



# Why Kubernetes?

A Deep Dive in Options, Benefits and Use Cases



**Bettina Bassermann**  
Solution Sales Specialist DevOps  
[Bettina.Bassermann@suse.com](mailto:Bettina.Bassermann@suse.com)

**Rania Mohamed**  
Solution Architect, Services Consulting  
[Rania.Mohamed@suse.com](mailto:Rania.Mohamed@suse.com)

# Who is SUSE?

- Founded in 1992
- Largest independent open source vendor as of March 2019
- Technology company
- Our Mission is to help customers to master the digital transformation through Open Source technology
- Innovating with Partners and communities
- Enterprise-Grade Support



# Series about modern Application Development

- Software Development, Microservices & Container Management, a SUSE webinar series on modern Application Development
- Please find all SUSE Webinars here

<https://www.suse.com/de-de/events/webinars>

<https://www.suse.com/c/author/rasadus/>

**Microservices –  
Is it the Holy  
Grain? A  
Perspective of a  
Developer**

**Container and  
Cloud Native  
Technologies –  
Why do we need  
them and what is  
so great about it?**

**Why Kubernetes?  
A Deep Dive in  
Options, Benefits  
and Usecases**

**About making  
Choices –  
CaaSPv4 as  
SUSE's  
empowering of  
Kubernetes**

....stay tuned for the 2020 sessions with the Chamelion



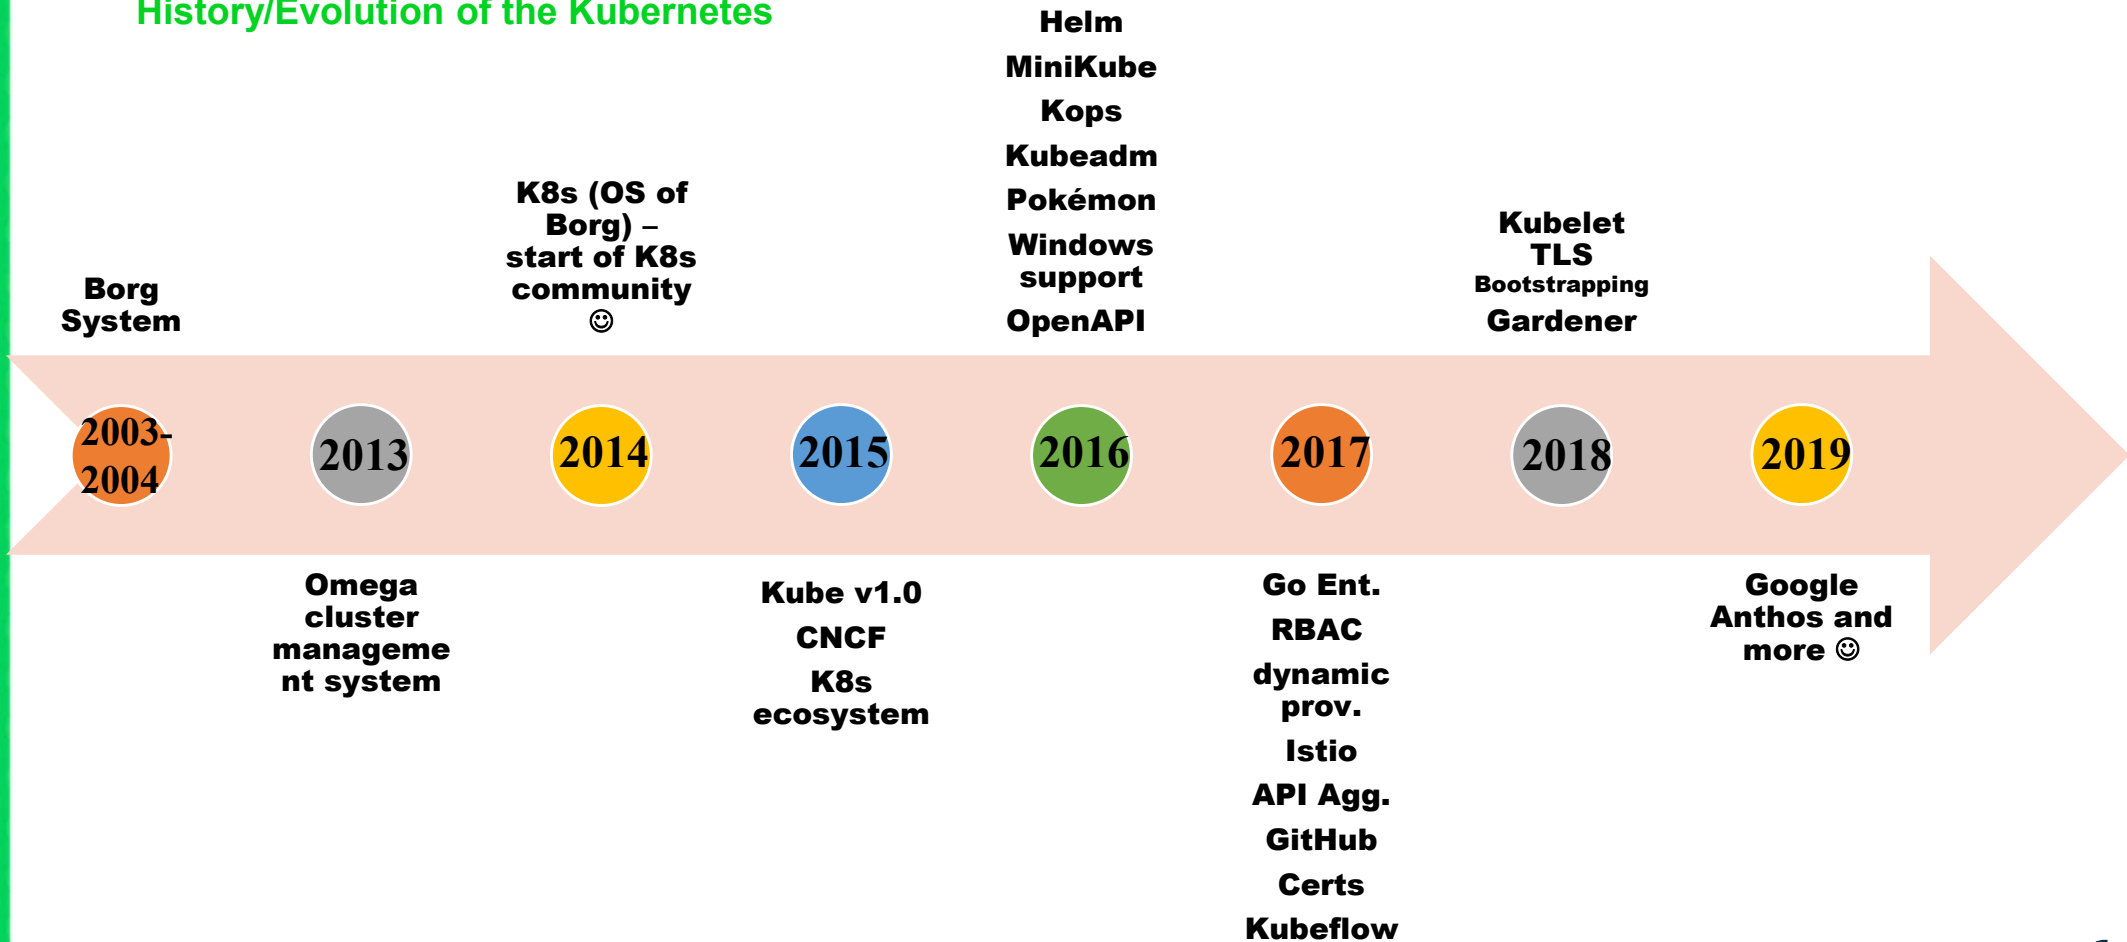


# Agenda

- **Kubernetes evolution**
- **Kubernetes Architecture**
- **Kubernetes Components**
- **Deep Dive in Kubernetes Network component**
- **Deep Dive in Kubernetes container engine component**
- **Deep Dive in monitoring component**
- **Deep Dive in Kubernetes service Mesh Component**
- **Kubernetes Usecases**

# Basics about Kubernetes

## History/Evolution of the Kubernetes



# Kubernetes Architecture

# Kubernetes Architecture – Main Soln Design Principles

- **API Centric**
- **Separation of concerns** → divide and conquer 😊
- **Pluggability**
- Flexibility
- Well-defined **State management** → following MSA 😊
- **Extensibility**
- Scalability
- **Automation**
- **Simplicity**
- Standardization
- **Design for failure**
- **PaaS**
- Repeatability → **reconstructable** by **observation**.
- **Self-healing**
- Implements **Event** processing and even complex event processing
- **Graceful degradation**
- Target **autonomous**
- **Manage** dependencies
- **Transparency**
- Design API by **SLO** rather than implementation
- **High** Availability
- **Multi-tenancy**
- **Decentralizing** more distributed

