

# FAST AND FURIOUS: CAPTURING EDGE COMPUTING DATA WITH ORACLE 19C FAST INGEST

- › [MASH Program:](#)  
First Steps for Speakers
- › [APEX Series Part 8:](#)  
APEX 21.1: What's new?
- › [What's Your Passion?](#)  
Janny Ekelson Talks  
about Running





What's your Passion with Janny Ekelson

page 7



Capturing Edge Computing Data With Oracle 19c Fast Ingest

page 12



With Great APIs Comes Great Responsibility...

page 22

Editorial

Submit Your Article!

## I. Work & Life

Comic: The Secret of a Successful Code Review

What's Your Passion?  
Janny Ekelson: "For me Running was Crucial in Surviving the Lockdown"

3

## II. Techs & Nerds

4

Capturing Edge Computing Data With Oracle 19c Fast Ingest **12**

With Great APIs Comes Great Responsibility ... Making your APIs Safe **22**

5

Oracle Application Express (Part 8): APEX 21.1: What's new? **26**

7

New Kid on the Block: OCI Language Service **32**

Mounting an Oracle Cloud Object Storage Bucket as a File System on Linux **39**

## III. Users & Groups

MASH – Mentor and Speaker Hub **44**

A Look Back at APEX connect 2021 [online] **48**

## IV. Past & Future

Call for Papers **50**

Events **51**

Contact us **52**

Legal notice **52**

# Editorial

Dear ORAWORLD Readers.

The first issue of the ORAWORLD magazine was published in May 2016. You now have the 25th edition in front of you.

This time Oracle ACE Director Jim Czuprynski shows which possibilities the Fast Ingest features of the database version 19c show in the focus of IOT. In the last issue we had Phil Wilkins's article about GraphQL. Now Phil continues with the great responsibility of making APIs safe.

The APEX user community is sure to be excited about the new possibilities Carsten Czarski has highlighted for the current APEX 21.1 release. This article is followed by a post of Adrian Png about Oracle Cloud Infrastructure Language Service.

I was particularly pleased about Kim Berg Hansen's article on the MASH Program. I hope that we can win some newcomers for lectures at user group conferences. A long-time companion in the Oracle community is Janny Ekelson from the Benelux User Group. But this time not with technical input but rather talking about having become a long-distance runner.

I can still remember the user group meeting when Ami Aharonovich from the Israel Oracle User Group proposed to publish a newsletter for better communication and identification with one another. Soon afterwards, Fried Saacke from the German Oracle User Group took up this idea and initiated ORAWORLD. Since then, I have been editor-in-chief of



this magazine. With this issue I am now retiring and handing this task over to Mirela Ardelean from the Romanian User Group.

I am very happy that Mirela takes over this responsibility. Mirela not only has many years of deep knowledge of Oracle. As a member of the EOUC Board, she is also a recognized personality within the EMEA user groups and has a worldwide network. I would like to wish her and her team all the best for the future. At the same time, I would like to thank everyone who has supported me in the past in creating the magazines.

