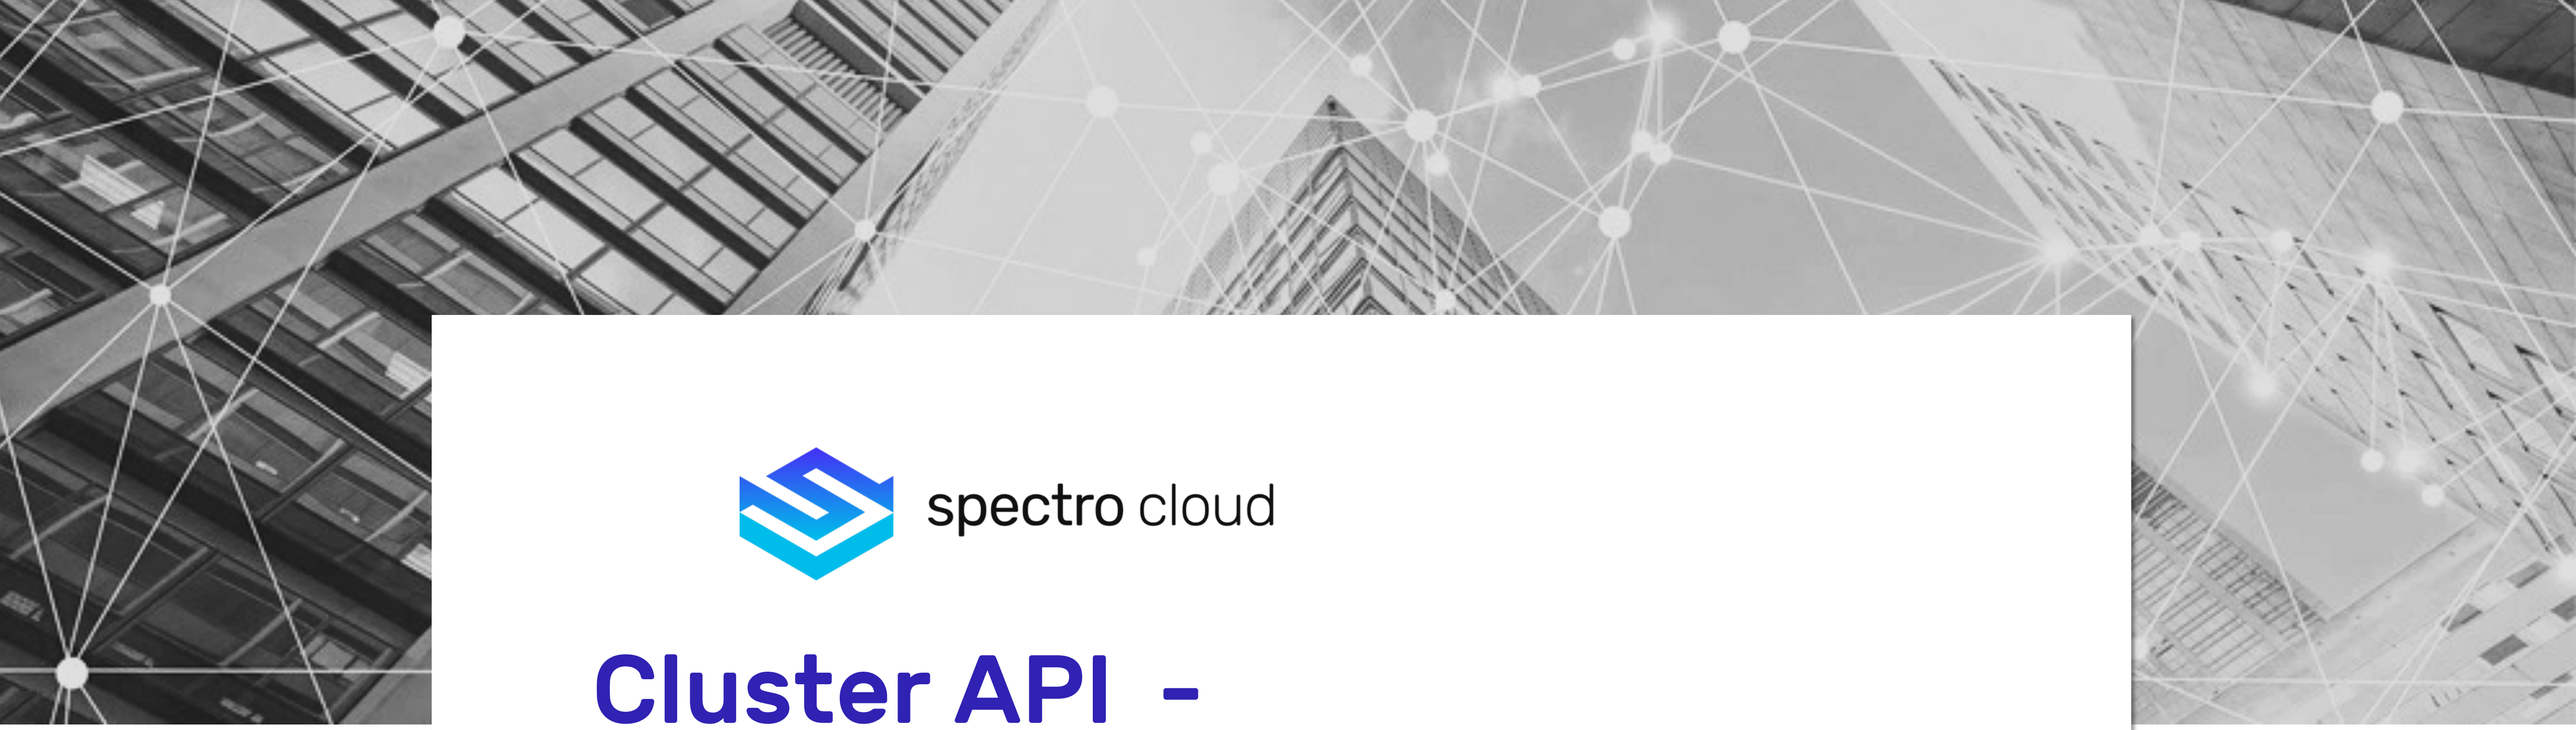

spectro cloud

Cluster API - Yesterday, Today, Tomorrow

July 2020

Saad Malik

CTO & Co-Founder @ Spectro Cloud

Jun Zhou

Chief Architect @ Spectro Cloud

spectro cloud

Agenda

- Cluster API --
- Cluster API
- Cluster API ++

Popular Tools

- kube-up.sh
- Kubespray Oct/2015
- Kubeadm Sep/2016
- Kops Oct/2016
- Cluster API Mar/2019

kube-up.sh

v1.0.0 **kubernetes** / **cluster** /

mbforbes committed `739cb2f` on Jul 9, 2015

..	
addons	Assigned emp
aws	Merge pull re
azure	Fix cluster mc
gce	Robustly clea
gke	GKE upgrade
images	Stop exposing
juju	Updating scri
libvirt-coreos	Change confi
ovirt	Add support f
rackspace	Remove unus
saltbase	Merge pull re
ubuntu	update kubec
vagrant	Missing ca cr
vsphere	Remove unus
README.md	Add ga-beac

v1.3.0 **kubernetes** / **cluster** /

k8s-github-robot authored and **eparis** committed `d0e521b` on Jun 30, 2016

..	
addons	Added PetSet support to addon manager.
aws	Merge pull request #28030 from nikhiljindal/rever
azure	azure: azkube v0.0.5 + deploy kube-system
centos	Merge pull request #23829 from derekwaynecarr/
gce	Merge pull request #28132 from madhusudancs/f
gke	Merge pull request #27803 from fabioy/fix-multiz
images	Merge pull request #28087 from luxas/fix_hyperku
juju	cluster/juju: Updated the url for the getting starte
kubemark	Fix default arguments in kubemark
lib	Handle multiple MIGs (single-zone) properly in GK
libvirt-coreos	Remove the restart-kube-proxy and restart-apise
local	Add local/util.sh
mesos/docker	Merge pull request #28132 from madhusudancs/f
openstack-heat	Merge pull request #28132 from madhusudancs/f
ovirt	Add support for oVirt cloud provider
photon-controller	Merge pull request #28132 from madhusudancs/f
rackspace	fix for #13511
saltbase	Merge pull request #28207 from mwielgus/ca-0.2
skeleton	Use a skeleton provider for unimplemented functi
ubuntu	Merge pull request #28040 from ibm-contribs/fix
vagrant	Merge pull request #28132 from madhusudancs/f
vsphere	Merge pull request #28132 from madhusudancs/f
OWNERS	Remove myself from a bunch of OWNERS files, as

master **kubernetes** / **cluster** /

ixdy committed `089a1af` yesterday

..	
addons	Merge pull req
gce	Merge pull req
images	Move ixdy to e
kubemark	*.sh: cleanup a
log-dump	*.sh: cleanup a
pre-existing	Merge pull req
skeleton	/cluster: add /
BUILD	Switch to stati

spectro cloud

kubeadm

- First release: Sep/2016
- Domain expertise, Only cluster bootstrapping / Configuration
- Kubeadm Init + Kubeadm Join
- Kubeadm upgrade
- Very focused scope, do the minimum, and do it best

kubeadm phases

Kubeadm init

- preflight
- kubelet-start
- certs
- kubeconfig
- control-plane
- etcd
- upload-config
- upload-certs
- mark-control-plane
- bootstrap-token
- kubelet-finalize
- dns/kube-proxy

