

COURSE CATALOG

POSTGRESQL & GIS COURSES



OUR COURSE PORTFOLIO

INTERMEDIATE COURSES

POSTGRESQL ADMINISTRATION & PERFORMANCE TUNING 6

ADVANCED COURSES

POSTGRESQL GEO-SPATIAL DATA MANAGEMENT WITH POSTGIS 10

POSTGRESQL HIGH AVAILABILITY ARCHITECTURE WITH PATRONI... 16

POSTGRESQL CONTAINERIZED DEPLOYMENT IN KUBERNETES..... 22

GENERAL INFORMATION

Elevate your skills by taking an expert-led PostgreSQL and GIS training session, designed to deepen your knowledge and empower you to tackle complex data challenges with confidence.

📍 LOCATIONS

- CYBERTEC HQ (Austria)
- Remote
- At your company

🌐 LANGUAGES

- English (EN)
- German (DE)

DIFFICULTIES

Level 

No prior knowledge required

Level 

Basic SQL and server administration knowledge required

Level 

Experienced PostgreSQL knowledge required

COURSE MATERIALS AND ENVIRONMENT

- We will provide access to a dedicated tab environment and virtual machines running a Linux distribution.
- Upon completion of the course, participants will receive a PDF of the slide deck that was used during the training sessions.
- All scripts and configuration files present on the virtual machines during training will be available for download and use.
- Upon completion of the course, participants will receive a complimentary copy of the most recent edition of Hans-Jürgens Schönig's book Mastering PostgreSQL.

CERTIFICATION

At the end of the course, you will receive a certificate to validate your knowledge.



SHORT INFO

Level:

Duration: 5 Days (6h/per day)
4 Days (8h/per day)

Language: EN, DE

Location: Online, on-site

TARGET AUDIENCE

- ▶ Database Administrators
- ▶ Developers
- ▶ IT Specialists

PREREQUISITES

- Basic knowledge of SQL
- Basic Linux skills (e.g., using commands and a text editor)

DESCRIPTION

This three-day training offers you a comprehensive introduction to the use of PostgreSQL, learning basic and advanced concepts from installation and architecture to high-availability setups and performance tuning. This training is aimed at database administrators, developers and IT specialists, wanting to deepen their knowledge and operate PostgreSQL securely in productive environments.

OBJECTIVES

- ✓ Understanding the PostgreSQL architecture and community
- ✓ Mastering basic administration, transaction management and backup strategies
- ✓ Database replication and its applications
- ✓ Efficient troubleshooting and tuning for optimal performance

TRAINER



Hans-Jürgen Schöning
CEO & Founder



Christoph Berg
Project Lead



Laurenz Albe
Senior Consultant

WHAT YOU WILL LEARN

- **PostgreSQL setup and operation:**
Installation, configuration and daily administration of PostgreSQL databases.
- **Backup and restore for secure operation:**
Implementation of logical and physical backups as well as point-in-time recovery for data restoration.
- **Efficient transaction and lock management:**
Handling transactions and isolation (optimistic and pessimistic locking) to ensure data integrity and optimal performance.
- **Performance tuning and monitoring:**
Using of EXPLAIN, parameter tuning and monitoring tools to optimise query performance and resource usage
- **Understanding PostgreSQL architecture and processes:**
Gaining knowledge of the core background processes of PostgreSQL and how their interaction to manage connections, storage and data.
- **Maintenance and optimisation:**
Application of VACUUM and Autovacuum to optimise database performance and disk space management.
- **Security and access management:**
Managing user rights, authentication methods and remote access to ensure database security.

These skills enable participants to run PostgreSQL efficiently, automate maintenance procedures, and optimize database performance and security.

MODULE 1: INTRODUCTION TO POSTGRESQL & BASIC ADMINISTRATION

- The PostgreSQL project, community and version management
- PostgreSQL architecture overview
- Installation, starting and stopping the server
- Connection to the ,psql' client
- Introduction to indexing (B-tree-indexes)
- First steps with transactions: atomicity, isolation and locking

MODULE 2: TRANSACTION MANAGEMENT & DATABASE MAINTENANCE

- Deepening transaction management: Optimistic and pessimistic locking mechanisms
- Multiversion Concurrency Control (MVCC)
- VACUUM and Autovacuum: database maintenance and fragmentation
- Backup strategies: logical backups, data export and import with pg_dump
- Physical backups and transaction log archiving
- Point-in-Time Recovery (PITR)

MODULE 3: HIGH AVAILABILITY & REPLICATION

- Streaming replication and synchronous replication
- Hot standby and conflict management
- Failover mechanisms for reliability
- Concept and use of Patroni for high availability solutions
- User administration and authentication
- Remote access and access rights (permissions)

