 2n-tx2.

Latency by percentile distribution plots are used to show packet latency percentiles at different packet rate load levels: i) No-Load latency streams only, ii) Low-Load at 10% PDR, iii) Mid-Load at 50% PDR and iv) High-Load at 90% PDR.

For more details, see *Packet Latency* (page 45).

Additional information about graph data:

1. **Graph Title**: describes tested DUT packet path.
2. **X-axis Labels**: percentile of packets.
3. **Y-axis Labels**: measured one-way packet latency values in [uSec].
4. **Graph Legend**: list of latency tests at different packet rate load level.
5. **Hover Information**: packet rate load level, stream direction (East-West, West-East), percentile, one-way latency.

Note: Test results are stored in build logs from FD.io vpp performance job 2n-icx[139], build logs from FD.io vpp performance job 3n-icx[140], build logs from FD.io vpp performance job 2n-aws[141], build logs from FD.io vpp performance job 2n-skx[142], build logs from FD.io vpp performance job 3n-skx[143], build logs from FD.io vpp performance job 2n-clx[144], build logs from FD.io vpp performance job 2n-zn2[145], build logs from FD.io vpp performance job 3n-alt[146], build logs from FD.io vpp performance job 3n-tsh[147] and build logs from FD.io vpp performance job 2n-tx2[148] with RF result files csit-vpp-perf-2206-*.zip archived here.

---

[139] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-icx
[140] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-icx
[141] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-aws
[142] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx
[143] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-skx
[144] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-clx
[145] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-zn2
[146] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-alt
[147] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-tsh
[148] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-tx2

2.5.1 L2 Ethernet Switching

CSIT source code for the test cases used for plots can be found in CSIT git repository[149].

[149] https://git.fd.io/csit/tree/tests/vpp/perf/l2?h=rls2206

2n-icx-xxv710

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

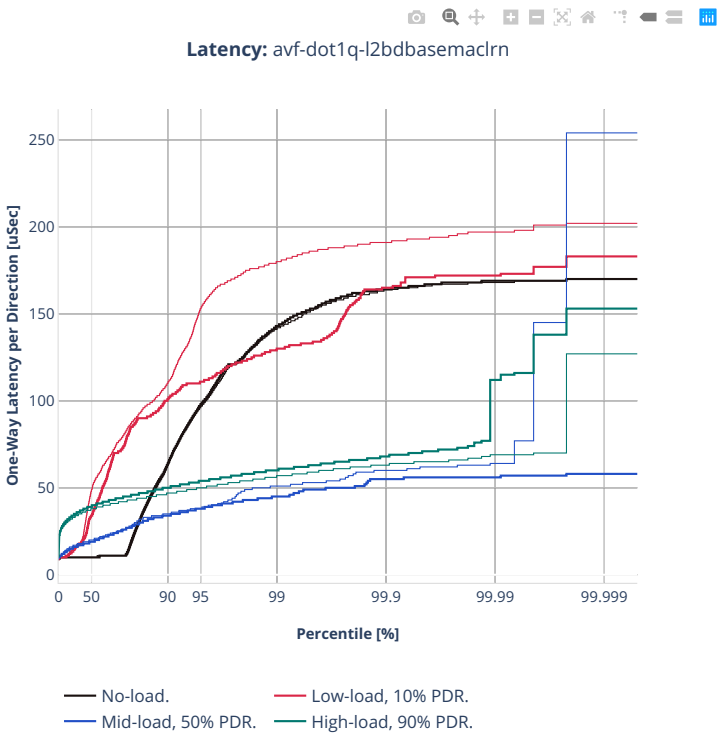

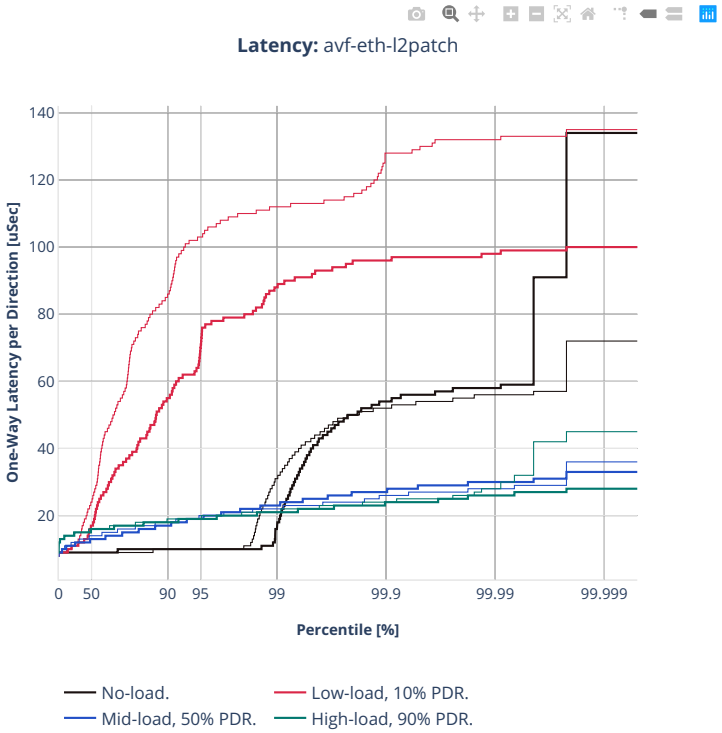
Latency: avf-eth-l2patch
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

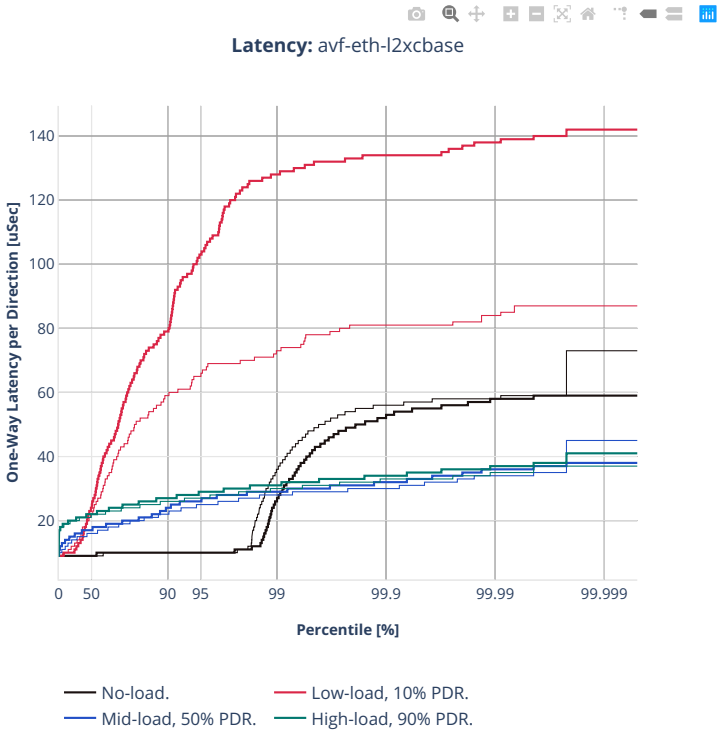
Latency: avf-eth-l2xcbase
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

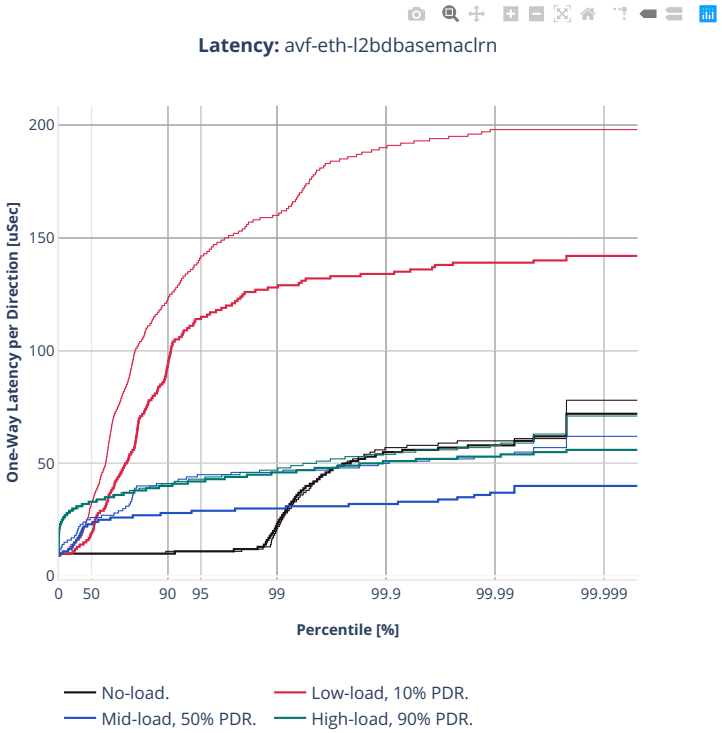
Latency: avf-eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

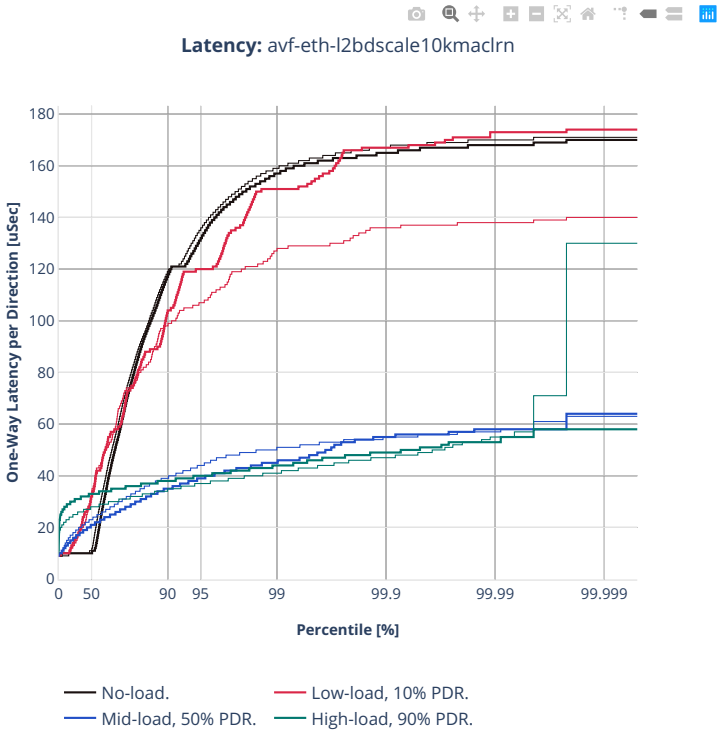
Latency: avf-eth-l2bdscale10kmaclrn
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

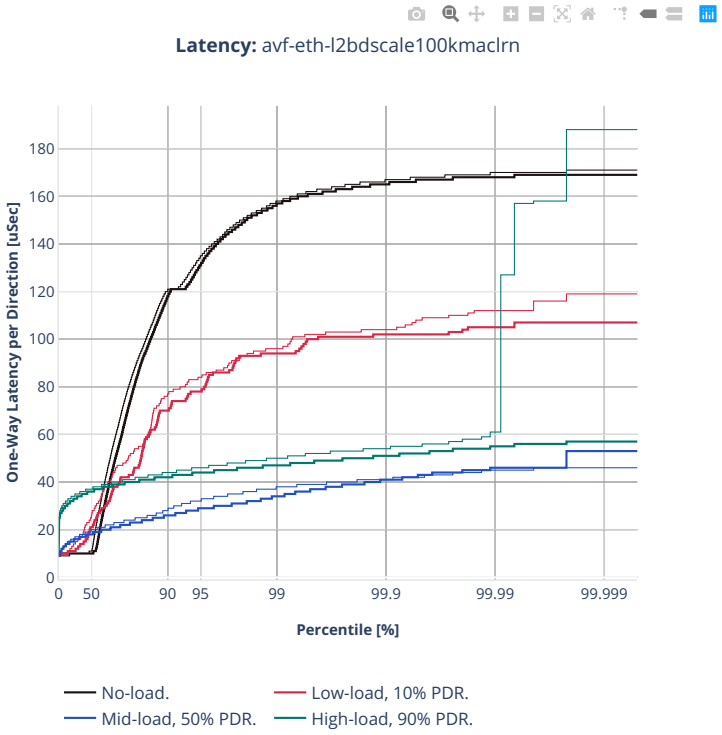
Latency: avf-eth-l2bdscale100kmaclrn
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

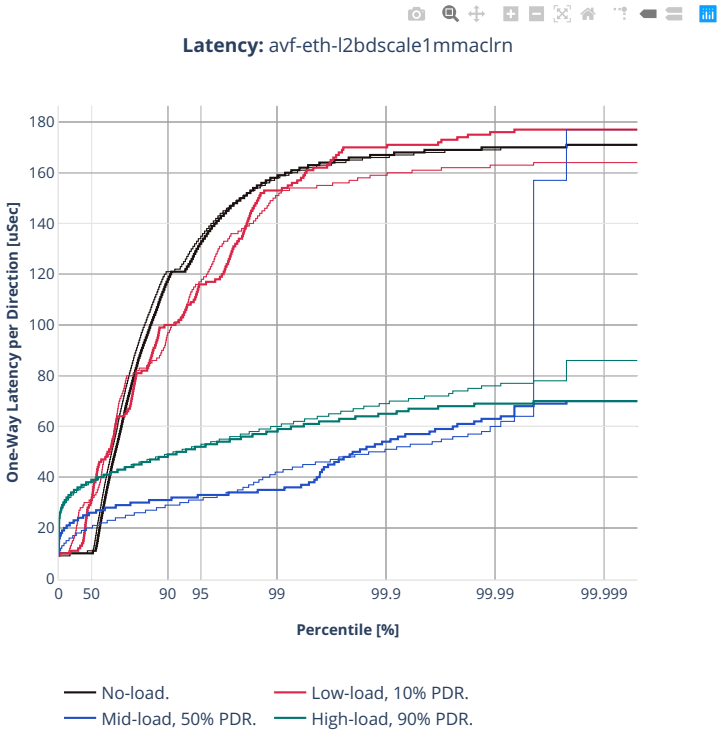
Latency: avf-eth-l2bdscale1mmaclrn
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

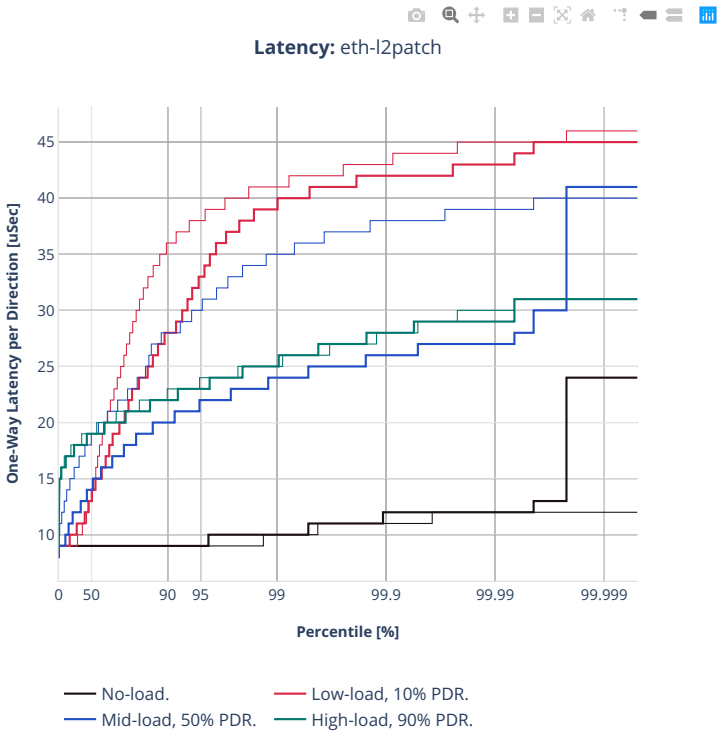

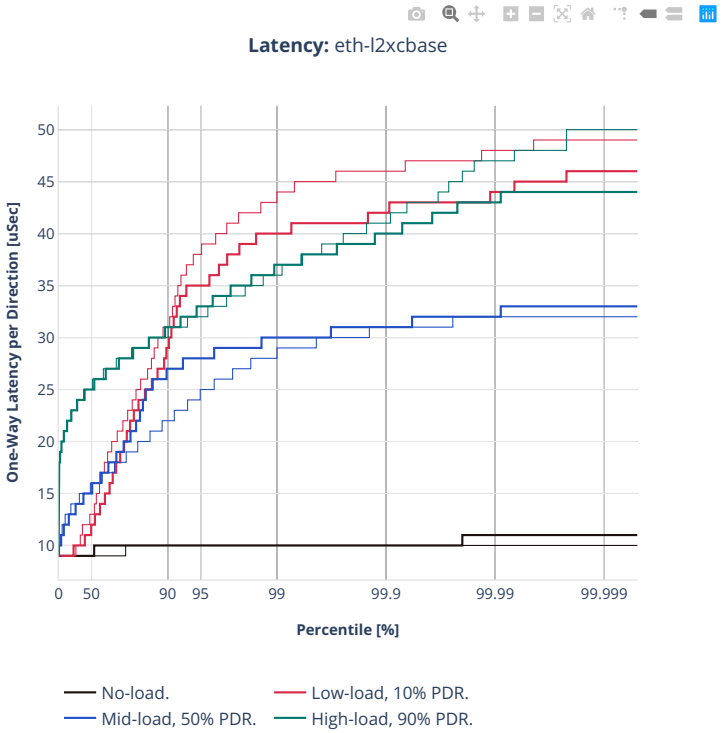
Latency: eth-l2xcbase
One-Way Latency per Direction [uSec]
50
45
40
35
30
25
20
15
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

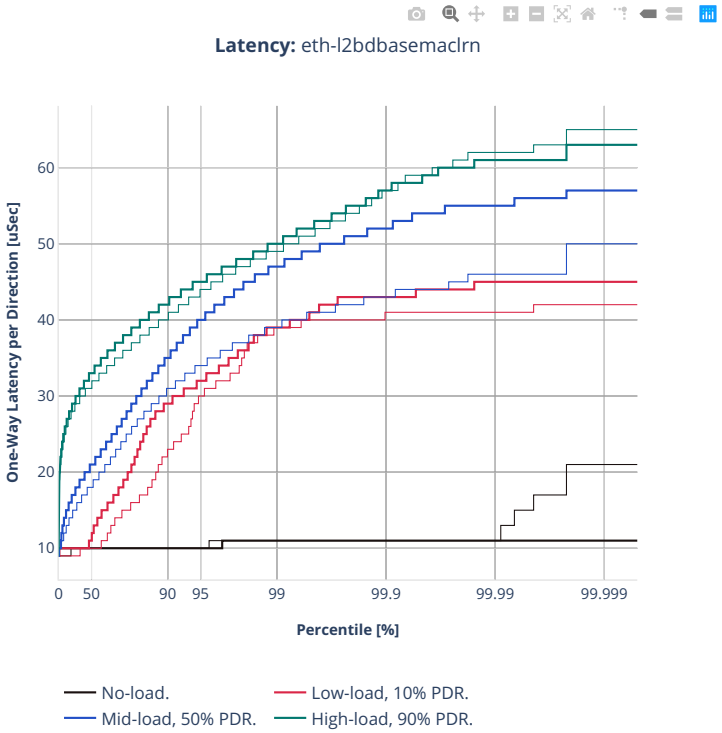
Latency: eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

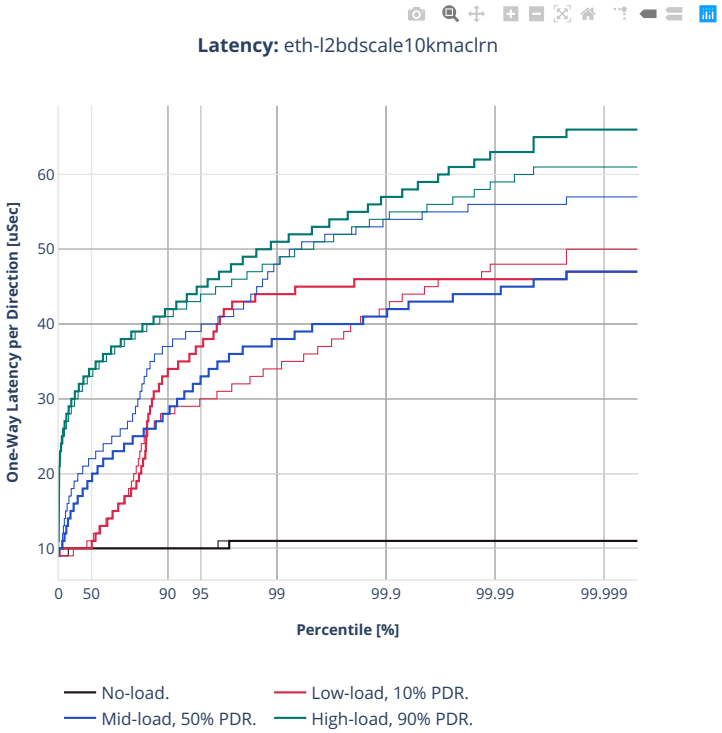
Latency: eth-l2bdscale10kmaclrn
One-Way Latency per Direction [uSec]
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

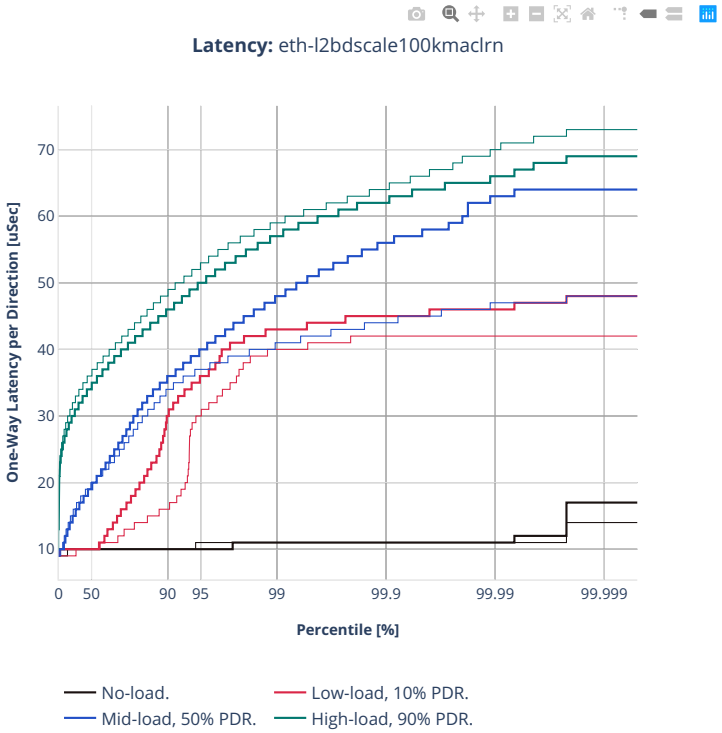
Latency: eth-l2bdscale100kmaclrn
One-Way Latency per Direction [uSec]
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

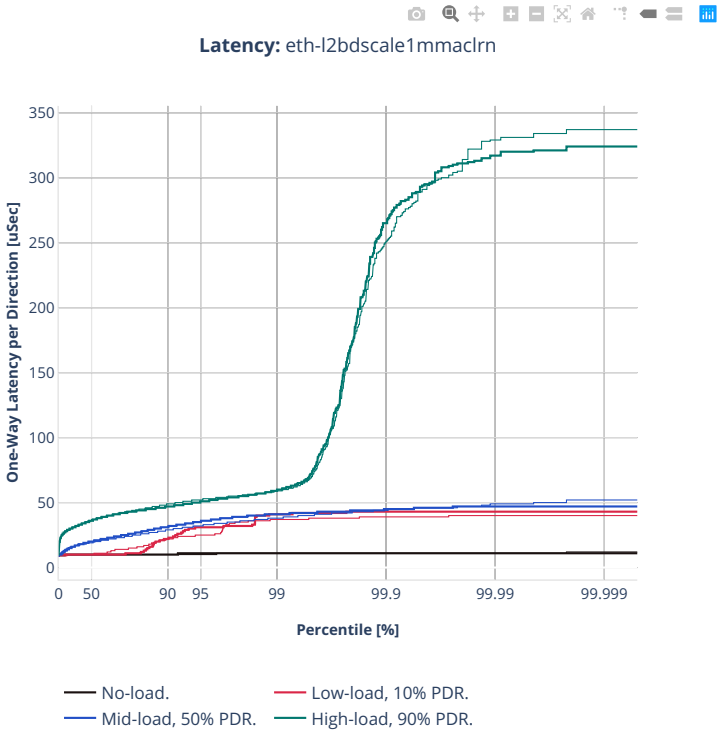
Latency: eth-l2bdscale1mmaclrn
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-icx-xxv710

64b-2t1c-l2switching-base-avf

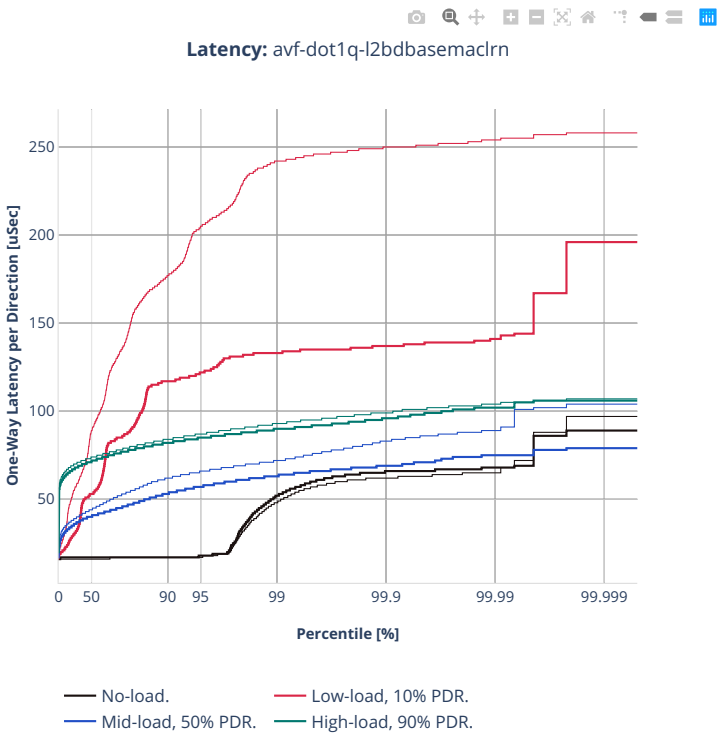

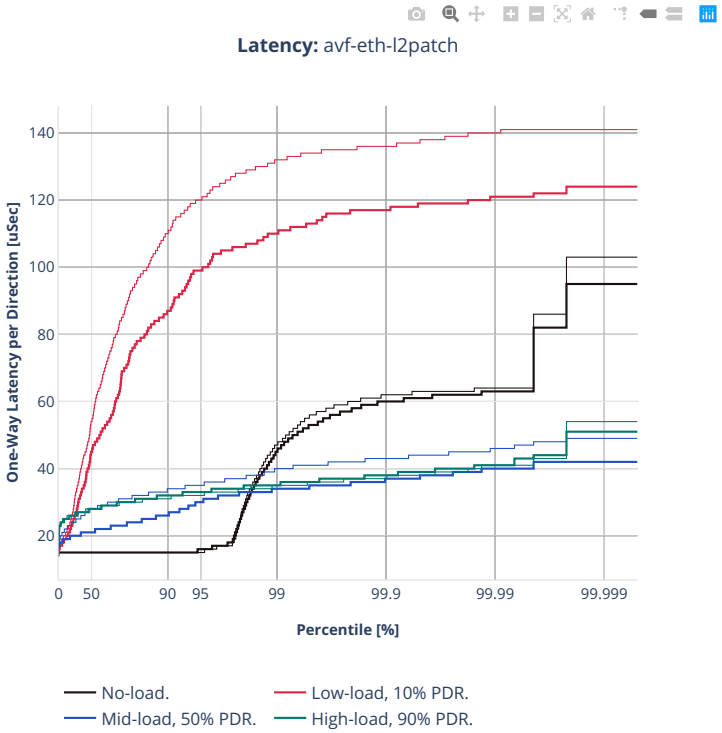
Latency: avf-eth-l2patch
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

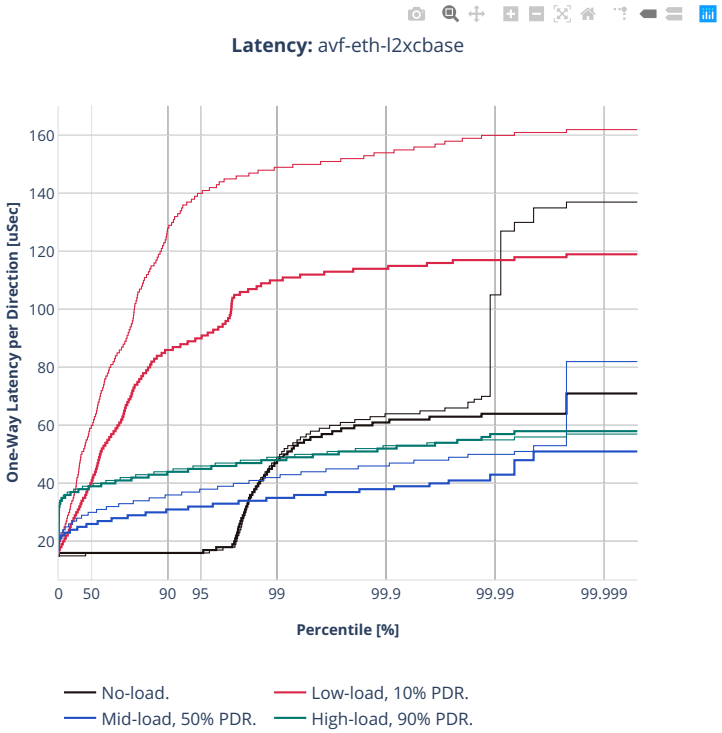
Latency: avf-eth-l2xcbase
One-Way Latency per Direction [uSec]
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

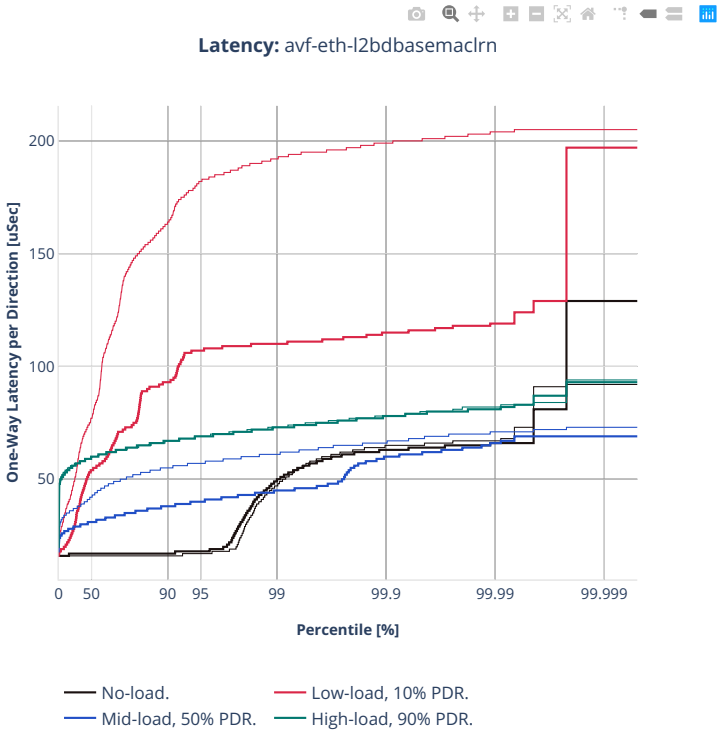
Latency: avf-eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
200
150
100
50
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-l2switching-base-dpdk

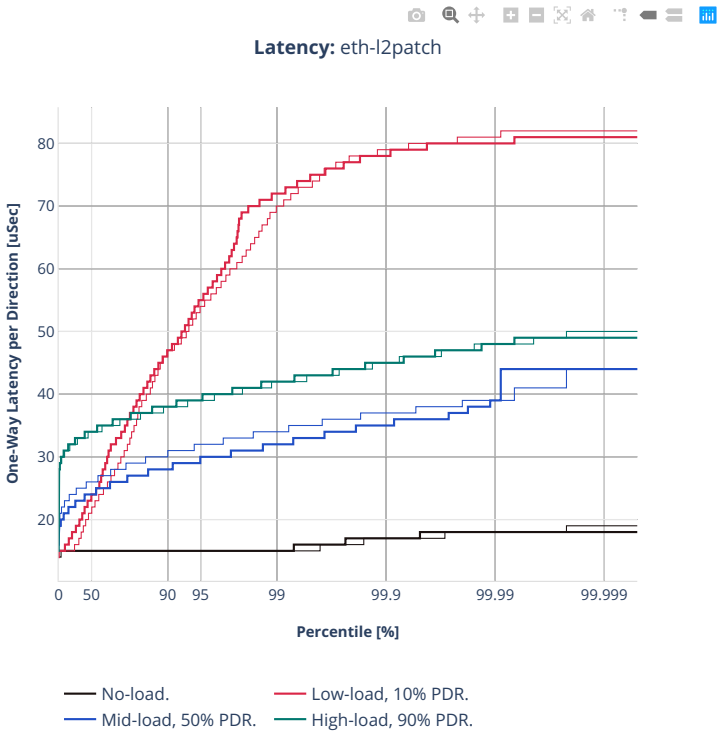

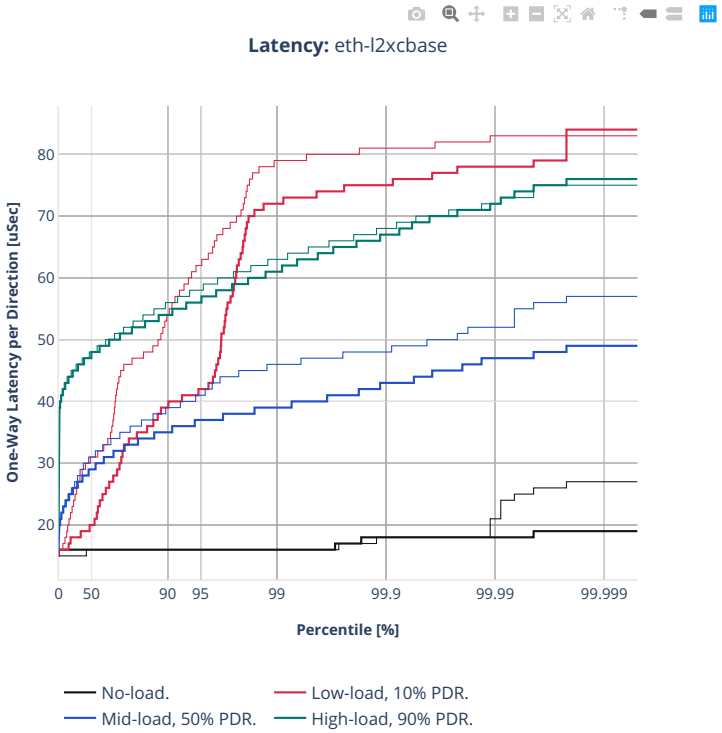
Latency: eth-l2xcbase
One-Way Latency per Direction [uSec]
80
70
60
50
40
30
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

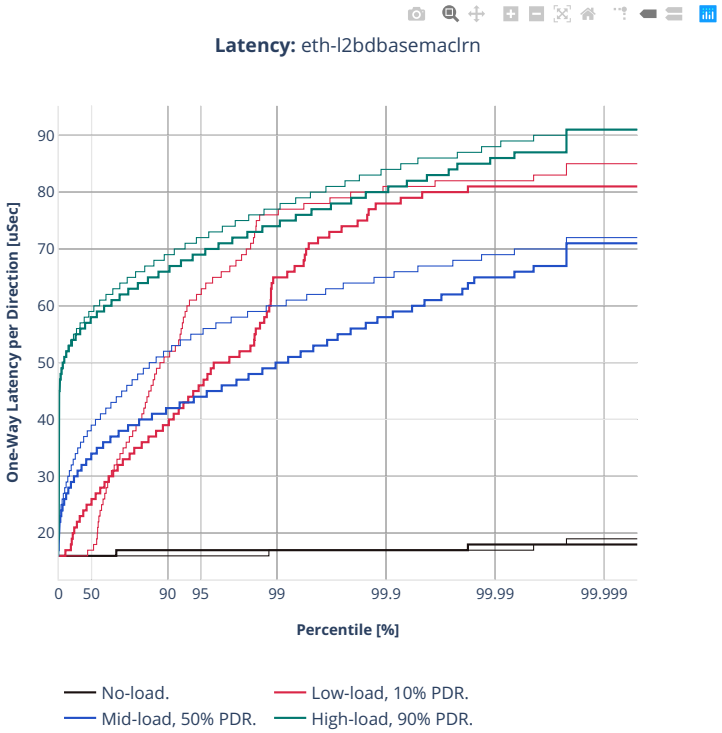
Latency: eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
90
80
70
60
50
40
30
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-skx-xxv710

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

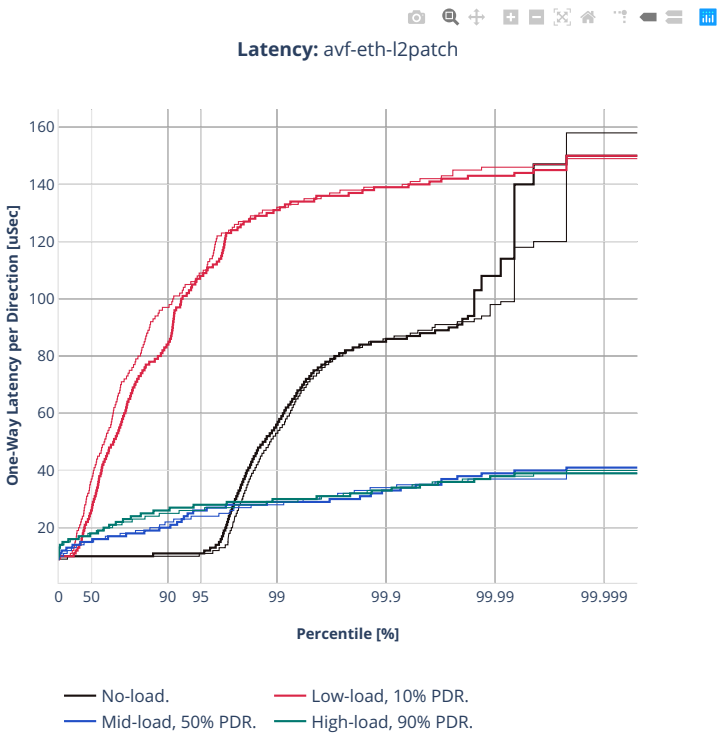

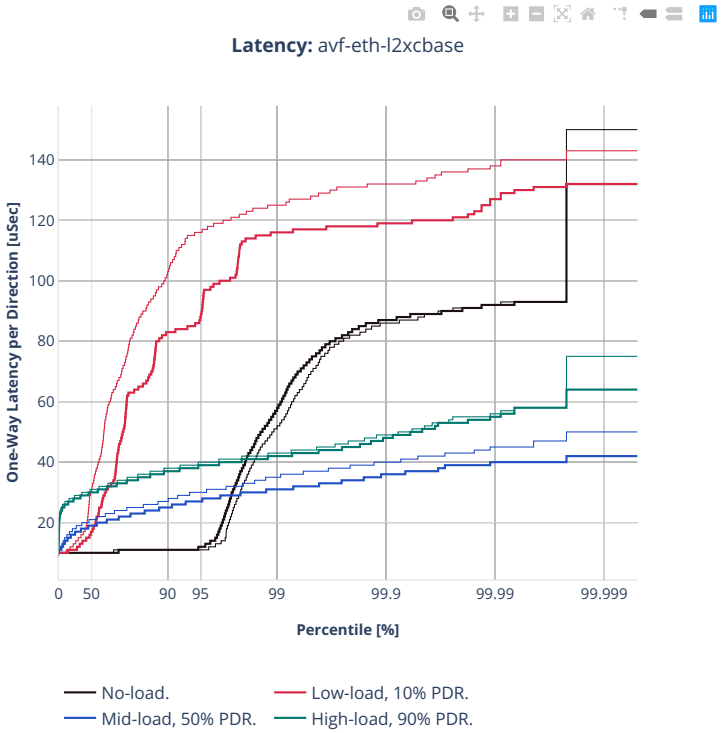
Latency: avf-eth-l2xcbase
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

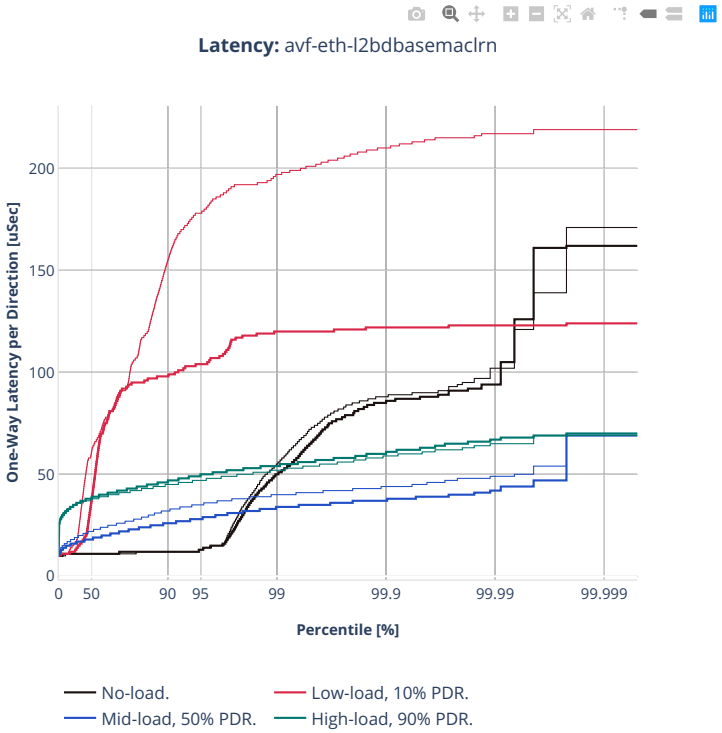
Latency: avf-eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

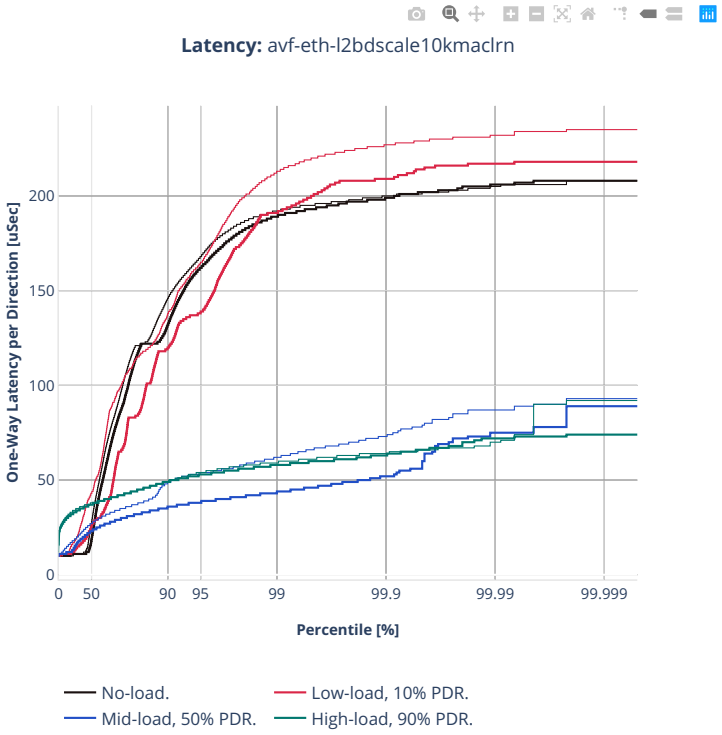
Latency: avf-eth-l2bdscale10kmaclrn
One-Way Latency per Direction [uSec]
0
50
100
150
200
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

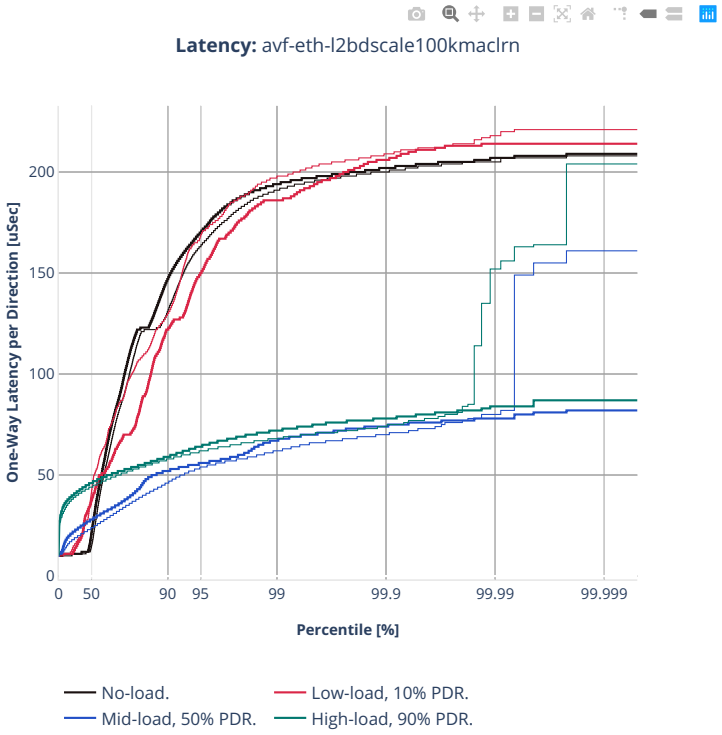
Latency: avf-eth-l2bdscale100kmaclrn
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

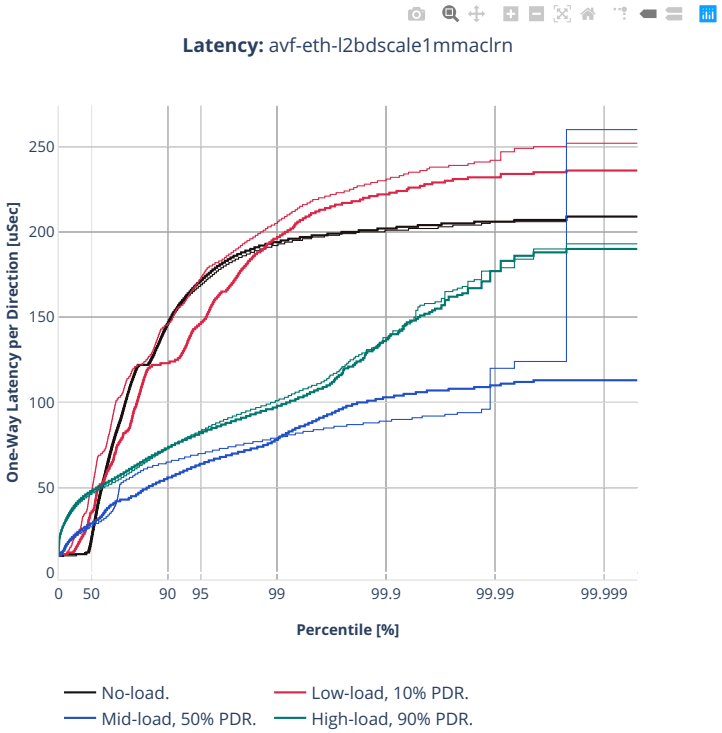
Latency: avf-eth-l2bdscale1mmaclrn
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

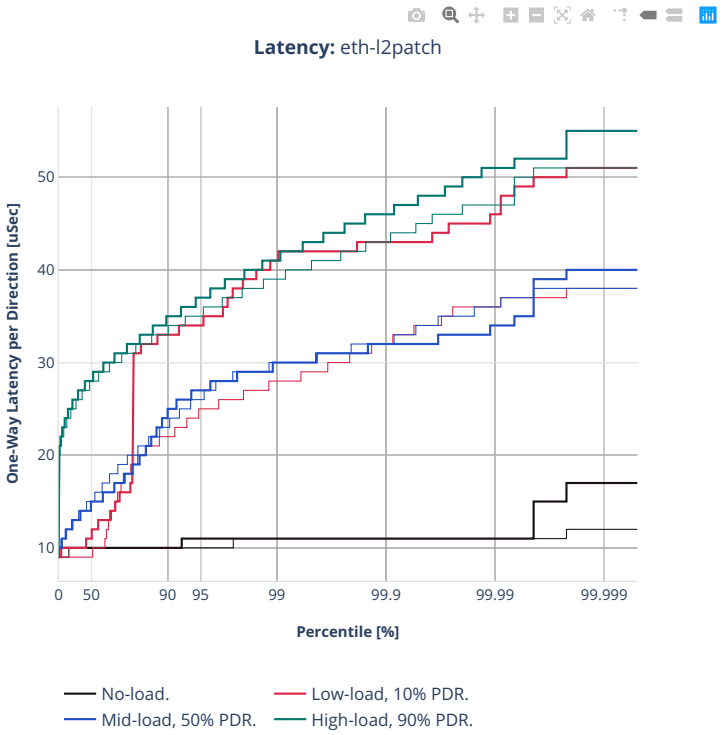

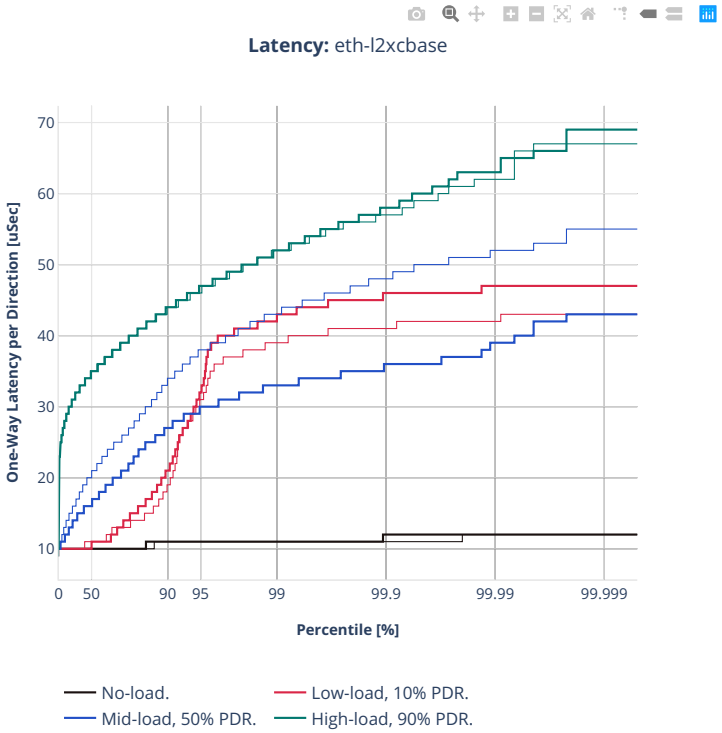
Latency: eth-l2xcbase
One-Way Latency per Direction [uSec]
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

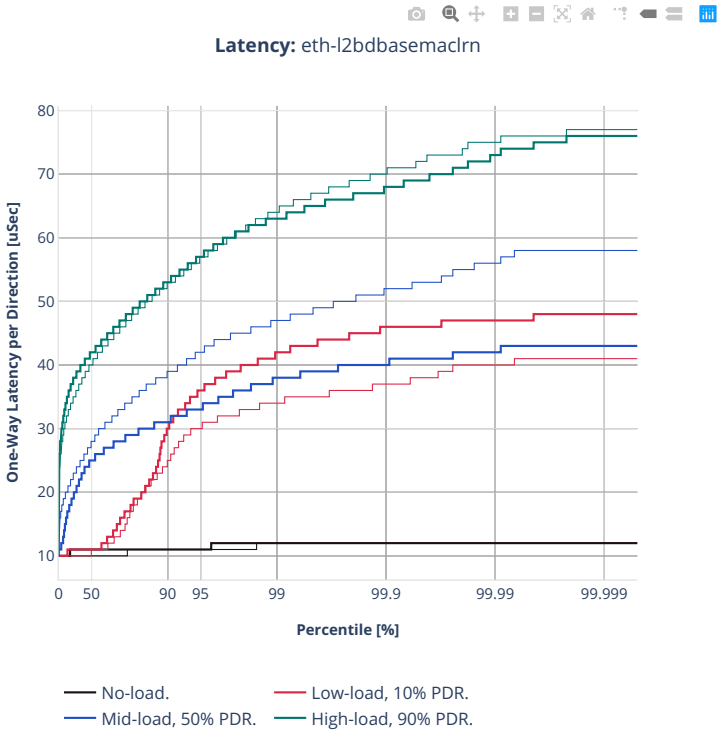
Latency: eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

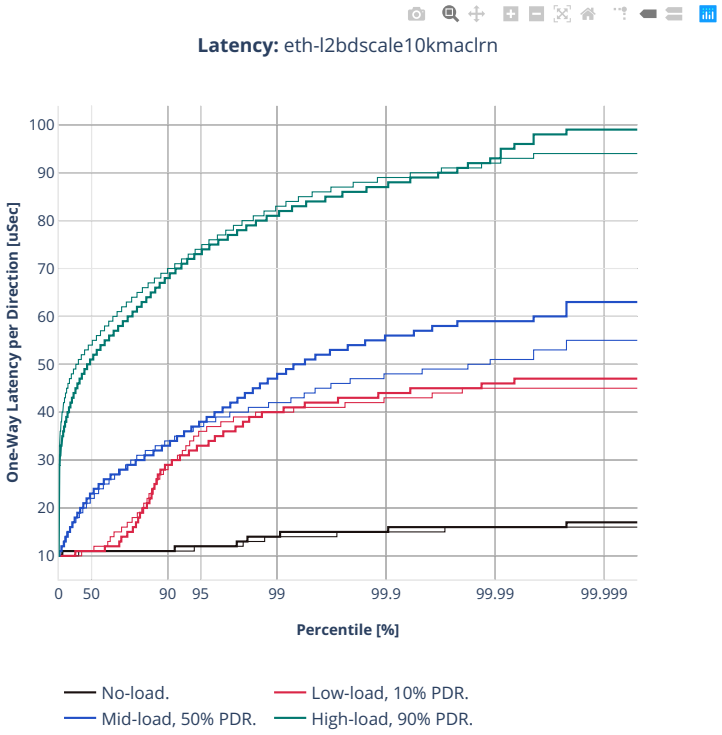
Latency: eth-l2bdscale10kmaclrn
One-Way Latency per Direction [uSec]
100
90
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

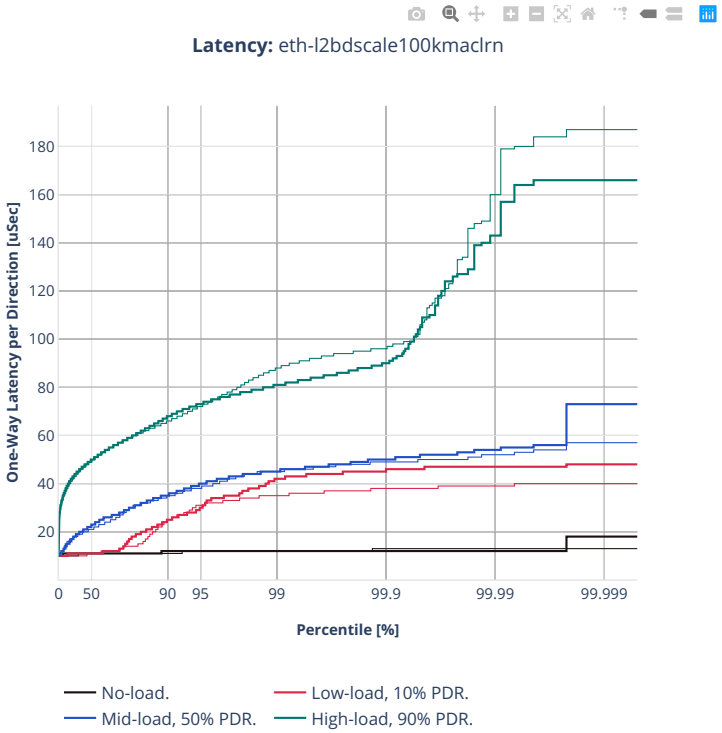
Latency: eth-l2bdscale100kmaclrn
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

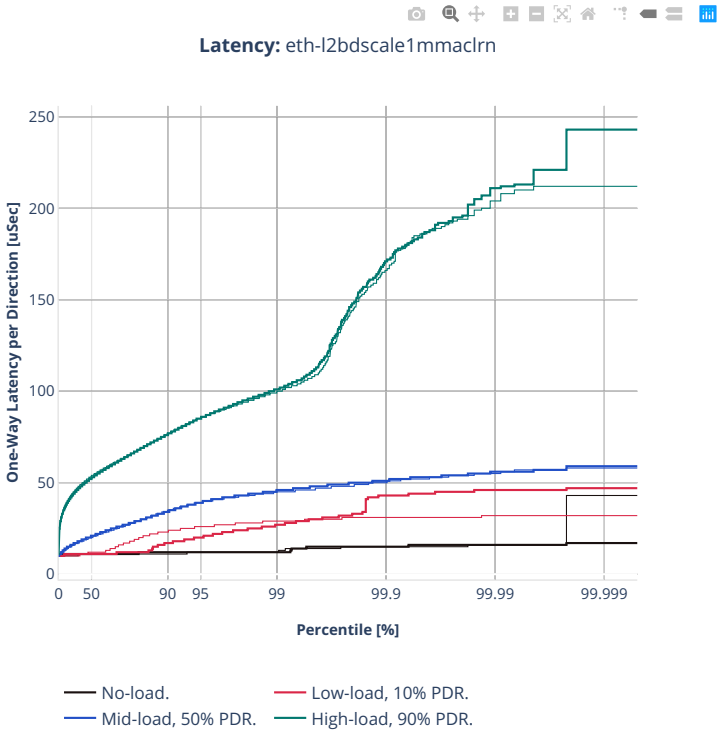
Latency: eth-l2bdscale1mmaclrn
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-skx-xxv710

64b-2t1c-l2switching-base-avf

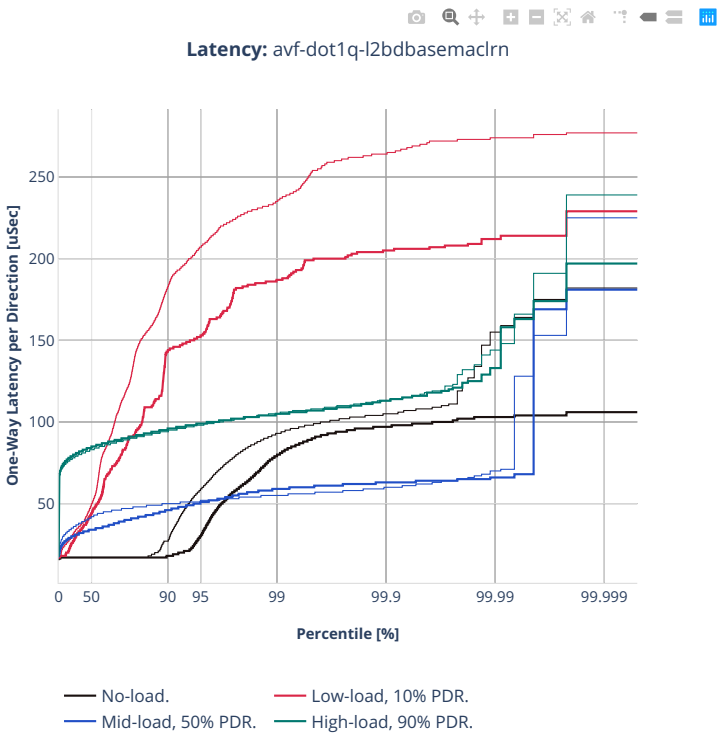

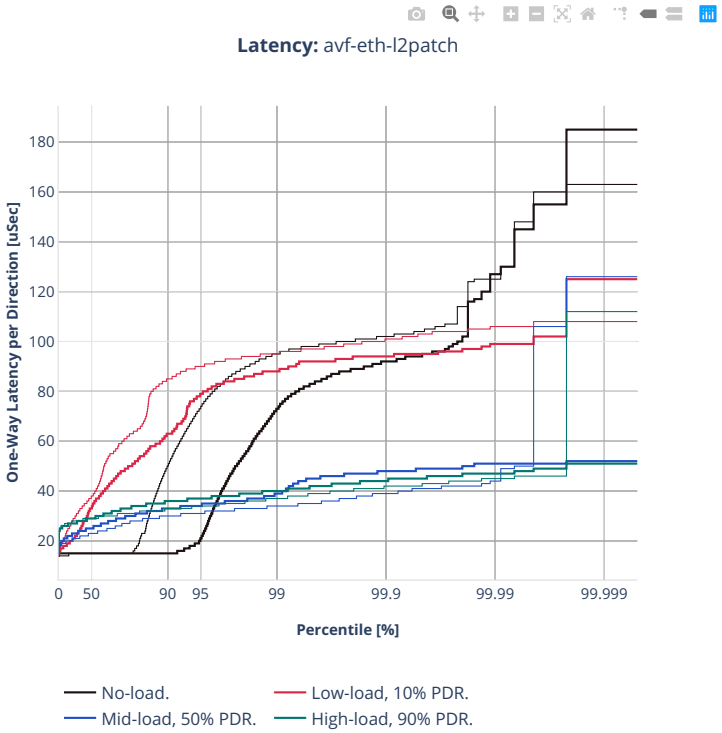
Latency: avf-eth-l2patch
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

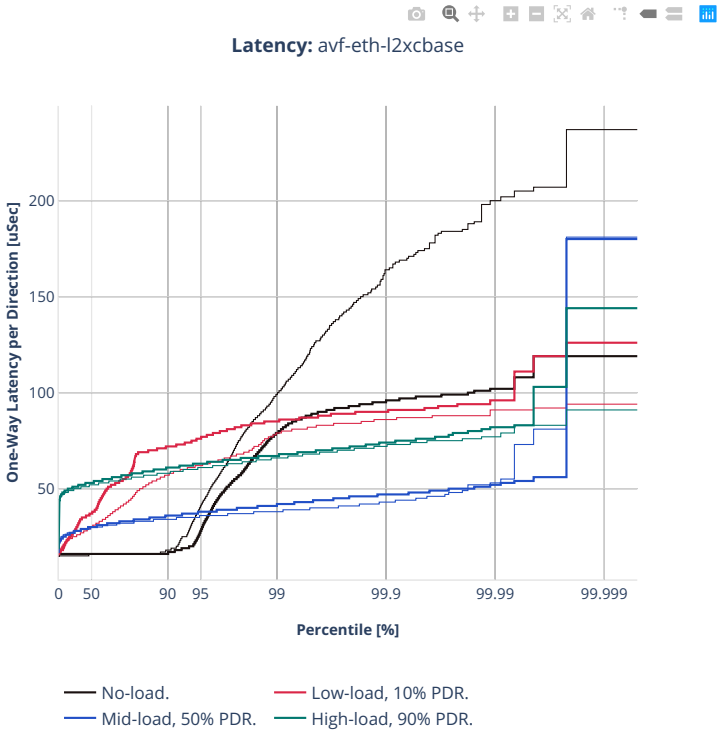
Latency: avf-eth-l2xcbase
One-Way Latency per Direction [uSec]
200
150
100
50
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

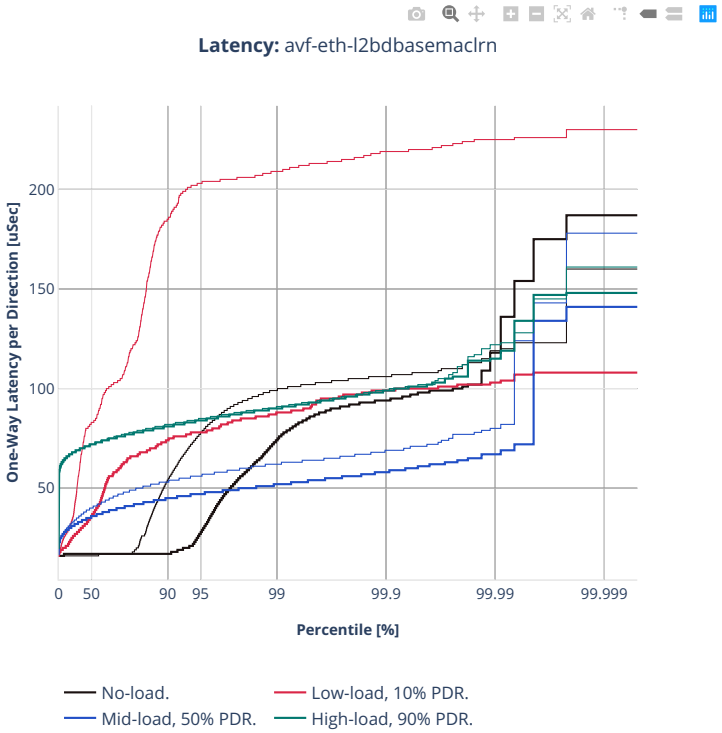
Latency: avf-eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
200
150
100
50
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-l2switching-base-dpdk

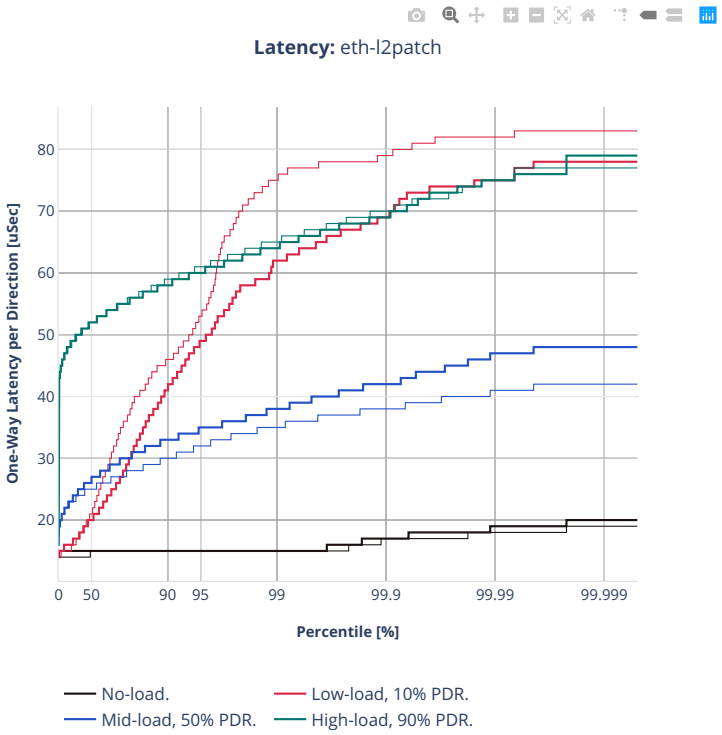

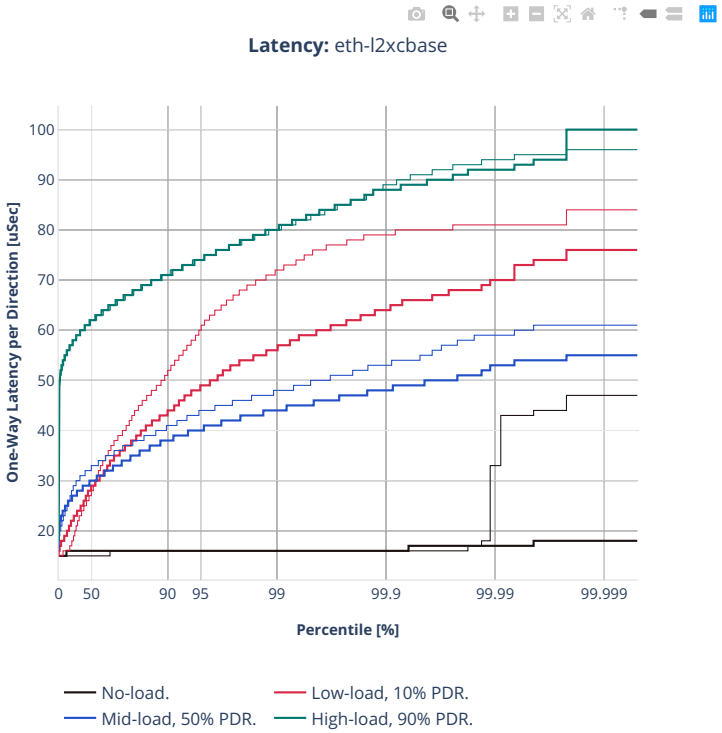
Latency: eth-l2xcbase
One-Way Latency per Direction [uSec]
100
90
80
70
60
50
40
30
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

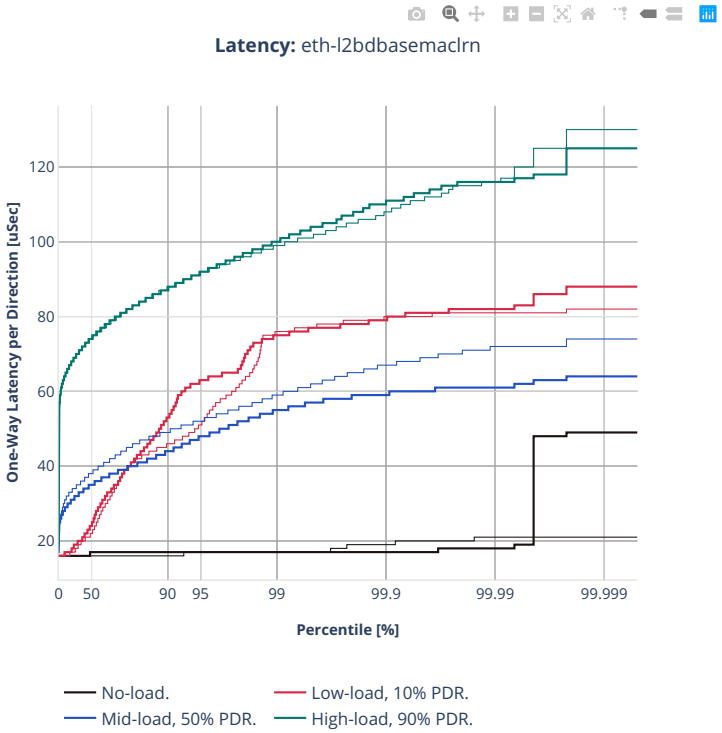
Latency: eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-clx-xxv710

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

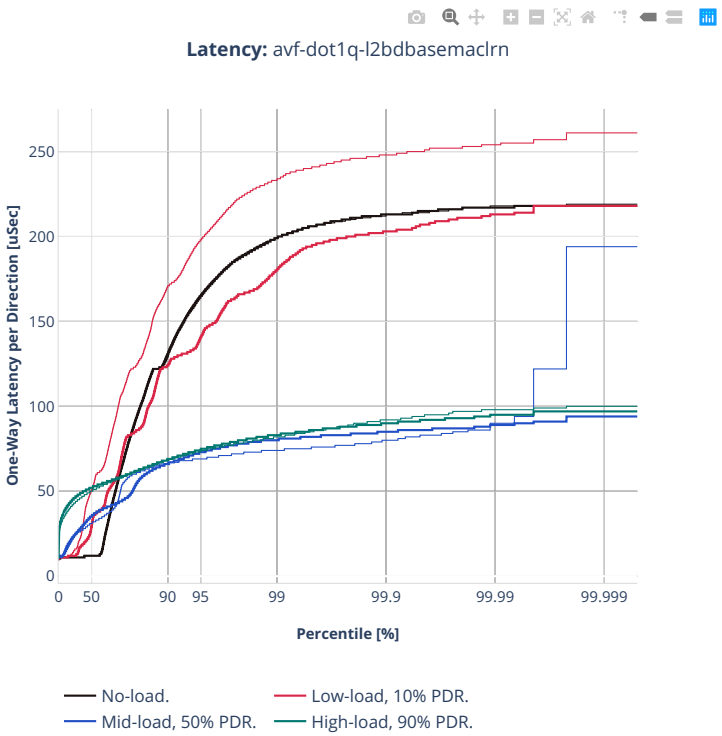

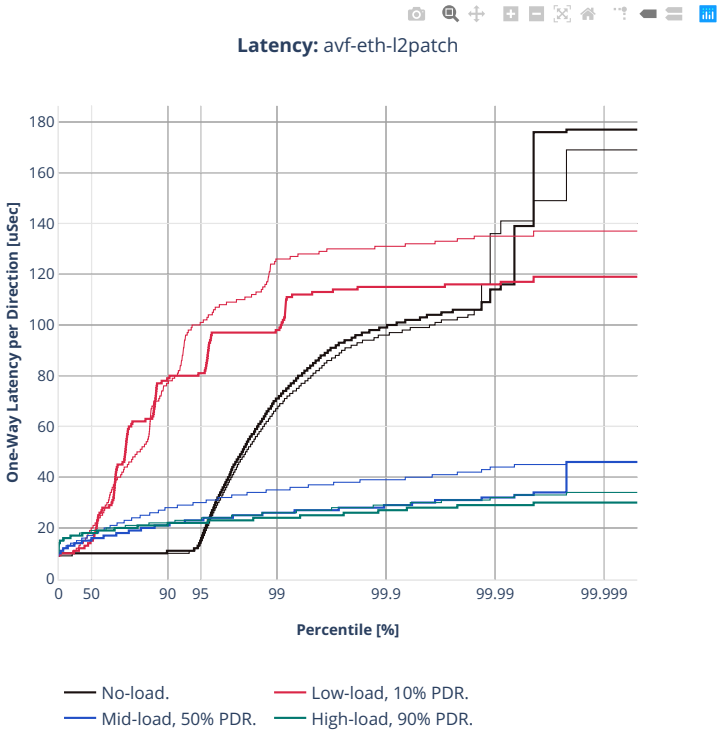
Latency: avf-eth-l2patch
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

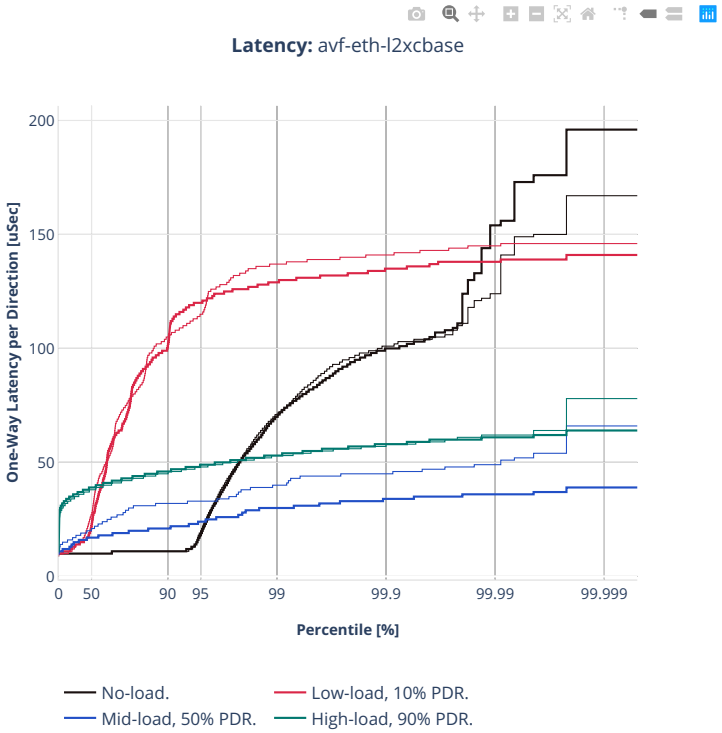
Latency: avf-eth-l2xcbase
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

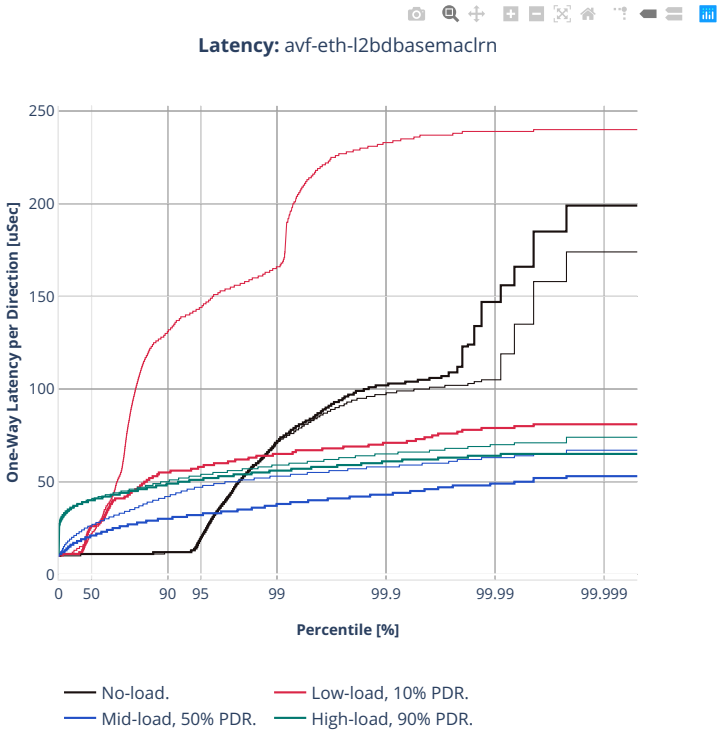
Latency: avf-eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

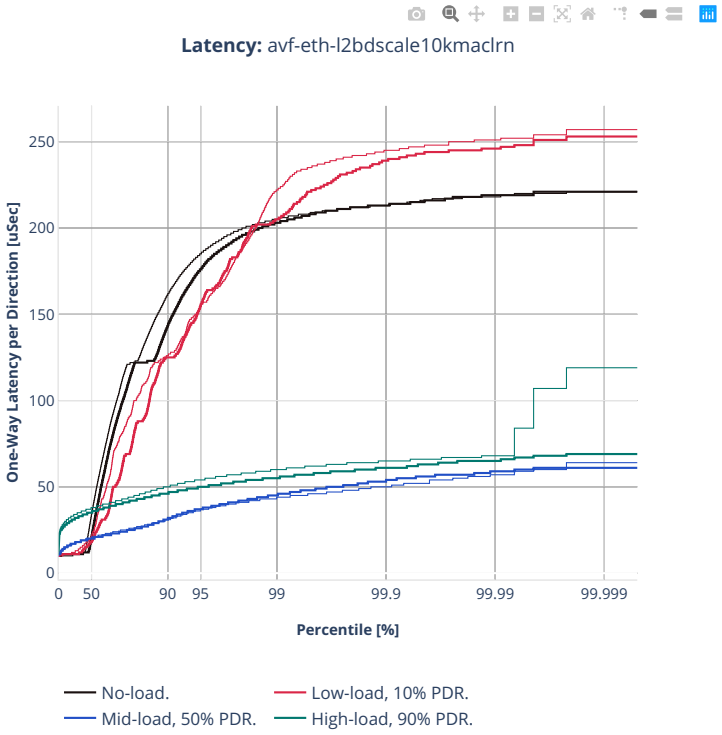
Latency: avf-eth-l2bdscale10kmaclrn
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

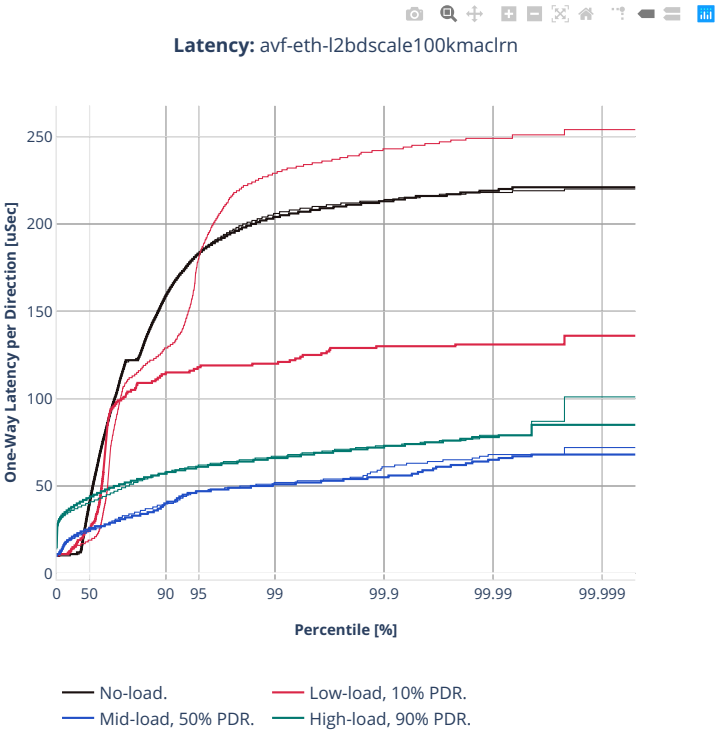
Latency: avf-eth-l2bdscale100kmaclrn
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

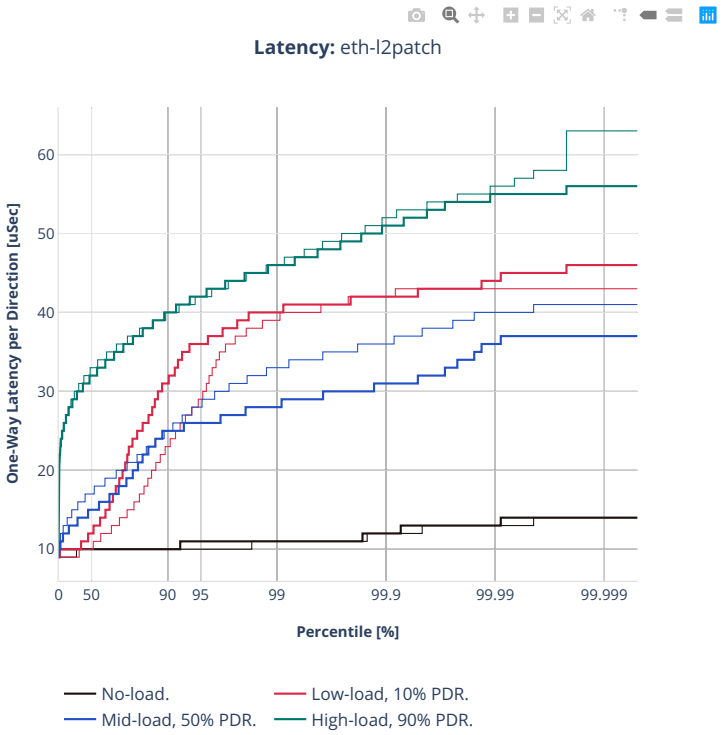

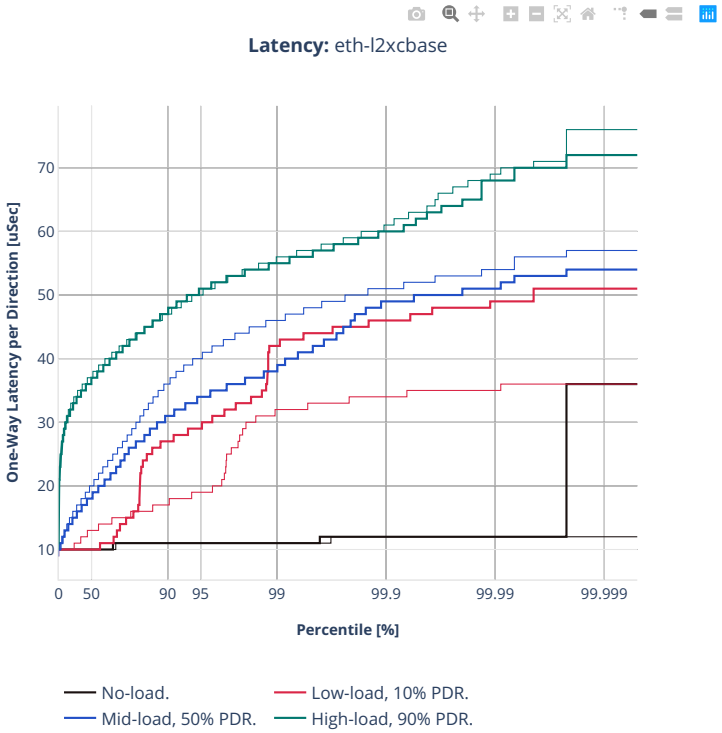
Latency: eth-l2xcbase
One-Way Latency per Direction [uSec]
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

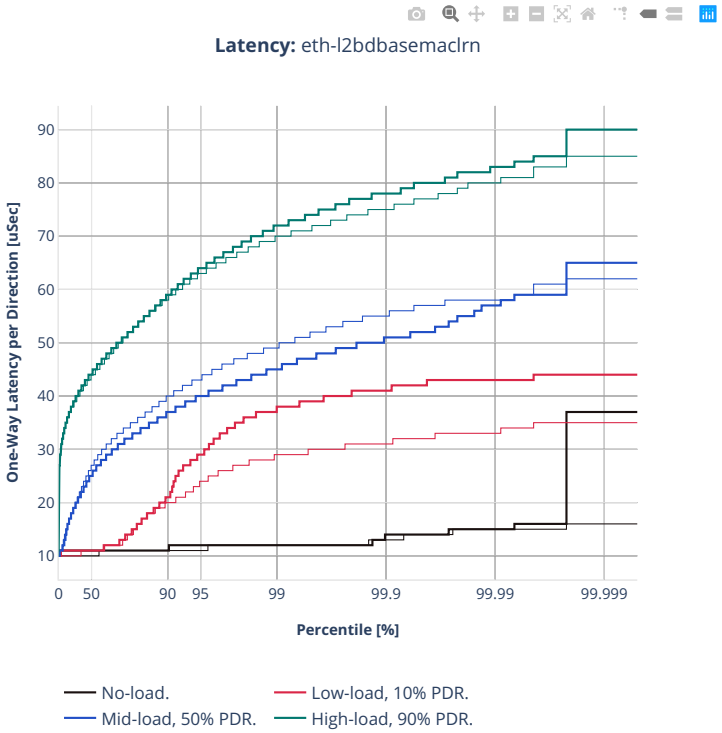
Latency: eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
90
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

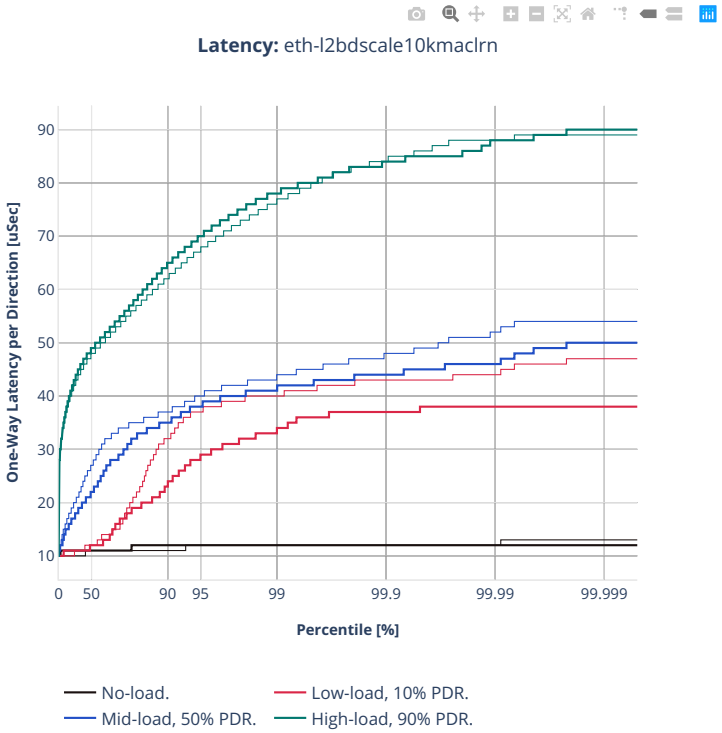
Latency: eth-l2bdscale10kmaclrn
One-Way Latency per Direction [uSec]
90
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

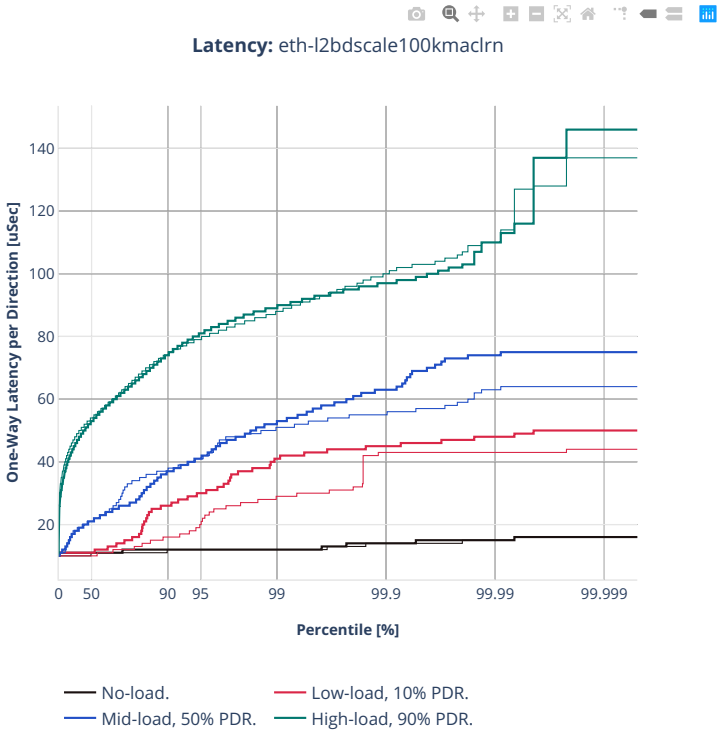
Latency: eth-l2bdscale100kmaclrn
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

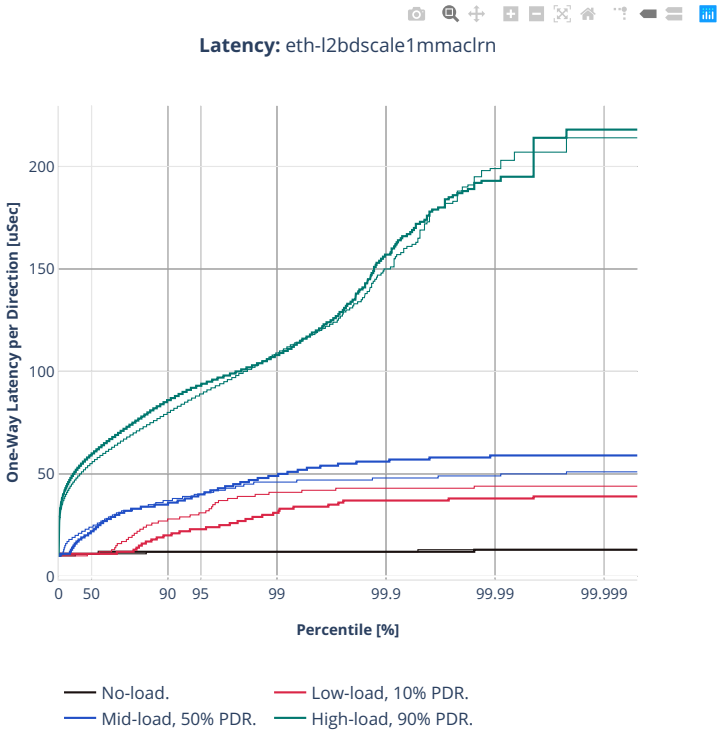
Latency: eth-l2bdscale1mmaclrn
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-clx-cx556a

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

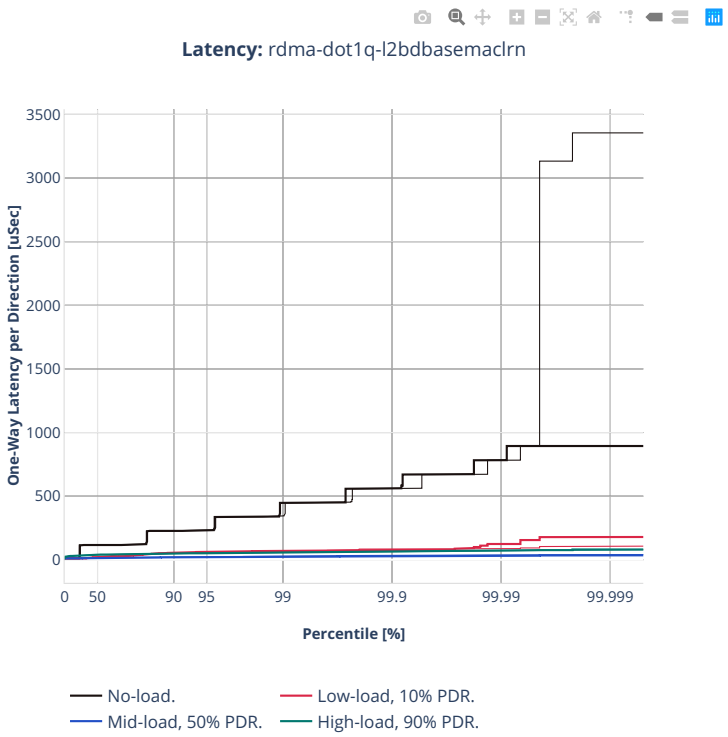

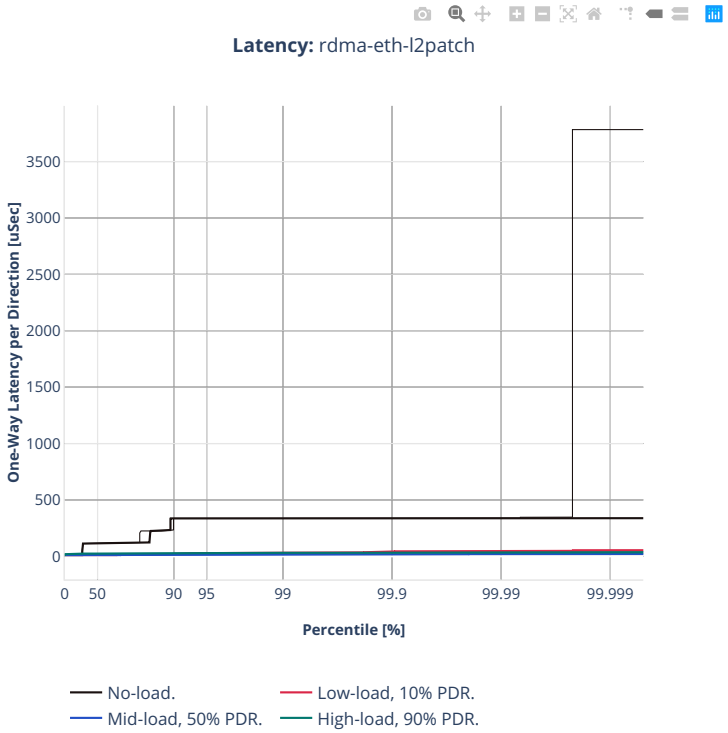
Latency: rdma-eth-l2patch
One-Way Latency per Direction [uSec]
3500
3000
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

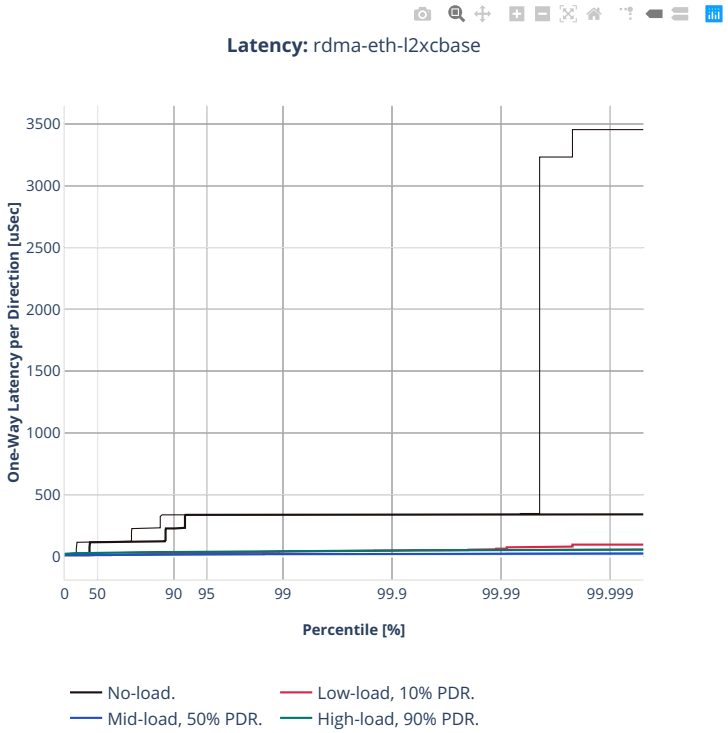
Latency: rdma-eth-l2xcbase
One-Way Latency per Direction [uSec]
3500
3000
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

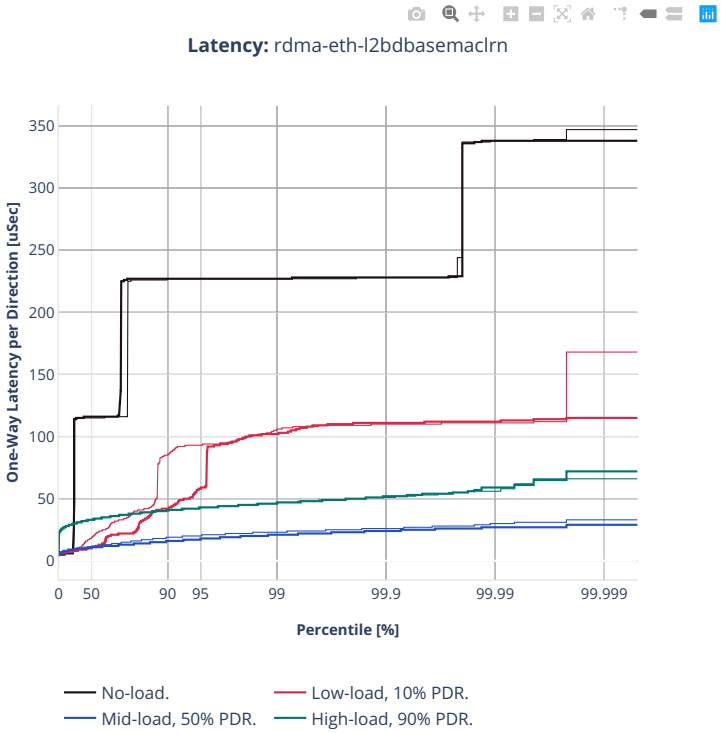
Latency: rdma-eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

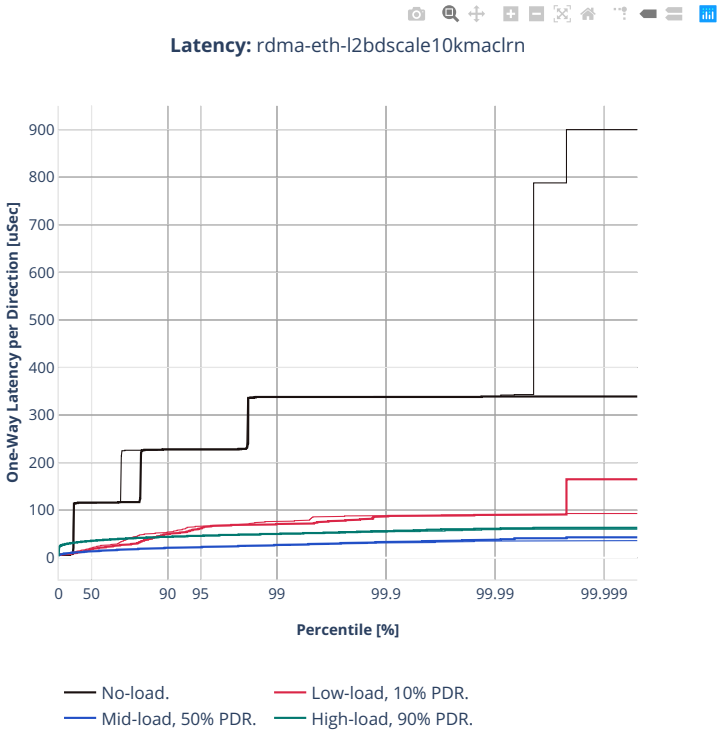
Latency: rdma-eth-l2bdscale10kmaclrn
One-Way Latency per Direction [uSec]
900
800
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

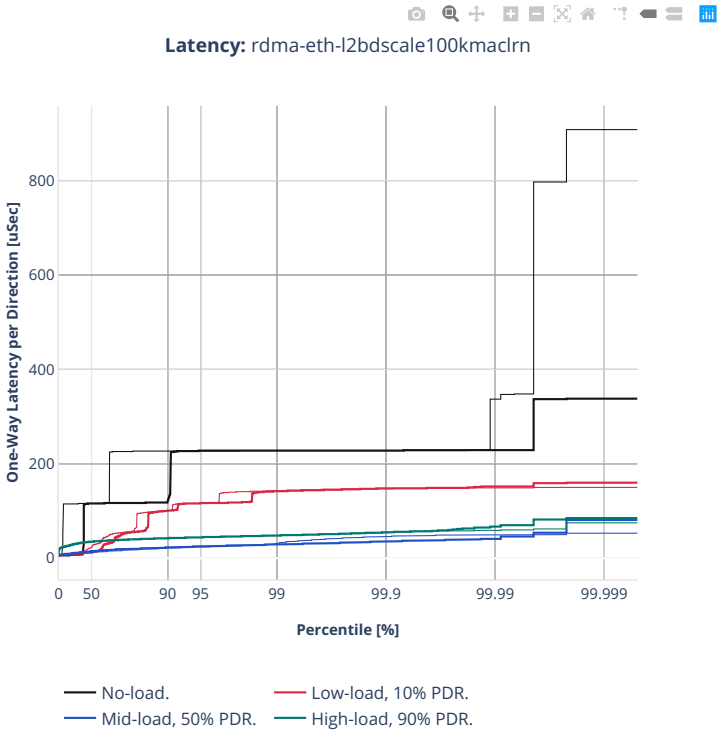
Latency: rdma-eth-l2bdscale100kmaclrn
One-Way Latency per Direction [uSec]
800
600
400
200
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

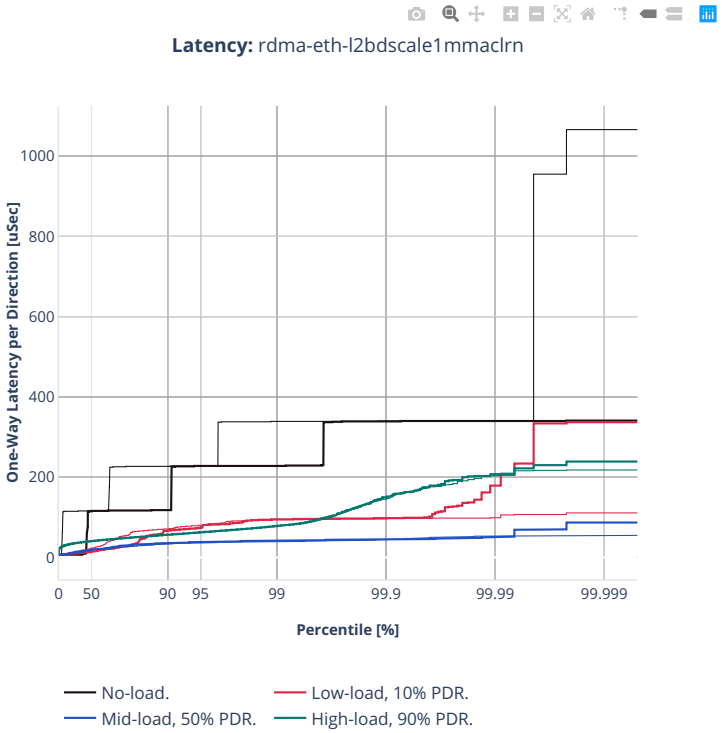
Latency: rdma-eth-l2bdscale1mmaclrn
One-Way Latency per Direction [uSec]
1000
800
600
400
200
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-clx-e810cq

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

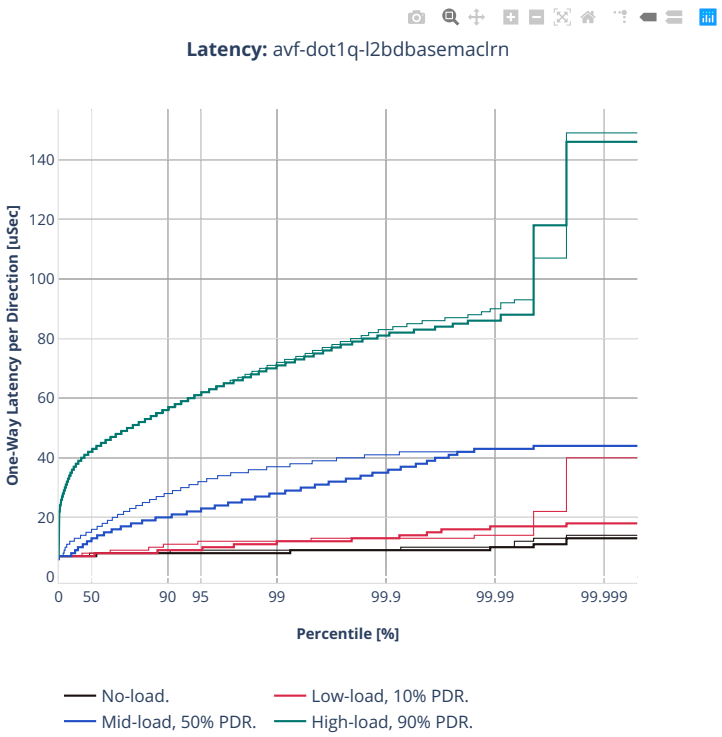

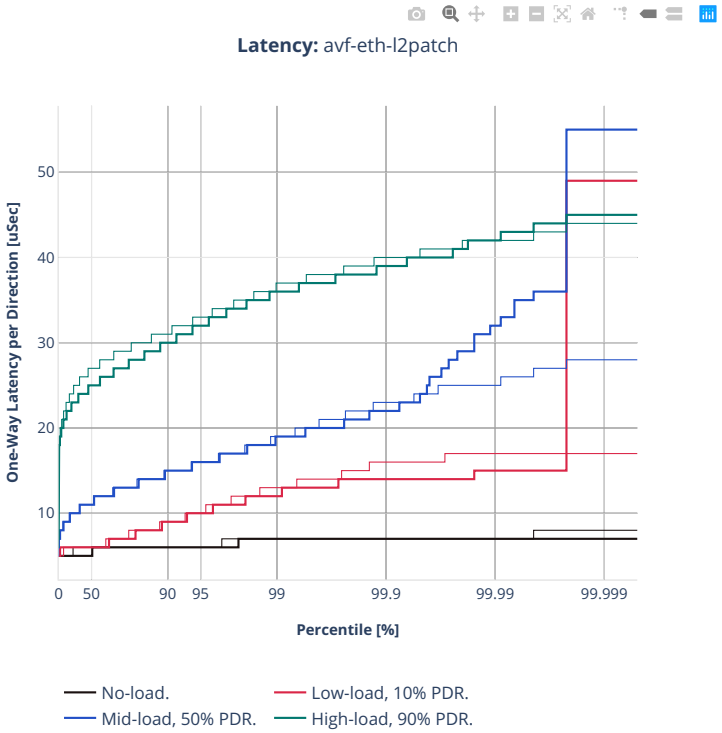
Latency: avf-eth-l2patch
One-Way Latency per Direction [uSec]
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

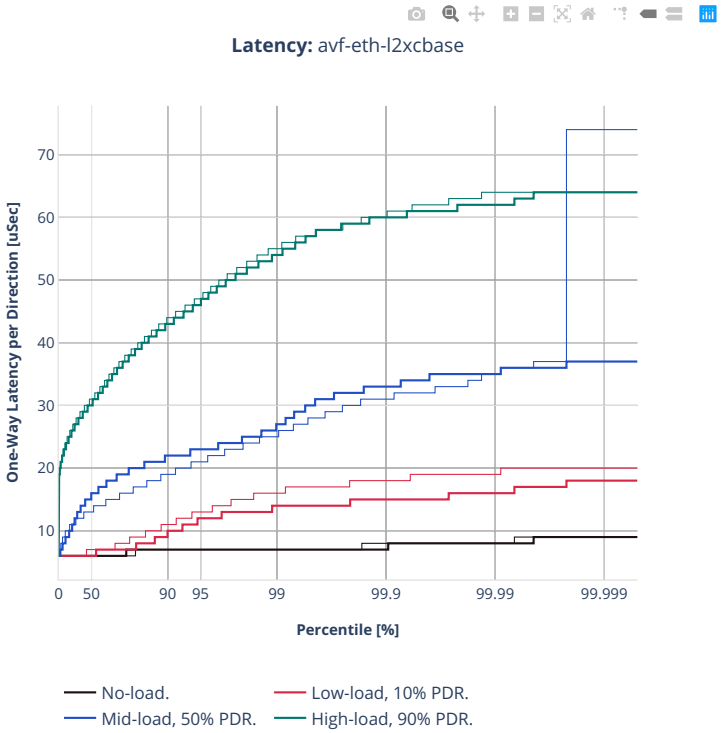
Latency: avf-eth-l2xcbase
One-Way Latency per Direction [uSec]
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

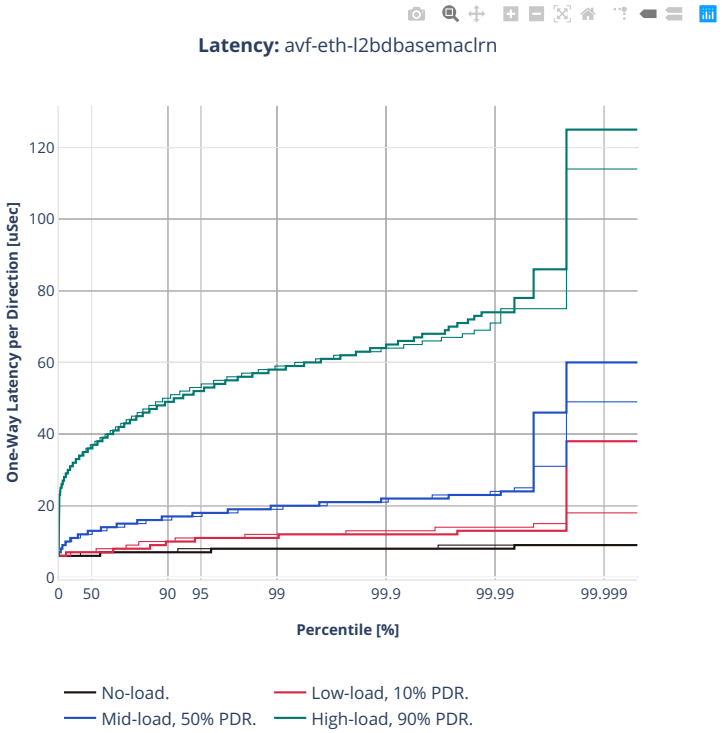
Latency: avf-eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

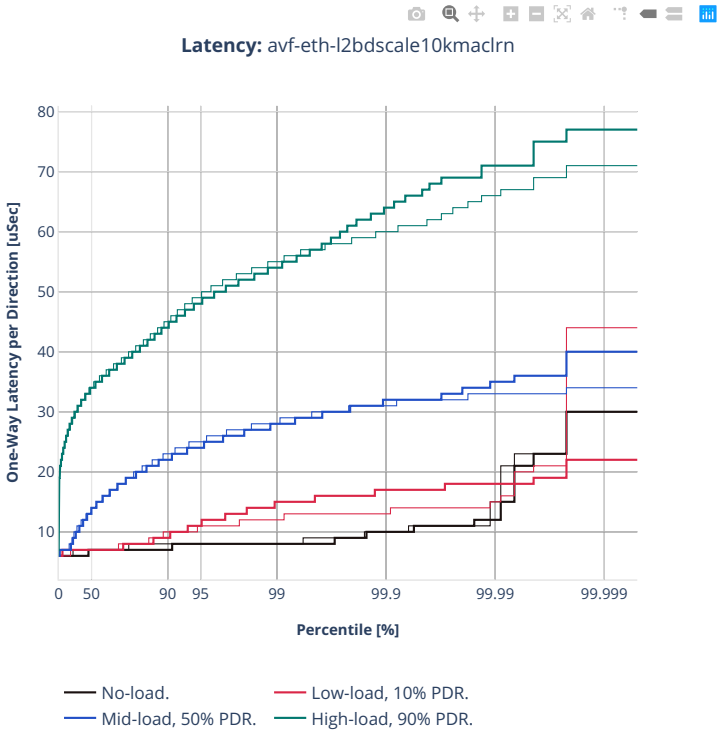
Latency: avf-eth-l2bdscale10kmaclrn
One-Way Latency per Direction [uSec]
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

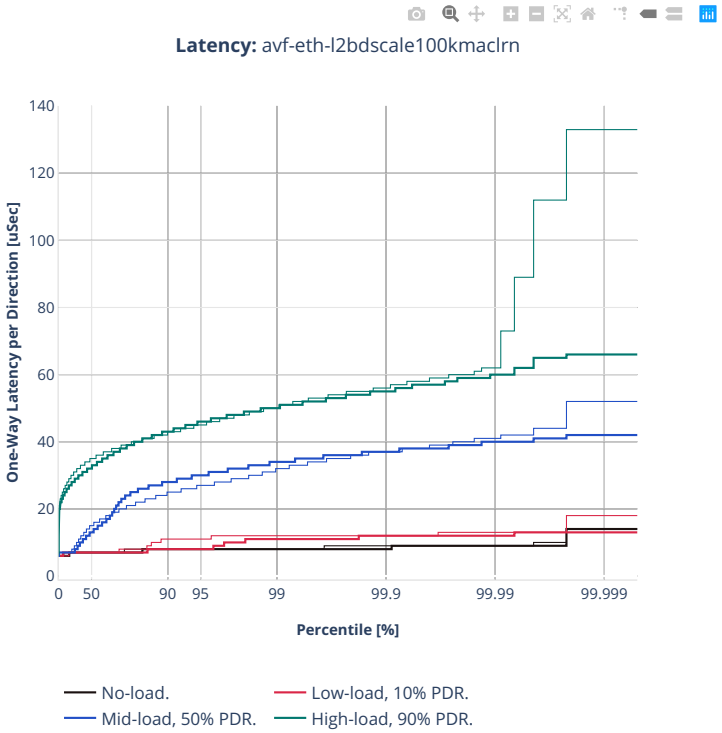
Latency: avf-eth-l2bdscale100kmaclrn
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

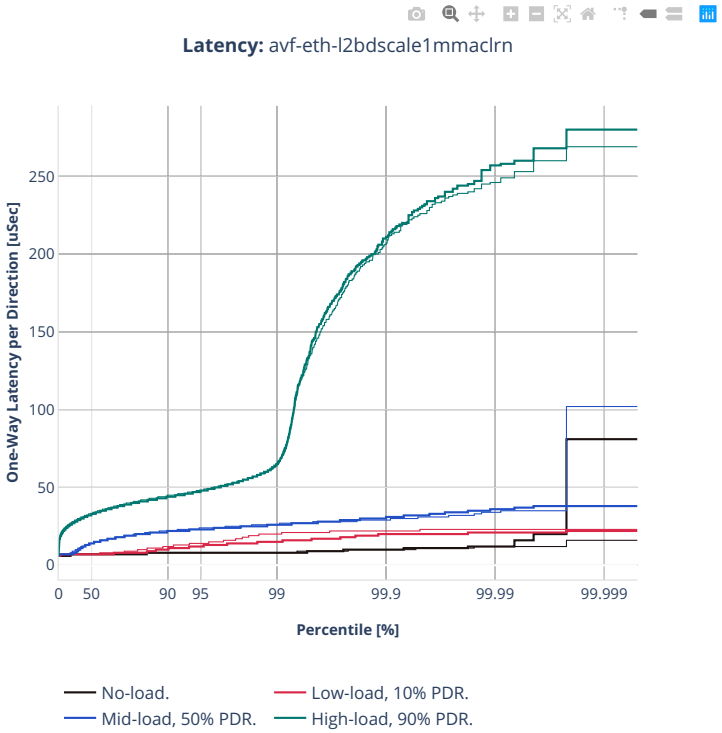
Latency: avf-eth-l2bdscale1mmaclrn
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

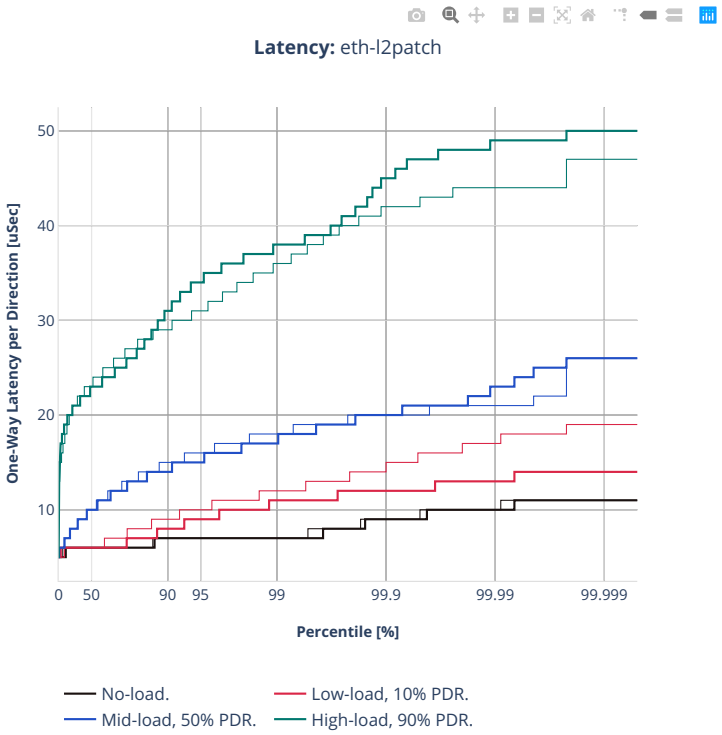

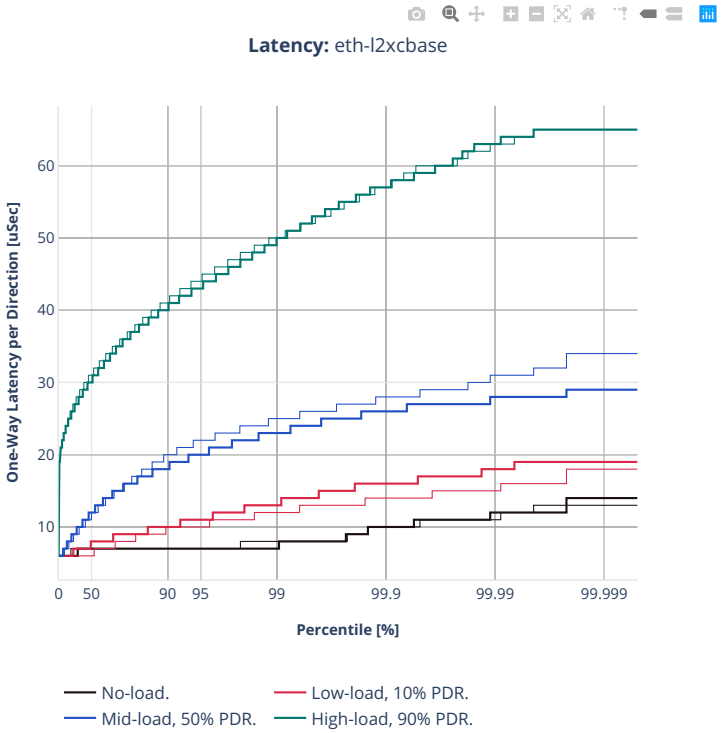
Latency: eth-l2xcbase
One-Way Latency per Direction [uSec]
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

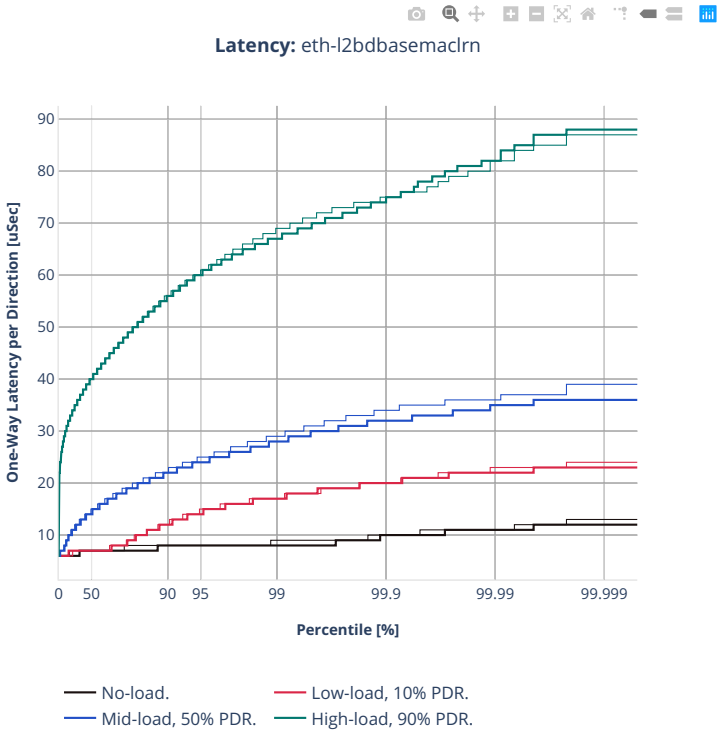
Latency: eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
90
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

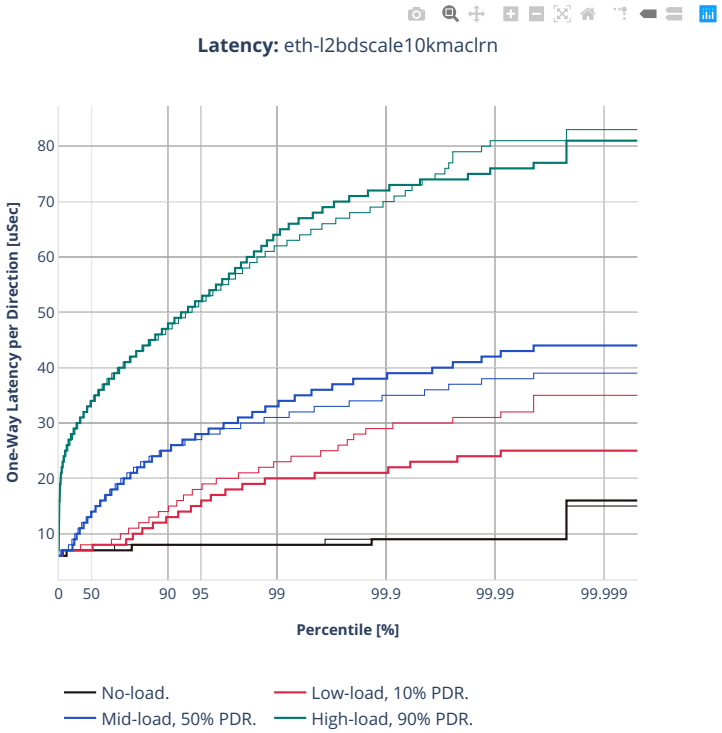
Latency: eth-l2bdscale10kmaclrn
One-Way Latency per Direction [uSec]
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-zn2-xxv710

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

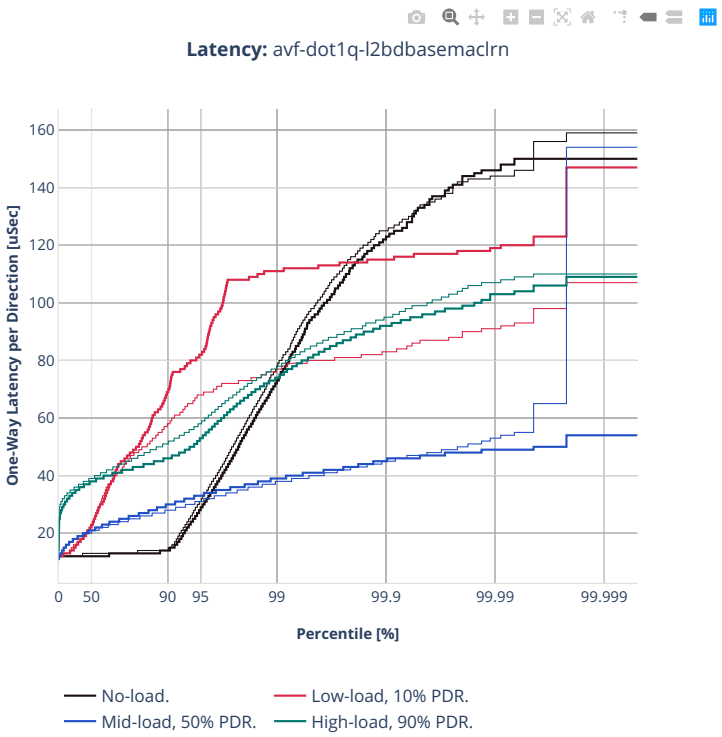

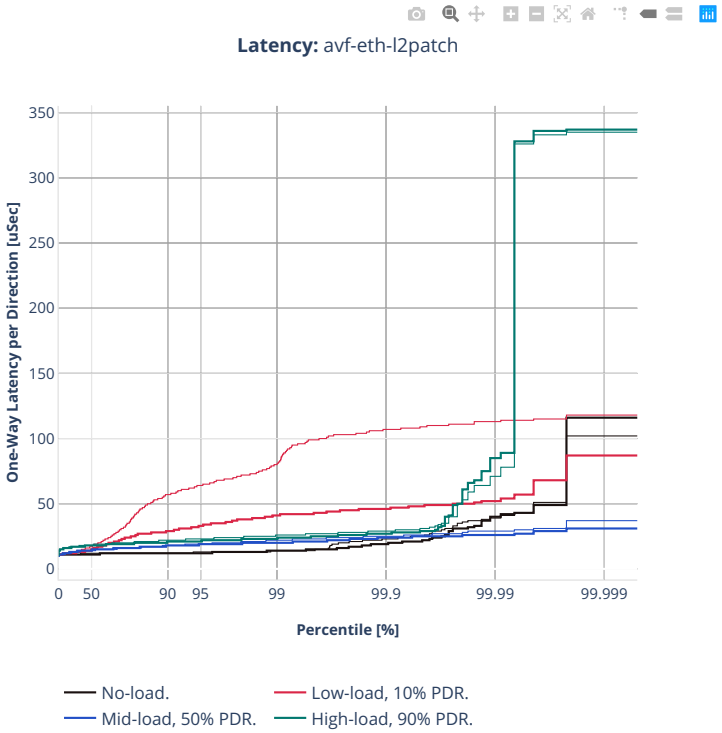
Latency: avf-eth-l2patch
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

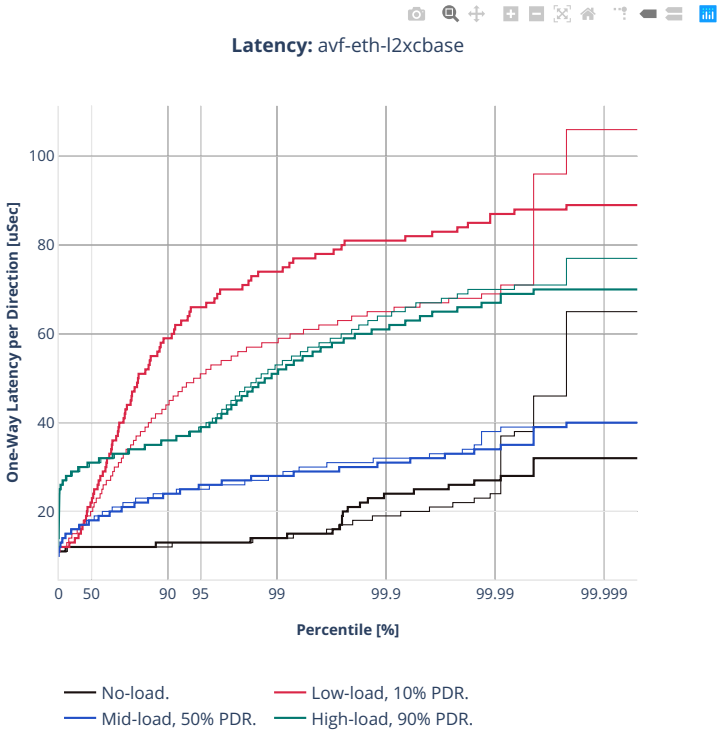
Latency: avf-eth-l2xcbase
One-Way Latency per Direction [uSec]
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

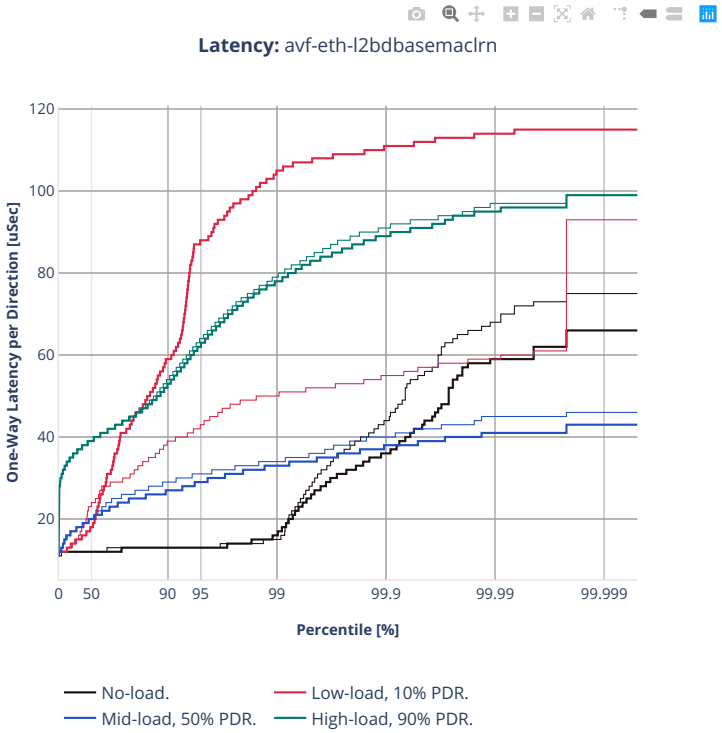
Latency: avf-eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

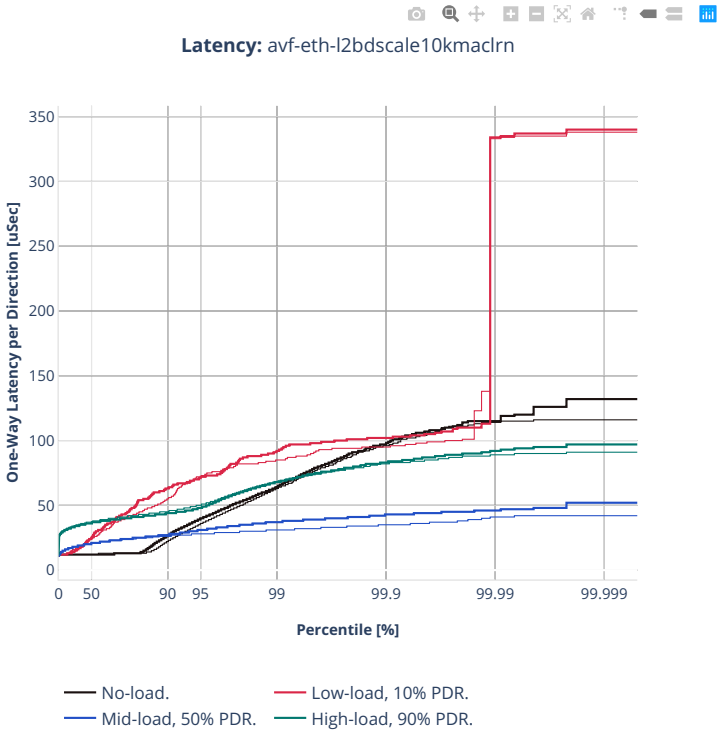
Latency: avf-eth-l2bdscale10kmaclrn
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

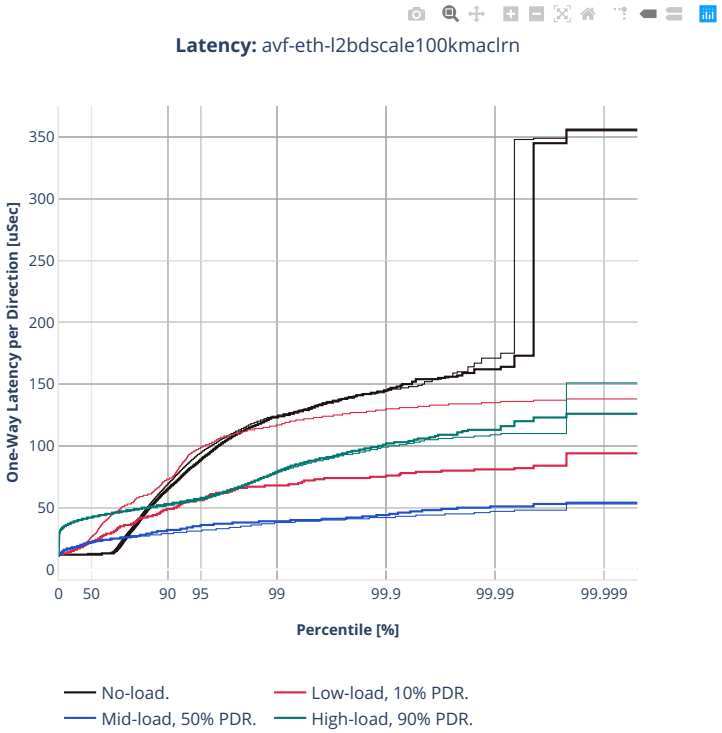
Latency: avf-eth-l2bdscale100kmaclrn
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

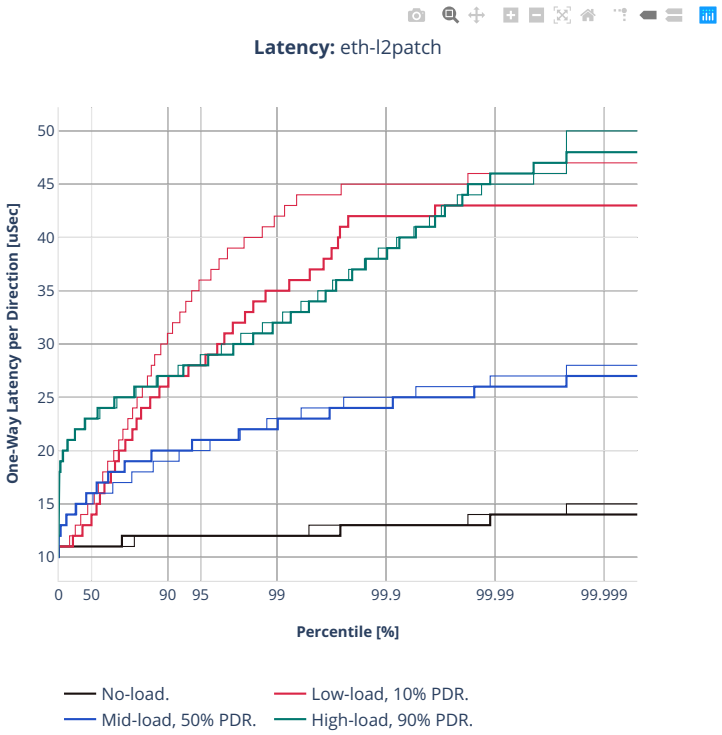

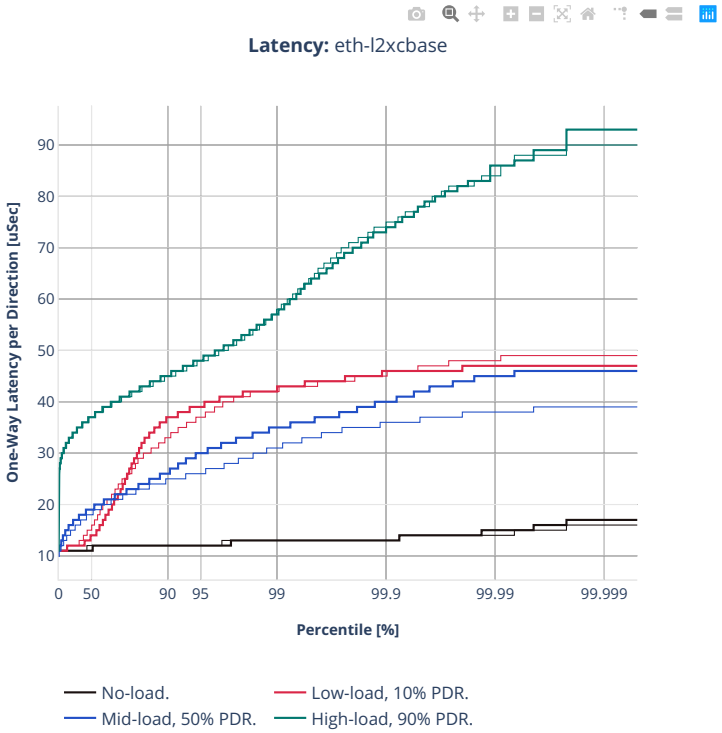
Latency: eth-l2xcbase
One-Way Latency per Direction [uSec]
90
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

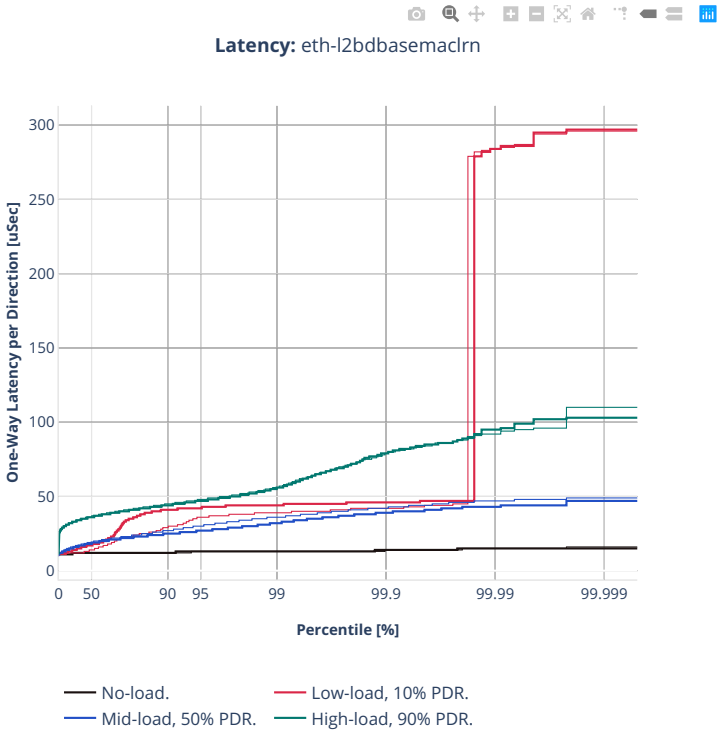
Latency: eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

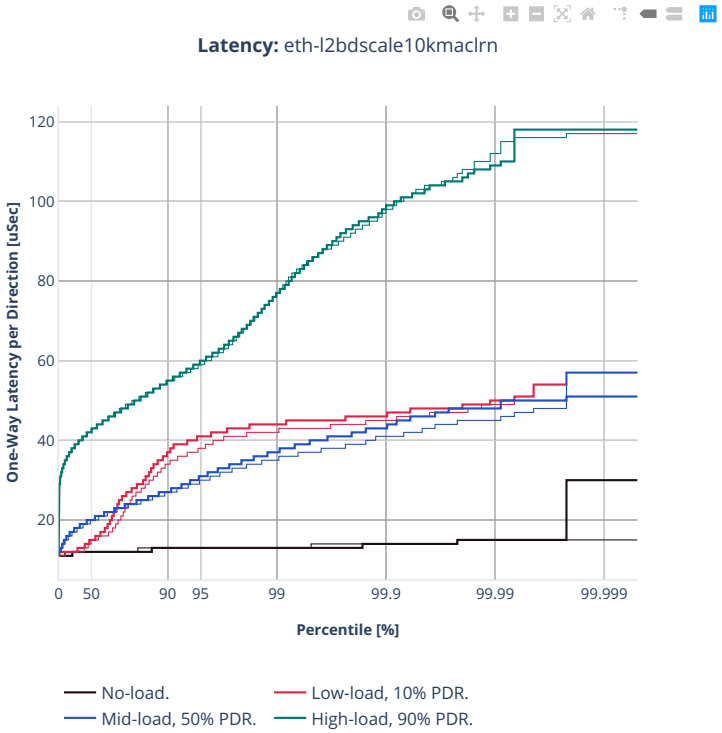
Latency: eth-l2bdscale10kmaclrn
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

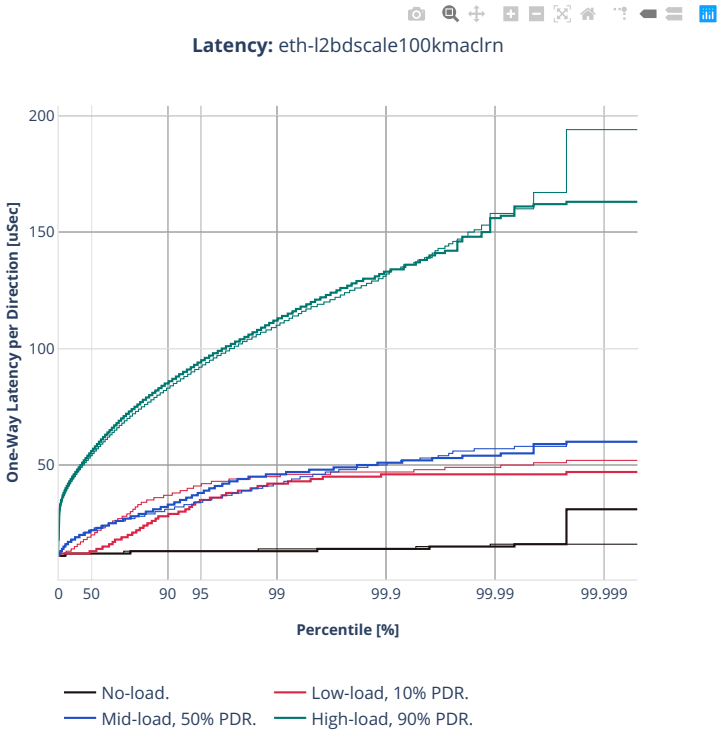
Latency: eth-l2bdscale100kmaclrn
One-Way Latency per Direction [uSec]
200
150
100
50
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-alt-xl710

64b-1t1c-l2switching-base-scale

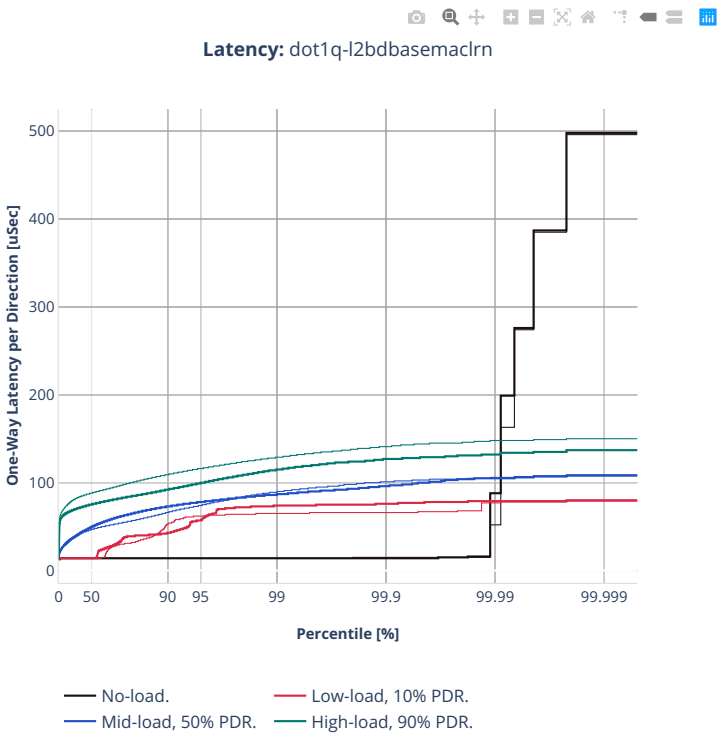

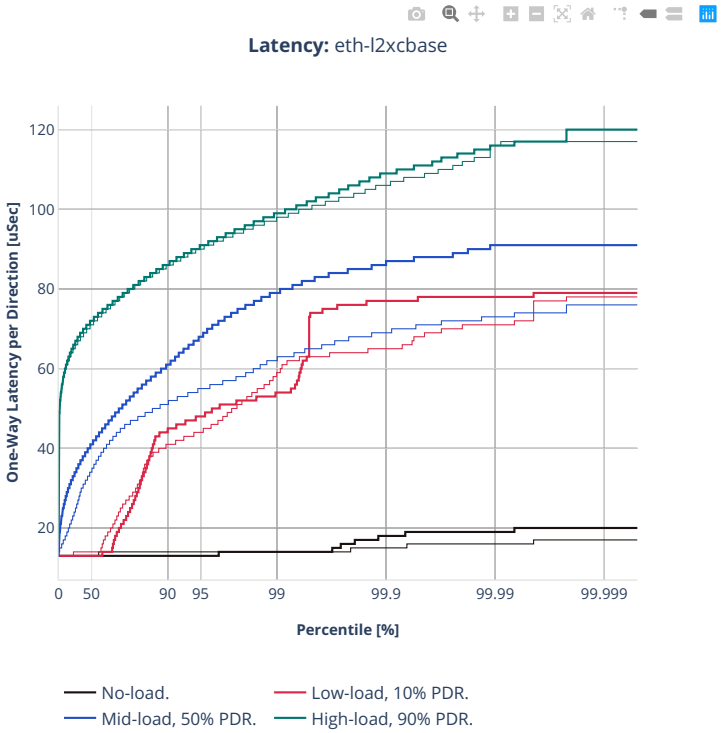
Latency: eth-l2xcbase
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

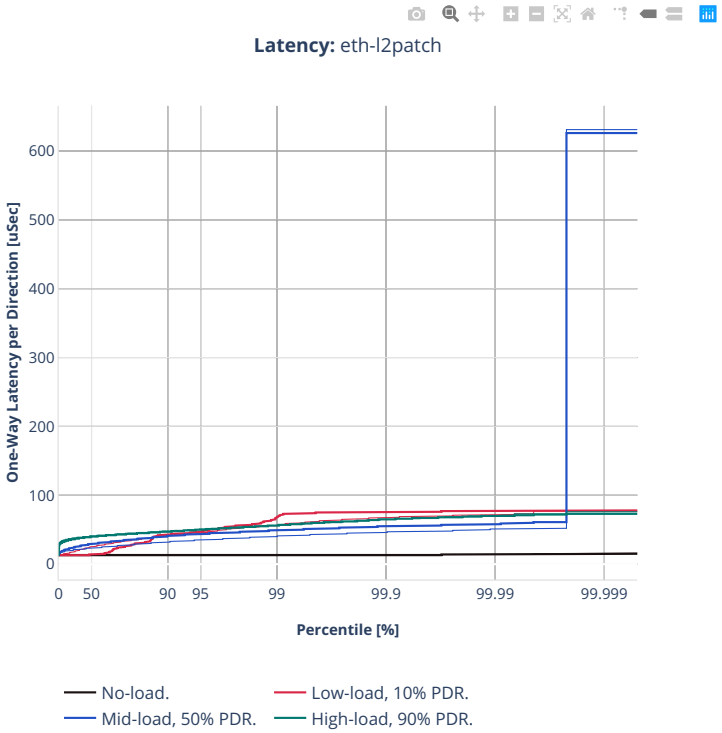
Latency: eth-l2patch
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

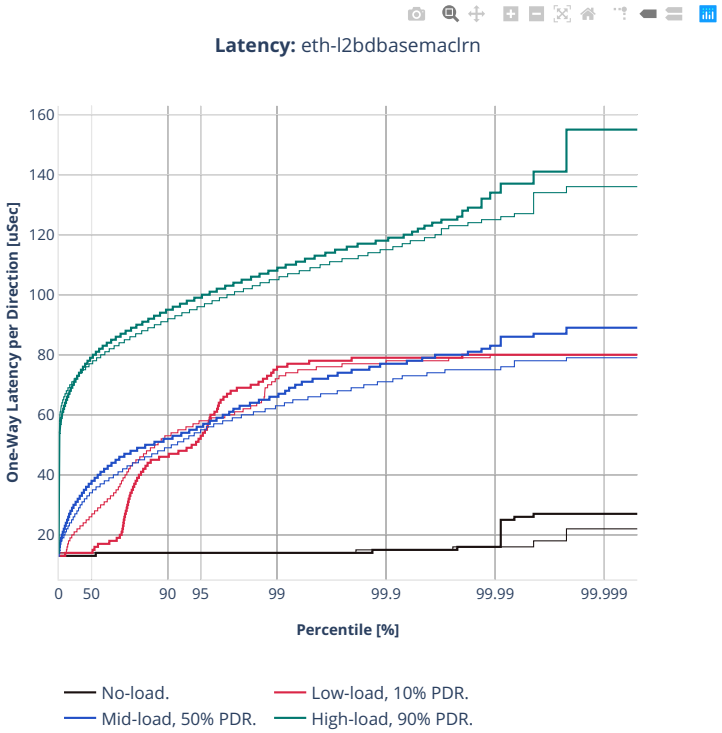
Latency: eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

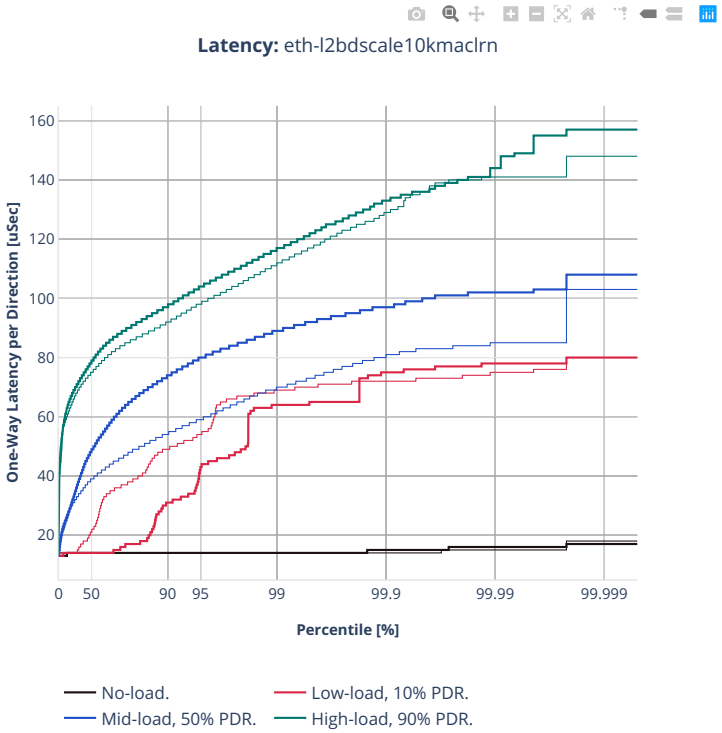
Latency: eth-l2bdscale10kmaclrn
One-Way Latency per Direction [uSec]
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

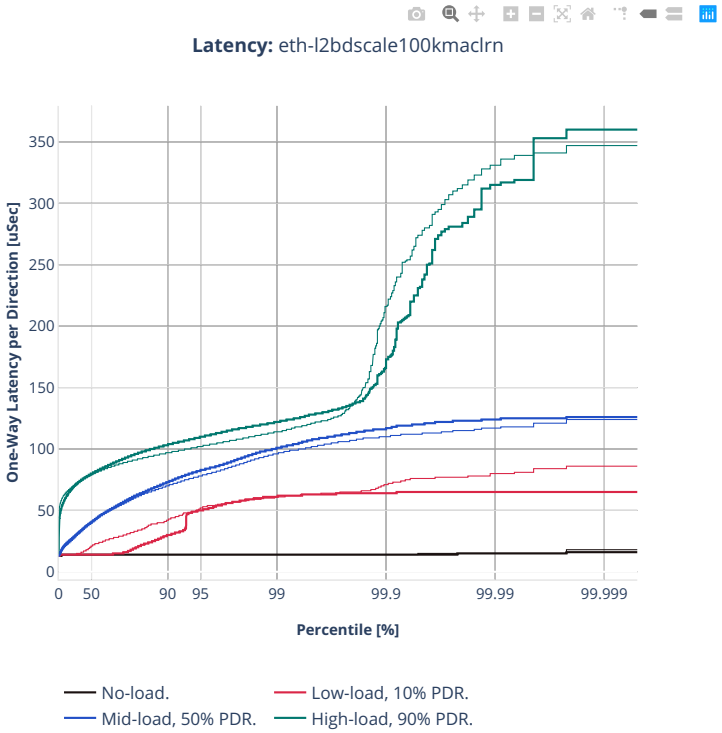
Latency: eth-l2bdscale100kmaclrn
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-1t1c-features-l2switching-base

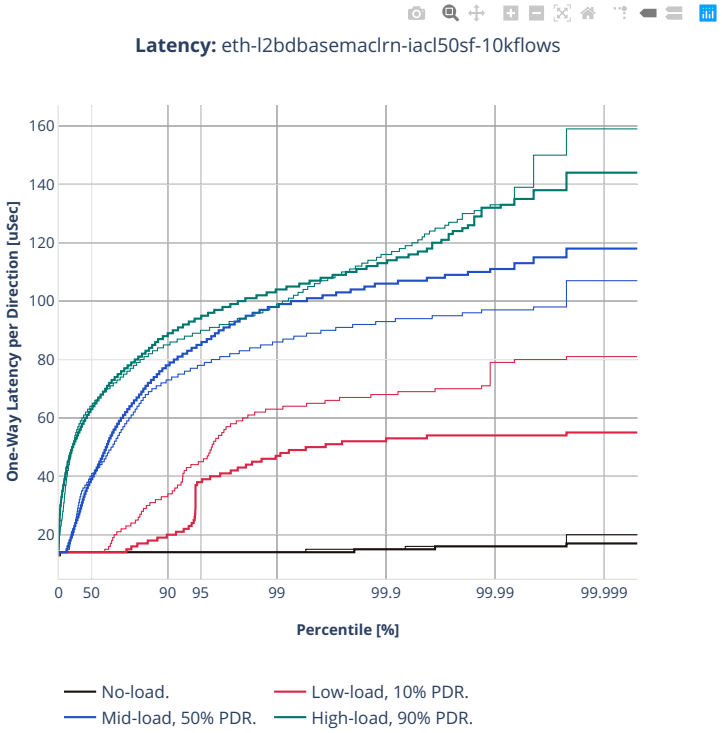

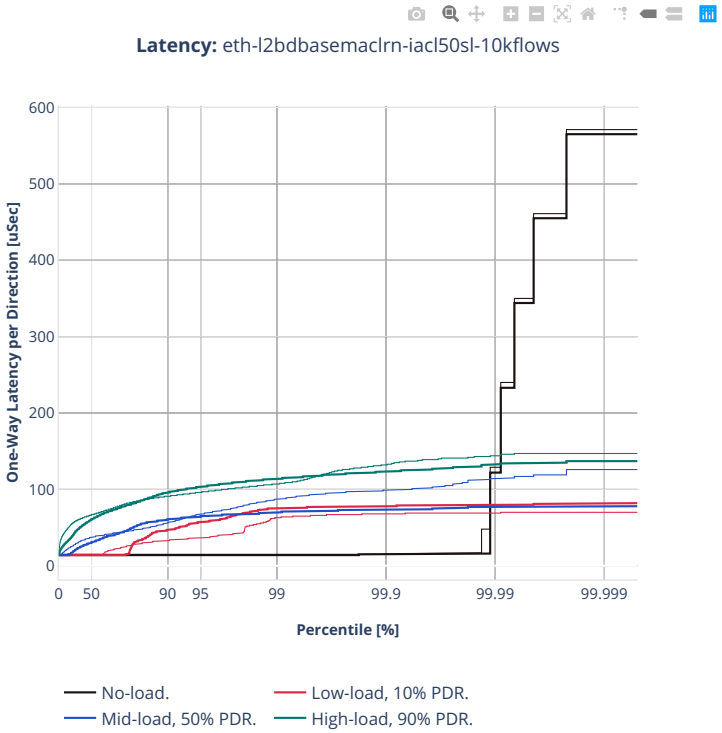
Latency: eth-l2bdbasemaclrn-iacl50sl-10kflows
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

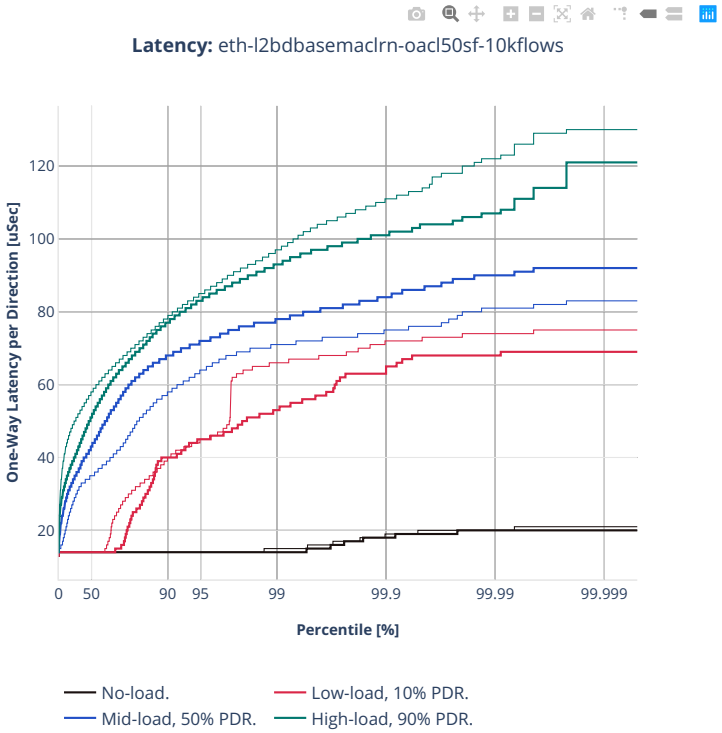
Latency: eth-l2bdbasemaclrn-oacl50sf-10kflows
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

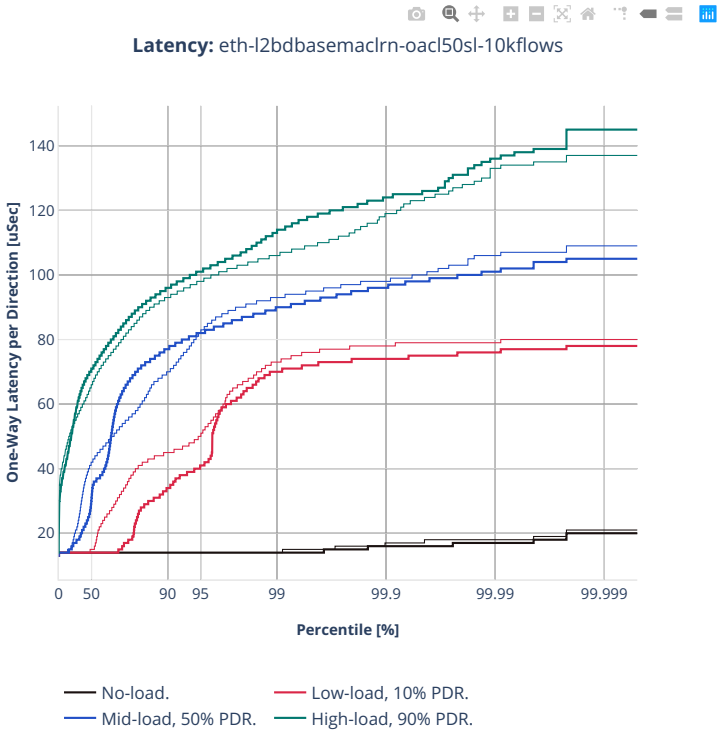
Latency: eth-l2bdbasemaclrn-oacl50sl-10kflows
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

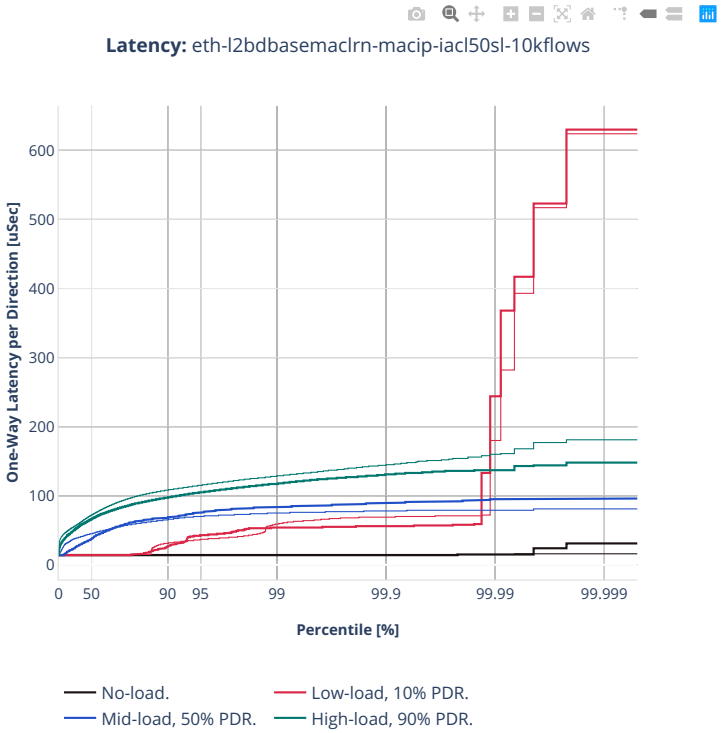
Latency: eth-l2bdbasemaclrn-macip-iacl50sl-10kflows
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-tsh-x520

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

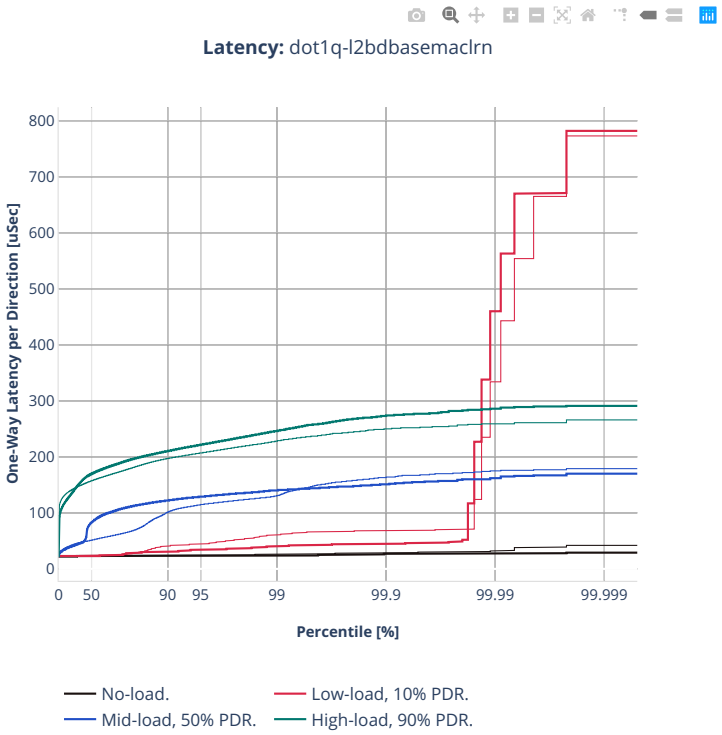

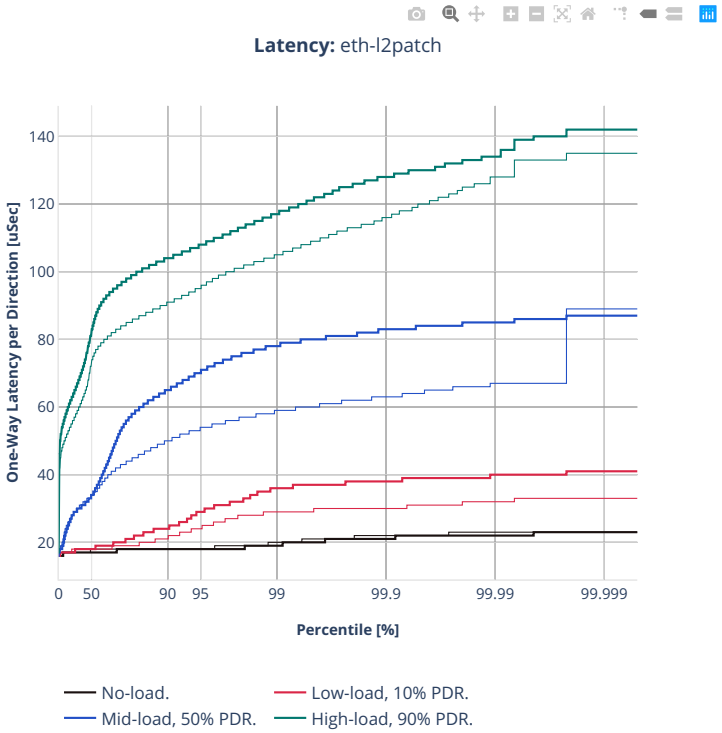
Latency: eth-l2patch
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

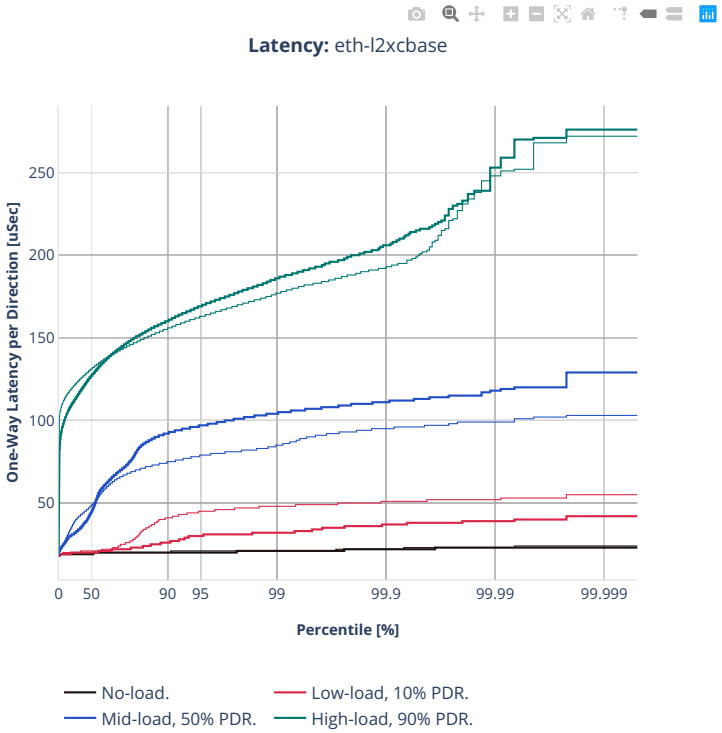
Latency: eth-l2xcbase
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

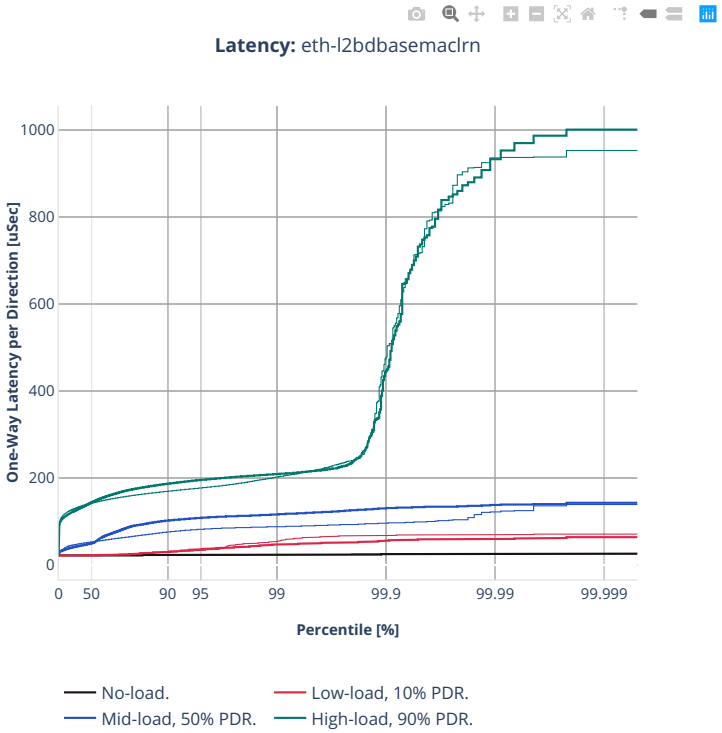
Latency: eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
1000
800
600
400
200
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-1t1c-features-l2switching-base-ixgbe

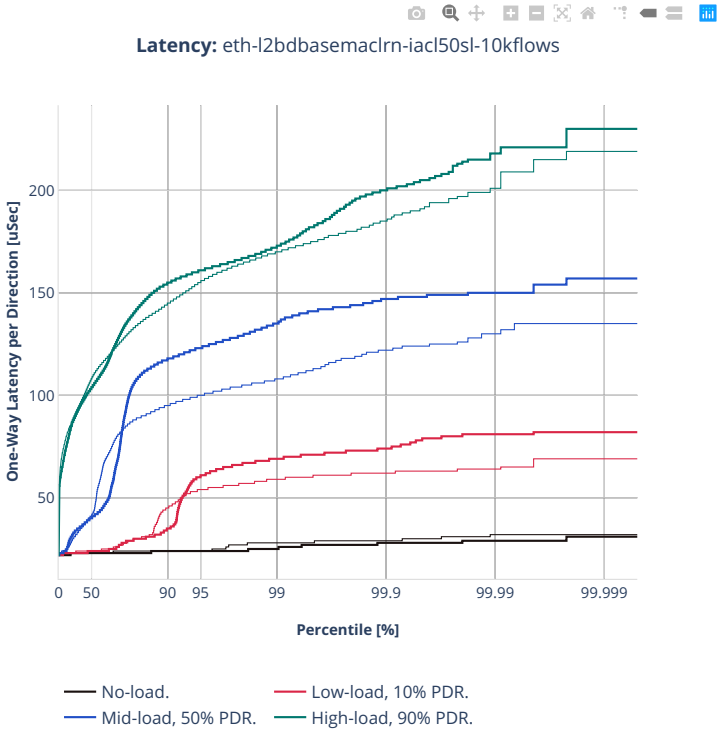

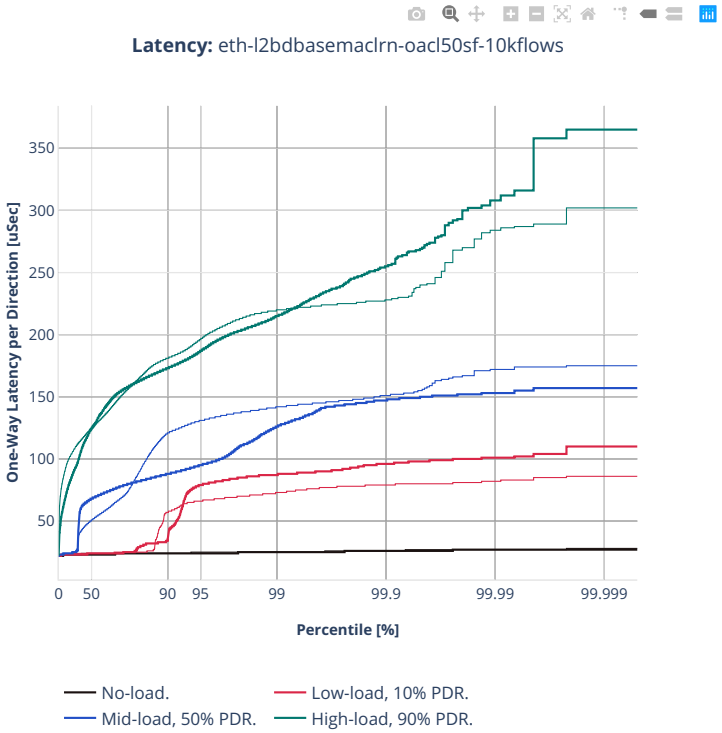
Latency: eth-l2bdbasemaclrn-oacl50sf-10kflows
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

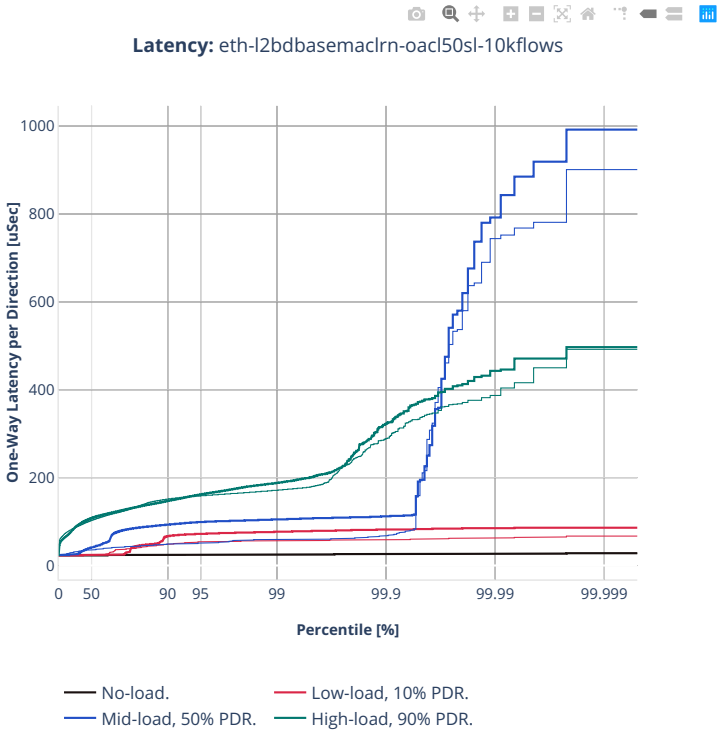
Latency: eth-l2bdbasemaclrn-oacl50sl-10kflows
One-Way Latency per Direction [uSec]
1000
800
600
400
200
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-tx2-xl710

64b-1t1c-l2switching-base-dpdk

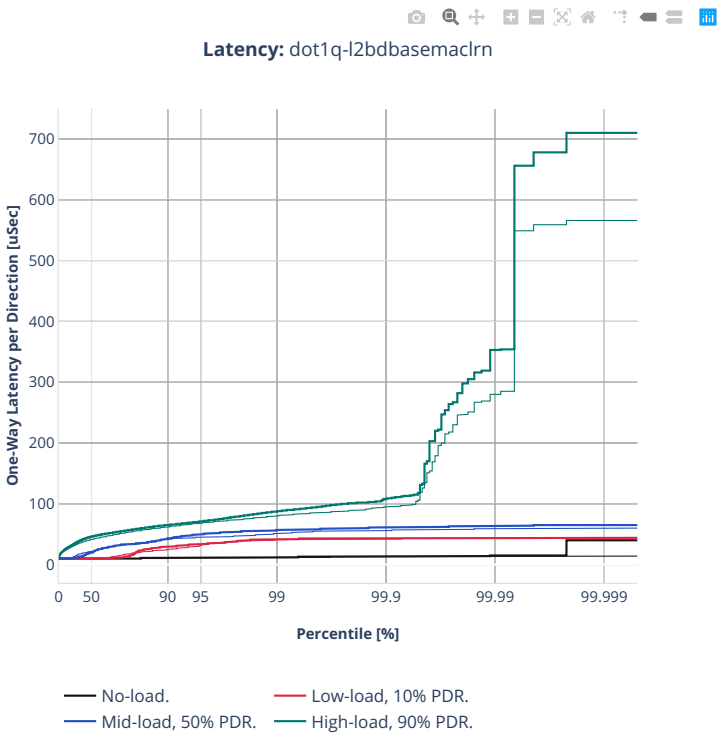

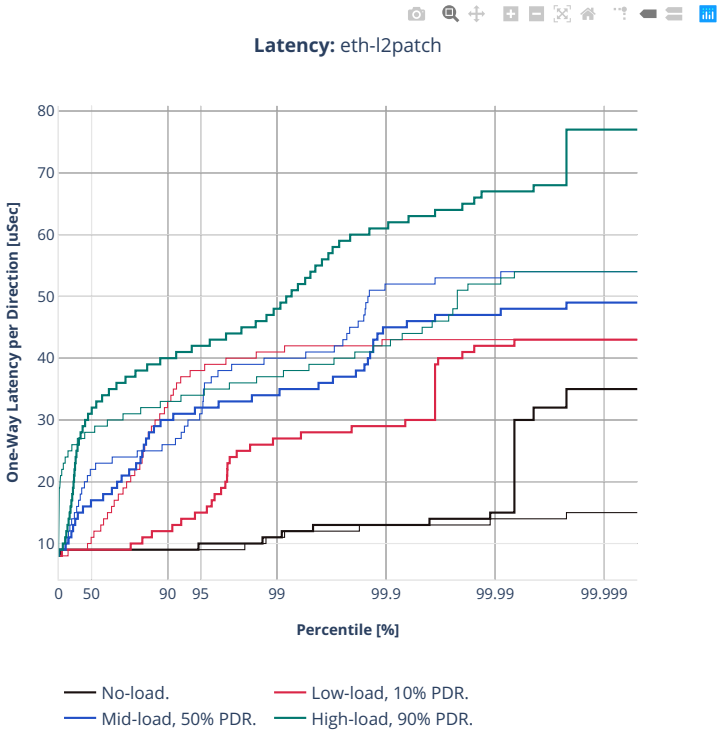
Latency: eth-l2patch
One-Way Latency per Direction [uSec]
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

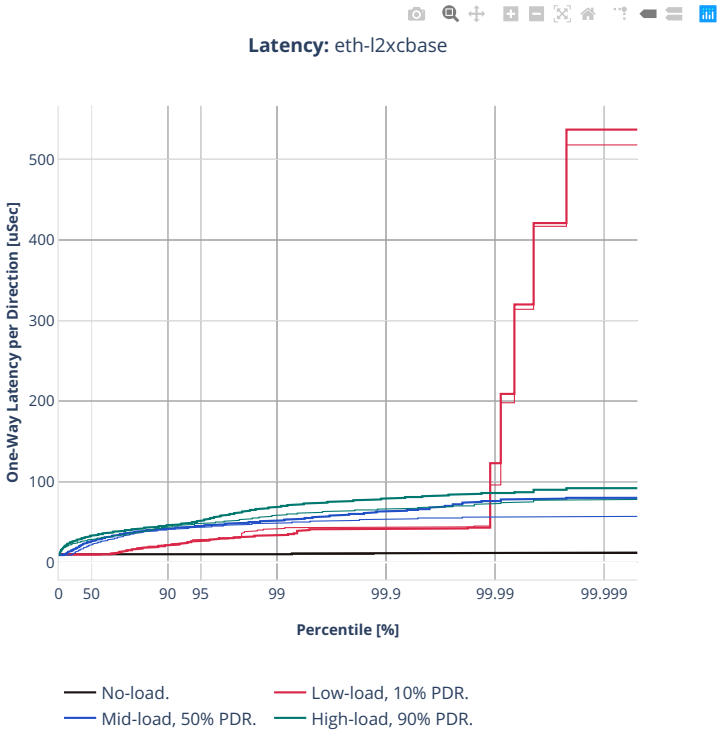
Latency: eth-l2xcbase
One-Way Latency per Direction [uSec]
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

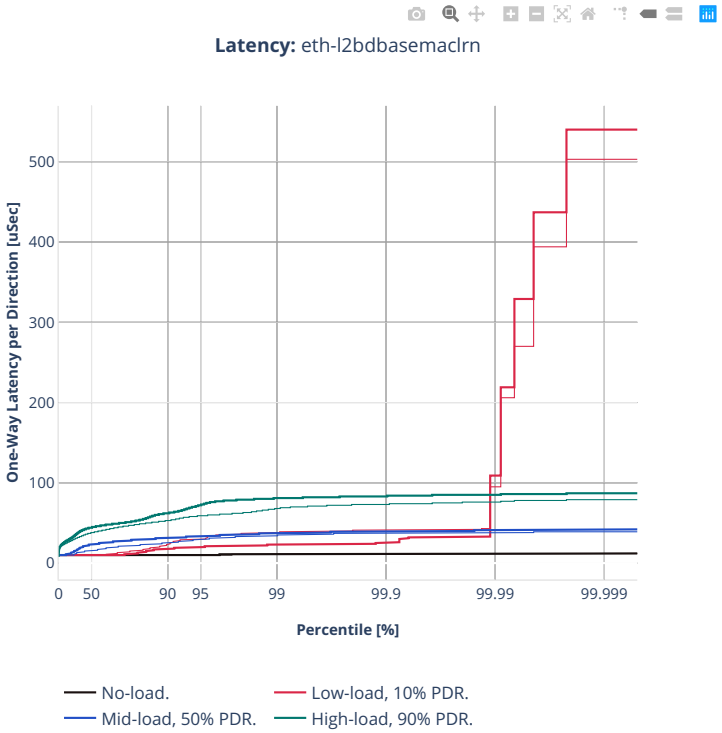
Latency: eth-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-1t1c-l2switching-scale-dpdk

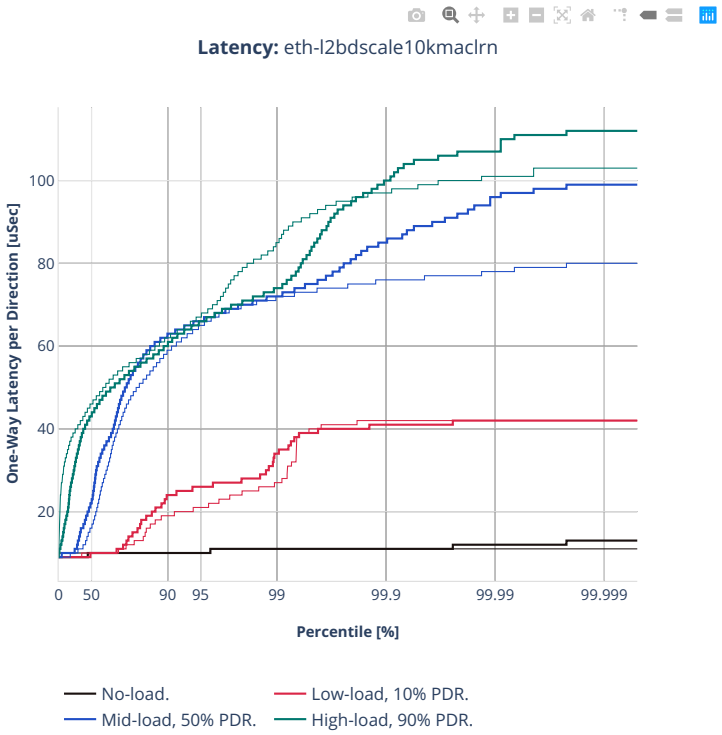

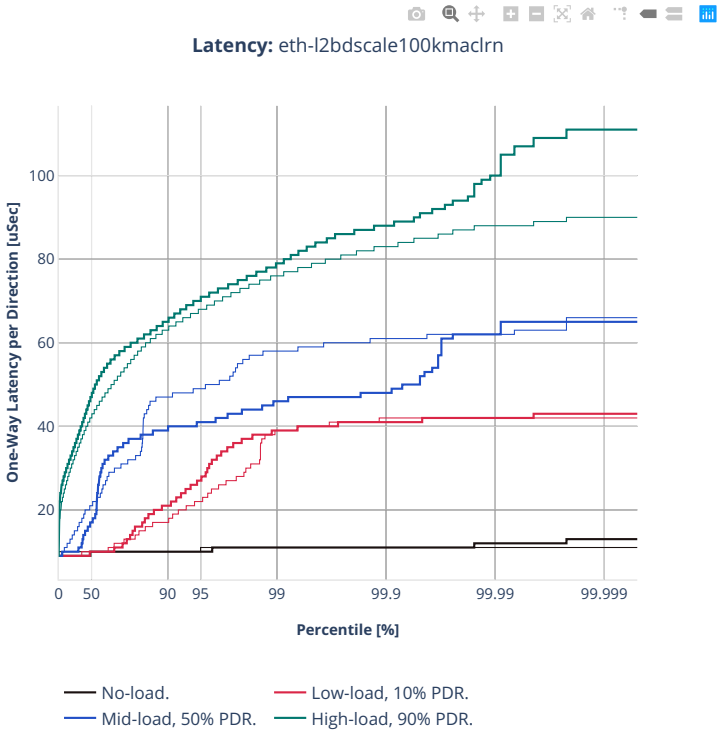
Latency: eth-l2bdscale100kmaclrn
One-Way Latency per Direction [uSec]
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-1t1c-features-l2switching-base-dpdk

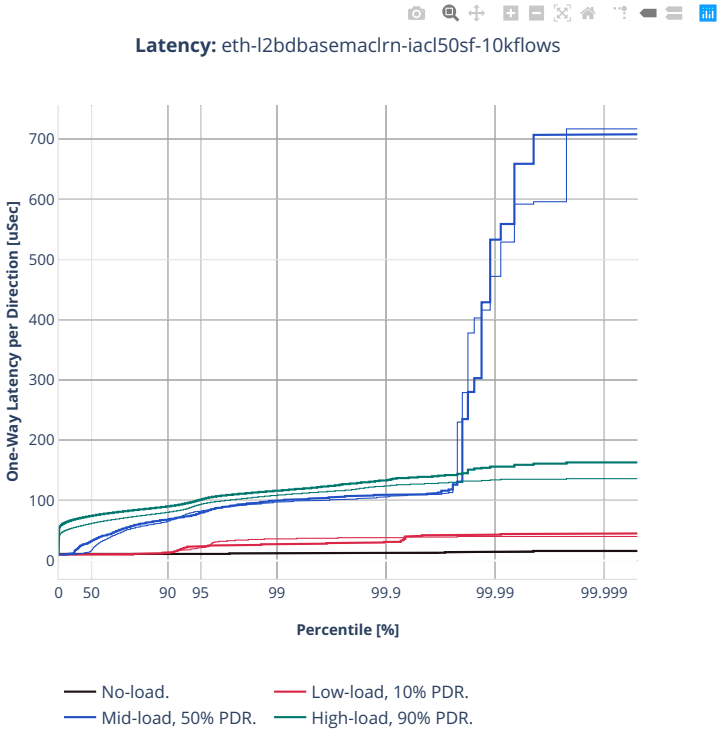

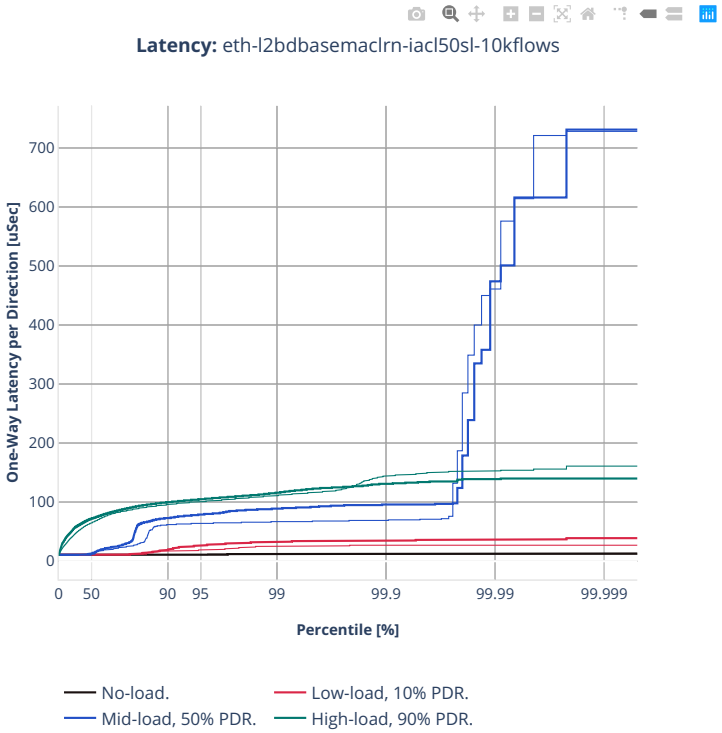
Latency: eth-l2bdbasemaclrn-iacl50sl-10kflows
One-Way Latency per Direction [uSec]
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

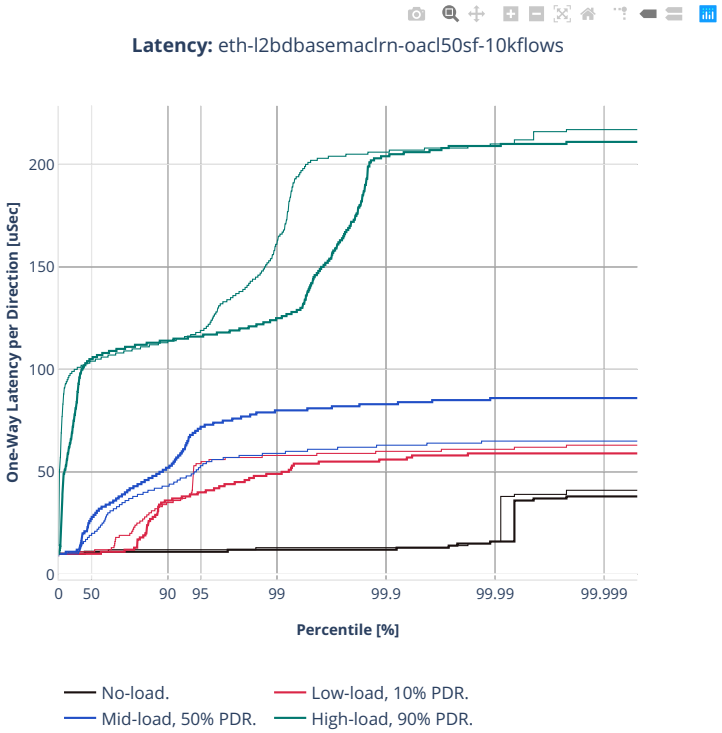
Latency: eth-l2bdbasemaclrn-oacl50sf-10kflows
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

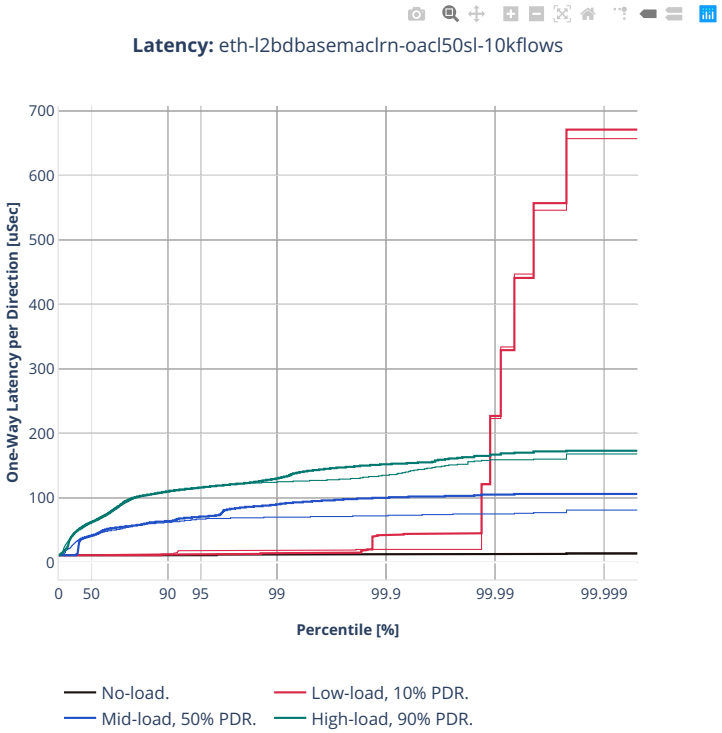
Latency: eth-l2bdbasemaclrn-oacl50sl-10kflows
One-Way Latency per Direction [uSec]
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

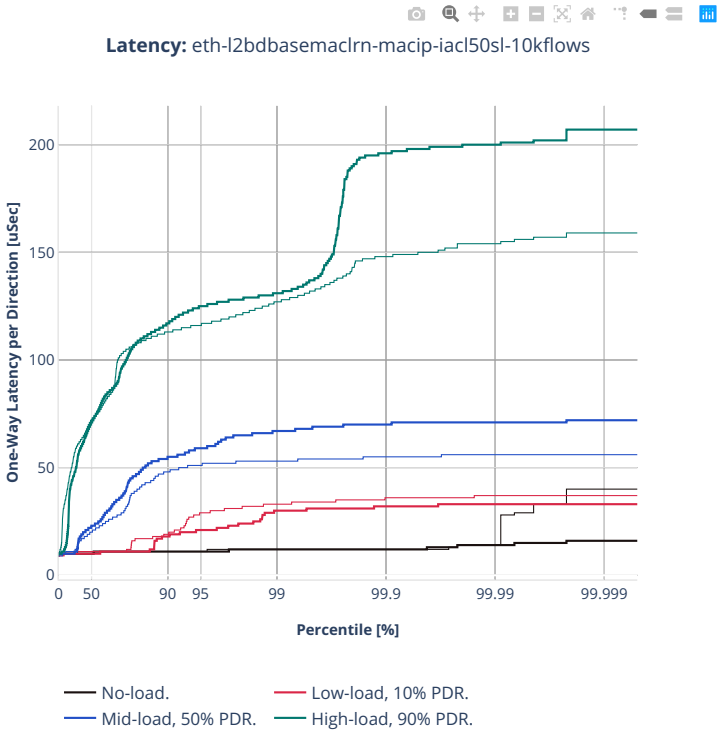
Latency: eth-l2bdbasemaclrn-macip-iacl50sl-10kflows
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2.5.2 IPv4 Routing

CSIT source code for the test cases used for plots can be found in CSIT git repository[150].

