

# This Time, with Tentacles

Containerized Samba on Ceph

sambaXP 2024

John Mulligan


IBM

# Me @ sambaXP

- Last time I spoke at sambaXP (live) I was at Red Hat, now with IBM
  - [phlogistonjohn@asynchrone.us](mailto:phlogistonjohn@asynchrone.us)
  - [jmulliga@redhat.com](mailto:jmulliga@redhat.com) / [jmulliga@ibm.com](mailto:jmulliga@ibm.com)
- 2021, Co-Presenter with Michael Adam on Samba-Operator
- 2022, Presented on Samba Containers and Samba Operator
- 2023, Co-Presenter (via a pre-recorded demo) on Samba Containers and Samba Operator

# Let's Talk Kubernetes

Let's Talk Kubernetes

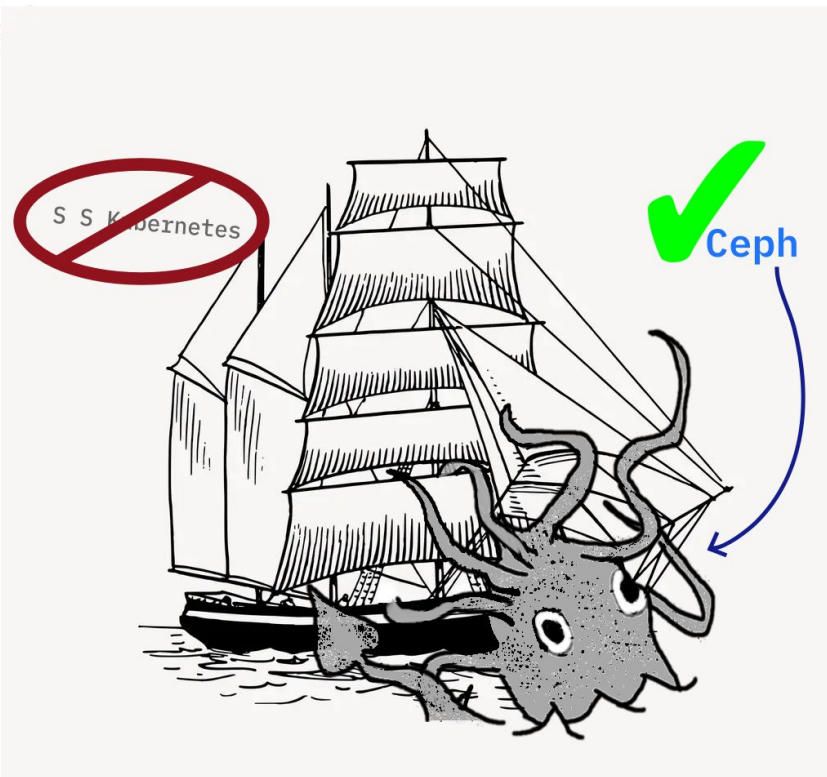


Eh, let's not...

Еу' 161,2 uof'''



# It's All About Ceph This Time



I'm on a mission to integrate Samba with Ceph to provide native SMB Shares management.

PD/CC0 images based on:  
<https://www.rawpixel.com/image/6302993>  
<https://www.rawpixel.com/image/8664576>

# Integration - What Does it Mean?

We're bringing together Samba's SMB file system layer, using some of the same components we've spoken about in previous years, under Ceph's native Orchestration layer. Shares and "clusters" can be managed using Ceph command line tools and GUI. Our workflow should be roughly similar to the existing NFS support on Ceph from a User's point-of-view.

# Reuse & Recycle

## Reuse & Recycle

- samba-containers project - the exact same images and repos on quay.io
- sambacc - same glue layer, just with a few different options
- Testing platforms
  - The SIT suite
  - See the talk by Anoop & Sachin
- What we learned doing the samba-operator
  - Not the code but some ideas can be “stolen” wholesale
  - Saves time designing things from scratch

# A One-Slide Introduction to Ceph

<https://ceph.io> - an open-source, distributed storage system

## Protocols:

- RADOS - object based storage layer that all other layers build upon
- RBD - Ceph's block device protocol
- **CephFS** - Ceph's file system protocol
- RGW - Ceph's HTTP/Object protocol

## Server Processes:

- MON - daemons that define the cluster and establish quorum
- OSD - daemons that expose storage to the cluster
- MDS - metadata server for the file system
- **MGR** - management daemon



