LARA

AWS Reference Architecture

LABYRINTH LABS

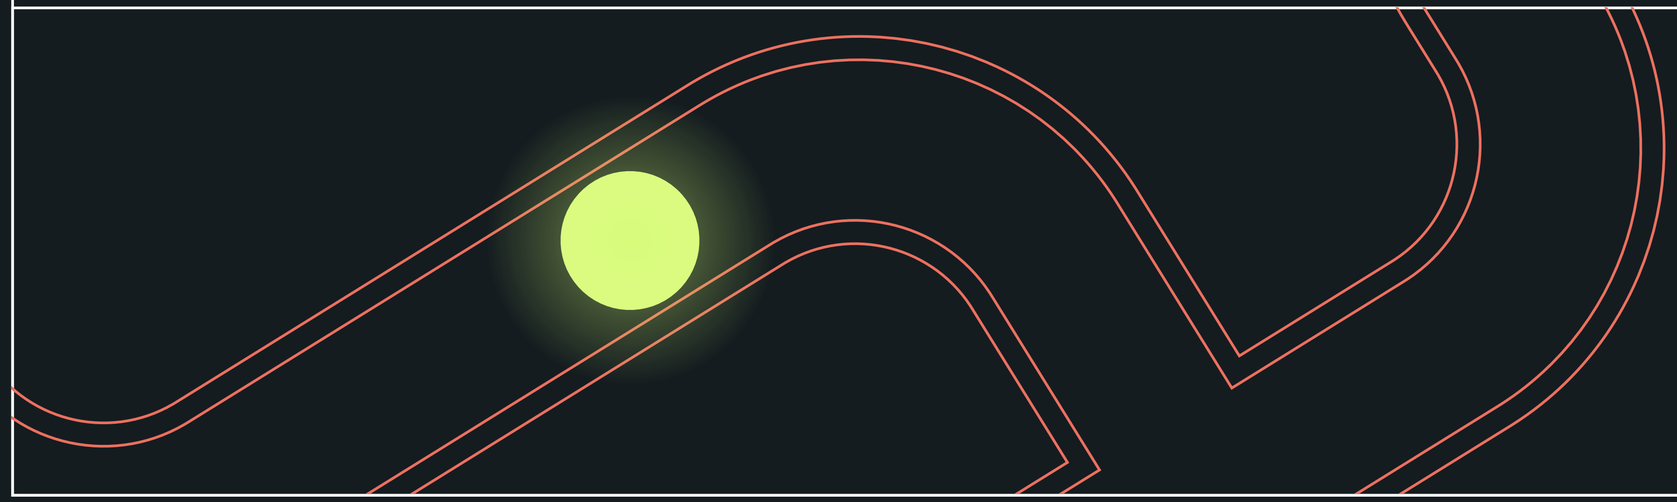

Content

Technical Overview

LARA is designed to help companies build infrastructure for hosting their services quickly & reliably, according to AWS best practices and based on open source.

All of the componentes mentioned below are defined using Infrastructure as Code, allowing for further customization and reusability in different environments.

LARA is composed of multiple components and modules, making it easily extendable.

