Evolution of Repeatability in Benchmarking: Fraser Plugfest (Summary for IETF BMWG)

Sridhar Rao, Spirent Communications

Al Morton, AT&T Labs

VSPERF Project Team

Problem Statement

- For a given Device Under Test (DUT) there are Multiple Dimensions of Test Results Repeatability (and issues)
  - Over Time (repeated tests using the same conditions)
  - Over Minor Changes
  - Over “Identical” Nodes implementing the same DUT config
  - Using different Test Management Tools
  - Using different Test Equipment
    - Evaluated at Danube Plugfest/OPNFV Summit 2017
- This campaign covers Time, with multiple Nodes if possible, and multiple Tools if possible.
- Initial focus on Search Algorithms and their contribution to Repeatability.

Test Combinations

- Select test configurations with apparent issues
- Start with VSPERF CI Parameters, but using Trex Traffic Gen
- Evaluate consistency when applying different (new) Search Algorithm(s)
- Background – Search Algorithms are evolving to satisfy multiple roles:
  - Highly repetitive testing (CI or CSIT)
  - Multiple Search Goals
  - Reliable Results in purely Experimental Scenario
    - Zero loss Throughput level is the basis for many other benchmarks

VSPERF LTD Supported Deployment Scenarios

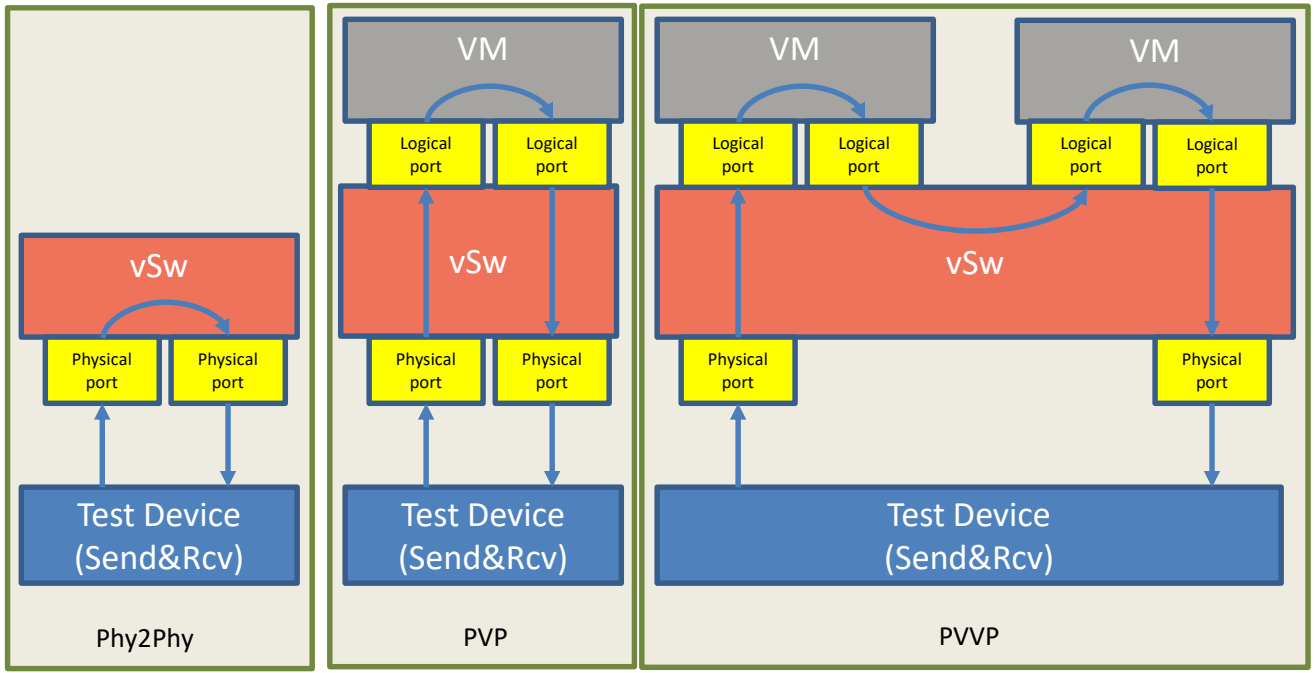

TST009: Hierarchy of Method, Sets, Tests, and Trials

Method (iterate over multiple frame sizes)

