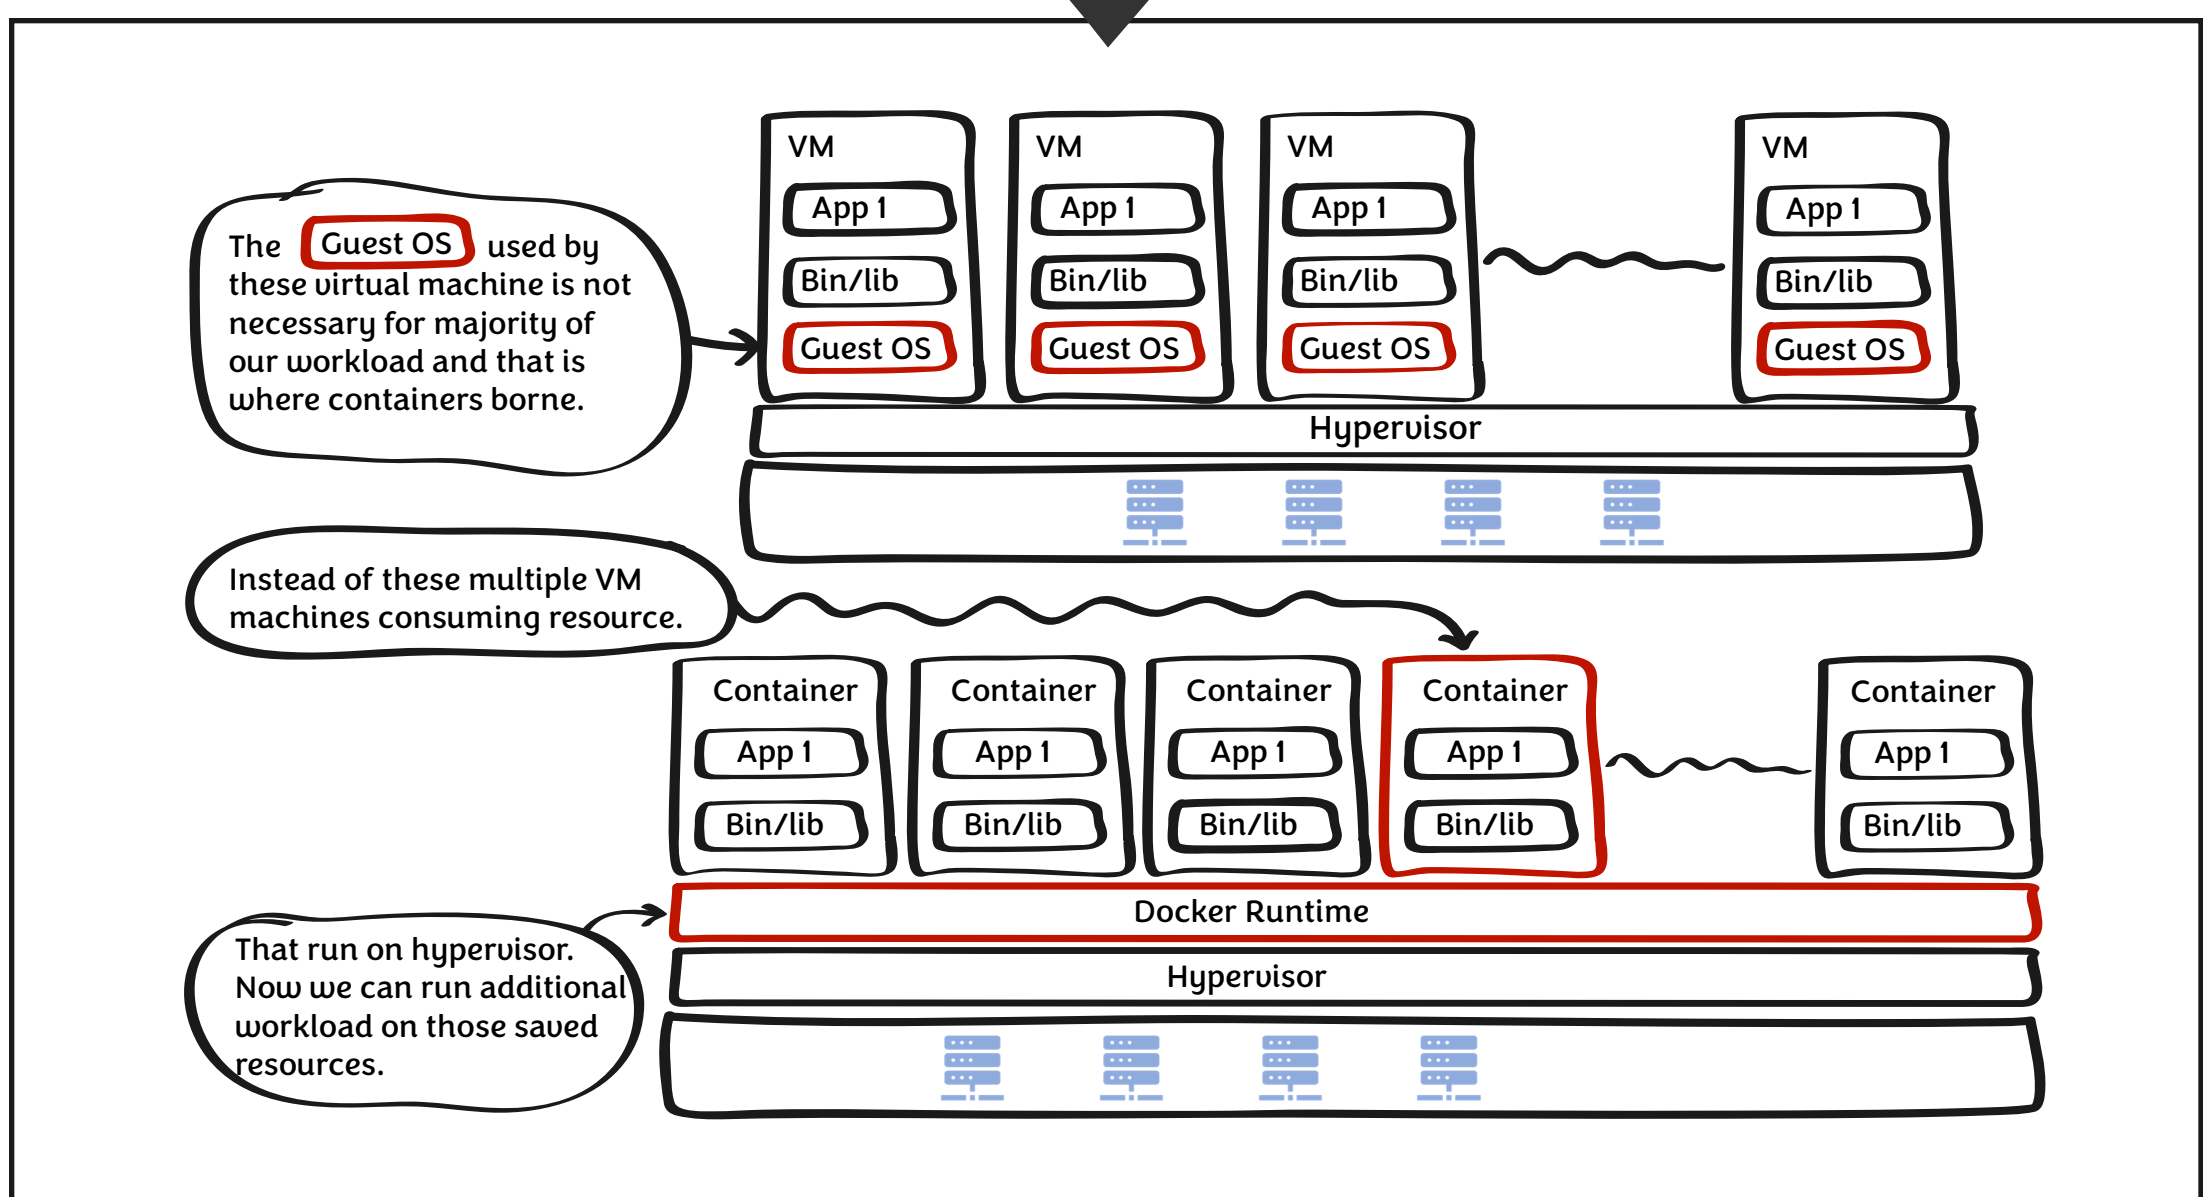
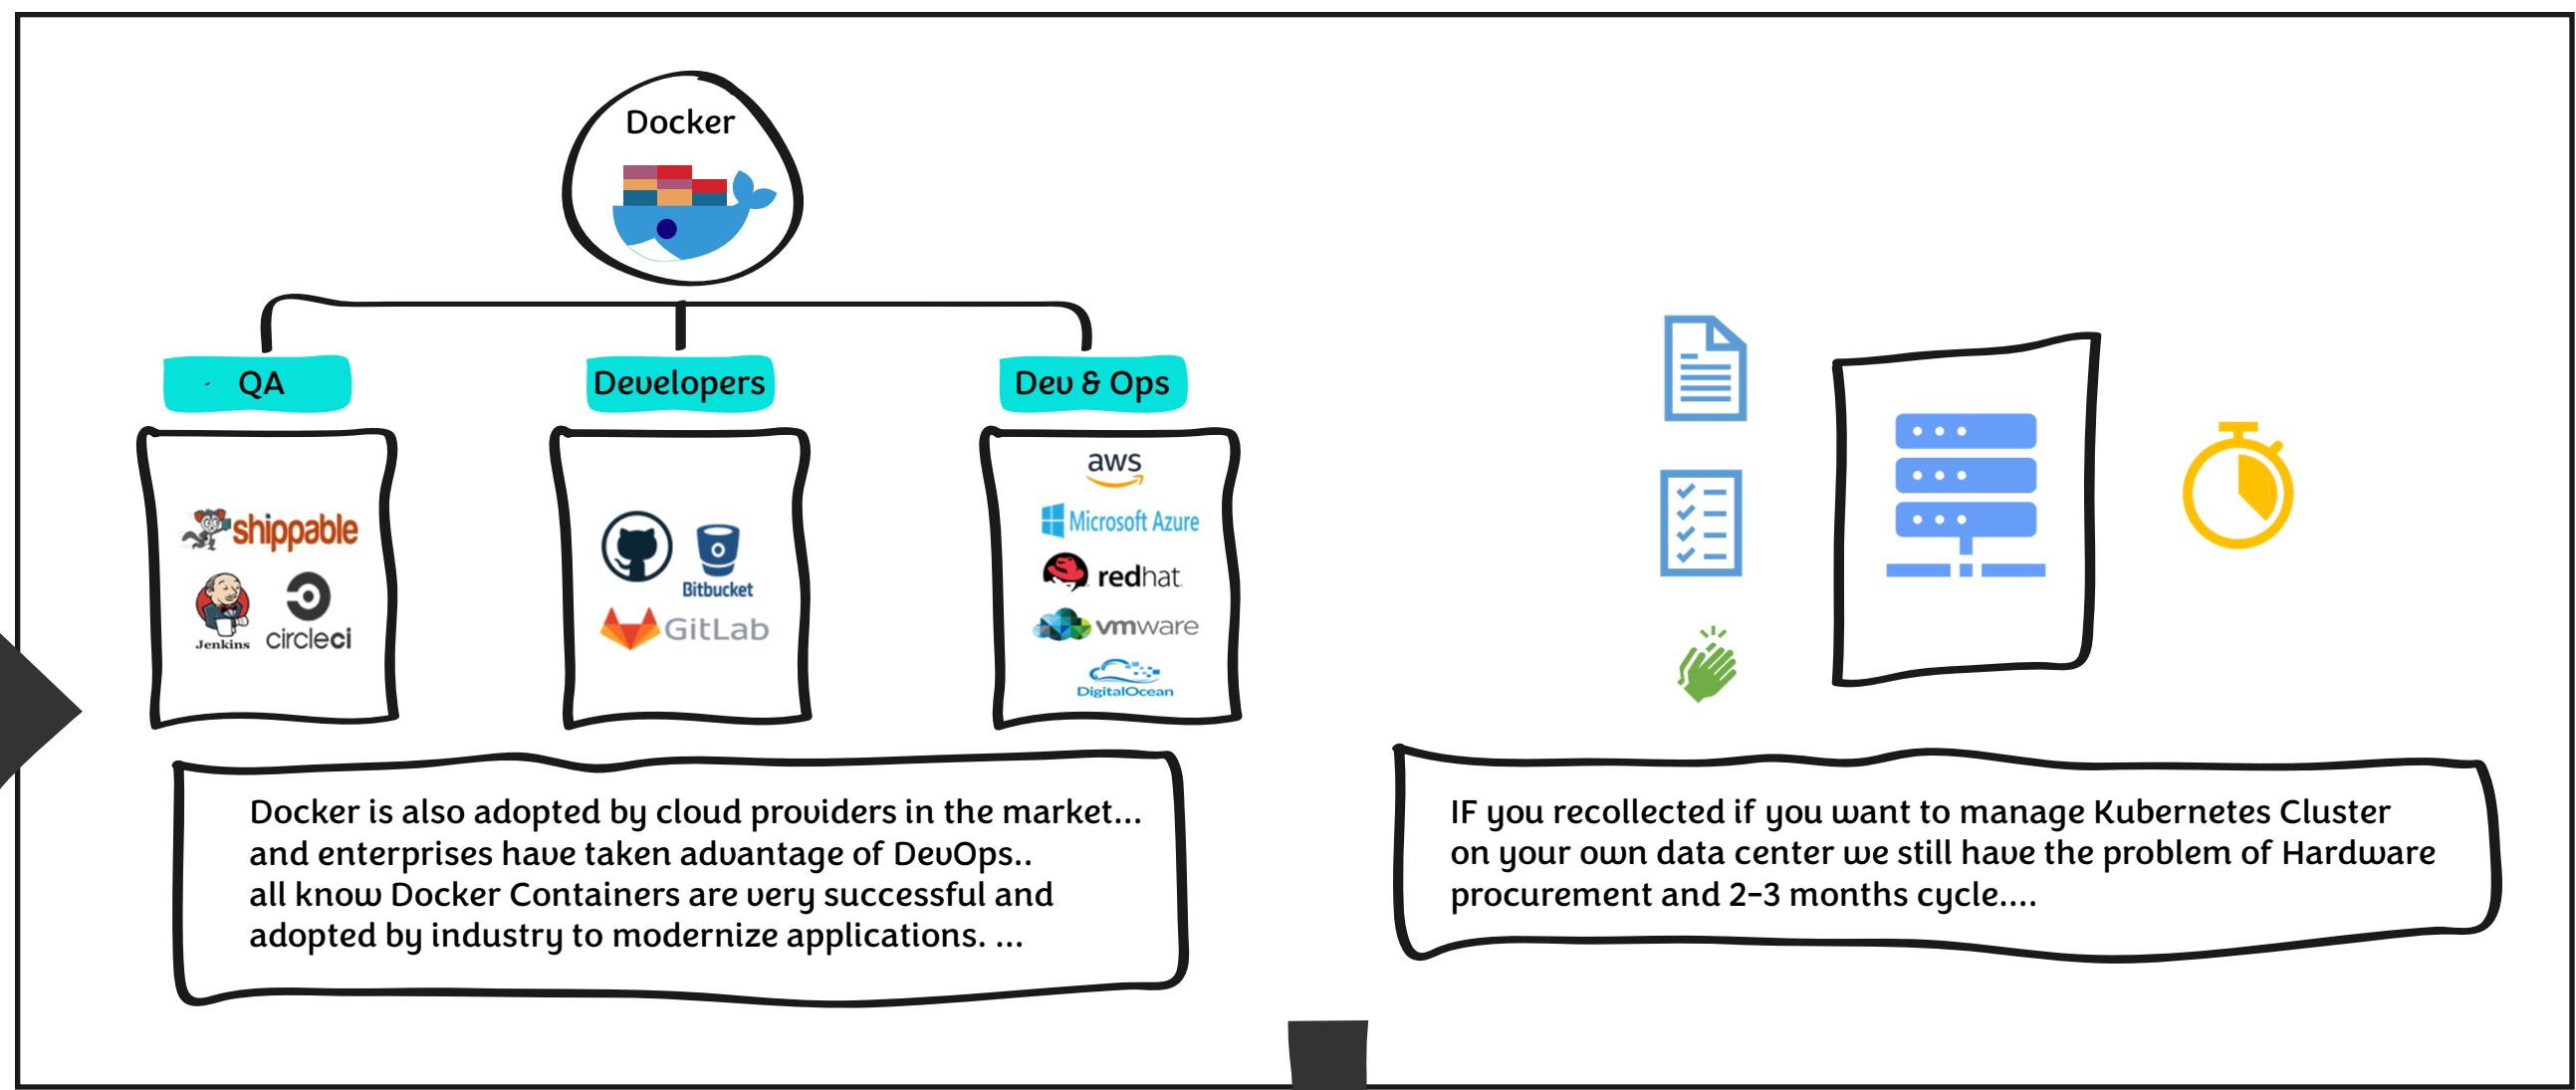