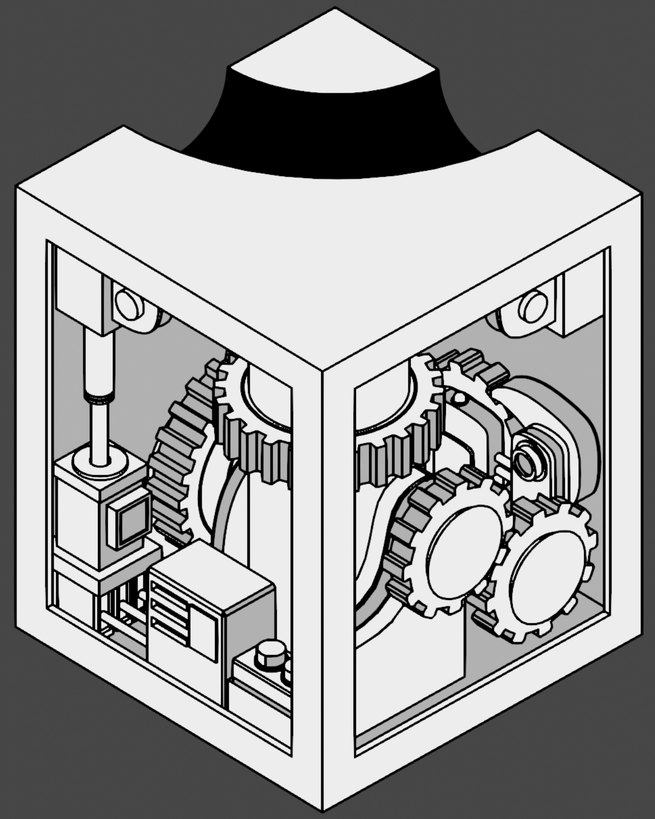

LARA's general overview

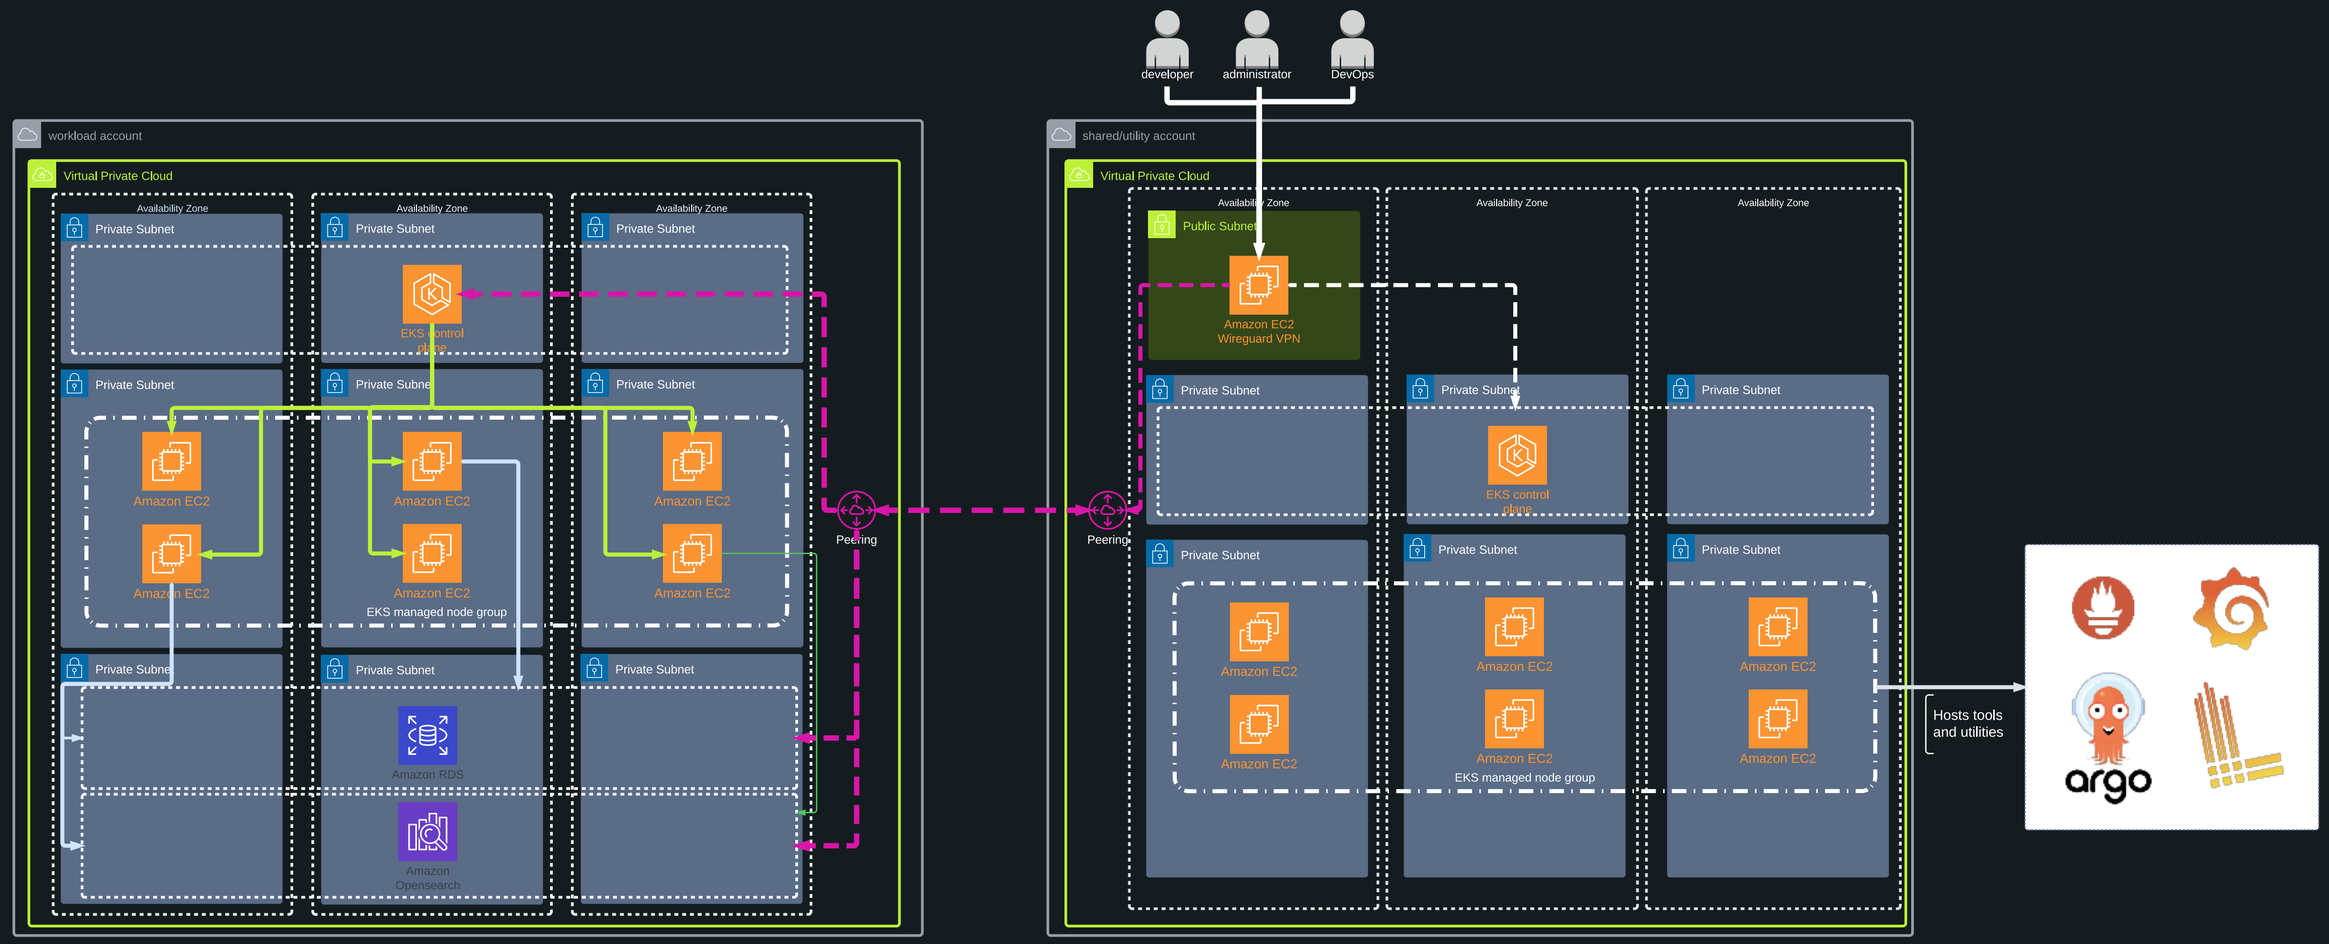

AWS Organization and Access Management

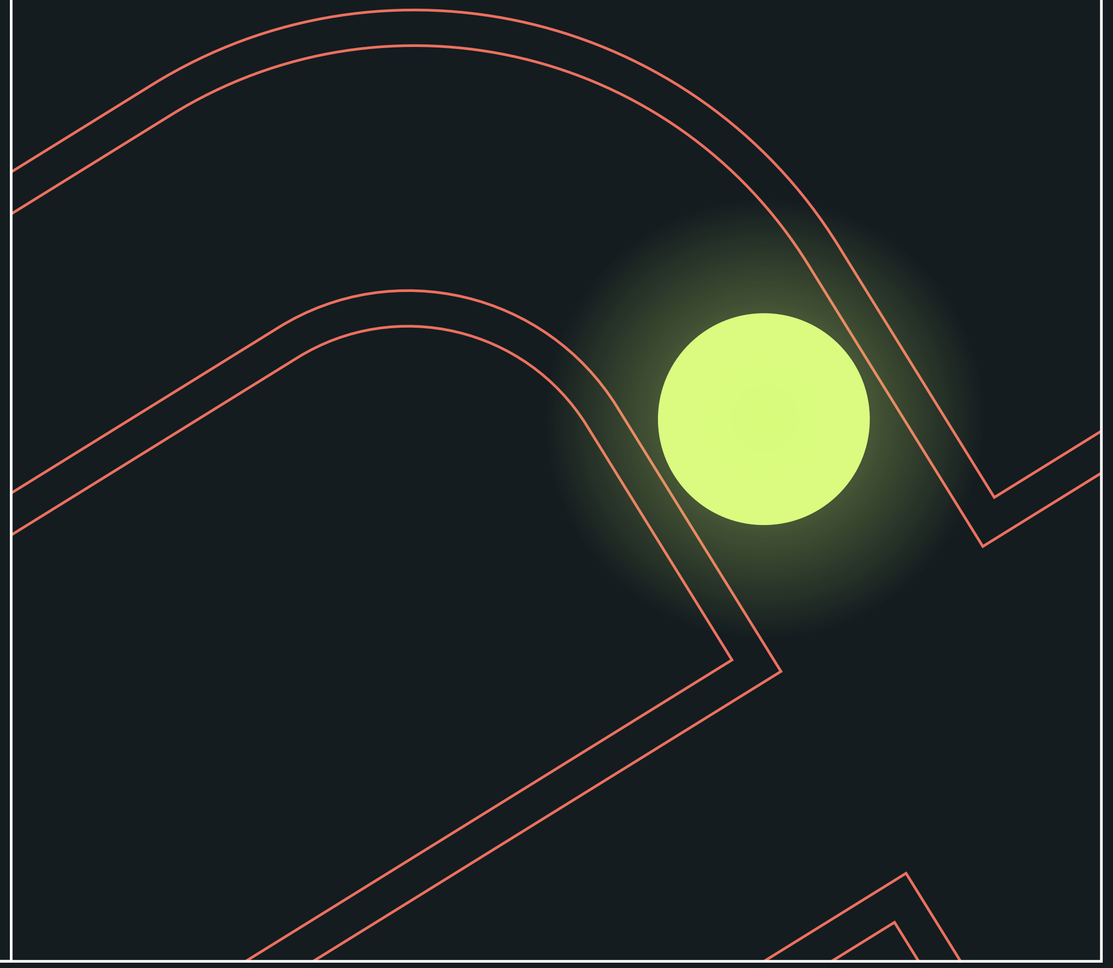

Centrally govern all your environments, manage access rights, permissions and policies

AWS organization composed from multiple AWS accounts, which are used to separate workloads from different environments and projects.

AWS SSO is used to manage access for individual users. SSO can be used with external Identity Providers such as AzureAD, Google Workspace, etc. or used as a standalone solution for user management within AWS.

With AWS Cloudtrail, we're able to gather all of the actions performed in the AWS organization and archive them for auditing purposes.