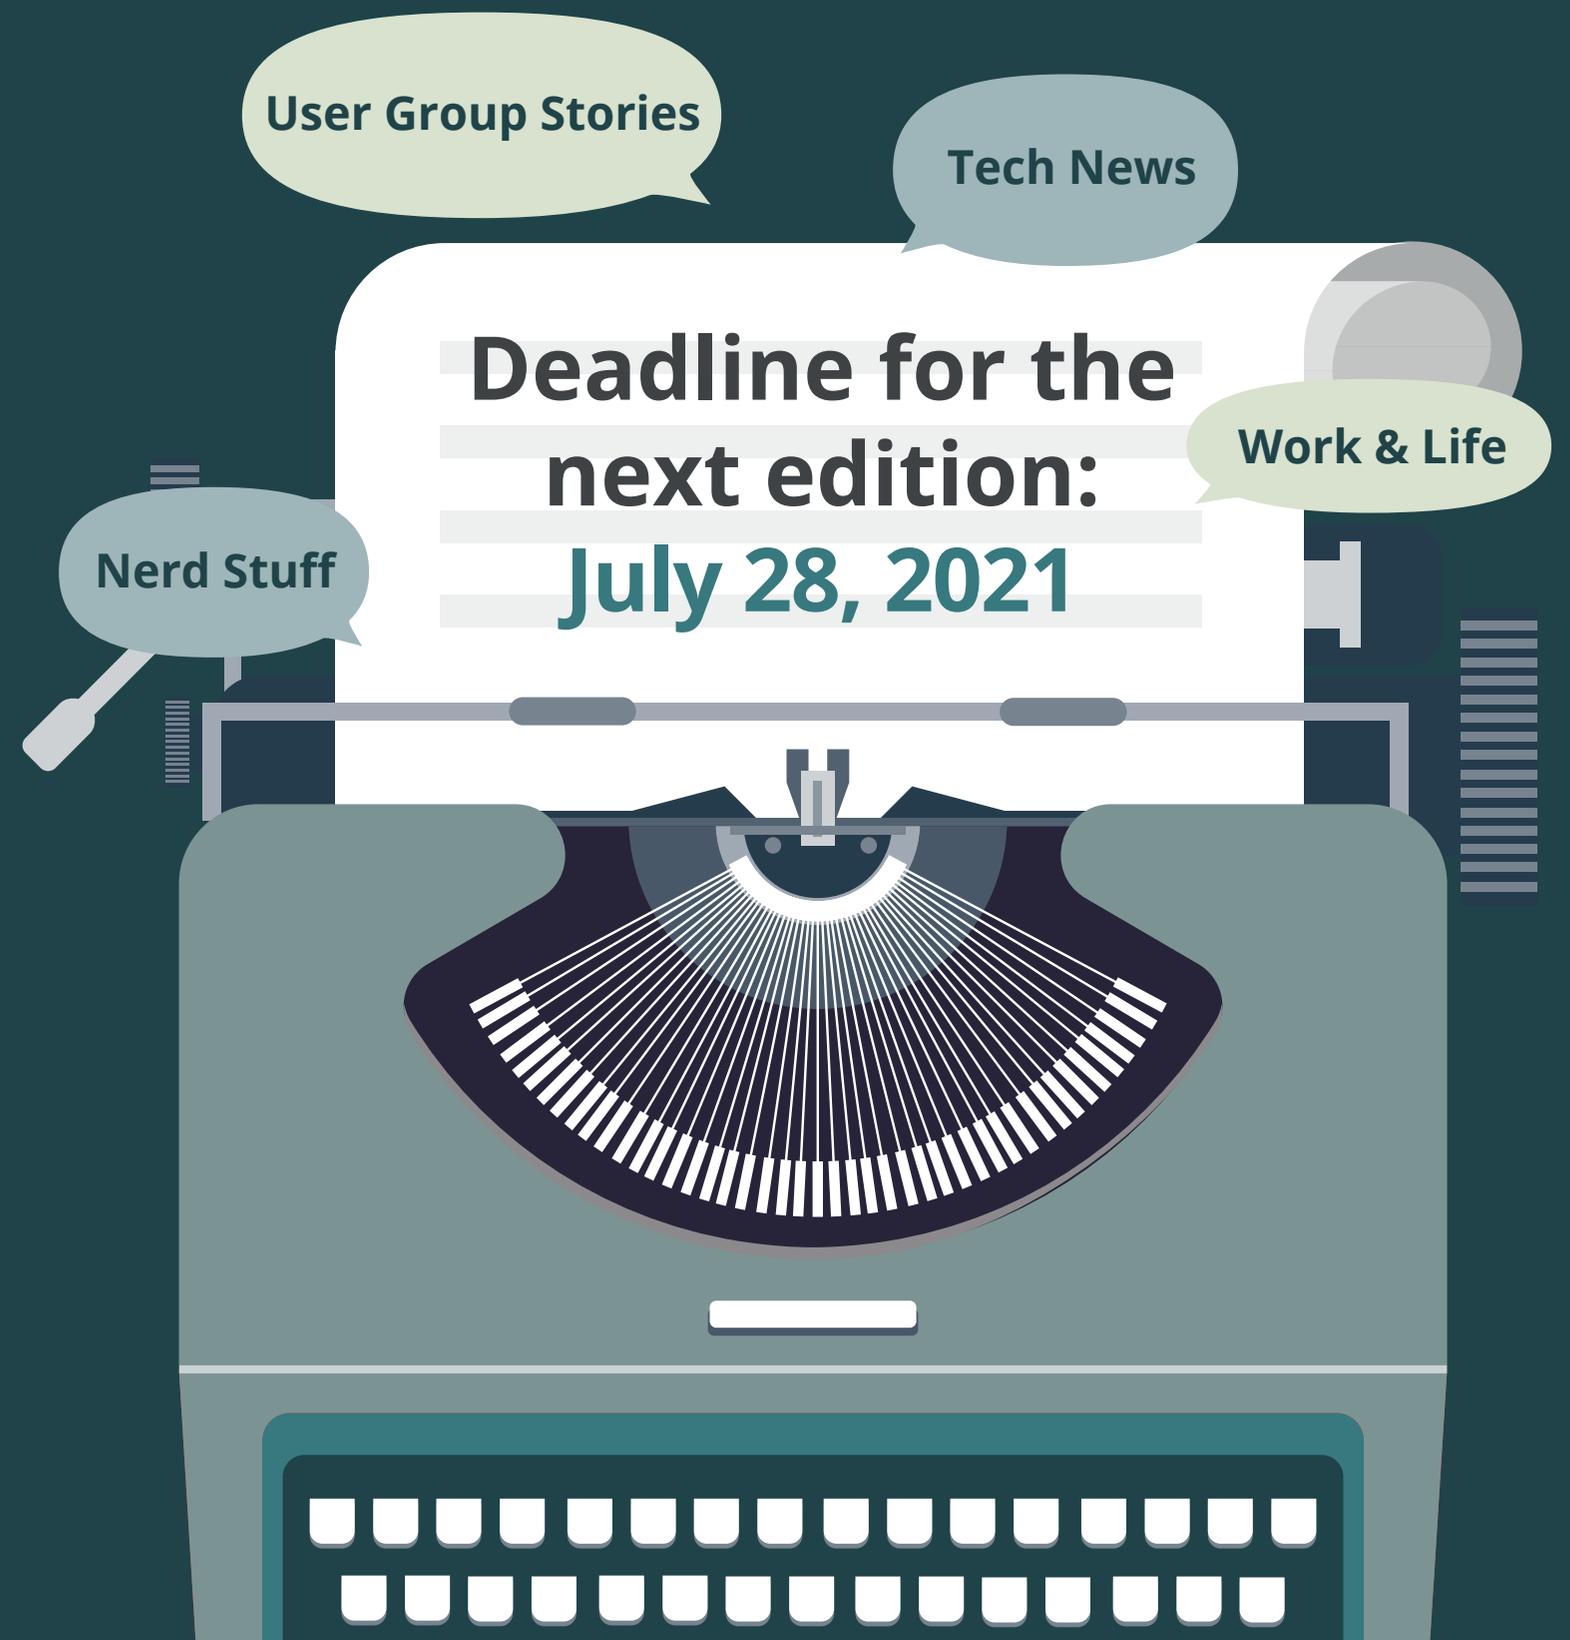
Good luck  
Dietmar Neugebauer



# Submit Your Article!

You have an interesting topic to publish in ORAWORLD Magazine? Then submit your article and be part of it!

Please e-mail us your article via the online form at [www.oraworld.org](http://www.oraworld.org).