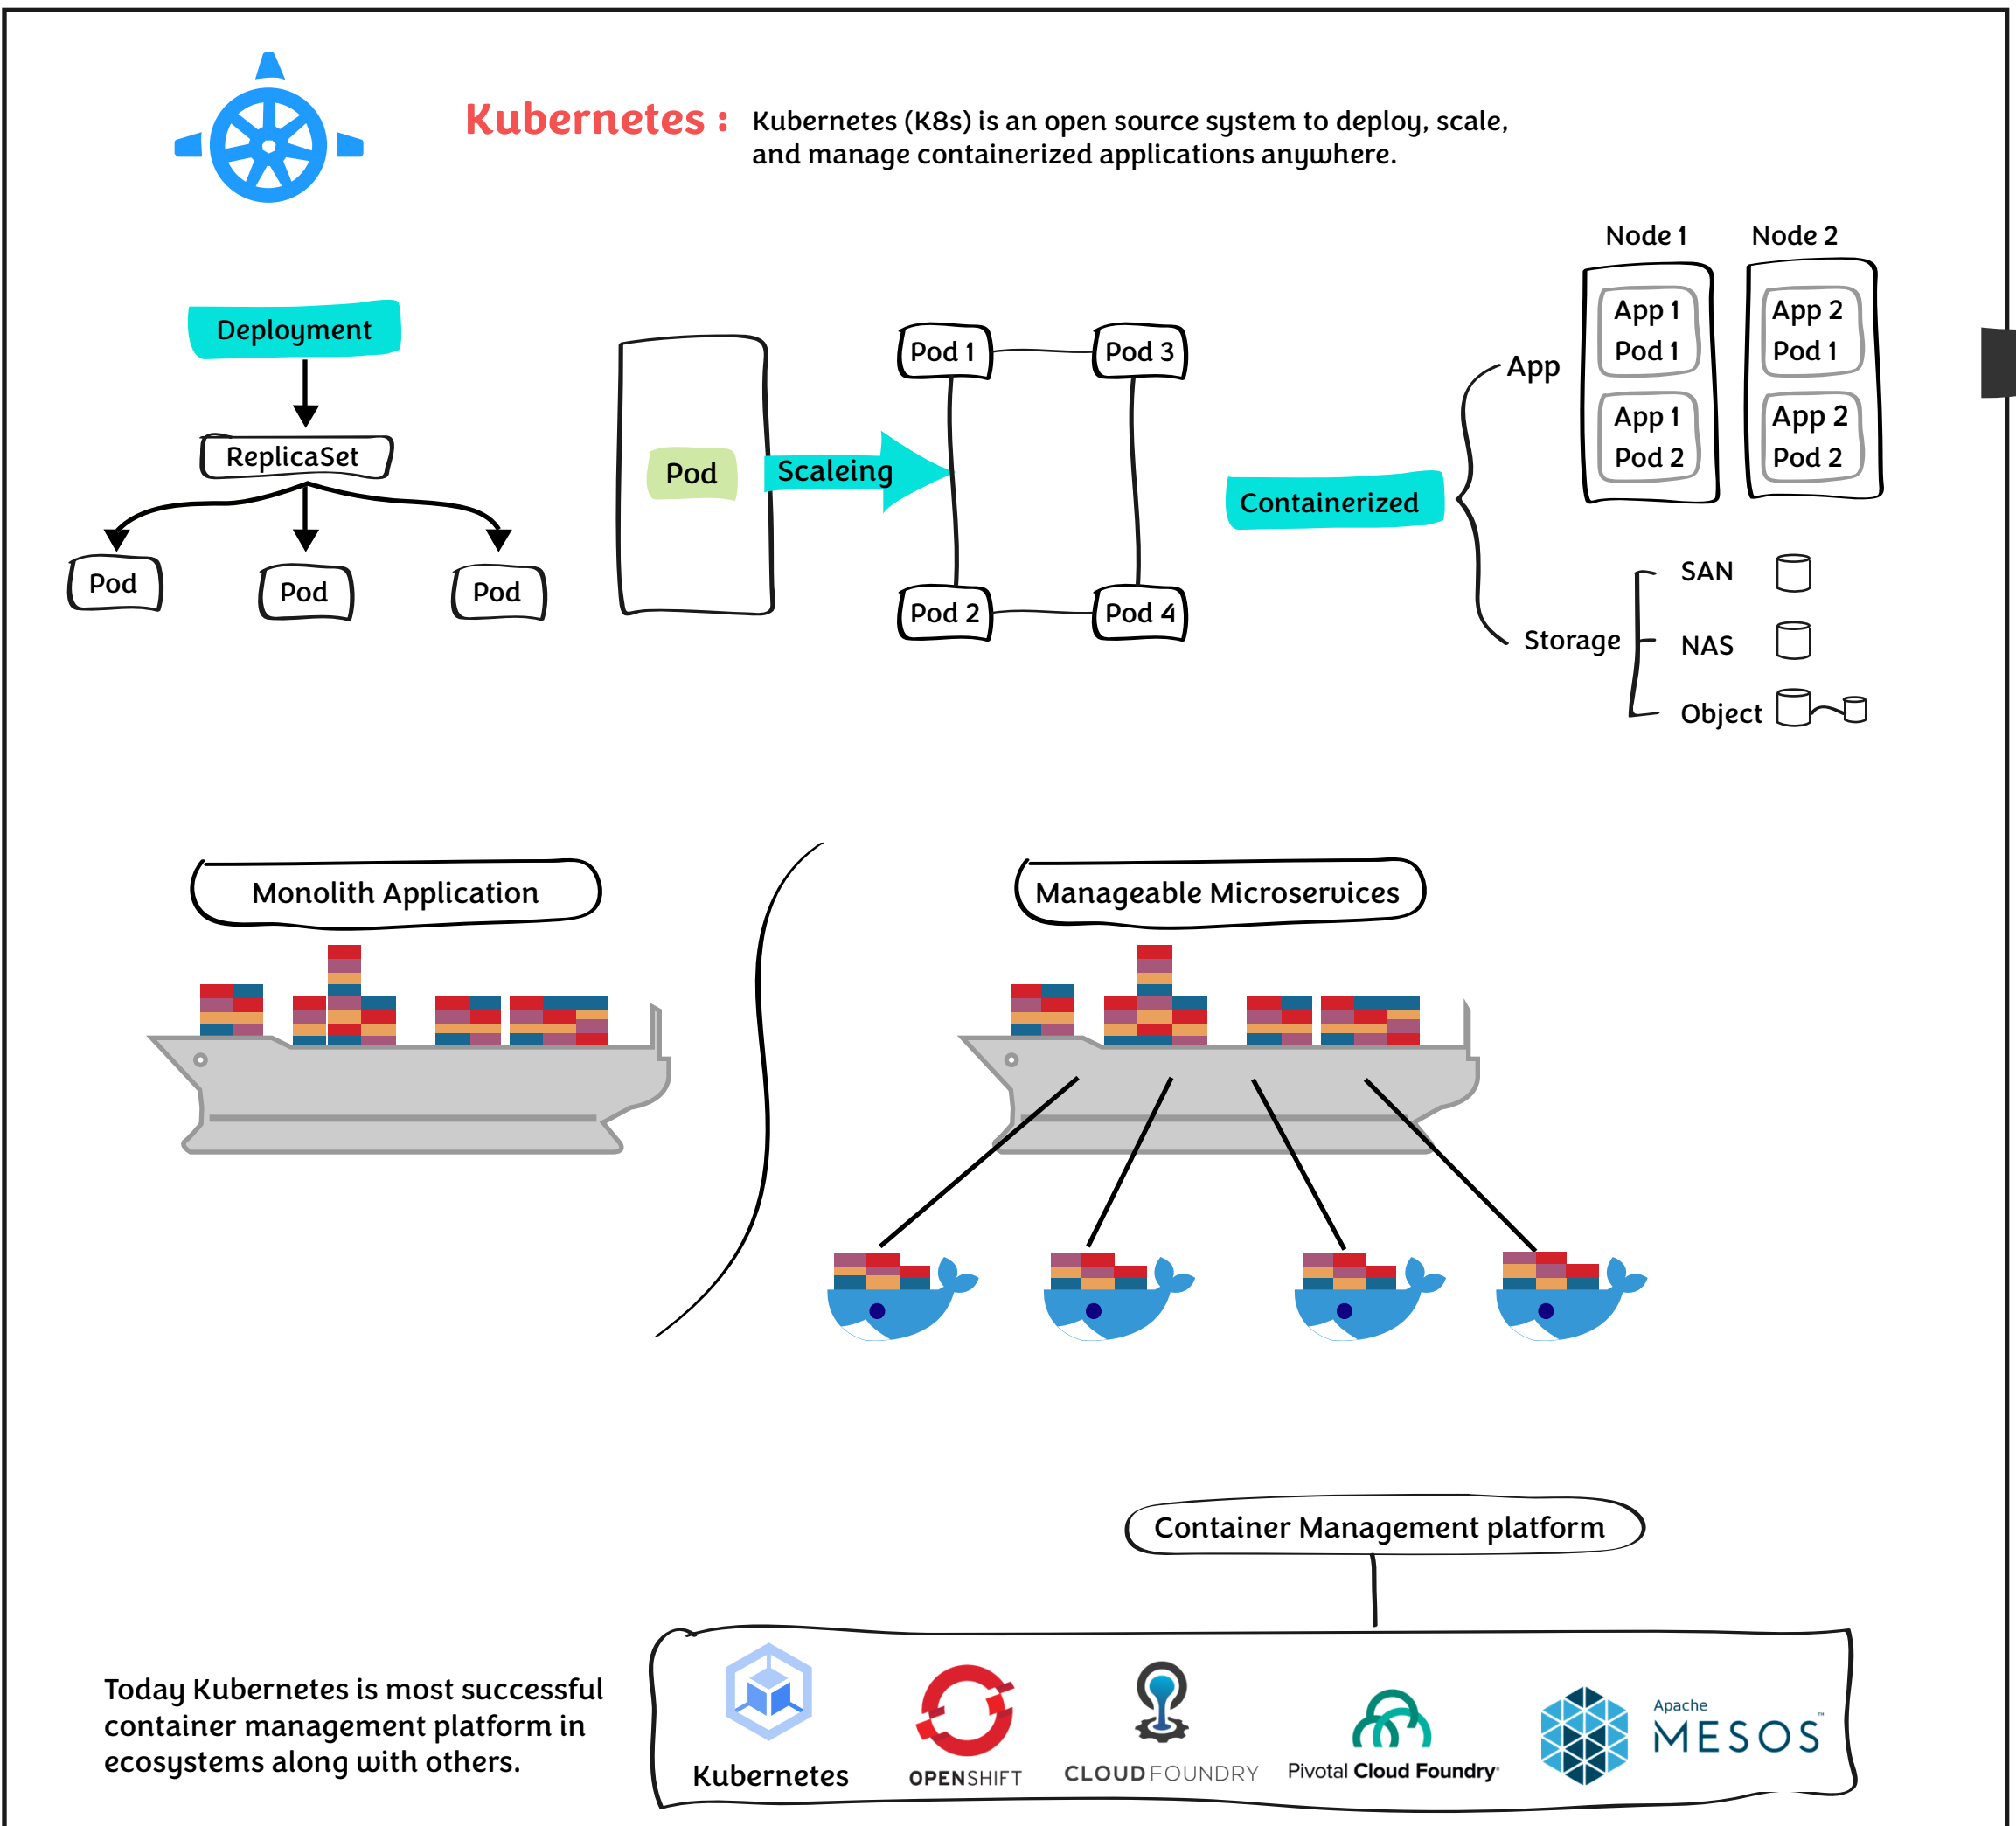

