

Read Online Postgresql Documentation 91

Right here, we have countless books **Postgresql Documentation 91** and collections to check out. We additionally present variant types and along with type of the books to browse. The up to standard book, fiction, history, novel, scientific research, as competently as various further sorts of books are readily within reach here.

As this Postgresql Documentation 91, it ends happening being one of the favored book Postgresql Documentation 91 collections that we have. This is why you remain in the best website to look the unbelievable ebook to have.

BOG03Z - SINGH SOSA

The Oracle Solaris DTrace feature revolutionizes the way you debug operating systems and applications. Using DTrace, you can dynamically instrument software and quickly answer virtually any question about its behavior. Now, for the first time, there's a comprehensive, authoritative guide to making the most of DTrace in any supported UNIX environment—from Oracle Solaris to OpenSolaris, Mac OS X, and FreeBSD. Written by key contributors to the DTrace community, DTrace teaches by example, presenting scores of commands and easy-to-adapt, downloadable D scripts. These concise examples generate answers to real and useful questions, and serve as a starting point for building more complex scripts. Using them, you can start making practical use of DTrace immediately, whether you're an administrator, developer, analyst, architect, or support professional. The authors fully explain the goals, techniques, and output associated with each script or command. Drawing on their extensive experience, they provide strategy suggestions, checklists, and functional diagrams, as well as a chapter of advanced tips and tricks. You'll learn how to Write effective scripts using DTrace's D language Use DTrace to thoroughly understand system performance Expose functional areas of the operating system, including I/O, filesystems, and protocols Use DTrace in the application and database development process Identify and fix security problems with DTrace Analyze the operating system kernel Integrate DTrace into source code Extend DTrace with other tools This book will help you make the most of DTrace to solve problems more quickly and efficiently, and build systems that work faster and more reliably.

Learn the Ansible automation technology with some real-life examples. Every successful IT department needs automation nowadays for bare metal servers, virtual machines, cloud, containers, and edge computing. Automate your IT journey with Ansible automation technology. I'm going to teach you example by example how to accomplish the most common Dabatse Administrator tasks. Each of the lessons summarizes a module: from the most important parameter to some live demo of code and real-life usage. Each code is battle proved in the real life. Console interaction and verification are included in every lesson. Mundane activities like installing the database management system in Red Hat (Red Hat Enterprise Linux, CentOS, Amazon Linux, AlmaLinux, Rocky Linux, etc.) and Debian (Ubuntu, Scientific Linux) like systems, creating a database file, creating a user, granting the permission to a user to connect to the database, backup ad restore a single database could be automated with some lines of code and these are only some of the long lists included in the course. There are some Ansible codes usable in all the Linux systems, some specific for RedHat-like, Debian-like systems. The Ansible troubleshooting lesson teaches you how to read the error message, how to reproduce it, and the process of troubleshooting and resolution. Are you ready to automate your day with Ansible? Examples in the book are tested with the latest version of Ansible 2.9+ and Ansible Core 2.11+.

Starting your first project with Spring Boot can be a bit daunting given the vast options that it provides. This book will guide you step-by-step along the way to be a Spring Boot hero in no time. The book covers: * Setup of your project * Security and user management for your application * Writing REST endpoints * Connecting with a database from your application * Unit and integration testing for all aspects * Writing documentation for your REST endpoints * Support file upload from your REST API

The open source PostgreSQL database is soaring in popularity, as thousands of database and web professionals discover its powerful features, transaction support, performance, and industrial-strength scalability. In this book, a founding member of the PostgreSQL development team introduces everything you need to know to succeed with PostgreSQL, from basic SQL commands through database administration and optimization. PostgreSQL assumes no previous database expertise: it establishes a firm foundation of basic concepts and commands before turning to PostgreSQL's advanced, innovative capabilities. Bruce Momjian walks readers step-by-step from their first database queries through the complex queries needed to solve real-world problems. He presents proper query syntax, then explores the value and use of each key SQL commands in working applications. Learn to manipulate and update databases, customize queries, work with SQL aggregates, use joins, combine SELECTs with subqueries, work with triggers and transactions, import and export data, use PostgreSQL query tools, and more. Discover PostgreSQL techniques for server-side programming and multi-user control, and master PostgreSQL's interfaces to C, C++, ODBC, JDBC, Perl, and Tcl/TK. You'll also find detailed coverage of PostgreSQL administration, including backups, troubleshooting, and access configuration.

