



Database Cloud Day - Abstracts

Keynote

The Future of Data-Centric AppDev – Dominic Giles

In this session, Dominic explores the transformative influence of cutting-edge Transformers like ChatGPT on our daily lives. Unveiling the technology driving these innovations, he navigates the convergence with Oracle database capabilities. Discover how the integration of Transformers and the latest Vector search functionality propels enterprise value to unprecedented heights.

Track 1

Game, Set, and Match – Patrick Barel

The tagline of the Game of Set is: The Family Game of Visual Perception. It is a game of finding matching cards in a selected group of cards. Being a SQL guy, I read this as: SET: The GAME of MATCHing cards. In any database we can select SETs using MATCHing criteria. Where is the GAME? Well, for me, it is a game of trying to solve this puzzle using SQL.

You will see how you can use the different SET operators in Oracle, how you can use JOINS to create a SET of cards, from which you can SELECT a puzzle and also find the SETs in that puzzle.

Leverage APEX for management and self-service of OCI – Robert Pastijn

Giving colleagues and customers access to the OCI console for self-service might be overwhelming or expose too many options. How about simply writing your own Self Service application with your rules, layouts and restrictions? This presentation will demonstrate how easy it is to connect your APEX application to the Oracle Cloud Infrastructure pushing requests or retrieving information using the various APIs.

Why the new 23c database feature JSON relation Duality Views may change the life of developers – Peter de Vaal

In the past 20 years the architecture of applications working with a relational database has changed from client-server with persistent database connections to web technology or service oriented technology using connection pools.

The majority of the applications are being built using object-oriented languages or frameworks. The

mapping of the OO-model to the relational model in the database (OR-mapping or ORM) has always been a challenge, and proved to be a bottleneck regarding performance, refactoring data model changes and effectively maintaining data integrity.

The cause of such problems was a common one: the OR-mapping was usually done in the application layer, not in the database layer. In some solutions this pattern was circumvented by making use of XML (since Oracle 9i) or JSON (since Oracle 12.2) documents: Both the OO-application and the database do understand these, and the database can easily query the XML or JSON and convert it to a relational model. By using stored procedures that consumed the XML or JSON as input, while other stored procedures produced XML/JSON as output it was possible to do all the mapping in the database. This method, however, had some downsides: it did not support ACID transactions, and it required some low level programming.

With the new JSON Relational Duality Views both problems have been solved: The mapping is now done in the database in a declarative way, and it is possible to support ACID transactions again.

Oracle Database 23c for the APEX Developer – Roel Hartman

Oracle Database 23c is labeled as the first release primarily aimed at the developer. The release includes many new and updated features that makes developing applications using Oracle database easier and faster. But what does that mean for the APEX Developer? What features are of particular interest and benefit of using this low code tool for developing web applications?

This session will show you how you can benefit of this "Developer Release" in order to build APEX applications even faster, more secure and more scalable than before!

Track 2

Azure AD oauth token Oracle database authentication – Martijn Pronk

As of Oracle database release update 19.18 it is possible to use Oauth token based authentication. In this session I will present the steps involved on both Azure AD and the database side to make this work. Topics include getting TCPS (Oracle net with TLS encryption) to work, configuring Azure AD, configuring the database and mapping Azure identities to a database user account.

Oracle True Cache: caching simplified – Dominic Giles

Join Dominic in this session as he unveils Oracle Database 23c's game-changing feature: True Cache. Dive into how True Cache revolutionizes application performance effortlessly, requiring minimal or no code adjustments. Dominic outlines key use cases, demonstrating how this breakthrough technology simplifies the implementation of what was once a complex database caching process.

Migrating to Autonomous Database – Toon Koppelaars

Oracle Autonomous Database takes care of many tasks that would once have been the responsibility of developers and DBA's. This makes it even easier for you to develop and deploy new applications. When it comes to moving existing applications, simply lifting and shifting does not always result in the best performance. Right now your applications could be using an "exciting" database: one with underscore parameters and other non-default settings, and possibly certain patches. Oracle

Autonomous Database is boring by design: there are few deviations from defaults, and most are chosen with reliability in mind. And there are other areas that will differ: think, statistics collection, resource management, privileges, to just name a few. The available Cloud configurations are flexible but do not cater well for exciting databases. In this session we'll share our experiences of the past few years assisting customers, and internal development groups, with their application migration to the Cloud.

One PDB to go, please! – Christian Gohmann

Starting with 12c Release 1, Oracle introduced a completely new architecture concept for its database - the Container Database. With this new architecture, new challenges came up but with the same breath a wide branch of new opportunities.

The presentation will address the capabilities to create fast and easy new (test) databases or clones for a running production database. Five different ways will be discussed.