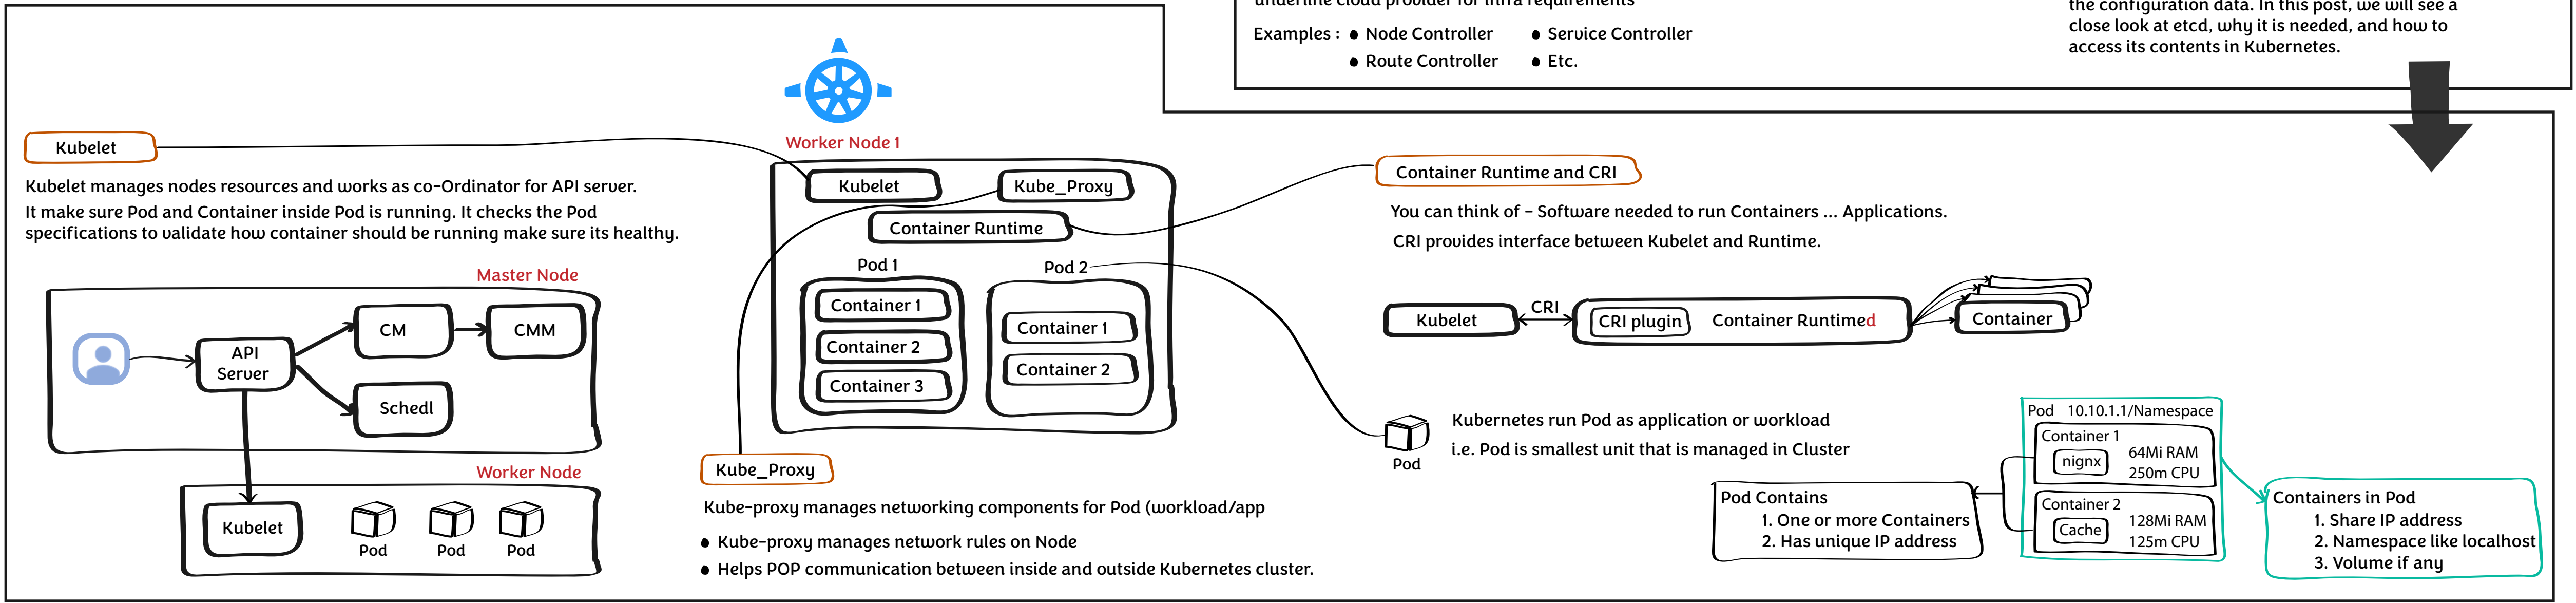
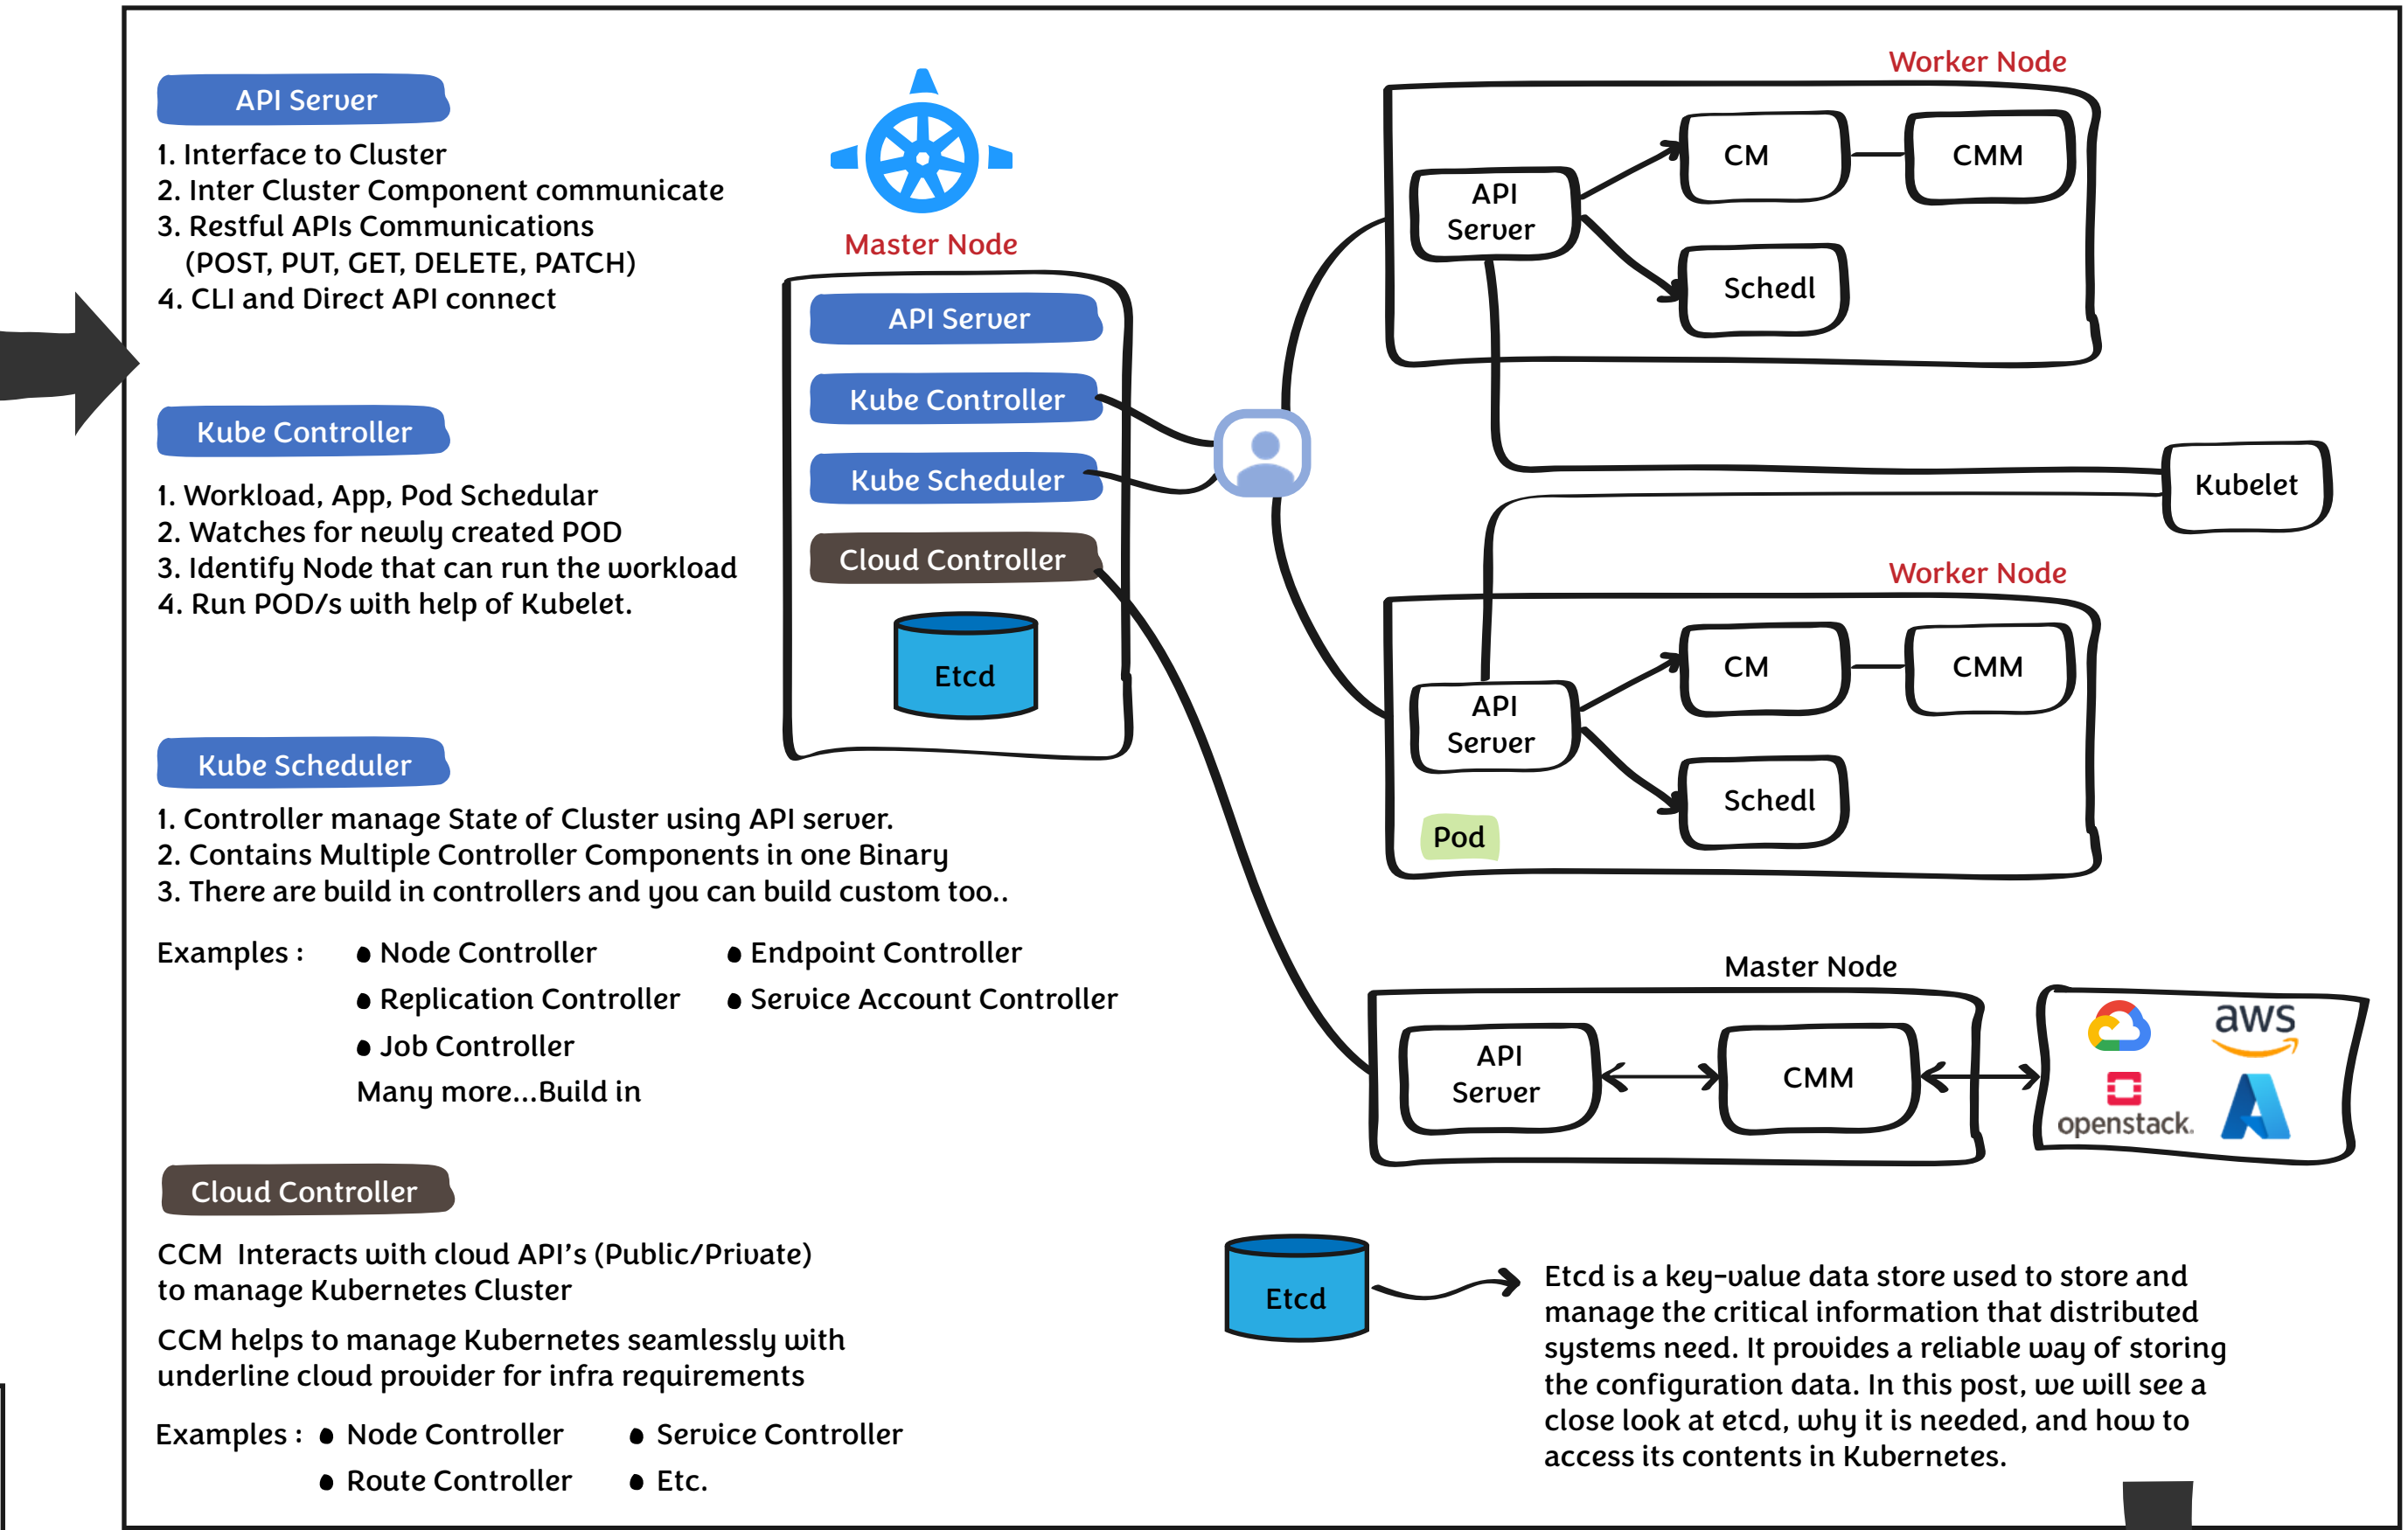
