Oracle vs. MySQL:
First impression

Oracle Community
Tour EMEA

Dani Schnider and
Andrea Kennel
about beer brewing

EDITION-BASED REDEFINITION

And why it is essential for your
Continuous Database Delivery Setup

EOUC

What's Your Passion? Beer brewing with Dani Schnider and Andrea Kennel

Oracle vs. MySQL: First impression by Liron Amitzi

Oracle Community Tour EMEA – Looking back at 2022, Planning ahead for 2023

Have a comment? Suggestion? Want to submit an article proposal? **Here's how to do it.**

Editorial

Dear ORAWORLD Readers,

Nowadays, it is crucial to stay informed and up-to-date on the latest developments and trends in our field. As members of EOUC, we have a unique opportunity to share our knowledge and expertise with one another, and to learn from each other's experiences and insights.

In this issue of our magazine, we have included a variety of articles that we hope will be of interest and value to you.

Jasmin Fluri is explaining why Edition Based redefinition is important for the Continuous Database Delivery Setup.

Read about how to do the analysis of financial statements using Financial Consolidation and Close Cloud Services in the article written by Mohit Jain and Megha Gupta.

The article written by Rishin Mitra is revealing some insights about how secured the reports on Oracle Cloud SaaS are.

If you are wondering which one is a better database – Oracle or MySQL – then read the fair comparison Liron Amitzi is doing in his article.

The importance of collaboration and community is one of the most important themes running constantly through ORAWORLD issues. This issue is not an exception. In a world where technology can sometimes feel isolating and overwhelming, it is crucial that we come together and support one another. Whether it's through sharing our knowledge and experience, or simply offering a friendly word of encouragement, we can all play a role in building a strong and vibrant community. Check the article about the new Symposium 42 community to find out the what's, why and who are behind it.

Reading an article about the Oracle Community Tour EMEA can be an excellent way to gain insight about which communities took part in the tour in 2022 and the plan for this year.

Andrea Kennel and Dani Schnider are a known presence in our community, presenting at many conferences all over the world. But this time they will share another side of their story in the interview by Christian Luda: how they started brewing their own beers almost three decades ago.

As always, we welcome your feedback and contributions to our magazine. If you have a story to share or an idea for an article, please don't hesitate to get in touch.
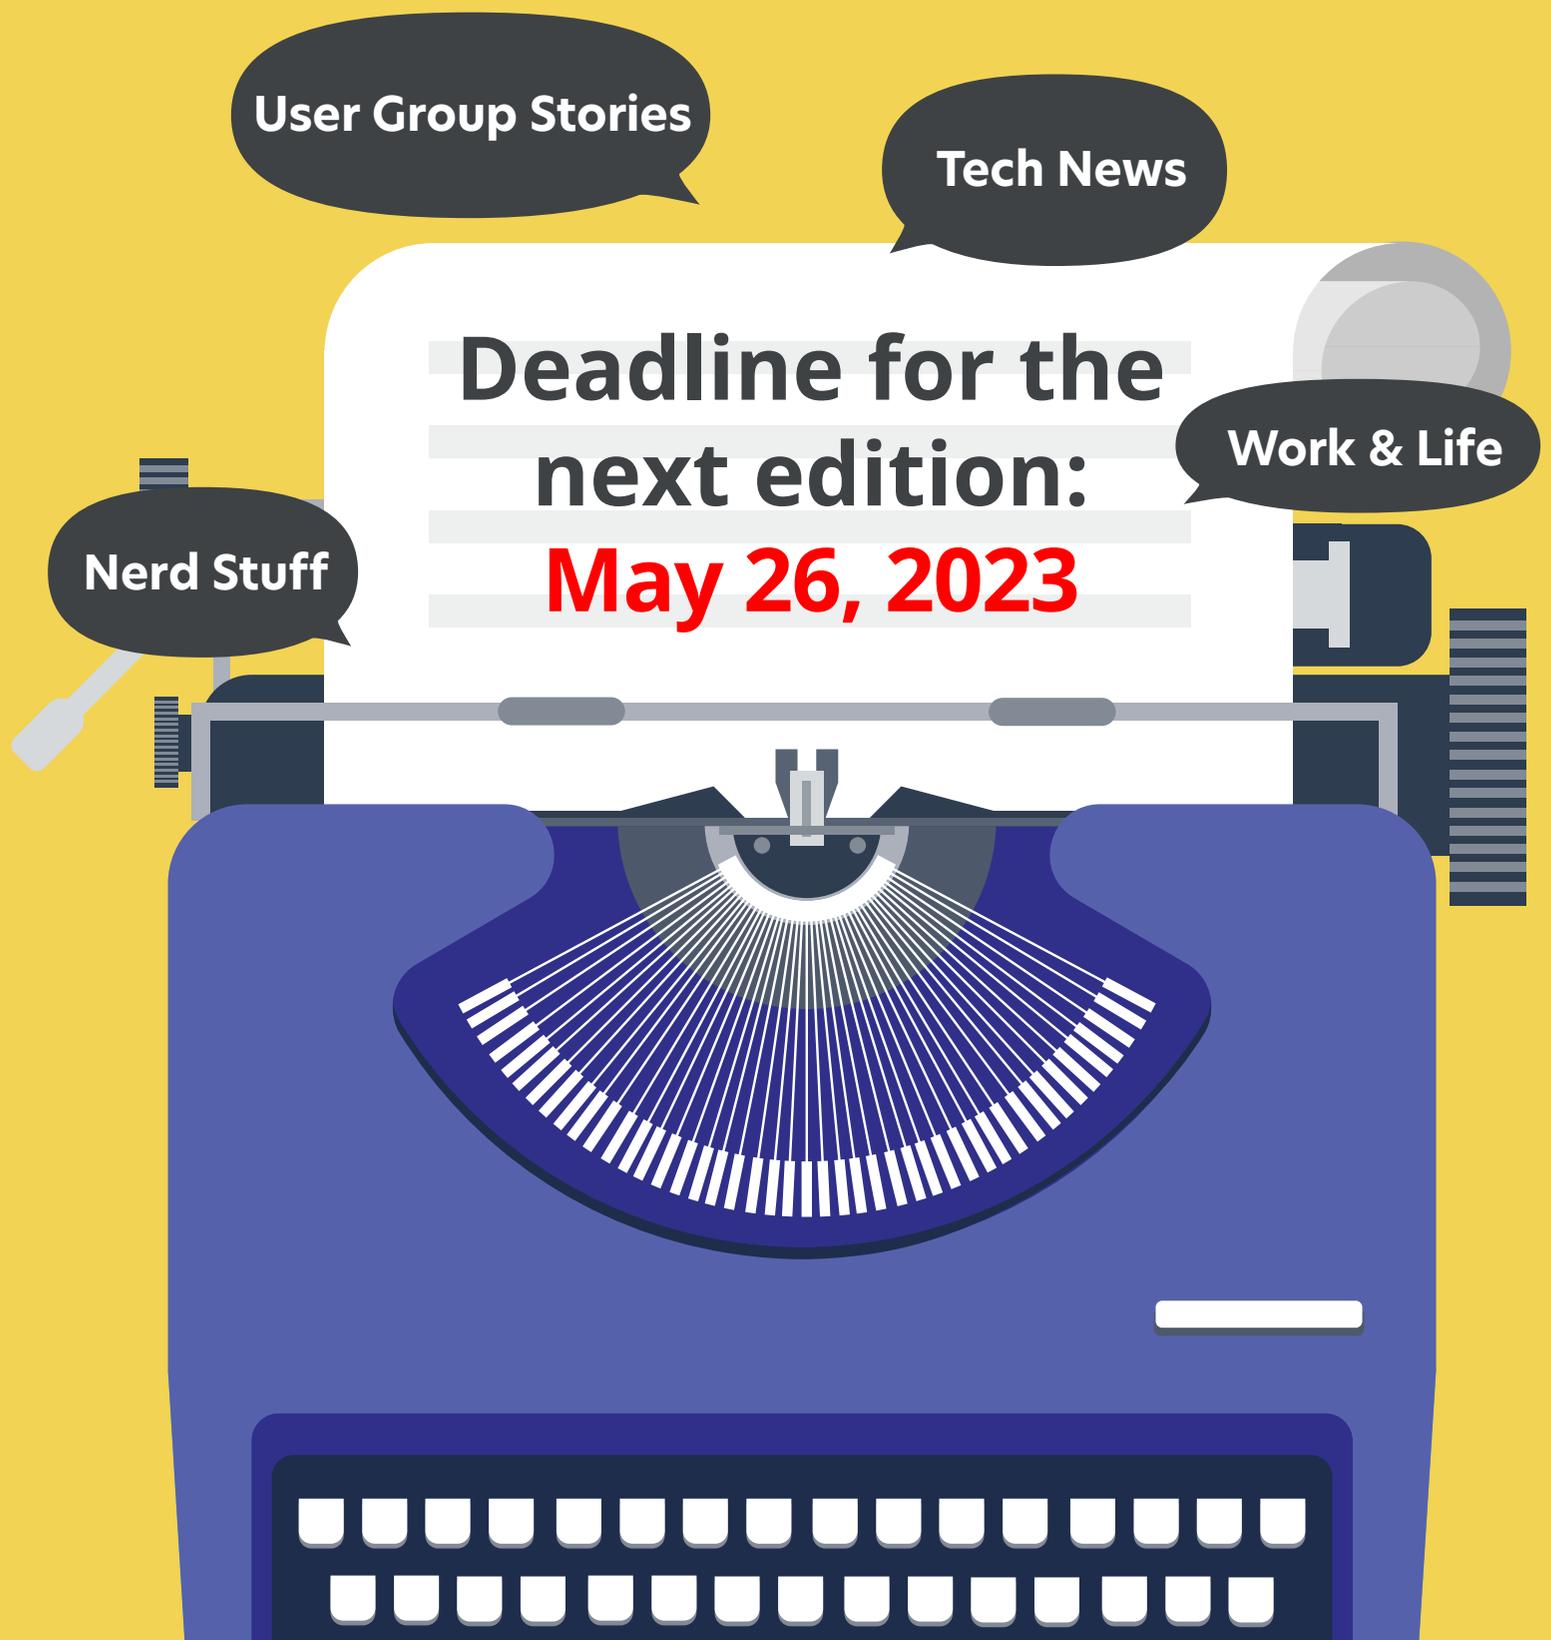
I hope you enjoy this issue of our magazine.

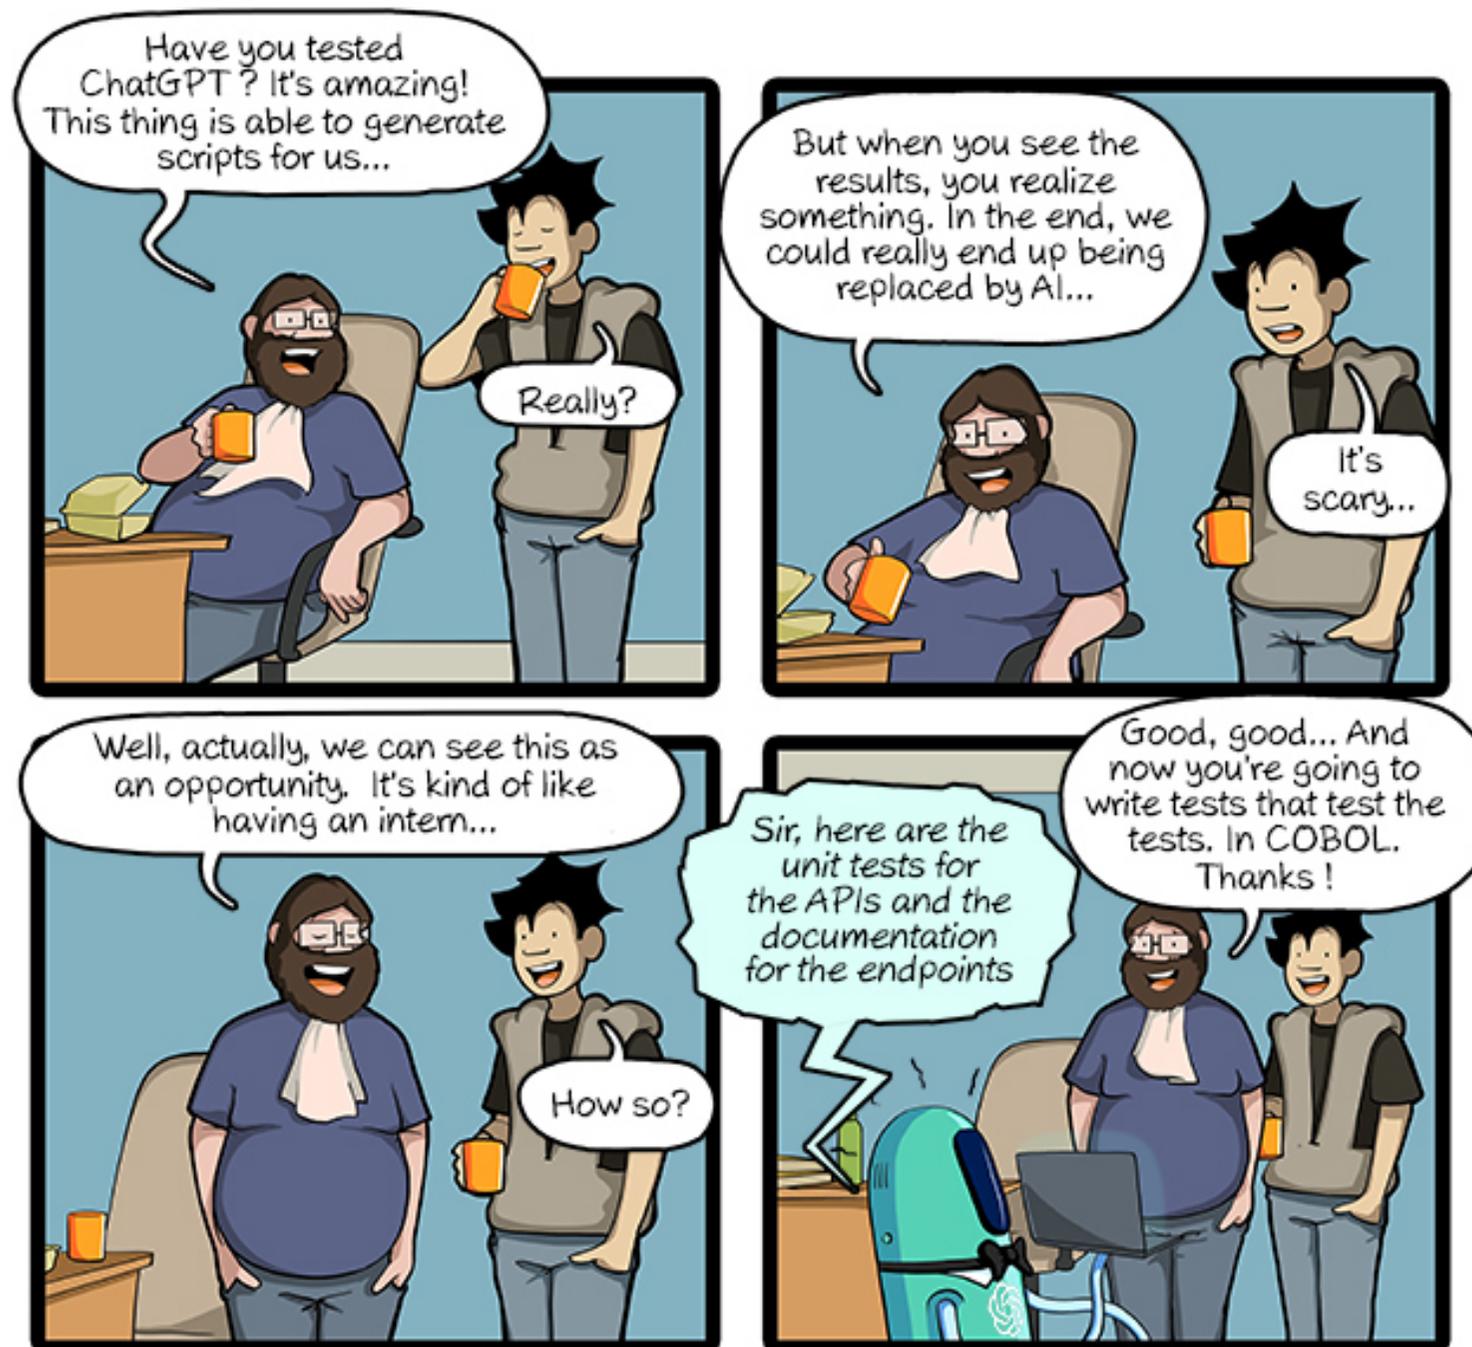Yours,
Mirela Ardelean
Oracle ACE Pro, President of EOUC