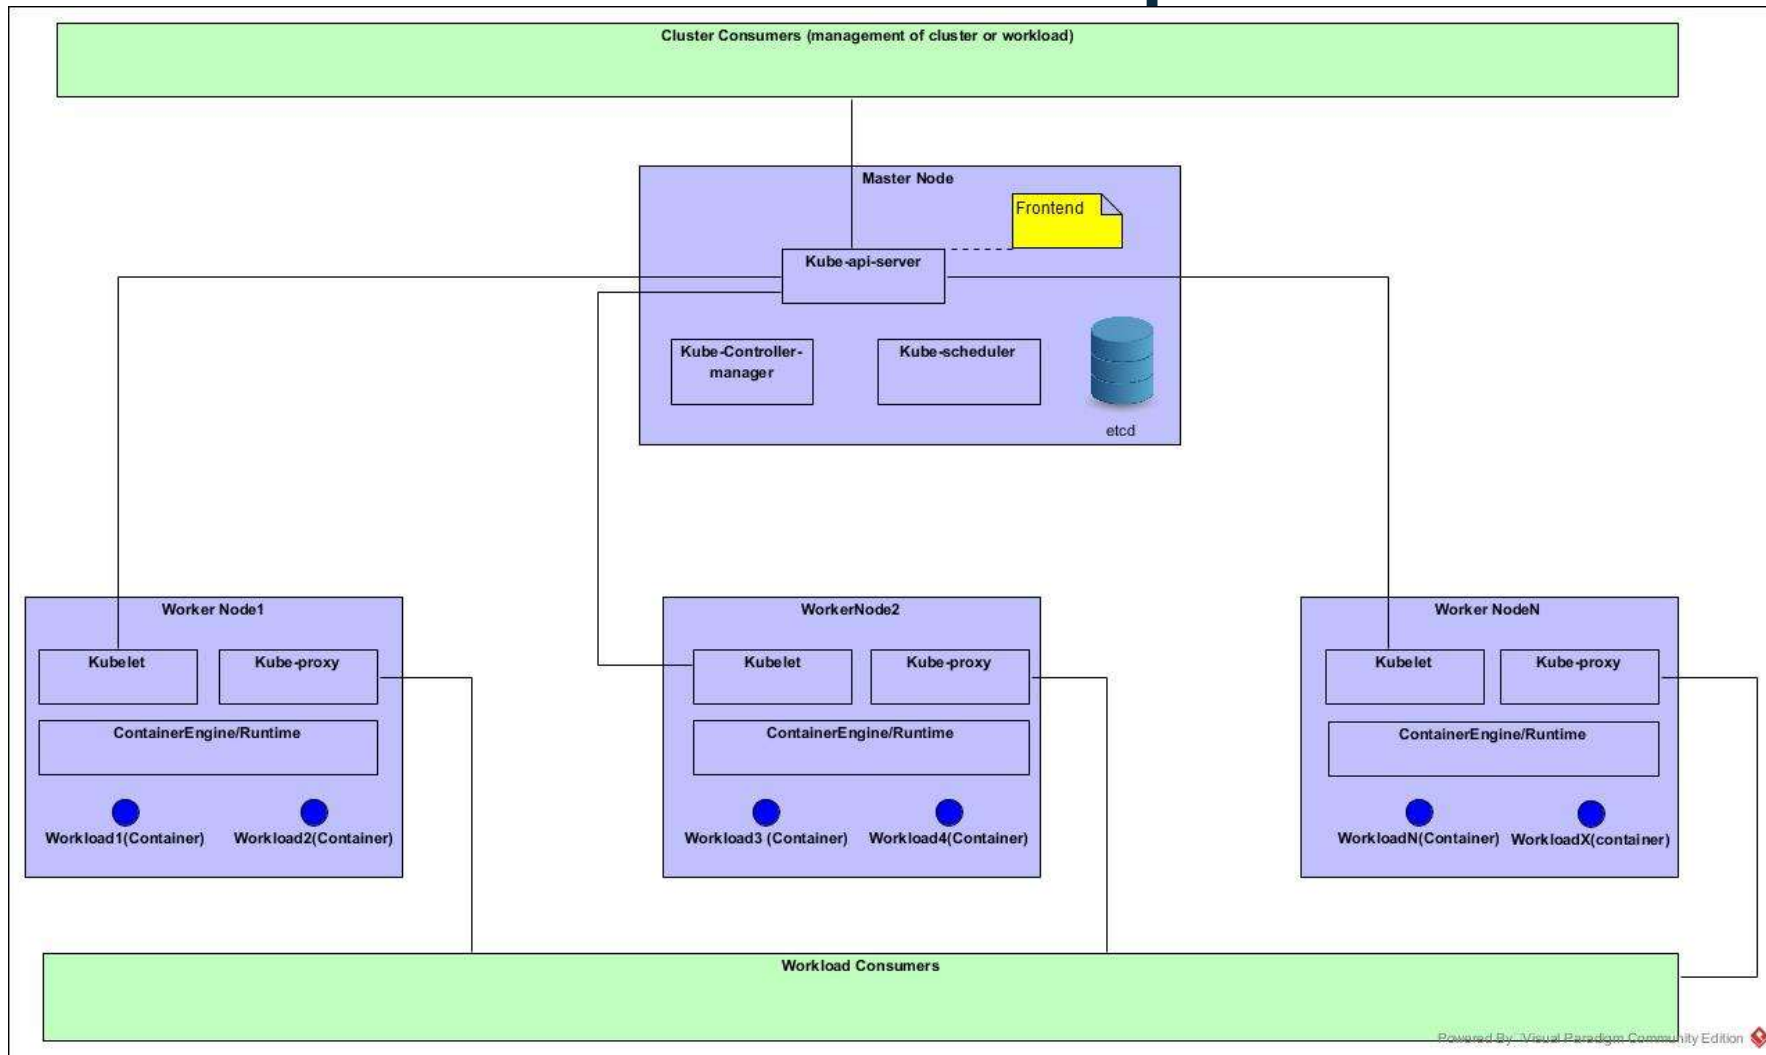


# Kubernetes Architecture – Main Target Requirements

- **Orchestrate & Manage** Containers
- Handle **Variable** load **efficiently**
- Enable **integration** with external world
- Enable **apps/workloads integration**
- Enable **business continuity**
- Support running **distributed** apps/workload and **MSA**
- Enable **hybrid/multi-cloud**
- Enable **CNA** development
- Support/offer **DevOps**
- Support of **pushing** workload **updates** with **no disruption**, ex blue-green
- Enable **Modernization**
- Enable **12factor** development BP
- Enable integration with **marketplace**
- Workloads **Log aggregation** and **analytics**
- Enable **ease** of **monitoring** distributed workloads and gathering data
- Enable **service discovery** for the running workloads
- Enabling **load balancing** between workload instances
- Enable **authentication** and **authorization** for the workload

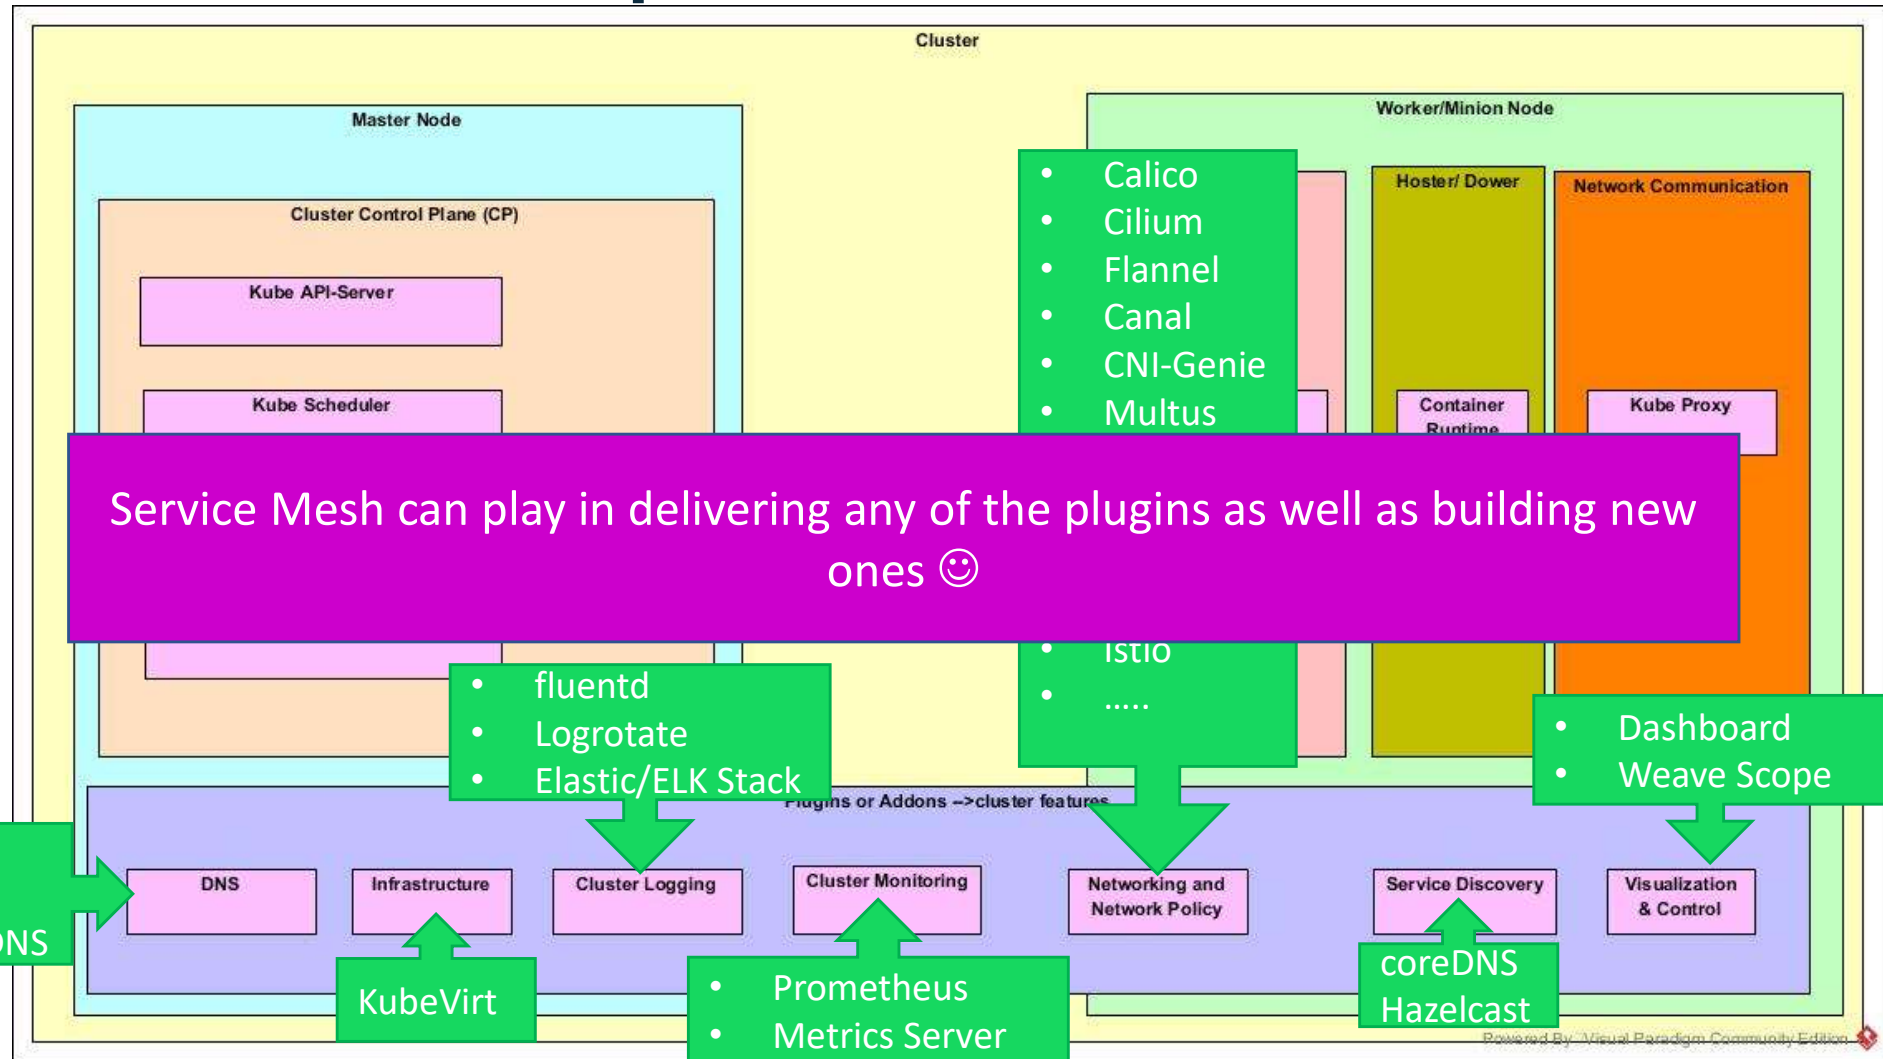
**Build and facilitate enterprise environment for containers (same as an application server does for a Java app) 😊**

# Kubernetes Architecture - Blueprint



# Kubernetes Components & Objects

# Kubernetes Components





# Kubernetes Objects – Main

- **Node** – a **machine** running in the cluster
- **Pod** – **smallest deployable** unit, may have 1+ containers
- **ReplicaSet** – defines **number** of **running** instances of a pod
- **Service** – defines **access** to one or more workload **API** (types: **ClusterIP**, **NodePort** & **LoadBalancer**) → sub-objects: svc, endpoint, iptables and IPVS
- **Volume** – **storage** of a container (supporting **stateful workloads**) → sub-objects: **PV** and **PVC** and **StorageClass**
- **Namespace** – resources **grouping**
- **Deployment** – declarative **packaging** of a set of pods and services
- **DaemonSet** – run a pod **instance** on **all nodes**
- **StatefulSet** – manages **Pods** that are based on **stateful** workloads.
- **Job** – creates one or more **Pods** / **service**.
- **Ingress** – manages **external access** to the services in a cluster

# Deep Diving in Kubernetes world 😊



# Network Component(s)

# Network Component –Main Target Requirements

- **Pods** can communicate **without** using **NAT**
- **Nodes** can communicate with any **pod** **without** using **NAT**
- **Pod** internal **IP** is always seen by the same no matter the angle 😊
- MSA **Gateway** (controlling routing)
- **Securing** communication between workloads using policies
- Ability to use **different network layer** (2, 7, 3 and 4) isolation
- Ability to **mix** and **match** different **network policies**
- **Standardize** network communication with container engine, kube api server and other nodes/pods → **CNI plugins**

