# Real Security for Services on Kubernetes

Eric Wang, Yun Zhang

Dec 8, 2017

databricks

# Who are we?

- Software Engineers on Cloud Platform team@Databricks
- Provide Infrastructure Services to other engineering teams
  - Deployment Platform (Kubernetes)
  - Monitoring System
  - Permission & Credential Management
- Use a lot Open Source Tools
  - Kubernetes, Prometheus, Bazel, Hashicorp Terraform & Vault, etc

# Agenda

- What is Databricks
- Databricks Security Concerns
- Kubernetes Access Control
- Kubernetes Secret Management
- Kubernetes Audit Logging

# What is Databricks?

- Creators of Apache Spark
  - https://spark-summit.org/2017/events/apache-spark-on-kubernetes/
- Unified Analytics Platform
  - Databricks Notebook
  - SaaS product on AWS & Azure
  - Provides Better Performance, Collaboration, **Security** around Apache Spark

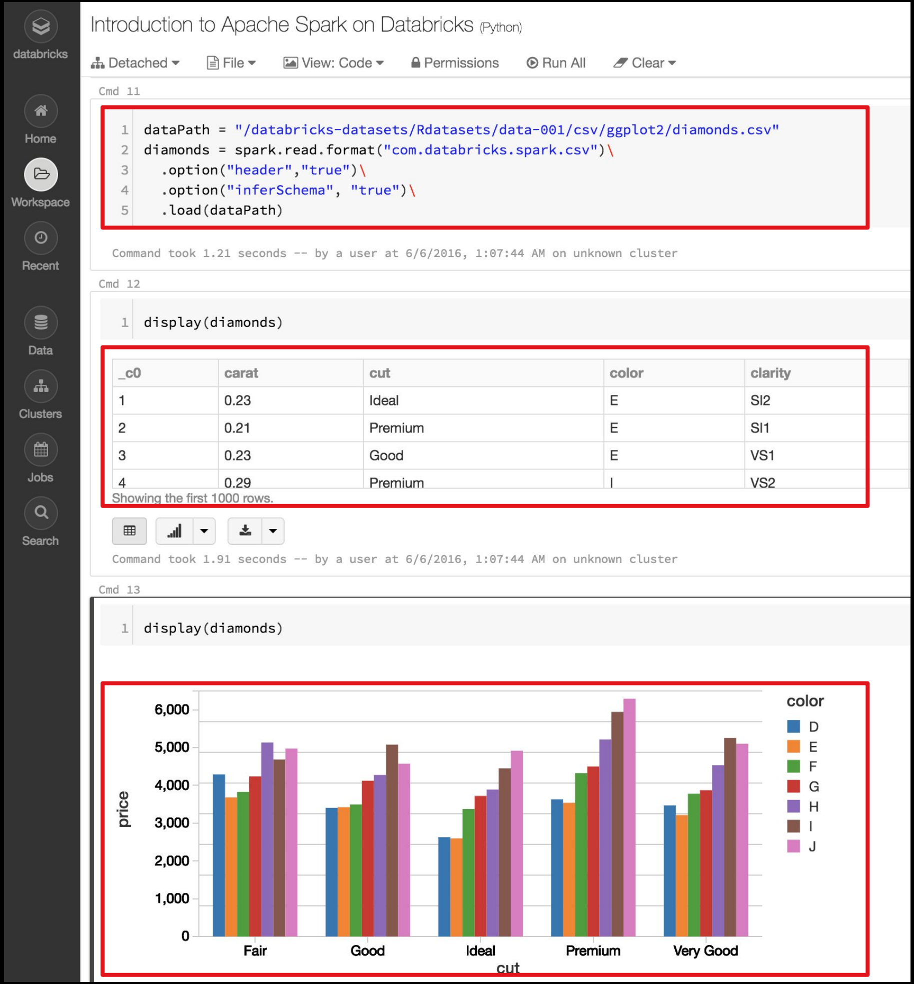

Cmd 11

```
dataPath = "/databricks-datasets/Rdatasets/data-001/csv/ggplot2/diamonds.csv"
diamonds = spark.read.format("com.databricks.spark.csv")\
  .option("header","true")\
  .option("inferSchema", "true")\
  .load(dataPath)
```

```
Command took 1.21 seconds -- by a user at 6/6/2016, 1:07:44 AM on unknown cluster
```

Cmd 12

```
display(diamonds)
```

| _c0 | carat | cut | color | clarity |
|---|---|---|---|---|
| 1 | 0.23 | Ideal | E | SI2 |
| 2 | 0.21 | Premium | E | SI1 |
| 3 | 0.23 | Good | E | VS1 |
| 4 | 0.29 | Premium | I | VS2 |

Showing the first 1000 rows.

```
Command took 1.91 seconds -- by a user at 6/6/2016, 1:07:44 AM on unknown cluster
```

Cmd 13

```
display(diamonds)
```

# How does Databricks work?

- Databricks Environment and Customer Environment

- Control Plane Services deployed to Kubernetes clusters

- Security of Kubernetes clusters is **critical**!

