Tips and Tricks for Managing and Administering Ceph Clusters

Michael Hackett

January 23rd, 2021

Quick background on me

Introduction:

- Working in data storage for over 15 years
- Specifically on Ceph for 5+ years
- Presented at Cephalocon and Red Hat Summit for multiple years.
- Have co-authored two books on deploying and managing Ceph clusters
- Current position is Support Product Lead for Ceph in the Product Experience Team at Red Hat

Red Hat

What are we going to be talking about

Prerequisites:

- A want to know more about Ceph!
- No deep Ceph knowledge required

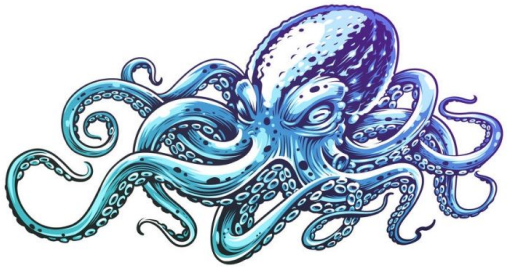

Agenda:

- Will cover the Octopus release of Ceph
  - Current stable release is Octopus (15.2.7)
- An overview Cephadm, and usage basics
- A glance of our new embedded Ceph Dashboard

Upstream Code:

- https://github.com/ceph/ceph

Note:

Will not be discussing Rook Ceph, or the Ceph Operator for Kubernetes

What is Cephadm?

- Introduced in Ceph Octopus (v15.2.0).
- New orchestration API intended to become preferred installation and management tool for Ceph.
- Intended to replace Ansible, Puppet and other tools.
- Remove the dependency of understanding external tools in order to deploy Ceph.
- Centralizes configuration information enabling advanced management functionality.
- Intended to be future of installation and management for “bare-metal” Ceph deployments. We will use Rook Ceph for Kubernetes environments.

What's the goal here?

- Focus on solving the core problem of making Ceph easier to install.
- Deployment of all components inside containers
- Inside containers limits OS dependencies
- Integrate with the orchestrator API
- No further dependency on external management tools
- Ability to isolate Ceph clusters on same hardware from eachother
- Automate the upgrade process
- Allow migration from legacy deployment tools

It’s really that simple!

All starts from a single host!

Prerequisite: Ensure the following dependencies are met: python3, lvm2, podman, systemd, ntp.

```
Step 1 - Fetch the most recent version of the script
# curl --silent --remote-name --location
https://github.com/ceph/ceph/raw/octopus/src/cephadm/cephadm
Step 2(optional) -  Install cephadm binary on the host
# ./cephadm install
Step 3 - Bootstrap your first node
# cephadm bootstrap --mon-ip <monitor_ip> --allow-fqdn-hostname
Step 4(prerequisite) - Copy the clusters public key over to additional
nodes
# ssh-copy-id -f -i /etc/ceph/ceph.pub root@ceph7
Step 5 - Add the new host(s) to ceph orchestrator
# ceph orch host add ceph6.cee.lab.eng.rdu2.redhat.com
Step 6 - Allow the orchestrator to discover and deploy your devices
# ceph orch apply osd --all-available-devices
```

How Cephadm works?

Deploys and manages Ceph cluster using SSH connections to interact with hosts. Services are deployed using container images from Docker Hub

