
# Postgres sur Kubernetes pour le DBA réticent

**Karen Jex | Senior Solutions Architect @ Crunchy Data**

**PGDay FR | Mons | juin 2025**

# whoami

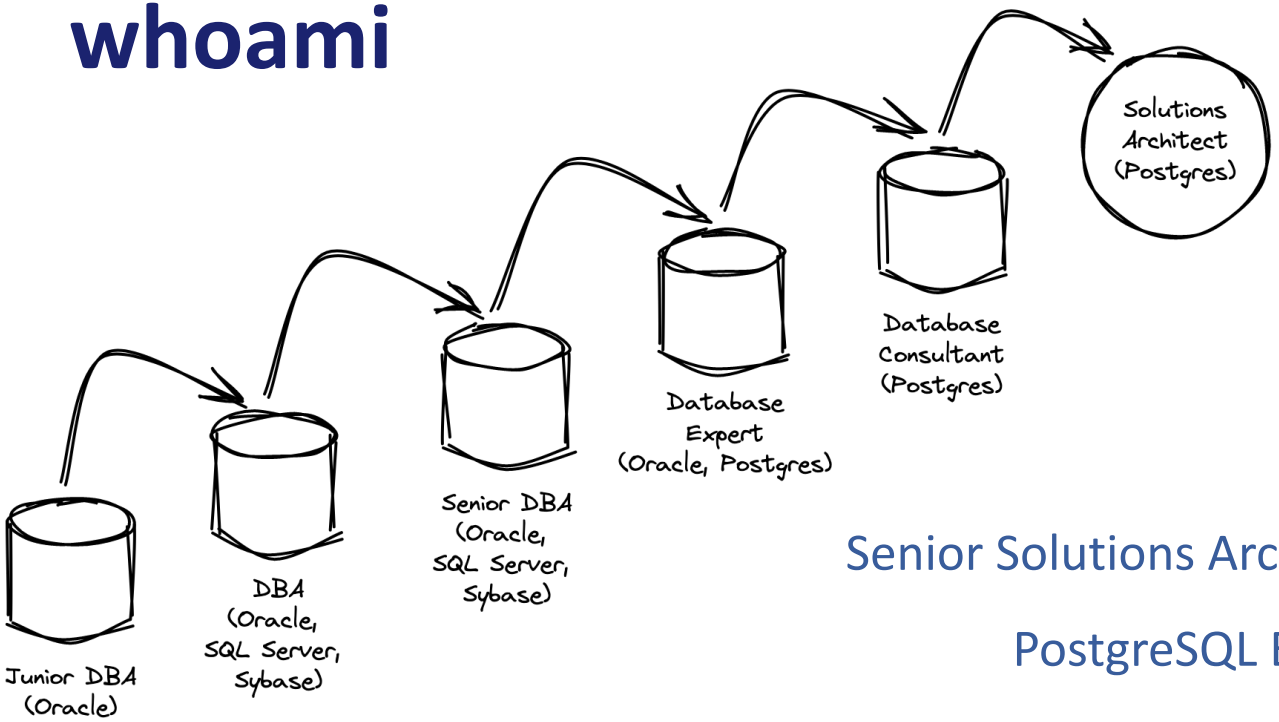


Senior Solutions Architect @ Crunchy Data

PostgreSQL Europe Board Member

PostgreSQL Europe Diversity Committee Chair



crunchy data

# Introduction


Image: https://commons.wikimedia.org/wiki/File:Largetape.jpg



crunchy data

# Agenda

1. Background
2. Audience Participation
3. DBA Concerns, Worries, Fears
4. Challenges of Databases on Kubernetes
5. Strengths of Databases on Kubernetes
6. Getting Started and Building Confidence

# Agenda

1. **Background**
2. Audience Participation
3. DBA Concerns, Worries, Fears
4. Challenges of Databases on Kubernetes
5. The Positive Side - Strengths of Kubernetes
6. Getting Started and Building Confidence

