# LF NETWORKING

## LFN Developer & Testing Forum

LF NETWORKING

LFN Developer & Testing Forum

# FD.io CSIT Performance Benchmarking with Anomaly Detection and Results Tracking

Presented by CSIT Committers: Tibor Frank, Maciek Konstantynowicz, Peter Mikus, Vratko Polák

csit-dev@lists.fd.io

FD.io

# Topics



- FD.io Overview
  - VPP and CSIT core projects (10 min)
- FD.io CSIT
  - Daily Trending and Release Reports (10 min)
  - Anomaly Detection (10 min)
  - Data Processing Infrastructure (10 min)
- Wrap-up (2 min)
- Q&A (15 min)

# Topics

LF NETWORKING
LFN Developer & Testing Forum

- FD.io Overview
  - VPP and CSIT core projects
- FD.io CSIT
  - Daily Trending and Release Reports
  - Anomaly Detection
  - Data Processing Infrastructure
- Wrap-up
- Q&A

# FastData.io Core-Project: VPP

LF NETWORKING
LFN Developer & Testing Forum

## Vector Packet Processing

- Performance
- Portability
- Platforms
- FD.io LF NETWORKING
- SDN
- Efficiency

## Fast packet processing software runing in Linux user space

VPP on x86 servers **outperforms** specialized packet processing HW.

VPP runs on COTS hardware: AMD ARM intel OpenPOWER POWER8

VPP runs in any environment: **bare-metal, VM, containers**.

Platform for building forwarding applications and network functions with plug-ins as first class citizens and with configurable graph architecture.

Core project in FD.io, an open source collaborative project in LF Networking.

Software **programmable**, **extendable** and **flexible.**

Allows ability to **upscale** and **downscale.**

# VPP Multiple Deployment Models

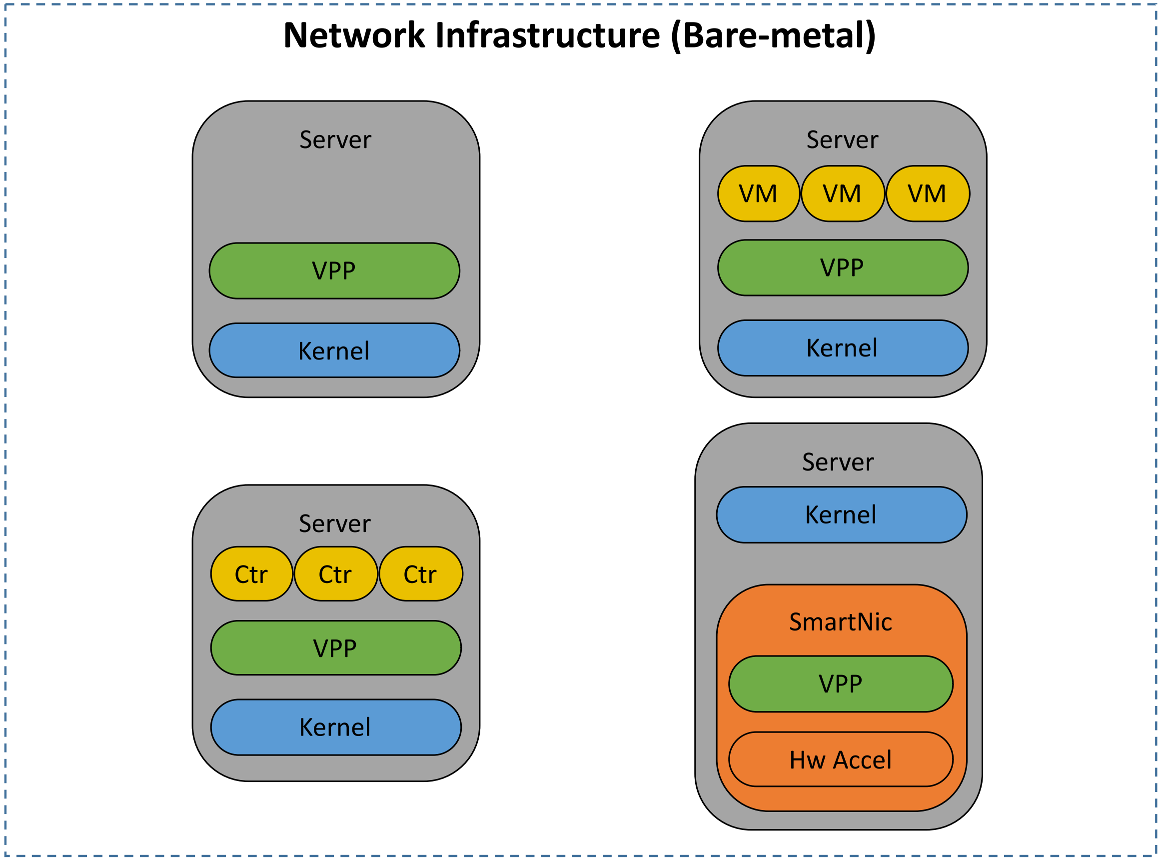


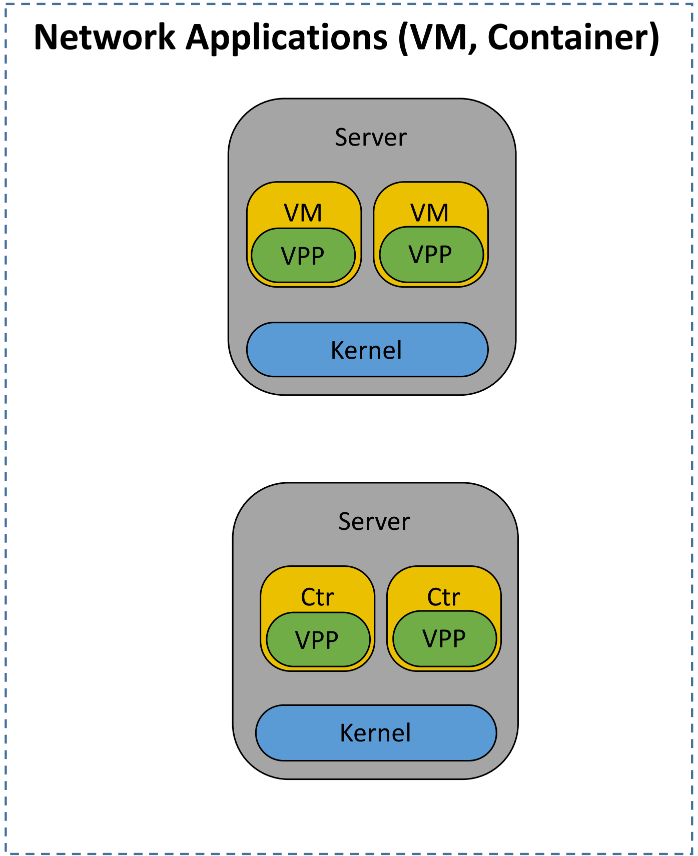

# VPP Infrastructure and Nodes

LF NETWORKING
LFN Developer & Testing Forum

**VPP provides infrastructure** for packet processing:

- **high performance** - tens of Mpps per core
- **low latency** - single/double digit usec
- **small memory footprint** - 1GB including buffers
- **scalable** - millions of prefixes, unlimited flows
- **flexible** - static/dynamic processing node graph
- **modular** - small node as dev & deployment unit
- **extensible** - custom plugins (no need to upstream)
- **in user space** - no need to rebuild kernel
- **under popular distros** - Ubuntu, CentOS, Debian
- **for major CPU families** - x86 and ARM
- **open source** - supported by large community

**VPP includes ready-to-use nodes**, libraries and entire data-planes for:

- L2 switching
- MPLS switching
- IPv4 and IPv6 routing
- all sorts of tunneling
- encryption/decryption IPsec, WireGuard
- host stack including TCP, TLS, QUIC
- numerous features

# Magic of VPP

LFN Networking
LFN Developer & Testing Forum

**VPP makes an extensive use of**:

- **avoidance of instruction cache thrashing**
  packet vector processing in small graph nodes
- **just-in-time data prefetch to cache**
  several loop iterations b/w data prefetch and use
- **instruction level parallelism**
  processing of 4-8 packets in single loop iteration
- **Intel Direct Data I/O (DDIO), ARM stashing**
  packet exchange with NIC directly from L3 cache
- **branching avoidance**
  increased IPC efficiency
- **advanced vector extensions**
  AVX512 and AVX2
- **hyper-threading**
  up to 50% better performance

