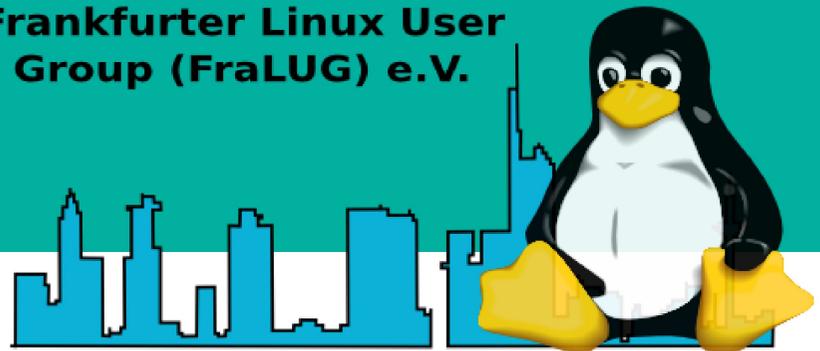


K8s - Service Meshes

Istio und Network Service Mesh

**Frankfurter Linux User
Group (FraLUG) e.V.**



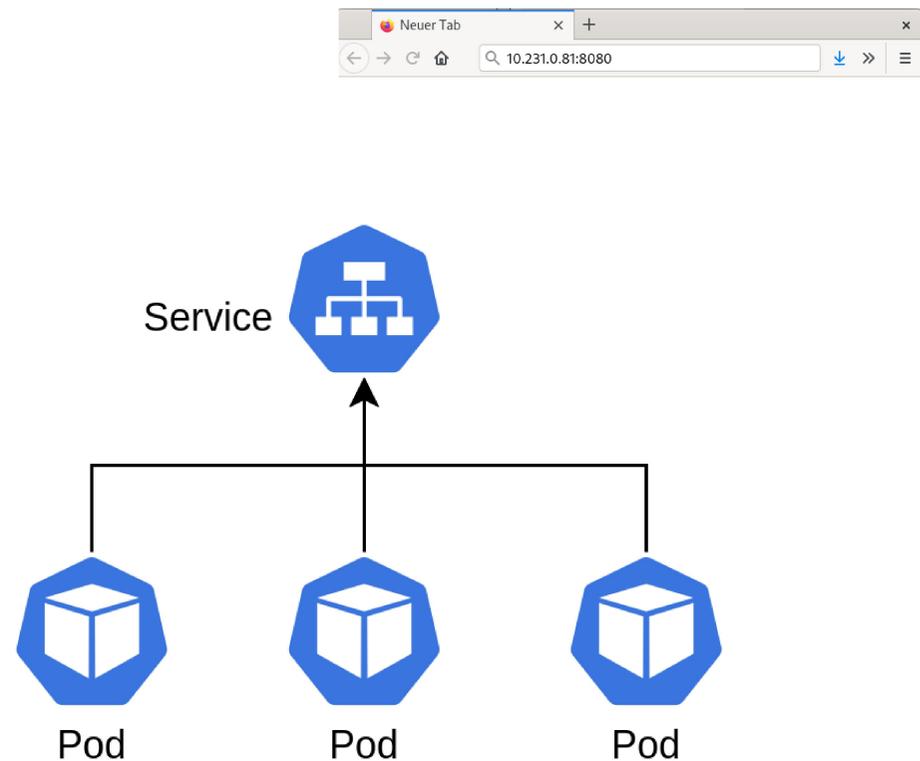
Kubernetes - Networking

Networking-Herausforderungen von Microservice-Architekturen

- Application Exposure
- Microservices müssen untereinander kommunizieren
- Kommunikation muss sicher sein
- Metrics
- L3 intradomain Connectivity
- L3 interdomain Connectivity

Basic Application Exposure

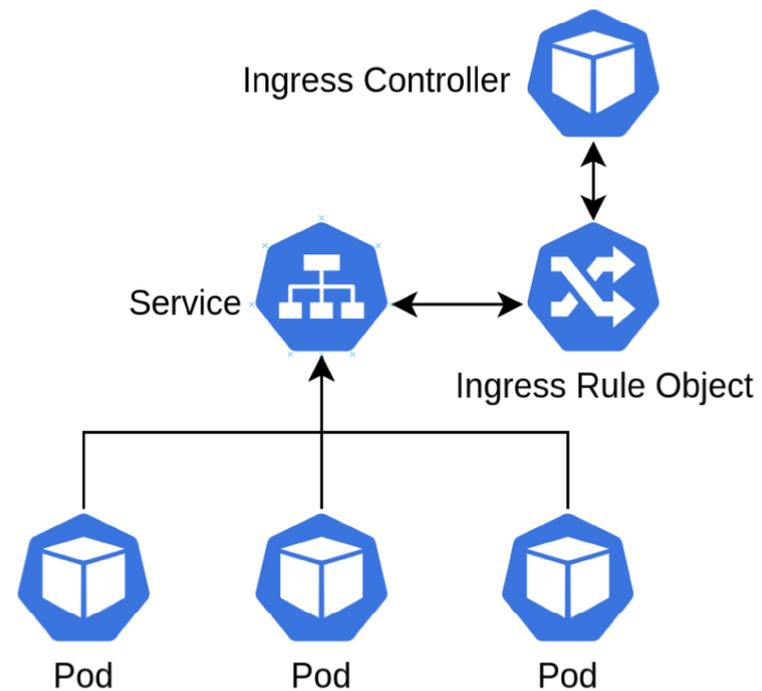
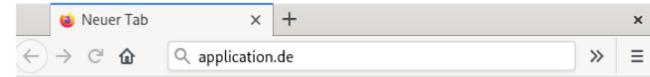
- Deployments/Stateful Sets sind nur lokal erreichbar
- Services machen Deployments/Stateful Sets via IP/Port verfügbar



Basic Application Exposure

INGRESS

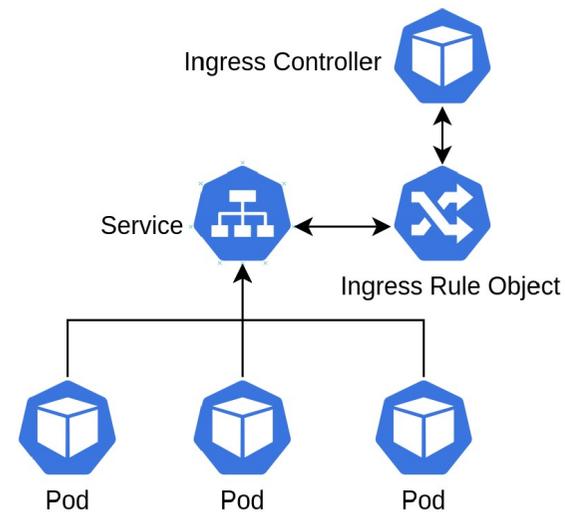
- INGRESS mappt Hostnames und URLs auf Services



Basic Application Exposure

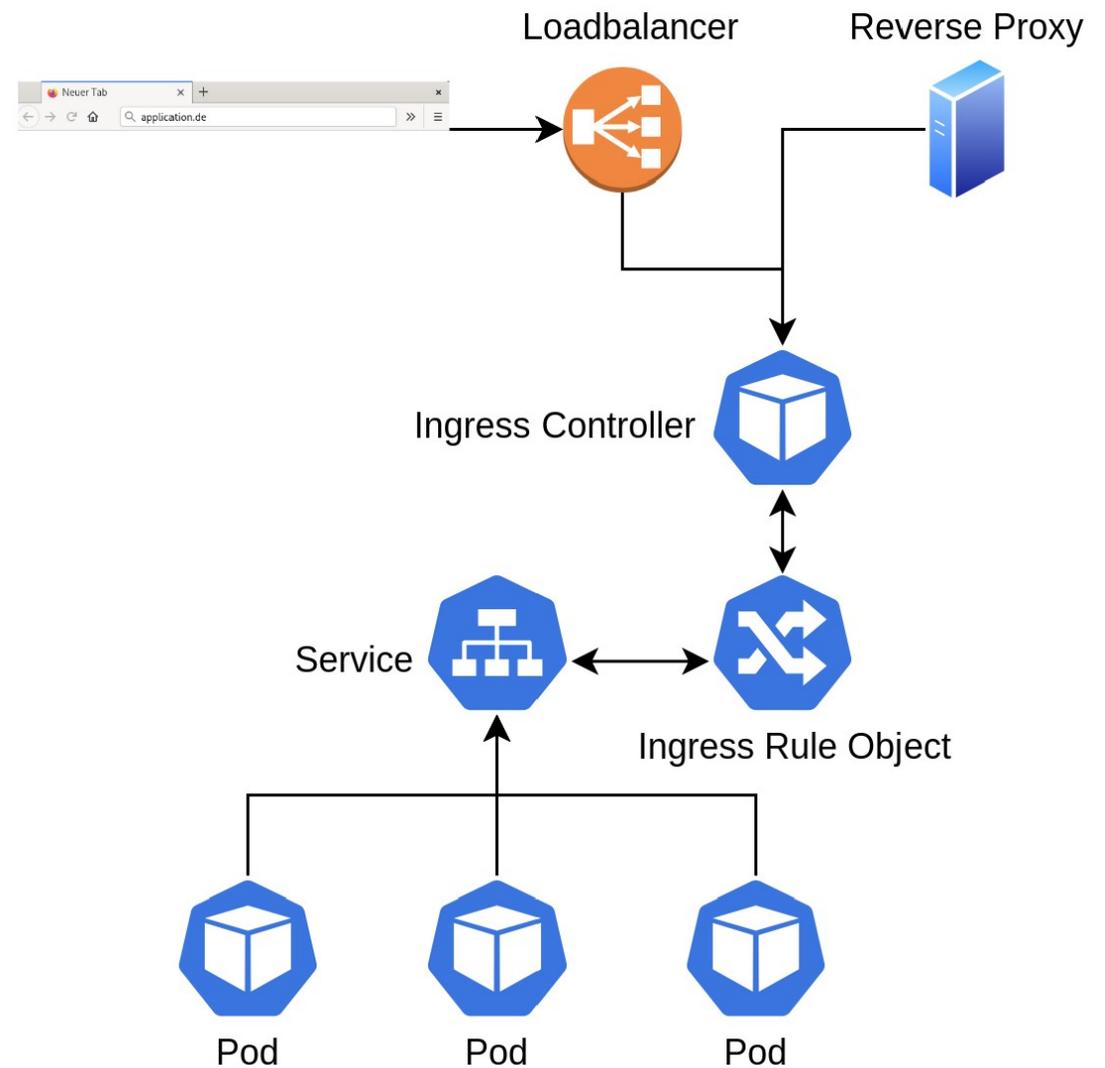
INGRESS

- ```
apiVersion: networking.k8s.io/v1
kind: Ingress
metadata:
 name: application-ingress
spec:
 rules:
 - host: application.de
 http:
 paths:
 backend:
 serviceName: nginx-service
 servicePort: 8080
```