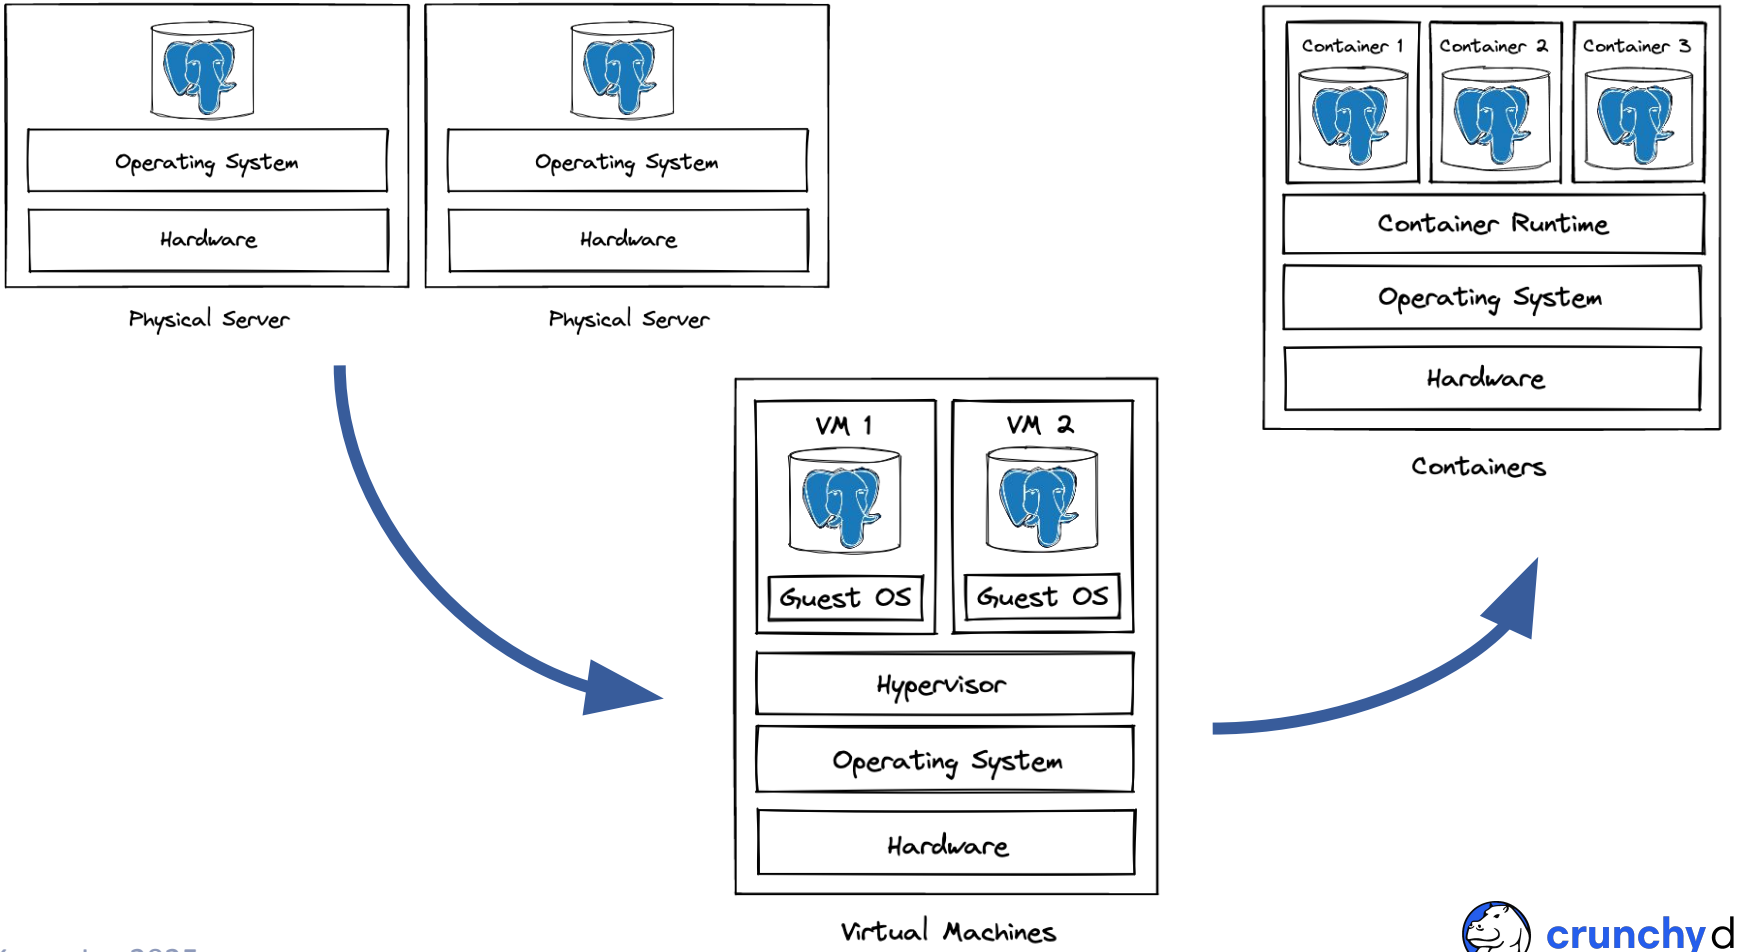




crunchy data

# Background

Kubernetes is 10 years old!

| | |
|---|---|
| 2016 | Kubernetes Operator and Stateful Set support in 2016 |
| 2017 | Crunchy Data Postgres Operator (PGO) |
| 2020 | Data on Kubernetes Community (DoKC) launched |
| 2022 | 70% of companies running stateful workloads K8s in production |

# Dominance of Databases on Kubernetes

*Database Dominance: **Databases remain the #1 workload** on Kubernetes for the third year. Even as the ecosystem expands into more advanced use cases like AI/ML, streaming, and analytics, databases continue to serve as the backbone of DoK deployments. This consistency demonstrates the platform's **reliability for mission-critical workloads**.*

2024 Data on Kubernetes Report



crunchy data

# Agenda

1. Background
2. **Audience Participation**
3. DBA Concerns, Worries, Fears
4. Challenges of Databases on Kubernetes
5. Strengths of Databases on Kubernetes
6. Getting Started and Building Confidence

