RED HAT SUMMIT

A Day in the Life of an Open Source Performance Engineering Team

More than just A/B testing

Aakarsh Gopi, Anaga Vijaykumar, and Peter Portante
Engineers from Red Hat's Performance and Scale Team

Introduction & Background

- A presentation of the OpenSource.com article, “A Day in the Life of a Performance Engineering Team”
    - https://opensource.com/article/19/5/life-performance-engineer
- Remember “Space Balls”?
    - “The Open Source Community Republic” - the good guys
        - “Jedi Performance Council” works to maintain the Republic
        - “Resistance Collaborators” are those in the community that collaborate together
    - “The Complex Federation” - the bad guys
        - The Sith Lord Darth Complexity, and his apprentice, Darth Slowness, work to destroy the Republic

THE COMPLEXITY MENACE

Performance Wars - Episode XLII

[illegible]

[illegible] OpenStack [illegible]
Gather data and distill it to
understand component
interactions, and take action.

Master Obi Wan, and her
apprentice, Luke Skywalker, have
been charged with the mission to
save the Alderaan OpenShift
cluster using tools to gather the
data to combat complexity and
slowness.

What does a Performance Engineering Team do?

A/B testing is less than 10% of what we do.

Our “Performance and Scale Team”, as we are named at Red Hat, has the following mission statement:

- Establish performance and scale leadership of the Red Hat portfolio
  - The scope includes component level, system, and solution analysis
- Collaborate with engineering, product management, product marketing, and CEE, as well as hardware and software partners
- Deliver public-facing guidance, internal enablement, and continuous integration tests

How does our team fulfill that mission?

Our team’s major responsibilities

- Working with product teams to set performance goals
  - Develop performance tests, deploy, test, collect data and measure against goals
  - Work to re-run performance tests to catch regressions
- Developing open-source tooling to help us gather the data we need
  - Making them available to the communities from which the products are derived
- We work to be transparent and open about how we do performance engineering
  - We share methods and approaches to benefit the communities, allowing them to re-use our work, and that benefits us by letting us leverage the work they contribute in turn with these tools

Key attributes of the team

Attributes responsible for our team’s success

- Collaboration
  - We have to work together in our team, with our teams at Red Hat, and with partners and customers
  - This is the main foundation
- Knowledge
  - Our team has people with both broad, deep, and diverse knowledge, spanning Red Hat’s production portfolio, and through all layers of the software stack
- Tooling
  - Leverage open-source tools to help us gather the data we need to understand these complex product solutions

The Complexity Battle

Where does all this complexity come from?

- Product solutions are complex
  - Hardware components are growing in complexity
  - Software solutions on that hardware involve multiple layers
- For example, an OpenShift Cluster
  - Kubernetes ...
  - On OpenStack
  - RHEL used in both over- and under-cloud nodes
  - Using software-defined netorking (ovs) and storage (ceph)
  - All on multi-socket, multi-nic, hardware
- Performance problems can be at any level, or in any component, or a combination

We need data to help us attack complexity

Data is the answer, but that leads to more to more questions?

- The only way to battle the rising complexity is to gather data and distill it
- We need tools to help us do that
    - Gather data in a comprehensive fashion, efficiently
    - Make it easy to gather that data in a repeatable manner
    - Capturing
        - Configuration data
        - Workload data
        - Telemetry at as many levels of software and hardware as possible
        - Adding log and / or trace data to flesh out behaviors
- This is a tall order that no one team can achieve on their own …

Importance of open-source tools

Open-Source Tooling is no longer a luxury, but a need for today's performance engineering teams.

- A performance engineering team can't spend all of its time only on developing tools, since that will prevent the team from effectively collecting data. By developing its tools in a collaborative manner, a team can leverage work from the community to make further progress, while still being able to generate results on which they will be measured.
- Tooling is the backbone of our performance engineering team, and we've strived to use the tools already available upstream.
  - Where there were no tools available in the community that fit our needs, we've taken the initiative to build tools that help us achieve our goals and made them available to the community.
  - Open-sourcing our tools has helped us immensely because we've received contributions from our competitors and partners as well, as we're able to collectively solve problems through collaboration.

