

AWS
re:Invent



GPS307

Containers on AWS: Architecting software delivery platforms

Mikhail Shapirov

Sr. Partner Solutions Architect, Containers

AWS

Agenda

Enterprise setup – organization and environments

Day one – provisioning

Tenant onboarding

Workload onboarding

Operational strategies (accounts and regions)

What is next?

Enterprise DevOps



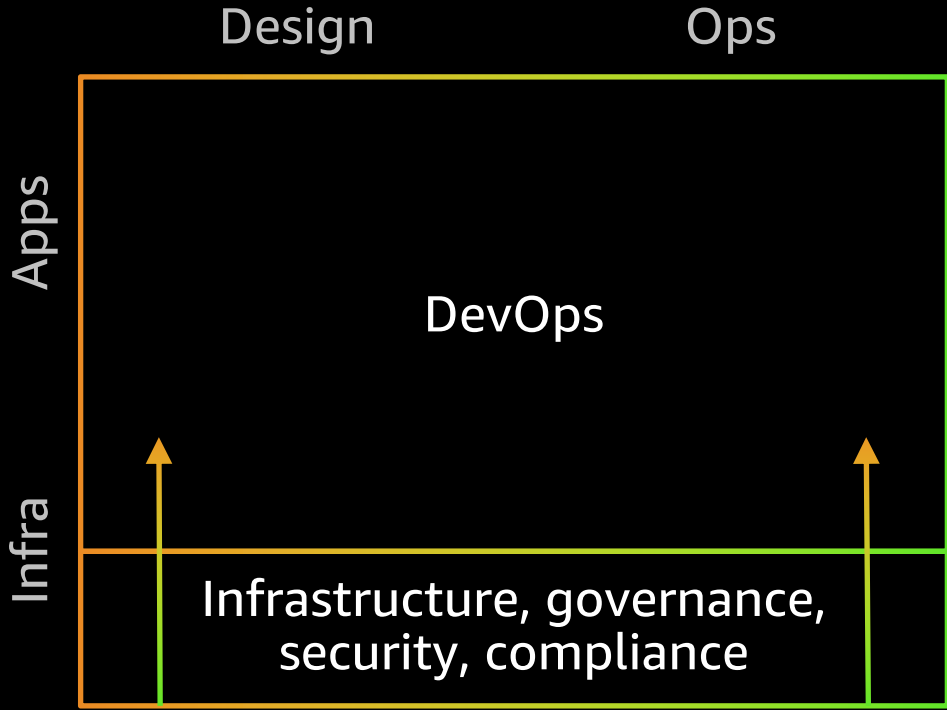
Traditional IT



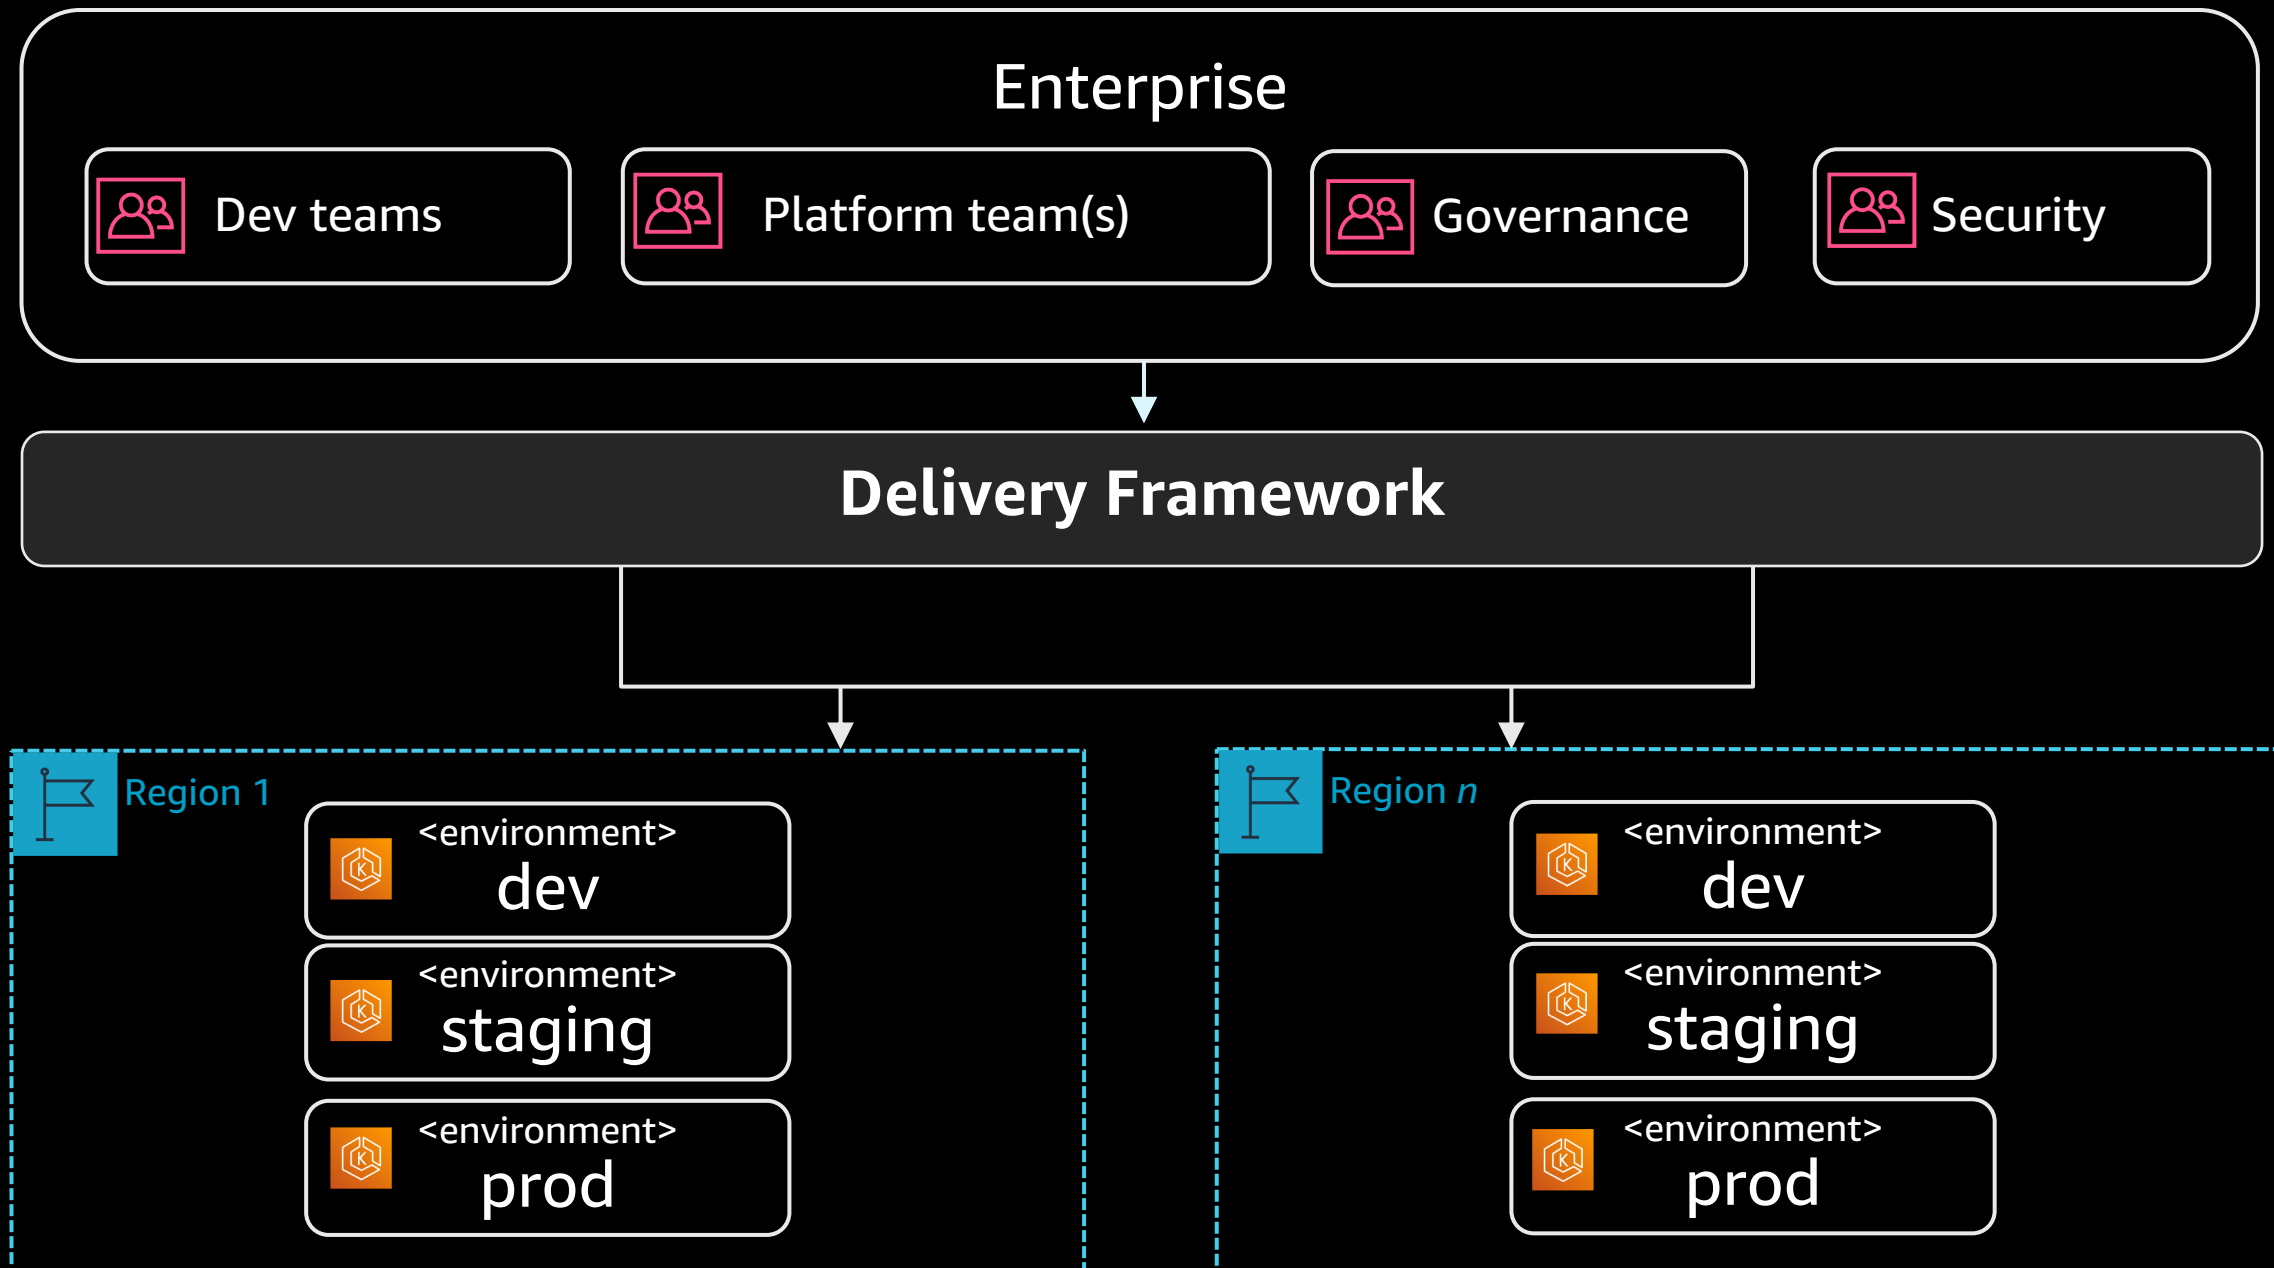
Enterprise DevOps



DevOps (perception)