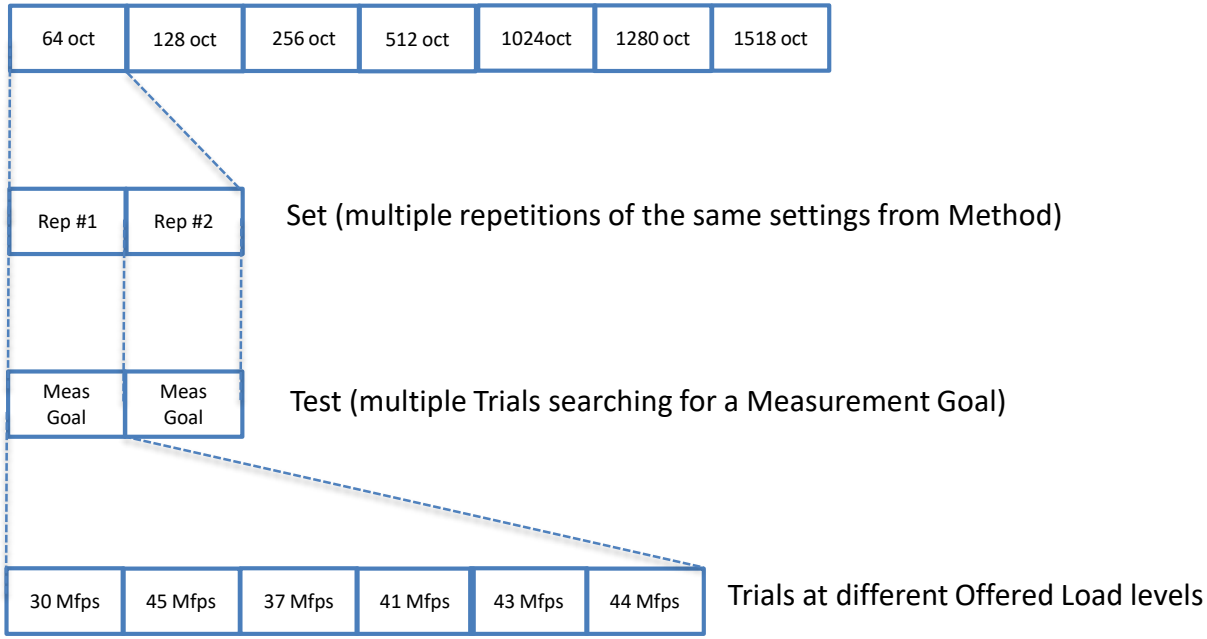

Day in the Life of a Classic “Benchmarking” Binary Search

Trial #	%MaxRate	rx_fps	lost frames	loss %ge	Duration
1	100	6882195	426573461	72.83	23.131
2	50	6903766	189253220	64.582	15.106
3	25	6914303	42679111	29.134	15.103
4	12.5	4881245	237	0	15.102
5	6.25	2440822	694	0.002	15.092
6	3.125	1190880	0	0	15.12
7	4.688	1822732	436	0.002	15.1
8	3.906	1525229	0	0	15.091
9	4.297	1675920	392	0.002	15.114
10	4.102	1599537	0	0	15.091
11	4.199	1637255	0	0	15.101
12	4.248	1656431	0	0	15.102

TST009 Annex B: Mitigating background processes that cause Errors (loss)

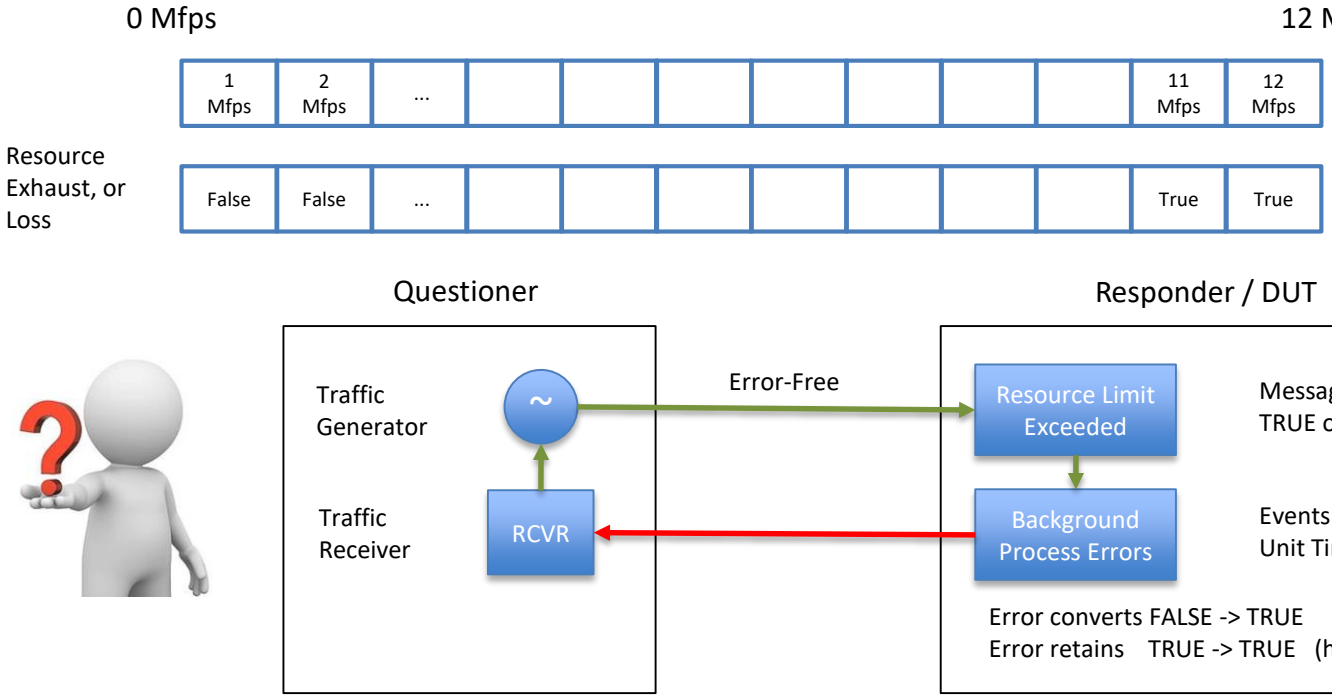

Andrzej Pelc, "Searching games with errors— fifty years of coping with liars ", Theoretical Computer Science 270 (2002) 71–109. Available from https://www.gwern.net/docs/statistics/comparison/2002-pelc.pdf

Timing of Transients and avoiding them during Searches

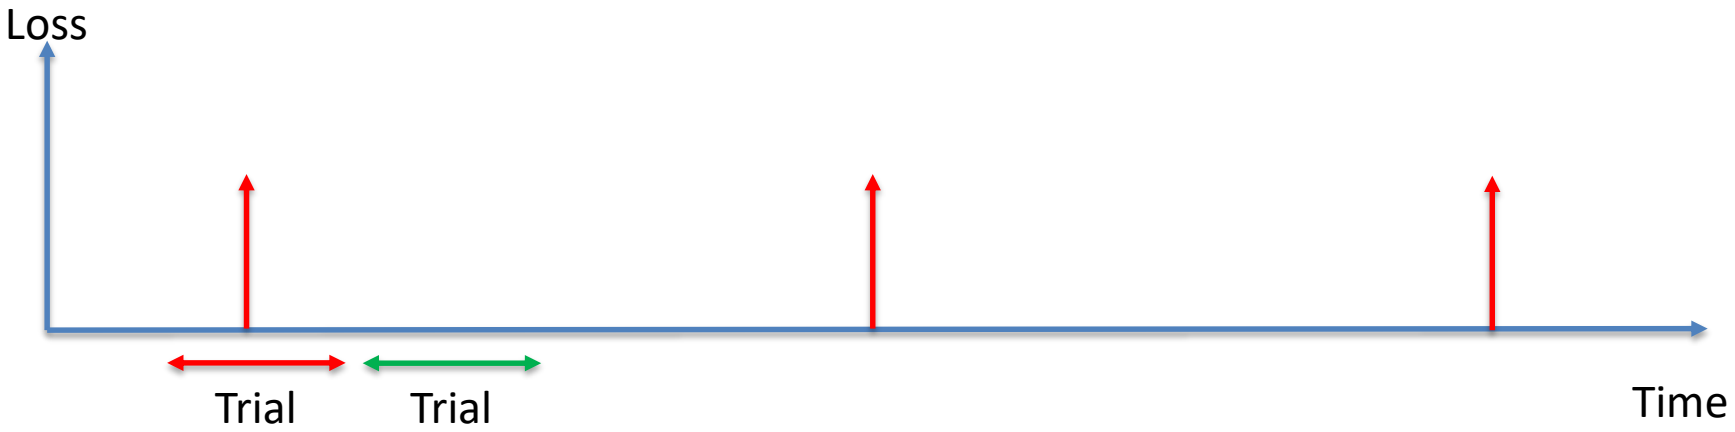

Binary Search with Loss Verification (BSwLV) Parameters

see https://docbox.etsi.org/ISG/NFV/Open/Drafts/TST009_NFVI_Benchmarks (12.3) for version 0013, the latest as of July 2018. also https://jira.opnfv.org/browse/VSPERF-583