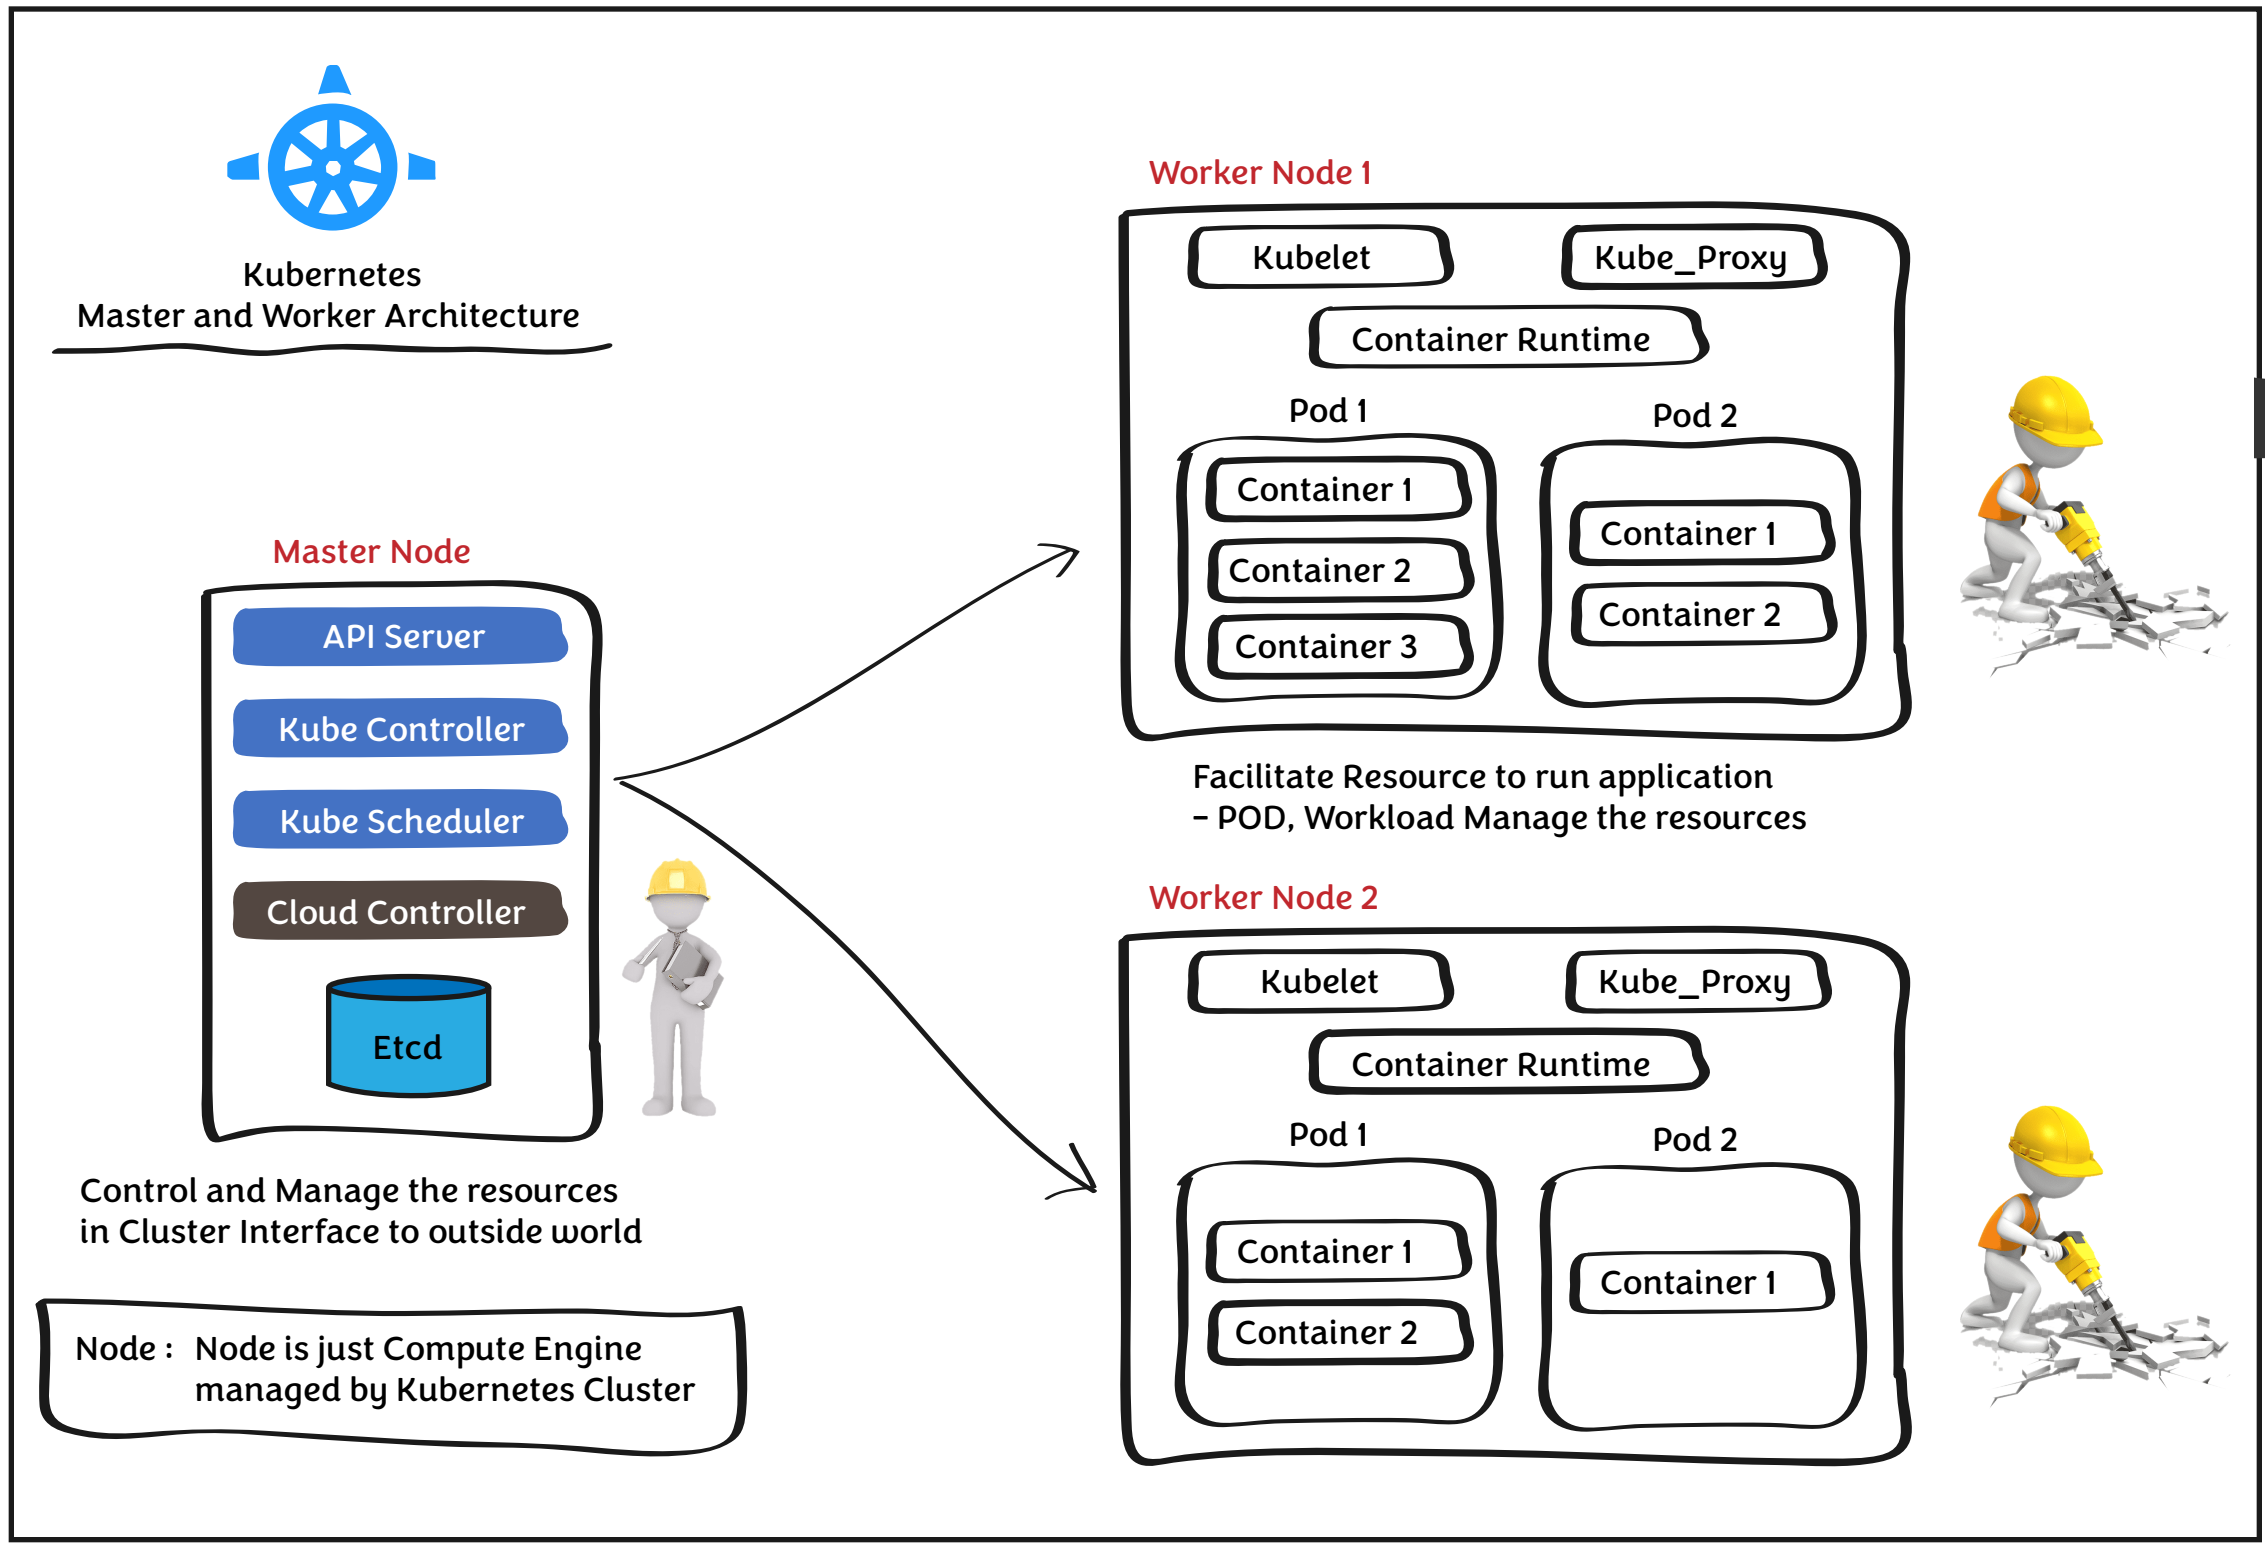