# slido.com
# #2138903

1. Do you run databases on Kubernetes
2. If yes: what are your challenges
3. If no: what are your worries/concerns



crunchy data

# Est-ce que vous avez déjà des BDD sur Kubernetes ?

# Réponses:

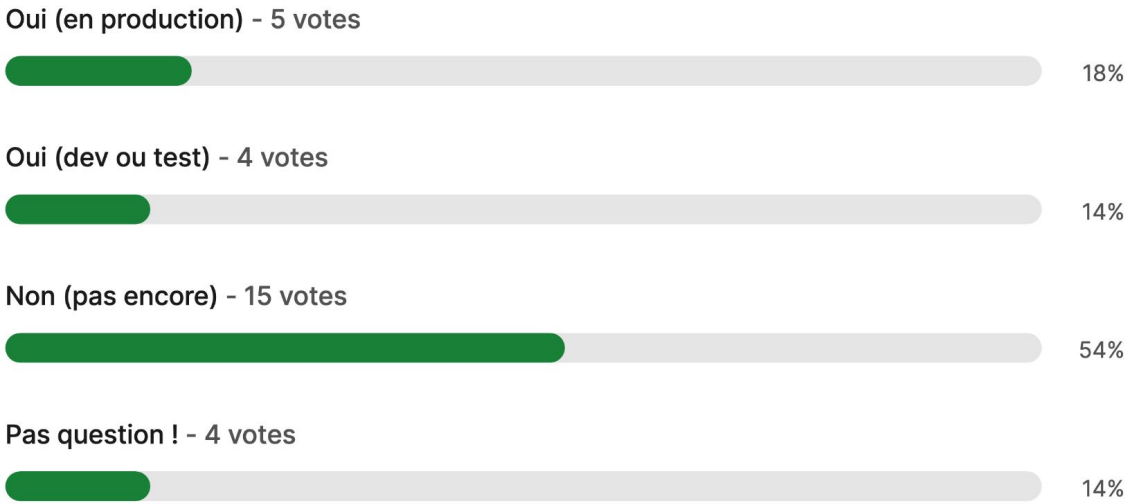




crunchy data

# Si vous avez des BDD sur Kubernetes: Quels sont vos défis (actuels ou par le passé) ?

# Réponses:

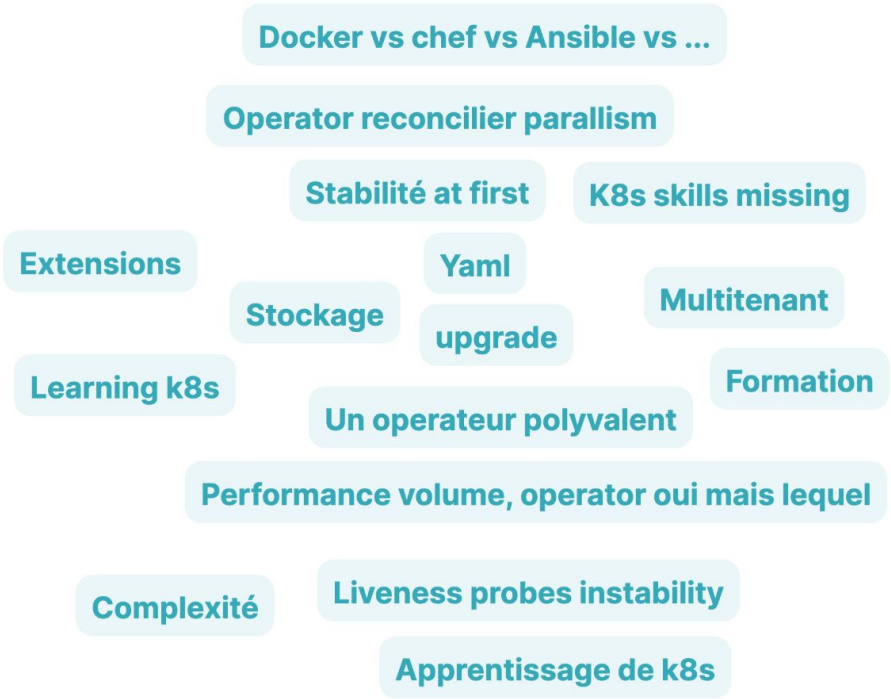




crunchy data

# Si vous n’avez pas de BDD sur Kubernetes: Qu’est-ce qui vous retient ? Quelles sont vos inquiétudes ?

# Réponses:




crunchy data

# Agenda

1. Background
2. Audience Participation
3. **DBA Concerns, Worries, Fears**
4. Challenges of Databases on Kubernetes
5. Strengths of Databases on Kubernetes
6. Getting Started and Building Confidence

# Scary Stuff

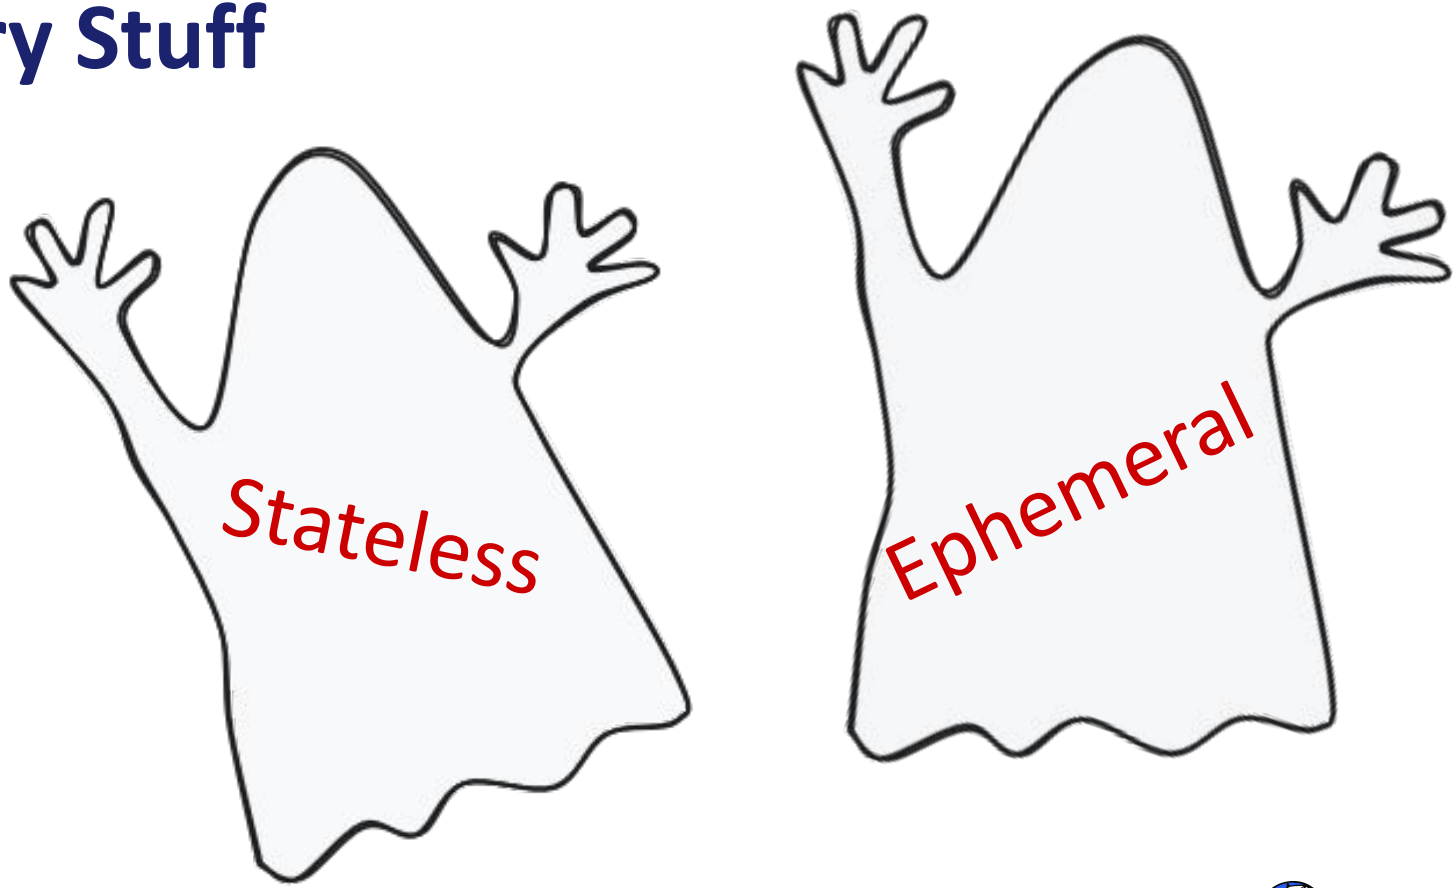

# Concerns

- Will we still need DBAs?
- Kubernetes is just for stateless apps.
- Will my data be safe?
- What about the skills I've developed over all these years?
- It's yet another thing to learn.
- I don't know where to start.
- Where do I go for help?