- Minimum Step size, or nominal load level change (0.05%)
- Max number of Trial repetitions (2) and Loss Thresh,v2 (2k)
- Trial Duration (30s, 15s, but best determined after long-term test discovers frequency of transient interruptions)
- Frame Sizes (64, 128, try IMIX in future)
- Number of repeated tests (4 max)

For this testing campaign, we'll also collect/calculate:

- Number of repeated tests where loss outcome changed
- Metric on consistency

Long Duration Tests (planned)

- Conduct at (or near) Zero-Loss Throughput Load Level
- Characterize Transient Events:
  - Ave Event Frequency and Period
  - Extent of Loss During an Event
  - Consider that there may be multiple Event Loss Signatures
- Inter-Event Period > Trial Duration for BSwLV

Comparison: Trex + OVS_DPDK + Pod12_Node4_DUT
30 Sec Trial Duration; Throughput in Frames per Second

Setup	Search Alg	64 octet	64 octet	64 octet	128 octet	128 octet	128 octet
p2p	BS rx	23,038,250	23,480,830	21,861,407	15,630,265	16,872,340	16,889,285
p2p	BS %	59.082	60.205	56.201	81.201	86.475	86.768
p2p	BS #it	12	12	12	12	12	12
p2p	BSwLV rx	23,378,172	22,867,154	23,383,808	16,851,926	16,744,952	16,889,797
p2p	BSwLV %	60.205	60.205	60.205	90.625	86.719	100
p2p	BSwLV #it	19	19	20	21	21	2
p2p	BSwLV reversal	0	0	0	1	2	1

"reversal" means that repeated trial result was zero loss

Comparison: Trex + OVS_DPDK + Pod12_Node4_DUT
30 Sec Trial Duration; Throughput in Frames per Second

Setup	Search Alg	64 octet	64 octet	64 octet	128 octet	128 octet	128 octet
pvp	BS rx, fps	1,617,550.8	1,636,372.6	1,640,444	1,646,015	2031533	1,719,593
pvp	BS %	4.15	4.199	4.199	8.447	10.449	8.838
pvp	BS #it	12	12	12	12	12	12
pvp	BSwLV rx	3,762,113	4396882	2,415,518	4,980,638	4,215,393	4,320,085
pvp	BSwLV %	9.668	11.279	6.201	25.488	21.631	22.119
pvp	BSwLV #it	23	21	21	22	22	21
pvp	BSwLV reversal	3	3	4	2	5	2

"reversal" means that repeated trial result was zero loss

Comparison: Trex + OVS_DPDK + Pod12_Node4_DUT 30 Sec Trial Duration; 64 octet Frames

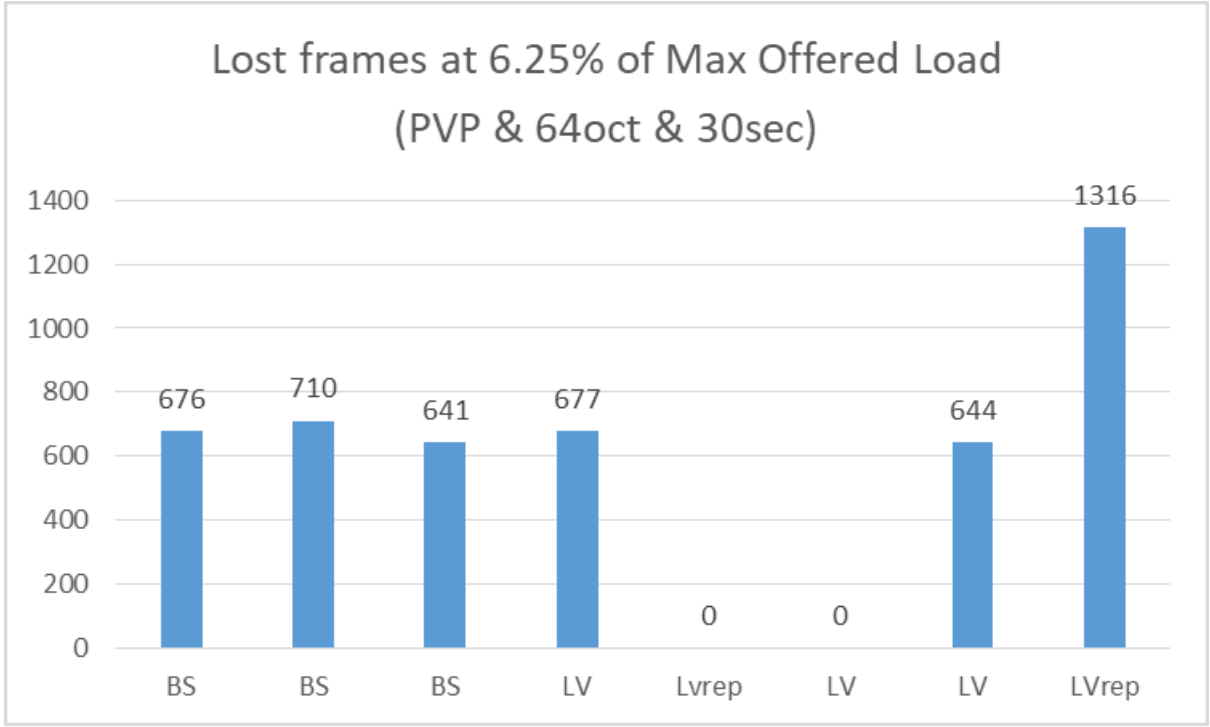

Binary Search=BS, Loss Verification=LV, repeated trial = LVrep,
All trials at 6.25% of Max Offered Load

Comparison: Trex + OVS_DPDK + Pod12_Node4_DUT
15 Sec Trial Duration; 64 octet Frames

