Karen Jex | Senior solutions Architect
PGConf.EU 2024 | Athens

Crunchy Postgres for Kubernetes: Your Virtual DBA

Introduction

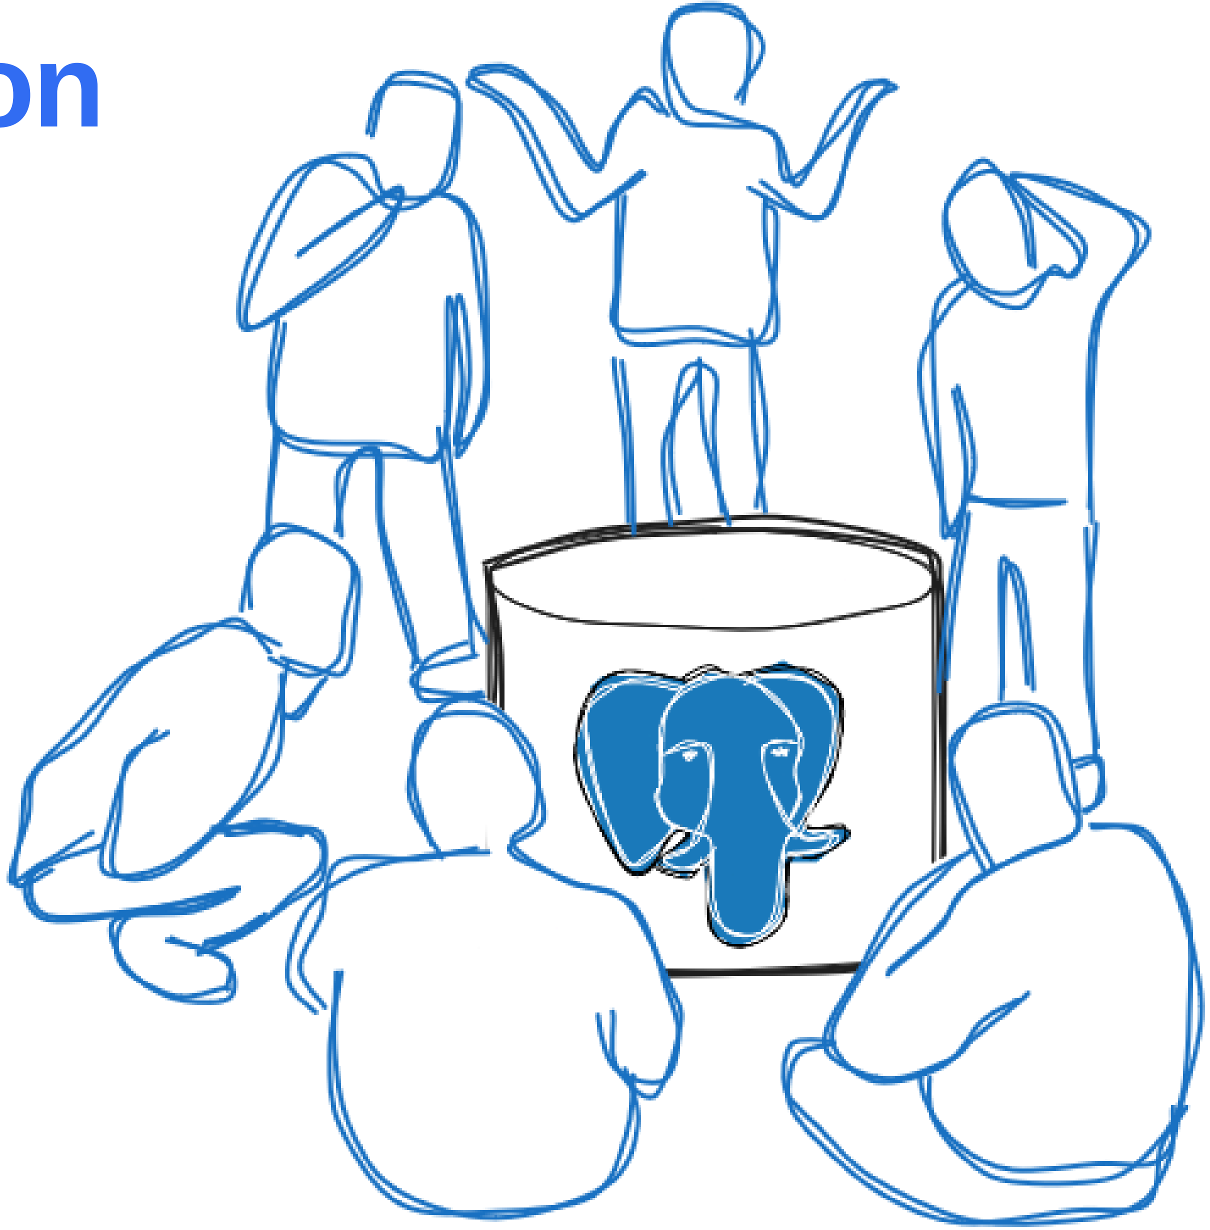

whoami

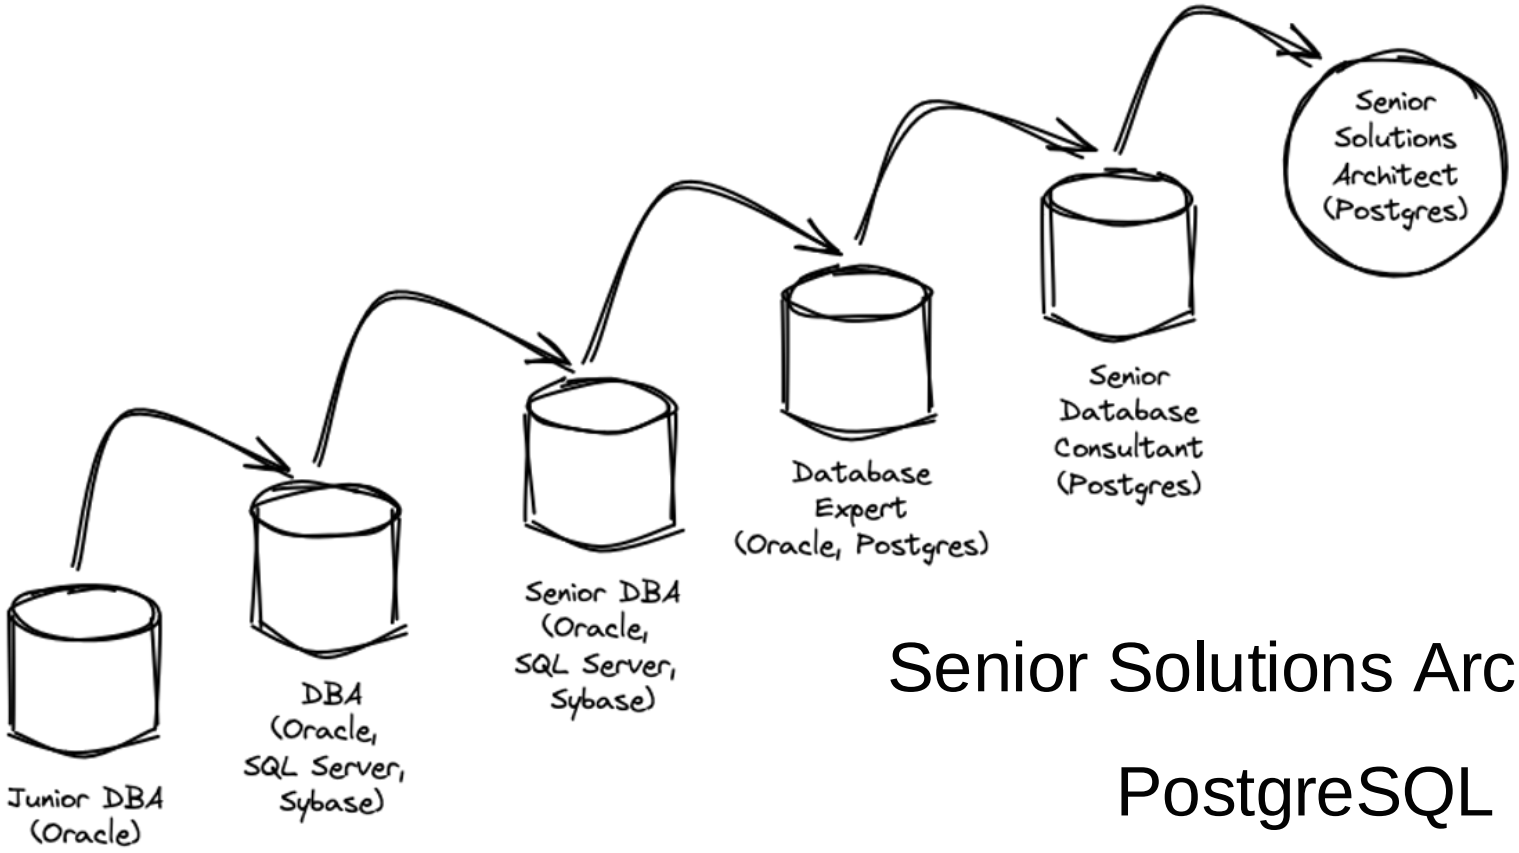

Senior Solutions Architect @ Crunchy Data

PostgreSQL Europe Board Member

PostgreSQL Europe Diversity Committee Chair

crunchydata

Postgres Anywhere

BARE METAL, VMS, CLOUD

Crunchy Postgres

Crunchy Certified PostgreSQL is production ready Postgres.

INCLUDES:

- ✓ Backups
- ✓ Disaster recovery
- ✓ High availability
- ✓ Monitoring
- ✓ Automation
- ✓ Self managed

KUBERNETES

Crunchy Postgres for Kubernetes

Cloud Native Postgres on Kubernetes powered by Crunchy Postgres Operator.

INCLUDES:

- ✓ Simple provisioning
- ✓ Backups and DR included
- ✓ High availability
- ✓ Seamless upgrades
- ✓ Scale from 1 to 1000s of DBs
- ✓ Self managed

FULLY MANAGED CLOUD

Crunchy Bridge

The fully managed Postgres option on your choice of Cloud provider.

INCLUDES:

- ✓ AWS, Azure or GCP
- ✓ Continuous protection
- ✓ Backups included
- ✓ Point in Time Recovery
- ✓ Pay for what you use
- ✓ The developer experience you want

crunchy data

Agenda

1. Database Architecture History
2. Container Orchestration/Kubernetes Overview
3. Kubernetes Superpowers
4. What does a DBA do anyway?
5. Postgres Operators for Kubernetes
6. What do you want from an Operator?
7. What does CPK do for you?
8. Getting Started with CPK

Agenda

1. Database Architecture History
2. Container Orchestration/Kubernetes Overview
3. Kubernetes Superpowers
4. What does a DBA do anyway?
5. Postgres Operators for Kubernetes
6. What do you want from an Operator?
7. What does CPK do for you?
8. Getting Started with CPK