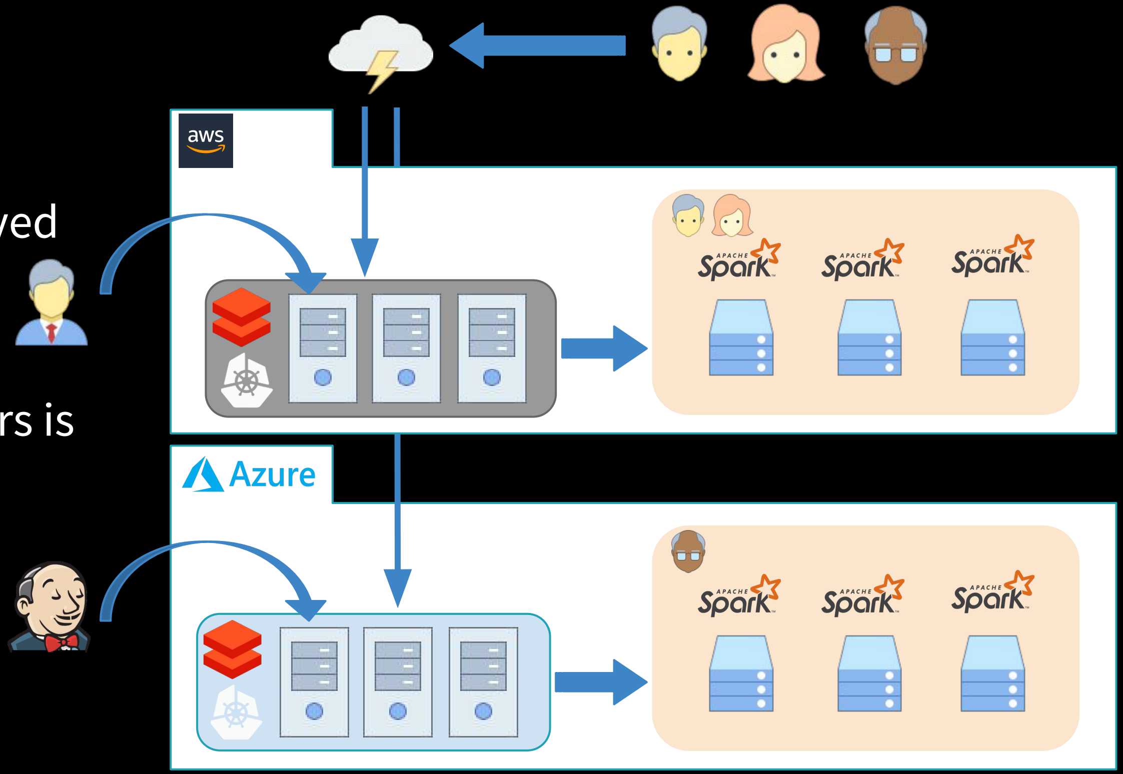

databricks

# Databricks Security Concerns

**Customer:**

- Customer data remains private
- Security Compliance, e.g. HIPAA, SOC2

**AppSec:**

- Defense in Depth
- Limited and audited production access

**Engineer:**

- Security by default
- Easy to extend

✓ **Access Control**

✓ **Secret Management**

✓ **Audit Logging**

# Agenda

- What is Databricks
- Databricks Security Concerns
- Kubernetes Access Control
- Kubernetes Secret Management
- Kubernetes Audit Logging

# How do we do Access Control?

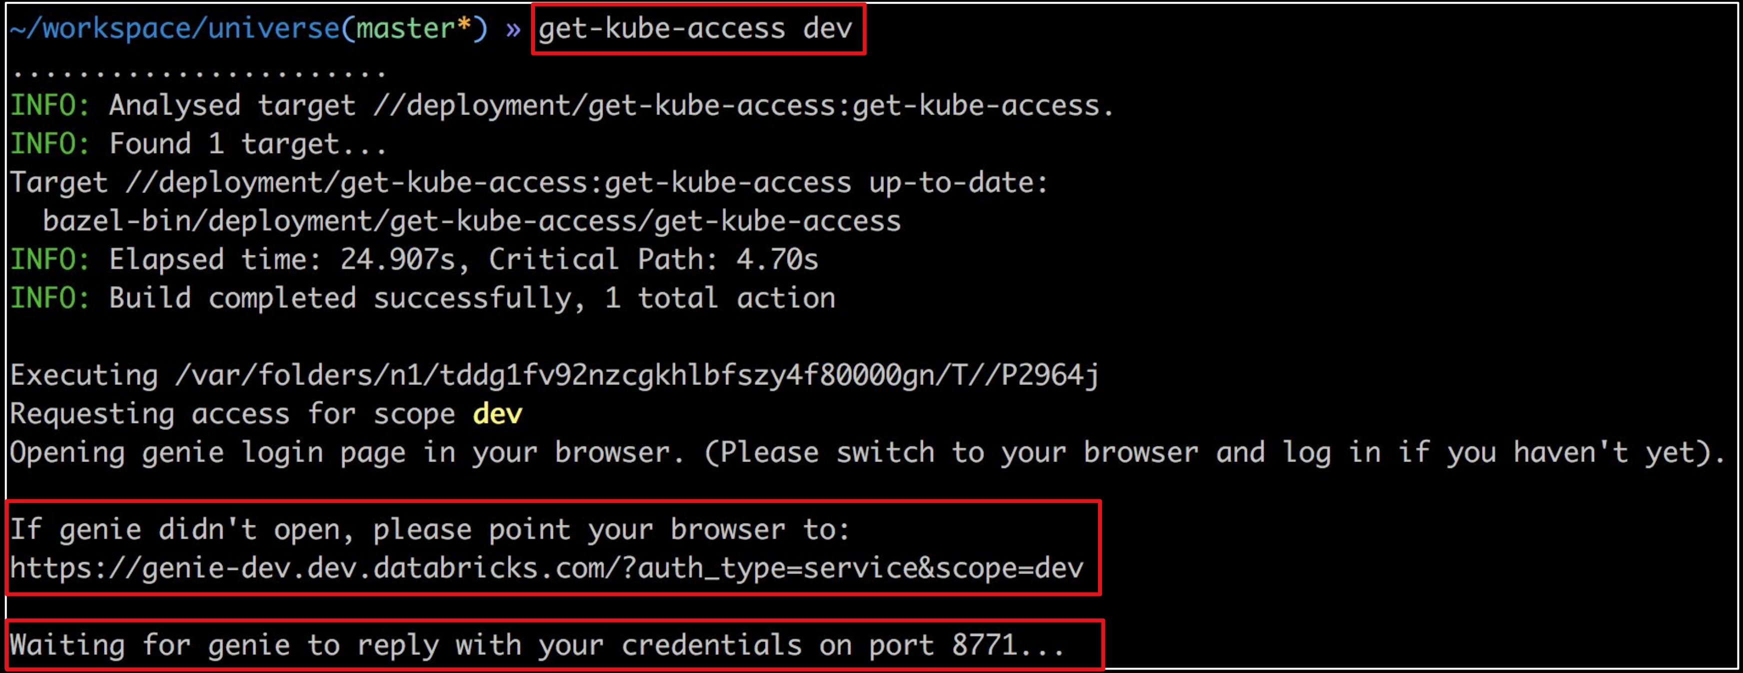

```
~/workspace/universe(master*) » get-kube-access dev
........................
INFO: Analysed target //deployment/get-kube-access:get-kube-access.
INFO: Found 1 target...
Target //deployment/get-kube-access:get-kube-access up-to-date:
  bazel-bin/deployment/get-kube-access/get-kube-access
INFO: Elapsed time: 24.907s, Critical Path: 4.70s
INFO: Build completed successfully, 1 total action

Executing /var/folders/n1/tddg1fv92nzcgkhlbfszy4f80000gn/T//P2964j
Requesting access for scope dev
Opening genie login page in your browser. (Please switch to your browser and log in if you haven't yet).

If genie didn't open, please point your browser to:
https://genie-dev.dev.databricks.com/?auth_type=service&scope=dev

Waiting for genie to reply with your credentials on port 8771...
```