# „The Secret of a Successful Code Review“

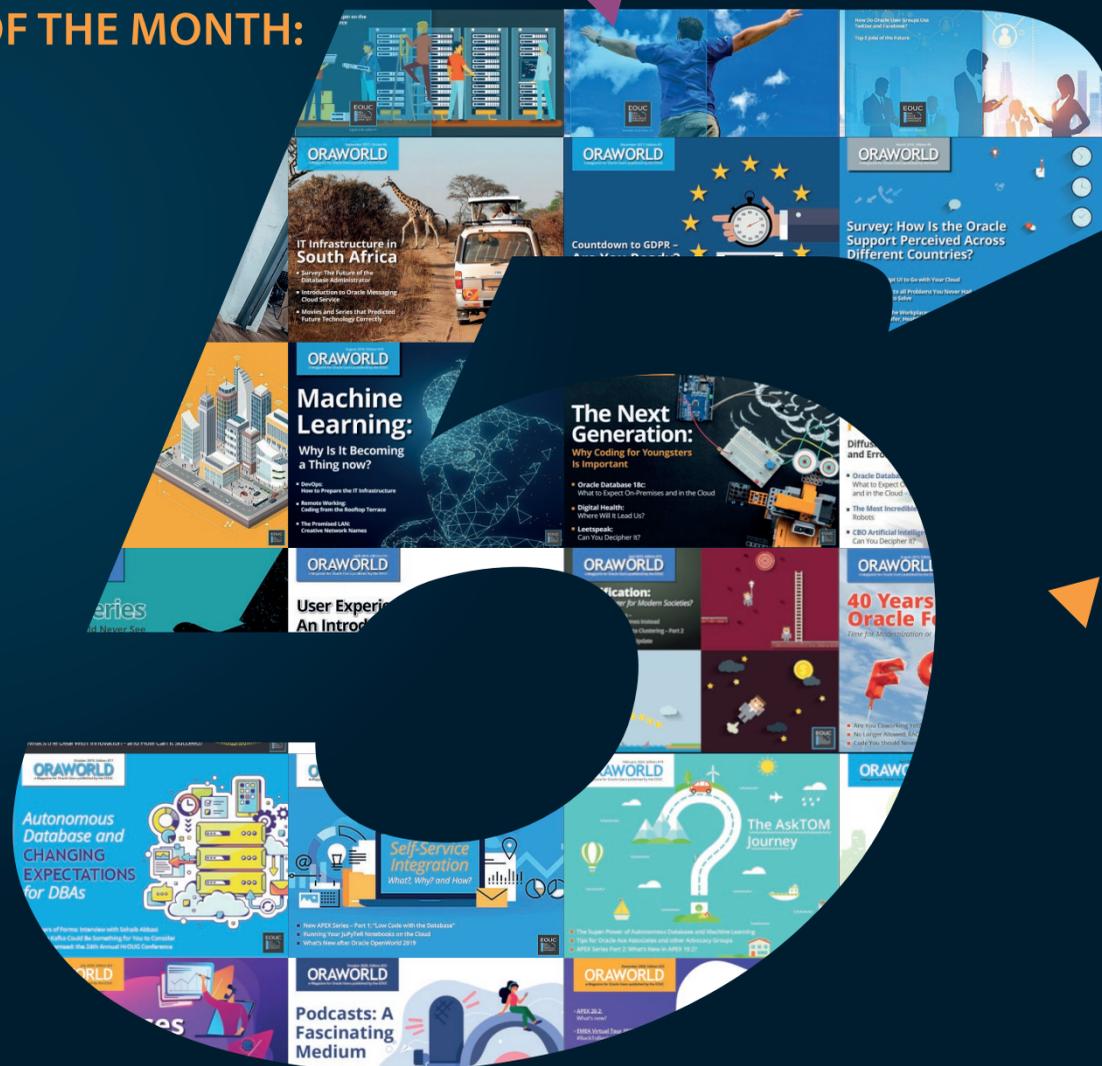
www.commitstrip.com



CommitStrip.com

CommitStrip is a daily strip recounting funny anecdotes of life as a coder mixed up with a dash of topical tech news. Find more comics here: [www.commitstrip.com](http://www.commitstrip.com)

NUMBER OF THE MONTH:



— YEARS OF ORAWORLD —

half a decade ago now, in May 2016, the first issue of ORAWORLD magazine was published. Before, there was the EOUC Newsletter, which eventually became ORAWORLD five years ago. Since then, our eMagazine has regularly provided readers with exciting stories and articles from the Oracle cosmos and beyond.

The sections "Tech & Nerds" and "Users & Groups", which still exist today, were already introduced with issue #1. With the third issue, the section "Work & Life" was added, in which we look beyond the technical horizon and personal stories are given space – as in our new interview series "What's Your Passion?", in which well-known community representatives talk about their hobbies and passions.

Celebrating 5 years and 25 issues, we went back and rummaged through the **first ORAWORLD**: Our premiere issue was dedicated to the topic "Women in IT" and took a critical look at the low percentage of women within the Oracle User Groups – also on the basis of a specially conducted survey.

In addition, the authors dealt with the status quo of Oracle Forms and explored the question of why the Twitter hashtag #letswreckthistgether was becoming increasingly popular within the APEX community. Furthermore, ITOUG, the then still very young Italian User Group founded in 2014, introduced itself.

*Interview by Christian Luda*

# What's Your Passion?

## Janny Ekelson: “For me Running was Crucial in Surviving the Lockdown”



*In the previous ORAWORLD issue our interview with Mike Dietrich kicked off our new series “What’s Your Passion?” with known people from the Oracle community telling us about their passions outside of work. This time, ORAWORLD gets sporty with Janny Ekelson, president of the Oracle Benelux User Group (OBUG), who talked to us about his new-found love of running and other passions.*



**Janny, you're a passionate runner.  
When did this passion start for you?**

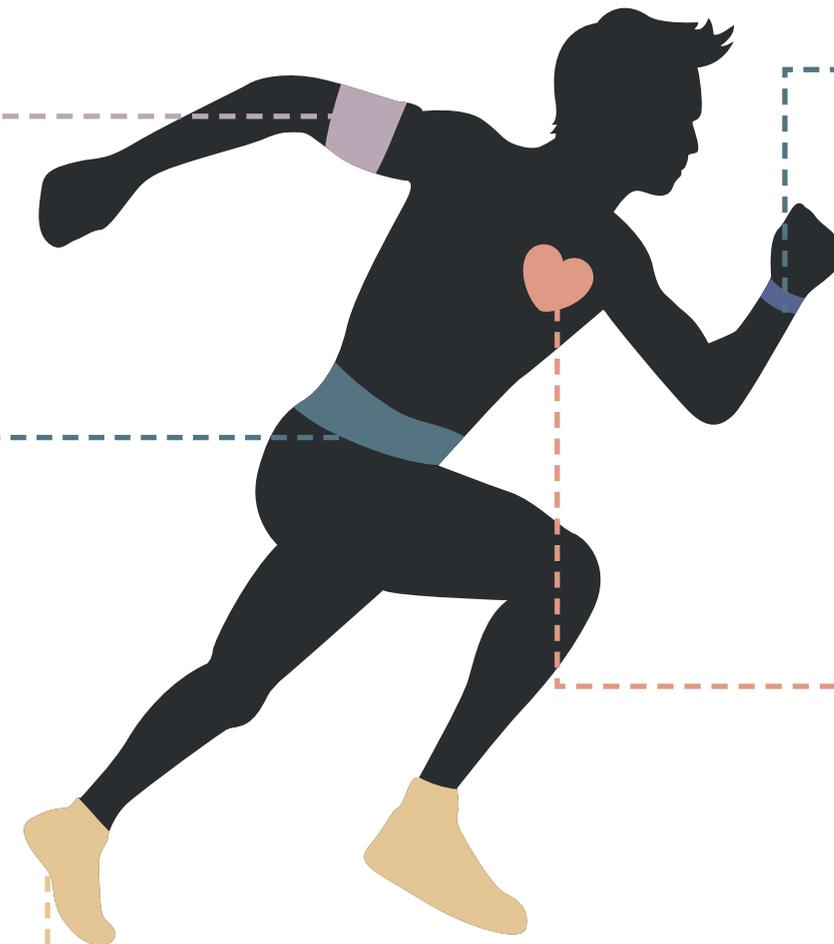
Around seven years ago. Before that, I was a squash player. When my regular squash partners had to stop because of injuries I needed a new sport. And I said to myself: "Well, the easiest one is running, you don't need anybody, and you can do it almost anywhere in the world. All you need is a good pair of shoes and time to do it." So, I gave it a try. In Belgium, we have the program "Start to run" where people motivate others to start running. It helps you building up your stamina depending on your experience. For example: run a kilometer, then walk a few minutes, then run another kilometer and so on. But eventually I decided to build my own program because when I started, my watch told me to stop running and walk while I could still turn around and see my house, and I thought: "This is not really motivating." So, I adopted a schedule based on how I felt and ran the first time for an hour with regular intervals of breaks – running five to ten minutes, then I walked a bit and when I felt relaxed, I picked up again. A bit like the official "Start to run" program but at my own feeling. So that was my first experience and everybody at home called me crazy because before that I had always stated "I hate running".