Enterprise software delivery




Platform as a service (PaaS)


- Provider manages entire platform infrastructure
- You are abstracted from lower-level details of the environment




API wrapper




Workload manager




Customer endpoint



Deployment pipeline



Monitoring service



Metering service

Region scope



Amazon API Gateway



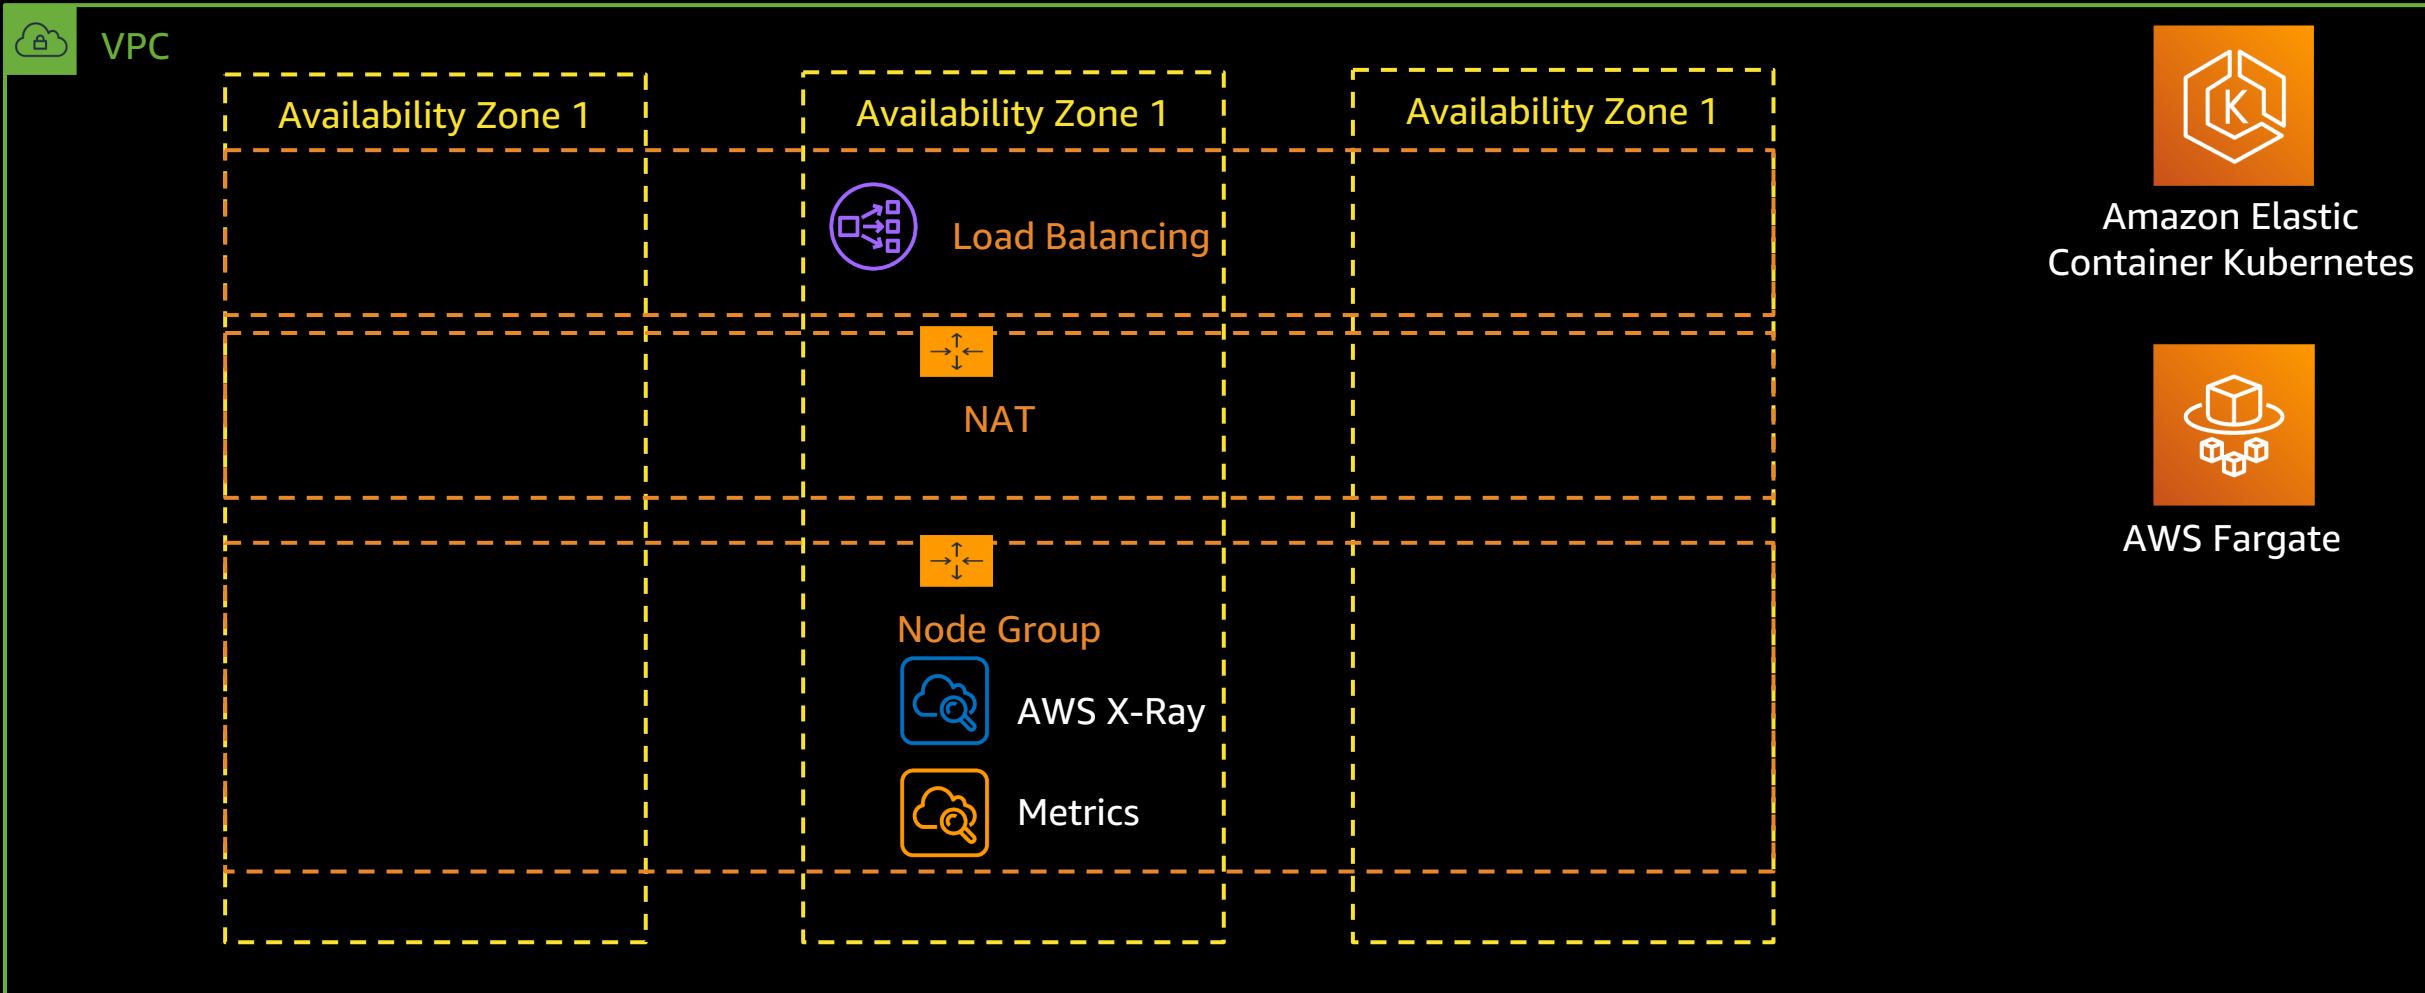
Amazon Elastic Container Registry



AWS App Mesh



Amazon Route 53



Amazon Elastic Container Kubernetes

Amazon CloudWatch

AWS Fargate

AWS Certificate Manager

Cluster scope (EKS)

Cluster management

IaC

Cluster Autoscaler

Backup / recovery

Storage

Observability

Metrics server

Container Insights

X-Ray daemon

Security

Network policy

Runtime security

OPA Gatekeeper

Tenancy management

Onboarding

Network isolation

Compute isolation

Storage isolation

Tenant (team) scope

Tenant-specific (Team = Namespace)

Build

Deploy

Support

Source



CodeCommit



GitHub



Bitbucket



AWS
CodePipeline

CI



Tekton



Resource control



Ingress control

Deployment control
(blue / green, canary)



CloudWatch



Access to X-Ray



Cluster Access (read-only)

Day one: Provisioning



Infrastructure as code (IaC)

- Actual code
- Define stack
