

# Building a fully cloud-native ATLAS Tier 2 on Kubernetes

## CA-VICTORIA-WESTGRID-T2 Site Report

Ryan Taylor

on behalf of UVic Research Computing Services



**University  
of Victoria**



# UVic site background

## Physical

## Virtual

- |   |   |                                    |                                      |
|---|---|------------------------------------|--------------------------------------|
| batch                                       | ● Bare metal batch cluster and WLCG ATLAS T2 commissioned | 2010                               | ● Cloud technology experimentation   |
|   | ○ Serial & parallel partitions, GbE & IB networks         | 2011                               | ○ Nimbus, OpenNebula, Oracle cloud   |
|   | ● Serial cluster expansion                                | 2012                               | ● Nimbus cloud deployment (Synnefo)  |
| <hr style="border-top: 1px dashed black;"/> |   |                                    |                                      |
| cloud                                       | ● Cloud funding, dedicated hardware                       | 2014                               | ● Virtualized batch cluster in cloud |
|   | ○ first national cloud service offering                   | 2016                               | ● OpenStack cloud (Nephos/West)      |
|   | ○ major national cloud site                               |                                    | ● National cloud site (Arbutus)      |
|   | ● Cloud hardware expansion                                | 2018                               | ● Kubernetes experimentation         |
|   | ● Cloud hardware expansion                                | 2019                               | ● CA-VICTORIA-K8S-T2 in production   |
|   | 2020  | ○ gaining k8s experience for ATLAS |                                      |
|   | 2021  | ● T2 compute entirely k8s-native   |                                      |
|   |   | ○ CREAM CE decommissioning         |                                      |



# Arbutus

General-purpose scientific cloud at UVic



ARBUSCLOUD

**OPEN FOR RESEARCH**

vCPUs vGPUs MASSIVE STORAGE

<https://computeCanada.ca/research-portal/national-services/compute-canada-cloud>

compute Canada | calcul Canada | University of Victoria

- ~ 3000 users
- ~ 1000 projects



- 44,000 vCPUs
- 160 TB RAM
- 416 vGPUs



- 18 PB (usable) Ceph
- RBD, S3/object, CephFS

<b>Horizon Dashboard</b>	
<b>Openstack CLI</b>	
<b>Openstack API</b>	
Storage Virtualization <b>Ceph</b>	Storage API <b>Cinder/Glance</b>
Network Virtualization <b>Neutron</b>	Network API <b>Neutron Server</b>
Compute Virtualization <b>QEMU / KVM</b>	Compute API <b>Nova</b>

