

# Declarative Networking

in Declarative World

in 2025

Mateusz Kowalski

Principal Software Engineer

# \$whoami

- ▶ Based in Switzerland
- ▶ Background in academia, banking and telco
- ▶ Buzzwords
  - Cloud
  - Metal
  - Network security
  - ~~Artificial intelligence~~



# Multiple network interfaces

Why?

- Network equipment, routers, switches
- Telco
- Separate storage networks
- ...



# Network configuration

At "boot"-time

```
root@fedora-nuc:/etc/NetworkManager/system-connections# cat ens192.nmconnection
[connection]
id=ens192
uuid=a0aeaa9a-03cb-3a23-ac47-a608e1555115
type=ethernet
autoconnect-priority=-999
interface-name=ens192
permissions=
timestamp=1628204789

[ethernet]
mac-address-blacklist=

[ipv4]
dns-search=
method=auto

[ipv6]
addr-gen-mode=eui64
dns-search=
method=auto

[proxy]
```

- ▶ NetworkManager uses using static nmconnection files
- ▶ Changes not applied automatically
  - You may break your config...
  - ... and not notice for ages

# Network configuration

At runtime

```
[root@worker-6 system-connections]# cat default_connection.nmconnection | grep address1
address1=192.99.99.123/24,192.99.99.1
[root@worker-6 system-connections]# nmcli con mod 1969ebe8-b753-40fb-9e29-a73b7d075737 \
> ipv4.addresses 10.10.99.99/24 \
> ipv4.gateway 192.168.111.111
Error: failed to modify ipv4.addresses: invalid IP address: Invalid IPv4 address '10.10.99.9924'.
[root@worker-6 system-connections]# nmcli con mod 1969ebe8-b753-40fb-9e29-a73b7d075737 \
> ipv4.addresses 10.10.99.99/24 \
> ipv4.gateway 192.168.111.111
Warning: There is another connection with the name 'Wired Connection'. Reference the connection by its
uid '1969ebe8-b753-40fb-9e29-a73b7d075737'
[root@worker-6 system-connections]# cat default_connection.nmconnection | grep address1
address1=10.10.99.99/24,192.168.111.111
[root@worker-6 system-connections]#
```

- ▶ Nmcli to modify the configuration
- ▶ Imperative way
- ▶ Basic protection against stupid things

# Network configuration

At runtime, but declarative

```
$ sudo nmstatectl show
---
dns:
  config:
    server:
      - 192.0.2.1
    search:
      - example.org
routes:
  config:
    - destination: 0.0.0.0/0
      next-hop-interface: eth1
      next-hop-address: 192.0.2.1
interfaces:
  - name: eth1
    type: ethernet
    description: Main-NIC
    state: up
    ipv4:
      enabled: true
      dhcp: false
      address:
        - ip: 192.0.2.9
          prefix-length: 24
    ipv6:
      enabled: false
```

- ▶ Nmstatectl to apply the configuration
- ▶ Based on yaml
- ▶ Declarative

# Network configuration

Declarative, k8s-managed

```
apiVersion: nmstate.io/v1beta1
kind: NodeNetworkConfigurationPolicy
metadata:
  name: bond0-eth1-eth2-policy
spec:
  nodeSelector:
    kubernetes.io/hostname: <node01>
  desiredState:
    interfaces:
      - name: bond0
        description: Bond enslaving eth1 and eth2
        type: bond
        state: up
        ipv4:
          dhcp: true
          enabled: true
        link-aggregation:
          mode: active-backup
          options:
            mimon: '140'
          slaves:
            - eth1
            - eth2
        mtu: 1450
```

- ▶ Kubernetes to apply the configuration
- ▶ Based on CRD
- ▶ Declarative

# Demo

...



# NMstate

Today

- ▶ Written in Rust
- ▶ NetworkManager as a backend
- ▶ Kubernetes Operator live and proven in action
- ▶ Usable from Rust, Golang, Python

# What's new

2024

```
sh-5.1# nmstatectl statistic /tmp/state2 2>/dev/null
topology:
- auto_ip4,auto_ip6 -> ethernet
- static_ip4,auto_ip6 -> ethernet
- static_ip4,static_ip6 -> loopback
- auto_ip4,auto_ip6 -> ovs-bridge -> ethernet
- auto_ip4,auto_ip6 -> ovs-interface -> ovs-bridge -> ethernet
features:
- lldp
- ovn-mapping
- ovs-db-global
- ovs-db-interface
- static-dns-name-server
- static-dns-search
- static-hostname
```

- ▶ In-cluster usage metrics
- ▶ Performance improvements
  - E.g. applying 100s of manifests for cluster with 100s of nodes

```
sh-5.1# curl -sS http://10.129.0.149:8089/metrics | grep features
# HELP kubernetes_nmstate_features_applied Number of nmstate features applied labeled by its name
# TYPE kubernetes_nmstate_features_applied gauge
kubernetes_nmstate_features_applied{name="static-dns-name-server"} 1
kubernetes_nmstate_features_applied{name="static-dns-search"} 1
```

# What's coming

2025

- ▶ Rollback/revert when deleting NNCP
- ▶ Custom operator configuration (e.g. for DNS health checks)
  
- ▶ (maybe) Syntax improvements for telco use cases
- ▶ (maybe) Reducing K8s API calls even more
  
- ▶ Publishing as Helm chart (apart from OLM today)
- ▶ Joining [K8s Network Plumbing Working Group](#)

# The End

@mko - kubernetes.slack.com

github.com/mkowalski

linkedin.com/in/mateuszkowalski



[linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)



[youtube.com/user/RedHatVideos](https://youtube.com/user/RedHatVideos)



[facebook.com/redhatinc](https://facebook.com/redhatinc)



[twitter.com/RedHat](https://twitter.com/RedHat)