Kubernetes Engine

Udemy GCP Gurus









Pods, Deployments & Replicaset


Pod and Container Lifecycles



Pod Lifecycles




Container Lifecycles



Deployment

Deployment in Kubernetes describe “Desired state of object”.
Deployment controller try to achieve desired state from current state with controlled rate
Deployment is “Way to manage your application”




Application Management

1. Create Deployment to create rollout Replicaset.
2. Declare new state of Object.
3. Rollout new release or rollback to earlier version.
4. Pause rollout.
5. Scale up and Scale down.
6. Cleanup Replicaset.

Replicaset

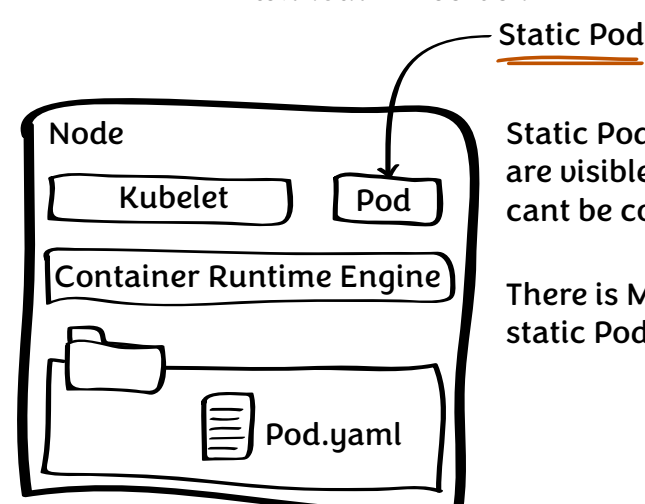
You also “Link” and “unlink” Pod based on need.



If this happens – Replicaset always try to maintain desired replicas of Pod.

Static Pod

Static Pods are managed by Kubelet Daemon without API server.



Static Pod

Static Pod running on the Node are visible to the API server but cant be controlled by API server



There is Mirror Pod for each static Pod

Init Containers

Containers that run before application containers.

You can setup init containers for utilities or initial setup scripts if can not be included in app image.

Resource Requests are handled differently.



Objects & Namespace

There are system NS & Default NS


- ☞ Applicable only to namespaced objects
e.g. Deployments, services
- ☞ Not applicable for cluster scoped resources
e.g. Storage Class, Node, PV

Resource Quota :

Kubernetes allows us to configure resource Quota.

