

# Multi-tenant Multi-cluster Kubernetes Datapocalypse is Coming

Live Webcast  
June 23, 2020  
10:00 am PT



# Today's Presenters



**Paul Burt**  
**Technical Product Marketing Engineer**  
**NetApp**



**Jim Fister**  
**Principal**  
**The Decision Place**

# SNIA Legal Notice

- The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

# SNIA-At-A-Glance



**185**

industry leading  
organizations



**2,000**

active contributing  
members



**50,000**

IT end users & storage  
pros worldwide



# What We Do



**Educate** vendors and users on cloud storage, data services and orchestration



**Support & promote** business models and architectures: OpenStack, Software Defined Storage, Kubernetes, Object Storage



**Understand** Hyperscaler requirements  
Incorporate them into standards and programs



**Collaborate** with other industry associations

# Agenda

1. A bank goes down on a Friday  
An example of failure while running just one big cluster
2. A few good ~~men-ants~~ tenants  
Multi-tenancy and other reasons why multi-cluster is unavoidable
3. A 7-unit multi-tenant investment property  
Emerging solutions for multi-cluster, and anticipating the problems multi-cluster brings



"A good science fiction story  
should be able to predict not the  
automobile, but the traffic jam."

Frederik Pohl



# A bank goes down on a Friday

“Don't go spending it all on... I don't know. Caroline, what do these people buy? Tattered hats? Beard dirt?”



# Monzo's 2018 KubeCon Keynote

Just a minor conflict  
between etcd and Java



YouTube

Search



SIGN IN



KubeCon



CloudNativeCon

Europe 2018

# Keynote: Anatomy of a Production Kubernetes Outage

Oliver Beattie, Head of Engineering,  
Monzo Bank



0:01 / 17:44

Keynote: Anatomy of a Production Kubernetes Outage - Oliver Beattie, Head of Engineering, Monzo Bank

31,725 views • May 4, 2018

Like 366    4    SHARE    SAVE    ...





## **RESOLVED: Current account payments may fail - Major Outage (27/10/2017)**

 Help

- **15:13:** All linkerd pods are restarted, but services that process thousands of requests per second are now receiving no traffic. At this point, customers are completely unable to refresh their feed or balance in the Monzo app and our internal COps (“Customer Operations” 🧑) tools stop working. **The issue has now escalated to a full platform outage, and no services are able to serve requests.** As you can probably imagine, practically all of our automated alerts started triggering. 🖥️🔥
- **15:27:** We notice that linkerd is logging `NullPointerException` <sup>69</sup> when it is attempting to parse the service discovery response from the Kubernetes apiserver. We discover that this is an `incompatibility` <sup>422</sup> between the versions of Kubernetes and linkerd that we’re running, and specifically is a failure to parse empty services.

Because we have been testing an updated version of linkerd in our staging environment for several weeks which contains a fix for the incompatibility, engineers from the Platform team begin deploying a new version of linkerd in an attempt to roll forward.

A few good ~~men-ants~~  
**ten** -ants

“Apparently he's not very happy down here  
in Shangri-la because he's written letters to  
everyone but Santa Claus asking for a  
**transfer...**”



# We've been doing it all wrong

Security, multi-tenancy, and most of Kubernetes  
is approached incorrectly by the vast majority of us





# Challenge: multi-tenancy

How multiple teams are able to work in harmony



# Hard Multi-Tenancy in Kubernetes

Friday, May 18, 2018

**EDIT:** See my post on a [design doc for a multi-tenant orchestrator](#) instead. I wrote this when an internal requirement was to use Kubernetes but I do not personally think you should use Kubernetes for this use case.

Kubernetes is the new kernel. We can refer to it as a “cluster kernel” versus the typical operating system kernel. This means a lot of great things for users trying to deploy applications. It also leads to a lot of the same challenges we have already faced with operating system kernels. One of which being privilege isolation. In Kubernetes, we refer to this as multi-tenancy, or the



# Challenge: security

How to manage clusters without becoming a news story





KUBERNETES / SECURITY

# Kubernetes 'Billion Laughs' Vulnerability Is No Laughing Matter

9 Oct 2019 8:11am, by [Jack Wallen](#)

A new vulnerability has been discovered within the Kubernetes API. This flaw is centered around the parsing of YAML manifests by the Kubernetes API server. During this process, the API server is open to potential Denial of Service (DoS) attacks. The issue ([CVE-2019-11253](#) — which has yet to have any details

CRYPTOCURRENCY JACKING —

# Tesla cloud resources are hacked to run cryptocurrency-mining malware

Crooks find poorly secured access credentials, use them to install stealth miner.

