

Managing Gen3 Deployments

Gen3 Community Forum
March 4, 2026

- A CSOC dashboard for managing Gen3 deployment and monitoring - Jawad Qureshi, Center for Translational Data Science, University of Chicago
- Developing Kro as a solution for infrastructure setup and Gen3 deployment - Alan Walsh and Jimi Adeyemi, Indiana University
- Bootstrapping Gen3 with AWS CDK, ArgoCD, and Helm - Guerdon Mukama, Australian BioCommons
- Building and Managing Multiple Data Commons in OCC's CSOC - Stephen Dranger and Urvi Sheth, Open Commons Consortium
- Using Gen3 as a platform and running additional services such as GEARBOX - Luca Graglia, Data for the Common Good, University of Chicago

A CSOC dashboard for managing Gen3 deployment and monitoring

Jawad Qureshi, CTDS, University of Chicago

Gen3 CSOC Dashboard

A comprehensive admin dashboard and platform for managing **Gen3 deployments** across multiple Kubernetes clusters with enterprise-scale infrastructure management capabilities.

 [Open Source on GitHub](#)

The challenge

Managing multiple Gen3 deployments across different environments (production, development, testing) is complex and time-consuming

The solution

A centralized platform providing system administrators with comprehensive tools to **deploy, monitor, and manage** Gen3 ecosystems effectively.

Agent/Server Architecture

Central Server

- > Go-based API server
- > REST endpoints
- > gRPC communication

Gen3 Agents

- > Deployed per cluster
- > Secure proxy layer
- > Cluster communication

Security

- > **mTLS Encryption**
Secure agent-server communication
- > **Certificate Based Auth**
Automated certificate provisioning

Frontend

- > Next.js application
- > Unified interface
- > REST API integration

External Integrations

- > KeyCloak (Auth)
- > Grafana (Monitoring)
- > ArgoCD (GitOps)

Frontend

Next JS

Pages Router for SSR/SSG

Mantine UI

Modern React component library

Rest API

Communication with backend

Backend

Go + Gin framework

High-performance API server

gRPC with mTLS

Secure agent communication

Kubernetes API

Native K8s integration

Infrastructure

- Kubernetes-Native
- Cloud Agnostic
- Helm Charts

Multi Environment Support

Manage **Production, development and testing** environments from a single interface

Helm based deployments

Built on top of helm for consistent Gen3 deployments.

Deployment Wizard

Guided setup process for new Gen3 installations with step-by-step configuration.

At-a-glance environment health

Quickly monitor system health, services, and deployment status in one view

Postgres UI

Browse and manage Gen3 relational data directly in PostgreSQL

Run queries, inspect schemas, and troubleshoot data issues

Simplifies database access for developers and operators

Built on-top of pgweb

Elasticsearch UI

Explore indexed Gen3 data through Elasticsearch

Run searches and inspect indexed documents and mappings

Monitor search performance and debug indexing issues

Multi-Cluster support

Manage multiple Kubernetes clusters from one dashboard. View all your clusters in a unified interface.

Resource Monitoring

Real-time visibility into pods, services, deployments, and ConfigMaps etc across all clusters.

Cluster Health

Monitor node health, resource usage (CPU, memory, storage), and system metrics in real-time.

RBAC Integration

Kubernetes Role-Based Access Control integration for secure access management across teams.

Infrastructure-as-code

Deploy infrastructure using Terraform, CDK, or any language—as long as it can be packaged and run as a container.

Multi-Cloud Support

EKS, GKE, Azure

Bring Your Own Infrastructure

Bring Your Own Infrastructure - support for existing infrastructure and on-premise deployments (Openshift etc).

Developing Kro as a solution for infrastructure setup and Gen3 deployment

Alan Walsh and Jimi Adeyemi, Indiana University

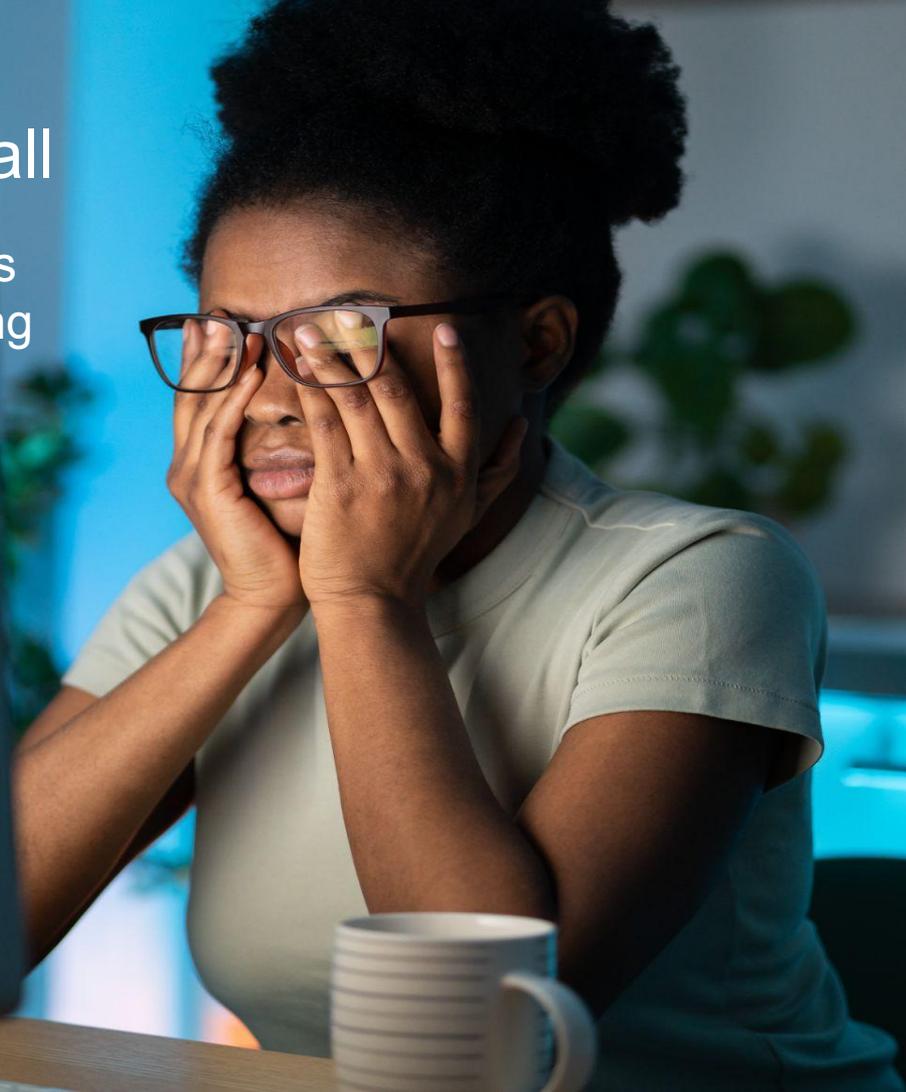
Developing Kro as a solution for infrastructure setup and Gen3 deployment



Alan Walsh and Jimi Adeyemi

The complexity wall

- Brittle Helm templates
- Terraform state locking
- “Glue code”