[150] https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls2206

2n-icx-xxv710

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

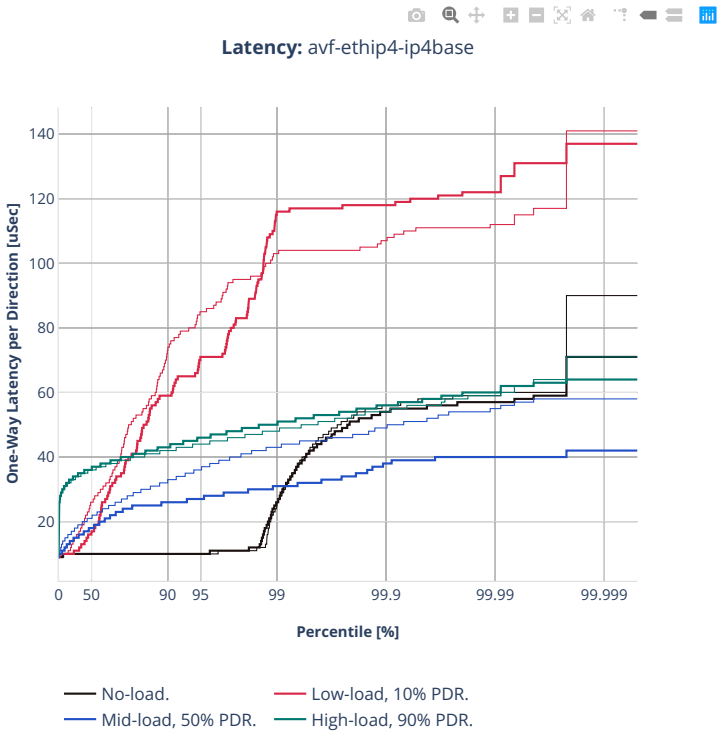

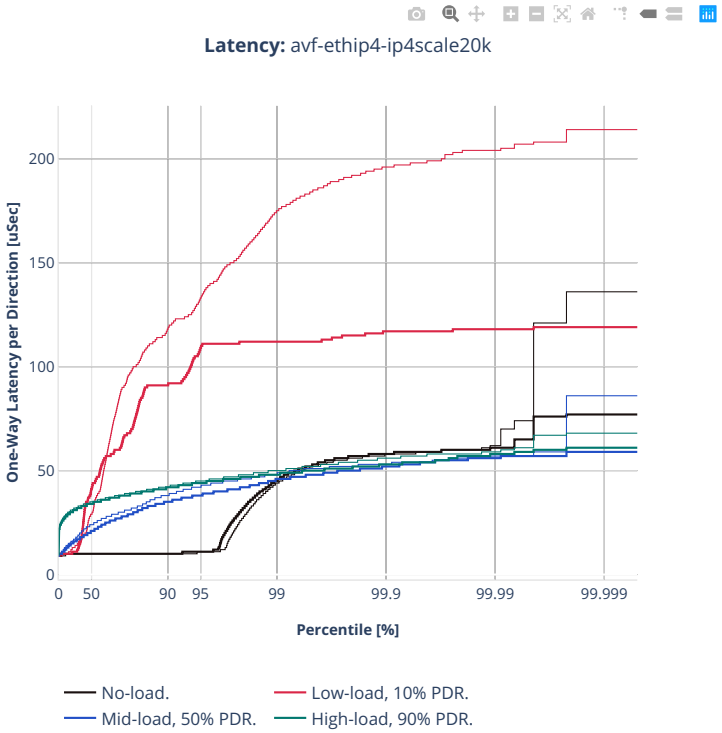
Latency: avf-ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

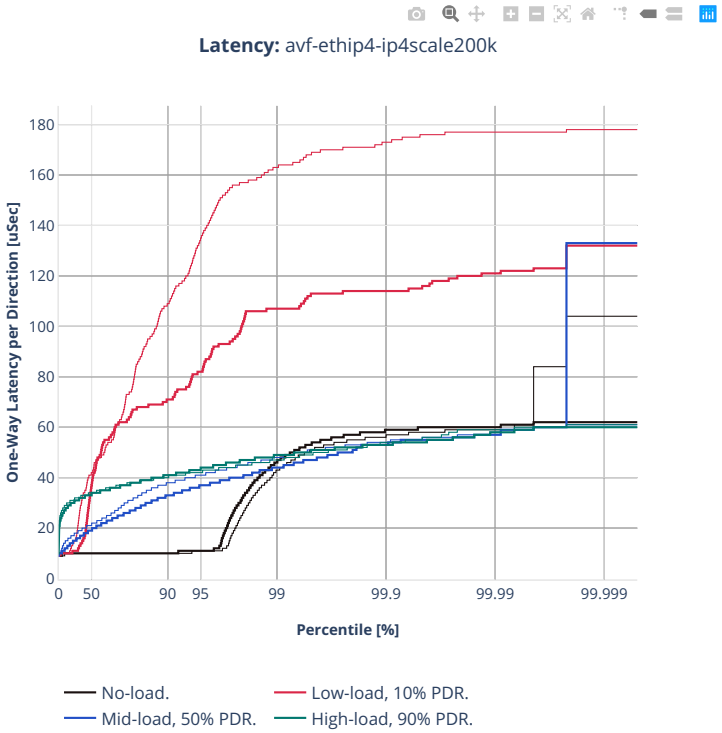
Latency: avf-ethip4-ip4scale200k
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

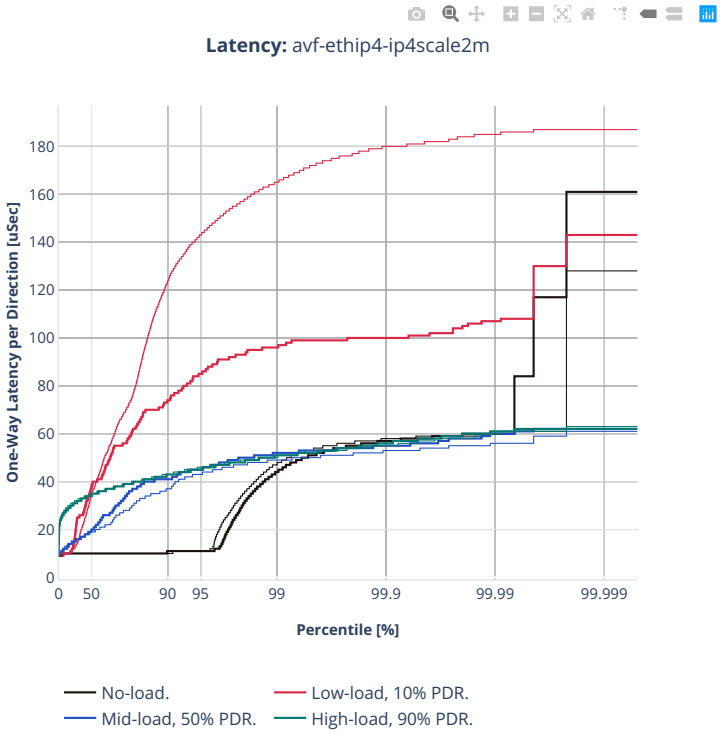
Latency: avf-ethip4-ip4scale2m
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

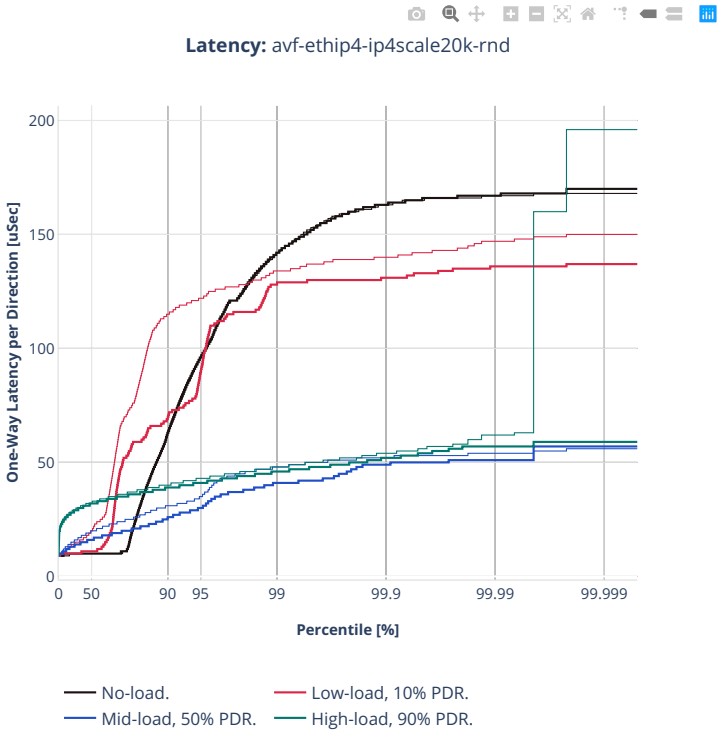
Latency: avf-ethip4-ip4scale20k-rnd
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

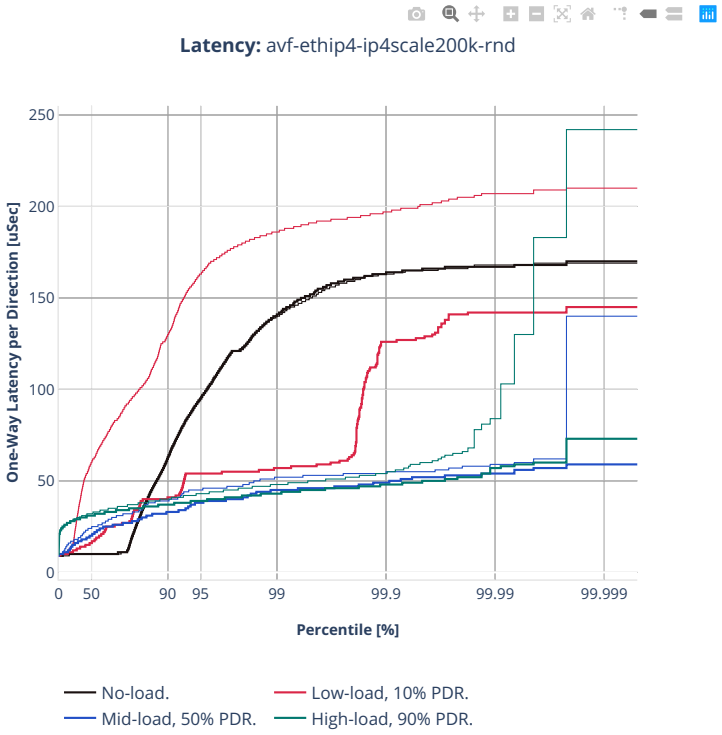
Latency: avf-ethip4-ip4scale200k-rnd
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

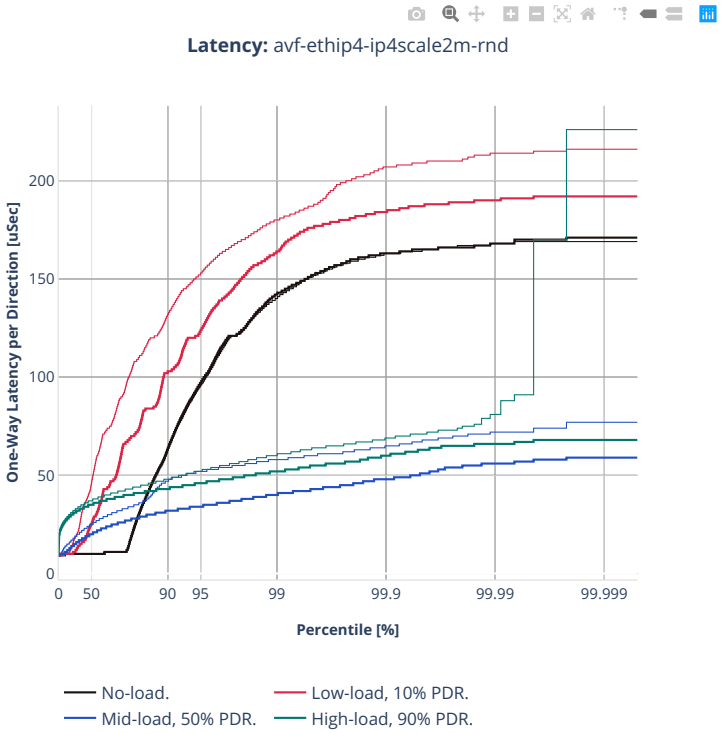
Latency: avf-ethip4-ip4scale2m-rnd
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

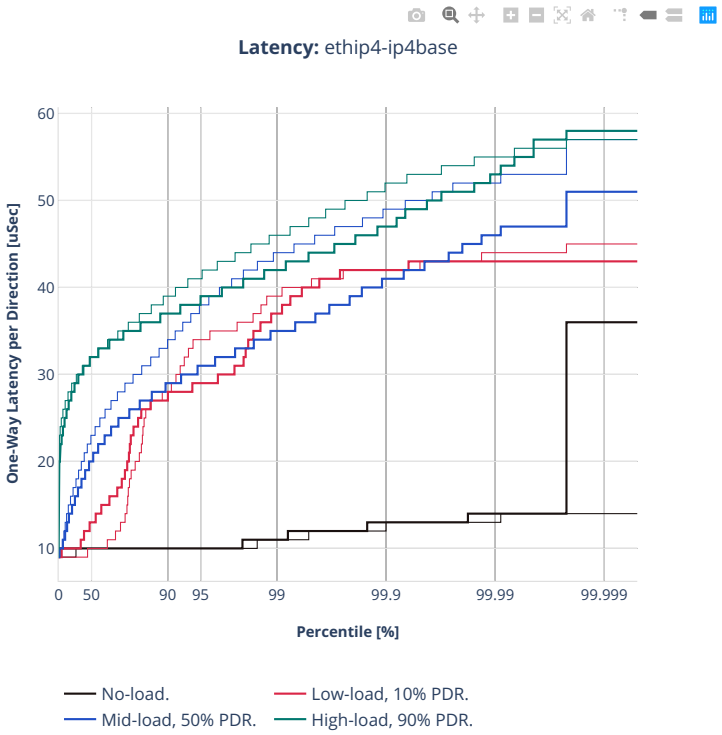

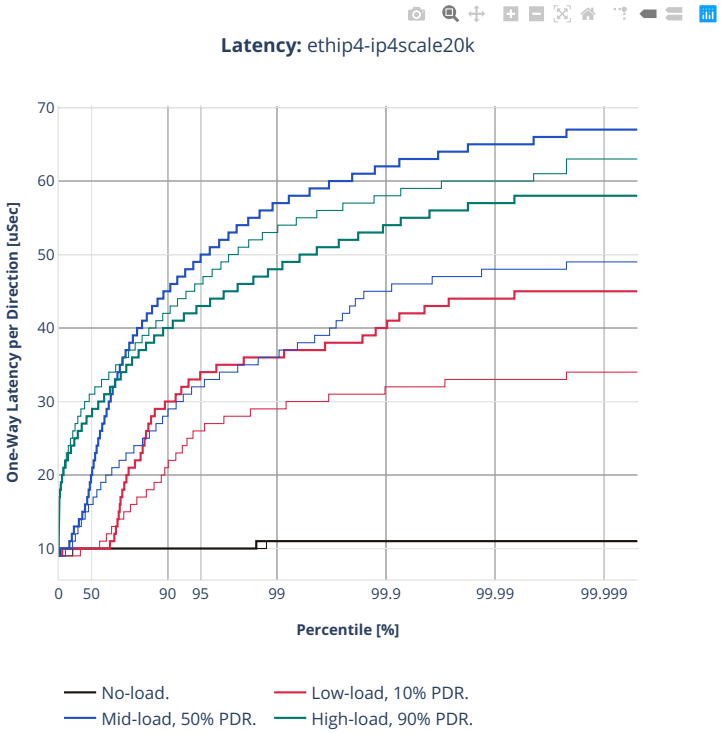
Latency: ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

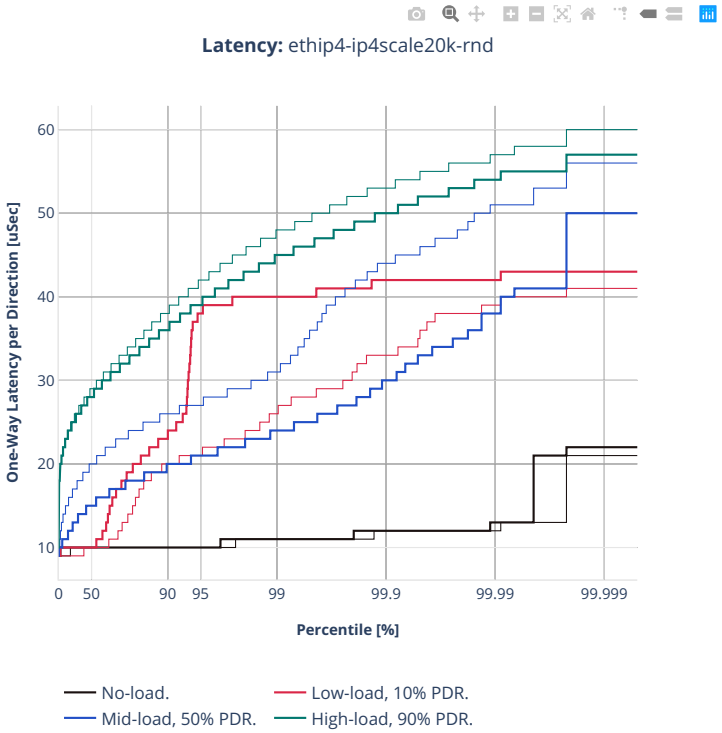
Latency: ethip4-ip4scale20k-rnd
One-Way Latency per Direction [uSec]
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-ip4routing-base-scale-af-xdp

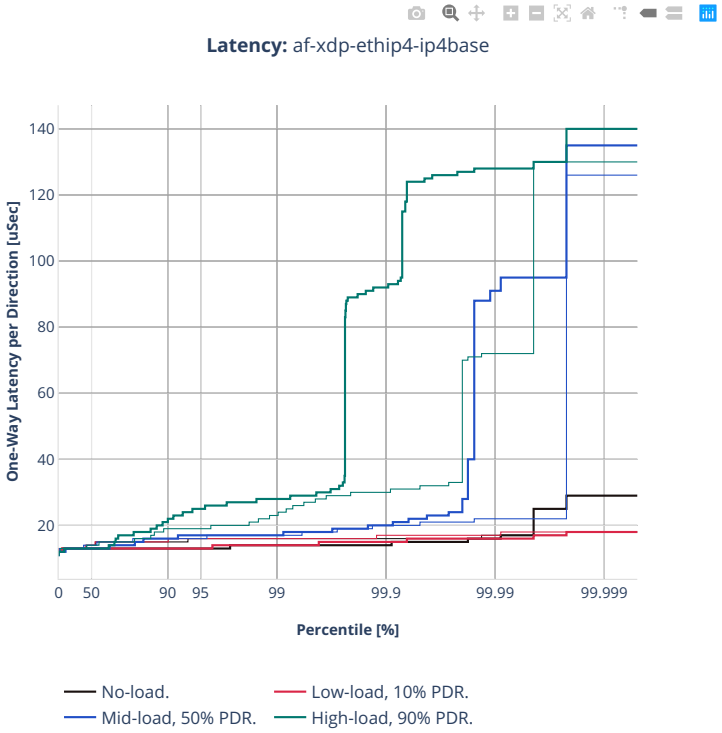

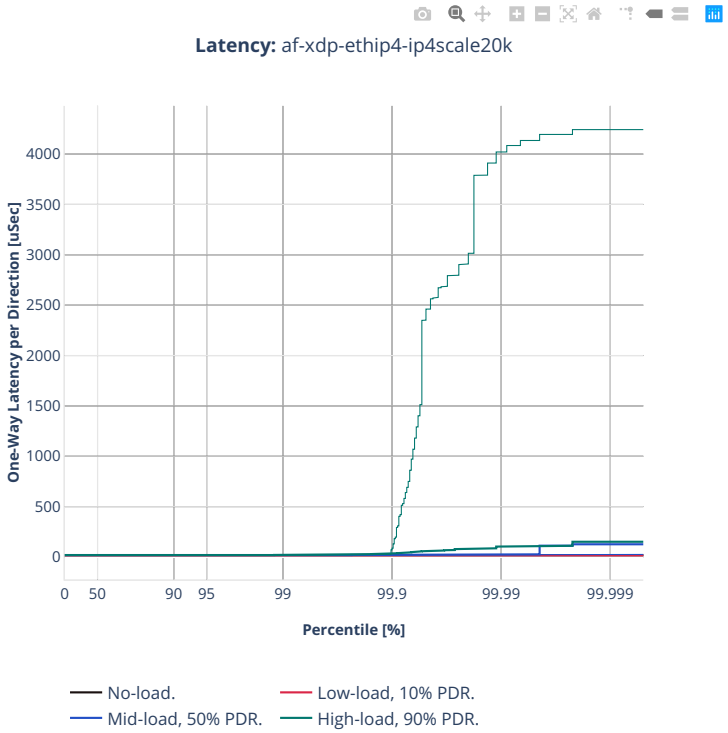
Latency: af-xdp-ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
4000
3500
3000
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

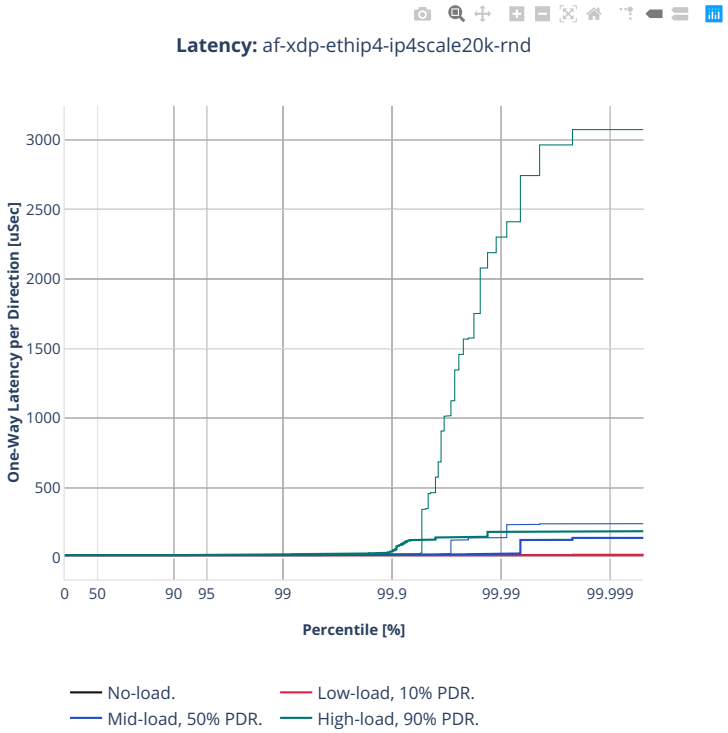
Latency: af-xdp-ethip4-ip4scale20k-rnd
One-Way Latency per Direction [uSec]
3000
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-ip4routing-features-avf

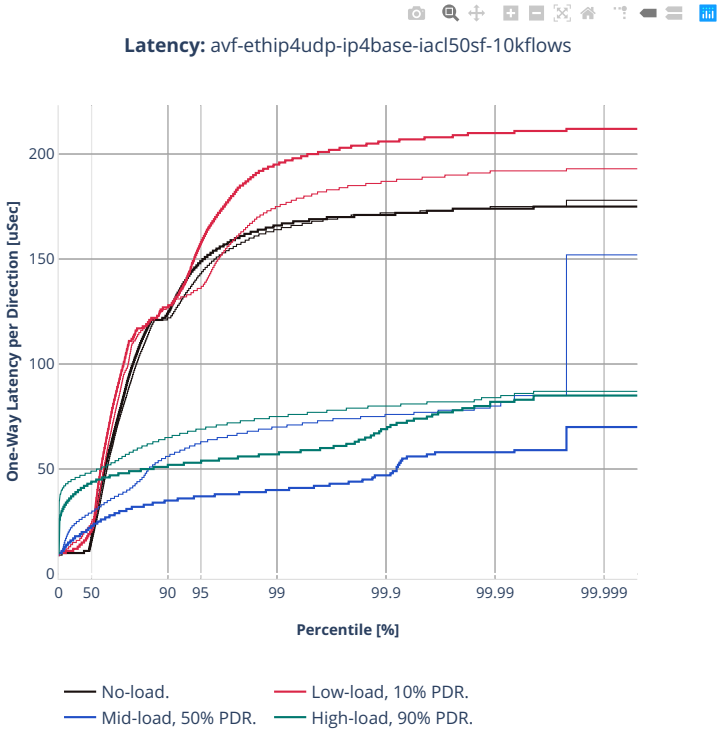

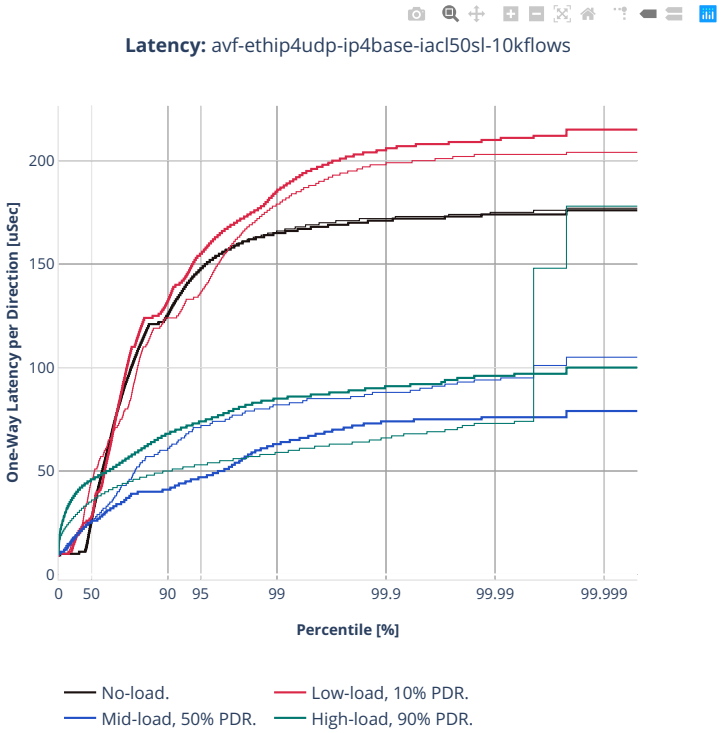
Latency: avf-ethip4udp-ip4base-iacl50sl-10kflows
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

Latency: avf-ethip4udp-ip4base-oacl50sf-10kflows

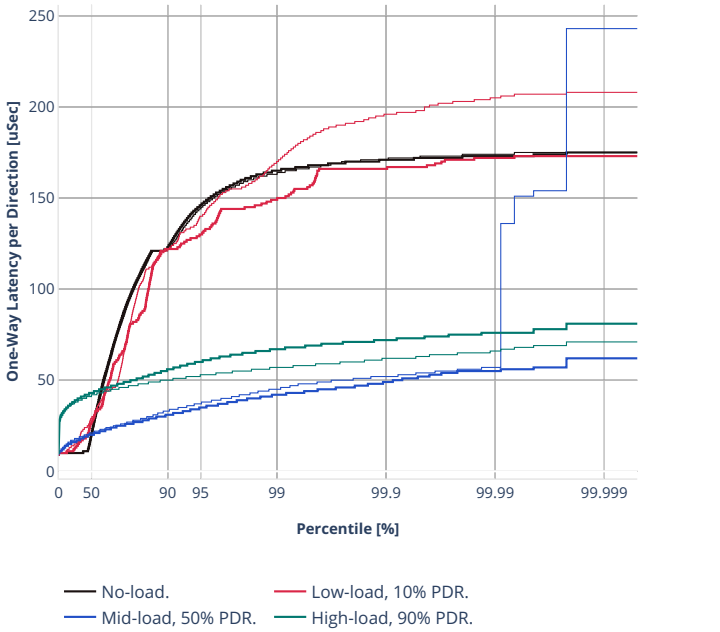

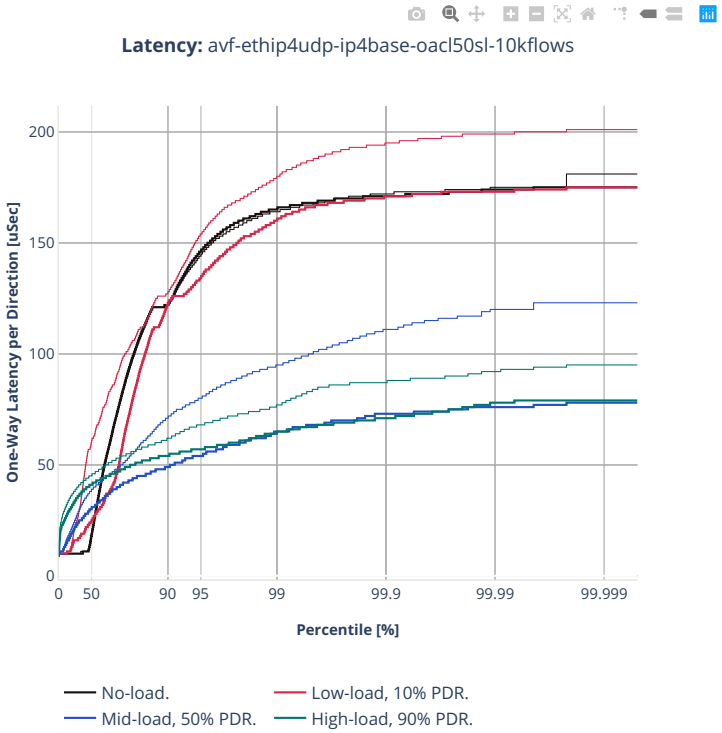
Latency: avf-ethip4udp-ip4base-oacl50sl-10kflows
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-icx-xxv710

64b-2t1c-ip4routing-base-avf

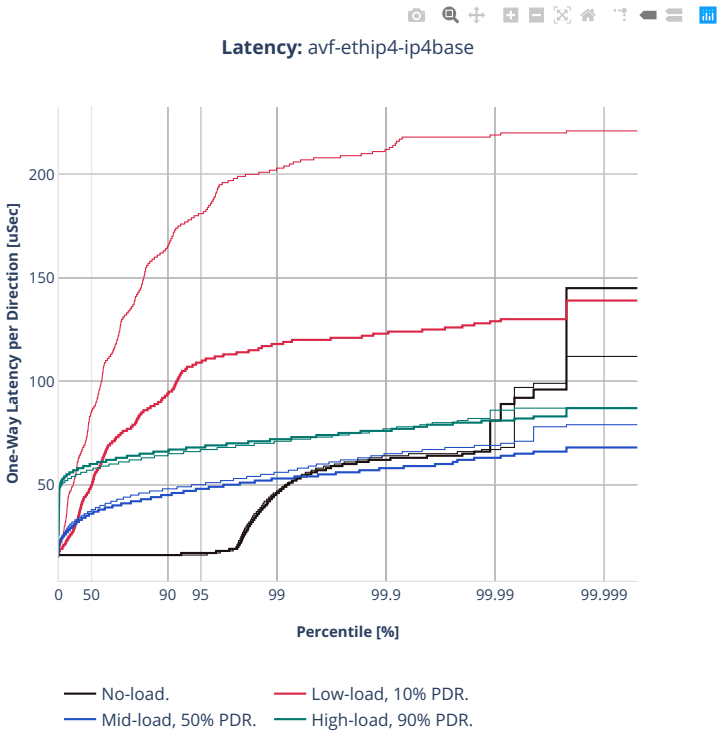

64b-2t1c-ip4routing-base-dpdk

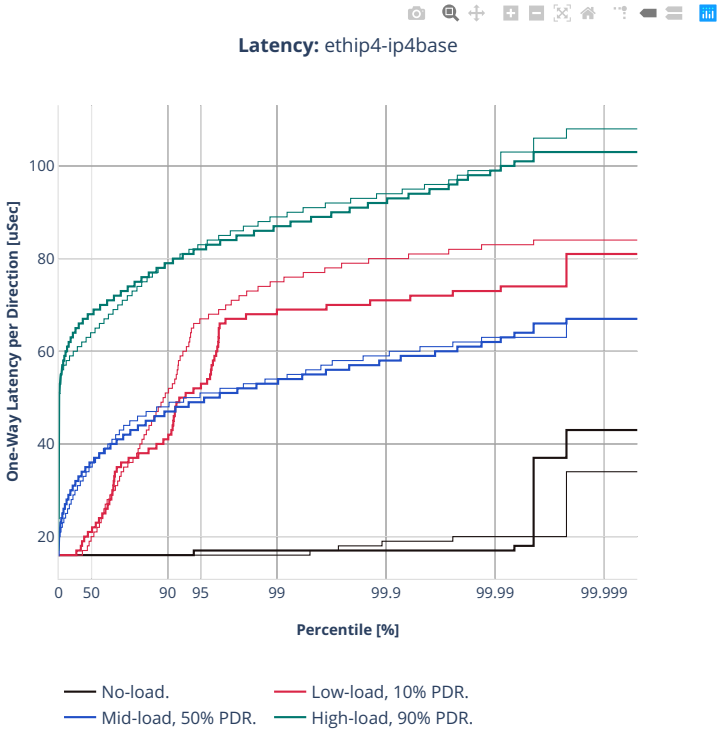

2n-skx-xxv710

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

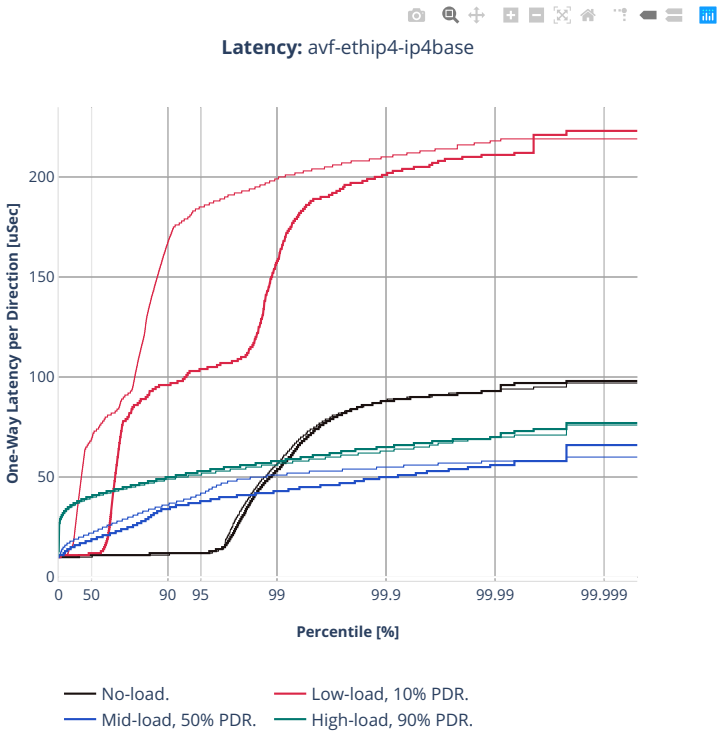

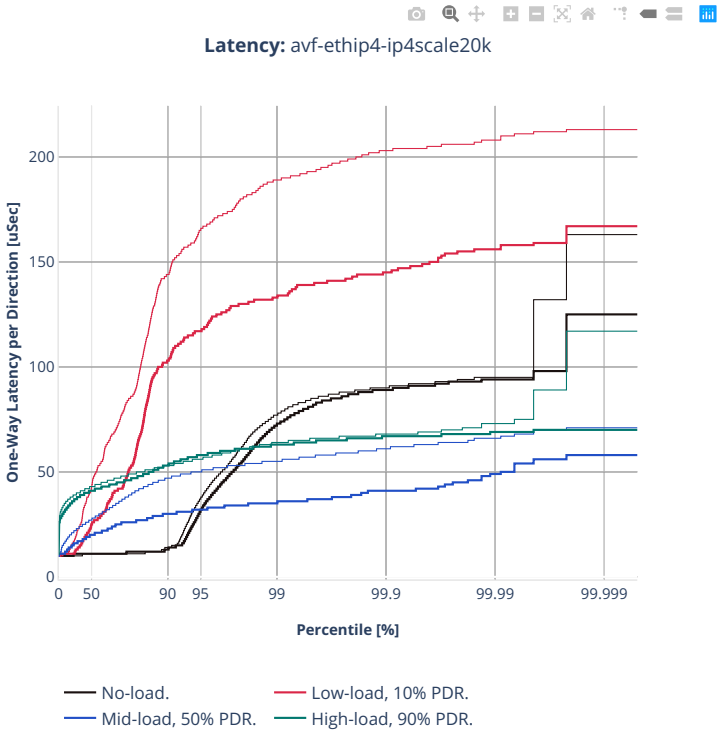
Latency: avf-ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

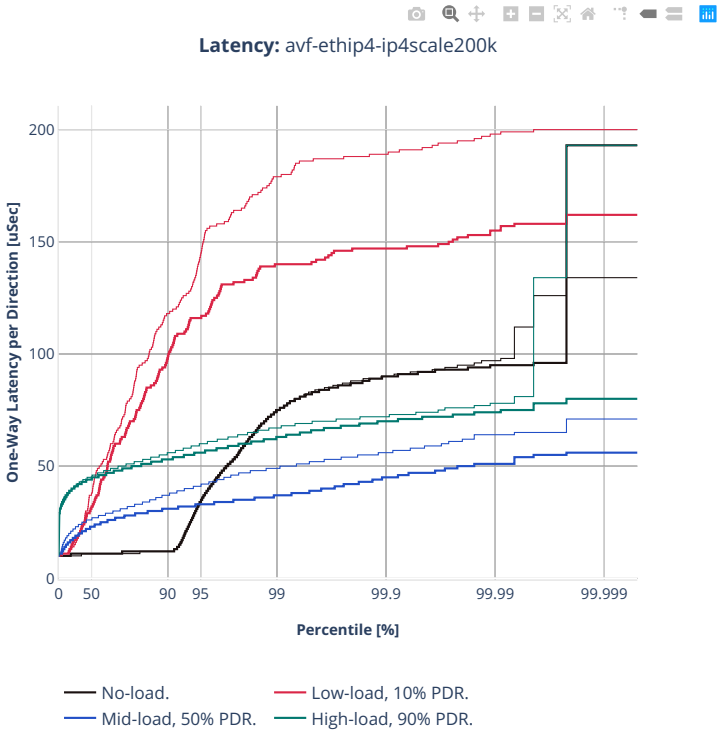
Latency: avf-ethip4-ip4scale200k
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

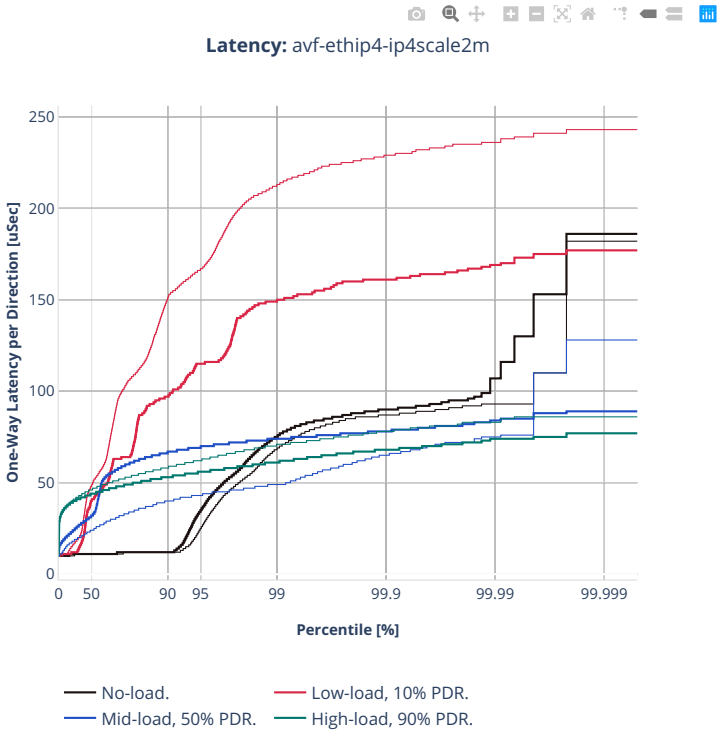
Latency: avf-ethip4-ip4scale2m
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

Latency: avf-ethip4-ip4scale20k-rnd

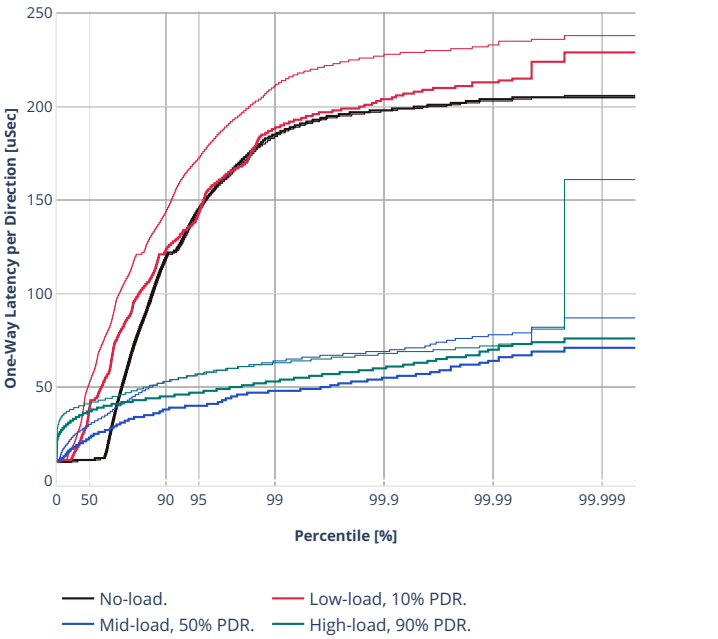

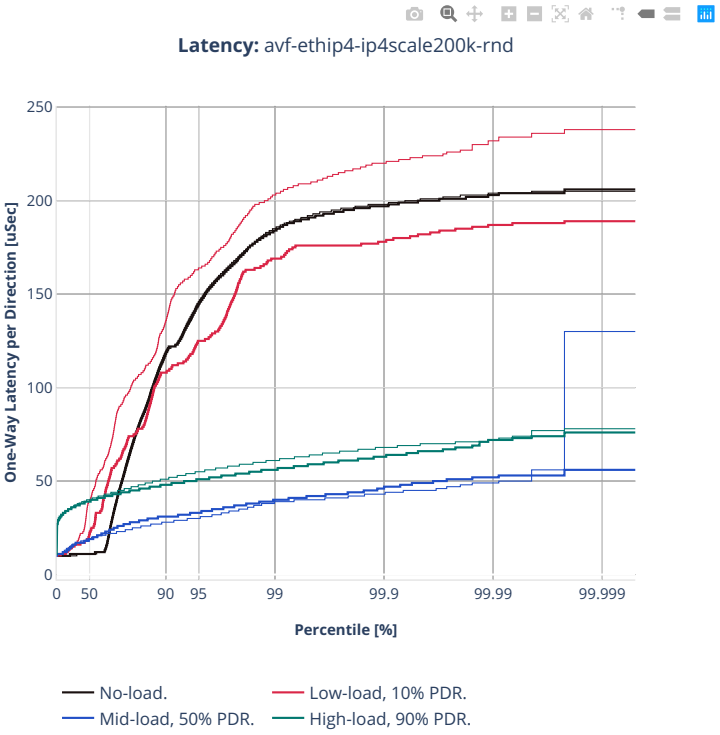
Latency: avf-ethip4-ip4scale200k-rnd
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

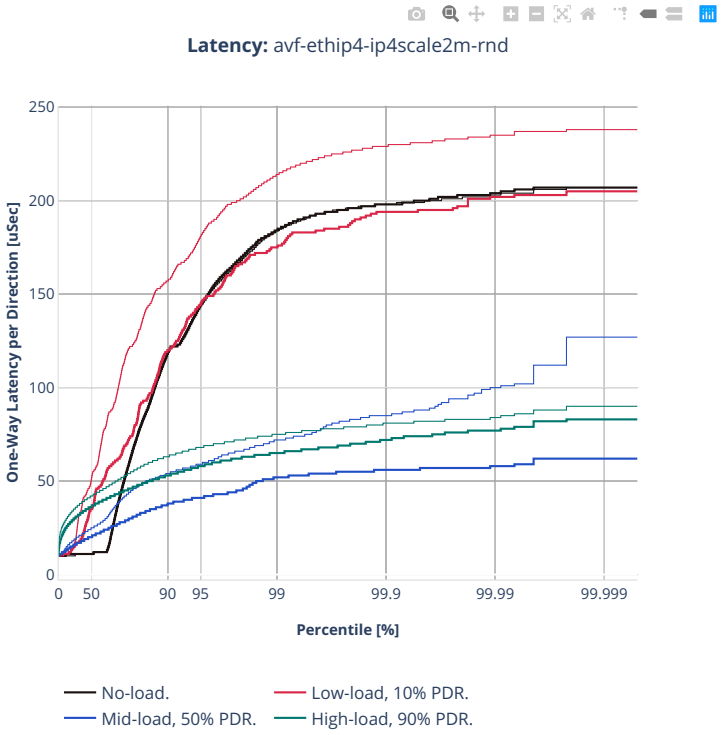
Latency: avf-ethip4-ip4scale2m-rnd
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

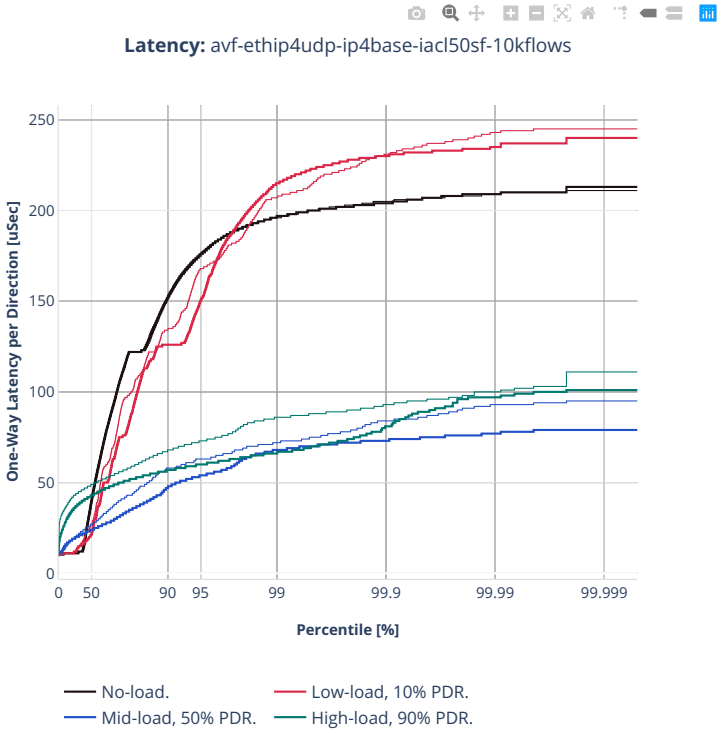

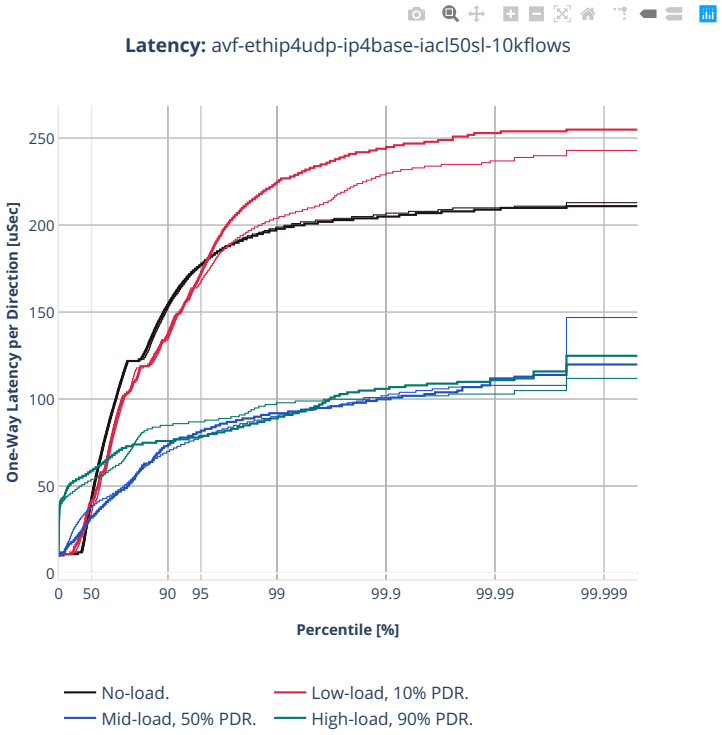
Latency: avf-ethip4udp-ip4base-iacl50sl-10kflows
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

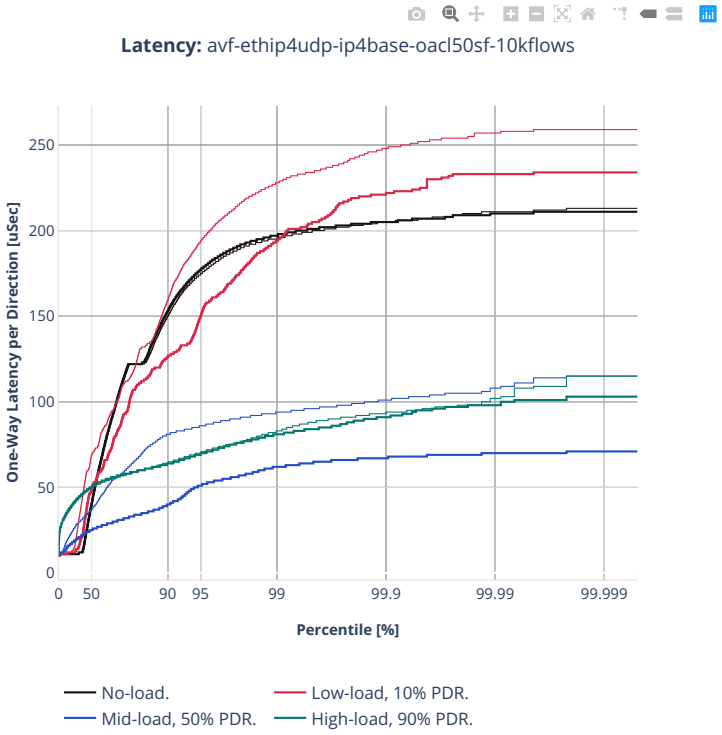
Latency: avf-ethip4udp-ip4base-oacl50sf-10kflows
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

Latency: avf-ethip4udp-ip4base-oacl50sl-10kflows

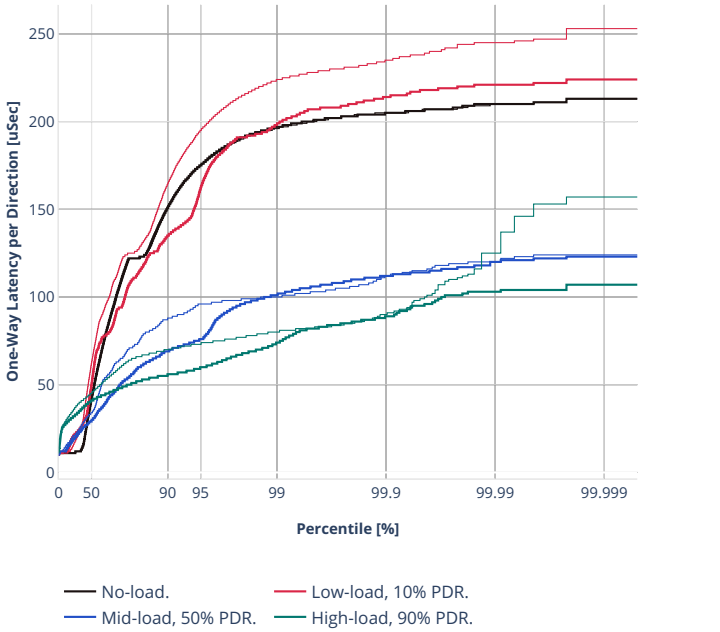

64b-2t1c-ip4routing-base-scale-af-xdp

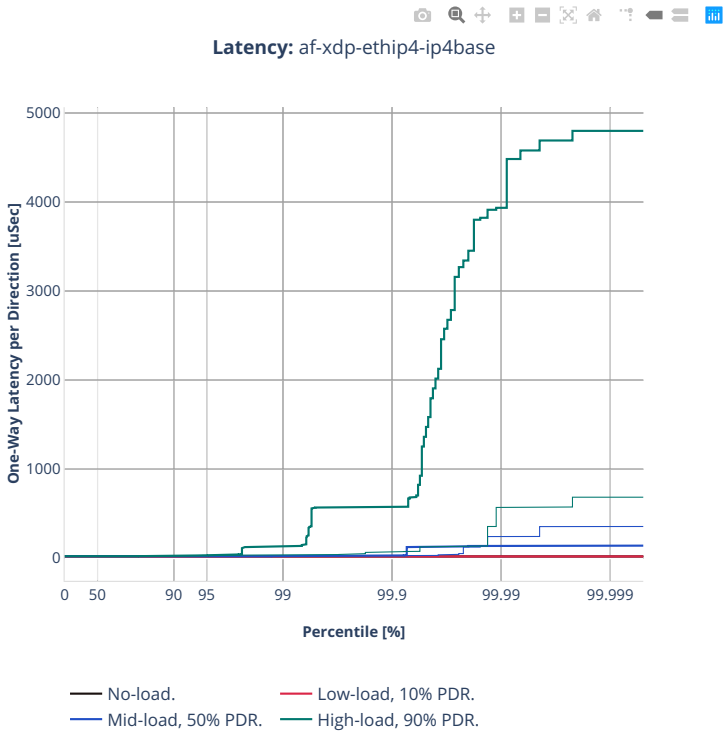

Latency: af-xdp-ethip4-ip4scale20k

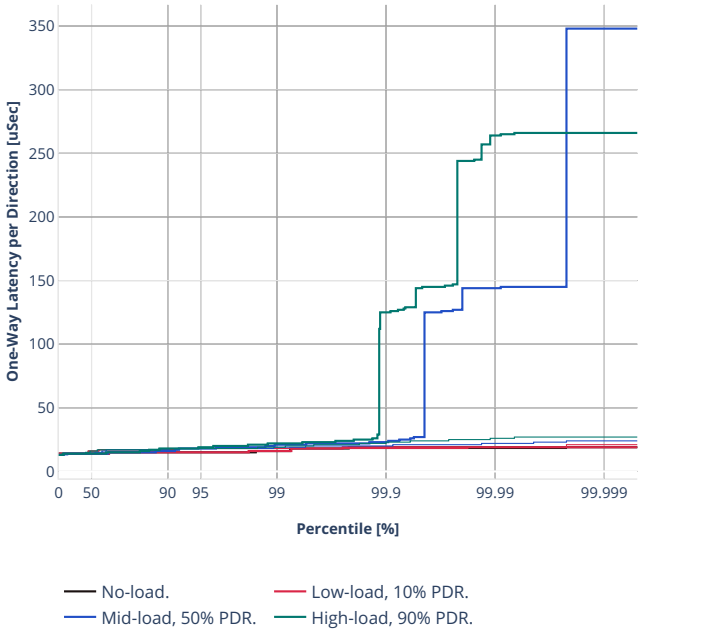

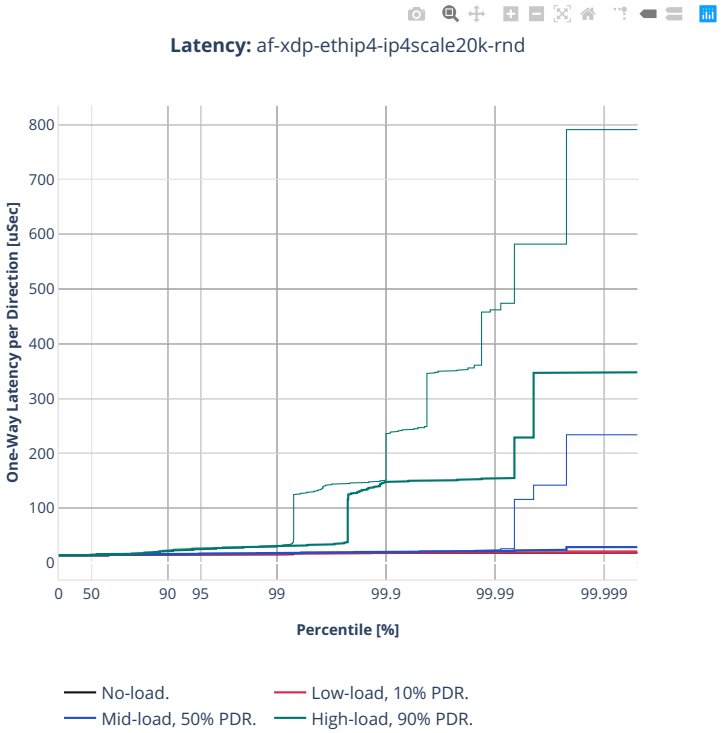
Latency: af-xdp-ethip4-ip4scale20k-rnd
One-Way Latency per Direction [uSec]
800
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

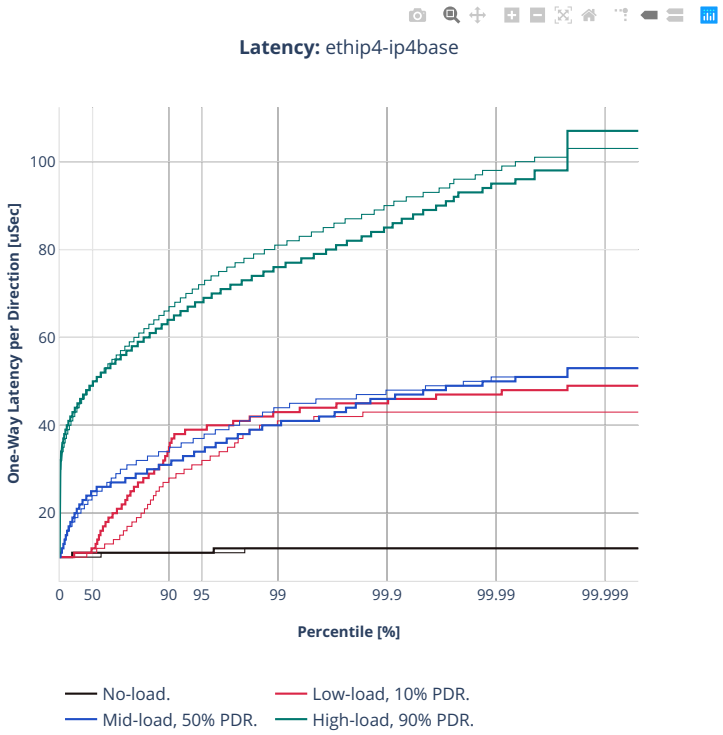

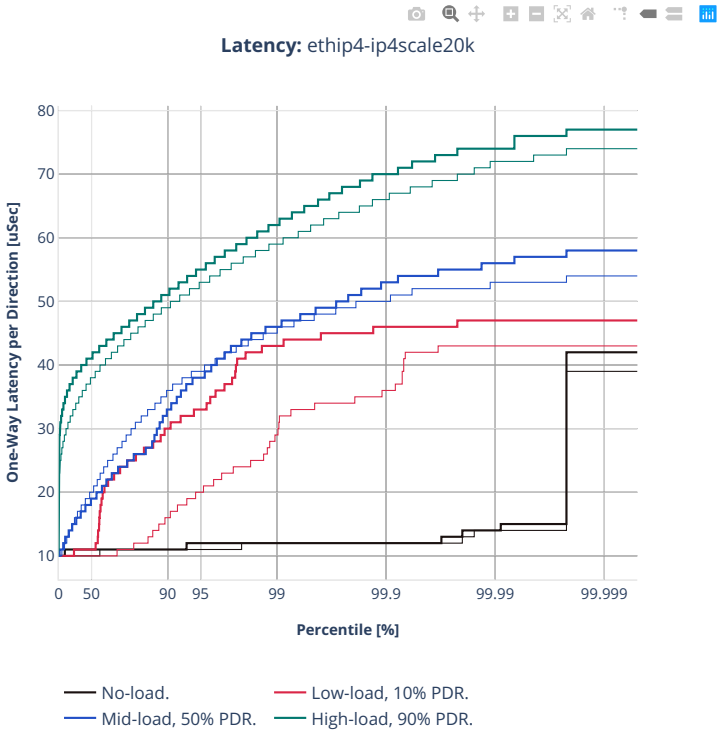
Latency: ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

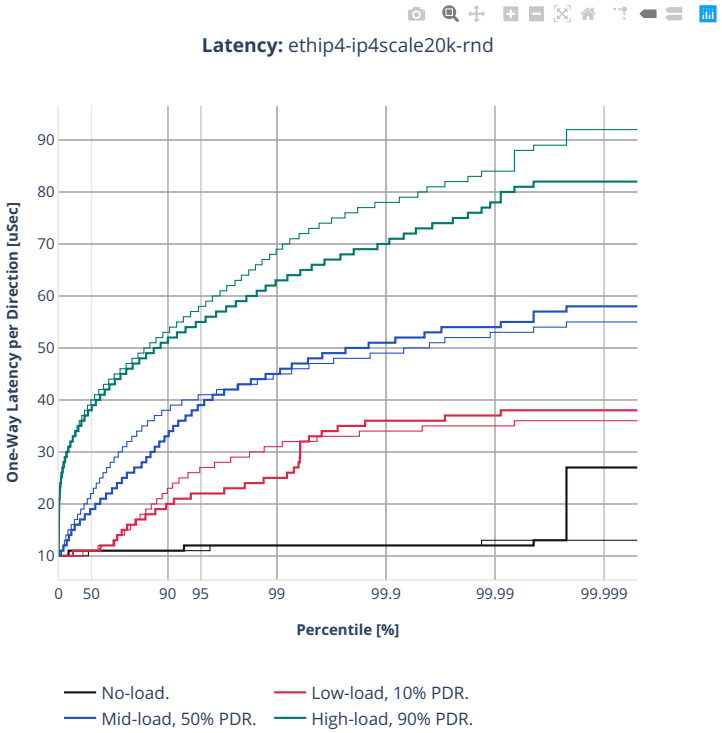
Latency: ethip4-ip4scale20k-rnd
One-Way Latency per Direction [uSec]
90
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-skx-xxv710

64b-2t1c-ip4routing-base-avf

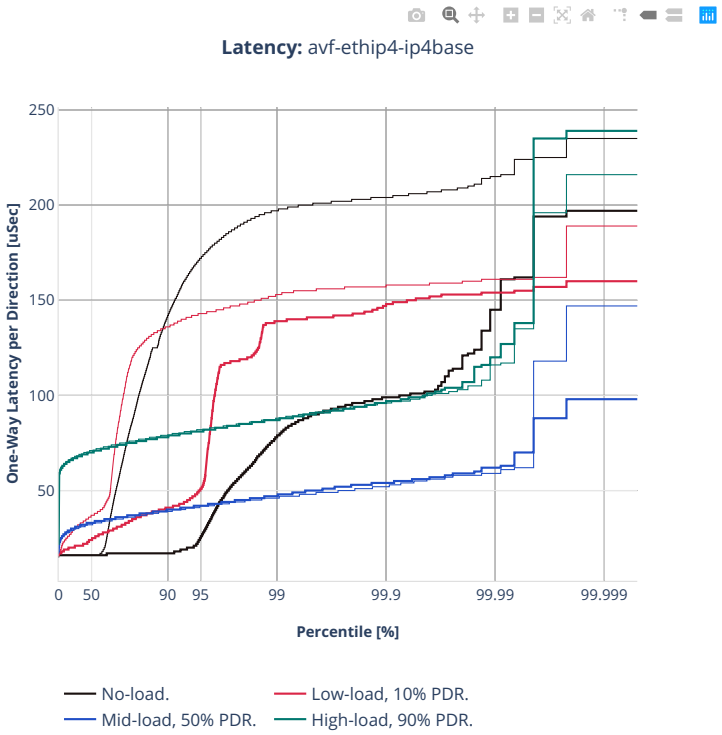

64b-2t1c-ip4routing-base-dpdk

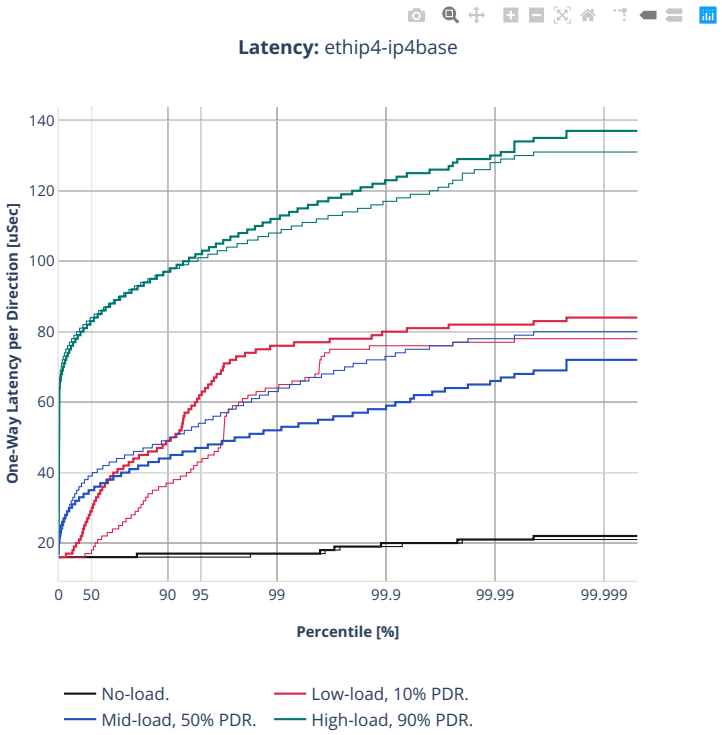

2n-clx-xxv710

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

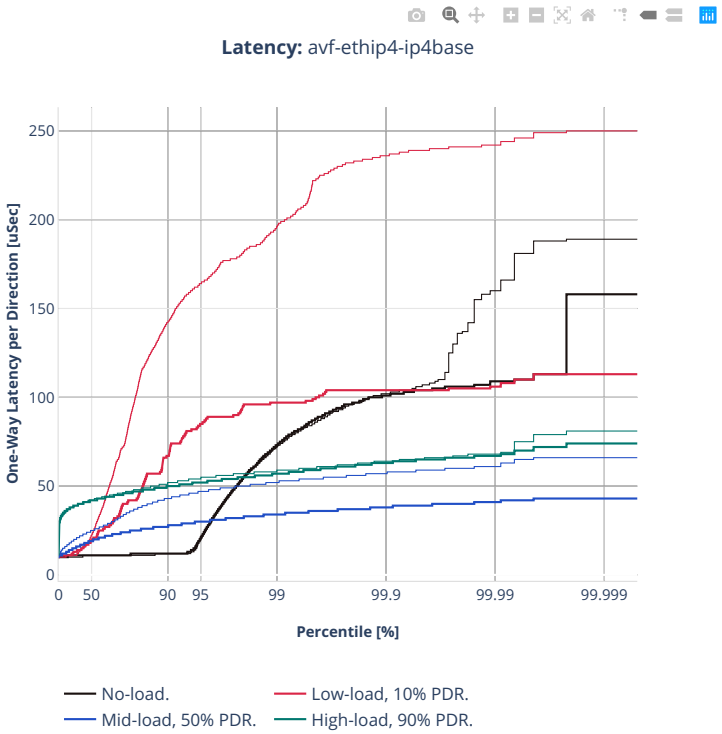

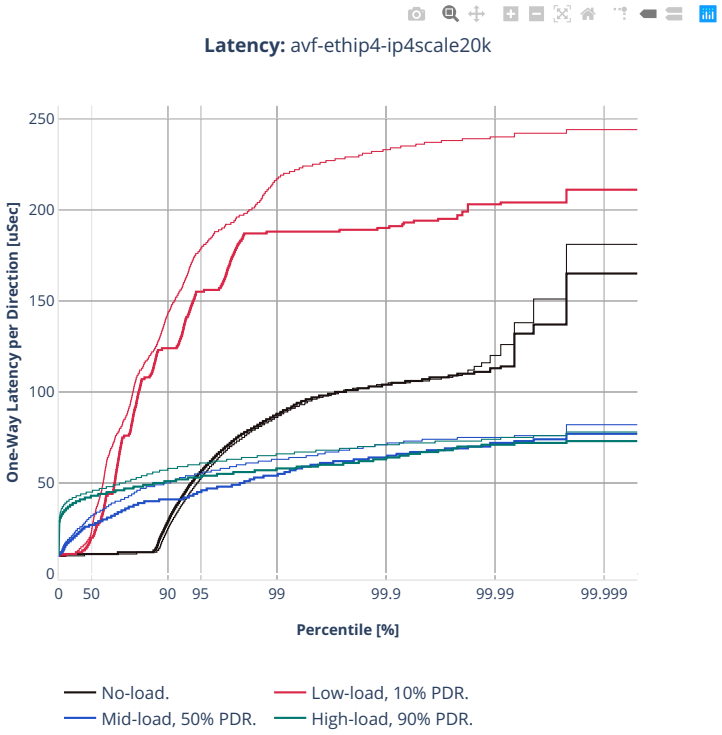
Latency: avf-ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

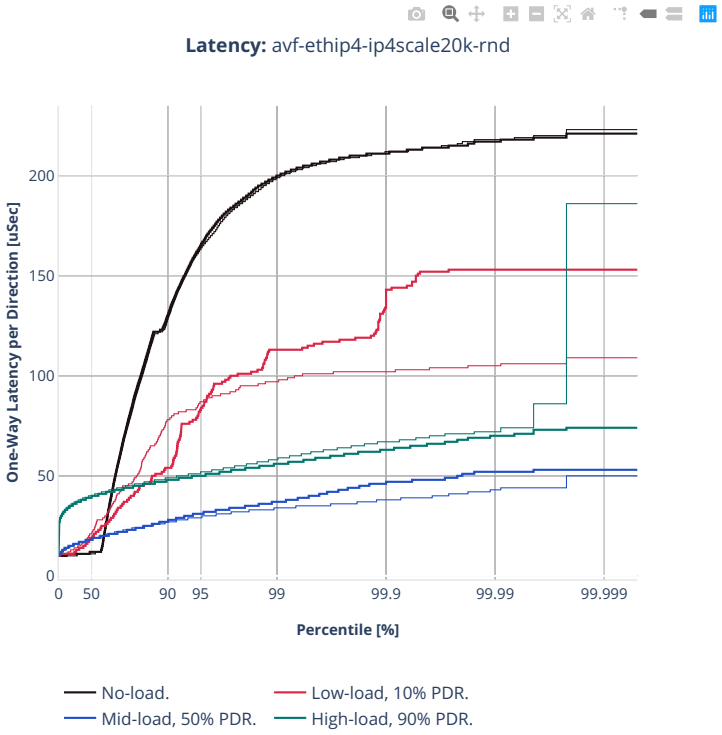
Latency: avf-ethip4-ip4scale20k-rnd
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

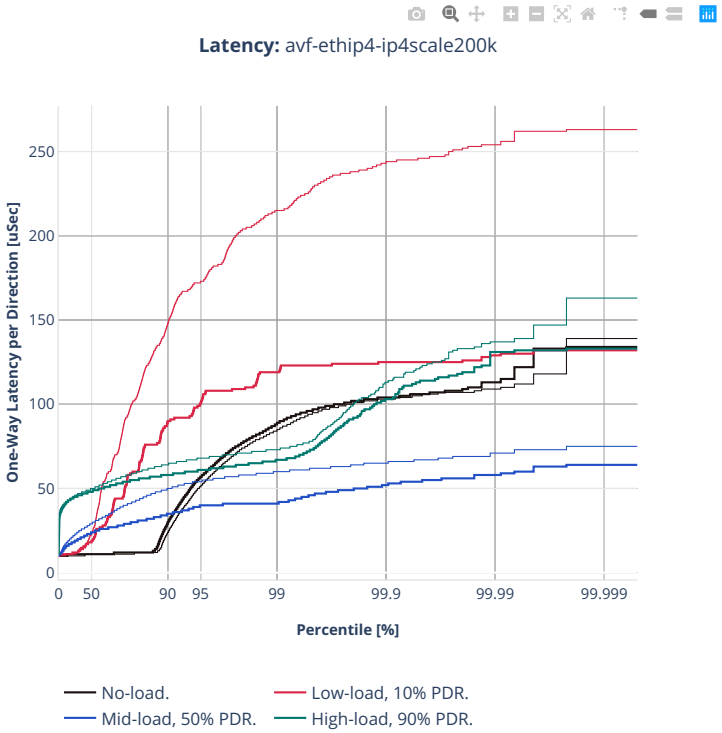
Latency: avf-ethip4-ip4scale200k
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

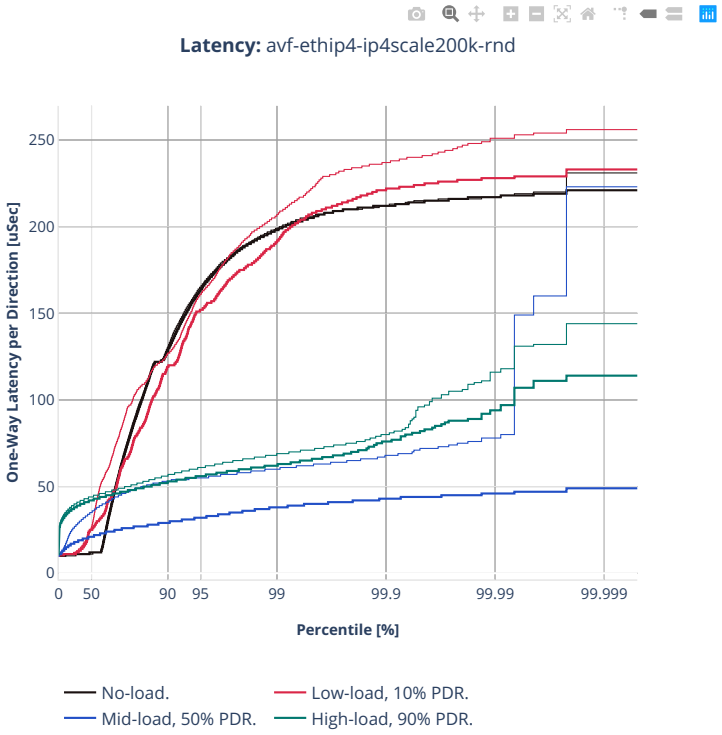
Latency: avf-ethip4-ip4scale200k-rnd
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

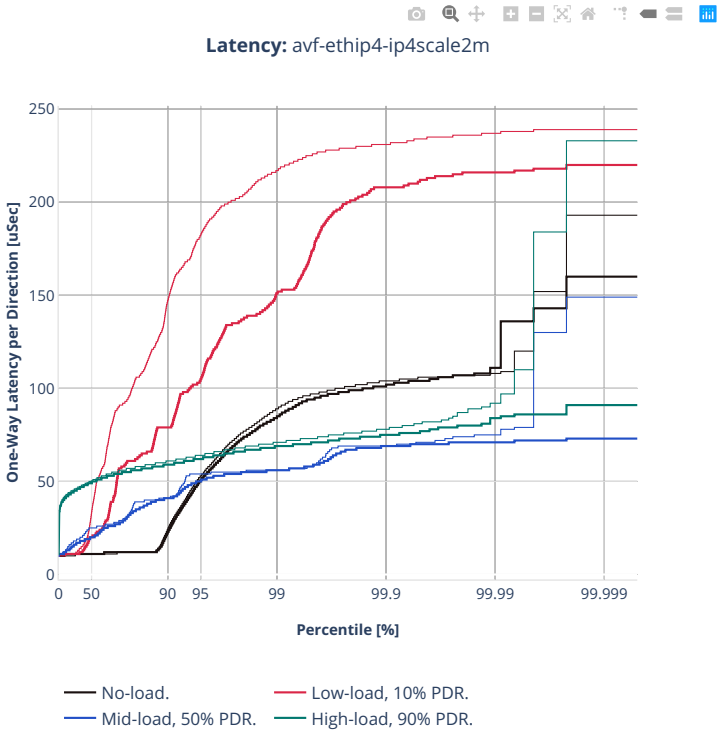
Latency: avf-ethip4-ip4scale2m
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

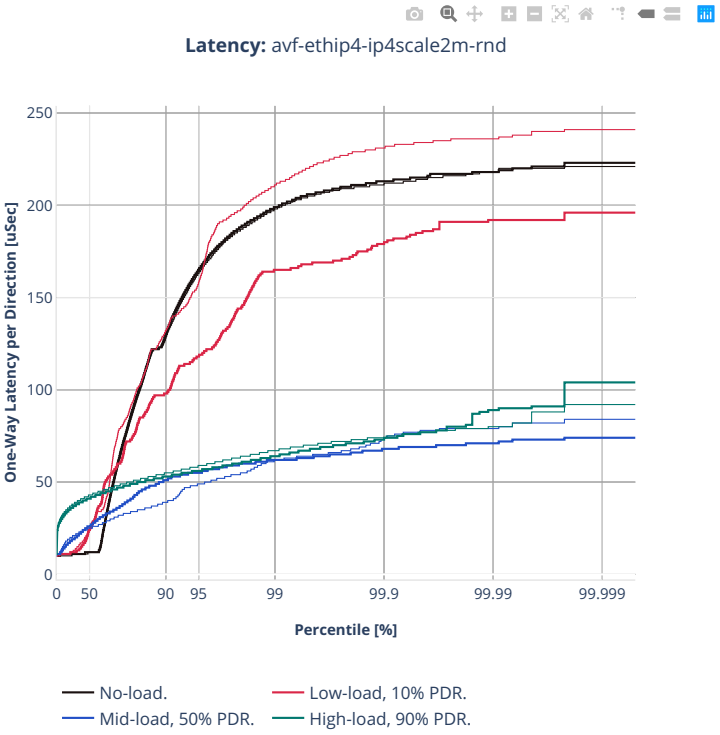
Latency: avf-ethip4-ip4scale2m-rnd
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-ip4routing-base-scale-af-xdp

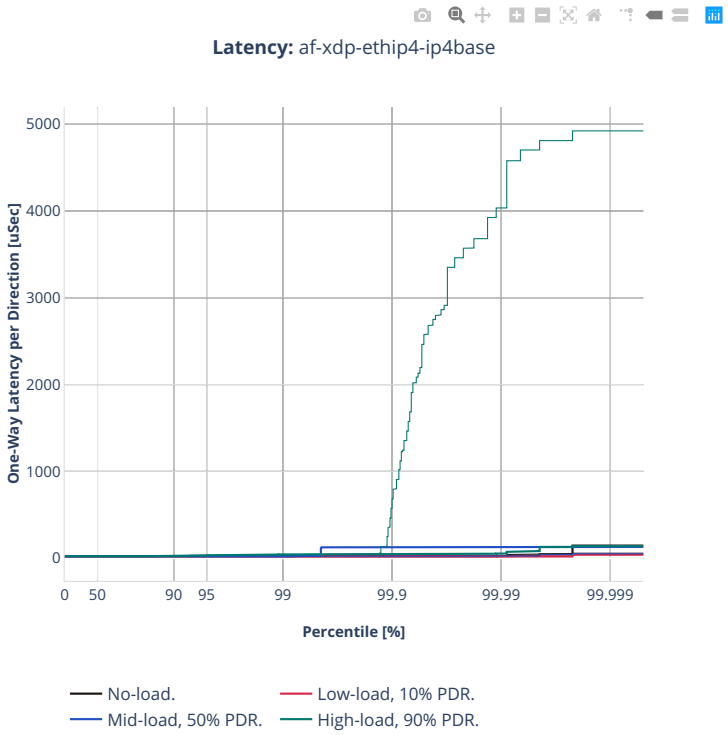

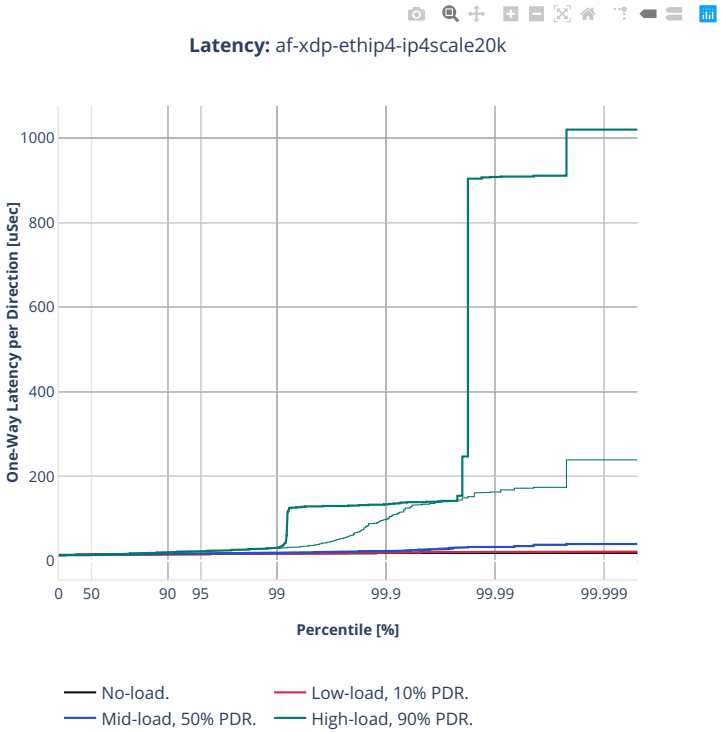
Latency: af-xdp-ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
1000
800
600
400
200
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

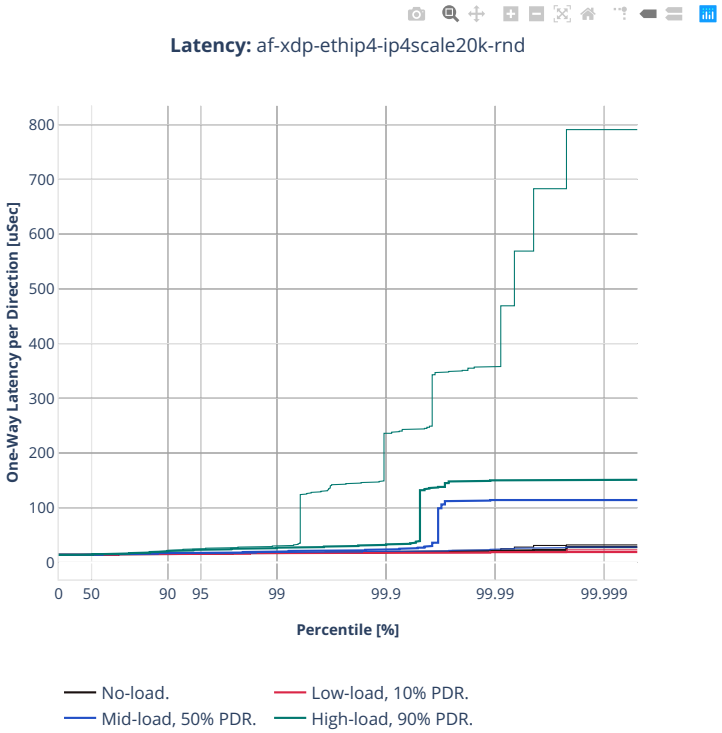
Latency: af-xdp-ethip4-ip4scale20k-rnd
One-Way Latency per Direction [uSec]
800
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-ip4routing-features-avf

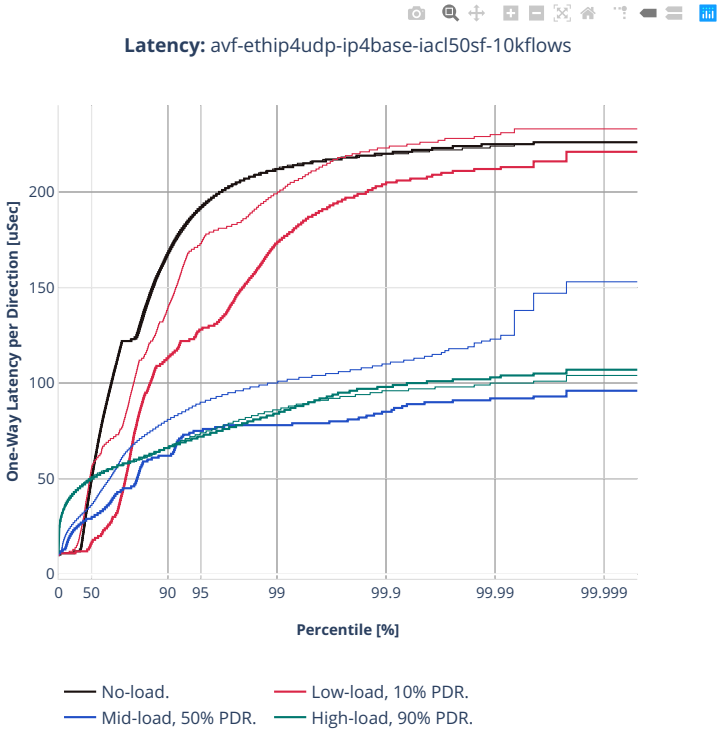

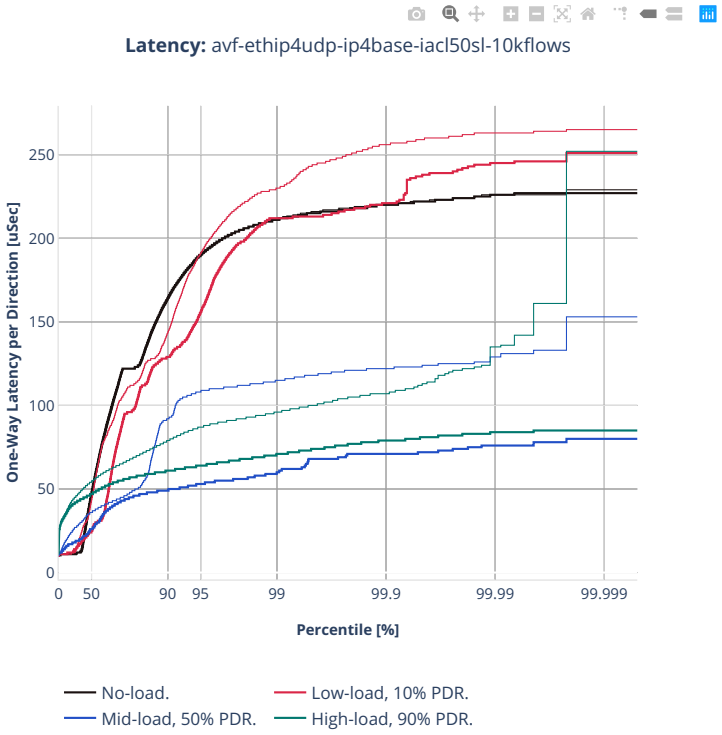
Latency: avf-ethip4udp-ip4base-iacl50sl-10kflows
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

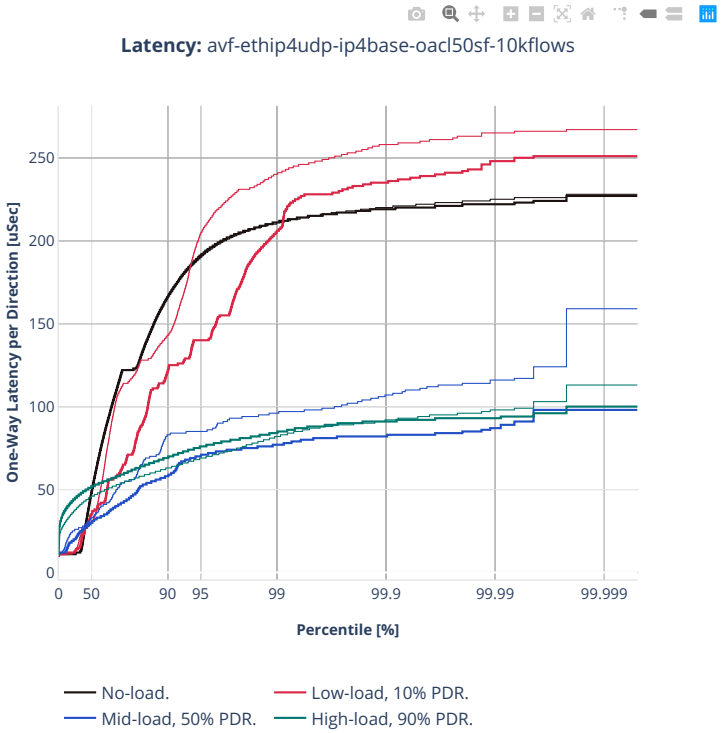
Latency: avf-ethip4udp-ip4base-oacl50sf-10kflows
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

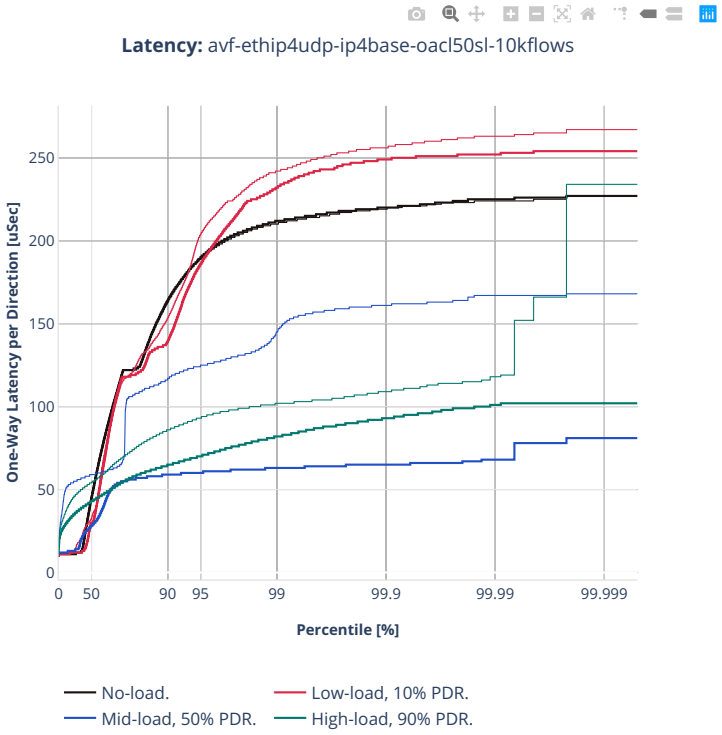
Latency: avf-ethip4udp-ip4base-oacl50sl-10kflows
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

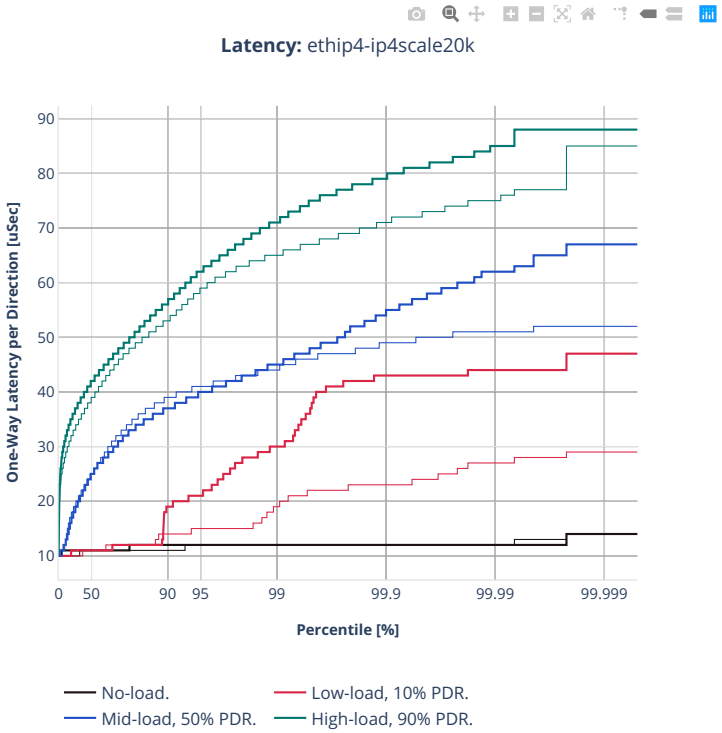
Latency: ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
90
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

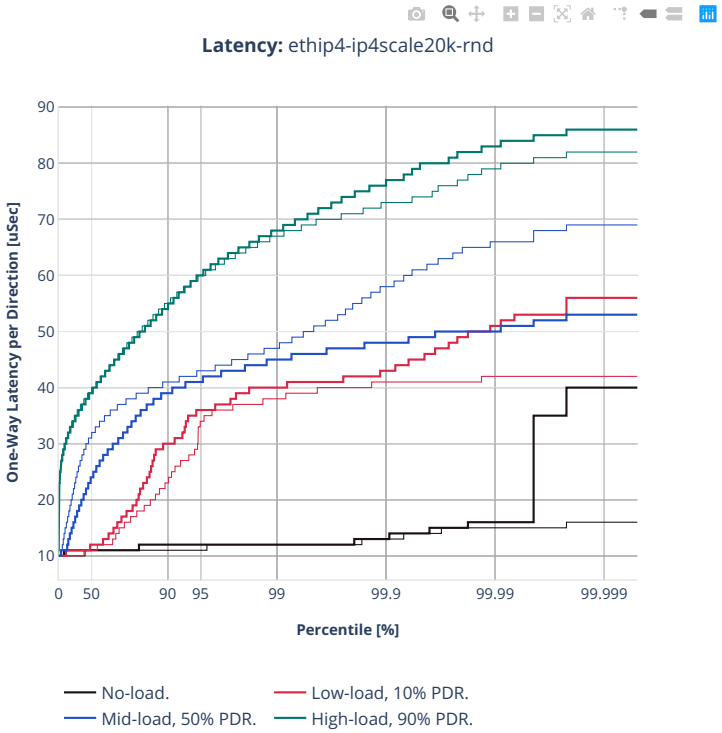
Latency: ethip4-ip4scale20k-rnd
One-Way Latency per Direction [uSec]
90
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-clx-cx556a

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

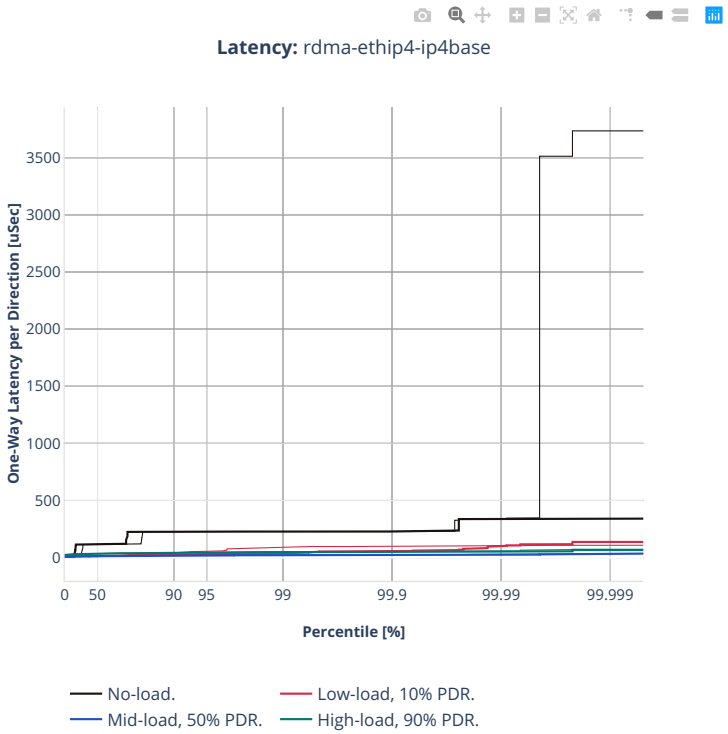

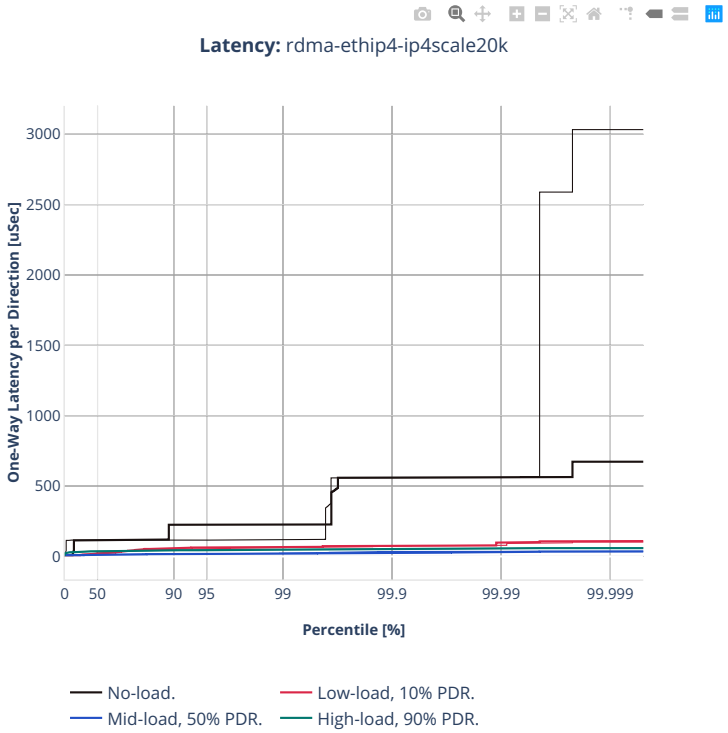
Latency: rdma-ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
3000
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

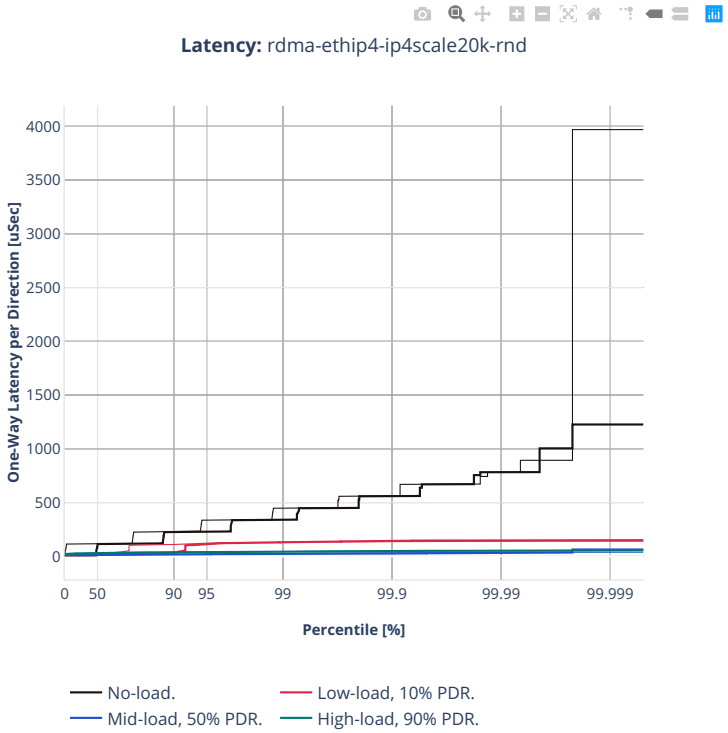
Latency: rdma-ethip4-ip4scale20k-rnd
One-Way Latency per Direction [uSec]
4000
3500
3000
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-ip4routing-features-rdma

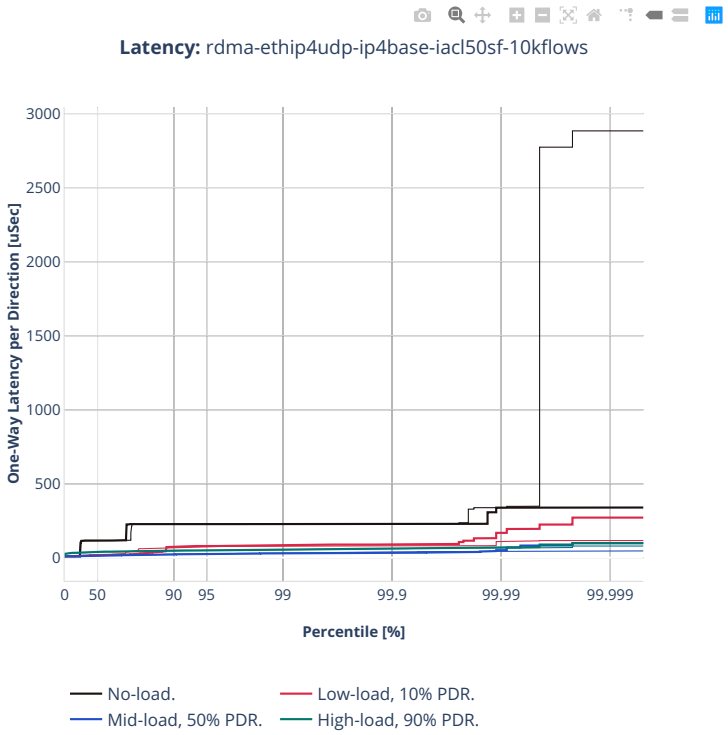

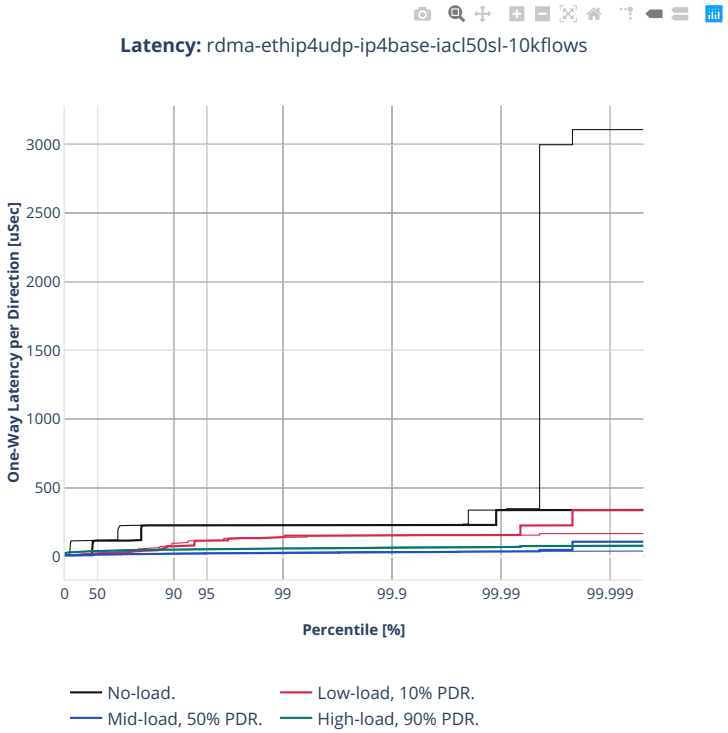
Latency: rdma-ethip4udp-ip4base-iacl50sl-10kflows
One-Way Latency per Direction [uSec]
3000
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

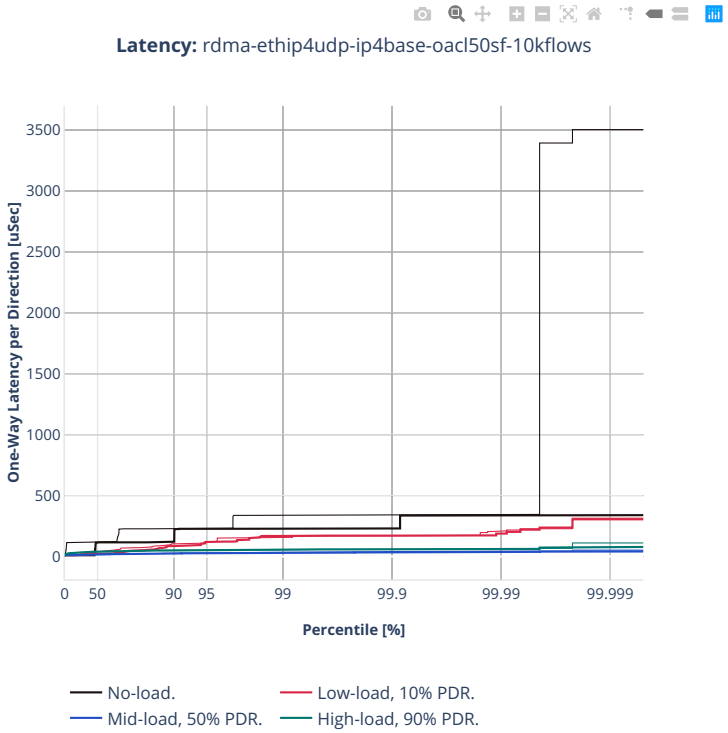
Latency: rdma-ethip4udp-ip4base-oacl50sf-10kflows
One-Way Latency per Direction [uSec]
3500
3000
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

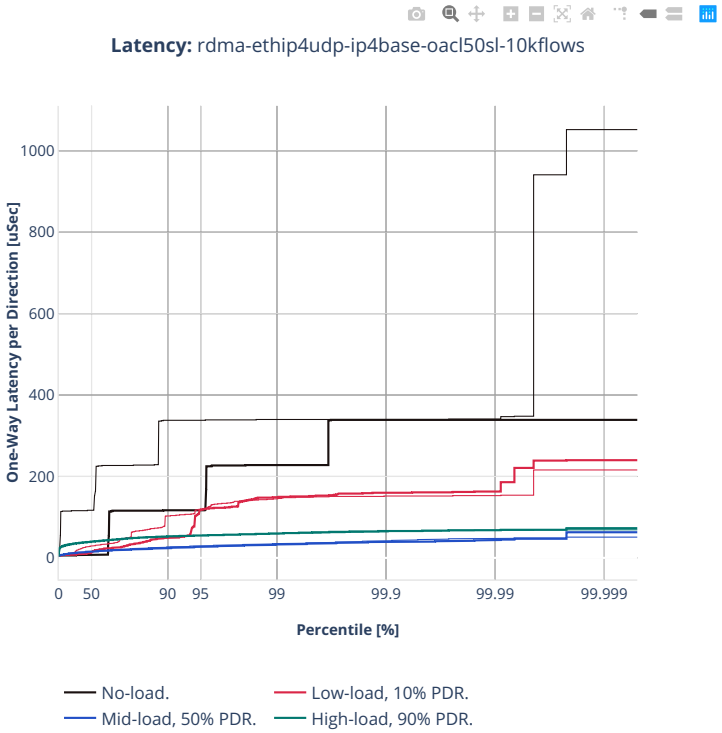
Latency: rdma-ethip4udp-ip4base-oacl50sl-10kflows
One-Way Latency per Direction [uSec]
1000
800
600
400
200
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-clx-e810cq

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

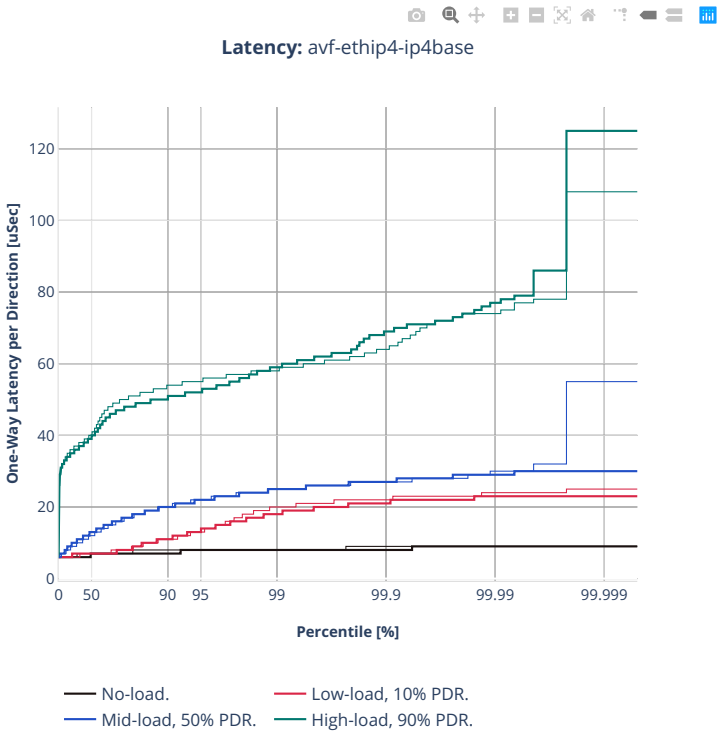

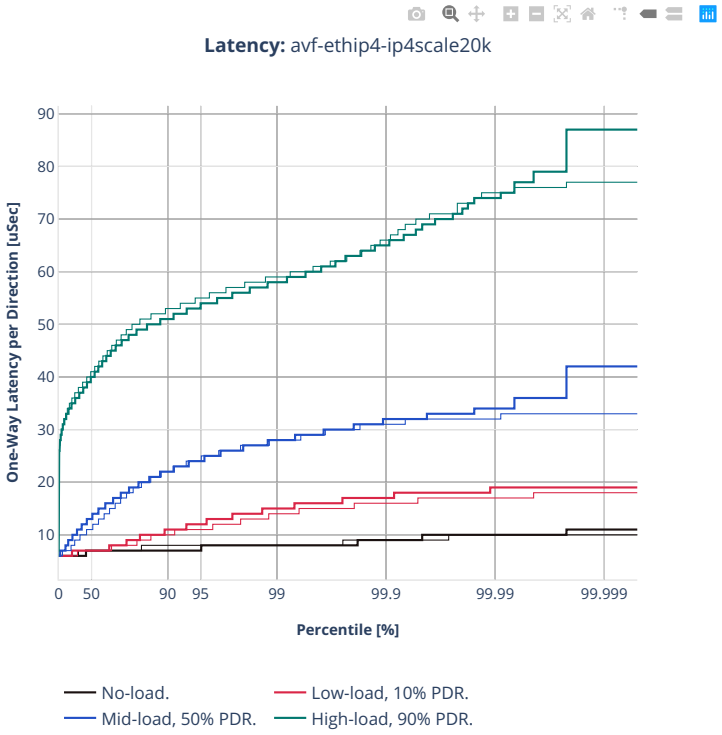
Latency: avf-ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
90
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

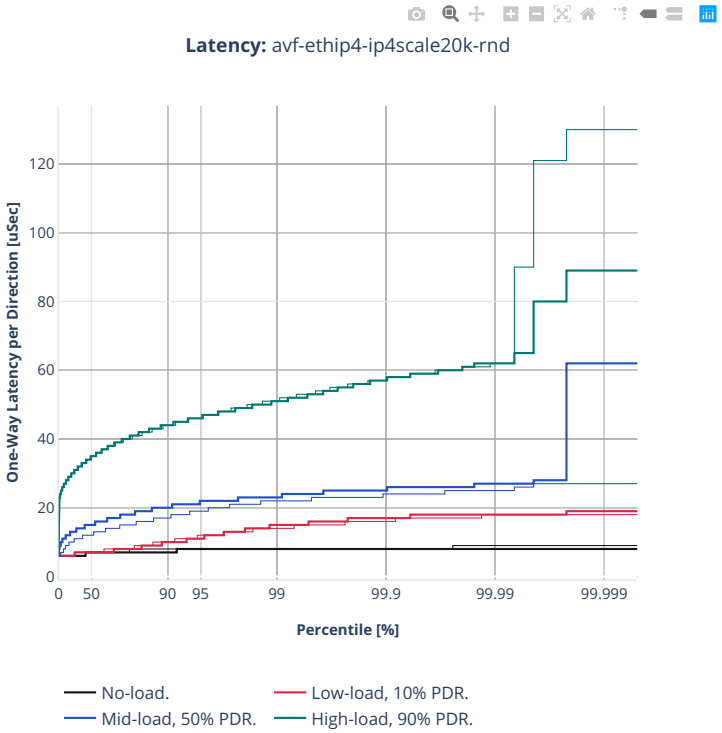
Latency: avf-ethip4-ip4scale20k-rnd
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-ip4routing-features-avf

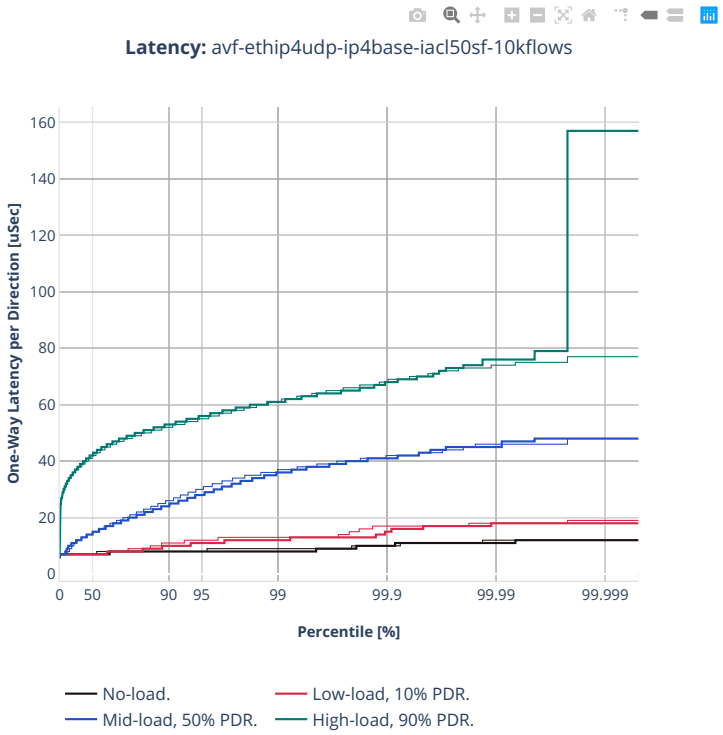

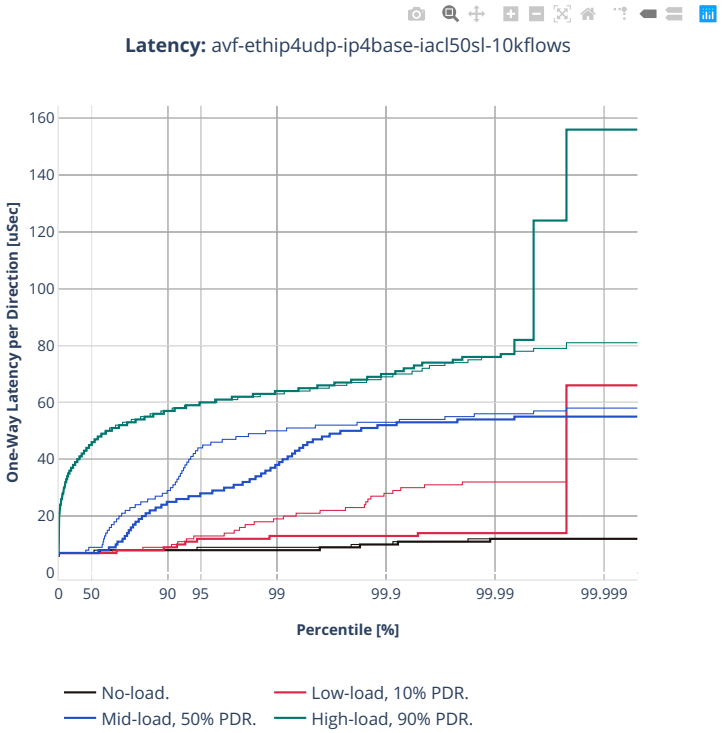
Latency: avf-ethip4udp-ip4base-iacl50sl-10kflows
One-Way Latency per Direction [uSec]
160
140
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

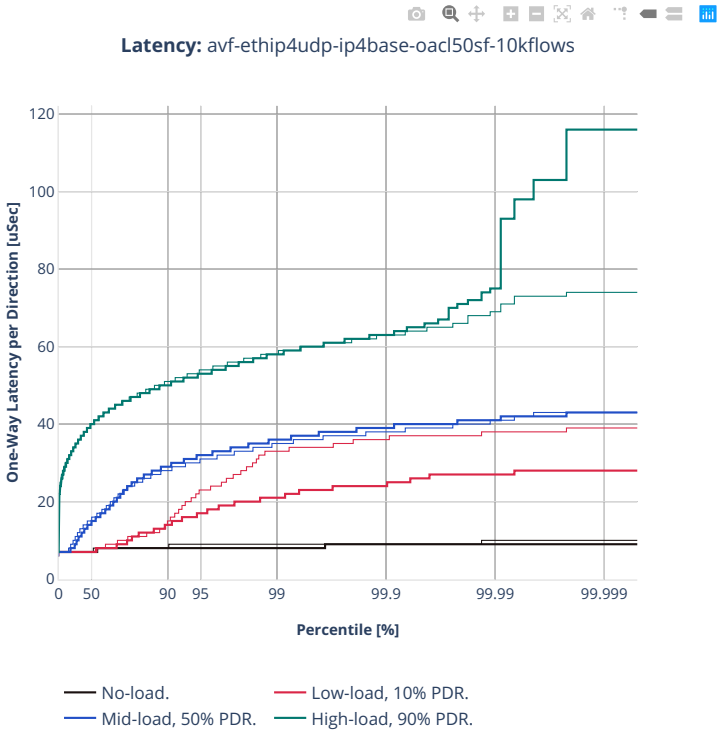
Latency: avf-ethip4udp-ip4base-oacl50sf-10kflows
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

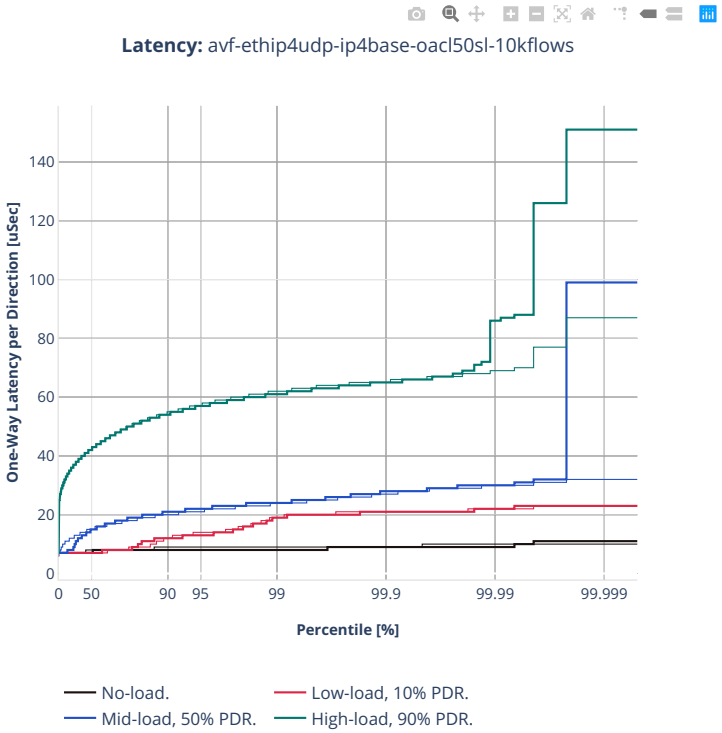
Latency: avf-ethip4udp-ip4base-oacl50sl-10kflows
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

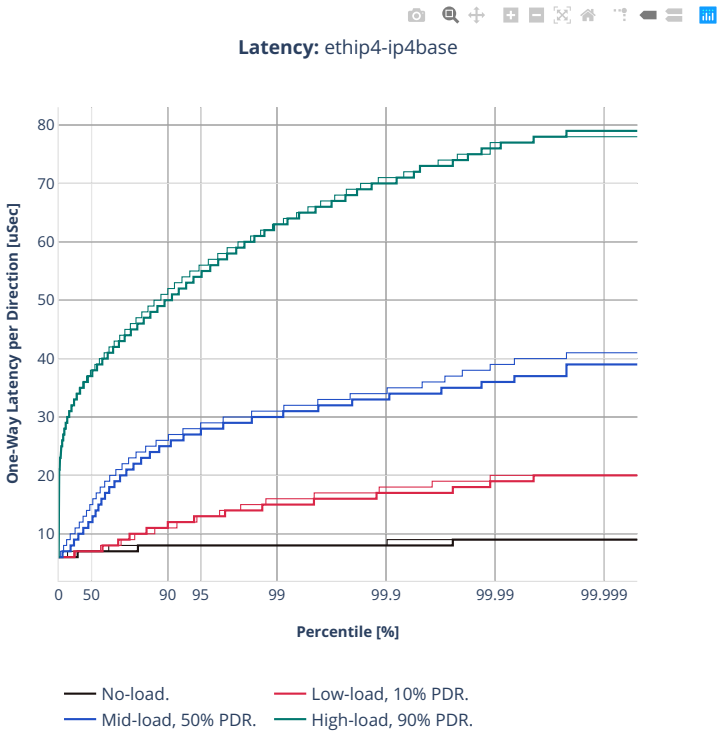

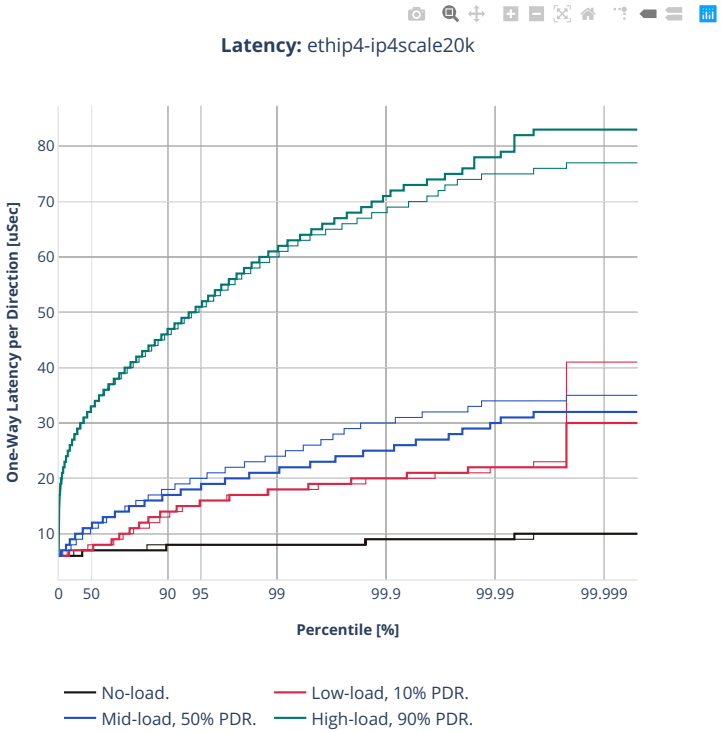
Latency: ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

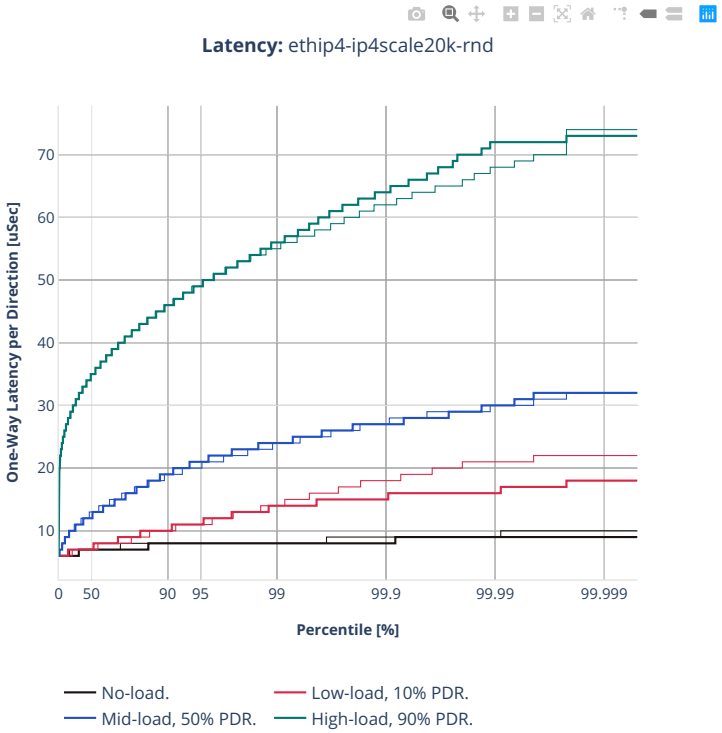
Latency: ethip4-ip4scale20k-rnd
One-Way Latency per Direction [uSec]
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-zn2-xxv710

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

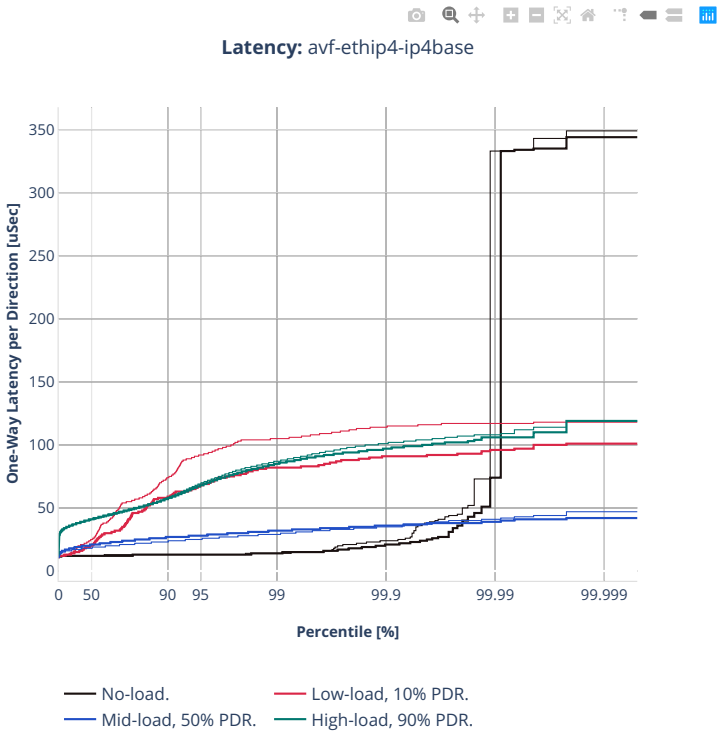

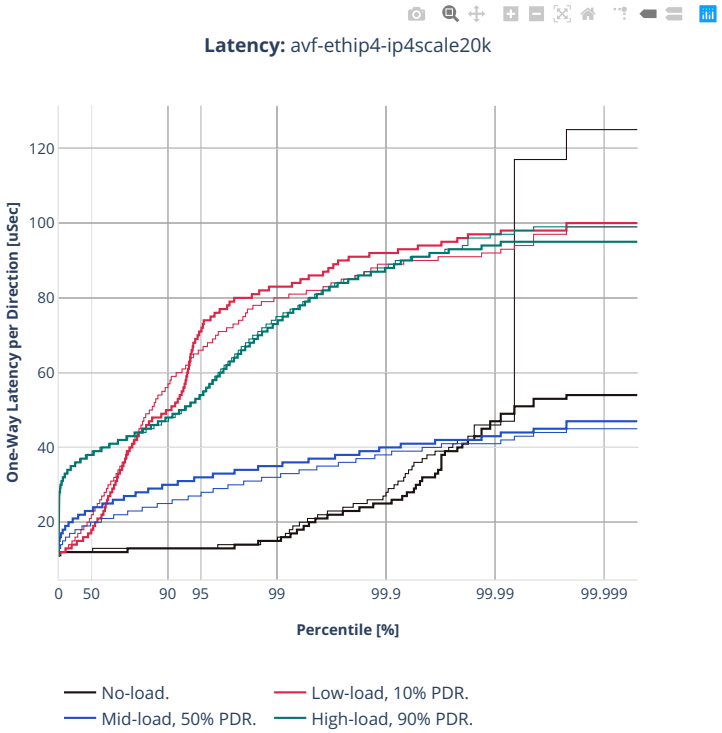
Latency: avf-ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

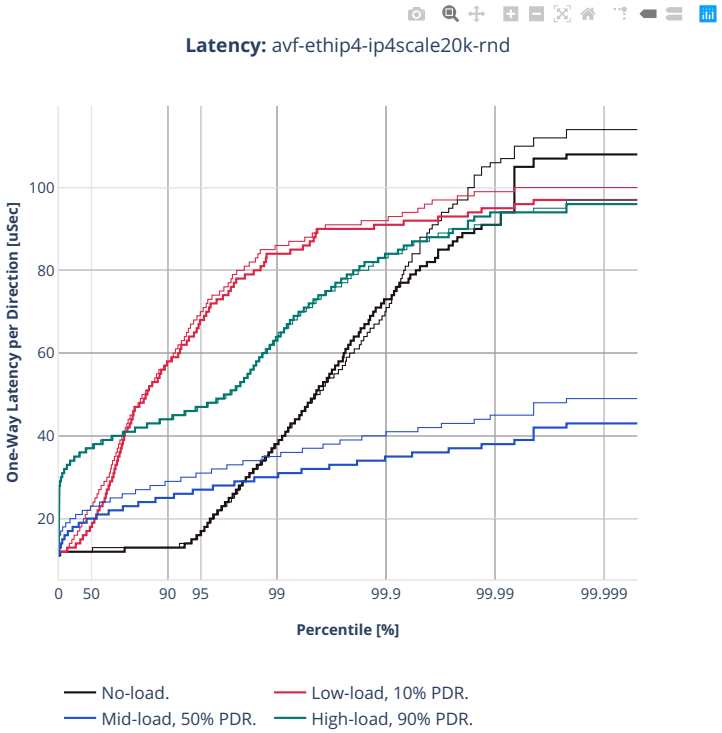
Latency: avf-ethip4-ip4scale20k-rnd
One-Way Latency per Direction [uSec]
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-ip4routing-features-avf

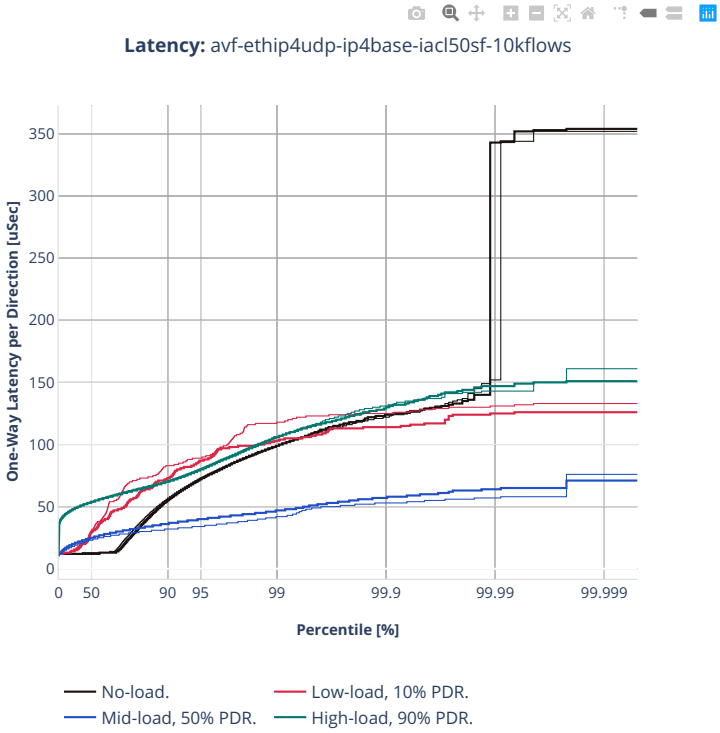

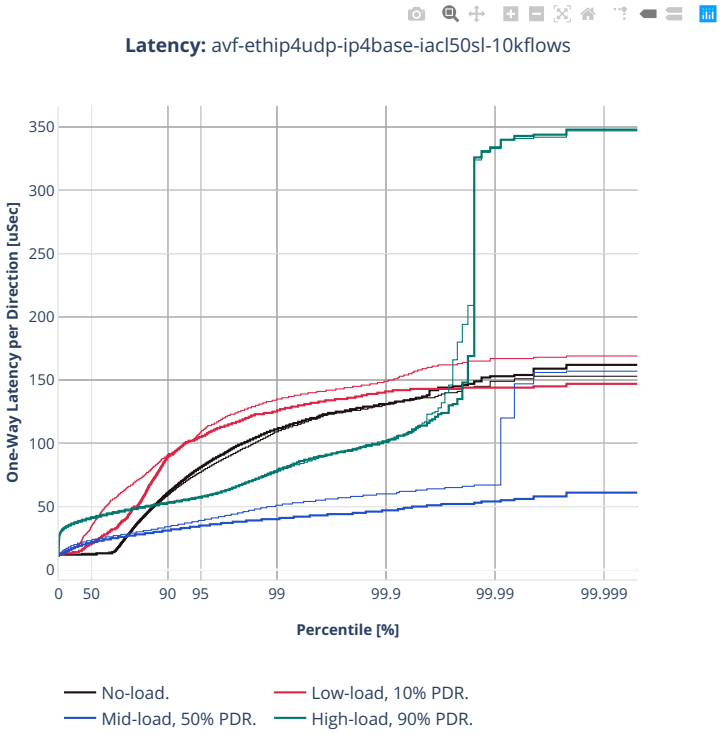
Latency: avf-ethip4udp-ip4base-iacl50sl-10kflows
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

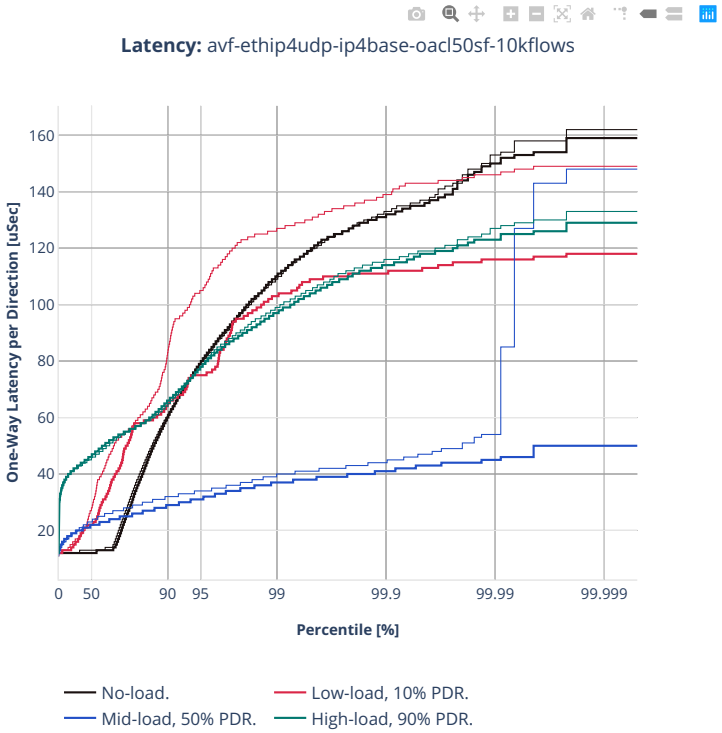
Latency: avf-ethip4udp-ip4base-oacl50sf-10kflows
One-Way Latency per Direction [uSec]
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

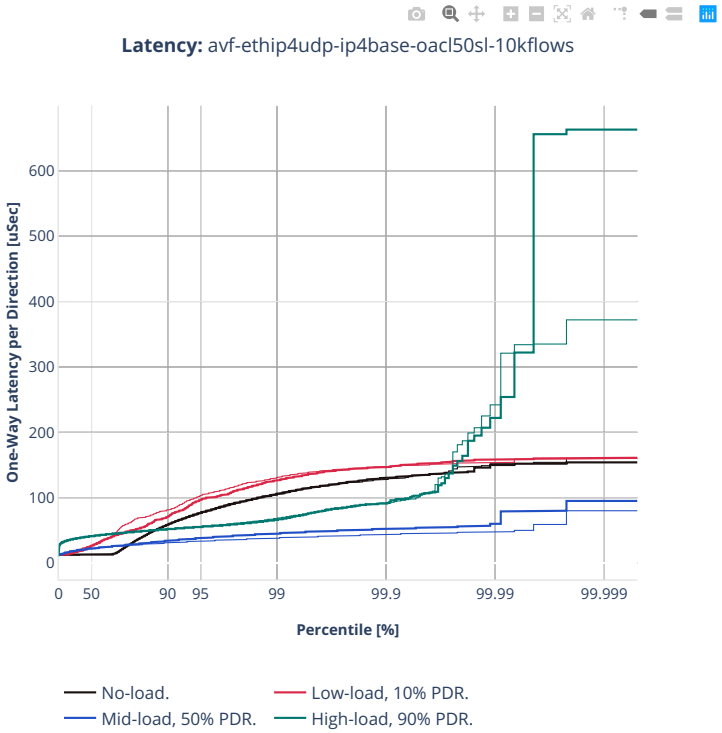
Latency: avf-ethip4udp-ip4base-oacl50sl-10kflows
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

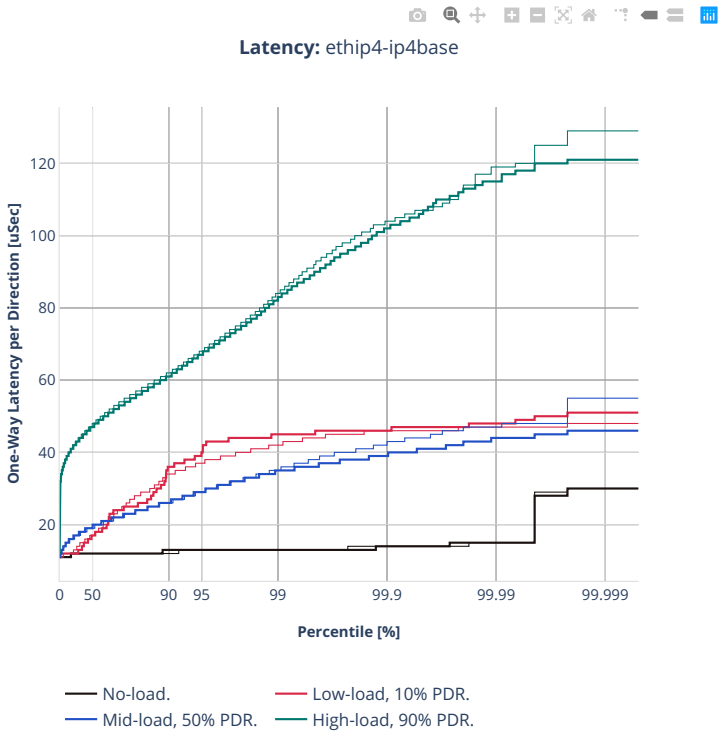

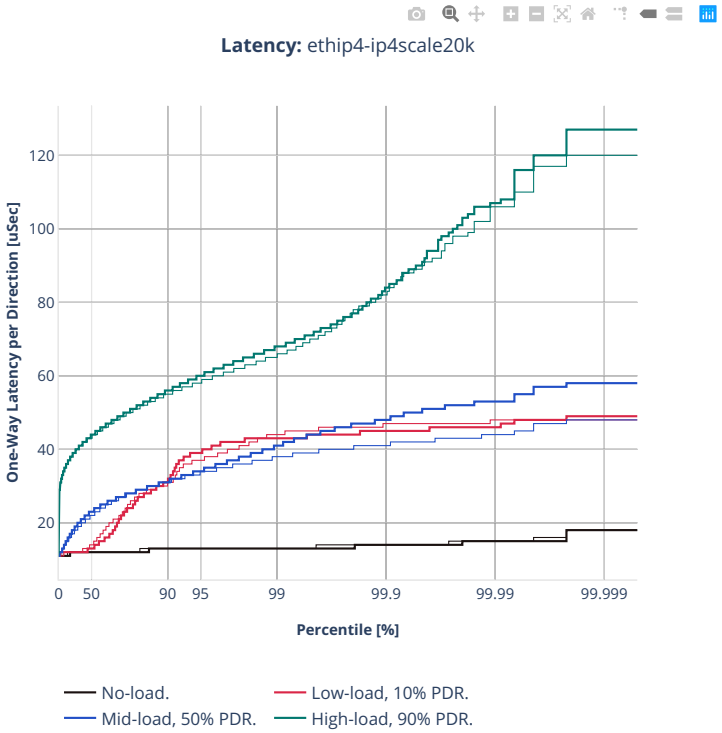
Latency: ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

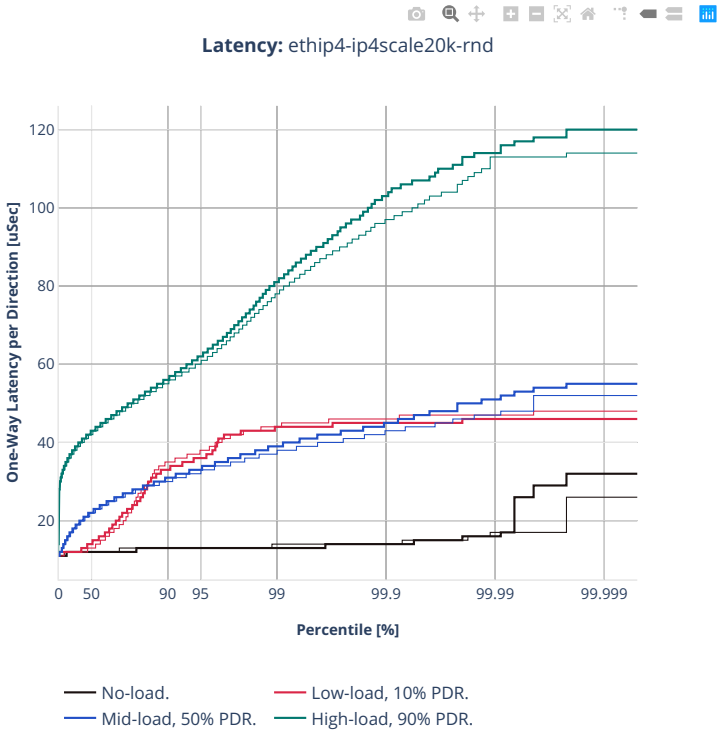
Latency: ethip4-ip4scale20k-rnd
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-alt-xl710

64b-1t1c-ip4routing-base-scale

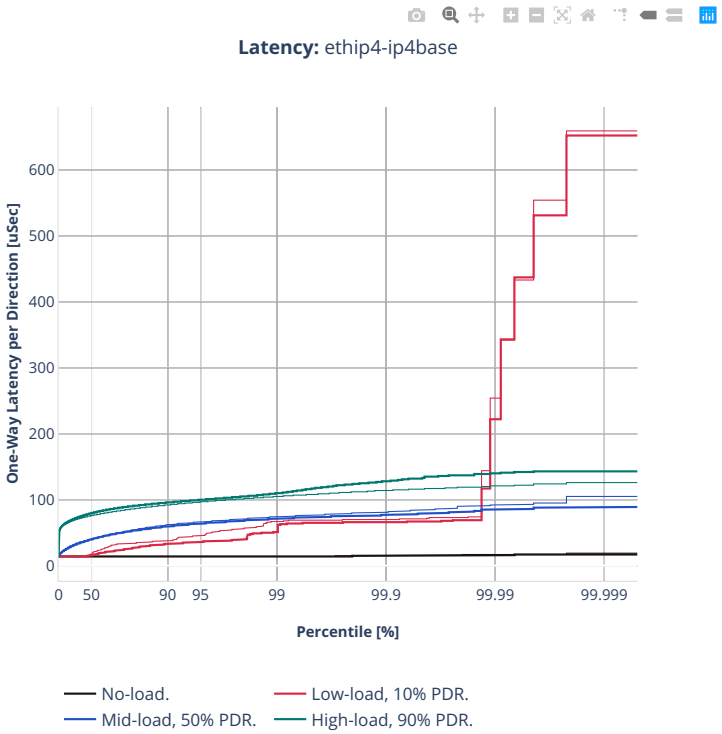

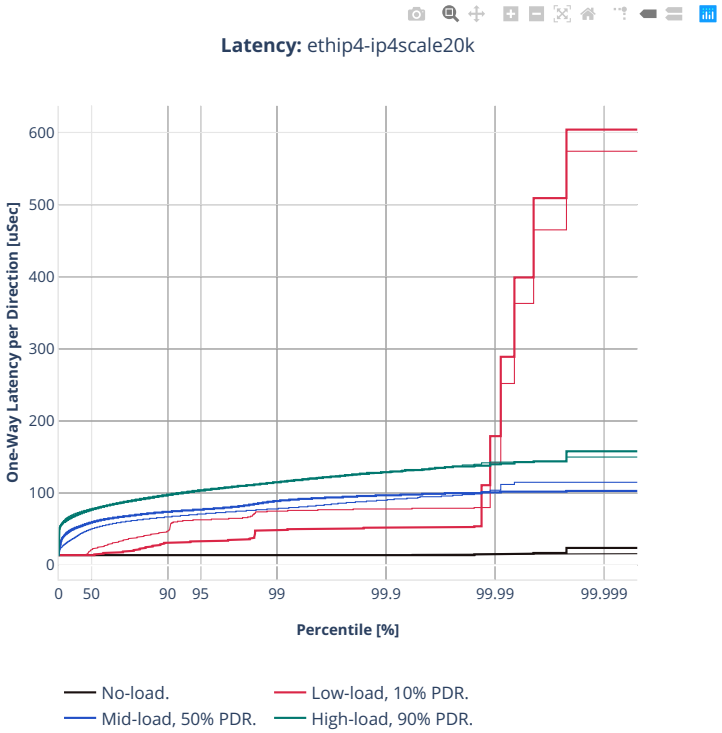
Latency: ethip4-ip4scale20k
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

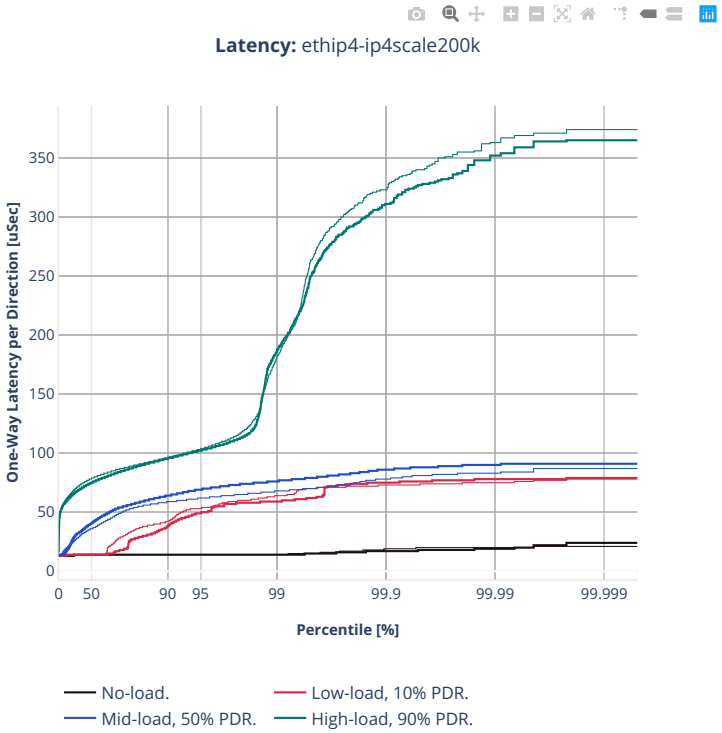
Latency: ethip4-ip4scale200k
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-1t1c-ip4routing-features

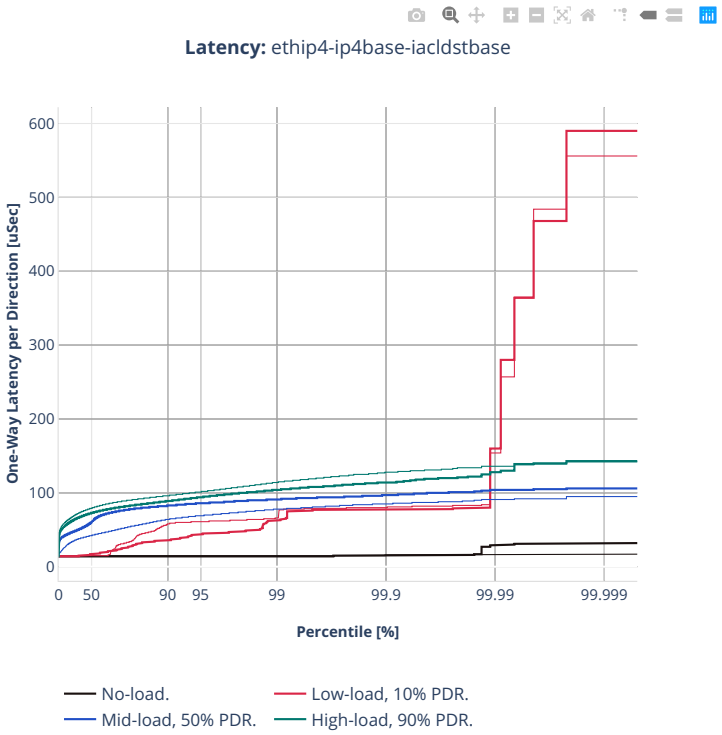

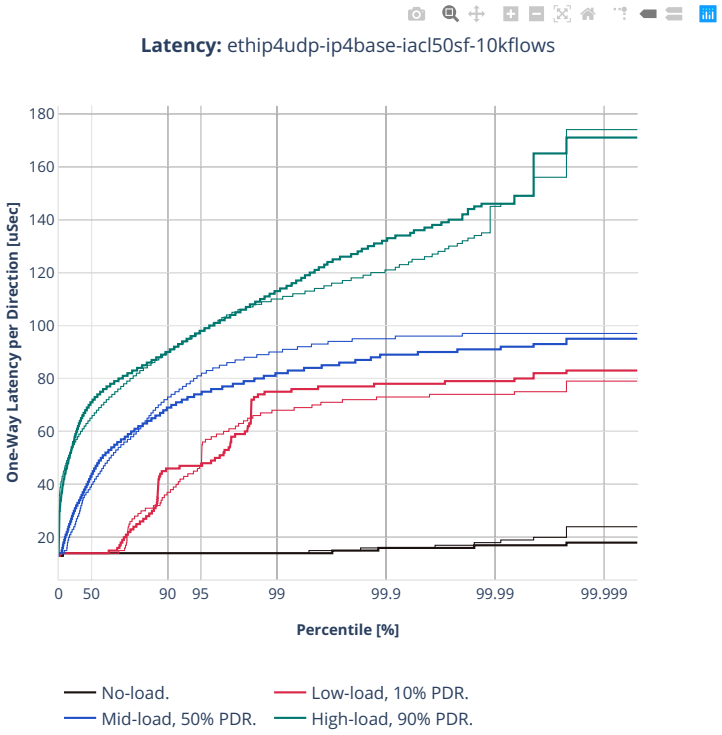
Latency: ethip4udp-ip4base-iacl50sf-10kflows
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

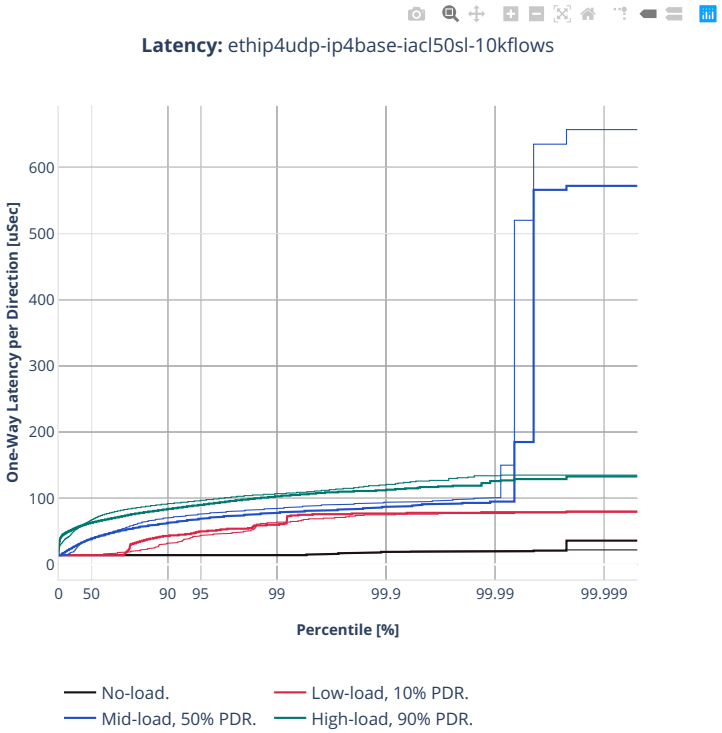
Latency: ethip4udp-ip4base-iacl50sl-10kflows
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

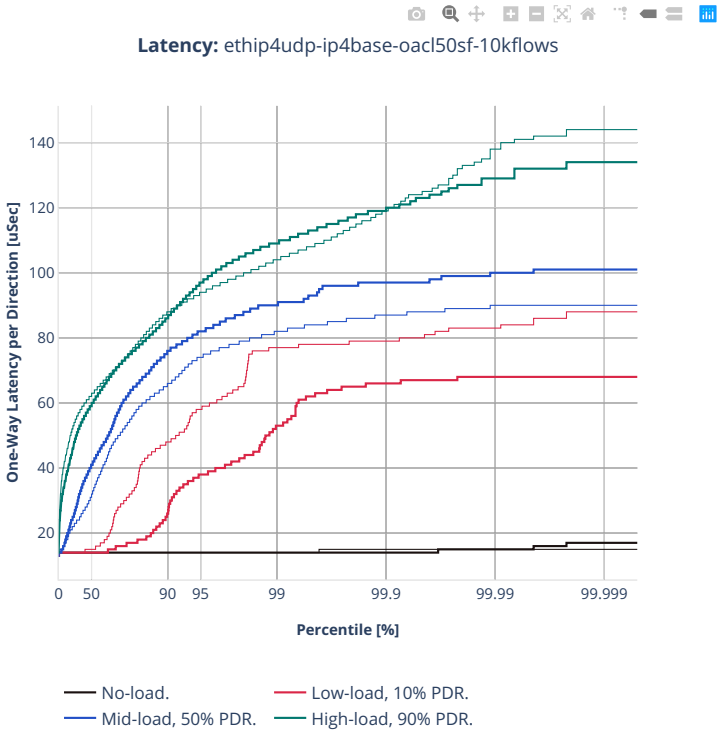
Latency: ethip4udp-ip4base-oacl50sf-10kflows
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

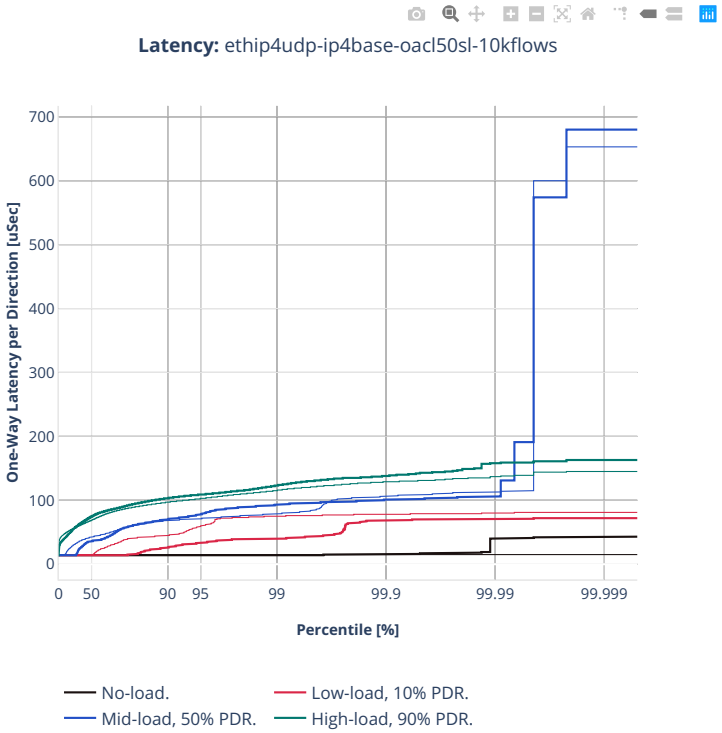
Latency: ethip4udp-ip4base-oacl50sl-10kflows
One-Way Latency per Direction [uSec]
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-tsh-x520

64b-1t1c-ip4routing-features-ixgbe

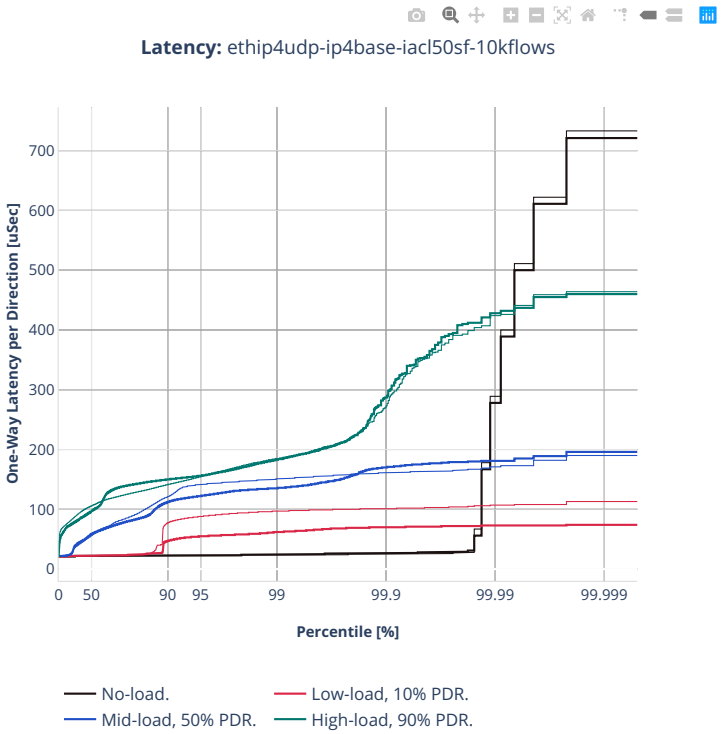

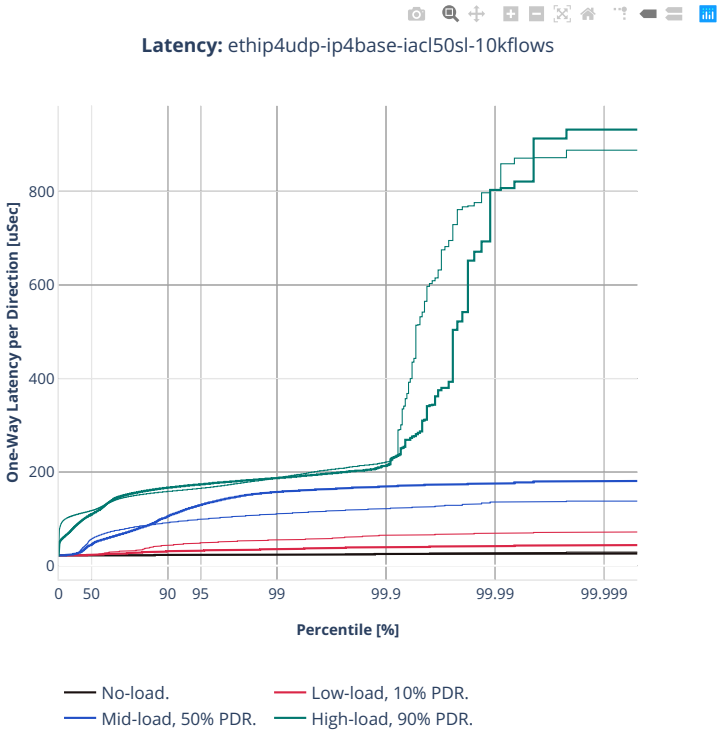
Latency: ethip4udp-ip4base-iacl50sl-10kflows
One-Way Latency per Direction [uSec]
800
600
400
200
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

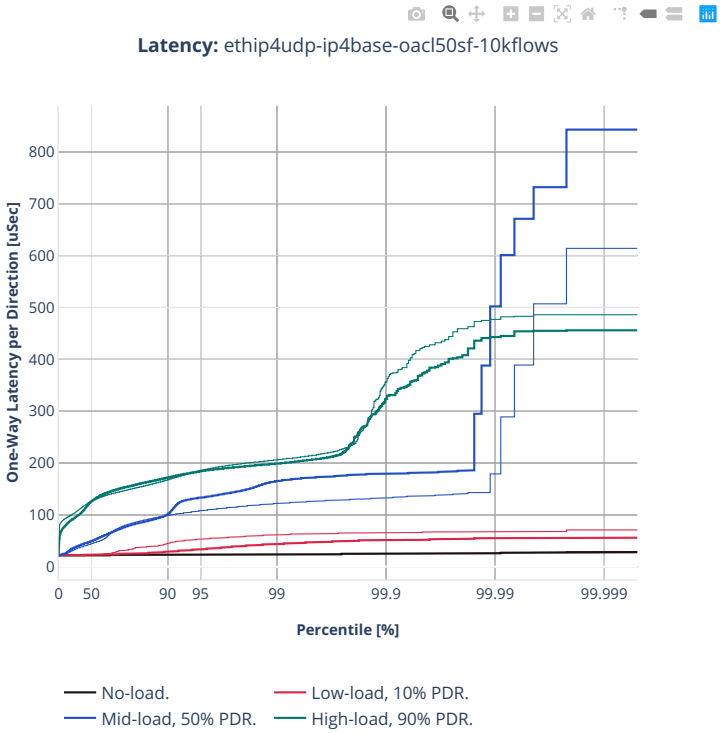
Latency: ethip4udp-ip4base-oacl50sf-10kflows
One-Way Latency per Direction [uSec]
800
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-tx2-xl710

64b-1t1c-ip4routing-base-dpdk

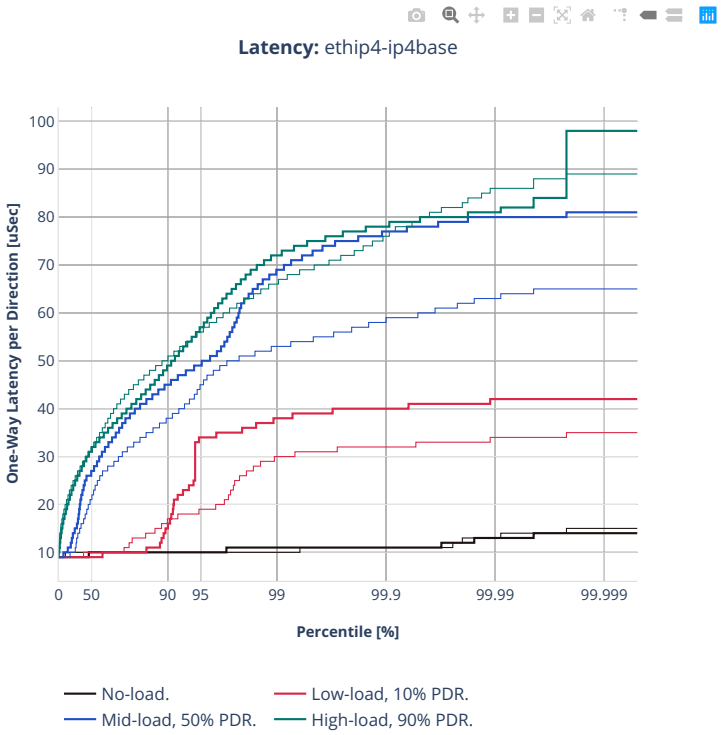

64b-1t1c-ip4routing-scale-dpdk

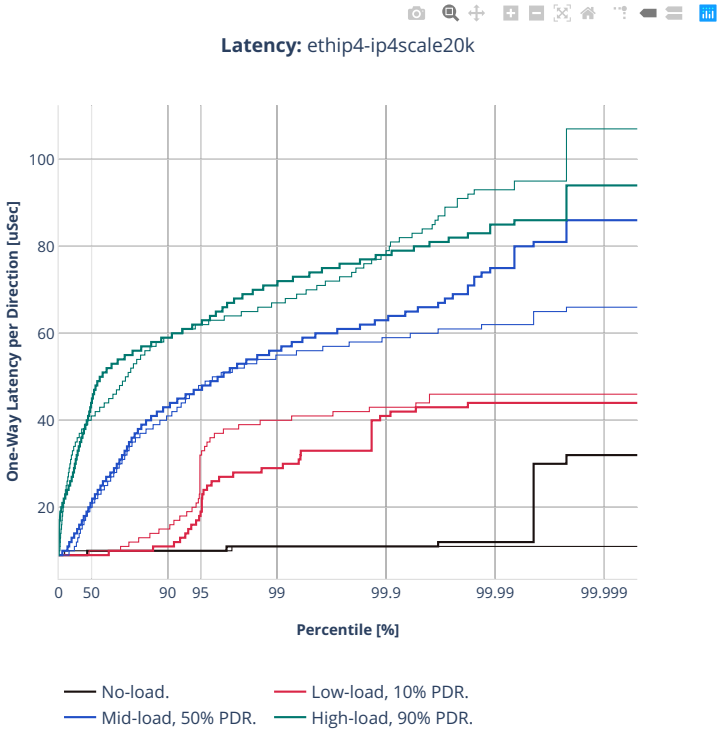

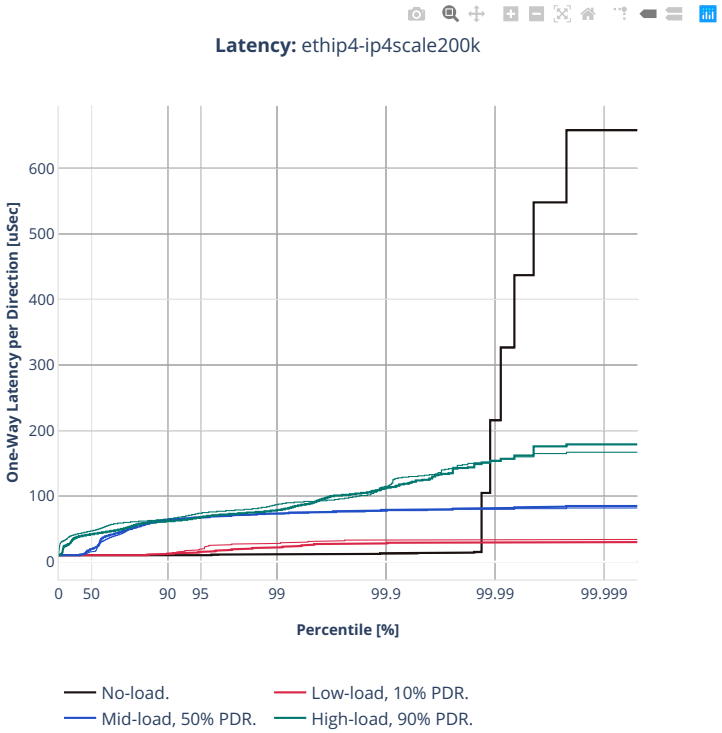
Latency: ethip4-ip4scale200k
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-1t1c-features-ip4routing-base-dpdk

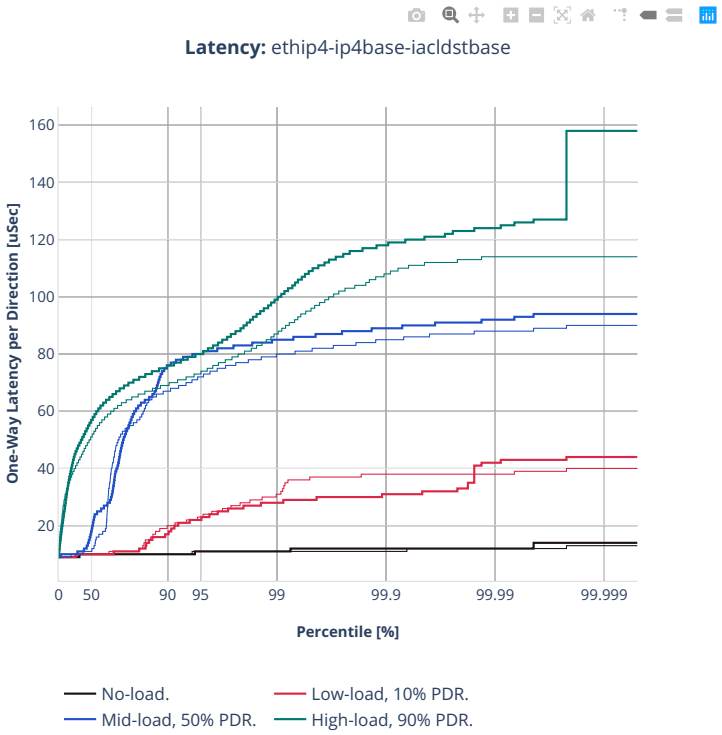

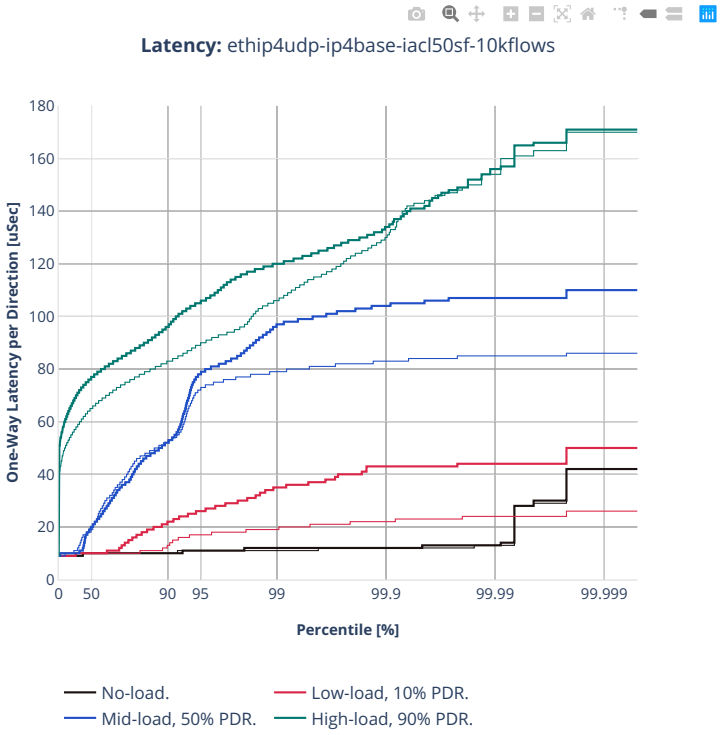
Latency: ethip4udp-ip4base-iacl50sf-10kflows
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

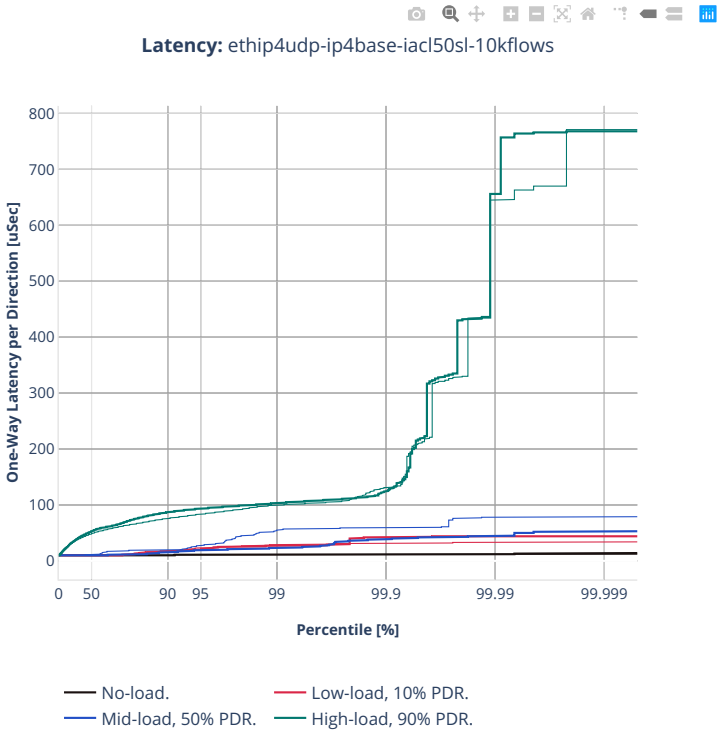
Latency: ethip4udp-ip4base-iacl50sl-10kflows
One-Way Latency per Direction [uSec]
800
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

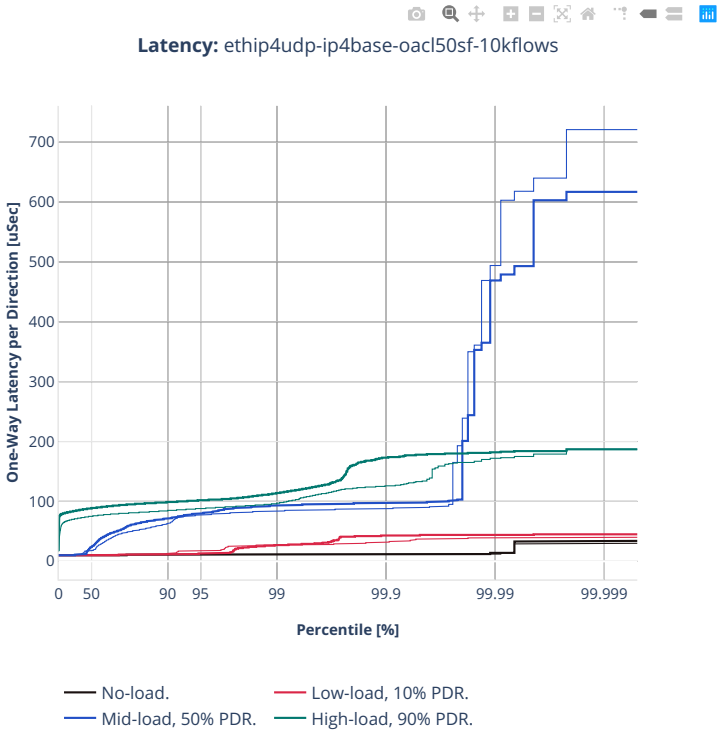
Latency: ethip4udp-ip4base-oacl50sf-10kflows
One-Way Latency per Direction [uSec]
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

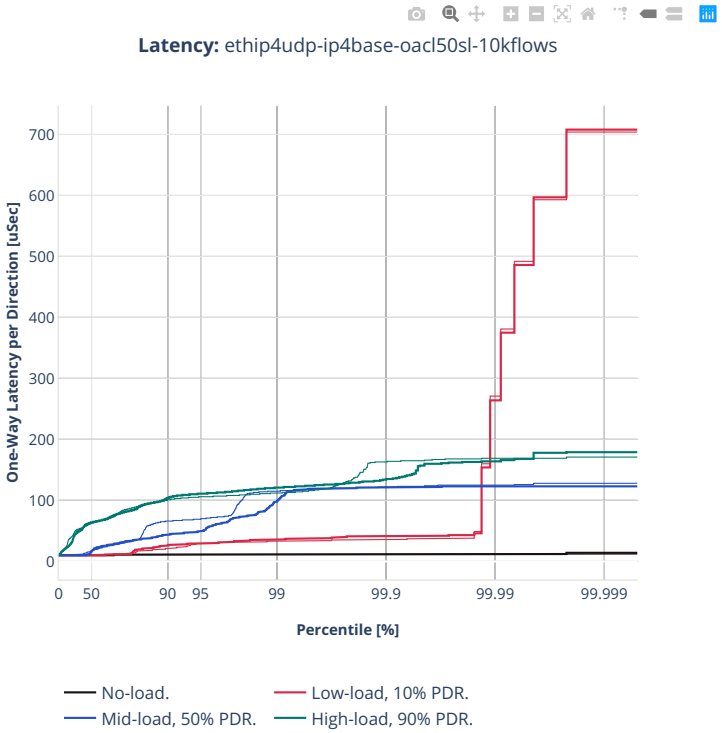
Latency: ethip4udp-ip4base-oacl50sl-10kflows
One-Way Latency per Direction [uSec]
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2.5.3 IPv6 Routing

CSIT source code for the test cases used for plots can be found in CSIT git repository[151].