DAN GOODIN - 2/20/2018, 2:21 PM



## Public Cloud Security Guide

See How The Major Cloud Providers Stack Up



Get The Guide

Keeping clusters up to date is important



# Challenge: multi-cloud

Real companies are building on multiple infrastructure providers



## It's not just Google — Snap has a \$1 billion cloud services deal with Amazon, too

Snap disclosed the deal in a revised version of its S-1 made public Thursday.

By **Karl Wagner** | Feb 9, 2017, 6:45am EST

[f](#) [t](#) [s](#) SHARE



Ass Mathat for Recode

*NEW YORK*

# Get Lost

ON

## One Great Story

Every day, delivered  
right to your inbox.



GET THE STORY

# Introducing Anthos: An entirely new platform for managing applications in today's multi-cloud world



Urs Hölzle

VP, Google Cloud Platform, Google Cloud

Let's face it, even in the best of cases, enterprise IT can be rigid, complex and expensive.





# Storage Scalability in Hybrid Cloud and Multicloud Environments

As data growth in enterprises continues to skyrocket, datacenter cloud scalability, whether on premises, in hybrid cloud or in multicloud deployments, is key for businesses. So, what are enterprise IT organizations supposed to do, given that 'run anything anywhere' is becoming more important than ever?

[...more](#)

---

Recorded Jun 11 2020 | 59 mins

## Presented by

Robert Cone, Intel; Piyush Chaudhary, IBM; Pekon Gupta, SMART Modular

---

Multi-cloud means multiple clusters  
(maybe even heterogeneous clusters)

# Challenge: large scale

Kubernetes has limits





# Quotas and limits

[Send feedback](#)[Contents](#)[Resource quotas](#)

This page explains the quotas and limits for Google Kubernetes Engine clusters, nodes, and GKE API requests.

GKE's per-project limits are:

- Maximum of 50 clusters per zone, plus 50 [regional clusters](#) per region.

GKE's per-cluster limits are:

- Maximum of 5000 nodes per cluster.
- Maximum of 1000 nodes per [node pool](#).
- Maximum of 1000 nodes per cluster, if you use the [GKE ingress controller](#) [↗](#).
- Maximum of 110 Pods per node.



# Even Kubernetes has limits



That's a lot of challenges.  
Are there solutions?  
(yes)

# A 7-unit multi-tenant investment property

“I'm sorry, quick question. From the time you guys started talking, Bear Stearns stock has fallen more than 38%.

Would you still buy more?”



# Service meshes

The thing that connects clusters and services

# Architecting for Multicluster Kubernetes



Thomas Rampelberg

February 17, 2020 • 6 min read

Recently, the Linkerd community has been spending time tackling the challenge of multicluster Kubernetes. How can we apply features like Linkerd's zero-config automatic mTLS, or traffic splitting, across multiple Kubernetes clusters? What should the service mesh do, and more importantly: what should the service mesh not do?

Like any good engineering project, the best way to start is by getting clear on the requirements. In this blog post, we outline the minimal requirements of a multi-cluster solution that makes cross-cluster traffic more reliable, secure and observable. In subsequent blog posts, we'll address some of the implementation choices.

## From whence dost thou multi-cluster, sirrah?

Kubernetes clusters are like Pringles - you can't just have one! In fact, some organizations such as [Zalando](#) have around 100 of them. By running multiple clusters, it becomes possible to keep the concerns of each cluster separate. Instead of having to solve the constraints of every application and solution, the problem space shrinks. This ends up being a fantastic tool to allow for architecture solutions in a more flexible, simple way.

