ENABLING DATA SCIENTISTS... DURING COVID

DATA ANALYTICS AS A SERVICE USING CNCF TECHNOLOGIES

THE TEAM

- A collaboration of teams at Statistics Canada
  - Cloud Services Enablement Project (CSEP)
  - Data Analytics as a Service (DAaaS)
  - Data Science Division (DSD)
  - Digital Innovation

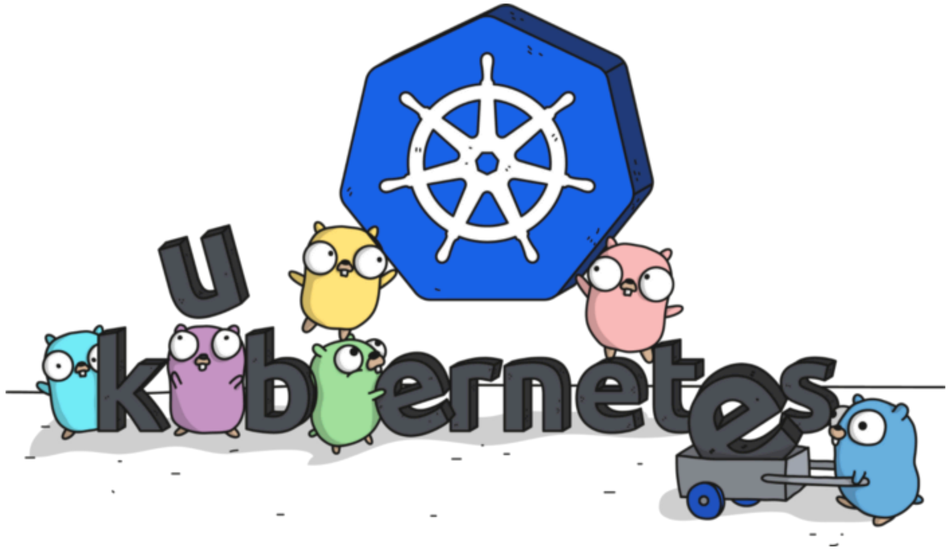

ARTWORK BY @ASHLEYMCNAMARA

- Close collaboration behind the technical teams building the system and the users of the system. Allowed for quick development of core system tools.

COMMUNITY

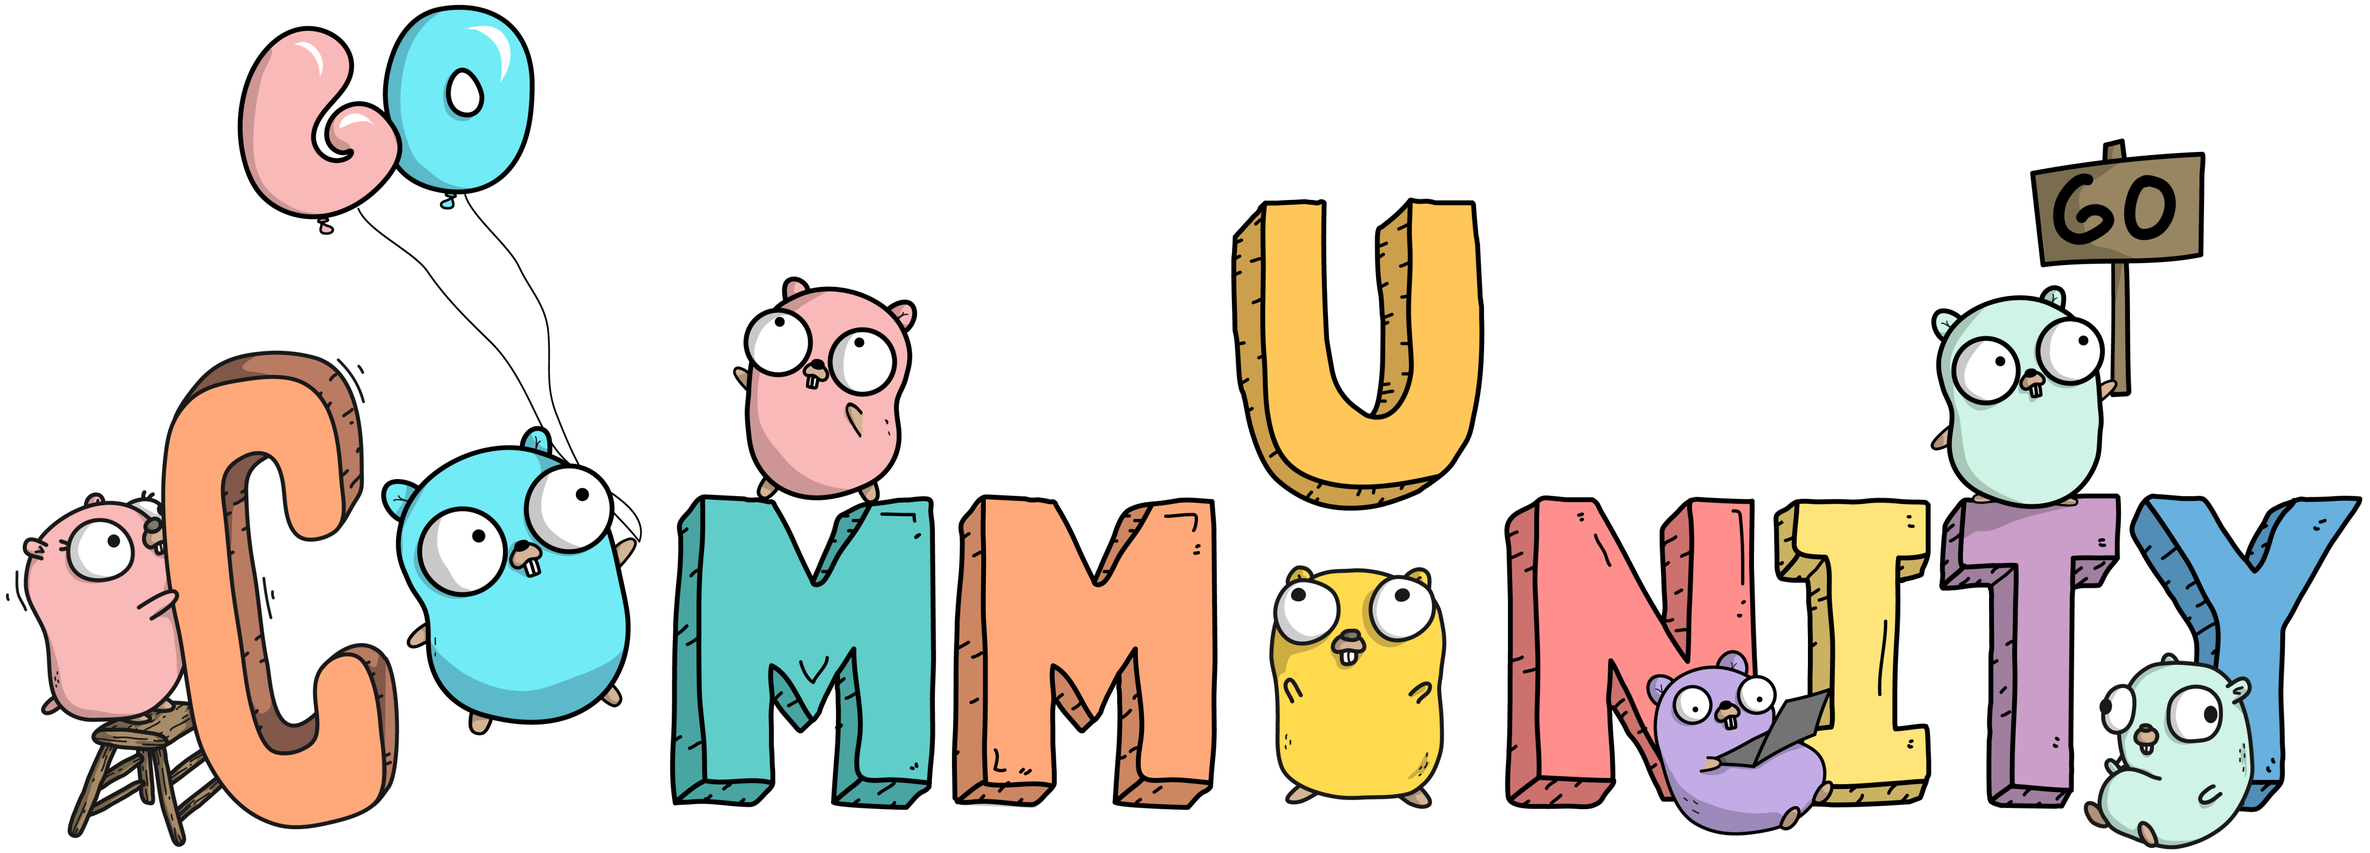

ARTWORK BY @ASHLEYMCNAMARA

THE CHALLENGE

- Enable Data Scientists (both internal and external) to perform data and analysis in a public cloud environment quickly
- Security as a foundational principle
- Utilize components already constructed/investigated at Statistics Canada
  - CSEP: Kubernetes platform (Documentation, Presentation)
  - DAaaS: Analytics tools (Kubeflow, Jupyter notebooks, Seldon, Shiny, …)