[151] https://git.fd.io/csit/tree/tests/vpp/perf/ip6?h=rls2206

2n-icx-xxv710

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

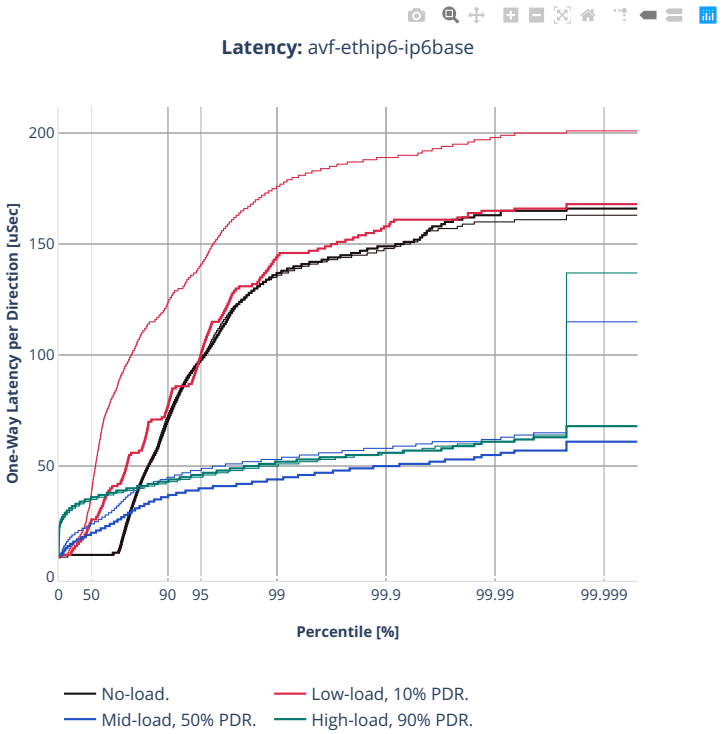

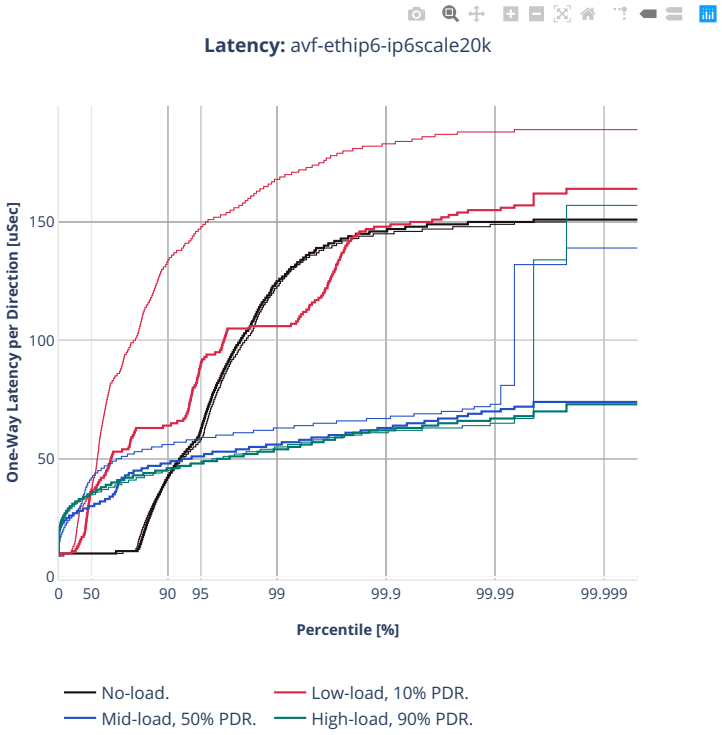
Latency: avf-ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

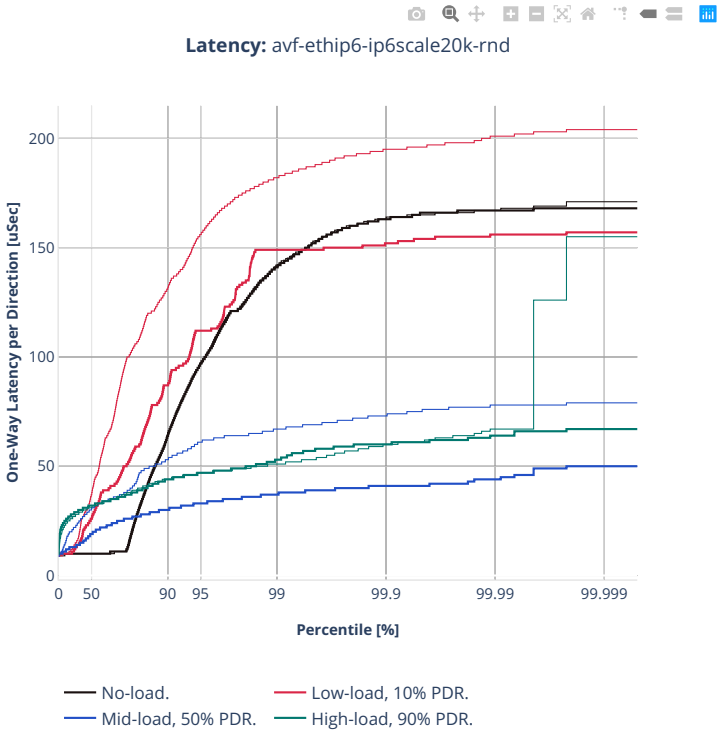
Latency: avf-ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

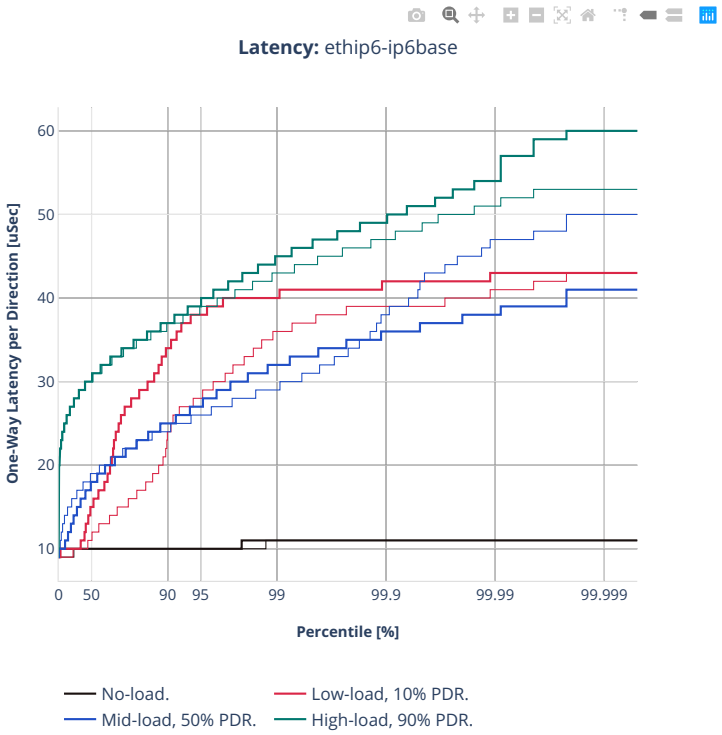

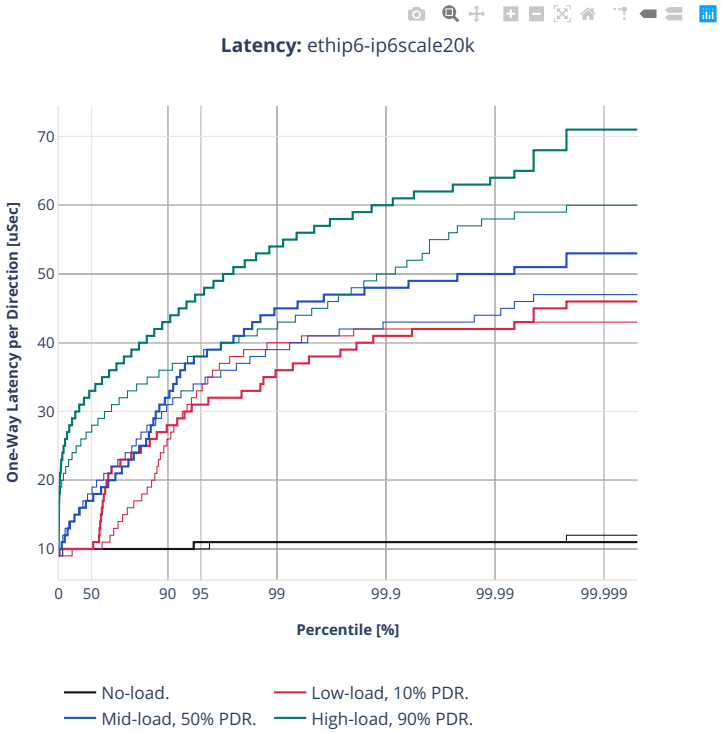
Latency: ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

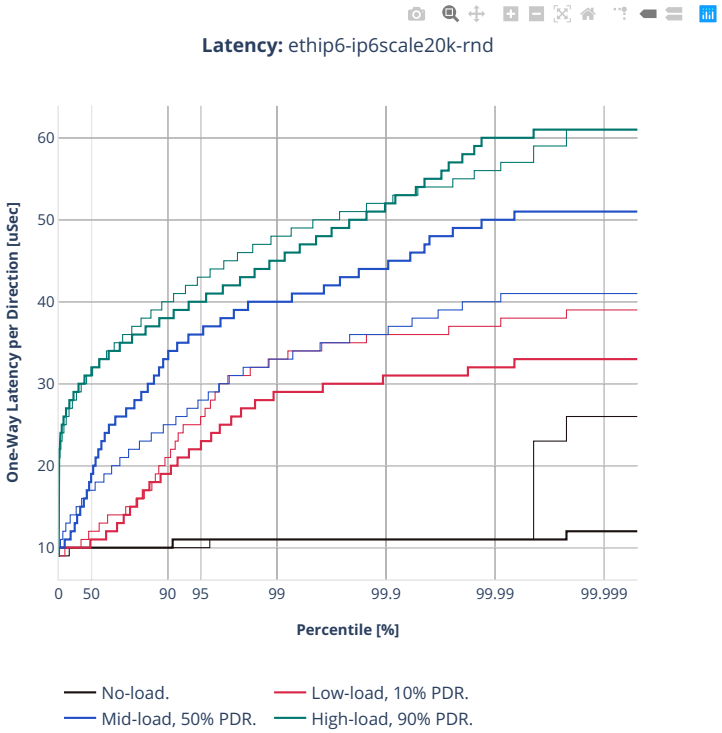
Latency: ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

78b-2t1c-ip6routing-base-scale-af-xdp

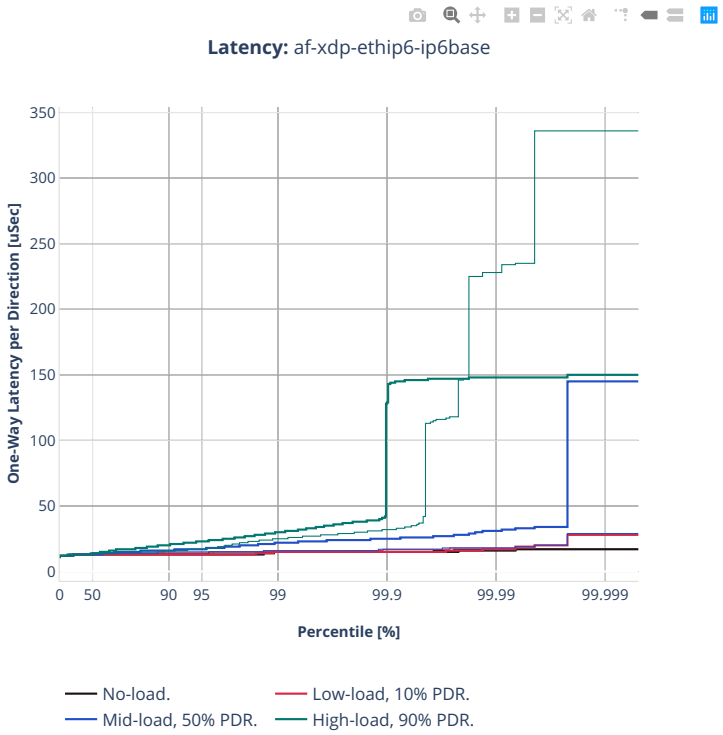

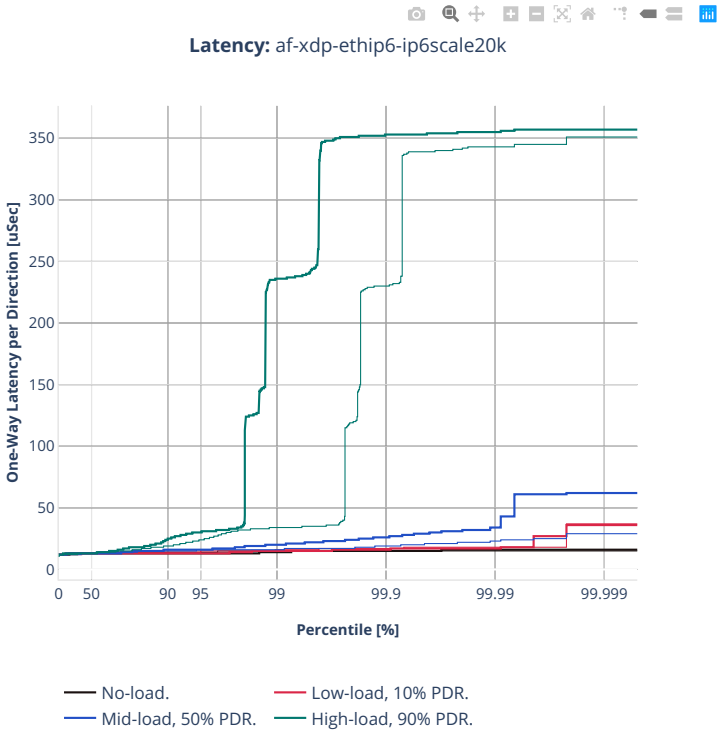
Latency: af-xdp-ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

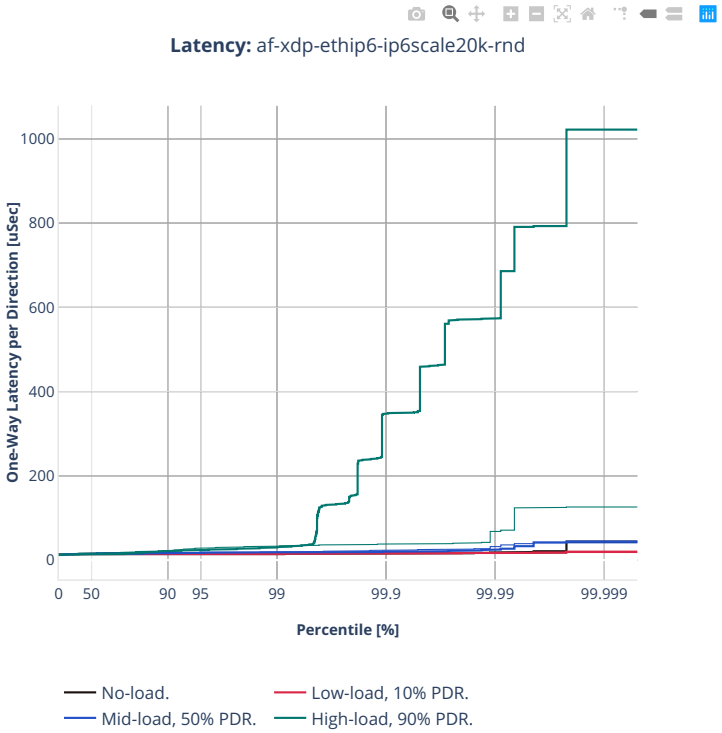
Latency: af-xdp-ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
1000
800
600
400
200
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-icx-xxv710

78b-2t1c-ip6routing-base-avf

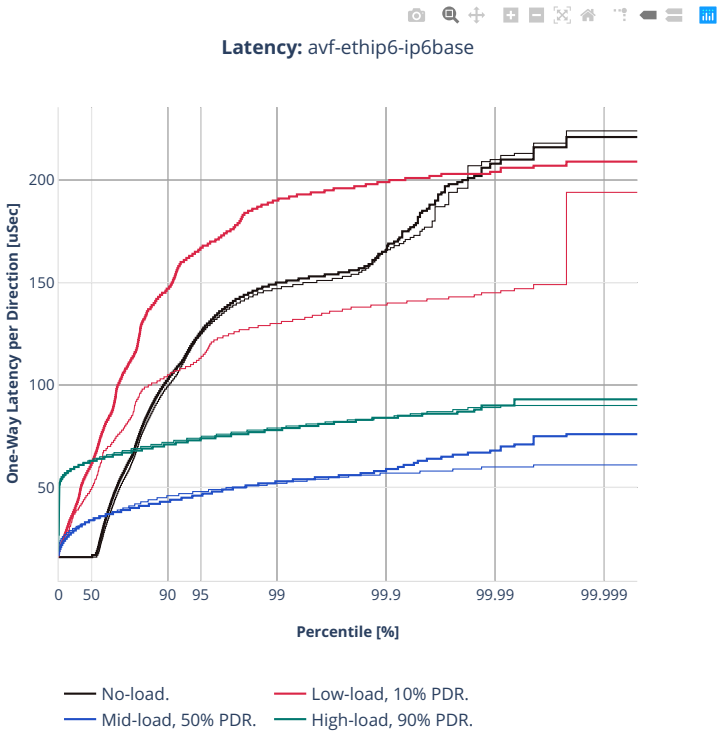

78b-2t1c-ip6routing-base-dpdk

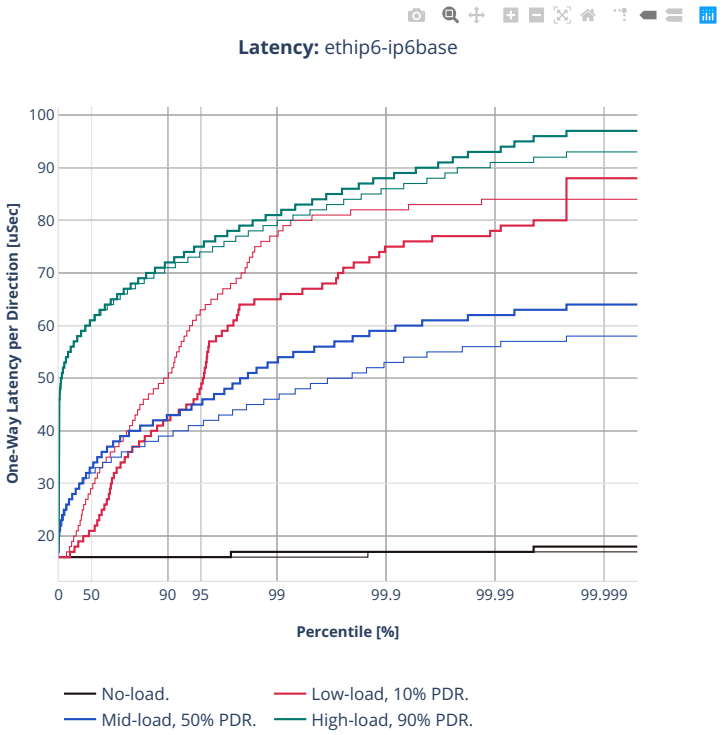

2n-skx-xxv710

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

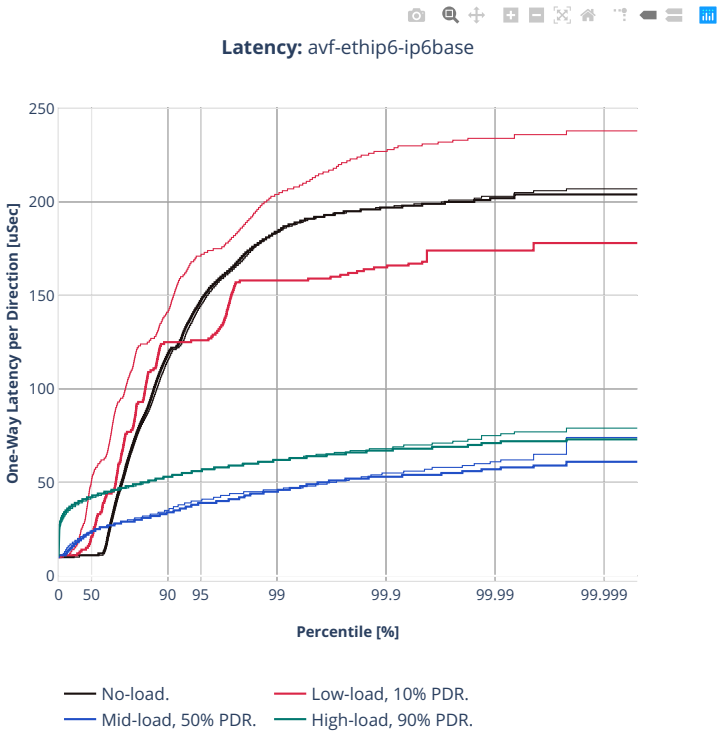

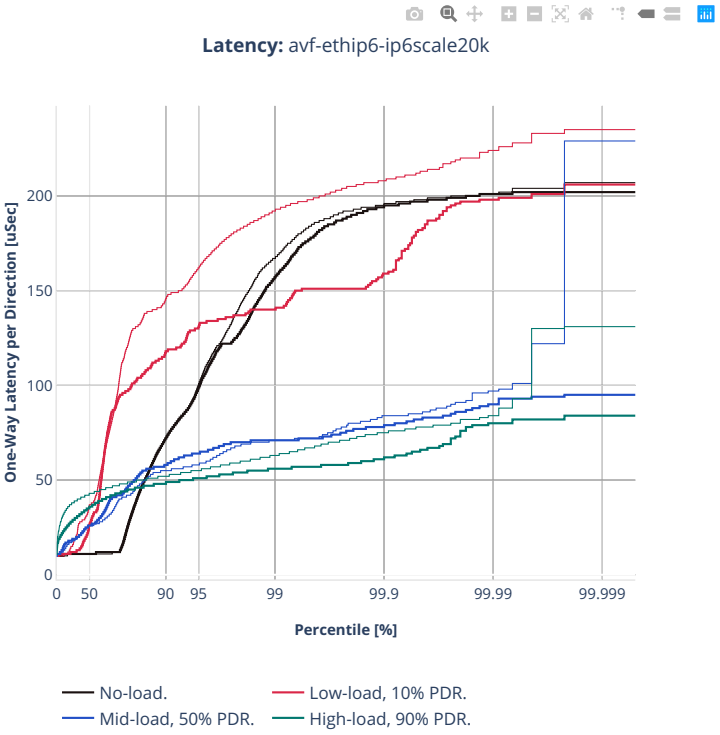
Latency: avf-ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
0
50
100
150
200
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

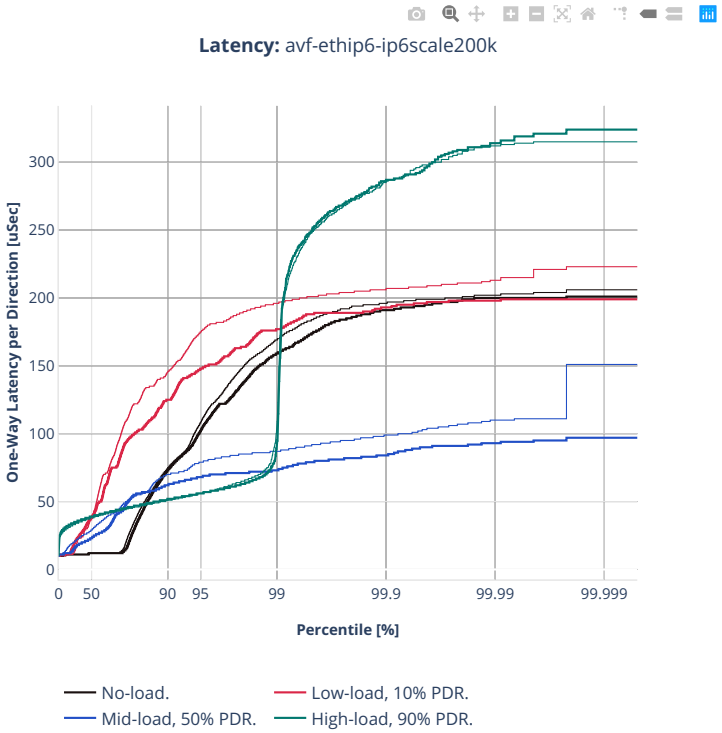
Latency: avf-ethip6-ip6scale200k
One-Way Latency per Direction [uSec]
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

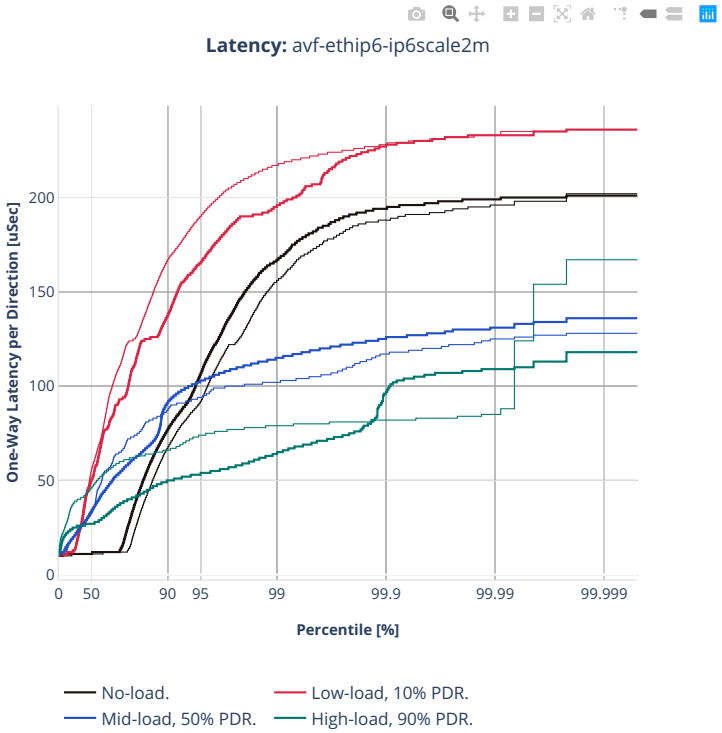
Latency: avf-ethip6-ip6scale2m
One-Way Latency per Direction [uSec]
0
50
100
150
200
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

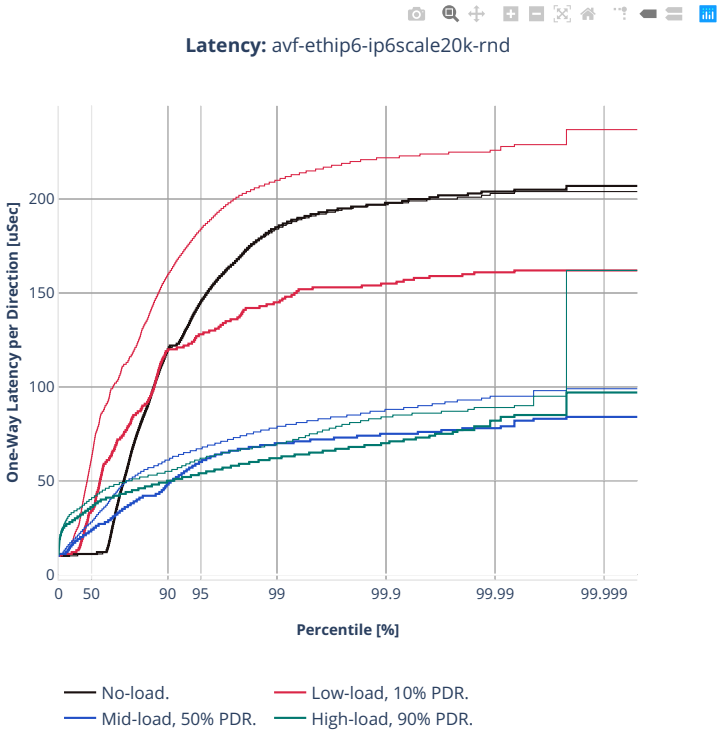
Latency: avf-ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

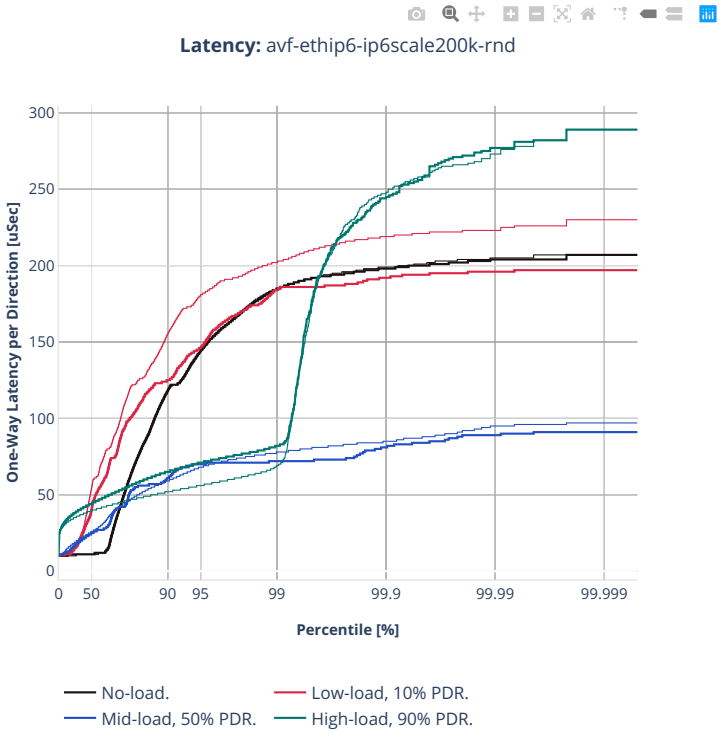
Latency: avf-ethip6-ip6scale200k-rnd
One-Way Latency per Direction [uSec]
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

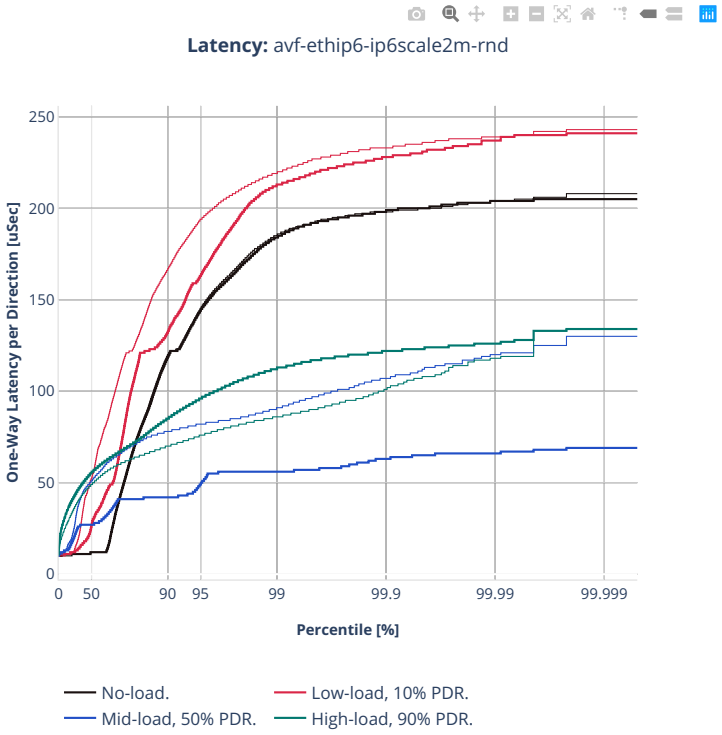
Latency: avf-ethip6-ip6scale2m-rnd
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

78b-2t1c-ip6routing-base-scale-af-xdp

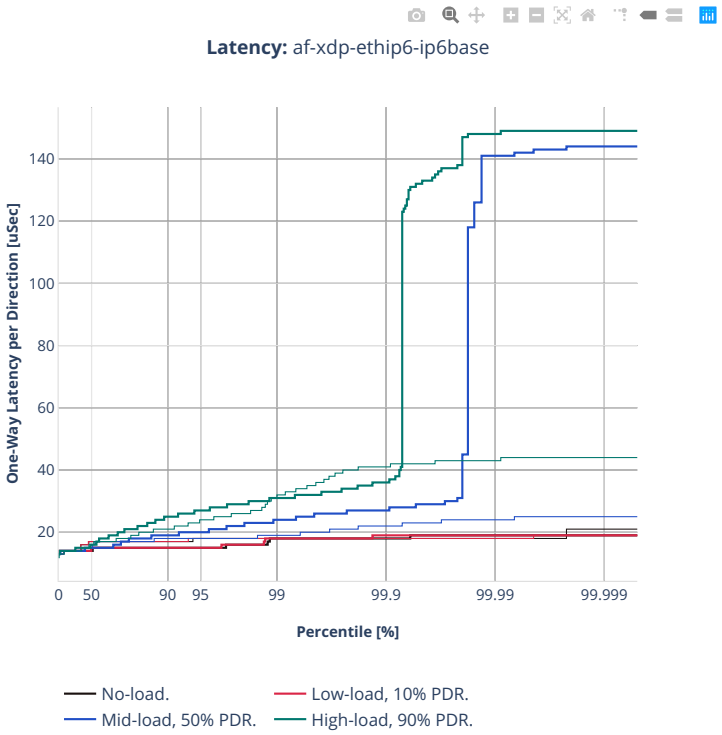

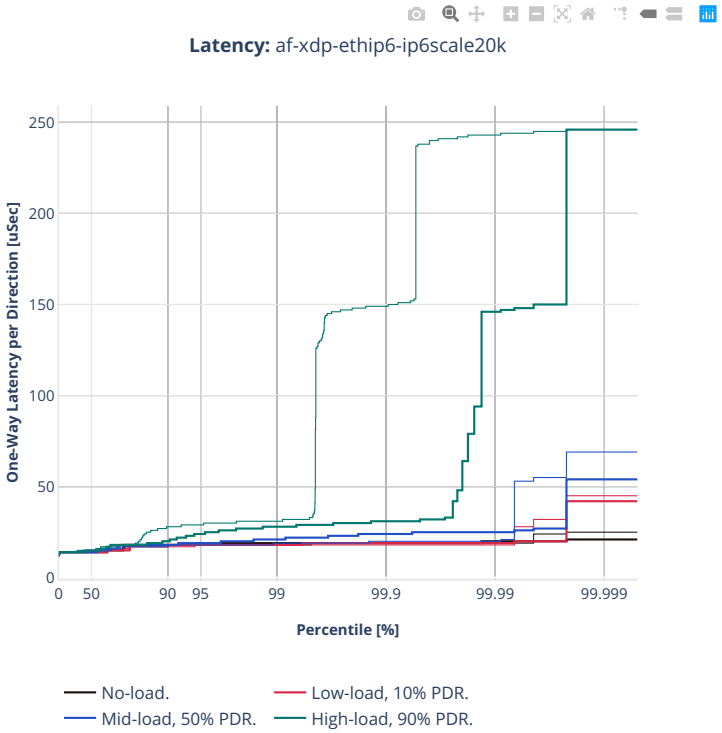
Latency: af-xdp-ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

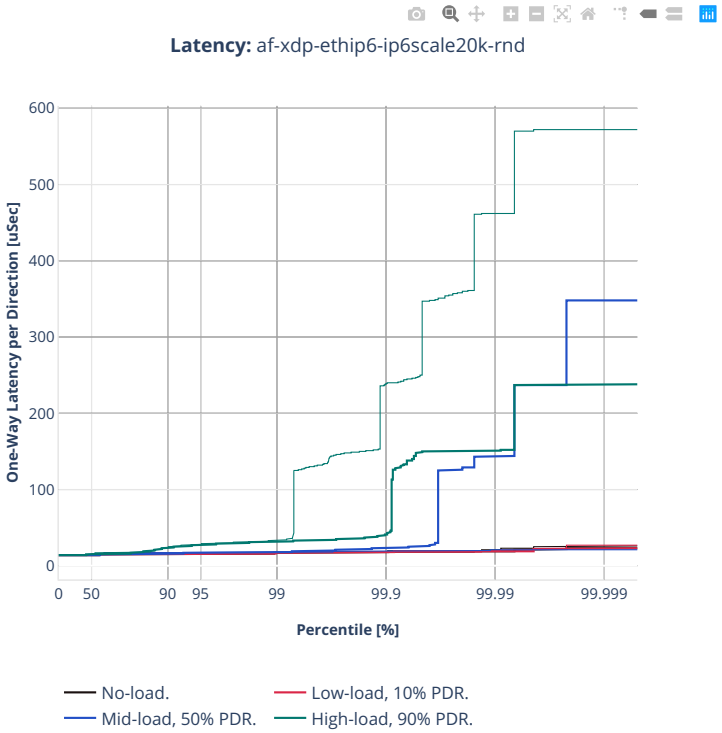
Latency: af-xdp-ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

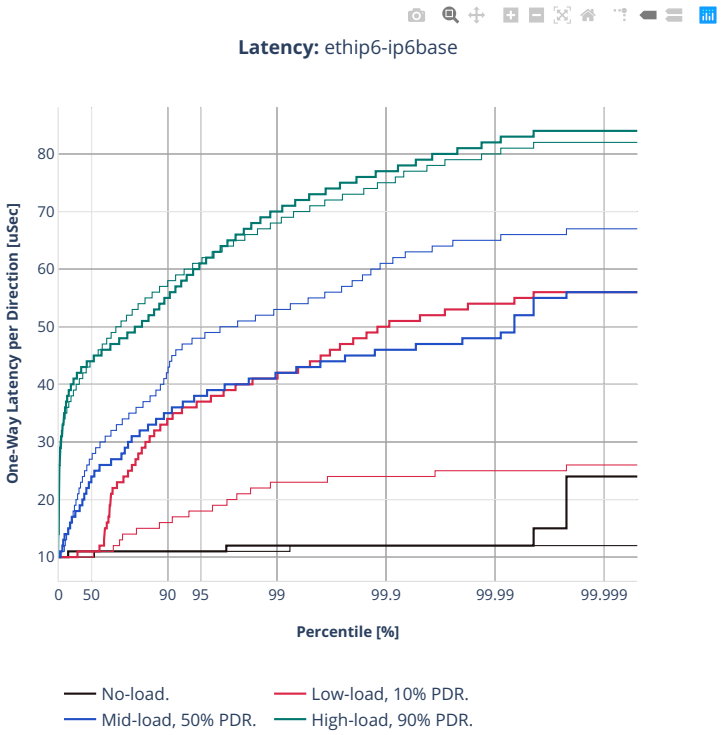

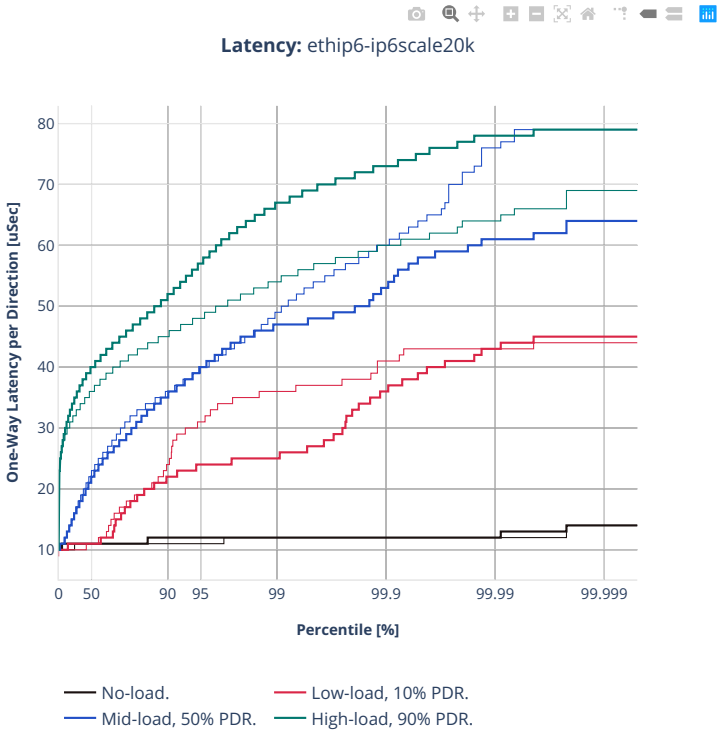
Latency: ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

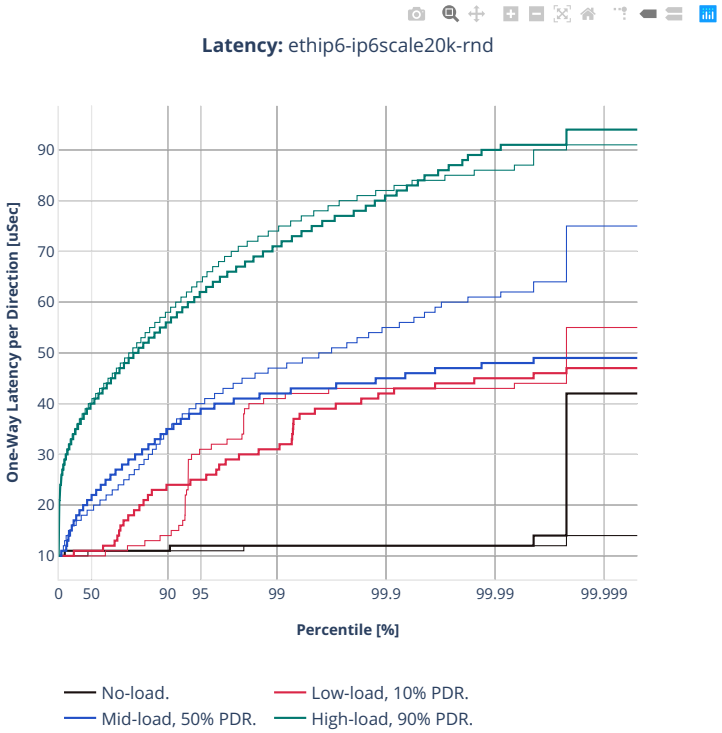
Latency: ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
90
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-skx-xxv710

78b-2t1c-ip6routing-base-avf

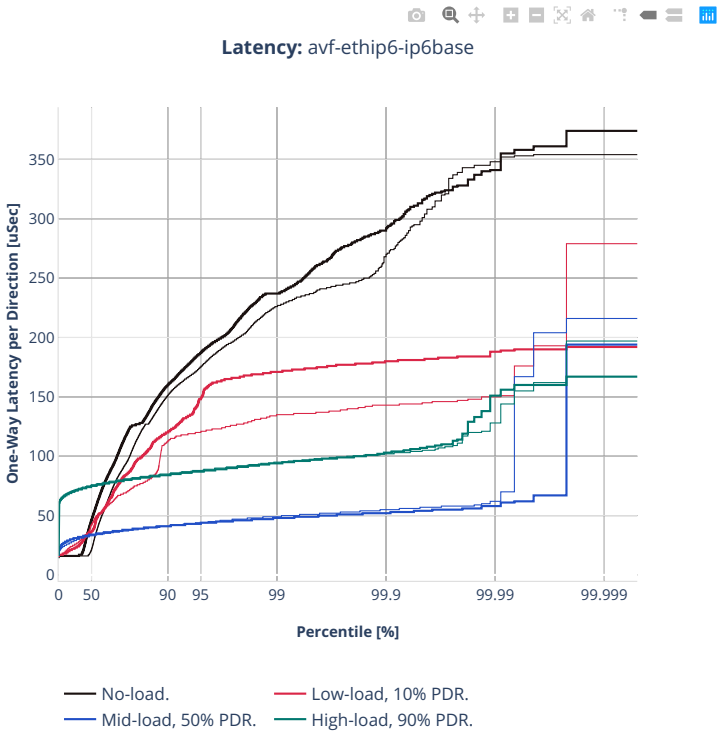

78b-2t1c-ip6routing-base-dpdk

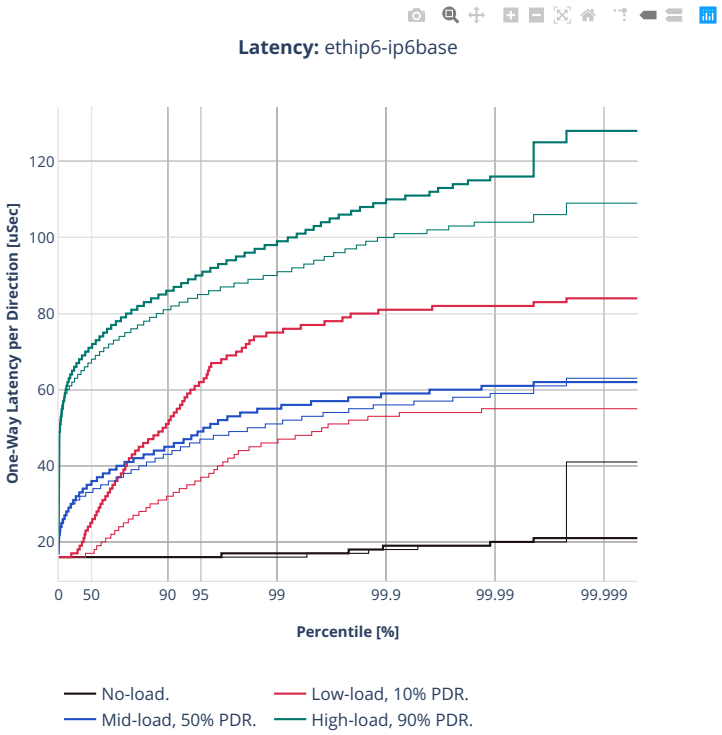

2n-clx-xxv710

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

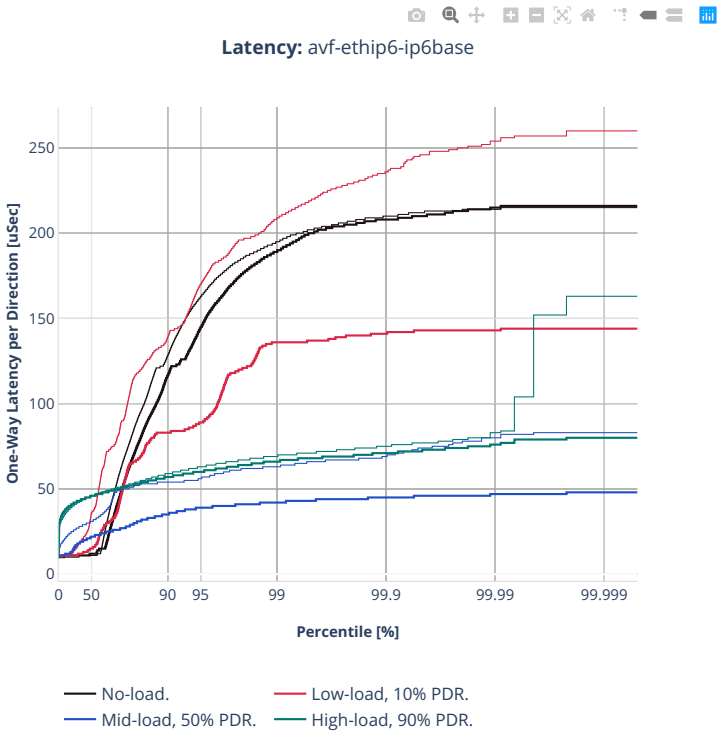

Latency: avf-ethip6-ip6scale20k

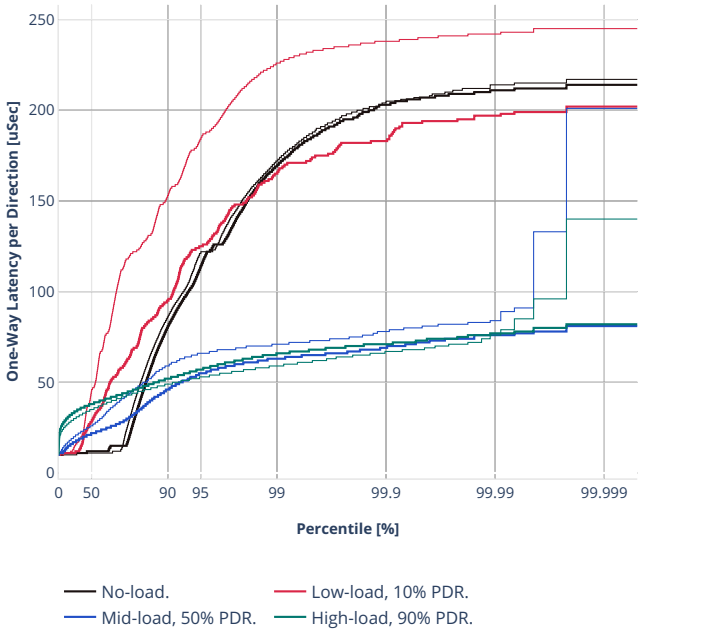

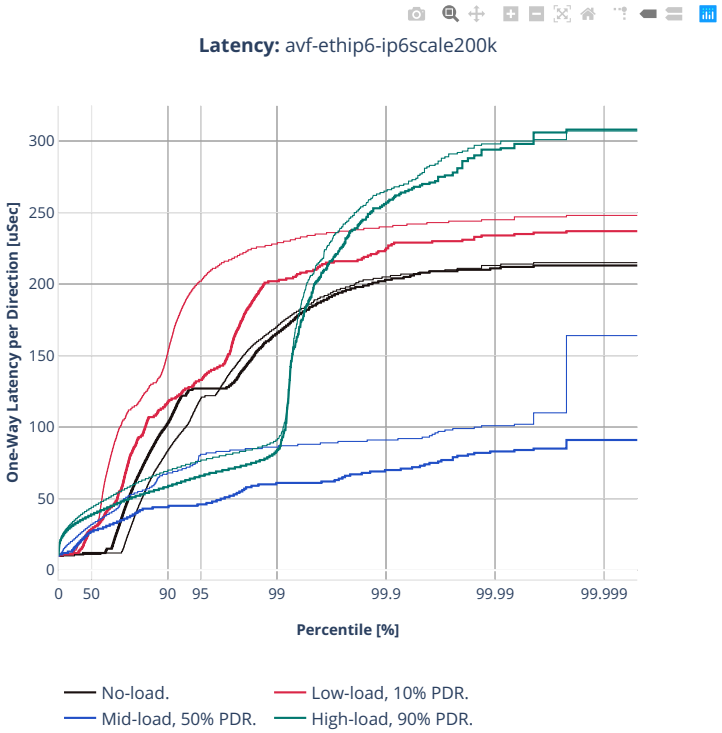
Latency: avf-ethip6-ip6scale200k
One-Way Latency per Direction [uSec]
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

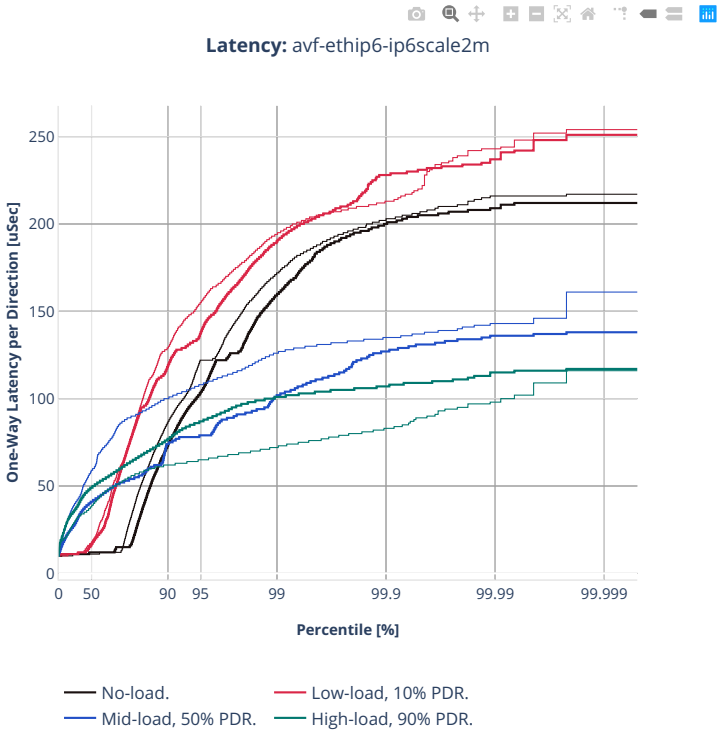
Latency: avf-ethip6-ip6scale2m
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

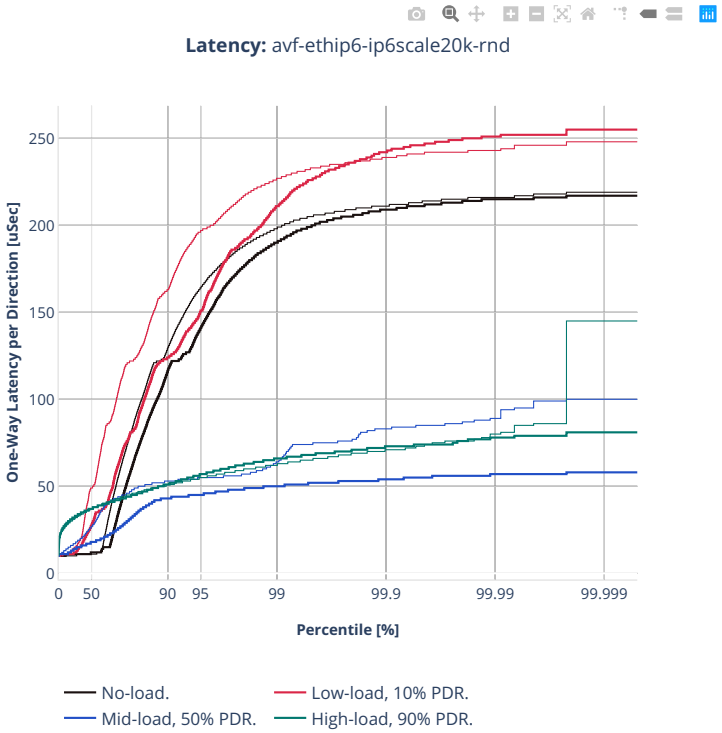
Latency: avf-ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

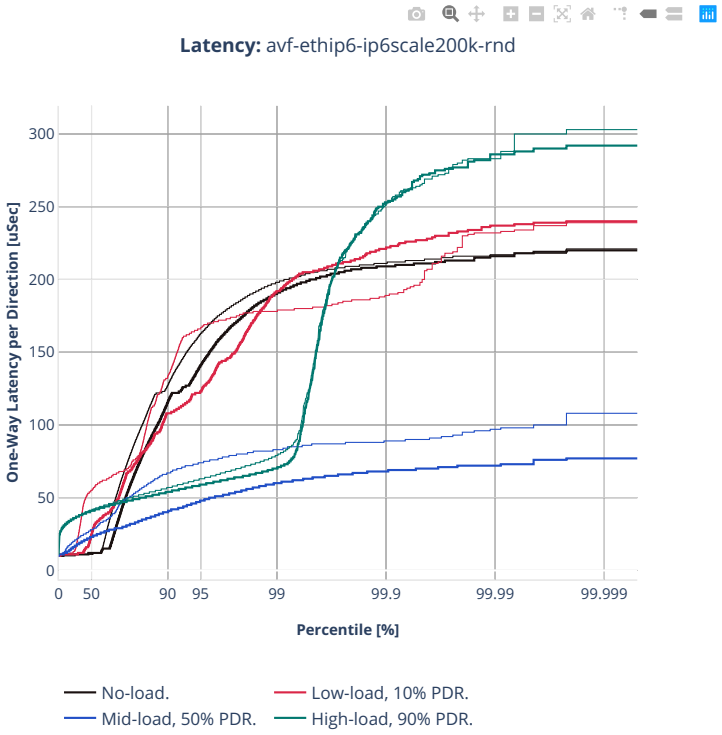
Latency: avf-ethip6-ip6scale200k-rnd
One-Way Latency per Direction [uSec]
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

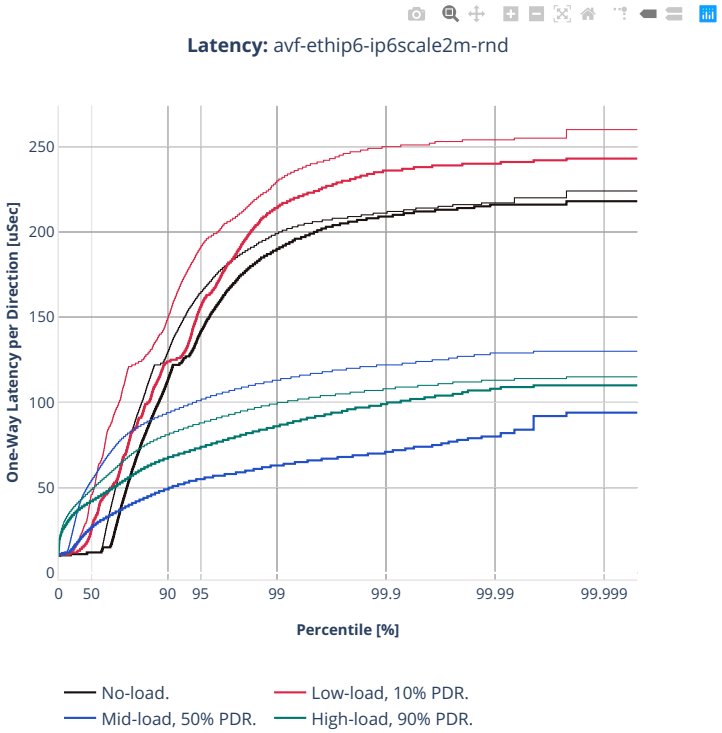
Latency: avf-ethip6-ip6scale2m-rnd
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

78b-2t1c-ip6routing-base-scale-af-xdp

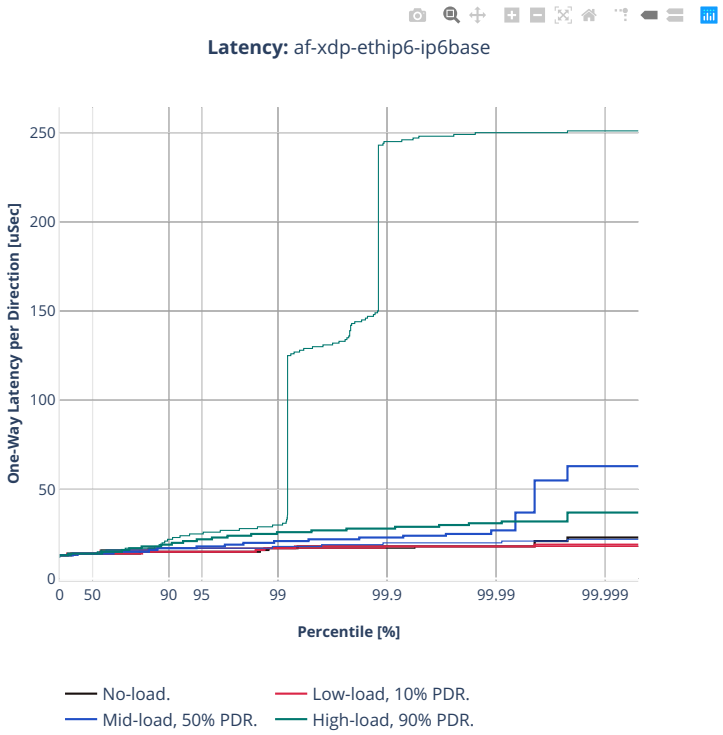

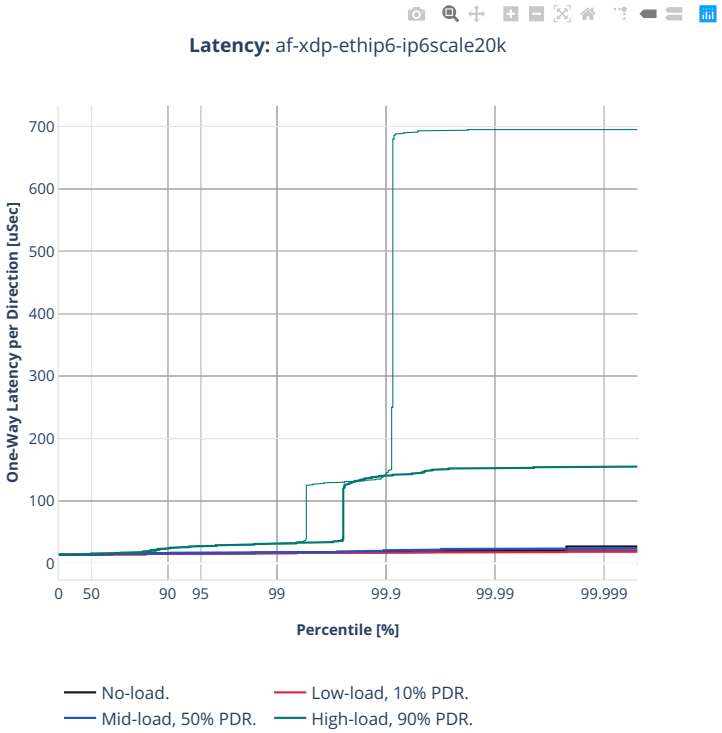
Latency: af-xdp-ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

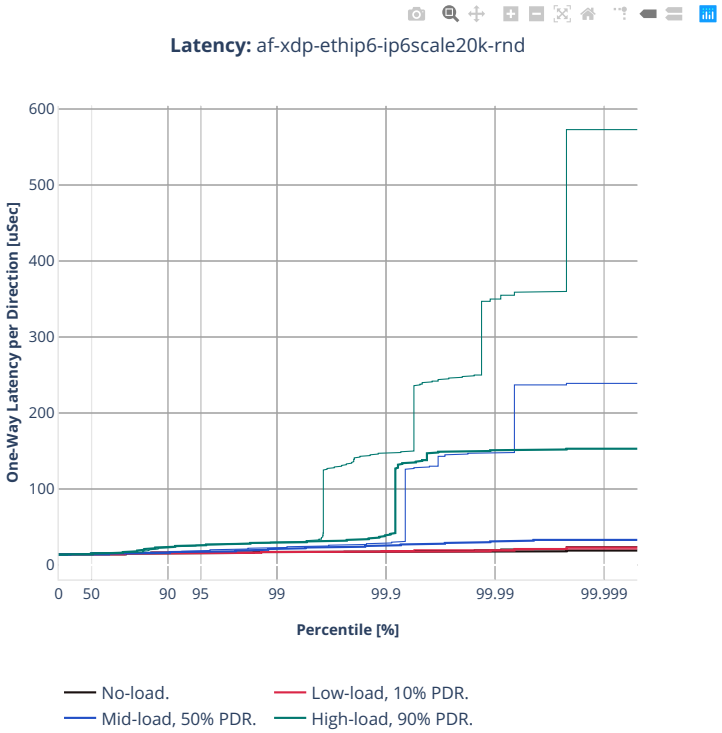
Latency: af-xdp-ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

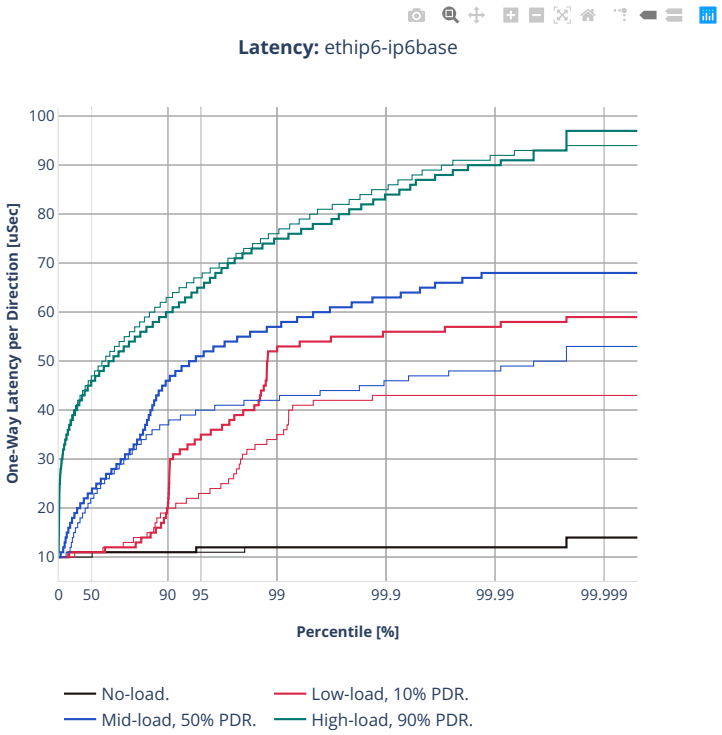

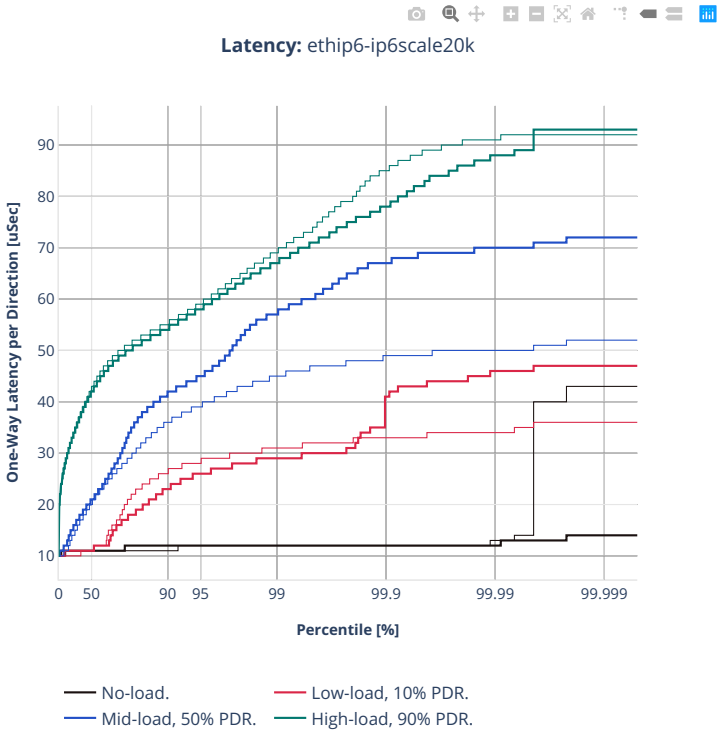
Latency: ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
90
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

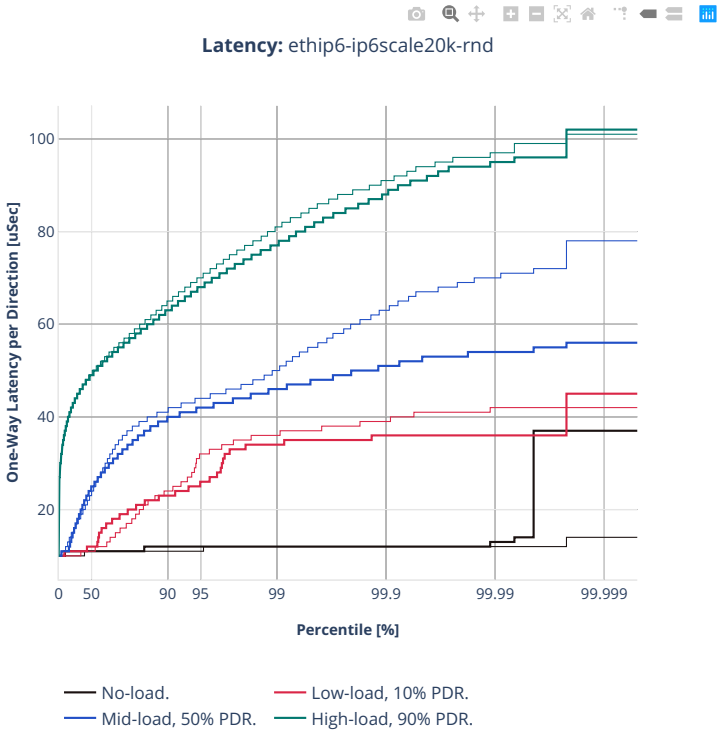
Latency: ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-clx-cx556a

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

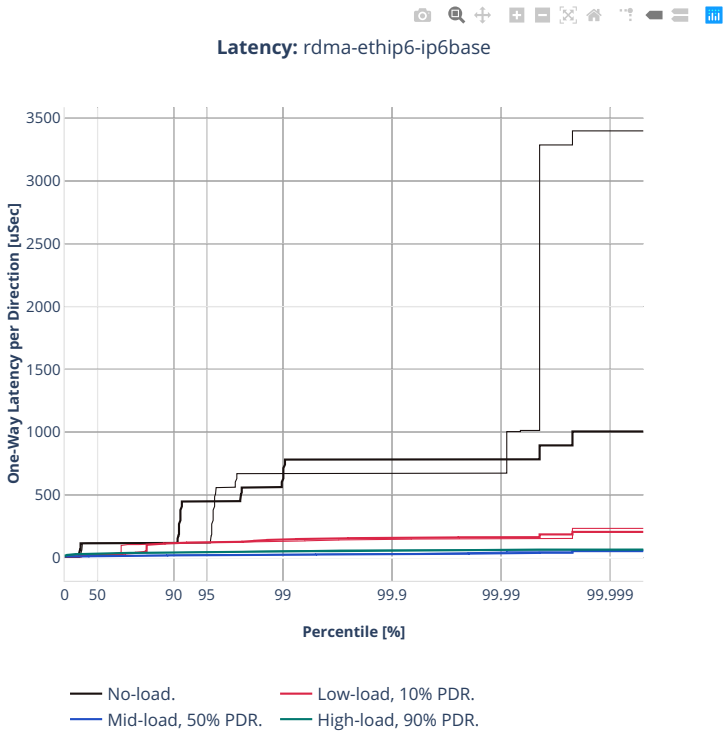

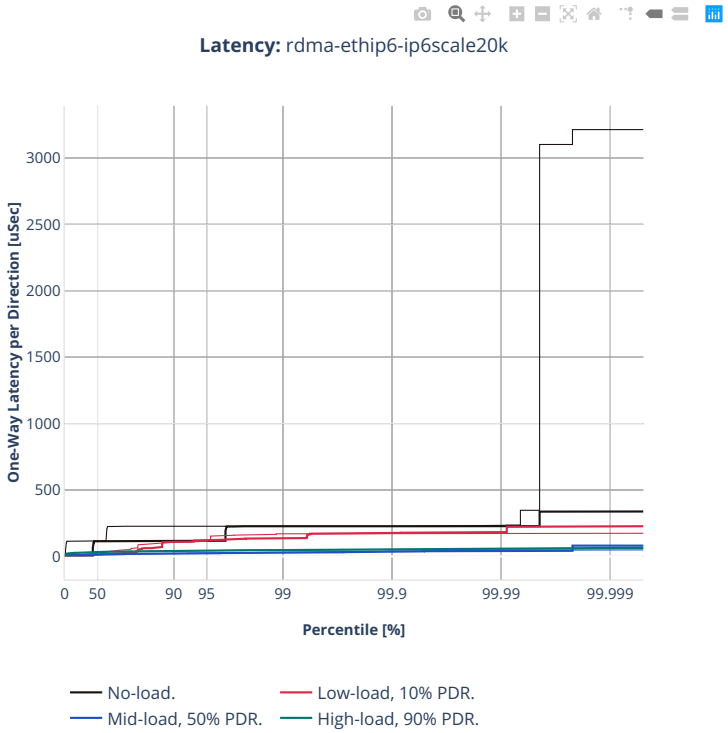
Latency: rdma-ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
3000
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

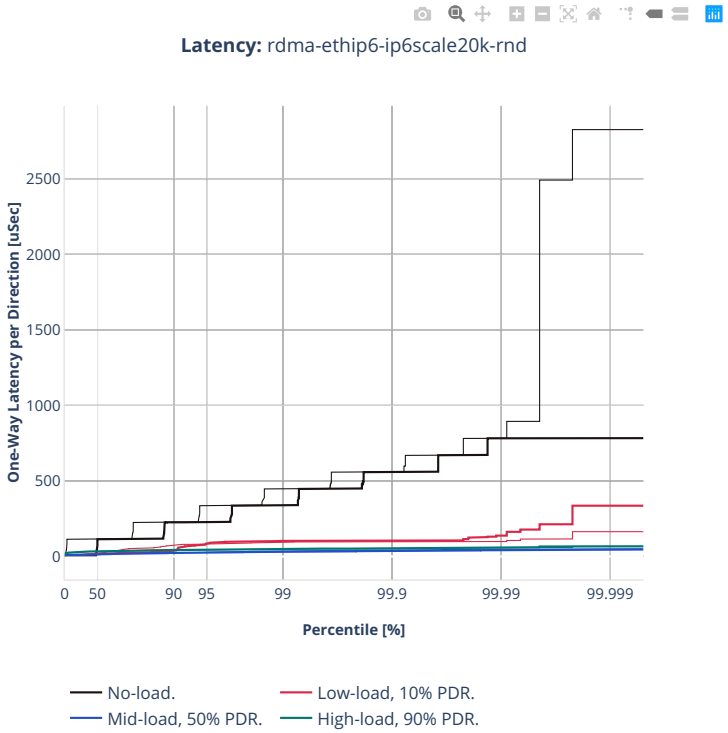
Latency: rdma-ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-clx-e810cq

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

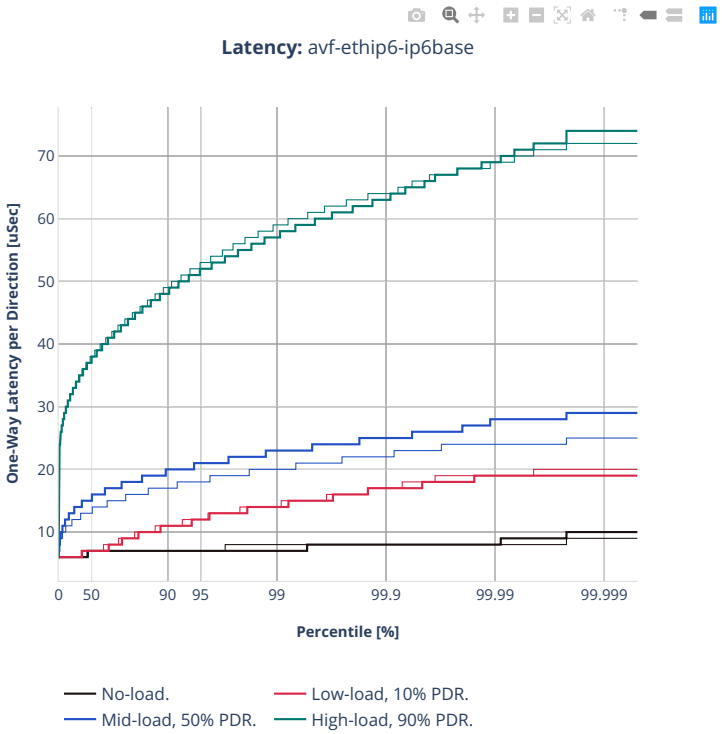

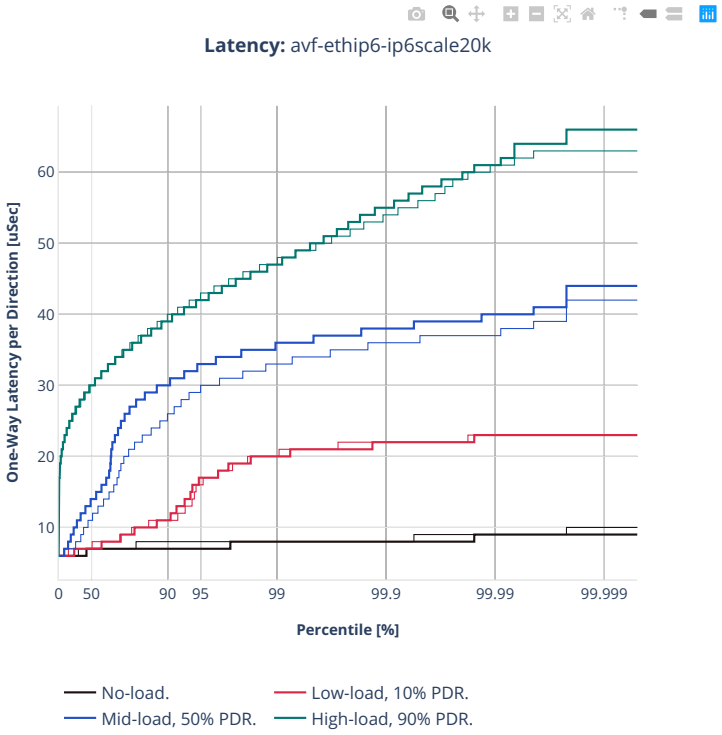
Latency: avf-ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

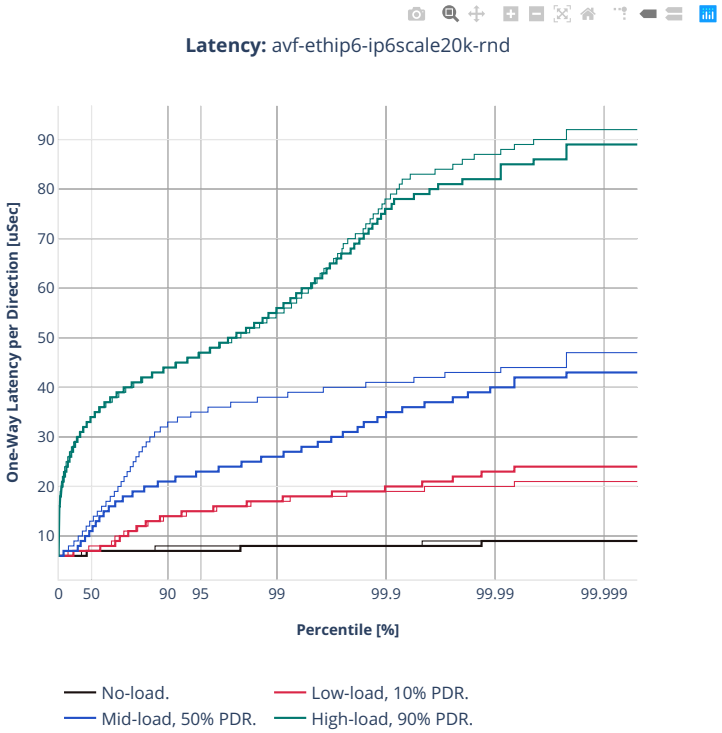
Latency: avf-ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
90
80
70
60
50
40
30
20
10
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

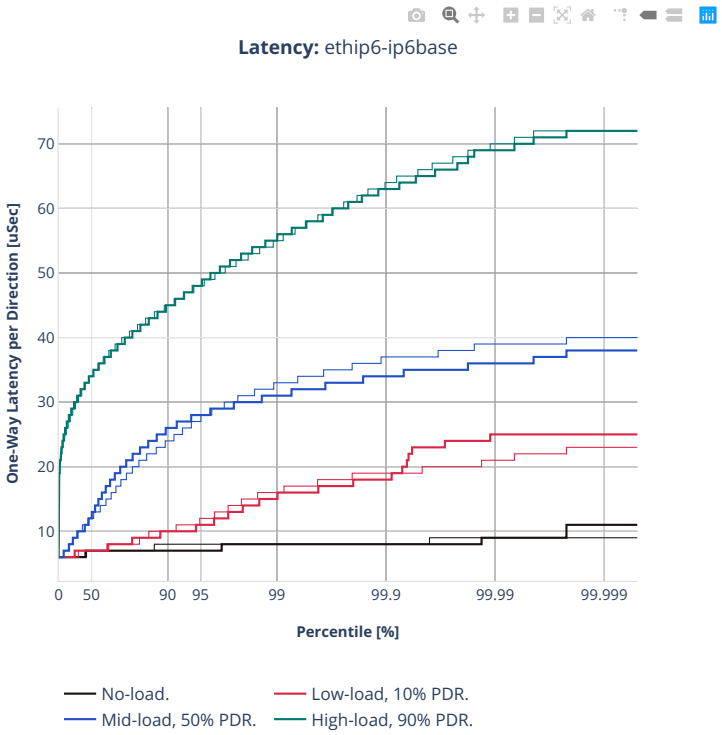

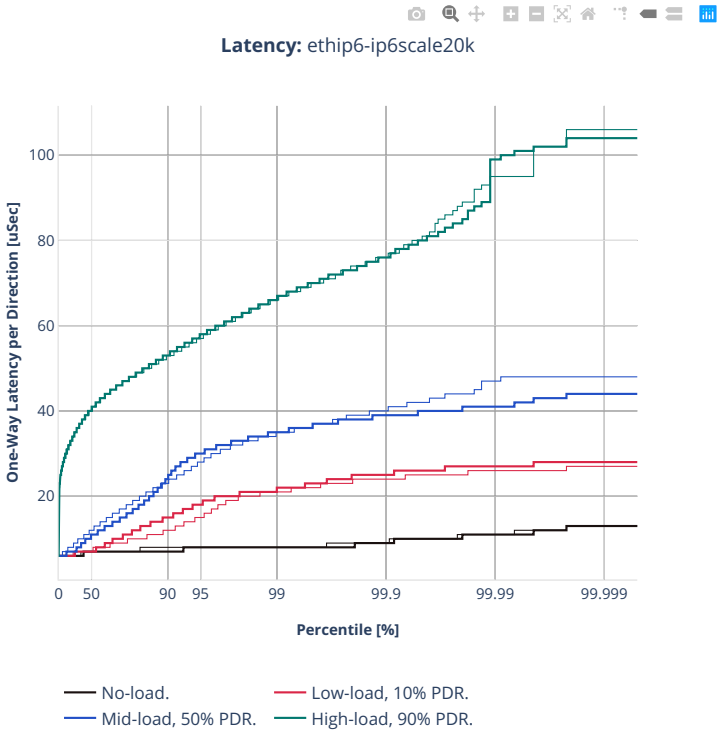
Latency: ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

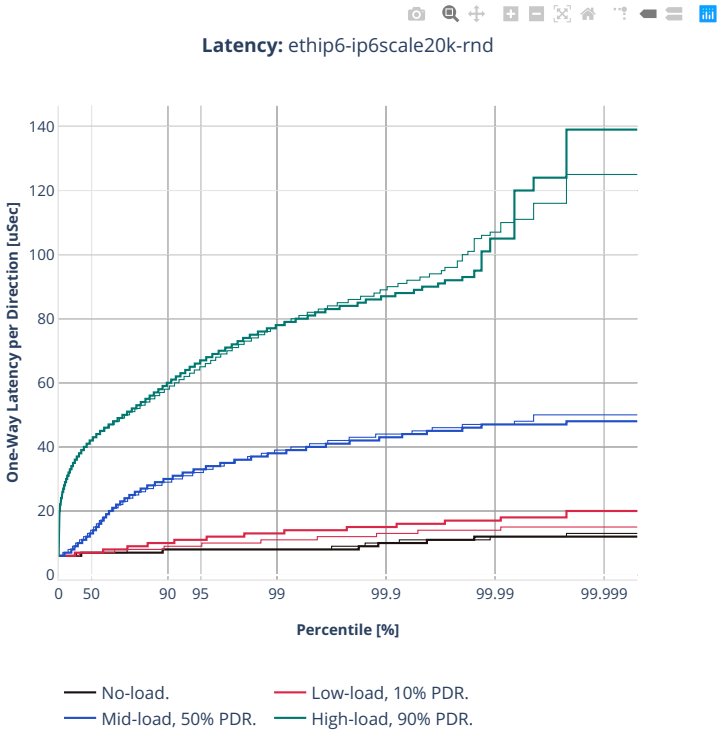
Latency: ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-zn2-xxv710

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

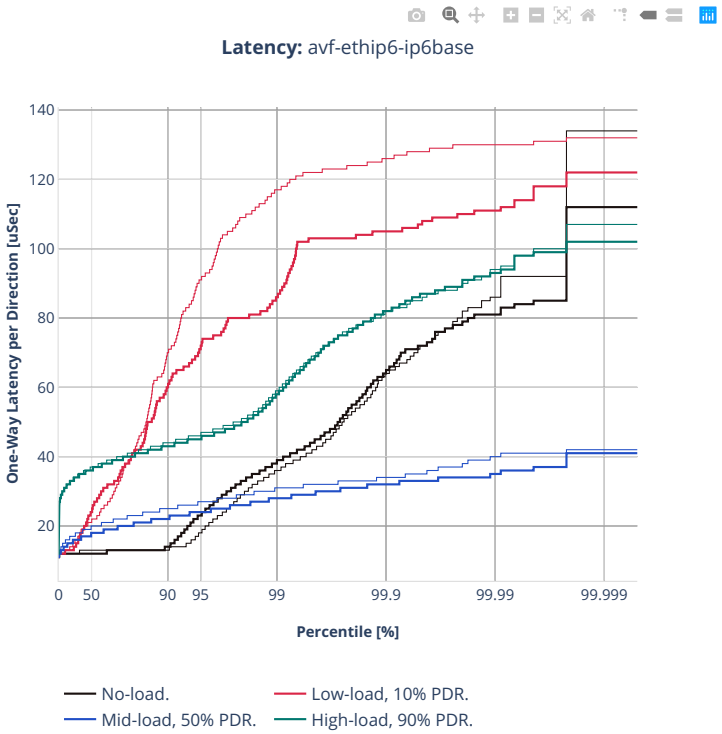

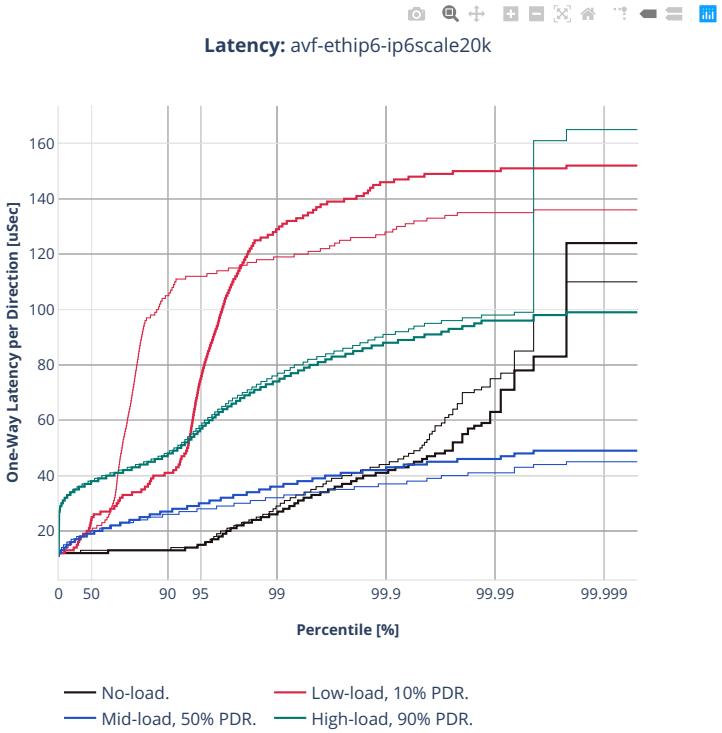
Latency: avf-ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

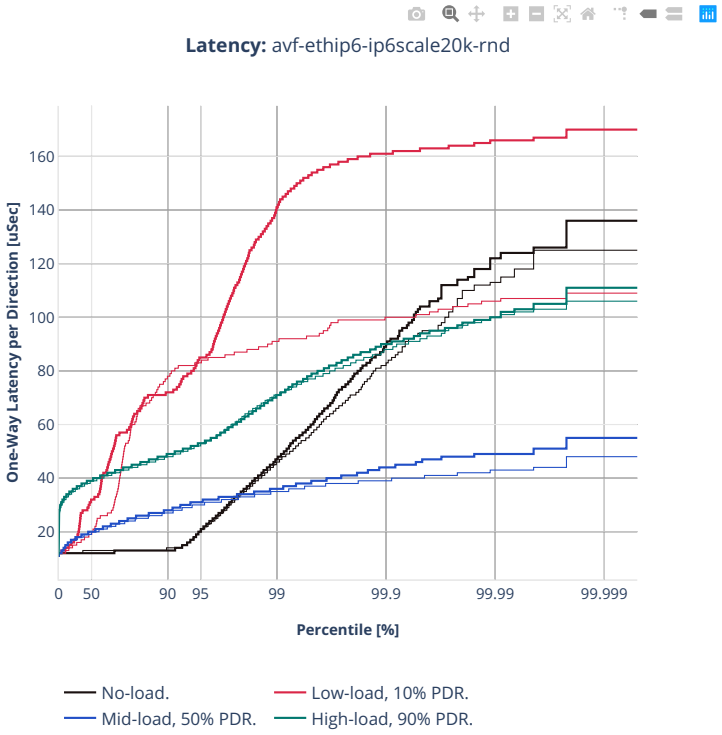
Latency: avf-ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

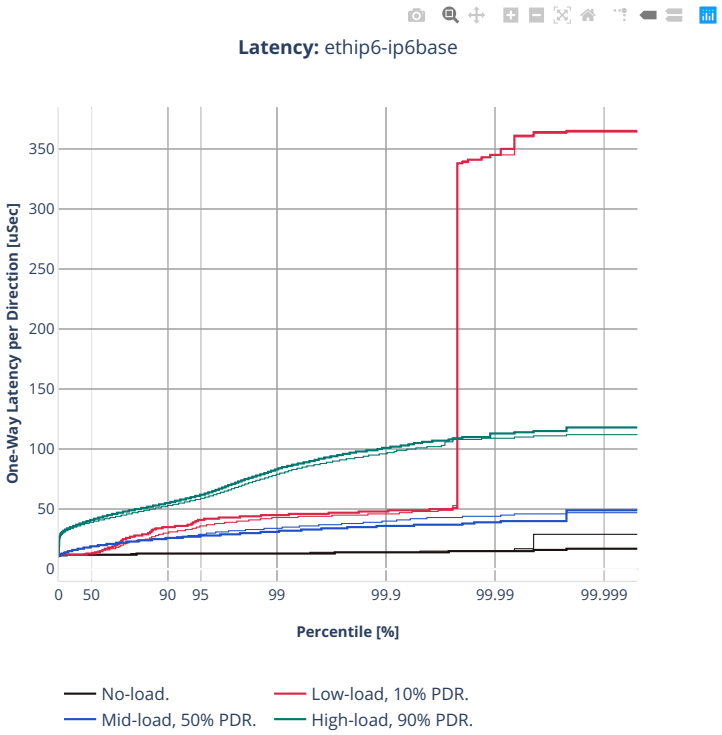

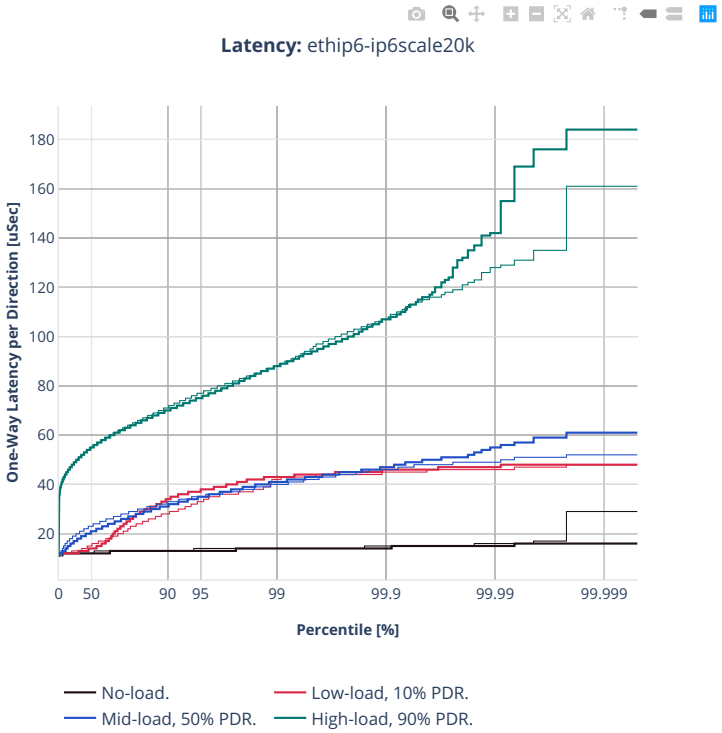
Latency: ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

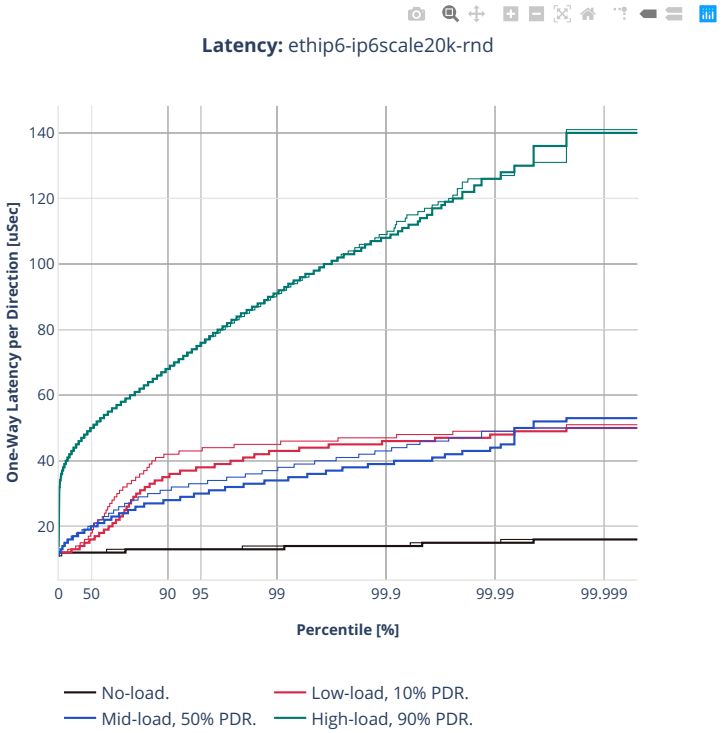
Latency: ethip6-ip6scale20k-rnd
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-alt-xl710

78b-1t1c-ip6routing-base-scale

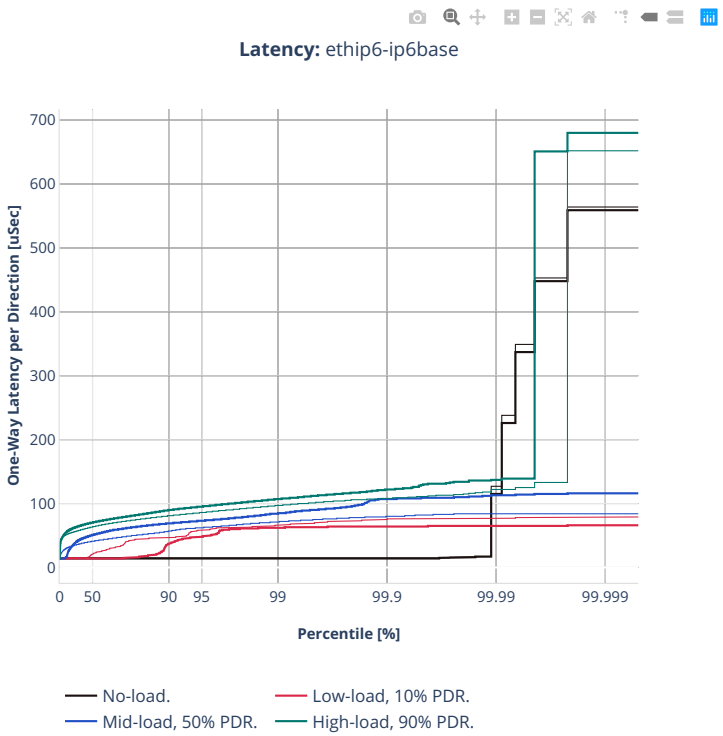

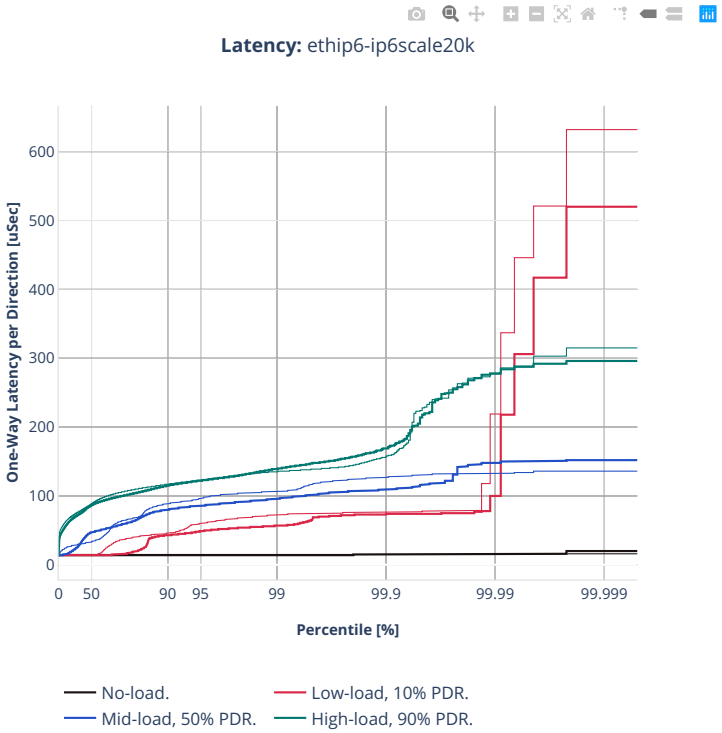
Latency: ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

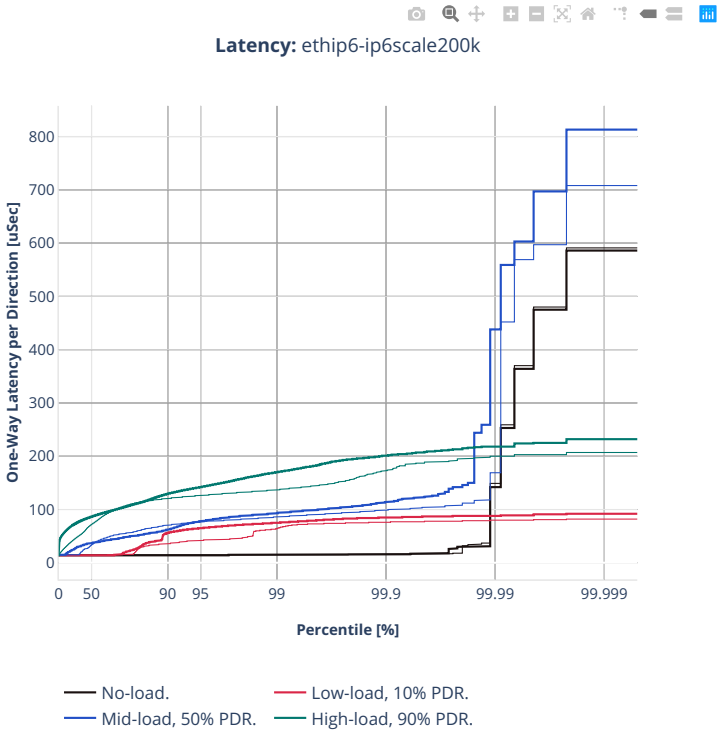
Latency: ethip6-ip6scale200k
One-Way Latency per Direction [uSec]
800
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

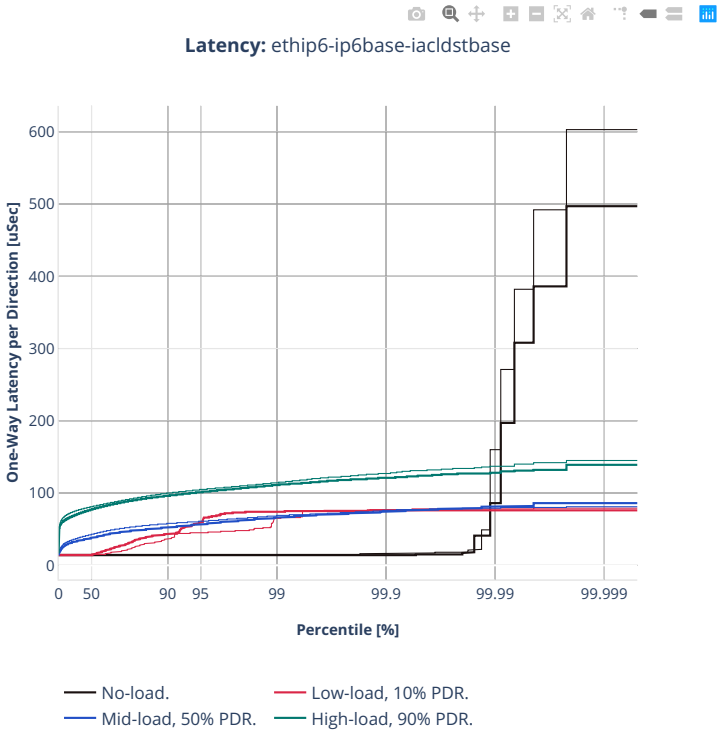
Latency: ethip6-ip6base-iacldstbase
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-tsh-x520

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

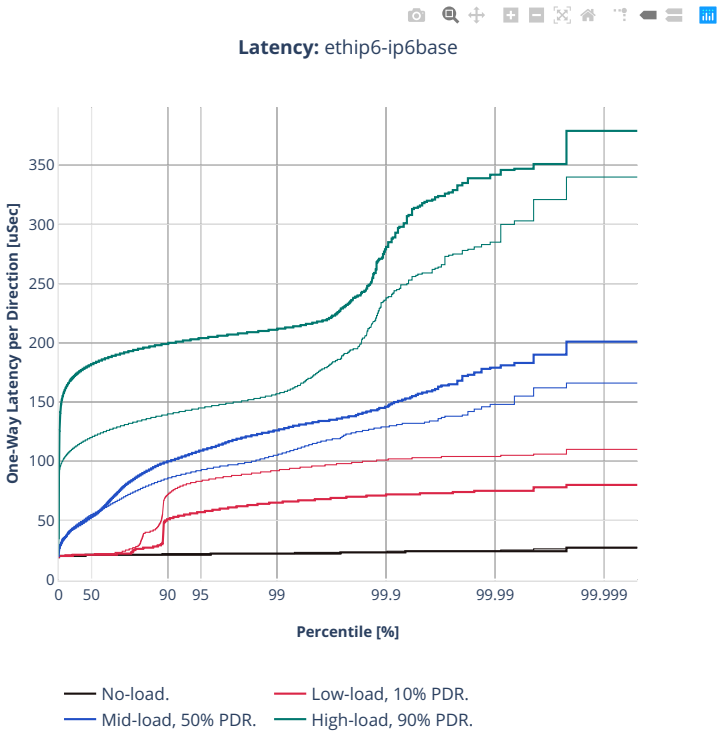

2n-tx2-xl710

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

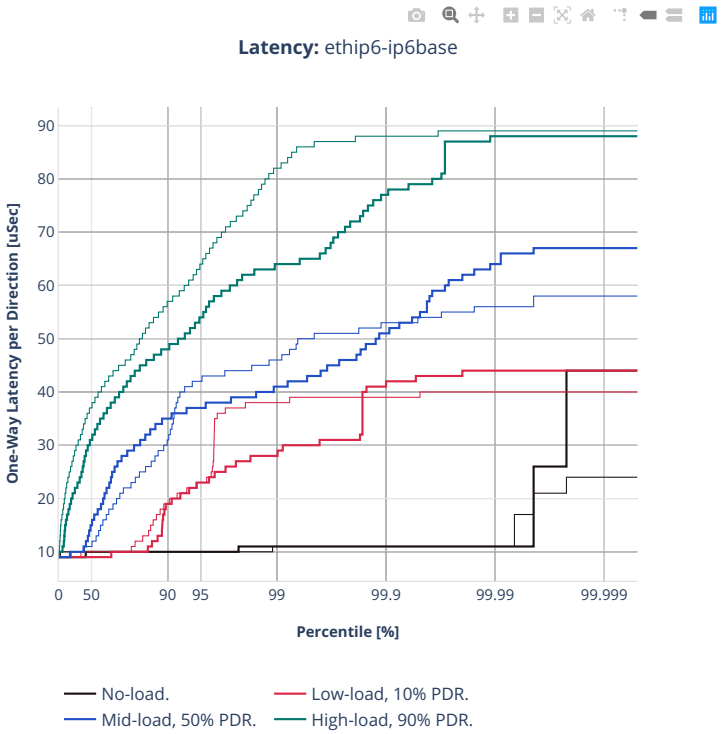

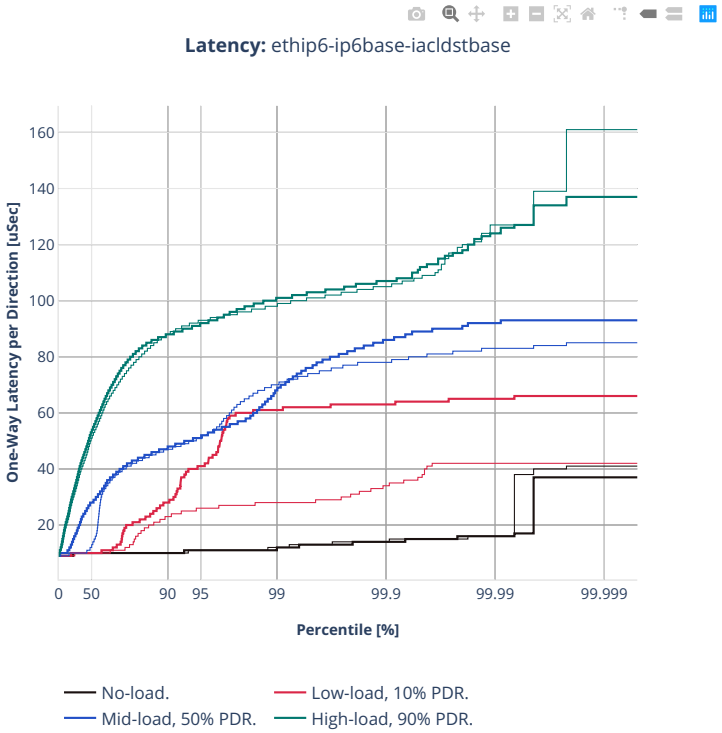
Latency: ethip6-ip6base-iacldstbase
One-Way Latency per Direction [uSec]
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

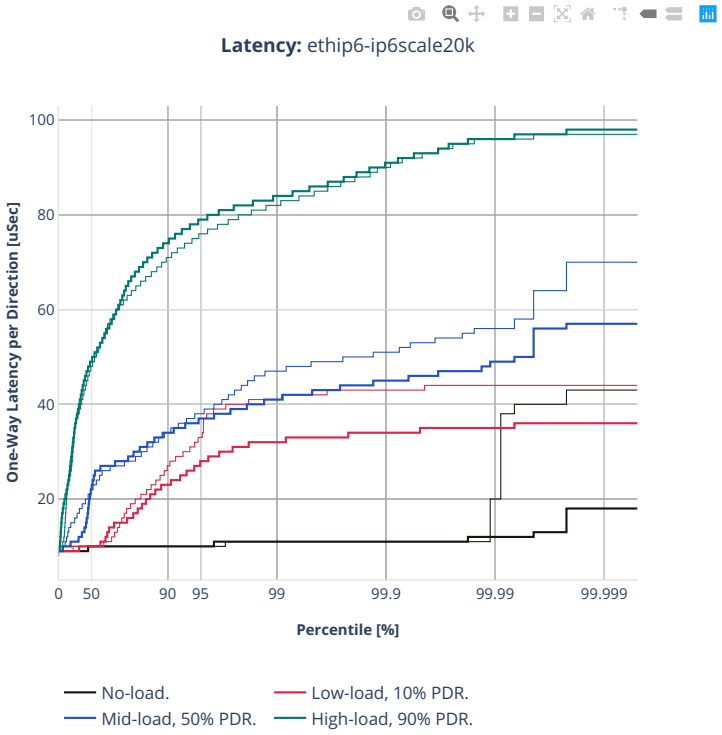
Latency: ethip6-ip6scale20k
One-Way Latency per Direction [uSec]
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.


Latency: ethip6-ip6scale200k
One-Way Latency per Direction [uSec]
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2.5.4 SRv6 Routing

CSIT source code for the test cases used for plots can be found in CSIT git repository[152].

[152] https://git.fd.io/csit/tree/tests/vpp/perf/srv6?h=rls2206

3n-icx-xxv710

78b-2t1c-srv6-ip6routing-base-avf

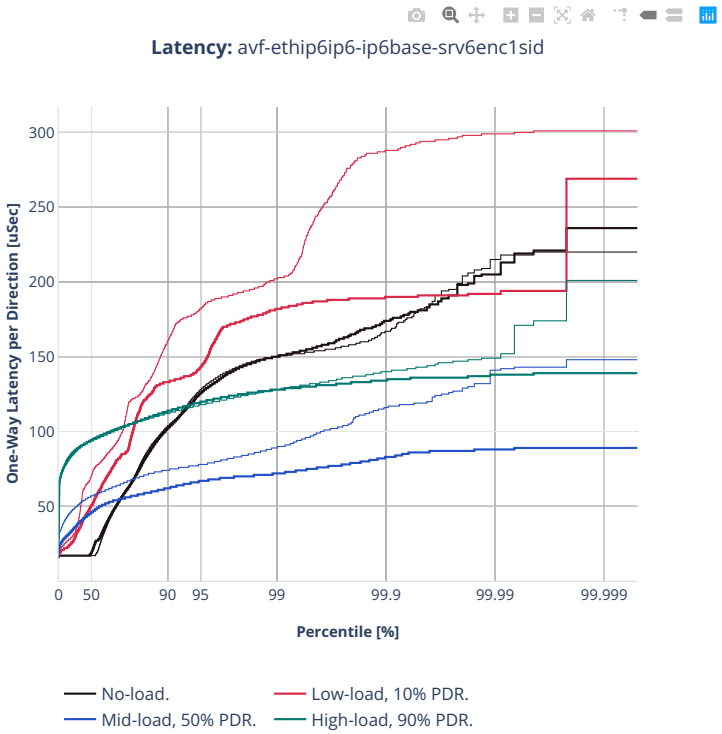

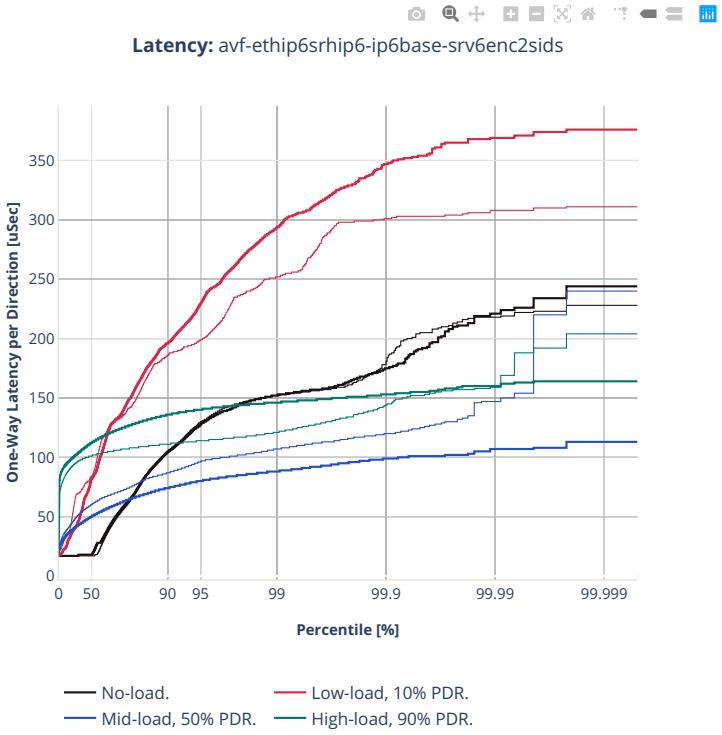
Latency: avf-ethip6srhip6-ip6base-srv6enc2sids
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-skx-xxv710

78b-2t1c-srv6-ip6routing-base-dpdk

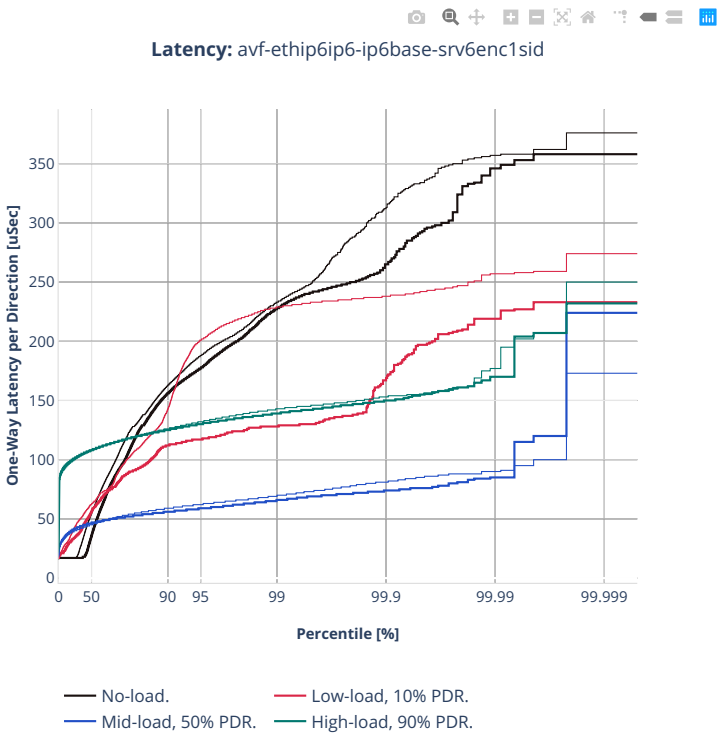

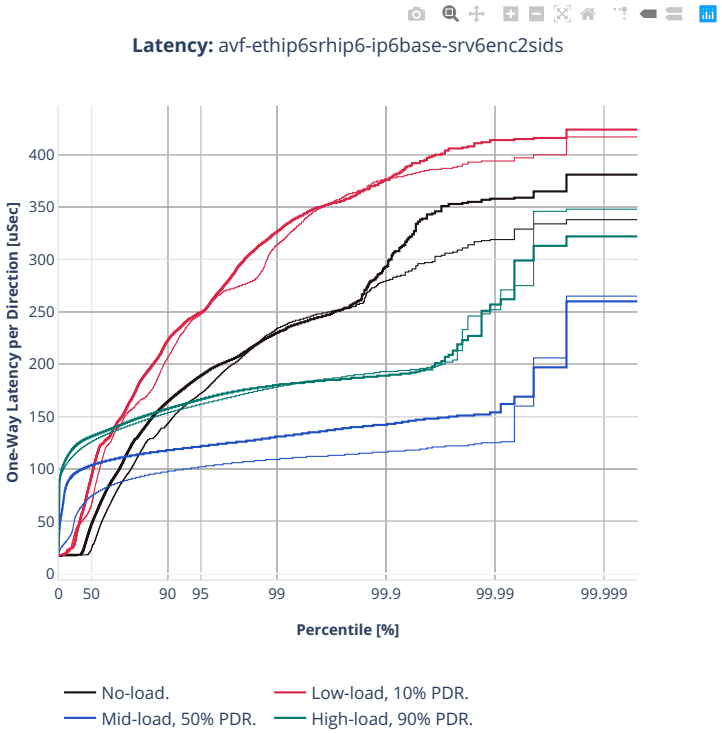
Latency: avf-ethip6srhip6-ip6base-srv6enc2sids
One-Way Latency per Direction [uSec]
400
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-tsh-x520

78b-1t1c-srv6-ip6routing-base-ixgbe

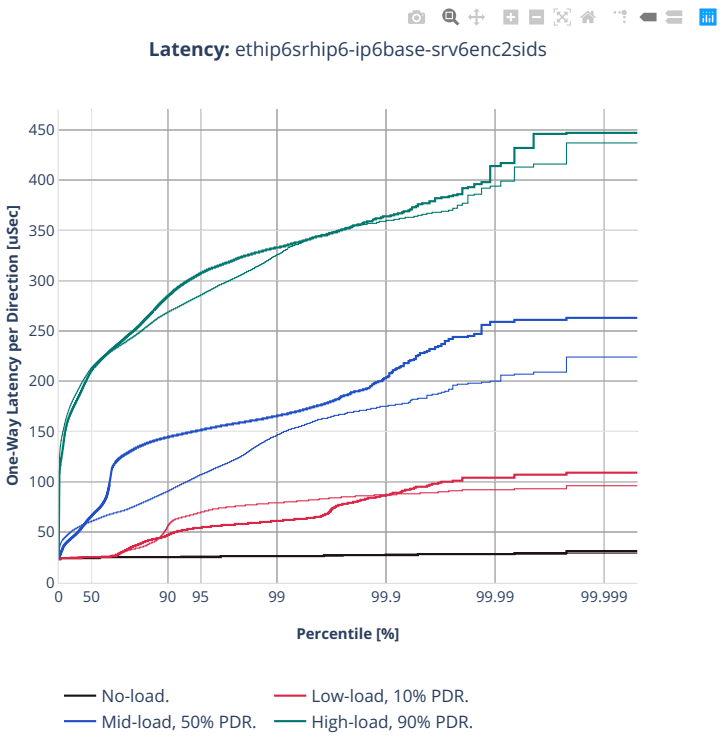

2.5.5 IPv4 Tunnels

CSIT source code for the test cases used for plots can be found in CSIT git repository[153].

[153] https://git.fd.io/csit/tree/tests/vpp/perf/ip4_tunnels?h=rls2206

3n-icx-xxv710

64b-2t1c-ip4tunnel-base-avf

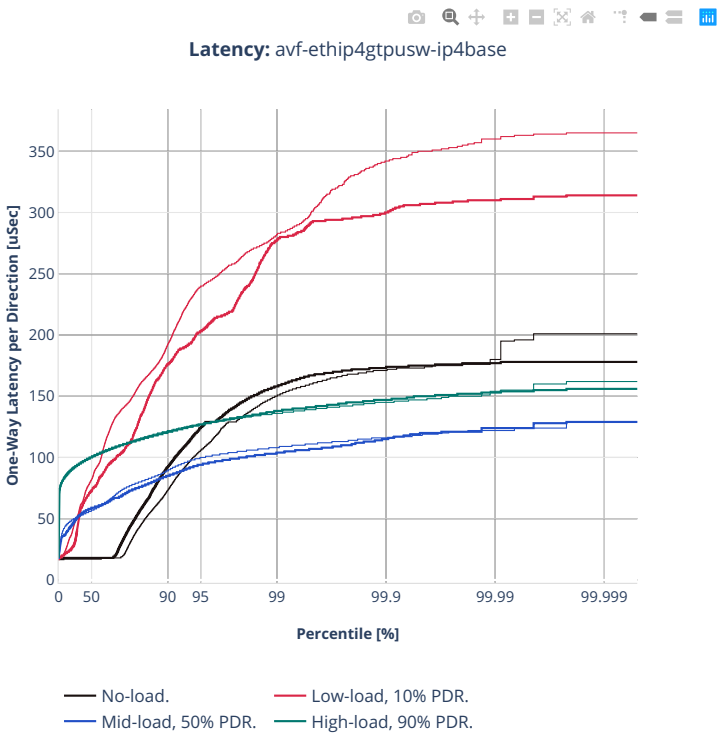

Latency: avf-ethip4vxlan-l2bdbasemaclrn

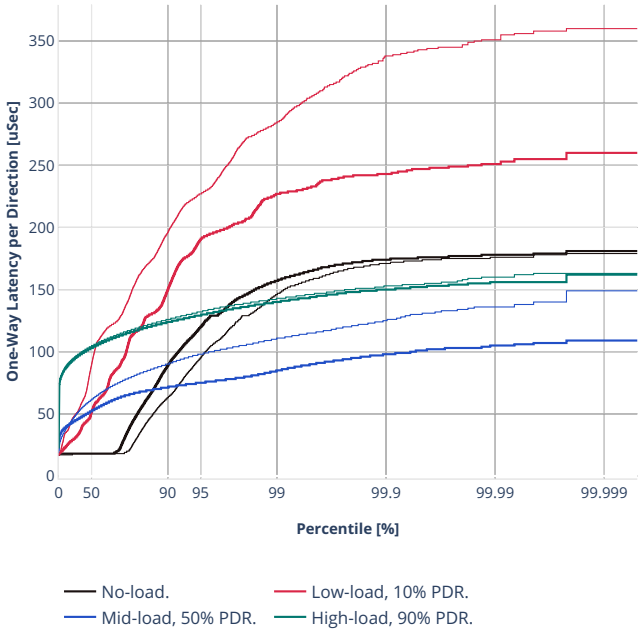

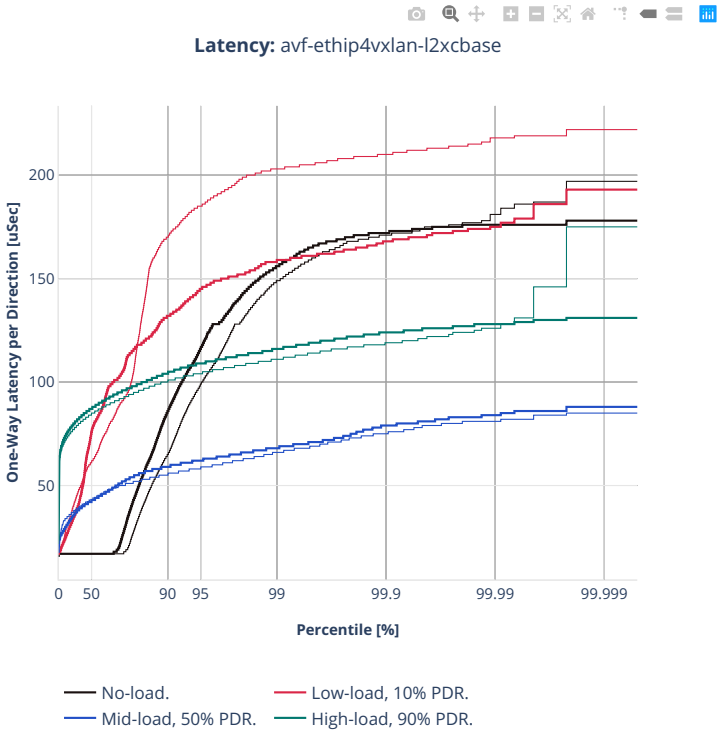
Latency: avf-ethip4vxlan-l2xcbase
One-Way Latency per Direction [uSec]
200
150
100
50
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-ip4tunnel-base-dpdk

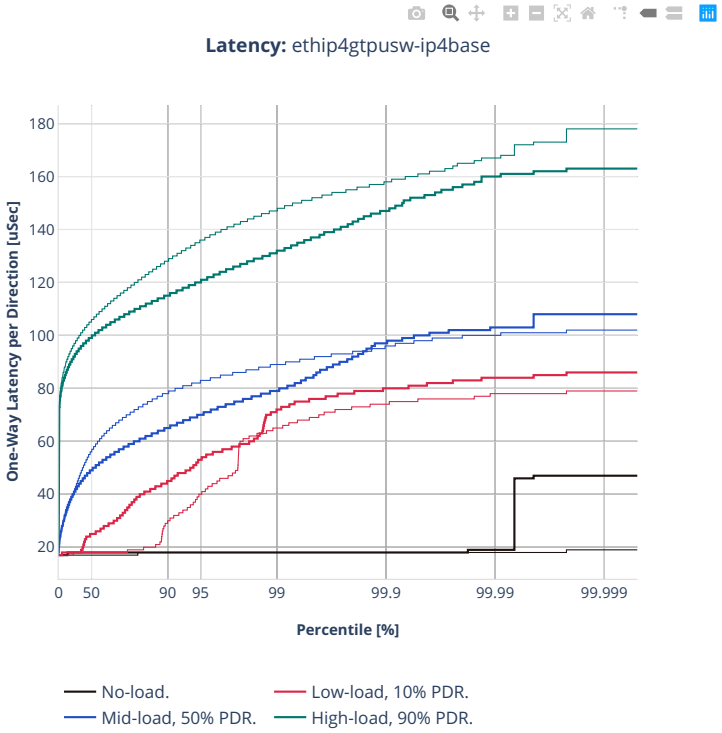

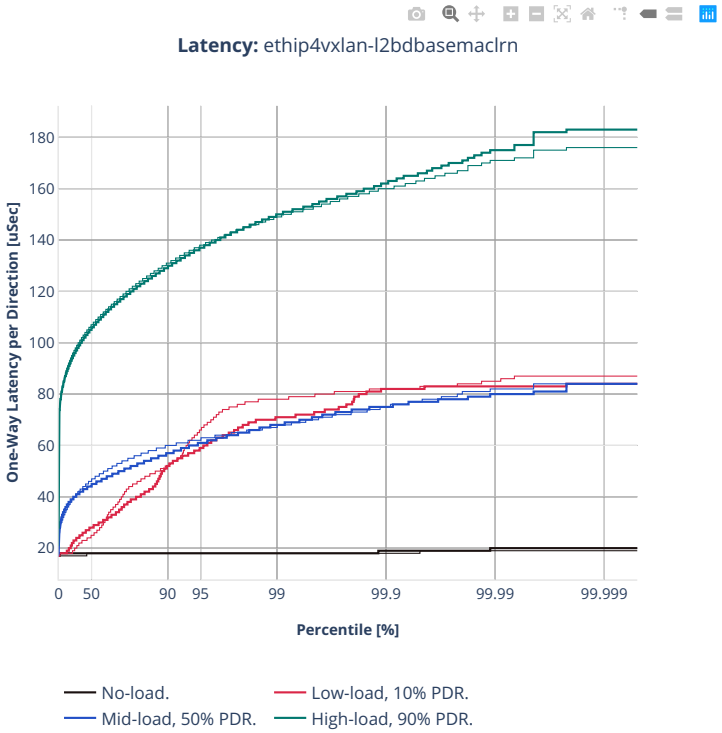
Latency: ethip4vxlan-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

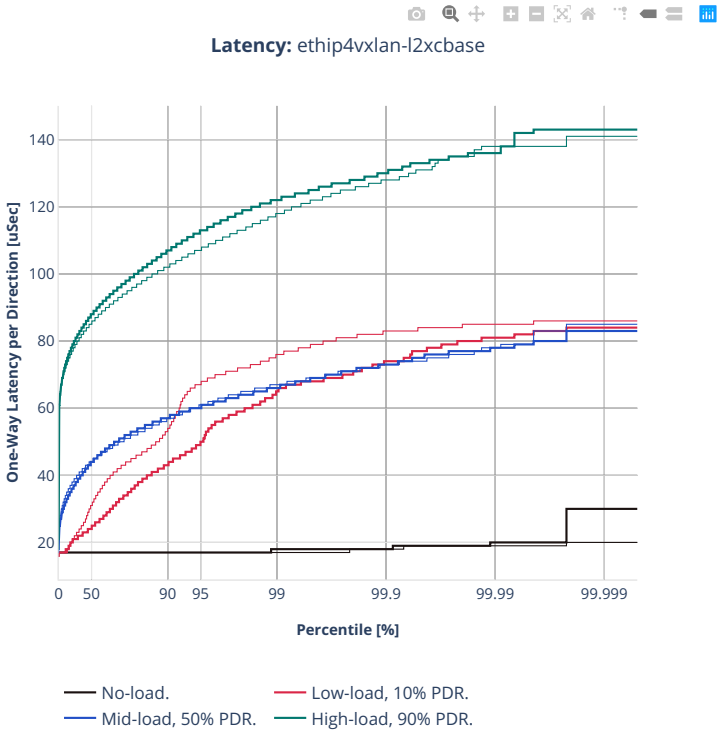
Latency: ethip4vxlan-l2xcbase
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-skx-xxv710

64b-2t1c-ip4tunnel-base-avf

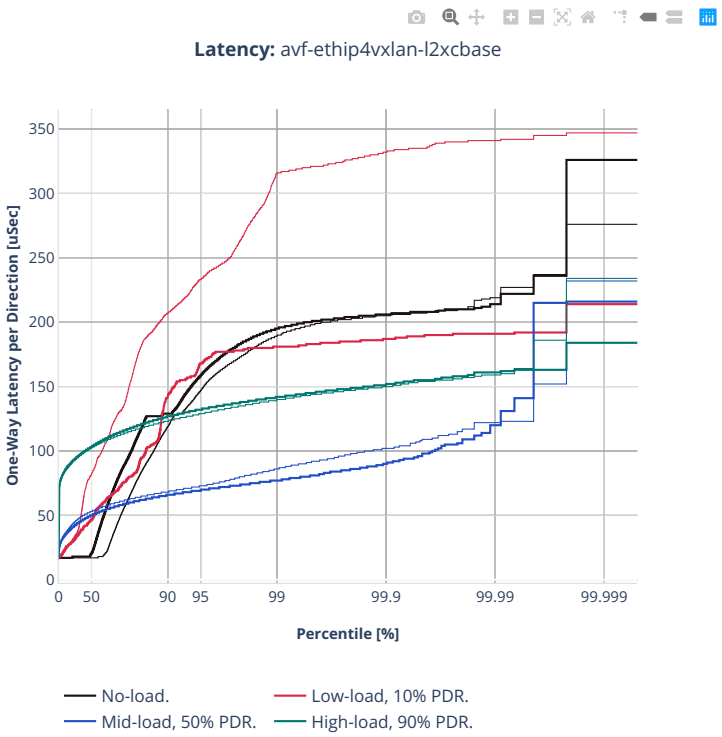

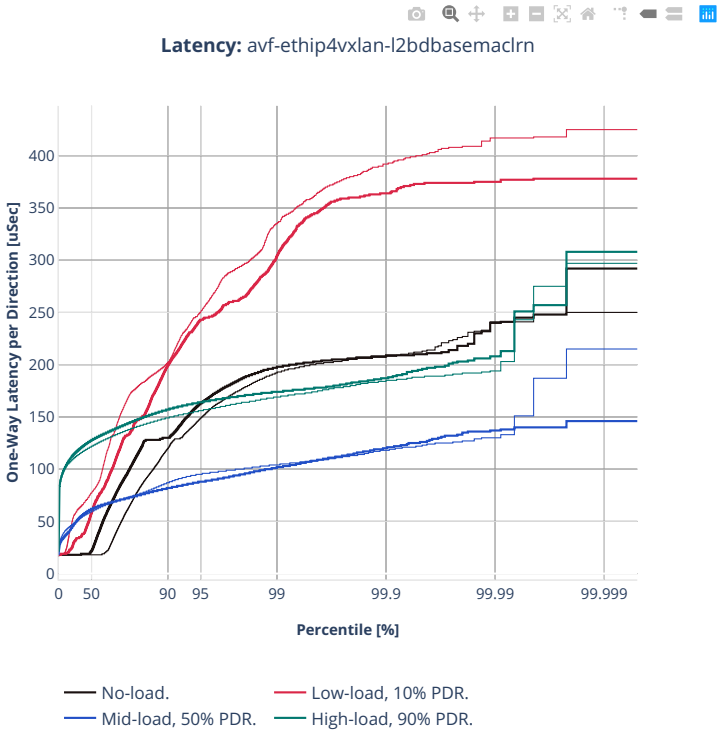
Latency: avf-ethip4vxlan-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
400
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-ip4tunnel-base-dpdk

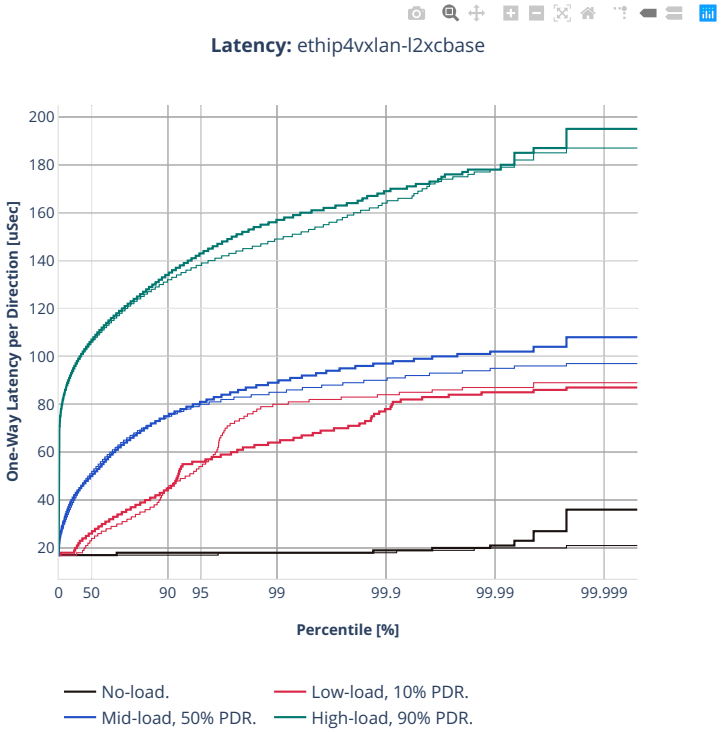

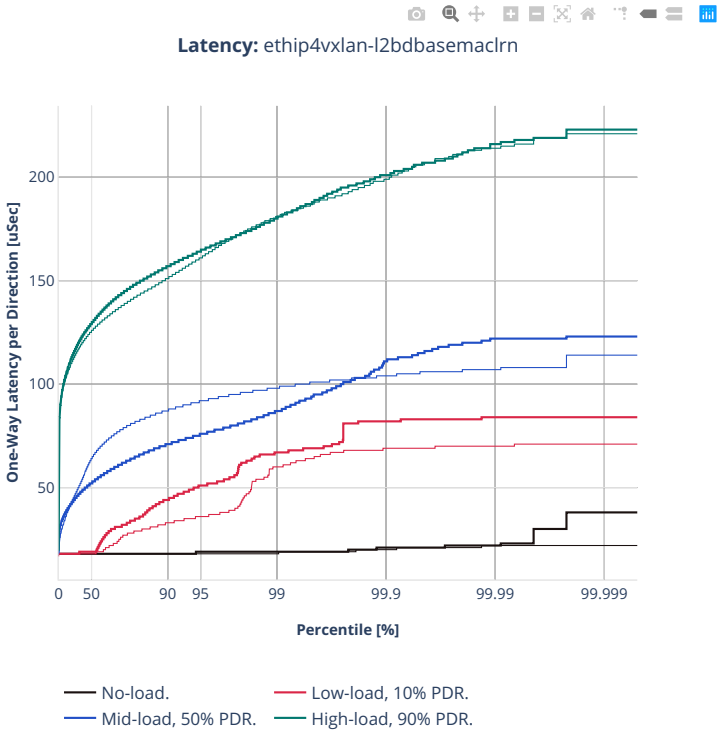
Latency: ethip4vxlan-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
200
150
100
50
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-alt-xl710

64b-1t1c-ip4tunnel-base

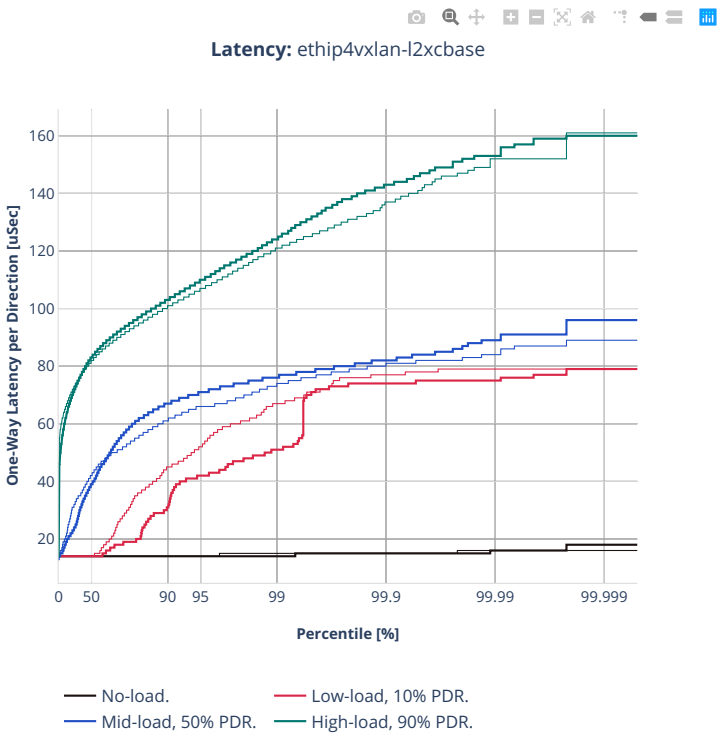

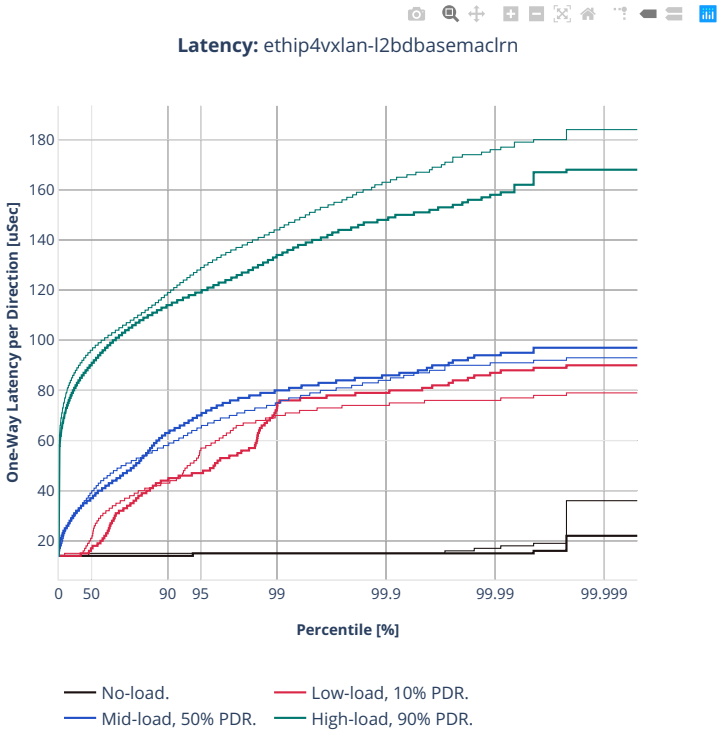
Latency: ethip4vxlan-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-tsh-x520

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

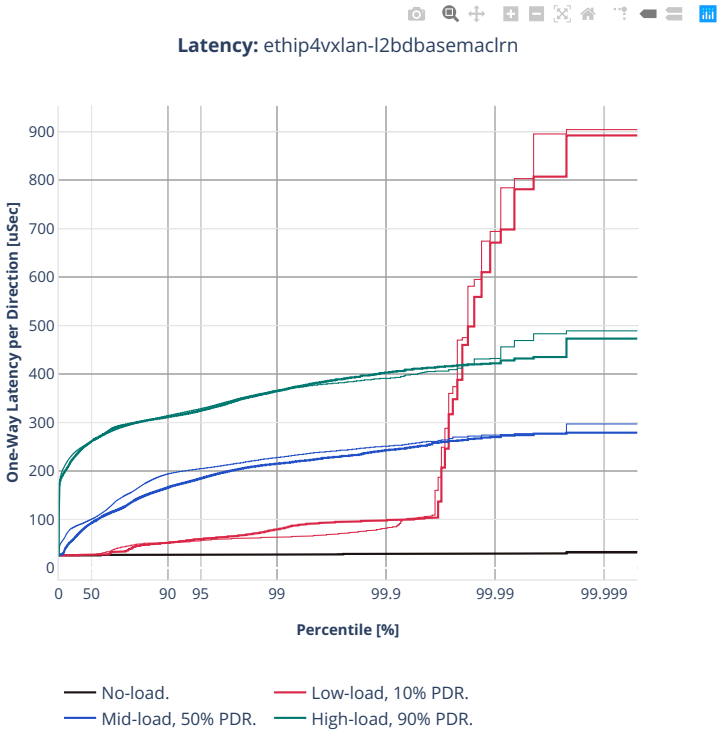
Latency: ethip4vxlan-l2bdbasemaclrn
One-Way Latency per Direction [uSec]
900
800
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2.5.6 NAT44 IPv4 Routing

CSIT source code for the test cases used for plots can be found in CSIT git repository[154].

[154] https://git.fd.io/csit/tree/tests/vpp/perf/ip4?h=rls2206

2n-icx-xxv710

64b-2t1c-ethip4udp-nat44det-avf

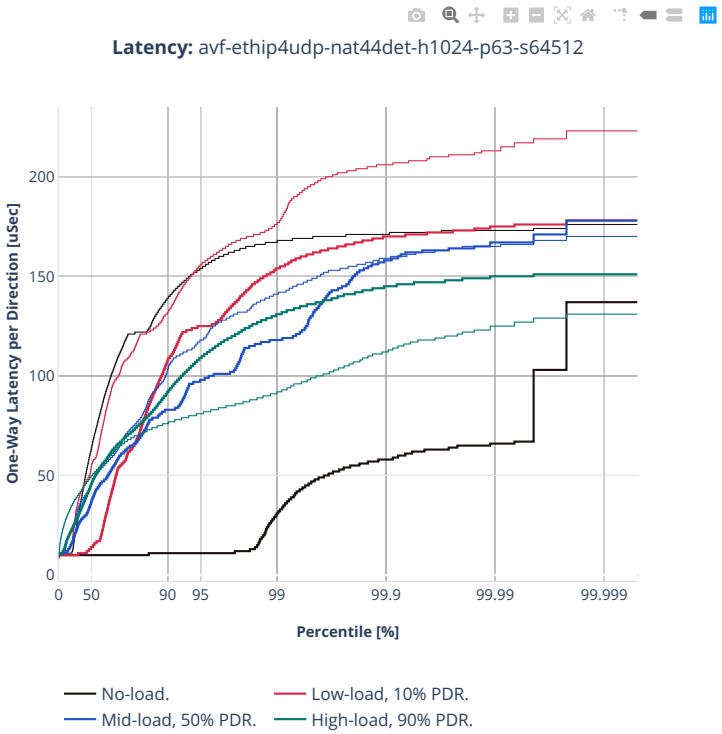

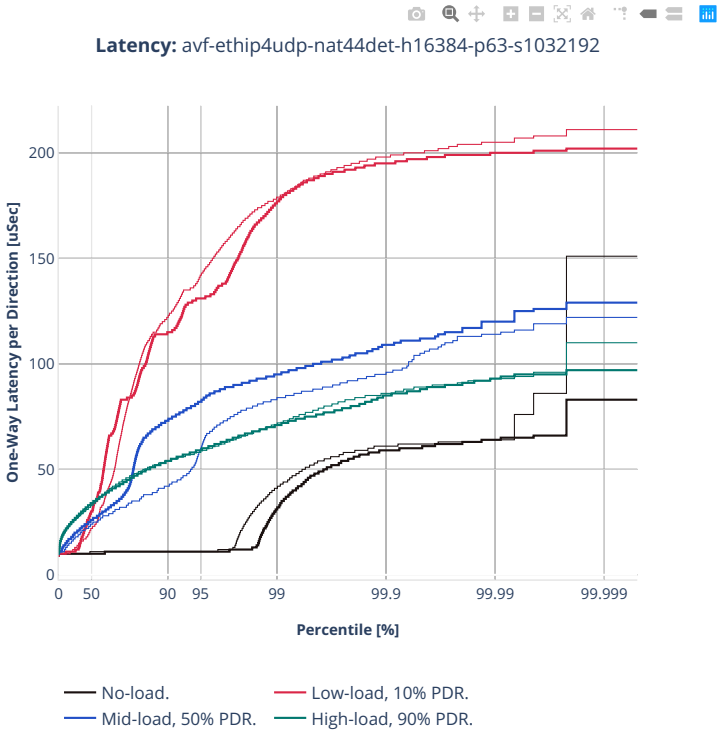
Latency: avf-ethip4udp-nat44det-h16384-p63-s1032192
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

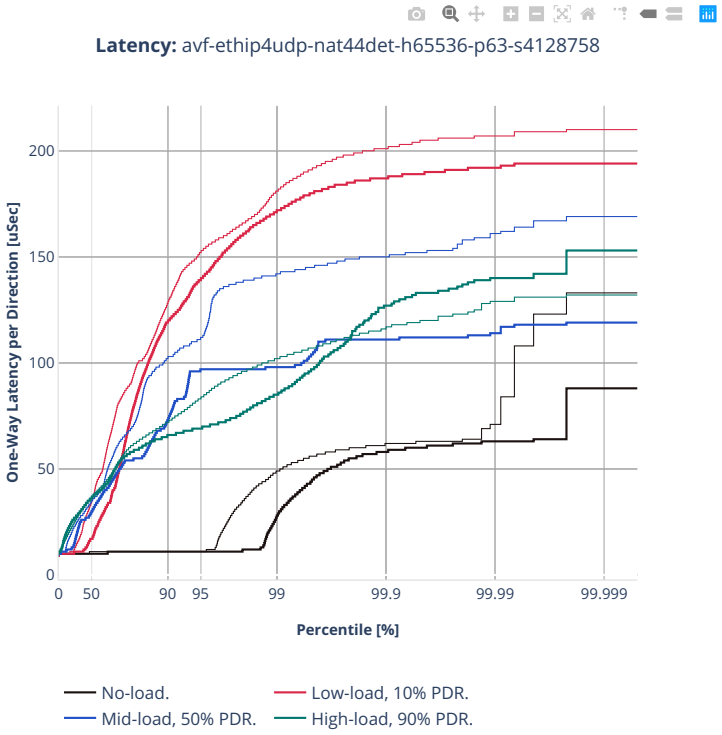
Latency: avf-ethip4udp-nat44det-h65536-p63-s4128758
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-ethip4udp-nat44ed-udir-avf

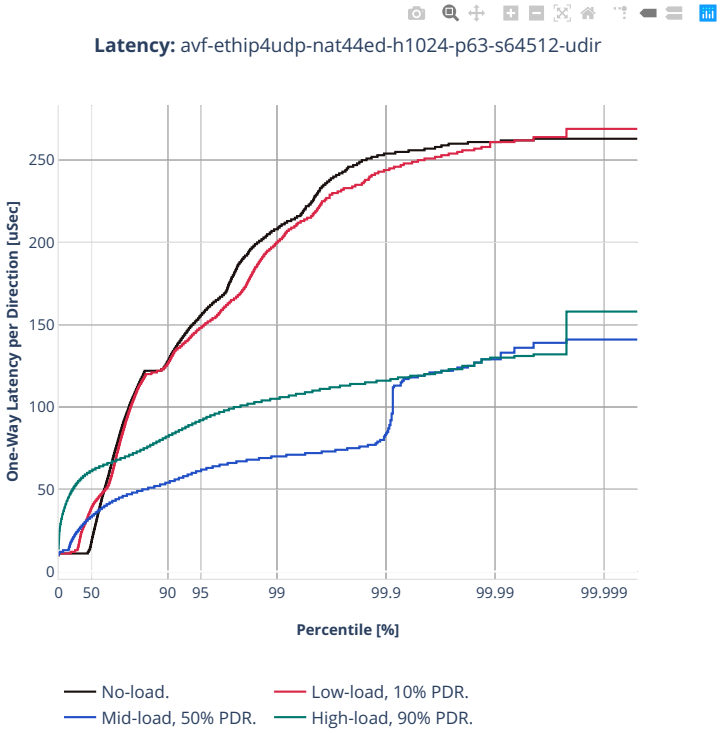

2n-skx-xxv710

64b-2t1c-ethip4udp-nat44det-avf

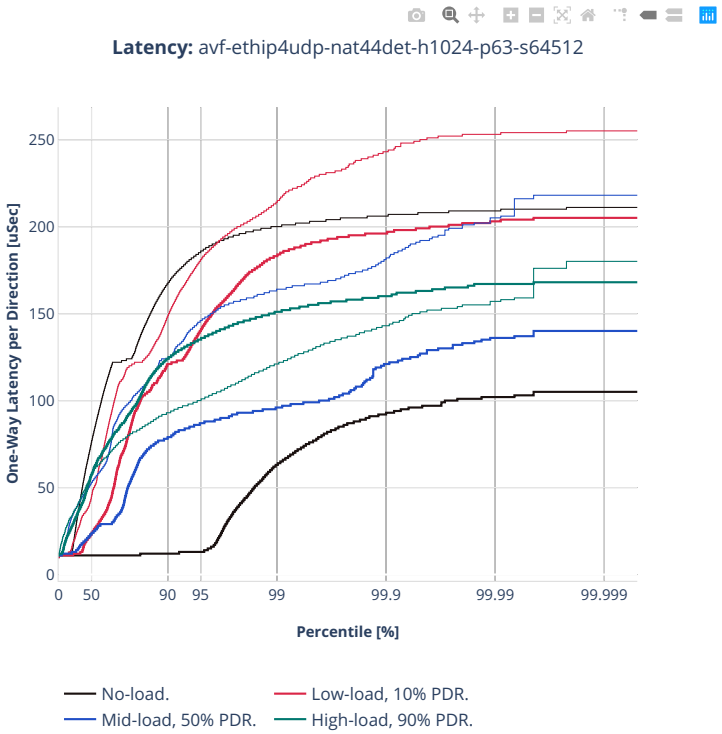

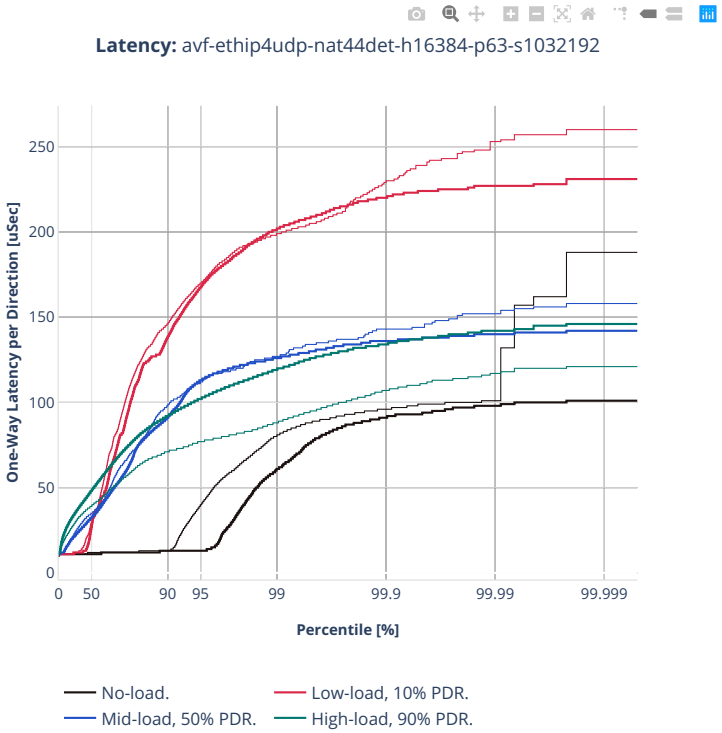
Latency: avf-ethip4udp-nat44det-h16384-p63-s1032192
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

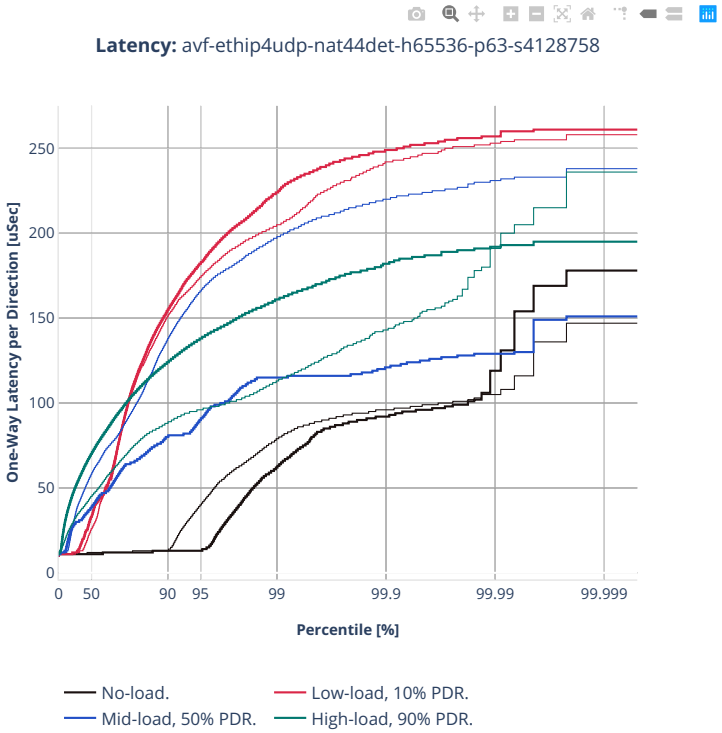
Latency: avf-ethip4udp-nat44det-h65536-p63-s4128758
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-ethip4udp-nat44ed-avf

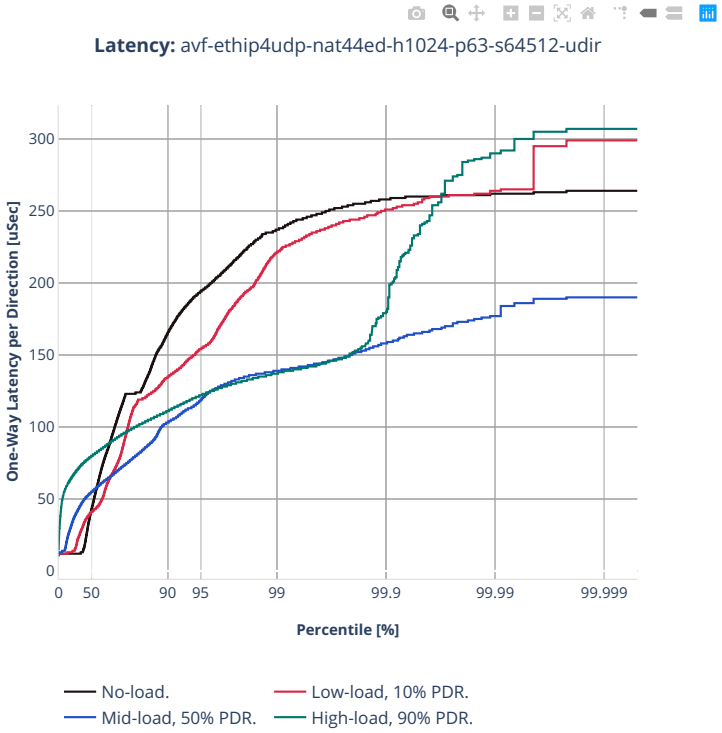

2n-clx-xxv710

64b-2t1c-ethip4udp-nat44det-avf

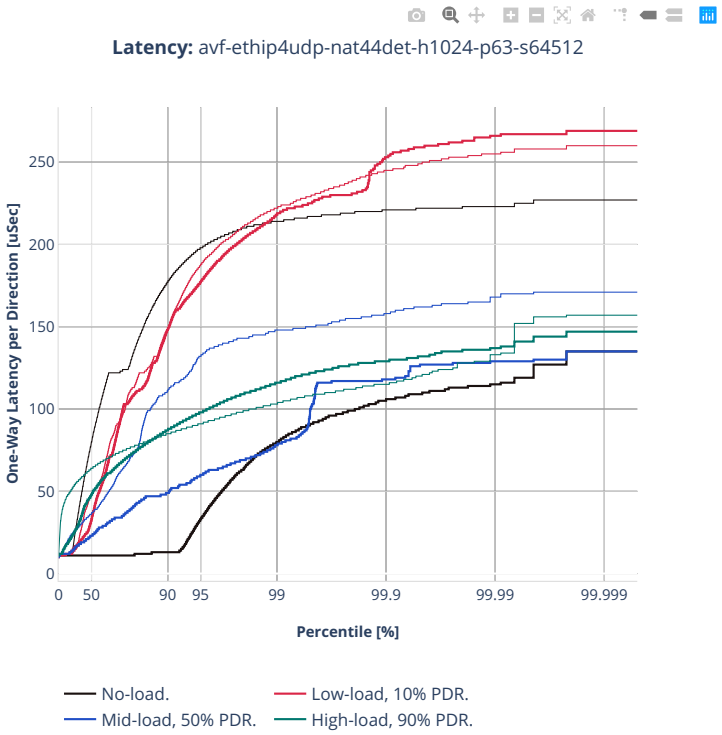

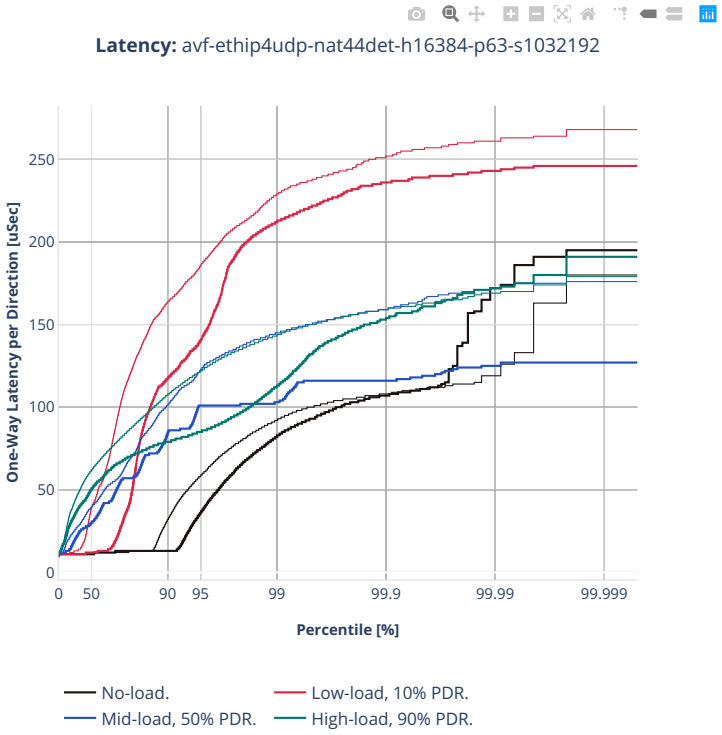
Latency: avf-ethip4udp-nat44det-h16384-p63-s1032192
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

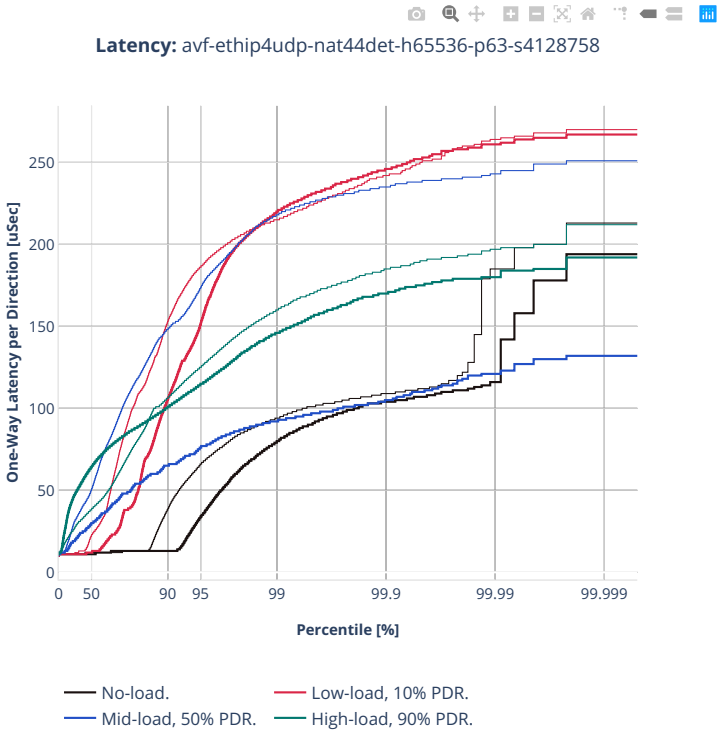
Latency: avf-ethip4udp-nat44det-h65536-p63-s4128758
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-ethip4udp-nat44ed-avf

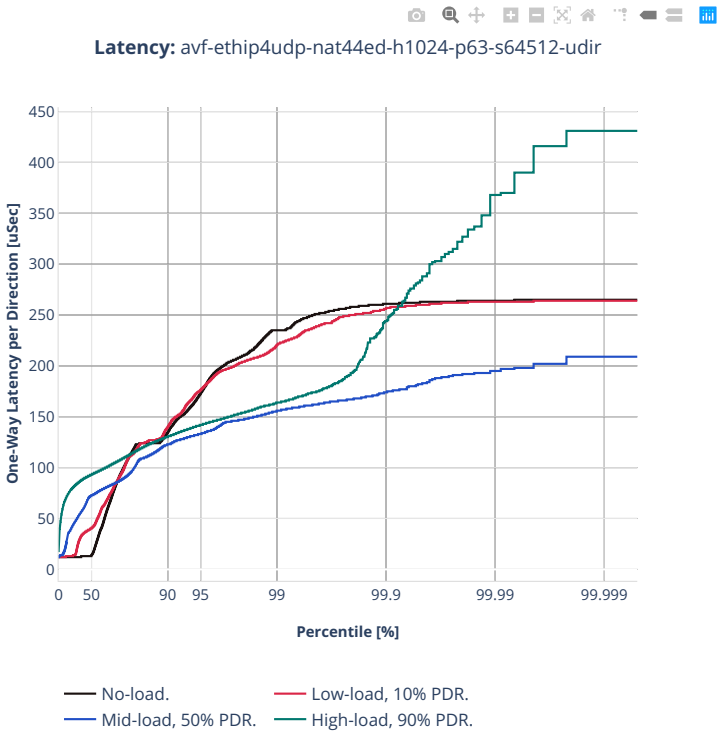

2n-zn2-xxv710

64b-2t1c-ethip4udp-nat44det-avf

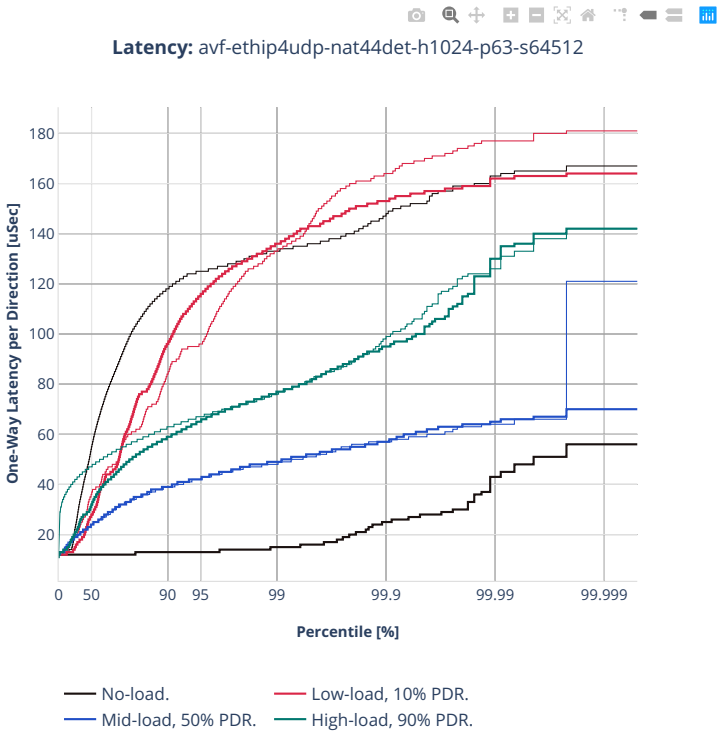

2.5.7 KVM VMs vhost-user

CSIT source code for the test cases used for plots can be found in CSIT git repository[155].

