redhat.

The Path to Cloud-Native Trading Platforms

STAC Spring 2018 Summit NYC

Jeremy Eder, Red Hat Performance Engineering

Containers and OpenShift adopted across industries

OPENSHIFT IS GAINING MOMENTUM

OPENSHIFT CUSTOMER GROWTH IS ACCELERATING

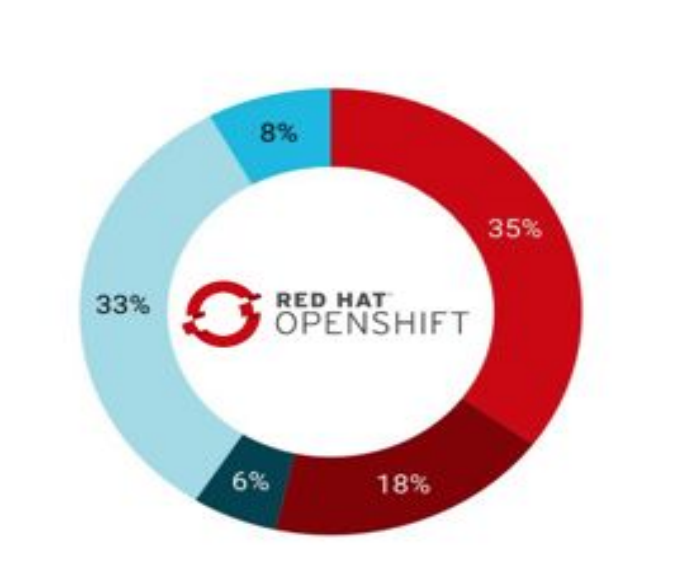

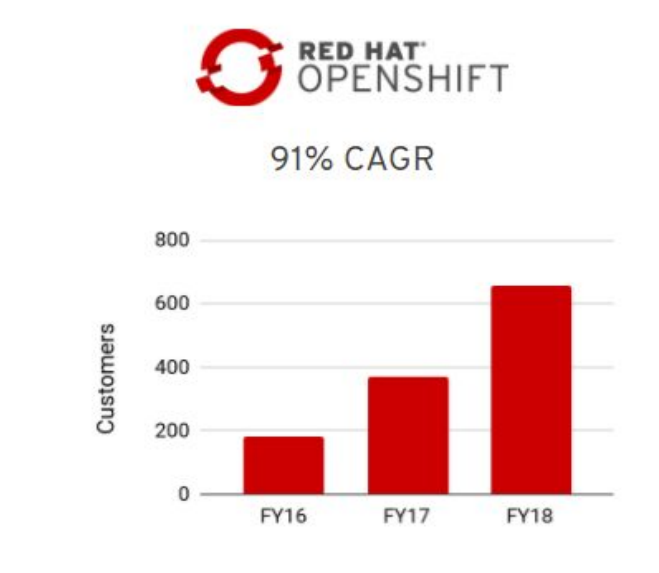

Financial services | Government | Telco | Media/technology | Other

Banking on OpenShift Panel - FINOS

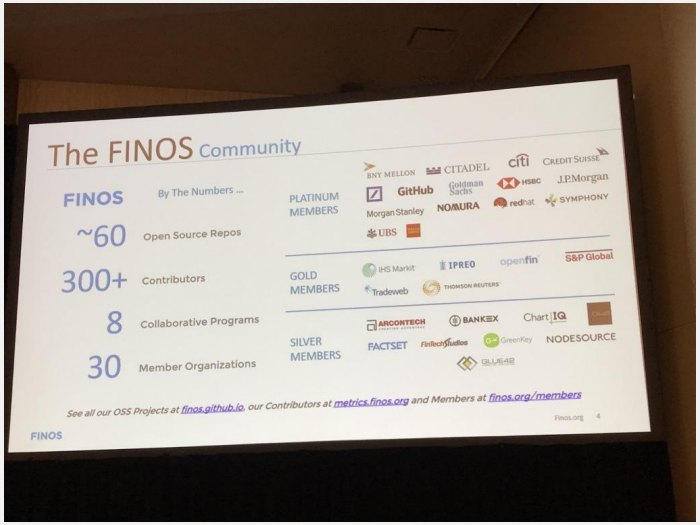

- Highlighting the recently announced FINOS community that is using OpenShift to run the online developer environment
- BNZ - SMB customers were taking 6hrs to set up loan accounts. Developers had bad tooling. Quarterly releases with limited trust in the process. Needed a platform that could manage the velocity of change. OpenShift took them from 1 release/month to 3 releases/week. 10,000 pods in production.
- Santander - 133 projects in production on OpenShift
- Nordea - “largest transformation in banking in Europe at the moment” - moving Core Banking to OpenShift - working closely with Temenos and Accenture

View the Banking on OpenShift Panel OpenShift Commons Gathering video.

