

# KUBERNETES UNDER THE HOOD

Richard Szabo  
Duelbox  
Co-Founder, Software Engineer



# WHY BOTHER?



# WHY BOTHER?

Understand the basics of the infrastructure



# WHY BOTHER?

Understand the basics of the infrastructure

Debugging / fixing issues



# WHY BOTHER?

Understand the basics of the infrastructure

Debugging / fixing issues

Helpful in design decisions



# WHY BOTHER?

Understand the basics of the infrastructure

Debugging / fixing issues

Helpful in design decisions

Learn from patterns of a complex distributed system



# WHY BOTHER?

Understand the basics of the infrastructure

Debugging / fixing issues

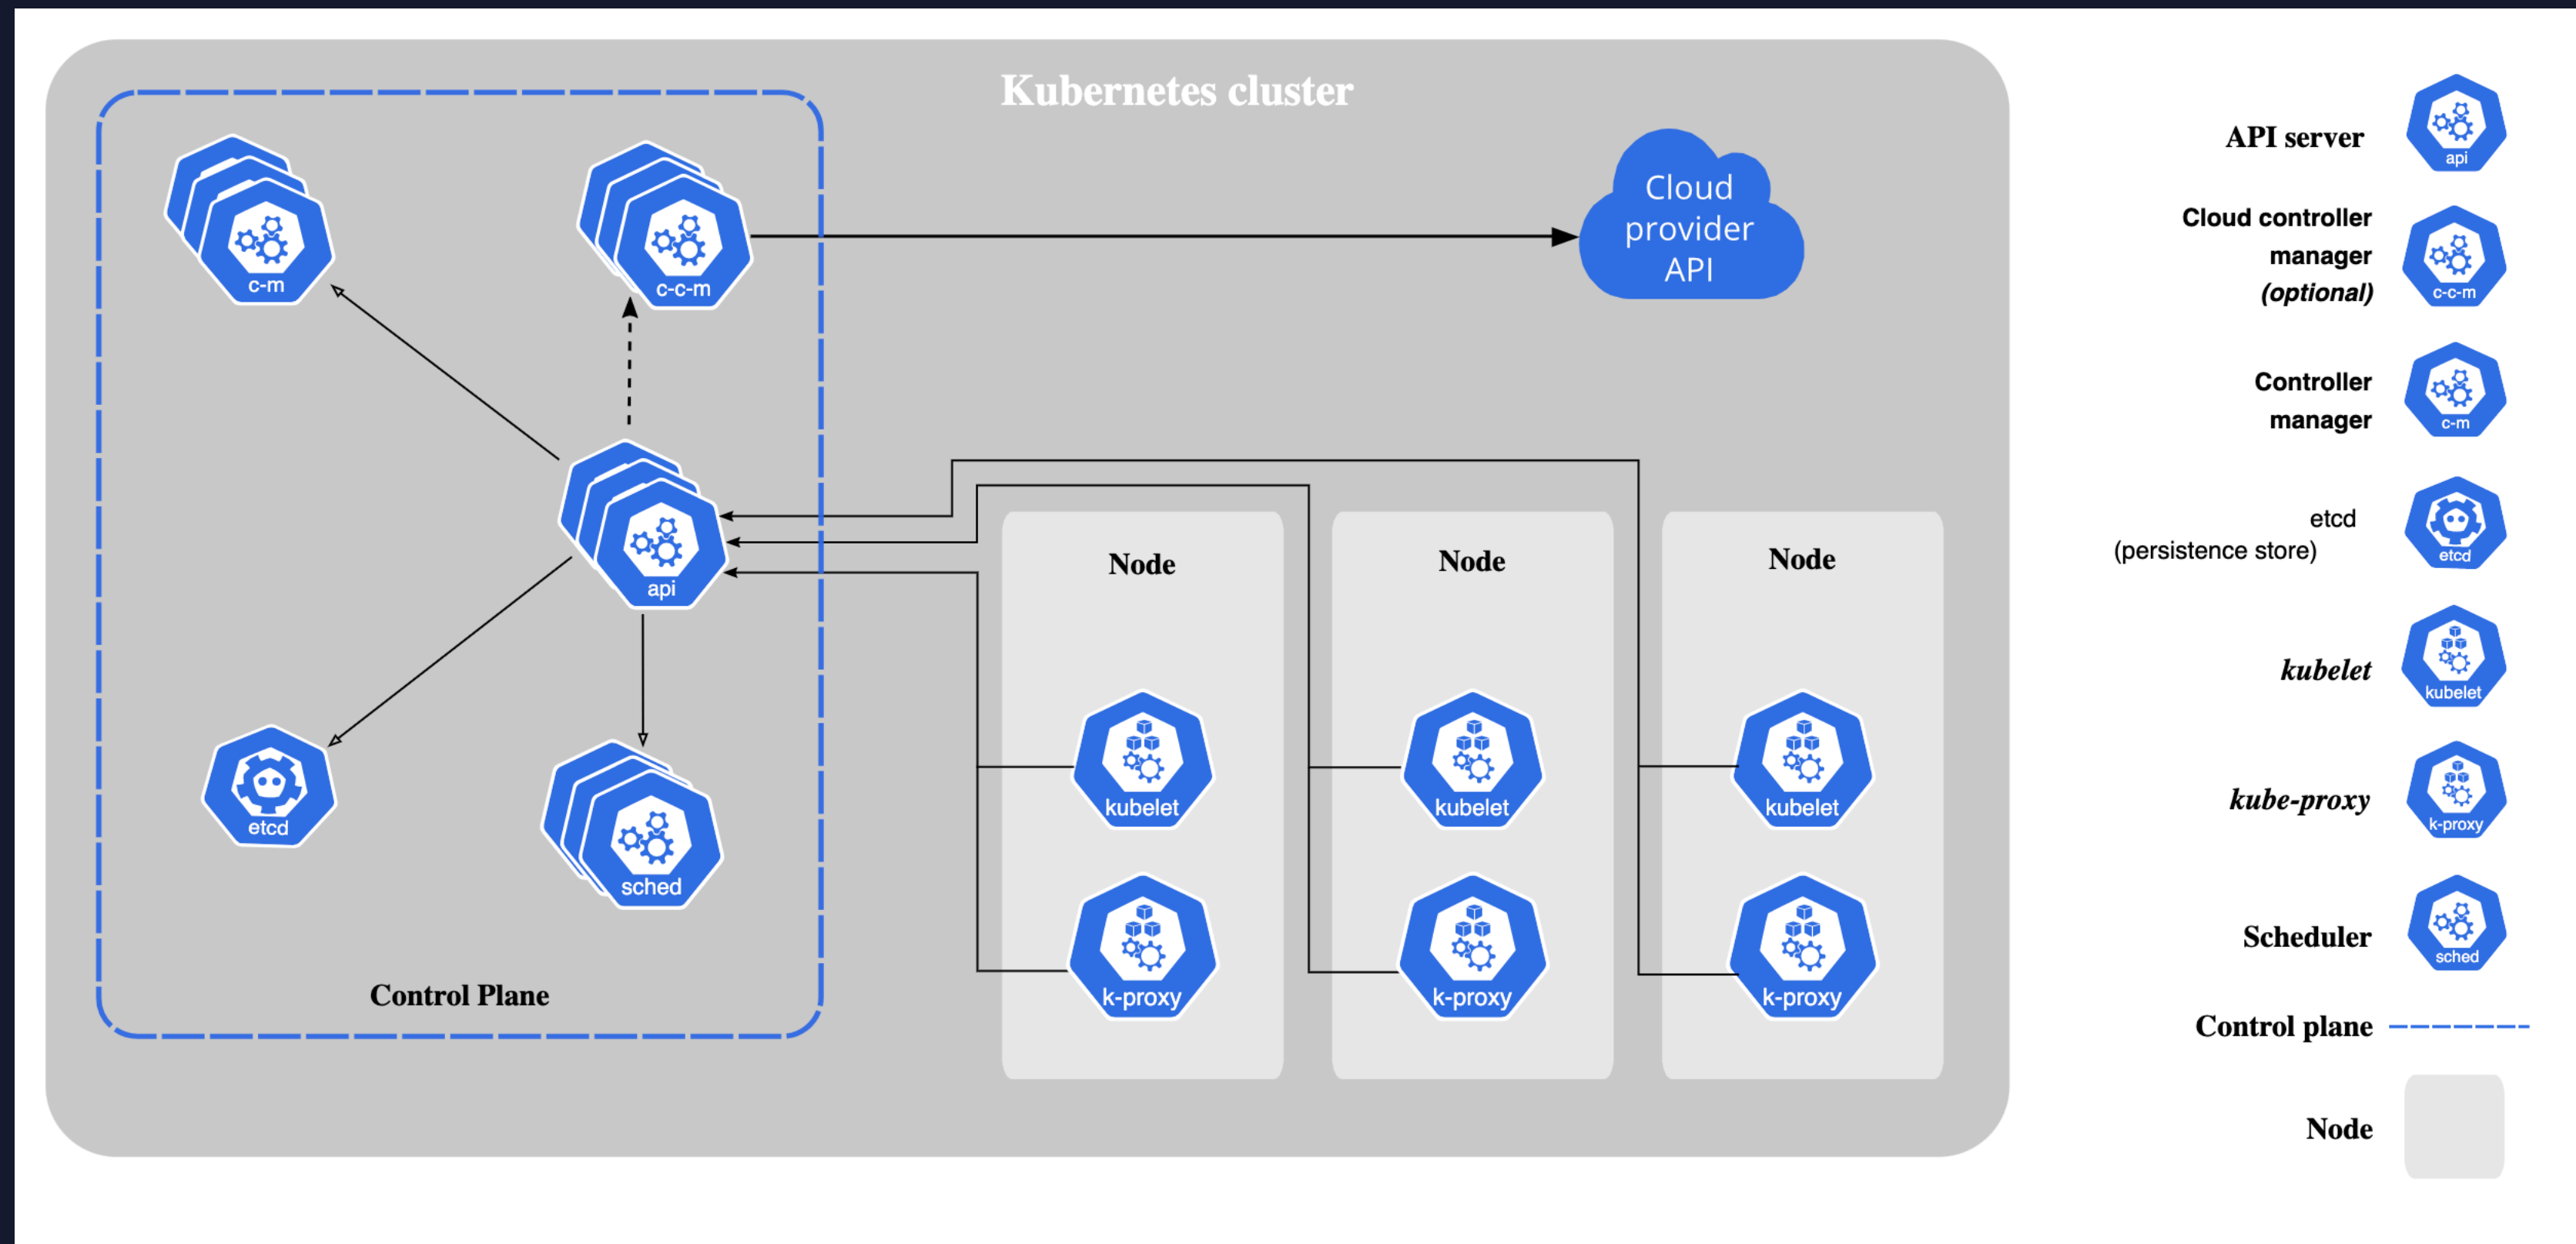