AWS SCP policies can be used to further manage permissions within the organization.

VPC and Related Services

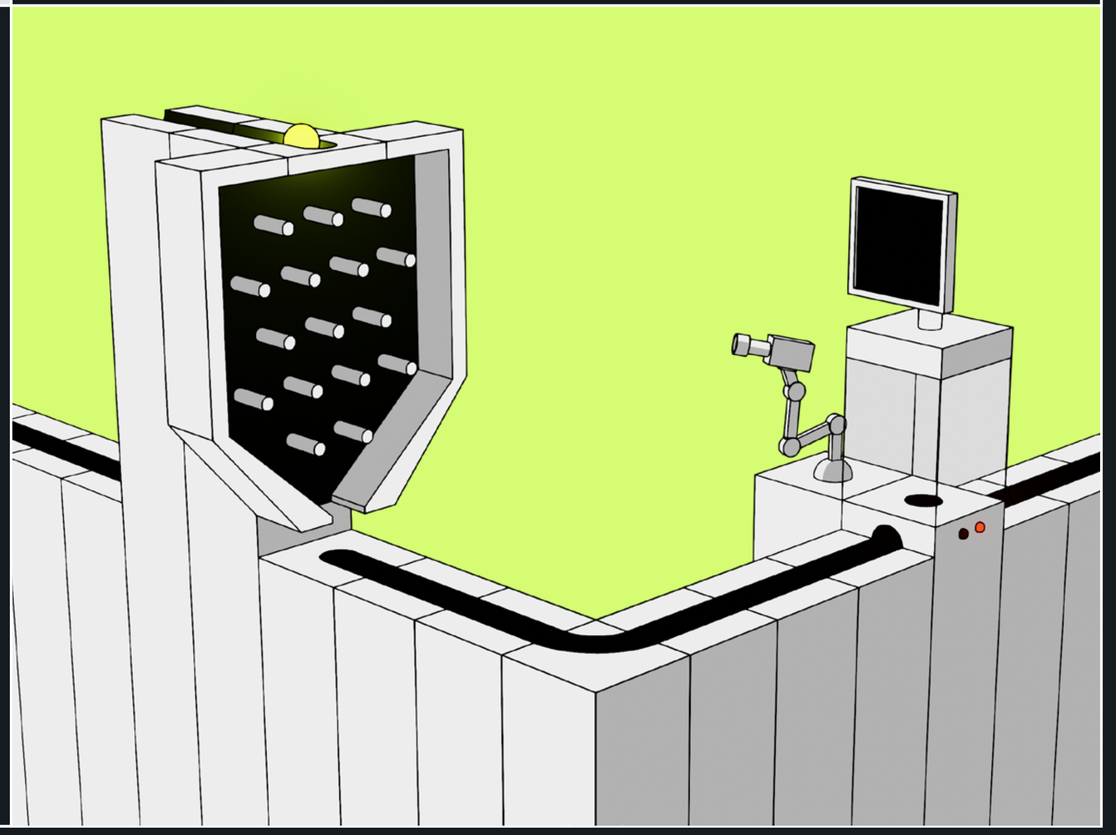

Network foundation, app and edge networking, hybrid connectivity and security

VPCs are used as a base for the components further below. The setup includes all of the necessary VPC components such as IGW, NAT, separate subnets for different types of workload (private, public, EKS control plane, other services)

The component allows spawning of VPCs in different regions with the option to interconnect them using VPC peering or AWS Transit Gateways.

Wiregard VPN with automatic provisioning is used to allow users to connect to the AWS internal resources hosted in private subnets securely and with little effort.