# Basic Application Exposure

## INGRESS





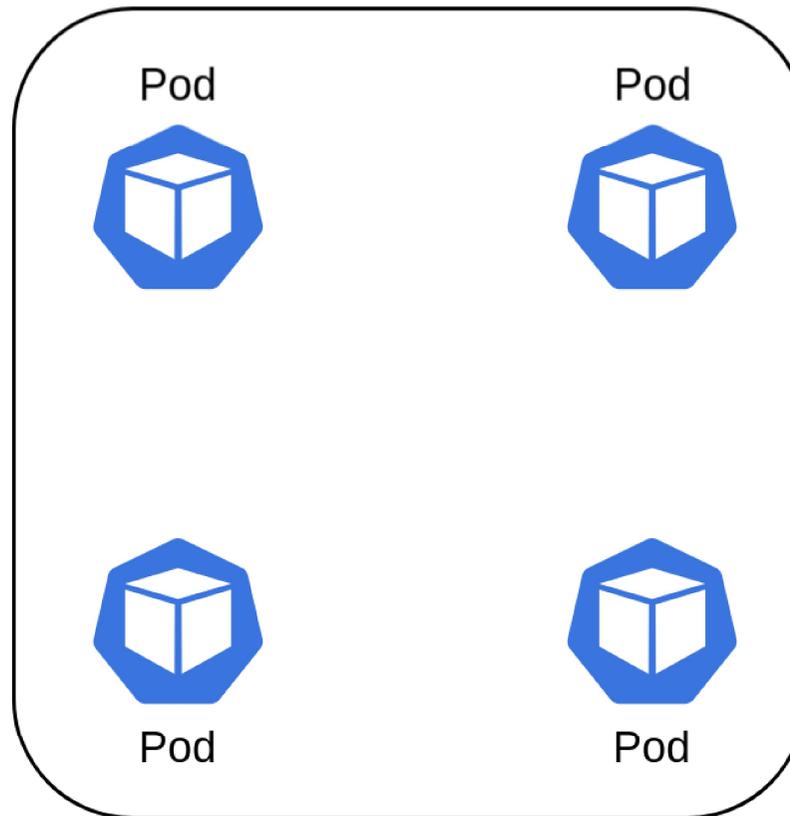
# Istio

L4/L7 Application Service Mesh

# Microservice- Challenges

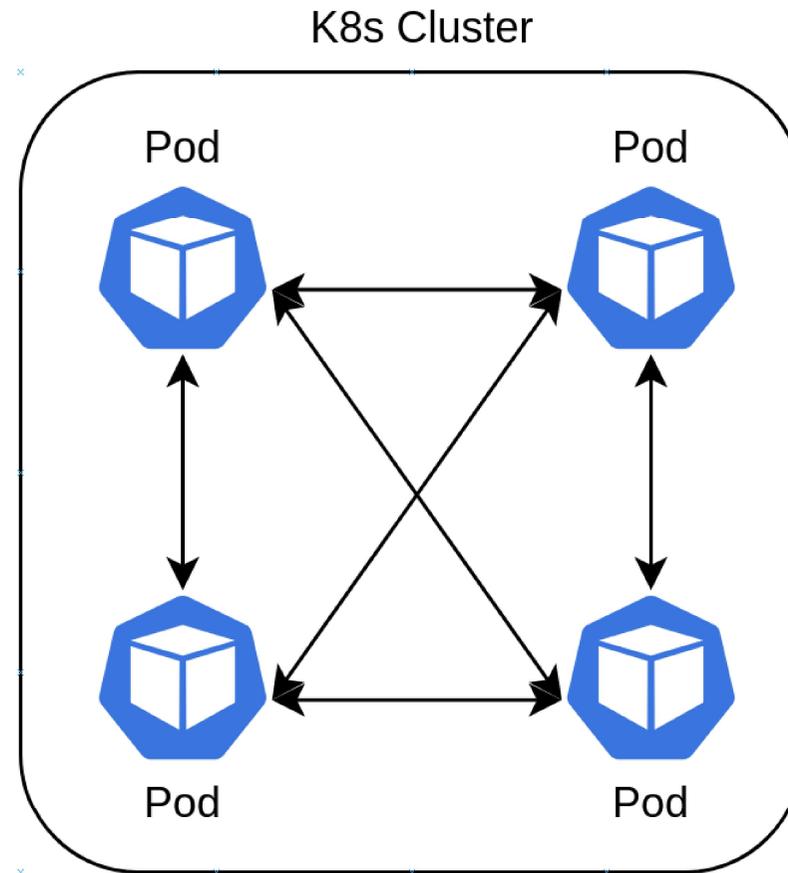
- Keine Traffic Restriktionen

K8s Cluster



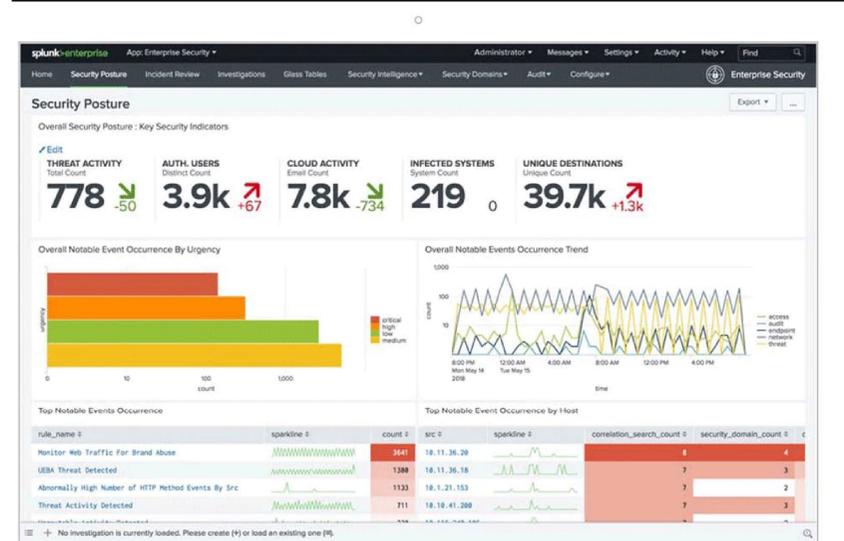
# Microservice- Challenges

- Keine Traffic Restriktionen



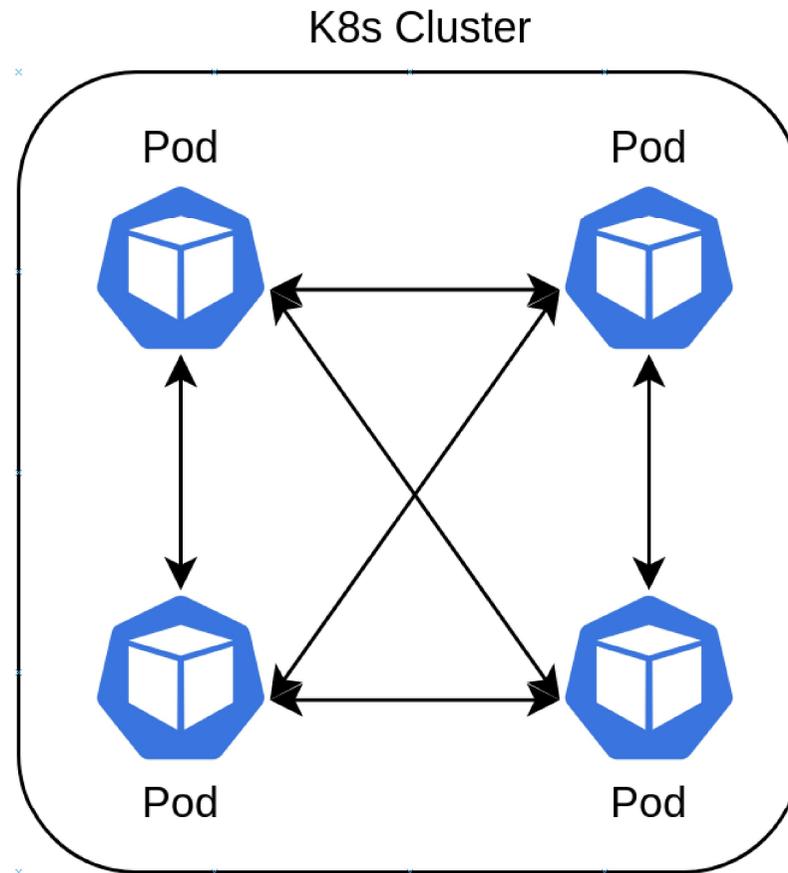
# Microservice- Challenges

- Keine Metriken



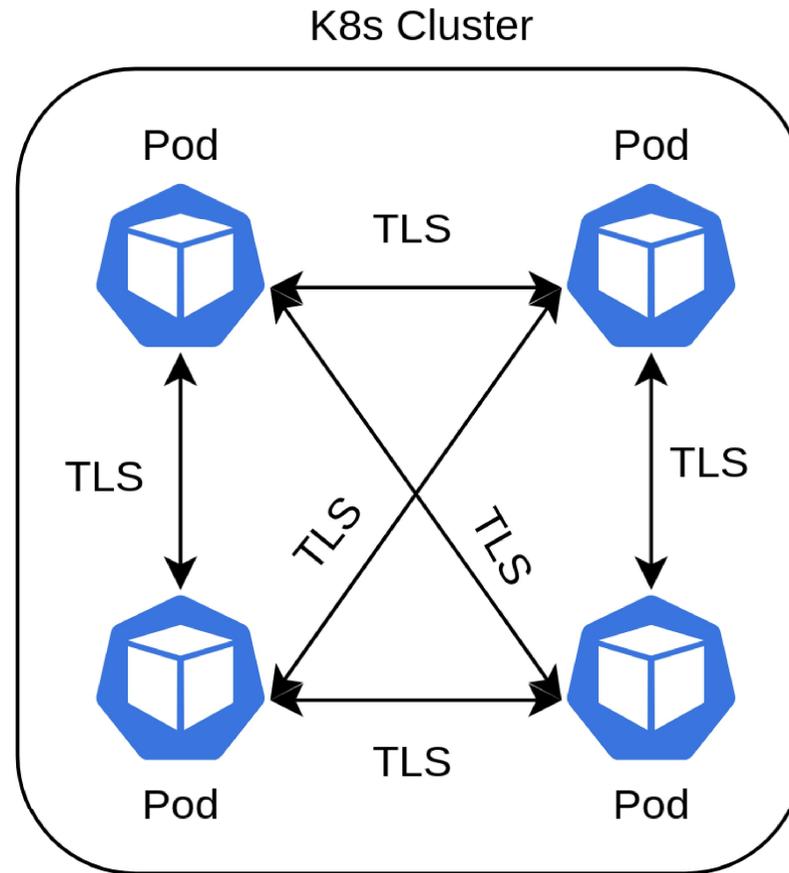
# Microservice- Challenges

- Traffic unverschlüsselt



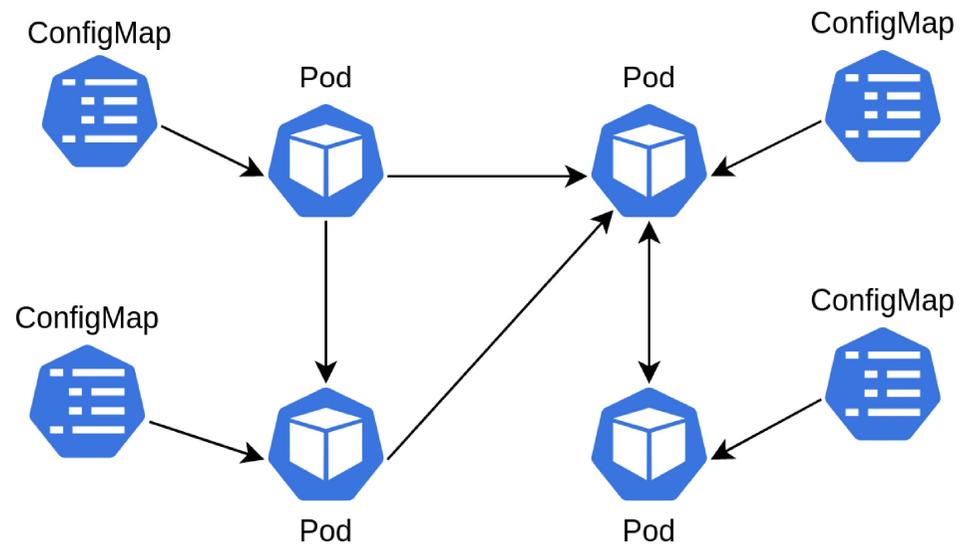
# Microservice- Challenges

- Traffic unverschlüsselt



# Microservice- Challenges

- komplizierte per pod endpoint Definition



The graphic features a teal square background. In the center, there is a blue, stylized geometric shape resembling a three-pointed star or a cluster of triangles. Overlaid on this shape is the text "Microservice-Challenges" in a white, sans-serif font, with "Microservice-" on the top line and "Challenges" on the bottom line.

# Microservice- Challenges

- Service Endpoint Definition komplex
- Retry Logik/Resilienz
- Unsichere Kommunikation
  - in einem Standard Kubernetes Cluster kann jeder Service mit jedem kommunizieren
  - zusätzliches Policing erforderlich (FW/Identy)
  - Anzahl der Services erhöht die Komplexität
- Traffic ist in der Regel unverschlüsselt
  - komplexe Encryption Setups auf per-Service/Pod Basis
- Traffic sollte zwischen workloads vermittelt werden
- Sammeln von Metriken kompliziert
- Traffic Splitting



- Multivendor Application Service Mesh  
Unterstützung für Kubernetes, Cisco CCP, VMWare NSX (und mehr)
- Istio erweitert die Kubernetes API
  - benutzt YAML Files  