# How do we do Access Control?

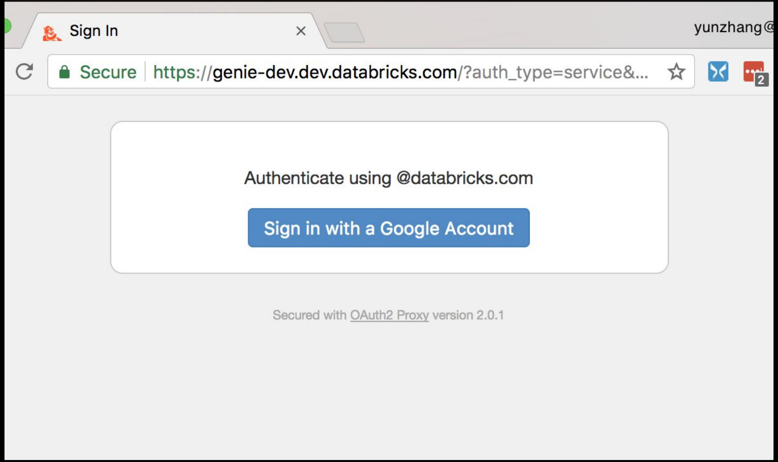

Authenticate using @databricks.com

Sign in with a Google Account

Secured with OAuth2 Proxy version 2.0.1

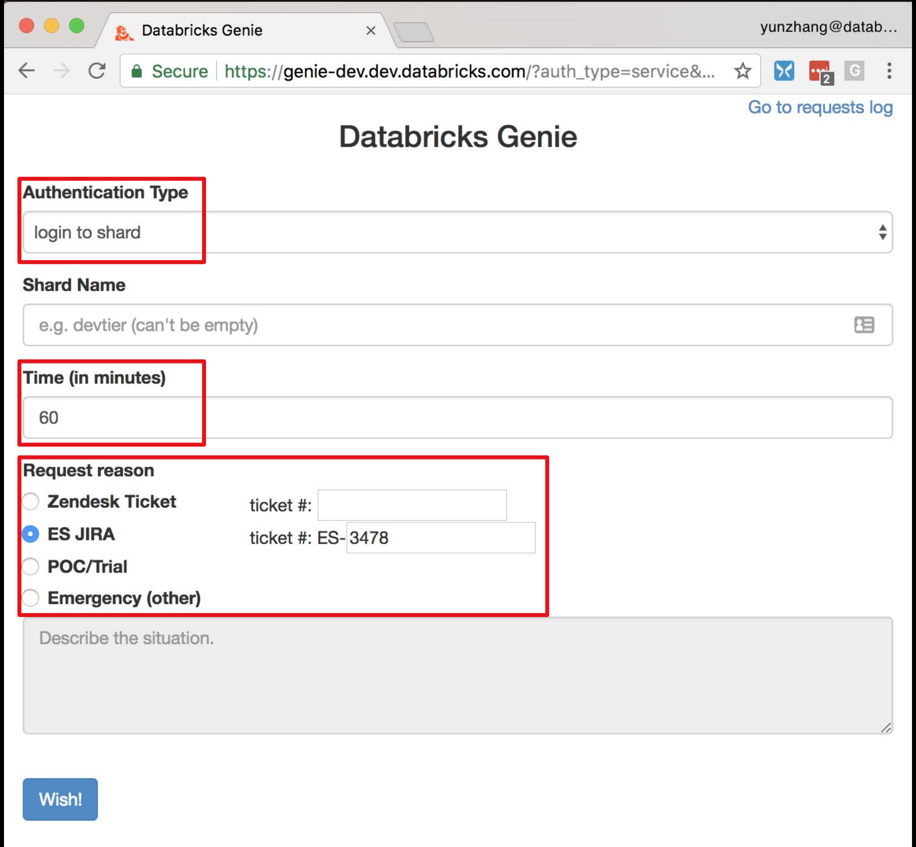

Go to requests log

## Databricks Genie

**Authentication Type**

login to shard

**Shard Name**

e.g. devtier (can't be empty)

**Time (in minutes)**

60

**Request reason**

- ○ **Zendesk Ticket** ticket #:
- ● **ES JIRA** ticket #: ES-3478
- ○ **POC/Trial**
- ○ **Emergency (other)**

Describe the situation.

Wish!