Helpful in design decisions

Learn from patterns of a complex distributed system

Extensions, CRD



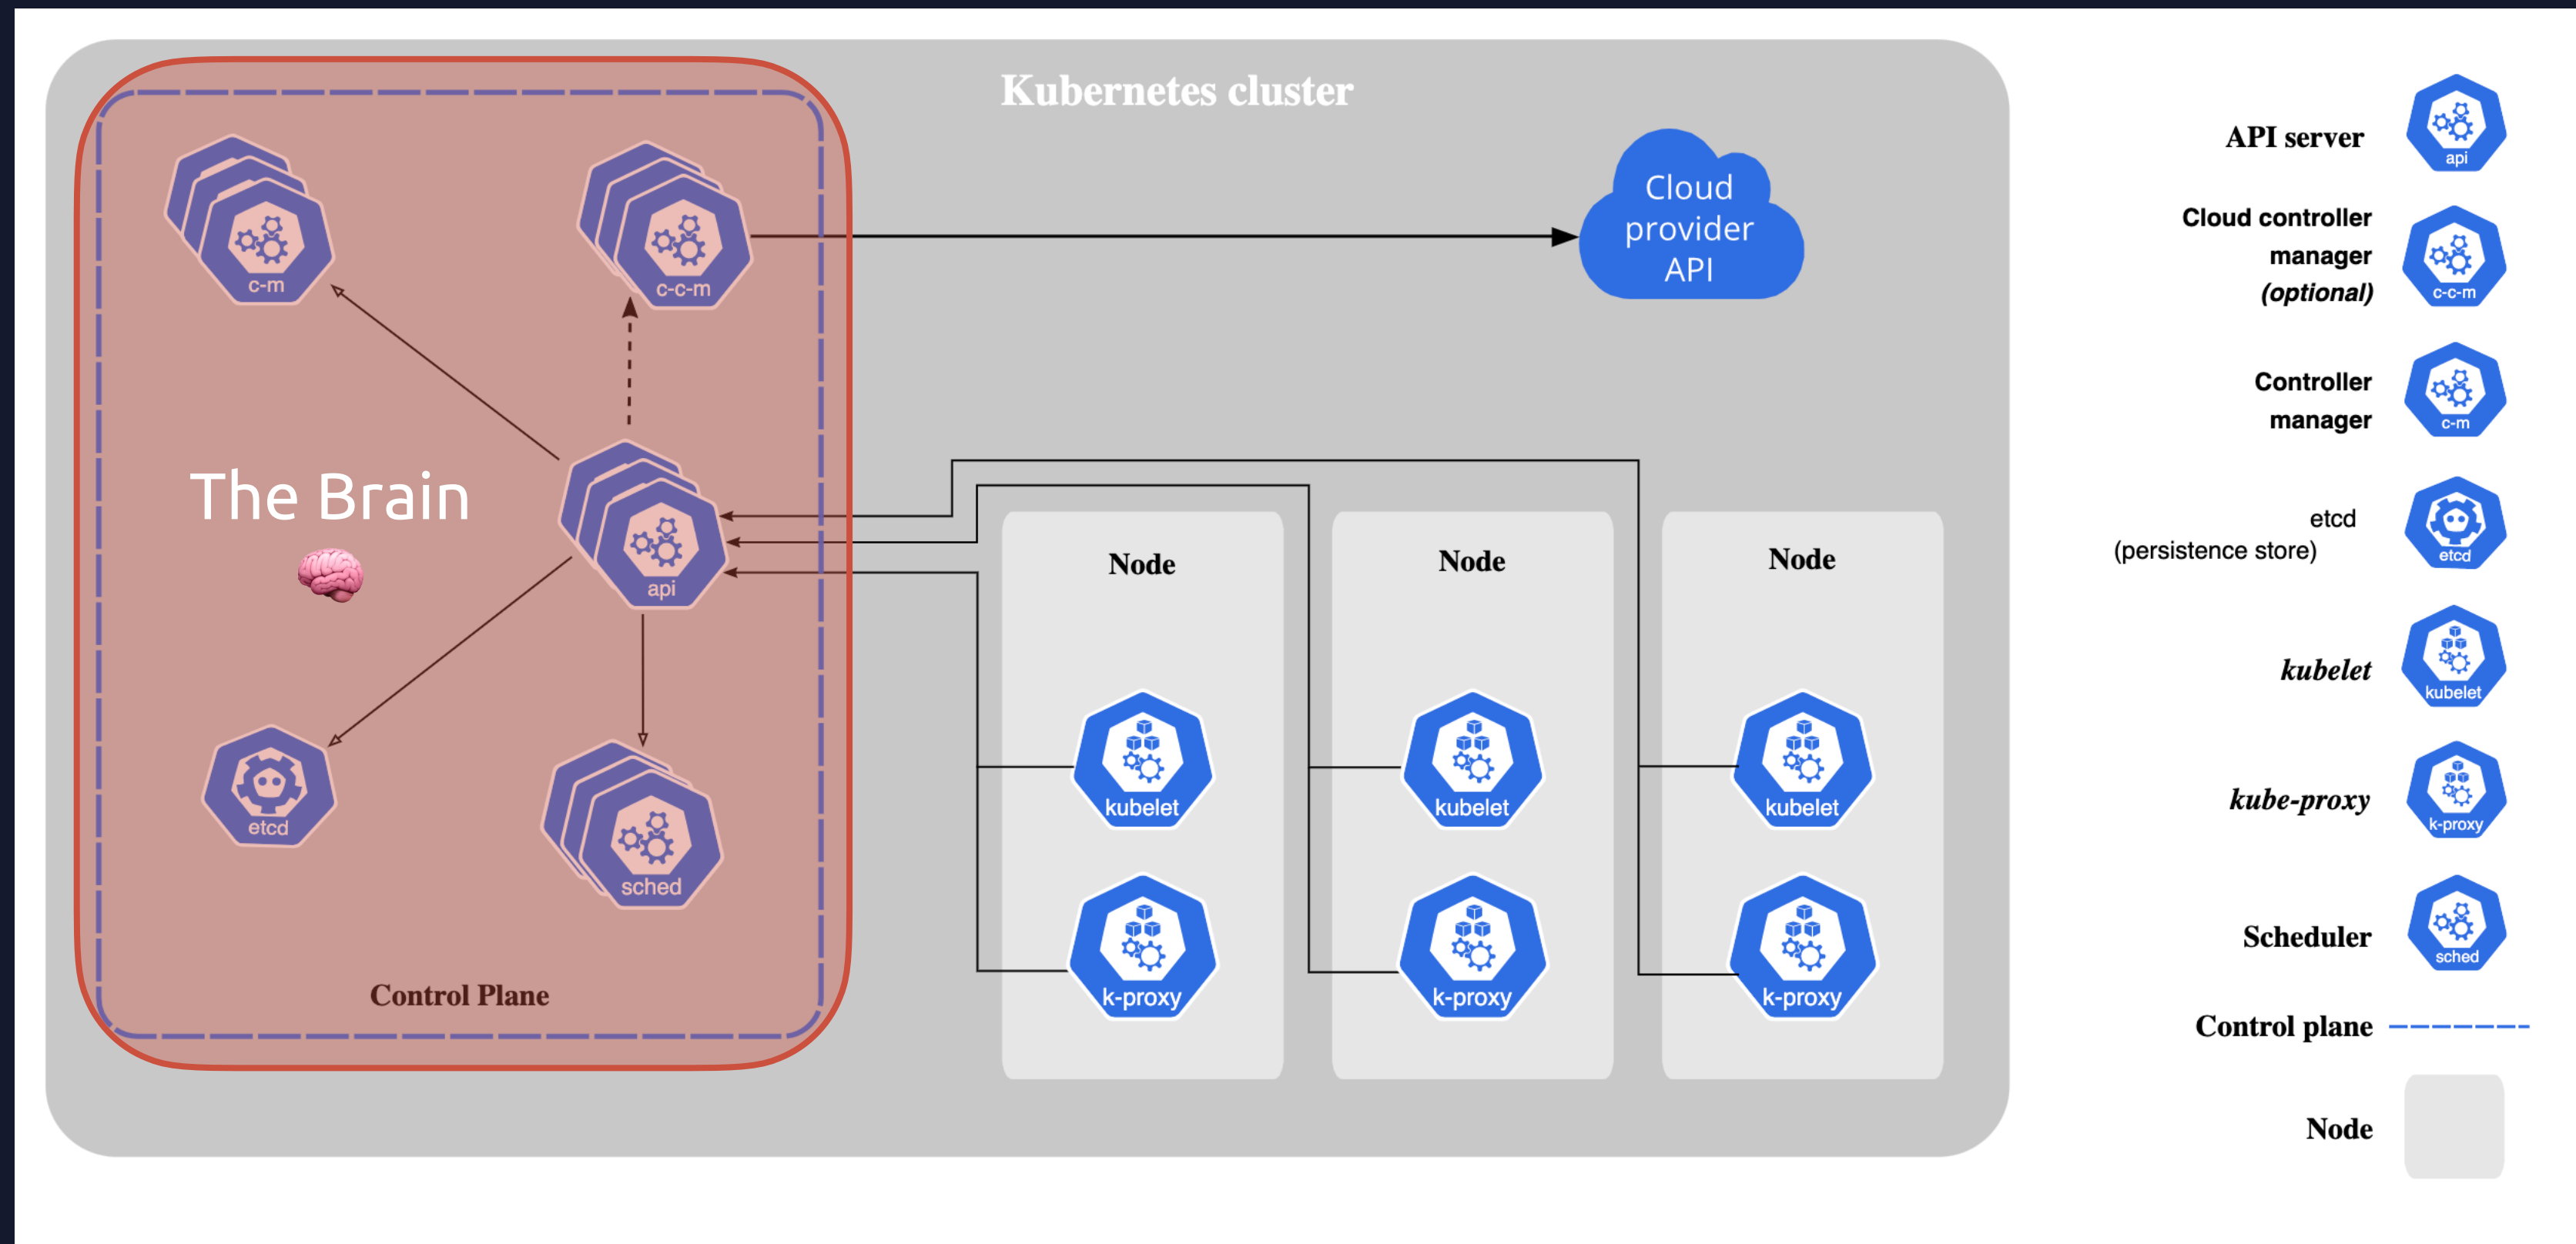
# KUBERNETES ARCHITECTURE



<https://kubernetes.io/docs/concepts/overview/components/>

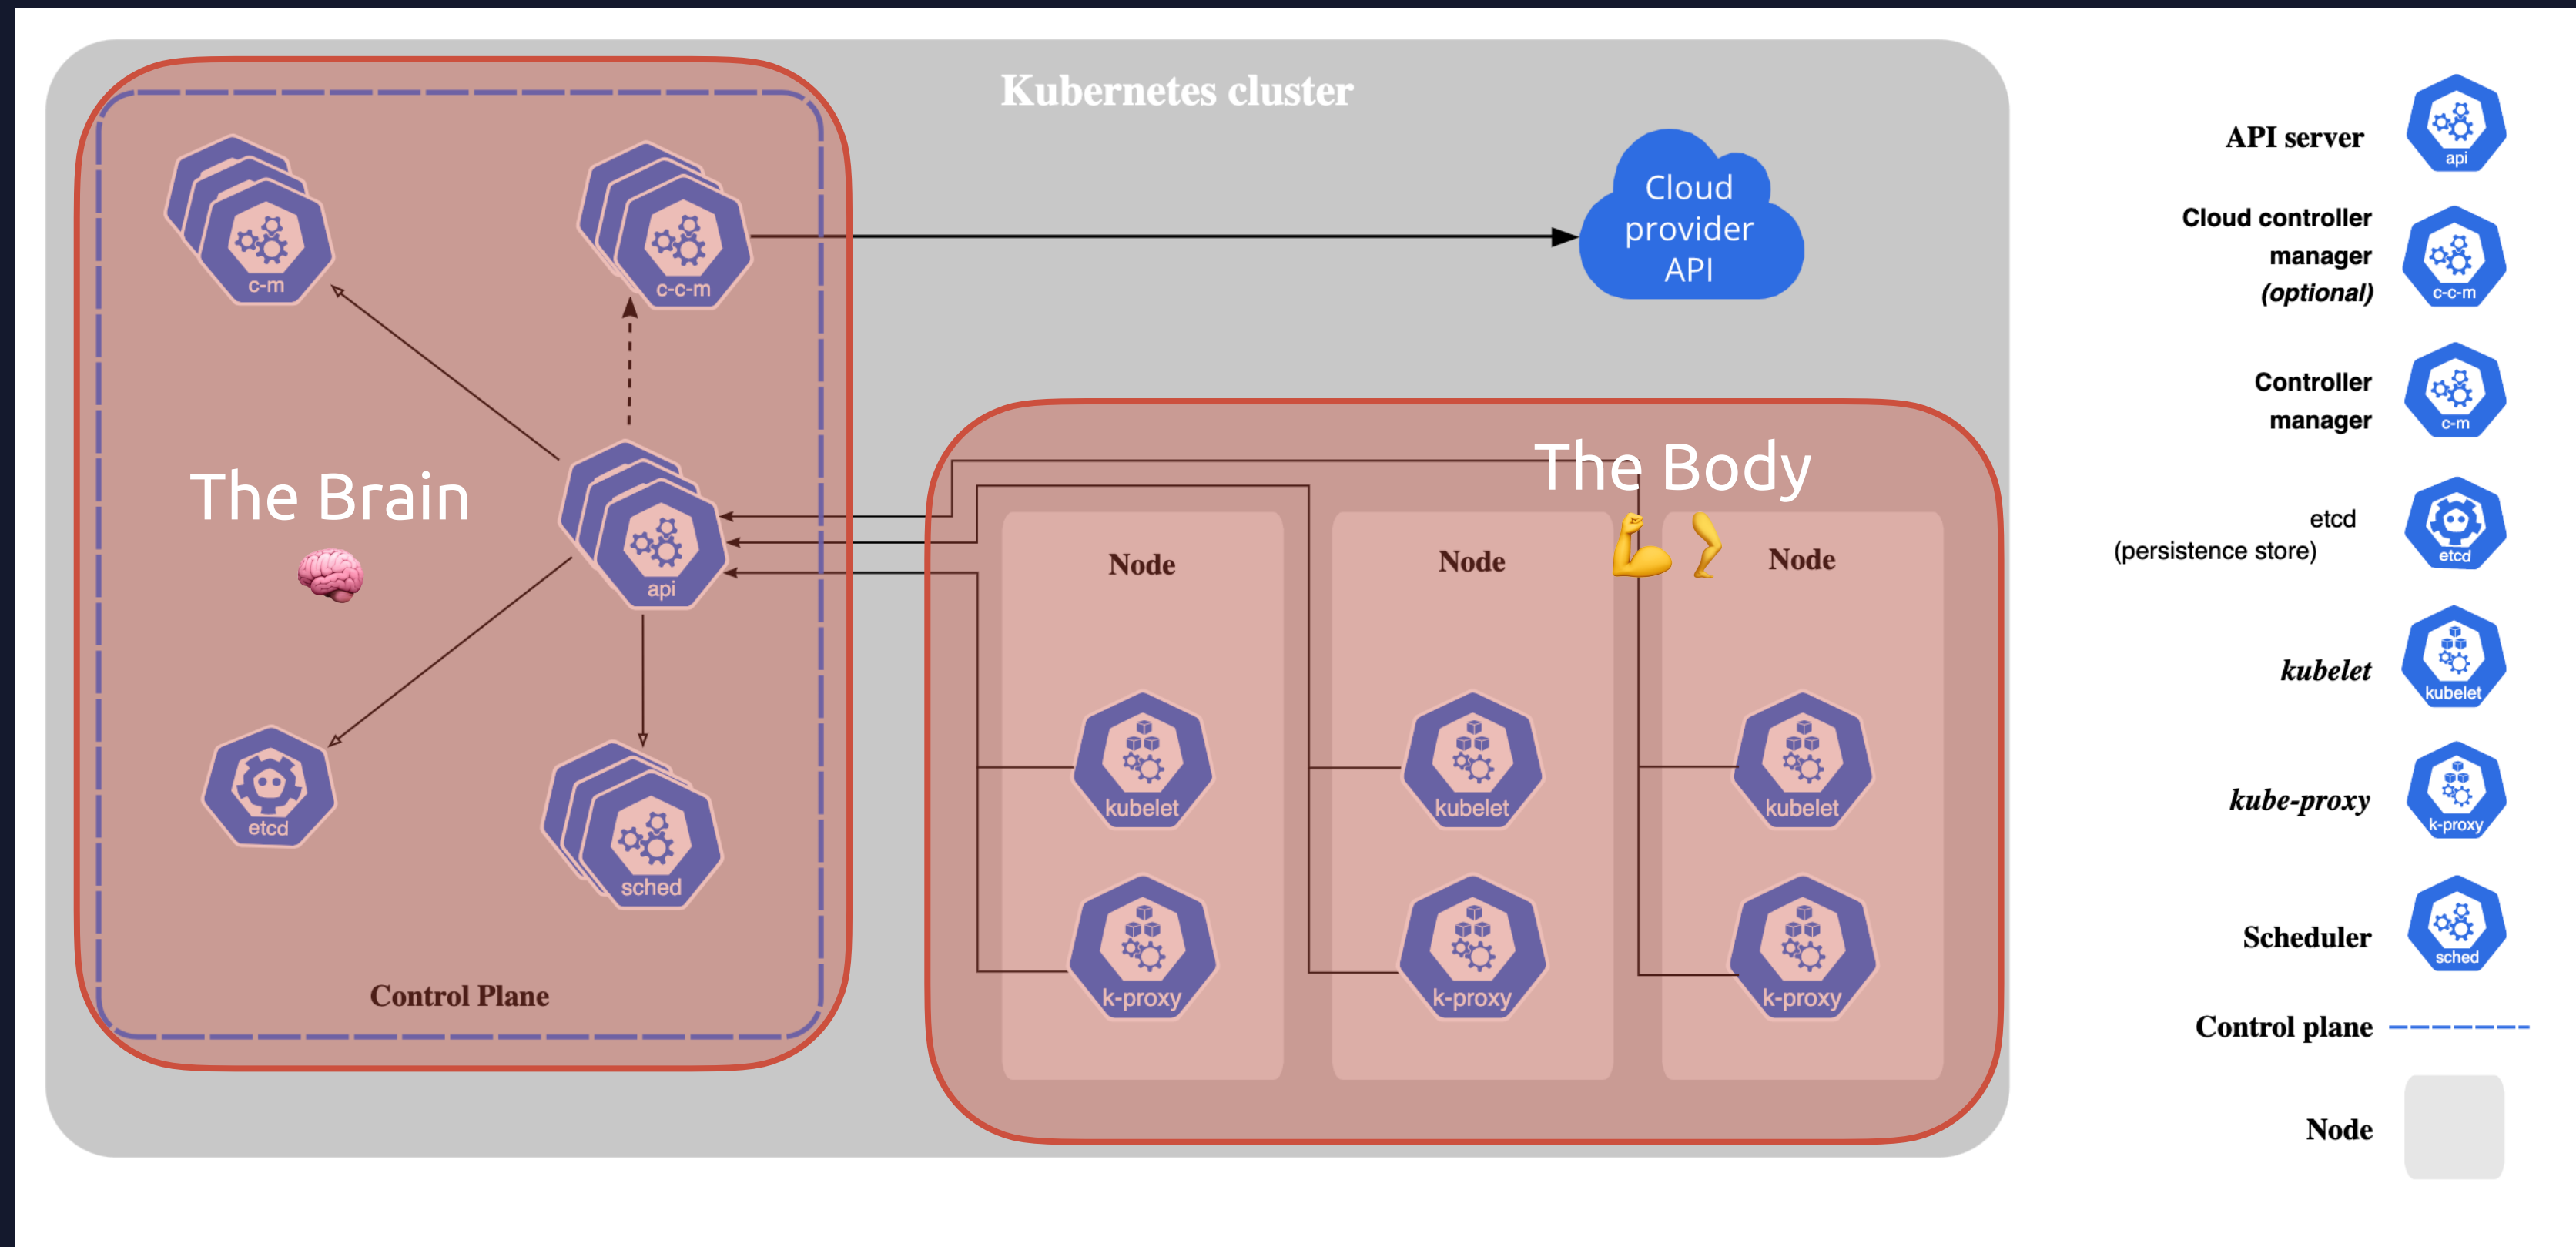


# KUBERNETES ARCHITECTURE



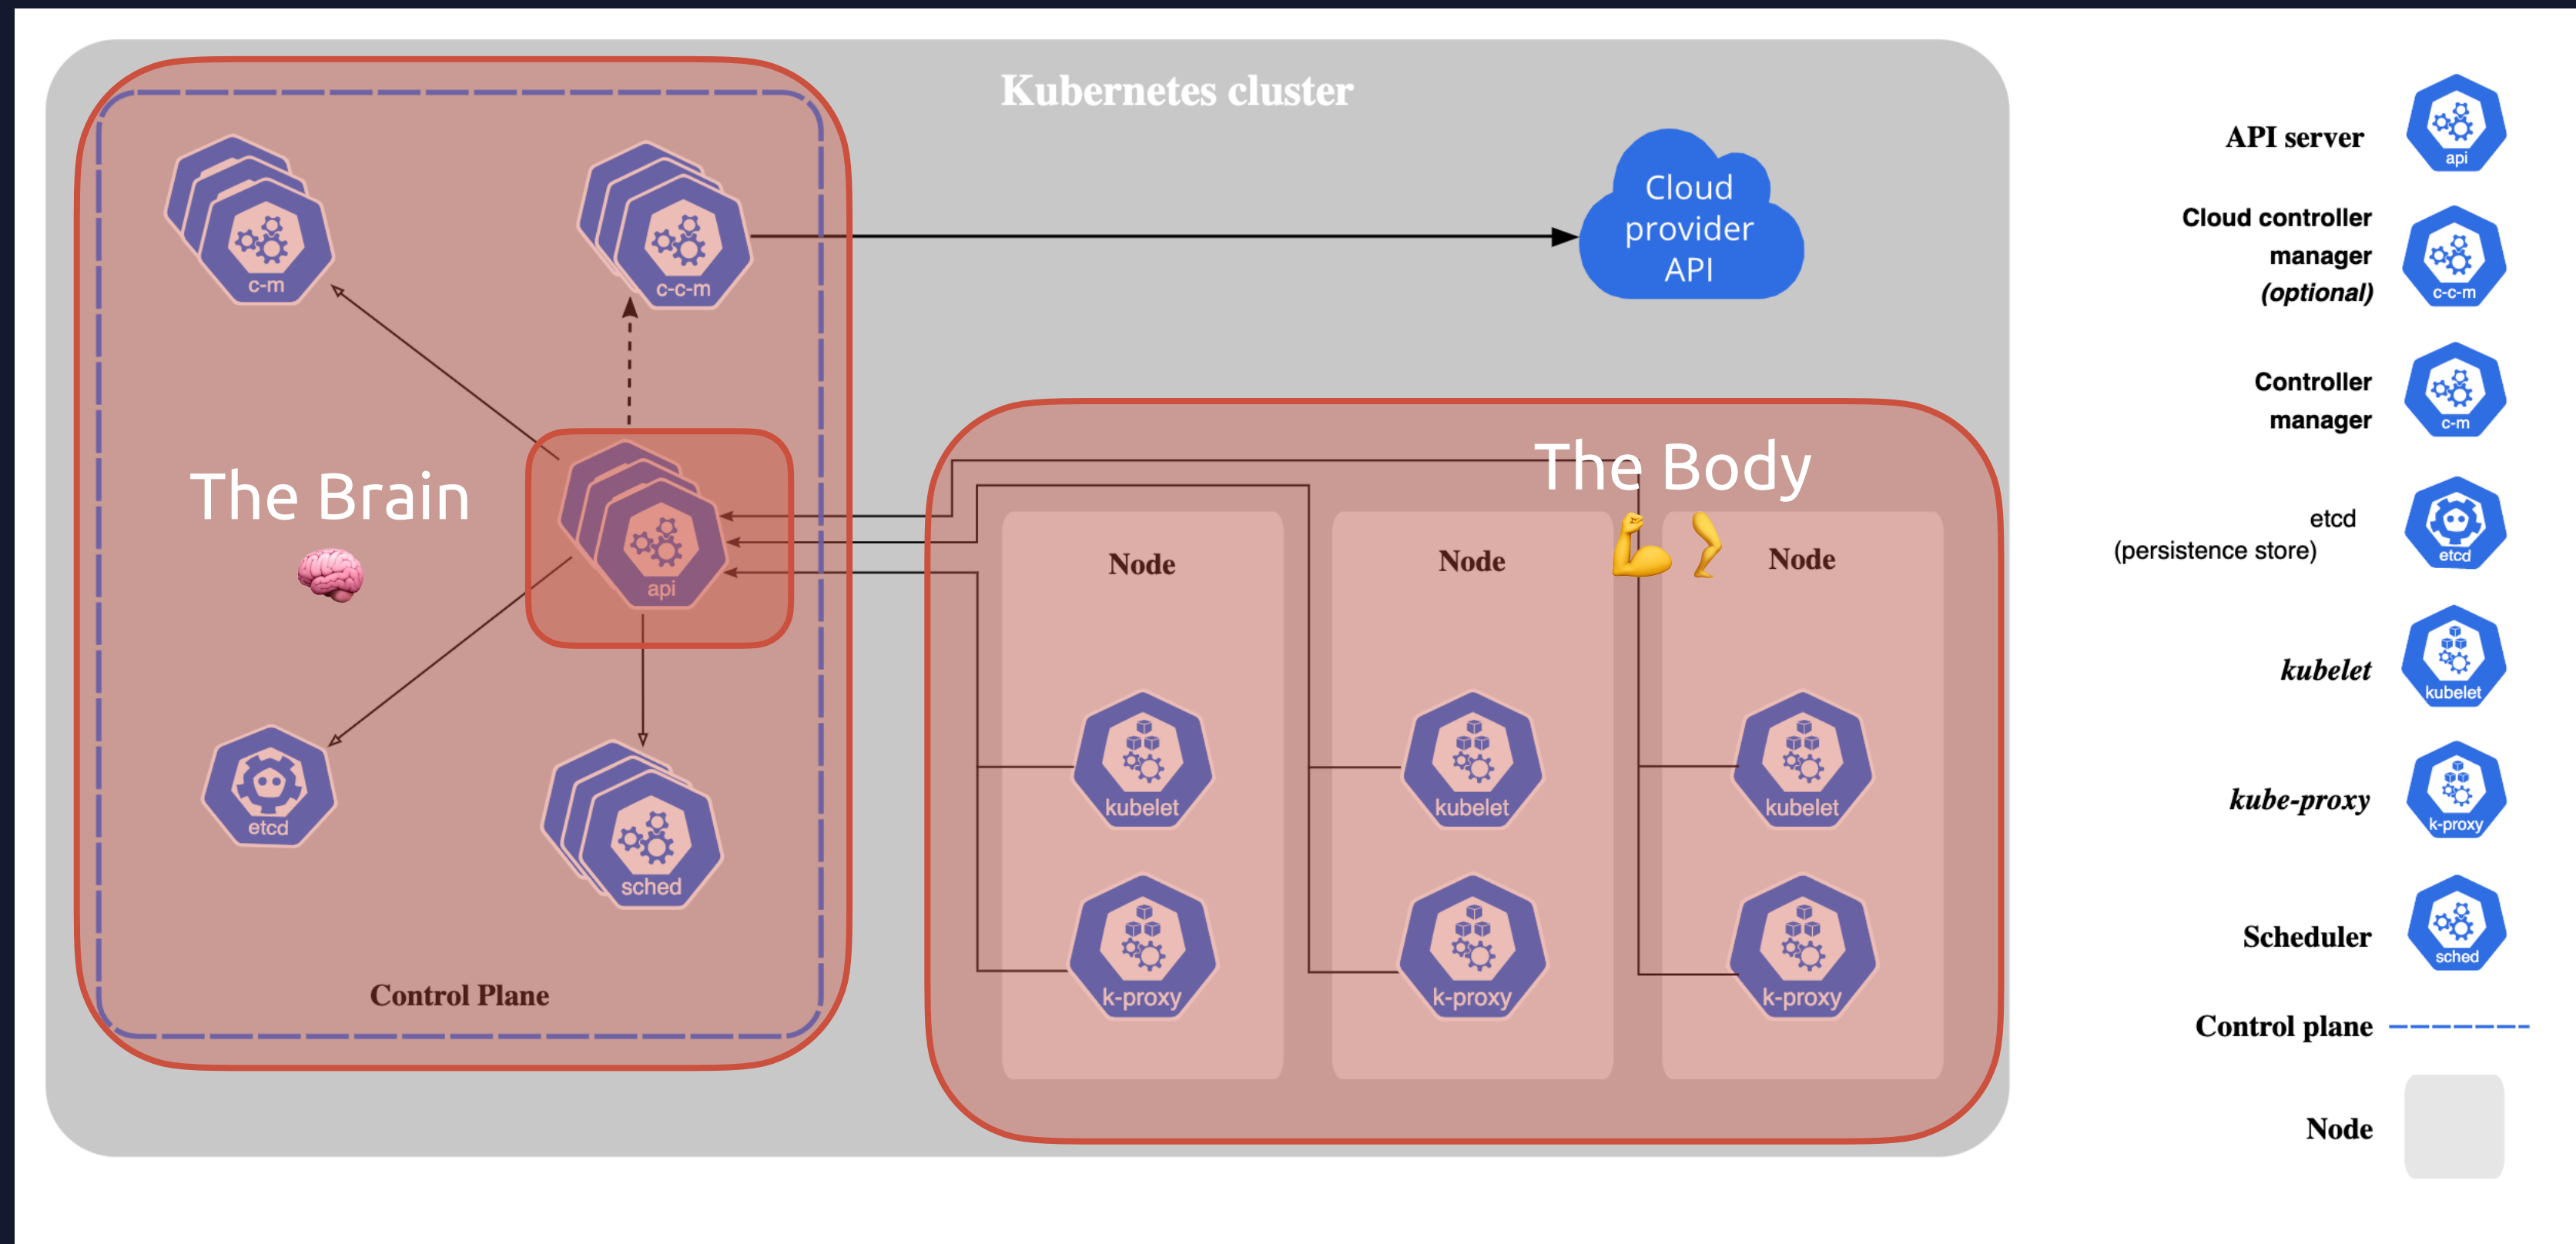
<https://kubernetes.io/docs/concepts/overview/components/>

# KUBERNETES ARCHITECTURE



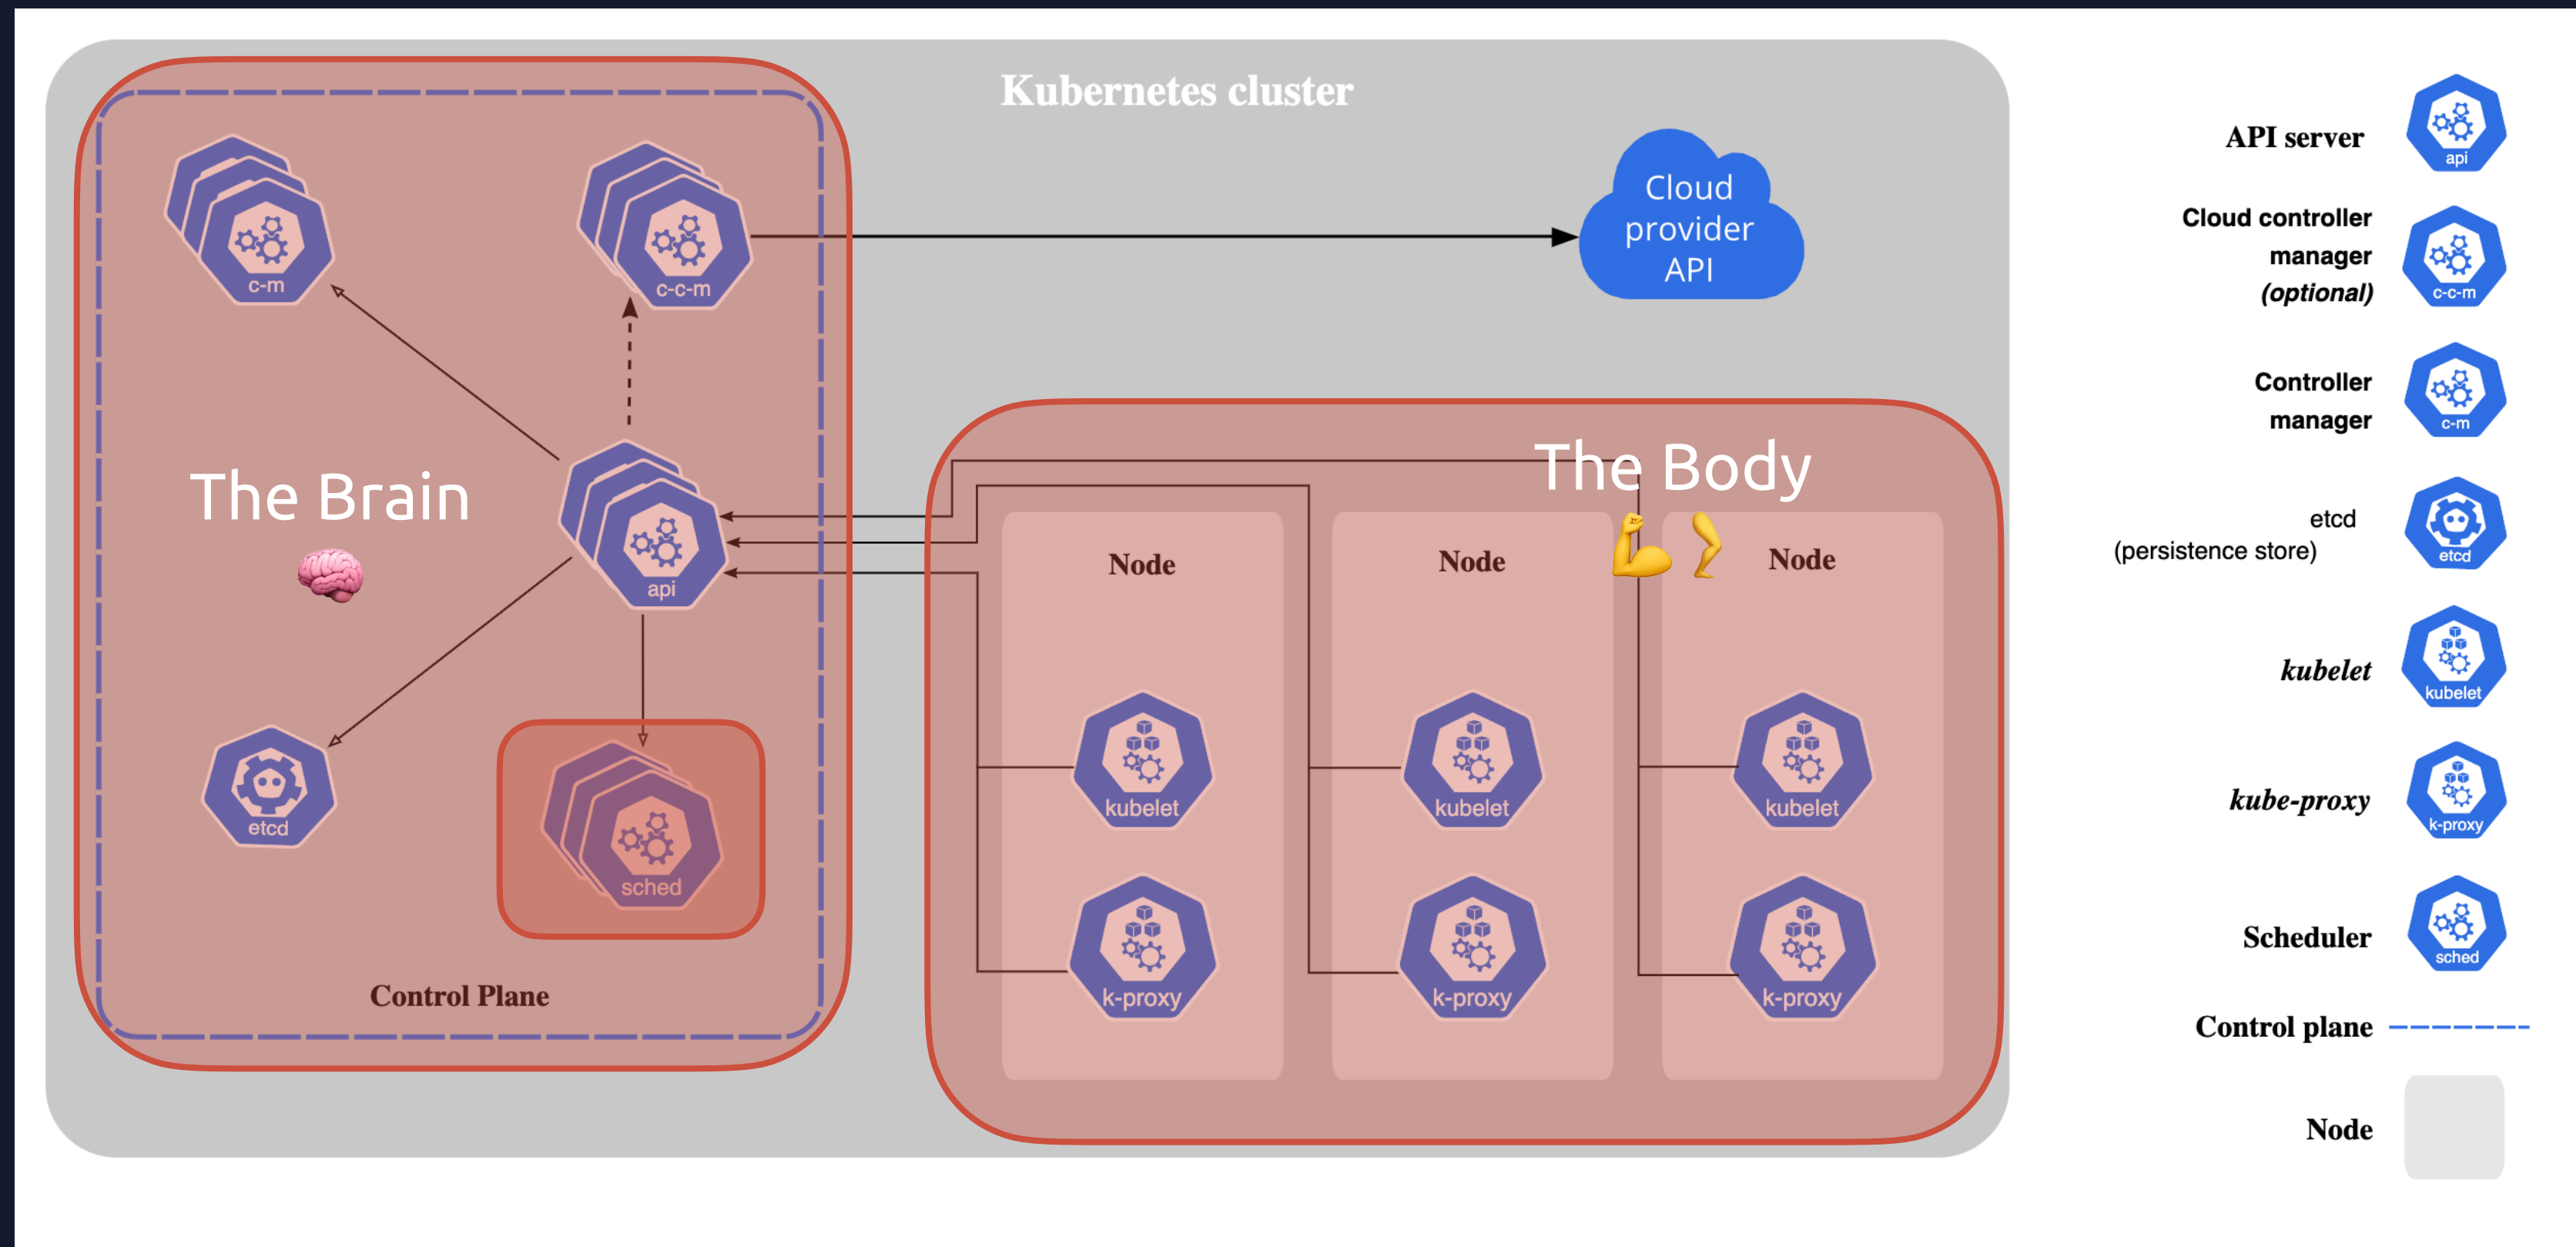
<https://kubernetes.io/docs/concepts/overview/components/>

# KUBERNETES ARCHITECTURE



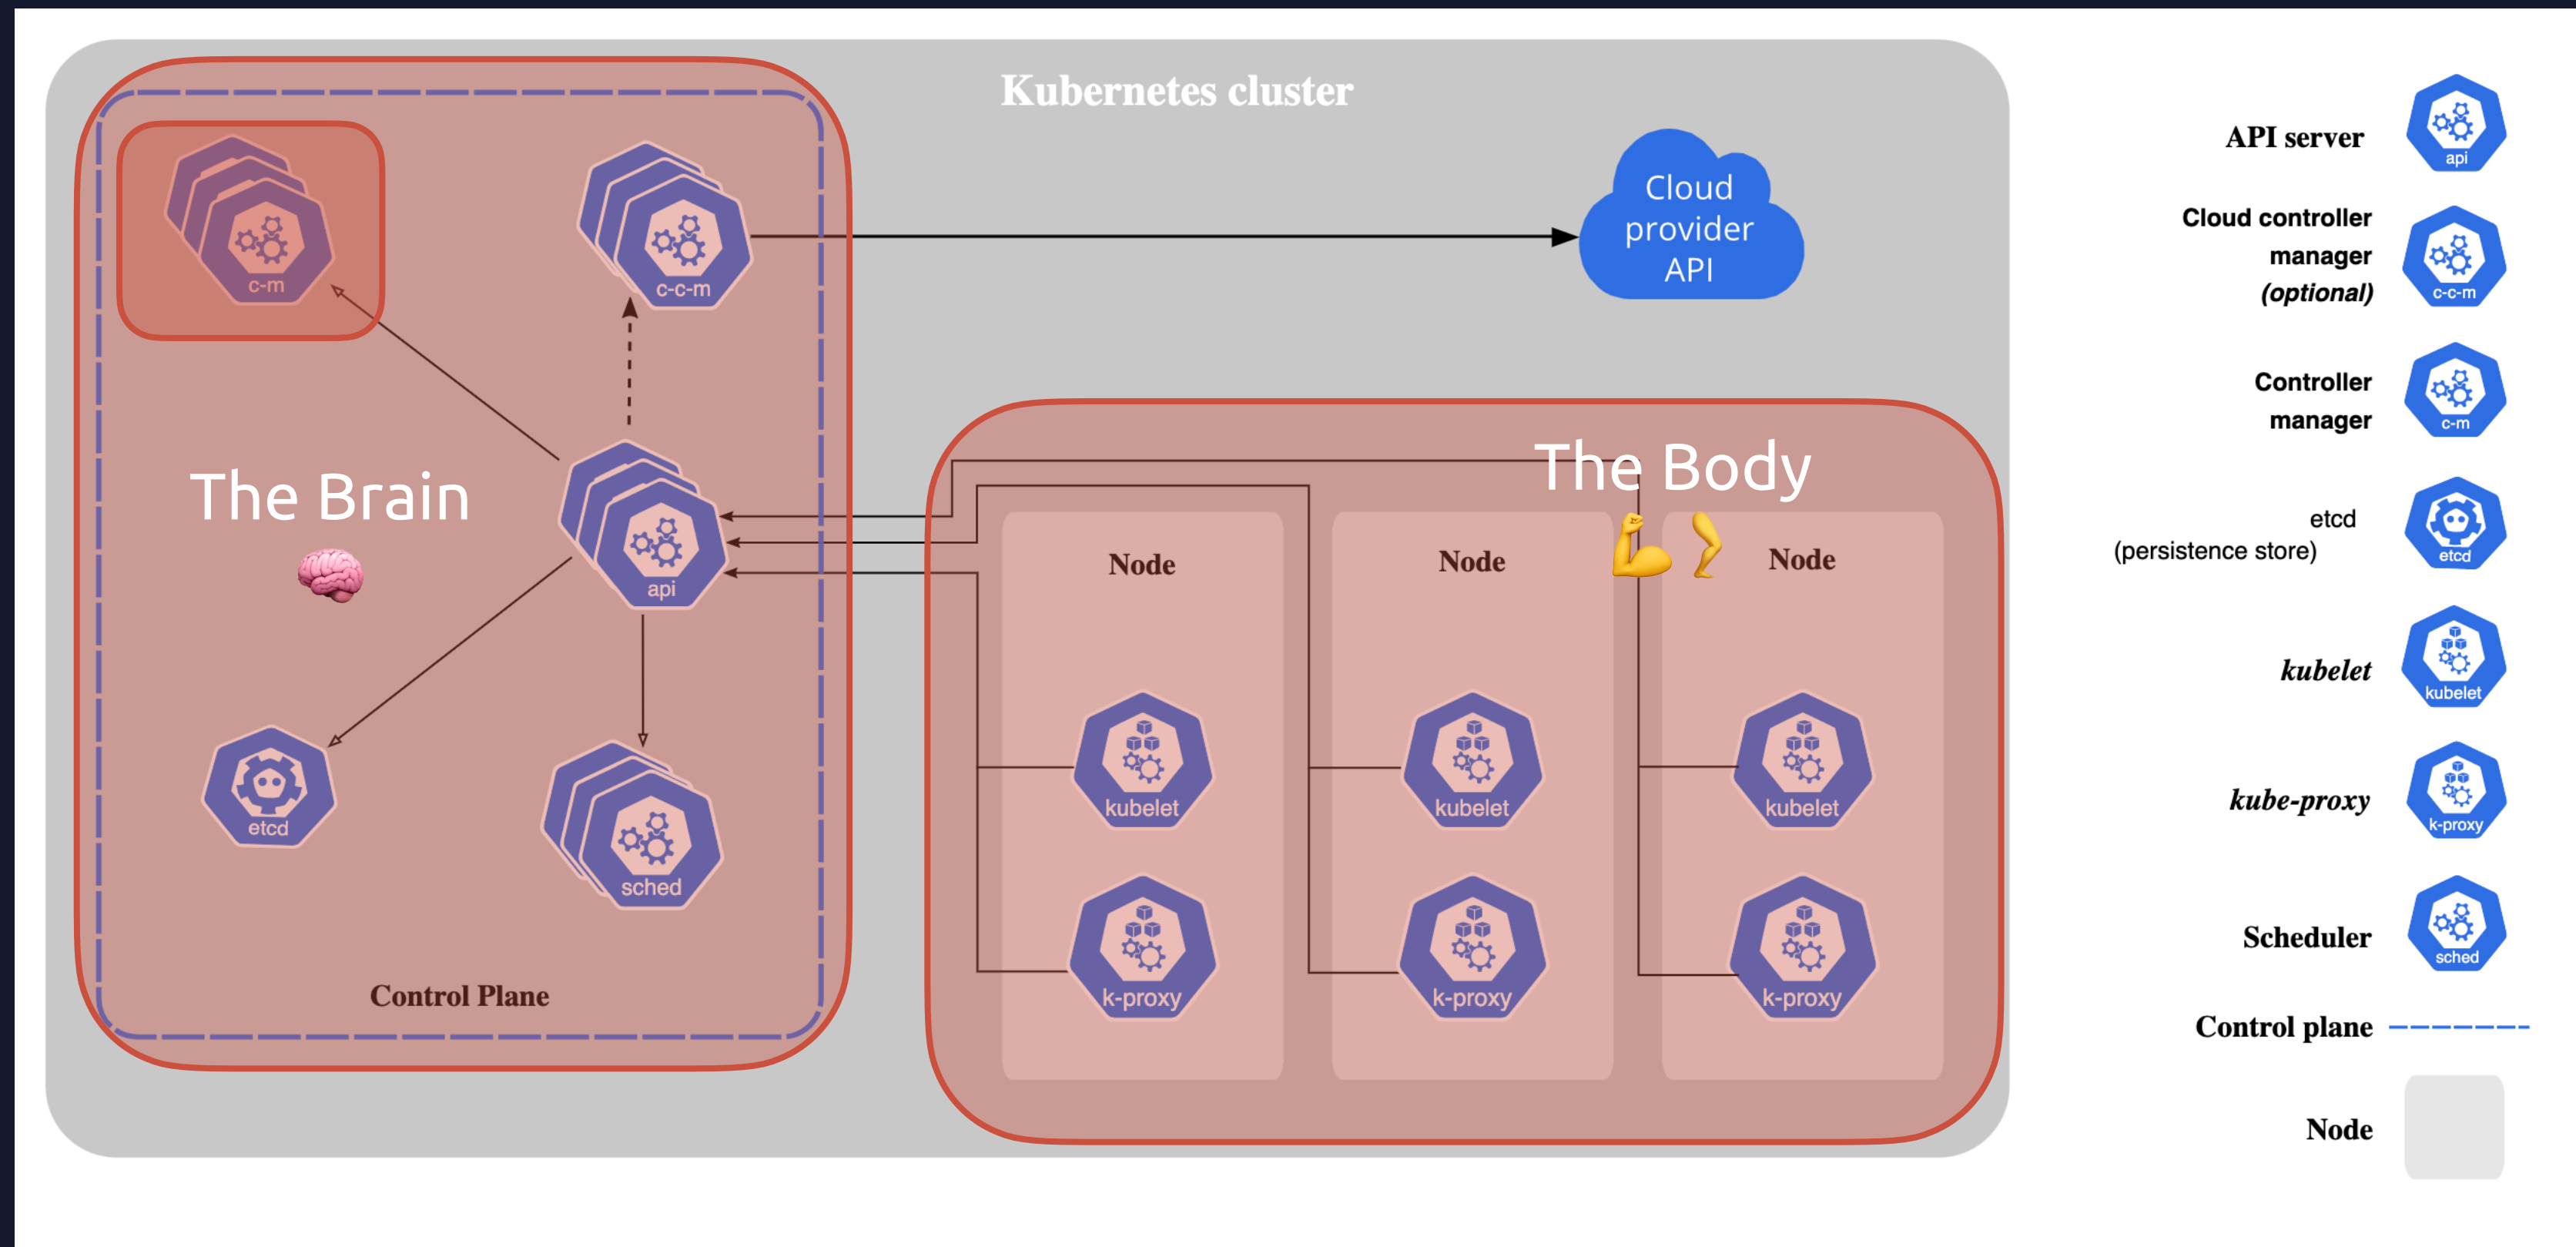
<https://kubernetes.io/docs/concepts/overview/components/>