deklarativer Approach ohne Kenntnis von spezifischen DSLs
  - integriert in kubectl als CRD  
(custom resource definition)
  - agiert als Helfer
- ControlPlane - Istiod
  - Traffic Management API
  - Certificate Authority
  - Service Discovery
  - speichert Performance Metriken

The logo for Istio Architecture features a teal square background. In the center, there is a stylized blue sailboat with three sails. The word "Istio" is written in white, sans-serif font across the top two sails, and the word "Architecture" is written in white, sans-serif font across the bottom sail.

# Istio Architecture

- Envoy Proxy
  - Loadbalancing
  - SSL/TLS Termination
  - Service Discovery
  - Resilience/health checking
  - traffic splitting(staged Rollouts)
  - MetrikCollector
  - Support für TCP/http/gRPC/Websockets
  - Service Routing
- Istio Ingress/Egress Gateway
  - verwaltet ein-/ausgehenden Traffic zum/vom Service Mesh (edge device)
  - Ingress Loadbalancing
  - CRSF Protection/CORS Support

The logo for Istio Architecture is a teal square containing a stylized blue sailboat. The sailboat has a large, dark blue sail and a smaller, lighter blue sail. The word "Istio" is written in white, sans-serif font on the left side of the sail, and "Architecture" is written in white, sans-serif font across the middle of the sail.

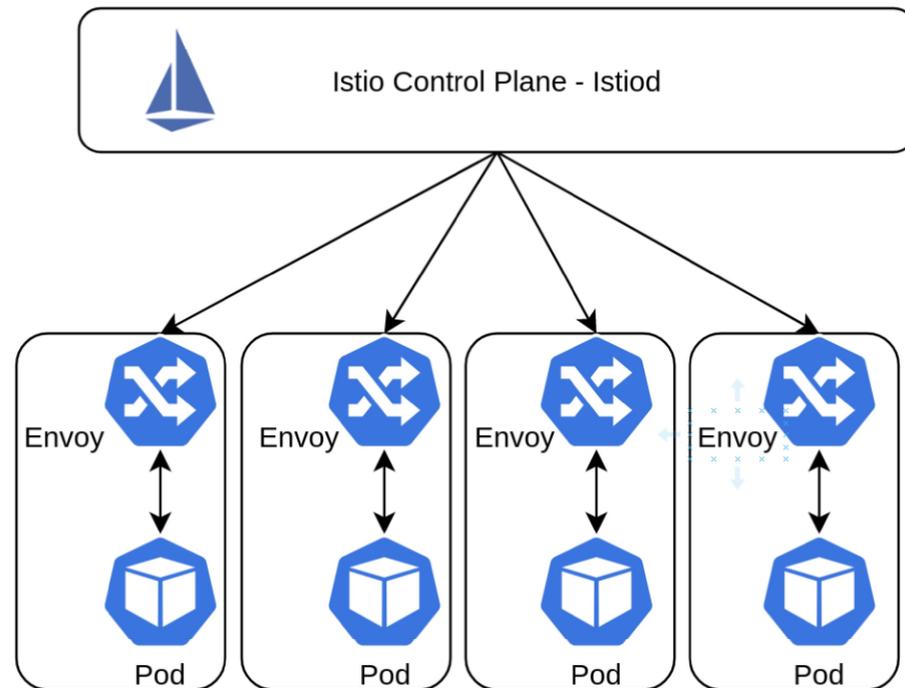
# Istio Architecture

## Wie funktioniert?

- Istiod importiert die Kubernetes Service Registry
- Virtual Service Creation
  - definiert wie Traffic zu einem Service geroutet wird
- Destination Rule Creation
  - definiert was mit diesem Traffic geschieht
- Rules & Services werden zu den Envoy Proxies deployed