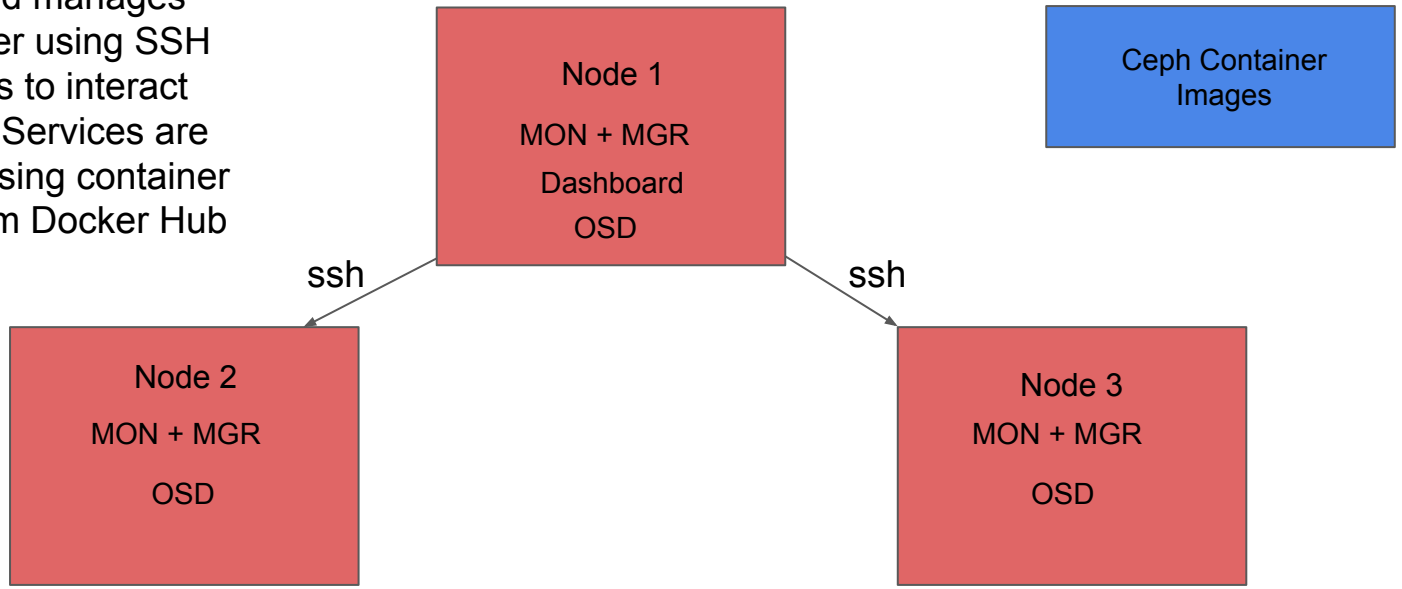

Want More on Cephadm!

- https://docs.ceph.com/en/latest/cephadm/
- https://docs.ceph.com/en/latest/cephadm/install/
- https://www.redhat.com/en/blog/red-hat-ceph-storage-5-introducing-cephadm
- https://ceph.io/ceph-management/introducing-cephadm/

What is the Ceph Dashboard?

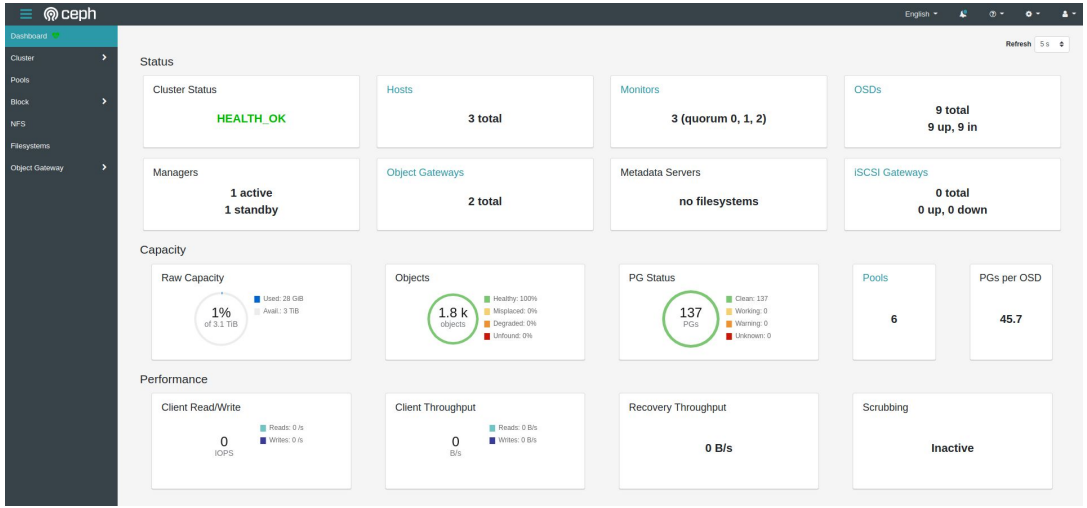

- A built-in web based management and monitoring application that is part of the Open Source Ceph distribution.
- Originally shipped with Ceph Luminous and has vastly improved since.
- Goal was to make day to day management (and life!) easier for Ceph Administrators

How Does it Work?

- Requires a running Ceph Octopus cluster.
- Supported and tested on Chrome and Firefox browsers
- Enabled via a ceph-mgr module (Dashboard) and is derived from the openATTIC Ceph Management and Monitoring Tool.
- Utilizes other ceph-mgr modules for collection and reporting.
- Provides a graphical representation of statistics and cluster information through a web server hosted by the ceph-mgr
- Manages and reports on Ceph clients in the cluster
- Configurable alerting system
- If deploying the cluster via Cephadm, by default the Dashboard is configured and enabled.