Alderaan and ClusterLoader

- Alderaan:
  - Alderaan Scale-CI pipeline automates installation of OCP on a OpenStack (~50 overcloud nodes) cluster to run performance and scale tests
  - Alderaan allows other teams to run their tests on OCP cluster at scale
  - Alderaan supports collection of Pbench data
  - Repositories:
    - https://github.com/openshift/svt
    - https://github.com/openshift/aos-cd-jobs
    - https://github.com/distributed-system-analysis/pbench

- ClusterLoader:
  - ClusterLoader can stress the cluster by loading it with namespaces, pods and such to simulate user env
  - Repositories:
    - https://github.com/openshift/origin/tree/master/test/extended/cluster#cluster-loader
    - https://github.com/kubernetes/perf-tests/tree/master/clusterloader

OpenStack and Browbeat

- OpenStack
  - OpenStack is a combination of open source tools (known as projects) that use pooled virtual resources to build and manage private and public clouds. Six of these projects handle the core cloud-computing services of compute, networking, storage, identity, and image services, while more than a dozen optional projects can be bundled together to create unique, deployable clouds.

- Browbeat
  - Browbeat is a performance analysis tool that facilitates running benchmarks against OpenStack environments
  - Browbeat’s power lies in its ability to collect comprehensive data ranging from logs to system metrics about an openstack cluster while it orchestrates workloads
  - Repositories:
    - https://github.com/openstack/browbeat

Importance of Collaboration

Intra-team collaboration is as important as inter-team collaboration for a performance engineering team

- Most performance engineers tend to create a niche for themselves in one or more sub-disciplines of performance engineering viz system tool development, performance analysis, systems knowledge, systems configuration and such.

- Our team is composed of engineers with knowledge of how to setup/configure systems across the product stack, and those who possess knowledge of how a configuration option would affect the system’s performance, and so on.
    - Thus our team’s success is heavily reliant on effective collaboration between performance engineers in the team.

Ceph, CBTool, and Pbench

- Ceph:
  - Ceph is an open, massively scalable software defined storage solution for modern workloads like cloud infrastructure, data analytics, media repositories, and backup and restore systems.
- CBTool:
  - CBTool is a framework that automates cloud benchmarking, used to drive fio workloads against Ceph
  - Repositories: https://github.com/ibmcb/cbtool
- Pbench:
  - Pbench is a swiss army knife as it allows the user to run benchmarks such as fio, uperf or custom user defined tests
  - Pbench gathers metrics through running tools such as sar, iostat, pidstat.
  - Pbench provides a dashboard UI to help review and analyze the data collected.
  - Repositories: https://github.com/distributed-system-analysis/pbench

Perf c2c

- Perf:
  - Perf is a performance analyzing tool in Linux.
  - Perf is capable of statistical profiling of the entire system (both kernel and userland code).

- Perf c2c:
  - C2c is a perf tool
  - At a high level, “perf c2c” will show you:
    - The cache lines where false sharing was detected.
    - The readers and writers to those cache lines, and the offsets where those accesses occurred.
    - The pid, tid, instruction addr, function name, binary object name for those readers and writers.
    - The source file and line number for each reader and writer.
    - The average load latency for the loads to those cache lines.
    - Which numa nodes the samples a cache line came from and which cpus were involved.
  - Blogposts:
    - https://joemario.github.io/blog/2016/09/01/c2c-blog/

Scale and Alias Labs & QUADS

Infrastructure Automation Tooling

- The “Scale” of the Performance and Scale team comes from the work we do to operate our products “at scale”
    - To accomplish that, we operate two labs automated to provide large scale deployments for testing - Scale and ALIAS labs

- QUADS
    - A framework for providing automated future scheduling, documentation, end-to-end provisioning and assignment of servers and networks
    - Used primarily in the Red Hat Performance Engineering R&D labs to manage large sets of constantly changing bare-metal systems and networks

- Badfish
    - A vendor-agnostic, redfish-based API tool used to consolidate management of IPMI and out-of-band interfaces for common server hardware vendors.