Kubeadm join controlplane

- preflight
- download-certs
- certs
- kubeconfig
- control-plane
- kubelet-start
- etcd
- update-status
- mark-control-plane

Kubeadm join node

- preflight
- kubelet-start

kubeadm HA

kubeadm HA topology - stacked etcd

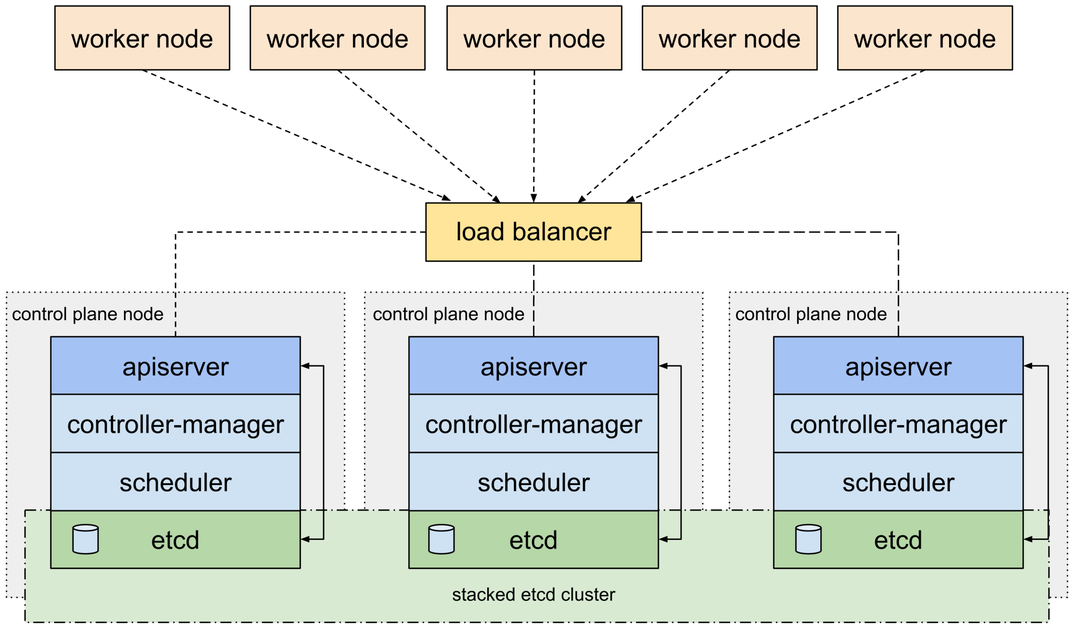

kubeadm HA topology - external etcd

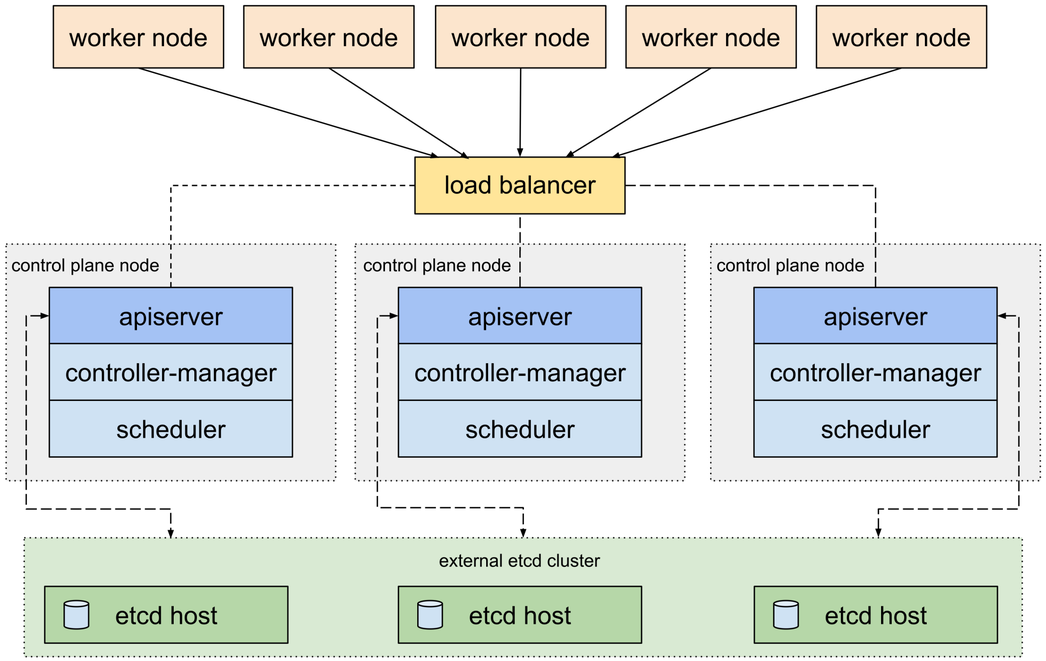

Image from k8s.io doc

kubeadm - what's missing

- Doesn't have:
  - Infrastructure provisioning
  - Critical addons: CNI/StorageClass
  - Non-critical addons: monitoring, logging, auth
  - cloud provider integrations

kubespray

- V1.0.0. Oct/2015
- Anywhere: AWS,GCE,Azure,OpenStack,vSphere,Packet,OCI ,Baremetal
- Any OS: CoreOS,Debian,Ubuntu,CentOS/RHEL,FedoraCoreOS
- Fedora,OpenSUSE,Oracle Linux
- Baremetal first
- Bring your own Machine
- Ansible, OS centric, generic configuration management

kubespray HA – without external LB

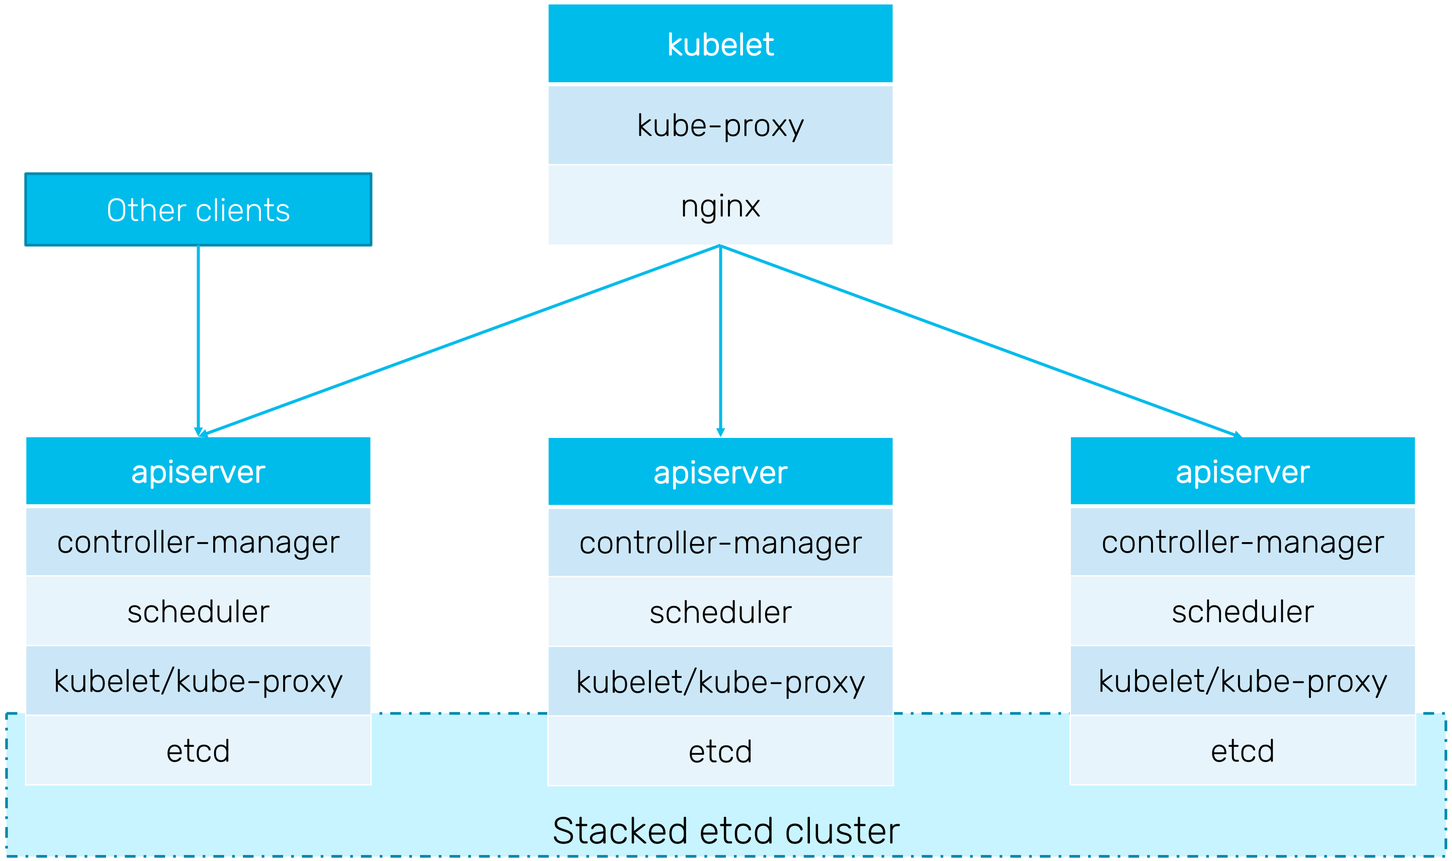

kubespray

- Supports kubeadm for cluster configuration since v2.3 (Oct/2017)