# Database on Kubernetes Users

*“There are two kinds of people [who want to run databases on Kubernetes]:*
***DBAs that need to be convinced that Kubernetes is stable enough** and won't mess up their databases, and*
***developers who are ready to run everything in Kubernetes** but might not know the specificities of databases and why DBAs are so conservative.”*

Lætitia Avrot, DoKC Town Hall Panel, March 2025

# Agenda

1. Background
2. Audience Participation
3. DBA Concerns, Worries, Fears
4. **Challenges of Databases on Kubernetes**
5. Strengths of Databases on Kubernetes
6. Getting Started and Building Confidence

# Challenges

- Kubernetes is new(ish) and complex.
- You may not (yet) have experienced/confident staff
- Things need to be done in (slightly) different ways.
- “lift and shift” approach may not always work.
- The business (or you yourself) may need additional reassurances.
- Kubernetes doesn’t know how to manage a database.

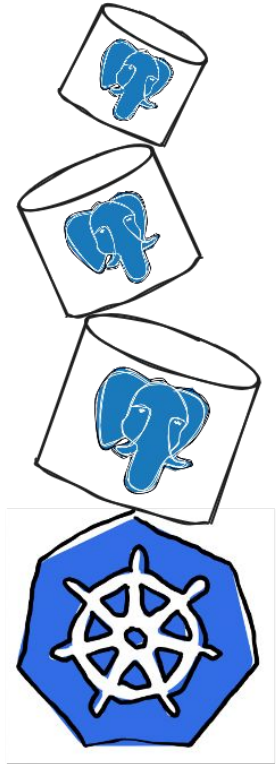

# Focus on your Strengths

- You don’t need to be a K8s expert
- Let someone else worry about it!
- Consider a K8s platform

# Start Small

- Don’t do it all at once
- Not all databases need to be migrated to Kubernetes
- Consider starting with a small, new project
- Build confidence with non-critical database applications

# Agenda

1. Background
2. Audience Participation
3. DBA Concerns, Worries, Fears
4. Challenges of Databases on Kubernetes
5. **Strengths of Databases on Kubernetes**
6. Getting Started and Building Confidence

# 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

# 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/

# Challenges without Kubernetes

- Not enough time to automate tasks
- Firefighting mode
- Middle of the night calls!
- Upgrades are scary and time-consuming
- It’s hard to recruit DBAs