This book constitutes the refereed proceedings of the 11th International Conference on Data Warehousing and Knowledge Discovery, DaWak 2009 held in Linz, Austria in August/September 2009. The 36 revised full papers presented were carefully reviewed and selected from 124 submissions. The papers are organized in topical sections on data warehouse modeling, data streams, physical design, pattern mining, data cubes, data mining applications, analytics, data mining, clustering, spatio-temporal mining, rule mining, and OLAP recommendation.

This book offers a collection of papers presented at the First International Conference "Information Technologies in the Research of Biodiversity" that was held from 11-14 September 2018 in Irkutsk (Russia). Papers in this book cover areas of interaction of knowledge on biodiversity and information technologies. The main topics include: methods, models, software systems for the analysis of biodiversity data; global data portals; information and analytical systems on biodiversity; application of remote methods in vegetation mapping; theoretical fundamentals and organization technologies of the information and telecommunications infrastructures.

Welcome to the "PostgreSQL 8.4 Official Documentation - Volume V. Internals and Appendixes"! After many years of development, PostgreSQL has become feature-complete in many areas. This release shows a targeted approach to adding features (e.g., authentication, monitoring, space reuse), and adds capabilities defined in the later SQL standards.

A comprehensive survey of the foundational models and recent research trends in access control models and mechanisms for database management systems.

Get the definitive guide on Gatsby, the JavaScript framework for building blazing fast websites and applications. Used by Nike, Costa Coffee, and other companies worldwide, Gatsby is emerging as one of the key technologies in the Jamstack (JavaScript, APIs, and markup) ecosystem. With this comprehensive guide, you'll learn how to architect, build, and deploy Gatsby sites independently or with CMSs, commerce systems, and other data sources. Author Preston So begins by showing you how to set up a Gatsby site from scratch. From there, you'll learn ways to use Gatsby's declarative ren-

dering and GraphQL API, build complex offline-enabled sites, and continuously deploy Gatsby sites on a variety of platforms, including Gatsby Cloud. Discover how Gatsby integrates with many data sources and plug-ins Set up, configure, and architect Gatsby sites using Gatsby's CLI, React, JSX, and GraphQL with high performance out of the box Build an independent Gatsby site based on Markdown and data- and content-driven Gatsby sites that integrate with CMSs and commerce platforms Deploy Gatsby sites with full CI/CD and test coverage on a variety of platforms, including Netlify, Vercel, and Gatsby Cloud

Achieve awesome user experiences and performance with simple, maintainable code! Embrace the full stack of web development, from styling with Bootstrap, building an interactive user interface with Angular 4, to storing data quickly and reliably in PostgreSQL. With this fully revised new edition, take a holistic view of full-stack development to create usable, high-performing applications with Rails 5.1. Rails is a great tool for building web applications, but it's not the best at everything. Embrace the features built into your database. Learn how to use front-end frameworks. Seize the power of the application stack through Angular 4, Bootstrap, and PostgreSQL. When used together, these powerful and easy-to-use tools will open you to a new world of possibilities. This second edition is updated to cover Angular - a completely reworked front-end framework - and dives into new Postgres 9.6 features such as UPSERT. Also new is Webpack coverage, to develop the front-end code for your Rails application. Create a usable and attractive login form using Bootstrap's styles, while ensuring the database table backing it is secure using Postgres' check constraints. See how creating an advanced Postgres index for a case-insensitive search speeds up your back end - enabling you to create a dynamic user experience using Angular 4. Create reusable components that bring Bootstrap and Angular together and effectively use materialized views for caching within Postgres. Get your front end working with Webpack, use Postgres' features from migrations, and write unit tests for all of it. All of this within Rails 5.1. You'll gain the confidence to work at every level of the application stack, bringing the right solution to every problem. What You Need: This book covers Postgres 9.5, Rails 5, and Ruby 2.3. You should have some experience with basic Rails concepts and a cursory understanding of JavaScript, CSS, and SQL, but by no means need to be an expert. You'll learn how to install Postgres on your computer or use a free version of it in the cloud.