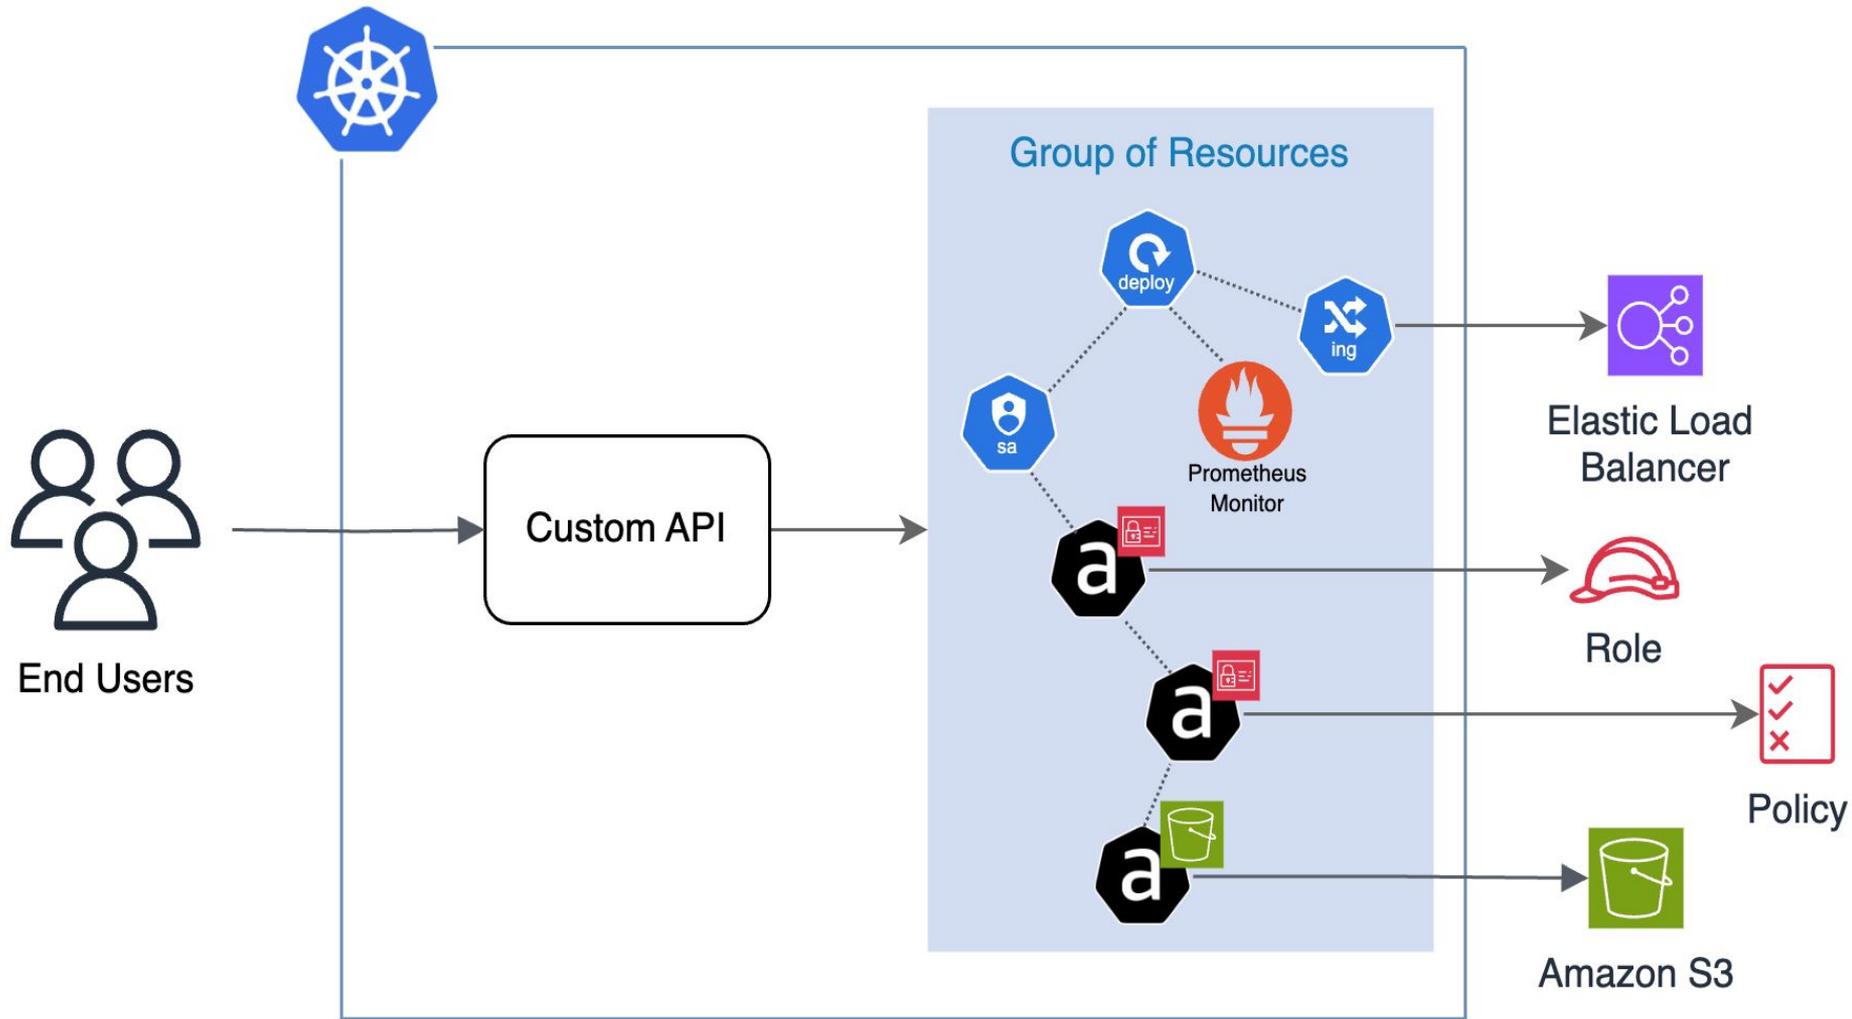
Kubernetes as a universal control plane

Kube Resource Orchestrator (Kro)

Open-source, Kubernetes-native project: <https://github.com/kubernetes-sigs/kro>

Lightweight engine that runs inside K8s to create custom APIs

Collaboration between AWS, Google, and Microsoft





gen3-kro

<https://github.com/indiana-university/gen3-kro>

In progress:

- AWS deployment for “hub” (Kro, ACK)
- ResourceGraph(s) for Gen3 infra
- ArgoCD integration

Future development:

- ResourceGraph(s) for Gen3 application
- Multi-cloud (GCP, Azure)

Demo

Agro CD Dashboard- Final State

APP HEALTH

Healthy

SYNC STATUS

Synced to Version2 (caac650)

Auto sync is enabled.

Author: Babasanmi Adeyemi <boadeyem@iu.edu>

Comment: fix: add Automode policies to EKS cluster role + ECR ...

LAST SYNC

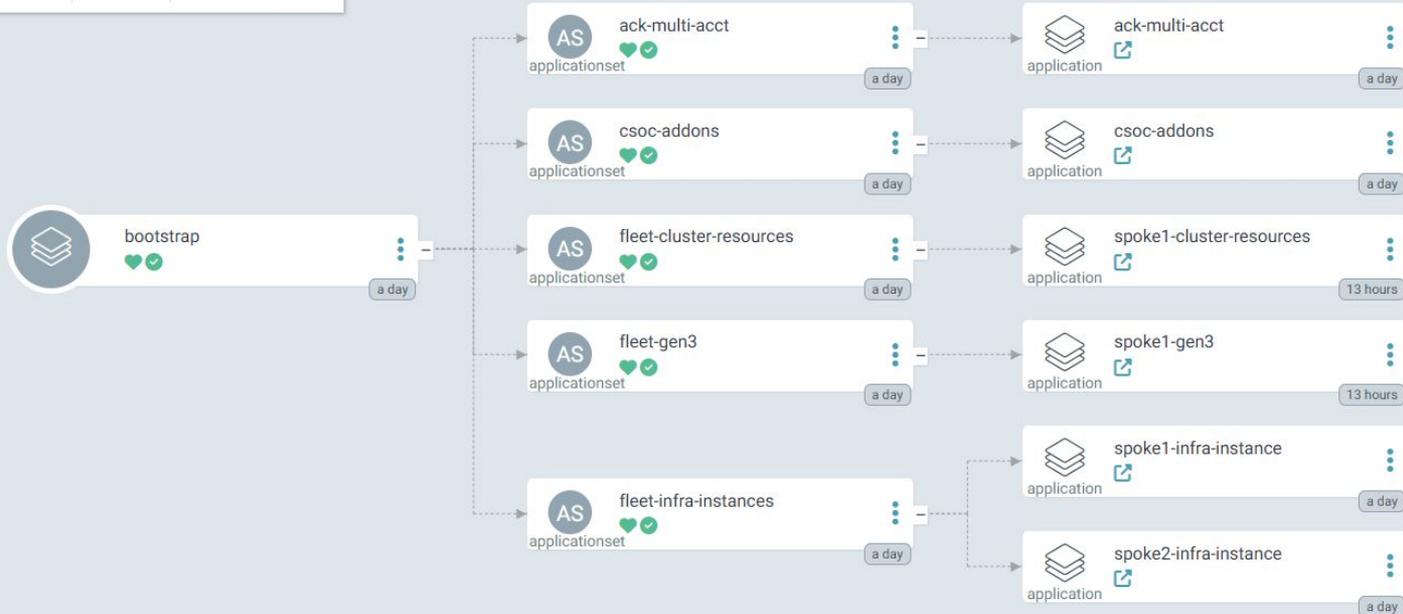
Sync OK to 87ea860

Succeeded a day ago (Sun Mar 01 2026 06:09:06 GMT-0500)

Author: Babasanmi Adeyemi <boadeyem@iu.edu>

Comment: feat: add IAM permissions for user policy

Navigation icons: list, expand, view, zoom in, zoom out, search, close, 100%



Agro CD Dashboard- Stage 1

APP HEALTH ●
♥ Healthy

SYNC STATUS ●
✔ Synced to [Version2 \(caac650\)](#)

Auto sync is enabled.
Author: Babasanmi Adeyemi <boadeyem@iu.edu> -
Comment: fix: add Automode policies to EKS cluster role + ECR ...

LAST SYNC ●
✔ Sync OK to [87ea860](#)

Succeeded a day ago (Sun Mar 01 2026 06:09:06 GMT-0500)
Author: Babasanmi Adeyemi <boadeyem@iu.edu> -
Comment: feat: add IAM permissions for user policy

☰ ⚙️ 🖼️ + - 🔍 🗨️ 100%

📄 bootstrap
✔✔
a day



Shared

- CSOC Multi-acct (Shared) - IAM mapping per namespace within CSOC for cross-account management
- CSOC Addons (Shared) - AWS Controllers for Kubernetes (ACKs), Google Cloud Controllers (GCCs), Azure Service Operator (ASO), KRO Controller, KRO Resource Graph Definitions

Agro CD Dashboard- Stage 2

APP HEALTH

Healthy

SYNC STATUS

Synced to [Version2 \(caac650\)](#)

Auto sync is enabled.

Author: Babasanmi Adeyemi <boadeyem@iu.edu>

Comment: fix: add Automode policies to EKS cluster role + ECR ...

LAST SYNC

Sync OK to [87ea860](#)

Succeeded a day ago (Sun Mar 01 2026 06:09:06 GMT-0500)

Author: Babasanmi Adeyemi <boadeyem@iu.edu>

Comment: feat: add IAM permissions for user policy

Navigation icons: list, full screen, search, zoom in, zoom out, zoom to 100%

bootstrap application icon with health status and sync frequency (a day)

- Fleet Infrastructure Instances - Manage KRO Instances (Kind Object) deployed in the CSOC per spoke namespace - these instances create the cloud infrastructure.

fleet-infra-instances applicationset icon with health status and sync frequency (a day)

spoke1-infra-instance and spoke2-infra-instance application icons with health status and sync frequency (a day)

Spoke Infra

Agro CD Dashboard- Stage 3

APP HEALTH Healthy

SYNC STATUS Synced to [Version2 \(caac650\)](#)
Auto sync is enabled.
Author: Babasanmi Adeyemi <boadeyem@iu.edu> -
Comment: fix: add Automode policies to EKS cluster role + ECR ...

LAST SYNC Sync OK to [87ea860](#)
Succeeded a day ago (Sun Mar 01 2026 06:09:06 GMT-0500)
Author: Babasanmi Adeyemi <boadeyem@iu.edu> -
Comment: feat: add IAM permissions for user policy

☰ ⚙️ 🖼️ + - 🔍 🗨️ 100%

- Fleet Cluster Resources - Deploy the cluster-resources Helm Chart as per current Gen3 requirements.



Spoke App

- Fleet Gen3 - Deploy the Gen3 Application into the built Spoke Cluster.