BS THPT, %max
4.541
9.131
4.443
9.473

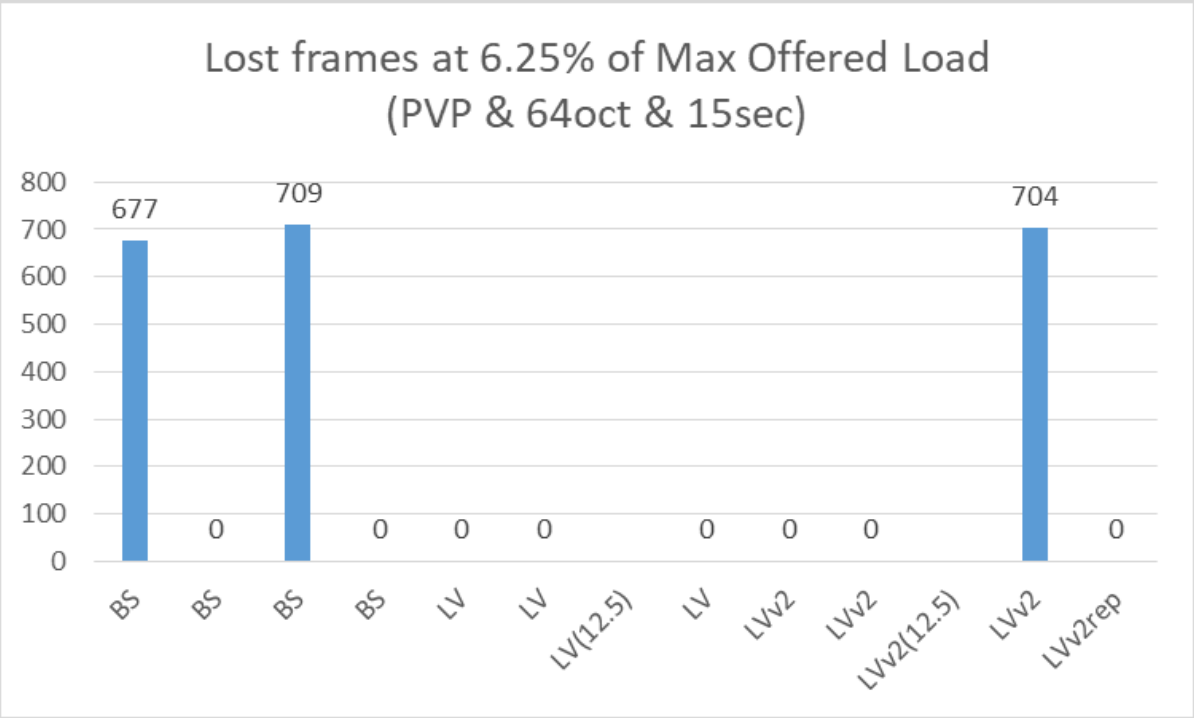

Binary Search=BS, Loss Verification=LV and LVv2, repeated trial = LVv2rep,
All trials at 6.25% of Max Offered Load

Comparison: Trex + OVS_DPDK + Pod12_Node4_DUT 30 Sec Trial Duration; 128 octet Frames

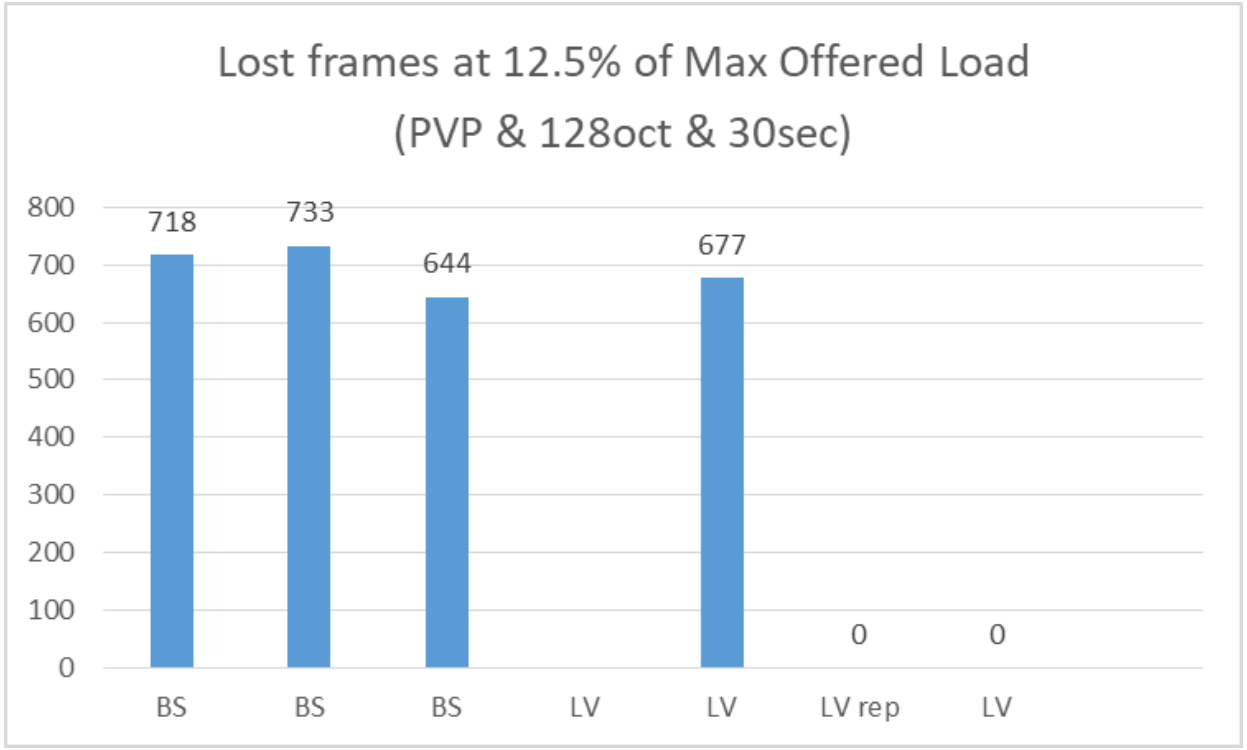

Binary Search=BS, Loss Verification=LV, repeated trial = LVrep,
All trials at 12.5% of Max Offered Load: Is “mode” size-dependent?

