Multi-cluster data replication with NooBaa

Guy Margalit, Red Hat
Erin Boyd, Red Hat
Emerging Technology Track - Office of the CTO
May 9, 2019

LEGAL DISCLAIMER

The content set forth herein does not constitute in any way a binding or legal agreement or impose any legal obligation or duty on Red Hat. This information is provided for discussion purposes only and is subject to change for any or no reason.

Why should I care about multi-cluster ?

- ★ Performance
- ★ Fault tolerance
- ★ Specialized HW or services

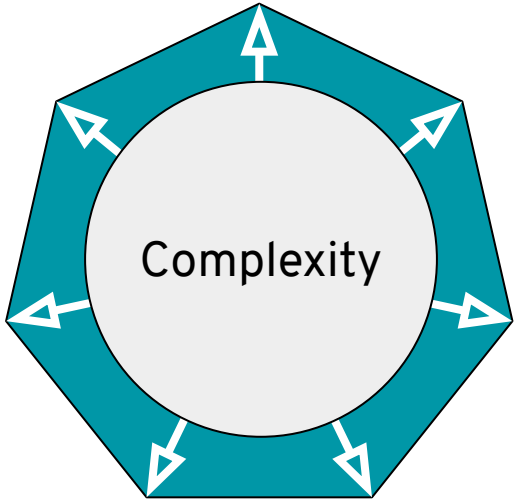

- ★ Vendor lock-in
- ★ Regulation
- ★ Collaboration

How applications data is shared in multi-cluster ?

Main Data Services

- ★ Object-store
- ★ Database
- ★ Filesystem
- ★ Queue

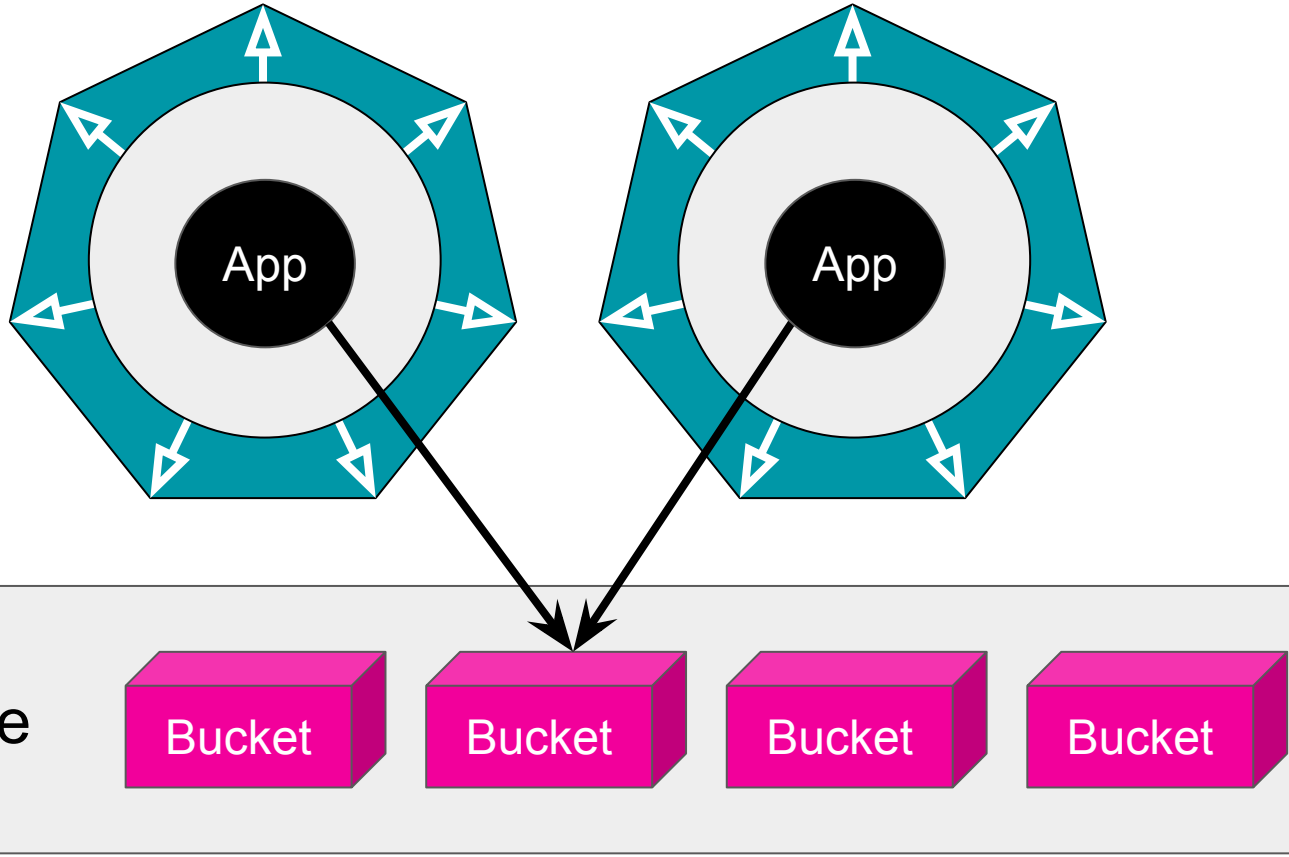
App
App
Object-store
Bucket
Bucket
Bucket
Bucket

