LF NETWORKING
Developer & Testing Forum

FD.io CSIT Performance Dashboard

Presented on behalf of CSIT project by:
Tibor Frank, Maciek Konstantynowicz, Peter Mikus, Vratko Polak

https://lfnetworking.org

Topics

FD.io Overview

- VPP and CSIT core projects

CSIT Dashboard - csit.fd.io

- Infrastructure
- UI and presentation layer

Usability Examples

- Performance and efficiency comparisons
- Failures, Anomalies and Root Cause Analysis

Discussion

FD.io VPP and CSIT
Core Projects in LFN FastData.io

LF NETWORKING
LFN Developer & Testing Forum

LF NETWORKING fd.io	**VPP** Vector Packet Processing	**CSIT** Continuous System Integration and Testing
Performance	Feature rich networking and host stack. VPP on COTS servers in many cases **outperforms** packet processing HW.	Continuous **benchmarking** of VPP and DPDK.
Portability	VPP runs on COTS hardware: AMD ARM intel OpenPOWER POWER8 VPP runs in any environment: **bare-metal, VM, containers**.	Performance testbeds with **Xeon**, **Arm**, **AMD**, **Atom** HW. **Bare-metal**, **VM**, **container** test environments.
Efficiency	Allows ability to **upscale** and **downscale.**	Executing **2,900 benchmarking tests** daily.
SDN	Software **programmable**, **extendable** and **flexible.**	Open-source **CI/CD infrastructure for benchmarking** of SW data-planes, **test data analytics** and **presentation**.

CDash

AWS S3 storage

- used for all data.

ETL

- from JSON to Parquet.
- running on premise.

AWS beanstalk

- application load balancing.
- EC2 instances (t3a.2xlarge).
- scalability.

Plot.ly Dash

- data dashboard.
- interactive UI.
- loads data frame partitions from S3 compatible storage.