# From Monolith to Multicloud Micro-Services: Inside Snap's Service Mesh

By Snap Engineering

on Wednesday, March 04, 2020





# KubeFed

The “official” approach to multi-cluster

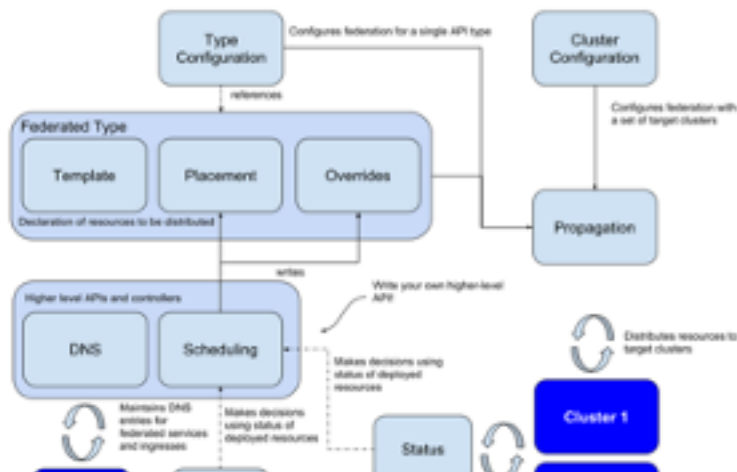
build [building](#) [go report](#) [3.4](#) [container](#) [image](#) [license](#) [apache2.0](#) [release](#) [v0.3.0](#)

## Kubernetes Cluster Federation

Kubernetes Cluster Federation (KubeFed for short) allows you to coordinate the configuration of multiple Kubernetes clusters from a single set of APIs in a hosting cluster. KubeFed aims to provide mechanisms for expressing which clusters should have their configuration managed and what that configuration should be. The mechanisms that KubeFed provides are intentionally low-level, and intended to be foundational for more complex multicluster use cases such as deploying multi-geo applications and disaster recovery.

KubeFed is currently **alpha** and moving rapidly toward its initial **beta release**.

### Concepts



## Verifying API type is installed on all member clusters

If the API type is not installed on one of your member clusters, you will see a repeated `controller-manager` log error similar to the one reported in [issue 314](#). At this time, you must manually verify that the API type is installed on each of your clusters as the `controller-manager` log error is the only indication.

For an example API type `bars.example.com`, you can verify that the API type is installed on each of your clusters by running:

```
CLUSTER_CONTEXTS="cluster1 cluster2"
for c in ${CLUSTER_CONTEXTS}; do
    echo ----- ${c} -----
    kubectl --context=${c} api-resources --api-group=example.com
done
```

The output should look like the following:

```
----- cluster1 -----
NAME      SHORTNAMES  APIGROUP      NAMESPACED  KIND
bars                               example.com   true         Bar
----- cluster2 -----
NAME      SHORTNAMES  APIGROUP      NAMESPACED  KIND
```





# Custom or commercial k8s

I give you money. You make the hard problems go away.

Our intention is to open the codebase for Red Hat Advanced Cluster Management for Kubernetes in chunks. Components will be released as separate projects, with priority given to those projects that we judge to be most suitable for community contribution. Currently there are 20 repositories available. Among the most interesting are the Application Lifecycle projects:

- <https://github.com/open-cluster-management/multicloud-operators-subscription>
- <https://github.com/open-cluster-management/multicloud-operators-placemetrule>
- <https://github.com/open-cluster-management/multicloud-operators-subscription-release>
- <https://github.com/open-cluster-management/multicloud-operators-channel>
- Demonstrations: <https://github.com/open-cluster-management/demo-subscription-gitops>

# Harvest the Benefits of a Multi-cluster Kubernetes Architecture with VMware Tanzu Mission Control

APRIL 30, 2020 | NING GE

When architecting the Kubernetes platform for your enterprise, among the fundamental questions you need to address first and foremost are "How many clusters will I need?" and "Do I need a few big clusters, or many smaller ones?"



# Problem solved?

I have some bad news, friend

# On Infrastructure at Scale: A Cascading Failure of Distributed Systems

Posted by Dan Woods

Jan 14, 2019

[distributed systems](#)[kafka](#)[consul](#)[kubernetes](#)[chaos](#)

At Target, we run a heterogeneous infrastructure in our datacenters (and many other places), where we have multiple different backend hosting infrastructure for workloads. Most of this is a legacy artifact of putting infrastructure into production for different use-cases and application development and deployment patterns. The [Target Application Platform \(TAP\)](#) provides a common interface for running and managing workloads, where we can spread workloads across different hosting infrastructure transparently to applications. We use this as a tool to better manage capacity, fully utilize the infrastructure we have available, and to move applications from one hosting provider to another without requiring application teams to re-platform

A number of questions need to be  
answered

- How is traffic routed across clusters?
- How do we move apps/services between clusters?
  - Stateless apps?
  - Stateful apps?
- How are security policies kept consistent?

*Service meshes or cloud managed load balancers?*

- How is traffic routed across clusters?
- How do we move apps/services between clusters?