ACK-multi-acct

- Helm-Templated Kind which deploys an IAM Role Selector per spoke namespace

```
11 {{- $clusters := .Values.clustersJson | fromJson }}
12 {{- range $alias, $accountId := $clusters }}
13 ---
14 apiVersion: services.k8s.aws/v1alpha1
15 kind: IAMRoleSelector
16 metadata:
17   name: {{ $alias }}-ack-role
18   namespace: {{ $alias }}
19   annotations:
20     | argocd.argoproj.io/sync-options: SkipDryRunOnMissingResource=true
21 spec:
22   arn: {{ printf "arn:aws:iam:%s:role/%s-spoke-role" $accountId $alias | quote }}
23   namespaceSelector:
24     names:
25     | - {{ $alias | quote }}
26 {{- end }}
27
```

Basic Resource Graph Definition: Overview

```
1  apiVersion: kro.run/v1alpha1
2  kind: ResourceGraphDefinition
3  metadata:
4    name: awsgen3infraflat
5    annotations:
6      argocd.argoproj.io/sync-options: SkipDryRunOnMissingResource=true
7      argocd.argoproj.io/sync-wave: "0"
8  spec:
9    schema:
10     apiVersion: v1alpha1
11     kind: AwsGen3InfraFlat
12 >   spec: ...
97 >   status: ...
166   # # Cache Status (ElastiCache)
167
168   resources:
169   # =====
170   # Layer 0: Namespace References
171   # Read spoke namespace to get spoke account ID (CARM annotation)
172 >   - id: spokeNamespace ...
178   # Read argocd namespace to get CSOC account ID (Terraform annotation)
179 >   - id: argocdNamespace ...
185   # =====
186   # SECTION 2: NETWORKING (VPC & Subnets)
187   # =====
188   # Layer 1: VPC
189 >   - id: vpc ...
218   # Layer 1: Elastic IP 1 (Required)
219 >   - id: eip1 ...
243   # Layer 2: Internet Gateway
244 >   - id: igw ...
272   # Layer 3: Public Route Table
273 >   - id: publicRouteTable ...
301   # Layer 4: Public Subnet 1 (Required)
302 >   - id: publicsubnet1 ...
336   # Layer 4: Public Subnet 2 (Required)
```

Basic Resource Graph Definition: Schema section

```
8 spec:
9   schema:
10    apiVersion: v1alpha1
11    kind: AwsGen3Infra1Flat
12    spec:
13      # Basic metadata
14      name: string | required=true
15      namespace: string | default="default"
16      region: string | default="us-east-1"
17      environment: string | default="dev"
18      project: string | default="kro-managed"
19      adoptionPolicy: string | default="adopt-or-create"
20      deletionPolicy: string | default="delete"
21
22      # ===== Application Configuration =====
23      # Hostname for the Gen3 deployment (used in fence-config BASE_URL)
24      hostname: string | required=true
25      # Default login IDP (e.g. google, cognito)
26      defaultLoginIDP: string | default="google"
27
```

Basic Basic Resource Graph Definition: Resource section

- a simple resource defined by the graph

```
429 - id: dbRouteTable
430   readyWhen:
431     - ${dbRouteTable.status.?routeTableID.orValue('null')} != 'null'}
432     - ${dbRouteTable.status.?conditions.orValue([]).exists(c, c.type == "ACK.ResourceSynced" && c.status == "True")}
433   template:
434     apiVersion: ec2.services.k8s.aws/v1alpha1
435     kind: RouteTable
436     metadata:
437       name: ${schema.spec.name}-db-rt
438       namespace: ${schema.spec.namespace}
439       annotations:
440         services.k8s.aws/region: ${schema.spec.region}
441         services.k8s.aws/adoption-policy: ${schema.spec.adoptionPolicy}
442         services.k8s.aws/deletion-policy: ${schema.spec.deletionPolicy}
443     spec:
444       vpcID: ${vpc.status.?vpcID}
445       routes:
446         - destinationCIDRBlock: "0.0.0.0/0"
447           natGatewayID: ${natGateway1.status.?natGatewayID}
448       tags:
449         - key: Name
450           value: ${schema.spec.name}-db-rt
451         - key: Environment
452           value: ${schema.spec.environment}
453         - key: ManagedBy
454           value: Kro
455         - key: Project
456           value: ${schema.spec.project}
```

Basic Infrastructure Graph: Output Resource

- describes to the ArgoCD, access to the spoke cluster for gen3 deployment

```
1517 - id: infraOutputsSecret
1518   readyWhen:
1519     - ${infraOutputsSecret.metadata.?name.orValue('null')} != 'null'}
1520   template:
1521     apiVersion: v1
1522     kind: Secret
1523     metadata:
1524       name: ${schema.spec.name}-infra-outputs
1525       namespace: ${schema.spec.namespace}
1526       labels:
1527         app.kubernetes.io/managed-by: kro
1528         gen3.io/infra-outputs: "true"
1529       annotations:
1530         kro.run/aurora-endpoint: ${auroraCluster.status.?endpoint}
1531         kro.run/eks-cluster: ${eksCluster.status.?ackResourceMetadata.?arn}
1532   type: Opaque
1533   stringData:
1534     # Aurora
1535     aurora_endpoint: ${auroraCluster.status.?endpoint}
1536     aurora_reader_endpoint: ${auroraCluster.status.?readerEndpoint}
1537     aurora_cluster_arn: ${auroraCluster.status.?ackResourceMetadata.?arn}
1538     aurora_master_username: ${schema.spec.auroraMasterUsername}
1539     aurora_database_name: ${schema.spec.auroraDatabaseName}
1540     # EKS
1541     eks_cluster_name: ${eksCluster.metadata.name}
1542     eks_endpoint: ${eksCluster.status.?endpoint}
1543     eks_ca_data: ${eksCluster.status.?certificateAuthority.?data}
1544     eks_cluster_arn: ${eksCluster.status.?ackResourceMetadata.?arn}
```

Infrastructure Graph Status

- for user monitoring and linking with dependent graphs

```
spec:
  schema:
    apiVersion: v1alpha1
    kind: AwsGen3InfraFlat
    spec: ...
  status:
    # Security Status (KMS)
    loggingKeyARN: ${loggingKey.status.?ackResourceMetadata.?arn.orValue("loading")}
    databaseKeyARN: ${databaseKey.status.?ackResourceMetadata.?arn.orValue("loading")}
    searchKeyARN: ${searchKey.status.?ackResourceMetadata.?arn.orValue("loading")}
    platformKeyARN: ${platformKey.status.?ackResourceMetadata.?arn.orValue("loading")}

    # Network Status
    vpcID: ${vpc.status.?vpcID.orValue("loading")}
    internetGatewayID: ${igw.status.?internetGatewayID.orValue("loading")}
    publicSubnetIDs:
      - ${publicsubnet1.status.?subnetID.orValue("loading")}
      - ${publicsubnet2.status.?subnetID.orValue("loading")}
    privateSubnetIDs:
      - ${privatesubnet1.status.?subnetID.orValue("loading")}
      - ${privatesubnet2.status.?subnetID.orValue("loading")}
    dbSubnetIDs:
      - ${dbsubnet1.status.?subnetID.orValue("loading")}
      - ${dbsubnet2.status.?subnetID.orValue("loading")}
```

Agro CD Dashboard- Stage 2

APP HEALTH

Healthy

SYNC STATUS

Synced to [Version2 \(caac650\)](#)

Auto sync is enabled.

Author: Babasanmi Adeyemi <boadeyem@iu.edu>

Comment: fix: add Automode policies to EKS cluster role + ECR ...

LAST SYNC

Sync OK to [87ea860](#)

Succeeded a day ago (Sun Mar 01 2026 06:09:06 GMT-0500)

Author: Babasanmi Adeyemi <boadeyem@iu.edu>

Comment: feat: add IAM permissions for user policy



bootstrap
Healthy
a day

- Fleet Infrastructure Instances - Manage KRO Instances (Kind Object) deployed in the CSOC per spoke namespace - these instances create the cloud infrastructure.

AS fleet-infra-instances
applicationset
Healthy
a day

spoke1-infra-instance
application
a day

spoke2-infra-instance
application
a day

Spoke Infra

Spoke - Infrastructure Instance: Part 2

- describes for one account, what we want to deploy (uses CEL)

```
33 privateSubnetCIDRs:
34   - "10.1.0.0/20"
35   - "10.1.16.0/20"
36 dbSubnetCIDRs:
37   - "10.1.32.0/24"
38   - "10.1.33.0/24"
39 availabilityZones:
40   - "us-east-1a"
41   - "us-east-1b"
42
43 # --- S3 Buckets ---
44 loggingBucketName: gen3-spoke1-dev-logging
45 dataBucketName: gen3-spoke1-dev-data
46 uploadBucketName: gen3-spoke1-dev-upload
47
48 # --- Aurora PostgreSQL ---
49 auroraEngine: aurora-postgresql
50 auroraEngineMode: provisioned
51 auroraEngineVersion: "14.19"
52 auroraMasterUsername: postgres
53 auroraDatabaseName: app
54 auroraInstanceClass: db.r6g.large
55 auroraInstanceCount: 2
56
57 # --- EKS ---
58 eksVersion: "1.35"
59
60 # --- KMS ---
61 kmsKeyDescription: gen3-spoke1 encryption key
62 kmsKeyAlias: gen3-spoke1-key
```

Agro CD Dashboard- Stage 3

APP HEALTH Healthy

SYNC STATUS Synced to [Version2 \(caac650\)](#)
Auto sync is enabled.
Author: Babasanmi Adeyemi <boadeyem@iu.edu> -
Comment: fix: add Automode policies to EKS cluster role + ECR ...

LAST SYNC Sync OK to [87ea860](#)
Succeeded a day ago (Sun Mar 01 2026 06:09:06 GMT-0500)
Author: Babasanmi Adeyemi <boadeyem@iu.edu> -
Comment: feat: add IAM permissions for user policy

☰ ⚙️ 🖼️ + - 🔍 🗨️ 100%

bootstrap
🟢🟢
a day

- Fleet Cluster Resources - Deploy the cluster-resources Helm Chart as per current Gen3 requirements.

fleet-cluster-resources
applicationset
🟢🟢
a day

fleet-gen3
applicationset
🟢🟢
a day

spoke1-cluster-resources
application
🔗
13 hours

spoke1-gen3
application
🔗
13 hours

Spoke App

- Fleet Gen3 - Deploy the Gen3 Application into the built Spoke Cluster.

Spoke - Gen3 Application

```
163   apiVersion: argoproj.io/v1alpha1
164   kind: ApplicationSet
165   metadata:
166     name: fleet-gen3
167     namespace: argocd
168 >   annotations: ...
170   spec:
171 >   syncPolicy: ...
173   goTemplate: true
174 >   goTemplateOptions: ...
176 >   generators: ...
185   template:
186     metadata:
187       name: '{{index .metadata.annotations "gen3.io/spoke-alias"}}-gen3'
188       labels:
189         environment: '{{index .metadata.labels "environment"}}'
190         fleet_spoke: '{{index .metadata.annotations "gen3.io/spoke-alias"}}'
191     spec:
192       project: default
193       # — Ignore auto-generated secrets (gen3-helm generates random passwords) —
194       # These Secrets regenerate random values on each Helm render, causing
195       # permanent OutOfSync. Ignore the data field so ArgoCD doesn't fight
196       # the initial random values already applied to the spoke cluster.
197 >     ignoreDifferences: ...
203     sources:
204 >     - repoURL: '{{.values.repoUrl}}' ...
207       # gen3-helm umbrella chart - deployed directly (no intermediate chart)
208 >     - repoURL: 'https://github.com/uc-cdis/gen3-helm' ...
224     destination:
225       # Deploy directly to the spoke cluster
226       name: '{{.name}}'
227       namespace: default
```

Agro CD Dashboard- Final state

APP HEALTH Healthy

SYNC STATUS Synced to [Version2 \(caac650\)](#)

Auto sync is enabled.
Author: Babasanmi Adeyemi <boadeyem@iu.edu> -
Comment: fix: add Automode policies to EKS cluster role + ECR ...