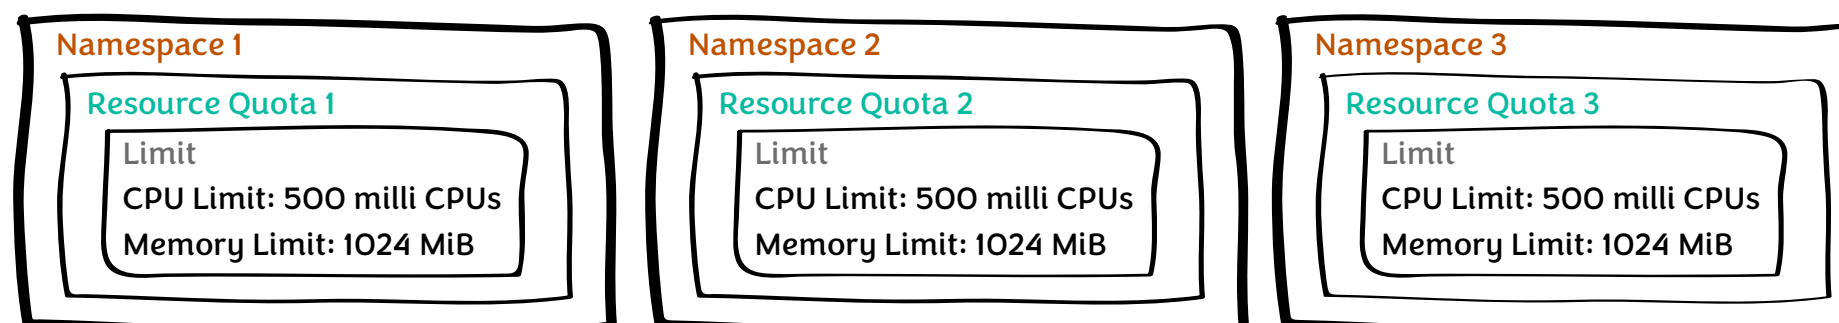
Namespace

Namespaces are used in Kubernetes to isolate objects in single cluster
Name of the resources should be unique in NS but not across clusters



Limit :

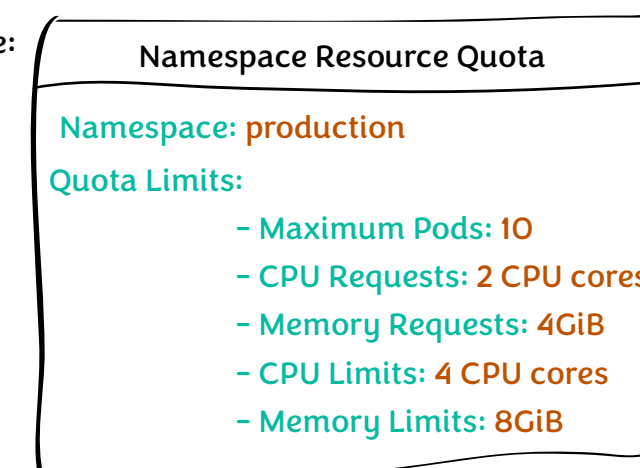
Limit the "testing" namespace to using 1 core and 1GiB RAM.
Let the "production" namespace use any amount.



Example:

Request : Minimum committed by cluster
Limit : Max Allowed by namespace resources.

Resources can be : Memory Size , Number or % of CPUs,
Number of objects etc




Objects

Pod is not the only objects in Kubernetes ...
there are many that helps Kubernetes cluster run itself as well as workload.
Kubernetes objects are persistent entities in the Kubernetes system !!!

Kubernetes Objects describes

1. Applications are running (and on which nodes)
2. The resources available to those applications
3. The policies around how those applications behave,
4. such as restart policies, upgrades, and fault-tolerance

Objects Types



Objects Managements Commands

Imperative Commands


```
kubectl create deployment nginx  
-image nginx
```

Imperative Object Config

```
kubectl create  
-f nginx.yaml
```

Declarative Object Config

```
kubectl apply  
-f configs.yaml
```



Labels & Selectors


Label




Labels are key/value pairs that are attached to objects (Deployment, POD, Namespace etc)

Selectors

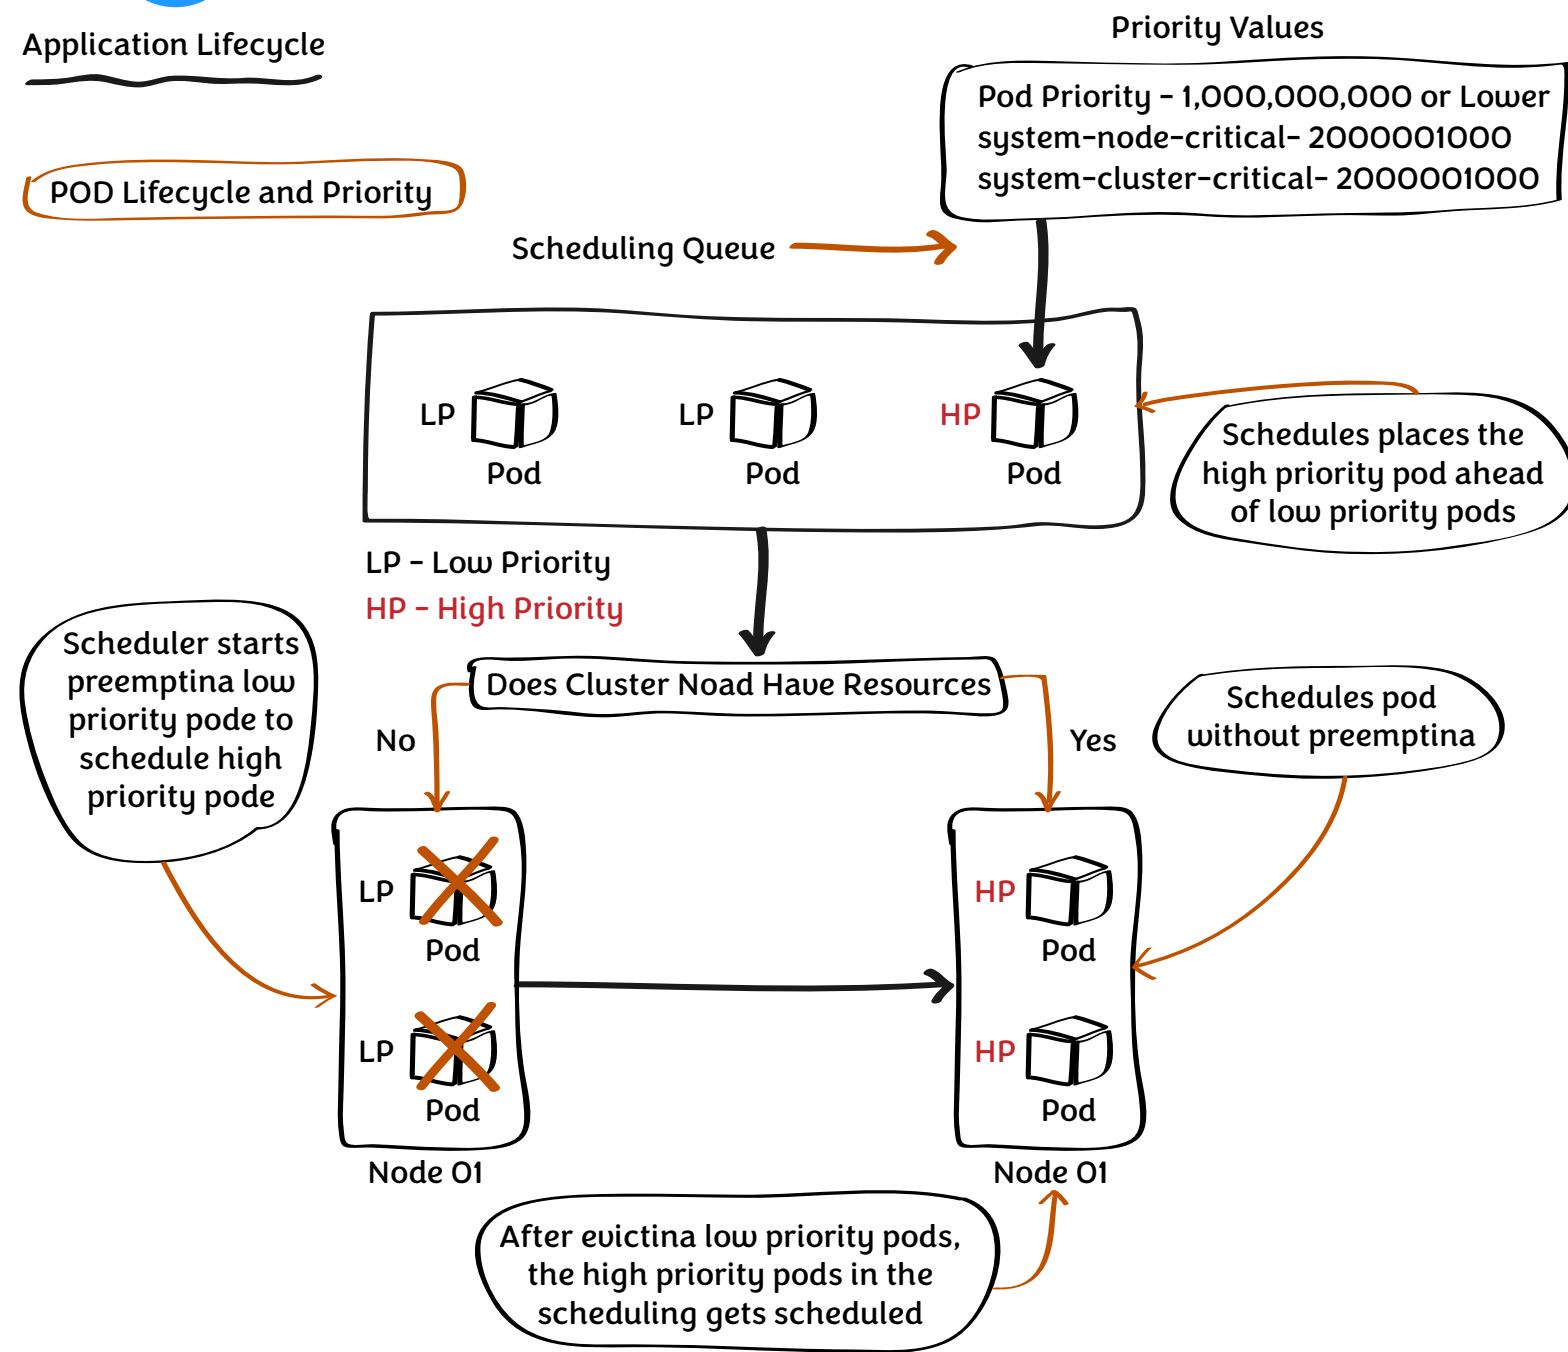
- ☞ Using labels Clients or Users can identify a set of objects
- ☞ Core Grouping of objects in Kubernetes
- ☞ You can have Equality based Label Selectors as well as Set Based