[155] https://git.fd.io/csit/tree/tests/vpp/perf/vm_vhost?h=rls2206

2n-icx-xxv710

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

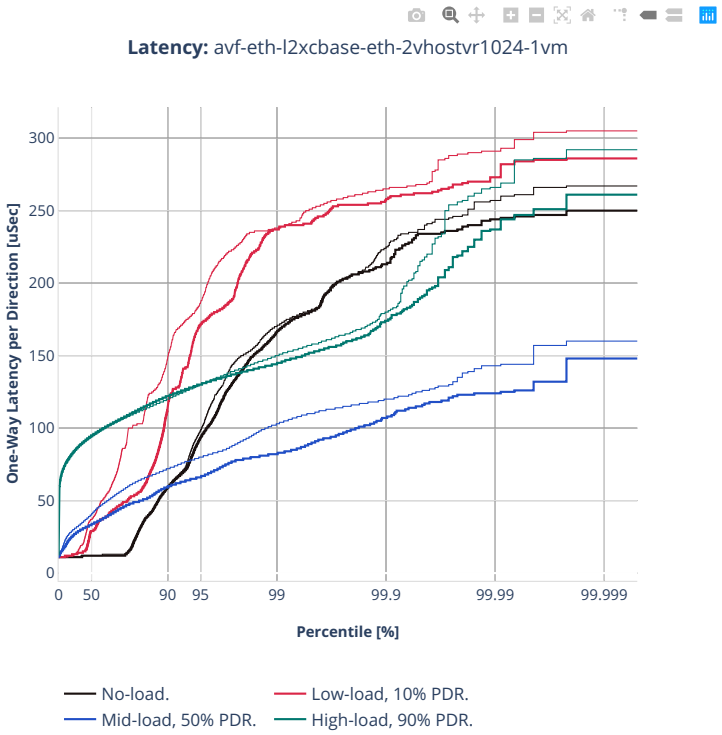

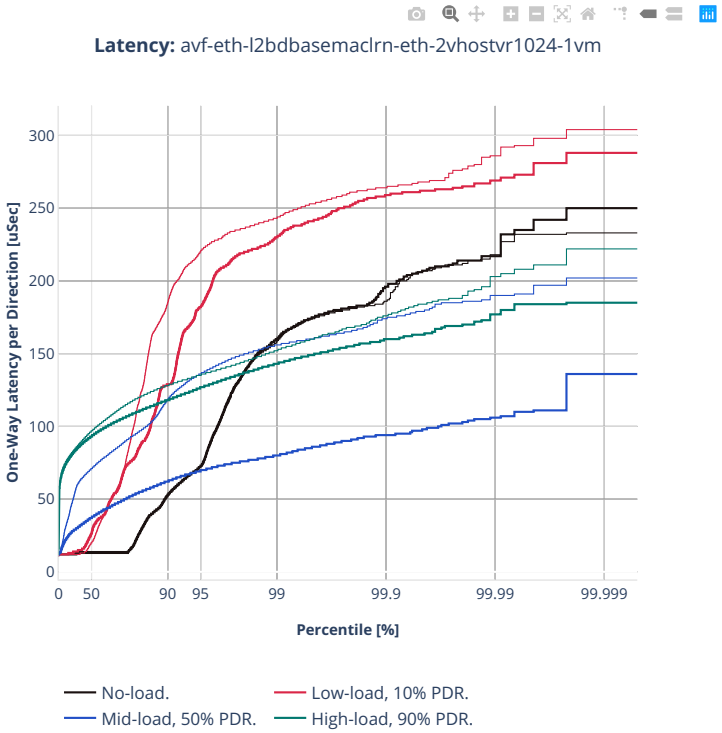
Latency: avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
One-Way Latency per Direction [uSec]
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

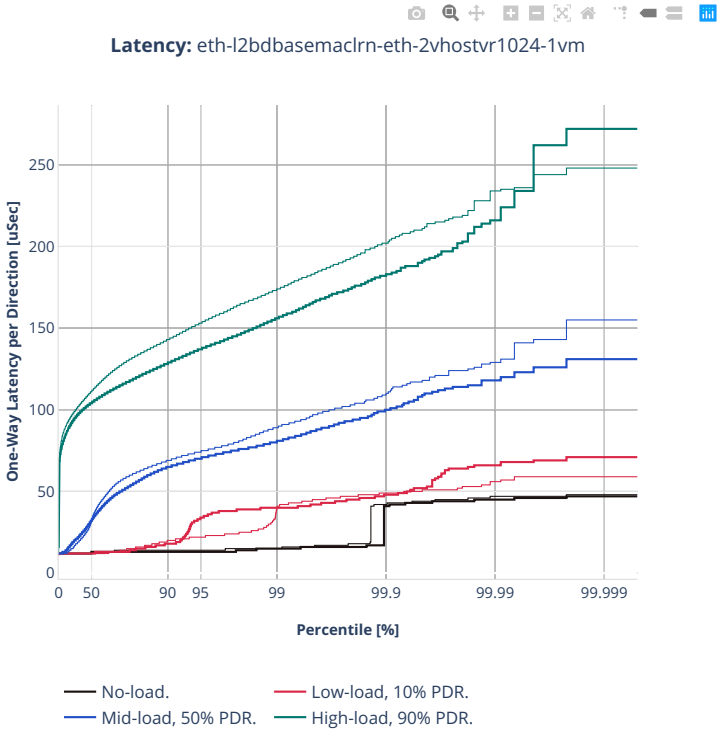

64b-2t1c-vhost-base-avf-vpp

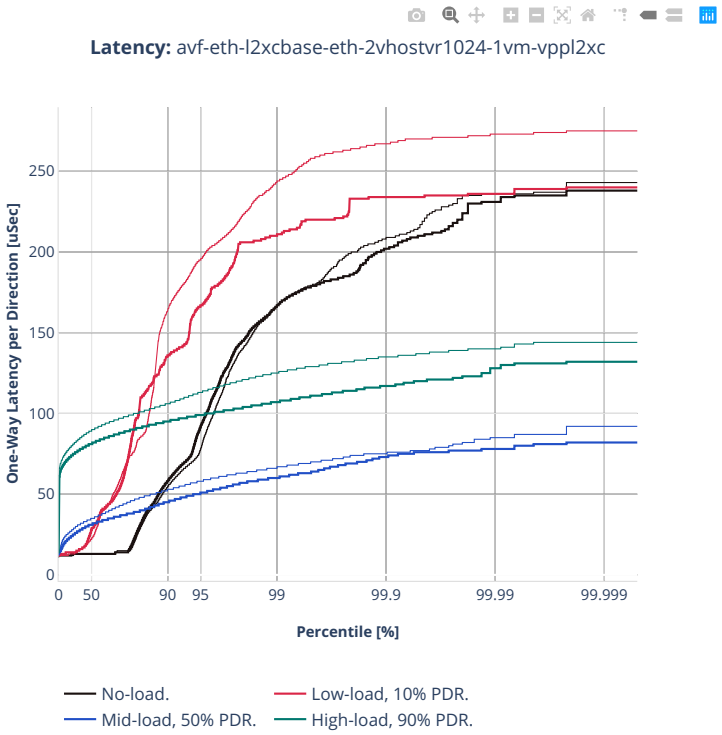

Latency: avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc

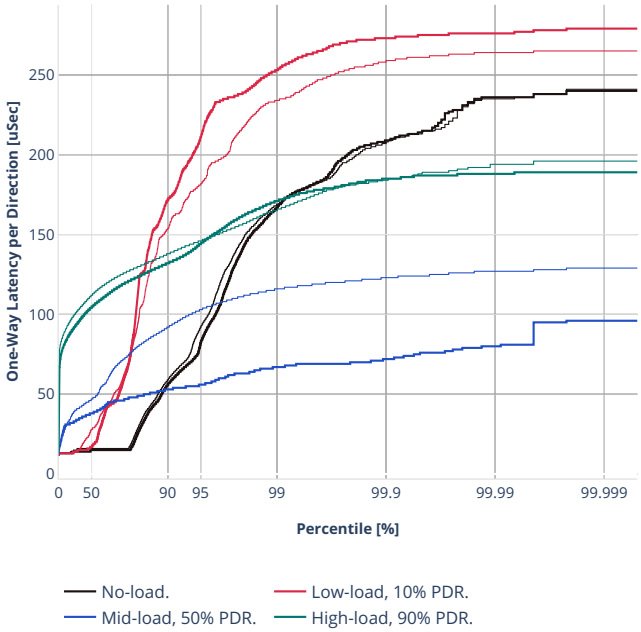

64b-2t1c-vhost-base-dpdk-vpp

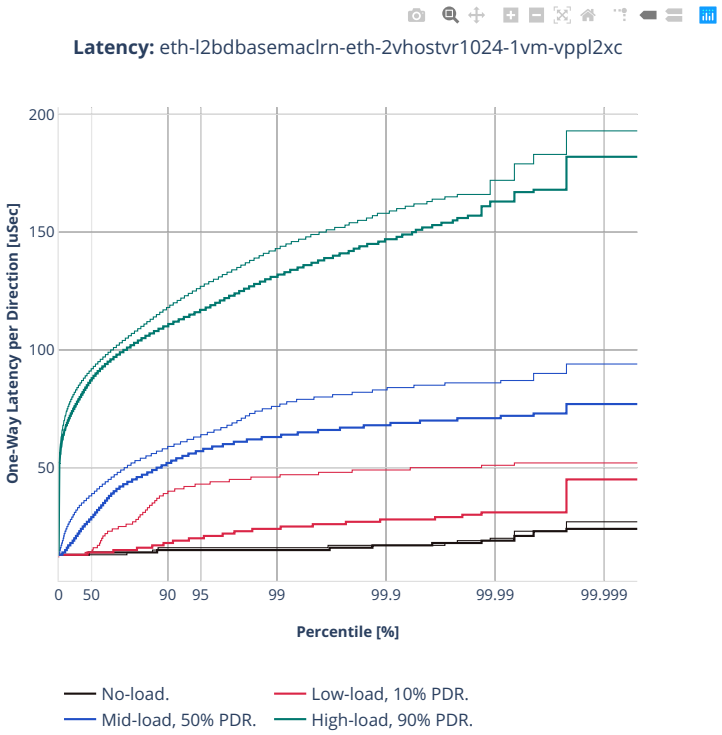

2n-skx-xxv710

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

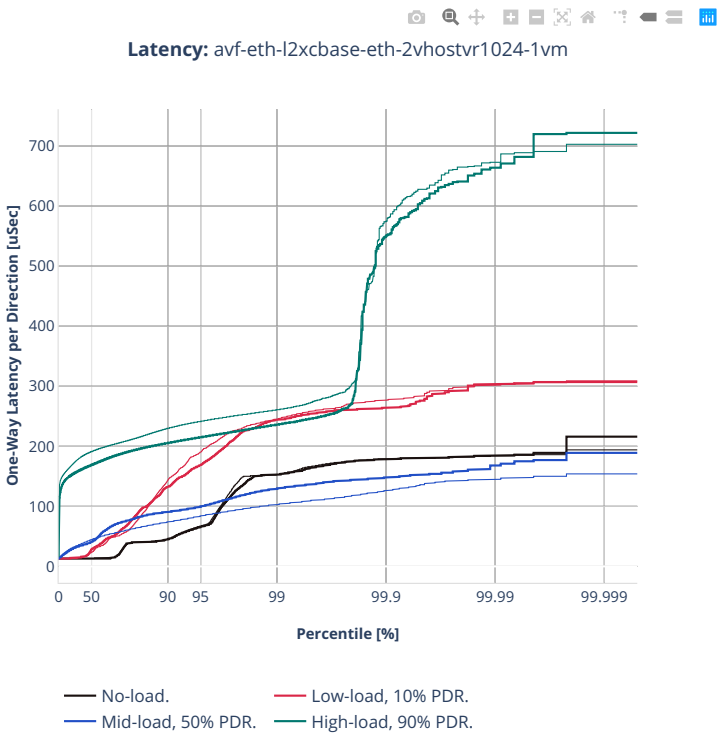

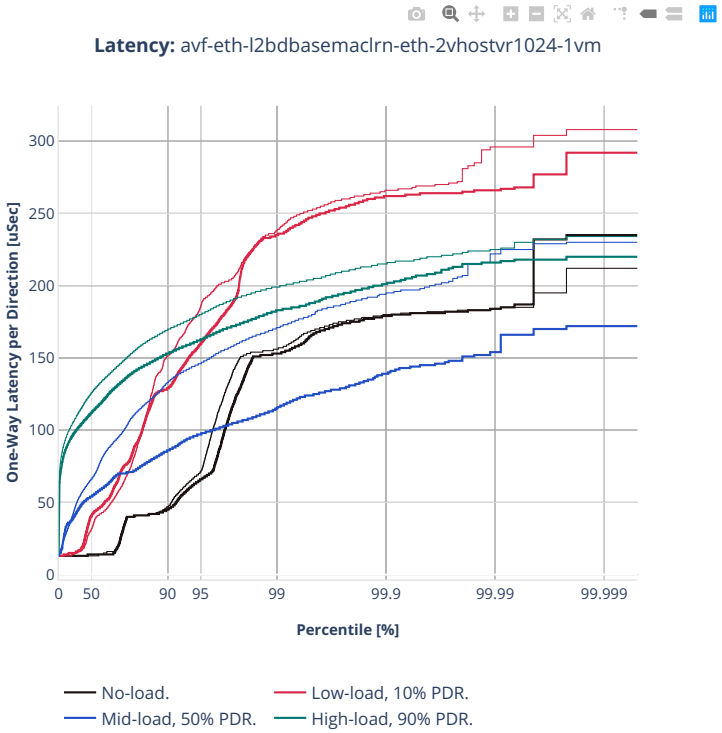
Latency: avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
One-Way Latency per Direction [uSec]
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

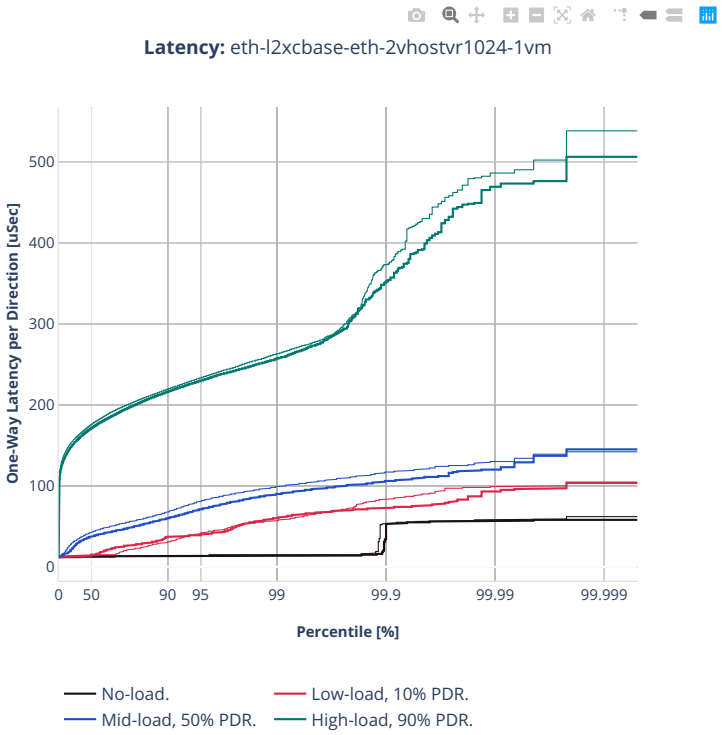

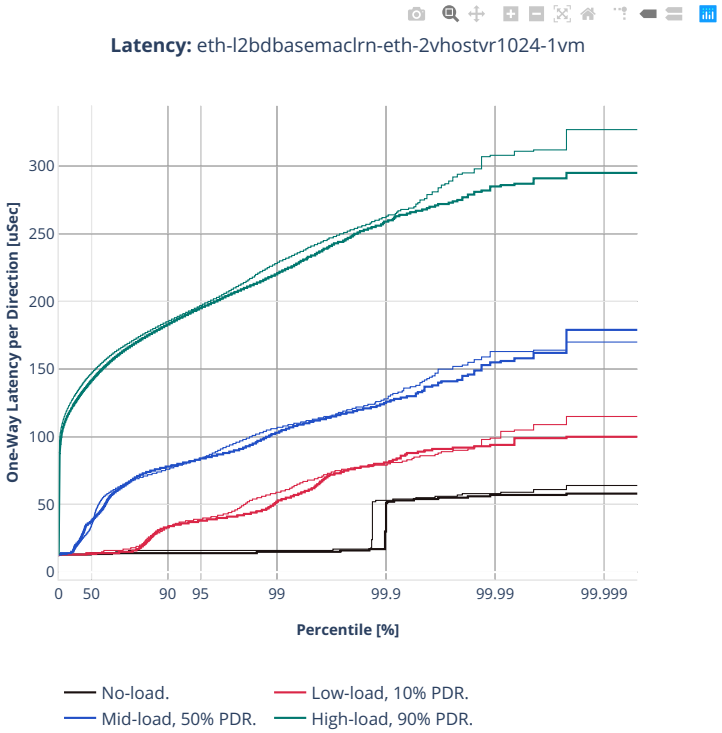
Latency: eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
One-Way Latency per Direction [uSec]
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

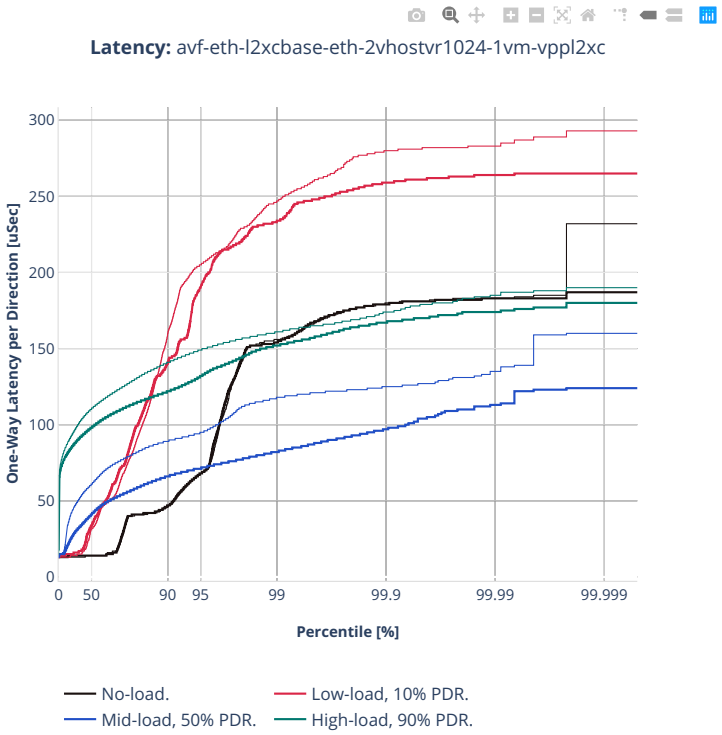

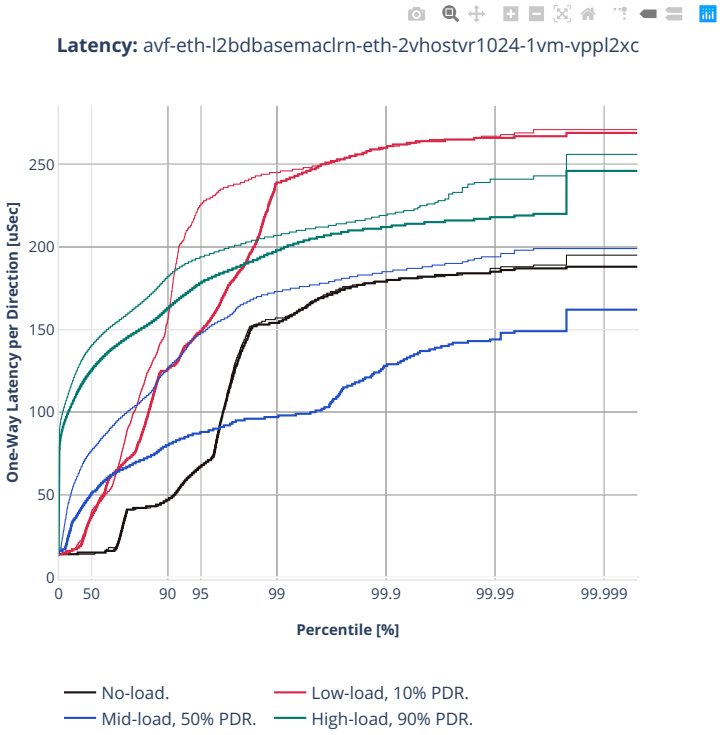
Latency: avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
One-Way Latency per Direction [uSec]
0
50
100
150
200
250
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

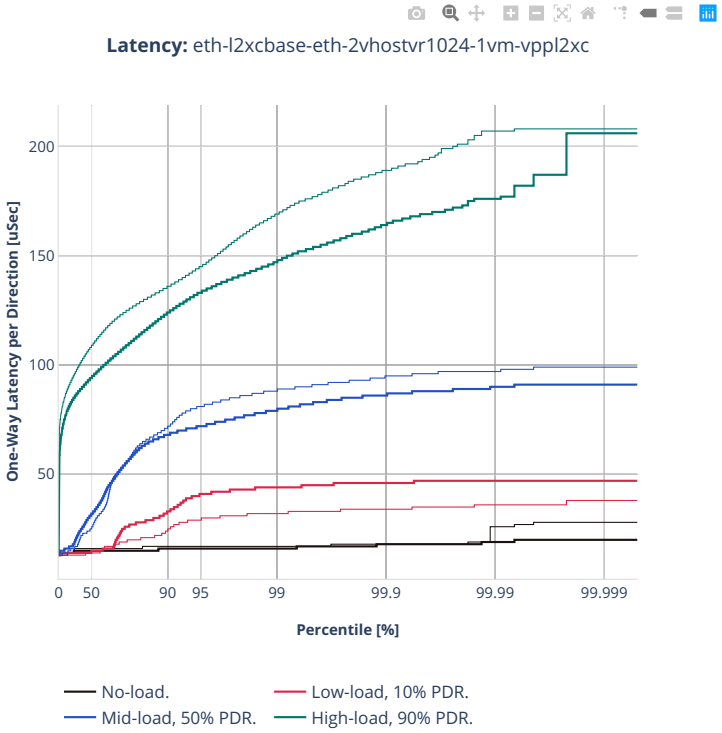

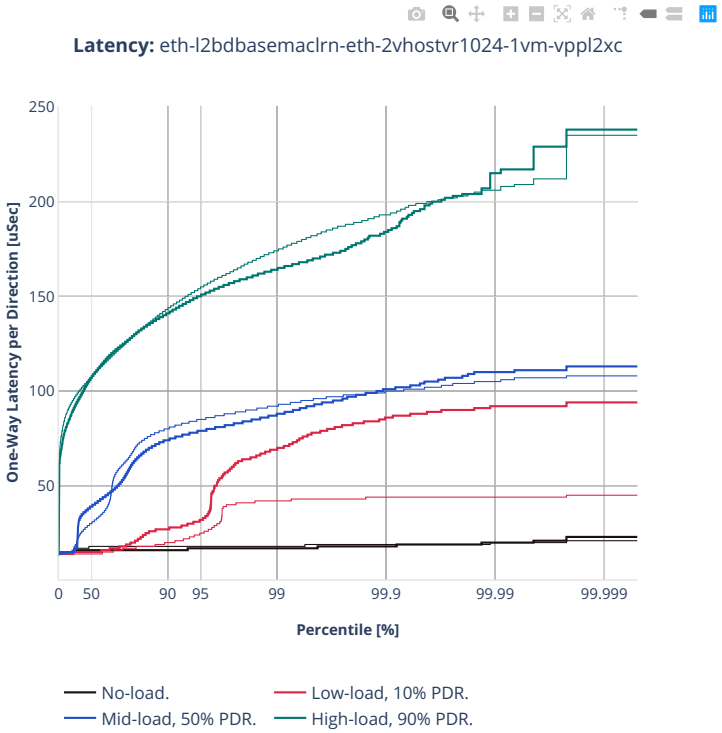
Latency: eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-clx-xxv710

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

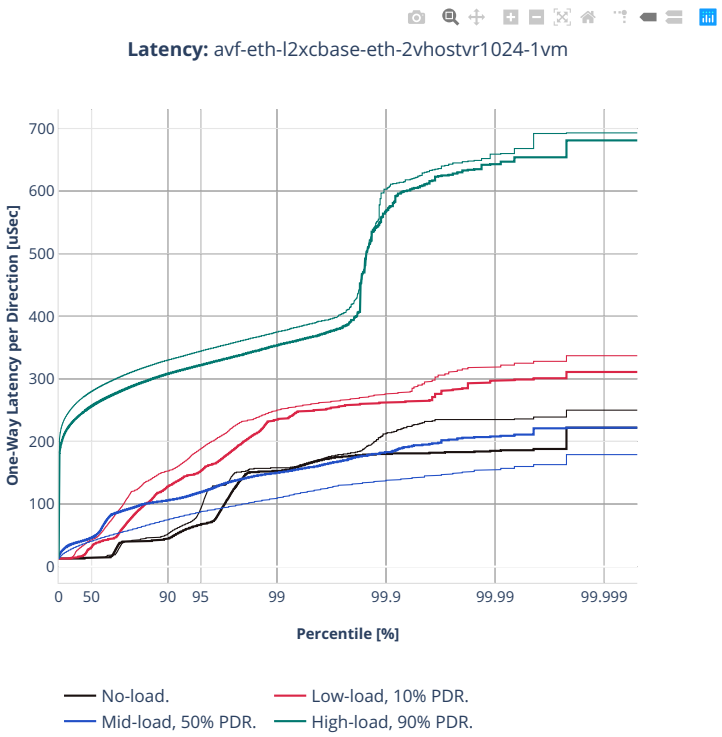

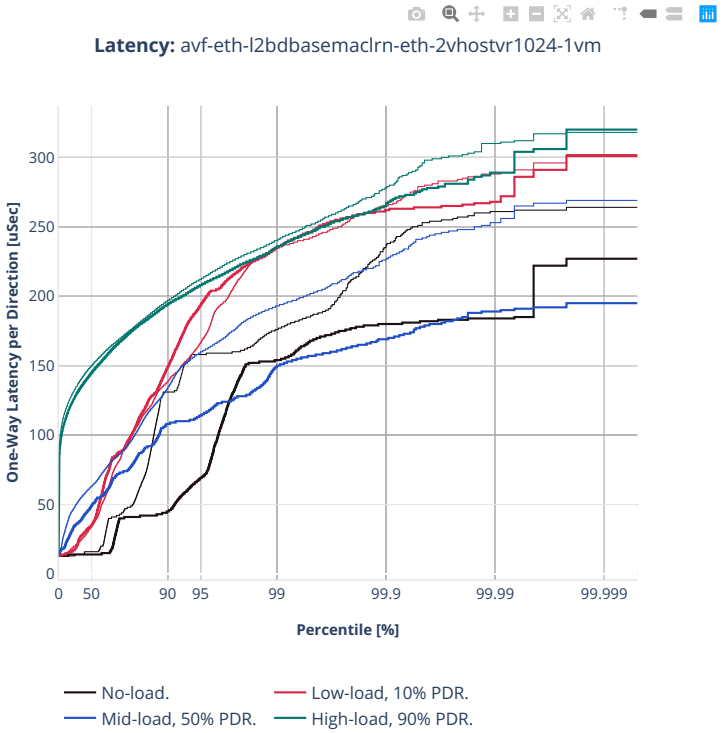
Latency: avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
One-Way Latency per Direction [uSec]
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

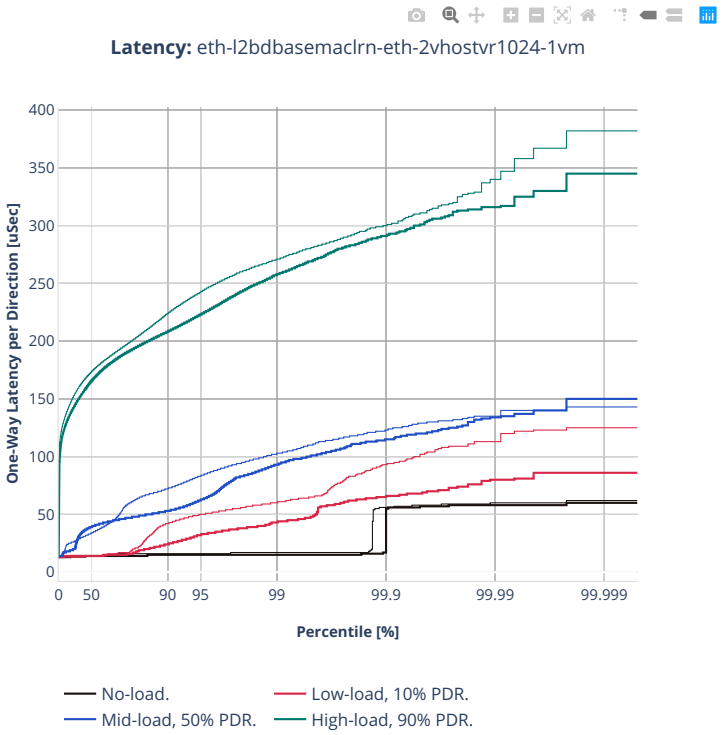

64b-2t1c-vhost-base-avf-vpp

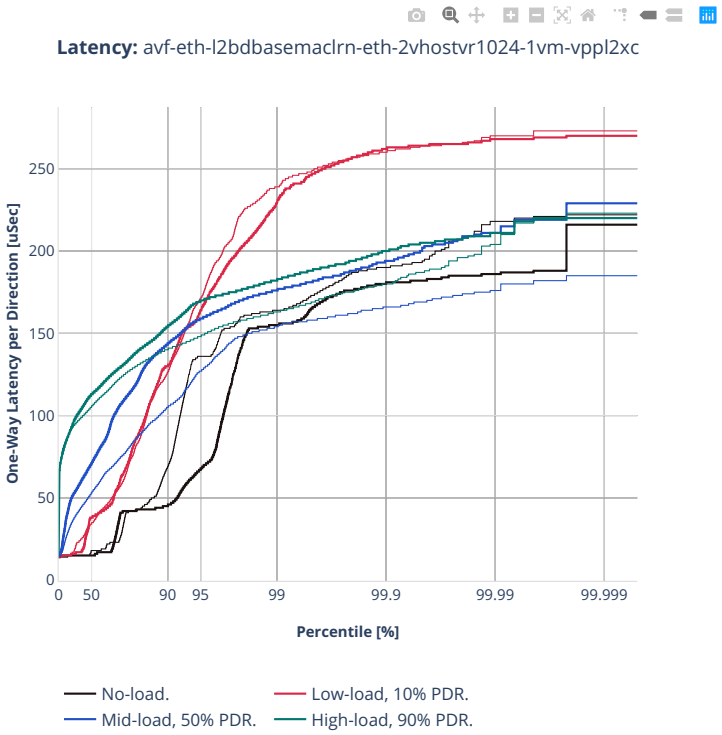

64b-2t1c-vhost-base-dpdk-vpp

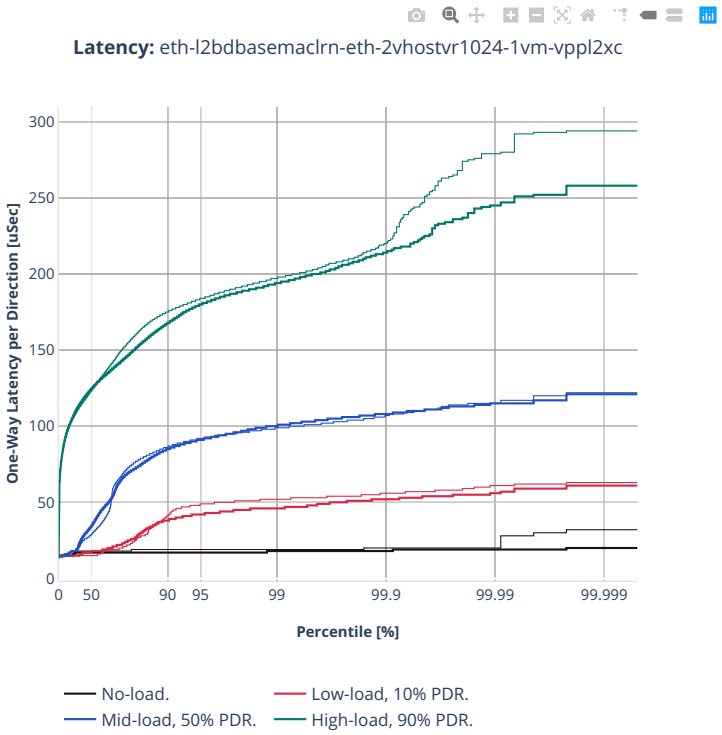

2n-clx-cx556a

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

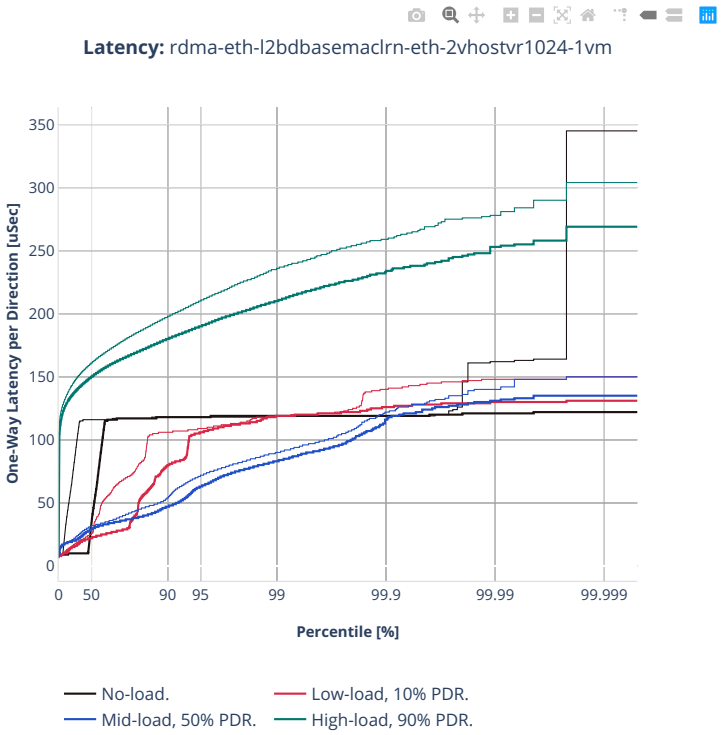

64b-2t1c-vhost-base-rdma-vpp

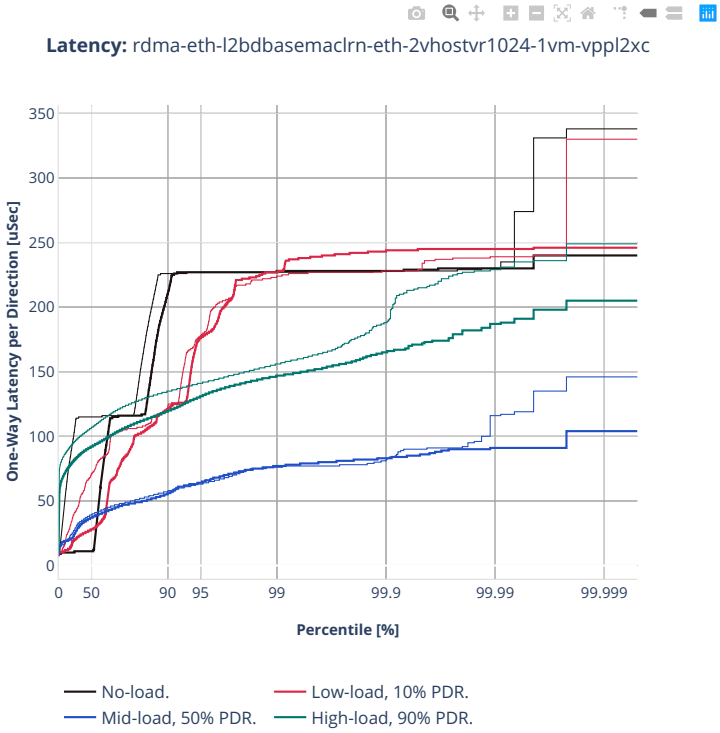

2n-clx-e810cq

64b-2t1c-vhost-base

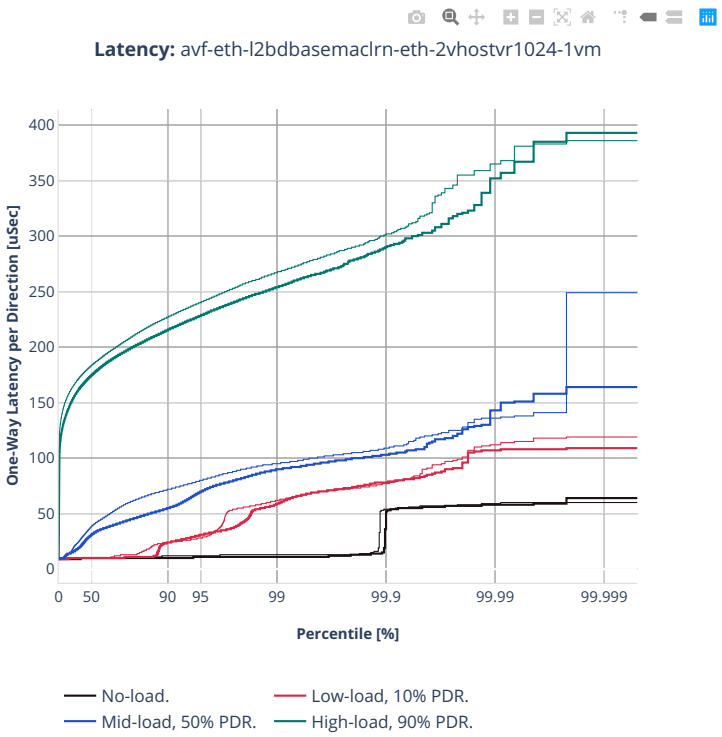

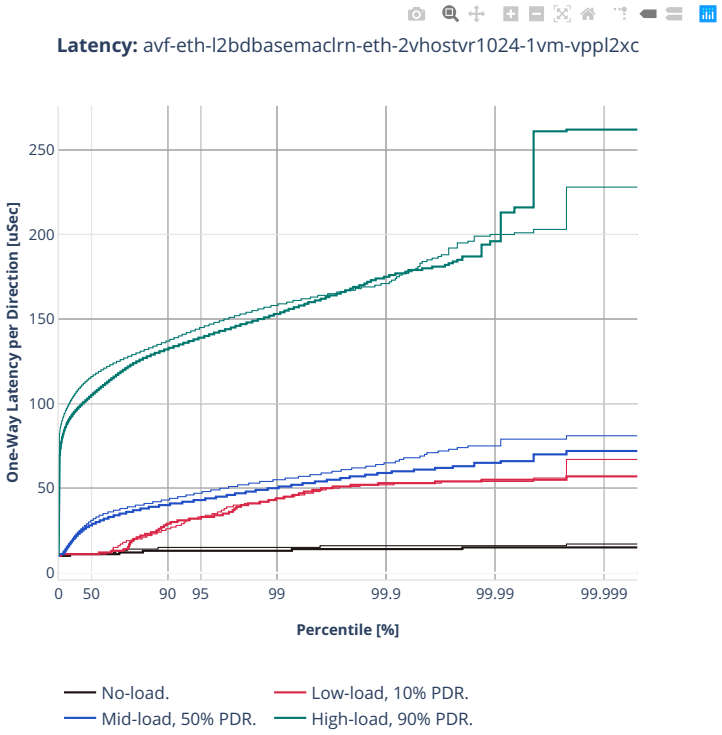
Latency: avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

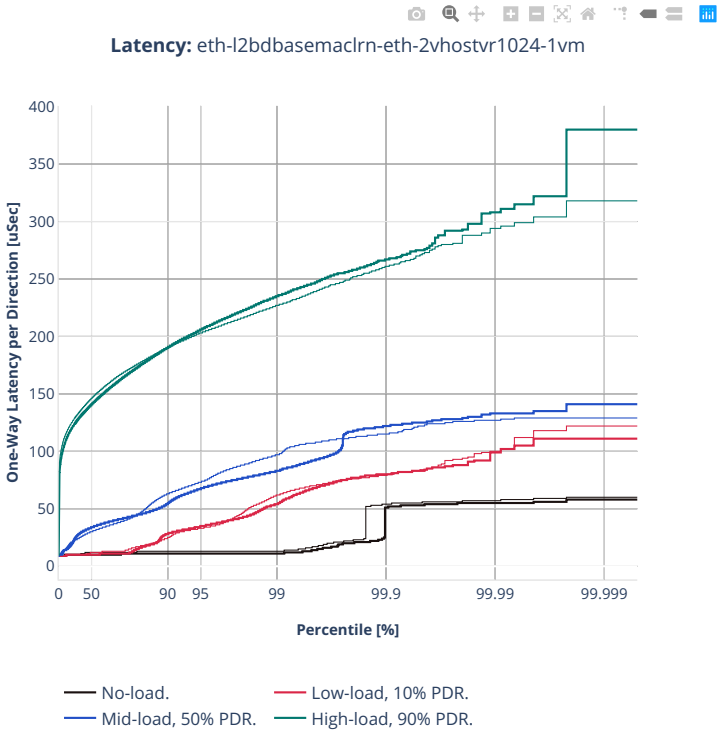
Latency: eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
One-Way Latency per Direction [uSec]
400
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

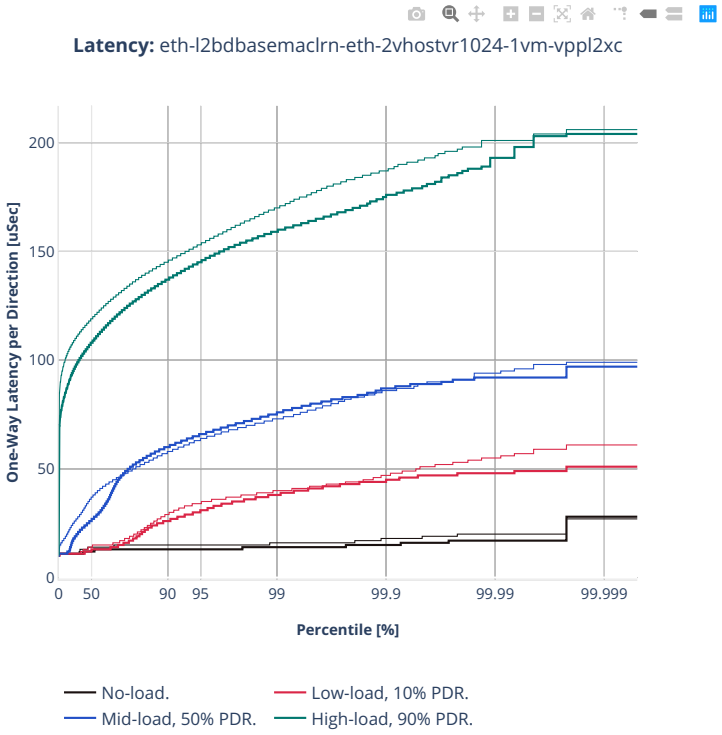
Latency: eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-zn2-xxv710

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

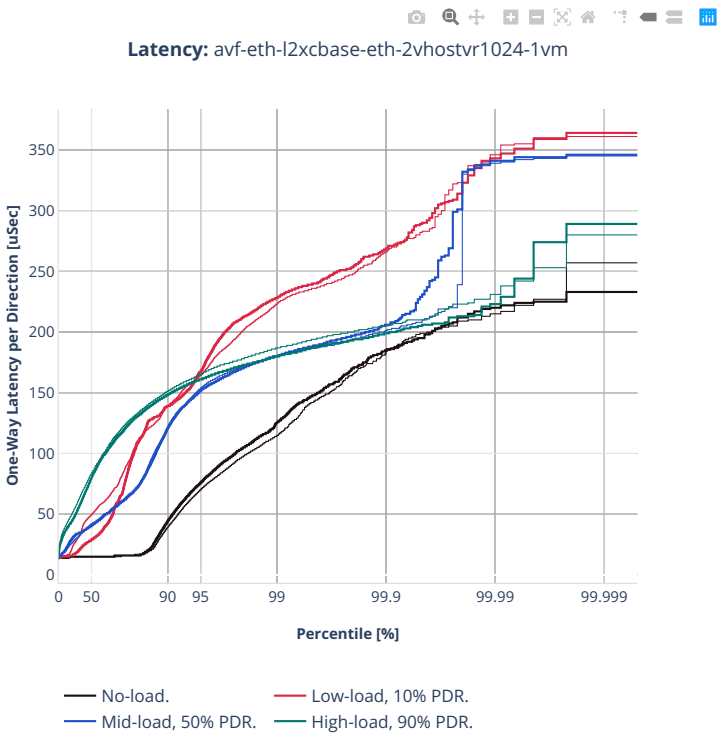

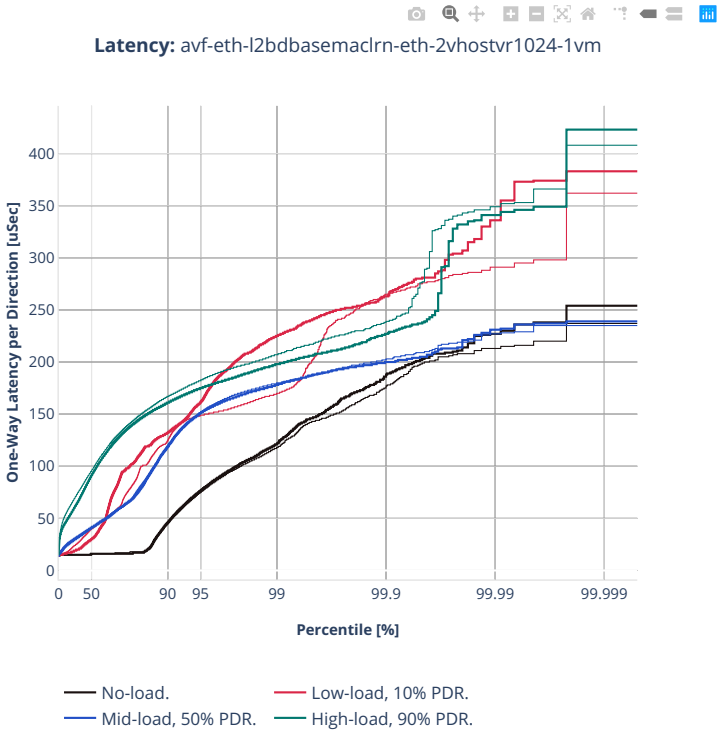
Latency: avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm
One-Way Latency per Direction [uSec]
400
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

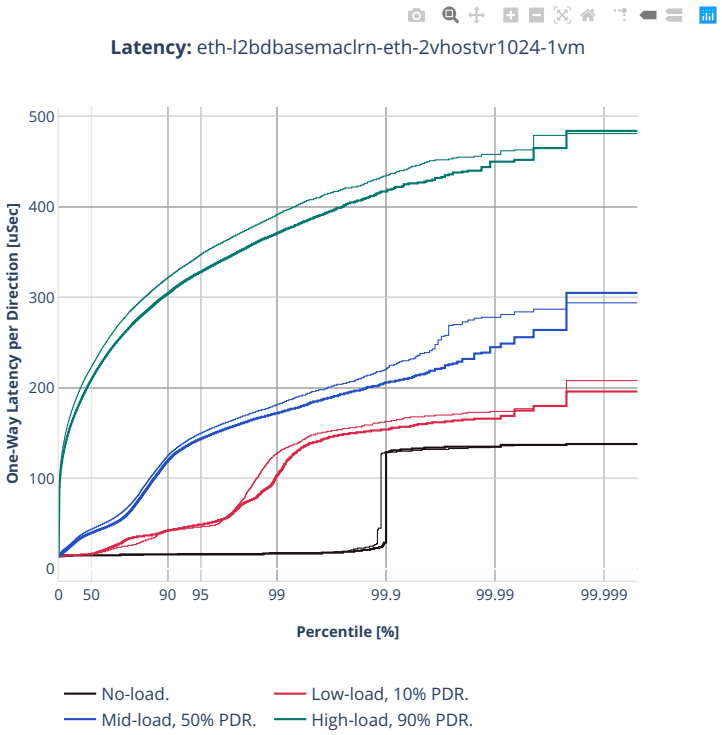

64b-2t1c-vhost-base-avf-vpp

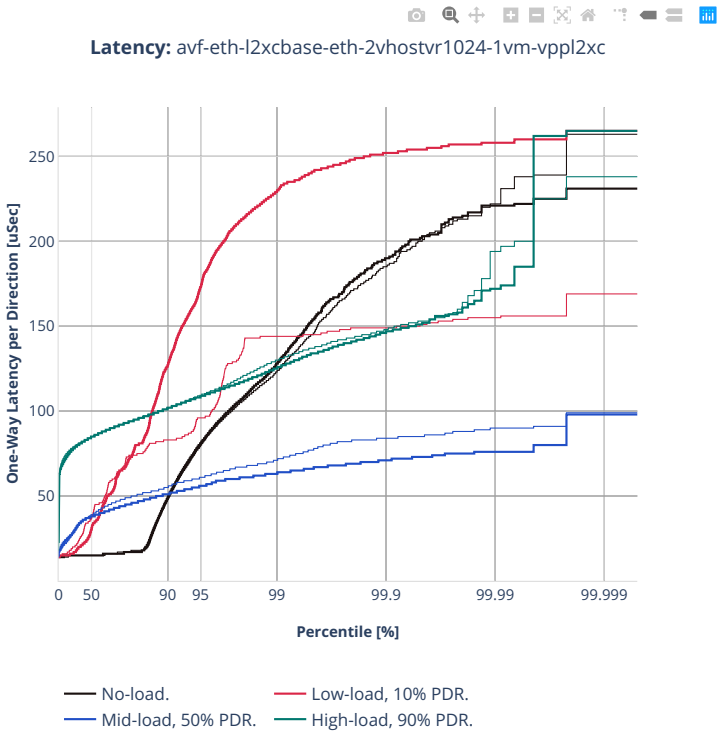

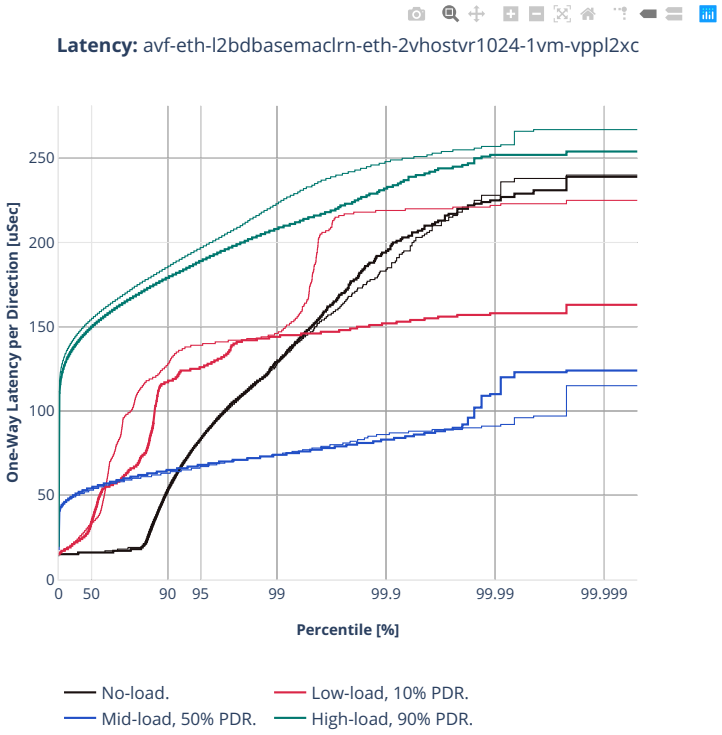
Latency: avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

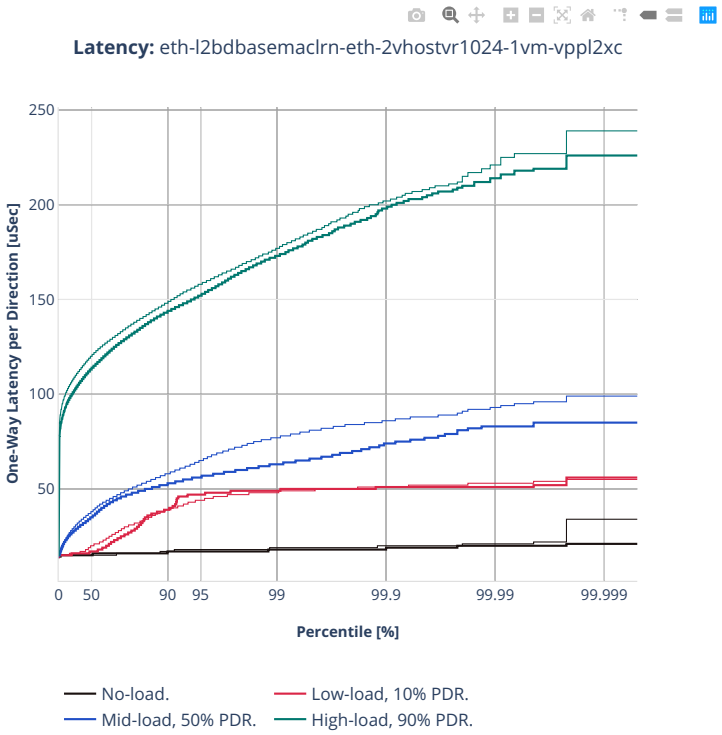

3n-alt-xl710

64b-1t1c-vhost-base

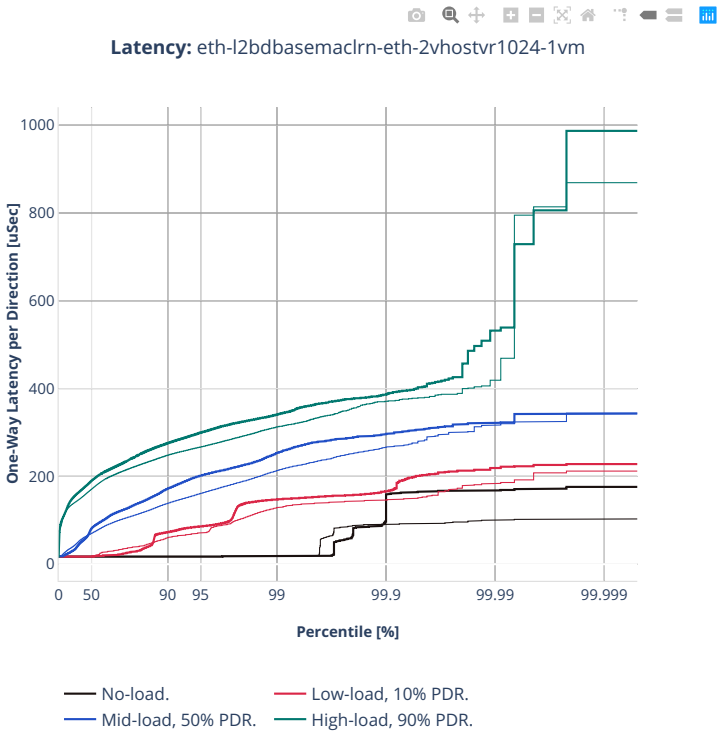

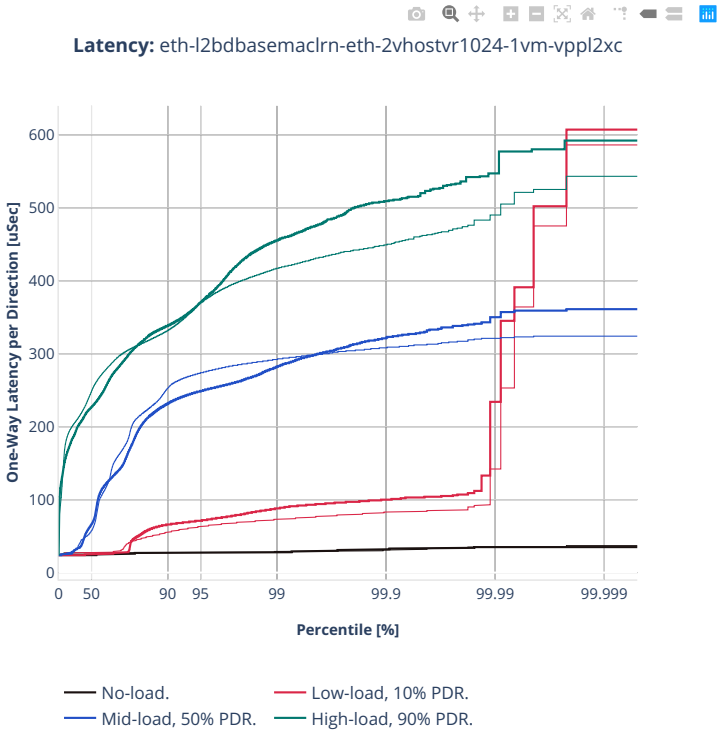
Latency: eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

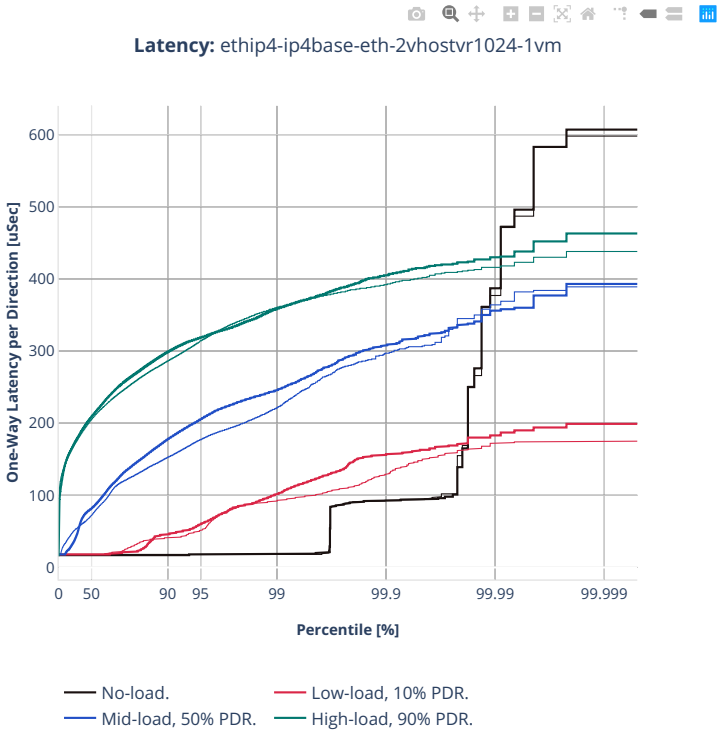
Latency: ethip4-ip4base-eth-2vhostvr1024-1vm
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

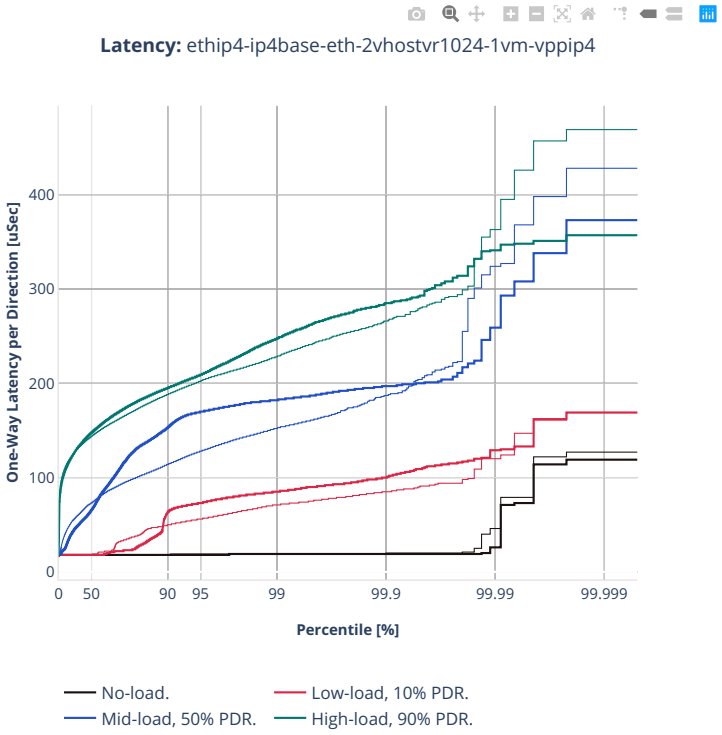
Latency: ethip4-ip4base-eth-2vhostvr1024-1vm-vppip4
One-Way Latency per Direction [uSec]
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-tsh-x520

64b-1t1c-vhost-base-ixgbe

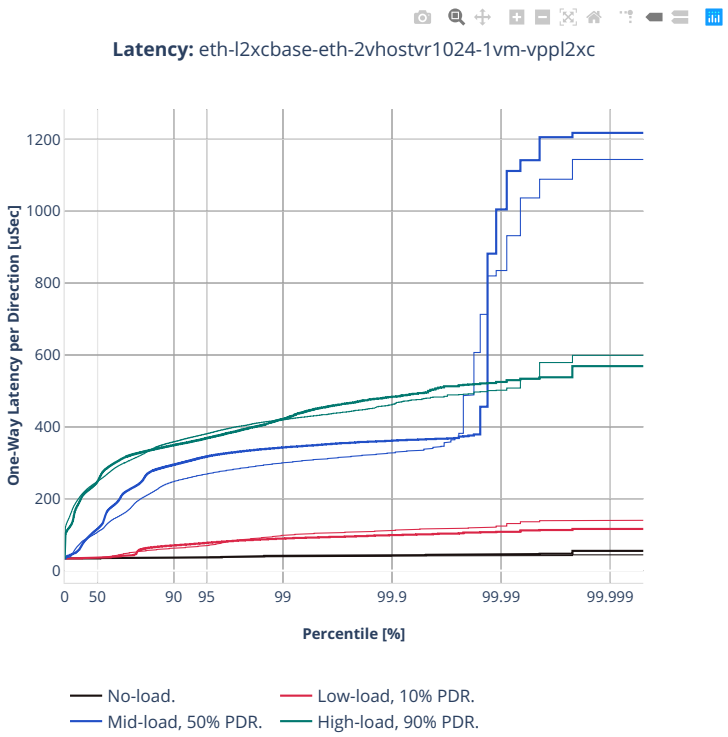

Latency: ethip4-ip4base-eth-2vhostvr1024-1vm-vppip4

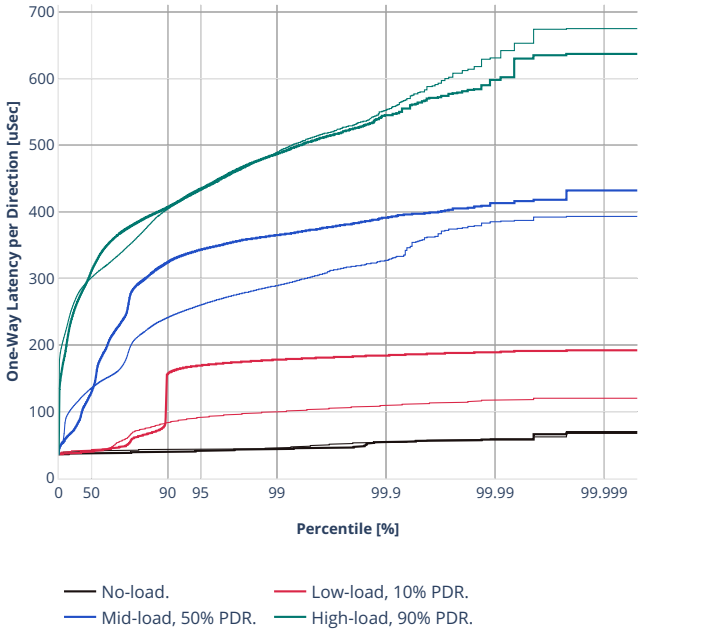

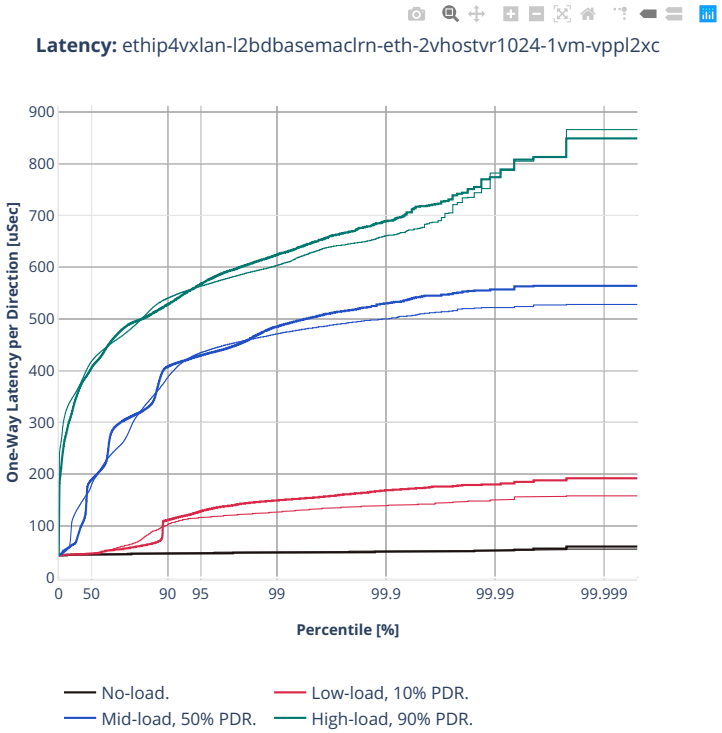
Latency: ethip4vxlan-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc
One-Way Latency per Direction [uSec]
900
800
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2.5.8 LXC/DRC Container Memif

CSIT source code for the test cases used for plots can be found in CSIT git repository[156].

[156] https://git.fd.io/csit/tree/tests/vpp/perf/container_memif?h=rls2206

2n-icx-xxv710

64b-2t1c-memif-base-avf

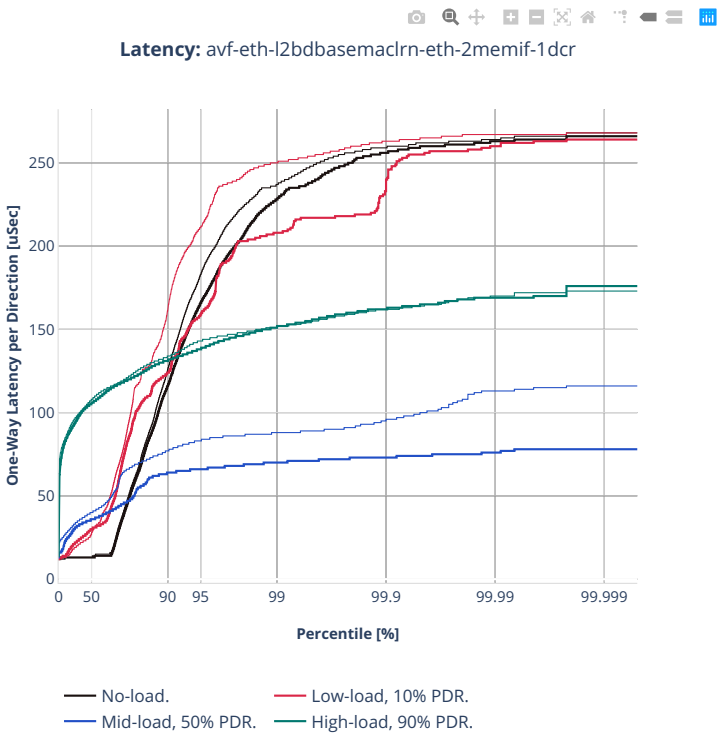

Latency: avf-eth-l2xcbase-eth-2memif-1dcr

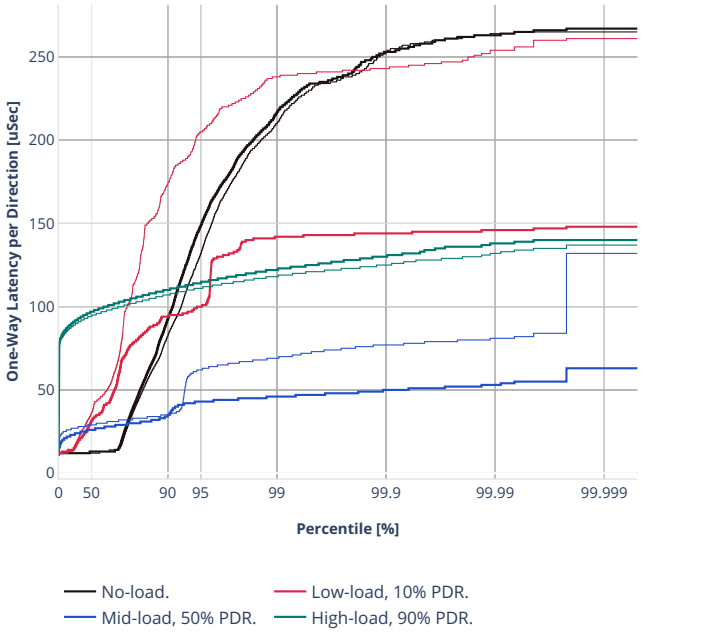

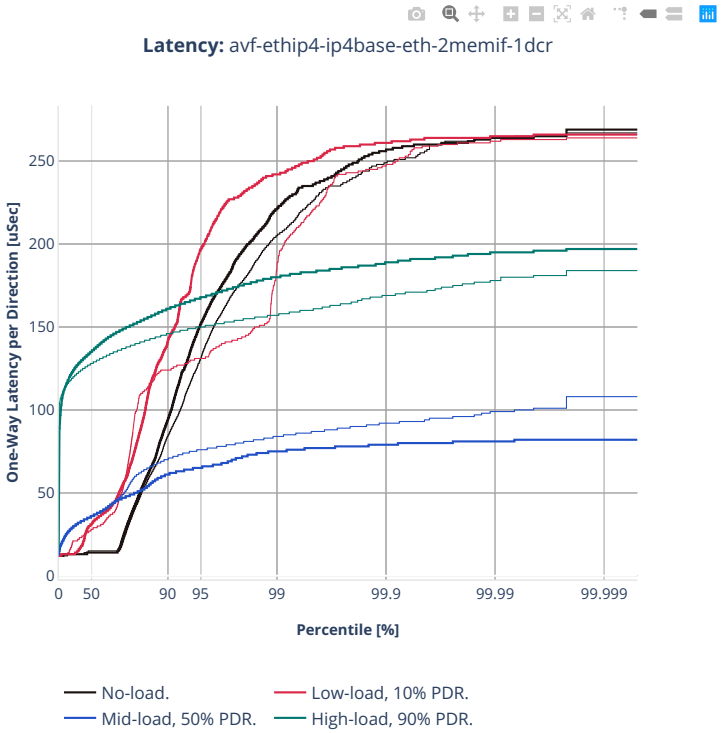
Latency: avf-ethip4-ip4base-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-memif-base-dpdk

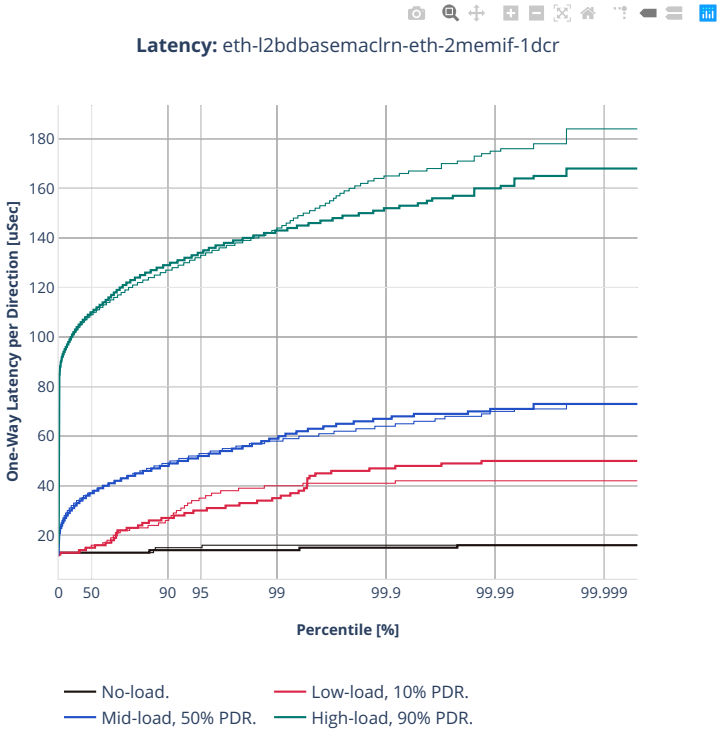

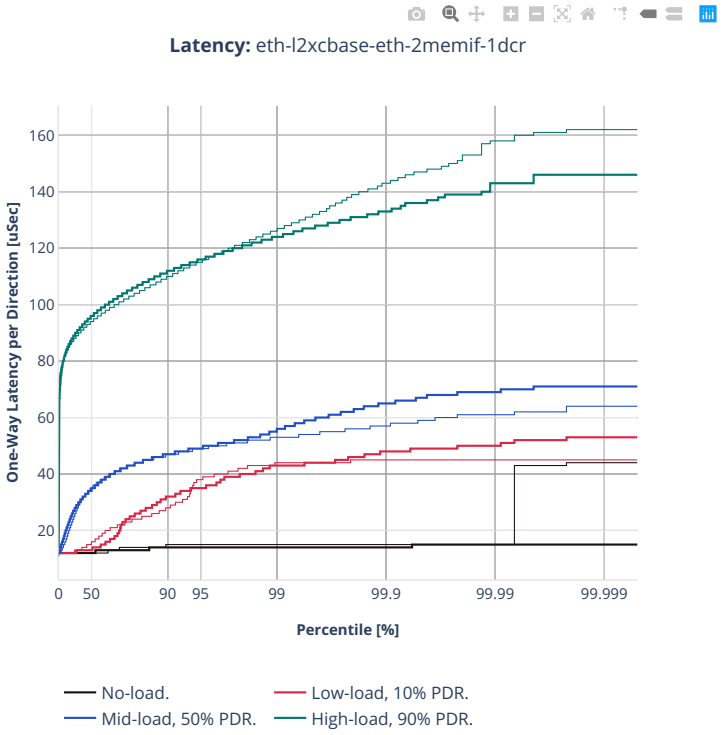
Latency: eth-l2xcbase-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

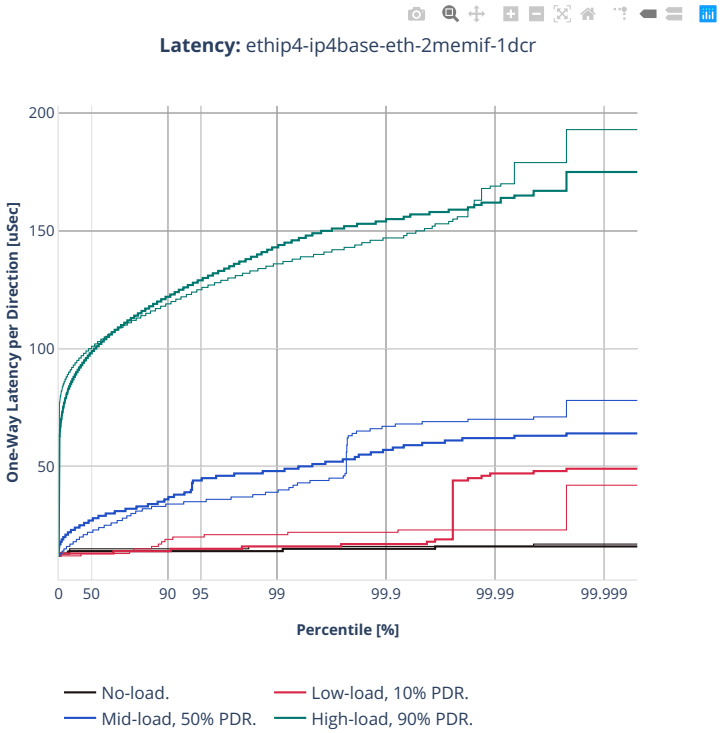
Latency: ethip4-ip4base-eth-2memif-1dcr
200
150
100
50
One-Way Latency per Direction [uSec]
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-skx-xxv710

64b-2t1c-memif-base-avf

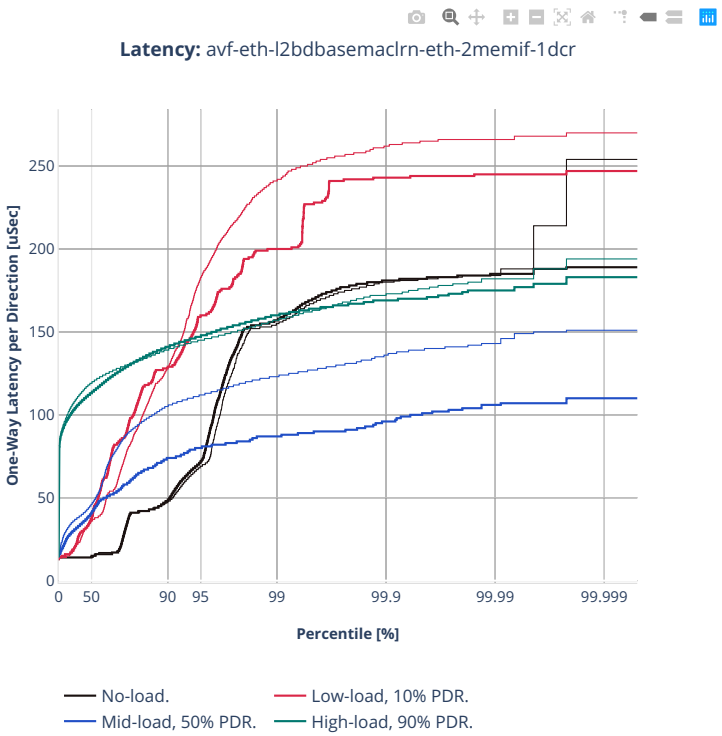

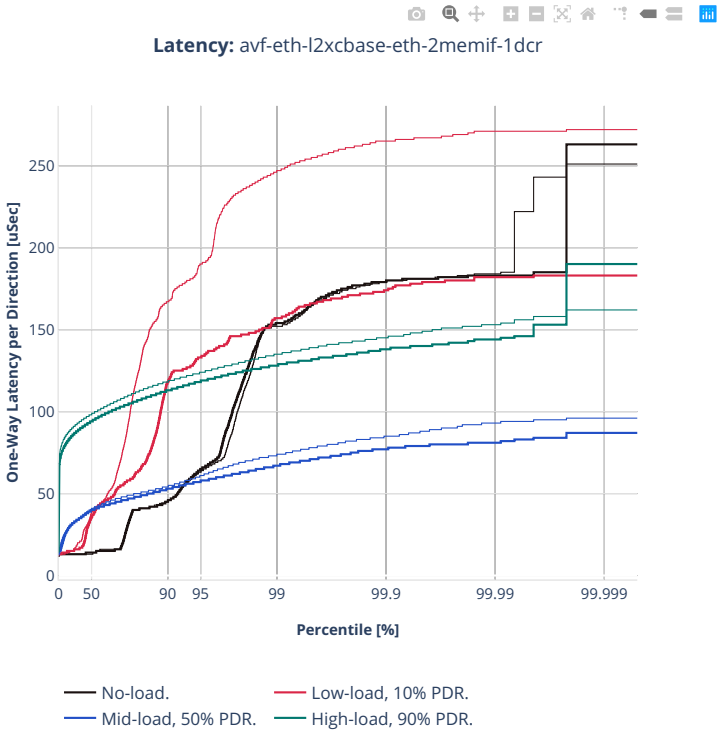
Latency: avf-eth-l2xcbase-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

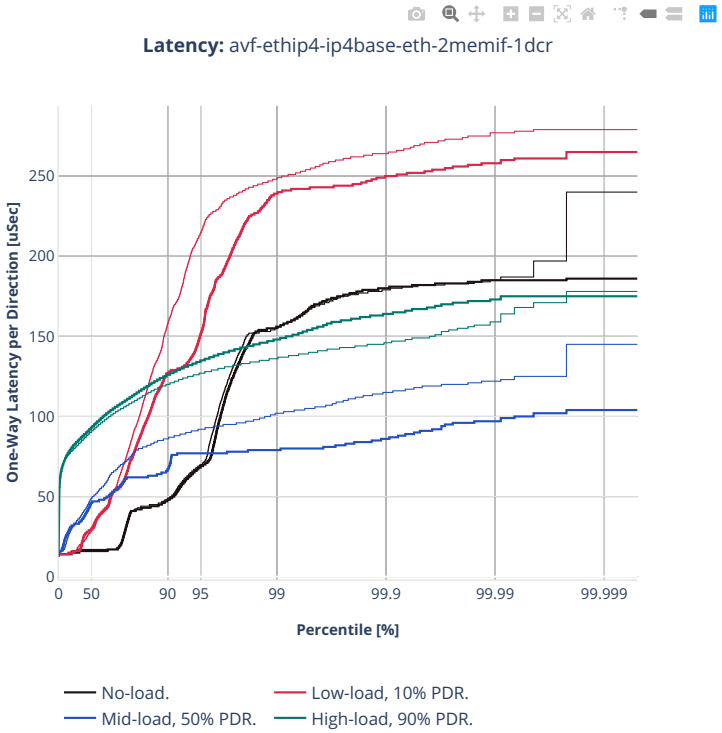
Latency: avf-ethip4-ip4base-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-memif-base-dpdk

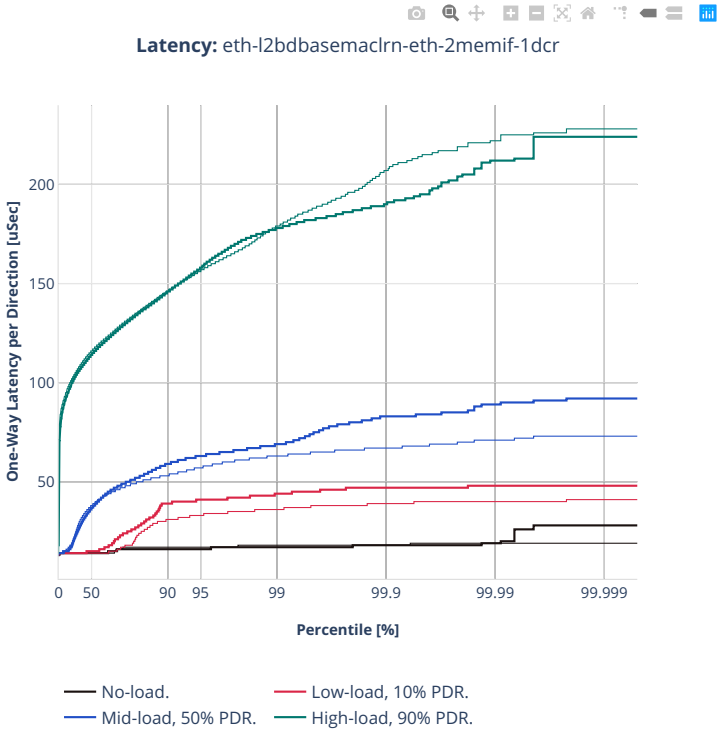

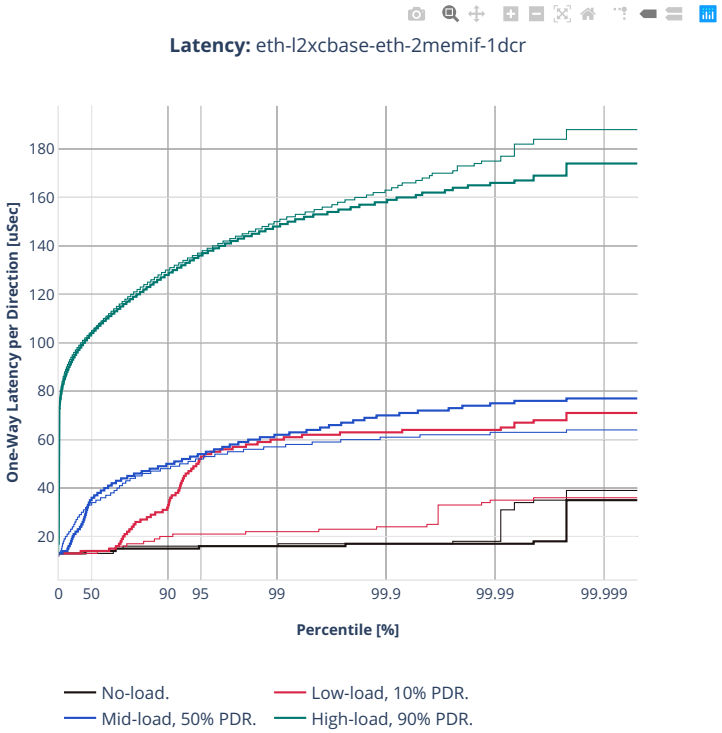
Latency: eth-l2xcbase-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

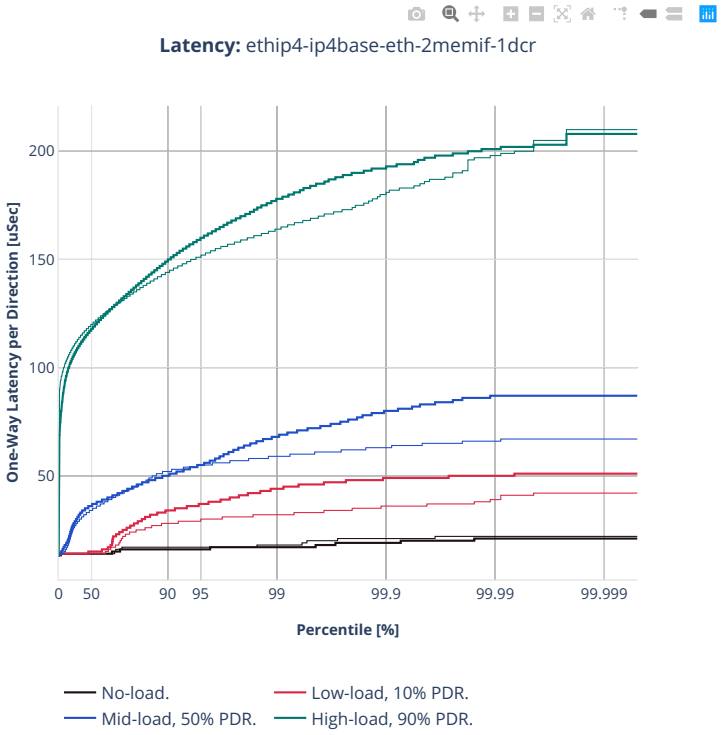
Latency: ethip4-ip4base-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
200
150
100
50
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-clx-xxv710

64b-2t1c-memif-base-avf

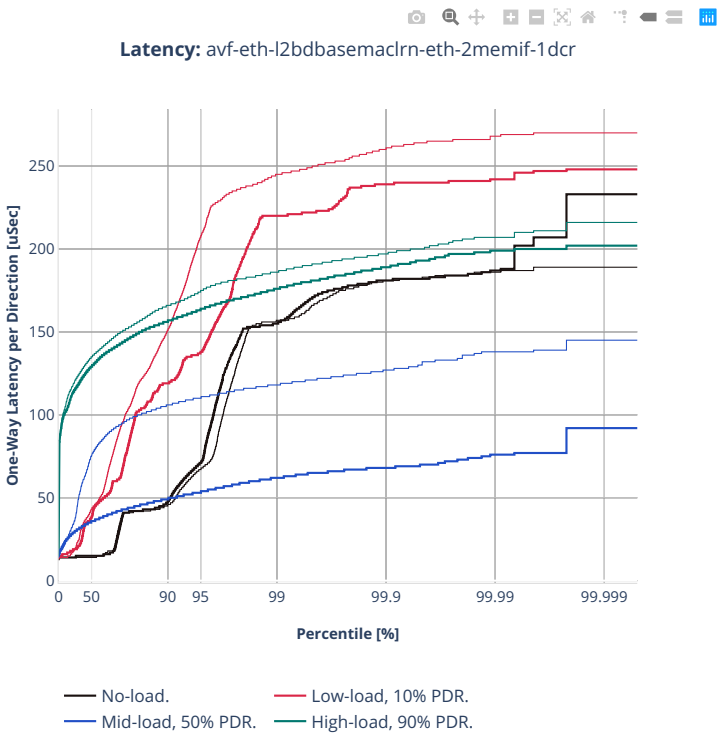

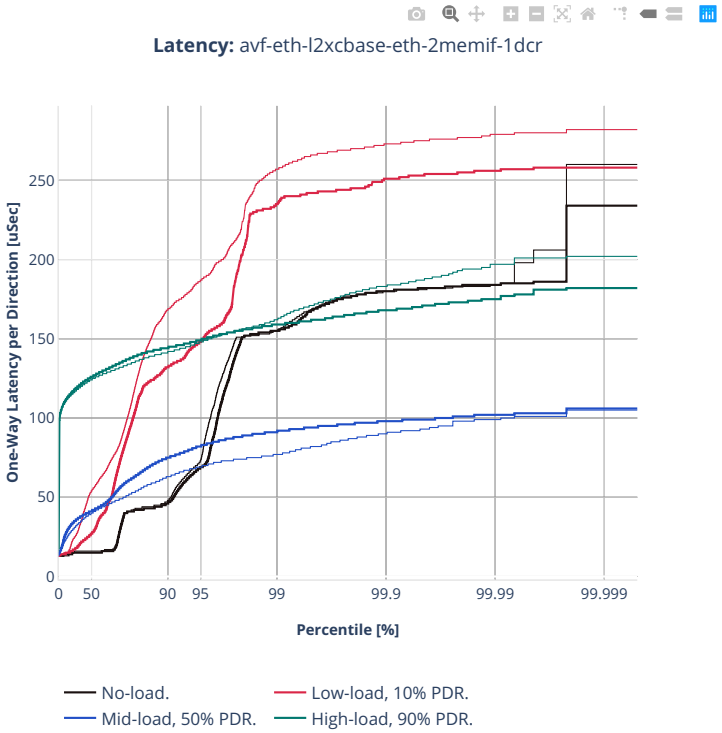
Latency: avf-eth-l2xcbase-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
0
50
100
150
200
250
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

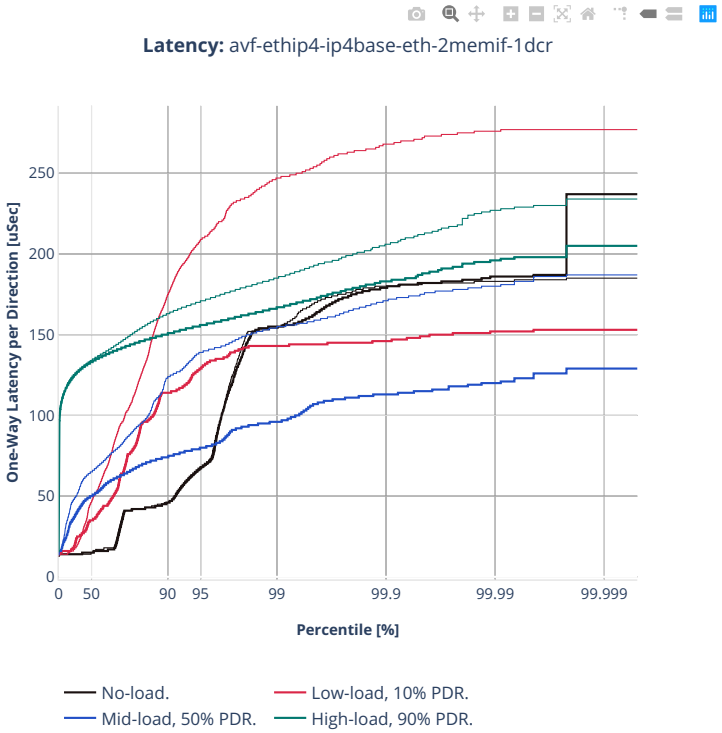
Latency: avf-ethip4-ip4base-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-memif-base-dpdk

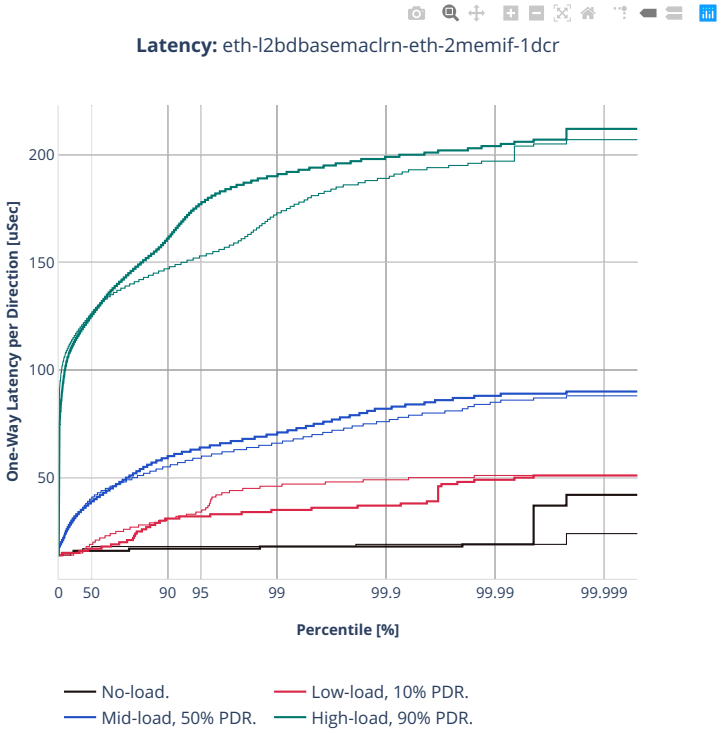

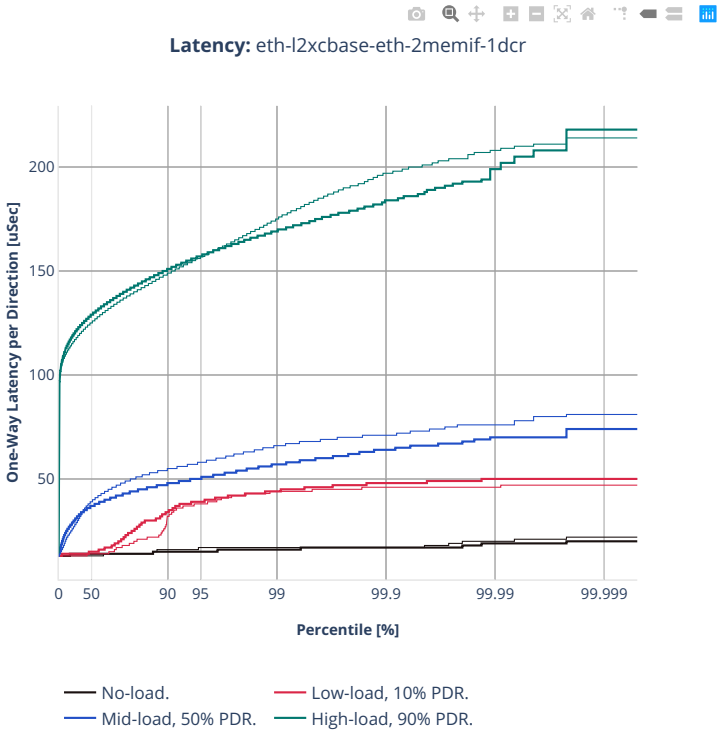
Latency: eth-l2xcbase-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
200
150
100
50
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

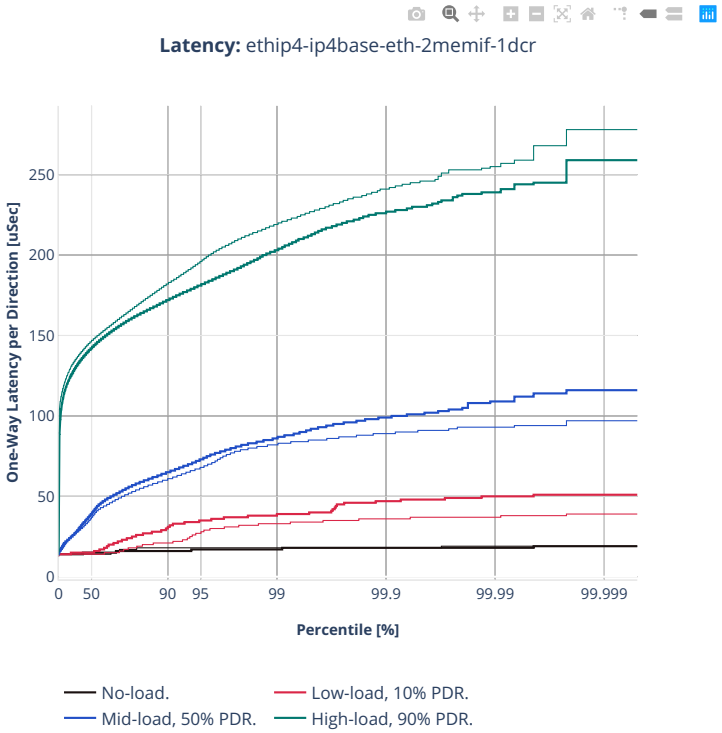
Latency: ethip4-ip4base-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-clx-cx556a

64b-2t1c-memif-base-rdma

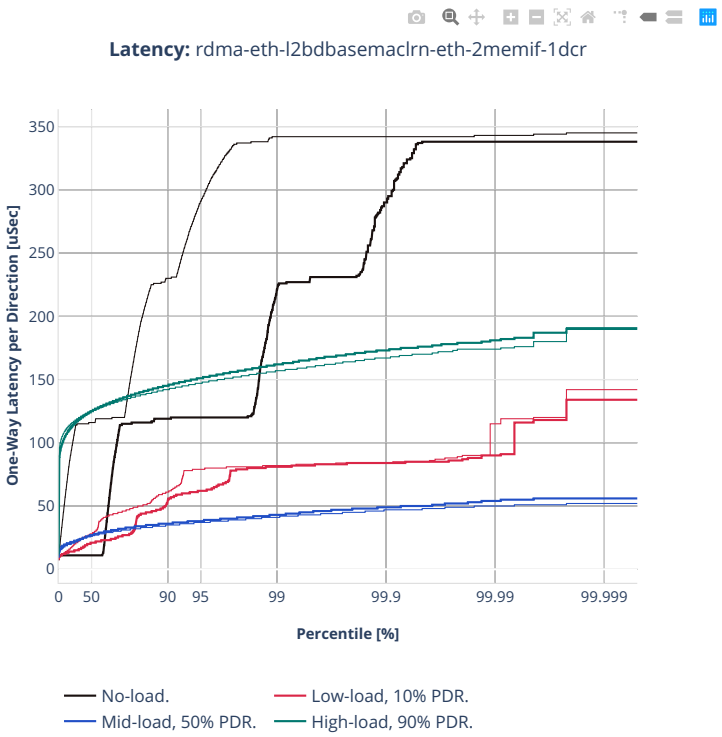

Latency: rdma-eth-l2xcbase-eth-2memif-1dcr

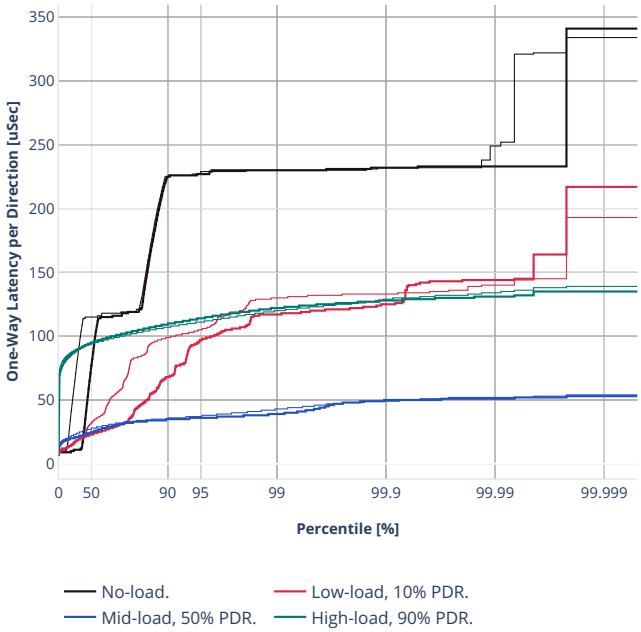

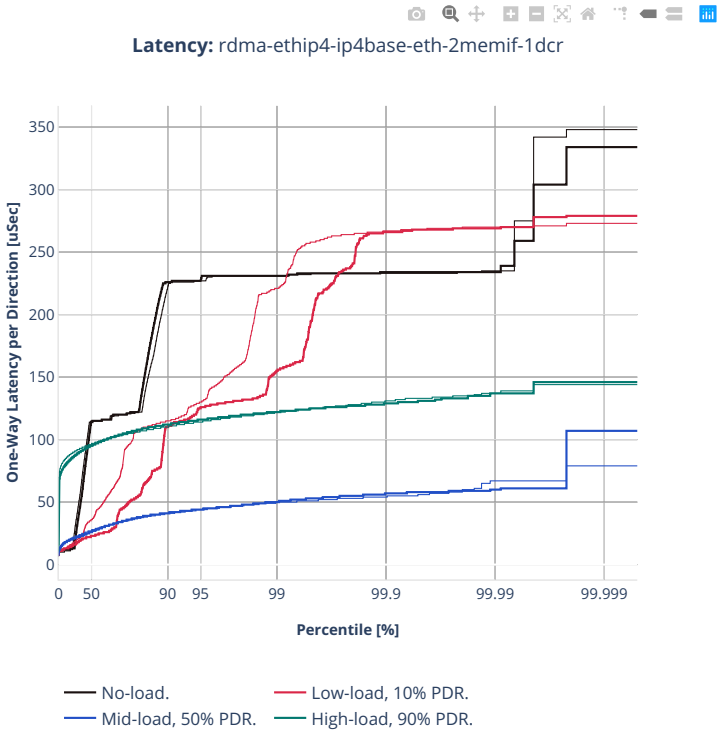
Latency: rdma-ethip4-ip4base-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-clx-e810cq

64b-2t1c-memif-base-rdma

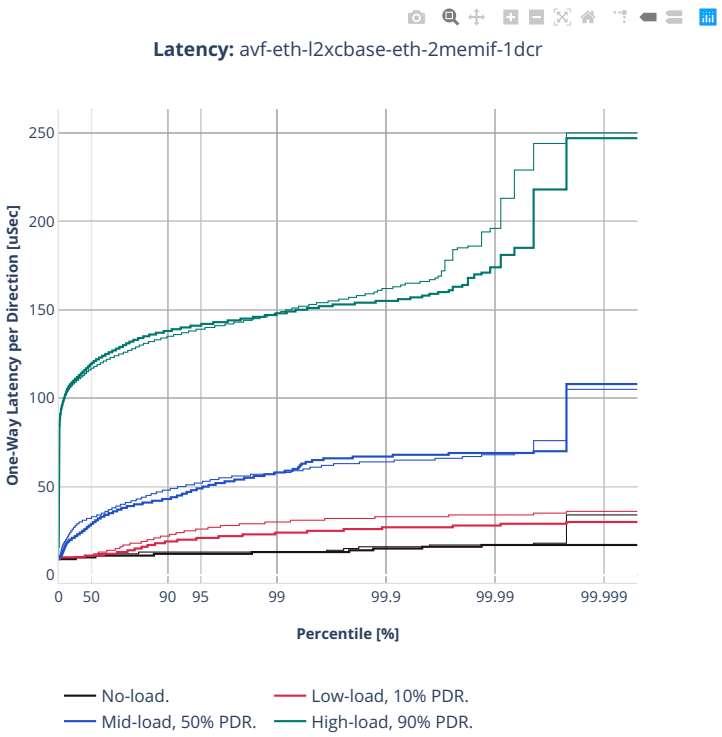

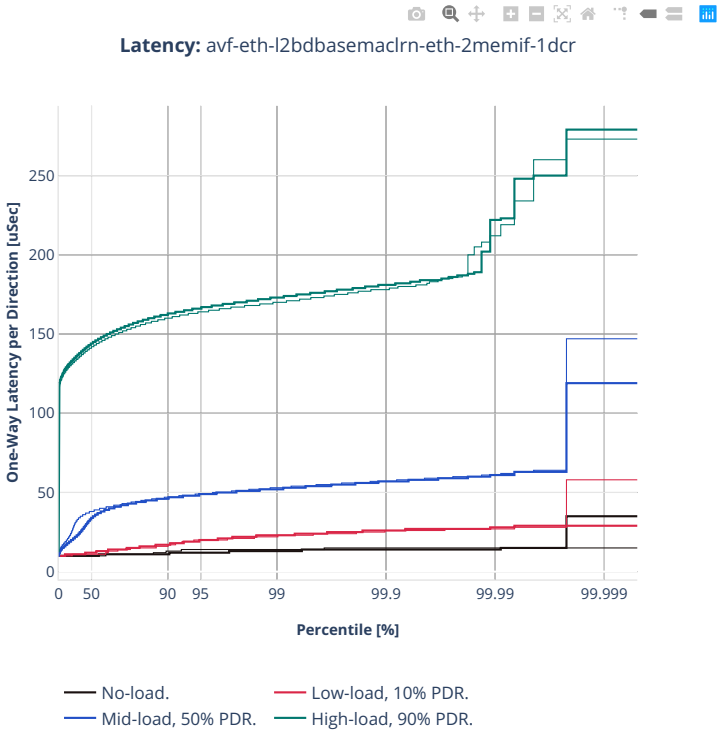
Latency: avf-eth-l2bdbasemaclrn-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

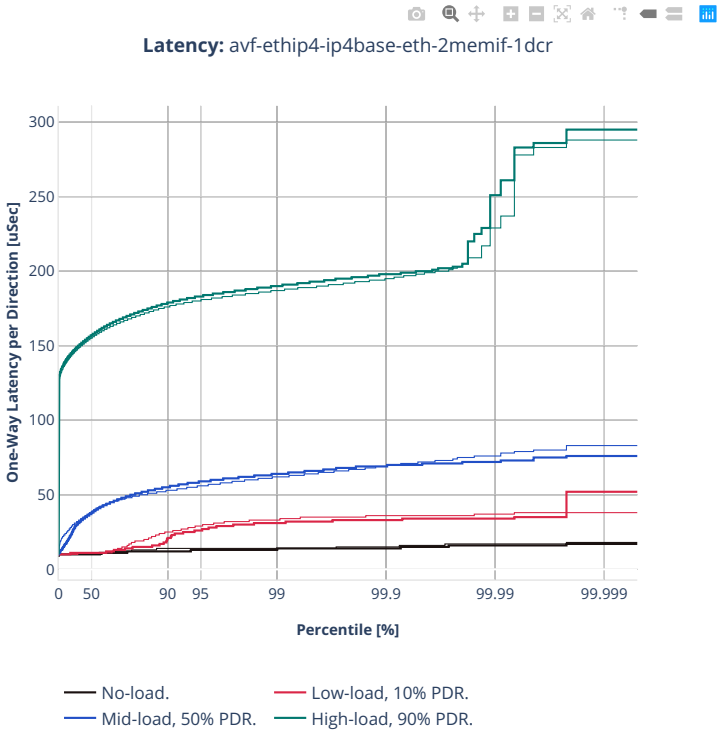
Latency: avf-ethip4-ip4base-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

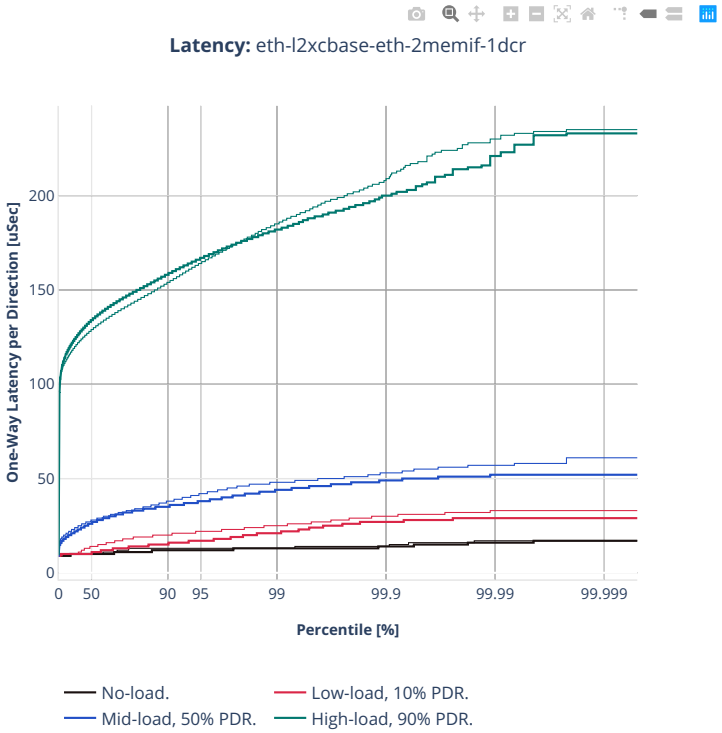
Latency: eth-l2xcbase-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
0
50
100
150
200
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

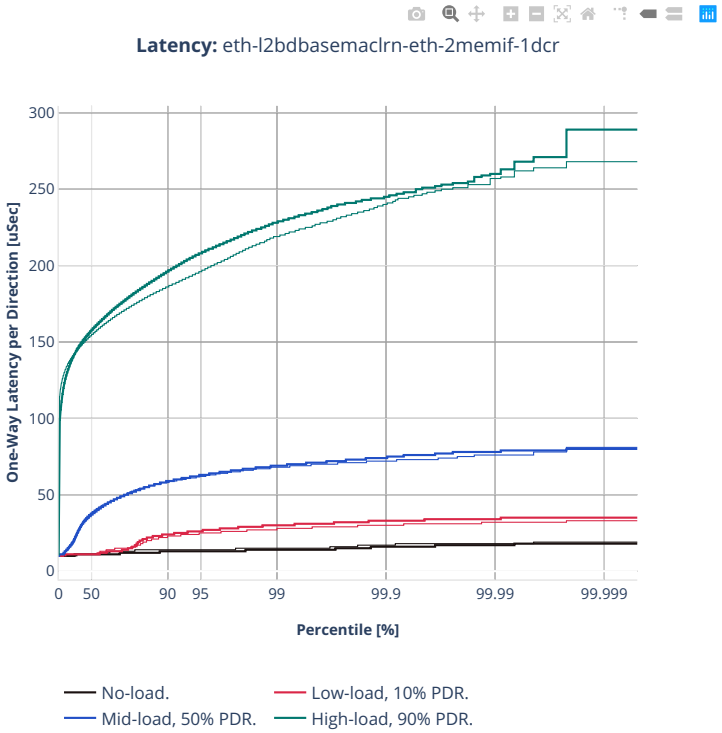
Latency: eth-l2bdbasemaclrn-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

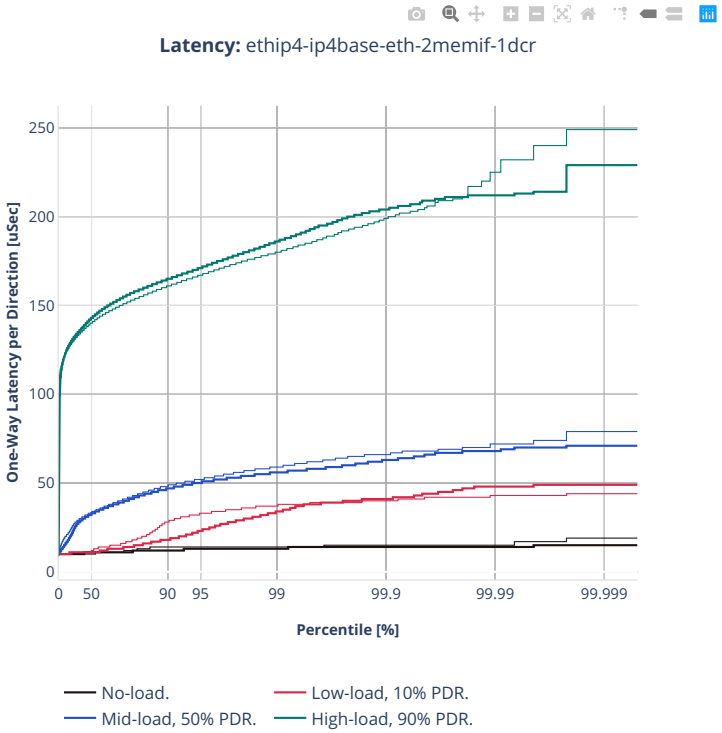
Latency: ethip4-ip4base-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-tx2-xl710

64b-1t1c-memif-base-dpdk

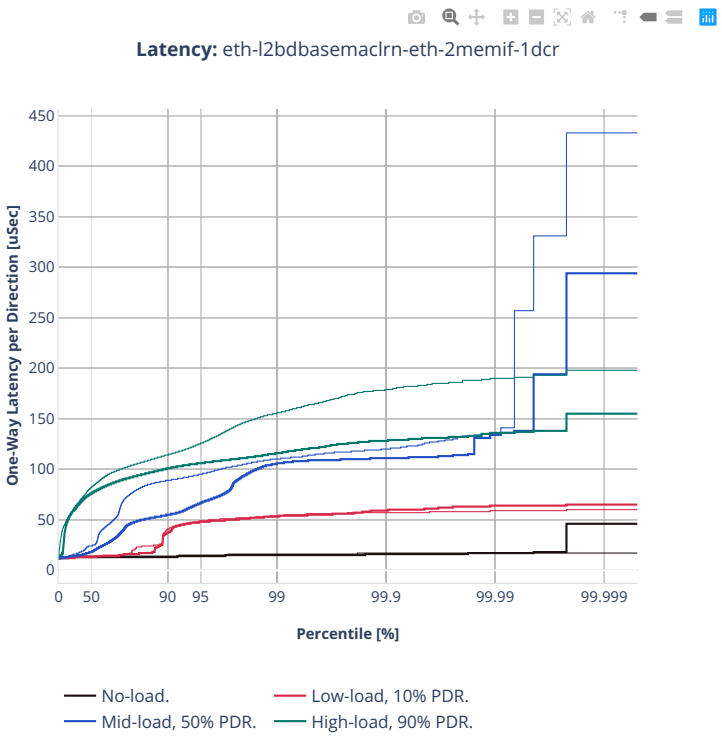

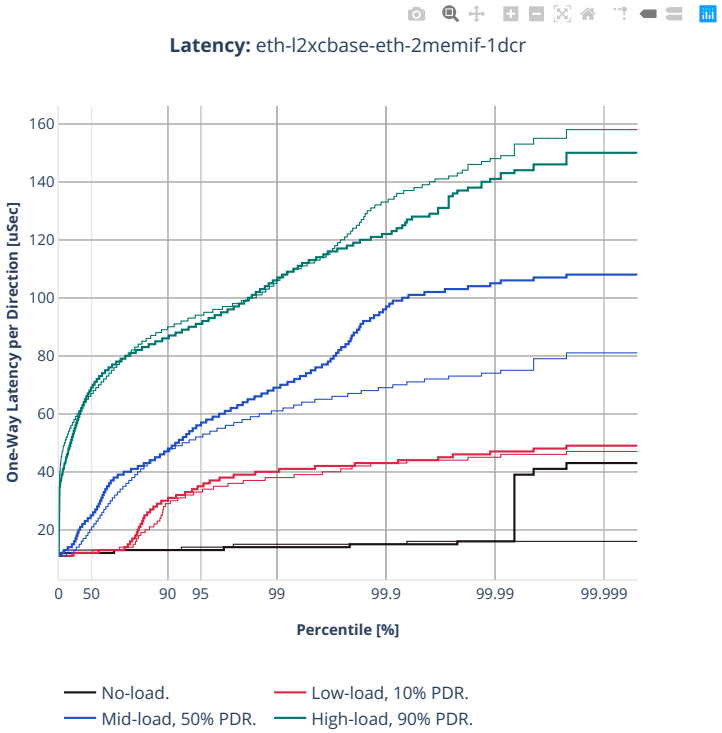
Latency: eth-l2xcbase-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

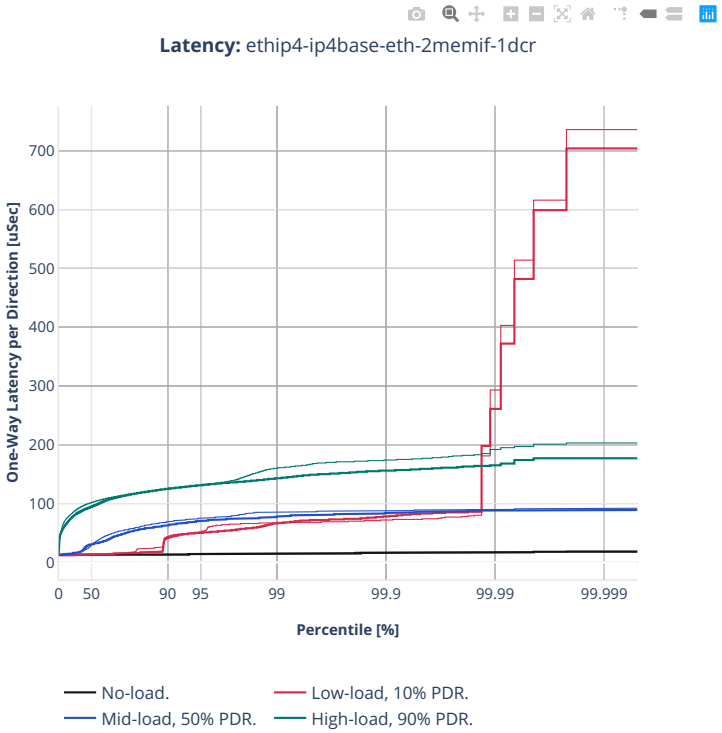
Latency: ethip4-ip4base-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2n-zn2-xxv710

64b-2t1c-memif-base-avf

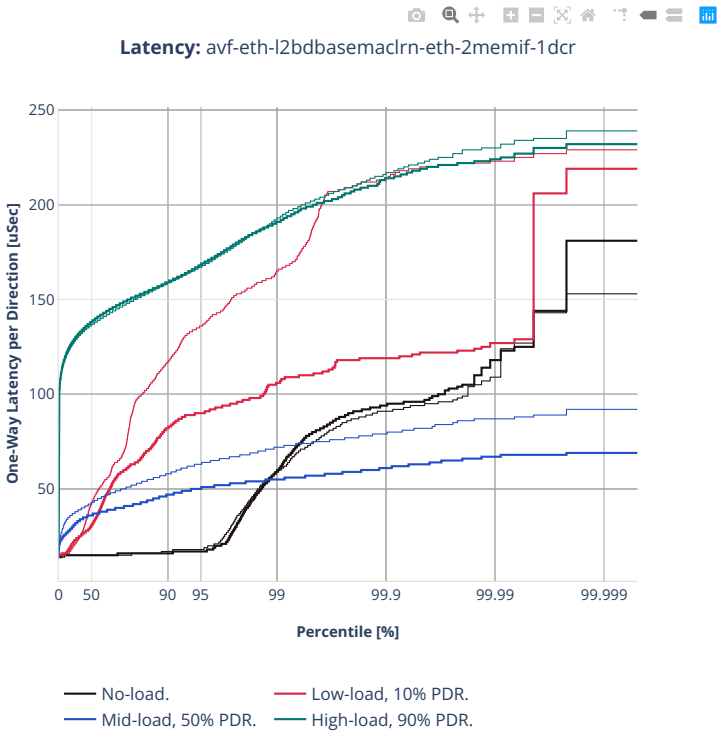

Latency: avf-eth-l2xcbase-eth-2memif-1dcr

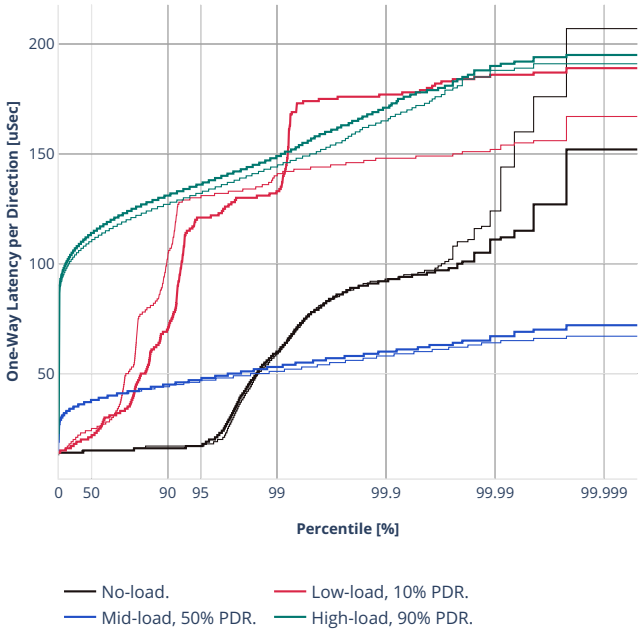

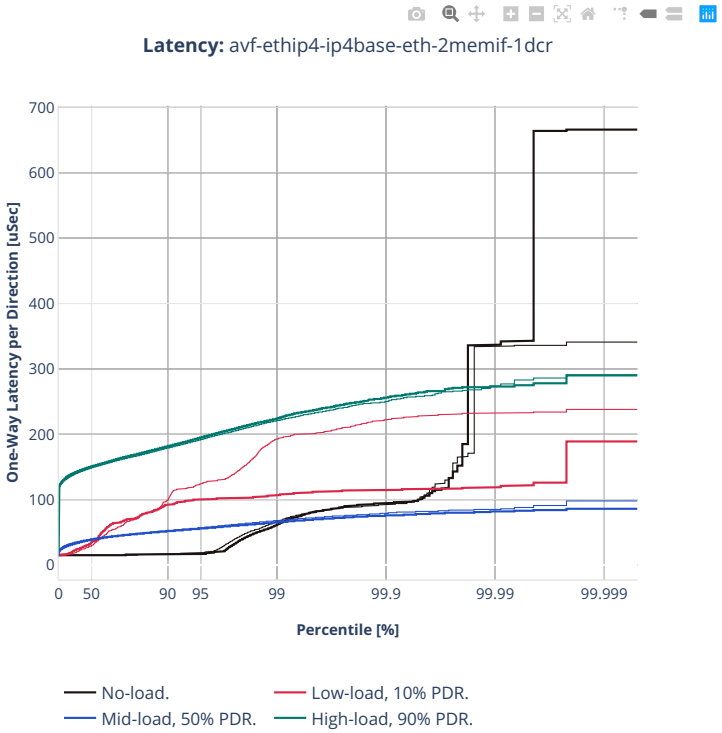
Latency: avf-ethip4-ip4base-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-2t1c-memif-base-dpdk

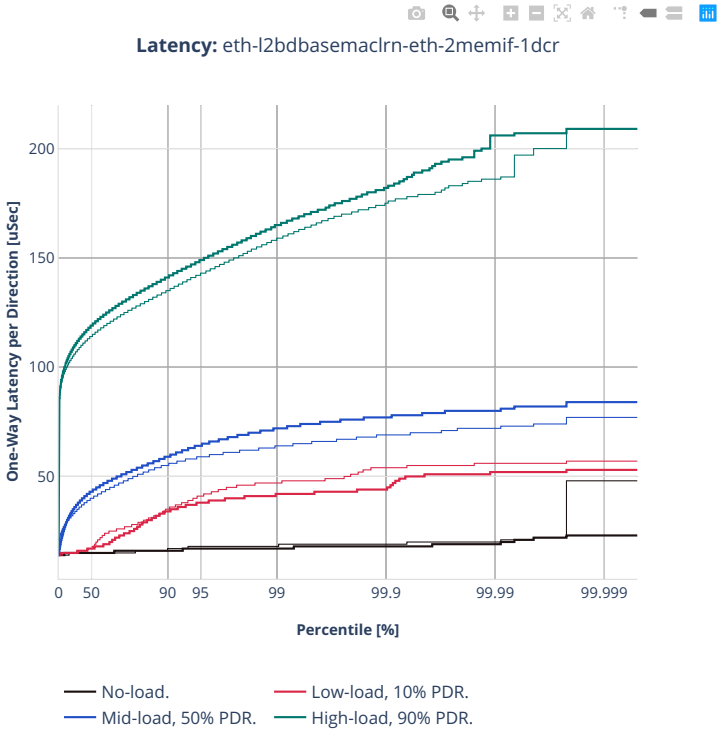

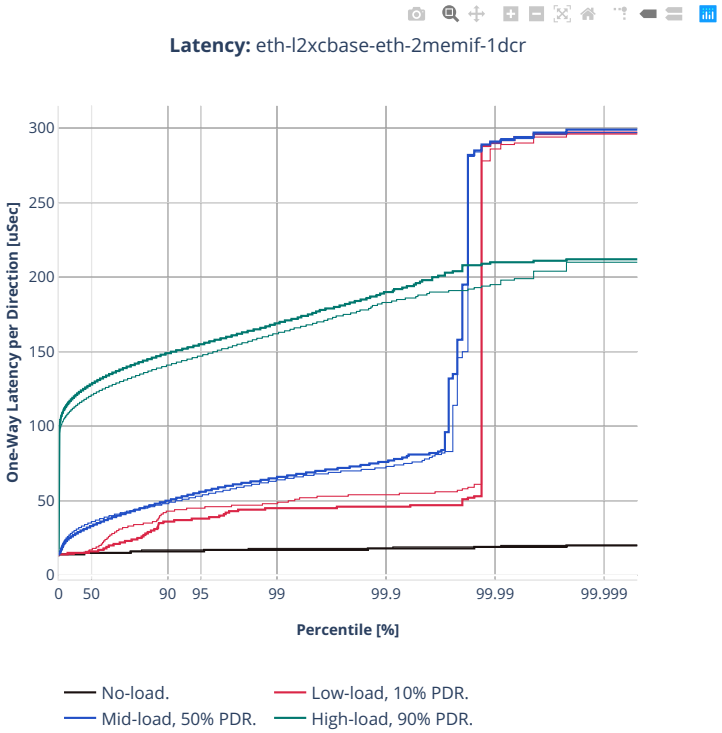
Latency: eth-l2xcbase-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

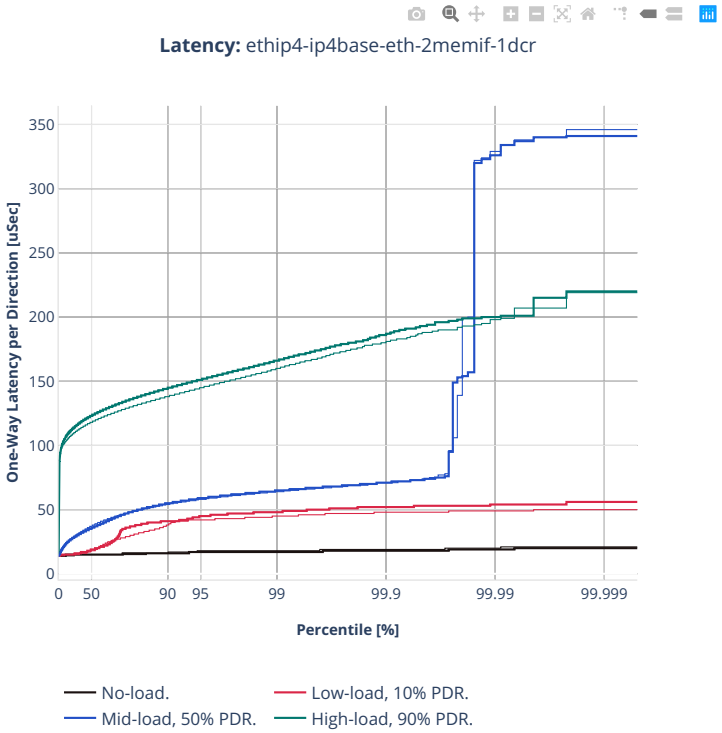
Latency: ethip4-ip4base-eth-2memif-1dcr
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2.5.9 IPSec IPv4 Routing

CSIT source code for the test cases used for plots can be found in CSIT git repository[157].

[157] https://git.fd.io/csit/tree/tests/vpp/perf/crypto?h=rls2206

3n-icx-xxv710

1518b-2t1c-ipsec-ip4routing-scale-sw-avf

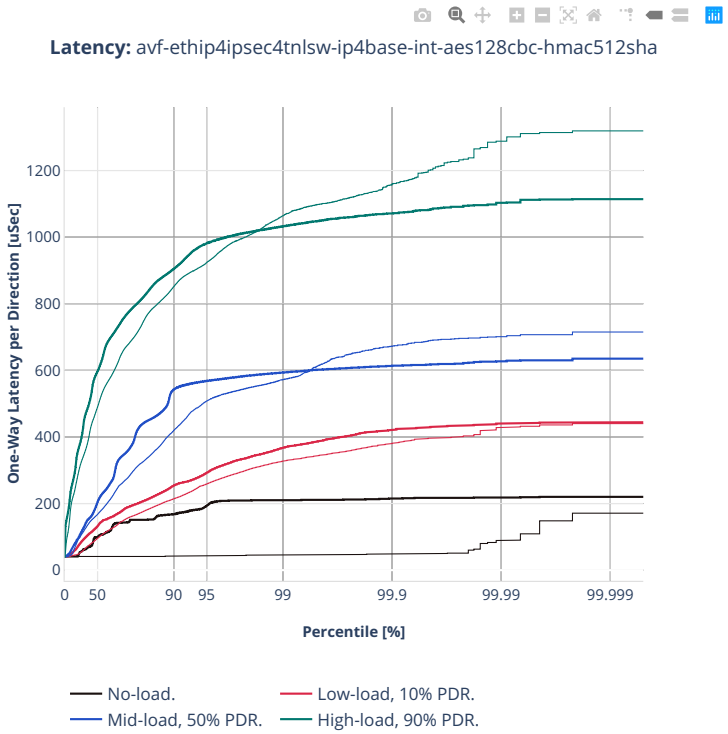

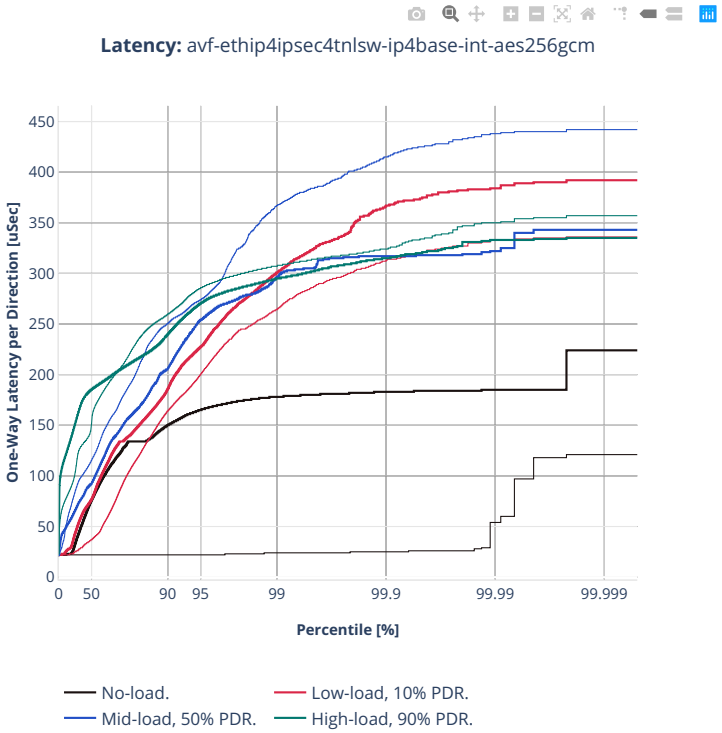
Latency: avf-ethip4ipsec4tnlsw-ip4base-int-aes256gcm
One-Way Latency per Direction [uSec]
450
400
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

Latency: avf-ethip4ipsec40tnlsw-ip4base-int-aes128cbc-hmac512sha

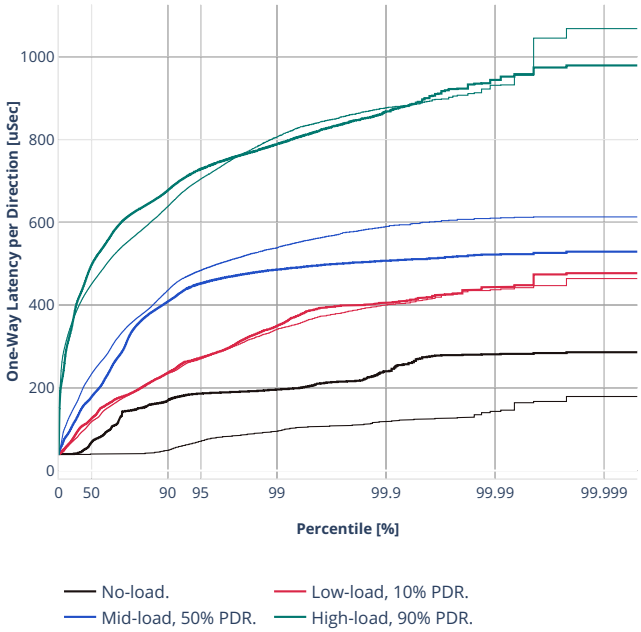

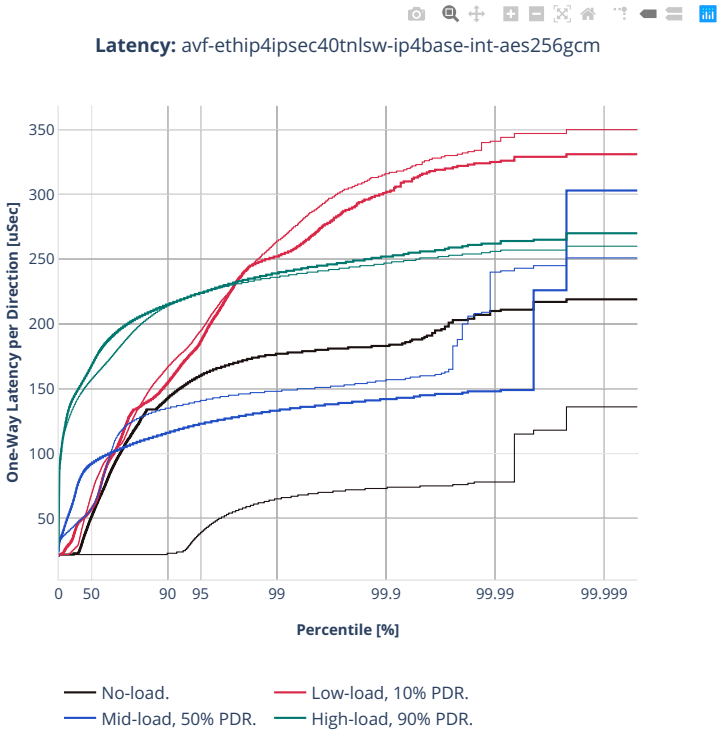
Latency: avf-ethip4ipsec40tnlsw-ip4base-int-aes256gcm
One-Way Latency per Direction [uSec]
350
300
250
200
150
100
50
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

Latency: avf-ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha

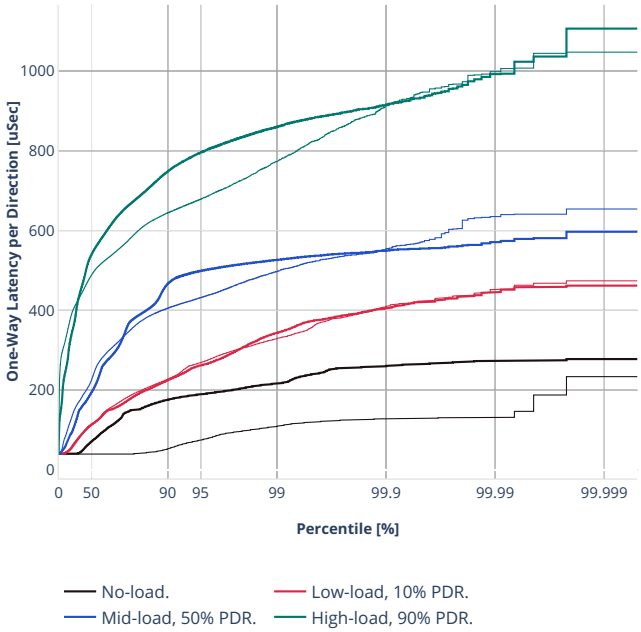


Latency: avf-ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
One-Way Latency per Direction [uSec]
400
350
300
250
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

Latency: avf-ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha

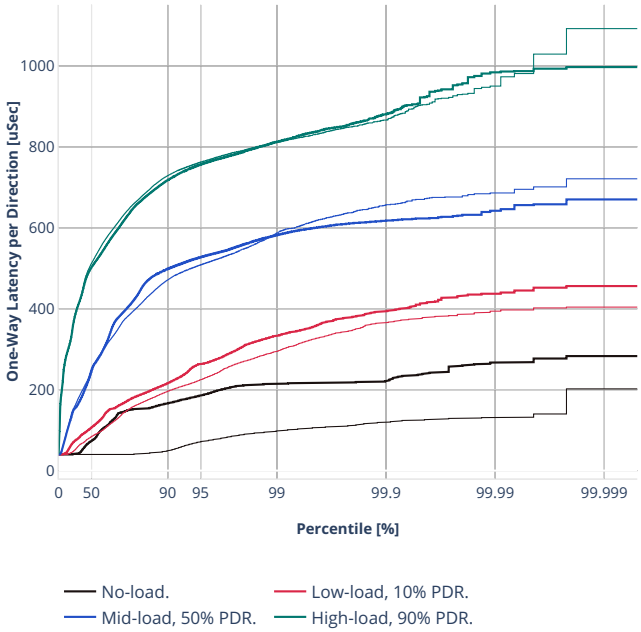

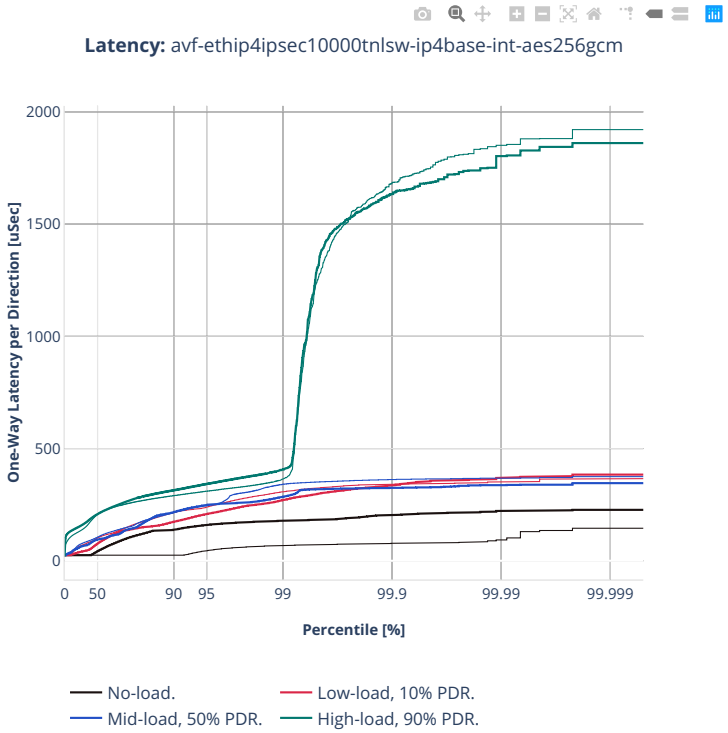
Latency: avf-ethip4ipsec10000tnlsw-ip4base-int-aes256gcm
One-Way Latency per Direction [uSec]
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

1518b-2t1c-ipsec-ip4routing-scale-sw-dpdk

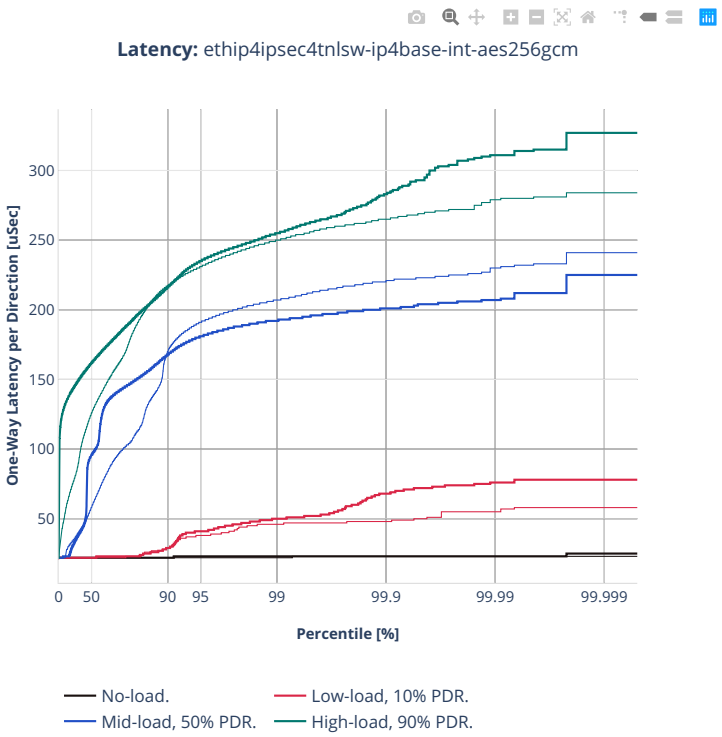


Latency: ethip4ipsec40tnlsw-ip4base-int-aes256gcm
One-Way Latency per Direction [uSec]
300
250
200
150
100
50
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

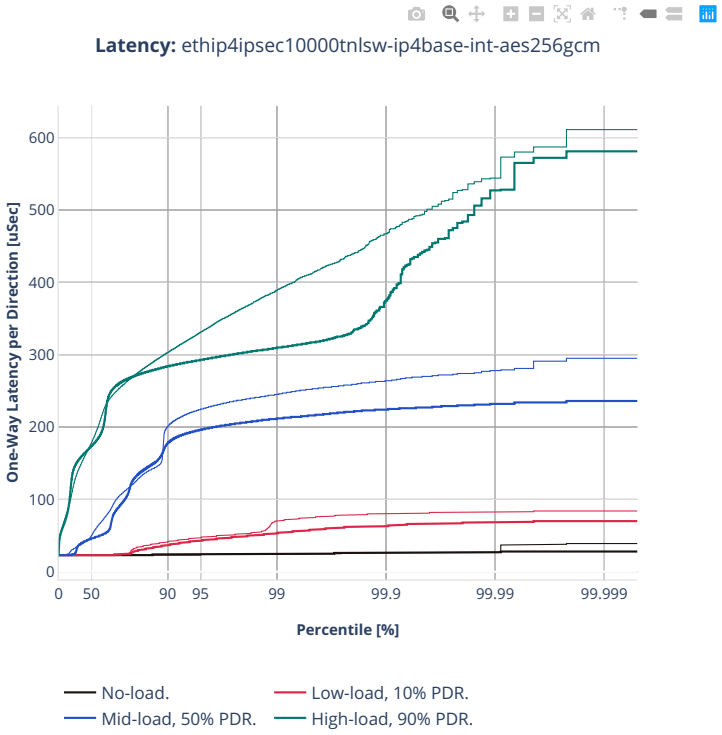
Latency: ethip4ipsec10000tnlsw-ip4base-int-aes256gcm
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-skx-xxv710

1518b-2t1c-ipsec-ip4routing-base-scale-sw-avf

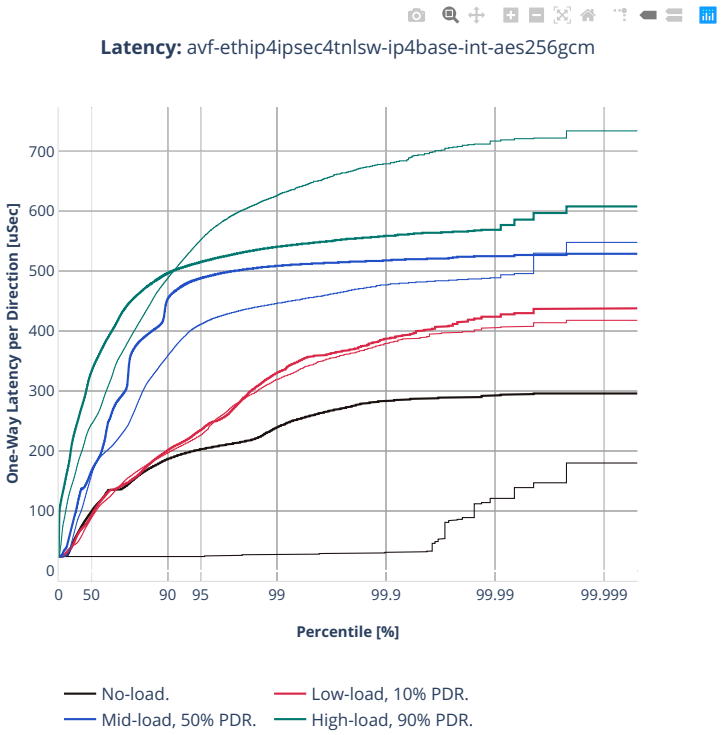

Latency: avf-ethip4ipsec1000tnlsw-ip4base-int-aes256gcm

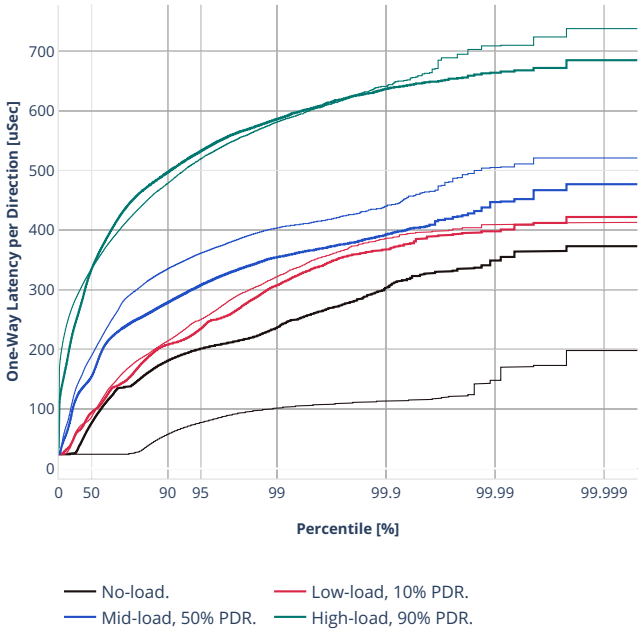

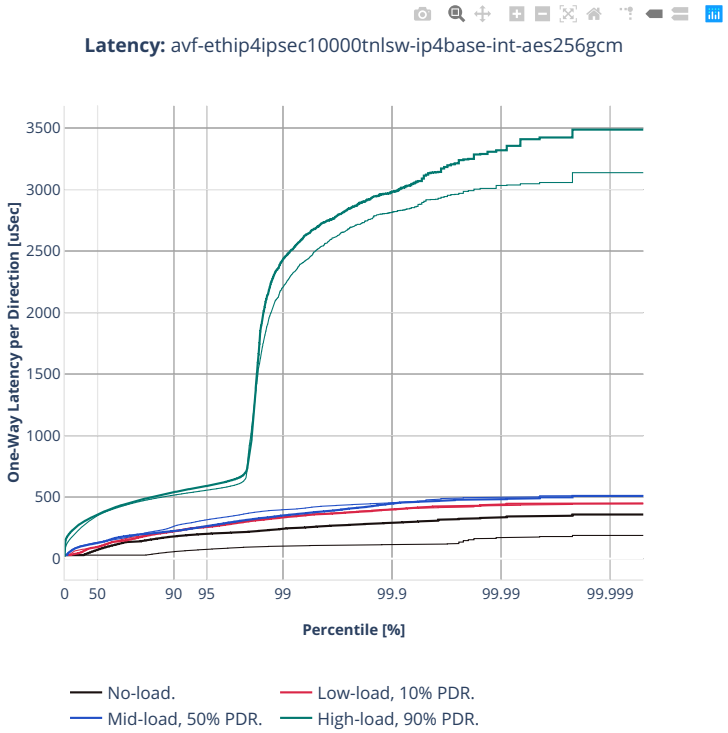
Latency: avf-ethip4ipsec10000tnlsw-ip4base-int-aes256gcm
One-Way Latency per Direction [uSec]
3500
3000
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

Latency: avf-ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha

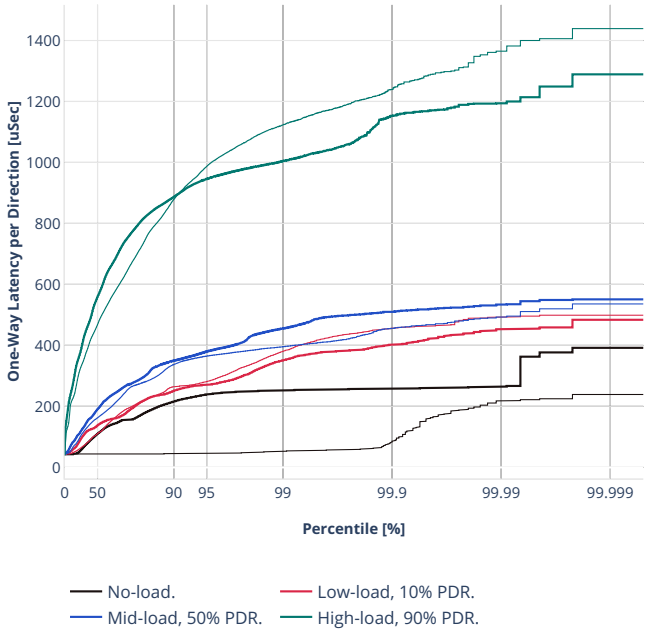

Latency: avf-ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha

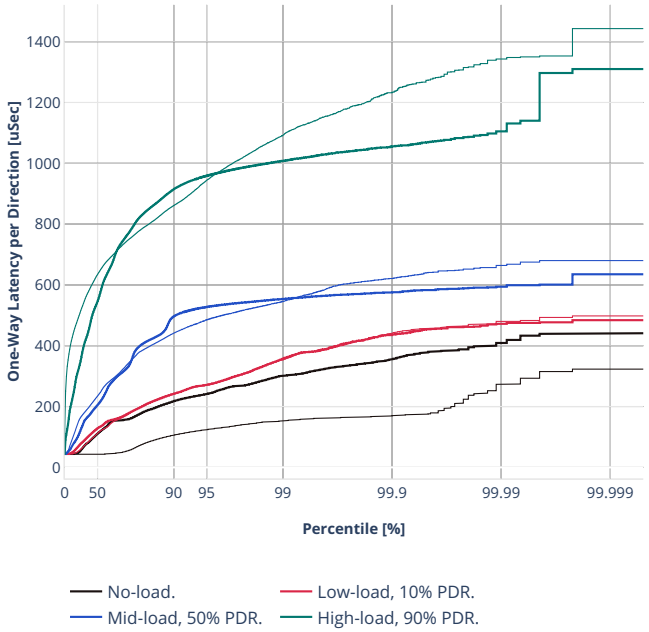

Latency: avf-ethip4ipsec10000tnlsw-ip4base-int-aes128cbc-hmac512sha

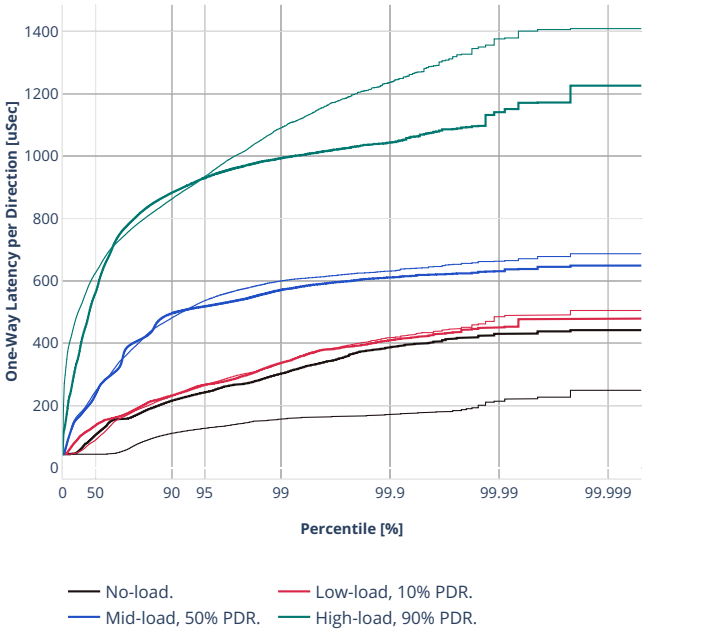

1518b-2t1c-ipsec-ip4routing-base-scale-sw-dpdk

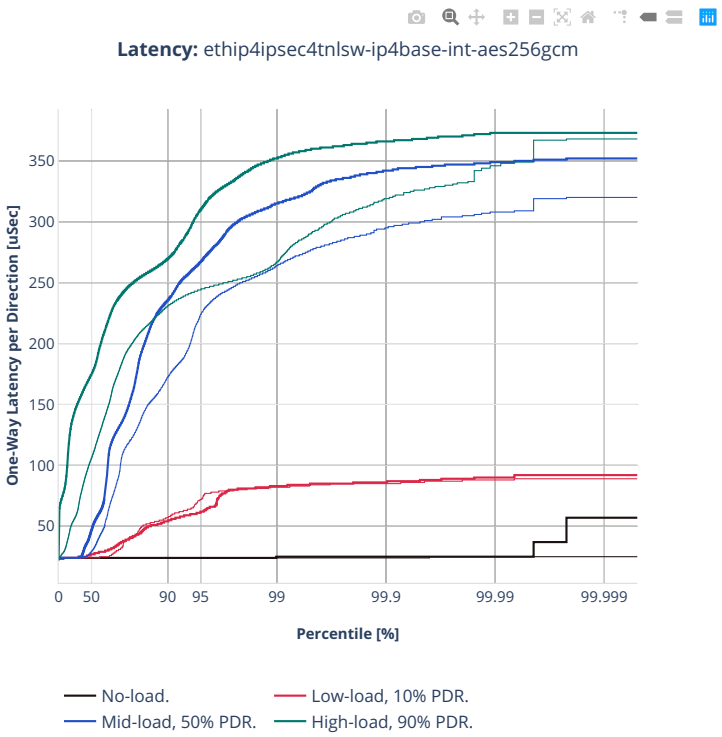


Latency: ethip4ipsec10000tnlsw-ip4base-int-aes256gcm
One-Way Latency per Direction [uSec]
900
800
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3n-alt-xl710

1518b-1t1c-ipsec-ip4routing-base-scale


Latency: ethip4ipsec10000tnlsw-ip4base-int-aes128gcm
One-Way Latency per Direction [uSec]
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.


Latency: ethip4ipsec10000tnlsw-ip4base-int-aes256gcm
One-Way Latency per Direction [uSec]
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

Latency: ethip4ipsec1000tnlsw-ip4base-int-aes128cbc-hmac512sha

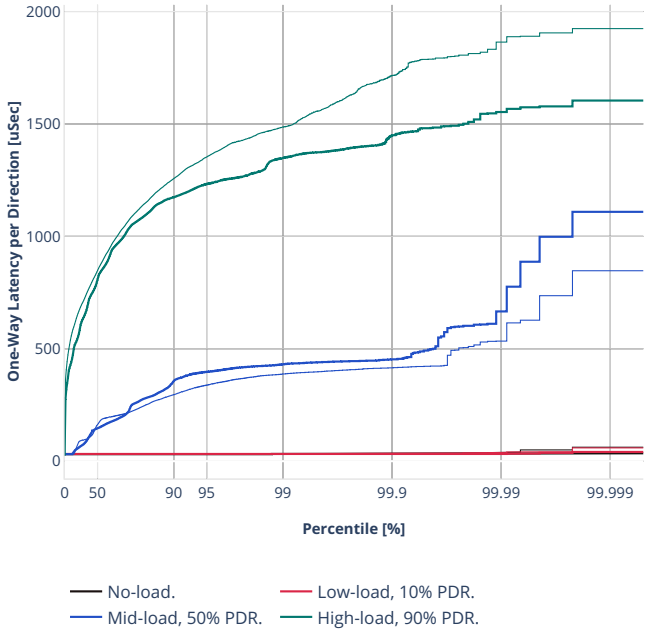

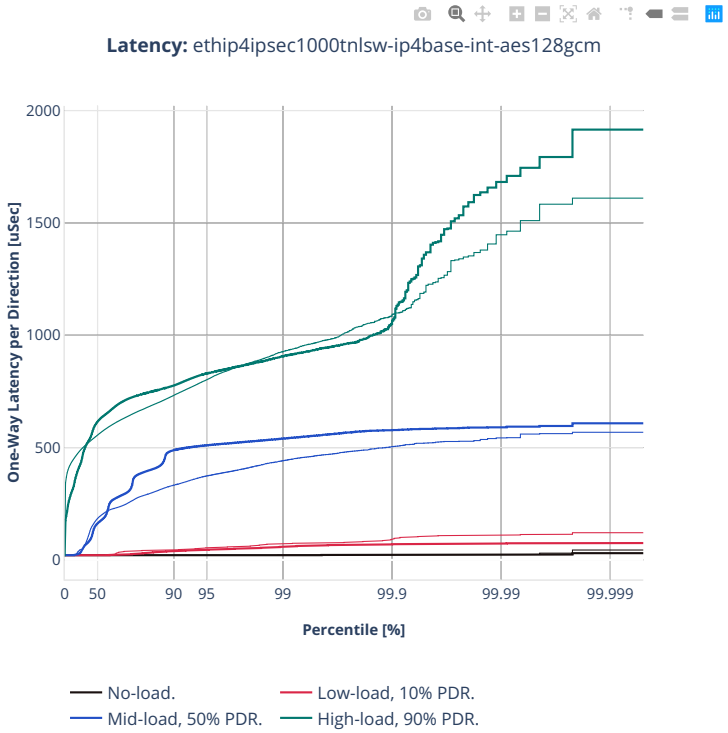
Latency: ethip4ipsec1000tnlsw-ip4base-int-aes128gcm
One-Way Latency per Direction [uSec]
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

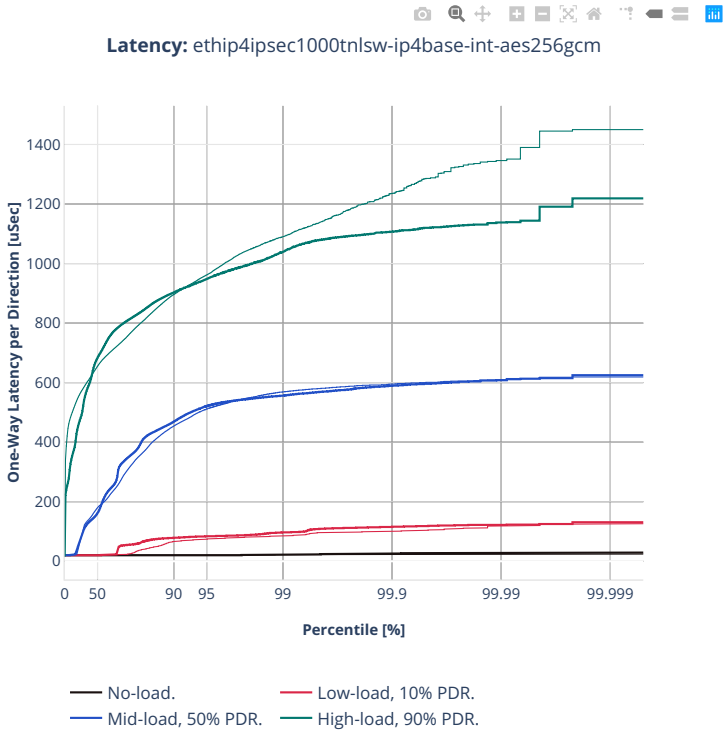
Latency: ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
One-Way Latency per Direction [uSec]
1400
1200
1000
800
600
400
200
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.


Latency: ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha
One-Way Latency per Direction [uSec]
1600
1400
1200
1000
800
600
400
200
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

Latency: ethip4ipsec4tnlsw-ip4base-int-aes128gcm

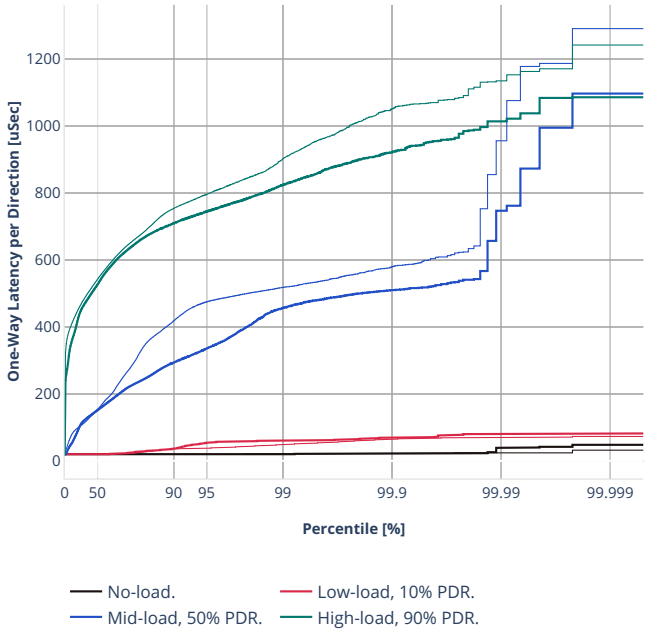

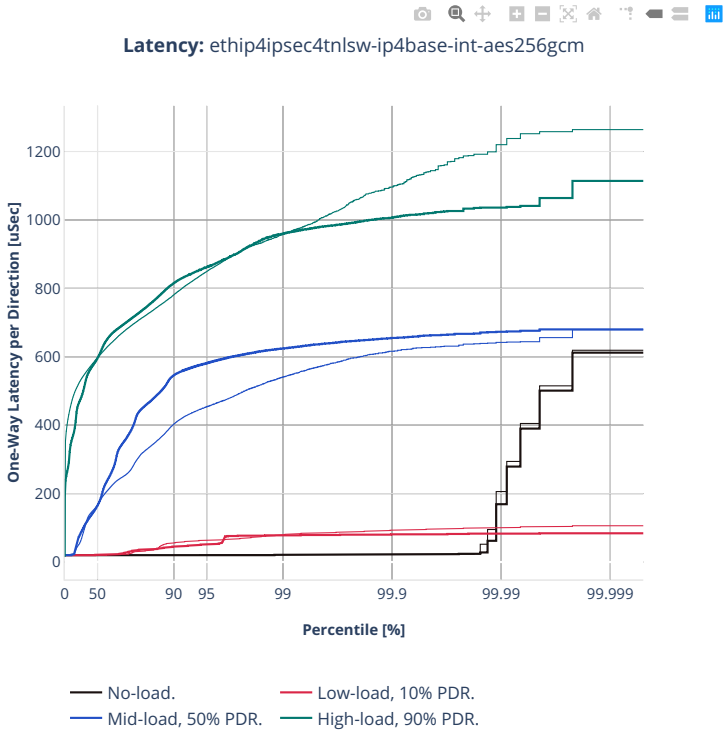
Latency: ethip4ipsec4tnlsw-ip4base-int-aes256gcm
One-Way Latency per Direction [uSec]
1200
1000
800
600
400
200
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

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

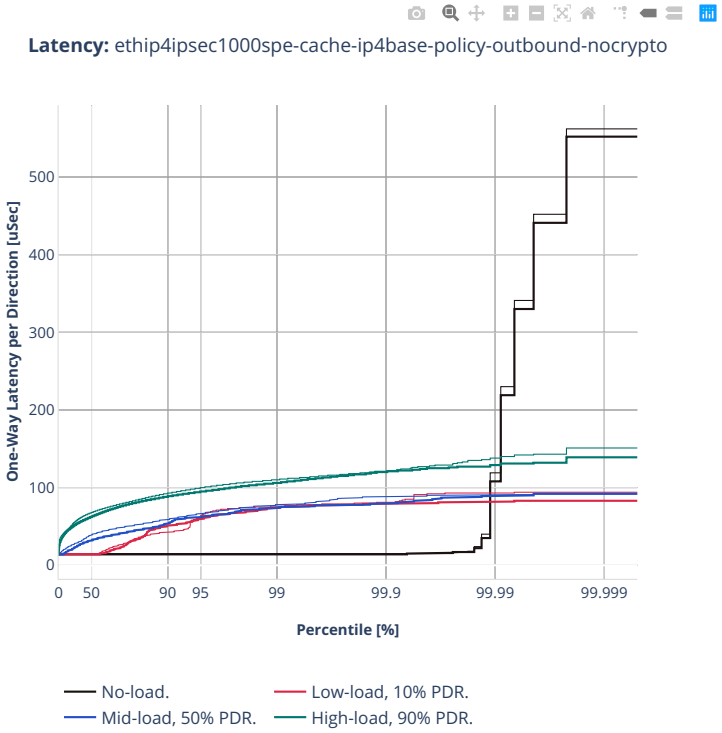

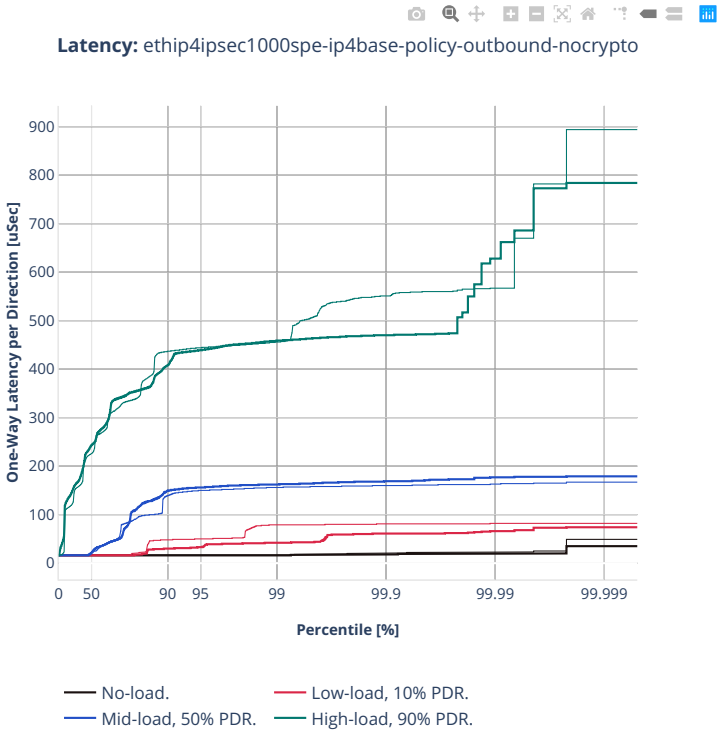
Latency: ethip4ipsec1000spe-ip4base-policy-outbound-nocrypto
One-Way Latency per Direction [uSec]
900
800
700
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.


