



STAC Update Fast Data

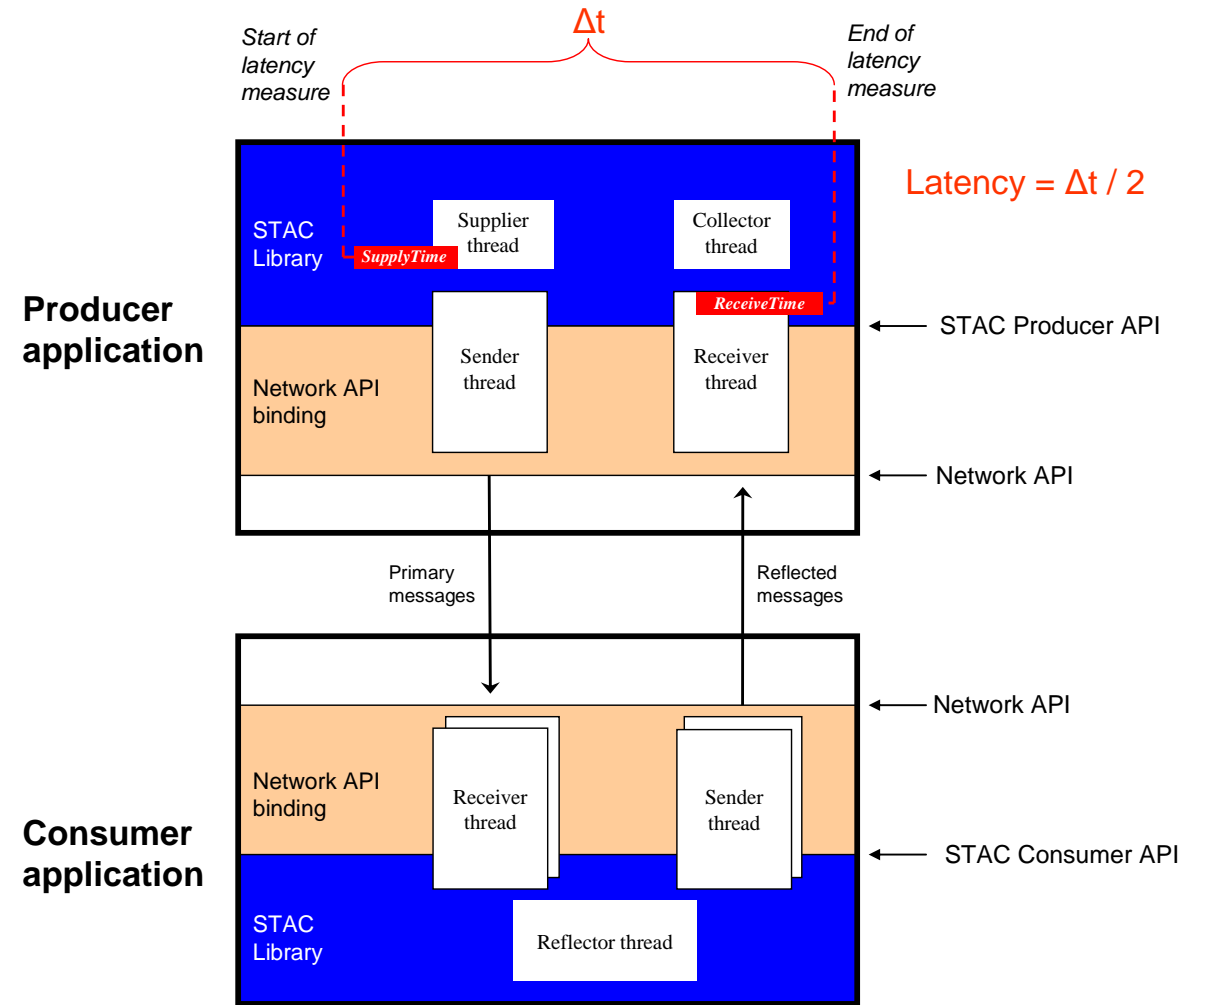
Jack Gidding
CEO, STAC

17 October 2024

jack.gidding@STACresearch.com

STAC-N1

- Measures the performance of a host network stack (server, OS, drivers, host adapter)
- Round-trip software timestamping
- Market data style workload
- Network API to network API
- No middleware, feed handlers, etc.



RedHat CoreOS and OCP with AMD EPYC 9374F and HPE

SUT Details

STAC-N1

- Tested on the same SUT as AMD230414
- STAC-N1 for UDP-TCP binding (udp-tcp-sock-1.0.3)
- 2 x HPE ProLiant DL325 Gen11 Servers, each with
- 1 x 32-core AMD EPYC™ 9374F @ 3.85 GHz (4.3 GHz Boost)
- AMD Xilinx XtremeScale™ X2522-25G-PLUS Adapter
- Red Hat Enterprise Linux CoreOS 9.4
- OpenShift Container Platform 4.16.1
- OpenOnload 8.1.2.26
- KubernetesOnload 3.0
- 25GbE (via cross-over cable, FEC off)



