



# Kubernetes Multi-Tenancy & User-Management

By Fabian Kramm

 @FabianKramm

 DevSpace

# Hi! I am Fabian.



- CTO @ DevSpace Technologies Inc.
- We help companies to provision secure Kubernetes sandbox environments
- Open-Source Contributor
- Prev. worked at SAP & sovanta AG
- Trivia: Born in Mannheim



[@FabianKramm](https://twitter.com/FabianKramm)

# What is Multi-Tenancy?



Source: platform9.com

4

 @FabianKramm

 DevSpace

# Why Multi-Tenancy?

# Separate Clusters vs. Cluster Sharing



- **Higher Compute Cost**

- Redundant cluster-wide resources  
(e.g. control planes, operators, storage and networking services)
- Less efficient cluster auto-scaling

- **Higher Maintenance Cost**

- Number of clusters explodes with adoption spreading across teams and applications
- Hard to keep everything in sync

# Separate Clusters vs. Cluster Sharing



- **Lower Compute Cost**
  - Shared cluster-wide resources
  - Less efficient cluster auto-scaling
  
- **Lower Maintenance Cost**
  - Reduce number of clusters (sharing if it makes sense)

# Hard Tenancy vs Soft Tenancy

- **Hard Tenancy:**
  - No trust between tenants
  - Multiple users from multiple different places are sharing the cluster  
=> Lots of configuration necessary
  
- **Soft Tenancy:**
  - Tenants usually part of the same organization
  - *Small* chance of malicious behaviour
  - “prevent accidents” rather than assuming the worst  
=> Rather easy to achieve

# Overview

- **Basics**

- Access Control
- Namespaces

- **Authentication**

- OpenID Connect
- Service Accounts
- Rancher

- **Authorization**

- **Admission Control**

- Resource Quotas
- Limit Ranges
- PodSecurityPolicies
- Open Policy Agent

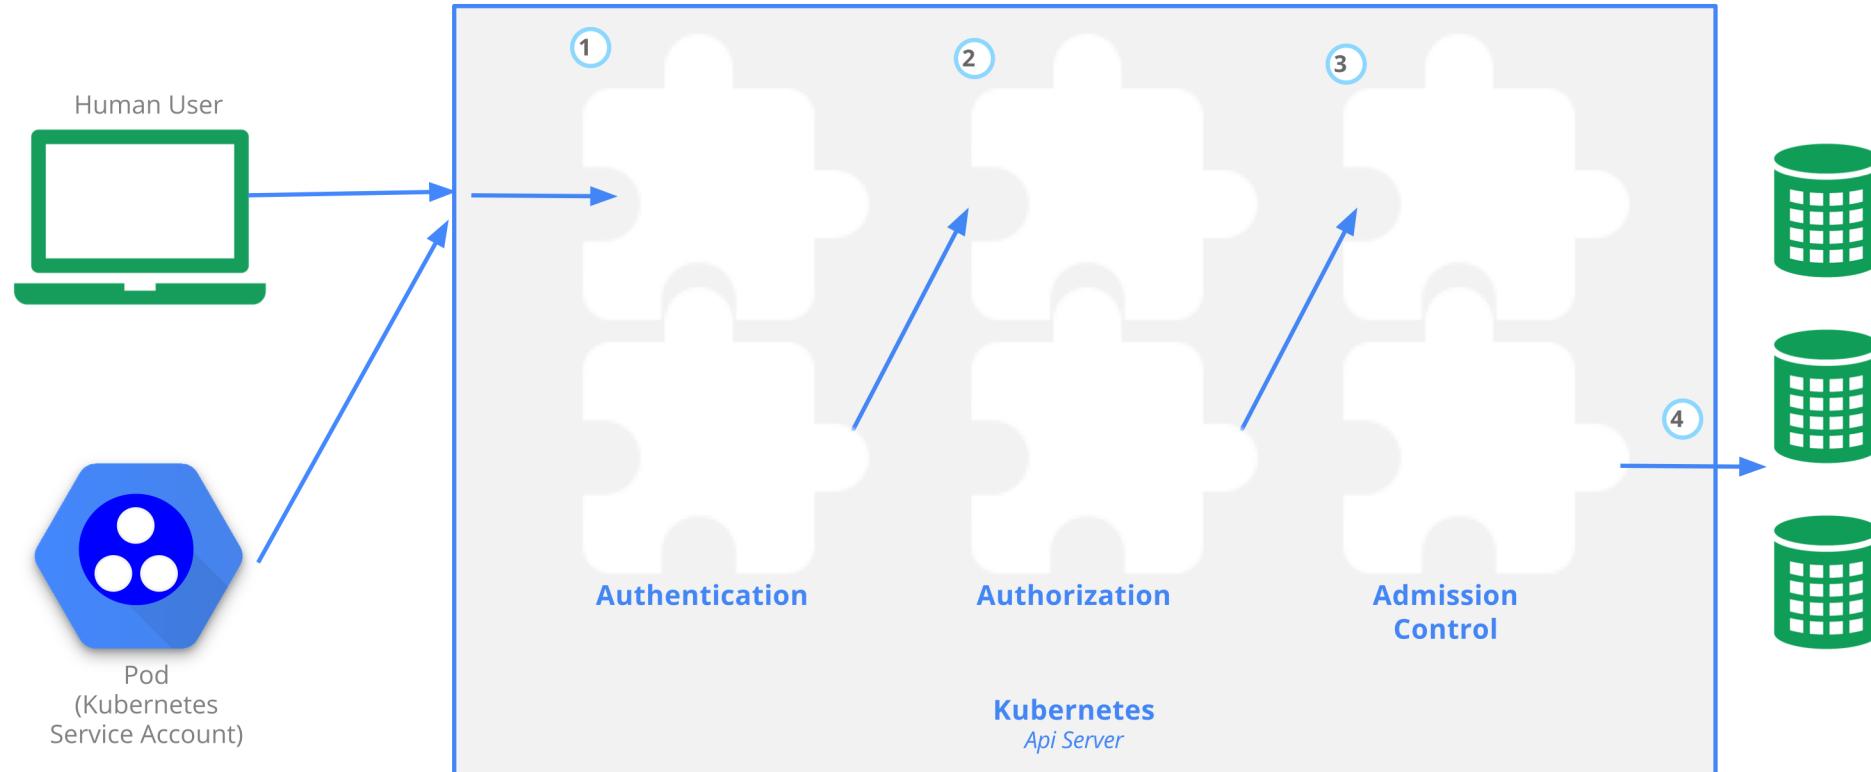
- **Advanced Isolation**

- Network Policies
- Sandboxing Containers

# Namespaces

- Namespace = Scope for names of Kubernetes resources
- Some resources are cluster-wide (Persistent Volume)
- Cannot be nested (yet?)
- Reserved namespaces:
  - kube-system (control plane)
  - kube-public (currently used for kubeadm)
  - default