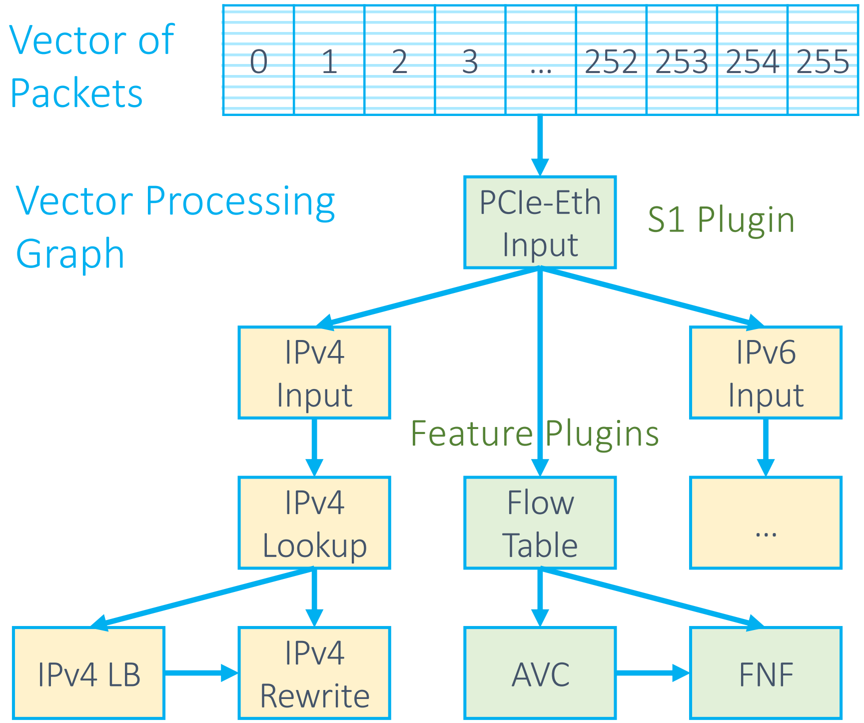

# FastData.io Core-Project: CSIT

LF NETWORKING
LFN Developer & Testing Forum

## Continuous System and Integration Testing

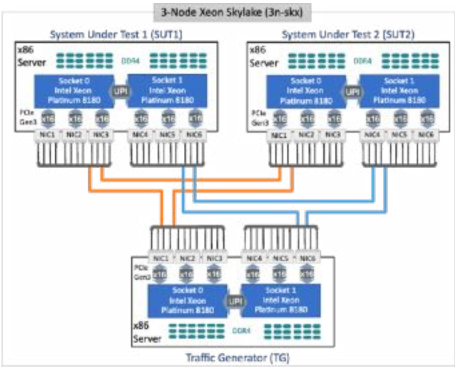


Physical Testbeds in LFN FD.io Labs

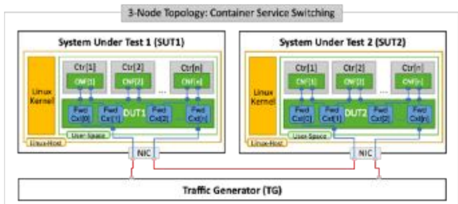


Logical Test Topologies

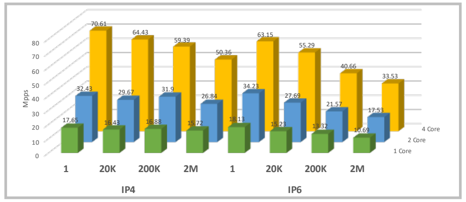


Benchmark Results

## Automated benchmarking of networking SW

- **Applicability to use cases**
  Understand why we measure and what the data tells us
- **Excellence in tooling and methodology**
  Be proficient in how we measure, continuously improving fidelity
- **Driving collaboration across companies and industries**
  Foster and engage in tight collaboration, keep learning from each other
- **Leading in communication**
  Striving to present the right data, at the right time and in the right way
- Trending and Release Reports, Anomaly Detection, Automated Test Data Processing

**“Above all else, show the data.” – Edward R. Tufte**

# CSIT Benchmark Areas and Methodologies

LF NETWORKING
LFN Developer & Testing Forum

## Benchmark Test Areas

- L2 Ethernet Switching (up to 1mil MACs)
- IPv4, IPv6 Routing (up to 2mil routes)
- IPsec IPv4 Routing (up to 60k IPsec tunnels)
- SRv6 Routing
- Features: ACLs, NAT44-EI/ED, Policer, …
- IPv4, IPv6 Tunnels
- KVM VMs vHost-user (up to box-full)
- LXC.DRC Container Memif (up to box-full)

## Test Methodologies

- Packet Throughput
- Stateless and stateful
- Speedup Multi-Core
- Packet Latency
- Soak Tests
- Reconfiguration Tests
- NFV Service Density
- Host-stack Testing

## Tools

Jenkins, Read the Docs, plotly | Dash, ROBOT FRAMEWORK, python, TRex Realistic Traffic Generator

## Physical and Logical Network Topologies

2-Node Topology: NIC-to-NIC Switching

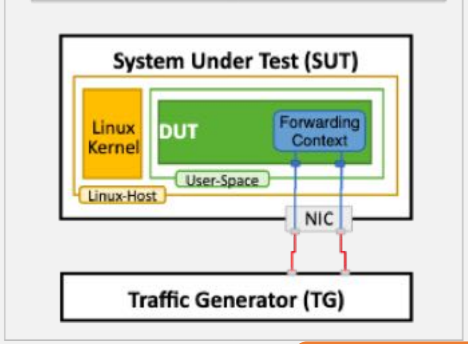


2-Node Topology: VM Service Switching

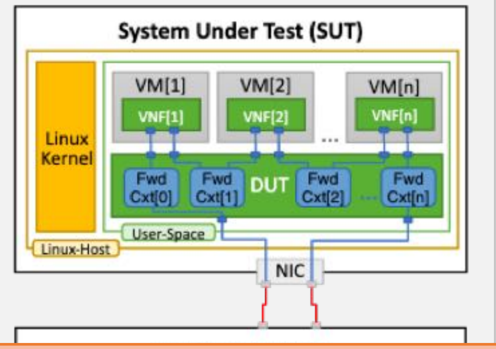


2-Node Topology: Container Service Switching

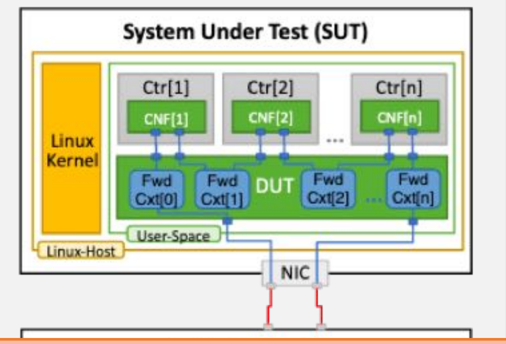


**CPUs**

skylake (skx)
cascadelake (clx)
**icelake (icx)**
denverton (dnv)
**snowridge (snr)**
thunderX2 (tx2)
zen2 (zn2)
**arm n1**
**aws ec2**

**NICs**

10 GbE (x5??)
10 GbE (x710)
25 GbE (xxv710)
100 GbE (cx556a)
**100 GbE (e810)**
**50 GbE (ena)**

"For Better or Worse, Benchmarks Shape a Field." – David Patterson

# Test Methodologies: MRR, MLRsearch, PLRsearch

LF NETWORKING
LFN Developer & Testing Forum

## Beyond RFC2544: Innovating in benchmarking methodologies

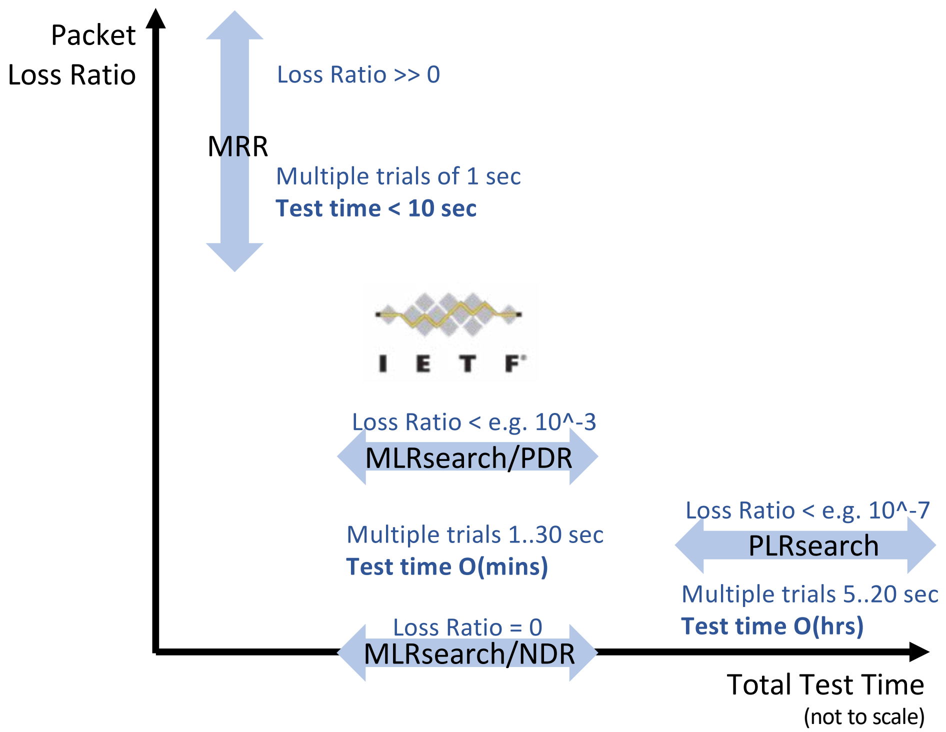


- **Maximum Receive Rate (MRR)**
  - Measures packet throughput under maximum load regardless of packet loss.
  - Used for daily trending and automated anomaly detection in CSIT.
- **Multiple Loss Ratio search (MLRsearch)**
  - Discovers multiple packet rates (NDR, PDR[1]) in a single search.
  - IETF BMWG draft-ietf-bmwg-mlrsearch .
- **Probabilistic Loss Ratio search (PLRsearch)**
  - Soak[2] testing for data planes over an extended period of time.
  - IETF BMWG draft-vpolak-bmwg-plrsearch .
- **All used in production CI for stateless and stateful benchmarks**
  - Stateless L2, IP4, IP6, clear-text and crypto tunnels tests for bits/sec, packets/sec.
  - Stateful UDP and TCP tests for bits/sec, connections/sec, packets/sec.
- **Why does one need MRR, PDR, NDR, Soak?**
  - A well behaving system should have MRR, PDR, NDR and Soak rates close to each other!

[1] Packet rate and bandwidth measurements: Non-Drop Rate (NDR, with zero packet loss), ii) Partial Drop Rate (PDR, with packet loss rate not greater than configured non-zero packet-loss-ratio).
[2] CSIT Soak tests use PLRsearch with packet-loss-ratio set to 10^-7.

# Topics

LF NETWORKING
LFN Developer & Testing Forum

- FD.io Overview
  - VPP and CSIT core projects
- FD.io CSIT
  - Daily Trending and Release Reports
  - Anomaly Detection
  - Data Processing Infrastructure
- Wrap-up
- Q&A

# Presentation and Analytics Layer

LF NETWORKING
LFN Developer & Testing Forum

- Trending Daily and Weekly
- Reporting per Release

fd.io

# Performance Trending Overview

LF NETWORKING
LFN Developer & Testing Forum

- Trending
  - Trending graphs
  - Dashboard
  - Alerting
- https://s3-docs.fd.io/csit/master/trending

fd.io

# Trending: Dashboard Graphs (web page)

LF NETWORKING
LFN Developer & Testing Forum

- The trend of packet throughput over the chosen time period.
- The data points (average of 10 samples) are displayed as dots.
- The trend lines are calculated by JumpAvg algorithm.
- The changes in the trend are marked them by circles: red for regression and green for progression.