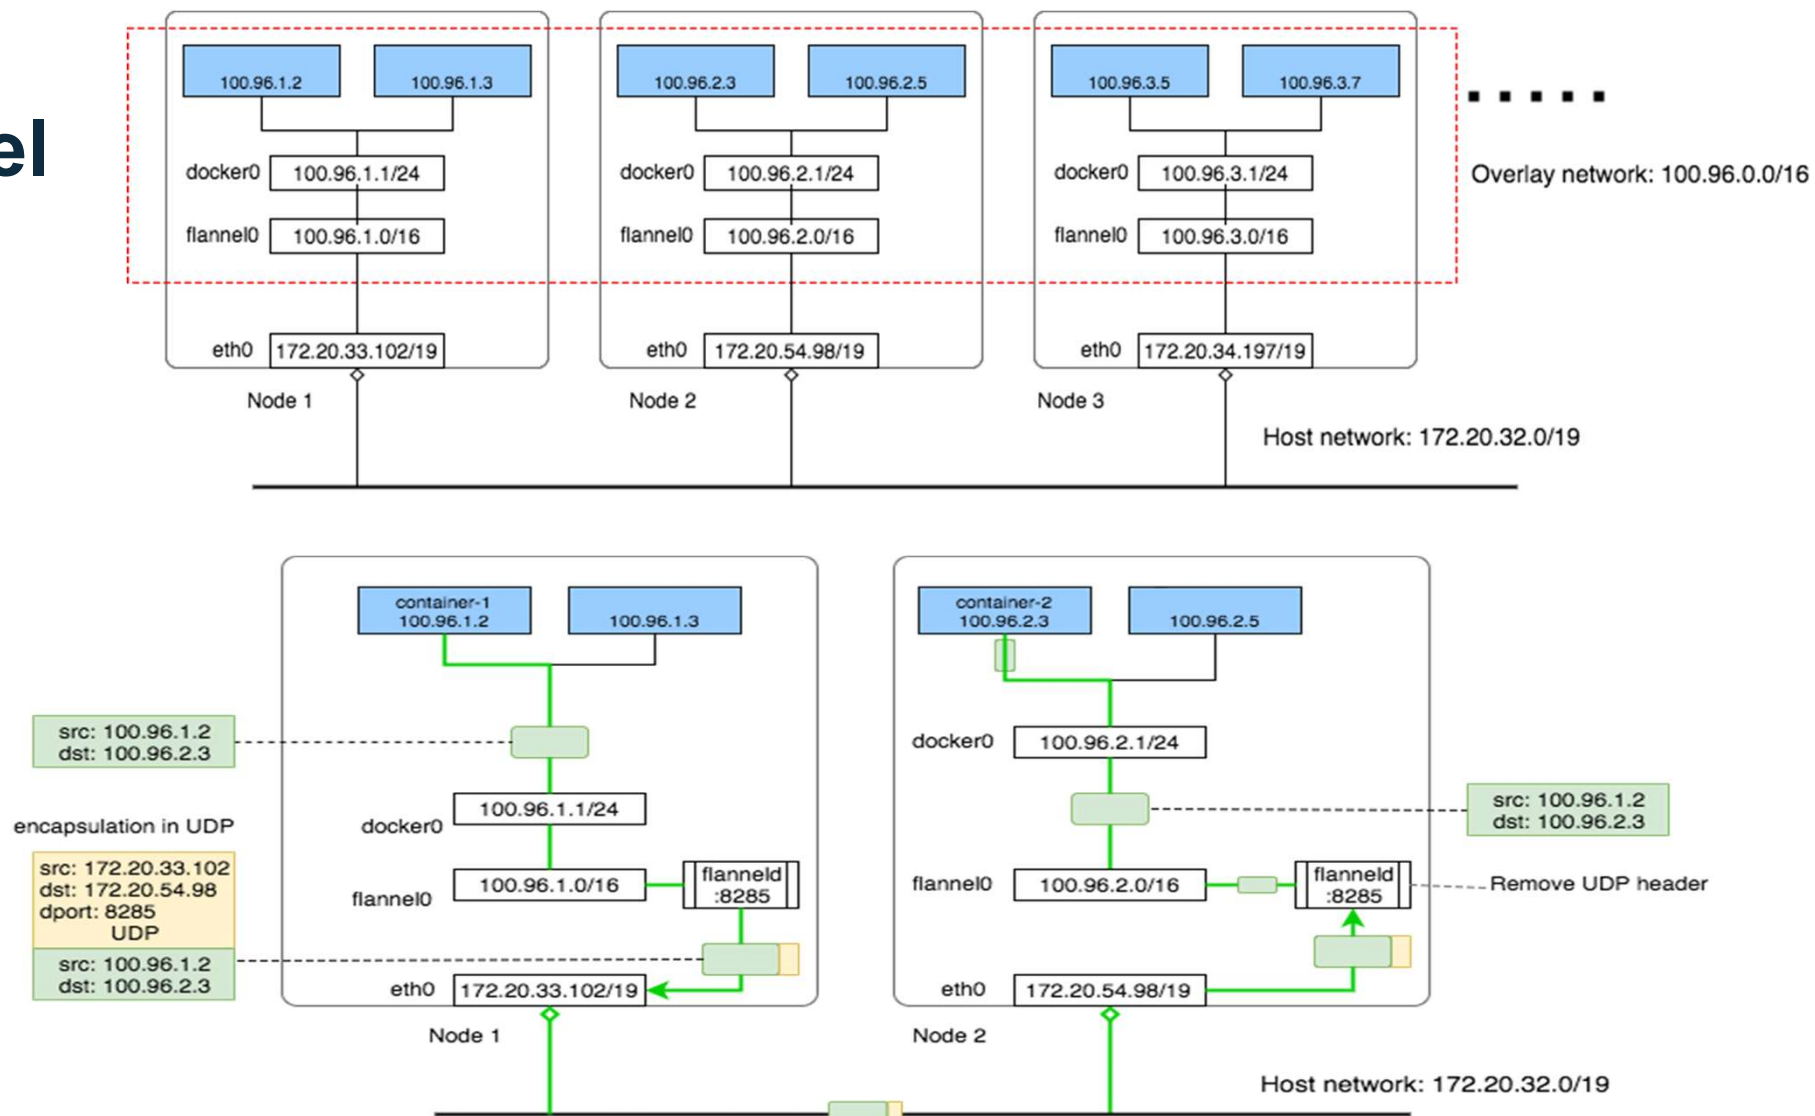
**K8s doesn't provide any default network implementation, It is a plugin😊**  
**Aim to support secured multi-tenancy 😊**



# Network Component – a comparison

Feature	Flannel	Calico	Multus	Cilium	Weave Net
<b>Brief features</b>	<ul style="list-style-type: none"> <li>- <b>Simple</b></li> <li>- <b>Flat</b>/overlay network</li> <li>- Each pod has <b>unique</b> ip, all pods run on the same network</li> <li>- Runs a <b>daemon</b> process flannel to update the ip routing table</li> <li>- Mapping of the subnet to host info stored in etcd</li> </ul>	<ul style="list-style-type: none"> <li>- Connects pods using the same IP networking principles as the internet</li> <li>- <b>Interoperable</b></li> <li>- <b>Flexible</b></li> <li>- Enable security enforcement (<b>self workload firewall</b>)</li> <li>- True <b>Cloud native</b> scalability</li> <li>- leverage best practice cloud-native design patterns</li> <li>- Enables hybrid soln using <b>BGP</b></li> </ul>	<ul style="list-style-type: none"> <li>- It enables attaching <b>multiple network</b> interfaces to the pods by creating homed pod, it is <b>magic</b> 😊</li> <li>- It is a <b>meta-plugin</b></li> <li>- It highly support <b>multitenancy</b></li> </ul>	<ul style="list-style-type: none"> <li>- Supports <b>MSA &amp; CNA</b></li> <li>- It works by <b>network policies</b></li> <li>- Supports <b>lightweight</b> protocols, such as HTTP, gRPC, Kafka</li> <li>- It is an <b>API-aware</b> network security <b>filtering</b>.</li> <li>- It uses Linux kernel technology called <b>BPF</b></li> <li>- <b>Simple</b></li> <li>- <b>Efficient</b></li> <li>- Enables building <b>gateway policies</b> which can be enforced network-layer and <b>application-layer</b> security <b>policies</b></li> <li>- <b>Scalability</b></li> <li>- <b>Multi-tenancy</b> 😊</li> <li>- L3 <b>Encryption</b> enforcement</li> </ul>	<ul style="list-style-type: none"> <li>- Creates a mesh <b>overlay</b> network <b>between</b> each of the nodes in the cluster</li> <li>- <b>Flexible</b> in the communication</li> <li>- <b>Simple</b></li> <li>- Enables service discovery using <b>micro DNS</b></li> <li>- <b>Encryption</b></li> <li>- Supports <b>multi-cast</b></li> <li>- Enables <b>portability</b></li> </ul>
<b>Net. Layering</b>	L3 network fabric	L3 & 7	N/A	L3 & 7	L3
<b>Stability</b>	Very high	Very high		High	High
<b>Service meshing</b>	Doesn't integrate and doesn't allow any network policy implementation	Integrates and enables defining rich network policy models	N/A	Integrates and enables defining rich network policy models	No
<b>Gateways?</b>	No	Yes	N/A	Yes	Yes
<b>Perf.</b>	Good	Very Good	N/A	Very Good	Very Good

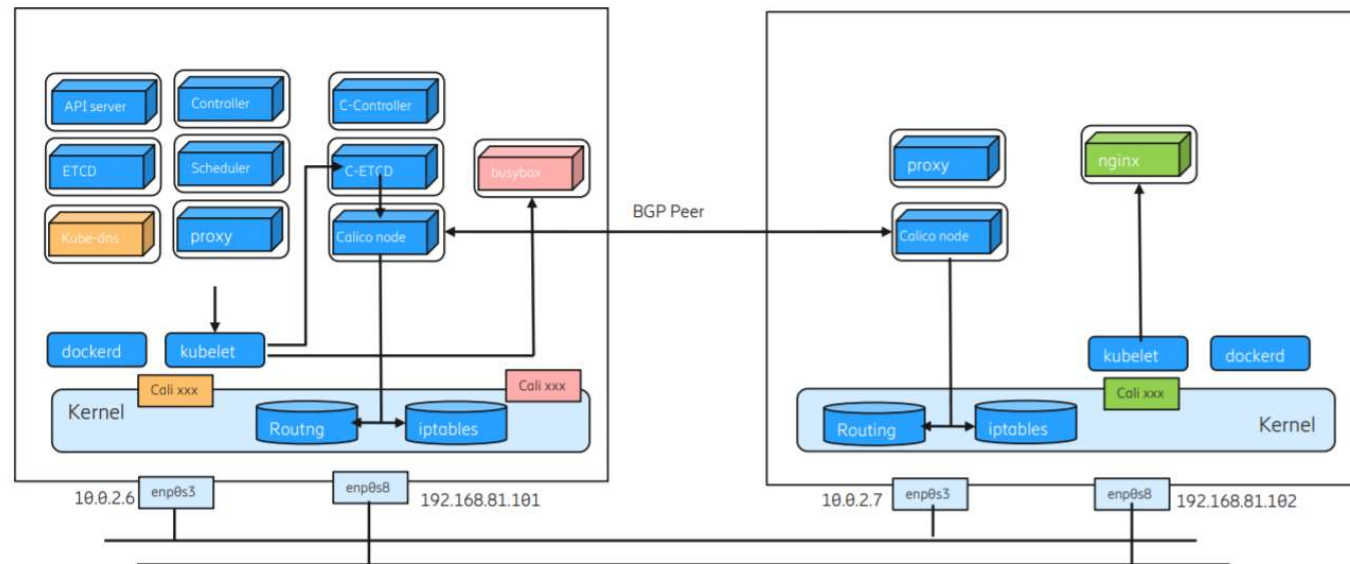
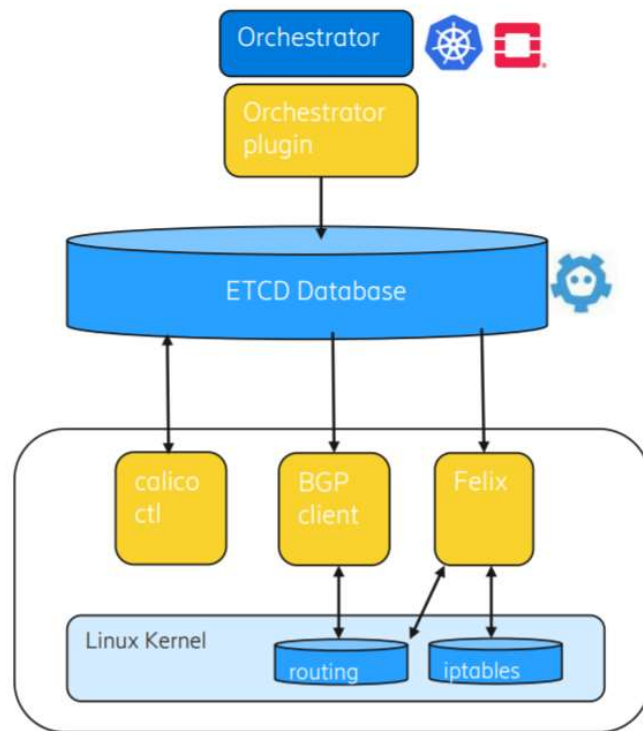
# Flannel



# Network Component – a comparison

Feature	Flannel	Calico	Multus	Cilium	Weave Net
<b>Brief features</b>	<ul style="list-style-type: none"> <li>- <b>Simple</b></li> <li>- <b>Flat</b>/overlay network</li> <li>- Each pod has <b>unique</b> ip, all pods run on the same network</li> <li>- Runs a <b>daemon</b> process flannel to update the ip routing table</li> <li>- Mapping of the subnet to host info stored in etcd</li> </ul>	<ul style="list-style-type: none"> <li>- Connects pods using the same IP networking principles as the internet</li> <li>- <b>Interoperable</b></li> <li>- <b>Flexible</b></li> <li>- Enable security enforcement (<b>self workload firewall</b>)</li> <li>- True <b>Cloud native</b> scalability</li> <li>- leverage best practice cloud-native design patterns</li> <li>- Enables hybrid soln using <b>BGP</b></li> </ul>	<ul style="list-style-type: none"> <li>- It enables attaching <b>multiple network</b> interfaces to the pods by creating homed pod, it is <b>magic</b> 😊</li> <li>- It is a <b>meta-plugin</b></li> <li>- It highly support <b>multitenancy</b></li> </ul>	<ul style="list-style-type: none"> <li>- Supports <b>MSA &amp; CNA</b></li> <li>- It works by <b>network policies</b></li> <li>- Supports <b>lightweight</b> protocols, such as HTTP, gRPC, Kafka</li> <li>- It is an <b>API-aware</b> network security <b>filtering</b>.</li> <li>- It uses Linux kernel technology called <b>BPF</b></li> <li>- <b>Simple</b></li> <li>- <b>Efficient</b></li> <li>- Enables building <b>gateway policies</b> which can be enforced network-layer and <b>application-layer</b> security <b>policies</b></li> <li>- <b>Scalability</b></li> <li>- <b>Multi-tenancy</b> 😊</li> <li>- L3 <b>Encryption</b> enforcement</li> </ul>	<ul style="list-style-type: none"> <li>- Creates a mesh <b>overlay</b> network <b>between</b> each of the nodes in the cluster</li> <li>- <b>Flexible</b> in the communication</li> <li>- <b>Simple</b></li> <li>- Enables service discovery using <b>micro DNS</b></li> <li>- <b>Encryption</b></li> <li>- Supports <b>multi-cast</b></li> <li>- Enables <b>portability</b></li> </ul>
<b>Net. Layering</b>	L3 network fabric	L3 & 7	N/A	L3 & 7	L3
<b>Stability</b>	Very high	Very high		High	High
<b>Service meshing</b>	Doesn't integrate and doesn't allow any network policy implementation	Integrates and enables defining rich network policy models	N/A	Integrates and enables defining rich network policy models	No
<b>Gateways?</b>	No	Yes	N/A	Yes	Yes
<b>Perf.</b>	Good	Very Good	N/A	Very Good	Very Good