Red Hat
OpenShift



www.STACresearch.com/AMD240924

RedHat CoreOS and OCP with AMD EPYC 9374F and HPE

STAC-N1

Key Results

Compared to all publicly disclosed STAC-N1 results to date with Enterprise Class CPUs that used UDP, this solution exhibited:

- *Tie for the lowest mean and median at the base rate of 100k messages per second (STAC.N1.β1.PINGPONG.LAT1)*
- *Tie for the lowest mean and median at the highest rate tested, 1.4 million messages per second (STAC.N1.β1.PINGPONG.LAT2), across SUTs with the same or lower highest rate tested*

Compared the previous solution with the same server, CPU, memory, and network card, but a previous generation OS (bare metal), and OpenOnload (AMD230414), this solution had:

- *Equivalent performance of the mean, median, and standard deviation at the base rate of 100k messages per second (STAC.N1.β1.PINGPONG.TPUT1)*
- *A 37% reduction in max latency at the base rate of 100k messages per second (STAC.N1.β1.PINGPONG.TPUT1)*



Red Hat OpenShift



www.STACresearch.com/AMD240924

STAC  [®]
STRATEGIC TECHNOLOGY ANALYSIS CENTER

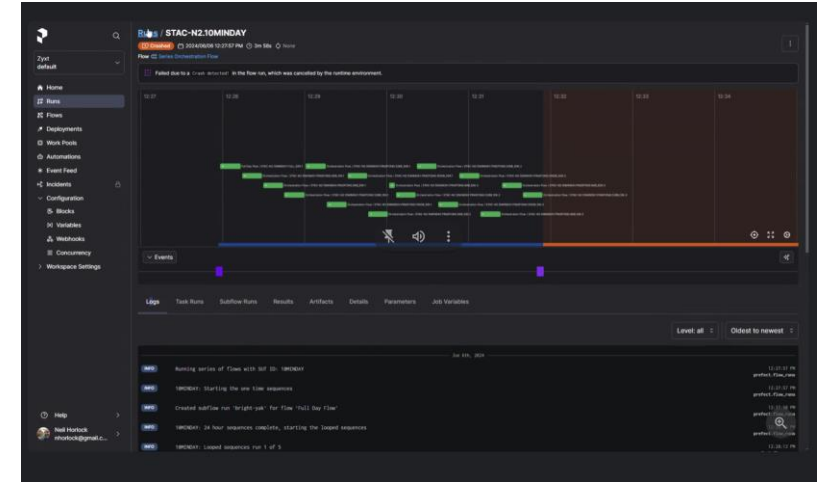
STAC-N1: What's next?

- STAC-N1 created in XXXX
- Has been used on TCP, UDP, and non-standard protocols
 - Overclocked & non-overclocked
 - Liquid cooled
 - 10G & 25G
 - Bare metal & containers
- Liquid Markets Project
 - Strong results published in May
 - New project to refresh those results to target much higher rates
 - Have not been able to hit their target rates on PCIe5/CXL -> library?
- Next up
 - Refresh on STAC-N1 Library
 - Updates to STAC-N1 Harness
 - Reassemble the Working Group to discuss STAC-N1

STAC-N2: Data Center Communications

Measure Cloud to Cloud, and Data Center to Cloud

- Designed to measure the latency and throughput of communication paths between data centers, both on-prem to cloud and cloud-cloud
- Based on STAC-N2 Benchmark Spec Rev D
- Tests orchestrated with [Prefect.io](https://prefect.io)
- Tests involve:
 - 3 packet sizes (66/528/1000B)
 - 1x 24-hour duration test each size @ 1mps
 - 5x each size, 1M messages @ 3 rates (10k/50k/max)
 - 5 times of day for 1M msg tests
 - Can run custom packet sizes
 - Orchestrated parallel execution
- Reliable and/or unreliable protocols
- Audit data forthcoming
- Available NOW



www.STACresearch.com/stac-n2-test-harness

WGs & SIGs

Working groups and special interest groups

Working groups are the bedrock of the STAC community

- Define benchmark specifications
- Ensure business relevancy of workloads
- Bring industry practitioners and technology vendors together to discuss topics beyond benchmarks

STAC Network I/O SIG

- STAC-N1 Updates

STAC Cloud Communications

- Benchmark & Tools
- Cloud Timing



How to get involved

Join the Network working groups



STACresearch.com/nio



STACresearch.com/cloudsig



Get access to this domain

If you'd like to obtain privileged materials from this domain, or if you would like to participate in this group, please click the button below.

[Enable me! >](#)



Job position: *

E-mail address: *

Phone number:

Please include country code.

Website: *

Join this group

Access privileged materials in this domain

[Submit](#)

Data Center SIG

- Challenges in Data Centers
 - Generational upgrades require more power in racks already at capacity
 - More heat being generated, challenging density
 - Top of the bin procs are hotter and require cooling
 - AI Workloads increase the power requirements
 - Need for speed -> need for power
 - Data Center SIG
 - Discuss challenges, solutions, and tradeoffs
 - Engage with providers of solutions
 - Discuss tools
 - Select a workload for performing benchmarks
 - Conduct testing, discuss in the working group
 - Interested? Join the Working Group
 - Click the link on the STAC Exhibitor Page in Agorify