The setup also includes all of the resources to enable Encrypted SSM access with audit logging, to allow you to connect to individual EC2s without having to expose anything to the internet

Kubernetes

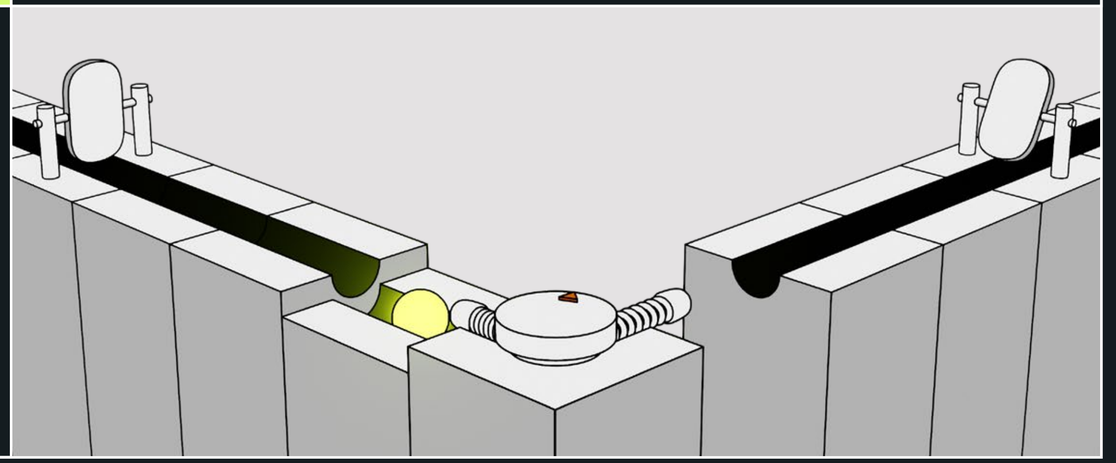

Operate kubernetes on your own infrastructure, ultimate scaling and cost savings

Battle hardened EKS setup for hosting containerized workloads without additional complexity.

Options to run Kubernetes nodes as either self-managed node group, EKS managed node group or Fargate - all depending on your needs.

Depending on your workload, spot instances are used to run all of the Kubernetes workload to save as much as ~60% of the EC2 costs

Infinite scaling options for your workloads to adapt to the usage patterns of your apps and ever growing user base.

Complex addon system

Addons are used to complement the default EKS functionality by introducing a further level of automation to solve common challenges when deploying customer services.

Addons are managed using ArgoCD deployed within the same cluster. This gives the operator an option to manage all of the addons in one place, seeing their status and changes that are to be made to them.

All of the addons mentioned below come as individual modules that can be added/ removed based on your current needs.

Dependent resources that are needed for the addon to function properly (IAM roles and permissions, etc.) are bundled in the modules. You don't have to worry about setting up any additional infrastructure.

Addons that we offer:

Cluster-autoscaler

Cluster autoscaler is a tool to automatically adjust the size of the cluster based on current demand - adding and removing nodes as needed.

External-dns

External-dns is used to automatically provision DNS records based on Ingress resources created within the cluster.

Cert-manager

Similar to external-dns, cert-manager is used to provision and manage SSL certificates used to secure in-transit traffic between the clients and your services.

Complex addon system

Addons that we offer:

Metrics-server

Metrics Server is a scalable, efficient source of container resource metrics for Kubernetes built-in autoscaling pipelines.

Node-problem-detector

node-problem-detector is a daemon running on each node that makes various node problems visible to the upstream layers in the cluster management stack.

External-secrets

External-secrets can be used to pull secrets from various secrets stores such as AWS Secrets Manager, ParameterStore, Hashicorp Vault and more with the ability to configure fine grained access policies.

Vector agent

Vector is used to gather all of the necessary logs from the infrastructure itself and your services and ship them to a log collection tool of your choice.

Linkerd

Ultra light service mesh that adds more security, observability and reliability to your services within the Kubernetes cluster.

ArgoCD + Argo rollouts

Powerful CD tool for managing deployments of your services to the cluster with the option to perform advanced deployment strategies such as canary deployments or blue-green deployments.

Gitlab CI runners

A module to dynamically create self-hosted Gitlab CI runners within your Kubernetes environment for running jobs in your CI/CD pipelines.

Github Actions runners

A module to create and manage self-hosted Github Actions runners with the ability to dynamically create and scale runners based on current workload or schedule.

Complex addon system

Addons that we offer:

Crossplane

Allows you to manage parts of the infrastructure using Kubernetes resources. Can be used to bundle creation of some infrastructure resources (database users, IAM roles, secrets) in the application deployment process.

Keda

Event driven Kubernetes autoscaler. Can be used for custom auto scaling configuration, based on events, cron schedules and more. Can be particularly useful if you need to pre-scale the infrastructure before planned events or schedules.

Efs-csi-driver and Ebs-csi-driver

Kubernetes CSI drivers to enable use of EFS and EBS within the cluster with the ability to dynamically provision persistent storage for your application workload.