# Access Control in Kubernetes

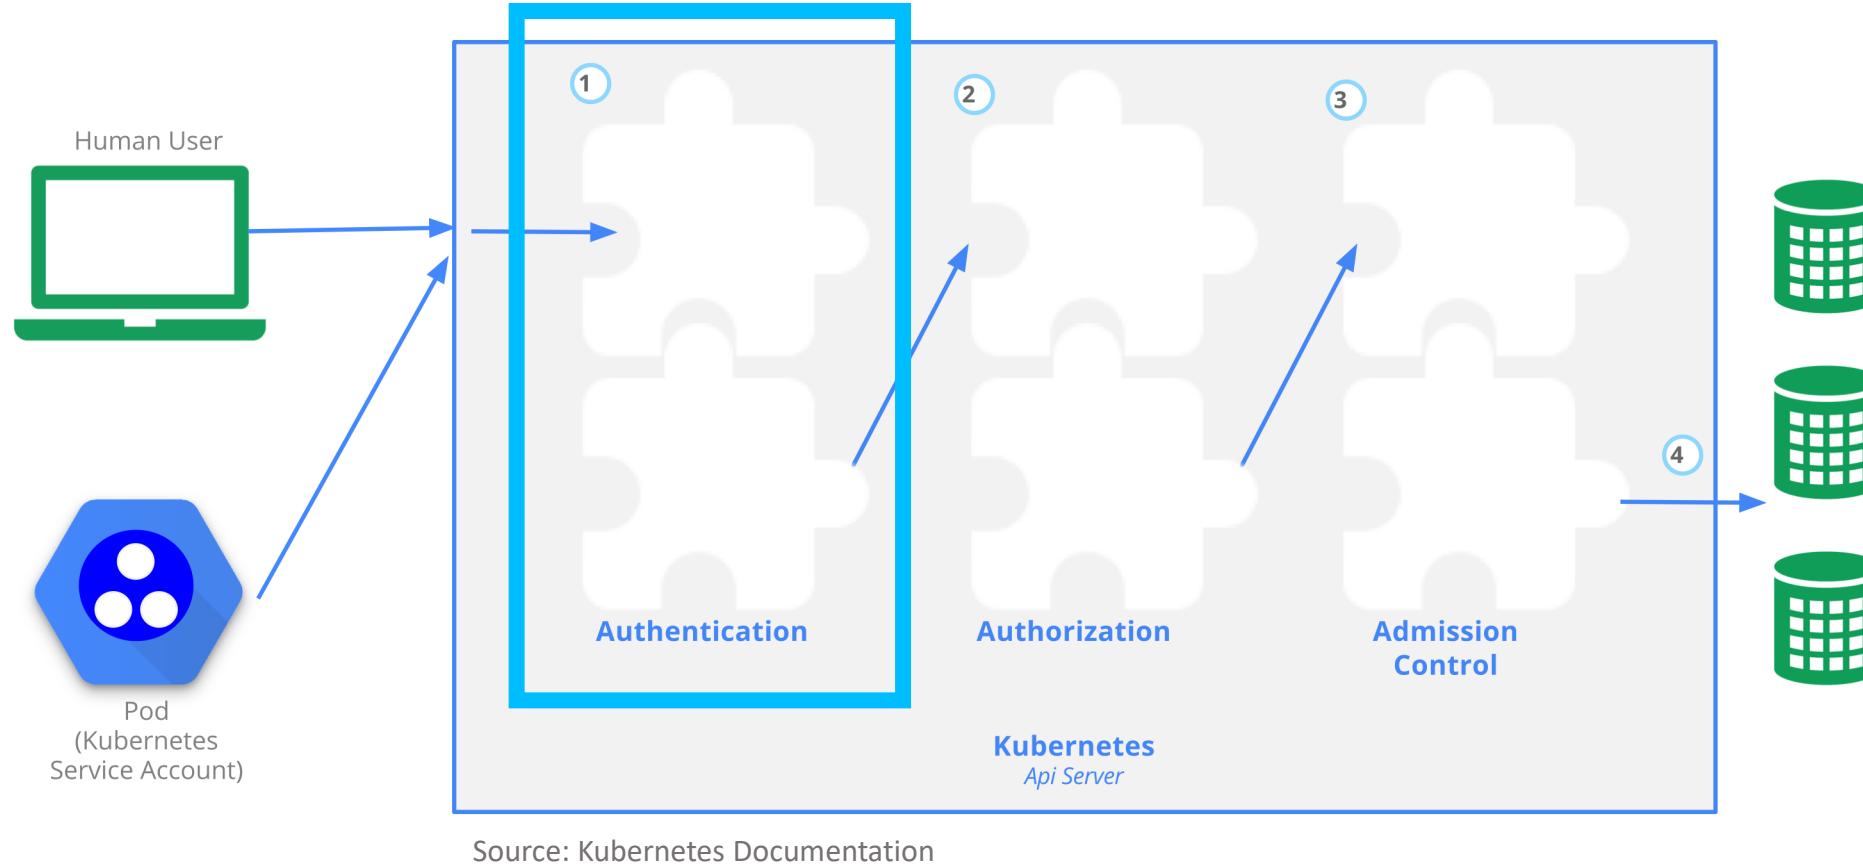


Source: Kubernetes Documentation

# Overview

- Basics
  - Access Control
  - Namespaces
- Authentication
  - OpenID Connect
  - Service Accounts
  - Rancher
- Authorization
- Admission Control
  - Resource Quotas
  - Limit Ranges
  - PodSecurityPolicies
  - Open Policy Agent
- Advanced Isolation
  - Network Policies
  - Sandboxing Containers

# Where are we?



# User Authentication

## ■ Several Methods

- Client Certificates (e.g. with flag --client-ca-file=CA-FILE.pem)
- HTTP Basic Auth (e.g. with flag --basic-auth-file=USER-FILE.csv)
- Bearer Tokens
  - using static token file (with flag --token-auth-file=TOKEN-FILE.csv)
  - using OIDC (OAuth2) => Dex project
  - using token verification via webhook
  - using service account tokens
- Authentication Proxy

⇒ Requires Kubernetes api server access

eyJhbGciOiJSUzI1NiIsImtpZCI6IiJ9.eyJpc3M  
 i0iJrdWJlc5ldGVzL3NlcnZpY2VhY2NvdW50Iiw  
 ia3ViZXJuZXRlc5pb9zZXJ2aWN1YWNjb3VudC9  
 uYW1lc3BhY2UiOijYyIsImt1YmVybmV0ZXMuaw8  
 vc2VydmljZWfjY291bnQvc2VjcmV0Lm5hbWUiOij  
 kZWZhdWx0LXRva2VuLTJoa2ZkIiwia3ViZXJuZX  
 lc5pb9zZXJ2aWN1YWNjb3VudC9zZXJ2aWN1LWF  
 jY291bnQubmFtZSI6ImR1ZmF1bHQiLCJrdWJlc5  
 1dGVzLmlvL3NlcnZpY2VhY2NvdW50L3NlcnZpY2U  
 tYWNjb3VudC51aWQiOii2MGQyMTg2MS0wNjQwLTE  
 xZWEtOGY1ZC0wMjUwMDAwMDAwMDEiLCJzdWIiOij  
 zeXN0ZW06c2VydmljZWfjY291bnQ6Y2M6ZGVmYXV  
 sdCJ9.rDmdpTKruM1KRB3Y3ygxS0-  
 7SMBzQ08wQBaSWJ7-  
 K4mTyekUKBvyLvk0qfkUdHwAljntk\_CeuMfcBnC6  
 Aed1j8PHwYBr8MJWpL376V2FEQPB90YJWYcH98bW  
 AL\_kw\_QHuaERU8inc0R9cqvg5jjef-  
 won6VpsLWnUeOLfhQo7sNm\_VIcunnUH0k6VUeR3a  
 Etxx7yPuGP2DT\_a0eeUGBGAc94QWdQxssu38pxhY  
 5Pe8nyaH28e2kHcE0fTck0z82nHBseJDwEmfWTMX

