

Introduction to Metal³ - Kubernetes Native Bare Metal Provisioning

Sunnatillo Samadov
Open Source Developer @Ericsson


About Me

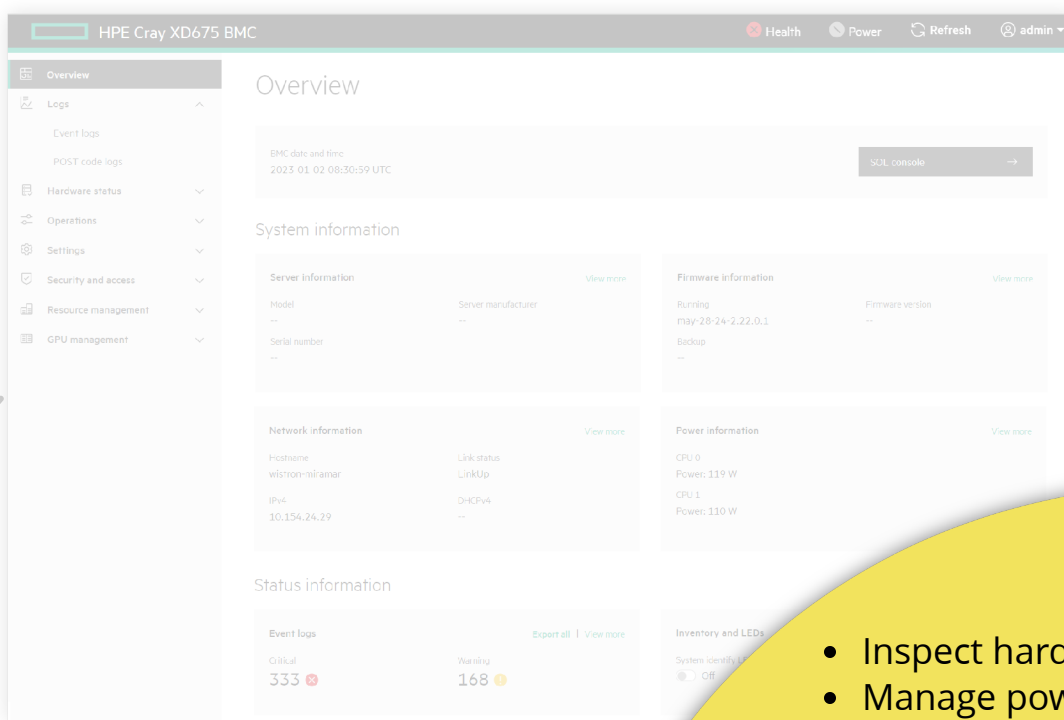
- Open Source Developer @Ericsson
- Maintainer of Metal3.io
- 5 years in open source
- Cluster API release team lead for release v1.9 release



**Sunnatillo
Samadov**

Why run Kubernetes on Baremetal?





- Inspect hardware inventory (CPU, RAM, disks, NICs)
- Manage power state (on/off/reboot)
- Configure boot settings (PXE, UEFI, boot device)
- Provision OS images to physical nodes
- Inject configuration (cloud-init)
- Select and partition root disk
- Deprovision and wipe nodes
- Perform disk cleaning and secure erase
- Configure BIOS/firmware settings
- Update firmware (BIOS/BMC where supported)
- Configure RAID before provisioning

```
1 apiVersion: metal3.io/v1alpha1
2 kind: BareMetalHost
3 metadata:
4   name: eselda13u31s05
5 spec:
6   online: true
7   bootMACAddress: 80:c1:6e:7a:5a:a8
8   bootMode: legacy
9   bmc:
10    address: ipmi://192.168.1.14:623
11    credentialsName: bml-ilo-login-secret-05
12    disableCertificateVerification: true
```

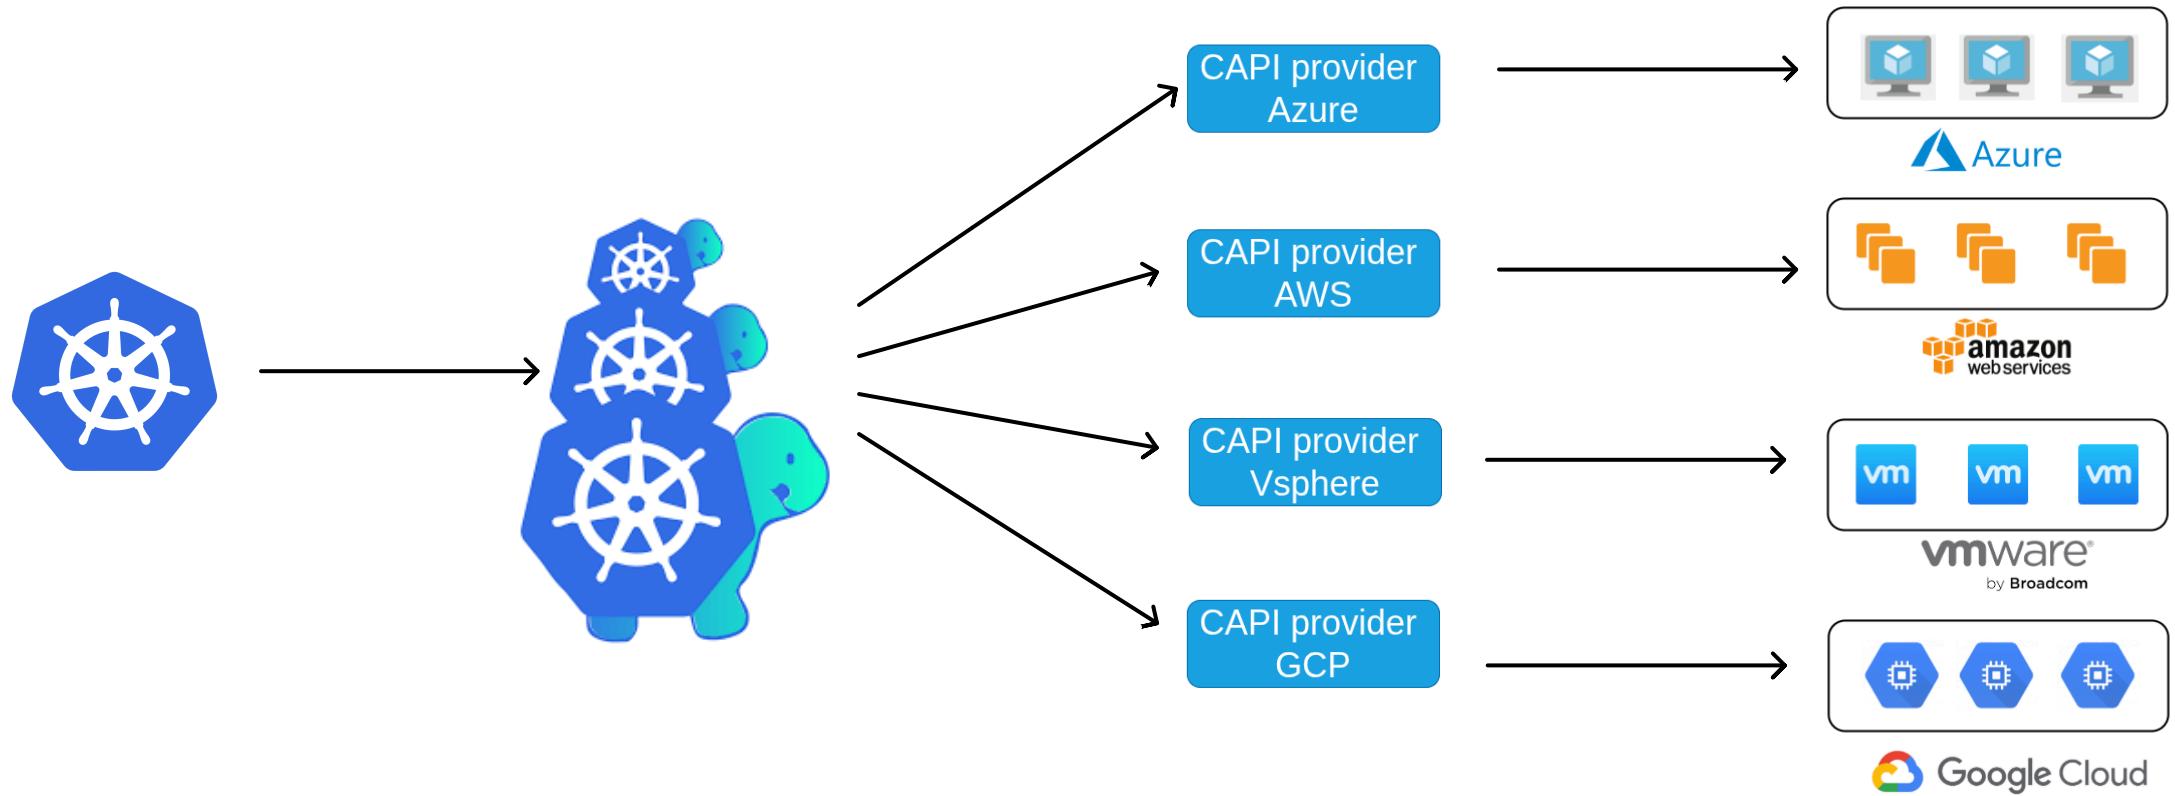
What is Metal3?