# KUBERNETES ARCHITECTURE



<https://kubernetes.io/docs/concepts/overview/components/>

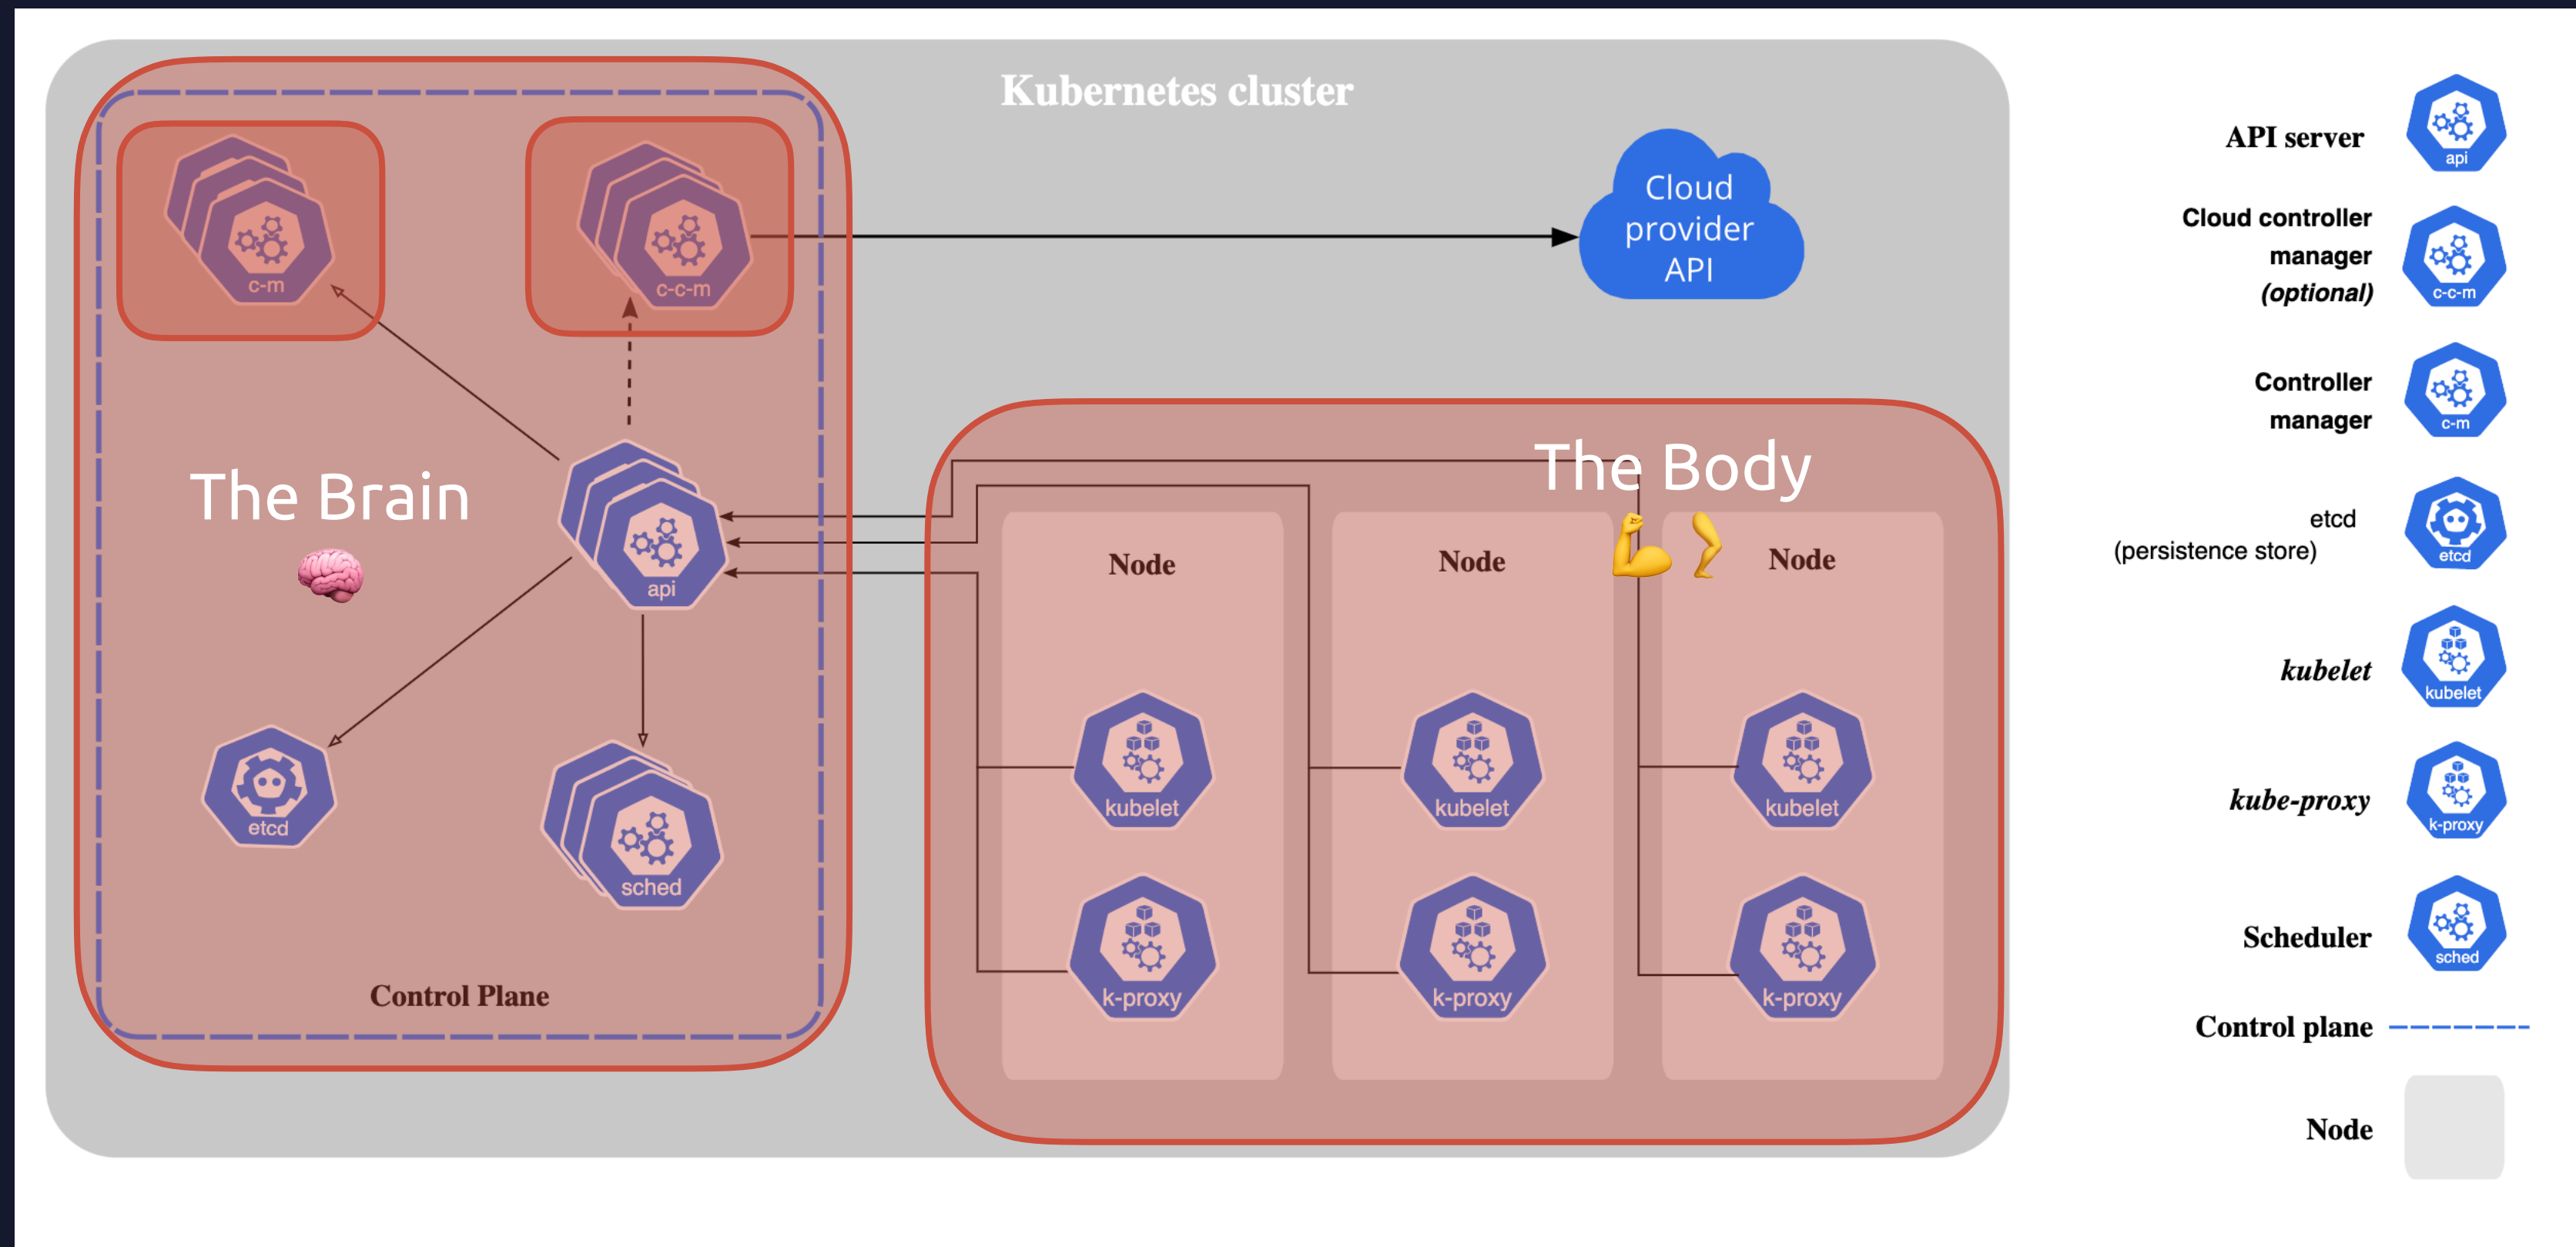
# KUBERNETES ARCHITECTURE



<https://kubernetes.io/docs/concepts/overview/components/>

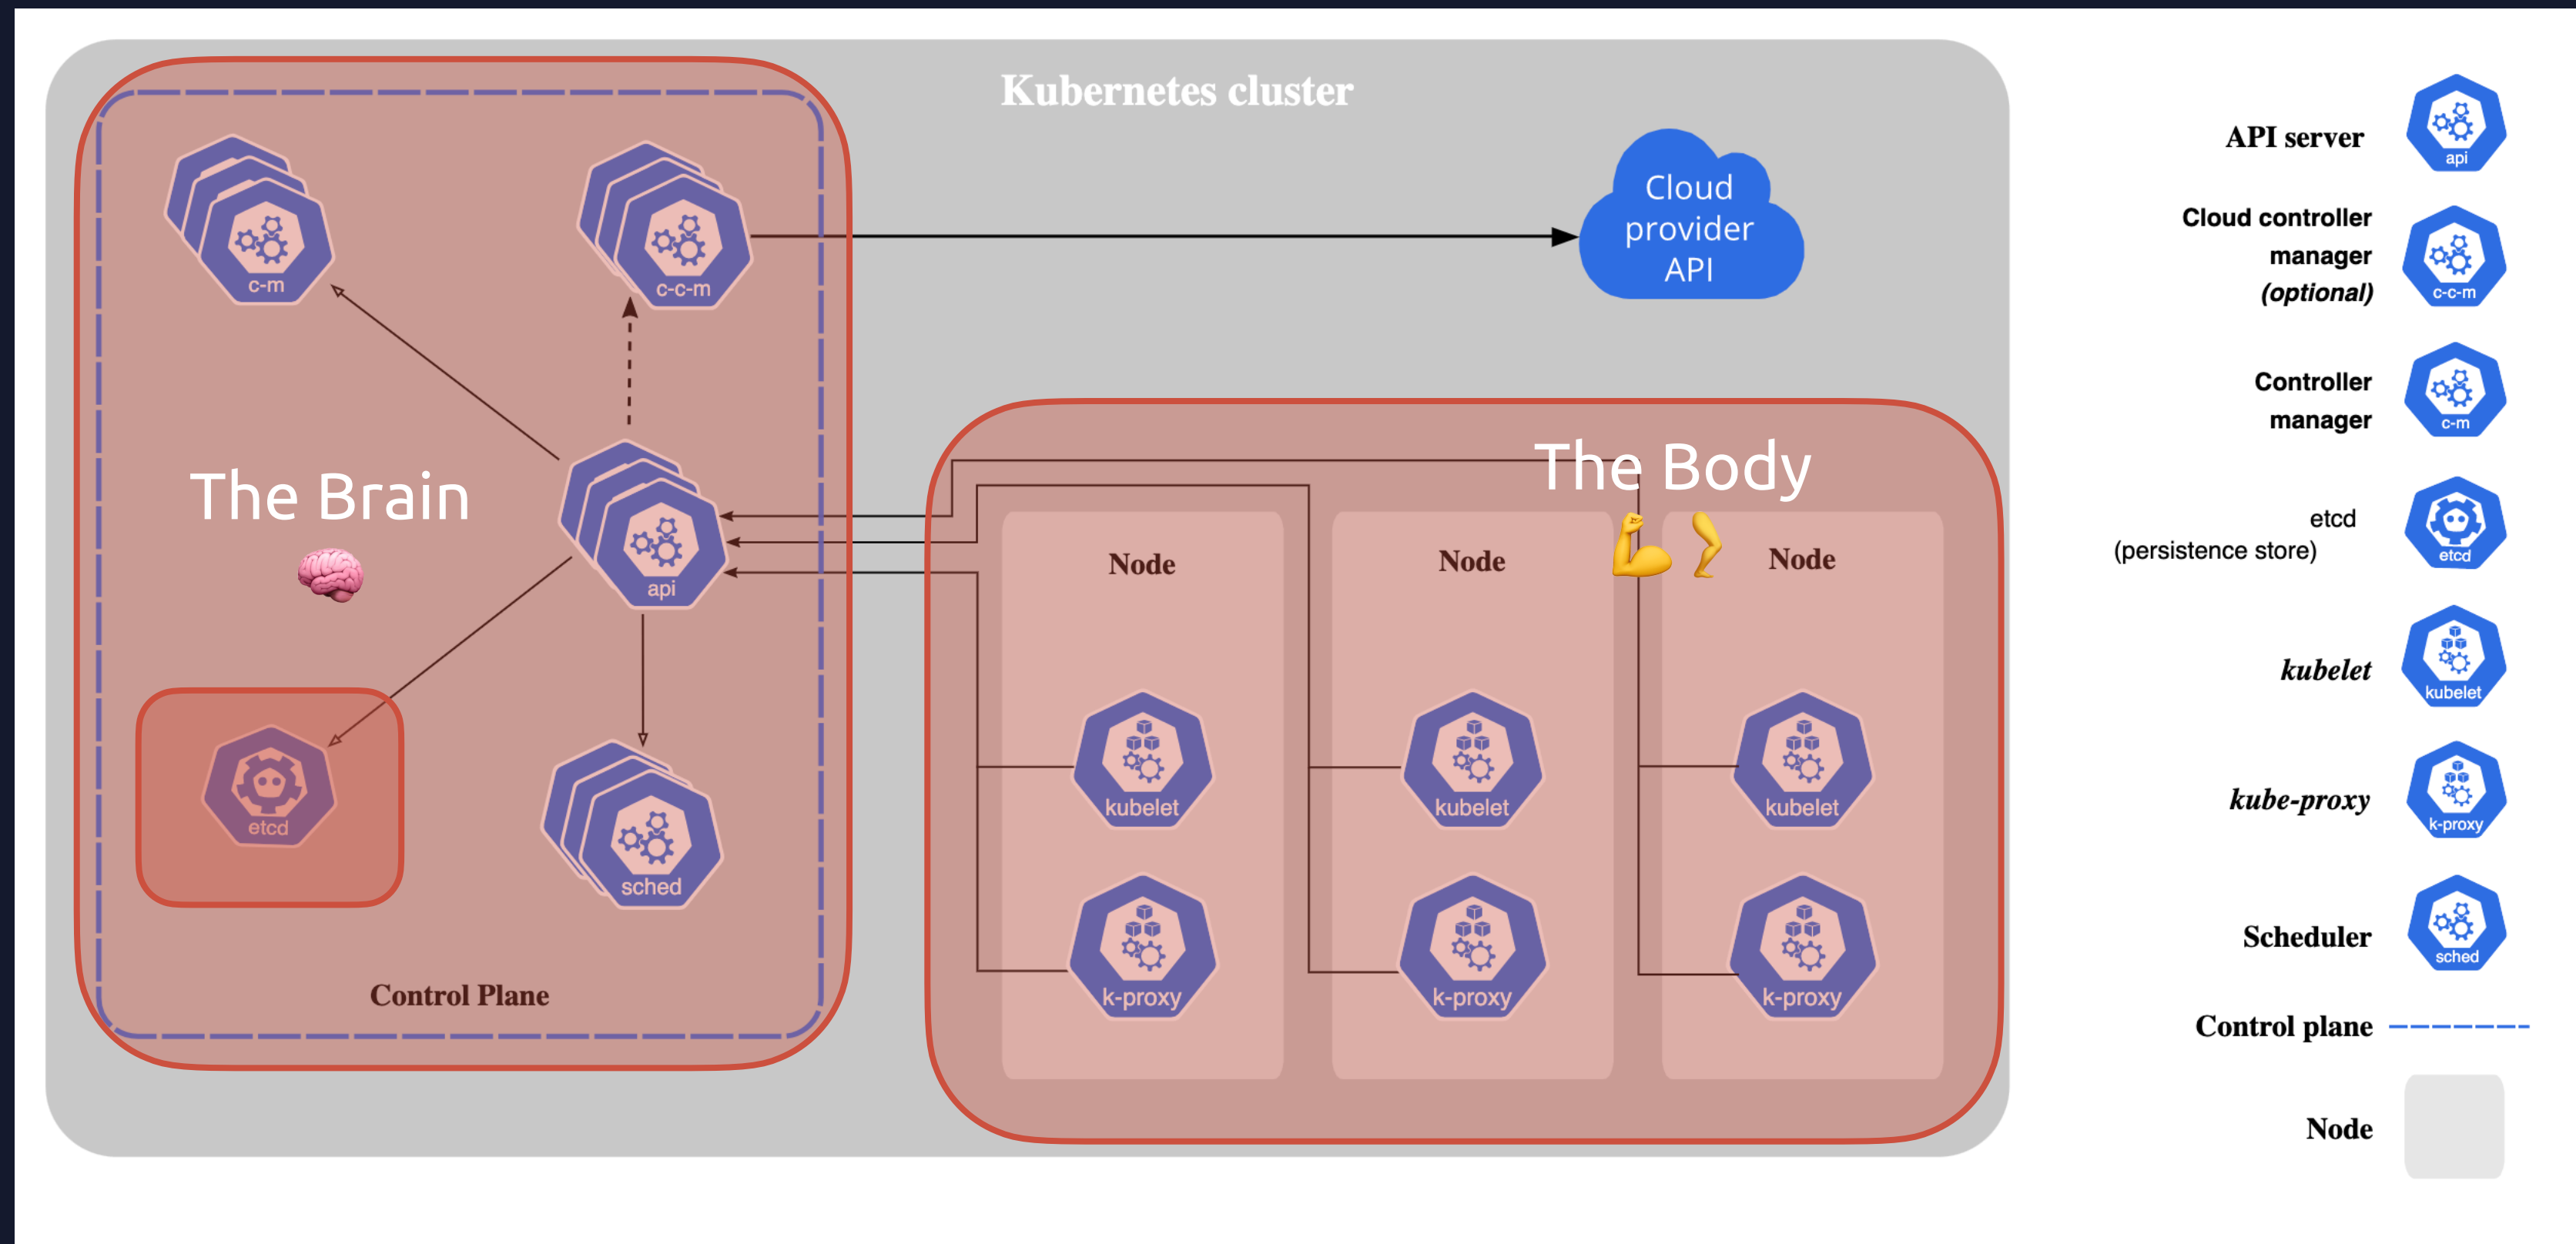


# KUBERNETES ARCHITECTURE



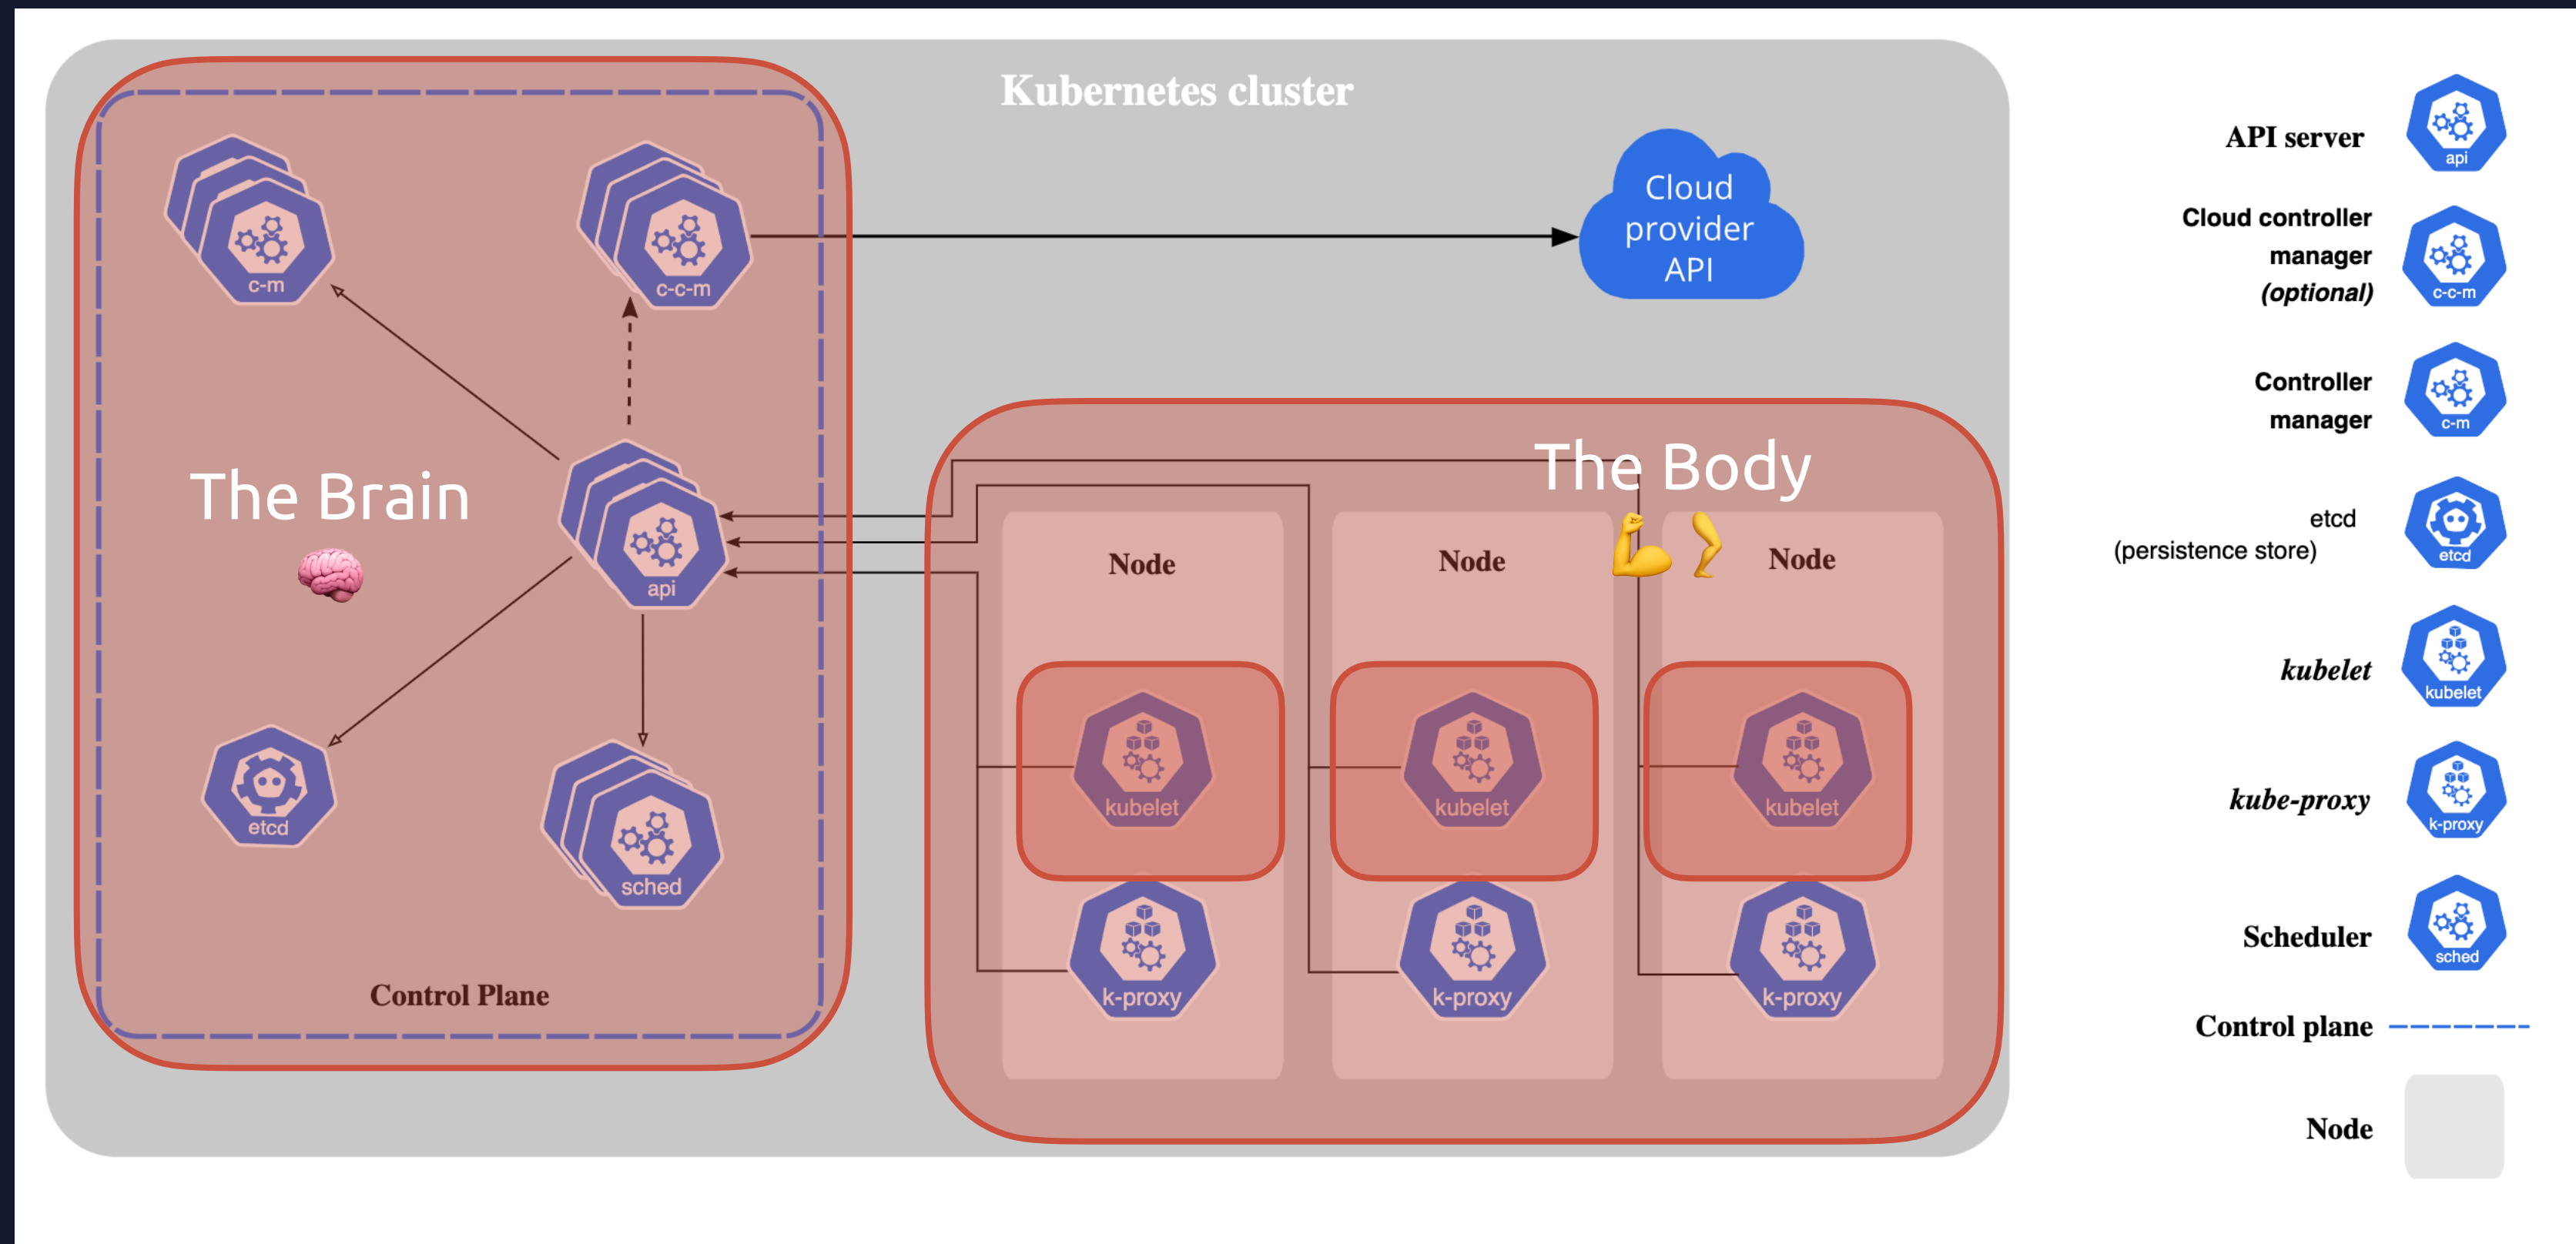
<https://kubernetes.io/docs/concepts/overview/components/>