*KubeFed? Knative?*

- Stateless apps?
- Stateful apps?

*Project Velero? Portworx?  
Project Astra? Minio?  
Rook.io? OpenEBS?*

- How are security policies kept consistent?

*Calico or Cilium, mixed with KubeFed and a lot of crying?*





Kubernetes is the king of container management. Between October 2015 and October 2019—a mere four-year timespan—the share of Kubernetes job searches increased by 2,125%. Kubernetes-related jobs per million also grew by 2,141% in that same time period, according to a [blog post](#) from Seen by Indeed.

**SEE: [What is Kubernetes? \(free PDF\)](#) (TechRepublic)**

"Kubernetes gained dominance by bridging the gap between development and application deployment," said Thomas Hatch, CTO and co-founder of SaltStack, a security operations provider.

#### Must-read developer content

- [Learn COBOL with these online training courses and tutorials](#)
- [Python: Cheat sheet \(free PDF\)](#)
- [Top 5 books developers should read](#)
- [3 essential hiring kits for key developer jobs](#)



{ ~~Service meshes or cloud managed load balancers?~~

- How is traffic routed across clusters?
- How do we move apps/services between clusters?

*KubeFed? Knative?*

- Stateless apps?
- Stateful apps?

*Project Velero? Portworx?  
Project Astra? Minio?  
Rook.io? OpenEBS?*



- How are security policies kept consistent?

*Calico or Cilium, mixed with KubeFed and a lot of crying?*



*Service meshes or cloud managed load balancers?*

- How is traffic routed across clusters?
- How do we move apps/services between clusters?

*KubeFed? Knative?*

- Stateless apps?
- **Stateful apps?**

*Project Velero? Portworx?  
Project Astra? Minio?  
Rook.io? OpenEBS?*

- How are security policies kept consistent?

*Calico or Cilium, mixed with KubeFed and a lot of crying?*

*Service meshes or cloud managed load balancers?*

- How is traffic routed across clusters?
- How do we move apps/services between clusters?

*KubeFed? Knative?*

- Stateless apps?
- Stateful apps?

*Project Velero? Portworx?  
Project Astra? Minio?  
Rook.io? OpenEBS?*

- How are security policies kept consistent?

*Calico or Cilium, mixed with KubeFed and a lot of crying?*



"A good science fiction story  
should be able to predict not the  
automobile, but the traffic jam."

Frederik Pohl





# What should we do?

Follow early solutions, recognize we don't need multi-cluster multi-tenant architectures (yet),  
and **frame the problem in terms of queryability.**



# Follow early solutions

KubeFed is about to go from alpha to Beta.

Everyone and their grandma seems to have a new service mesh these days.



# A standard interface for service meshes on Kubernetes.

## Service Mesh Interface provides:

- A standard interface for service meshes on Kubernetes
- A basic feature set for the most common service mesh use cases
- Flexibility to support new service mesh capabilities over time
- Space for the ecosystem to innovate with service mesh technology

[View the Spec](#)



# We don't need multi-cluster multi-tenant (yet)

Most of us running Kubernetes are only just dipping our toes into the container waters.

Some hyper-scale companies are at the point where they need multi-cluster multi-tenant rigor. We should all watch them, and see what their experiences are.

# Think in terms of queryability

One of the superpowers of Kubernetes is the ability to query.

We need that same superpower for multi-cluster management.

We also need it for storage management!

name@env~:\$ kubectl get storageclasses

NAME	PROVISIONER	RECLAIMPOLICY	VOLUMEBINDINGMODE
ALLOWVOLUMEEXPANSION	AGE		
netapp-cvs-extreme	csi.trident.netapp.io	Delete	Immediate
26h			true
netapp-cvs-premium (default)	csi.trident.netapp.io	Delete	Immediate
26h			true
netapp-cvs-standard	csi.trident.netapp.io	Delete	Immediate
26h			true
standard	kubernetes.io/gce-pd	Delete	Immediate
27h			true

## Verifying API type is installed on all member clusters

If the API type is not installed on one of your member clusters, you will see a repeated `controller-manager` log error similar to the one reported in issue [314](#). At this time, you must manually verify that the API type is installed on each of your clusters as the `controller-manager` log error is the only indication.