- Create VPC
- Create cluster
- Create node group(s)

```
export class CdkEksBlueprintStack extends cdk.Stack {  
  ...  
  const vpc = new ec2.Vpc(this, clusterName + "-vpc");  
  
  const cluster = new eks.Cluster(this, clusterName, {  
    vpc: vpc,  
    clusterName: clusterName,  
    outputClusterName: true,  
    defaultCapacity: 0,  
    version: eks.KubernetesVersion.V1_18,  
  });  
  
  const ng = cluster.addNodegroupCapacity(ngName, {  
    instanceType: new ec2.InstanceType(instanceType),  
    minSize: 1,  
    maxSize: 4  
  });  
};
```



Cluster scope

- CNI and global network policy
- Metrics server (HPA)
- Cluster Autoscaler
- CloudWatch Container Insights
- NGINX Ingress Controller
- ArgoCD
- Many more

```
const addons: Array<ClusterAddon> = [  
  new CalicoNetworkPolicyAddon,  
  new MetricsServerAddon,  
  new ClusterAutoScalerAddon,  
  new ContainerInsightsAddon,  
  new NginxAddon,  
  new FluxCDAddon,  
  new ArgoCDAddon,  
  ...  
];  
  
for(let addon of addons) { // strict order  
  addon.deploy(this);  
}
```

Tenant onboarding



Tenant onboarding

- Create Namespace
- Set up network ingress / egress rules
- Set quotas
- Service accounts and IRSA
- Managed services setup
- Other setup