„A Whole New World"

www.commitstrip.com

Interview by Christian Luda

What's your Passion?

"The First Beer is Always the Best"
Dani Schnider and Andrea Kennel
about their love for beer brewing

In our "What's your Passion?" series we ask people from the Oracle community about their passions aside from work. For our latest edition we talked to Dani Schnider and Andrea Kennel about their passion: The Swiss couple started brewing their own beers almost three decades ago and since 2009, they are running their own small brewery "Monsterbräu" in Dübendorf near Zurich.

Andrea and Dani, you make craft beer. What exactly is craft beer and how does it differ from, let's say, Pilsner or wheat beer?

Dani: Pilsner and wheat beer can be craft beer, too. In principle, craft beer is simply the umbrella term for beer made by hand - just as there are craft sweaters or other things. We don't have a big machine where we press a button and then 100,000 hectoliters of beer are produced. We still stir, add hops and weigh – with technical aids, but a lot of manual work.

Andrea: I always say our beer is "handgehopselt" ("hand-hopped"). It's not industrially produced, but hand-brewed in small batches with a lot of love. And when you brew by hand, it's much easier to change the recipe. That's why with craft beer the variety is much greater.

Could you also make a Pilsner?

Dani: Pilsner is a beer we cannot make. Beer is divided into top-fermented and bottom-fermented. Most hobby brewers like us make top-fermented beer for the practical reason that it does not require refrigeration equipment or large

cold storage rooms. Top-fermented beers ferment at room temperature. That's why we make wheat beer, for example. That was one of the first beer styles we started with. Then also English beer styles like stout. IPA, India Pale Ale is very popular. So, it's first of all a technical reason why we make top-fermented beers. And the other reason is that Pilsners and lagers – these classic bottom-fermented beers – are produced by the big breweries. There's so much of that in the world already, we don't need to get in on the action. We prefer to concentrate on the less frequently produced beers.

Why is Pilsner the dominant beer in many parts of the world?

Andrea: Craft beers are usually very aromatic or have a lot of alcohol. These are not beers that you drink by the liter, but rather sip them and enjoy them by the glass at most. That's why you automatically have smaller amounts, and when you're thirsty, you down a Pilsner.

Dani: We also do beer seminars from time to time, and that's how we explain it there: Beer has four main ingredients: Malt, hops, yeast, and water. Depending on the beer style, these are important in different ways. A stout, for example, tends to be more malt-driven. Generally, with the dark strong beers, the malt is the flavor carrier. An IPA and, in principle, a Pilsner are more hop-driven. This is the reason for the bitter or the tart flavor, depending on whether it is an IPA or a Pilsner. Wheat beer is rather yeast-emphasized. The typical taste of wheat beer comes not from the wheat malt, but from the yeast. The common industrial beers, which are produced in large quantities, have a water emphasis (laughs). And that's not meant disparagingly. If you have a large brewery and want to sell millions of liters of beer, then you have to produce a beer that will be drunk in large quantities and not one where people have already had enough after one bottle.

Do you have a favorite beer?

Dani: One favorite beer, no. That would not be what we want to convey to the people. We want to show that there are just different types of beer. It really depends on the situation. I like IPAs, because they are relatively fruity and often bitter - you can enjoy them all year round. In the summer, I tend to drink maybe a wheat beer. And the darker, malty beers are more for the cold winter days. Or even a stronger Belgian Dubbel or Tripel. You don't drink those at a barbecue in the middle of summer. Coincidentally, these are all beers that we brew as well.

Andrea: What you also like - and so do I - are the sour beers, especially on a summer evening. Then I like Belgian Tripel, paired with a piece of meat. Depending on the time of day, depending on the season, depending on what food, there's a beer that goes with it. When it comes to our own beers, I think we can be proud of the Belgian Lazariter. And the one that's doing best at the moment is our IPA which is well received. In Switzerland you often hear from women: "Beer is too bitter for me." Then I usually serve them an IPA, the Blue Pirate, and then they say: "Oh yeah, I like that, that's bitter, but a different kind of bitter." It's a single hop with a fruity bitter flavor reminiscent of grapefruit. That one is quite popular.

When you create a beer, do you always have the food already in mind to go with it?

Dani: Food pairing is an important topic. We did a beer sommelier training course a few years ago, and that was an important aspect. It's also part of beer tastings, that when you describe a beer, you also say what food it goes with. And vice versa: We have a certain meal - what beer can we drink with it?

You started out in 1994 with a ready-mix. Which of you was the first to come up with the idea to start brewing back then?

Dani: I believe that was you, Andrea.

Andrea: Yes, I was in a store where you could buy beer, and I saw a fermenting vat and a can next to it. And I thought: "Ah, you can brew your own beer." It was a concentrate, a syrup. The recipe said you have to boil 4 liters of water and then add 16 liters of cold water to have the right temperature,

mix it in there and then add the yeast to ferment. That was the first attempt. Then we slipped into it step by step and attended a seminar on how you can mash it yourself.

How long did it take before you felt confident to do it yourself?

Andrea: Our course instructor put it well: It's like making pizza. In the beginning, you have the ready-made pizza that you throw in the oven. That's the syrup. At some point, you buy a pizza crust and topple it yourself. Then you start to vary on that. We're at the stage now where we're also making the pizza crust ourselves. You really learn that step by step. You can also buy kits where the malt is already ground, where the right mix is there, with a recipe, that's relatively easy. At some point you think: "Good, now I've tried the pizzas, I know how to do it, now I'll make my own recipe." It's not like you have to practice for a year, you can make something yourself very quickly with the recipes that are out there. And the first beer is always the best.

Dani: In the beginning, we stuck very closely to the standard recipes, with the idea that if you deviate a just a tiny bit from them, it's no longer a beer. But actually, it starts to get exciting when you bring in your imagination.

If you are invited to a party where they serve normal standard beer, is that a problem for you?

Dani: Depends on the party.

Andrea: Depending on the beer, I say I'll drink a glass of water first to quench my thirst. I somehow don't like lager so much anymore.

Dani: When you've tried a lot of different varieties and know different styles of beer, you get picky over time. It's like wine. Once you've had a Bordeaux, you don't drink the cheap fusel

wine anymore. Or with whiskey: I don't drink Ballantine's or Johnnie Walker anymore.

You supply bars in your area. Do you plan to expand even further?

Andrea: We have an acquaintance who diligently advertises for us. At one point we said: "OK, this is getting to be too much." We invited her to brew and bottle beer with us so that she can see what kind of effort it is. We were already maxed out with what we could produce. Craft beer means you bottle each bottle individually. Then you close the clip. Then stick the labels on. And when the bottle comes back, you wash the bottle and scrape off the label. A lot of manual labor. It's a hobby. And if you had to go to the brewery twice a week, it would definitely be too much.

Dani: But the manual work can be combined quite well with other activities. My job is mainly filling. That's actually hard

work. You stand in front of the bottling line and spend hours bottling. But I find it quite relaxing in parts. For example, when I come back from a long trip, I still go to the brewery and bottle beer while thinking about the project. During Corona, it was a reason to get out of the home office. For example, I attended the KScope conference remotely and with the time difference, the talks ran from 6 pm to midnight talks. So, I took the laptop to the brewery and listened to the talks while bottling. That was quite relaxing.

Does the limitation also make the appeal of your product?

Dani: That is certainly an important aspect of the whole craft beer scene. It makes it special, a bit of a luxury that you afford yourself. In terms of price, we can't compete with industrial beer, so we try to do it through the unique, the local. And that certainly goes down well with people. That's a completely different market segment. Our customers aren't the ones who go out and get a case of beer to get drunk or party and play beer pong. Instead, they take one bottle – we have half-liter bottles that two people can share.

Andrea: And that's also the special thing about craft beer, that some of it is really only available locally. When you're traveling, you can find a local brewery in one place and find beers that are not available anywhere else. I can remember a trip to St. Petersburg. There, if you look for the local beer bars, you'll find yourself in backyards and neighborhoods that aren't touristy at all anymore. That's fascinating. You really meet people from there.

Your job takes you to many places in the world. Is there one place where the beer scene is particularly good?

Dani: The hotspots are certainly Belgium, in Germany more the traditional varieties, especially Bavaria and Franconia, and also the USA. When I was in the US for the first time, I had the feeling that only Budweiser and Millers were there. But the US is the cradle of the whole craft beer scene, especially on the West Coast. OpenWorld in San Francisco had a pub called the Thirsty Bear.

Andrea: I remember a few years ago, we went to Stockholm and deliberately drank a little beer beforehand, because we thought it would be difficult beer-wise there. But then we passed a restaurant and saw the people have beer in front of their nose and each glass had a different color. Strange. We sat down and saw they had different beers. Great! By chance we found the beer pub. That's when you noticed, this had also developed in Sweden. I've been surprised in many places how big the craft beer scene is. There are many small breweries in many countries. That's where the density has grown in many countries. And that's where I have to advertise Switzerland now: The density of breweries per capita is the highest in Switzerland.

Dani: We have well over 1,000 breweries per 8 million inhabitants. But that's not because the Swiss are the world's record holders in beer drinking; that's still the Czechs. The legislation in Switzerland is very simple. If you start a brewery, you have to register it, you have to pay beer tax, and you have to comply with food regulations. But it's much easier legally

every 200 students. That meant that as a woman you had to assert yourself a bit. Which is what I did and what earned me a nickname: Monster.

What significance does this passion for beer brewing have for your life?

Andrea: Hobbies are very important for work-life balance. And when you're doing IT, you're at the screen. With beer, you really have the product in your hand. All the senses that are involved. Software, it doesn't smell. It's something completely different. For me, it's important that I have a hobby that doesn't just take place in the virtual world, but that I have something that I can perceive haptically.

Dani: I'll go along with that. I think that's typical: of the colleagues we know who brew beer, a large proportion are in IT or other office jobs. I don't think that's a coincidence. A craftsman is less likely to think of brewing a beer in the evening or on the weekend. Work-life balance is also extremely important for me. It's definitely a great balance.

Andrea and Dani, thanks a lot for taking your time.

than in Germany, for example. That's why there are so many of them here. Of the 1,000 breweries, there are 3 or 4 huge breweries and maybe 20 medium-sized breweries, and then there are lots of breweries like ours.

Where does your name Monsterbräu come from?

Dani: The question always comes up when we hold beer seminars. I always give the short answer: I'm the "Bräu". Then there's a few seconds of silence and then comes the long answer from Andrea.

Andrea: The long answer actually has to do with computer science which we both studied. I started studying in 1983 and back then the percentage of women was 3.5 %. 7 women for

About Dr. Andrea Kennel

Andrea studied Computer Science at the ETH (Swiss Federal Institute of Technology) and complemented this with a PhD in Computer Science. She has over 20 years of professional experience mainly in Oracle database and DWH projects as Lead Data Architect and also as developer or project manager. Since 2014, in addition to her consulting work, she is also a lecturer at the University of Applied Sciences in Brugg/Windisch (www.fhnw.ch) and teaches databases as well as project management in the Computer Science and Data Science courses. She regularly gives talks at conferences and writes articles.