ARTWORK BY @ASHLEYMCNAMARA

AN OVERVIEW

- All tools and processes discussed have been Open Sourced on the official Statistics Canada GitHub organization (https://github.com/StatCan?q=daaas)
- Development occurs in an Agile fashion on GitHub:
  - 1 to 2 week sprints
  - Daily standup
  - Weekly Retrospective
- We presented this platform to the Chief Statistician and the Assistant Chief Statisticians, who alongside with the Data Scientists who use it daily were impressed and expressed their interest in seeing this platform grow across government

ARTWORK BY @ASHLEYMCNAMARA

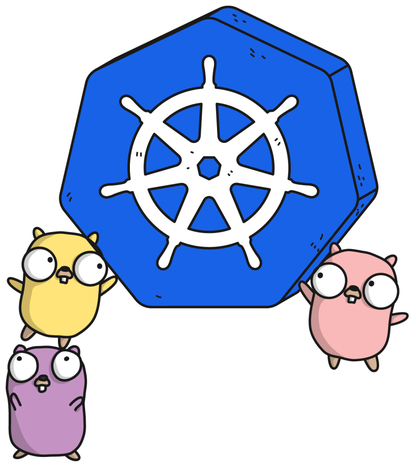

ARTWORK BY @ASHLEYMCNAMARA

WHAT'S BEHIND THE SERVICE?

TECHNOLOGY STACK


kubernetes


Open Policy Agent
Istio
envoy
HashiCorp
Vault
Security

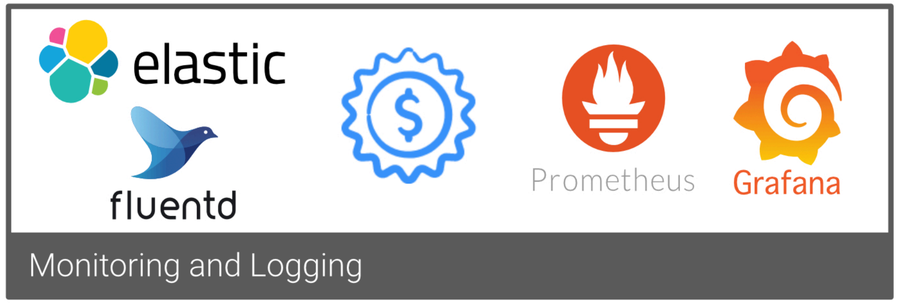
elastic
fluentd
Prometheus
Grafana
Monitoring and Logging

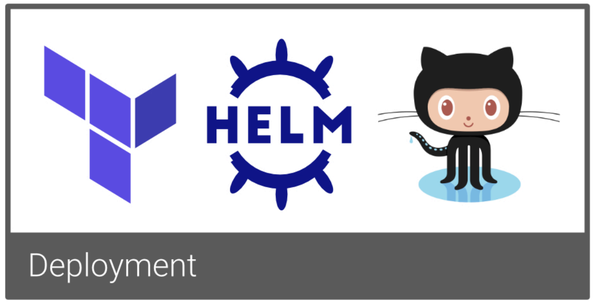
HELM
Deployment


Kubeflow
mlflow
jupyter
Shiny
Data Science


docker
Packaging


Platform as a Service

KUBERNETES

- Use the managed Azure Kubernetes Service (AKS), reducing cluster management overhead
  - Kubernetes is a desired-state configuration system, which integrates well with our automation processes
- Support for user isolation via namespaces
- Kubernetes is a platform to build platforms
  - Extensible via the use of custom controllers and custom resource definitions (CRDs)
    - Profiles controller extension (configure user profiles with credentials, configuration options, etc.)
    - GPU toleration injector (automatically configure tolerations for GPU workloads)
    - Goofys injector (automatically configure notebooks with MinIO storage configuration)
- Support for a wide-range of languages (if it compiles, it’s supported!)
  - We’ve used: Python, Go, Node.JS, TypeScript.JS, .NET Core, Java, Julia and R

AUTOMATION

- Infrastructure is:
  - Deployed using Infrastructure as Code (Terraform and Helm)
    - Pull requests provide a review process, as the CI/CD system outputs a GitHub comment on the PR with the planned change
  - Deployed using a CI/CD process (GitHub Actions)
    - Code linting
    - Container scanning (identify CVEs in container images) [Trivy]
    - Build and pushed to the Azure Container Registry (ACR)
    - Deployed to the Kubernetes environment
- CI/CD processes with GitHub Actions extend into the technology stack:
  - Jupyter Notebook containers running under Kubeflow
  - “self-service” (but gated) deployment of R-Shiny dashboards
  - Web portal / Other Data Science related workloads

DID YOU KNOW?

We tested rebuilding the entire environment! We rebuilt it in a day using Terraform, and migrated over user data using VMWARE VELERO!

SECURITY

- Azure Active Directory integration
  - Kubernetes RBAC, Vault, Kubeflow, Portal and other web-based services (OpenID Connect)
- Gatekeeper (from the Open Policy Agent) provides an enforcement of a range of policies:
  - Authorized container sources, container permissions, volumes, no public load balancers), and more! (GateKeeper Policies)
- All inbound traffic is TLS protected (Let's Encrypt certs automatically issued by Cert Manager)
  - Mutual TLS authentication for most inter-service communication (Istio)

SECURITY

- Full audit logging provided by the Kubernetes control plane
- All cluster and application logs are aggregated into an Elasticsearch instance
- Daily automated cluster assessment of CIS benchmarks and best practices (Starboard)
- Dynamic credential generation and auditing provided by Hashicorp Vault
- Internal version of the platform has received an interim Authority to Operate (IATO) and will in weeks be given an ATO (ATO)

COST CONTROL

- Use of reserved instances to reduce the cost of base compute needs (~33%/year)
- Support for auto-scaling of compute infrastructure (scale up and down)
- Limit workloads assigned to GPU nodes to those who require GPU
- Breakdown of resource utilization per-user
  - Available through an API for future billing integration (cost-recovery)
  - Alert users when they reach a certain cost threshold

k8s-cancentral-02-covid-aks ov

kubecost.covid.cloud.statcan.ca/overview.html

AKS

- Overview
- Cost Allocation
- Savings
- Health
- Reports
- Notifications

Switch cluster
k8s-cancentral-02-covi...