# Calico



```
default via 10.0.2.1 dev enp0s3
10.0.2.0/24 dev enp0s3 proto kernel scope link src 10.0.2.6
192.168.81.0/24 dev enp0s8 proto kernel scope link src 192.168.81.101
blackhole 192.200.59.192/26 proto bird
192.200.59.193 dev calidf072d3c423 scope link
192.200.59.198 dev cali0aa3720a2c7 scope link
192.200.203.0/26 via 192.168.81.102 dev tunl0 proto bird onlink
```

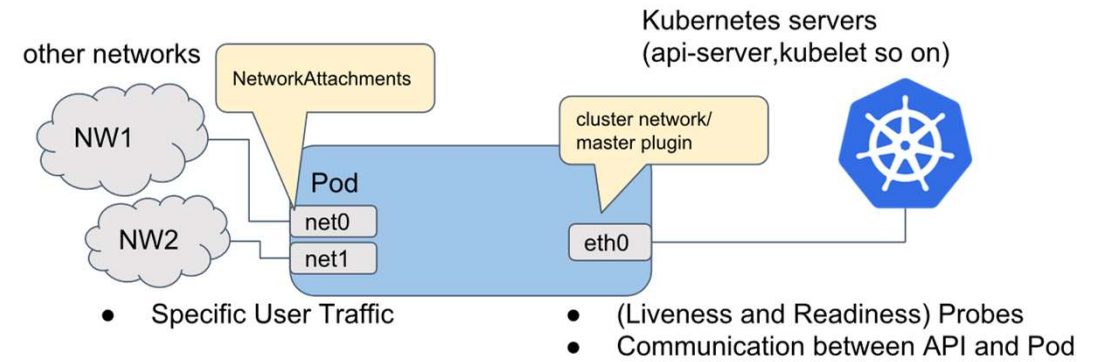
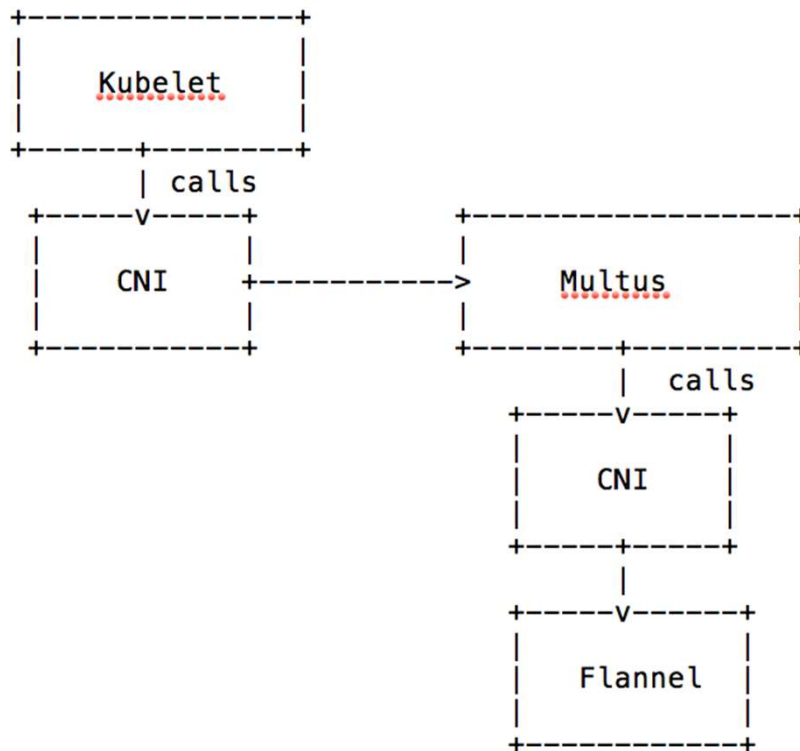
```
default via 10.0.2.1 dev enp0s3
10.0.2.0/24 dev enp0s3 proto kernel scope link src 10.0.2.7
192.168.81.0/24 dev enp0s8 proto kernel scope link src 192.168.81.102
192.200.59.192/26 via 192.168.81.101 dev tunl0 proto bird onlink
blackhole 192.200.203.0/26 proto bird
192.200.203.4 dev cali7bb4560a7c2 scope link
```



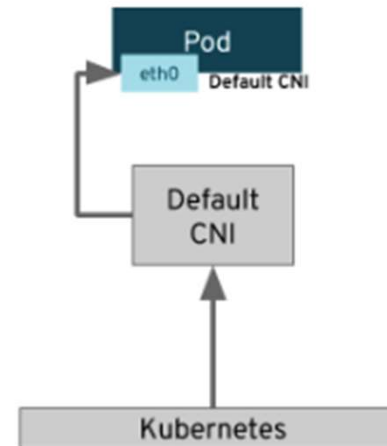
# Network Component – a comparison

Feature	Flannel	Calico	Multus	Cilium	Weave Net
<b>Brief features</b>	<ul style="list-style-type: none"> <li>- <b>Simple</b></li> <li>- <b>Flat</b>/overlay network</li> <li>- Each pod has <b>unique</b> ip, all pods run on the same network</li> <li>- Runs a <b>daemon</b> process flannel to update the ip routing table</li> <li>- Mapping of the subnet to host info stored in etcd</li> </ul>	<ul style="list-style-type: none"> <li>- Connects pods using the same IP networking principles as the internet</li> <li>- <b>Interoperable</b></li> <li>- <b>Flexible</b></li> <li>- Enable security enforcement (<b>self workload firewall</b>)</li> <li>- True <b>Cloud native</b> scalability</li> <li>- leverage best practice cloud-native design patterns</li> <li>- Enables hybrid soln using <b>BGP</b></li> </ul>	<ul style="list-style-type: none"> <li>- It enables attaching <b>multiple network</b> interfaces to the pods by creating homed pod, it is <b>magic</b> 😊</li> <li>- It is a <b>meta-plugin</b></li> <li>- It highly support <b>multitenancy</b></li> </ul>	<ul style="list-style-type: none"> <li>- Supports <b>MSA &amp; CNA</b></li> <li>- It works by <b>network policies</b></li> <li>- Supports <b>lightweight</b> protocols, such as HTTP, gRPC, Kafka</li> <li>- It is an <b>API-aware</b> network security <b>filtering</b>.</li> <li>- It uses Linux kernel technology called <b>BPF</b></li> <li>- <b>Simple</b></li> <li>- <b>Efficient</b></li> <li>- Enables building <b>gateway policies</b> which can be enforced network-layer and <b>application-layer</b> security <b>policies</b></li> <li>- <b>Scalability</b></li> <li>- <b>Multi-tenancy</b> 😊</li> <li>- L3 <b>Encryption</b> enforcement</li> </ul>	<ul style="list-style-type: none"> <li>- Creates a mesh <b>overlay</b> network <b>between</b> each of the nodes in the cluster</li> <li>- <b>Flexible</b> in the communication</li> <li>- <b>Simple</b></li> <li>- Enables service discovery using <b>micro DNS</b></li> <li>- <b>Encryption</b></li> <li>- Supports <b>multi-cast</b></li> <li>- Enables <b>portability</b></li> </ul>
<b>Net. Layering</b>	L3 network fabric	L3 & 7	N/A	L3 & 7	L3
<b>Stability</b>	Very high	Very high		High	High
<b>Service meshing</b>	Doesn't integrate and doesn't allow any network policy implementation	Integrates and enables defining rich network policy models	N/A	Integrates and enables defining rich network policy models	No
<b>Gateways?</b>	No	Yes	N/A	Yes	Yes
<b>Perf.</b>	Good	Very Good	N/A	Very Good	Very Good

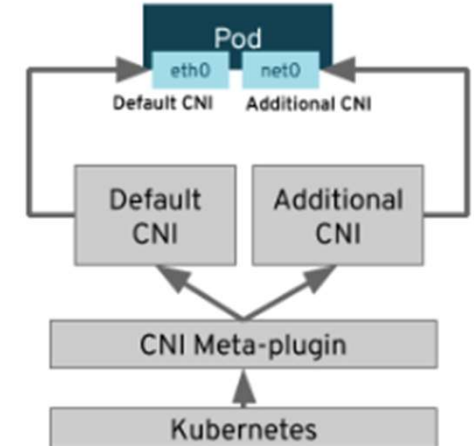
# Multus



Pod without CNI meta-plugin



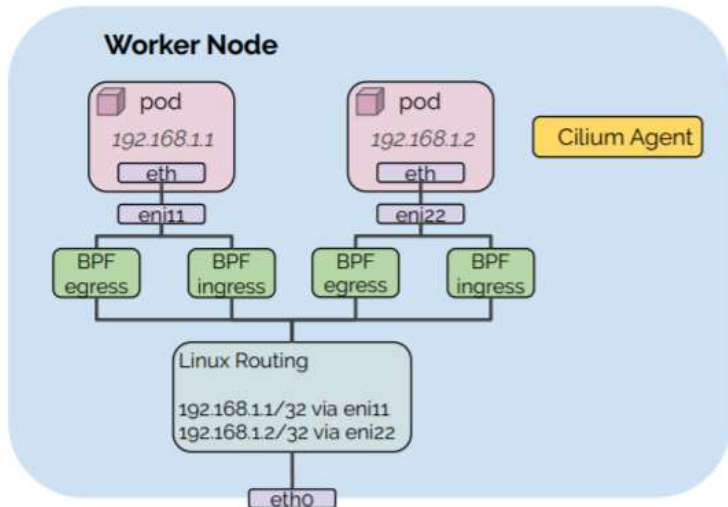
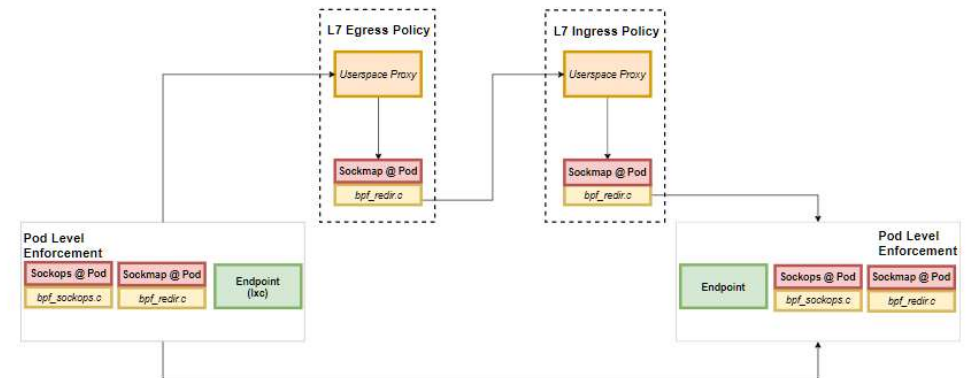
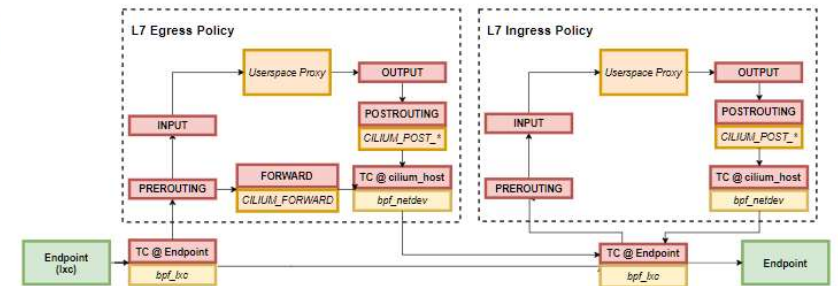
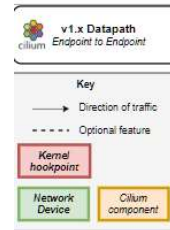
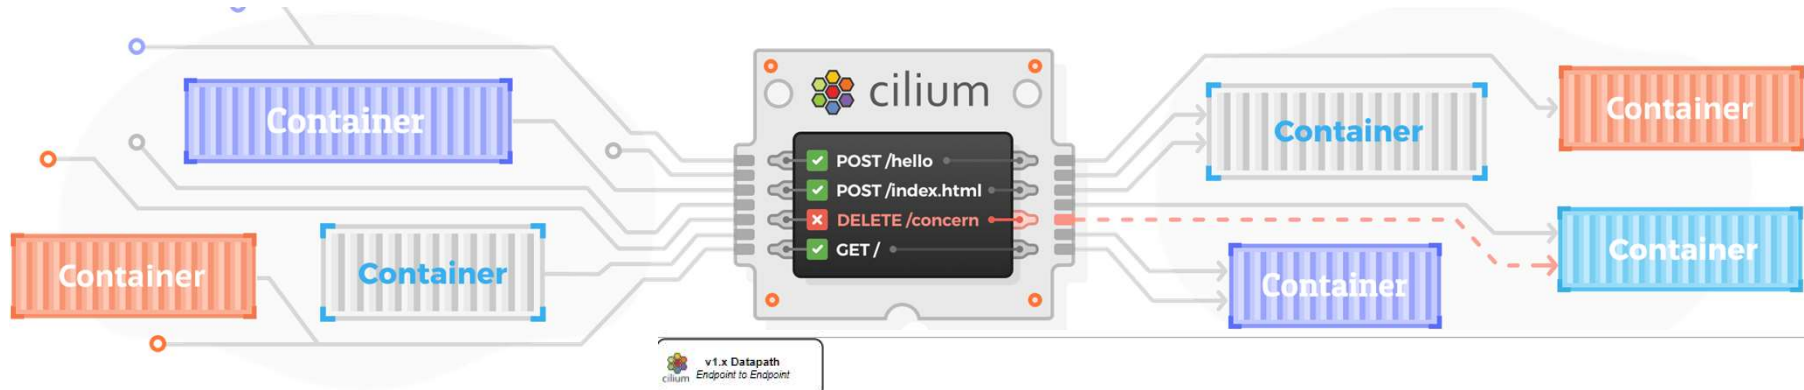
Pod with CNI meta-plugin