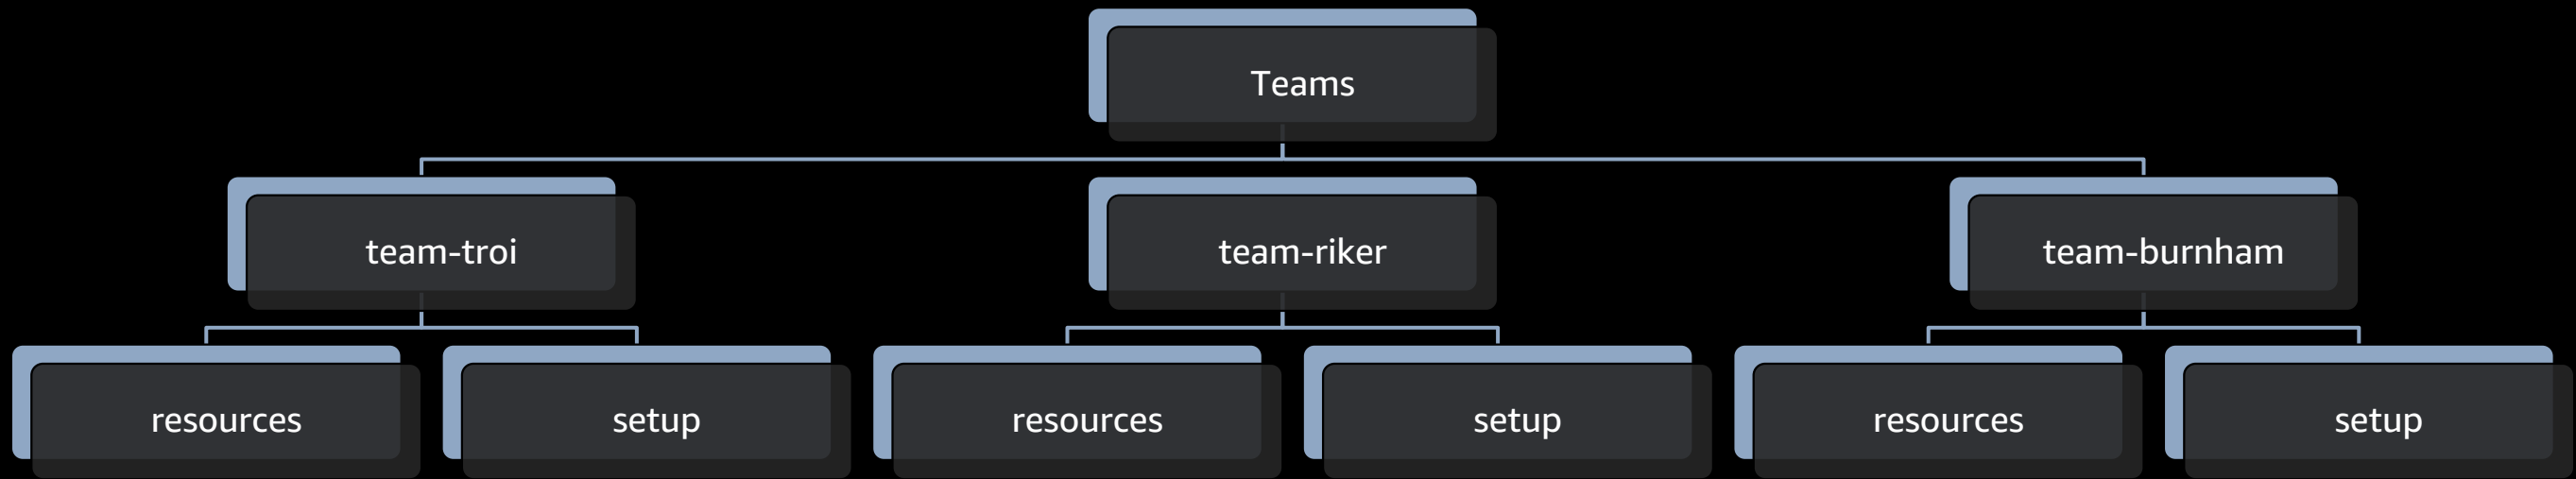
```
export class TeamTroisetup implements TeamSetup {
  setup(cluster: eks.Cluster, stack: cdk.Stack) {
    const namespace = cluster.addManifest(this.teamName, {
      apiVersion: 'v1',
      kind: 'Namespace',
      metadata: { name: this.teamName }
    });

    this.setupNamespacePolicies(cluster);

    const sa = cluster.addServiceAccount('inf-backend',
      { name: 'inf-backend', namespace: this.teamName});
    sa.node.addDependency(namespace);
    const bucket = new s3.Bucket(stack, 'inf-inbox');
    bucket.grantReadWrite(sa);

    ...
  }
}
```