**Did you set goals for yourself at the start –  
like reaching certain kilometers within a certain time?**

Not in the beginning. But then I had a good friend who lived in Belgium at that time and liked running too. She motivated me to go for longer distances and we ran together sometimes. It was motivating to have someone stimulate you. And after a couple of months, I had the goal to do a first race. I wanted to see where I was at and what I could do – not to win anything but just to get the feeling of being part of a group. When you're at the start of a race, you feel that adrenaline. I looked around at the other people and said to myself: "I wonder who I can beat here". The competition makes you a bit more eager to go faster.

**Do you remember how many kilometers your first race had?**

Nothing wild. Around six kilometers. I had already done more than that so I knew I could do the distance. It was more the competitive part for me. And the next time, I went for 12 kilometers which was already a bit more challenging. After that, I started training regularly. Before it was once or twice a week depending on time and mood. But then I had the goal to do a half-marathon. And if you want to do a longer distance you should have a regular rhythm in your body to make sure you don't hurt yourself. Frequency is quite important and to build up gradually and not force it. So, I finally registered for the Half Marathon of Brussels. Before, I had only done local races with a few hundred people and now it was a few thousand people. It's a totally different atmosphere. Everywhere along the route



there's people making music and cheering for you – even if they don't know you. It was very stimulating. I did the same race three times and really enjoyed it. Then, the ultimate goal was the marathon which is very different. So, I visited my doctor, my practitioner for advice. And I was in luck: my doctor has a group practice and one of the doctors is partner to a semi-professional marathon-runner. She advised me how to build up my rhythm. Let's say if on average you run 50 kilometers a week and you want to do a marathon you need to go for 60 to 70 kilometers a week, but every third week take it slower to let your body recover.

***Did you also change your nutrition in preparation of the marathon?***

Not really. Just the last week, I didn't drink alcohol because it was something generally recommended. I don't know if it would have made a difference. It probably has an effect on your body. But I just ran for my own fun and not to win any price. For me it was all about seeing that I could do it. To achieve it and finish it – that was the goal for me.

***And you did finish it. Do you remember your time?***

I wanted to do the 42 kilometers below 4 hours and ideally at 3:45 hours – and I did it in 3:40 hours. So, I was even below my best estimated time which was very satisfying. And I didn't really suffer afterwards. Of course, I felt in my muscles that I'd done a huge effort, but nothing serious. I then didn't run for two or three days and picked up again slowly with short distances to not stress my body too much. But yeah, it was quite fun, but also long. At kilometer 30 I was thinking to myself: "What the hell? Why am I doing this?" (laughs) And then at 40: "Please, show me the last kilometer sign, please, I'm



tired". The last kilometers are tough but then again there's a lot of public which helps. It's important to keep the same pace and rhythm. When you leave at the start with 5,000 people, everybody is rushing forward and going much too fast, and of course you don't want to be over-run. But that's the danger, then you suffer at the end. Which really helped were the pacers: professional or semi-professional runners that have either a flag or a balloon with a specific color for different running times. So, there is one pacer at 3:30 hours and another one at 4 hours, and if you run in between those two, you know that you're at the right pace and not over-forcing yourself.

***How much time was between the point when you decided to start running and the marathon?***

Probably five years. Until then, I hadn't done more than a half-marathon, maybe 25 kilometers. It wasn't that I couldn't do it but I wasn't feeling mentally ready. When I finally did it, I registered just the week before the race when I really felt comfortable doing it. It was a last-minute decision. Until now, that also was the only time for me running that distance. But I run a lot of half-marathons, every weekend about 20 to 25 kilometers.

***Have you done half-marathons outside of Belgium?***

I've run everywhere in the world. Until Covid, I travelled a lot for work and have run everywhere where I've been, for example at the DOAG Conference in Nuremberg. But I've only done one race abroad, the Star Wars Half Marathon in Orlando. At 5 am in the morning – which was good because in Orlando it's very humid and hot. It was a huge group and the only time I won a price. I got a plaque for being among the





top 5 in my age group. While abroad, I always have my shoes with me. I think I've run on every continent except for Oceania and Antarctica. I like the challenge of running at new places. Unfortunately, right now with the lockdown, there is less option for variation and for me variation is very stimulating.

***Do you have some tips and for someone that hasn't run before and wants to start?***

First of all, get decent shoes. At the beginning, I just took the shoes I had at home. It's like playing guitar but not having a decent guitar. You can do it, but it's maybe not the best thing. When I noticed after a few weeks that I kind of enjoyed running, I went to a sports shop and got some decent shoes. When they saw what I was wearing they said: "Don't tell me you're running with that!" (laughs). It's a different feeling if you have good shoes on your feet and it hurts your knees less being softer on the ground. And then, running is a mental thing. Everybody without physical limitations can do it. You need to find out what you like in running. Some people like long distances, other people like short distances but fast. That's something you have to find out. And then at the beginning take it slow and short. Don't overdo it. Don't force it. Otherwise, you might end up injuring yourself or not liking it because you're suffering afterwards. All the known sports brands have training programs that guide you how to start and build up depending on your goals. Also, if you like company, you could find a partner to run with. And if you and your partner can still speak while running, you know that you're at the right pace.

***If you decided to run another marathon, would you have a favorite city?***

Maybe Madrid, and also New York City or London, because of the nice scenery. But then again, I might even prefer a marathon outside of a city as I'm switching more and more to off-road running. The only drawback

is the weather sometimes – you not only have to shower yourself but also your shoes (laughs).

***What are doing when you're not working or running?***

I love travelling because I really enjoy discovering new cultures, new people. I really enjoy South America. I've been on vacation once in Peru, and for business in Mexico and Chile a few times. And I like languages. I learned Spanish just by interest and later on, I could use it on my travels. Both passions, traveling and languages, stimulated the other. I also have a couple of friends in South America that I do regular calls or video calls with. That's replacing the absence of travelling at the moment a bit.

***How many languages do you speak?***

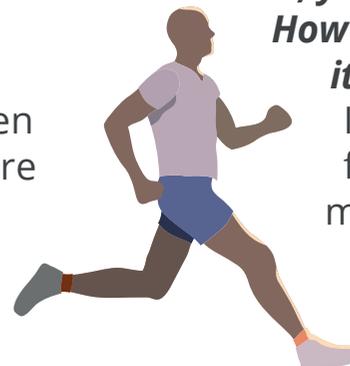
Four and a half. Dutch is my native language, then English, then French because it's the second language in Belgium, Spanish out of interest, and the half part is German. It's basic, I can book a hotel and get a dinner, but I can't do a business call in German (laughs).

***Is there another language you would like to learn?***

Yes, I've been talking about it recently. I'd like to learn a Scandinavian language, like Swedish. It's such a nice language. I'm also tempted to say Italian but it's too close to French and Spanish. I'm looking for something a bit different from what I know. I have been to Stockholm and caught some words and I think it would be an interesting challenge to learn Swedish but I'm not sure if I'll find the time.