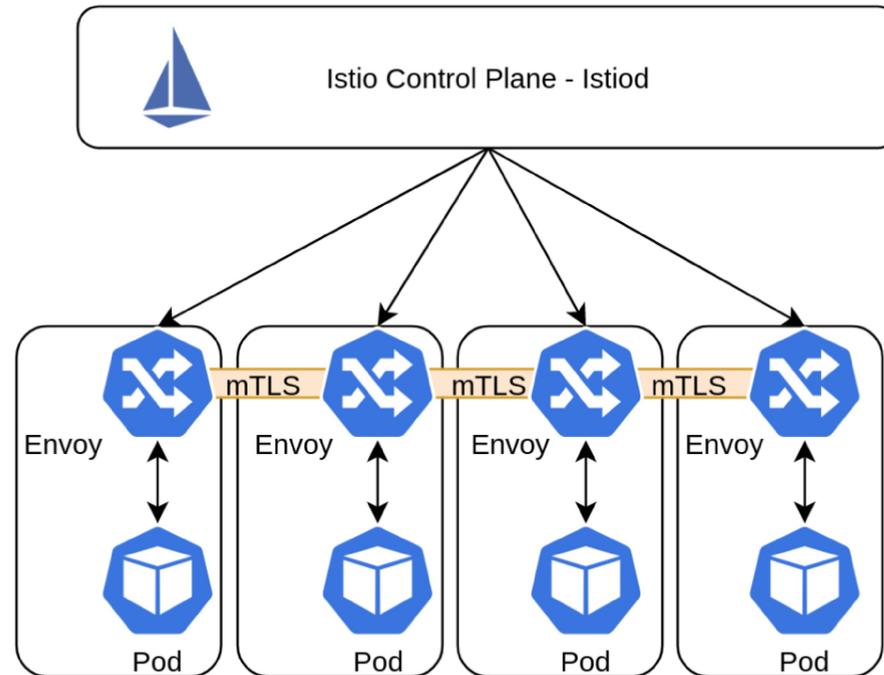
# Istio Architecture

- Istio injiziert einen Envoy Proxy in jeden pod - diese envoy proxies bilden die Dataplane



# Istio Architecture

- traffic zwischen den Services wird verschlüsselt über die Envoy Proxies geroutet



The logo for Istio Architecture features a teal square background with a blue geometric shape resembling a stylized sail or a cluster of triangles. The word "Istio" is written in white above the word "Architecture", both in a sans-serif font.

# Istio Architecture

- Virtual Service Definition

---

```
- apiVersion: networking.istio.io/v1
kind: VirtualService
metadata:
name: vservice
spec:
 hosts:
 - db.cluster.local
 tcp:
 - match:
 - port: 3300
 route:
 - destination:
 host: db.cluster.local
 port:
 number: 3309

- apiVersion: networking.istio.io/v1
kind: VirtualService
metadata:
 name: webservice
spec:
hosts:
 - web.cluster.local
http:
 - name: webservice-route
 match:
 - uri:
 prefix: "/"
 route:
 - destination:
 host: web.cluster.local
```

The logo for Istio Architecture, featuring a teal square background with a blue geometric shape resembling a stylized sail or a cluster of triangles. The word "Istio" is written in white on the left side, and "Architecture" is written in white on the right side, both in a sans-serif font.

# Istio Architecture

- Destination Rule Definition

```
apiVersion: networking.istio.io/v1
kind: DestinationRule
metadata:
 name: webservice-destination-rule
spec:
 host: web.cluster.local
 trafficPolicy:
 loadBalancer:
 simple: LEAST_CONN
```

The logo for Istio Architecture features a teal square background. In the center, there is a blue, three-dimensional geometric shape resembling a cube or a complex polyhedron. The word "Istio" is written in white, sans-serif font across the top-left face of the shape, and the word "Architecture" is written in white, sans-serif font across the bottom face of the shape.

# Istio Architecture

- Ingress Gateway Definition

```
apiVersion: networking.istio.io/v1
kind: DestinationRule
metadata:
 name: webgw
spec:
 selector:
 istio: ingressgateway
 servers:
 - port:
 number: 80
 name: http
 protocol: HTTP
 hosts:
 - "*"

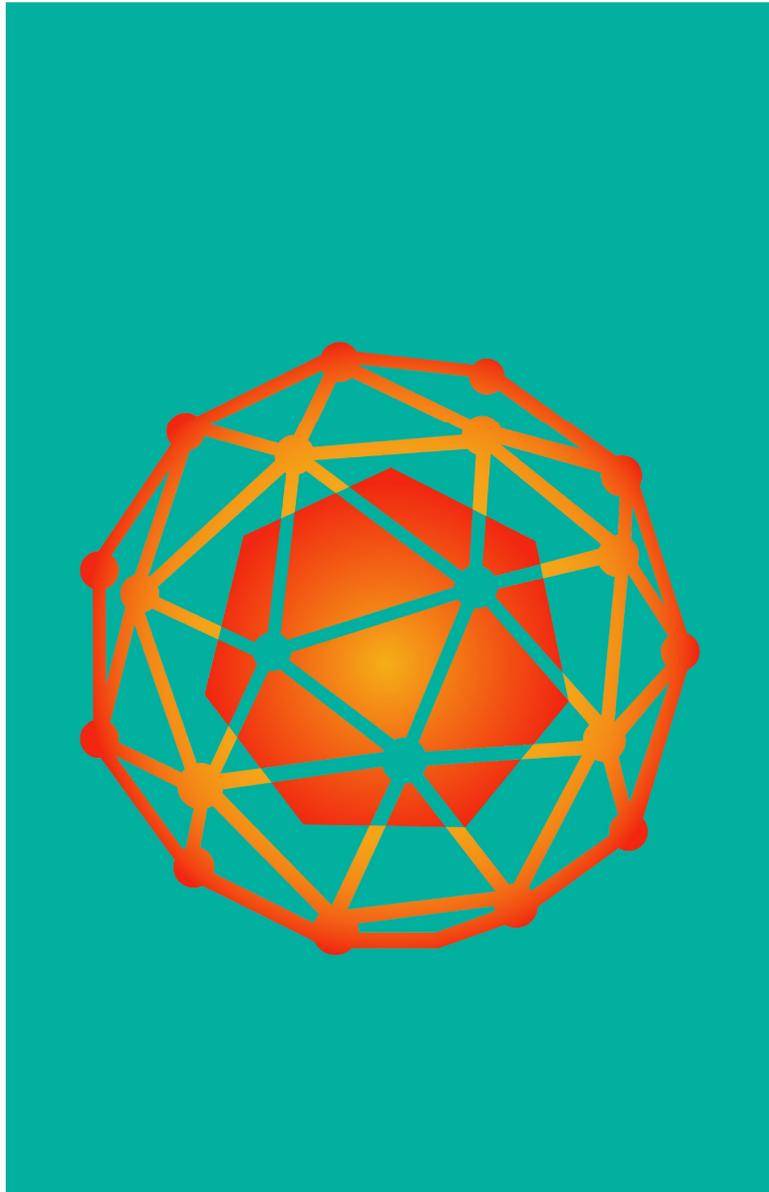
```

The logo for Istio Architecture, featuring a teal background with a blue geometric shape resembling a stylized 'I' or a sail. The word 'Istio' is written in white above the word 'Architecture', both in a sans-serif font.

# Istio Architecture

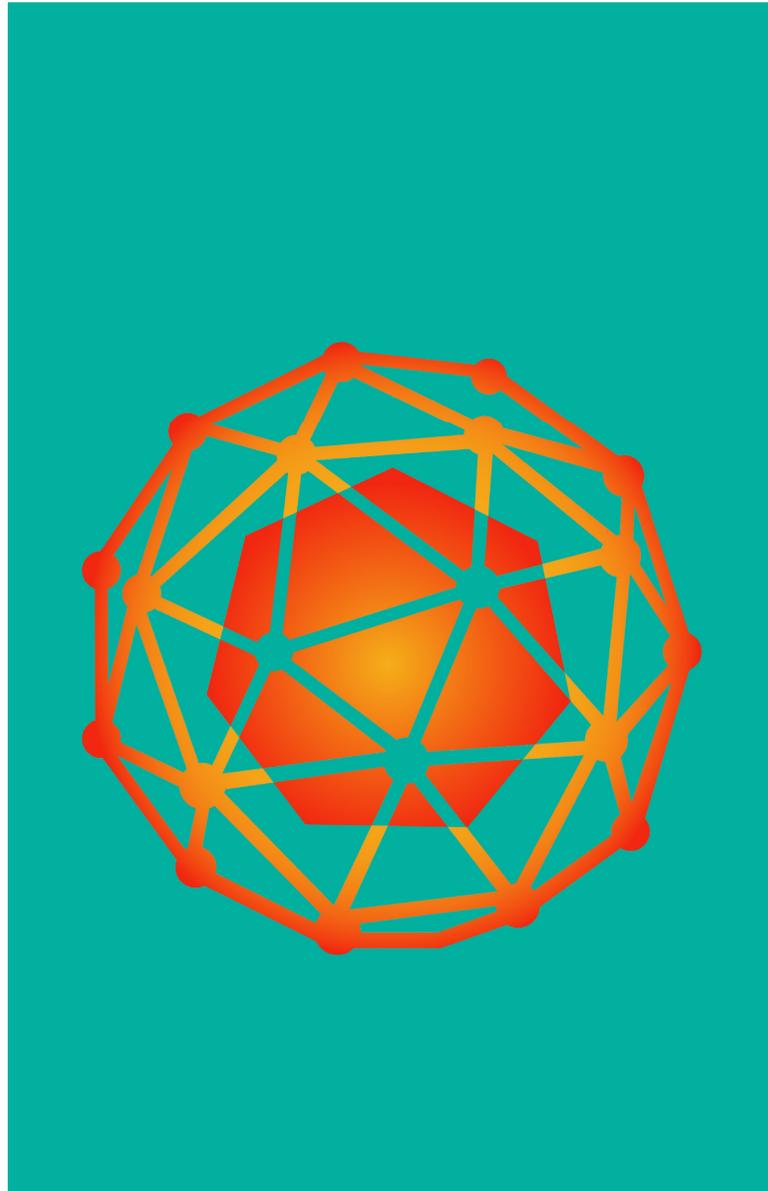
- Ingress Gateway Einbinden

```
- apiVersion: networking.istio.io/v1
kind: VirtualService
metadata:
 name: webservice
spec:
 hosts:
 - "*"
 gateways:
 - webgw
 http:
 - name: webservice-route
 match:
 - uri:
 prefix: "/"
 route:
 - destination:
 host: web.cluster.local
 port:
 number: 80
```