## HEADER: ALGORITHM &amp; TOKEN TYPE

```
{
  "alg": "RS256",
  "kid": ""
}
```

## PAYLOAD: DATA

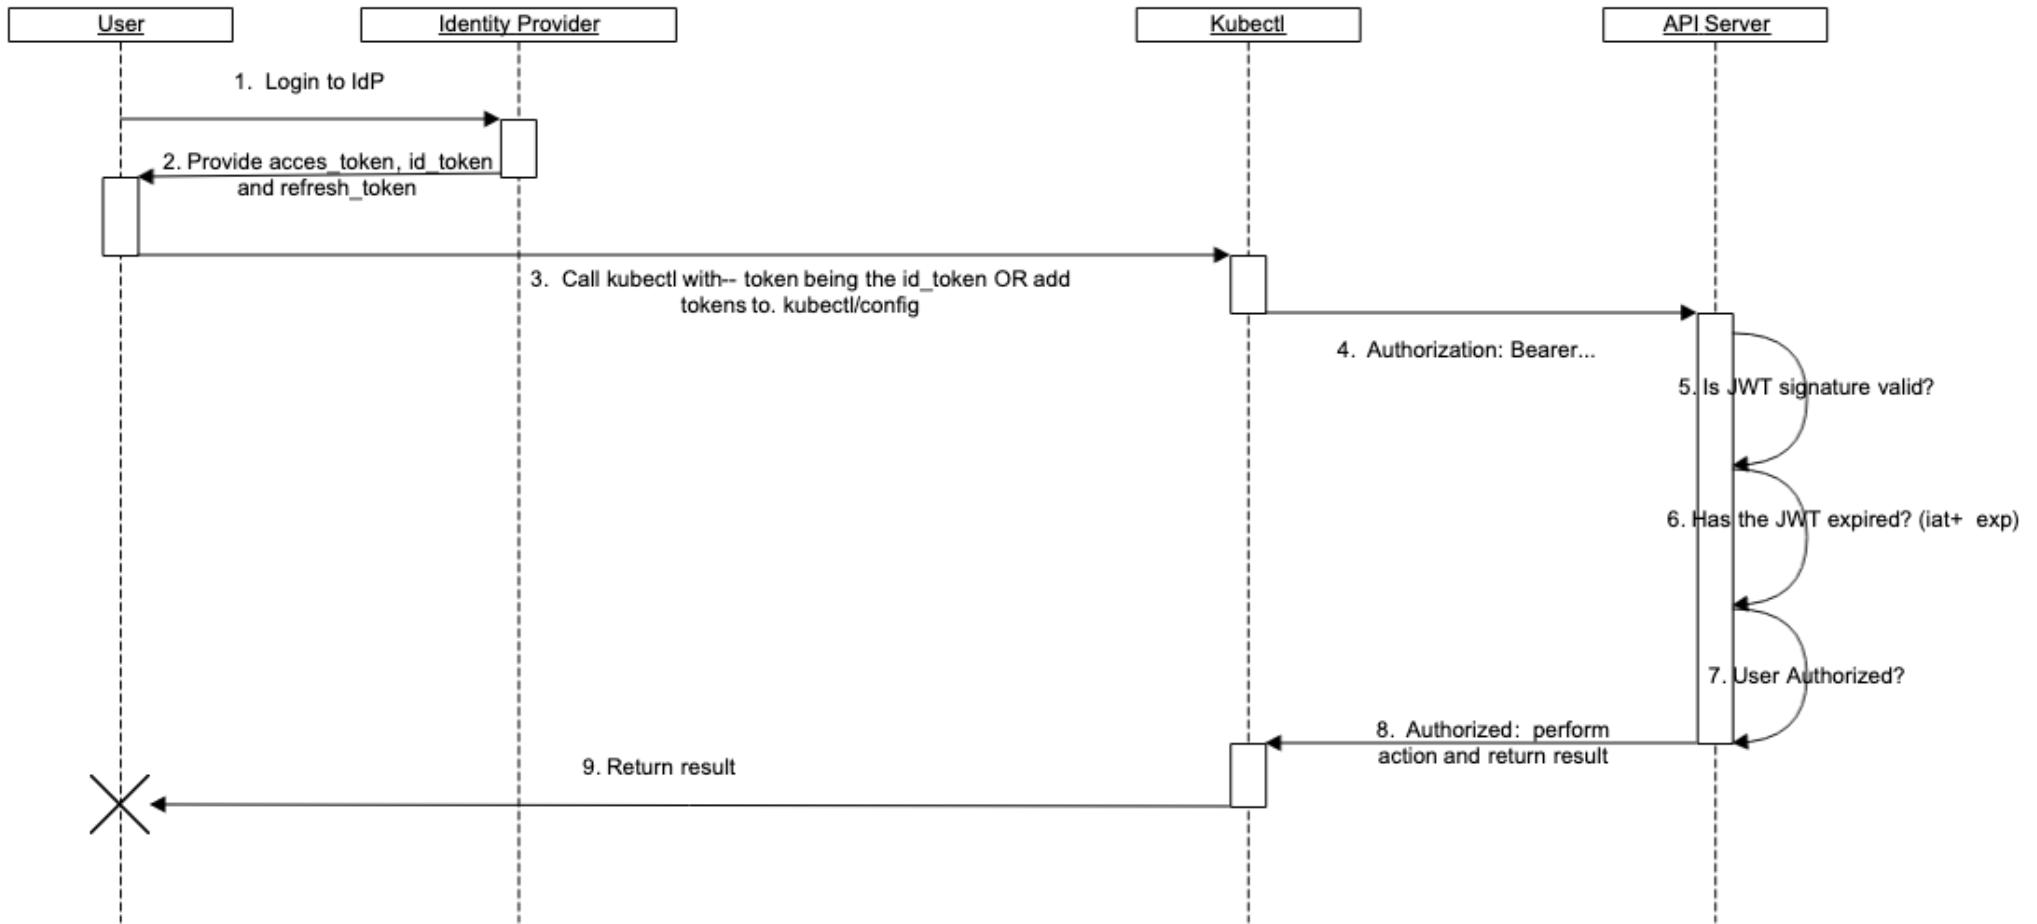
```
{
  "iss": "kubernetes/serviceaccount",
  "kubernetes.io/serviceaccount/namespace": "cc",
  "kubernetes.io/serviceaccount/secret.name": "default-token-2hkfd",
  "kubernetes.io/serviceaccount/service-account.name": "default",
  "kubernetes.io/serviceaccount/service-account.uid": "60d21861-0640-11ea-8f5d-025000000001",
  "sub": "system:serviceaccount:cc:default"
}
```

## VERIFY SIGNATURE

```
RSASHA256(
  base64UrlEncode(header) + "." +
  base64UrlEncode(payload),
  Public Key or Certificate. Ent
```

# OpenID Connect

Source: Kubernetes Documentation



# ServiceAccounts

- Usually used by in cluster services
- Kubernetes auto-generates ServiceAccount “default”
- Every container mounts a volume with the token of the ServiceAccount referenced by the container’s pod (can be disabled)

