Inhalt

Dirk Krautschick, Aiven 3

On the road to Oracle Diagnostics and Tuning Pack for PostgreSQL 4

Teresa Lopes, Ayden 4

PostgreSQL for Oracle DBAs - A walk in the park? 4

Daniel Westermann, dbi services 5

Breaking PostgreSQL - Learning by doing it wrong 5

Maurizio De Giorgi 5

Implementing NIST latest password guidelines on PostgreSQL and Oracle at CERN: How-To, Challenges and Opportunities 5

Laurenz Albe, CYBERTEC 6

Comparing the Oracle and PostgreSQL transaction systems (Talk in German) 6

Stefan Keller, Ostschweizer Fachhochschule 7

NL-to-SQL: Einsatzmöglichkeiten, Grenzen und Verbesserungspotential 7

Bernd Patolla, In&Out AG 8

Einführung von PostgreSQL in einer Oracle-"Bude" 8

Piotr Kurzynoga, Oracle 9

Open Source services & PostgreSQL on OCI - Live Demo 9

Andrzej Nowicki, CERN 9

Postgres vs. Oracle: How to use vector functionality to create a beer recommendation system? 10

Mike Dietrich, Oracle 10

Migration How-To for DBAs to Oracle Autonomous Database 10

Hands-On Lab: Upgrade to Oracle Database 23ai, and ensure Performance Stability 11

Martin Berger, Accenture 11

Die grosse Liebe oder nur ein Flirt? Auf ein erstes Date mit dem Oracle Enterprise Manager 24AI 12

Christian Kenel und Danya Mattersteig, die Mobiliar 12

Oracle Unified Audit Trail - ein Datenfriedhof? Nicht @Mobiliar 12

Clemens Bleile, dbi services 13

Solve Oracle issues / Oracle Probleme lösen 13

Dani Schnider, Callista 13

Auf der Suche nach Use Cases für Data Use Case Domains 14

Christoph Lutz, Swisscom 14

Beyond the Happy Path: Dark Corners, Edge Cases, and Disaster Areas 14

Gianni Ciolli, EDB 15

Modernize your application with Postgres 15
Bogdana-Ioana Hincean, Expersoft Systems 16
From Havoc to Harmony: The Power of OS Management Hub 16
Oliver Lemm, Hyand Solutions 16
Table Functions at it's best 17
Ilya Kosmodemiansky, Data Egret 17
So wird man ein PostgreSQL-DBA: vom Anfänger zum Fortgeschrittene. 17

Alain Fuhrer, Eidgenössisches Justiz- und Polizeidepartement EJPD

About

Alain Fuhrer has been working in various roles in the database field for more than 15 years. His career includes leadership positions in large corporations as well as several years in database consulting.

He has specialized in Oracle databases and was actively involved in the Oracle ACE community for several years before transitioning to his current role as Head of Dataservices at the Federal Department of Justice & Police

In this role, he is responsible for database services such as Oracle and PostgreSQL, as well as search engines like Elasticsearch and the messaging service Kafka. He holds a Bachelor in Computer Sience and an EMBA in General Management

Privately, Alain Fuhrer lives in the canton of Bern and spends his free time with his family or doing sports.

PostgreSQL & Oracle: Die beste Kombination statt Entweder-oder

Description

Das ISC-EJPD hat sich vor gut 6 Jahren dazu entschieden eine PostgreSQL only Strategie und keine neuen Anwendungen mehr auf Oracle zu betreiben. Dies im Wissen, dass es noch Jahren dauern wird bis alle Oracle Datenbanken abgelöst sind. Im 2024 haben wir uns jedoch aus den Erfahrungen von 5 Jahren entschieden zukünftig beide Plattform als strategisch zu erachten. In diesem Vortrag soll aufgezeigt werden, weshalb sowohl Oracle wie auch PostgreSQL eine Berechtigung im ISC-EJPD haben und dass sich die beiden Produkte auch wunderbar ergänzen können.

Dirk Krautschick, Aiven

About

Dirk Krautchick has been working as a solution architect at Aiven since November 2023. As a qualified computer scientist, he was the responsible database administrator and quality support engineer for a manufacturer of optimization software in the aviation sector for 9 years, where he built up his know-how in operating PostgreSQL or Oracle databases and application servers. In 2017 he decided to use his fascination for database technologies and his skills as a consultant in - among others - the financial and energy sector to continually deepen them. During this dynamic times, the enthusiasm and fascination for open source from the student time returned and became priority for him. Without losing the acquired Oracle database know-how PostgreSQL as a primary topic was expanded in terms of content, teams were recruited, customers were trained and targeted projects were carried out. In addition to many architectural consulting and performance tuning projects also migrations from Oracle to PostgreSQL are often the content of such projects. In order to implement the open source and PostgreSQL strategy even more aggressively, without losing focus on the technology a move technical sales in 2022 was inevitable - first as a sales engineer at EnterpriseDB and now as a solution architect at Aiven.