Data Services are siloed by default

The problems with siloed data are the same as we described for a single cluster.

Reminder

- ★ Performance
- ★ Fault tolerance
- ★ Specialized HW or services

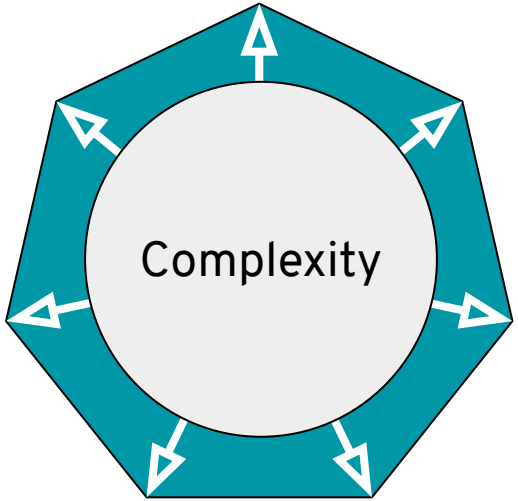

- ★ Vendor lock-in
- ★ Regulation
- ★ Collaboration

What is the solution of NooBaa’s data service ?

NooBaa is the non siloed Object-store

- ★ Top layer - Applicative
  - Buckets, Accounts, Permissions
- ★ Middle layer - Policy
  - Mirror, Tiering, Spread
- ★ Bottom layer - Storage
  - PVs and Cloud storage
  - Location aware

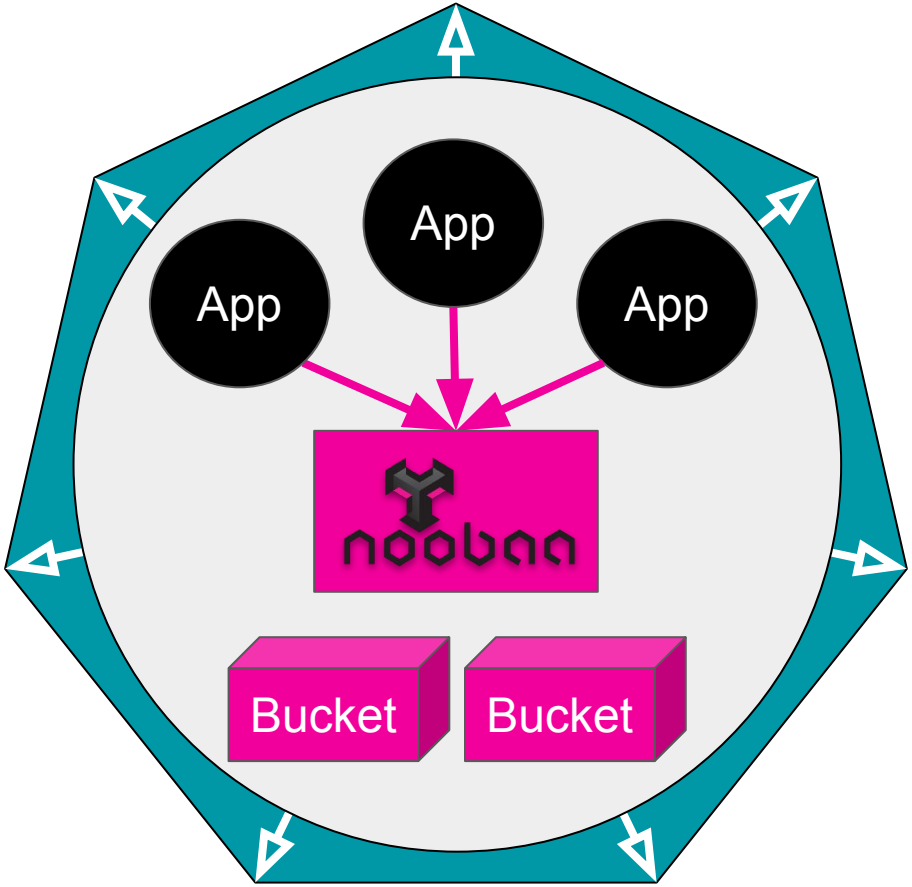

NooBaa's top layer is providing a scalable S3 service for applications.