Check Out:
https://docs.ceph.com/en/latest/mgr/dashboard/

Hosts Overall Performance View

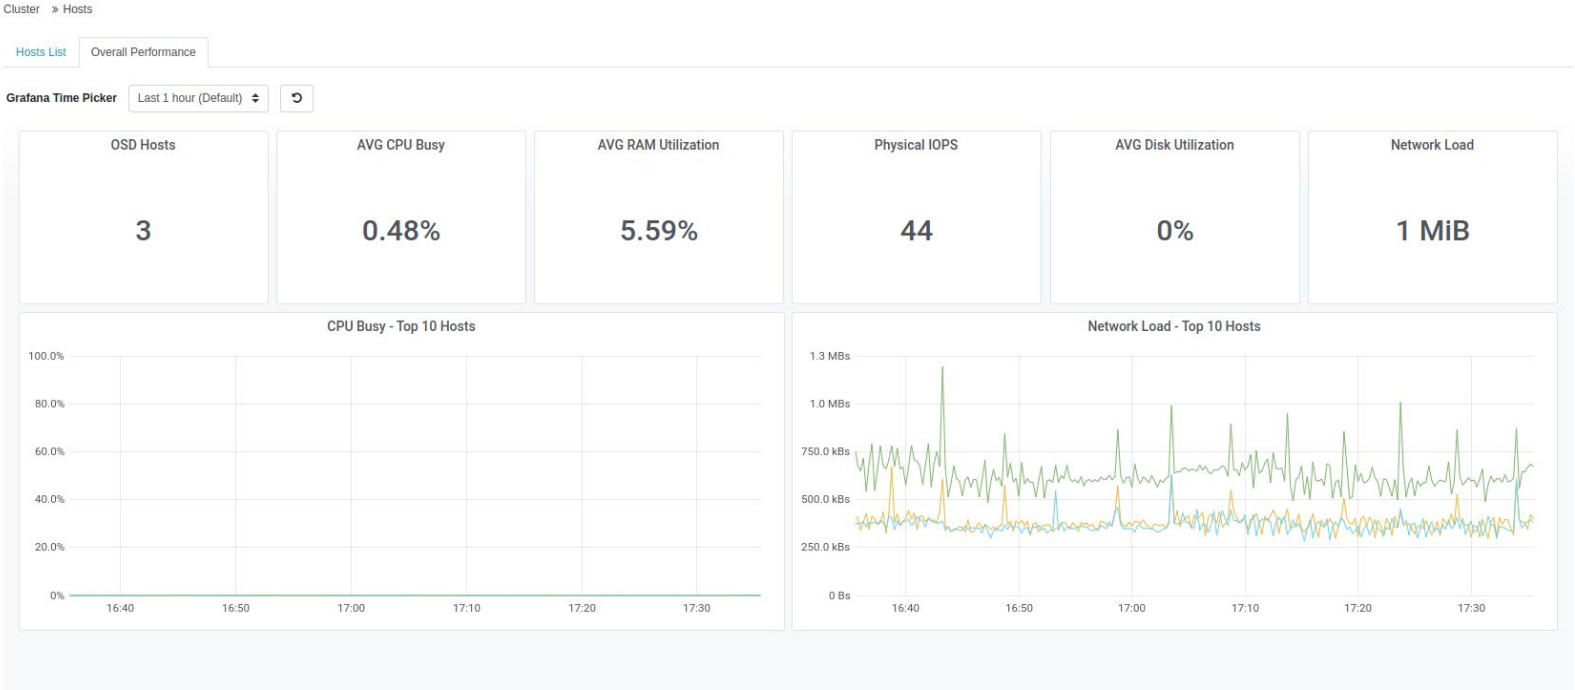

Ceph Dashboard

Host Drilled Down

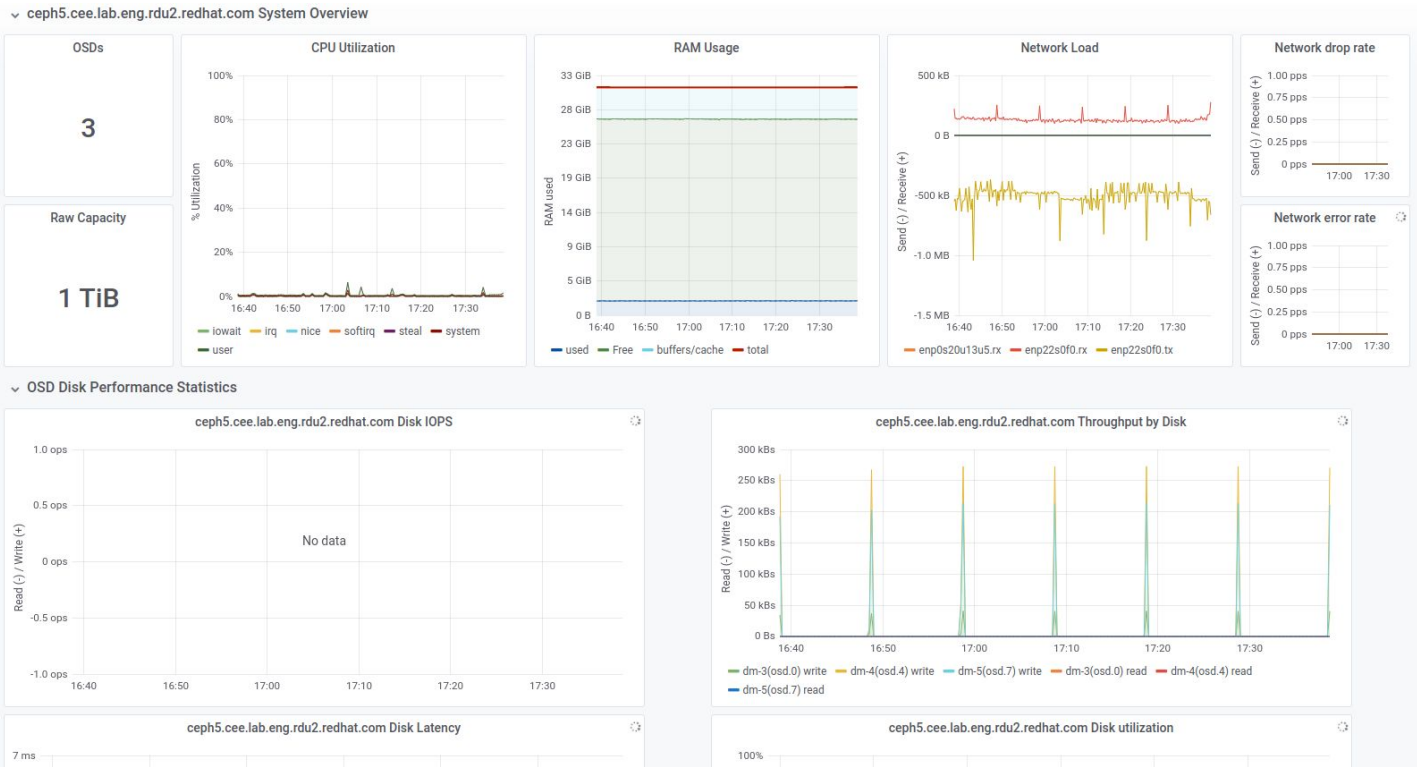

Ceph Dashboard

Pool Performance Stats

Pools

Pools List | Overall Performance