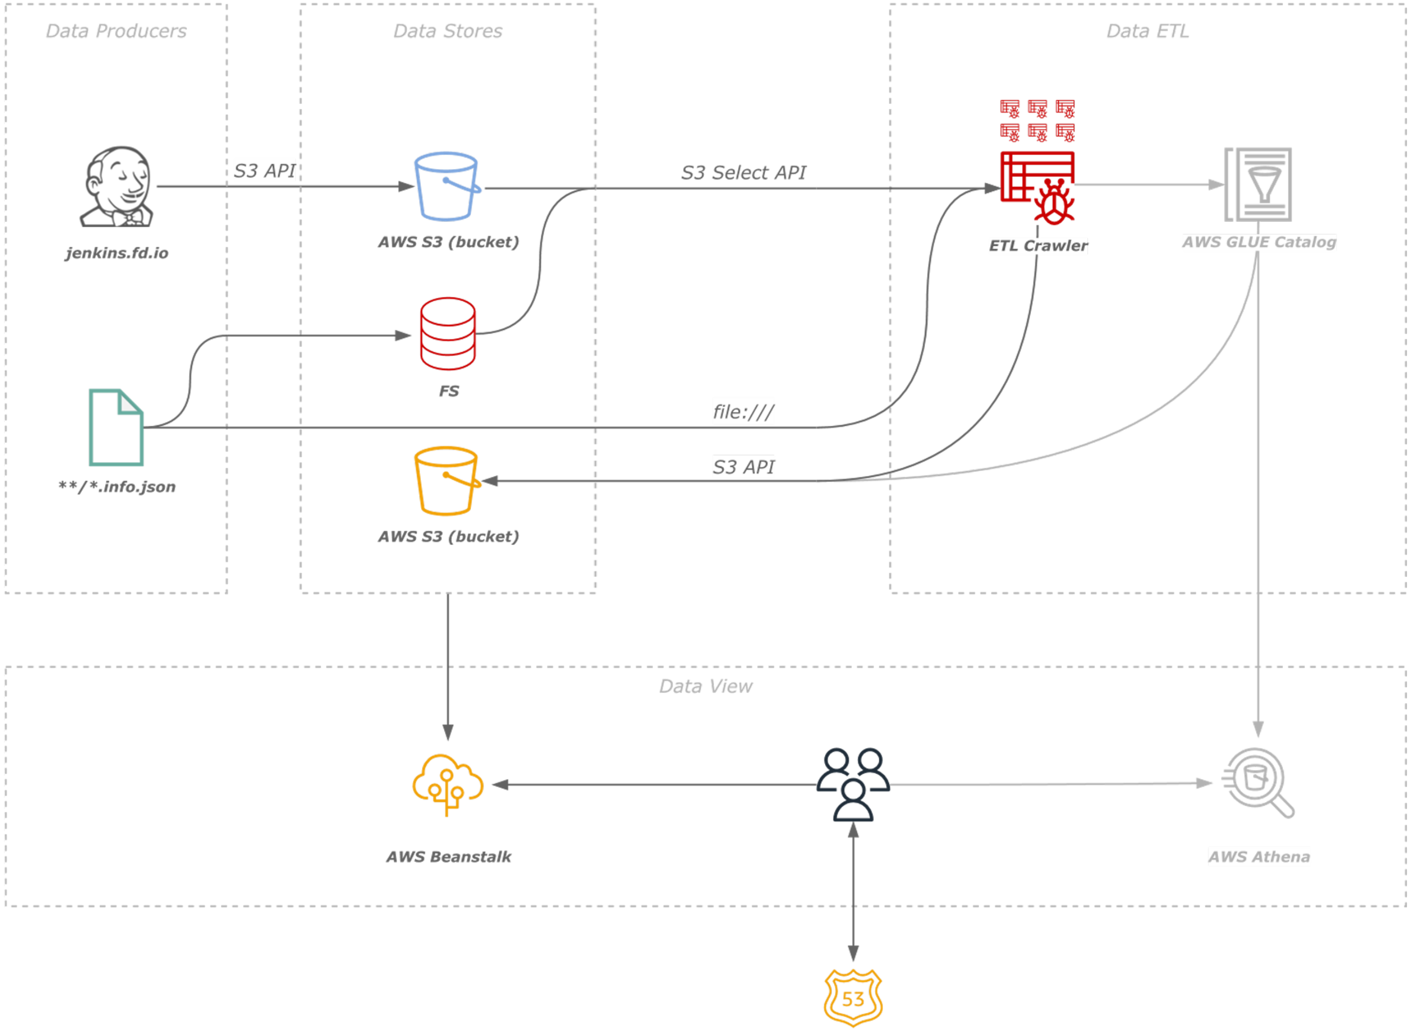

CDash

17
ETL pipelines

JSON
model

~2.9k
tests daily

~50k
performance tests per release

optimization
8GB RAM

Parquet

~17.5k
tests weekly

rls2302
rls2306

2021 ...

