How Does Google Release Kubernetes in GKE

KubeCon Barcelona
May 2019

Google Cloud

Josh Hoak

Senior Software Engineer,
Google Cloud

Kobi Magnezi

Product Manager - Google Kubernetes Engine,
Google

certified kubernetes

KubeCon

Google Cloud

Agenda

1. An Intro to Google Kubernetes Engine
2. Hassle-free Kubernetes
3. Kubernetes Release Management at Google
4. Managing Your Risk and Disruption
5. Your clusters everywhere

Google Cloud

An Intro to Google Kubernetes Engine

Google Cloud

Google Kubernetes Engine Overview

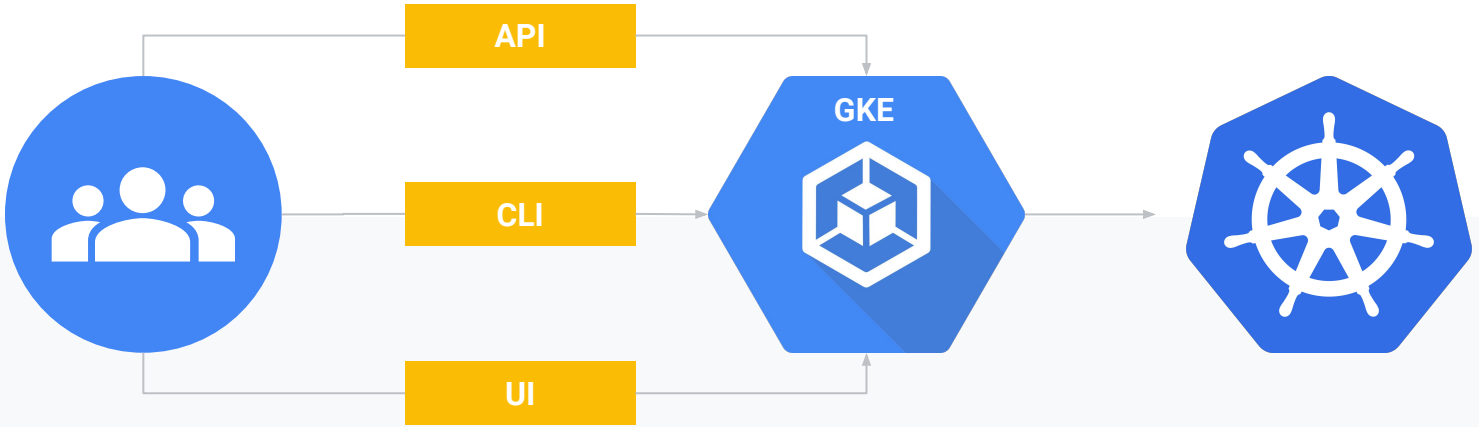

Google Cloud

After Cluster Creation

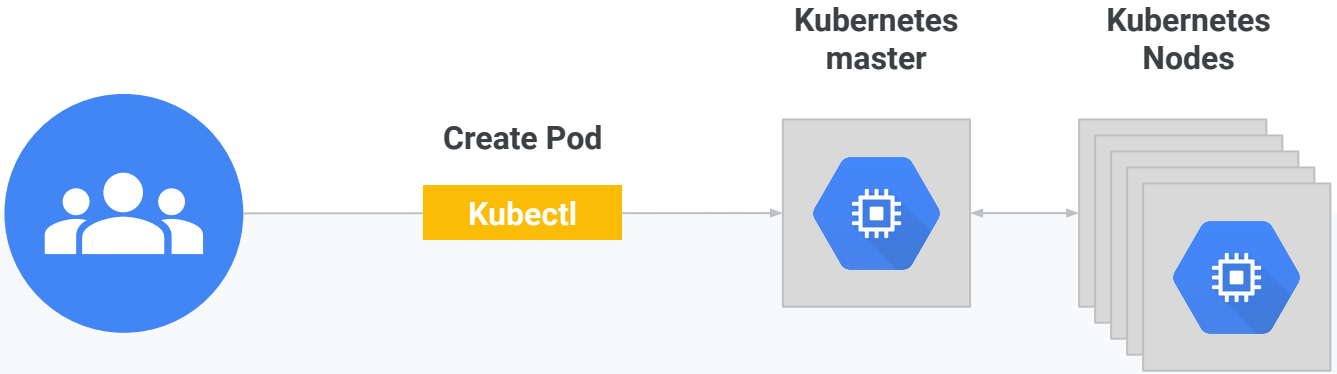

Google Cloud

Master Components

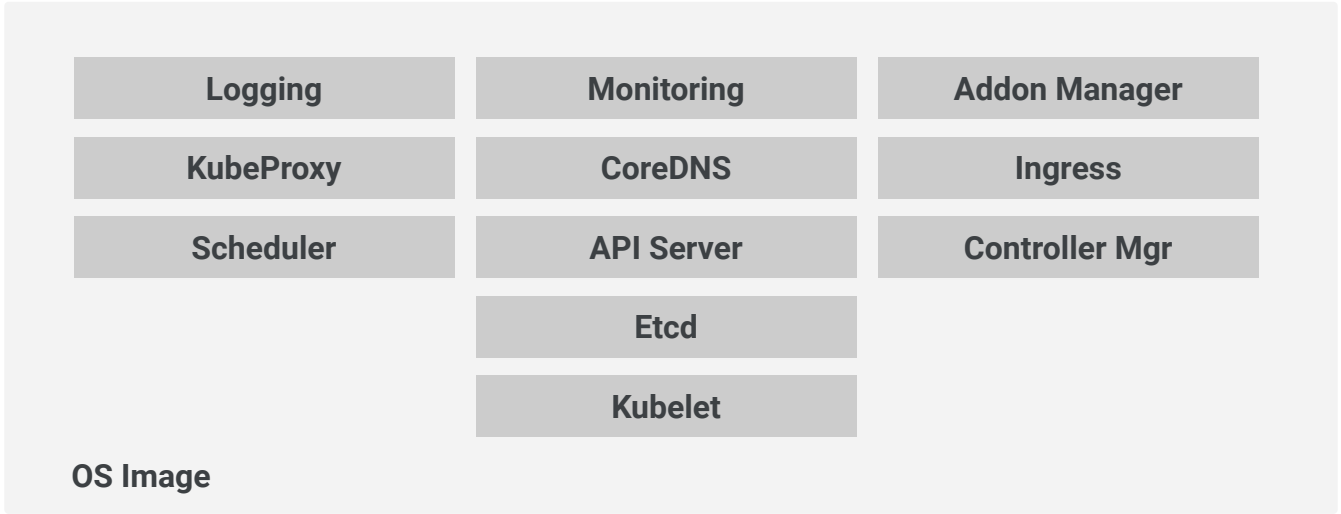

Node Components

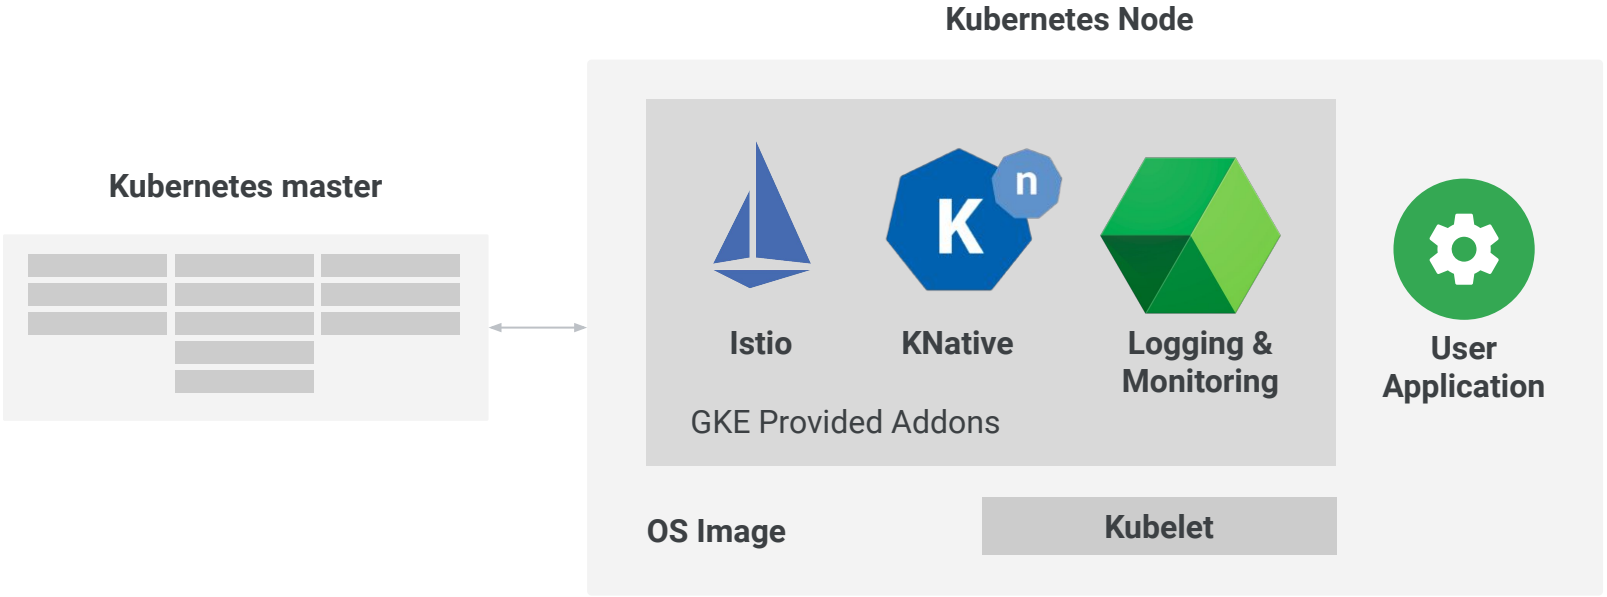

Google Cloud

Kubernetes

A microservices framework built on microservices!

But how do we release and qualify it?

Google Cloud

Hassle-free Kubernetes

Google Cloud

Kubernetes versions in GKE *(2018-12-3)*

```
gloud container get-server-config
--zone=us-central1-b
Fetching server config for us-central1-b
defaultClusterVersion: 1.9.7-gke.11
defaultImageType: COS
validImageTypes:
- UBUNTU
- COS_CONTAINERD
- COS
validMasterVersions:
- 1.11.3-gke.18
- 1.11.2-gke.20
- 1.11.2-gke.18
- 1.10.9-gke.7
- 1.10.9-gke.5
- 1.10.7-gke.13
- 1.10.7-gke.11
- 1.10.6-gke.13
- 1.10.6-gke.11
- 1.9.7-gke.11
```

```
validNodeVersions:
- 1.11.3-gke.18
- 1.11.2-gke.20
- 1.11.2-gke.18
- 1.11.2-gke.15
- 1.11.2-gke.9
- 1.10.9-gke.7
- 1.10.9-gke.5
- 1.10.9-gke.3
- 1.10.9-gke.0
- 1.10.7-gke.13
- 1.10.7-gke.11
- 1.10.7-gke.9
- 1.10.7-gke.6
- 1.10.7-gke.2
- 1.10.7-gke.1
- 1.10.6-gke.13
- 1.10.6-gke.11
- 1.10.6-gke.9
- 1.10.6-gke.6
- 1.10.6-gke.4
- 1.10.6-gke.3
- 1.10.6-gke.2
- 1.10.6-gke.1
```

```
- 1.10.5-gke.4
- 1.10.5-gke.3
- 1.10.5-gke.2
- 1.10.5-gke.0
- 1.10.4-gke.3
- 1.10.4-gke.2
- 1.10.4-gke.0
- 1.10.2-gke.4
- 1.10.2-gke.3
- 1.10.2-gke.1
- 1.9.7-gke.11
- 1.9.7-gke.7
- 1.9.7-gke.6
- 1.9.7-gke.5
- 1.9.7-gke.4
- 1.9.7-gke.3
- 1.9.7-gke.1
- 1.9.7-gke.0
- 1.9.6-gke.2
- 1.9.6-gke.1
- 1.9.3-gke.0
- 1.8.12-gke.3
- 1.8.12-gke.2
```

```
- 1.8.12-gke.1
- 1.8.12-gke.0
- 1.8.10-gke.2
- 1.8.10-gke.0
- 1.8.9-gke.1
- 1.8.8-gke.0
- 1.7.15-gke.0
- 1.7.12-gke.2
- 1.6.13-gke.1
```