# How do we do Access Control?

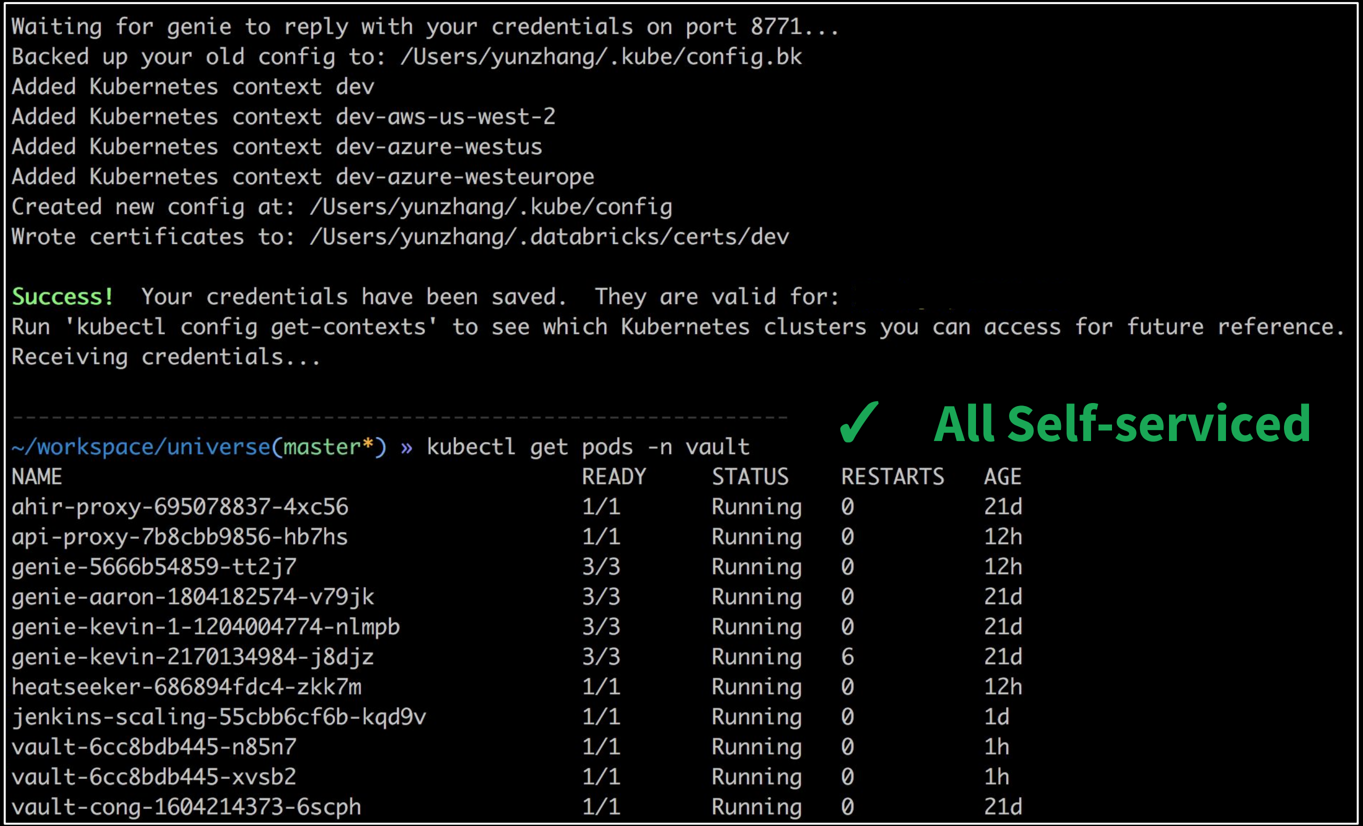

```
Waiting for genie to reply with your credentials on port 8771...
Backed up your old config to: /Users/yunzhang/.kube/config.bk
Added Kubernetes context dev
Added Kubernetes context dev-aws-us-west-2
Added Kubernetes context dev-azure-westus
Added Kubernetes context dev-azure-westeurope
Created new config at: /Users/yunzhang/.kube/config
Wrote certificates to: /Users/yunzhang/.databricks/certs/dev

Success!  Your credentials have been saved.  They are valid for:
Run 'kubectl config get-contexts' to see which Kubernetes clusters you can access for future reference.
Receiving credentials...

-------------------------------------------------------------
~/workspace/universe(master*) » kubectl get pods -n vault
NAME                                   READY     STATUS    RESTARTS   AGE
ahir-proxy-695078837-4xc56             1/1       Running   0          21d
api-proxy-7b8cbb9856-hb7hs             1/1       Running   0          12h
genie-5666b54859-tt2j7                 3/3       Running   0          12h
genie-aaron-1804182574-v79jk           3/3       Running   0          21d
genie-kevin-1-1204004774-nlmpb         3/3       Running   0          21d
genie-kevin-2170134984-j8djz           3/3       Running   6          21d
heatseeker-686894fdc4-zkk7m            1/1       Running   0          12h
jenkins-scaling-55cbb6cf6b-kqd9v       1/1       Running   0          1d
vault-6cc8bdb445-n85n7                 1/1       Running   0          1h
vault-6cc8bdb445-xvsb2                 1/1       Running   0          1h
vault-cong-1604214373-6scph            1/1       Running   0          21d
```

✓ **All Self-serviced**

# What Just Happened?

Genie

- “The” service for access control
- Integrates with Google Group
- Integrated with Hashicorp Vault for certificate generation
- Issues certificates applicable to multiple services