LARA's addon ecosystem

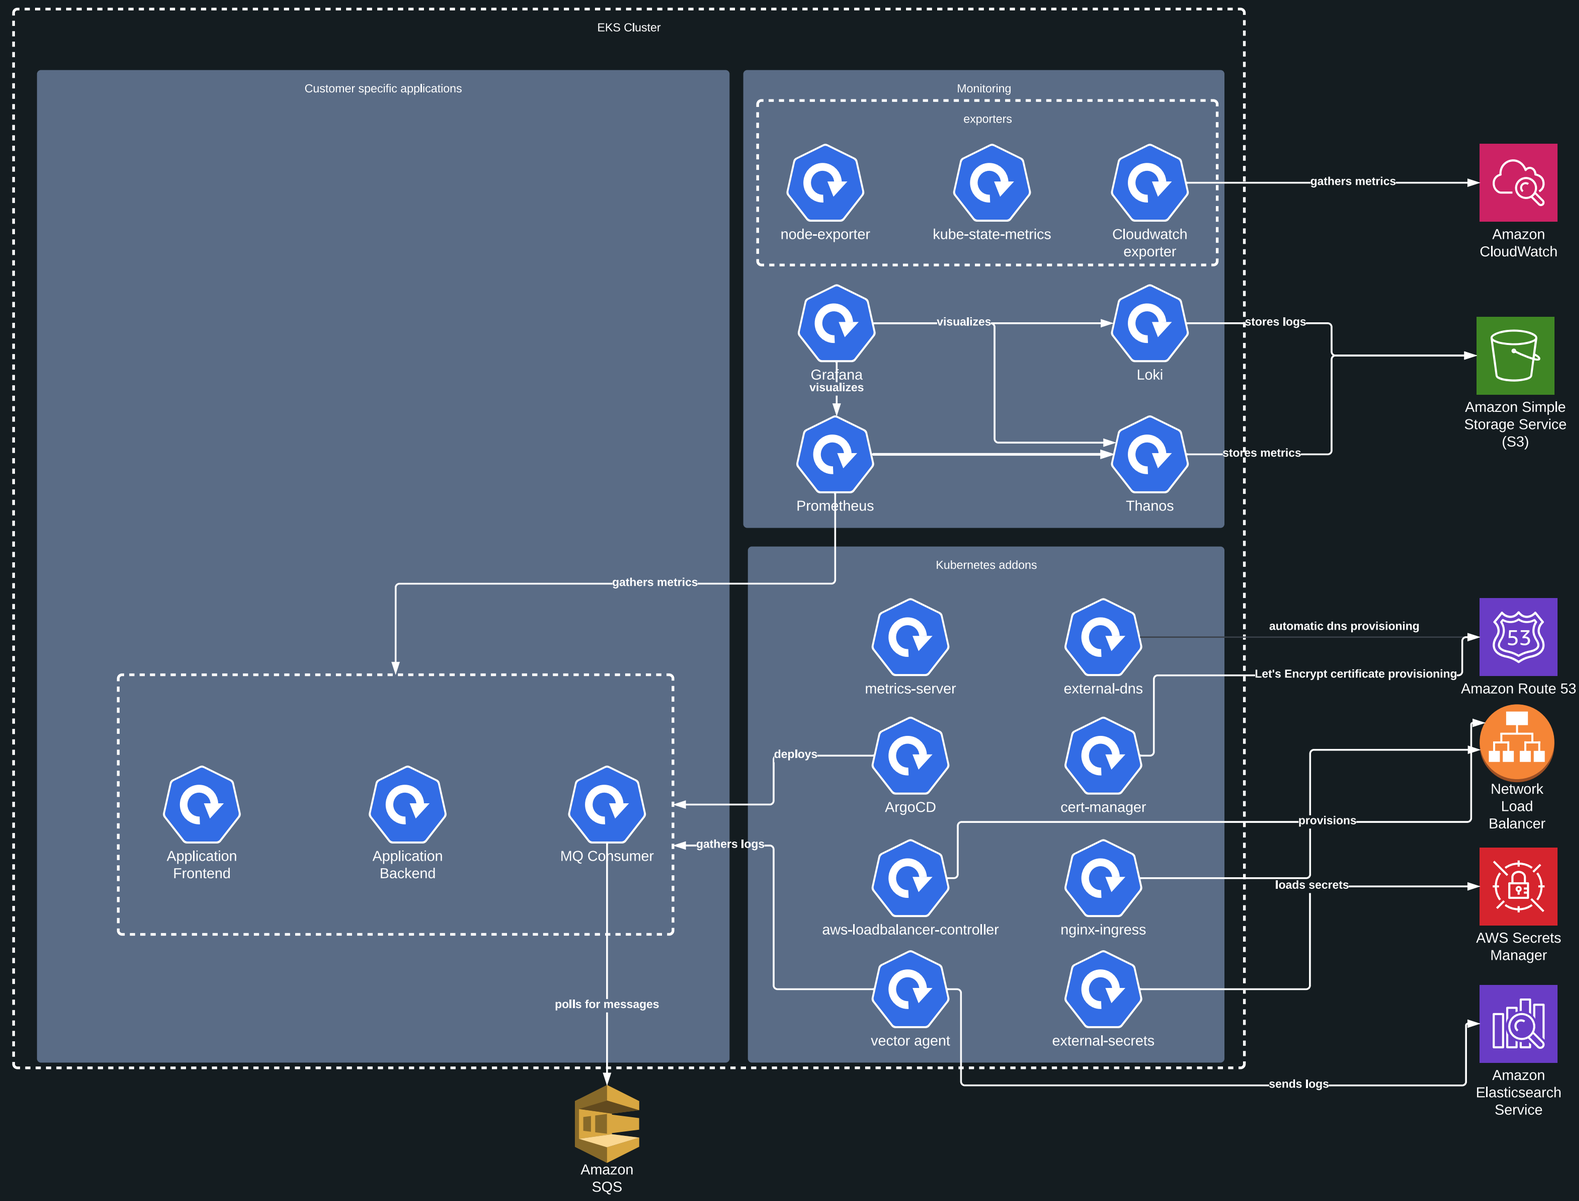

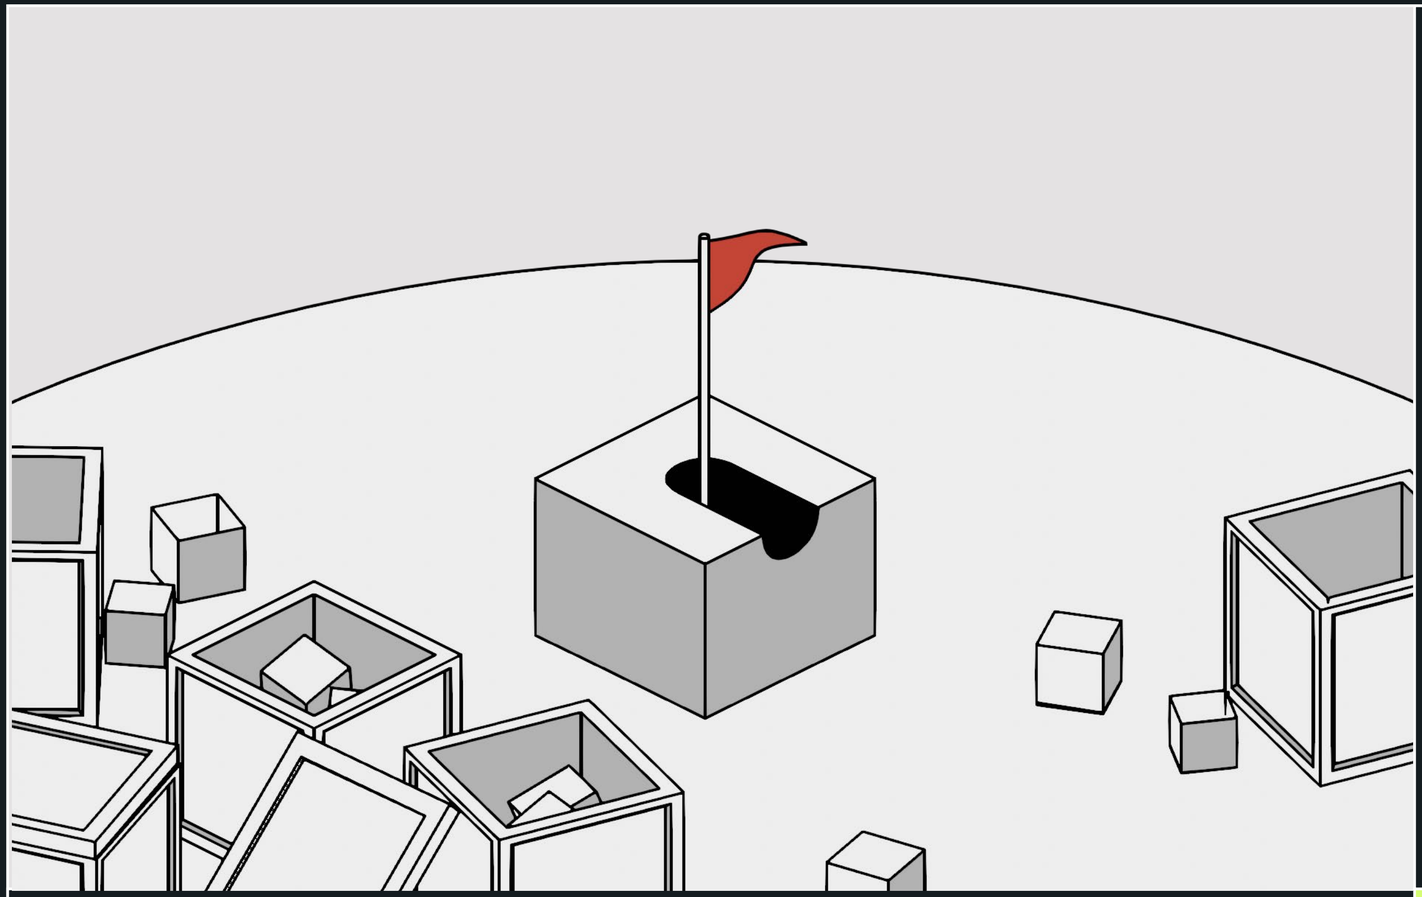

Observability

Grafana, Loki, Thanos, Tempo, Elastic. You name it, we have it.

LARA offers a complex observability system for gathering metrics, logs and traces. It allows you to gain a global view of the state of all infra-structure resources as well as deep dive view in case you need to trace or debug some issues.

Grafana is used as the visualization and alerting tool for all of the gathered metrics, logs and traces, offering a single point-of-view for all of your environments and applications.

LARA's observability stack comes with a set of ready made Grafana dashboards and alerts that can be further extended with the views of your application workload for giving you as much info as you might need.

Prometheus with Thanos is used to gather all of the metric data from individual infrastructure components and your services.