- The Metal³ provides components for bare metal host management with Kubernetes.
- You can enroll your bare metal machines, provision OS and deploy Kubernetes clusters to them.
- Metal3 itself is Kubernetes application

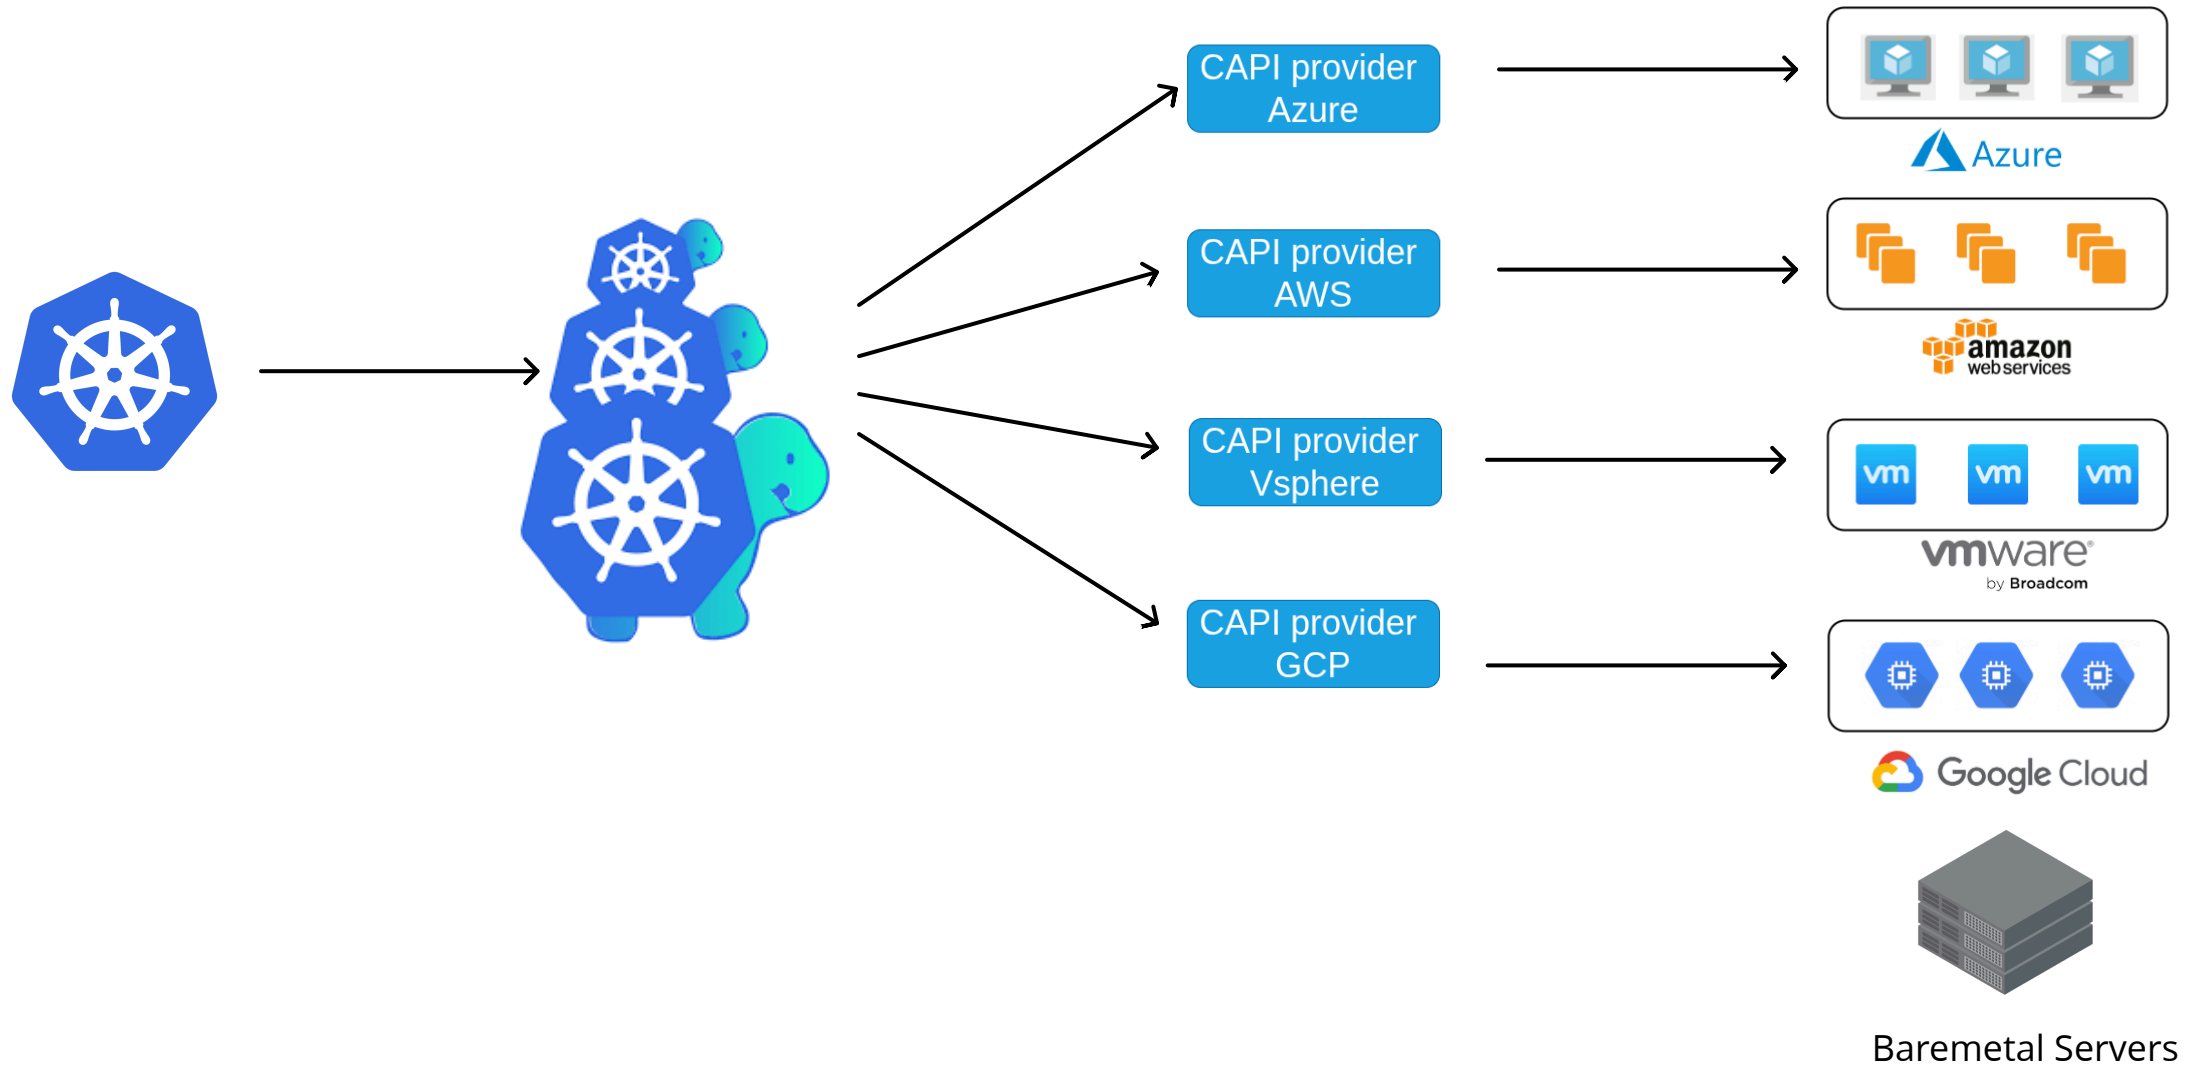


Cluster API

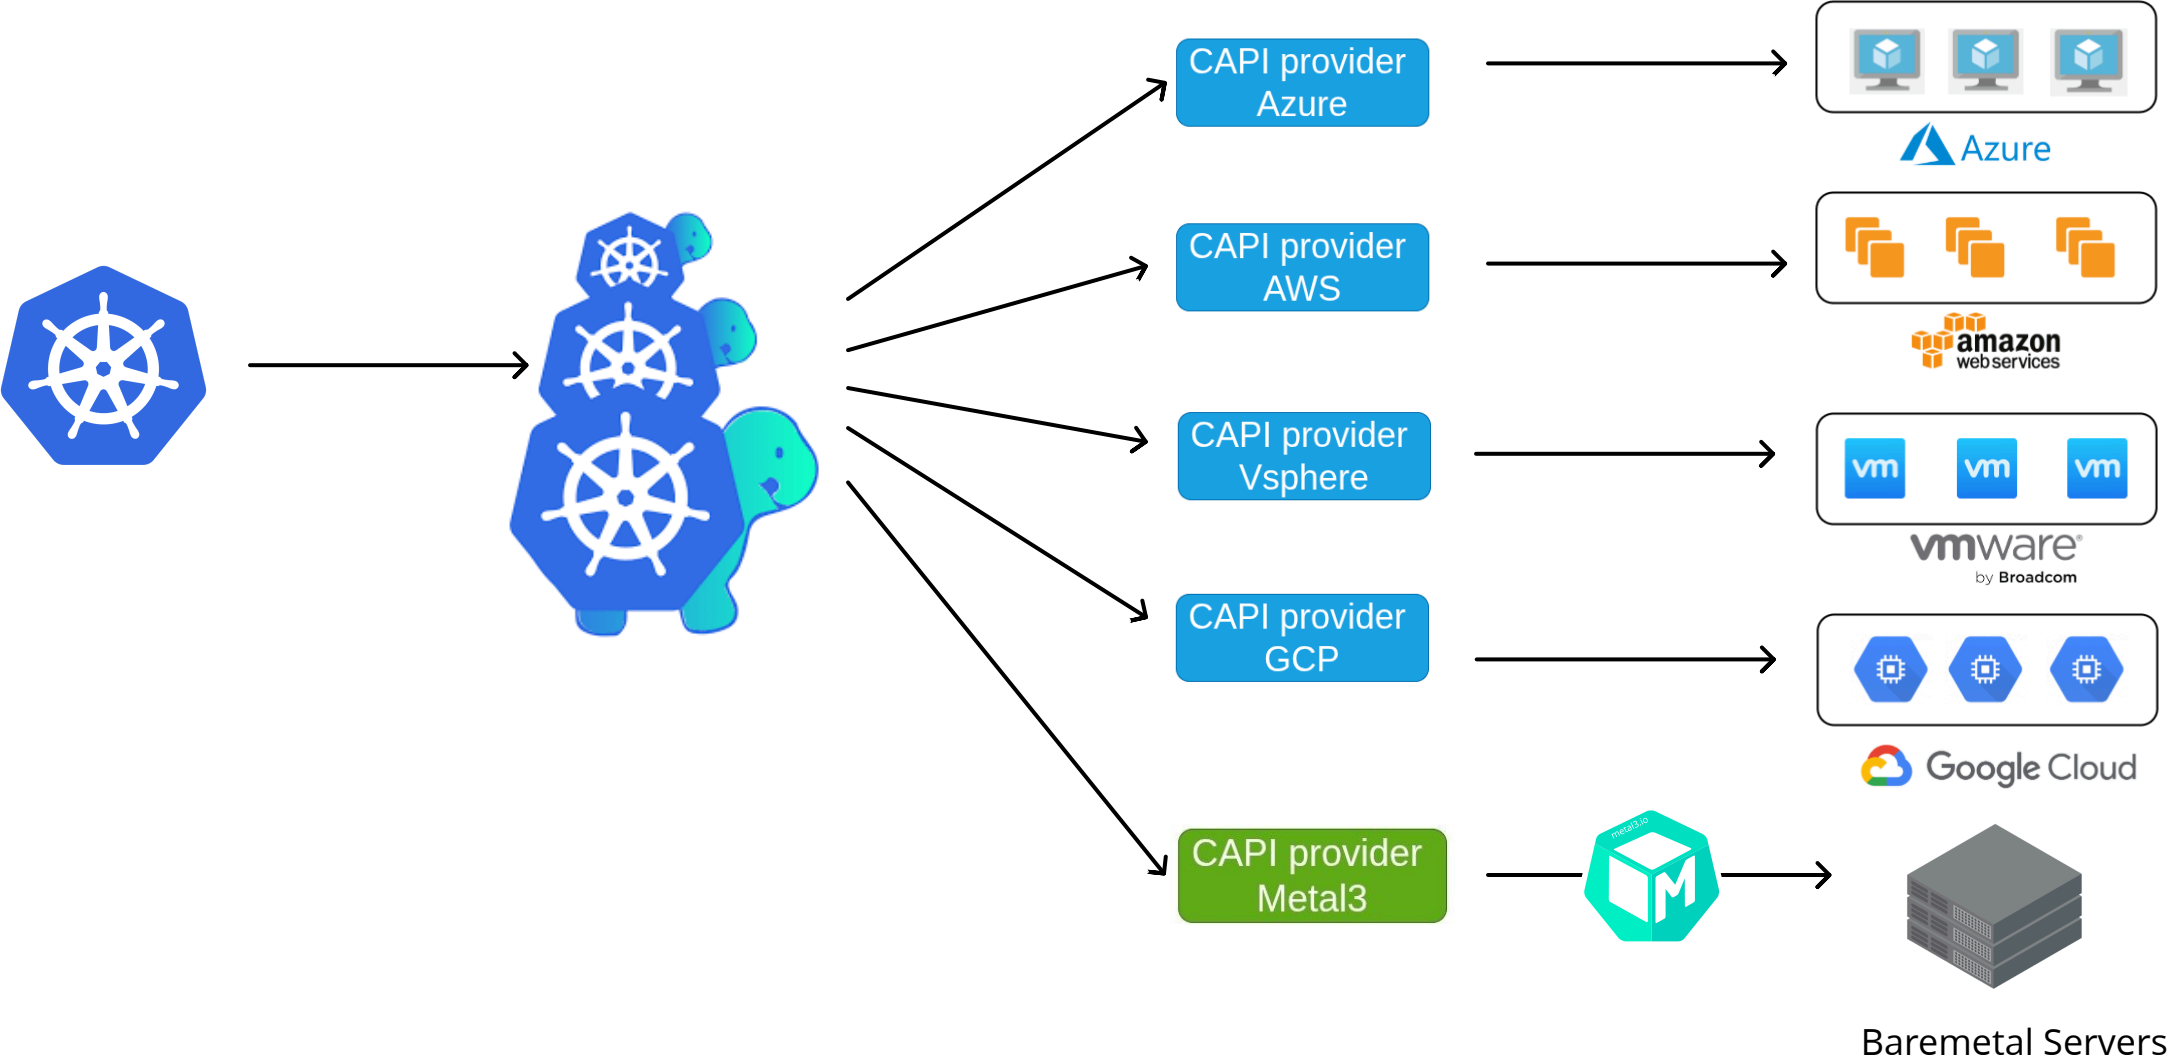
Cluster API



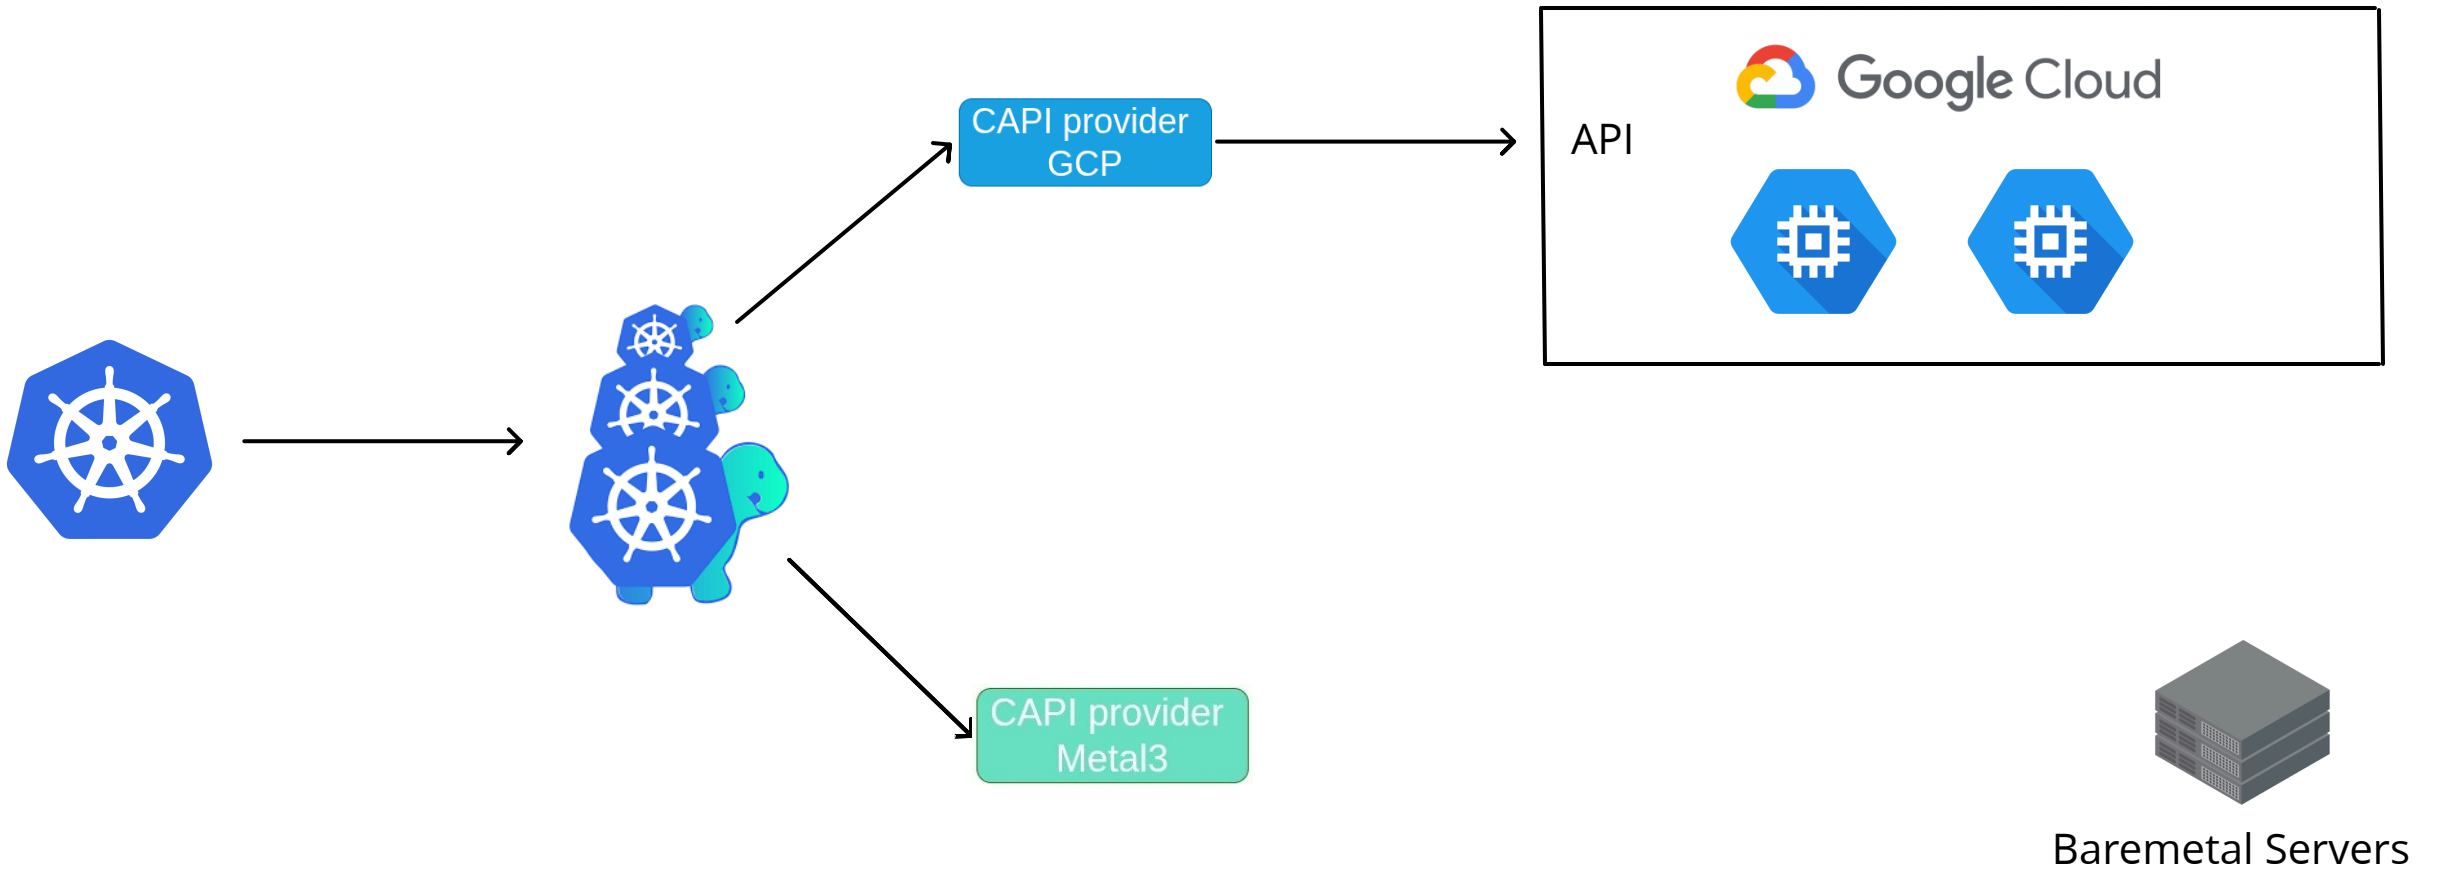
Cluster API



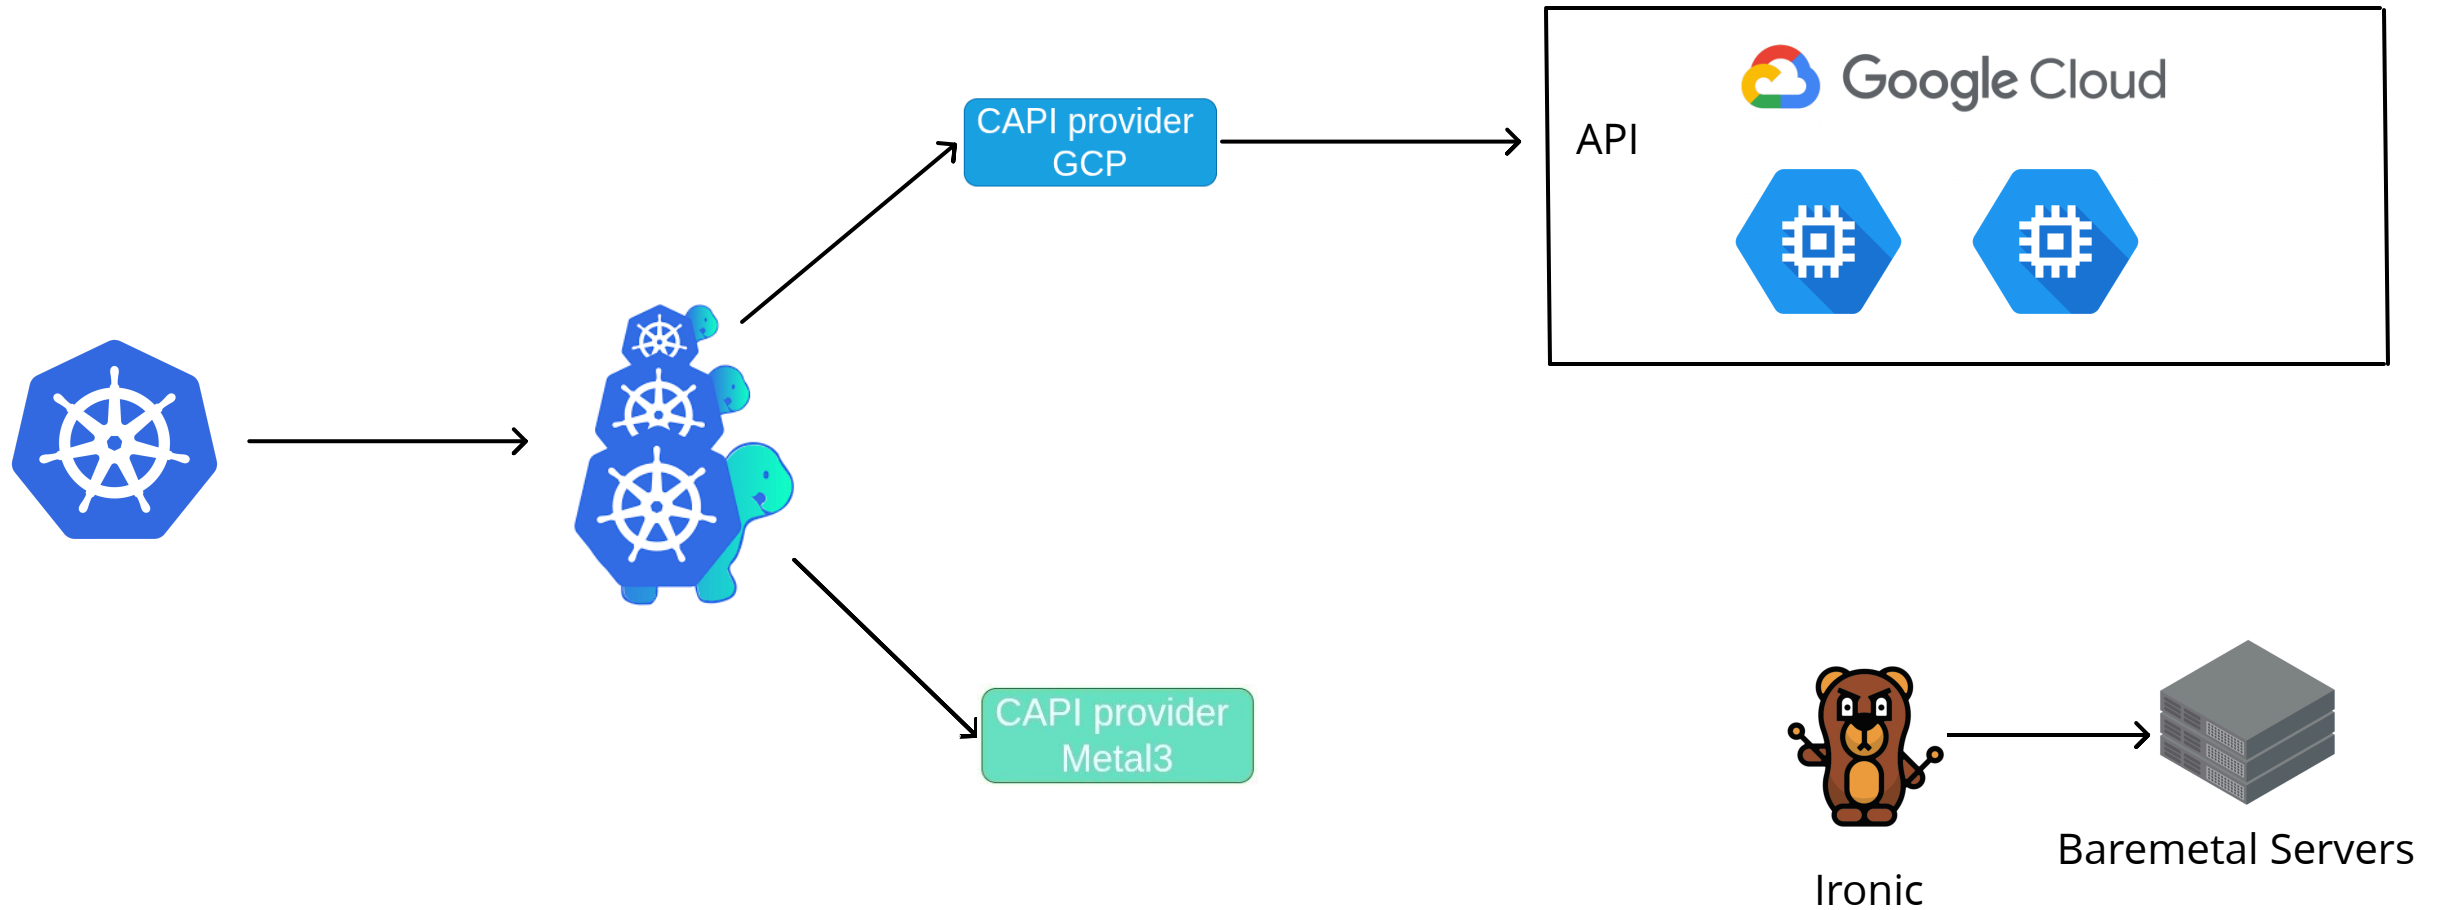
Cluster API