# Database Architecture

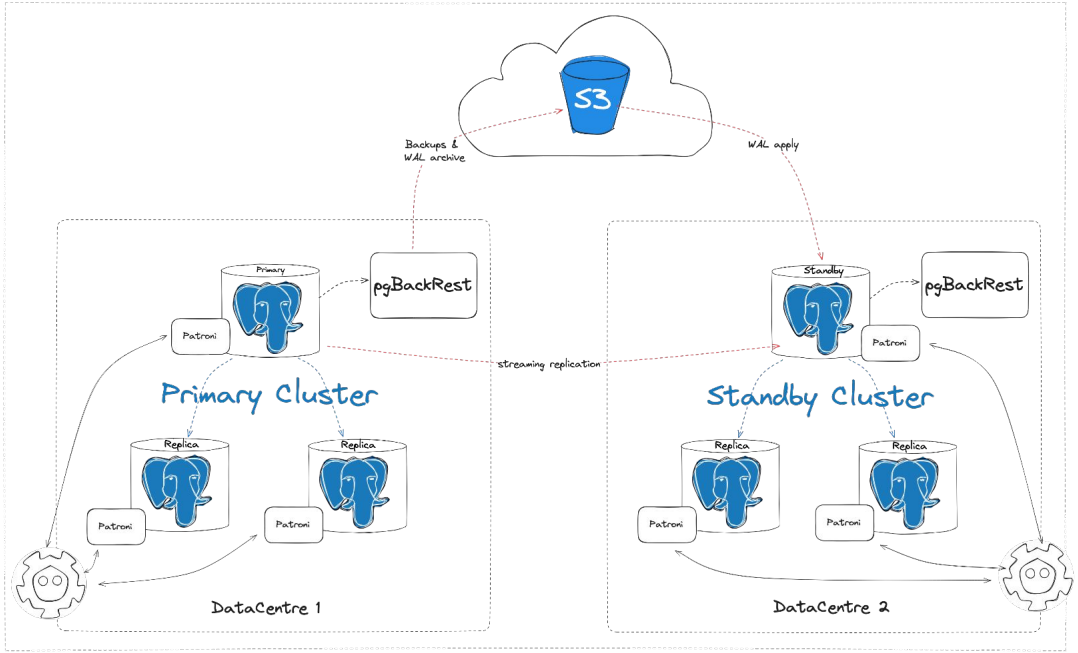




crunchy data

# Database Architecture

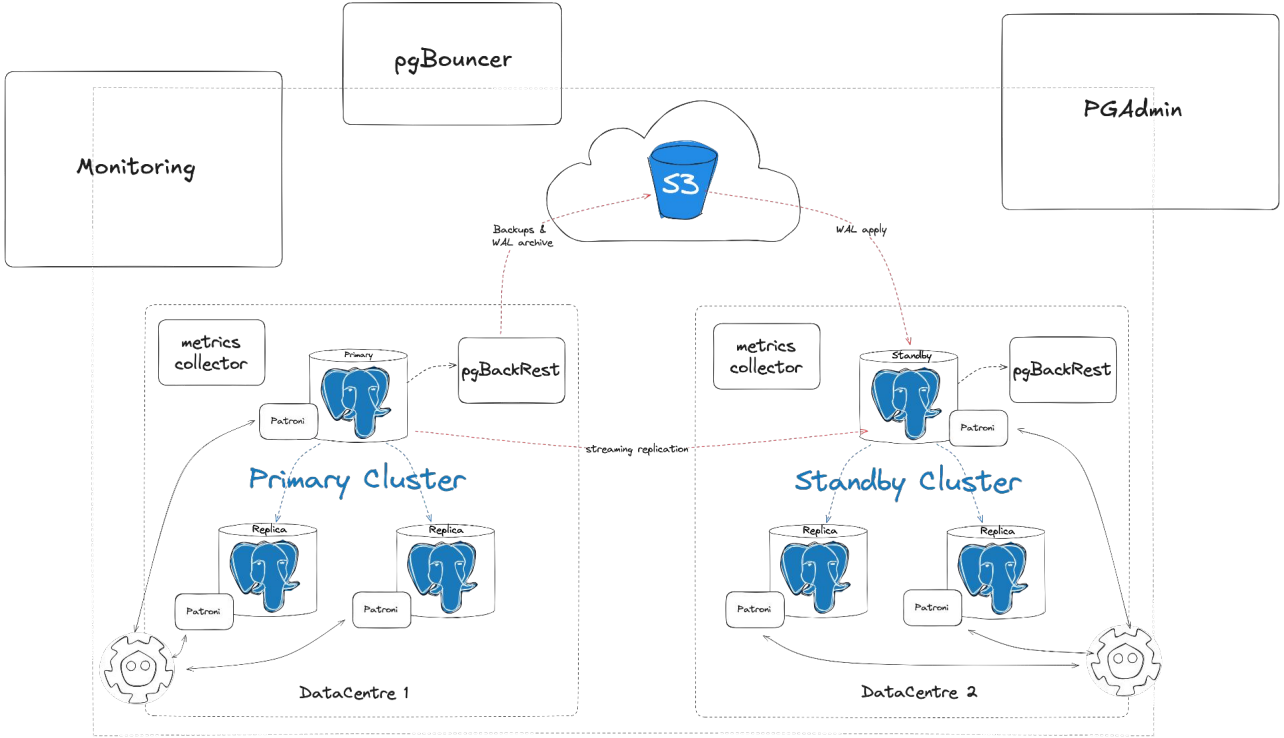




crunchy data

## ~~DBA~~ Operator Responsibilities

# Monitoring

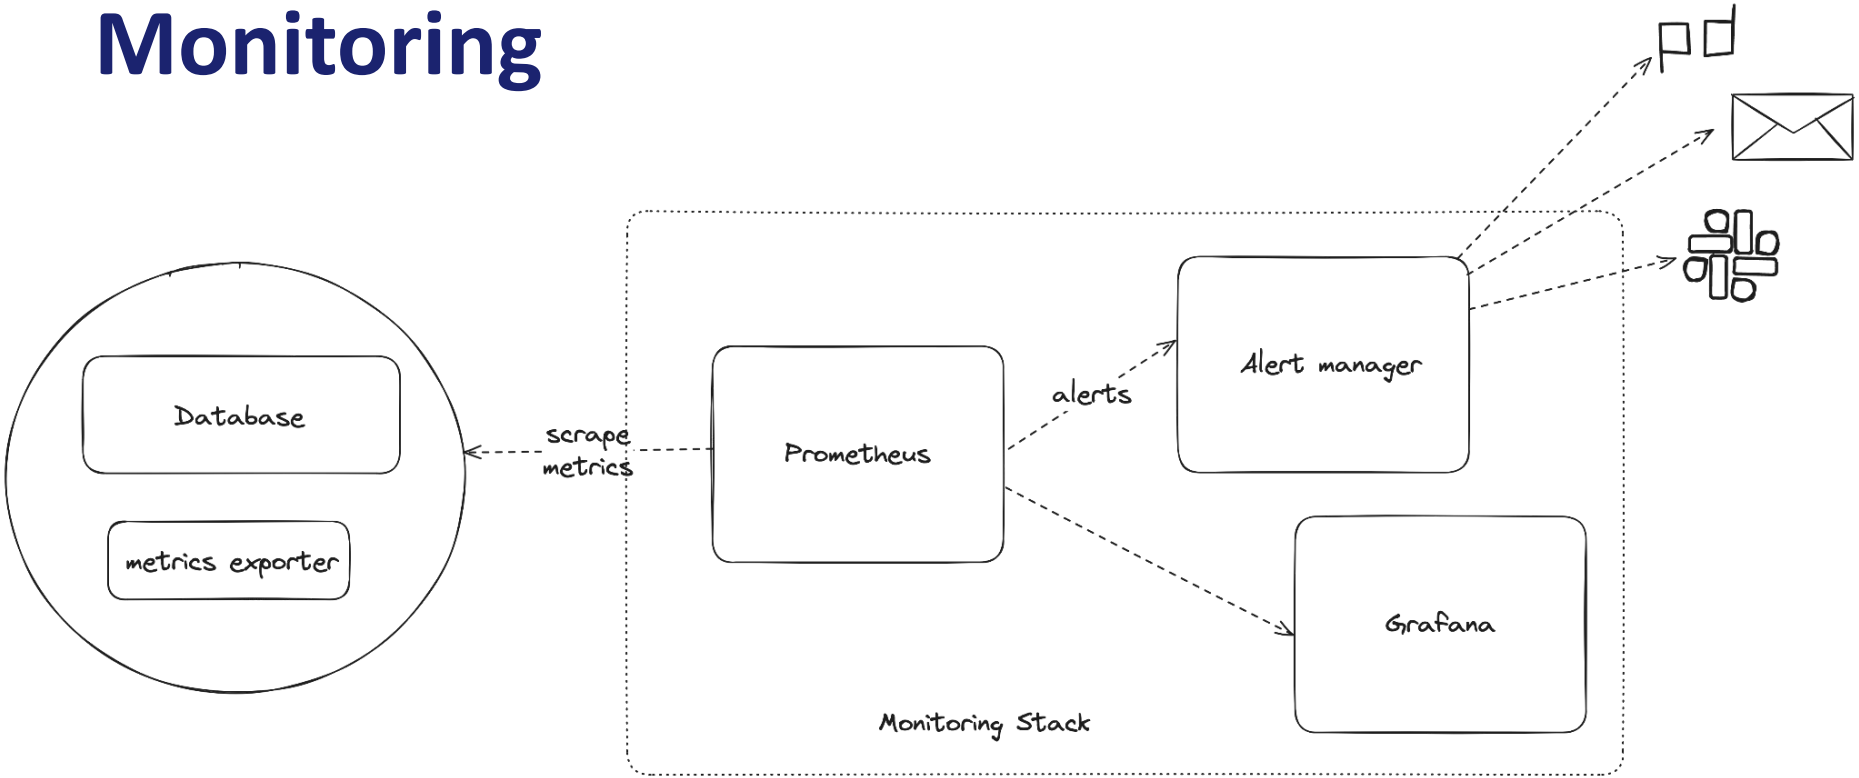




crunchy data

# Kubernetes Features


- Provisioning & Deployment
- Configuration
- Scheduling
- Scaling up & down
- Self-healing
- Services
- Storage
- Resource allocation
- Load Balancing & Networking
- Security
- Stateful Sets
- Sidecars