About Dani Schnider

Dani works as a Principal Oracle & DWH Consultant at Callilsta, a Swiss IT consulting company. He has been working in database consulting for over 20 years, mainly in the area of data warehousing. 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 with Oracle" (2011) and "Data Warehouse Blueprints" (2016), both published by Hanser-Verlag.

Number of the Month: 42

In Douglas Adams' book "The Hitchhiker's Guide to the Galaxy," 42 was described as the "Answer to the Ultimate Question of Life, the Universe, and Everything."

Chat GPT answered the question "what's the meaning of the number 42?" with
" ... *However, beyond this fictional reference, the number 42 itself does not hold any special significance or meaning in mathematics, science, or other fields. It is simply a natural number that comes after 41 and before 43.*"

No special significance? Really?
What **if** "42" **has** a significance to a bunch of people ?
I invite you to read the article "A new community was born" in the current magazine to find out!

Techs & Nerds

Analysis of Financial Statements using Financial Consolidation & Close Cloud Services (FCCS)
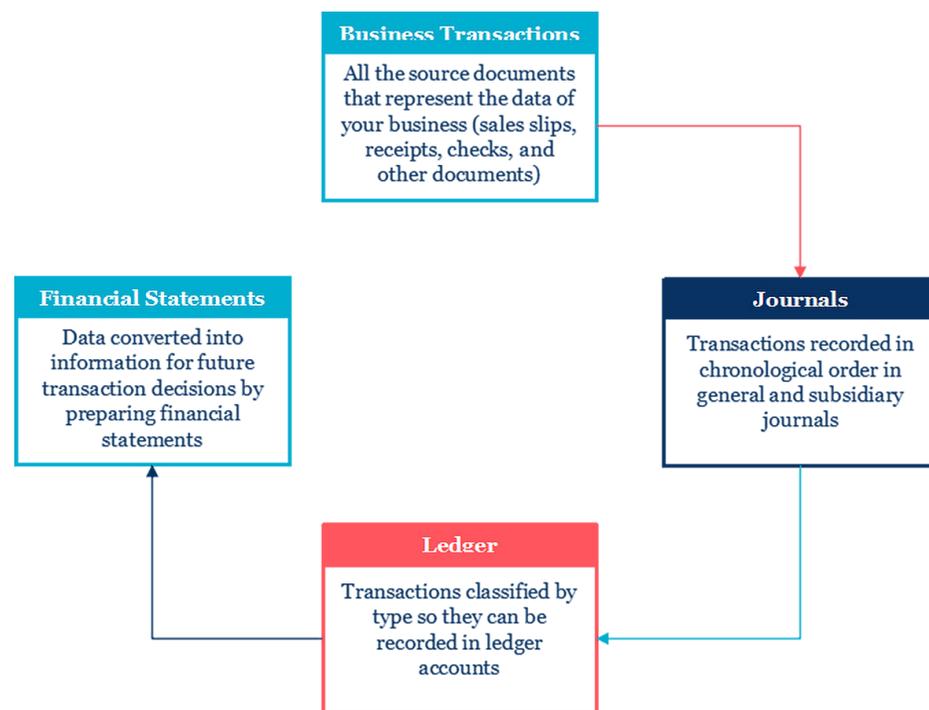
Mohit Jain and Megha Gupta
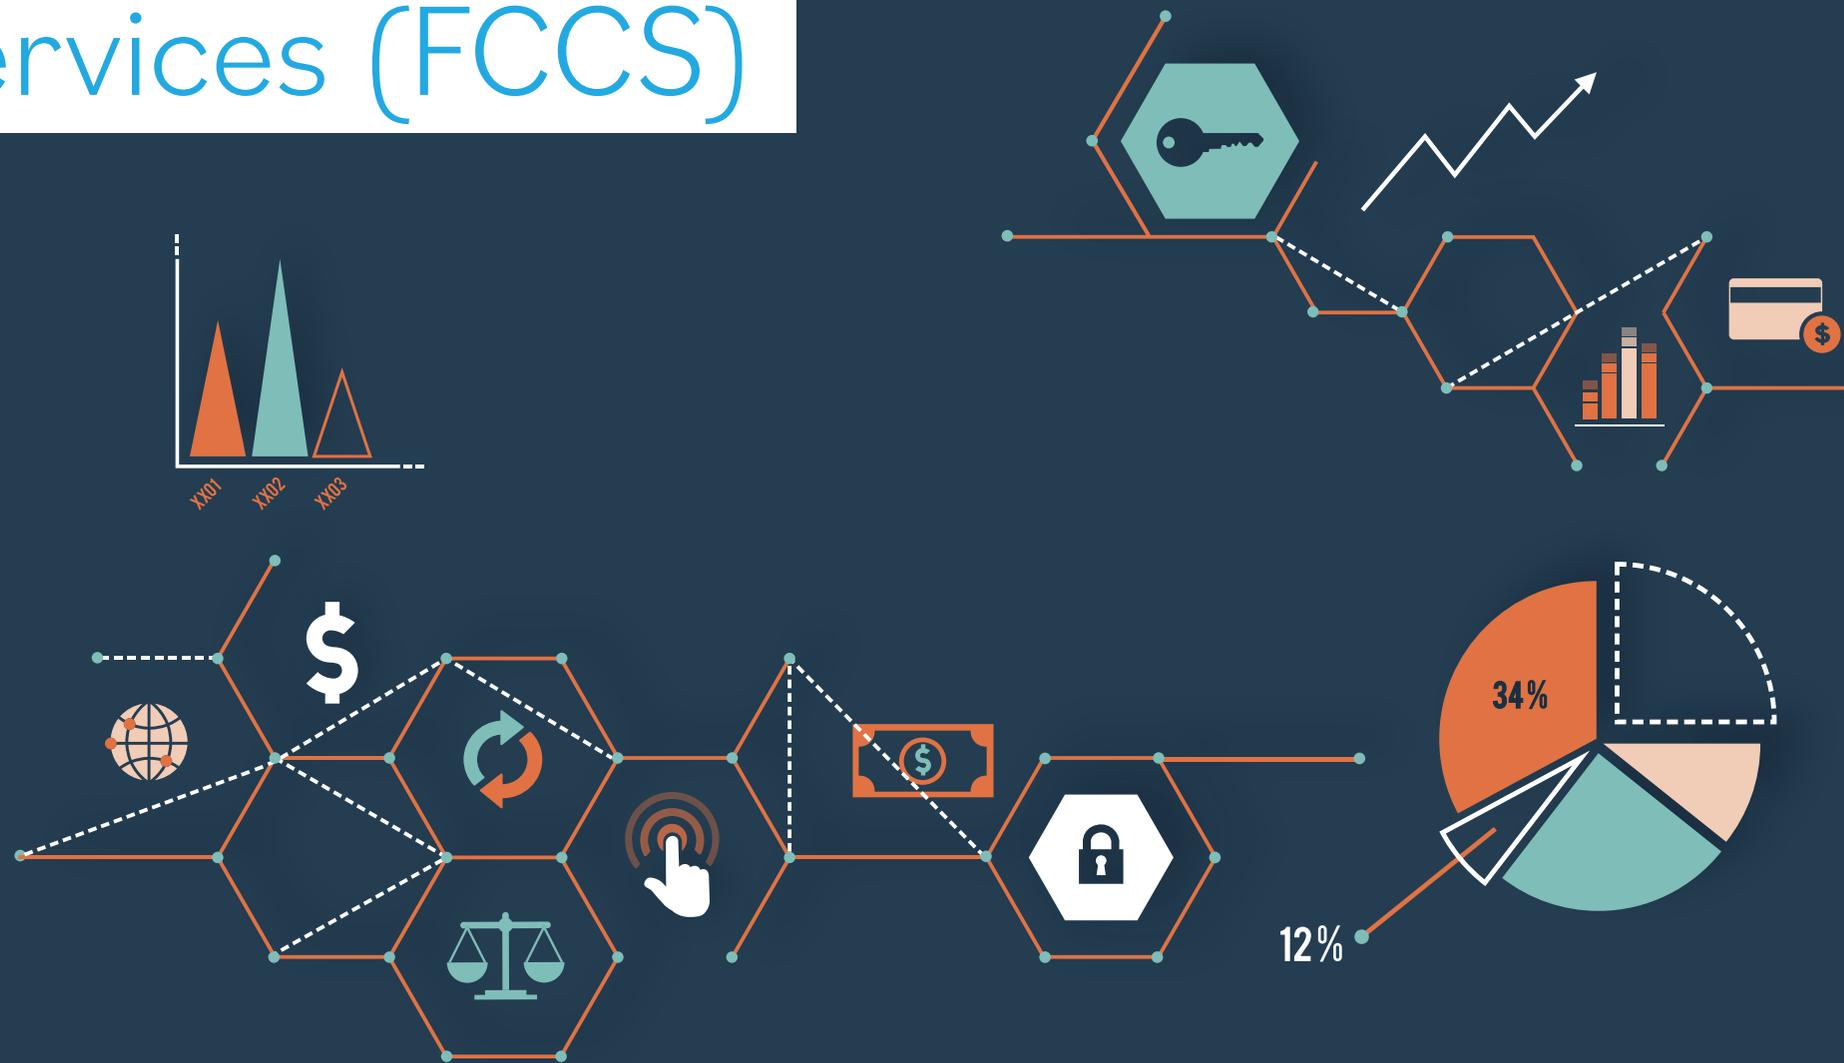
Financial statements of a business are 'Summarized Financial Reports' which provide the 'operating results and financial position' of a company and the detailed information contained therein is useful for assessing the financial soundness and operational efficiency of a company.

> *Financial Statements are the complete reports on the health of a business.*

The financial statement determines –
- If a business has the ability to repay loans
- If it has the cash flow to meet bills and purchase stock
- From where the business is generating cash and where the cash goes
- If the business is profitable and if it will stay profitable
- If there are any large problems approaching

And other important insights...

Business Transactions
All the source documents that represent the data of your business (sales slips, receipts, checks, and other documents)

Journals
Transactions recorded in chronological order in general and subsidiary journals

Ledger
Transactions classified by type so they can be recorded in ledger accounts

Financial Statements
Data converted into information for future transaction decisions by preparing financial statements

There are three basic financial statements:
- Balance sheet
- Income statement
- Cash Flow statement

In this article, we will have an overview of analysing financial statements using Oracle EPM Cloud tool named "Financial Consolidation & Close cloud Services (FCCS)"

Financial Consolidation & Close cloud Services (FCCS)

For many organizations managing the financial consolidation and close process is a burden that takes time and resources that could better be used for analysis and decision-making. Financial Consolidation and Close is a purpose-built business process available in the Oracle Cloud EPM for both effectively and efficiently managing the consolidation and close, end to end. It supports:

- A configurable and upgradable consolidation and close framework to tailor the solution to specific requirements.
- An end-to-end consolidation and close solution - allows business to configure consolidation application and monitor the business close process.
- Powerful analytics, dashboards, reports and analysis for unparalleled insights.
- No infrastructure investments required – it is a SaaS based solution.
- Best practices out-of-the box to get up and running quickly - dramatically reduces the time and cost to complete a close process.
- Complex Consolidations – out of the box, yet configurable, complex consolidation support to re-classify, adjust and eliminate reported data.
- Supports Intercompany Elimination, Currency Translation,

Automated Cash flow reporting, robust journal module, ownership module, multi-GAAP reporting, Detailed data source tracking and other features.
- Using Oracle Reports, Narrative Reporting, Smart View for Microsoft Office and in-built dashboards, Financial Consolidation and Close has robust reporting capabilities to meet every requirement.
- Provides significant out-of-the-box functionality and built-in task orchestration for management of the overall close process.
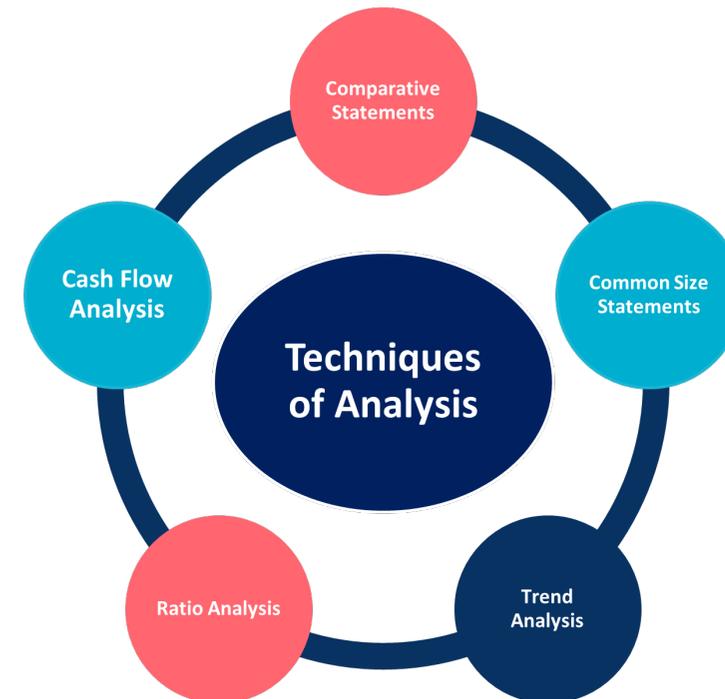
Analysis of Financial Statements

Financial statement analysis is the process of analyzing a company's financial statements for decision-making purposes. It consists of the evaluation of the financial condition and operating performance of a business firm, an industry, or even the economy, and the forecasting of its future condition and performance. External stakeholders use it to understand the overall health of an organization as well as to evaluate financial performance and business value.