- In addition to kubeadm
  - CNI / storage class
  - Cloud-provider config
  - Proxy support
  - Air-gap environment
  - Other CRI runtimes

- What's missing
  - Infrastructure management

kops

- First release: Oct/2016
- Infrastructure + k8s + Addons
- Cloud First
- Deeply integrated with Cloud features
  - AWS: DNS / Auto Scaling Group / ELB / EBS / KMS / S3 / IAM

kops – unique features

- DNS
- Etcd
- Auto Scaling Group
- Terraform compatible configuration
- Horizontal Pod Autoscaling
- AWS IAM Authenticator

kops tooling

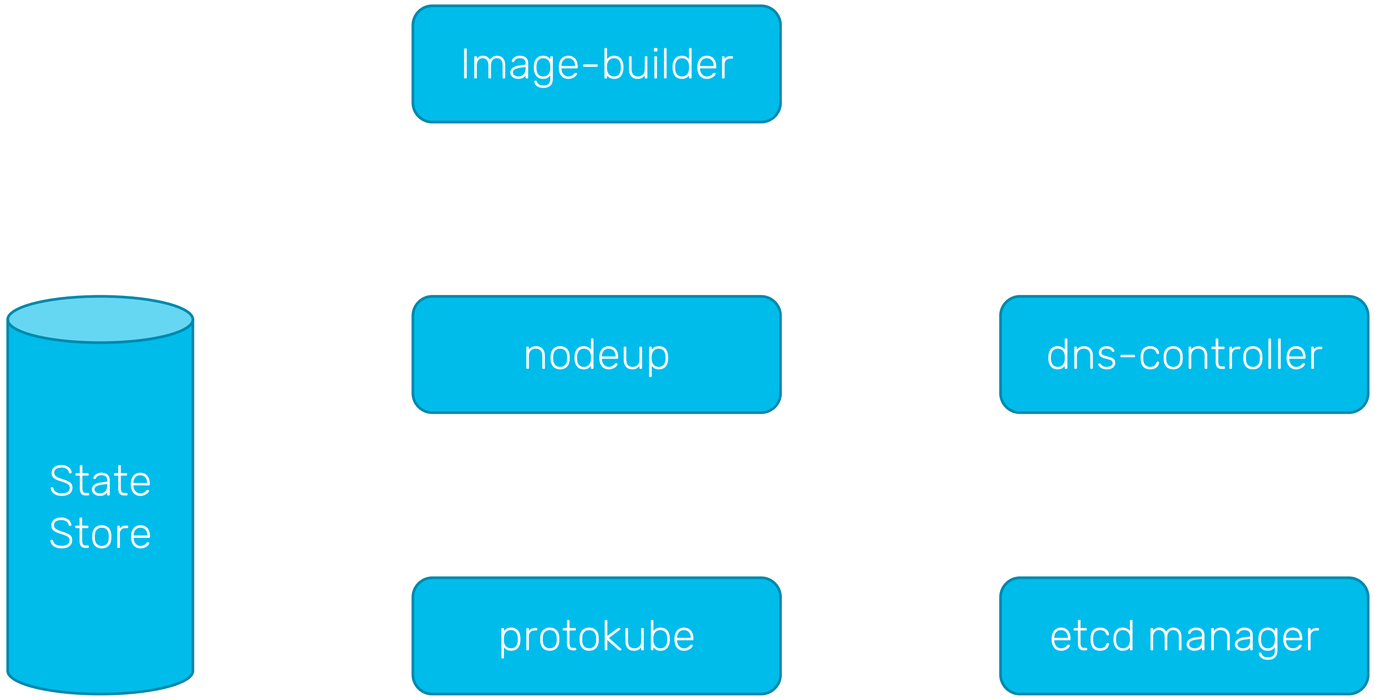

kops HA on AWS

- Resources
  - AIG
  - IAM
  - EBS
  - ELB
  - VPC
  - Subnet
  - NAT Gateway
  - Internet Gateway
  - Route Table
  - Security Group
  - Elastic IP