crunchy data

# Postgres Operators for Kubernetes

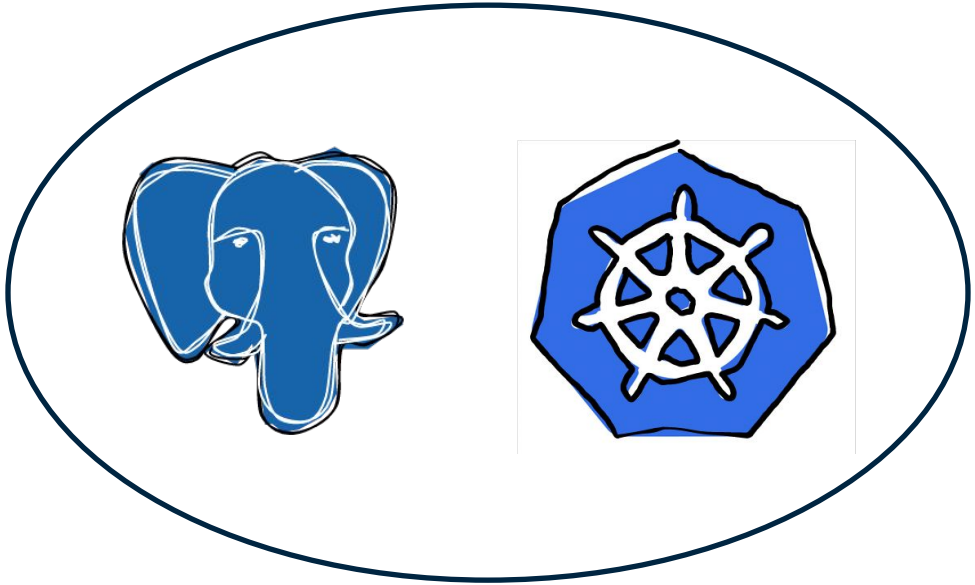



crunchy data

# Database Architecture

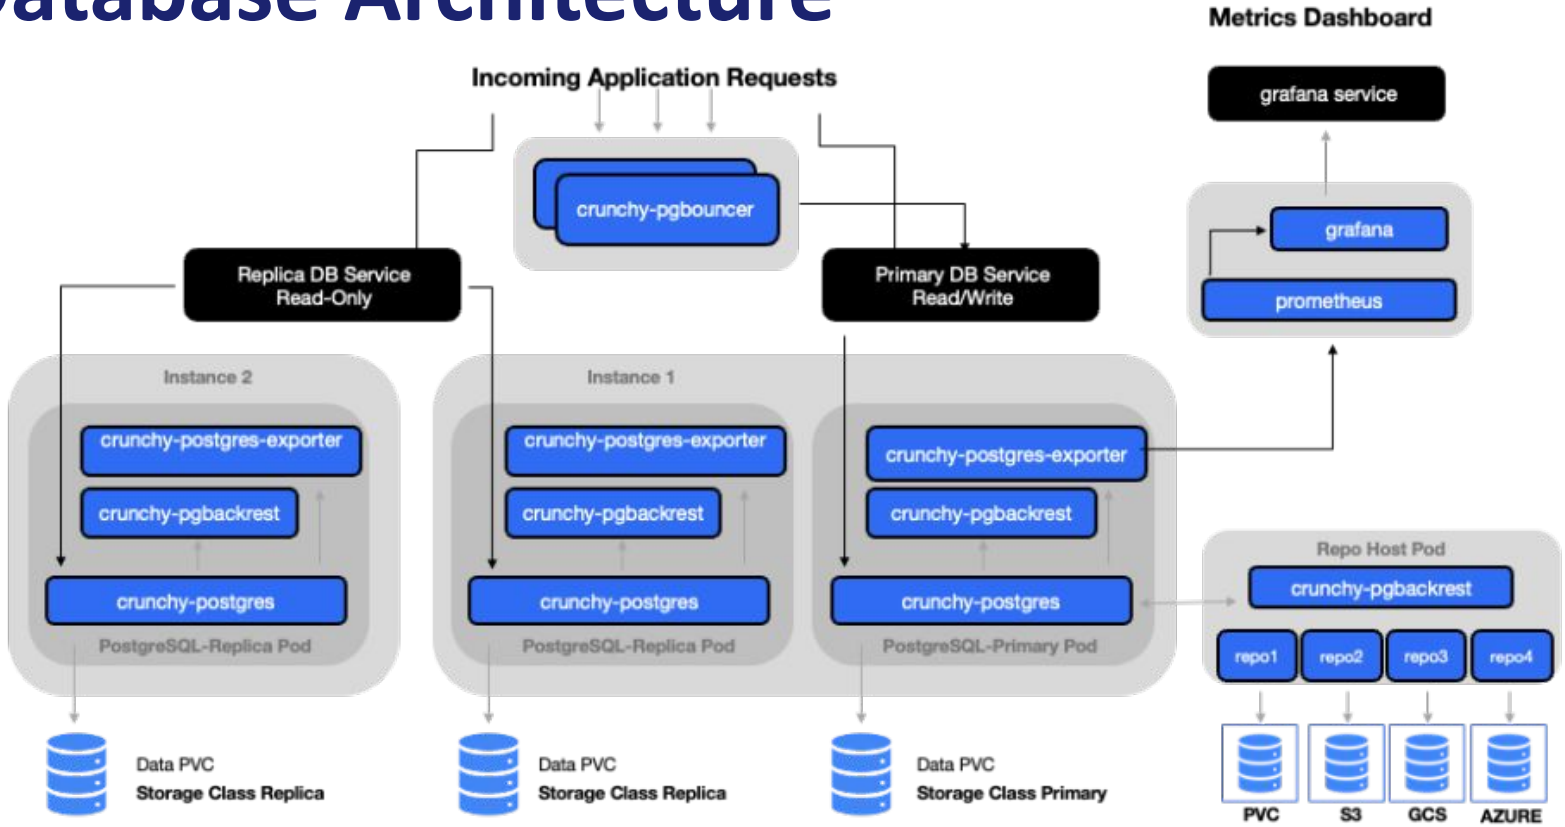




crunchy data

# More Operator Features

- Security
  - pg_hba.conf, user creation
  - password encryption, secrets
  - SSL/TLS, managing certificates
- Upgrades
  - Automated, rolling, minor upgrades
  - Automated major upgrades

# Reassurance!

- Databases on Kubernetes is no longer new
- Mission-critical, multi-terabyte databases on Kubernetes
- Tried and tested in production for multiple years
- Teams can do more with less
- Internal DBaaS

# Case Studies

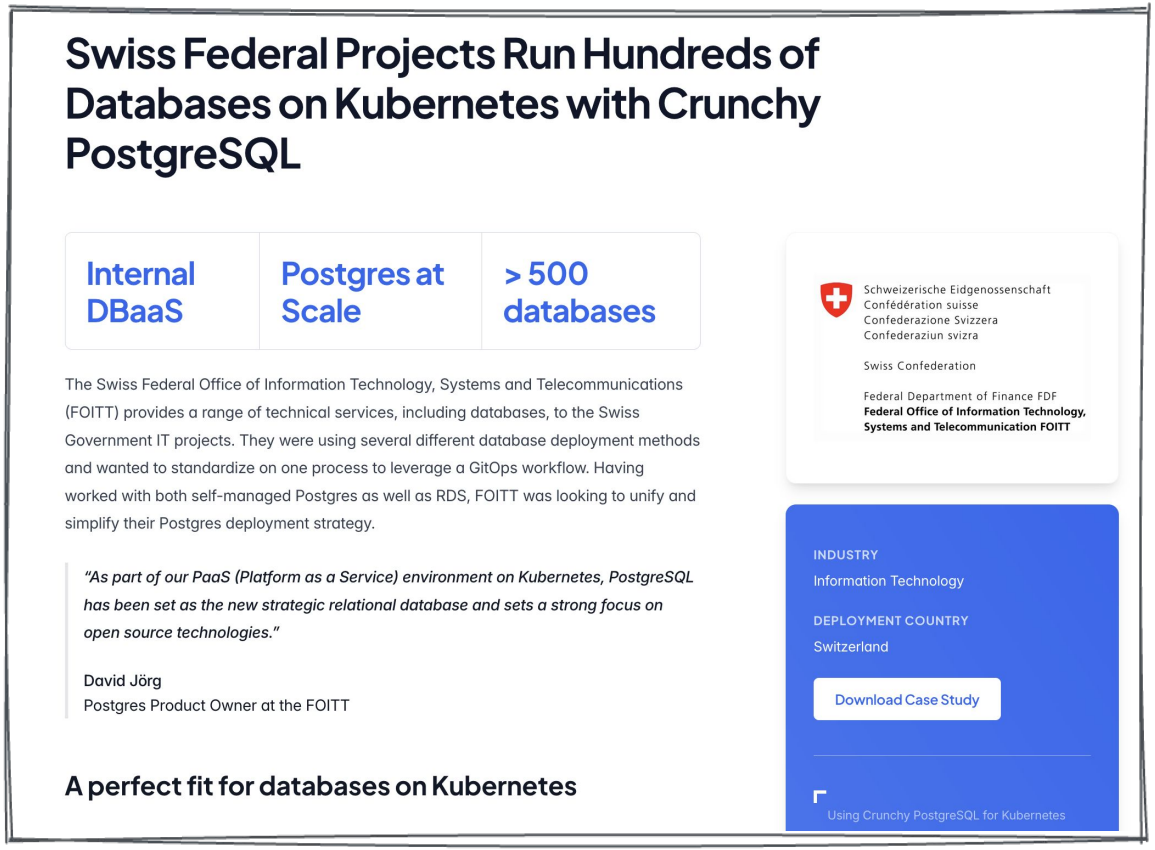

## Swiss Federal Projects Run Hundreds of Databases on Kubernetes with Crunchy PostgreSQL

| Internal DBaaS | Postgres at Scale | > 500 databases |
|---|---|---|

The Swiss Federal Office of Information Technology, Systems and Telecommunications (FOITT) provides a range of technical services, including databases, to the Swiss Government IT projects. They were using several different database deployment methods and wanted to standardize on one process to leverage a GitOps workflow. Having worked with both self-managed Postgres as well as RDS, FOITT was looking to unify and simplify their Postgres deployment strategy.

> *"As part of our PaaS (Platform as a Service) environment on Kubernetes, PostgreSQL has been set as the new strategic relational database and sets a strong focus on open source technologies."*
>
> David Jörg
> Postgres Product Owner at the FOITT

### A perfect fit for databases on Kubernetes

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Finance FDF
**Federal Office of Information Technology, Systems and Telecommunication FOITT**

INDUSTRY
Information Technology

DEPLOYMENT COUNTRY
Switzerland

Download Case Study

Using Crunchy PostgreSQL for Kubernetes

https://www.crunchydata.com/case-studies/swiss-fiott



crunchy data

# Case Studies

## World Wide Technology Starts Leveraging Postgres with Crunchy Postgres for Kubernetes

World Wide Technology (WWT) is a global technology services provider based in St. Louis, Missouri specializing in enterprise architecture, IT services, and digital strategy. They have a large footprint across the U.S. deploying technology solutions and custom application development to businesses needing digital innovation.