Latency: ethip4ipsec1000tnlsw-ip4base-policy-aes256gcm
One-Way Latency per Direction [uSec]
2500
2000
1500
1000
500
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

Latency: ethip4ipsec100spe-cache-ip4base-policy-outbound-nocrypto

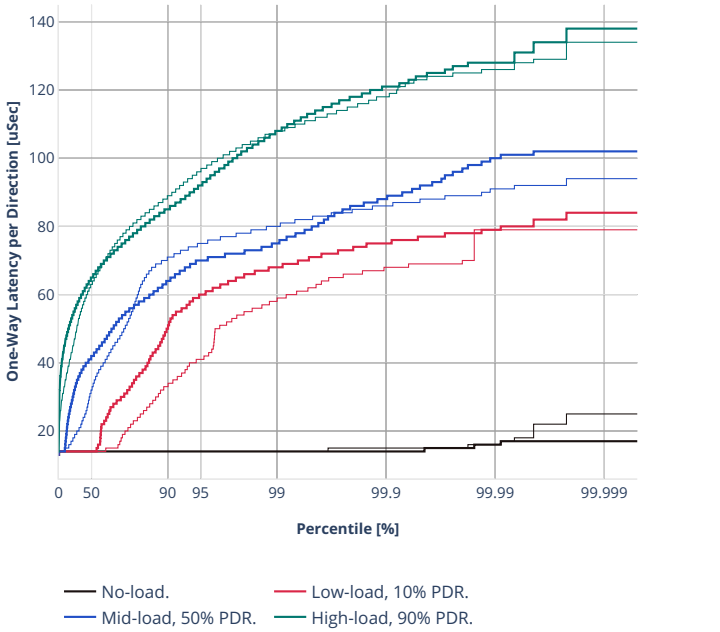

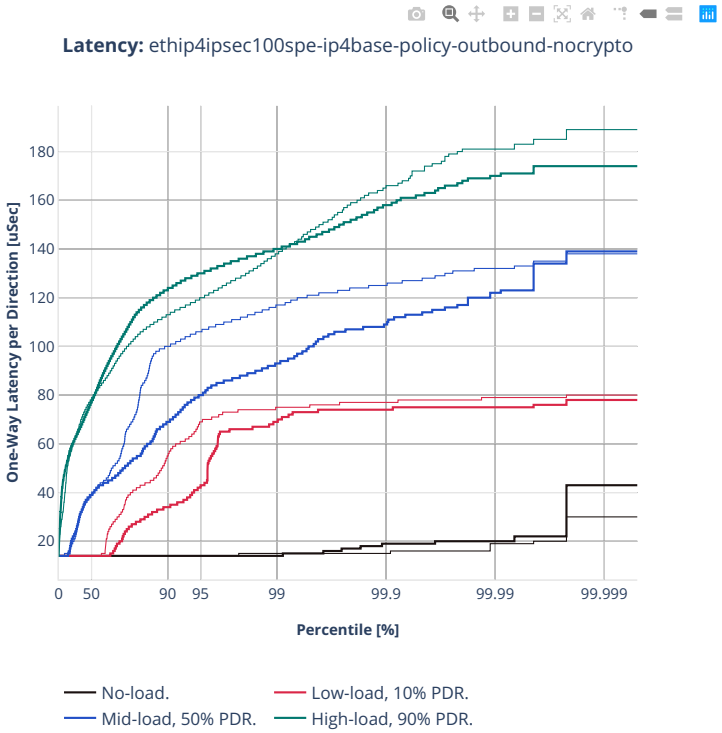
Latency: ethip4ipsec100spe-ip4base-policy-outbound-nocrypto
One-Way Latency per Direction [uSec]
180
160
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

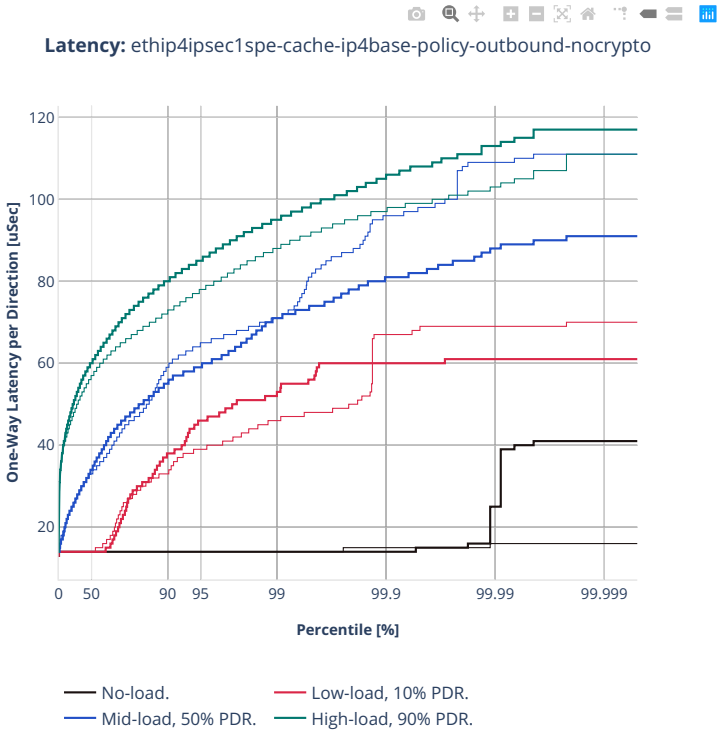
Latency: ethip4ipsec1spe-cache-ip4base-policy-outbound-nocrypto
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

Latency: ethip4ipsec1spe-ip4base-policy-outbound-nocrypto

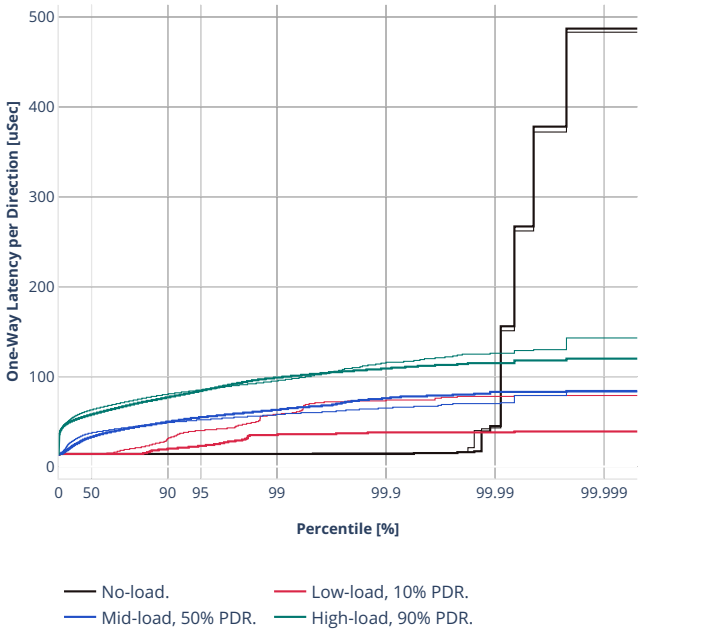

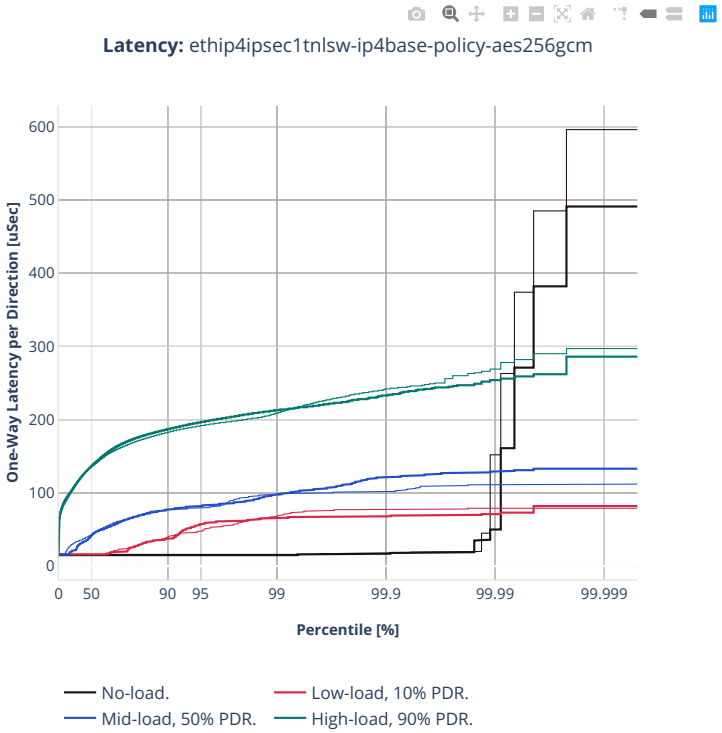
Latency: ethip4ipsec1tnlsw-ip4base-policy-aes256gcm
One-Way Latency per Direction [uSec]
600
500
400
300
200
100
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

Latency: ethip4ipsec40tnlsw-ip4base-policy-aes256gcm

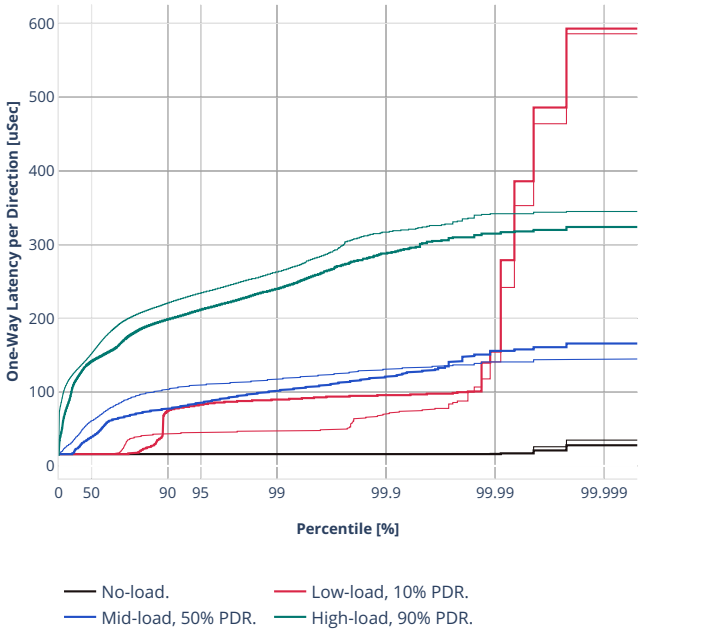

3n-tsh-x520

1518b-1t1c-ipsec-ip4routing-base-scale-sw-ixgbe

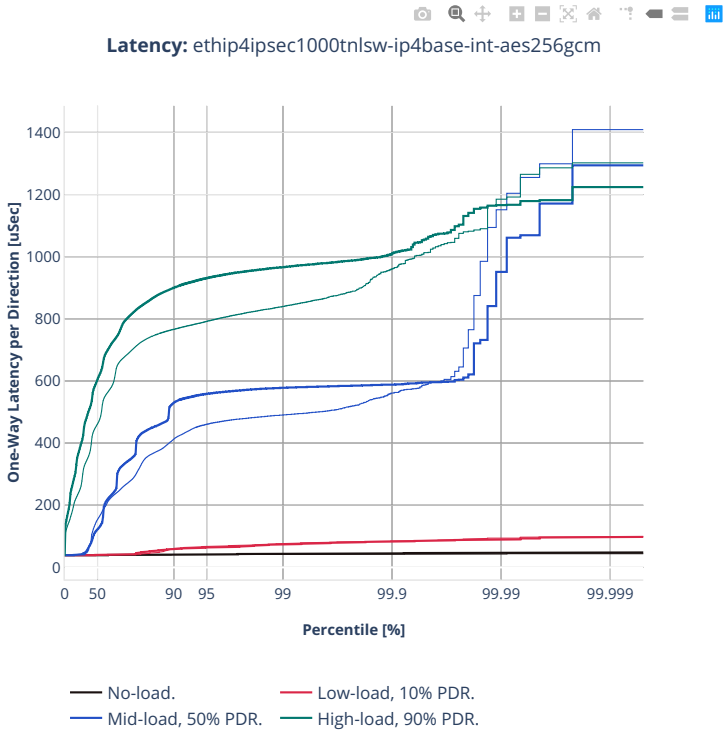
Latency: ethip4ipsec1000tnlsw-ip4base-int-aes256gcm
One-Way Latency per Direction [uSec]
1400
1200
1000
800
600
400
200
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.


Latency: ethip4ipsec10000tnlsw-ip4base-int-aes256gcm
One-Way Latency per Direction [uSec]
1600
1400
1200
1000
800
600
400
200
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

2.6 Soak Tests

Long duration (30 minutes per test) soak tests are executed using *PLRsearch* (page 28) algorithm. As the tests take long time, only 12 test cases were executed, two runs each.

Additional information about graph data:

1. **Graph Title**: describes type of tests and soak test duration.
2. **X-axis Labels**: indices of test suites.
3. **Y-axis Labels**: estimated lower bounds for critical rate value in [Mpps].
4. **Graph Legend**: list of X-axis indices with CSIT test cases.
5. **Hover Information**: in general lists minimum, first quartile, median, third quartile, and maximum. As only two samples are used, minimum and maximum are not distinguished from quartiles.

Note: Test results are stored in build logs from FD.io vpp performance job 2n-icx[158], build logs from FD.io vpp performance job 2n-skx[159] and build logs from FD.io vpp performance job 2n-clx[160] with RF result files csit-vpp-perf-2206-*.zip archived here.

[158] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-icx
[159] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx
[160] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-clx

2.6.1 2n-icx

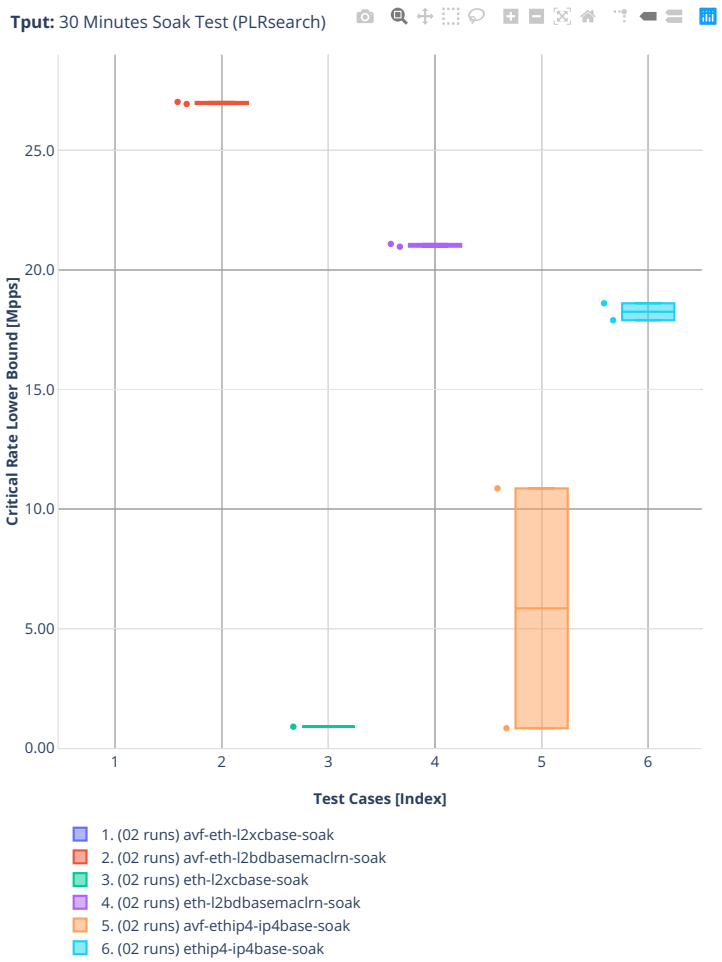
Tput: 30 Minutes Soak Test (PLRsearch)
Critical Rate Lower Bound [Mpps]
25.0
20.0
15.0
10.0
5.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1. (02 runs) avf-eth-l2xcbase-soak
2. (02 runs) avf-eth-l2bdbasemaclrn-soak
3. (02 runs) eth-l2xcbase-soak
4. (02 runs) eth-l2bdbasemaclrn-soak
5. (02 runs) avf-ethip4-ip4base-soak
6. (02 runs) ethip4-ip4base-soak

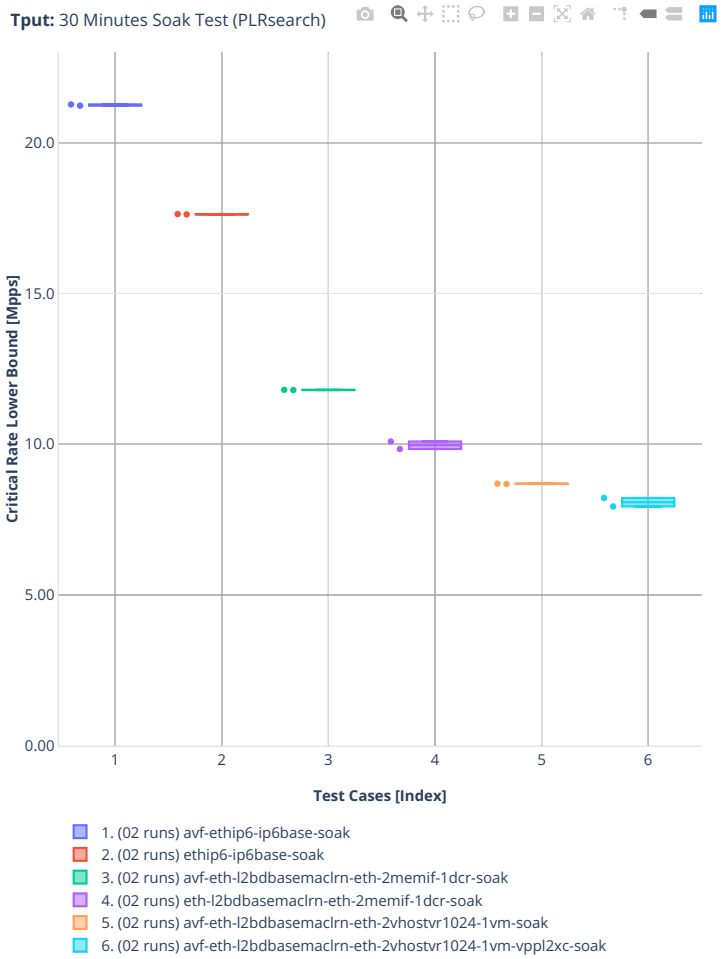
Tput: 30 Minutes Soak Test (PLRsearch)
Critical Rate Lower Bound [Mpps]
20.0
15.0
10.0
5.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1. (02 runs) avf-ethip6-ip6base-soak
2. (02 runs) ethip6-ip6base-soak
3. (02 runs) avf-eth-l2bdbasemaclrn-eth-2memif-1dcr-soak
4. (02 runs) eth-l2bdbasemaclrn-eth-2memif-1dcr-soak
5. (02 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-soak
6. (02 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc-soak

2.6.2 2n-clx

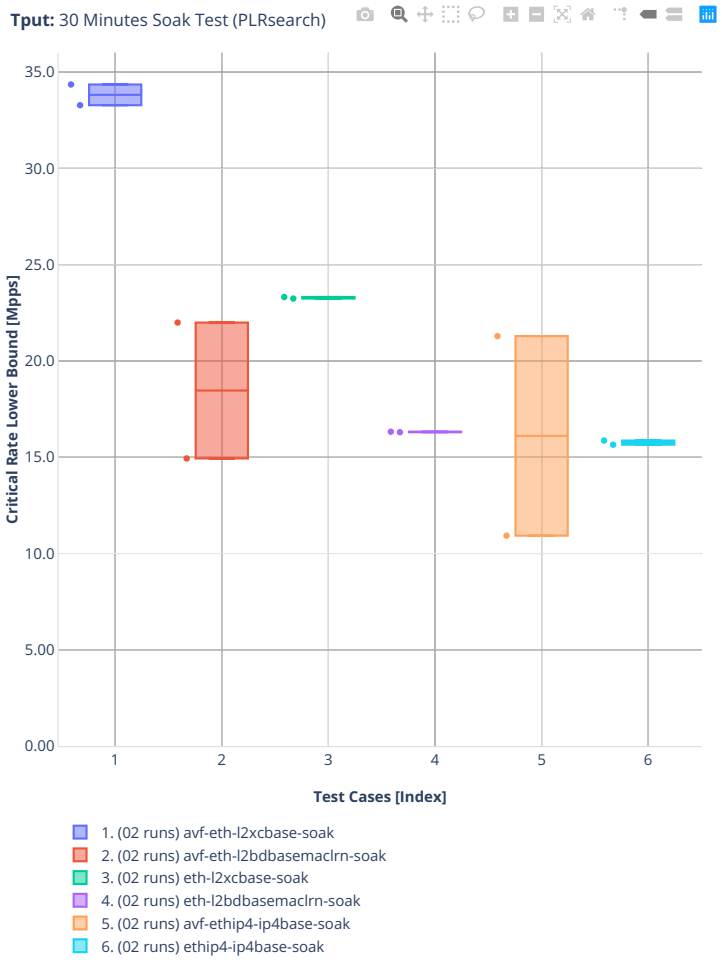
Tput: 30 Minutes Soak Test (PLRsearch)
Critical Rate Lower Bound [Mpps]
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1. (02 runs) avf-eth-l2xcbase-soak
2. (02 runs) avf-eth-l2bdbasemaclrn-soak
3. (02 runs) eth-l2xcbase-soak
4. (02 runs) eth-l2bdbasemaclrn-soak
5. (02 runs) avf-ethip4-ip4base-soak
6. (02 runs) ethip4-ip4base-soak

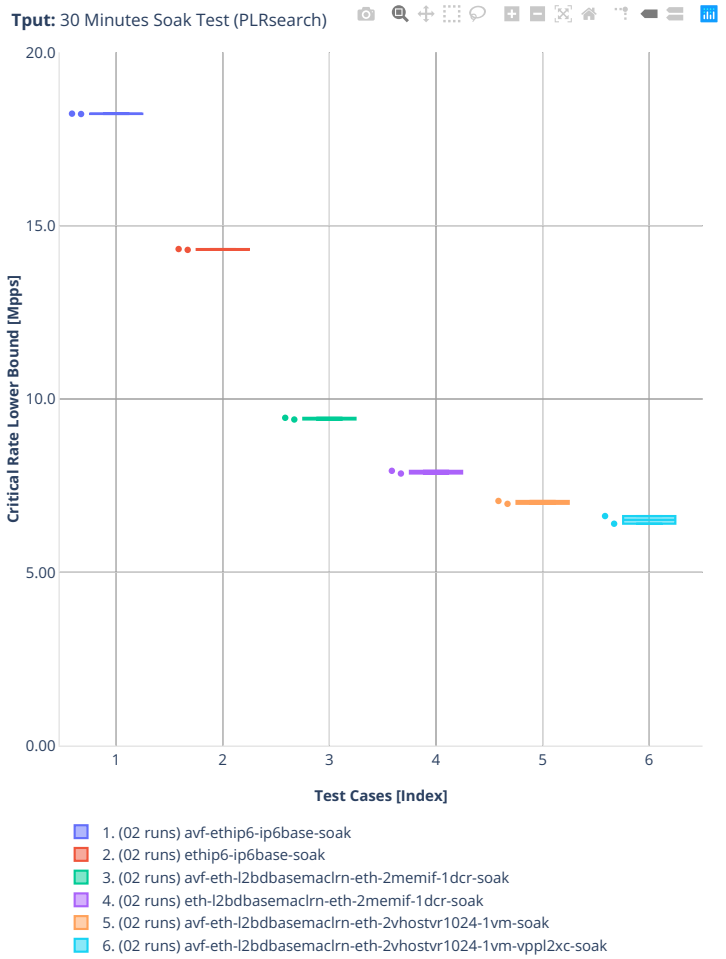
Tput: 30 Minutes Soak Test (PLRsearch)
20.0
15.0
10.0
5.00
0.00
Critical Rate Lower Bound [Mpps]
1
2
3
4
5
6
Test Cases [Index]
1. (02 runs) avf-ethip6-ip6base-soak
2. (02 runs) ethip6-ip6base-soak
3. (02 runs) avf-eth-l2bdbasemaclrn-eth-2memif-1dcr-soak
4. (02 runs) eth-l2bdbasemaclrn-eth-2memif-1dcr-soak
5. (02 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-soak
6. (02 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc-soak

2.6.3 2n-skx

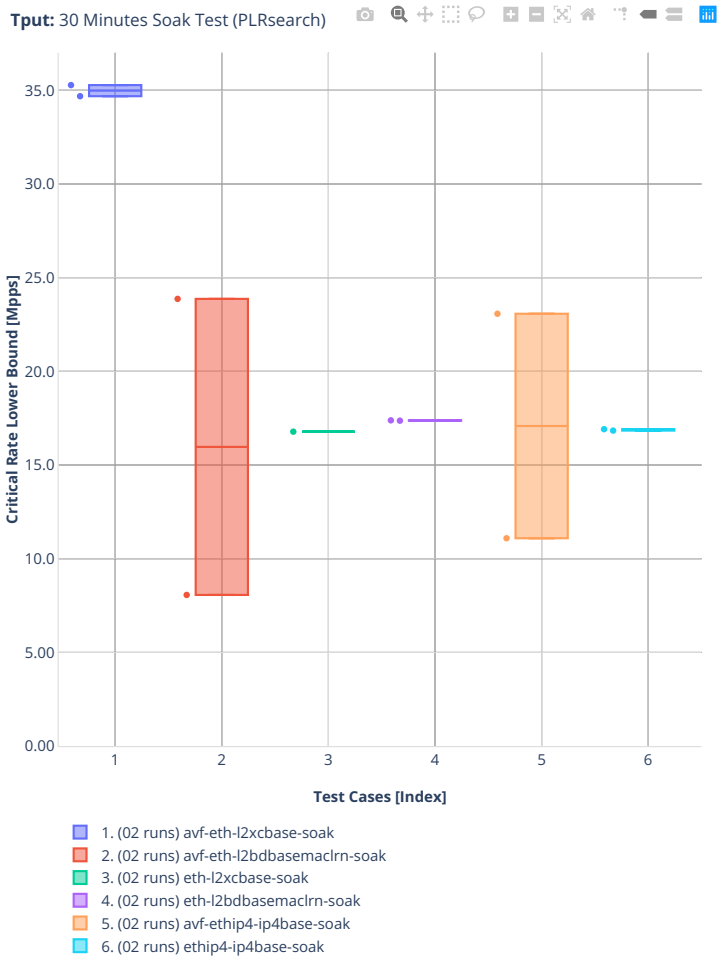
Tput: 30 Minutes Soak Test (PLRsearch)
Critical Rate Lower Bound [Mpps]
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1. (02 runs) avf-eth-l2xcbase-soak
2. (02 runs) avf-eth-l2bdbasemaclrn-soak
3. (02 runs) eth-l2xcbase-soak
4. (02 runs) eth-l2bdbasemaclrn-soak
5. (02 runs) avf-ethip4-ip4base-soak
6. (02 runs) ethip4-ip4base-soak

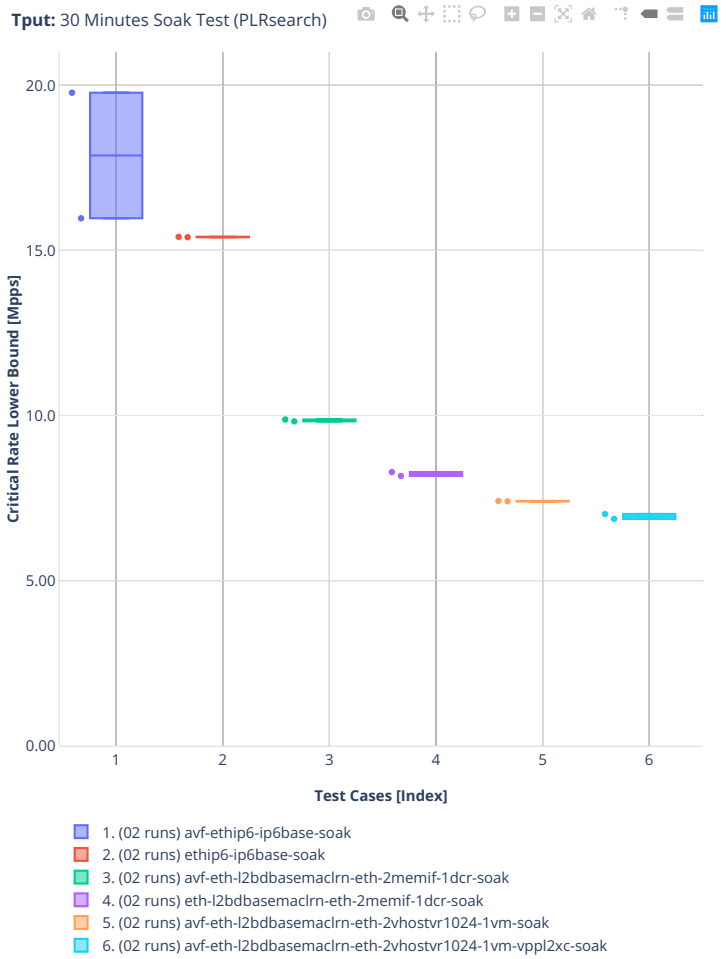
Tput: 30 Minutes Soak Test (PLRsearch)
Critical Rate Lower Bound [Mpps]
20.0
15.0
10.0
5.00
0.00
1
2
3
4
5
6
Test Cases [Index]
1. (02 runs) avf-ethip6-ip6base-soak
2. (02 runs) ethip6-ip6base-soak
3. (02 runs) avf-eth-l2bdbasemaclrn-eth-2memif-1dcr-soak
4. (02 runs) eth-l2bdbasemaclrn-eth-2memif-1dcr-soak
5. (02 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-soak
6. (02 runs) avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc-soak

2.7 Reconfiguration Tests

See *Reconfiguration Tests* (page 54) for methodology description of this test type.

2.7.1 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 `dot1qip4vlxan`,
   - VPP forwarding mode `l2bd`,
   - total number {Y} of initial service chains `{Y}ch`,
   - total number of additional chains being reconfigured `1ach`,
   - total number of initial 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 (both initial and final) 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 Effective Blocked Time [s] values.
4. **Graph Legend**: lists X-axis indices with associated CSIT test suites executed to generate graphed test results and the average value of packet loss (measured in 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 are stored in build logs from FD.io vpp performance job 2n-icx[161], build logs from FD.io vpp performance job 2n-skx[162], build logs from FD.io vpp performance job 2n-clx[163] with RF result files csit-vpp-perf-2206-*.zip archived here.

---

[161] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-icx
[162] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx
[163] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-clx

2n-icx-xxv710

imix-2t1c-dot1qip4vxlan-l2bd

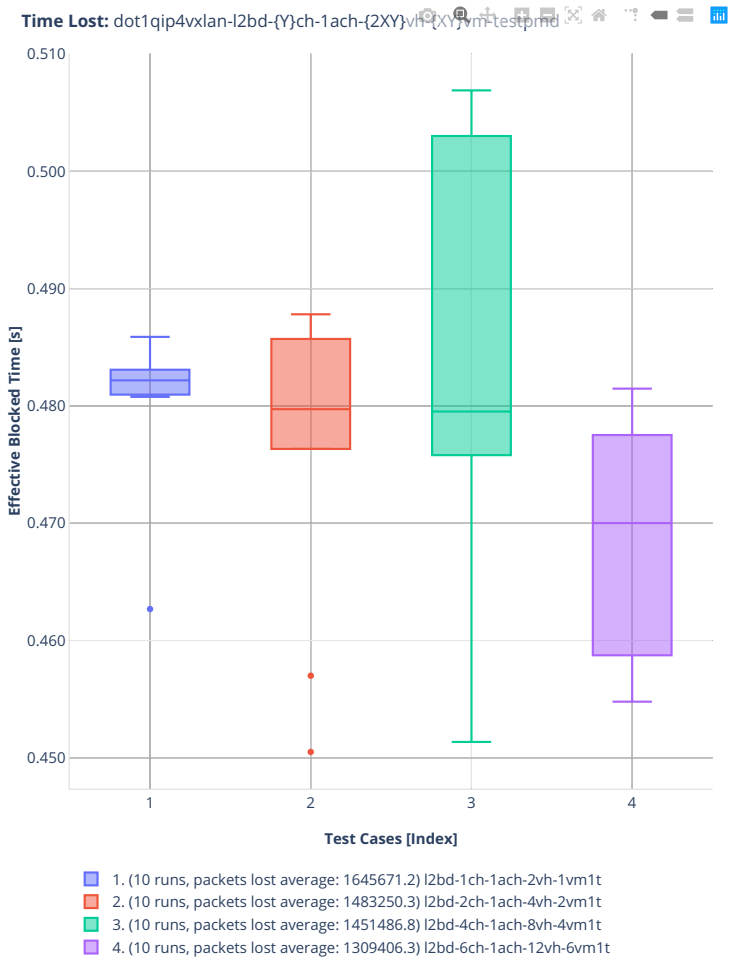

imix-4t2c-dot1qip4vxlan-l2bd

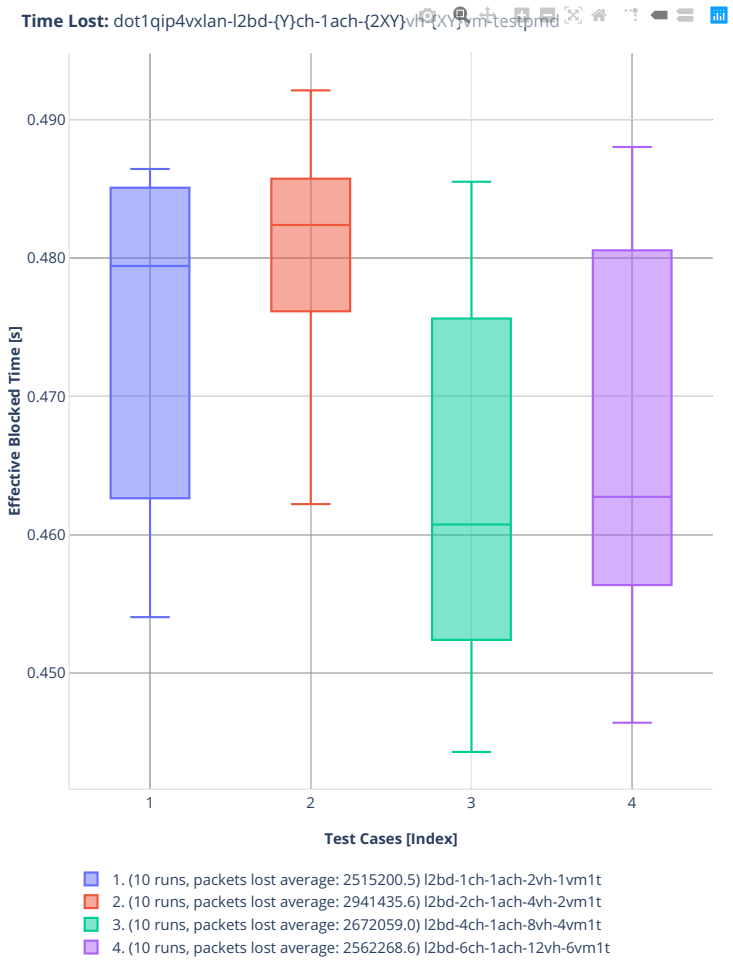

imix-8t4c-dot1qip4vxlan-l2bd

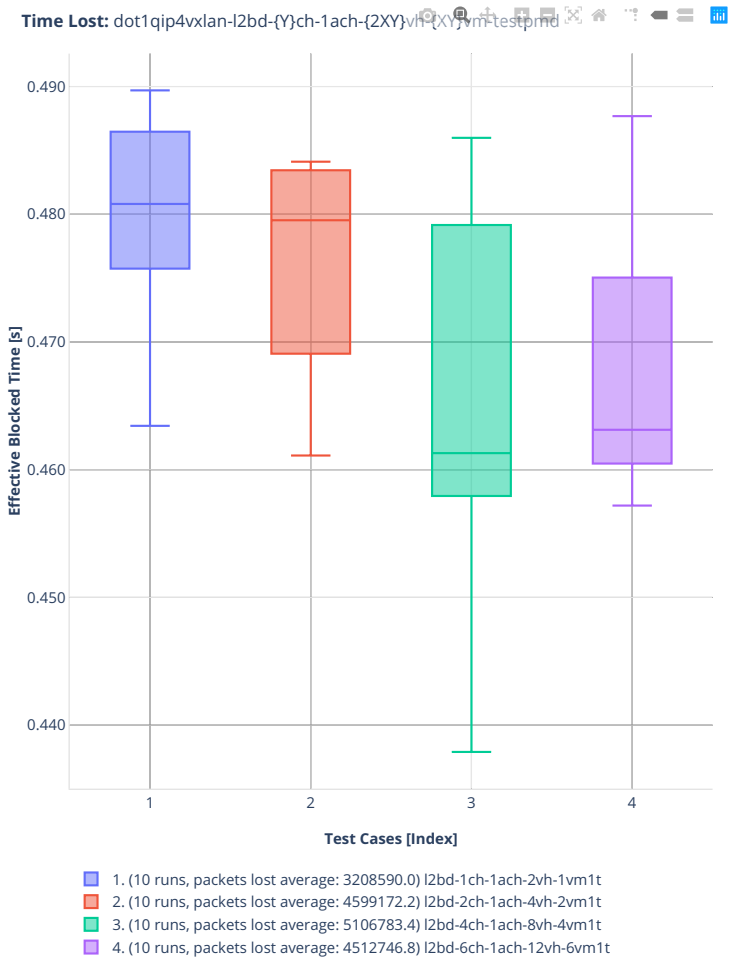

2n-skx-xxv710

imix-2t1c-dot1qip4vxlan-l2bd

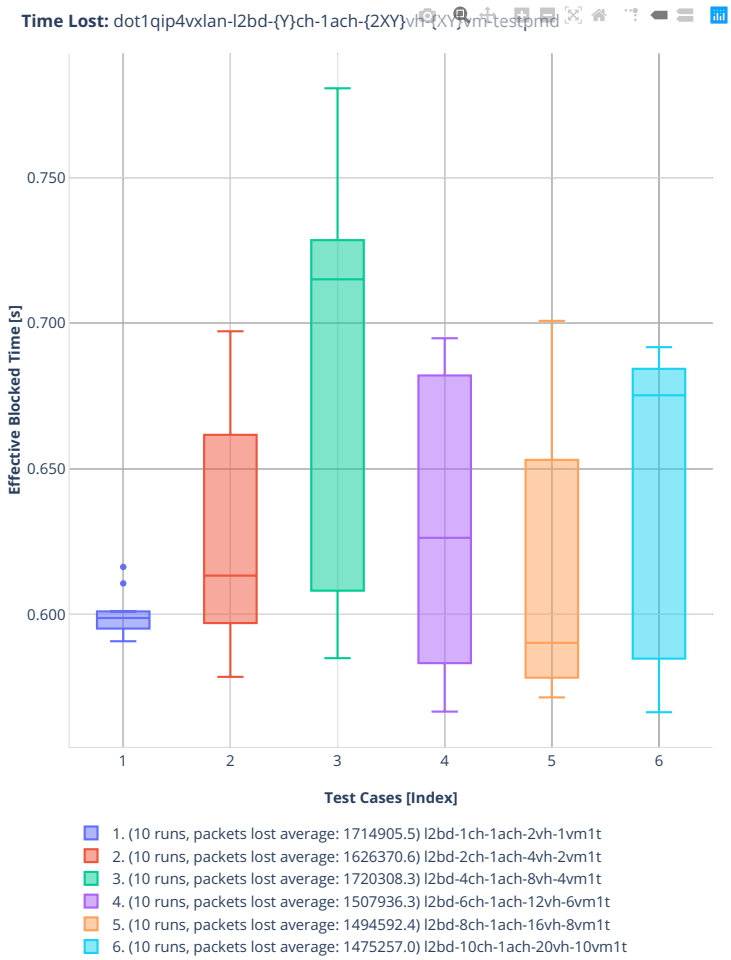

imix-4t2c-dot1qip4vxlan-l2bd

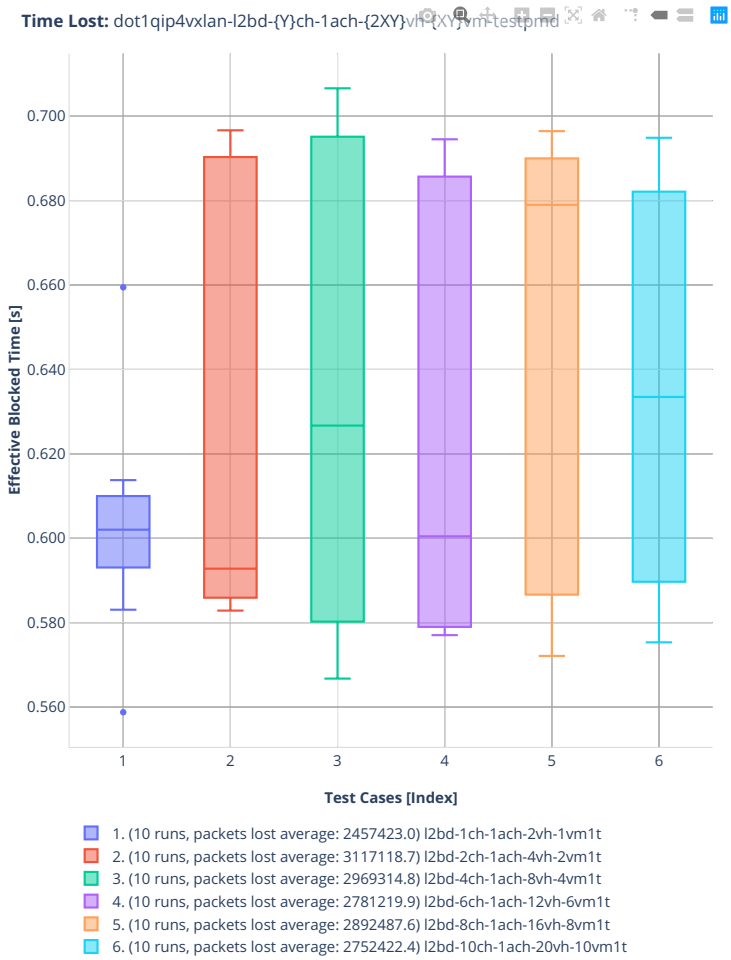

imix-8t4c-dot1qip4vxlan-l2bd

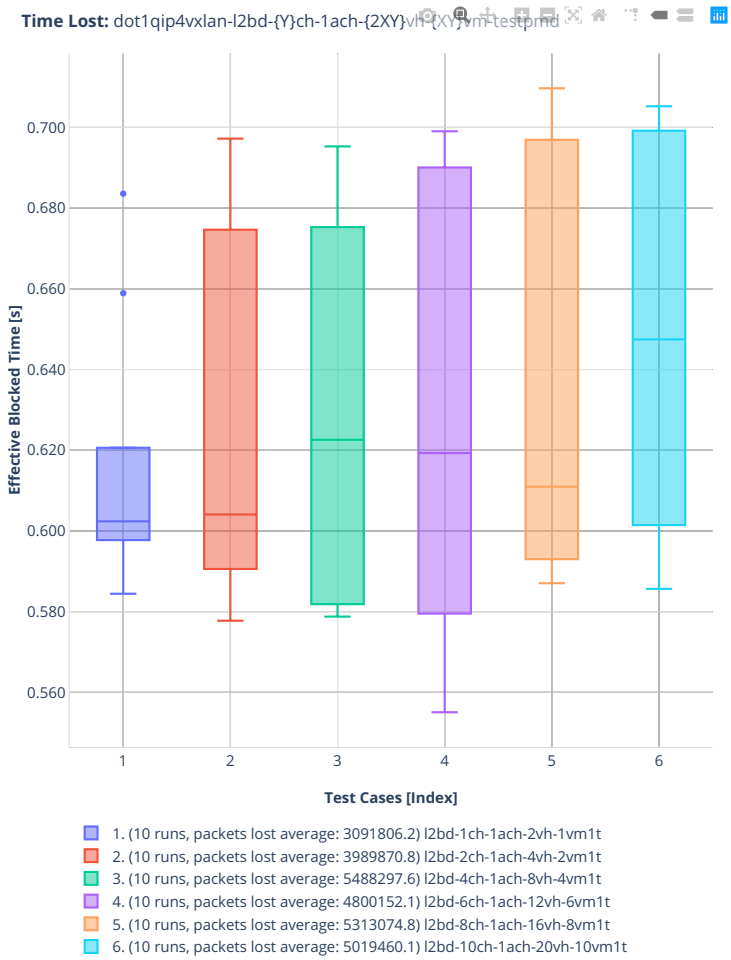

2n-clx-xxv710

imix-2t1c-dot1qip4vxlan-l2bd

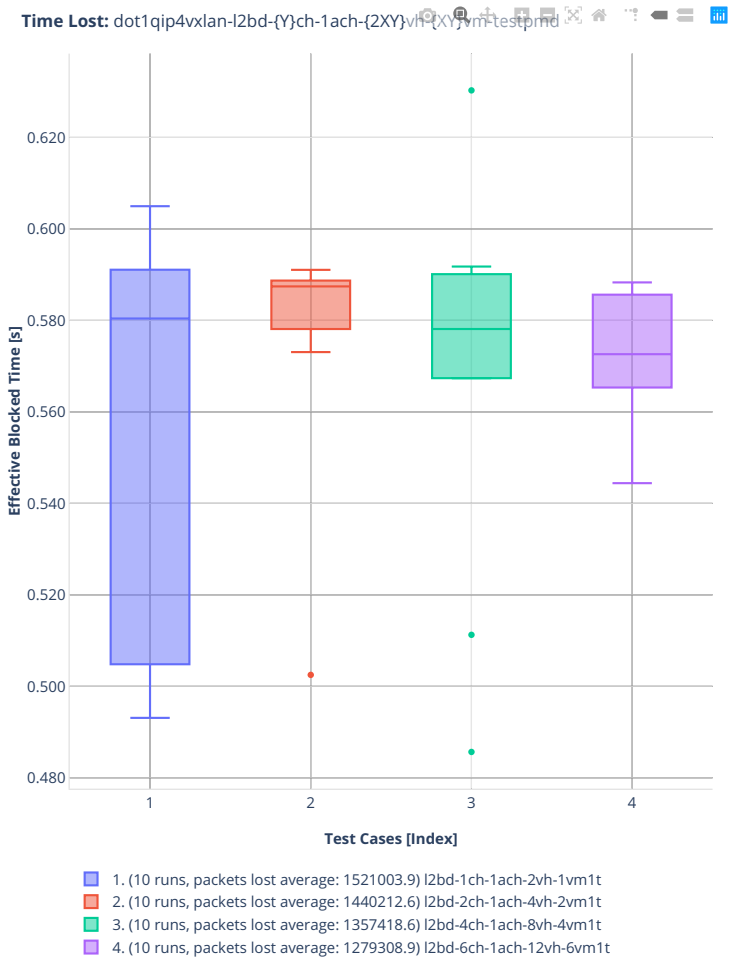

imix-4t2c-dot1qip4vxlan-l2bd

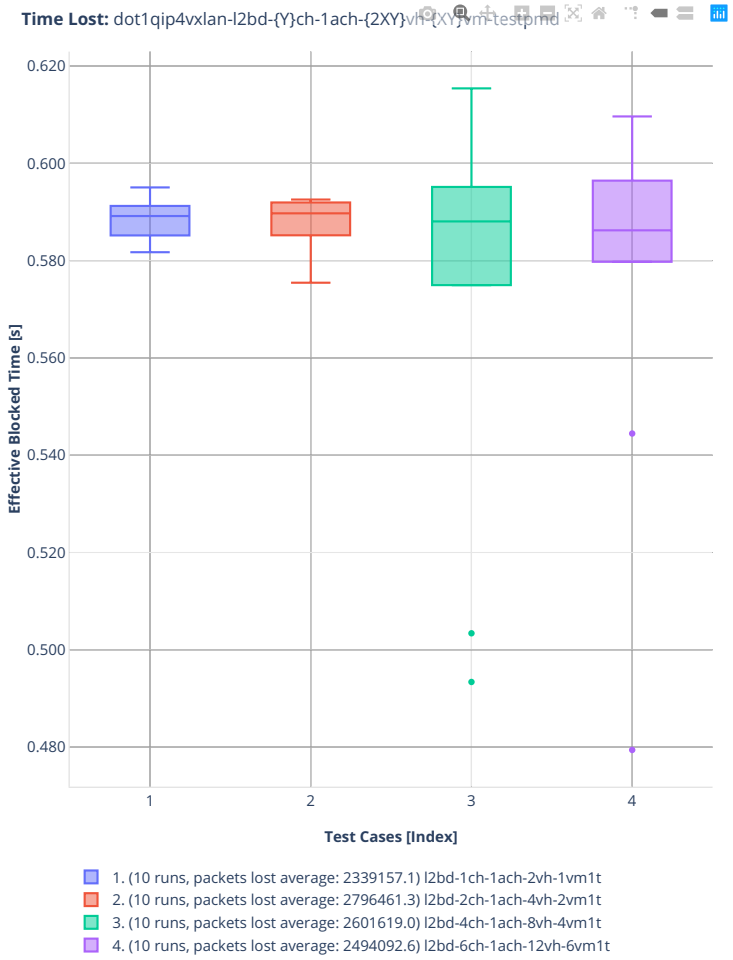

imix-8t4c-dot1qip4vxlan-l2bd

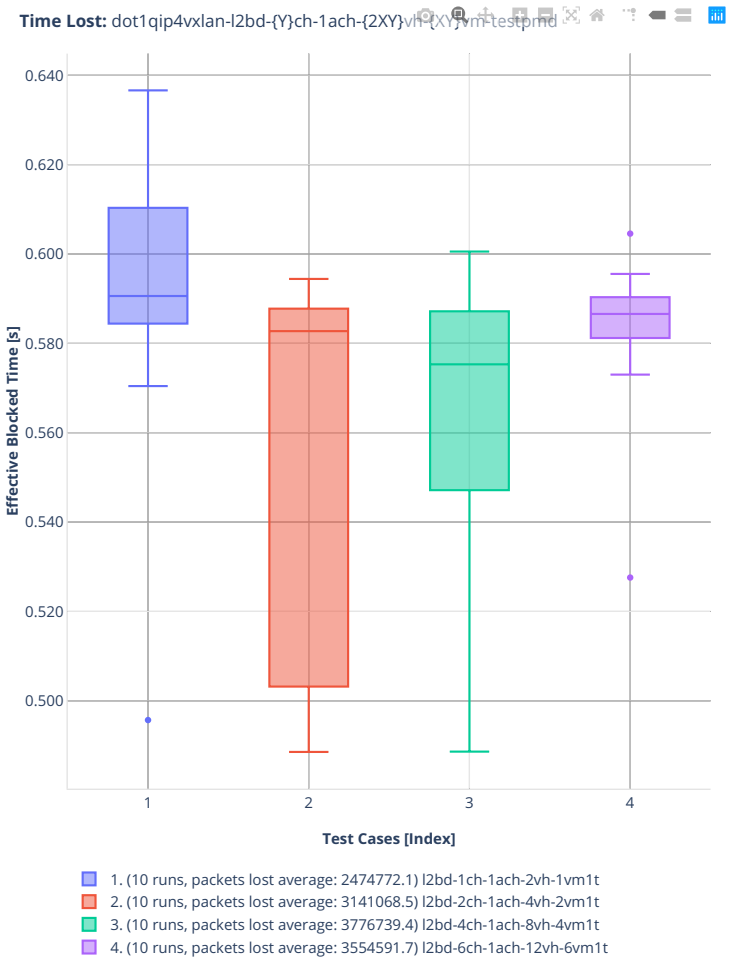

2.8 NFV Service Density

NFV Service Density is benchmarked in three distinct NF service configurations:

- VNF Service Chains Routing
- CNF Service Chains Routing
- CNF Service Pipelines Routing
- VNF Service Chains Tunnels

Each configuration is tested in a number of service density combinations [Number of Service Instances] x [Number of NFs per Service Instance]. The actual tested range is based on available CPU physical core resources.

2.8.1 VNF Service Chains Routing

Throughput graphs for VNF service chains are generated by multiple executions of tests covering a range of VNF service densities defined as [Number of Service Chains] x [Number of VNFs per Service Chain]. The results are presented in the service density graph. Each graph includes the results of both configurations: one NF per physical core and two NFs per physical core and their relative difference.

Additional information about graph data:

1. **Graph Title**: describes tested packet path including VNF workload running in each VM.
2. **X-axis Labels**: VNFs per service chain.
3. **Y-axis Labels**: number of service chains.
4. **Z-axis Color Scale**: lists 64B/IMIX Packet Throughput (mean MRR/NDR/PDR value) in Mpps or the Relative Difference.
5. **Hover Information**: specific test substring listing vhost-chain-vm combinations, number of runs executed, mean MRR/NDR/PDR throughput in Mpps, standard deviation for both configurations and their relative difference.

Note: Test results are stored in build logs from FD.io vpp performance job 2n-icx[164], build logs from FD.io vpp performance job 2n-skx[165] and build logs from FD.io vpp performance job 2n-clx[166] with RF result files csit-vpp-perf-2206-*.zip archived here.

---

[164] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-icx
[165] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx
[166] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-clx

2n-icx-xxv710-mrr

imix-2t1c-eth-l2bd

imix-4t2c-eth-l2bd

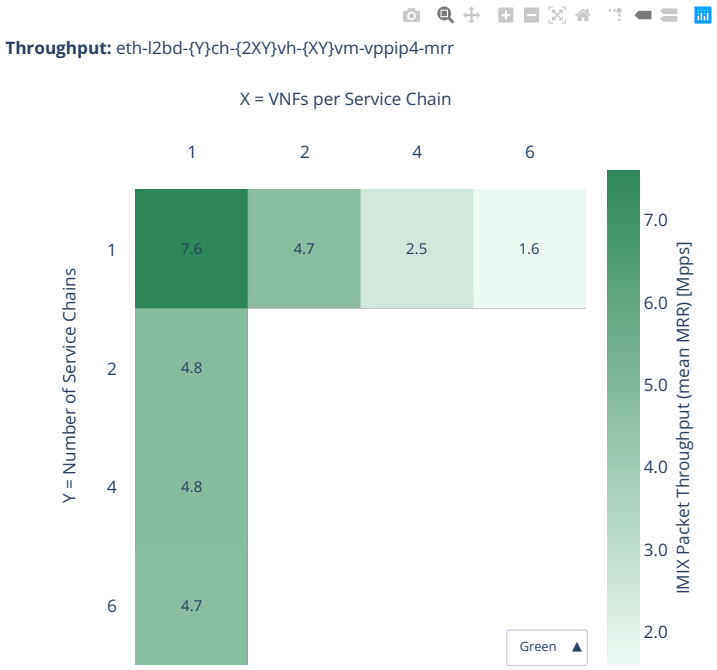

imix-8t4c-eth-l2bd

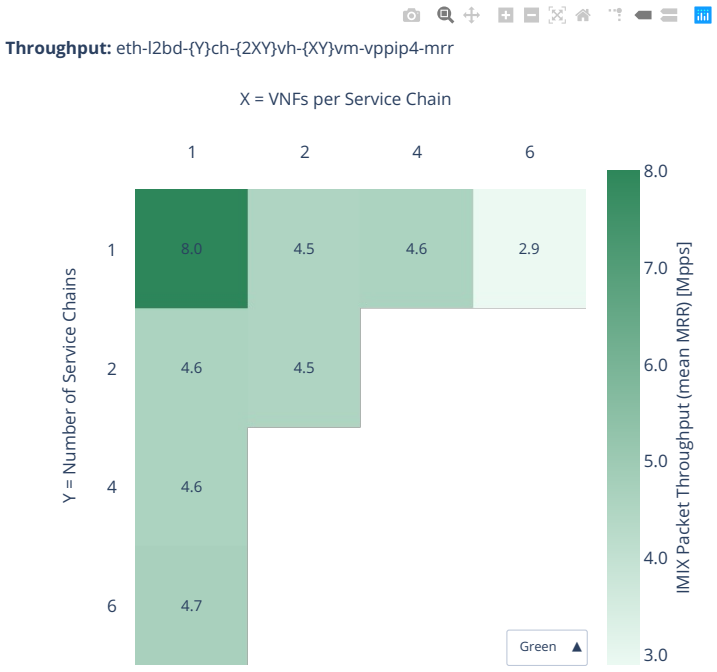

2n-icx-xxv710-ndr

imix-2t1c-eth-l2bd

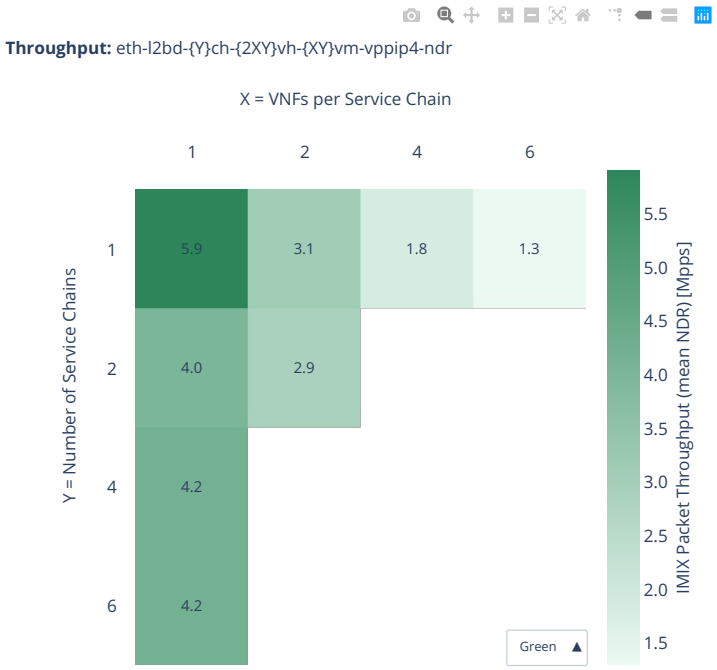

imix-4t2c-eth-l2bd

imix-8t4c-eth-l2bd

2n-icx-xxv710-pdr

imix-2t1c-eth-l2bd

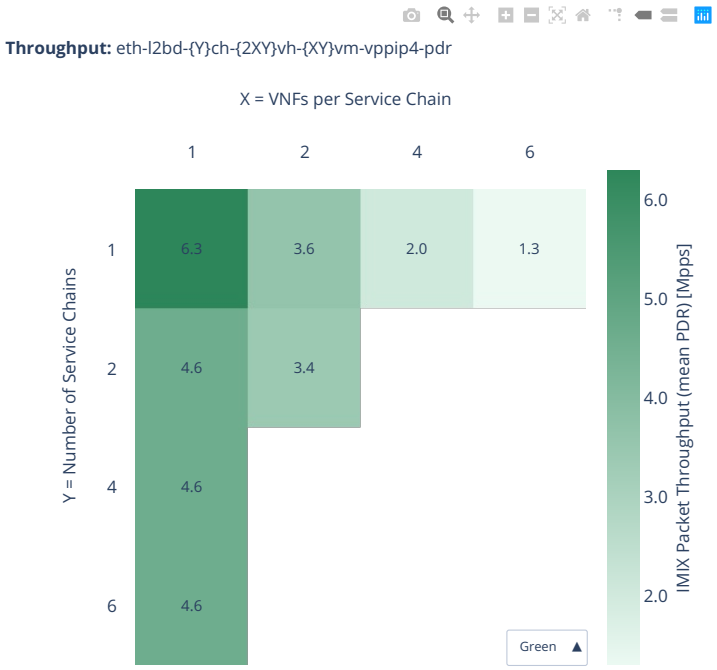

imix-4t2c-eth-l2bd

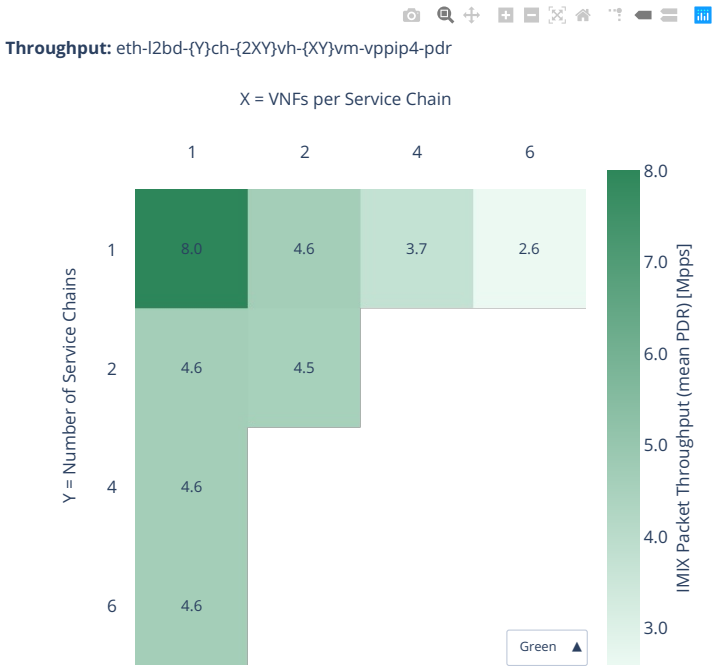

imix-8t4c-eth-l2bd

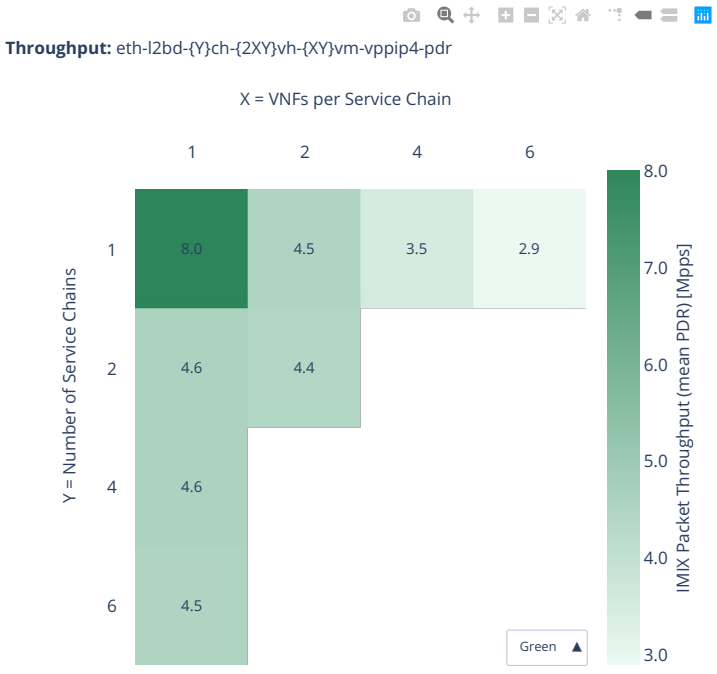

2n-skx-xxv710-mrr

imix-2t1c-eth-l2bd

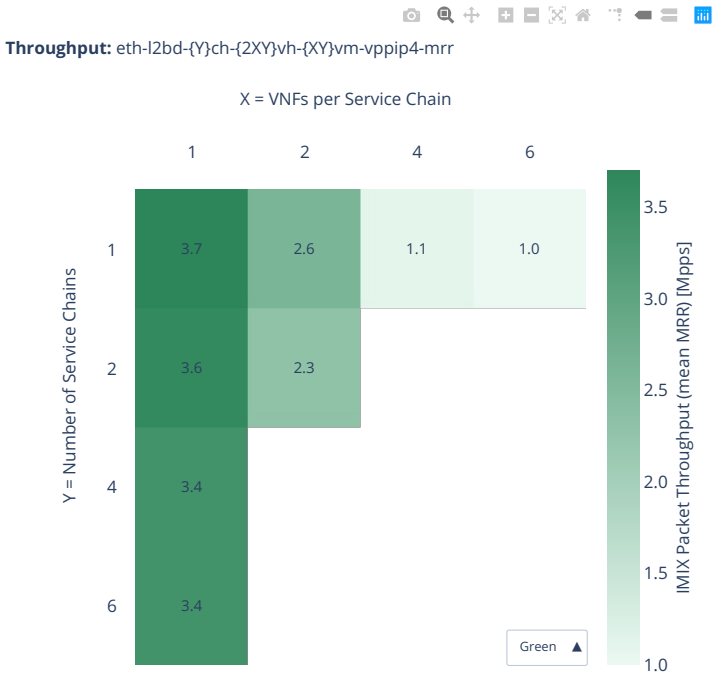

imix-4t2c-eth-l2bd

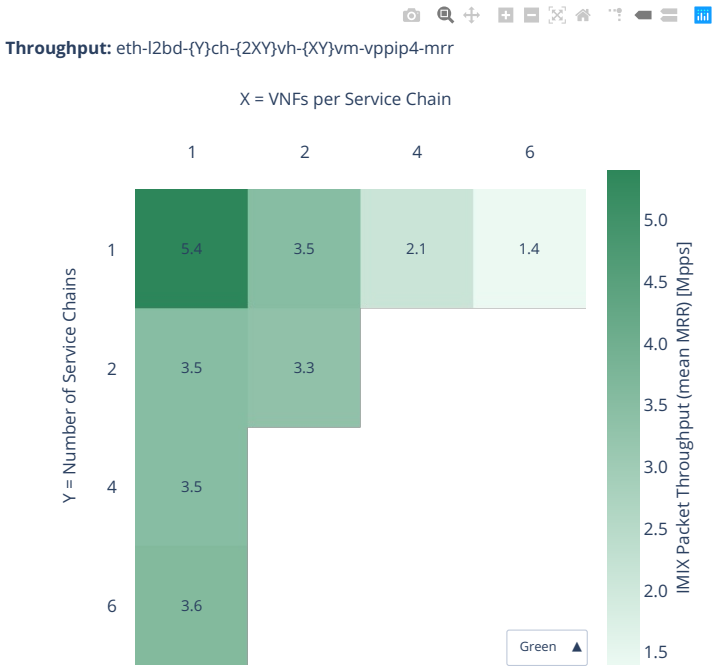

imix-8t4c-eth-l2bd

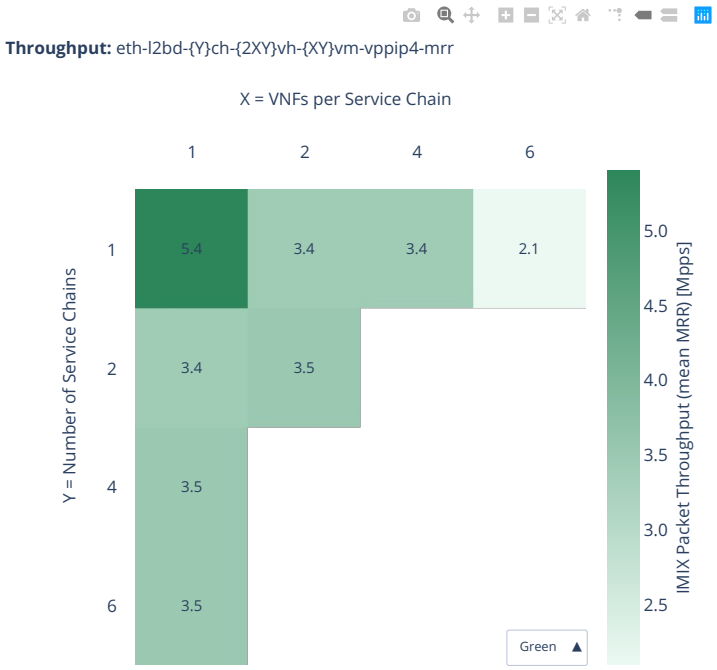

2n-skx-xxv710-ndr

imix-2t1c-eth-l2bd

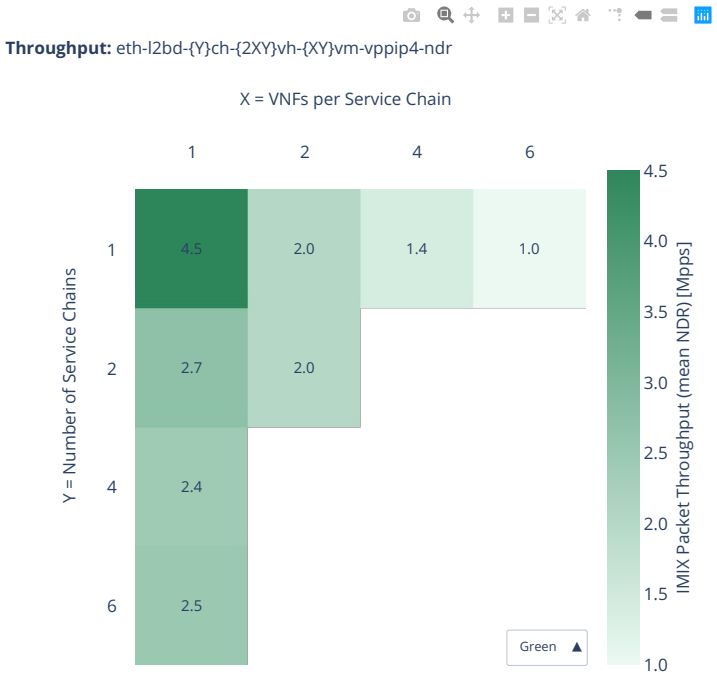

imix-4t2c-eth-l2bd

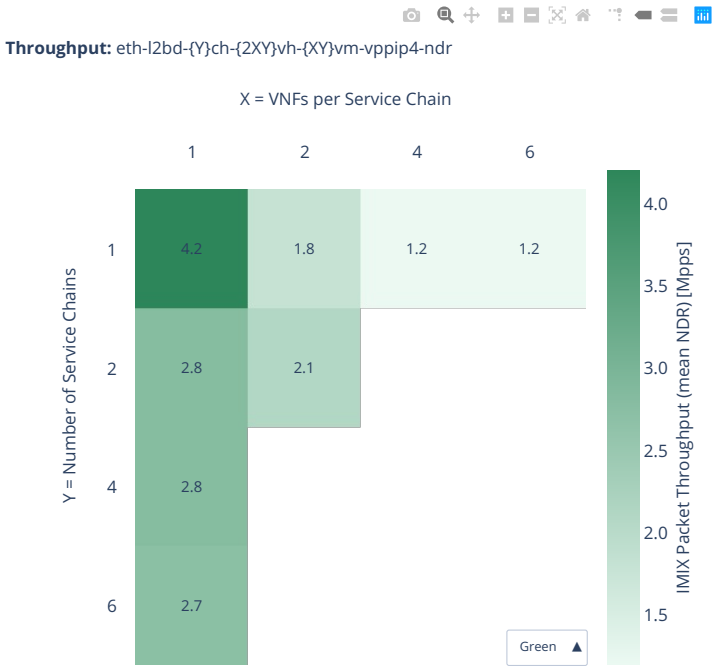

imix-8t4c-eth-l2bd

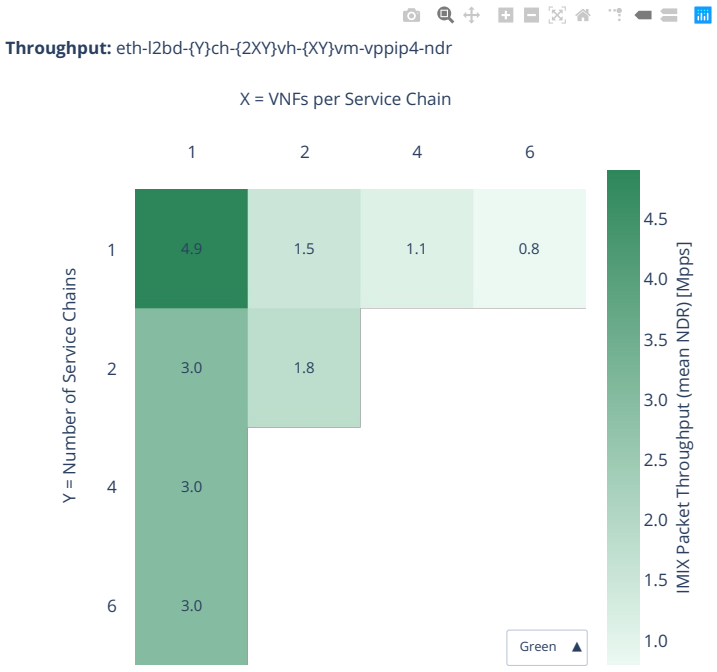

2n-skx-xxv710-pdr

imix-2t1c-eth-l2bd

imix-4t2c-eth-l2bd

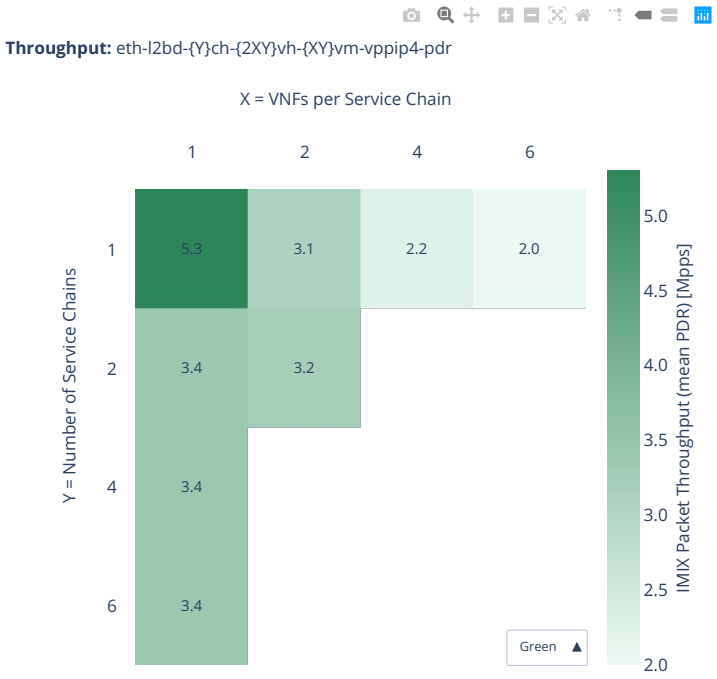

imix-8t4c-eth-l2bd

2n-clx-xxv710-mrr

imix-2t1c-eth-l2bd

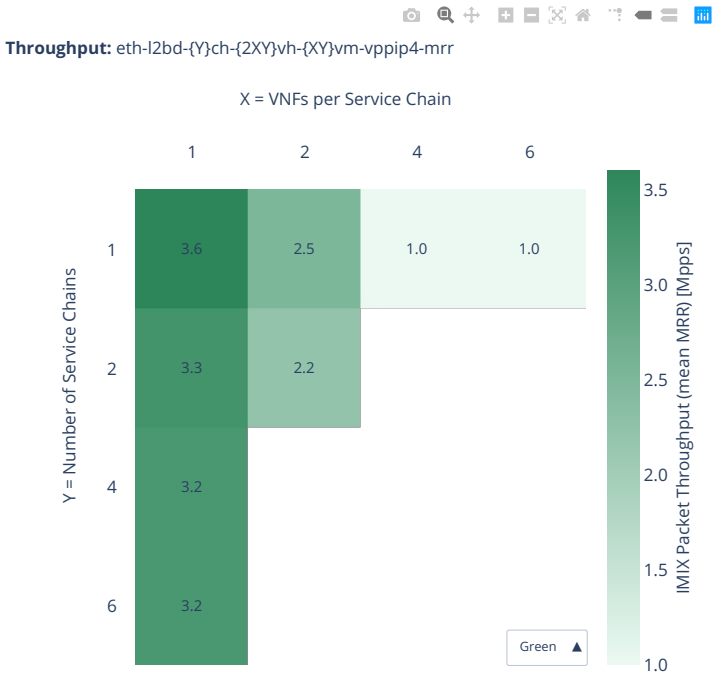

imix-4t2c-eth-l2bd

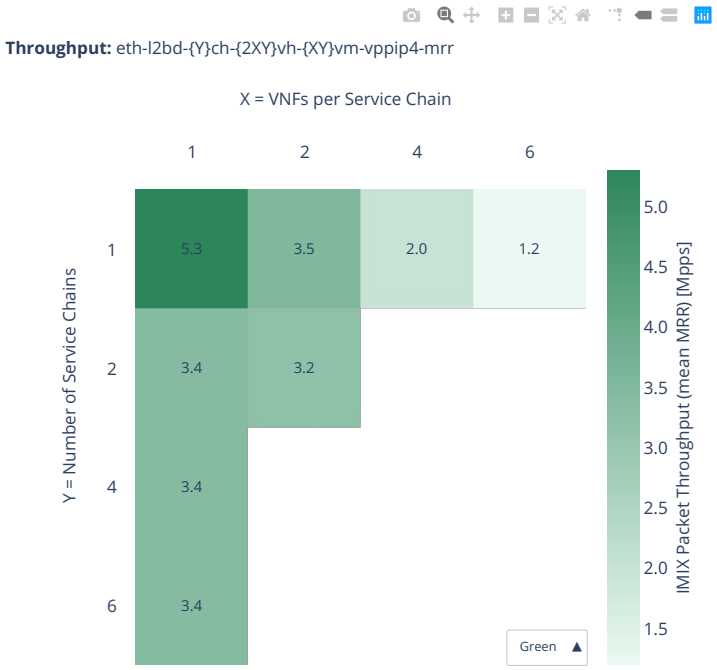

imix-8t4c-eth-l2bd

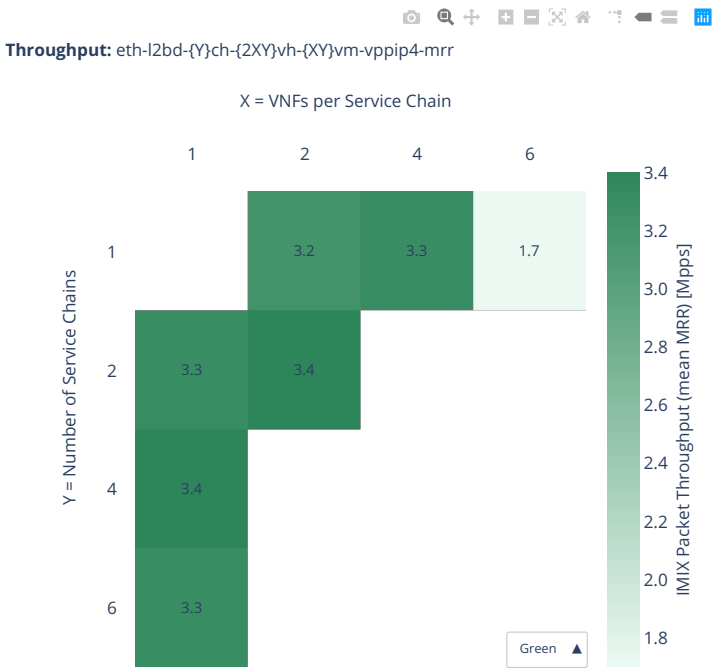

2n-clx-xxv710-ndr

imix-2t1c-eth-l2bd

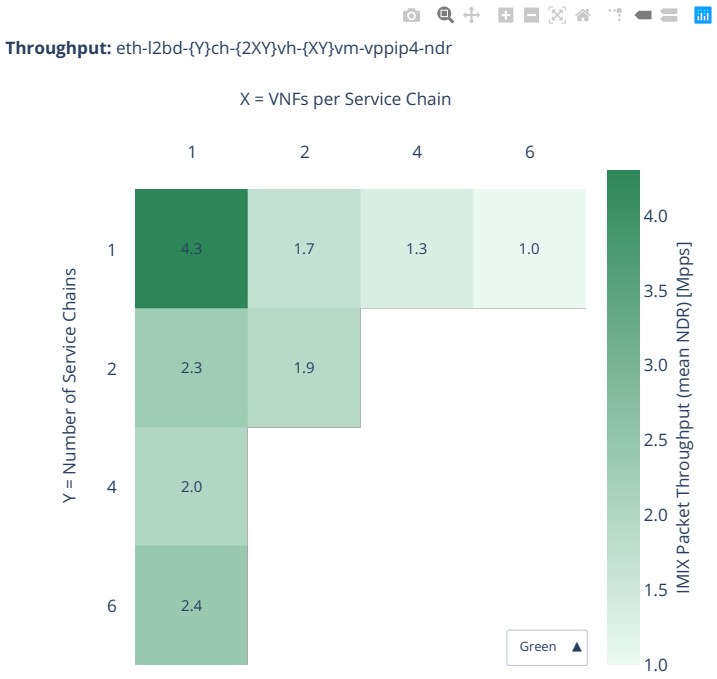

imix-4t2c-eth-l2bd

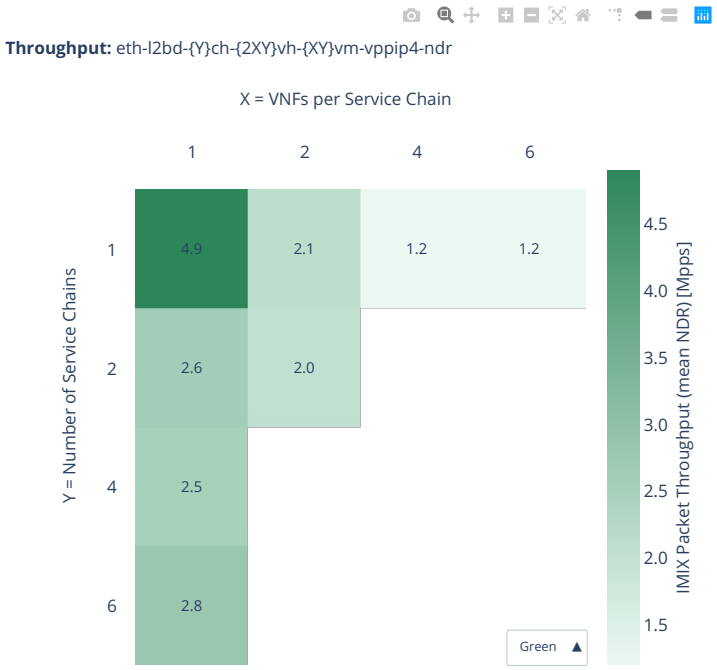

imix-8t4c-eth-l2bd

2n-clx-xxv710-pdr

imix-2t1c-eth-l2bd

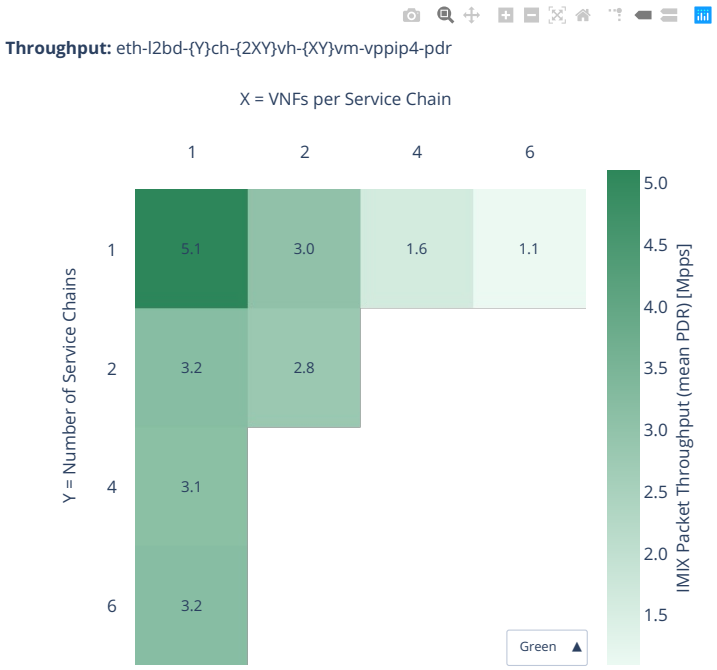

imix-4t2c-eth-l2bd

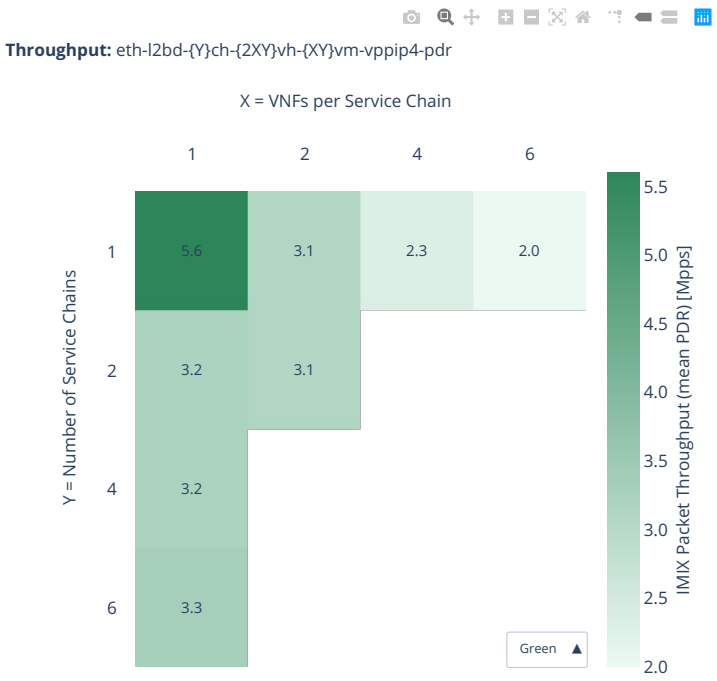

imix-8t4c-eth-l2bd

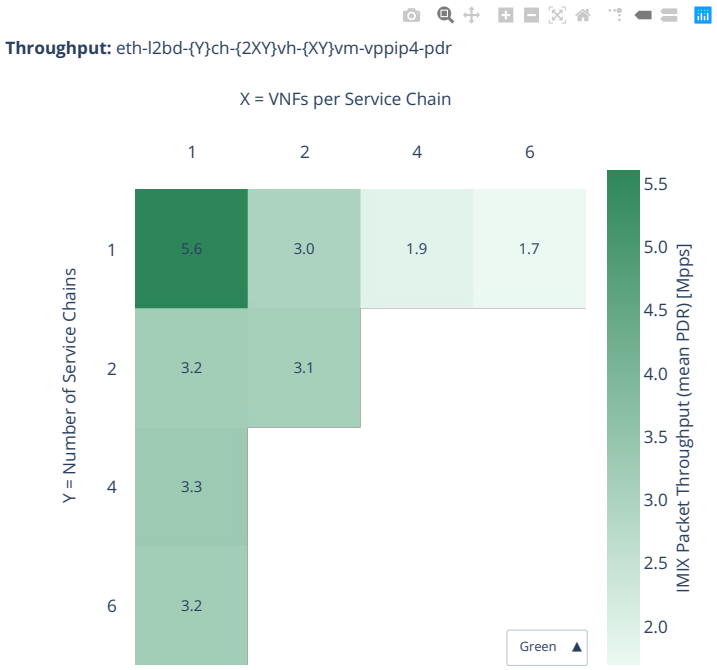

2.8.2 CNF Service Chains Routing

Throughput graphs for CNF service chains are generated by multiple executions of tests covering a range of CNF service densities defined as [Number of Service Chains] x [Number of CNFs per Service Chain]. The results are presented in the service density graph. Each graph includes the results of both configurations: one NF per physical core and two NFs per physical core and their relative difference.

Additional information about graph data:

1. **Graph Title**: describes tested packet path including CNF workload running in each Docker Container.
2. **X-axis Labels**: CNFs per service chain.
3. **Y-axis Labels**: number of service chains.
4. **Z-axis Color Scale**: lists 64B/IMIX Packet Throughput (mean MRR/NDR/PDR value) in Mpps or the Relative Difference.
5. **Hover Information**: specific test substring listing memif-chain-docker_container combinations, number of runs executed, mean MRR/NDR/PDR throughput in Mpps, standard deviation for both configurations and their relative difference.

Note: Test results are stored in build logs from FD.io vpp performance job 2n-icx[167], build logs from FD.io vpp performance job 2n-skx[168] and build logs from FD.io vpp performance job 2n-clx[169] with RF result files csit-vpp-perf-2206-*.zip archived here.

---

[167] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-icx
[168] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx
[169] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-clx

2n-icx-xxv710-mrr

imix-2t1c-eth-l2bd

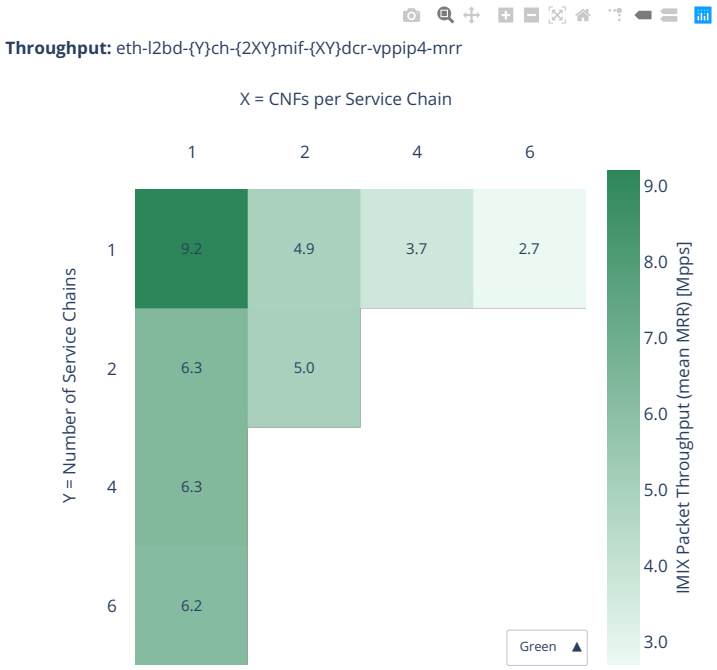

imix-4t2c-eth-l2bd

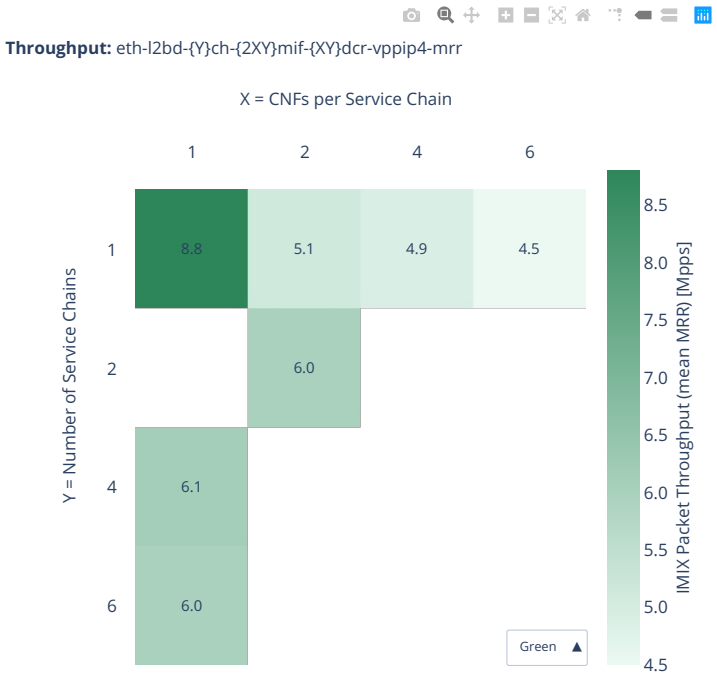

imix-8t4c-eth-l2bd

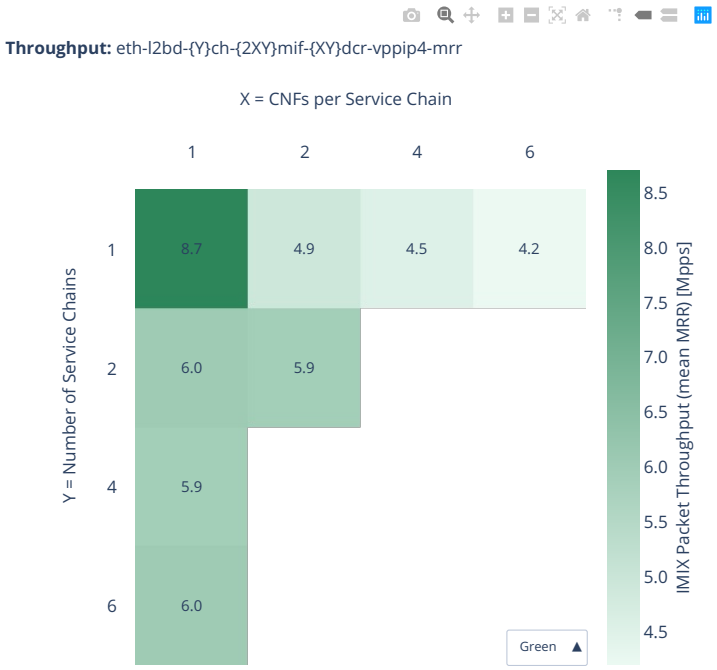

2n-icx-xxv710-ndr

imix-2t1c-eth-l2bd

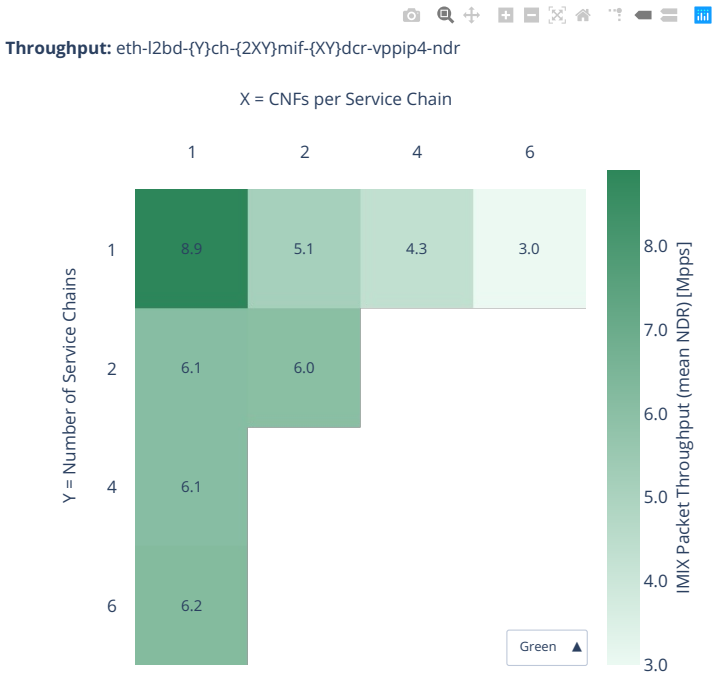

imix-4t2c-eth-l2bd

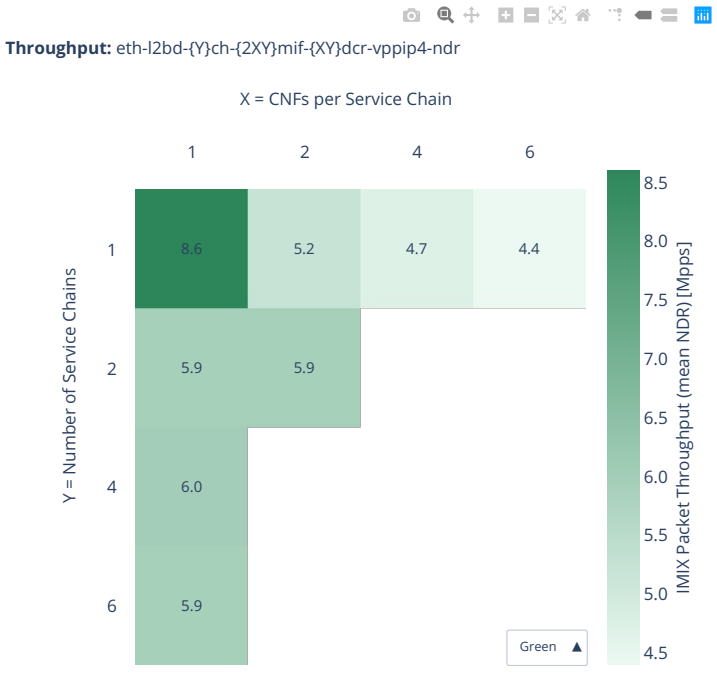

imix-8t4c-eth-l2bd

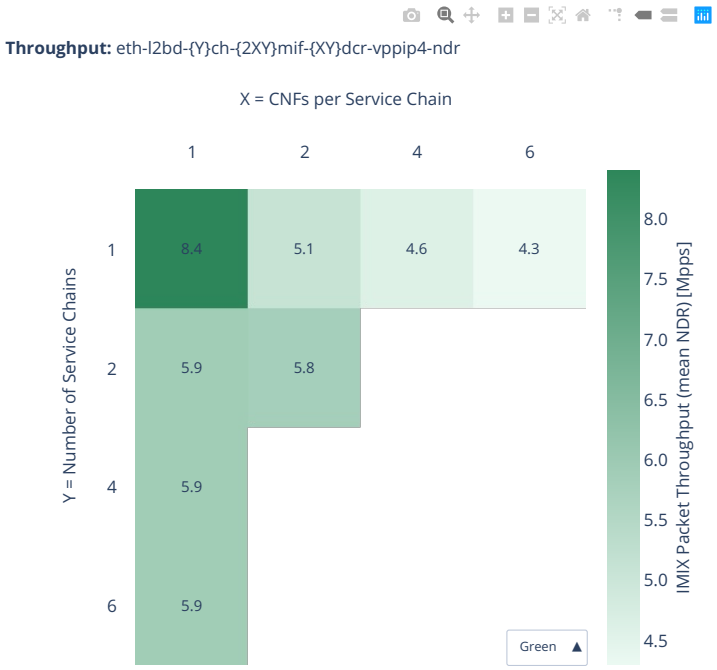

2n-icx-xxv710-pdr

imix-2t1c-eth-l2bd

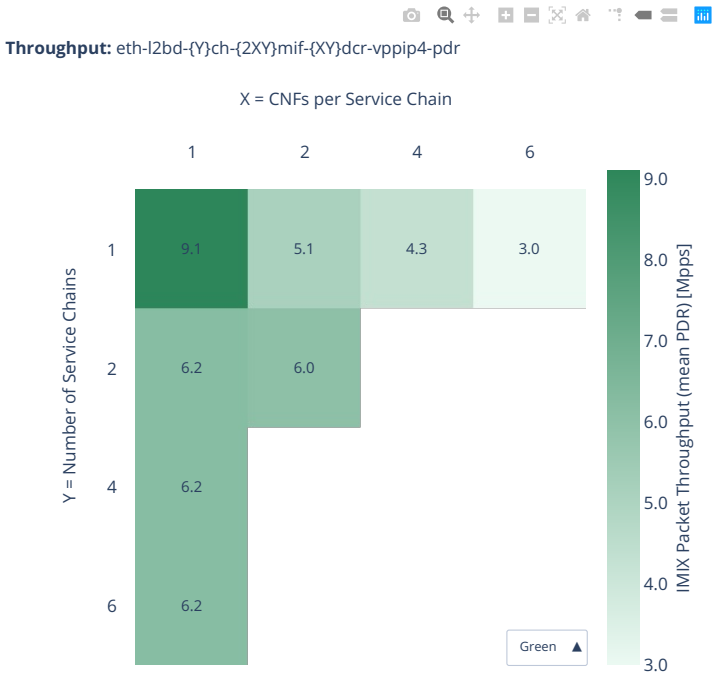

imix-4t2c-eth-l2bd

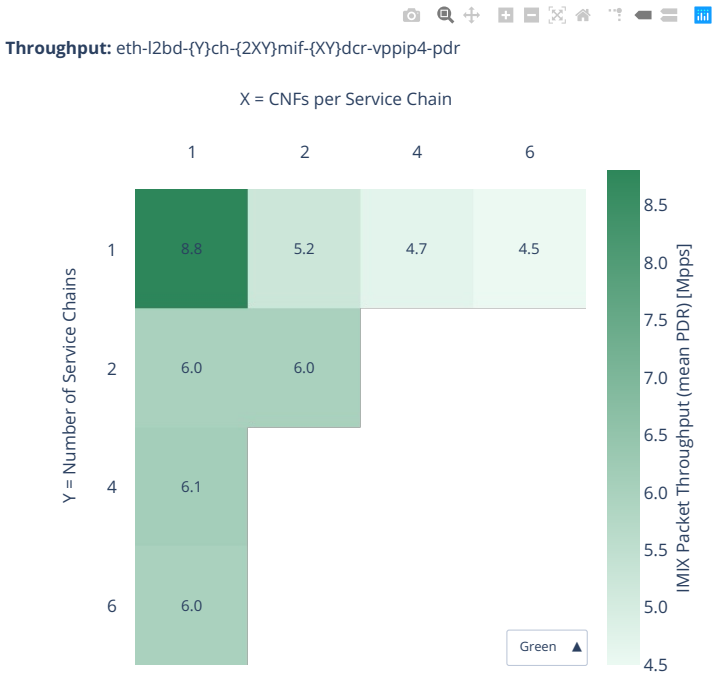

imix-8t4c-eth-l2bd

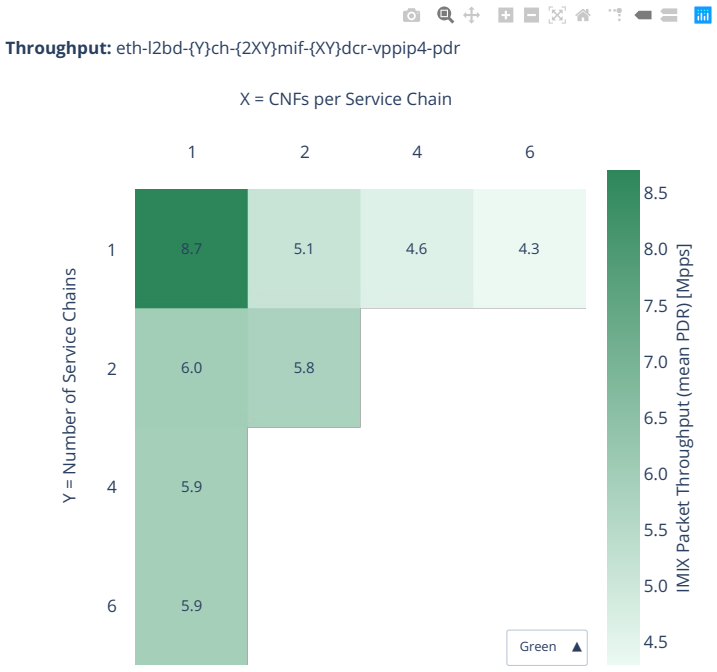

2n-skx-xxv710-mrr

imix-2t1c-eth-l2bd

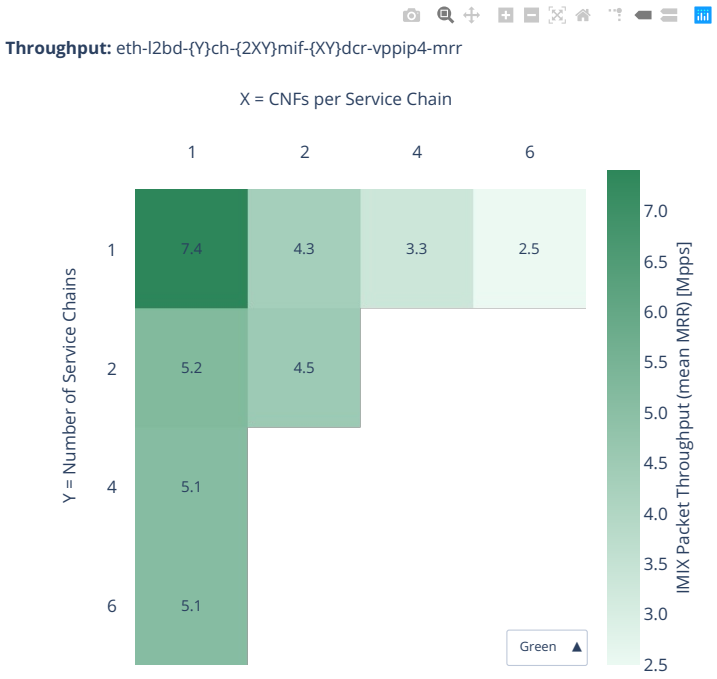

imix-4t2c-eth-l2bd

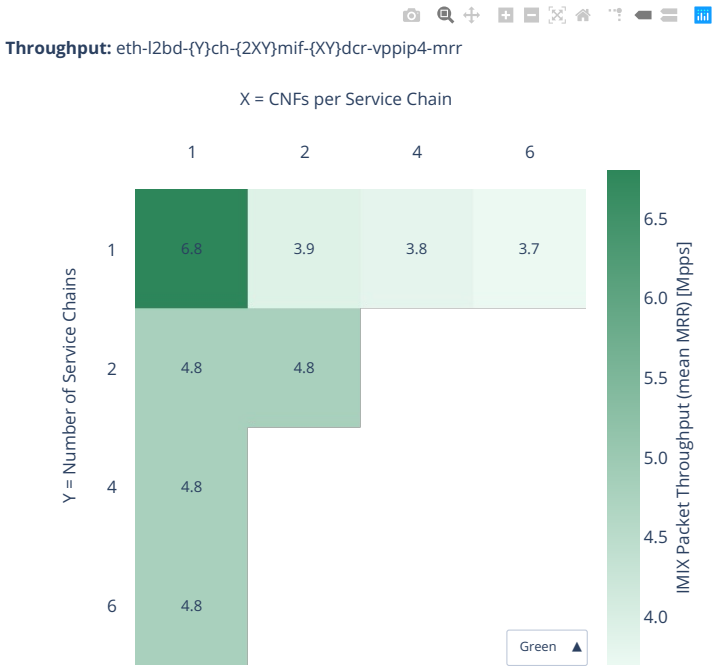

imix-8t4c-eth-l2bd

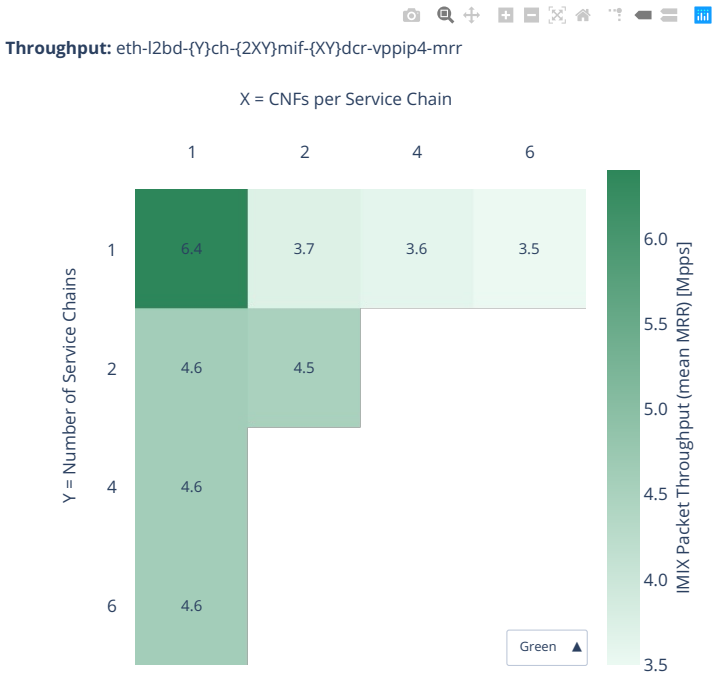

2n-skx-xxv710-ndr

imix-2t1c-eth-l2bd

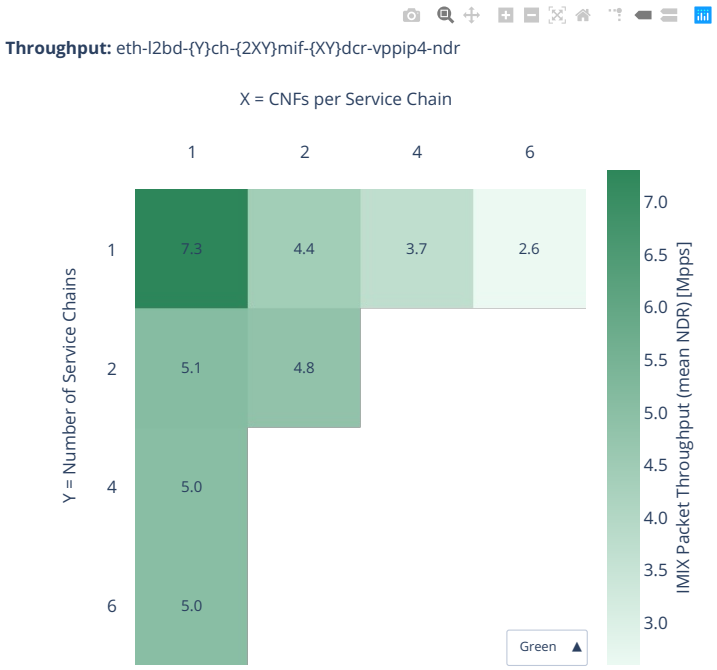

imix-4t2c-eth-l2bd

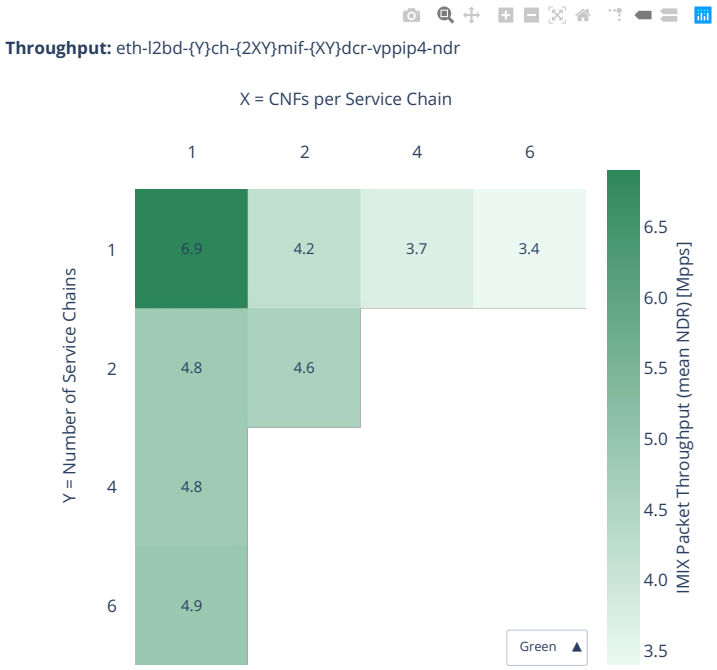

imix-8t4c-eth-l2bd

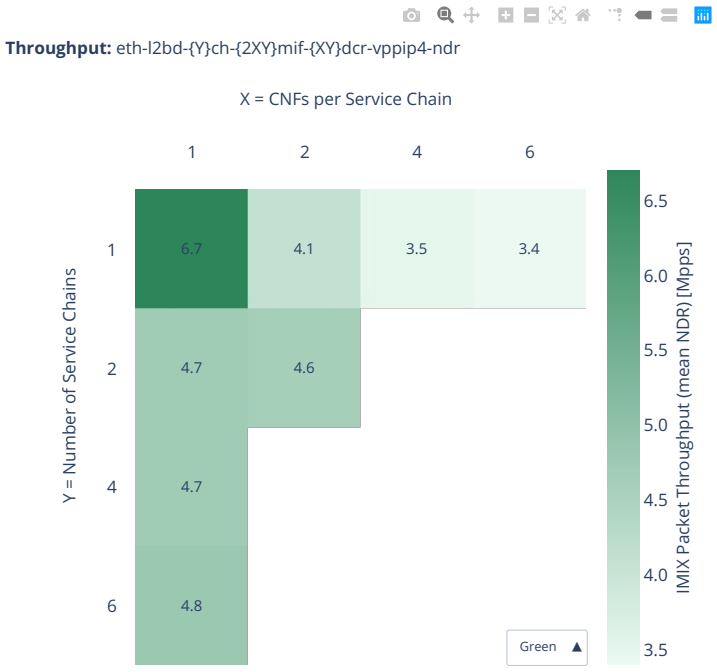

2n-skx-xxv710-pdr

imix-2t1c-eth-l2bd

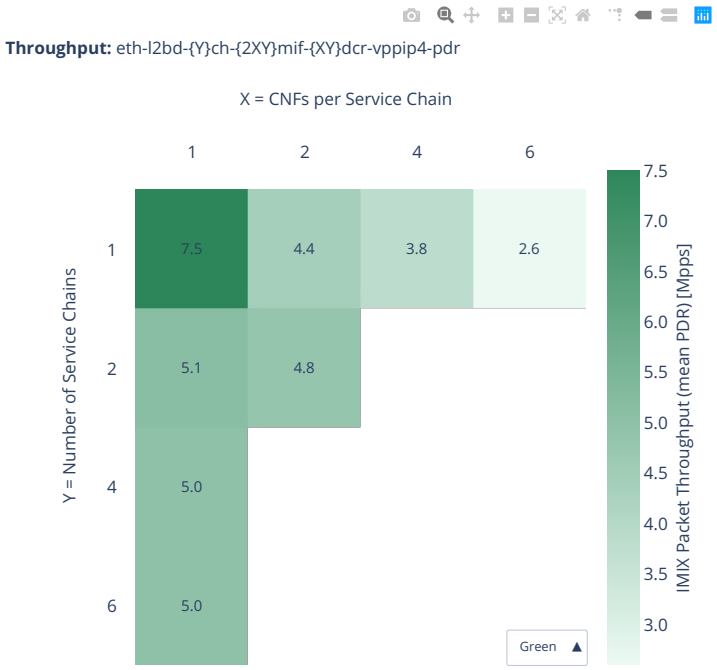

imix-4t2c-eth-l2bd

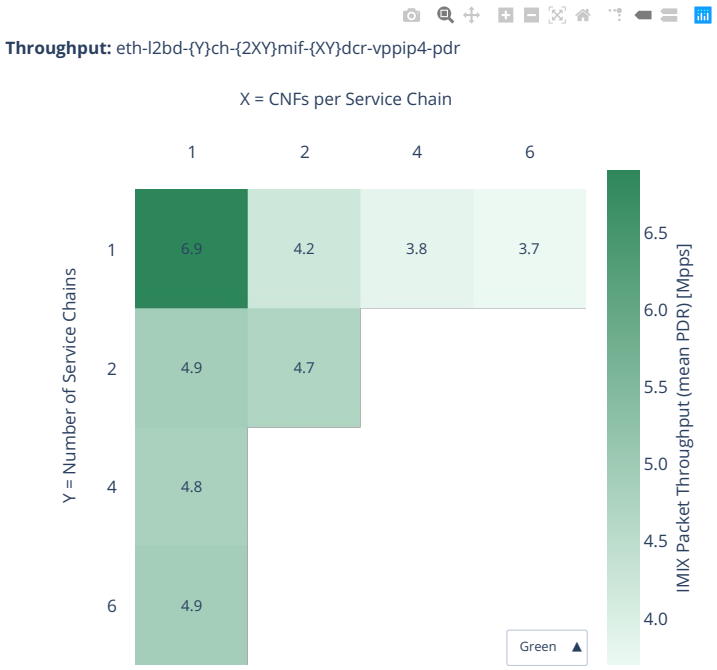

imix-8t4c-eth-l2bd

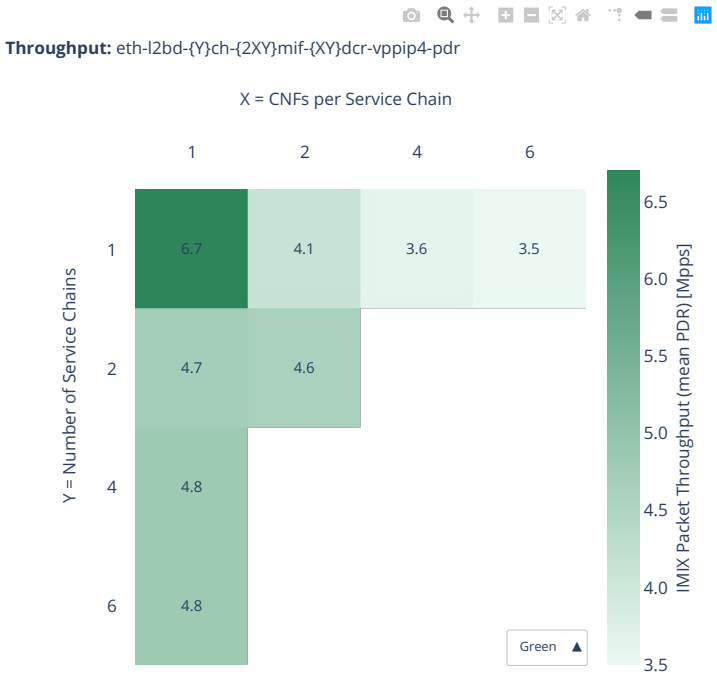

2n-clx-xxv710-mrr

imix-2t1c-eth-l2bd

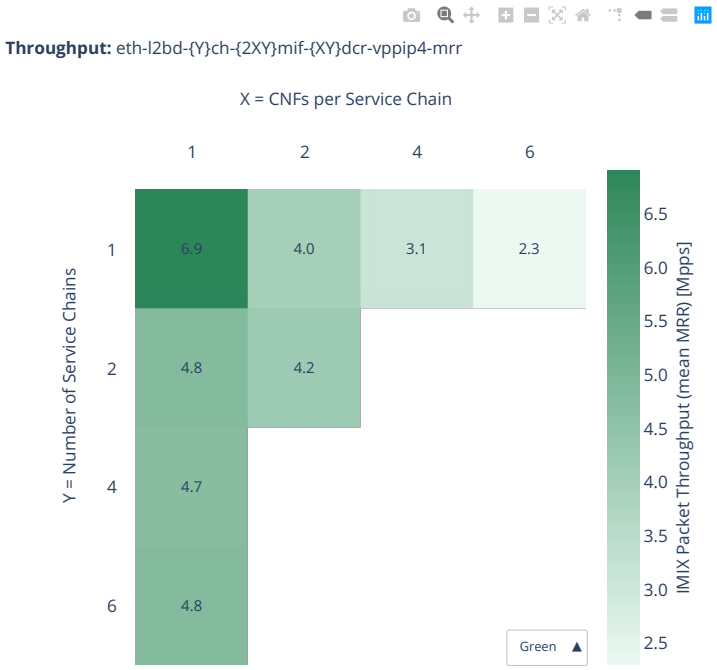

imix-4t2c-eth-l2bd

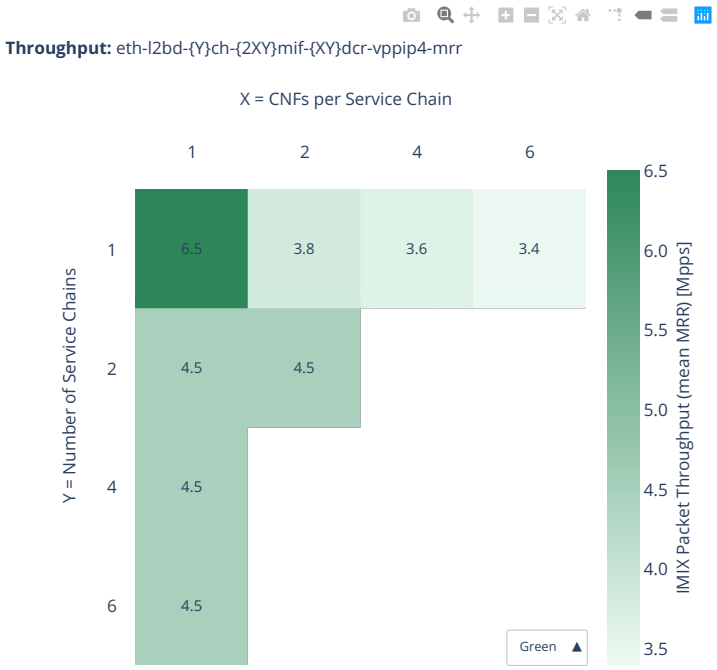

imix-8t4c-eth-l2bd

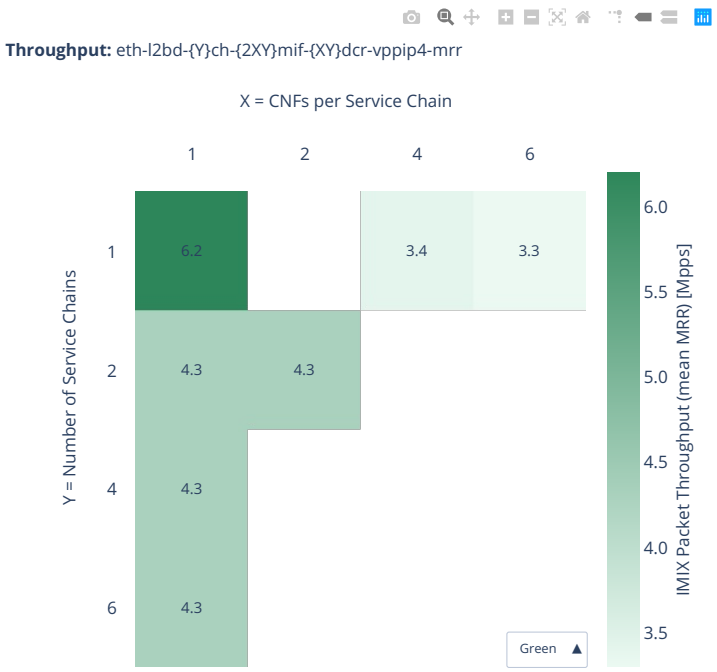

2n-clx-xxv710-ndr

imix-2t1c-eth-l2bd

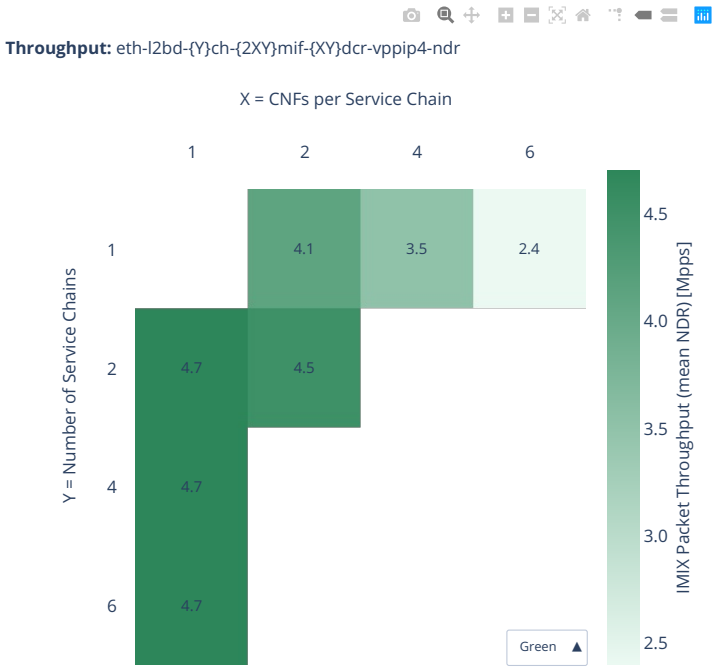

imix-4t2c-eth-l2bd

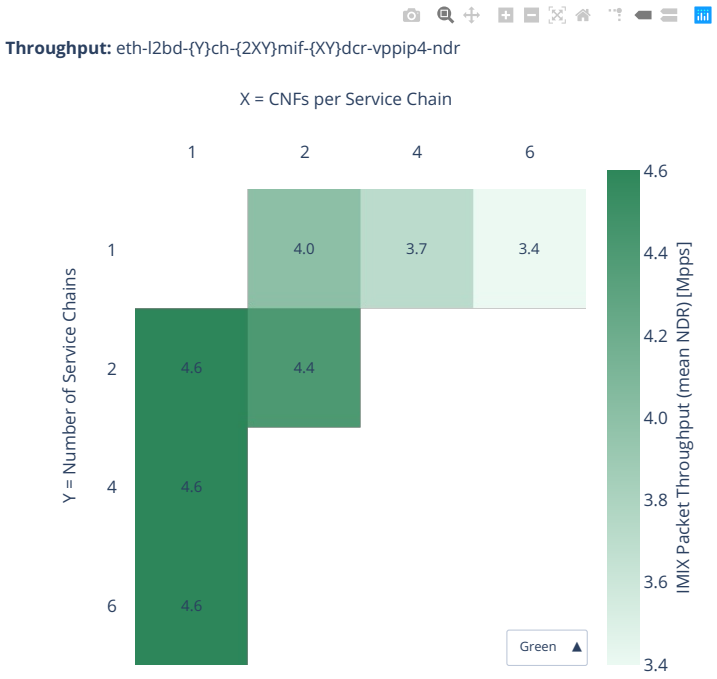

imix-8t4c-eth-l2bd

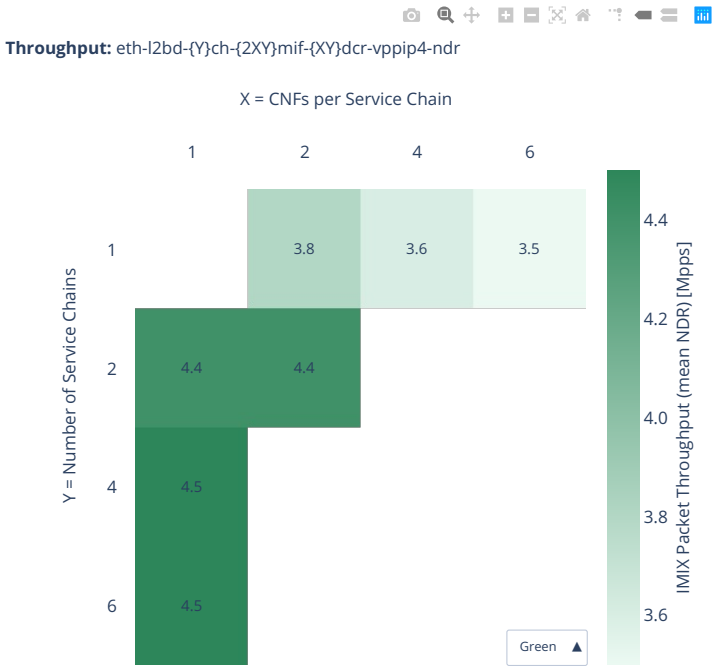

2n-clx-xxv710-pdr

imix-2t1c-eth-l2bd

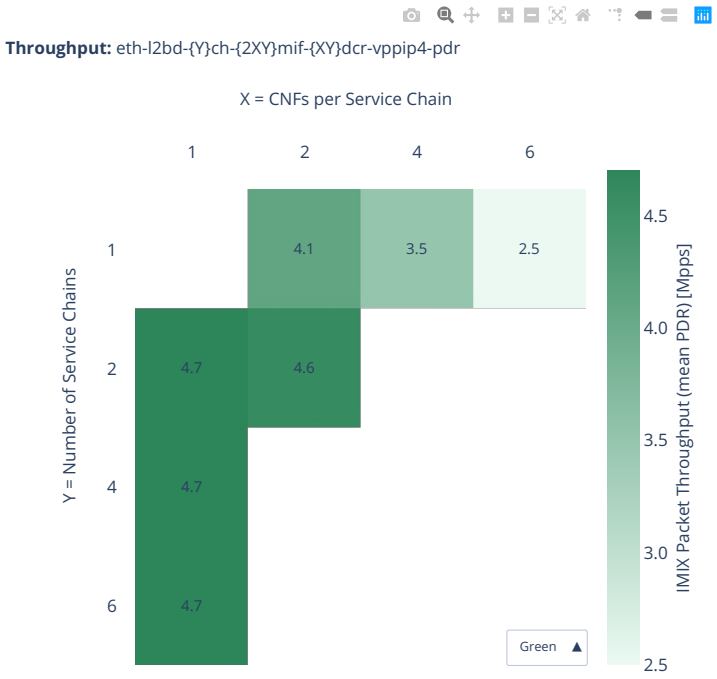

imix-4t2c-eth-l2bd

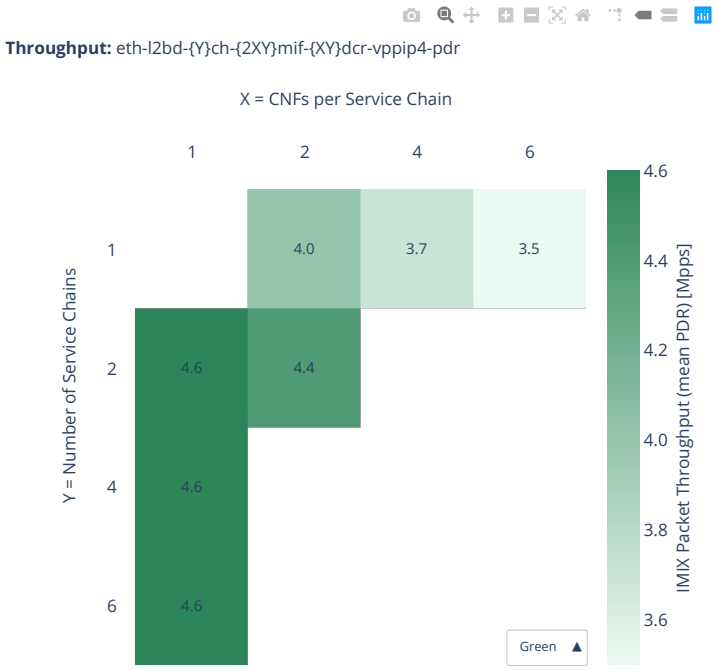

imix-8t4c-eth-l2bd

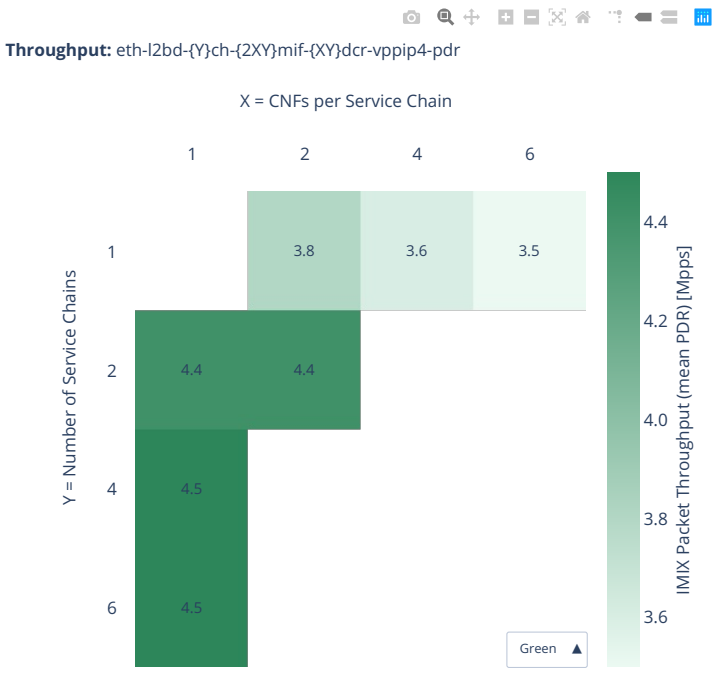

2.8.3 CNF Service Pipelines Routing

Throughput graphs for CNF service pipelines are generated by multiple executions of tests covering a range of CNF service densities defined as [Number of Service Pipelines] x [Number of CNFs per Service Pipeline]. The results are presented in the service density graph. Each graph includes the results of both configurations: one NF per physical core and two NFs per physical core and their relative difference.

Additional information about graph data:

1. **Graph Title**: describes tested packet path including CNF workload running in each Docker Container.
2. **X-axis Labels**: CNFs per service pipeline.
3. **Y-axis Labels**: number of service pipelines.
4. **Z-axis Color Scale**: lists 64B/IMIX Packet Throughput (mean MRR/NDR/PDR value) in Mpps or the Relative Difference.
5. **Hover Information**: specific test substring listing memif-pipeline-docker_container combinations, number of runs executed, mean MRR/NDR/PDR throughput in Mpps, standard deviation for both configurations and their relative difference.

Note: Test results are stored in build logs from FD.io vpp performance job 2n-icx[170], build logs from FD.io vpp performance job 2n-skx[171] and build logs from FD.io vpp performance job 2n-clx[172] with RF result files csit-vpp-perf-2206-*.zip archived here.

---

[170] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-icx
[171] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx
[172] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-clx

2n-icx-xxv710-mrr

imix-2t1c-eth-l2bd

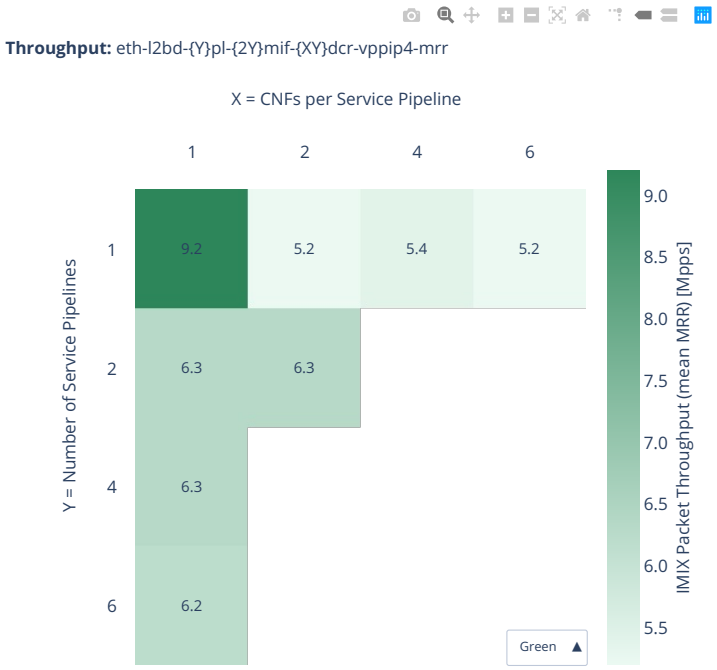

imix-4t2c-eth-l2bd

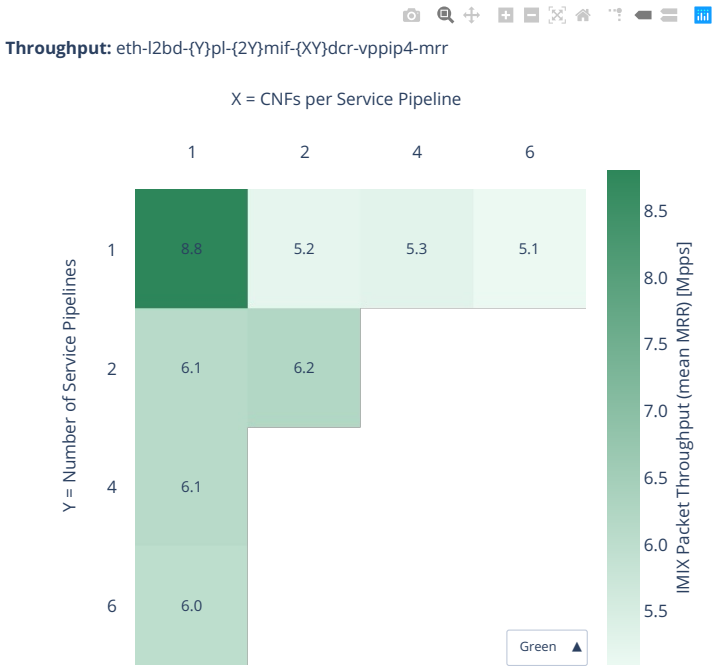

imix-8t4c-eth-l2bd

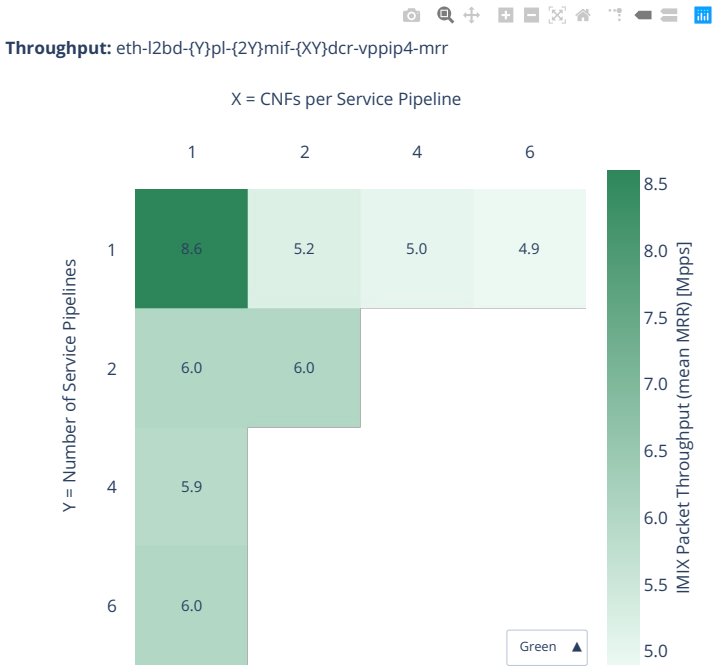

2n-icx-xxv710-ndr

imix-2t1c-eth-l2bd

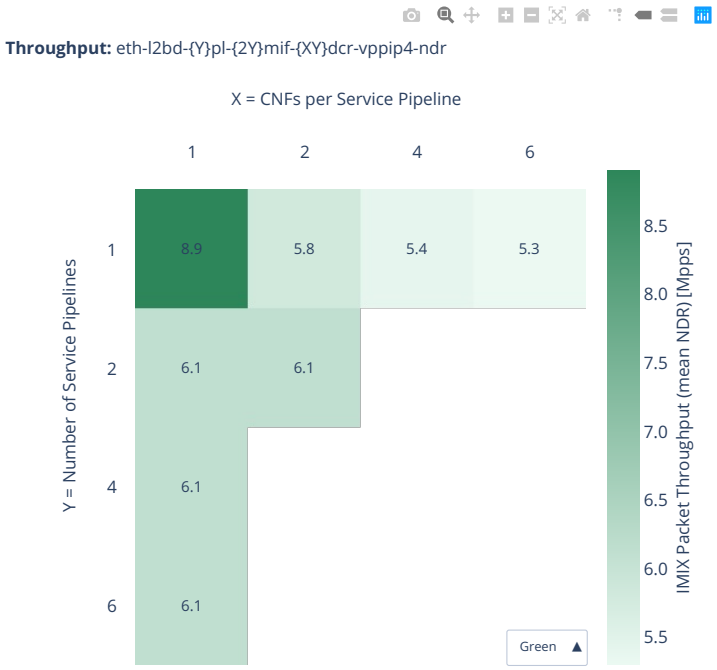

imix-4t2c-eth-l2bd

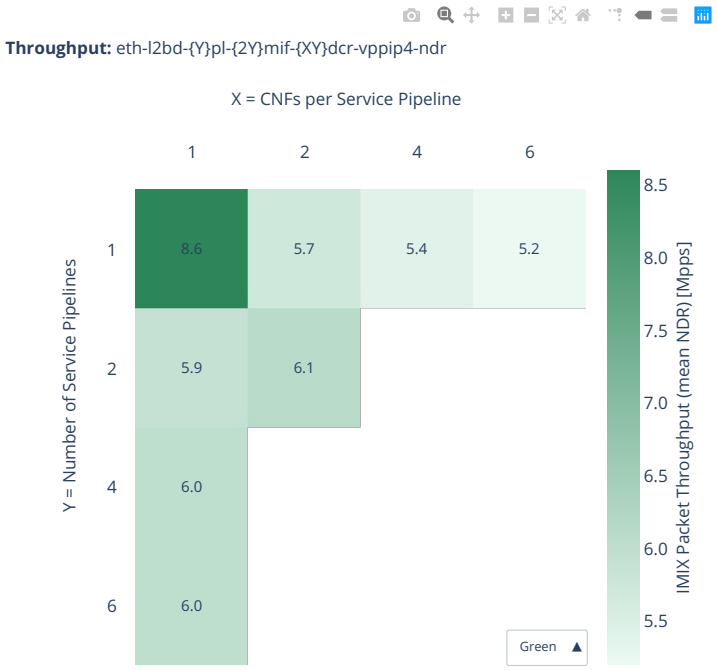

imix-8t4c-eth-l2bd

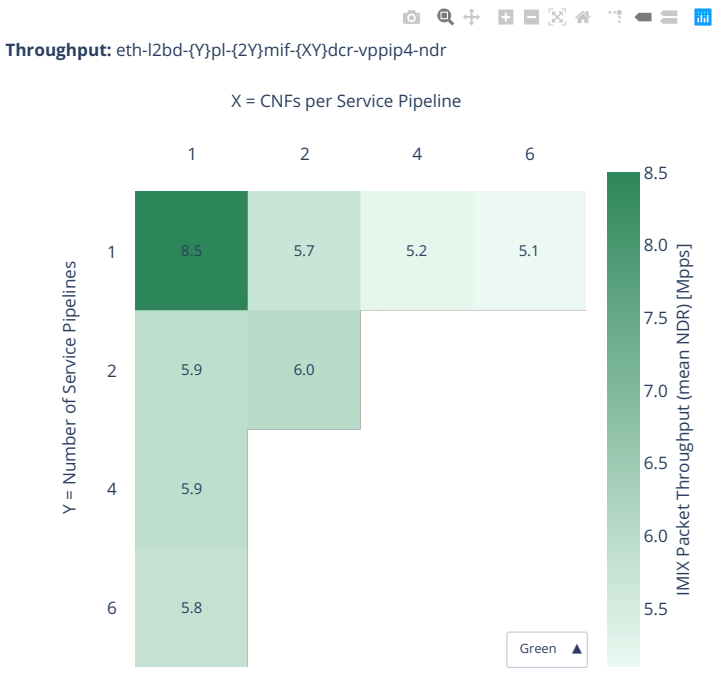

2n-icx-xxv710-pdr

imix-2t1c-eth-l2bd

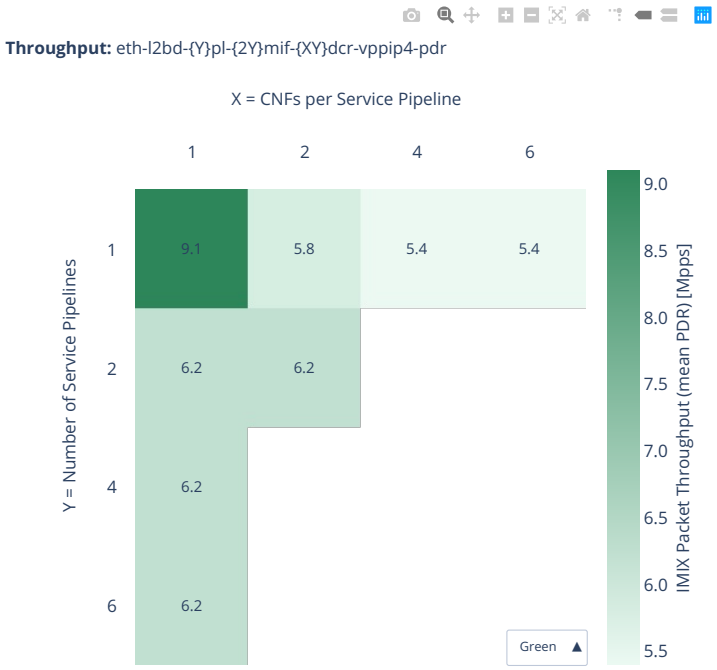

imix-4t2c-eth-l2bd

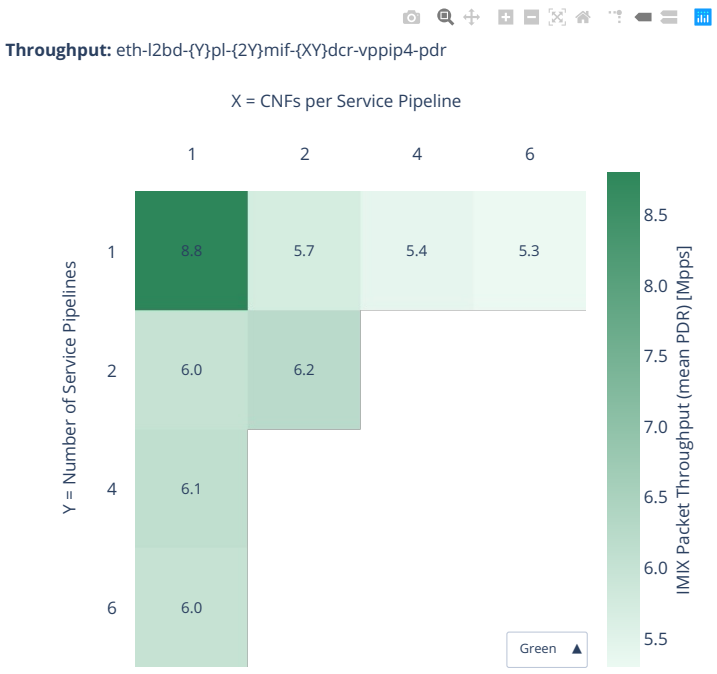

imix-8t4c-eth-l2bd

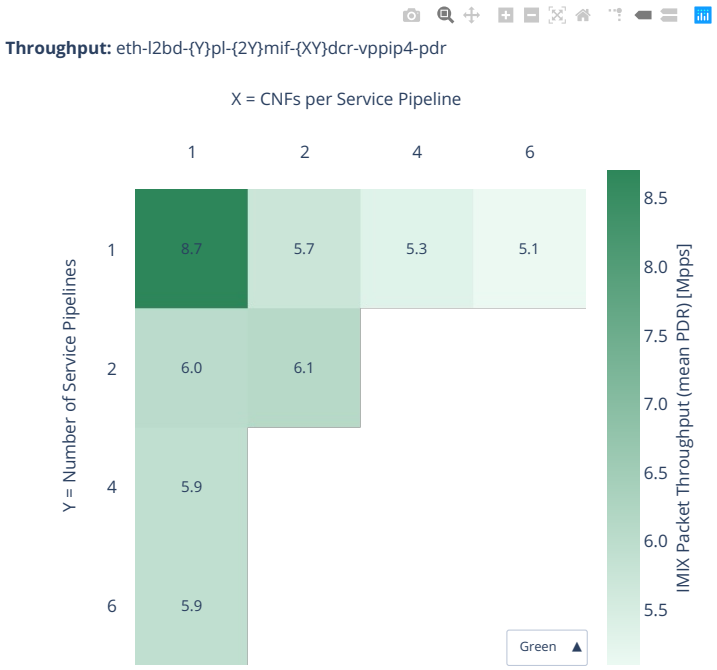

2n-skx-xxv710-mrr

imix-2t1c-eth-l2bd

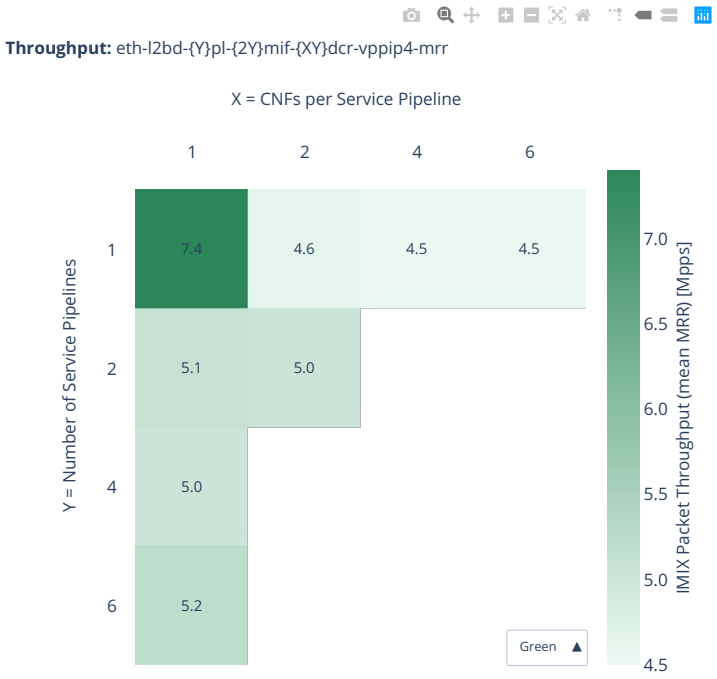

imix-4t2c-eth-l2bd

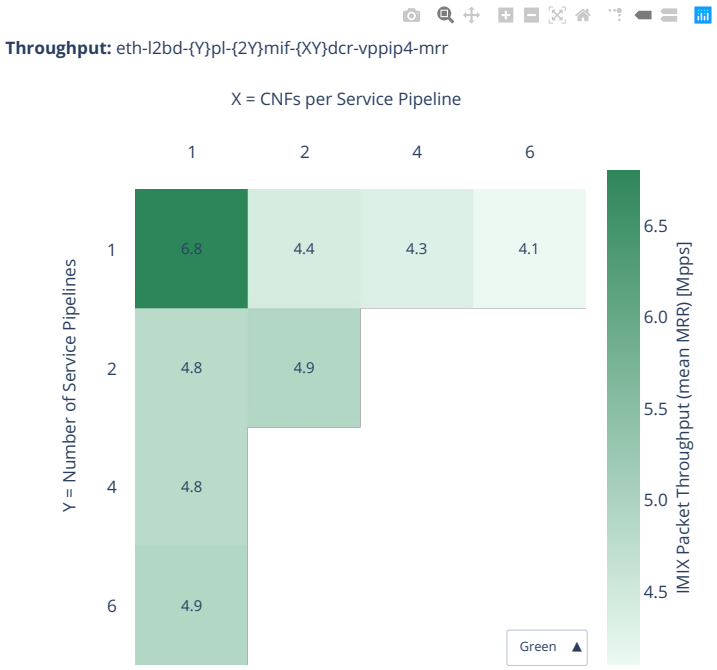

imix-8t4c-eth-l2bd

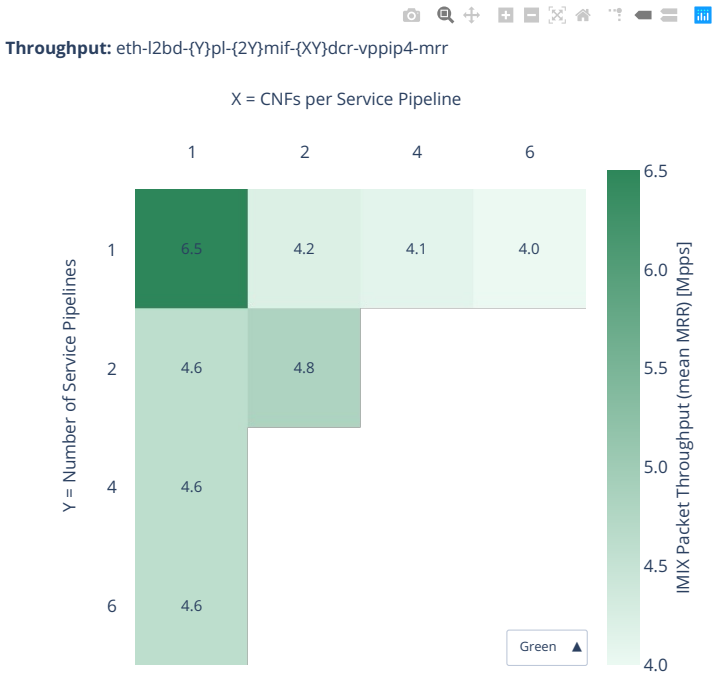

2n-skx-xxv710-ndr

imix-2t1c-eth-l2bd

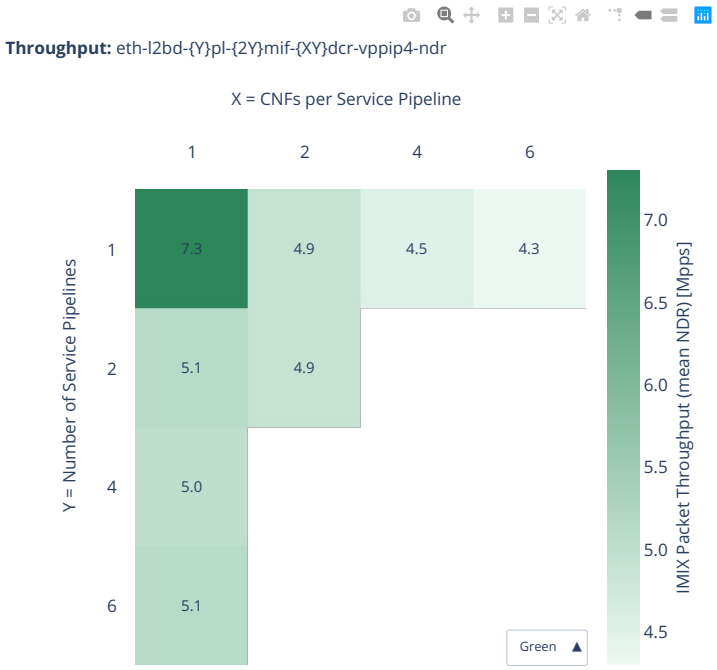

imix-4t2c-eth-l2bd

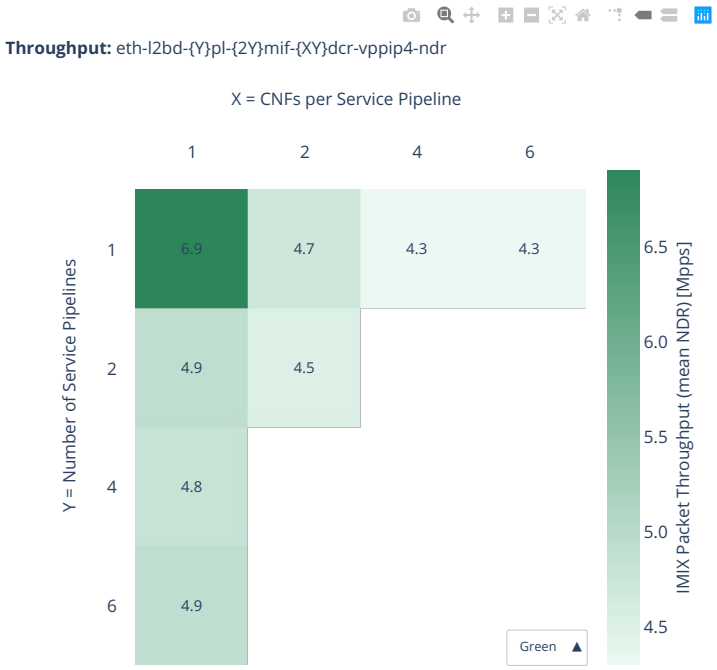

imix-8t4c-eth-l2bd

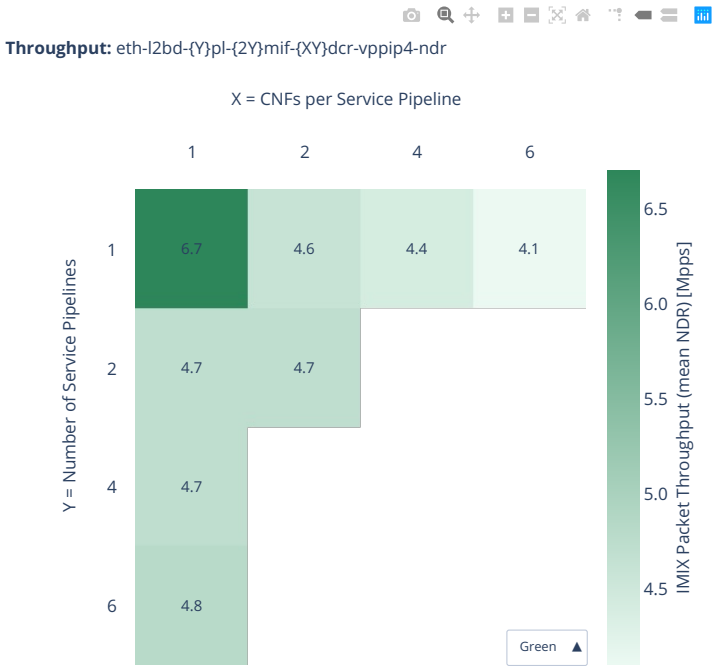

2n-skx-xxv710-pdr

imix-2t1c-eth-l2bd

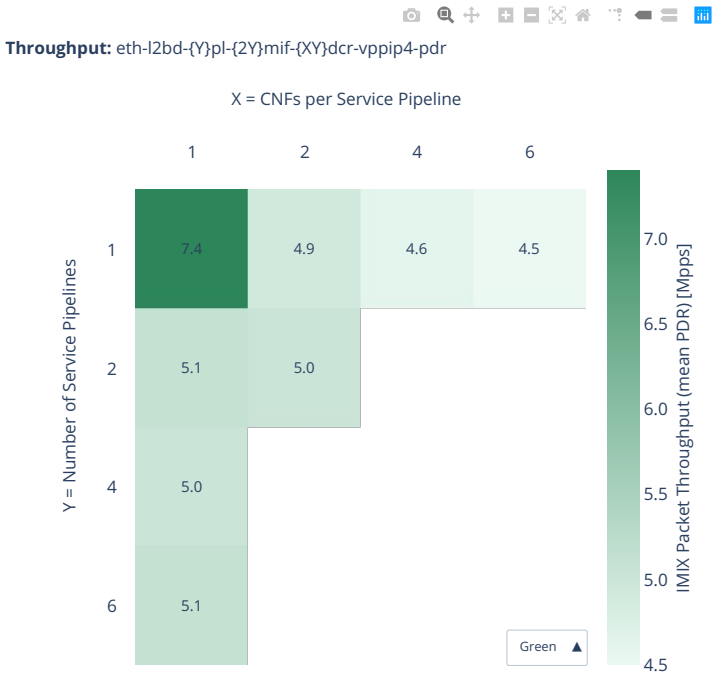

imix-4t2c-eth-l2bd

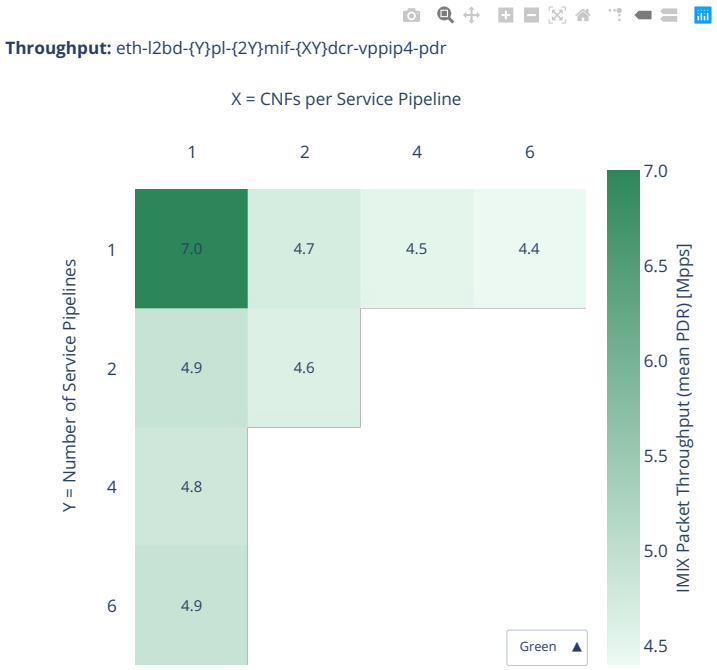

imix-8t4c-eth-l2bd

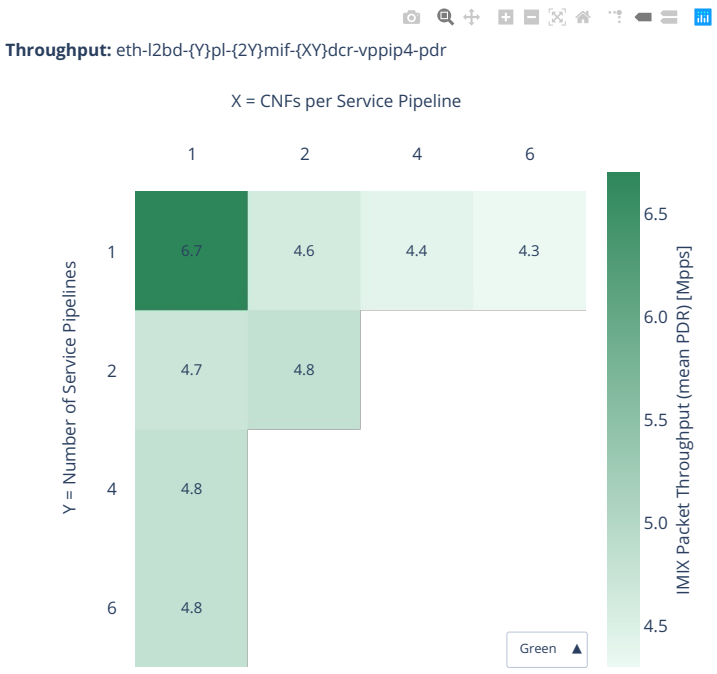

2n-clx-xxv710-mrr

imix-2t1c-eth-l2bd

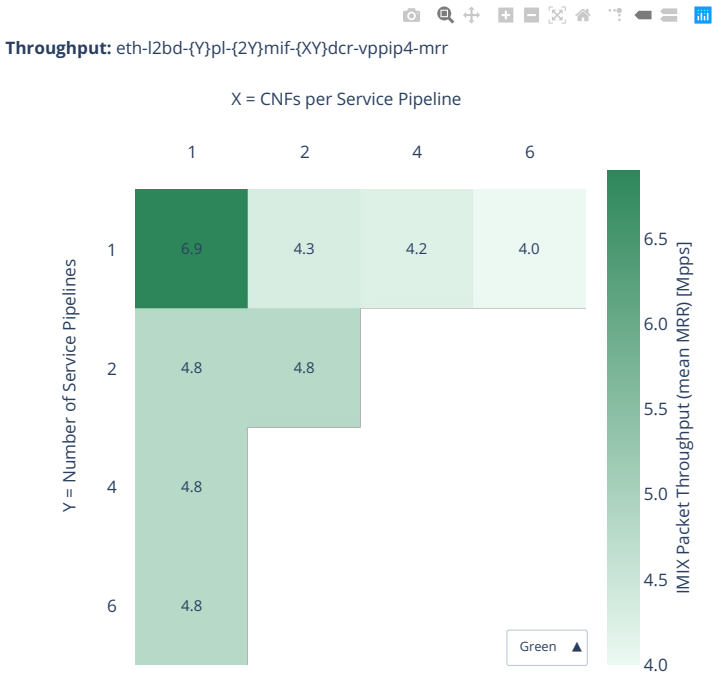

imix-4t2c-eth-l2bd

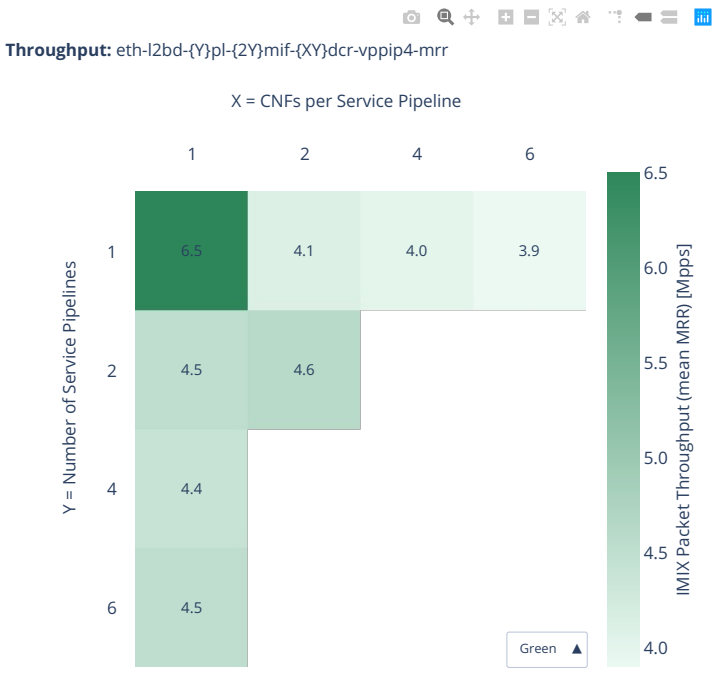

imix-8t4c-eth-l2bd

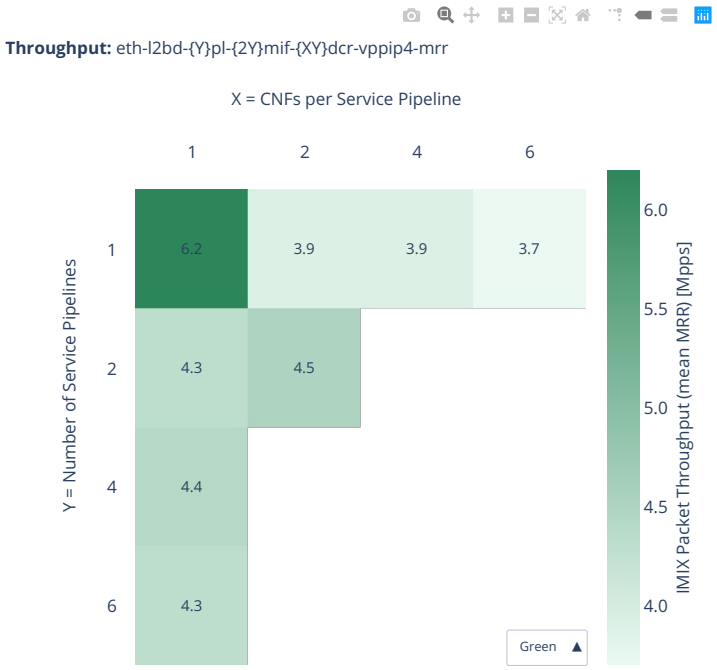

2n-clx-xxv710-ndr

imix-2t1c-eth-l2bd

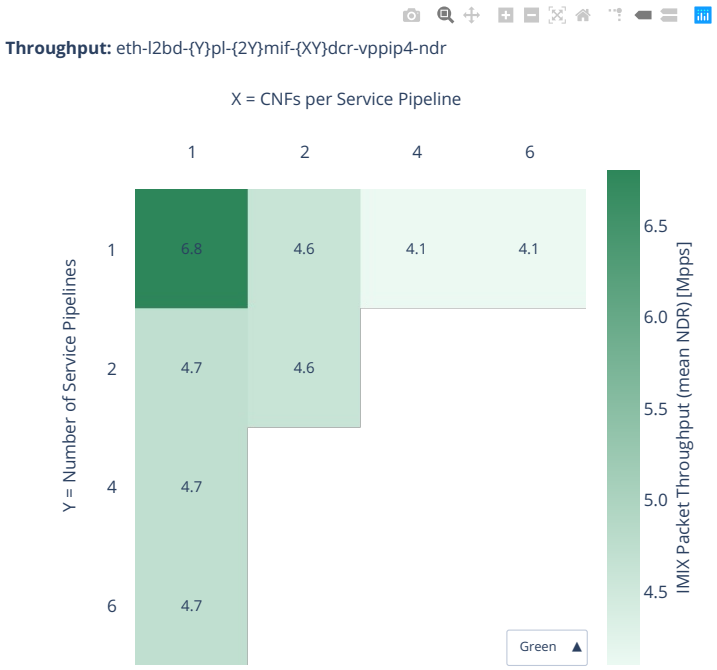

imix-4t2c-eth-l2bd

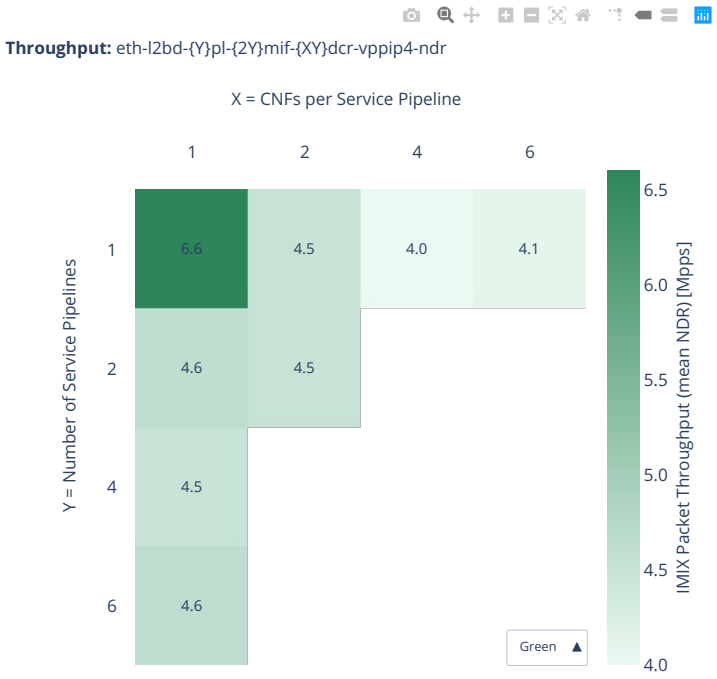

imix-8t4c-eth-l2bd

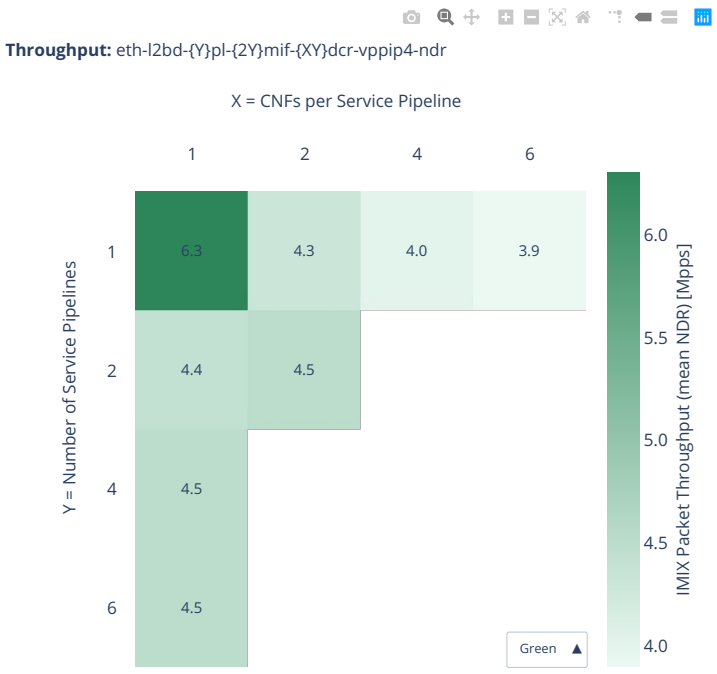

2n-clx-xxv710-pdr

imix-2t1c-eth-l2bd

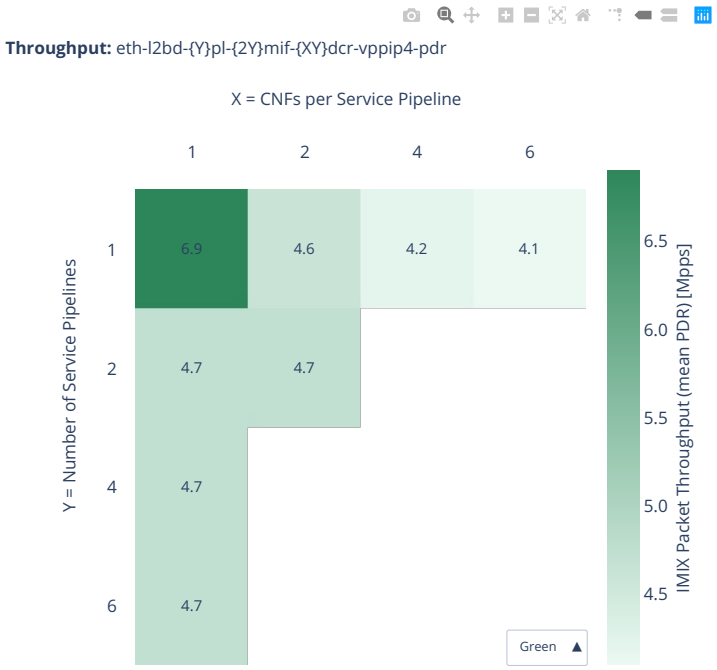

imix-4t2c-eth-l2bd

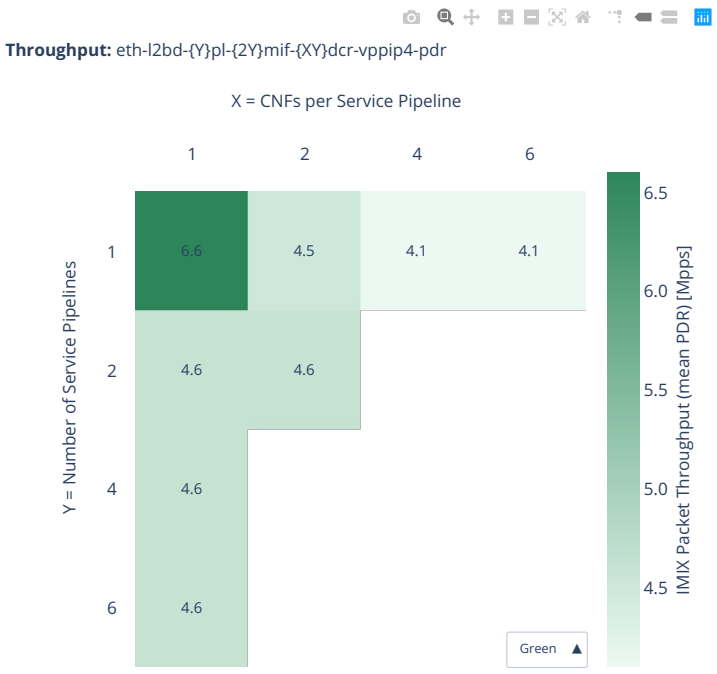

imix-8t4c-eth-l2bd

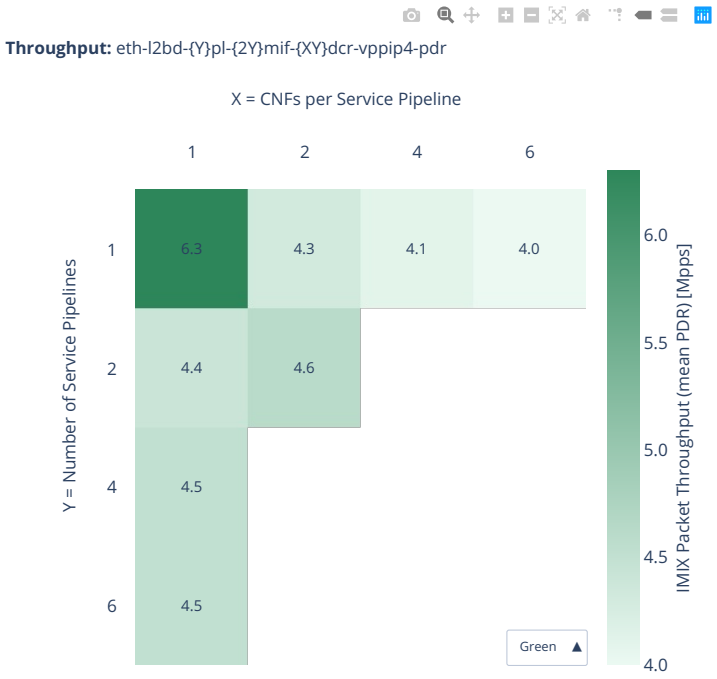

2.8.4 VNF Service Chains Tunnels

Additional information about graph data:

1. **Graph Title**: describes tested packet path including VNF workload running in each VM.
2. **X-axis Labels**: VNFs per service chain.
3. **Y-axis Labels**: number of service chains.
4. **Z-axis Color Scale**: lists 64B/IMIX Packet Throughput (mean MRR/NDR/PDR value) in Mpps or the Relative Difference.
5. **Hover Information**: specific test substring listing vhost-chain-vm combinations, number of runs executed, mean MRR/NDR/PDR throughput in Mpps, standard deviation for both configurations and their relative difference.

Note: Test results are stored in build logs from FD.io vpp performance job 2n-icx[173], build logs from FD.io vpp performance job 2n-skx[174] and build logs from FD.io vpp performance job 2n-clx[175] with RF result files csit-vpp-perf-2206-*.zip archived here.