> *The process of critical evaluation that aims to estimate current and past financial position from the financial information contained in the financial statements, in order to understand and make decisions regarding the operations of a company is called 'Financial Statement Analysis'*

Techniques of Analysis

There are different techniques/methods to analyze the

financial statements -

Comparative Statements - Compare the financial statements of a firm for different periods of time to see the profitability and financial position of a firm

Common Size Statements - Shows the relationship between different items in the same financial statement, is identified by expressing all amounts as a percentage of total amount.

Trend Analysis - It is a technique of studying the operational results and financial position over a series of years.

Ratio Analysis - Accounting ratios measure the comparative significance of the individual items of the income and balance sheet statements.

Cash Flow Analysis - It includes the analysis of actual movement of cash into and out of a company

Analysis of Financial Statements using FCCS

FCCS provides Powerful analytics, dashboards, reports, and analysis for unparalleled insight. It provides different pre-seeds financial reports, dashboards, Data forms to analysis financial statements. It allows to configure existing financial reports, dashboards, Data forms or create new based on requirements.

Moreover, FCCS is seamlessly integrated with Microsoft Office products using Smart View to allow users to analysis the financial statements, with dashboards in Microsoft Office products (Excel, Word and Power point).

Conclusion
Why FCCS for Financial Statement Analysis
- FCCS, a means to have all the data in one place in same format, readily available for analysis.

- Consolidate strategic ERP Financial data into one repository, single version of truth.
- Reduce the time report out post business acquisition.
- Strategic for enterprise with diversified business.
- Enable reporting and analytics across ERP's.
- Enable outcomes-based analytics and dashboards to support processes.
- It includes close management as well as Supplemental Data Management.
- Streamlines Finance and Reporting by creating rule-based mappings.
- Customizations if any required.
- Create a fast works culture around data management, rapid prototyping & delivery.
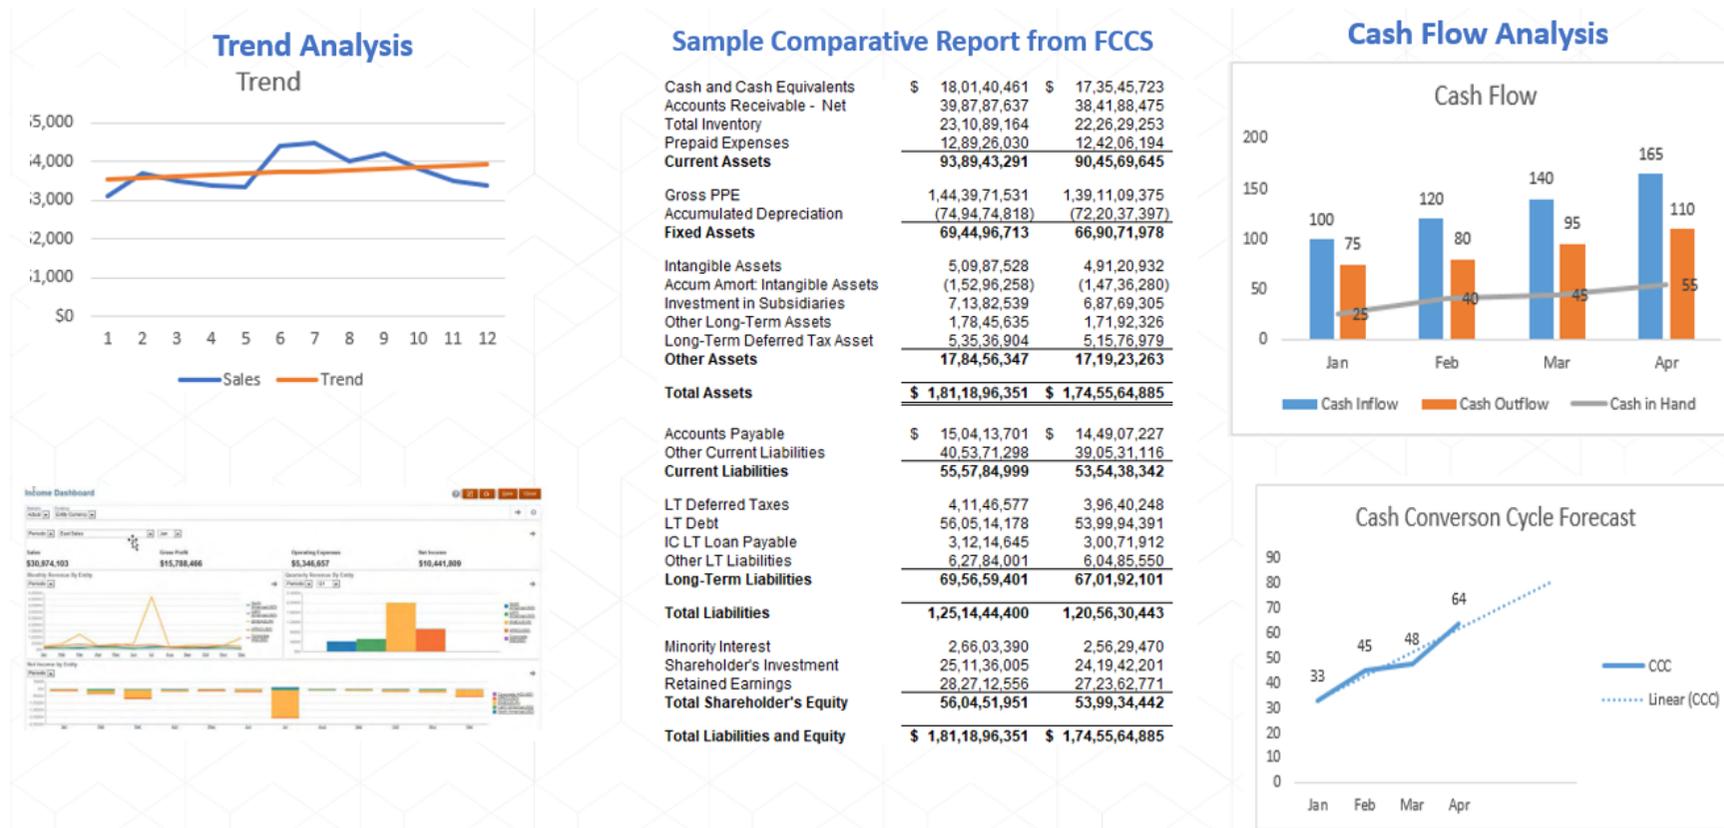
Example of Analysis of Financial Statements using FCCS

About Mohit Jain

Mohit is an Oracle ACE Pro and ODTUG EPM community Lead with 15+ years of finance & IT experience and a focus on the Oracle EPM technology stack and Business Intelligence.

He specializes in architecting, developing, and training of EPM set of tools for Controllership and FP&A applications, primarily in Oracle EPM Cloud tools (ARCS, FCCS, PBCS), Essbase, Planning, FDMEE, DRM and Reporting. He is a Hyperion certified consultant, an avid reader, and data enthusiast with a passion for open source, Machine Learning and AI ecosystem.

Mohit was a presenter at ODTUG Kscope21, Kscope20, UKOUG20, AUSOUG22, AUSOUG21 and INNOVATE 21 and has presented webinars and written articles for ODTUG on a couple of topics. In his free time, he loves to spend time with his wife and 2 kids.

About Megha Gupta

Megha is an Oracle ACE Pro, ODTUG EPM community Lead, having 16 years of experience, with proven mixture of business intelligence and technical knowledge on consolidation and Reporting tools and Oracle EBS to successfully capture requirements from the business perspective and translate those into designs that are efficient, effective and logical.

Megha was a presenter at ODTUG Kscope21, Kscope20, Kscope19, UKOUG20, UKOUG22, AUSOUG22, AUSOUG21 and INNOVATE 21 and has presented webinars and written articles for ODTUG on a couple of topics. In her free time she loves to play with her six years old daughter and go for walk.

Jasmin Fluri

Edition-based Redefinition – and why it is essential for your Continuous Database Delivery Setup

Continuous Delivery is the new de facto standard in software development. Continuous releasing is important in database development in order to continuously bring small changes into production and to keep feedback cycles small. This continuous releasing is based on the fact that interfaces are provided versioned and application components are decoupled. This decoupling must also be implemented between the database and the consuming applications. This article explains how Oracle Edition-based Redefinition (EBR) can support this.

Motivation

In order to establish continuous releasing of changes in the long term, it must be ensured that database changes can be brought into production without causing application downtime or redeployment. This decoupling between the database and the application can be implemented with the versioning of the database access layer, as is also known from the versioning of REST endpoints. Editioning the database access layer primarily enables the access layer to be changed, even if it is still in use.

But what is a database access layer? Since only PL/SQL objects can occur versioned within the database and no persistent objects like tables or sequences, the access layer acts as an abstraction layer of the persistent objects and consists of PL/SQL objects like views, packages, procedures, functions, triggers, libraries, synonyms and types. All these PL/SQL objects are editionable within the Oracle database.

Decoupled Release Cycles

Decoupling the database from the application with an editioned abstraction layer makes it possible to roll out database changes without redeploying them simultaneously. We thus have a versioned API, which is consumed by the application and can upgrade to new API versions. This is illustrated in the following image.

With a versioned API on the database side, applications can switch to newer versions on their own release cycle and thus aren't coupled to database changes.

This decoupling creates separate release cycles and allows two teams to implement features independently without having to coordinate adjustments.

EBR - What is it?

In the Oracle database an object is defined with its name, the associated user (schema) and its edition. The purpose of the edition is to support multiple copies (versions) of a PL/SQL object. The feature that enables the versioning of objects is called Oracle Edition-Based Redefinition and allows to implement database schema changes in a highly available way and to provide versioned changes in new editions of the application.

Preconditions

Before you can use EBR in the project, several preconditions must be met:

- The user who will create editions and perform schema changes needs the appropriate permissions. For the edition management at least CREATE EDITION and DROP EDITION grants are required.
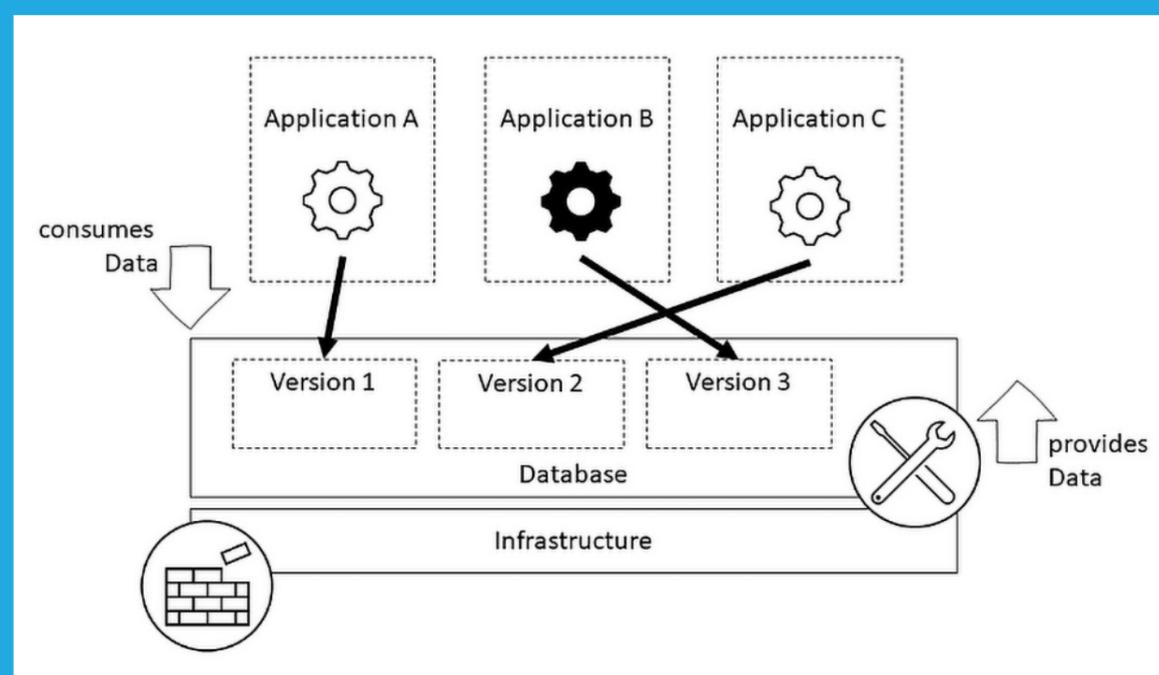
Database access layer with different versions (editions) which are used by the consuming applications for database access.

- The database is accessed through an access layer, which consists of PL/SQL objects (views, procedures etc.). Only PL/SQL objects are editionable. There must be no direct access to tables, because this would bypass the editioning and break the access if those objects change.
- Editioning must be enabled for all affected schemas.

When shouldn't you use EBR?

EBR introduces complexity as every additional feature does. Remember that your use case might not necessarily need it after all and be careful in your decision to adopt a new feature. If your application and your database have the same release cycle and downtimes aren't an issue, chances are you might not need EBR as an additional feature in your project.

When should you use EBR? / What problems does it solve?

Modern applications today are required to "run forever". This requirement extends from the hardware up to the application. Thus, every maintenance must be online, i.e. involve zero downtime. EBR provides the following key benefits to minimize downtime for an application upgrade:

- Allow arbitrary changes to a set of artefacts implementing the application's database of record.
- Utilize both pre-upgrade and post-upgrade applications simultaneously (hot rollover).
- Maintain uninterrupted availability of the application across editions (live operation).
- Ensure no noticeable negative impact on performance.
- Deployment of changes in the privacy of a new edition without exposing them directly.

When deploying changes into production, the use cases for EBR arise when the consuming applications either cannot be deployed simultaneously and cause downtime or if multiple applications are consumers of the same database application that cannot adopt the changes simultaneously.

Another use case is simply when your application requires a zero-downtime database change.

EBR in development projects

When we use EBR in our development project, we need to consider how we create editions and how editioning is handled in general. EBR can basically be used in two ways: Manual assignment of editions in the migration scripts or automatic assignment of editions in the pipeline previously to the deployment of migration scripts.