Google Cloud

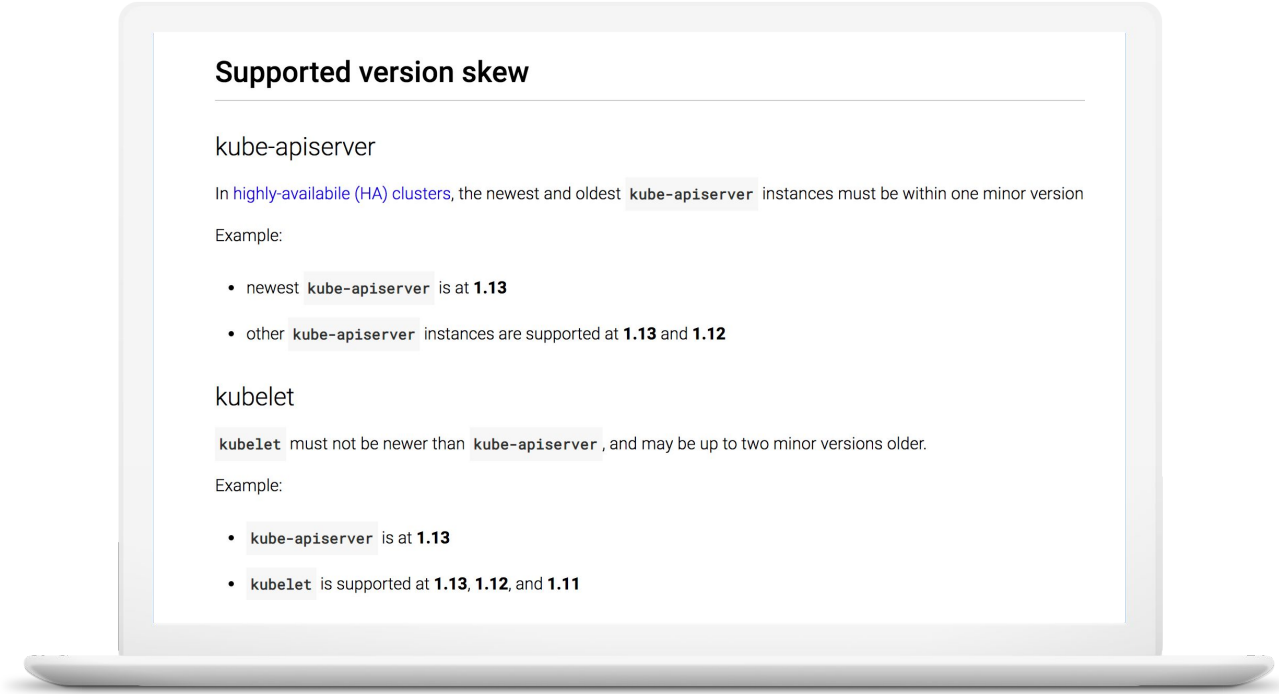

Google Cloud

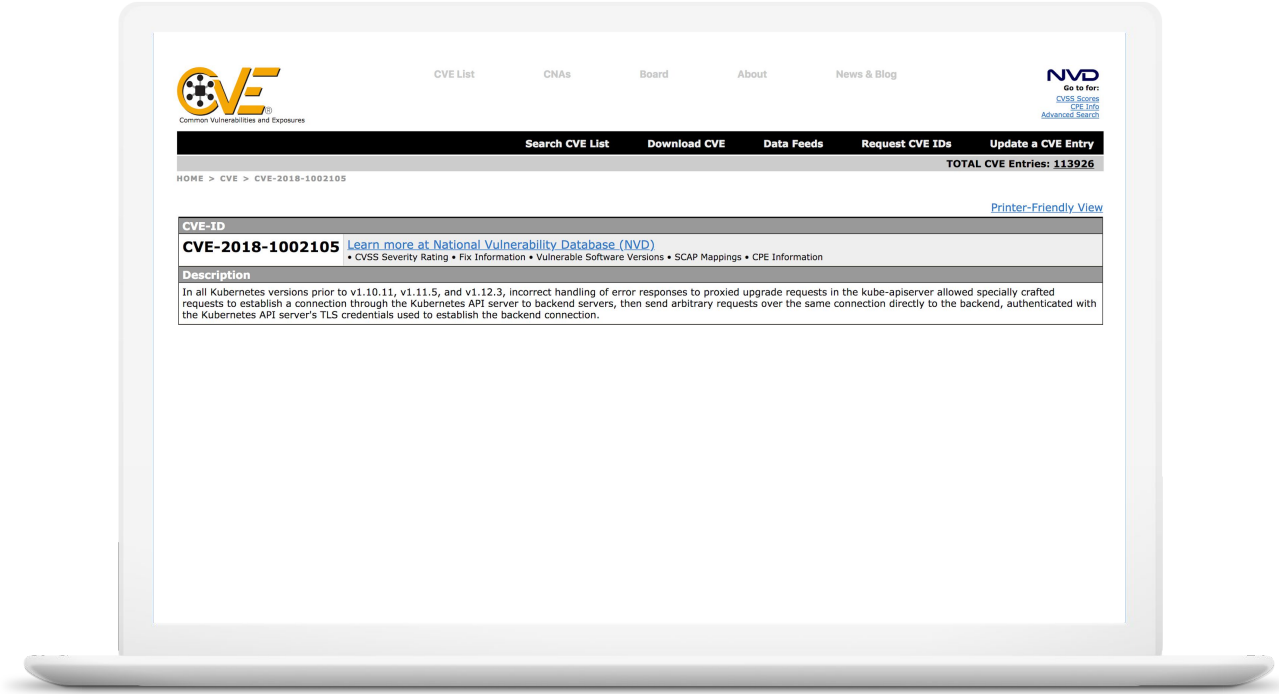

Google Cloud

Managing Kubernetes at scale in GKE

Google Cloud

Open Source to Google

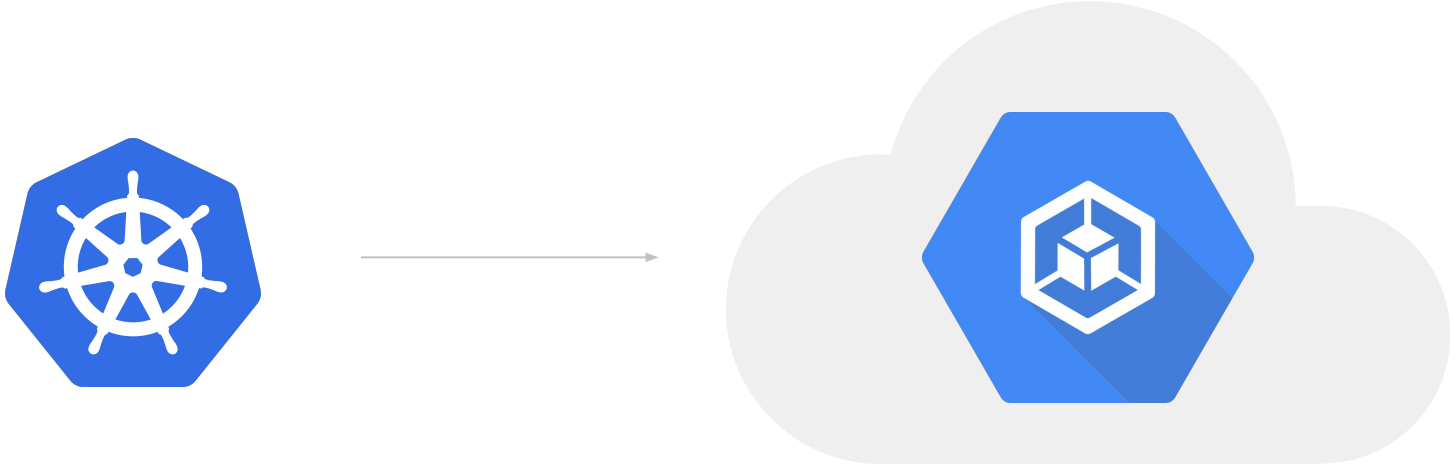

Open Source to Google

1. A new OSS Kubernetes release is cut
2. Import the release and apply patches
3. Qualify the versions internally
4. Cut a release candidate and qualify
5. Roll out to production

Prod Rollout

https://cloud.google.com/kubernetes-engine/docs/release-notes

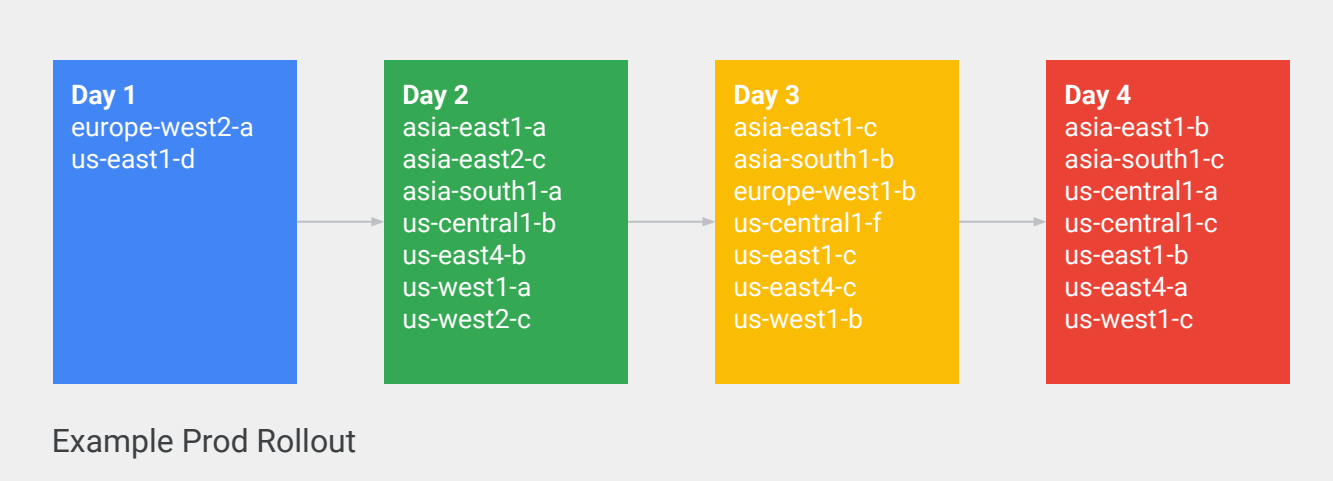

Example Prod Rollout

What's in a Version?

Google Cloud

What's in a Version?

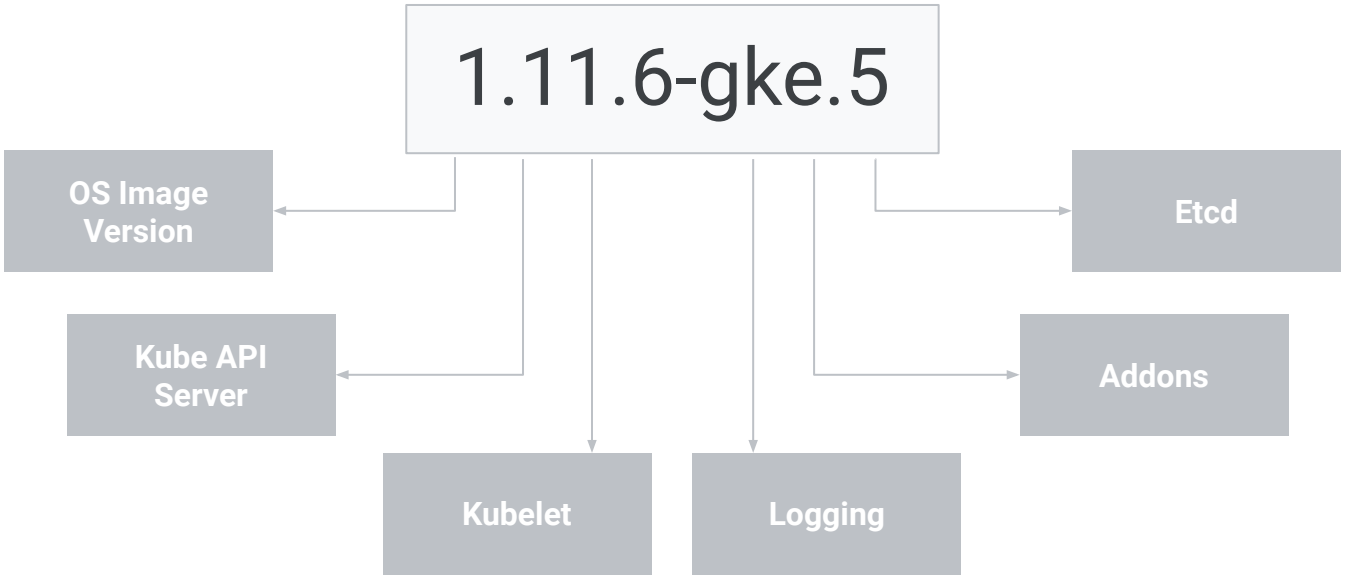

Google Cloud

Cluster Version

GKE Clusters have two versions

Master Version

The version of the Kubernetes on the Master. This encapsulates most of the Kubernetes software.

Node Version

The version of Kubernetes on the Node. In practice, this is mostly the version of the OS Image and the version of the Kubelet.