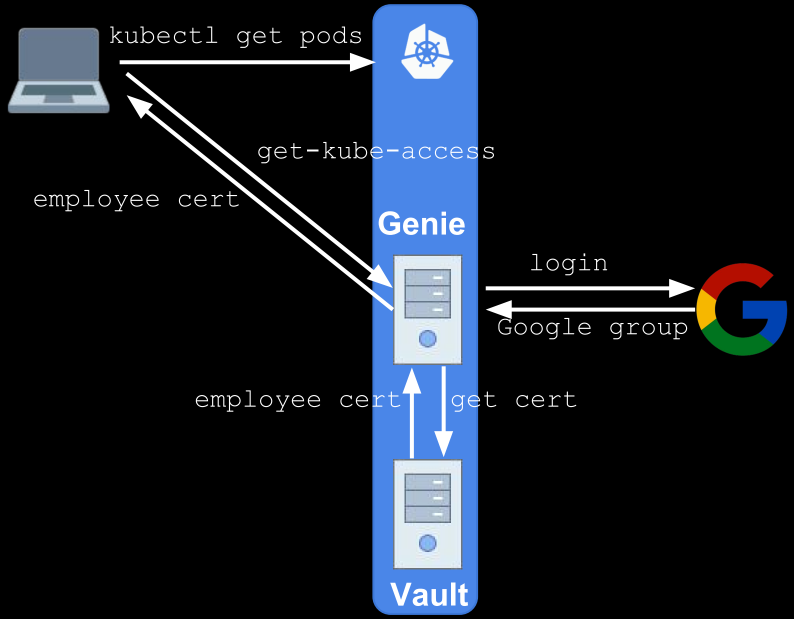

# Employee Certificate

- Signed by global Employee CA
- Short validity
- Identity: Employee’s email address in Common Name (CN)
- Role: Google Group in Subject (Organization) & Subject Alternative Names (SAN)

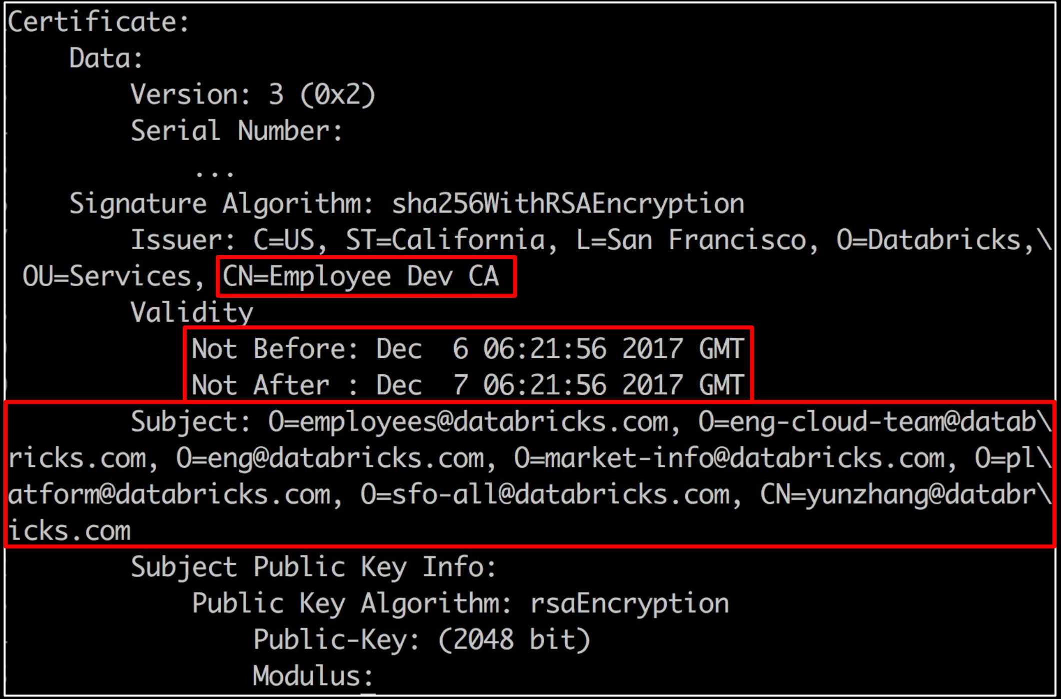

```
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number:
            ...
    Signature Algorithm: sha256WithRSAEncryption
        Issuer: C=US, ST=California, L=San Francisco, O=Databricks,\
 OU=Services, CN=Employee Dev CA
        Validity
            Not Before: Dec  6 06:21:56 2017 GMT
            Not After : Dec  7 06:21:56 2017 GMT
        Subject: O=employees@databricks.com, O=eng-cloud-team@datab\
ricks.com, O=eng@databricks.com, O=market-info@databricks.com, O=pl\
atform@databricks.com, O=sfo-all@databricks.com, CN=yunzhang@databr\
icks.com
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (2048 bit)
                Modulus:
```