Manual Editioning
If we choose manual editioning, we define ourselves for which changes new editions will be created and for which not - maybe this is not always necessary. With manual editioning, the new edition is assigned manually in the database migration script in each case. The advantage of this variant is that it is guaranteed that the same editions exist over all environments, which is not guaranteed with the automated editioning. A precondition for manual editioning is that the database migration scripts are created manually and not generated by a framework.

Automated Editioning
With automated editioning, a new edition is created for each database schema migration and the change is deployed to this new edition, which is automatically created in the pipeline and then set as the new default. The developers have nothing more to do with the editioning in that it is handled automatically before and after deploying their database changes through the automated pipeline.

Old editions can be cleaned up either nightly or by a predefined interval by housekeeping jobs so that only a few editions are available. However, the number of editions within a database is unlimited.

Providing EBR functionality as a package

To make EBR easier to use, offering the EBR functionalities as a package in a database utility schema is usually worthwhile. In this way, the application schemas don't need to be highly privileged and grants can be restricted to the specific functions and procedures provided.

Switching Versions

With Oracles' EBR, there are different ways to define the edition when accessing the database over the access layer. I present here four that can be used alone or in combination. Before an edition can be accessed, it must be made available to users — **learn more in the Oracle documentation about it**.

1. Default Database Edition

Once an edition is created inside the database, the easiest way to make all applications use it is to make it the default edition. The default edition is used for all incoming database connections that do not explicitly specify the edition they want to access.

```
ALTER DATABASE DEFAULT EDITION = myEditionName
```

2. Database Service

If several versions need to be accessed simultaneously, one possibility is to create different database services with preset editions. This can be done with Oracles' DBMS_SERVICE. CREATE_SERVICE procedures edition parameter.

```
DBMS_SERVICE.CREATE_SERVICE(
    service_name            IN VARCHAR2,
    network_name            IN VARCHAR2,
    goal                    IN NUMBER DEFAULT NULL,
    dtp                     IN BOOLEAN DEFAULT NULL,
    aq_ha_notifications     IN BOOLEAN DEFAULT NULL,
    failover_method         IN VARCHAR2 DEFAULT NULL,
    failover_type           IN VARCHAR2 DEFAULT NULL,
    failover_retries        IN NUMBER DEFAULT NULL,
    failover_delay          IN NUMBER DEFAULT NULL,
    clb_goal                IN NUMBER DEFAULT NULL,
    edition                 IN VARCHAR2 DEFAULT NULL);
```

Oracle 19c Documentation — DBMS_SERVICE.CREATE_SERVICE package definition - **source**

If the new database service was created with the edition, it could be used by the application to establish connections to the database. By default, it always connects to the set edition.

3. JPA Datasource with Edition

The application side can also set the edition via JPA data sources. The name of the edition can be specified in the property oracle.jdbc.editionName. This edition will then be used when establishing a connection through the JPA data source.

```
properties.setProperty("oracle.jdbc.editionName", "myEditionName");
```

An example of a data source with the editionName property is shown below.

```java
@Bean
    @Primary
    protected DataSource dataSource() {
        DriverManagerDataSource dataSource = new DriverManagerDataSource();

        dataSource.setDriverClassName(driver);
        dataSource.setUsername(user);
        dataSource.setPassword(pwd);
        dataSource.setUrl(url);

        Properties properties = new Properties();
        properties.setProperty("oracle.jdbc.editionName", edition);
        dataSource.setConnectionProperties(properties);

        log.debug("dataSource:{}", dataSource);
        return dataSource;
    }
```

The application can use multiple editions simultaneously with multiple JPA data sources with different editions set.

4. Manually setting the edition

An edition can also be set manually during a session. To do this, an ALTER SESSION command is issued in the session.

```
ALTER SESSION SET EDITION = myEditionName
```

After the ALTER SESSION command, no commit is necessary; the new edition is immediately available.

Conclusion

The only way to achieve high available database deployments is to decouple your database from your application and to build a versioned access layer that allows to deploy objects as new editions. The Oracle database offers those functionalities out of the box.

The application and use of editions can be designed in many different ways. Therefore, it is essential to use the method that best harmonizes with your release cycle and work methodology to avoid possible barriers.

Resources
PL/SQL Packages and Types Reference
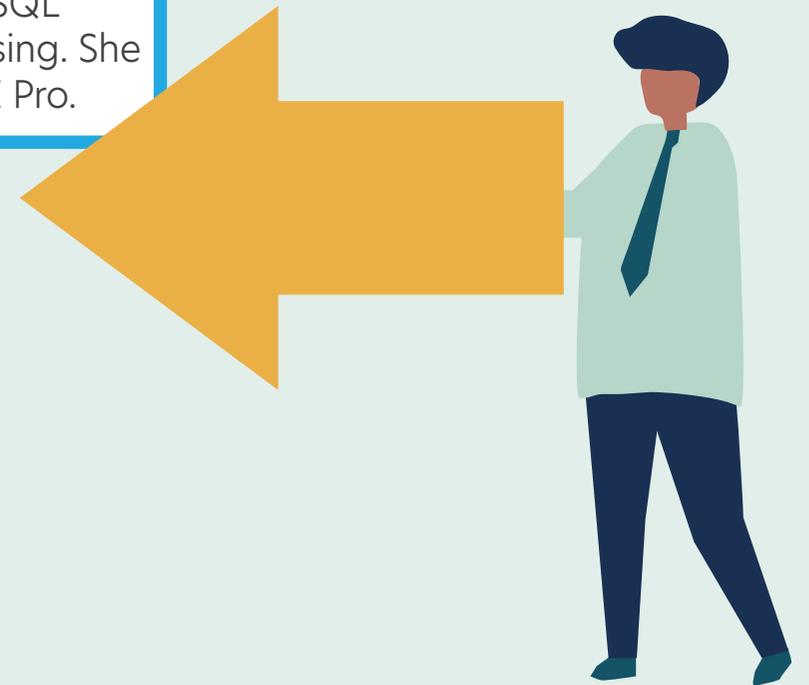**My EBR Blog Post Series - @DBoriented**
Database Development Guide
OracleConnection (Oracle JDBC API Reference)

About Jasmin Fluri

Jasmin works as a Database Automation Engineer at Schaltstelle GmbH and lectures on software engineering at the University of Applied Sciences North-Western Switzerland. Her focus as a database developer lies on continuous integration and delivery pipelines, PL/SQL development, data engineering and data warehousing. She is part of the Oracle ACE Program as an Oracle ACE Pro.

How Secured are your Reports on Oracle Cloud SaaS?

Rishin Mitra

*Lately, more and more customers have been moving away from their legacy on-premises systems and implementing their applications on Cloud SaaS solutions, Oracle Cloud SaaS being the preferred option for most of the customers. With that, there comes a need to develop reports on Oracle Cloud, which can facilitate business with their day-to-day activities. Oracle already provides many reports, out-of-the-box, for end-users. These reports are fully secured. However, these reports might not always be sufficient for business. Hence, you either need to customize these reports or develop reports from scratch. But how do we ensure that a Finance user does not have access to Salary Reports with Oracle HCM, or an Human Resource user can't access any Finance users. Moreover, how do we confirm that we are not exposing unwanted data to users. In this article, we'll discuss the basics of how we can ensure **Object Level Security** and **Data Security** for the reports in Oracle Cloud SaaS.*

Before we discuss how the security comes into play for Reports, let us have a look at how the security works within Oracle Fusion Cloud. Security in Oracle Fusion is controlled through **Role Based Access Control**, otherwise known as **RBAC**. **RBAC** enables organizations to create **R**oles based on specific job functions performed by different users within the organization, and then these roles are assigned to users with the appropriate permissions.

Fundamentally, there are two types of security - **Function Security** and **Data Security**.

- **Function Security**: Function Security is defined more from the UI perspective, as in, what kind of screen a user can see, what kind of task the user can perform, what kind of functions a user can access, what kind of actions users can do, all these come under Function security.
- **Data Security**: Data Security deals with what kind of data a user can see. In other words, on what kind of entity or object the privilege would be imposed on - Ledger, Business Unit, Operating Unit, etc.

Image 1 shows the different components of implementing Security for OTBI Reports within Oracle Fusion Financials.

We can summarize Security in Oracle Fusion in a single statement – "**WHO** can do **WHAT** on **WHICH** set of data" **(see chart on the right)**.

This, within Oracle Fusion is controlled or managed through Roles and Data Access.

Typically, there are different types of Roles present within Oracle Fusion Applications. There are as follows:
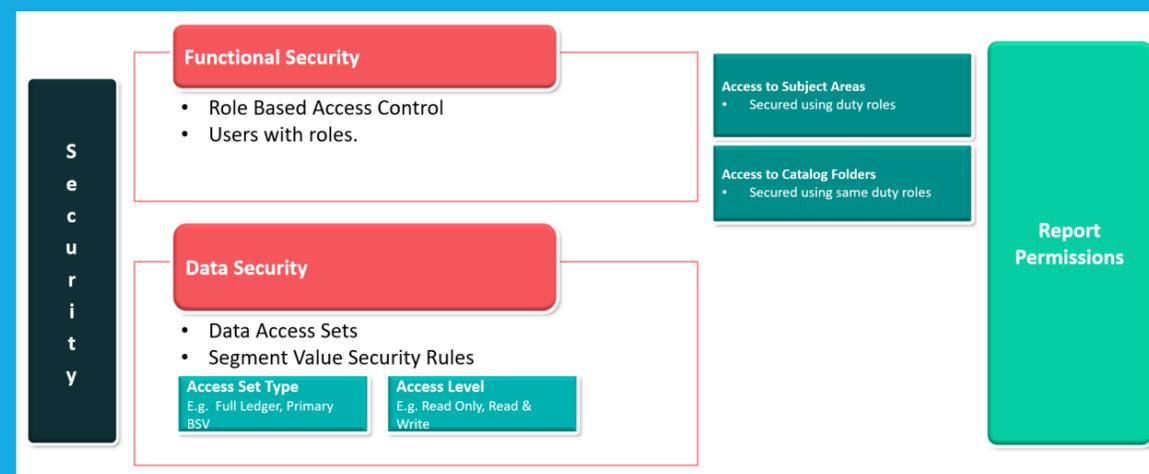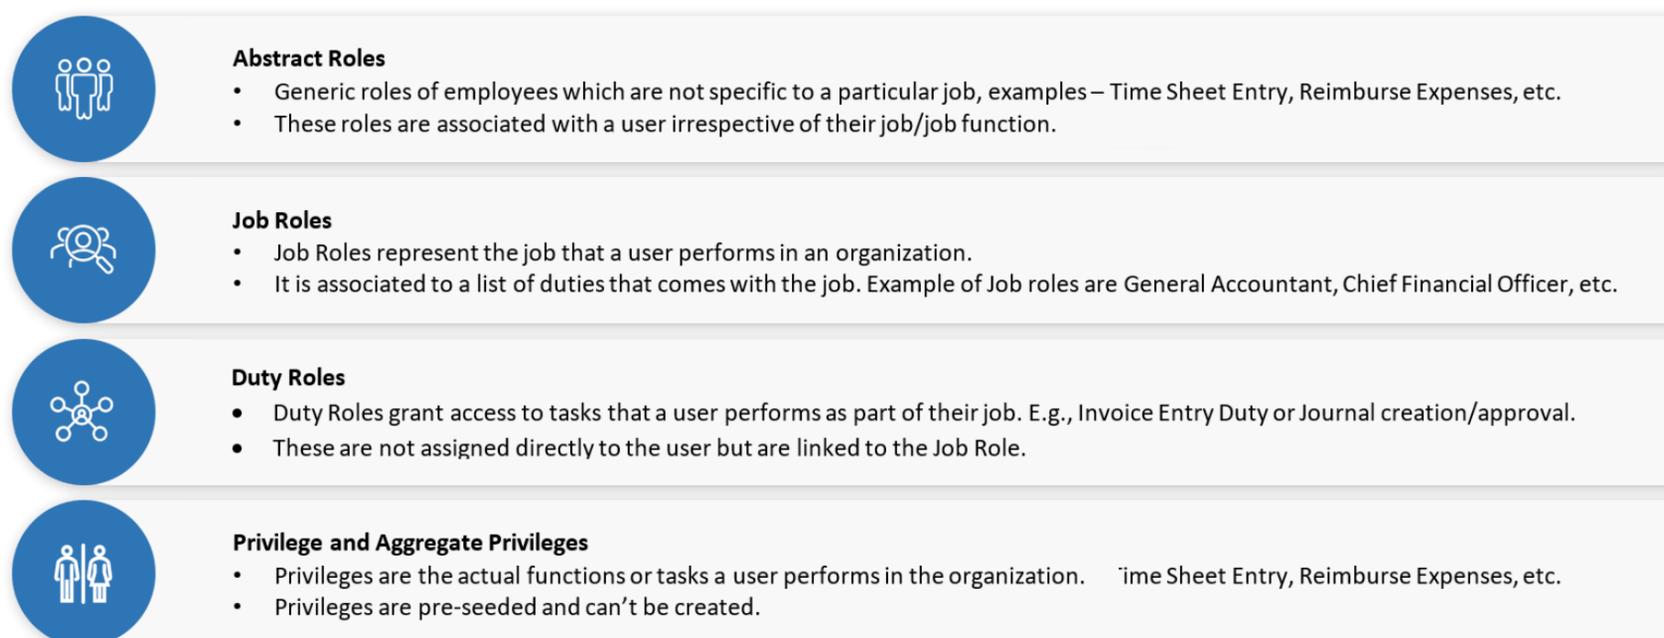
Image 1: components of implementing Security for OTBI Reports within Oracle Fusion Financials.

WHO: Individuals/ End-Users
Andrew, Darryl, John, etc.