***So, you have a lot of passions and are also very busy at work. How do you find the time and balance? And: how important is it to have passions in your free time?***

It's very important to have one or more passions in your free time. It gives you something to look forward to. For me running was crucial in surviving the lockdown. It helps



you relax. My schedule is fully booked from early morning until late in the day. So, in my calendar, I always book time for running each day within the next two weeks. One hour for running plus 15 minutes for a shower, just to make sure it's in there. Often, I have to shift it because you can't say no to some things in work relations. But if you don't do it, your day is booked. It's important to block time for your passion as it's easy to drop your own things for other professional reasons.

***When the pandemic is over, do you have one special travelling goal on your bucket list?***

That's a long list but number one is Argentina – I want to go all the way down to the end, to the Antarctica. That was planned for last year, but didn't happen. It might take another year, but that's the next one. I'll get there. That's for sure.

***Janny, I hope, you'll be there soon. Thanks a lot for taking your time.***



## About Janny Ekelson

**Janny** is IT Manager at the FedEx Express International, he is responsible for planning, organizing and implementing the IT applications for FedEx across the region, more specifically the IT projects for Revenue Operations and Ship-to-Collect. Main focus is the digital transformation of the finance application domain towards eInvoicing, online payments and increasing Customer Experience. Janny joined FedEx in 1998 as a project leader within the HR team and was in charge of the EMEA wide roll-out of PeopleSoft HRMS project. In 2001, he became the IT manager of the Applications and Architecture team, which is his current role.

Janny has a University degree in Commercial Engineer at the KULeuven, the University of Leuven and he completed an additional year for Business Intelligence at EHSAL Brussels in 2004.

Janny is married since almost 30 years and has twins of almost 25 years.

*Jim Czuprynski*

# ***FAST & FURIOUS:***

**CAPTURING EDGE COMPUTING DATA**

**WITH ORACLE 19C FAST INGEST**

**The Internet of Things (IoT) has penetrated every corner of our modern world, from the smart meters that monitor our home's electricity usage to the smart home technology that lets us turn on lights or watch our favorite shows with a simple voice command. All these devices generate enormous quantities of data, and capturing that information for analysis while it's still streaming means we need a database that can handle transaction volumes of immense density and diversity. In the first part of this article series, Oracle ACE Director Jim Czuprynski shows how Oracle Database 19c's Fast Ingest features are a perfect match for capturing streaming data so it can be retained and analyzed in real time.**

## IoT, Everywhere You Look

The Internet of Things (IoT) is already making an enormous penetration into our daily lives, well beyond our daily involvement with internet-enabled devices like video doorbells and voice-activated portals, and this expansion shows no signs of abating in the future. For example, we just recently permitted our local power utility to install a smart meter to monitor our daily usage of electrical power. In return for improving our off-peak power

usage, we're given a moderate discount on the cost of electricity. Our utility is also promoting the installation of photovoltaic (PV) solar power panels for homes in our area that will allow us to feed electricity created via alternative energy back into the power grid.

So, let's assume that you're the DBA for a forward-looking public utility that recently implemented an aggressive approach to capturing energy from solar power sources. Your IT team therefore needs to capture the huge quantity of data from smart meter technology monitoring all those solar panels installed across our customer base spanning small at-home businesses as well as large industrial facilities. Of course, your IT team will need to answer all kinds of complex business questions using these data, including which types of customers are benefiting the most from the utility's solar power investment, and I'll cover how to tackle those analyses in the next article in this series. But for now, the main thing you need to accomplish as DBA is to get these data captured into your Oracle database.

*Figure 1* shows the basic entity-relationship diagram for the data we'll be exploring during this article series; the code to construct this schema and its related tables is available in this [Github repository link](#).

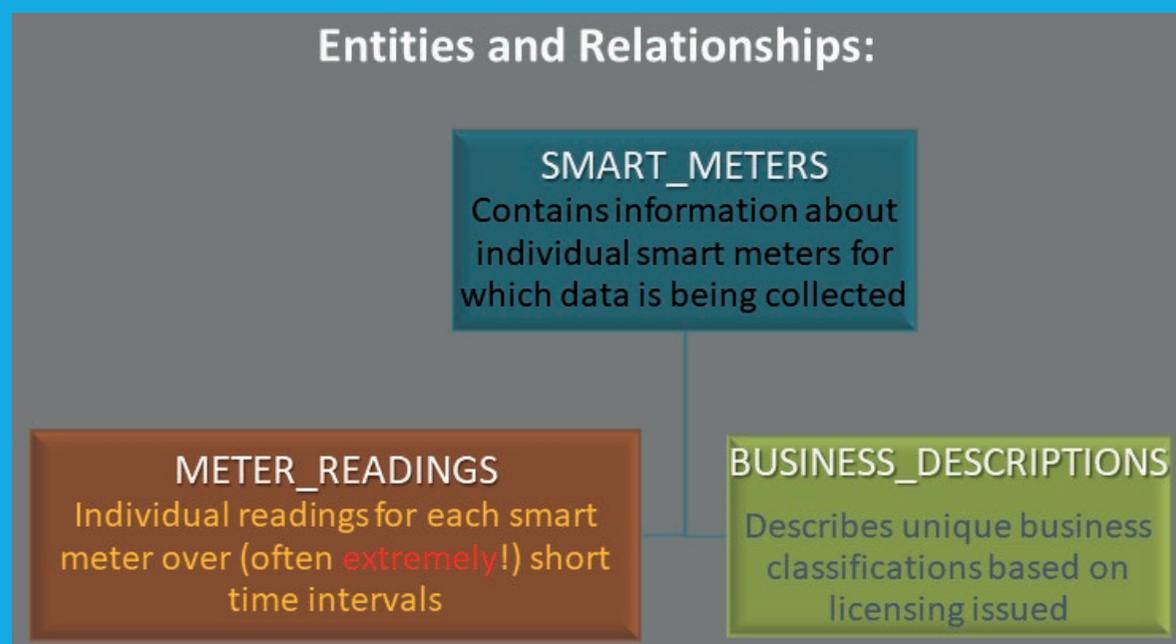


Figure 1: ERD for Smart Meters entities

## Fast Ingest: The Big Picture

Let's take a look at the elements of Fast Ingest from the perspective of what happens within an Oracle database instance. I've outlined its key elements in *Figure 2* below.

Here's the basics of how the Fast Ingest transaction strategy differs from normal Oracle database transaction processing:

- Before we can enable a table for Fast Ingest, a special subarea of the database instance's memory structure – the

*Memoptimized Row Store (MORS)* – must be initialized within the instance's Large Pool. The size of the MORS is set by the **MEMOPTIMIZE\_POOL\_SIZE** initialization parameter, and as I'll demonstrate shortly, setting this parameter to a non-zero value does require restarting the database instance to initialize the MORS.