For an example API type `bars.example.com`, you can verify that the API type is installed on each of your clusters by running:

```
CLUSTER_CONTEXTS="cluster1 cluster2"
for c in ${CLUSTER_CONTEXTS}; do
    echo ----- ${c} -----
    kubectl --context=${c} api-resources --api-group=example.com
done
```

The output should look like the following:

```
----- cluster1 -----
NAME      SHORTNAMES  APIGROUP      NAMESPACED  KIND
bars                               example.com   true         Bar
----- cluster2 -----
NAME      SHORTNAMES  APIGROUP      NAMESPACED  KIND
```

# Visualize your Kubernetes workloads

Octant is an open source developer-centric web interface for Kubernetes that lets you inspect a Kubernetes cluster and its applications.

[LATEST RELEASE INFORMATION](#)[DOWNLOAD OCTANT](#)

## Visualization

Provide a visual interface to managing Kubernetes that complements and extends existing tools like kubectl and kustomize.



## Extensibility

Add information to your cluster views through Octant's plug-in system.



## Versatility


Support a variety of debugging features such as filtering labels and streaming container logs to be part of the Kubernetes development toolkit.

# Run (mostly) homogenous clusters

We might want to strive for mostly homogenous Kubernetes clusters to simplify things.

We'll win in the short-run, and lose in the long-run.











