Kubernetes

Architecture, Setup, Usage

Kubernetes Architecture

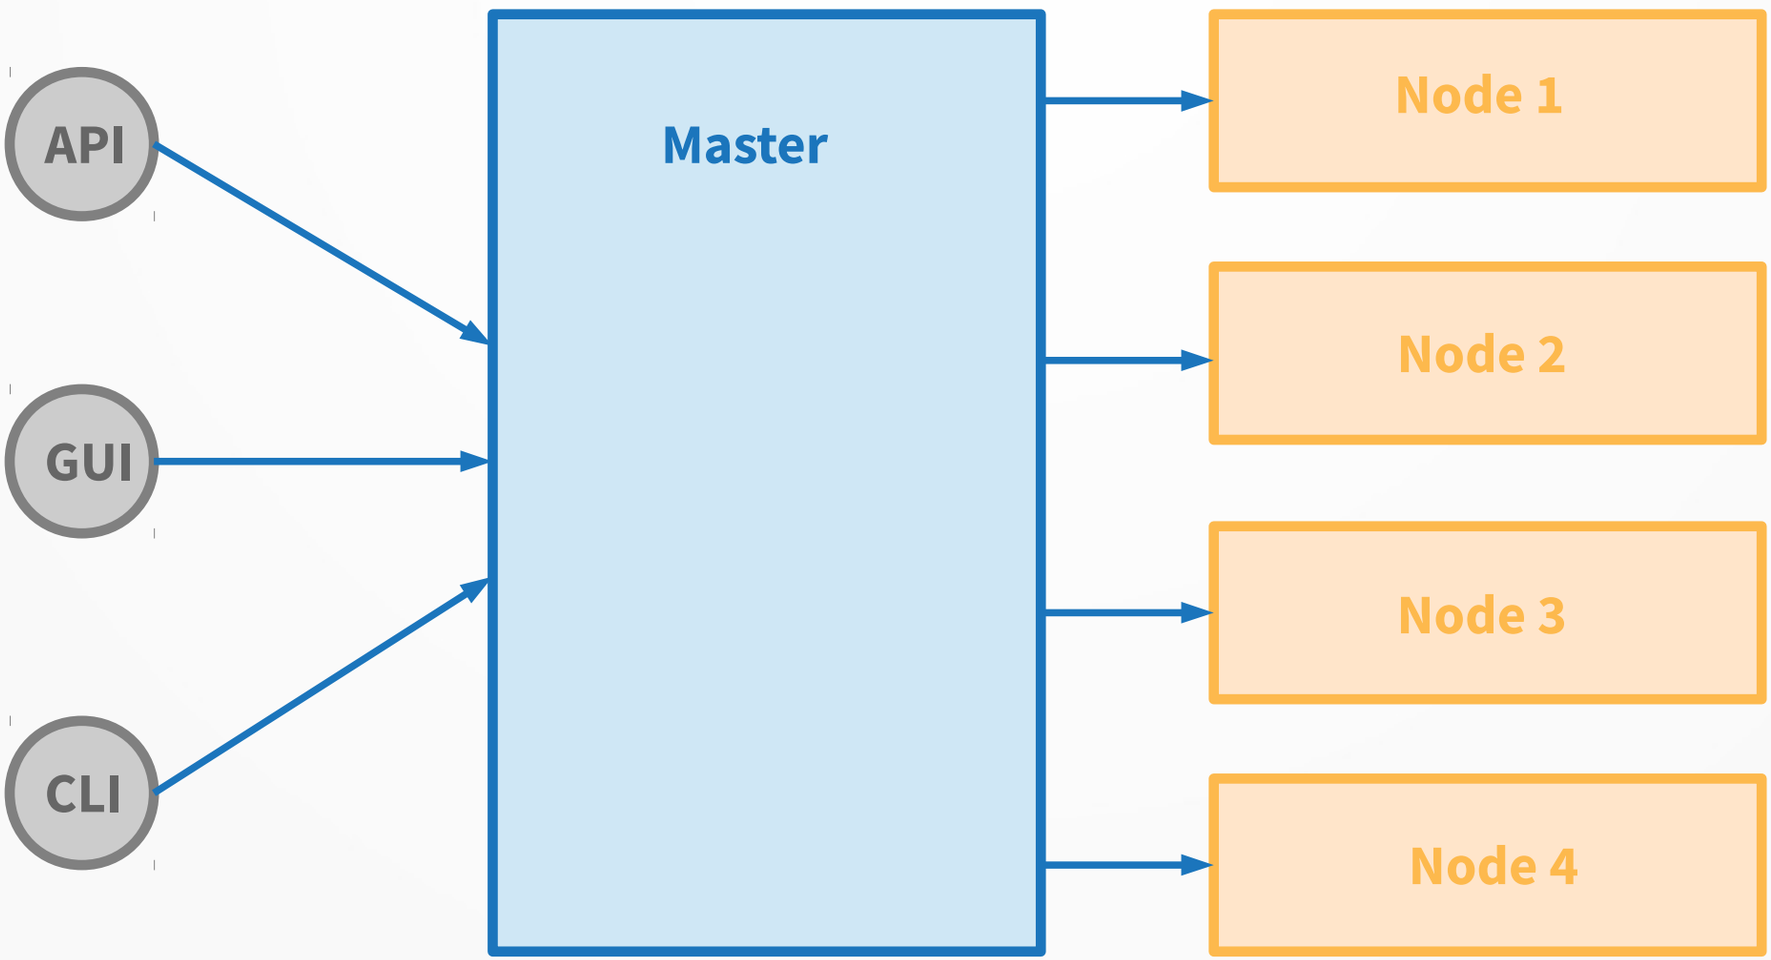

Kubernetes Architecture

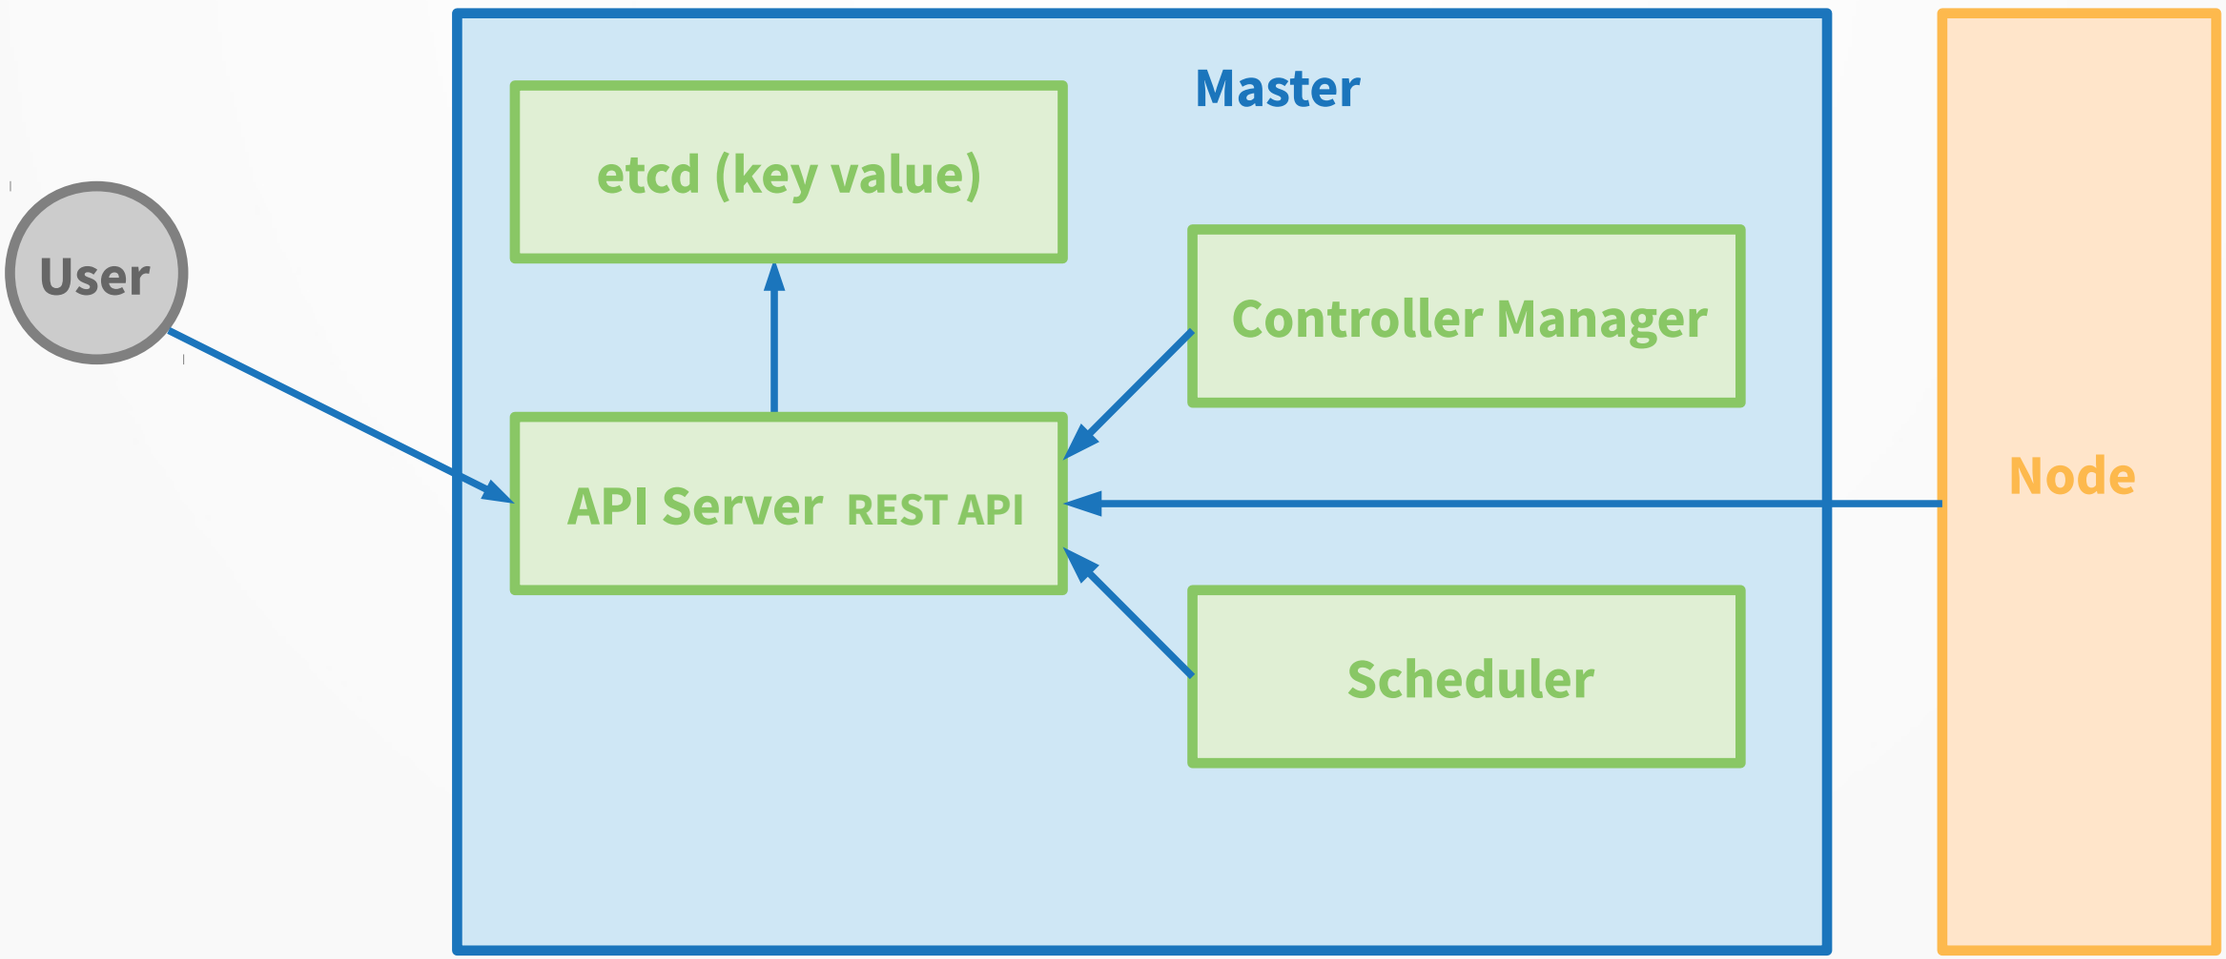

Kubernetes Architecture

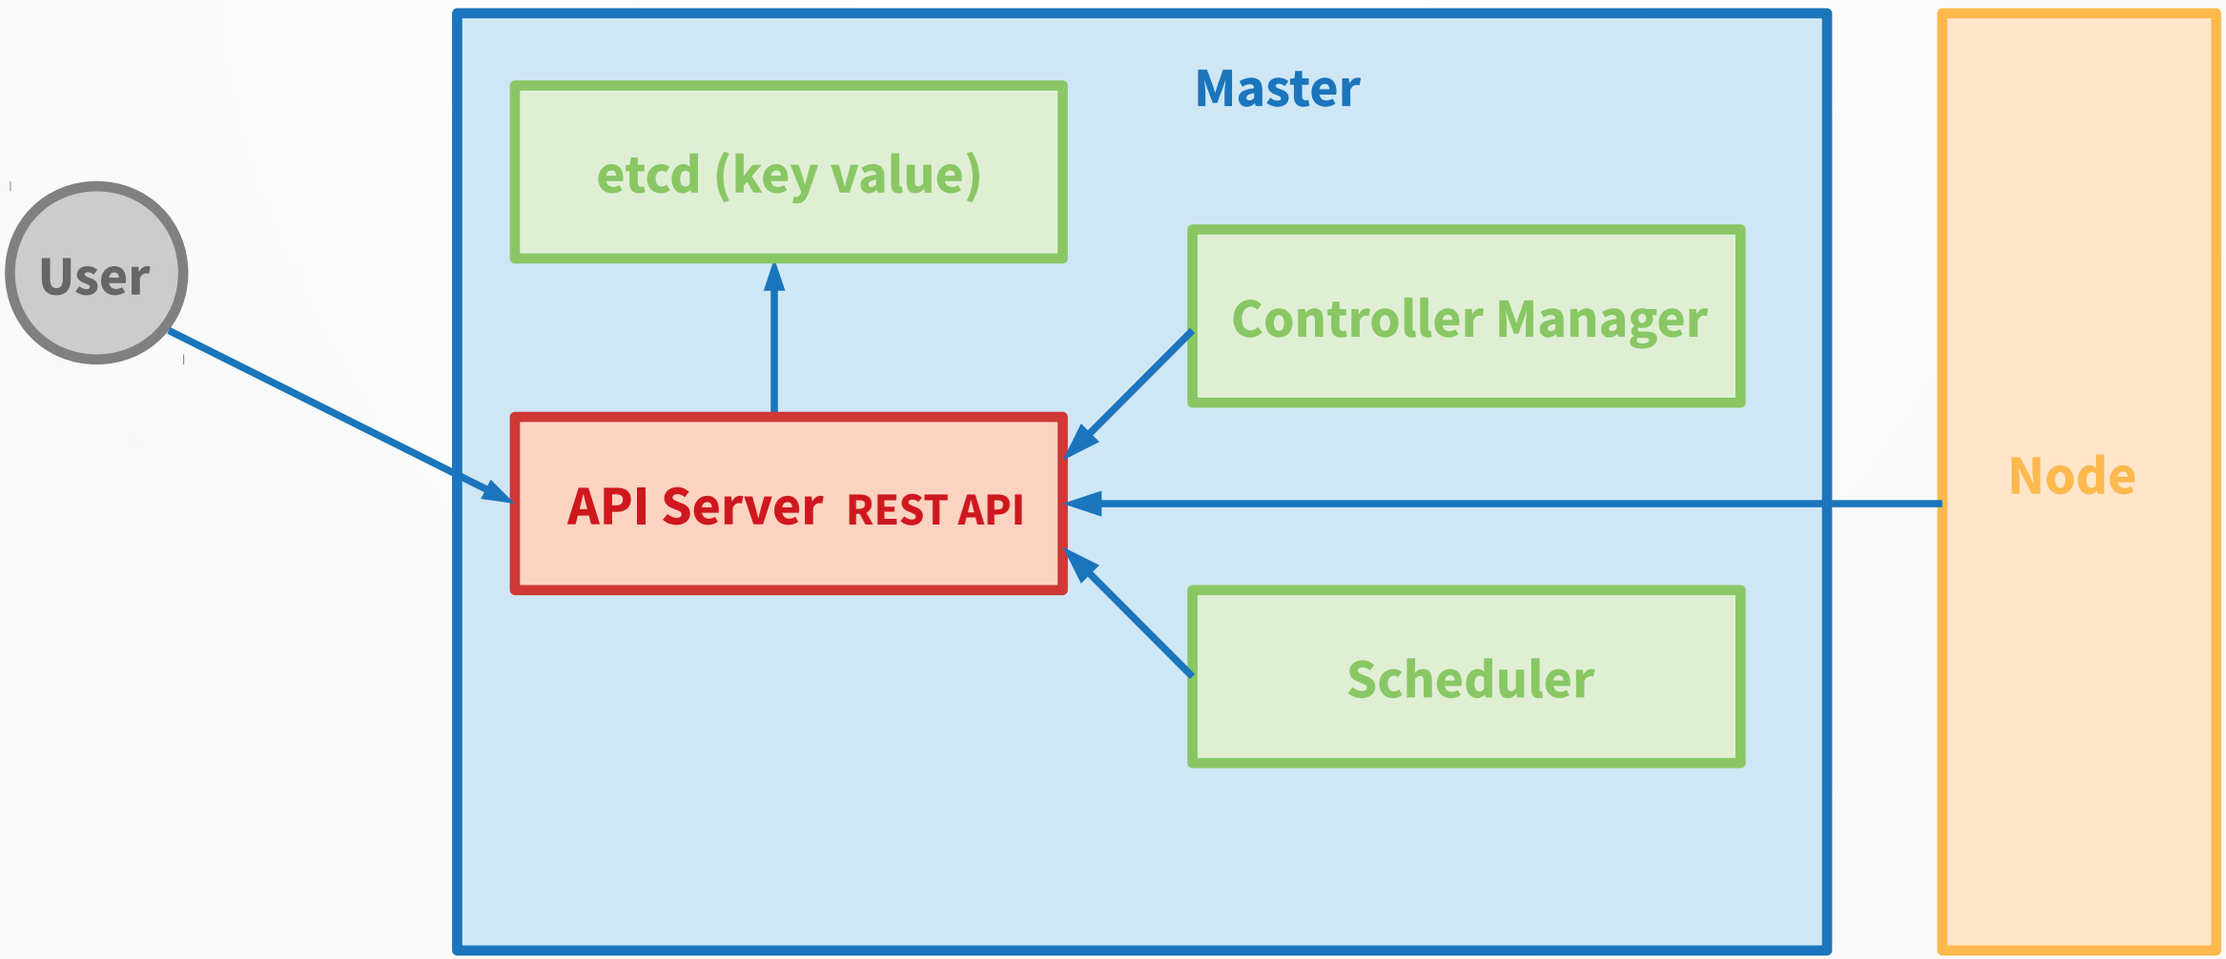

Kubernetes Architecture

API Server

- The API Server is the main management point of the entire cluster.
- It processes REST operations, validates them, and updates the corresponding objects in etcd.
- The API Server is the only Kubernetes component that connects to etcd, all the other components must go through the API Server to work with the cluster state.