* Test naming convention:
https://s3-docs.fd.io/csit/master/report/csit_framework_documentation/csit_test_naming.html

fd.io

# Trending: Dashboard Tables (web page)

LF NETWORKING
LFN Developer & Testing Forum

| Test Case * | Trend [Mpps] | Number of runs [#] | Trend Change [%] | Regressions [#] | Progressions [#] |
|---|---|---|---|---|---|
| 25ge2p1xxv710-64b-4t2c-ethip4vxlan-l2bdbasemaclrn-mrr | 18.63 | 4 | -11.05 | 1 | 0 |
| 25ge2p1xxv710-64b-4t2c-eth-l2bdbasemaclrn-mrr | 32.98 | 2 | -10.93 | 1 | 0 |
| 25ge2p1xxv710-64b-4t2c-ethip4vxlan-l2xcbase-mrr | 22.21 | 4 | -10.72 | 1 | 0 |
| 25ge2p1xxv710-78b-4t2c-ethip6-ip6base-mrr | 28.29 | 2 | -9.8 | 1 | 0 |
| 25ge2p1xxv710-64b-4t2c-ethip4-ip4base-mrr | 32.33 | 2 | -8.93 | 1 | 0 |
| 25ge2p1xxv710-64b-4t2c-ethip4gtpusw-ip4base-mrr | 17.64 | 4 | -8.34 | 1 | 0 |
| 25ge2p1xxv710-78b-4t2c-avf-ethip6-ip6base-mrr | 37.34 | 12 | -0.07 | 1 | 0 |
| 25ge2p1xxv710-64b-4t2c-avf-ethip4-ip4scale20k-mrr | 37.37 | 9 | 0.4 | 0 | 1 |
| 25ge2p1xxv710-64b-4t2c-avf-ethip4-ip4scale200k-mrr | 37.36 | 9 | 0.41 | 0 | 1 |
| 25ge2p1xxv710-64b-4t2c-avf-eth-l2bdscale10kmaclrn-mrr | 37.32 | 13 | 0.57 | 0 | 1 |
| 25ge2p1xxv710-1518b-4t2c-avf-ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha-mrr | 0.96 | 108 | -20.41 | 0 | 0 |
| 25ge2p1xxv710-imix-4t2c-avf-ethip4ipsec4tnlsw-ip4base-int-aes256gcm-mrr | 4.68 | 109 | -13.82 | 0 | 0 |
| 25ge2p1xxv710-imix-4t2c-avf-ethip4ipsec4tnlsw-ip4base-int-aes128cbc-hmac512sha-mrr | 2.27 | 107 | -11.15 | 0 | 0 |

- Test Case: name of FD.io CSIT test case.
- Trend [Mpps]: Trend (population average) of the last group of runs.
- Number of runs [#]: Number of runs belonging to the last sample group classified under the last trend value.
- Trend Change [%]: Relative change of last Trend vs maximum of trend values over the last quarter except last week.
- Regressions [#]: Number of regressions detected.
- Progressions [#]: Number of progressions detected.

* Test naming convention: https://s3-docs.fd.io/csit/master/report/csit_framework_documentation/csit_test_naming.html

fd.io

# Alerting: Failures (email)

LF NETWORKING
LFN Developer & Testing Forum

From: noreply@jenkins.fd.io
Subject: [csit-report] master: Failed tests as of 2022-xx-xx xx:xx:xx UTC #email
To: Fdio+Csit-Report via Email Integration <csit-report@lists.fd.io>
Reply-To: csit-report@lists.fd.io

Following tests failed in the last trending job runs, listed per testbed type.
...
vpp-2n-skx-ndrpdr,
5 tests failed, 313 tests passed, duration: 28:10,
CSIT build: https://jenkins.fd.io/view/csit/job/csit-vpp-perf-ndrpdr-weekly-master-2n-skx/177,
vpp version: 22.02-rc0~461-ged5be4700~b2120

Legend:

| test-name | nic | frame-size | core-config |
|---|---|---|---|
| af-xdp-ethip4-ip4scale20k | 25ge2p1xxv710 | 64b | 4t2c[1] |
| avf-ethip6-ip6scale2m | 25ge2p1xxv710 | 78b | 8t4c[0] |

Failure Legend:
[0] - Failed to create AVF interface on host xxx.xxx.xxx.
[1] - Minimal rate loss ratio 1.851646113394808e-06 does not reach target 0.0.1 packets lost.

fd.io

# Alerting: Regressions (email)

LF NETWORKING
LFN Developer & Testing Forum

From: noreply@jenkins.fd.io
Subject: [csit-report] master: Regressions as of 2022-xx-xx xx:xx:xx UTC #email
To: Fdio+Csit-Report via Email Integration <csit-report@lists.fd.io>
Reply-To: csit-report@lists.fd.io

Following regressions occured in the last trending job runs, listed per testbed type.
...

vpp-3n-dnv-mrr,
CSIT build: https://jenkins.fd.io/view/csit/job/csit-vpp-perf-mrr-daily-master-3n-dnv/1106,
vpp version: 22.02-rc0~477-g2b4f74ff9~b2136

Legend:
test-name [ Last trend in Mpps | Number of runs for last trend | Long-Term Change ]

10ge2p1x553-78b-1t1c-ethip6-ip6scale20k-mrr [ 2.55M | 1 | -13.12% ]
10ge2p1x553-imix-4t4c-ethip4ipsec4tnlsw-ip4base-int-aes256gcm-mrr [ 0.58M | 1 | -72.54% ]

fd.io

# Alerting: Progressions (email)



From: noreply@jenkins.fd.io
Subject: [csit-report] master: Progressions as of 2022-xx-xx xx:xx:xx UTC #email
To: Fdio+Csit-Report via Email Integration <csit-report@lists.fd.io>
Reply-To: csit-report@lists.fd.io

Following progressions occured in the last trending job runs, listed per testbed type.
...

vpp-2n-dnv-mrr,
CSIT build: https://jenkins.fd.io/view/csit/job/csit-vpp-perf-mrr-daily-master-2n-dnv/1092,
vpp version: 22.02-rc0~477-g2b4f74ff9~b2136

Legend:
test-name [ Last trend in Mpps | number of runs for last trend | Long-Term Change ]

10ge2p1x553-64b-2t2c-eth-l2patch-mrr [ 18.55M | 1 | 27.8% ]

# Release Report Overview

LF NETWORKING
LFN Developer & Testing Forum

- Report
  - Box and whiskers for iterative results
  - Histograms for latency
  - Heat-map for service density
  - CPU scaling (speedup) graphs
  - Tabular data for coverage results
  - Comparison tables
- https://s3-docs.fd.io/csit/master/report

# Report: Box and Whiskers for Iterative Results

LF NETWORKING

LFN Developer & Testing Forum

- The plot provides us all information about statistical data - minimum, first quartile, median, third quartile, maximum and outliers. This information is displayed in the hover box.
- The X-axis lists indices of individual test suites as listed in Graph Legend.
- The Y-axis presents measured Packet throughput values [Mpps].
- The tests in each plot are grouped and ordered by chosen criteria. This grouping is written in the plot title (topology, processor architecture, NIC, frame size, number of cores, area, test type, measured property).

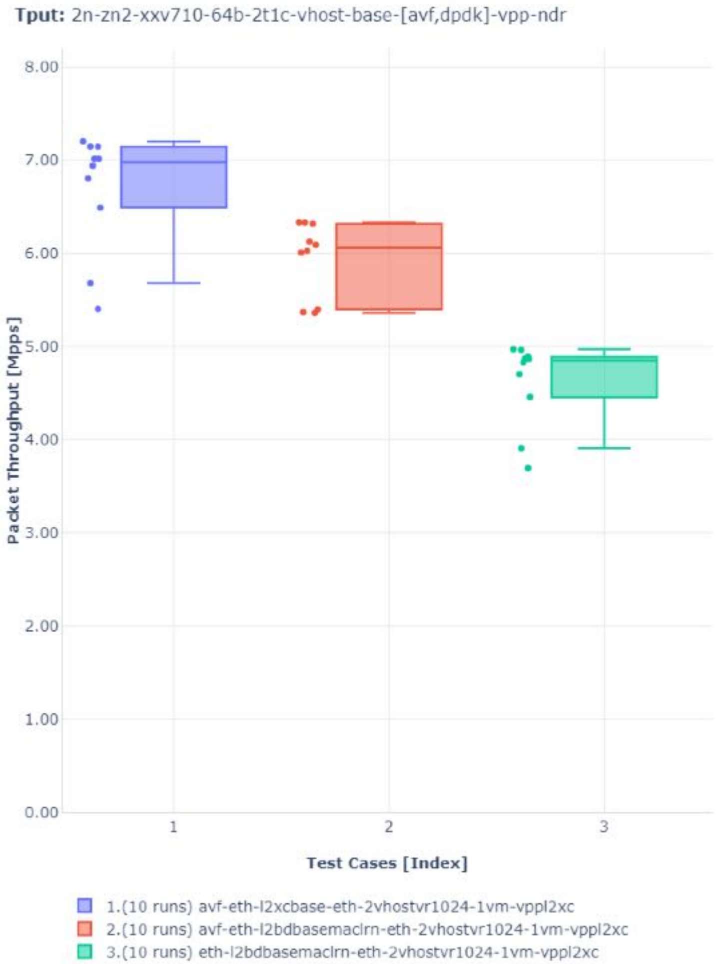


.io

# Report: HDR Histograms for Latency



- TRex integrates High Dynamic Range Histogram (HDRH) functionality and reports per packet latency distribution for latency streams sent in parallel to the main load packet 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.

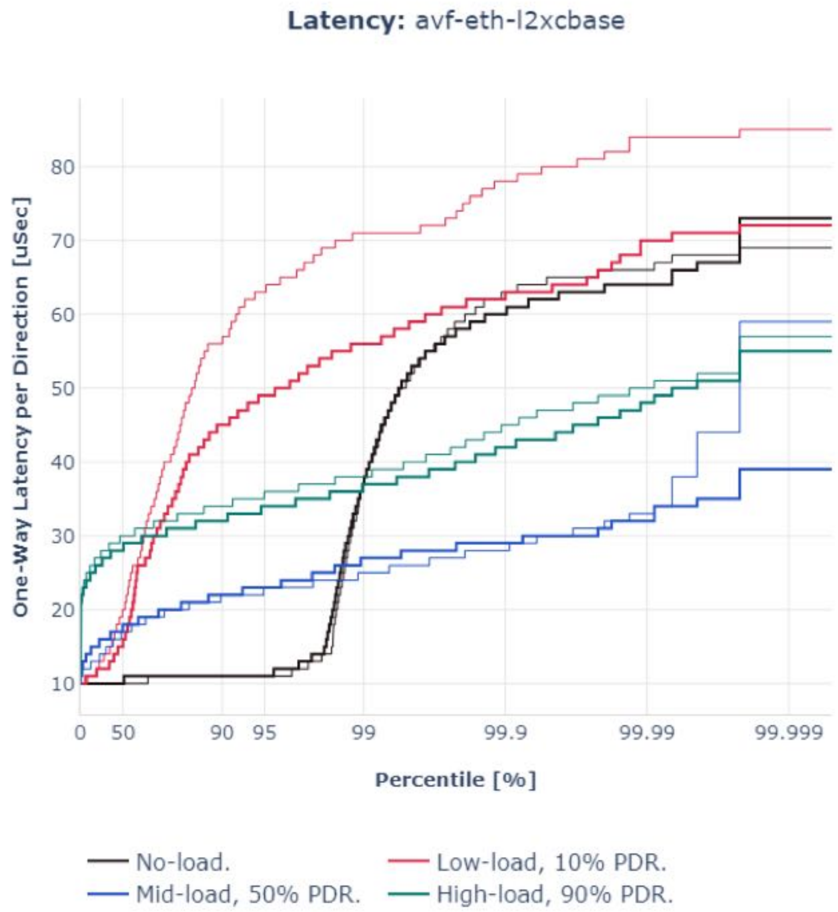

# Report: Heat-map for Service Density

- 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.
- 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).