LAST SYNC Sync OK to [87ea860](#)

Succeeded a day ago (Sun Mar 01 2026 06:09:06 GMT-0500)
Author: Babasanmi Adeyemi <boadeyem@iu.edu> -
Comment: feat: add IAM permissions for user policy

bootstrap applicationset (a day)

- ack-multi-acct** applicationset (a day) → **ack-multi-acct** application (a day) [Shared]
- csoc-addons** applicationset (a day) → **csoc-addons** application (a day) [Shared]
- fleet-cluster-resources** applicationset (a day) → **spoke1-cluster-resources** application (13 hours) [Spoke App]
- fleet-gen3** applicationset (a day) → **spoke1-gen3** application (13 hours) [Spoke App]
- fleet-infra-instances** applicationset (a day) → **spoke1-infra-instance** application (a day) [Spoke Infra]
- fleet-infra-instances** applicationset (a day) → **spoke2-infra-instance** application (a day) [Spoke Infra]

Code Review/Questions

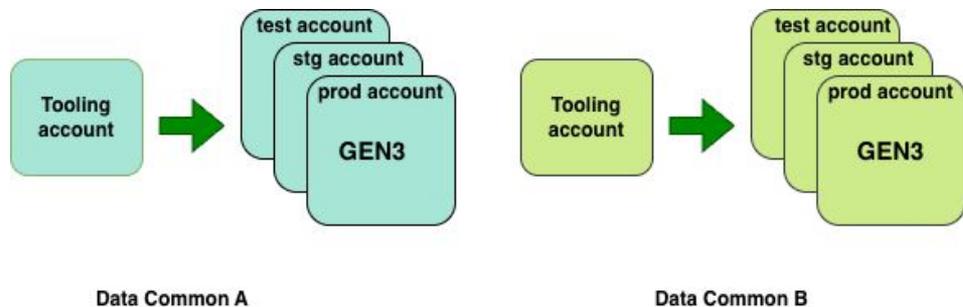
Bootstrapping Gen3 with AWS CDK, ArgoCD, and Helm

Guerdon Mukama, Australian BioCommons

Our Operating Model

Operating Multiple Gen3 Projects

- Multiple Gen3 Data commons
- Independent AWS accounts per project, environments
- No shared EKS clusters



The Scaling Challenge

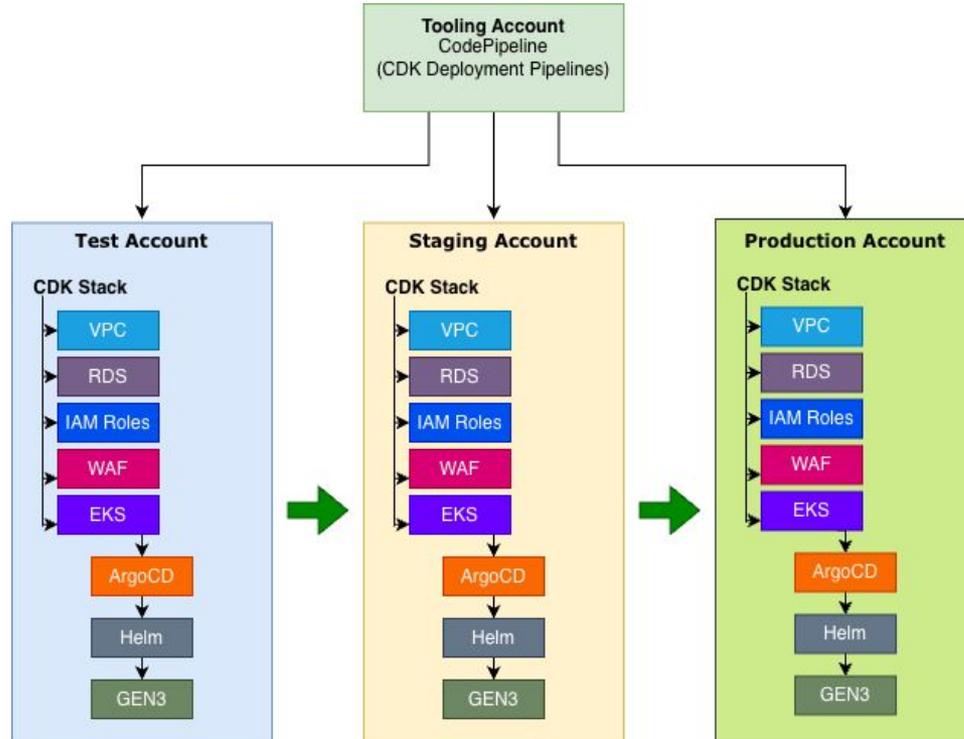
How Do You Scale Without Scaling Engineers?

- Each Gen3 requires:
 - Networking
 - EKS
 - RDS
 - IAM Roles
 - AWS Secrets
 - WAF
 - GitOps bootstrap
 - Environment promotion
 - ArgoCD
- Manual provisioning does not scale
- Copy-paste infrastructure causes drift
Security posture becomes inconsistent

We needed a repeatable pattern — not bespoke deployments.

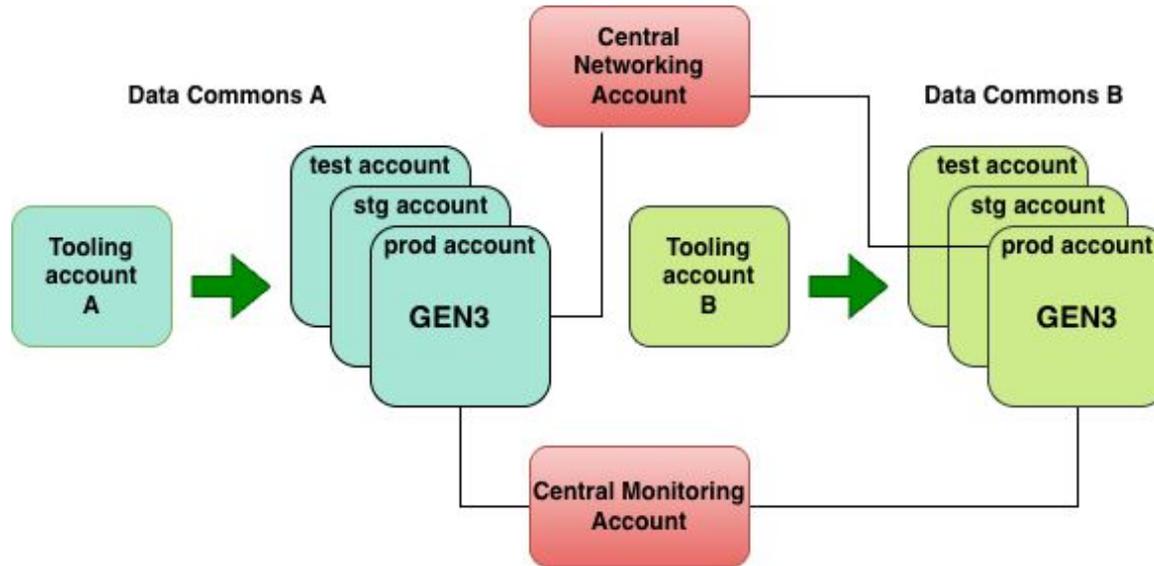
High-Level Architecture

Layered Multi-Account Deployment Model



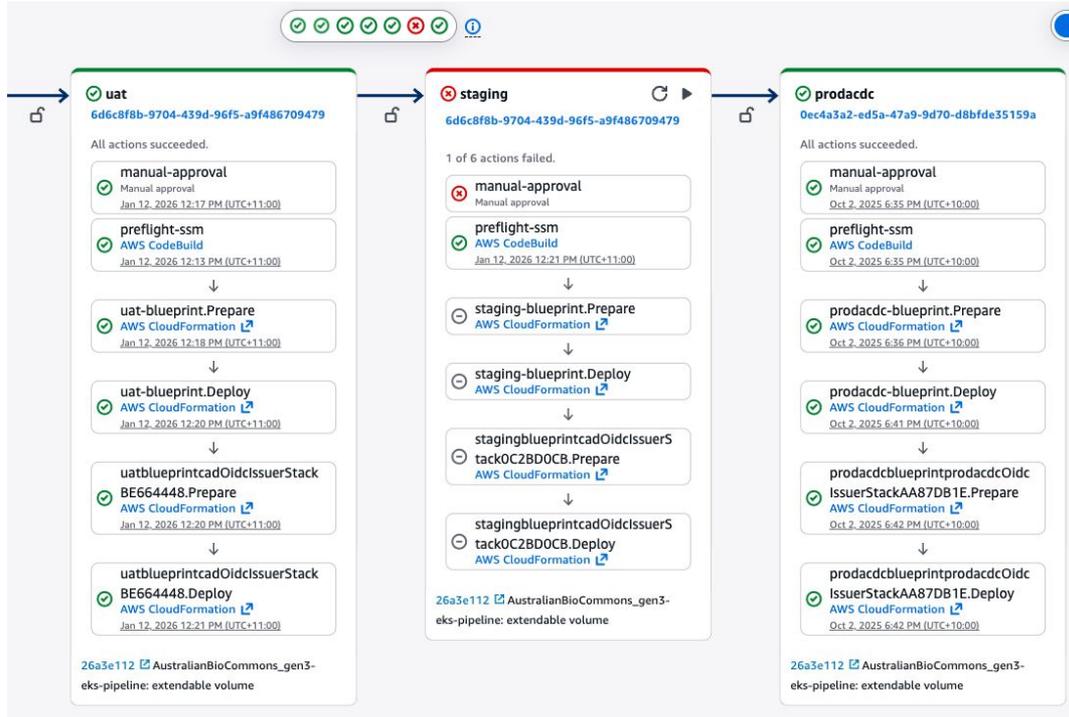
High-Level Architecture

Multi-Commons AWS Account Model



Tooling Account & CI/CD Pipeline

Infrastructure Pipeline for a Data Commons



- Tooling account hosts infra pipeline per data commons
- Promotes stacks through UAT/Staging/Prod
- Ensures reproducible deployments and consistent security