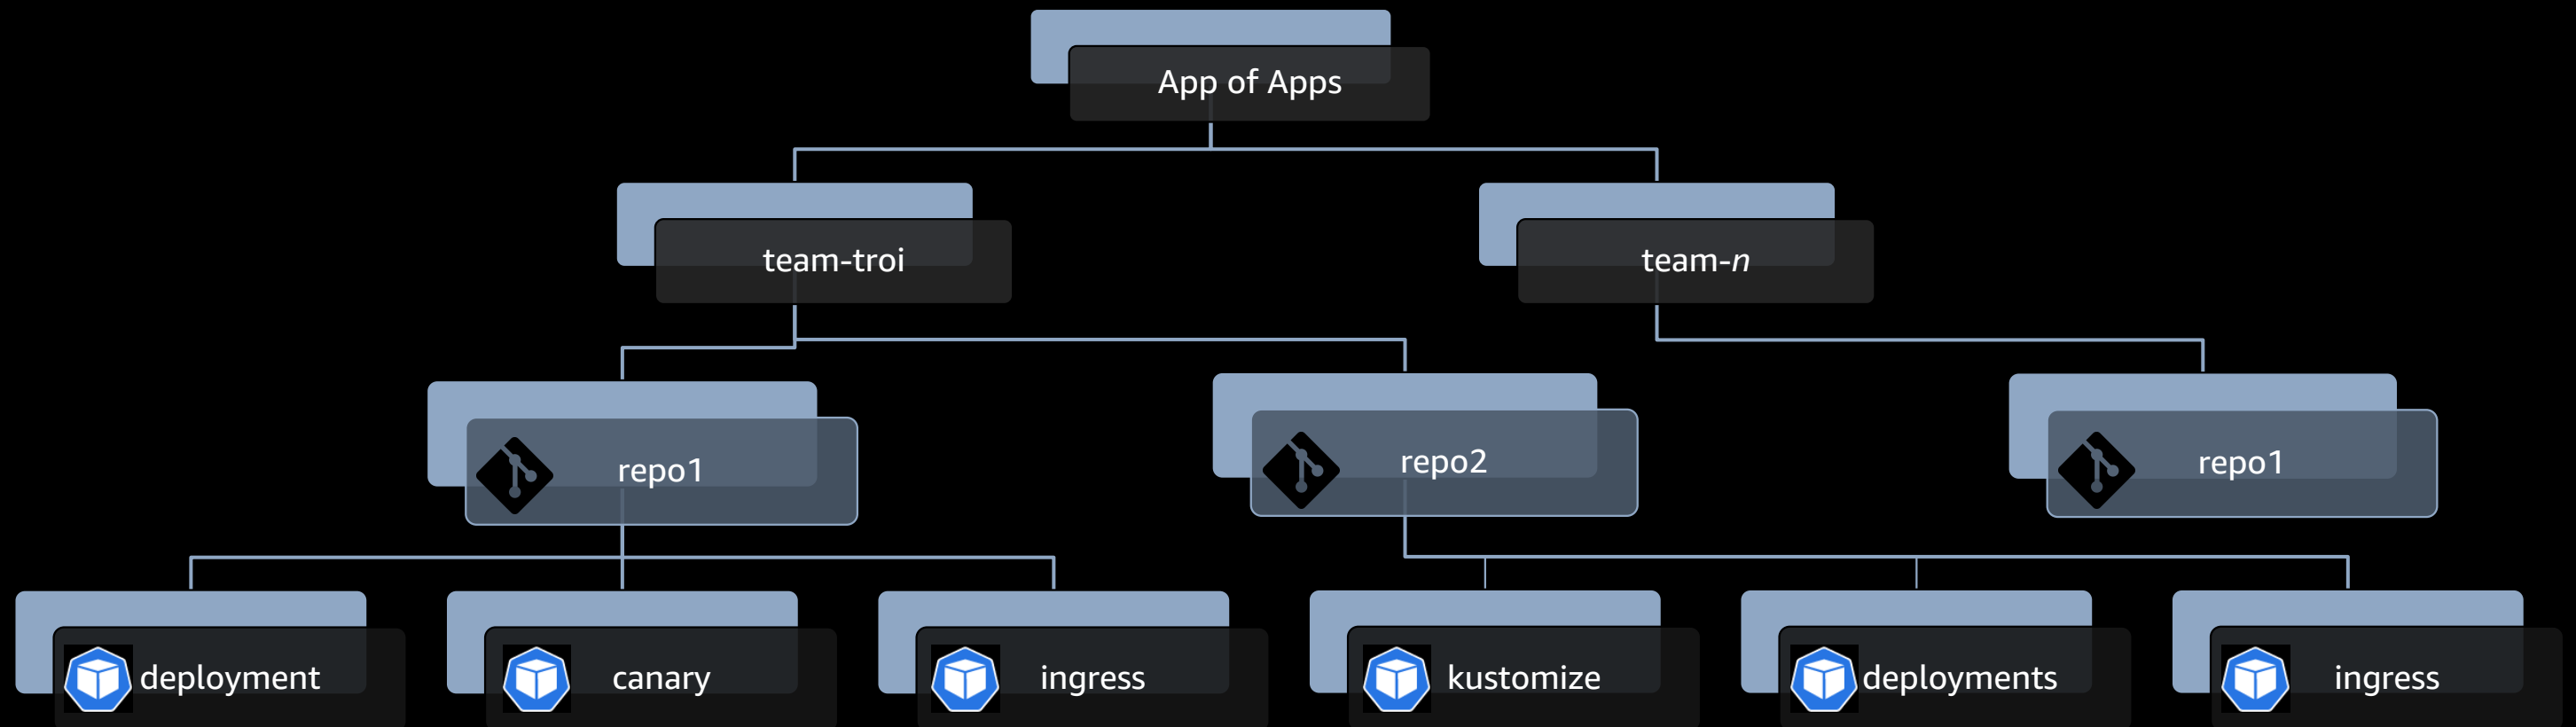
Tenant onboarding



```
const teams: Array<TeamSetup> = [  
  new TeamTroisSetup(),  
  new TeamRikerSetup(),  
  new TeamBurnhamSetup(),  
  ...  
];  
  
teams.forEach(setup => setup.setup(cluster, this));
```