# T2 Computing



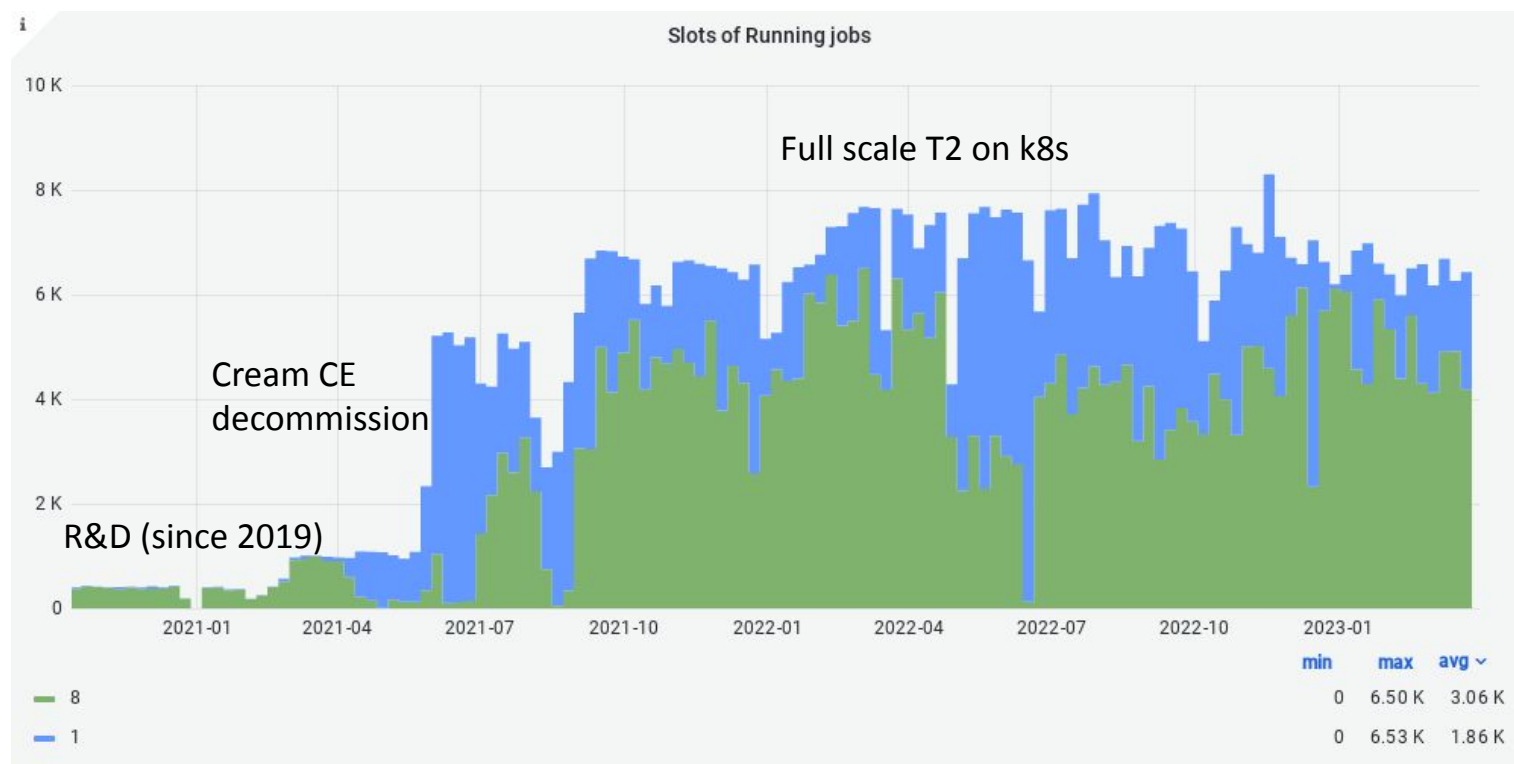
## Using Kubernetes as an ATLAS computing site

*Fernando Barreiro Megino, Jeffrey Ryan Albert, Frank Berghaus, Danika MacDonell, Tadahshi Maeno, Ricardo Brito Da Rocha, Rolf Seuster, Ryan P. Taylor, Ming-Jyuan Yang on behalf of the ATLAS experiment  
CHEP 2019, Adelaide, Australia*



[CHEP 2019 presentation](#)

[CHEP 2023 presentation](#)



CA-VICTORIA-WESTGRID-T2 uses Kubernetes for container-native batch computing. Harvester submits ATLAS grid jobs to k8s API, which runs them as pods. No traditional batch system or Compute Element.

# The eventual goal: a fully k8s-native T2


## Installable with Helm



- Helm: application manager for Kubernetes
  - One command to install/upgrade everything
  - Comprehensive configuration via one YAML file
- **helm install T2Site**
  - (K)APEL accounting done
  - frontier-squid done
  - compute done (static YAML)
  - EOS SE in progress
  - CVMFS-CSI optional
  - ~~Compute Element~~ built-in
  - ~~Batch system~~ built-in