130 days
sliding window

... 2023

CSIT-Dash

Customizable Views on Performance Data

Performance Trending

- Daily MRR* Data
- Weekly NDR PDR** Data
- Packet Latency
- Telemetry

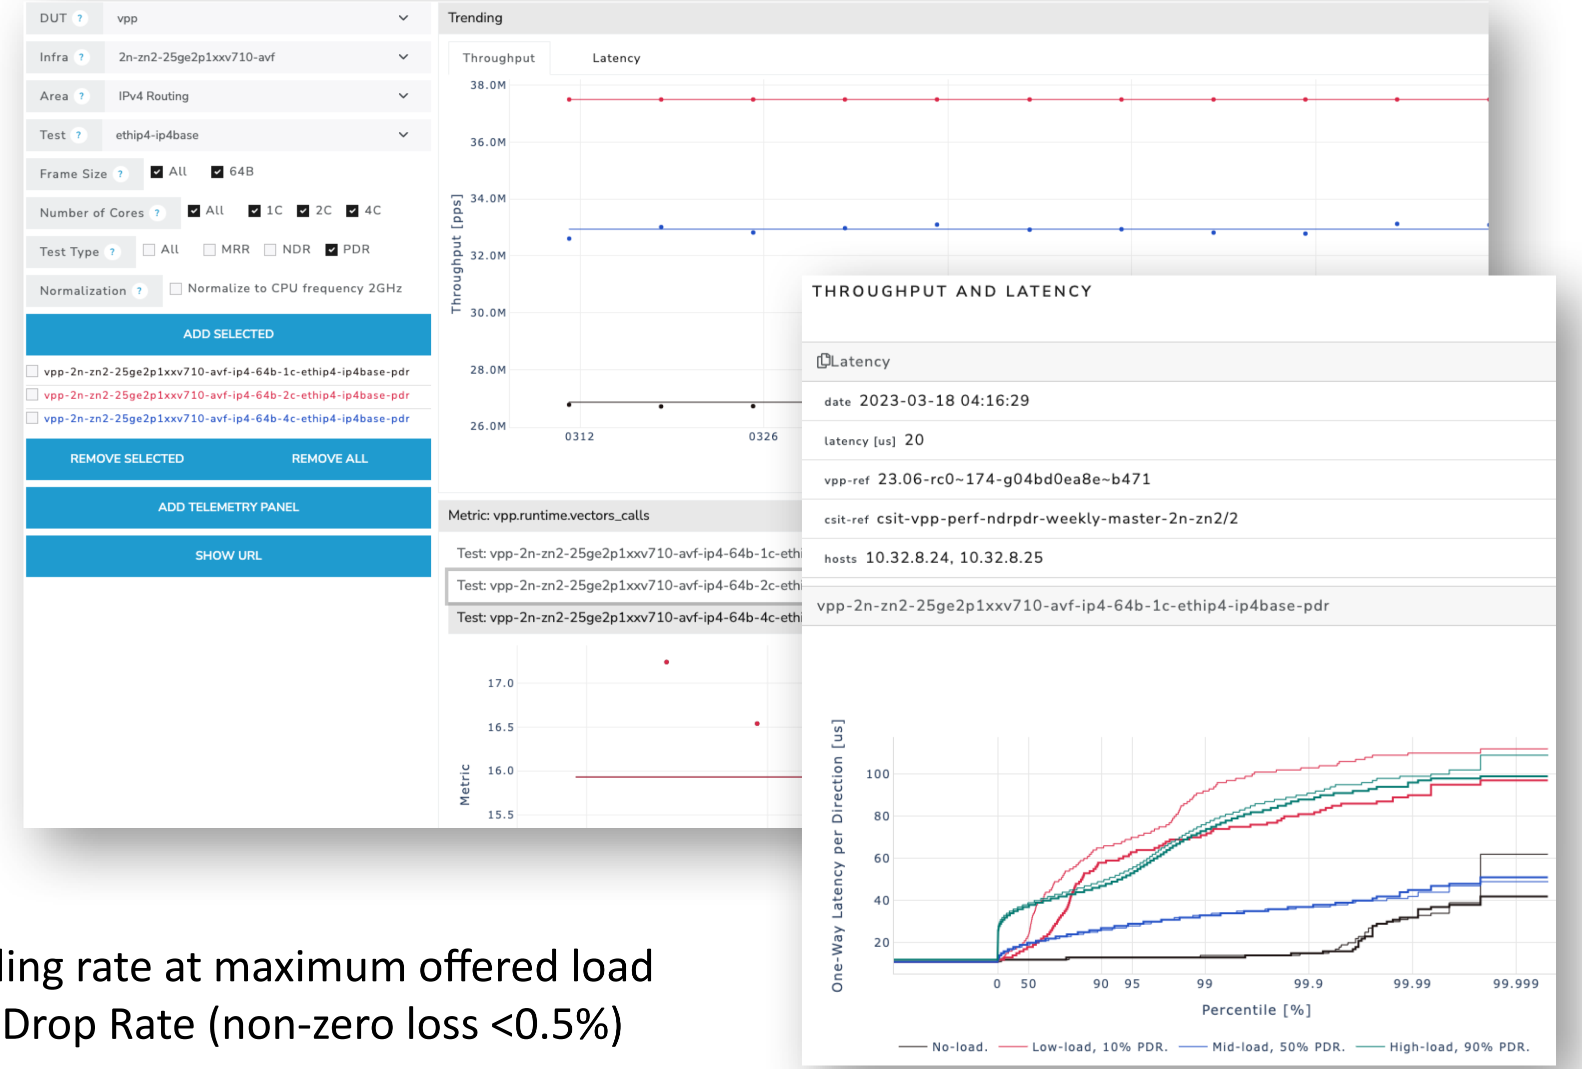

* MRR – Maximum Receive Rate a.k.a. forwarding rate at maximum offered load

** NDR PDR – Non-Drop Rate (no loss), Partial Drop Rate (non-zero loss <0.5%)

CSIT-Dash

Customizable Views on Performance Data

Release Performance Data

- Iterative and Coverage Data
- Packet Throughput
- Packet Latency
- Results Statistics
- Comparison Tables