The same problem for k8s infra



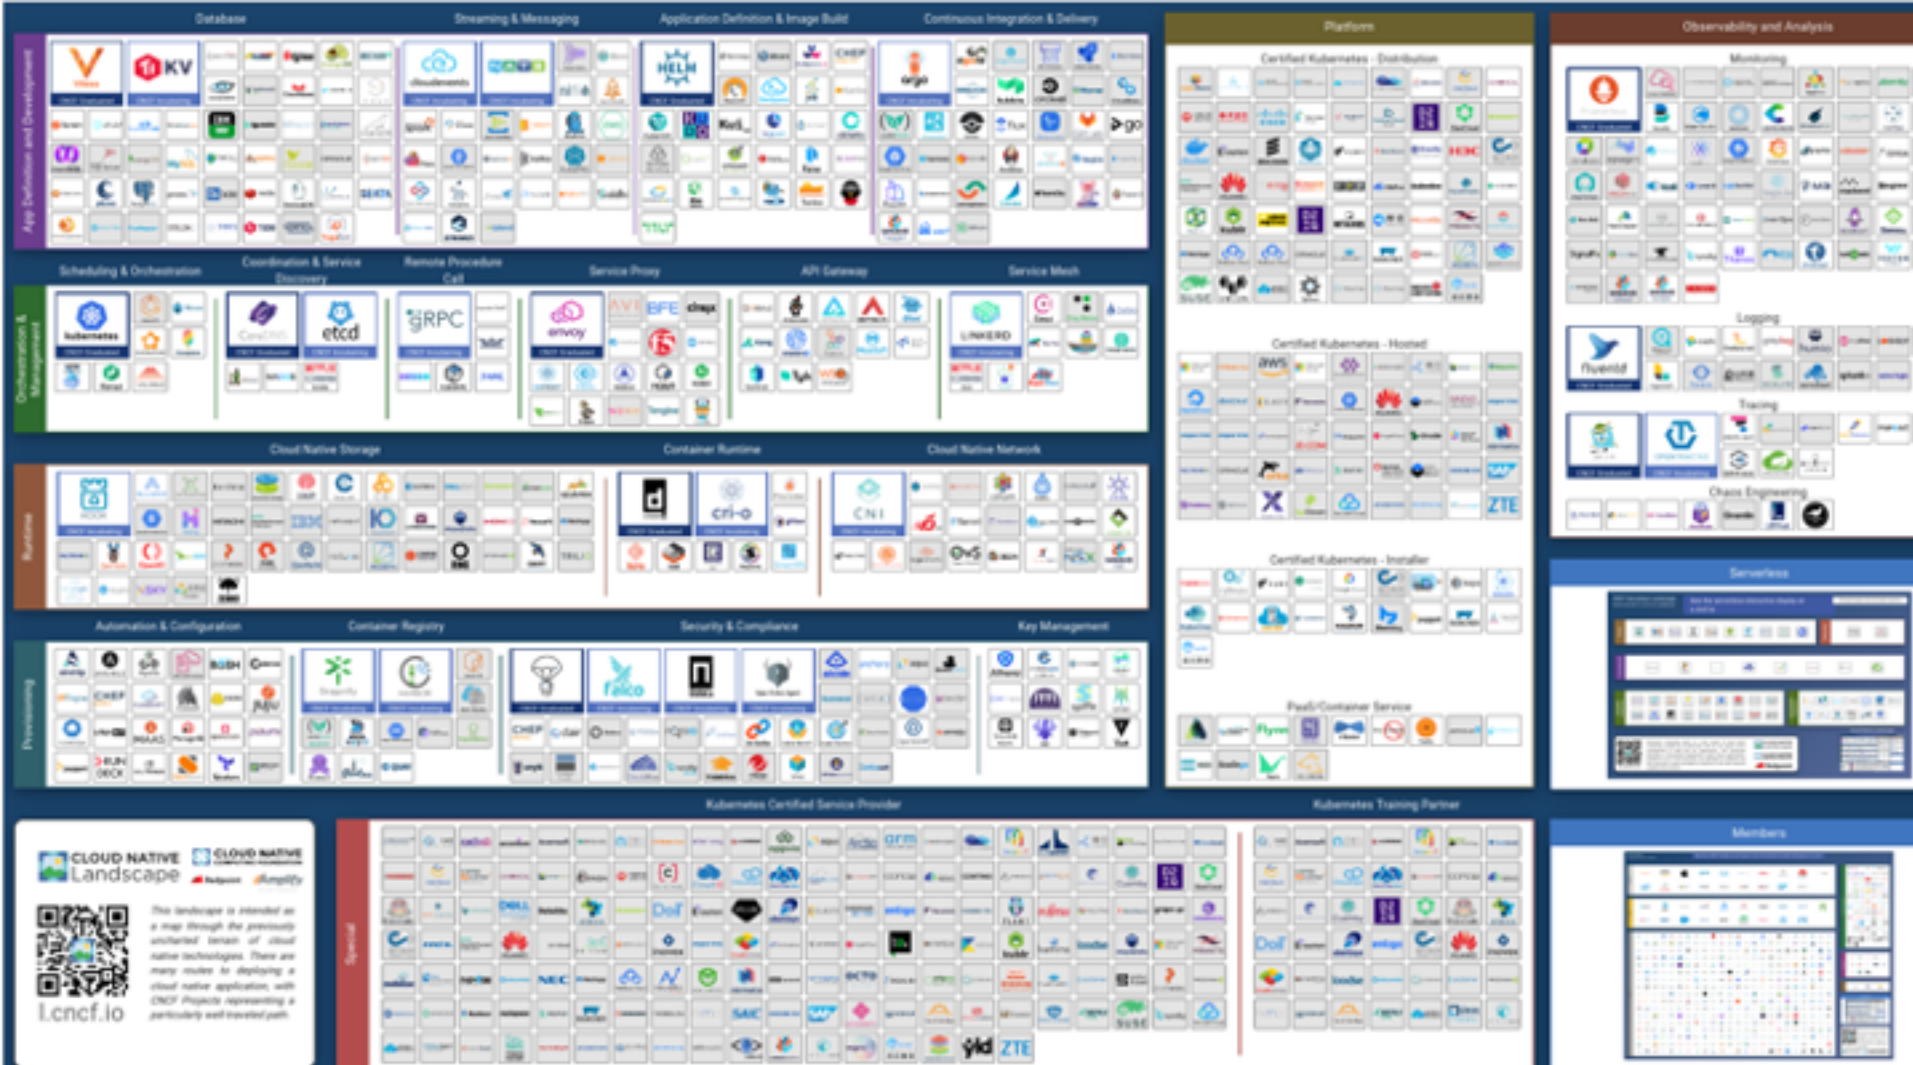
## RESOLVED: Current account payments may fail - Major Outage (27/10/2017)

■ Help

- **15:13:** All linkerd pods are restarted, but services that process thousands of requests per second are now receiving no traffic. At this point, customers are completely unable to refresh their feed or balance in the Monzo app and our internal COps ("Customer Operations" 🧑) tools stop working. **The issue has now escalated to a full platform outage, and no services are able to serve requests.** As you can probably imagine, practically all of our automated alerts started triggering. 🖥️🔥
- **15:27:** We notice that linkerd is logging `NullPointerException` <sup>69</sup> when it is attempting to parse the service discovery response from the Kubernetes `apiserver`. We discover that this is an `incompatibility` <sup>422</sup> between the versions of Kubernetes and linkerd that we're running, and specifically is a failure to parse empty services.