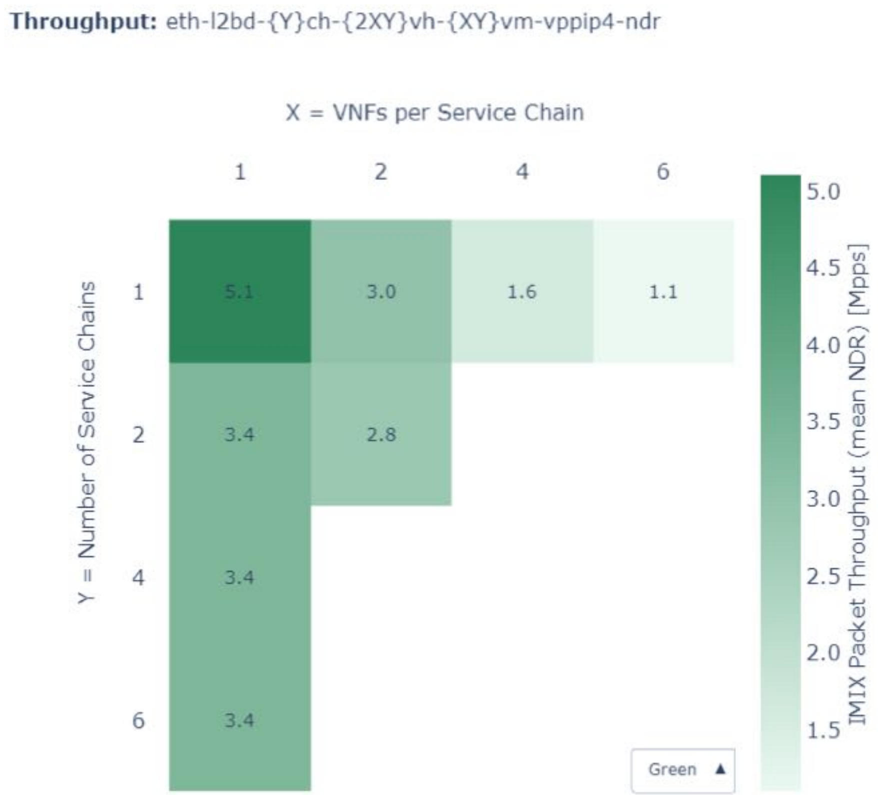


LF NETWORKING
LFN Developer & Testing Forum
FD.io

# Report: CPU Scaling Graphs

- Grouped bars illustrate the 64B/78B packet throughput speedup ratio for 2- and 4-core multi-threaded VPP configurations relative to 1-core configurations.

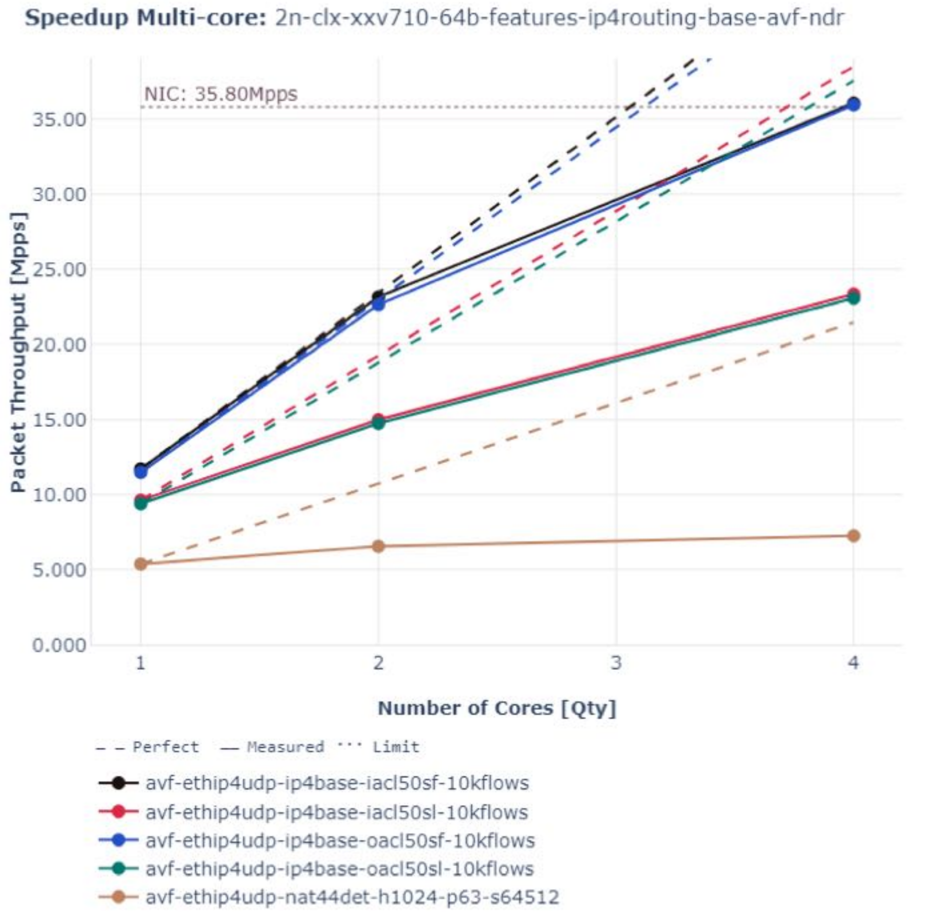

# Report: Tables with Coverage Test Results



- Detailed test results
- Test configuration
- Test operational data

# Report: Detailed Test results (web-page)

LF NETWORKING
LFN Developer & Testing Forum

| Test Name | Throughput:<br>1. Mpps Gbps (NDR)<br>2. Mpps Gbps (PDR)<br><br>One-Way Latency Percentiles in uSec at %PDR load, one set per each direction:<br>3. P50 P90 P99 P50 P90 P99 (10% PDR)<br>4. P50 P90 P99 P50 P90 P99 (50% PDR)<br>5. P50 P90 P99 P50 P90 P99 (90% PDR) |
|---|---|
| 64b-2t1c-avf-dot1q-l2bdbasemaclrn-ndrpdr | 1. 15.75 11.09<br>2. 15.93 11.22<br><br>3. 34 104 119 31 117 157<br>4. 26 36 50 26 38 49<br>5. 45 54 62 45 55 62 |

fd.io

# Report: Test Configuration

LF NETWORKING
LFN Developer & Testing Forum

| Test Name | VPP API Test Commands History - Commands Used Per Test Case |
|---|---|
| 64b-0c-ethip4-l2xcbase-scapy | **DUT1:**<br>cli_inband(cmd='show logging')<br>show_version()<br>sw_interface_dump(name_filter_valid=False,name_filter='')<br>cli_inband(cmd='trace add dpdk-input 50')<br>cli_inband(cmd='trace add vhost-user-input 50')<br>cli_inband(cmd='trace add memif-input 50')<br>cli_inband(cmd='trace add avf-input 50')<br>sw_interface_set_flags(sw_if_index=1,flags=1)<br>hw_interface_set_mtu(sw_if_index=1,mtu=9200)<br>sw_interface_set_flags(sw_if_index=2,flags=1)<br>hw_interface_set_mtu(sw_if_index=2,mtu=9200)<br>sw_interface_dump(name_filter_valid=False,name_filter='')<br>sw_interface_dump(name_filter_valid=False,name_filter='')<br>sw_interface_dump(name_filter_valid=False,name_filter='')<br>sw_interface_dump(name_filter_valid=False,name_filter='')<br>sw_interface_set_l2_xconnect(rx_sw_if_index=1,tx_sw_if_index=2,enable=True)<br>sw_interface_set_l2_xconnect(rx_sw_if_index=2,tx_sw_if_index=1,enable=True) |

.io

# Report: Test Operational Data

LF NETWORKING
LFN Developer & Testing Forum

| 64b-2t1c-avf-ethip4-ip4scale2m-rnd-ndrpdr | | | | | |
|---|---|---|---|---|---|
| Host IP: 10.32.8.18, Socket: /run/vpp/stats.sock | | | | | |
| **worker_1** | | | | | |
| **Name** | **Nr of Vectors** | **Nr of Packets** | **Suspends** | **Cycles per Packet** | **Average Vector Size** |
| avf-input | 168612128 | 1316884084 | 0 | 44.09 | 7.81 |
| ip4-input-no-checksum | 8712235 | 1316884084 | 0 | 36.72 | 151.15 |
| ip4-rewrite | 8712235 | 1316884084 | 0 | 41.13 | 151.15 |
| ip4-lookup | 8712235 | 1316884084 | 0 | 158.83 | 151.15 |
| ethernet-input | 8712236 | 1316884244 | 0 | 21.17 | 151.15 |
| unix-epoll-input | 174602 | 0 | 0 | 60351.98 | 0.00 |
| avf-0/3b/a/0-tx | 8712235 | 1316884084 | 0 | 29.19 | 151.15 |
| avf-0/3b/a/0-output | 8712235 | 1316884084 | 0 | 13.25 | 151.15 |
| **worker_2** | | | | | |
| **Name** | **Nr of Vectors** | **Nr of Packets** | **Suspends** | **Cycles per Packet** | **Average Vector Size** |
| avf-input | 169800849 | 1447719319 | 0 | 42.01 | 8.53 |
| ip4-input-no-checksum | 8215579 | 1316856436 | 0 | 36.96 | 160.29 |
| ip4-rewrite | 8215579 | 1316856436 | 0 | 41.10 | 160.29 |
| ip4-lookup | 8215579 | 1316856436 | 0 | 159.28 | 160.29 |
| ethernet-input | 8215579 | 1316856436 | 0 | 21.27 | 160.29 |
| unix-epoll-input | 174431 | 0 | 0 | 60421.28 | 0.00 |
| avf-0/3b/2/0-tx | 8215579 | 1316856436 | 0 | 28.98 | 160.29 |
| avf-0/3b/2/0-output | 8215579 | 1316856436 | 0 | 13.17 | 160.29 |

.io

# Report: Comparison Tables Overview



Comparison between different:

- VPP releases in the same CSIT/lab environment
- CSIT/lab environments for the same VPP release
- CPU architectures
- NIC models, drivers

