openstack.

May 22, 2018

OpenStack Kuryr

Network Policy Support, OpenStack Summit Vancouver

Pino de Candia - @pinodeca

Daniel Mellado Area - dmellado

What is Kuryr?

- Brings OpenStack networking and storage to containers
- Kubernetes Neutron Networking
- Native OpenStack infrastructure for mixed workloads

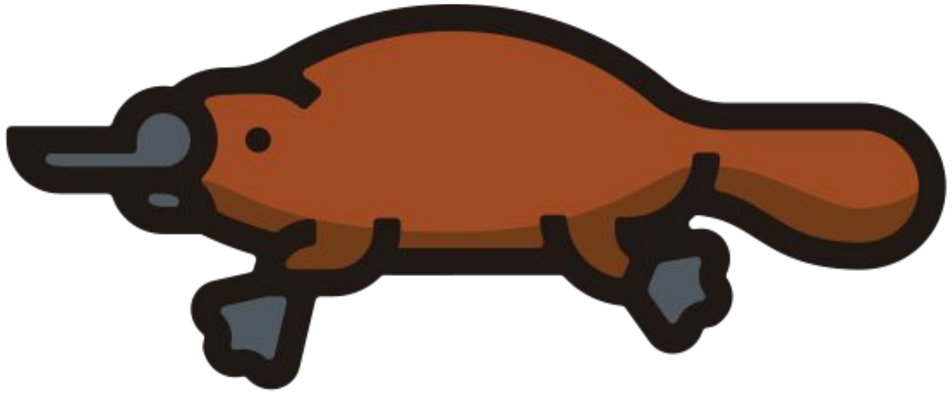

KURYR

an OpenStack Community Project

OpenStack Rocky: Kuryr-Kubernetes

Network Policy API

- Specify which connectivity to allow
- Ingress and egress policy types

Kuryr Network Policy support

- Mapped to Neutron Security Groups
- Extend Kuryr Controller
  - Network Policy Handler
  - Enforce policies on label match
  - Add support to the port pool drivers