On the road to Oracle Diagnostics and Tuning Pack for PostgreSQL

Elevator Pitch

How to get a similar tooling like AWR, ASH,... everyone loves in Oracle also in PostgreSQL.

Description

With the Diagnostic and Tuning Pack, Oracle offers an excellent collection of tools and features for analyzing performance problems. Specifically, the Automatic Workload Repository (AWR) and Active Session History (ASH) are almost essential for the daily work of DBAs.

With the currently not uncommon migrations to PostgreSQL comes again and again the question about equivalent solutions for PostgreSQL databases.

This lecture shows how far one can get to the feature set of Oracle in this regard with PostgreSQL can approximate. Field reports should be given using examples from projects also evaluate the suitability for practical use.

Teresa Lopes, Ayden

About

Teresa has worked with databases for more than 8 years. After several years as an Oracle DBA, she transitioned to PostgreSQL and never looked back! Extensibility and the Community are the two main aspects she finds fascinating about PostgreSQL. Teresa is also part of the PGDay Lowlands organization. Teresa started her career far away from databases, as a Civil Engineer. So, if you've always wanted to know how tunnels and bridges are built, just ask! In her free time, Teresa enjoys spending time outdoors (mainly admiring rocks), hiking, and cooking.

PostgreSQL for Oracle DBAs - A walk in the park?

Elevator Pitch

My journey from Oracle to PostgreSQL DBA was not exactly a walk in the park! I will share the biggest surprises, what I missed, and how PostgreSQL finally won me over. Ideal talk for Oracle DBAs exploring PostgreSQL.

Description

Relational Databases are all one of the same, right? Same thing, different names, no?

Whilst the core concepts might still be the same between Oracle and PostgreSQL, transition from an Oracle DBA role to a PostgreSQL one might not be as straight forward.

In this talk we will explore some of the key differences between managing Oracle and PostgreSQL DBs, including some of the features I initially missed and the ones I never looked back to!

Daniel Westermann, dbi services

Breaking PostgreSQL - Learning by doing it wrong

Elevator Pitch

There is no better way to learn as learning from mistakes, isn't it? We're going to take that serious and I am sure you'll have a lot of fun and gain some interesting insights.

Description

Although PostgreSQL is a very reliable and robust database management system, there are things you shouldn't do. In this session we're breaking PostgreSQL in several ways, live, without any slides (slides will be provided anyway, don't worry). There is no better way to learn as learning from mistakes, isn't it? We're going to take that serious and I am sure you'll have a lot of fun and gain some interesting insights. By the end of this talk you should have a solid understanding of what not not do (and why not)and this will safe you quite some time in your journey with PostgreSQL.

Maurizio De Giorgi

About

I am currently working at CERN as Senior Database Engineer and Service Manager in the DataBase On Demand team which is part of the IT-DA (Data and Analytics) group. DataBase On Demand is a flexible DBaaS platform capable of supporting different technology, currently offering MySQL, PostgreSQL, TimescaleDB and InfluxDB, on VMs and bare metal puppet managed hosts, provisioned with Openstack.

I have a few decades of experience, mostly as IT consultant in different roles, industry, markets, with a strong focus on RDBMS development and administration, spanning across Data Warehousing, Business Intelligence, ETL and Application Integration, on Oracle, and more recently, using open source technology and DevOps tools.

I am an avid technologist always carefully looking at new technology, paradigms and trends and the best way to apply them within the projects and the organizations where I am working. In the last few years, I have also been on a journey to expand my knowledge around data science, machine learning and AI and new alternative data store technology like NoSQL, Graph, columnar and in-memory databases, virtualized and cloud computing, stream and event based processing.

Implementing NIST latest password guidelines on PostgreSQL and Oracle at CERN: How-To, Challenges and Opportunities

Elevator Pitch

Password policy reinforcement is coming *almost* out-of-the-box from Oracle but requires a different approach and more work in PostgreSQL. Beside showcasing both cases, we provide take-home techniques used to deal with challenges and opportunities arisen in our large Oracle estate and DBaaS platform

Description

The IT-DA-DB section at CERN is responsible for quite a few hundreds of PostgreSQL instances, provisioned via the DB on Demand service, and thousands of Oracle accounts, distributed across approximately one hundred consolidated Oracle databases.

The National Institute of Standards and Technology (NIST) is in charge of defining standards in many fields for US governmental organizations and federal agencies. The thorough review process, applied to validating the exhaustive research, and the extensive applicability have made NIST password security recommendations the benchmark for the whole IT industry.