Keeping your cluster up to date and risk-free

Our customers can focus on delivering value for their end users while we keep their clusters up and running with the latest security patches and bug fixes.

There are two types of upgrades: **Upgrading the Masters** and **Upgrading the Nodes**.

Upgrading a Master

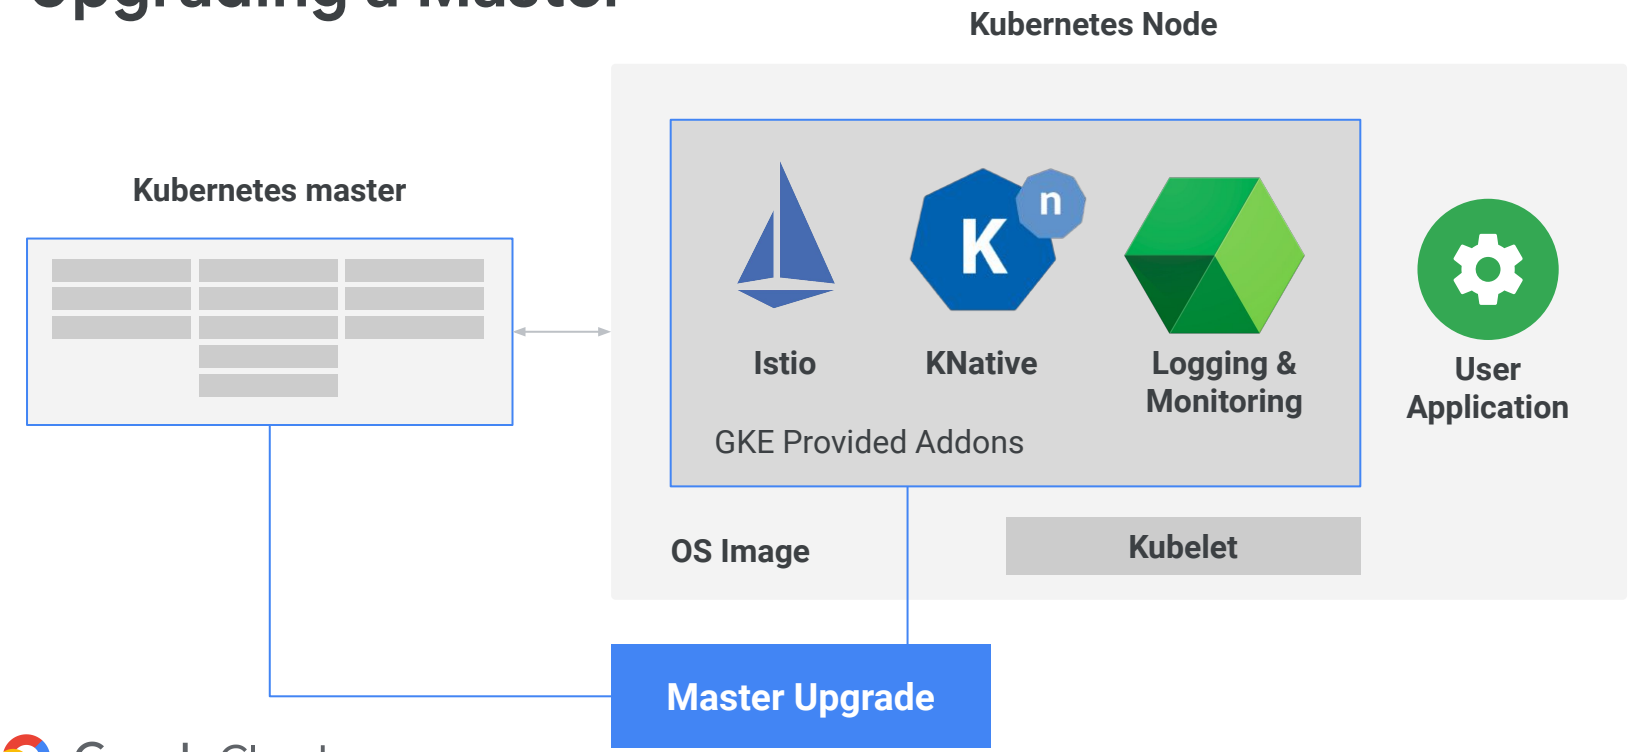

Google Cloud

Upgrade Masters: Zonal/Regional

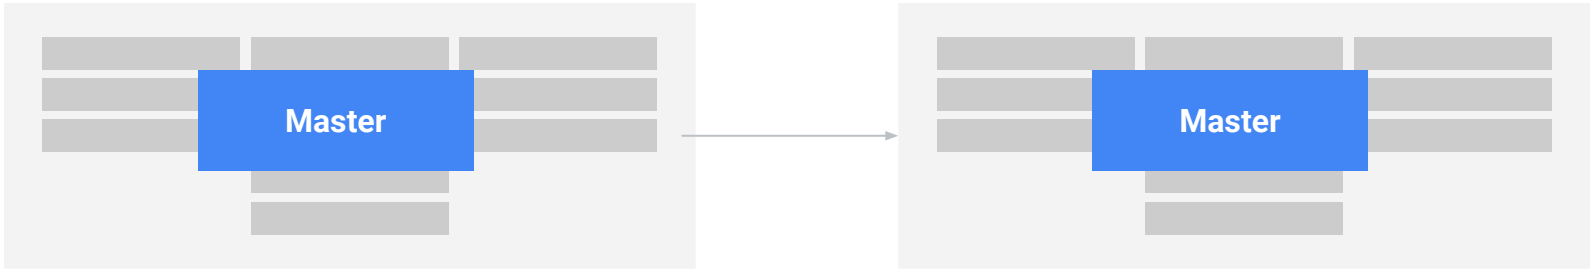

Upgrading a Node

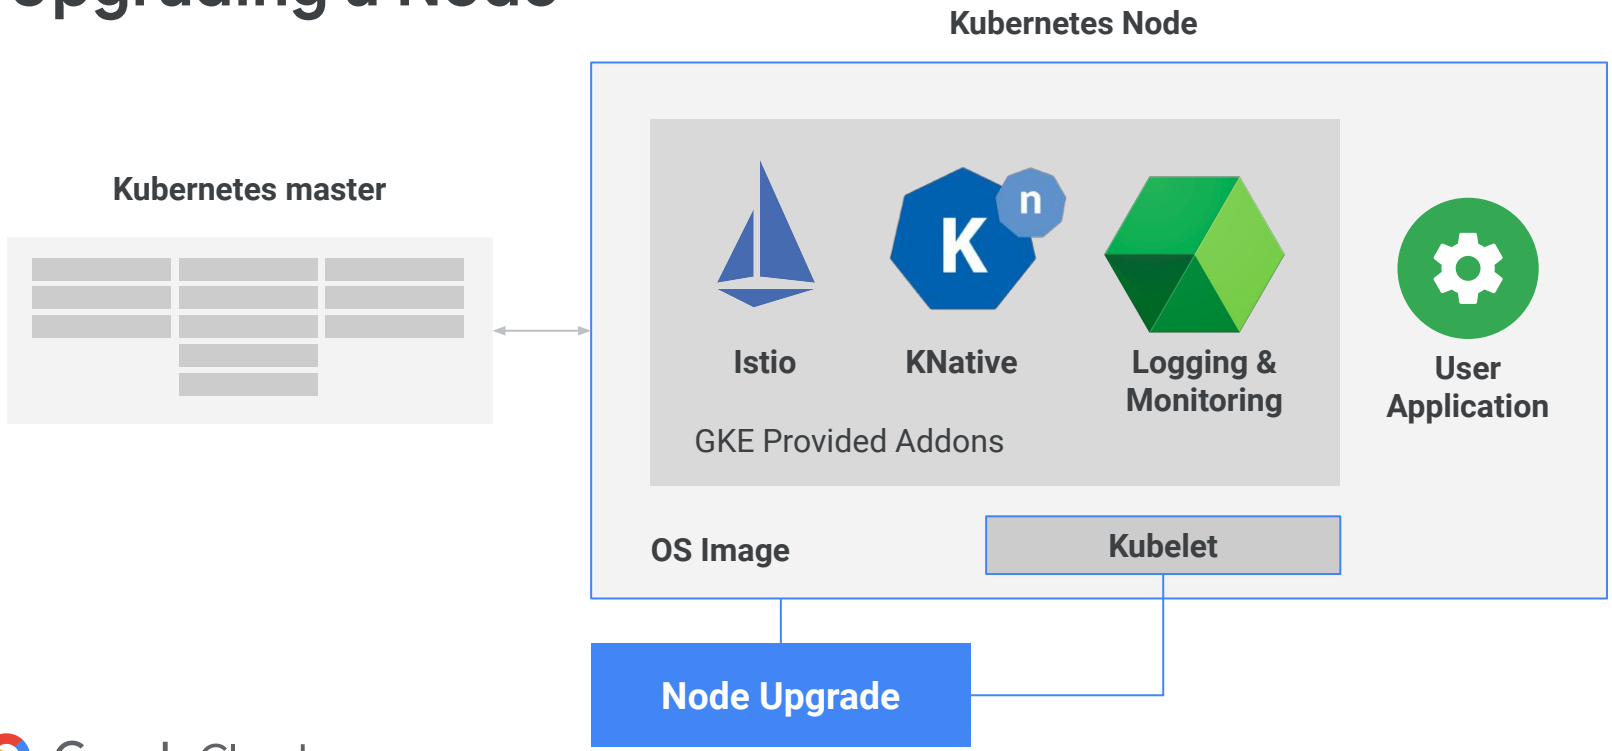

Google Cloud

Version Lifecycle

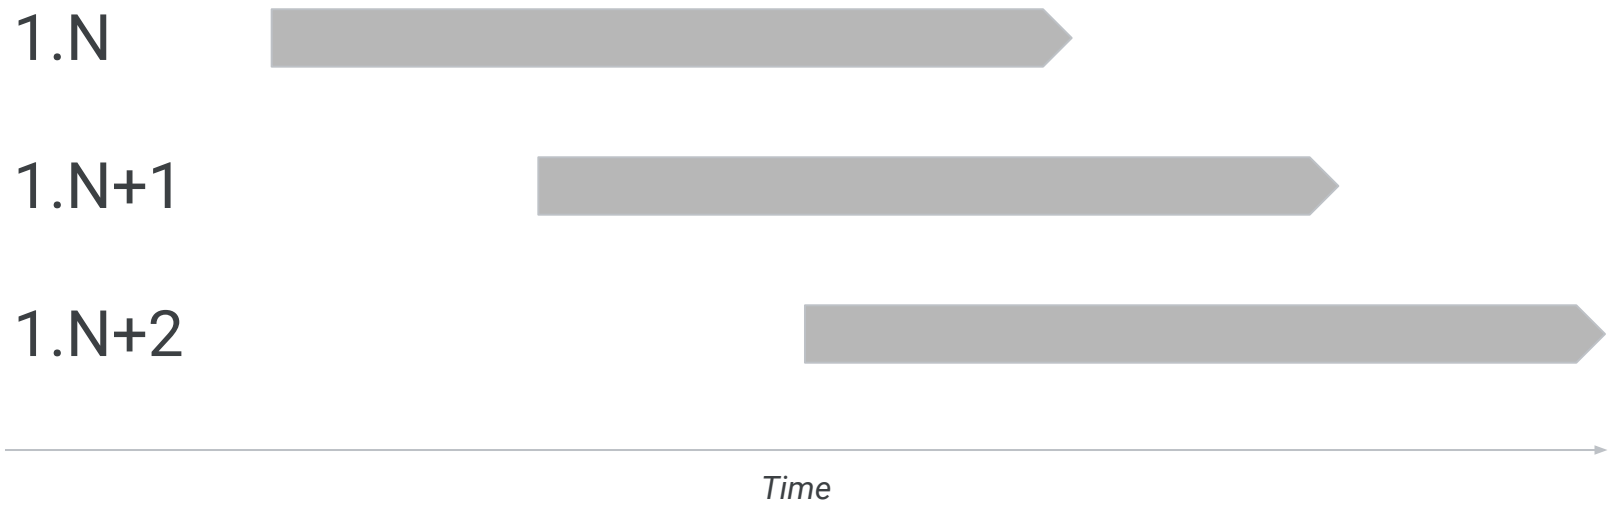

Google Cloud

Version Lifecycle

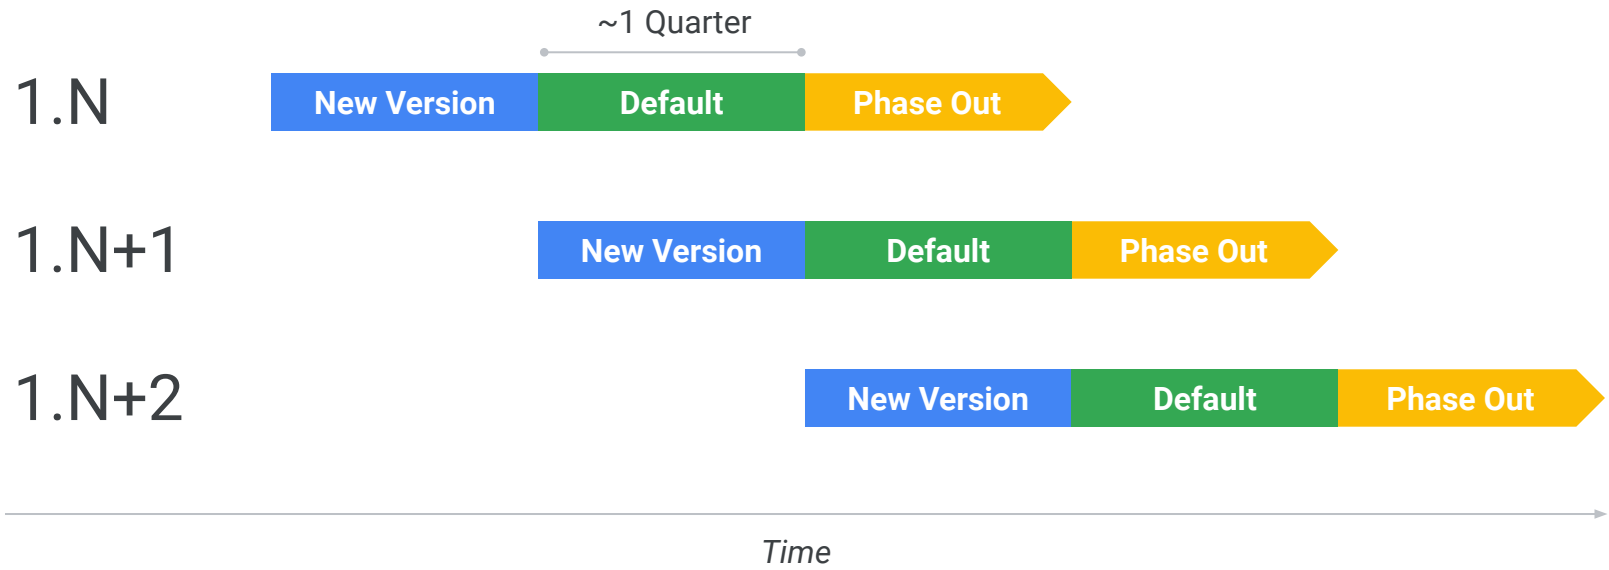

Upgrade Schedule: Masters

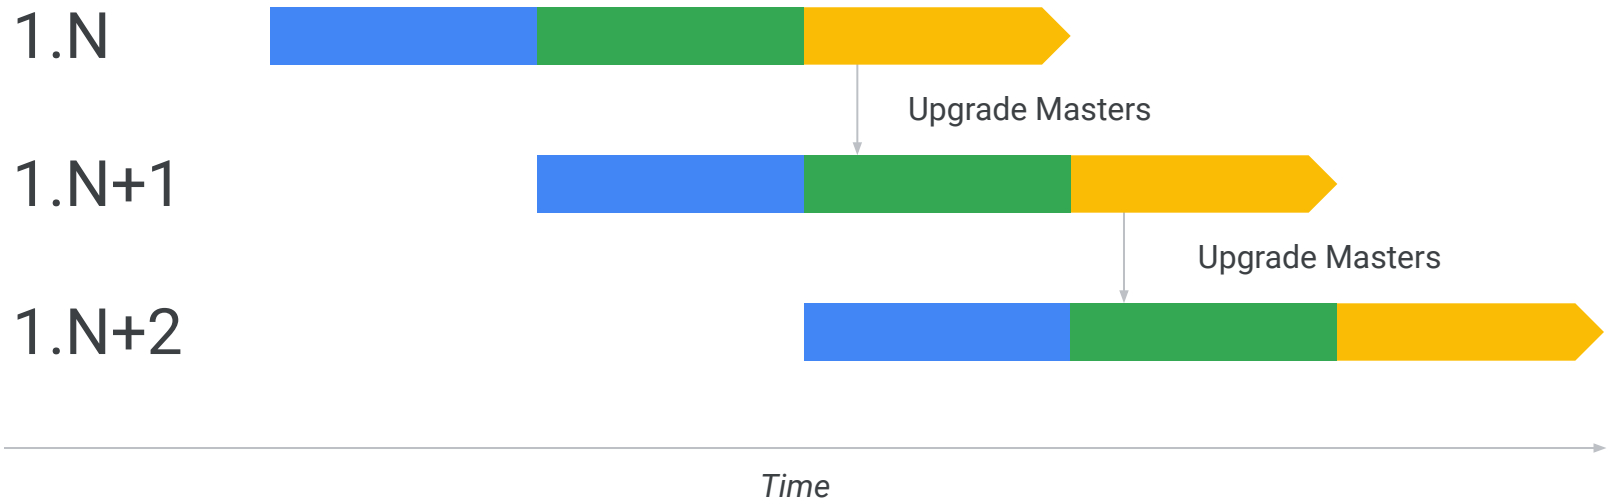

Upgrade Schedule: Nodes

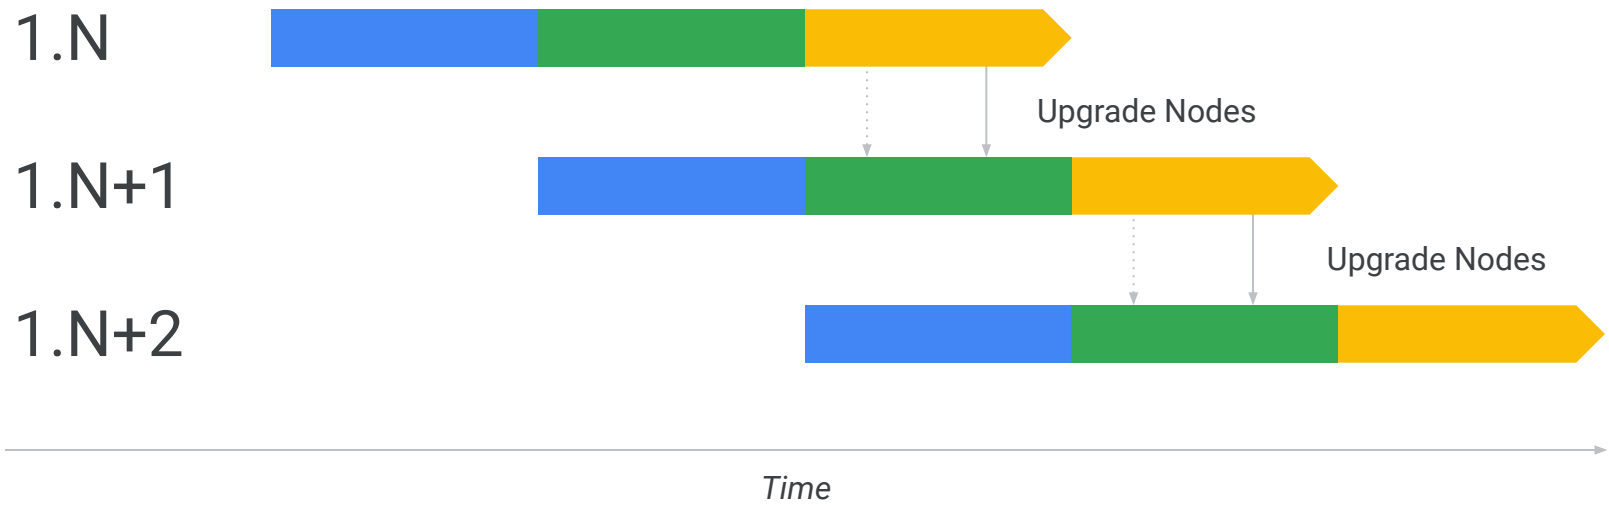

Google Cloud

Version Lifecycle: Patches

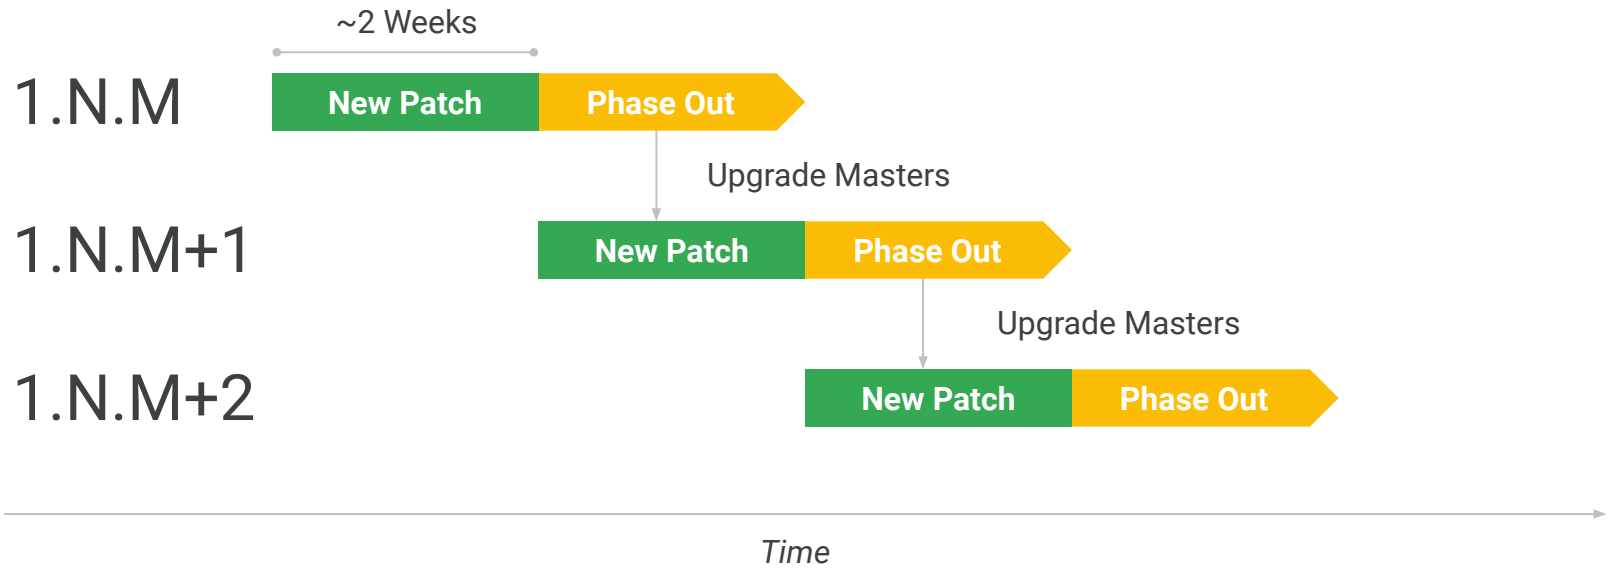

Google Cloud

Managing Your Risk

04

Customer Personas - One size doesn't fit all

