openstack.

November 2019

OpenStack Provider for Virtual Kubelet: A Nodeless Approach for Kubernetes

Kevin Zhao
Hongbin Lu

openstack.

Agenda

- **Introduction**
- OpenStack Zun
- Virtual Kubelet
- OpenStack Provider for Virtual Kubelet
- Demo

VMs or Containers?

“The future of Kubernetes is Virtual Machines, not Containers”?

“Kubernetes Will Start to Replace The Hypervisor”?

[1] https://tech.paulcz.net/blog/future-of-kubernetes-is-virtual-machines/

[2] https://chrisshort.net/2018-learnings-2019-expectations/

Kubernetes and OpenStack

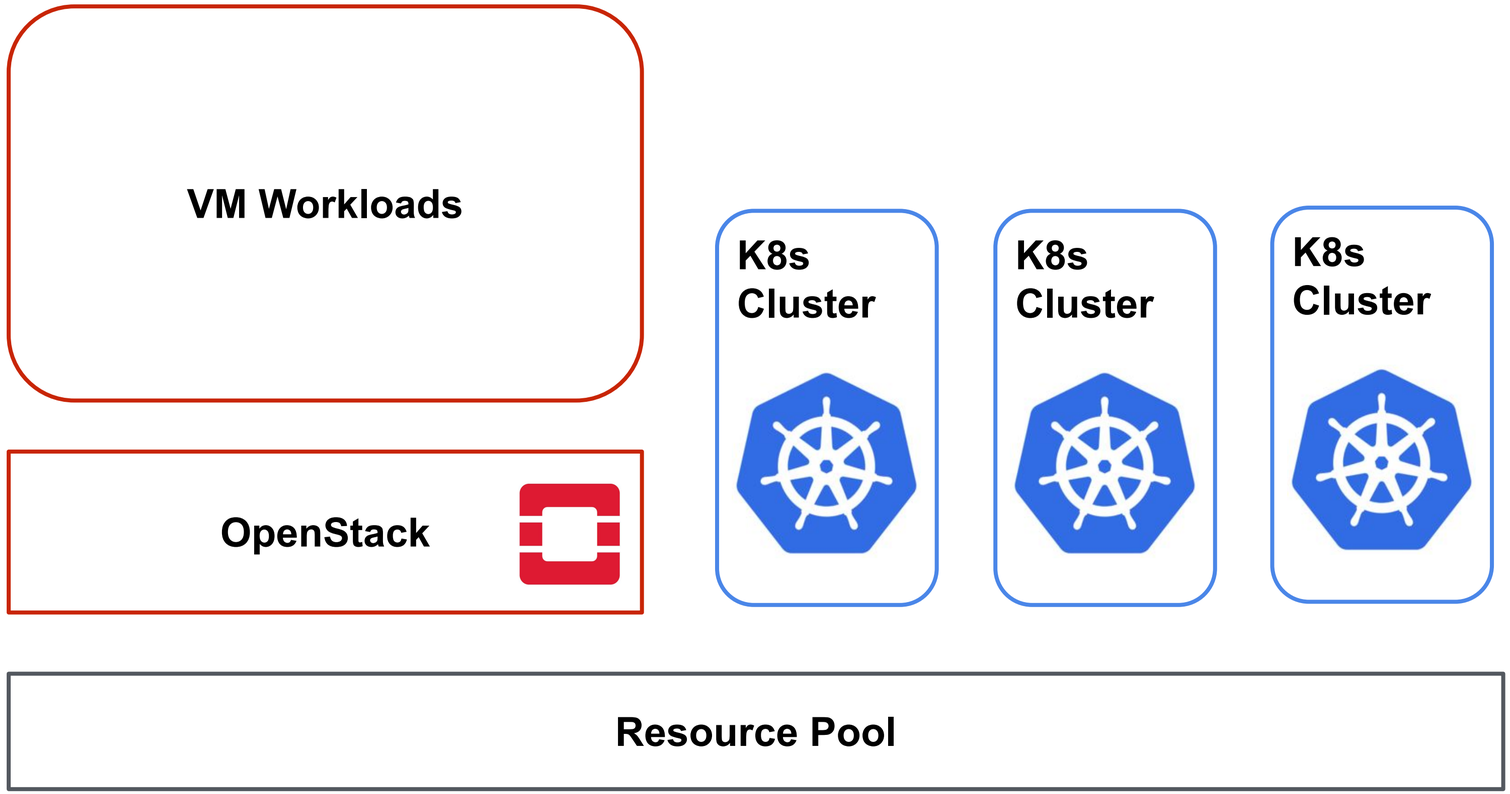

- Kubernetes OpenStack side-by-side
  - Statical partitioning of resource pool
- Kubernetes on OpenStack
  - Virtualization Overhead (or use Ironic)
  - Network
  - Kubesprawl

openstack.