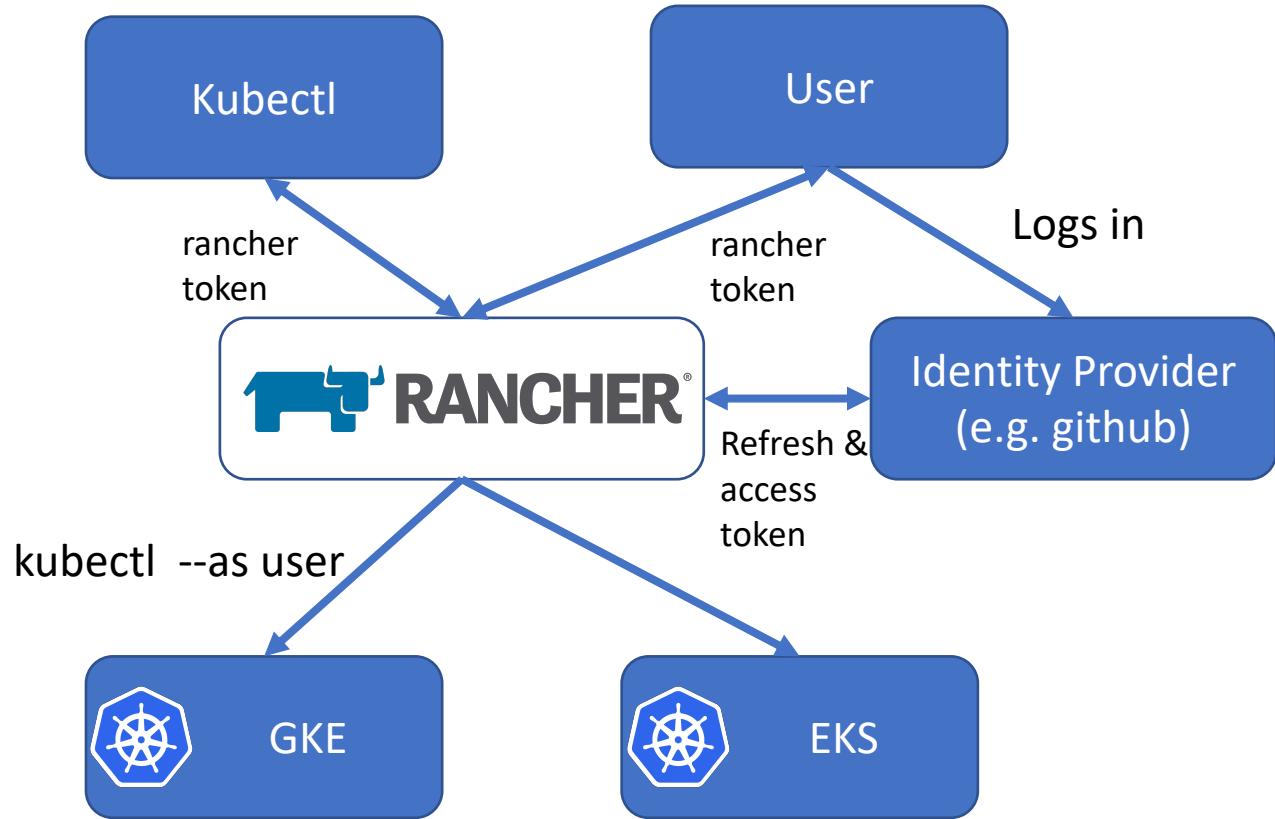
/var/run/secrets/kubernetes.io/serviceaccount

- Only Kubernetes managed account resource  
(Kubernetes does not know a user resource)

```
1  apiVersion: v1
2  kind: Config
3  clusters:
4    - cluster:
5      certificate-authority-data: LS0tLS1CRUdJTiBDRVJUSUZJQ0FURS0tLS0tCk1JSU
6      server: https://kubernetes.docker.internal:6443
7      name: docker
8  contexts:
9    - context:
10      cluster: docker
11      user: docker
12      name: docker
13  current-context: docker
14  users:
15    - name: docker
16      user:
17        token: eyJhbGciOiJSUzI1NiIsImtpZCI6IiJ9.eyJpc3MiOiJrdWJlcmtldGVzL3Nlcn
```

# Rancher Authentication Proxy

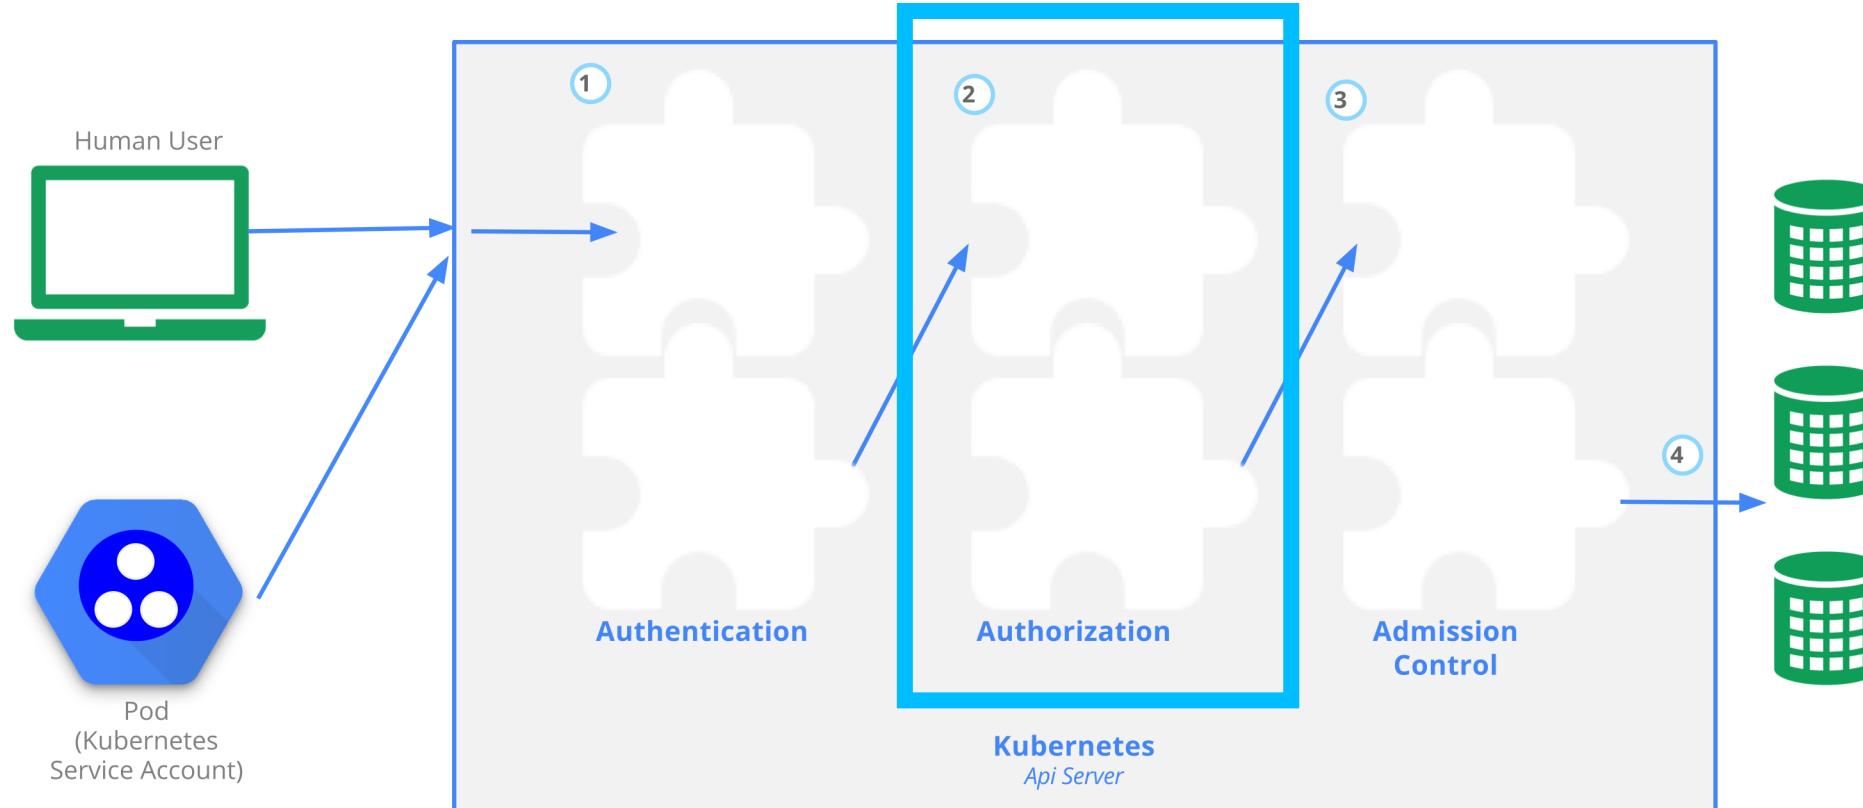
- Rancher forwards user requests to Kubernetes using a service account and impersonation
- Supports local or external authentication (e.g. github, ldap etc.)
- Users can access multiple clusters with one single set of credentials