# cephadm: Ceph's Native Orchestration Layer

- A container orchestration system, dedicated to Ceph's needs
- Runs on host systems with Podman or Docker, systemd, and python
- Server side component: a module within the Ceph MGR
- Per-node component: the cephadm "binary"
- Fewer constraints on what we can do compared to Kubernetes
  - Especially around networking
- Fewer pre-existing "building blocks" too
  - We have to build certain abstractions ourselves
  - No "guardrails", especially around networking ;-)

# They're Always Comparing Me to NFS

- Ceph has built-in management of NFS servers and exports
- A NFS module in the MGR enables GUI and CLI commands:
  - `nfs cluster create foo ...`
  - `nfs export create foo bar /path/to/dir`
  - `nfs export rm foo bar`
- The NFS module uses orchestration APIs to deploy Ceph container images that contain `nfs-ganesha` servers
- Servers can listen on the standard `nfs` ports directly or be put behind an “ingress” layer of Keepalived and HAProxy

# New Building Blocks for cephadm

## NEW Building Blocks for cephadm

- Based on our experience with samba-operator for Kubernetes we had a pretty good idea already of how we wanted to configure the Samba Containers
- Needed to add "init containers" - containers that run before daemon containers and exit
- Needed to add "sidecar containers" - containers that work alongside a primary demon container
- Enable namespace sharing for sidecar containers
- Various improvements and a bunch of refactoring along the way

# Teaching cephadm how to Samba

- Initializing the configuration & general setup
- Joining Active Directory
- Monitoring & fetching changed configuration
- Winbind for AD members
- Attaching to a CephFS file system
- Creating a Ceph "Service Spec" for deploying smb services:
  - Tell cephadm what containers to launch
  - Define the ceph "users" that will be needed
  - Locations (pseudo-URLs) where to find configuration items - mostly opaque to cephadm

# A Manager Module Made Manifest

- OK, so we can create servers, let's make it all manageable
- Commands/APIs to create/update/remove Clusters and Shares
- A declarative side to the Commands/API based on "resources":
  - Cluster
  - Share
  - Users & Groups Config
  - Join Auth Information
- "Who orchestrates the orchestrators?"
  - The smb mgr module ends up being the "brains behind the operation"
  - Understands how clusters and shares work together and instructs the cephadm layer what to deploy while taking care of the configuration cephadm is otherwise unaware of
  - Updates the share-level configuration even when orchestration changes are not needed

# DEMO

I'll be showing you a brief tour of what's available so far in the smb mgr module.

Work-in-progress PR:

<https://github.com/ceph/ceph/pull/56350>

Work-in-progress Docs:

<https://ceph--56350.org.readthedocs.build/en/56350/mgr/smb/>

Discussion thread:

<https://lists.ceph.io/hyperkitty/list/ceph-users@ceph.io/thread/PXGXDWJBPTIVLCM3Q6SKS5IN64ME3DZ/>

# Major Planned To-Do Items

- Clustering / CTDB integration
- Metrics & monitoring
- Configurability
  - How much is too much?

# Open Questions and Future Plans

## Networking approaches:

- We have a lot of flexibility but that leaves many questions
- Ceph NFS has integration with HAProxy/Keepalived called “ingress”
  - Reuse that or do our own thing?
- Limited options with need to use port 445

## Integration with NFS

- Currently nothing is planned, but we all know it's going to be requested



# Shout Out to samba-operator

- Maybe I fibbed, I have a slide on samba-operator
- We learned a lot with this effort, inspiring some of the Ceph work
- I'm fond of the operator but I have very little time to spend on it currently
- Are **you** interested in the samba-operator
- Can **you** join in and perhaps, one day, help maintain it?

<https://github.com/samba-in-kubernetes/samba-operator>

# Thank You!

## How can you participate?

- Try it out
- Review the docs
- Join our conversation

## Questions?

## John Mulligan

- [phlogistonjohn@asynchrone.us](mailto:phlogistonjohn@asynchrone.us)
- [jmulligan@redhat.com](mailto:jmulligan@redhat.com)
- [jmulliga@ibm.com](mailto:jmulliga@ibm.com)
- <http://asynchrone.us>

## Ceph Project:

- <http://ceph.io>
- <https://github.com/ceph/ceph>
- Community Calendar:  
<https://ceph.io/en/community/meetups/>