crunchy data

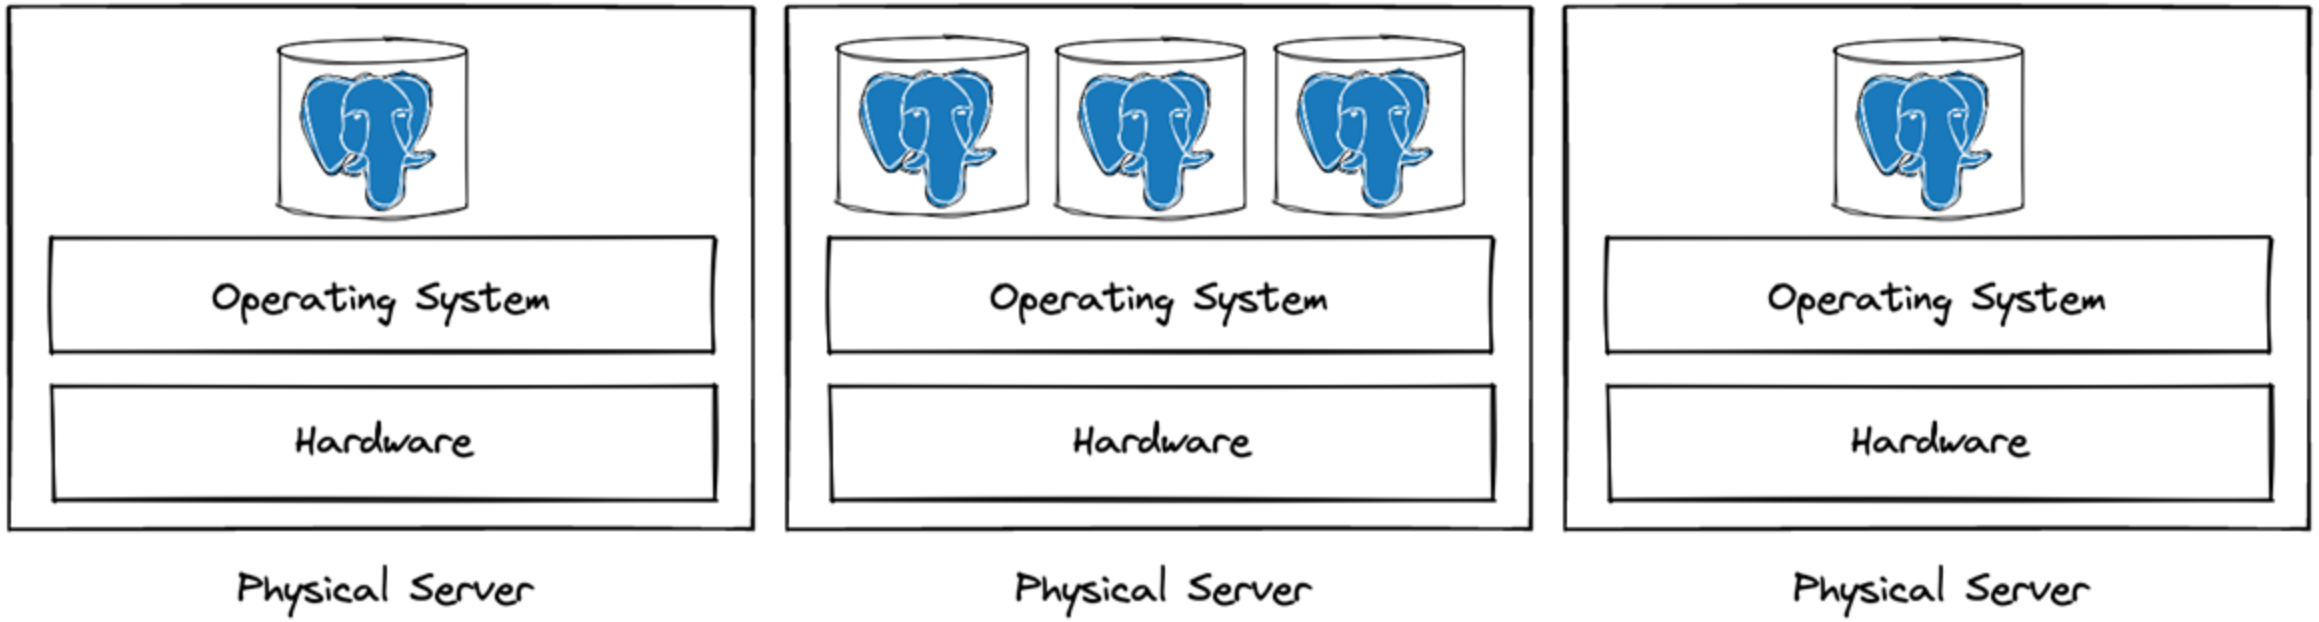

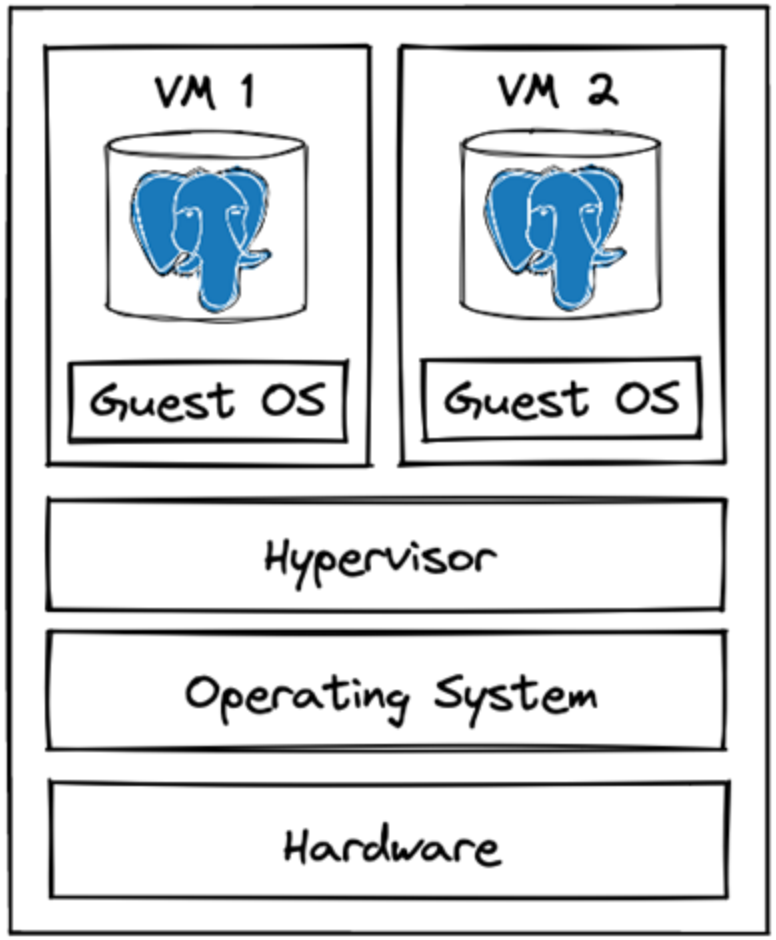

Virtual Machines

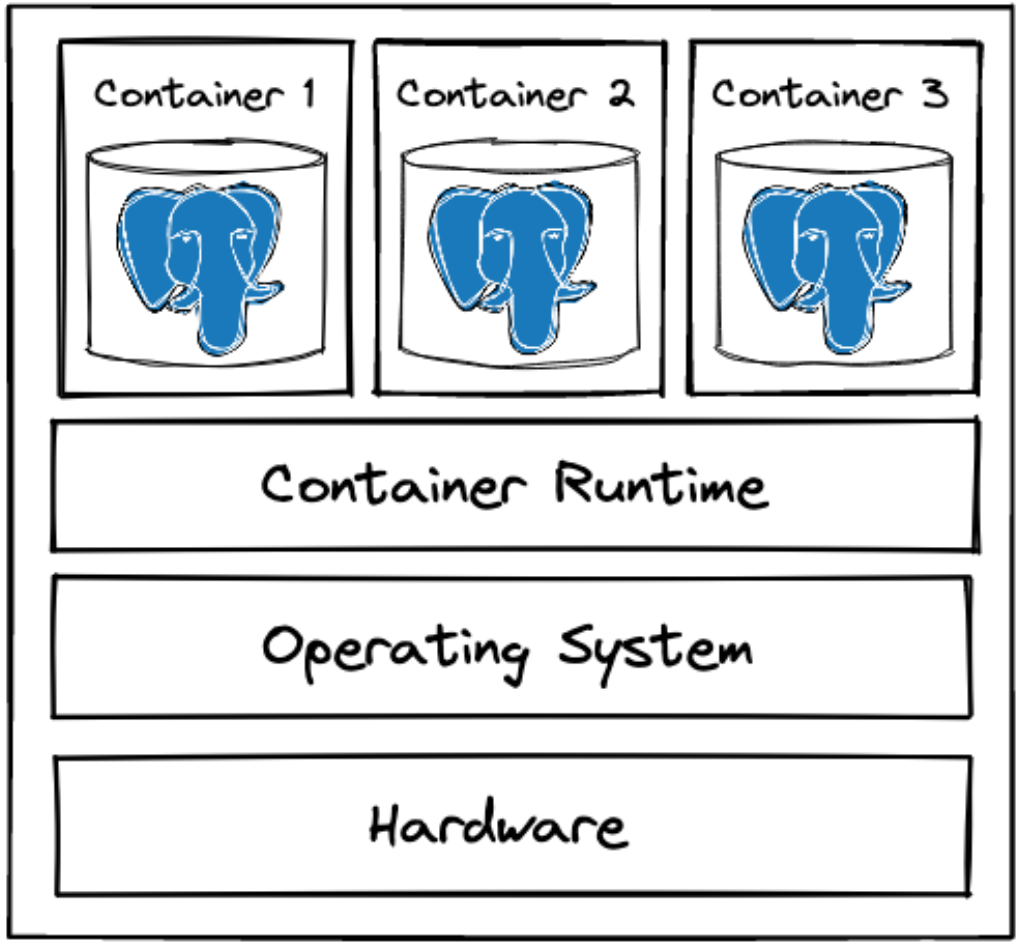

Containers

Agenda

1. Database Architecture History
2. Container Orchestration/Kubernetes Overview
3. Kubernetes Superpowers
4. What does a DBA do anyway?
5. Postgres Operators for Kubernetes
6. What do you want from an Operator?
7. What does CPK do for you?
8. Getting Started with CPK

Features of Containers

- Lightweight
- Scalable
- Portable
- Isolated

- Stateless
- Ephemeral

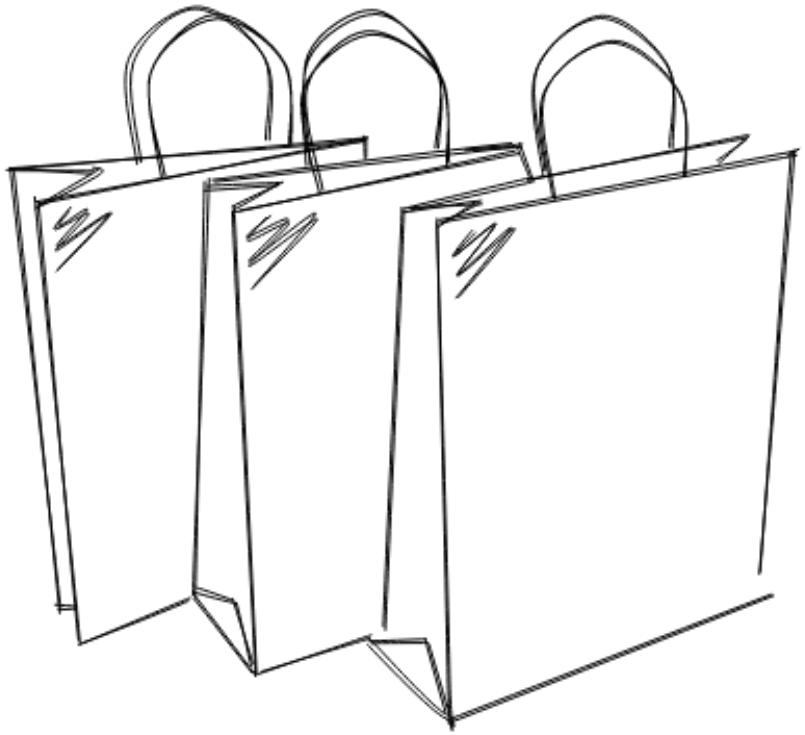

crunchydata

Herding Cats (Managing Containers)

Container Orchestration

- Tools such as Kubernetes
- Manage many containers
- Automate container lifecycle
- Integrate with DevOps tools

Container Orchestration Tasks

- Provisioning
- Deployment
- Configuration
- Scheduling
- Scaling up and down
- Self-healing
- Services
- Storage
- Resource allocation
- Load Balancing
- Networking
- Security

What is Kubernetes?

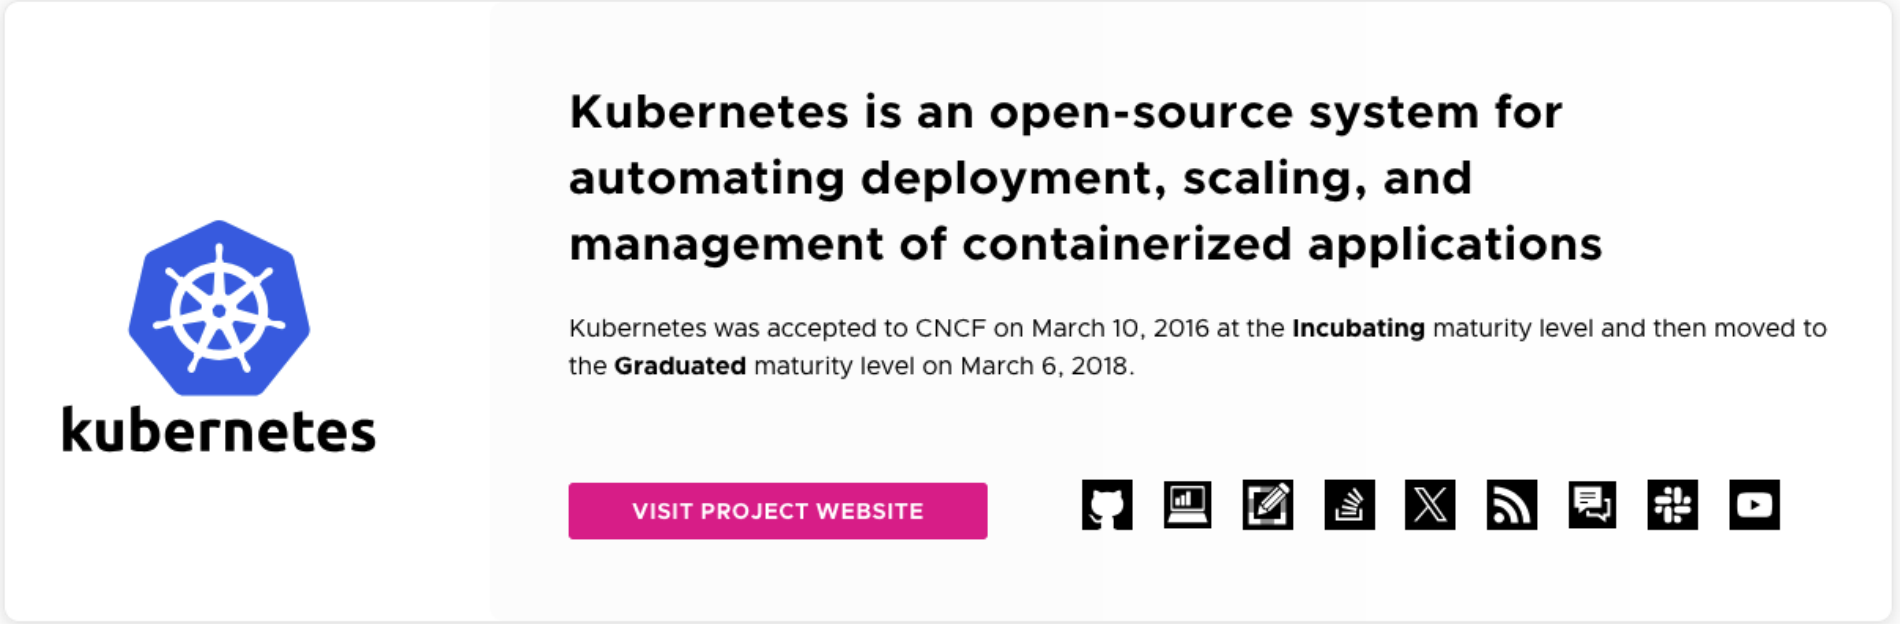

https://www.cncf.io/projects/kubernetes/

Why Postgres on Kubernetes?