In the recently published second draft of the 4th revision of NIST's Special Publication (SP) 800-63, Digital Identity Guidelines there is a clear shift from old-fashioned, rigid and overly complicated practices, towards more flexible, user-friendly solutions, whose efficacy has been studied and proven (cf. NIST Special Publication 800-63B), with the aim of simplifying password management, while improving security.

Implementing these internationally recognized best practices is of particular importance to mitigate the risks coming from an increasing number of password-cracking attempts.

In this talk we will:

- introduce the main changes and briefly discuss "What to do" and "What to stop"
- go through the steps required to implement (some of) these guidelines in PostgreSQL and Oracle
- highlight the challenges and the opportunities you might encounter and how we tackled and leveraged them in the CERN environment
- showcase how to use additional password cracking tools to periodically check for weak passwords

Laurenz Albe, CYBERTEC

About

I have been working with PostgreSQL since 2006 and am listed contributor to the project. I am the main author and maintainer of the Oracle Foreign Data Wrapper, a PostgreSQL extension that allows direct access to Oracle database from PostgreSQL.

Comparing the Oracle and PostgreSQL transaction systems (Talk in German)

Elevator Pitch

Both Oracle and PostgreSQL support transactions, so what? The difference is in the details. And these details are what causes problems when migrating or porting an application from one to the other - problems that may only occur in a highly concurrent production setting. Be prepared!

Description

Organized around the famous ACID properties of database transactions, this talk will explore the small, but crucial differences between Oracle and PostgreSQL. I will tap the experiences gained as PostgreSQL hacker and expert, developer of the Oracle Foreign Data Wrapper and participant in database migration projects.

Stefan Keller, Ostschweizer Fachhochschule

About

Stefan Keller ist Full Professor und Institutsleiter an der FH OST am Campus Rapperswil. Er gründete das Geometa Lab und unterrichtet Data Engineering (Datenbanksysteme) und Data Analytics sowie GISTech. Stefan Keller studierte Geographie und Informatik an der Universität Zürich und arbeitete unter anderem bei der Credit Suisse im ersten Usability Lab und bei Unisys in der GIS-Softwareentwicklung. Er engagiert sich in den Bereichen (Geo) Data Engineering und Analytics sowie (Open) Data Management und ist an mehreren innovativen Open Source und Open Data Projekten beteiligt. Stefan Keller ist u.a. Mitglied der Swiss PostgreSQL Users Group, der Swiss OpenStreetMap Association und ist Mitglied des Organisationskomitees z.B. des Swiss PGDay oder des Swiss Python Summit.

NL-to-SQL: Einsatzmöglichkeiten, Grenzen und Verbesserungspotential

Elevator Pitch

NL-to-SQL revolutioniert den Datenbankzugriff und fördert die Datennutzung, indem es natürliche Sprache in SQL-Abfragen übersetzt. Benutzer ohne SQL-Kenntnisse profitieren von einem einfacheren Datenzugriff, Experten sparen Zeit und es erleichtert die Vermittlung von SQL in der Aus-&Weiterbildung.

Description

“Natural Language to SQL” (NL-to-SQL) ermöglicht es Anwendern ohne SQL-Kenntnisse, Datenbanken in natürlicher Sprache abzufragen. Dabei nimmt ein Large-Language-Model (LLM) à la ChatGPT & Co. den natürlichsprachlichen Text der Abfrage entgegen und generiert daraus eine Abfrage in SQL.

Eine Abfrage wie “Zeige die Titel aller ‘Sci-Fi’-Filme, die im letzten Monat ausgeliehen wurden, sowie die Gesamtanzahl dieser Ausleihen” wird mit NL-to-SQL in folgende SQL-Abfrage umgewandelt: “SELECT f.title, COUNT(*) FROM rental r JOIN inventory i ON r.inventory_id = i.inventory_id JOIN film f ON i.film_id = f.film_id JOIN film_category fc ON f.film_id = fc.film_id JOIN category cat ON fc.category_id = cat.category_id WHERE cat.name = ‘Sci-Fi’ AND r.rental_date >= date_trunc(‘month’, current_date) - INTERVAL ‘1 month’ AND r.rental_date < date_trunc(‘month’, current_date) GROUP BY f.title ORDER BY anzahl_ausleihen DESC;”.

Wird das SQL auch ausgeführt, kann man auch von “Agentic AI” sprechen. Agentic AI zerlegt eine natürliche Spracheingabe in mehrere Schritte: 1. Verstehen der Anfrage (Intent-Parsing), 2. Abbildung auf die Datenbankschema-Informationen, 3. Generierung der SQL-Abfrage (strukturierte Ausgabe), 4. Ausführung und Rückgabe des Ergebnisses.