# Network Service Mesh

L2/L3 interdomain connectivity



## Network Service Mesh

- Konnektivität zwischen verschiedenen Workloads bei
  - onPremise Datacentern
  - verschiedenen Cloudanbietern
  - anderen Clustern
- ermöglicht L2/L3 Connectivity
- Cloud Native Approach
  - Entwickler statt Operator fokussiert
  - deklarative an Stelle von imperativer Konfiguration
- Einbindung via Custom Resource Definition (CRD)

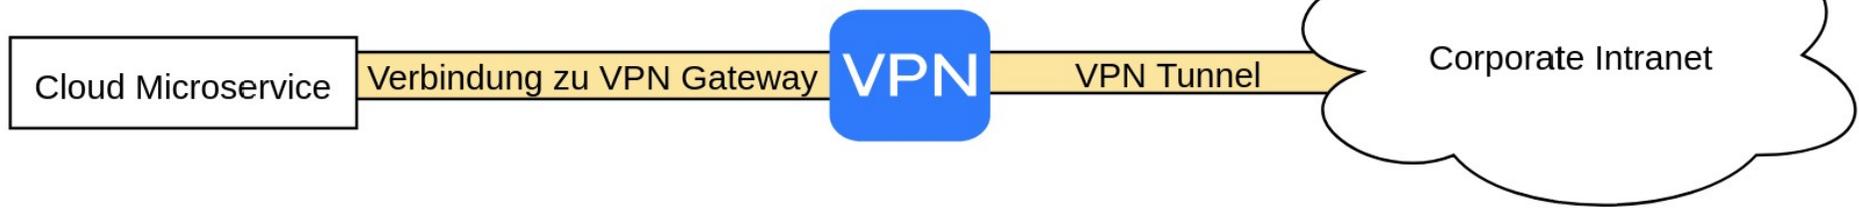
Ein Beispiel...

Sarahs Pod



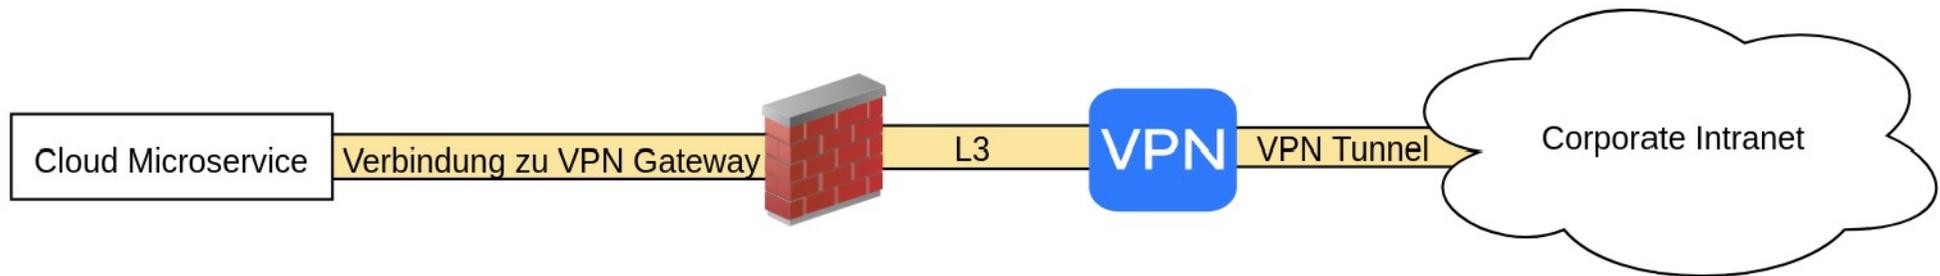
# Definitionshölle für Entwickler

Welches Subnet brauch ich eigentlich?  
Wie groß muss das sein?  
Muss ich auch den VPN Pod konfigurieren?  
Muss ich Routen für Corporate IPs in meinem Pod definieren?  
Woher bekomme ich die?  
Was wenn die sich ändern?



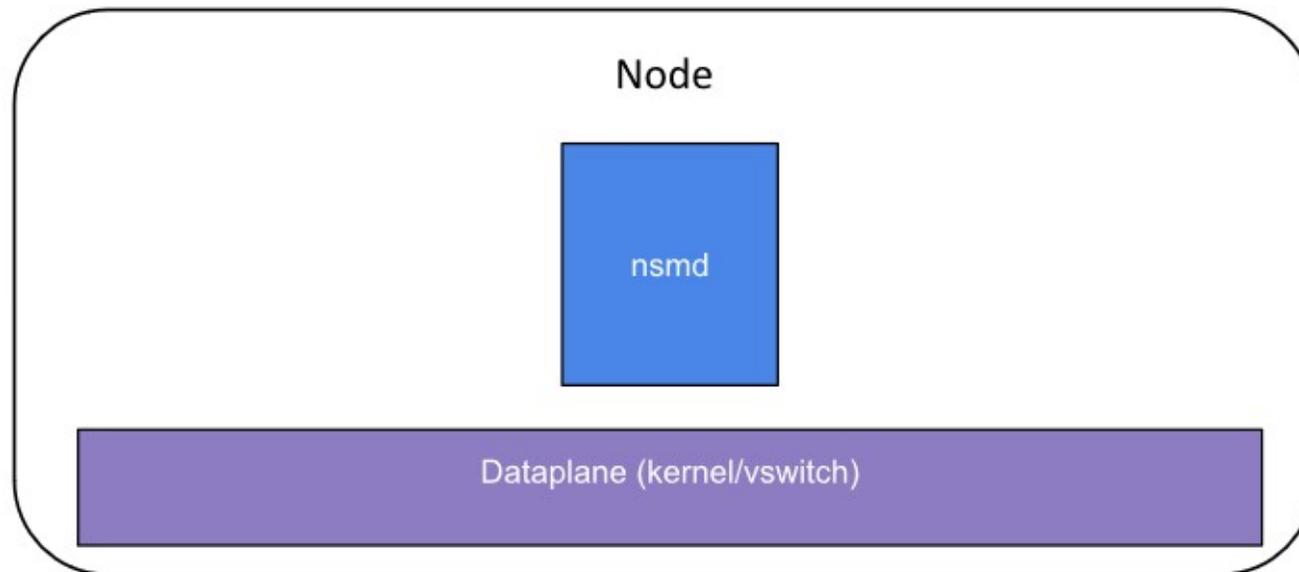
# Definitionshölle für Entwickler die Zweite

Operations sagt ohne Firewall geht garnichts...  
Muss ich die jetzt konfigurieren?  
Was muss ich konfigurieren?  
Warum beschäftige ich mich jetzt mit Netzwerk Setups?  
Was wenn sich was ändert?