OPENSHIFT + CoreOS UNIFIED PLATFORM

Infra Monitoring

Service Cat/Operators
etcd, Prometheus, Red Hat Middleware, ISVs

Cluster Services
Metering/Chargeback, Metrics, Logging, Registry

Developer Productivity
Build Automation, Image Streams, CI/CD

Operator Lifecycle Manager

Automated Operations

Hybrid Cloud | Install / Upgrade | Network / CNI

Ops & Dev Consoles | Security / Auth | Storage / CSI

Kubernetes

Red Hat Enterprise Linux or Red Hat CoreOS

App Monitoring

Best Operator Experience — CaaS ⟷ PaaS — **Best Developer Experience**

Performance-Sensitive Applications

Going beyond generic web hosting workloads

- Identify requirement overlap across verticals
- Plumb enhancements generically
- Allow flexibility

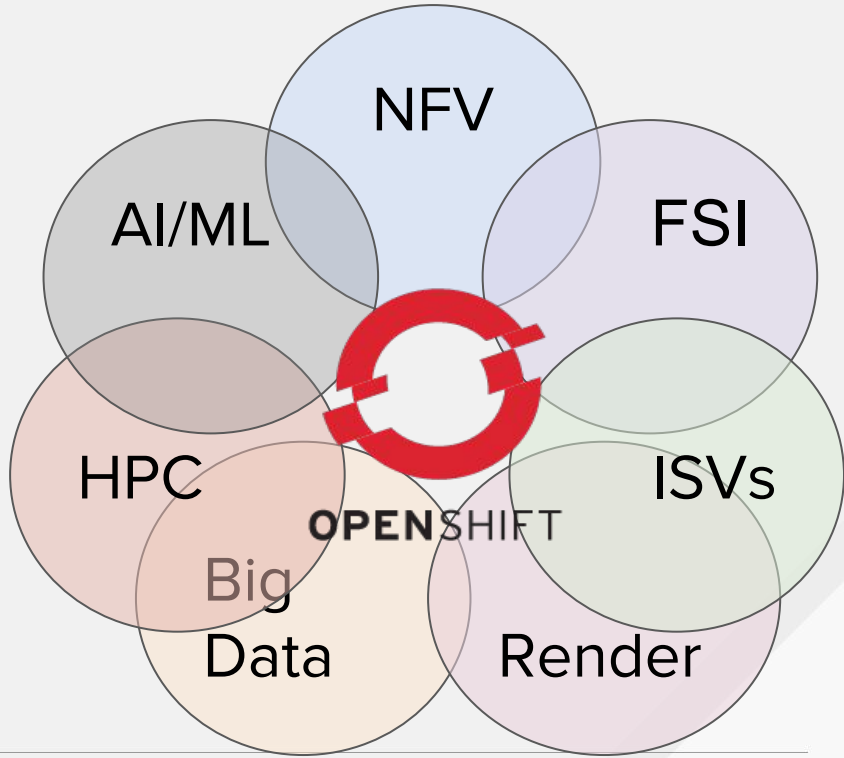

Upstream First: Kubernetes Working Groups

- Resource Management Working Group
  - Features Delivered (all are GA in OCP 3.10)
    - Device Plugins (GPU/Bypass/FPGA)
    - CPU Manager (exclusive cores)
    - Huge Pages Support
  - Extensive Roadmap
- Intel, IBM, Google, NVIDIA, Red Hat, many more…

Upstream First: Kubernetes Working Groups

- Network Plumbing Working Group
  - Formalized Dec 2017
- Implement an out of tree, pseudo-standard collection of CRDs for multiple networks
- Separate control- and data-plane, Fast Data-plane
- IBM, Intel, Red Hat, Huawei, Cisco, Tigera...at least.

Progress Report

What has been done in the last year?