Cluster API



Cluster API



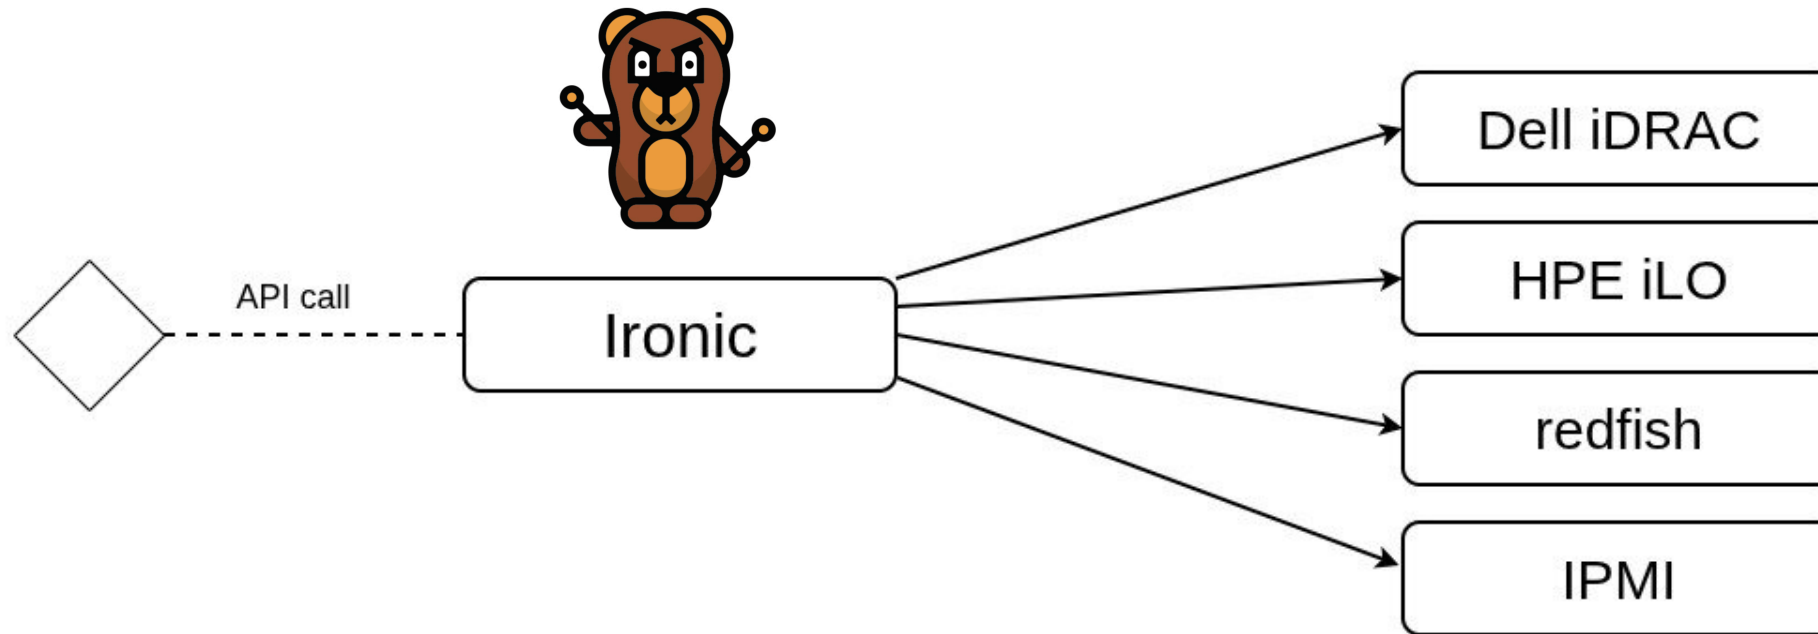
Ironic

- "Ironic is an open source project that fully manages bare metal infrastructure. It discovers bare-metal nodes, catalogs them in a management database, and manages the entire server lifecycle including enrolling, provisioning, maintenance, and decommissioning."
- Ironic is an open source project, written in Python, that manages baremetal infrastructures.

