



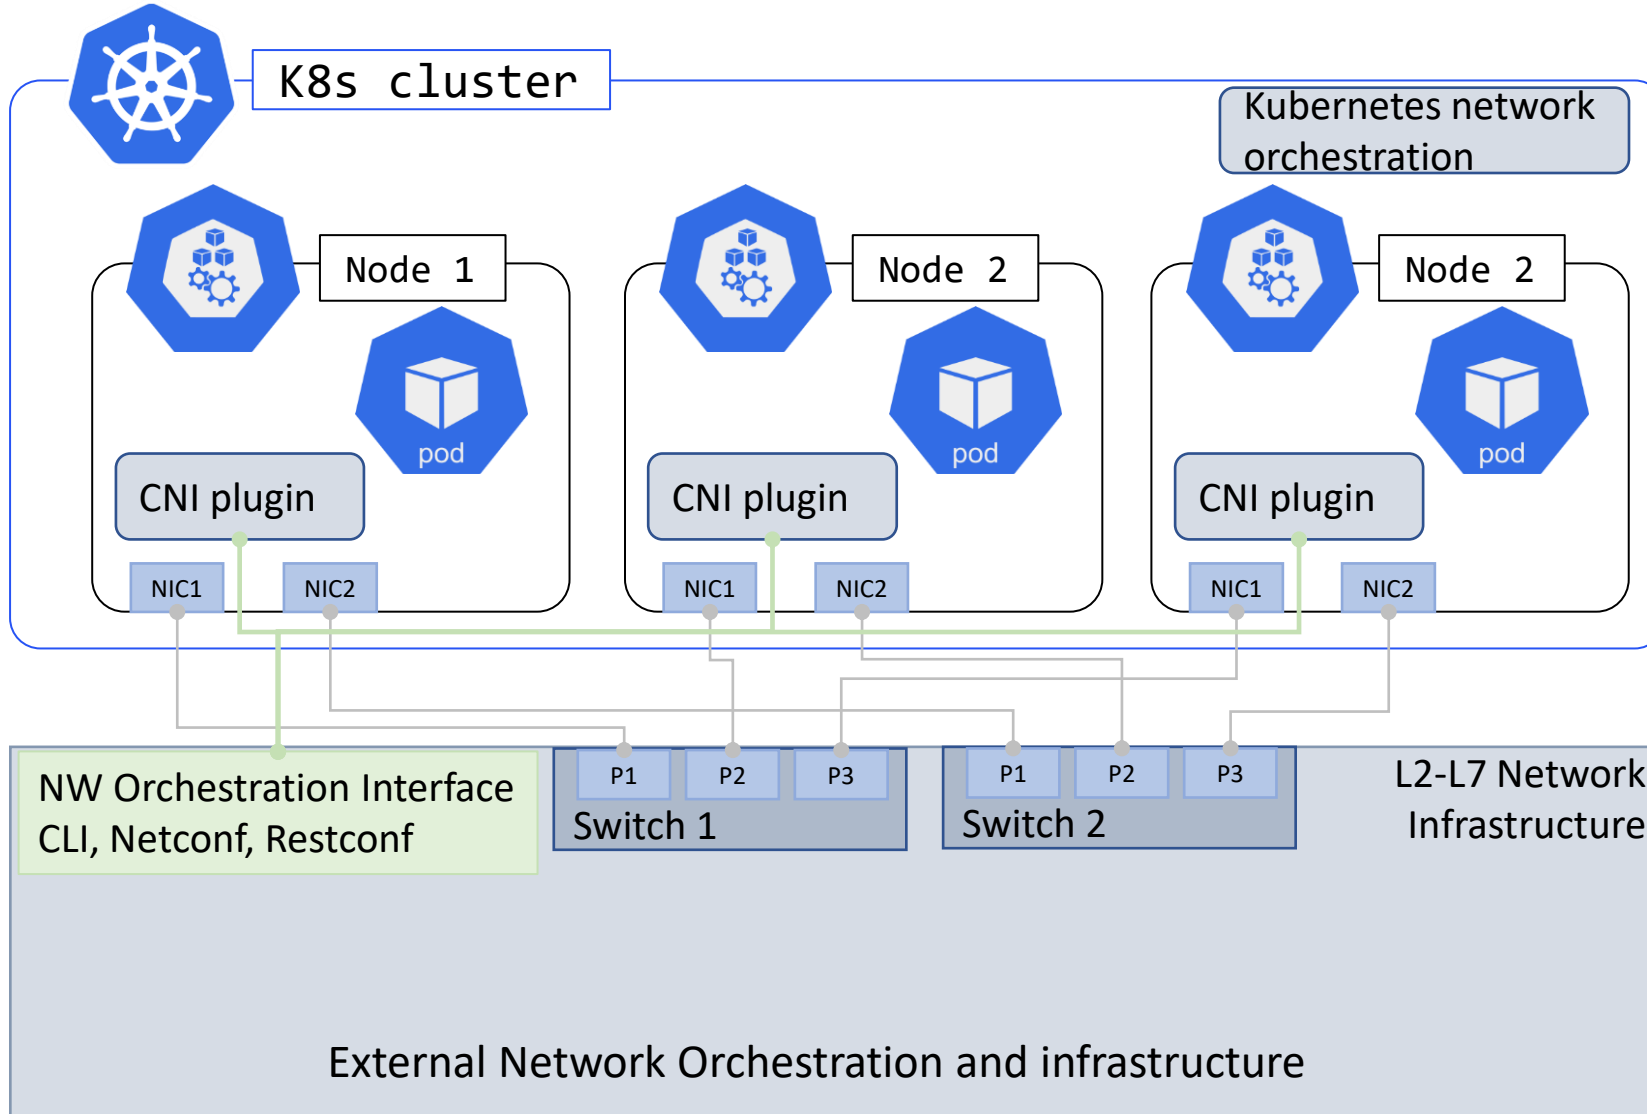
Kubernetes: Network orchestration

- Kubernetes proper does not support multi network orchestration
- Bare minimal NW orchestration provided by CNI plugins such as [Multus](#)
- Kubernetes Network plumbing group has defined the [Kubernetes Network Custom Resource Definition De-facto Standard](#)
- This de-facto standard defines
 - NetworkAttachmentDefinition: Makes a network known to the “Kubernetes” system
 - It is possible to add plugin specific information to the definition such as a network identifier from the network infrastructure serving the K8s cluster
 - Network Attachment Selection: Selects one or more secondary networks that a pod should be attached to.
 - This is specified inside an annotation inside the pod specification. Kubernetes proper is unaware of these networks and pod network attachments

```
apiVersion: "k8s.cni.cncf.io/v1"
kind: NetworkAttachmentDefinition
metadata:
  name: an-awesome-network
spec:
  config: '{
    "cniVersion": "0.3.0",
    "type": "awesome-plugin"
  }'
```