Thinking of migrating to PostgreSQL? This clear, fast-paced introduction helps you understand and use this open source database system. Not only will you learn about the enterprise class features in versions 9.2, 9.3, and 9.4, you'll also discover that PostgreSQL is more than a database system—it's also an impressive application platform. With examples throughout, this book shows you how to achieve tasks that are difficult or impossible in other databases. This second edition covers LATERAL queries, augmented JSON support, materialized views, and other key topics. If you're a current PostgreSQL user, you'll pick up gems you may have missed before. Learn basic administration tasks such as role management, database creation, backup, and restore Apply the psql command-line utility and the pgAdmin graphical administration tool Explore PostgreSQL tables, constraints, and indexes Learn powerful SQL constructs not generally found in other databases Use several different languages to write database functions Tune your queries to run as fast as your hardware will allow Query external and variegated data sources with foreign data wrappers Learn how use built-in replication filters to replicate data

This book constitutes the refereed proceedings of the 5th International Workshop on Ambient Assisted Living, IWAAL 2013, held in Carillo, Costa Rica, in December 2013. The 13 full papers and 7 short papers were carefully reviewed and selected from 23 submissions. The focus of the papers is on following topics: AAL environments, user interaction and assistive solutions, sensing and activity recognition and key application domains.

Arguably the most capable of all the open source databases, PostgreSQL is an object-relational database management system first developed in 1977 by the University of California at Berkeley. In spite of its long history, this robust database suffers from a lack of easy-to-use documentation. Practical PostgreSQL fills that void with a fast-paced guide to installation, configuration, and usage. This comprehensive new volume shows you how to compile PostgreSQL from source, create a database, and configure PostgreSQL to accept client-server connections. It also covers the many advanced features, such as transactions, versioning, replication, and referential integrity that enable developers and DBAs to use PostgreSQL for serious business applications. The thorough introduction to PostgreSQL's PL/pgSQL programming language explains how you can use this very useful but under-documented feature to develop stored procedures and triggers. The book includes a complete command reference, and database administrators will appreciate the chapters on user management, database maintenance, and backup & recovery. With Practical PostgreSQL, you will discover quickly why this open source database is such a great open source alternative to proprietary products from Oracle, IBM, and Microsoft.

A comprehensive guide to understanding key techniques for architecture and hardware planning, monitoring, replication, backups, and decoupling Key Features Newly updated edition, covering the latest PostgreSQL 12 features with hands-on industry-driven recipes Create a PostgreSQL cluster that stays online even when disaster strikes Learn how to avoid costly downtime and data loss that can ruin your business Book Description Databases are nothing without the data they store. In the event of an outage or technical catastrophe, immediate recovery is essential. This updated edition ensures that you will learn the important concepts related to node architecture design, as well as techniques such as using repmgr for failover automation. From cluster layout and hardware selection to software stacks and horizontal scalability, this PostgreSQL cookbook will help you build a PostgreSQL cluster that will survive crashes, resist data corruption, and grow smoothly with customer demand. You'll start by understanding how to plan a PostgreSQL database architecture that is resistant to outages and scalable, as it is the scaffolding on which everything rests. With the bedrock established, you'll cover the topics that PostgreSQL database administrators need to know to manage a highly available cluster. This includes configuration, troubleshooting, monitoring and alerting, backups through proxies, failover automation, and other considerations that are essential for a healthy PostgreSQL cluster. Later, you'll learn to use multi-master replication to maximize server availability. Later chapters will guide you through managing major version upgrades without downtime. By the end of this book, you'll have learned how to build an efficient and adaptive PostgreSQL 12 database cluster. What you will learn Understand how to protect data with PostgreSQL replication tools Focus on hardware planning to ensure that your database runs efficiently Reduce database resource contention with connection pooling Monitor and visualize cluster activity with Nagios and the TIG (Telegraf, InfluxDB, Grafana) stack Construct a robust software stack that can detect and avert outages Use

multi-master to achieve an enduring PostgreSQL cluster. Who this book is for: This book is for PostgreSQL administrators and developers who are looking to build and maintain a highly reliable PostgreSQL cluster. Although knowledge of the new features of PostgreSQL 12 is not required, a basic understanding of PostgreSQL administration is expected.

Welcome to the "PostgreSQL 8.4 Official Documentation - Volume IV. Reference"! After many years of development, PostgreSQL has become feature-complete in many areas. This release shows a targeted approach to adding features (e.g., authentication, monitoring, space reuse), and adds capabilities defined in the later SQL standards.