Static Partitioning of Resource Pool

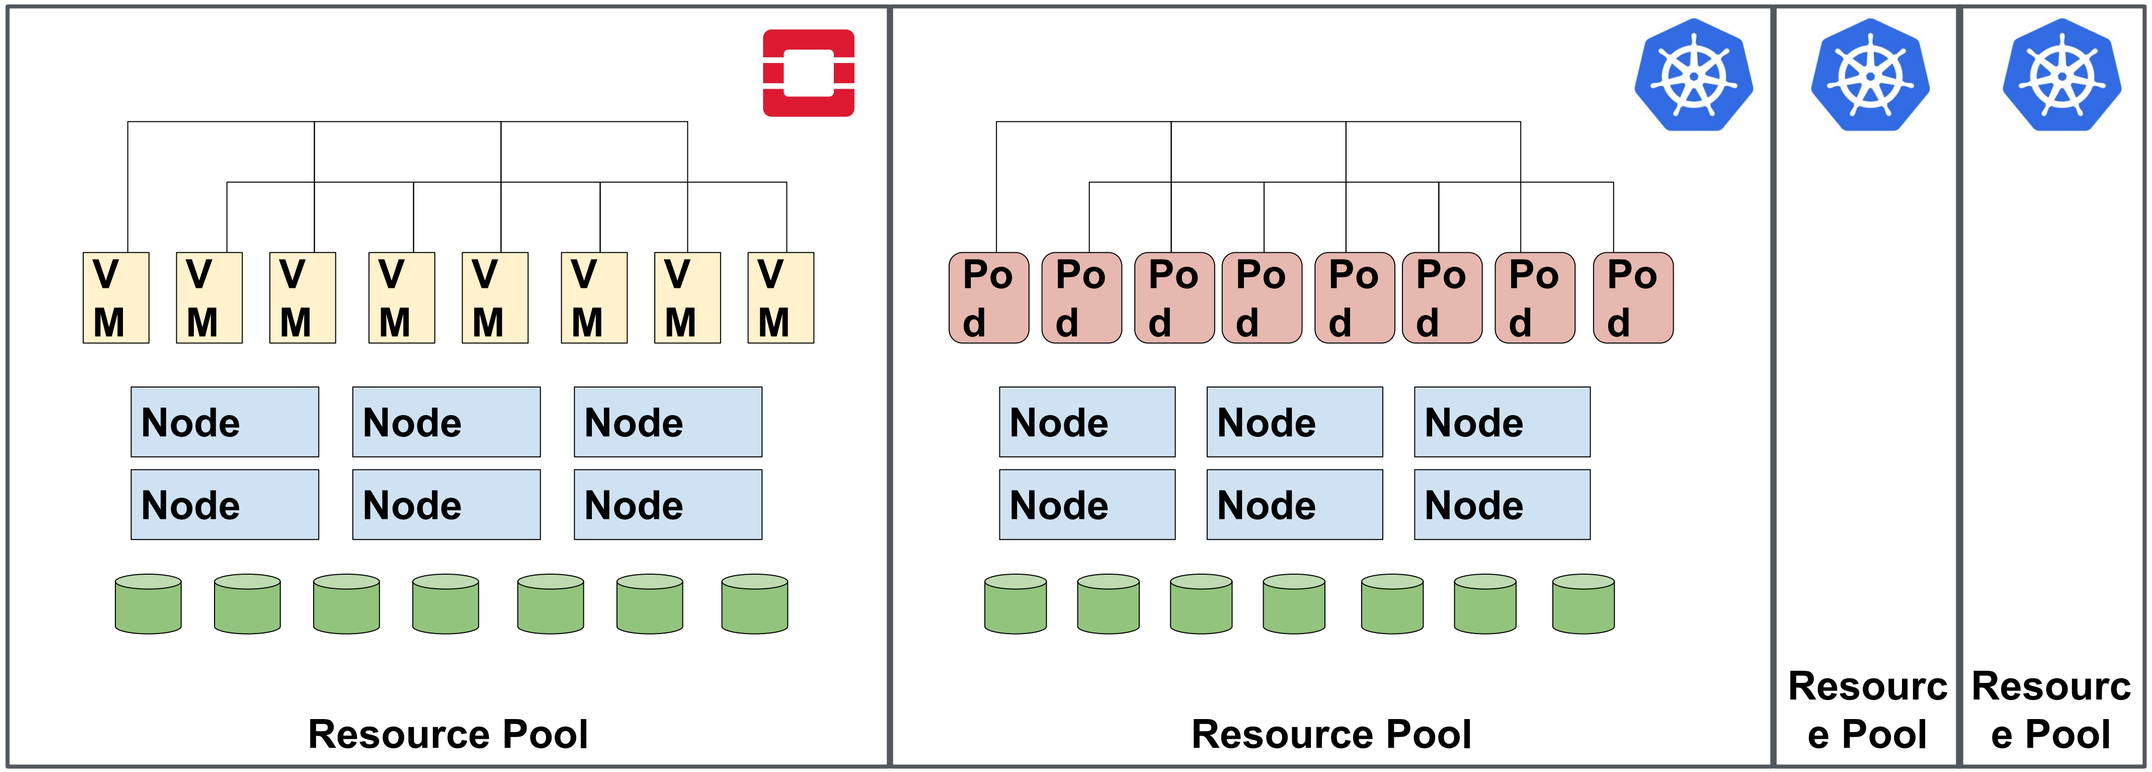

Kubernetes on OpenStack

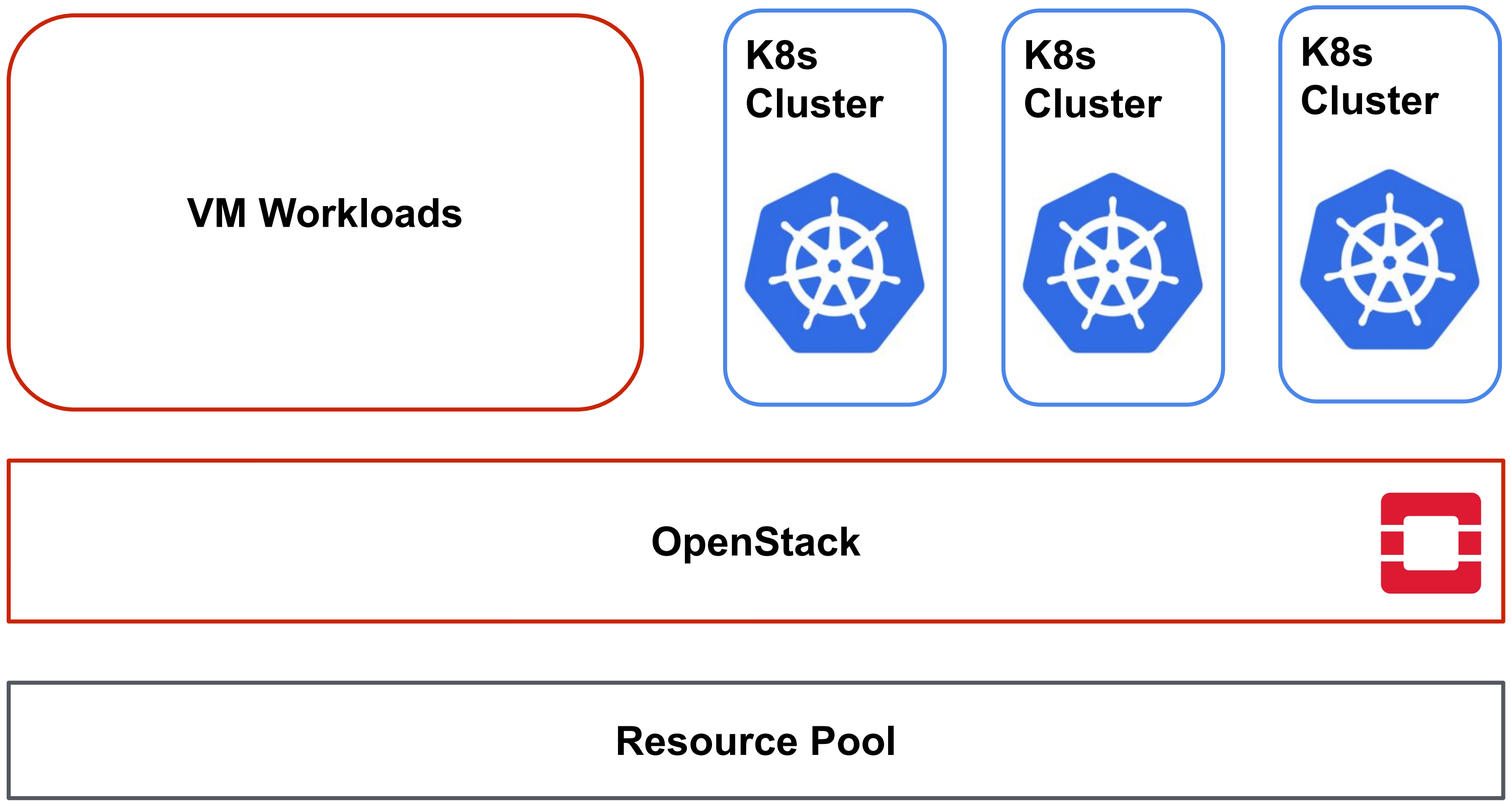

- Kubernetes OpenStack side-by-side
  - Statical partitioning of resource pool
- Kubernetes on OpenStack
  - Virtualization Overhead (or use Ironic)
  - Network
  - Kubesprawl

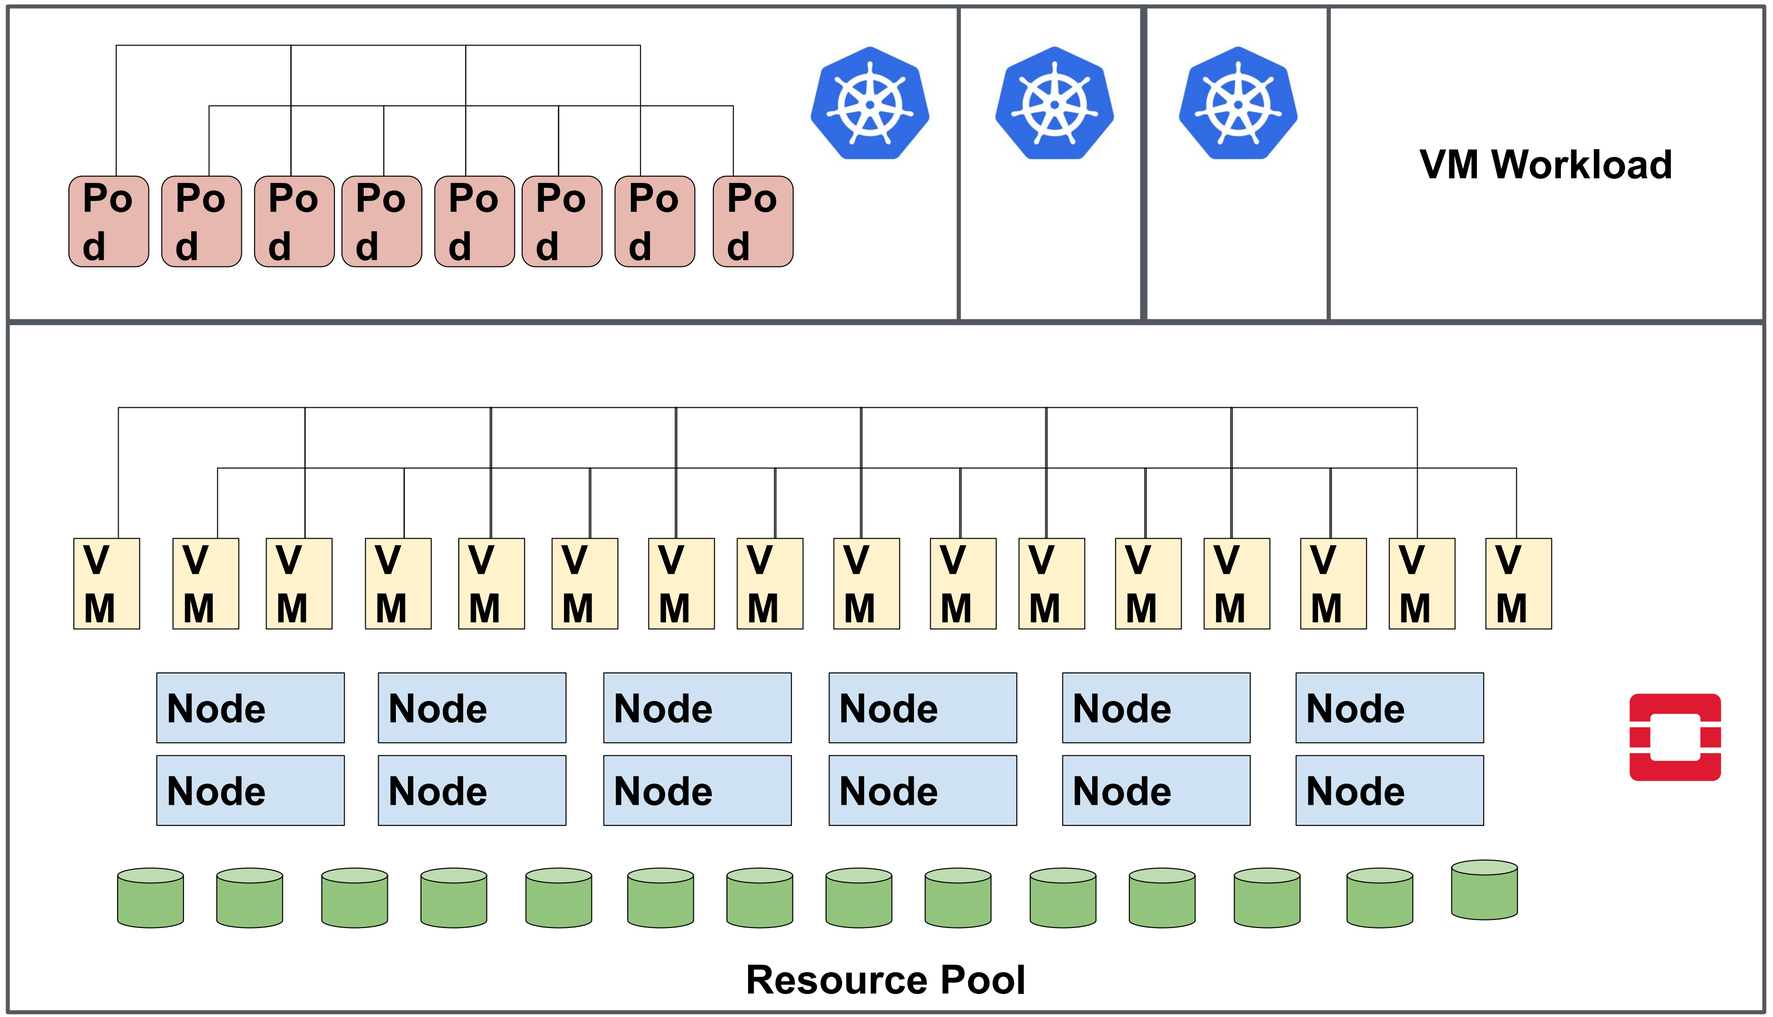
Pod
Pod
Pod
Pod
Pod
Pod
Pod
Pod
VM Workload
VM
VM
VM
VM
VM
VM
VM
VM
VM
VM
VM
VM
VM
VM
VM
Node
Node
Node
Node
Node
Node
Node
Node
Node
Node
Node
Node
Resource Pool

“Your Cloud Utilization Sucks!”