Because we have been testing an updated version of linkerd in our staging environment for several weeks which contains a fix for the incompatibility, engineers from the Platform team begin deploying a new version of linkerd in an attempt to roll forward.

So just let anyone run anything?  
No, that's pure chaos.



We want clusters with good defaults,  
and flexibility to deviate when truly needed

We want clusters with good defaults,  
and flexibility to deviate when truly needed

*Note:*

*Deviations are why queryability is extra important!*



# Recognize that data will be the hardest part

Security, scalability, tenancy... all of those considerations multiply the challenge of managing the data side of Kubernetes.

Some of the simple stateful solutions will work just fine for now. Complexity will 100% be needed later.

April 22, 2019 | IT Life

## Moving databases between servers with minimal downtime

By mafields

Recently we decided to roll out new database infrastructure for many of our backed systems. One of the challenges we faced after deploying the new infrastructure was how to migrate the data from the current production database server to the new server without having to take the system offline while we moved the data. The easiest way to move the data would have been to take the service offline, move the data over, and then point clients at the new server; however, this outage is a massive inconvenience for customers, and so we wanted to



**Grant Shipley**

@gshipley



My kids just requested I give them an outage window when I reboot the router going forward so they can plan around it. What a world we live in.

11:15 AM · Jun 16, 2020 · [Twitter Web App](#)

---

**179** Retweets   **1.2K** Likes



Simple tools, good for starters





Complex tools, for as you mature




*Note:  
Look for queryability!*



# In Conclusion

Everything is different, but the same...  
things are more modern than before...  
bigger, and yet smaller... it's computers...  
San Dimas High School football rules!!



# DistSys is our Quantum Mechanics

Learn the homogenizing tools today



## Search results

### Portworx Volumes on Docker Practice Lab

BY PORTWORK

In this lab, we will learn more about mounting PX volumes on Docker

Start Scenario

### Portworx Volumes Practice Lab

BY PORTWORK

In this lab, we will get hands on with the creating, updating and inspecting PX volumes

Start Scenario

### Portworx Volumes Practice Lab

BY PORTWORK

In this lab, we will get hands on with the creating, updating and inspecting PX volumes

Start Scenario

### Portworx Storage Pool Practice Lab

BY PORTWORK

In this lab, we will get hands on with the Portworx Storage Pools

Start Scenario

### Portworx Shared Volumes Practice Lab

BY PORTWORK

In this lab, we will learn more about shared volumes

Start Scenario

### Portworx Replication Sets Practice Lab

BY PORTWORK

In this lab, we will learn more about Replication Sets

Start Scenario

### Portworx Lighthouse Practice Lab

BY PORTWORK

In this lab, we will get to explore the Portworx Lighthouse

Start Scenario

### Portworx Installation Docker Standalone

BY PORTWORK

In this lab, we will get hands on with the Portworx installation on a single node

Start Scenario

# More SNIA Resources

- [Kubernetes in the Cloud \(Part 1\)](#)
- [Kubernetes in the Cloud \(Part 2\)](#)
- [Kubernetes in the Cloud \(Part 3\) Stateful Workloads](#)
- [Stateful Kubernetes in the Cloud](#)

# After This Webcast

- Please rate this webcast and provide us with feedback
- This webcast and a copy of the slides will be available at the SNIA Educational Library <https://www.snia.org/educational-library>
- A Q&A from this webcast will be posted to the SNIA Cloud blog: [www.sniacloud.com/](http://www.sniacloud.com/)
- Follow us on Twitter @SNIACloud



Thank you!

# A Multi-tenant Multi-cluster Kubernetes Datapocalypse is Coming

Or, how I learned to stop worrying and love Federation

The END. Thank you!