Develop programmatic functions to create powerful database applications. About This Book: Write complex SQL queries and design a robust database design that fits your application's need. Improve database performance by indexing, partitioning tables, and query optimizing. A comprehensive guide covering the advanced PostgreSQL concepts without any hassle. Who This Book Is For: If you are a PostgreSQL developer with a basic knowledge of PostgreSQL development and you want deeper knowledge to develop applications, then this book is for you. As this book does not cover basic installation and configurations, you should have PostgreSQL installed on your machine as a prerequisite. What You Will Learn: Write more complex queries with advanced SQL queries. Design a database that works with the application exactly the way you want. Make the database work in extreme conditions by tuning, optimizing, partitioning, and indexing. Develop applications in other programming languages such as Java and PHP. Use extensions to get extra benefits in terms of functionality and performance. Build an application that does not get locked by data manipulation. Explore in-built db functions and data type conversions. In Detail: PostgreSQL is the most advanced open source database in the world. It is easy to install, configure, and maintain by following the documentation; however, it's difficult to develop applications using programming languages and design databases accordingly. This book is what you need to get the most out of PostgreSQL. You will begin with advanced SQL topics such as views, materialized views, and cursors, and learn about performing data type conversions. You will then perform trigger operations and use trigger functions in PostgreSQL. Next we walk through data modeling, normalization concepts, and the effect of transactions and locking on the database. The next half of the book covers the types of indexes, constraints, and the concepts of table partitioning, as well as the different mechanisms and approaches available to write efficient queries or code. Later, we explore PostgreSQL Extensions and Large Object Support in PostgreSQL. Finally, you will perform database operations in PostgreSQL using PHP and Java. By the end of this book, you will have mastered all the aspects of PostgreSQL development. You will be able to build efficient enterprise-grade applications with PostgreSQL by making use of these concepts. Style and approach: Every chapter follows a step by step approach that first explains the concept, then shows you how to execute it practically so that you can implement them in your application.

Welcome to the "PostgreSQL 8.4 Official Documentation - Volume III. Server Programming"! After many years of development, PostgreSQL has become feature-complete in many areas. This release shows a targeted approach to adding features (e.g., authentication, monitoring, space reuse), and adds capabilities defined in the later SQL standards.

Emerging Trends in Computing, Informatics, Systems Sciences, and Engineering includes a set of rigorously reviewed world-class manuscripts addressing and detailing state-of-the-art research projects in the areas of Industrial Electronics, Technology & Automation, Telecommunications and Networking, Systems, Computing Sciences and Software Engineering, Engineering Education, Instructional Technology, Assessment, and E-learning. This book includes the proceedings of the International Joint Conferences on Computer, Information, and Systems Sciences, and Engineering (CISSE 2010). The proceedings are a set of rigorously reviewed world-class manuscripts presenting the state of international practice in Innovative Algorithms and Techniques in Automation, Industrial Electronics and Telecommunications.

This book is composed of a selection of articles from The 2021 World Conference on Information Systems and Technologies (WorldCIST'21), held online between 30 and 31 of March and 1 and 2 of April 2021 at Hangra de Heroismo, Terceira Island, Azores, Portugal. WorldCIST is a global forum for researchers and practitioners to present and discuss recent results and innovations, current trends, professional experiences and challenges of modern information systems and technologies research, together with their technological development and applications. The main topics covered are: A) Information and Knowledge Management; B) Organizational Models and Information Systems; C) Software and Systems Modeling; D) Software Systems, Architectures, Applications and Tools; E) Multimedia Systems and Applications; F) Computer Networks, Mobility and Pervasive Systems; G) Intelligent and Decision Support Systems; H) Big Data Analytics and Applications; I) Human-Computer Interaction; J) Ethics, Computers & Security; K) Health Informatics; L) Information Technologies in Education; M) Information Technologies in Radiocommunications; N) Technologies for Biomedical Applications.