Proposed Architecture

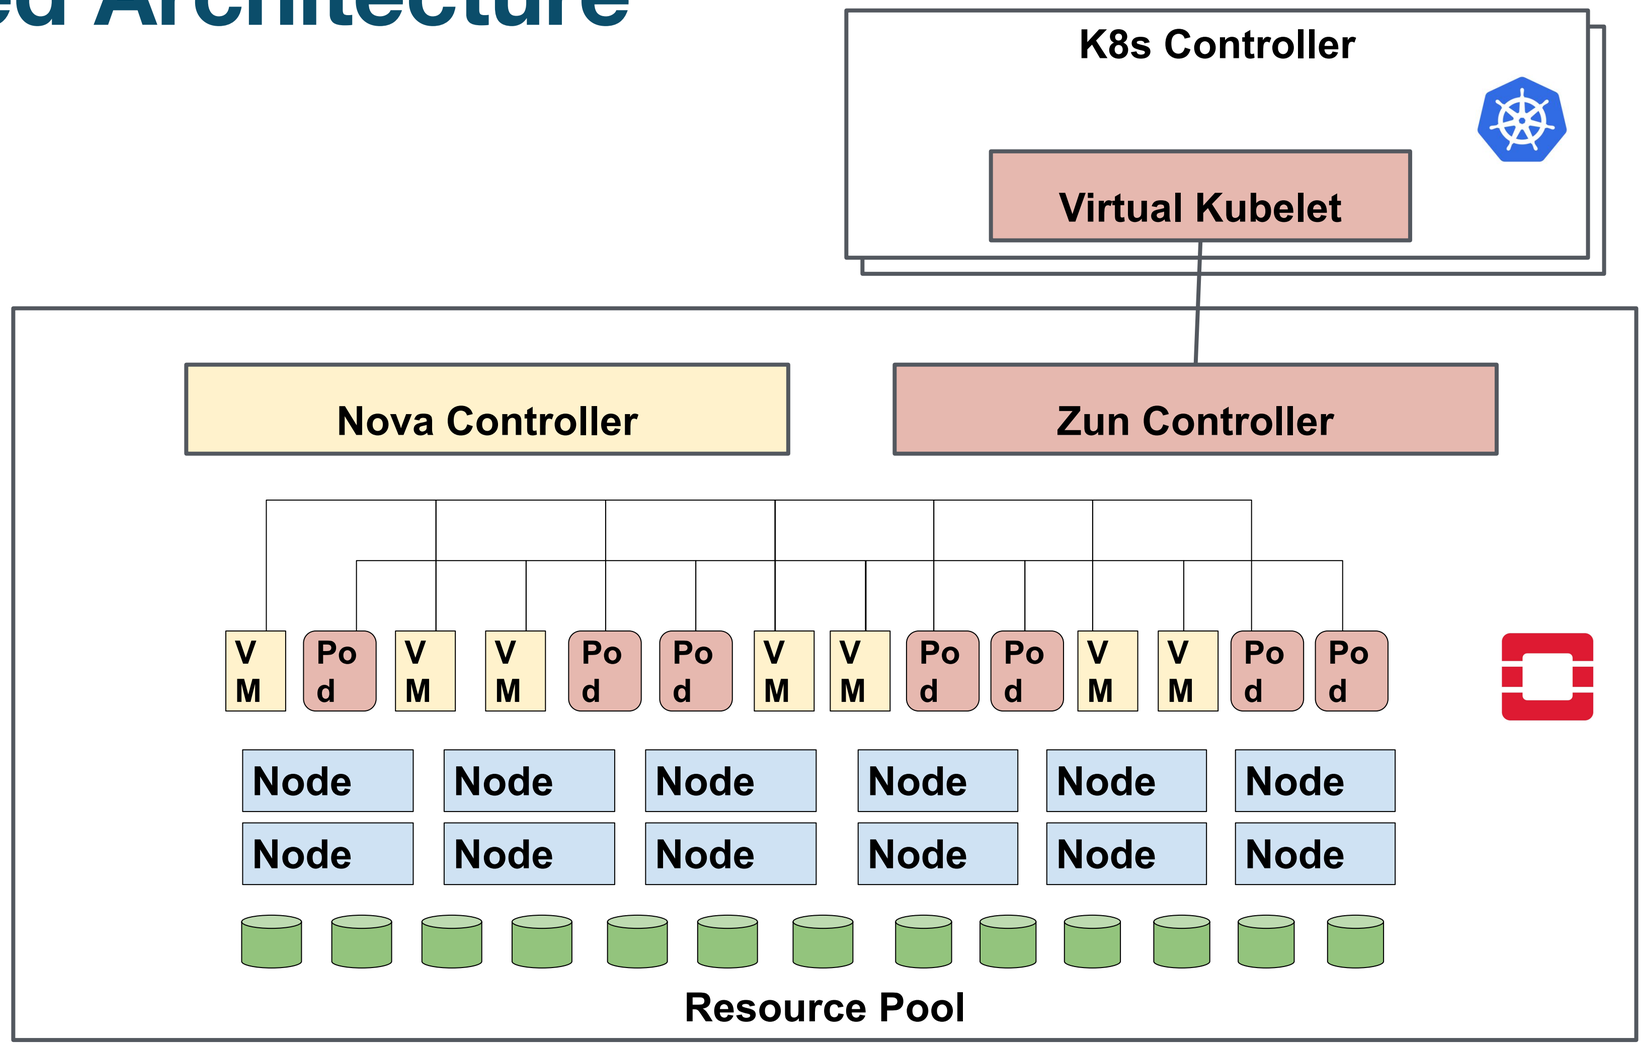

openstack.

Agenda

- ➡️ Introduction
- ➡️ **OpenStack Zun**
- ➡️ Virtual Kubelet
- ➡️ OpenStack Provider for Virtual Kubelet
- ➡️ Demo

Introduce Zun

- OpenStack Container service
- Provide API for provisioning and managing containers without VMs
- Reuse the infrastructure for containers
  - Shared compute nodes
  - Shared Neutron L2 network
  - Shared Cinder block storage pool