# Report: Current vs Previous Release (e.g. 21.10 vs 21.06)

LF NETWORKING
LFN Developer & Testing Forum

| Test Case | 2101.0-7 | 2106.0-7 | 2106.0-8 | 2110.0-8 | 2106.0-8 vs 2106.0-7 | 2110.0-8 vs 2106.0-8 |
|---|---|---|---|---|---|---|
| 64b-2t1c-avf-ethip4udp-nat44ed-h1024-p63-s64512-udir | 6.07 +- 0.02 | 7.19 +- 0.06 | 7.16 +- 0.05 | 8.07 +- 0.02 | -0.34 +- 1.06 | +12.67 +- 0.85 |
| 64b-2t1c-avf-ethip4udp-nat44ed-h1024-p63-s64512-tput | 6.46 +- 0.33 | 6.75 +- 0.3 | 6.7 +- 0.34 | 6.98 +- 0.32 | -0.68 +- 6.67 | +4.12 +- 7.05 |
| 64b-2t1c-avf-ethip4udp-nat44det-h65536-p63-s4128758 | 5.2 +- 0.02 | 5.51 +- 0.09 | 5.43 +- 0.08 | 5.65 +- 0.02 | -1.53 +- 2.1 | +4.0 +- 1.55 |
| 64b-2t1c-avf-ethip4udp-nat44det-h1024-p63-s64512 | 8.34 +- 0.77 | 6.18 +- 0.17 | 6.12 +- 0.21 | 6.28 +- 0.03 | -0.98 +- 4.34 | +2.72 +- 3.6 |
| 64b-2t1c-avf-ethip4udp-nat44ed-h65536-p63-s4128768-tput | 6.57 +- 0.36 | 6.52 +- 0.33 | 6.45 +- 0.18 | 6.63 +- 0.27 | -1.02 +- 5.69 | +2.7 +- 5.08 |
| 78b-2t1c-ethip6-ip6scale20k | 13.73 +- 0.18 | 12.41 +- 0.19 | 12.35 +- 0.09 | 12.67 +- 0.25 | -0.46 +- 1.71 | +2.54 +- 2.14 |
| 64b-2t1c-avf-ethip4tcp-nat44ed-h65536-p63-s4128768-cps | 0.49 +- 0.0 | 0.5 +- 0.0 | 0.5 +- 0.0 | 0.52 +- 0.0 | -0.05 +- 0.23 | +2.45 +- 0.28 |
| 64b-2t1c-avf-ethip4tcp-nat44ed-h16384-p63-s1032192-cps | 0.55 +- 0.0 | 0.53 +- 0.0 | 0.53 +- 0.0 | 0.55 +- 0.0 | +0.06 +- 0.33 | +2.3 +- 0.43 |
| 64b-2t1c-avf-ethip4udp-ip4base-h65536-p63-s4128768-cps | 3.99 +- 0.09 | 3.7 +- 0.09 | 3.63 +- 0.09 | 3.7 +- 0.06 | -1.94 +- 3.49 | +1.97 +- 3.1 |
| | | … | | | | |
| 64b-2t1c-eth-l2bdbasemaclrn | 19.71 +- 0.06 | 19.39 +- 0.16 | 19.23 +- 0.21 | 18.69 +- 0.22 | -0.83 +- 1.35 | -2.8 +- 1.56 |
| 64b-2t1c-avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm | 7.62 +- 0.06 | 8.12 +- 0.17 | 8.07 +- 0.04 | 7.84 +- 0.03 | -0.7 +- 2.09 | -2.89 +- 0.58 |
| 64b-2t1c-avf-ethip4udp-ip4base-h1024-p63-s64512-cps | 4.6 +- 0.22 | 4.73 +- 0.17 | 4.83 +- 0.11 | 4.69 +- 0.34 | +2.3 +- 4.42 | -3.06 +- 7.35 |
| 64b-2t1c-avf-eth-l2xcbase-eth-2vhostvr1024-1vm | 9.68 +- 0.72 | 10.57 +- 0.17 | 10.56 +- 0.24 | 10.18 +- 0.05 | -0.15 +- 2.79 | -3.61 +- 2.24 |
| 64b-2t1c-avf-eth-l2xcbase-eth-2vhostvr1024-1vm-vppl2xc | 8.92 +- 1.18 | 10.11 +- 0.02 | 10.1 +- 0.02 | 9.73 +- 0.02 | -0.11 +- 0.25 | -3.67 +- 0.25 |
| 64b-2t1c-ethip4-ip4base | 17.97 +- 0.1 | 17.93 +- 0.41 | 18.09 +- 0.4 | 17.36 +- 0.05 | +0.9 +- 3.21 | -4.06 +- 2.14 |
| 64b-2t1c-avf-ethip4tcp-nat44ed-h1024-p63-s64512-cps | 0.64 +- 0.04 | 0.7 +- 0.0 | 0.7 +- 0.0 | 0.67 +- 0.0 | +0.16 +- 0.54 | -4.17 +- 0.3 |
| 64b-2t1c-avf-eth-l2bdbasemaclrn-eth-2vhostvr1024-1vm-vppl2xc | 7.16 +- 0.39 | 8.18 +- 0.31 | 8.57 +- 0.27 | 7.94 +- 0.12 | +4.81 +- 5.15 | -7.37 +- 3.21 |
| 64b-2t1c-avf-ethip4udp-nat44ed-h1024-p63-s64512-cps | 1.41 +- 0.01 | 1.54 +- 0.01 | 1.54 +- 0.01 | 1.32 +- 0.05 | -0.19 +- 1.28 | -13.83 +- 3.24 |

```
Test Case:             Test cases executed for: PDR throughput on 2n-skx with xxv710.
YYMM.V-E:              [Mpps +- StDev] results for VPP (YY year, MM month, V minor version) tested in CSIT (E environment version, "7" for CSIT rls2106, "8" for CSIT rls2110).
YYMM.V-E vs YYMM.V-E:  [% +- StDev] Relative difference, latter is the base. [N] for Root Cause Analysis (RCA) reference, see RCA list below.
2106.0-8 vs 2106.0-7:  Compares CSIT environment current vs previous with previous VPP code.
2110.0-8 vs 2106.0-8:  Compares VPP code current vs previous in the current CSIT environment.
```

# Report: Testbed Topologies (2-node vs 3-node)

LF NETWORKING
LFN Developer & Testing Forum

| Test Case | 3n-skx | 2n-skx | Diff |
| --- | --- | --- | --- |
| 78b-1c-avf-ethip6-ip6base | 19.82 +- 0.12 | 20.31 +- 0.2 | +2.5 +- 1.17 |
| 78b-1c-ethip6-ip6base | 15.39 +- 0.08 | 15.72 +- 0.21 | +2.14 +- 1.49 |
| 64b-1c-ethip4-ip4base | 17.12 +- 0.05 | 17.36 +- 0.05 | +1.37 +- 0.44 |
| 64b-1c-eth-l2bdbasemaclrn | 18.52 +- 0.12 | 18.69 +- 0.22 | +0.9 +- 1.36 |
| 64b-1c-avf-eth-l2patch | 37.32 +- 0.01 | 37.5 +- 0.0 | +0.49 +- 0.03 |
| 64b-1c-eth-l2xcbase | 25.5 +- 0.17 | 25.57 +- 0.18 | +0.27 +- 0.96 |
| 64b-1c-avf-eth-l2bdbasemaclrn | 26.2 +- 0.06 | 26.25 +- 0.11 | +0.22 +- 0.5 |
| 64b-1c-avf-eth-l2xcbase | 37.5 +- 0.0 | 37.5 +- 0.0 | +0.0 +- 0.0 |
| 64b-1c-eth-l2patch | 34.49 +- 0.06 | 34.36 +- 0.1 | -0.38 +- 0.34 |
| 64b-1c-avf-ethip4-ip4base | 24.14 +- 0.04 | 23.95 +- 0.02 | -0.78 +- 0.18 |
| 64b-1c-avf-dot1q-l2bdbasemaclrn | 19.13 +- 0.03 | 15.9 +- 0.06 | -16.88 +- 0.36 |

3n-skx, 2n-skx - topologies tested.
Diff - difference between 3n-skx and 2n-skx topologies.

# Report: CPU Models (e.g. skx vs clx)

LF NETWORKING
LFN Developer & Testing Forum

| Test Case | 2n-skx | 2n-clx | Diff |
|---|---|---|---|
| 64b-1c-af-xdp-ethip4-ip4scale20k | 0.09 +- 0.02 | 0.1 +- 0.03 | +11.8 +- 42.09 |
| 78b-1c-af-xdp-ethip6-ip6scale20k-rnd | 0.1 +- 0.02 | 0.12 +- 0.03 | +11.65 +- 37.75 |
| 64b-1c-avf-ethip4tcp-ip4base-h16384-p63-s1032192-cps | 1.44 +- 0.01 | 1.53 +- 0.0 | +6.15 +- 0.89 |
| 64b-1c-avf-ethip4tcp-ip4base-h1024-p63-s64512-cps | 2.09 +- 0.03 | 2.21 +- 0.02 | +6.06 +- 1.65 |
| 78b-1c-af-xdp-ethip6-ip6base | 0.11 +- 0.03 | 0.12 +- 0.01 | +5.7 +- 27.81 |
| 64b-1c-avf-ethip4tcp-ip4base-h65536-p63-s4128768-cps | 1.36 +- 0.0 | 1.43 +- 0.0 | +5.09 +- 0.32 |
| 64b-1c-avf-ethip4udp-ip4base-h1024-p63-s64512-cps | 4.69 +- 0.34 | 4.91 +- 0.3 | +4.82 +- 9.91 |
| … | | | |
| 64b-1c-avf-eth-l2bdbasemaclrn | 26.25 +- 0.11 | 24.16 +- 0.05 | -7.98 +- 0.45 |
| 64b-1c-avf-dot1q-l2bdbasemaclrn | 15.9 +- 0.06 | 14.62 +- 0.05 | -8.05 +- 0.48 |
| 64b-1c-avf-eth-l2bdscale1mmaclrn | 15.94 +- 0.42 | 14.66 +- 0.22 | -8.07 +- 2.77 |
| 78b-1c-avf-ethip6-ip6scale200k-rnd | 12.93 +- 0.11 | 11.88 +- 0.05 | -8.13 +- 0.9 |
| 64b-1c-avf-eth-l2xcbase-eth-2memif-1dcr | 15.41 +- 0.04 | 14.12 +- 0.02 | -8.42 +- 0.26 |
| 64b-1c-avf-ethip4udp-nat44det-h1024-p63-s64512 | 6.28 +- 0.03 | 5.57 +- 0.13 | -11.31 +- 2.17 |

2n-skx, 2n-clx - testbeds and NICs tested.
Diff - difference between 2n-skx and 2n-clx testbeds.