Settings

Overview / k8s-cancentral-02-covid-aks

Monthly savings of CA$6,235.94 identified

LEARN MORE >

Monthly cost
CA$16,873.86

Cost Efficiency
45.5%

Monthly cluster costs
Monthly run rate expenses based on resource prices

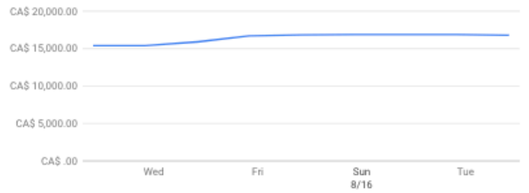

Total cost

CLUSTER METRICS >

Resource Efficiency
Based on currently provisioned resources and last 24h utilization

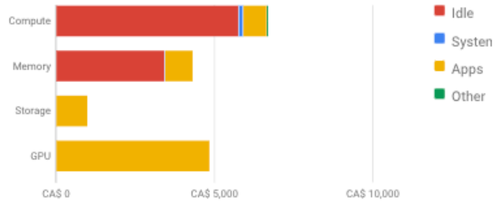

REAL-TIME ASSETS >

Deployment Allocation

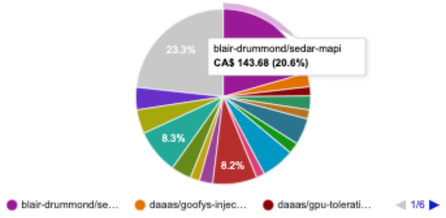

Product Allocation

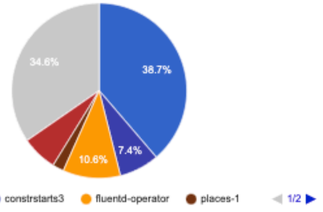

COST ALLOCATION >

Namespace	Monthly Cost	Efficiency
christian-ritter	CA$1,669.65	22
daaas	CA$1,443.83	19
monitoring	CA$1,262.60	100
najeeb-gazi	CA$848.46	20
spencer-karau	CA$786.66	4
minio	CA$596.63	100