# CA Trust Chain

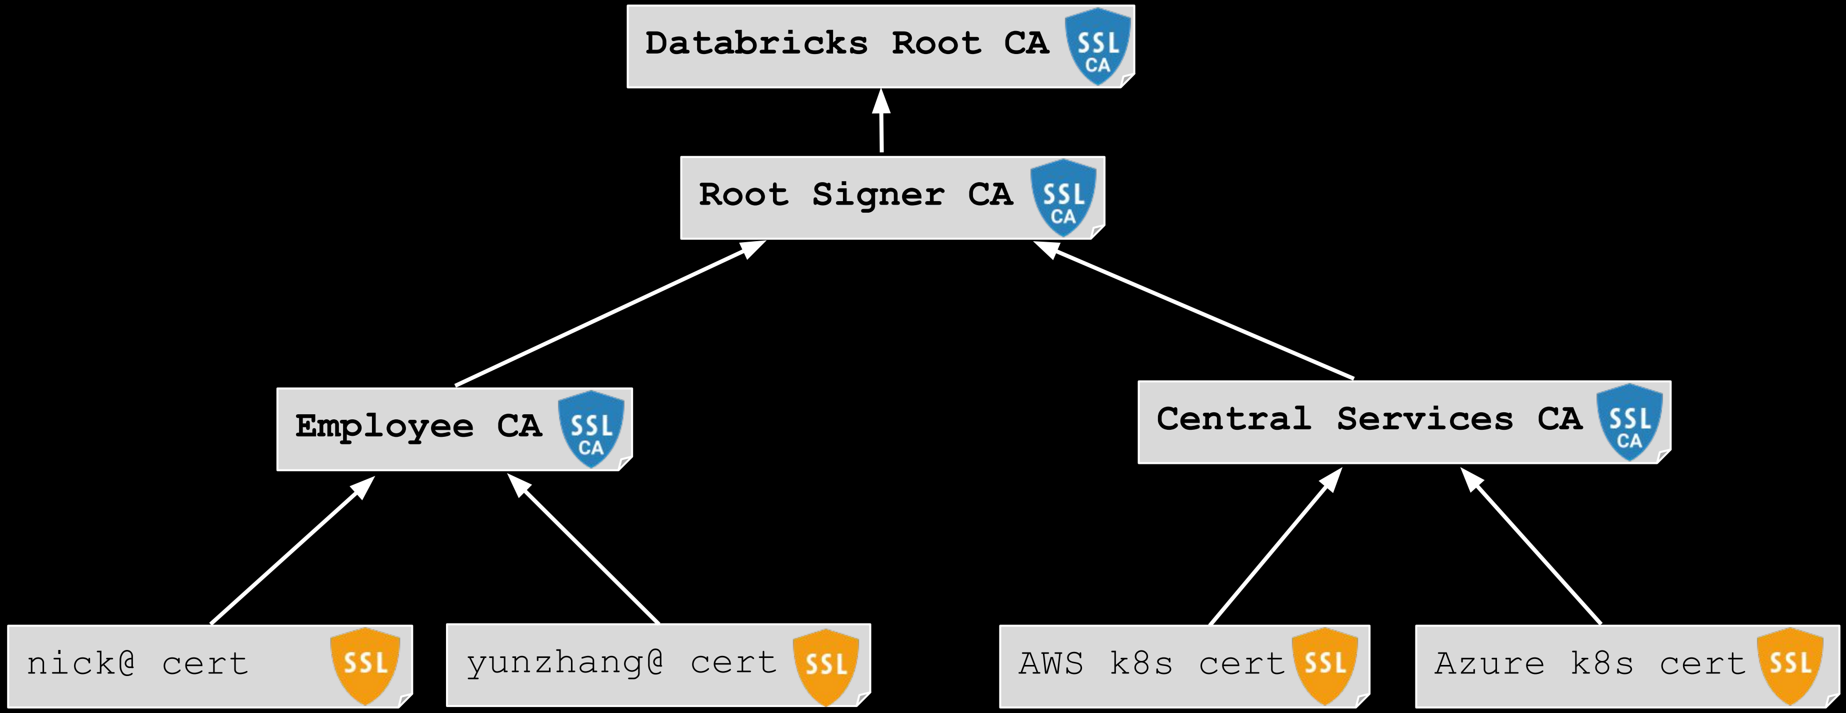

# Kubernetes Auth Strategies

- Client TLS Certificate ✓
- Token-based
- Password-based
- AuthProxy
- ...

Why TLS?

✓ Widely supported

✓ Applicable to multi-services

✓ Easy to control expiration

✓ Easy to integrate with Kubernetes RBAC

# Kubernetes Role Bindings

- Integrated with Google Group
- Access restricted to namespace
- Easy to transfer service ownership

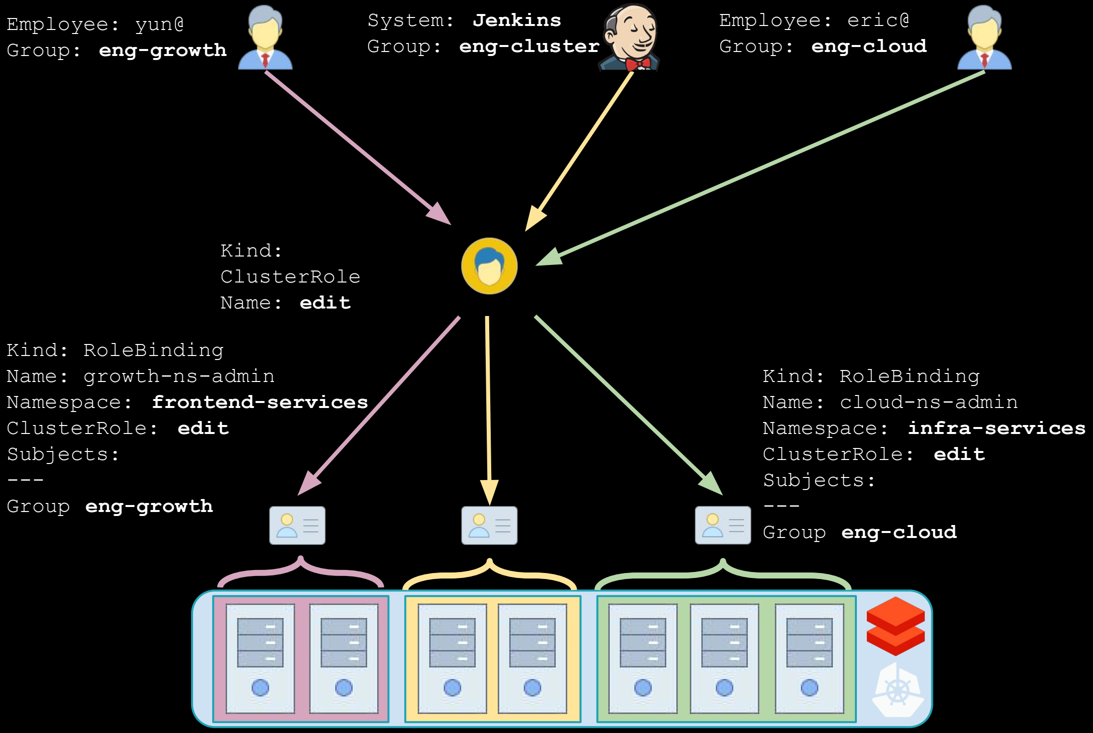

```
Employee: yun@
Group: eng-growth
```

```
System: Jenkins
Group: eng-cluster
```

```
Employee: eric@
Group: eng-cloud
```

```
Kind:
ClusterRole
Name: edit
```

```
Kind: RoleBinding
Name: growth-ns-admin
Namespace: frontend-services
ClusterRole: edit
Subjects:
---
Group eng-growth
```

```
Kind: RoleBinding
Name: cloud-ns-admin
Namespace: infra-services
ClusterRole: edit
Subjects:
---
Group eng-cloud
```

databricks

# Permission Impersonation

- Needed for an adhoc debugging
- Case: Alice from Team A needs to access a service pod belongs to Team B