---

[173] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-icx
[174] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx
[175] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-clx

2n-icx-xxv710-mrr

imix-2t1c-eth-l2bd

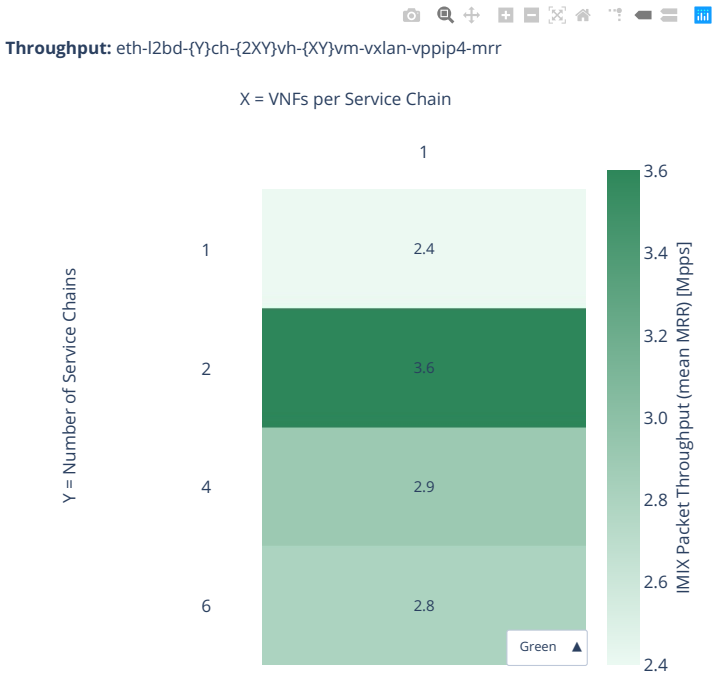

imix-4t2c-eth-l2bd

imix-8t4c-eth-l2bd

2n-icx-xxv710-ndr

imix-2t1c-eth-l2bd

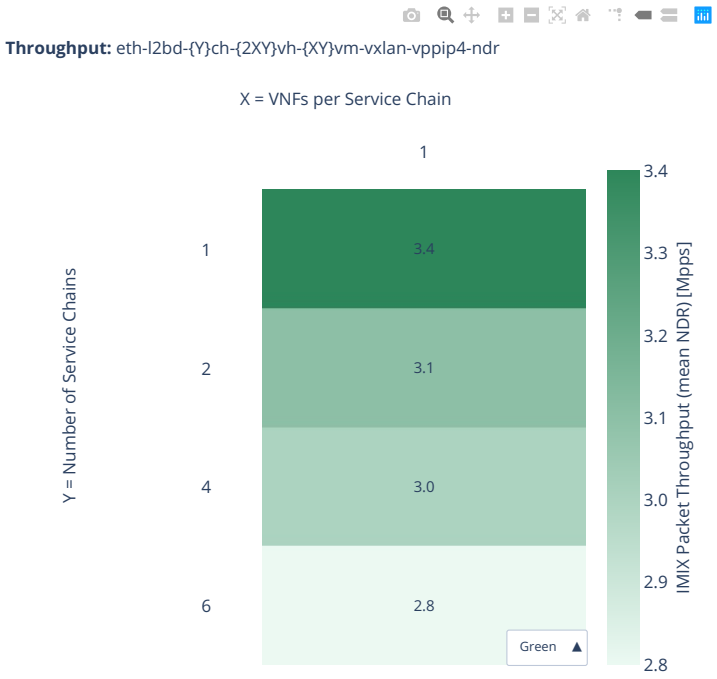

imix-4t2c-eth-l2bd

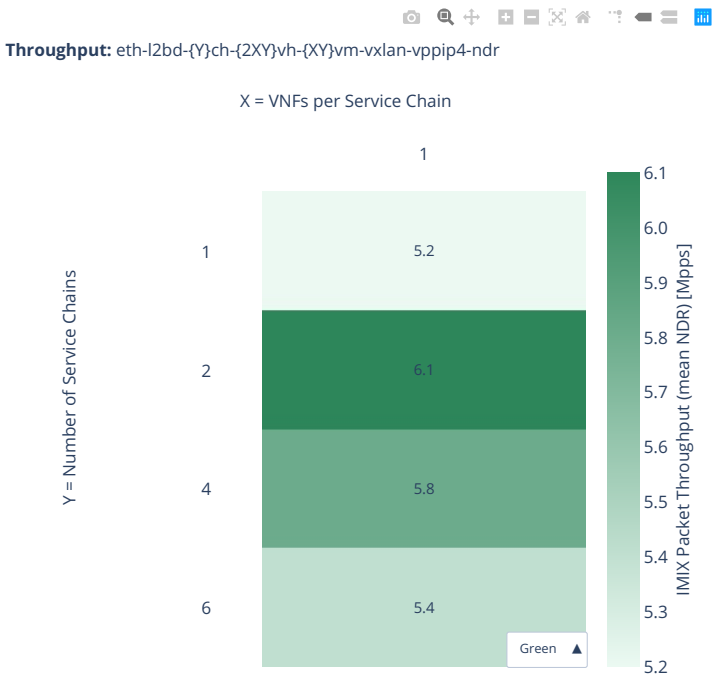

imix-8t4c-eth-l2bd

2n-icx-xxv710-pdr

imix-2t1c-eth-l2bd

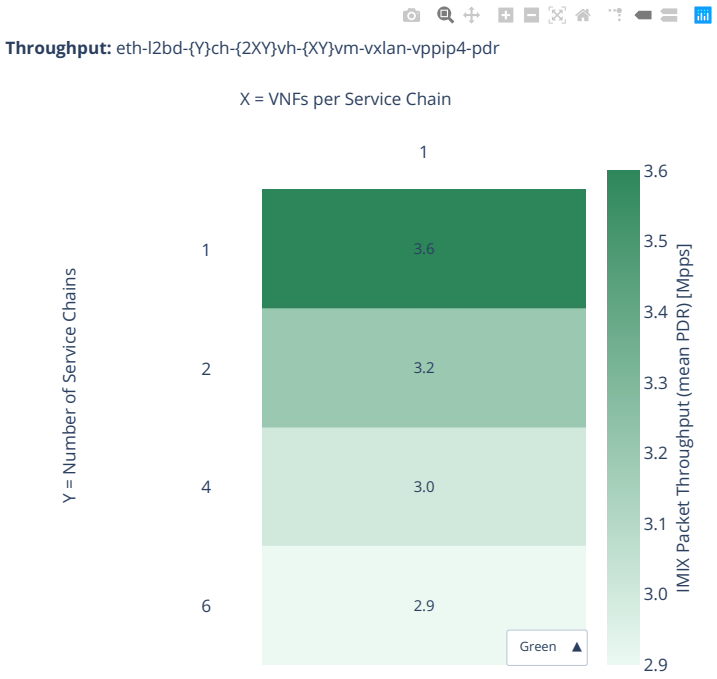

imix-4t2c-eth-l2bd

imix-8t4c-eth-l2bd

2n-skx-xxv710-mrr

imix-2t1c-eth-l2bd

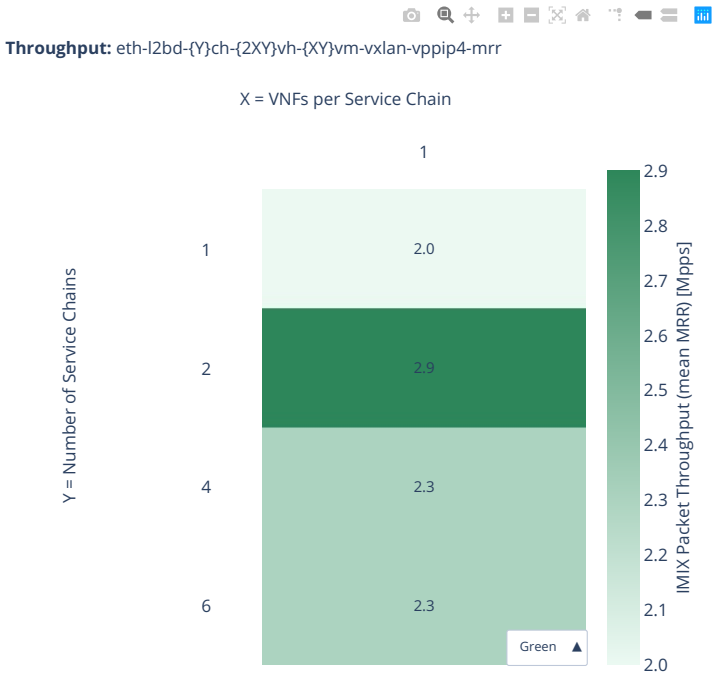

imix-4t2c-eth-l2bd

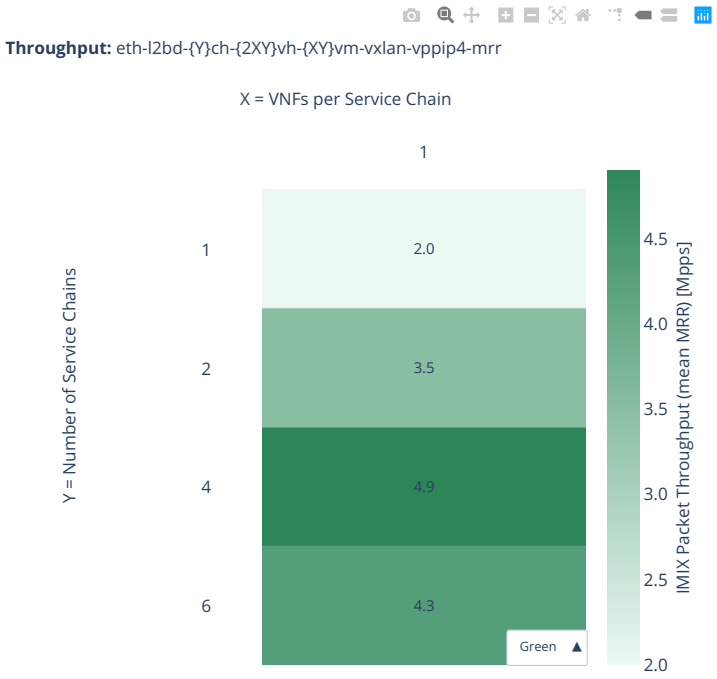

imix-8t4c-eth-l2bd

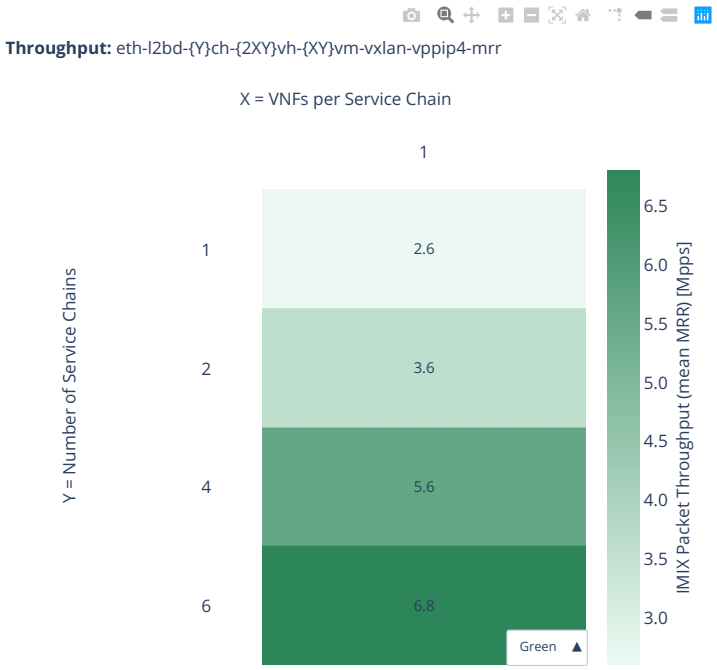

2n-skx-xxv710-ndr

imix-2t1c-eth-l2bd

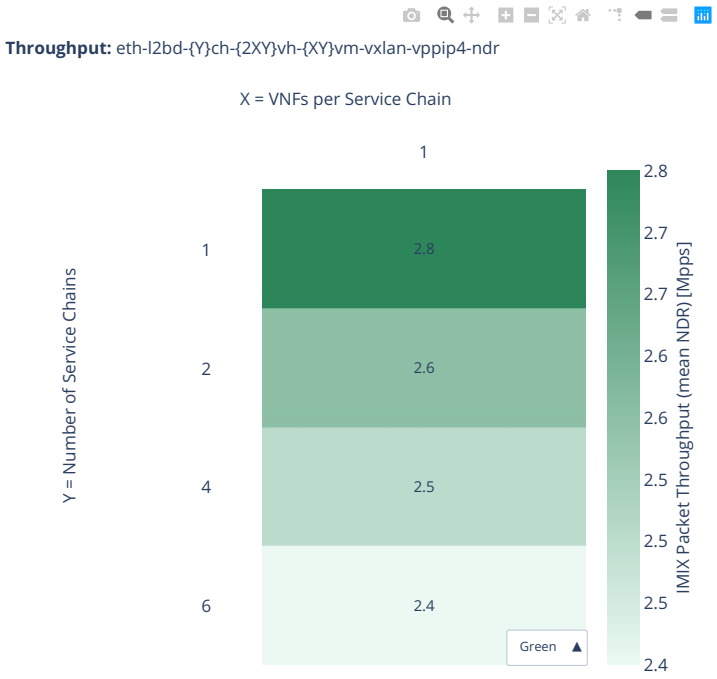

imix-4t2c-eth-l2bd

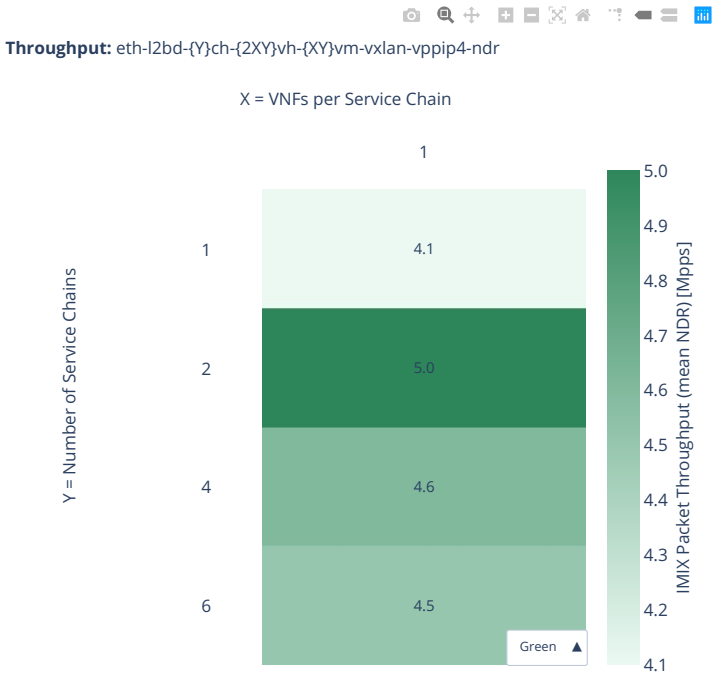

imix-8t4c-eth-l2bd

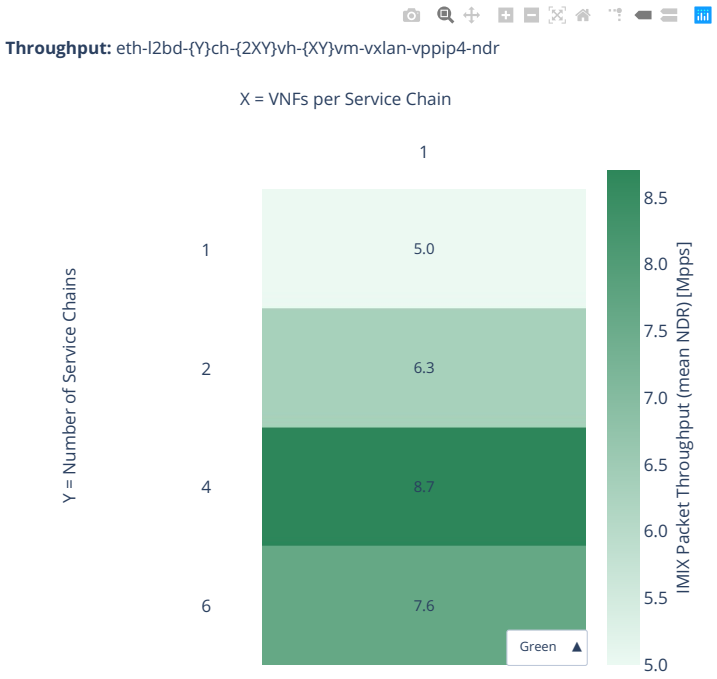

2n-skx-xxv710-pdr

imix-2t1c-eth-l2bd

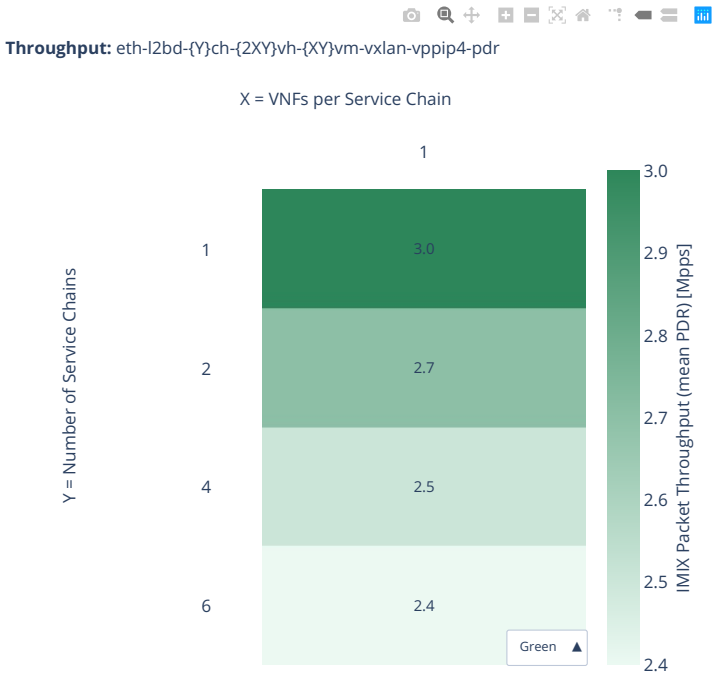

imix-4t2c-eth-l2bd

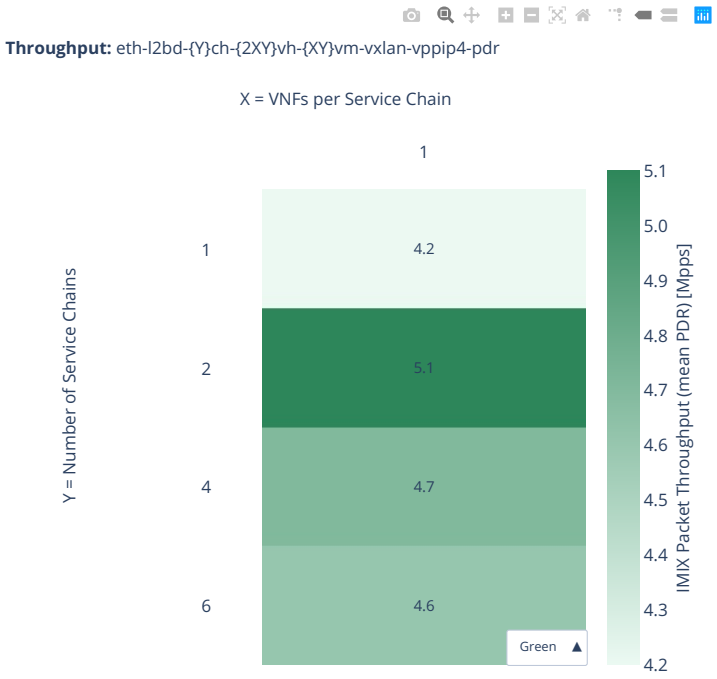

imix-8t4c-eth-l2bd

2n-clx-xxv710-mrr

imix-2t1c-eth-l2bd

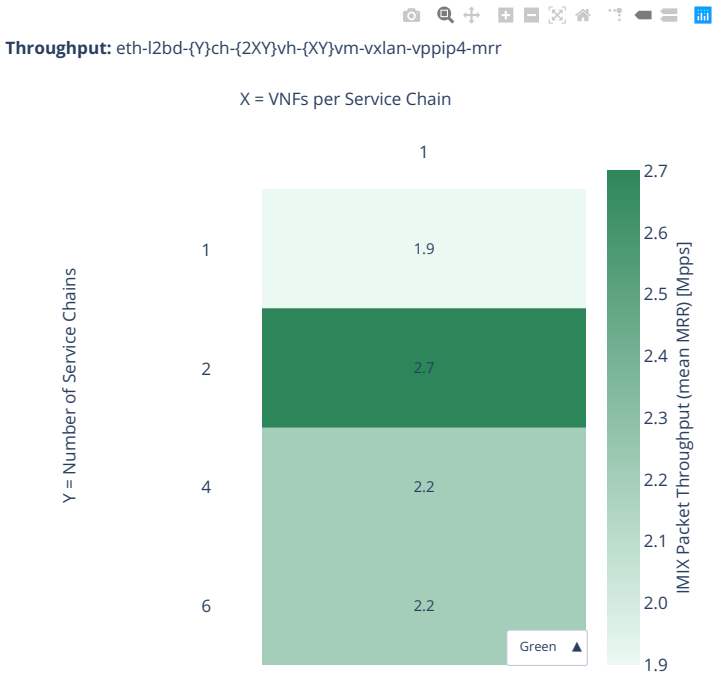

imix-4t2c-eth-l2bd

imix-8t4c-eth-l2bd

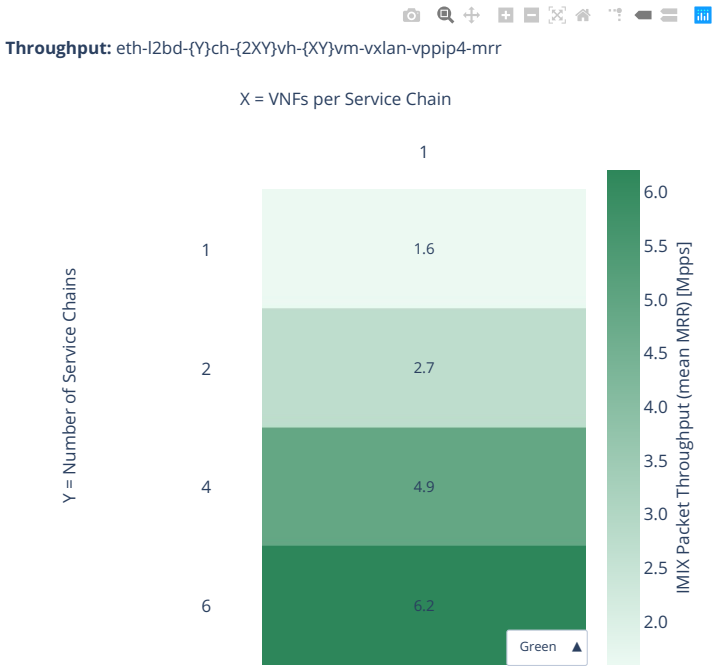

2n-clx-xxv710-ndr

imix-2t1c-eth-l2bd

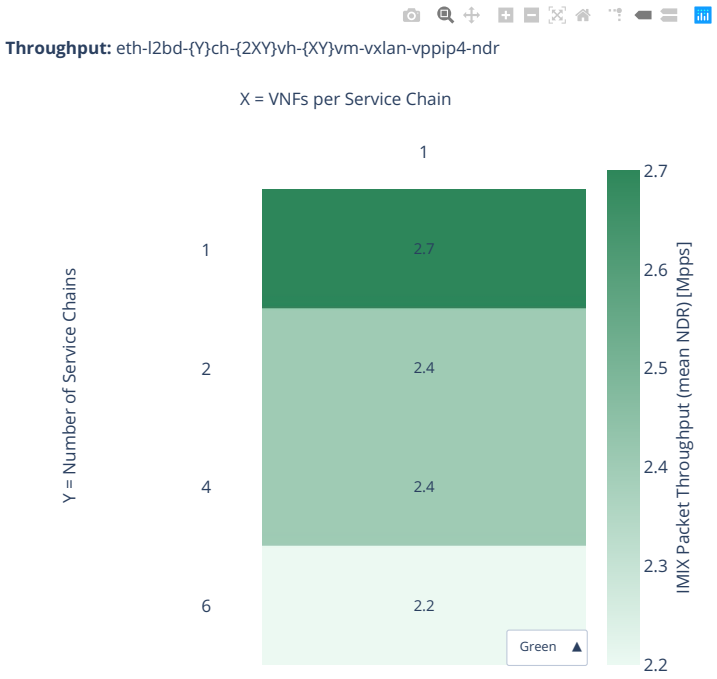

imix-4t2c-eth-l2bd

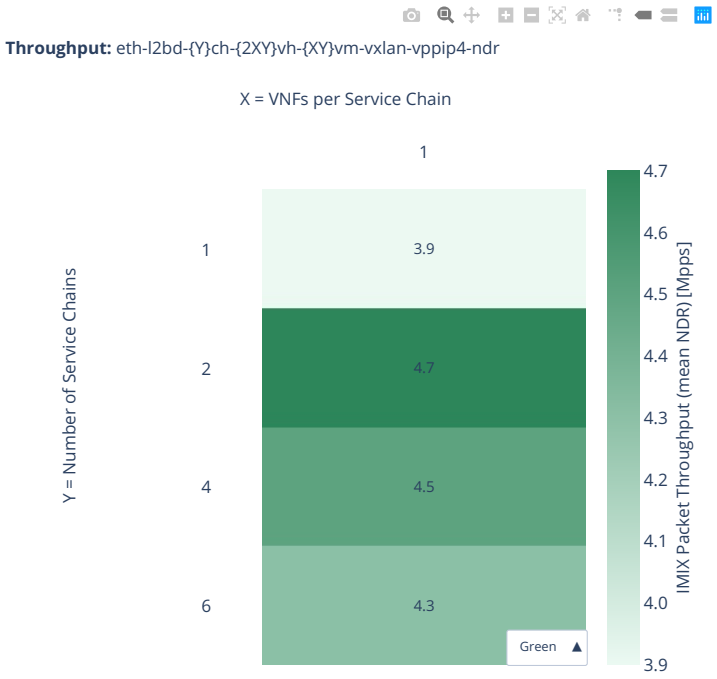

imix-8t4c-eth-l2bd

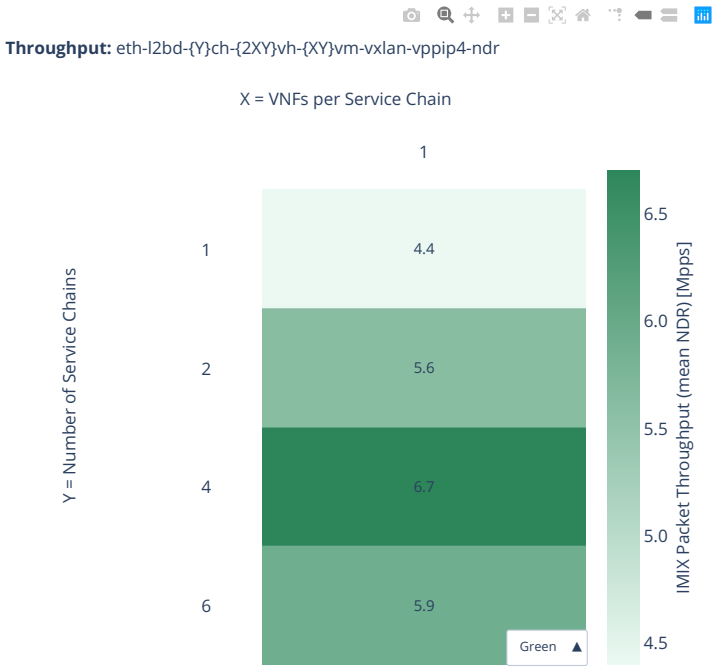

2n-clx-xxv710-pdr

imix-2t1c-eth-l2bd

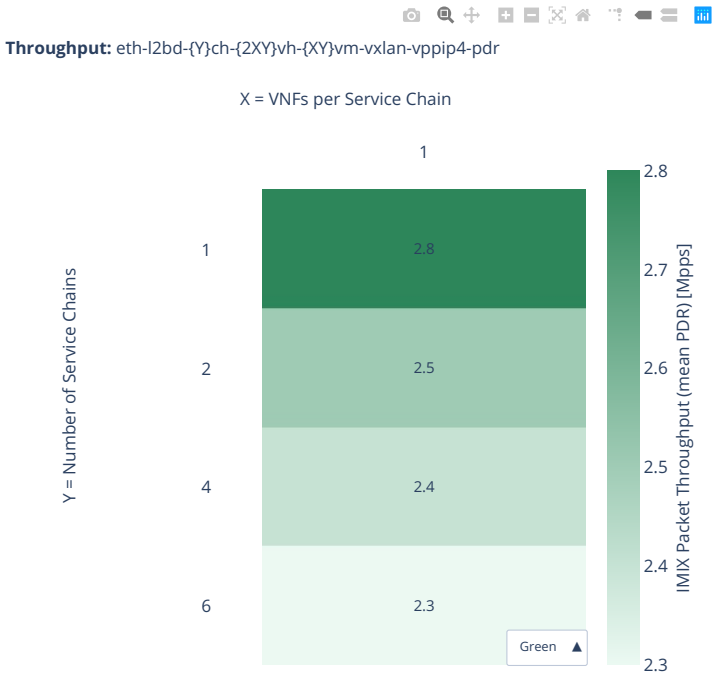

imix-4t2c-eth-l2bd

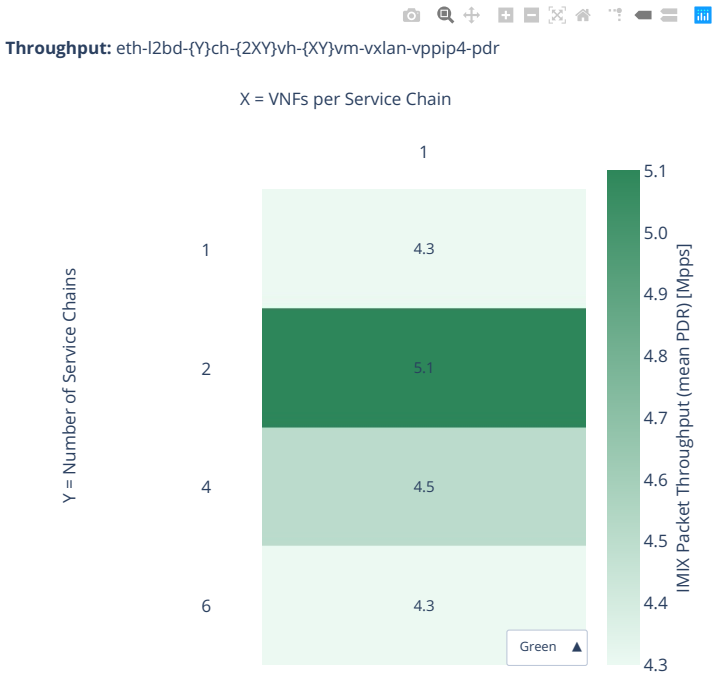

imix-8t4c-eth-l2bd

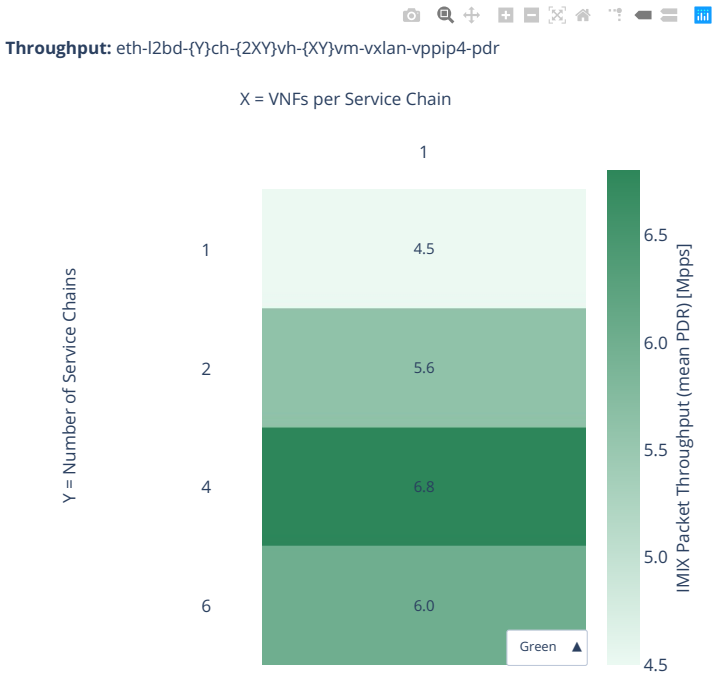

2.9 Hoststack Testing

2.9.1 TCP/IP with iperf3

1460b-2t1c-xxv710-ip4tcp-base-scale

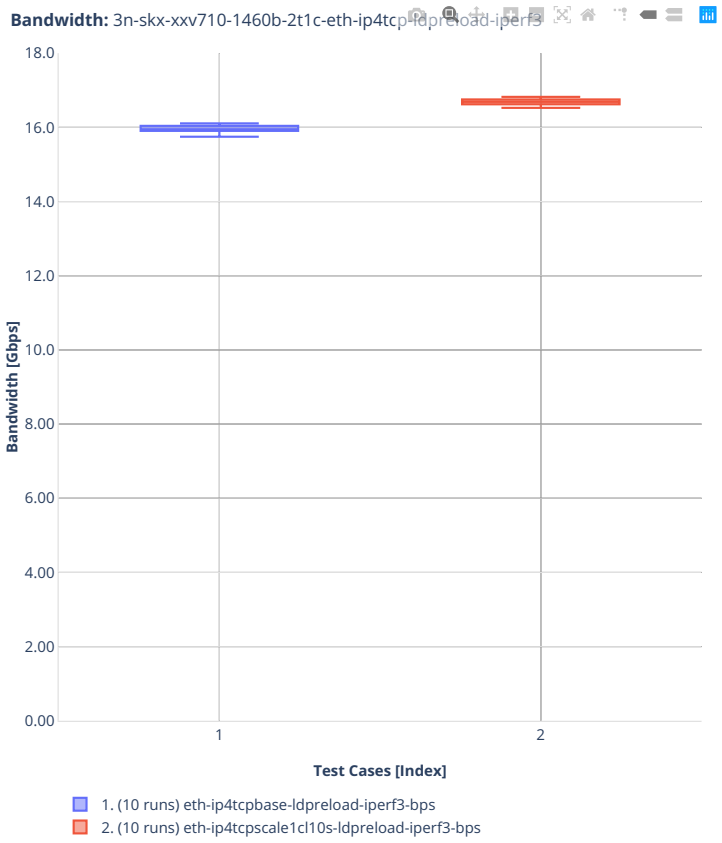

2.9.2 UDP/IP with iperf3

1460b-2t1c-xxv710-ip4udp-base-scale

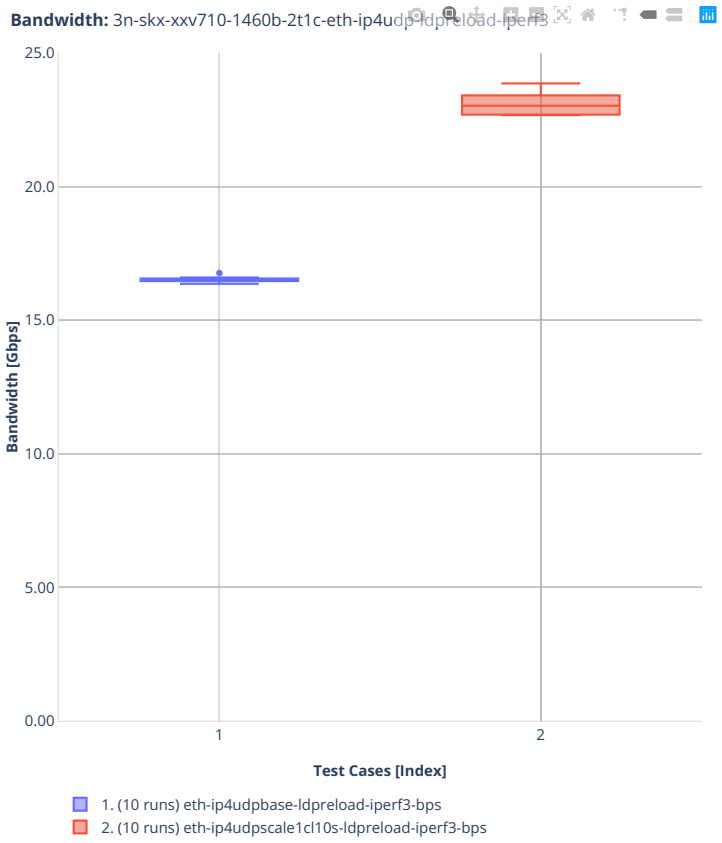

2.9.3 QUIC/UDP/IP with vpp_echo

1280b-2t1c-xxv710-base-scale

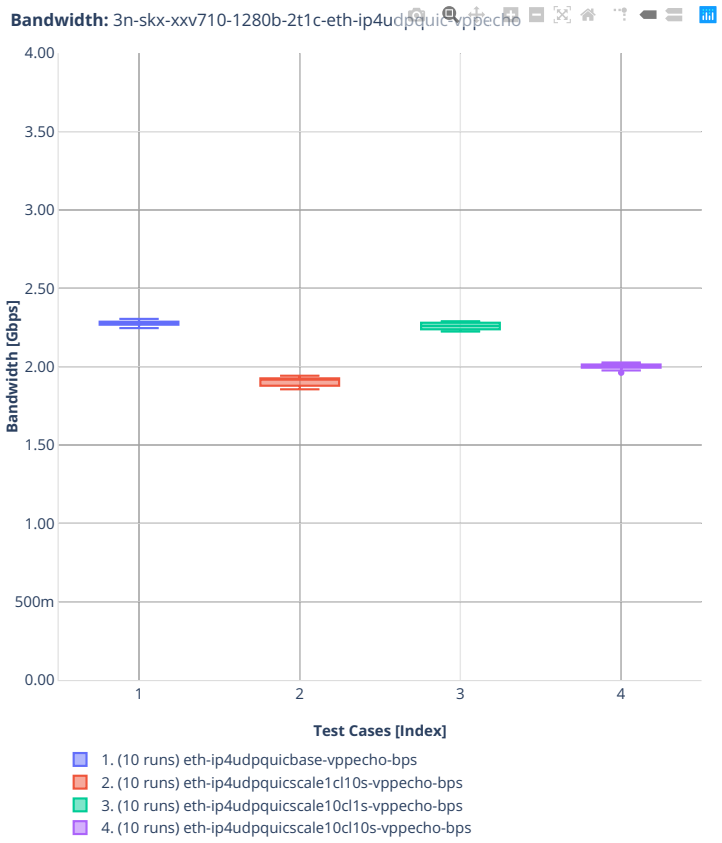

2.9.4 VSAP with ldpreload

2t1c-cx556a-rdma-base-cps

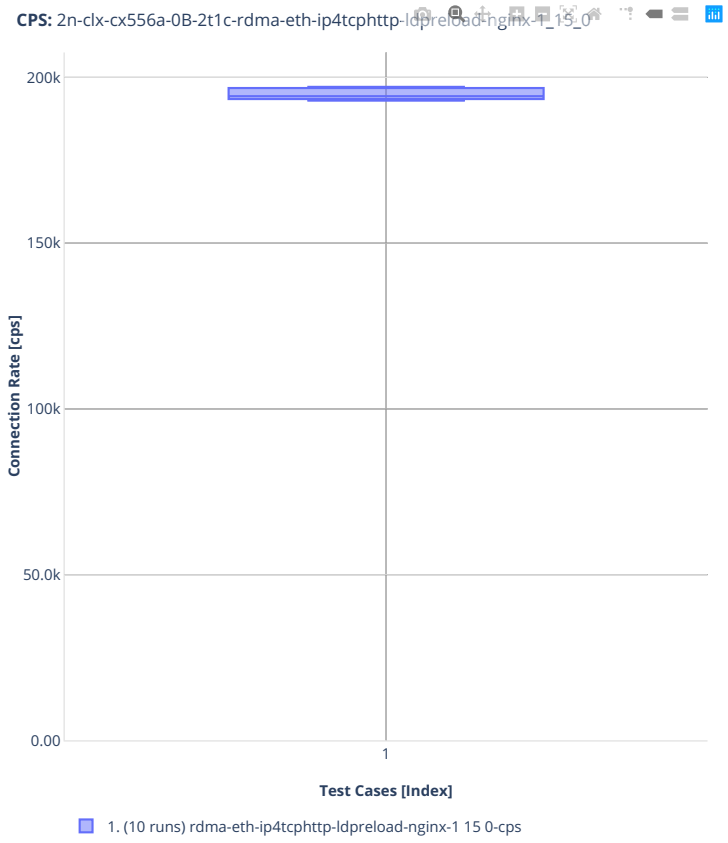

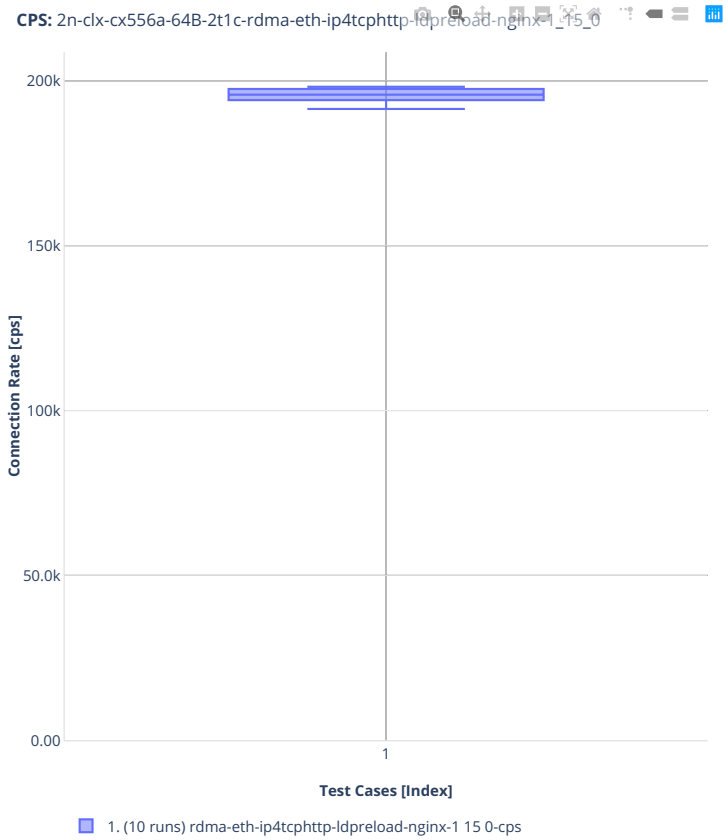
CPS: 2n-clx-cx556a-64B-2t1c-rdma-eth-ip4tcphttp-ldpreload-nginx-1_15_0
200k
150k
100k
50.0k
0.00
Connection Rate [cps]
1
Test Cases [Index]
1. (10 runs) rdma-eth-ip4tcphttp-ldpreload-nginx-1 15 0-cps

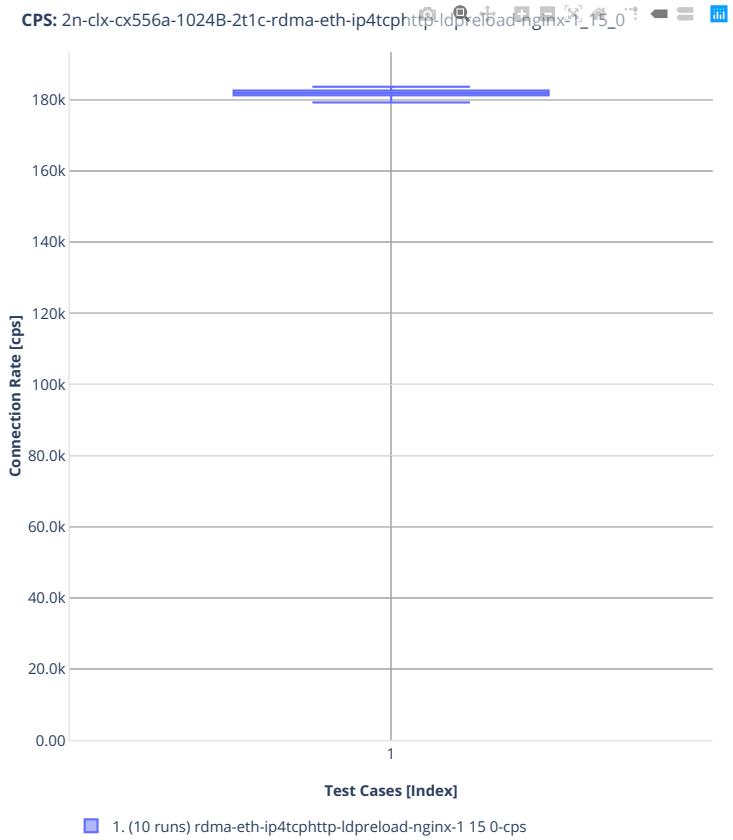
CPS: 2n-clx-cx556a-1024B-2t1c-rdma-eth-ip4tcphttp-ldpreload-nginx-1_15_0
180k
160k
140k
120k
100k
80.0k
60.0k
40.0k
20.0k
0.00
Connection Rate [cps]
1
Test Cases [Index]
1. (10 runs) rdma-eth-ip4tcphttp-ldpreload-nginx-1 15 0-cps

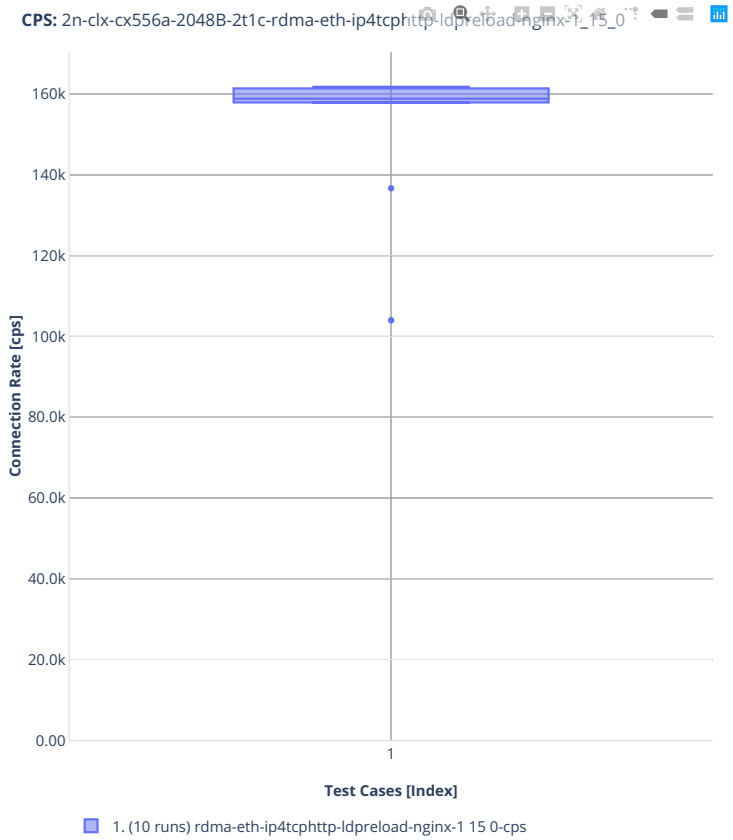
CPS: 2n-clx-cx556a-2048B-2t1c-rdma-eth-ip4tcphttp-ldpreload-nginx-1_15_0
160k
140k
120k
100k
80.0k
60.0k
40.0k
20.0k
0.00
Connection Rate [cps]
1
Test Cases [Index]
1. (10 runs) rdma-eth-ip4tcphttp-ldpreload-nginx-1 15 0-cps

2t1c-cx556a-rdma-base-rps

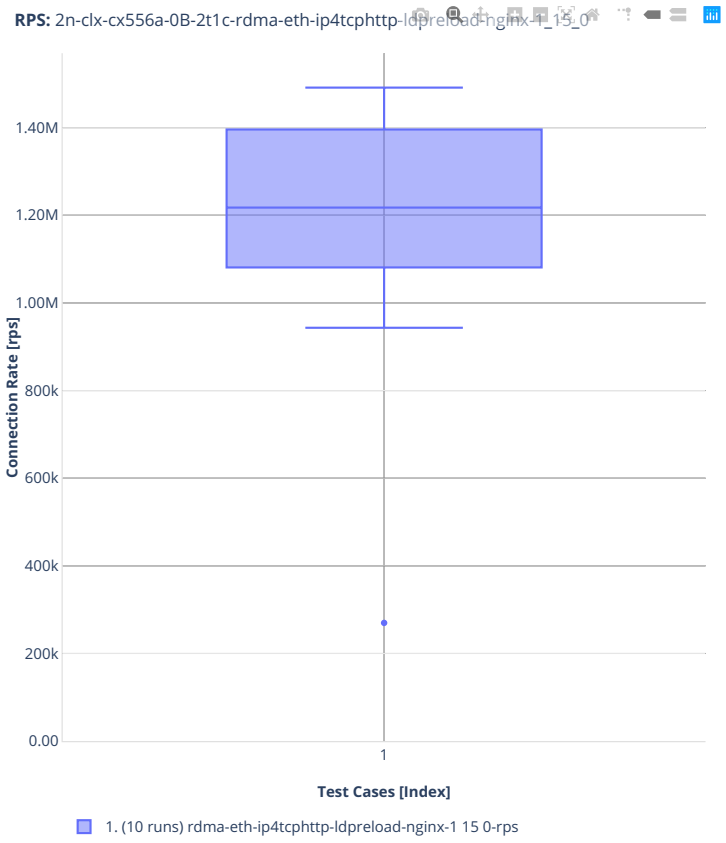

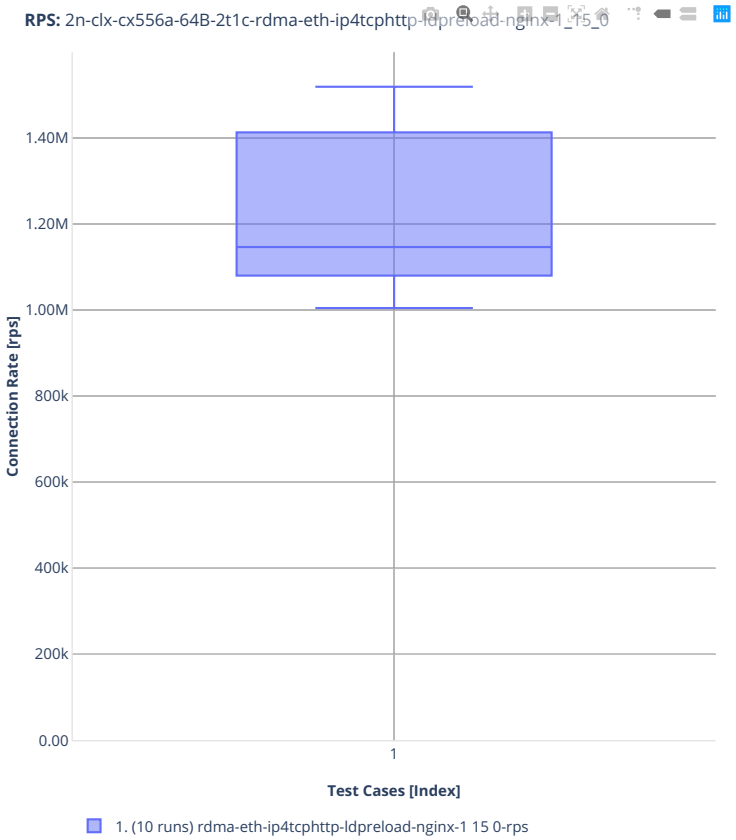
RPS: 2n-clx-cx556a-64B-2t1c-rdma-eth-ip4tcphttp-ldpreload-nginx-1_15_0
1.40M
1.20M
1.00M
800k
600k
400k
200k
0.00
Connection Rate [rps]
1
Test Cases [Index]
1. (10 runs) rdma-eth-ip4tcphttp-ldpreload-nginx-1 15 0-rps

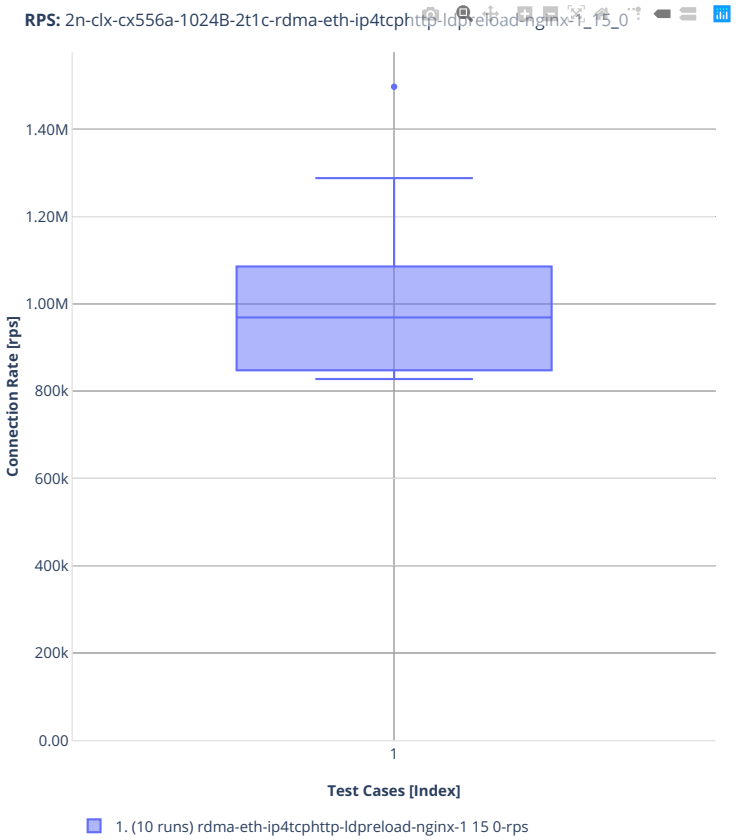
RPS: 2n-clx-cx556a-1024B-2t1c-rdma-eth-ip4tcphttp-ldpreload-nginx-1_15_0
1.40M
1.20M
1.00M
800k
600k
400k
200k
0.00
Connection Rate [rps]
1
Test Cases [Index]
1. (10 runs) rdma-eth-ip4tcphttp-ldpreload-nginx-1 15 0-rps

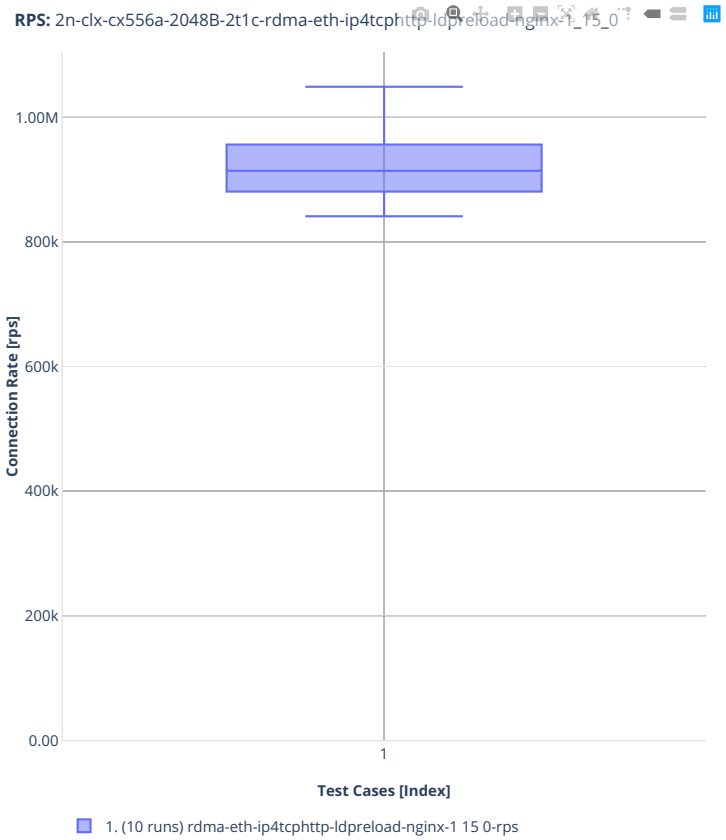
RPS: 2n-clx-cx556a-2048B-2t1c-rdma-eth-ip4tcphttp-ldpreload-nginx-1_15_0
1.00M
800k
600k
400k
200k
0.00
Connection Rate [rps]
1
Test Cases [Index]
1. (10 runs) rdma-eth-ip4tcphttp-ldpreload-nginx-1 15 0-rps

2t1c-e810cq-base-cps

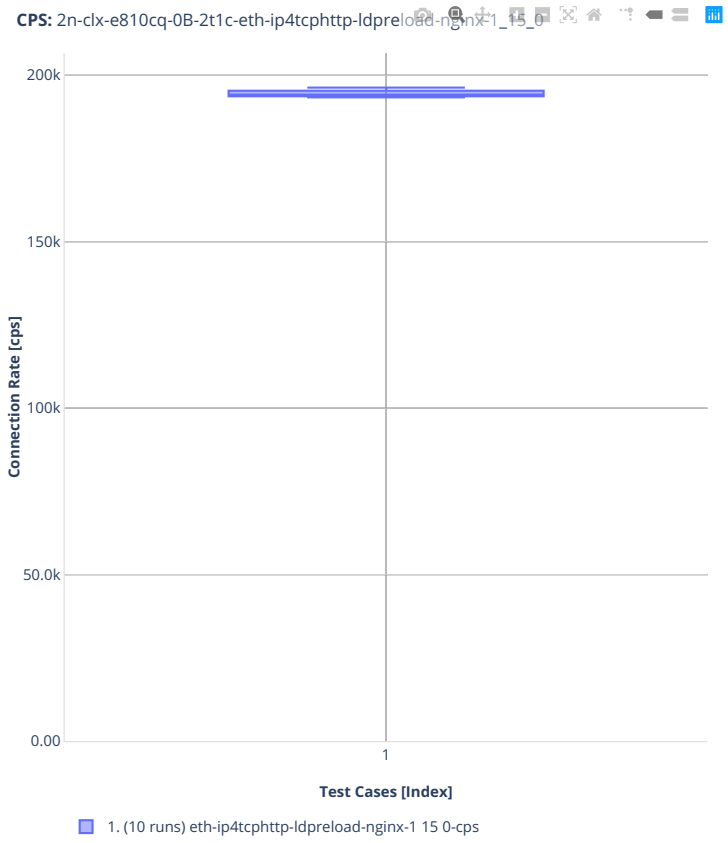

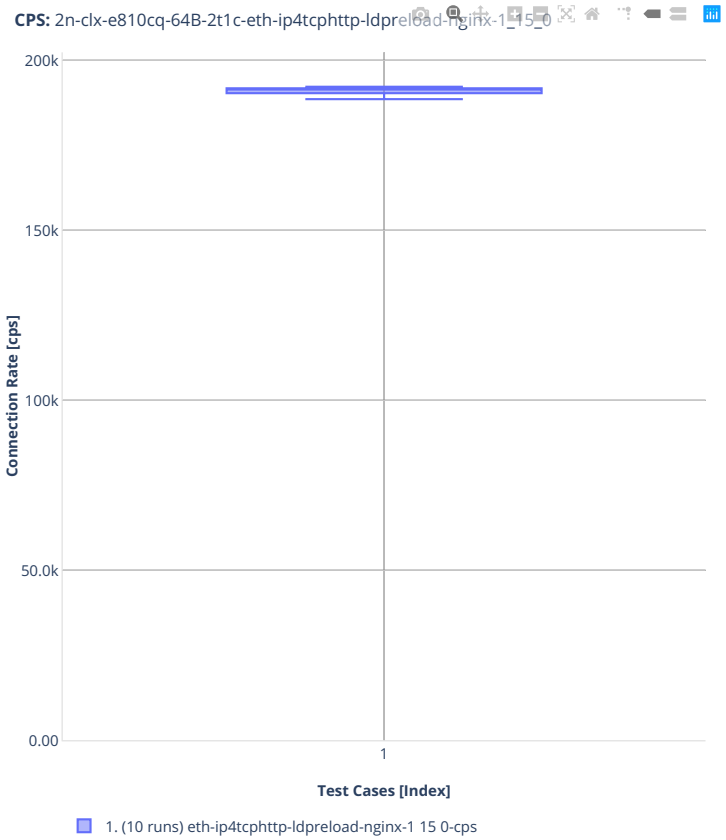
CPS: 2n-clx-e810cq-64B-2t1c-eth-ip4tcphttp-ldpreload-nginx-1_15_0
200k
150k
100k
50.0k
0.00
Connection Rate [cps]
1
Test Cases [Index]
1. (10 runs) eth-ip4tcphttp-ldpreload-nginx-1 15 0-cps

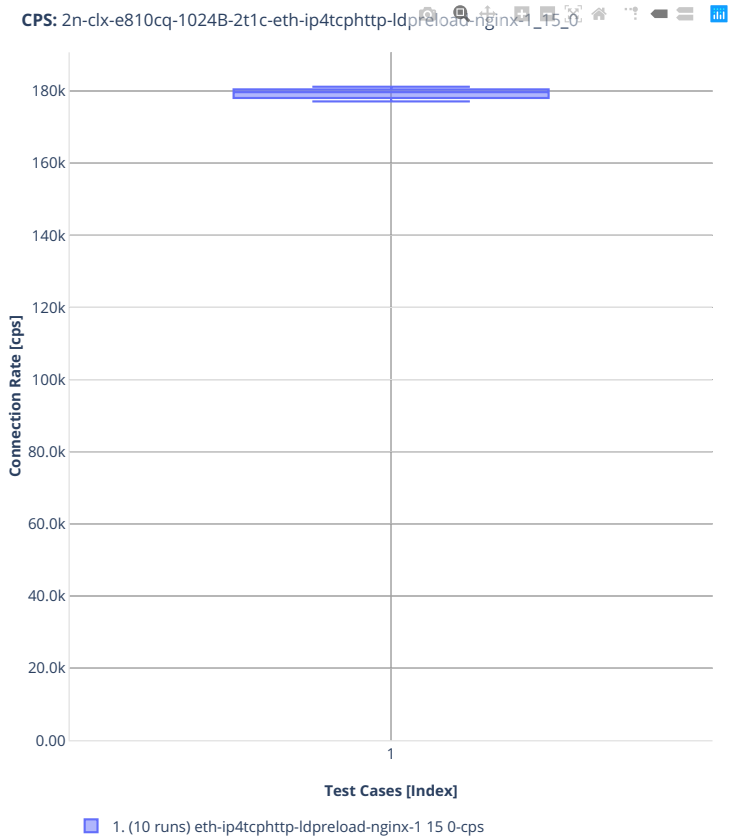
CPS: 2n-clx-e810cq-1024B-2t1c-eth-ip4tcphttp-ldpreload-nginx-1_15_0
180k
160k
140k
120k
100k
80.0k
60.0k
40.0k
20.0k
0.00
Connection Rate [cps]
1
Test Cases [Index]
1. (10 runs) eth-ip4tcphttp-ldpreload-nginx-1 15 0-cps

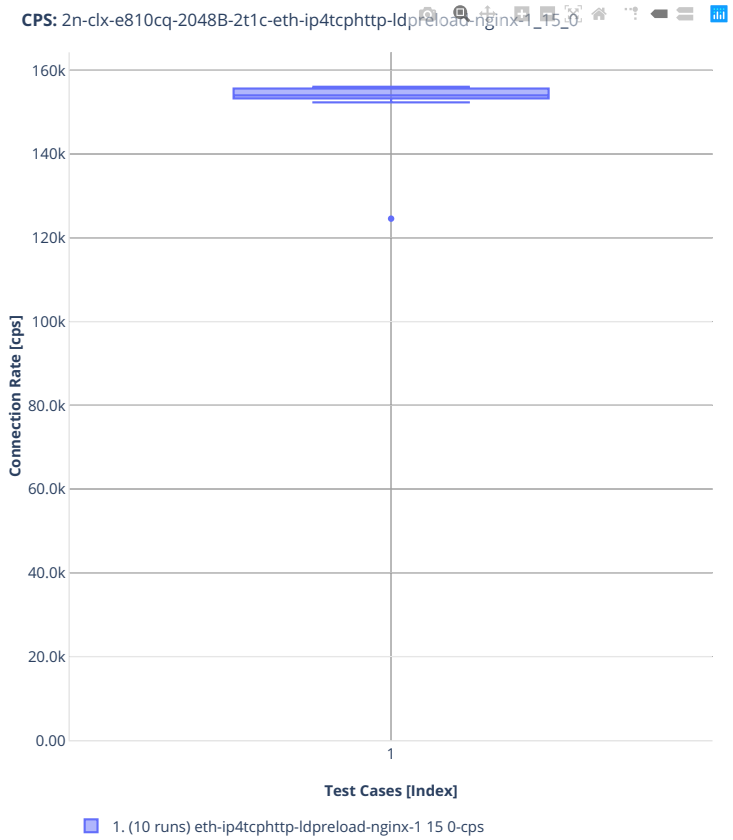
CPS: 2n-clx-e810cq-2048B-2t1c-eth-ip4tcphttp-ldpreload-nginx-1_15_0
160k
140k
120k
100k
80.0k
60.0k
40.0k
20.0k
0.00
Connection Rate [cps]
1
Test Cases [Index]
1. (10 runs) eth-ip4tcphttp-ldpreload-nginx-1 15 0-cps

2t1c-e810cq-base-rps

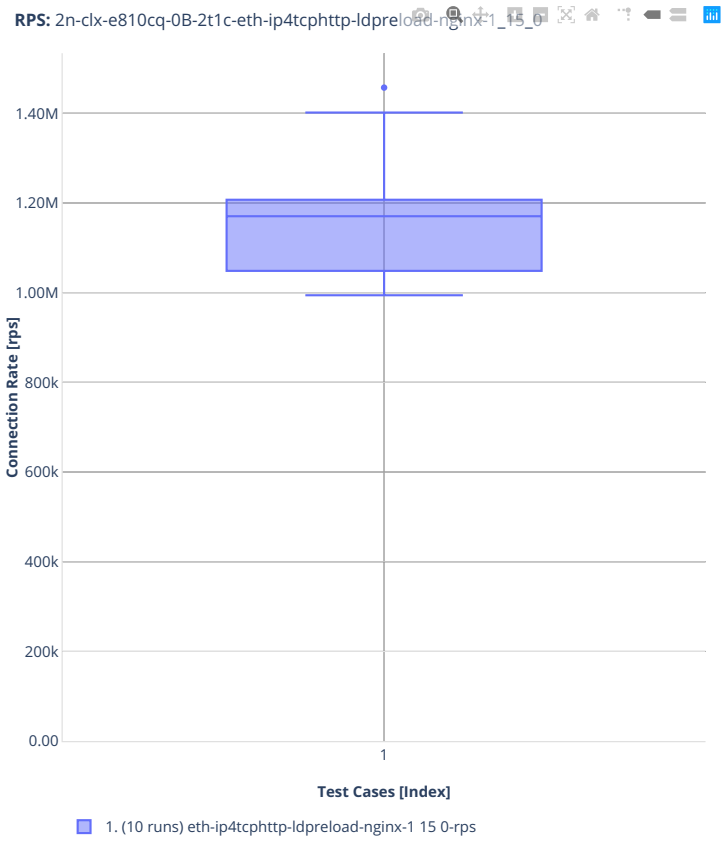

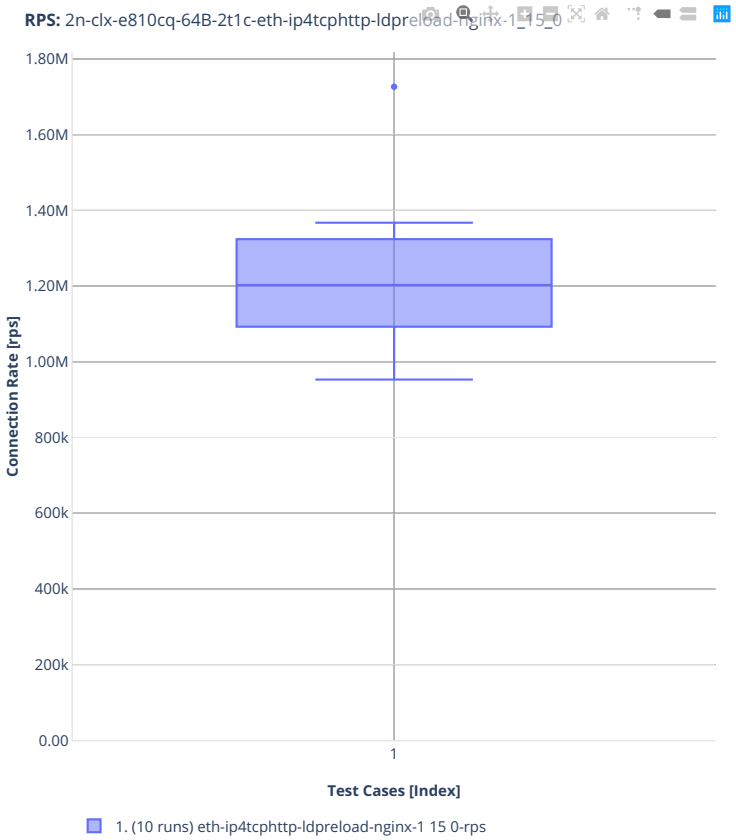
RPS: 2n-clx-e810cq-64B-2t1c-eth-ip4tcphttp-ldpreload-nginx-1_15_0
1.80M
1.60M
1.40M
1.20M
1.00M
800k
600k
400k
200k
0.00
Connection Rate [rps]
1
Test Cases [Index]
1. (10 runs) eth-ip4tcphttp-ldpreload-nginx-1 15 0-rps

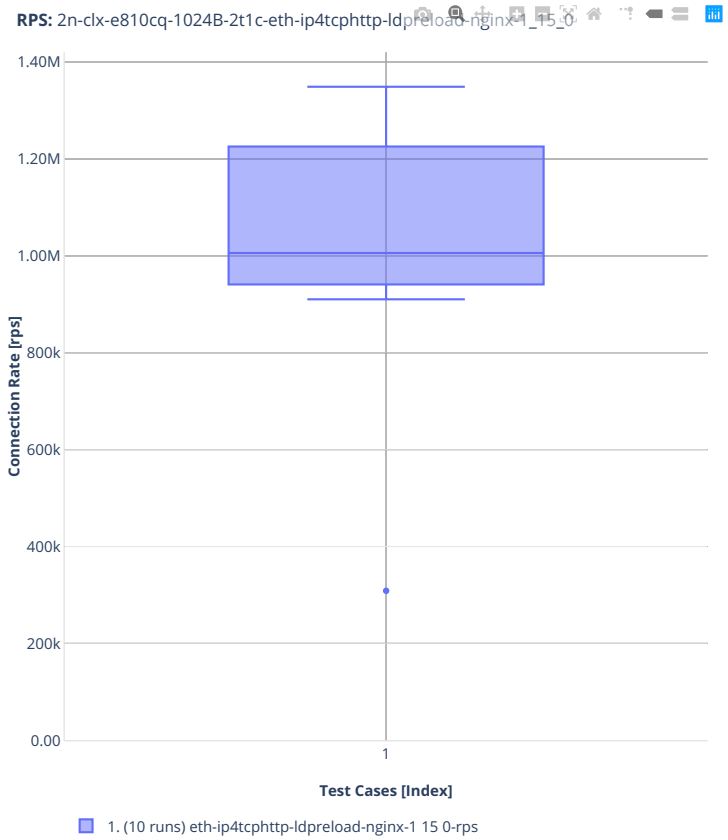
RPS: 2n-clx-e810cq-1024B-2t1c-eth-ip4tcphttp-ldpreload-nginx-1_15_0
1.40M
1.20M
1.00M
800k
600k
400k
200k
0.00
Connection Rate [rps]
1
Test Cases [Index]
1. (10 runs) eth-ip4tcphttp-ldpreload-nginx-1 15 0-rps

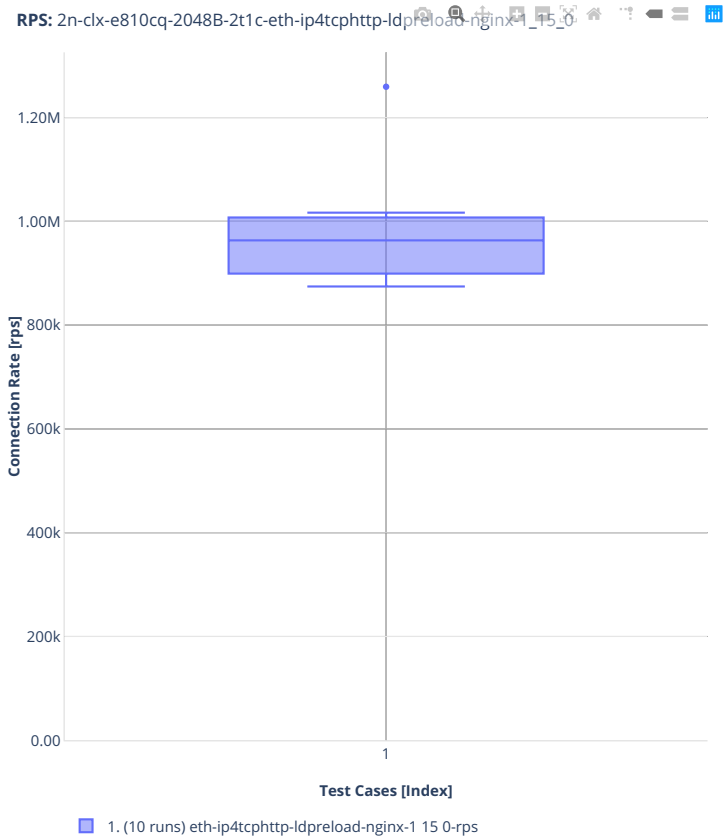
RPS: 2n-clx-e810cq-2048B-2t1c-eth-ip4tcphttp-ldpreload-nginx-1_15_0
1.20M
1.00M
800k
600k
400k
200k
0.00
Connection Rate [rps]
1
Test Cases [Index]
1. (10 runs) eth-ip4tcphttp-ldpreload-nginx-1 15 0-rps

2.10 GSO Testing

2.10.1 2n-icx

1t1c

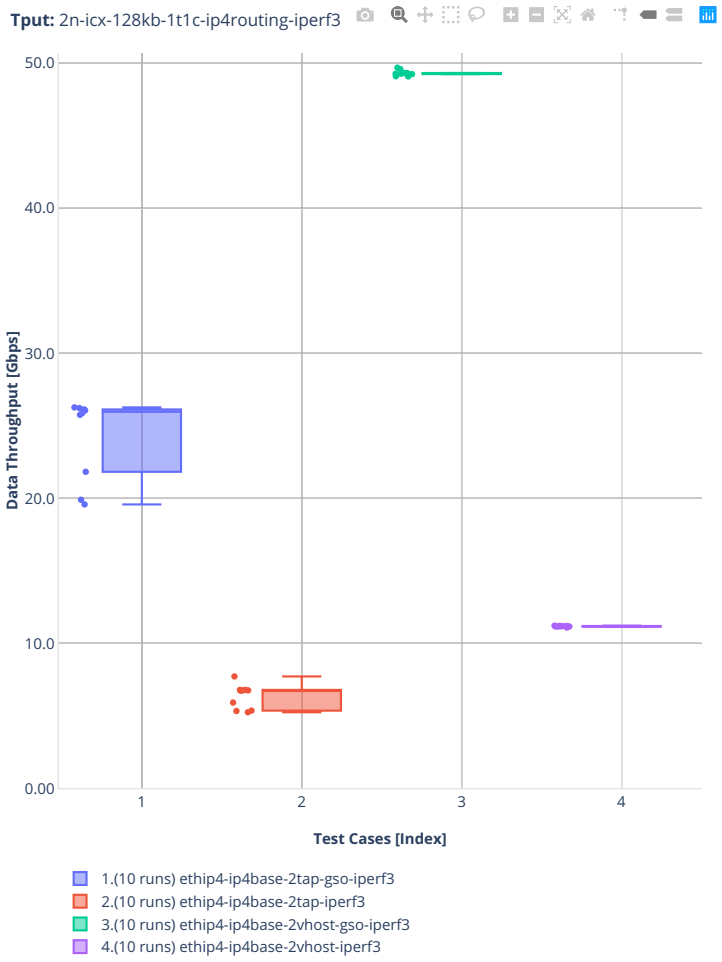
Tput: 2n-icx-128kb-1t1c-ip4routing-iperf3
50.0
40.0
30.0
20.0
10.0
0.00
Data Throughput [Gbps]
1
2
3
4
Test Cases [Index]
1.(10 runs) ethip4-ip4base-2tap-gso-iperf3
2.(10 runs) ethip4-ip4base-2tap-iperf3
3.(10 runs) ethip4-ip4base-2vhost-gso-iperf3
4.(10 runs) ethip4-ip4base-2vhost-iperf3

2.10.2 2n-skx

1t1c

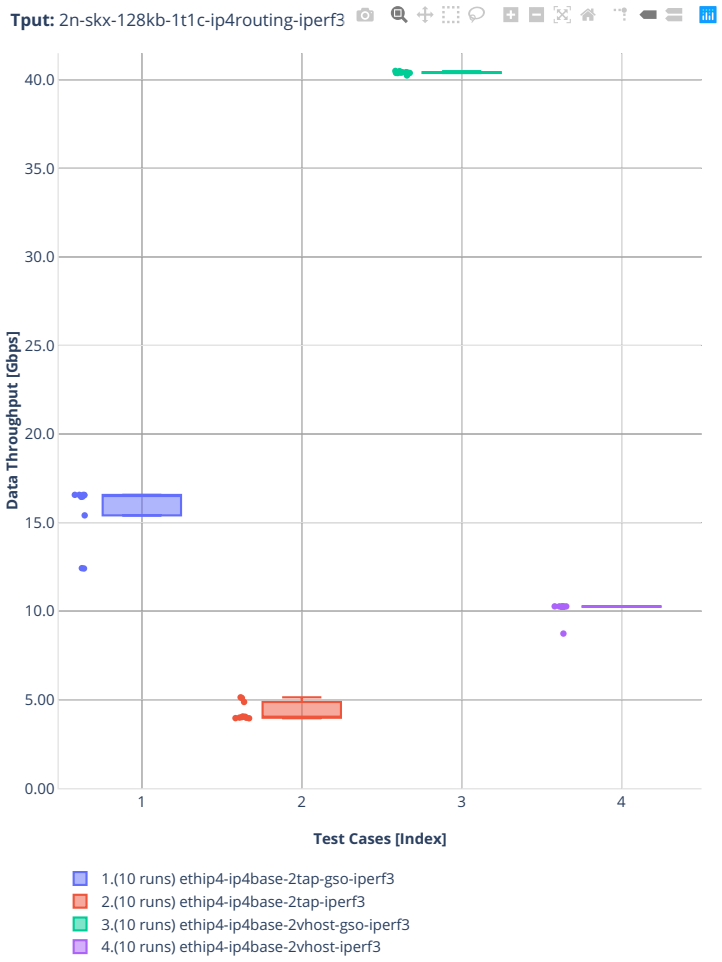
Tput: 2n-skx-128kb-1t1c-ip4routing-iperf3
Data Throughput [Gbps]
40.0
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
1
2
3
4
Test Cases [Index]
1.(10 runs) ethip4-ip4base-2tap-gso-iperf3
2.(10 runs) ethip4-ip4base-2tap-iperf3
3.(10 runs) ethip4-ip4base-2vhost-gso-iperf3
4.(10 runs) ethip4-ip4base-2vhost-iperf3

2.10.3 2n-clx

1t1c

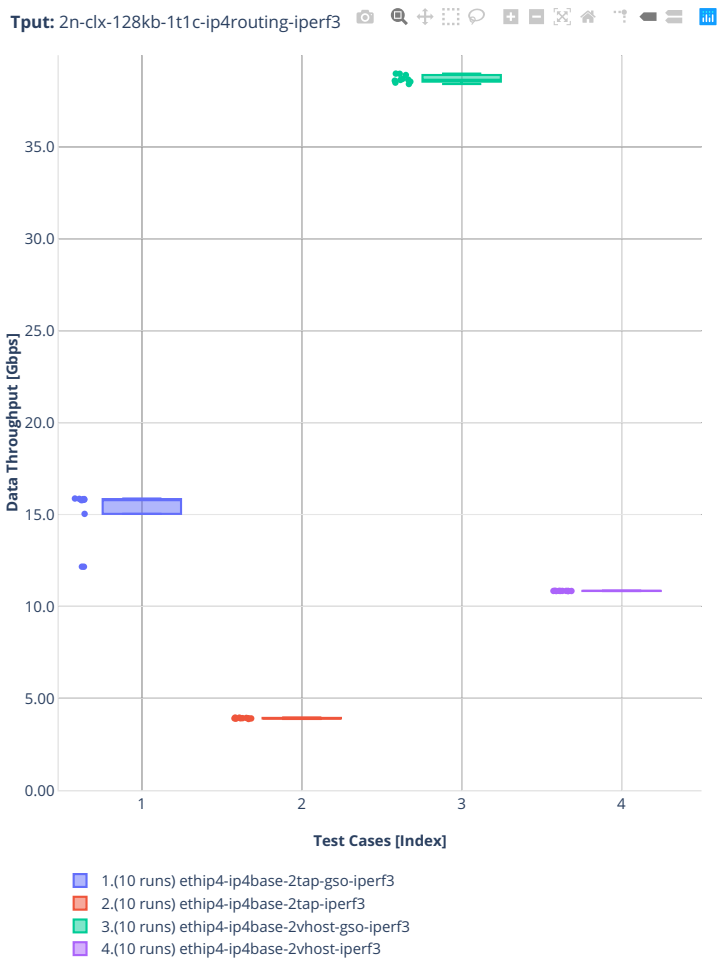
Tput: 2n-clx-128kb-1t1c-ip4routing-iperf3
Data Throughput [Gbps]
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
1
2
3
4
Test Cases [Index]
1.(10 runs) ethip4-ip4base-2tap-gso-iperf3
2.(10 runs) ethip4-ip4base-2tap-iperf3
3.(10 runs) ethip4-ip4base-2vhost-gso-iperf3
4.(10 runs) ethip4-ip4base-2vhost-iperf3

2.10.4 2n-zn2

1t1c

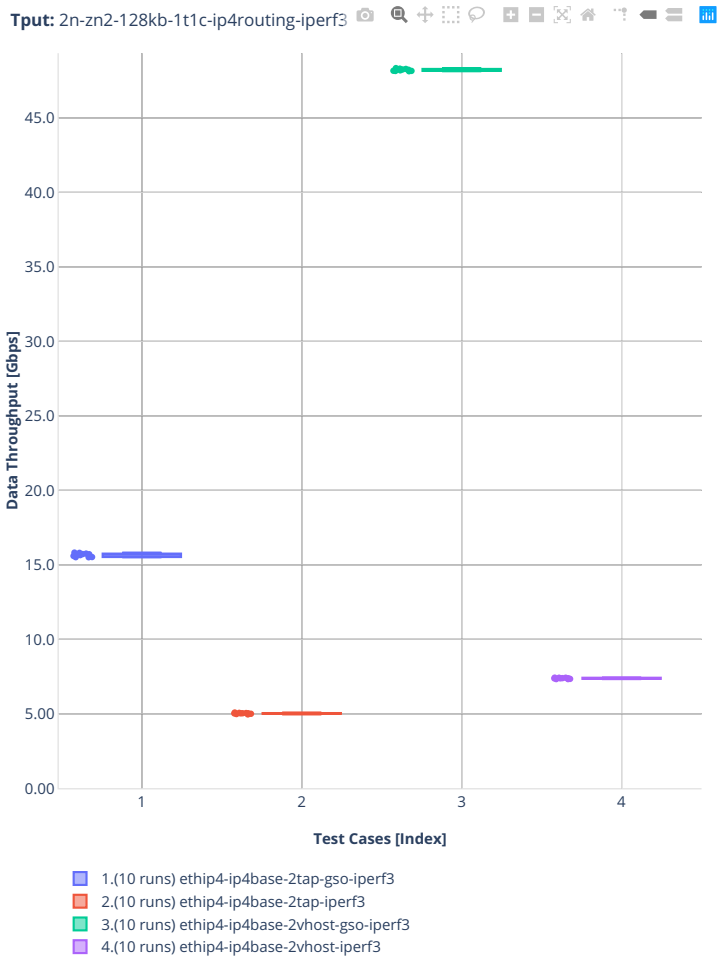
Tput: 2n-zn2-128kb-1t1c-ip4routing-iperf3
Data Throughput [Gbps]
45.0
40.0
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
1
2
3
4
Test Cases [Index]
1.(10 runs) ethip4-ip4base-2tap-gso-iperf3
2.(10 runs) ethip4-ip4base-2tap-iperf3
3.(10 runs) ethip4-ip4base-2vhost-gso-iperf3
4.(10 runs) ethip4-ip4base-2vhost-iperf3

2.11 Comparisons

2.11.1 Current vs Previous Release

Relative comparison of VPP packet throughput (NDR, PDR and MRR) between VPP-22.06 release and VPP-22.02 release (measured for CSIT-2206 and CSIT-2202 respectively) is calculated from results of tests running on 2-node Intel Xeon Skylake (2n-skx), 3-node Intel Xeon Skylake (3n-skx), 2-node Intel Atom Denverton (2n-dnv), 3-node Intel Atom Denverton (3n-dnv), 3-node Arm TaiShan (3n-tsh) testbeds, 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 are stored in

- build logs from FD.io vpp performance job 2n-icx[176],
- build logs from FD.io vpp performance job 3n-icx[177],
- build logs from FD.io vpp performance job 2n-skx[178],
- build logs from FD.io vpp performance job 3n-skx[179],
- build logs from FD.io vpp performance job 2n-clx[180],
- build logs from FD.io vpp performance job 2n-zn2[181],
- build logs from FD.io vpp performance job 2n-dnv[182],
- build logs from FD.io vpp performance job 3n-dnv[183],
- build logs from FD.io vpp performance job 3n-tsh[184],
- build logs from FD.io vpp performance job 2n-tx2[185],
- build logs from FD.io vpp performance job 2n-aws[186],

with RF result files csit-vpp-perf-2206-*.zip archived here.

---

[176] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-icx
[177] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-icx
[178] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx
[179] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-skx
[180] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-clx
[181] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-zn2
[182] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-dnv
[183] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-dnv
[184] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-tsh
[185] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-tx2
[186] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-aws

2n-icx-xxv710

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c NDR comparison
- HTML 4t2c NDR comparison
- ASCII 2t1c NDR comparison
- ASCII 4t2c NDR comparison
- CSV 2t1c NDR comparison
- CSV 4t2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR comparison
- HTML 4t2c PDR comparison
- ASCII 2t1c PDR comparison
- ASCII 4t2c PDR comparison
- CSV 2t1c PDR comparison
- CSV 4t2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c MRR comparison
- HTML 4t2c MRR comparison
- HTML 8t4c MRR comparison
- ASCII 2t1c MRR comparison
- ASCII 4t2c MRR comparison
- ASCII 8t4c MRR comparison
- CSV 2t1c MRR comparison
- CSV 4t2c MRR comparison
- CSV 8t4c MRR comparison

Latency Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR50, direction1, average value comparison
- HTML 2t1c PDR90, direction1, average value comparison
- HTML 2t1c PDR90, direction1, max value comparison
- ASCII 2t1c PDR50, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, max value comparison
- CSV 2t1c PDR50, direction1, average value comparison
- CSV 2t1c PDR90, direction1, average value comparison
- CSV 2t1c PDR90, direction1, max value comparison

3n-icx-xxv710

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c NDR comparison
- HTML 4t2c NDR comparison
- ASCII 2t1c NDR comparison
- ASCII 4t2c NDR comparison
- CSV 2t1c NDR comparison
- CSV 4t2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR comparison
- HTML 4t2c PDR comparison
- ASCII 2t1c PDR comparison
- ASCII 4t2c PDR comparison
- CSV 2t1c PDR comparison
- CSV 4t2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c MRR comparison
- HTML 4t2c MRR comparison
- HTML 8t4c MRR comparison
- ASCII 2t1c MRR comparison
- ASCII 4t2c MRR comparison
- ASCII 8t4c MRR comparison
- CSV 2t1c MRR comparison
- CSV 4t2c MRR comparison
- CSV 8t4c MRR comparison

Latency Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR50, direction1, average value comparison
- HTML 2t1c PDR90, direction1, average value comparison
- HTML 2t1c PDR90, direction1, max value comparison
- ASCII 2t1c PDR50, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, max value comparison
- CSV 2t1c PDR50, direction1, average value comparison
- CSV 2t1c PDR90, direction1, average value comparison
- CSV 2t1c PDR90, direction1, max value comparison

2n-skx-xxv710

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c NDR comparison
- HTML 4t2c NDR comparison
- ASCII 2t1c NDR comparison
- ASCII 4t2c NDR comparison
- CSV 2t1c NDR comparison
- CSV 4t2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR comparison
- HTML 4t2c PDR comparison
- ASCII 2t1c PDR comparison
- ASCII 4t2c PDR comparison
- CSV 2t1c PDR comparison
- CSV 4t2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c MRR comparison
- HTML 4t2c MRR comparison
- HTML 8t4c MRR comparison
- ASCII 2t1c MRR comparison
- ASCII 4t2c MRR comparison
- ASCII 8t4c MRR comparison
- CSV 2t1c MRR comparison
- CSV 4t2c MRR comparison
- CSV 8t4c MRR comparison

Latency Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR50, direction1, average value comparison
- HTML 2t1c PDR90, direction1, average value comparison
- HTML 2t1c PDR90, direction1, max value comparison
- ASCII 2t1c PDR50, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, max value comparison
- CSV 2t1c PDR50, direction1, average value comparison
- CSV 2t1c PDR90, direction1, average value comparison
- CSV 2t1c PDR90, direction1, max value comparison

3n-skx-xxv710

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c NDR comparison
- HTML 4t2c NDR comparison
- ASCII 2t1c NDR comparison
- ASCII 4t2c NDR comparison
- CSV 2t1c NDR comparison
- CSV 4t2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR comparison
- HTML 4t2c PDR comparison
- ASCII 2t1c PDR comparison
- ASCII 4t2c PDR comparison
- CSV 2t1c PDR comparison
- CSV 4t2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c MRR comparison
- HTML 4t2c MRR comparison
- HTML 8t4c MRR comparison
- ASCII 2t1c MRR comparison
- ASCII 4t2c MRR comparison
- ASCII 8t4c MRR comparison
- CSV 2t1c MRR comparison
- CSV 4t2c MRR comparison
- CSV 8t4c MRR comparison

Latency Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR50, direction1, average value comparison
- HTML 2t1c PDR90, direction1, average value comparison
- HTML 2t1c PDR90, direction1, max value comparison
- ASCII 2t1c PDR50, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, max value comparison
- CSV 2t1c PDR50, direction1, average value comparison
- CSV 2t1c PDR90, direction1, average value comparison
- CSV 2t1c PDR90, direction1, max value comparison

Hoststack

Comparison tables in HTML, ASCII and CSV formats:

- HTML Hoststack comparison
- ASCII Hoststack comparison
- CSV Hoststack comparison

2n-clx-xxv710

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c NDR comparison
- HTML 4t2c NDR comparison
- ASCII 2t1c NDR comparison
- ASCII 4t2c NDR comparison
- CSV 2t1c NDR comparison
- CSV 4t2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR comparison
- HTML 4t2c PDR comparison
- ASCII 2t1c PDR comparison
- ASCII 4t2c PDR comparison
- CSV 2t1c PDR comparison
- CSV 4t2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c MRR comparison
- HTML 4t2c MRR comparison
- HTML 8t4c MRR comparison
- ASCII 2t1c MRR comparison
- ASCII 4t2c MRR comparison
- ASCII 8t4c MRR comparison
- CSV 2t1c MRR comparison
- CSV 4t2c MRR comparison
- CSV 8t4c MRR comparison

Latency Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR50, direction1, average value comparison
- HTML 2t1c PDR90, direction1, average value comparison
- HTML 2t1c PDR90, direction1, max value comparison
- ASCII 2t1c PDR50, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, max value comparison
- CSV 2t1c PDR50, direction1, average value comparison
- CSV 2t1c PDR90, direction1, average value comparison
- CSV 2t1c PDR90, direction1, max value comparison

VSAP

Comparison tables in HTML, ASCII and CSV formats:

- HTML VSAP comparison
- ASCII VSAP comparison
- CSV VSAP comparison

2n-clx-cx556a

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c NDR comparison
- HTML 4t2c NDR comparison
- ASCII 2t1c NDR comparison
- ASCII 4t2c NDR comparison

- CSV 2t1c NDR comparison
- CSV 4t2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR comparison
- HTML 4t2c PDR comparison
- ASCII 2t1c PDR comparison
- ASCII 4t2c PDR comparison
- CSV 2t1c PDR comparison
- CSV 4t2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c MRR comparison
- HTML 4t2c MRR comparison
- HTML 8t4c MRR comparison
- ASCII 2t1c MRR comparison
- ASCII 4t2c MRR comparison
- ASCII 8t4c MRR comparison
- CSV 2t1c MRR comparison
- CSV 4t2c MRR comparison
- CSV 8t4c MRR comparison

Latency Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR50, direction1, average value comparison
- HTML 2t1c PDR90, direction1, average value comparison
- HTML 2t1c PDR90, direction1, max value comparison
- ASCII 2t1c PDR50, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, max value comparison
- CSV 2t1c PDR50, direction1, average value comparison
- CSV 2t1c PDR90, direction1, average value comparison
- CSV 2t1c PDR90, direction1, max value comparison

2n-zn2-xxv710

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c NDR comparison
- HTML 4t2c NDR comparison
- ASCII 2t1c NDR comparison
- ASCII 4t2c NDR comparison
- CSV 2t1c NDR comparison
- CSV 4t2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR comparison
- HTML 4t2c PDR comparison
- ASCII 2t1c PDR comparison
- ASCII 4t2c PDR comparison
- CSV 2t1c PDR comparison
- CSV 4t2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c MRR comparison
- HTML 4t2c MRR comparison
- HTML 8t4c MRR comparison
- ASCII 2t1c MRR comparison
- ASCII 4t2c MRR comparison
- ASCII 8t4c MRR comparison
- CSV 2t1c MRR comparison
- CSV 4t2c MRR comparison
- CSV 8t4c MRR comparison

Latency Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR50, direction1, average value comparison
- HTML 2t1c PDR90, direction1, average value comparison
- HTML 2t1c PDR90, direction1, max value comparison
- ASCII 2t1c PDR50, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, average value comparison
- ASCII 2t1c PDR90, direction1, max value comparison
- CSV 2t1c PDR50, direction1, average value comparison
- CSV 2t1c PDR90, direction1, average value comparison
- CSV 2t1c PDR90, direction1, max value comparison

2n-dnv-x553

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c NDR comparison
- HTML 2t2c NDR comparison
- ASCII 1t1c NDR comparison
- ASCII 2t2c NDR comparison
- CSV 1t1c NDR comparison
- CSV 2t2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c PDR comparison
- HTML 2t2c PDR comparison
- ASCII 1t1c PDR comparison
- ASCII 2t2c PDR comparison
- CSV 1t1c PDR comparison
- CSV 2t2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c MRR comparison
- HTML 2t2c MRR comparison
- HTML 4t4c MRR comparison
- ASCII 1t1c MRR comparison
- ASCII 2t2c MRR comparison
- ASCII 4t4c MRR comparison
- CSV 1t1c MRR comparison
- CSV 2t2c MRR comparison
- CSV 4t4c MRR comparison

3n-dnv-x553

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c NDR comparison
- HTML 2t2c NDR comparison
- ASCII 1t1c NDR comparison
- ASCII 2t2c NDR comparison
- CSV 1t1c NDR comparison
- CSV 2t2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c PDR comparison
- HTML 2t2c PDR comparison
- ASCII 1t1c PDR comparison
- ASCII 2t2c PDR comparison
- CSV 1t1c PDR comparison
- CSV 2t2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c MRR comparison
- HTML 2t2c MRR comparison
- HTML 4t4c MRR comparison
- ASCII 1t1c MRR comparison
- ASCII 2t2c MRR comparison
- ASCII 4t4c MRR comparison
- CSV 1t1c MRR comparison
- CSV 2t2c MRR comparison
- CSV 4t4c MRR comparison

3n-tsh-x520

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c NDR comparison
- HTML 2t2c NDR comparison
- ASCII 1t1c NDR comparison
- ASCII 2t2c NDR comparison
- CSV 1t1c NDR comparison
- CSV 2t2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c PDR comparison
- HTML 2t2c PDR comparison
- ASCII 1t1c PDR comparison
- ASCII 2t2c PDR comparison
- CSV 1t1c PDR comparison
- CSV 2t2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c MRR comparison
- HTML 2t2c MRR comparison
- HTML 4t4c MRR comparison
- ASCII 1t1c MRR comparison
- ASCII 2t2c MRR comparison
- ASCII 4t4c MRR comparison
- CSV 1t1c MRR comparison
- CSV 2t2c MRR comparison
- CSV 4t4c MRR comparison

Latency Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c PDR50, direction1, average value comparison
- HTML 1t1c PDR90, direction1, average value comparison
- HTML 1t1c PDR90, direction1, max value comparison
- ASCII 1t1c PDR50, direction1, average value comparison
- ASCII 1t1c PDR90, direction1, average value comparison
- ASCII 1t1c PDR90, direction1, max value comparison
- CSV 1t1c PDR50, direction1, average value comparison
- CSV 1t1c PDR90, direction1, average value comparison
- CSV 1t1c PDR90, direction1, max value comparison

2n-tx2-xl710

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c NDR comparison
- HTML 2t2c NDR comparison
- ASCII 1t1c NDR comparison
- ASCII 2t2c NDR comparison
- CSV 1t1c NDR comparison
- CSV 2t2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c PDR comparison
- HTML 2t2c PDR comparison
- ASCII 1t1c PDR comparison
- ASCII 2t2c PDR comparison
- CSV 1t1c PDR comparison
- CSV 2t2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c MRR comparison
- HTML 2t2c MRR comparison
- HTML 4t4c MRR comparison
- ASCII 1t1c MRR comparison
- ASCII 2t2c MRR comparison
- ASCII 4t4c MRR comparison
- CSV 1t1c MRR comparison
- CSV 2t2c MRR comparison
- CSV 4t4c MRR comparison

Latency Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1t1c PDR50, direction1, average value comparison
- HTML 1t1c PDR90, direction1, average value comparison
- HTML 1t1c PDR90, direction1, max value comparison
- ASCII 1t1c PDR50, direction1, average value comparison
- ASCII 1t1c PDR90, direction1, average value comparison
- ASCII 1t1c PDR90, direction1, max value comparison
- CSV 1t1c PDR50, direction1, average value comparison
- CSV 1t1c PDR90, direction1, average value comparison
- CSV 1t1c PDR90, direction1, max value comparison

2n-aws-nitro50g

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c NDR comparison
- HTML 4t2c NDR comparison
- ASCII 2t1c NDR comparison
- ASCII 4t2c NDR comparison
- CSV 2t1c NDR comparison
- CSV 4t2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c PDR comparison
- HTML 4t2c PDR comparison
- ASCII 2t1c PDR comparison
- ASCII 4t2c PDR comparison
- CSV 2t1c PDR comparison
- CSV 4t2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c MRR comparison
- HTML 4t2c MRR comparison
- ASCII 2t1c MRR comparison
- ASCII 4t2c MRR comparison
- CSV 2t1c MRR comparison
- CSV 4t2c MRR comparison

2.11.2 2n-Icx vs 2n-Clx Testbeds

Relative comparison of VPP-22.06 release packet throughput (NDR, PDR and MRR) is calculated for the same tests executed on 2-Node Cascadelake (2n- clx) and 2-Node Icelake (2n-icx) physical testbed types, in 1-core, 2-core and 4-core configurations.

Note: Test results are stored in build logs from FD.io vpp performance job 2n-icx[187] and build logs from FD.io vpp performance job 2n-clx[188] with RF result files csit-vpp-perf-2206-*.zip archived here.

---

[187] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-icx
[188] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-clx

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c NDR comparison
- HTML 2c NDR comparison
- ASCII 1c NDR comparison
- ASCII 2c NDR comparison
- CSV 1c NDR comparison
- CSV 2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c PDR comparison
- HTML 2c PDR comparison
- ASCII 1c PDR comparison
- ASCII 2c PDR comparison
- CSV 1c PDR comparison
- CSV 2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c MRR comparison
- HTML 2c MRR comparison
- HTML 4c MRR comparison
- ASCII 1c MRR comparison
- ASCII 2c MRR comparison
- ASCII 4c MRR comparison
- CSV 1c MRR comparison
- CSV 2c MRR comparison
- CSV 4c MRR comparison

2.11.3 2n-Icx vs 2n-Skx Testbeds

Relative comparison of VPP-22.06 release packet throughput (NDR, PDR and MRR) is calculated for the same tests executed on 2-Node Cascadelake (2n- skx) and 2-Node Icelake (2n-icx) physical testbed types, in 1-core, 2-core and 4-core configurations.

Note: Test results are stored in build logs from FD.io vpp performance job 2n-icx[189] and build logs from FD.io vpp performance job 2n-skx[190] with RF result files csit-vpp-perf-2206-*.zip archived here.

---

[189] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-icx
[190] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c NDR comparison
- HTML 2c NDR comparison
- ASCII 1c NDR comparison
- ASCII 2c NDR comparison
- CSV 1c NDR comparison
- CSV 2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c PDR comparison
- HTML 2c PDR comparison
- ASCII 1c PDR comparison
- ASCII 2c PDR comparison
- CSV 1c PDR comparison
- CSV 2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c MRR comparison
- HTML 2c MRR comparison
- HTML 4c MRR comparison
- ASCII 1c MRR comparison
- ASCII 2c MRR comparison
- ASCII 4c MRR comparison
- CSV 1c MRR comparison
- CSV 2c MRR comparison
- CSV 4c MRR comparison

2.11.4 2n-Clx vs 2n-Skx Testbeds

Relative comparison of VPP-22.06 release packet throughput (NDR, PDR and MRR) is calculated for the same tests executed on 2-Node Skylake (2n- skx) and 2-Node Cascade Lake (2n-clx) physical testbed types, in 1-core, 2-core and 4-core configurations.

Note: Test results are stored in build logs from FD.io vpp performance job 2n-skx[191] and build logs from FD.io vpp performance job 2n-clx[192] with RF result files csit-vpp-perf-2206-*.zip archived here.

---

[191] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx
[192] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-clx

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c NDR comparison
- HTML 2c NDR comparison
- ASCII 1c NDR comparison
- ASCII 2c NDR comparison
- CSV 1c NDR comparison
- CSV 2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c PDR comparison
- HTML 2c PDR comparison
- ASCII 1c PDR comparison
- ASCII 2c PDR comparison
- CSV 1c PDR comparison
- CSV 2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c MRR comparison
- HTML 2c MRR comparison
- HTML 4c MRR comparison
- ASCII 1c MRR comparison
- ASCII 2c MRR comparison
- ASCII 4c MRR comparison
- CSV 1c MRR comparison
- CSV 2c MRR comparison
- CSV 4c MRR comparison

2.11.5 3n-Icx vs 3n-Skx Testbeds

Relative comparison of VPP-22.06 release packet throughput (NDR, PDR and MRR) is calculated for the same tests executed on 3-Node Skylake (3n- skx) and 3-Node Icelake (3n-icx) physical testbed types, in 1-core, 2-core and 4-core configurations.

Note: Test results are stored in build logs from FD.io vpp performance job 3n-skx[193] and build logs from FD.io vpp performance job 3n-icx[194] with RF result files csit-vpp-perf-2206-*.zip archived here.

[193] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-skx
[194] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-icx

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c NDR comparison
- HTML 2c NDR comparison
- ASCII 1c NDR comparison
- ASCII 2c NDR comparison
- CSV 1c NDR comparison
- CSV 2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c PDR comparison
- HTML 2c PDR comparison
- ASCII 1c PDR comparison
- ASCII 2c PDR comparison
- CSV 1c PDR comparison
- CSV 2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c MRR comparison
- HTML 2c MRR comparison
- HTML 4c MRR comparison
- ASCII 1c MRR comparison
- ASCII 2c MRR comparison
- ASCII 4c MRR comparison
- CSV 1c MRR comparison
- CSV 2c MRR comparison
- CSV 4c MRR comparison

2.11.6 3n-Skx vs 2n-Skx Testbeds

Relative comparison of VPP-22.06 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 are stored in build logs from FD.io vpp performance job 3n-skx[195] and build logs from FD.io vpp performance job 2n-skx[196] with RF result files csit-vpp-perf-2206-*.zip archived here.

[195] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-skx
[196] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c NDR comparison
- HTML 2c NDR comparison
- ASCII 1c NDR comparison
- ASCII 2c NDR comparison
- CSV 1c NDR comparison
- CSV 2c NDR comparison

PDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c PDR comparison
- HTML 2c PDR comparison
- ASCII 1c PDR comparison
- ASCII 2c PDR comparison
- CSV 1c PDR comparison
- CSV 2c PDR comparison

MRR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 1c MRR comparison
- HTML 2c MRR comparison
- HTML 4c MRR comparison
- ASCII 1c MRR comparison
- ASCII 2c MRR comparison
- ASCII 4c MRR comparison
- CSV 1c MRR comparison
- CSV 2c MRR comparison
- CSV 4c MRR comparison

2.11.7 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 3n-skx, 2n-skx testbeds.

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 are stored in

- build logs from FD.io vpp performance job 3n-skx[197],
- build logs from FD.io vpp performance job 2n-skx[198]

with RF result files csit-vpp-perf-2206-*.zip archived here.

3n-skx

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c NDR Intel-x710 and Intel-xxv710 comparison
- HTML 4t2c NDR Intel-x710 and Intel-xxv710 comparison
- 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 HTML, ASCII and CSV formats:

- HTML 2t1c PDR Intel-x710 and Intel-xxv710 comparison
- HTML 4t2c PDR Intel-x710 and Intel-xxv710 comparison
- 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 HTML, ASCII and CSV formats:

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

[197] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-3n-skx
[198] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx

2n-skx

NDR Comparison

Comparison tables in HTML, ASCII and CSV formats:

- HTML 2t1c NDR Intel-x710 and Intel-xxv710 comparison
- HTML 4t2c NDR Intel-x710 and Intel-xxv710 comparison
- 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 HTML, ASCII and CSV formats:

- HTML 2t1c PDR Intel-x710 and Intel-xxv710 comparison
- HTML 4t2c PDR Intel-x710 and Intel-xxv710 comparison
- 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 HTML, ASCII and CSV formats:

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

2.11.8 Soak Tests vs NDR Tests

Relative comparison of VPP-22.06 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 are stored in build logs from FD.io vpp performance job 2n-icx[199], build logs from FD.io vpp performance job 2n-skx[200], build logs from FD.io vpp performance job 2n-clx[201] with RF result files csit-vpp-perf-2206-*.zip archived here.

Comparison tables in ASCII and CSV formats:

2n-icx

- ASCII Soak vs NDR comparison
- CSV Soak vs NDR comparison

2n-skx

- ASCII Soak vs NDR comparison
- CSV Soak vs NDR comparison

2n-clx

- ASCII Soak vs NDR comparison
- CSV Soak vs NDR comparison

2.12 Throughput Trending

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

1. Performance Dashboard[202]: per VPP test case throughput trend, trend compliance and summary of detected anomalies.
2. Trending Methodology[203]: throughput test metrics, trend calculations and anomaly classification (progression, regression).
3. VPP Trendline Graphs[204]: per VPP build MRR throughput measurements against the trendline with anomaly highlights and associated CSIT test jobs.

[199] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-icx
[200] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-skx
[201] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-vpp-perf-report-iterative-2206-2n-clx
[202] https://s3-docs.fd.io/csit/master/trending/introduction/dashboard.html
[203] https://s3-docs.fd.io/csit/master/trending/methodology/index.html
[204] https://s3-docs.fd.io/csit/master/trending/index.html

2.13 Test Environment

2.13.1 Environment Versioning

CSIT test environment versioning has been introduced to track modifications of the test environment.

Any benchmark anomalies (progressions, regressions) between releases of a DUT application (e.g. VPP, DPDK), are determined by testing it in the same test environment, to avoid test environment changes clouding the picture. To beter distinguish impact of test environment changes, we also execute tests without any SUT (just with TRex TG sending packets over a link looping back to TG).

A mirror approach is introduced to determine benchmarking anomalies due to the test environment change. This is achieved by testing the same DUT application version between releases of CSIT test system. This works under the assumption that the behaviour of the DUT is deterministic under the test conditions.

CSIT test environment versioning scheme ensures integrity of all the test system components, including their HW revisions, compiled SW code versions and SW source code, within a specific CSIT version. Components included in the CSIT environment versioning include:

- **HW** Server hardware firmware and BIOS (motherboard, processsor, NIC(s), accelerator card(s)), tracked in CSIT branch in `./docs/lab/<server_platform_name>_hw_bios_cfg.md`, e.g. Xeon Skylake servers[205].
- **Linux** Server Linux OS version and configuration, tracked in CSIT Reports in SUT Settings[206] and Pre-Test Server Calibration[207].
- **TRex** TRex Traffic Generator version, drivers and configuration tracked in TG Settings[208].
- **CSIT** CSIT framework code tracked in CSIT release branches.

Following is the list of CSIT versions to date:

- Ver. 1 associated with CSIT rls1908 branch (HW[209], Linux[210], TRex[211], CSIT[212]).
- Ver. 2 associated with CSIT rls2001 branch (HW[213], Linux[214], TRex[215], CSIT[216]).
- Ver. 4 associated with CSIT rls2005 branch (HW[217], Linux[218], TRex[219], CSIT[220]).
- Ver. 5 associated with CSIT rls2009 branch (HW[221], Linux[222], TRex[223], CSIT[224]).
  - The main change is TRex data-plane core resource adjustments: increase from 7 to 8 cores and pinning cores to interfaces[225] for better TRex performance with symmetric traffic profiles.

[205] https://git.fd.io/csit/tree/docs/lab/testbeds_sm_skx_hw_bios_cfg.md#n556
[206] https://s3-docs.fd.io/csit/master/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[207] https://s3-docs.fd.io/csit/master/report/vpp_performance_tests/test_environment.html#id21
[208] https://s3-docs.fd.io/csit/master/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[209] https://git.fd.io/csit/tree/docs/lab?h=rls1908
[210] https://docs.fd.io/csit/rls1908/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[211] https://docs.fd.io/csit/rls1908/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[212] https://git.fd.io/csit/tree/?h=rls1908
[213] https://git.fd.io/csit/tree/docs/lab?h=rls2001
[214] https://docs.fd.io/csit/rls2001/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[215] https://docs.fd.io/csit/rls2001/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[216] https://git.fd.io/csit/tree/?h=rls2001
[217] https://git.fd.io/csit/tree/docs/lab?h=rls2005
[218] https://docs.fd.io/csit/rls2005/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[219] https://docs.fd.io/csit/rls2005/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[220] https://git.fd.io/csit/tree/?h=rls2005
[221] https://git.fd.io/csit/tree/docs/lab?h=rls2009
[222] https://docs.fd.io/csit/rls2009/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[223] https://docs.fd.io/csit/rls2009/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[224] https://git.fd.io/csit/tree/?h=rls2009
[225] https://gerrit.fd.io/r/c/csit/+/28184

- Ver. 6 associated with CSIT rls2101 branch (HW[226], Linux[227], TRex[228], CSIT[229]).
  - The main change is TRex version upgrade: increase from 2.82 to 2.86[230].
- Ver. 7 associated with CSIT rls2106 branch (HW[231], Linux[232], TRex[233], CSIT[234]).
  - TRex version upgrade: increase from 2.86 to 2.88[235].
  - Ubuntu upgrade: upgrade from 18.04 LTS to 20.04.2 LTS[236].
- Ver. 8 associated with CSIT rls2110 branch (HW[237], Linux[238], TRex[239], CSIT[240]).
  - Intel NIC 700/800 series firmware upgrade based on DPDK compatibility matrix.
- Ver. 9 associated with CSIT rls2202 branch (HW[241], Linux[242], TRex[243], CSIT[244]).
  - Intel NIC 700/800 series firmware upgrade based on DPDK compatibility matrix.
- Ver. 10 associated with CSIT rls2206 branch (HW[245], Linux[246], TRex[247], CSIT[248]).
  - Intel NIC 700/800 series firmware upgrade based on DPDK compatibility matrix.
  - Mellanox 556A series firmware upgrade based on DPDK compatibility matrix.
  - Intel IceLake all core turbo frequency turned off. Current base frequency is 2.6GHz.

To identify performance changes due to VPP code development between previous and current VPP release version, both have been tested in CSIT environment of latest version and compared against each other. All substantial progressions and regressions have been marked up with RCA analysis. See *Comparisons* (page 1429) and *Known Issues* (page 81).

2.13.2 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 Icelake-SP, Intel Xeon Skylake-SP, Intel Xeon Cascade Lake-SP, Arm, Intel Atom. More detailed description is provided in *Performance Physical Testbeds* (page 4). Tested logical topologies are described in *Logical Topologies* (page 73).

[226] https://git.fd.io/csit/tree/docs/lab?h=rls2101
[227] https://docs.fd.io/csit/rls2101/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[228] https://docs.fd.io/csit/rls2101/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[229] https://git.fd.io/csit/tree/?h=rls2101
[230] https://gerrit.fd.io/r/c/csit/+/29980
[231] https://git.fd.io/csit/tree/docs/lab?h=rls2106
[232] https://s3-docs.fd.io/csit/rls2106/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[233] https://s3-docs.fd.io/csit/rls2106/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[234] https://git.fd.io/csit/tree/?h=rls2106
[235] https://gerrit.fd.io/r/c/csit/+/31652
[236] https://gerrit.fd.io/r/c/csit/+/31290
[237] https://git.fd.io/csit/tree/docs/lab?h=rls2110
[238] https://s3-docs.fd.io/csit/rls2110/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[239] https://s3-docs.fd.io/csit/rls2110/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[240] https://git.fd.io/csit/tree/?h=rls2110
[241] https://git.fd.io/csit/tree/docs/lab?h=rls2202
[242] https://s3-docs.fd.io/csit/rls2202/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[243] https://s3-docs.fd.io/csit/rls2202/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[244] https://git.fd.io/csit/tree/?h=rls2202
[245] https://git.fd.io/csit/tree/docs/lab?h=rls2206
[246] https://s3-docs.fd.io/csit/rls2206/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[247] https://s3-docs.fd.io/csit/rls2206/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[248] https://git.fd.io/csit/tree/?h=rls2206

2.13.3 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 Cascade Lake[249], FD.io CSIT testbeds - Xeon Skylake, Arm, Atom[250].

2.13.4 SUT Settings - Linux

System provisioning is done by combination of PXE boot unattented install and Ansible[251] described in CSIT Testbed Setup[252].

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.

[249] https://git.fd.io/csit/tree/docs/lab/testbeds_sm_clx_hw_bios_cfg.md?h=rls2206
[250] https://git.fd.io/csit/tree/docs/lab/testbeds_sm_skx_hw_bios_cfg.md?h=rls2206
[251] https://www.ansible.com
[252] https://git.fd.io/csit/tree/fdio.infra.ansible?h=rls2206

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.

2.13.5 DUT Settings - VPP

VPP Version

VPP-22.06 release

VPP Compile Parameters

FD.io VPP compile job[253]

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*, *$$UIO_DRIVER*, and *$$NO_MULTI_SEG* parameter. List of plugins to enable is driven by test requirements. Default template is provided below:

```
ip
{
  heap-size 4G
}
statseg
{
  size 4G
  per-node-counters on
}
unix
{
  cli-listen /run/vpp/cli.sock
  log /tmp/vpe.log
  nodaemon
  full-coredump
}
socksvr {
  socket-name /run/vpp/api.sock
}
ip6
{
  heap-size 4G
  hash-buckets 2000000
}
heapsize 4G
plugins
{
  plugin default
```

(continues on next page)

[253] https://jenkins.fd.io/view/vpp/job/vpp-merge-2206-ubuntu2004-x86_64/

(continued from previous page)

```
  {
    disable
  }
  plugin <$$test_requirement>_plugin.so
  {
    enable
  }
}
cpu
{
  corelist-workers $$CORELIST_WORKERS
  main-core 1
}
buffers
{
  buffers-per-numa 215040
}

# Below: in case of dpdk based drivers (vfio-pci) only
dpdk
{
  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 20).

2.13.6 TG Settings - TRex

TG Version

TRex v2.97

DPDK Version

DPDK v21.02

TG Installation

T-Rex installation is managed via Ansible role.

TG Startup Configuration

```
$ sudo -E -S sh -c 'cat << EOF > /etc/trex_cfg.yaml
- version: 2
  c: 8
  limit_memory: 8192
  interfaces: ["${pci1}","${pci2}"]
  port_info:
      - dest_mac: [${dest_mac1}]
        src_mac: [${src_mac1}]
      - dest_mac: [${dest_mac2}]
        src_mac: [${src_mac2}]
  platform :
      master_thread_id: 0
      latency_thread_id: 9
      dual_if:
          - socket: 0
            threads: [1, 2, 3, 4, 5, 6, 7, 8]
EOF'
```

TG Startup Command (Stateless Mode)

```
$ sudo -E -S sh -c "cd '${trex_install_dir}/scripts/' && \
  nohup ./t-rex-64 -i --prefix $(hostname) --hdrh --no-scapy-server \
  --mbuf-factor 32 > /tmp/trex.log 2>&1 &" > /dev/null
```

Also, Python client is now starting traffic with:

```
core_mask=STLClient.CORE_MASK_PIN
```

TG Startup Command (Stateful Mode)

```
$ sudo -E -S sh -c "cd '${trex_install_dir}/scripts/' && \
  nohup ./t-rex-64 -i --prefix $(hostname) --astf --hdrh --no-scapy-server \
  --mbuf-factor 32 > /tmp/trex.log 2>&1 &" > /dev/null
```

TG API Driver

TRex driver[254]

2.13.7 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[255].
2. Memory bandwidth - measure bandwidth with Intel MLC tool[256].
3. Memory latency - measure memory latency with Intel MLC tool.

[254] https://git.fd.io/csit/tree/GPL/tools/trex/trex_stl_profile.py?h=rls2206
[255] https://git.fd.io/pma_tools/tree/jitter
[256] https://software.intel.com/en-us/articles/intelr-memory-latency-checker

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 testbeds.

Ice Lake

Following sections include sample calibration data measured on s71-t212-sut1 server running in one of the Intel Xeon Ice Lake testbeds as specified in `FD.io CSIT testbeds - Xeon Ice Lake`_.

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

Linux cmdline

```
$ cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-5.4.0-65-generic root=UUID=3250758a-9bb6-48c8-9c36-ecb6a269223f ro audit=0␣
↪default_hugepagesz=2M hugepagesz=1G hugepages=32 hugepagesz=2M hugepages=32768 hpet=disable intel_
↪idle.max_cstate=1 intel_iommu=on intel_pstate=disable iommu=pt isolcpus=1-31,33-63,65-95,97-127␣
↪mce=off nmi_watchdog=0 nohz_full=1-31,33-63,65-95,97-127 nosoftlockup numa_balancing=disable␣
↪processor.max_cstate=1 rcu_nocbs=1-31,33-63,65-95,97-127 tsc=reliable console=ttyS0,115200n8 quiet
```

Linux uname

```
$ uname -a
Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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 1.9
Iterations=20
The pragram will execute a dummy function 80000 times
Display is updated every 20000 displayUpdate intervals
Thread affinity will be set to core_id:7
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 calcualted 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
126082,133950,7868,126094,126082,133950,3829268480,2524167454,1
126082,134696,8614,126094,126082,134696,1778253824,2524273022,2
126082,136092,10010,126094,126082,136092,4022206464,2524203296,3
126082,135094,9012,126094,126082,136092,1971191808,2524274302,4
126082,136482,10400,126094,126082,136482,4215144448,2524318496,5
```

(continues on next page)

(continued from previous page)

```
126082,134990,8908,126094,126082,136482,2164129792,2524155038,6
126082,134710,8628,126092,126082,136482,113115136,2524215228,7
126082,135080,8998,126092,126082,136482,2357067776,2524168906,8
126082,134470,8388,126094,126082,136482,306053120,2524163312,9
126082,135246,9164,126092,126082,136482,2550005760,2524394986,10
126082,132662,6580,126094,126082,136482,498991104,2524163156,11
126082,132954,6872,126094,126082,136482,2742943744,2524154386,12
126082,135340,9258,126092,126082,136482,691929088,2524222386,13
126082,133036,6954,126094,126082,136482,2935881728,2524150132,14
126082,137776,11694,126094,126082,137776,884867072,2524239346,15
126082,137850,11768,126094,126082,137850,3128819712,2524342944,16
126082,133000,6918,126094,126082,137850,1077805056,2524160062,17
126082,133332,7250,126094,126082,137850,3321757696,2524158804,18
126082,133234,7152,126092,126082,137850,1270743040,2524174400,19
126082,152552,26470,126094,126082,152552,3514695680,2524857280,20
```

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[257].

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 x86_64
CPU is Intel(R) Xeon(R) Platinum 8358 CPU @ 2.60GHz

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)
  * Microarchitectural 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: YES
  * 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 explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): YES
  * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): YES
  * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): YES
```

(continues on next page)

[257] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
    * TSX_CTRL MSR indicates TSX RTM is disabled: YES
    * TSX_CTRL MSR indicates TSX CPUID bit is cleared: YES
  * CPU supports Transactional Synchronization Extensions (TSX): NO
  * CPU supports Software Guard Extensions (SGX): YES
  * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
  * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x6a stepping 0x6␣
↪ucode 0xd000280 cpuid 0x606a6)
  * CPU microcode is the latest known available version: NO (latest version is 0xd0002a0 dated 2021/
↪04/25 according to builtin firmwares DB v191+i20210217)
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): YES
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): YES
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __user␣
↪pointer sanitization)
> STATUS: UNKNOWN (/sys vulnerability interface use forced, but its not available!)

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: YES (Mitigation: Enhanced IBRS, IBPB: conditional, RSB␣
↪filling)
> STATUS: VULNERABLE (IBRS+IBPB or retpoline+IBPB is needed to mitigate the vulnerability)

CVE-2017-5754 aka Variant 3, Meltdown, rogue data cache load
* Mitigated according to the /sys interface: YES (Not affected)
* 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)
> 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: YES
> STATUS: NOT VULNERABLE (your CPU microcode mitigates the vulnerability)

CVE-2018-3620 aka Foreshadow-NG (OS), L1 terminal fault
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (Not affected)
```

(continues on next page)

(continued from previous page)

```
CVE-2018-3646 aka Foreshadow-NG (VMM), L1 terminal fault
* Information from the /sys interface: Not affected
> STATUS: NOT VULNERABLE (your kernel 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)
> 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)
> 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)
> 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)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (Not affected)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

> SUMMARY: CVE-2017-5753:?? CVE-2017-5715:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

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[258].

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

Linux cmdline

```
$ cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-5.4.0-65-generic root=UUID=55d44abd-94d6-4b26-9d93-5877a8658016 ro audit=0␣
↪hpet=disable intel_idle.max_cstate=1 intel_iommu=on intel_pstate=disable iommu=pt isolcpus=1-27,
↪29-55,57-83,85-111 mce=off nmi_watchdog=0 nohz_full=1-27,29-55,57-83,85-111 nosoftlockup numa_
↪balancing=disable processor.max_cstate=1 rcu_nocbs=1-27,29-55,57-83,85-111 tsc=reliable␣
↪console=ttyS0,115200n8 quiet
```

[258] https://git.fd.io/csit/tree/docs/lab/testbeds_sm_skx_hw_bios_cfg.md?h=rls2206

Linux uname

```
$ uname -a
Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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 pragram will execute a dummy function 80000 times
Display is updated every 20000 displayUpdate intervals
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 calcualted 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
    160022     171330      11308     160022     160022     171330    2538733568 3204142750         ␣
↪1
    160022     167294       7272     160026     160022     171330     328335360 3203873548         ␣
↪2
    160022     167560       7538     160026     160022     171330    2412904448 3203878736         ␣
↪3
    160022     169000       8978     160024     160022     171330     202506240 3203864588         ␣
↪4
    160022     166572       6550     160026     160022     171330    2287075328 3203866224         ␣
↪5
    160022     167460       7438     160026     160022     171330      76677120 3203854632         ␣
↪6
    160022     168134       8112     160024     160022     171330    2161246208 3203874674         ␣
↪7
    160022     169094       9072     160022     160022     171330    4245815296 3203878798         ␣
↪8
    160022     172460      12438     160024     160022     172460    2035417088 3204112010         ␣
↪9
    160022     167862       7840     160030     160022     172460    4119986176 3203856800         ␣
↪10
    160022     168398       8376     160024     160022     172460    1909587968 3203854192         ␣
↪11
    160022     167548       7526     160024     160022     172460    3994157056 3203847442         ␣
↪12
    160022     167562       7540     160026     160022     172460    1783758848 3203862936         ␣
↪13
    160022     167604       7582     160024     160022     172460    3868327936 3203859346         ␣
↪14
    160022     168262       8240     160024     160022     172460    1657929728 3203851120         ␣
↪15
    160022     169700       9678     160024     160022     172460    3742498816 3203877690         ␣
↪16
    160022     170476      10454     160026     160022     172460    1532100608 3204088480         ␣
↪17
```

(continues on next page)

(continued from previous page)

```
    160022     167798       7776     160024     160022     172460    3616669696 3203862072
↪18
    160022     166540       6518     160024     160022     172460    1406271488 3203836904
↪19
    160022     167516       7494     160024     160022     172460    3490840576 3203848120
↪20
```

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     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[259].

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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)
  * Microarchitectural Data Sampling
    * VERW instruction is available: YES (MD_CLEAR feature bit)
  * 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
```

(continues on next page)

[259] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
  * CPU explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): NO
  * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): NO
  * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): NO
  * CPU supports Transactional Synchronization Extensions (TSX): YES (RTM feature bit)
  * CPU supports Software Guard Extensions (SGX): NO
  * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
  * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x55 stepping 0x4␣
↪ucode 0x2000065 cpuid 0x50654)
  * CPU microcode is the latest known available version: NO (latest version is 0x2006b06 dated 2021/
↪03/08 according to builtin firmwares DB v191+i20210217)
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): YES
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): YES
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪YES
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): YES
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪YES
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): YES
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): YES
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __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
* Kernel has array_index_nospec (arm64): NO
> STATUS: NOT VULNERABLE (Mitigation: usercopy/swapgs barriers and __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:␣
↪conditional, IBRS_FW, STIBP: conditional, 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)
  * 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)
```

(continues on next page)

(continued from previous page)

```
* 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 (boltd fwupd irqbalance systemd-
↪journald systemd-logind systemd-networkd systemd-resolved systemd-timesyncd systemd-udevd)
> 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
> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2018-12126 aka Fallout, microarchitectural store buffer data sampling (MSBDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2018-12130 aka ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2018-12127 aka RIDL, microarchitectural load port data sampling (MLPDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
```

(continues on next page)

(continued from previous page)

```
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2019-11091 aka RIDL, microarchitectural data sampling uncacheable memory (MDSUM)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
> STATUS: NOT VULNERABLE (Mitigation: Clear CPU buffers; SMT vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (KVM: Mitigation: Split huge pages)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
* iTLB Multihit mitigation enabled and active: YES (KVM: Mitigation: Split huge pages)
> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: YES (found SRBDS implementation evidence in␣
↪kernel image. Your kernel is up to date for SRBDS mitigation)
* SRBDS mitigation control is enabled and active: 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

Cascade Lake

Following sections include sample calibration data measured on s32-t27-sut1 server running in one of the Intel Xeon Skylake testbeds as specified in FD.io CSIT testbeds - Xeon Cascade Lake[260].

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

Linux cmdline

```
$ cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-5.4.0-65-generic root=UUID=b1f0dc29-1d4f-4777-b37d-a5e26e233d55 ro audit=0␣
↪hpet=disable intel_idle.max_cstate=1 intel_iommu=on intel_pstate=disable iommu=pt isolcpus=1-27,
↪29-55,57-83,85-111 mce=off nmi_watchdog=0 nohz_full=1-27,29-55,57-83,85-111 nosoftlockup numa_
↪balancing=disable processor.max_cstate=1 rcu_nocbs=1-27,29-55,57-83,85-111 tsc=reliable␣
↪console=ttyS0,115200n8 quiet
```

[260] https://git.fd.io/csit/tree/docs/lab/testbeds_sm_clx_hw_bios_cfg.md?h=rls2206

Linux uname

```
$ uname -a
Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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 1.9
Iterations=30
The pragram will execute a dummy function 80000 times
Display is updated every 20000 displayUpdate intervals
Thread affinity will be set to core_id:7
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 calcualted 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
160022,167590,7568,160026,160022,167590,2057568256,3203711852,1
160022,170628,10606,160024,160022,170628,4079222784,3204010824,2
160022,169824,9802,160024,160022,170628,1805910016,3203812064,3
160022,168832,8810,160030,160022,170628,3827564544,3203792594,4
160022,168248,8226,160026,160022,170628,1554251776,3203765920,5
160022,167834,7812,160028,160022,170628,3575906304,3203761114,6
160022,167442,7420,160024,160022,170628,1302593536,3203769250,7
160022,169120,9098,160028,160022,170628,3324248064,3203853340,8
160022,170710,10688,160024,160022,170710,1050935296,3203985878,9
160022,167952,7930,160024,160022,170710,3072589824,3203733756,10
160022,168314,8292,160030,160022,170710,799277056,3203741152,11
160022,169672,9650,160024,160022,170710,2820931584,3203739910,12
160022,168684,8662,160024,160022,170710,547618816,3203727336,13
160022,168246,8224,160024,160022,170710,2569273344,3203739052,14
160022,168134,8112,160030,160022,170710,295960576,3203735874,15
160022,170230,10208,160024,160022,170710,2317615104,3203996356,16
160022,167190,7168,160024,160022,170710,44302336,3203713628,17
160022,167304,7282,160024,160022,170710,2065956864,3203717954,18
160022,167500,7478,160024,160022,170710,4087611392,3203706674,19
160022,167302,7280,160024,160022,170710,1814298624,3203726452,20
160022,167266,7244,160024,160022,170710,3835953152,3203702804,21
160022,167820,7798,160022,160022,170710,1562640384,3203719138,22
160022,168100,8078,160024,160022,170710,3584294912,3203716636,23
160022,170408,10386,160024,160022,170710,1310982144,3203946958,24
160022,167276,7254,160024,160022,170710,3332636672,3203706236,25
160022,167052,7030,160024,160022,170710,1059323904,3203696444,26
160022,170322,10300,160024,160022,170710,3080978432,3203747514,27
160022,167332,7310,160024,160022,170710,807665664,3203716210,28
160022,167426,7404,160026,160022,170710,2829320192,3203700630,29
160022,168840,8818,160024,160022,170710,556007424,3203727658,30
```

Memory Bandwidth

```
$ sudo /home/testuser/mlc --bandwidth_matrix
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --bandwidth_matrix

Using buffer size of 100.000MiB/thread for reads and an additional 100.000MiB/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        122097.7     51327.9
       1        51309.2      122005.5
```

```
$ sudo /home/testuser/mlc --peak_injection_bandwidth
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --peak_injection_bandwidth

Using buffer size of 100.000MiB/thread for reads and an additional 100.000MiB/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        :      243159.4
3:1 Reads-Writes :      219132.5
2:1 Reads-Writes :      216603.1
1:1 Reads-Writes :      203713.0
Stream-triad like:      193790.8
```

```
$ sudo /home/testuser/mlc --max_bandwidth
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --max_bandwidth

Using buffer size of 100.000MiB/thread for reads and an additional 100.000MiB/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        :      244114.27
3:1 Reads-Writes :      219441.97
2:1 Reads-Writes :      216603.72
1:1 Reads-Writes :      203679.09
Stream-triad like:      214902.80
```

Memory Latency

```
$ sudo /home/testuser/mlc --latency_matrix
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --latency_matrix

Using buffer size of 2000.000MiB
Measuring idle latencies (in ns)...
                Numa node
Numa node            0       1
       0          81.2   130.2
       1         130.2    81.1
```

```
$ sudo /home/testuser/mlc --idle_latency
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --idle_latency

Using buffer size of 2000.000MiB
Each iteration took 186.1 core clocks ( 80.9    ns)
```

```
$ sudo /home/testuser/mlc --loaded_latency
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --loaded_latency

Using buffer size of 100.000MiB/thread for reads and an additional 100.000MiB/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  233.86   243421.9
 00002  230.61   243544.1
 00008  232.56   243394.5
 00015  229.52   244076.6
 00050  225.82   244290.6
 00100  161.65   236744.8
 00200  100.63   133844.0
 00300   96.84    90548.2
 00400   95.71    68504.3
 00500   95.68    55139.0
 00700   88.77    39798.4
 01000   84.74    28200.1
 01300   83.08    21915.5
 01700   82.27    16969.3
 02500   81.66    11810.6
 03500   81.98     8662.9
 05000   81.48     6306.8
 09000   81.17     3857.8
 20000   80.19     2179.9
```

L1/L2/LLC Latency

```
$ sudo /home/testuser/mlc --c2c_latency
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --c2c_latency

Measuring cache-to-cache transfer latency (in ns)...
Local Socket L2->L2 HIT  latency        55.5
Local Socket L2->L2 HITM latency        55.6
Remote Socket L2->L2 HITM latency (data address homed in writer socket)
                      Reader Numa Node
Writer Numa Node     0       1
            0        -   115.6
            1    115.6       -
Remote Socket L2->L2 HITM latency (data address homed in reader socket)
                      Reader Numa Node
Writer Numa Node     0       1
            0        -   178.2
            1    178.4       -
```

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[261].

```
Spectre and Meltdown mitigation detection tool v0.44+

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)
   * Microarchitectural 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: YES
   * 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 explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): NO
   * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): NO
   * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): YES
     * TSX_CTRL MSR indicates TSX RTM is disabled: YES
```

(continues on next page)

[261] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
     * TSX_CTRL MSR indicates TSX CPUID bit is cleared: YES
   * CPU supports Transactional Synchronization Extensions (TSX): NO
   * CPU supports Software Guard Extensions (SGX): NO
   * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
   * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x55 stepping 0x7␣
↪ucode 0x500002c cpuid 0x50657)
   * CPU microcode is the latest known available version: NO (latest version is 0x5003102 dated␣
↪2021/03/08 according to builtin firmwares DB v191+i20210217)
 * CPU vulnerability to the speculative execution attack variants
   * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
   * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
   * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
   * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
   * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
   * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
   * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
   * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
   * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
   * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
   * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
   * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory␣
↪(MDSUM)): NO
   * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
   * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): YES
   * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka  Spectre Variant 1, bounds check bypass
 * Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __user␣
↪pointer sanitization)
 > STATUS: UNKNOWN (/sys vulnerability interface use forced, but it  s not available!)

CVE-2017-5715 aka  Spectre Variant 2, branch target injection
 * Mitigated according to the /sys interface: YES (Mitigation: Enhanced IBRS, IBPB: conditional,␣
↪RSB filling)
 > STATUS: VULNERABLE (IBRS+IBPB or retpoline+IBPB+RSB filling, is needed to mitigate the␣
↪vulnerability)

CVE-2017-5754 aka  Variant 3, Meltdown, rogue data cache load
 * Mitigated according to the /sys interface: YES (Not affected)
 * 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)
 > 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)
 > STATUS: NOT VULNERABLE (Not affected)
```

(continues on next page)

(continued from previous page)

```
CVE-2018-3646 aka  Foreshadow-NG (VMM), L1 terminal fault
 * Information from the /sys interface: Not affected
 > STATUS: NOT VULNERABLE (your kernel 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)
 > 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)
 > 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)
 > 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)
 > STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2019-11135 aka  ZombieLoad V2, TSX Asynchronous Abort (TAA)
 * Mitigated according to the /sys interface: YES (Mitigation: TSX disabled)
 > STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka  No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
 * Mitigated according to the /sys interface: YES (KVM: Mitigation: Split huge pages)
 > STATUS: NOT VULNERABLE (KVM: Mitigation: Split huge pages)

CVE-2020-0543 aka  Special Register Buffer Data Sampling (SRBDS)
 * Mitigated according to the /sys interface: YES (Not affected)
 > STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

 > SUMMARY: CVE-2017-5753:?? CVE-2017-5715:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 x86_64
CPU is Intel(R) Xeon(R) Gold 6252N CPU @ 2.30GHz

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)
  * Microarchitectural Data Sampling
    * VERW instruction is available: YES (MD_CLEAR feature bit)
  * Enhanced IBRS (IBRS_ALL)
```

(continues on next page)

(continued from previous page)

```
    * CPU indicates ARCH_CAPABILITIES MSR availability: YES
    * ARCH_CAPABILITIES MSR advertises IBRS_ALL capability: YES
  * 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 explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): NO
  * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): NO
  * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): YES
    * TSX_CTRL MSR indicates TSX RTM is disabled: YES
    * TSX_CTRL MSR indicates TSX CPUID bit is cleared: YES
  * CPU supports Transactional Synchronization Extensions (TSX): NO
  * CPU supports Software Guard Extensions (SGX): NO
  * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
  * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x55 stepping 0x7␣
↪ucode 0x500002c cpuid 0x50657)
  * CPU microcode is the latest known available version: NO (latest version is 0x5003102 dated 2021/
↪03/08 according to builtin firmwares DB v191+i20210217)
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): YES
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __user␣
↪pointer sanitization)
> STATUS: UNKNOWN (/sys vulnerability interface use forced, but its not available!)

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: YES (Mitigation: Enhanced IBRS, IBPB: conditional, RSB␣
↪filling)
> STATUS: VULNERABLE (IBRS+IBPB or retpoline+IBPB+RSB filling, is needed to mitigate the␣
↪vulnerability)

CVE-2017-5754 aka Variant 3, Meltdown, rogue data cache load
* Mitigated according to the /sys interface: YES (Not affected)
* 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)
```

(continues on next page)

(continued from previous page)

```
> 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)
> STATUS: NOT VULNERABLE (Not affected)

CVE-2018-3646 aka Foreshadow-NG (VMM), L1 terminal fault
* Information from the /sys interface: Not affected
> STATUS: NOT VULNERABLE (your kernel 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)
> 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)
> 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)
> 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)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Mitigation: TSX disabled)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (KVM: Mitigation: Split huge pages)
> STATUS: NOT VULNERABLE (KVM: Mitigation: Split huge pages)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

> SUMMARY: CVE-2017-5753:?? CVE-2017-5715:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

Denverton

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

A 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-5.4.0-65-generic root=UUID=26ca7b0f-904a-462d-a1c6-98c420c29515 ro audit=0␣
↪hpet=disable intel_idle.max_cstate=1 intel_iommu=on intel_pstate=disable iommu=pt isolcpus=1-5␣
↪mce=off nmi_watchdog=0 nohz_full=1-5 nosoftlockup numa_balancing=disable processor.max_cstate=1␣
↪rcu_nocbs=1-5 tsc=reliable console=tty0 console=ttyS0,115200n8
```

Linux uname

```
$ uname -a
Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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 pragram will execute a dummy function 80000 times
Display is updated every 20000 displayUpdate intervals
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 calcualted 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
    177530     196100      18570     177530     177530     196100    4156751872 3556820054         ␣
↪1
    177530     200784      23254     177530     177530     200784     321060864 3556897644         ␣
↪2
    177530     196346      18816     177530     177530     200784     780337152 3556918674         ␣
↪3
    177530     195962      18432     177530     177530     200784    1239613440 3556847928         ␣
↪4
    177530     195960      18430     177530     177530     200784    1698889728 3556860214         ␣
↪5
    177530     198824      21294     177530     177530     200784    2158166016 3556854934         ␣
↪6
    177530     198522      20992     177530     177530     200784    2617442304 3556862410         ␣
↪7
    177530     196362      18832     177530     177530     200784    3076718592 3556851636         ␣
↪8
    177530     199114      21584     177530     177530     200784    3535994880 3556870846         ␣
↪9
    177530     197194      19664     177530     177530     200784    3995271168 3556933584        ␣
↪10
    177530     198272      20742     177536     177530     200784     159580160 3556869044        ␣
↪11
```

(continues on next page)

(continued from previous page)

```
    177530     197586      20056     177530     177530     200784     618856448 3556903482
↪12
    177530     196072      18542     177530     177530     200784    1078132736 3556825540
↪13
    177530     196354      18824     177530     177530     200784    1537409024 3556881664
↪14
    177530     195906      18376     177530     177530     200784    1996685312 3556839924
↪15
    177530     199066      21536     177530     177530     200784    2455961600 3556860220
↪16
    177530     196968      19438     177530     177530     200784    2915237888 3556871890
↪17
    177530     195896      18366     177530     177530     200784    3374514176 3556855338
↪18
    177530     196020      18490     177530     177530     200784    3833790464 3556839820
↪19
    177530     196030      18500     177530     177530     200784    4293066752 3556889196
↪20
```

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
        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
```

(continues on next page)

(continued from previous page)

```
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
```

(continues on next page)

(continued from previous page)

```
 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[262].

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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)
  * Microarchitectural Data Sampling
    * VERW instruction is available: YES (MD_CLEAR feature bit)
  * 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 explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): NO
  * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): NO
  * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): NO
  * CPU supports Transactional Synchronization Extensions (TSX): YES (RTM feature bit)
```

(continues on next page)

[262] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
  * CPU supports Software Guard Extensions (SGX): NO
  * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
  * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x55 stepping 0x4␣
↪ucode 0x2000065 cpuid 0x50654)
  * CPU microcode is the latest known available version: NO (latest version is 0x2006b06 dated 2021/
↪03/08 according to builtin firmwares DB v191+i20210217)
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): YES
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): YES
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪YES
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): YES
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪YES
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): YES
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): YES
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __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
* Kernel has array_index_nospec (arm64): NO
> STATUS: NOT VULNERABLE (Mitigation: usercopy/swapgs barriers and __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:␣
↪conditional, IBRS_FW, STIBP: conditional, 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)
  * 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)
```

(continues on next page)

(continued from previous page)

```
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 (boltd fwupd irqbalance systemd-
↪journald systemd-logind systemd-networkd systemd-resolved systemd-timesyncd systemd-udevd)
> 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
> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2018-12126 aka Fallout, microarchitectural store buffer data sampling (MSBDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2018-12130 aka ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2018-12127 aka RIDL, microarchitectural load port data sampling (MLPDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)
```

(continues on next page)

(continued from previous page)

```
CVE-2019-11091 aka RIDL, microarchitectural data sampling uncacheable memory (MDSUM)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
> STATUS: NOT VULNERABLE (Mitigation: Clear CPU buffers; SMT vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (KVM: Mitigation: Split huge pages)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
* iTLB Multihit mitigation enabled and active: YES (KVM: Mitigation: Split huge pages)
> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: YES (found SRBDS implementation evidence in␣
↪kernel image. Your kernel is up to date for SRBDS mitigation)
* SRBDS mitigation control is enabled and active: 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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: NO
  * L1 data cache invalidation
    * FLUSH_CMD MSR is available: NO
    * CPU indicates L1D flush capability: NO
  * Microarchitectural Data Sampling
    * VERW instruction is available: NO
  * 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
```

(continues on next page)

(continued from previous page)

```
  * 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 explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): NO
  * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): NO
  * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): NO
  * CPU supports Transactional Synchronization Extensions (TSX): NO
  * CPU supports Software Guard Extensions (SGX): NO
  * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
  * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x5f stepping 0x1␣
↪ucode 0x20 cpuid 0x506f1)
  * CPU microcode is the latest known available version: NO (latest version is 0x34 dated 2020/10/
↪23 according to builtin firmwares DB v191+i20210217)
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): NO
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): NO
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): NO
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __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
* Kernel has array_index_nospec (arm64): NO
> STATUS: NOT VULNERABLE (Mitigation: usercopy/swapgs barriers and __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:␣
↪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)
```

(continues on next page)

(continued from previous page)

```
* Kernel supports Page Table Isolation (PTI): YES
  * PTI enabled and active: NO
  * Reduced performance impact of PTI: NO (PCID/INVPCID 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: NO
> STATUS: VULNERABLE (an up-to-date CPU microcode is needed to mitigate this vulnerability)

CVE-2018-3639 aka Variant 4, speculative store bypass
* Mitigated according to the /sys interface: NO (Vulnerable)
* Kernel supports disabling speculative store bypass (SSB): YES (found in /proc/self/status)
* SSB mitigation is enabled and active: NO
> STATUS: VULNERABLE (Your CPU doesnt support SSBD)

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)
* Kernel supports using MD_CLEAR mitigation: YES (found md_clear implementation evidence in kernel␣
↪image)
* 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 (found md_clear implementation evidence in kernel␣
↪image)
* 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 (found md_clear implementation evidence in kernel␣
↪image)
* Kernel mitigation is enabled and active: NO
* SMT is either mitigated or disabled: NO
```

(continues on next page)

(continued from previous page)

```
> 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 (found md_clear implementation evidence in kernel␣
↪image)
* 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-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Not affected)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (Not affected)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
* iTLB Multihit mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: YES (found SRBDS implementation evidence in␣
↪kernel image. Your kernel is up to date for SRBDS mitigation)
* SRBDS mitigation control is enabled and active: 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:KO CVE-2018-3639:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

Altra

Following sections include sample calibration data measured on s62-t34-sut1 server running in one of the Altra testbeds.

Linux cmdline

```
$ cat /proc/cmdline
BOOT_IMAGE=/vmlinuz-5.4.0-65-generic root=/dev/mapper/ubuntu--vg-ubuntu--lv ro audit=0 default_
↪hugepagesz=2M hugepagesz=1G hugepages=32 hugepagesz=2M hugepages=32768 iommu.passthrough=1␣
↪isolcpus=1-40,81-120 nmi_watchdog=0 nohz_full=1-40,81-120 nosoftlockup processor.max_cstate=1 rcu_
↪nocbs=1-40,81-120
```

Linux uname

```
$ uname -a
Linux s62-t34-sut1 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:27:25 UTC 2021 aarch64 aarch64␣
↪aarch64 GNU/Linux
```

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[263].

```
Spectre and Meltdown mitigation detection tool v0.45

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:27:25 UTC 2021 aarch64
CPU is ARM v8 model 0xd0c

Hardware check
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): NO
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): NO
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): NO
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): NO
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): NO
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: __user pointer sanitization)
> STATUS: UNKNOWN (/sys vulnerability interface use forced, but it's not available!)

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2017-5754 aka Variant 3, Meltdown, rogue data cache load
* Mitigated according to the /sys interface: YES (Not affected)
* Running as a Xen PV DomU: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-3640 aka Variant 3a, rogue system register read
* CPU microcode mitigates the vulnerability: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-3639 aka Variant 4, speculative store bypass
* Mitigated according to the /sys interface: YES (Mitigation: Speculative Store Bypass disabled via␣
↪prctl)
```

(continues on next page)

[263] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
> STATUS: NOT VULNERABLE (Mitigation: Speculative Store Bypass disabled via prctl)

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 affected)

CVE-2018-3620 aka Foreshadow-NG (OS), L1 terminal fault
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-3646 aka Foreshadow-NG (VMM), L1 terminal fault
* Information from the /sys interface: Not affected
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-12126 aka Fallout, microarchitectural store buffer data sampling (MSBDS)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-12130 aka ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-12127 aka RIDL, microarchitectural load port data sampling (MLPDS)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2019-11091 aka RIDL, microarchitectural data sampling uncacheable memory (MDSUM)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

> SUMMARY: CVE-2017-5753:?? 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

TaiShan

Following sections include sample calibration data measured on s17-t33-sut1 server running in one of the Cortex-A72 testbeds.

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

Linux cmdline

```
$ cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-5.4.0-65-generic root=UUID=7d1d0e77-4df0-43df-9619-a99db29ffb83 ro audit=0␣
↪intel_iommu=on isolcpus=1-27,29-55 nmi_watchdog=0 nohz_full=1-27,29-55 nosoftlockup processor.max_
↪cstate=1 rcu_nocbs=1-27,29-55 console=ttyAMA0,115200n8 quiet
```

Linux uname

```
$ uname -a
Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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 1.9
Iterations=30
The pragram will execute a dummy function 80000 times
Display is updated every 20000 displayUpdate intervals
Thread affinity will be set to core_id:7
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 calcualted 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
    160022     172254      12232     160042     160022     172254    1903230976 3204401362      ␣
↪1
    160022     173148      13126     160044     160022     173148     814809088 3204619316      ␣
↪2
    160022     169460       9438     160044     160022     173148    4021354496 3204391306      ␣
↪3
    160024     170270      10246     160044     160022     173148    2932932608 3204385830      ␣
↪4
    160022     169660       9638     160044     160022     173148    1844510720 3204387290      ␣
↪5
    160022     169410       9388     160040     160022     173148     756088832 3204375832      ␣
↪6
    160022     169012       8990     160042     160022     173148    3962634240 3204378924      ␣
↪7
    160022     169556       9534     160044     160022     173148    2874212352 3204374882      ␣
↪8
    160022     171684      11662     160042     160022     173148    1785790464 3204394596      ␣
↪9
    160022     171546      11524     160024     160022     173148     697368576 3204602774      ␣
↪10
    160022     169248       9226     160042     160022     173148    3903913984 3204401676      ␣
↪11
    160022     168458       8436     160042     160022     173148    2815492096 3204256350      ␣
↪12
```

(continues on next page)

(continued from previous page)

```
    160022     169574       9552     160044     160022     173148    1727070208 3204278116  ␣
↪13
    160022     169352       9330     160044     160022     173148     638648320 3204327234  ␣
↪14
    160022     169100       9078     160044     160022     173148    3845193728 3204388132  ␣
↪15
    160022     169338       9316     160042     160022     173148    2756771840 3204380724  ␣
↪16
    160022     170828      10806     160046     160022     173148    1668349952 3204430452  ␣
↪17
    160022     173162      13140     160026     160022     173162     579928064 3204611318  ␣
↪18
    160022     170482      10460     160042     160022     173162    3786473472 3204389896  ␣
↪19
    160024     170704      10680     160044     160022     173162    2698051584 3204422126  ␣
↪20
    160024     169302       9278     160044     160022     173162    1609629696 3204397334  ␣
↪21
    160022     171848      11826     160044     160022     173162     521207808 3204389818  ␣
↪22
    160022     169438       9416     160042     160022     173162    3727753216 3204395382  ␣
↪23
    160022     169312       9290     160042     160022     173162    2639331328 3204371202  ␣
↪24
    160022     171368      11346     160044     160022     173162    1550909440 3204440464  ␣
↪25
    160022     171998      11976     160042     160022     173162     462487552 3204609440  ␣
↪26
    160022     169740       9718     160046     160022     173162    3669032960 3204405826  ␣
↪27
    160022     169610       9588     160044     160022     173162    2580611072 3204390608  ␣
↪28
    160022     169254       9232     160044     160022     173162    1492189184 3204399760  ␣
↪29
    160022     169386       9364     160046     160022     173162     403767296 3204417762  ␣
↪30
```

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[264].

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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)
```

(continues on next page)

[264] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
  * 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)
  * Microarchitectural Data Sampling
    * VERW instruction is available: YES (MD_CLEAR feature bit)
  * 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 explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): NO
  * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): NO
  * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): NO
  * CPU supports Transactional Synchronization Extensions (TSX): YES (RTM feature bit)
  * CPU supports Software Guard Extensions (SGX): NO
  * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
  * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x55 stepping 0x4␣
↪ucode 0x2000065 cpuid 0x50654)
  * CPU microcode is the latest known available version: NO (latest version is 0x2006b06 dated 2021/
↪03/08 according to builtin firmwares DB v191+i20210217)
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): YES
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): YES
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪YES
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): YES
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪YES
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): YES
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): YES
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __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
* Kernel has array_index_nospec (arm64): NO
> STATUS: NOT VULNERABLE (Mitigation: usercopy/swapgs barriers and __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:␣
↪conditional, IBRS_FW, STIBP: conditional, RSB filling)
```

(continues on next page)

(continued from previous page)

```
* 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 (boltd fwupd irqbalance systemd-
↪journald systemd-logind systemd-networkd systemd-resolved systemd-timesyncd systemd-udevd)
> 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
> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2018-12126 aka Fallout, microarchitectural store buffer data sampling (MSBDS)
```

(continues on next page)

(continued from previous page)

```
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2018-12130 aka ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2018-12127 aka RIDL, microarchitectural load port data sampling (MLPDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2019-11091 aka RIDL, microarchitectural data sampling uncacheable memory (MDSUM)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
> STATUS: NOT VULNERABLE (Mitigation: Clear CPU buffers; SMT vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (KVM: Mitigation: Split huge pages)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
* iTLB Multihit mitigation enabled and active: YES (KVM: Mitigation: Split huge pages)
> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: YES (found SRBDS implementation evidence in␣
↪kernel image. Your kernel is up to date for SRBDS mitigation)
* SRBDS mitigation control is enabled and active: 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:27:25 UTC 2021 aarch64
CPU is ARM v8 model 0xd08
```

(continues on next page)

(continued from previous page)

```
Hardware check
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): NO
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): NO
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)): ␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)): ␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): NO
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): 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: NO
* Kernel has the Red Hat/Ubuntu patch: NO
* Kernel has mask_nospec64 (arm64): NO
* Kernel has array_index_nospec (arm64): NO
* Checking count of LFENCE instructions following a jump in kernel... NO (only 0 jump-then-lfence␣
↪instructions found, should be >= 30 (heuristic))
> STATUS: NOT VULNERABLE (Mitigation: __user pointer sanitization)

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: NO (Vulnerable)
* Mitigation 1
  * Kernel is compiled with IBRS support: YES
    * IBRS enabled and active: NO
  * Kernel is compiled with IBPB support: NO
    * IBPB enabled and active: NO
* Mitigation 2
  * Kernel has branch predictor hardening (arm): YES
  * Kernel compiled with retpoline option: NO
> STATUS: NOT VULNERABLE (Branch predictor hardening mitigates 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/INVPCID 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: NO
> STATUS: VULNERABLE (an up-to-date CPU microcode is needed to mitigate this vulnerability)

CVE-2018-3639 aka Variant 4, speculative store bypass
* Mitigated according to the /sys interface: NO (Vulnerable)
* Kernel supports disabling speculative store bypass (SSB): YES (found in /proc/self/status)
* SSB mitigation is enabled and active: NO
> STATUS: VULNERABLE (Your CPU doesnt support SSBD)
```

(continues on next page)

(continued from previous page)

```
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: NO
* 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: N/A (the kvm_intel module is not loaded)
* Mitigation 2
  * L1D flush is supported by kernel: NO
  * L1D flush enabled: NO
  * Hardware-backed L1D flush supported: NO (flush will be done in software, this is slower)
  * Hyper-Threading (SMT) is enabled: UNKNOWN
> 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)
* Kernel supports using MD_CLEAR mitigation: NO
* 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: NO
* 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: NO
* 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: NO
* 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-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Not affected)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (Not affected)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
```

(continues on next page)

(continued from previous page)

```
* iTLB Multihit mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: NO
* SRBDS mitigation control is enabled and active: 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:KO CVE-2018-3639:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK

Need more detailed information about mitigation options? Use --explain
A false sense of security is worse than no security at all, see --disclaimer
ok: [10.30.51.37] =>
spectre_meltdown_poll_results.stdout_lines:
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:27:25 UTC 2021 aarch64
CPU is ARM v8 model 0xd08

Hardware check
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): NO
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): NO
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): NO
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): 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: NO
* Kernel has the Red Hat/Ubuntu patch: NO
* Kernel has mask_nospec64 (arm64): NO
* Kernel has array_index_nospec (arm64): NO
* Checking count of LFENCE instructions following a jump in kernel... NO (only 0 jump-then-lfence␣
↪instructions found, should be >= 30 (heuristic))
> STATUS: NOT VULNERABLE (Mitigation: __user pointer sanitization)

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: NO (Vulnerable)
* Mitigation 1
  * Kernel is compiled with IBRS support: YES
    * IBRS enabled and active: NO
  * Kernel is compiled with IBPB support: NO
    * IBPB enabled and active: NO
```

(continues on next page)

(continued from previous page)

```
* Mitigation 2
  * Kernel has branch predictor hardening (arm): YES
  * Kernel compiled with retpoline option: NO
> STATUS: NOT VULNERABLE (Branch predictor hardening mitigates 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/INVPCID 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: NO
> STATUS: VULNERABLE (an up-to-date CPU microcode is needed to mitigate this vulnerability)

CVE-2018-3639 aka Variant 4, speculative store bypass
* Mitigated according to the /sys interface: NO (Vulnerable)
* Kernel supports disabling speculative store bypass (SSB): YES (found in /proc/self/status)
* SSB mitigation is enabled and active: NO
> STATUS: VULNERABLE (Your CPU doesnt support SSBD)

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: NO
* 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: N/A (the kvm_intel module is not loaded)
* Mitigation 2
  * L1D flush is supported by kernel: NO
  * L1D flush enabled: NO
  * Hardware-backed L1D flush supported: NO (flush will be done in software, this is slower)
  * Hyper-Threading (SMT) is enabled: UNKNOWN
> 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)
* Kernel supports using MD_CLEAR mitigation: NO
* 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: NO
* 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)
```

(continues on next page)

(continued from previous page)

```
* Mitigated according to the /sys interface: YES (Not affected)
* Kernel supports using MD_CLEAR mitigation: NO
* 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: NO
* 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-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Not affected)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (Not affected)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
* iTLB Multihit mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: NO
* SRBDS mitigation control is enabled and active: 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:KO CVE-2018-3639:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

ThunderX2

Following sections include sample calibration data measured on s27-t211-sut1 server running in one of the ThunderX2 testbeds.

Linux cmdline

```
$ cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-5.4.0-65-generic root=UUID=7d1d0e77-4df0-43df-9619-a99db29ffb83 ro audit=0␣
↪intel_iommu=on isolcpus=1-27,29-55 nmi_watchdog=0 nohz_full=1-27,29-55 nosoftlockup processor.max_
↪cstate=1 rcu_nocbs=1-27,29-55 console=ttyAMA0,115200n8 quiet
```

Linux uname

```
$ uname -a
Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 x86_64 x86_64 x86_64 GNU/Linux
```

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[265].