- Once the MORS is initialized and the desired table(s) have been enabled for Fast Ingest through the **ALTER TABLE ... MEMOPTIMIZE FOR WRITE** command, data that has been inserted via an **INSERT** statement containing the

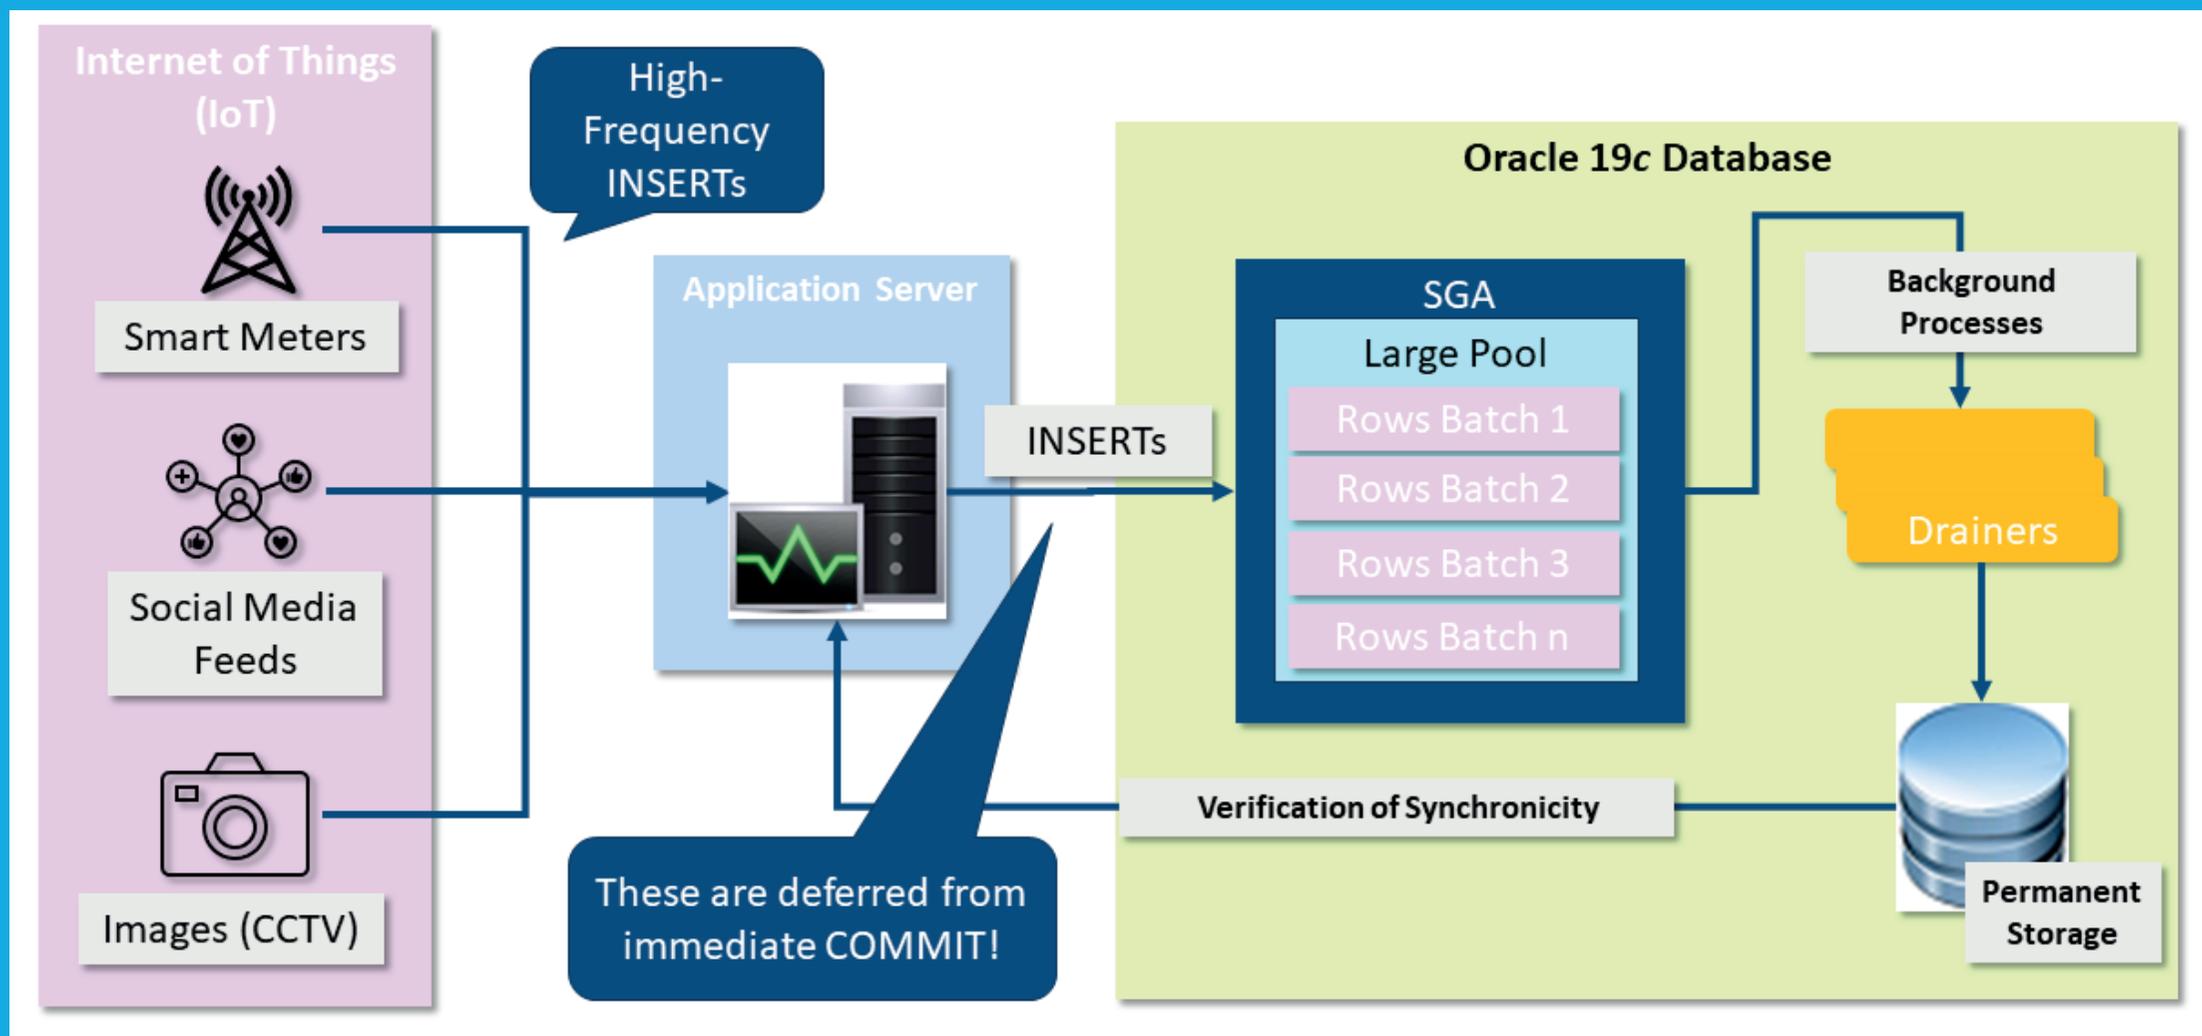


Figure 2: Fast Ingest: The big picture

**+MEMOPTIMIZE\_WRITE** hint will automatically be routed to the MORS for Fast Ingest instead of normal transaction processing.