Kubernetes Architecture

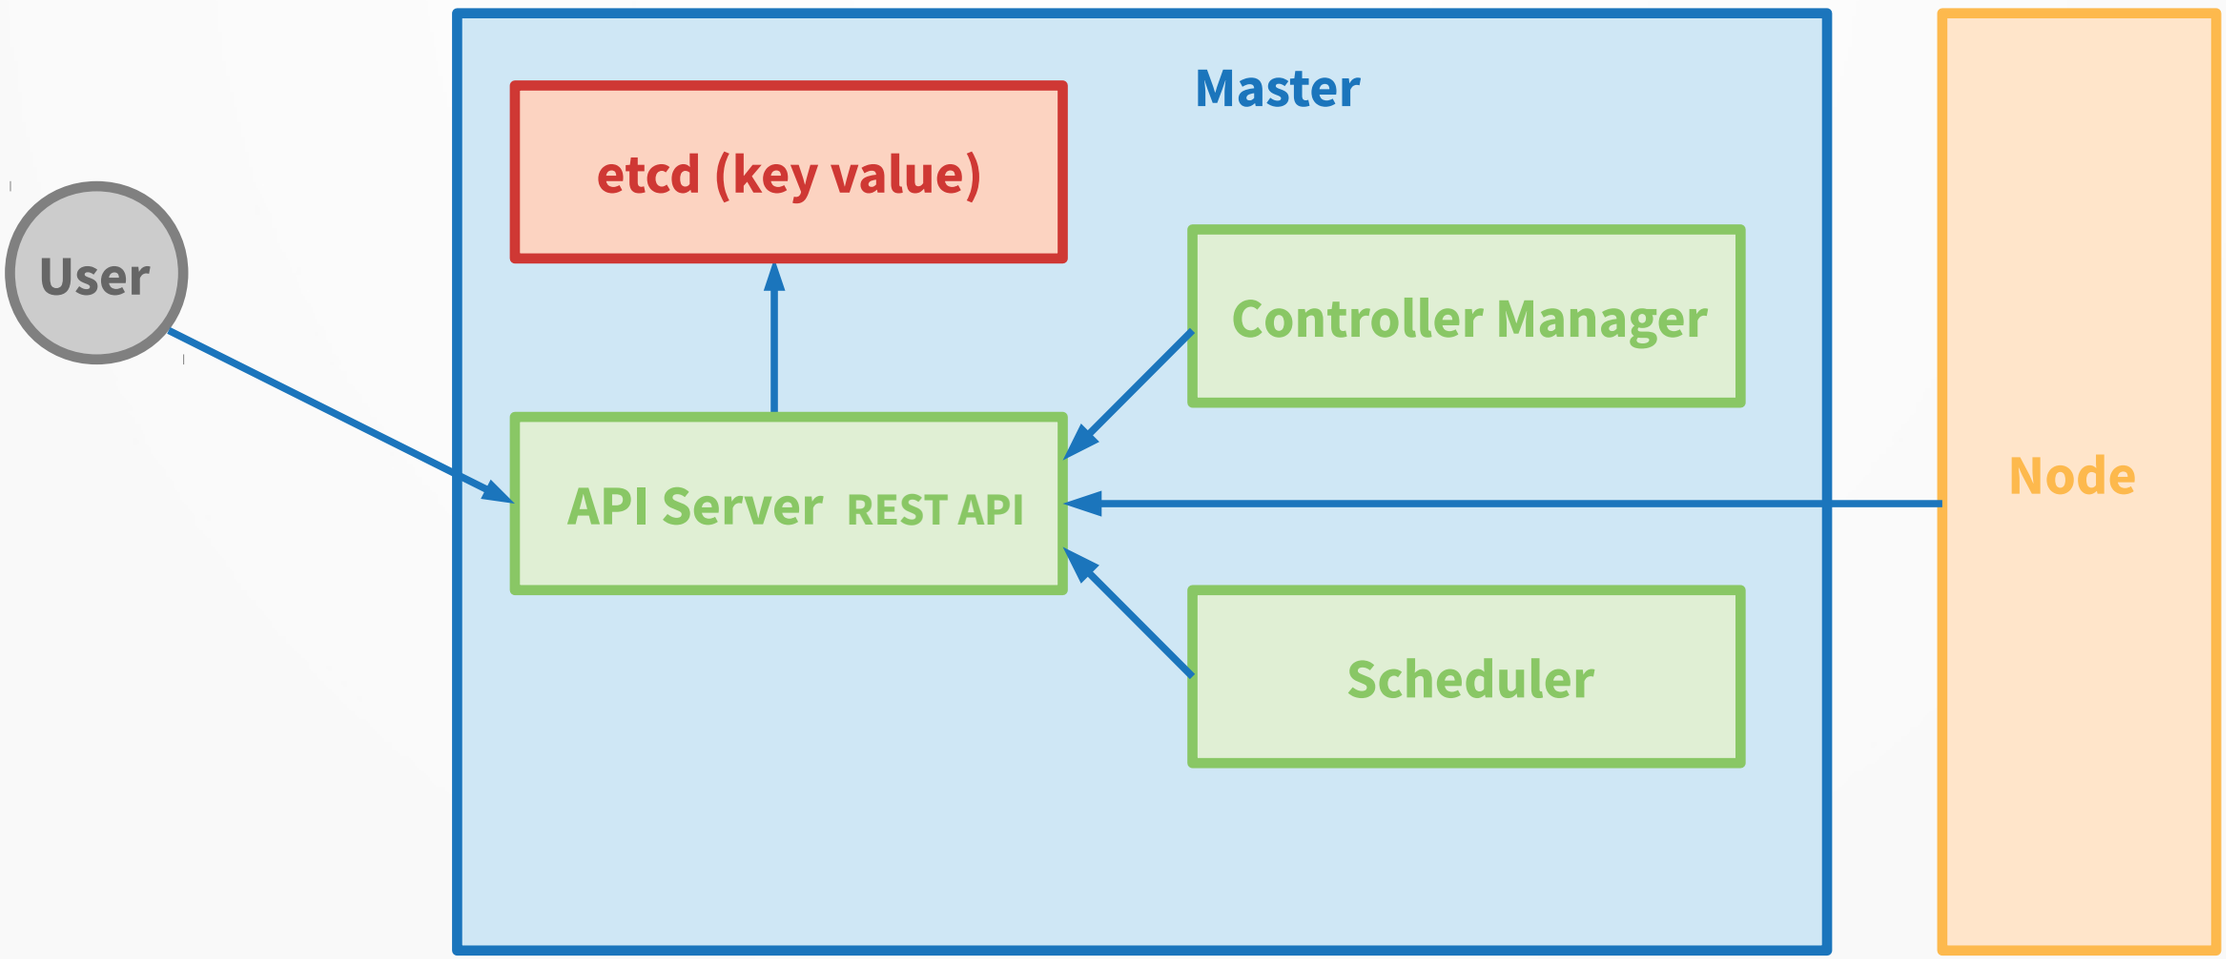

Kubernetes Architecture

etcd

- distributed, consistent key-value store used for configuration management, service discovery, and coordinating distributed work.
- etcd reliably stores the configuration data of the Kubernetes cluster, representing the state of the cluster.
- etcd is written in Go and uses the Raft consensus algorithm to manage a highly-available replicated log.
- http://play.etcd.io/play
- http://kanaka.github.io/raft.js/

Kubernetes Architecture

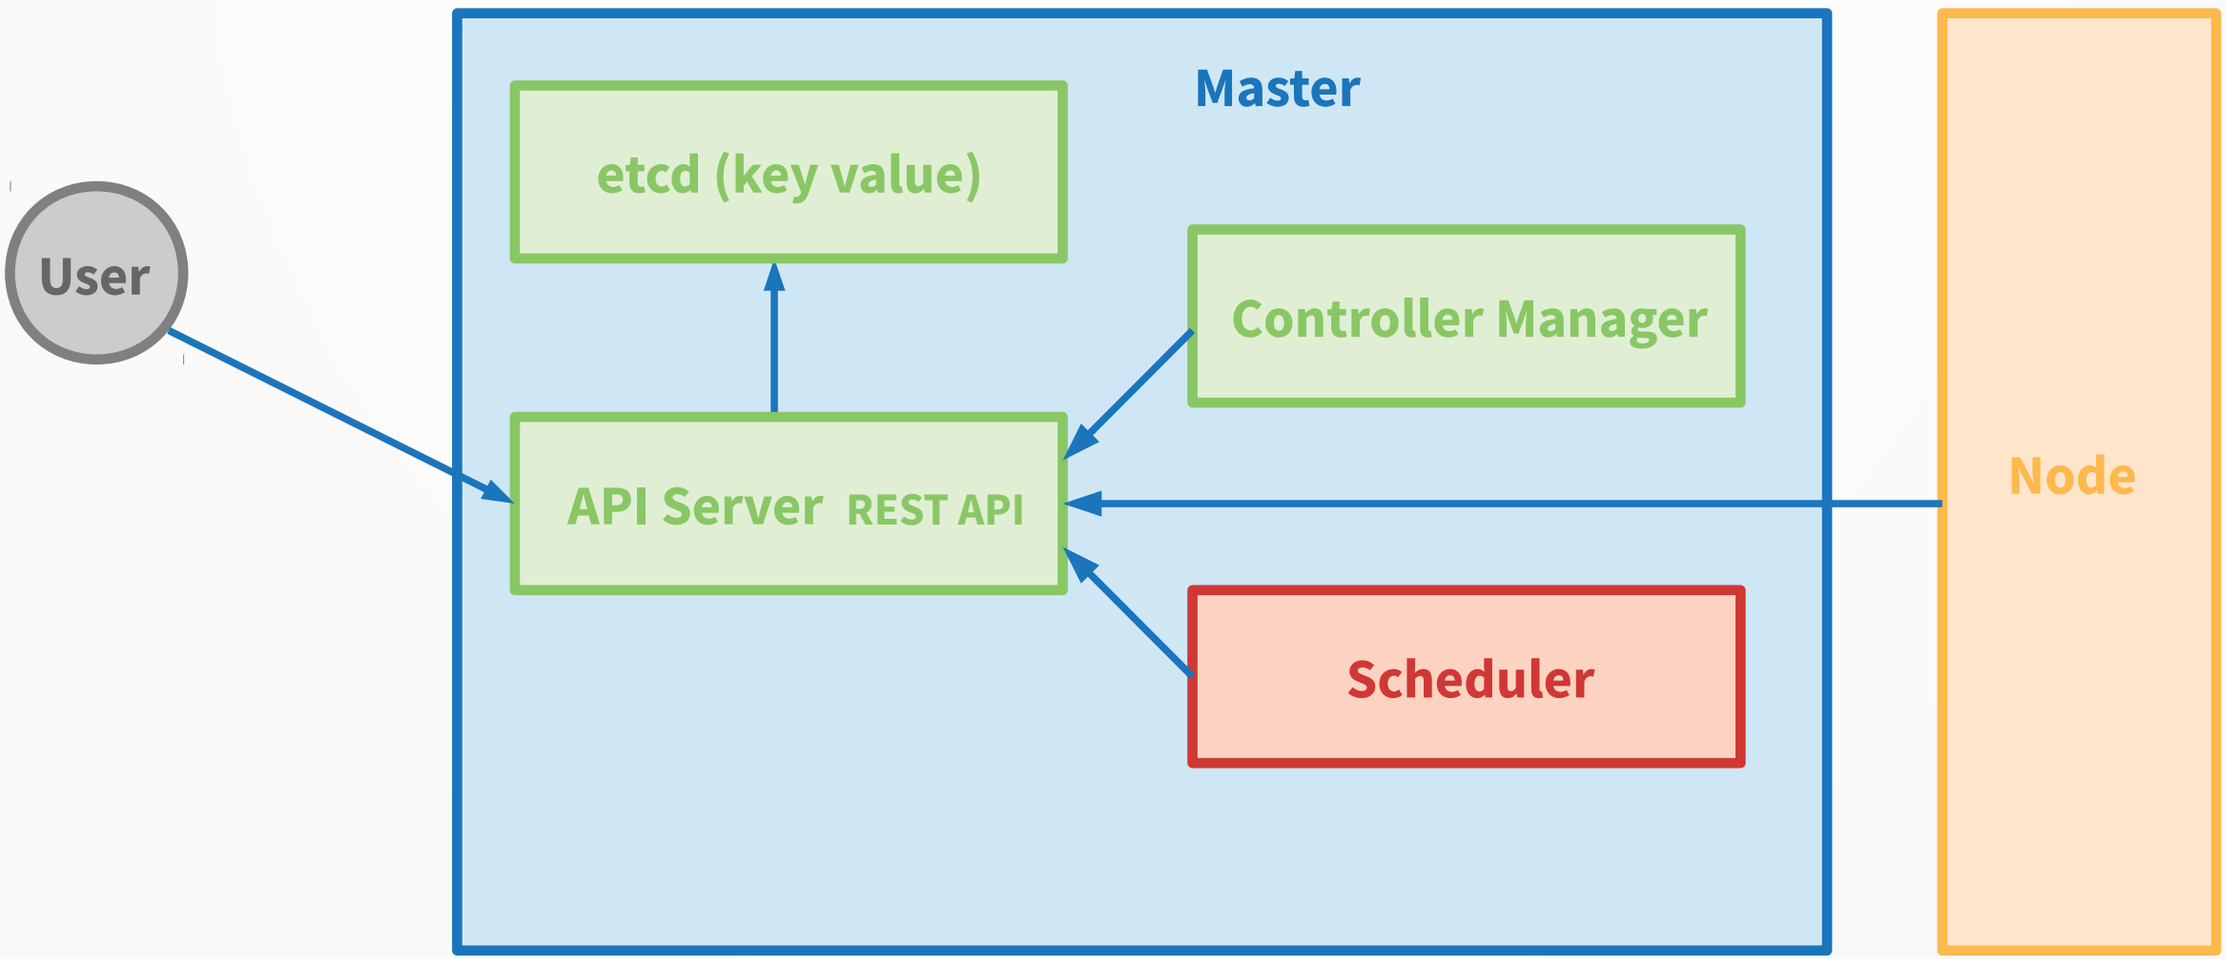

Kubernetes Architecture

Scheduler

- The Scheduler watches for unscheduled pods and binds them to nodes.
- The Scheduler considers the availability of the requested resources, quality of service requirements and other constraints.
- The scheduler watches the state in etcd over the API Server and writes the changes back into etcd.