MODULE 4: PERFORMANCE TUNING & SECURITY

- Configuration parameters and logging for optimised performance
- Resource and parameter tuning
- Security concepts and user rights
- Upgrade strategies
- Database monitoring: Statistics and monitoring
- Query optimisation with EXPLAIN (ANALYZE, BUFFERS) and extended statistics

SHORT INFO

Level: ■ ■ ■ ■

Duration: 2 Days (5h/per day)

Language: EN, DE

Location: Online

TARGET AUDIENCE

- ▶ Database Administrator
- ▶ GIS Analysts
- ▶ Developers

DESCRIPTION

In this comprehensive course, participants will learn the basics and advanced techniques of spatial data processing with PostGIS, a powerful extension for PostgreSQL that makes it possible to efficiently store, manage and analyse geographic information. The course combines theoretical knowledge with practical exercises to give participants a deep understanding of spatial data processing.

OBJECTIVES

- ✓ Successfully install and configure PostGIS
- ✓ Use and apply spatial data types and indices
- ✓ Use complex spatial functions to analyse and process data
- ✓ Gain practical experience in handling spatial data sets


TRAINER



Florian Nadler
GIS Consultant

PREREQUISITES

- **Basic knowledge of geodata:**
Participants should have a basic understanding of geographic information systems (GIS) and spatial data. This includes knowledge of different types of geospatial data, coordinate systems and spatial relationships.
- **Familiarity with PostgreSQL:**
A basic knowledge of PostgreSQL is required, especially with regard to database structures, queries and basic SQL skills. Participants should be familiar with the PostgreSQL database environment.
- **Recommendation: Participation in the course PostgreSQL Administration & Performance Tuning:**
It is recommended to complete the PostgreSQL Administration and Performance Tuning course beforehand. This ensures that participants have the necessary administrative skills and knowledge.





WHAT YOU WILL LEARN

- **Proficiency in the installation and maintenance of PostGIS:**
Students will be able to install, configure and regularly maintain PostGIS in PostgreSQL environments to ensure system performance and stability.
- **Competence in using spatial data types and indices:**
Ability to identify and efficiently use various spatial data types, such as Geometry and Geography, and implement GIST, SP-GIST, and BRIN indexes, to optimise query speed.
- **Demonstrate advanced skills in the application of spatial functions:**
Master the use of a variety of spatial functions for data analysis, including topological relationships, measurement functions and special functions for geospatial data processing.
- **Practical experience in processing and validating spatial data:**
Ability to import, validate and clean spatial data to ensure that data integrity and quality is maintained.
- **Competence to analyse and visualise spatial data:**
Learn to analyse spatial data and effectively present the results, using appropriate visualisation techniques to gain insights from the data and communicate these clearly.
- **Architecture and functionality of PostGIS:**
Gain knowledge of the internal architecture of PostGIS, including the integration of spatial data types and indexes into the PostgreSQL database, and their interaction with other GIS tools and applications.
- **Application of spatial functions for data analysis:**
Become familiar with the various spatial functions available in PostGIS and understand how these functions can be used to solve complex geographic problems.

MODULE 1: INTRODUCTION TO POSTGIS AND SPATIAL DATA PROCESSING

- Introduction to PostGIS and its importance in spatial data processing
- Installation and configuration of PostGIS in a PostgreSQL environment
- Overview of tools for import, export and ETL (Extract, Transform, Load)

MODULE 2: SPATIAL DATA TYPES, INDICES AND PROJECTIONS

- Introduction to the various spatial data types:
 - Geometry (2D and 3D)
 - Geography (geodetic data)
 - Raster
- Introduction to projections and transformations:
 - Using functions such as ST_Transform, to project spatial data into different coordinate systems.
 - Introduction to the use of SRIDs (Spatial Reference IDs) for the definition of projections
- Understanding and implementation of spatial indices:
 - GIST (Generalised Search Tree)
 - SP-GIST and BRIN (Block Range Index)

MODULE 3: RASTER DATA PROCESSING

- Introduction to the raster data type and its properties
- Dealing with raster files:
 - Importing raster data (e.g. GeoTIFF)
 - Management of raster data in PostGIS
 - Processing of raster data

MODULE 4: SPATIAL FUNCTIONS

- Introduction to the core functions of PostGIS:
 - Output, construction, access and setter functions
 - Measurement, composition, decomposition and simplification functions
- Application of topological functions:
 - Bounding boxes, equality checks and spatial relationships
 - Spatial joins for linking data sets
- Special functions for processing and analysis:
 - Spatial aggregation, clipping, splitting, tessellation, segmentation, as well as translation, scaling and rotation