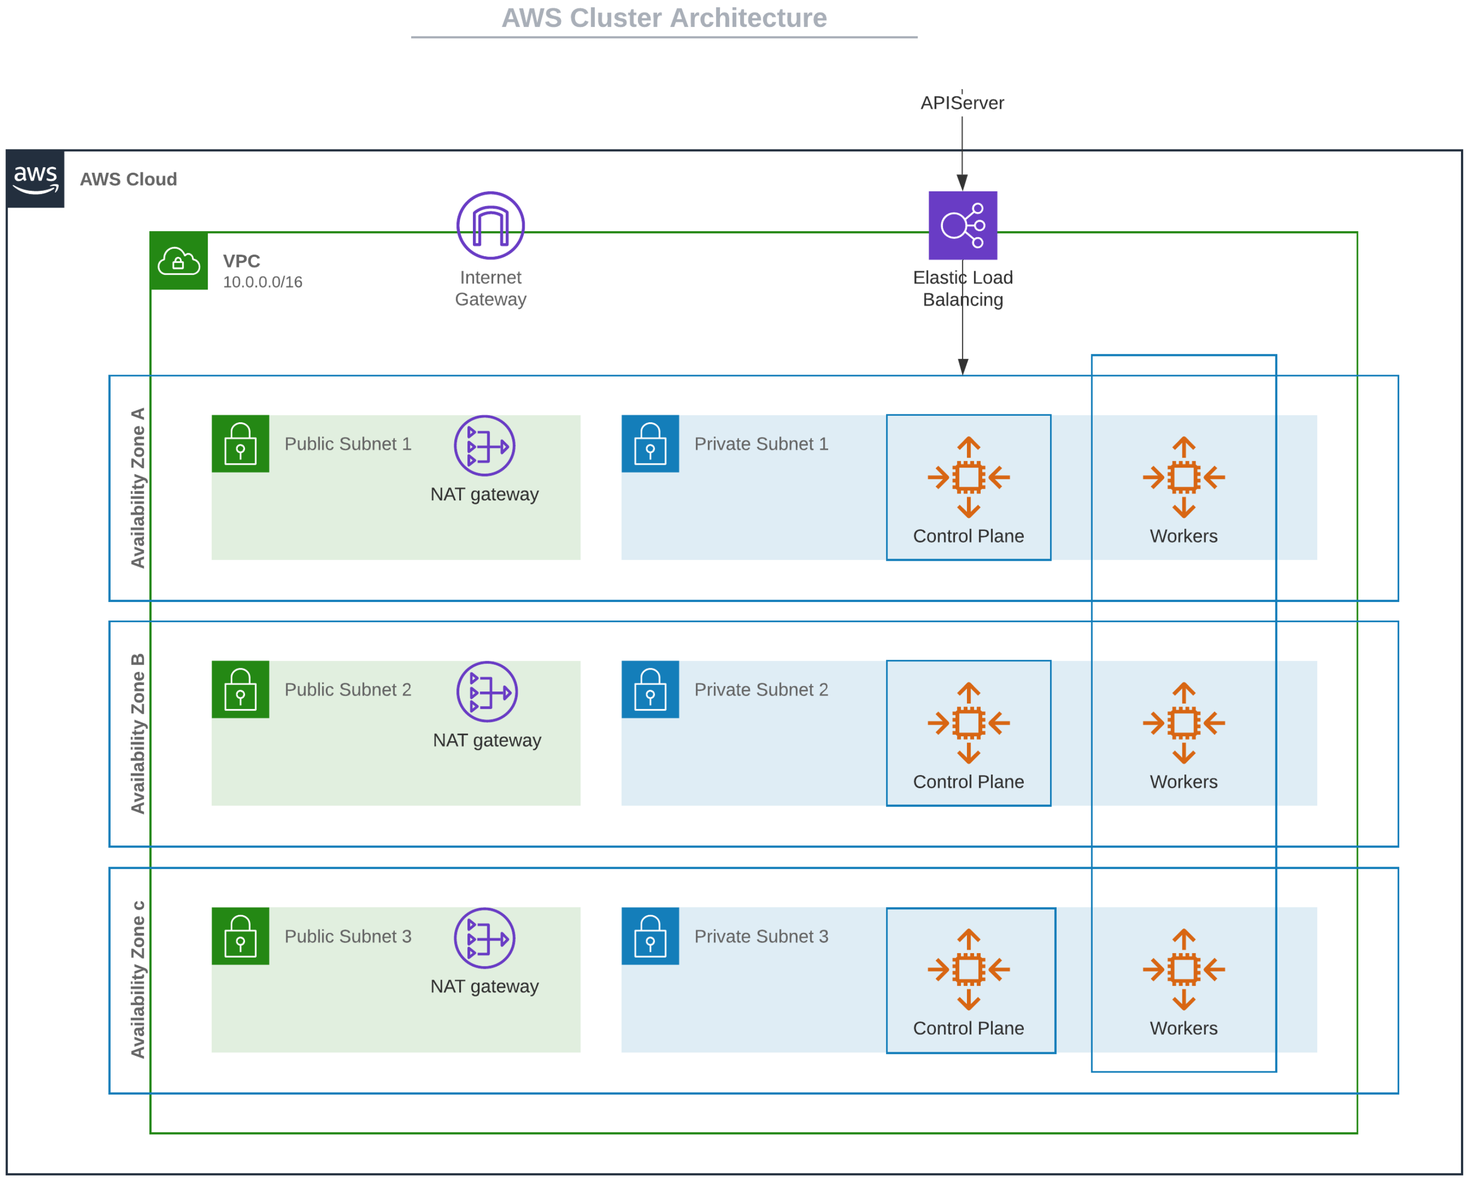

kops – possible improvements

- CLI only, no server reconciliation
- Critical Addons are bundled within kops with fixed version
- No baremetal/vsphere support
- Lag behind upstream release.
  - Latest release: v1.17.0 on 05/31/2020
  - Upstream 1.17.0 released on 12/09/2019

Summary – challenges need to be addressed by installer

- Public/Private
- BareMetal/Edge
- Networking
- Storage

1. Infrastructure

- Debian/Ubuntu
- CentOS/RHEL
- Fedora/CoreOS
- Windows

2. Operation System

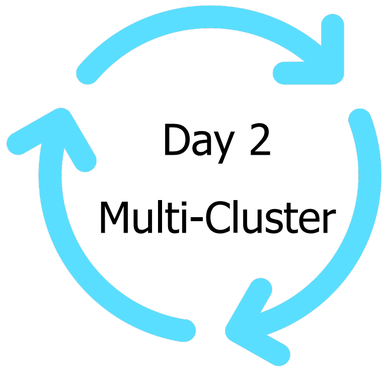

- CNI/StorageClass
- LB/Ingress/
- Auth/Log/Monitor
- Security

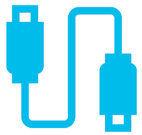

4. Integration

- Single master / HA
- Etcd/Failure Domain
- 100+ config options
- Security best practice

3. Configuration

What is Cluster API

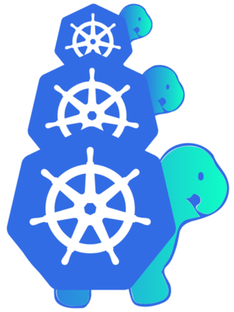

- The Cluster API is a Kubernetes project to bring declarative, Kubernetes-style APIs to cluster creation, configuration, and management. It provides optional, additive functionality on top of core Kubernetes to manage the lifecycle of a Kubernetes cluster. -- https://cluster-api.sigs.k8s.io

- Current state:
  - v1alpha3
  - 200+ contributors from VMware/Google/IBM/Red hat/Microsoft/…
  - AWS/VSphere/Azure/GCP/Openstack/Digital Ocean/Packet/Alibaba/Tencent/Metal3/…

- Cluster API CNCF webinar : https://www.youtube.com/watch?v=A2BBuKx1Yhk

How does Cluster API work

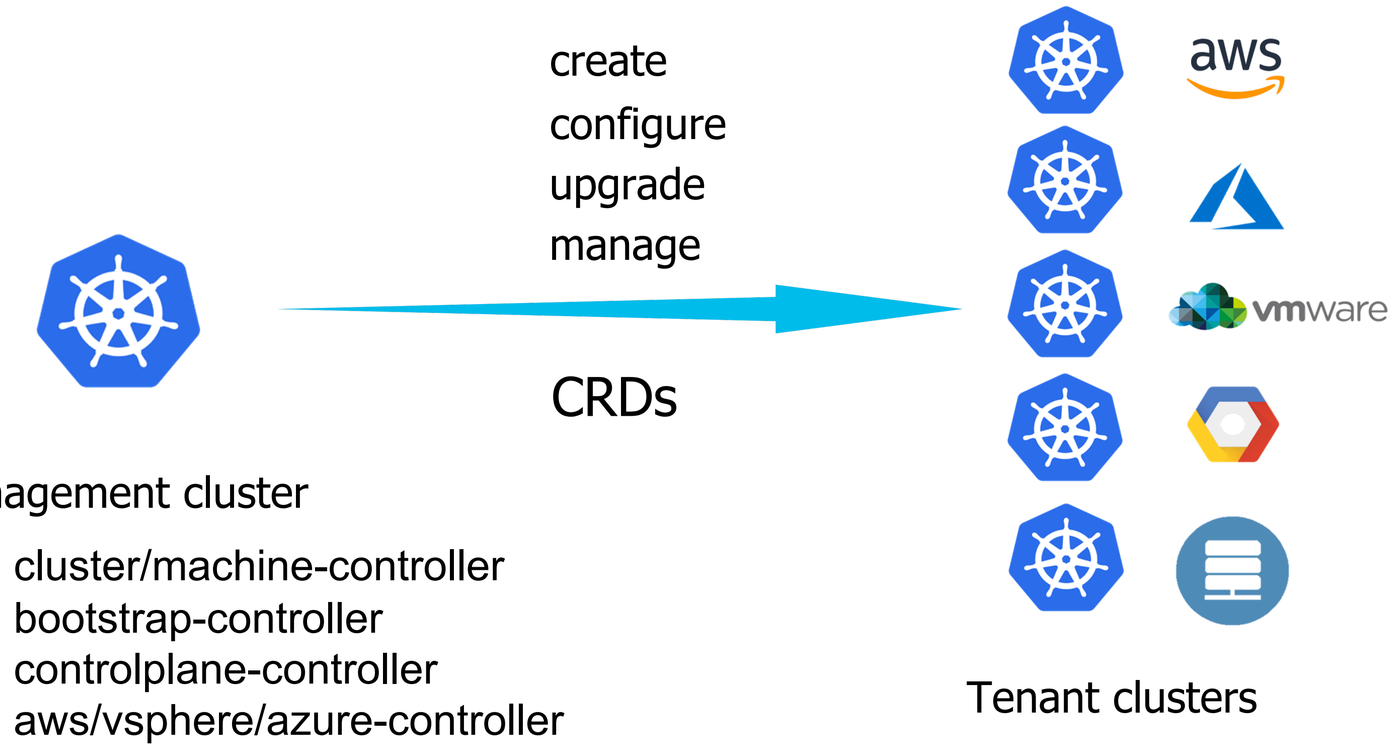

- cluster/machine-controller
- bootstrap-controller
- controlplane-controller
- aws/vsphere/azure-controller