Analyze data like a pro, even if you're a beginner. Practical SQL is an approachable and fast-paced guide to SQL (Structured Query Language), the standard programming language for defining, organizing, and exploring data in relational databases. Anthony DeBarros, a journalist and data analyst, focuses on using SQL to find the story within your data. The examples and code use the open-source database PostgreSQL and its companion pgAdmin interface, and the concepts you learn will apply to most database management systems, including MySQL, Oracle, SQLite, and others. * You'll first cover the fundamentals of databases and the SQL language, then build skills by analyzing data from real-world datasets such as US Census demographics, New York City taxi rides, and earthquakes from US Geological Survey. Each chapter includes exercises and examples that teach even those who have never programmed before all the tools necessary to build powerful databases and access information quickly and efficiently. You'll learn how to: Create databases and related tables using your own data. Aggregate, sort, and filter data to find patterns. Use functions for basic math and advanced statistical operations. Identify errors in data and clean them up. Analyze spatial data with a geographic information system (PostGIS). Create advanced queries and automate tasks. This updated second edition has been thoroughly revised to reflect the latest in SQL features, including additional advanced query techniques for wrangling data. This edition also has two new chapters: an expanded set of instructions on for setting up your system plus a chapter on using PostgreSQL with the popular JSON data interchange format. Learning SQL doesn't have to be dry and complicated. Practical SQL delivers clear examples with an easy-to-follow approach to teach you the tools you need to build and manage your own databases. * Microsoft SQL Server employs a variant of the language called T-SQL, which is not covered by Practical SQL.

Discover practical recipes to help you efficiently monitor enterprise IT infrastructure for Windows, Linux, and networking. Key Features: Find out how you can leverage some of the most exciting features of Zabbix. Perform professional IT infrastructure and application monitoring on multiple platforms. Discover easy-to-follow, practical solutions to problems in network monitoring with Zabbix. Book Description: Zabbix offers useful insights into your infrastructure performance and issues and enables you to enhance your monitoring setup with its variety of powerful features. This book covers hands-on, easy-to-follow recipes for using Zabbix 5 for effectively monitoring the performance of devices and applications over networks. The book starts by guiding you through the installation of Zabbix and using the Zabbix frontend. You'll then work your way through the most prominent features of

Zabbix and make the right design choices for building a scalable and easily manageable environment. The book contains recipes for building items and triggers for different types of monitoring, building templates, and using Zabbix proxies. As you advance, you'll learn how to use the Zabbix API for customization and manage your Zabbix server and database efficiently. Finally, you'll find quick solutions to the common and not-so-common problems that you may encounter in your everyday Zabbix monitoring work. By the end of this Zabbix book, you'll have learned how to use Zabbix for all your monitoring needs and be able to build a solid Zabbix setup by leveraging its key functionalities. What you will learn: Explore the different types of monitoring available in Zabbix. Find out how to build your own Zabbix templates. Use Zabbix proxies for effective load balancing/scaling. Work with custom integrations and the Zabbix API. Set up triggers and alerting with Zabbix. Maintain your Zabbix setup for scaling, backups, and upgrades. Discover how to perform advanced Zabbix database management. Monitor cloud-based products such as Amazon Web Services (AWS), Azure, and Docker. Who this book is for: This book is for IT engineers who want to get started with Zabbix and anyone with an intermediate understanding of Zabbix looking to extend their knowledge. Although not necessary, prior experience with Zabbix will help you to make the most of this book.

This book constitutes revised tutorial lectures of the 7th European Business Intelligence and Big Data Summer School, eBISS 2017, held in Bruxelles, Belgium, in July 2017. The tutorials were given by renowned experts and covered advanced aspects of business intelligence and big data. This summer school, presented by leading researchers in the field, represented an opportunity for postgraduate students to equip themselves with the theoretical, practical, and collaboration skills necessary for developing challenging business intelligence applications.

Python is a highly expressive language that makes it easy to write sophisticated programs. Combining high-quality geospatial data with Python geospatial libraries will give you a powerful toolkit for solving a range of geospatial programming tasks. The book begins with an introduction to geospatial analysis and programming and explains the ideas behind geospatial data. You will explore Python libraries for building your own geospatial applications. You will learn to create a geospatial database for your application using PostGIS and the psycopg2 library, and see how the Mapnik library can be used to create attractive and useful maps. Finally, you will learn to use the Shapely and NetworkX libraries to create, analyze, and manipulate complex geometric objects, before implementing a system to match GPS recordings against a database of roads to produce a heatmap of the most frequently used roads.

The numeric values retrieved from a data warehouse may be difficult for business users to interpret, and may even be interpreted incorrectly. Therefore, in order to better understand numeric values, business users may require an interpretation in meaningful, non-numeric terms. However, if the transition between non-numeric terms is crisp, true values cannot be measured and a smooth transition between classes may no longer be possible. This book addresses this problem by presenting a fuzzy classification-based approach for a data warehouse. Moreover, it introduces a modeling approach for fuzzy data warehouses that makes it possible to integrate fuzzy linguistic variables in a meta-table structure. The essence of this structure is that fuzzy concepts can be integrated into the dimensions and facts of an existing classical data warehouse without affecting its core. This allows a simultaneous analysis, both fuzzy and crisp. A case study of a movie rental company underlines and exemplifies the proposed approach.