# KUBERNETES ARCHITECTURE



<https://kubernetes.io/docs/concepts/overview/components/>

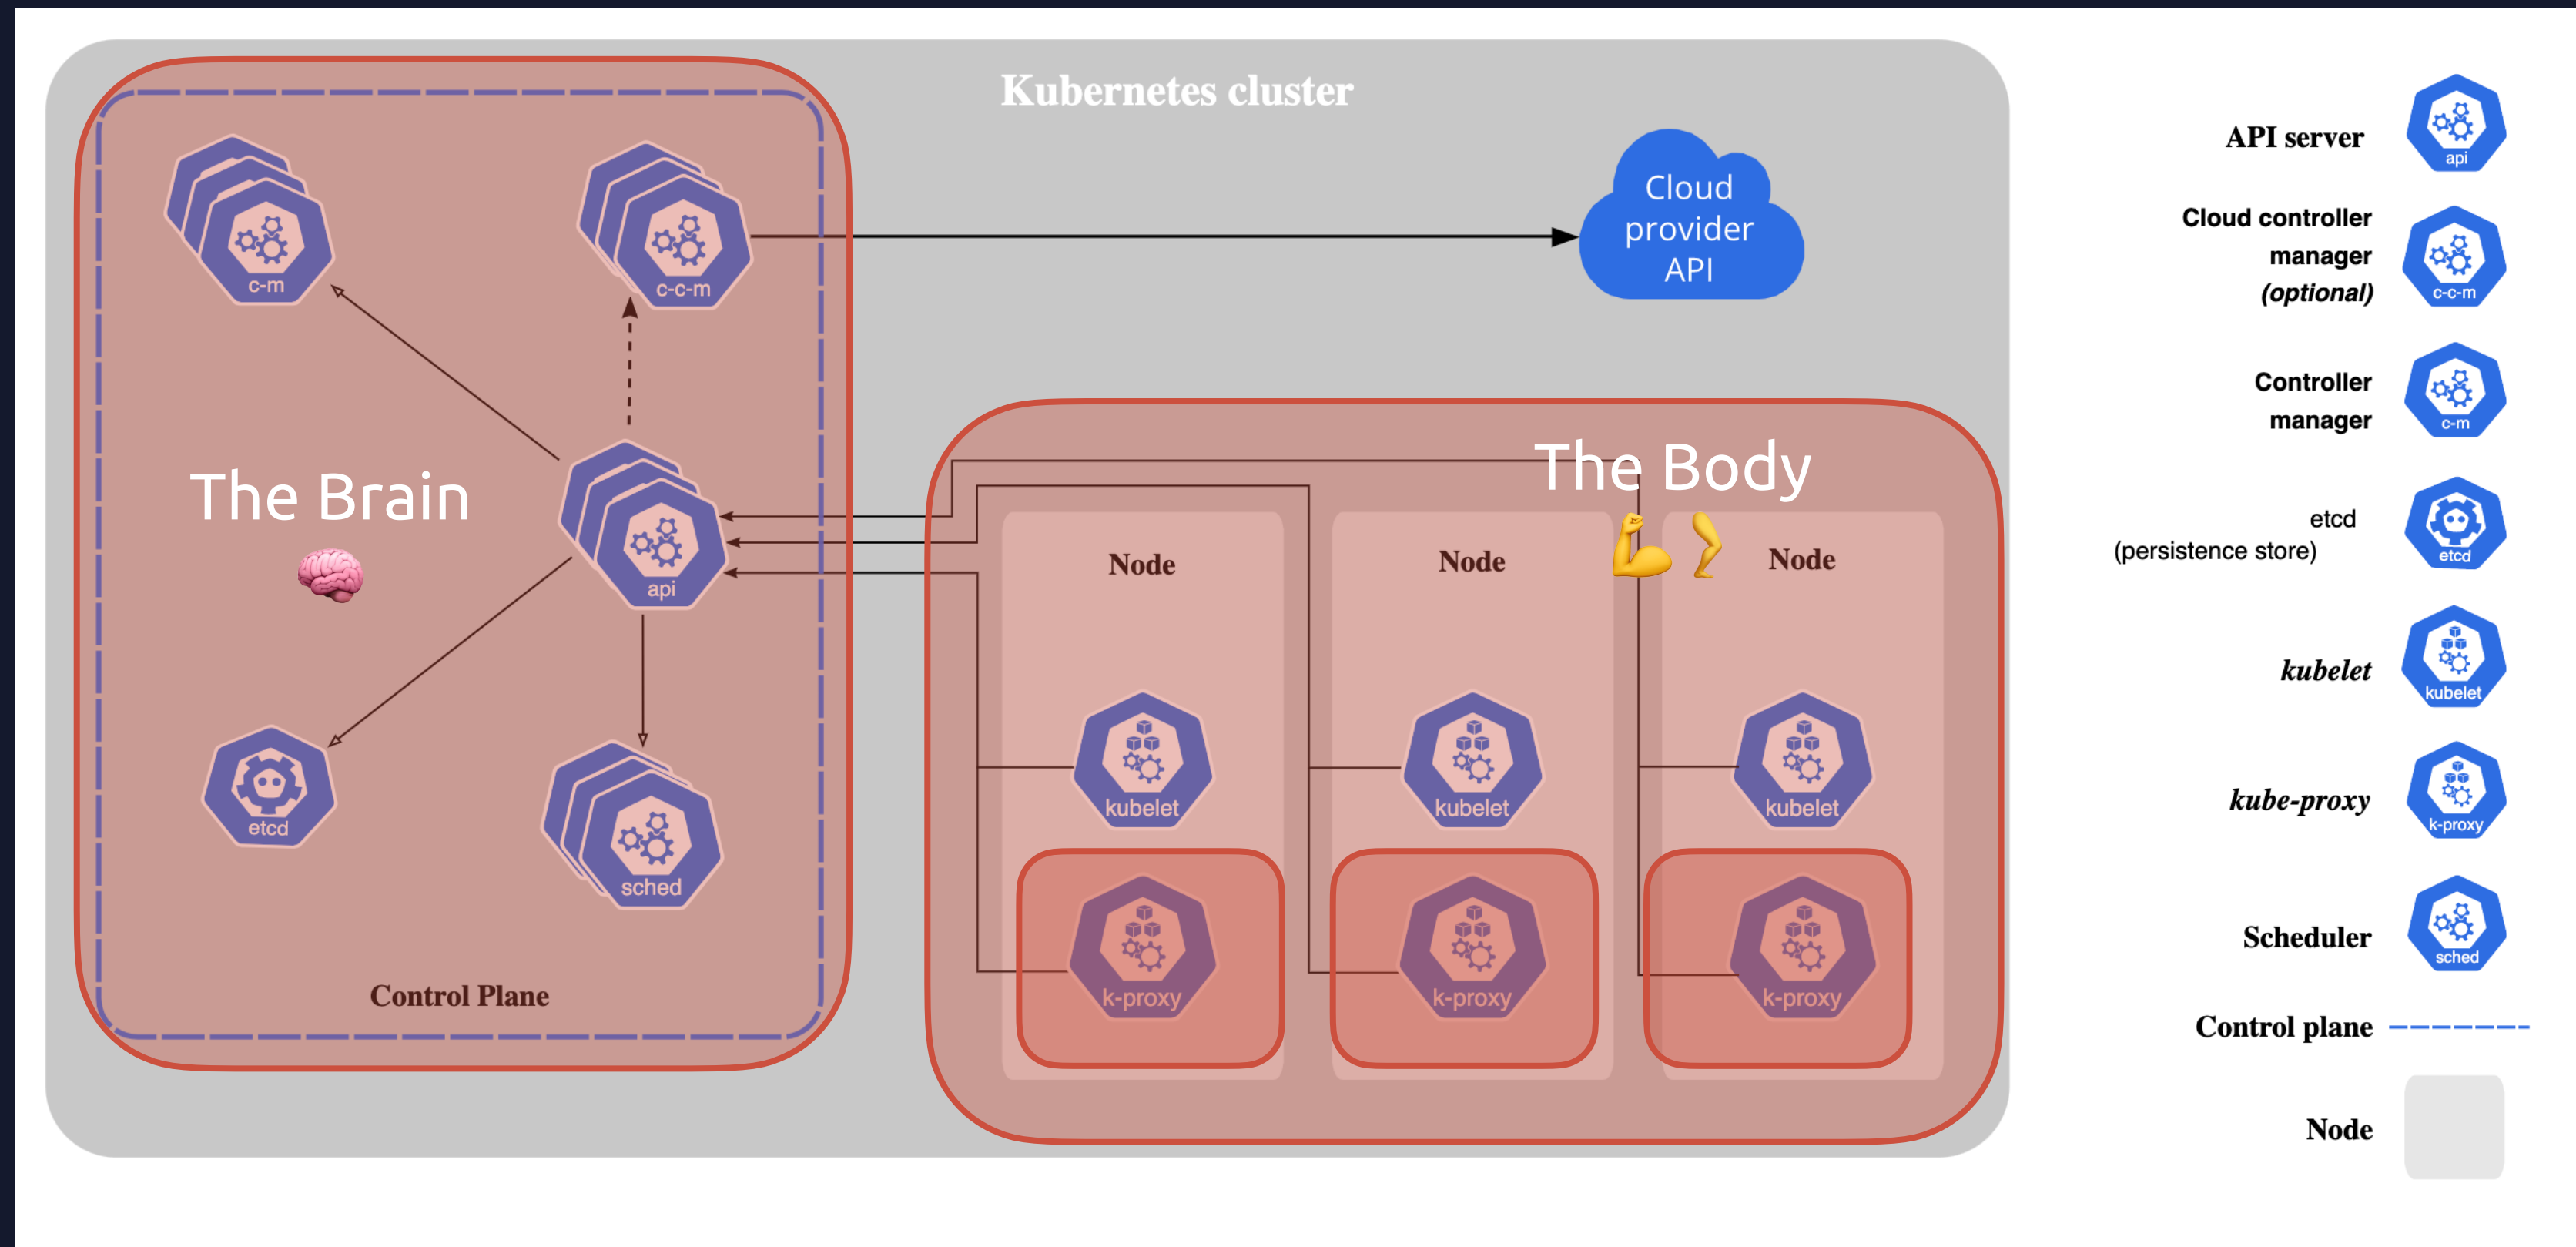
# KUBERNETES ARCHITECTURE



<https://kubernetes.io/docs/concepts/overview/components/>

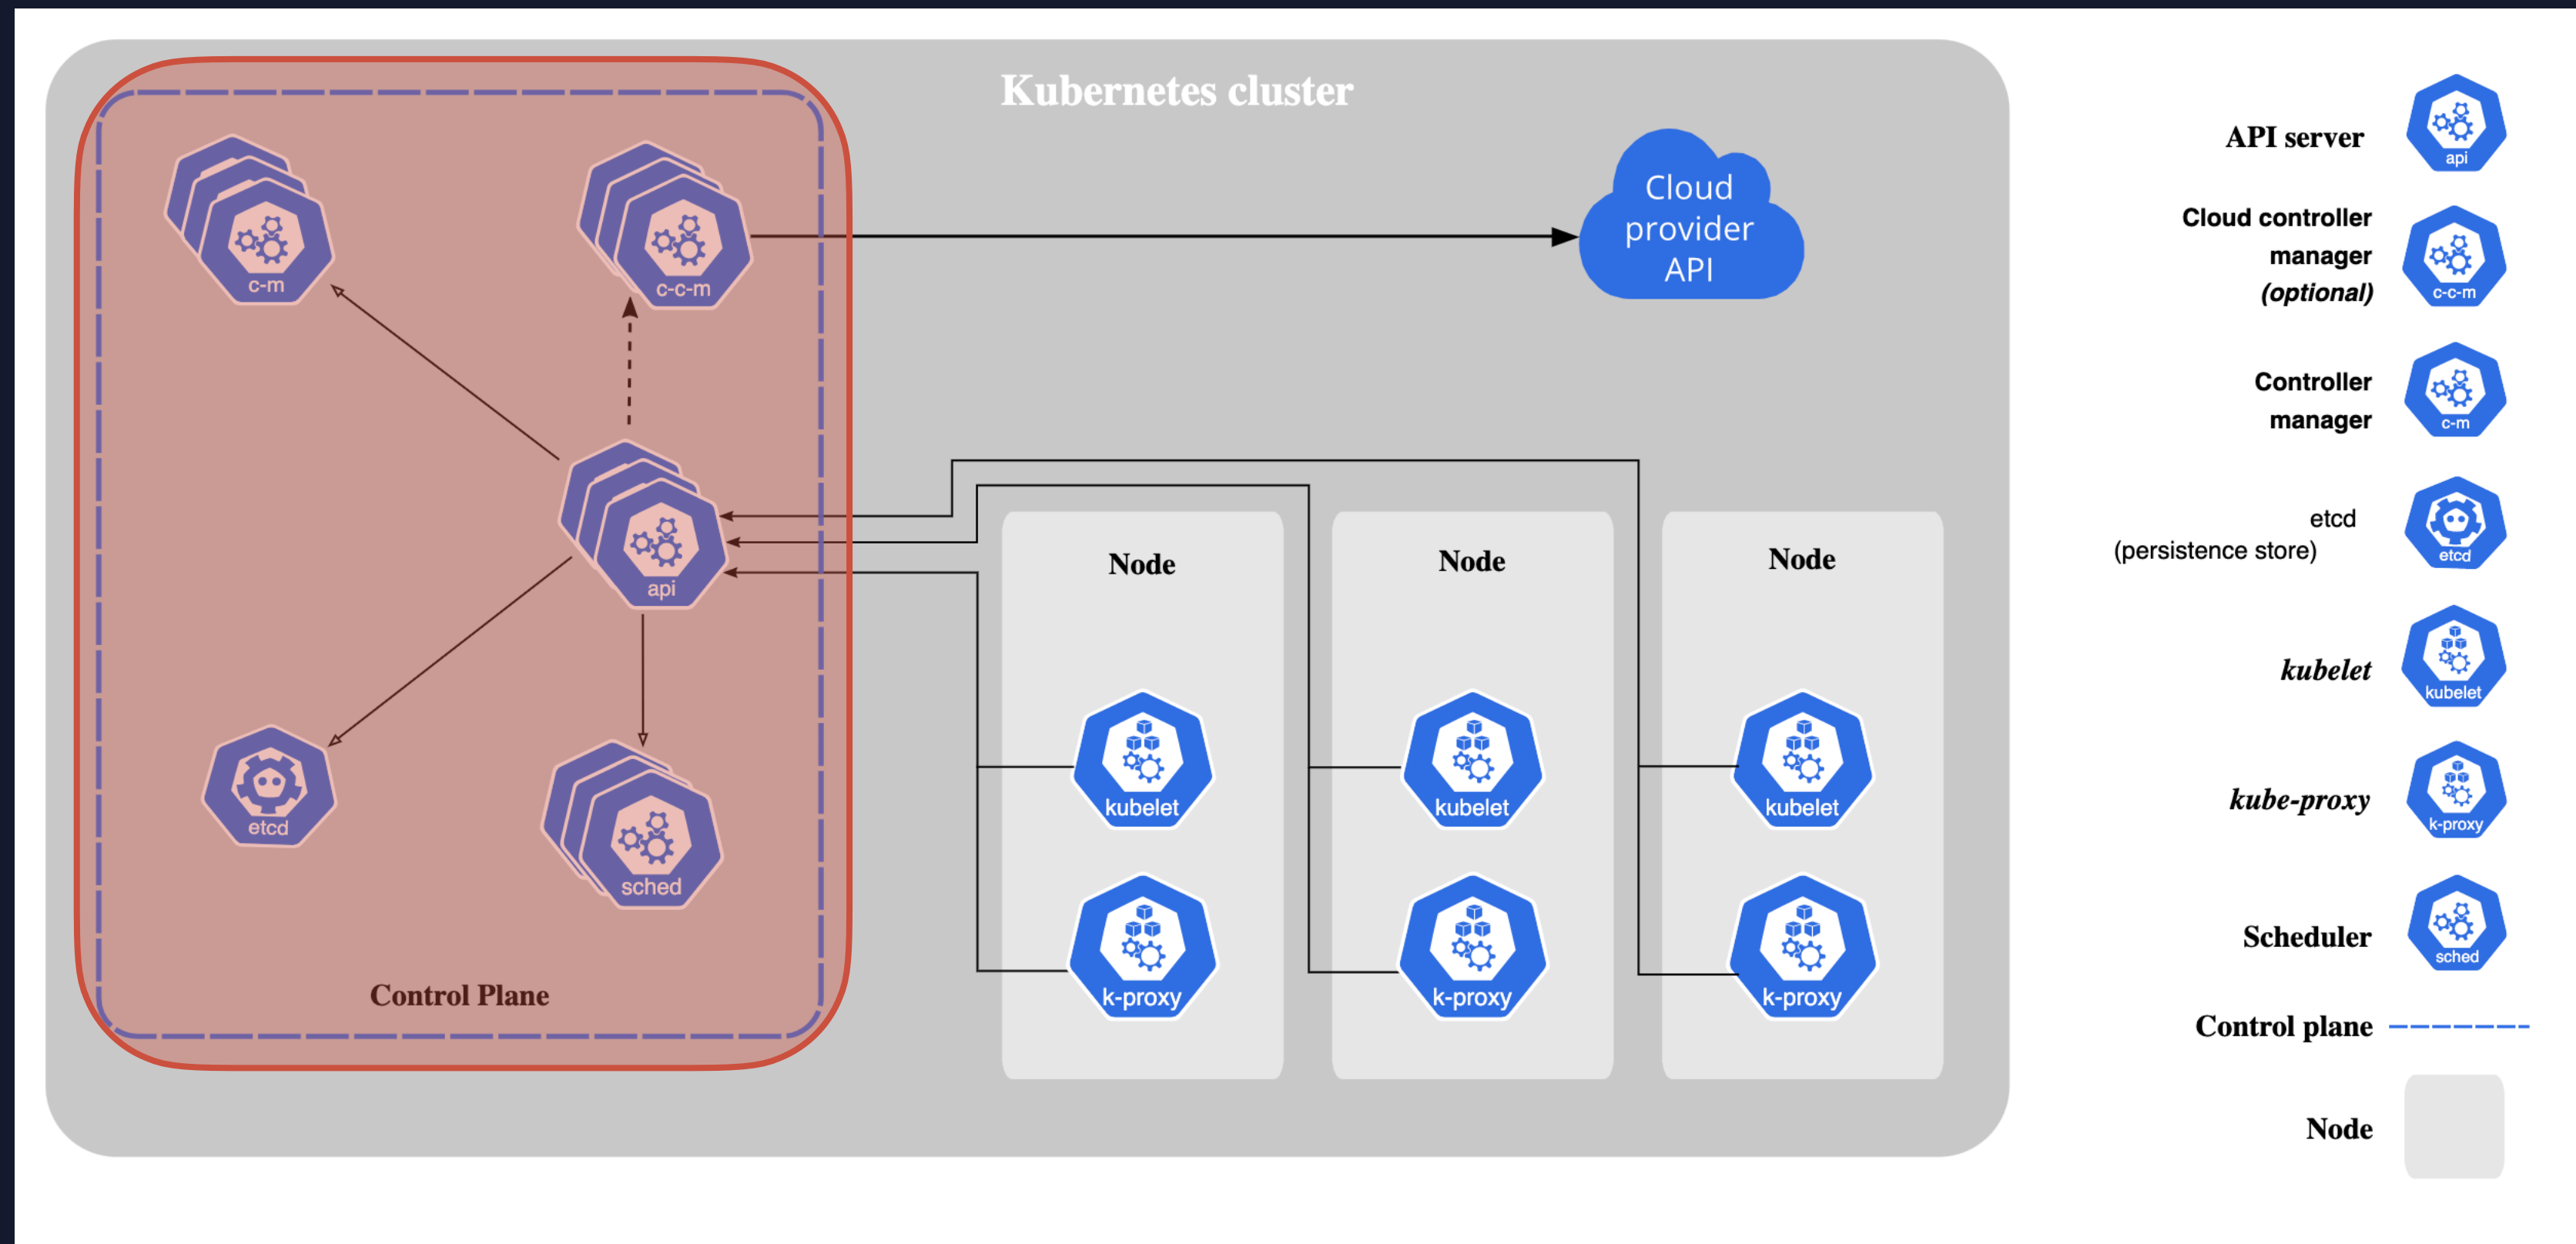


# KUBERNETES ARCHITECTURE

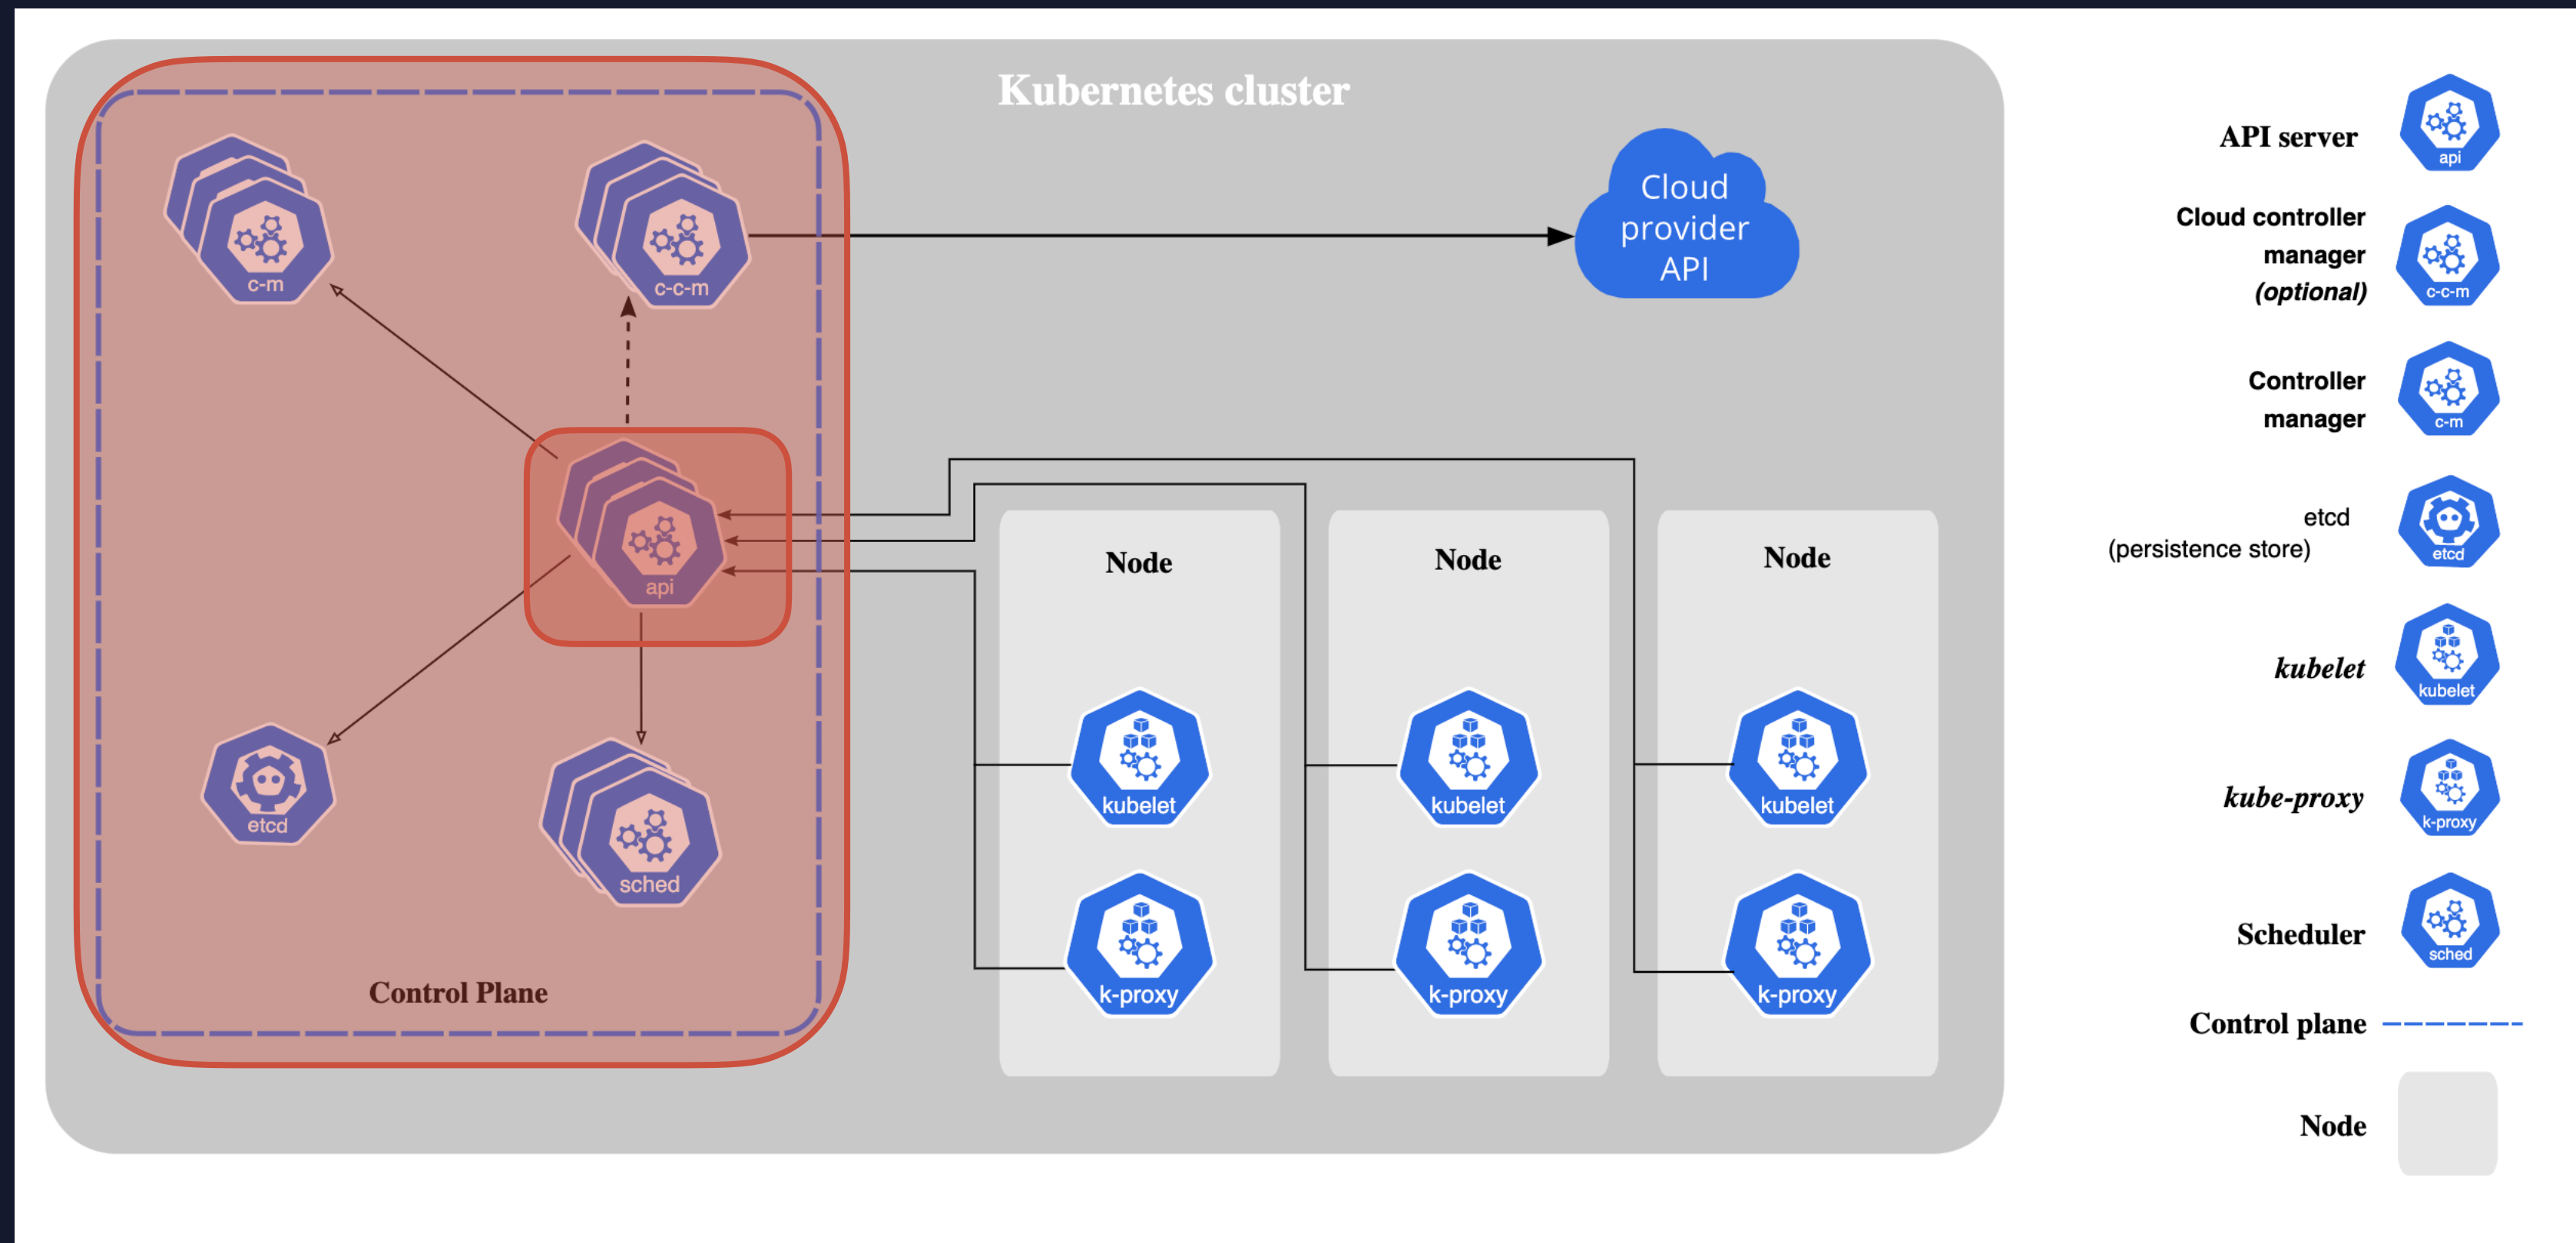


<https://kubernetes.io/docs/concepts/overview/components/>

# CONTROL PLANE



# API SERVER



# API SERVER

kube-apiserver

The main entry point for cluster administration, the frontend for the Kubernetes control plane.



# API SERVER

kube-apiserver

The main entry point for cluster administration, the frontend for the Kubernetes control plane.

REST API



# API SERVER

kube-apiserver

The main entry point for cluster administration, the frontend for the Kubernetes control plane.

REST API

Request validation



# API SERVER

kube-apiserver

The main entry point for cluster administration, the frontend for the Kubernetes control plane.

REST API

Request validation

Scaled horizontally





# API SERVER

kube-apiserver


The main entry point for cluster administration, the frontend for the Kubernetes control plane.

REST API

Request validation

Scaled horizontally

Cluster endpoint with load balancer





# API SERVER

kube-apiserver

The main entry point for cluster administration, the frontend for the Kubernetes control plane.

REST API

Request validation

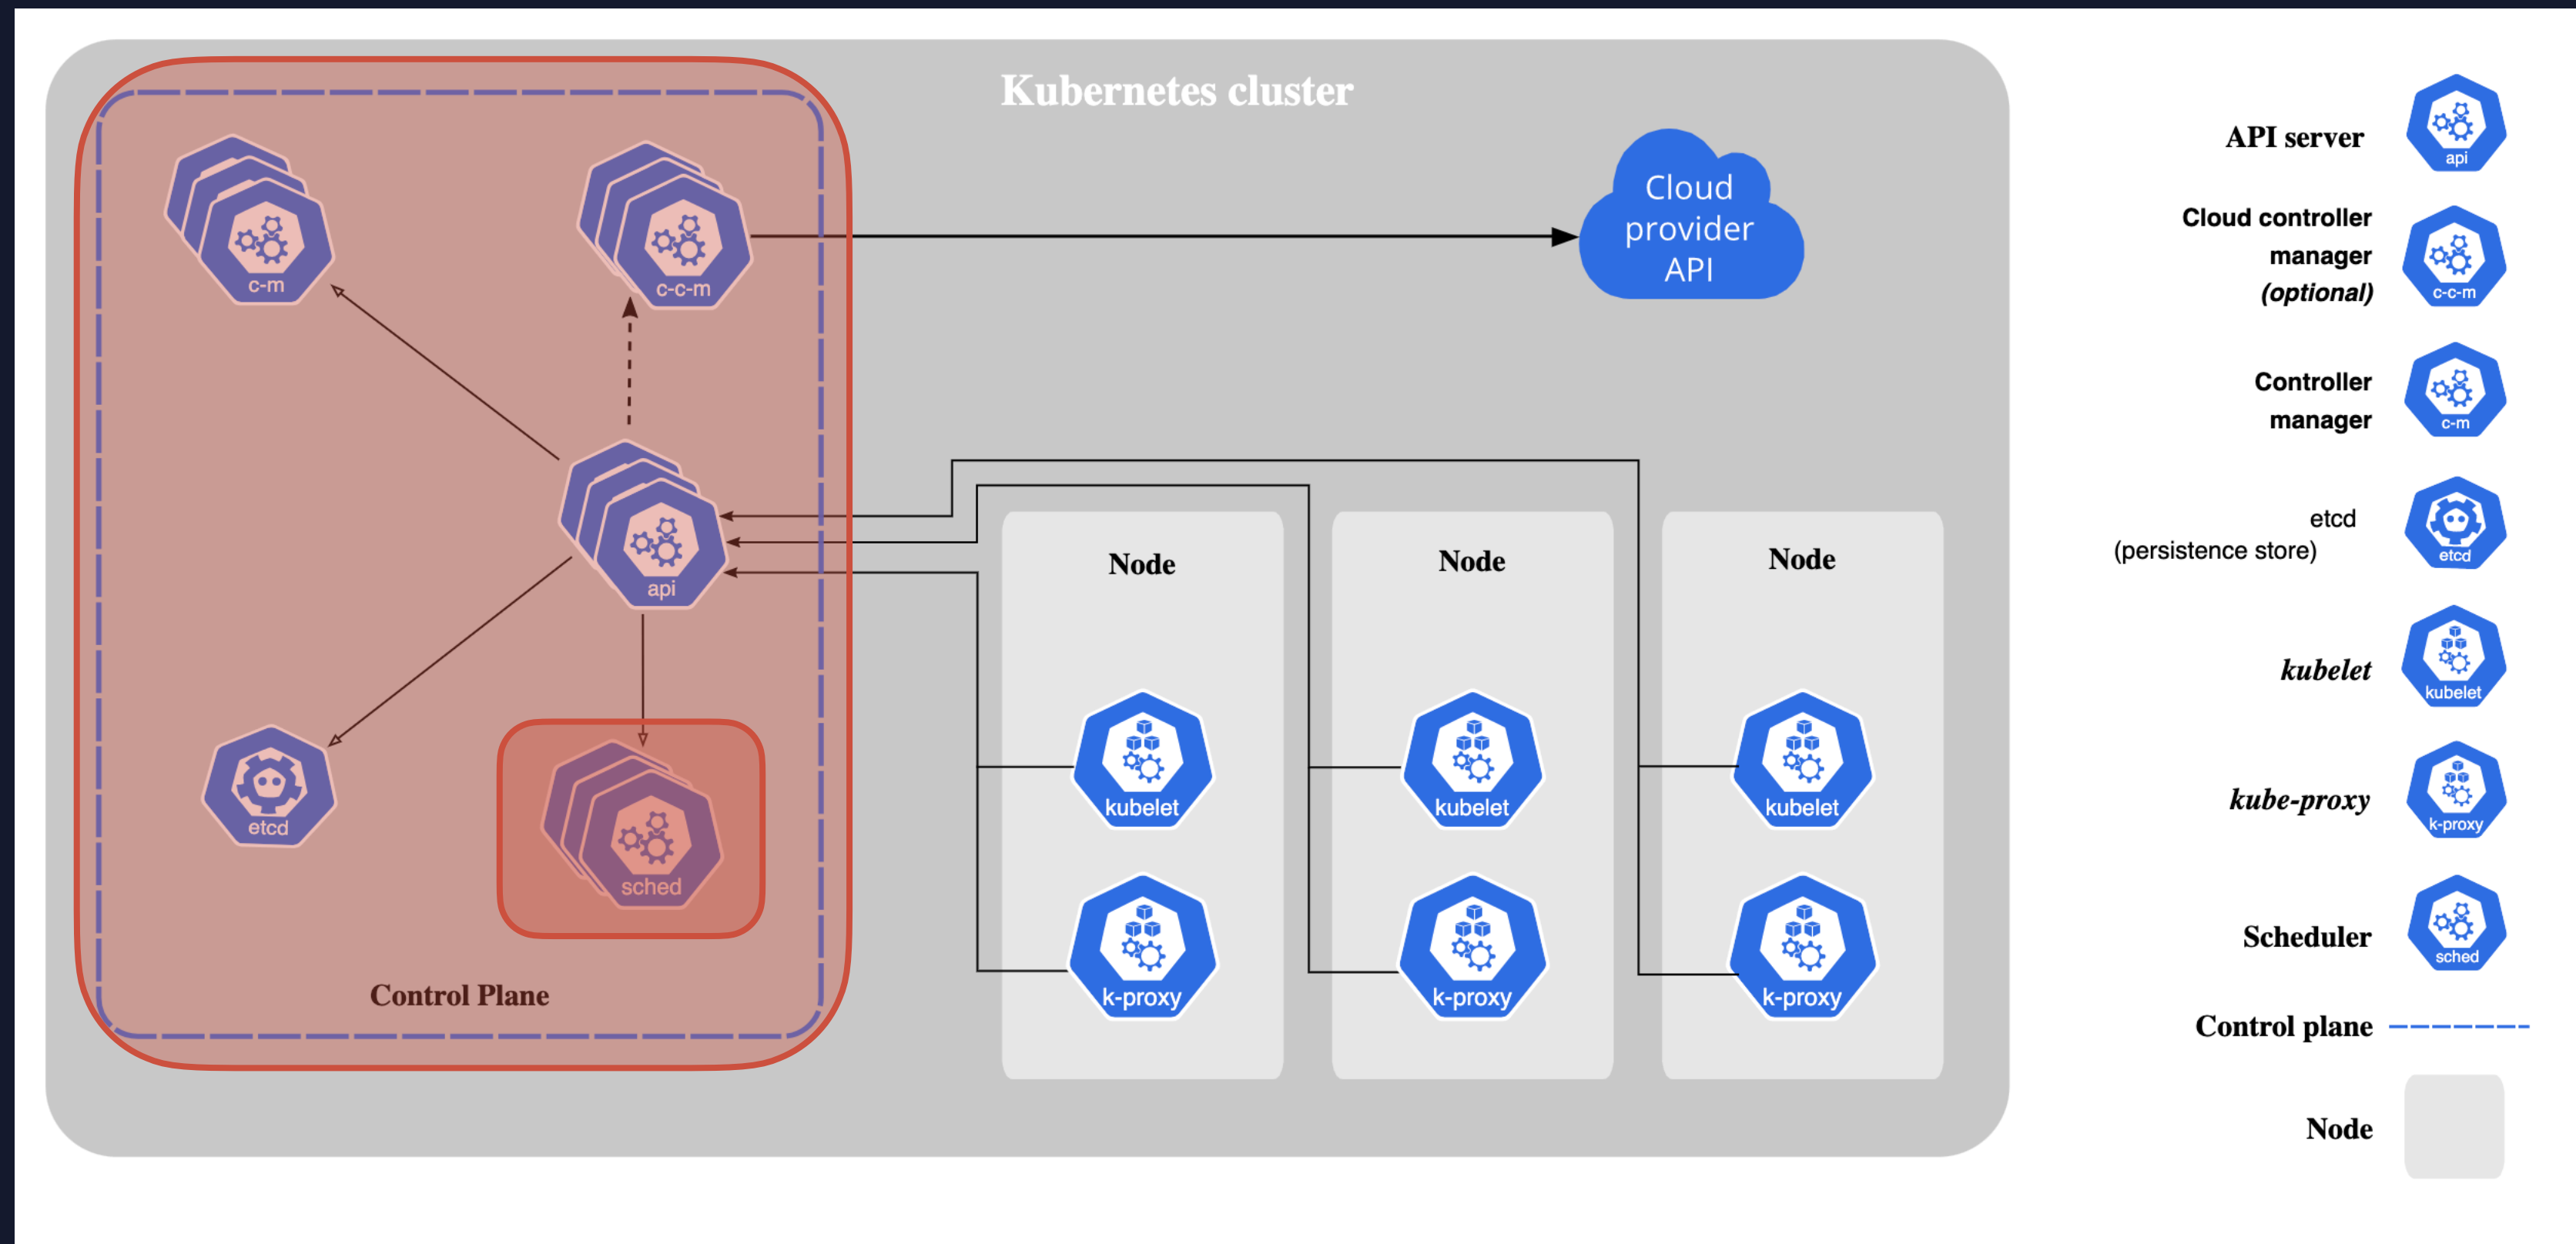
Scaled horizontally

Cluster endpoint with load balancer

kubectl



# SCHEDULER



# SCHEDULER

kube-scheduler



# SCHEDULER

kube-scheduler

Schedules Pods onto *appropriate* nodes of the cluster.



# SCHEDULER

kube-scheduler

Schedules Pods onto *appropriate* nodes of the cluster.

Matches supply and demand:



# SCHEDULER

kube-scheduler

Schedules Pods onto *appropriate* nodes of the cluster.

Matches supply and demand:

Resource requirements / limits (CPU, memory)



# SCHEDULER

kube-scheduler

Schedules Pods onto *appropriate* nodes of the cluster.