WHAT: Actions performed by Users
Post Journals, Make Payment, Close Period

WHICH: Entity on which Users perform Actions
US Ledger, UK Business Unit, UK Asset Book

Abstract Roles
- Generic roles of employees which are not specific to a particular job, examples – Time Sheet Entry, Reimburse Expenses, etc.
- These roles are associated with a user irrespective of their job/job function.

Job Roles
- Job Roles represent the job that a user performs in an organization.
- It is associated to a list of duties that comes with the job. Example of Job roles are General Accountant, Chief Financial Officer, etc.

Duty Roles
- Duty Roles grant access to tasks that a user performs as part of their job. E.g., Invoice Entry Duty or Journal creation/approval.
- These are not assigned directly to the user but are linked to the Job Role.

Privilege and Aggregate Privileges
- Privileges are the actual functions or tasks a user performs in the organization.   Time Sheet Entry, Reimburse Expenses, etc.
- Privileges are pre-seeded and can't be created.

Image 2: Users are assigned Job Roles and Abstract Roles.
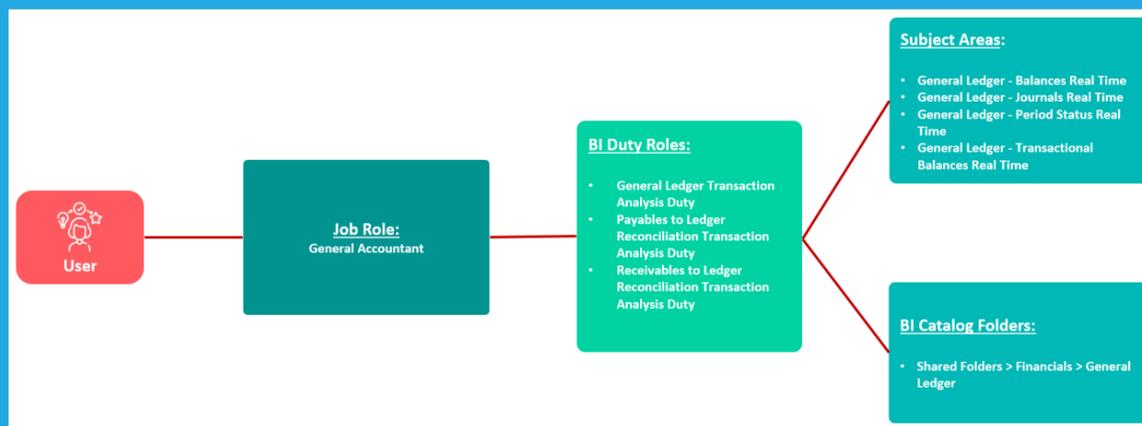
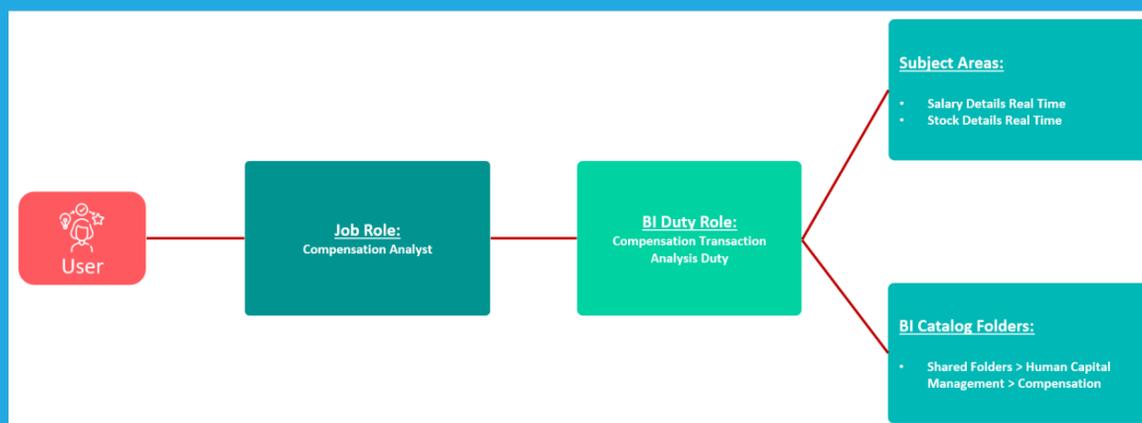*Image 3: A User is assigned the "General Accountant" Job Role.*

Image 4: A User is assigned the "Compensation Analyst" Job Role.

Role Hierarchy

In Oracle Fusion, each role can be linked to other roles in a parent-child format to form a hierarchy of roles. Users are assigned Job Roles and Abstract Roles, which inherit Duty Roles and their associated Privileges. Duty Roles in turn can inherit Privileges from subordinate Duty Roles. All roles are pre-seeded. One can either copy or customize roles/create from scratch as per customer requirement.

Oracle Fusion Transactional Business Intelligence Security

Oracle Transactional Business Intelligence or OTBI is a real-time, self-serviced reporting tool, using which business users can create analyses or dashboards to support decision making. Users can use the predefined or seeded Subject Areas to create analyses and save them within the BI catalog folder structure. Within Oracle Fusion, the **Subject Areas**, the Business Intelligence Catalog Folders, and the Business Intelligence Reports available to the users are controlled by **OTBI Transactional Analysis Duty Roles**. As mentioned above, since Duty Roles can't be directly assigned to users, they are inherited by Job Roles which are then assigned to the users. For examples, a Finance user is assigned the **"General Accountant"** Job Role **(see image 3)**, which inherits 3 Duty Roles - **General Ledger Transaction Analysis Duty, Payables to Ledger Reconciliation Transaction Analysis Duty and Receivables to Ledger Reconciliation Transaction Analysis Duty**. This gives the user access to the Subject Areas - **General Ledger - Balances Real Time, General Ledger - Journals Real Time, General Ledger - Period Status Real Time** and **General Ledger - Transactional Balances Real Time**, using which the user can create analyses.

Detailed list of the Seeded Subject Areas, Job Roles and the corresponding OTBI Duty Roles for Oracle Fusion Financials can be found **here**.

Like Oracle Fusion Financials, Oracle Fusion HCM also secures the Subject Areas, BI Catalog Folders and BI Reports through the OTBI Transactional Analysis Duty Roles. **Image 4** shows how the role inheritance can control the access for **"Compensation Analyst"** Job Role.

Oracle Transactional Business Intelligence Data Security

The data that a user can see when he runs an OTBI Report is same as the data that he/she has access to with the application. This data access is managed by Duty Roles linked to **Security Profiles**. The Transactional Business

Intelligence Duty Roles are granted with one or more security policies, which in turn gives access to data. So, when creating a custom job role for users who needs access to OTBI Reports, then the correct duty role must be assigned to the job role.

For example, if a custom Job Role is created for users to access Compensation reports, then both **HCM** and **OBI** versions of the Transactional Analysis Duty Roles for the Compensation module as shown in the table below.

Duty Role	Role Code	Role Category	Version
Compensation Transaction Analysis	ORA_FBI_COMPENSATION_TRANSAC-TION_ANALYSIS_DUTY_HCM	BI - Duty Roles	HCM
Compensation Transaction Analysis Duty	FBI_COMPENSATION_TRANSACTION_ANALYSIS_DUTY	HCM Analytics OTBI	OBI

Important Note: *It is highly recommended that the seeded Transactional Analysis Duty Roles should be added to the Custom Job Roles, instead of creating copies of the roles. If the roles are copied, then additional steps needs to be performed to manually secure the reports and folders within BI Catalog, using the copied roles, from the report/folder "Permissions".*

Oracle Fusion Financials also follow the same rule for role or mechanism for Data Access.

Business Intelligence Roles

Oracle Business Intelligence roles apply to all the objects present within the Oracle BI Catalog. These Permissions/Roles grant access to Business Intelligence functionality, such as the ability to run or author reports. Users need one or more of these roles in addition to the roles that grant access to reports, subject areas, Business Intelligence catalog folders, and data. There are four BI Roles available. They are as shown below on this page.

We can further control what actions each role or user can perform for the objects in the BI Catalog through the **Custom Permissions**. The different options available are given in **image 5**. Hopefully this would clarify how one needs to secure the report objects and the data being fetched in the reports.
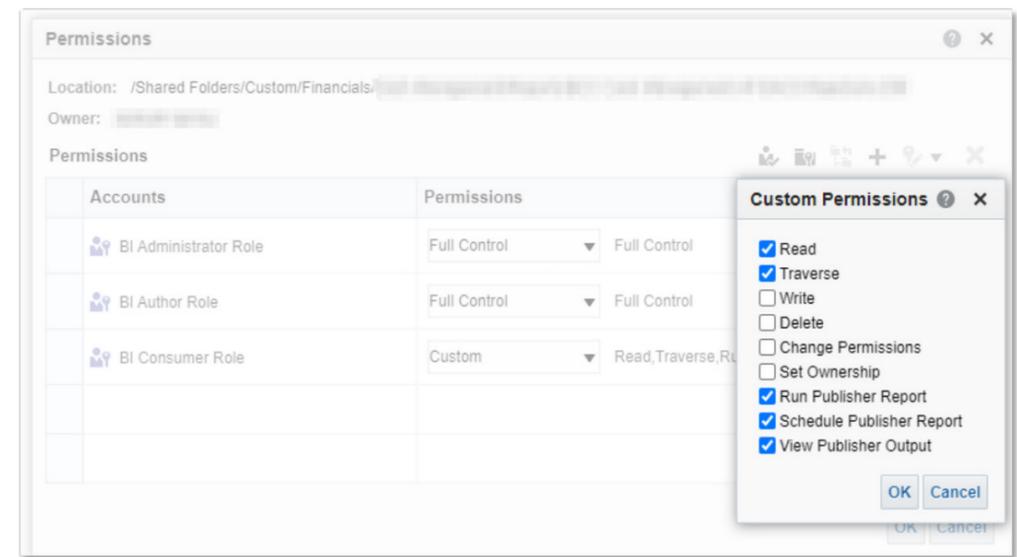
Image 5: Custom Permissions.

BI Consumer
• Allows users to Run Business Intelligence Reports
• OTBI Transaction Analysis Duty roles is inherited by predefined OTBI Transaction Analysis Duty Roles.

BI Author
• Allows users to create, edit, and run OTBI reports.
• BI Author Role inherits the BI Consumer Role.

BI Administrator
• BI Administrator Role is a superuser role, giving administrative permissions.
• BI Administrator Role inherits the BI Author Role.
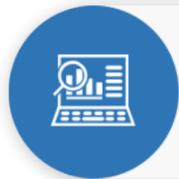
BI Publisher Data Model Developer
• Gives permissions to create and manage Data Models for BI Publisher Reports.
• Inherited by the Application Developer role, which is inherited by the Application Implementation Consultant role.

Techs & Nerds

Data Security for BI Publisher Reports

Since Oracle BI Publisher otherwise known as Oracle Analytics Publisher issues SQL queries directly to the database, the data returned is not secured. Hence, the number of users who can create Data Models should be restricted. It is highly recommended that **BI Publisher Data Model Developer Role** should only be assigned to technical team who can write and issue queries directly against the database.

Now, to restrict the data being fetched in the BI Publisher reports, for **Oracle Fusion HCM**, one needs to make use of the available **Secured List Views**. If the select statement is joined with one of the secured list views, then the data returned will be restricted based on the security profiles that are assigned to the roles of the user or the Areas of Responsibilities of the user.

For **Oracle Fusion Financials**, for each module, the technique of restricting the data in the reports are different. While for General Ledger, the data is controlled through incorporating **Data Access Sets** or **Segment value security rules** by making use of **Flexfields** within the BI Publisher Data Model, whereas for Payables and Receivable, the data security is controlled through **Fine Grain Access Control**.
Details of implementing data-security for custom BI Publisher reports can be found in the documents below:

- *Segment Value Security for Chart of Accounts in Custom BI Publisher report (Doc ID 2520305.1)*
- *Fusion Security For Custom FA Reports In Financials Cloud (Doc ID 2315270.1)*
- *Data Restriction Based On Procumerement Agent Access When User Runs Custom BIP Report (Doc ID 2514012.1)*

Summary:

- OTBI Reports security leverages **Oracle Fusion Users** and **Job Roles**.
- **OTBI Duty Roles** control access to Subject Areas and objects in the BI Catalog.
- Uses same data security as Oracle Fusion applications. OTBI Users will only be able to access data from assigned Data Access Sets Roles.
- Data returned in BI Publisher reports should in controlled through **Secured List Views** in **HCM** and through

References:

- ***Security and Reporting (HCM)***
- ***Overview of Financial Reporting Security***
- *Segment Value Security for Chart of Accounts in Custom BI Publisher report (Doc ID 2520305.1)*
- *Fusion Security For Custom FA Reports In Financials Cloud (Doc ID 2315270.1)*
- *Query To Extract Open Purchase Orders And In-Transit Shipments In SCC (Doc ID 2857823.1)*
- *Data Restriction Based On Procurement Agent Access When User Runs Custom BIP Report (Doc ID 2514012.1)*
- *Fusion SaaS R12: Implementing Data Security for Oracle Fusion Financials in Custom BI Publisher Reports*

About Rishin Mitra

Rishin is an Oracle ACE Pro, having around 11+ years of industry experience in Applications Development. Currently he is working as Oracle Cloud Technical Lead with Version 1. He has got extensive knowledge of working on Oracle E-Business Suite, Oracle ERP Cloud, OBIEE and Informatica, and has been involved in development and large-scale global delivery of various projects in products and services. Before joining Version 1, he worked as a Senior Applications Engineer at Oracle and has worked on multiple legislative projects for EMEA as part of the Oracle E-Business Suite Globalization team. He is very active in the community and regularly participates in OUG conferences as a speaker as well as a volunteer. He is a member of the MASH Program.