# Network Component – a comparison

Feature	Flannel	Calico	Multus	Cilium	Weave Net
<b>Brief features</b>	<ul style="list-style-type: none"> <li>- <b>Simple</b></li> <li>- <b>Flat</b>/overlay network</li> <li>- Each pod has <b>unique</b> ip, all pods run on the same network</li> <li>- Runs a <b>daemon</b> process flannel to update the ip routing table</li> <li>- Mapping of the subnet to host info stored in etcd</li> </ul>	<ul style="list-style-type: none"> <li>- Connects pods using the same IP networking principles as the internet</li> <li>- <b>Interoperable</b></li> <li>- <b>Flexible</b></li> <li>- Enable security enforcement (<b>self workload firewall</b>)</li> <li>- True <b>Cloud native</b> scalability</li> <li>- leverage best practice cloud-native design patterns</li> <li>- Enables hybrid soln using <b>BGP</b></li> </ul>	<ul style="list-style-type: none"> <li>- It enables attaching <b>multiple network</b> interfaces to the pods by creating homed pod, it is <b>magic</b> 😊</li> <li>- It is a <b>meta-plugin</b></li> <li>- It highly support <b>multitenancy</b></li> </ul>	<ul style="list-style-type: none"> <li>- Supports <b>MSA &amp; CNA</b></li> <li>- It works by <b>network policies</b></li> <li>- Supports <b>lightweight</b> protocols, such as HTTP, gRPC, Kafka</li> <li>- It is an <b>API-aware</b> network security <b>filtering</b>.</li> <li>- It uses Linux kernel technology called <b>BPF</b></li> <li>- <b>Simple</b></li> <li>- <b>Efficient</b></li> <li>- Enables building <b>gateway policies</b> which can be enforced network-layer and <b>application-layer</b> security <b>policies</b></li> <li>- <b>Scalability</b></li> <li>- <b>Multi-tenancy</b> 😊</li> <li>- L3 <b>Encryption</b> enforcement</li> </ul>	<ul style="list-style-type: none"> <li>- Creates a mesh <b>overlay</b> network <b>between</b> each of the nodes in the cluster</li> <li>- <b>Flexible</b> in the communication</li> <li>- <b>Simple</b></li> <li>- Enables service discovery using <b>micro DNS</b></li> <li>- <b>Encryption</b></li> <li>- Supports <b>multi-cast</b></li> <li>- Enables <b>portability</b></li> </ul>
<b>Net. Layering</b>	L3 network fabric	L3 & 7	N/A	L3 & 7	L3
<b>Stability</b>	Very high	Very high		High	High
<b>Service meshing</b>	Doesn't integrate and doesn't allow any network policy implementation	Integrates and enables defining rich network policy models	N/A	Integrates and enables defining rich network policy models	No
<b>Gateways?</b>	No	Yes	N/A	Yes	Yes
<b>Perf.</b>	Good	Very Good	N/A	Very Good	Very Good

# Cilium



**Cilium**

- Load-balancing
- Network policy
- Encryption
- Multi-cluster
- Visibility



# Network Component – a comparison

Feature	Flannel	Calico	Multus	Cilium	Weave Net
<b>Brief features</b>	<ul style="list-style-type: none"> <li>- <b>Simple</b></li> <li>- <b>Flat</b>/overlay network</li> <li>- Each pod has <b>unique</b> ip, all pods run on the same network</li> <li>- Runs a <b>daemon</b> process flannel to update the ip routing table</li> <li>- Mapping of the subnet to host info stored in etcd</li> </ul>	<ul style="list-style-type: none"> <li>- Connects pods using the same IP networking principles as the internet</li> <li>- <b>Interoperable</b></li> <li>- <b>Flexible</b></li> <li>- Enable security enforcement (<b>self workload firewall</b>)</li> <li>- True <b>Cloud native</b> scalability</li> <li>- leverage best practice cloud-native design patterns</li> <li>- Enables hybrid soln using <b>BGP</b></li> </ul>	<ul style="list-style-type: none"> <li>- It enables attaching <b>multiple network</b> interfaces to the pods by creating homed pod, it is <b>magic</b> 😊</li> <li>- It is a <b>meta-plugin</b></li> <li>- It highly support <b>multitenancy</b></li> </ul>	<ul style="list-style-type: none"> <li>- Supports <b>MSA &amp; CNA</b></li> <li>- It works by <b>network policies</b></li> <li>- Supports <b>lightweight</b> protocols, such as HTTP, gRPC, Kafka</li> <li>- It is an <b>API-aware</b> network security <b>filtering</b>.</li> <li>- It uses Linux kernel technology called <b>BPF</b></li> <li>- <b>Simple</b></li> <li>- <b>Efficient</b></li> <li>- Enables building <b>gateway policies</b> which can be enforced network-layer and <b>application-layer</b> security <b>policies</b></li> <li>- <b>Scalability</b></li> <li>- <b>Multi-tenancy</b> 😊</li> <li>- L3 <b>Encryption</b> enforcement</li> </ul>	<ul style="list-style-type: none"> <li>- Creates a mesh <b>overlay</b> network <b>between</b> each of the nodes in the cluster</li> <li>- <b>Flexible</b> in the communication</li> <li>- <b>Simple</b></li> <li>- Enables service discovery using <b>micro DNS</b></li> <li>- <b>Encryption</b></li> <li>- Supports <b>multi-cast</b></li> <li>- Enables <b>portability</b></li> </ul>
<b>Net. Layering</b>	L3 network fabric	L3 & 7	N/A	L3 & 7	L3
<b>Stability</b>	Very high	Very high		High	High
<b>Service meshing</b>	Doesn't integrate and doesn't allow any network policy implementation	Integrates and enables defining rich network policy models	N/A	Integrates and enables defining rich network policy models	No
<b>Gateways?</b>	No	Yes	N/A	Yes	Yes
<b>Perf.</b>	Good	Very Good	N/A	Very Good	Very Good

# Container Component(s)

# Container Runtime Component – Main Target Requirements

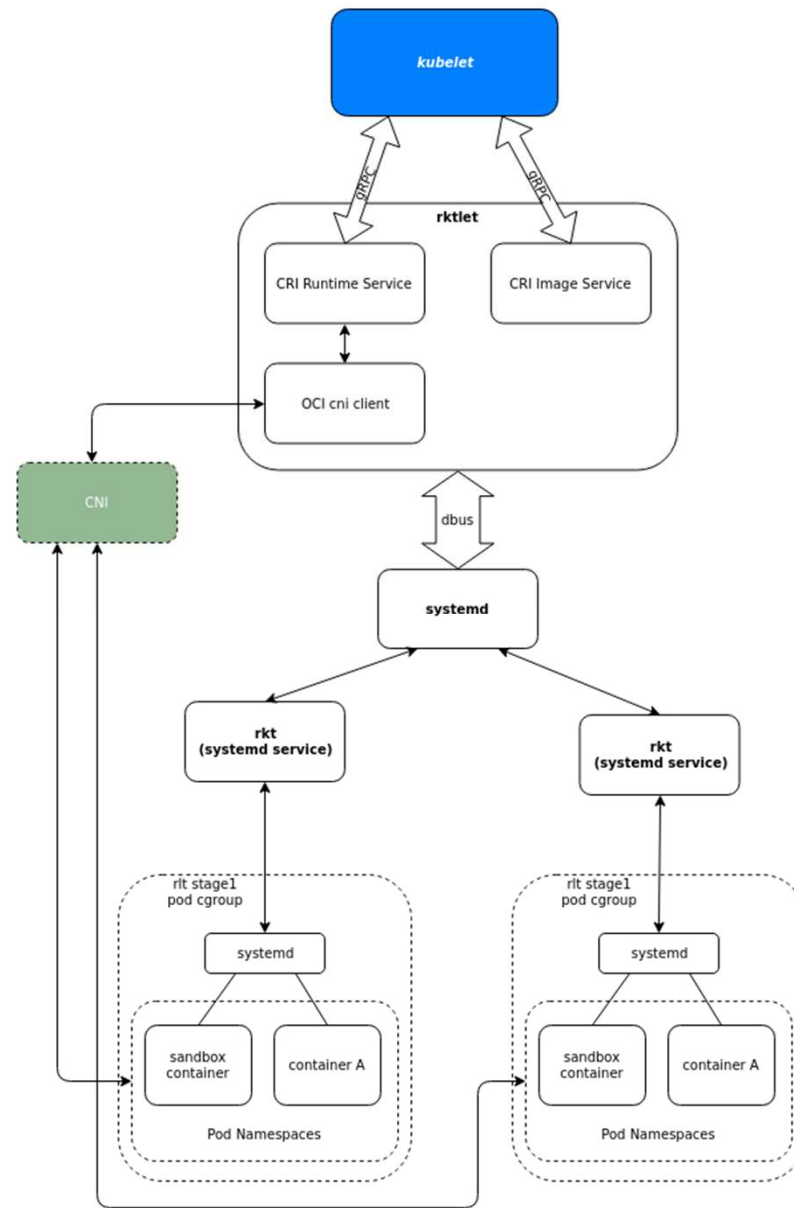
- **Standardize** the communication with Container engine, Container Runtime Interface **CRI**.
- Manages the **namespace isolation** and **resource allocation** at the OS levels using Linux cgroup and namespaces
- Creates & build a **container** using an **Image**
- It the **runtime** the **container** run above
- **Abstract** the **container** from the hosting **OS**
- **Integrates** with image registry
- **Supports High level** and **low level** container runtime
- Manages **containers' lifecycle**
- Support both running **stateful** (storage is a must here) and **stateless** containers
- Support **logging** and **troubleshooting** of a running container