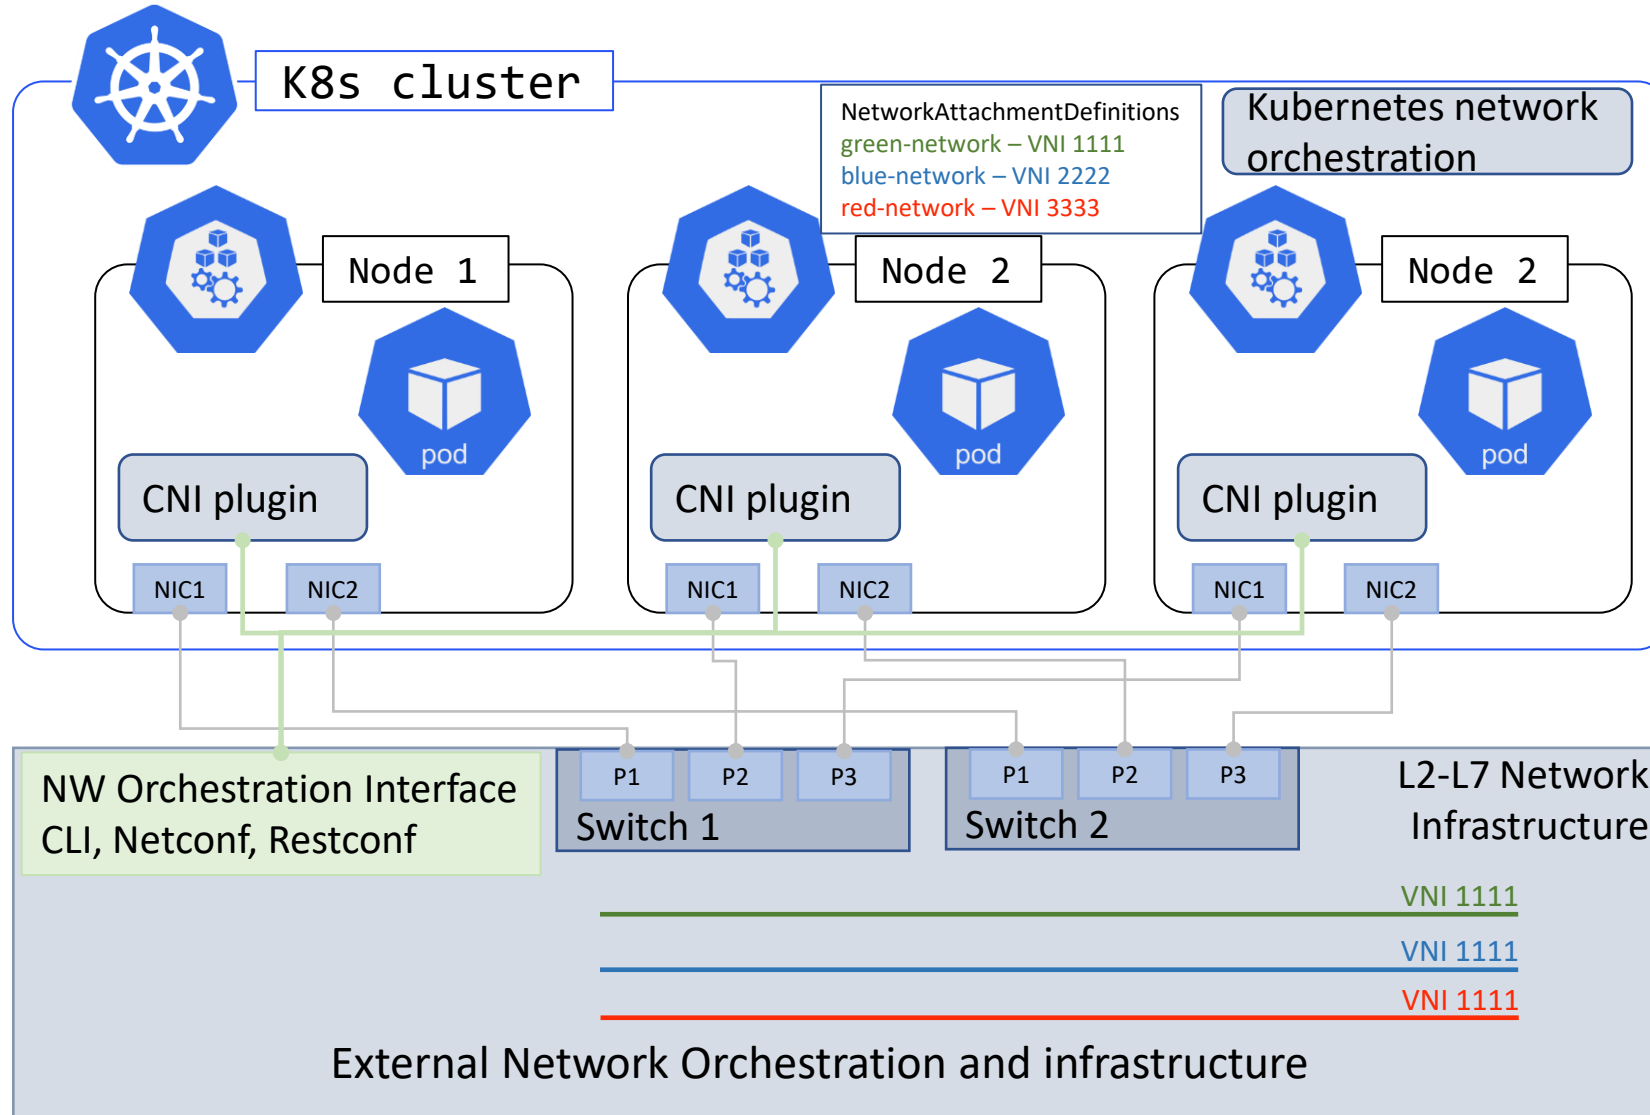
```
kind: Pod
metadata:
  name: my-pod
  namespace: my-namespace
  annotations:
    k8s.v1.cni.cncf.io/networks: net-a,an-awesome-network,other-ns/net-c
```