Customer A have a multi- weeks certification cycles for every release. They are looking to align their internal certification process with GKE releases.

Google Cloud

Customer Personas - One size doesn't fit all

Customer B upgrade manually today for better control on what they get in every release. They're excited about the ability to test early versions of K8s and get a better time-to-market with new capabilities.

Google Cloud

Customer Personas - One size doesn't fit all

Customer C know how to operate Kubernetes, and want to adopt new features quickly, but roll them out to multiple clusters carefully. They rather validate upgrade on staging environment and control the roll out to different clusters targeting different users.

Google Cloud

Release Notes

We maintain extensive release notes on both the Kubernetes release page and on the GKE release page. This is your best ally in learning about new features, bugs, breaking changes, rollout schedule, and deprecated versions.

Kubernetes Releases
https://github.com/kubernetes/kubernetes/releases

GKE Releases:
https://cloud.google.com/kubernetes-engine/docs/release-notes

Release Notes

New versions available for upgrades and new clusters

The following Kubernetes versions are now available for new clusters and for opt-in master upgrades and node upgrades for existing clusters:

V1.11.X | V1.12.X | V1.13.X | RAPID CHANNEL

v1.11.9-gke.13

- Improvements to Vertical Pod Autoscaler
- Improvements to Cluster Autoscaler
- Cloud Run for GKE now uses the default Istio sidecar injection behavior
- Fix an issue that prevented the kubelet from seeing all GPUs available to nodes using the Ubuntu node image.

Google Cloud

Picking a Location

Later zones receive updates later. Outages are often detected in Day 1 and Day 2 locations first.

Strategy: Serve production workloads in later locations, canary in earlier locations.