.io

# Report: NICs Comparison (e.g. xxv710 vs x710)

| Test Case | x710 | xxv710 | Diff |
|---|---|---|---|
| 64b-2t1c-avf-eth-l2patch | 29.76 +- 0.0 | 37.5 +- 0.0 | +26.0 +- 0.0 |
| 64b-2t1c-avf-eth-l2xcbase | 29.76 +- 0.0 | 37.5 +- 0.0 | +26.0 +- 0.0 |
| 64b-2t1c-eth-l2patch | 29.76 +- 0.0 | 34.36 +- 0.1 | +15.45 +- 0.35 |
| 64b-2t1c-eth-l2xcbase | 25.39 +- 0.2 | 25.57 +- 0.18 | +0.7 +- 1.05 |
| 78b-2t1c-ethip6-ip6base | 15.67 +- 0.26 | 15.72 +- 0.21 | +0.29 +- 2.16 |
| 64b-2t1c-avf-ethip4-ip4base | 23.94 +- 0.04 | 23.95 +- 0.02 | +0.05 +- 0.18 |
| 64b-2t1c-avf-eth-l2bdbasemaclrn | 26.26 +- 0.06 | 26.25 +- 0.11 | -0.03 +- 0.49 |
| 78b-2t1c-avf-ethip6-ip6base | 20.39 +- 0.19 | 20.31 +- 0.2 | -0.39 +- 1.35 |
| 64b-2t1c-eth-l2bdbasemaclrn | 18.79 +- 0.25 | 18.69 +- 0.22 | -0.54 +- 1.78 |
| 64b-2t1c-ethip4-ip4base | 17.83 +- 0.49 | 17.36 +- 0.05 | -2.66 +- 2.67 |
| 64b-2t1c-af-xdp-ethip4-ip4base | 0.12 +- 0.02 | 0.12 +- 0.02 | -3.52 +- 20.27 |
| 64b-2t1c-avf-eth-l2bdscale1mmaclrn | 16.58 +- 0.54 | 15.94 +- 0.42 | -3.86 +- 4.01 |
| 78b-2t1c-af-xdp-ethip6-ip6base | 0.12 +- 0.02 | 0.11 +- 0.03 | -4.74 +- 28.13 |
| 64b-2t1c-eth-l2bdscale1mmaclrn | 13.2 +- 0.2 | 12.47 +- 0.35 | -5.46 +- 3.02 |

x710, xxv710 - NICs tested.
Diff - difference between x710 and xxv710 NIC.

# Topics

LF NETWORKING
LFN Developer & Testing Forum

- FD.io Overview
  - VPP and CSIT core projects
- FD.io CSIT
  - Daily Trending and Release Reports
  - Anomaly Detection
  - Data Processing Infrastructure
- Wrap-up
- Q&A

# Anomaly Detection: Problem



- For each tracked test case, we have samples (sequence of floats).
  - Samples come from periodic (daily, weekly) job runs.
  - A sample can be MRR, NDR, PDR, average of multiple trials, …
  - CSIT code and lab environment can change in time.
  - Results are noisy, distribution is not normal.
    - This is true even after we minimized most sources of noisiness.
    - E.g. decreased CPU jitter via carefully choosing bios options and avoiding kernel context switches.
- Problem: Determine a point where the underlying true performance changes.
  - Performance of System Under Test (SUT).
  - More investigation needed to see if the cause is in VPP or in lab.

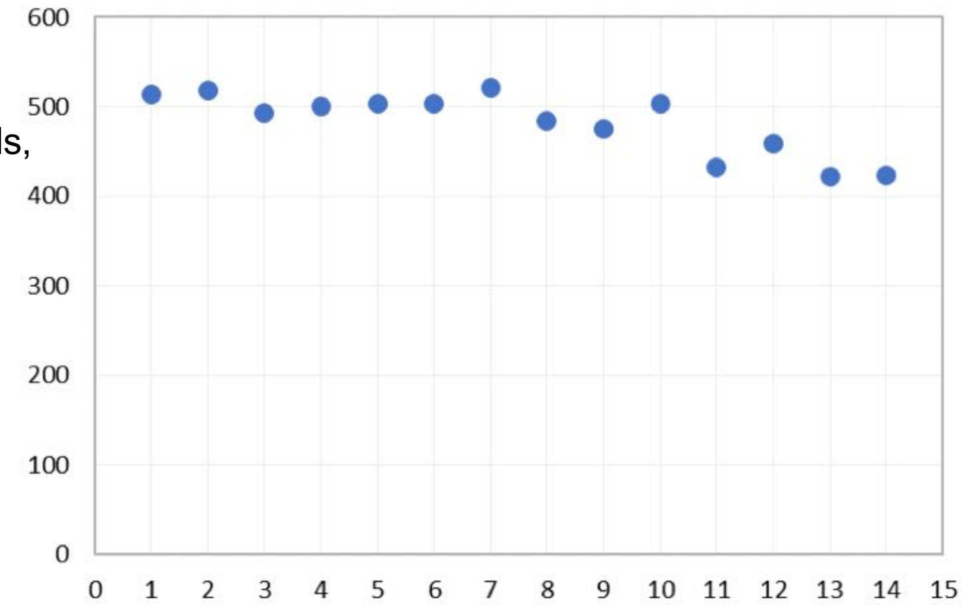

# Anomaly Detection: Solution

LF NETWORKING
LFN Developer & Testing Forum

- Restrictions for the solution algorithm:
  - No global state (no user labels, no learning).
  - Should be based on serious probability theory.
  - Should agree with human evaluation in most cases.
  - Occam's razor:
    - Start with the simplest underlying ideas.
    - Only add tweaks when needed to avoid "bad" evaluation.
  - Practicality:
    - Stop polishing as soon as evaluations are good enough.
- Solution: jumpavg.
  - Home-grown Python library.
  - Available at https://pypi.org/project/jumpavg/

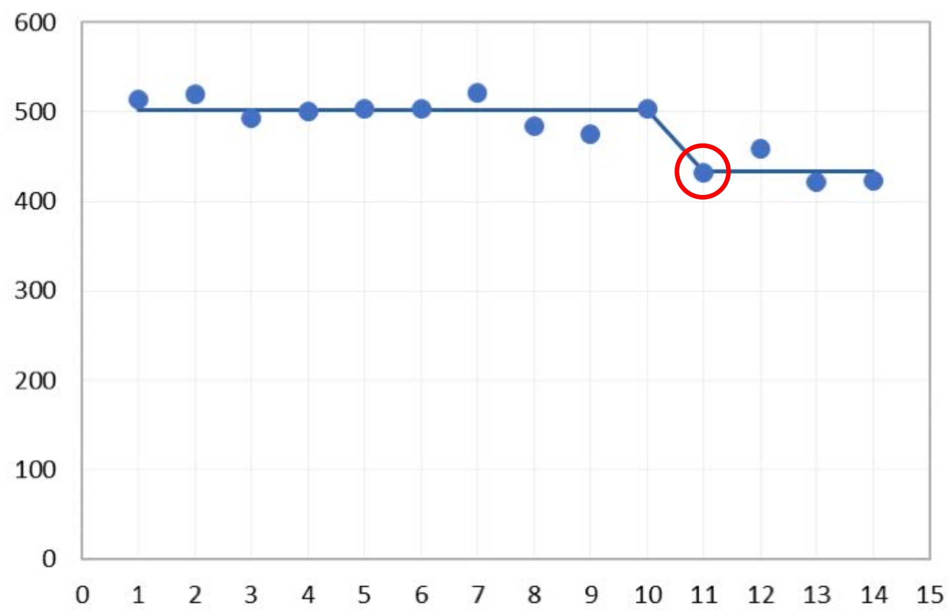


.io

# Anomaly Detection: Underlying Ideas

LF NETWORKING
LFN Developer & Testing Forum

- Partition samples into subsequent groups.
  - Model samples within a group as coming from a normal distribution.
  - Specific group sample positions are an information content, can be measured in bits.
  - The average and standard deviation of the distribution can also be encoded in bits.
  - Choose partition with least overall bit length (Minimal Description Length).
    - Another way of looking at it: Find a partition that compresses the best.
- Average of the group is the trend.
- Regression is when subsequent group has smaller trend.