# Container Runtime Component – a comparison

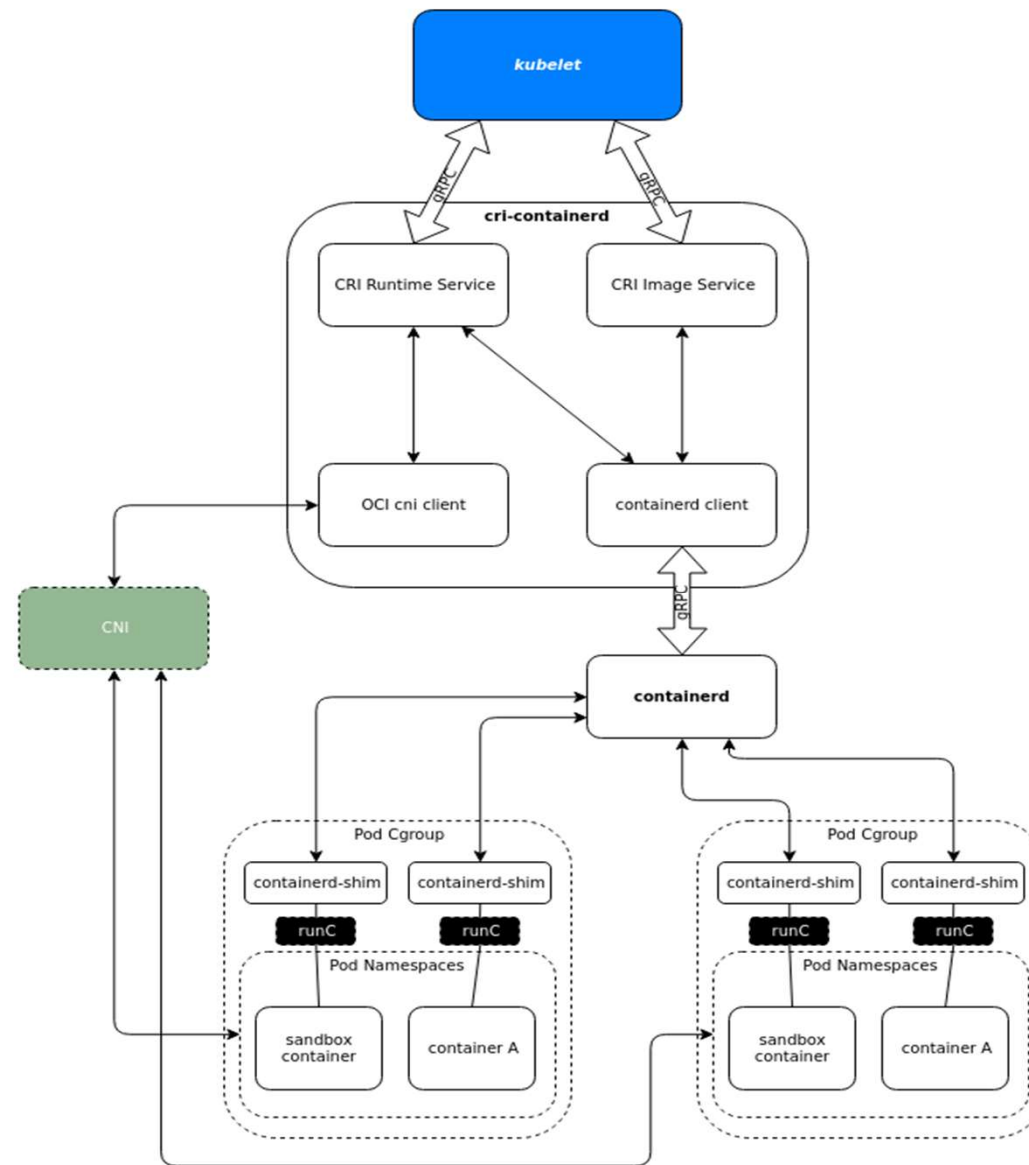
Feature	RunC	Rkt	CRI-O	Docker	Containerd
<b>Brief features</b>	<ul style="list-style-type: none"> <li>- <b>Low-level</b> container runtime</li> <li>- Implements <b>OCI</b></li> <li>- Requires <b>expertise</b> of the underlying host OS and configuration.</li> <li>- Does <b>not verify</b> container images or prepare the FS</li> <li>- <b>No centralized daemon</b></li> </ul>	<ul style="list-style-type: none"> <li>- <b>High level</b> container engine</li> <li>- Built as an <b>alternative</b> to docker in K8s</li> <li>- Can group containers/apps in a shared context (<b>pod</b>)</li> <li>- <b>No centralized</b> init <b>daemon</b></li> <li>- Support different container/pod <b>configurations</b> (like isolation parameters)</li> <li>- Better <b>isolation</b> Each pod runs in a different process</li> <li>- Supports <b>OCI</b></li> </ul>	<ul style="list-style-type: none"> <li>- <b>High level</b> container engine</li> <li>- support <b>OCI</b> and implements <b>CRI</b></li> <li>- It uses <b>runC</b> by default</li> <li>- Can plug any <b>OCI runtime</b></li> <li>- <b>Light weight</b> (lots of small components, with defined roles &amp; collaborating flows)</li> <li>- <b>Decentralized architecture</b></li> <li>- <b>Secured</b> by as CRI-O containers are children of the process that spawned it</li> <li>- Fully compatible with <b>K8s Roadmap</b> and community</li> <li>- Implements <b>CNI</b> which make it more standard from a network setup</li> <li>- <b>Fast</b></li> <li>- Can run Docker images</li> </ul>	<ul style="list-style-type: none"> <li>- <b>Not Standard</b></li> <li>- <b>Heavyweight</b>/fat daemon</li> <li>- Central architecture</li> <li>- Has <b>security constraints</b></li> <li>- Has no <b>limitation</b> 😊</li> </ul>	<ul style="list-style-type: none"> <li>- Runs as a <b>daemon</b></li> <li>- Implements <b>CRI</b></li> </ul>
<b>Security</b>	Yes	Yes	Yes	No	No
<b>Perf.</b>	Good	Very Good	Very Good	Very Good	Very Good
<b>Standard</b>	Yes	Yes	Yes	No	Yes
<b>Stability</b>	Very high	High	Very High	Very high	High



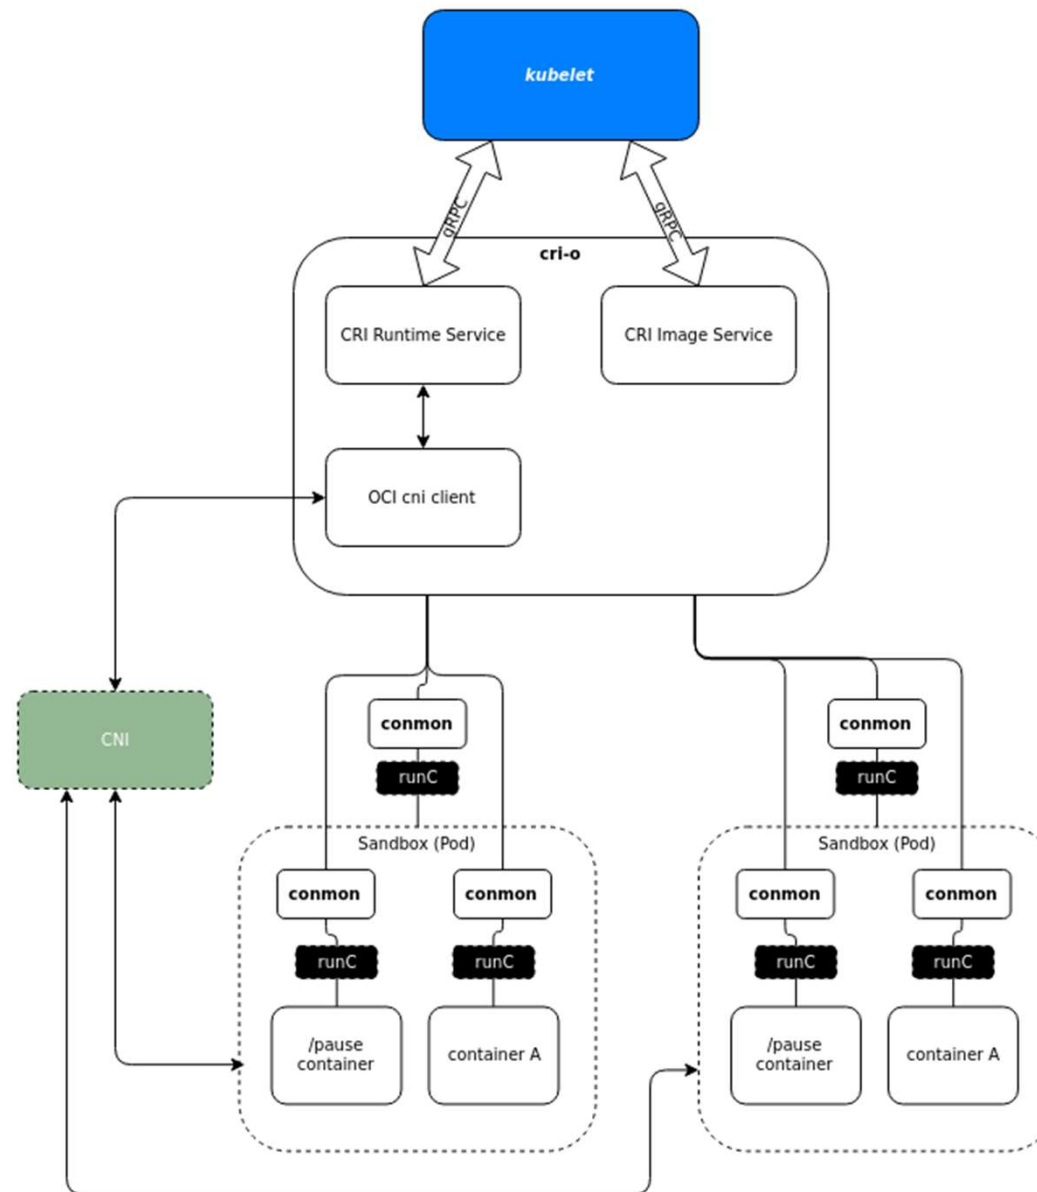
# rkt



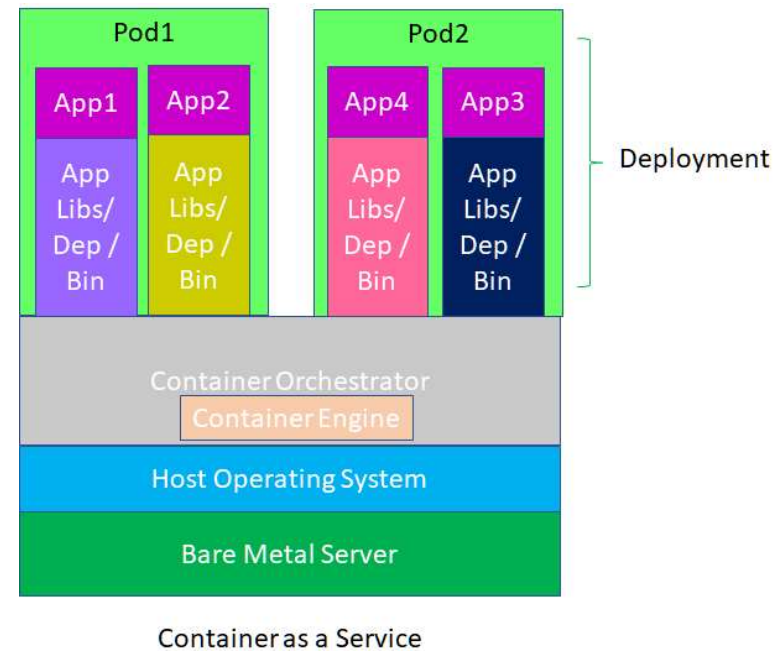
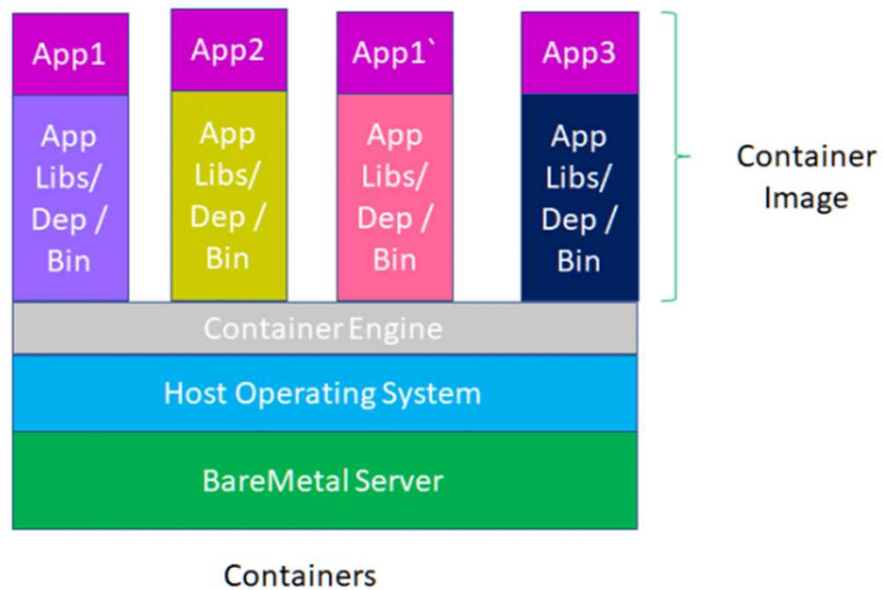
# containerd



# CRIO



# Container engine vs Containers Orchestrator





# Monitoring Component(s)

# Monitoring Component – Main Target Requirements

- Monitor both cluster metrics as well as app/workload metrics
- Monitor the health of the cluster
- Monitor resources consumptions/utilization (node & pods)
- Availability (node and K8s objects)
- Gather k8s, apps and container metrics

# Monitoring Component – solutions

Heapster, InfluxDB, & Grafana	Prometheus & Grafana	Heapster & ELK Stack	Datadog	Dynatrace
<ul style="list-style-type: none"> <li>- Heapster, is a uniform platform which push monitoring metrics to a external tool to process</li> <li>- InfluxDB is used to store the collected metrics</li> <li>- Grafana is used to visualize the collected info</li> <li>- Simple</li> </ul>	<ul style="list-style-type: none"> <li>- Prometheus is a platform to gather metrics</li> <li>- Grafana is used to visualize the collected info</li> <li>- Simple</li> <li>- Flexible</li> </ul>	<ul style="list-style-type: none"> <li>- Heapster, is a uniform platform which push monitoring metrics to a external tool to process</li> <li>- Use ELK or Elastic stack which includes Elasticsearch, Logstash, and Kibana, which define the data pipeline.</li> <li>- Central logging and dashboard and can hold some sort of analytics on the gathered data.</li> <li>- Powerful in analytics</li> <li>- Flexible</li> </ul>	<ul style="list-style-type: none"> <li>- Simple</li> <li>- Flexible data pipeline</li> <li>- Uses DaemonSet agent</li> </ul>	<ul style="list-style-type: none"> <li>- Uses DaemonSet agent</li> <li>- Not flexible in the gather metrics</li> <li>- complex</li> </ul>



# Service Mesh / Service Collaboration Component(s)



# Service Mesh / Service Collaboration Component

**What is Service Mesh?**



# Service Mesh / Service Collaboration Component

## – main challenges before service mesh

- MSA is small and fabulous but hard to control, watch and govern
- CNA is awesome but hard to troubleshoot
- MSA and CNA is about hybrid development and agility targeting time to market so how to balance that keeping the aspects of governance & quality → hard balance 😊
- Security
- Monitoring
- Managing dependencies

# Service Mesh / Service Collaboration Component

## – target requirements

- Services or app must be **self**:
  - **Governed** – follow all policies and notify/get notified for changes.
  - **Secured** – not only the 2 authes but ability to defend itself
  - **Monitored** – metrics gathering
  - **Logging** – metrics gathering
  - **Gateway** (service/app discovery)
- Not compromising to the freedom of the developer 😊
- Must be **FAST** and **lightweight**
- **Observer** and **enforcer** and not just an **reactor**
- Supports well known protocols as HTTP2, gRPC ...