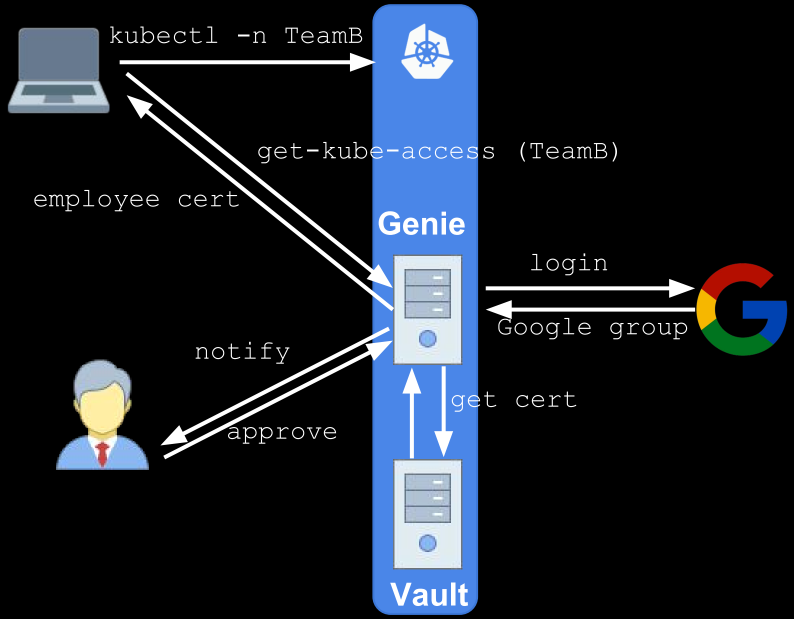

# Continuous Deployment

Jenkins worker authentication

- Talk to “Genie” backend directly
- Uses a long -lived token
- “Genie” Backend generate TLS certificate given the token

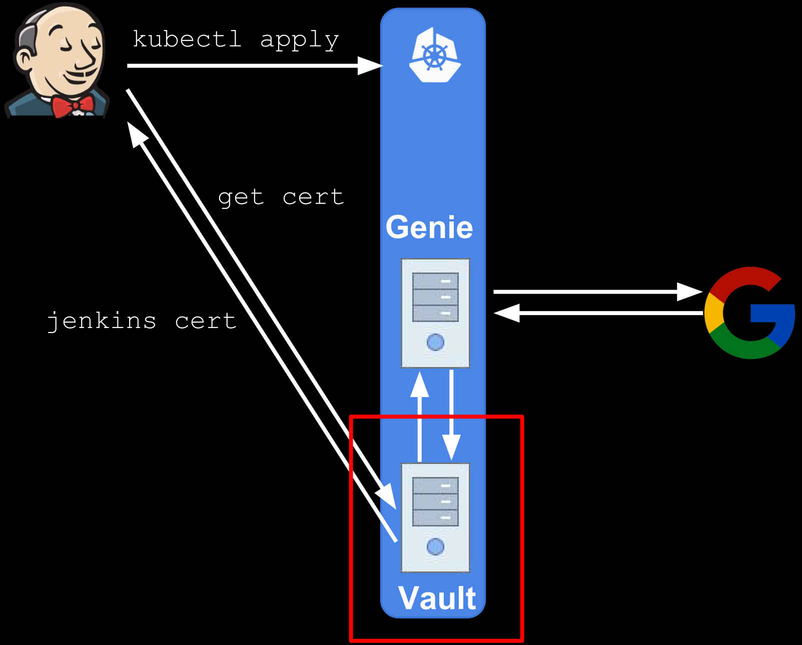

# Agenda

- What is Databricks
- Databricks Security Concerns
- Kubernetes Access Control
- Kubernetes Secret Management
- Kubernetes Audit Logging

# Hashicorp Vault

- Open source secret management solution
- Can generate TLS certs
  - Really useful for short lived (makes it easy)
  - This is how Jenkins talks to our k8s cluster
- Also stores "generic" secret material

```
apiVersion: v1
kind: Secret
metadata:
  annotations:
    name: mysrv-scrtkeys
    namespace: development
data:
  secret-config: TkFVR0hUWQo=
type: Opaque
```

# Secrets in your k8s Deployments

- Deployment Secrets are hard
  - Difficult to audit
  - Difficult to rotate
  - Can't check them into source control

# Secret Templates

```
scrtkeys = { "scrtconfig": "NOTAPASSWORD" }
```

```
apiVersion: v1alpha1
kind: SecretTemplate
metadata:
  name: mysrv-scrtkeys
  namespace: development
context:
  name: dev
  url: https://[url]:8200
```

```
vault_data:
- name: scrtkeys
  path: v1/svcs/mysrv/scrtkeys
spec:
  secret-config:
    format: scrtkeys.scrtconfig
```

# Secret Templates

- Internal tooling
  - pull data from Vault
  - automatically turn into k8s Secret
- Solves previous issues
  - Difficult to audit → Everything in Vault
  - Difficult to rotate → Update in Vault, re-apply
  - Can't check them into source control → Templates contain no secrets