NAMESPACE ALLOCATION >

Infrastructure health

CLUSTER HEALTH >

Have questions? We're on Slack or email at team@kubecost.com

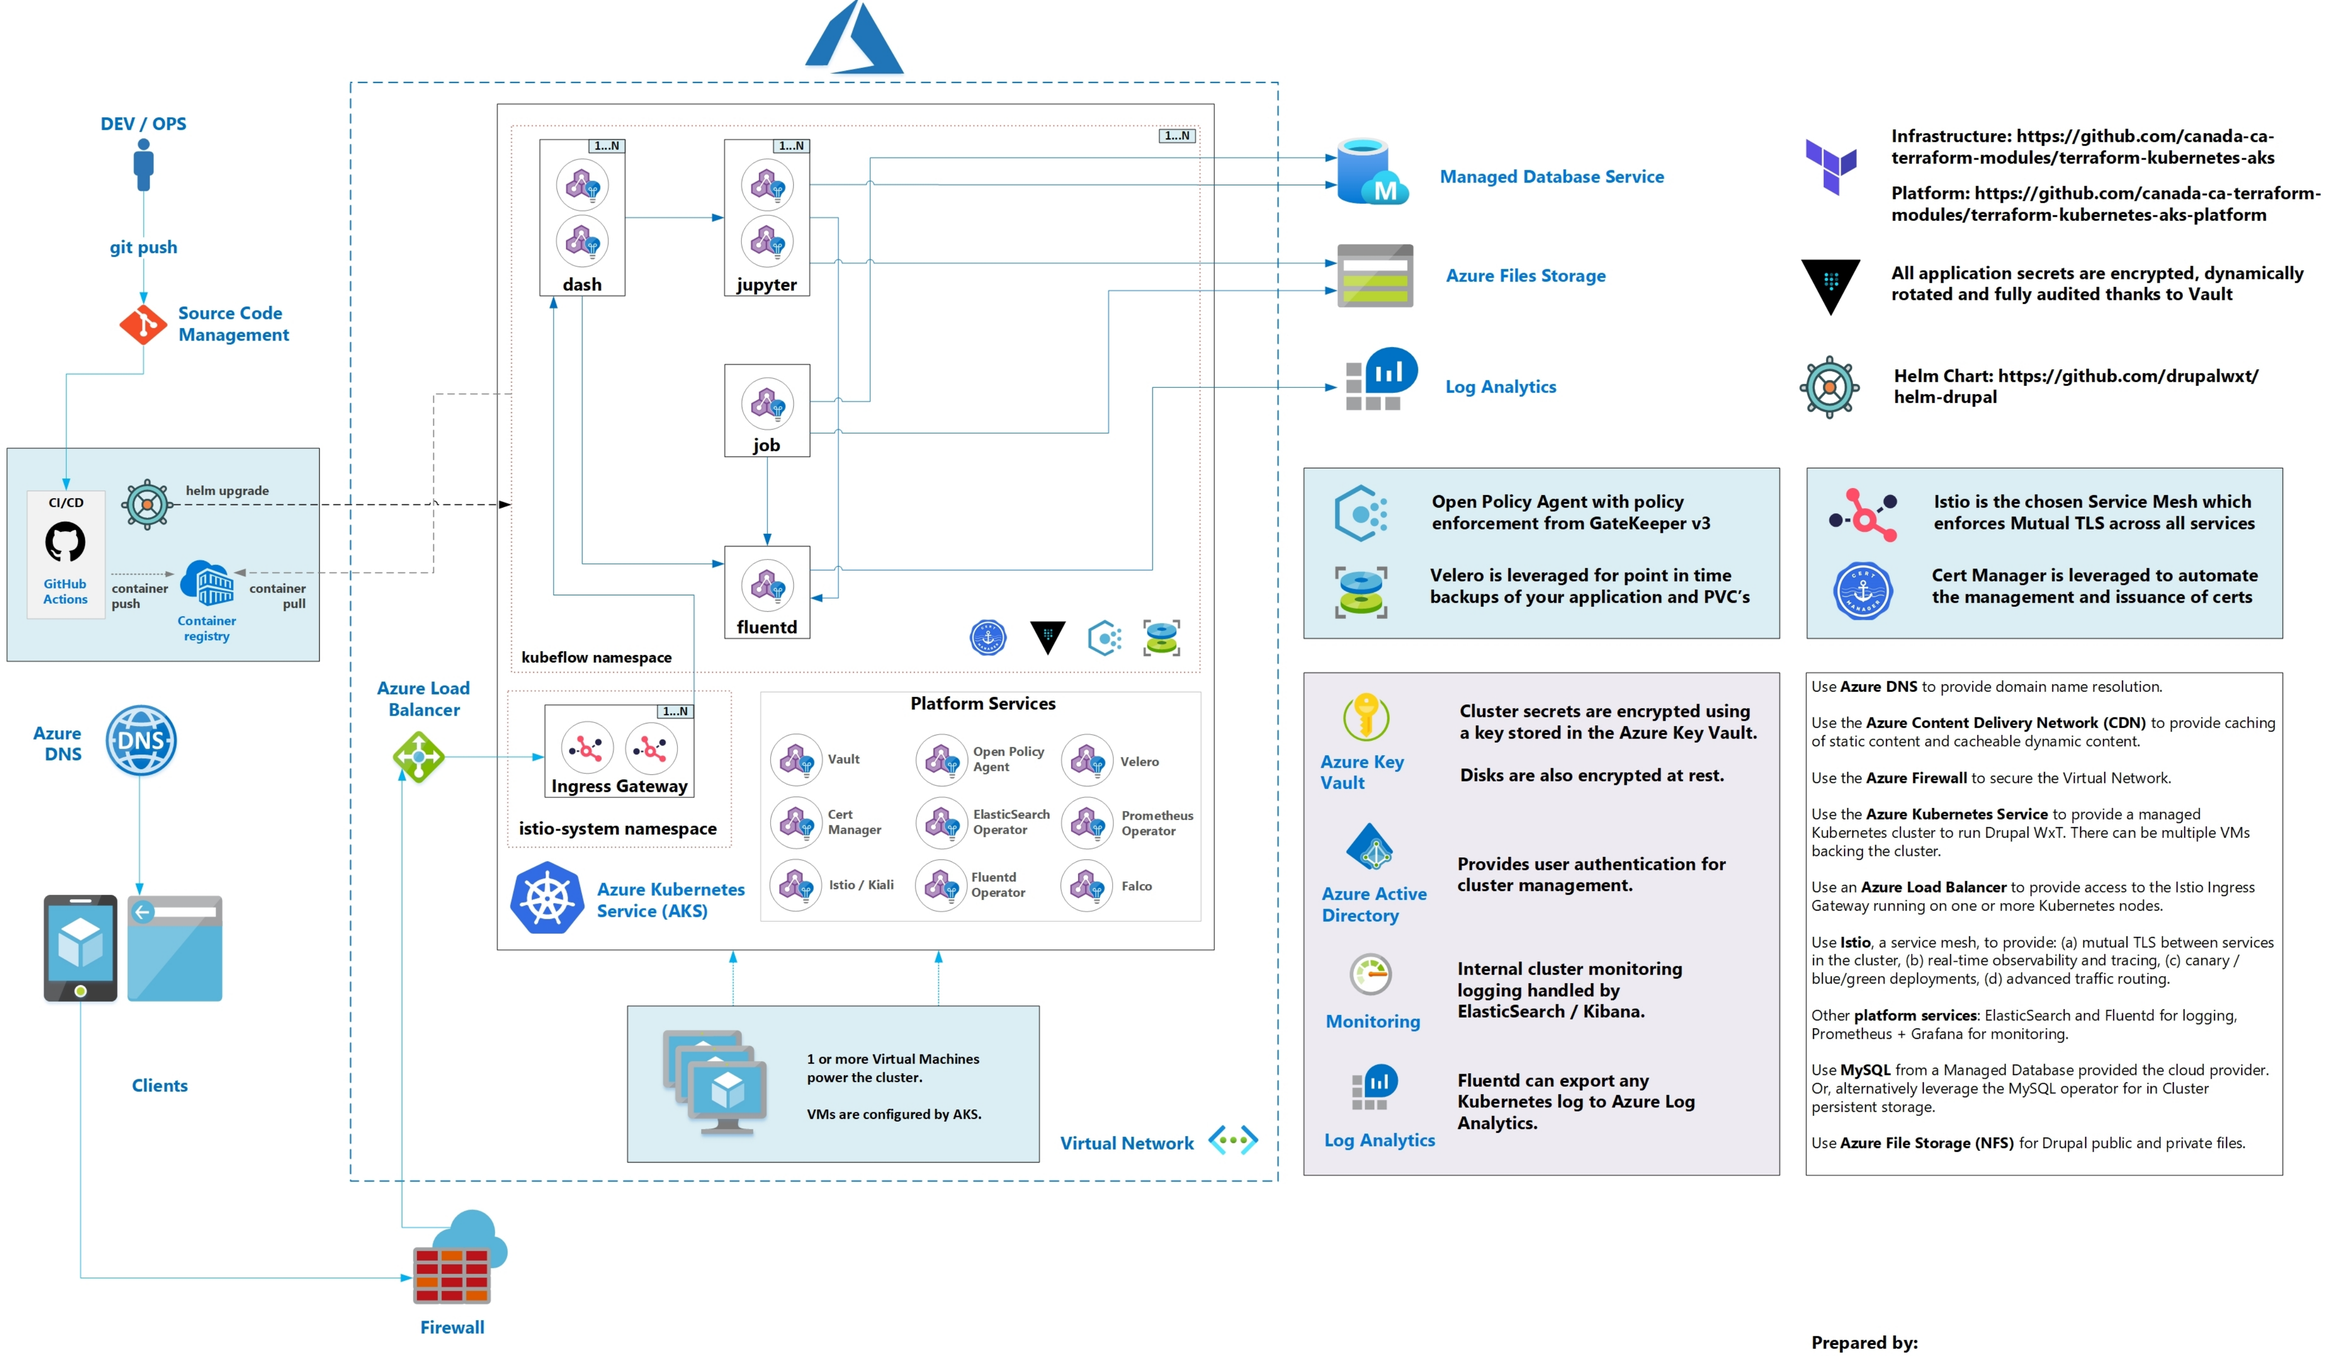

Infrastructure: https://github.com/canada-ca-terraform-modules/terraform-kubernetes-aks

Platform: https://github.com/canada-ca-terraform-modules/terraform-kubernetes-aks-platform

All application secrets are encrypted, dynamically rotated and fully audited thanks to Vault

Helm Chart: https://github.com/drupalwxt/helm-drupal

Open Policy Agent with policy enforcement from GateKeeper v3

Velero is leveraged for point in time backups of your application and PVC's

Istio is the chosen Service Mesh which enforces Mutual TLS across all services

Cert Manager is leveraged to automate the management and issuance of certs

Azure Key Vault

Cluster secrets are encrypted using a key stored in the Azure Key Vault.

Disks are also encrypted at rest.

Azure Active Directory

Provides user authentication for cluster management.

Monitoring

Internal cluster monitoring logging handled by ElasticSearch / Kibana.

Log Analytics

Fluentd can export any Kubernetes log to Azure Log Analytics.

Use **Azure DNS** to provide domain name resolution.

Use the **Azure Content Delivery Network (CDN)** to provide caching of static content and cacheable dynamic content.

Use the **Azure Firewall** to secure the Virtual Network.

Use the **Azure Kubernetes Service** to provide a managed Kubernetes cluster to run Drupal WxT. There can be multiple VMs backing the cluster.

Use an **Azure Load Balancer** to provide access to the Istio Ingress Gateway running on one or more Kubernetes nodes.

Use **Istio**, a service mesh, to provide: (a) mutual TLS between services in the cluster, (b) real-time observability and tracing, (c) canary / blue/green deployments, (d) advanced traffic routing.

Other **platform services**: ElasticSearch and Fluentd for logging, Prometheus + Grafana for monitoring.

Use **MySQL** from a Managed Database provided the cloud provider. Or, alternatively leverage the MySQL operator for in Cluster persistent storage.

Use **Azure File Storage (NFS)** for Drupal public and private files.

Azure Reference Architecture

Prepared by:

William Hearn (william.hearn@canada.ca)
Zachary Seguin (zachary.seguin@canada.ca)