- CPU manager (static pinning) - GA in 3.10
- HugePages - GA in 3.10
- Device Plugins (GPU, etc.) - GA in 3.10
- Sysctl support - TP in 3.10
- Extended Resources - GA in 3.9

Roadmap

Red Hat continues to invest in evolving support

Topic areas

- NUMA
- Co-located device scheduling
- External device monitoring
- Resource API V2

STAC-N1 Project Overview

STAC-N1: Bare Metal

- Solarflare XtremeScale X2522 Adapters
- Supermicro SYS-1029UX-LL1-S16 Servers
- Red Hat Enterprise Linux 7.5

STAC-N1: Containerized/Kubernetes

- Solarflare XtremeScale X2522 Adapters
- Supermicro SYS-1029UX-LL1-S16 Servers
- Red Hat Enterprise Linux 7.5
- Red Hat OpenShift 3.10 (pre-release)

Cluster Topology

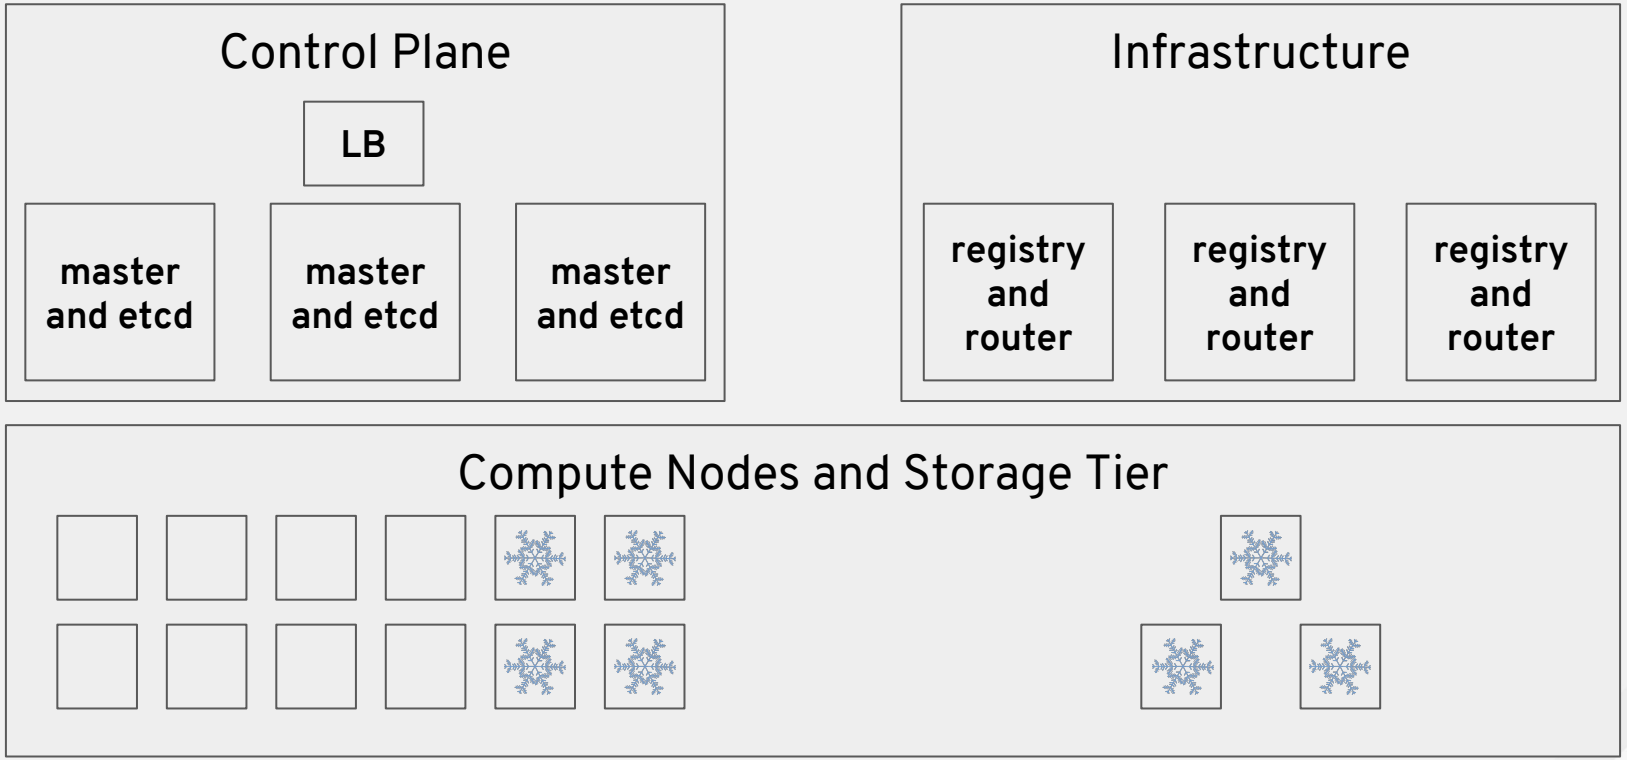

OpenShift Deployment for STAC-N1

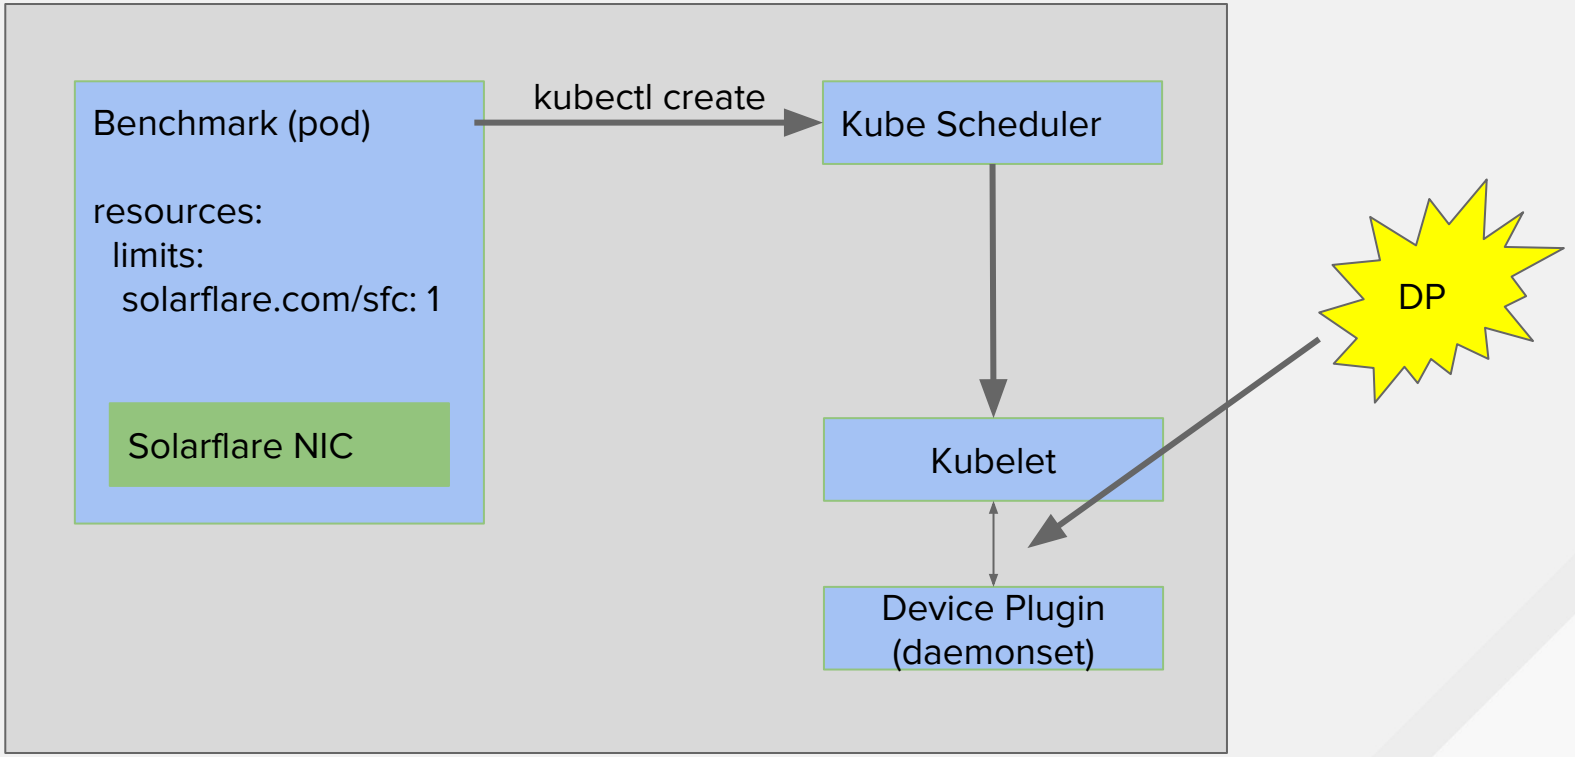