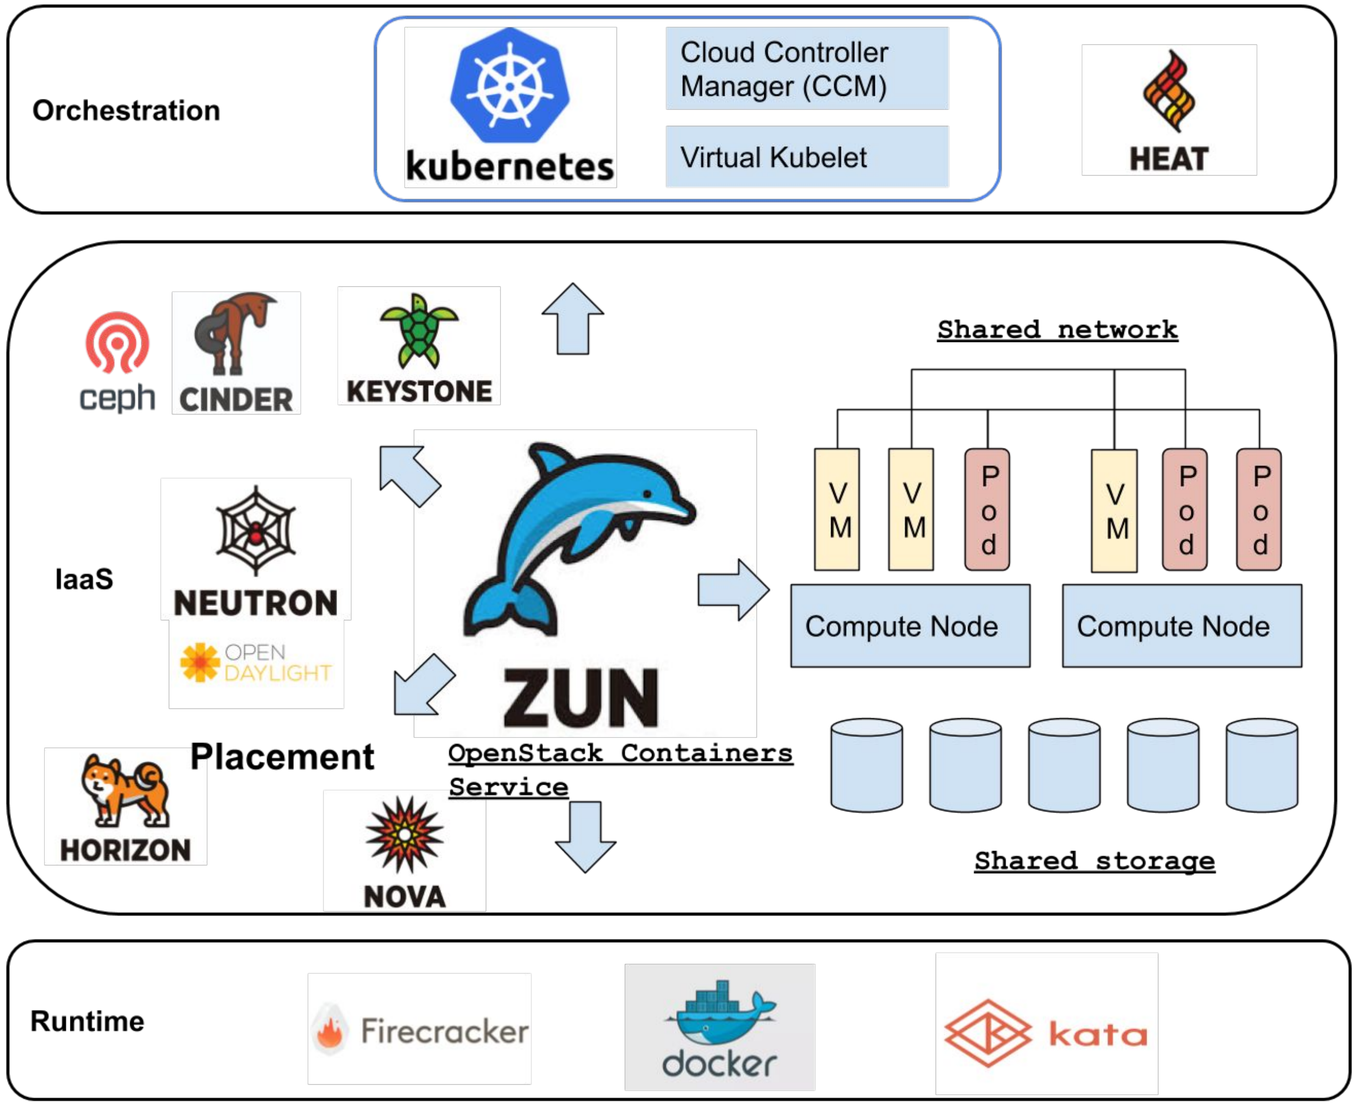

openstack.

Abstraction

➡️ **Container**: A single container
- Properties: image, cpu, memory, disk, command, environment, addresses, security_groups, ...
- Actions: create, update, delete, start, stop, kill, network-attach, exec, attach, log, ...

➡️ **Capsule**: A group of containers that are co-located, have shared network and volumes.
- Properties: containers, cpu, memory, ...
- Actions: create, delete, …

openstack.

Zun Container

➡️ Zun container/pod has a (or more) Neutron port
- Private IP address
- Floating IP address
- Security groups

➡️ Multi-tenancy
- Isolated by Keystone Project
- Kata Container

➡️ Image
- Public Registry (i.e. Docker Hub)
- Private Registry
- Glance (tar archive formate)

Zun Architecture

- Zun API
  - Provide REST APIs
  - Manage all compute nodes
  - Scheduling containers
- Zun Compute
  - Compute node agent
  - Manage local containers
  - Track compute resources

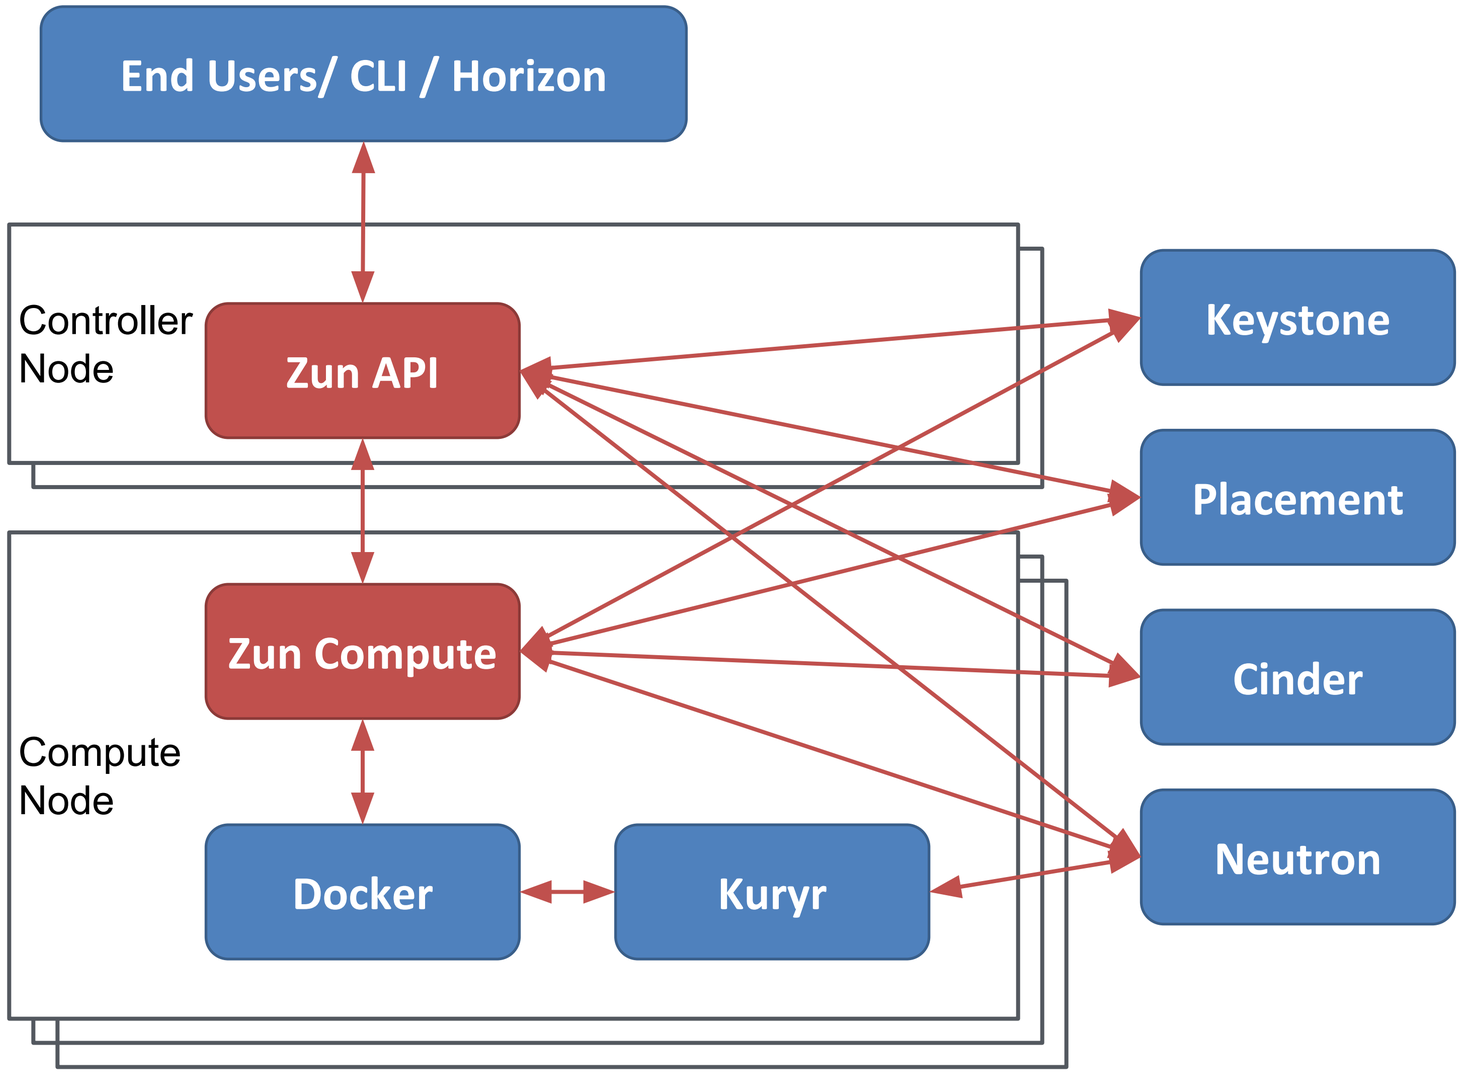

Scheduling

1. Report local resources (i.e. how much CPUs, memory, etc.)
2. Create a container
3. Get allocation candidates
4. Select a compute host (filters and weighters)
5. Claim resource allocation for the container.
6. Run the container in the compute host