Our rollout plan is part of our Release Notes:
https://cloud.google.com/kubernetes-engine/docs/release-notes

Picking a Location

Date	Available zones	Available regions
2019-05-13	europe-west2-a, us-east1-d	europe-west3, us-east1
2019-05-14	asia-east1-a, asia-east2-c, asia-northeast1-a, asia-northeast2-c, asia-south1-a, asia-southeast1-a, australia-southeast1-a, europe-north1-c, europe-west1-c, europe-west3-a, europe-west4-a, europe-west6-c, northamerica-northeast1-c, southamerica-east1-a, us-central1-b, us-east4-b, us-west1-a, us-west2-c	asia-east1, asia-southeast1, europe-west6, northamerica-northeast1, us-east4, us-west2
2019-05-15	asia-east1-c, asia-east2-b, asia-northeast1-b, asia-northeast2-b, asia-south1-b, asia-southeast1-b, australia-southeast1-b, europe-north1-b, europe-west1-b, europe-west2-b, europe-west3-b, europe-west4-c, europe-west6-b, northamerica-northeast1-b, southamerica-east1-b, us-central1-f, us-east1-c, us-east4-c, us-west1-b, us-west2-b	asia-east2, asia-northeast1, australia-southeast1, europe-west1, europe-west2, southamerica-east1, us-west1
2019-05-16	asia-east1-b, asia-east2-a, asia-northeast1-c, asia-northeast2-a, asia-south1-c, asia-southeast1-c, australia-southeast1-c, europe-north1-a, europe-west1-d, europe-west2-c, europe-west3-c, europe-west4-b, europe-west6-a, northamerica-northeast1-a, southamerica-east1-c, us-central1-a, us-central1-c, us-east1-b, us-east4-a, us-west1-c, us-west2-a	asia-northeast2, asia-south1, europe-north1, europe-west4, us-central1

Google Cloud

Picking a Version

New versions are inherently more unstable. The **default** version is where we believe the majority of the fleet should be.

Strategy: Use the **default** version for most workloads and use new versions to try out new features and test new versions.

```
gcloud container get-server-config --zone=[ZONE]
```

Release Channels

We now offer ways for customers to opt-into **release channels**, which are streams of Chrome-like, automated updates. Release channels enable customers to choose a release cadence and feature set to match their risk preference.

The Goal: If a change is available in later release channel, it was qualified in a previous release channel.

Google Cloud

Release Channel Example

To create a cluster on a Release Channel, it's as simple as:

```
gcloud alpha container clusters create [CLUSTER_NAME] --release-channel rapid
```

Result: A cluster whose master and nodes are maintained at the newest release in GKE.

The **Rapid** Release Channel will be available in **Alpha** in Q2 2019, with the **Regular** and **Stable** channels following later this year!

Node Auto Upgrades

Node Auto Upgrades allow GKE to auto-upgrade customer nodes

Strategy: We recommend turning on **node auto upgrades** for most clusters.

```
gcloud container clusters create [NAME] --zone [ZONE] --enable-autoupgrade
```

More at https://cloud.google.com/kubernetes-engine/docs/how-to/node-auto-upgrades

Google Cloud

Control disruptions with GKE

Google Cloud

Surge Upgrade

- Securing resources **before** upgrade starts
- Ensure successful completion of node upgrade
- Protect customer capacity
- Reduce disruption
- Reduce downtime