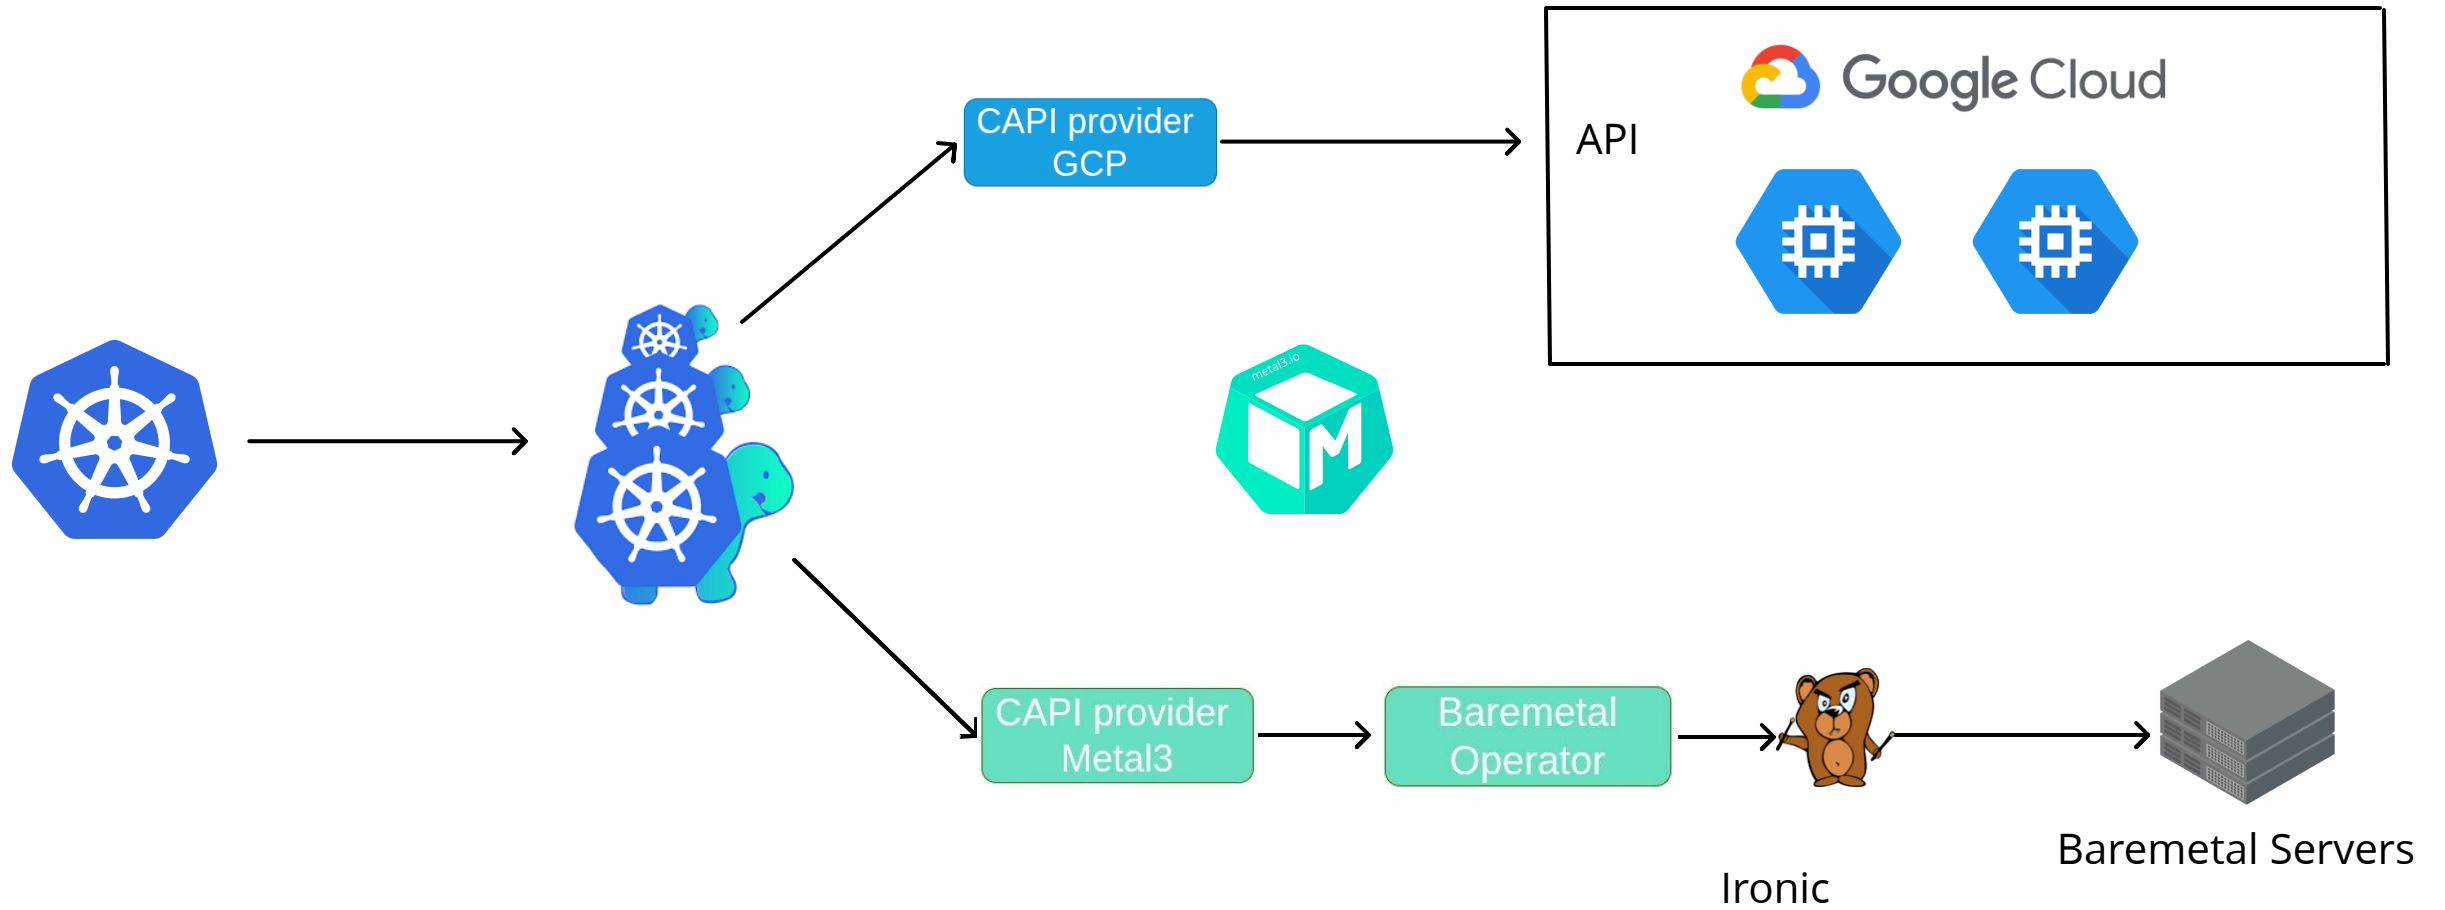


Ironic

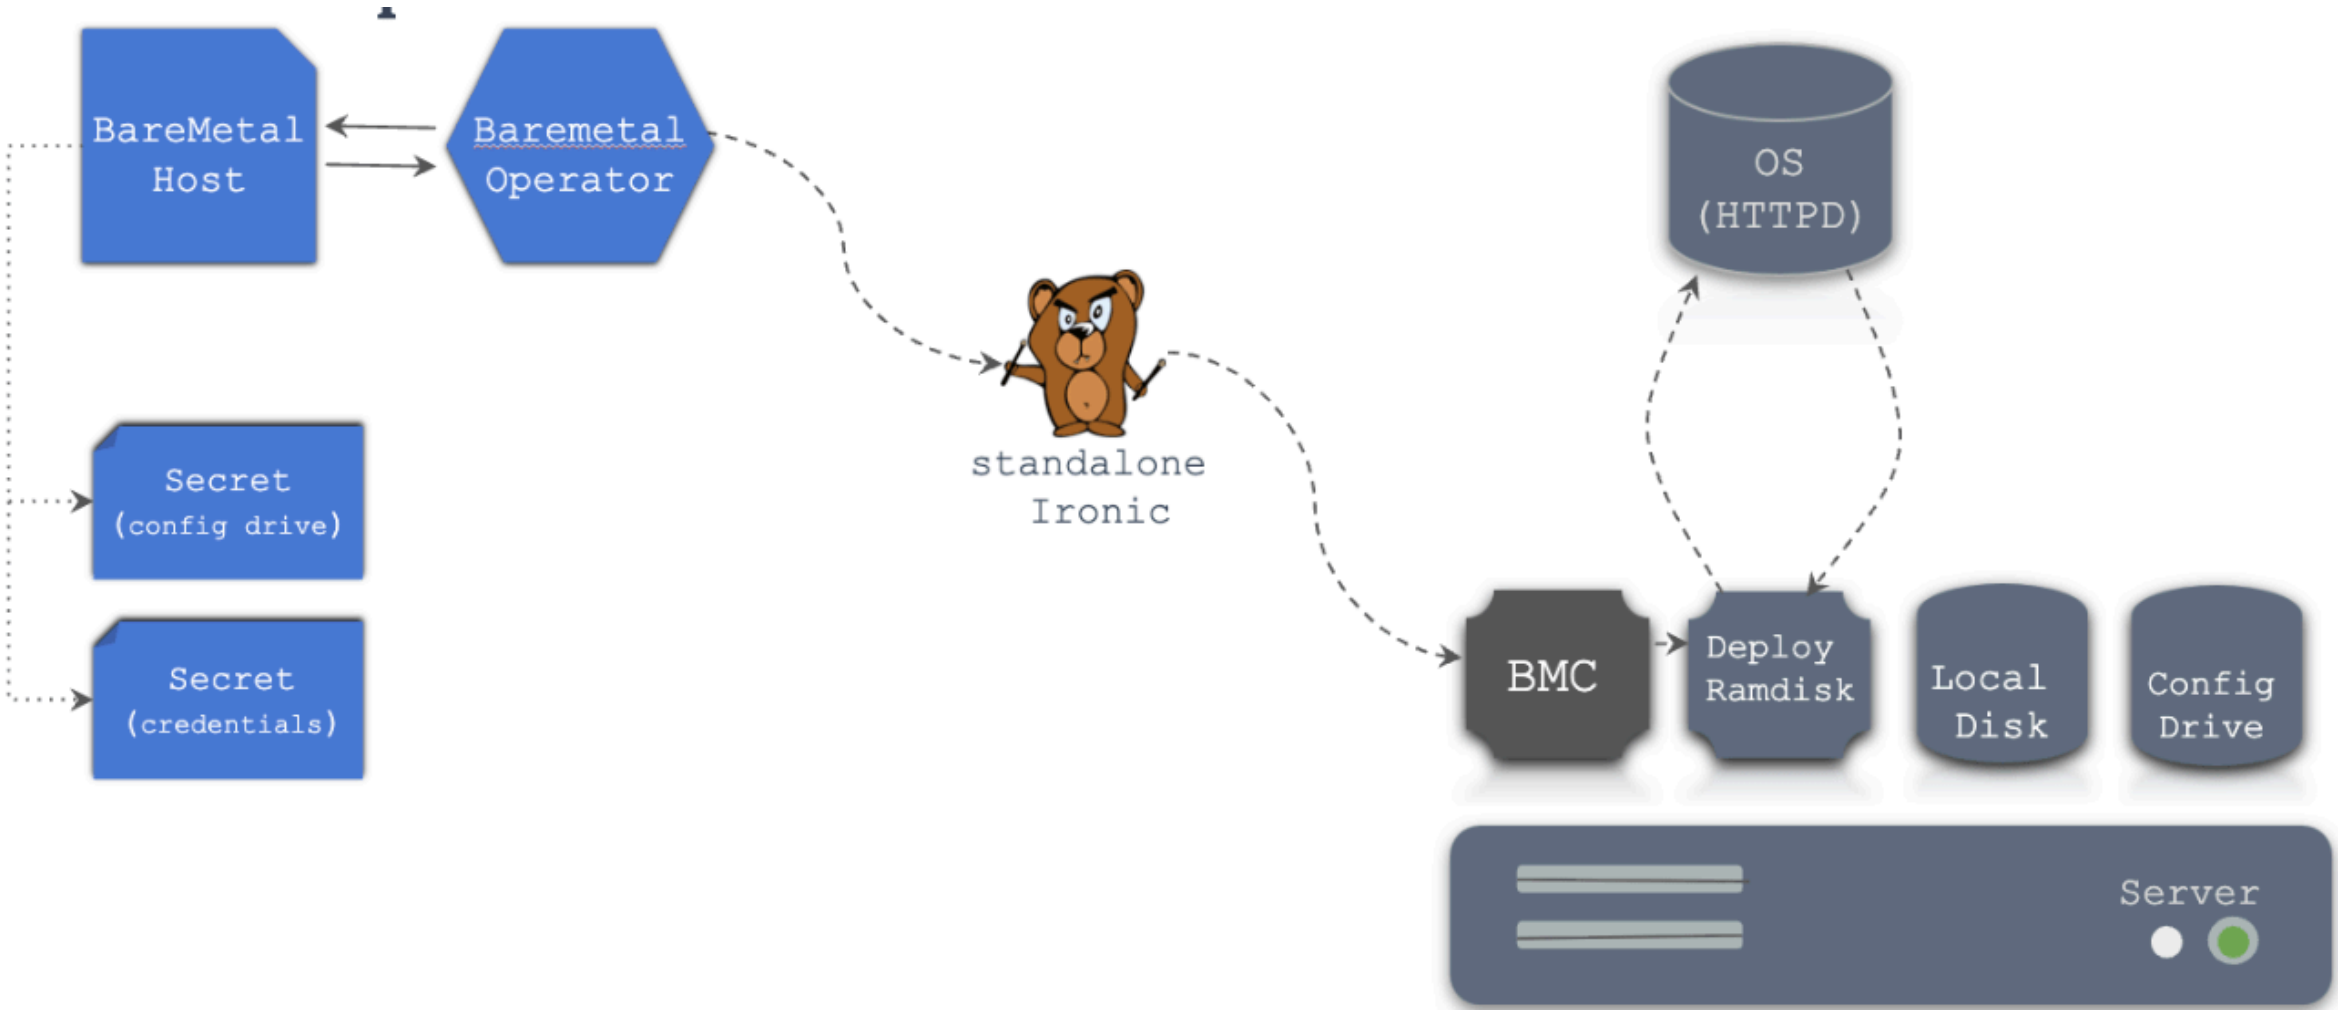
Ironic is API-driven and API-first, and it includes a full set of RESTful APIs that provide a common vendor agnostic interface, allowing provisioning and management of bare metal machines for their entire lifecycle, from enrollment to retirement.