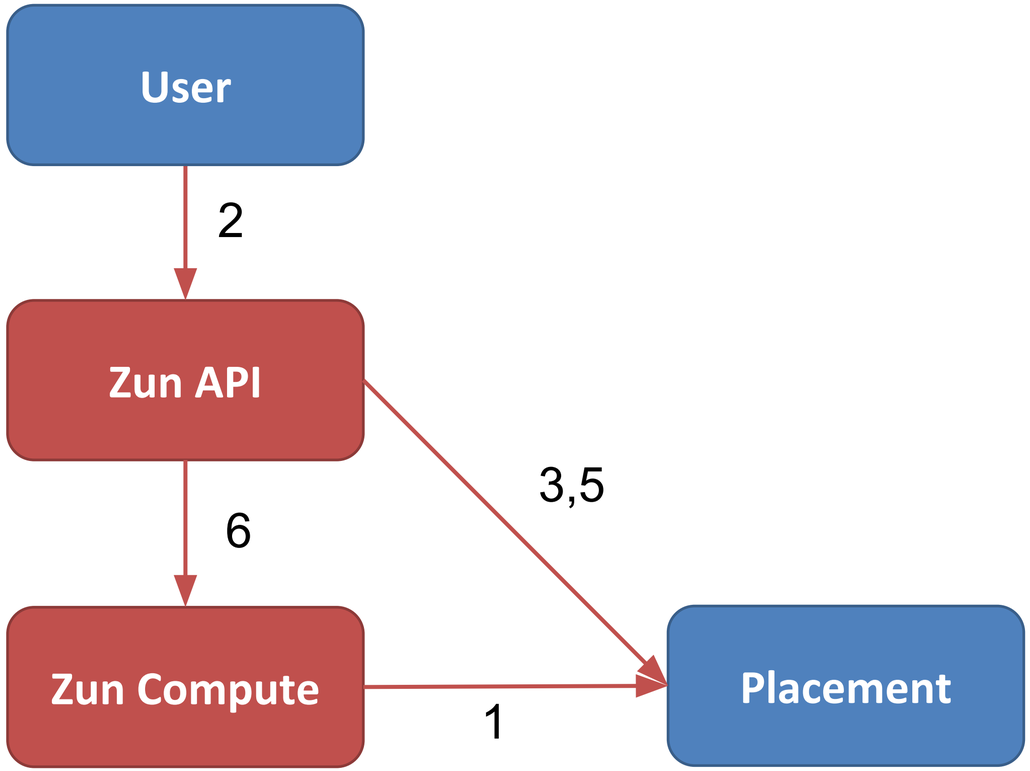

Shared Compute Node with Nova

- Zun and Nova works on the same resource provider
  - Nova Compute creates the resource provider
  - Nova Compute reports inventories
  - Zun Compute reports traits
- Zun and Nova claim allocations on the same resource provider
  - Nova claims allocations for VMs
  - Zun claims allocations for containers/pods

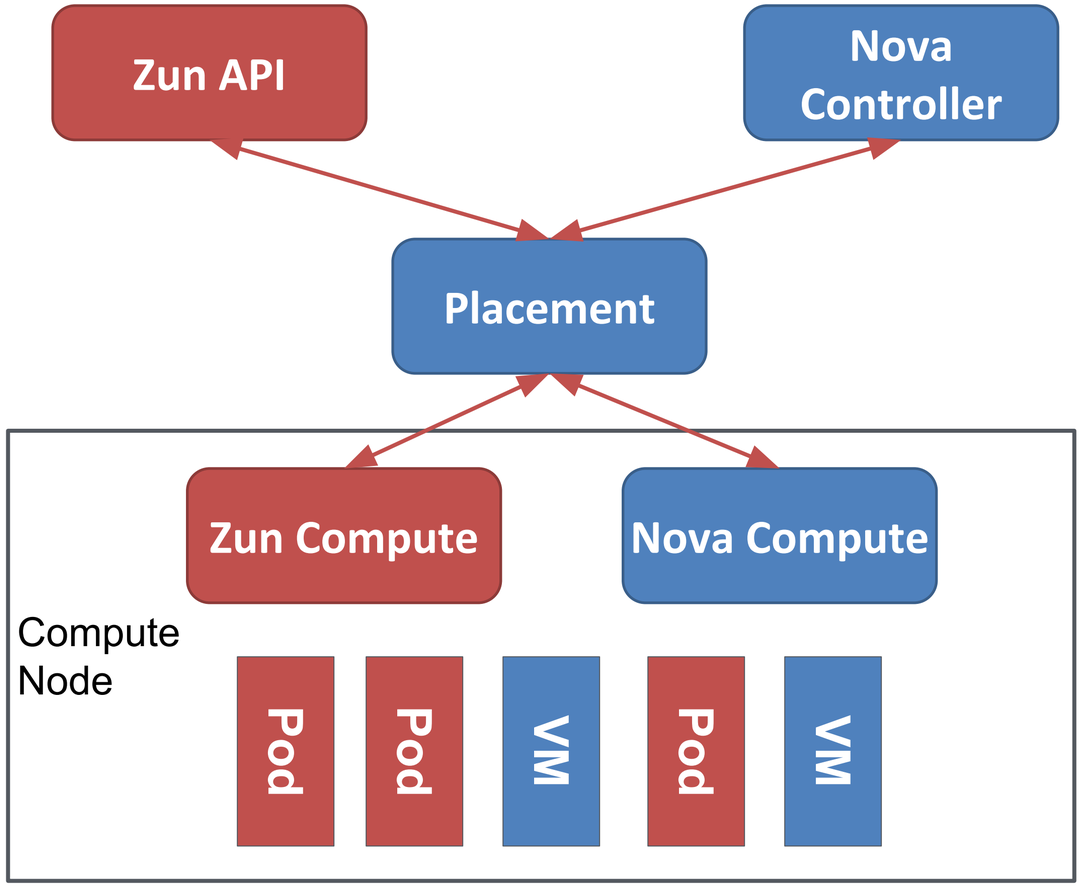

openstack.

Agenda

- Introduction
- OpenStack Zun
- **Virtual Kubelet**
- OpenStack Provider for Virtual Kubelet
- Demo

Virtual Kubelet

- Kubelet implementation, masquerades container service as Kubelet node.
- **Kubernetes on top, programmed back.**
- Intermediary to map Kubernetes requests and resource to container service

Virtual Kubelet

Virtual Kubelet Framework
- PodLifeCycle Controller
- NodeLifeCycle Controller

Provider Implementation:
- Virutal Node Registration
- Pod/Node method implementation

openstack.