<https://github.com/AustralianBioCommons/gen3-eks-pipeline>

What CDK Standardises

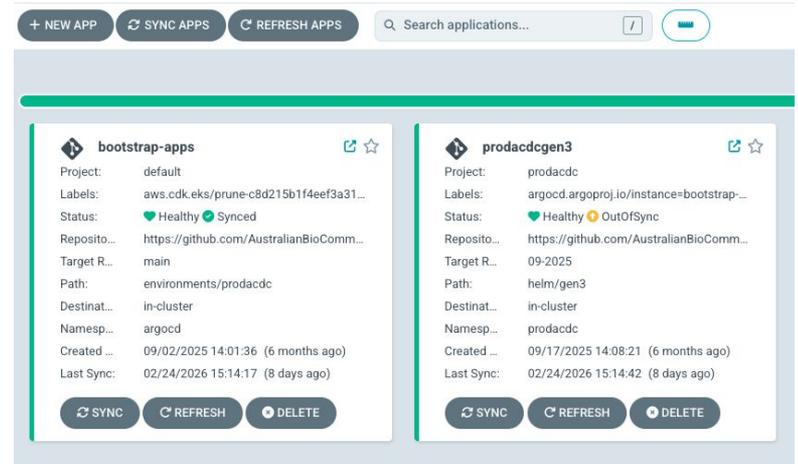
CDK Provides Baseline Infrastructure

- Networking (VPC)
- RDS
- EKS clusters + node groups
- IAM roles
- WAF
- ArgoCD bootstrap

GitOps Layer

ArgoCD + Helm for Application Lifecycle

- CDK bootstraps ArgoCD
- ArgoCD watches Git repositories
- Helm deploys Gen3 services
- Declarative state + drift correction
- Separation of infrastructure vs application lifecycle



Why No Shared Clusters

Isolation by Design

- Independent governance and policies
- Reduced blast radius
- Simplified RBAC
- Easier auditing

Scaling Gen3 Through Standardisation

- Onboard a new data commons without redesigning infrastructure
- Predictable security baseline across environments
- Clear separation of infrastructure and application lifecycle
- Reduced operational drift
- Controlled promotion across UAT → Staging → Production

Questions

Source code:

EKS Pipeline: <https://github.com/AustralianBioCommons/gen3-eks-pipeline>

Infra Bootstrap: <https://github.com/AustralianBioCommons/gen3-infra-bootstrap>

Building and Managing Multiple Data Commons in OCC's CSOC

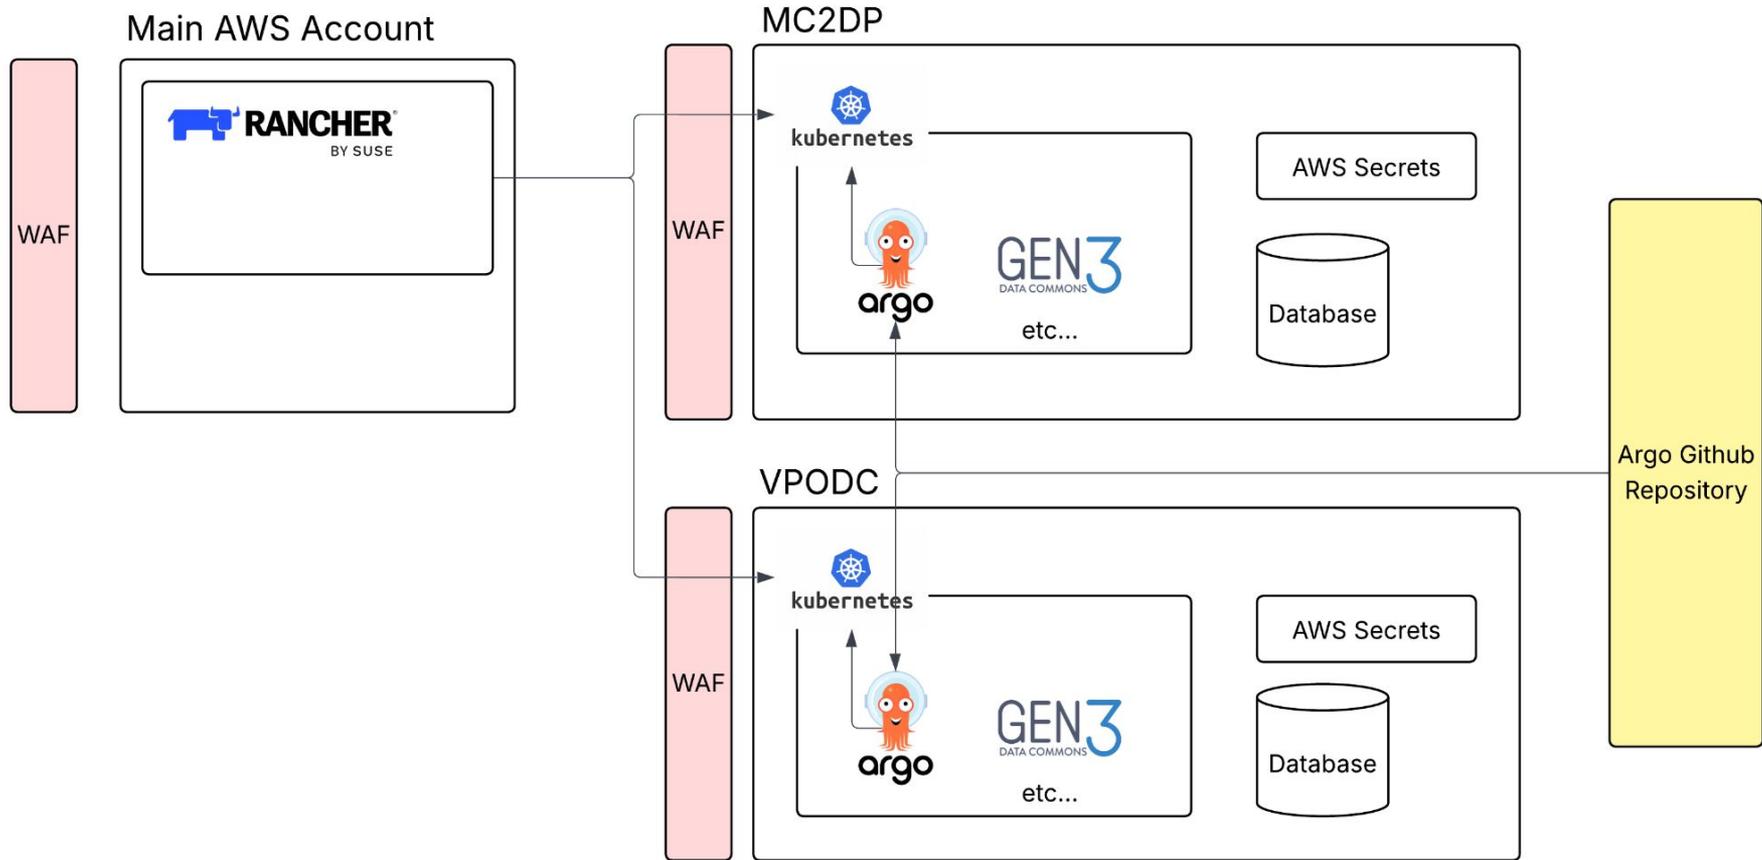
Stephen Dranger and Urvi Sheth,
Sr. Software Engineers
Open Commons Consortium

Objectives

- Central management and monitoring of commons within a CSOC
- Automated deployment of fully-functional Gen3 environment
- Automated version upgrades, including infrastructure code
- CI/CD process for new configuration changes
- Enforce best security practices
- Leverage existing code as much as possible



Architecture

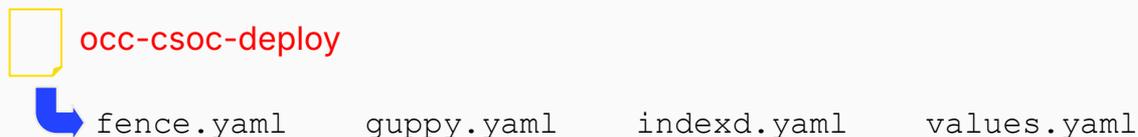




Demo

Continuous Deployment with Argo

Argo watches a Github repository for changes, then deploys as needed



Other Services in Environment

Argo CD

Certificate Manager

External DNS

External Secrets

Jaeger (monitoring)

Loki & Promtail (logging)

AWS LB Controller

PGAdmin



occ-windtunnel-dev

All Namespaces

- Cluster
- Workloads
 - CronJobs (1)
 - DaemonSets (8)
 - Deployments (38)**
 - Jobs (4)
 - StatefulSets (7)
 - Pods (81)
- Apps
- Service Discovery
- Storage
- Policy
- More Resources

| Namespace | Deployment Name | Image | Replicas | Ready | Unavailable | Age | Health | Actions |
|---------------------------------|-------------------------------|--|----------|-------|-------------|-----|---------|---------|
| Namespace: jaeger-system | | | | | | | | |
| <input type="checkbox"/> Active | jaeger-collector | jaegertracing/jaeger-collector:1.53.0 | 1/1 | 1 | 1 | 5 | 34 days | ⋮ |
| <input type="checkbox"/> Active | jaeger-query | jaegertracing/jaeger-query:1.53.0 <small>+ 1 more</small> | 1/1 | 1 | 1 | 0 | 34 days | ⋮ |
| Namespace: kube-system | | | | | | | | |
| <input type="checkbox"/> Active | aws-load-balancer-controller | public.ecr.aws/eks/aws-load-balancer-controller:v2.12.0 | 1/1 | 1 | 1 | 0 | 34 days | ⋮ |
| <input type="checkbox"/> Active | coredns | 602401143452.dkr.ecr.us-east-1.amazonaws.com/eks/coredns:v1.11.4-eksbuild.22 | 2/2 | 2 | 2 | 0 | 45 days | ⋮ |
| <input type="checkbox"/> Active | ebs-csi-controller | 602401143452.dkr.ecr.us-east-1.amazonaws.com/eks/aws-ebs-csi-driver:v1.48.0 <small>+ 5 more</small> | 2/2 | 2 | 2 | 0 | 45 days | ⋮ |
| <input type="checkbox"/> Active | efs-csi-controller | 602401143452.dkr.ecr.us-east-1.amazonaws.com/eks/aws-efs-csi-driver:v2.1.11 <small>+ 2 more</small> | 2/2 | 2 | 2 | 0 | 45 days | ⋮ |
| Namespace: pgadmin4 | | | | | | | | |
| <input type="checkbox"/> Active | pgadmin4 | dpape/pgadmin4:9.1 | 1/1 | 1 | 1 | 0 | 34 days | ⋮ |
| Namespace: windtunnel-occ | | | | | | | | |
| <input type="checkbox"/> Active | ambassador-deployment | quay.io/datawire/ambassador:1.4.2 | 1/1 | 1 | 1 | 0 | 27 days | ⋮ |
| <input type="checkbox"/> Active | arborist-deployment | quay.io/cdis/arborist:2025.09 | 1/1 | 1 | 1 | 0 | 27 days | ⋮ |
| <input type="checkbox"/> Active | fence-deployment | quay.io/cdis/fence:2025.09 | 1/1 | 1 | 1 | 0 | 27 days | ⋮ |
| <input type="checkbox"/> Active | frontend-framework-deployment | quay.io/cdis/vpodc-data-commons:main | 1/1 | 1 | 1 | 0 | 27 days | ⋮ |
| <input type="checkbox"/> Active | hatchery-deployment | quay.io/cdis/hatchery:2025.09 | 1/1 | 1 | 1 | 0 | 27 days | ⋮ |
| <input type="checkbox"/> Active | indexd-deployment | quay.io/cdis/indexd:master | 1/1 | 1 | 1 | 0 | 27 days | ⋮ |
| <input type="checkbox"/> Active | manifestservice-deployment | quay.io/cdis/manifestservice:master | 1/1 | 1 | 1 | 0 | 27 days | ⋮ |
| <input type="checkbox"/> Active | metadata-deployment | quay.io/cdis/metadata-service:feat_es-7 | 1/1 | 1 | 1 | 0 | 27 days | ⋮ |
| <input type="checkbox"/> Active | peregrine-deployment | quay.io/cdis/peregrine:2025.09 | 1/1 | 1 | 1 | 0 | 27 days | ⋮ |