This book explains relational theory in practice, and demonstrates through two projects how you can apply it to your use of PostgreSQL and SQLite databases. This book covers the important requirements of teaching databases with a practical and progressive perspective. This book offers the straightforward, practical answers you need to help you do your job. This hands-on tutorial/reference/guide to PostgreSQL and SQLite is not only perfect for students and beginners, but it also works for experienced developers who aren't getting the most from both databases. In designing a GUI and as an IDE, you will make use Qt Designer. In the first chapter, you will learn to use several widgets in PyQt5: Display a welcome message; Use the Radio Button widget; Grouping radio buttons; Displays options in the form of a check box; and Display two groups of check boxes. In chapter two, you will learn to use the following topics: Using Signal / Slot Editor; Copy and place text from one Line Edit widget to another; Convert data types and make a simple calculator; Use the Spin Box widget; Use scrollbars and sliders; Using the Widget List; Select a number of list items from one Widget List and display them on another Widget List widget; Add items to the Widget List; Perform operations on the Widget List; Use the Combo Box widget; Displays data selected by the user from the Calendar Widget; Creating a hotel reservation application; and Display tabular data using Table Widgets. In chapter three, you will learn: How to create the initial three tables project in the School database: Teacher, Class, and Subject tables; How to create database configuration files; How to create a Python GUI for inserting and editing tables; How to create a Python GUI to join and query the three tables. In chapter four, you will learn how to: Create a main form to connect all forms; Create a project will add three more tables to the school database: Student, Parent, and Tuition tables; Create a Python GUI for inserting and editing tables; Create a Python GUI to join and query over the three tables. In chapter five, you will join the six classes, Teacher, TClass, Subject, Student, Parent, and Tuition and make queries over those tables. In chapter six and chapter seven, you will get introduction of postgresql. And then, you will learn querying data from the postgresql using Python including establishing a database connection, creating a statement object, executing the query, processing the resultset object, querying data using a statement that returns multiple rows, querying data using a statement that has parameters, inserting data into a table using Python, updating data in postgresql database using Python, calling postgresql stored function using Python, deleting data from a postgresql table using Python, and postgresql Python transaction. In chapter eight, you will create dan configure PostgreSQL database. In this chapter, you will create Suspect table in crime database. This table has eleven columns: suspect_id (primary key), suspect_name, birth_date, case_date, report_date, suspect_status, arrest_date, mother_name, address, telephone, and photo. You will also create GUI to display, edit, insert, and delete for this table. In chapter nine, you will create a table with the name Feature_Extraction, which has eight columns: feature_id (primary key), suspect_id (foreign key), feature1, feature2, feature3, feature4, feature5, and feature6. The six fields (except keys) will have a VARCHAR data type (200). You will also create GUI to display, edit, insert, and delete for this table. In chapter ten, you will create two tables, Police and Investigator. The Police table has six columns: police_id (primary key), province, city, address, telephone, and photo. The Investigator table has eight columns: investigator_id (primary key), investigator_name, rank, birth_date, gender, address, telephone, and photo. You will also create GUI to display, edit, insert, and delete for both tables. In chapter eleven, you will create two tables, Victim and Case_File. The Victim table has nine columns: victim_id (primary key), victim_name, crime_type, birth_date, crime_date, gender, address, telephone, and photo. The Case_File table has seven columns: case_file_id (primary key), suspect_id (foreign key), police_id (foreign key), investigator_id (foreign key), victim_id (foreign key), status, and description. You will create GUI to display, edit, insert, and delete for both tables as well.

Welcome to the PostgreSQL 8.4 Official Documentation - Volume I. The SQL Language! After many years of development, PostgreSQL has become feature-complete in many areas. This release shows a targeted approach to adding features (e.g., authentication, monitoring, space reuse), and adds capabilities defined in the later SQL standards.