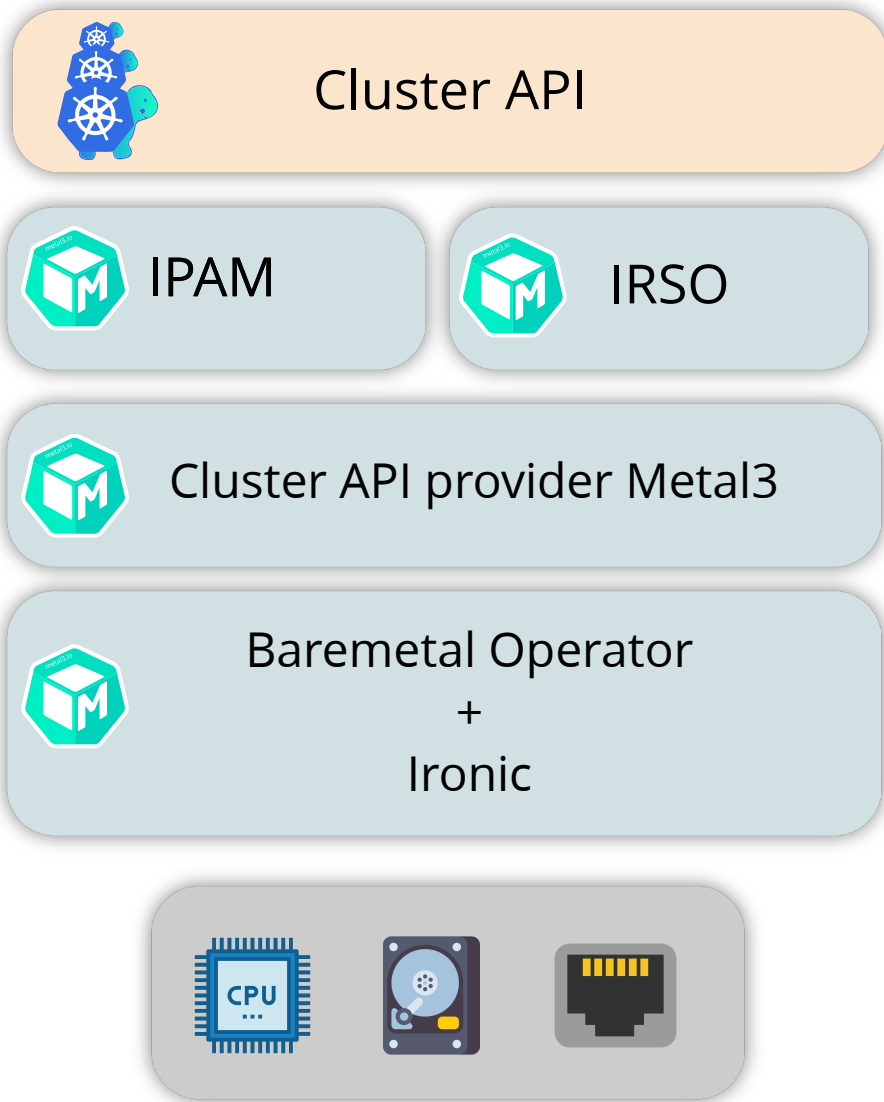
Cluster API



Baremetal Operator



Metal3 Stack



Kubernetes Node abstraction across layers

Cluster API: Machine

CAPM3: Metal3Machine

BMO: BareMetalHost

Ironic: IronicNode



Deployment workflow with Metal3

- Bootstrap cluster
- Metal3 stack deployment (on Bootstrap Cluster)
- Enroll Bare Metal Servers
- Apply Custer, Control Plane and Workers resources
- Bare metal machines becomes provisioned
- Cluster is ready



Metal3 Features

- Remediation through reboot
- CAPI based upgrades: chained upgrade, in-place upgrade
- Self-management (Pivoting): move controllers into the managed cluster
- Node re-use: Use the same host after upgrade/changes
- Automated cleaning
- Verified scalability up to 500 nodes
- Top of the Rack switch configuration (WIP)



Metal3 Adopters

- **Ericsson:** Integrating Metal³ project for bare metal deployments and for bare metal cluster LCM.
- **RedHat:** Red Hat's OpenShift distribution includes Metal³ as part of its solution for automating the deployment of baremetal clusters.
- **Fujitsu:** As a server vendor, they are developing and using Metal³ to provide Fujitsu servers as Kubernetes bare metal nodes.
- **DT Technik:** "Das SCHIFF" uses parts of Metal³ to handle bare metal provisioning in its CAPI-Based Infrastructure.
- **IKEA IT AB:** IKEA IT AB uses Metal³ to handle Bare Metal provisioning and lifecycle management in its CAPI-Based bare metal cloud infrastructure.
- **SUSE:** Metal³ is used for automated bare metal deployment as part of the SUSE Edge solution.
- **Mistral AI:** Mistral AI uses Metal3 for baremetal k8s cluster management in its Mistral Compute cloud offering.
- **Mirantis:** Mirantis uses Metal3 inside it's Mirantis Container Cloud (MCC) product to provision bare metal. Also, planned to use inside k0rdent product particularly for particularly for AI-based GPU workloads.


Metal3 Community

 Slack channel #cluster-api-baremetal at the Kubernetes slack

 Documentation: <https://book.metal3.io/>

 Mailing list: <https://groups.google.com/g/metal3-dev>

 Community meeting every Wednesday 14.00 UTC on Zoom


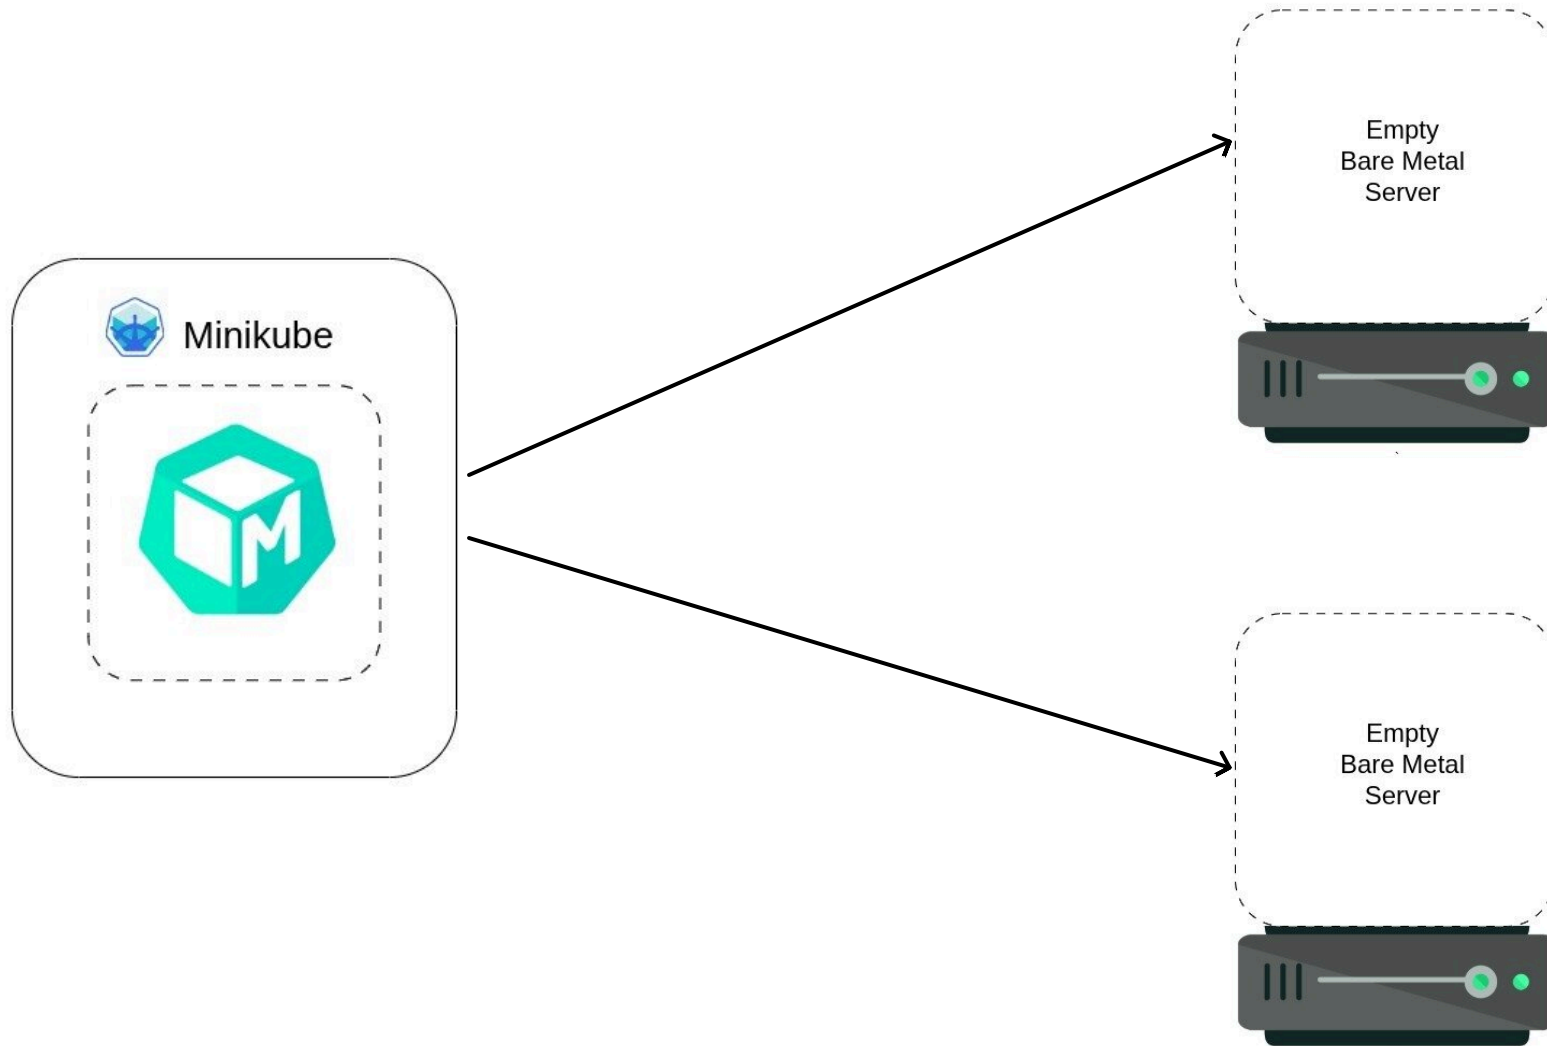
 YouTube: Metal³ YouTube channel

 Github: <https://github.com/metal3-io>


 Website: <https://metal3.io>



Demo



Demo



Questions