# Overview

- Basics
  - Access Control
  - Namespaces
- Authentication
  - OpenID Connect
  - Service Accounts
  - Rancher
- Authorization
- Admission Control
  - Resource Quotas
  - Limit Ranges
  - PodSecurityPolicies
  - Open Policy Agent
- Advanced Isolation
  - Network Policies
  - Sandboxing Containers

# Where are we?



# RBAC (Role-based access control)

- Stable since Kubernetes 1.8
- Successor of ABAC
  - (using kubectl, via Kubernetes resources)

```
1  apiVersion: rbac.authorization.k8s.io/v1
2  kind: Role # Namespace-scoped
3  metadata:
4    namespace: default
5    name: pod-reader
6  rules:
7    - apiGroups: [""]
      indicates the core API group
8      resources: ["pods"]
9      verbs: ["get", "watch", "list"]
```

- Roles define what type of access is allowed on resources
  - Roles for namespace scoped authorization
  - ClusterRoles for cluster scoped authorization

# RBAC - RoleBindings

- **Role bindings assign Users / Groups / ServiceAccounts to roles**
  - ClusterRoleBinding binds ClusterRoles
  - RoleBinding binds ClusterRoles / Roles
- **Multiple subjects per binding possible**

```
1  apiVersion: rbac.authorization.k8s.io/v1
2  # This role binding allows "fabian" to
3  # read pods in the "default" namespace.
4  kind: RoleBinding
5  metadata:
6  | name: read-pods
7  | namespace: default
8  subjects:
9  | kind: User
10 | name: fabian
11 | apiGroup: rbac.authorization.k8s.io
12 roleRef:
13 | # Role or ClusterRole
14 | kind: Role
15 | # this must match the name of the Role
16 | # or ClusterRole you wish to bind to
17 | name: pod-reader
18 | apiGroup: rbac.authorization.k8s.io
```