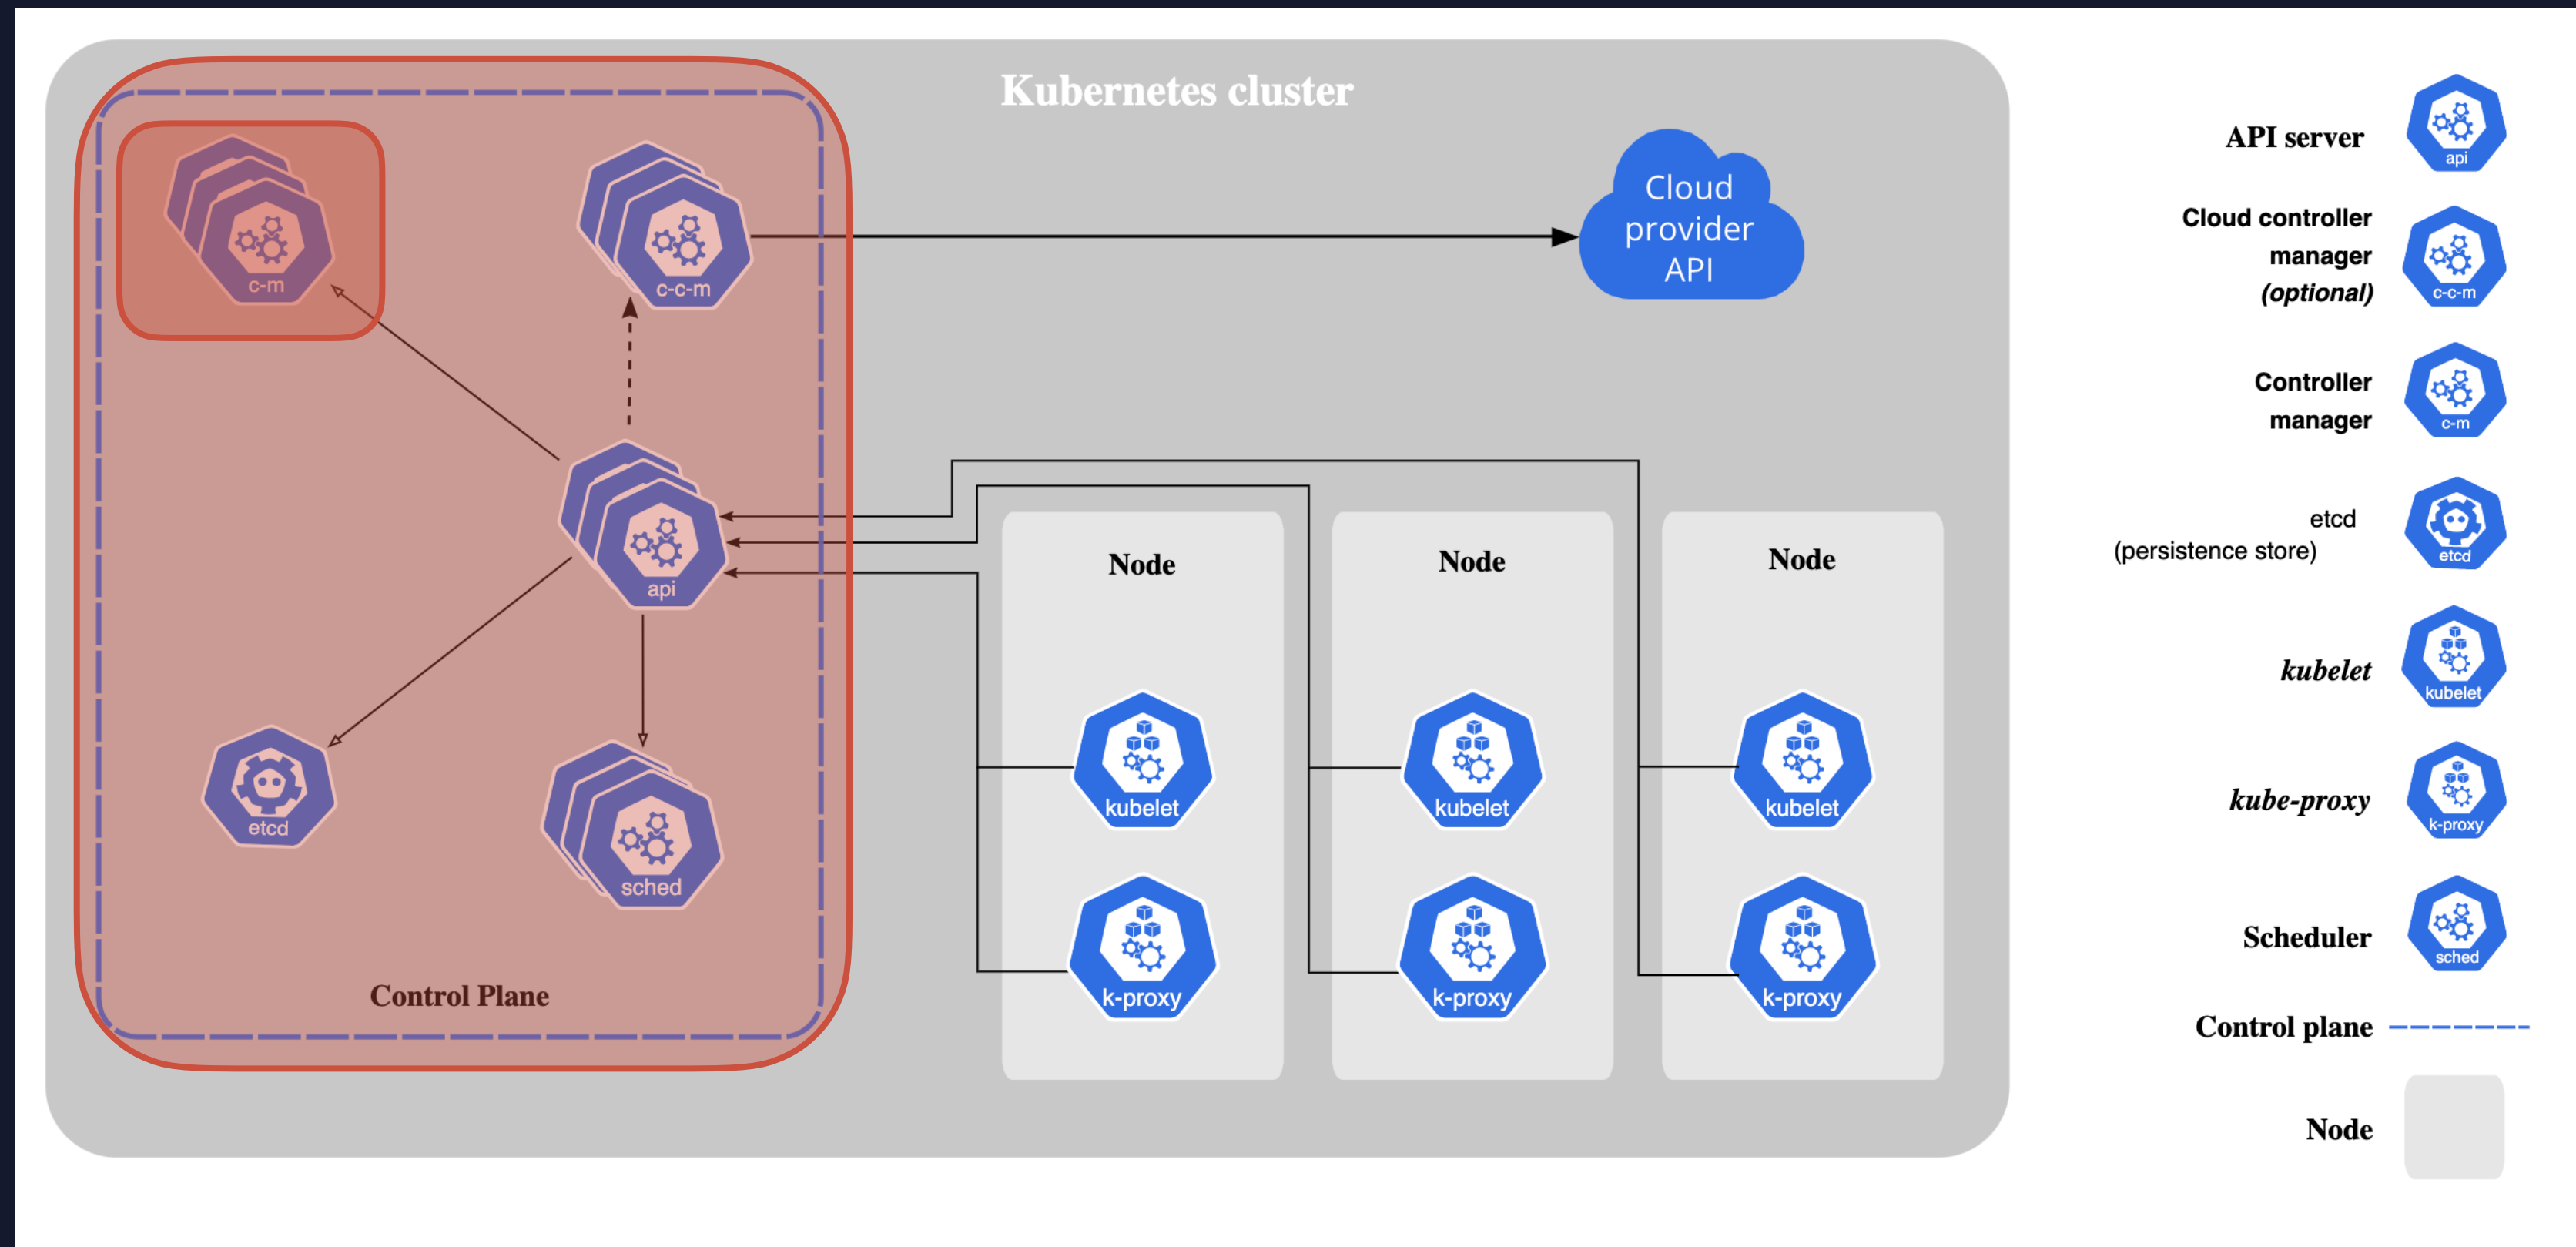
Matches supply and demand:

Resource requirements / limits (CPU, memory)

nodeSelector / affinity / taints



# CONTROLLER MANAGER





# CONTROLLER MANAGER

kube-controller-manager

Manages various controllers.

# WHAT IS A CONTROLLER?



# WHAT IS A CONTROLLER?

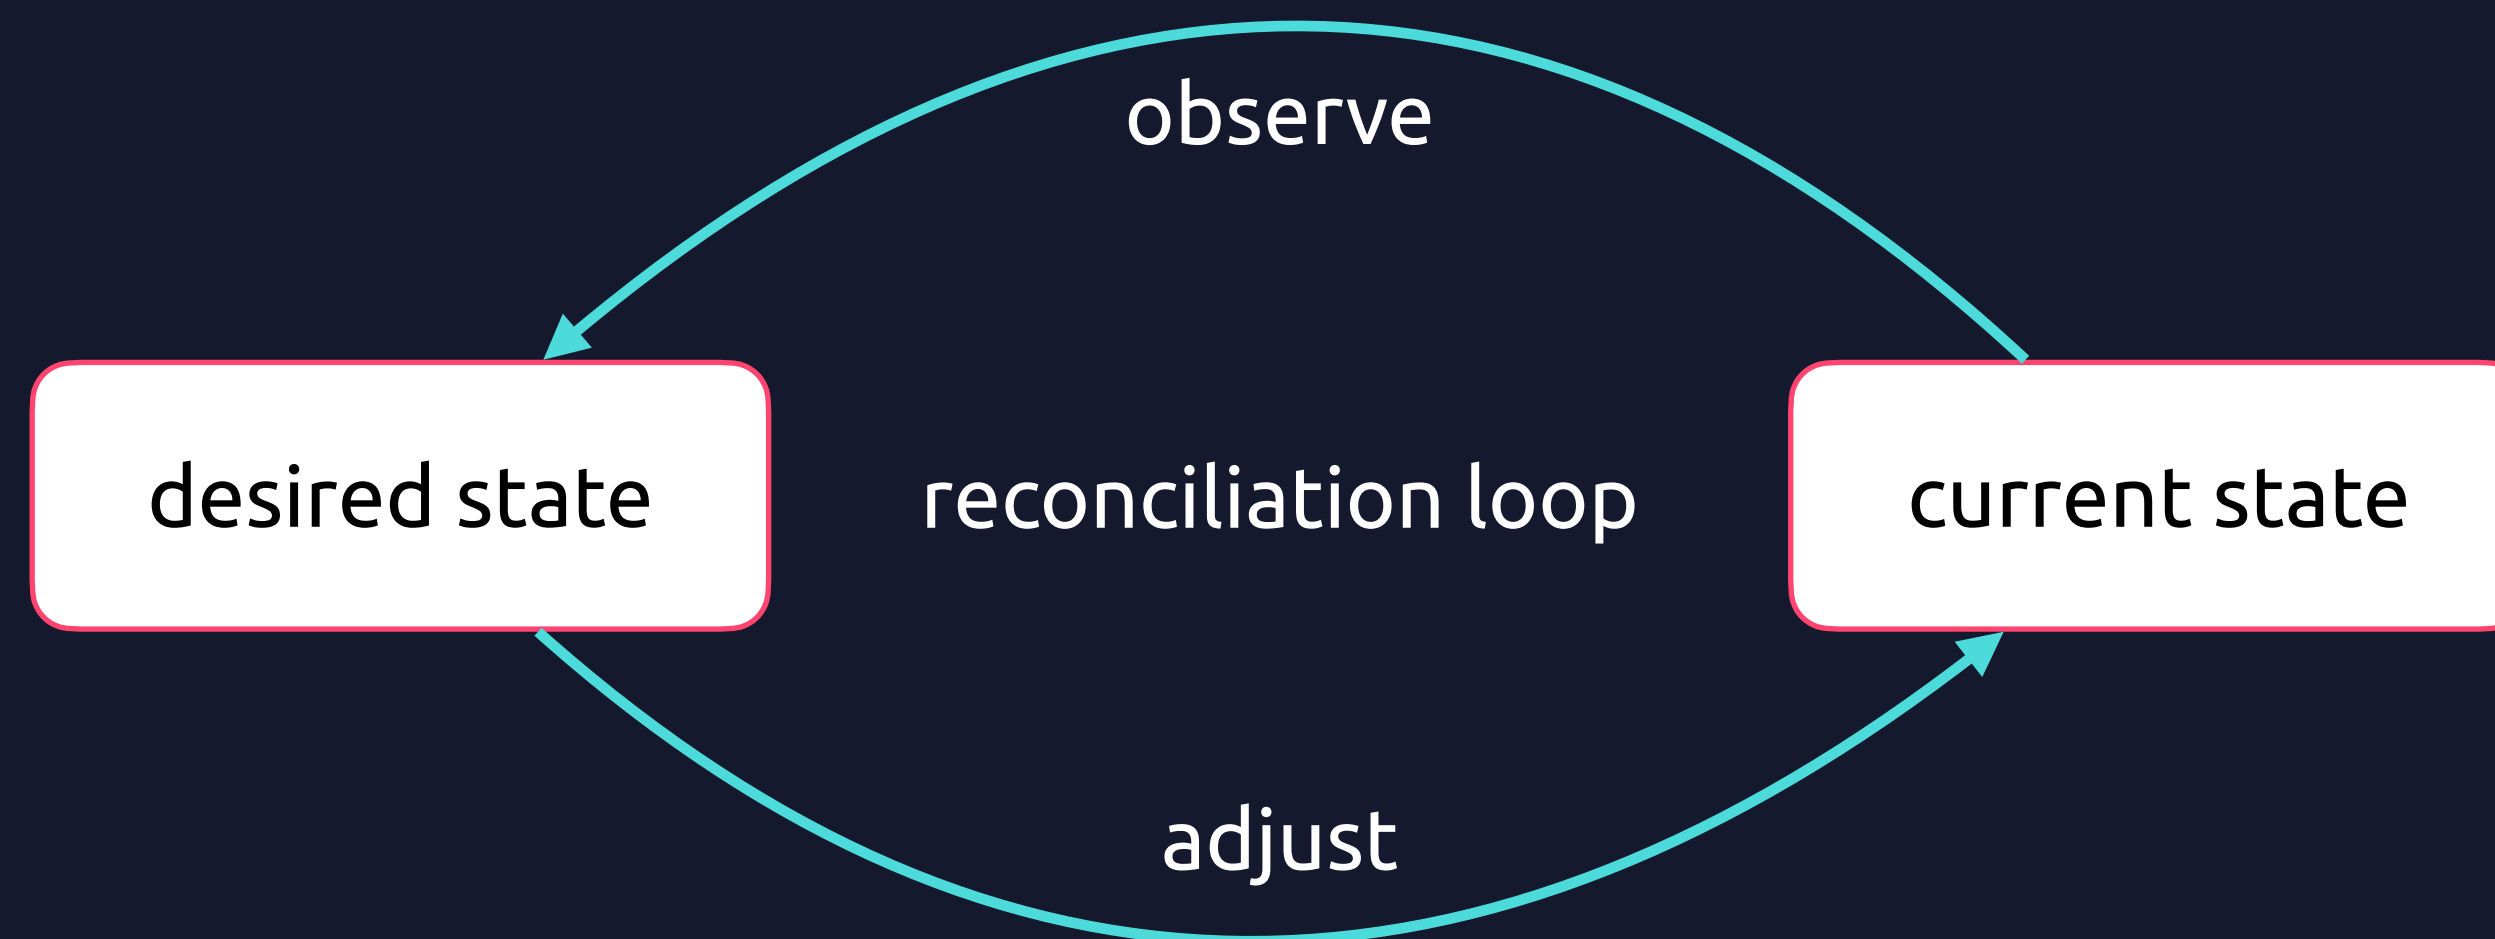
Controllers are responsible for ensuring the components of the cluster are in the *desired* state.



# WHAT IS A CONTROLLER?

Controllers are responsible for ensuring the components of the cluster are in the *desired* state.

Reconciliation / control loop:



# CONTROLLER MANAGER

Manages various controllers.



# CONTROLLER MANAGER

Manages various controllers.

ReplicaSet / Deployment controller

# CONTROLLER MANAGER

Manages various controllers.

ReplicaSet / Deployment controller

Node controller

# CONTROLLER MANAGER

Manages various controllers.

ReplicaSet / Deployment controller

Node controller

Endpoints controller



# CONTROLLER MANAGER

Manages various controllers.

ReplicaSet / Deployment controller

Node controller

Endpoints controller

CRD controller

# + CLOUD CONTROLLER MANAGER

Cloud-specific control logic:



# + CLOUD CONTROLLER MANAGER

Cloud-specific control logic:

Node-related information



# + CLOUD CONTROLLER MANAGER

Cloud-specific control logic:

Node-related information

Storage



# + CLOUD CONTROLLER MANAGER

Cloud-specific control logic:

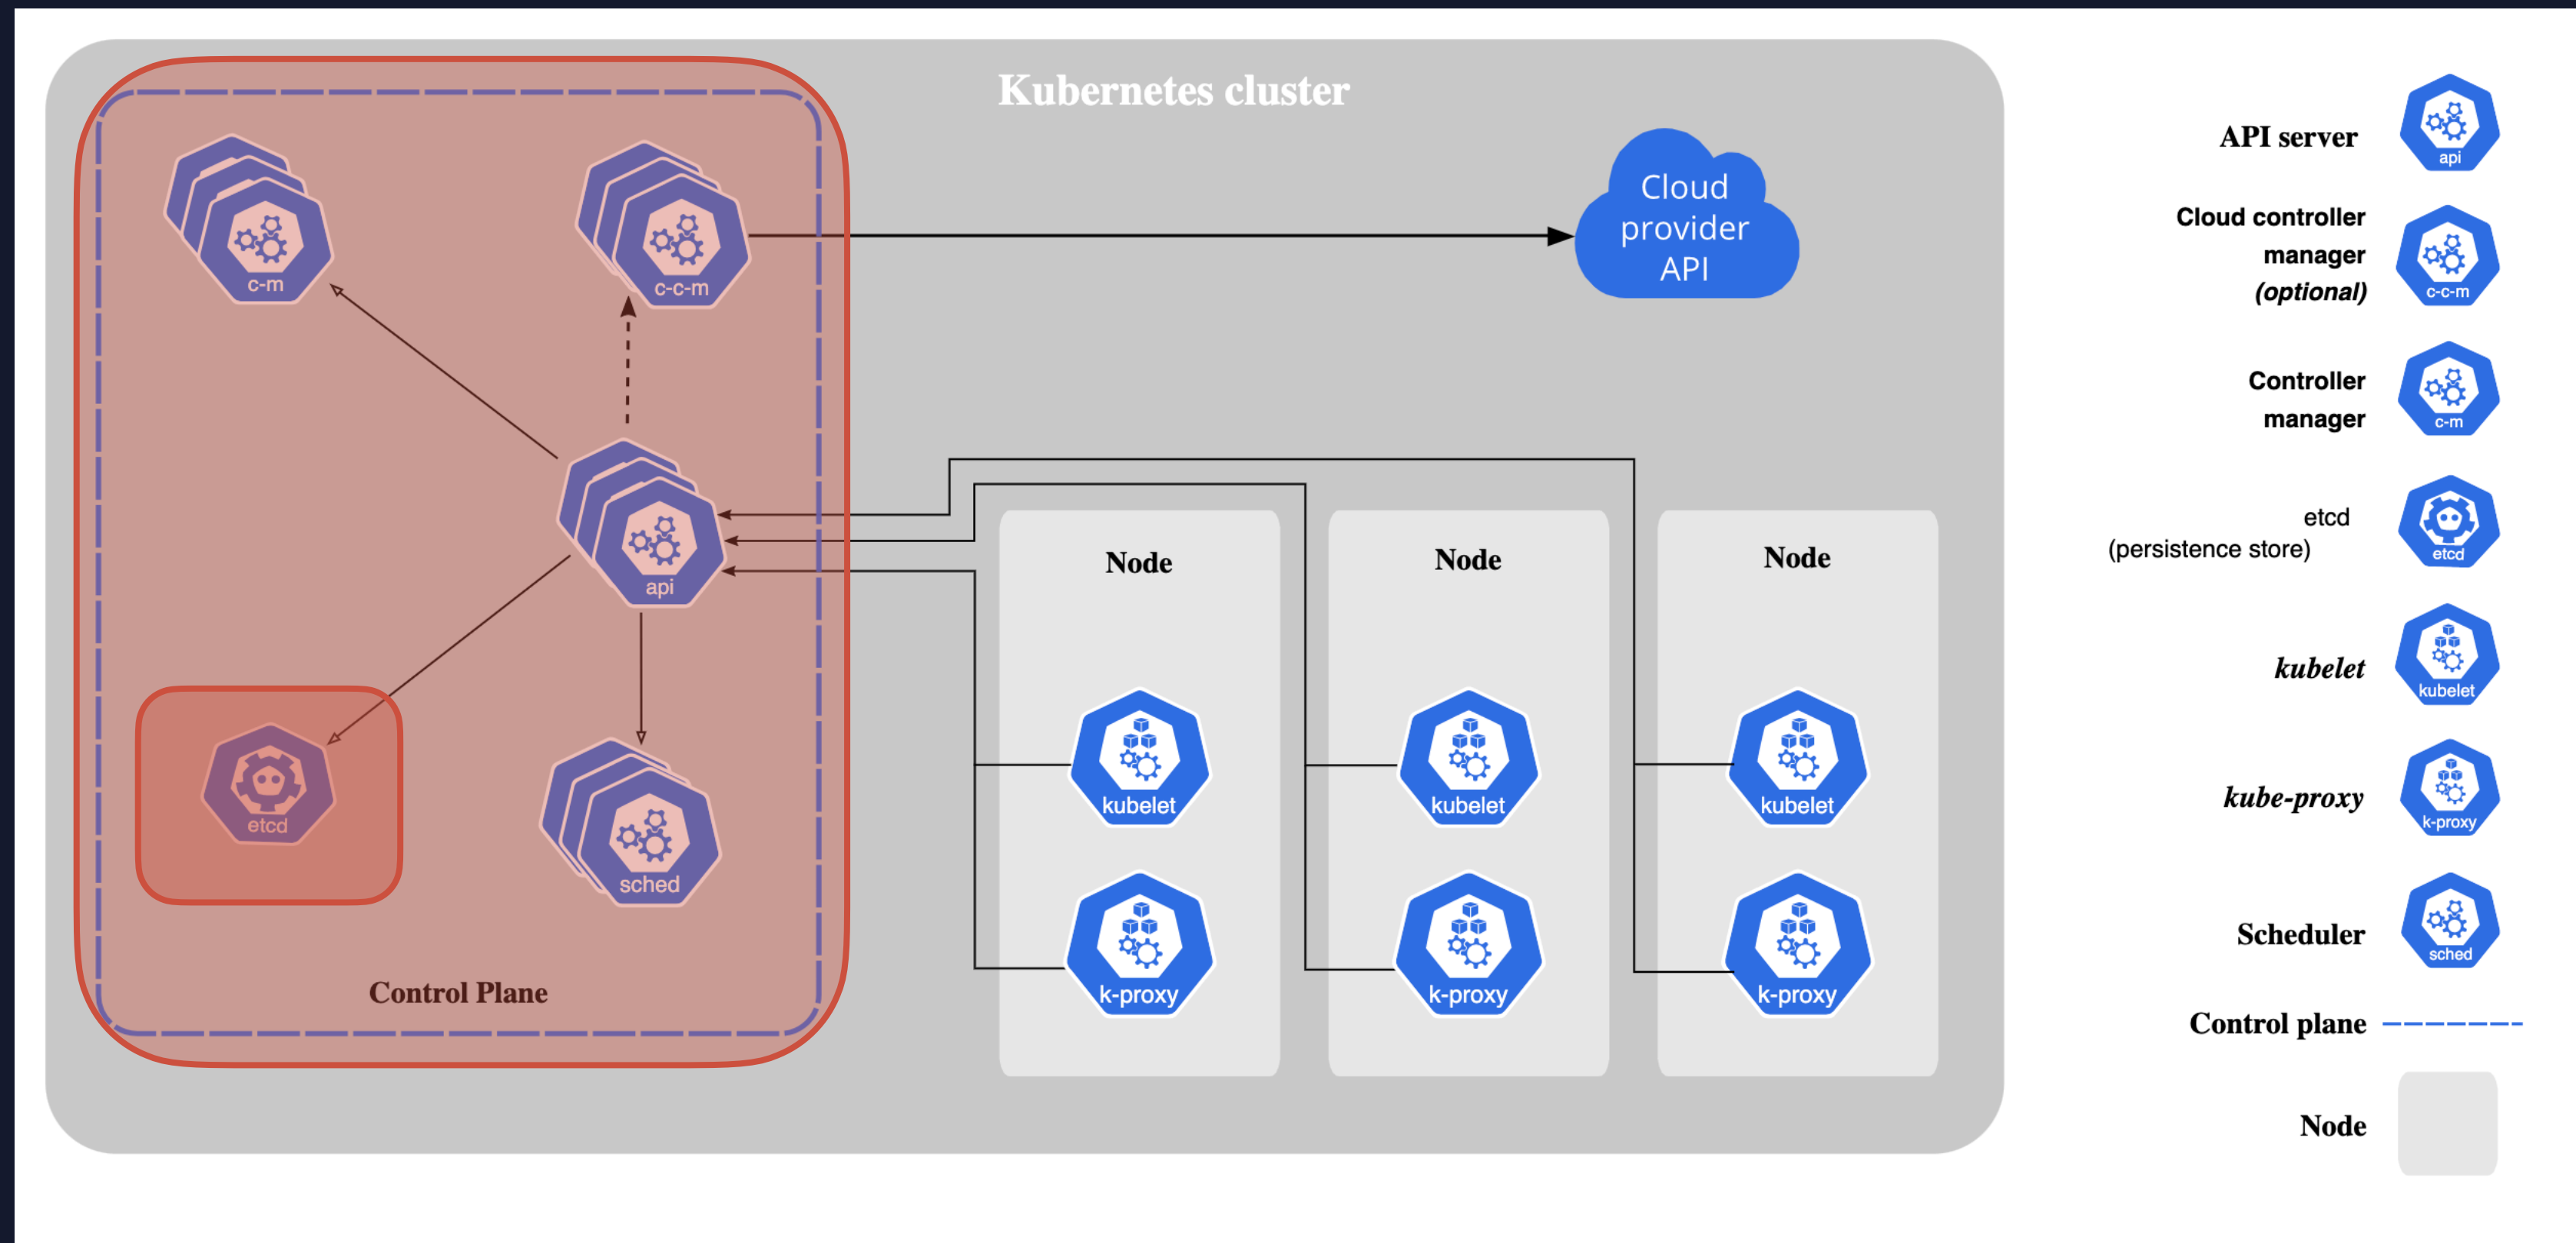
Node-related information

Storage

Services, Networking



# ETCD



# ETCD

○ The single source of truth – the database for storing all cluster data.

# ETCD

○ The single source of truth – the database for storing all cluster data.

HA, distributed, key-value storage



# ETCD

○ The single source of truth – the database for storing all cluster data.

HA, distributed, key-value storage

Kubernetes API objects: Pods, Services, Nodes, etc.



# ETCD

The single source of truth – the database for storing all cluster data.

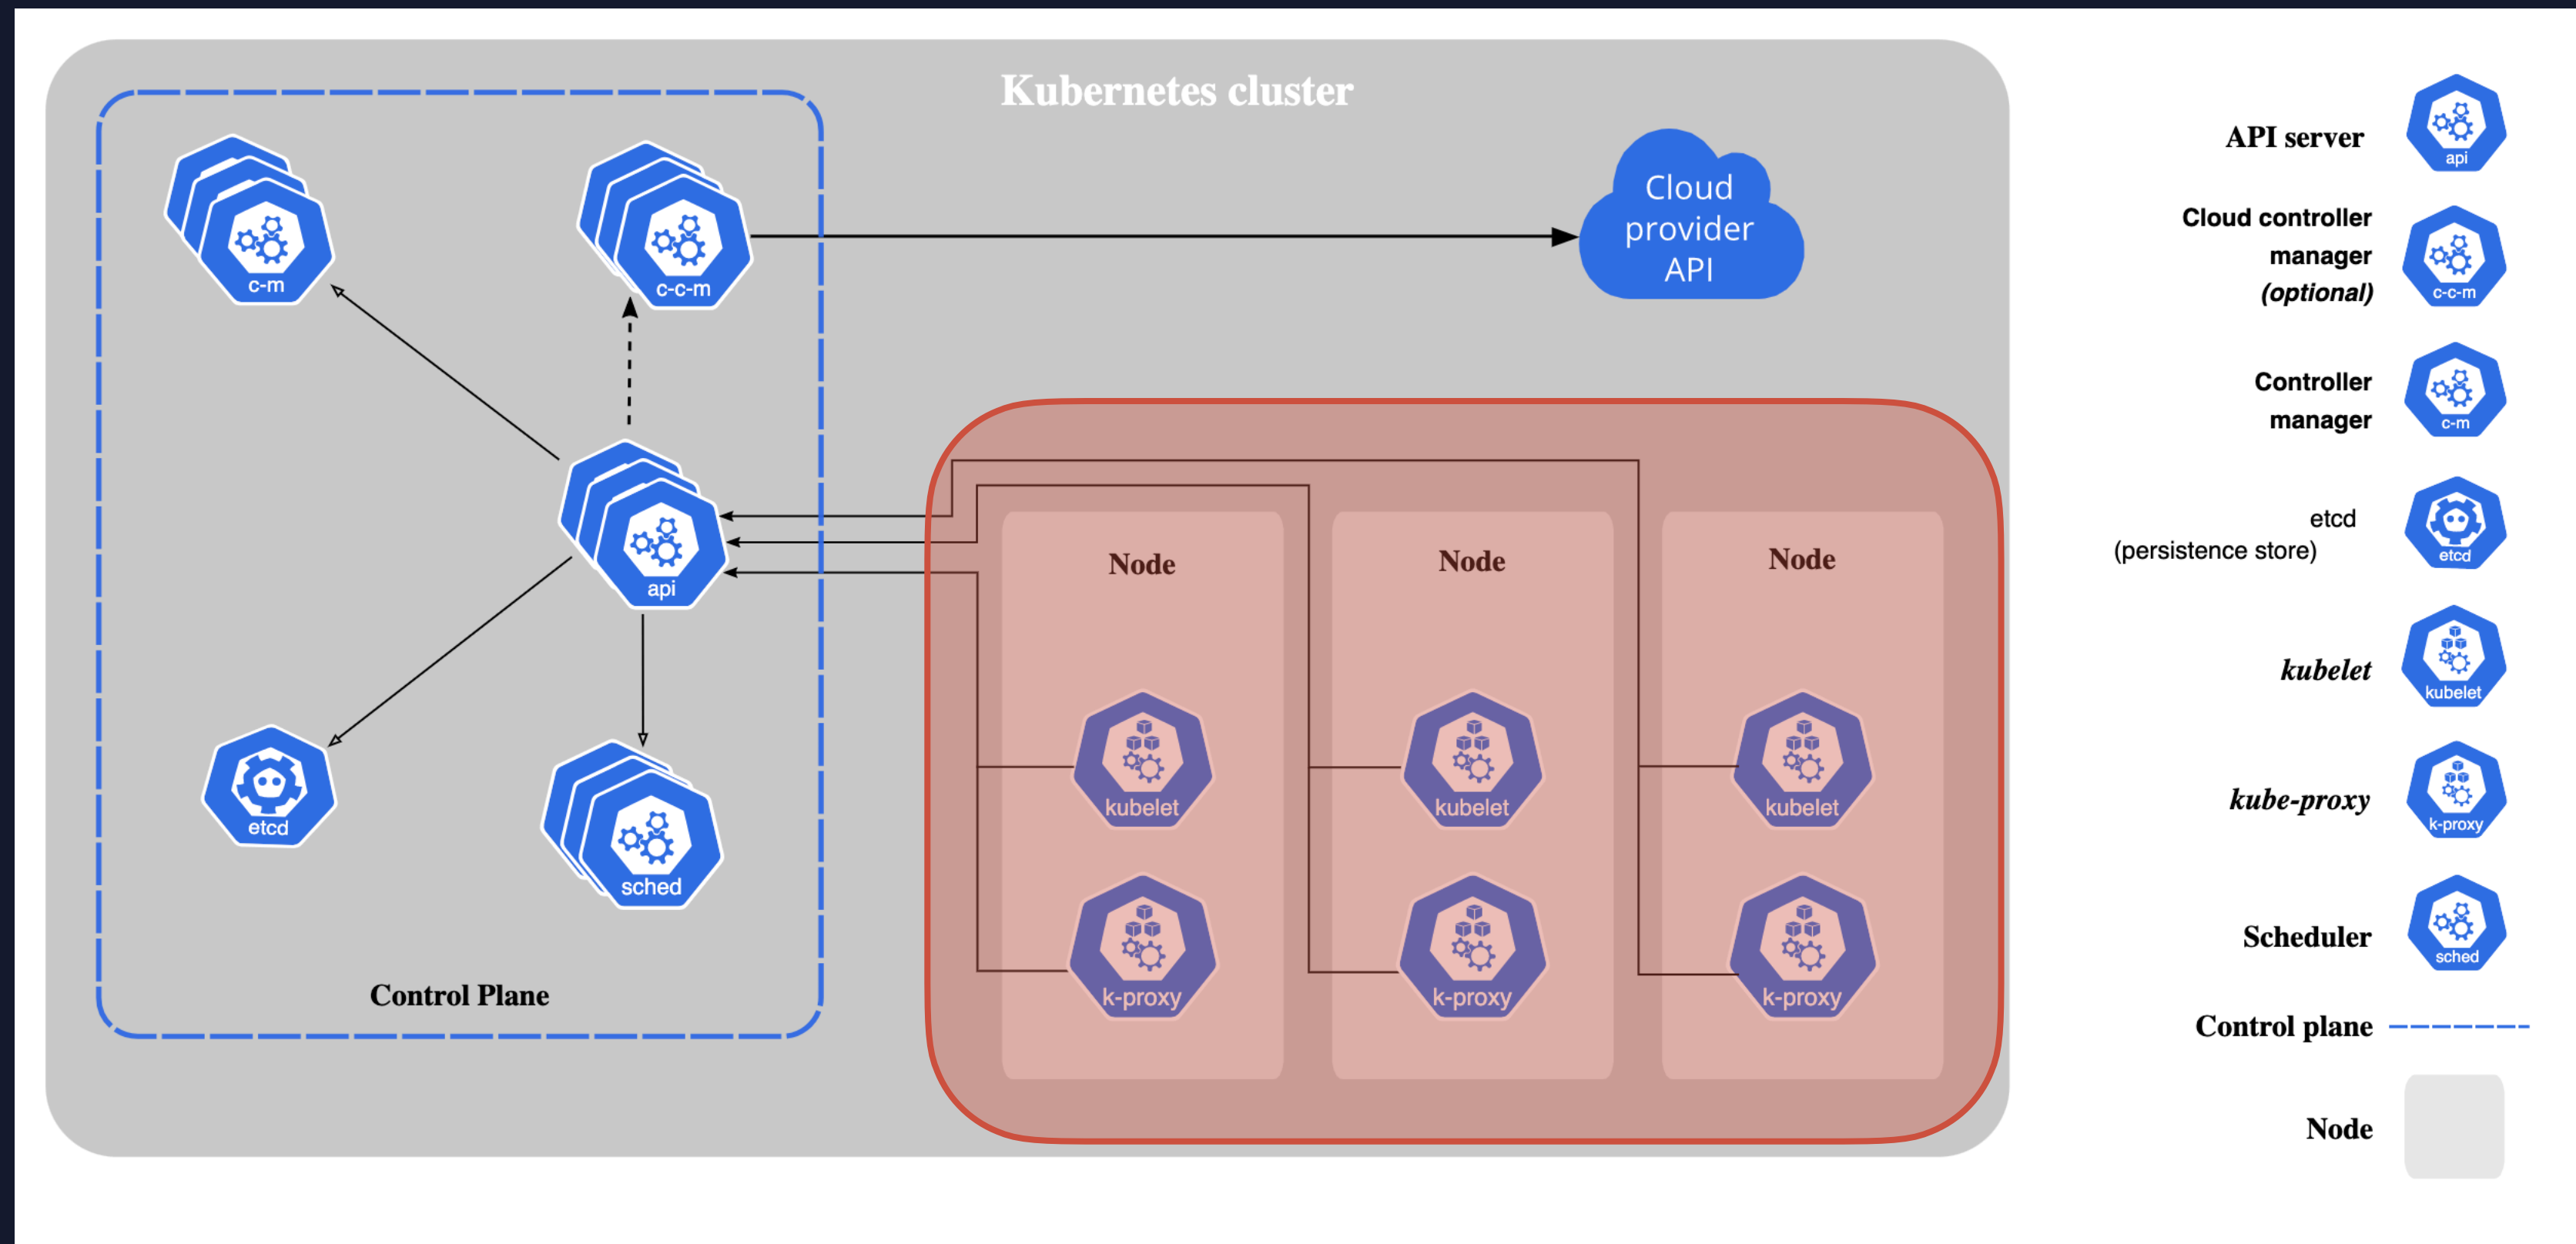
HA, distributed, key-value storage

Kubernetes API objects: Pods, Services, Nodes, etc.