Cluster API Deep Dive - CRD

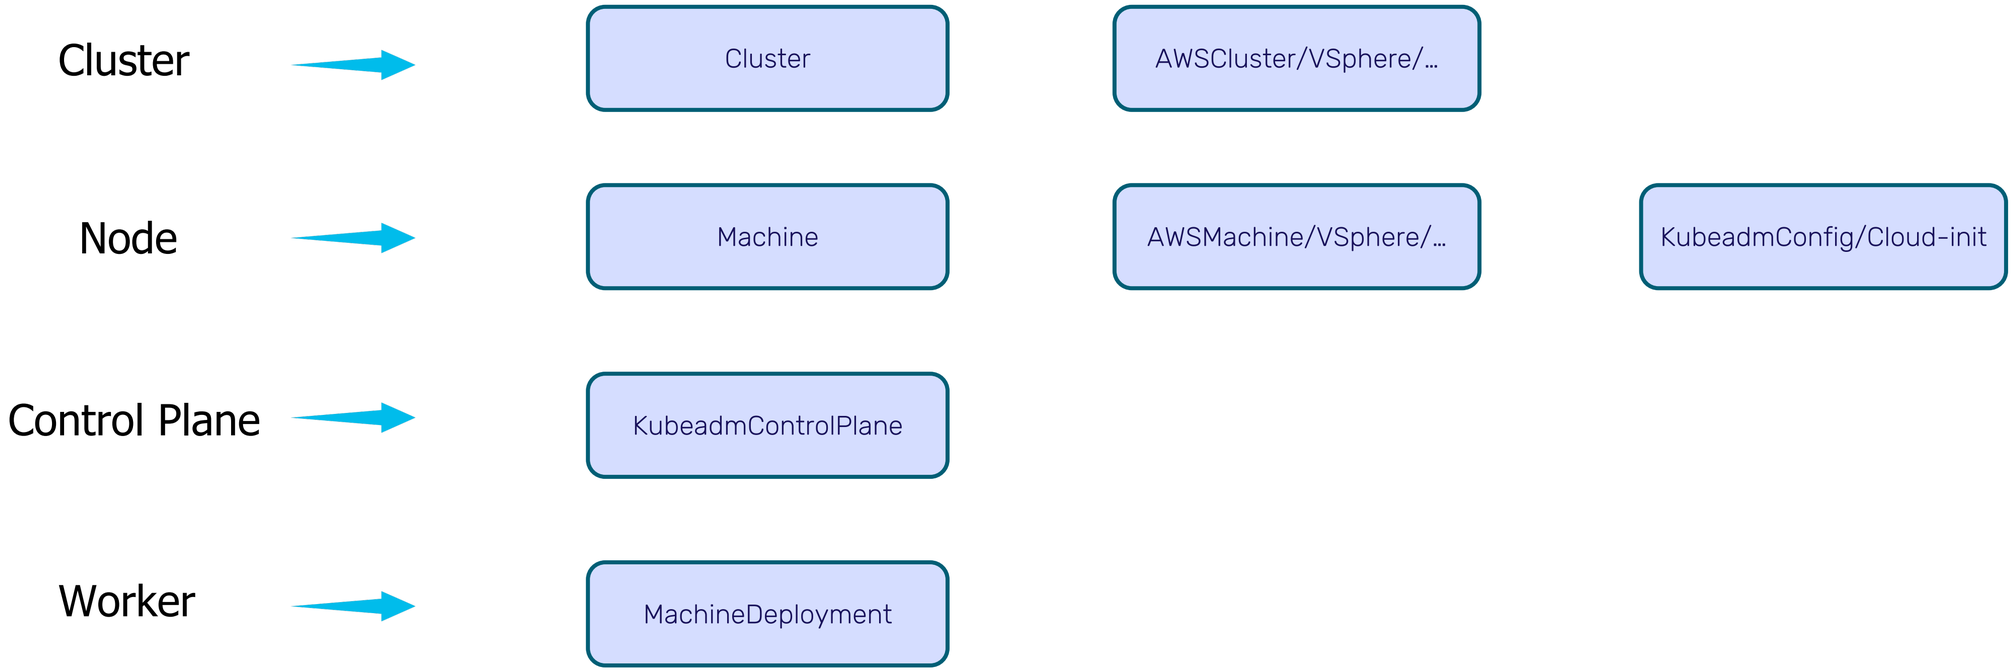

Cluster on AWS

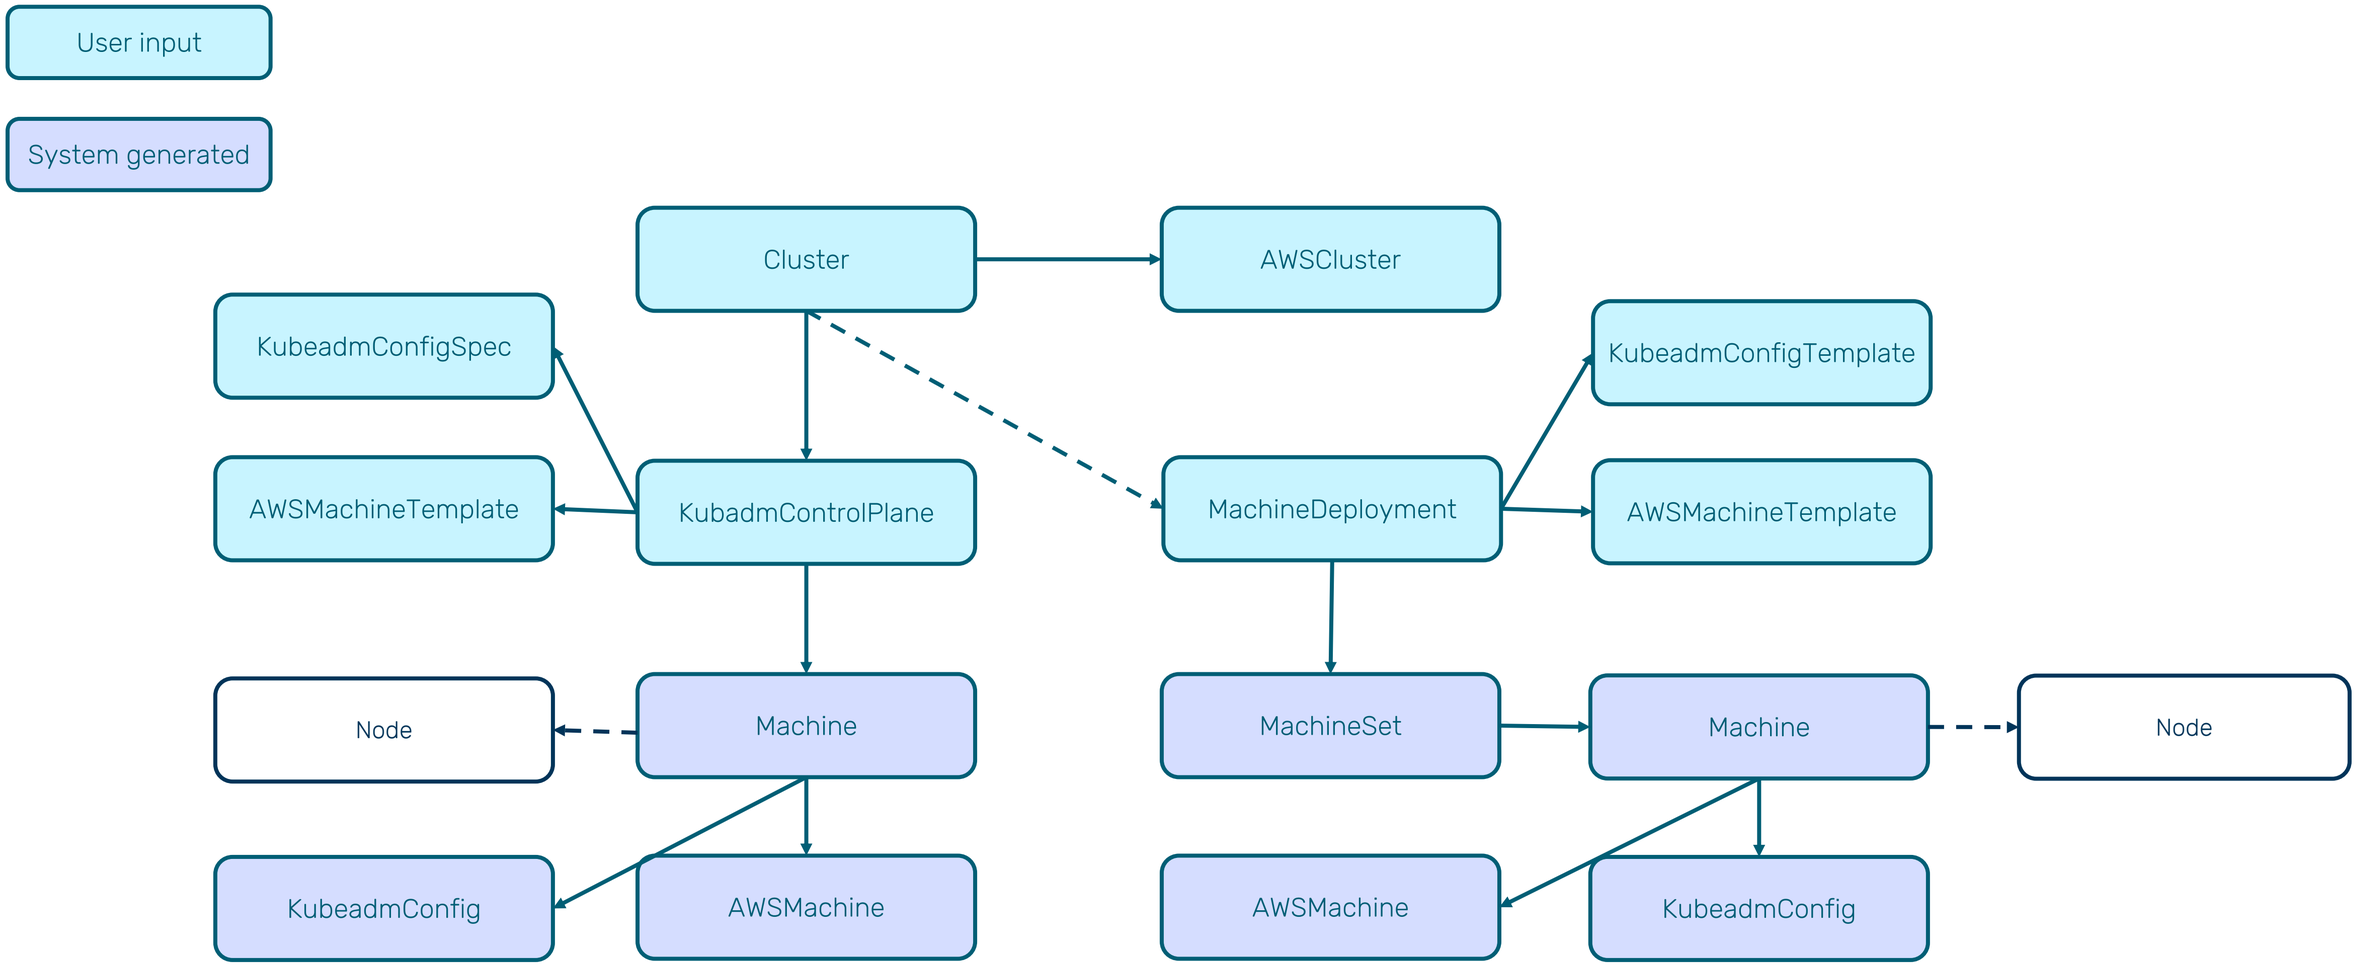

Cluster API on AWS

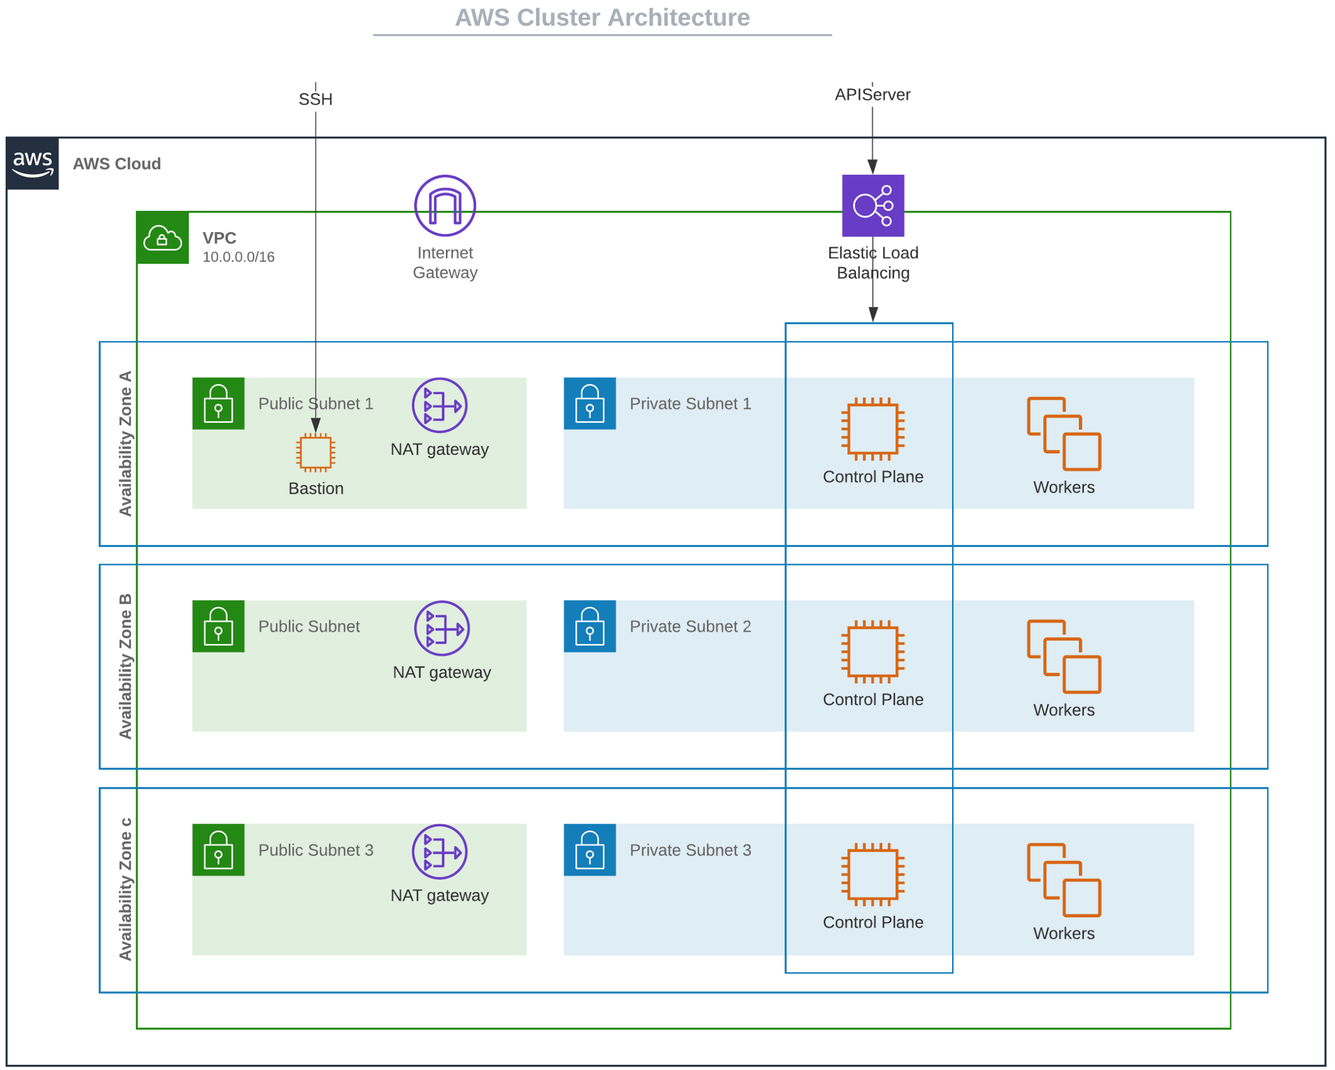

spectro cloud

Cluster API Roadmap

- V0.3.0. -- March/2020
  - KubeadmControlPlane
  - MachineHealthCheck (MHC) for worker nodes
  - VSphere HA with HAProxyLoadBalancer
- V0.3.+ -- July/2020
  - Control-plane robustness
  - Autoscaler integration
  - Spot instances support
  - Machine pre-delete hooks
  - ClusterResourceSet
- V0.4 -- Q4/2020
  - Pluggable Machine load balancers
  - Bootstrap failure detection

Kops VS Cluster API

	Kops	Cluster API
Clouds		+
OS	+	+
Bootstrapping	+	+
Integration	+	
Day 2	+	+
MultiCluster		+

Kops + Cluster API

Will this meet enterprise requirements?