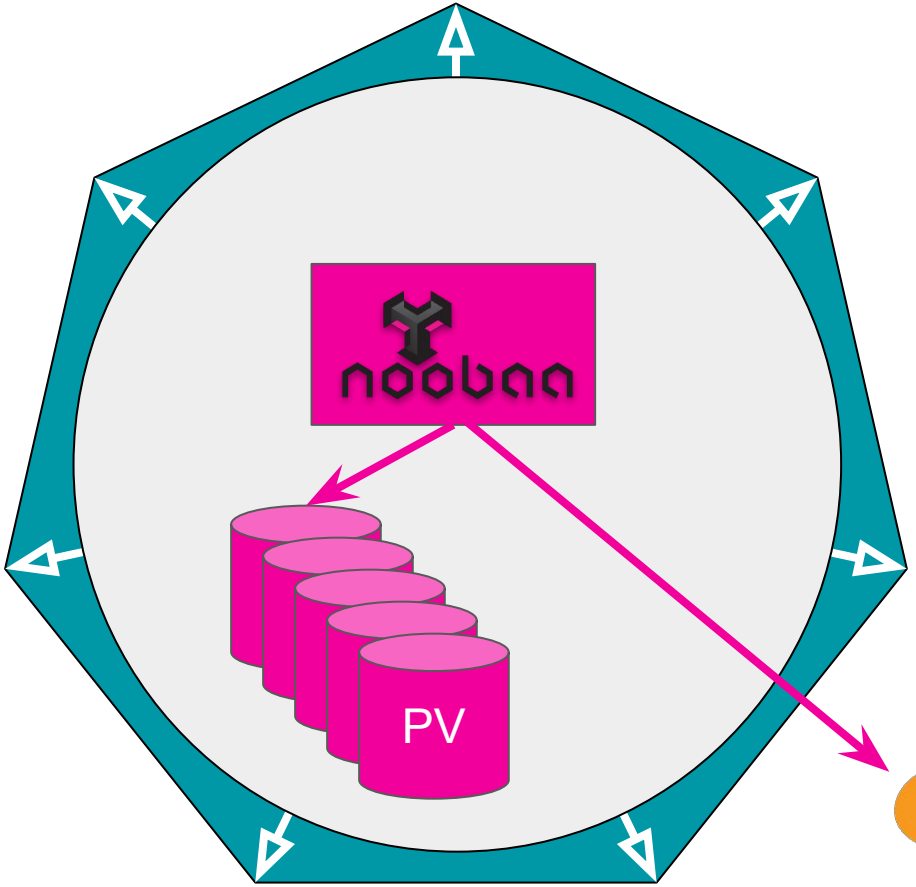

NooBaa's bottom layer connects to any storage such as PVs and cloud providers object storage outside of the cluster.