(Might be replaced with other storage engine: e.g. SQLite)



# WORKER NODES



# KUBELET

The Kubernetes agent that runs on each node of the cluster.



# KUBELET

The Kubernetes agent that runs on each node of the cluster.

Registers node with the API server.



# KUBELET

The Kubernetes agent that runs on each node of the cluster.

Registers node with the API server.

Takes PodSpecs from the API server and ensures that containers of Pods are running and healthy.



# KUBELET

The Kubernetes agent that runs on each node of the cluster.

Registers node with the API server.

Takes PodSpecs from the API server and ensures that containers of Pods are running and healthy.

Reports the status of the node itself.



# + CONTAINER RUNTIME

Docker

containerd

CRI-O





# PROXY

kube-proxy

A network proxy running on each node.



# PROXY

kube-proxy

A network proxy running on each node.

Maintains network rules to allow network communication to Pods from inside or outside of the cluster.



# PROXY

kube-proxy

A network proxy running on each node.

Maintains network rules to allow network communication to Pods from inside or outside of the cluster.

Key component for Services.



# + OTHER

DNS

Web UI (Dashboard)

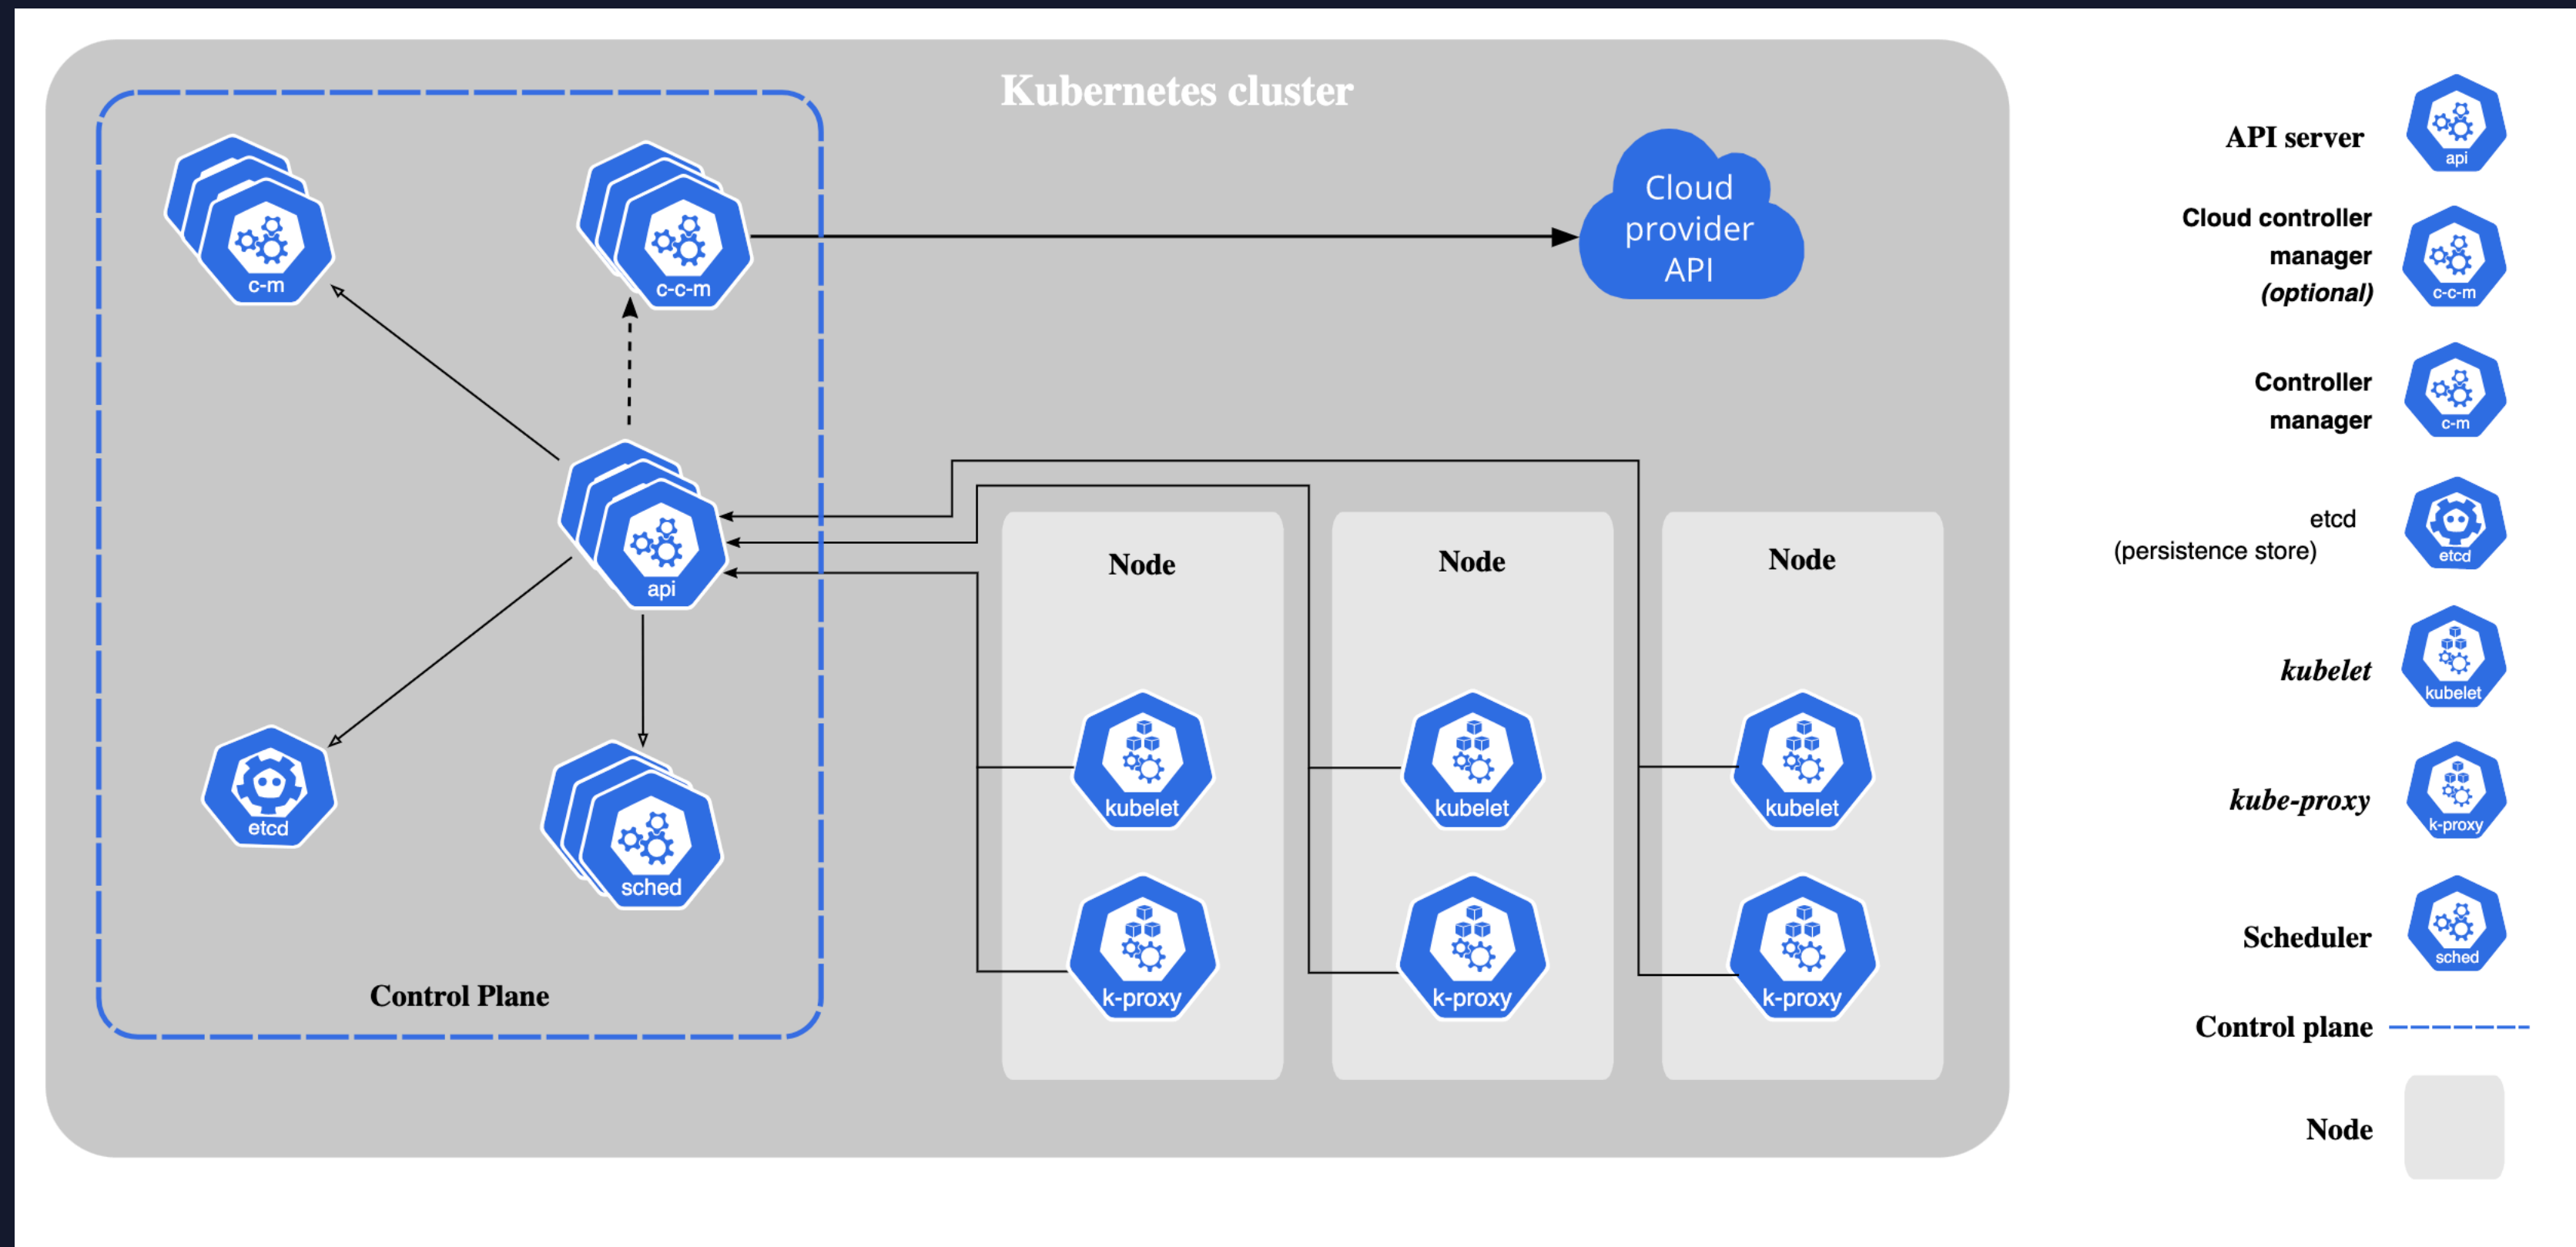
Cluster-level logging



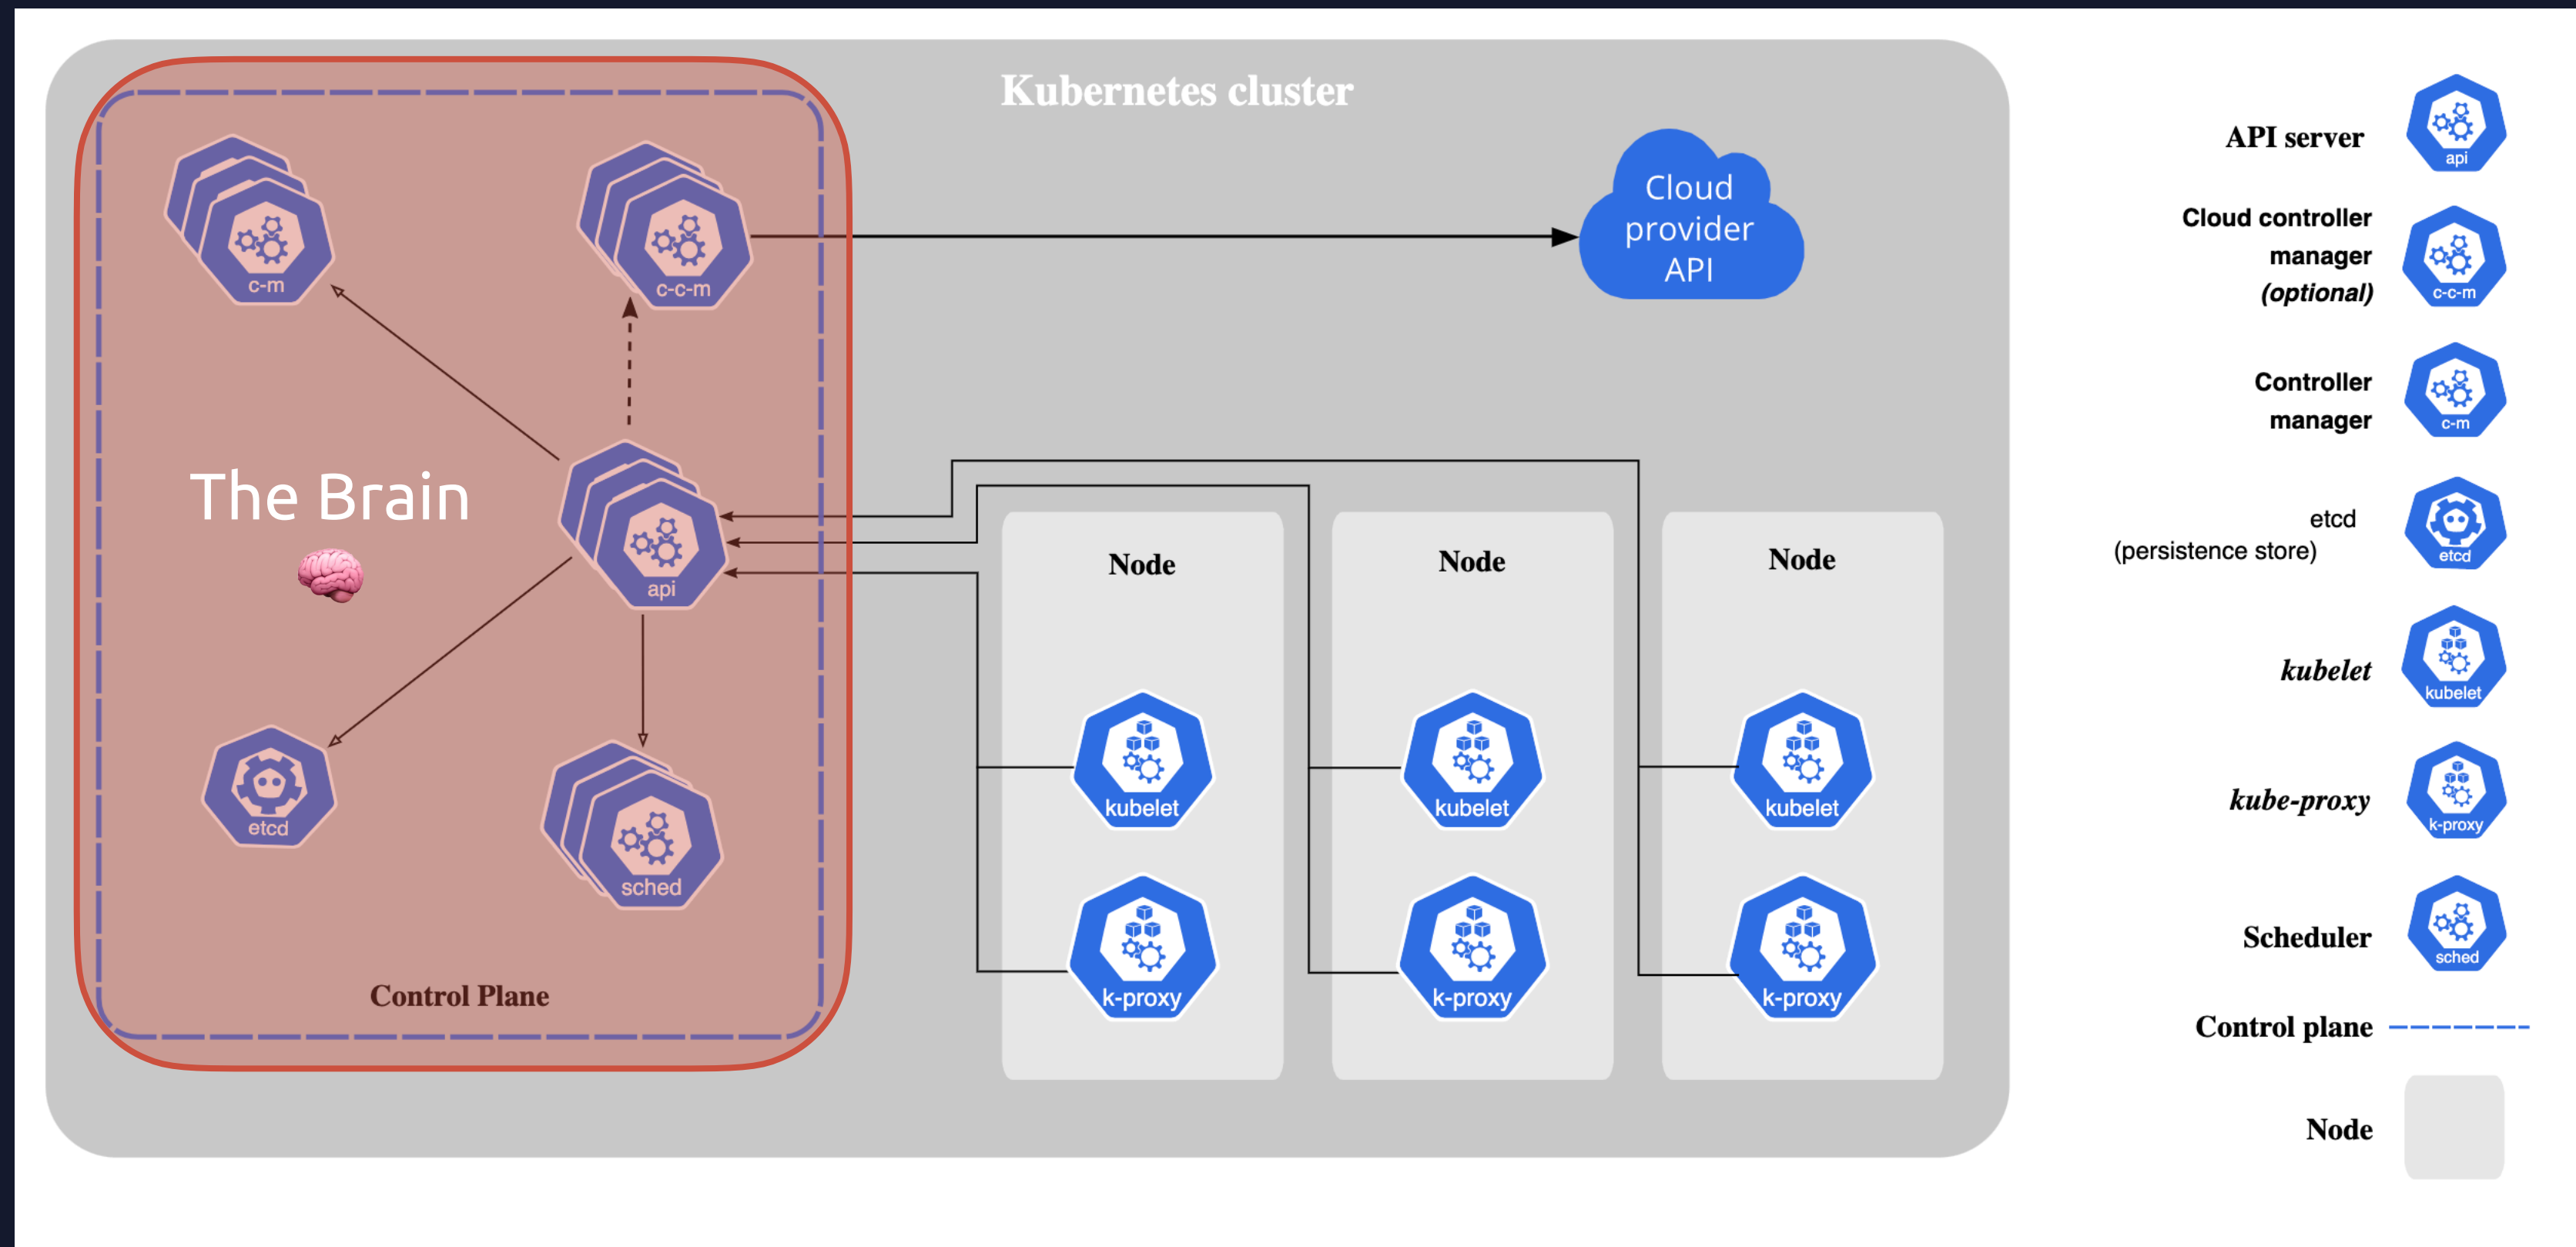