Application Lifecycle

POD Lifecycle and Priority



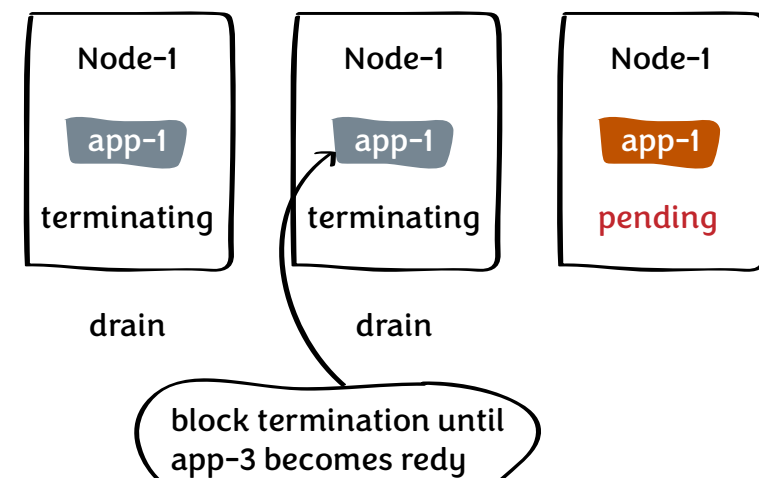
Horizontal Pod Autoscaler

HorizontalPodAutoscaler automatically updates a workload resource to match the demand.
HPA is implemented as Kubernetes APIs and Controller resource



Pod Disruption Budget

Pod Eviction and termination of Pod happens voluntary or involuntary

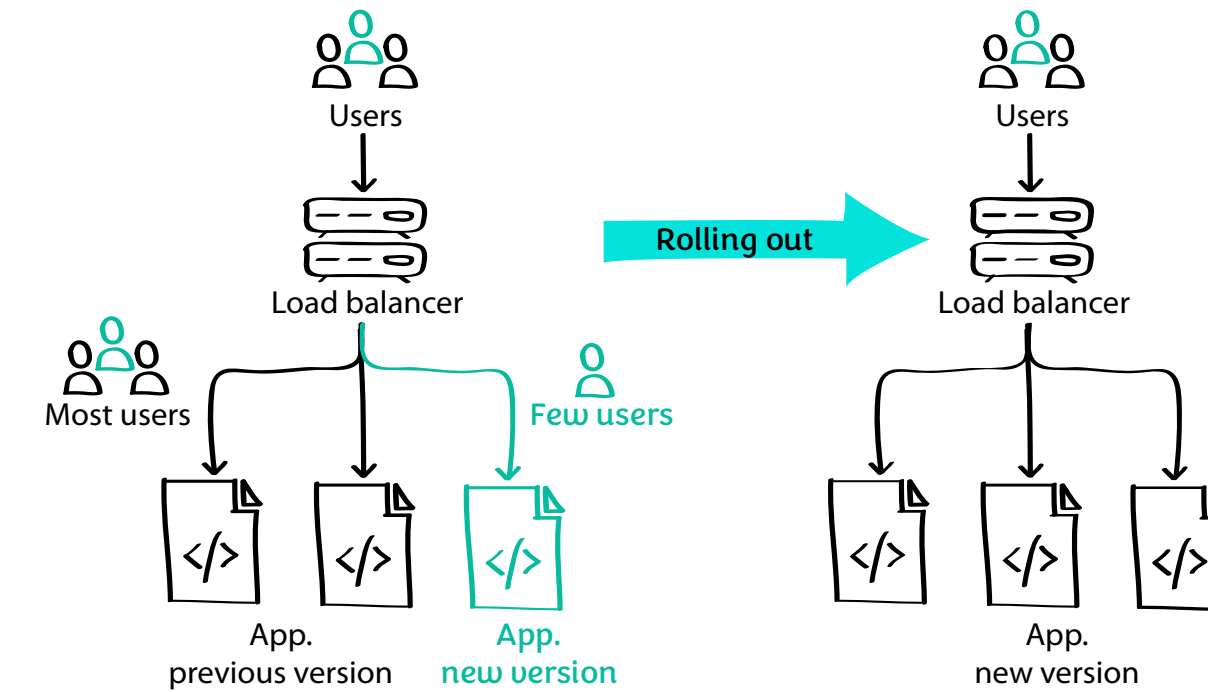


Rolling updates and constraints


Rolling updates allow Deployments' update to take place with zero downtime by incrementally updating Pods instances with new ones.

Rollback policy:


1. Recreate : Recreate will kill existing Pods and create new pods.
2. Rolling Updates : Rolling updates will update Pod one by one - Means update one at a time.



Rolling Updates:




Blue Green Deployment




Canary Deployment:

A canary deployment consists of routing a subset of users to a new functionality. Based on confidence or applications telemetry, new versions will gradually take complete load.




A/B Testing - Canary Application



We creates two deployments with same labels. Replicasets are used to controls number of Pods of each version running.

A/B Testing really checking business decisions on statistics rather than deployment strategy. canary deployment consists of routing a subset of users to a new functionality



Scaling Deployment

There are multiple ways you can scale Pods / Deployment / Application if cluster has resources.

High level - There are two methods -> Manual and Automatic

Manual Pod Scaling

kubecti scale deployment/nginx-deployment--replicas=10

Automatic (HPA)

kubecti autoscale deployment/nginx-deployment--min=10 --max=15 --cpu-percent=80

Cluster Autoscaler

GCP Kubernetes Engine use cluster autoscaler to add or remove nodes based on workload demand.



How aggressively you want Kubernetes Engine to scale up and scale down and rearrange pods?


Autoscaling Profiles :

Balanced: The default profile. Optimize-utilization: Prioritize optimizing utilization over keeping spare resources in the cluster. When selected, the cluster autoscaler scales down the cluster more aggressively: it can remove more nodes, and remove nodes faster.

Node Pool

Node Pool in Kubernetes Engine on GCP configures group of nodes (as part of managed instance group) to managed it together. It allows us to configure autoscaling properties e.g. min and max size of Node pool.





Clients want to connect directly to one or two pod

HeadLess

Nodeport exposes the service on each node IP at static port.

NodePort

Cluster IP is default Access only available within cluster

ClusterIP

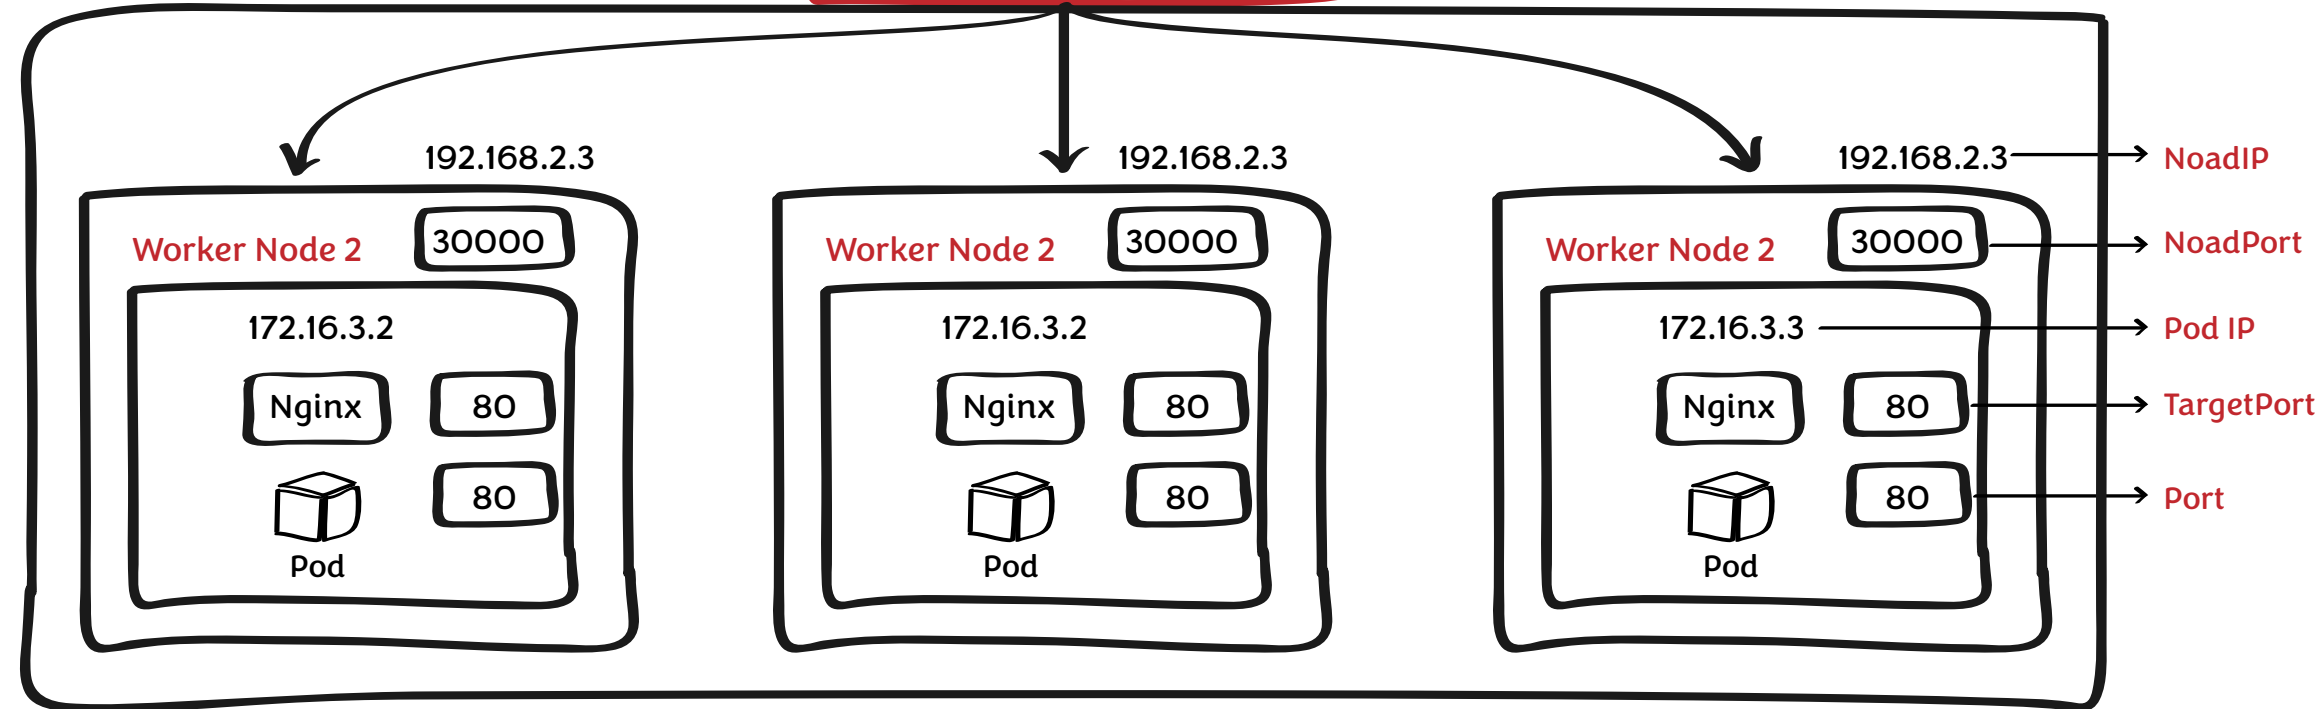
Service is accessible via LoadBalancer and outside cluster.

LoadBalancer

This type of Service maps a Service to a DNS name by using the contents of the externalName field by returning a CNAME record with its value

ExternalName

Service



Ingress



Loadbalancer is associated to IP address and if you need to expose multiple services – you need to have multiple load balancer and hence increase the cost.

We can not use Nodeport in productions.

Ingress exposes HTTP and HTTPS routes from outside the cluster to services within the cluster. Traffic routing is controlled by rules defined on the Ingress resource.