Over 150 recipes to help you administer your PostgreSQL database more efficiently About This Book Get to grips with the capabilities of PostgreSQL 9.6 to administer your database more efficiently Monitor, tune, secure and protect your database A step-by-step, recipe-based guide to help you tackle any problem in PostgreSQL administration with ease Who This Book Is For This book is for system administrators, database administrators, data architects, developers, and anyone with an interest in planning for, or running, live production databases. This book is most suited to those who have some technical experience. What You Will Learn Implement PostgreSQL features for performance and reliability Harness the power of the latest PostgreSQL 9.6 features Manage open source PostgreSQL versions 9.5 and 9.6 on various platforms Discover advanced technical tips for experienced users Explore best practices for planning and designing live databases Select and implement robust backup and recovery techniques Explore concise and clear guidance on replication and high availability See the latest details on Logical Replication and Bi-Directional Replication In Detail PostgreSQL is a powerful opensource database management system; now recognized as the expert's choice for a wide range of applications, it has an enviable reputation for performance and stability. PostgreSQL provides an integrated feature set comprising relational database features, object-relational, text search, Geographical Info Systems, analytical tools for big data and JSON/XML document management. Starting with short and simple recipes, you will soon dive into core features, such as configuration, server control, tables, and data. You will tackle a variety of problems a database administrator usually encounters, from creating tables to managing views, from improving performance to securing your database, and from using monitoring tools to using storage engines. Recipes based on important topics such as high availability, concurrency, replication, backup and recovery, as well as diagnostics and troubleshooting are also given special importance. By the end of this book, you will have all the knowledge you need to run, manage, and maintain PostgreSQL efficiently. Style and approach This book takes a step-by-step, recipe-based approach, where each recipe focuses on a particular challenge faced by a PostgreSQL administrator while administering his/her database. Explained in a very easy to follow manner, every task is supported with best practices, tips and tricks. This book combines elementary theory from computer science with real-world challenges in global geodetic observation, based on examples from the Geodetic Observatory Wettzell, Germany. It starts with a step-by-step introduction to developing stable and safe scientific software to run successful software projects. The use of software toolboxes is another essential aspect that leads to the application of generative programming. An example is a generative network middleware that simplifies communication. One of the book's main focuses is on explaining a potential strategy involving autonomous production cells for space geodetic techniques. The complete software design of a satellite laser ranging system is taken as an example. Such automated systems are then combined for global interaction using secure communication tunnels for remote access. The network of radio telescopes is used as a reference. Combined observatories form coordinated multi-agent systems and offer solutions for operational aspects of the Global Geodetic Observing System (GGOS) with regard to "Industry 4.0".

A step-by-step guide to building microservices using Python and Docker, along with managing and orchestrating them with Kubernetes Key Features Learn to use Docker containers to create, operate, and deploy your microservices Create workflows to manage independent deployments on coordinating services using CI and GitOps through GitHub, Travis CI, and Flux Develop a REST microservice in Python using the Flask framework and Postgres database Book Description Microservices architecture helps create complex systems with multiple, interconnected services that can be maintained by independent teams working in parallel. This book guides you on how to develop these complex systems with the help of containers. You'll start by learning to design an efficient strategy for migrating a legacy monolithic system to microservices. You'll build a RESTful microservice with Python and learn how to encapsulate the code for the services into a container using Docker. While developing the services, you'll understand how to use tools such as GitHub and Travis CI to ensure continuous delivery (CD) and continuous integration (CI). As the systems become complex and grow in size, you'll be introduced to Kubernetes and explore how to orchestrate a system of containers while managing multiple services. Next, you'll configure Kubernetes clusters for production-ready environments and secure them for reliable deployments. In the concluding chapters, you'll learn how to detect and debug critical problems with the help of logs and metrics. Finally, you'll discover a variety of strategies for working with multiple teams dealing with different microservices for effective collaboration. By the end of this book, you'll be able to build production-grade microservices as well as orchestrate a complex system of services using containers. What you will learn Discover how to design, test, and operate scalable microservices Coordinate and deploy different services using Kubernetes Use Docker to construct scalable and manageable applications with microservices Understand how to monitor a complete system to ensure early detection of problems Become well versed with migrating from an existing monolithic system to a microservice one Use load balancing to ensure seamless operation between the old monolith and the new service Who this book is for This book is for developers, engineers, or software architects who are trying to move away from traditional approaches

for building complex multi-service systems by adopting microservices and containers. Although familiarity with Python programming is assumed, no prior knowledge of Docker is required.