Comparison: Trex + OVS_DPDK + Pod12_Node4_DUT
15 Sec Trial Duration; 128 Octet Frames

BS THPT, %max
23.633
15.625
18.164
16.992

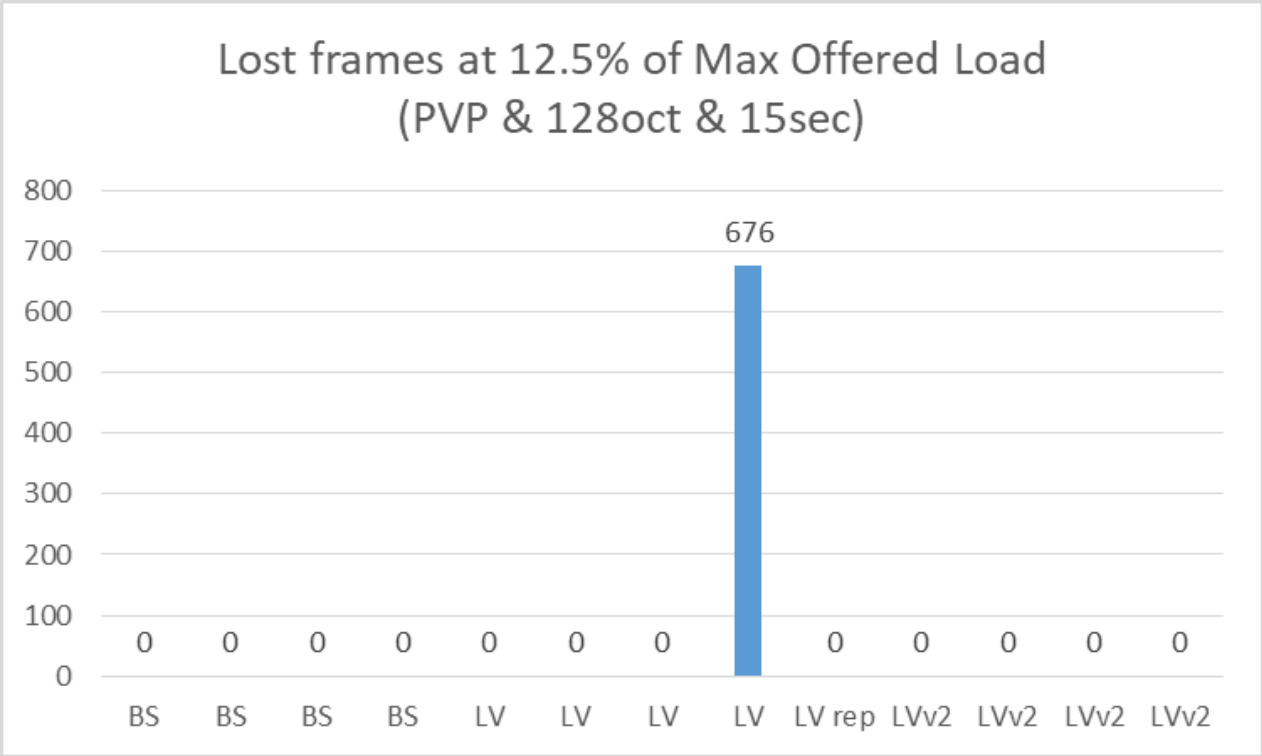

Binary Search=BS, Loss Verification=LV and LVv2, repeated trial = LVv2rep,
All trials at 12.5% of Max Offered Load: Is “mode” size-dependent?

Comparison: Trex + OVS_DPDK + Pod12_Node4_DUT
30 Sec Trial Duration; Throughput in Frames per Second

Blue:
% of Max fps

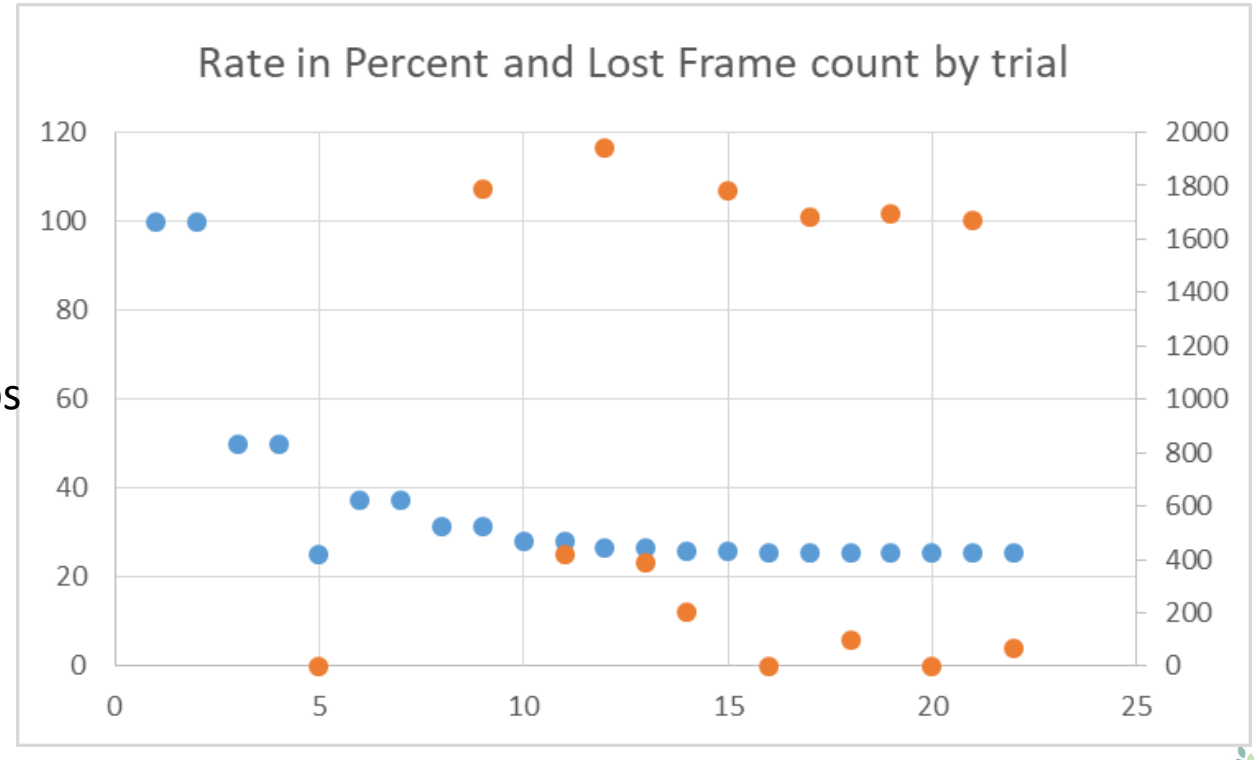

Orange:
Lost Frame Count,
(some >2000)