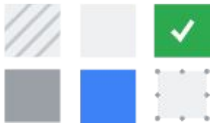

Google Cloud

Upgrading Nodes

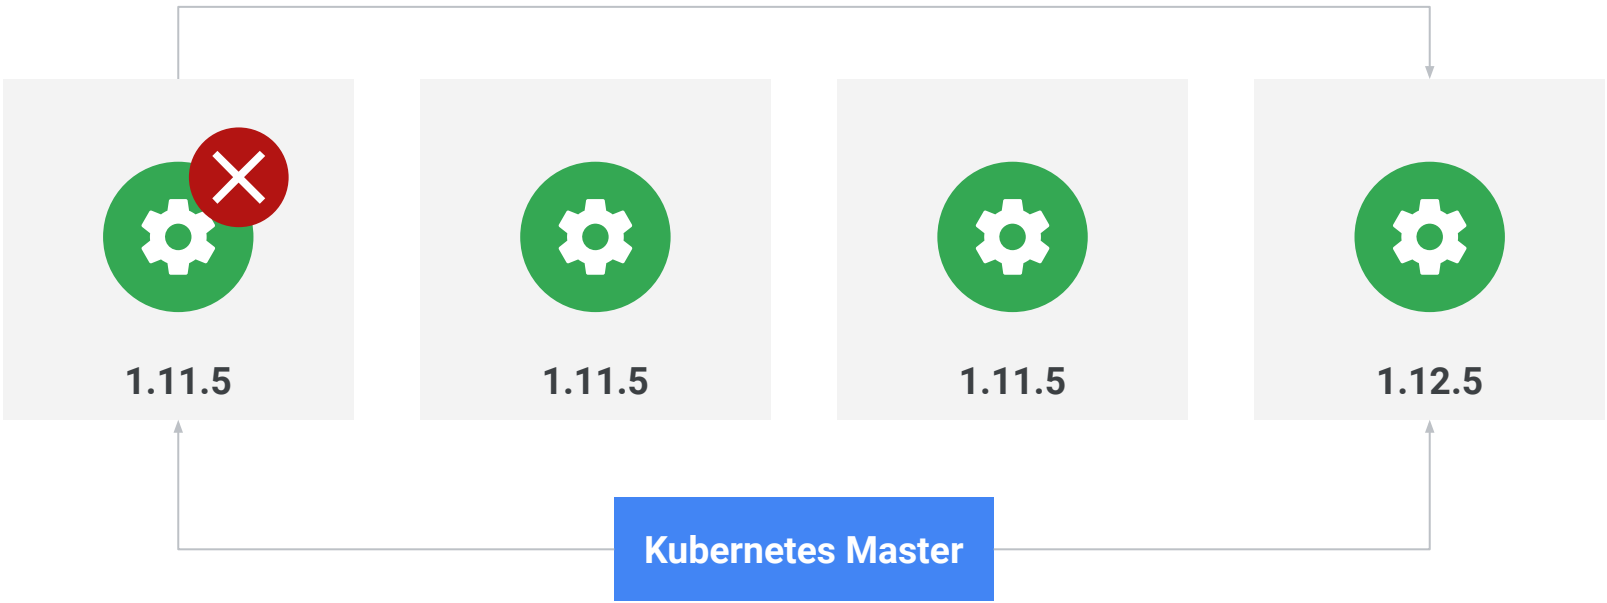

Google Cloud

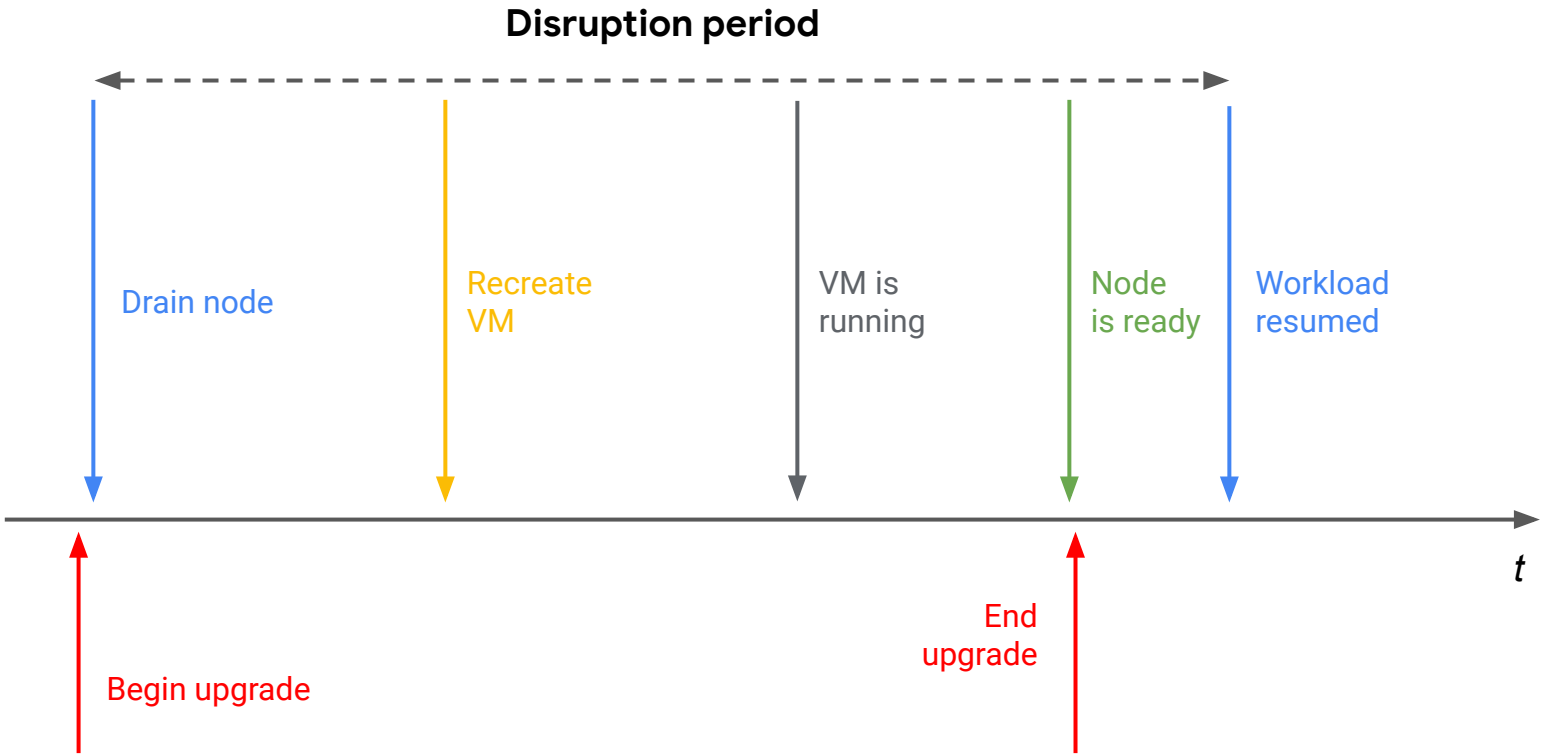

Google Cloud

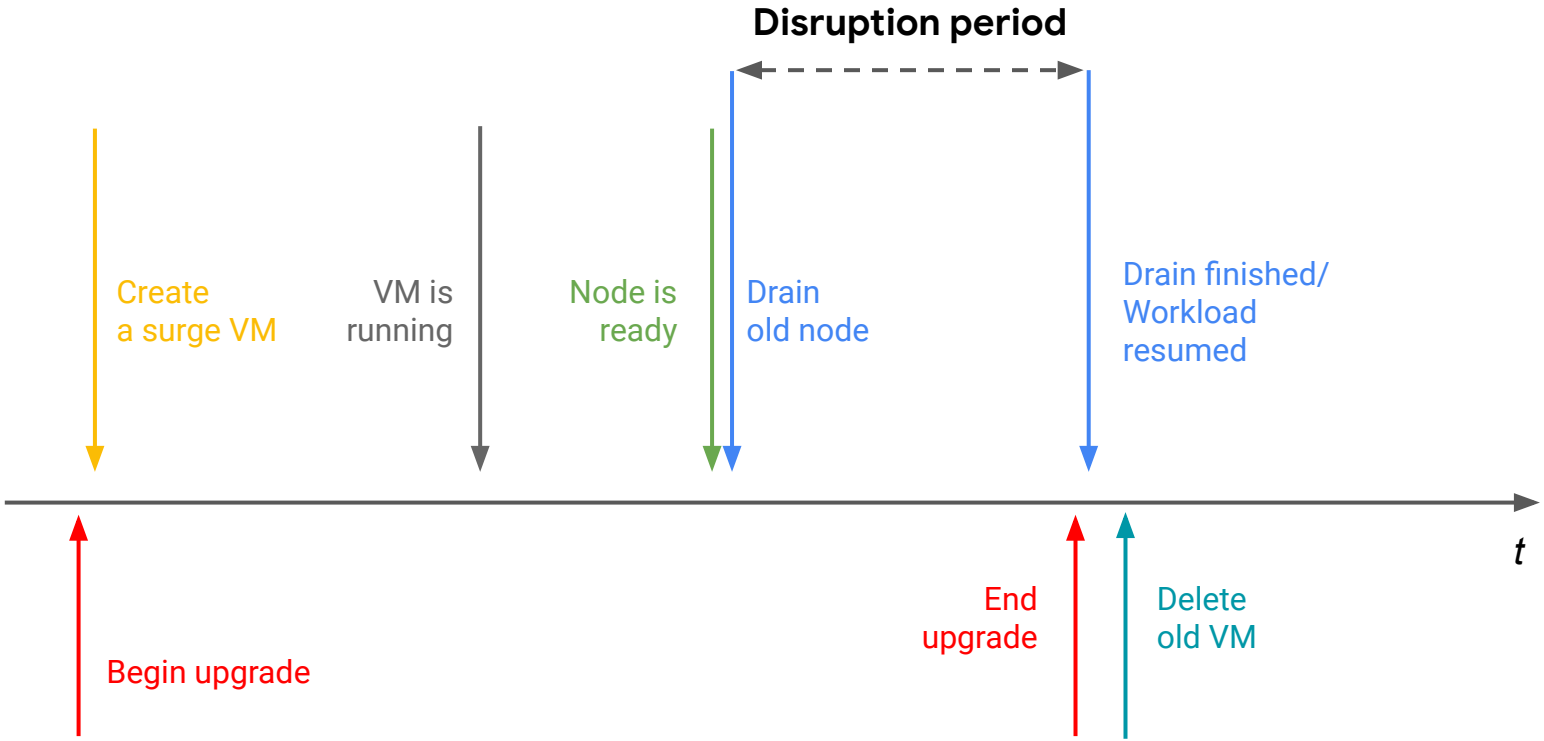

Google Cloud

Maintenance Windows

Maintenance windows allow customers to indicate to GKE when master and node upgrades should occur during the day, within a 4-hour window.

Strategy: If you have times of the day that make more sense for GKE to perform upgrades, check out Maintenance Windows.

```
gcloud container clusters create [CLUSTER_NAME] --maintenance-window [HH:MM]
```

More at https://cloud.google.com/kubernetes-engine/docs/how-to/maintenance-window

Maintenance Windows

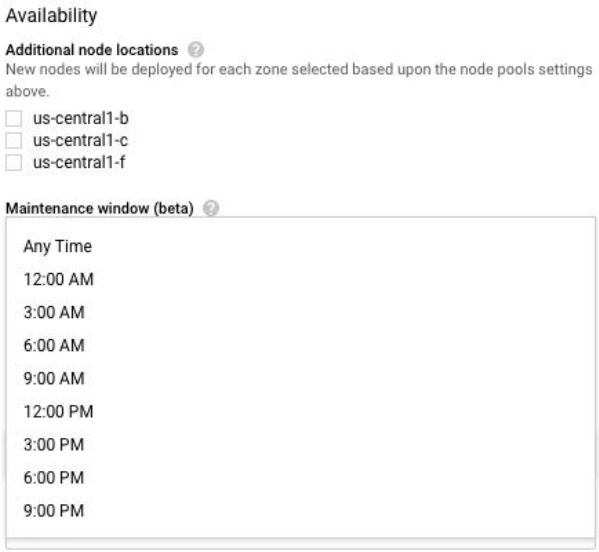

Google Cloud

Enhanced Maintenance Windows

Enhanced Maintenance windows will allow customers broader flexibility to indicate to GKE when master and node upgrades should be avoided.

Use case 1

If you prefer to avoid upgrades during weekends.

Use case 2

You're a retailer, and prefer to avoid upgrades during specific dates (e.g. Black Friday / Cyber Monday)

Use case 3