- Automation
- Deploying at Scale
- Multitenancy
- Microservices
- Kubernetes already in use
- DBaaS

Agenda

1. Database Architecture History
2. Container Orchestration/Kubernetes Overview
3. Kubernetes Superpowers
4. What does a DBA do anyway?
5. Postgres Operators for Kubernetes
6. What do you want from an Operator?
7. What does CPK do for you?
8. Getting Started with CPK

Pods & Deployments

Pod contains 1 or more containers

Deployment consists of 1 or more replica pods

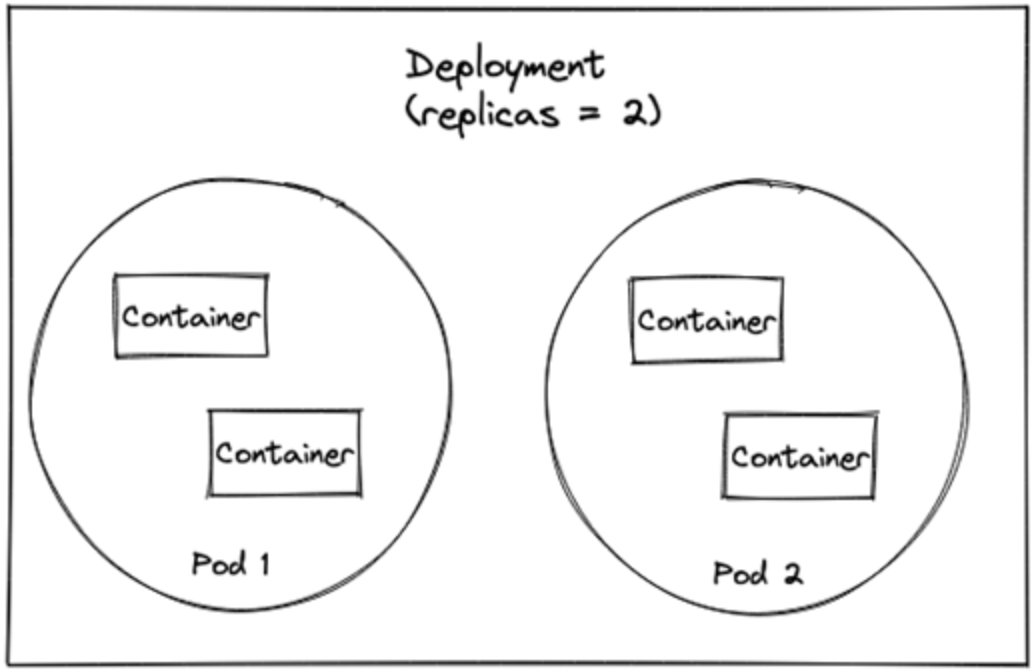

crunchydata

Persistent Volumes (PVs)

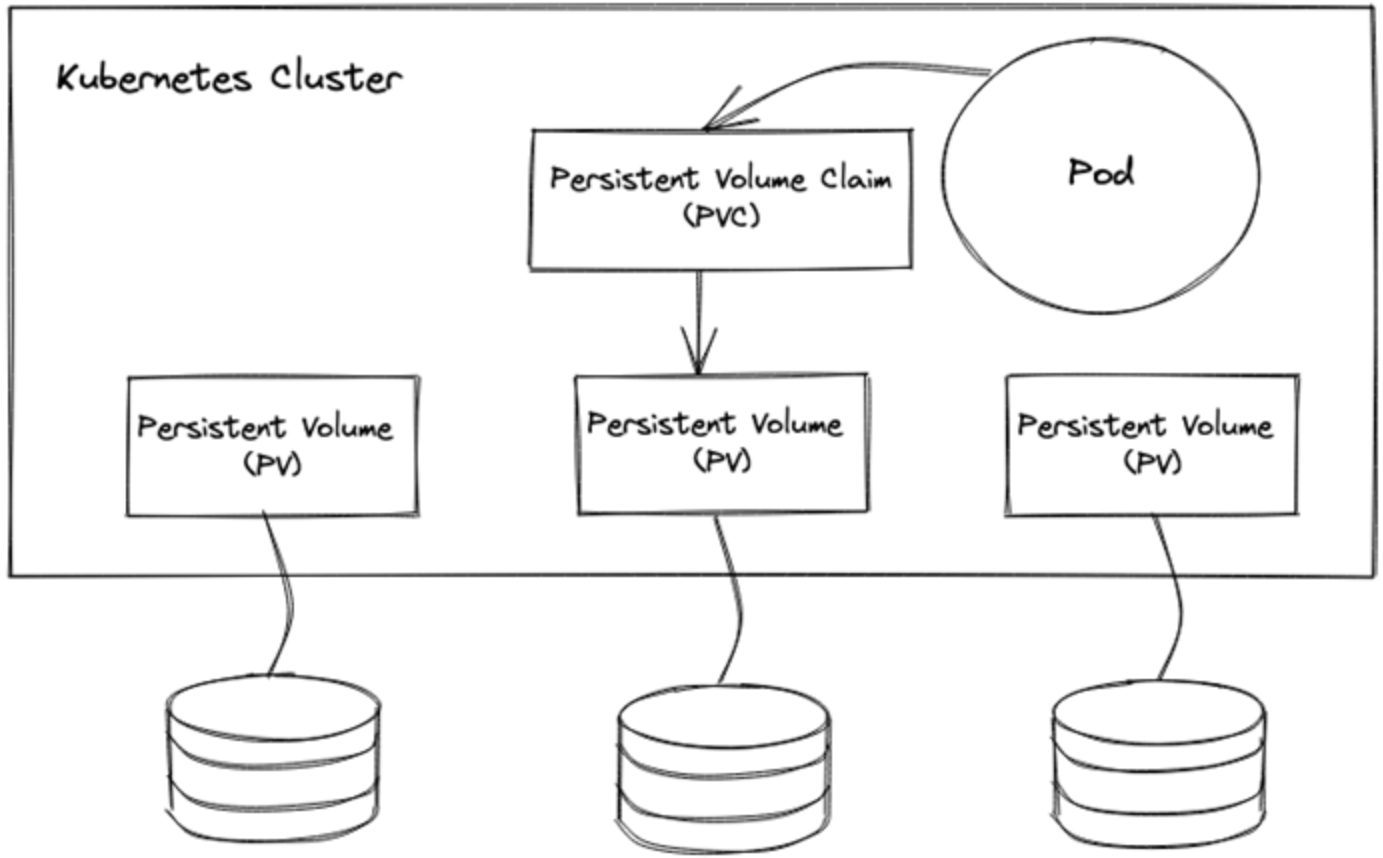

What about Standby Databases?

- Pods in a deployment are interchangeable
- What about a HA environment?
- Primary and standby databases aren’t interchangeable
- Need ordered startup/shutdown

StatefulSets

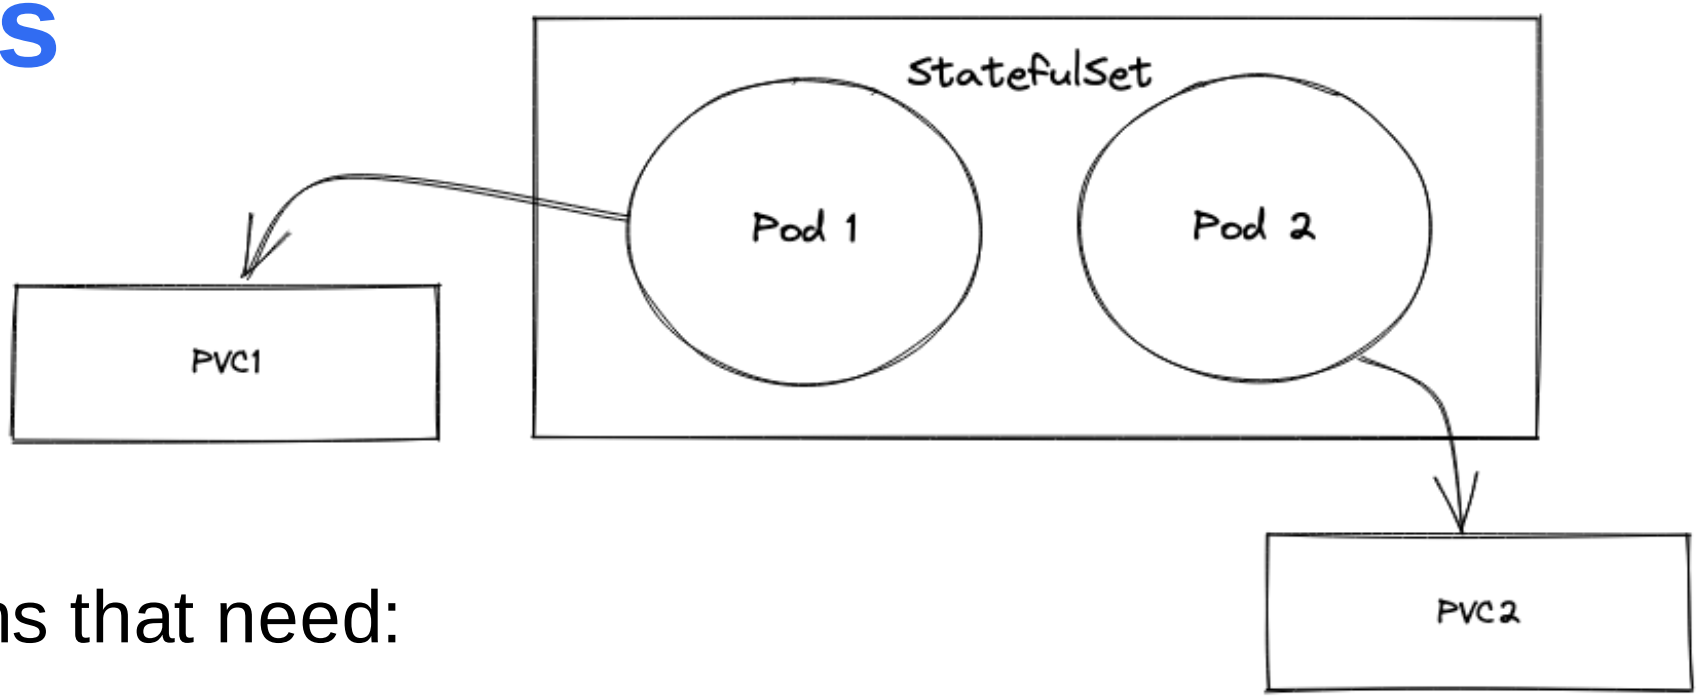

Useful for applications that need:

- Stable, persistent storage
- Ordered, graceful deployment and scaling
- Ordered, automated rolling updates

Sidecars

- “Helper” container
- Tightly coupled with main container
- eg: metrics exporter, backup tool

Agenda

1. Database Architecture History
2. Container Orchestration/Kubernetes Overview
3. Kubernetes Superpowers
4. What does a DBA do anyway?
5. Postgres Operators for Kubernetes
6. What do you want from an Operator?
7. What does CPK do for you?
8. Getting Started with CPK

What is a DBA?