First test for PVP setup, 128 octet Frames and BS with Loss Verification

Comparison: Trex + OVS_DPDK + Pod12_Node4_DUT
30 Sec Trial Duration; Throughput in Frames per Second

Blue:
% of Max fps

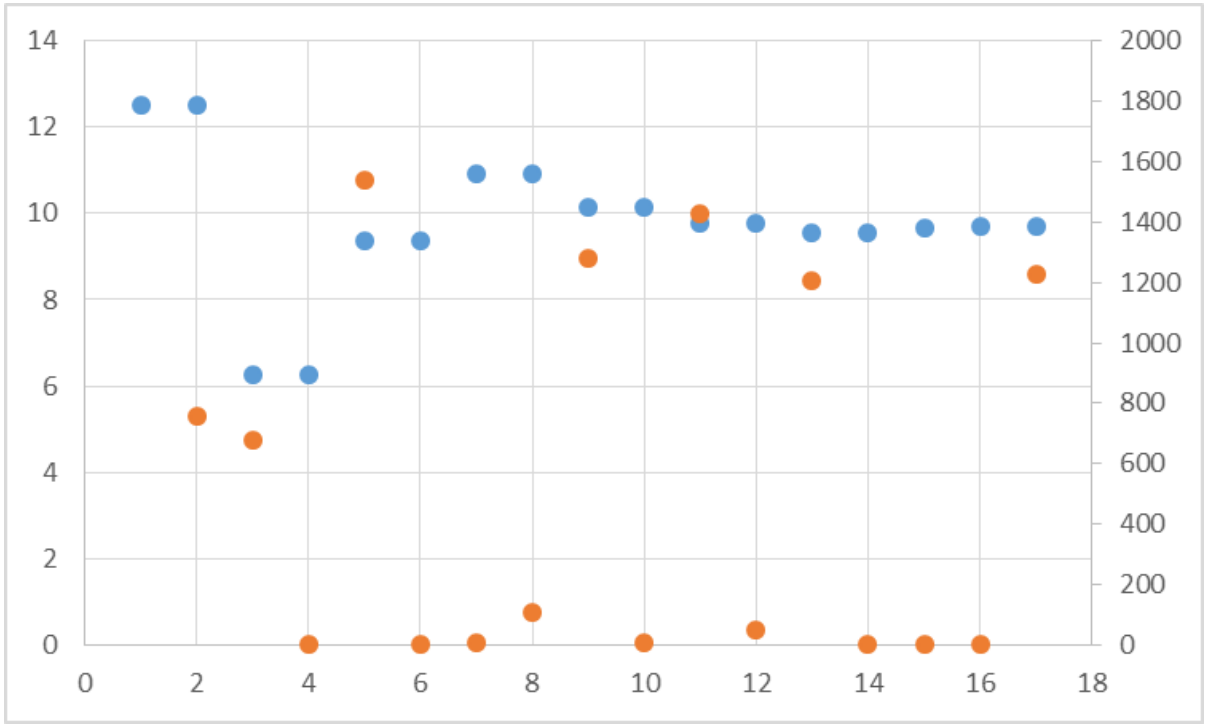

Orange:
Lost Frame Count,
(some >2000)

First test for PVP setup, 64 octet Frames and BS with Loss Verification

Comparison: Trex + OVS_DPDK + Pod12_Node4_DUT
30 Sec Trial Duration; Throughput in Frames per Second

Setup	Search Alg	64 octet	64 octet	64 octet	128 octet	128 octet	128 octet
pvvp	BS rx	1,026,199.3	286,159.8	1,617,446.	1,631,199.	2,635,543.	2,591,939.
pvvp	BS %	2.637	0.732	4.15	8.35	13.525	13.33
pvvp	BS #it	12	12	12	12	12	12
pvvp	BSwLV rx	2,411,807.4	3,206,343.7	2,412,862.	2,425,272	2,901,232.	2,850,596.
pvvp	BSwLV %	6.201	8.203	6.201	12.451	14.893	14.648
pvvp	BSwLV #it	21	22	22	22	24	23
pvvp	BSwLV reversal	4	1	5	6	4	3

"reversal" means that repeated trial result was zero loss

Comparison: Trex + OVS_DPDK + Pod12_Node4_DUT
15 Sec Trial Duration; Throughput in Frames per Second

Setup	Search Alg + val	64 octet	64 octet	64 octet	64 octet	Consistency, STDEV
pvp	BS %	4.541	9.131	4.443	9.473	2.41
pvp	BS rx fps	1,765,869	3,560,964	1,727,143	3,693,305	941,577
pvp	BS #trials	12	12	12	12	
pvp	LV %	11.719	10.596	12.793	10.303	0.99
pvp	LV rx fps	4,559,088	4,119,579	4,989,678	3,970,250	398,741
pvp	LV #trials	21	19	21	20	
pvp	LVv2 %	11.865	10.547	12.5	10.84	0.78
pvp	LVv2 rxfps	4,627,818	4,118,182	4,881,696	4,215,487	309,279
pvp	LVv2#trials	17	18	20	18	

LVv2 indicates a Loss Verification Search with a loss threshold for repeats

Comparison: Trex + OVS_DPDK + Pod12_Node4_DUT
15 Sec Trial Duration; Throughput in Frames per Second

Setup	Search Alg + val	128 octet	128 octet	128 octet	128 octet	Consistency, STDEV
pvp	BS %	23.633	17.529	18.164	16.992	2.66
pvp	BS rx fps	4,607,641	3,417,108	3,532,929	3,303,548	521,534
pvp	BS #trials	12	12	12	12	
pvp	LV %	19.092	19.58	22.559	18.701	1.52
pvp	LV rx fps	3,711,719	3,809,377	4,394,346	3,642,627	297,423
pvp	LV #trials	21	22	20	21	
pvp	LVv2 %	23.389	21.191	21.045	20.312	1.15
pvp	LVv2 rxfps	4,560,212	4,132,494	4,108,020	3,964,987	222,322
pvp	LVv2#trials	18	19	17	19	