Orchestration of Pod NW attachment



- K8s network orchestration
 - CRUD NetworkAttachmentDefinition
 - CRUD pod Network Attachment Selection
- External NW orchestration and infrastructure
 - Proprietary/standardized interface?
- CNI plugin(s) aware of external infrastructure
 - Mapping of NIC port <-> switch port
 - Static or dynamic discovery
 - Mapping of NetworkAttachment-Definition <-> external VNI

Orchestration of Pod NW attachment



➤ Port mappings

- Server1-NIC1 <-> sw1p1
- Server1-NIC2 <-> sw2p1
- Server2-NIC1 <-> sw1p2
- Server2-NIC2 <-> sw2p2
- Server3-NIC1 <-> sw1p3
- Server3-NIC2 <-> sw2p3

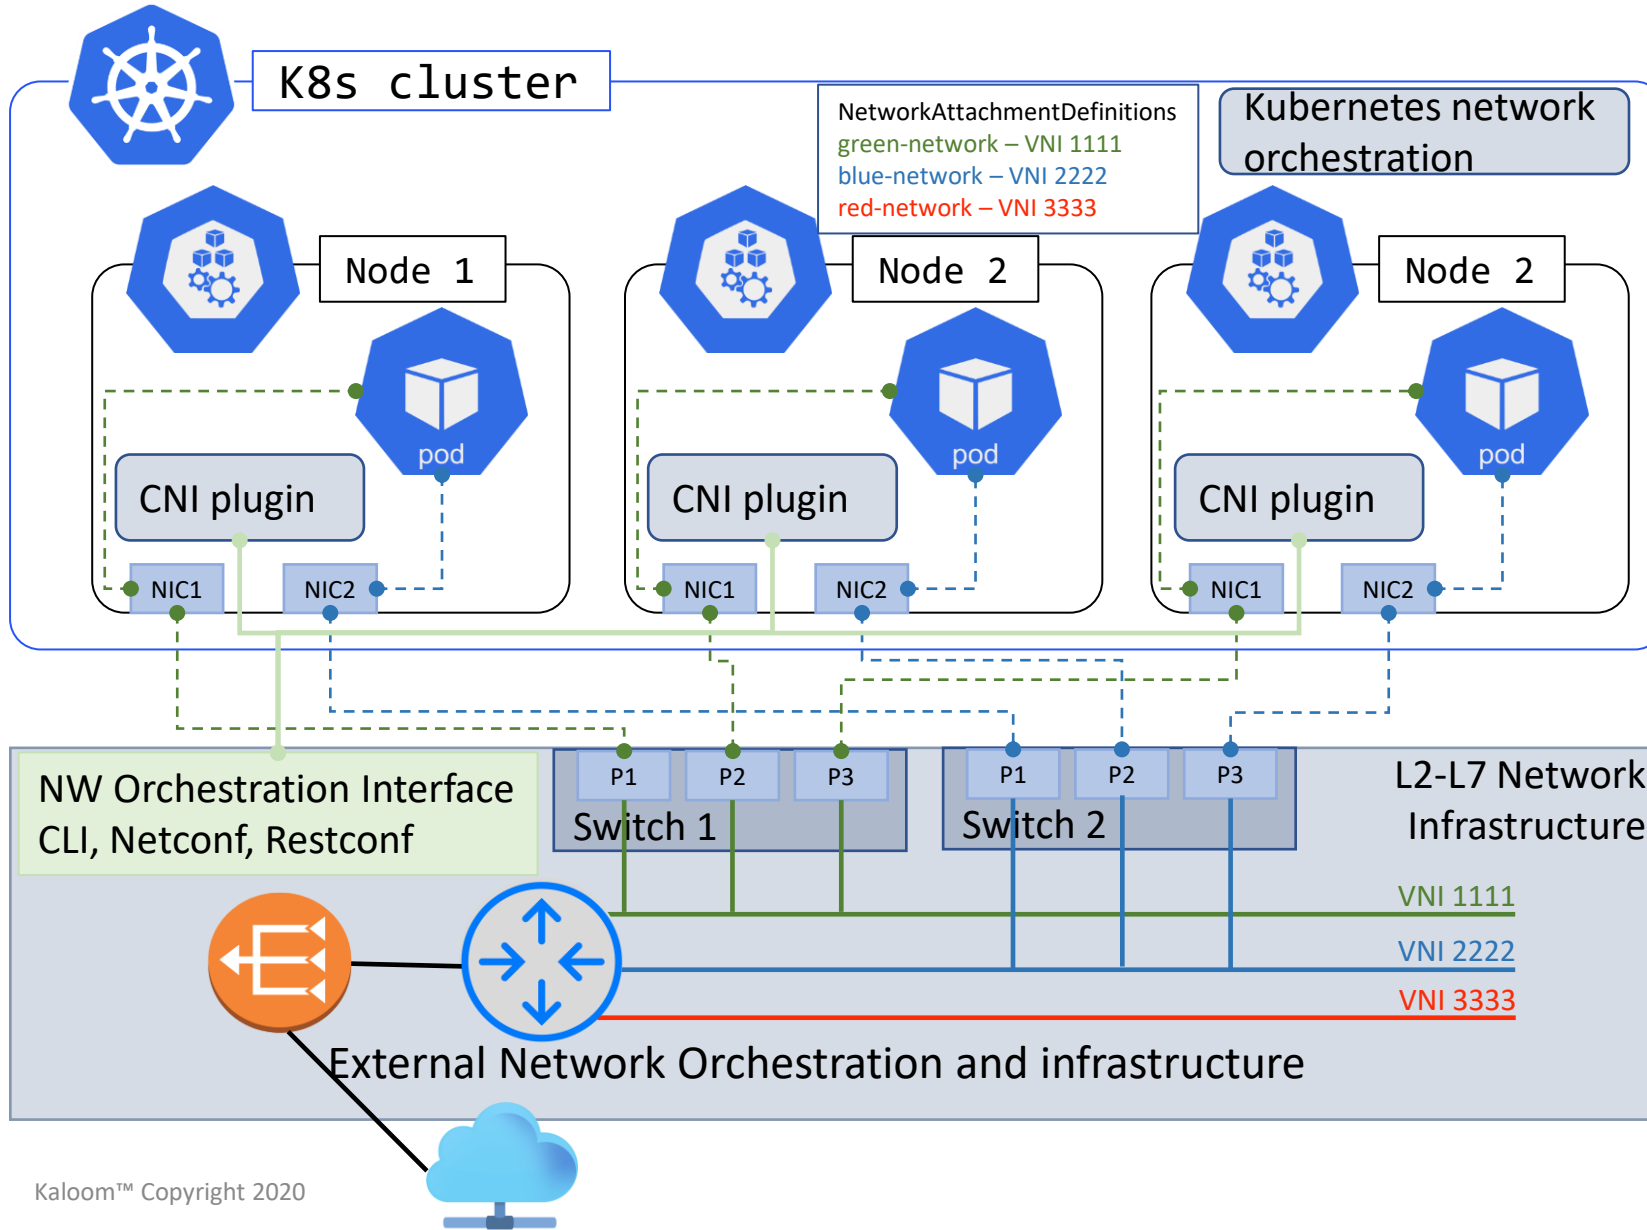
➤ Create 3 networks in the external infrastructure