# Overview

- **Basics**

- Access Control
- Namespaces

- **Authentication**

- OpenID Connect
- Service Accounts
- Rancher

- **Authorization**

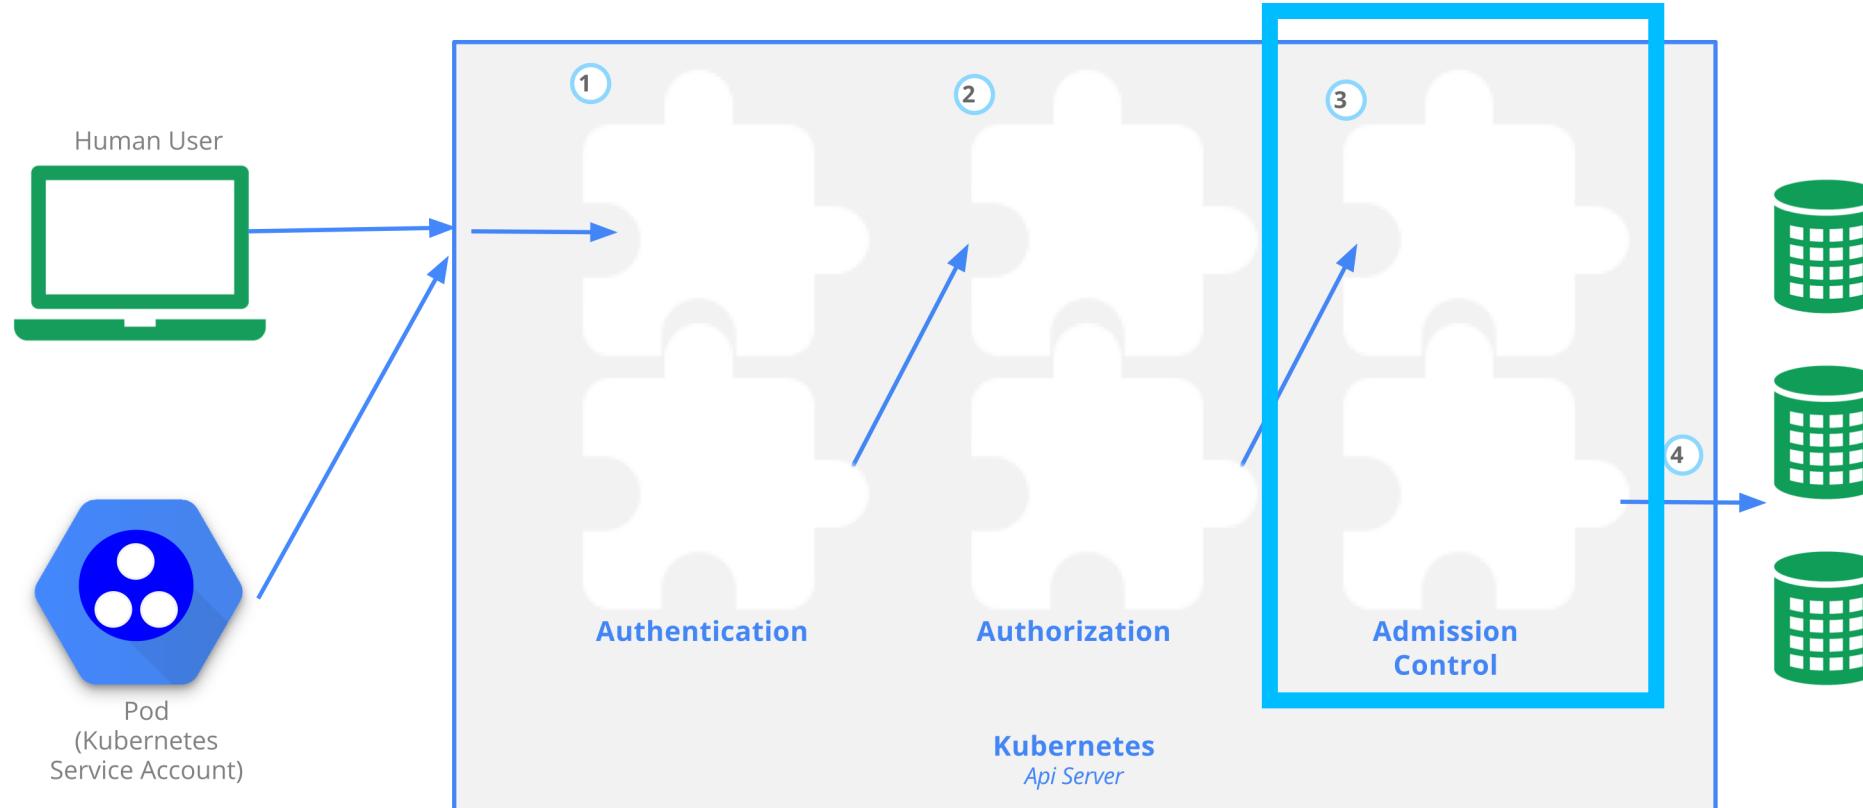
- **Admission Control**

- Resource Quotas
- Limit Ranges
- PodSecurityPolicies
- Open Policy Agent

- **Advanced Isolation**

- Network Policies
- Sandboxing Containers

# Where are we?

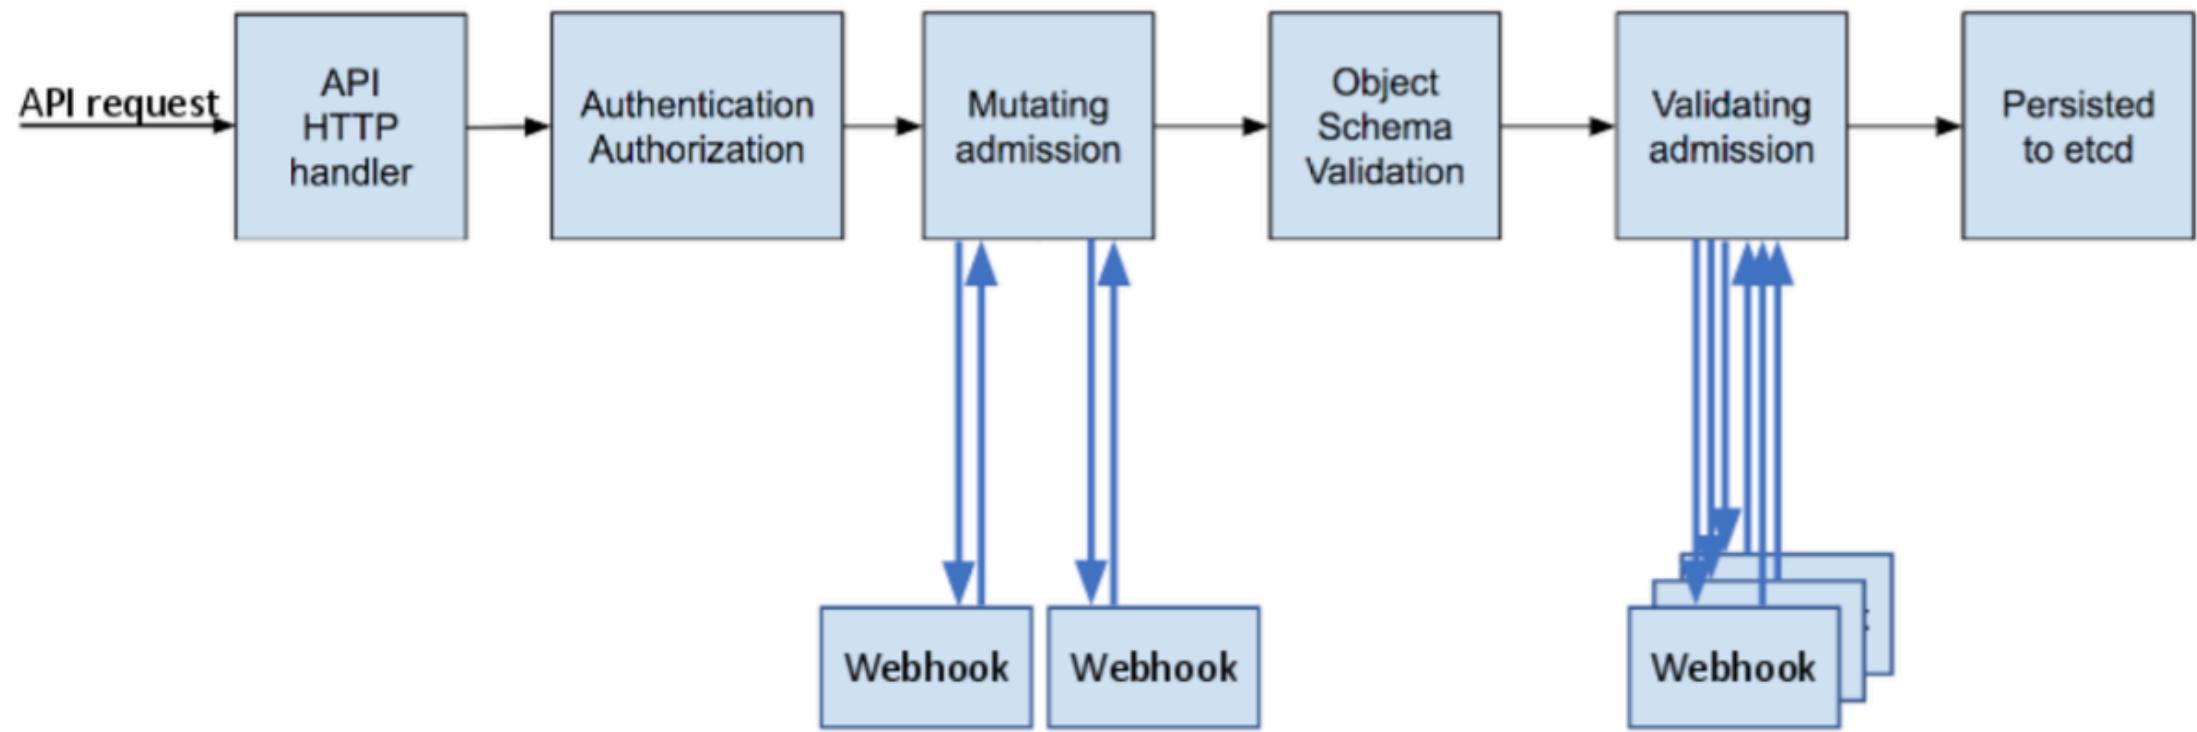


Source: Kubernetes Documentation

# Admission Controllers

- Intercept requests to the API server
- Types:
  - Mutating – Change the resources sent within the request (called first)
  - Validating – Allow or deny the resources sent within the request
- 30+ prebuilt admission controllers available
- MutatingAdmissionWebhook and ValidatingAdmissionWebhook allow to add dynamic admission control via webhooks

# Admission Controllers



Source: Kubernetes Blog

# ResourceQuota

- Enforced as admission controller
- Limits aggregate resource consumption per namespace
- Common cases:
  - Limit total cpu/memory requests/limits
  - Limit total persistent storage or ephemeral storage
  - Limit total number of pods/services/secrets/configmaps/...
  - Limit total number of resource quotas ;-)

```
1  apiVersion: v1
2  kind: ResourceQuota
3  metadata:
4    name: quota
5    namespace: myspace
6  spec:
7    hard:
8      count/roles.rbac.authorization.k8s.io: "5"
9      pods: "100"
10     requests.cpu: "8"
11     limits.cpu: "16"
```