CSIT-Dash

Customizable Views on Performance Data

LF NETWORKING
LFN Developer & Testing Forum

Additional Information

- Test Job Statistics
- Failures and Anomalies
- Documentation

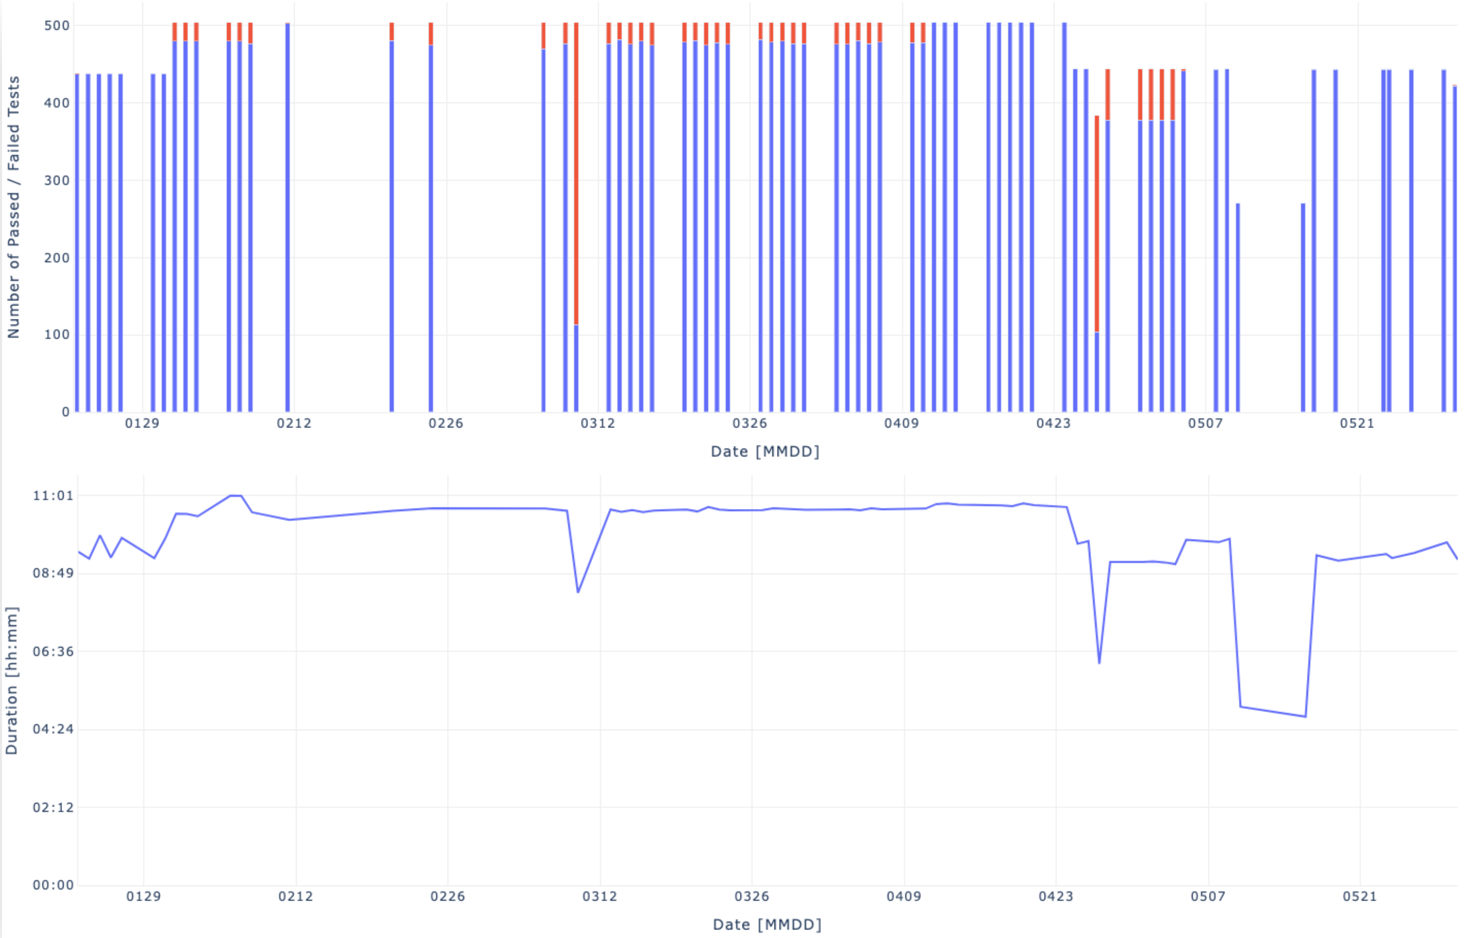

CSIT Benchmark Areas and Methodologies

Tests

Benchmark Test Areas

- L2 Ethernet Switching
- IPv4, IPv6 Routing
- IPsec, Wireguard with IPv4 Routing
- SRv6 Routing
- Features: ACLs, NAT44-EI/ED, Policer, …
- IPv4, IPv6 Tunnels
- KVM VMs vHost-user
- Docker Container Memif
- Drivers: DPDK, AVF, RDMA, AF_XDP

Test Methodologies

- Packet Throughput and Latency
- Stateful NAT44ed
- Stateful Host-stack
- Speedup Multi-Core
- Soak Tests
- Reconfiguration Tests