# Agenda

- What is Databricks
- Databricks Security Concerns
- Kubernetes Access Control
- Kubernetes Secret Management
- Kubernetes Audit Logging

# Auditing & Logging

- Genie records details of all types of authorizations it grants
- Enable k8s api-server audit logging

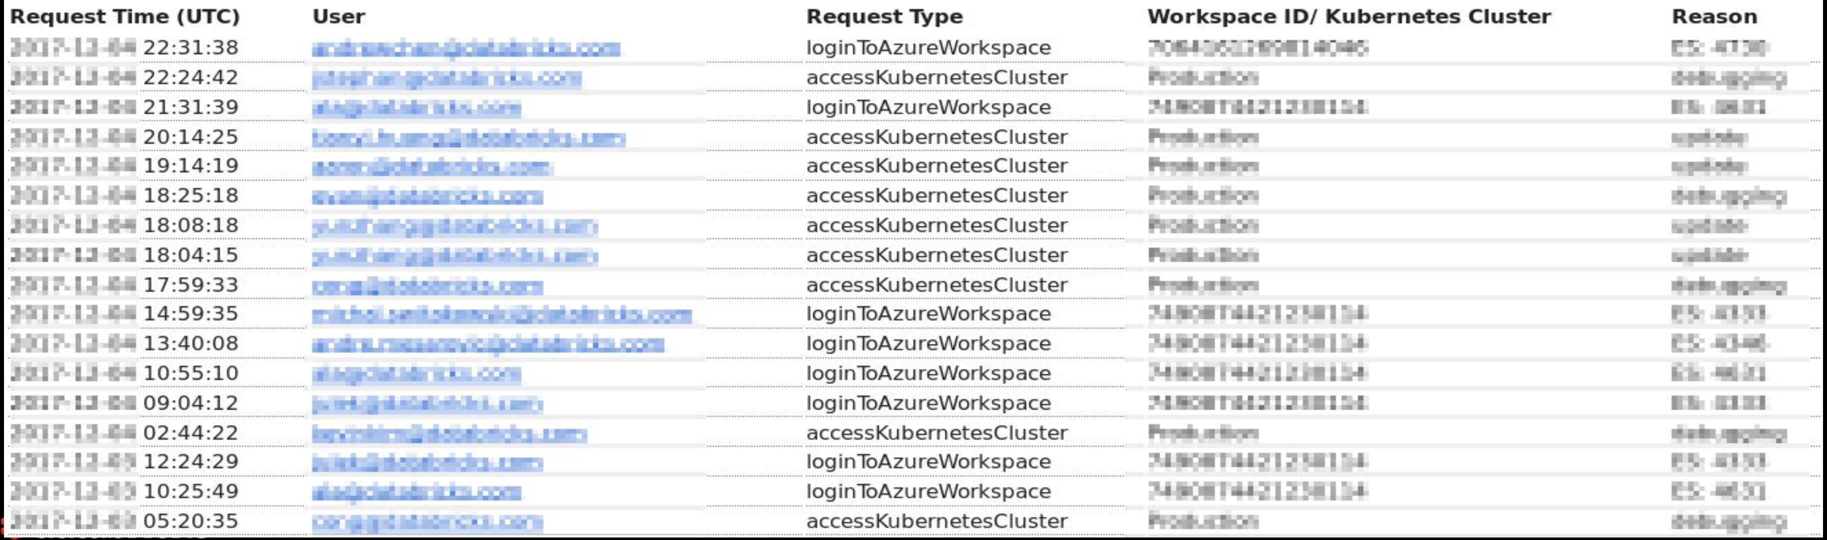

| Request Time (UTC) | User | Request Type | Workspace ID/ Kubernetes Cluster | Reason |
|---|---|---|---|---|
| [redacted] 22:31:38 | [redacted] | loginToAzureWorkspace | [redacted] | [redacted] |
| [redacted] 22:24:42 | [redacted] | accessKubernetesCluster | [redacted] | [redacted] |
| [redacted] 21:31:39 | [redacted] | loginToAzureWorkspace | [redacted] | [redacted] |
| [redacted] 20:14:25 | [redacted] | accessKubernetesCluster | [redacted] | [redacted] |
| [redacted] 19:14:19 | [redacted] | accessKubernetesCluster | [redacted] | [redacted] |
| [redacted] 18:25:18 | [redacted] | accessKubernetesCluster | [redacted] | [redacted] |
| [redacted] 18:08:18 | [redacted] | accessKubernetesCluster | [redacted] | [redacted] |
| [redacted] 18:04:15 | [redacted] | accessKubernetesCluster | [redacted] | [redacted] |
| [redacted] 17:59:33 | [redacted] | accessKubernetesCluster | [redacted] | [redacted] |
| [redacted] 14:59:35 | [redacted] | loginToAzureWorkspace | [redacted] | [redacted] |
| [redacted] 13:40:08 | [redacted] | loginToAzureWorkspace | [redacted] | [redacted] |
| [redacted] 10:55:10 | [redacted] | loginToAzureWorkspace | [redacted] | [redacted] |
| [redacted] 09:04:12 | [redacted] | loginToAzureWorkspace | [redacted] | [redacted] |
| [redacted] 02:44:22 | [redacted] | accessKubernetesCluster | [redacted] | [redacted] |
| [redacted] 12:24:29 | [redacted] | loginToAzureWorkspace | [redacted] | [redacted] |
| [redacted] 10:25:49 | [redacted] | loginToAzureWorkspace | [redacted] | [redacted] |
| [redacted] 05:20:35 | [redacted] | accessKubernetesCluster | [redacted] | [redacted] |

# Take aways

- Security should not be a tax on developer productivity
- A generic solution is usually the more secure one
- Remember everything, until you are allowed to forget
- Don't reinvent the wheel -- OSS is your friend

# Thank you!

ericwang@databricks.com
yunzhang@databricks.com