↳ Thanos allows you to view metrics from multiple Prometheus instances and solves some common problems that come with Prometheus such as High availability and limited metric retention.

For log aggregation LARA supports both Grafana Loki and AWS OpenSearch, choosing the solution that is the best fit for your needs.

Distributed traces can be ingested in various formats by Grafana Tempo and visualized in Grafana with the ability to correlate the data with metrics and logs.

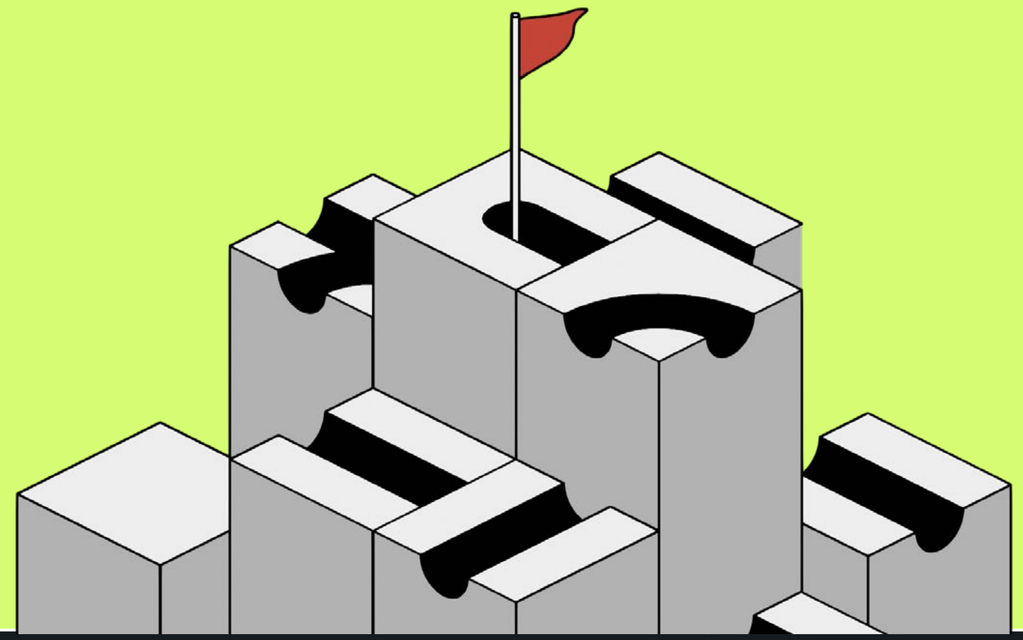

Databases and Storage

Choose from plenty of purpose-built database engines to achieve the best efficiency

LARA includes modules for setting up databases and different storage solutions.

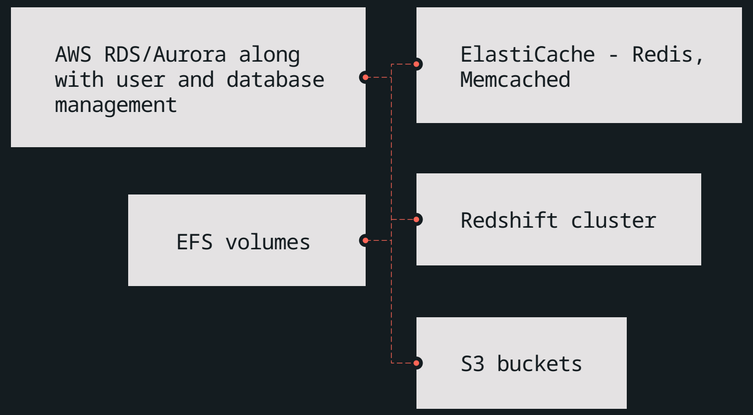

We utilise native AWS solutions for providing means of backing your data up using various strategies - point-in-time restores for your databases, object versioning or regular scheduled based backups. Everything encrypted in-transit and at-rest.