"Database administrators create, organise and look after computer systems that store data for a company." https://nationalcareers.service.gov.uk/job-profiles/database-administrator

"A database administrator, or DBA, is responsible for maintaining, securing, and operating databases and also ensures that data is correctly stored and retrieved." https://www.oracle.com/database/what-is-a-dba/

"Database administrators use specialist software to organise and maintain a secure database."
https://www.prospects.ac.uk/job-profiles/database-administrator

"A Database Administrator, or Database Manager is responsible for managing computer systems that store and organise data for companies. Their duties include… securing data and identifying areas for improvement with the infrastructure."
https://www.indeed.com/career/database-administrator

crunchydata

DBA Responsibilities

- Database (High) Availability
- Backup and Recovery
- Security & Data Protection
- Monitoring
- DB Design/Data Modelling
- Support/Troubleshooting
- DB Software install/upgrade
- Database Expertise
- Performance Tuning
- Capacity Planning
- Database Creation
- Database Maintenance

Agenda

1. Database Architecture History
2. Container Orchestration/Kubernetes Overview
3. Kubernetes Superpowers
4. What does a DBA do anyway?
5. Postgres Operators for Kubernetes
6. What do you want from an Operator?
7. What does CPK do for you?
8. Getting Started with CPK

crunchy data

Running Postgres on Kubernetes

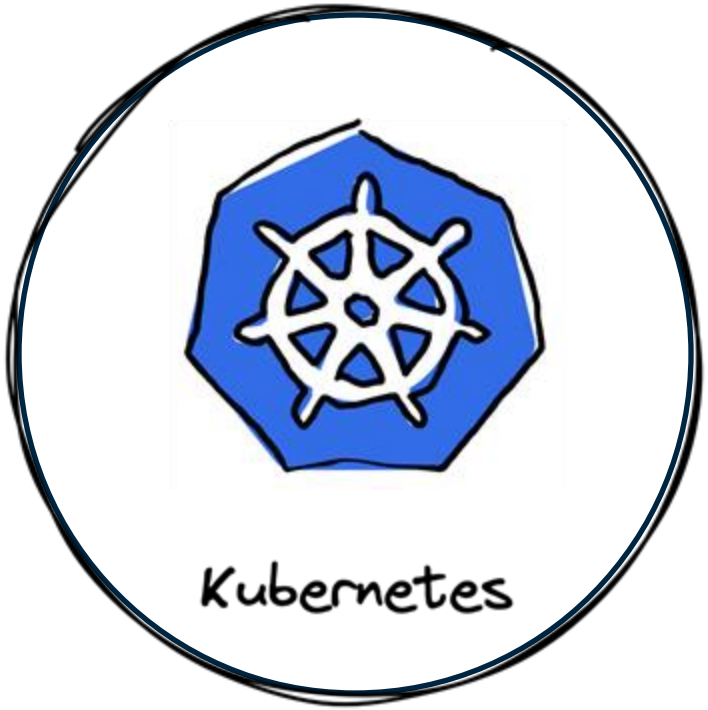

crunchydata

Kubernetes Operators

“Operators are software extensions to Kubernetes that make use of **custom resources** to manage applications and their components. Operators follow Kubernetes principles, notably **the control loop**.”

https://kubernetes.io/docs/concepts/extend-kubernetes/operator/

Postgres Operators for Kubernetes

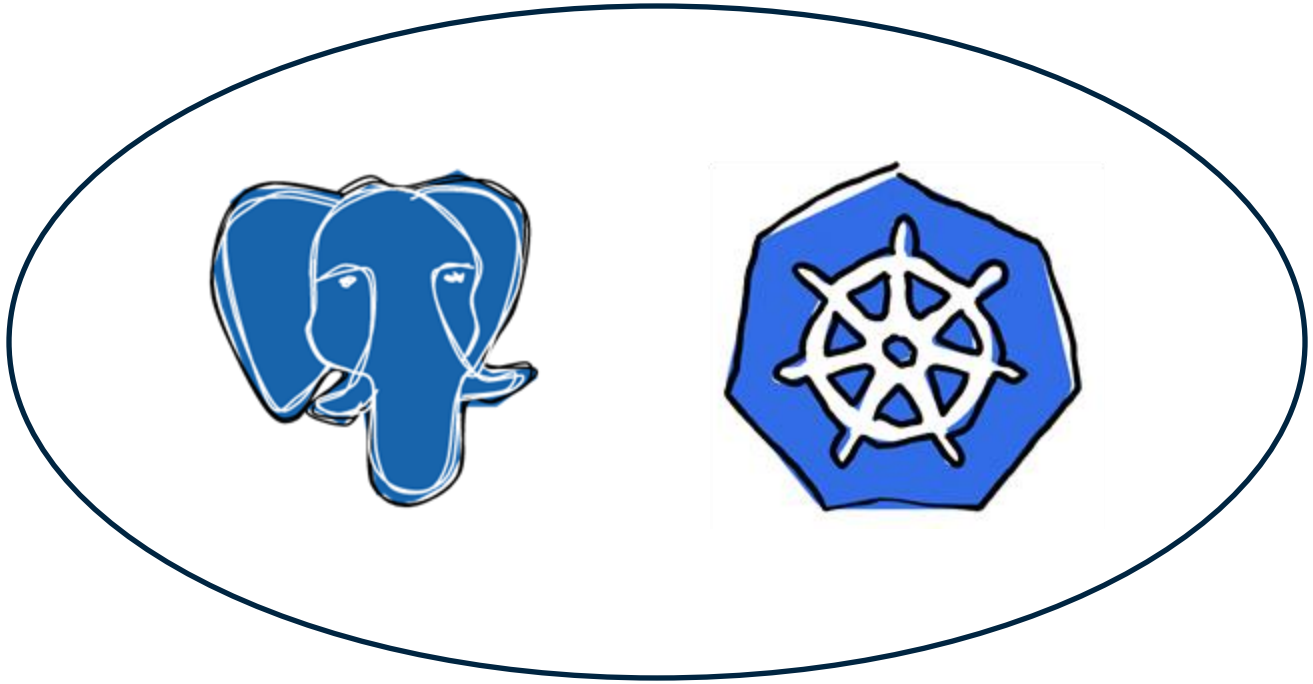

crunchydata

Agenda

1. Database Architecture History
2. Container Orchestration/Kubernetes Overview
3. Kubernetes Superpowers
4. What does a DBA do anyway?
5. Postgres Operators for Kubernetes
6. What do you want from an Operator?
7. What does CPK do for you?
8. Getting Started with CPK

~~DBA~~ Operator Responsibilities

- Database (High) Availability
- Backup and Recovery
- Security & Data Protection
- Monitoring
- ~~DB Design/Data Modelling~~
- Support/Troubleshooting
- DB Software install/upgrade
- Database Expertise
- Performance Tuning
- Capacity Planning
- Database Creation
- Database Maintenance

Agenda

1. Database Architecture History
2. Container Orchestration/Kubernetes Overview
3. Kubernetes Superpowers
4. What does a DBA do anyway?
5. Postgres Operators for Kubernetes
6. What do you want from an Operator?
7. What does CPK do for you?
8. Getting Started with CPK

crunchy data

What does CPK do for you?

- Database (High) Availability
- Backup and Recovery
- Security & Data Protection
- Monitoring
- ~~DB Design/Data Modelling~~
- Support/Troubleshooting
- DB Software install/upgrade
- Database Expertise
- Performance Tuning
- Capacity Planning
- Database Creation
- Database Maintenance