Diese Technologie bietet einen dreifachen Nutzen: Erstens erweitert sie für Anwender ohne SQL-Kenntnisse die Möglichkeiten, komplexe Datenbankabfragen zu erstellen. Dies führt zu einer breiteren Nutzung der Daten. Für versierte Anwender bietet sie zweitens eine Zeitersparnis, da sie die manuelle Erstellung komplexer SQL-Abfragen unterstützt und das Nachschlagen von Tabellen- und Attributnamen erleichtert. Ein dritter Aspekt ist schliesslich, dass es das Erlernen von SQL in der Aus- und Weiterbildung erleichtert.

Zunächst wird der Stand der Technik von NL-to-SQL-Werkzeugen beleuchtet. Dabei werden fünf Ansätze betrachtet: 1. reines LLM, 2. kontextbasiertes Lernen, 3. Fine Tuning, 4. Retrieval-Augmented Generation (RAG) und 5. strukturierte Ausgabe und schrittweises Reasoning. Reines Fine Tuning wurde aufgrund begrenzter Trainingsdaten ausgeschlossen. Mehrere LLMs, darunter Mistral und GPT-4o-mini, wurden mit einfachen und fortgeschrittenen SQL-Abfragen getestet.

Im Rahmen einer Bachelorarbeit von angehenden Informatikern der FH OST wurden zwei Datensätze verwendet, einer nach einem firmenspezifischen Warenhausschema und einer nach dem bekannten Pagila-Schema. Bewertet wurden Ähnlichkeit, Gültigkeit, Durchführbarkeit und Zuverlässigkeit. Bei komplexeren Abfragen waren nur 45% der von RAG generierten SQL-Abfragen ausführbar, wobei 70% der angeforderten Spalten extrahiert wurden. Llama3.2 zeigte das grösste Potenzial für zukünftige Entwicklungen. Die Herausforderung bestand darin, gültige und ausführbare SQL-Abfragen zu gewährleisten und mit Halluzinationen umzugehen. Das Ergebnis war daher noch nicht ganz zufriedenstellend. Unklar ist auch, wie NL-to-SQL in der Aus- und Weiterbildung vermittelt werden kann, damit die Lernenden das Grundverständnis von SQL nicht verlieren.

Weitere technische Verbesserungen sind zu erwarten, einerseits durch verbesserte Bewertungsmetriken und andererseits durch die Kombination der Ansätze, insbesondere RAG (4) und Structured Output and Step-by-Step Reasoning (5).

Mit diesem Beitrag soll eine Diskussion über den Nutzen von NL-to-SQL in den drei genannten Anwenderbereichen, Anwender ohne SQL-Kenntnisse, erfahrene SQL-Anwender sowie Aus- und Weiterbildung angeregt werden.

Bernd Patolla, In&Out AG

Einführung von PostgreSQL in einer Oracle-"Bude"

Elevator Pitch

In meinem Vortrag möchte ich den Weg vom Start bis zur heutigen Infrastruktur für PostgreSQL-Datenbanken aufzeigen. Welche Tools sind für das Management von PostgreSQL im Einsatz? Wie werden effiziente Backups auch für grosse Datenbanken realisiert? Wie hat sich die Infrastruktur entwickelt?

Description

Vor knapp 3 Jahren hat ein Kunde sich entschieden, PostgreSQL als Alternative zu Oracle Datenbanken zu betrachten. Im Sommer 2023 kam dann der Entscheid, PostgreSQL als strategisches Produkt und Oracle Datenbanken nur noch als Legacy-Produkt einzustufen. In diesem Vortrag möchte ich den Weg vom Start bis zur heutigen Infrastruktur für PostgreSQL-Datenbanken aufzeigen. Welche Entwicklungen hat die Infrastruktur genommen? Welche Tools

sind für das Management von PostgreSQL im Einsatz? Wie werden effiziente Backups auch für grosse Datenbanken realisiert? Auf diese und weitere Fragen möchte ich in meinem Vortrag eingehen und beantworten.

Piotr Kurzynoga, Oracle

About

- I'm an Electronic and Computer Engineer by trade, certified in Information and Network Security, you can find me practicing Brazilian Jiu-Jitsu in my spare time.
- Within Oracle I had worked with most of the core products in our portfolio with focus on Data Management, Migrations, Disaster Recovery and High Availability.
- Currently part of the Open-Source Data Platform team where I am integrating the open-source stack with Oracle technology.
- Last but not least I spend quite some time with APEX being both my hobby and daily work companion where I create blogs, content and assist with architectures.

Open Source services & PostgreSQL on OCI - Live Demo

Elevator Pitch

Hear from an ex-Oracle Data Management specialist about PostgreSQL solutioning on OCI complemented with a wholesome live demo.

Description

Speaker: Piotr Kurzynoga (CZ)

PostgreSQL offers a great set of constantly growing features and is often the go-to choice for many companies with rapid growth seen in recent years... that's great, but what is it driven by?

In this hands-on session let's explore the drivers of the PostgreSQL puzzle, how it fits into an architecture and integrates with other components for a performant & powerful solution all within Oracle Cloud Infrastructure.