Containerized STAC-N1 Benchmark

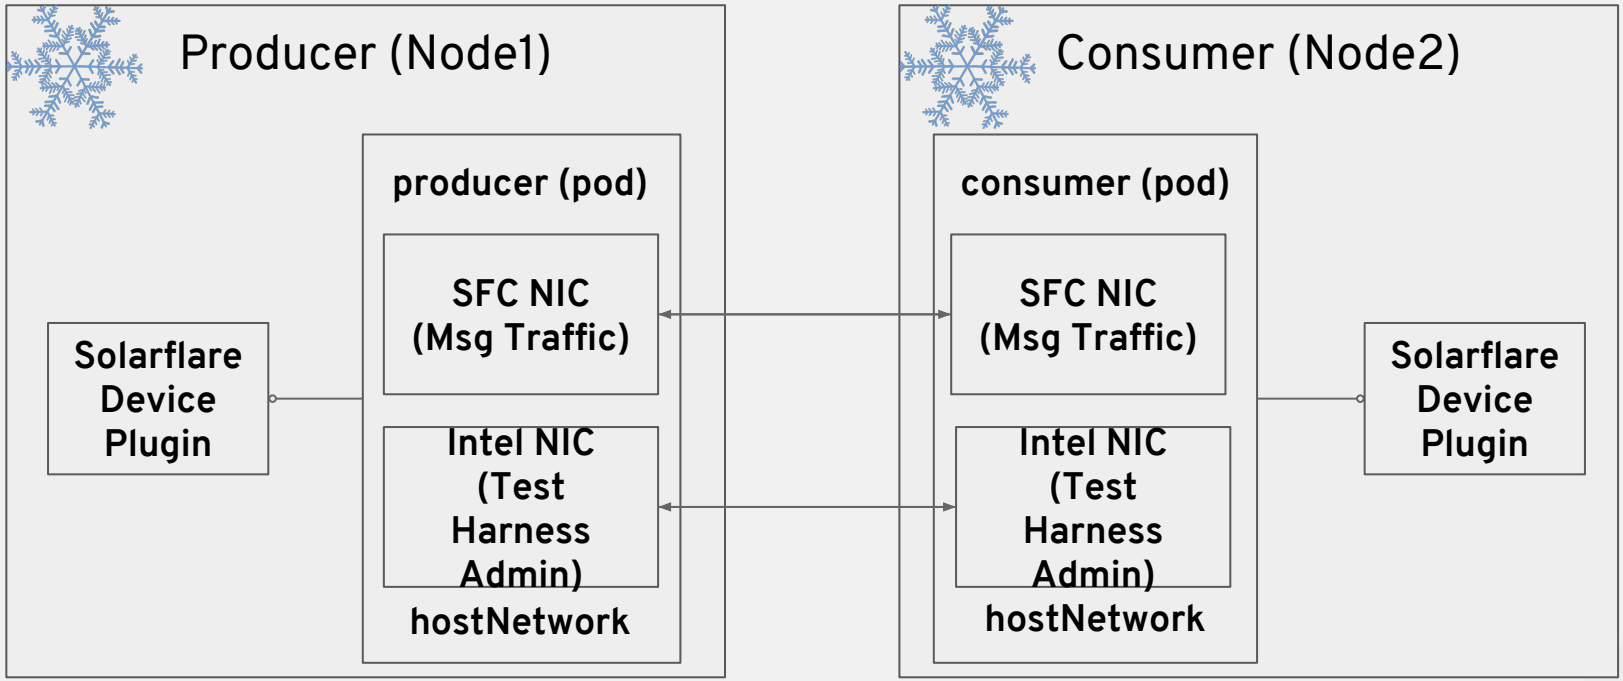

STAC-N1 Bare Metal Benchmark SUT: SFC180604a

- Lowest mean latency (2.3us) for 264-byte messages at both the base rate of 100K messages per second and the highest rate tested of 1 million messages per second.
- Max latency at 100K messages per second was the lowest of any system using sockets (10us)