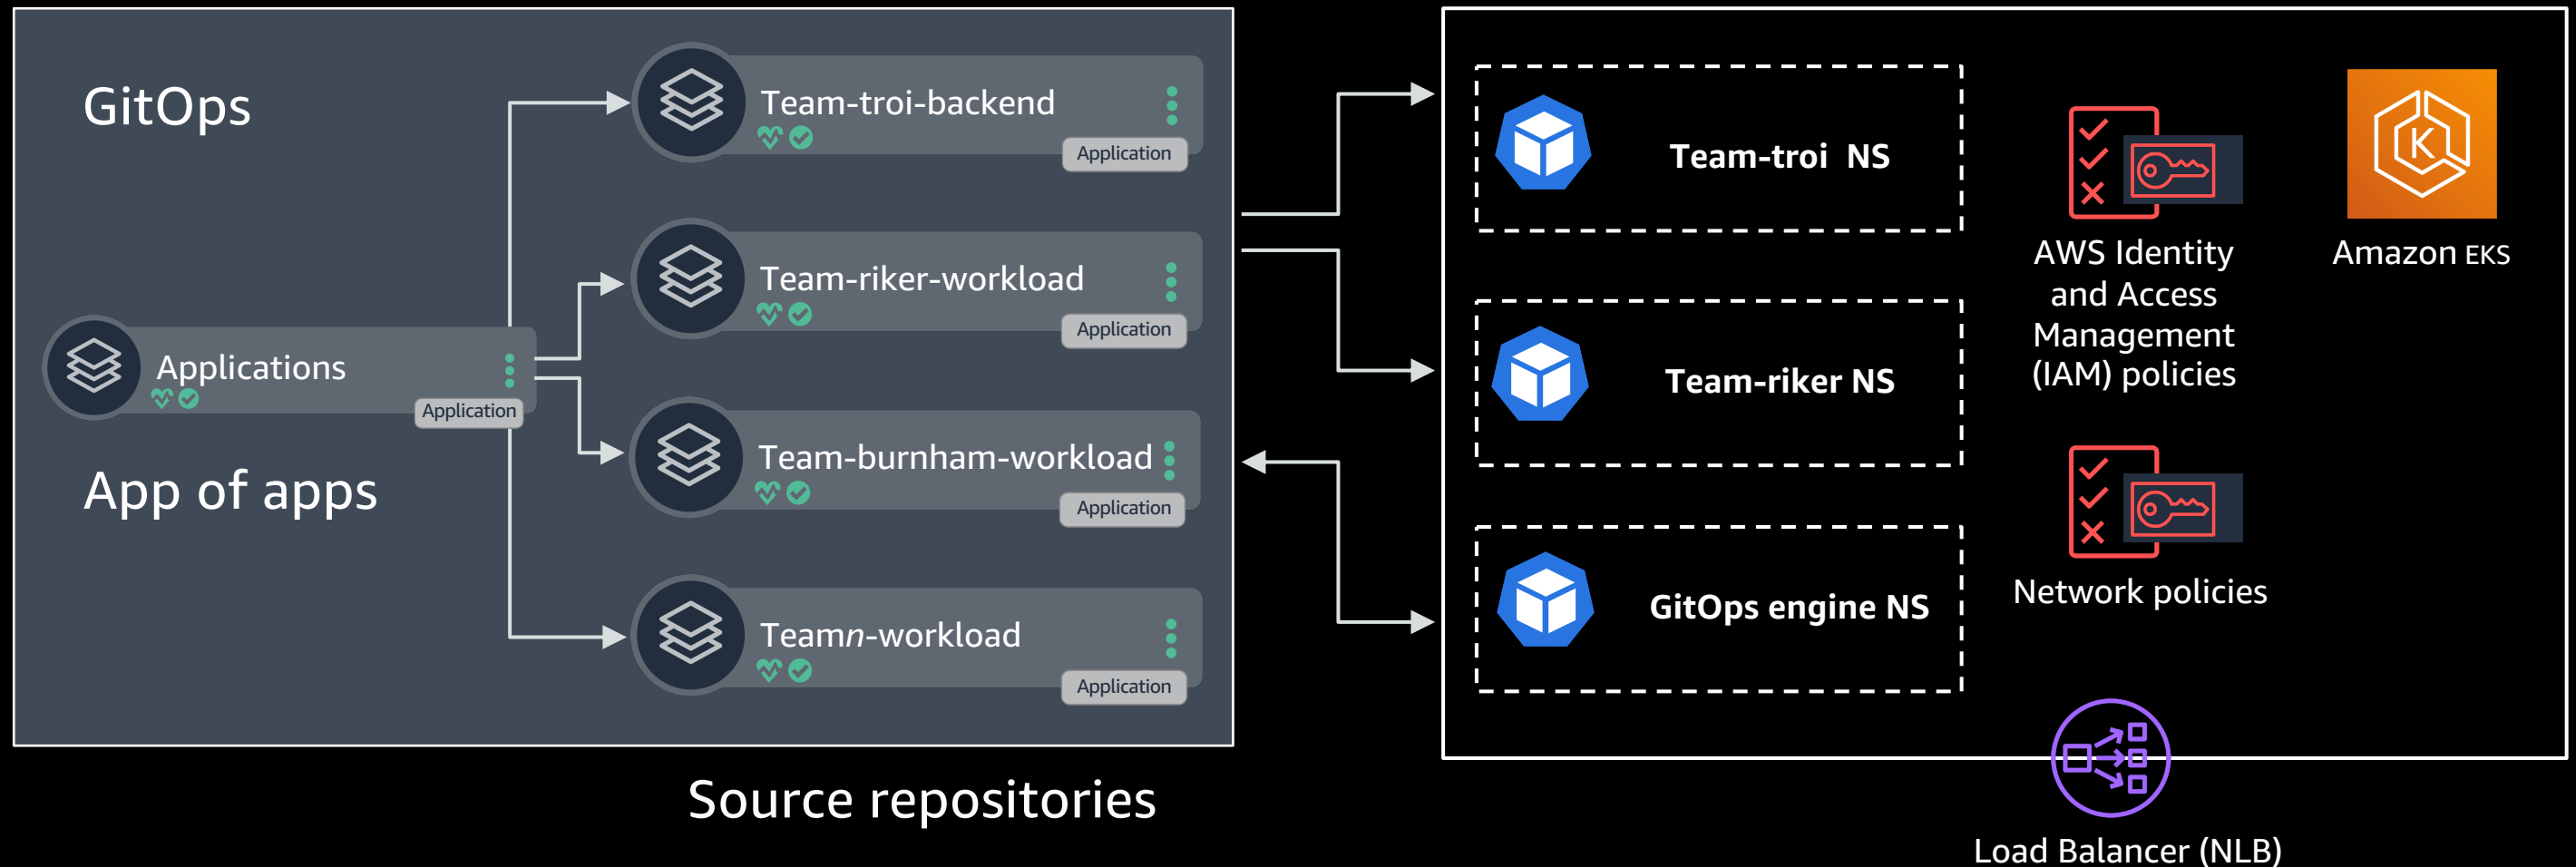
Tenant workloads

Multi-tenant GitOps with ArgoCD



Workload onboarding

Cluster bootstrapping



Bootstrapping: How?