- You can learn more about these tools at https://quads.dev

Glance at all tools leveraged by our team

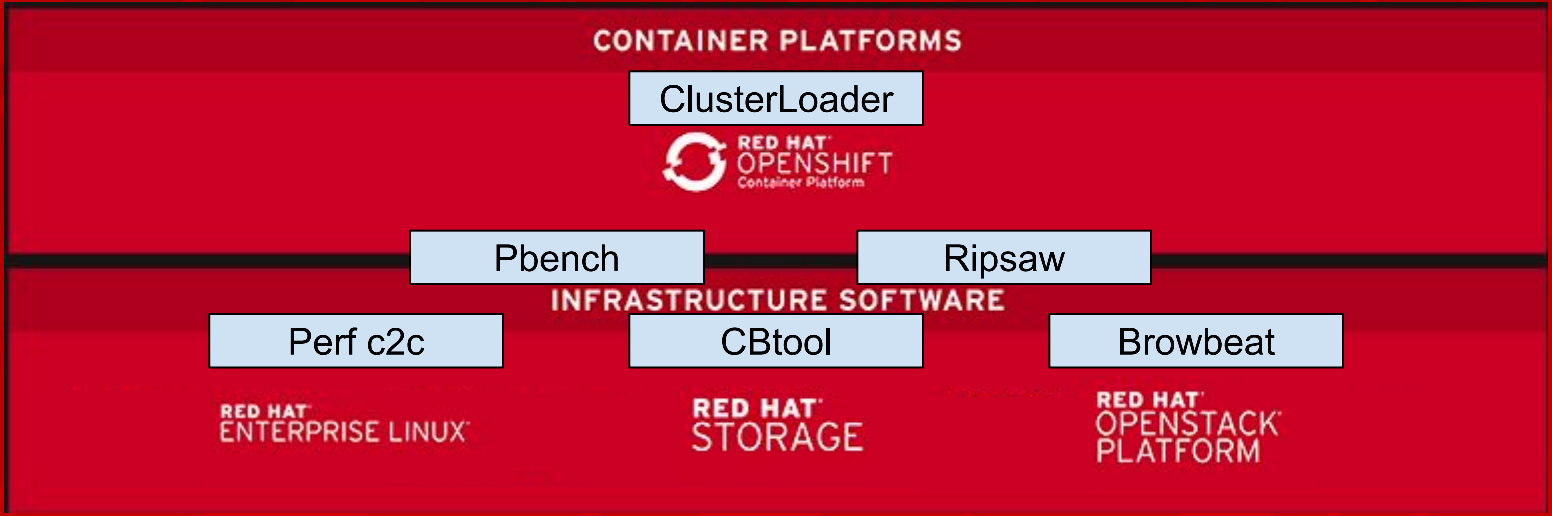

ClusterLoader

Pbench

Ripsaw

Perf c2c

CBtool

Browbeat

Bottom line

Performance & Scale Team operational success

- Collaboration is our key to accomplishing our performance analysis goals
- Using that collaborative spirit to utilize the diverse knowledge of the team
- Leverage open-source tools that help us gather and distill data about complex systems

Credits

- Cast
  - “Obi Wan” - Anaga
  - “Luke” - Peter
- Writers
  - Anaga, Aakarsh, Peter
- Directors
  - Anaga, Aakarsh, Peter
- Special Effects
  - Anaga, Aakarsh, Peter

RED HAT SUMMIT

THANK YOU

in linkedin.com/company/Red-Hat

YouTube youtube.com/user/RedHatVideos

f facebook.com/RedHatinc

twitter.com/RedHat