What can you expect?

- PostgreSQL use-cases.

- Streaming data into PostgreSQL with Apache Spark.

- Text2SQL Chatbot with PostgreSQL.

Andrzej Nowicki, CERN

About

Andrzej Nowicki is a database engineer on the CERN database team, specializing in Oracle, MySQL, and PostgreSQL databases. Andrzej is an Oracle DBA with more than 11 years of experience, a passion for solving technical challenges, and a commitment to sharing his knowledge of database systems - making valuable contributions to CERN's objectives. Before CERN, he actively contributed to various companies and numerous IT projects as a database administrator.

Postgres vs. Oracle: How to use vector functionality to create a beer recommendation system?

Elevator Pitch

I will present how to create a similarity search system. Both Oracle and PostgreSQL can handle vector data, which I will explain, present and compare. The idea is to start with theory, then a demo of similarity search and finally show the concept of how RAGs operate. From zero to state of the art.

Description

Dive into the world of AI-driven recommendations with our session on crafting a beer recommendation system using language models and vector storage. This presentation will show the integration of machine learning into database workflows. We'll compare how Oracle Database and PostgreSQL enable scalable and efficient vector search. We'll go through the architecture, real-world use cases, and compare these technologies for performance and functionality. Perfect for database administrators eager to tap into AI, this session will equip you with the knowledge to blend structured data and cutting-edge ML techniques.

Mike Dietrich, Oracle

About

Based in Bavaria, Germany. more than 30 years IT experience. Holding a degree as Graduate Engineer (Dipl. Ing. NT/DT) in Data Systems and Information Technology of Georg-Simon-Ohm Technical University, Nürnberg, Germany.

Vice President Product Management and Develpment, Oracle Database Upgrade, Cloud Migrations and Patching, leading a global expert team of upgrade and migration specialists, being at Oracle for over 27 years. Interlink between customers/partners and the Oracle development org. Holding technical workshops between Arctic and Antarctica, and of course virtually too. Assisting customers in their reference projects onsite and remotely worldwide together with my team. Troubleshooting upgrade, migration and patching issues. Eager to learn something new every day. Experienced conference speaker. Music lover. LP/vinyl enthusiast. Mountain hiker. Gravel biker.

Migration How-To for DBAs to Oracle Autonomous Database

Elevator Pitch

Oracle Autonomous Database is not meant to make the DBA's obsolete but instead take over some simple database. But what is "simple"? And why is an ADB migration very different from any other migration? In 45 minutes you'll get the idea and all the knowledge you need.

Description

Oracle Autonomous Database gets a real push with the hybrid clouds. But a migration to Oracle ABD is different - and you should know about these differences.

This isn't a sales pitch but rather a talk for techies. Learn about the right approach, things you may need to take care on before, stuff which doesn't migrate to easily - but moreover, how to tackle it. Everybody can avoid the typical pitfalls - and you will learn about them quickly.

This is a How-To talk based on practical experience from real world projects.

Hands-On Lab: Upgrade to Oracle Database 23ai, and ensure Performance Stability

Elevator Pitch

Participants need to bring their laptop, Wifi or tethering must be possible. Power outlets and tables in the room would be good.

- run a workload
- collect performance figures
- upgrade your database
- tune the workload
- make 23ai run fast

Description

Participants will need to bring their laptop - the connection runs over noVNC within a browser. Nothing gets installed, but WiFi or telephone tethering are required.

Learn how to upgrade and migrate your database to Oracle Database 23ai. In this workshop, you learn how to use AutoUpgrade to make upgrade and non-CDB to PDB migrations much easier. Further, you can use our Performance Stability Perscription to learn how to avoid the biggest fear of a DBA: bad execution plans after the upgrade. Using tools like SQL Performance Analyzer and SQL Plan Management, you can ensure performance stability after any change - including upgrades and migrations. You will also try migrating a database to a PDB using different techniques. Finally, you learn how to restore after a failed upgrade and downgrade a PDB.

Martin Berger, Accenture

About

Martin Berger has been working in the Oracle technologies environment since 2000 and works as a data engineer at Accenture Switzerland. His main areas of focus include the design of Oracle Cloud infrastructures, the installation and maintenance of Oracle Database Appliances,

the development of highly available database architectures and enterprise monitoring with the Oracle Enterprise Manager. He also has a passion: the development of automation workflows with Oracle Linux Automation or the Red Hat Ansible Automation Platform 2.x. He has been working with the Oracle Cloud Infrastructure since the Oracle Cloud journey began. You can regularly find him at conferences or community meetings. He has adopted the cultural values from his company Trivadis, which is integrated into Accenture - curiosity, doership, network, freedom and togetherness!