In 2018 WWT set out to modernize their infrastructure and application development environments. As they evaluated technologies for this modernization effort, WWT looked to Kubernetes as their cloud native infrastructure abstraction. In selecting a database technology for this new cloud native infrastructure, WWT selected Crunchy Postgres for Kubernetes as the best combination open source RDBMS and Kubernetes native


https://www.crunchydata.com/case-studies/wwt



crunchy data

# Human DBA Expertise

- Strategic considerations
- Architecture
- Data modelling
- Training/tutoring/mentoring
- Performance tuning

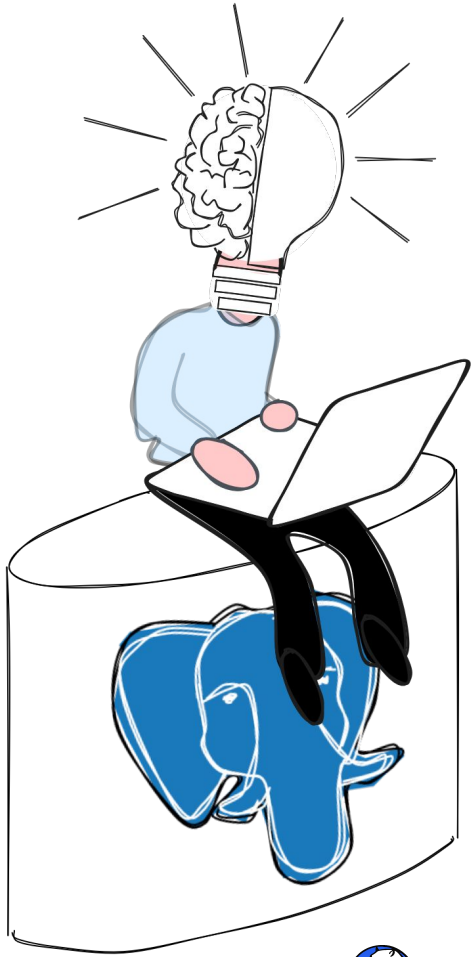

# Agenda

1. Background
2. Audience Participation
3. DBA Concerns, Worries, Fears
4. Challenges of Databases on Kubernetes
5. Strengths of Databases on Kubernetes
6. **Getting Started and Building Confidence**

# Getting Started

*“There’s nothing to it but to do it!”*

Frances Thai, DoKC Town Hall Panel, March 2025

# Getting Started: Kubernetes

- Learn “just enough” about Kubernetes

  https://kubernetes.io/docs/concepts/architecture

- Try creating a cluster for yourself

  https://kubernetes.io/docs/tutorials

# Getting Started: Postgres Operator

- Try out a Postgres Operator for Kubernetes
- Watch the magic and have fun!
- Docs, tutorials, videos
- Gain confidence
- Ask questions - there’s a community out there wanting to help

https://www.crunchydata.com/developers/get-started/postgres-operator

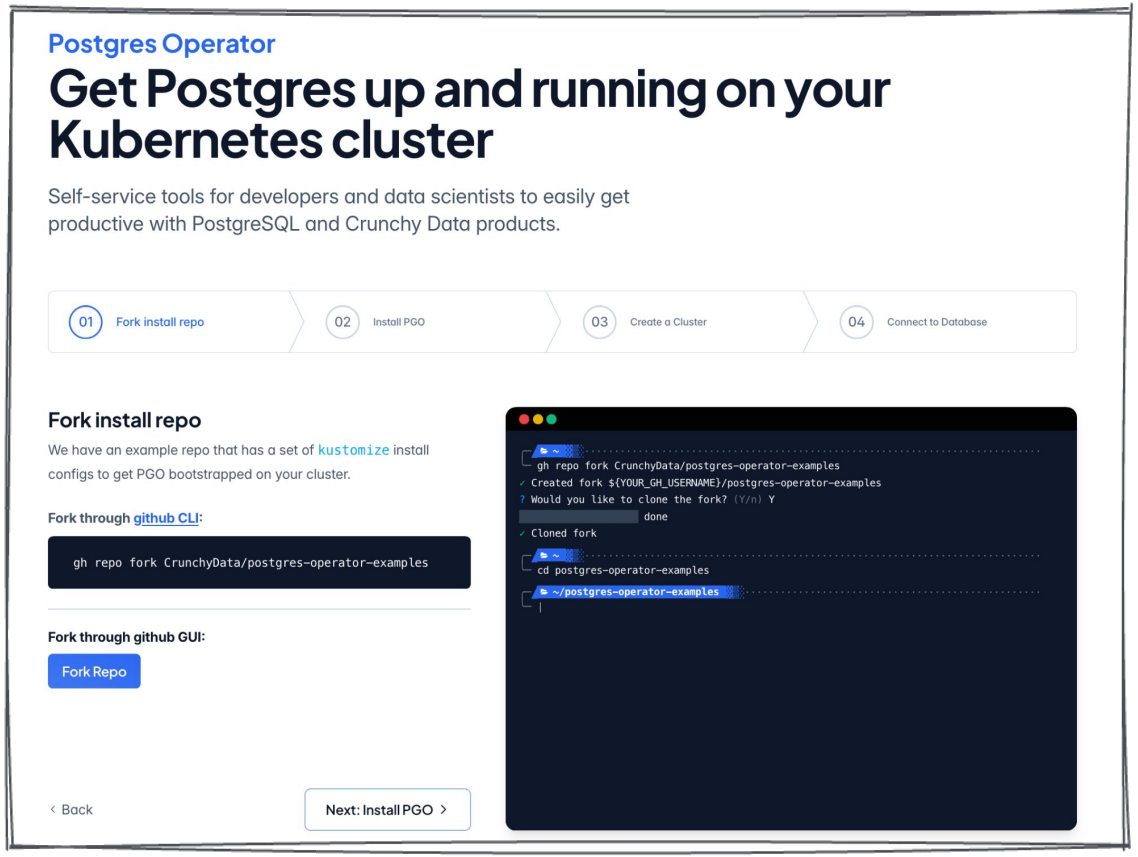




crunchy data

# Conclusions

- Data on Kubernetes is tried and tested
- Your database expertise is still relevant and necessary
- Do more with less
- Fewer middle of the night calls
- You get to do the fun parts of database administration

crunchy data

# Merci !

Karen Jex | @karenhjex | karen.jex@crunchydata.com