**It is just like Aspect Oriented Programming, AOP and IoC, Injection and Inversion of Controller 😊**

# Service Mesh / Service Collaboration Component

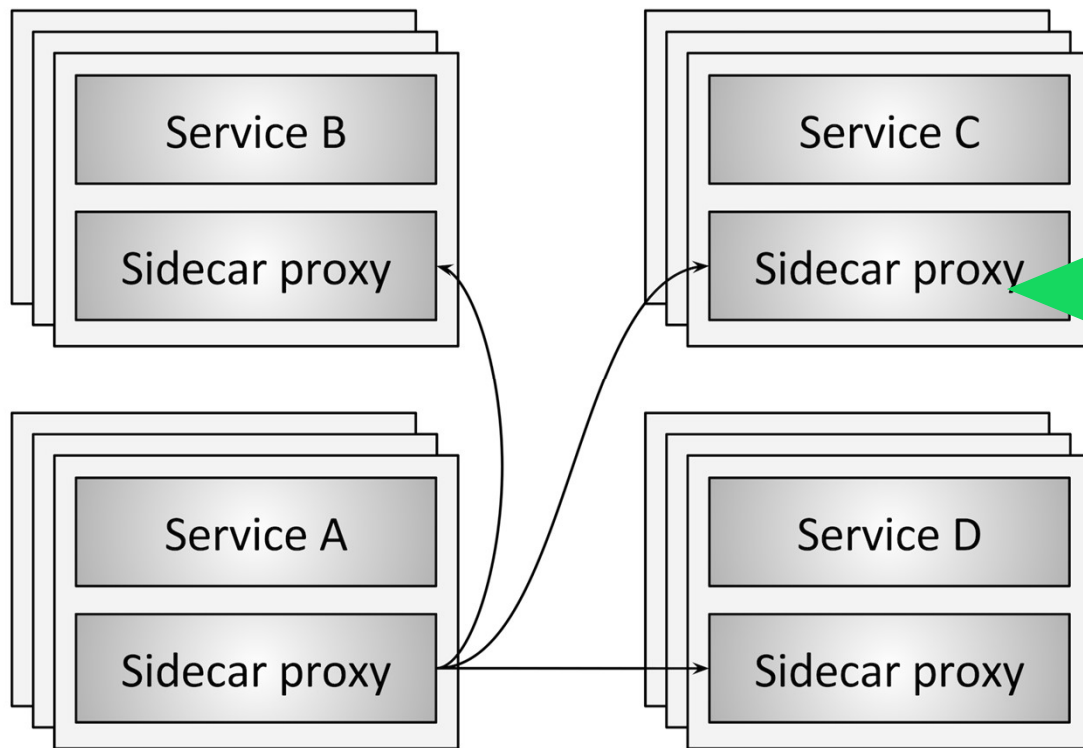
## – available Service Mesh FWs

- **Istio**, super powerful implementation 😊 → it uses envoy as a data plane, think of it as an extension.
- **Google Service Mesh**, it uses Istio but with more visualization capabilities and troubleshooting
- **AWS app Mesh**, it uses envoy it only provides parts of the requirements as the proxying and the monitoring metrics
- **Envoy**, it only focus on the data more than the control (i.e. what to do with the data)
- **Azure Service Mesh**, it is fully built by Microsoft, it covers most of the requirements



# Service Mesh / Service Collaboration Component

## – More into Service Mesh – Istio &



Data Plane (mainly gathering data transforming or converting it to events and forward it):

1. The 2 auths
2. Service discovery
3. Monitoring metrics and Health checking
4. Load balancing and Routing
5. Observation to any metrics or events
6. Enables lots of MSA patterns such as Circuit breakers
7. K8s and CNA/MSA deployment strategies

**Service mesh works on the network packet level so it is an INF layer still**  
**Sidecars are always stateless 😊**

# Service Mesh / Service Collaboration Component

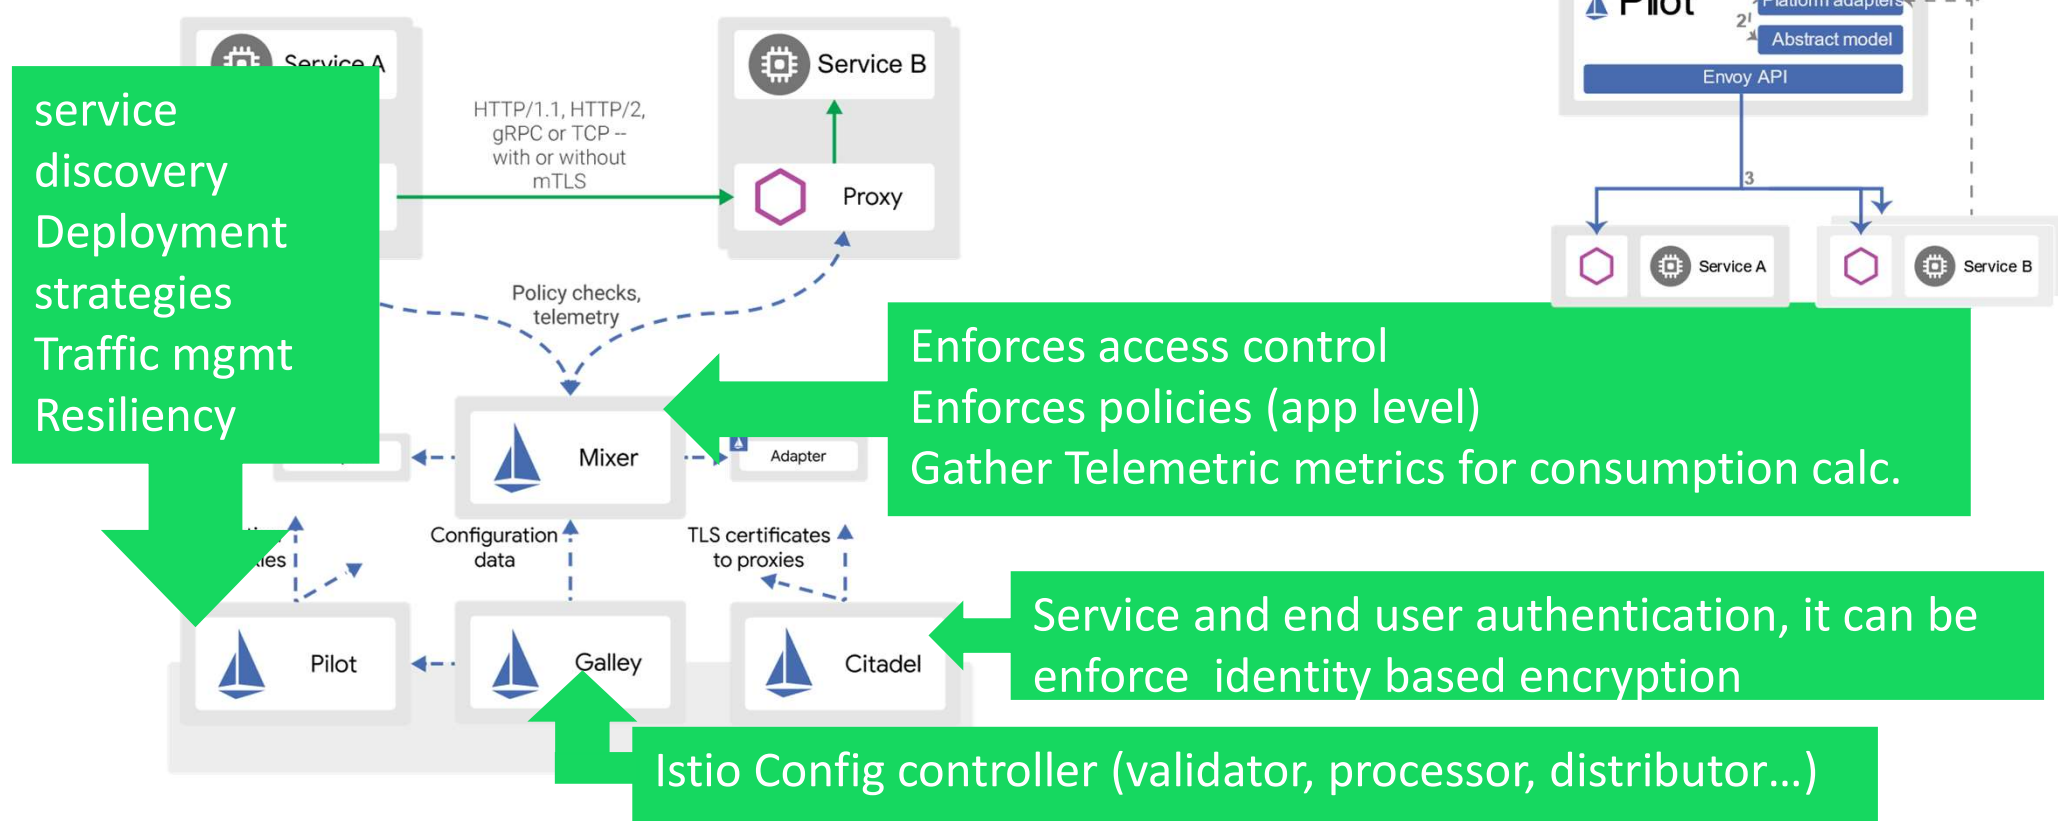
- More into Service Mesh – Istio & Envoy ☺