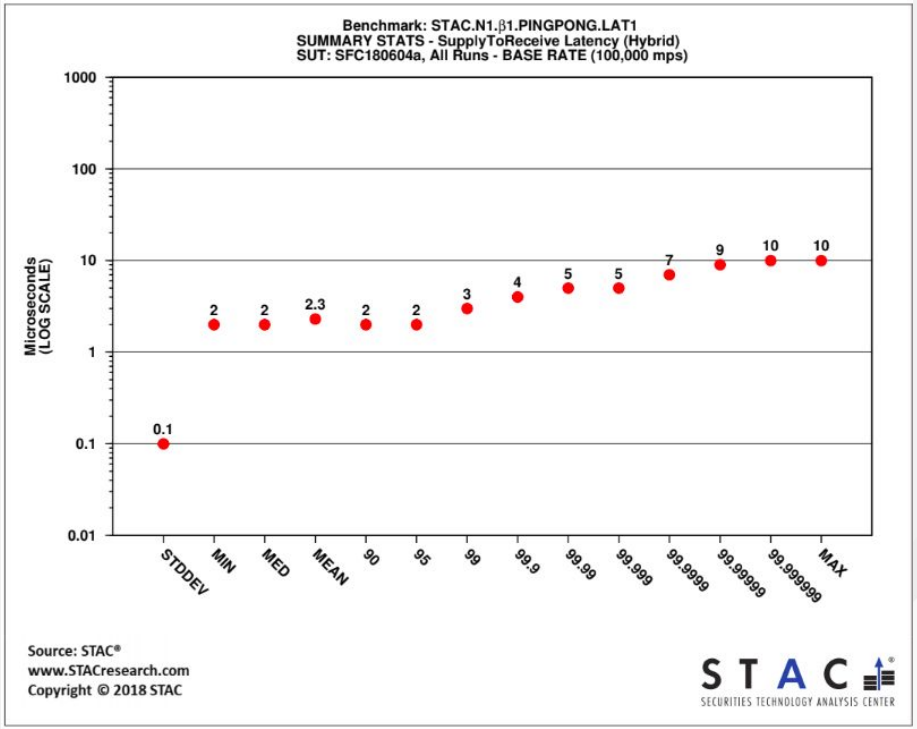

STAC-N1 OpenShift Benchmark SUT: SFC180604b

- Mean and 99th percentile latency were the ***same as bare metal @ 100k and 1M/s rate***

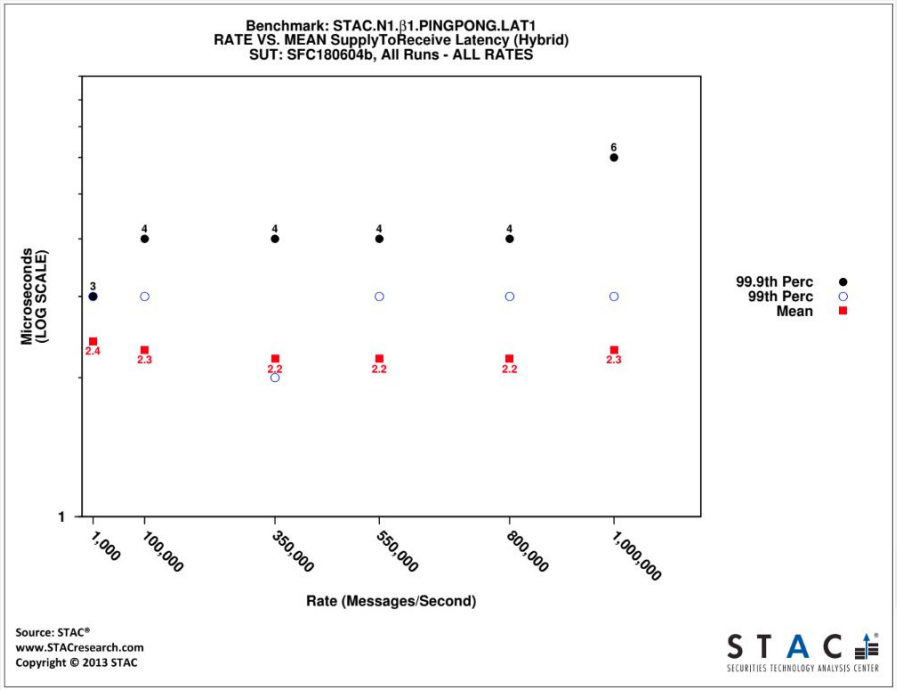

redhat.