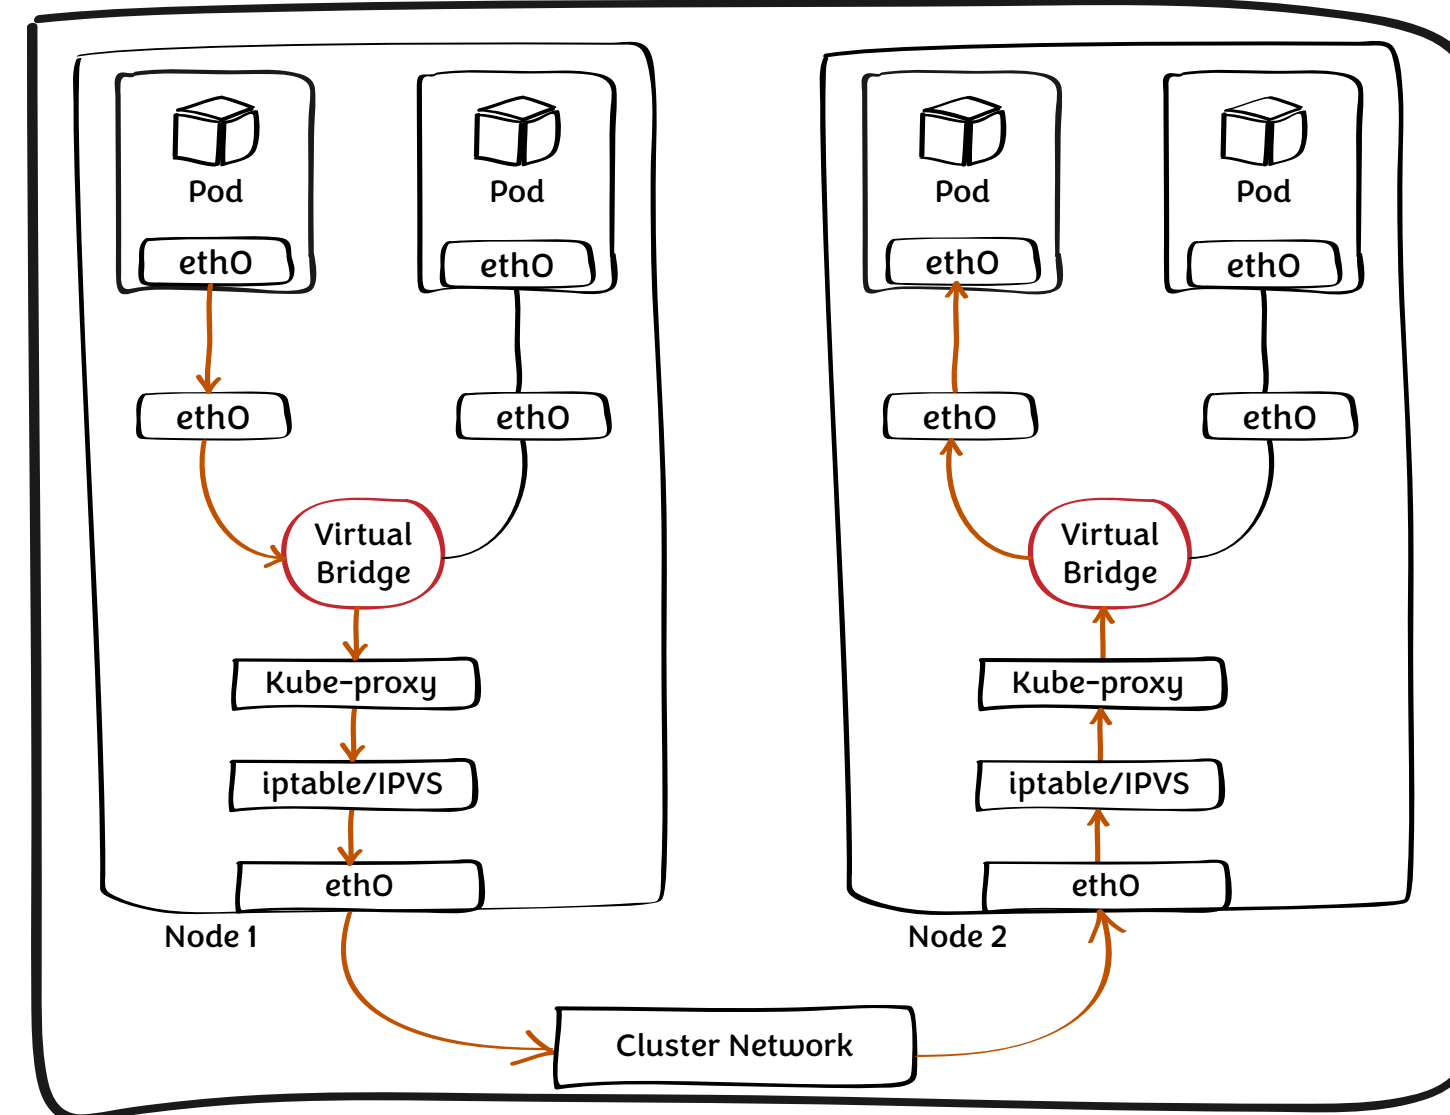

Ingress Controller


- Ingress Resource does not do any request forwarding...
- Ingress Controller does all the heavy lifting...
- Watches API Server for change in ingress resources
- Deployed as Pod and daemonset or normal deployment is used to deploy ingress object
- Not Started automatically with cluster.
- GCP uses nginx controller in GKE



Networking

- Highly-coupled container-to-container
- Pod-to-Pod
- Pod-to-Service - Already Covered
- External-to-internal - Already Covered






Scheduler


The Kubernetes scheduler is a control plane process which assigns Pods to Nodes .

Cluster



Kube Scheduler


1. Identifi Unscheduled pods



2. Node Selection is based on 2 step operation


1. Filtering :

Filtering out Node-1
cpu < 10)




```
Pod.yamll
apiVersion: v1
kind: Pod
metadata:
  name: <pod_name>
  :
spec :
  :
  containers:
    - name: <container_name>
      image: <image>
      resources:
        requests:
          cpu: 10
```

2. Scoring :



3. Binding :

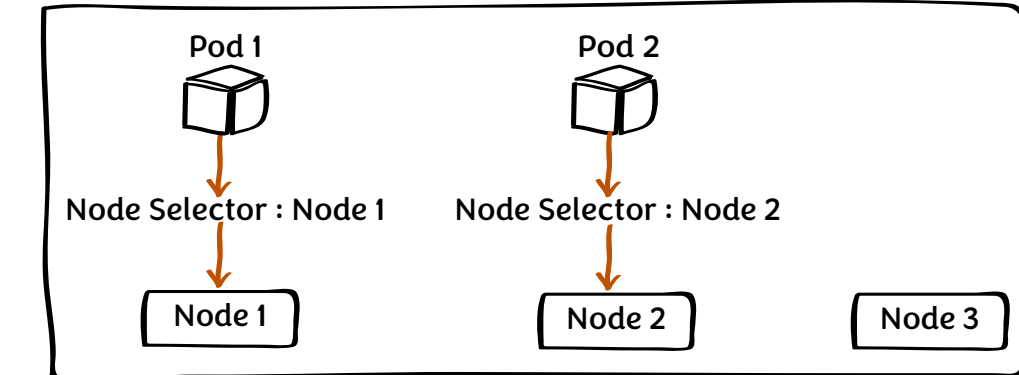


update **nodename**
with the selected node


```
pod.yamll
apiVersion: v1
kind: Pod
metadata:
  name: <pod_name>
  :
containers:
  - name: <container_name>
    image: <image>
    resources:
      requests:
        cpu: 10
        nodename: node 2
```

Kube_scheduler does scoring and filtering provides other criteria are met ?


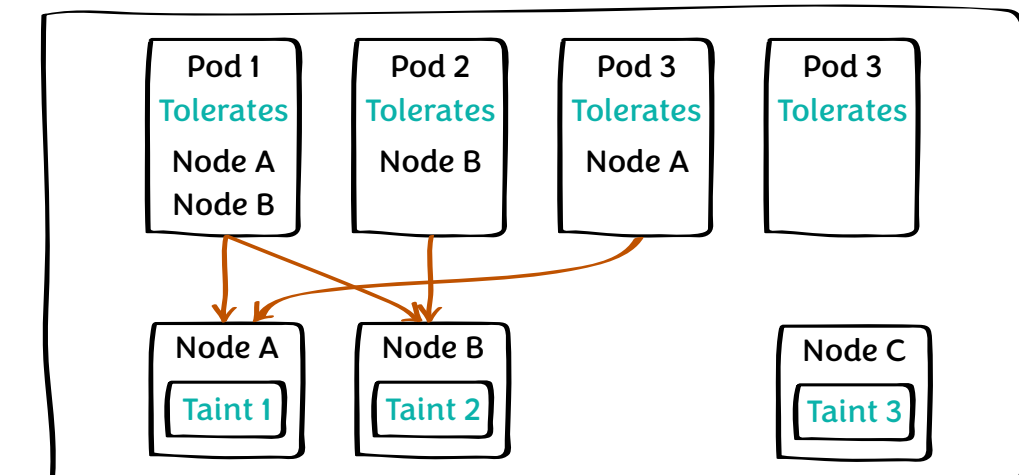
1. Node Selector :



2. Affinity:



2. Taint/Toleration:




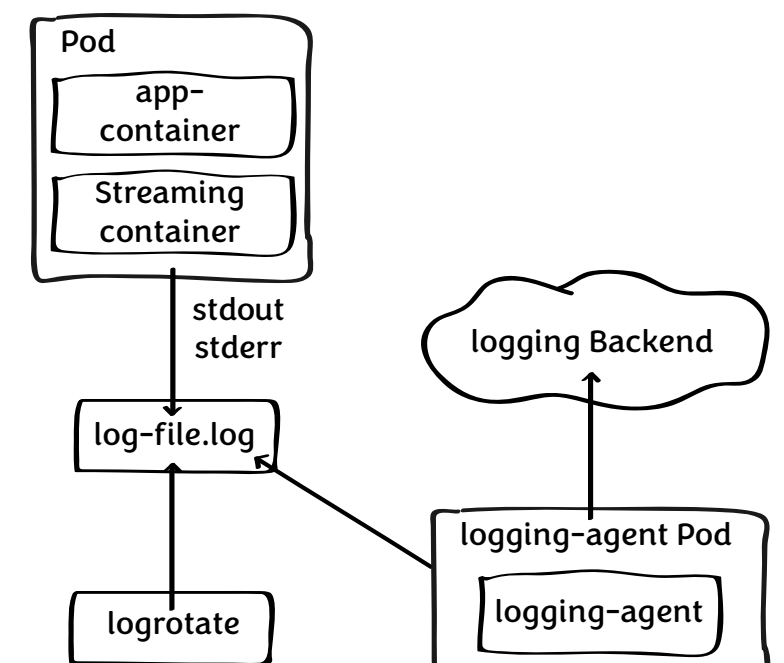
Logging and Monitoring

1. POD/Container Logging

Logging and Monitoring is essential of activity to keep application running in production
Applications container runs in POD as well as System components also runs in POD(container)
Kubelet and Container runtime writes to journalD



2. Node Logging

All stdout and stderr is handled and redirected by container engine
Kublete on node handles request for all the logs on node



GCP - Kubernetes Engine


Nodes and Pods are ephemeral and we may loose logs if node died by autoscaler or any other unplanned activity
Enable Logging and Monitoring of Kubernetes Engine Cluster.



StatefulSets, Jobs and DaemonSets

1. StatefulSets

Not all the workload that you run is stateless and have flexibility to scale up and scale down and Schedule as and when we need.



Not all the workload that you run is stateless and have flexibility to scale up and scale down and Schedule as and when we need.

2. DaemonSet

What we can do - if we need to run one application on all nodes.. ?

- A DaemonSet ensures that all (or some) Nodes run a copy of a Pod.
- As nodes are added to the cluster, Pods are added to them.
- As nodes are removed from the cluster, those Pods are garbage collected.
- Deleting a DaemonSet will clean up the Pods it created.
- While Creating - Define highest priority so that Pods are not evicted.

Uses of a DaemonSet are:

- Running a cluster storage daemon on every node
- Running a logs collection daemon on every node
- Running a node monitoring daemon on every node