- The insertion of these batched rows into their corresponding tables is handled by a series of “worker” background processes (**W00n**). Because normal transaction processing is not used, potentially huge volumes of data can be inserted efficiently *en masse* instead of being processed on a row-by-row basis.
- To further increase insertion speed, the worker processes add rows to corresponding indexes and perform constraint checking only when the data is written from the MORS to database storage.
- And here’s the reason Fast Ingest is so much faster than normal transaction processing: *No redo log entries are written to the transaction log, and neither are any UNDO segments used for transaction processing.* In fact, **COMMIT** and **ROLLBACK** statements are simply *ignored* for tables using Fast Ingest.

## Wait. What Was That About COMMIT and ROLLBACK?

If you’re a reasonably experienced DBA, your reaction to that last bullet point probably echoed mine: *Since our job is to never lose any data, how can we possibly guarantee against the potential loss of transactions?* Here are a few other considerations along that tack:

- Fortunately, the database instance does keep track of both a *low water mark* (LWM) for all rows that have been inserted into the database already at an instance level, as well as an individual session-level *high water mark* (HWM) reflecting the last row inserted for each session. When the database’s LWM exceeds the session’s HWM, it’s then safe for the application session to discard any record of the inserted rows. (I will demonstrate how to check these watermarks to verify successful Fast Ingest processing a bit later in this article.)
- Constraint checking for primary key violations happens while the **W00n** processes are writing data to disk, so it’s also the

application’s responsibility to verify that any potential PK violations are trapped *before* rows are batched for Fast Ingest processing.

- Also, since there’s no **UNDO** blocks from which to read data, it’s not possible to query data from the MORS until the **W00n** processes have already flushed batched rows to storage.
- Finally, should the database instance crash unexpectedly before all Fast Ingest data has been written to the database, *then it is indeed possible to lose data.*

When these limitations seem frightening, I find comfort in remembering why we’d be interested in using Fast Ingest in the first place. In the use case I’ve described above, I’ve been tasked with capturing *extremely* large numbers of individual rows generated from potentially *tens of thousands* of disparate sources in *extremely tight time frames*, and likely to be used for *immediate* analysis of power utilization trends.

So, if I am capturing data from 15,000 Smart Meters that are each generating a few thousand data points every minute, would a transaction loss ratio of less than, say, 0.01% be acceptable, as long as I realize the resulting accuracy? My answer in this scenario is likely to be a qualified yes, as long as my business unit customers understand the limitations. On the other hand, I would absolutely *reject* using Fast Ingest to capture complex stock trading transactions or to determine the precise position of an aircraft for an air traffic control system.

## Exploring Fast Ingest Features: A Sample Implementation

To explore Fast Ingest’s features, I searched diligently for a streaming source whose payload I could ramp up so that it would be sufficiently intense to demonstrate its potential performance gains. While I did locate a few promising sources, I was disappointed at their inability to scale beyond