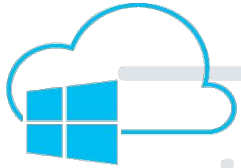

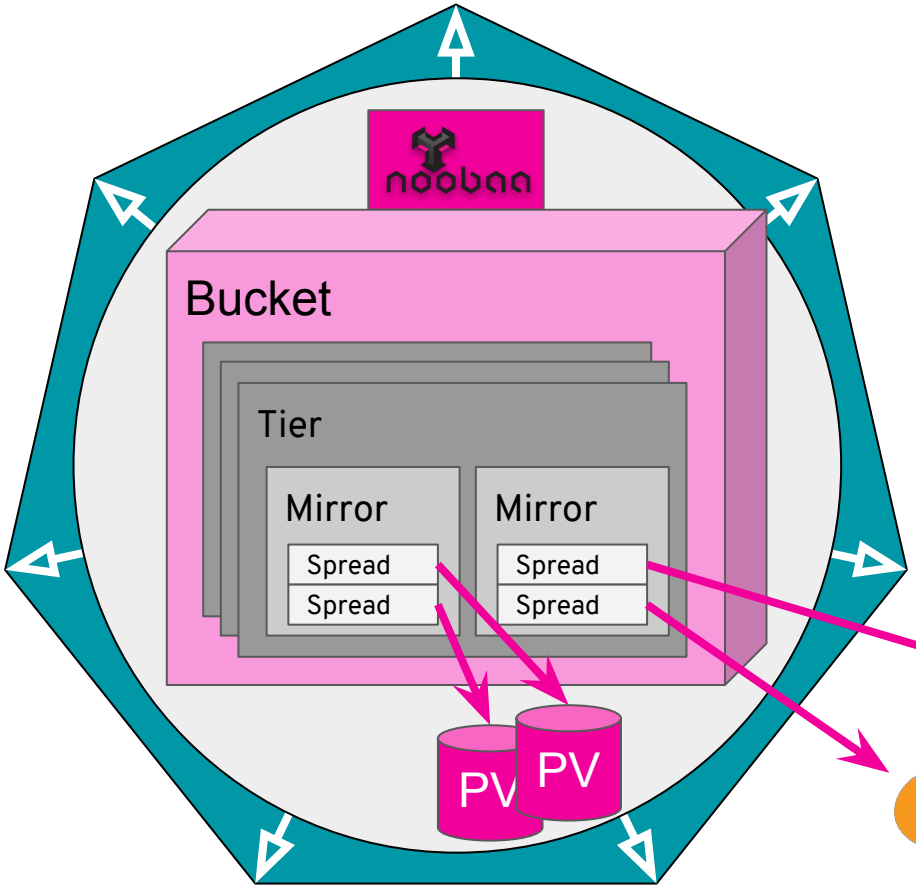

NooBaa's middle layer defines a flexible placement policy per bucket - **tiering**, **mirror** and **spread** over any set of resources.

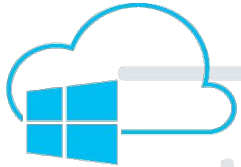

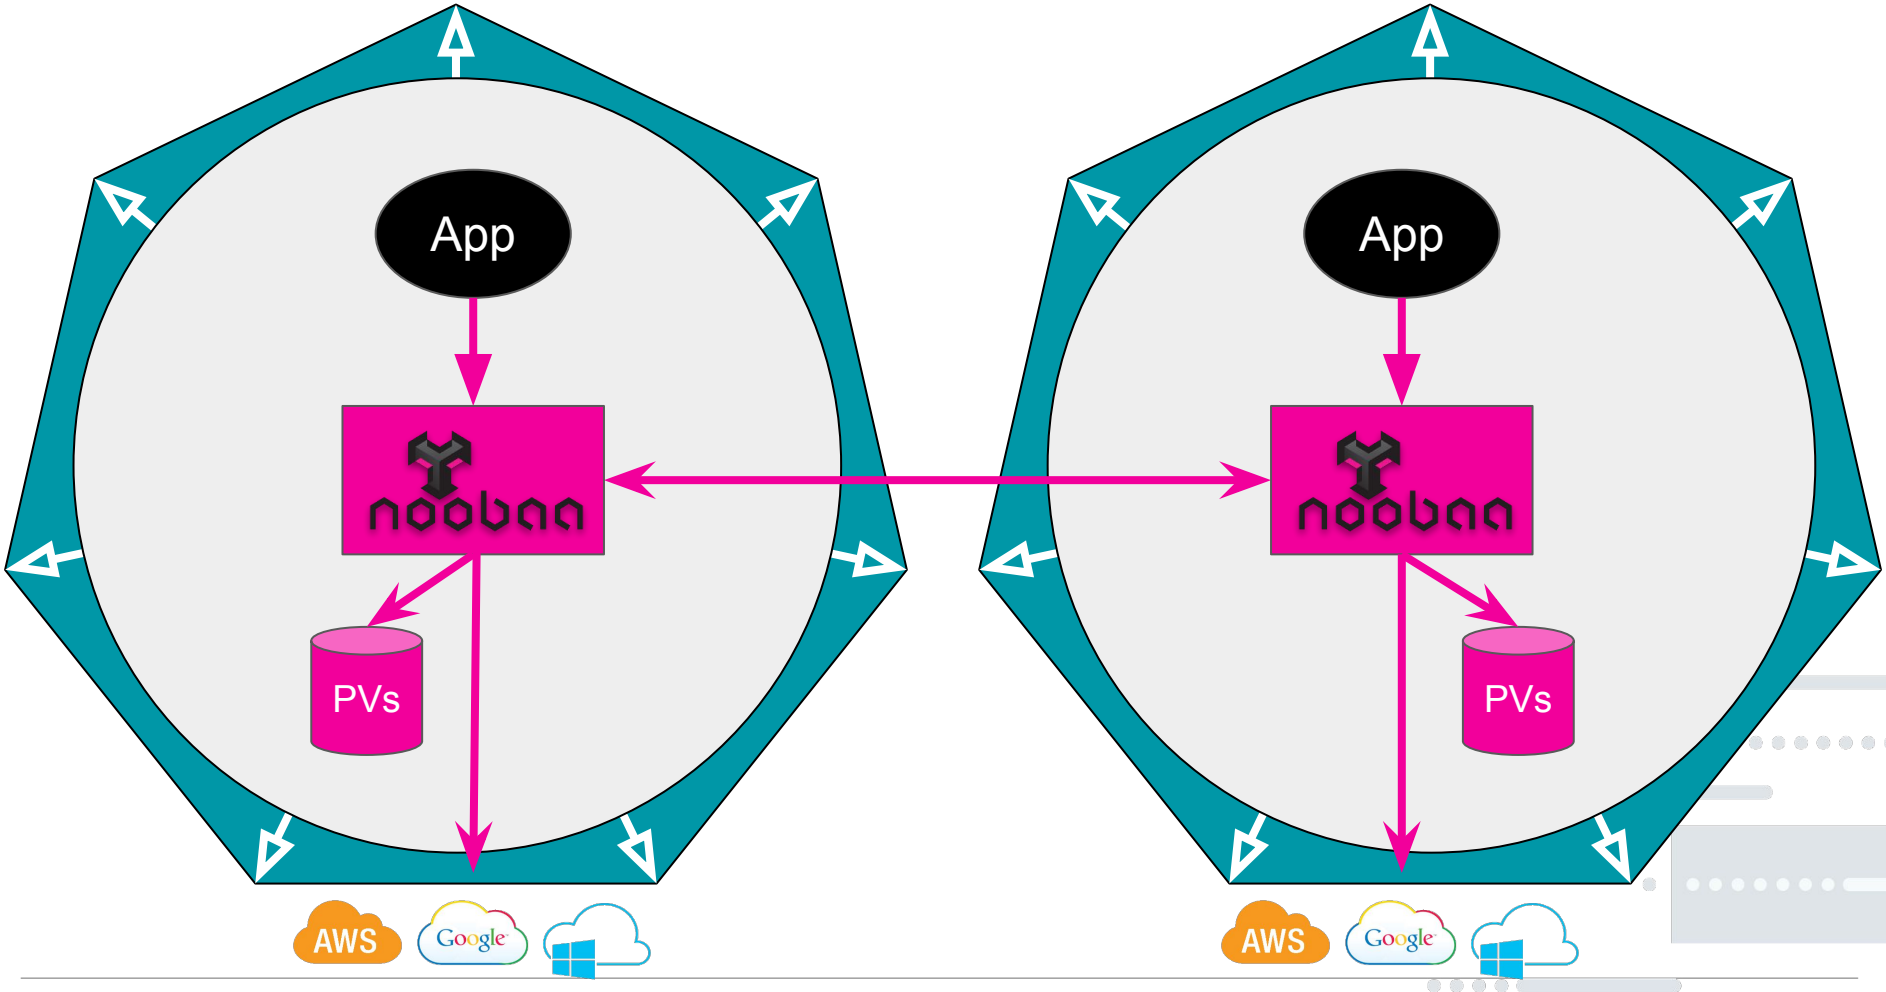
App
noobaa
PVs
AWS
Google
App
noobaa
PVs
AWS
Google