Interesting in learning more?

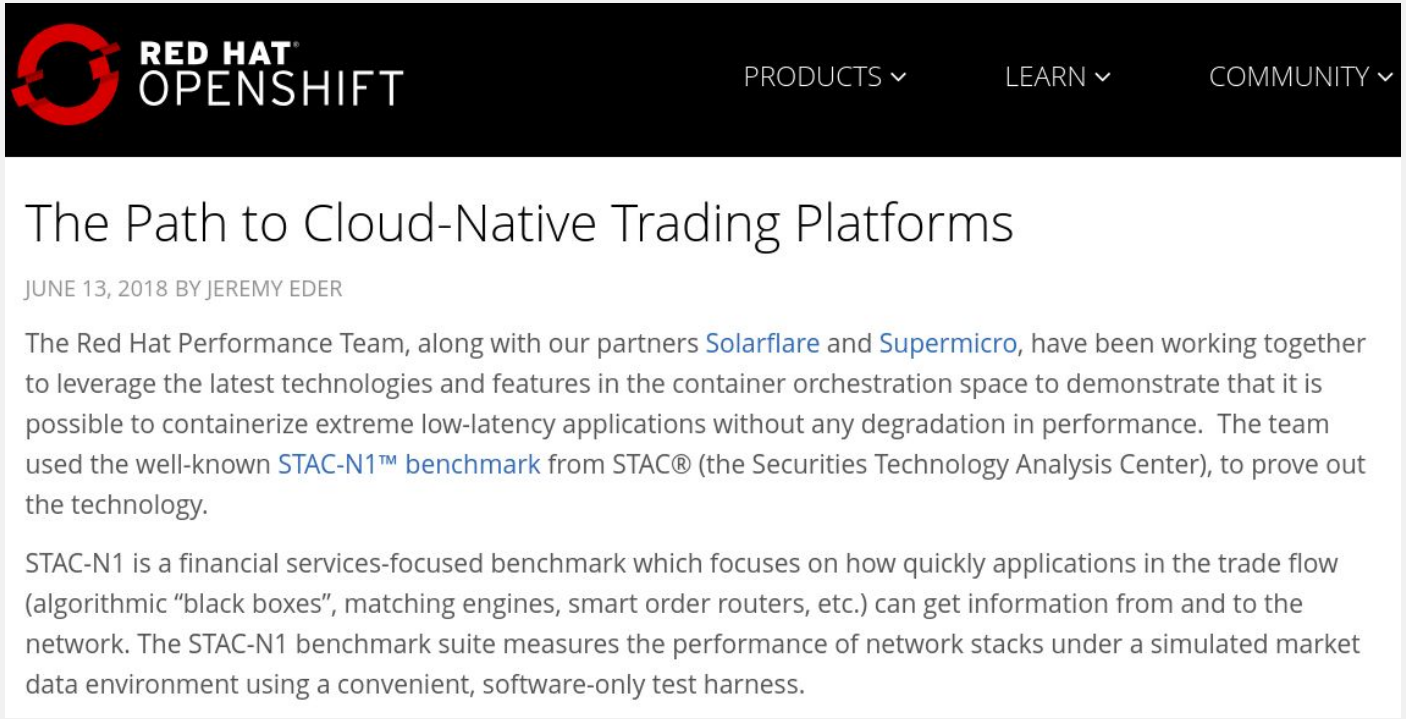

https://blog.openshift.com/the-path-to-cloud-native-trading-platforms

Kubernetes Deployment for STAC-A2

- CUDA 9
- 8 x NVIDIA Tesla V100 (Volta) GPUs
- HPE Apollo 6500 w/XL270d Gen9
- Red Hat Enterprise Linux 7.4
- Kubernetes 1.8 (setup info)
- nvidia-smi --applications-clocks=877,1380

- All-in-One Kubernetes Installation
- (hack/local-up-cluster.sh)
- Node labeled
- Containers:
  - RHEL7+CUDA9
  - RHEL7+CUDA9+-DEVICE-PLUGIN
  - RHEL7+CUDA9+STAC-A2

- https://rhelblog.redhat.com/2017/11/21/red-hat-and-partners-deliver-new-performance-records-on-prominent-risk-analytics-benchmark/

- https://news.developer.nvidia.com/a-new-stac-a2-record/