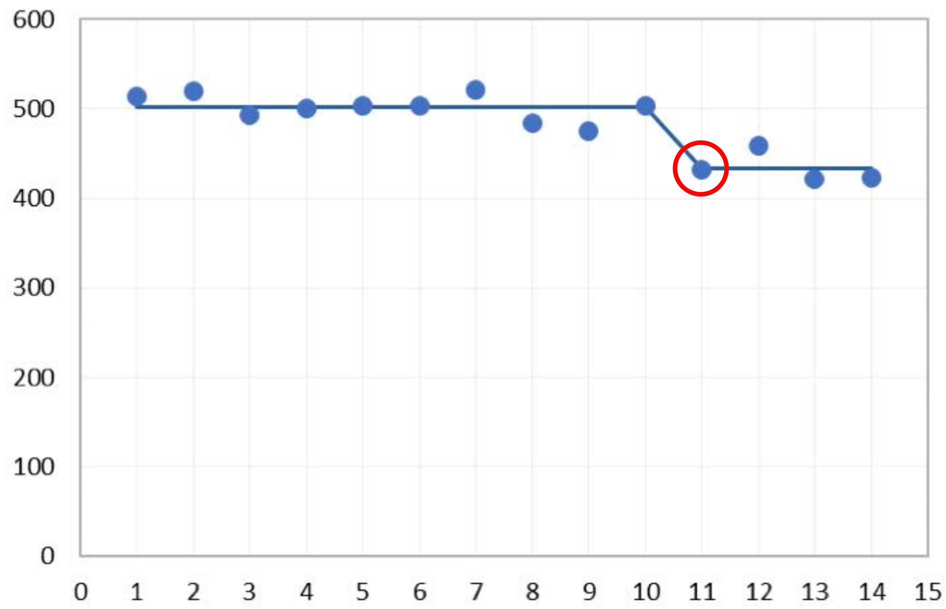


.io

# Patterns: Single Regression

- Single medium regression

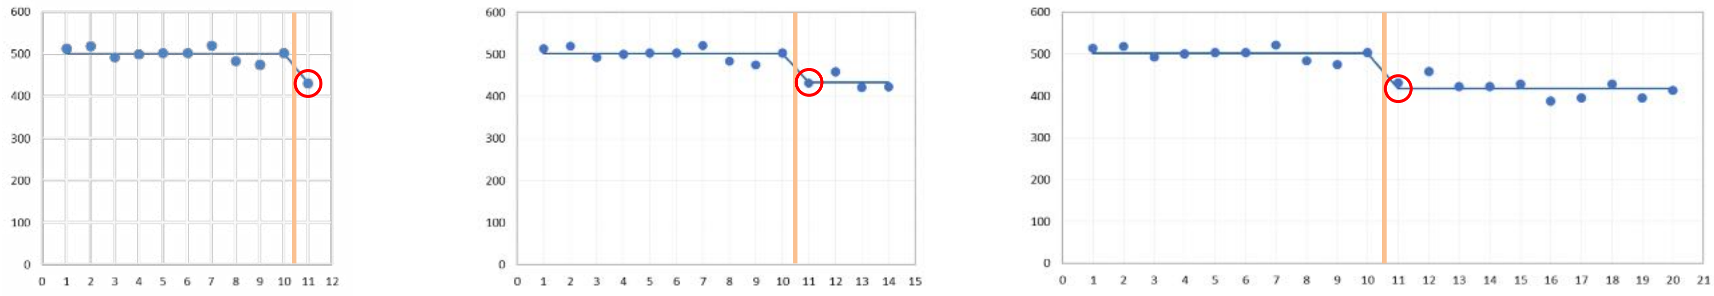

- Single small regression

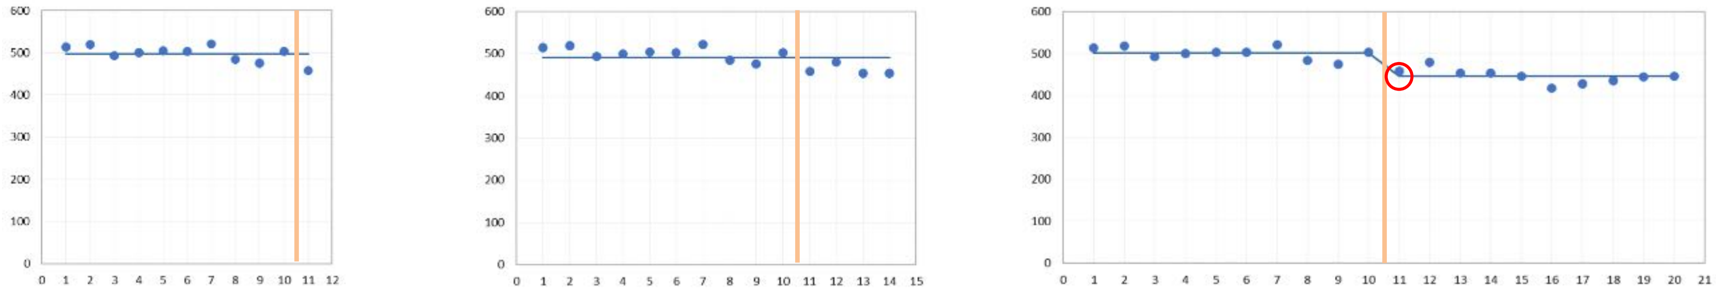

# Patterns: Reversed Regression



- Reversed big regression

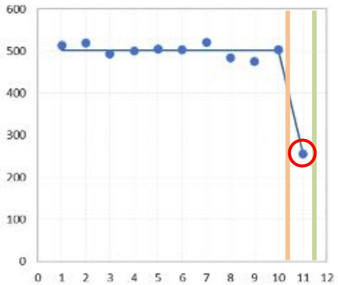


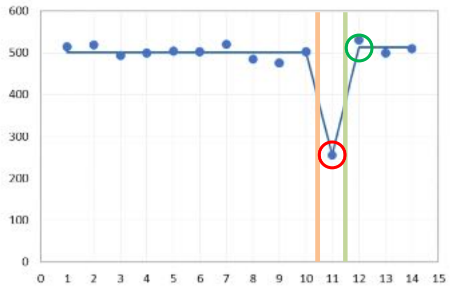


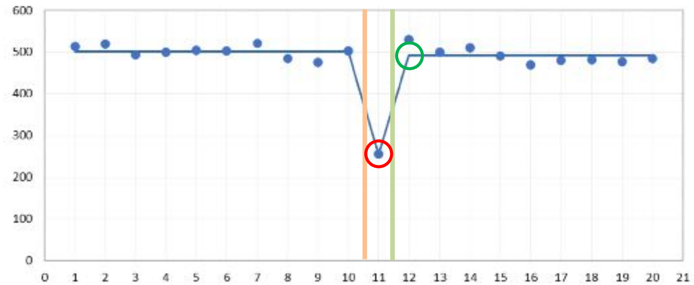


- Reversed medium regression

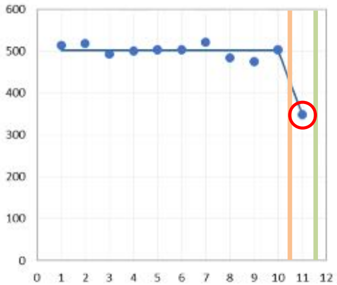


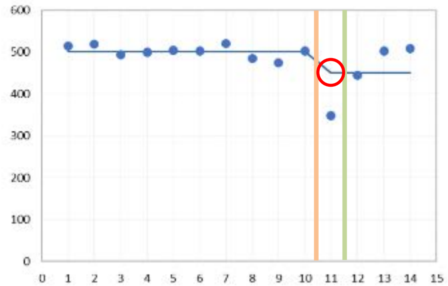


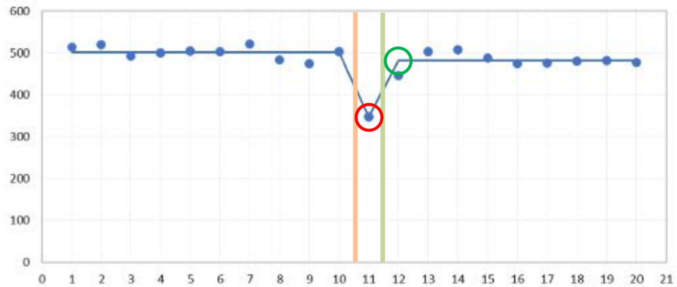


.io

# Topics

LF NETWORKING
LFN Developer & Testing Forum

- FD.io Overview
  - VPP and CSIT core projects
- FD.io CSIT
  - Daily Trending and Release Reports
  - Anomaly Detection
  - Data Processing Infrastructure
- Wrap-up
- Q&A

# Infrastructure: Data Producers



- CI/CD pipeline
  - daily, weekly, ad-hoc verify jobs
  - multiple performance testbeds
  - multiple releases
  - thousands of unique performance tests executions per day
- local execution
  - ad-hoc manual verify
  - distinct testbeds
- **json formatted output**
  - schema/model defined
  - stored in nested directories

# Infrastructure: Data Stores

LF NETWORKING
LFN Developer & Testing Forum

- **AWS S3 cloud storage**
  - migrated from LFN IT managed Nexus
  - AWS CloudFront
- S3 local storage
  - gateway cache
  - backup
  - NGINX proxied
- local filesystem

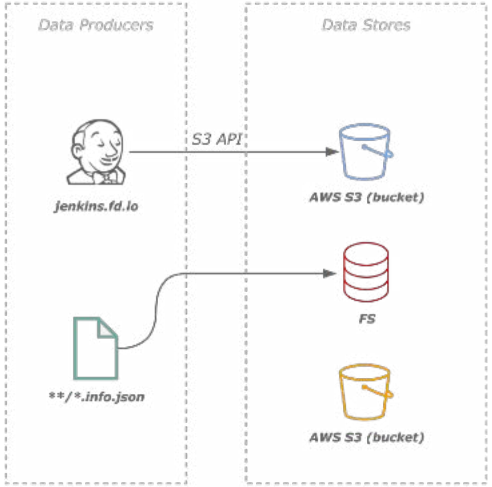


fd.io

# Infrastructure: Data ETL

LF NETWORKING
LFN Developer & Testing Forum

- **ETL scripts using pySpark**
  - extract, transform, load
  - pre-process semi-formatted data into partitioned data frames
  - stateless
  - eliminate the need of real-time processing of the same data twice
- data loaded
  - directly from filesystem
  - S3 compatible storage via S3 API
- optionally use AWS services
  - AWS GLUE

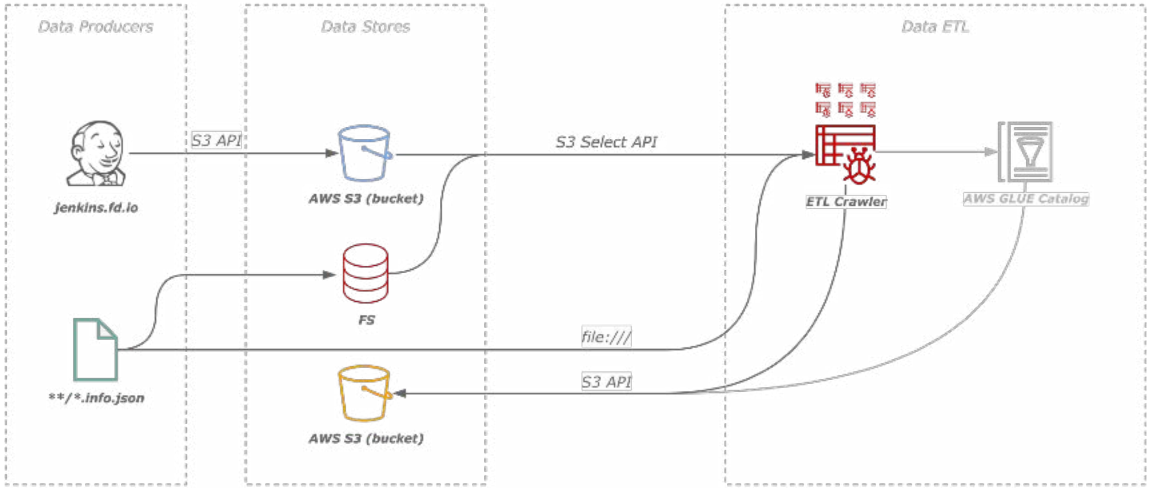


.io

# Infrastructure: Data View

LF NETWORKING
LFN Developer & Testing Forum

- **AWS beanstalk**
  - application load balancing
  - EC2 instance
- **Plot.ly Dash**
  - data dashboard
  - interactive UI
  - loads data frame partitions from S3 compatible storage
  - loads data frame partitions from local storage
- optionally use AWS services
  - AWS Athena