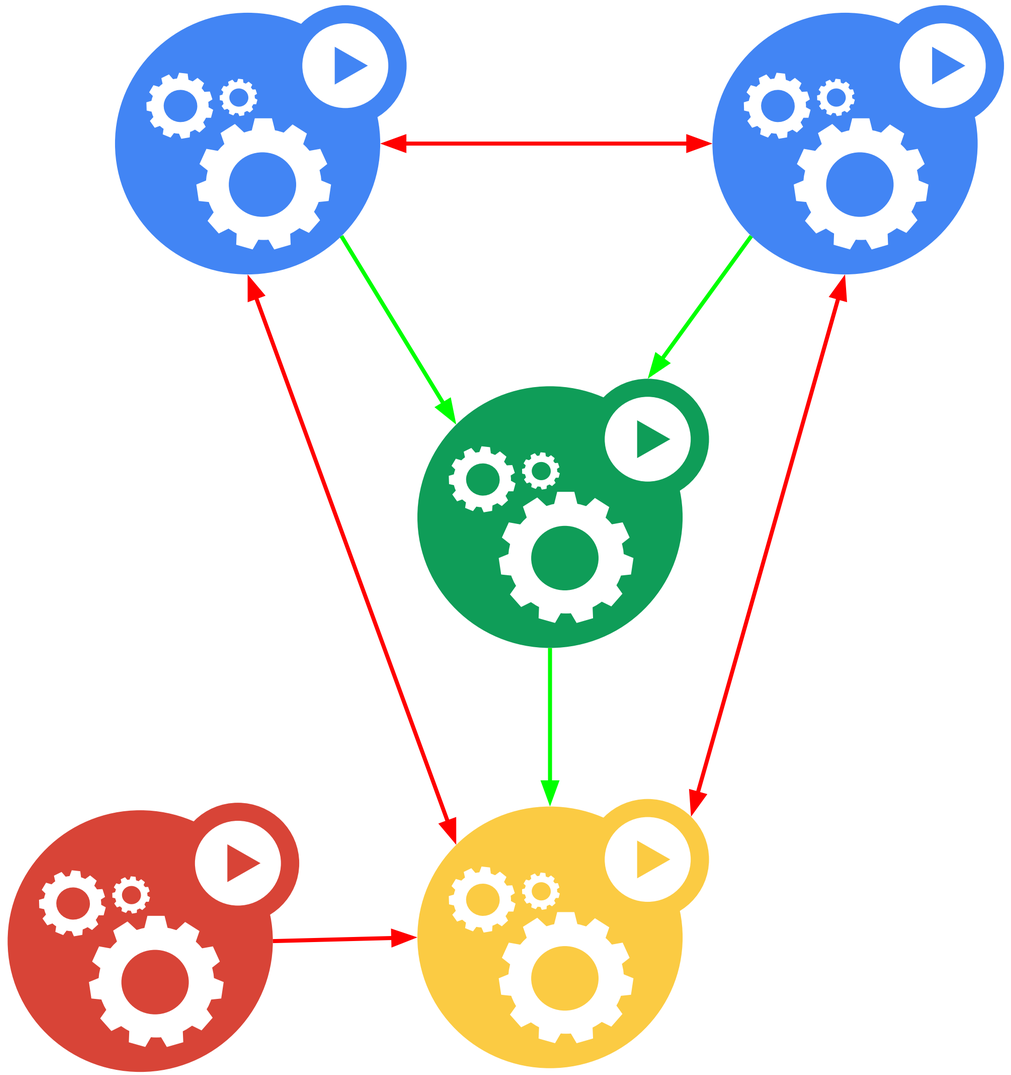

Kubernetes Network Policy (stable since 1.7)

```
spec:
  podSelector:
    matchLabels:
      role: db
  policyTypes:
  - Ingress
  - Egress
  ingress:
  - from:
    - ipBlock:
        cidr: 172.17.0.0/16
        except:
        - 172.17.1.0/24
    - namespaceSelector:
        matchLabels:
          project: myproject
    - podSelector:
        matchLabels:
          role: frontend
    ports:
    - protocol: TCP
      port: 6379
```

For the purposes of our talk:

- call the top one the **protected pod selector**
- call the bottom one the **remote pod selector**

Review: OpenStack Security Groups

- A membership set of Neutron network ports (vNICs)
- A set of rules each describing **allowed** traffic
  - ingress/egress
  - IPv4/IPv6 prefix **OR** Security Group ID
  - destination protocol and port range (or icmp type/code)
- Many-to-many relationship between network ports and SGs

```
openstack security group rule create SG_NAME --protocol PROTO \
     --dst-port FROM:TO [--remote-ip CIDR | --remote-group] [--egress]
```

Compared to K8s Network Policy:

- Empty SGs are similar to labels
- Similar use of ingress/egress, protocol, ports, IP prefixes
- **“except”** in K8s **ipBlock** requires expansion to multiple SG rules

Kuryr policy translation to Neutron

```
spec:
  podSelector:
    matchLabels:
      role: db
  policyTypes:
  - Ingress
  ingress:
  - from:
    - ipBlock:
        cidr: 1.1.1.0/24
        except:
        - 1.1.1.0/26
    - namespaceSelector:
        matchLabels:
          project: myproject
    - podSelector:
        matchLabels:
          role: frontend
    ports:
    - protocol: TCP
      port: 6379
```

Neutron Translate:

- Create SG that applied on pods with “role:db”
- No Egress policy: allow all egress;
  Ingress according to spec
- Translated to set of remote ip prefix : (1.1.1.128/25 , 1.1.1.64/26)
- Create SG and use as remote_group_id
- Create SG and use as remote_group_id
- Each rule above must match this protocol and port

K8s watches:

- Watch all pods “role: db” Watch callback - annotate pods with sg-id
- Watch on namespace that matches to query “project: my project”
- Watches on pods that matches to query “role: frontend”