Test Topologies

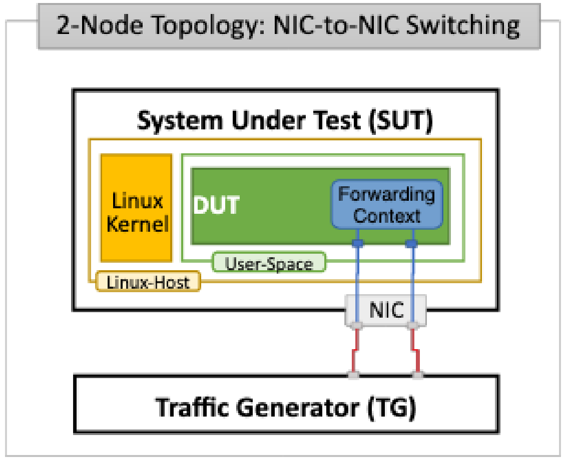

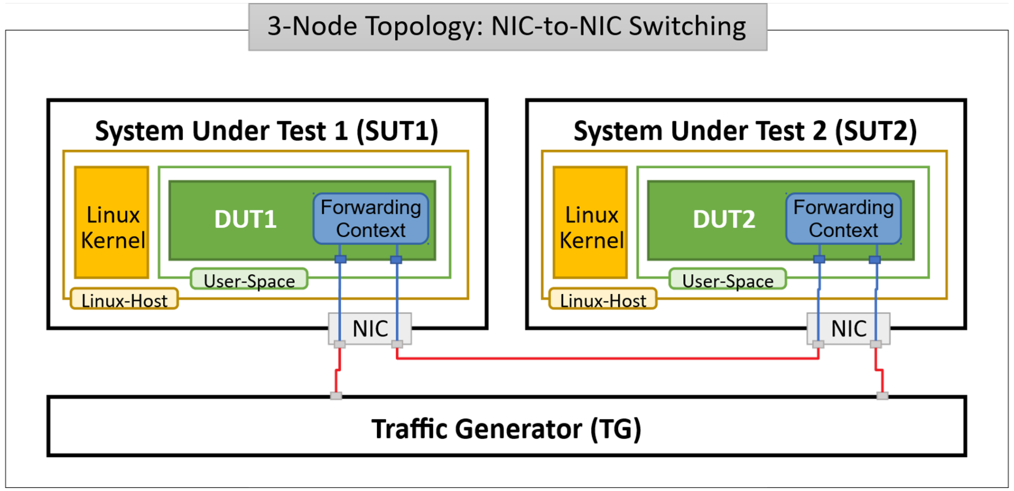

Performance Testbed Variants*

2n–clx
2n–icx
2n–spr
2n–tx2
2n–zn2

3n–icx
3na–spr
3nb–spr
3n–alt
3n–snr
3n–tsh

*Testbed Topology – SUT Processor Model

Performance testing: List of compute platforms

LF NETWORKING
LFN Developer & Testing Forum

Processor Family	Model	Cores per Socket	Base Frequency GHz	L3 Cache (LLC) MB	Testbeds	NICs
CLX Intel Xeon Cascade Lake	8280	28	2.6	38.5	3 x 2n-clx	x710-4p10GE, xxv710-2p25GE, e810-2p100GE, cx5-2p100GE
ICX Intel Xeon Ice Lake	8358	32	2.6	48	4 x 2n-icx 2 x 3n-icx	xxv710-2p25GE, e810-4p25GE, e810-2p100GE, cx7-2p200GE* xxv710-2p25GE, e810-4p25GE, e810-2p100GE, cx6-2p200GE*
SPR Intel Sapphire Rapids	8462Y+	32	2.8	60	2 x 2n-spr 1 x 3na-spr 1 x 3nb-spr	e810-4p25GE, e810-2p100GE, cx7-2p200GE cx7-2p200GE e810-4p25GE, e810-2p100GE
ZN2 AMD EPYC Zen2	7532	32	2.4	256	1 x 2n-zn2	x710-4p10GE, xxv710-2p25GE, cx5-2p100GE
TX2 Cavium ThunderX2 ARMv8.1	CN9975	28	2.0	28	1 x 2n-tx2	xl710-2p40GE
ALT Ampere Altra N1	Q80-30	80	3.0	32	1 x 3n-alt	xl710-2p40GE cx7-2p200GE*
SNR Intel Atom Snowridge	P5362B	24	2.2	27	1 x 3n-snr	e810-4p25GE
TSH Huawei TaiShan 2280	hip07-d05	32	2.2	-	1 x 3n-tsh	x520-2p10GE, cx4-2p25GE

* Being added

VPP release performance comparison views
Performance: Xeon **ICX** vs **CLX**

LF NETWORKING
LFN Developer & Testing Forum

Comparison VPP v23.02 ICX vs CLX:

COMPARISON FOR: VPP-RLS2302-23.02-RELEASE-1C-64B|78B-MRR

	2N-CLX-100GE2P1E810CQ-AVF		2N-ICX-100GE2P1E810CQ-AVF		RELATIVE CHANGE	
TEST NAME	MEAN [MPPS]	STDEV [MPPS]	MEAN [MPPS]	STDEV [MPPS]	MEAN [%]	STDEV [%]
ip4	CLX		ICX		ICX vs CLX	
64b-1c-ethip4-ip4base-eth-2memif-1dcr-mrr	7.49	0.18	8.80	0.03	17.57	2.82
64b-1c-ethip4-ip4scale20k-mrr	16.22	0.03	18.75	0.16	15.57	1.04
64b-1c-ethip4udp-ip4base-oacl50sf-10kflows-mrr	10.08	0.02	11.63	0.05	15.36	0.53
64b-1c-ethip4-ip4scale20k-rnd-mrr	16.57	0.07	18.80	[illegible]	13.51	0.52
64b-1c-ethip4udp-ip4base-iacl50sf-10kflows-mrr	10.61	0.20	11.62	[illegible]	9.57	2.52
64b-1c-ethip4-ip4base-mrr	18.99	0.15	20.68	0.02	8.91	0.89
64b-1c-ethip4udp-ip4base-oacl50sl-10kflows-mrr	8.34	0.01	8.99	0.05	7.79	0.66
64b-1c-ethip4udp-ip4base-iacl50sl-10kflows-mrr	8.64	0.15	9.06	0.10	4.88	2.19

Show URL Download Data

COMPARISON FOR: VPP-RLS2302-23.02-RELEASE-1C-64B|78B-MRR

	2N-CLX-100GE2P1E810CQ-AVF		2N-ICX-100GE2P1E810CQ-AVF		RELATIVE CHANGE	
TEST NAME	MEAN [MPPS]	STDEV [MPPS]	MEAN [MPPS]	STDEV [MPPS]	MEAN [%]	STDEV [%]
ip6	CLX		ICX		ICX vs CLX	
78b-1c-ethip6-ip6scale20k-mrr	13.01	0.45	14.84	0.09	14.01	3.99
78b-1c-ethip6-ip6scale20k-rnd-mrr	12.96	0.64	14.54	[illegible]	12.22	5.55
78b-1c-ethip6-ip6base-mrr	15.88	0.03	16.96	0.03	6.81	0.28

Show URL Download Data

Similar views for Xeon **SPR** vs **ICX** – results just arrived from VPP v23.06-rc1 pre-release testing.