Kubernetes Architecture

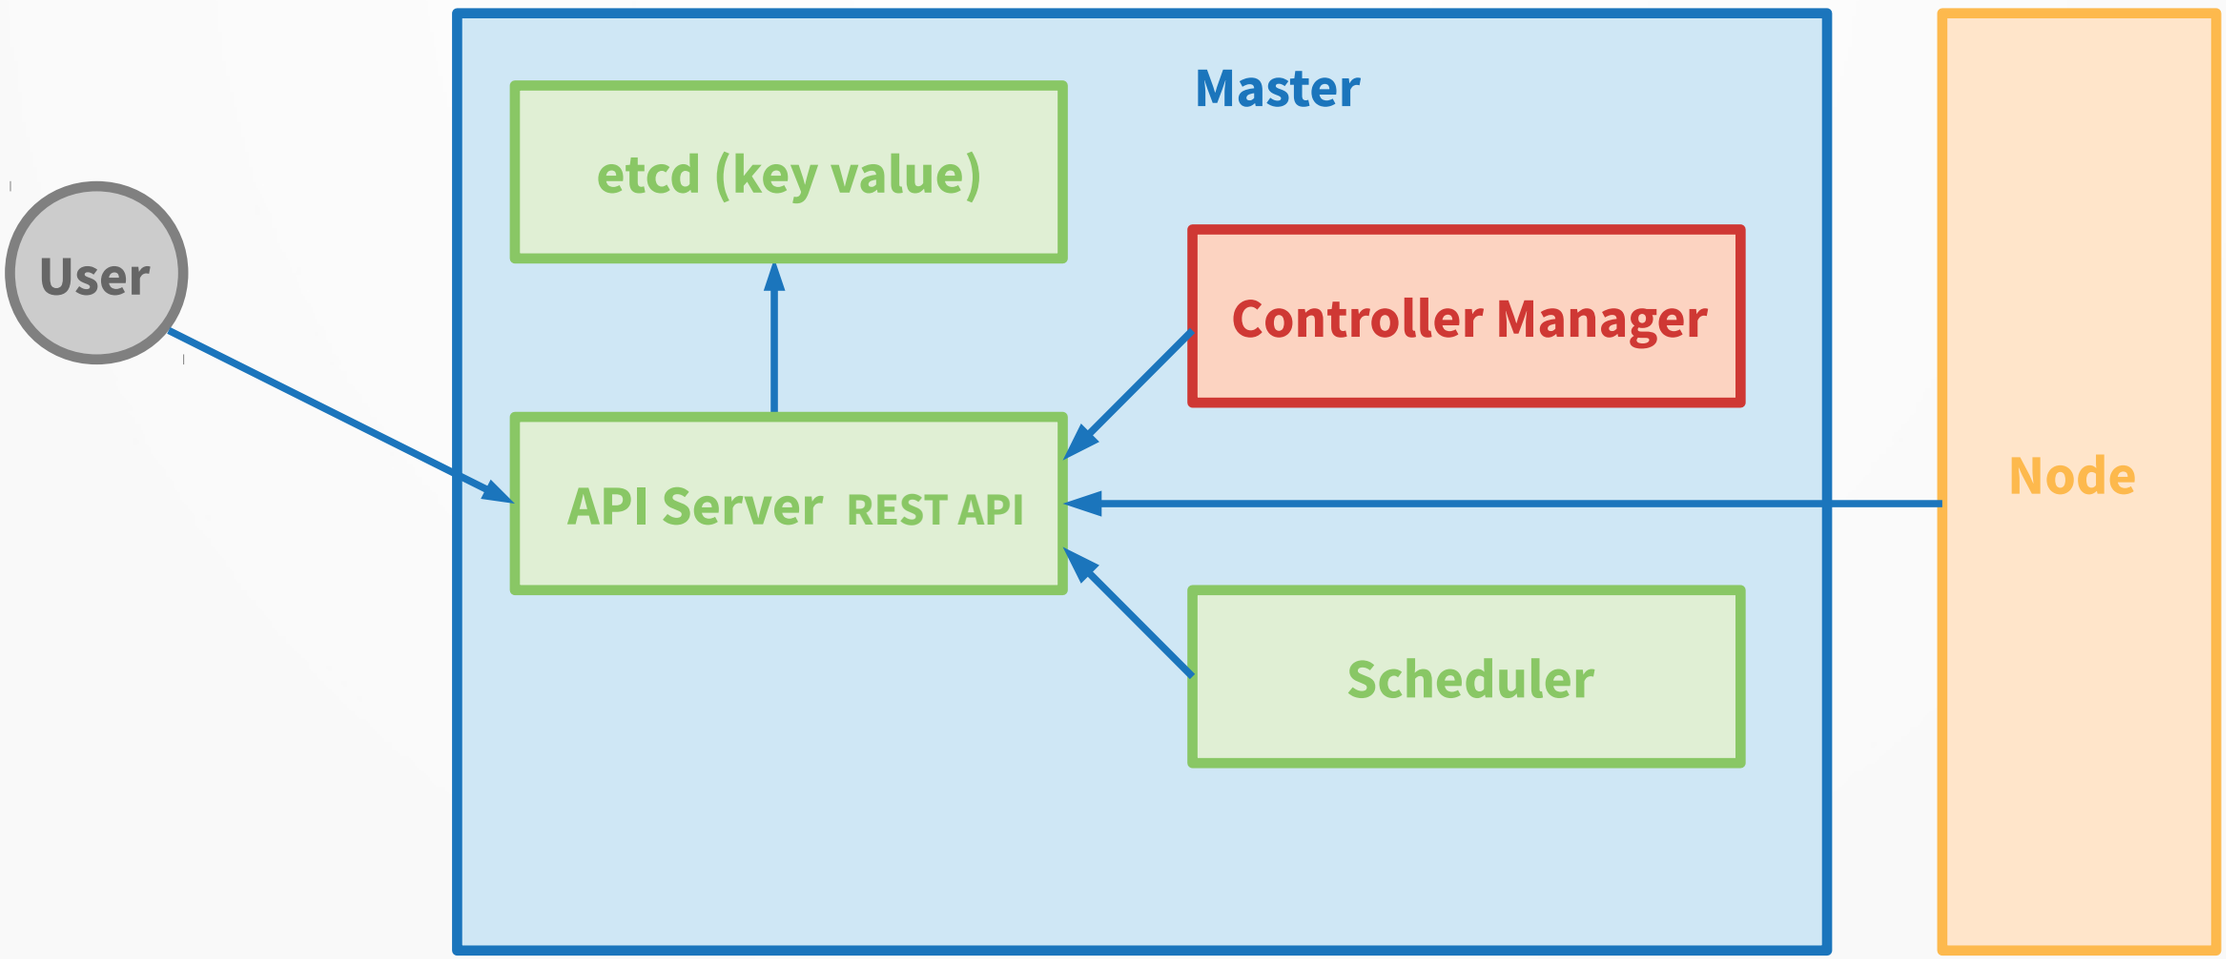

Kubernetes Architecture

Controller Manager

- Similar to the scheduler, but applies more sophisticated concepts.
- core control loops
- In Kubernetes, a controller watches the shared state of the cluster through the apiserver and makes changes to reach the desired state.