Statistics Canada DevSecOps Platform Architecture*

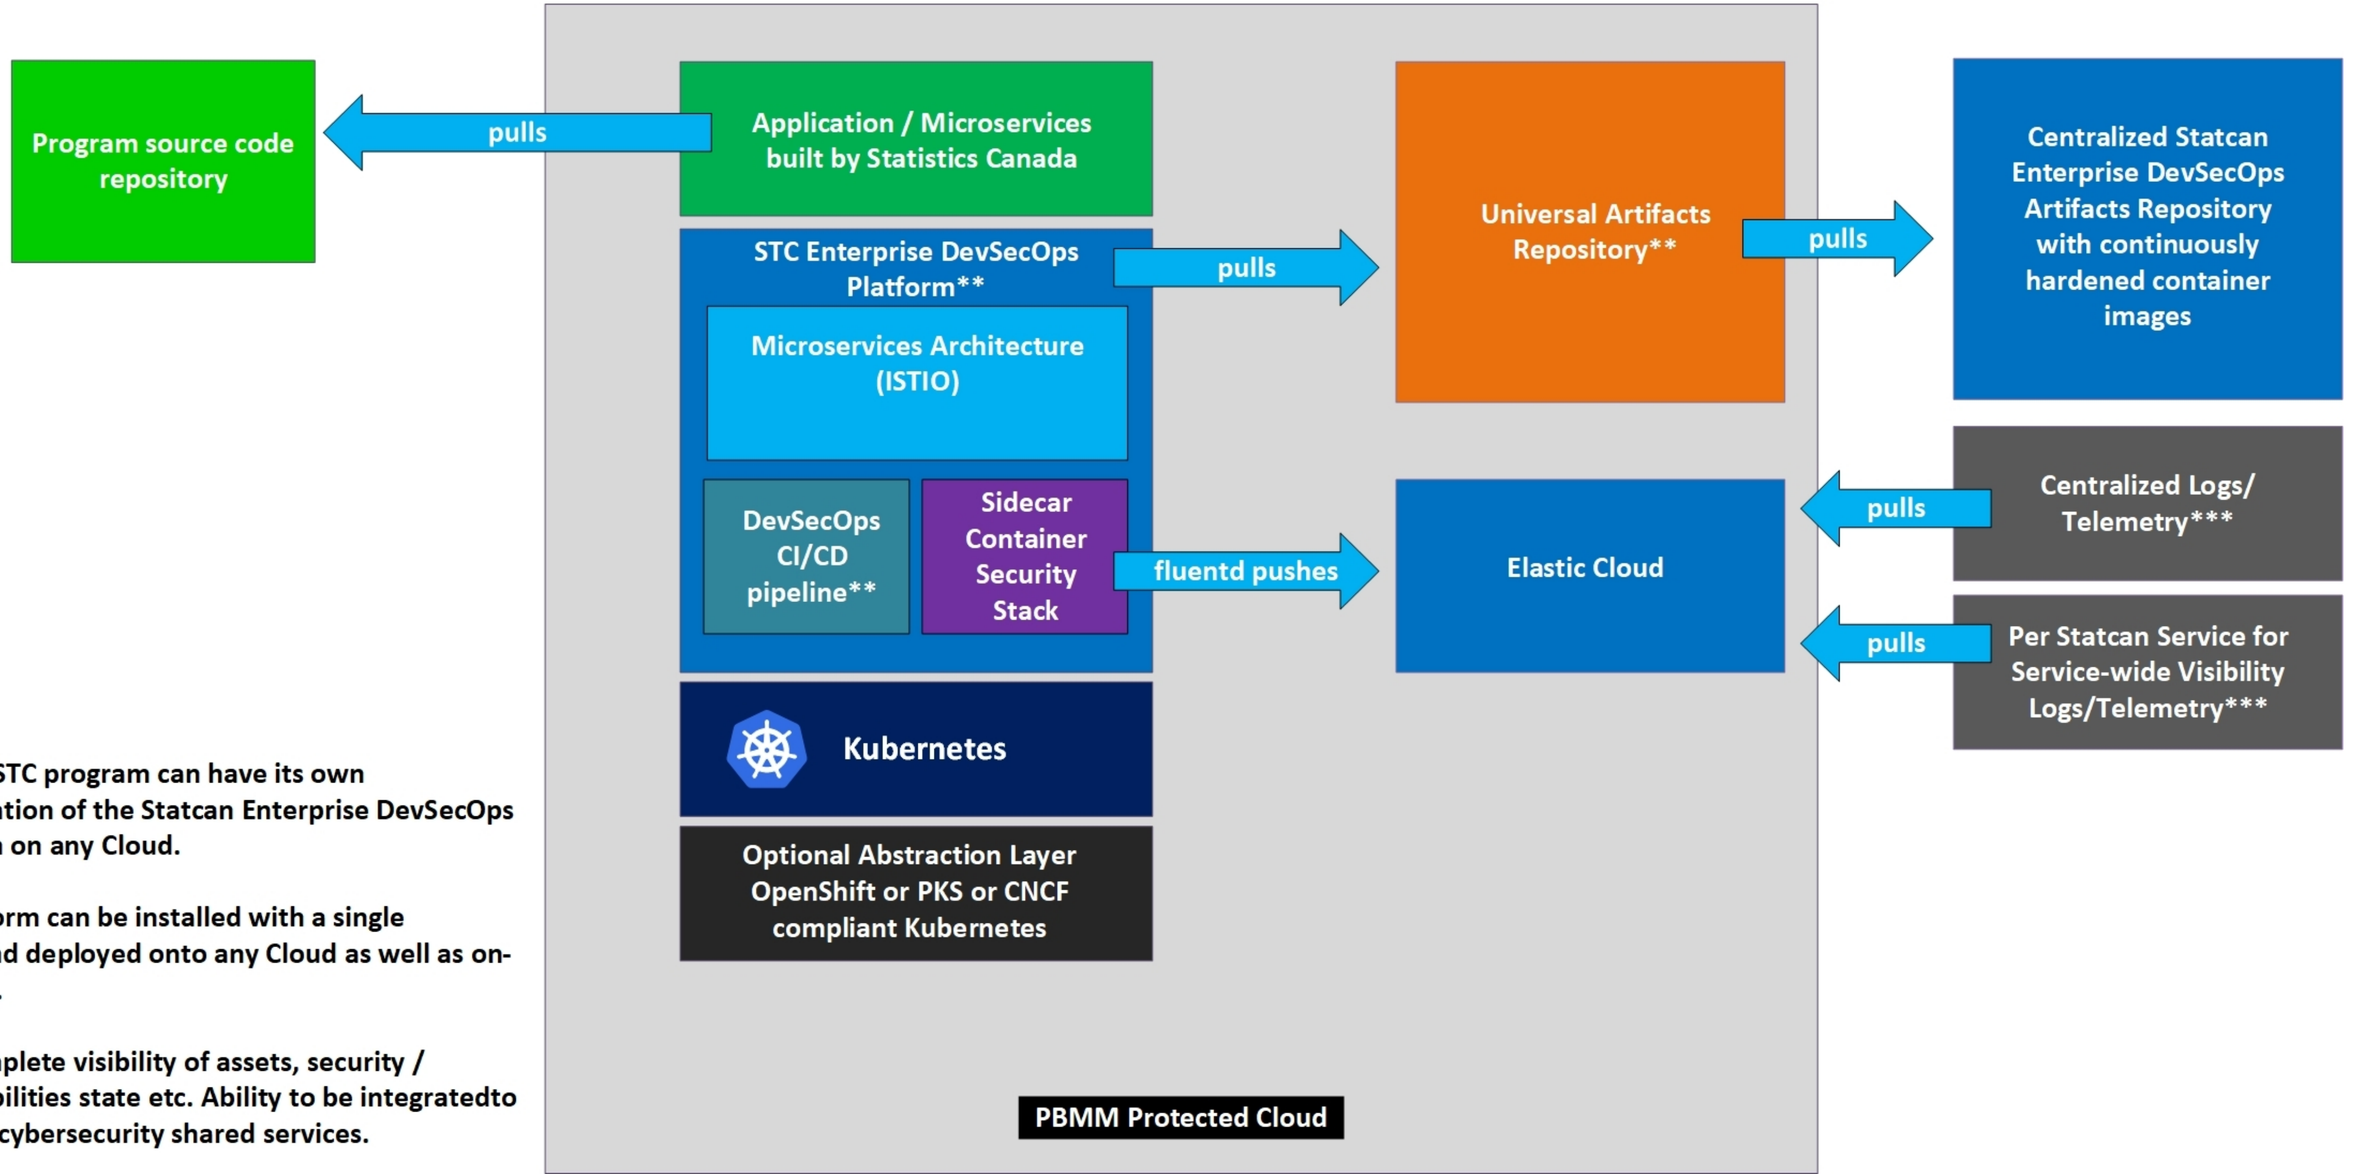

* Every STC program can have its own instantiation of the Statcan Enterprise DevSecOps Platform on any Cloud.

** Platform can be installed with a single command deployed onto any Cloud as well as on-premise.

*** Complete visibility of assets, security / vulnerabilities state etc. Ability to be integratedto existing cybersecurity shared services.

github-actions bot commented 19 days ago

`terraform plan` Success

▼ Show Output

```
An execution plan has been generated and is shown below.
Resource actions are indicated with the following symbols:
  + create
-/+ destroy and then create replacement

Terraform will perform the following actions:

  # local_file.cert_wildcard will be created
  + resource "local_file" "cert_wildcard" {
      + content  = <<~EOT
            apiVersion: cert-manager.io/v1alpha2
            kind: Certificate
            metadata:
              name: wildcard
              namespace: istio-system
              labels:
                use-azuredns-solver: 'true'
            spec:
              secretName: wildcard-tls
              commonName: "*.covid.cloud.statcan.ca"
              dnsNames:
              - "*.covid.cloud.statcan.ca"
              issuerRef:
                name: letsencrypt
                kind: ClusterIssuer
        EOT
      + filename = "./generated/wildcard.yaml"
      + id       = (known after apply)
    }

  # local_file.eck_daaas will be created
  + resource "local_file" "eck_daaas" {
      + content  = <<~EOT
            apiVersion: elasticsearch.k8s.elastic.co/v1
            kind: Elasticsearch
            metadata:
              name: daaas
              namespace: daaas
            spec:
              version: 7.8.1
              nodeSets:
              - config:
                  node.data: true
                  node.ingest: true
                  node.master: true
                count: 2
                name: nodes
                podTemplate:
                  metadata:
                    annotations:
```