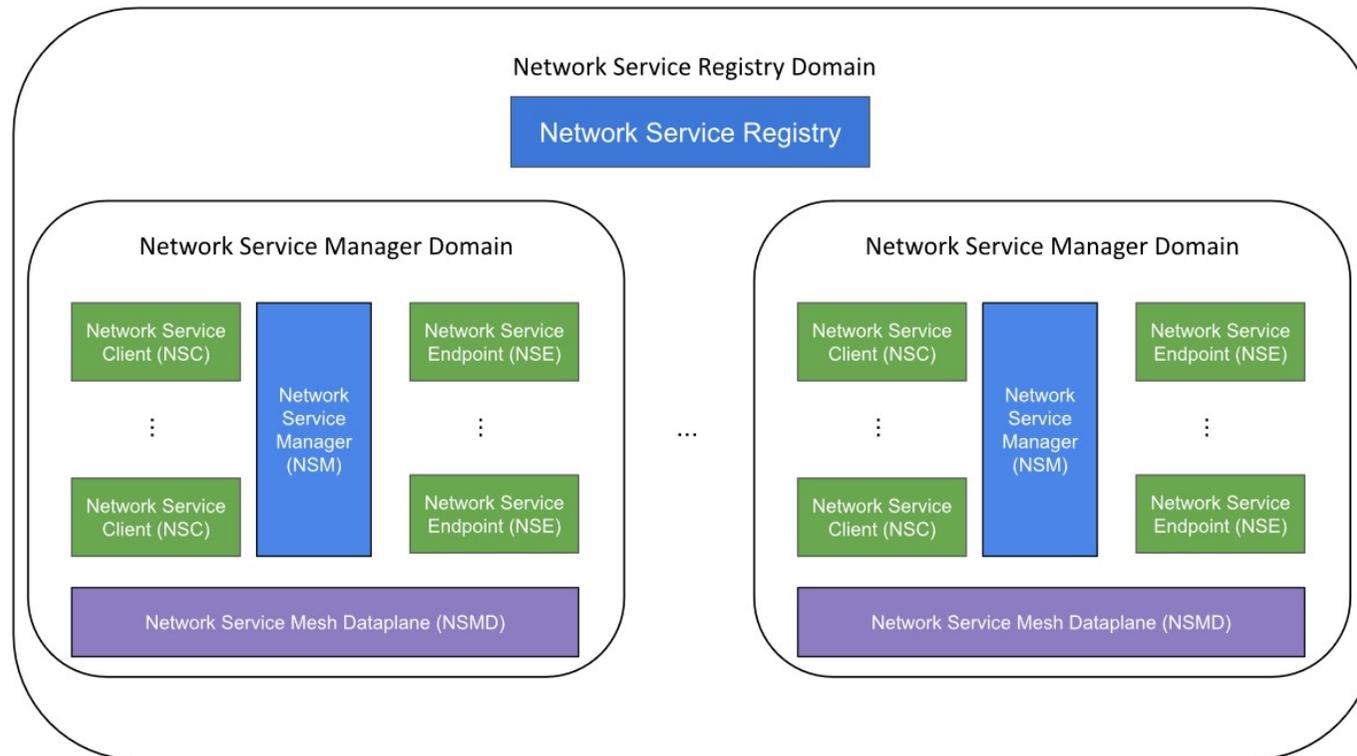
# Architecture

- jeder Worker Node hat zusätzlich einen nsmd service und bildet eine Network Service Manager Domain

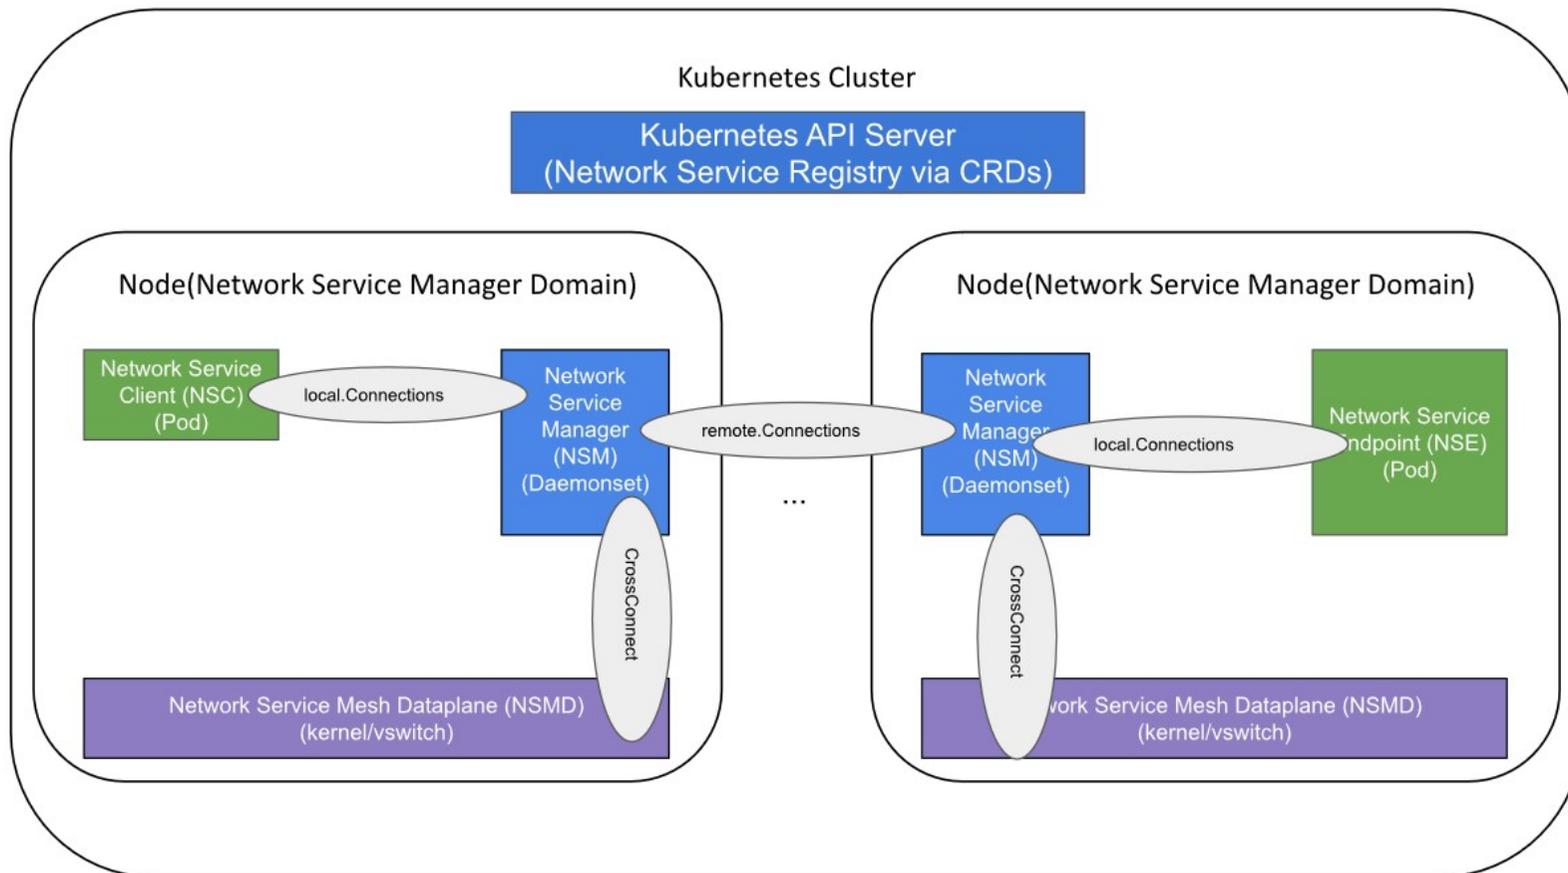


# Architecture

- eine NSM Domain kann Clients und Endpoints enthalten
  - Clients stellen Applikationspods dar (Sarahs Pod zum Bsp.)
  - jedem Client Pod wird ein NSM initContainer hinzugefügt
  - Endpoints sind pods die externe Connectivity herstellen