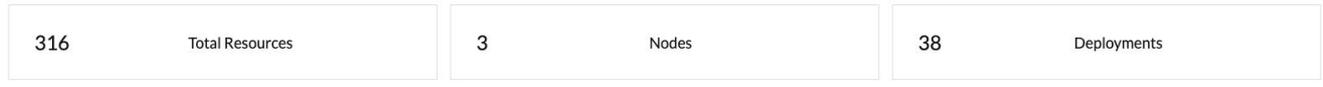
v2.10.2

Cluster Dashboard

occ-windtunnel-dev-eks-cluster

Provider: Amazon EKS Kubernetes Version: v1.32.9-113c36 Architecture: amd64 Created: 45 days ago

Cluster Tools



Capacity



- Etcd
- Scheduler
- Controller Manager
- Cattle
- Fleet

| Reason | Object | Message | Name | Date |
|--------------|--|---|--|------------------------------|
| BackOff | Pod portal-deployment-59cd49cc86-d4f6z | Back-off restarting failed container portal in pod portal-deployment-59cd49cc86-d4f6z_windtunnel-occ(9e4d93c4-40bc-4319-b903-fdbbd1c730f) | portal-deployment-59cd49cc86-d4f6z.186a705c79c0f4cc3 | Mon, Oct 27 2025 10:48:48 pm |
| Valid | ClusterSecretStore aws-secretsmanager | store validated | aws-secretsmanager.1867fb912082d2b | Mon, Oct 27 2025 10:48:04 pm |
| Valid | SecretStore gen3-secret-store | store validated | gen3-secret-store.186a28a35797a0aa | Mon, Oct 27 2025 10:48:04 pm |
| Pulled | Pod portal-deployment-59cd49cc86-d4f6z | Container image "quay.io/cdis/data-portal:master" already present on machine | portal-deployment-59cd49cc86-d4f6z.186a705c7d48806 | Mon, Oct 27 2025 10:46:53 pm |
| Unhealthy | Pod loki-write-1 | Readiness probe failed: HTTP probe failed with statuscode: 503 | loki-write-1.186a26893170ad83 | Mon, Oct 27 2025 10:46:18 pm |
| UpdateFailed | ExternalSecret indexd-service-creds | error processing spec.dataFrom[0].extract, err: AccessDeniedException: User: arn:aws:sts::976193266170:assumed-role/occ-gen3-occ-windtunnel-dev-external-secrets-role/external-secrets-provider-aws is not authorized to perform: secretsmanager:GetSecretValue on resource: indexd-service-creds because no identity-based policy allows the secretsmanager:GetSecretValue action on resource: ARN | indexd-service-creds.186a28a391575501 | Mon, Oct 27 2025 10:46:14 pm |

occ-windtunnel-dev

All Namespaces

Namespace: windtunnel-occ

| | | | | | | | | |
|--------------------------|--------|----------------------------|---------------------|--------------------------------|---|------------|---------|---|
| <input type="checkbox"/> | Active | ambassador-admin | 172.20.152.140:8877 | ↑ 8877/TCP | app=ambassador app.kubernetes.io/instance=windtunnel-occ app.kubernetes.io/name=ambassador release=production | Cluster IP | 27 days | ⋮ |
| <input type="checkbox"/> | Active | ambassador-service | proxy | ↑ 8080/TCP | app=ambassador app.kubernetes.io/instance=windtunnel-occ app.kubernetes.io/name=ambassador release=production | Cluster IP | 27 days | ⋮ |
| <input type="checkbox"/> | Active | arborist-service | http | ↑ http/TCP | app=arborist app.kubernetes.io/instance=windtunnel-occ app.kubernetes.io/name=arborist release=production | Cluster IP | 27 days | ⋮ |
| <input type="checkbox"/> | Active | fence-service | http | ↑ http/TCP | app=fence app.kubernetes.io/instance=windtunnel-occ app.kubernetes.io/name=fence release=production | Cluster IP | 27 days | ⋮ |
| <input type="checkbox"/> | Active | frontend-framework-service | http | ↑ 3000/TCP https ↑ 3000/TCP | app=frontend-framework app.kubernetes.io/instance=windtunnel-occ app.kubernetes.io/name=frontend-framework release=dev | Cluster IP | 27 days | ⋮ |
| <input type="checkbox"/> | Active | hatchery-service | http | ↑ http/TCP | app=hatchery app.kubernetes.io/instance=windtunnel-occ app.kubernetes.io/name=hatchery release=production | Cluster IP | 27 days | ⋮ |
| <input type="checkbox"/> | Active | indexd-service | http | ↑ http/TCP | app=indexd app.kubernetes.io/instance=windtunnel-occ app.kubernetes.io/name=indexd release=production | Cluster IP | 27 days | ⋮ |
| <input type="checkbox"/> | Active | manifestservice-service | http | ↑ 80/TCP | app=manifestservice app.kubernetes.io/instance=windtunnel-occ app.kubernetes.io/name=manifestservice | Cluster IP | 27 days | ⋮ |

v2.10.2

argo
v2.13.1+mf44fb

Applications

- Settings
- User info
- Documentation

Favorites Only

SYNC STATUS

- Unknown 0
- Synced 14
- OutOfSync 2

HEALTH STATUS

- Unknown 0
- Progressing 1
- Suspended 0
- Healthy 14
- Degraded 1
- Missing 0

LABELS

PROJECTS

CLUSTERS

NAMESPACES

AUTO SYNC

Applications

+ NEW APP SYNC APPS REFRESH APPS

Search applications... [7] [refresh]

app-of-apps ☆

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Repository: git@github.com:occ-data/occ-cscc-deploy.git

Target Rev... HEAD

Path: windtunnel/deployments

Destination: in-cluster

Namespa... argo

Created At: 09/23/2025 13:12:33 (a month ago)

Last Sync: 09/30/2025 15:14:04 (a month ago)

[SYNC] [REFRESH] [DELETE]

argo-cd ☆

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Repository: https://argoproj.github.io/argo-helm

Target Rev... 7.7.5

Chart: argo-cd

Destination: in-cluster

Namespa... argo

Created At: 09/23/2025 13:12:35 (a month ago)

Last Sync: 09/30/2025 14:30:57 (a month ago)

[SYNC] [REFRESH] [DELETE]

aws-load-balancer-controller ☆

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Repository: https://aws.github.io/eks-charts

Target Rev... 1.12.0

Chart: aws-load-balancer-controller

Destination: in-cluster

Namespa... kube-system

Created At: 09/23/2025 13:12:35 (a month ago)

Last Sync: 09/23/2025 13:13:22 (a month ago)

[SYNC] [REFRESH] [DELETE]

cert-manager ☆

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Repository: https://charts.jetstack.io

Target Rev... v1.17.0

Chart: cert-manager

Destination: in-cluster

Namespa... cert-manager

Created At: 09/23/2025 13:12:35 (a month ago)

Last Sync: 09/23/2025 13:13:28 (a month ago)

[SYNC] [REFRESH] [DELETE]

cert-manager-manifests ☆

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Repository: git@github.com:occ-data/occ-cscc-deploy.git

Target Rev... HEAD

Path: windtunnel/platform/cert-manager

Destination: in-cluster

Namespa...

Created At: 09/23/2025 13:12:35 (a month ago)

Last Sync: 09/23/2025 13:13:45 (a month ago)

[SYNC] [REFRESH] [DELETE]

external-dns ☆

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Repository: https://kubernetes-sigs.github.io/external-dns/

Target Rev... 1.15.2

Chart: external-dns

Destination: in-cluster

Namespa... external-dns

Created At: 09/23/2025 13:12:35 (a month ago)

Last Sync: 09/30/2025 13:15:17 (a month ago)

[SYNC] [REFRESH] [DELETE]

external-secrets ☆

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Repository: https://charts.external-secrets.io

Target Rev... 0.14.3

Chart: external-secrets

Destination: in-cluster

Namespa... external-secrets

Created At: 09/23/2025 13:12:35 (a month ago)

Last Sync: 09/23/2025 13:13:21 (a month ago)

[SYNC] [REFRESH] [DELETE]

external-secrets-manifests ☆

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Repository: git@github.com:occ-data/occ-cscc-deploy.git

Target Rev... HEAD

Path: windtunnel/platform/external-secrets

Destination: in-cluster

Namespa...

Created At: 09/23/2025 13:12:35 (a month ago)

Last Sync: 09/30/2025 13:54:43 (a month ago)

[SYNC] [REFRESH] [DELETE]

grafana-manifests ☆

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Repository: git@github.com:occ-data/occ-cscc-deploy.git

Target Rev... HEAD

Path: windtunnel/platform/grafana

Destination: in-cluster

Namespa...

Created At: 09/23/2025 13:12:35 (a month ago)

Last Sync: 09/23/2025 13:12:38 (a month ago)

[SYNC] [REFRESH] [DELETE]

jaeger ☆

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Repository: https://jaegertracing.github.io/helm-charts

Target Rev... 3.4.0

Chart: jaeger

Destination: in-cluster

Namespa... jaeger-system

Created At: 09/23/2025 13:12:35 (a month ago)

Last Sync: 09/23/2025 13:12:46 (a month ago)

[SYNC] [REFRESH] [DELETE]

loki ☆

Project: default

Labels:

Status: ○ Progressing ✔ Synced

Repository: https://grafana.github.io/helm-charts

Target Rev... 6.25.1

Chart: loki

Destination: in-cluster

Namespa... cattle-monitoring-system

Created At: 09/23/2025 13:12:35 (a month ago)

Last Sync: 09/30/2025 14:30:41 (a month ago)

[SYNC] [REFRESH] [DELETE]

loki-manifests ☆

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Repository: git@github.com:occ-data/occ-cscc-deploy.git

Target Rev... HEAD

Path: windtunnel/platform/loki

Destination: in-cluster

Namespa...