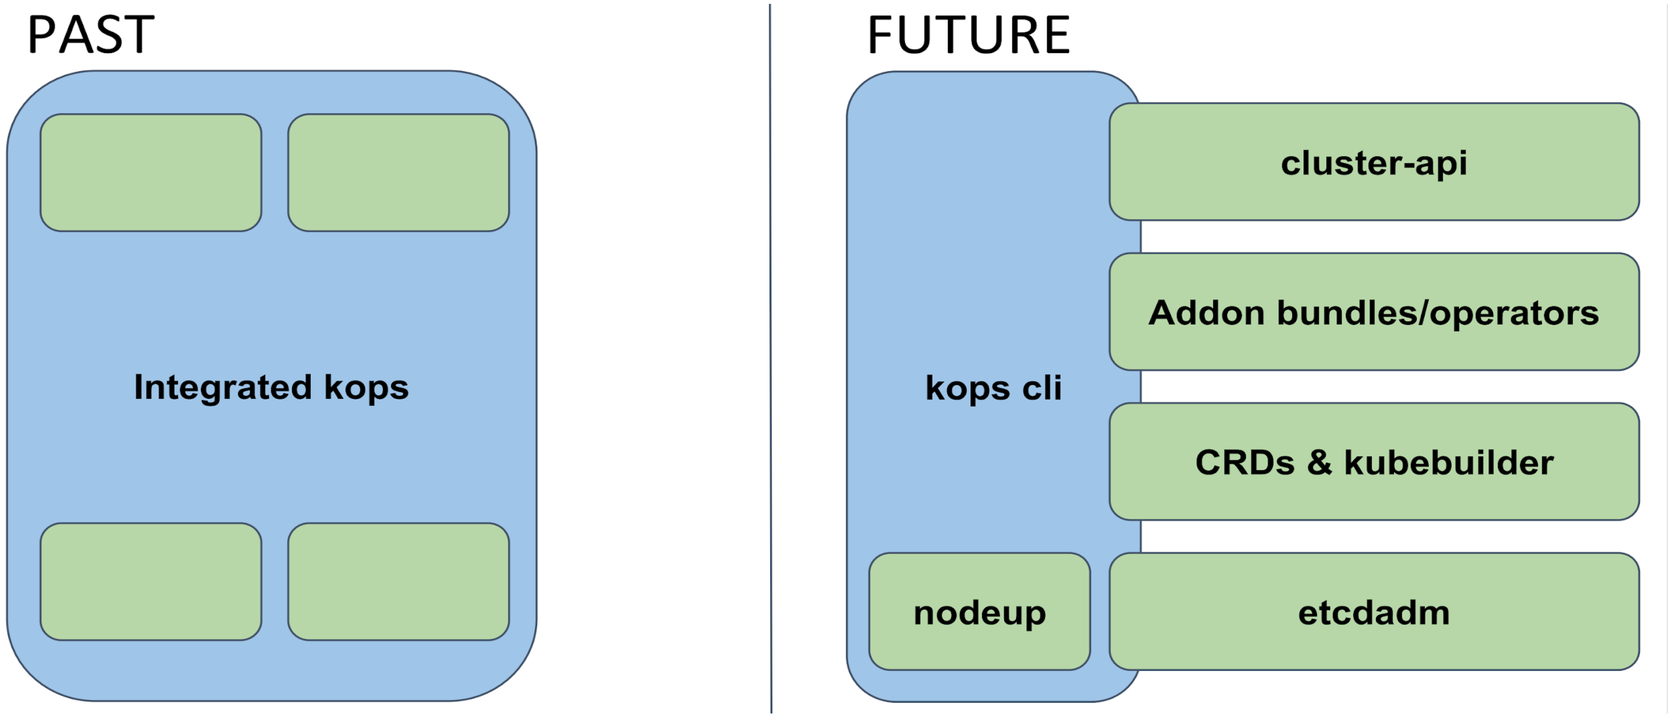

From justinsb's kubecon 2019 kops talk

Enterprise responsibilities

Security

- Base OS, containers, and integrations
- Ingress network policies
- Access to clusters

Logging & Observability

- Logging, Monitoring, Tracing
- Alerting and notifications
- For both platform & application workloads

Config Management

- Namespace management
- Resource Quotas

Business Continuity

- HA for control planes
- Backup, Restore, and Snapshots
- Control plane upgrades with limited impact to workloads

Enterprise requirements

Ease-of-use

- Higher-level abstraction
- Stable APIs
- Combining best of all tools!

Manageability

- Multi-cluster/cloud: public, private, and even BM environments
- Consistency across different clusters

Flexibility

- Any cloud of your choice
- OS: pre-built or dynamically generated
- Any K8s integrations
- Default configuration, but overridable

Higher-Level Orchestrator Features

Focus on value-add capabilities such as:

- Namespace and resource management
- Secure and protect cluster and workloads
- Policy, Audit, Alert, Compliance
- Governance, RBAC, and Quota Control
- CICD and Deployment
- Service Mesh, Observability, and Logging
- Archive, Backup, Restore

- Multi-cluster lifecycle management
- Day-2 upgrades
- Support for 10+ public and private clouds

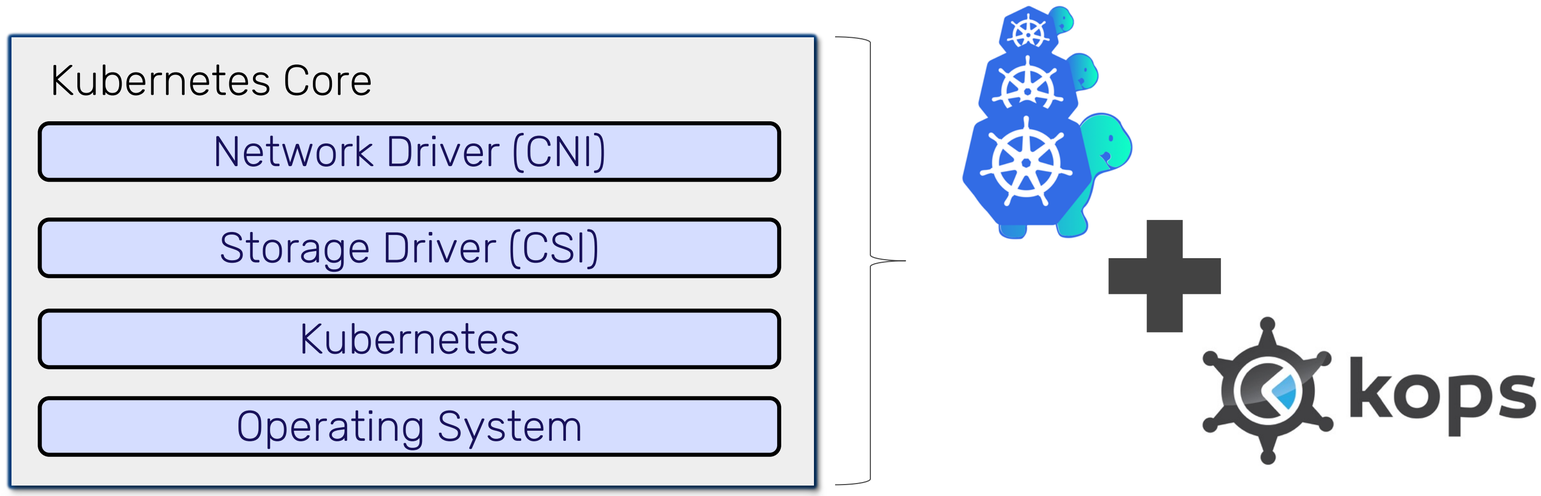

Add-On Integrations

- API Gateway
- Load Balancers
- Ingress
- Authentication
- Monitoring
- Logging
- Service Mesh
- Tracing
- CICD
- Build Security
- Image Build
- Runtime Security