You need to postpone (snooze) a schedule upgrade for a short period of time.

Canary Clusters

Canary clusters enable customers to create a **canary cluster** and link it with their respective **production cluster(s).** GKE will roll out new versions to canary clusters first, allowing customers to run automated tests on the new version. Roll out to production clusters will be conditioned on the successful pass of the automated test.

The Goal: Mitigate risk of rolling our new versions to existing clusters.

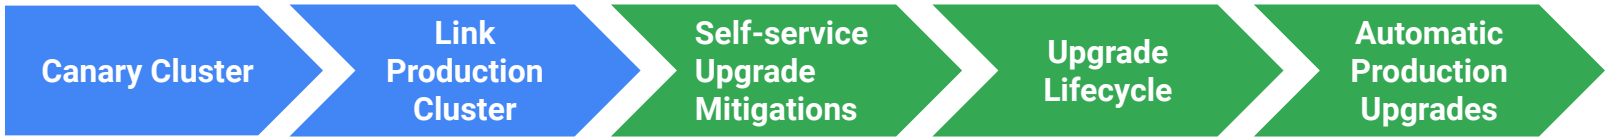

* Details TBD. Subject to change

Component Based Releases

The **Cluster Bundle** is an open source project developed by the GKE team to provide tooling and infrastructure for building better GKE Releases for components. We are currently using the project for both **GKE** and **GKE On-Prem.** Check it out if you're interested in how we think about releases!

See more at https://github.com/GoogleCloudPlatform/k8s-cluster-bundle

Component Based Releases

What does that mean for customers?

More stable releases

Improved time-to-market for components

Better release notes

Thank you

Questions?

Google Cloud