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:27:25 UTC 2021 aarch64
CPU is

Hardware check
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): NO
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): NO
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): NO
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): NO
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): 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: NO
* Kernel has the Red Hat/Ubuntu patch: NO
* Kernel has mask_nospec64 (arm64): NO
* Kernel has array_index_nospec (arm64): NO
* Checking count of LFENCE instructions following a jump in kernel... NO (only 0 jump-then-lfence␣
↪instructions found, should be >= 30 (heuristic))
> STATUS: NOT VULNERABLE (Mitigation: __user pointer sanitization)

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: NO (Vulnerable)
* Mitigation 1
  * Kernel is compiled with IBRS support: YES
    * IBRS enabled and active: NO
  * Kernel is compiled with IBPB support: NO
    * IBPB enabled and active: NO
* Mitigation 2
  * Kernel has branch predictor hardening (arm): YES
  * Kernel compiled with retpoline option: NO
```

(continues on next page)

[265] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
> STATUS: NOT VULNERABLE (Branch predictor hardening mitigates 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/INVPCID 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: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-3639 aka Variant 4, speculative store bypass
* Mitigated according to the /sys interface: NO (Vulnerable)
* Kernel supports disabling speculative store bypass (SSB): YES (found in /proc/self/status)
* SSB mitigation is enabled and active: NO
> STATUS: VULNERABLE (Your CPU doesnt support SSBD)

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: NO
* 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: N/A (the kvm_intel module is not loaded)
* Mitigation 2
  * L1D flush is supported by kernel: NO
  * L1D flush enabled: NO
  * Hardware-backed L1D flush supported: NO (flush will be done in software, this is slower)
  * Hyper-Threading (SMT) is enabled: UNKNOWN
> 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)
* Kernel supports using MD_CLEAR mitigation: NO
* 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: NO
* 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: NO
* Kernel mitigation is enabled and active: NO
```

(continues on next page)

(continued from previous page)

```
* 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: NO
* 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-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Not affected)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (Not affected)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
* iTLB Multihit mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: NO
* SRBDS mitigation control is enabled and active: 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:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

CHAPTER

THREE

DPDK PERFORMANCE

3.1 Overview

DPDK performance test results are reported for all physical testbed types present in FD.io labs and installed NIC models. For description of physical testbeds used for DPDK performance tests please refer to *Performance Physical Testbeds* (page 4).

3.1.1 Logical Topologies

CSIT DPDK performance tests are executed on physical testbeds described in *Performance Physical Testbeds* (page 4). Based on the packet path through server SUTs, one distinct logical topology type is used for DPDK DUT data plane testing: NIC-to-NIC switching topology.

NIC-to-NIC Switching

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

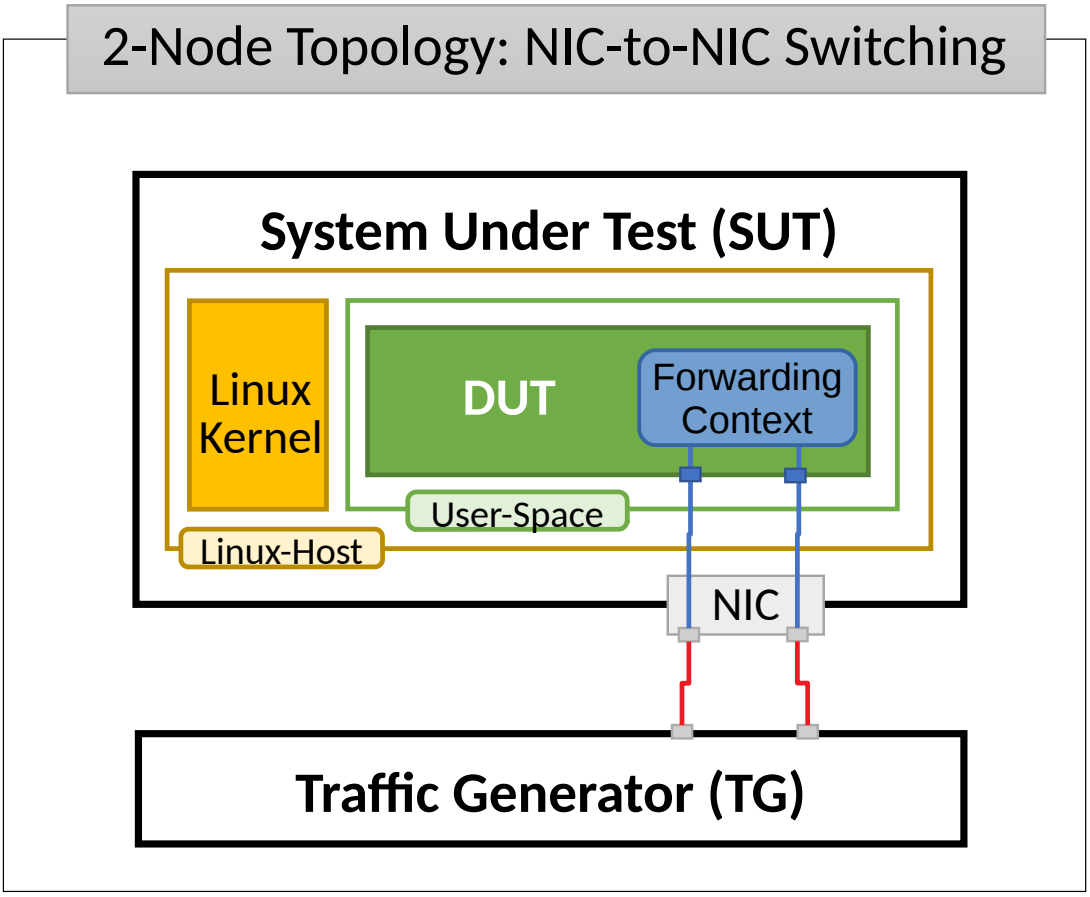

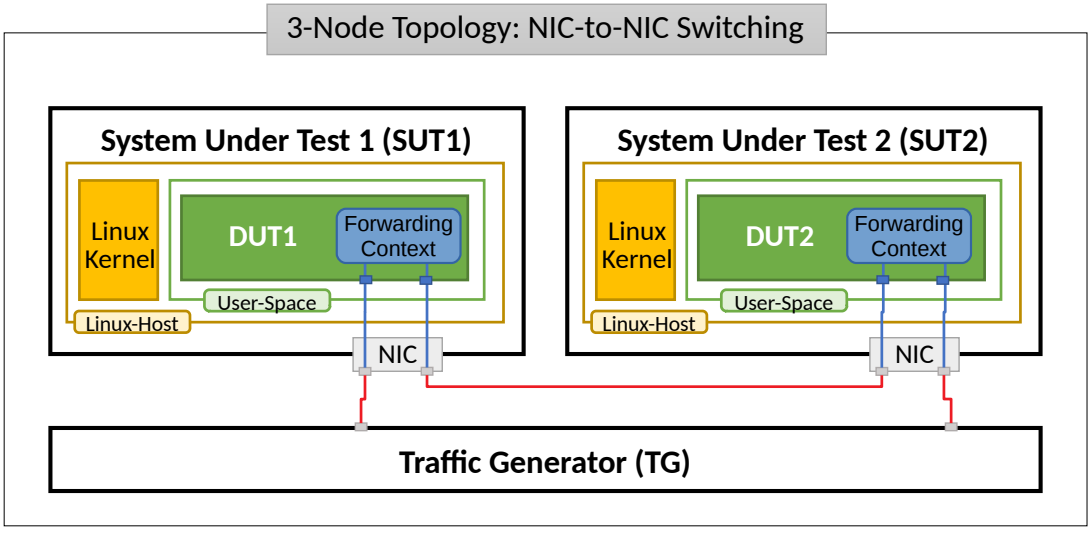

Server Systems Under Test (SUT) run DPDK Testpmd or L3fwd application in Linux user-mode as a Device Under Test (DUT). Server Traffic Generator (TG) runs T-Rex application. Physical connectivity between SUTs and TG is provided using different drivers and NIC models that need to be tested for performance (packet/bandwidth throughput and latency).

From SUT and DUT perspectives, all performance tests involve forwarding packets between two 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 limiations. 100GE NICs are not supported in PCIe Gen3 x8 slots.

Note that reported DPDK 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 DPDK packet thoughput 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.

3.1.2 Performance Tests Coverage

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

- Packet Throughput: measured in accordance with **RFC 2544**[266], 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): measured 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-2206 includes following DPDK Testpmd and L3fwd data plane functionality performance tested across a range of NIC drivers and NIC models:

Functionality	Description
L2IntLoop	L2 Interface Loop forwarding all Ethernet frames between two Interfaces.
IPv4 Routed Forwarding	Longest Prefix Match (LPM) L3 IPv4 forwarding of Ethernet frames between two Interfaces, with two /8 prefixes in lookup table.

3.2 Release Notes

3.2.1 Changes in CSIT-2206

1. TEST FRAMEWORK
   - **CSIT test environment** version has been updated to ver. 10, see *Environment Versioning* (page 1453).
2. DPDK PERFORMANCE TESTS
3. DPDK RELEASE VERSION CHANGE
   - CSIT-2206 tested DPDK-22.03, as used by VPP-22.06 release.

[266] https://tools.ietf.org/html/rfc2544.html

3.2.2 Known Issues

List of known issues in CSIT-2206 for DPDK performance tests:

#	JiraID	Issue Description
1	CSIT-1762[267]	TRex reports link DOWN in case of dpdk testpmd tests on FD.io CSIT Denverton systems (2n-dnv and 3n-dnv).

New

List of new issues in CSIT-2206 for DPDK performance tests:

#	JiraID	Issue Description
1	CSIT-1848[268]	2n-clx, 2n-skx, 3n-alt, 3n-skx: sporadic testpmd/l3fwd tests fail with no or low traffic.

[267] https://jira.fd.io/browse/CSIT-1762
[268] https://jira.fd.io/browse/CSIT-1848

3.3 Packet Throughput

Throughput graphs are generated by multiple executions of the same performance tests across physical testbeds hosted LF FD.io labs: 2n-icx, 3n-icx, 2n-skx, 3n-skx, 2n-clx, 2n-zn2, 3n-alt, 3n-tsh, 2n-tx2. Box-and-Whisker plots are used to display variations in measured throughput values, without making any assumptions of the underlying statistical distribution.

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 DPDK 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×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 are stored in build logs from FD.io dpdk performance job 2n-icx[269], build logs from FD.io dpdk performance job 3n-icx[270], build logs from FD.io dpdk performance job 2n-skx[271], build logs from FD.io dpdk performance job 3n-skx[272], build logs from FD.io dpdk performance job 2n-clx[273], build logs from FD.io dpdk performance job 2n-zn2[274], build logs from FD.io dpdk performance job 3n-alt[275], build logs from FD.io dpdk performance job 3n-tsh[276], build logs from FD.io dpdk performance job 2n-tx2[277] with RF result files csit-dpdk-perf-2206-*.zip archived here. Required per test case data set size is **10** and for DPDK tests this is the actual size, as all scheduled test executions completed successfully.

---

[269] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-icx
[270] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-icx
[271] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-skx
[272] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-skx
[273] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-clx
[274] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-zn2
[275] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-alt
[276] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-tsh
[277] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-tx2

3.3.1 2n-icx-xxv710

Following sections include summary graphs of Phy-to-Phy performance with packet 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[278].

[278] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

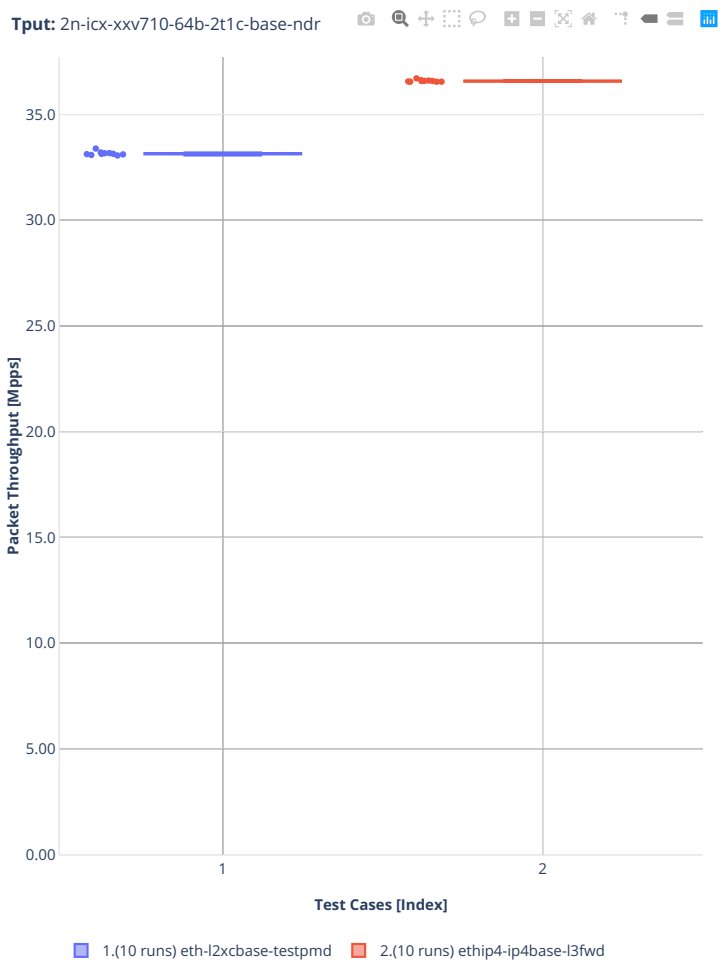

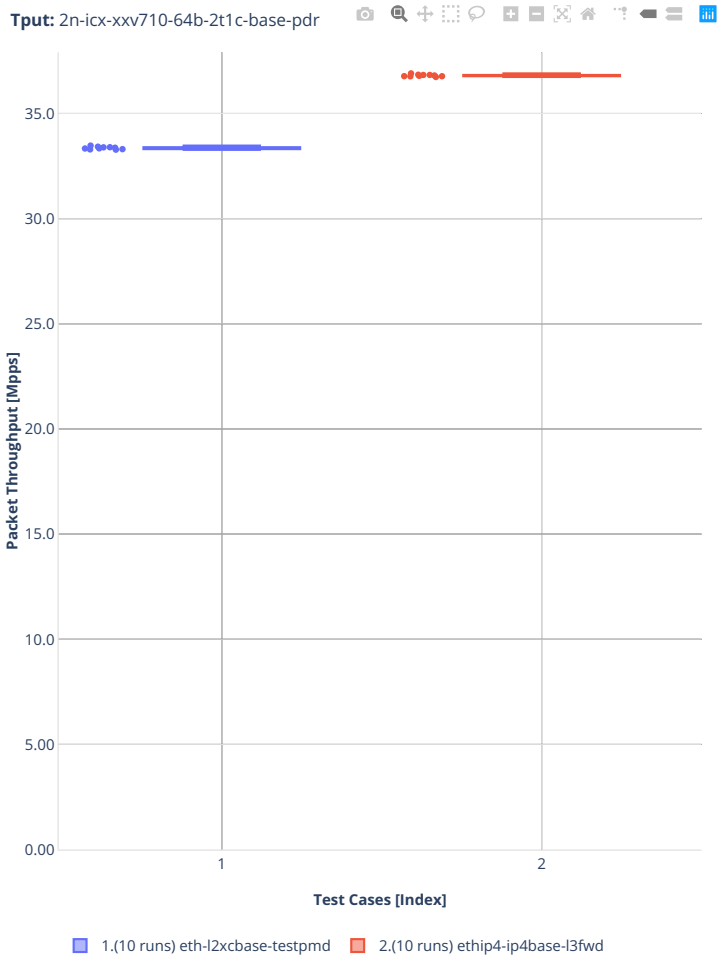
Tput: 2n-icx-xxv710-64b-2t1c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

64b-4t2c-base

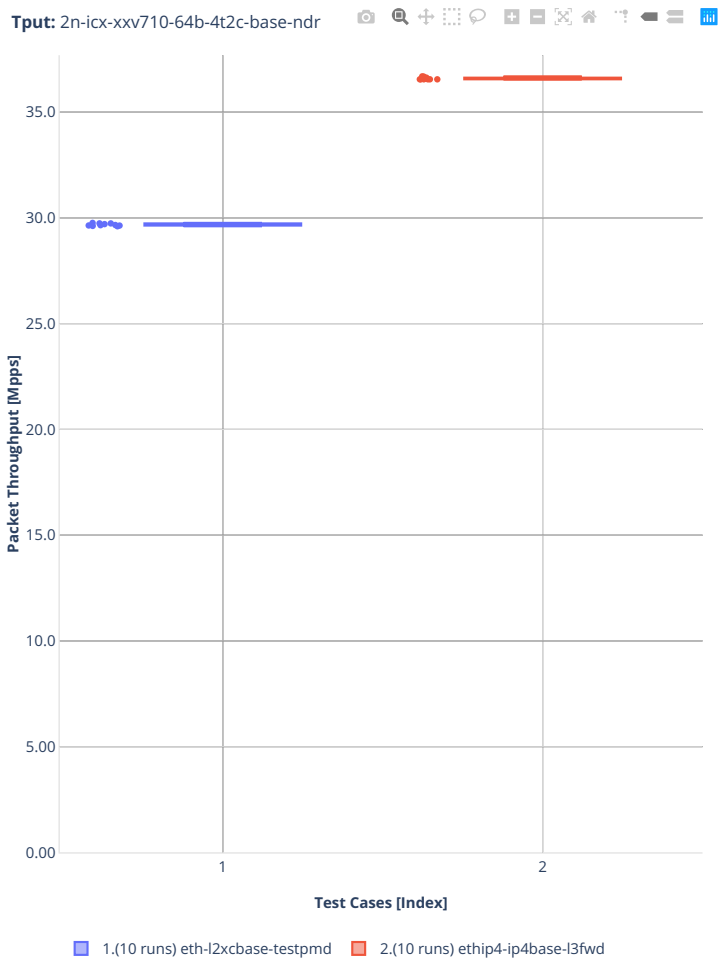

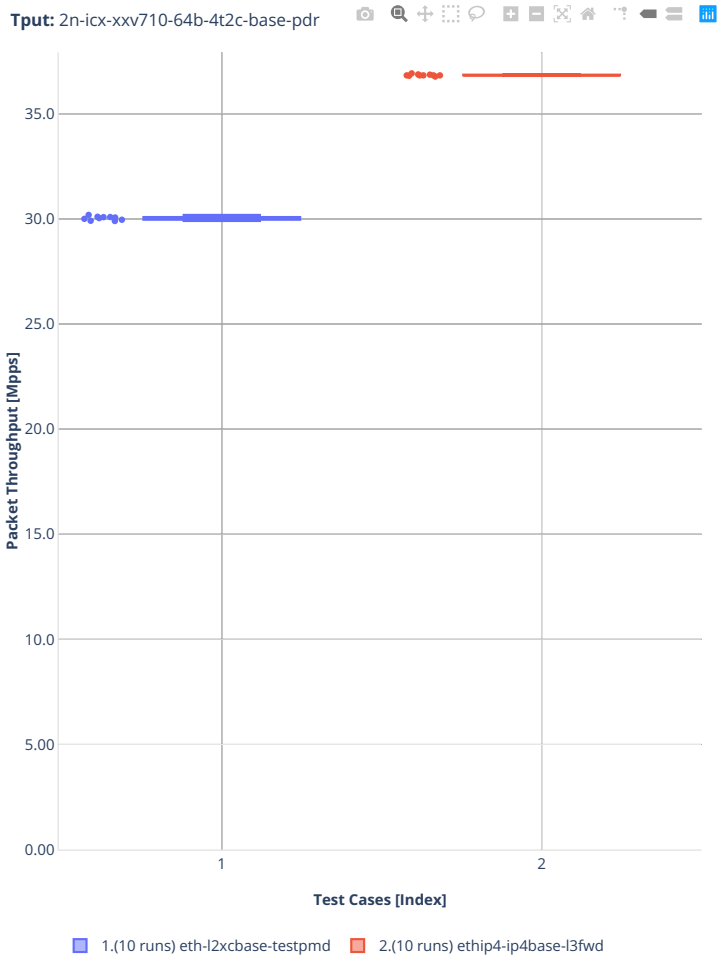
Tput: 2n-icx-xxv710-64b-4t2c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

3.3.2 3n-icx-xxv710

Following sections include summary graphs of Phy-to-Phy performance with packet 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[279].

[279] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

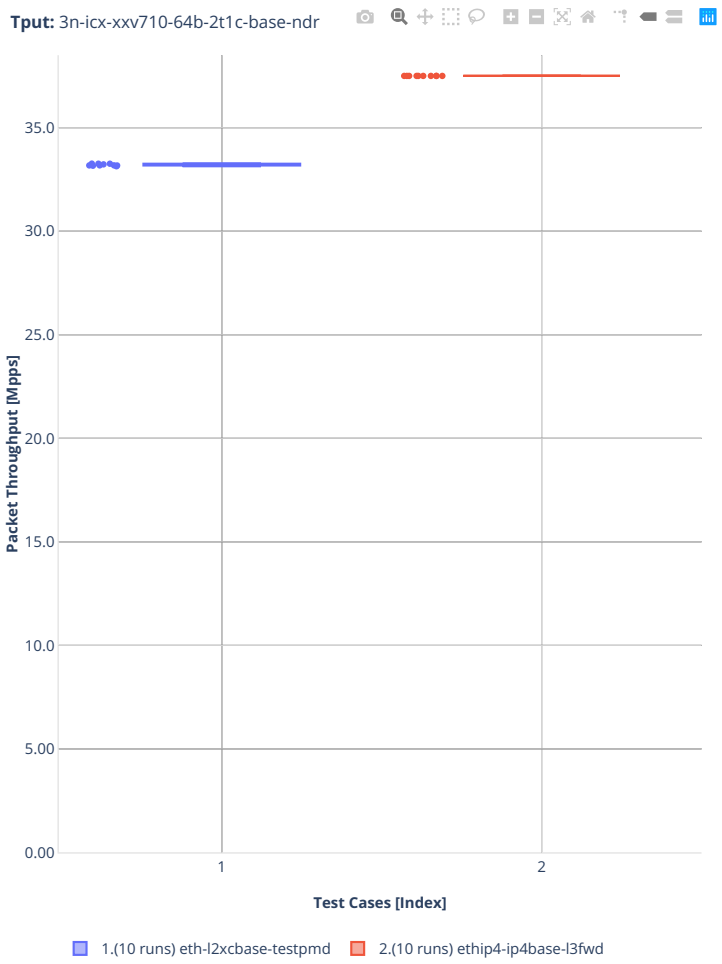

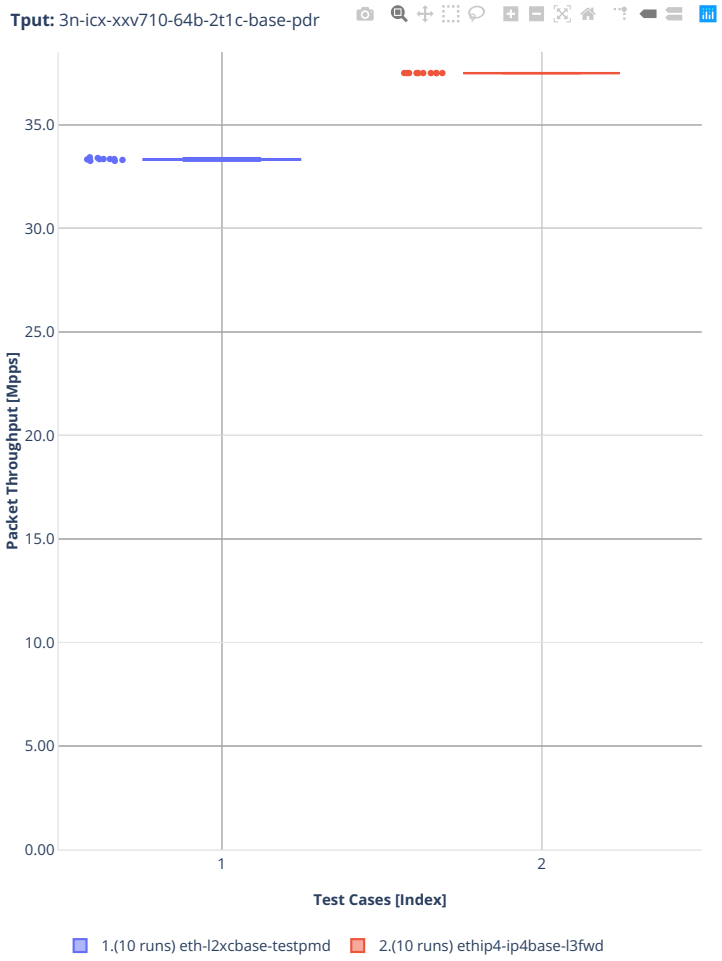
Tput: 3n-icx-xxv710-64b-2t1c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

64b-4t2c-base

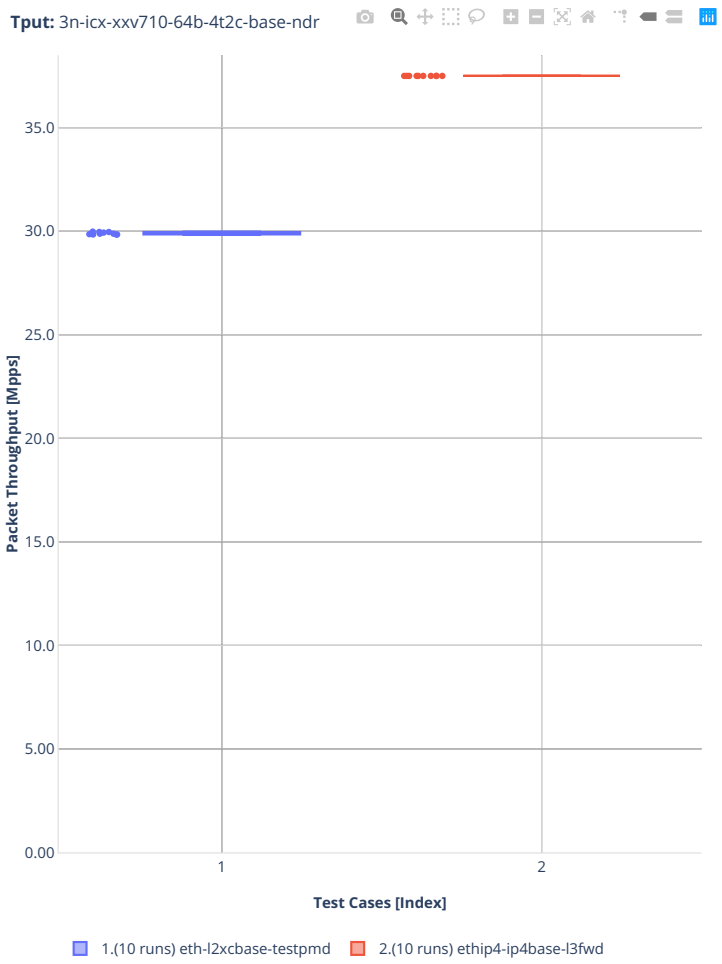

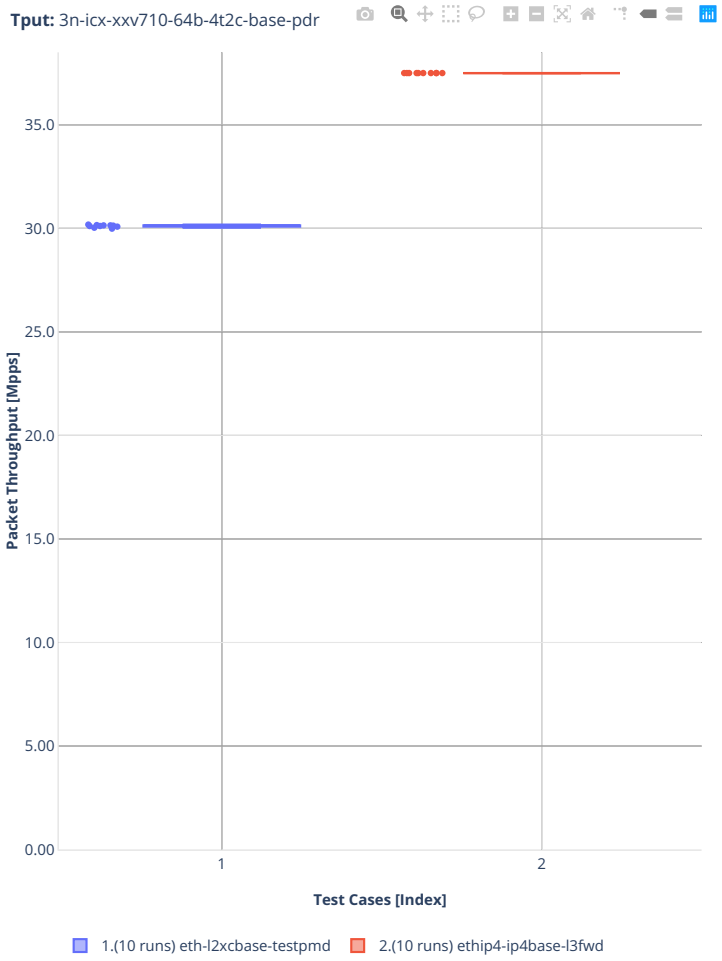
Tput: 3n-icx-xxv710-64b-4t2c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

3.3.3 2n-skx-xxv710

Following sections include summary graphs of Phy-to-Phy performance with packet 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[280].

[280] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

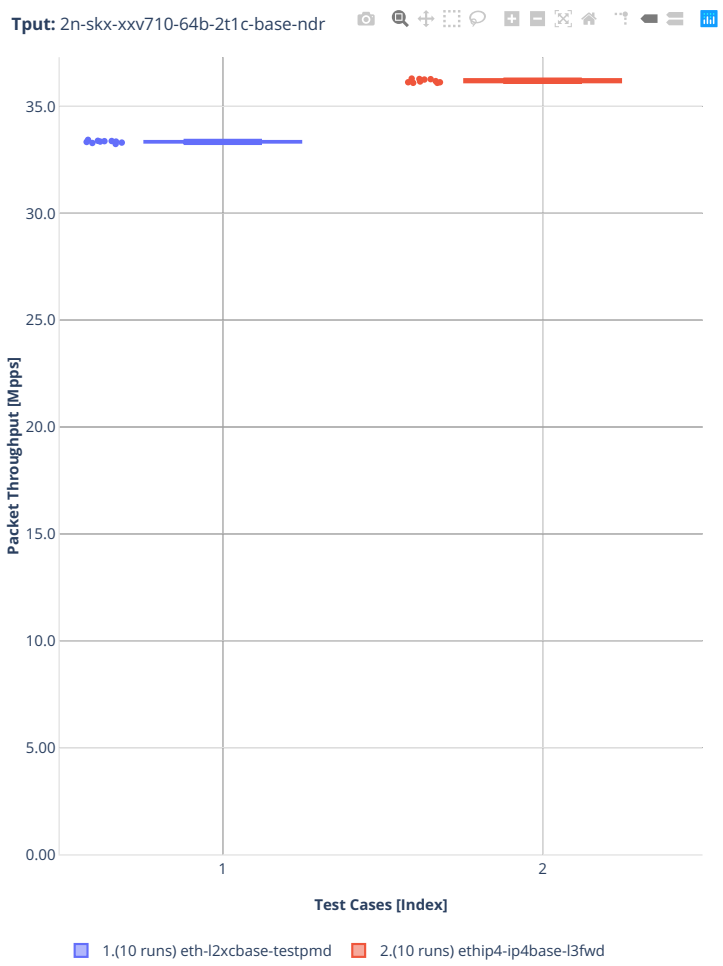

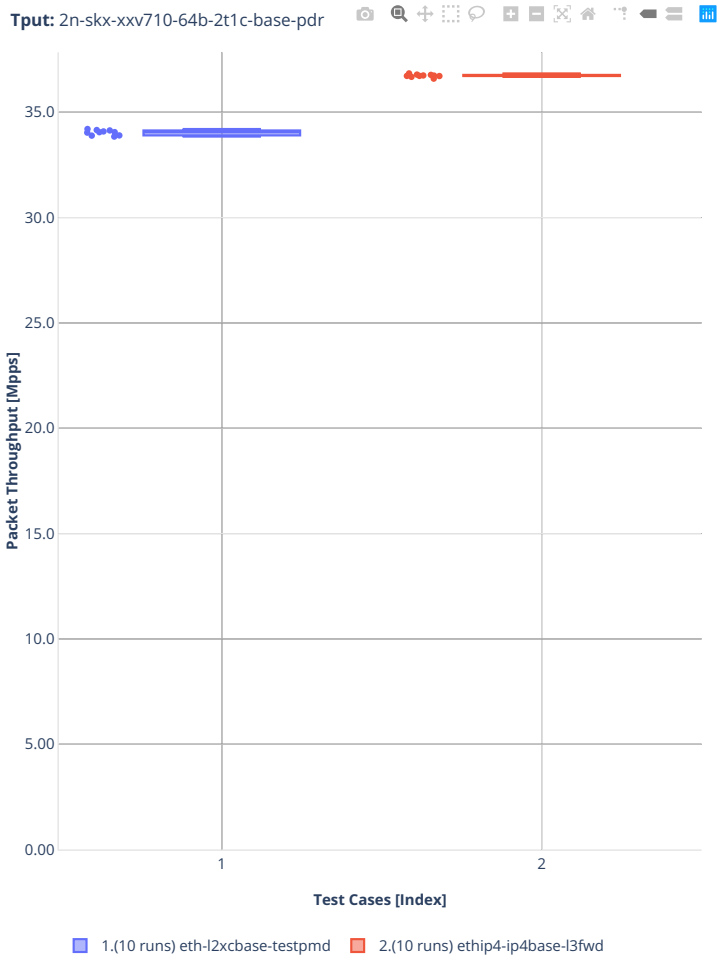
Tput: 2n-skx-xxv710-64b-2t1c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

64b-4t2c-base

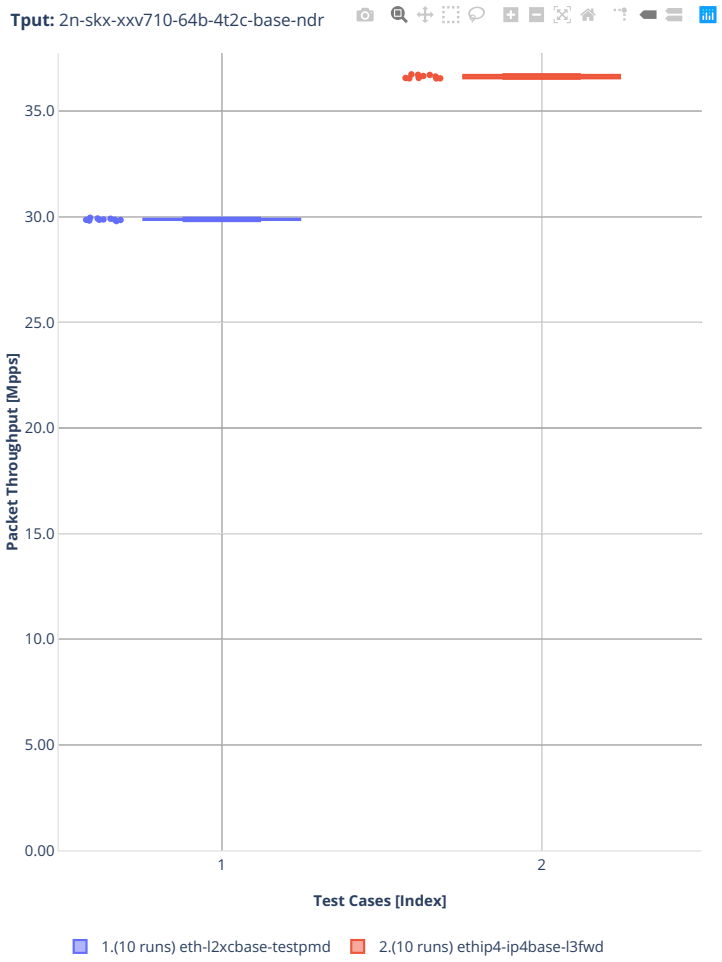

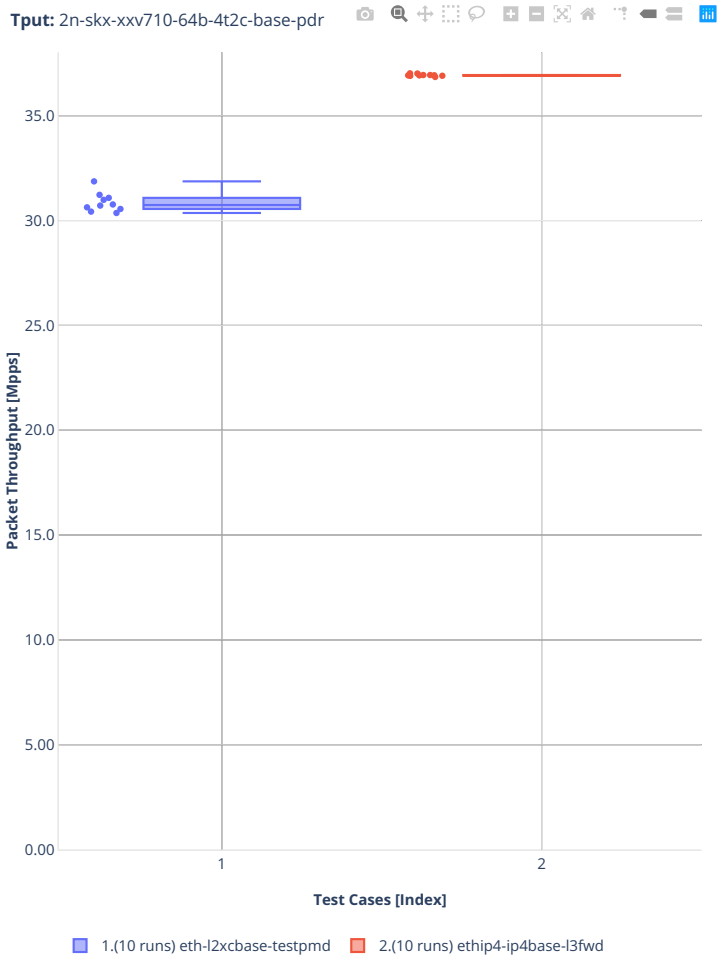
Tput: 2n-skx-xxv710-64b-4t2c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

3.3.4 2n-skx-x710

Following sections include summary graphs of Phy-to-Phy performance with packet 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[281].

[281] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

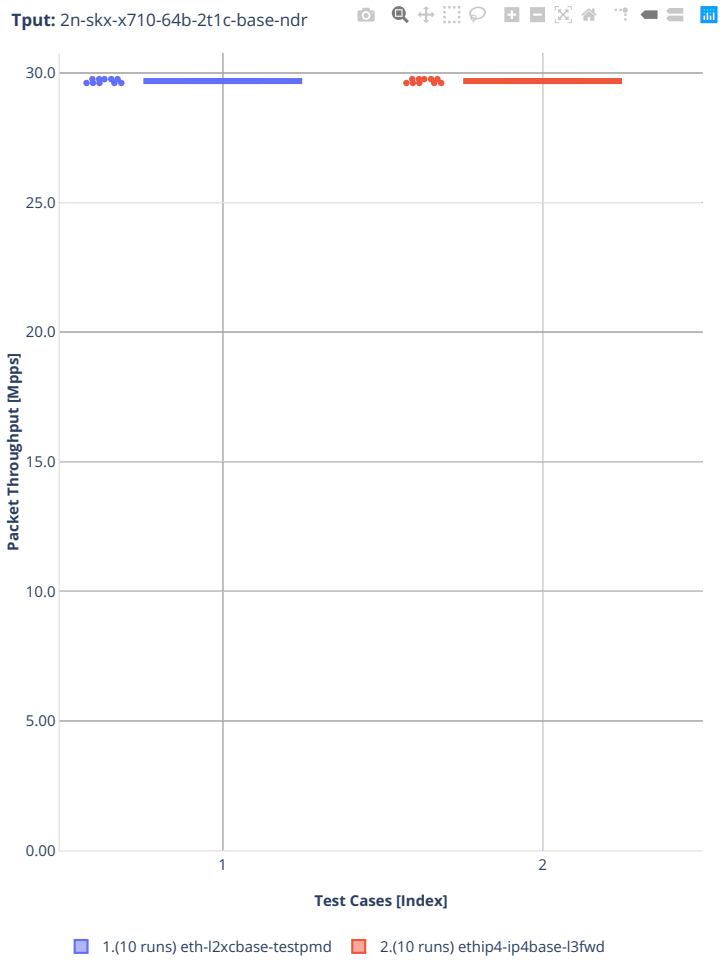

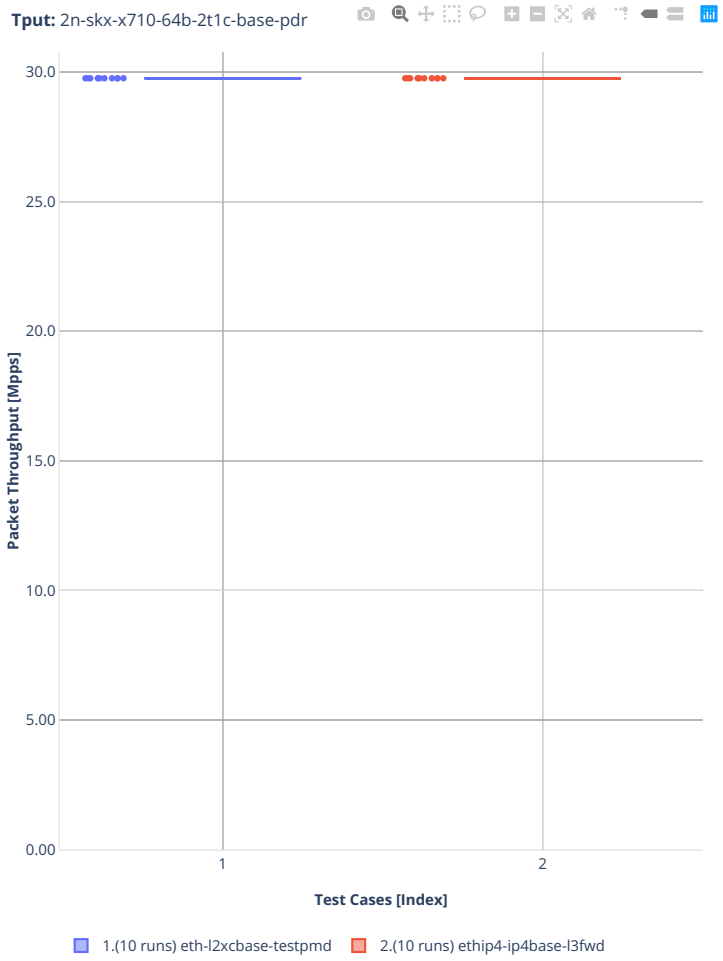
Tput: 2n-skx-x710-64b-2t1c-base-pdr
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

64b-4t2c-base

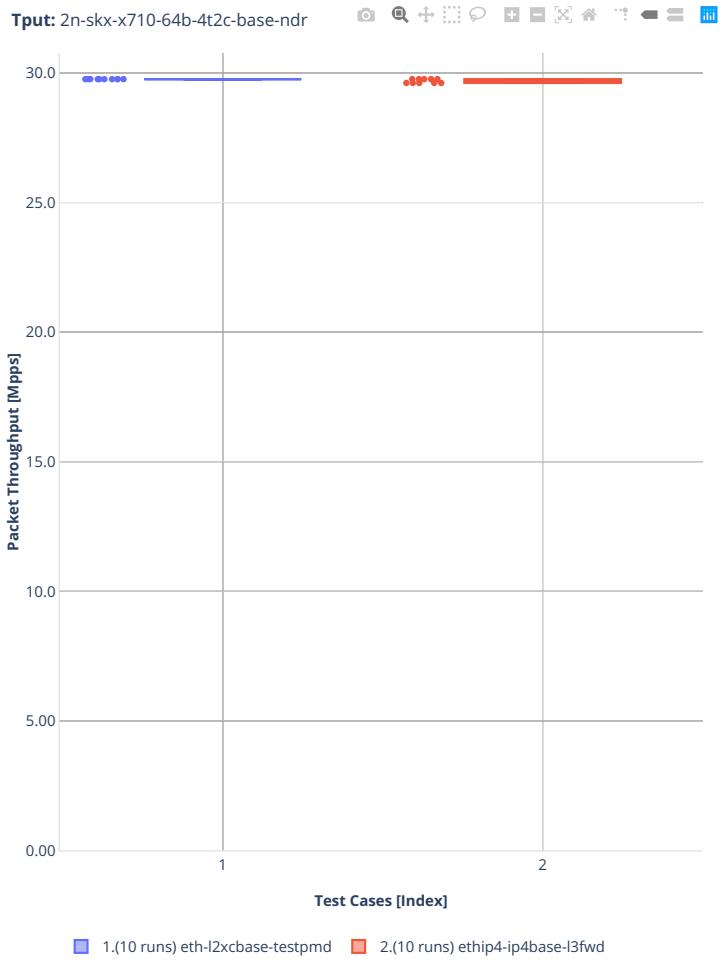

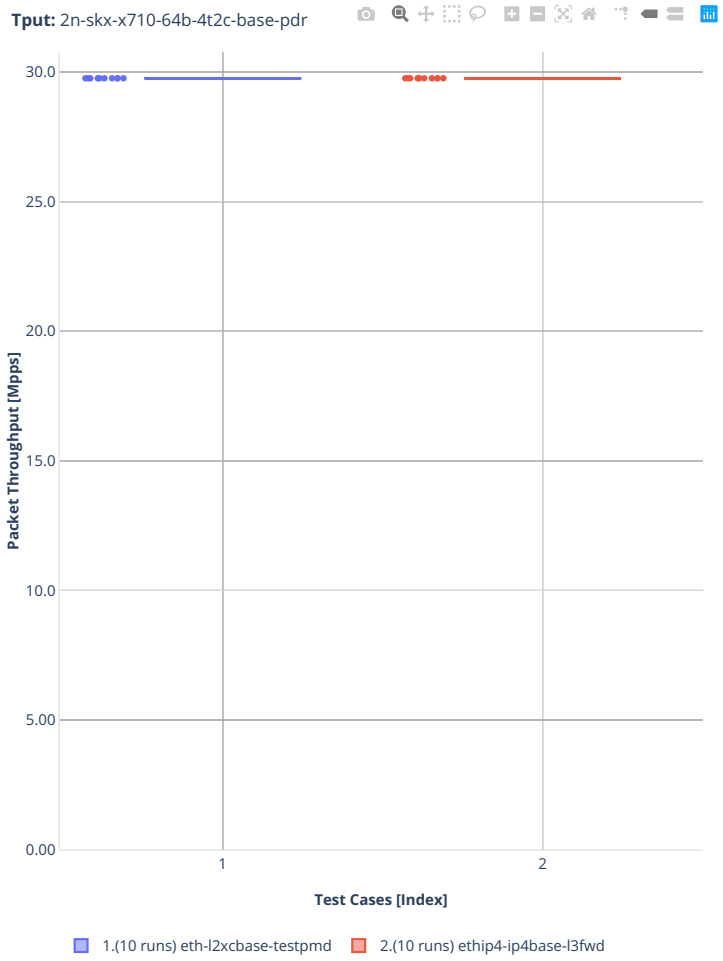
Tput: 2n-skx-x710-64b-4t2c-base-pdr
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

3.3.5 3n-skx-xxv710

Following sections include summary graphs of Phy-to-Phy performance with packet 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[282].

[282] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

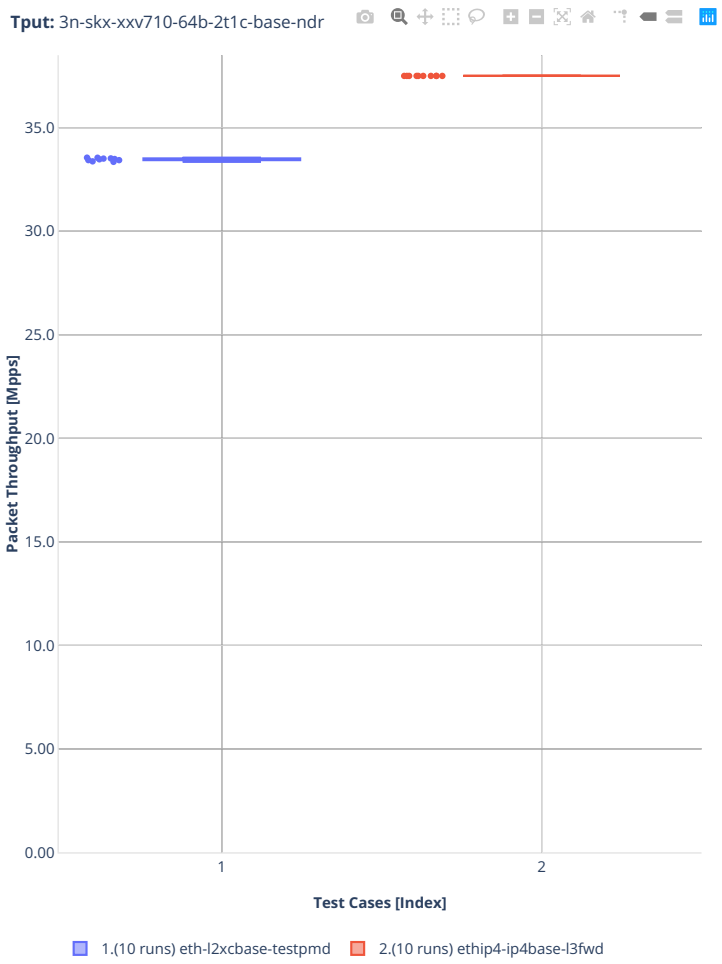

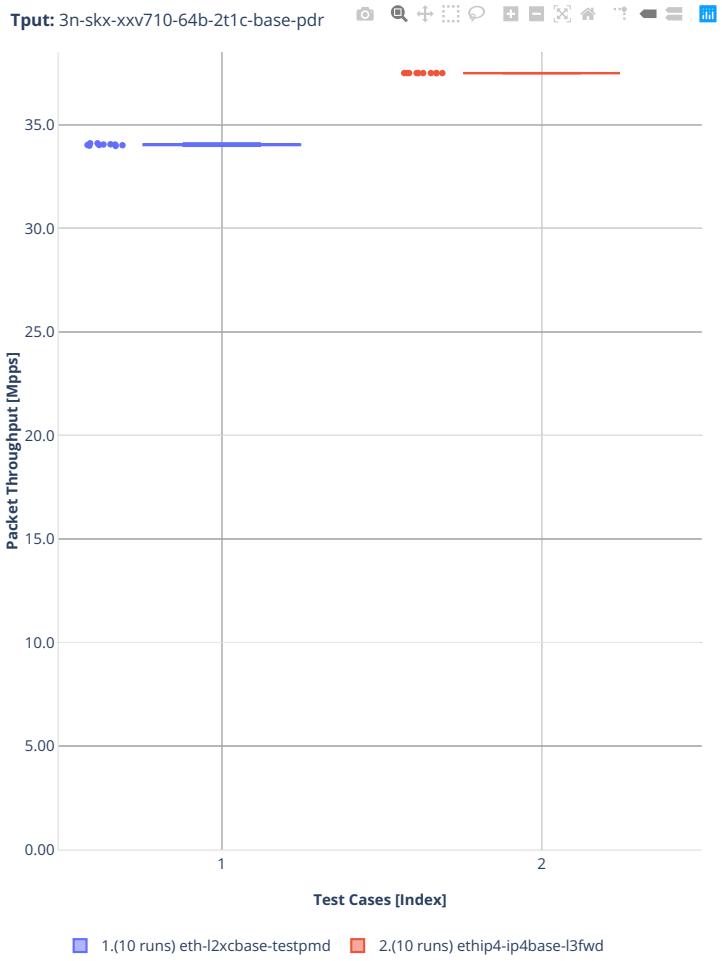
Tput: 3n-skx-xxv710-64b-2t1c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

64b-4t2c-base

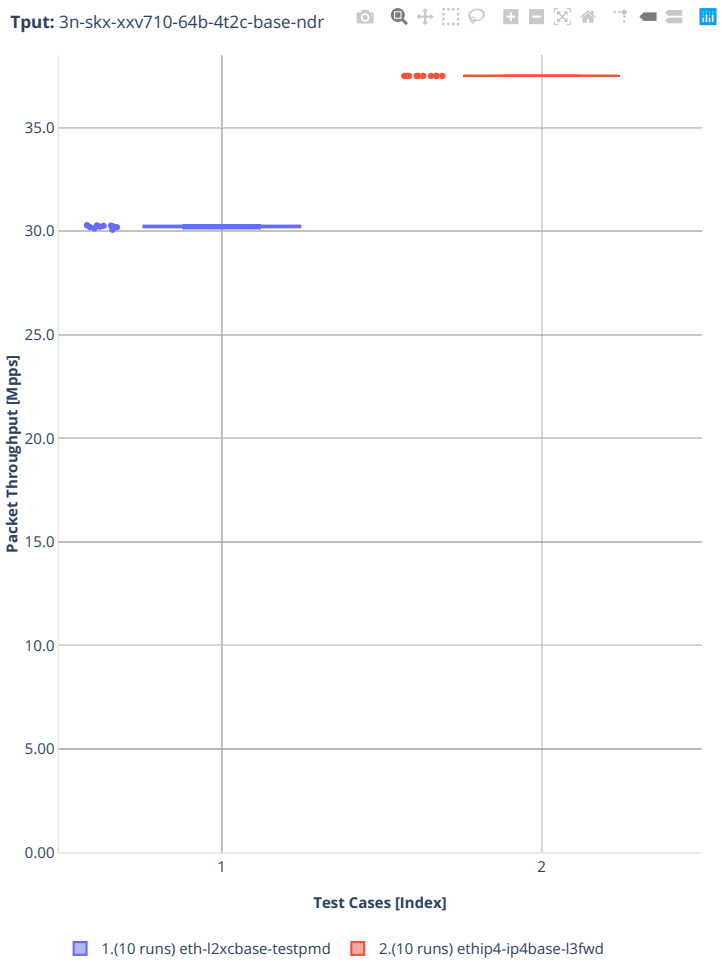

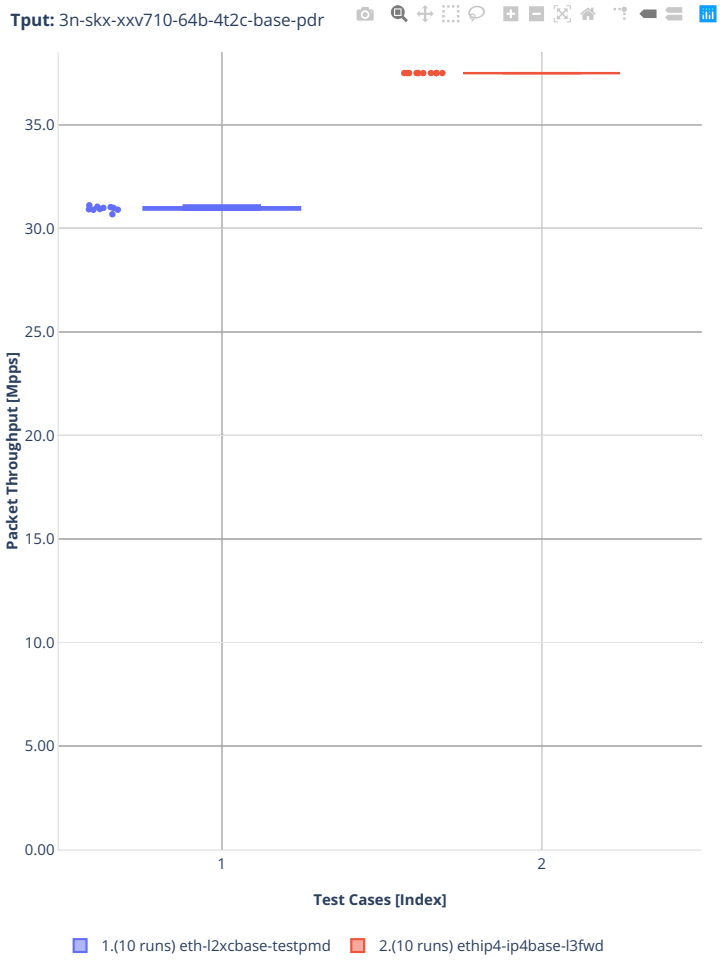
Tput: 3n-skx-xxv710-64b-4t2c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

3.3.6 3n-skx-x710

Following sections include summary graphs of Phy-to-Phy performance with packet 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[283].

[283] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

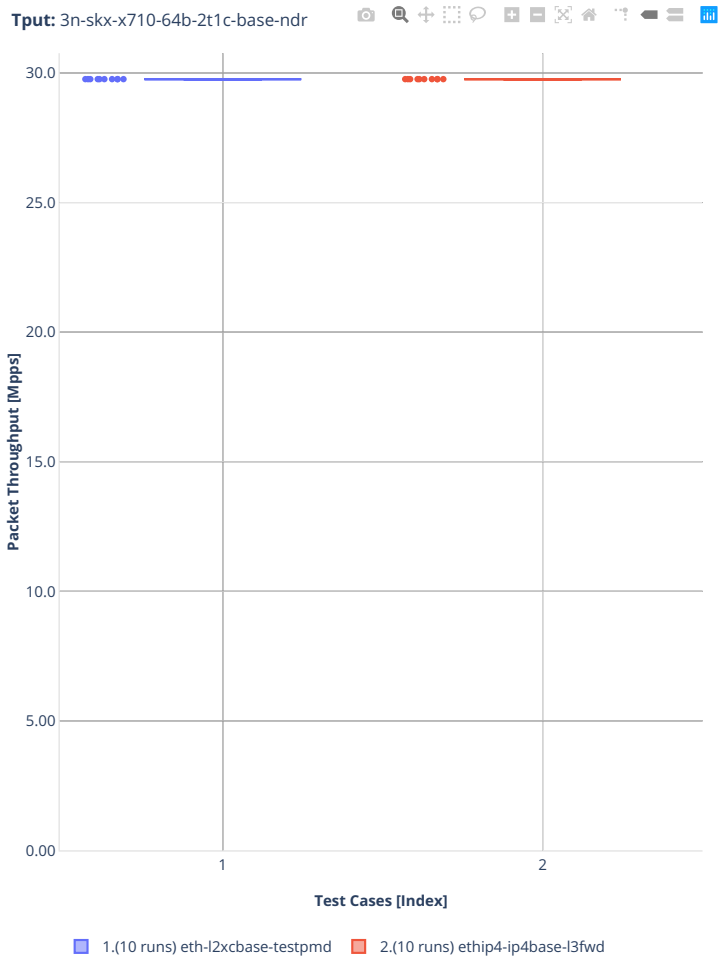

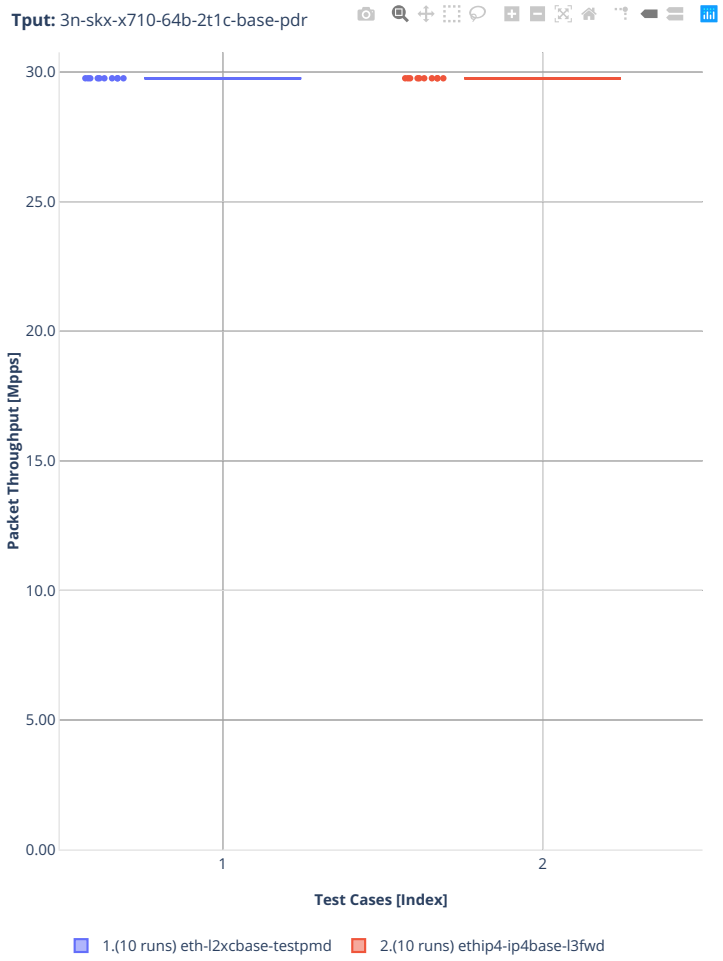
Tput: 3n-skx-x710-64b-2t1c-base-pdr
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

64b-4t2c-base

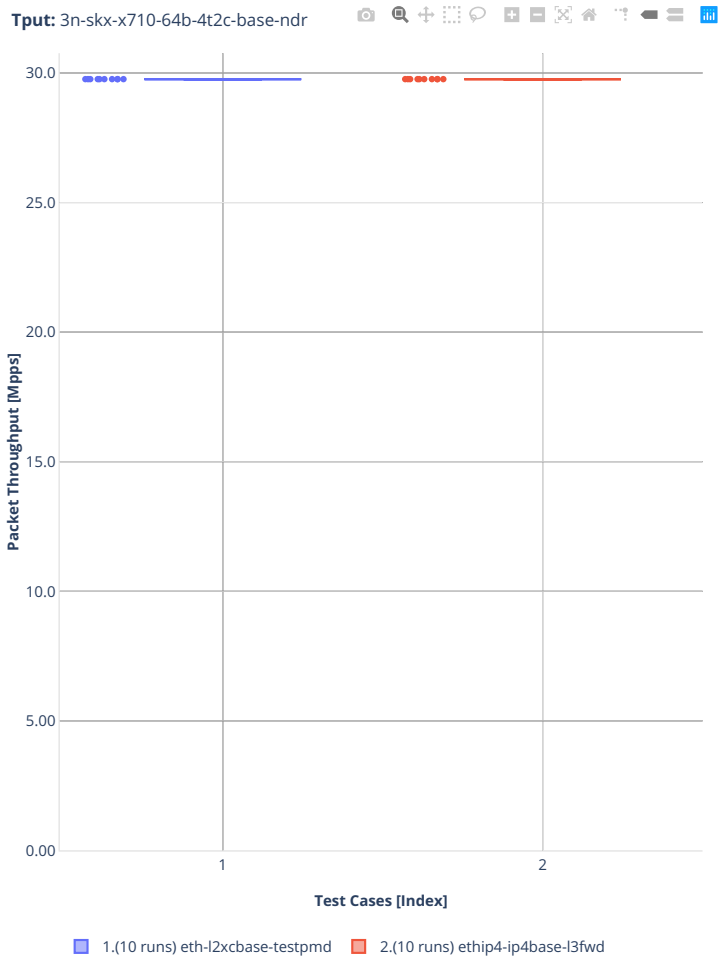

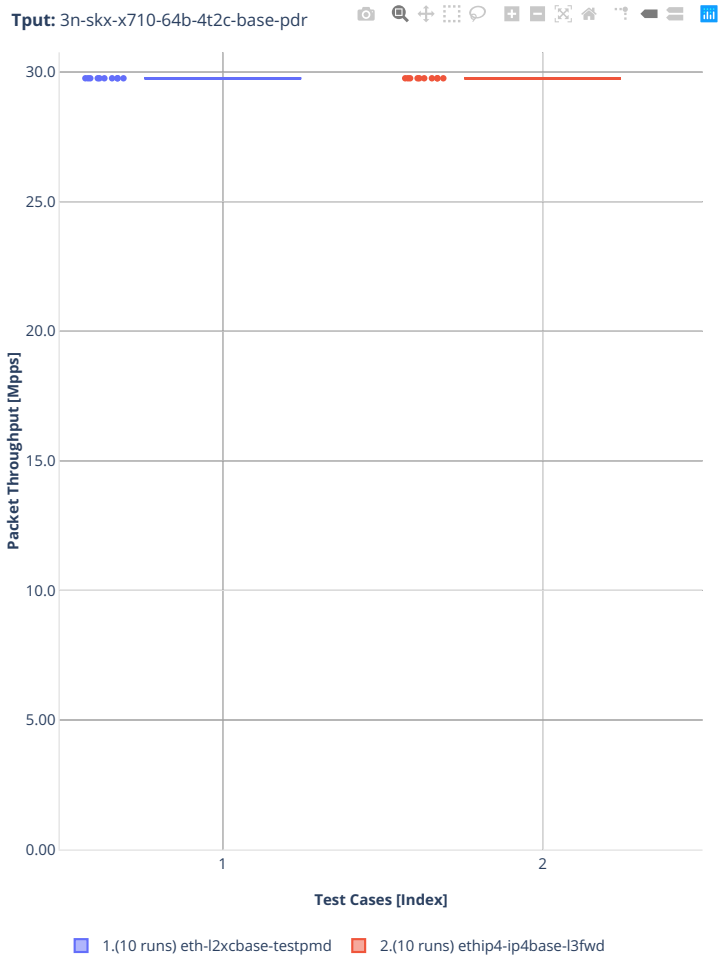
Tput: 3n-skx-x710-64b-4t2c-base-pdr
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

3.3.7 2n-clx-xxv710

Following sections include summary graphs of Phy-to-Phy performance with packet 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[284].

[284] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

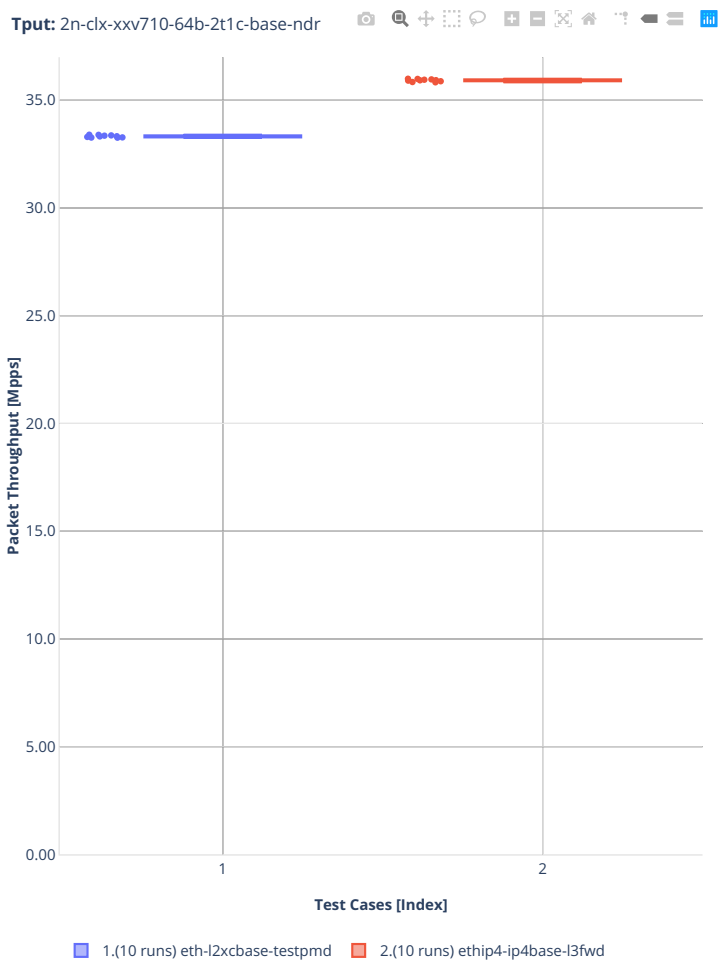

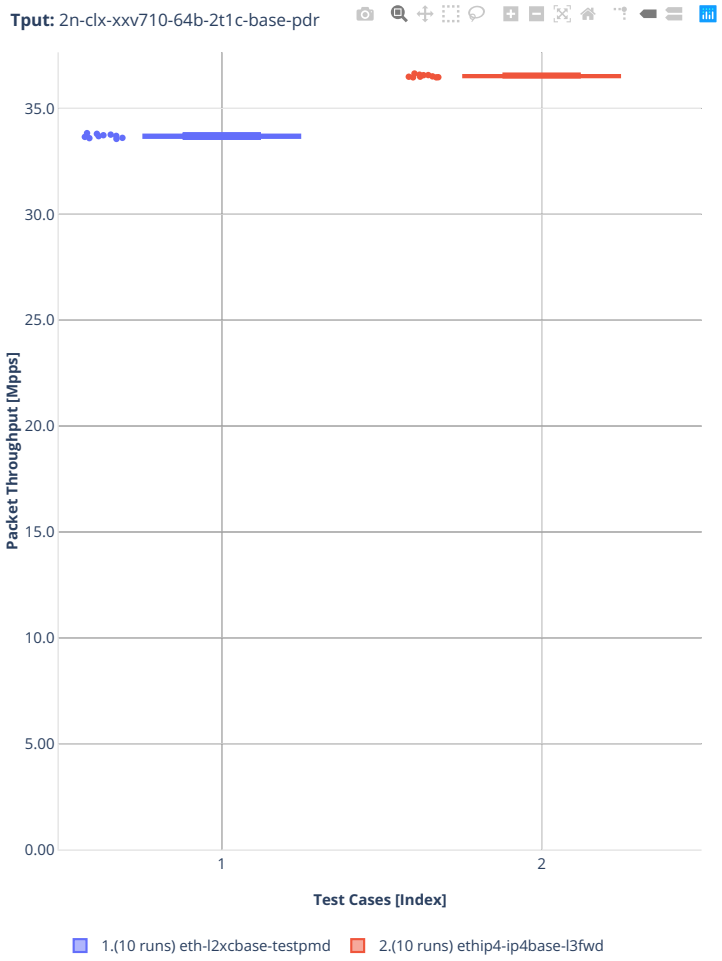
Tput: 2n-clx-xxv710-64b-2t1c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

64b-4t2c-base

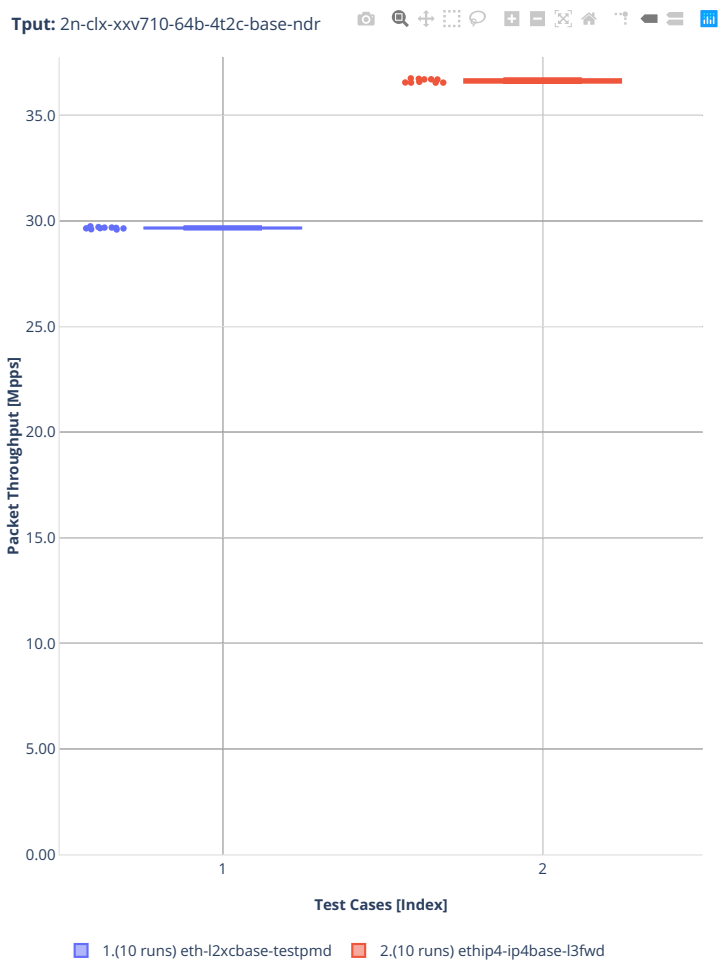

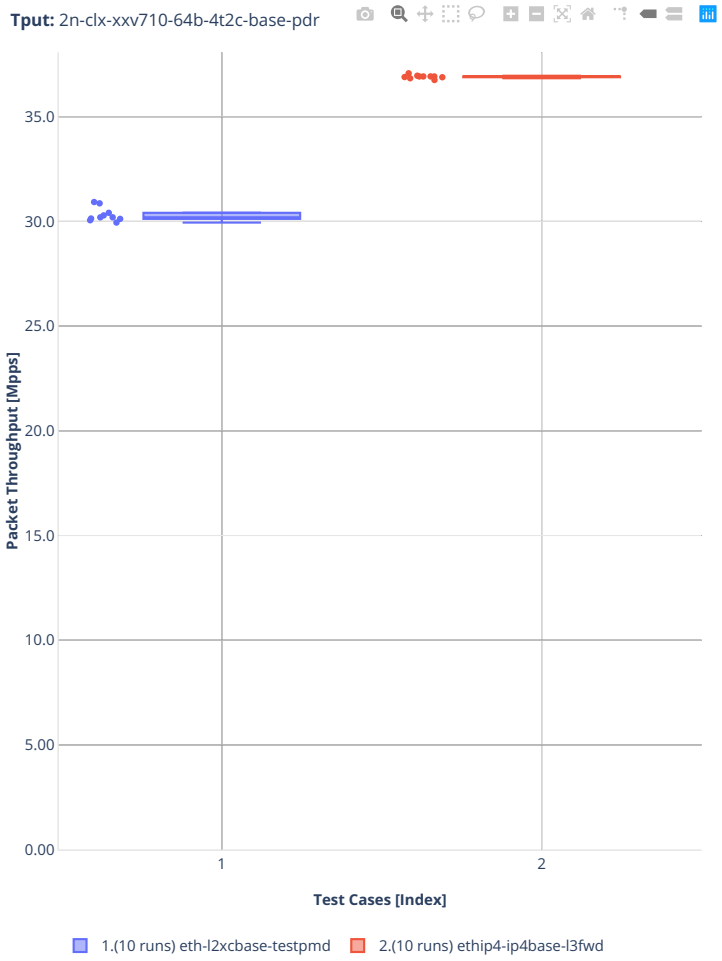
Tput: 2n-clx-xxv710-64b-4t2c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

3.3.8 2n-clx-x710

Following sections include summary graphs of Phy-to-Phy performance with packet 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[285].

[285] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

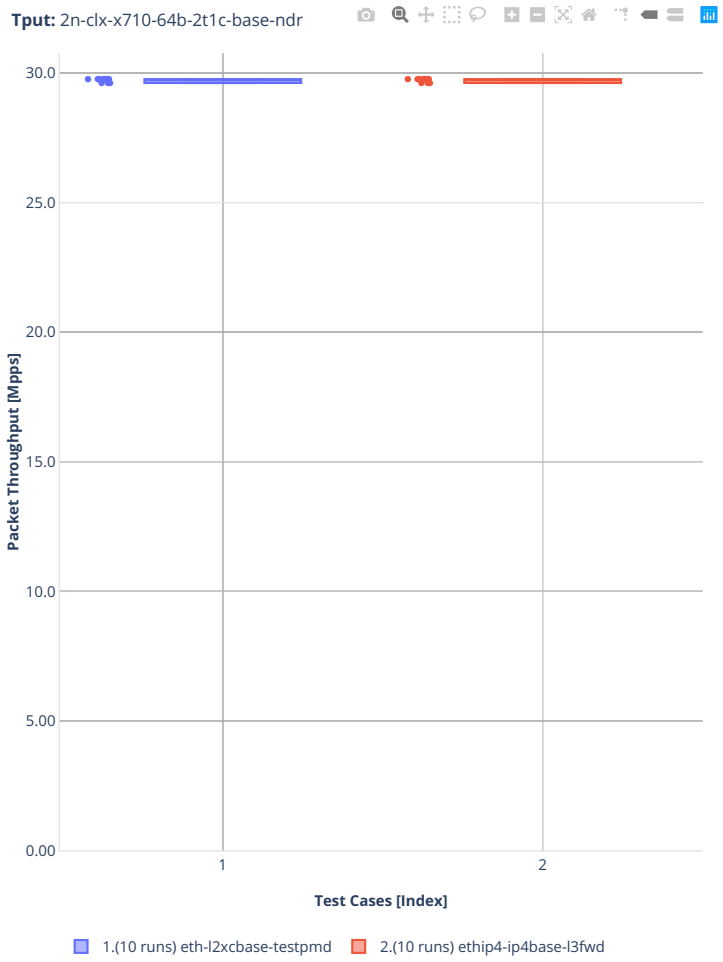

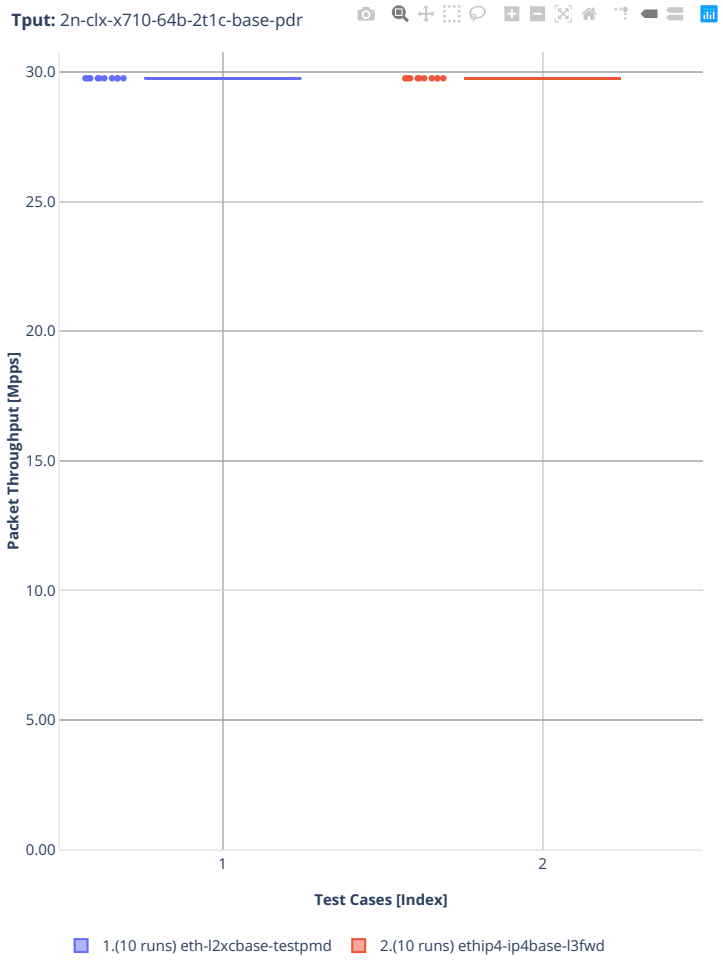
Tput: 2n-clx-x710-64b-2t1c-base-pdr
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

64b-4t2c-base

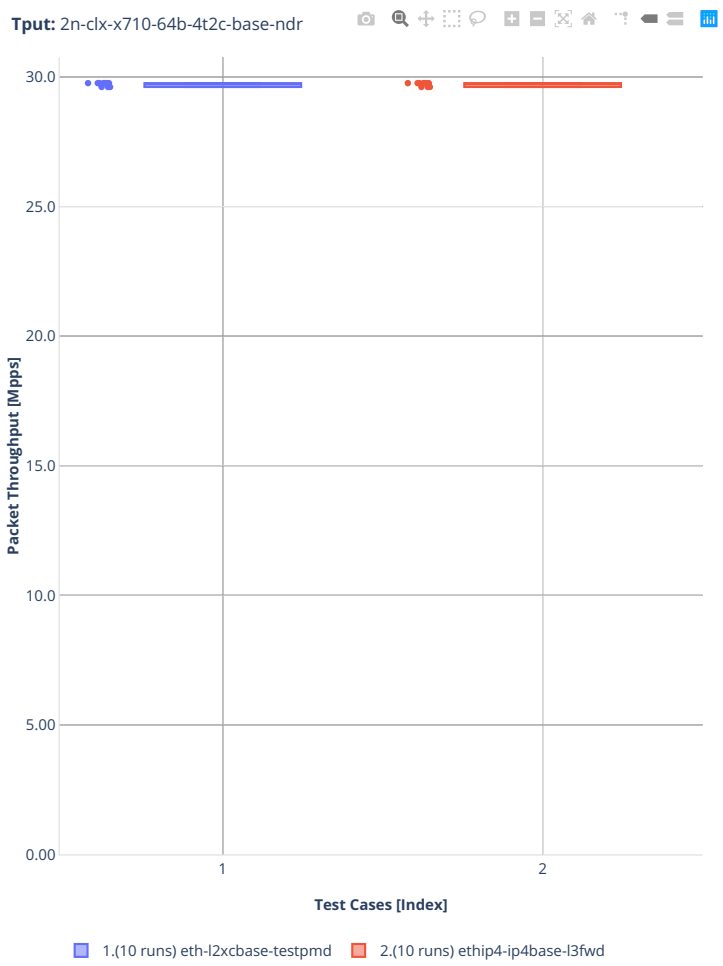

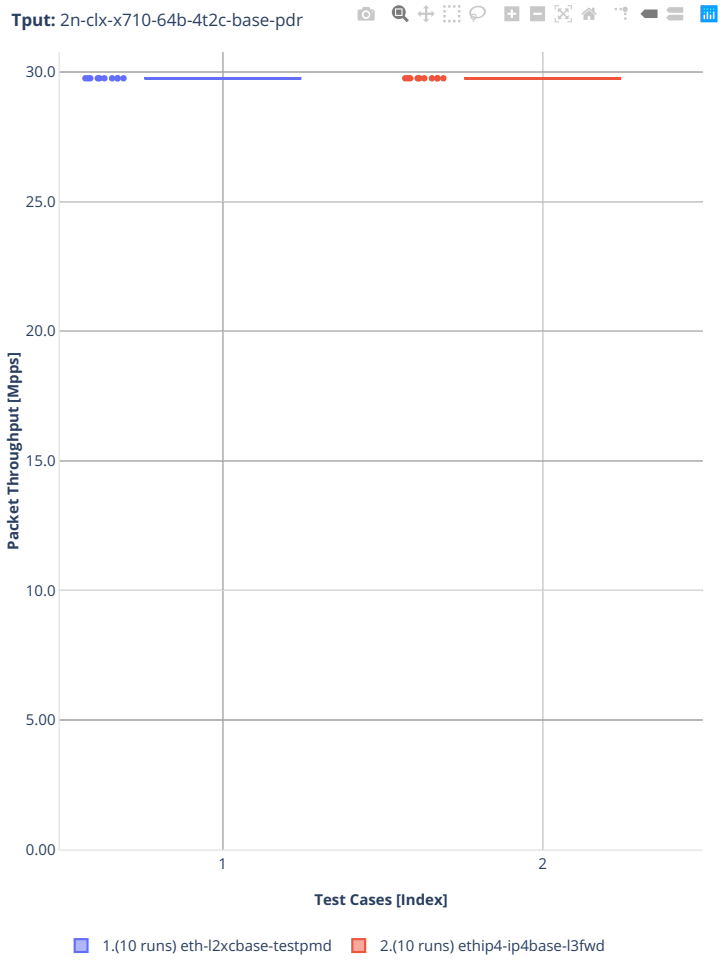
Tput: 2n-clx-x710-64b-4t2c-base-pdr
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

3.3.9 2n-zn2-xxv710

Following sections include summary graphs of Phy-to-Phy performance with packet 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[286].

[286] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

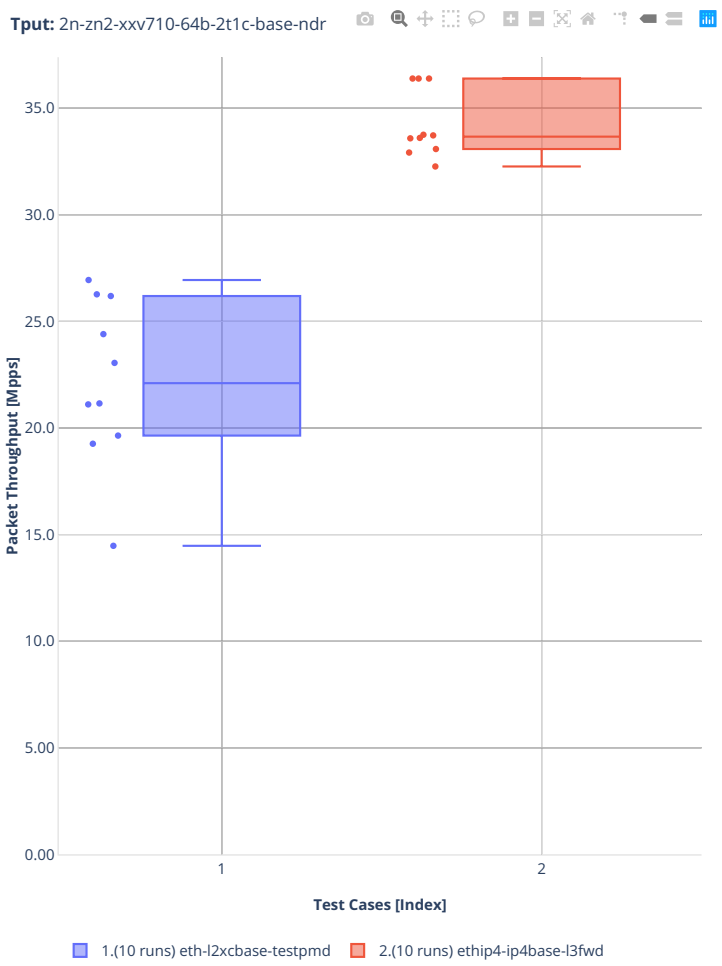

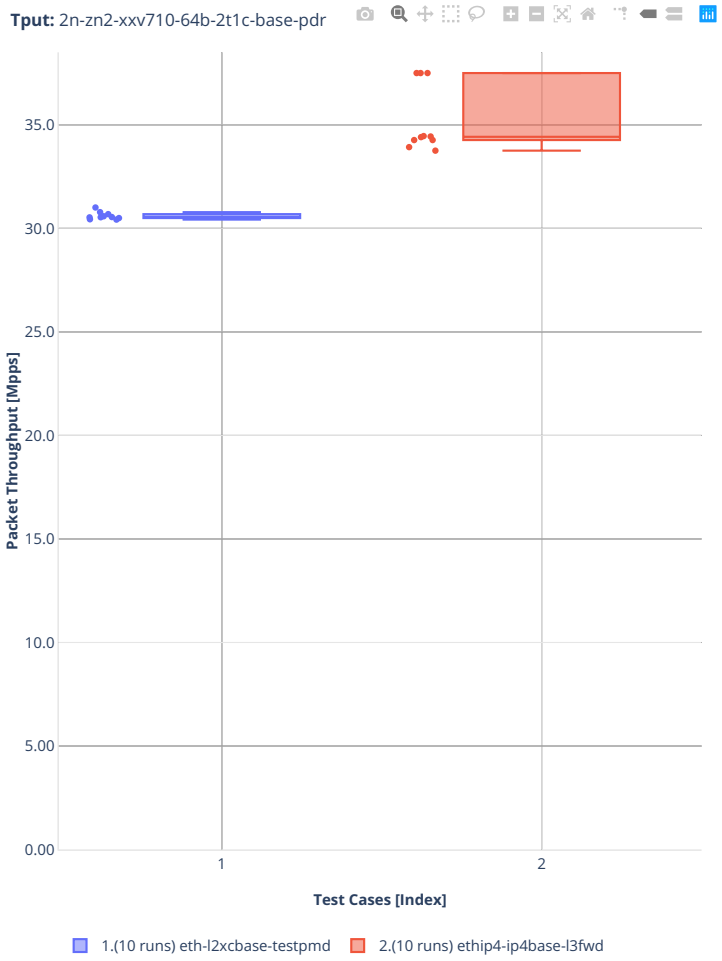
Tput: 2n-zn2-xxv710-64b-2t1c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

64b-4t2c-base

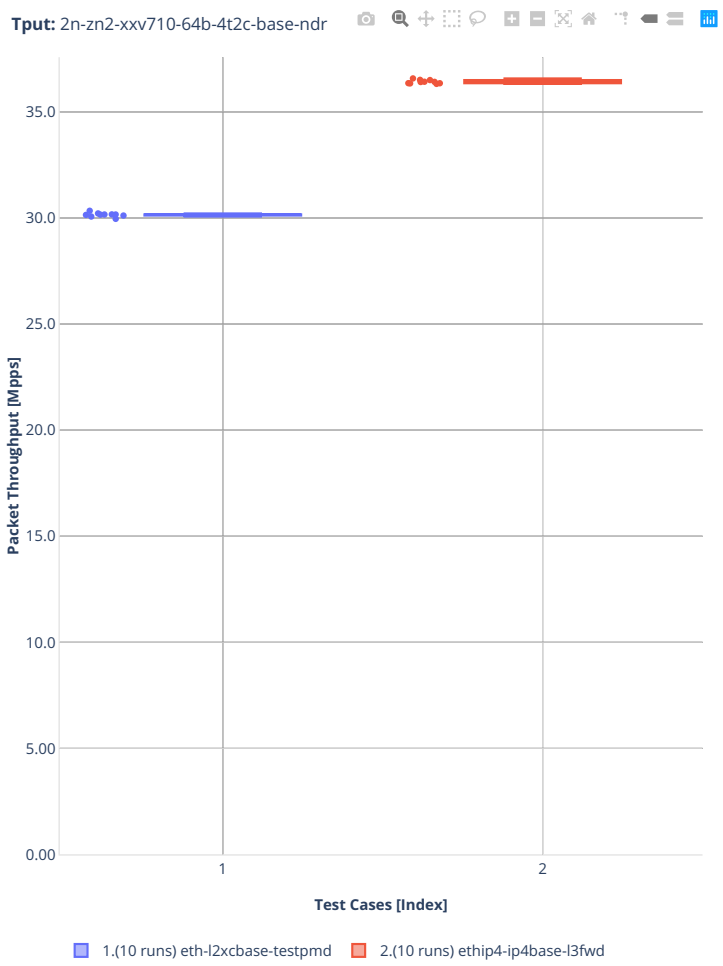

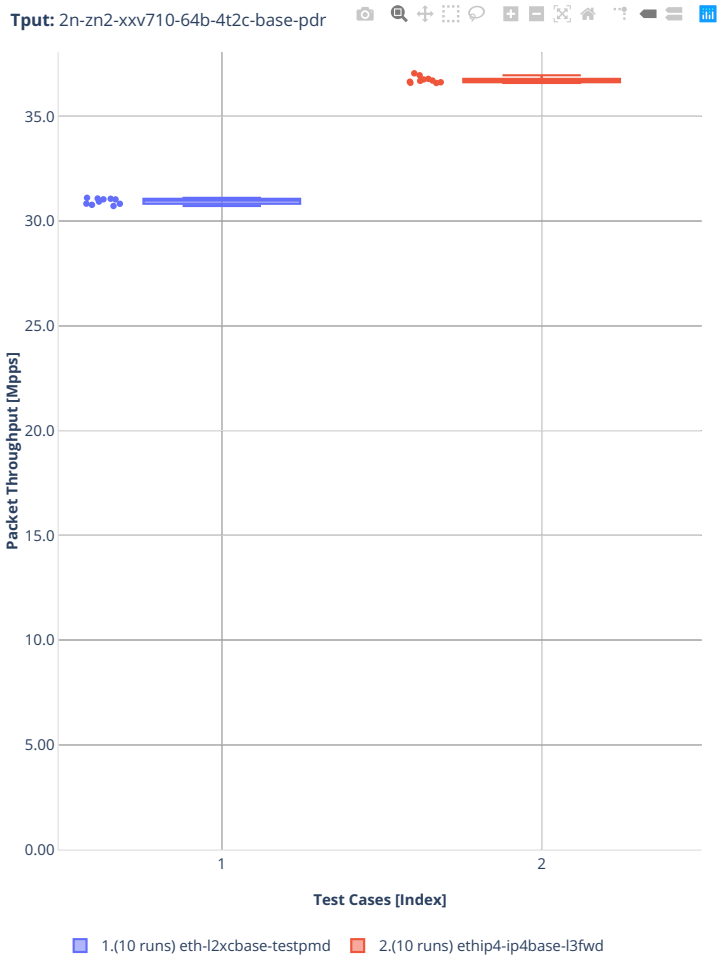
Tput: 2n-zn2-xxv710-64b-4t2c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

3.3.10 2n-zn2-x710

Following sections include summary graphs of Phy-to-Phy performance with packet 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[287].

[287] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

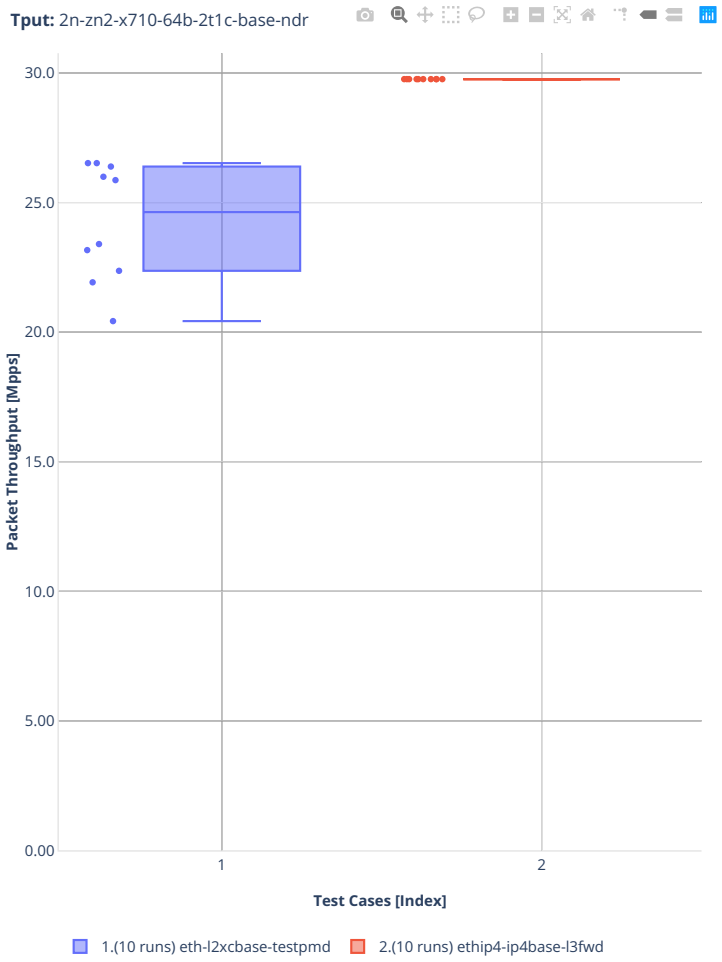

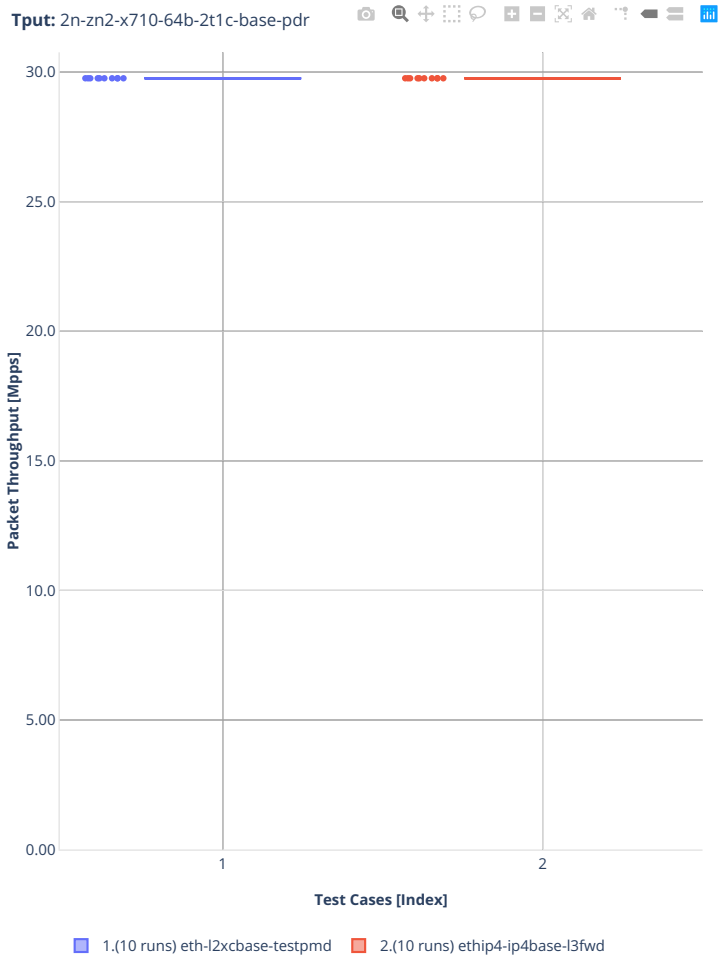
Tput: 2n-zn2-x710-64b-2t1c-base-pdr
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

64b-4t2c-base

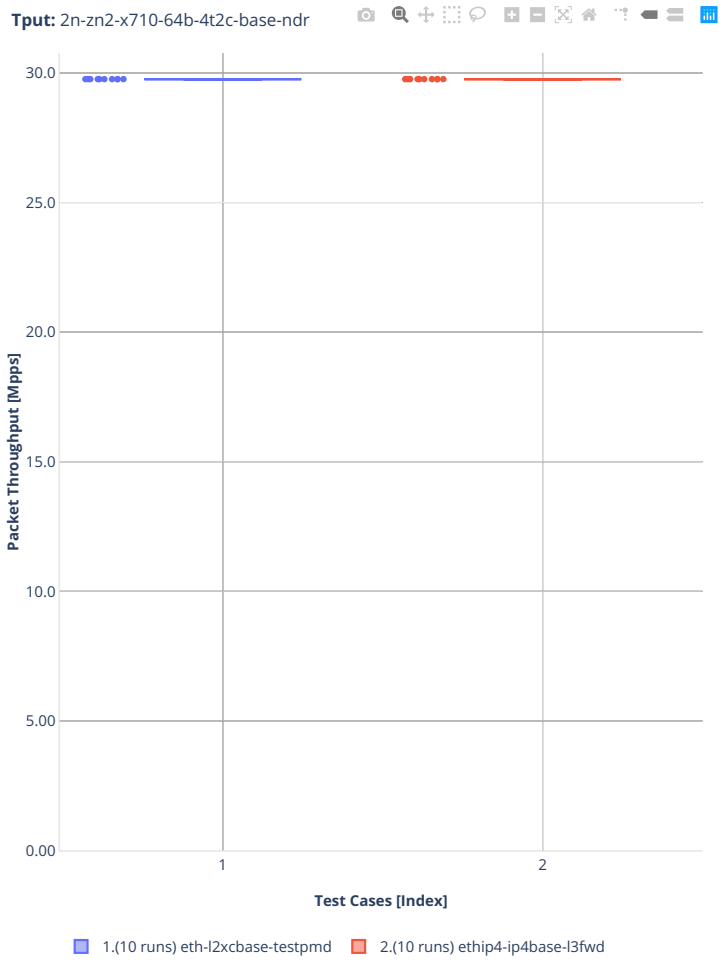

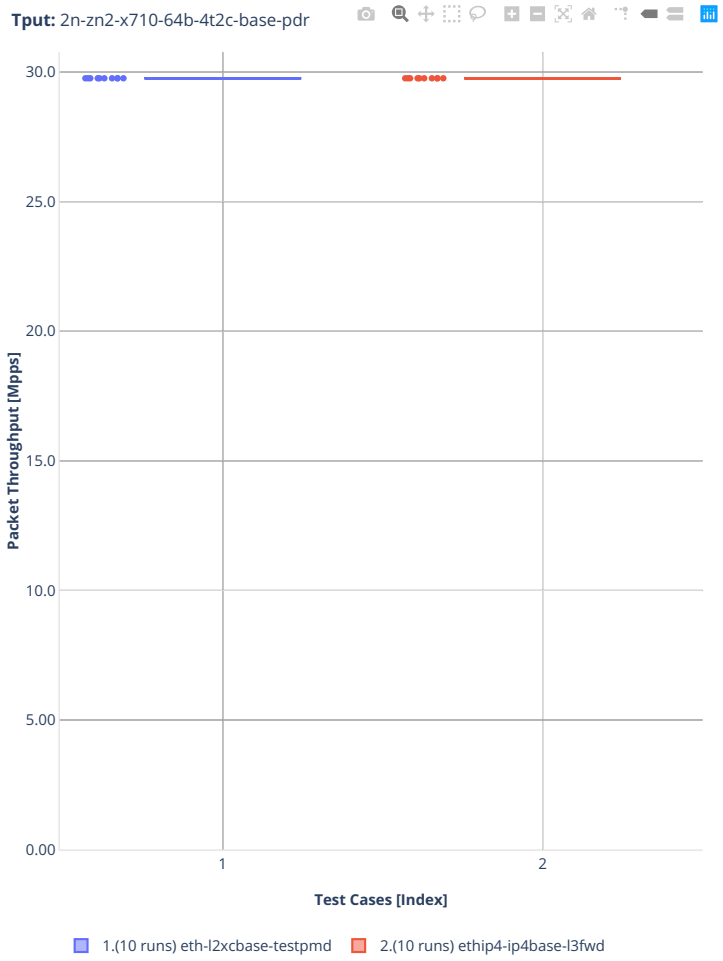
Tput: 2n-zn2-x710-64b-4t2c-base-pdr
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

3.3.11 3n-alt-xl710

Following sections include summary graphs of Phy-to-Phy performance with packet 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[288].

[288] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-1t1c-base

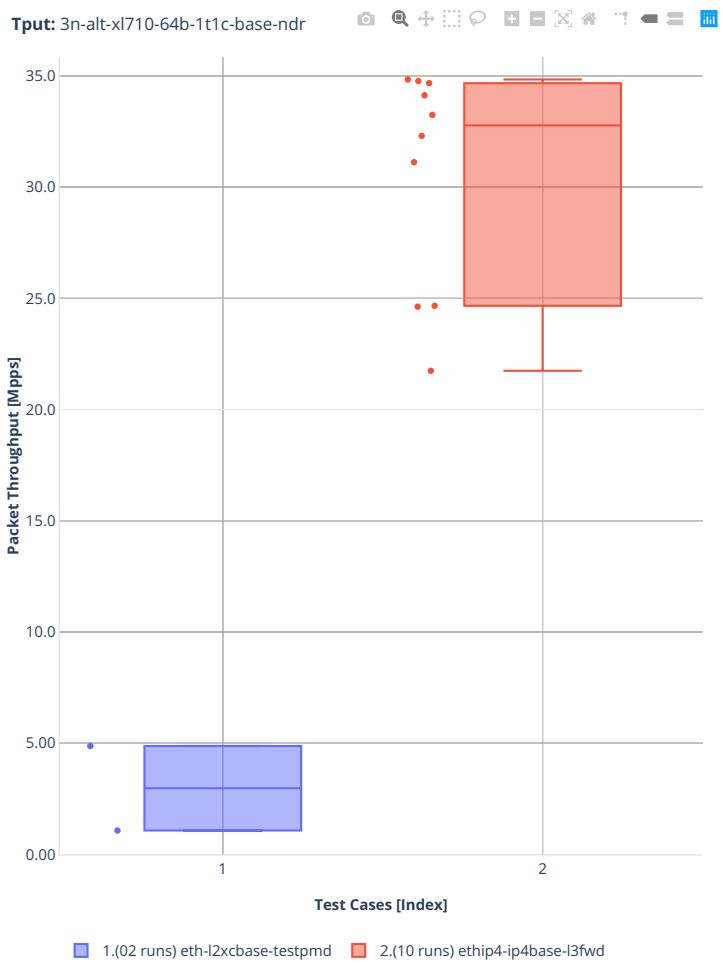

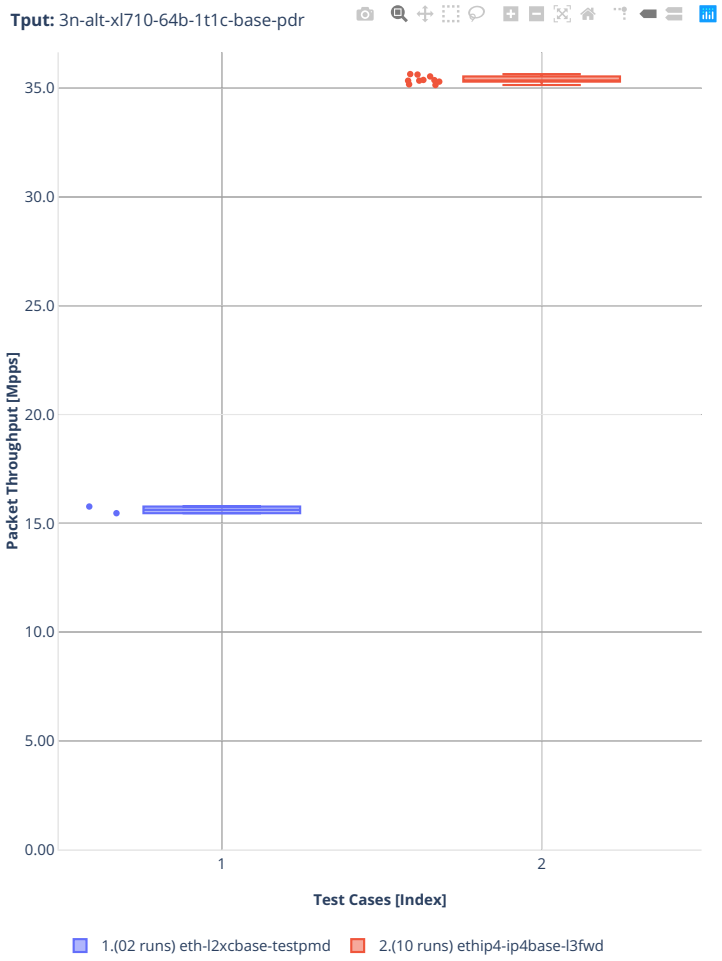
Tput: 3n-alt-xl710-64b-1t1c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(02 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

64b-2t2c-base

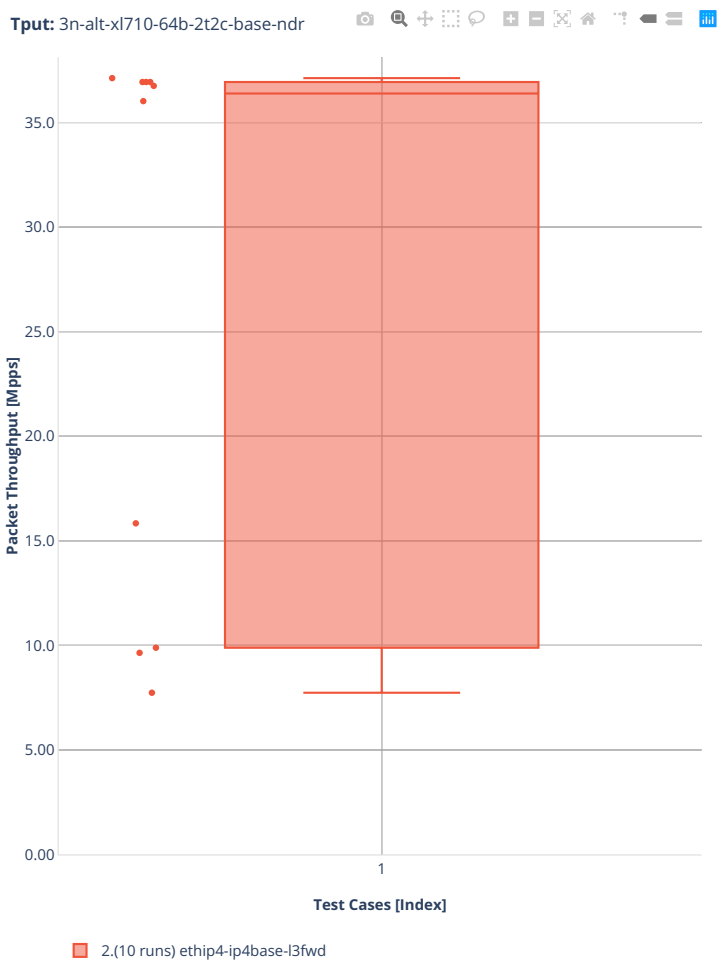

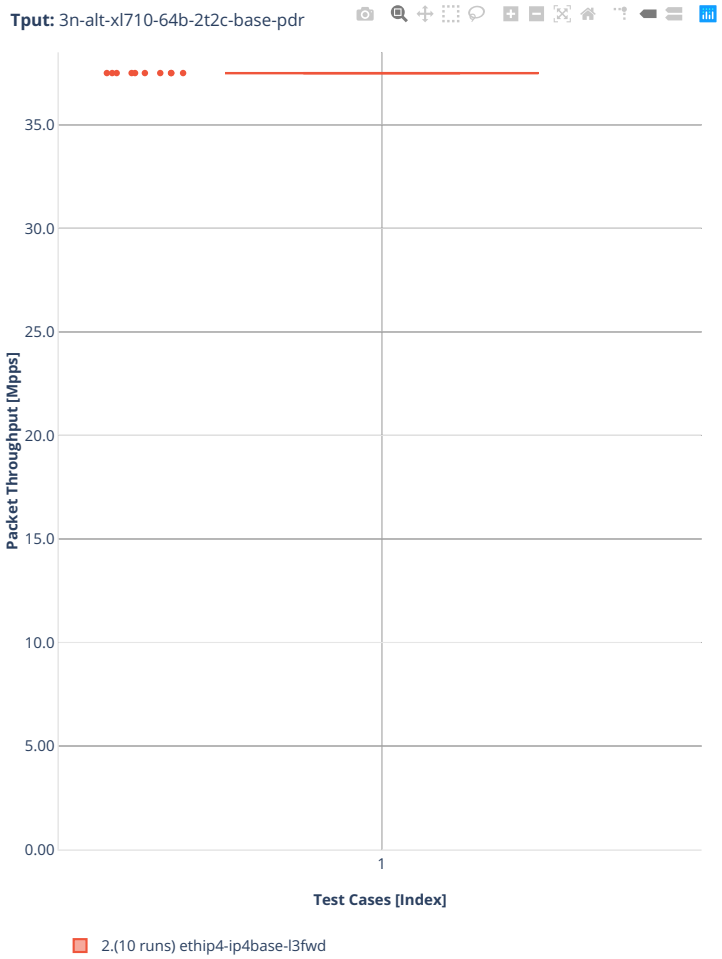
Tput: 3n-alt-xl710-64b-2t2c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
Test Cases [Index]
2.(10 runs) ethip4-ip4base-l3fwd

3.3.12 3n-tsh-x520

Following sections include summary graphs of Phy-to-Phy performance with packet 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[289].

[289] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-1t1c-base

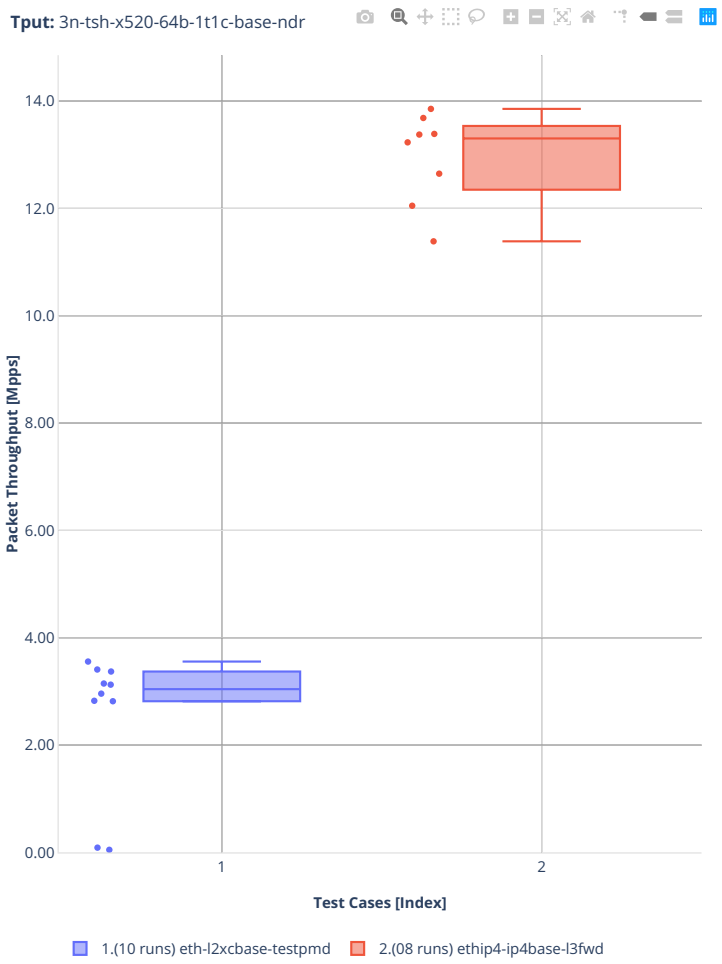

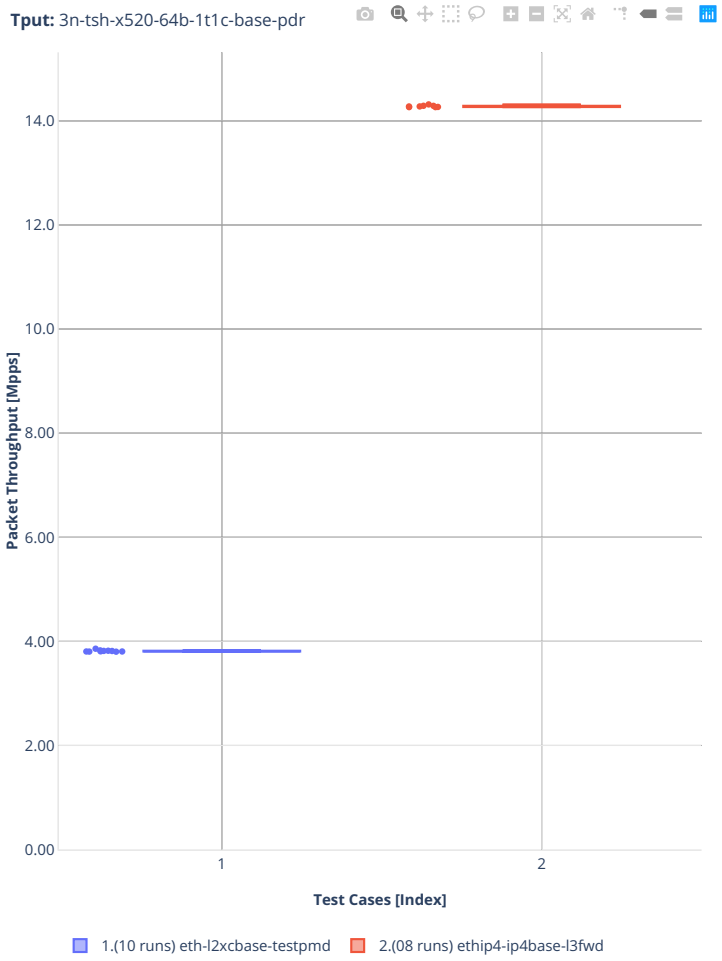
Tput: 3n-tsh-x520-64b-1t1c-base-pdr
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(08 runs) ethip4-ip4base-l3fwd

64b-2t2c-base

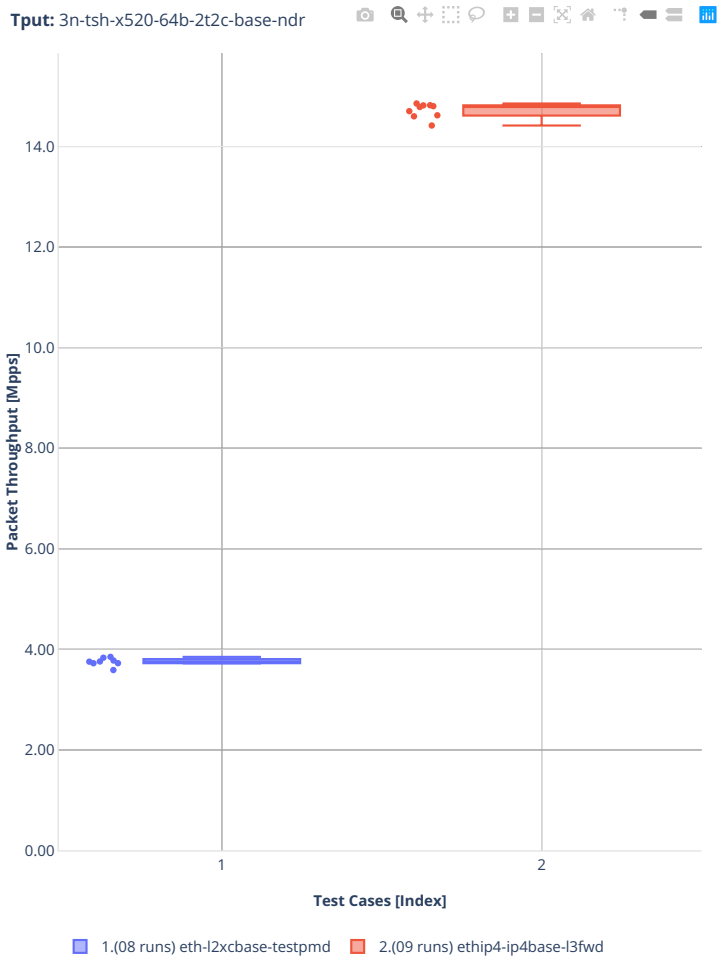

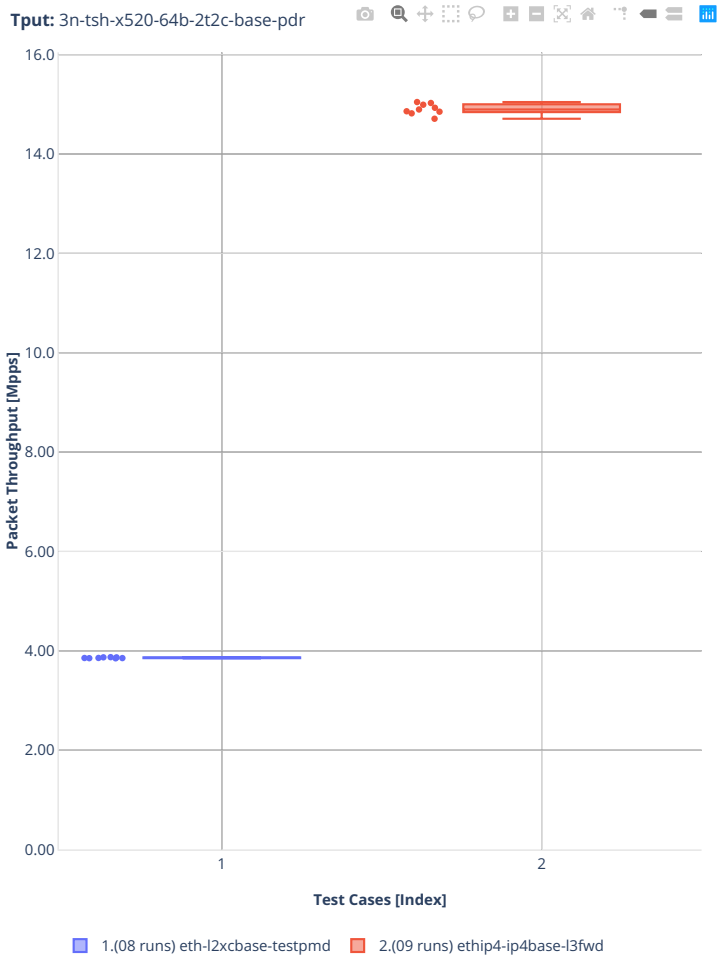
Tput: 3n-tsh-x520-64b-2t2c-base-pdr
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(08 runs) eth-l2xcbase-testpmd
2.(09 runs) ethip4-ip4base-l3fwd

3.3.13 2n-tx2-xl710

Following sections include summary graphs of Phy-to-Phy performance with packet 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[290].

[290] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-1t1c-base

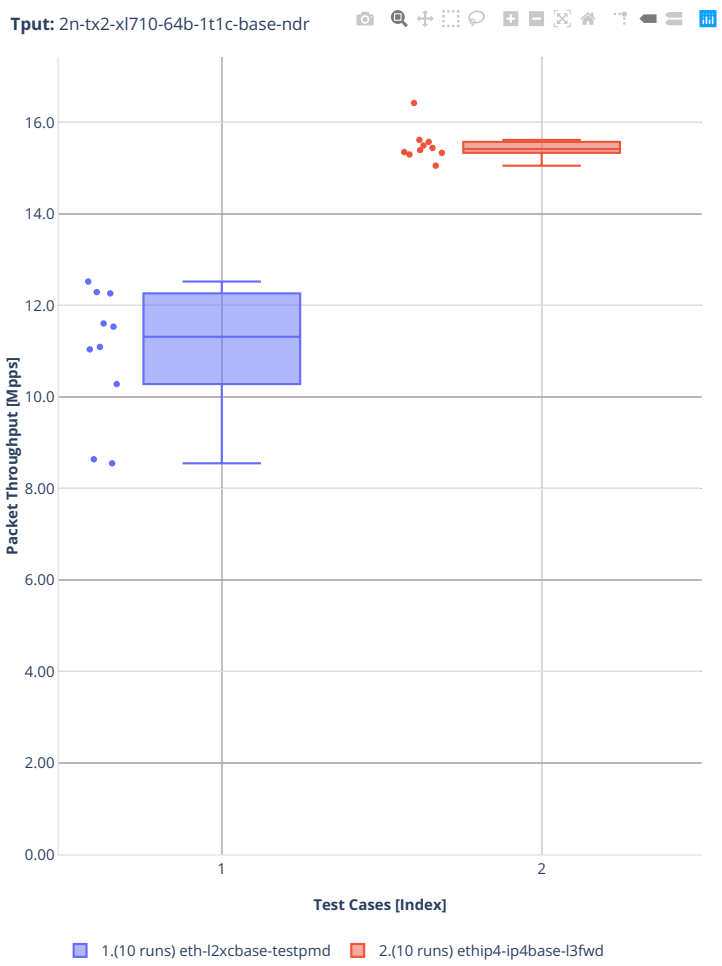

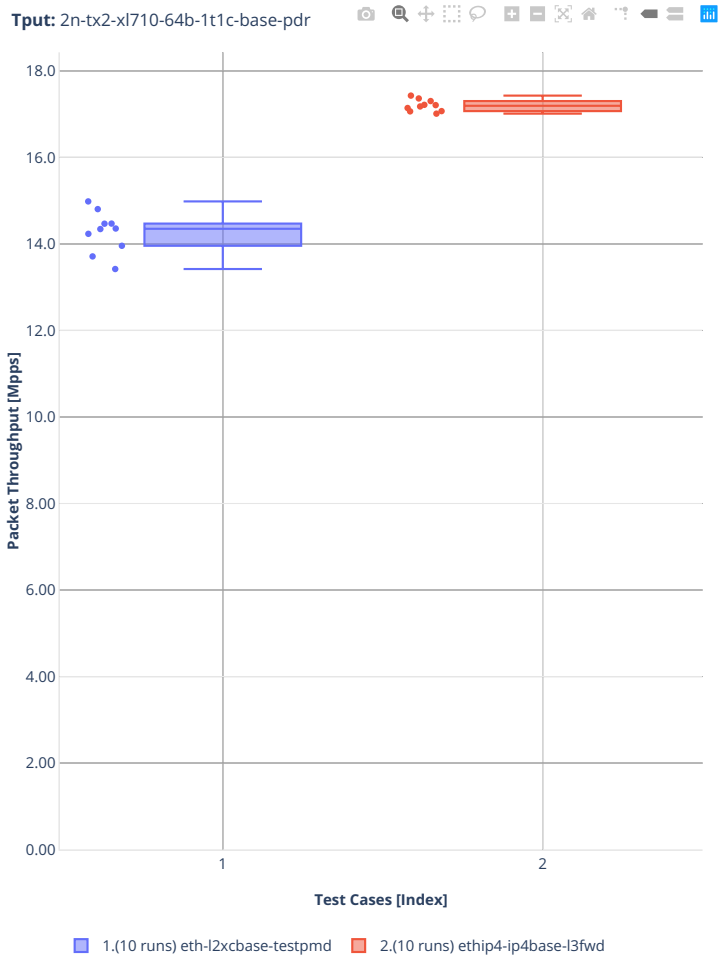
Tput: 2n-tx2-xl710-64b-1t1c-base-pdr
18.0
16.0
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

64b-2t2c-base

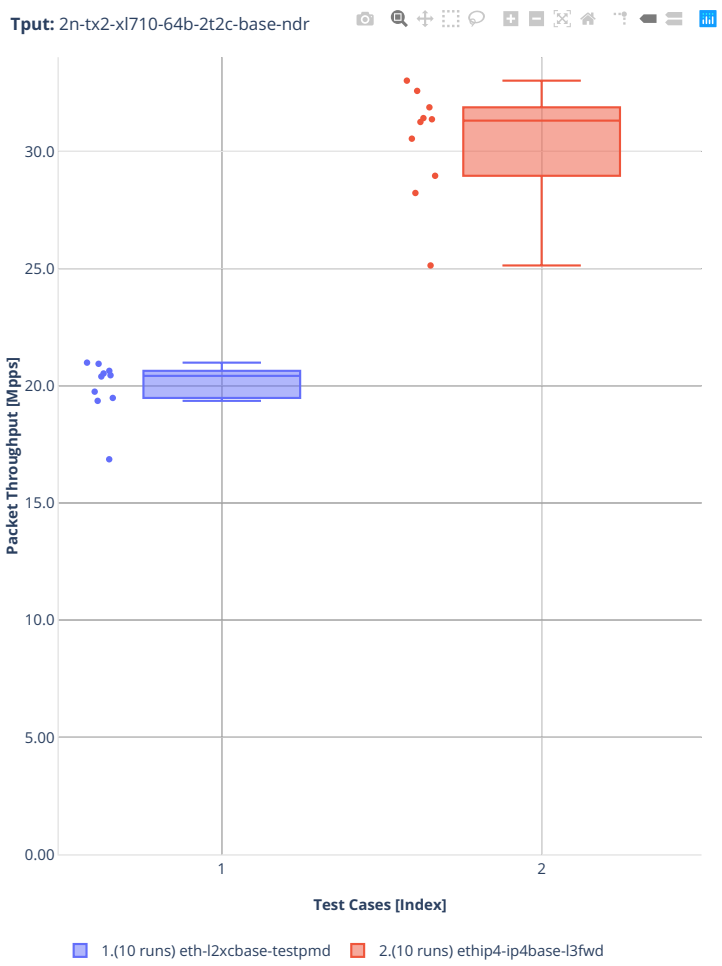

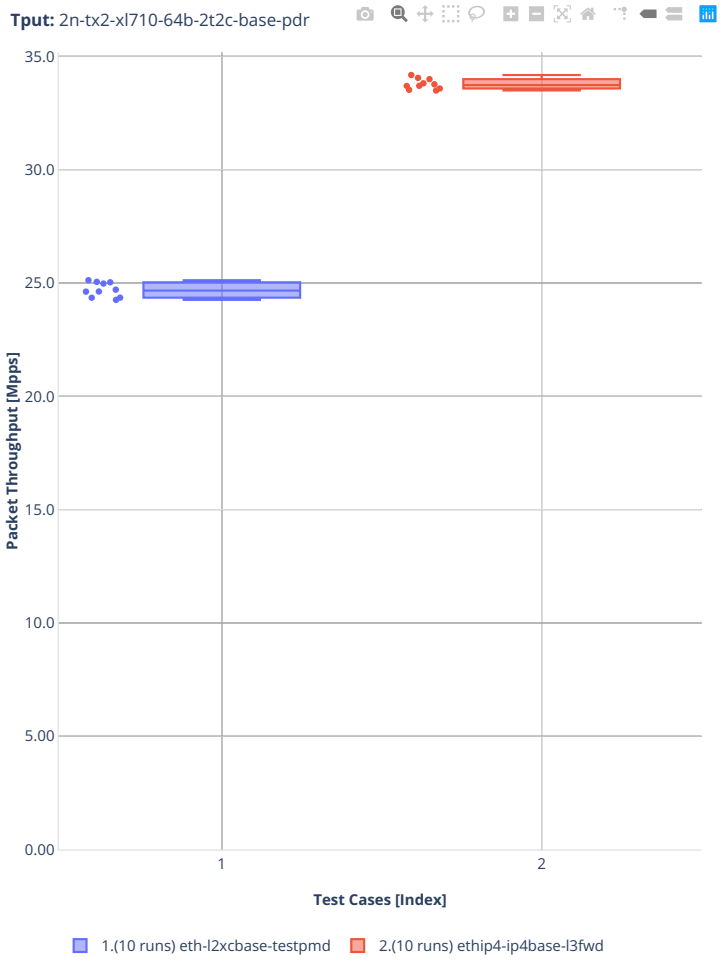
Tput: 2n-tx2-xl710-64b-2t2c-base-pdr
35.0
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) eth-l2xcbase-testpmd
2.(10 runs) ethip4-ip4base-l3fwd

3.4 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: 2n-skx, 3n-skx, 2n-clx, 3n-tsh, 2n-tx2, 2n-zn2. Grouped bars illustrate the 64B packet throughput speedup ratio for 2- and 4-core multi-threaded DPDK 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 are stored in build logs from FD.io dpdk performance job 2n-icx[291], build logs from FD.io dpdk performance job 3n-icx[292], build logs from FD.io dpdk performance job 2n-skx[293], build logs from FD.io dpdk performance job 3n-skx[294], build logs from FD.io dpdk performance job 2n-clx[295], build logs from FD.io dpdk performance job 2n-zn2[296], build logs from FD.io dpdk performance job 3n-alt[297], build logs from FD.io dpdk performance job 3n-tsh[298], build logs from FD.io dpdk performance job 2n-tx2[299] with RF result files csit-vpp-perf-2206-*.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.

---

[291] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-icx
[292] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-icx
[293] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-skx
[294] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-skx
[295] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-clx
[296] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-zn2
[297] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-alt
[298] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-tsh
[299] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-tx2

3.4.1 2n-icx-xxv710

CSIT source code for the test cases used for plots can be found in CSIT git repository[300].

[300] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-base

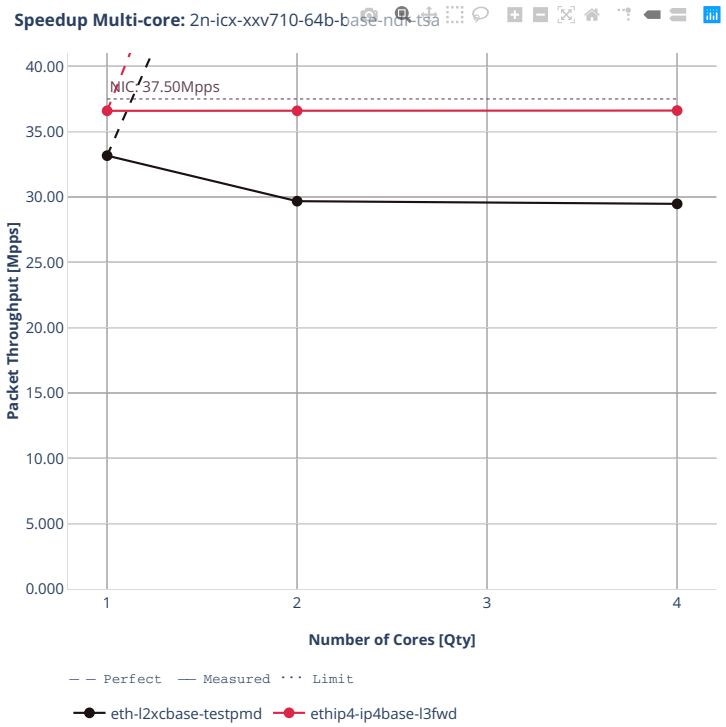
Speedup Multi-core: 2n-icx-xxv710-64b-base-ndr-tsa
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

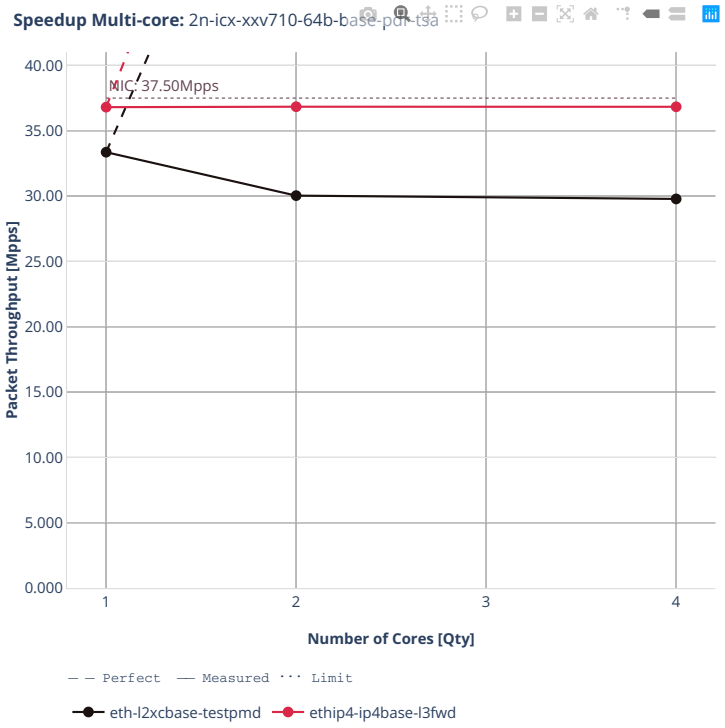
Speedup Multi-core: 2n-icx-xxv710-64b-base-pdr-tsa
MC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

3.4.2 3n-icx-xxv710

CSIT source code for the test cases used for plots can be found in CSIT git repository[301].

[301] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-base

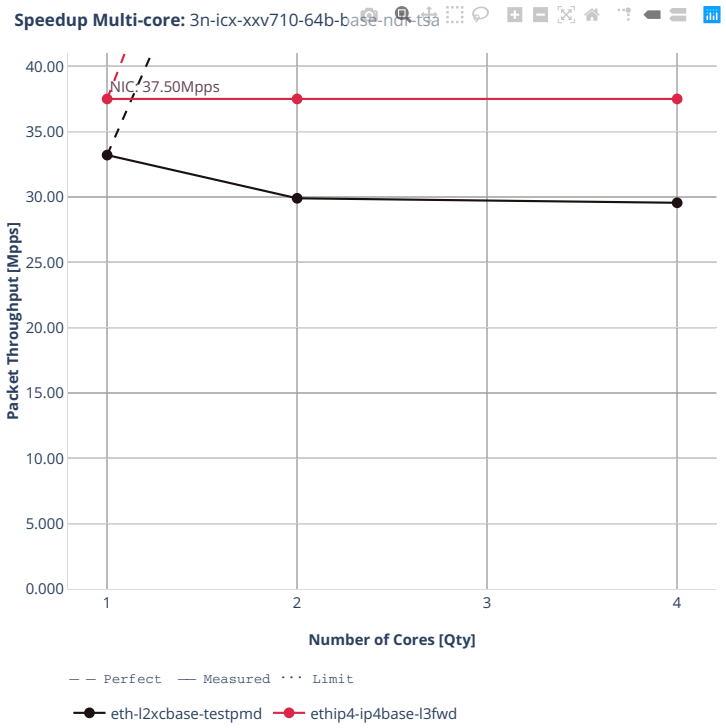
Speedup Multi-core: 3n-icx-xxv710-64b-base-ndr-tsa
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

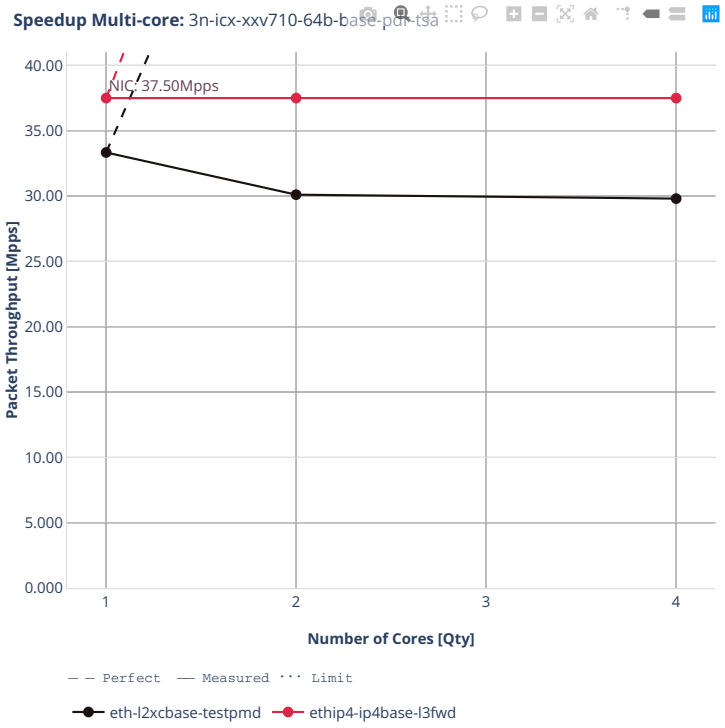
Speedup Multi-core: 3n-icx-xxv710-64b-base-pdr-tsa
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

3.4.3 2n-skx-xxv710

CSIT source code for the test cases used for plots can be found in CSIT git repository[302].

[302] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-base

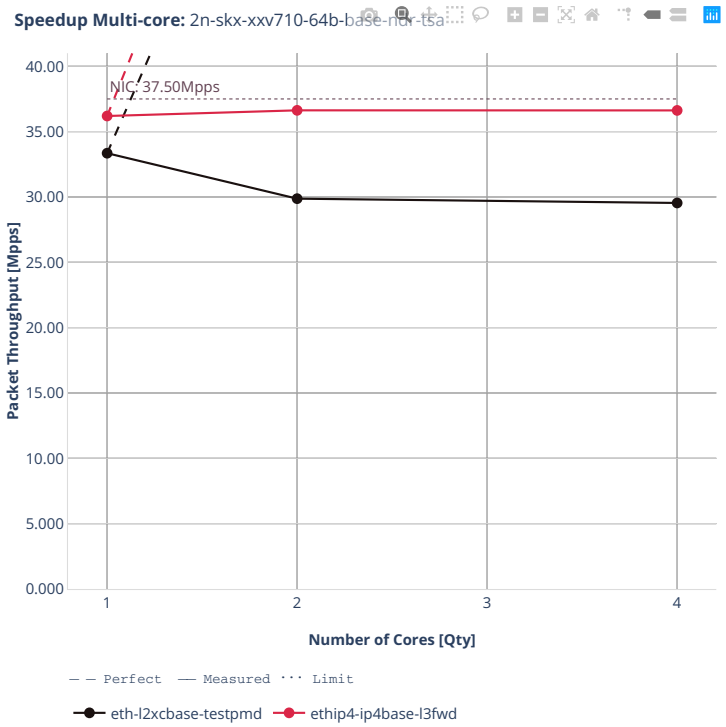
Speedup Multi-core: 2n-skx-xxv710-64b-base-ndr-tsa
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

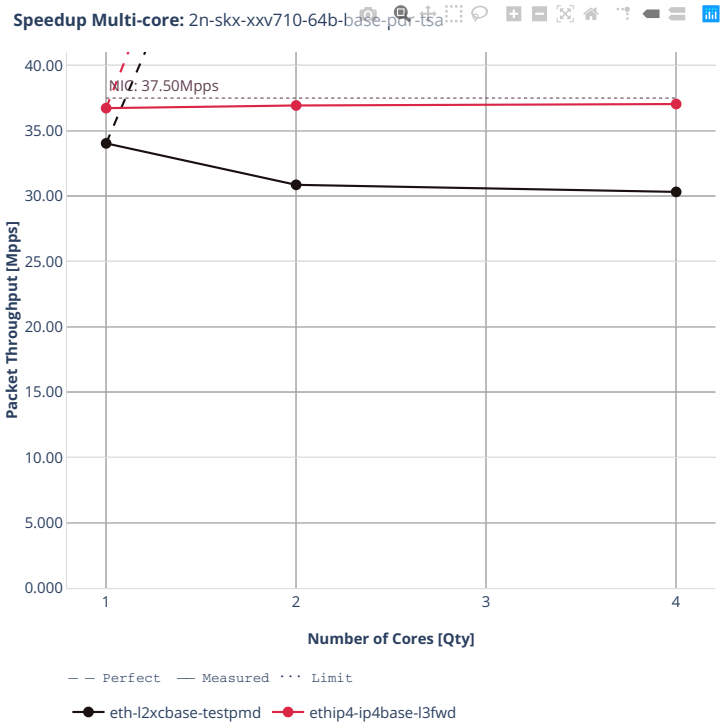
Speedup Multi-core: 2n-skx-xxv710-64b-base-pdr-tsa
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
37.50Mpps
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

3.4.4 2n-skx-x710

CSIT source code for the test cases used for plots can be found in CSIT git repository[303].

[303] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-base

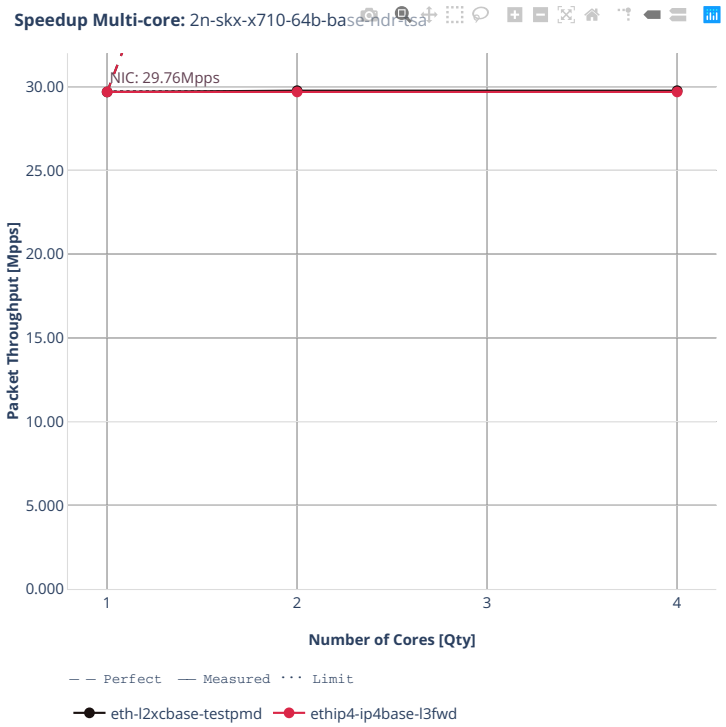
Speedup Multi-core: 2n-skx-x710-64b-base-ndrpsa
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

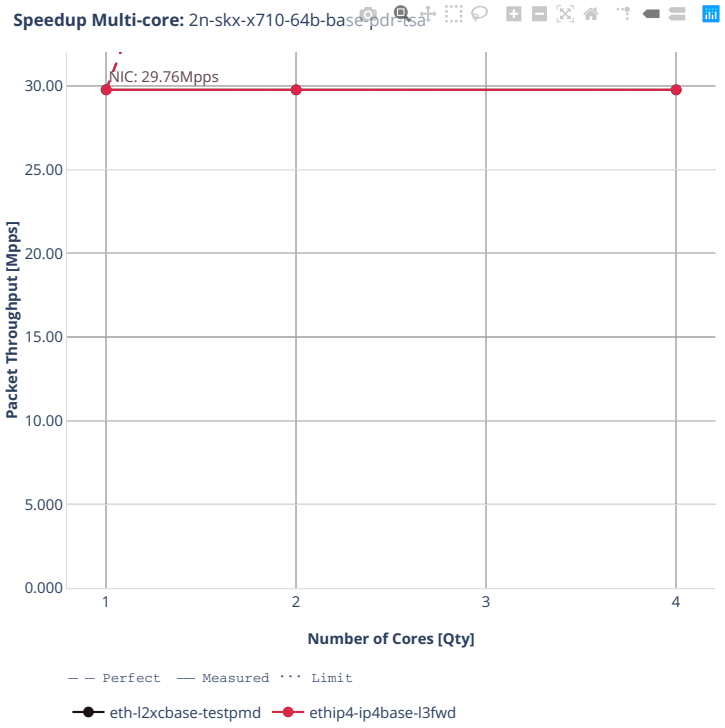
Speedup Multi-core: 2n-skx-x710-64b-base-pdr-tsa
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

3.4.5 3n-skx-xxv710

CSIT source code for the test cases used for plots can be found in CSIT git repository[304].

[304] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-base

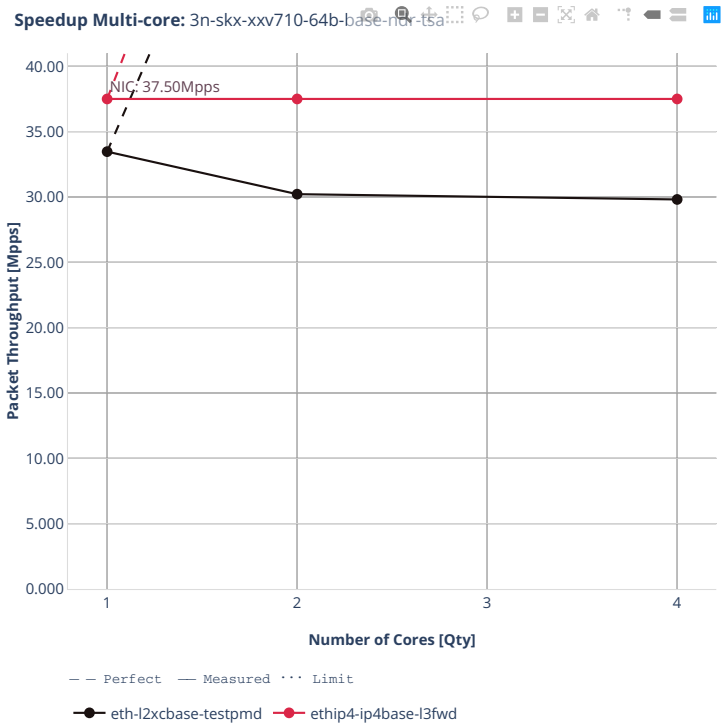
Speedup Multi-core: 3n-skx-xxv710-64b-base-ndr-tsa
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

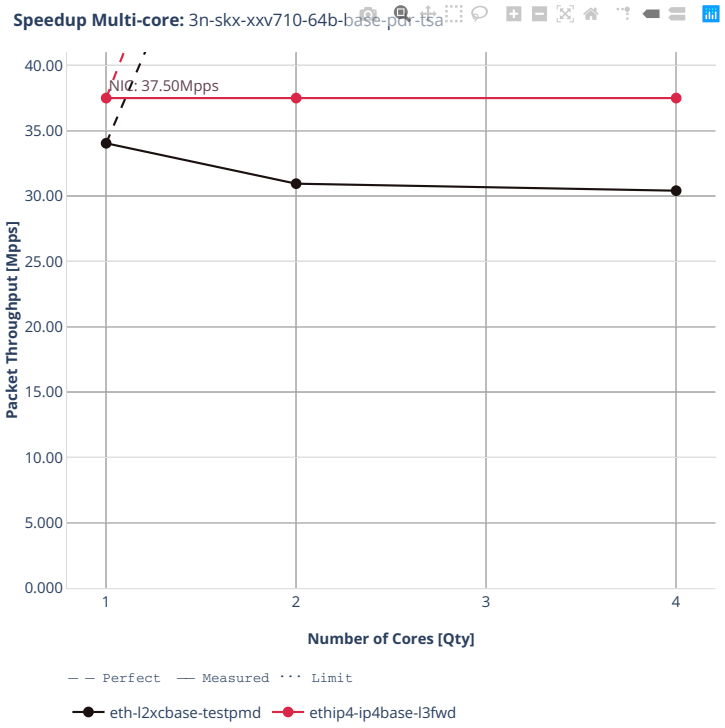
Speedup Multi-core: 3n-skx-xxv710-64b-base-pdr-tsa
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

3.4.6 3n-skx-x710

CSIT source code for the test cases used for plots can be found in CSIT git repository[305].

[305] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-base

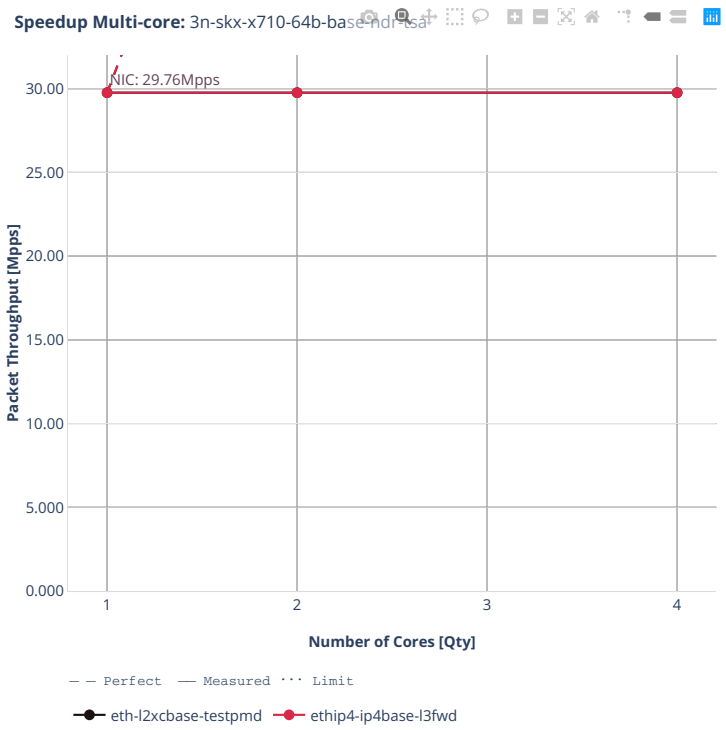
Speedup Multi-core: 3n-skx-x710-64b-base-ndrpsa
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

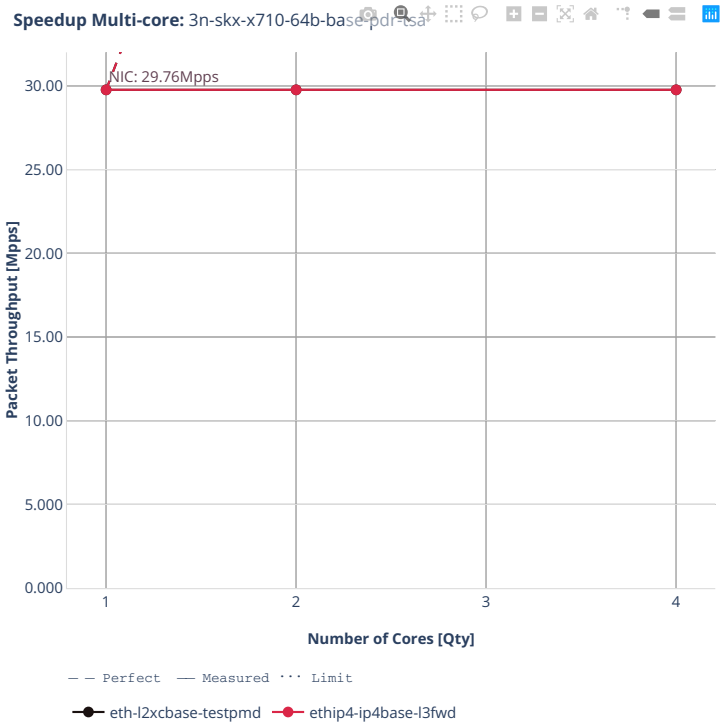
Speedup Multi-core: 3n-skx-x710-64b-base-pdr-tsa
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

3.4.7 2n-clx-xxv710

CSIT source code for the test cases used for plots can be found in CSIT git repository[306].

[306] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-base

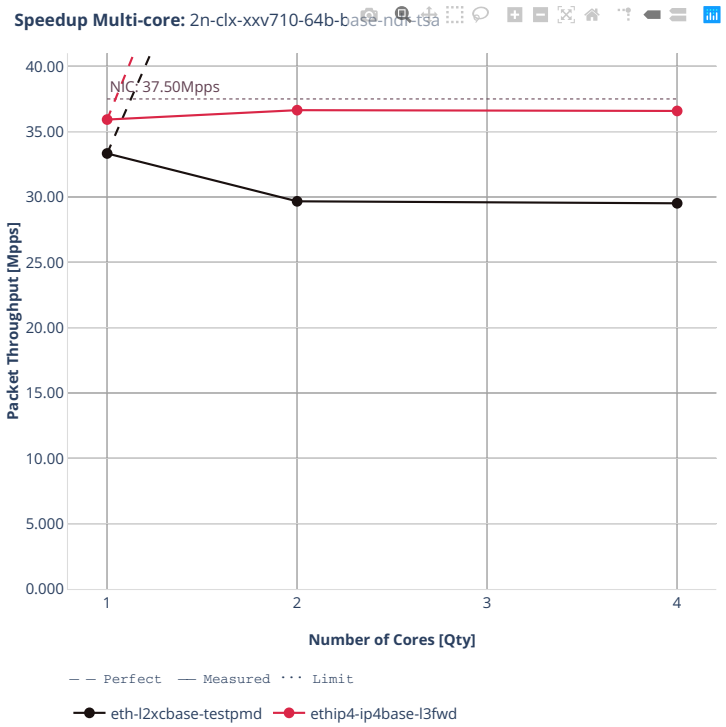
Speedup Multi-core: 2n-clx-xxv710-64b-base-ndr-tsa
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd


Speedup Multi-core: 2n-clx-xxv710-64b-base-pdr-tsa
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

3.4.8 2n-clx-x710

CSIT source code for the test cases used for plots can be found in CSIT git repository[307].

[307] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-base

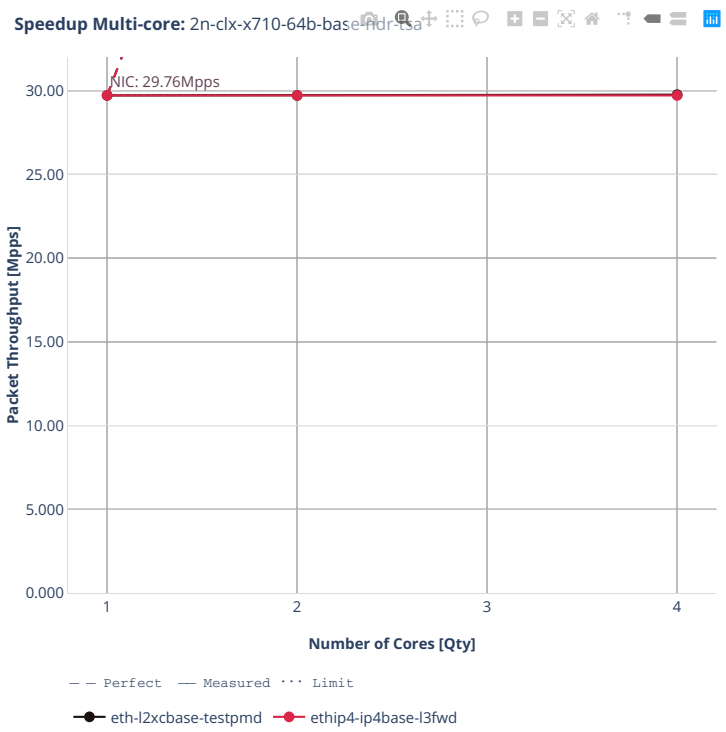
Speedup Multi-core: 2n-clx-x710-64b-base-ndr-tsa
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

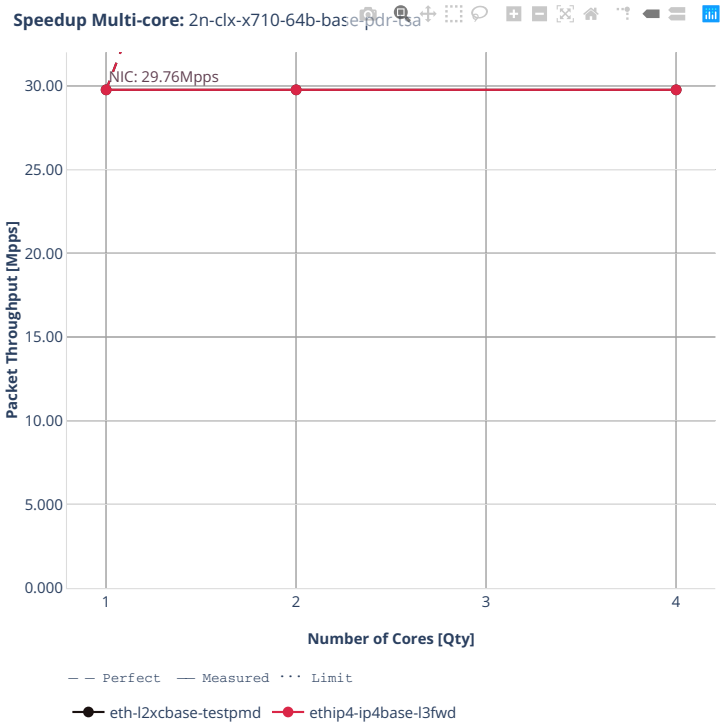
Speedup Multi-core: 2n-clx-x710-64b-base-pdr-tsa
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

3.4.9 2n-zn2-xxv710

CSIT source code for the test cases used for plots can be found in CSIT git repository[308].

[308] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-base

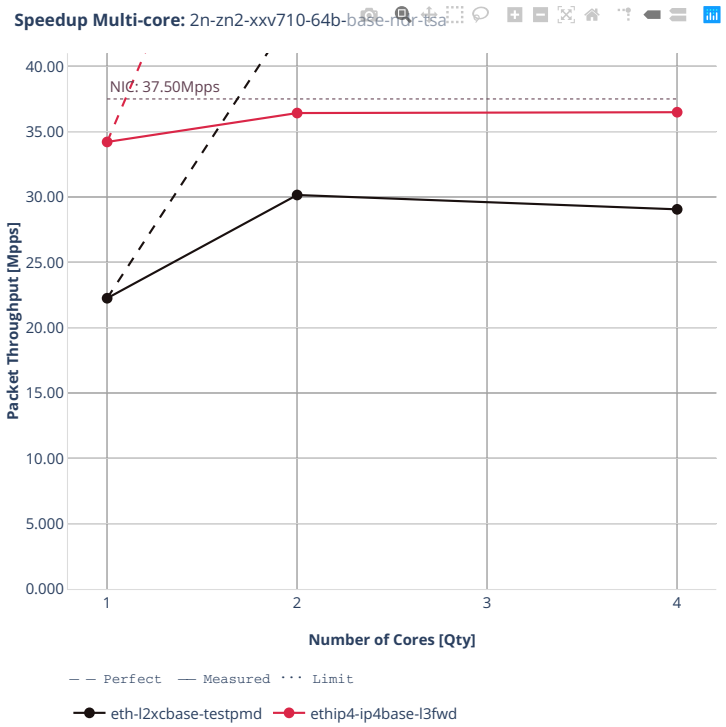
Speedup Multi-core: 2n-zn2-xxv710-64b-base-mrr-tsa
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

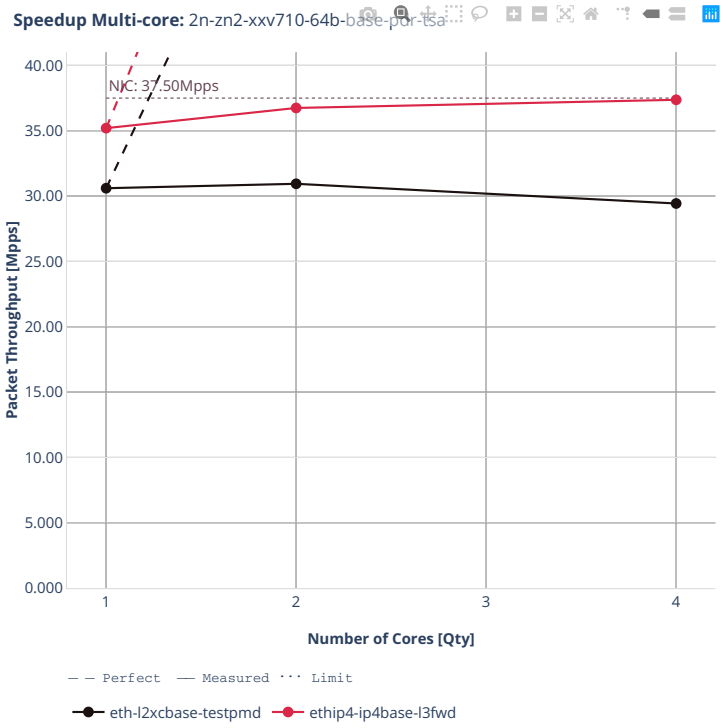
Speedup Multi-core: 2n-zn2-xxv710-64b-base-por-tsa
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

3.4.10 2n-zn2-x710

CSIT source code for the test cases used for plots can be found in CSIT git repository[309].

[309] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-base

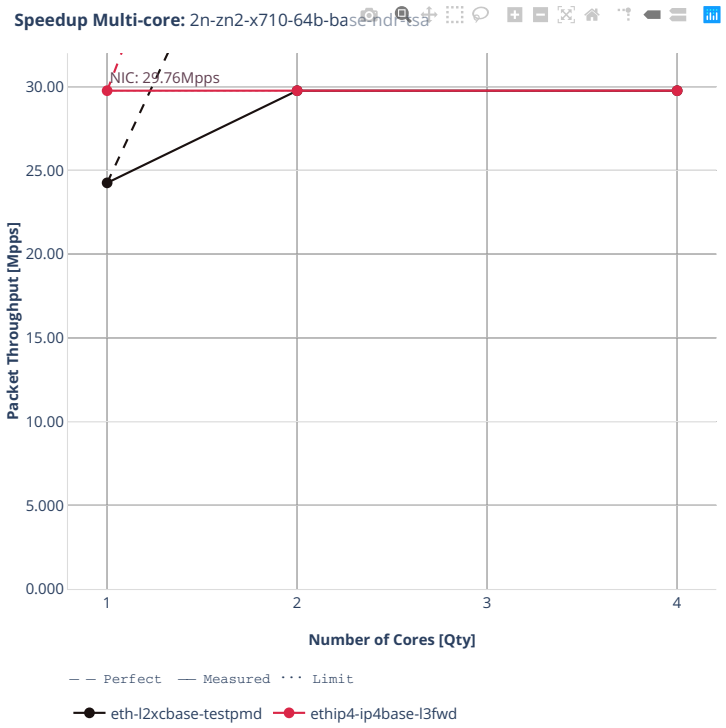
Speedup Multi-core: 2n-zn2-x710-64b-base-ndrpsa
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

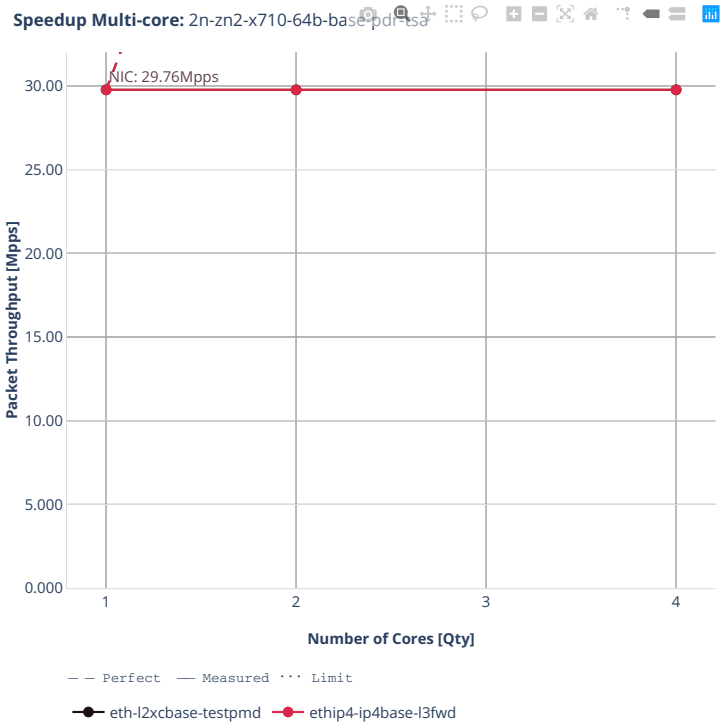
Speedup Multi-core: 2n-zn2-x710-64b-base-pdr-tsa
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

3.4.11 3n-alt-xl710

CSIT source code for the test cases used for plots can be found in CSIT git repository[310].

[310] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-base

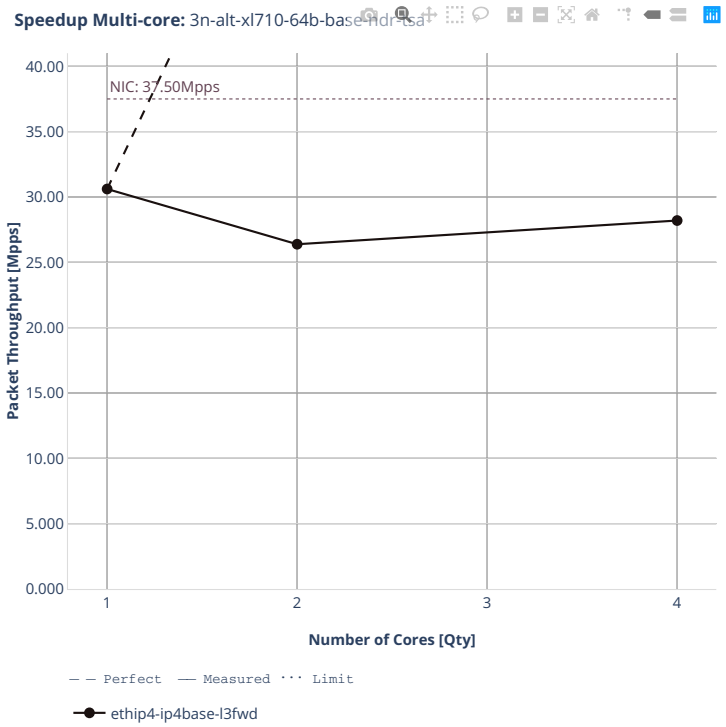
Speedup Multi-core: 3n-alt-xl710-64b-base-ndr-lsa
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base-l3fwd

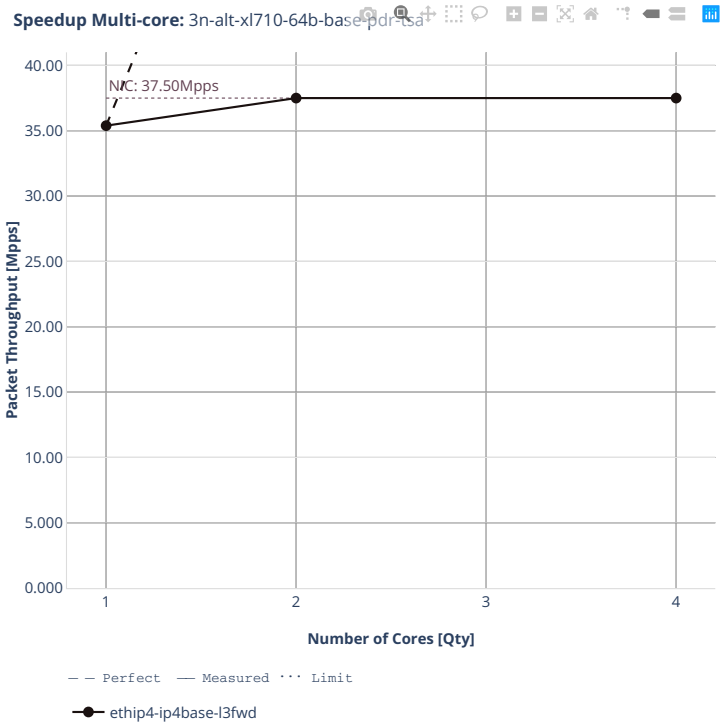
Speedup Multi-core: 3n-alt-xl710-64b-base-pdr-lsa
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
ethip4-ip4base-l3fwd

3.4.12 3n-tsh-x520

CSIT source code for the test cases used for plots can be found in CSIT git repository[311].

[311] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-base

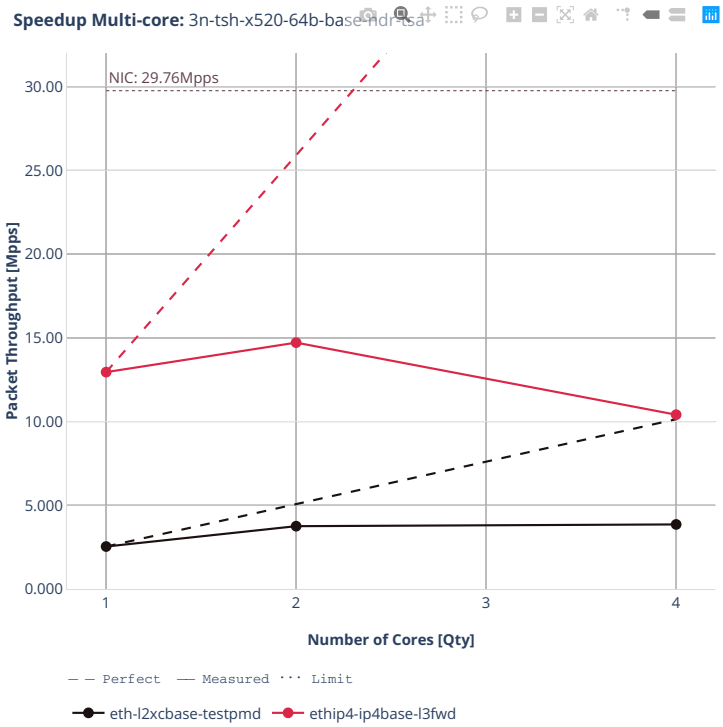
Speedup Multi-core: 3n-tsh-x520-64b-base-ndr-lsa
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

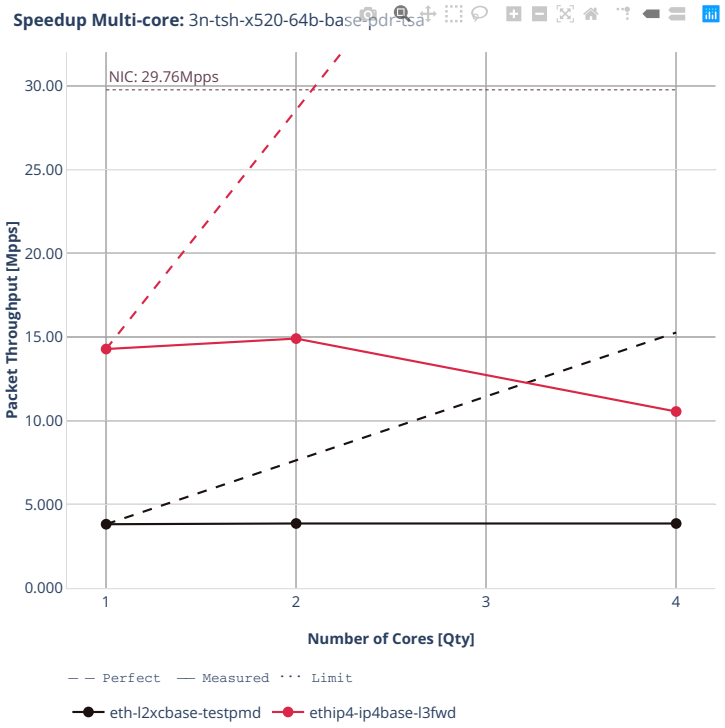
Speedup Multi-core: 3n-tsh-x520-64b-base-pdr-lsa
NIC: 29.76Mpps
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

3.4.13 2n-tx2-xl710

CSIT source code for the test cases used for plots can be found in CSIT git repository[312].

[312] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-base

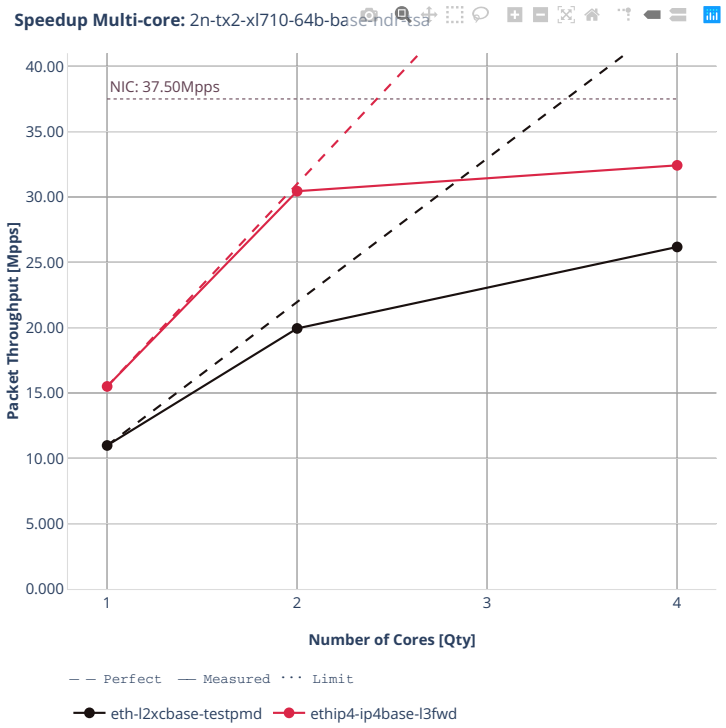
Speedup Multi-core: 2n-tx2-xl710-64b-base-ndrpsa
NIC: 37.50Mpps
Packet Throughput [Mpps]
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

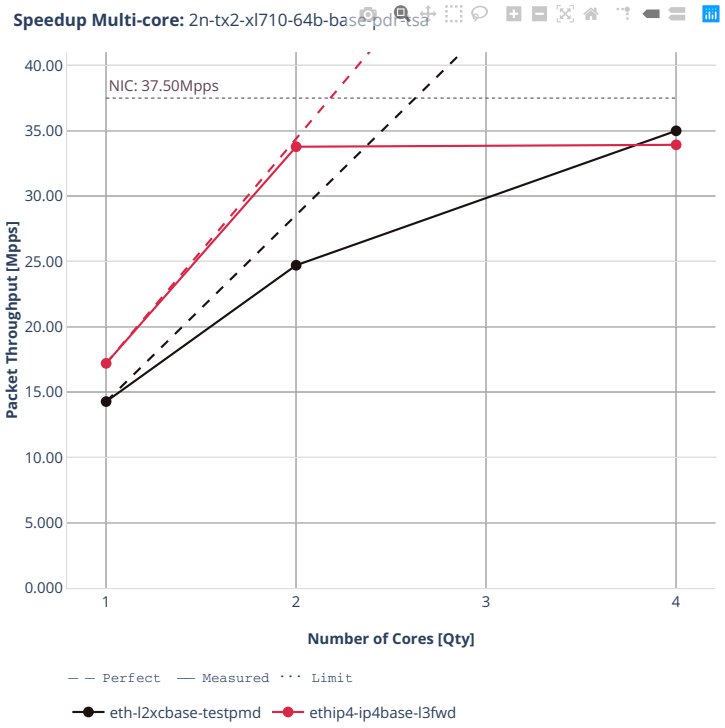
Speedup Multi-core: 2n-tx2-xl710-64b-base-pdr-tsa
NIC: 37.50Mpps
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.000
0.000
Packet Throughput [Mpps]
1
2
3
4
Number of Cores [Qty]
Perfect
Measured
Limit
eth-l2xcbase-testpmd
ethip4-ip4base-l3fwd

3.5 Packet Latency

DPDK Testpmd and L3fwd latency results are generated based on the test data obtained from CSIT-2206 NDR-PDR throughput tests executed across physical testbeds hosted in LF FD.io labs: 2n-icx, 3n-icx, 2n-skx, 3n-skx, 2n-clx, 2n-zn2, 3n-alt, 3n-tsh, 2n-tx2.

Latency by percentile distribution plots are used to show packet latency percentiles at different packet rate load levels: i) No-Load latency streams only, ii) Low-Load at 10% PDR, iii) Mid-Load at 50% PDR and iv) High-Load at 90% PDR.

For more details, see *Packet Latency* (page 45).

Additional information about graph data:

1. **Graph Title**: describes tested DUT packet path.
2. **X-axis Labels**: percentile of packets.
3. **Y-axis Labels**: measured one-way packet latency values in [uSec].
4. **Graph Legend**: list of latency tests at different packet rate load level.
5. **Hover Information**: packet rate load level, stream direction (East-West, West-East), percentile, one-way latency.

Note: Test results are stored in build logs from FD.io dpdk performance job 2n-icx[313], build logs from FD.io dpdk performance job 3n-icx[314], build logs from FD.io dpdk performance job 2n-skx[315], build logs from FD.io dpdk performance job 3n-skx[316], build logs from FD.io dpdk performance job 2n-clx[317], build logs from FD.io dpdk performance job 2n-zn2[318], build logs from FD.io dpdk performance job 3n-alt[319], build logs from FD.io dpdk performance job 3n-tsh[320] and build logs from FD.io dpdk performance job 2n-tx2[321] with RF result files csit-dpdk-perf-2206-*.zip archived here.

---

[313] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-icx
[314] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-icx
[315] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-skx
[316] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-skx
[317] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-clx
[318] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-zn2
[319] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-alt
[320] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-tsh
[321] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-tx2

3.5.1 2n-icx-xxv710

CSIT source code for the test cases used for plots can be found in CSIT git repository[322].

[322] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

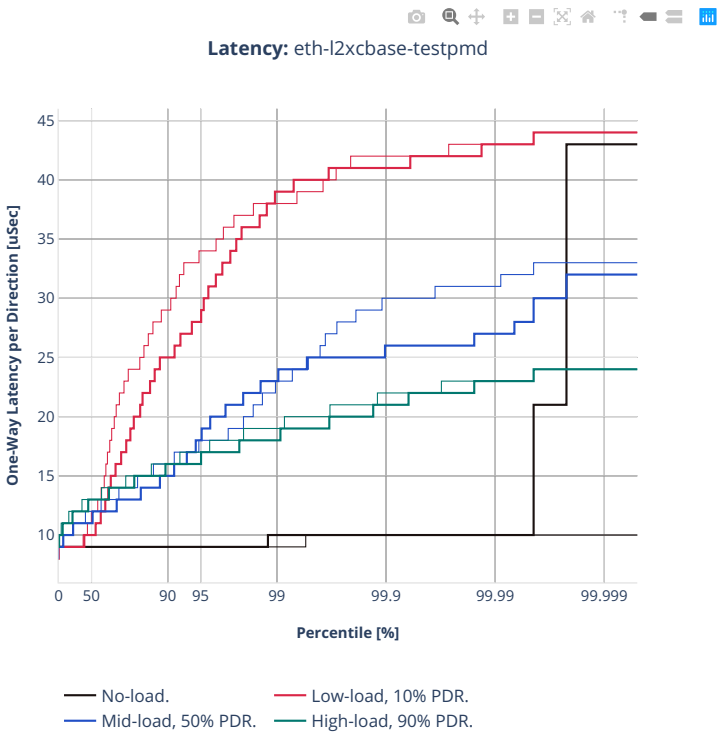

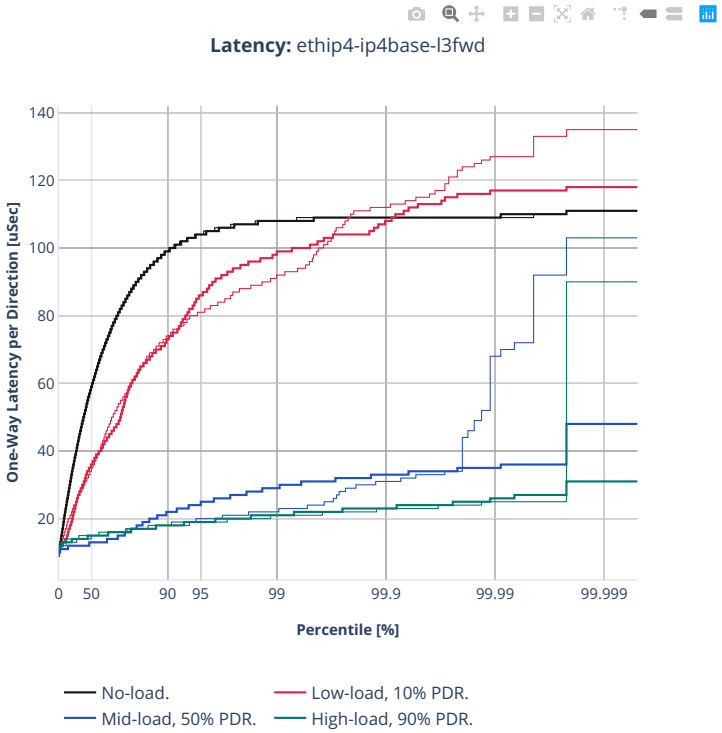
Latency: ethip4-ip4base-l3fwd
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3.5.2 3n-icx-xxv710

CSIT source code for the test cases used for plots can be found in CSIT git repository[323].

[323] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

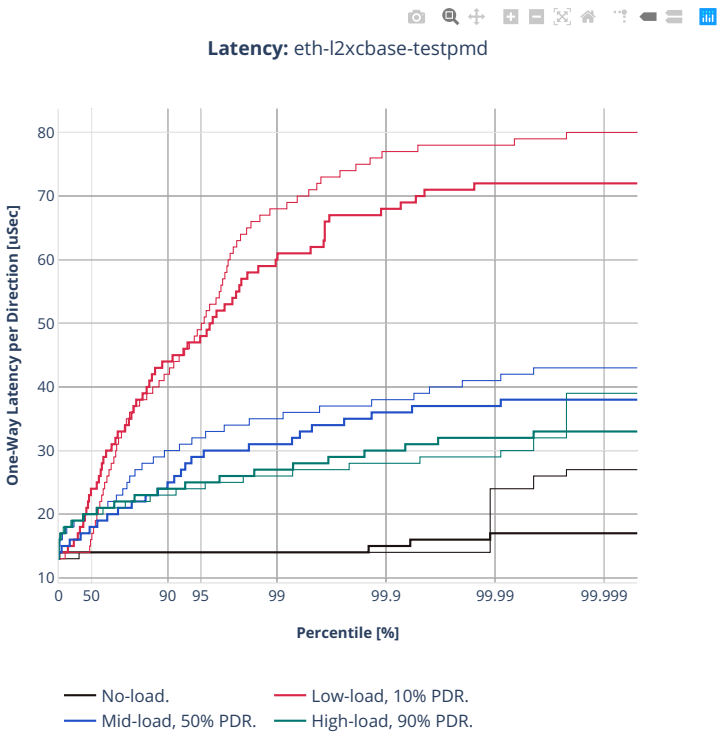

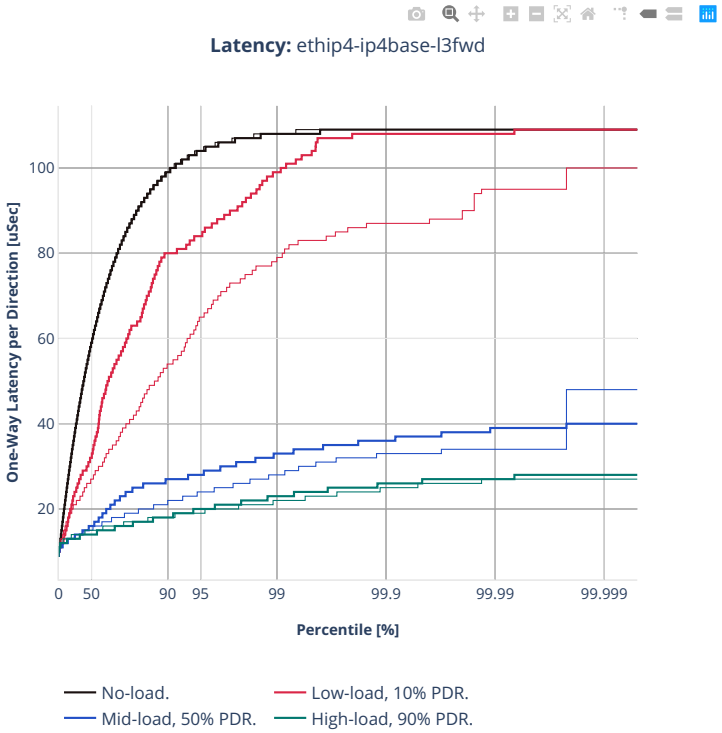
Latency: ethip4-ip4base-l3fwd
One-Way Latency per Direction [uSec]
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3.5.3 2n-skx-xxv710

CSIT source code for the test cases used for plots can be found in CSIT git repository[324].

[324] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

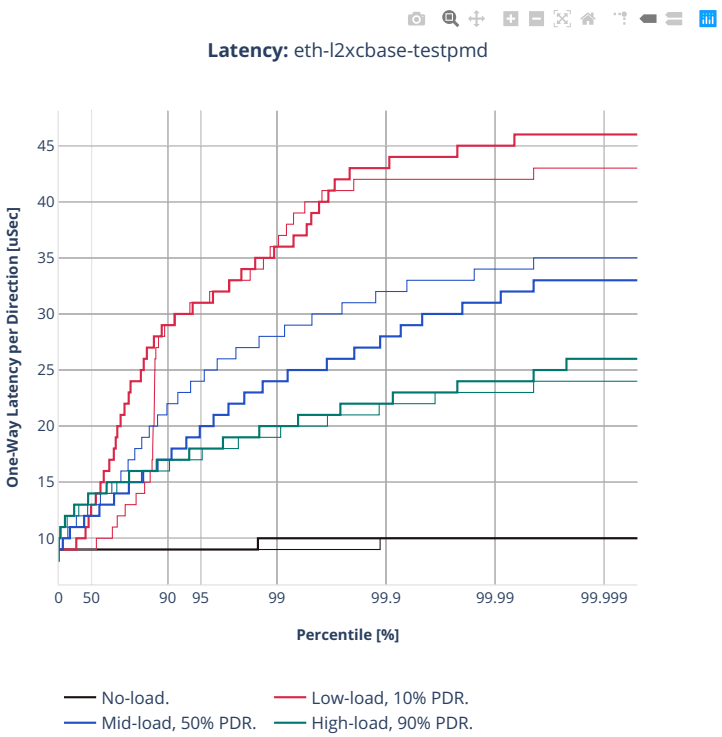

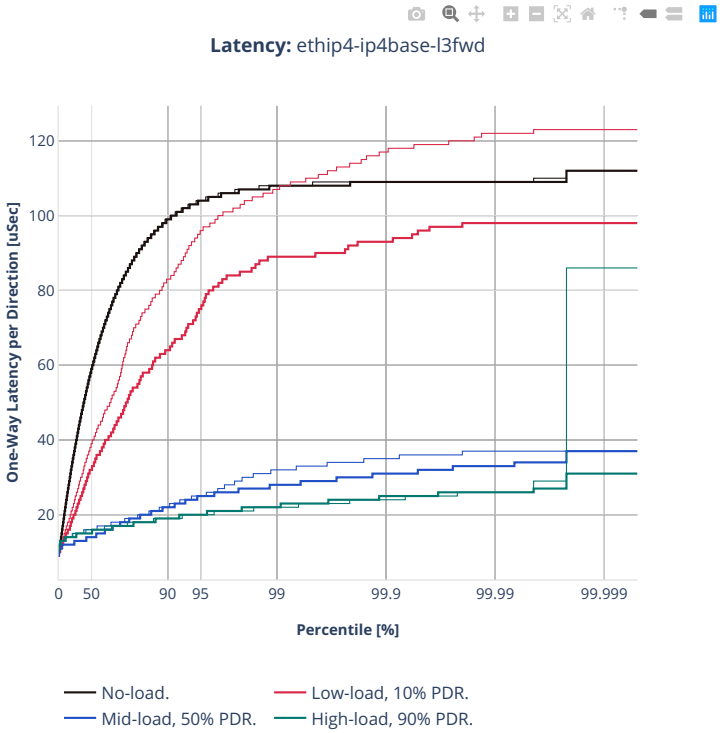
Latency: ethip4-ip4base-l3fwd
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3.5.4 3n-skx-xxv710

CSIT source code for the test cases used for plots can be found in CSIT git repository[325].

[325] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

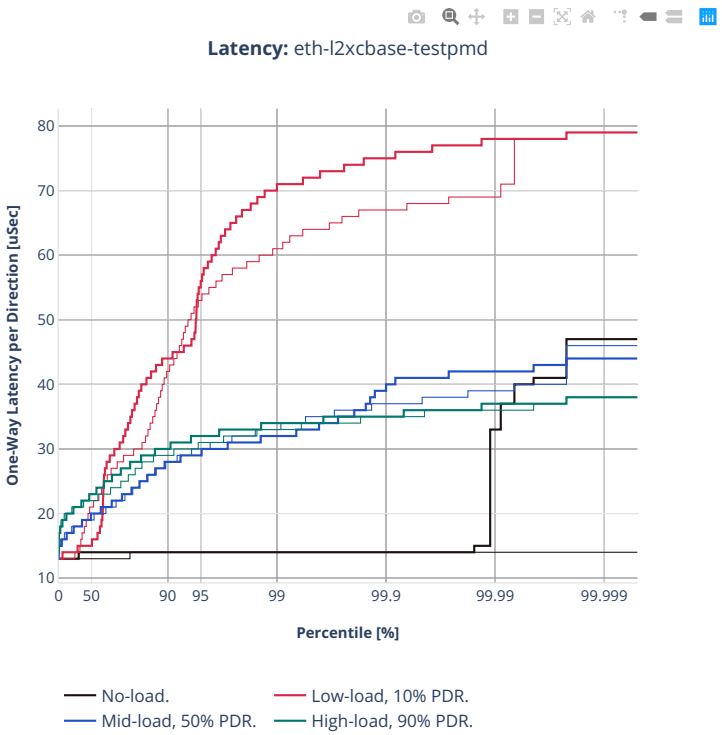

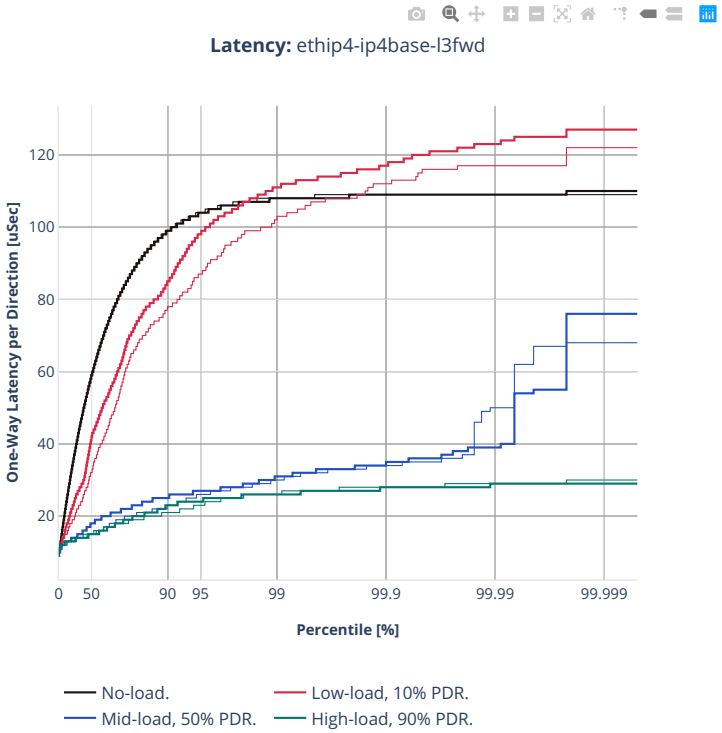
Latency: ethip4-ip4base-l3fwd
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3.5.5 2n-clx-xxv710

CSIT source code for the test cases used for plots can be found in CSIT git repository[326].

[326] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

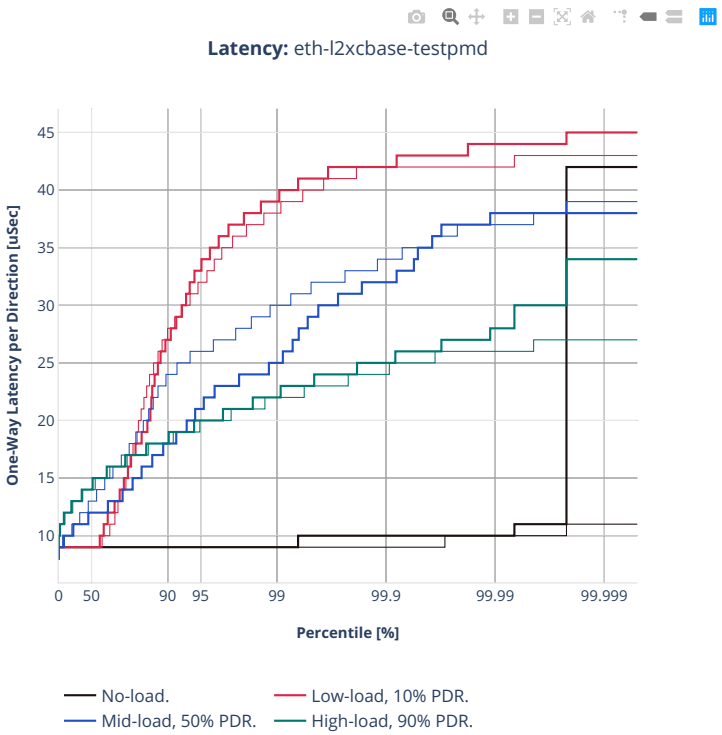

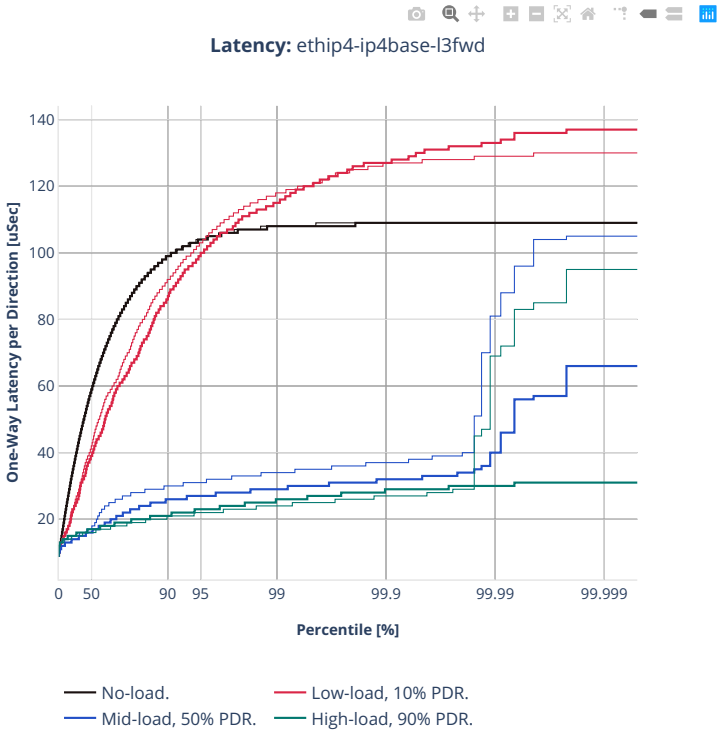
Latency: ethip4-ip4base-l3fwd
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3.5.6 2n-zn2-xxv710

CSIT source code for the test cases used for plots can be found in CSIT git repository[327].

[327] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-2t1c-base

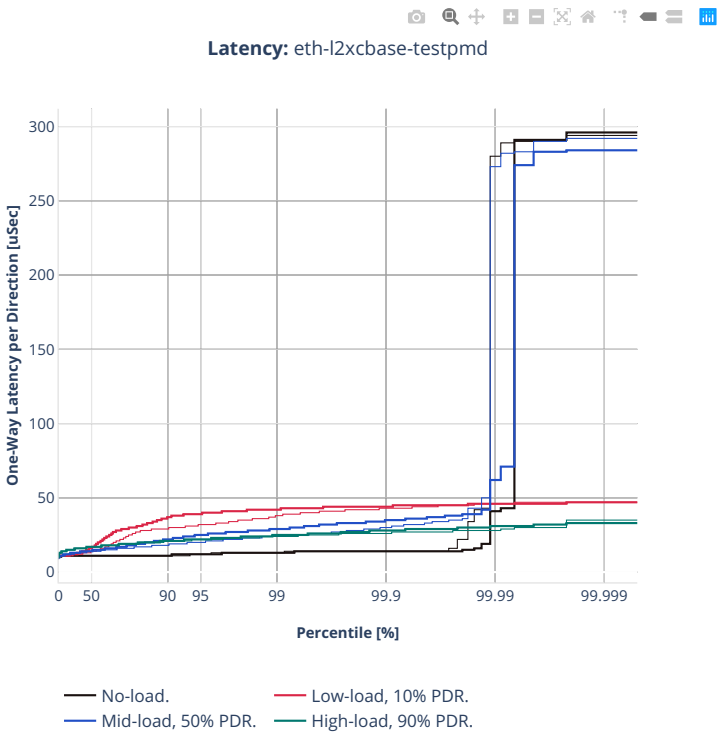

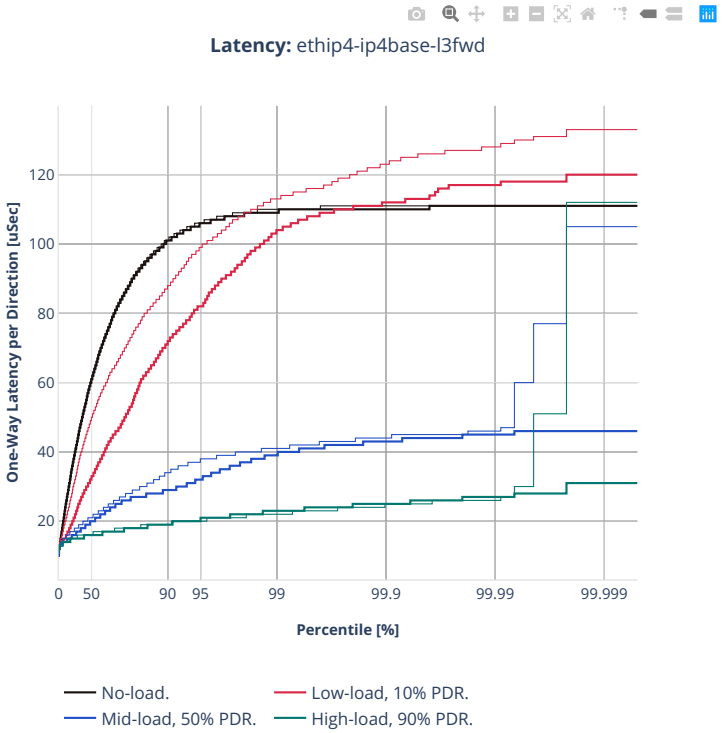
Latency: ethip4-ip4base-l3fwd
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3.5.7 3n-alt-xl710

CSIT source code for the test cases used for plots can be found in CSIT git repository[328].

[328] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-1t1c-base

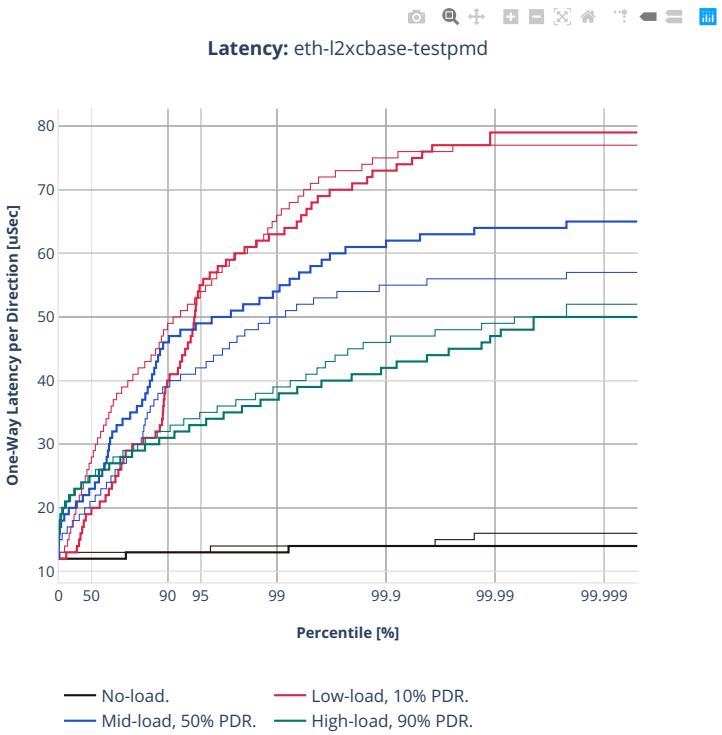

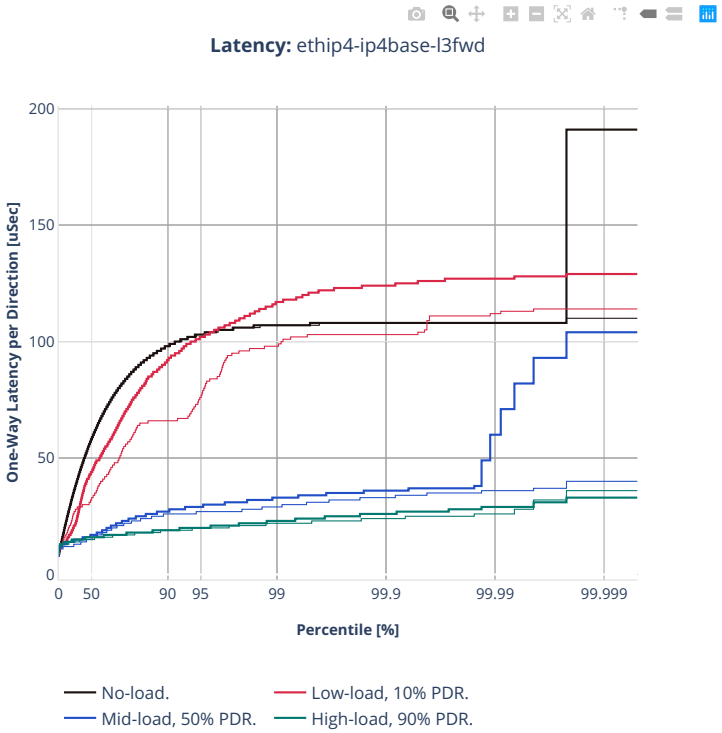
Latency: ethip4-ip4base-l3fwd
One-Way Latency per Direction [uSec]
200
150
100
50
0
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3.5.8 3n-tsh-x520

CSIT source code for the test cases used for plots can be found in CSIT git repository[329].

[329] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-1t1c-base

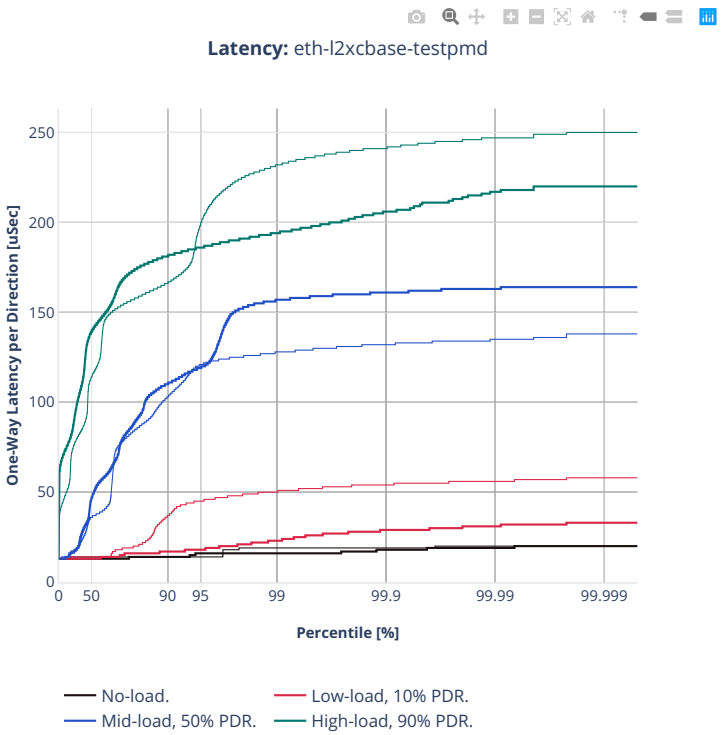

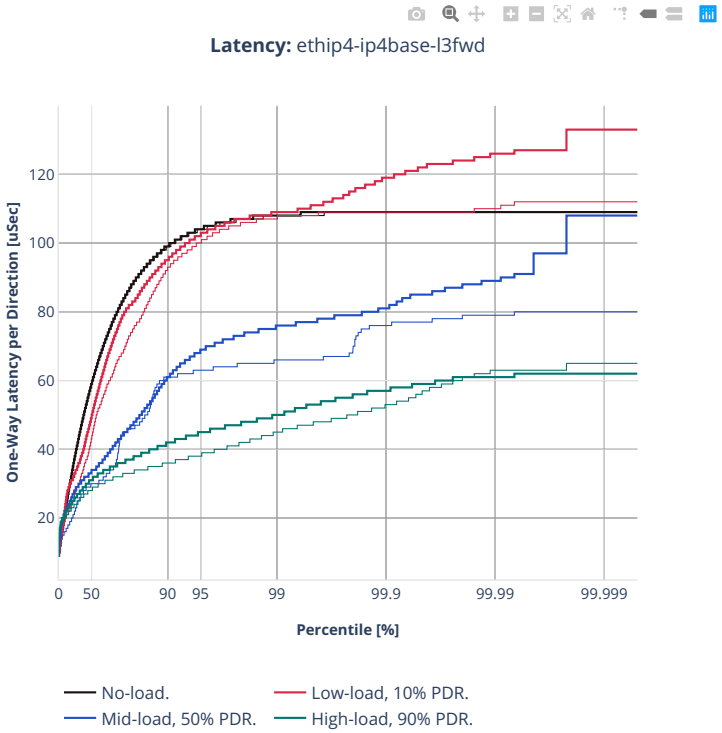
Latency: ethip4-ip4base-l3fwd
One-Way Latency per Direction [uSec]
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3.5.9 2n-tx2-xl710

CSIT source code for the test cases used for plots can be found in CSIT git repository[330].

[330] https://git.fd.io/csit/tree/tests/dpdk/perf?h=rls2206

64b-1t1c-base

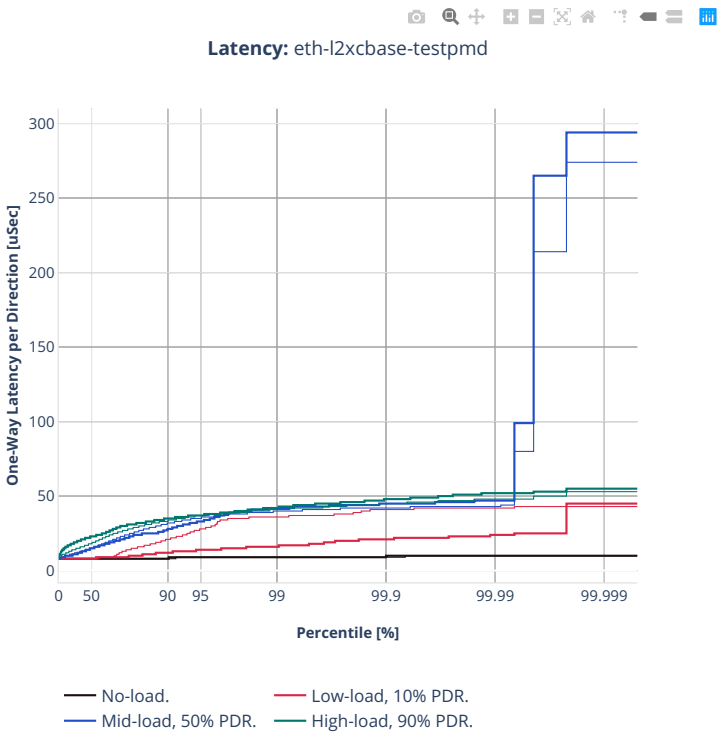

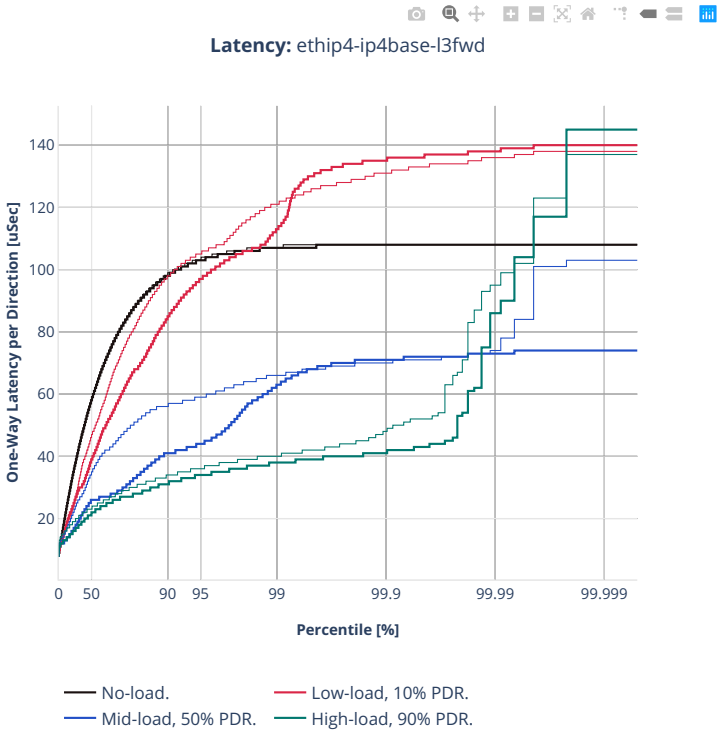
Latency: ethip4-ip4base-l3fwd
One-Way Latency per Direction [uSec]
140
120
100
80
60
40
20
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

3.6 Comparisons

3.6.1 Current vs Previous Release

Relative comparison of DPDK Testpmd and L3fwd packet throughput (NDR, PDR and MRR) between DPDK-22.03 and DPDK-21.11 (measured for CSIT-2206 and CSIT-2202 respectively) is calculated from results of tests running on 3n-skx, 2n-skx, 2n-clx, 2n-zn2 testbeds in 1-core and 2-core configurations.

Listed mean and standard deviation values are computed based on a series of the same tests executed against respective DPDK releases to verify test results repeatability, with percentage change calculated for mean values.

Note: Test results are stored in build logs from FD.io dpdk performance job 3n-icx[331], build logs from FD.io dpdk performance job 2n-icx[332], build logs from FD.io dpdk performance job 3n-skx[333], build logs from FD.io dpdk performance job 2n-skx[334], build logs from FD.io dpdk performance job 2n-clx[335], build logs from FD.io dpdk performance job 2n-zn2[336], build logs from FD.io dpdk performance job 3n-tsh[337], build logs from FD.io dpdk performance job 2n-tx2[338], with RF result files csit-dpdk-perf-2206-*.zip archived here.

3n-icx-xxv710

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

[331] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-icx
[332] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-icx
[333] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-skx
[334] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-skx
[335] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-clx
[336] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-zn2
[337] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-tsh
[338] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-tx2

2n-icx-xxv710

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

3n-skx-xxv710

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

2n-skx-xxv710

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

2n-clx-xxv710

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

2n-zn2-xxv710

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

3n-tsh-x520

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

2n-tx2-xl710

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

3.6.2 2n-Icx vs 2n-Clx Testbeds

Relative comparison of DPDK-22.03 Testpmd and L3fwd packet throughput (NDR, PDR) is calculated for the same tests executed on 2-Node Cascadelake (2n-clx) and 2-Node Ice Lake (2n-icx) physical testbed types, in 1-core, 2-core and 4-core configurations.

Note: Test results are stored in build logs from FD.io dpdk performance job 2n-clx[339] and build logs from FD.io dpdk performance job 2n-icx[340] with RF result files csit-dpdk-perf-2206-*.zip archived here.

NDR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII NDR comparison
- CSV NDR comparison

PDR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII PDR comparison
- CSV PDR comparison

[339] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-clx
[340] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-icx

3.6.3 2n-Icx vs 2n-Skx Testbeds

Relative comparison of DPDK-22.03 Testpmd and L3fwd packet throughput (NDR, PDR) is calculated for the same tests executed on 2-Node Cascadelake (2n-skx) and 2-Node Ice Lake (2n-icx) physical testbed types, in 1-core, 2-core and 4-core configurations.

Note: Test results are stored in build logs from FD.io dpdk performance job 2n-skx[341] and build logs from FD.io dpdk performance job 2n-icx[342] with RF result files csit-dpdk-perf-2206-*.zip archived here.

NDR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII NDR comparison
- CSV NDR comparison

PDR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII PDR comparison
- CSV PDR comparison

3.6.4 3n-Icx vs 3n-Skx Testbeds

Relative comparison of DPDK-22.03 Testpmd and L3fwd packet throughput (NDR, PDR) is calculated for the same tests executed on 3-Node Skylake (3n-skx) and 3-Node Ice Lake (3n-icx) physical testbed types, in 1-core, 2-core and 4-core configurations.

Note: Test results are stored in build logs from FD.io dpdk performance job 3n-skx[343] and build logs from FD.io dpdk performance job 3n-icx[344] with RF result files csit-dpdk-perf-2206-*.zip archived here.

NDR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII NDR comparison
- CSV NDR comparison

[341] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-skx
[342] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-icx
[343] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-skx
[344] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-icx

PDR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII PDR comparison
- CSV PDR comparison

3.6.5 2n-Clx vs 2n-Skx Testbeds

Relative comparison of DPDK-22.03 Testpmd and L3fwd packet throughput (NDR, PDR) is calculated for the same tests executed on 2-Node Skylake (2n-skx) and 2-Node Cascade Lake (2n-clx) physical testbed types, in 1-core, 2-core and 4-core configurations.

Note: Test results are stored in build logs from FD.io dpdk performance job 2n-skx[345] and build logs from FD.io dpdk performance job 2n-clx[346] with RF result files csit-dpdk-perf-2206-*.zip archived here.

NDR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII NDR comparison
- CSV NDR comparison

PDR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII PDR comparison
- CSV PDR comparison

3.6.6 3n-Skx vs 2n-Skx Testbeds

Relative comparison of DPDK-22.03 Testpmd and L3fwd packet throughput (NDR, PDR) 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 are stored in build logs from FD.io dpdk performance job 3n-skx[347] and build logs from FD.io dpdk performance job 2n-skx[348] with RF result files csit-dpdk-perf-2206-*.zip archived here.

---

[345] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-skx
[346] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-clx
[347] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-3n-skx
[348] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-dpdk-perf-report-iterative-2206-2n-skx

NDR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII NDR comparison
- CSV NDR comparison

PDR Comparison

Comparison tables in ASCII and CSV formats:

- ASCII PDR comparison
- CSV PDR comparison

3.7 Throughput Trending

In addition to reporting throughput comparison between DPDK releases, CSIT provides regular performance trending for DPDK release branches:

1. Performance Dashboard[349]: per DPDK test case throughput trend, trend compliance and summary of detected anomalies.
2. Trending Methodology[350]: throughput test metrics, trend calculations and anomaly classification (progression, regression).
3. DPDK Trendline Graphs[351]: weekly DPDK Testpmd and L3fwd MRR throughput measurements against the trendline with anomaly highlights and associated CSIT test jobs.

[349] https://s3-docs.fd.io/csit/master/trending/introduction/dashboard.html
[350] https://s3-docs.fd.io/csit/master/trending/methodology/index.html
[351] https://s3-docs.fd.io/csit/master/trending/trending/dpdk.html

3.8 Test Environment

3.8.1 Environment Versioning

CSIT test environment versioning has been introduced to track modifications of the test environment.

Any benchmark anomalies (progressions, regressions) between releases of a DUT application (e.g. VPP, DPDK), are determined by testing it in the same test environment, to avoid test environment changes clouding the picture. To beter distinguish impact of test environment changes, we also execute tests without any SUT (just with TRex TG sending packets over a link looping back to TG).

A mirror approach is introduced to determine benchmarking anomalies due to the test environment change. This is achieved by testing the same DUT application version between releases of CSIT test system. This works under the assumption that the behaviour of the DUT is deterministic under the test conditions.

CSIT test environment versioning scheme ensures integrity of all the test system components, including their HW revisions, compiled SW code versions and SW source code, within a specific CSIT version. Components included in the CSIT environment versioning include:

- **HW** Server hardware firmware and BIOS (motherboard, processsor, NIC(s), accelerator card(s)), tracked in CSIT branch in `./docs/lab/<server_platform_name>_hw_bios_cfg.md`, e.g. Xeon Skylake servers[352].
- **Linux** Server Linux OS version and configuration, tracked in CSIT Reports in SUT Settings[353] and Pre-Test Server Calibration[354].
- **TRex** TRex Traffic Generator version, drivers and configuration tracked in TG Settings[355].
- **CSIT** CSIT framework code tracked in CSIT release branches.

Following is the list of CSIT versions to date:

- Ver. 1 associated with CSIT rls1908 branch (HW[356], Linux[357], TRex[358], CSIT[359]).
- Ver. 2 associated with CSIT rls2001 branch (HW[360], Linux[361], TRex[362], CSIT[363]).
- Ver. 4 associated with CSIT rls2005 branch (HW[364], Linux[365], TRex[366], CSIT[367]).
- Ver. 5 associated with CSIT rls2009 branch (HW[368], Linux[369], TRex[370], CSIT[371]).
  - The main change is TRex data-plane core resource adjustments: increase from 7 to 8 cores and pinning cores to interfaces[372] for better TRex performance with symmetric traffic profiles.

[352] https://git.fd.io/csit/tree/docs/lab/testbeds_sm_skx_hw_bios_cfg.md#n556
[353] https://s3-docs.fd.io/csit/master/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[354] https://s3-docs.fd.io/csit/master/report/vpp_performance_tests/test_environment.html#id21
[355] https://s3-docs.fd.io/csit/master/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[356] https://git.fd.io/csit/tree/docs/lab?h=rls1908
[357] https://docs.fd.io/csit/rls1908/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[358] https://docs.fd.io/csit/rls1908/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[359] https://git.fd.io/csit/tree/?h=rls1908
[360] https://git.fd.io/csit/tree/docs/lab?h=rls2001
[361] https://docs.fd.io/csit/rls2001/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[362] https://docs.fd.io/csit/rls2001/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[363] https://git.fd.io/csit/tree/?h=rls2001
[364] https://git.fd.io/csit/tree/docs/lab?h=rls2005
[365] https://docs.fd.io/csit/rls2005/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[366] https://docs.fd.io/csit/rls2005/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[367] https://git.fd.io/csit/tree/?h=rls2005
[368] https://git.fd.io/csit/tree/docs/lab?h=rls2009
[369] https://docs.fd.io/csit/rls2009/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[370] https://docs.fd.io/csit/rls2009/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[371] https://git.fd.io/csit/tree/?h=rls2009
[372] https://gerrit.fd.io/r/c/csit/+/28184

- Ver. 6 associated with CSIT rls2101 branch (HW[373], Linux[374], TRex[375], CSIT[376]).
  - The main change is TRex version upgrade: increase from 2.82 to 2.86[377].
- Ver. 7 associated with CSIT rls2106 branch (HW[378], Linux[379], TRex[380], CSIT[381]).
  - TRex version upgrade: increase from 2.86 to 2.88[382].
  - Ubuntu upgrade: upgrade from 18.04 LTS to 20.04.2 LTS[383].
- Ver. 8 associated with CSIT rls2110 branch (HW[384], Linux[385], TRex[386], CSIT[387]).
  - Intel NIC 700/800 series firmware upgrade based on DPDK compatibility matrix.
- Ver. 9 associated with CSIT rls2202 branch (HW[388], Linux[389], TRex[390], CSIT[391]).
  - Intel NIC 700/800 series firmware upgrade based on DPDK compatibility matrix.
- Ver. 10 associated with CSIT rls2206 branch (HW[392], Linux[393], TRex[394], CSIT[395]).
  - Intel NIC 700/800 series firmware upgrade based on DPDK compatibility matrix.
  - Mellanox 556A series firmware upgrade based on DPDK compatibility matrix.
  - Intel IceLake all core turbo frequency turned off. Current base frequency is 2.6GHz.

3.8.2 SUT Settings - Linux

System provisioning is done by combination of PXE boot unattented install and Ansible[396] described in CSIT Testbed Setup[397].

[373] https://git.fd.io/csit/tree/docs/lab?h=rls2101
[374] https://docs.fd.io/csit/rls2101/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[375] https://docs.fd.io/csit/rls2101/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[376] https://git.fd.io/csit/tree/?h=rls2101
[377] https://gerrit.fd.io/r/c/csit/+/29980
[378] https://git.fd.io/csit/tree/docs/lab?h=rls2106
[379] https://s3-docs.fd.io/csit/rls2106/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[380] https://s3-docs.fd.io/csit/rls2106/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[381] https://git.fd.io/csit/tree/?h=rls2106
[382] https://gerrit.fd.io/r/c/csit/+/31652
[383] https://gerrit.fd.io/r/c/csit/+/31290
[384] https://git.fd.io/csit/tree/docs/lab?h=rls2110
[385] https://s3-docs.fd.io/csit/rls2110/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[386] https://s3-docs.fd.io/csit/rls2110/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[387] https://git.fd.io/csit/tree/?h=rls2110
[388] https://git.fd.io/csit/tree/docs/lab?h=rls2202
[389] https://s3-docs.fd.io/csit/rls2202/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[390] https://s3-docs.fd.io/csit/rls2202/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[391] https://git.fd.io/csit/tree/?h=rls2202
[392] https://git.fd.io/csit/tree/docs/lab?h=rls2206
[393] https://s3-docs.fd.io/csit/rls2206/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[394] https://s3-docs.fd.io/csit/rls2206/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[395] https://git.fd.io/csit/tree/?h=rls2206
[396] https://www.ansible.com
[397] https://git.fd.io/csit/tree/fdio.infra.ansible?h=rls2206

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.

3.8.3 DUT Settings - DPDK

DPDK Version

DPDK-22.03

DPDK Compile Parameters