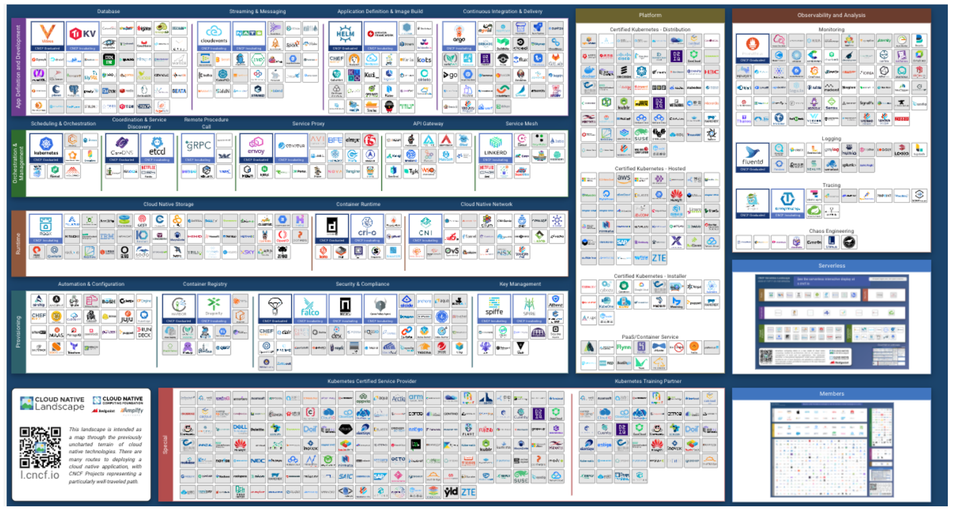

CLOUD NATIVE Landscape

Kubernetes Core

Add-On Integrations

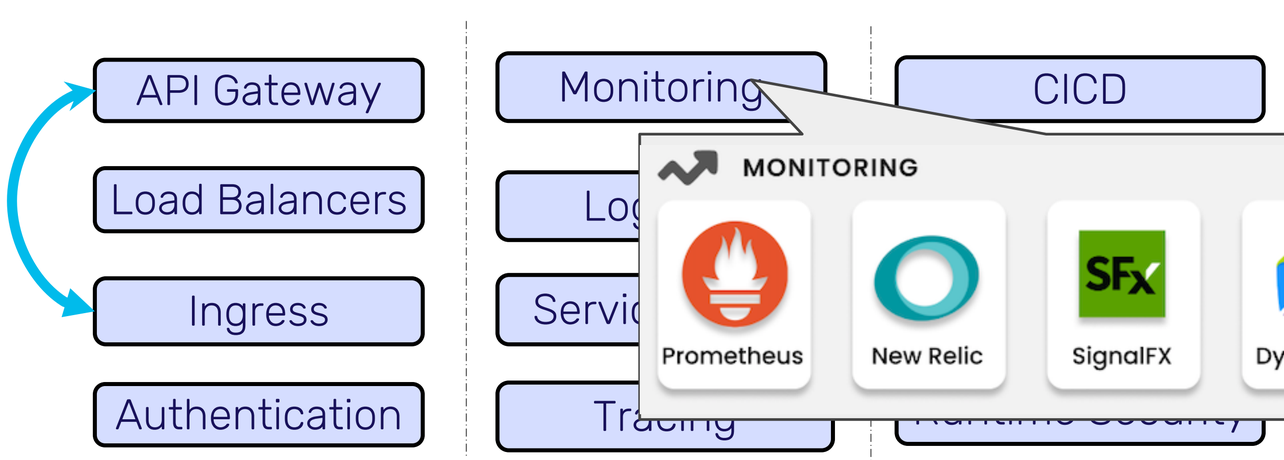

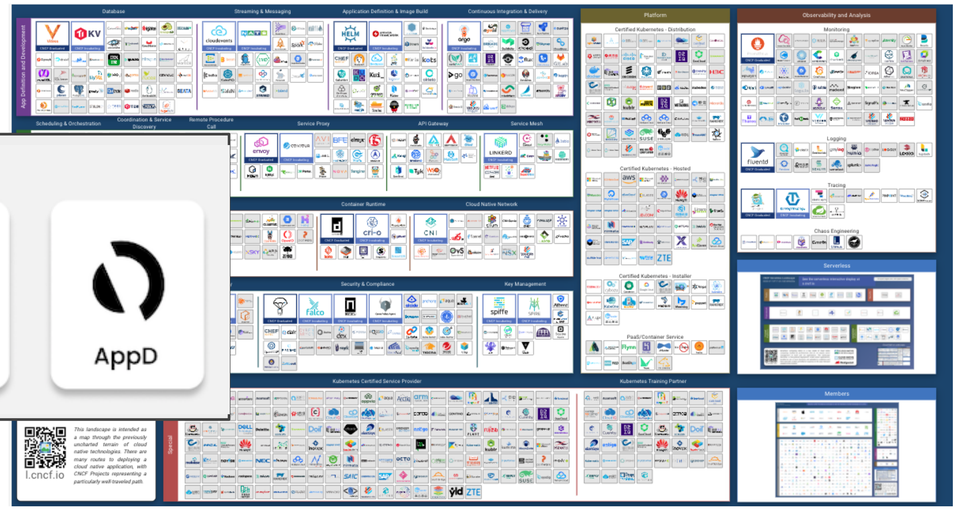

CLOUD NATIVE Landscape

Kubernetes Core

Q & A

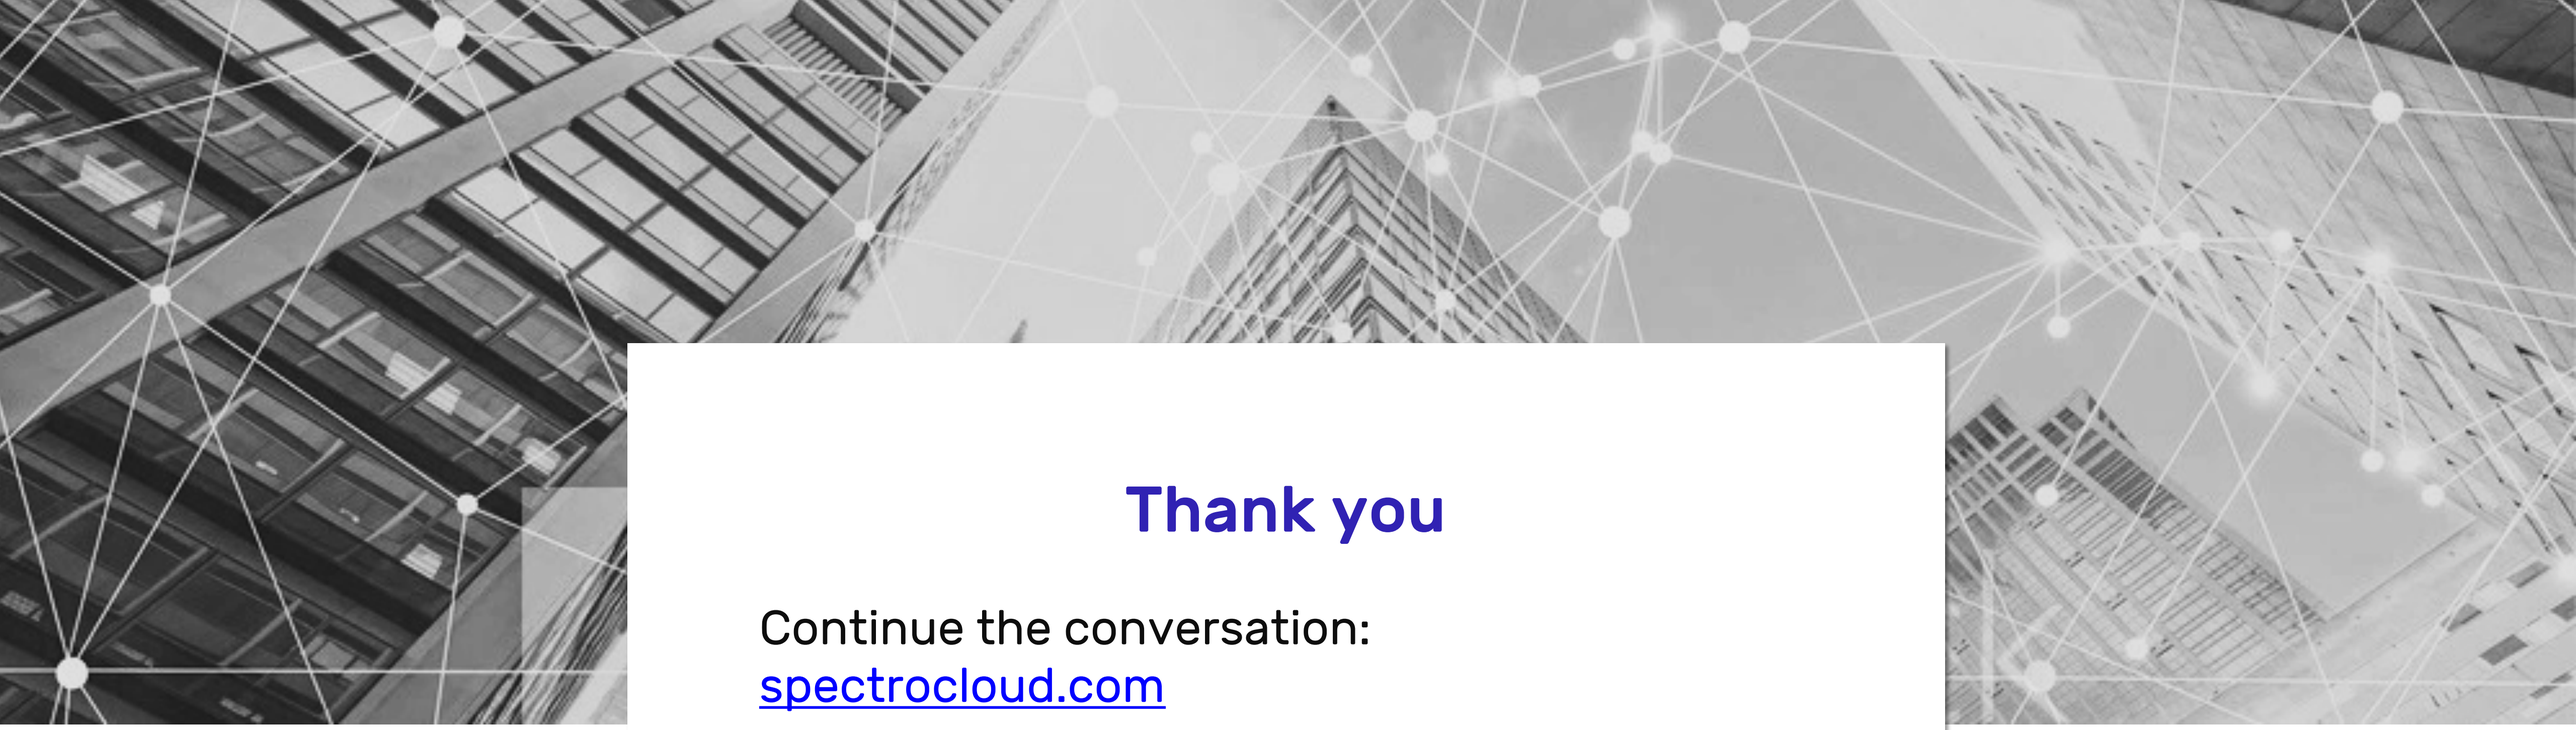

Thank you

Continue the conversation:
spectrocloud.com

Join our community:
slack.spectrocloud.com