1. Define apps using Kubernetes DSL (CRD) in single platform app repository

```
apiVersion: argoproj.io/v1alpha1
kind: Application
metadata:
  name: inf-backend-service
. . .
spec:
  destination:
    namespace: team-troi
  source:
    repoURL: git@{{troi-repo}}
    path: team-troi
. . .
```

```
apiVersion: argoproj.io/v1alpha1
kind: Application
metadata:
  name: authorization-service
. . .
spec:
  destination:
    namespace: team-riker
  source:
    repoURL: git@{{riker-repo}}
    path: team-riker
. . .
```

```
apiVersion: argoproj.io/v1alpha1
kind: Application
metadata:
  name: inventory-service
. . .
spec:
  destination:
    namespace: team-burnham
  source:
    repoURL: git@{{burnham-repo}}
    path: team-burnham
. . .
```

2. Bootstrap

```
$ argocd app create apps --repo ${APP_REPO} ...
$ argocd app sync apps
```

Deployment strategies

Capability	Argo Rollouts	Flagger	Service Mesh Direct	Kubernetes Direct
Blue / green	Yes	Yes	Yes	Yes
Canary	Yes	Yes	Yes	No
Metrics-driven decisions	Automatic	Automatic	Manual	Manual
Rollback	Automatic	Automatic	Manual	Manual
Analysis	Built-in	Built-in	Manual	Manual

Progressive delivery with Flagger

```
apiVersion: flagger.app/v1beta1
kind: Canary
metadata:
  name: inf-backend
spec:
  # service mesh provider can be: kubernetes, istio, appmesh, nginx, gloo
  provider: nginx
  # deployment reference
  targetRef:
    apiVersion: apps/v1
    kind: Deployment
    name: inf-backend
  # the maximum time in seconds before rollback
  progressDeadlineSeconds: 60
  # HPA reference (optional)
  autoscalerRef:
    apiVersion: autoscaling/v2beta1
    kind: HorizontalPodAutoscaler
    name: inf-backend
```

Multi-cluster, multi-region, and multi-account



Multi-cluster, multi-region

Deploy a single dev cluster

```
new CdkEksBlueprintStack(app, 'dev-stack',  
    "dev-cluster");
```

```
$ cdk deploy dev-stack
```

Deploy a few staging clusters
to multiple regions

```
new CdkEksBlueprintStack(app, 'east-staging-stack',  
    "staging-cluster", { env: {region: 'us-east-2'}});
```

```
new CdkEksBlueprintStack(app, 'west-staging-stack',  
    "staging-cluster", { env: {region: 'us-west-2'}});
```

Multi-account

Split dev and prod accounts

```
new CdkEksBlueprintStack(app, 'east-dev-stack',  
{  
  env: {  
    region: 'us-east-2',  
    account: 'DEV_ACCOUNT'  
  }  
});
```

```
new CdkEksBlueprintStack(app, 'east-prod-stack',  
{  
  env: {  
    region: 'us-east-2',  
    account: 'PROD_ACCOUNT'  
  }  
});
```

IaC pipeline

```
// Do this as many times as necessary
// with any account and region
pipeline.addApplicationStage(
  new FactoryApplication(this, 'dev', {
    env: {
      account: '${DEV_ACCOUNT}',
      region: 'us-east-2',
    }
  })
);

pipeline.addApplicationStage(
  new FactoryApplication(this, 'staging', {
    env: {
      account: '${ST_ACCOUNT}',
      region: 'us-east-2'
    }
  })
);
```

The screenshot displays the AWS CodePipeline console for a pipeline execution. The pipeline execution ID is 942ca1aa-18db-4e4c-a34a-e58233259b49. The pipeline consists of three stages, all of which have succeeded:

- Source** (Succeeded):
 - Provider: GitHub (Version 1)
 - Duration: 6 minutes ago
 - Job ID: e88d1094
 - Log: [e88d1094](#) GitHub: changed region for pipeline stack stages, readme updates
- Build** (Succeeded):
 - Provider: Synth (AWS CodeBuild)
 - Duration: 5 minutes ago
 - Job ID: e88d1094
 - Log: [e88d1094](#) GitHub: changed region for pipeline stack stages, readme updates
- UpdatePipeline** (Succeeded):
 - Provider: SelfMutate (AWS CodeBuild)
 - Duration: 2 minutes ago
 - Job ID: e88d1094
 - Log: [e88d1094](#) GitHub: changed region for pipeline stack stages, readme updates

Between the stages, there are buttons labeled "Disable transition" with a downward arrow, indicating that the transitions between stages are disabled.

References

- <https://github.com/shapiro103/cdk-eks-blueprint>
- <https://github.com/shapiro103/argo-apps>
- <https://argoproj.github.io/argo-cd/>
- <https://github.com/weaveworks/flagger>
- <https://argoproj.github.io/argo-rollouts>
- https://d1.awsstatic.com/events/reinvent/2019/Architecting_multi-tenant_PaaS_offerings_with_Amazon_EKS_GPSTEC337.pdf
- <https://www.slideshare.net/AmazonWebServices/enterprise-devops-patterns-of-efficiency-ent311r1-aws-reinvent-2018>
- <https://www.gartner.com/en/documents/3989064/hype-cycle-for-platform-as-a-service-2020>
- <https://searchcloudcomputing.techtarget.com/opinion/Why-the-PaaS-market-failed-to-live-up-to-the-hype>

Thank you!





Please complete
the session survey