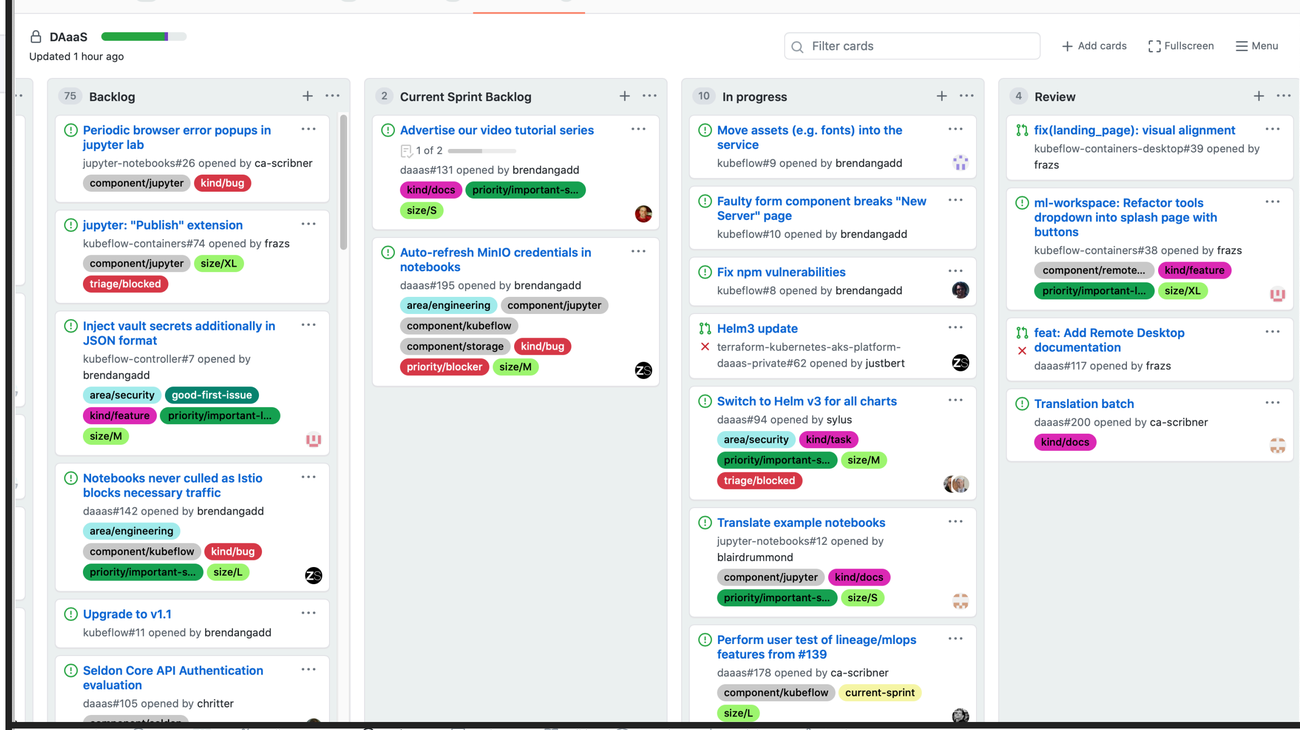

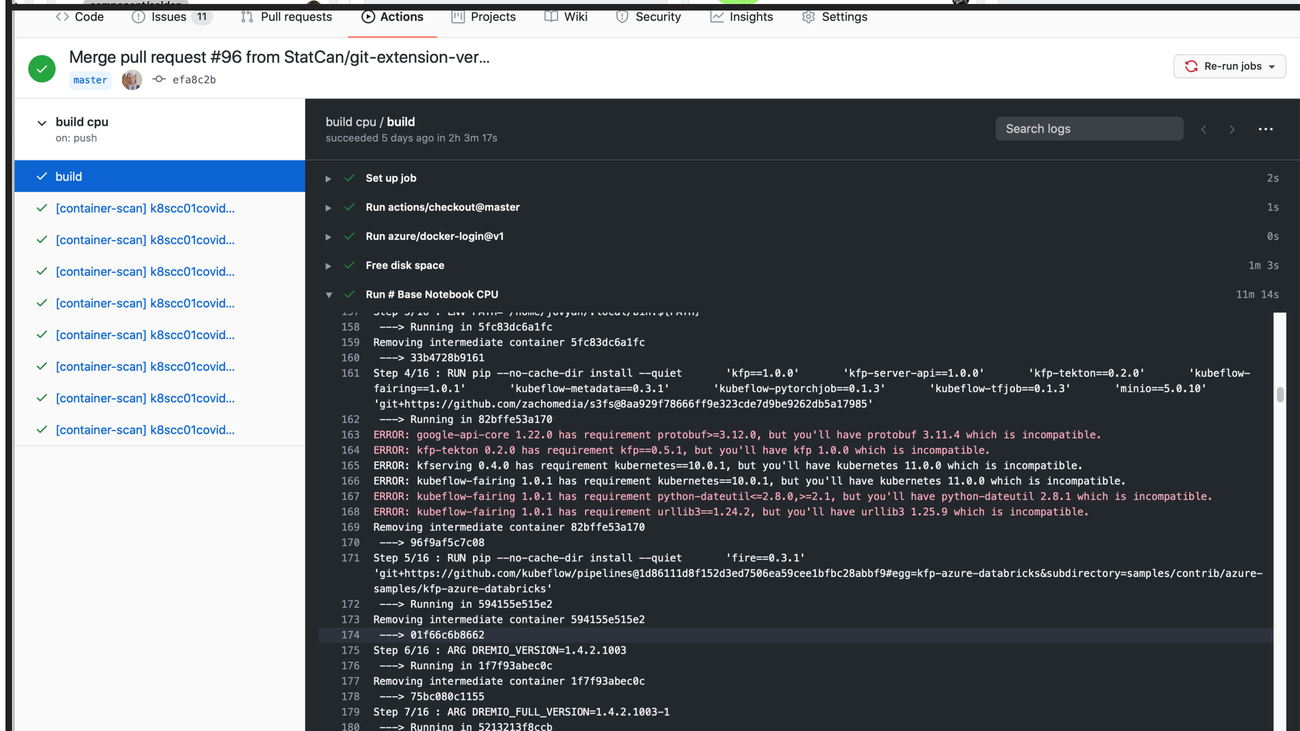

ARTWORK BY @ASHLEYMCNAMARA

WHAT'S BUILT ONTO OF THE SERVICE?

ANALYTICS STACK

KUBEFLOW

- Kubeflow is the current entrypoint to our platform
- Self-service enrolment and management
  - Profile (and Kubernetes namespace) is provisioned on first login
  - Manage your namespace contributors
- A machine learning toolkit
  - Provides pipelines, experiments, notebooks, and more!
- Usable in any cloud environment

JUPYTER NOTEBOOKS (PART OF KUBEFLOW)

- A variety of notebook containers available to users
  - Minimal, Geomatics, Machine Learning and R Studio [CPU + GPU]
  - "Desktop"-based images for Geomatics, R
- Notebook containers are maintained using a public GitHub repository
- Users have the ability to add software to their notebook instances ("self-service")
- Support for various storage mechanisms: Raw disk (PVC), S3-compatible mount, Azure Files, Azure Data Lake, Azure Blob Storage

PLUGGABLE ANALYTICS COMPONENTS (MLOPS)

- Kubeflow (and Kubernetes) provide an analytics orchestration system
  - All analytics components are pluggable, allowing for integration with any open source or Platform as a Service (PaaS) offering
    - TensorFlow
    - MinIO
    - MLFlow
    - Azure Machine Learning
    - Azure Databricks