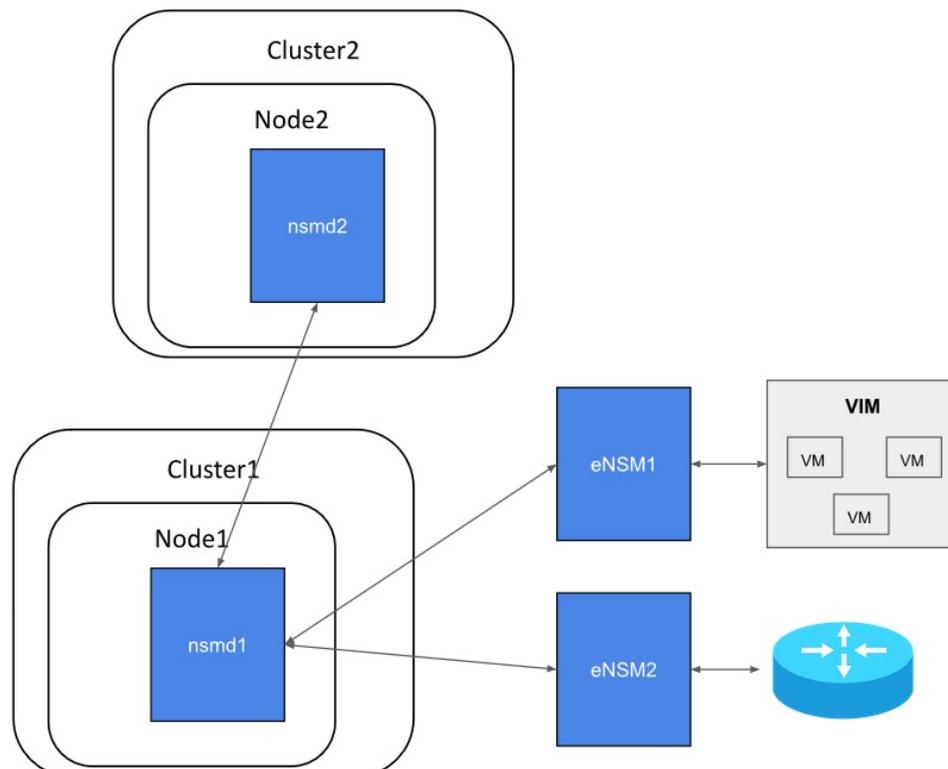


# Architecture



# Architecture

eNSMs bilden die Schnittstelle zu externen Devices



# Architecture - eNSM Connectivity

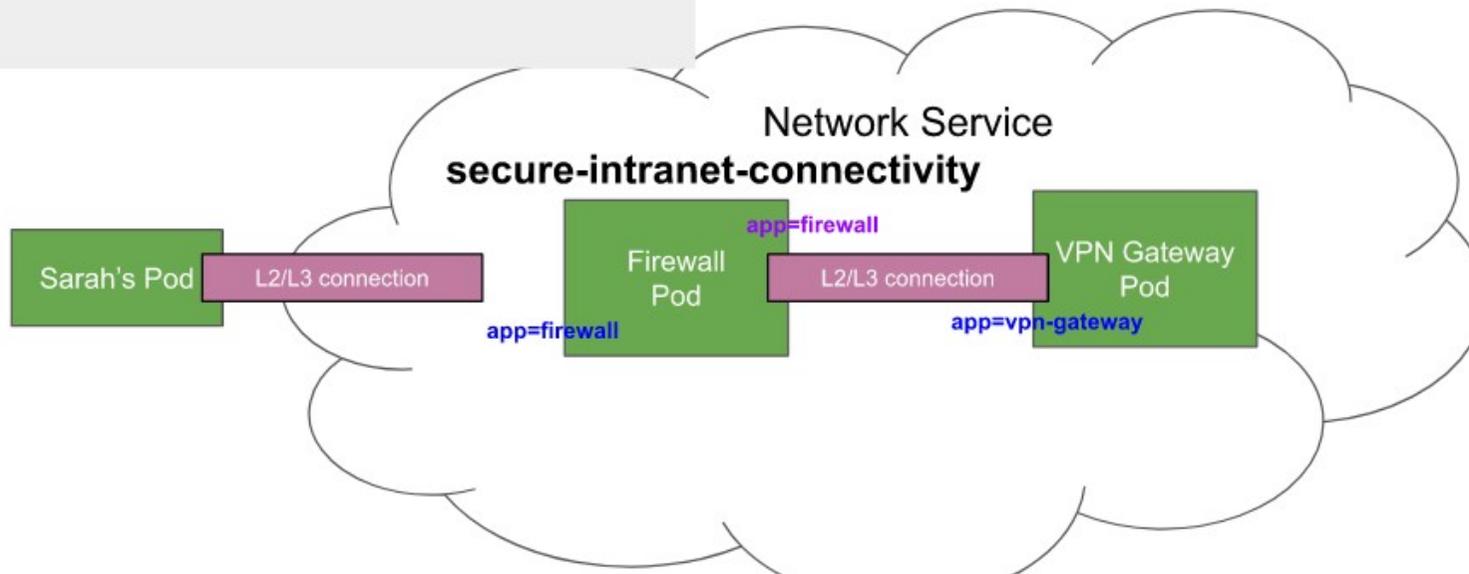
- Externe NSMs werden anhand von DNS Records identifiziert
  - service records (SRV)
- Authentication/Authorization erfolgt via SPIFFE/SPIRE (<https://spiffe.io/>)

```
_dienst._proto.name. TTL Klasse SRV Priorität Gewicht Port Ziel.
_sip._tcp.example.com. 86400 IN SRV 10 0 5060 sip.example.com.
_nsm._tcp.example.com 86400 IN SRV 10 0 8000 nsm.example.com.
```

# Konfiguration - der deklarative Approach

## Source und Destination Selektoren

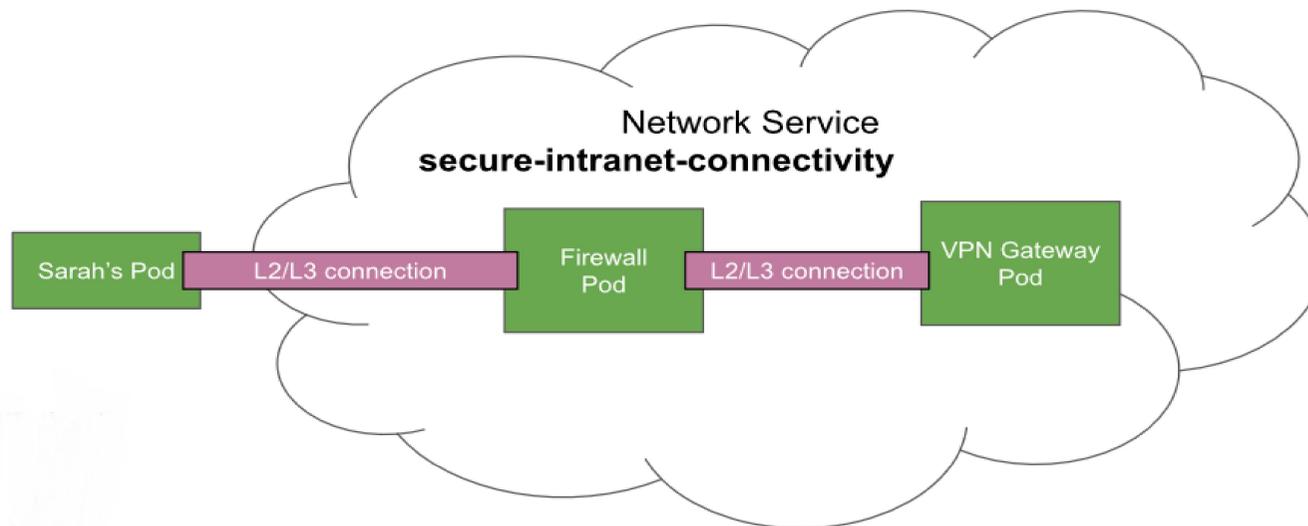
```
kind: NetworkService
apiVersion: V1
metadata:
 name: secure-intranet-connectivity
spec:
 payload: IP
 matches:
 - match:
 sourceSelector:
 app:firewall
 route:
 - destination:
 destinationSelector:
 App:vpn-gateway
 - match:
 route:
 - destination:
 destinationSelector:
 app:firewall
```



# Serviceeinbindung - Annotationlabel

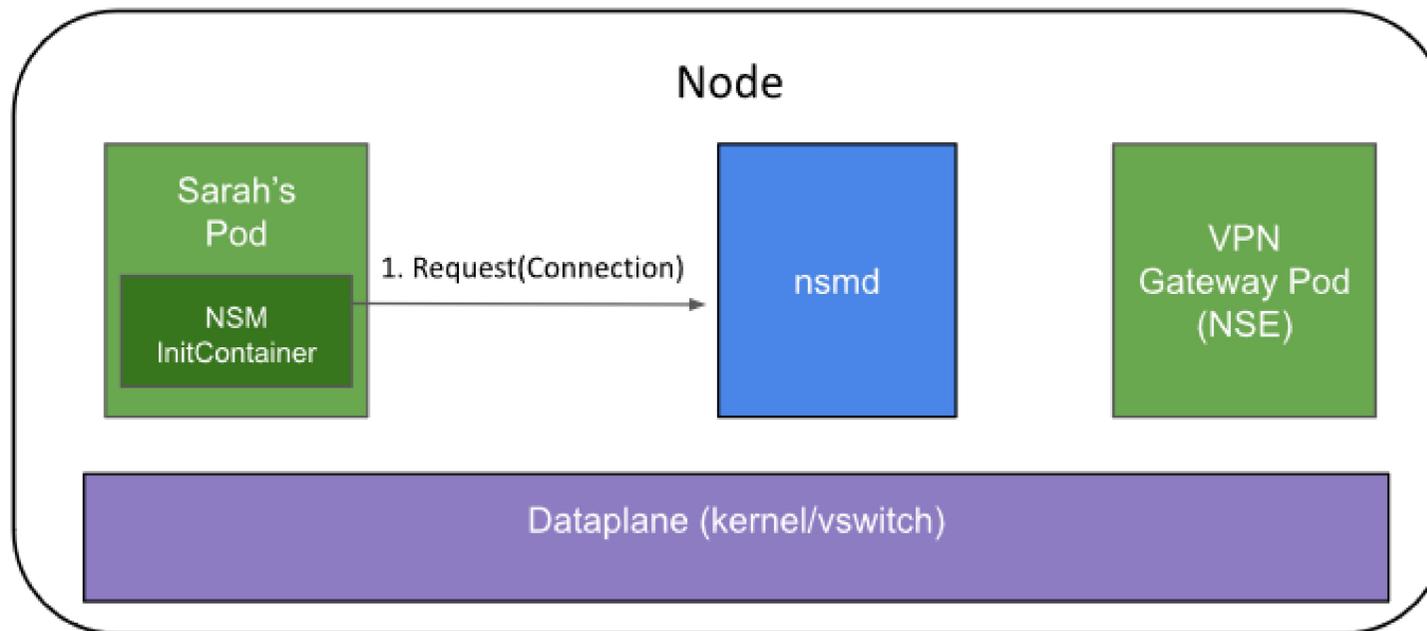
```
apiVersion: v1
kind: Deployment
metadata:
 name: saras-app
 annotations:
 ns.networkservicemesh.io: secure-intranet-connectivity
spec:
 containers:
 - name: myapp
 image: myapp:1.0.0