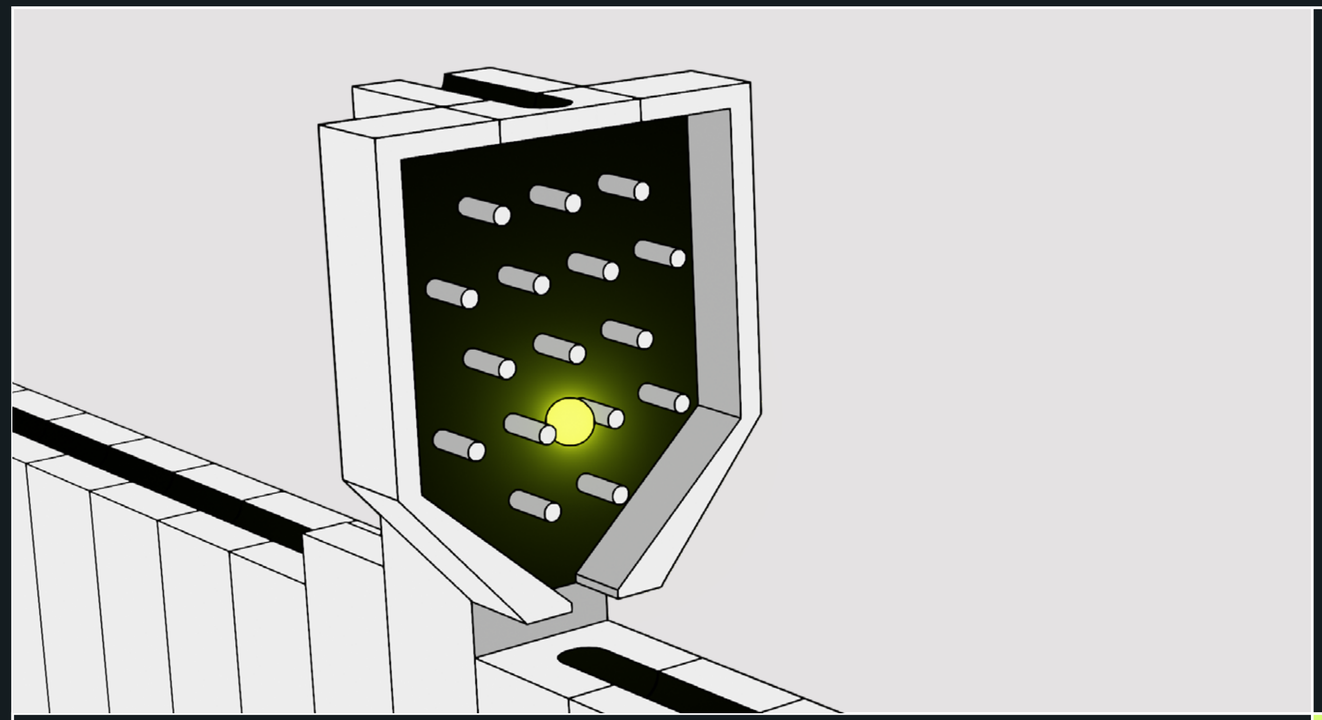

Messaging & Streaming

Managed services for messaging and streaming for modern application architecture.

LARA comes with modules for different auto scalable messaging and steaming platforms supporting high throughputs such as:

`RabbitMQ` `AWS MSK Kafka` `SQS/SNS`

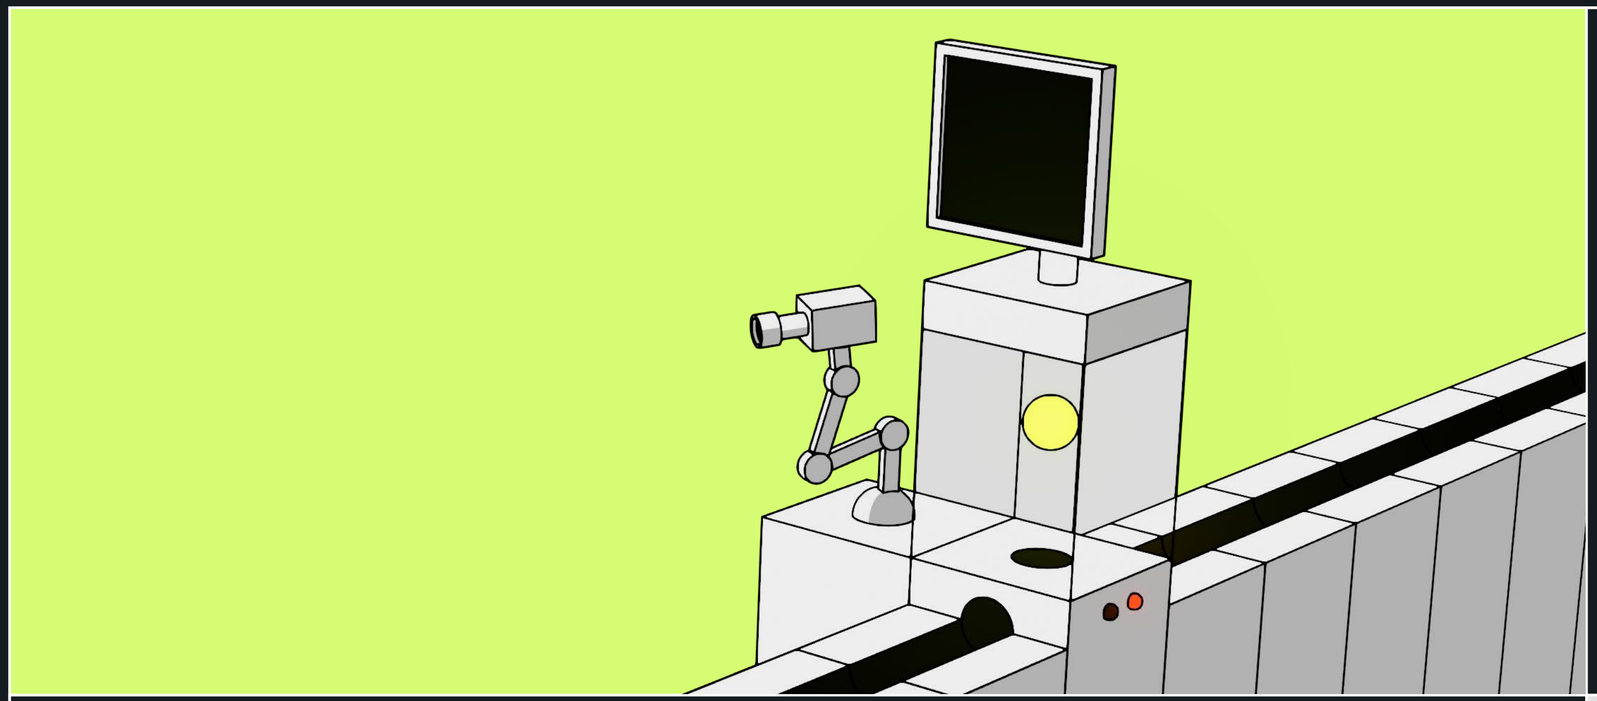

Security and Compliance

Workload isolation, access policies, data governance, security scans.