# Links (1/4)

- A bank has a kubernetes outage <https://www.youtube.com/watch?v=OUYTNywPk-s>
- Jess Frazz multi-tenancy kubernetes research <https://blog.jessfraz.com/post/hard-multi-tenancy-in-kubernetes/>
- Multi-cloud environments SNIA webinar <https://www.brighttalk.com/webcast/663/407296>
- Kubernetes vulnerabilities like the thousand laughs attack <https://thenewstack.io/kubernetes-billion-laugh-vulnerability-is-no-laughing-matter/>
- Tesla has their kubernetes clusters hacked <https://arstechnica.com/information-technology/2018/02/tesla-cloud-resources-are-hacked-to-run-cryptocurrency-mining-malware/>
- Snapchat is one example of multi-cloud <https://www.vox.com/2017/2/9/14555202/snap-amazon-google-cloud-services-deal>
- Cloud providers like Google (Anthos) also do multi-cloud <https://cloud.google.com/blog/topics/hybrid-cloud/new-platform-for-managing-applications-in-todays-multi-cloud-world>
- There are limits to kubernetes, such as cloud quotas <https://cloud.google.com/kubernetes-engine/quotas>
- Real world architects suggest 500 nodes is a good limit per cluster <https://www.youtube.com/watch?v=Rb0JMVVZ1N8>
- Service meshes are one way to do multi-cluster <https://linkerd.io/2020/02/17/architecting-for-multicloud-kubernetes/>
- Snapchat built their own service mesh <https://eng.snap.com/monolith-to-multicloud-microservices-snap-service-mesh/>
- Kubefed is another option, but it's alpha <https://github.com/kubernetes-sigs/kubefed>

# Links (2/4)

- The UI/UX of multicluster is still being figured out <https://github.com/kubernetes-sigs/kubefed/blob/master/docs/userguide.md>
- Commercial kubernetes like Red Hat or VMware do multi-cloud, too <https://www.redhat.com/en/blog/open-sourcing-red-hat-advanced-cluster-management-kubernetes?source=bloglisting>  
and <https://tanzu.vmware.com/content/blog/harvest-the-benefits-of-a-multi-cluster-kubernetes-architecture-with-vmware-tanzu-mission-control>
- Even multiple clusters can fail, such as what happened to Target <https://tech.target.com/2019/01/14/cascading-failure-of-distributed-systems.html>
- Kubernetes is an in demand job skill <https://www.techrepublic.com/article/why-kubernetes-job-searches-grew-by-more-than-2000-in-4-years/>  
and <https://www.beseen.com/blog/talent/kubernetes-career-trends/>  
and <https://containerjournal.com/features/indeed-report-finds-kubernetes-job-market-hot/>
- Service mesh interface is worth following <https://smi-spec.io/>
- Project Octant is leading on visualizing multiple resources <https://octant.dev/>

# Links (3/4)

- Monzo Banks retrospective on their outage as an argument for homogeneity <https://community.monzo.com/t/resolved-current-account-payments-may-fail-major-outage-27-10-2017/26296/95>

and <https://github.com/linkerd/linkerd/issues/1219>

- The CNCF landscape is both an argument for and against heterogeneity <https://landscape.cncf.io/>

- Database migration with minimal downtime are always important <https://it.engr.ncsu.edu/2019/04/22/moving-databases-between-servers-with-minimal-downtime/>

- Everyone is very concerned about availability these days, just try rebooting your home router <https://twitter.com/gshipley/status/1272910652282191883>

- Simple storage options should be used until you feel the pain of scaling <https://cloud.google.com/storage>

and <https://azure.microsoft.com/en-us/free/storage/>

and <https://aws.amazon.com/s3/>

and <https://velero.io/>

and <https://min.io/>



# Links (4/4)

- Complex storage tools are the answer to scaling storage, because of queryability <https://rook.io/>  
and <https://portworx.com/>  
and <https://cloud.netapp.com/project-astra>  
and <https://openebs.io/>
- Katacoda is great for playing with options, with a free kubernetes sandbox <https://www.katacoda.com/courses/kubernetes/playground>

# Abstract

At the 2018 KubeCon keynote [1], Monzo showed us the cost of running a single massive cluster. A minor conflict between etcd and Java led to an outage during one of their busiest business days. The lesson most took from this was to start considering blast radius. Or, “Could one of our clusters go down, and our business would keep functioning?”

Dealing with blast radius means thinking in terms of multiple clusters. This multi-cluster ideal was further reinforced when Jessie Frazelle presented her research to the multitenancy working group. Her findings [2]? For ideal security and isolation, Kubernetes ought to run in a one team per cluster model.

How do we manage multiple clusters? Kubefed (federation) and multi-cluster service meshes are possible solutions, but still in their early days. Meanwhile, major cloud hosted kubernetes providers impose limits [3] like 110 pods per node, or 1,000 nodes per cluster. The multi-cluster Kubernetes reality is rapidly moving from “best practice” to “must have”.

In this talk we'll dive into more detail around the history of multi-cluster Kubernetes, and explore how multi-cluster setups could affect data heavy workloads (such as multiple microservices backed by independent data stores). Let's get prepared for the coming datapocalypse.