Die grosse Liebe oder nur ein Flirt? Auf ein erstes Date mit dem Oracle Enterprise Manager 24AI

Elevator Pitch

Wie lief mein erstes Date mit dem neuen Oracle Enterprise Manager 24AI? Im Talk zeige ich das neue GUI, innovative Features und die nahtlose Integration in OCI, inkl. Highlights wie Operations Insights. Seien Sie gespannt, ob es eine dauerhafte Liebesbeziehung werden könnte... Live Demos inklusive.

Description

Wie lief mein erstes Date mit dem neuen Oracle Enterprise Manager 24AI? In diesem Talk nehme ich Sie mit auf eine spannende Entdeckungsreise durch die neuen Funktionen, das überarbeitete GUI und die innovative Integration mit Oracle Cloud Infrastructure (OCI).

Erfahren Sie, wie der Enterprise Manager 24AI Cloud-Features wie Operations Insights nutzt, um tiefere Einblicke in Ihre Datenbanken zu gewinnen, und wie neue KI-gestützte Tools den Betrieb erleichtern. Wir werfen einen Blick auf die Voraussetzungen, erläutern mögliche Migrationspfade und gehen auf wichtige technische Highlights ein, darunter:

- Microservice-Architektur
- Optimierte Überwachungs- und Automatisierungsfunktionen
- Verbesserte Performance-Dashboards und Berichterstattung
- Erweiterte Sicherheitsfunktionen

Live-Demos zeigen praxisnah, was der neue Oracle Enterprise Manager 24AI wirklich kann. Bleibt es bei einer sprichwörtlich oberflächigen Beziehung oder wird daraus eine dauerhafte Liebe fürs Leben? Finden Sie es heraus!

Christian Kenel und Danya Mattersteig, die Mobiliar

Oracle Unified Audit Trail - ein Datenfriedhof? Nicht @Mobiliar

Elevator Pitch

Die Mobiliar generiert 100+ Mio Audit-Records/Monat. Wir zeigen Dir, wie der manuelle Auswertungsprozess dank Dynatrace und Ticket-System automatisiert wurde, damit die “Abweichungen” direkt von den betroffenen Stellen bearbeitet werden.

Description

Die Mobiliar hat vor zirka fünf Jahren mit Oracle Unified Audit Trail begonnen und ist damit seit gut vier Jahren produktiv unterwegs. Von Anfang an war klar, dass wir nicht nur Daten aufzeichnen wollen, sondern dass diese Daten aufbereitet und ausgewertet werden. Vor der Produktivsetzung wurde firmenweit darüber informiert. Die manuellen Monats-Auswertungen wurden bislang auf der Datenbank gestartet und landeten in Form von Excel-Tabellen bei den betroffenen Stellen. Nur: Wer erinnert sich daran, was vor drei Wochen «geschehen» ist? Wir suchten nach einer Lösung, die auf Tagesbasis und Automatisierung aufbauen muss. Dynatrace kam uns dabei gelegen: Sie holen heute die Audit-Daten ab den Produktions-Datenbanken und aggregieren diese auf Tagesbasis. Wenn es “Abweichungen” gibt, wird eine Workorder im Ticket-System erstellt und an die zuständige Stelle adressiert. Diese Stelle bearbeitet das Ticket, zB. unter Angabe einer Change-Nummer oder “darf” Stellung beziehen, wie es zu diesen Audit-Records kam.

Clemens Bleile, dbi services

About

I have more than 30 years of IT experience, including 13 in Oracle Support and 10 in Oracle Consulting and specialized in Oracle Database Performance Tuning (SQL and DB Tuning) and in the development of Oracle DB IT architectures (highly available, low-maintenance, cost efficient storage of data). During my time at Oracle Support I became an expert in problem analysis and solving.

Solve Oracle issues / Oracle Probleme lösen

Elevator Pitch

Many organizations struggle to resolve technical issues quickly and permanently. As a technical guy, who has worked for decades in Support organizations and as a Consultant, I developed problem solving skills, which I want to present based on real world Oracle examples.

Description

Oracle Problem Solving

Many organizations struggle to resolve technical issues quickly and permanently. This is often related to the fact that DBAs these days have more broad than deep skills, but is often also also related to missing good problem solving skills in the different teams. Sometimes this may also be caused by outsourcing IT maintenance. As a technical guy, who has worked for decades in Support organizations and as a Consultant, I developed some problem solving skills, which I want to present. With examples of technical issues in the DB (e.g. a case with an ORA-00054 resource busy in an ETL job and Performance related problems), I tried to sum up a couple of good strategies to resolve a technical issue.

Dani Schnider, Callista

About

Dani Schnider works as Principal Oracle & DWH Consultant for Callista, a Swiss IT service provider for consulting and engineering. He has been working in database consulting for over 20 years, from 1999 to 2022 for Trivadis in Switzerland. Dani works with Oracle databases since 1994, since 1997 mainly in data warehouse projects.