Created At: 09/23/2025 13:12:35 (a month ago)

Last Sync: 09/23/2025 13:12:36 (a month ago)

[SYNC] [REFRESH] [DELETE]

pgadmin4 ☆

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Repository: https://helm.runix.net

Target Rev... 1.36.0

Chart: pgadmin4

Destination: in-cluster

Namespa... pgadmin4

Created At: 09/23/2025 13:12:35 (a month ago)

Last Sync: 09/30/2025 15:34:50 (a month ago)

[SYNC] [REFRESH] [DELETE]

promtail ☆

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Repository: https://grafana.github.io/helm-charts

Target Rev... 6.16.6

Chart: promtail

Destination: in-cluster

Namespa... cattle-monitoring-system

Created At: 09/23/2025 13:12:35 (a month ago)

Last Sync: 09/23/2025 13:12:42 (a month ago)

[SYNC] [REFRESH] [DELETE]

windtunnel-occ ☆

Project: default

Labels:

Status: ♥ Degraded ○ OutOfSync ✖ Sync failed

Repository: https://github.com:uoc-ucds/gen3-helm.git

Target Rev... gen3-0.1.63

Path: helm/gen3

Destination: in-cluster

Namespa... windtunnel-occ

Created At: 09/30/2025 15:14:04 (a month ago)

Last Sync: 10/21/2025 14:47:03 (6 days ago)

[SYNC] [REFRESH] [DELETE]

argo
v2.13.1+14f54ef8

- Applications
- Settings
- User Info
- Documentation

NAME

NAME

KINDS

KINDS

SYNC STATUS

- Synced 105
- OutOfSync 20

HEALTH STATUS

- Healthy 65
- Progressing 2
- Degraded 8
- Suspended 0
- Missing 0
- Unknown 0

- DETAILS
- DIFF
- SYNC
- SYNC STATUS
- HISTORY AND ROLLBACK
- DELETE
- REFRESH

APP HEALTH @

Degraded

SYNC STATUS @

OutOfSync from gen3-0.1.63 (bdf2866) and (1) more

Auto sync is enabled.
Author: EliasCastro23 <109446148-EliasCastro23@users.noreply.github.com>
Comment: adding 'gen3 integration tests' to helm charts. (#241)

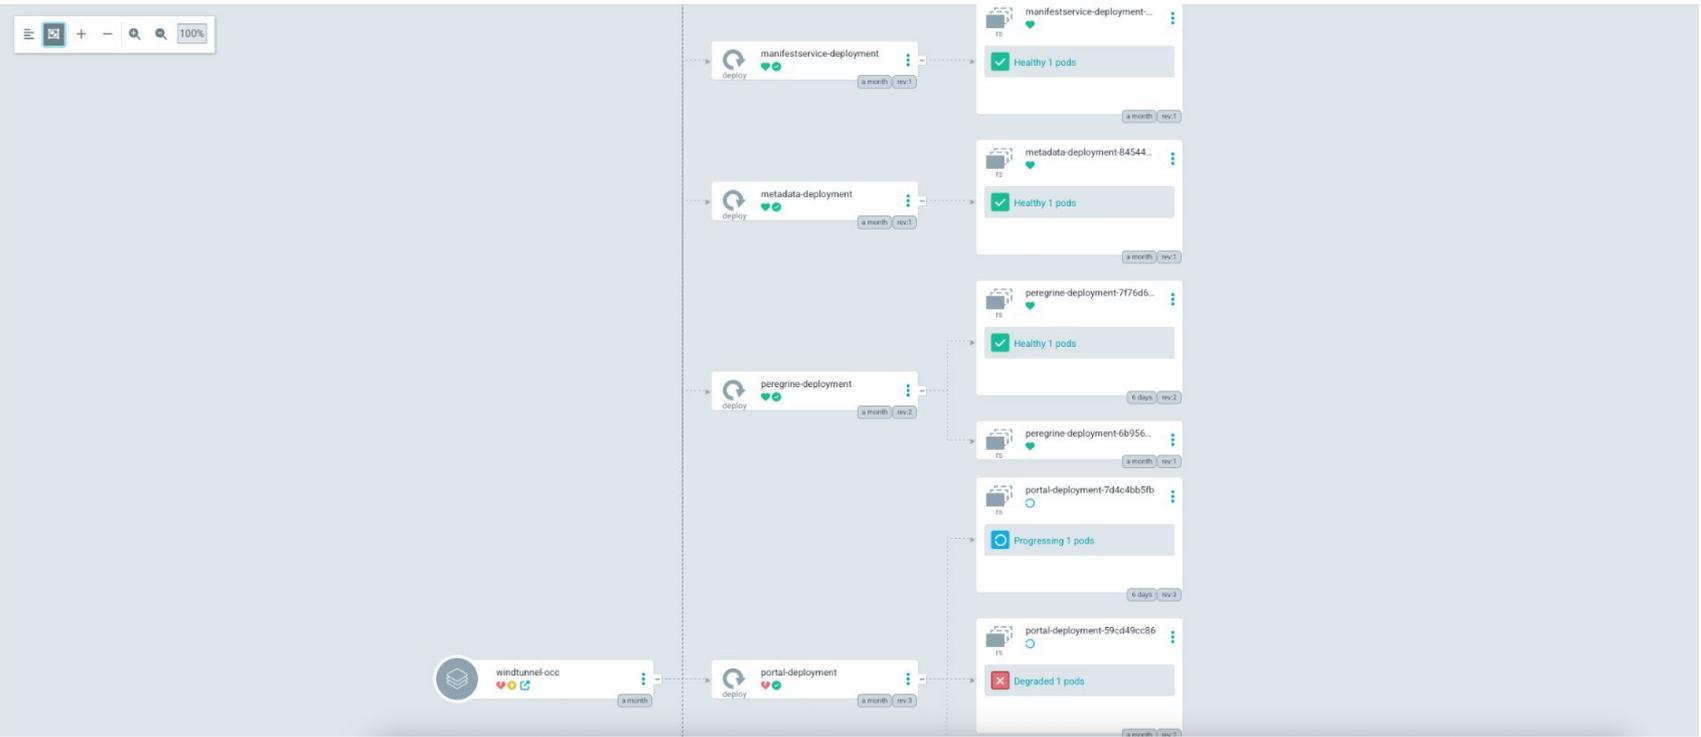
LAST SYNC @

Sync failed to bdf2866 and (1) more

Failed 6 days ago (Tue Oct 21 2025 14:47:03 GMT-0500)
Author: EliasCastro23 <109446148-EliasCastro23@users.noreply.github.com>
Comment: adding 'gen3 integration tests' to helm charts. (#241)

APP CONDITIONS

1 Error



Key Benefits

- Works with OCC paradigm of separate AWS accounts
- Developed with best security practices; ensures each commons is properly secured
- WAF allows simple control over who can access server
- Reduces complexity; Rancher gives central view of all data commons
- Argo CD built-in, continuous deployment for all data commons
- Uses existing Gen3 Helm charts
- Logging/monitoring built into each commons environment (in development)

Challenges

- Automating secrets management
- Microservice-specific configurations
- Implementation of lesser-known parts of Gen3 e.g. ssjdispatcher
- Incorporating latest Gen3 developments
- Improving developer documentation of Gen3 ecosystem
- Covers similar ground as gen3-terraform
- Currently only works with AWS

Next Steps

- Record further steps to create SOP document
- Ongoing Commons development & configuration (BLOODPAC)
 - Setting up a development environment
 - Data transfer from old production environment to development
 - Smooth transition live service from old environment to new setup
- Gradually onboard Commons into CSOC framework
- Establish unified governance model across commons

Contact support@occ-data.org for any questions and concerns

Using Gen3 as a platform and running additional services such as GEARBOX

Paul Murdoch, Data for the Common Good,
University of Chicago

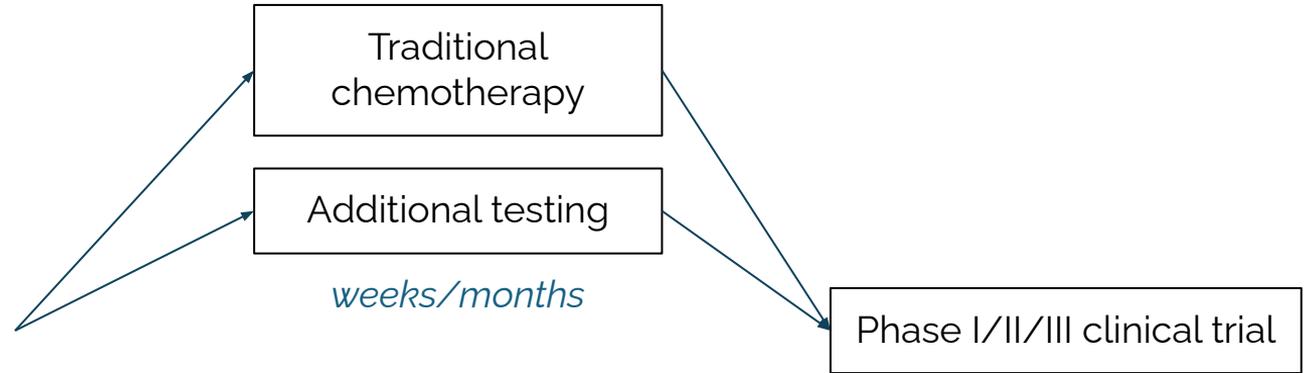
What is GEARBOX?