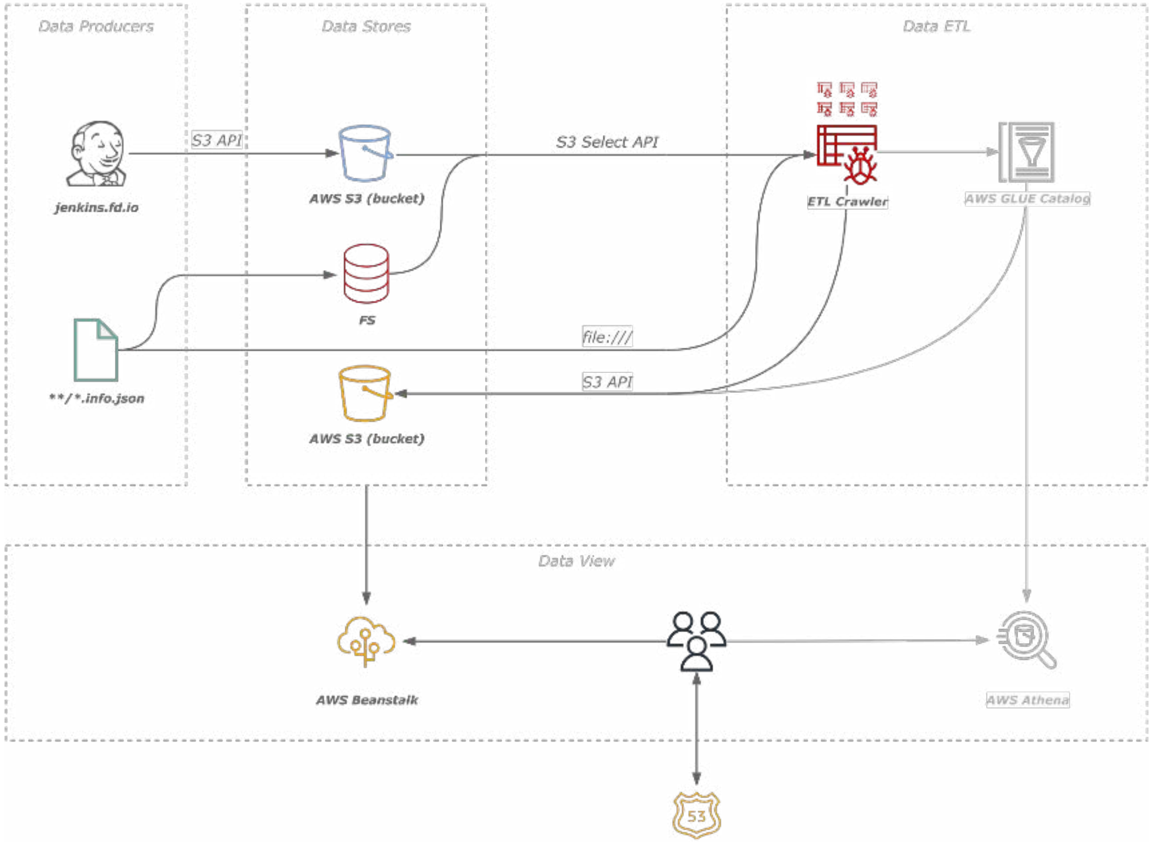

# Topics

LF NETWORKING
LFN Developer & Testing Forum

- FD.io Overview
  - VPP and CSIT core projects
- FD.io CSIT
  - Daily Trending and Release Reports
  - Anomaly Detection
  - Data Processing Infrastructure
- Wrap-up
- Q&A

fd.io

# Wrap-up

- **VPP** and **CSIT FD.io** projects provide technologies needed to build Terabit-level Network Functions and Services
  - **VPP** achieves best-in-class performance on general purpose compute servers.
  - **CSIT** provides end-to-end open source tooling for testing and benchmarking Terabit-level Network Functions.
- **FD.io** is an open and collaborative community
  - We are eager to work with you!

# CSIT Resources

LF NETWORKING
LFN Developer & Testing Forum

- **Project**
  - Wiki pages: https://wiki.fd.io/view/CSIT
  - Meetings: https://wiki.fd.io/view/CSIT/Meeting
  - Mailing list: csit-dev@lists.fd.io
- **Performance Trending**
  - Daily and Weekly: https://s3-docs.fd.io/csit/master/trending/index.html
- **Release Reports**
  - Latest CSIT-2110: https://s3-docs.fd.io/csit/rls2110/report/
  - Previous: https://wiki.fd.io/view/CSIT#Test_Reports
- **Source Code**
  - Git repo: https://git.fd.io/csit
  - Github mirror: https://github.com/FDio/csit
  - Gerrit reviews: https://gerrit.fd.io

# FD.io CSIT Resources



- **Technical Papers**
  - “Benchmarking Software Data Planes Intel® Xeon® Skylake vs. Broadwell” (2019)
    - https://www.lfnetworking.org/wp-content/uploads/sites/55/2019/03/benchmarking_sw_data_planes_skx_bdx_mar07_2019.pdf
  - “Benchmarking and Analysis of Software Data Planes” (2017)
    - https://fd.io/docs/whitepapers/performance_analysis_sw_data_planes_dec21_2017.pdf
- **Technical Marketing Video Clips**
  - “VPP: A Terabit Secure Network Data-plane” (Intel Xeon Icelake 07-APR-2021)
    - https://www.youtube.com/watch?v=ipQQmjzE_g0
  - “FD.io: A Universal Terabit Network Dataplane" (Intel Xeon Skylake, 11-JUL-2017)
    - https://www.youtube.com/watch?v=aLJ0XLeV3V4
- **FD.io Presentations**
  - https://wiki.fd.io/view/Presentations
- **Other FD.io Materials**
  - https://fd.io/
  - https://fd.io/latest/whitepapers/

# Topics

LF NETWORKING
LFN Developer & Testing Forum

- FD.io Overview
  - VPP and CSIT core projects
- FD.io CSIT
  - Daily Trending and Release Reports
  - Anomaly Detection
  - Data Processing Infrastructure
- Wrap-up
- Q&A

.io

LF NETWORKING

LFN Developer & Testing Forum

# FD.io CSIT Performance Benchmarking, Anomaly Detection, Results Tracking.

FD.io

# Backup slide: AD Compression

LF NETWORKING
LFN Developer & Testing Forum

- Computing bits for a group:
  - How many samples?
    - Probability of N samples is 1/N – 1/(N+1)
      - N=0 is not allowed
      - Sum of probabilities is 1 as it should be.
  - What is the new average?
    - Probability has “inverse triangle” distribution.
      - Probability density increases linearly by absolute difference from previous group average.
        - This is the non-obvious tweak.
        - Without this, the algorithms sees too many tiny (and false) regressions and progressions.
      - For first group, density is uniform from zero to max value.
  - What is the new standard deviation?
    - Probability density is uniform from zero to max value.
      - Dumb but safe.
  - The positions of each sample.
    - Luckily constant when size average and stdev are fixed.
      - It is an “area” of multi-dimensional sphere.
- Caveats:
  - If samples are integers, average and stdev are not.
  - If samples are floats, number of significant bits matter.
  - Algorithm assumes all is integer for probability density purposes.
    - But promptly uses formulas only valid for floats, e.g. for sample position information.
  - Upscale your samples to avoid errors due to this integer/float mismatch.

.io

# FD.io VPP on Intel® Xeon® Processors

LF NETWORKING
LFN Developer & Testing Forum

Increased I/O and Cipher Processing Improving Crypto Packet Throughput

## Intel Xeon "Cascadelake"

YESTERDAY

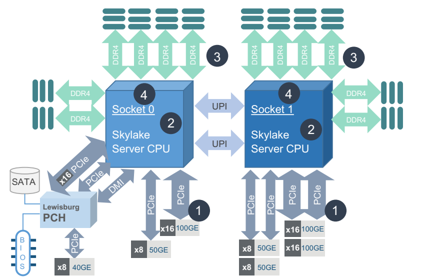

1. Network I/O PCIe Gen3: **320 Gbps**
2. Per Core Caches: L1d **32 KB**, L2 **1.0 MB**
3. Memory: **6-channels** 2666 MHz
4. Max power: **150 W** (TDP)

Intel® Xeon® 6252N
24 Cores, 2.3 GHz, 33 MB Cache

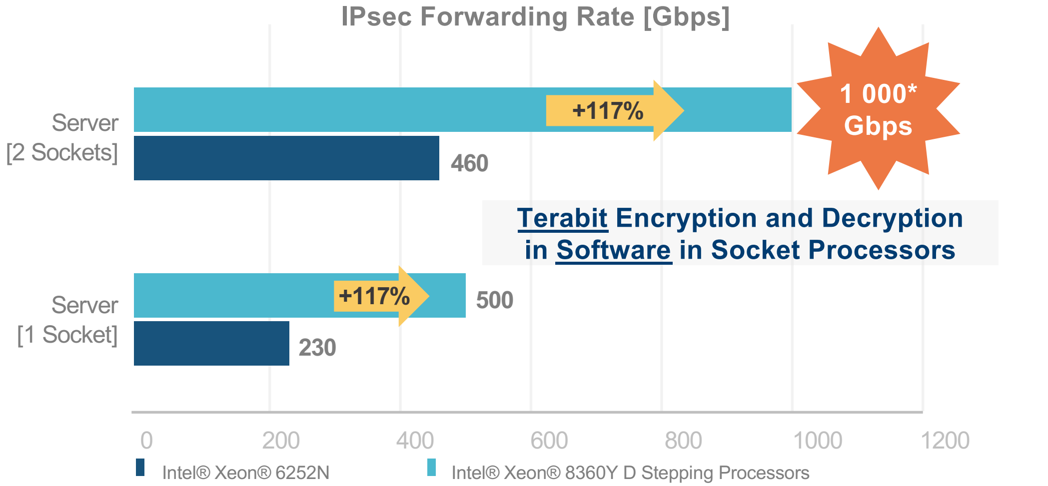


* IPsec forwarding rate bound by multiple instruction execution factors. Measured for 1024B cleartext frame size.

## Intel Xeon "Icelake"

TODAY

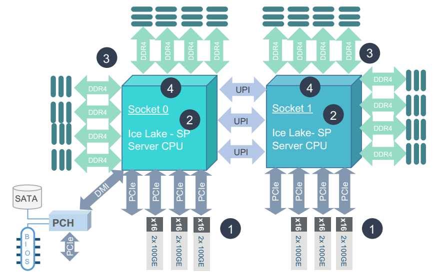

1. Network I/O PCIe Gen4: **600 Gbps**
2. Per Core Caches: L1d **48 KB**, L2 **1.25 MB**
3. Memory: **8-channels** 3200 MHz
4. Max power: **250 W** (TDP)

Intel® Xeon® 8360Y D Stepping
36 Cores, 2.4 GHz, 54 MB Cache

**Breaking the Terabit Barrier of Secure NaaS Services**
**on a Single 2-Socket Intel® Xeon® Server !**

intel XEON PLATINUM

"VPP: A Terabit Secure Network Data-plane" https://www.youtube.com/watch?v=ipQQmjzE_g0 (3min 21sec clip)