Grafana Time Picker: Last 1 hour (Default)

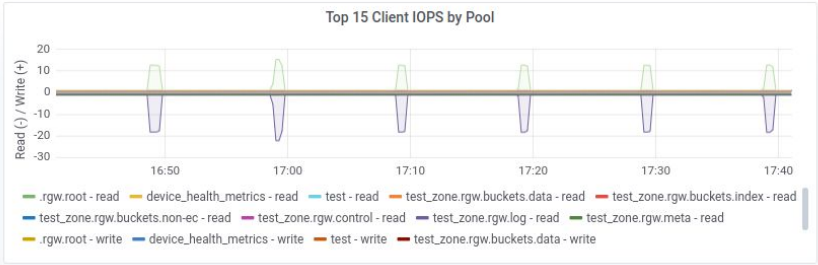

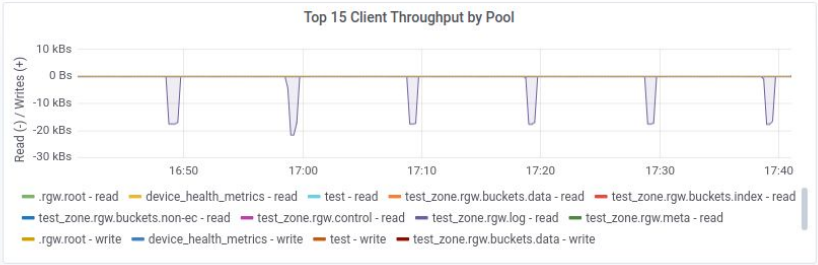

Top 15 Pools by Client IOPS

Pool Name	Pool ID	IOPS (R+W)
test_zone.rgw.buckets.non-ec	9.00	1
test_zone.rgw.buckets.data	8.00	1
test_zone.rgw.buckets.index	7.00	1
test	6.00	1
test_zone.rgw.meta	5.00	1