Virtual Kubelet

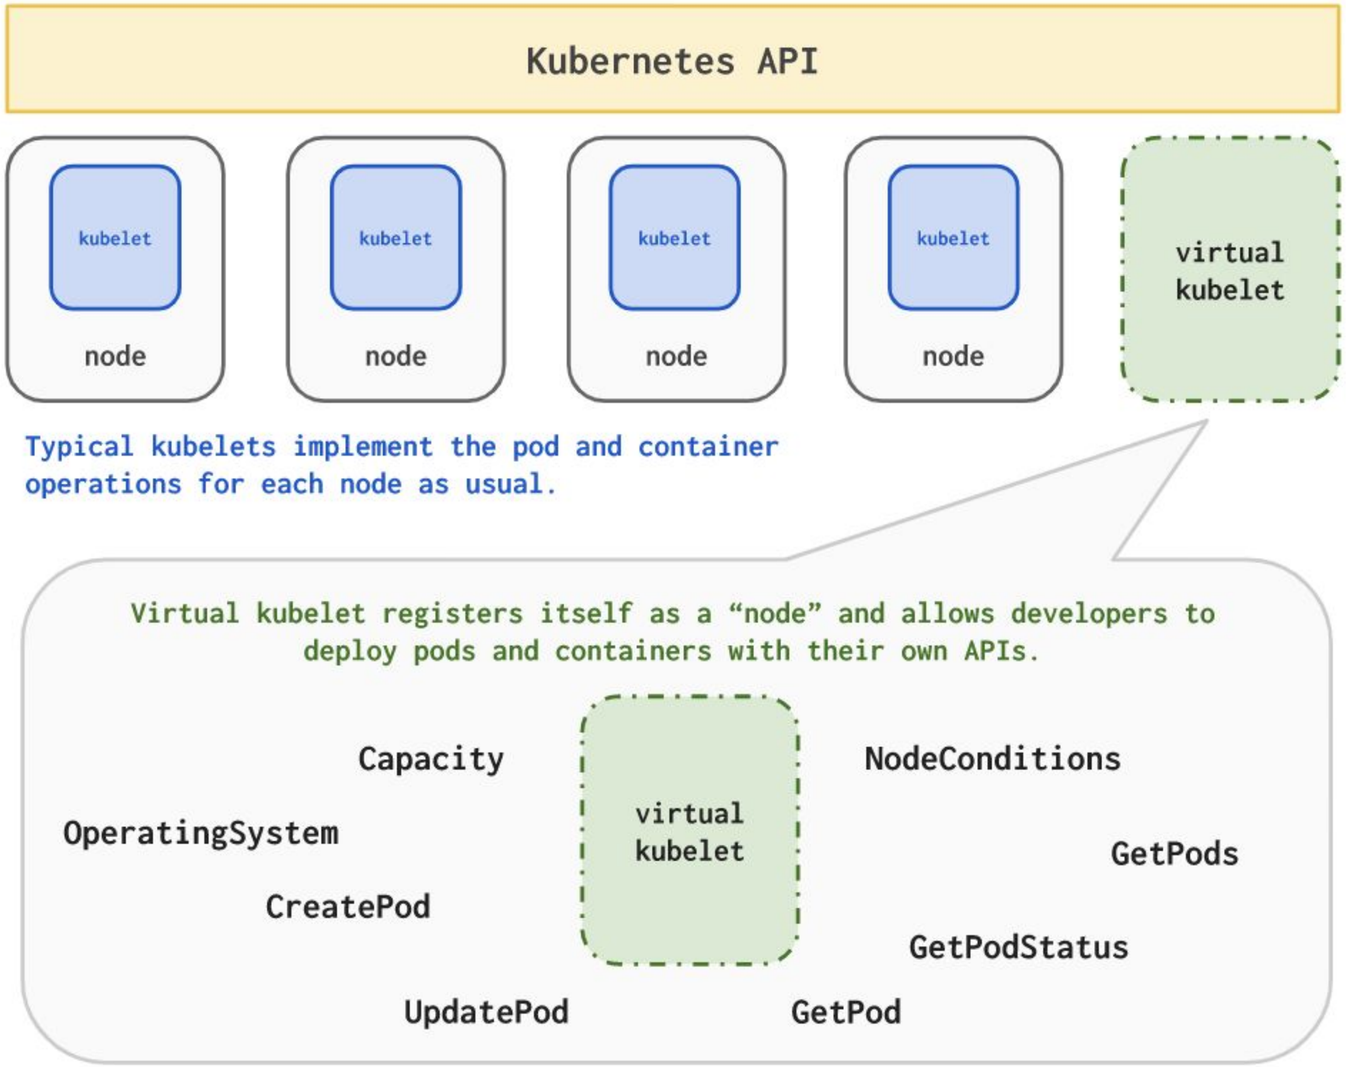

Virtual Kubelet Pod Lifecycle

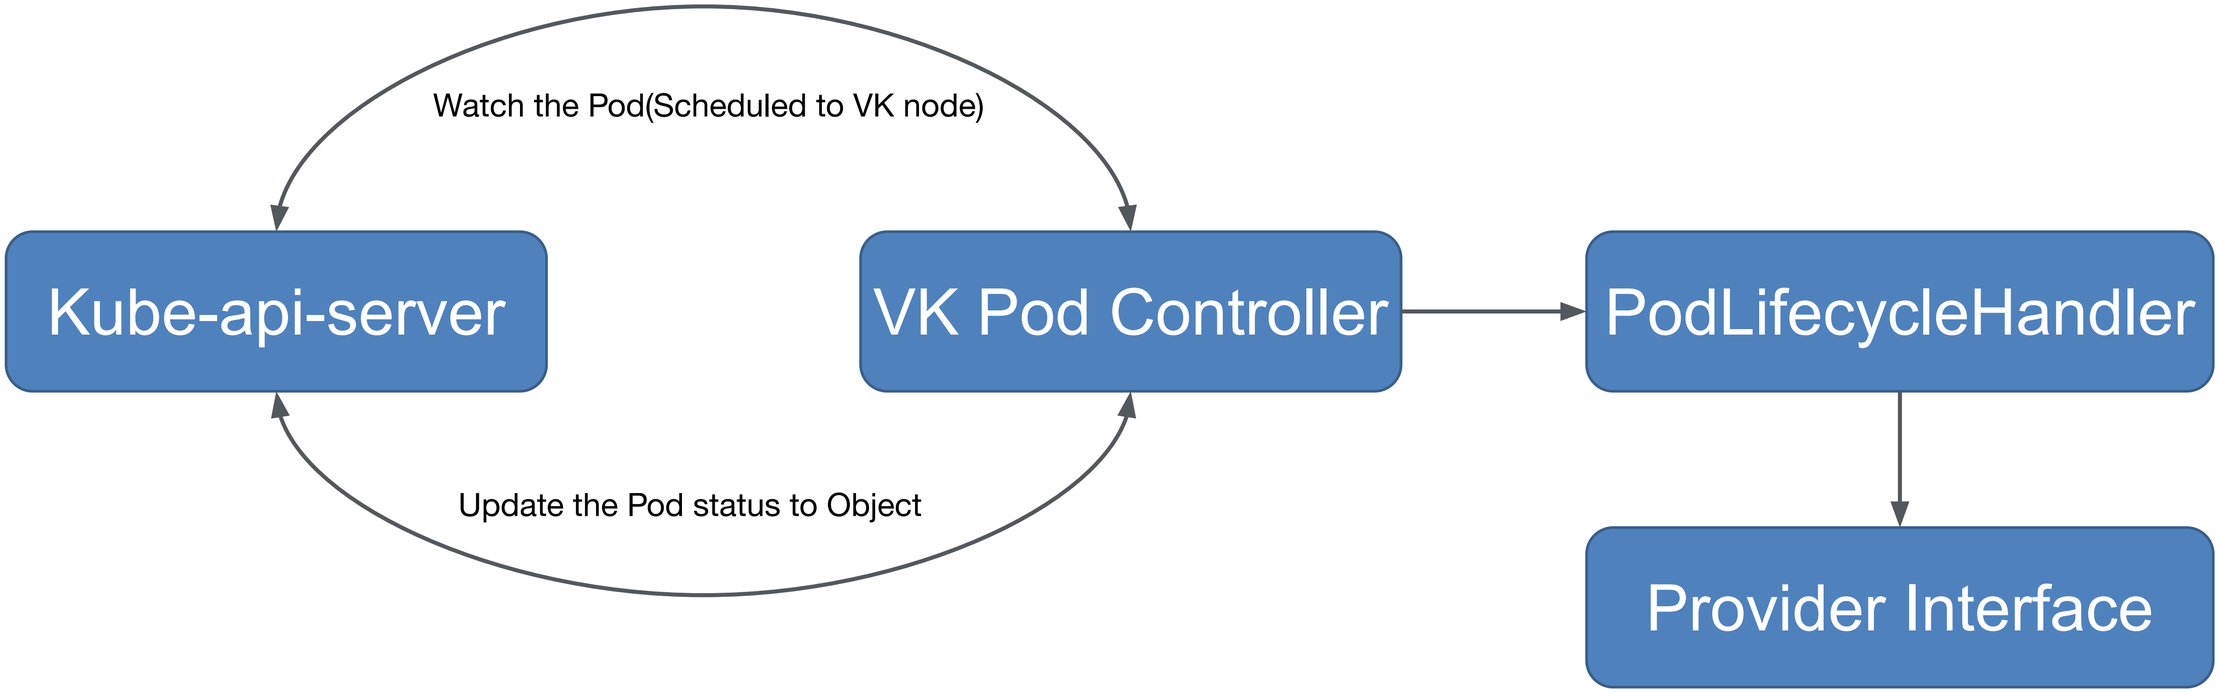

openstack.

PodLifycycleHandler

```
type PodLifecycleHandler interface {
    // CreatePod takes a Kubernetes Pod and deploys it within the provider.
    CreatePod(ctx context.Context, pod *corev1.Pod) error

    // UpdatePod takes a Kubernetes Pod and updates it within the provider.
    UpdatePod(ctx context.Context, pod *corev1.Pod) error

    // DeletePod takes a Kubernetes Pod and deletes it from the provider.
    DeletePod(ctx context.Context, pod *corev1.Pod) error

    // GetPod retrieves a pod by name from the provider (can be cached).
    GetPod(ctx context.Context, namespace, name string) (*corev1.Pod, error)

    // GetPodStatus retrieves the status of a pod by name from the provider.
    GetPodStatus(ctx context.Context, namespace, name string) (*corev1.PodStatus, error)

    // GetPods retrieves a list of all pods running on the provider (can be cached).
    GetPods(context.Context) ([]*corev1.Pod, error)
}
```

Virtual Kubelet Node Lifecycle

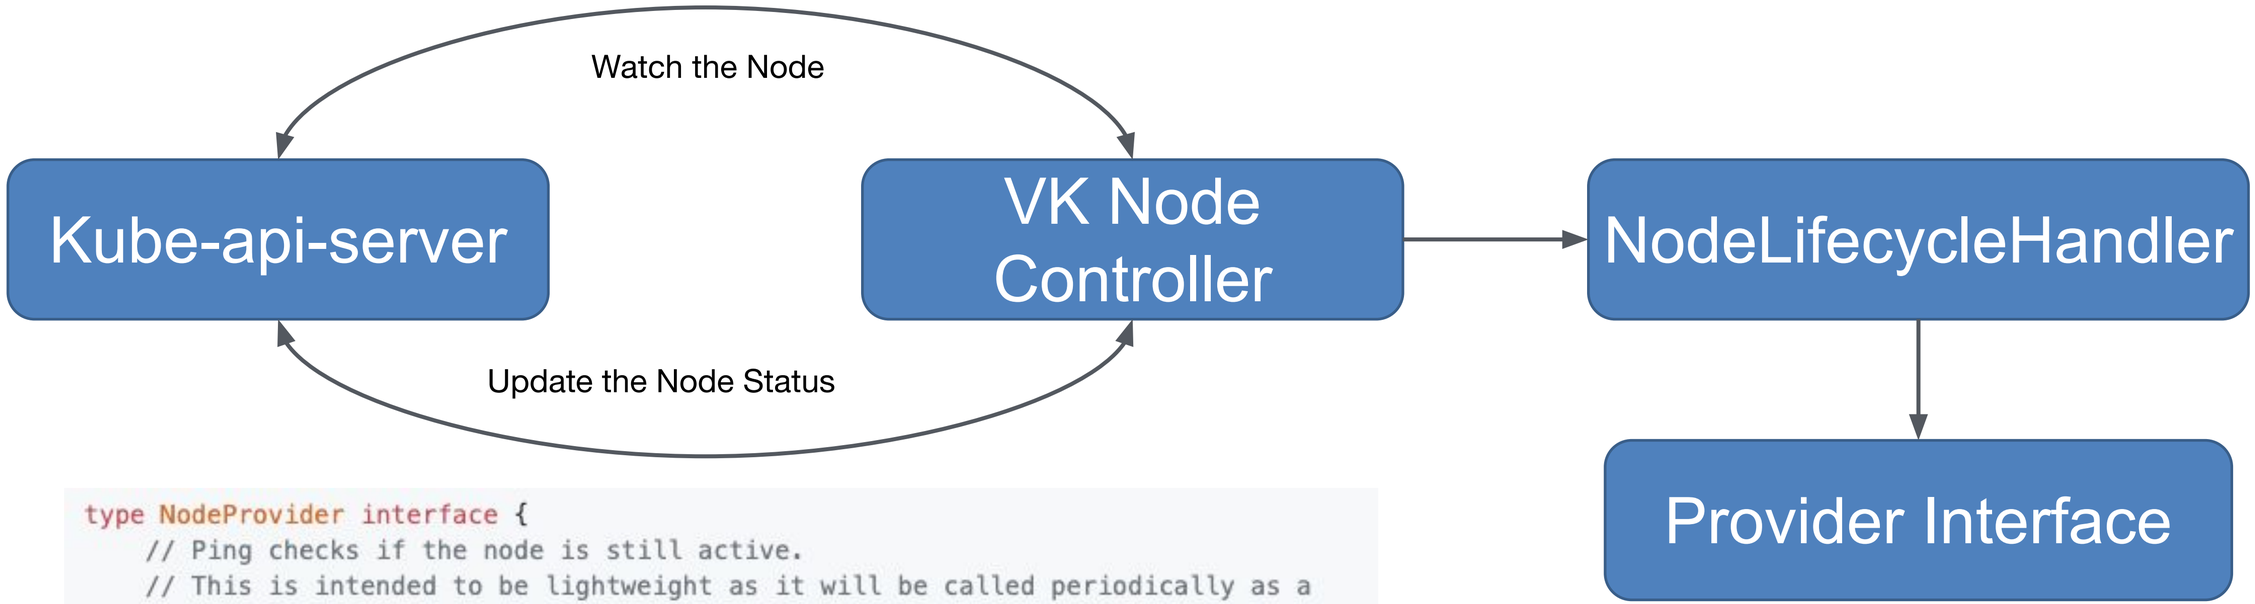


```
type NodeProvider interface {
    // Ping checks if the node is still active.
    // This is intended to be lightweight as it will be called periodically as a
    // heartbeat to keep the node marked as ready in Kubernetes.
    Ping(context.Context) error

    // NotifyNodeStatus is used to asynchronously monitor the node.
    // The passed in callback should be called any time there is a change to the
    // node's status.
    // This will generally trigger a call to the Kubernetes API server to update
    // the status.
    //
    // NotifyNodeStatus should not block callers.
    NotifyNodeStatus(ctx context.Context, cb func(*corev1.Node))
}
```