Introducing VPP Per Node Metrics

LF NETWORKING
LFN Developer & Testing Forum

VPP per node metrics collected in CSIT performance tests

VPP "show runtime"

Calls
Clocks
Suspends
Vectors
Vectors/Call

VPP perfmon bundles

Intel platforms: vpp git repo /src/plugins/perfmon/bundle/intel/

inst-and-clock
- Calls
- Packets
- Packets/Call,
- Clocks/Packet
- Instructions/Packet
- IPC

cache-hierarchy
- L1 hit/pkt
- L1 miss/pkt
- L2 hit/pkt
- L2 miss/pkt
- L3 hit/pkt
- L3 miss/pkt

Arm platforms: vpp git repo /src/plugins/perfmon/bundle/arm/

Currently not running in CSIT
Perfmon access to Arm counters requires newer kernel 5.17+ – work in progress …

NOTE: VPP perfmon is VPP Developer Tool
Heavily dependant on Processor PMU counters
Proceed with caution !

Introducing VPP Telemetry Views

LF NETWORKING
LFN Developer & Testing Forum

https://csit.fd.io/trending/

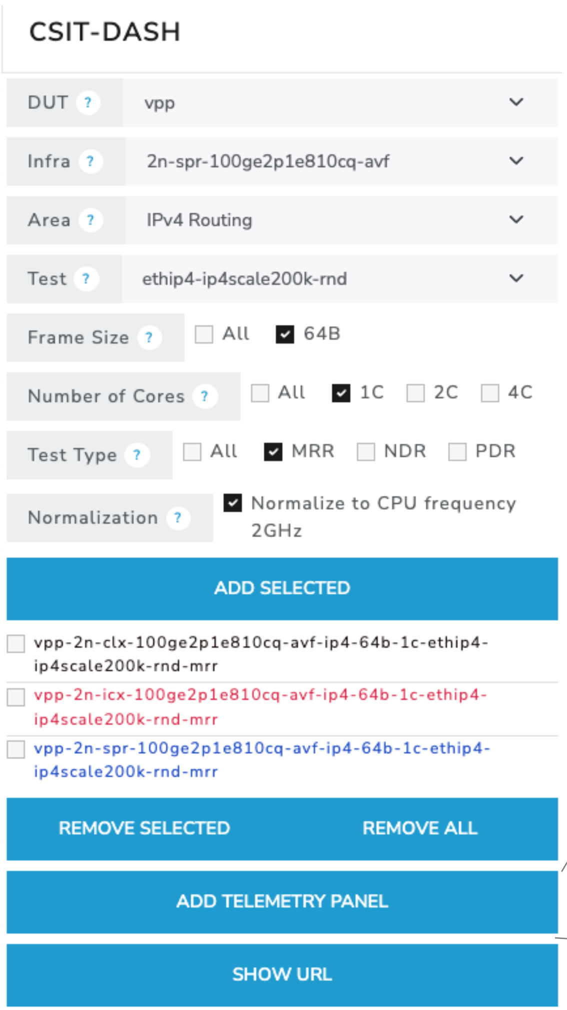

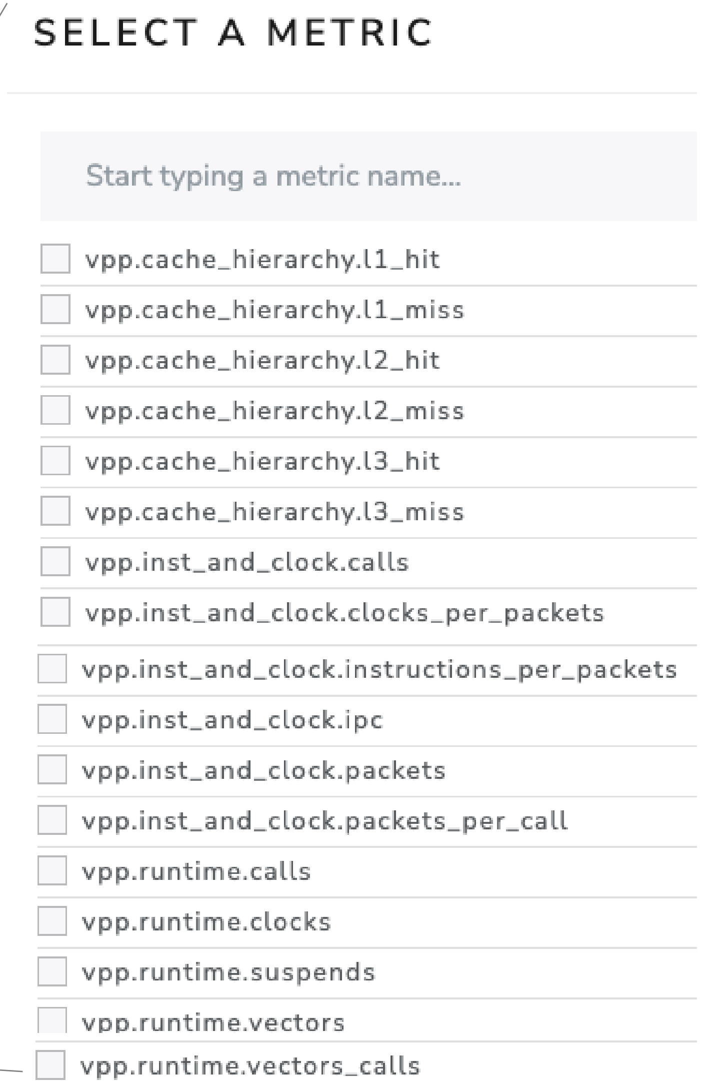

Performance trending – Throughput and efficiency metrics Xeon **SPR** vs **ICX** vs **CLX**

LF NETWORKING
LFN Developer & Testing Forum

Test	Measurement	VPP Nodes of Interest	Node Telemetry
ip4scale200k-rnd	MRR	ip4-lookup	runtime.clocks inst_and_clock.clocks_per_packet cache_hierarchy.{l3|l2|l1}_{hit} cache_hierarchy.{l3|l2|l1}_{miss}
ip6scale200k-rnd	MRR	ip6-lookup	runtime.clocks inst_and_clock.clocks_per_packet cache_hierarchy.{l3|l2|l1}_{hit} cache_hierarchy.{l3|l2|l1}_{miss}
ipsec40tnlsw-aes256gcm	MRR	esp4-encrypt-tun esp4-decrypt-tun	runtime.clocks inst_and_clock.clocks_per_packet cache_hierarchy.{l3|l1}_{hit} cache_hierarchy.{l3|l1}_{miss} runtime.vectors_calls

Performance trending – Throughput and efficiency metrics VPP **IPv6** and **IPV4 FIB scale** tests

LF NETWORKING
LFN Developer & Testing Forum