crunchy data

CPK:
High Availability

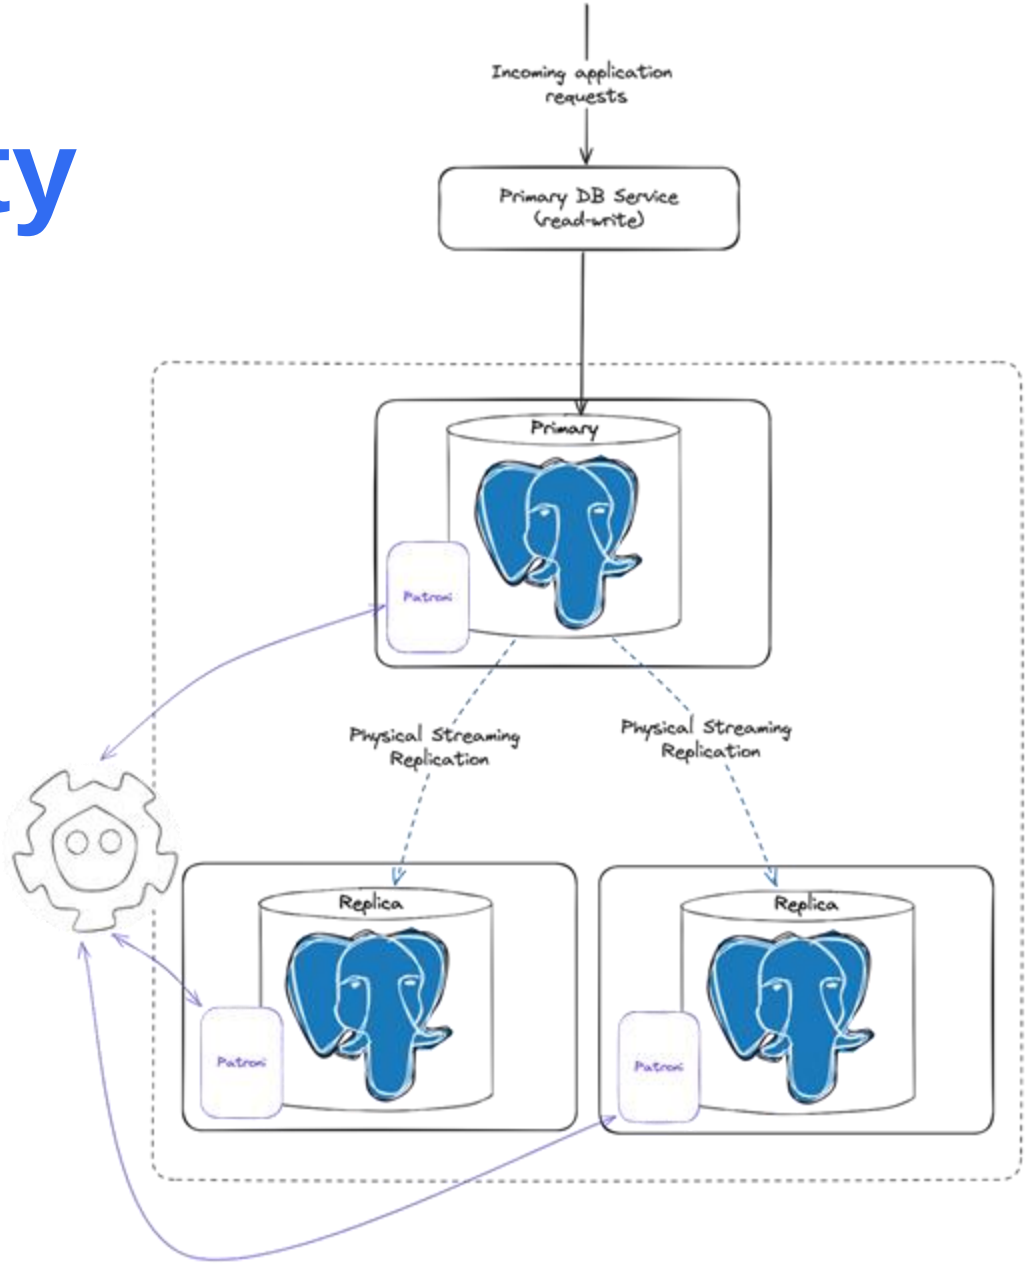

crunchydata

CPK:
High Availability

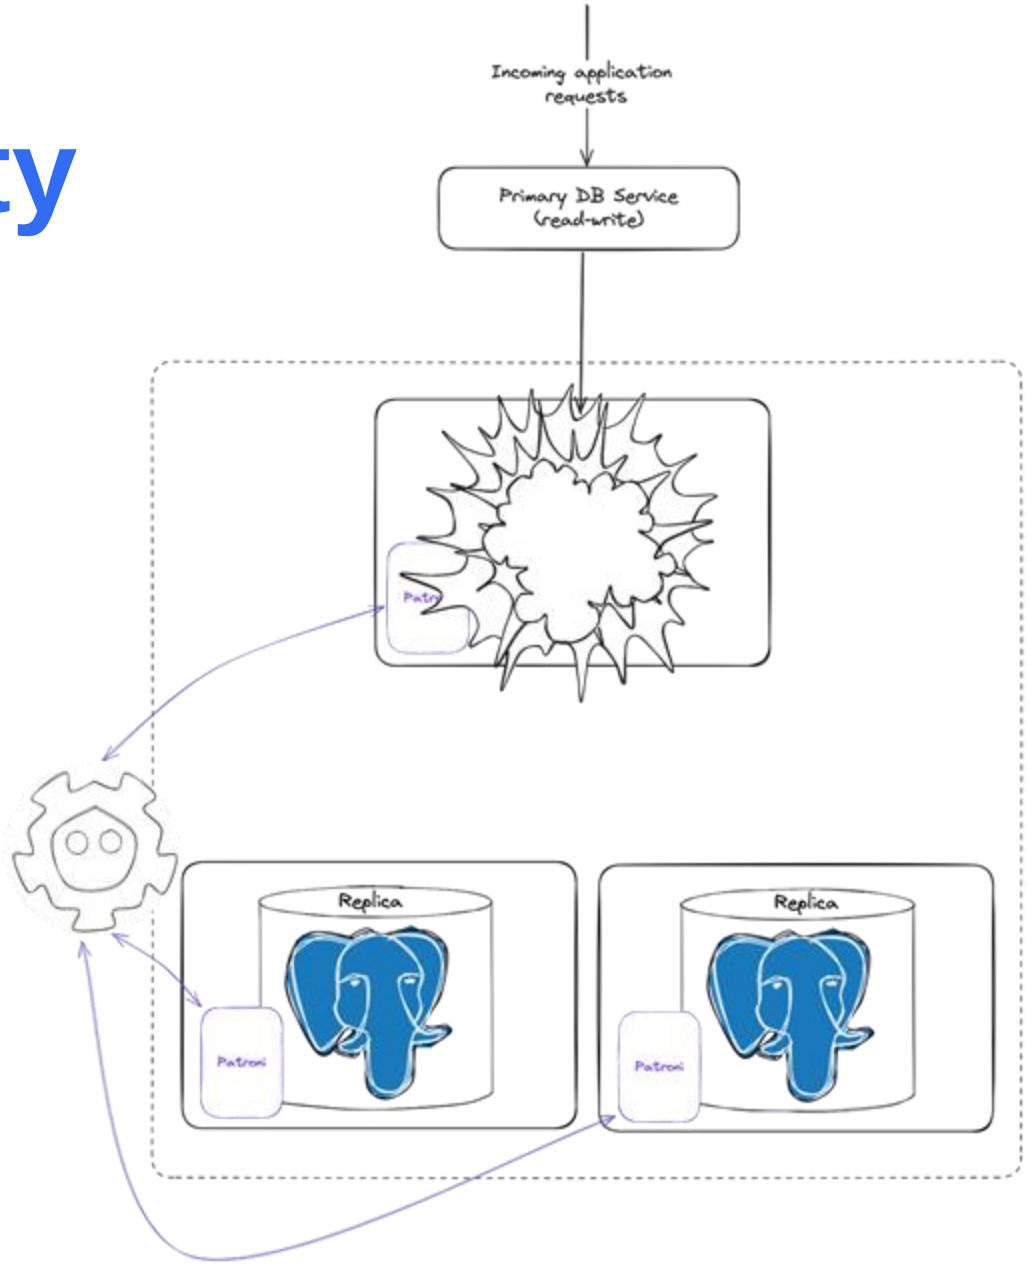

crunchydata

CPK:
High Availability

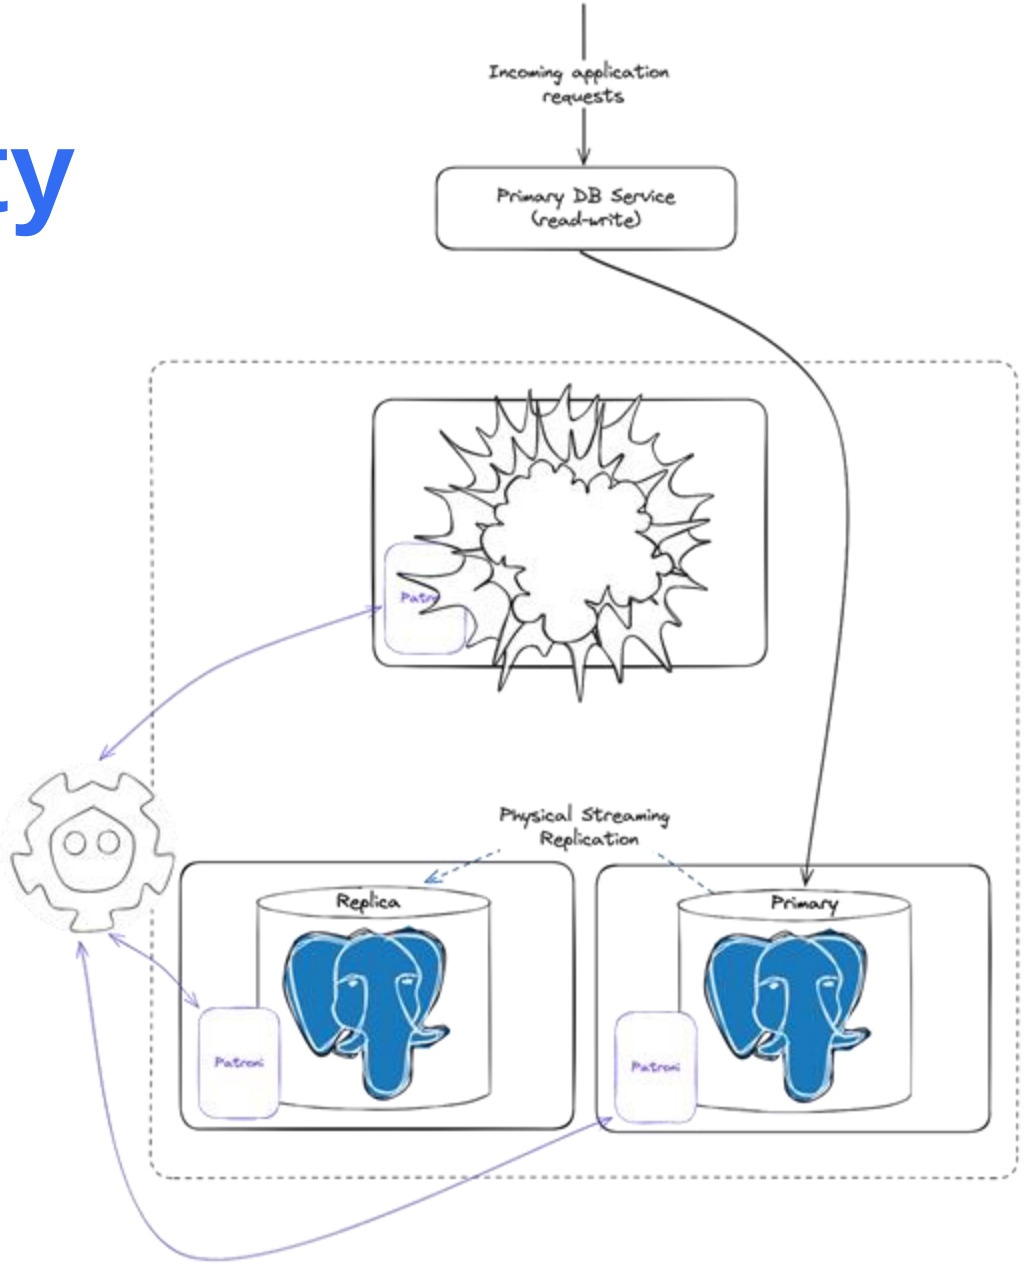

crunchydata

CPK:
High Availability

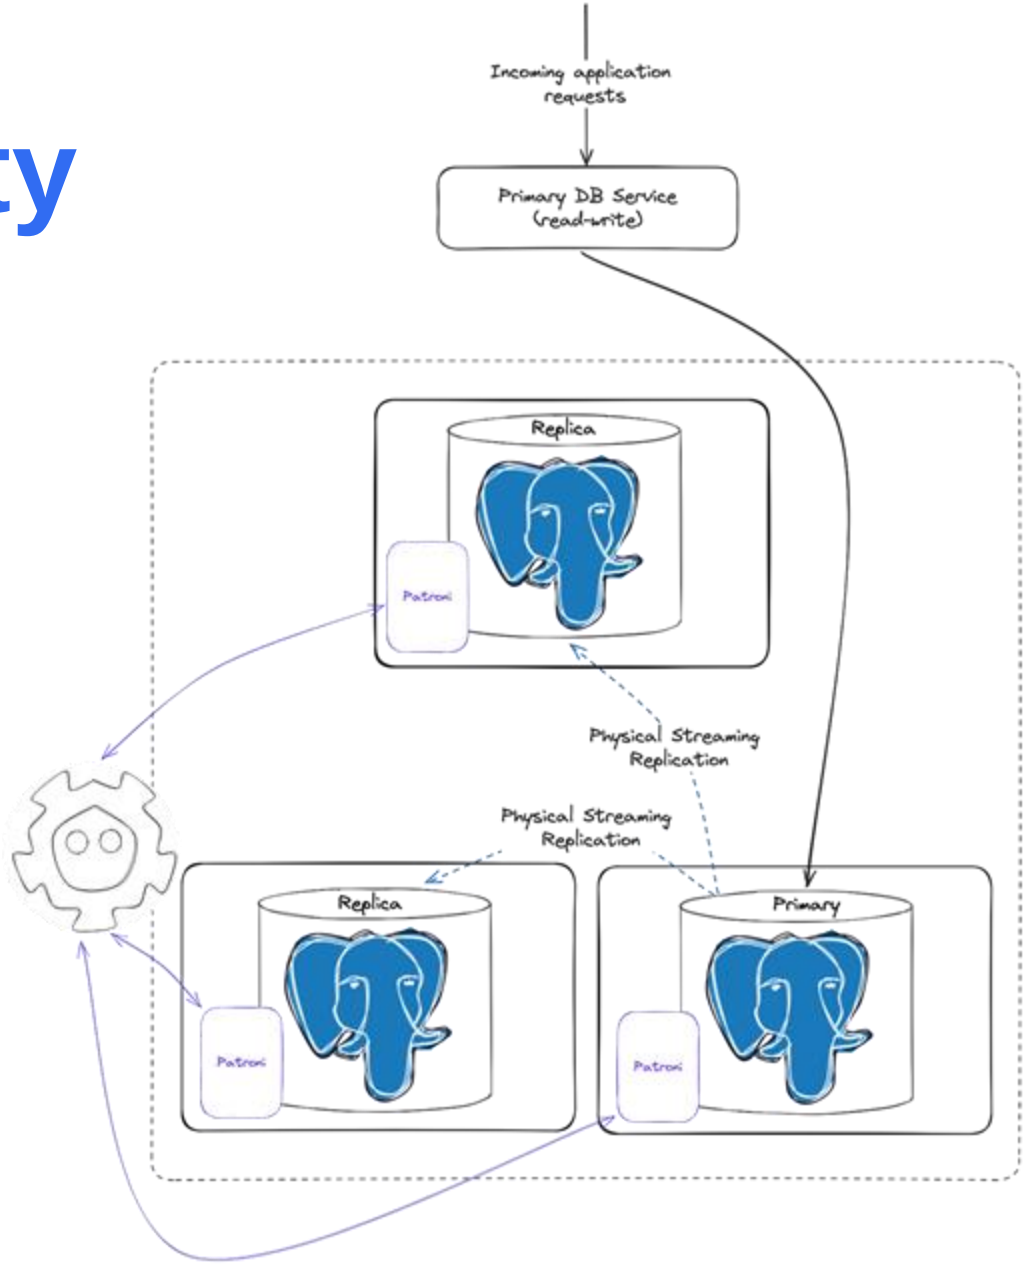

crunchydata

CPK:
High Availability

- Kubernetes
- Operator (PostgresCluster)
- Patroni

CPK:

Backup & Recovery

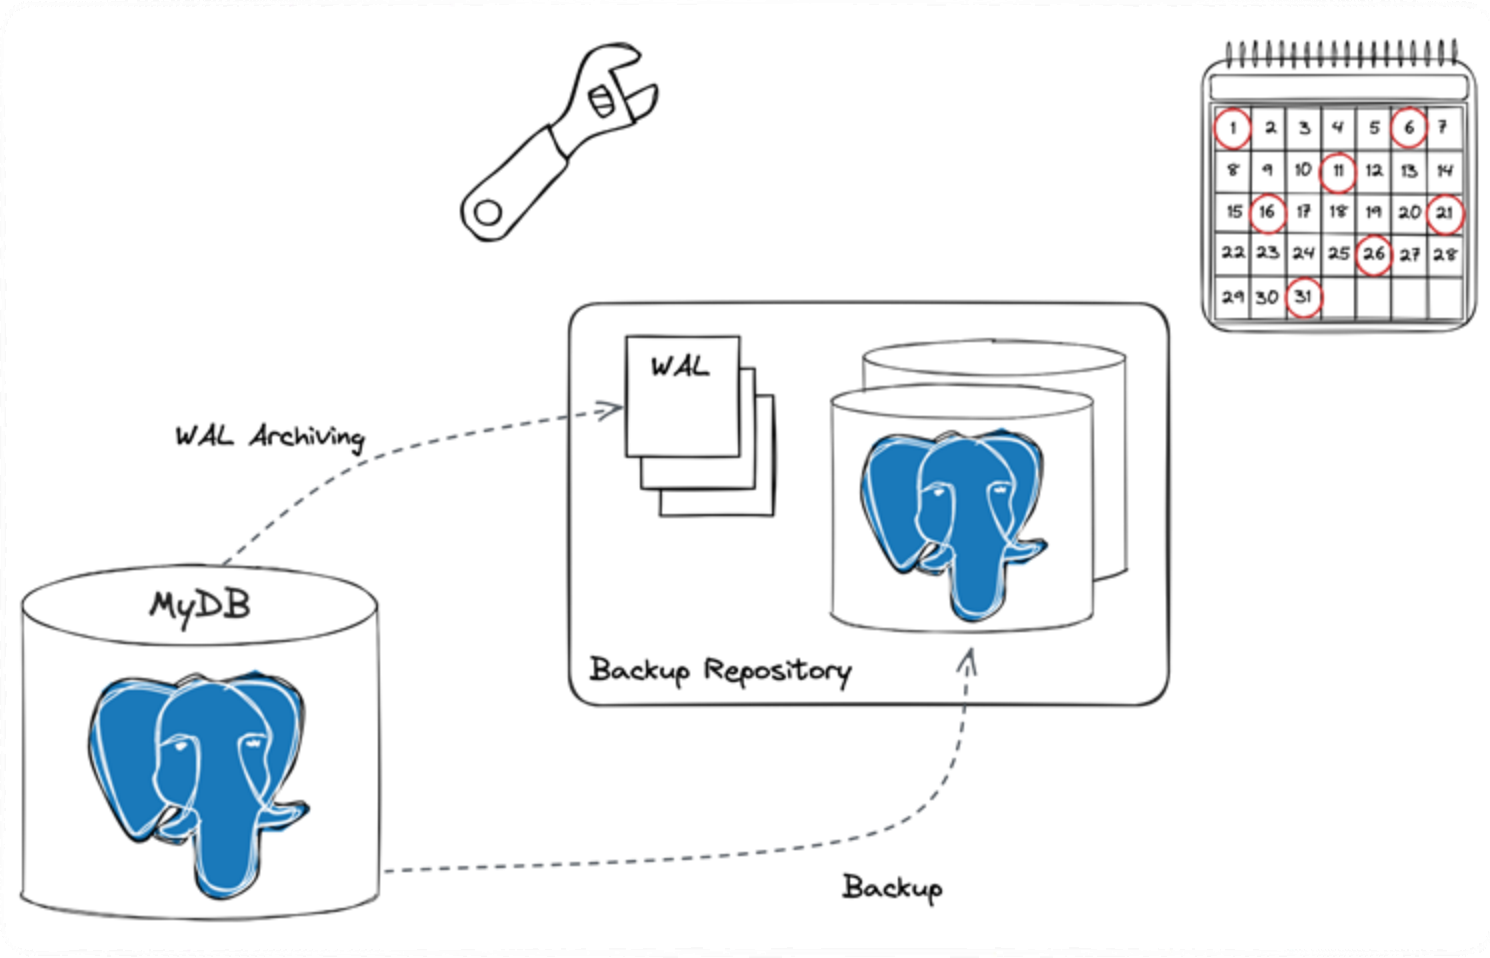

crunchydata

CPK:
Backup & Recovery

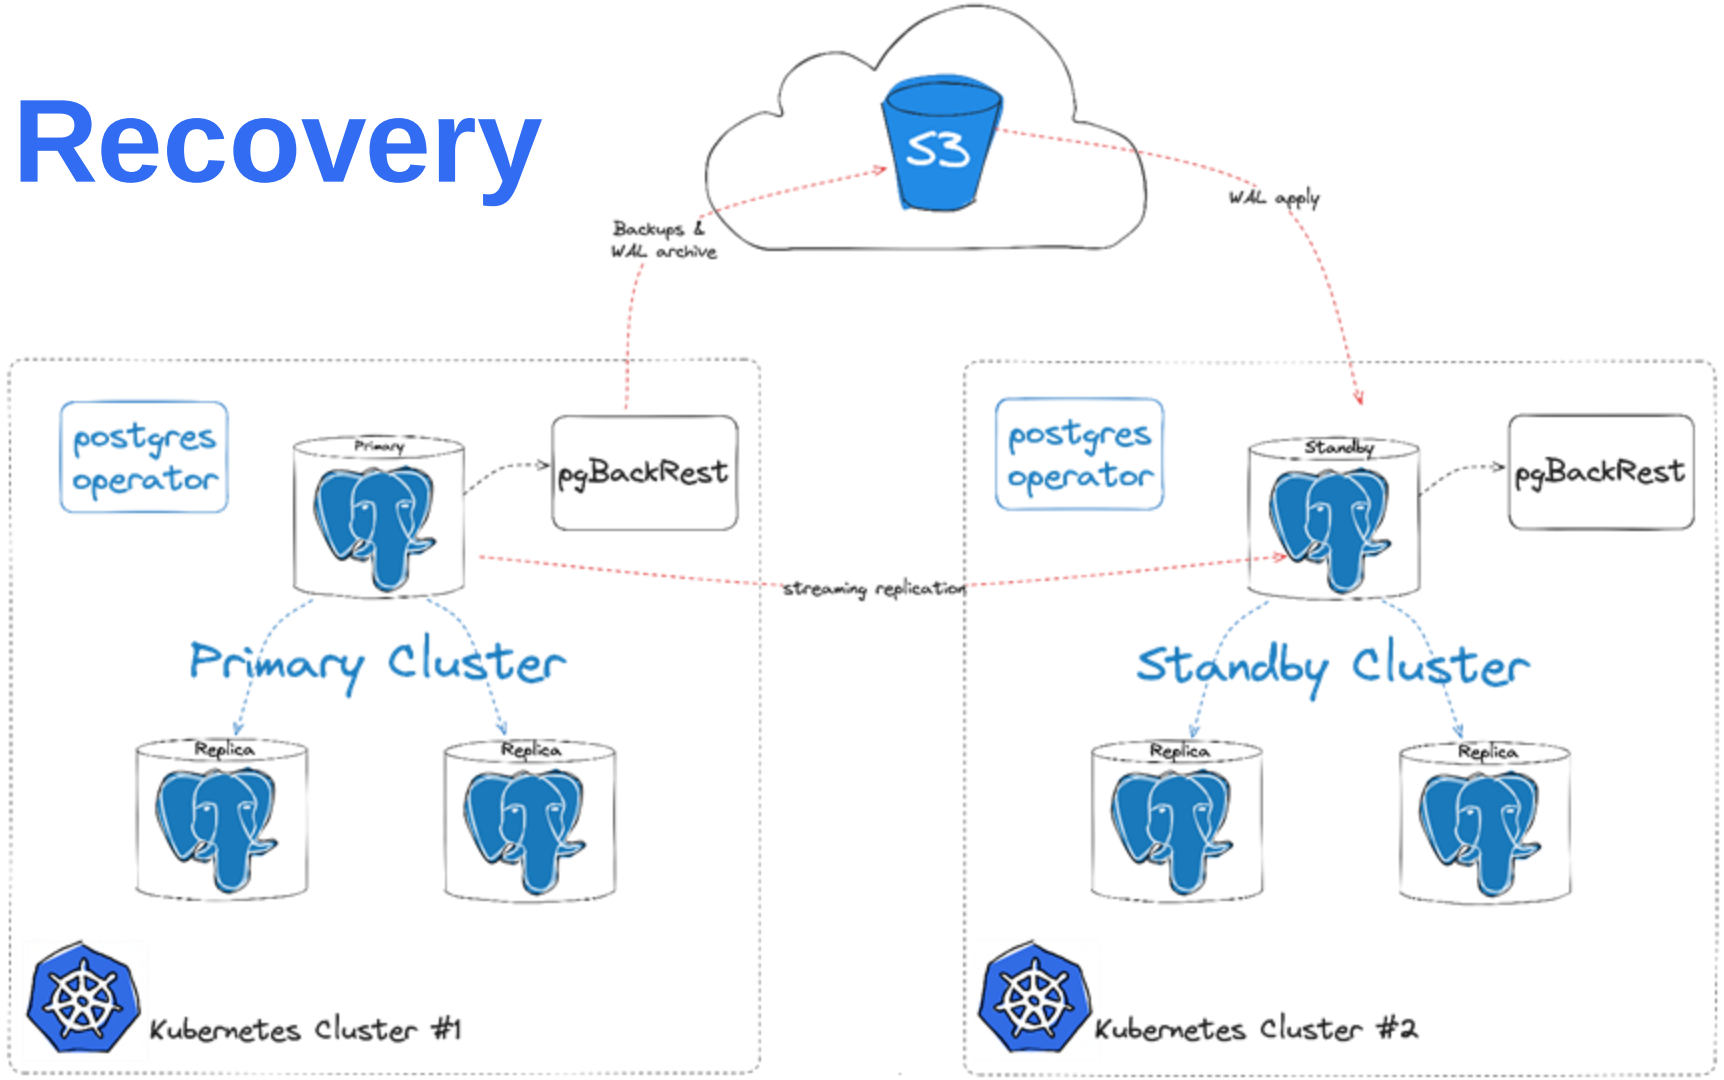

crunchydata

CPK: Security

- Database Access
- Password Encryption
- Secrets
- SSL/TLS
- Certificates

crunchy data

CPK:
Monitoring

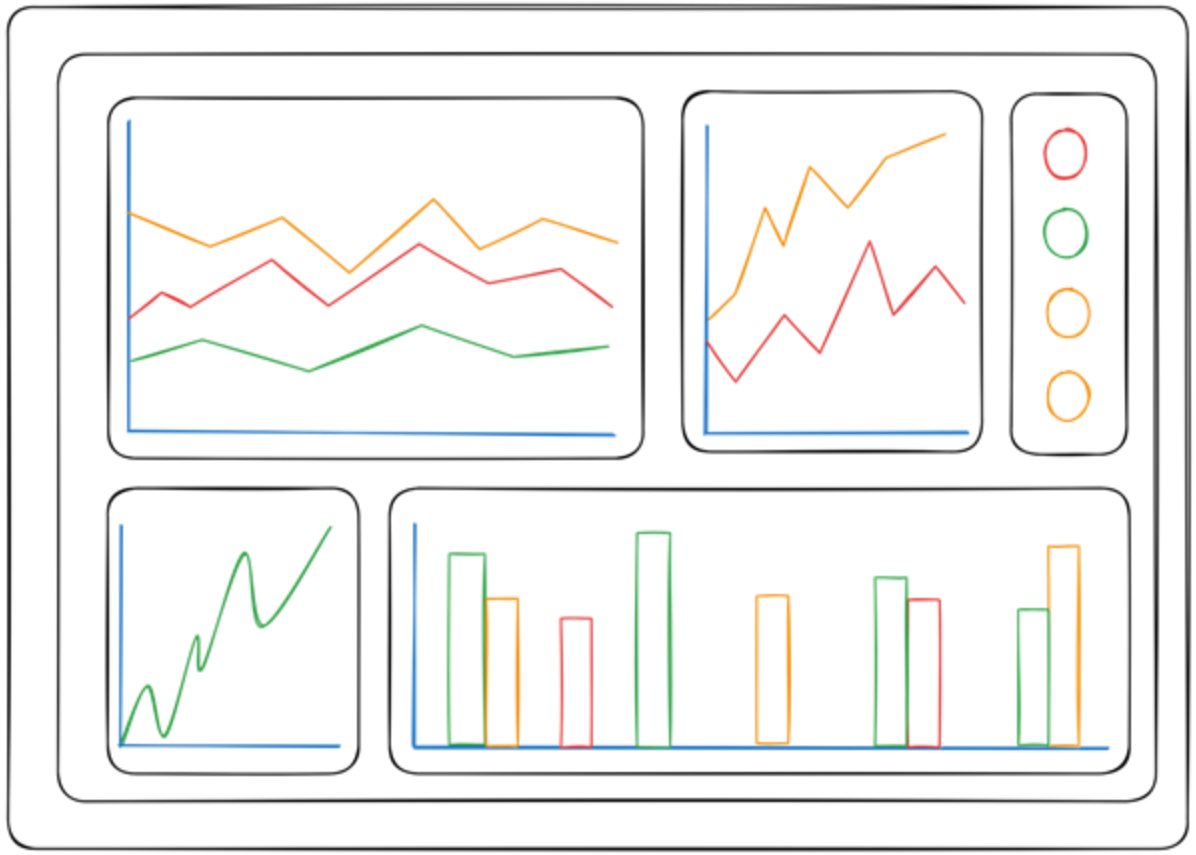

crunchydata

CPK:
Monitoring

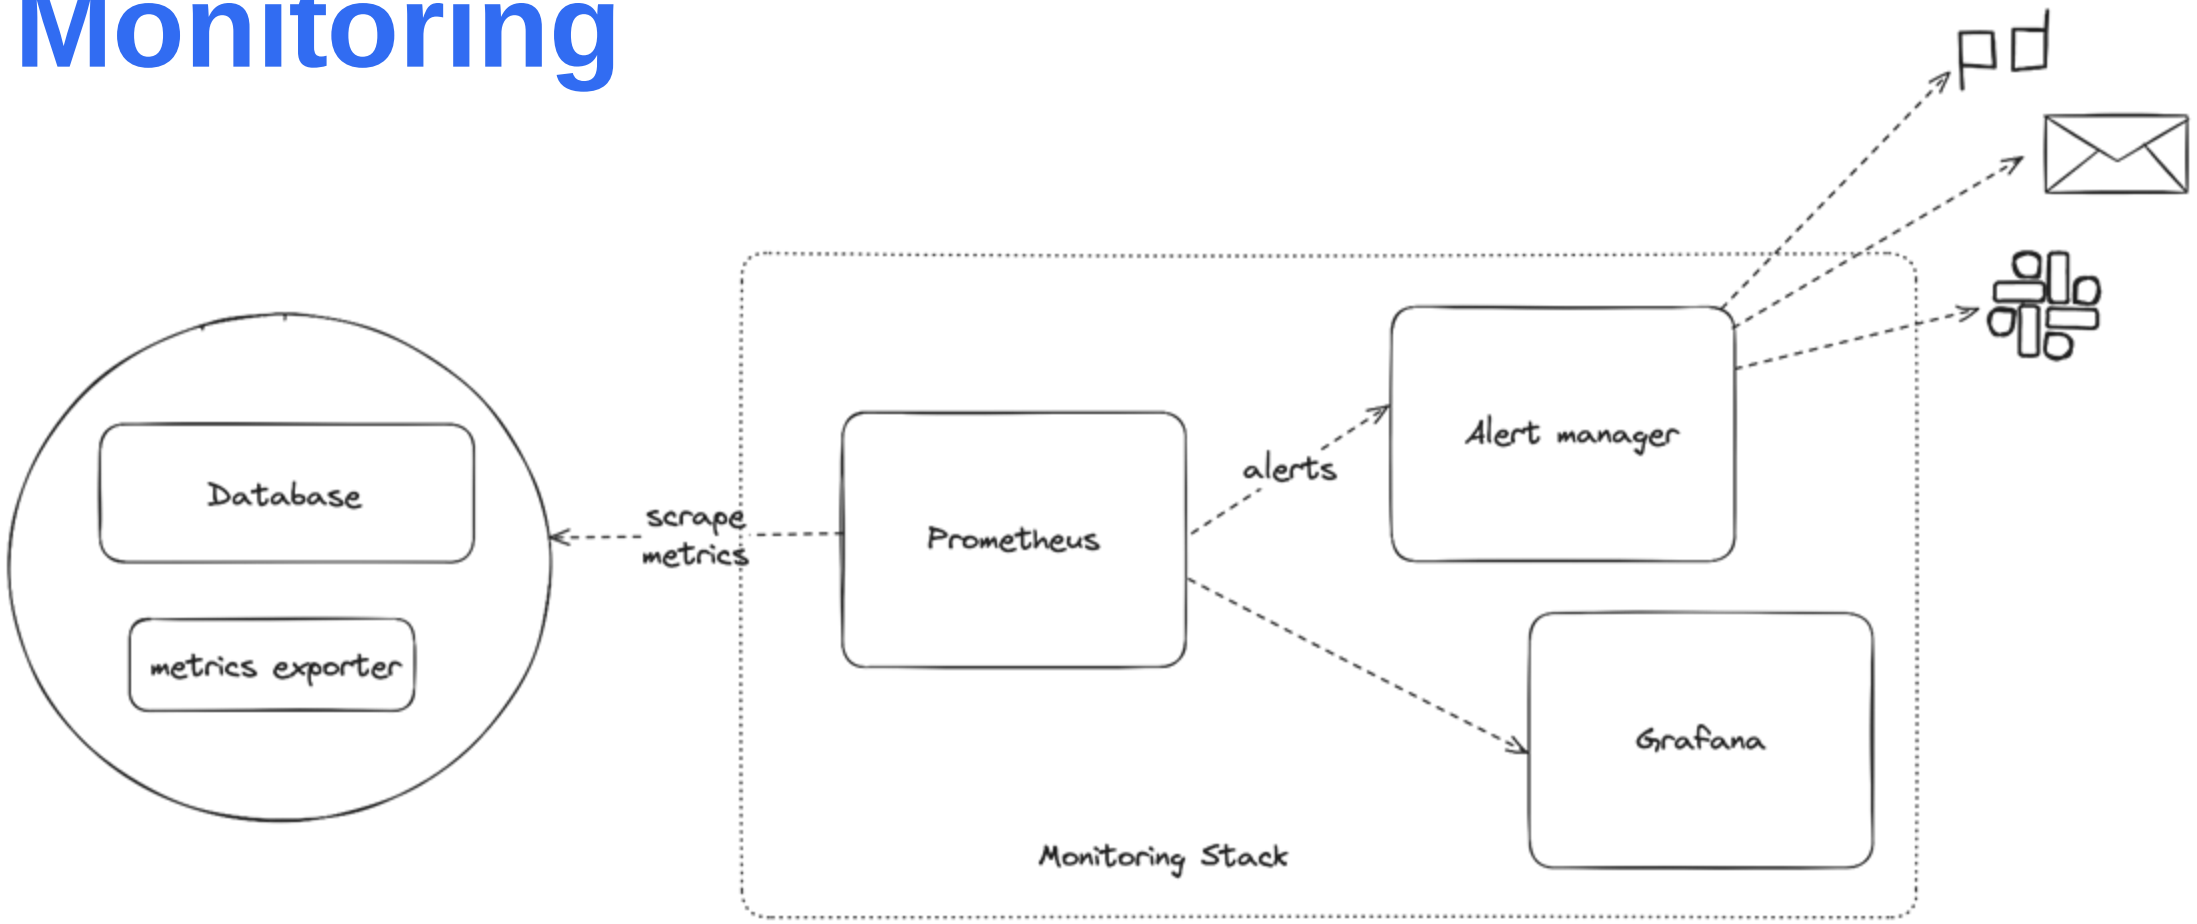

crunchydata

CPK:
Support/Troubleshooting

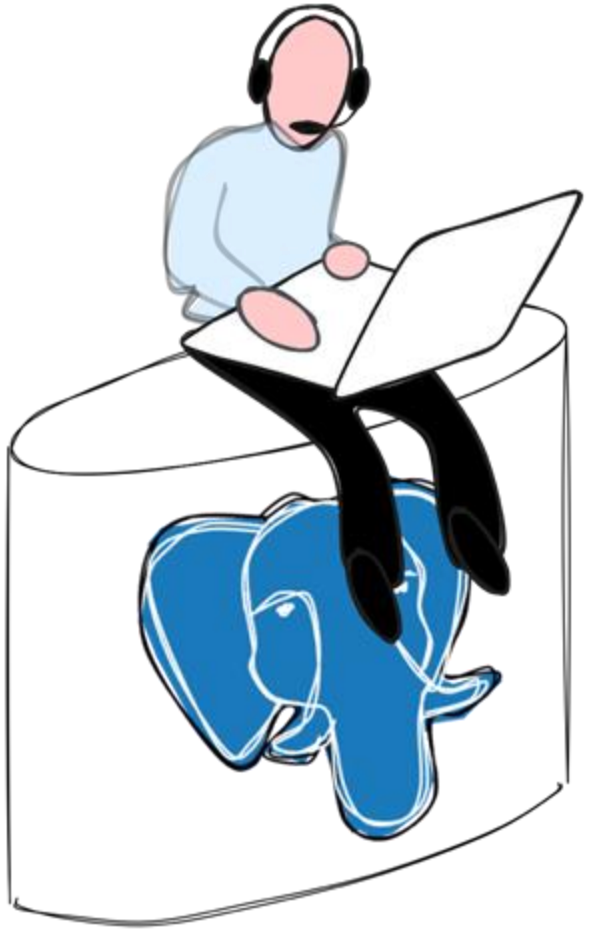