```



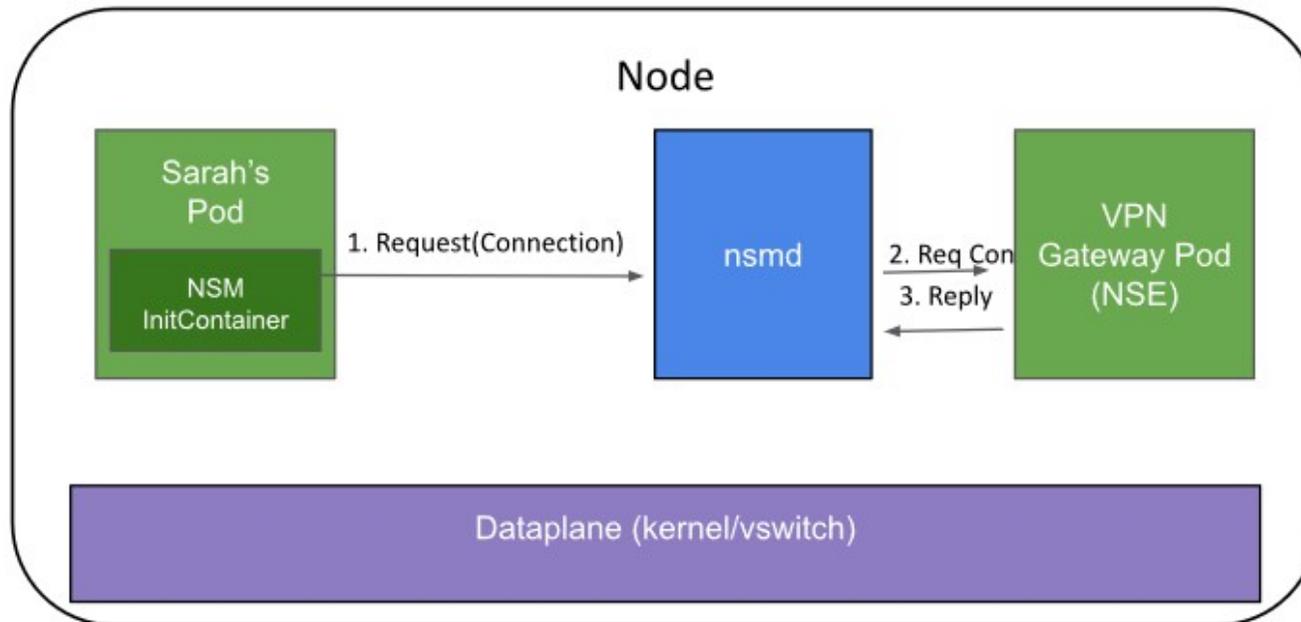
# Verbindungsaufbau

- wenn der Application Pod startet wird via NSM InitContainer ein Request an den nsmd geschickt



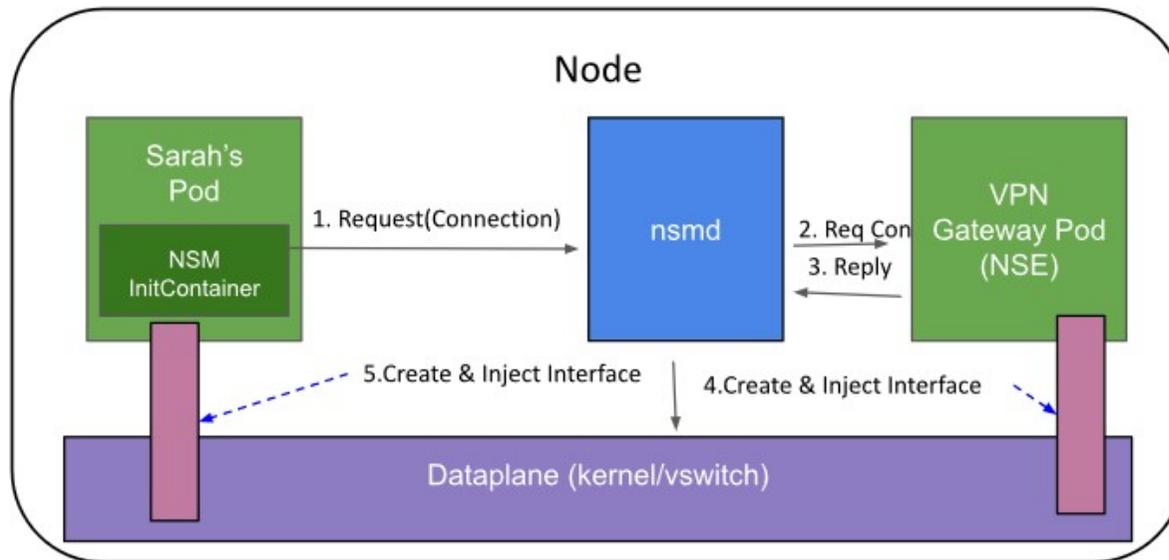
# Verbindungsaufbau

- der nsmd kümmert sich um die Einbindung und Konfiguration des Endpoints



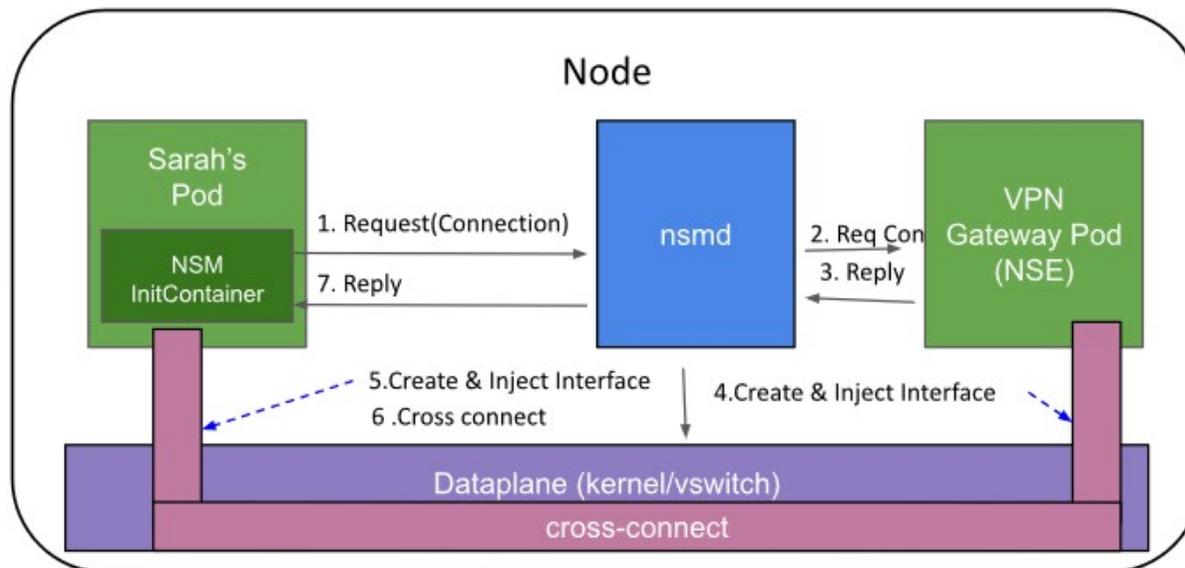
# Verbindungsaufbau

- anschliessend wird die Konnektivität zwischen NSE und NSC hergestellt



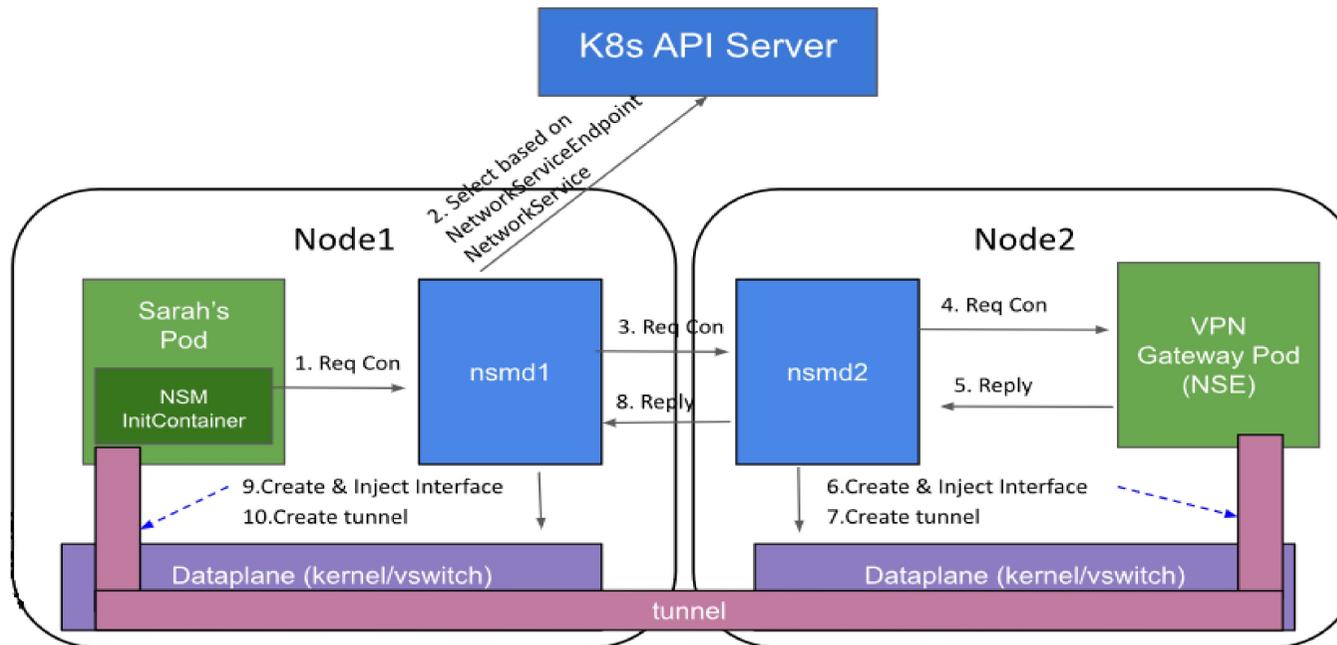
# Verbindungsaufbau

- abschliessend wird der Client über den erfolgreichen Verbindungsaufbau informiert



# Verbindungsaufbau

- NSC und NSE müssen sich nicht auf dem selben Client befinden



Ende

**Frankfurter Linux User  
Group (FraLUG) e.V.**