I worked with Oracle databases for more than 20 years and I know it pretty well, from deep understanding, through installation and upgrading, to tuning and configuring complex environments such as Data Guard and RAC. However, I made a huge change in my life lately, and as part of it I started a new job where I mostly work with MySQL. This article is a summary of my (somewhat short) experience with MySQL while comparing it to my (rather long) experience with Oracle database.

1 Purpose

MySQL and Oracle serve slightly different purposes, and this will explain some of the stuff I'm about to describe. While both are RDBMS (well, kind of), Oracle is a fully-featured-all-purpose database, while MySQL was created as a database for the web. I added the "kind of" comment in the previous sentence, because MySQL architecture is different than Oracle. It has different engines that we can use and each engine works in a different way - it saves and processes data in a different way. Some of these engines are not really (or at all) relational. For this article I'll only refer to the INNODB engine, which is the default MySQL engine and is completely relational.

2 Installation and Maintenance

MySQL is a very lightweight database, while Oracle is really not. The installation of MySQL is super easy and fast, while Oracle installation is a more complex task, and you should know what you're doing (Oracle did a good job with Oracle XE to make it more "ready out-of-the-box", but they are still not there with their Standard/Enterprise editions).

MySQL also feels lightweight while working with it. It starts up faster and feels quicker in general. Oracle on the other hand takes its time.

In our environment we are working with micro-services and some of them have databases. The fact that MySQL is simple to install and use makes it easy to manage dozens or hundreds of these databases or even more. Managing a large number of Oracle databases is a bit tricky, but as I said, the ways you work with these databases are different. Large organizations do have many Oracle databases, but each is usually large and complex, many of them serve many services and run a large amount of different queries and use-cases. With MySQL and the micro-service architecture, each MySQL database may run only a handful of different queries, and serve a very specific set of use-cases.

3 Cost

This is probably the deal breaker here for many companies. It doesn't matter what I explain later on in this article, and even if I show that Oracle is superior in every aspect (it's not), the cost is a huge issue. MySQL (the community version) is free. You can install it in your own data center as much as you'd like (selling it with your product will cost you though). Use (almost) any feature you like as much as you'd like and you'll be fine, only very few features are "enterprise" and come at a cost (please note: I'm not a license person, if you intend to use MySQL with your product, make sure you comply with the license and seek legal advice).

Oracle database, on the other hand, is expensive (Oracle XE excluded, which is free but also quite limited), and on top of that, many features are at extra cost (partitioning, some security features, some compression features, etc.). The Oracle license is also per CPU, so as you want to scale (up as you increase your server power, or out by adding more servers) it will cost significantly more.

But even though I know that cost is a huge factor for many, let's leave cost aside and keep the discussion purely about the technology.

4 Performance

Performance is a big topic and I can talk about it a lot. I'll only focus here on a few aspects as examples.

Index Scans
Being a database for the web (websites and website-based-application), MySQL focuses on index-based queries, and these are blazing fast (especially seek by the primary key). Seeing tons of queries (and simple inserts) running for much less than 1ms each is quite impressive. I never did a head-to-head Oracle / MySQL comparison myself, but I've heard people who did say that MySQL was faster in these scenarios.

More Complex Queries
What about other types of queries? What about joins? Here MySQL is far behind Oracle. Its optimizer is much simpler (I'll talk about it later on) and complex access methods do not exist (there is no index skip scan, OR expansion, and until MySQL 8 it didn't even have hash join). Why? Because

these are not the use-cases MySQL is meant to solve. It needs to serve websites and web applications, and it does this really well.

With micro-service architecture, the concept is that each service has its own database. Does it have to be MySQL? Of course not. MySQL will be great for the very specific OLTP-style operations. You want to have more complex queries? The relevant micro-service can have PostgreSQL, Oracle, Snowflake, or whatever the right solution is for its use-cases. In comparison, Oracle database is a much more full-featured database for the enterprise, implementing tons of options and optimizations to support anything you throw at it, and it works well. So, if you want to work with a central database for many different use-cases and combine OLTP and analytical usage, no sweat at all.

Shared Pool and the Optimizer
While MySQL has a cache for data dictionary objects, it doesn't have what Oracle calls "SQL area". Every time a query is sent, MySQL will parse and optimize it in order to decide how to execute it. In Oracle database, there is a specific cache for holding parsed SQL statements, so the second time a query is sent, it can be found in the cache, eliminating (in some cases) the need to re-parse and optimize it.

On one hand, Oracle is smarter, having a cache reduces resource consumption. However, I already mentioned that MySQL is focused on specific use-cases, and that its optimizer is actually much less complicated than Oracle's. While Oracle needs to search the cache for the query (and needs to protect the memory and manage it), compare it (there are many sophisticated steps here like estimating if the values of bind variables will benefit from a different plan, and checking if the same object names are actually the same object, or maybe just the same name for objects in different

Techs & Nerds
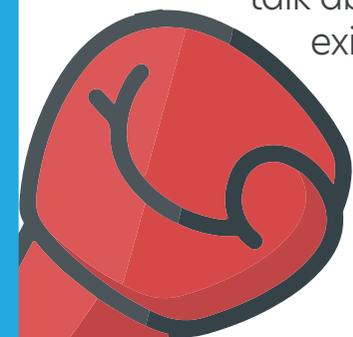
schemas, etc.) and then decide if it can reuse the same plan, MySQL already parsed the query, optimized it with its much simpler optimizer and executed it.

And mentioning bind variables, if you don't have a cache for the query text in MySQL, you don't really care about bind variables, do you? MySQL has bind variables and prepared statements, but these are mainly used to make developers' lives easier, and to fight SQL injections. There is a small performance benefit of using prepared statements, but again, it's a simple database with simple architecture, so everything is fast anyway and the benefits are rather small.

5 Dropping a Table

This simply blew my mind. While with Oracle I never even thought about something like "drop table" to be a problem, apparently in MySQL it is (at least until 8.0.23, which was released in early 2021). Apparently, when you drop a table in MySQL (until 8.0.23, which I'll stop mentioning now), the server scans the entire cache for blocks of this table, and cleans them up. Then it deletes the file from the filesystem (in MySQL every table is a file). Until these two operations are finished, the entire MySQL server waits (there are locks and stuff involved), and this can completely hang your application if the table being dropped is large. I wouldn't even have thought about this unless I had heard that this is actually a problem. If you want to read a bit more about this, I liked this **blog post**.

I'm sure there are more examples, some will show that MySQL behaves better, and some will demonstrate that Oracle

makes more sense, but since I encountered this specific one, I had to add it here.

6 Replication / Standby Database

Comparing this is a bit tricky, because it's a bit like comparing apples to oranges, but I'll try.

Like everything in MySQL, replication is quite simple as well. The concept is similar to Oracle's logical standby database (which I've never used as in my opinion it always had a bad reputation. I might be wrong, though). There are a few options here which I'm not going to get into, but in general it's super easy to create a MySQL replica database and maintain it. You can also create quite complex scenarios with multiple replicas.

In Oracle, standby database (a physical one at least, which is the one I've worked with a lot), is a super robust solution that works remarkably well. Again, I know that I'm not really comparing apples to apples (because I'm comparing Oracle's physical standby to MySQL logical replication), but I'll show some examples of the different behavior hereinafter.

In MySQL, the replicas get the changes from the source, and it's a one-way communication. This means that the source can and will delete the binlogs (logs with changes that need to be copied to the replicas in order to be applied there), even if they are needed. Oracle, when using RMAN (Recovery Manager, which is Oracle's utility for backup, recovery, and some maintenance), will never delete archive logs (more

Techs & Nerds
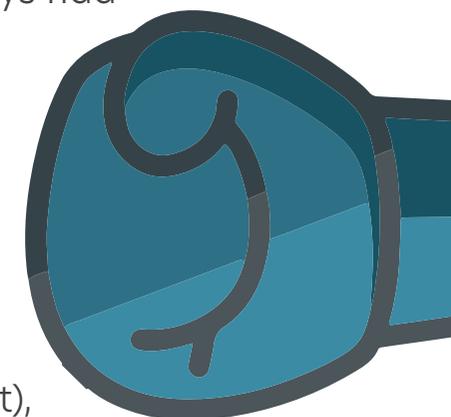
or less the equivalent of binlogs for this purpose) that are needed (unless it's in order to avoid a production issue).

In Oracle, all of the databases in the architecture know about each other (if using the DG broker, which is very common) and the configuration is automated. This means that when you failover to one of the standby databases, all the other databases will automatically adjust and become standby databases of the new primary. In MySQL, it's our responsibility to reconfigure the replica databases, and, depending on your configuration, it's not always possible. There are some tools that can help here, but nothing out of the box.

Since in MySQL the replication is logical and simple, it has a few advantages. One is fixing the replication in case something broke it. Again, I'm not familiar with Oracle logical standby and don't know if this is possible, but with physical standby, as it recovers binary data, you cannot fix anything manually. There are features and tools to help when something happens, but it's usually not that easy. Another thing with MySQL is two-way replication. This means that you can have two databases active and online and getting synchronized (one way or both); in this case, moving the application between them is easy and quick, so migrations are very easy and require no downtime. With Oracle database, you will have to use tools like GoldenGate to achieve this functionality. These tools are much more complex (and expensive).

7 Some Small Cool Stuff

There are some things that are, though small, super nice and convenient with MySQL. The command line tool is nice, the output is neat, you can use the arrows to navigate between queries, and more. SQL*Plus feels so old compared to it. I know that Oracle uses SQLcl now, I have to admit I have never used it, so my comparison is still to SQL*Plus.

Regarding the data dictionary, both are quite good, but when you want to see something simple like a table and its indexes, MySQL's "show create table" is simply elegant. Much more than trying to figure out how to call dbms_ddl package, or use dba_tables and dba_indexes in Oracle.

8 Diagnostics

Last thing I'm going to mention here is diagnostics. Oracle's AWR is an amazing tool and everything I'll say about it will still not give it the praise it deserves. MySQL has the performance schema which includes diagnosis information, but I think Oracle's AWR is much more extensive. MySQL, on the other hand, has a few tools Oracle doesn't have like the "slow log" that logs all queries that took more than X seconds to run. This can be very useful, even though Oracle people will cringe when hearing something like this, as the load and amount of data it will write will be huge with Oracle databases.

I'm still quite new to diagnosing MySQL issues, so this is what I have to write so far, but I'm sure I'll return to this topic and write more about diagnosing different issues and performance problems in the future.

9 Conclusion

To conclude this article, I'll start by saying that I'm still biased. I love the technology Oracle brings, and it's a very powerful database that can do almost anything. It's very mature and full of great features to provide scalability, overall performance, and tools. But, and it's a big but, it's very expensive and not as flexible as MySQL (for example, developing a product that is based on many micro services, each with its own database, will not be easy with Oracle).

As I continue to work with MySQL (and soon other databases as well), I'm sure I'll start liking them and see the pros and cons each of them brings.

About Liron Amitzi

Liron has been working with databases for more than 20 years. He specialized in Oracle database for most of his career as a senior consultant, working with a large number of companies in various fields and managed an Oracle Professional Services Team. Lately he became the DB team leader at monday.com, where he works with MySQL, PostgreSQL, and NoSQL databases.

Liron is a former Oracle ACE Director and was active in Oracle User Groups (as a board member at ilOUG in Israel and the president of BCOUG in Canada). He is a well-known instructor, speaks at many conferences, and he also blogs on his website gotodba.com.
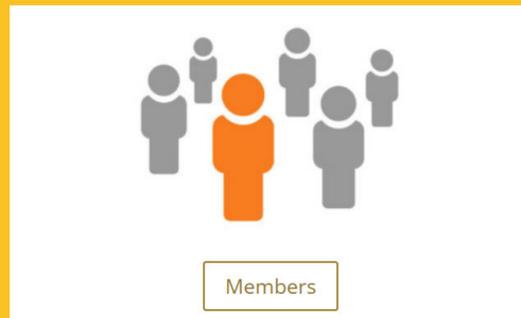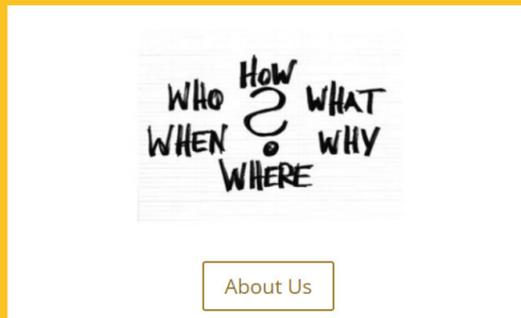
A New Community is Born

Mirela Ardelean

The Year 2020 brought Covid-19, lock-downs all over the world, an Oracle ACE Program restructured but nonfunctional and an Oracle User Group Program totally nonexistent.
On this complex - or complicated – landscape, a bunch of people, presenters and conference organizers, started having weekly discussions over zoom and daily conversations on a Twitter group.

It did not take very long until the idea of creating a totally independent group that would promote and support user groups and recognize those people having a valuable contribution on sharing their knowledge with the community became an almost constantly topic.

Discussion after discussion, Friday after Friday, we debated, we identified the values we want to recognize, the rules, and finally the name of the new group: Symposium 42 (or SYM42 for short).

SYM42 is a group totally independent of any vendor, with an initial focus on technologies that are part of Oracle, but open to expand in other directions as time goes by. It is run by a president, supported by a board. Martin Widlake, a well-known presenter and contributor to Oracle community, is the first President of Symposium 42 and the board members are Mirela Ardelean, Erik van Roon, Sabine Heimsath, Kamil

Stawiarski and Gianni Ceresa. Membership of SYM42 is by invitation and it is free.

SYM42 already celebrated its 1-year anniversary. We have 50 members, and we expect this number to increase soon. We plan to organize dedicated events but, until then, the first collaboration with a user group will happen on 12th of May, 2023: the POUG is organizing a one-day workshop where all presenters are members of SYM42. By the time when this article was sent to be published, there were only 12 seats left, so you need to hurry and secure your seat if you plan to attend this workshop in Warsaw.