Kubernetes Architecture

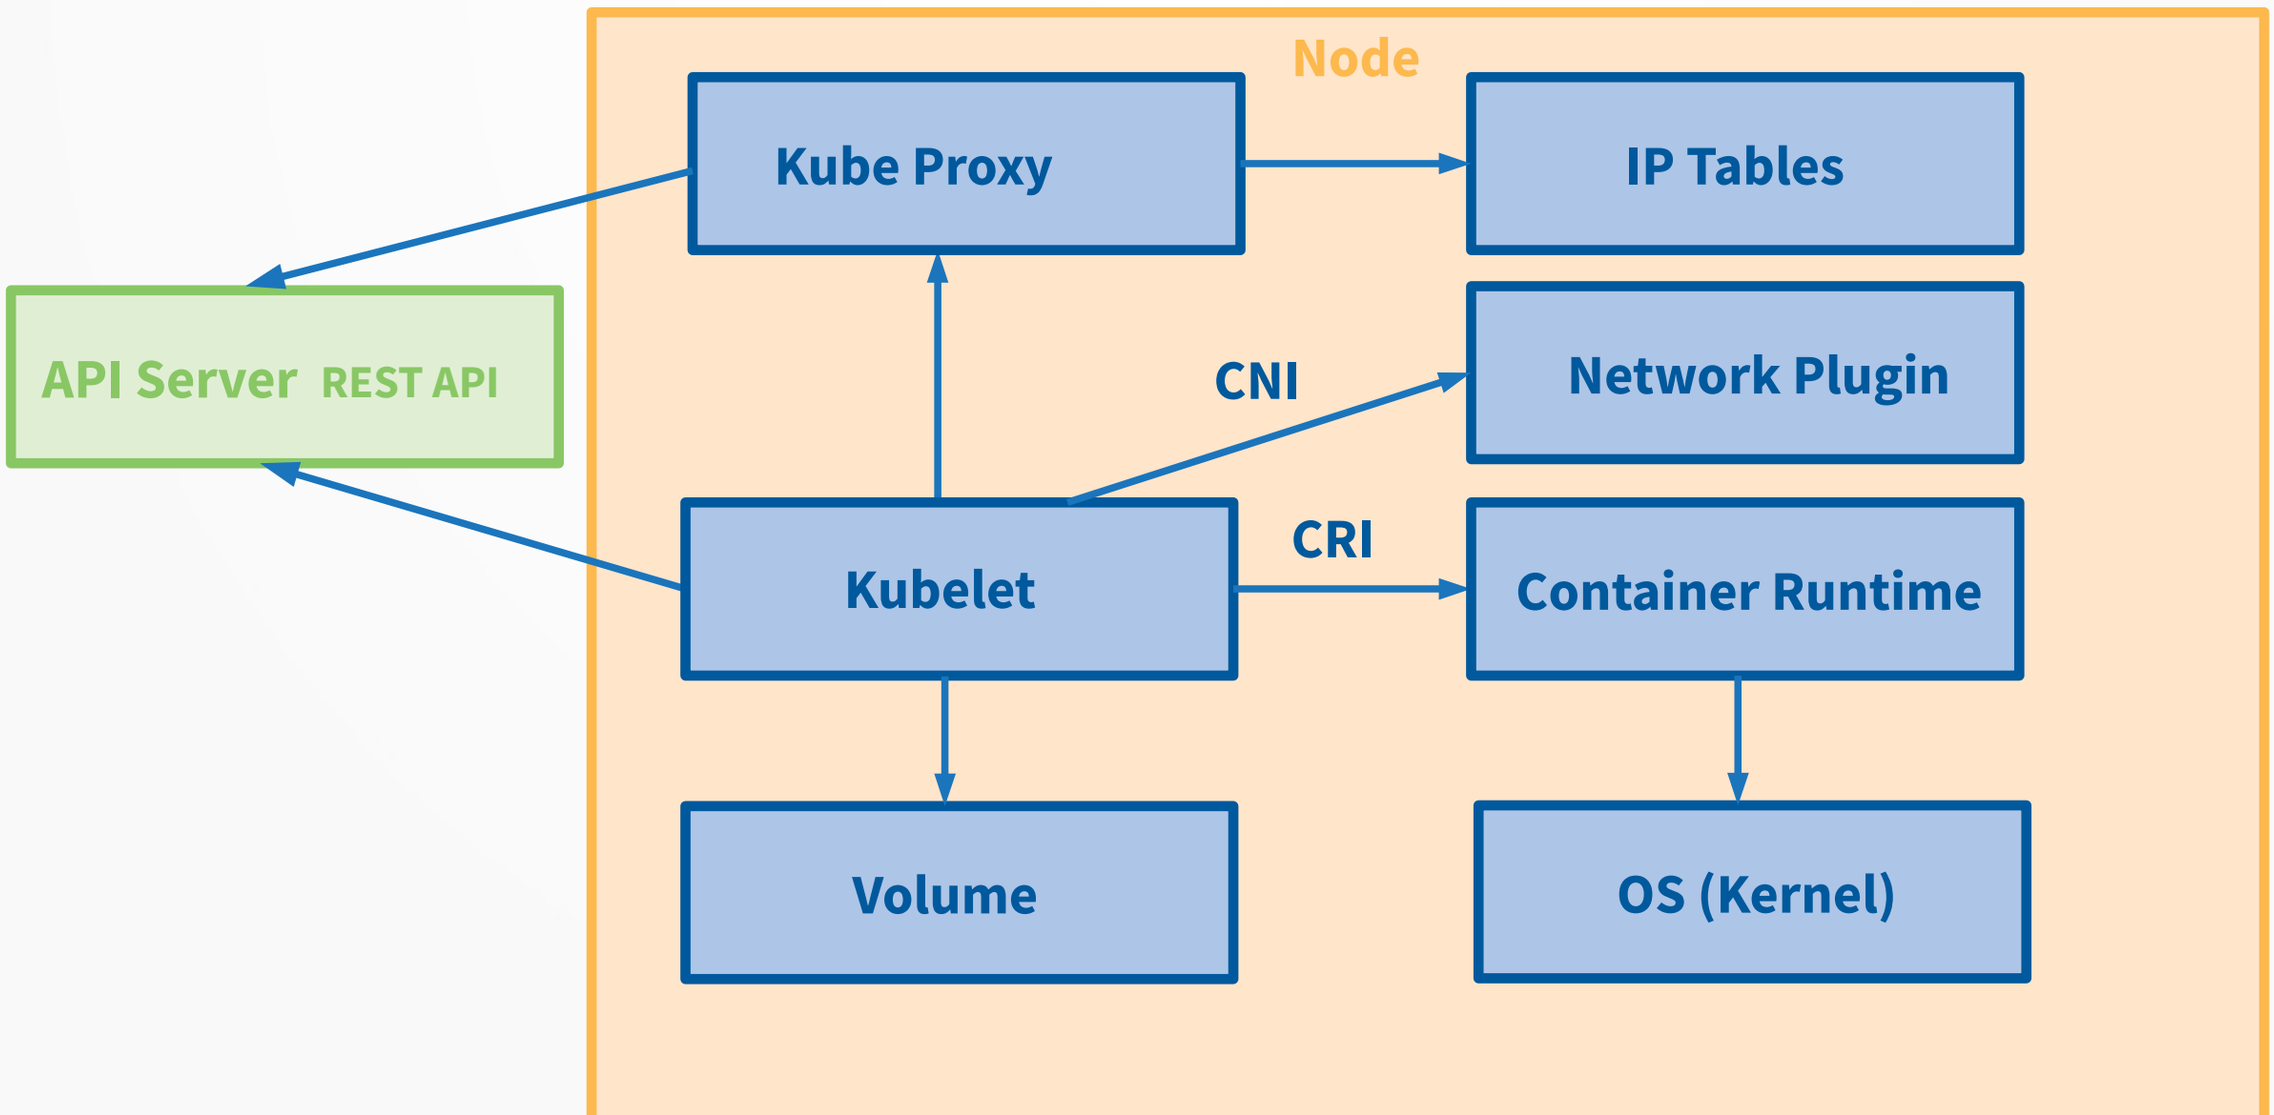

Kubernetes Architecture

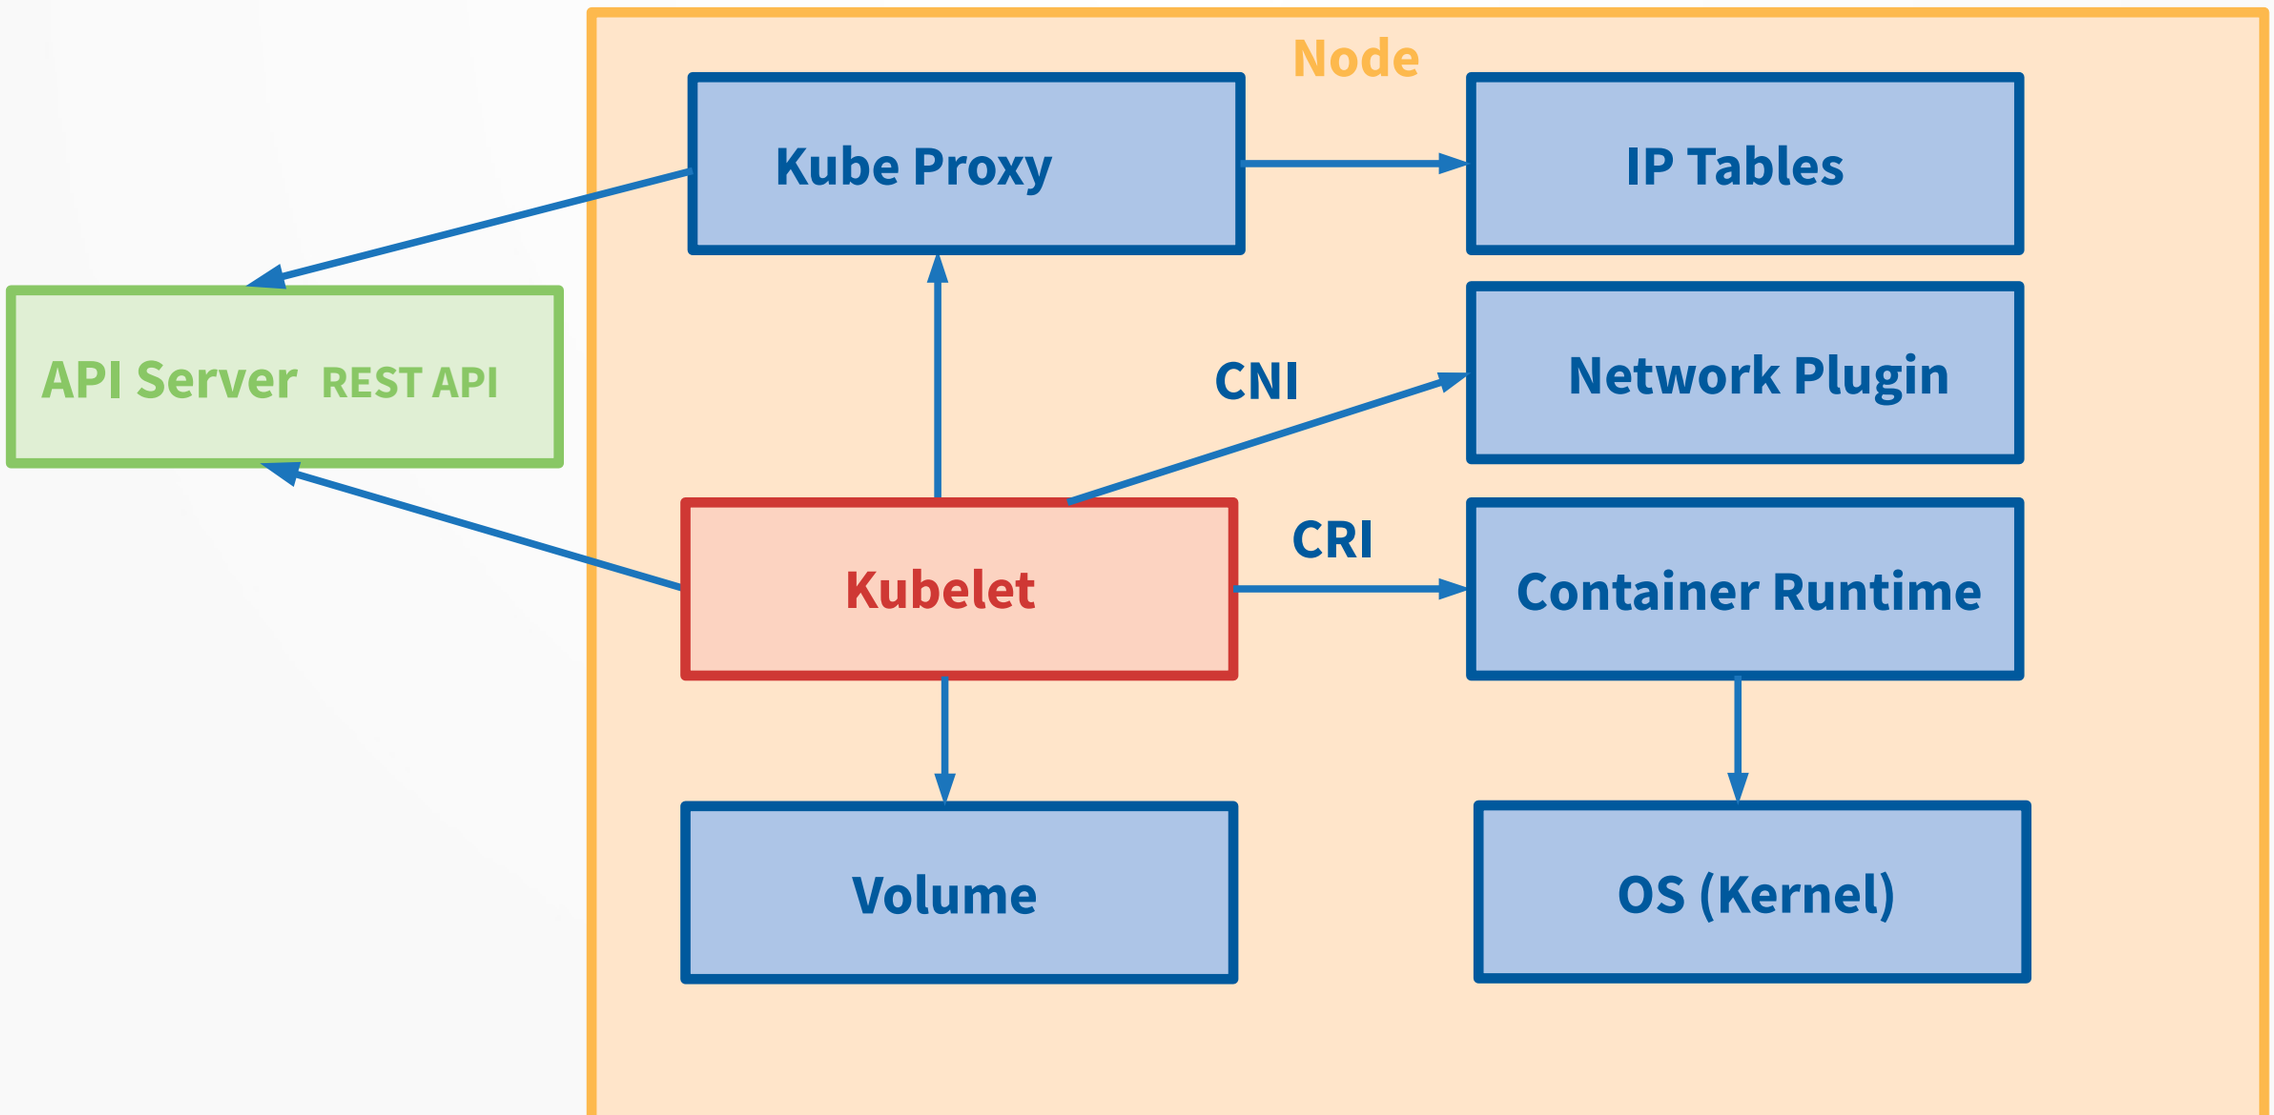

Kubernetes Architecture

Kubelet

- Watching the API Server for changes on pods, that are bound to its node.
- Report the status of the node and each pod to the API Server.
- Run container (pod) probes.