Top 15 Pools by Throughput

Pool Name	Pool ID	Throughput
test_zone.rgw.buckets.non-ec	9.00	1.00 B
test_zone.rgw.buckets.data	8.00	1.00 B
test_zone.rgw.buckets.index	7.00	1.00 B
test	6.00	1.00 B
test_zone.rgw.meta	5.00	1.00 B

Top 15 Pools By Capacity Used

Pool Name	Pool ID	Capacity Used
test	6.00	0.66%
test_zone.rgw.buckets.data	8.00	0.03%
.rgw.root	2.00	0.00%
test_zone.rgw.log	3.00	0.00%
test_zone.rgw.meta	5.00	0.00%

Red Hat

RGW Performance Stats

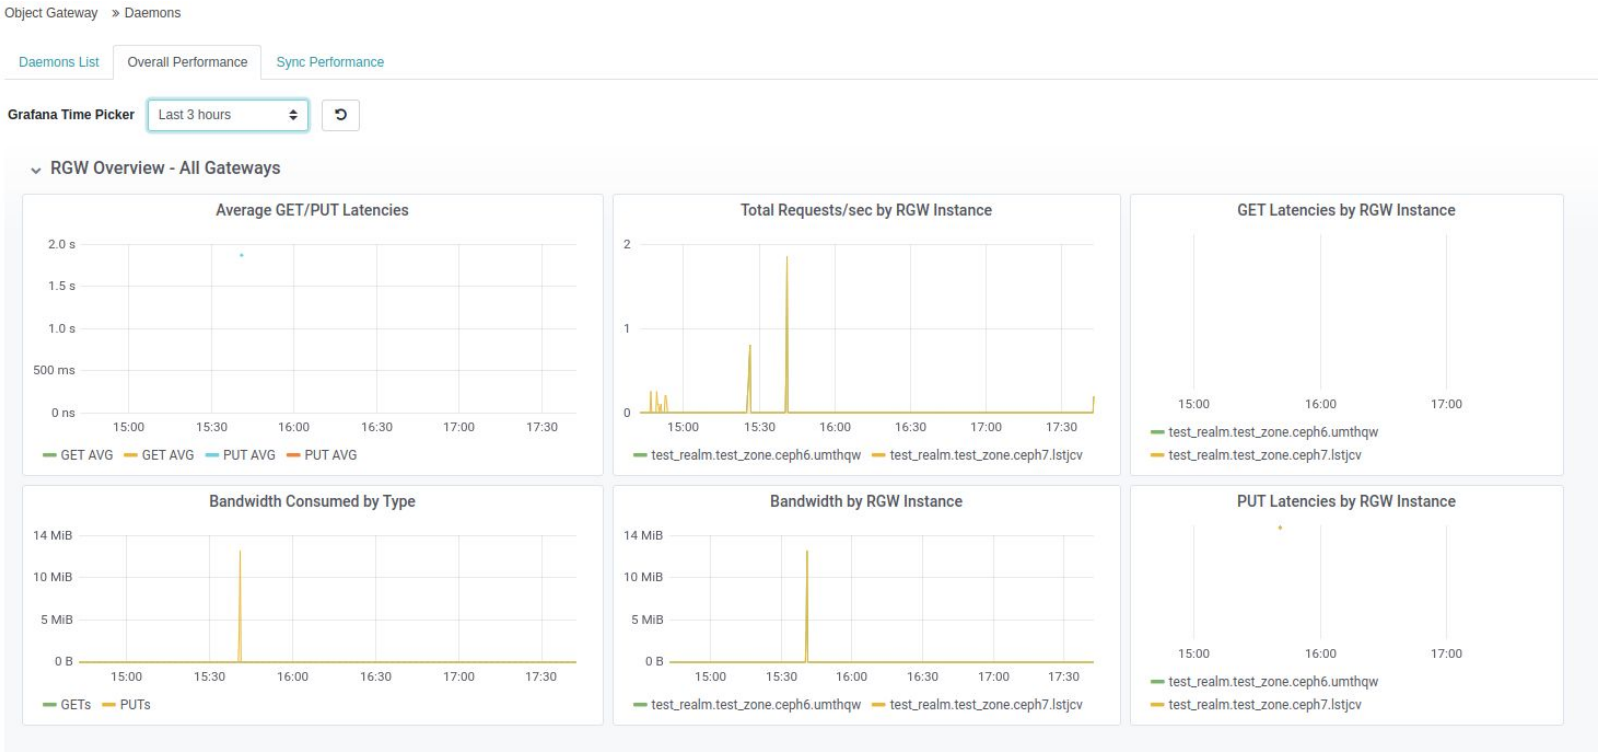

Thank you!