If you want to know more about the new SYM42, the how and why and whom, check our **website** and read the **posts** written by Martin Widlake, aka El Presidente ;).

🌐 https://sym42.org   🐦 @sym_42

Oracle Community Tour EMEA – Looking back at 2022, Planning ahead for 2023

Mirela Ardelean
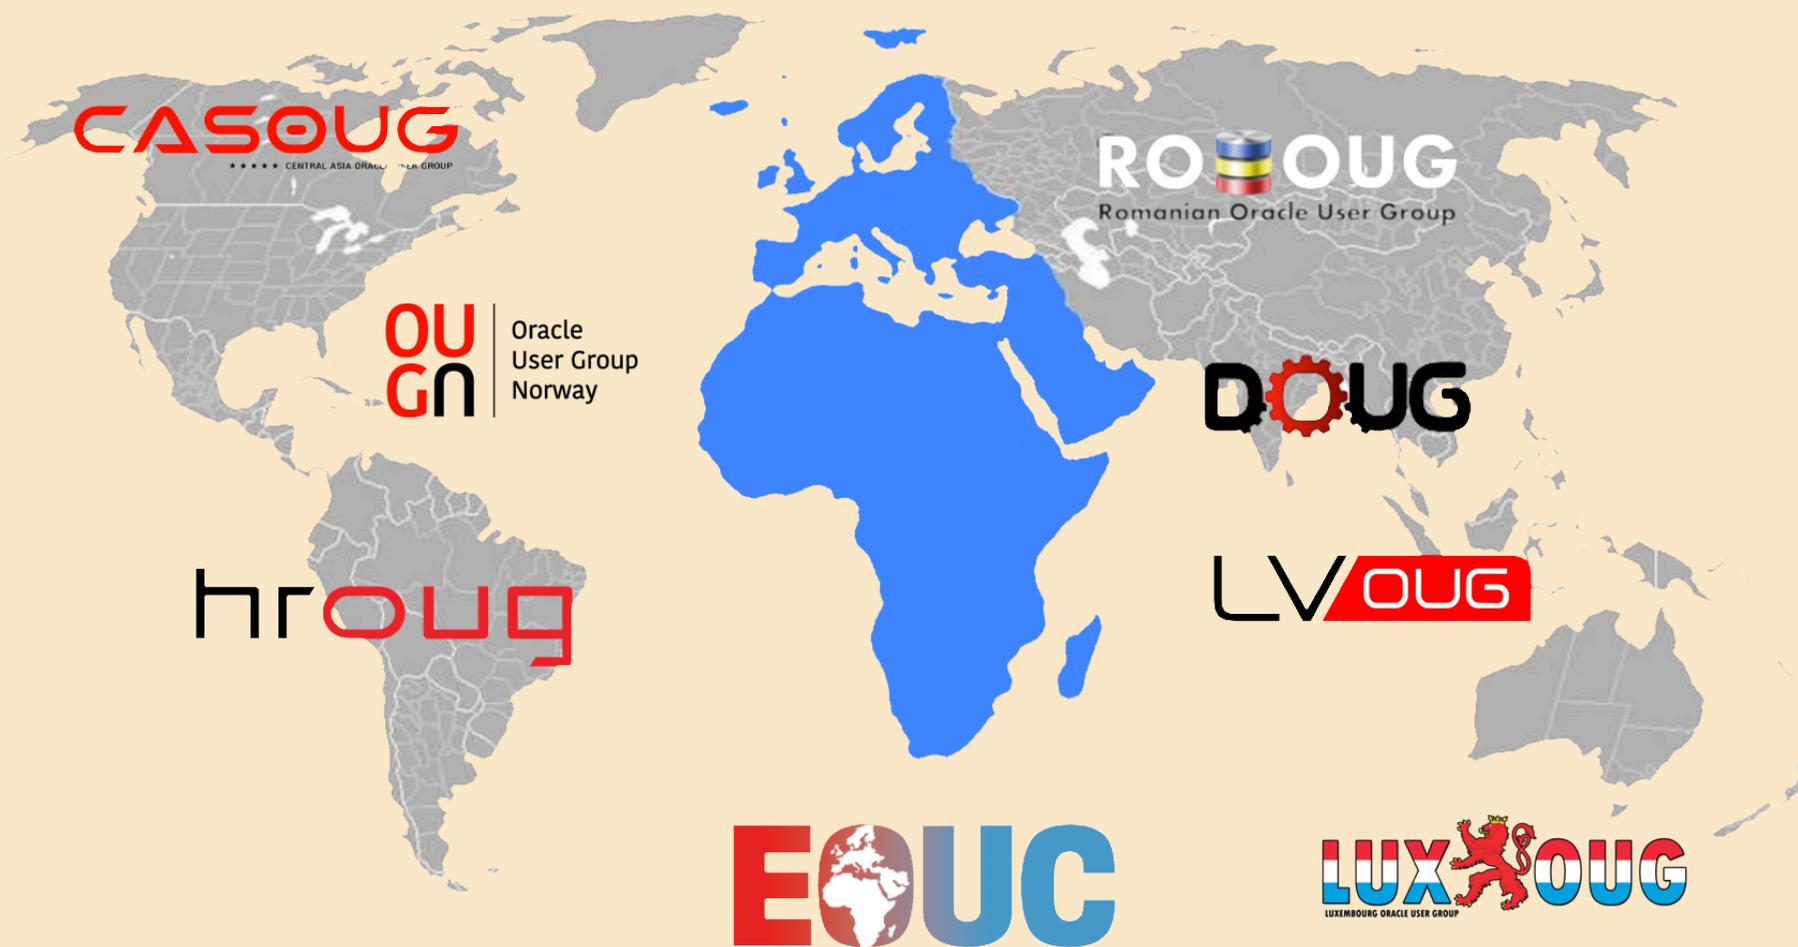
We all know that a community tour is a great way to bring together like-minded individuals and to provide opportunities for networking, education, and collaboration. But, to be honest, these benefits were not so obvious when the tours took place only online, in 2020 and in 2021.

Users & Groups

But after a 2-years break, the year 2022 brought back the possibility of organizing in-persons conferences. The organizing team that sits behind the Oracle Community Tour EMEA (formerly known under the name Oracle Groundbreakers Tour) thought that this perspective was something to consider so they sat together with 7 OUG leaders to build the tour route through EMEA.

The Oracle Community Tour EMEA 2022 took place between 26th of September and 13th of October and consisted of 4 online events, organized by CASOUG and LVOUG, one hybrid event organized by OUGN and 4 in-person events organized by DOUG in Copenhagen, HrOUG in Rovinj, LuxOUG in Luxembourg and RoOUG in Bucharest.

As always, the tour offered a unique opportunity to learn about the latest developments in Oracle technology, connect with experts in the field, and discover new ways to innovate and improve the business.

Read below what some organizers and presenters are saying about the 2022 tour:

Nelson Calero, Oracle Principal Consultant, Pythian | Oracle ACE Director | Presenter:

"As co-organizer of the Oracle Latin America Tour, I am aware of the hard work needed to organize a local user group event, and even more to coordinate that effort with other countries. And as a frequent speaker of our LATAM tour, I am also aware of the intensity and energy required to present at several events in a few days. Despite knowing that, I decided to submit sessions for this year's EMEA Oracle tour, for the first time. Because I also know it is an enriching experience, getting to know

great people from different cultures by the hand of the local organizers, exchanging knowledge with many local professionals interested in making good use of the same technology, and discovering beautiful places.

I made plans to skip a few events to make the schedule easier for me, so I feel sorry to have missed other beautiful events. But I am grateful to Oracle and the local user groups for making this tour happen because it was a fantastic experience. There was an excellent attendance at all the events I presented: Denmark, Luxembourg, Romania, and Croatia. People were engaged, asked questions, and were interested in the topics. The locations were really beautiful. All deserve to take more time than we had left before traveling to the next event.

The good thing is there will be another Oracle user group event around the corner where we can meet again, proving the Oracle community continues to be alive and exciting."

Rodrigo Mufalani, LuxOUG President, OCM & ACE | Organizer:

"The Luxembourg Oracle User Group was founded in 2019 and this year was very special for us, because we hosted, at Kyndryl Luxembourg HQ in Bertrange, our first face-to-face and it was amazing: we had 16 sessions for 2022 counting with a huge amount of community stars: 1 Oracle VP, 4 ACED and 4 ACE presenting a full day event about Oracle Cloud, tools, multi-cloud, Analytics, Development, DevOps and more.

The LuxOUG Day 2022 was part of EOUC EMEA Oracle Community Tour 2022, an event created by and for

Users & Groups

the Oracle community EMEA. We had 40 attendees, including the illustrious presence of Martin Bach (Oracle PM) and Pieter de Visser who came just to be part of the party, we had a lot of fun.

As the technology field is evolving very fast, we are always looking for ways to bring the best speakers with the best content to present at our events. With the Luxembourg Oracle User Group Day 2022, the speakers present new ideas and cutting-edge technology, and best practices on Oracle technologies. There were several countries well represented at the event and the audience exchanged a lot of ideas with the speakers during the sessions and during the breaks, while they tasted delicious snacks provided by our sponsor.

After the event, we had a great dinner, in a Brazilian BBQ restaurant in Luxembourg city center, again, another opportunity to celebrate the success of the event and the hard work and commitment of the speakers and organization. It's all about community spirit."

Kay Yu, Distinguished Engineer, Oracle Solution Engineering, Dell Technologies | Oracle ACE Director | Presenter:

"During the Oracle EMEA community tour 2022, I participated in the conferences at two cities: LuxOUG day in Luxembourg and ConTech by RoOUG 2022 in Bucharest, Romania. I was really overwhelmed by the organizers and the audiences of both conferences. Although it is the first LUXOUG conference, Rodrigo and his LUXOUG team did an great job organizing

the conference. From the early morning pickup to the beautiful speaker dinner at the fine restaurant at Luxemburg old town, I really appreciated the warm welcome and the gracious support that Rodrigo's LUXOUG team provided.

I also enjoyed the interactions and discussions with audiences at my two sessions.

Mirela and her RoOUG team prepared us a well-organized conference. I got almost a full house at my Oracle machine learning session and my other session on Oracle Sharding was also very well attended. After the session, we had to continue some of discussions outside of the room. Of course, the warm welcome by RoOUG at the speaker dinner and the second day's 9 course Oracle ACE dinner will be our forever sweet memory!

Really enjoyed this year's EMEA community tour, look forward to doing this again next year!"

If someone wonders what are the plans for 2023, we can easily answer: of course, the tour will happen!

The OUGs ready to host the tour are already known: AzerOUG (Azerbaijan), CASOUG (Kyrgyzstan, Tajikistan, Uzbekistan), DOUG (Denmark), HrOUG (Croatia), KSAOUG (Saudi Arabia), nlOUG (Netherlands), OUGF (Finland), SPOUG (Spain). The dates are still unset, but we aim to be "on the road" during middle of October.

Whether you are an Oracle user or simply interested in the world of technology, keep an eye on our announcement to be informed when the call for papers starts, when the registrations starts and when the tour will happen.

About Mirela Ardelean

Mirela is a Oracle Database Developer with more than 20 years of experience, with a dedicated focus on APEX. She is very much involved in Oracle User Groups communities, trying to build a strong collaboration between OUGs. She is the co-founder and the Vice President of RoOUG since 2012, a board member of SYM42 since 2020 and an Oracle ACE Pro.

Currently, Mirela became the President of EOUC, after 3 years of contribution as an EOUC board member.

She tweets as @Mirela_RoOUG about what's happening in the OUGs' world.

Call for Papers

LAOUC Community Tour 2023

April 15, 2023
https://laouc-2023-community-tour.mykommu.com/cfp

IT-Tage 2023

April 28, 2023
https://www.ittage.informatik-aktuell.de/call-for-papers.html

EMEA Community Tour 2023

June 1, 2023
http://ogbemea.com

HrOUG 2023

June 15, 2023
https://2023.hroug.hr/eng/Call-for-papers

Events

APEX World 2023

March 30 - 31, 2023
Zeist, Netherlands
https://www.nloug.nl/en/evenementen

SOUG Day

April 4, 2023
Winterthur, Switzerland
https://soug.ch/event/soug-day-4-april-2023

APEX Alpe Adria

April 21, 2023
Zagreb, Croatia
https://www.aaapeks.info/home

OUGN 2023 Spring Conference

April 24 - 25, 2023
Oslo, Norway
https://2023.ougn.no

APEX Connect 2023

May 3 - 4, 2023
Berlin, Germany
https://apex.doag.org/en/home

Blueprint 4D

May 8 - 11, 2023
Dallas, TX, USA
https://questoraclecommunity.org/events/
conferences/blueprint4d

POUG Workshop 2023

May 12, 2023
Warsaw, Poland
https://poug.org/en/edycja/poug-workshop-2023

DOAG Database with Exaday

May 24 - 25, 2023
Düsseldorf, Germany
https://datenbank.doag.org/en/home

MakeIT 2023

June 1 - 2, 2023
Portorož, Slovenia
https://2023.makeit.si/en

AppsConnected23

June 14, 2023
Fort Voordorp, Groenekan, Netherlands
https://theoraclecommunity.eu/event/save-the-date-june-14-2023-appsconnected23

CloudLand 2023

June 20 - 23, 2023
Phantasialand, Brühl, Germany
https://www.cloudland.org/en/home

ODTUG Kscope23

June 25 - 29, 2023
Aurora, CO, USA
https://kscope23.odtug.com

LAOUC Community Tour 2023

July 28 - August 15, 2023
https://laouc-2023-community-tour.mykommu.com

HrOUG 2023

October 17 - 20, 2023
Rovinj, Croatia
https://2023.hroug.hr/eng

IT-Tage 2023

December 11 - 14, 2023
Frankfurt, Germany
https://www.ittage.informatik-aktuell.de/konferenz/einfuehrung.html