Kubernetes Architecture

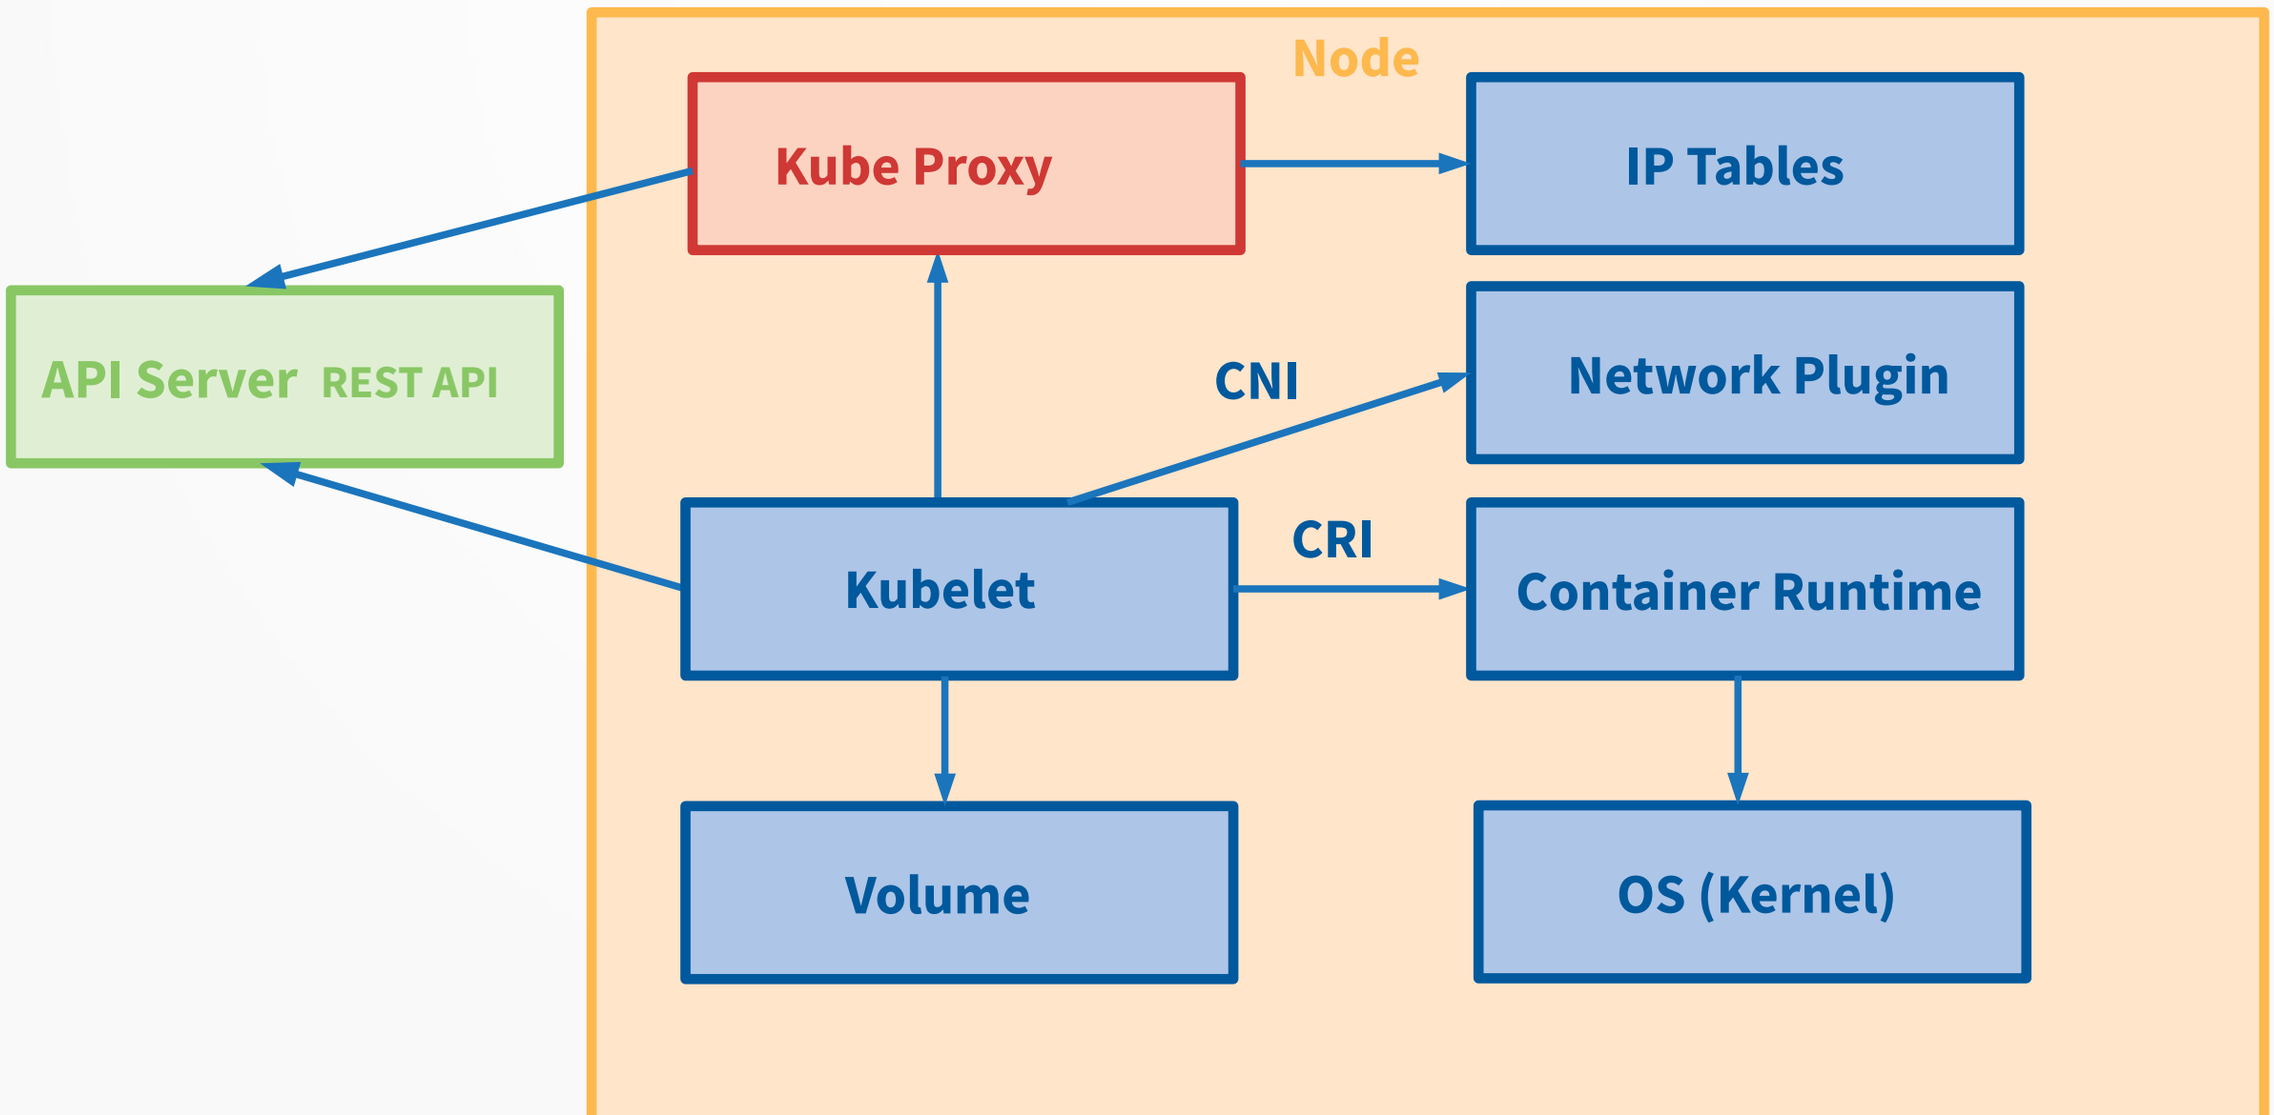

Kubernetes Architecture

Kube Proxy

- Exposes Kubernetes **services** and manipulates iptables rules to trap access to service IPs and redirect them to the correct pods (containers).
- Creates the iptables rules for loadbalancing (service → pod)
- Watching the API Server for changes on services and pods definitions to maintain the network configuration up to date.

Kubernetes Architecture

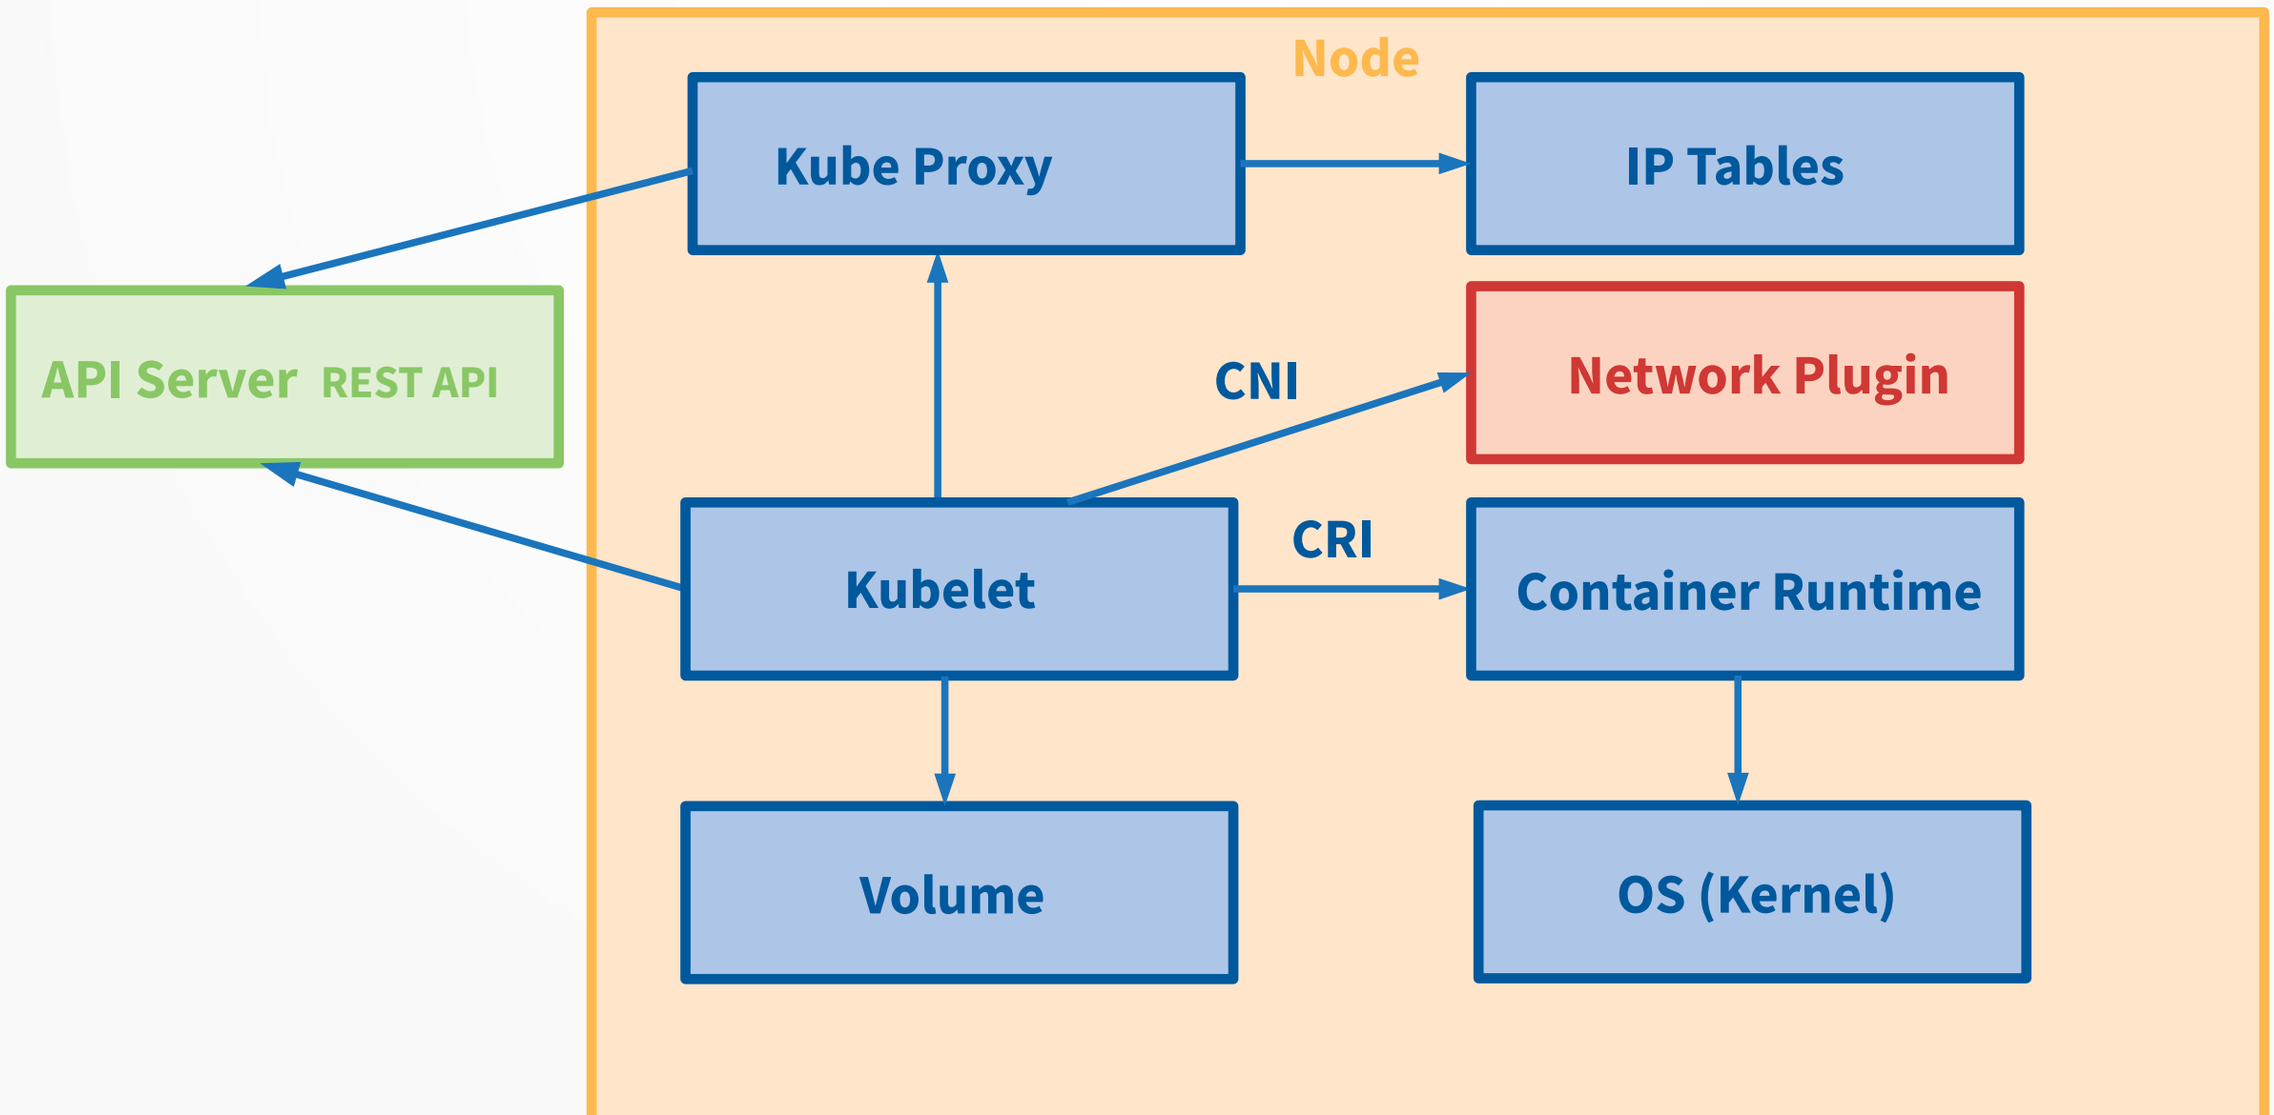

Kubernetes Architecture

Network Plugin (Overlay Network)

- Provides IPs for the Pods and connects them over an overlay network.
- Keep a stable mapping of nodes to subnets and keep every node in your cluster updated with that mapping. Do the right thing when nodes are added & removed.