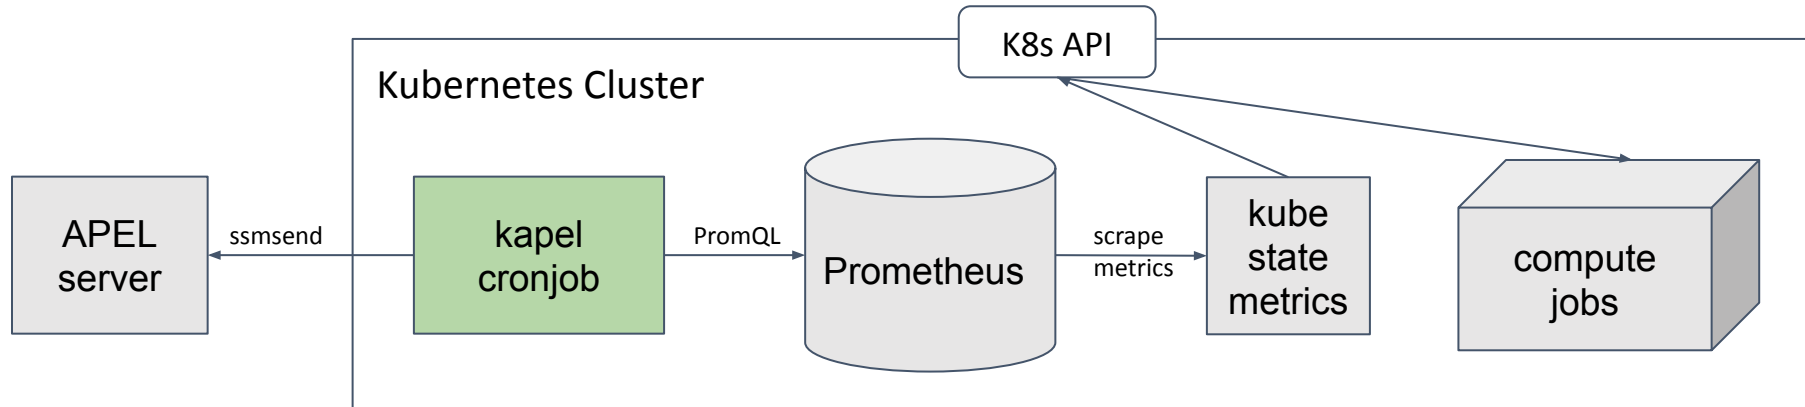


# Kubernetes deployment approach

- Infrastructure as Code:  **KUBESPRAY** (Ansible, Terraform)
- Openstack-related improvements pushed upstream
  - [predefine master floating IPs](#)
  - [master volume type support](#)
  - [separate Availability Zones for kubelets](#)
  - [extra groups for k8s nodes](#) etc...
- UVic custom extensions for multi-cluster management, integration of additional Ansible roles, Helm charts, etc.
- Identical dev and prod clusters for testing changes
  - Complete with test queue CA-VICTORIA-K8S-TEST-T2
- Can destroy and rebuild the cluster from scratch (git) in ~ hours
- Currently using Almalinux 8 VMs
  - Upgrade (rebuild) to new EL version typically quite easy