Testbed-Platform	Test	Measurement	VPP Nodes of Interest	Telemetry
2n-icx	ip6base ip6scale20k-rnd ip6scale200k-rnd ip6scale2m-rnd	MRR	ip6-lookup	runtime.clocks cache_hierarchy.{l3|l1}_{hit} cache_hierarchy.{l3|l1}_{miss}
2n-icx	ip4base ip4scale-2m ip4scale-2m-rnd	MRR	ip4-lookup	inst_and_clock.clocks_per_packet cache_hierarchy.{l3|l1}_{hit}

Failures, Anomalies and Root Cause Analysis

Link	Comment
Test cases failure	All test cases affected by an infra issue.
Test suites failure	Reduction in test cases executed can be caused by failure in suite setups.
Anomalies and testbeds	Different testbeds of the same type show unequal performance, depending on NIC and driver.
Anomalies and RSS	Random RSS generates noise. Peformance change depends on testbed type and scale.
Scaling and anomalies	Comparing scalability of a test, includes a transient regression and failures.

CSIT Resources

- **Technical Papers**
  - SPR 2Tbps IPsec (2023)
    - https://networkbuilders.intel.com/solutionslibrary/intel-avx-512-high-performance-ipsec-with-4th-gen-intel-xeon-scalable-processor-technology-guide
  - “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
- **Technology Demonstrator 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/

CSIT Resources

- **Project**
  - Wiki pages: https://wiki.fd.io/view/CSIT
  - Meetings: https://wiki.fd.io/view/CSIT/Meeting
  - Mailing list: csit-dev@lists.fd.io
- **CDash**
  - Dashboard: https://csit.fd.io
- **Previous Release Reports**
  - Wiki: 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
- **Standalone libraries**
  - Speeding up binary search using shorter measurements: https://pypi.org/project/MLRsearch/
  - Locating changes in time series by grouping results: https://pypi.org/project/jumpavg/

LF NETWORKING

Developer & Testing Forum

Q&A

https://lfnetworking.org