Kubernetes Architecture

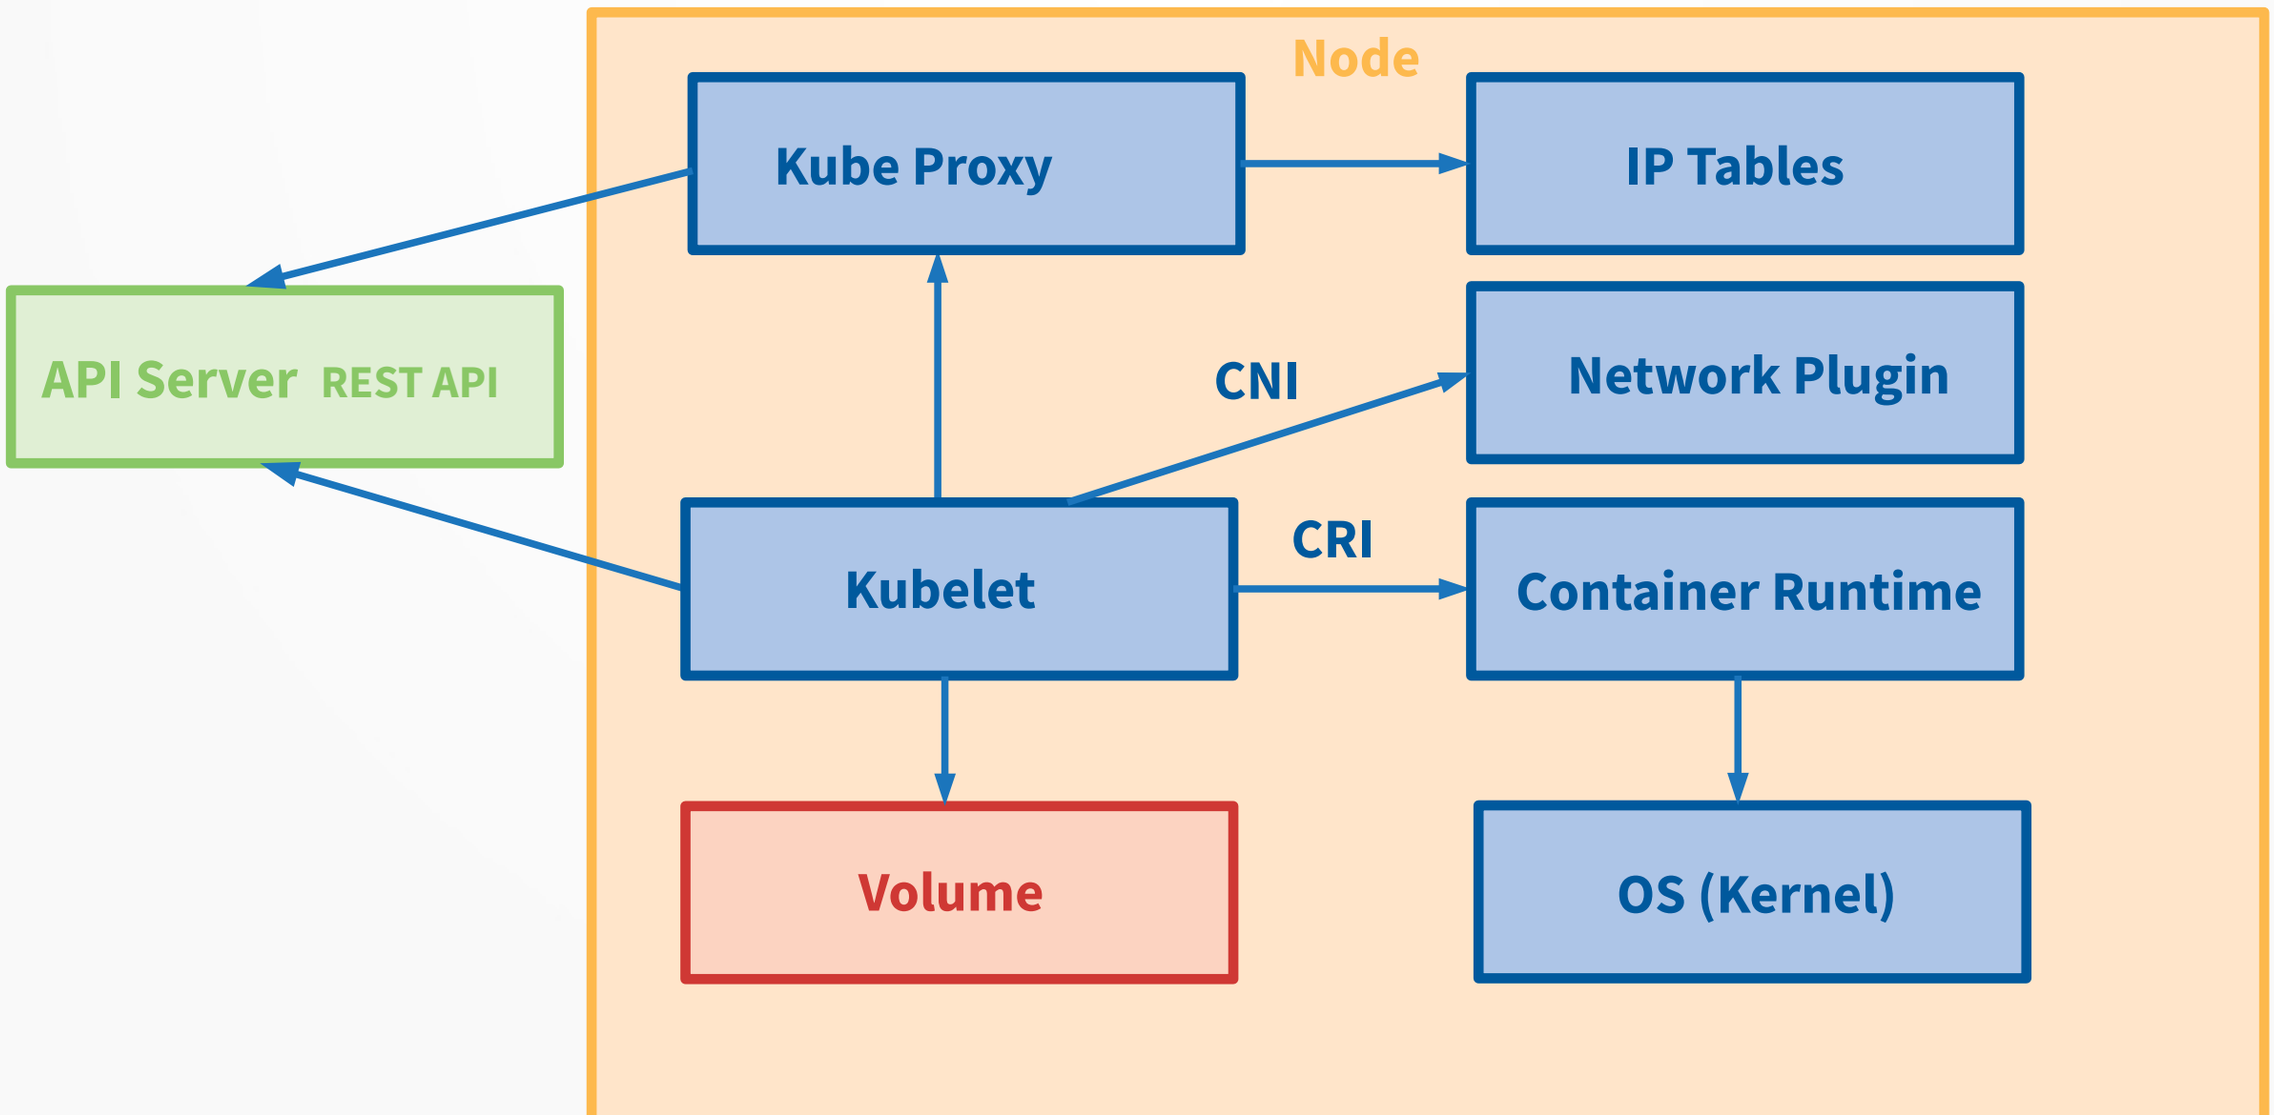

Kubernetes Architecture

Volume

- A directory accessible to all containers running in a pod.
- The data in volumes is preserved across container restarts.