```
make install T=<arch>-<machine>-linuxapp-gcc -j
```

Testpmd Startup Configuration

Testpmd startup configuration changes per test case with different settings for *$$INT*, *$$CORES*, *$$RXQ*, *$$RXD* and max-pkt-len parameter if test is sending jumbo frames. Startup command template:

```
testpmd -v -l $$CORE_LIST -w $$INT1 -w $$INT2 --master-lcore 0 --in-memory -- --forward-mode=io --
↪burst=64 --txd=$$TXD --rxd=$$RXD --txq=$$TXQ --rxq=$$RXQ --tx-offloads=0x0 --numa --auto-start --
↪total-num-mbufs=16384 --nb-ports=2 --portmask=0x3 --disable-link-check --max-pkt-len=$$PKT_LEN [--
↪mbuf-size=16384] --nb-cores=$$CORES
```

L3FWD Startup Configuration

L3FWD startup configuration changes per test case with different settings for *$$INT*, *$$CORES* and enable-jumbo parameter if test is sending jumbo frames. Startup command template:

```
l3fwd -v -l $$CORE_LIST -w $$INT1 -w $$INT2 --master-lcore 0 --in-memory -- --parse-ptype --eth-
↪dest="0,${adj_mac0}" --eth-dest="1,${adj_mac1}" --config="${port_config}" [--enable-jumbo] -P -L -
↪p 0x3
```

3.8.4 TG Settings - TRex

TG Version

TRex v2.97

DPDK Version

DPDK v21.02

TG Installation

T-Rex installation is managed via Ansible role.

TG Startup Configuration

```
$ sudo -E -S sh -c 'cat << EOF > /etc/trex_cfg.yaml
- version: 2
  c: 8
  limit_memory: 8192
  interfaces: ["${pci1}","${pci2}"]
  port_info:
      - dest_mac: [${dest_mac1}]
        src_mac: [${src_mac1}]
      - dest_mac: [${dest_mac2}]
        src_mac: [${src_mac2}]
  platform :
      master_thread_id: 0
      latency_thread_id: 9
      dual_if:
```

(continues on next page)

(continued from previous page)

```
          - socket: 0
            threads: [1, 2, 3, 4, 5, 6, 7, 8]
EOF'
```

TG Startup Command (Stateless Mode)

```
$ sudo -E -S sh -c "cd '${trex_install_dir}/scripts/' && \
  nohup ./t-rex-64 -i --prefix $(hostname) --hdrh --no-scapy-server \
  --mbuf-factor 32 > /tmp/trex.log 2>&1 &" > /dev/null
```

Also, Python client is now starting traffic with:

```
core_mask=STLClient.CORE_MASK_PIN
```

TG Startup Command (Stateful Mode)

```
$ sudo -E -S sh -c "cd '${trex_install_dir}/scripts/' && \
  nohup ./t-rex-64 -i --prefix $(hostname) --astf --hdrh --no-scapy-server \
  --mbuf-factor 32 > /tmp/trex.log 2>&1 &" > /dev/null
```

TG API Driver

TRex driver[398]

3.8.5 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[399].
2. Memory bandwidth - measure bandwidth with Intel MLC tool[400].
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 testbeds.

[398] https://git.fd.io/csit/tree/GPL/tools/trex/trex_stl_profile.py?h=rls2206
[399] https://git.fd.io/pma_tools/tree/jitter
[400] https://software.intel.com/en-us/articles/intelr-memory-latency-checker

Ice Lake

Following sections include sample calibration data measured on s71-t212-sut1 server running in one of the Intel Xeon Ice Lake testbeds as specified in `FD.io CSIT testbeds - Xeon Ice Lake`_.

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

Linux cmdline

```
$ cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-5.4.0-65-generic root=UUID=3250758a-9bb6-48c8-9c36-ecb6a269223f ro audit=0␣
↪default_hugepagesz=2M hugepagesz=1G hugepages=32 hugepagesz=2M hugepages=32768 hpet=disable intel_
↪idle.max_cstate=1 intel_iommu=on intel_pstate=disable iommu=pt isolcpus=1-31,33-63,65-95,97-127␣
↪mce=off nmi_watchdog=0 nohz_full=1-31,33-63,65-95,97-127 nosoftlockup numa_balancing=disable␣
↪processor.max_cstate=1 rcu_nocbs=1-31,33-63,65-95,97-127 tsc=reliable console=ttyS0,115200n8 quiet
```

Linux uname

```
$ uname -a
Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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 1.9
Iterations=20
The pragram will execute a dummy function 80000 times
Display is updated every 20000 displayUpdate intervals
Thread affinity will be set to core_id:7
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 calcualted 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
126082,133950,7868,126094,126082,133950,3829268480,2524167454,1
126082,134696,8614,126094,126082,134696,1778253824,2524273022,2
126082,136092,10010,126094,126082,136092,4022206464,2524203296,3
126082,135094,9012,126094,126082,136092,1971191808,2524274302,4
126082,136482,10400,126094,126082,136482,4215144448,2524318496,5
126082,134990,8908,126094,126082,136482,2164129792,2524155038,6
126082,134710,8628,126092,126082,136482,113115136,2524215228,7
126082,135080,8998,126092,126082,136482,2357067776,2524168906,8
126082,134470,8388,126094,126082,136482,306053120,2524163312,9
126082,135246,9164,126092,126082,136482,2550005760,2524394986,10
126082,132662,6580,126094,126082,136482,498991104,2524163156,11
126082,132954,6872,126094,126082,136482,2742943744,2524154386,12
126082,135340,9258,126092,126082,136482,691929088,2524222386,13
126082,133036,6954,126094,126082,136482,2935881728,2524150132,14
```

(continues on next page)

(continued from previous page)

```
126082,137776,11694,126094,126082,137776,884867072,2524239346,15
126082,137850,11768,126094,126082,137850,3128819712,2524342944,16
126082,133000,6918,126094,126082,137850,1077805056,2524160062,17
126082,133332,7250,126094,126082,137850,3321757696,2524158804,18
126082,133234,7152,126092,126082,137850,1270743040,2524174400,19
126082,152552,26470,126094,126082,152552,3514695680,2524857280,20
```

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[401].

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 x86_64
CPU is Intel(R) Xeon(R) Platinum 8358 CPU @ 2.60GHz

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)
  * Microarchitectural 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: YES
  * 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 explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): YES
  * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): YES
  * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): YES
    * TSX_CTRL MSR indicates TSX RTM is disabled: YES
    * TSX_CTRL MSR indicates TSX CPUID bit is cleared: YES
  * CPU supports Transactional Synchronization Extensions (TSX): NO
  * CPU supports Software Guard Extensions (SGX): YES
  * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
  * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x6a stepping 0x6␣
↪ucode 0xd000280 cpuid 0x606a6)
  * CPU microcode is the latest known available version: NO (latest version is 0xd0002a0 dated 2021/
↪04/25 according to builtin firmwares DB v191+i20210217)
* CPU vulnerability to the speculative execution attack variants
```

(continues on next page)

[401] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): YES
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): YES
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __user␣
↪pointer sanitization)
> STATUS: UNKNOWN (/sys vulnerability interface use forced, but its not available!)

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: YES (Mitigation: Enhanced IBRS, IBPB: conditional, RSB␣
↪filling)
> STATUS: VULNERABLE (IBRS+IBPB or retpoline+IBPB is needed to mitigate the vulnerability)

CVE-2017-5754 aka Variant 3, Meltdown, rogue data cache load
* Mitigated according to the /sys interface: YES (Not affected)
* 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)
> 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: YES
> STATUS: NOT VULNERABLE (your CPU microcode mitigates the vulnerability)

CVE-2018-3620 aka Foreshadow-NG (OS), L1 terminal fault
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (Not affected)

CVE-2018-3646 aka Foreshadow-NG (VMM), L1 terminal fault
* Information from the /sys interface: Not affected
> STATUS: NOT VULNERABLE (your kernel 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)
> 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)
```

(continues on next page)

(continued from previous page)

```
> 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)
> 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)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (Not affected)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

> SUMMARY: CVE-2017-5753:?? CVE-2017-5715:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

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[402].

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

Linux cmdline

```
$ cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-5.4.0-65-generic root=UUID=55d44abd-94d6-4b26-9d93-5877a8658016 ro audit=0␣
↪hpet=disable intel_idle.max_cstate=1 intel_iommu=on intel_pstate=disable iommu=pt isolcpus=1-27,
↪29-55,57-83,85-111 mce=off nmi_watchdog=0 nohz_full=1-27,29-55,57-83,85-111 nosoftlockup numa_
↪balancing=disable processor.max_cstate=1 rcu_nocbs=1-27,29-55,57-83,85-111 tsc=reliable␣
↪console=ttyS0,115200n8 quiet
```

Linux uname

```
$ uname -a
Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 x86_64 x86_64 x86_64 GNU/Linux
```

[402] https://git.fd.io/csit/tree/docs/lab/testbeds_sm_skx_hw_bios_cfg.md?h=rls2206

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 pragram will execute a dummy function 80000 times
Display is updated every 20000 displayUpdate intervals
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 calcualted 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
    160022     171330      11308     160022     160022     171330    2538733568 3204142750         ␣
↪1
    160022     167294       7272     160026     160022     171330     328335360 3203873548         ␣
↪2
    160022     167560       7538     160026     160022     171330    2412904448 3203878736         ␣
↪3
    160022     169000       8978     160024     160022     171330     202506240 3203864588         ␣
↪4
    160022     166572       6550     160026     160022     171330    2287075328 3203866224         ␣
↪5
    160022     167460       7438     160026     160022     171330      76677120 3203854632         ␣
↪6
    160022     168134       8112     160024     160022     171330    2161246208 3203874674         ␣
↪7
    160022     169094       9072     160022     160022     171330    4245815296 3203878798         ␣
↪8
    160022     172460      12438     160024     160022     172460    2035417088 3204112010         ␣
↪9
    160022     167862       7840     160030     160022     172460    4119986176 3203856800         ␣
↪10
    160022     168398       8376     160024     160022     172460    1909587968 3203854192         ␣
↪11
    160022     167548       7526     160024     160022     172460    3994157056 3203847442         ␣
↪12
    160022     167562       7540     160026     160022     172460    1783758848 3203862936         ␣
↪13
    160022     167604       7582     160024     160022     172460    3868327936 3203859346         ␣
↪14
    160022     168262       8240     160024     160022     172460    1657929728 3203851120         ␣
↪15
    160022     169700       9678     160024     160022     172460    3742498816 3203877690         ␣
↪16
    160022     170476      10454     160026     160022     172460    1532100608 3204088480         ␣
↪17
    160022     167798       7776     160024     160022     172460    3616669696 3203862072         ␣
↪18
    160022     166540       6518     160024     160022     172460    1406271488 3203836904         ␣
↪19
    160022     167516       7494     160024     160022     172460    3490840576 3203848120         ␣
↪20
```

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     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[403].

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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)
  * Microarchitectural Data Sampling
    * VERW instruction is available: YES (MD_CLEAR feature bit)
  * 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
```

(continues on next page)

[403] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
  * CPU explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): NO
  * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): NO
  * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): NO
  * CPU supports Transactional Synchronization Extensions (TSX): YES (RTM feature bit)
  * CPU supports Software Guard Extensions (SGX): NO
  * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
  * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x55 stepping 0x4␣
↪ucode 0x2000065 cpuid 0x50654)
  * CPU microcode is the latest known available version: NO (latest version is 0x2006b06 dated 2021/
↪03/08 according to builtin firmwares DB v191+i20210217)
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): YES
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): YES
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪YES
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): YES
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪YES
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): YES
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): YES
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __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
* Kernel has array_index_nospec (arm64): NO
> STATUS: NOT VULNERABLE (Mitigation: usercopy/swapgs barriers and __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:␣
↪conditional, IBRS_FW, STIBP: conditional, 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)
  * 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)
```

(continues on next page)

(continued from previous page)

```
* 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 (boltd fwupd irqbalance systemd-
↪journald systemd-logind systemd-networkd systemd-resolved systemd-timesyncd systemd-udevd)
> 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
> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2018-12126 aka Fallout, microarchitectural store buffer data sampling (MSBDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2018-12130 aka ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2018-12127 aka RIDL, microarchitectural load port data sampling (MLPDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
```

(continues on next page)

(continued from previous page)

```
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2019-11091 aka RIDL, microarchitectural data sampling uncacheable memory (MDSUM)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
> STATUS: NOT VULNERABLE (Mitigation: Clear CPU buffers; SMT vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (KVM: Mitigation: Split huge pages)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
* iTLB Multihit mitigation enabled and active: YES (KVM: Mitigation: Split huge pages)
> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: YES (found SRBDS implementation evidence in␣
↪kernel image. Your kernel is up to date for SRBDS mitigation)
* SRBDS mitigation control is enabled and active: 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

Cascade Lake

Following sections include sample calibration data measured on s32-t27-sut1 server running in one of the Intel Xeon Skylake testbeds as specified in FD.io CSIT testbeds - Xeon Cascade Lake[404].

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

Linux cmdline

```
$ cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-5.4.0-65-generic root=UUID=b1f0dc29-1d4f-4777-b37d-a5e26e233d55 ro audit=0␣
↪hpet=disable intel_idle.max_cstate=1 intel_iommu=on intel_pstate=disable iommu=pt isolcpus=1-27,
↪29-55,57-83,85-111 mce=off nmi_watchdog=0 nohz_full=1-27,29-55,57-83,85-111 nosoftlockup numa_
↪balancing=disable processor.max_cstate=1 rcu_nocbs=1-27,29-55,57-83,85-111 tsc=reliable␣
↪console=ttyS0,115200n8 quiet
```

[404] https://git.fd.io/csit/tree/docs/lab/testbeds_sm_clx_hw_bios_cfg.md?h=rls2206

Linux uname

```
$ uname -a
Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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 1.9
Iterations=30
The pragram will execute a dummy function 80000 times
Display is updated every 20000 displayUpdate intervals
Thread affinity will be set to core_id:7
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 calcualted 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
160022,167590,7568,160026,160022,167590,2057568256,3203711852,1
160022,170628,10606,160024,160022,170628,4079222784,3204010824,2
160022,169824,9802,160024,160022,170628,1805910016,3203812064,3
160022,168832,8810,160030,160022,170628,3827564544,3203792594,4
160022,168248,8226,160026,160022,170628,1554251776,3203765920,5
160022,167834,7812,160028,160022,170628,3575906304,3203761114,6
160022,167442,7420,160024,160022,170628,1302593536,3203769250,7
160022,169120,9098,160028,160022,170628,3324248064,3203853340,8
160022,170710,10688,160024,160022,170710,1050935296,3203985878,9
160022,167952,7930,160024,160022,170710,3072589824,3203733756,10
160022,168314,8292,160030,160022,170710,799277056,3203741152,11
160022,169672,9650,160024,160022,170710,2820931584,3203739910,12
160022,168684,8662,160024,160022,170710,547618816,3203727336,13
160022,168246,8224,160024,160022,170710,2569273344,3203739052,14
160022,168134,8112,160030,160022,170710,295960576,3203735874,15
160022,170230,10208,160024,160022,170710,2317615104,3203996356,16
160022,167190,7168,160024,160022,170710,44302336,3203713628,17
160022,167304,7282,160024,160022,170710,2065956864,3203717954,18
160022,167500,7478,160024,160022,170710,4087611392,3203706674,19
160022,167302,7280,160024,160022,170710,1814298624,3203726452,20
160022,167266,7244,160024,160022,170710,3835953152,3203702804,21
160022,167820,7798,160022,160022,170710,1562640384,3203719138,22
160022,168100,8078,160024,160022,170710,3584294912,3203716636,23
160022,170408,10386,160024,160022,170710,1310982144,3203946958,24
160022,167276,7254,160024,160022,170710,3332636672,3203706236,25
160022,167052,7030,160024,160022,170710,1059323904,3203696444,26
160022,170322,10300,160024,160022,170710,3080978432,3203747514,27
160022,167332,7310,160024,160022,170710,807665664,3203716210,28
160022,167426,7404,160026,160022,170710,2829320192,3203700630,29
160022,168840,8818,160024,160022,170710,556007424,3203727658,30
```

Memory Bandwidth

```
$ sudo /home/testuser/mlc --bandwidth_matrix
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --bandwidth_matrix

Using buffer size of 100.000MiB/thread for reads and an additional 100.000MiB/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        122097.7     51327.9
       1        51309.2      122005.5
```

```
$ sudo /home/testuser/mlc --peak_injection_bandwidth
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --peak_injection_bandwidth

Using buffer size of 100.000MiB/thread for reads and an additional 100.000MiB/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        :      243159.4
3:1 Reads-Writes :      219132.5
2:1 Reads-Writes :      216603.1
1:1 Reads-Writes :      203713.0
Stream-triad like:      193790.8
```

```
$ sudo /home/testuser/mlc --max_bandwidth
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --max_bandwidth

Using buffer size of 100.000MiB/thread for reads and an additional 100.000MiB/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        :      244114.27
3:1 Reads-Writes :      219441.97
2:1 Reads-Writes :      216603.72
1:1 Reads-Writes :      203679.09
Stream-triad like:      214902.80
```

Memory Latency

```
$ sudo /home/testuser/mlc --latency_matrix
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --latency_matrix

Using buffer size of 2000.000MiB
Measuring idle latencies (in ns)...
                Numa node
Numa node            0       1
       0          81.2   130.2
       1         130.2    81.1
```

```
$ sudo /home/testuser/mlc --idle_latency
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --idle_latency

Using buffer size of 2000.000MiB
Each iteration took 186.1 core clocks ( 80.9    ns)
```

```
$ sudo /home/testuser/mlc --loaded_latency
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --loaded_latency

Using buffer size of 100.000MiB/thread for reads and an additional 100.000MiB/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  233.86   243421.9
 00002  230.61   243544.1
 00008  232.56   243394.5
 00015  229.52   244076.6
 00050  225.82   244290.6
 00100  161.65   236744.8
 00200  100.63   133844.0
 00300   96.84    90548.2
 00400   95.71    68504.3
 00500   95.68    55139.0
 00700   88.77    39798.4
 01000   84.74    28200.1
 01300   83.08    21915.5
 01700   82.27    16969.3
 02500   81.66    11810.6
 03500   81.98     8662.9
 05000   81.48     6306.8
 09000   81.17     3857.8
 20000   80.19     2179.9
```

L1/L2/LLC Latency

```
$ sudo /home/testuser/mlc --c2c_latency
Intel(R) Memory Latency Checker - v3.7
Command line parameters: --c2c_latency

Measuring cache-to-cache transfer latency (in ns)...
Local Socket L2->L2 HIT  latency        55.5
Local Socket L2->L2 HITM latency        55.6
Remote Socket L2->L2 HITM latency (data address homed in writer socket)
                      Reader Numa Node
Writer Numa Node     0       1
            0        -   115.6
            1    115.6       -
Remote Socket L2->L2 HITM latency (data address homed in reader socket)
                      Reader Numa Node
Writer Numa Node     0       1
            0        -   178.2
            1    178.4       -
```

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[405].

```
Spectre and Meltdown mitigation detection tool v0.44+

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)
   * Microarchitectural 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: YES
   * 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 explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): NO
   * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): NO
   * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): YES
     * TSX_CTRL MSR indicates TSX RTM is disabled: YES
```

(continues on next page)

[405] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
     * TSX_CTRL MSR indicates TSX CPUID bit is cleared: YES
   * CPU supports Transactional Synchronization Extensions (TSX): NO
   * CPU supports Software Guard Extensions (SGX): NO
   * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
   * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x55 stepping 0x7␣
↪ucode 0x500002c cpuid 0x50657)
   * CPU microcode is the latest known available version: NO (latest version is 0x5003102 dated␣
↪2021/03/08 according to builtin firmwares DB v191+i20210217)
 * CPU vulnerability to the speculative execution attack variants
   * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
   * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
   * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
   * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
   * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
   * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
   * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
   * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
   * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
   * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
   * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
   * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory␣
↪(MDSUM)): NO
   * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
   * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): YES
   * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka  Spectre Variant 1, bounds check bypass
 * Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __user␣
↪pointer sanitization)
 > STATUS: UNKNOWN (/sys vulnerability interface use forced, but it  s not available!)

CVE-2017-5715 aka  Spectre Variant 2, branch target injection
 * Mitigated according to the /sys interface: YES (Mitigation: Enhanced IBRS, IBPB: conditional,␣
↪RSB filling)
 > STATUS: VULNERABLE (IBRS+IBPB or retpoline+IBPB+RSB filling, is needed to mitigate the␣
↪vulnerability)

CVE-2017-5754 aka  Variant 3, Meltdown, rogue data cache load
 * Mitigated according to the /sys interface: YES (Not affected)
 * 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)
 > 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)
 > STATUS: NOT VULNERABLE (Not affected)
```

(continues on next page)

(continued from previous page)

```
CVE-2018-3646 aka  Foreshadow-NG (VMM), L1 terminal fault
 * Information from the /sys interface: Not affected
 > STATUS: NOT VULNERABLE (your kernel 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)
 > 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)
 > 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)
 > 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)
 > STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2019-11135 aka  ZombieLoad V2, TSX Asynchronous Abort (TAA)
 * Mitigated according to the /sys interface: YES (Mitigation: TSX disabled)
 > STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka  No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
 * Mitigated according to the /sys interface: YES (KVM: Mitigation: Split huge pages)
 > STATUS: NOT VULNERABLE (KVM: Mitigation: Split huge pages)

CVE-2020-0543 aka  Special Register Buffer Data Sampling (SRBDS)
 * Mitigated according to the /sys interface: YES (Not affected)
 > STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

 > SUMMARY: CVE-2017-5753:?? CVE-2017-5715:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 x86_64
CPU is Intel(R) Xeon(R) Gold 6252N CPU @ 2.30GHz

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)
  * Microarchitectural Data Sampling
    * VERW instruction is available: YES (MD_CLEAR feature bit)
  * Enhanced IBRS (IBRS_ALL)
```

(continues on next page)

(continued from previous page)

```
    * CPU indicates ARCH_CAPABILITIES MSR availability: YES
    * ARCH_CAPABILITIES MSR advertises IBRS_ALL capability: YES
  * 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 explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): NO
  * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): NO
  * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): YES
    * TSX_CTRL MSR indicates TSX RTM is disabled: YES
    * TSX_CTRL MSR indicates TSX CPUID bit is cleared: YES
  * CPU supports Transactional Synchronization Extensions (TSX): NO
  * CPU supports Software Guard Extensions (SGX): NO
  * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
  * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x55 stepping 0x7␣
↪ucode 0x500002c cpuid 0x50657)
  * CPU microcode is the latest known available version: NO (latest version is 0x5003102 dated 2021/
↪03/08 according to builtin firmwares DB v191+i20210217)
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): YES
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __user␣
↪pointer sanitization)
> STATUS: UNKNOWN (/sys vulnerability interface use forced, but its not available!)

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: YES (Mitigation: Enhanced IBRS, IBPB: conditional, RSB␣
↪filling)
> STATUS: VULNERABLE (IBRS+IBPB or retpoline+IBPB+RSB filling, is needed to mitigate the␣
↪vulnerability)

CVE-2017-5754 aka Variant 3, Meltdown, rogue data cache load
* Mitigated according to the /sys interface: YES (Not affected)
* 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)
```

(continues on next page)

(continued from previous page)

```
> 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)
> STATUS: NOT VULNERABLE (Not affected)

CVE-2018-3646 aka Foreshadow-NG (VMM), L1 terminal fault
* Information from the /sys interface: Not affected
> STATUS: NOT VULNERABLE (your kernel 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)
> 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)
> 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)
> 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)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Mitigation: TSX disabled)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (KVM: Mitigation: Split huge pages)
> STATUS: NOT VULNERABLE (KVM: Mitigation: Split huge pages)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

> SUMMARY: CVE-2017-5753:?? CVE-2017-5715:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

Denverton

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

A 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-5.4.0-65-generic root=UUID=26ca7b0f-904a-462d-a1c6-98c420c29515 ro audit=0␣
↪hpet=disable intel_idle.max_cstate=1 intel_iommu=on intel_pstate=disable iommu=pt isolcpus=1-5␣
↪mce=off nmi_watchdog=0 nohz_full=1-5 nosoftlockup numa_balancing=disable processor.max_cstate=1␣
↪rcu_nocbs=1-5 tsc=reliable console=tty0 console=ttyS0,115200n8
```

Linux uname

```
$ uname -a
Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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 pragram will execute a dummy function 80000 times
Display is updated every 20000 displayUpdate intervals
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 calcualted 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
    177530     196100      18570     177530     177530     196100    4156751872 3556820054         ␣
↪1
    177530     200784      23254     177530     177530     200784     321060864 3556897644         ␣
↪2
    177530     196346      18816     177530     177530     200784     780337152 3556918674         ␣
↪3
    177530     195962      18432     177530     177530     200784    1239613440 3556847928         ␣
↪4
    177530     195960      18430     177530     177530     200784    1698889728 3556860214         ␣
↪5
    177530     198824      21294     177530     177530     200784    2158166016 3556854934         ␣
↪6
    177530     198522      20992     177530     177530     200784    2617442304 3556862410         ␣
↪7
    177530     196362      18832     177530     177530     200784    3076718592 3556851636         ␣
↪8
    177530     199114      21584     177530     177530     200784    3535994880 3556870846         ␣
↪9
    177530     197194      19664     177530     177530     200784    3995271168 3556933584        ␣
↪10
    177530     198272      20742     177536     177530     200784     159580160 3556869044        ␣
↪11
```

(continues on next page)

(continued from previous page)

```
    177530     197586      20056     177530     177530     200784      618856448 3556903482
↪12
    177530     196072      18542     177530     177530     200784     1078132736 3556825540
↪13
    177530     196354      18824     177530     177530     200784     1537409024 3556881664
↪14
    177530     195906      18376     177530     177530     200784     1996685312 3556839924
↪15
    177530     199066      21536     177530     177530     200784     2455961600 3556860220
↪16
    177530     196968      19438     177530     177530     200784     2915237888 3556871890
↪17
    177530     195896      18366     177530     177530     200784     3374514176 3556855338
↪18
    177530     196020      18490     177530     177530     200784     3833790464 3556839820
↪19
    177530     196030      18500     177530     177530     200784     4293066752 3556889196
↪20
```

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
        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
```

(continues on next page)

(continued from previous page)

```
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
```

(continues on next page)

(continued from previous page)

```
 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[406].

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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)
  * Microarchitectural Data Sampling
    * VERW instruction is available: YES (MD_CLEAR feature bit)
  * 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 explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): NO
  * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): NO
  * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): NO
  * CPU supports Transactional Synchronization Extensions (TSX): YES (RTM feature bit)
```

(continues on next page)

[406] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
  * CPU supports Software Guard Extensions (SGX): NO
  * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
  * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x55 stepping 0x4␣
↪ucode 0x2000065 cpuid 0x50654)
  * CPU microcode is the latest known available version: NO (latest version is 0x2006b06 dated 2021/
↪03/08 according to builtin firmwares DB v191+i20210217)
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): YES
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): YES
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪YES
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): YES
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪YES
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): YES
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): YES
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __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
* Kernel has array_index_nospec (arm64): NO
> STATUS: NOT VULNERABLE (Mitigation: usercopy/swapgs barriers and __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:␣
↪conditional, IBRS_FW, STIBP: conditional, 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)
  * 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)
```

(continues on next page)

(continued from previous page)

```
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 (boltd fwupd irqbalance systemd-
↪journald systemd-logind systemd-networkd systemd-resolved systemd-timesyncd systemd-udevd)
> 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
> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2018-12126 aka Fallout, microarchitectural store buffer data sampling (MSBDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2018-12130 aka ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2018-12127 aka RIDL, microarchitectural load port data sampling (MLPDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)
```

(continues on next page)

(continued from previous page)

```
CVE-2019-11091 aka RIDL, microarchitectural data sampling uncacheable memory (MDSUM)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
> STATUS: NOT VULNERABLE (Mitigation: Clear CPU buffers; SMT vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (KVM: Mitigation: Split huge pages)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
* iTLB Multihit mitigation enabled and active: YES (KVM: Mitigation: Split huge pages)
> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: YES (found SRBDS implementation evidence in␣
↪kernel image. Your kernel is up to date for SRBDS mitigation)
* SRBDS mitigation control is enabled and active: 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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: NO
  * L1 data cache invalidation
    * FLUSH_CMD MSR is available: NO
    * CPU indicates L1D flush capability: NO
  * Microarchitectural Data Sampling
    * VERW instruction is available: NO
  * 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
```

(continues on next page)

(continued from previous page)

```
  * 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 explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): NO
  * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): NO
  * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): NO
  * CPU supports Transactional Synchronization Extensions (TSX): NO
  * CPU supports Software Guard Extensions (SGX): NO
  * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
  * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x5f stepping 0x1␣
↪ucode 0x20 cpuid 0x506f1)
  * CPU microcode is the latest known available version: NO (latest version is 0x34 dated 2020/10/
↪23 according to builtin firmwares DB v191+i20210217)
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): NO
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): NO
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): NO
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __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
* Kernel has array_index_nospec (arm64): NO
> STATUS: NOT VULNERABLE (Mitigation: usercopy/swapgs barriers and __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:␣
↪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)
```

(continues on next page)

(continued from previous page)

```
* Kernel supports Page Table Isolation (PTI): YES
  * PTI enabled and active: NO
  * Reduced performance impact of PTI: NO (PCID/INVPCID 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: NO
> STATUS: VULNERABLE (an up-to-date CPU microcode is needed to mitigate this vulnerability)

CVE-2018-3639 aka Variant 4, speculative store bypass
* Mitigated according to the /sys interface: NO (Vulnerable)
* Kernel supports disabling speculative store bypass (SSB): YES (found in /proc/self/status)
* SSB mitigation is enabled and active: NO
> STATUS: VULNERABLE (Your CPU doesnt support SSBD)

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)
* Kernel supports using MD_CLEAR mitigation: YES (found md_clear implementation evidence in kernel␣
↪image)
* 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 (found md_clear implementation evidence in kernel␣
↪image)
* 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 (found md_clear implementation evidence in kernel␣
↪image)
* Kernel mitigation is enabled and active: NO
* SMT is either mitigated or disabled: NO
```

(continues on next page)

(continued from previous page)

```
> 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 (found md_clear implementation evidence in kernel␣
↪image)
* 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-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Not affected)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (Not affected)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
* iTLB Multihit mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: YES (found SRBDS implementation evidence in␣
↪kernel image. Your kernel is up to date for SRBDS mitigation)
* SRBDS mitigation control is enabled and active: 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:KO CVE-2018-3639:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

Altra

Following sections include sample calibration data measured on s62-t34-sut1 server running in one of the Altra testbeds.

Linux cmdline

```
$ cat /proc/cmdline
BOOT_IMAGE=/vmlinuz-5.4.0-65-generic root=/dev/mapper/ubuntu--vg-ubuntu--lv ro audit=0 default_
↪hugepagesz=2M hugepagesz=1G hugepages=32 hugepagesz=2M hugepages=32768 iommu.passthrough=1␣
↪isolcpus=1-40,81-120 nmi_watchdog=0 nohz_full=1-40,81-120 nosoftlockup processor.max_cstate=1 rcu_
↪nocbs=1-40,81-120
```

Linux uname

```
$ uname -a
Linux s62-t34-sut1 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:27:25 UTC 2021 aarch64 aarch64␣
↪aarch64 GNU/Linux
```

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[407].

```
Spectre and Meltdown mitigation detection tool v0.45

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:27:25 UTC 2021 aarch64
CPU is ARM v8 model 0xd0c

Hardware check
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): NO
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): NO
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): NO
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): NO
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): NO
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: __user pointer sanitization)
> STATUS: UNKNOWN (/sys vulnerability interface use forced, but it's not available!)

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2017-5754 aka Variant 3, Meltdown, rogue data cache load
* Mitigated according to the /sys interface: YES (Not affected)
* Running as a Xen PV DomU: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-3640 aka Variant 3a, rogue system register read
* CPU microcode mitigates the vulnerability: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-3639 aka Variant 4, speculative store bypass
* Mitigated according to the /sys interface: YES (Mitigation: Speculative Store Bypass disabled via␣
↪prctl)
```

(continues on next page)

[407] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
> STATUS: NOT VULNERABLE (Mitigation: Speculative Store Bypass disabled via prctl)

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 affected)

CVE-2018-3620 aka Foreshadow-NG (OS), L1 terminal fault
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-3646 aka Foreshadow-NG (VMM), L1 terminal fault
* Information from the /sys interface: Not affected
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-12126 aka Fallout, microarchitectural store buffer data sampling (MSBDS)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-12130 aka ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-12127 aka RIDL, microarchitectural load port data sampling (MLPDS)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2019-11091 aka RIDL, microarchitectural data sampling uncacheable memory (MDSUM)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not affected)

> SUMMARY: CVE-2017-5753:?? 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

TaiShan

Following sections include sample calibration data measured on s17-t33-sut1 server running in one of the Cortex-A72 testbeds.

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

Linux cmdline

```
$ cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-5.4.0-65-generic root=UUID=7d1d0e77-4df0-43df-9619-a99db29ffb83 ro audit=0␣
↪intel_iommu=on isolcpus=1-27,29-55 nmi_watchdog=0 nohz_full=1-27,29-55 nosoftlockup processor.max_
↪cstate=1 rcu_nocbs=1-27,29-55 console=ttyAMA0,115200n8 quiet
```

Linux uname

```
$ uname -a
Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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 1.9
Iterations=30
The pragram will execute a dummy function 80000 times
Display is updated every 20000 displayUpdate intervals
Thread affinity will be set to core_id:7
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 calcualted 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
    160022     172254      12232     160042     160022     172254    1903230976 3204401362      ␣
↪1
    160022     173148      13126     160044     160022     173148     814809088 3204619316      ␣
↪2
    160022     169460       9438     160044     160022     173148    4021354496 3204391306      ␣
↪3
    160024     170270      10246     160044     160022     173148    2932932608 3204385830      ␣
↪4
    160022     169660       9638     160044     160022     173148    1844510720 3204387290      ␣
↪5
    160022     169410       9388     160040     160022     173148     756088832 3204375832      ␣
↪6
    160022     169012       8990     160042     160022     173148    3962634240 3204378924      ␣
↪7
    160022     169556       9534     160044     160022     173148    2874212352 3204374882      ␣
↪8
    160022     171684      11662     160042     160022     173148    1785790464 3204394596      ␣
↪9
    160022     171546      11524     160024     160022     173148     697368576 3204602774      ␣
↪10
    160022     169248       9226     160042     160022     173148    3903913984 3204401676      ␣
↪11
    160022     168458       8436     160042     160022     173148    2815492096 3204256350      ␣
↪12
```

(continues on next page)

(continued from previous page)

```
    160022     169574       9552     160044     160022     173148    1727070208 3204278116 _
↪13
    160022     169352       9330     160044     160022     173148     638648320 3204327234 _
↪14
    160022     169100       9078     160044     160022     173148    3845193728 3204388132 _
↪15
    160022     169338       9316     160042     160022     173148    2756771840 3204380724 _
↪16
    160022     170828      10806     160046     160022     173148    1668349952 3204430452 _
↪17
    160022     173162      13140     160026     160022     173162     579928064 3204611318 _
↪18
    160022     170482      10460     160042     160022     173162    3786473472 3204389896 _
↪19
    160024     170704      10680     160044     160022     173162    2698051584 3204422126 _
↪20
    160024     169302       9278     160044     160022     173162    1609629696 3204397334 _
↪21
    160022     171848      11826     160044     160022     173162     521207808 3204389818 _
↪22
    160022     169438       9416     160042     160022     173162    3727753216 3204395382 _
↪23
    160022     169312       9290     160042     160022     173162    2639331328 3204371202 _
↪24
    160022     171368      11346     160044     160022     173162    1550909440 3204440464 _
↪25
    160022     171998      11976     160042     160022     173162     462487552 3204609440 _
↪26
    160022     169740       9718     160046     160022     173162    3669032960 3204405826 _
↪27
    160022     169610       9588     160044     160022     173162    2580611072 3204390608 _
↪28
    160022     169254       9232     160044     160022     173162    1492189184 3204399760 _
↪29
    160022     169386       9364     160046     160022     173162     403767296 3204417762 _
↪30
```

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[408].

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 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)
```

(continues on next page)

[408] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
  * 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)
  * Microarchitectural Data Sampling
    * VERW instruction is available: YES (MD_CLEAR feature bit)
  * 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 explicitly indicates not being vulnerable to TSX Asynchronous Abort (TAA_NO): NO
  * CPU explicitly indicates not being vulnerable to iTLB Multihit (PSCHANGE_MSC_NO): NO
  * CPU explicitly indicates having MSR for TSX control (TSX_CTRL_MSR): NO
  * CPU supports Transactional Synchronization Extensions (TSX): YES (RTM feature bit)
  * CPU supports Software Guard Extensions (SGX): NO
  * CPU supports Special Register Buffer Data Sampling (SRBDS): NO
  * CPU microcode is known to cause stability problems: NO (family 0x6 model 0x55 stepping 0x4␣
↪ucode 0x2000065 cpuid 0x50654)
  * CPU microcode is the latest known available version: NO (latest version is 0x2006b06 dated 2021/
↪03/08 according to builtin firmwares DB v191+i20210217)
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): YES
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): YES
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): YES
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): YES
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪YES
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): YES
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪YES
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): YES
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): YES
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): NO

CVE-2017-5753 aka Spectre Variant 1, bounds check bypass
* Mitigated according to the /sys interface: YES (Mitigation: usercopy/swapgs barriers and __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
* Kernel has array_index_nospec (arm64): NO
> STATUS: NOT VULNERABLE (Mitigation: usercopy/swapgs barriers and __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:␣
↪conditional, IBRS_FW, STIBP: conditional, RSB filling)
```

(continues on next page)

(continued from previous page)

```
* 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 (boltd fwupd irqbalance systemd-
↪journald systemd-logind systemd-networkd systemd-resolved systemd-timesyncd systemd-udevd)
> 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
> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2018-12126 aka Fallout, microarchitectural store buffer data sampling (MSBDS)
```

(continues on next page)

(continued from previous page)

```
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2018-12130 aka ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2018-12127 aka RIDL, microarchitectural load port data sampling (MLPDS)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2019-11091 aka RIDL, microarchitectural data sampling uncacheable memory (MDSUM)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* Kernel supports using MD_CLEAR mitigation: YES (md_clear found in /proc/cpuinfo)
* Kernel mitigation is enabled and active: YES
* SMT is either mitigated or disabled: NO
> STATUS: NOT VULNERABLE (Your microcode and kernel are both up to date for this mitigation, and␣
↪mitigation is enabled)

CVE-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: YES (Mitigation: Clear CPU buffers; SMT vulnerable)
> STATUS: NOT VULNERABLE (Mitigation: Clear CPU buffers; SMT vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (KVM: Mitigation: Split huge pages)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
* iTLB Multihit mitigation enabled and active: YES (KVM: Mitigation: Split huge pages)
> STATUS: NOT VULNERABLE (this system is not running a hypervisor)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: YES (found SRBDS implementation evidence in␣
↪kernel image. Your kernel is up to date for SRBDS mitigation)
* SRBDS mitigation control is enabled and active: 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:27:25 UTC 2021 aarch64
CPU is ARM v8 model 0xd08
```

(continues on next page)

(continued from previous page)

```
Hardware check
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): NO
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): NO
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): NO
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): 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: NO
* Kernel has the Red Hat/Ubuntu patch: NO
* Kernel has mask_nospec64 (arm64): NO
* Kernel has array_index_nospec (arm64): NO
* Checking count of LFENCE instructions following a jump in kernel... NO (only 0 jump-then-lfence␣
↪instructions found, should be >= 30 (heuristic))
> STATUS: NOT VULNERABLE (Mitigation: __user pointer sanitization)

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: NO (Vulnerable)
* Mitigation 1
  * Kernel is compiled with IBRS support: YES
    * IBRS enabled and active: NO
  * Kernel is compiled with IBPB support: NO
    * IBPB enabled and active: NO
* Mitigation 2
  * Kernel has branch predictor hardening (arm): YES
  * Kernel compiled with retpoline option: NO
> STATUS: NOT VULNERABLE (Branch predictor hardening mitigates 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/INVPCID 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: NO
> STATUS: VULNERABLE (an up-to-date CPU microcode is needed to mitigate this vulnerability)

CVE-2018-3639 aka Variant 4, speculative store bypass
* Mitigated according to the /sys interface: NO (Vulnerable)
* Kernel supports disabling speculative store bypass (SSB): YES (found in /proc/self/status)
* SSB mitigation is enabled and active: NO
> STATUS: VULNERABLE (Your CPU doesnt support SSBD)
```

(continues on next page)

(continued from previous page)

```
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: NO
* 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: N/A (the kvm_intel module is not loaded)
* Mitigation 2
  * L1D flush is supported by kernel: NO
  * L1D flush enabled: NO
  * Hardware-backed L1D flush supported: NO (flush will be done in software, this is slower)
  * Hyper-Threading (SMT) is enabled: UNKNOWN
> 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)
* Kernel supports using MD_CLEAR mitigation: NO
* 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: NO
* 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: NO
* 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: NO
* 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-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Not affected)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (Not affected)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
```

(continues on next page)

(continued from previous page)

```
* iTLB Multihit mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: NO
* SRBDS mitigation control is enabled and active: 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:KO CVE-2018-3639:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK

Need more detailed information about mitigation options? Use --explain
A false sense of security is worse than no security at all, see --disclaimer
ok: [10.30.51.37] =>
spectre_meltdown_poll_results.stdout_lines:
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:27:25 UTC 2021 aarch64
CPU is ARM v8 model 0xd08

Hardware check
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): YES
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): NO
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): NO
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): NO
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): 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: NO
* Kernel has the Red Hat/Ubuntu patch: NO
* Kernel has mask_nospec64 (arm64): NO
* Kernel has array_index_nospec (arm64): NO
* Checking count of LFENCE instructions following a jump in kernel... NO (only 0 jump-then-lfence␣
↪instructions found, should be >= 30 (heuristic))
> STATUS: NOT VULNERABLE (Mitigation: __user pointer sanitization)

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: NO (Vulnerable)
* Mitigation 1
  * Kernel is compiled with IBRS support: YES
    * IBRS enabled and active: NO
  * Kernel is compiled with IBPB support: NO
    * IBPB enabled and active: NO
```

(continues on next page)

(continued from previous page)

```
* Mitigation 2
  * Kernel has branch predictor hardening (arm): YES
  * Kernel compiled with retpoline option: NO
> STATUS: NOT VULNERABLE (Branch predictor hardening mitigates 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/INVPCID 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: NO
> STATUS: VULNERABLE (an up-to-date CPU microcode is needed to mitigate this vulnerability)

CVE-2018-3639 aka Variant 4, speculative store bypass
* Mitigated according to the /sys interface: NO (Vulnerable)
* Kernel supports disabling speculative store bypass (SSB): YES (found in /proc/self/status)
* SSB mitigation is enabled and active: NO
> STATUS: VULNERABLE (Your CPU doesnt support SSBD)

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: NO
* 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: N/A (the kvm_intel module is not loaded)
* Mitigation 2
  * L1D flush is supported by kernel: NO
  * L1D flush enabled: NO
  * Hardware-backed L1D flush supported: NO (flush will be done in software, this is slower)
  * Hyper-Threading (SMT) is enabled: UNKNOWN
> 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)
* Kernel supports using MD_CLEAR mitigation: NO
* 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: NO
* 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)
```

(continues on next page)

(continued from previous page)

```
* Mitigated according to the /sys interface: YES (Not affected)
* Kernel supports using MD_CLEAR mitigation: NO
* 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: NO
* 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-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Not affected)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (Not affected)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
* iTLB Multihit mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: NO
* SRBDS mitigation control is enabled and active: 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:KO CVE-2018-3639:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

ThunderX2

Following sections include sample calibration data measured on s27-t211-sut1 server running in one of the ThunderX2 testbeds.

Linux cmdline

```
$ cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-5.4.0-65-generic root=UUID=7d1d0e77-4df0-43df-9619-a99db29ffb83 ro audit=0␣
↪intel_iommu=on isolcpus=1-27,29-55 nmi_watchdog=0 nohz_full=1-27,29-55 nosoftlockup processor.max_
↪cstate=1 rcu_nocbs=1-27,29-55 console=ttyAMA0,115200n8 quiet
```

Linux uname

```
$ uname -a
Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:25:17 UTC 2021 x86_64 x86_64 x86_64 GNU/Linux
```

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[409].

```
Spectre and Meltdown mitigation detection tool v0.44+

Checking for vulnerabilities on current system
Kernel is Linux 5.4.0-65-generic #73-Ubuntu SMP Mon Jan 18 17:27:25 UTC 2021 aarch64
CPU is

Hardware check
* CPU vulnerability to the speculative execution attack variants
  * Affected by CVE-2017-5753 (Spectre Variant 1, bounds check bypass): YES
  * Affected by CVE-2017-5715 (Spectre Variant 2, branch target injection): YES
  * Affected by CVE-2017-5754 (Variant 3, Meltdown, rogue data cache load): NO
  * Affected by CVE-2018-3640 (Variant 3a, rogue system register read): NO
  * Affected by CVE-2018-3639 (Variant 4, speculative store bypass): YES
  * Affected by CVE-2018-3615 (Foreshadow (SGX), L1 terminal fault): NO
  * Affected by CVE-2018-3620 (Foreshadow-NG (OS), L1 terminal fault): NO
  * Affected by CVE-2018-3646 (Foreshadow-NG (VMM), L1 terminal fault): NO
  * Affected by CVE-2018-12126 (Fallout, microarchitectural store buffer data sampling (MSBDS)): NO
  * Affected by CVE-2018-12130 (ZombieLoad, microarchitectural fill buffer data sampling (MFBDS)):␣
↪NO
  * Affected by CVE-2018-12127 (RIDL, microarchitectural load port data sampling (MLPDS)): NO
  * Affected by CVE-2019-11091 (RIDL, microarchitectural data sampling uncacheable memory (MDSUM)):␣
↪NO
  * Affected by CVE-2019-11135 (ZombieLoad V2, TSX Asynchronous Abort (TAA)): NO
  * Affected by CVE-2018-12207 (No eXcuses, iTLB Multihit, machine check exception on page size␣
↪changes (MCEPSC)): NO
  * Affected by CVE-2020-0543 (Special Register Buffer Data Sampling (SRBDS)): 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: NO
* Kernel has the Red Hat/Ubuntu patch: NO
* Kernel has mask_nospec64 (arm64): NO
* Kernel has array_index_nospec (arm64): NO
* Checking count of LFENCE instructions following a jump in kernel... NO (only 0 jump-then-lfence␣
↪instructions found, should be >= 30 (heuristic))
> STATUS: NOT VULNERABLE (Mitigation: __user pointer sanitization)

CVE-2017-5715 aka Spectre Variant 2, branch target injection
* Mitigated according to the /sys interface: NO (Vulnerable)
* Mitigation 1
  * Kernel is compiled with IBRS support: YES
    * IBRS enabled and active: NO
  * Kernel is compiled with IBPB support: NO
    * IBPB enabled and active: NO
* Mitigation 2
  * Kernel has branch predictor hardening (arm): YES
  * Kernel compiled with retpoline option: NO
```

(continues on next page)

[409] https://github.com/speed47/spectre-meltdown-checker

(continued from previous page)

```
> STATUS: NOT VULNERABLE (Branch predictor hardening mitigates 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/INVPCID 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: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-3639 aka Variant 4, speculative store bypass
* Mitigated according to the /sys interface: NO (Vulnerable)
* Kernel supports disabling speculative store bypass (SSB): YES (found in /proc/self/status)
* SSB mitigation is enabled and active: NO
> STATUS: VULNERABLE (Your CPU doesnt support SSBD)

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: NO
* 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: N/A (the kvm_intel module is not loaded)
* Mitigation 2
  * L1D flush is supported by kernel: NO
  * L1D flush enabled: NO
  * Hardware-backed L1D flush supported: NO (flush will be done in software, this is slower)
  * Hyper-Threading (SMT) is enabled: UNKNOWN
> 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)
* Kernel supports using MD_CLEAR mitigation: NO
* 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: NO
* 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: NO
* Kernel mitigation is enabled and active: NO
```

(continues on next page)

(continued from previous page)

```
* 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: NO
* 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-2019-11135 aka ZombieLoad V2, TSX Asynchronous Abort (TAA)
* Mitigated according to the /sys interface: YES (Not affected)
* TAA mitigation is supported by kernel: YES (found tsx_async_abort in kernel image)
* TAA mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2018-12207 aka No eXcuses, iTLB Multihit, machine check exception on page size changes (MCEPSC)
* Mitigated according to the /sys interface: YES (Not affected)
* This system is a host running a hypervisor: NO
* iTLB Multihit mitigation is supported by kernel: YES (found itlb_multihit in kernel image)
* iTLB Multihit mitigation enabled and active: NO
> STATUS: NOT VULNERABLE (your CPU vendor reported your CPU model as not vulnerable)

CVE-2020-0543 aka Special Register Buffer Data Sampling (SRBDS)
* Mitigated according to the /sys interface: YES (Not affected)
* SRBDS mitigation control is supported by the kernel: NO
* SRBDS mitigation control is enabled and active: 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:KO 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 CVE-2019-11135:OK CVE-2018-12207:OK CVE-2020-0543:OK
```

CHAPTER

FOUR

TREX PERFORMANCE

4.1 Overview

TRex performance test results are reported for a range of processors. For description of physical testbeds used for TRex performance tests please refer to *Performance Physical Testbeds* (page 4).

4.1.1 Logical Topology

CSIT TRex performance tests are executed on physical testbeds described in *Performance Physical Testbeds* (page 4). Logical topology use 1 nic that has loopback connected ports. See figure below.

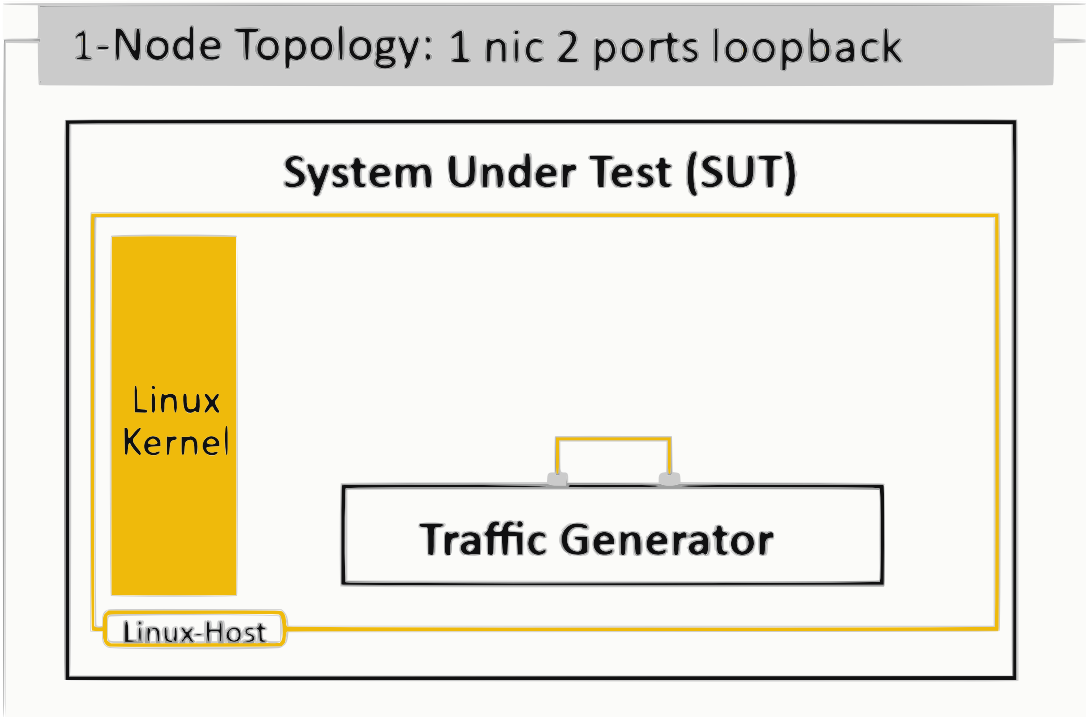

4.1.2 Performance Tests Coverage

Performance tests measure following metrics for tested TRex topologies and configurations:

- Packet Throughput: measured in accordance with RFC 2544[410], 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%.
- Two-way Packet Latency: measured both east-west and west-east at different offered packet loads:
  - 90% of discovered PDR throughput.
  - 50% of discovered PDR throughput.
  - 10% of discovered PDR throughput.
  - Minimal offered load.

CSIT-2206 includes tests using the following TRex traffic profiles (corresponding to data plane functionality when DUT is used) performance tested across a range of NIC drivers and NIC models:

Traffic profile	Corresponding dataplane functionality
IPv4 Base	IPv4 routing.
IPv4 Scale	IPv4 routing with 2M entries.
IPv6 Base	IPv6 routing.
IPv6 Scale	IPv6 routing with 2M entries.
L2BD Scale	L2 Bridge-Domain switching of untagged Ethernet frames.

4.1.3 Performance Tests Naming

FD.io CSIT-2206 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 1743).

4.2 Release Notes

4.2.1 Changes in CSIT-2206

1. TREX PERFORMANCE TESTS
   - No updates
2. TEST FRAMEWORK
3. TRex RELEASE VERSION - **TRex version used: 2.97**

[410] https://tools.ietf.org/html/rfc2544.html

4.2.2 Known Issues

List of known issues in CSIT-2206 for TRex performance tests:

#	JiraID	Issue Description
1		

4.3 Packet Throughput

Throughput graphs are generated by multiple executions of the same performance tests across physical testbeds hosted LF FD.io labs: 2n-skx. Box-and-Whisker plots are used to display variations in measured throughput values, without making any assumptions of the underlying statistical distribution.

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 DPDK 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×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 are stored in build logs from FD.io trex performance job 2n-skx[411]. Required per test case data set size is **10** and for TRex tests this is the actual size, as all scheduled test executions completed successfully.

---

[411] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-trex-perf-report-iterative-2206-2n-skx

4.3.1 2n-skx-x710

Following sections include summary graphs of Phy-to-Phy performance with packet 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[412].

[412] https://git.fd.io/csit/tree/tests/trex/perf?h=rls2206

64b-ip4routing-base-scale

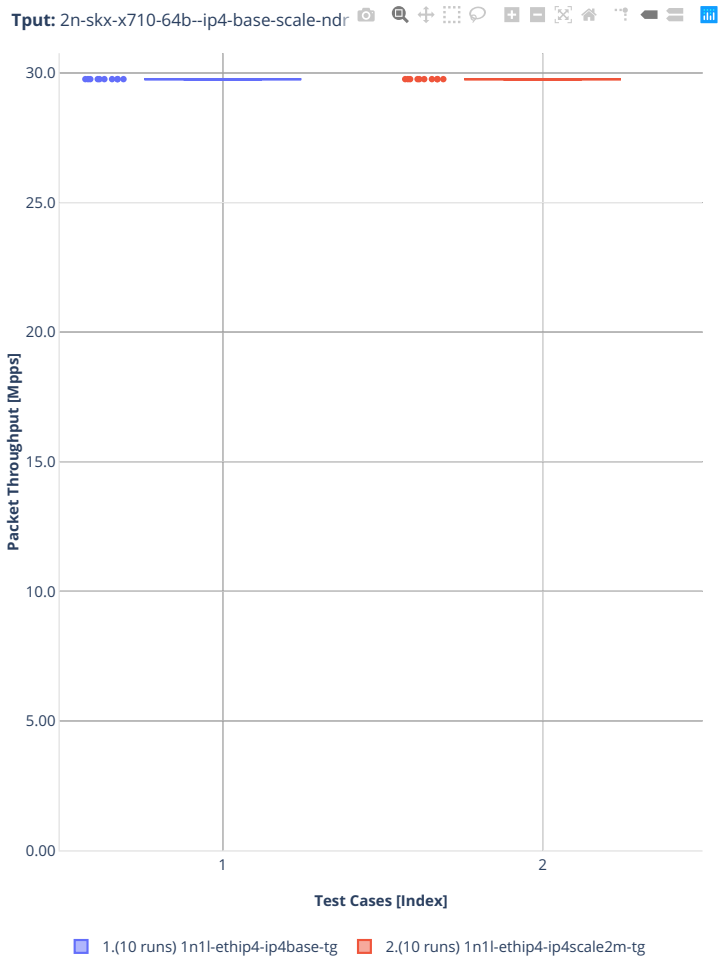

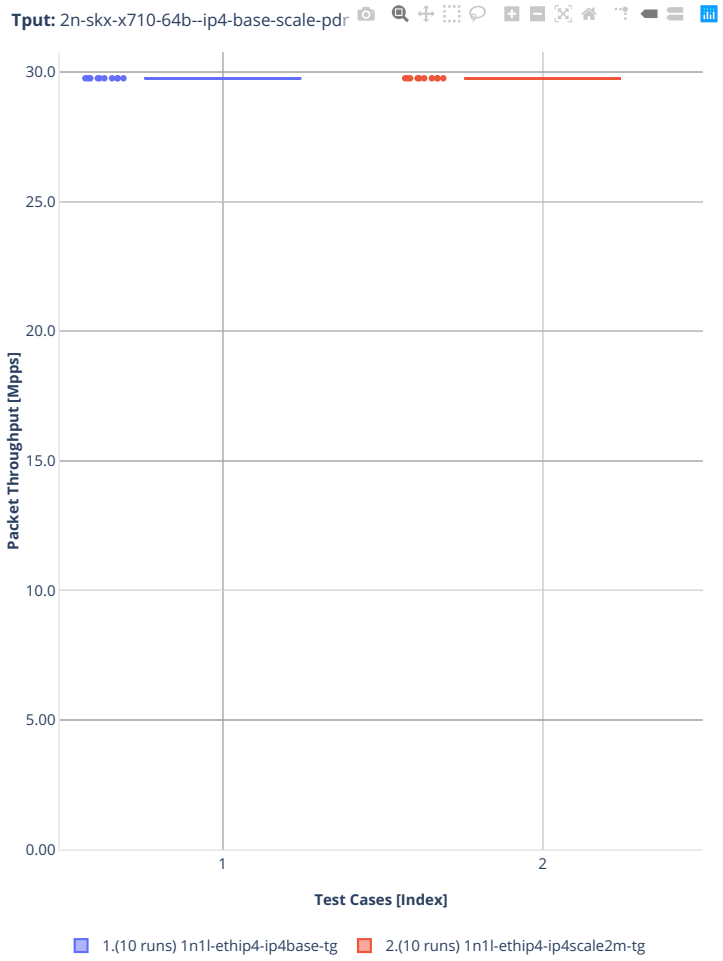
Tput: 2n-skx-x710-64b--ip4-base-scale-pdr
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) 1n1l-ethip4-ip4base-tg
2.(10 runs) 1n1l-ethip4-ip4scale2m-tg

64b-ip4routing-[udp|tcp]-stf-cps

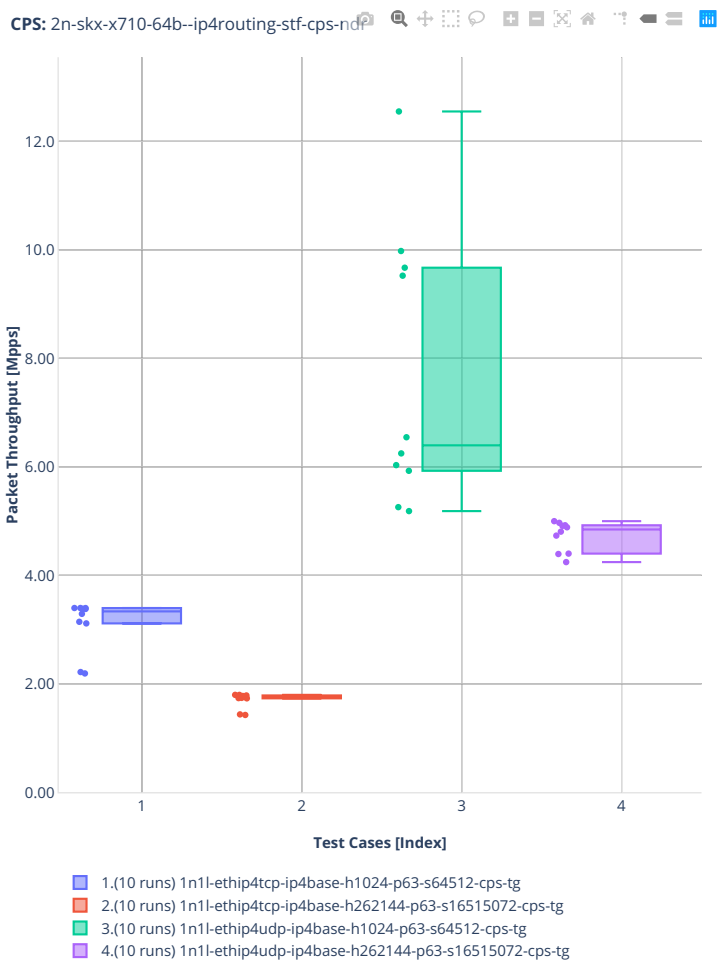

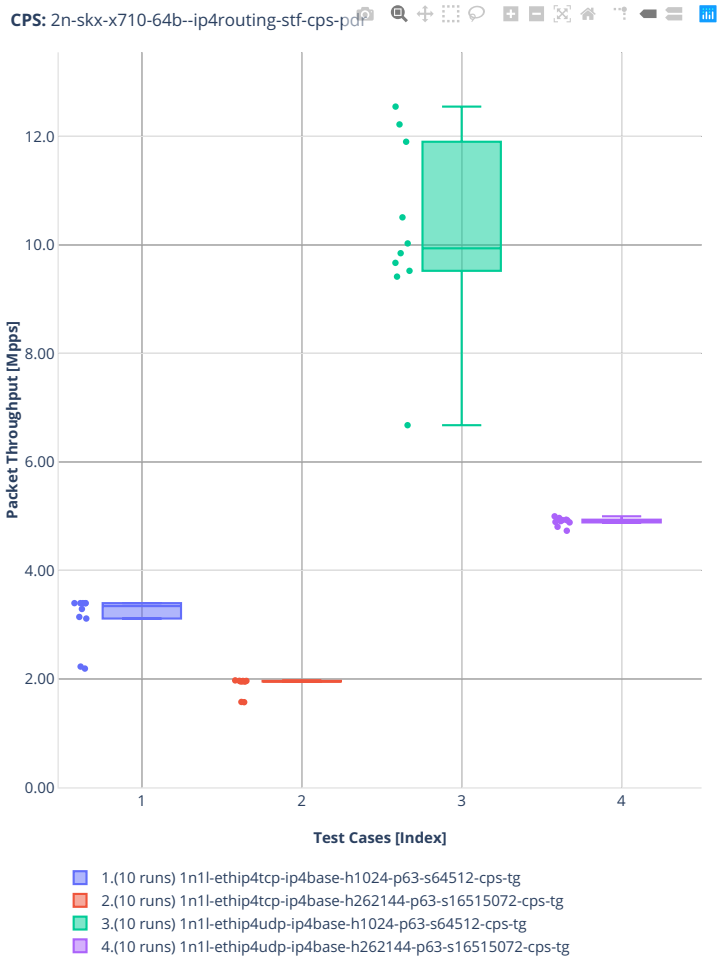
CPS: 2n-skx-x710-64b--ip4routing-stf-cps-pdr
Packet Throughput [Mpps]
12.0
10.0
8.00
6.00
4.00
2.00
0.00
1
2
3
4
Test Cases [Index]
1.(10 runs) 1n1l-ethip4tcp-ip4base-h1024-p63-s64512-cps-tg
2.(10 runs) 1n1l-ethip4tcp-ip4base-h262144-p63-s16515072-cps-tg
3.(10 runs) 1n1l-ethip4udp-ip4base-h1024-p63-s64512-cps-tg
4.(10 runs) 1n1l-ethip4udp-ip4base-h262144-p63-s16515072-cps-tg

64b-ip4routing-[udp|tcp]-stf-pps

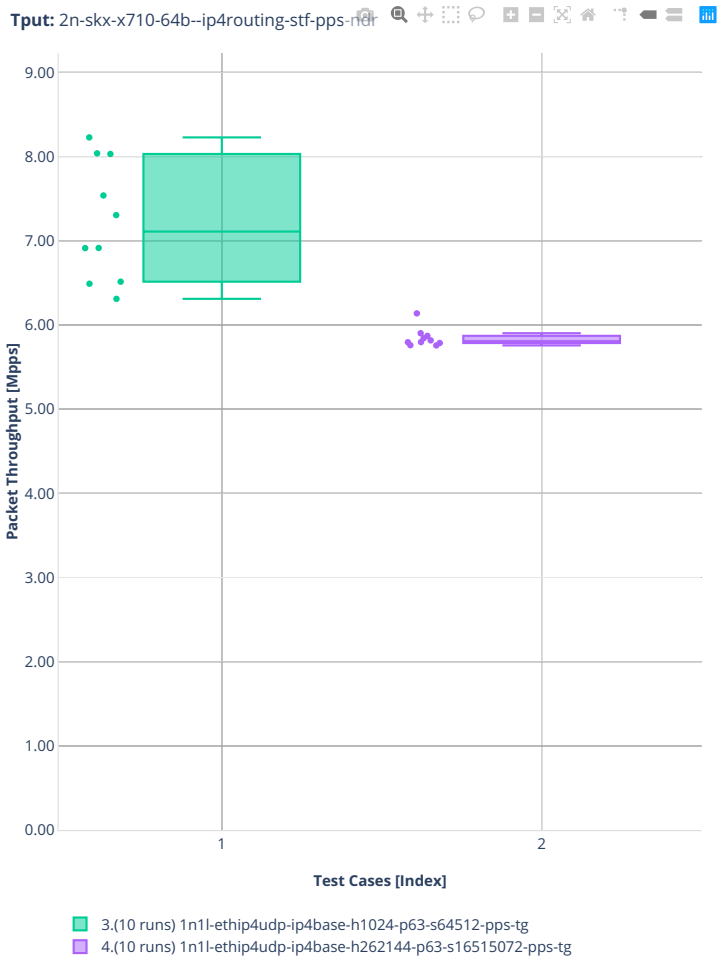

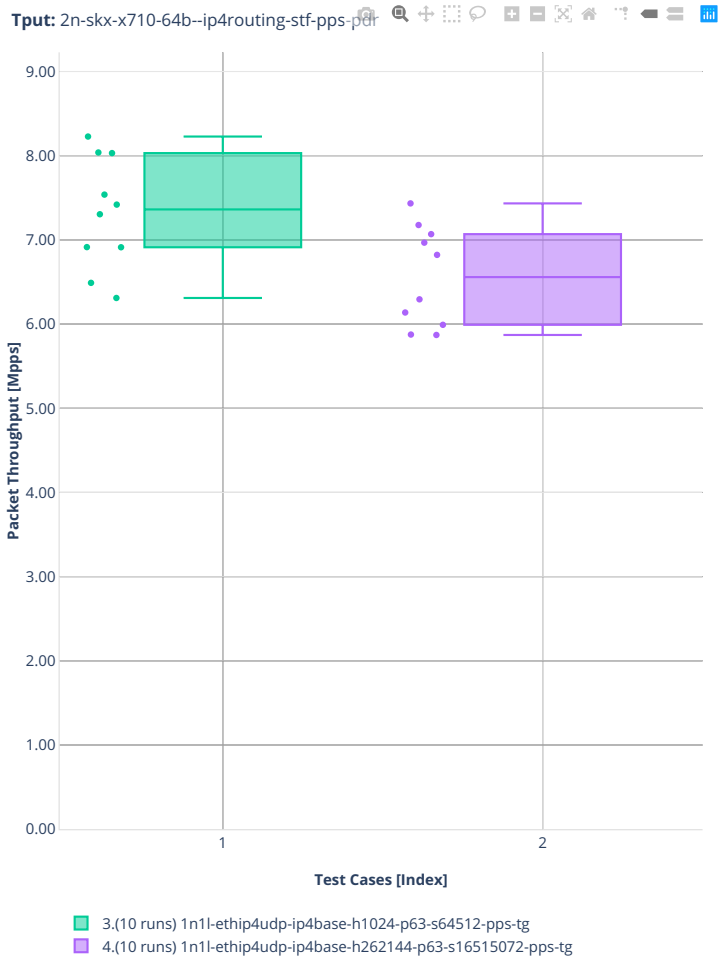
Tput: 2n-skx-x710-64b--ip4routing-stf-pps-pdr
9.00
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
3.(10 runs) 1n1l-ethip4udp-ip4base-h1024-p63-s64512-pps-tg
4.(10 runs) 1n1l-ethip4udp-ip4base-h262144-p63-s16515072-pps-tg

78b-ip6routing-base-scale

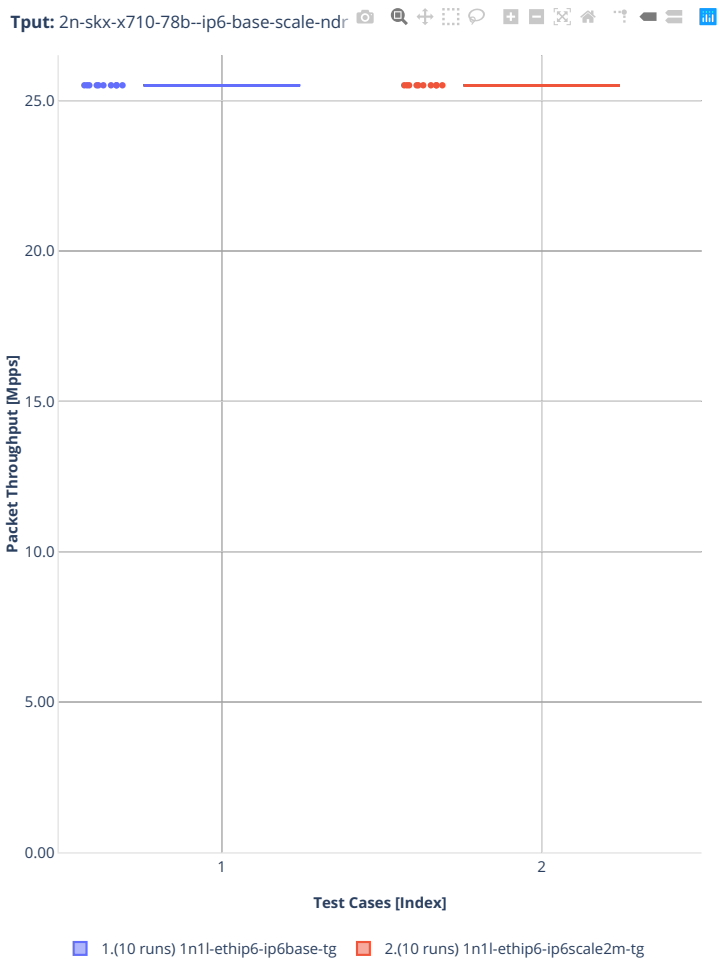

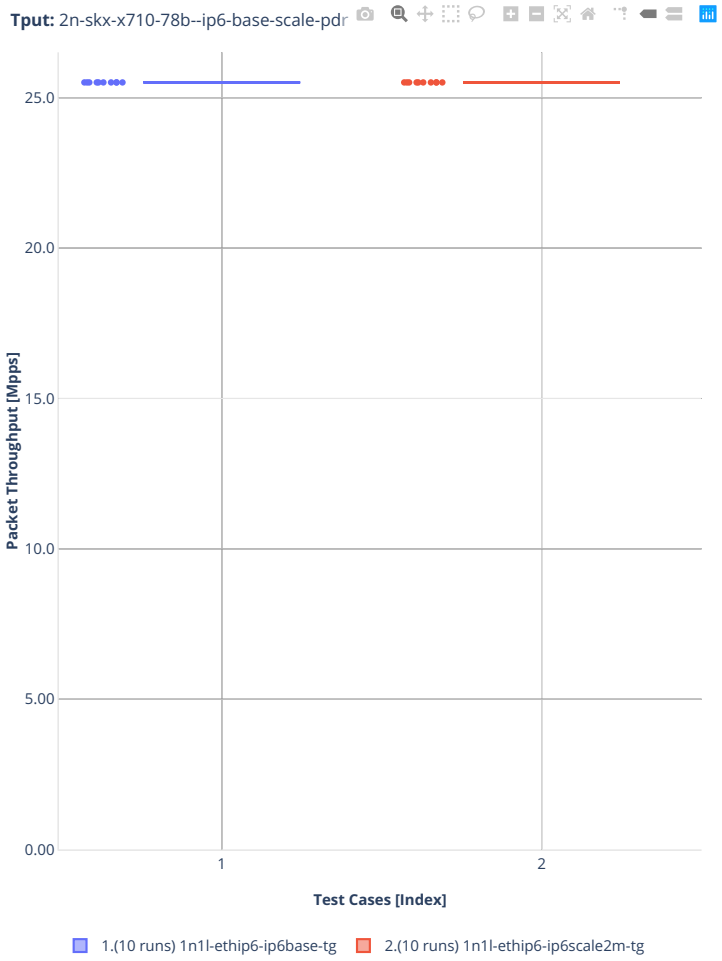
Tput: 2n-skx-x710-78b--ip6-base-scale-pdr
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
2
Test Cases [Index]
1.(10 runs) 1n1l-ethip6-ip6base-tg
2.(10 runs) 1n1l-ethip6-ip6scale2m-tg

64b-l2switching-scale

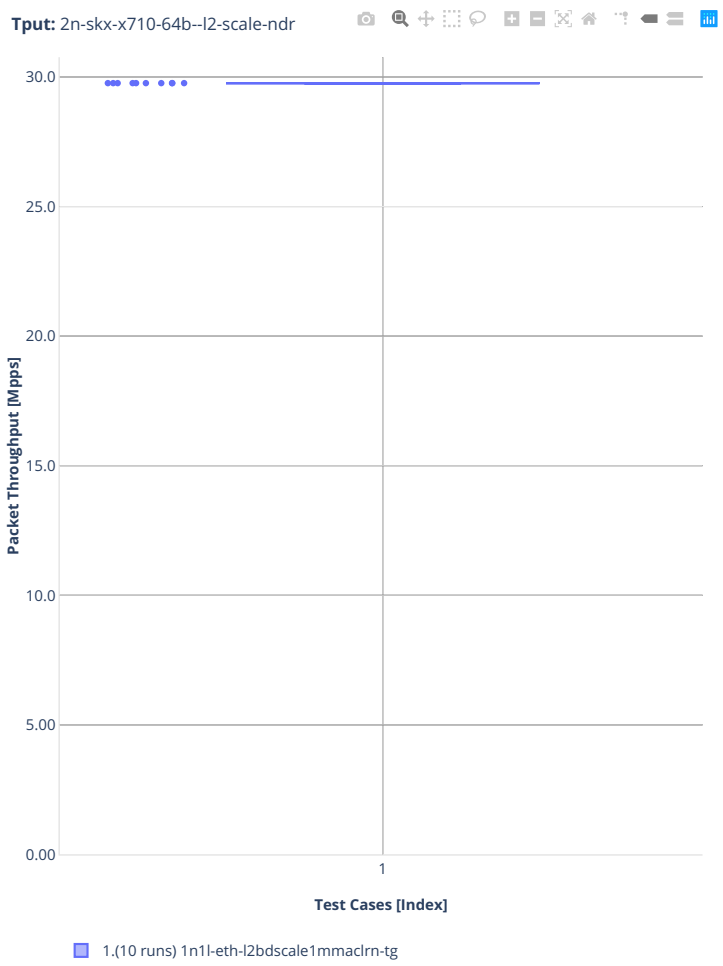

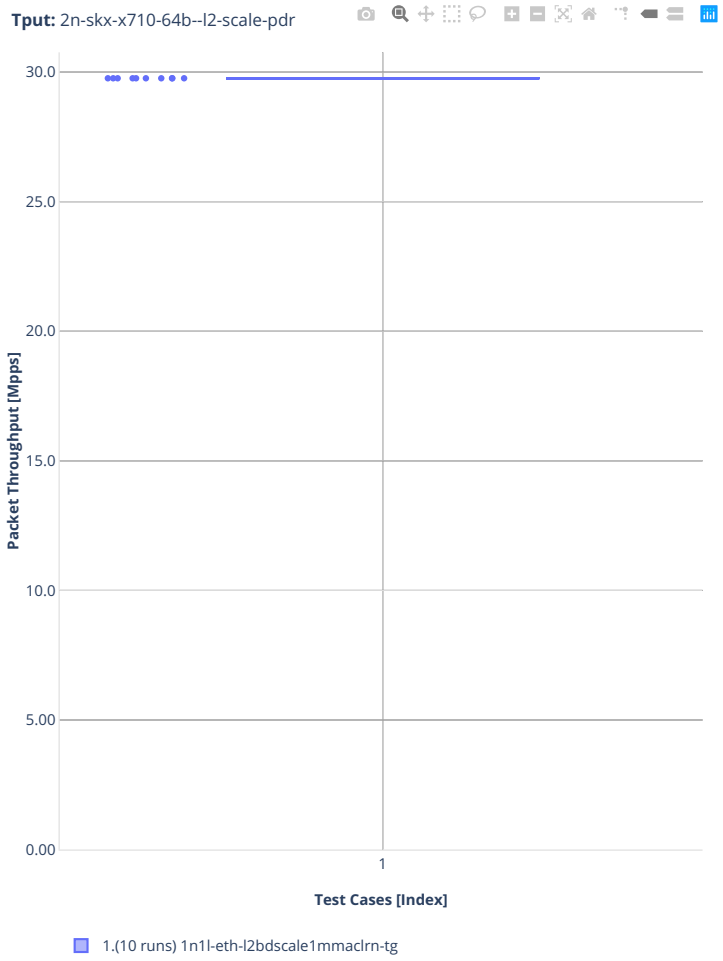
Tput: 2n-skx-x710-64b--l2-scale-pdr
30.0
25.0
20.0
15.0
10.0
5.00
0.00
Packet Throughput [Mpps]
1
Test Cases [Index]
1.(10 runs) 1n1l-eth-l2bdscale1mmaclrn-tg

4.4 Packet Latency

TRex latency results are generated based on the test data obtained from CSIT-2206 NDR-PDR throughput tests executed across physical testbeds hosted in LF FD.io labs: 2n-skx.

Latency by percentile distribution plots are used to show packet latency percentiles at different packet rate load levels: i) No-Load latency streams only, ii) Low-Load at 10% PDR, iii) Mid-Load at 50% PDR and iv) High-Load at 90% PDR.

For more details, see *Packet Latency* (page 45).

Additional information about graph data:

1. **Graph Title**: describes tested DUT packet path.
2. **X-axis Labels**: percentile of packets.
3. **Y-axis Labels**: measured one-way packet latency values in [uSec].
4. **Graph Legend**: list of latency tests at different packet rate load level.
5. **Hover Information**: packet rate load level, stream direction (East-West, West-East), percentile, one-way latency.

Note: Test results are stored in build logs from FD.io trex performance job 2n-skx[413].

[413] https://s3-logs.fd.io/vex-yul-rot-jenkins-1/csit-trex-perf-report-iterative-2206-2n-skx

4.4.1 2n-skx-x710

CSIT source code for the test cases used for plots can be found in CSIT git repository[414].

[414] https://git.fd.io/csit/tree/tests/trex/perf?h=rls2206

64b-ip4routing-base-scale

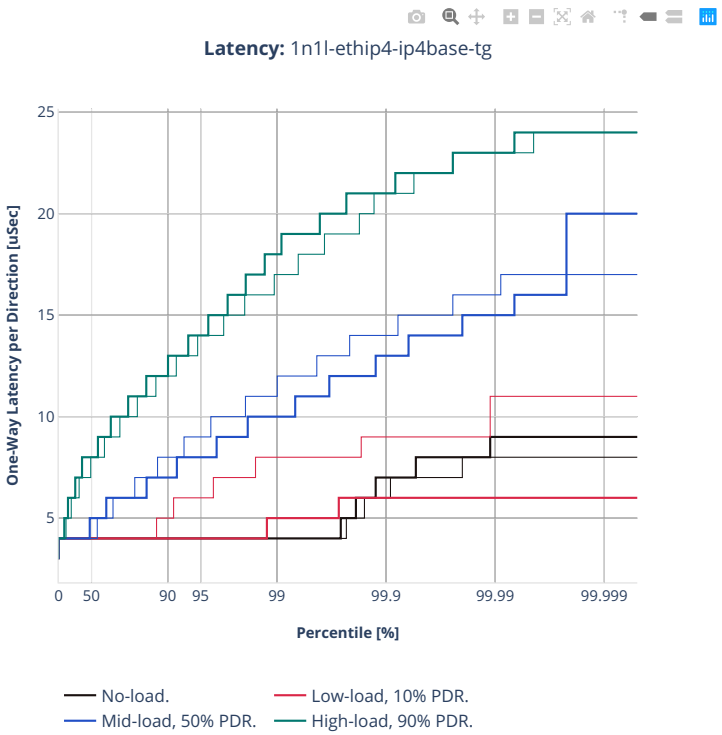

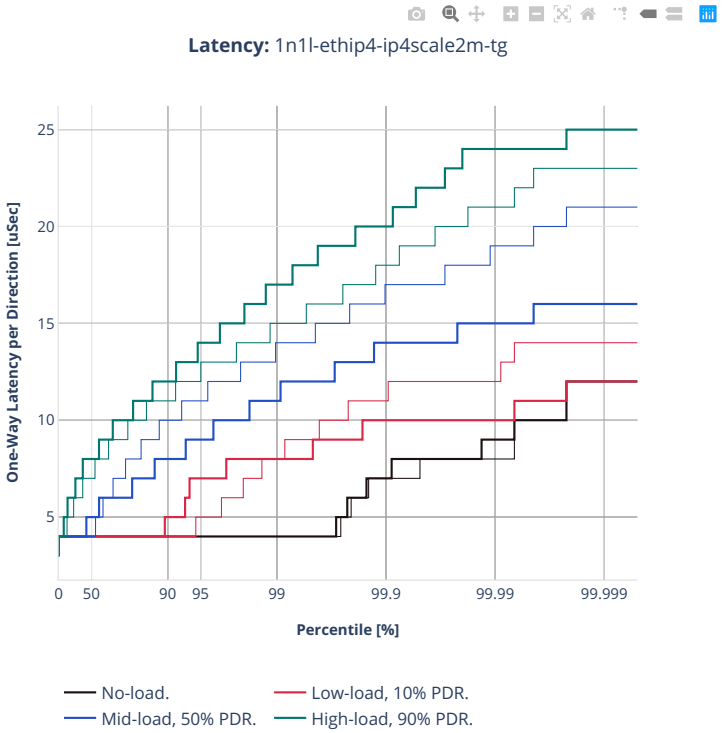
Latency: 1n1l-ethip4-ip4scale2m-tg
One-Way Latency per Direction [uSec]
25
20
15
10
5
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

78b-ip6routing-base-scale

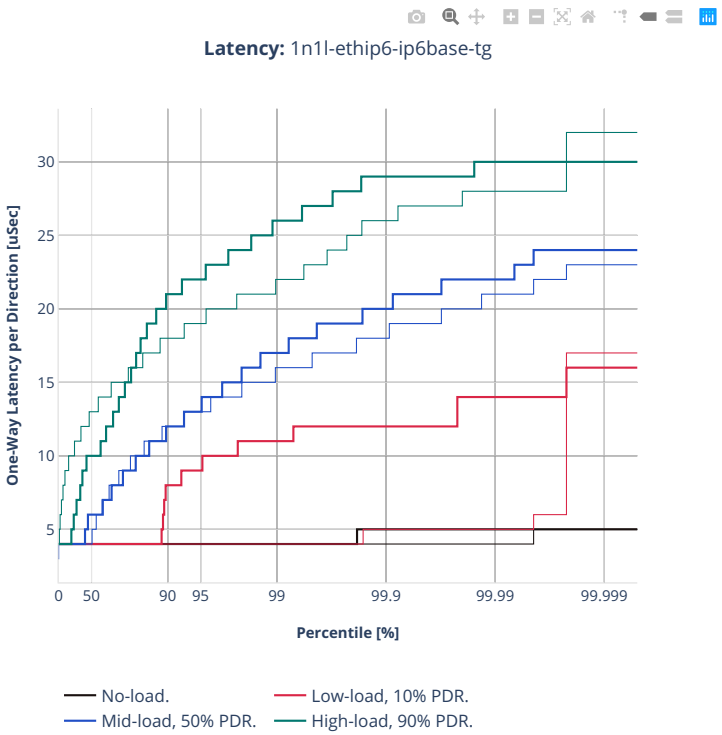

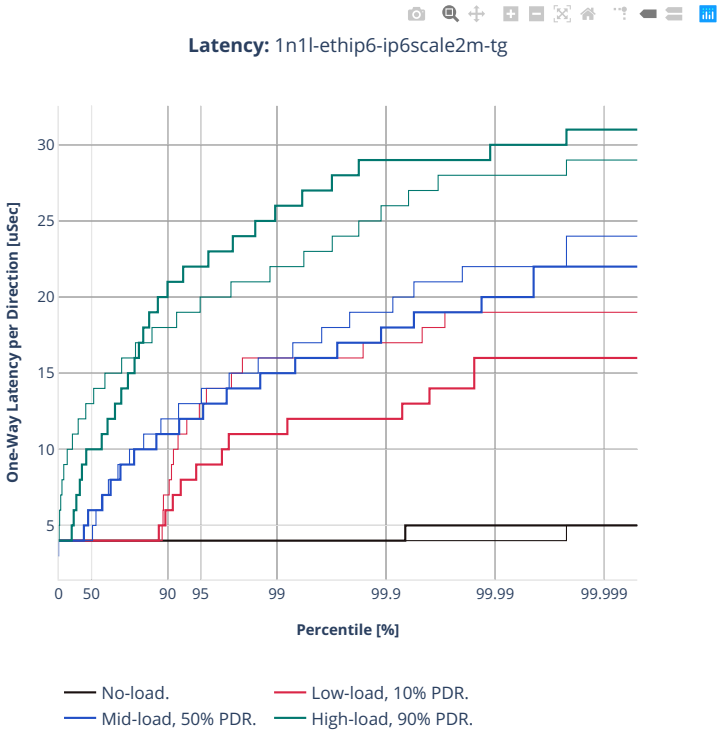
Latency: 1n1l-ethip6-ip6scale2m-tg
One-Way Latency per Direction [uSec]
30
25
20
15
10
5
0
50
90
95
99
99.9
99.99
99.999
Percentile [%]
No-load.
Low-load, 10% PDR.
Mid-load, 50% PDR.
High-load, 90% PDR.

64b-l2switching-scale

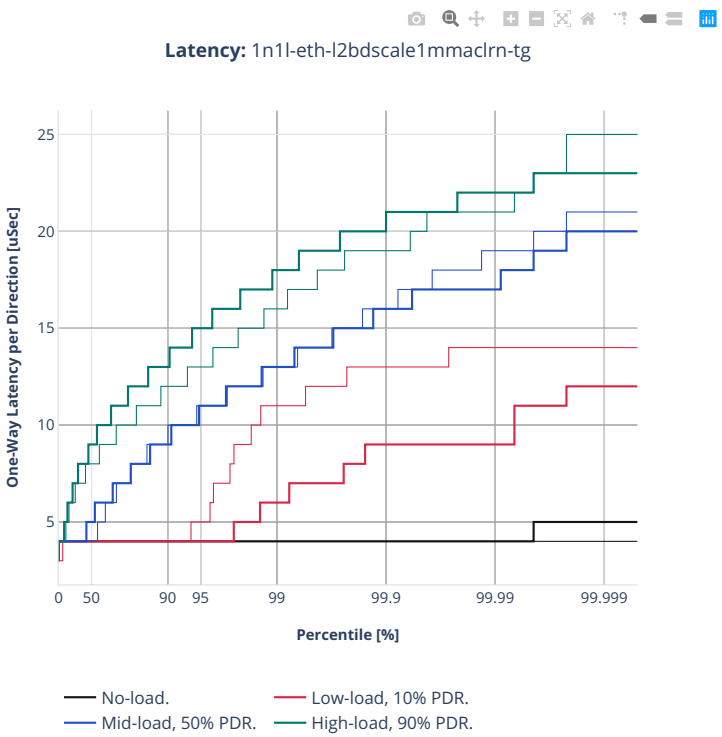

4.5 Throughput Trending

CSIT provides continuous performance trending for master branch:

1. TRex Trending Graphs[415]: per TRex test case throughput trend, trend compliance and summary of detected anomalies. We expect TRex to hit the curently used bps or pps limit, so no anomalies here (unless we change those limits in CSIT).
2. TRex Latency Graphs[416]: per TRex build NDRPDR latency measurements against the trendline. We have seen in past that the latency numbers can depend on TRex version, NIC firmware, or driver used.

[415] https://s3-docs.fd.io/csit/master/trending/ndrpdr_trending/trex.html
[416] https://s3-docs.fd.io/csit/master/trending/ndrpdr_latency_trending/trex.html

4.6 Test Environment

4.6.1 Environment Versioning

CSIT test environment versioning has been introduced to track modifications of the test environment.

Any benchmark anomalies (progressions, regressions) between releases of a DUT application (e.g. VPP, DPDK), are determined by testing it in the same test environment, to avoid test environment changes clouding the picture. To beter distinguish impact of test environment changes, we also execute tests without any SUT (just with TRex TG sending packets over a link looping back to TG).

A mirror approach is introduced to determine benchmarking anomalies due to the test environment change. This is achieved by testing the same DUT application version between releases of CSIT test system. This works under the assumption that the behaviour of the DUT is deterministic under the test conditions.

CSIT test environment versioning scheme ensures integrity of all the test system components, including their HW revisions, compiled SW code versions and SW source code, within a specific CSIT version. Components included in the CSIT environment versioning include:

- **HW** Server hardware firmware and BIOS (motherboard, processsor, NIC(s), accelerator card(s)), tracked in CSIT branch in `./docs/lab/<server_platform_name>_hw_bios_cfg.md`, e.g. Xeon Skylake servers[417].
- **Linux** Server Linux OS version and configuration, tracked in CSIT Reports in SUT Settings[418] and Pre-Test Server Calibration[419].
- **TRex** TRex Traffic Generator version, drivers and configuration tracked in TG Settings[420].
- **CSIT** CSIT framework code tracked in CSIT release branches.

Following is the list of CSIT versions to date:

- Ver. 1 associated with CSIT rls1908 branch (HW[421], Linux[422], TRex[423], CSIT[424]).
- Ver. 2 associated with CSIT rls2001 branch (HW[425], Linux[426], TRex[427], CSIT[428]).
- Ver. 4 associated with CSIT rls2005 branch (HW[429], Linux[430], TRex[431], CSIT[432]).
- Ver. 5 associated with CSIT rls2009 branch (HW[433], Linux[434], TRex[435], CSIT[436]).
  - The main change is TRex data-plane core resource adjustments: increase from 7 to 8 cores and pinning cores to interfaces[437] for better TRex performance with symmetric traffic profiles.

[417] https://git.fd.io/csit/tree/docs/lab/testbeds_sm_skx_hw_bios_cfg.md#n556
[418] https://s3-docs.fd.io/csit/master/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[419] https://s3-docs.fd.io/csit/master/report/vpp_performance_tests/test_environment.html#id21
[420] https://s3-docs.fd.io/csit/master/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[421] https://git.fd.io/csit/tree/docs/lab?h=rls1908
[422] https://docs.fd.io/csit/rls1908/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[423] https://docs.fd.io/csit/rls1908/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[424] https://git.fd.io/csit/tree/?h=rls1908
[425] https://git.fd.io/csit/tree/docs/lab?h=rls2001
[426] https://docs.fd.io/csit/rls2001/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[427] https://docs.fd.io/csit/rls2001/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[428] https://git.fd.io/csit/tree/?h=rls2001
[429] https://git.fd.io/csit/tree/docs/lab?h=rls2005
[430] https://docs.fd.io/csit/rls2005/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[431] https://docs.fd.io/csit/rls2005/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[432] https://git.fd.io/csit/tree/?h=rls2005
[433] https://git.fd.io/csit/tree/docs/lab?h=rls2009
[434] https://docs.fd.io/csit/rls2009/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[435] https://docs.fd.io/csit/rls2009/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[436] https://git.fd.io/csit/tree/?h=rls2009
[437] https://gerrit.fd.io/r/c/csit/+/28184

- Ver. 6 associated with CSIT rls2101 branch (HW[438], Linux[439], TRex[440], CSIT[441]).
  - The main change is TRex version upgrade: increase from 2.82 to 2.86[442].
- Ver. 7 associated with CSIT rls2106 branch (HW[443], Linux[444], TRex[445], CSIT[446]).
  - TRex version upgrade: increase from 2.86 to 2.88[447].
  - Ubuntu upgrade: upgrade from 18.04 LTS to 20.04.2 LTS[448].
- Ver. 8 associated with CSIT rls2110 branch (HW[449], Linux[450], TRex[451], CSIT[452]).
  - Intel NIC 700/800 series firmware upgrade based on DPDK compatibility matrix.
- Ver. 9 associated with CSIT rls2202 branch (HW[453], Linux[454], TRex[455], CSIT[456]).
  - Intel NIC 700/800 series firmware upgrade based on DPDK compatibility matrix.
- Ver. 10 associated with CSIT rls2206 branch (HW[457], Linux[458], TRex[459], CSIT[460]).
  - Intel NIC 700/800 series firmware upgrade based on DPDK compatibility matrix.
  - Mellanox 556A series firmware upgrade based on DPDK compatibility matrix.
  - Intel IceLake all core turbo frequency turned off. Current base frequency is 2.6GHz.

To identify performance changes due to TRex code development between previous and current TRex version, both have been tested in CSIT environment of latest version and compared against each other. All substantial progressions and regressions have been marked up with RCA analysis. See *Known Issues* (page 1700).

4.6.2 Physical Testbeds

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

- **1-Node Topology**: Consisting of TG with 1 NIC with 2 ports connected together - loopback connection.

[438] https://git.fd.io/csit/tree/docs/lab?h=rls2101
[439] https://docs.fd.io/csit/rls2101/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[440] https://docs.fd.io/csit/rls2101/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[441] https://git.fd.io/csit/tree/?h=rls2101
[442] https://gerrit.fd.io/r/c/csit/+/29980
[443] https://git.fd.io/csit/tree/docs/lab?h=rls2106
[444] https://s3-docs.fd.io/csit/rls2106/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[445] https://s3-docs.fd.io/csit/rls2106/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[446] https://git.fd.io/csit/tree/?h=rls2106
[447] https://gerrit.fd.io/r/c/csit/+/31652
[448] https://gerrit.fd.io/r/c/csit/+/31290
[449] https://git.fd.io/csit/tree/docs/lab?h=rls2110
[450] https://s3-docs.fd.io/csit/rls2110/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[451] https://s3-docs.fd.io/csit/rls2110/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[452] https://git.fd.io/csit/tree/?h=rls2110
[453] https://git.fd.io/csit/tree/docs/lab?h=rls2202
[454] https://s3-docs.fd.io/csit/rls2202/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[455] https://s3-docs.fd.io/csit/rls2202/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[456] https://git.fd.io/csit/tree/?h=rls2202
[457] https://git.fd.io/csit/tree/docs/lab?h=rls2206
[458] https://s3-docs.fd.io/csit/rls2206/report/vpp_performance_tests/test_environment.html#sut-settings-linux
[459] https://s3-docs.fd.io/csit/rls2206/report/vpp_performance_tests/test_environment.html#tg-settings-trex
[460] https://git.fd.io/csit/tree/?h=rls2206

4.6.3 SUT Settings - TRex

4.6.4 TG Settings - TRex

TG Version

TRex v2.97

DPDK Version

DPDK v21.02

TG Installation

T-Rex installation is managed via Ansible role.

TG Startup Configuration

```
$ sudo -E -S sh -c 'cat << EOF > /etc/trex_cfg.yaml
- version: 2
  c: 8
  limit_memory: 8192
  interfaces: ["${pci1}","${pci2}"]
  port_info:
      - dest_mac: [${dest_mac1}]
        src_mac: [${src_mac1}]
      - dest_mac: [${dest_mac2}]
        src_mac: [${src_mac2}]
  platform :
      master_thread_id: 0
      latency_thread_id: 9
      dual_if:
          - socket: 0
            threads: [1, 2, 3, 4, 5, 6, 7, 8]
EOF'
```

TG Startup Command (Stateless Mode)

```
$ sudo -E -S sh -c "cd '${trex_install_dir}/scripts/' && \
  nohup ./t-rex-64 -i --prefix $(hostname) --hdrh --no-scapy-server \
  --mbuf-factor 32 > /tmp/trex.log 2>&1 &" > /dev/null
```

Also, Python client is now starting traffic with:

```
core_mask=STLClient.CORE_MASK_PIN
```

TG Startup Command (Stateful Mode)

```
$ sudo -E -S sh -c "cd '${trex_install_dir}/scripts/' && \
  nohup ./t-rex-64 -i --prefix $(hostname) --astf --hdrh --no-scapy-server \
  --mbuf-factor 32 > /tmp/trex.log 2>&1 &" > /dev/null
```

TG API Driver

TRex driver[461]

[461] https://git.fd.io/csit/tree/GPL/tools/trex/trex_stl_profile.py?h=rls2206

CHAPTER

FIVE

VPP DEVICE

5.1 Overview

5.1.1 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.

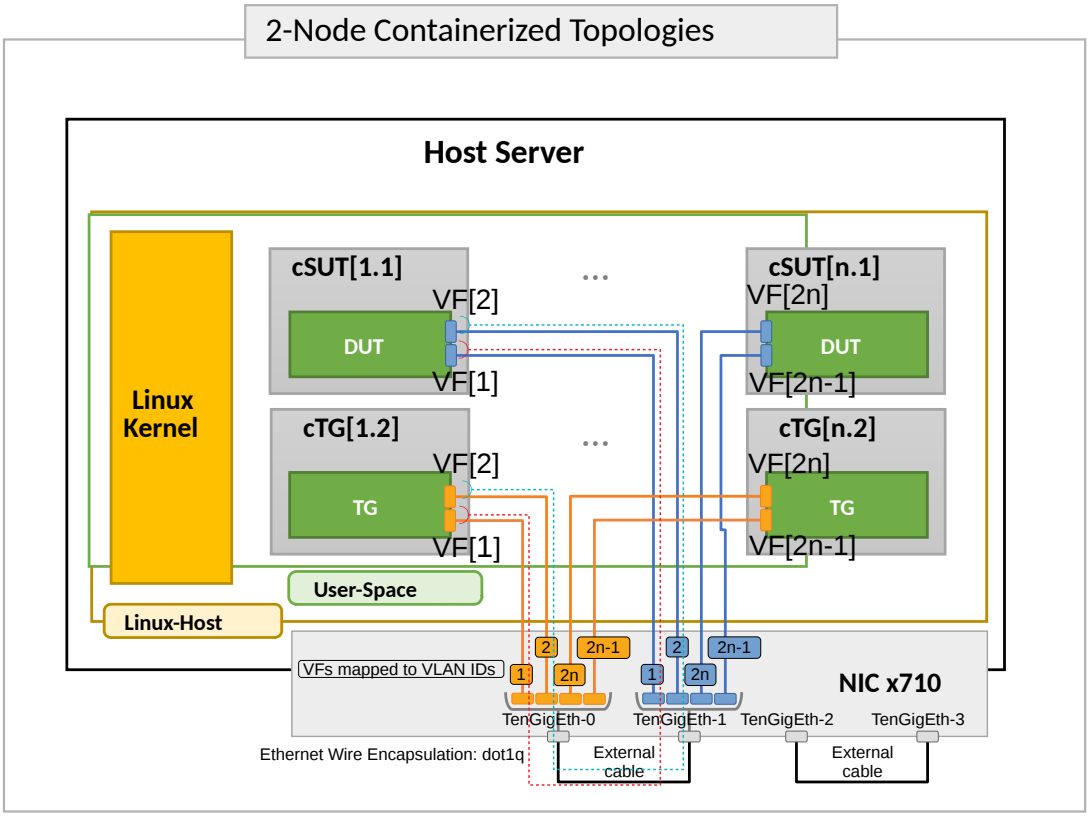

SUT1 is Docker Container (running Ubuntu, depending on the test suite), TG is a Traffic Generator (running Ubuntu Container). SUTs run VPP SW application in Linux user-mode as a Device Under Test (DUT) within the container. TG runs Scapy SW application as a packet Traffic Generator. Network connectivity 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

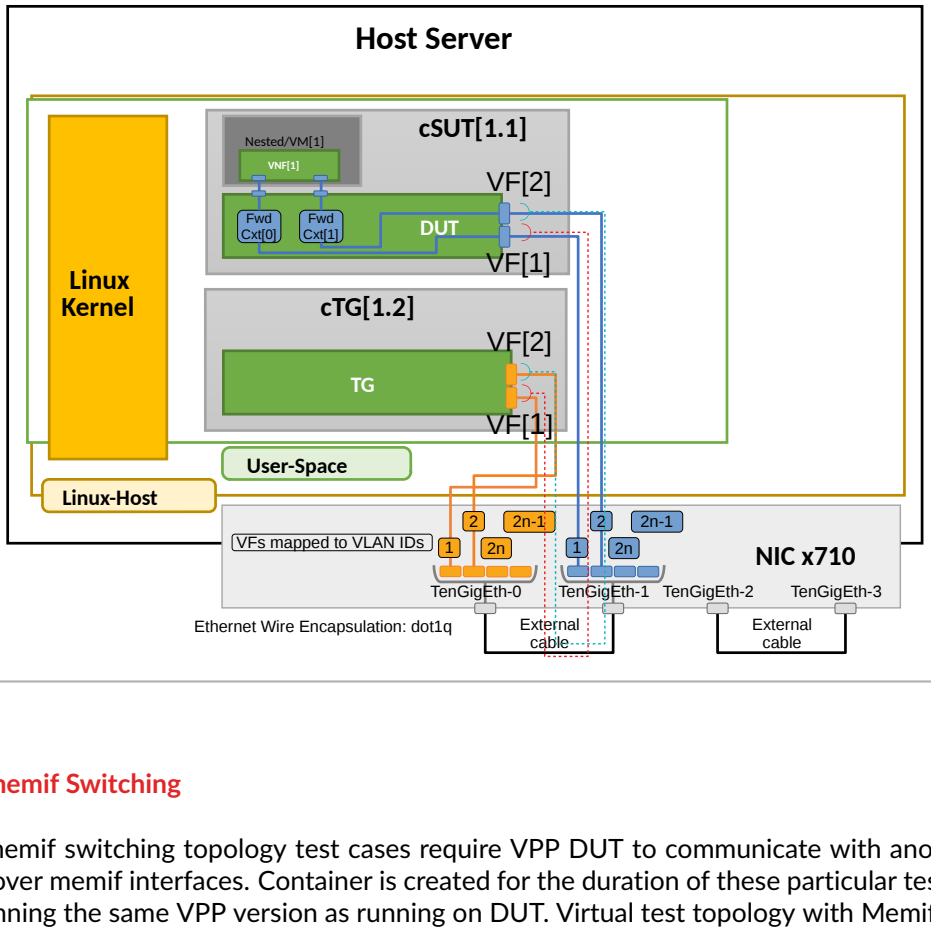

vfNIC-to-memif Switching

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.

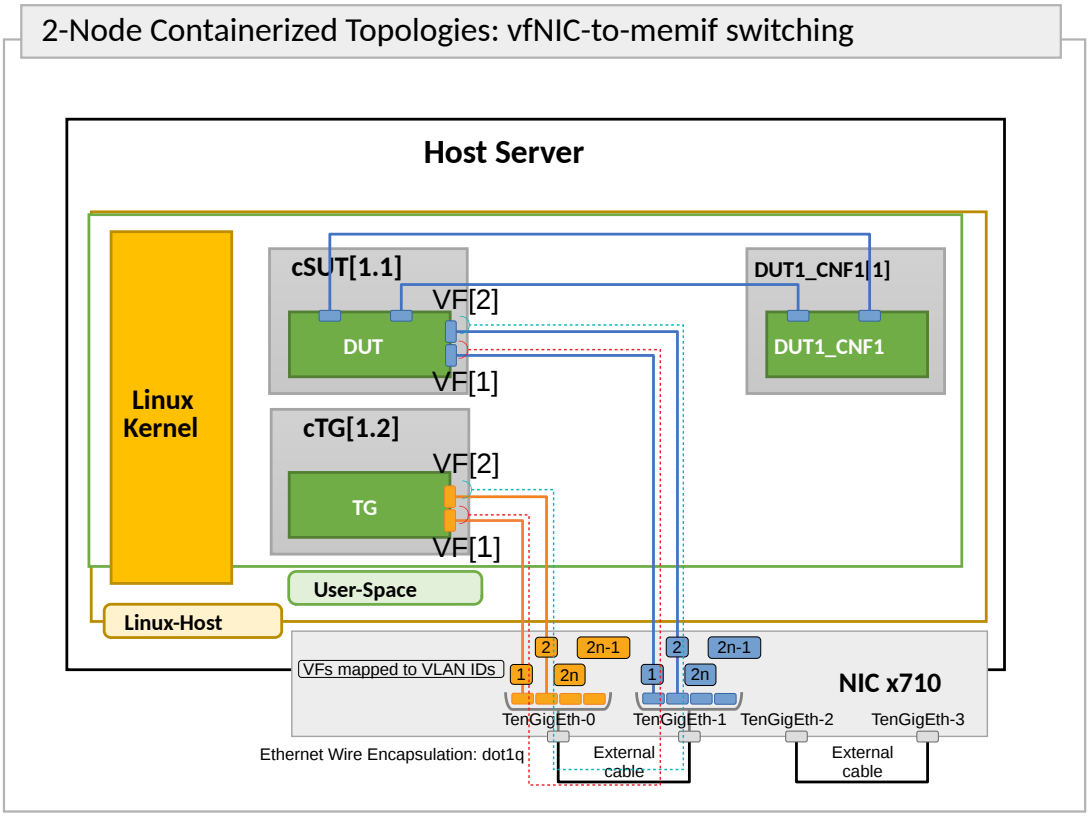

5.1.2 Functional Tests Coverage

CSIT-2206 includes following VPP functionality tested in VPP Device environment:

Functionality	Description
ACL (classify)	Ingress Access Control List security for L2 Bridge-Domain MAC switching, IPv4 routing, IPv6 routing.
ACL (acl_plugin)	Ingress and Egress Access Control List security in stateless and stateful mode for L2 Bridge-Domain MAC switching, IPv4 routing, IPv6 routing.
ADL	ADL address allow-list and block-list filtering for IPv4 and IPv6 routing.
GENEVE	GENEVE tunnels for IPv4 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.
MACIP (acl_plugin)	Ingress Access Control List security for L2 Bridge-Domain MAC switching based on mix of MAC and IP address matches.
Memif Interface	Baseline VPP memif interface tests.
NAT44	Network Address and Port Translation deterministic mode and endpoint-dependent mode tests for IPv4.
QoS Policer Metering	Ingress packet rate metering and marking for IPv4, IPv6.
SRv6	Segment routing over IPv6, base and proxy.
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.

5.1.3 Tests Naming

CSIT-2206 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 1743).

5.2 Release Notes

5.2.1 Changes in CSIT-2206

1. TEST FRAMEWORK
    - No updates

5.2.2 Known Issues

List of known issues in CSIT-2206 for VPP functional tests in VPP Device:

#	JiraID	Issue Description

5.3 Integration Tests

5.3.1 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**[462].

5.3.2 Overview

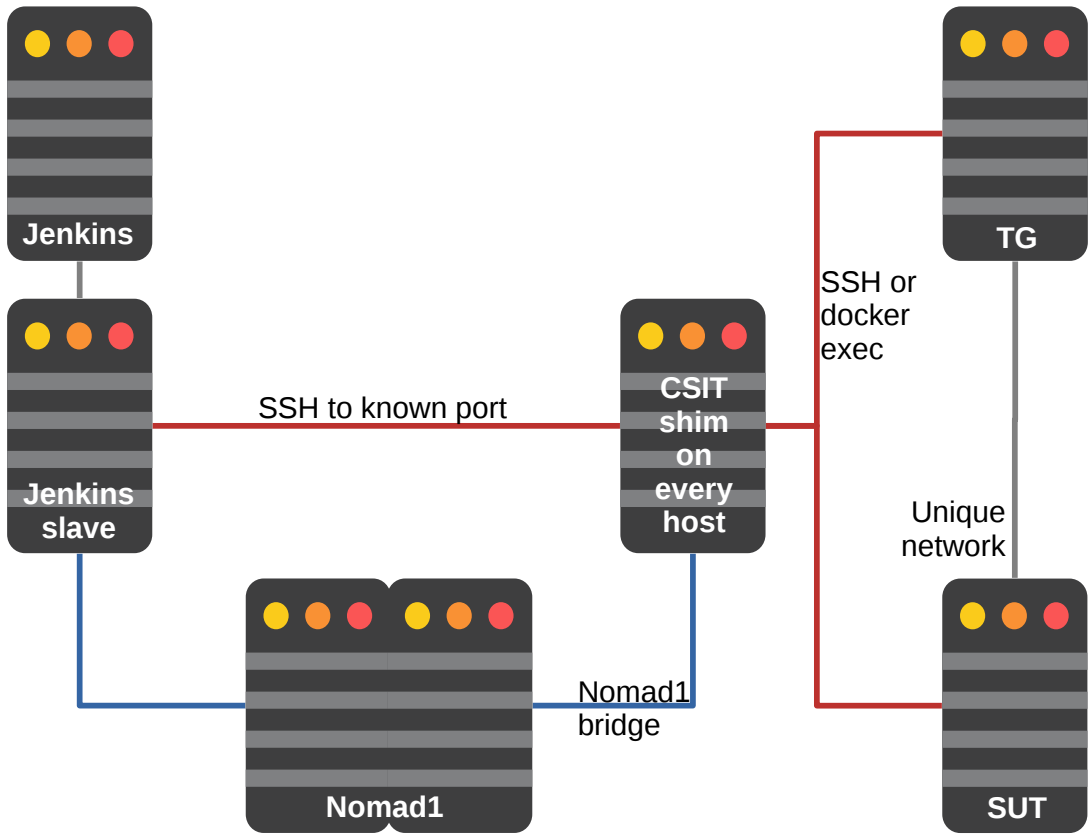

[462] https://tools.ietf.org/html/rfc8174.html

5.3.3 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.

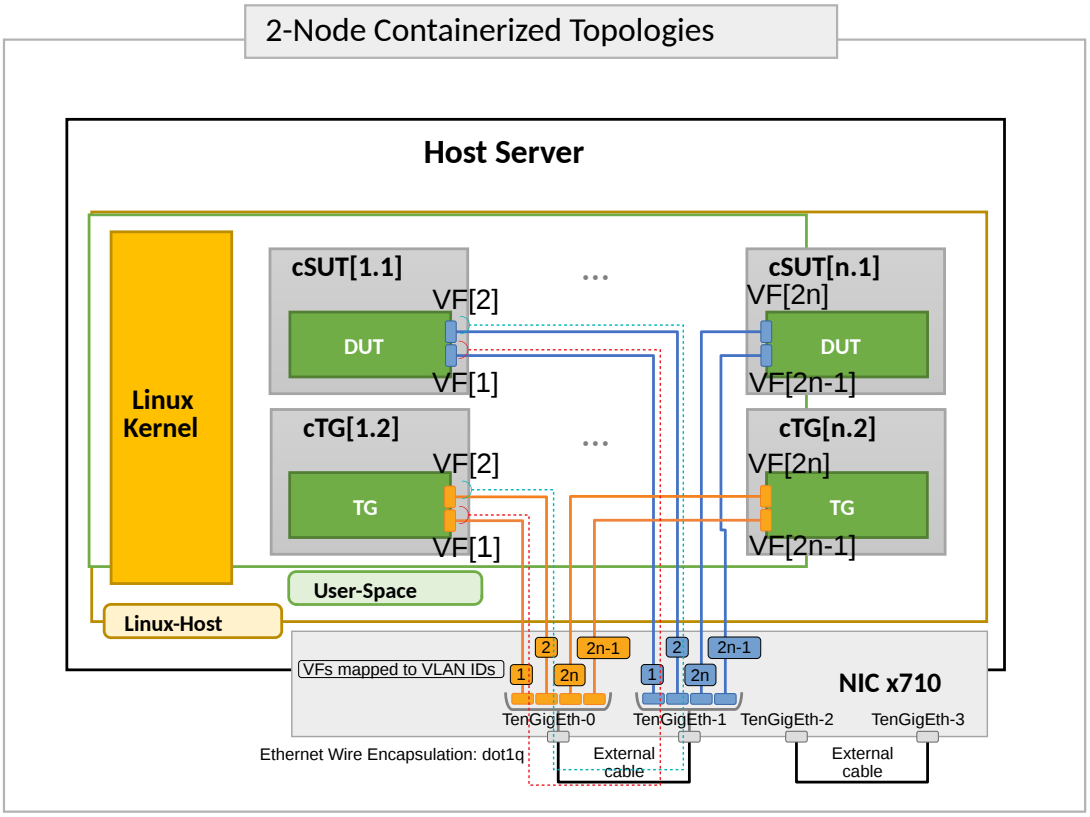

Server is populated with the following NIC models:

1. NIC-1: x710-da4 4p10GE Intel.
2. NIC-2: E810-2CQDA2 2p100GbE 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 vfio_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 Testbeds[463] 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 boostrap scripts). Example of Vagrant configuration:

```
Vagrant.configure(2) do |c|
  c.vm.network "private_network", type: "dhcp", auto_config: false,
      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.

5.3.4 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[464].

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[468]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[466]and[?] for automated build purposes and image hosting cannot be established with the trust and respectful to security of FD.io project. Unless adressed, DCR images will be placed in custom registry service[470]. Automated

[463] https://wiki.fd.io/view/CSIT/Testbeds:_Xeon_Skx,_Arm,_Atom.
[464] TWS[465]
[465] https://wiki.fd.io/view/CSIT/TWS
[468] FD.io/CSIT gerrit[469]
[469] https://gerrit.fd.io/r/CSIT
[466] Docker hub[467]
[467] https://hub.docker.com/
[470] FD.io regisy

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[?]. 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[471].
- *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[473].
- *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
```

[471] jenkins-slave-dcr-file[472]
[472] https://github.com/snergfdio/multivppcache/blob/master/ubuntu18/Dockerfile
[473] csit-shim-dcr-file[474]
[474] https://github.com/snergfdio/multivppcache/blob/master/csit-shim/Dockerfile

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[475].
- *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 bound interfaces. We cannot use
# --device=/dev/vfio as this does not see newly bound interfaces.
dcr_stc_params+="--volume /dev/vfio:/dev/vfio "
# Mount docker.sock to be able to use docker deamon 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 vfio work, access to `/dev/vfio` must be granted.

[475] csit-sut-dcr-file[476]

[476] https://github.com/snergfdio/multivppcache/blob/master/csit-sut/Dockerfile

5.3.5 Environment initialization

All 1-node servers are to be managed and provisioned via the[477] set of playbooks with *vpp-device* role. Full playbooks can be found under[479] 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 (Virtual Functions) 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 (useful for e.g. excluding management network interfaces).
- **pci_whitelist**: PCI addresses to be included for VFs initialization.

5.3.6 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. Mutal 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})
```

(continues on next page)

[477] ansible[478]
[478] https://www.ansible.com/
[479] Fd.io/CSIT ansible[480]
[480] https://git.fd.io/csit/tree/fdio.infra.ansible

(continued from previous page)

```
            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[483]. 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.

5.3.7 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.[481]

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.[?]

```
$ 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).[?]

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.[?]

- For promiscuous all:

[483] pci ids[484]
[484] http://pci-ids.ucw.cz/v2.2/pci.ids
[481] Intel i40e[482]
[482] https://downloadmirror.intel.com/26370/eng/readme.txt

```
$ 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.[?]

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 % (${#PCI_WHITELIST[@]}/2)))
                    fi
                    # 802.1Q base offset.
                    vlan_bs_off=1100
                    # 802.1Q PF PCI address offset.
                    vlan_pf_off=$(( 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})"
```

(continues on next page)

(continued from previous page)

```
                    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.

5.3.8 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.

5.3.9 Links

CHAPTER

SIX

CSIT FRAMEWORK

6.1 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.

6.1.1 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.

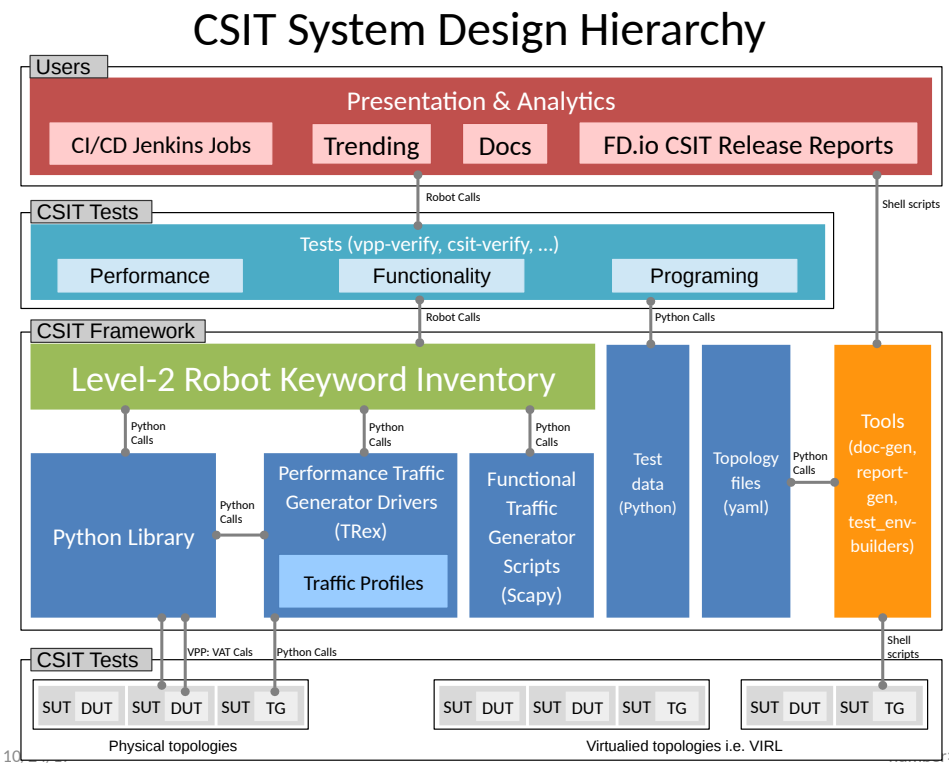

A brief bottom-up description is provided here:

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;
    - Performance tests using physical testbed environment:
        - VPP;
        - DPDK-Testpmd;
        - DPDK-L3Fwd;
    - Tools:
        - Documentation generator;
        - Report generator;

- Testbed environment setup ansible playbooks;
- Operational debugging scripts;

6.1.2 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 commmon 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;

6.1.3 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;

6.1.4 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.

6.1.5 Coding Guidelines

Coding guidelines can be found on Design optimizations wiki page[485].

6.2 Test Naming

6.2.1 Background

CSIT-2206 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[486]. Below few illustrative examples of the naming usage for test suites across CSIT performance, functional and Honeycomb management test areas.

6.2.2 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

[485] https://wiki.fd.io/view/CSIT/Design_Optimizations
[486] https://wiki.fd.io/view/CSIT/csit-test-naming

- 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.

6.2.3 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.
   - *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-apihc-apivat-func* => configuration of L2 Bridge-Domain with Honeycomb API and VAT API calls, functional tests.
   - *mgmt-oper-int-apihcnc-func* => reading status and operational data of interface with Honeycomb NetConf API calls, functional tests.
   - *mgmt-cfg-int-tap-apihcnc-func* => configuration of tap interfaces with Honeycomb NetConf API calls, functional tests.
   - *mgmt-notif-int-subint-apihcnc-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[487].

6.3 Presentation and Analytics

6.3.1 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[488].
  - 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

[487] https://wiki.fd.io/view/CSIT/csit-test-naming
[488] https://pandas.pydata.org/pandas-docs/stable/generated/pandas.Series.html

- versions: minimal, full (TODO: define the names and scope of versions)

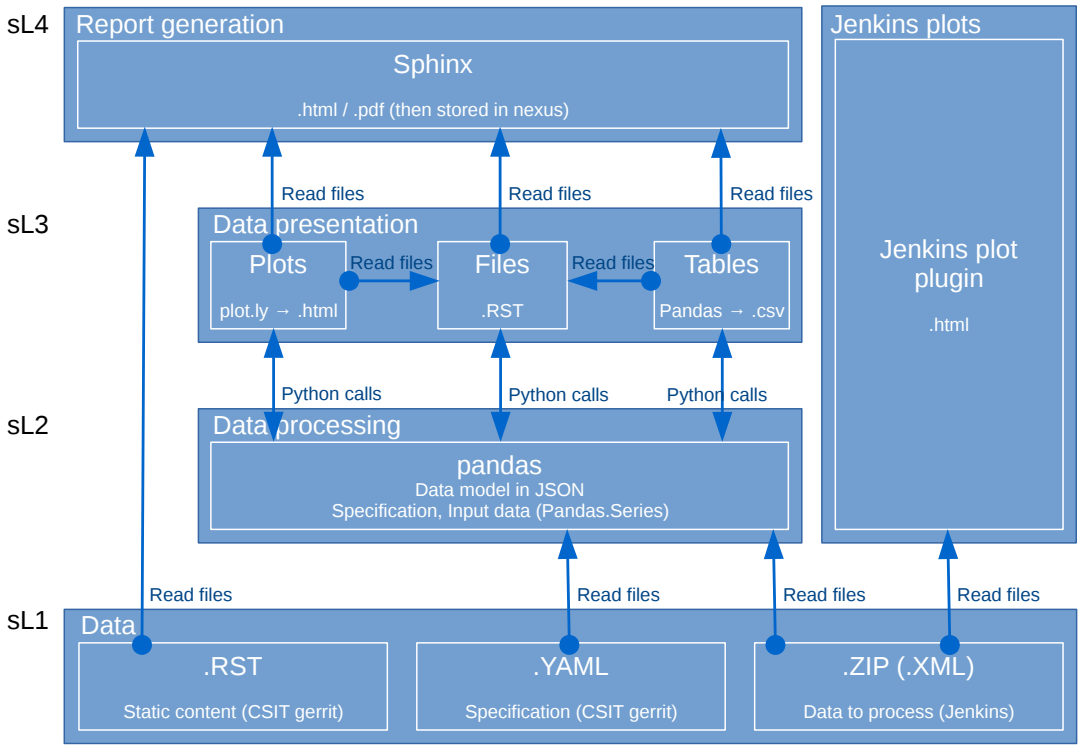

6.3.2 Data

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"
-
```

(continues on next page)

(continued from previous page)

```
  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)
```

(continues on next page)

(continued from previous page)

```
    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,FUNC,VPP]: "{DIR[DTR]}/vpp_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"
    DIR[CSS_PATCH_FILE2]: "{DIR[WORKING,SRC,STATIC]}/theme_overrides.css"

  urls:
    URL[JENKINS,CSIT]: "https://jenkins.fd.io/view/csit/job"
    URL[S3_STORAGE,LOG]: "https://logs.nginx.service.consul/vex-yul-rot-jenkins-1"
    URL[NEXUS,LOG]: "https://logs.fd.io/production/vex-yul-rot-jenkins-1"
    URL[NEXUS,DOC]: "https://docs.fd.io/csit"
    DIR[NEXUS,DOC]: "report/_static/archive"

  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
      - 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
```

(continues on next page)

(continued from previous page)

```
      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-sel1-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:
  - "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-vpp-functional-1707-ubuntu1604-virl:
    -
      build: lastSuccessfulBuild
      file: "csit-vpp-functional-1707-ubuntu1604-virl-lastSuccessfulBuild.xml"
    hc2vpp-csit-integration-1707-ubuntu1604:
    -
      build: lastSuccessfulBuild
      file: "hc2vpp-csit-integration-1707-ubuntu1604-lastSuccessfulBuild.xml"
    csit-vpp-perf-1707-all:
    -
      build: 16
      file: "csit-vpp-perf-1707-all__16__output.xml"
    -
      build: 17
      file: "csit-vpp-perf-1707-all__17__output.xml"
```

Section: Static

This section defines the static content which is stored in git and will be used as a source to generate the report.

This section has these parts:

- type: “static” - says that this section is the “static”.
- src-path - path to the static content.
- dst-path - destination path where the static content is copied and then processed.

```
-
  type: "static"
  src-path: "{DIR[RST]}"
  dst-path: "{DIR[WORKING,SRC]}"
```

Section: Input

This section defines the data used to generate elements. It is mandatory if the debug mode is not used.

This section has the following parts:

- type: “input” - says that this section is the “input”.
- general - parameters common to all builds:
  - file-name: file to be downloaded.
  - file-format: format of the downloaded file, “.zip” or “.xml” are supported.

  - download-path: path to be added to url pointing to the file, e.g.: "{job}/{build}/robot/report/*zip*/{filename}"; {job}, {build} and {filename} are replaced by proper values defined in this section.
  - extract: file to be extracted from downloaded zip file, e.g.: "output.xml"; if xml file is downloaded, this parameter is ignored.
- builds - list of jobs (keys) and numbers of builds which output data will be downloaded.

The structure of the section "Input" is as follows (example from 17.07 report):

```
-
  type: "input"  # Ignored in debug mode
  general:
    file-name: "robot-plugin.zip"
    file-format: ".zip"
    download-path: "{job}/{build}/robot/report/*zip*/{filename}"
    extract: "robot-plugin/output.xml"
  builds:
    csit-vpp-perf-1707-all:
    - 9
    - 10
    - 13
    - 14
    - 15
    - 16
    - 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-virl:
    - lastSuccessfulBuild
    hc2vpp-csit-perf-master-ubuntu1604:
    - 8
    - 9
    hc2vpp-csit-integration-1707-ubuntu1604:
    - lastSuccessfulBuild
```

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]}/tmpl_performance_improvements.csv"
  output-file-ext: ".csv"
  output-file: "{DIR[DTR,PERF,VPP,IMPRV]}/performance_improvements"
  columns:
  -
    title: "VPP Functionality"
    data: "template 1"
  -
    title: "Test Name"
    data: "template 2"
  -
    title: "VPP-16.09 mean [Mpps]"
    data: "template 3"
  -
    title: "VPP-17.01 mean [Mpps]"
    data: "template 4"
  -
    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"
```

(continues on next page)

(continued from previous page)

```
    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|l2bdbasemaclrn)-ndrdisc"
  algorithm: "plot_performance_box"
  output-file-type: ".html"
  output-file: "{DIR[STATIC,VPP]}/64B-1t1c-l2-sel1-ndrdisc"
  data:
    csit-vpp-perf-1707-all:
    - 9
    - 10
    - 13
    - 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|l2bdbasemaclrn)-ndrdisc"
    xaxis:
      autorange: True
      autotick: False
      fixedrange: False
      gridcolor: "rgb(238, 238, 238)"
      linecolor: "rgb(238, 238, 238)"
      linewidth: 1
      showgrid: True
      showline: True
      showticklabels: True
      tickcolor: "rgb(238, 238, 238)"
      tickmode: "linear"
      title: "Indexed Test Cases"
      zeroline: False
    yaxis:
      gridcolor: "rgb(238, 238, 238)'"
      hoverformat: ".4s"
      linecolor: "rgb(238, 238, 238)"
      linewidth: 1
      range: []
      showgrid: True
      showline: True
      showticklabels: True
      tickcolor: "rgb(238, 238, 238)"
      title: "Packets Per Second [pps]"
      zeroline: False
    boxmode: "group"
```

(continues on next page)

(continued from previous page)

```
    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
```

(continues on next page)

(continued from previous page)

```
    range: []
    showgrid: True
    showline: True
    showticklabels: True
    tickcolor: "rgb(238, 238, 238)"
    title: "Latency min/avg/max [uSec]"
    zeroline: False
  boxmode: "group"
  boxgroupgap: 0.5
  autosize: False
  margin:
    t: 50
    b: 20
    l: 50
    r: 20
  showlegend: True
  legend:
    orientation: "h"
  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_perf_box"
  output-file-type: ".html"
  output-file: "{DIR[STATIC,VPP]}/http-server-performance-cps"
  data:
    "plot-vpp-httlp-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\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.

6.3.3 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
    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"
    }
```

(continues on next page)

(continued from previous page)

```
    "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": {
                "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"
```

(continues on next page)

(continued from previous page)

```
    },
    "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
N_core_throughput_speedup = ------------------
                            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"
```

(continues on next page)

(continued from previous page)

```
  - "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 repeatibility 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_perf_comparison"
  output-file-ext: ".csv"
  output-file: "{DIR[DTR,PERF,VPP,IMPRV]}/vpp_performance_comparison"
```

(continues on next page)

(continued from previous page)

```
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.

6.3.4 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.ly[489] 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.

[489] https://plot.ly/

6.3.5 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.

6.3.6 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:
    1. By PT job at its completion.
    2. Manually from Jenkins UI.
2. Download and parse archived historical data and the new data:
    1. New data from latest PT job is evaluated against the rolling window of <N> sets of historical data.
    2. Download RF output.xml files and compressed archived data.
    3. 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:
    1. Calculate quartiles Q1, Q2, Q3.
    2. Trim outliers using IQR.
    3. Calculate TMA and TMSD.
    4. Calculate normal trending range per test case based on TMA and TMSD.
4. Evaluate new test data against trend metrics:
    1. If within the range of (TMA +/- 3*TMSD) => Result = Pass, Reason = Normal.

2. If below the range => Result = Fail, Reason = Regression.
3. If above the range => Result = Pass, Reason = Progression.

5. Generate and publish results
    1. Relay evaluation result to job result.
    2. Generate a new set of trend analysis summary graphs and drill-down graphs.
        1. Summary graphs to include measured values with Normal, Progression and Regression markers. MM shown in the background if possible.
        2. Drill-down graphs to include MM, TMA and TMSD.
    3. Publish trend analysis graphs in html format on https://s3-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 'L2XCFWD') and not 'VHOST' and not 'MEMIF'"
      parameters:
      - "result"
      layout: "plot-cpta-vpp"

    - title: "DPDK 4T4C IMIX MRR Trending"
      output-file-name: "dpdk-imix-4t4c-xl710"
      data: "plot-performance-trending-dpdk"
```

(continues on next page)

(continued from previous page)

```
    filter: "'NIC_Intel-XL710' and 'IMIX' and 'MRR' and '4T4C' and 'DPDK'"
    parameters:
    - "result"
    layout: "plot-cpta-dpdk"
```

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_perf_trending_dash"
  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
  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_perf_trending_dash_html"
  input-file: "{DIR[STATIC,VPP]}/performance-trending-dashboard-1t1c.csv"
  output-file: "{DIR[STATIC,VPP]}/performance-trending-dashboard-1t1c.rst"
```

6.3.7 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:
    1. Summary graphs to include measured values with Progression and Regression markers.
    2. 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
```

6.3.8 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
```

(continues on next page)

(continued from previous page)

```
            files
            static


input_data_parser.py

    class InputData

        Methods:
            read_data
            filter_data

        Getters:
            data
            metadata
            suites
            tests


environment.py

    Functions:
        clean_environment

    class Environment

        Methods:
            set_environment

        Getters:
            environment


input_data_files.py

    Functions:
        download_data_files
        unzip_files


generator_tables.py

    Functions:
        generate_tables

    Functions implementing algorithms to generate particular types of
    tables (called by the function "generate_tables"):
        table_details
        table_performance_improvements


generator_plots.py

    Functions:
        generate_plots

    Functions implementing algorithms to generate particular types of
    plots (called by the function "generate_plots"):
        plot_performance_box
        plot_latency_box
```

(continues on next page)

(continued from previous page)

```
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

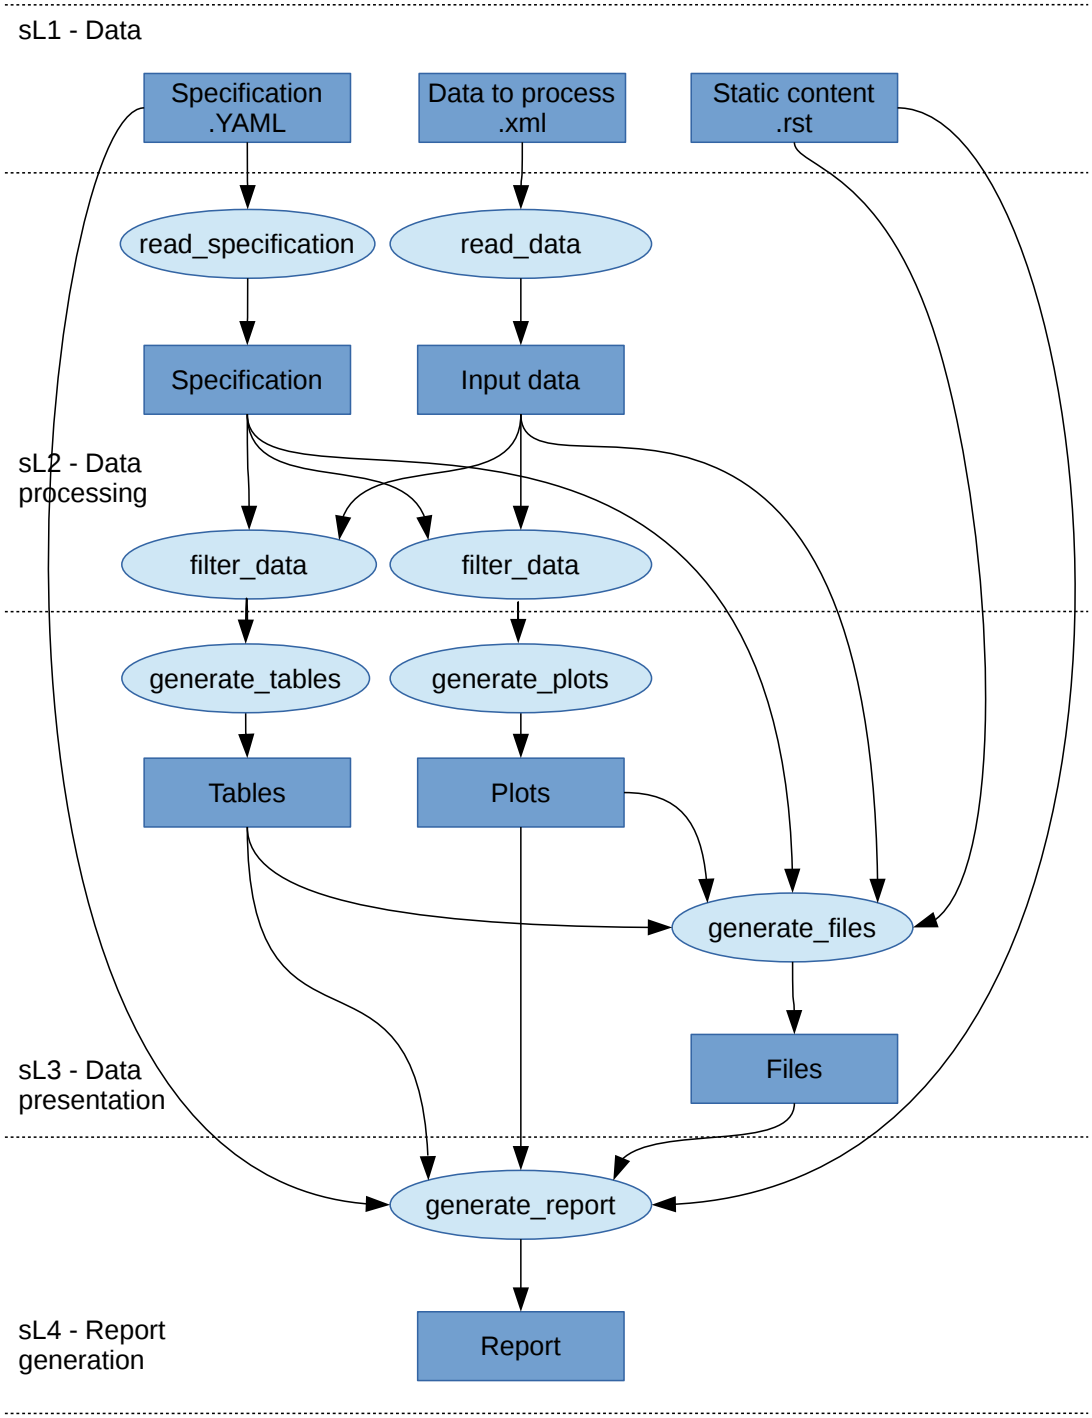

How to add an element

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.

6.4 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[490].

6.4.1 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.

6.4.2 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).

6.4.3 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.

[490] https://git.fd.io/csit/tree/docs/tag_documentation.rst?h=rls2206

6.4.4 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.

NIC_Amazon-Nitro-50G

Amazon EC2 ENA NIC.

6.4.5 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_{u}

Traffic flow with {u} unique IPs (users) in one direction. {u}=(1,10,100,1000,2000,4000).

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.

HOSTS_{h}

Stateless or stateful traffic stream with {h} client source IP4 addresses, usually with 63 flow differing in source port number. Could be UDP or TCP. If NAT is used, the clients are inside. Outside IP range can differ. {h}=(1024,4096,16384,65536,262144).

GENEVE4_{t}TUN

Test with {t} GENEVE IPv4 tunnel. {t}=(1,4,16,64,256,1024)

6.4.6 Test Category Tags

DEVICETEST

All vpp_device functional test cases.

PERFTEST

All performance test cases.

6.4.7 VPP Device Type Tags

SCAPY

All test cases that uses Scapy for packet generation and validation.

6.4.8 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.

6.4.9 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 ${overhead} value.

{b}B

{b} Bytes frames used for test.

IMIX

IMIX frame sequence (28x 64B, 16x 570B, 4x 1518B) used for test.

6.4.10 Test Type Tags

BASE

Baseline test cases, no encapsulation, no feature(s) configured in tests. No scaling whatsoever, beyond minimum needed for RSS.

IP4BASE

IPv4 baseline test cases, no encapsulation, no feature(s) configured in tests. Minimal number of routes. Other quantities may be scaled.

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. Other tags specify which quantities are scaled. Also applies if scaling is set on TG only (e.g. DUT works as IP4BASE).

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).

UDP

Tests which use any kind of UDP traffic (STL or ASTF profile).

TCP

Tests which use any kind of TCP traffic (STL or ASTF profile).

TREX

Tests which test trex traffic without any software DUTs in the traffic path.

UDP_UDIR

Tests which use unidirectional UDP traffic (STL profile only).

UDP_BIDIR

Tests which use bidirectional UDP traffic (STL profile only).

UDP_CPS

Tests which measure connections per second on minimal UDP pseudoconnections. This implies ASTF traffic profile is used. This tag selects specific output processing in PAL.

TCP_CPS

Tests which measure connections per second on empty TCP connections. This implies ASTF traffic profile is used. This tag selects specific output processing in PAL.

TCP_RPS

Tests which measure requests per second on empty TCP connections. This implies ASTF traffic profile is used. This tag selects specific output processing in PAL.

UDP_PPS

Tests which measure packets per second on lightweight UDP transactions. This implies ASTF traffic profile is used. This tag selects specific output processing in PAL.

TCP_PPS

Tests which measure packets per second on lightweight TCP transactions. This implies ASTF traffic profile is used. This tag selects specific output processing in PAL.

HTTP

Tests which use traffic formed of valid HTTP requests (and responses).

LDP_NGINX

LDP NGINX is un-modified NGINX with VPP via LD_PRELOAD.

NF_DENSITY

Performance tests that measure throughput of multiple VNF and CNF service topologies at different service densities.

6.4.11 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}DCRt}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.

6.4.12 Forwarding Mode Tags

L2BDMACSTAT

VPP L2 bridge-domain, L2 MAC static.

L2BDMACLRN

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.

LOADBALANCER_MAGLEV

VPP Load balancer maglev mode.

LOADBALANCER_L3DSR

VPP Load balancer l3dsr mode.

LOADBALANCER_NAT4

VPP Load balancer nat4 mode.

N2N

Mode, where NICs from the same physical server are directly connected with a cable.

6.4.13 Underlay Tags

IP4UNRLAY

IPv4 underlay.

IP6UNRLAY

IPv6 underlay.

MPLSUNRLAY

MPLS underlay.

6.4.14 Overlay Tags

L2OVRLAY

L2 overlay.

IP4OVRLAY

IPv4 overlay (IPv4 payload).

IP6OVRLAY

IPv6 overlay (IPv6 payload).

6.4.15 Tagging Tags

DOT1Q

All test cases with dot1q.

DOT1AD

All test cases with dot1ad.

6.4.16 Encapsulation Tags

ETH

All test cases with base Ethernet (no encapsulation).

LISP

All test cases with LISP.

LISPGPE

All test cases with LISP-GPE.

LISP_IP4o4

All test cases with LISP_IP4o4.

LISPGPE_IP4o4

All test cases with LISPGPE_IP4o4.

LISPGPE_IP6o4

All test cases with LISPGPE_IP6o4.

LISPGPE_IP4o6

All test cases with LISPGPE_IP4o6.

LISPGPE_IP6o6

All test cases with LISPGPE_IP6o6.

VXLAN

All test cases with Vxlan.

VXLANGPE

All test cases with VXLAN-GPE.

GRE

All test cases with GRE.

GTPU

All test cases with GTPU.

GTPU_HWACCEL

All test cases with GTPU_HWACCEL.

IPSEC

All test cases with IPSEC.

WIREGUARD

All test cases with WIREGUARD.

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.

GENEVE

All test cases with GENEVE.

GENEVE_L3MODE

All test cases with GENEVE tunnel in L3 mode.

FLOW

All test cases with FLOW.

FLOW_DIR

All test cases with FLOW_DIR.

FLOW_RSS

All test cases with FLOW_RSS.

NTUPLE

All test cases with NTUPLE.

L2TPV3

All test cases with L2TPV3.

6.4.17 Interface Tags

PHY

All test cases which use physical interface(s).

GSO

All test cases which uses Generic Segmentation Offload.

VHOST

All test cases which uses VHOST.

VHOST_1024

All test cases which uses VHOST DPDK driver with qemu queue size set to 1024.

VIRTIO

All test cases which uses VIRTIO native VPP driver.

VIRTIO_1024

All test cases which uses VIRTIO native VPP driver 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_DPDK

All test cases which uses DPDK 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_{n}L

All test cases which use {n} link(s) for link bonding.

DRV_{d}

All test cases which NIC Driver for DUT is set to {d}. Default is VFIO_PCI. {d}=(AVF, RDMA_CORE, VFIO_PCI, AF_XDP).

TG_DRV_{d}

All test cases which NIC Driver for TG is set to {d}. Default is IGB_UIO. {d}=(RDMA_CORE, IGB_UIO).

RXQ_SIZE_{n}

All test cases which RXQ size (RX descriptors) are set to {n}. Default is 0, which means VPP (API) default.

TXQ_SIZE_{n}

All test cases which TXQ size (TX descriptors) are set to {n}. Default is 0, which means VPP (API) default.

6.4.18 Feature Tags

IACLDST

iACL destination.

ADLALWLIST

ADL allowlist.

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.

6.4.19 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.

SCHEDULER

IPSec using crypto sw scheduler engine.

6.4.20 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.

6.4.21 Container Orchestration Tags

{n}VSWITCH

{n} VPP running in {n} Docker container(s) acting as a VSWITCH. {n}=(1).

{n}VNF

{n} VPP running in {n} Docker container(s) acting as a VNF work load. {n}=(1).

6.4.22 Multi-Threading Tags

STHREAD

Dynamic tag. All test cases using single poll mode thread.

MTHREAD

***Dynamic tag*.** All test cases using more then one poll mode driver thread.

{n}NUMA

All test cases with packet processing on {n} socket(s). {n}=(1,2).

{c}C

{c} worker thread pinned to {c} dedicated physical core; or if HyperThreading is enabled, {c}*2 worker threads each pinned to a separate logical core within 1 dedicated physical core. Main thread pinned to core 1. {t}=(1,2,4).

{t}T{c}C

***Dynamic tag*.** {t} worker threads pinned to {c} dedicated physical cores. Main thread pinned to core 1. By default CSIT is configuring same amount of receive queues per interface as worker threads. {t}=(1,2,4,8), {t}=(1,2,4).

BIBLIOGRAPHY

[lxc] Linux Containers[38]

[lxcnamespace] Resource management: Linux kernel Namespaces and cgroups[39].

[stgraber] LXC 1.0: Blog post series[40].

[lxcsecurity] Linux Containers Security[41].

[capabilities] Linux manual - capabilities - overview of Linux capabilities[42].

[cgroup1] Linux kernel documentation: cgroups[43].

[cgroup2] Linux kernel documentation: Control Group v2[44].

[selinux] SELinux Project Wiki[45].

[lxcsecfeatures] LXC 1.0: Security features[46].

[lxcsource] Linux Containers source[47].

[apparmor] Ubuntu AppArmor[48].

[seccomp] SECure COMPuting with filters[49].

[docker] Docker[50].

[k8sdoc] Kubernetes documentation[51].

[38] https://linuxcontainers.org/
[39] https://www.cs.ucsb.edu/~rich/class/cs293b-cloud/papers/lxc-namespace.pdf
[40] https://stgraber.org/2013/12/20/lxc-1-0-blog-post-series/
[41] https://linuxcontainers.org/lxc/security/
[42] http://man7.org/linux/man-pages/man7/capabilities.7.html
[43] https://www.kernel.org/doc/Documentation/cgroup-v1/cgroups.txt
[44] https://www.kernel.org/doc/Documentation/cgroup-v2.txt
[45] http://selinuxproject.org/page/Main_Page
[46] https://stgraber.org/2014/01/01/lxc-1-0-security-features/
[47] https://github.com/lxc/lxc
[48] https://wiki.ubuntu.com/AppArmor
[49] https://www.kernel.org/doc/Documentation/prctl/seccomp_filter.txt
[50] https://www.docker.com/what-docker
[51] https://kubernetes.io/docs/home/