# LimitRange

- Allows to set min | max | default values for cpu | memory | storage
- Prevents that a single pod/container consumes all available resources in a namespace (Resource Quota)
- Also available for persistent volume claims

```
1  apiVersion: v1
2  kind: LimitRange
3  metadata:
4    name: mylimitrange
5    namespace: myspace
6  spec:
7    # Limits the container
8    # resources to 1-4 cores
9    limits:
10   - max:
11     - cpu: "4"
12     min:
13     - cpu: "1"
14     type: Container
```

```
1  apiVersion: v1
2  kind: Pod
3  metadata:
4    name: nginx
5    labels:
6      name: nginx
7  spec:
8    containers:
9      name: nginx
10     image: nginx
11     ports:
12       containerPort: 80
13     resources:
14       requests:
15         cpu: "0"
16         memory: "0"
17       limits:
18         cpu: "0"
19         memory: "0"
```

```
1  apiVersion: v1
2  kind: Pod
3  metadata:
4    name: nginx
5    labels:
6      name: nginx
7  spec:
8    containers:
9      name: nginx
10     image: nginx
11     ports:
12       containerPort: 80
13     volumeMounts:
14       mountPath: /mnt
15         name: bad-volume
16     volumes:
17       name: bad-volume
18       hostPath:
19         path: /
```

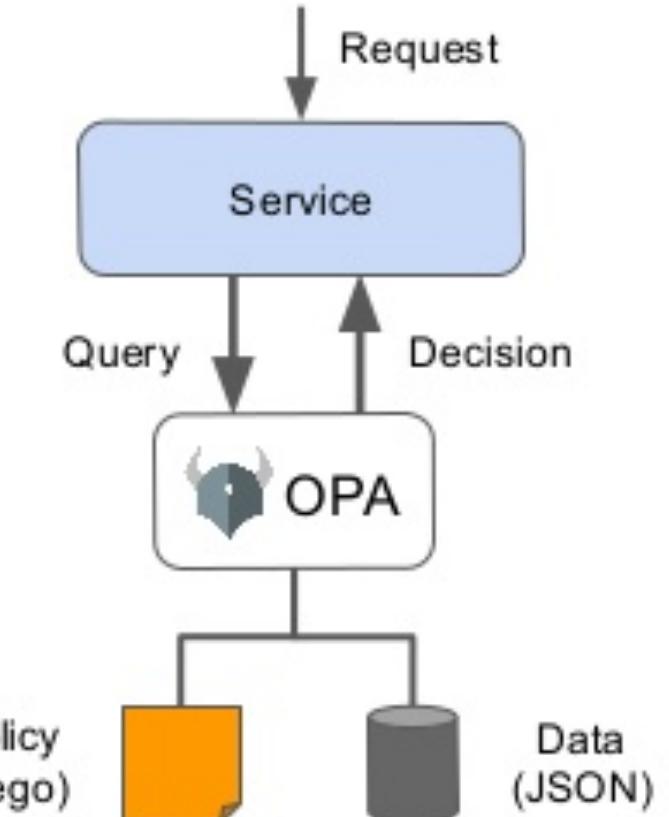
# PodSecurityPolicy

- Allows admins to limit the capabilities of pods
- Once enabled no pods can be deployed without active policy
- Common cases:
  - Disable privileged containers
  - Disable use of hostNetwork
  - Disable certain volume types (e.g. hostPath)
  - Restrict use of Linux capabilities (e.g. CAP\_CHOWN)
- Assignable to roles

```
1 apiVersion: extensions/v1beta1
2 kind: PodSecurityPolicy
3 metadata:
4   name: policy
5 spec:
6   allowPrivilegeEscalation: true
7   fsGroup:
8     rule: RunAsAny
9   hostPorts:
10    - max: 0
11    - min: 0
12   runAsUser:
13     rule: RunAsAny
14   seLinux:
15     rule: RunAsAny
16   supplementalGroups:
17     rule: RunAsAny
18   volumes:
19     persistentVolumeClaim
```

# Open Policy Agent (OPA)

- Dynamic admission controller
- Language for policy logic: Rego
- Allows to define admission rules as Kubernetes resources
- Audits
- Gatekeeper project adds custom resource definitions for policies