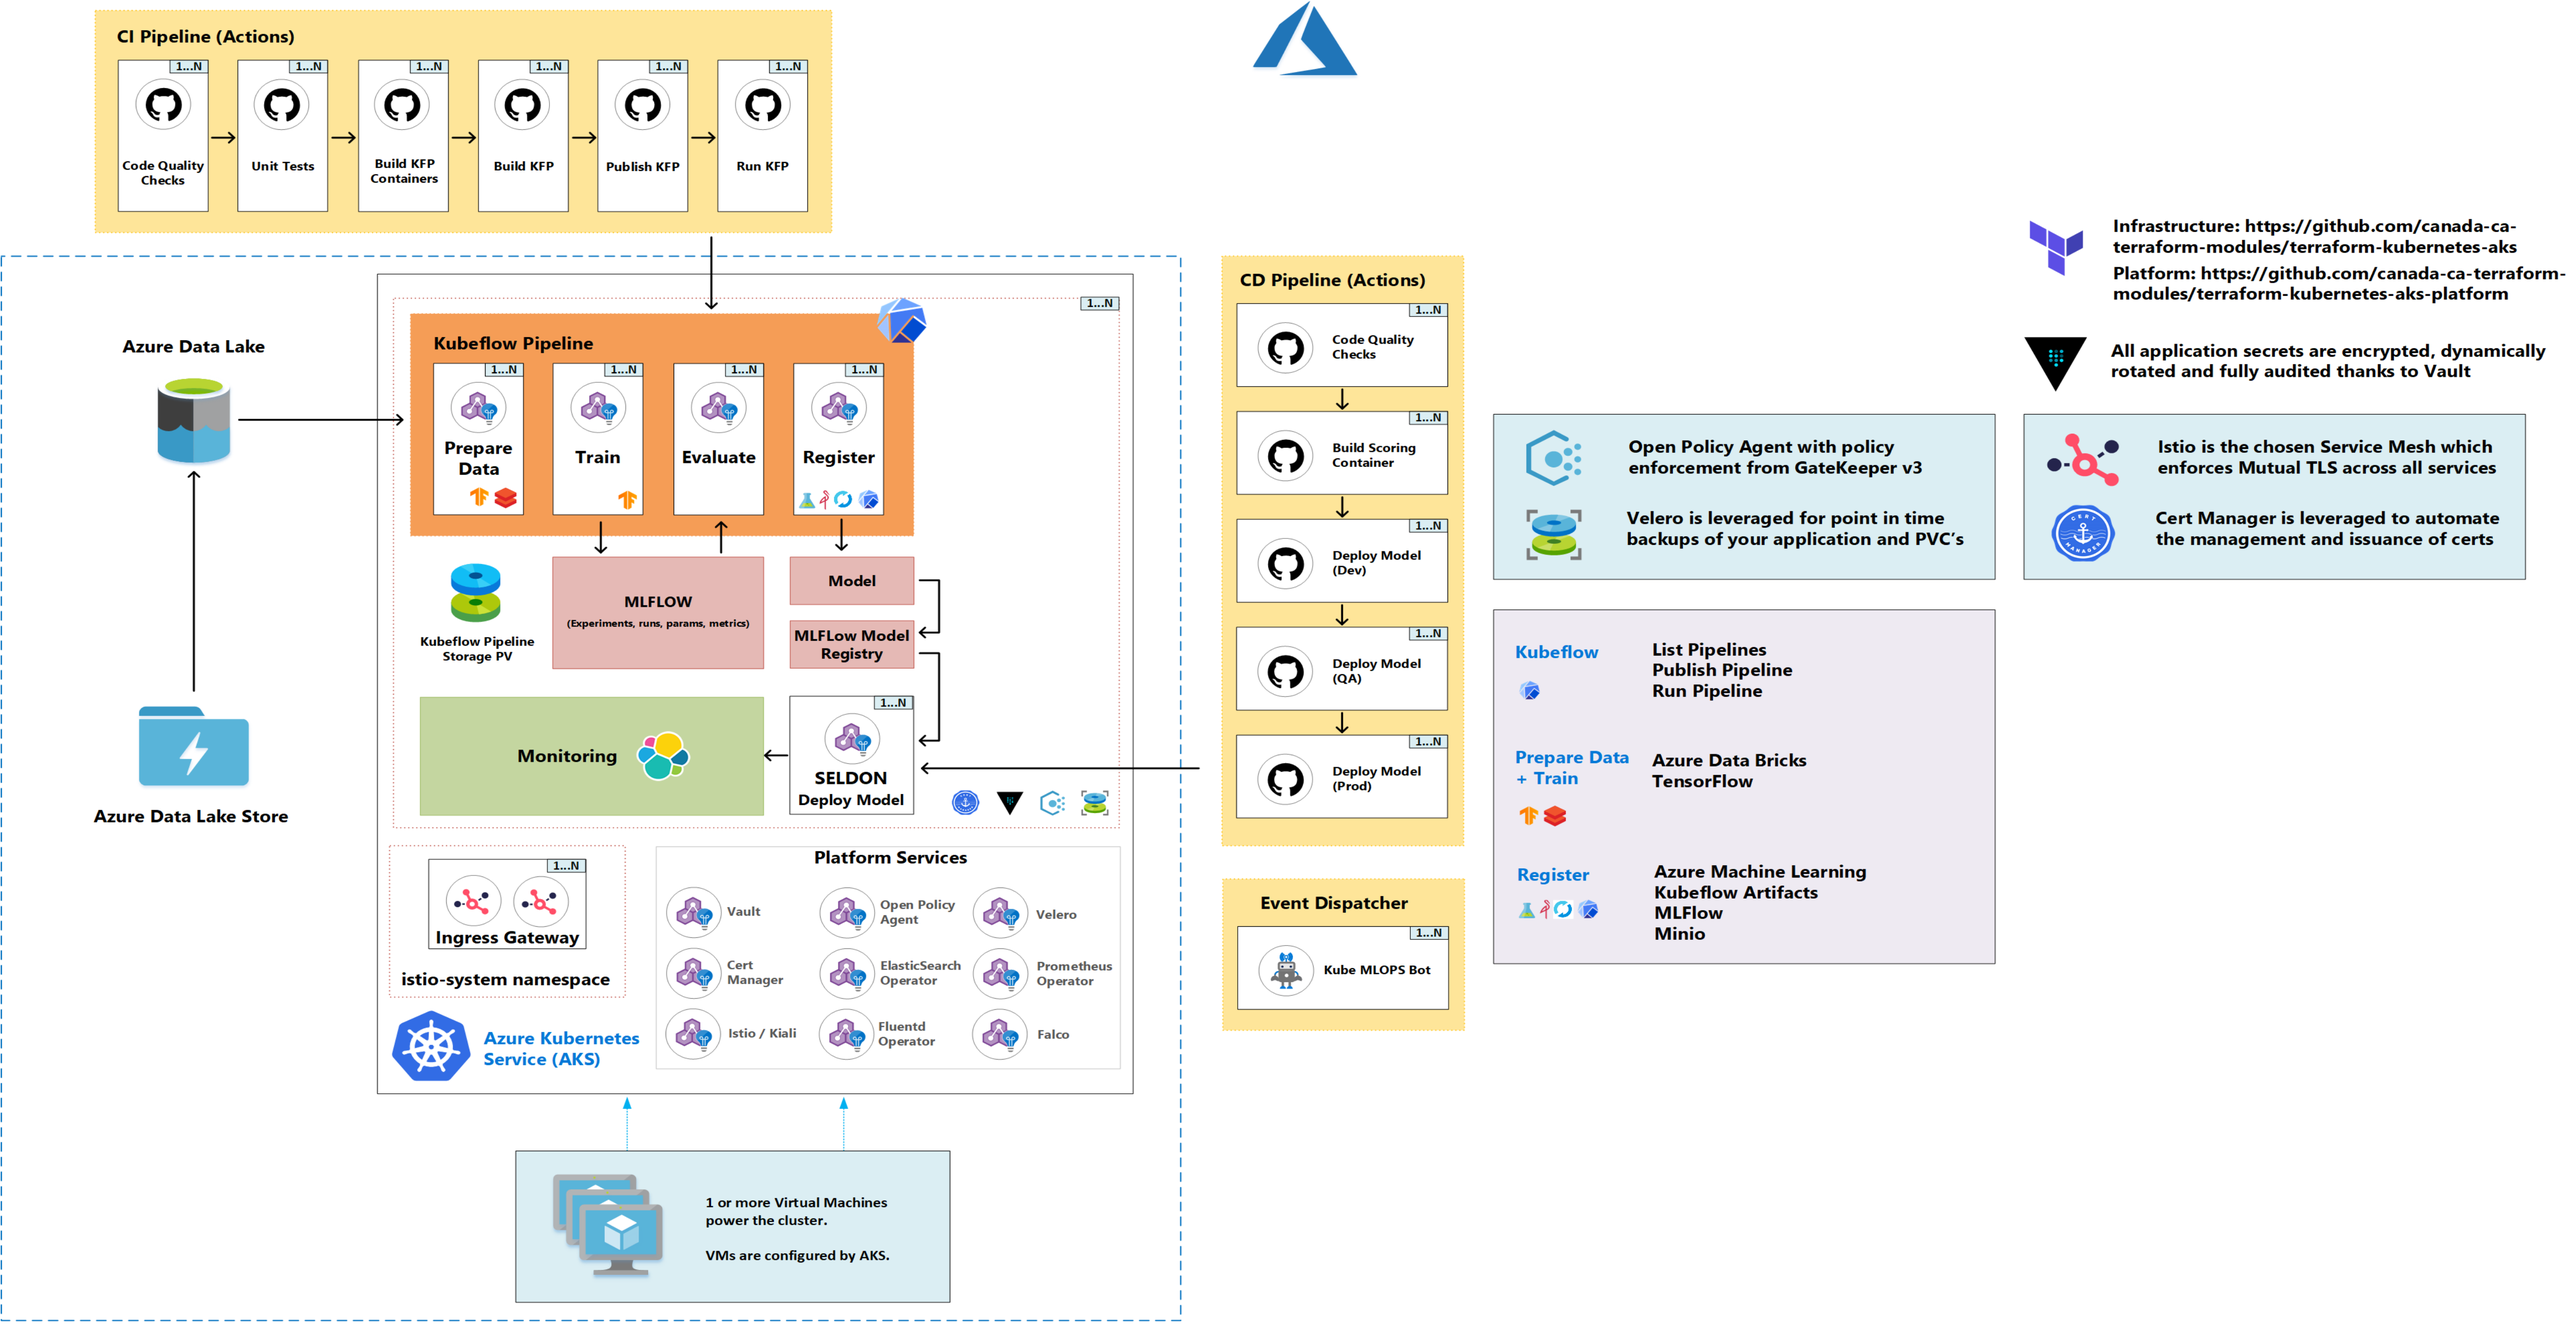

Kubeflow MLOPS

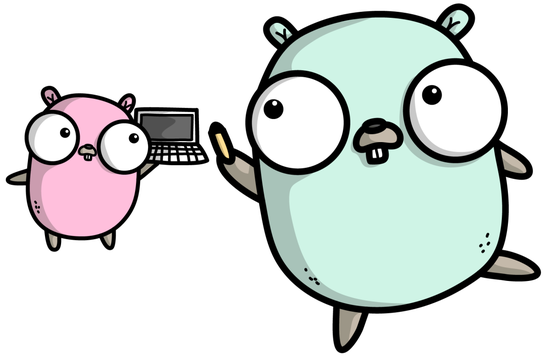

ARTWORK BY @ASHLEYMCNAMARA

WHAT DID WE LEARN?

EXPERIENCES

USER EXPERIENCE

- Web portal to introduce the platform
- Development of documentation to assist users with the platform
  - GitHub Pages
  - YouTube
- User support via Slack
- GitHub issues

POSITIVE LESSONS

- Automated scaling functions as designed
  - When extra compute is required, it is added. When it’s no longer needed, it is removed. No administrator intervention is required. Scaling to Zero for different node pools.
- Working closely with our users allowed us to build exactly what they needed
  - This was described as “the dream platform” by the Director of the Data Science Division
- Working with the open-source community allows for quick advancements
  - Re-use of existing work (we’re not the first ones to do data science on Kubernetes)
  - Contribution of our improvements into these systems / technologies (bug fixes, features, etc.)
  - Ability to adjust these tools for our specific use cases

POSITIVE LESSONS

- Use of automation via CI/CD for deployment of infrastructure and components
- Kubeflow Co-Founder and Chief OSS at Microsoft, David Aronchick, reviewed our work and agreed with the strategic direction of the platform
- Large organizations are picking similar toolsets
  - Microsoft, IBM, AWS working with Kubeflow
  - All cloud providers providing a managed Kubernetes service
- Self-service architecture enables users to accomplish their tasks without friction
  - For example, users can submit R Dashboards to a PR, and after a review, are published

WHAT COULD WE IMPROVE?

- We hit one notable issue with Kubeflow we are keen on resolving
  - Kubeflow is still building multi-user support, so some of its functionality is global and not user-isolated (pipelines, for example is coming in the next v1.1 release)
- Hiding some of the technical details of the platform, providing a better user experience for those unfamiliar with Kubernetes
