

## **Network Service Mesh**



A Narrative Introduction Frederick Kautz and Kyle Mestery







## Meet Sarah

Sarah is writing a Kubernetes app to be deployed in the public cloud

One of the Pods in Sarah's app needs secure access to her corporate intranet





From Sarah's point of view her needs look like this









Sarah still wants her normal Kubernetes Networking...

































































Sort of. You know all the cool things Istio does for you with TCP connections and HTTP?







I do that for IP, Ethernet, and other L2/L3 protocols. Tell me about your problems.







kind: NetworkService
apiVersion: V1
metadata:
 name: secure-intranet-connectivity
spec:
 selector:
 app: secure-intranet-gateway
 payload: IP

Cool, this is how you would do it with NSM.



## That looks a lot like a Service Resource in K8s.



Sarah

```
kind: NetworkService
apiVersion: V1
metadata:
   name: secure-intranet-connectivity
spec:
   selector:
        app: secure-intranet-gateway
        payload: IP
```

Yep! You use selectors on Pods to find the Pods providing the Network Service, and it exposes 'payloads' instead of ports and tcp.









kind: NetworkService
apiVersion: V1
metadata:
 name: secure-intranet-connectivity
spec:
 selector:
 app: secure-intranet-gateway
 payload: IP

Close! You would also want to insert our standard Network Service Mesh init-container and a Config Map telling it what Network Service you want to connect to into your Pod. That's it.



































L2/L3 connections are point to point cross connects between your Pod, and the Network Service you want. You don't have to think about subnets.

If you want a Bridge Domain, that's a Network



Sarah

What about subnets?





Generally, addresses and routes for an L2/L3 connection come from the **Network Service** And routes? My network guys added an intranet CIDR last **Endpoint**, like your VPN Gateway. They are in month and I had to update the routes on \*all\* of my Pods. a better position to know what they should be. It sucked! Advanced use cases can be done with more flexibility, but that's probably not what you want here. **Network Service** Sarah Example: secure-intranetconnectivity **Network Service Endpoint** Sarah's Pod L2/L3 connection Example: VPN Gateway Pod








```
kind: NetworkServiceWiring
apiVersion: V1
```

metadata:

```
name: secure-intranet-connectivity-wiring-1
spec:
```

target:

- secure-intranet-connectivity
qualifiers:
 source:
 sourceService:
 !secure-intranet-connectivity
action:
 route:
 -destination:

```
podSelector:
firewall=true
```

The 'secure-intranet-connectivity-wiring-1' **Network Service Wiring** says:

'target: secure-intranet-connectivity'

That means it applies to L2/L3 connections trying to reach the secure-intranet-connectivity Network Service.







```
kind: NetworkServiceWiring
apiVersion: V1
```

metadata:

```
name: secure-intranet-connectivity-wiring-1
```

spec:

target:

- secure-intranet-connectivity

qualifiers:

source:

```
sourceService:
```

!secure-intranet-connectivity

```
action:
```

route:

- -destination:
  - podSelector:
    firewall=true

-des

The 'secure-intranet-connectivity-wiring-1' **Network Service Wiring** then has a qualifier that only matches things **not** providing the secure-intranet-connectivity service themselves.

Things like your Pod.







```
kind: NetworkServiceWiring
apiVersion: V1
```

metadata:

```
name: secure-intranet-connectivity-wiring-1
```

spec:

target:

- secure-intranet-connectivity
qualifiers:

```
source:
```

```
sourceService:
```

```
!secure-intranet-connectivity
```

```
action:
```

```
route:
```

- -destination:
   podSelector:
  - firewall=true



And it has a 'route' that sends those connections to one of the Network Service Endpoints with label 'firewall=true'... like the Firewall Pod.









OK. How does the Firewall Pod get connected to the VPN Gateway Pod?

Sarah

kind: NetworkServiceWiring
apiVersion: V1
metadata:

name: secure-intranet-connectivity-wiring-2
spec:

target:

- secure-intranet-connectivity
qualifiers:
 source:
 podSelector:

firewall=true

#### action:

route:

-destination:

podSelector: vpngateway=true



You would have a second Network Service Wiring 'secure-intranet-connectivity-wiring-2'





```
kind: NetworkServiceWiring
apiVersion: V1
```

metadata:

```
name: secure-intranet-connectivity-wiring-2
```

spec:

target:

- secure-intranet-connectivity
qualifiers:

source:

podSelector:

firewall=true

#### action:

route:

- -destination:
  - podSelector: vpngateway=true



That has a qualifier that matches sources that have label firewall=true





```
kind: NetworkServiceWiring
apiVersion: V1
```

metadata:

```
name: secure-intranet-connectivity-wiring-2
```

spec:

target:

- secure-intranet-connectivity
qualifiers:
 source:
 podSelector:

firewall=true

#### action:

route:

-destination: podSelector: vpngateway=true

# And routes them to Pods with vpngateway=true























The NSM InitContainer reads your Config Map, and sends a GRPC call across a unix file socket to the NSM to Request an L2/L3 Connection to the secure-intranet-connectivity Network Service.







Sarah's Pod

NSM InitContainer Request Connection has any information needed to be clear about how you want the connection to look to your Pod locally, like that you want it to be a kernel interface, the interface name you'd prefer if you care, etc.







Let's talk first about the case where your VPN Gateway Pod is on the same Node.







The NSM sends a Request Connection to the VPN Gateway Pod, which sends back an Accept.







The NSM creates an interface for the connection for the VPN Gateway and injects it into the VPN Gateway Pod.







The NSM creates and injects an interface into Sarah's Pod







The NSM then cross connects the two interfaces in the dataplane.







Finally the NSM responds to your Pod's NSMInitContainer with an Accept.

Then you are ready to go.











# What if the VPN Gateway Pod is on a different Node.

Sarah









If the VPN Gateway Pod is on a different Node, it looks exactly the same to your Pod.

You send a Request Connection request to NSM using GRPC over a unix file socket.


























When the VPN Gateway Pod comes up, it Exposes to the NSM1 that it has a secureintranet-connectivity Network Service willing to accept connections.

K8s API Server

Sarah







Cool! This looks so much easier than trying to manually string together interfaces and subnets myself! Thank you!

You are welcome! Have fun!

Sarah





Thank you!