LVv2 indicates a Loss Verification Search with a loss threshold for repeats

Comparison: Trex + OVS_DPDK + Pod12_Node4_DUT

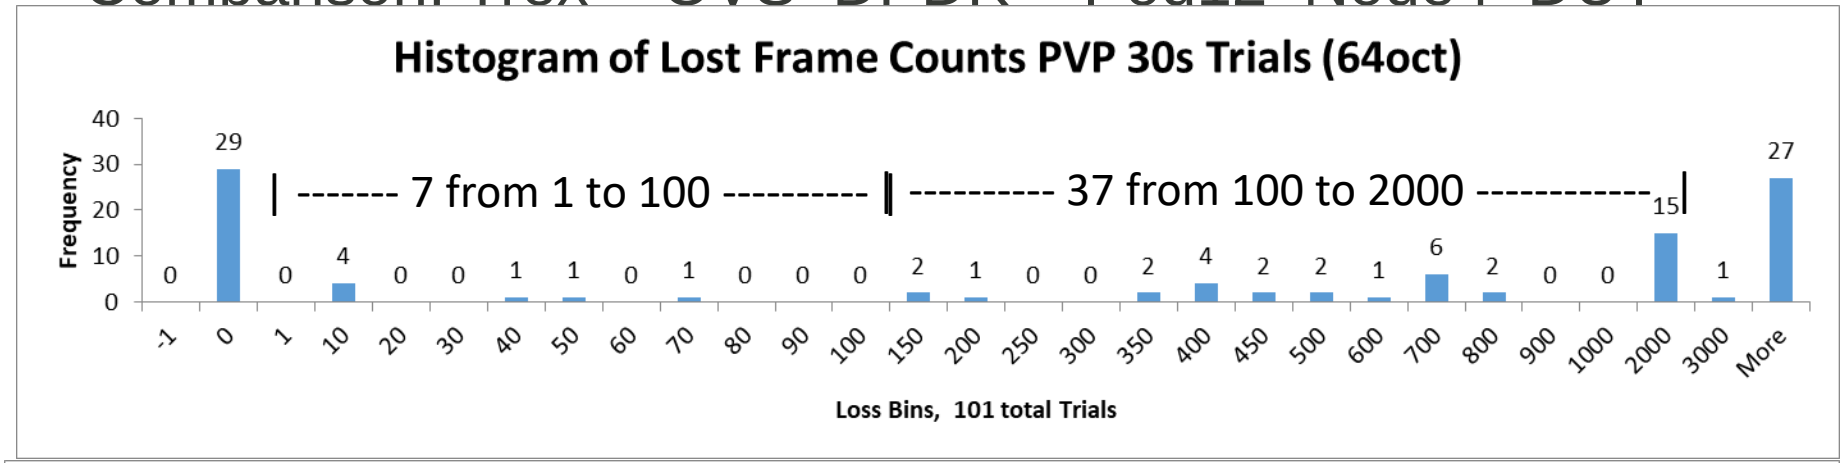

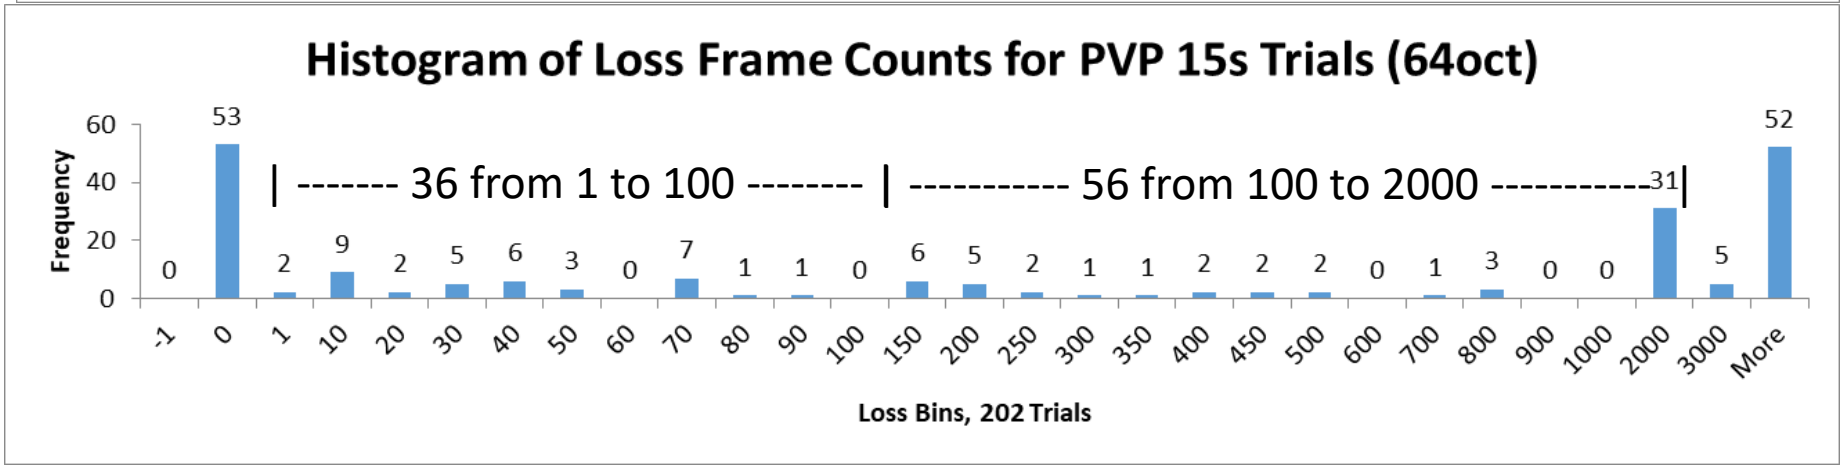

Comparison: Trex + OVS_DPDK + Pod12_Node4_DUT

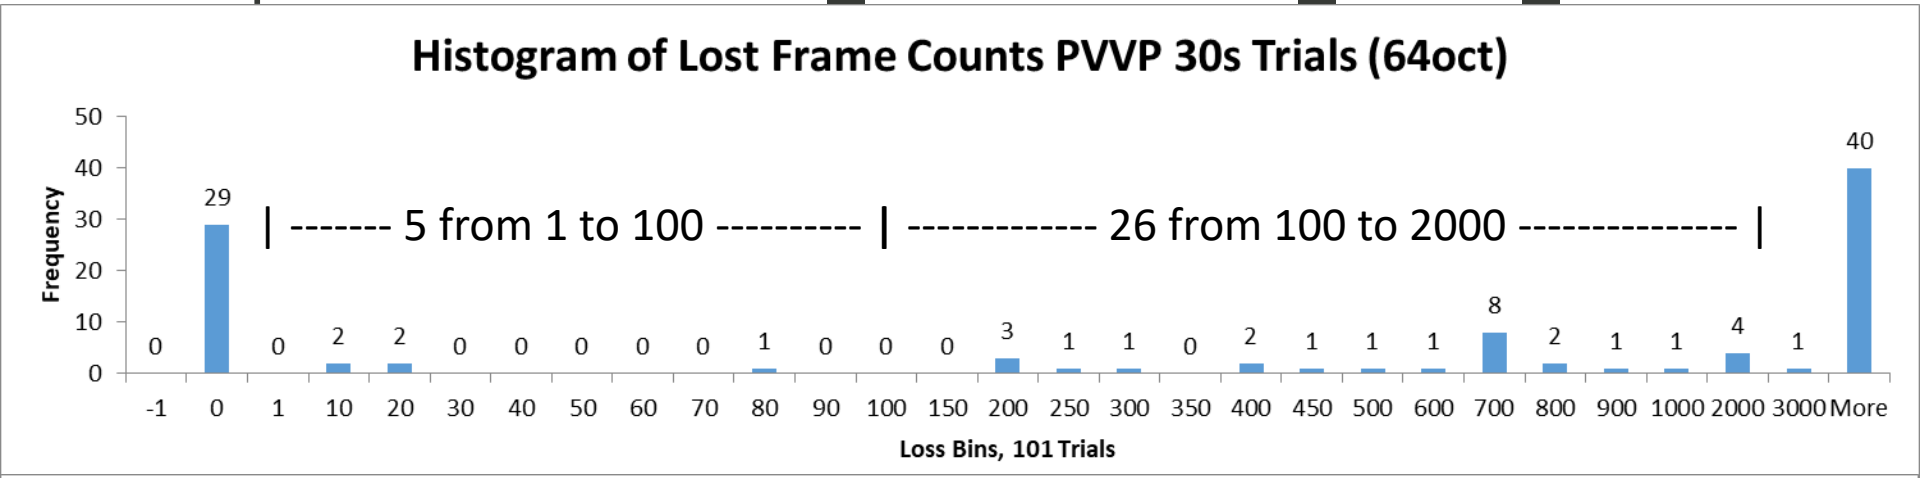

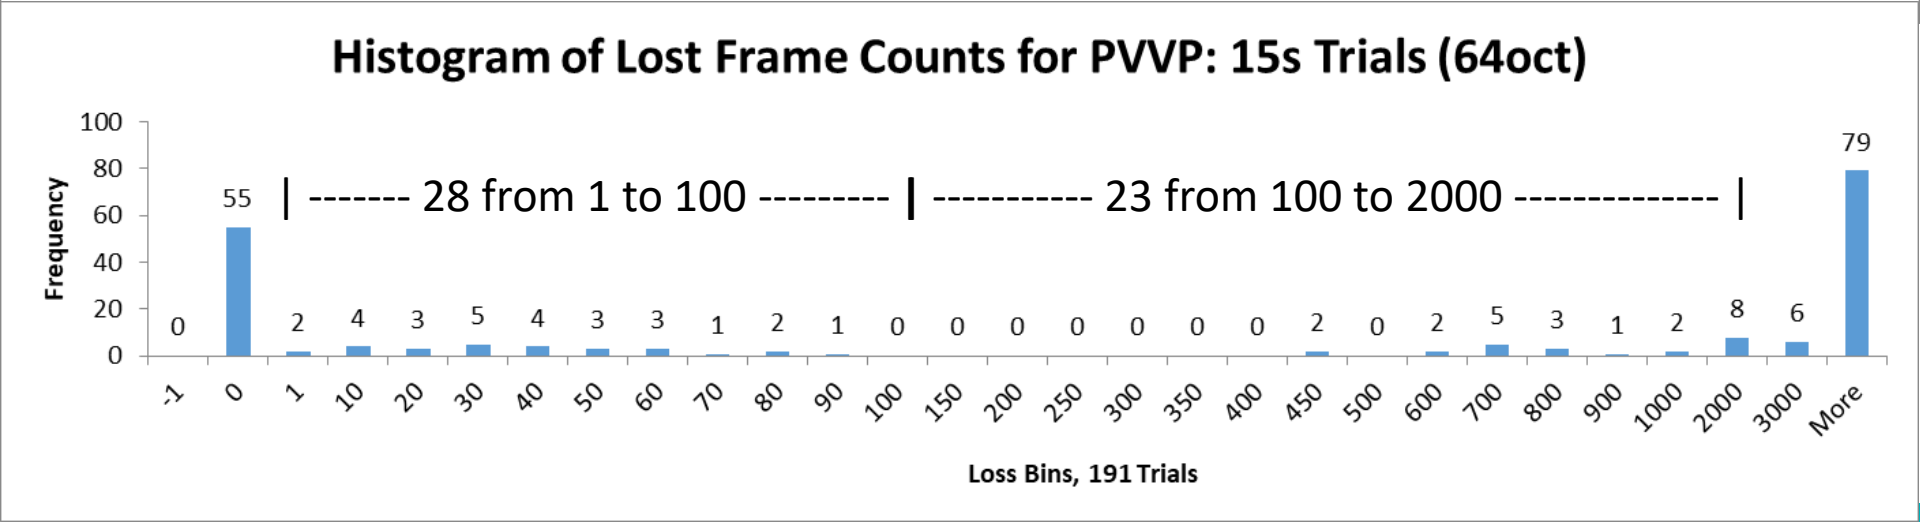

Next steps

- Identify processes that cause transient Loss
- Identify a measurement tool that can
  - characterize bursts of loss and
  - Timestamp loss events (even single loss) for correlation
  - Run in a Long-term mode
- Can we add a heartbeat check to the test VM loopback?
  - Timestamp interruptions of VM operation
- Observe Logical Interface drops and errors during tests.

BACKUP Slides follow, additional References below

NDR/PDR binary search page.

- http://docs.opnfv.org/en/latest/submodules/nfvbench/docs/development/design/ndrpdr.html

The source code of the binary search:

- https://github.com/opnfv/nfvbench/blob/master/nfvbench/traffic_client.py#L674

FD.io Multiple Drop Rate

- https://docs.fd.io/csit/rls1804/report/vpp_performance_tests/mdr_search.html