LARA makes enterprise level security features available to everyone

- Workload isolation on multiple levels
- Infrastructure security scans, SAST checks and more
- Fine grained access policies
- Data governance concepts

Advanced Networking

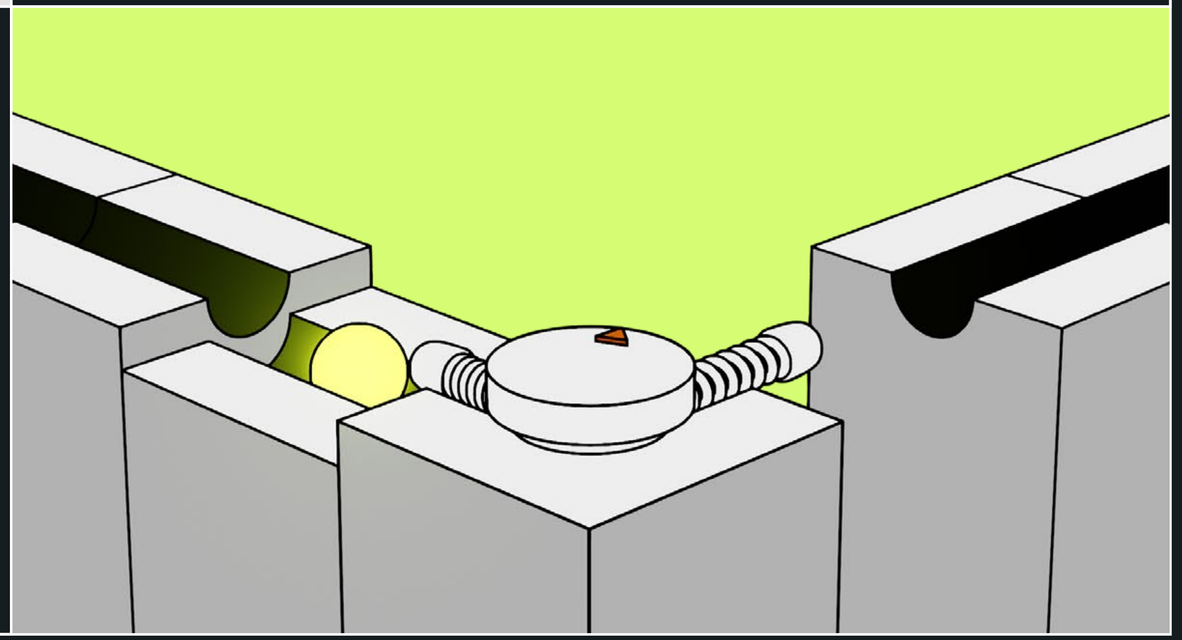

Enable service mesh, connect 3rd parties or on-prem sites and allow world-wide content delivery.

LARA includes advanced networking features such as:

Service mesh - giving you the ability to have greater control and security over the in-cluster traffic.

Fully automated and ready to use VPN using Wireguard.

Options to securely connect to your existing on-premises network.

Content Delivery Networks using Cloudfront/Cloudflare for cheaper and faster serving of static content and advanced protection from threats.

Continuous Deployment

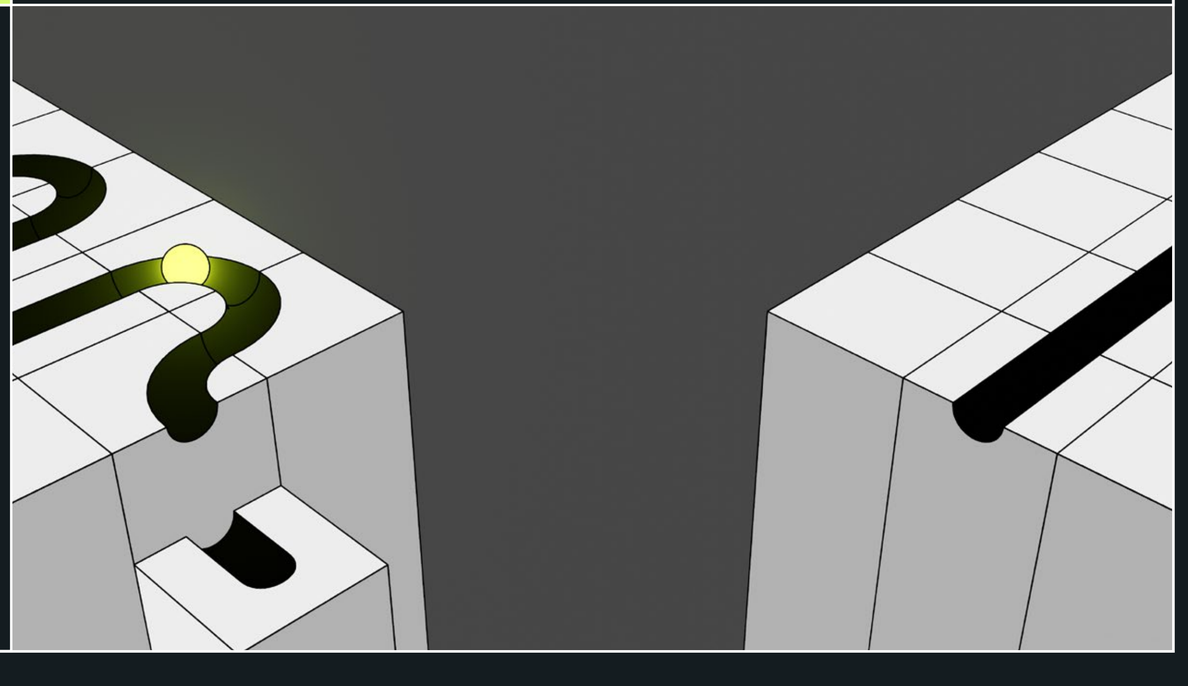

Provide your team with GitOps based workflows supporting advanced deployment methods

CI/CD automation is a major part of LARA. We've included these components to kickstart your Cloud deployments:

ArgoCD with Argo Rollouts for advanced GitOps based deployments.

Gitlab Actions and Gitlab CI private Kubernetes runners for improved security.

ECR repositories to store your Docker images.

Feedback Loop to your communication tools, so you always know what's happening with your deployments.

LARA

phone

+421 221 020 694

mail

contact@lablabs.io

web

lablabs.io

LABYRINTH LABS