Execution flow

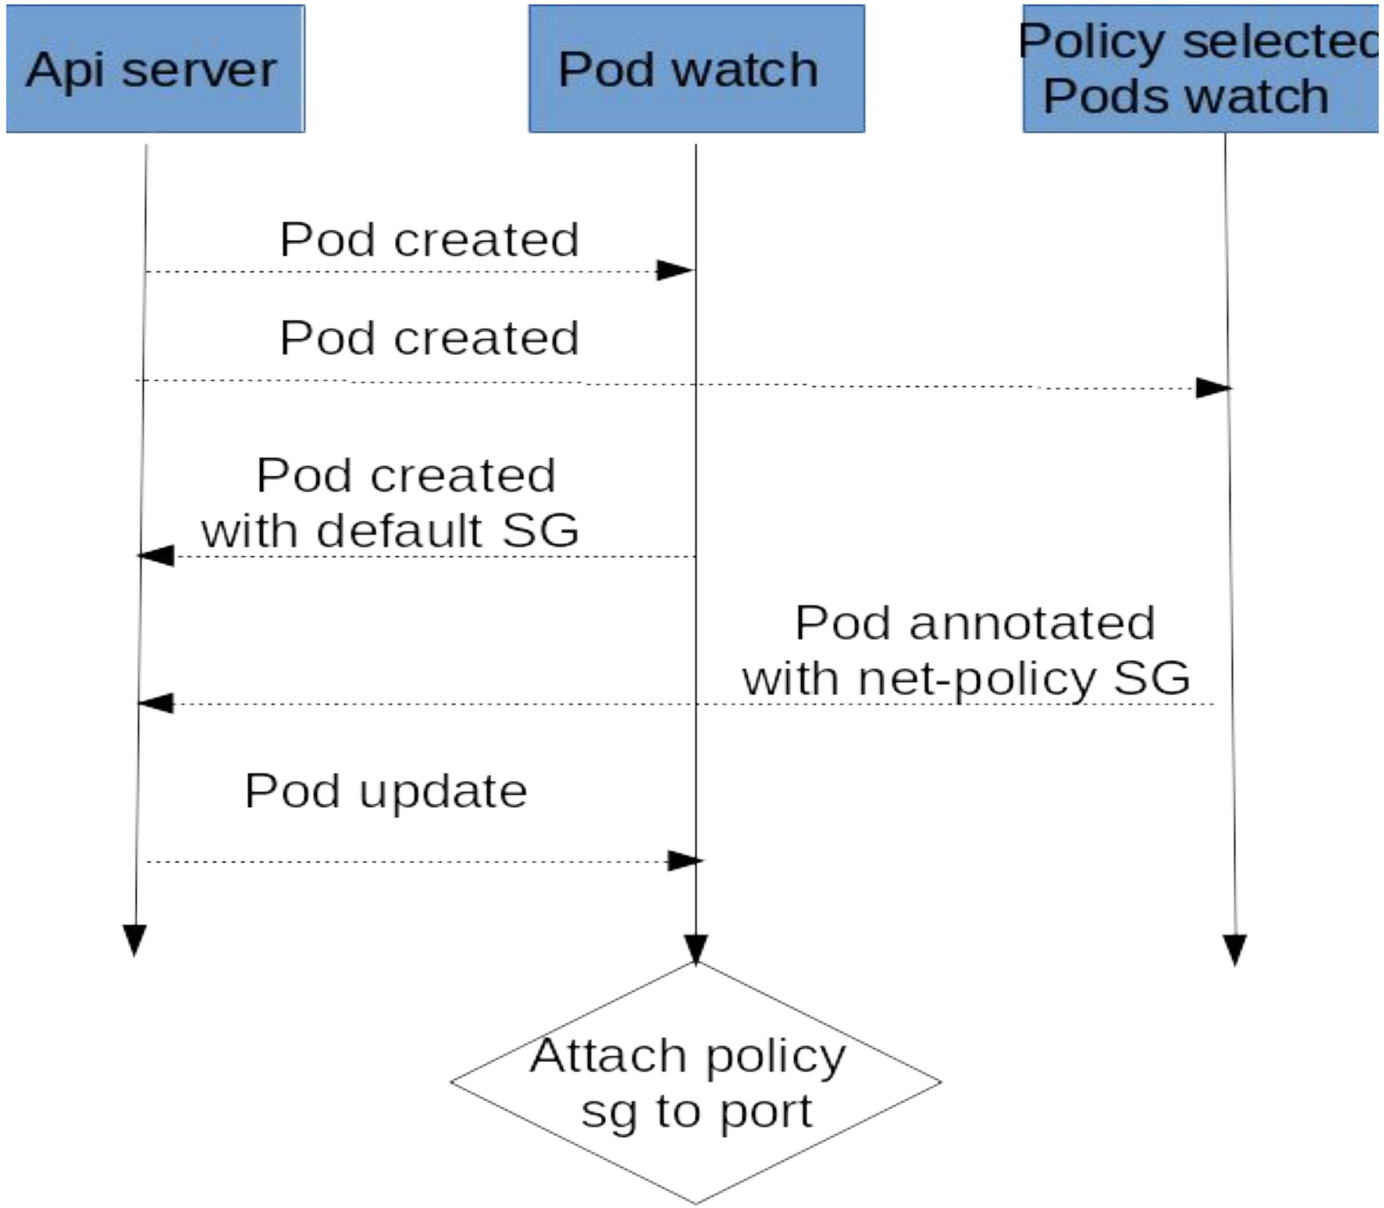

openstack.

Q&A

Thank you!

 openstack  @OpenStack  openstack  OpenStackFoundation