Making it super easy to manage

Zero touch - New Cluster

- ★ Federate NooBaa services to new clusters using Kubernetes Federation v2:
  - https://github.com/kubernetes-sigs/federation-v2/
- ★ Connect local cluster storage and “local” cloud storage defined by Federated Storage Classes
- ★ Ready to serve new/federated applications

Zero touch - New Application

- ★ With Object Bucket Claim (OBC) applications dynamic provisioner their buckets in kubernetes
- ★ The claim will be processed in every deployed cluster creating the object bucket native to that platform
- ★ NooBaa’s default reaction to a bucket claim is to update the bucket placement to mirror to local cluster resources, thus seamlessly reducing latency and increasing throughput.

Demo

Live demo

- ★ Multi-cluster setup
- ★ Storage resources
  - adding PVs, cloud, removing...
- ★ Bucket placement
  - lets make some changes

The NooBaa project

https://github.com/noobaa/noobaa-core

★ We invite you to Star and Watch to get notifications on releases.

★ Open issues for any question you have.

https://www.noobaa.com/try

https://www.noobaa.com/community

- ★ Subscribe to our newsletter.
- ★ We will also invite you to our Slack discussions.

Would love to talk

- https://twitter.com/NooBaaStorage
- support@noobaa.com
- You are welcome to reach out to me directly

  gmargali@redhat.com

Q&A

(Hallway track?)

THANK YOU

linkedin.com/company/Red-Hat

facebook.com/RedHatinc

youtube.com/user/RedHatVideos

twitter.com/RedHat