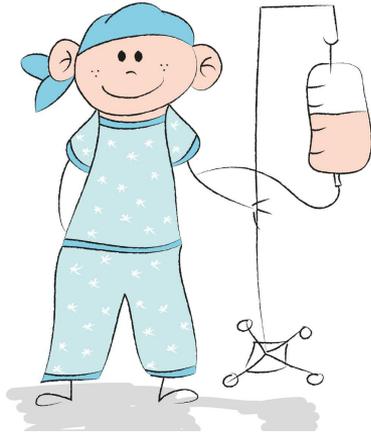
Relapsed patients struggle to find therapies



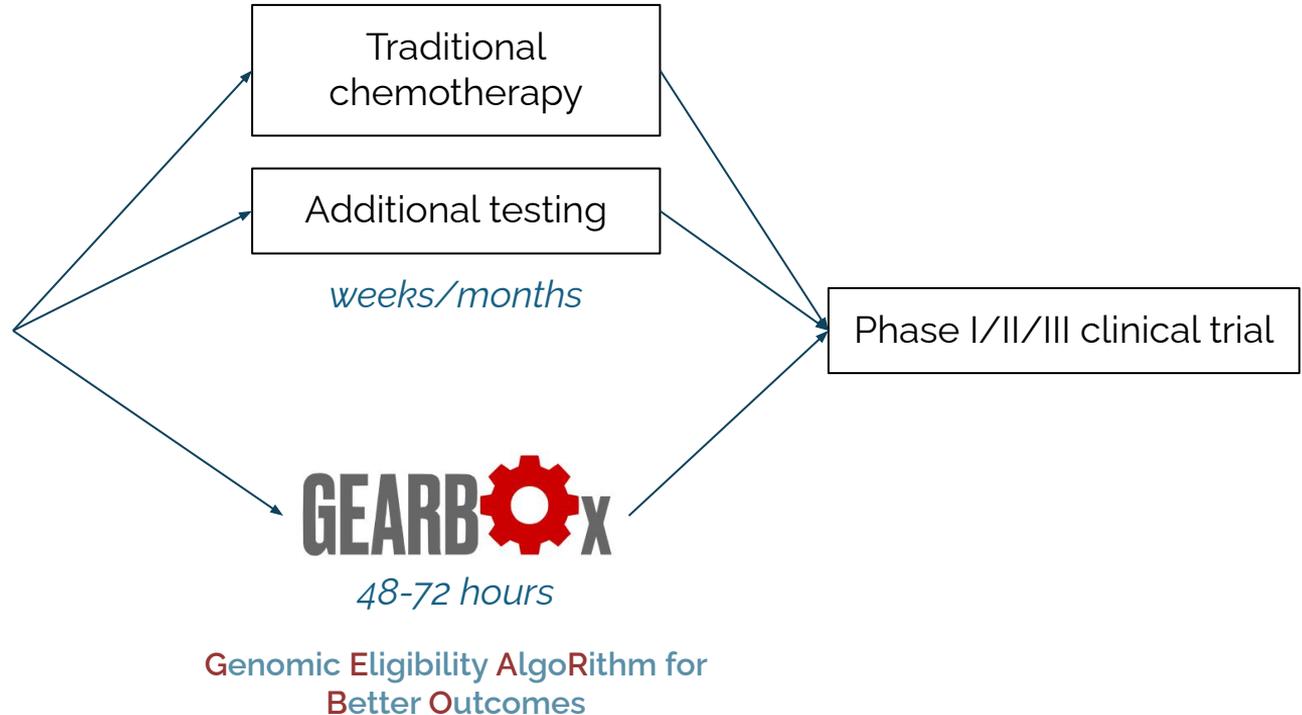
Child with relapsed AML



Relapsed patients struggle to find therapies



Child with relapsed AML



Goals for GEARBOx

1. Find a trial for a relapsed / refractory leukemia patient within 72 hours
2. Very quick and easy to use
3. For physicians and families
4. Designed to lower the friction to finding potentially compatible trials
5. Scalable and sustainable
6. Path to multisource data integration
7. Designed for on- and off-premises use

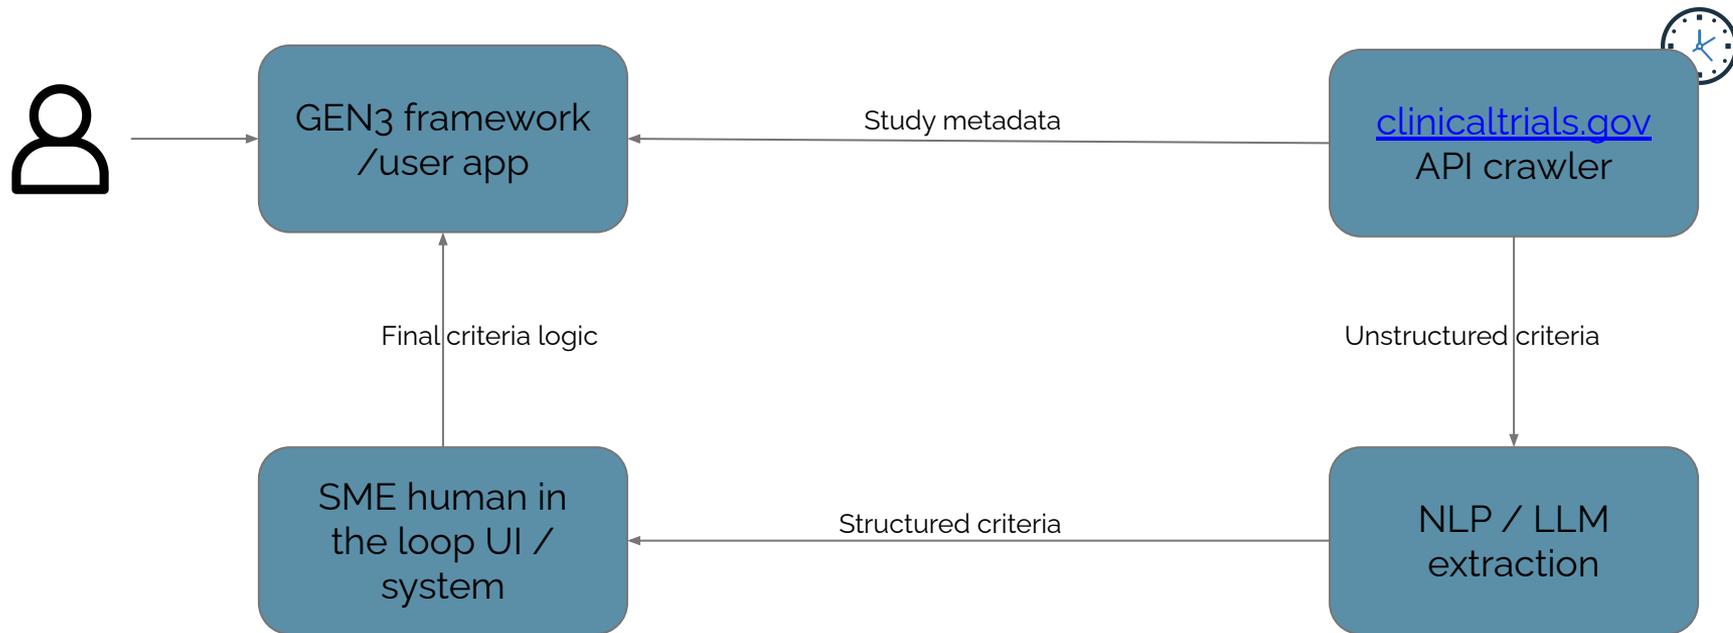
GEARBOX User Interface

Patient characteristics
Disease characteristics
Lab tests
Genomic testing

The screenshot shows a web browser window titled "GEARBOX by LLS PedAL". The page has a red header with the GEARBOX logo, navigation links for "ABOUT GEARBOX" and "USER GUIDE", and a user greeting "Hello, demo" with a "LOG OUT" link. The main content is divided into two columns. The left column, titled "PATIENT INFORMATION", contains several sections: "Demographics" with two input fields for "What is the patient's current age (in years)?" and "What is the patient's current weight (in kg)?"; "Disease" with a dropdown arrow; "Prior treatment" with a dropdown arrow; "Organ function" with a dropdown arrow; and "Biomarkers" with a dropdown arrow. A "RESET" button is located below these sections. The right column, titled "OPEN TRIALS", shows a list of trials with a "Matched (0)" and "Undetermined (9)" count at the top, and an "Unmatched (0)" count at the bottom. The trials listed are: APAL2020SC, AAML2112, APAL2020B, PEPN2113, APAL2020D, AAML2020E, APAL2020F, APAL2020G, and T2017-002. Each trial name is followed by a circular icon and a dropdown arrow. The footer of the page includes logos for "THE UNIVERSITY OF CHICAGO BIOLOGICAL SCIENCES", "PEDIATRIC CANCER DATA COMMONS", and "LLS PedAL".

Clinical trials
Information about enrollment
Study locations

GEARBOX System



GEN3 for GEARBOX



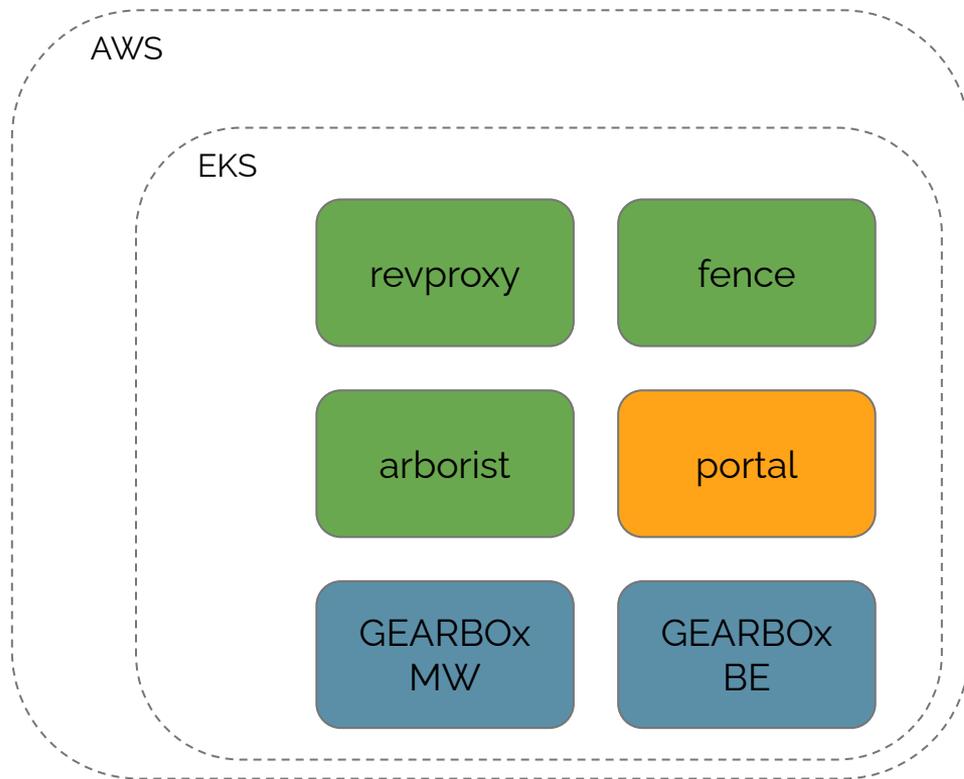
GEN3 for GEARBOX

Used:

- Fence
- Arborist
- Revproxy
- Main deployment frameworks (gen3-helm and cloud automation)

Updated / Added:

- New frontend / portal
- 2 new services to handle backend operations
- Deployment changes to support the new services



Adding Custom Services To Gen3