**Is it Enough?**



# Service Mesh / Service Collaboration Component

## – More into Service Mesh – Istio & Envoy ☺

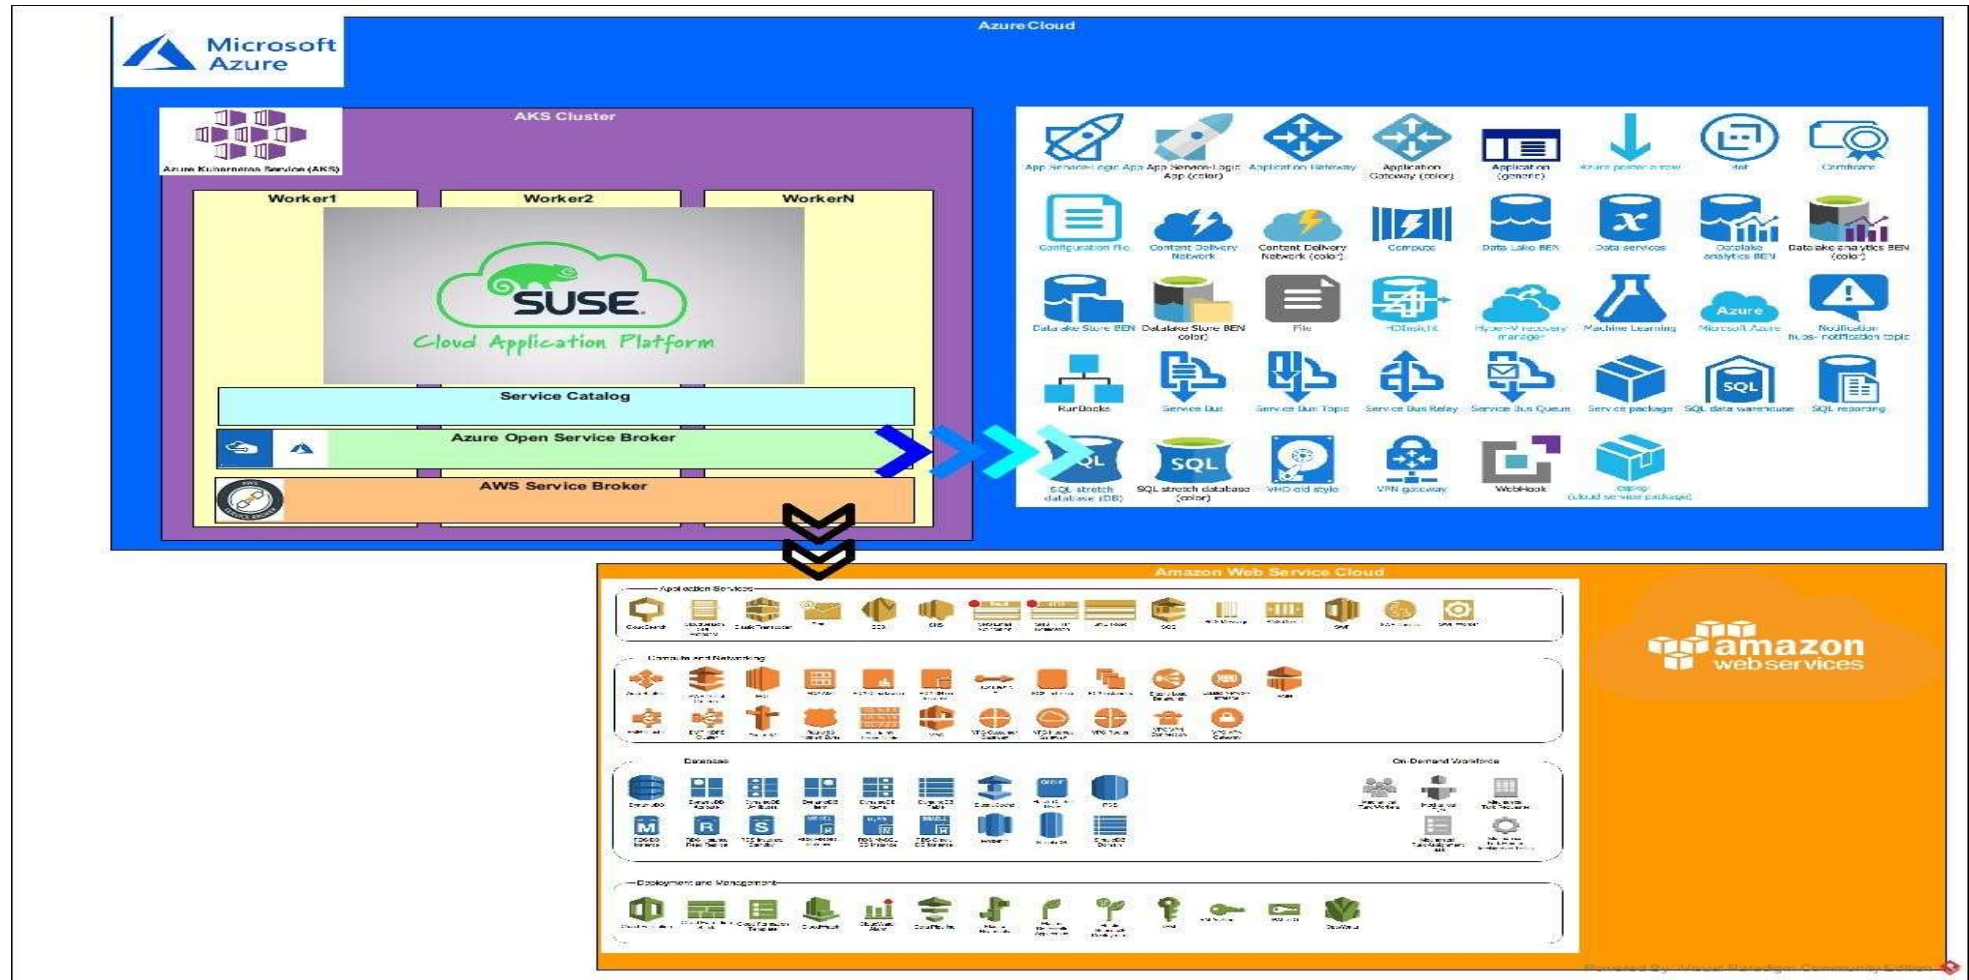


**Data Plane is Magic as it provides network abstraction and still works on fact but it is not enough it still require and a fairy and their magical wands ☺**

# Kubernetes UseCase 😊



# Multi/Hybrid – Cloud



<https://www.suse.com/c/cloud-native-applications-in-azure-supporting-hybrid-cloud/>





# Business Benefits Multicloud

- Elasticity/scalability

Pay as you grow/**On-demand availability**

cost effective and scalable

- More vendor agnostic/ higher flexibility

- Speed

**accelerate time to market**

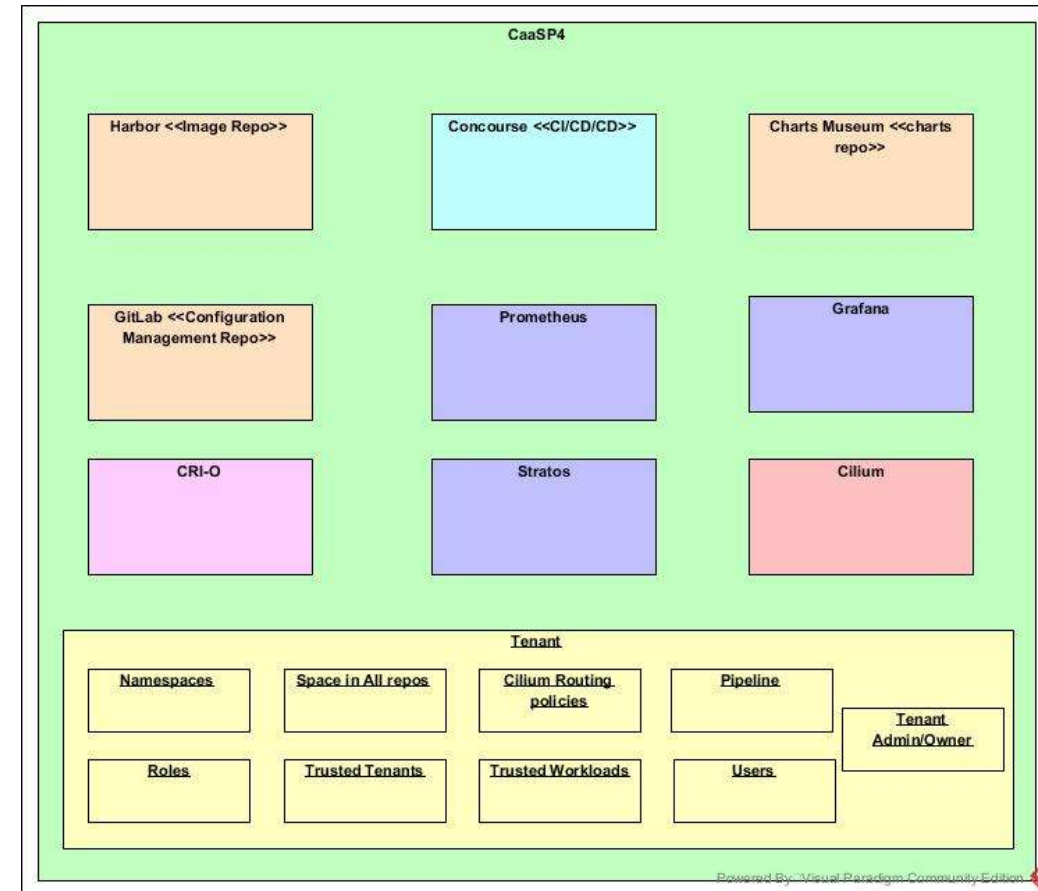
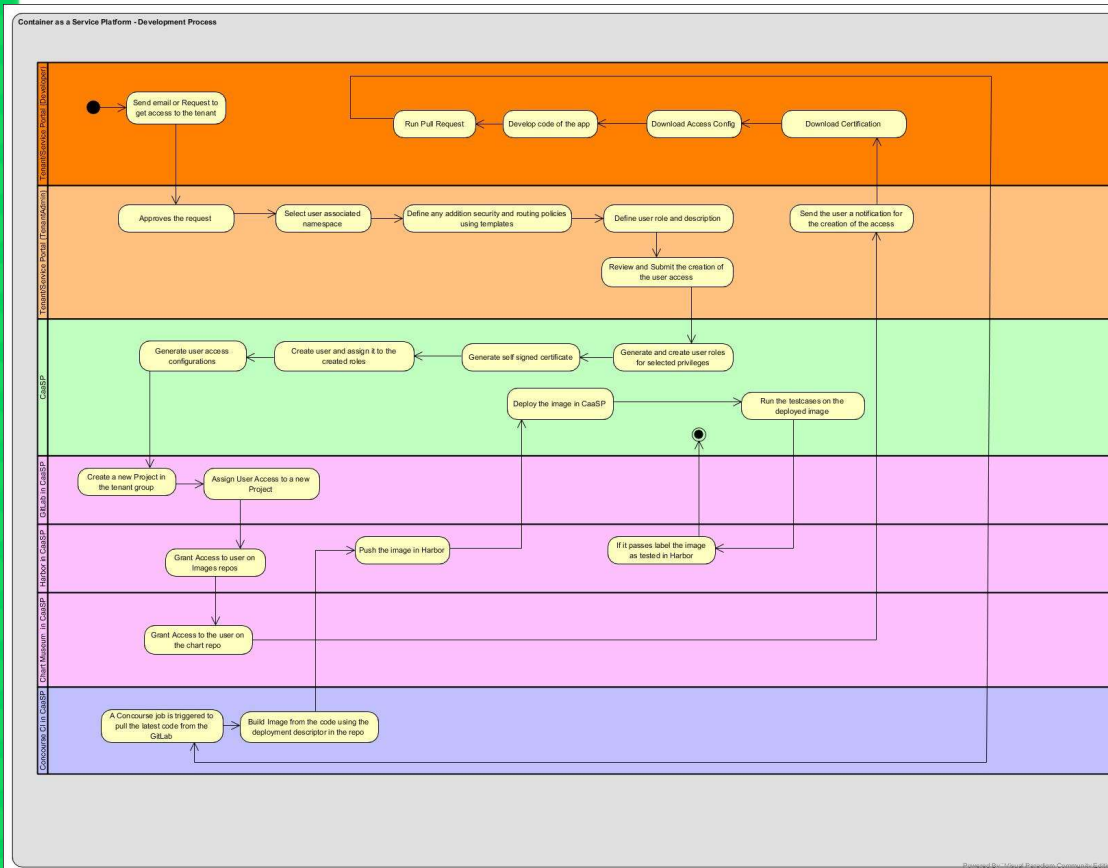
- Innovation (due to competition in Cloud)

- vast **variety of different services** (IaaS, SaaS, PaaS)

- Cloudmanagement and **optimization increase enhancements** in the process

- Greater Choice enables better cost control **BUT danger of Cloud Sprawl**  
**and need for a „single pane of glass“ to prevent Sprawl**

# DevOps



<https://www.suse.com/c/highly-automated-and-secured-multi-tenancy-using-suse-caas-platform-4/>

# Business Benefits DevOps

- Managing services not „IT assests“ – **business driven**
- **Faster delivery** of features or improvents in the sw
- Improved communication and **collaboration**
- More time to **innovate**
- Happier and **more productive teams**
- Faster recovery from incidents
- Lower change failure rate (due to smaller components to oversee)
- Results in **increased Customer Satisfaction**
- **Faster revenue** generation
- Increased efficiency and productivity

# Please join us on our next session:



January 17<sup>th</sup> 2020  
09:00 AM GMT

**About making Choices – CaaSPv4 as SUSE's  
empowering of Kubernetes**

# Q&A





Thank you



## **Unpublished Work of SUSE LLC. All Rights Reserved.**

This work is an unpublished work and contains confidential, proprietary and trade secret information of SUSE LLC. Access to this work is restricted to SUSE employees who have a need to know to perform tasks within the scope of their assignments. No part of this work may be practiced, performed, copied, distributed, revised, modified, translated, abridged, condensed, expanded, collected, or adapted without the prior written consent of SUSE. Any use or exploitation of this work without authorization could subject the perpetrator to criminal and civil liability.

## **General Disclaimer**

This document is not to be construed as a promise by any participating company to develop, deliver, or market a product. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. SUSE makes no representations or warranties with respect to the contents of this document, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. The development, release, and timing of features or functionality described for SUSE products remains at the sole discretion of SUSE. Further, SUSE reserves the right to revise this document and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes. All SUSE marks referenced in this presentation are trademarks or registered trademarks of Novell, Inc. in the United States and other countries. All third-party trademarks are the property of their respective owners.