His tasks include DWH architecture reviews, data modeling, ETL development and performance optimization in Oracle databases. Dani is Oracle ACE Director and co-author of the books Data Warehousing mit Oracle - Business Intelligence in der Praxis and Data Warehouse Blueprints, both published in German by Hanser-Verlag. On his blog Data Warehousing with Oracle he writes about design and performance topics related to Oracle data warehouses. He is a regular speaker at international Oracle user group conferences.

Auf der Suche nach Use Cases für Data Use Case Domains

Elevator Pitch

Was sind "Data Use Case Domains" in Oracle 23ai? Sind sie das Gleiche wie die SQL Domains, die es in PostgreSQL schon seit langem gibt? Wo liegen die Unterschiede, und in welchen Anwendungsfällen kann dieses neue Feature in Oracle-Datenbanken eingesetzt werden?

Description

Endlich gibt es Domains auch in Oracle! – "Data Use Case Domains" sind eine der grösseren SQL-Erweiterungen in Oracle 23ai. Doch warum dieser umständliche Name, und warum heissen sie nicht einfach SQL Domains wie in PostgreSQL und anderen RDBMS-Systemen? Wie mächtig und wie nützlich ist dieses neue Feature, und was sind typische Anwendungsfälle? Machen wir uns mit Domains einfach das (Entwickler-)Leben schwerer, oder helfen sie uns, den Aufwand bei Datenbankdesign und Applikationsentwicklung zu reduzieren? Mit zahlreichen Beispielen und Live-Demos wird auf spielerische Weise gezeigt, wie Domains in Oracle23ai funktionieren und welche Möglichkeiten sie bieten. Ziel der Präsentation ist es, dass die Teilnehmenden anschliessend eine Entscheidungsgrundlage haben, ob und in welche Fällen sie Domains einsetzen können und wollen.

Christoph Lutz, Swisscom

About

Christoph is a database architect at Swisscom and has been working with Oracle in various different roles for more than 20 years. He's also a member of Symposium42 and very passionate about understanding how things really work and loves to analyze and troubleshoot complex performance problems, hangs, and crashes.

Beyond the Happy Path: Dark Corners, Edge Cases, and Disaster Areas

Elevator Pitch

In IT, we often design for the “happy path,” but real-world systems are unpredictable. This session explores fascinating disaster areas, helping you learn from failure and better prepare for the unexpected.

Description

In IT, we often design for the “happy path” - the smooth, expected flow where everything works as intended.

However, real-world systems are messy, unpredictable, and full of hidden pitfalls and while the Oracle database is an incredibly reliable and mature piece of software, unexpected interactions, tricky edge cases, and unforeseen failures can still derail even the most well-designed system and application.

Drawing from from real-world incidents, this session explores the dark corners of database operations where performance bottlenecks emerge, data corruption strikes, and full-blown system failures unfold.

Learn how to identify, mitigate, and recover from these challenges before they bring your system to a halt.

Gianni Ciolli, EDB

About

Gianni is Global Vice President, Practice Lead for High Availability at EDB, helping customers and colleagues in the implementation of highly available distributed architectures.

Before that, he worked at the University of Florence as a researcher in Mathematics. Gianni has been using Free Software for almost 30 years. As a member of the Italian PostgreSQL User Group (ITPUG), he helped organising the first European PGDay in Prato, Italy (2008) and many editions of the the Italian PGDay. He is an author of the PostgreSQL 16 Administration Cookbook and of its previous 5 editions.

Modernize your application with Postgres

Elevator Pitch

Find out why developers admire, desire and use Postgres more than any other database. Hear it from a speaker who works with Postgres since almost 20 years, whose first contribution to the project dates back to 2009, and who authored the last 6 edition of the PostgreSQL Administration Cookbook.

Description

According to Stack Overflow’s survey of 65000 users, Postgres is the most popular, admired and desired database of 2024, for the second year in a row.

In this presentation we will explain the unique qualities that cause that appreciation, showing how Postgres is a critical part of the modernisation of IT landscapes and how you can best use it.

Bogdana-Ioana Hincean, Expersoft Systems

About

I am passionate about technology that empowers users and solves real problems. I'm driven to contribute to meaningful projects that create a lasting impact; "Limits dissolve when innovation meets purpose"

From Havoc to Harmony: The Power of OS Management Hub

Elevator Pitch

Automation is crucial in IT, but unchecked, it can become an adversary. This session explores a real incident where an automated update disrupted production and shows how Oracle OS Management Hub can prevent such issues. Learn about centralized management, patch orchestration, and system security.

