



SCaLE

Pasadena, CA

March 09-12, 2023

Life of a Postgres DevOps DBA

Welcome

Housekeeping Items



Slides and recording will be made available via email at:
doug.ortiz@enterprisedb.com



Questions are welcome to be asked, aiming for an interactive presentation.

If not possible, questions will be answered at the end

Vibhor Kumar

V.P. Performance Engineering And
Architecture Team

Authored tools -

- edb-ansible, postgres-deployment, pg_background, efm_extension, edb_user_login, edb-cloneschema, edb_block_commands

Expertise:

- Enterprise architecture, cloud technology, microservices, database technologies (Oracle, MySQL, PostgreSQL, DB2, EDB Postgres Advanced Server, MongoDB)

Leadership experience:

- Strong leader with expertise in defining security best practices, DevOps, Oracle migration and transformation, database and platform performance, and data security and governance.



Doug Ortiz

Senior Postgres DevOps
Engineer

Technologist

Original Co-Author of Open Source Projects:

- edb-deployment (aka postgres-deployment)
- edb-ansible

Experience in:

- Multi-Cloud
- Software Architecture and Development
- DevOps
- Microservices, Containerization, and K8s
- Automation
- Database Technologies



Agenda

- Introduction
- Postgres and DevOps philosophy
- Evolution of the DBA Role
- Postgres DevOps DBA
- Traditional DBA vs Postgres DevOps DBA
- Introduction to CloudNativePG
- Demo
- Takeaways
- Q&A

Definitions

Postgres	Database Administration	DevOps
<p>Free and open-source relational Database Management System emphasizing in extension and SQL compliance</p>	<ul style="list-style-type: none">• Specialized role within computer systems administration• Develop and operate the data storage ecosystem that digital business rely on	<ul style="list-style-type: none">• Philosophy for software development• Mix of culture, development, operations, and tools that allow for development, and improvement of products at a faster rate• Cross-disciplinary community of practices dedicated to the study of building, evolving, and operating rapidly-changing resilient systems at scale

Why Postgres



- As good as most commercial databases for both OLAP and OLTP workloads. It is a great general purpose database with performance benchmarks to prove its case
- No single commercial entity driving the roadmap
 - Vibrant, and growing community
 - Constantly adding new features and innovating
- Extremely extensible, supports multiple programming languages, and data models (multi-model)
 - SQL
 - Document (XML, JSONB)
 - Key value

Performant

Extensible, scalable, and
community driven


Built for speed

Evolution of DBA Role

DevOps - Principles



Systems thinking



Amplify feedback loops



Continuous learning and experimentation



DBA Role

How Postgres fits into the DevOps Framework

- Power and flexible open-source relational DBMS
- Popular among developer and used in a wide range of applications
- Easy integration with DevOps tools
 - Ansible
 - Jenkins
 - Chef
- Robust security features
 - SSL Encryption
 - Authentication
- Ability to handle large volumes of data with ease

Easy integration with
DevOps tools

Extensible, highly-
available, and scalable

Secure

Postgres DevOps DBA

- Specialized role that combines the responsibilities of a traditional database administrator with the principles of DevOps.
- Responsible for managing the DB's used by an organization's applications, while working closely with developers, testers, and other IT professionals to ensure that the DB is integrated into the overall DevOps process

Modern IT

Postgres DevOps DBA



Postgres

Database Administration

DevOps

Postgres DevOps DBA

Key skills and
experience required

- Expertise in Postgres
- DevOps tools and principles
- Database design principles
- Scripting and programming
- Security expertise
- Collaboration and communication
- Troubleshooting and problem-solving
- Familiarity with cloud technologies

Traditional DBA vs Postgres DevOps DBA

- Installing, configuring and upgrading Databases and software
- Monitoring and optimizing database performance
- Creating, maintaining, and testing backup and recovery procedures
- Designing and implementing data models and database schemas
- Managing database storage and capacity planning
- Ensuring data integrity and consistency
- Troubleshooting and resolving database issues
- Providing technical support to application developers and end-users
- Developing and implementing disaster recovery plans

- **Collaborates with developers and other IT professionals**
- **Automates database provisioning, deployment, and scaling using tools like Ansible, Chef, or Puppet**
- Designing and implementing high availability and disaster recovery solutions
- Implementing database monitoring and logging systems to ensure high performance and availability
- **Developing and maintaining CI/CD pipelines to continuously deploy database changes**
- **Managing and optimizing database infrastructure on cloud platforms like: AWS, Azure and GCP**
- **Managing and monitoring Postgres replication and sharding**
- Ensuring database security and compliance with regulator standards
- **Supporting and troubleshooting Postgres-related issues**

What is CloudNativePG (a.k.a CNPG)?