openstack.

Agenda

- ➡️ Introduction
- ➡️ OpenStack Zun
- ➡️ Virtual Kubelet
- ➡️ **OpenStack Provider for Virtual Kubelet**
- ➡️ Demo

OpenStack Provider with Zun

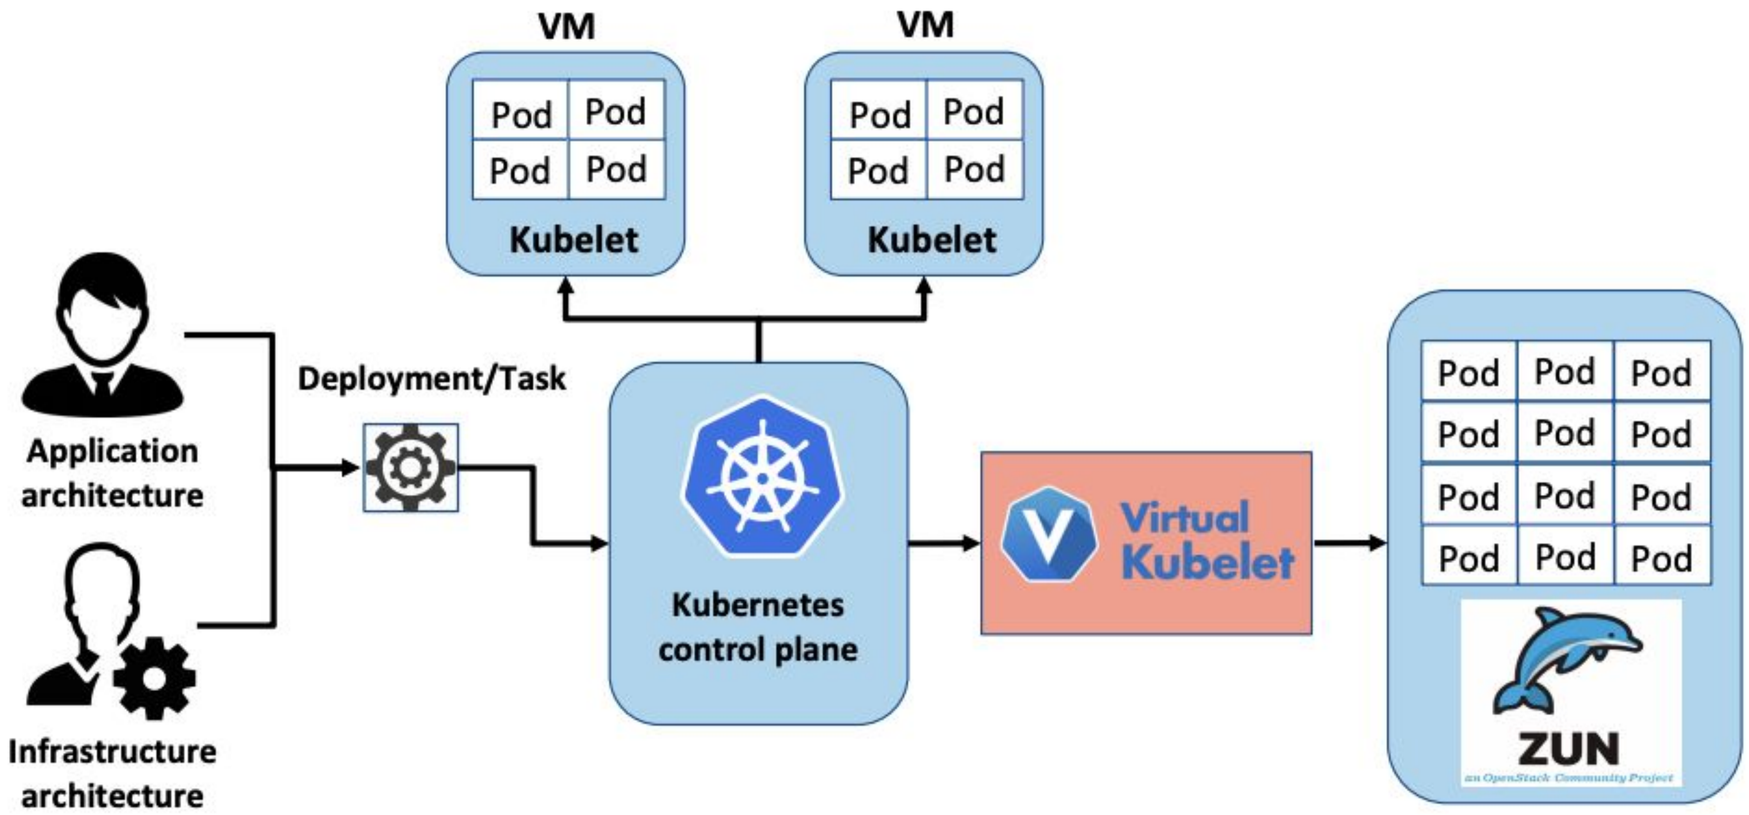

openstack.