# RECAP



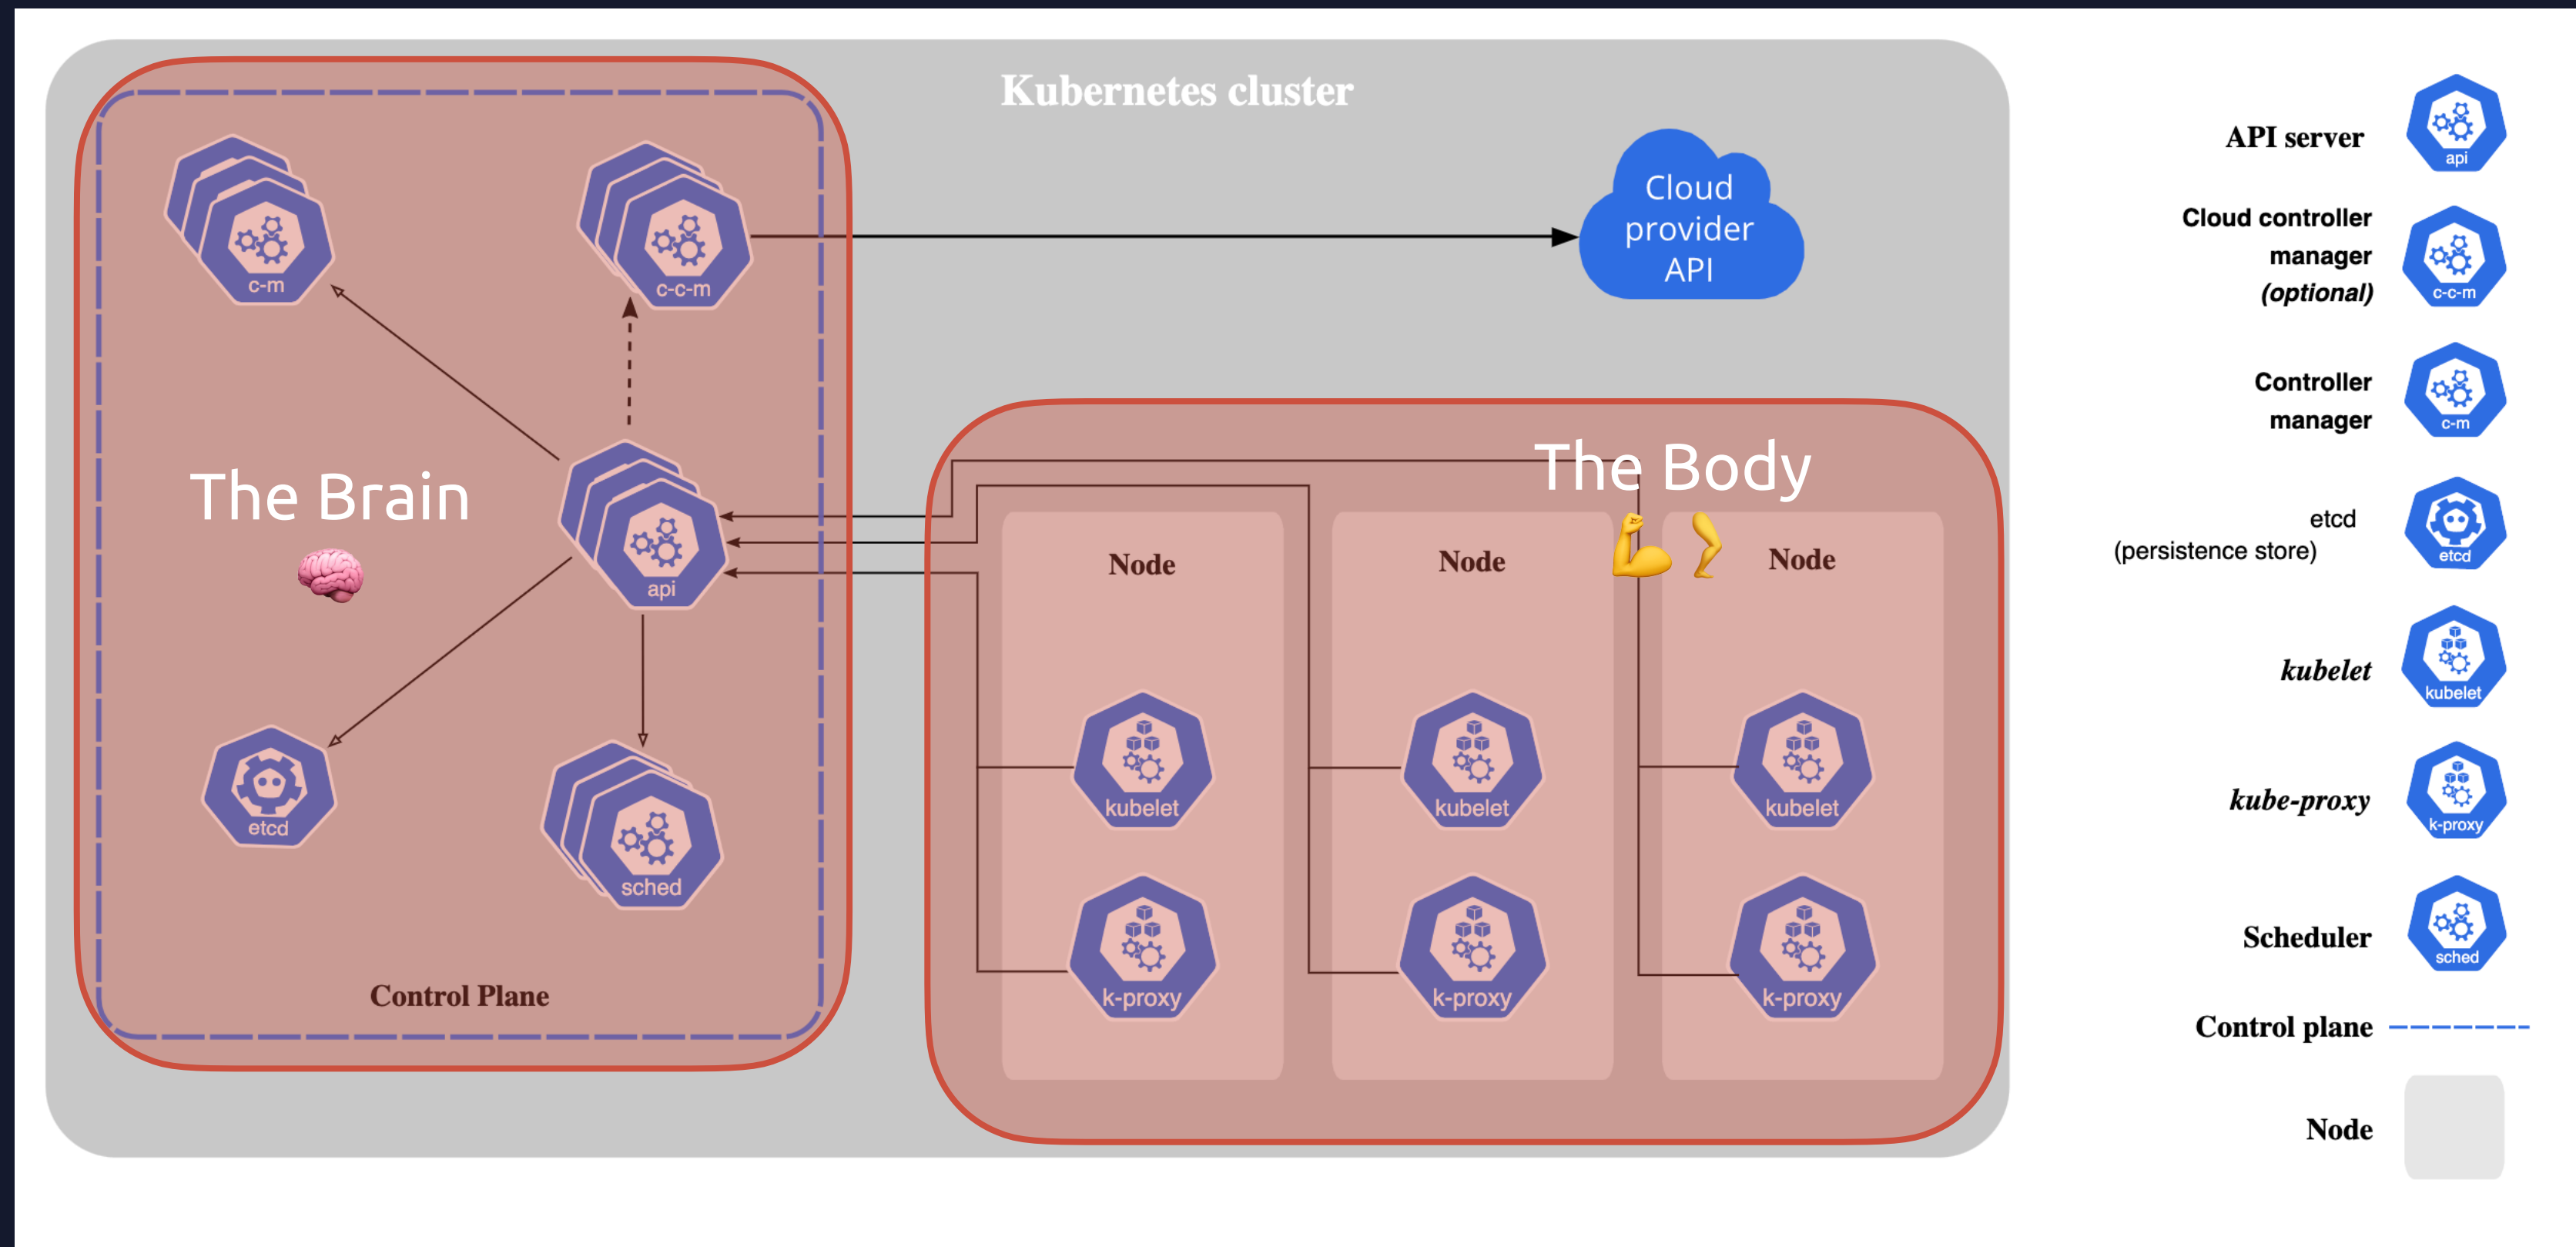
# KUBERNETES ARCHITECTURE



# KUBERNETES ARCHITECTURE

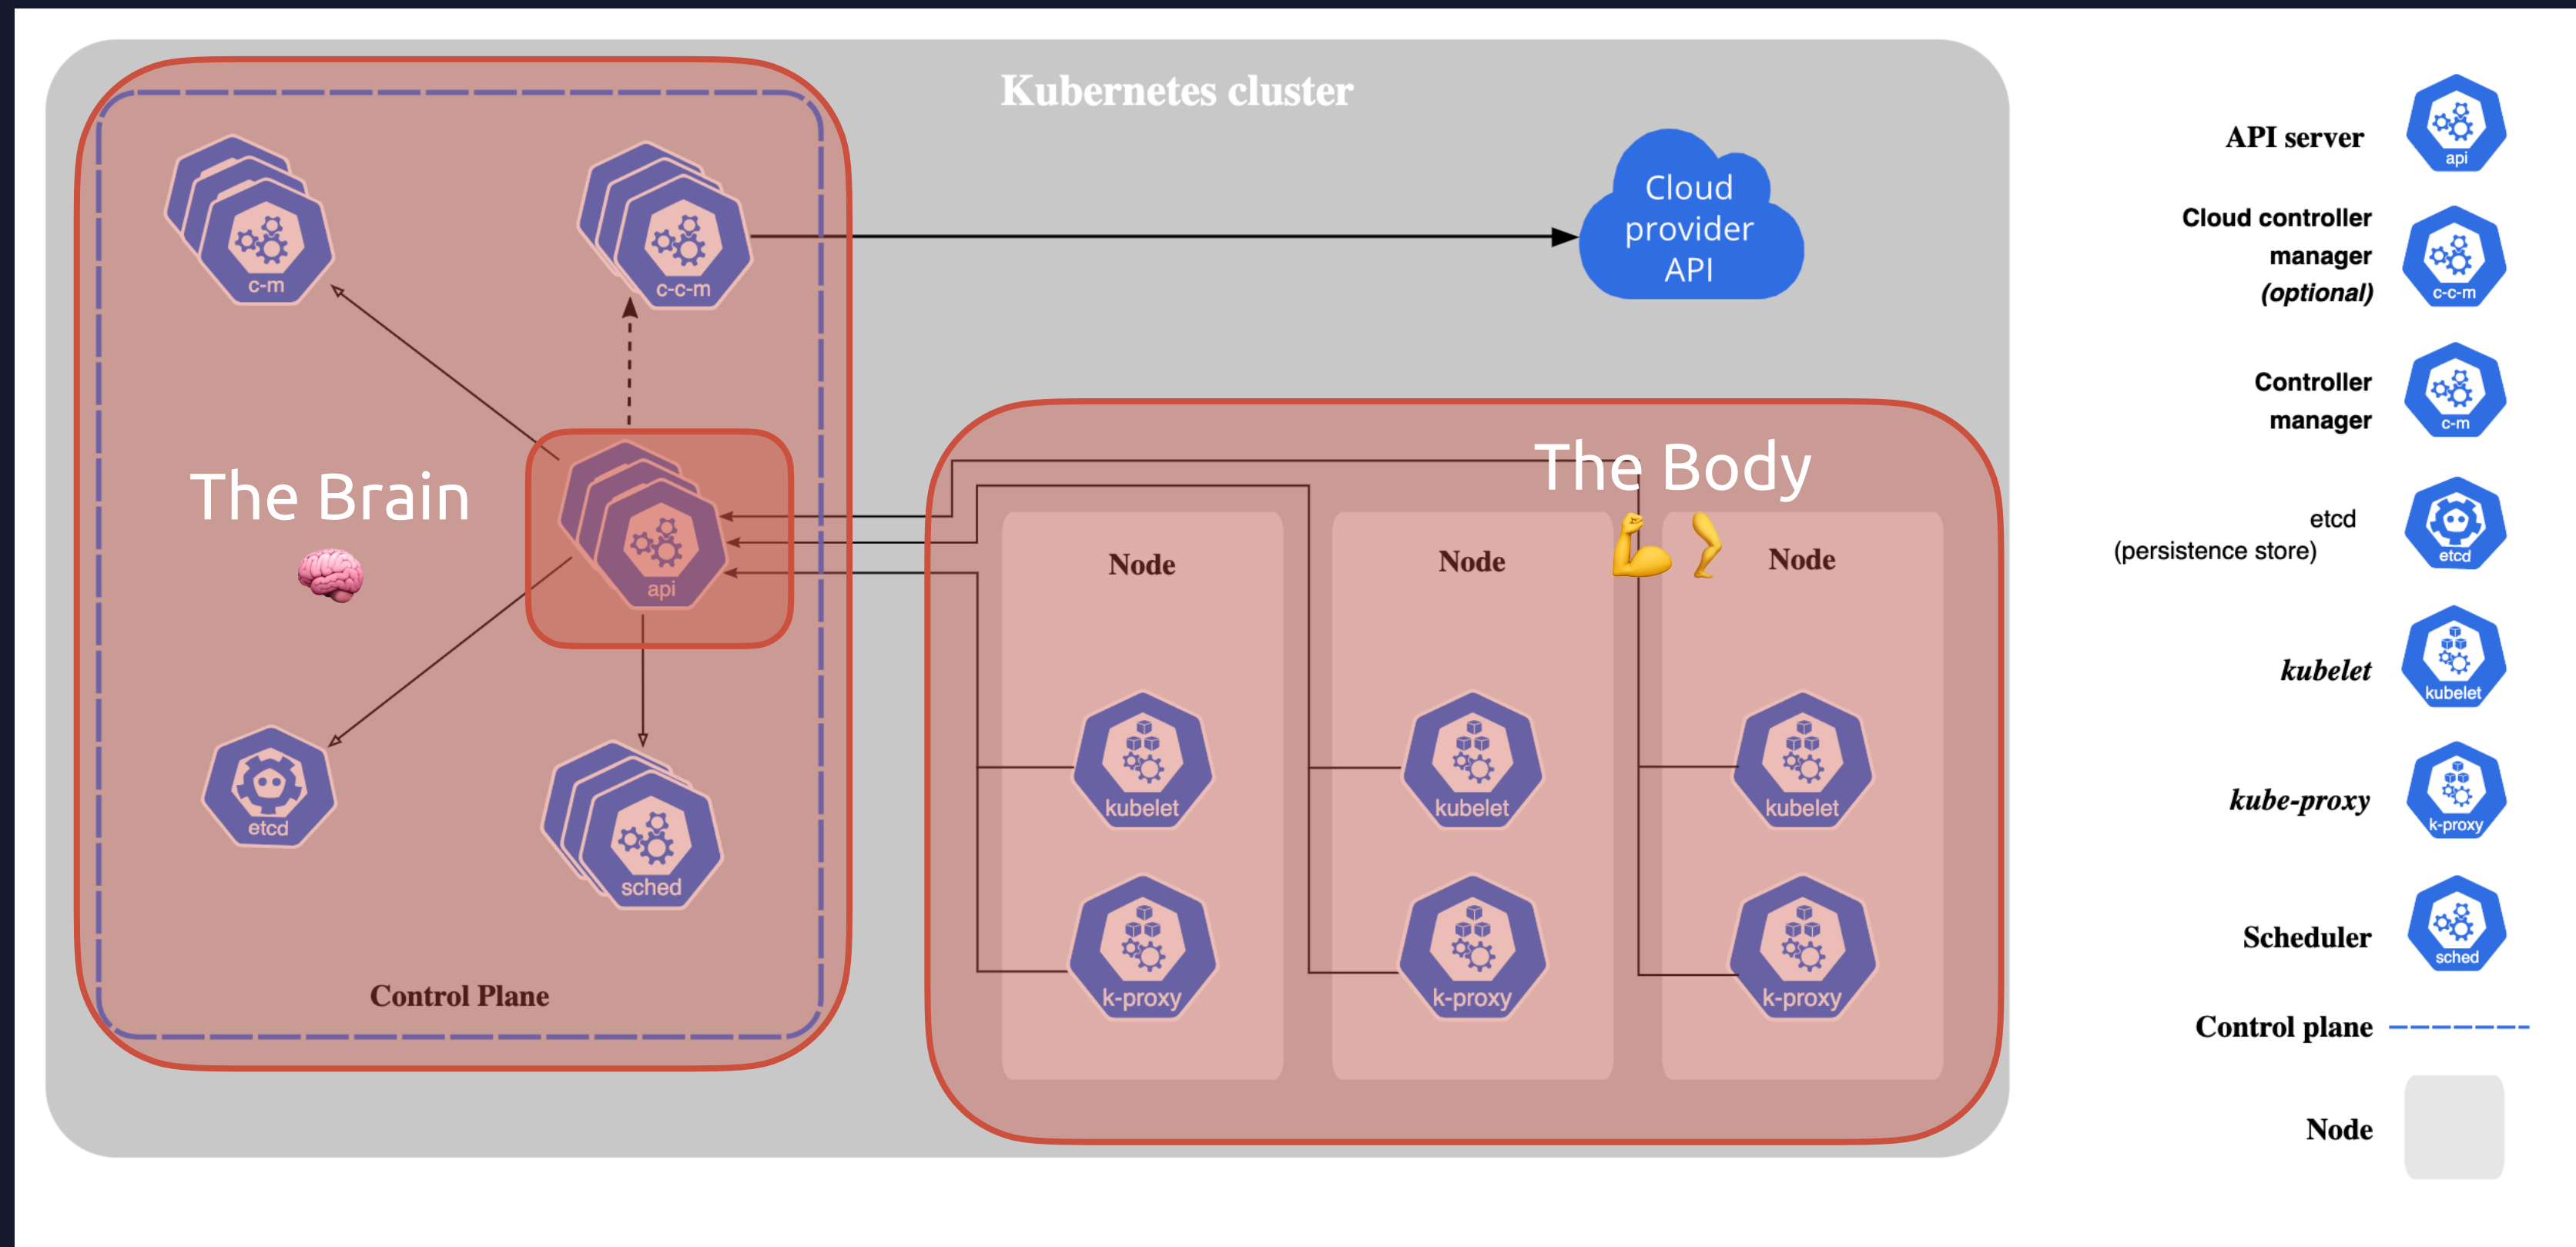


# KUBERNETES ARCHITECTURE

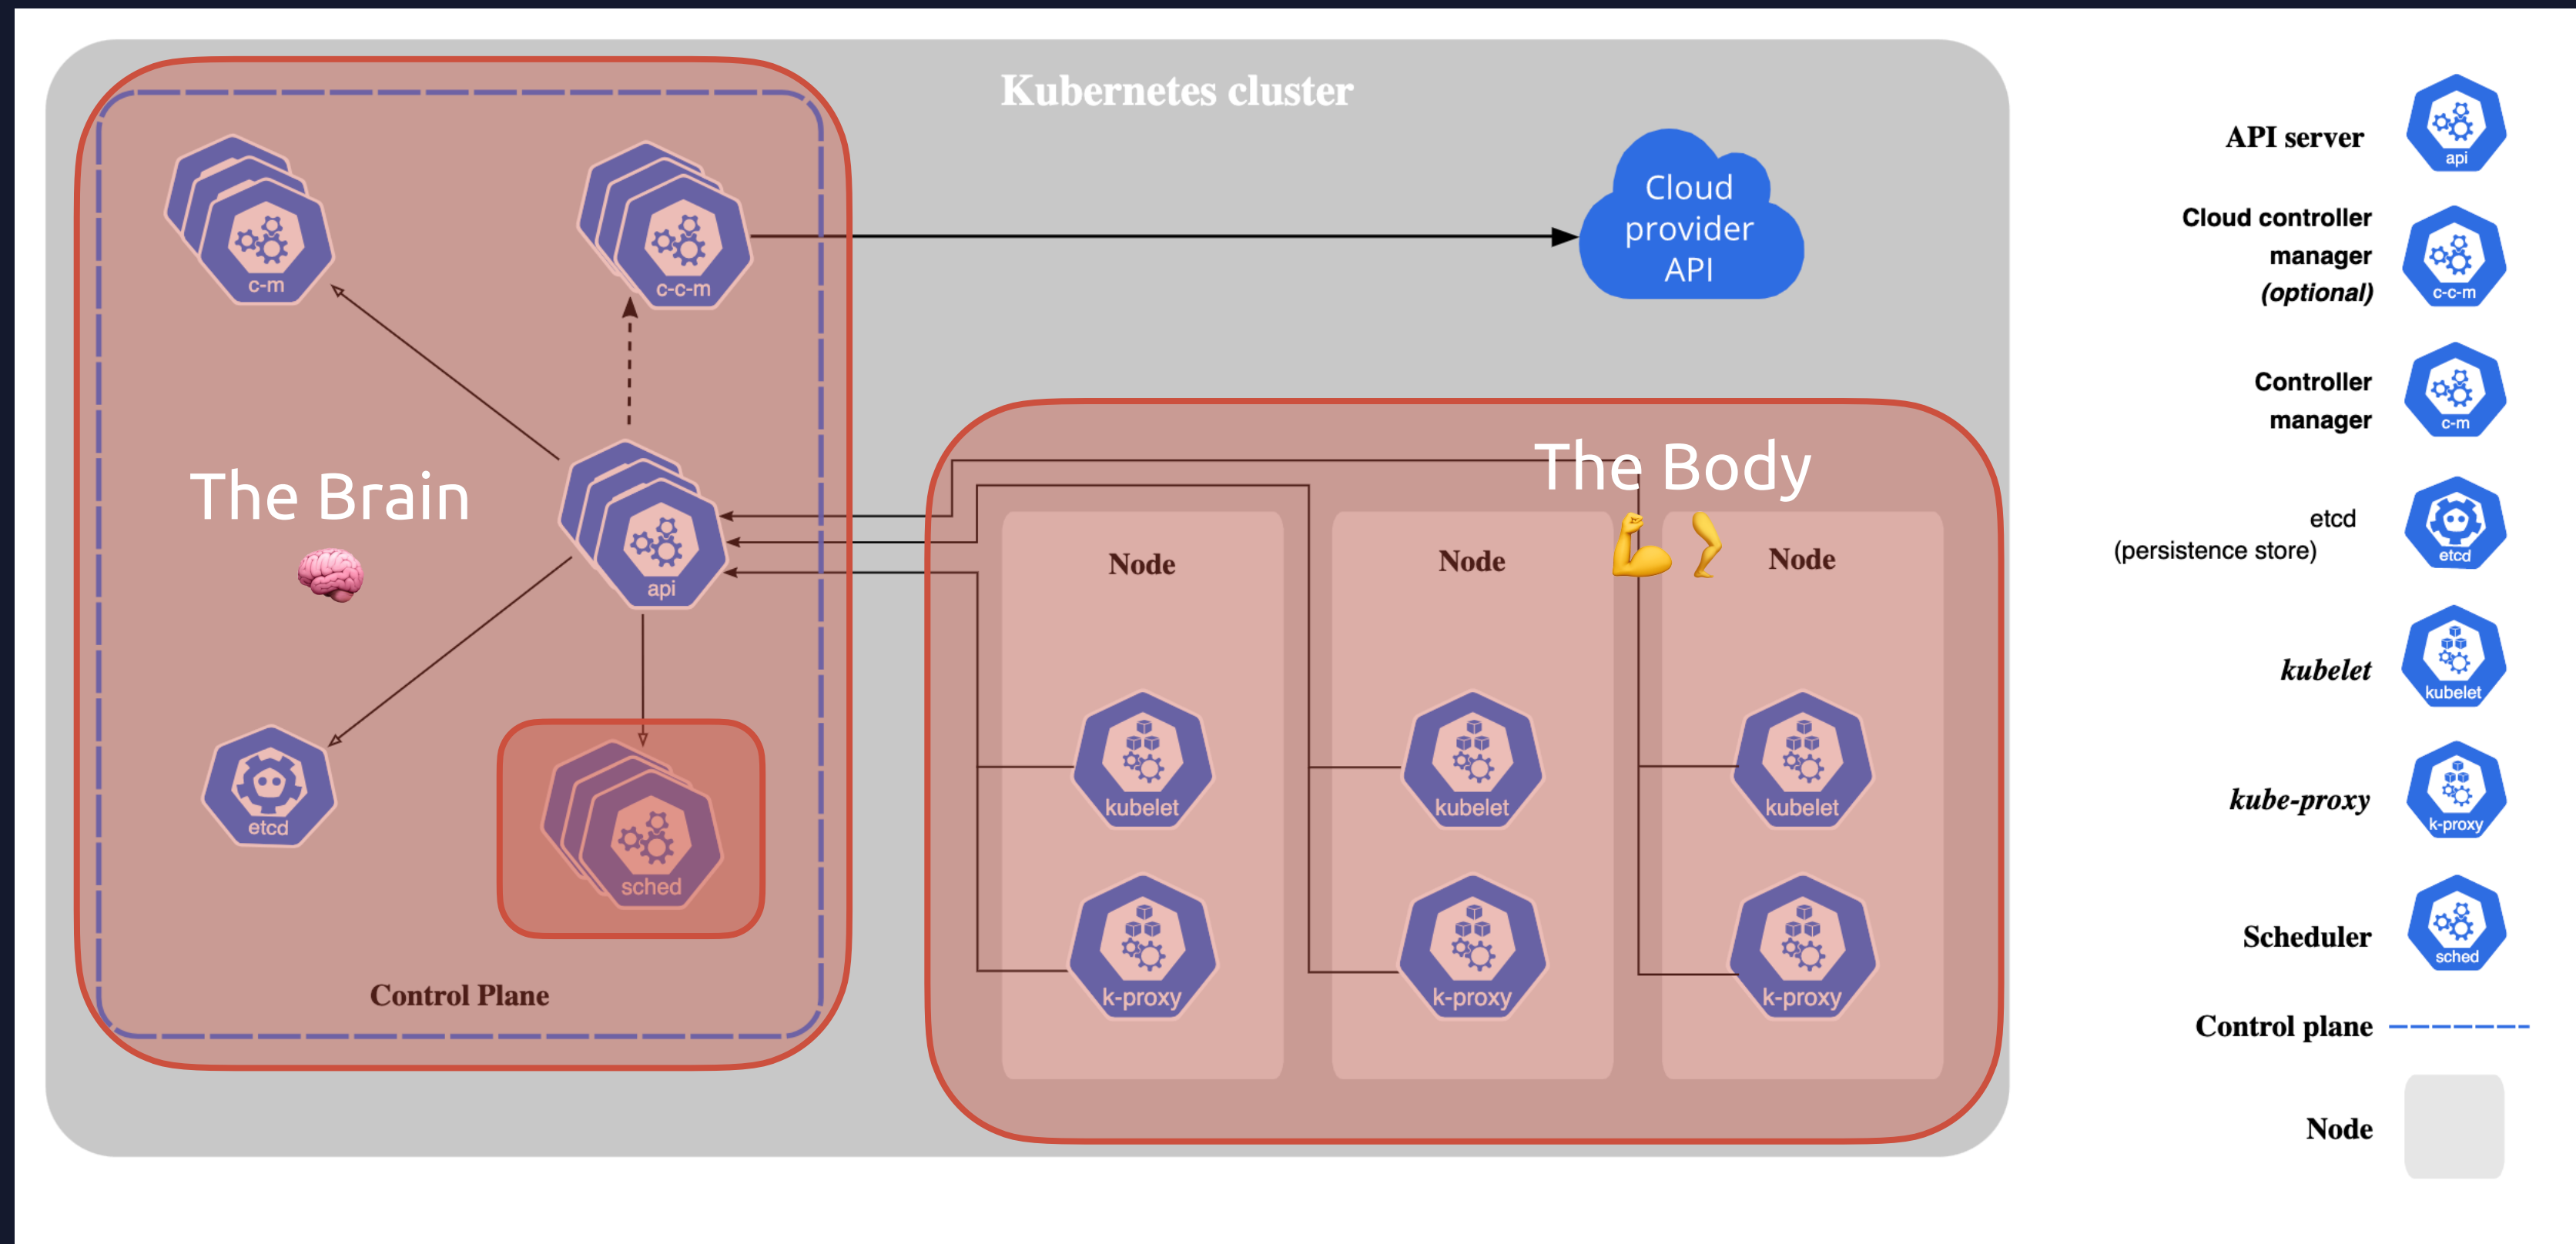




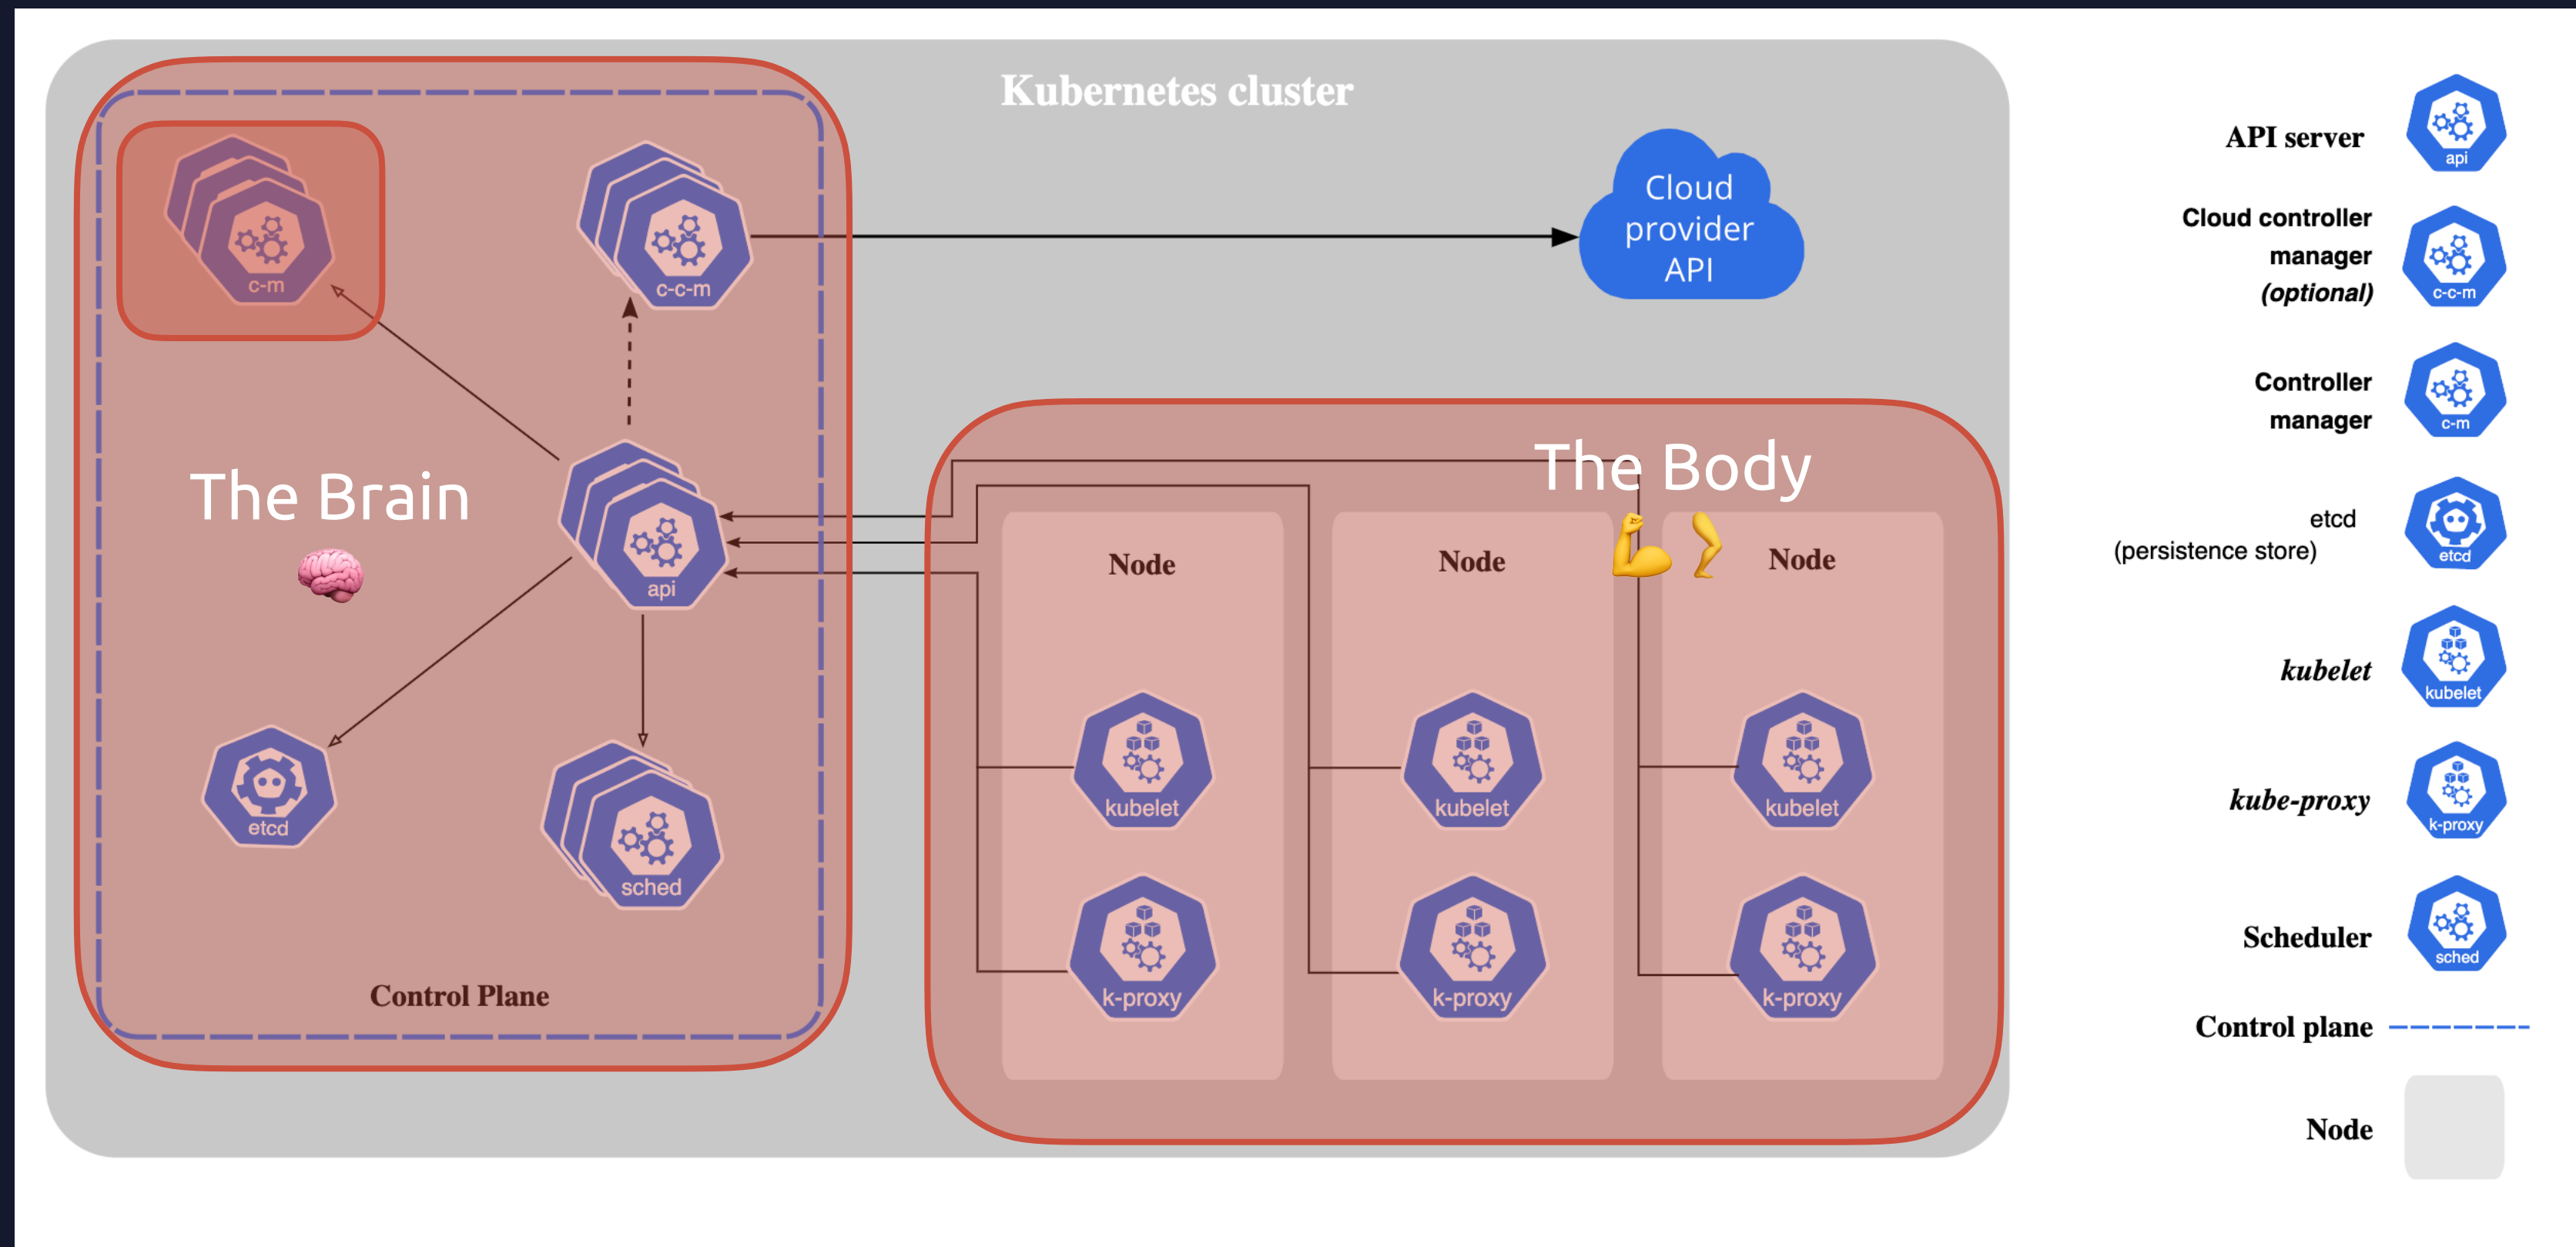
# KUBERNETES ARCHITECTURE



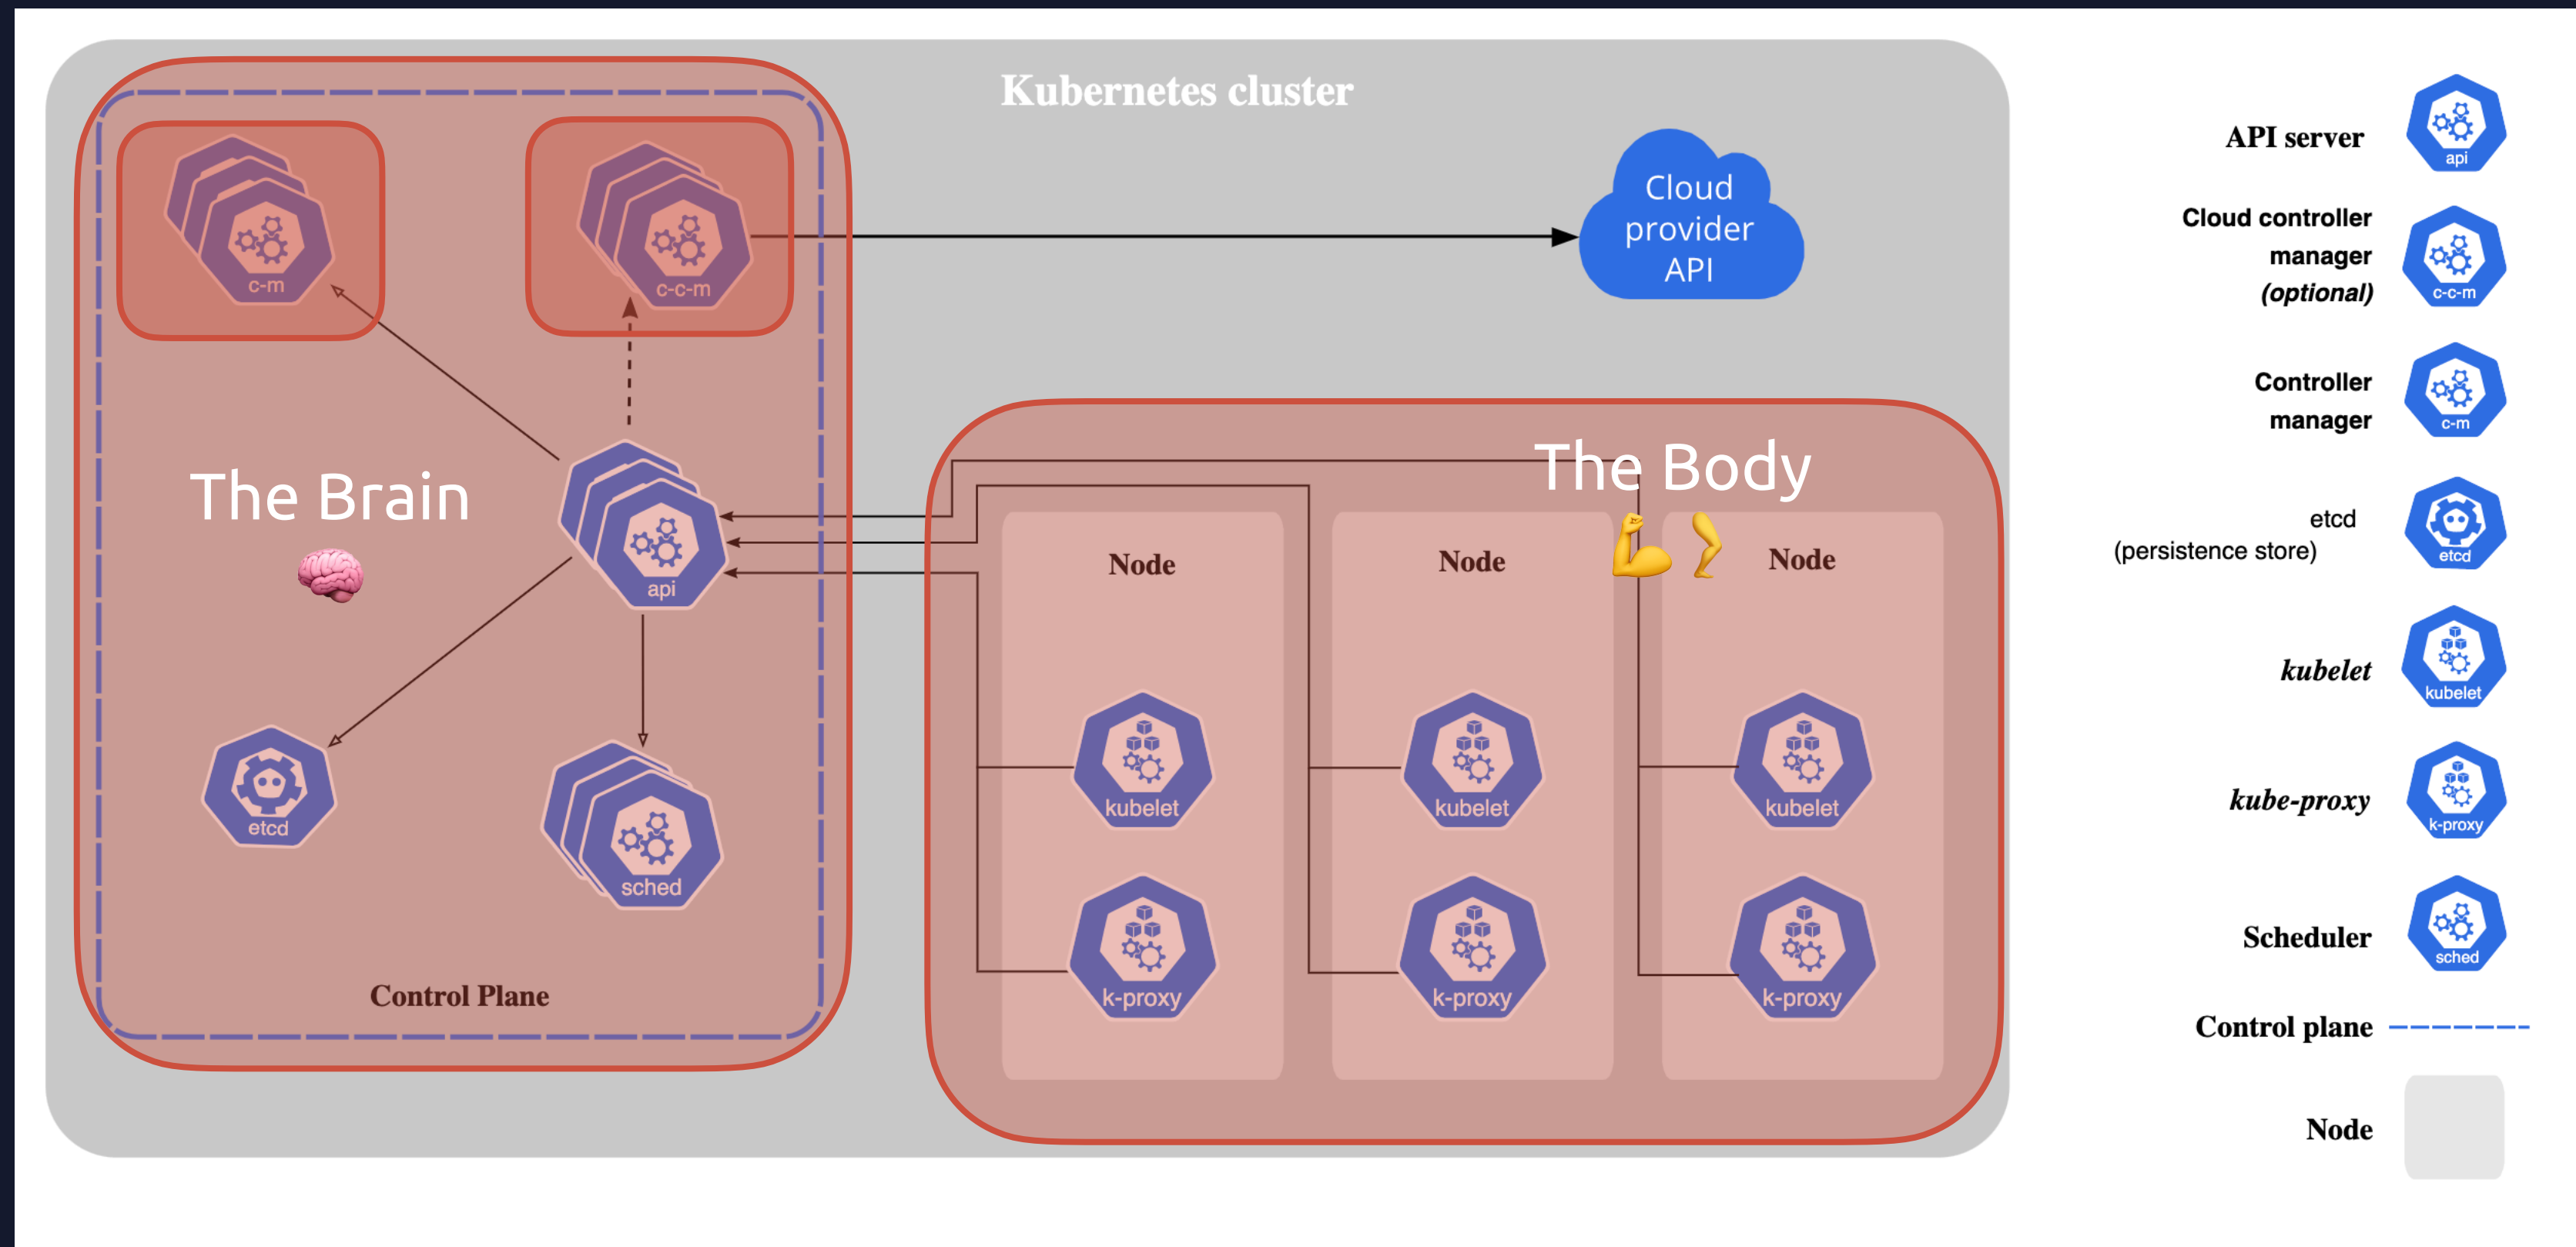
# KUBERNETES ARCHITECTURE



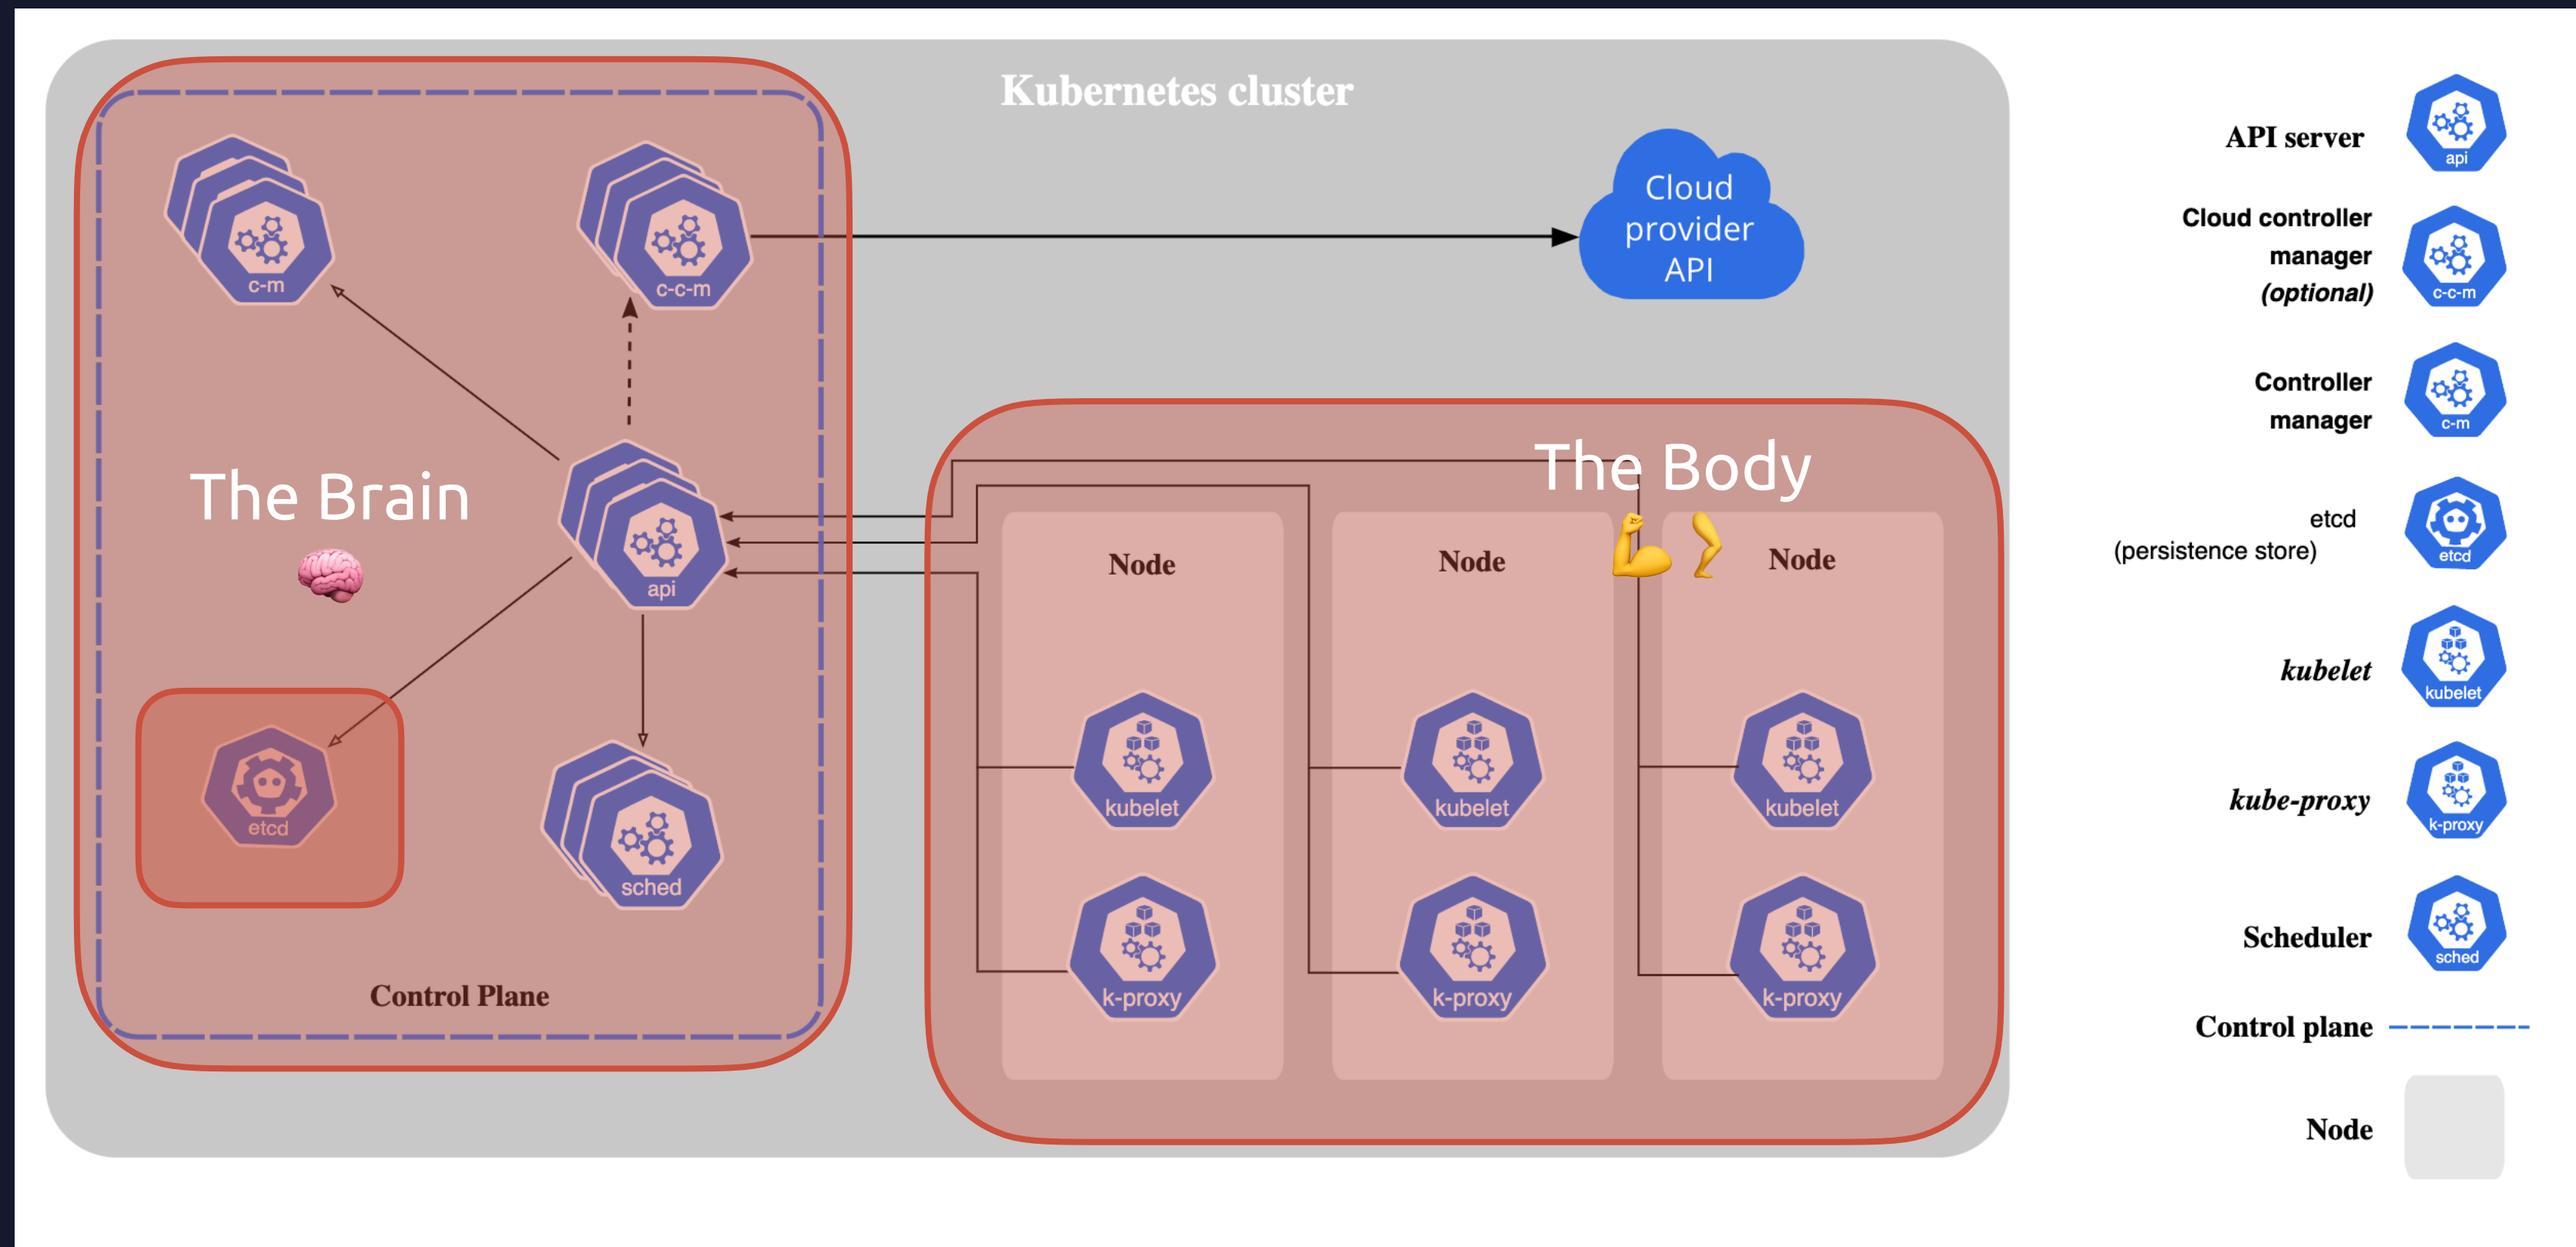
# KUBERNETES ARCHITECTURE



# KUBERNETES ARCHITECTURE

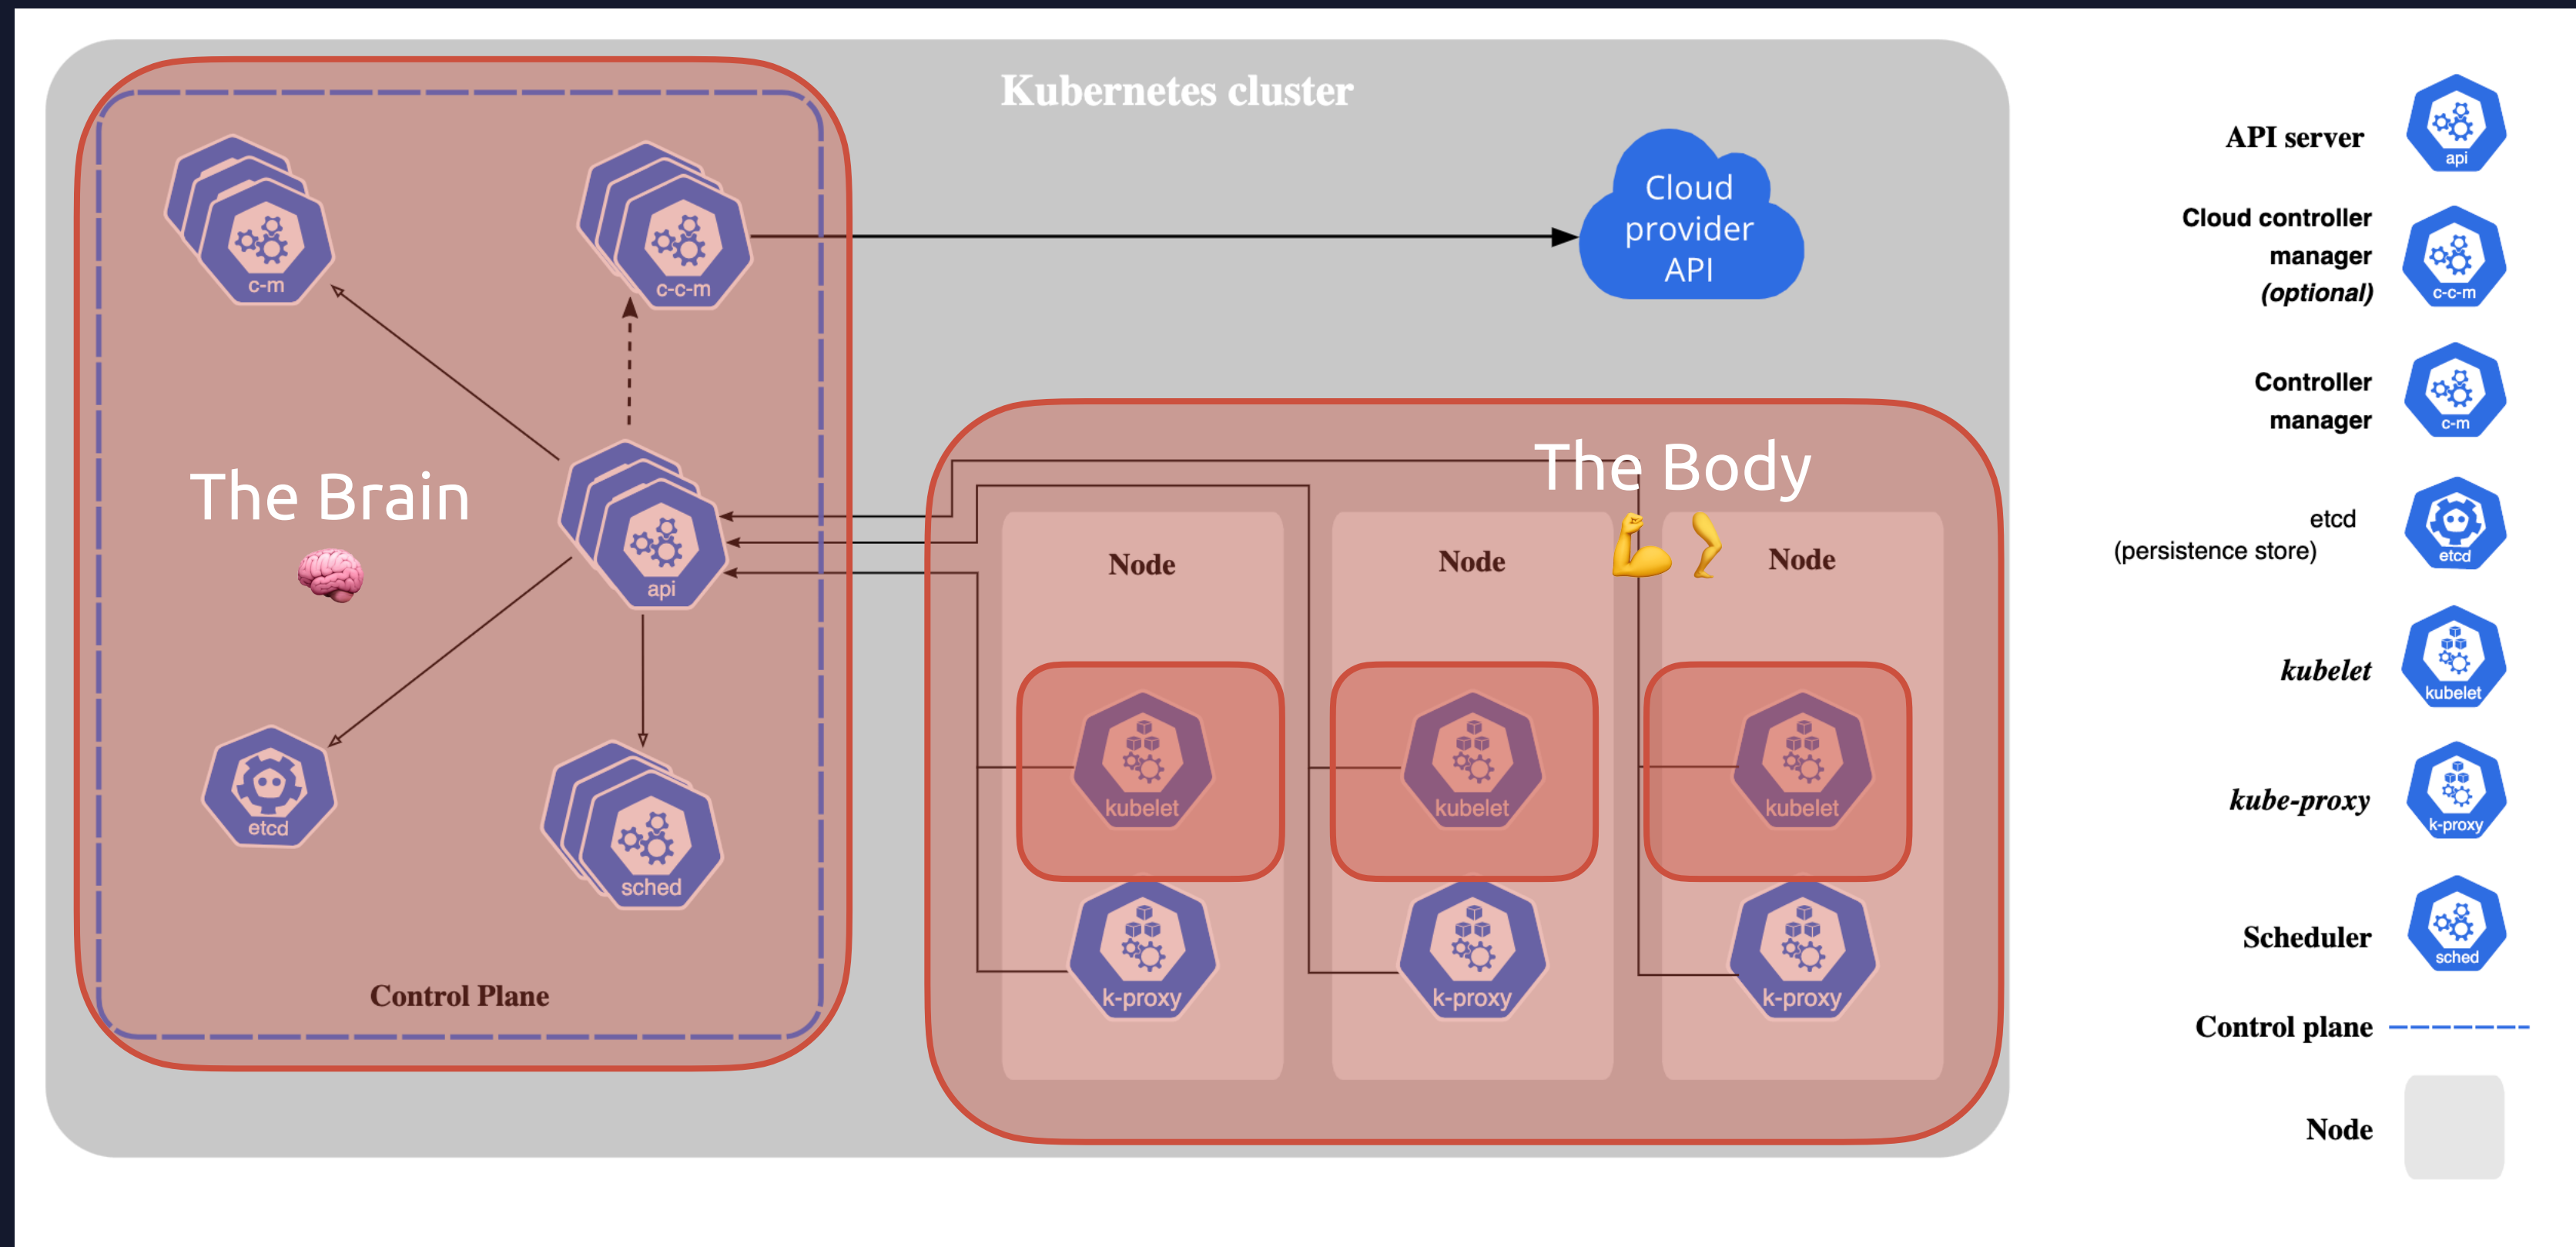


# KUBERNETES ARCHITECTURE

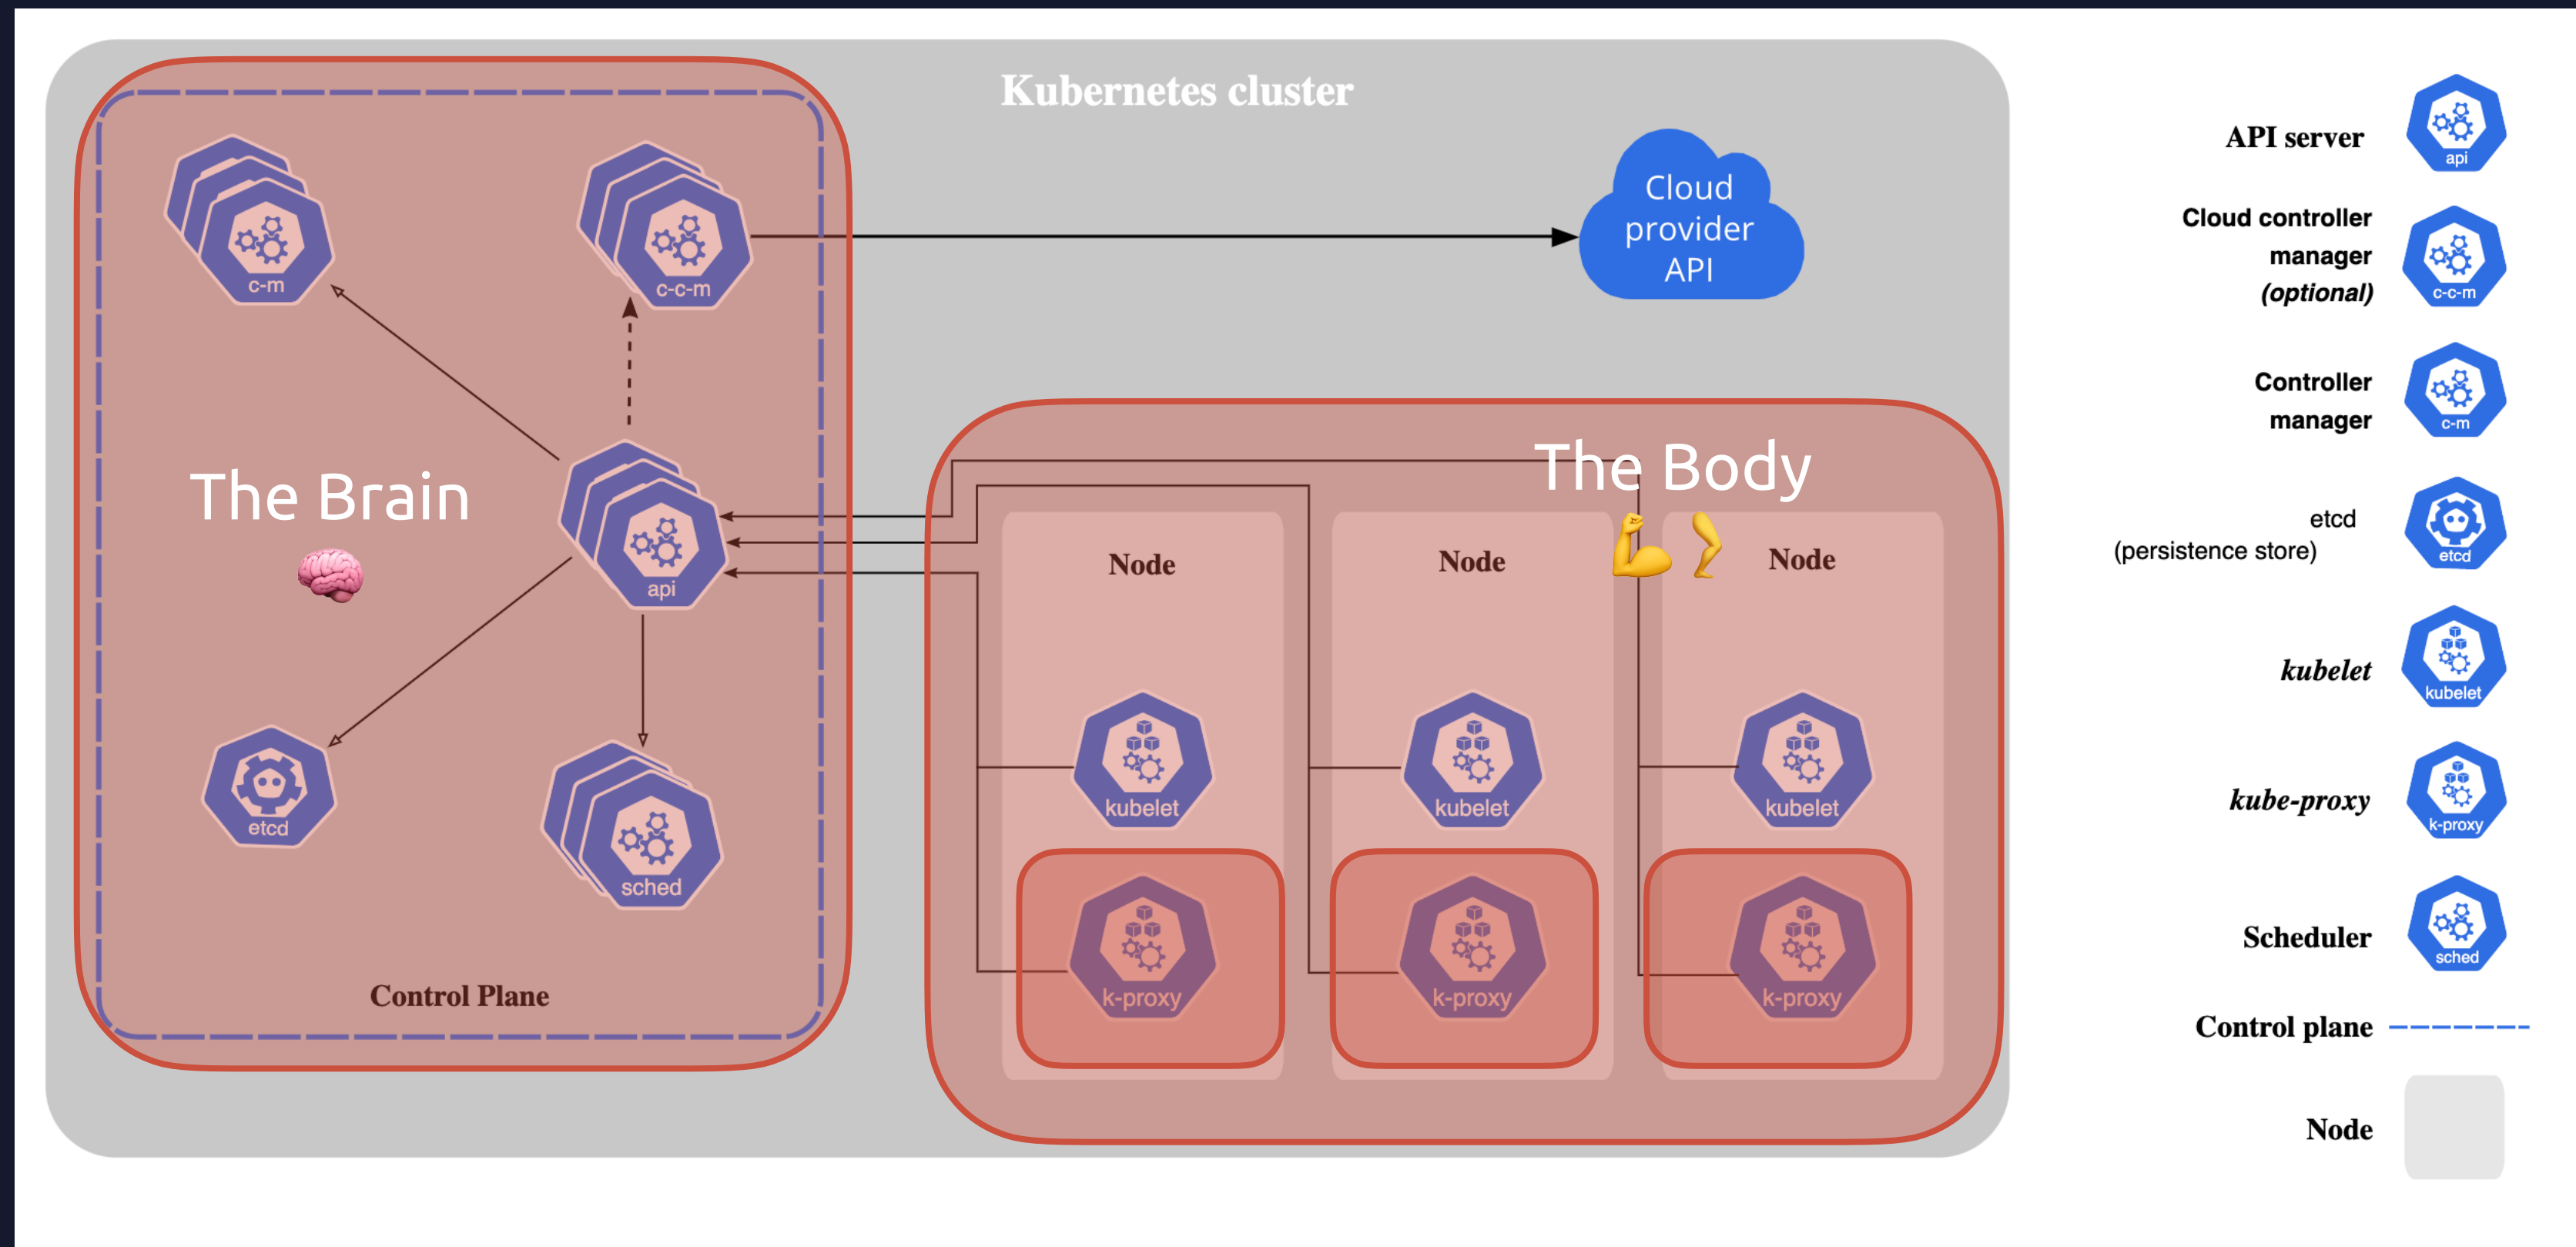




# KUBERNETES ARCHITECTURE



# KUBERNETES ARCHITECTURE



# THANK YOU!