- Using Local and Remote Cloning
- Using an Unplugged PDB (predefined master)
- Using Refreshable PDBs as a master for new (test) databases
- Snapshot Carousel

Another point on the agenda is the usage of the Snapshot features of ACFS and Direct NFS to speed up the creation process.

At the end of the presentation, a comparison of the advantages and disadvantages of each technique will be presented.

Track 3

DBMS/Utility-Packages - Explore the possibilities! – Axel vom Stein

When using the Oracle Database you get a bunch of PL/SQL-Packages for free. This includes the DBMS_ and also several other Utility-Packages. All these PL/SQL-Packages add extra functionality to the Oracle Database. Some of the extended features can't be used without knowledge of these PL/SQL-Packages. The extensions and PL/SQL-Packages are not only for Developers, even Database Administrators should have interests in them. This presentation gives a rough survey of the PL/SQL-Packages, categorizes them and some selected examples will be pointed out.

The presentation is aimed at beginners/ more or less independent of their respective job roles (Developer, DBA).

Learning Objective 1: An overview and classification of the PL/SQL-Packages which you will not get from the Oracle documentation.

Learning Objective 2: Understand what extended features you can use with PL/SQL-Packages.

Learning Objective 3: See the use of mostly unknown PL/SQL-Packages in some selected examples.

We've gone Serverless, this is what happened– Piet de Visser

Two Customers were fed up with the ticketing-SLA from the "service provider". They followed the agile-gurus and went "Serverless", taking a set of "legacy" systems to "The Cloud".

Going Serverless Definitely gave back control to the (IT department of) the customer, but there were some challenges. The deployment in the new situation is either via Kubernetes or Lambdas, and the databases are all DBaaS (and still mostly Oracle).

Because two different customers did take two different approaches, we are able to compare and learn something from both cases.

The overall experience has been very Positive and has given us a few lessons about our legacy-systems, software-tools and the data, in the process.

There are some interesting plans for the near future as well hence the content of this presentation may grow or adjust over time.

Oracle Data Pump Deep Dive into Internals – Daniel Overby Hansen

Oracle Data Pump is one of the most widely used DBA-oriented features of Oracle Database, combining powerful flexibility with high-speed movement of data and metadata into and between Oracle databases. Even if you already use Data Pump, have you ever wondered how it works internally, or why it behaves the way it does? If so, then this deep dive into Data Pump internals is for you! We will delve into the inner workings of this widely used utility to look at topics such as:

- How Data Pump works
- Parallel operations
- Data Pump Data Movement Methods
- Getting the best performance for LOBs
- Creating indexes faster
- Filtering
- DBMS_DATAPUMP API
- Troubleshooting

You'll have a better understanding of this important tool for managing your databases, including:

- Understand how Oracle Data Pump performance is improved with parallelism
- Learn how Data Pump provides flexibility and ease of use through object filtering and various import modes
- Grasp the key role played by the Data Pump Control Table

And the Saga continues... manage Microservice transaction wit Oracle MicroTX- Michel Schildmeijer

A Saga is a pattern for managing distributed transactions in a microservices architecture. It is essentially a sequence of local transactions that are coordinated by a Saga coordinator to ensure that the overall transaction is executed atomically. Oracle Database provides built-in support for implementing Saga patterns. By implementing the Saga pattern with Oracle Database, you can simplify microservice transactions and ensure that they are executed atomically, even in a distributed environment. The DBMS_SAGA package in Oracle Database can be used to implement the Saga coordinator, which manages the local transactions and ensures that they are executed in the correct order. The pattern also includes compensating transactions to undo the changes made by a failed transaction, which ensures that the overall transaction can be rolled back and the system remains consistent. This is one of the many patterns for handling Microservice Transactions, which can be done by the Oracle Transaction Manager.

Keynote

No Slide Zone - Database Patchings Insights – Daniel Overby Hansen

You don't like patching? It seems to complex? Or it takes too long? No worries. Tell us about your database patching experience, and we'll explain how you can optimize and improve it. We'll share insights about patching, secrets you may have never heard or seen before, and we'll even listen to your complaints. We'll do this interactively with flipcharts, pens, and a demo environment only. You don't need to bring anything except your questions.

Livelabs

Build a social media app using APEX– Kevin Lazarz

Create a social media app where users can upload images, and like other users' posts. Get hands-on experience creating pages, regions, customizing the UI and creating dynamic actions. Learn to get user device location by implementing Map region.

Upgrade and Migrate to Oracle Database 19c & 23c the easy way – Kevin Lazarz

This Hands-on Lab is designed to help you navigate the migration to the container database (CDB) architecture in Oracle Database 19c & 23c. You'll see how to automate the process using the existing AutoUpgrade capability, and try out various options with a real-world scenario. In addition to migrating a sample database to Oracle Database 19c, you'll also use the new performance stability features to simplify the process.