Source: OPA Documentation

# Overview

- **Basics**

- Access Control
- Namespaces

- **Authentication**

- OpenID Connect
- Service Accounts
- Rancher

- **Authorization**

- **Admission Control**

- Resource Quotas
- Limit Ranges
- PodSecurityPolicies
- Open Policy Agent

- **Advanced Isolation**

- Network Policies
- Sandboxing Containers

# NetworkPolicy

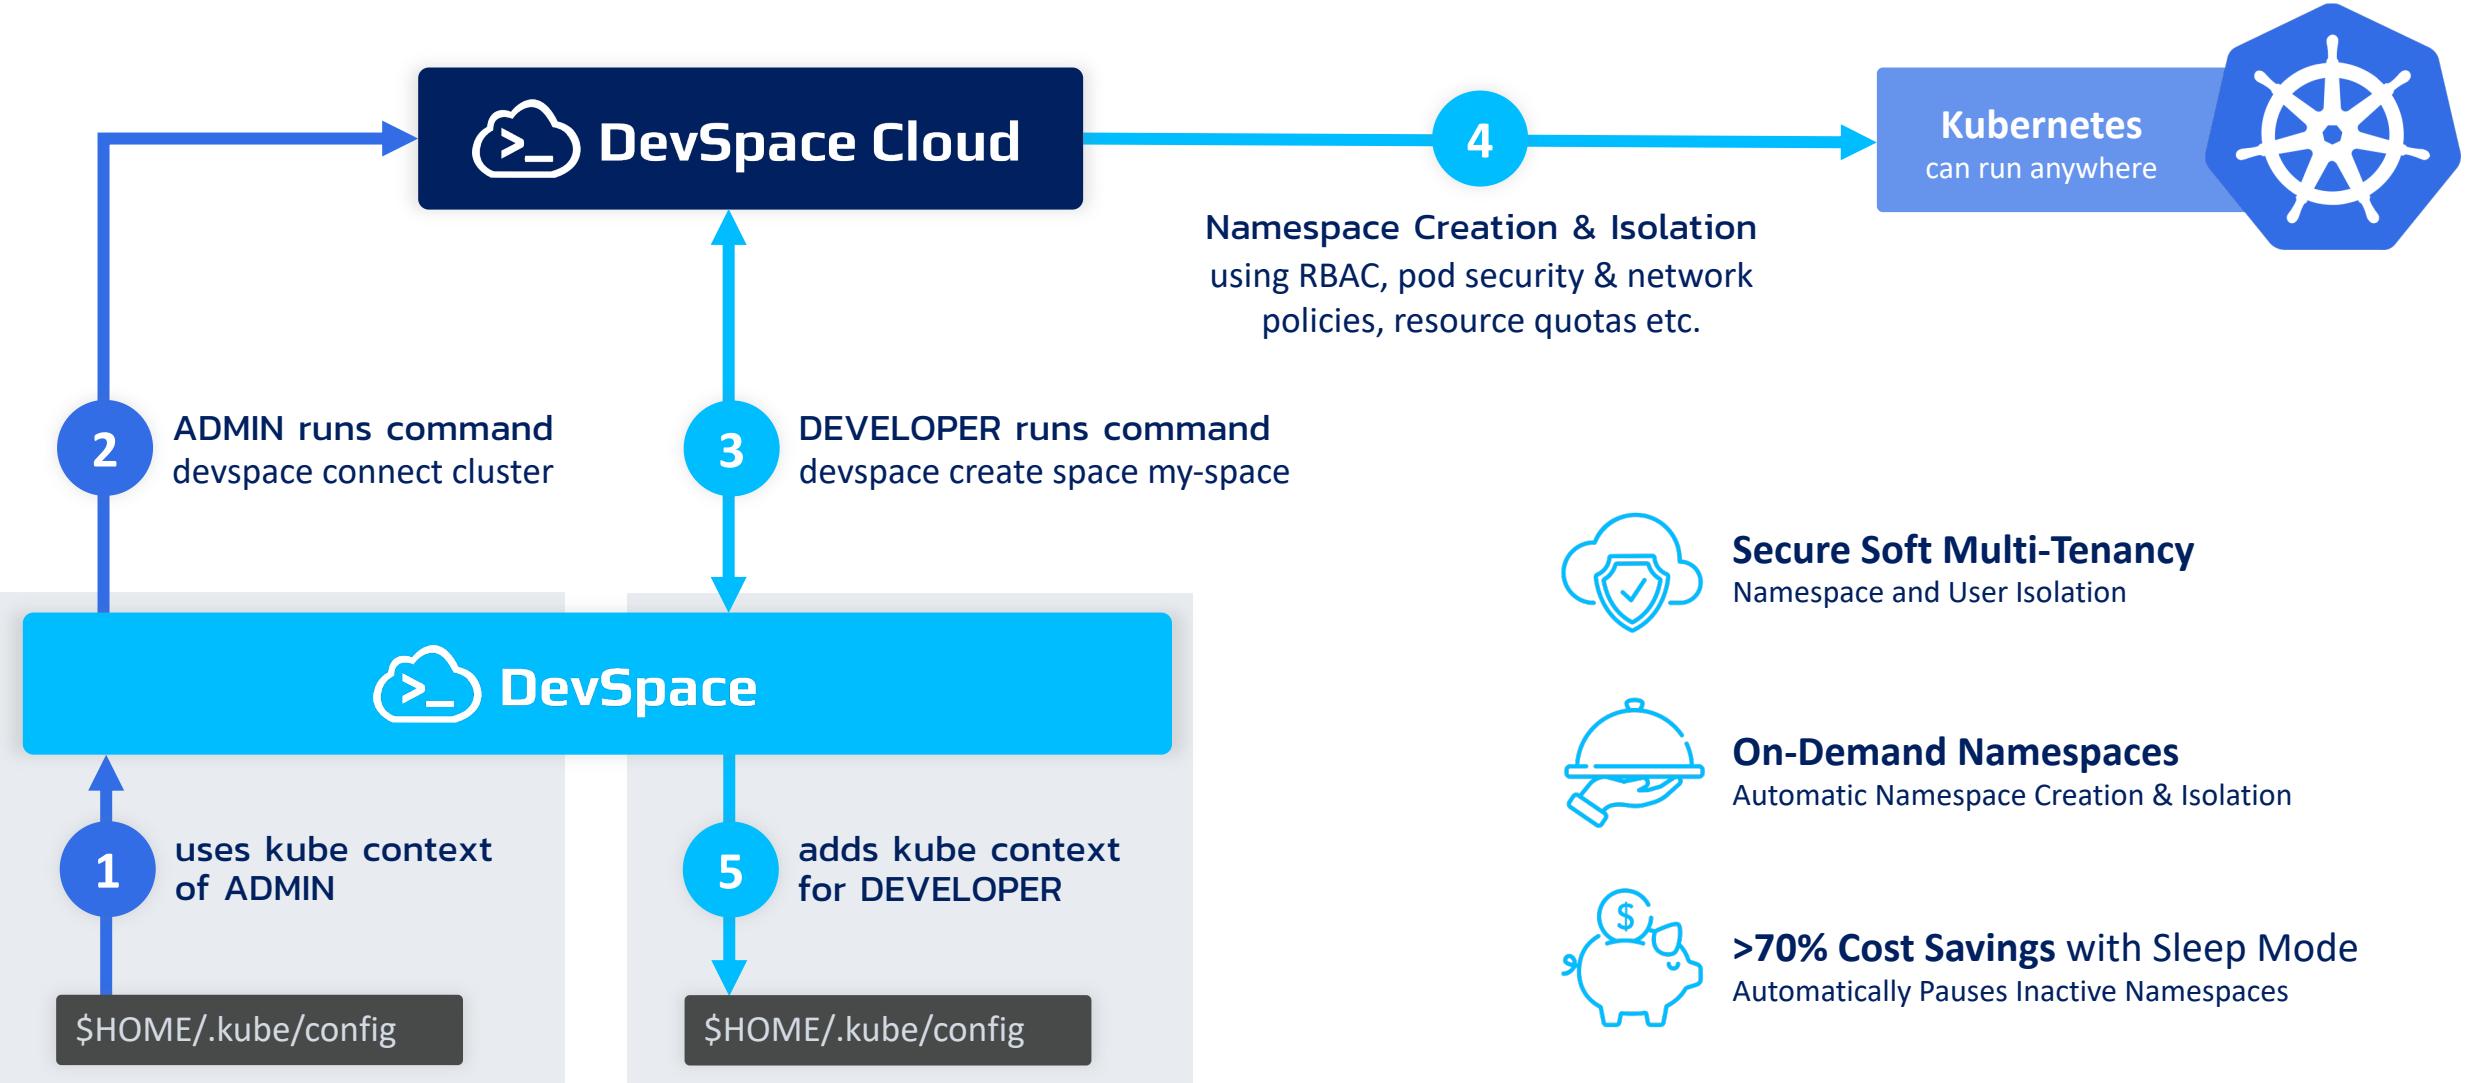
- **Restricts the incoming (ingress) and outgoing (egress) network traffic of pods**
- **Pods initially allow all ingress and egress traffic if they are not selected by a network policy**
- **Must be supported by the underlying network plugin!**

```
1  apiVersion: extensions/v1beta1
2  kind: NetworkPolicy
3  metadata:
4    name: networkpolicy
5    namespace: myspace
6  spec:
7    egress:
8      - to:
9        podSelector: {}
10       namespaceSelector:
11         matchLabels:
12           my.cloud.com/traffic-allowed: allowed
13       podSelector: {}
14     policyTypes:
15       - Egress
```

# Sandboxing Containers

- **Containers != VMs**
  - Sharing a kernel bears risks
  - Vulnerabilities can lead to programs escaping the “soft” boundaries of a container
- **Sandboxing Methods**
  - Container Runtimes for Isolation
    - [gVisor](#) (user-space kernel)
    - [Katacontainers](#) (lightweight vm)
  - Nodeless Kubernetes
    - [Virtual Kubelet](#)
    - [Elotl](#)

# How are we using all of these features in DevSpace Cloud?



# I'm here to help. Get in touch.



- Connect on Twitter: [@FabianKramm](https://twitter.com/FabianKramm) 
- Email me: [fk@devspace.cloud](mailto:fk@devspace.cloud)
- Check out our projects on GitHub:



Get the fastest development workflow for Kubernetes



DevSpace on GitHub



Turn any Kubernetes cluster into a powerful development platform



DevSpace Cloud on GitHub