Welcome to the "PostgreSQL 8.4 Official Documentation - Volume II. Server Administration"! After many years of development, PostgreSQL has become feature-complete in many areas. This release shows a targeted approach to adding features (e.g., authentication, monitoring, space reuse), and adds capabilities defined in the later SQL standards.

Leverage the power of SQL to perform geospatial analysis and increase your speed and efficiency working with a variety of spatial applications such as PostGIS and QGIS Key Features Follow along with actionable instructions with this practical guide Become well-versed in advanced spatial modeling and machine learning techniques Learn best practices for performing spatial analysis from an expert spatial data analyst Book Description Geospatial analysis is used in almost every industry to answer location-related questions. Combined with the power of SQL, which is becoming a popular choice for developers and analysts worldwide, this technology will help you to solve real-world spatial problems easily. This book shows you how to detect and quantify patterns in datasets through data exploration, data visualization, data engineering, and the application of analysis and spatial techniques. You'll begin by exploring the fundamentals of geospatial analysis and understand its importance along with vector and raster models, among other things. You'll then look at the framework for geospatial analysis using SQL and progress through the chapters to create a spatial database and analyze it. In the next part, you'll advance to learning about using SQL functions and building SQL queries. By the end of this book, you'll be able to make the most of open source libraries and frameworks such as PostGIS and QGIS for analyzing spatial information. What you will learn Understand geospatial fundamentals as a basis for learning spatial SQL Generate point, line, and polygon data with SQL Create geometry objects with WKT, WKB, and GeoJSON Use spatial data types to abstract and encapsulate spatial structures Work with open source GIS combined with plug-ins Visualize spatial data and expand QGIS functionality with Postgres Apply location data to leverage spatial analytics Perform single-layer and multiple-layer spatial analyses Who this book is for This book is for anyone looking to leverage their SQL knowledge to perform geospatial analysis. GIS analysts, data analysts, and data scientists with a basic understanding of both geospatial analysis and SQL will find this book useful.

Leverage the power of PostgreSQL 11 to build powerful database and data warehousing applications Key Features Monitor, secure, and fine-tune your PostgreSQL 11 database Learn client-side and server-side programming using SQL and PL/pgSQL Discover tips on implementing efficient database solutions Book Description PostgreSQL is one of the most popular open source database management systems in the world, and it supports advanced features included in SQL standards. This book will familiarize you with the latest features in PostgreSQL 11, and get you up and running with building efficient PostgreSQL database solutions from scratch. Learning PostgreSQL, 11 begins by covering the concepts of relational databases and their core principles. You'll explore the Data Definition Language (DDL) and commonly used DDL commands supported by ANSI SQL. You'll also learn how to create tables, define integrity constraints, build indexes, and set up views and other schema objects. As you advance, you'll come to understand Data Manipulation Language (DML) and server-side programming capabilities using PL/pgSQL, giving you a robust background to develop, tune, test, and troubleshoot your database application. The book will guide you in exploring NoSQL capabilities and connecting to your database to manipulate data objects. You'll get to grips with using data warehousing in analytical solutions and reports, and scaling the database for high availability and performance. By the end of this book, you'll have gained a thorough understanding of PostgreSQL 11 and developed the necessary skills to build efficient database solutions. What you will learn Understand the basics of relational databases, relational algebra, and data modeling Install a PostgreSQL server, create a database, and implement your data model Create tables and views, define indexes and stored procedures, and implement triggers Make use of advanced data types such as Arrays, hstore, and JSONB Connect your Python applications to PostgreSQL and work with data efficiently Identify bottlenecks to enhance reliability and performance of database applications Who this book is for This book is for you if you're interested in learning about PostgreSQL from scratch. Those looking to build solid database or data warehousing applications or wanting to get up to speed with the latest features of PostgreSQL 11 will also find this book useful. No prior knowledge of database programming or administration is required to get started.

This book is part of the PostgreSQL 9.0 documentation collection (up-to-date & full), published by Fultus Corporation. PostgreSQL 9.0 includes built-in, binary replication, and over a dozen other major features which will appeal to everyone from web developers to database hackers.

"PostgreSQL Developer's Handbook" provides a complete overview of the PostgreSQL database server and extensive coverage of its core features, including object orientation, PL/SQL, and the most important programming interfaces. The authors introduce the reader to the language and syntax of PostgreSQL and then move quickly into sophisticated programming topics.