MODULE 5: PRACTICAL APPLICATIONS AND EXERCISES

- Introduction to the data set and practical exercises:
 - Import and export of spatial and raster data
 - Validation and correction of spatial and raster data
 - Analysis of spatial data to gain insights

OPTIONAL, NOT INCLUDED BY DEFAULT

MODULE 6: CUSTOM DATA SETS

- Bring your own data for
 - Import
 - Analysis
 - Visualisation

SHORT INFO

Level: ■ ■ ■

Duration: 3 Days (5h/per day)

Language: EN, DE

Location: Online, on-site

TARGET AUDIENCE

- ▶ DB Administrators
- ▶ DevOps Engineers
- ▶ IT Professionals

PREREQUISITES

- Basic knowledge of SQL
- Basic Linux skills (e.g., using commands and a text editor)
- In addition the following knowledge is essential to be able to comprehend the content of the course „Administration and Performance Tuning“

DESCRIPTION

This course provides a comprehensive introduction to implementing High Availability (HA) in PostgreSQL with Patroni. Participants will learn how to create, manage and operate highly available PostgreSQL clusters using Patroni, etcd, pgBackRest and the VIP Manager for managing read and write access to the current primary server.

OBJECTIVES

- ✓ Understanding the concept and individual solutions required for the cluster
- ✓ Create highly available PostgreSQL clusters with Patroni, etcd and pgBackRest
- ✓ Manage Patroni clusters in operation and control failover processes
- ✓ Ensure that databases are operated in a highly available, fail-safe and efficient manner

TRAINER



Julian Markwort
Sr. Database Support
Engineer



Ants Aasma
Lead Database
Consultant

WHAT YOU WILL LEARN

- **Demonstrate technical understanding of the Patroni architecture and its components:** Deep understanding of how Patroni, etcd, PostgreSQL operate and interact to ensure high availability.
- **Become experienced in setting up and operating a Patroni cluster:** Hands-on experience setting up and managing a Patroni cluster, including configuring etcd, Patroni, and PostgreSQL.
- **Manage failover and reliability:** Mastering Patroni's automatic failover mechanisms to ensure operations, even during primary database failures.
- **Monitor and maintain highly available PostgreSQL environments:** Ability to monitor the health of the Patroni cluster, troubleshoot issues and perform regular maintenance to ensure a stable and resilient environment.
- **Knowledge of Patroni and its dependencies:** Understanding the architecture of Patroni and its dependencies, including etcd for cluster management and PostgreSQL for the database layer.
- **Best practices for monitoring and maintaining HA environments:** Learn best practices for monitoring and maintaining high-availability PostgreSQL environments to minimise downtime and optimise performance.

MODULE 1: INTRODUCTION TO HIGH AVAILABILITY WITH POSTGRESQL

- Importance and requirements of high availability (HA)
- Overview of Patroni as an HA solution
- Basics of cluster architecture

MODULE 2: BUILDING BLOCKS OF POSTGRESQL CLUSTERS

- Transaction Log (WAL) for Crash Recovery
- Backups to create replica PostgreSQL instances
- Configuring and monitoring PostgreSQL streaming replication

MODULE 3: PATRONI AND ITS COMPONENTS

- How Patroni works and its architecture
- Integration of etcd as a consensus mechanism
- Cluster management, leader, election, failover

MODULE 4: STRUCTURE AND CONFIGURATION OF A PATRONI CLUSTER

- Step-by-step instructions for setting up a Patroni cluster
- Configuration of etcd and PostgreSQL for HA
- Bootstrap a new Patroni cluster

MODULE 5: CLUSTER ADMINISTRATION

- Troubleshooting and monitoring Patroni clusters
- Runtime reconfiguration (PostgreSQL parameters, pg_hba)
- Cluster management (switchover, reinitialisation)
- Configuring failover and synchronous replication

MODULE 6: TANGENTIAL TOPICS

- Client connection routing
- Configuration and maintenance tasks
- Updates and upgrades
- Backups and restores
- Security and attack vectors of the patroni and etcd stack

OPTIONAL, NOT INCLUDED BY DEFAULT

MODULE 7: BACKUP AND RESTORE WITH PGBACKREST

- Introduction to pgBackRest and its functions
- Implementation of backup strategies
- Carrying out data restores and emergency procedures

MODULE 8: HAPROXY FOR LOAD BALANCING AND TRAFFIC MANAGEMENT

- Introduction to HAProxy and its role in highly available architectures
- Configuration of HAProxy for PostgreSQL clusters
- Load balancing and improving performance with HAProxy