Kubernetes Architecture

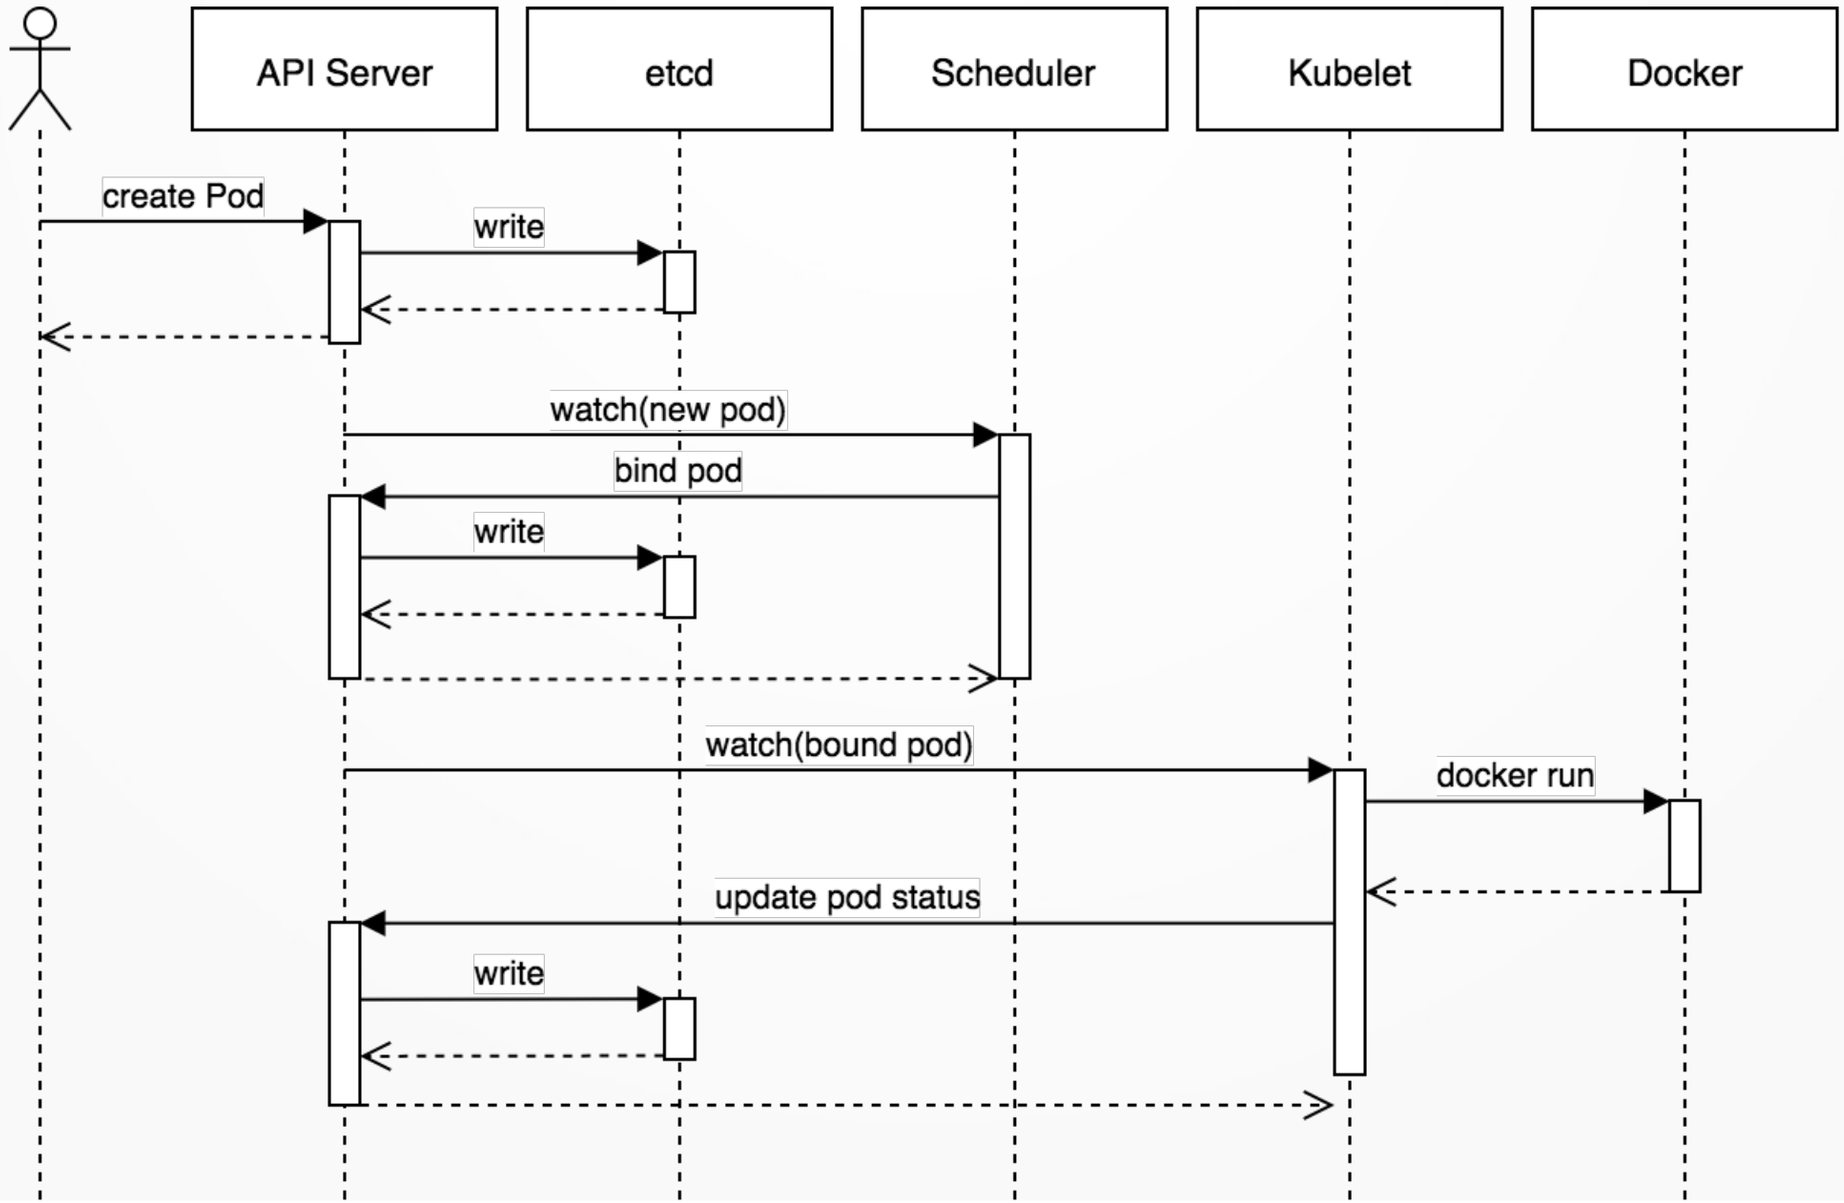

Kubernetes Setup

No «One Way» to setup Kubernetes

- Different scaling / HA models.
- Components (services) can be exchanged by alternative suppliers.
- Security setup.
- Components have a bazillion of configuration options.

Kubernetes Setup

Kubernetes the Hardway

- https://github.com/kelseyhightower/kubernetes-the-hard-way

Kubernetes Setup

This example setup

Kubernetes with HA and Security

Kubernetes Setup

PKI Infrastructure

- etcd, kube-apiserver, kube-controller-manager, kube-scheduler, kubelet, and kube-proxy need TLS Certificates and keys.

Kubernetes Setup

etcd

- 3 node cluster
- Failover is done by the clients
- Traffic encrypted with 2 way TLS

Kubernetes Setup

Bootstrapping the Kubernetes Control Plane

- Kubernetes API Server, Scheduler, and Controller Manager
- Exposes the API over the API Server
- Needs external Loadbalancer
- 3 compute instances for HA
- Install the components as systemd services, could also be installed as docker containers
- Traffic encrypted with 2 way TLS

Kubernetes Setup

Bootstrapping the Kubernetes Worker Nodes

- containerd, runc, gVisor, container networking plugins, kubelet, and kube-proxy.
- **containerd**: an industry-standard container runtime with an emphasis on simplicity, robustness and portability. Open sourced by docker.
- **runc**: tool for spawning and running containers according to the OCI specification. Used by containerd. Open sourced by docker.
- **gVisor**: provides an isolation boundary between the application and the host kernel.
- Point to loadbalancer to reach the API Server
- Traffic encrypted with 2 way TLS

Kubernetes Setup

Provisioning Pod Network Routes

- Pods scheduled to a node receive an IP address from the node's pod CIDR range.
- At this point pods can not communicate with other pods running on different nodes due to missing network routes..
- Create network routes for each worker instance.
- There are other ways to implement the Kubernetes networking model.
  A popular one is canal.

Kubernetes Setup

Deploying the DNS Cluster Add-on

- Deploy the DNS add-on which provides DNS based service discovery to applications running inside the Kubernetes cluster.
- `kubectl create -f https://storage.googleapis.com/kubernetes-the-hard-way/kube-dns.yaml`

Kubernetes Setup

Deploying the DNS Cluster Add-on

- Deploy the DNS add-on which provides DNS based service discovery to applications running inside the Kubernetes cluster.
- `kubectl create -f https://storage.googleapis.com/kubernetes-the-hard-way/kube-dns.yaml`

Kubernetes Example

https://github.com/floriankammermann/kubernetes-presentation/tree/master/example/kube-config

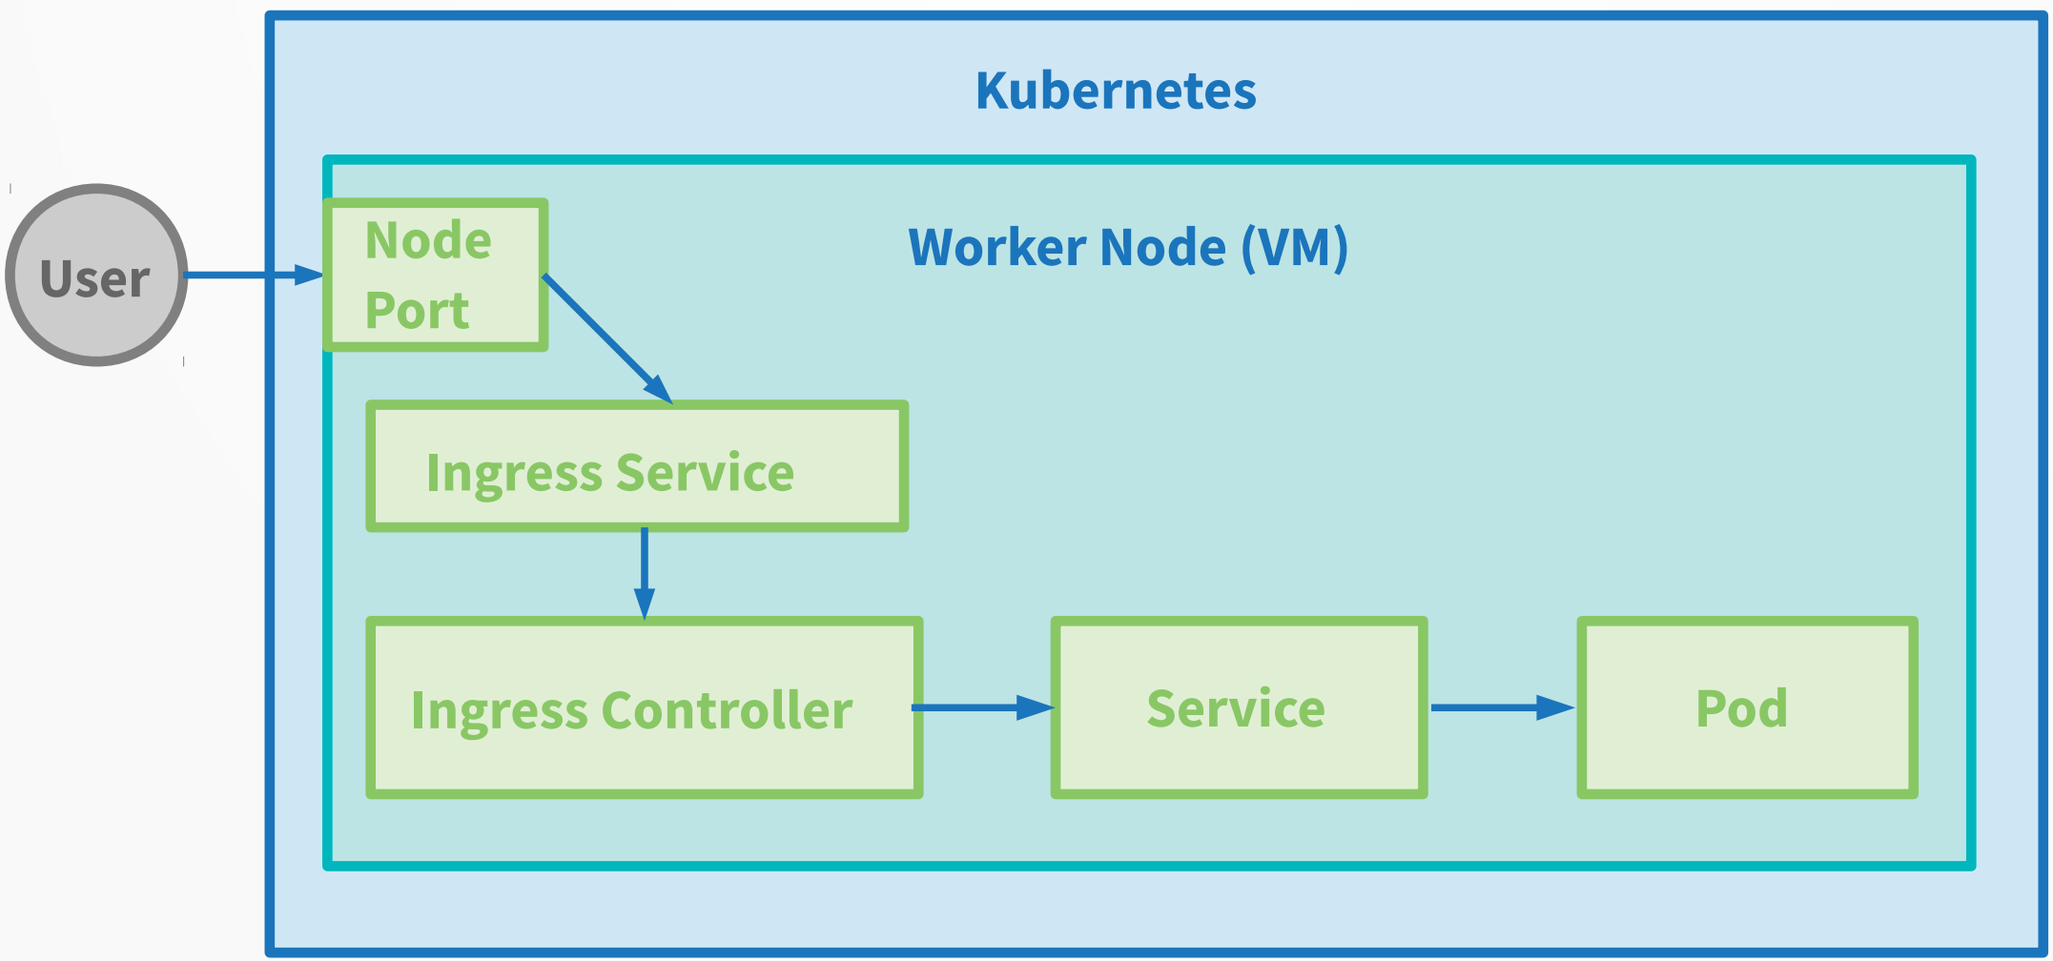

Kubernetes Conclusion

Kubernetes Benefits

- One consistent API
- Create Infrastracture Objects in seconds
- Problems solved on a higher infrastracture level,
  somewhere between Openstack (IaaS) and Cloudfoundry (PaaS)
- Custom Controller
- Enabler for GitOps

Kubernetes Conclusion

Kubernetes Benefits

- Backed by Cloud Native Computing Foundation
- Fast paced development, lots of Special Interest Groups
- More and more infrastracture problems get solved
  (storage, network, security, encryption)
- Interesting frameworks on top of kuberenetes like istio, conduit