# KAPEL

## Container-native APEL accounting for Kubernetes

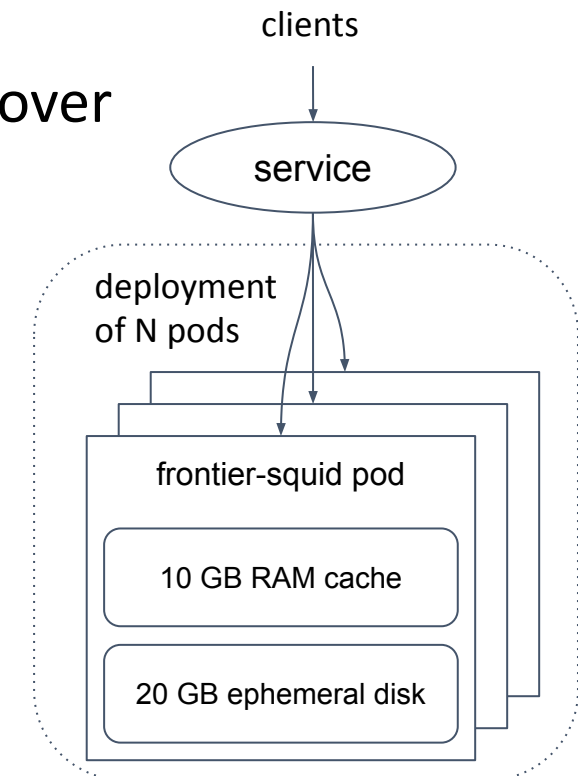


- Standard k8s add-ons do most of the work
  - k8s cron job instead of APEL node
  - Prometheus instead of MySQL DB for data collection and storage
  - PromQL for data querying, analytics
  - kube-state-metrics (KSM) instead of batch log parser
  - Only needed to write ~200 lines of python (and some YAML)
- Available as Helm chart: <https://github.com/rptaylor/kapel>

# Frontier-squid

## Deployed on Kubernetes

- Using frontier-squid [Helm chart](#) from CERN ScienceBox
  - Simple, lightweight, container-native approach
  - Trivial to scale, with automatic load-balancing and failover
- UVic contributed enhancements
  - Run as unprivileged squid user [#61](#)
  - Allow configuration of service details [#63](#)
  - Support for priorityClass and pod resource requests/limits [#64](#)
  - Send access logs to stdout [#69](#)
  - Configurable ACL activation [#72](#)
  - Harmonize configuration with upstream package [#73](#)
  - Add backup readiness probe URL for redundancy [#74](#)
  - Update ACLs for Frontier servers [#78](#)
  - Expand list of safe ports [#81](#)
- Suitable for new CVMFS proxy sharding feature





# EOS SE on k8s with CephFS



- Physical consolidation: all storage on Ceph
- Logical consolidation: services on k8s
- EOS can be installed on k8s via Helm chart
  - reproducible, single step deployment
  - easier to manage and maintain
  - easy to set up another instance, e.g. for dev
- Opportunity: [direct data access for jobs](#) on CephFS

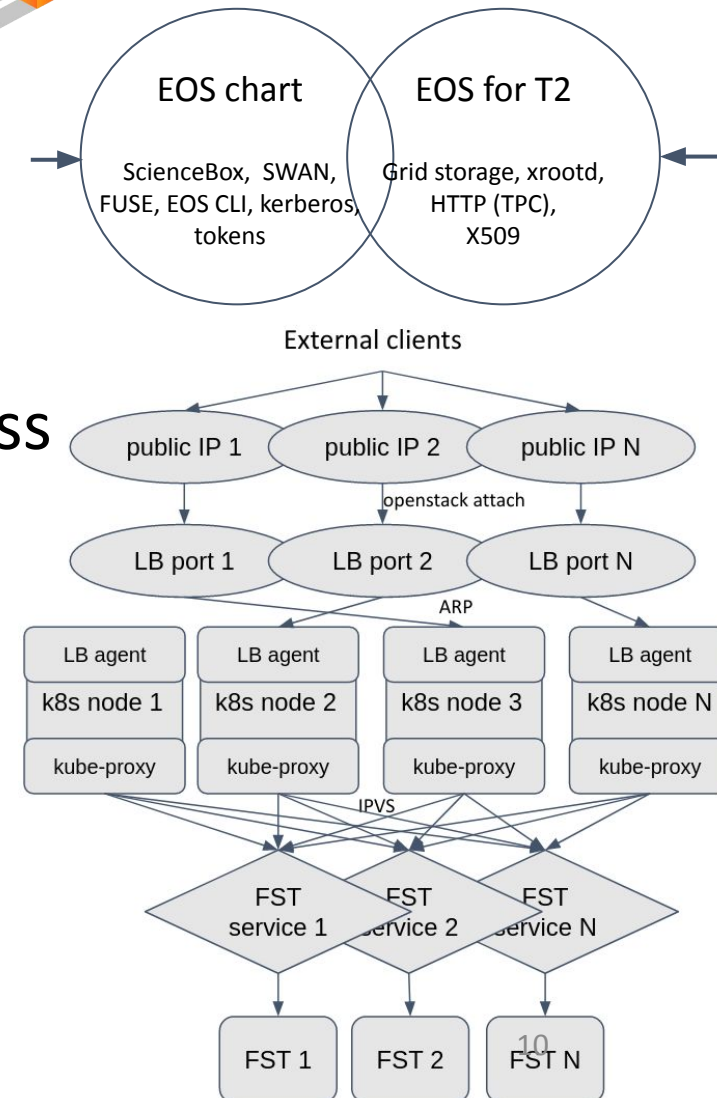
Also interesting:

[Deploying dCache in Kubernetes](#)

# EOS SE on k8s with CephFS



- Enhancements of Helm chart for T2 use case
  - VOMS authz/authn
  - Set up host certs, CAs in pods, run fetch-crl, etc.
- Kubernetes network architecture for external access
  - A LoadBalancer Service for each storage pod (FST)
- Same design should work for dCache on k8s too



# Summary

- CA-VICTORIA-WESTGRID-T2 running ~8K cores of ATLAS compute jobs on Kubernetes
- APEL accounting and Frontier-squid also deployed on Kubernetes
- Adaptation of EOS Helm chart and testing EOS SE deployment
- Enable streamlined replicable deployment of a full ATLAS T2

**From opening the queue to  
production in less than 2 weeks.**

[NET2: a first example of OpenShift/OKD for Tier 2 provisioning and cluster management in US ATLAS](#)

# What does cloud-native computing really mean?

or

## Why Kubernetes?

- We are a cloud site
- We need to run complex distributed applications at scale in a robust way
- Cloud + k8s provides:
  - Flexible & dynamic infrastructure
  - Resilience and automated remediation
  - Rapid application deployment
  - Application lifecycle management
  - Horizontal scalability



VMs as pets

Openstack



VMs as cattle

Openstack + ???



containers as  
cattle

Openstack + k8s

# CVMFS proxy sharding with k8s Squids

- New feature in CVMFS v2.10 to improve cache hit rates
- CVMFS understands round-robin DNS
  - dereferences multiple A records
- Solution using k8s Services: [headless ClusterIP](#)

```
service:
```

```
  clusterIP: None
```

- Should decrease CVMFS\_DNS\_MIN\_TTL to a small value
  - CVMFS default is 1 min
  - K8s deployment upgrade could be < 1 min (and DNS TTL is 5 s)
  - Details: [#97](#)

# Ingress and LBaaS

- Initial basic approach used keepalived and nginx-ingress to receive traffic from outside world into clusters
- Migrated to PureLB and Traefik
  - More maintainable/manageable, via Helm charts
  - Cohesive access to dashboards etc across all clusters
- PureLB: like MetalLB but simpler, lightweight
  - relies on Linux network stack of host
  - Programmable (LB -> LBaaS)
- Traefik Ingress controller
  - Widely used, full featured, nice web UI, CRDs
  - Better TCP and UDP support



- CephFS bug encountered: [55090](#)
  - Ceph fixes: [#46902](#) [#46905](#)



# Cloud Architecture

Generously-provisioned Control Plane - 10 nodes @ 56 cores / 256GB RAM

10G/25G Ethernet Backplane

Compute nodes aggregated by capability:

- High Memory (512GB/node, 1.5TB/node Phase 3 Optane)
- High Performance SSD Local Storage
- vGPU availability

Failure Domain-separated Availability Zones for Persistent Workloads