  - Finding tools that help improve the user experience
  - Lens provides a rich UX and broad overview of the Kubernetes resources that are available for a end user

DEMO

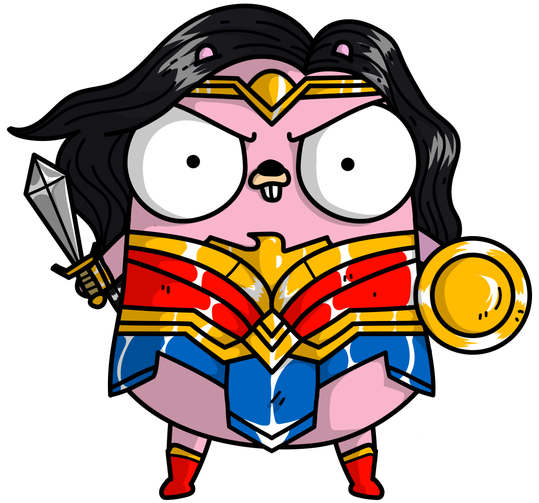

ARTWORK BY @ASHLEYMCNAMARA

DEMO

- Terraform GitHub automation
- KubeFlow
  - Setup profile
  - Launch a Jupyter Notebook (ML Workspace)
    - Automatic volume mounting with object stores
      - S3 compatible API’s (MinIO), Azure Blob Storage, Azure Files, and Azure Data Lakes
    - ML Workspace (VNC)
- KubeFlow MLOps (CI/CD workflow)
- Cost Tracking (KubeCost)

LOOKING FORWARD

- Continued development of the Data Analytics as a Service (DAaaS) platform
- Looking at expanding the platform towards other Government departments, as well as provincial / territorial governments and educational institutions
- Additional integrations for equivalent and newer services provided by other Cloud Providers (AWS, IBM, GKE, etc)
- Contributing some of our work back to the Kubeflow community (Boathouse, Jupyter APIS, Goofys Injector, Kubeflow Controller, etc)

ARTWORK BY @ASHLEYMCNAMARA

USEFUL LINKS (STATCAN)

- The following are links related to our implementation that should help you get started:

  - StatCan GitHub Org (DAaaS)
  - Advanced Analytics Workspace (User Manual)
  - AAW Platform (YouTube)
  - AAW Community (YouTube)
  - Brief Technical Summary
  - Kubeflow Manifests (GitHub)
  - Kubeflow Containers (GitHub)
  - Kubeflow MLOPS (GitHub)
  - Terraform AKS / Platform

USEFUL LINKS (GENERAL)

- The following are some links related to Kubeflow in general that we found very useful:
  - Kubeflow Documentation
  - Kubeflow (GitHub)
  - Kubeflow (Medium)
  - Kubeflow Dojo (GitHub)
  - Kubeflow 101 (YouTube)
  - ML in Production (YouTube)
  - Kubeflow for ML (Book)
  - Kubeflow Ops Guide (Book)

William Hearn

william.hearn@canada.ca

sylus

william_hearn

Questions?

Zachary Seguin

zachary.seguin@canada.ca

zachomedia

zachomedia

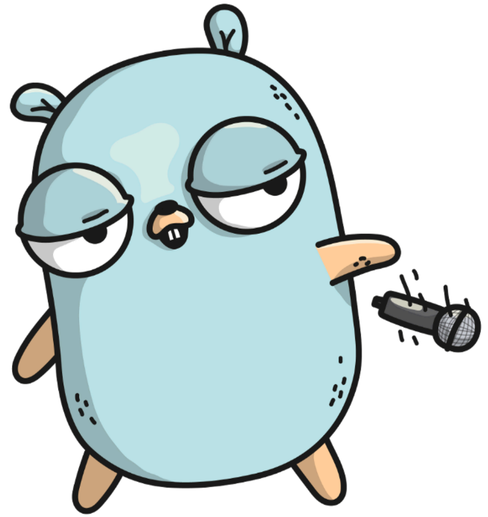

ARTWORK BY @ASHLEYMCNAMARA