OpenStack Provider - Registration

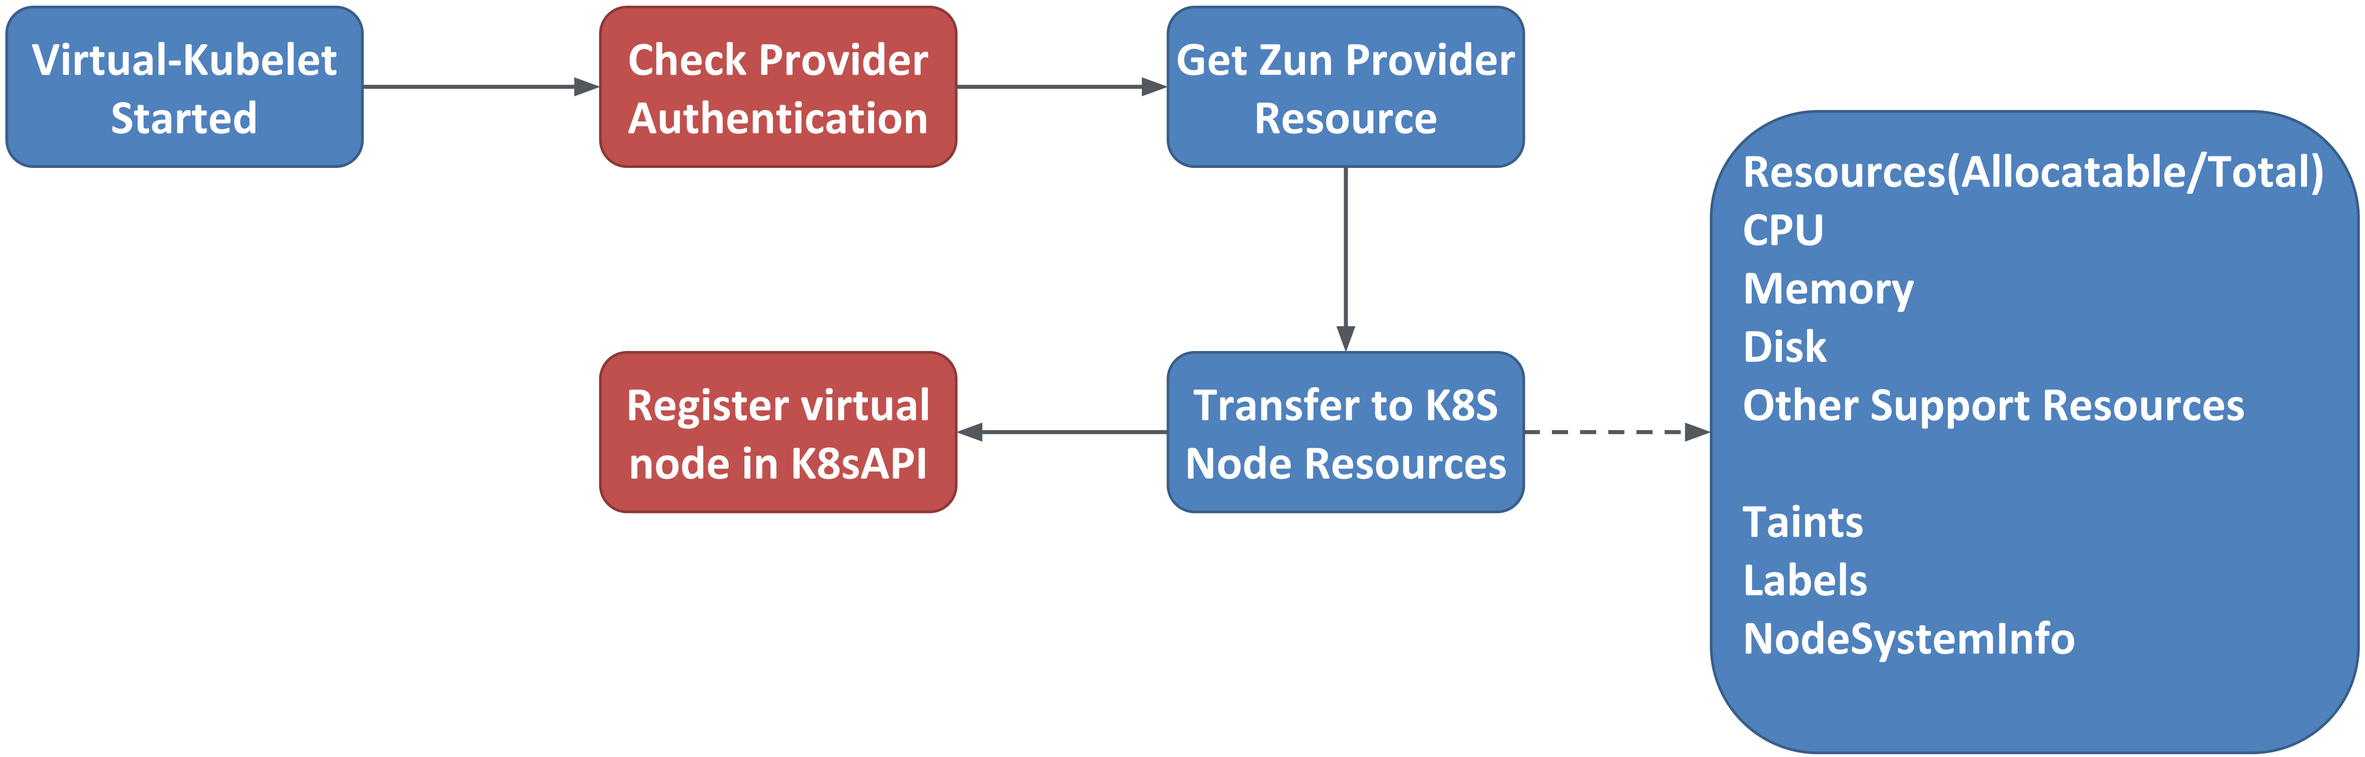

OpenStack Provider

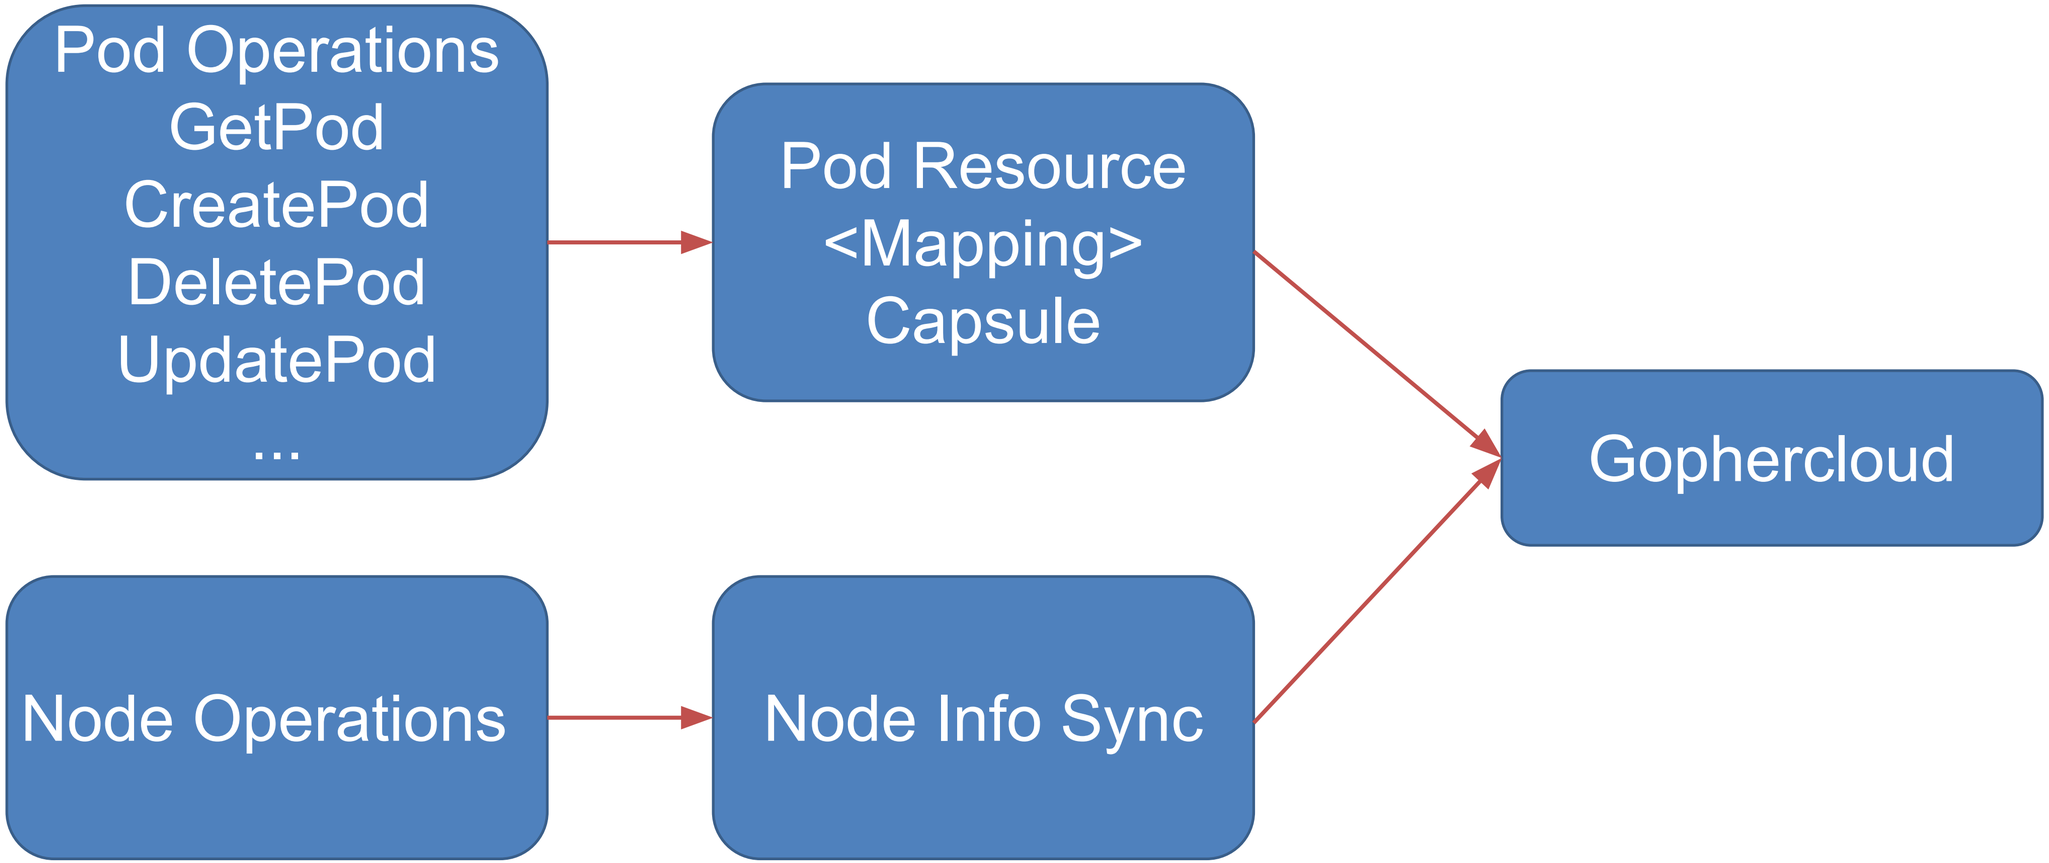

OpenStack Provider - Capsule

Container Capsule - Pod implementation in Zun

- One Sandbox container
- Multiple containers
- Multiple volumes Component

Features:

- Basic unit in Zun
- Co-Scheduled/Co-located
- Share the network namespace
- Share the resource limits

openstack.

Agenda

- Introduction
- OpenStack Zun
- Virtual Kubelet
- OpenStack Provider for Virtual Kubelet
- **Demo**

Demo

https://youtu.be/tC_f31Lse5k

openstack.

Q&A

Thank you!

 openstack  @OpenStack  openstack  OpenStackFoundation