MODULE 9: MONITORING AND MAINTENANCE OF HA ENVIRONMENTS

- Best practices for monitoring Patroni clusters
- Maintenance strategies to ensure stability and performance
- Identifying and resolving problems in HA setups



SHORT INFO

Level: ■ ■ ■ ■

Duration: 3 Days (5h/per day)

Language: EN

Location: Online

TARGET AUDIENCE

- ▶ DB Administrators
- ▶ DevOps Engineers
- ▶ System Administrators
- ▶ Software Developers
- ▶ IT Architects

DESCRIPTION

In this course, participants will learn how to deploy and manage PostgreSQL databases in a containerised Kubernetes environment. The training provides a comprehensive introduction to the concepts of Kubernetes, including container orchestration and architecture. Participants will become familiar with the setup and management of Kubernetes environments, the use of storage volumes and the implementation of the Postgres Operator and Patroni to efficiently manage PostgreSQL instances. Through practical exercises and examples, participants will acquire the necessary skills to successfully run PostgreSQL in Kubernetes.

OBJECTIVES

- ✓ Understanding the Kubernetes architecture
- ✓ Setting up and managing a Kubernetes development-environment
- ✓ Implement the PostgreSQL Operator
- ✓ Use of Patroni for cluster management

TRAINER



Ants Aasma
Lead Database Consultant

PREREQUISITES

- Basic knowledge of SQL
- Basic Linux skills (e.g., using commands and a text editor)
- Completion of the „PostgreSQL Administration and Performance Tuning“ course is required to fulfill this prerequisite
- Familiarity with Patroni: A basic understanding of Patroni and its role in ensuring the high availability of PostgreSQL is required.
- Completion of the „PostgreSQL High Availability Architecture with Patroni“ course fulfills this prerequisite
- Knowledge of container technologies: A foundational understanding of the concepts of containerization
- As containers and associated technologies are not PostgreSQL-specific, preparation on these areas are recommended



WHAT YOU WILL LEARN

- **Understanding Kubernetes architecture:**
Participants will develop a solid understanding of the Kubernetes architecture and container orchestration concepts
- **Setting up and managing Kubernetes environments:**
Participants will gain the skills to effectively set up and manage Kubernetes environments, including the use of tools such as Minikube and kubectl
- **Efficiently managing PostgreSQL with PostgreSQL Operator:**
Participants will learn how to configure and manage PostgreSQL Operator to efficiently run PostgreSQL databases in Kubernetes
- **Implementing Patroni for High Availability:**
Participants will be able to implement Patroni in Kubernetes to manage PostgreSQL clusters, including scaling, backups and configuration changes
- **Dealing with storage solutions in Kubernetes:**
Participants will learn how to deal with different storage options in Kubernetes, including persistent volumes and their configuration for PostgreSQL
- **Kubernetes Fundamentals:** Participants will understand the basic concepts of Kubernetes, including architecture, containerisation and orchestration
- **Administration with PostgreSQL Operator:**
Participants will have a comprehensive knowledge of how to configure and administer PostgreSQL Operator to manage PostgreSQL databases in Kubernetes



MODULE 1: INTRODUCTION TO KUBERNETES

- Containers: Understanding containers and their importance in software development
- Kubernetes architecture: overview of the main components and their functions, including resource management and networking
- Container orchestration: introduction to container orchestration and how the reconciliation loop works in Kubernetes

MODULE 2: SETTING UP THE KUBERNETES ENVIRONMENT

- Installation Methods and Minikube: various methods for installing Kubernetes, including the use of Minikube for local development
- Kubectl and CLI Tools: using kubectl and other CLI tools to manage Kubernetes resources
- Pods, ConfigMap, and Secrets: fundamentals of pod management and the use of ConfigMaps and Secrets for configuration and sensitive data

MODULE 3: STORAGE MANAGEMENT IN KUBERNETES

- Storage Volumes: introduction to different types of storage volumes, including emptyDir and HostPath
- Persistent Volumes and Claims: managing persistent storage solutions and provisioning them dynamically or statically

MODULE 4: POSTGRESQL OPERATOR

- Concept and configuration: introduction to the Postgres Operator and its customisation options
- Administration: management and monitoring of PostgreSQL instances, including backup and recovery strategies

MODULE 5: PATRONI IN KUBERNETES

- High Availability: understanding how Patroni works and how to use it to ensure availability.
- Database management: administration of PostgreSQL databases and roles as well as backup strategies (logical and continuous)



YOUR TRAINERS



**Hans-Jürgen
Schöning**
CEO & Founder



Ants Aasma
Lead Database
Consultant



Laurenz Albe
Senior Consultant



Christoph Berg
Project Lead



Julian Markwort
Sr. Database
Support



Florian Nadler
GIS Consultant