- “CloudNativePG is an open source operator designed to manage PostgreSQL workloads on any supported Kubernetes cluster running in private, public, hybrid, or multi-cloud environments.
- Originally built by EDB, released open source under Apache License 2.0 and submitted for CNCF Sandbox in April 2022
- Cloud Native PostgreSQL adheres to DevOps principles and concepts such as **declarative configuration** and **immutable infrastructure.**”
- Open source
 - <https://github.com/cloudnative-pg/cloudnative-pg>
- Documentation available at:
 - <https://cloudnative-pg.io/documentation/1.19/>

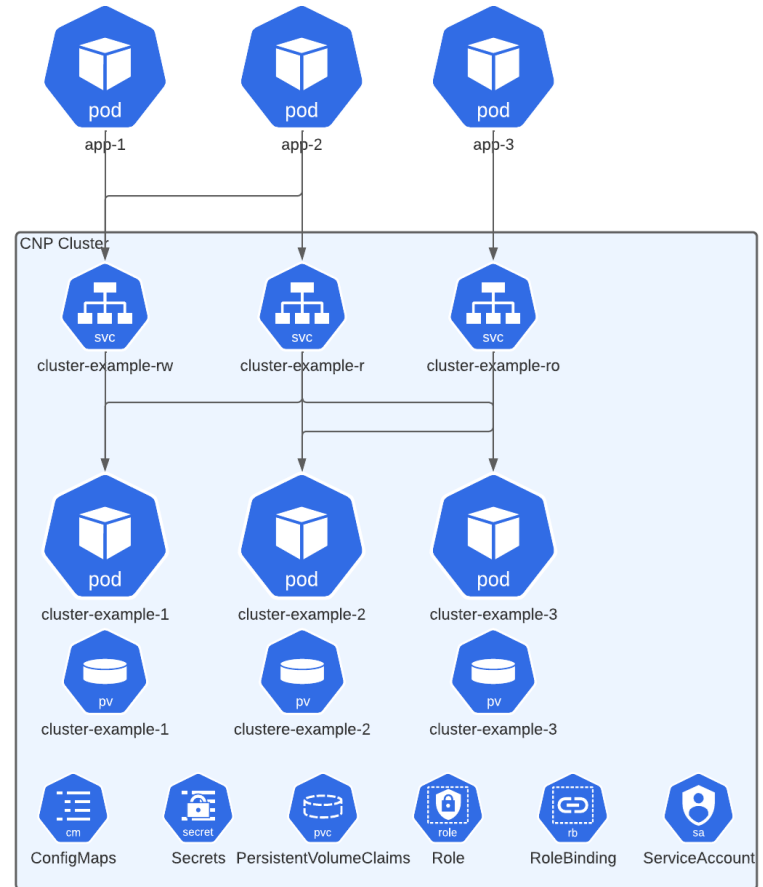
Why CloudNativePG?

- **Hard challenges and difficulties**
 - Manually deploying and managing Postgres
 - High availability
 - Disaster recovery
- **Kubernetes adoption is rising and becoming the de facto orchestration tool**
- **Handling Postgres clusters “the Kubernetes way” enables Cloud Native usage patterns**
 - Spinning up disposable clusters during tests
 - One cluster per microservice and one database per cluster
- **CloudNativePG attempts to encapsulate years of Postgres cluster management into an operator seeking to automate known management tasks**

CloudNativePG - How

```
# Example of a Basic PostgreSQL cluster
apiVersion: postgresql.cnpg.io/v1
kind: Cluster
metadata:
  name: cluster-example
spec:
  instances: 3

  storage:
    size: 1Gi
```



CloudNativePG Features

- Cloud Native
- Data persistent
- Designed for Kubernetes
- Security & TLS Certificates
- High availability
- Disaster Recovery
- Monitoring
- Advanced Architectures
- Postgres Operations

Demo

Takeaways

- ❖ Modern IT and DBA roles have evolved and integrated with DevOps
- ❖ Kubernetes is a tool to get familiarized with, and is no longer only for applications. K8s has matured to support databases
- ❖ Utilize
 - Kubernetes operator
 - Persistent volumes for data storage
- ❖ Automate as many tasks as possible
- ❖ Develop a monitoring strategy
- ❖ Test and validate your database backups

Resources

- **CloudNativePG Documentation**
 - <https://cloudnative-pg.io/documentation/1.19/>
- **edb-ansible**
 - <https://github.com/EnterpriseDB/edb-ansible>
- **Sandbox for CloudNativePG**
 - <https://github.com/EnterpriseDB/cnp-sandbox>
- **Setup Sandbox for CloudNativePG with edb-ansible**
 - https://github.com/dougortiz/edb-ansible/blob/master/playbook-examples/cnp/setup_cloudnativepg_sandbox.yml
- **Three ways DevOps Principles - Phoenix Project**
 - <https://www.techtarget.com/whatis/definition/The-Three-Ways>
- **The Phoenix Project**
 - <https://www.amazon.com/Phoenix-Project-DevOps-Helping-Business/dp/0988262592>

Q&A

THANK YOU