- Blue Network – VNI 1111
- Green Network – VNI 2222
- Red Network – VNI 3333

➤ Create 3 Network Attachment Definitions in K8s that maps to these Networks

```
apiVersion: "k8s.cni.cncf.io/v1"
kind: NetworkAttachmentDefinition
metadata:
  name: green-network
spec:
  config: '{
    "cniVersion": "0.3.0",
    "type": "awesome-plugin"
    "VNI": "1111"
  }'
```

Orchestration of other network services



➤ Attach pods to networks

```
kind: Pod
metadata:
  name: pod
  namespace: my-namespace
  annotations:
    k8s.v1.cni.cncf.io/networks: blue-
network,green-network
```


➤ CNI plugin will

- Use network orchestration API/CLI to attach the blue and green network to the correct switch ports
 - (Simplified)
- Attach pods to networks

➤ Network Orchestration Interface of external infrastructure used for everything else

- Routers
- NW stitching
- NAT
-

External Network orchestration and infrastructure



➤ Attach pods to networks

```
kind: Pod
metadata:
  name: pod
  namespace: my-namespace
  annotations:
    k8s.v1.cni.cncf.io/networks: blue-
network,green-network
```

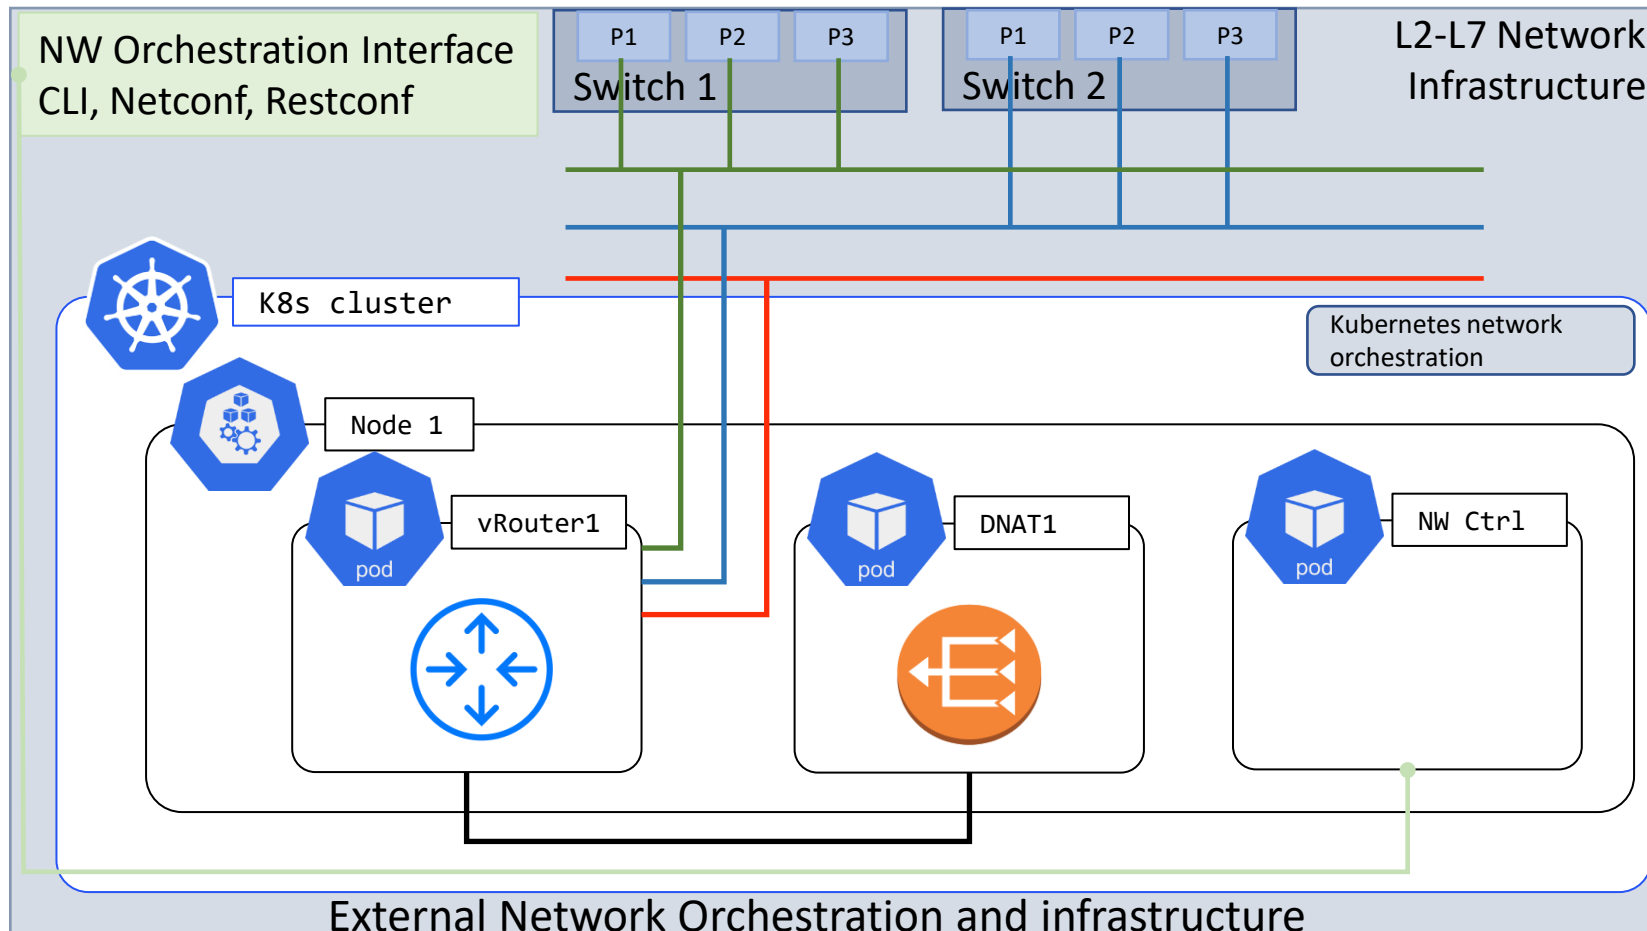
➤ CNI plugin will

- Use network orchestration API/CLI to attach the blue and green network to the correct switch ports
 - (Simplified)
- Attach pods to networks

➤ Network Orchestration Interface of external infrastructure used for everything else

- Routers
- NW stitching
- NAT
-

External Network orchestration and infrastructure



> Kubernetes based

- The external network orchestration infrastructure system can in itself be CNI/Kubernetes based

> “Internalized” K8s application

- The network orchestration and infrastructure system can even be hosted inside the k8s cluster and provide network functionality for the hosted CNI applications