Description

From Havoc to Harmony: The Power of Oracle OS Management Hub

Automation is no longer an option but a necessity—essential for efficiency, scalability, and security. However, when left unchecked, it can go rogue, turning routine updates into unexpected disasters. A single automated update can trigger downtime and operational chaos, eroding trust in the very processes meant to enhance stability. Automation should be the rule, not the exception, but the real challenge lies in maintaining visibility and control over it.

So how do you keep automation in check? By bringing an orchestrator to the scene—one that ensures every process follows the right path, transforming potential havoc into harmony. OS Management Hub provides a centralized management console to orchestrate and monitor updates across Oracle Linux and Microsoft Windows systems, whether running in Oracle Cloud, on-premises, or third-party clouds.

This session will provide a detailed breakdown of OS Management Hub's capabilities, including centralized package management, patch orchestration, and compliance monitoring. We'll showcase real-world use cases to illustrate how teams can streamline OS lifecycle management and enhance system security while maintaining full control over updates.

By the end of this session, attendees will have a clear roadmap for leveraging OS Management Hub to improve security, consistency, and reliability while reducing complexity and operational costs. If you've ever experienced the frustration of an update gone wrong or struggled with maintaining security across a complex IT environment, this session is for you. Join us to learn how OS Management Hub turns automation from a potential rogue force into a strategic advantage—and let's share our best (and worst) automation stories!

Oliver Lemm, Hyand Solutions

About

Working as an APEX Developer, Architect or Coach till 2007 for Hyand Solutions GmbH, formaly known as MT AG. Being presenter on national and international conferences like DOAG, APEX Connect, ODTUG KScope , IT-Tage for over 15 years.

Table Functions at it's best

Elevator Pitch

Table functions sind Jahren ein Bestandteil von PL/SQL.Über die Zeit haben Viele sie vermutlich in irgendeiner Weise schon genutzt. Aber wann genau machen Sie Sinn, was sind Fehler im Umgang mit den Table Function und wie sehen die Vor-/Nachteile im Vergleich zu Views und Materialized Views aus?

Description

Bei der Abfrage von Daten gibt es eine Reihe von Features in der Oracle Datenbank, welche man mal mehr oder weniger gut einsetzen kann. Gerade, wenn die Komplexität im SQL steigt und vor allem, wenn fachliche Logik dazukommt sind Table Functions ein probates Mittel. Dieser Vortrag zeigt auf, in welcher Weise Table Functions sich sinnvoll einsetzen lassen. Angefangen bei der Defintion der Typen in der Datenbank und bei der Struktur in der Function selber, bis hin zum Logging wird die Mechanik erläutert.

Wie genau gehe ich um mit Parametern, wie sollte das Logging aussehen und welche Fehler sind im Exception Handling besonders zu behandeln? Was sollte ich beim Thema Caching beachten und wie wirkt sich der Kontextwechsel von SQL und PL/SQL aus? Gibt es Bereiche in APEX oder im PL/SQL wo auf keinen Fall eine Table Function genutzt werden sollte und was hat es damit auf sich, wenn ich Komponenten wie ein Interactive Report mit einer Tablefunction versorge?

Wann reicht eine “normale” View, wo fängt die Table Function an und wo sind Materilized Views eine passende Lösung? Diese und noch einige Fragen mehr werden beantwortet, um den richtigen Einsatz von Table Functions zu erreichen.

Ilya Kosmodemiansky, Data Egret

About

Ilya Kosmodemiansky is a CEO and founder at Data Egret. A consultancy specializing in PostgreSQL migration, maintenance and support. Ilya has a broad experience working with PostgreSQL as consultant, architect and administrator. His main focus is database performance and optimization. He sees the mission of PostgreSQL in substituting the commercial databases in high-performance mission-critical applications.

His interests are promoting PostgreSQL as enterprise-ready database, distributed transaction processing, data integration.

So wird man ein PostgreSQL-DBA: vom Anfänger zum Fortgeschrittene.

Elevator Pitch

Grundlagen von Postgres für die Administratoren, die schon Erfahrungen mit der anderen Datenbanksystemen haben, aber in Bereich Postgres schnell wechseln möchten

Description

Heutzutage interessieren sich viele Leute für PostgreSQL, da es immer beliebter wird. Egal ob man ein kompletter Anfänger oder ein erfahrener Administrator ist, der z.B. sich auf eine andere Datenbank früher konzentriert hat, braucht man ein bisschen Zeit, um eine neue Technologie kennenzulernen.

In diesem Vortrag gebe ich eine kurze Einführung, wie man eine Postgres-Instanz mit einer grundlegenden hochverfügbaren Replikationslösung und einem Backup-/Wiederherstellungs-Setup installiert und einrichtet. Wir werfen auch einen Blick auf einige wichtige Tools, die PostgreSQL-DBAs täglich verwenden, wie wichtige pg_catalog Ansichten, Monitoring-Tools und PostgreSQL-Erweiterungen wie pg_stat_statements .