crunchydata

CPK: DB Software Install/Upgrade

Install

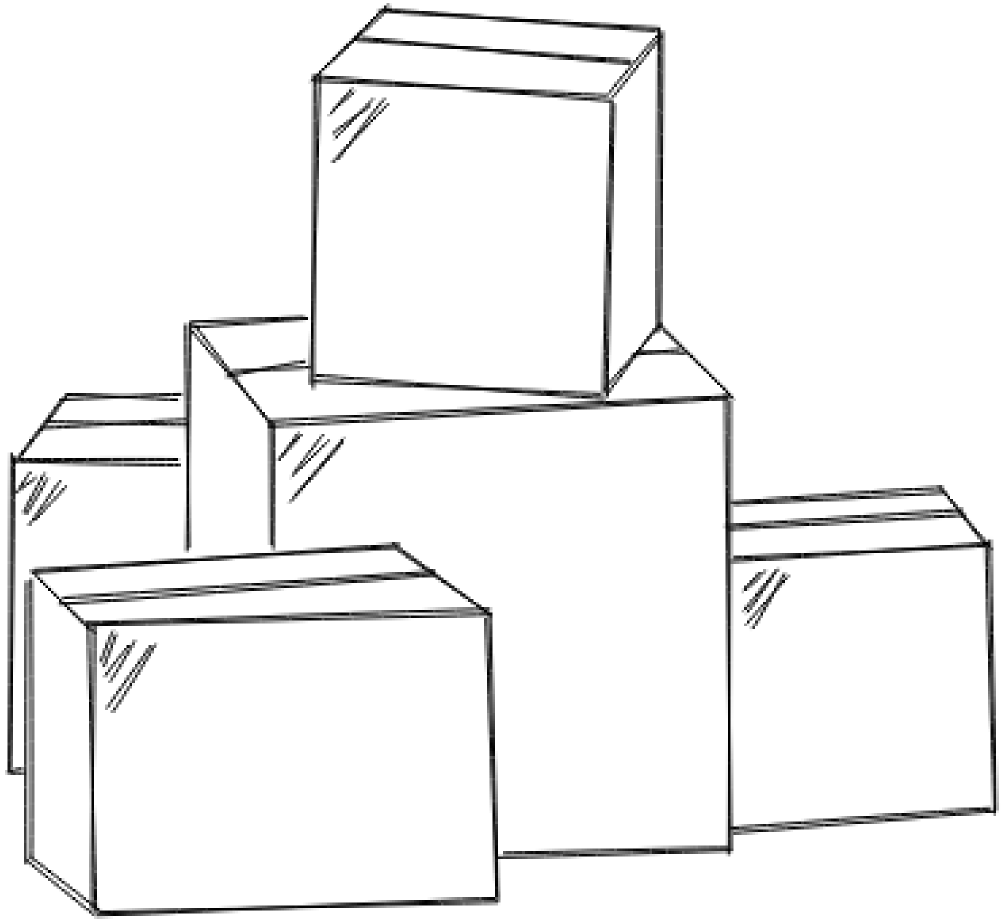

crunchydata

CPK: DB Software Install/Upgrade

Postgres Minor Upgrade

Ordered, automated rolling updates

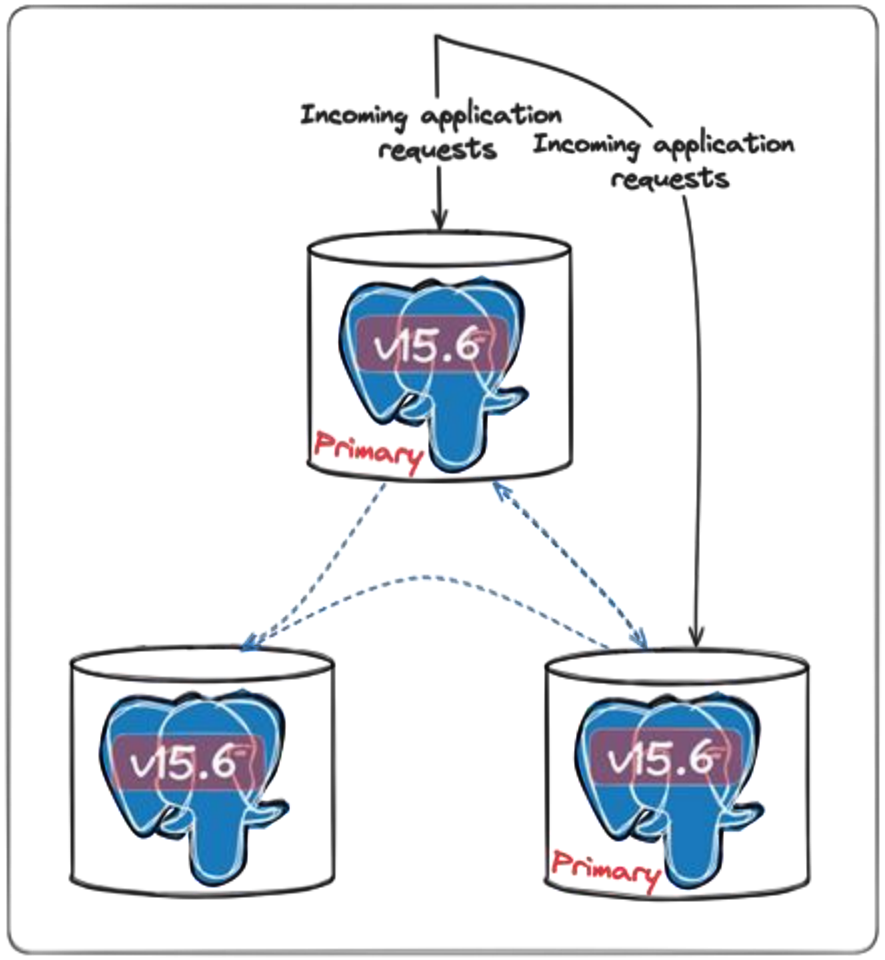

crunchydata

CPK: DB Software Install/Upgrade

Postgres Major Upgrade

- Automated by Operator
- Human planning & testing!

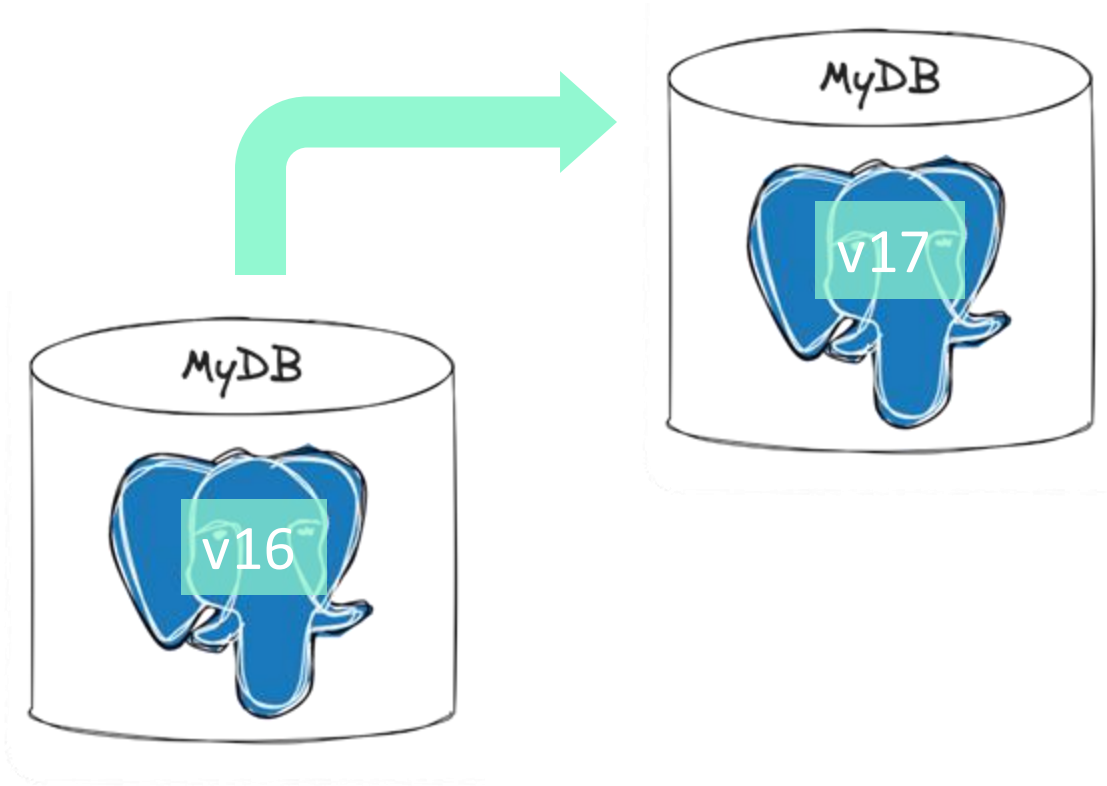

crunchy data

CPK:

Database Expertise

Operator:

- Built-in database expertise

Human expert:

- Strategic considerations
- Application needs
- Business requirements

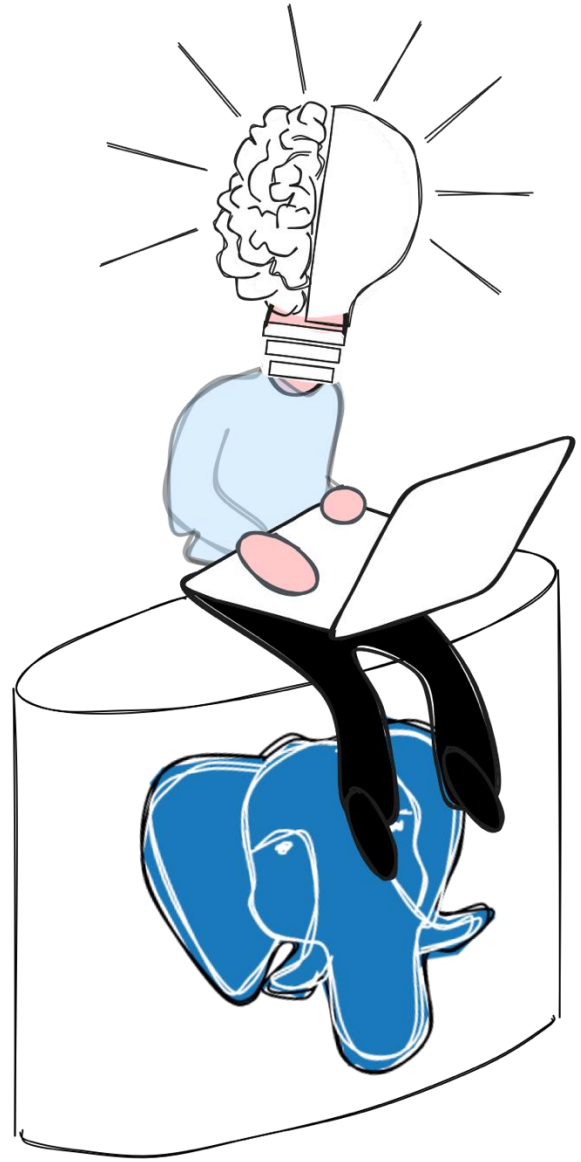

CPK:

Performance Tuning

- Configuration Parameters
- Connection pooling
- pg_stat_statements
- Logging
- Monitoring and alerting

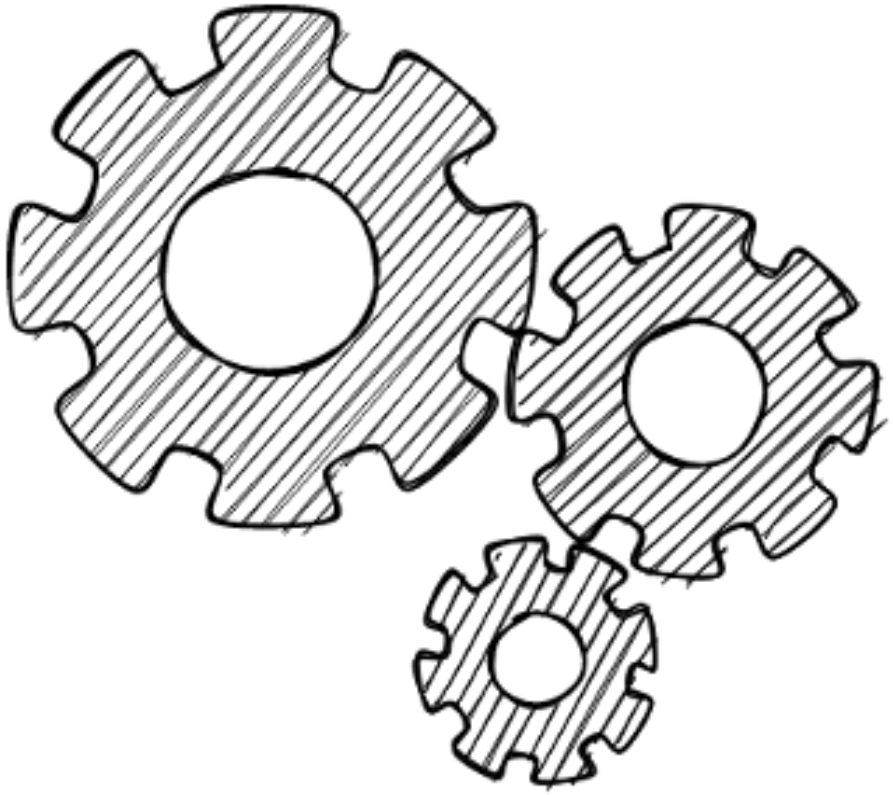

CPK:

Capacity Planning

```
dataVolumeClaimSpec:
        accessModes:
        - "ReadWriteOnce"
        resources:
          requests:
            storage: 1Gi
            storage: 10Gi
```

CPK:

Database Creation

```
users:
- name: slonik
  databases:
    - pgconfeu
```

CPK:

Database Maintenance

- Taking backups
- Creating users
- Rebuilding Indexes
- Gathering statistics
- Fixing errors
- Log management

Not an exhaustive list!

Agenda

1. Database Architecture History
2. Container Orchestration/Kubernetes Overview
3. Kubernetes Superpowers
4. What does a DBA do anyway?
5. Postgres Operators for Kubernetes
6. What do you want from an Operator?
7. What does CPK do for you?
8. Getting Started with CPK

crunchy data

Getting Started with CPK

1. **Deploy Operator from OperatorHub**

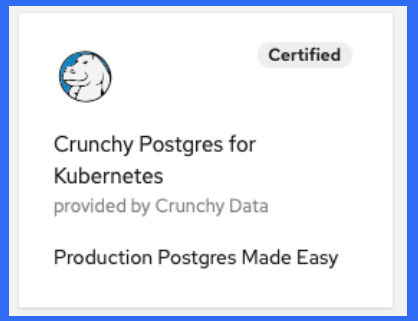

2. **Deploy using examples repository**
(https://github.com/CrunchyData/postgres-operator-examples)

3. **Learn PostgreSQL**
(https://www.crunchydata.com/developers/tutorials)

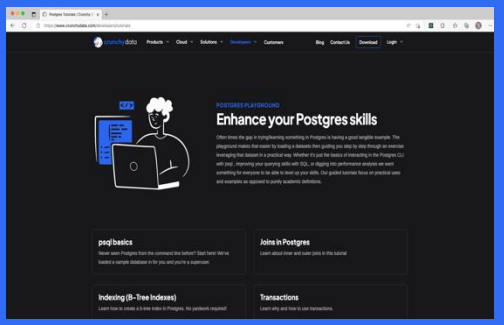

Getting Started with CPK

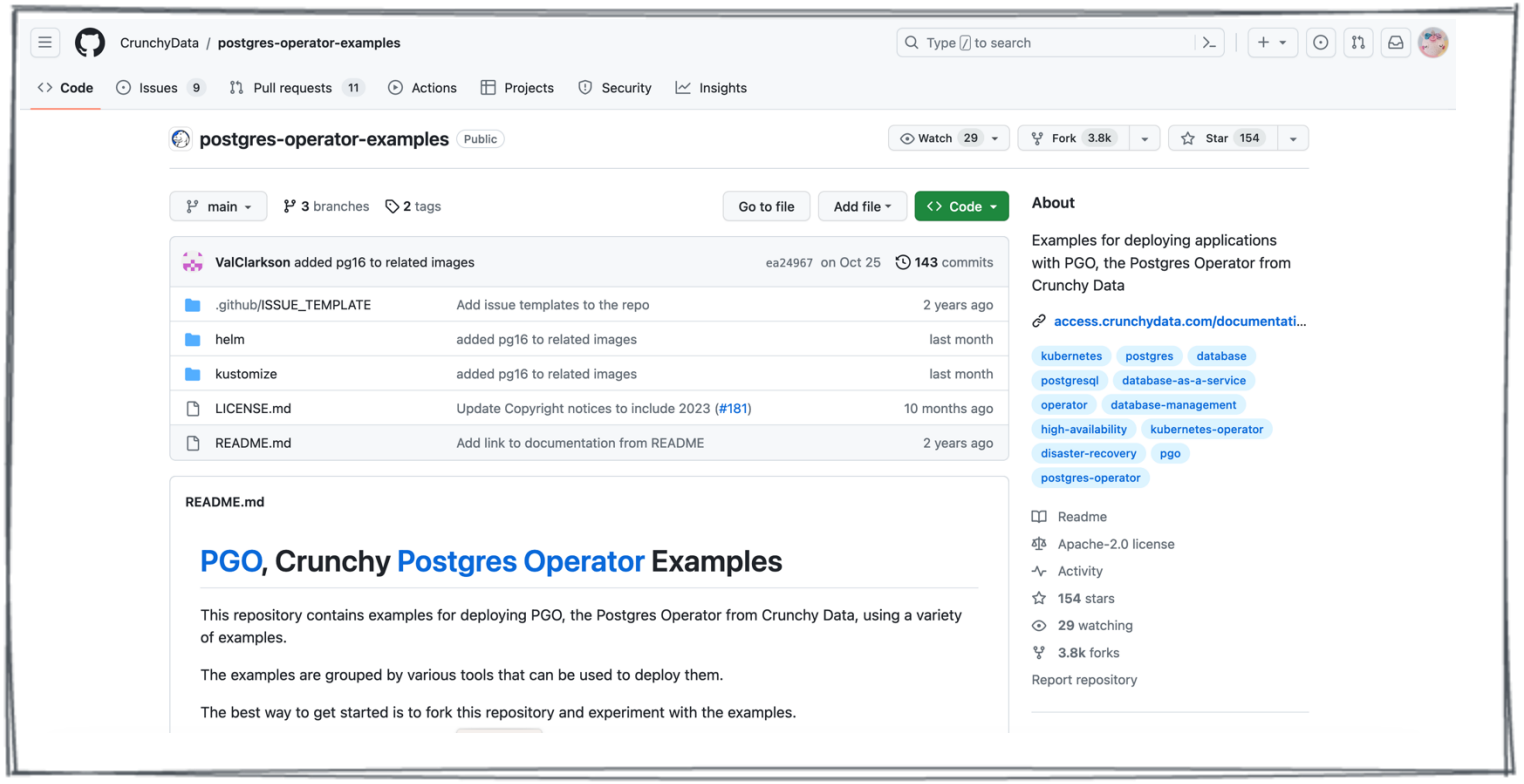

https://github.com/CrunchyData/postgres-operator-examples

crunchydata

Getting Started with CPK

Install the Operator

```
% git clone git@github.com:myuser/postgres-operator-examples.git
% cd postgres-operator-examples

% kubectl create namespace postgres-operator
% kubectl config set-context --current --namespace=postgres-operator

% kubectl apply --server-side -k kustomize/install/default
```

Getting Started with CPK

Install the Operator

```
customresourcedefinition.apiextensions.k8s.io/pgadmins.postgres-
operator.crunchydata.com serverside-applied
customresourcedefinition.apiextensions.k8s.io/pgupgrades.postgres-
operator.crunchydata.com serverside-applied
customresourcedefinition.apiextensions.k8s.io/postgresclusters.postgres-
operator.crunchydata.com serverside-applied
serviceaccount/pgo serverside-applied
clusterrole.rbac.authorization.k8s.io/postgres-operator serverside-applied
clusterrolebinding.rbac.authorization.k8s.io/postgres-operator serverside-
applied
deployment.apps/pgo serverside-applied
```

crunchydata

Getting Started with CPK

Create a PostgresCluster

```
% cp –r kustomize/postgres kustomize/pgconfeu

% vi kustomize/pgconfeu/postgres.yaml
```

Getting Started with CPK

PostgresCluster Manifest

```yaml
apiVersion: postgres-operator.crunchydata.com/v1beta1
kind: PostgresCluster
metadata:
  name: athens
spec:
postgresVersion: 17
  instances:
    - name: instance1
      replicas: 3
      dataVolumeClaimSpec:
        accessModes:
        - "ReadWriteOnce"
        resources:
          requests:
            storage: 1Gi
```

crunchydata

Getting Started with CPK

PostgresCluster Manifest (contd.)

```
backups:
   pgbackrest:
     repos:
     - name: repo1
       volume:
         volumeClaimSpec:
           accessModes:
           - "ReadWriteOnce"
           resources:
             requests:
               storage: 1Gi
```

Getting Started with CPK

Create a PostgresCluster

```
% kubectl apply -k athens

% kubectl get pods

NAME                                                READY   STATUS      RESTARTS   AGE
athens-backup-b7v4-krfkb                            0/1     Completed   0          13m
athens-instance1-b9nq-0                             4/4     Running     0          13m
athens-instance1-ljg8-0                             4/4     Running     0          13m
athens-instance1-m6mx-0                             4/4     Running     0          13m
athens-repo-host-0                                  2/2     Running     0          13m
pgadmin-611942b8-585c-4dea-886a-6774a1eccd7a-0      1/1     Running     0          45m
pgo-6fdd547449-6ghr9                                1/1     Running     0          46m
```

Conclusions

- CPK: Your Virtual DBA
- Postgres Expertise
- Automation
- Robust, Secure, Scalable Architecture
- Combined Strength: Postgres & Kubernetes
- Leaves the Fun Stuff to You!

Thank You!

Karen Jex | Crunchy Data | karen.jex@crunchydata.com