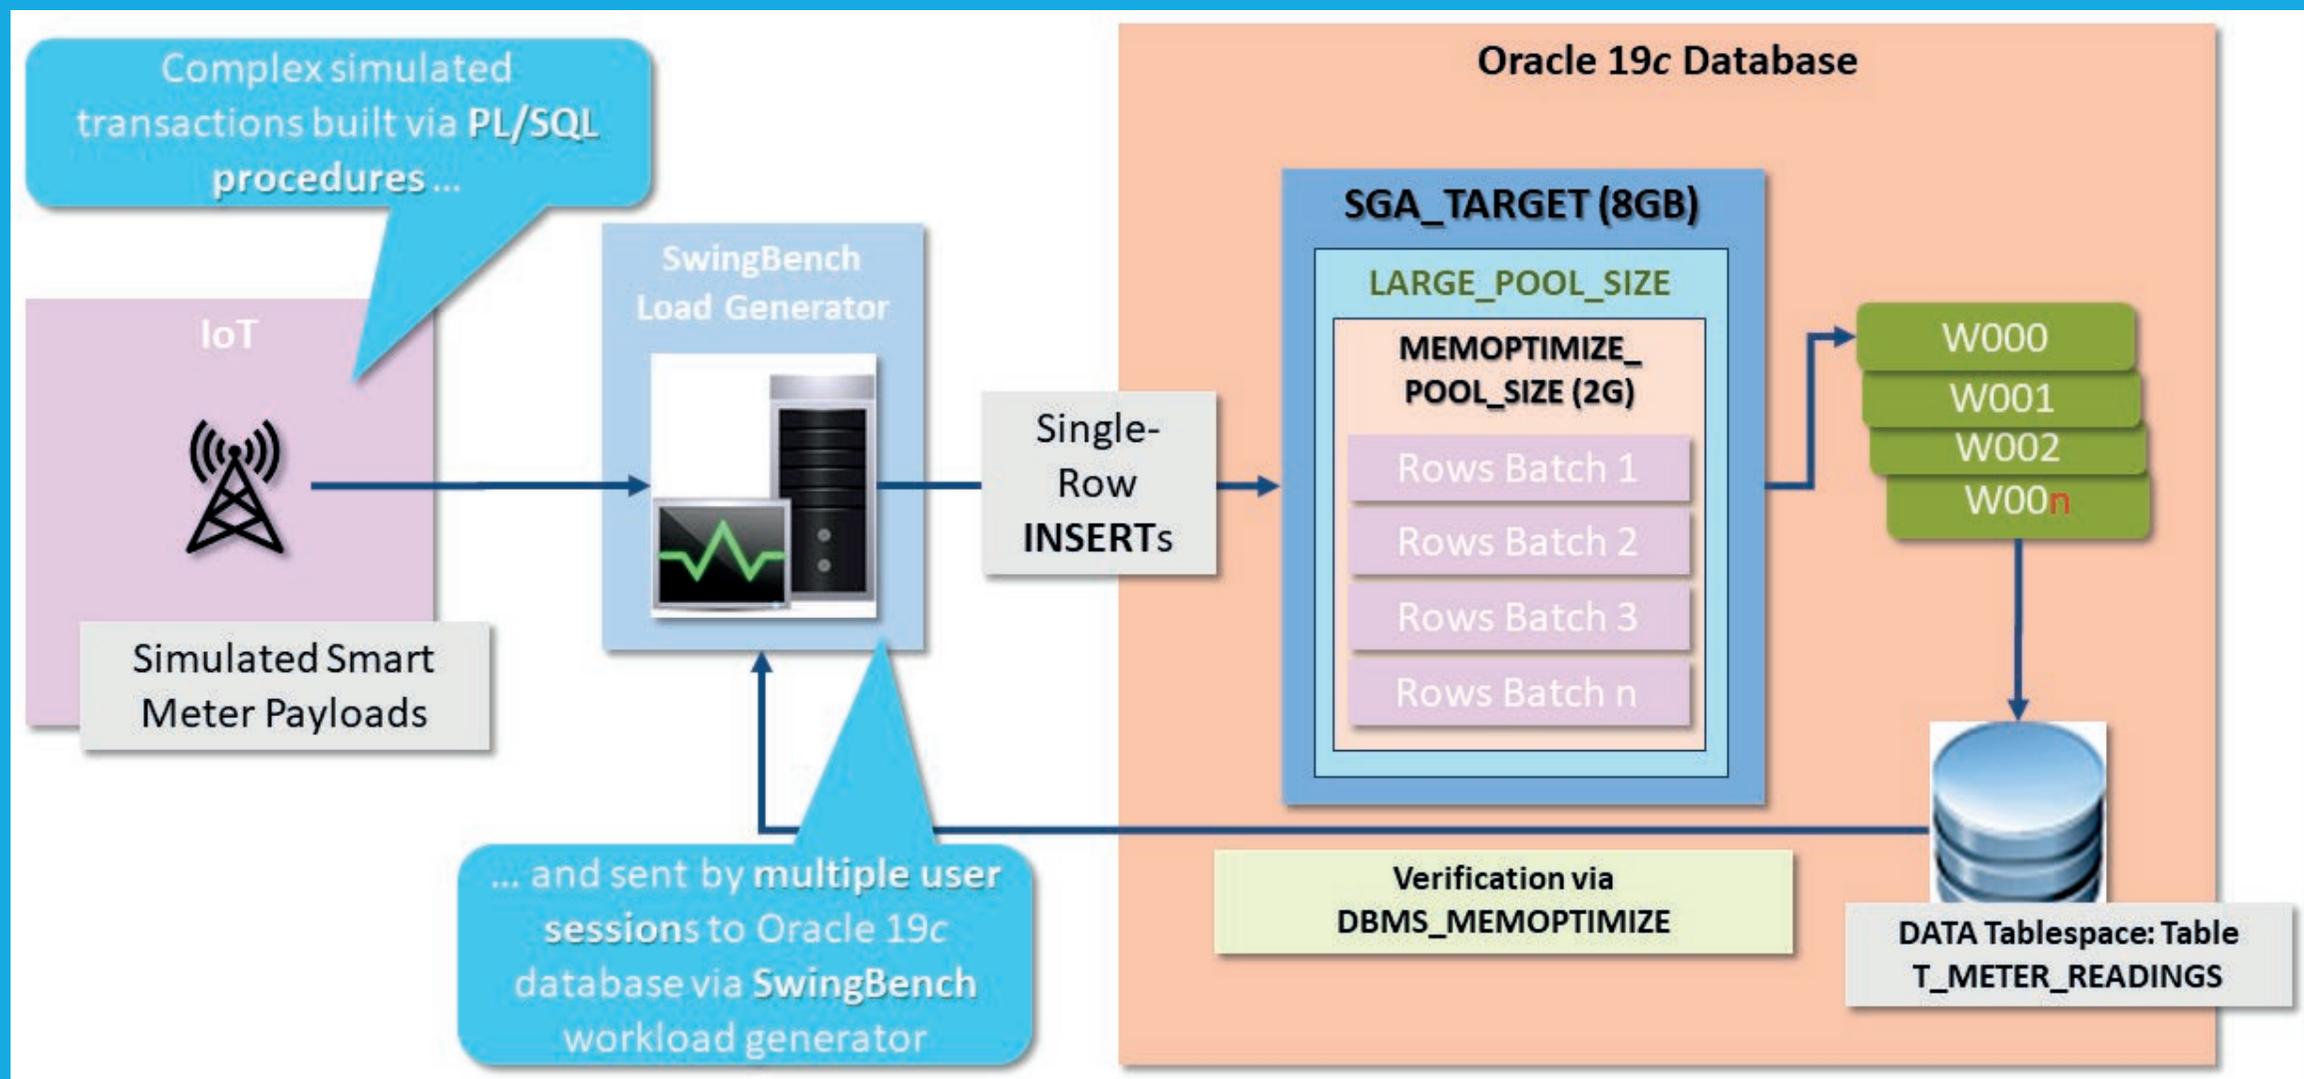


Figure 3. Demonstrating Fast Ingest features using SwingBench

the intensity my experimentation required, so I decided to build a simulated source that would supply a deluge of smart meter readings across multiple user sessions. I've laid out the elements of this experimental implementation in *Figure 3*.

- For my power utility's simulated customer base, I selected a subset of some 50,000 actual business licensees from the City of Chicago's public database. Since I'm planning to revisit these data in future articles, I wanted data that included real-life physical locations so I could demonstrate geographic information system (GIS) capabilities, and the City graciously

includes longitude and latitude positions for each licensee. Data from each individual licensee and their common business categorizations were then edited, recomposited, and loaded into the **T\_SMARTMETERS** and **T\_BUSINESS\_DESCRIPTIONS** tables, respectively.

- I built a package named **STREAM\_SIMULATOR** that generated and loaded several hundred thousand rows for the two historical partitions of **T\_METER\_READINGS** - **OLDEST** and **SINGLEDAY**. (The reasons for partitioning **T\_METER\_READINGS** will be made apparent as we explore another 19c feature called Fast Lookup in a future article

in this series.) See this [Github repository link](#) for that package's specification and body.

- Finally, I used the latest version of [SwingBench](#) and included the same logic in the **STREAM\_SIMULATOR** package to build two functions for random generation of Smart Meter readings. The first function inserted random rows directly into the **T\_METER\_READINGS** table *without* using the Fast Ingest hint, while the second function applied the **+MEMOPTIMIZE\_WRITE** hint to specifically request Fast Ingest processing. I was therefore able to invoke these functions from numerous database sessions and at a varying intensity based on time waited between transactions. This [Github repository link](#) shows the second function's code.

```
SQL> ALTER SYSTEM SET memoptimize_pool_size = 2G SCOPE=SPFILE;
System altered.
```

```
SQL> shutdown immediate;
```

```
Database closed.
Database dismounted.
ORACLE instance shut down.
```

```
SQL> startup;
```

```
ORACLE instance started.
Total System Global Area 1.4496E+10 bytes
Fixed Size                9702624 bytes
Variable Size             3892314112 bytes
Database Buffers          1.0570E+10 bytes
Redo Buffers              23851008 bytes
Database mounted.
Database opened.
```

```
SQL> show parameter pool_size
```

NAME	TYPE	VALUE
java_pool_size	big integer	0
large_pool_size	big integer	0
memoptimize_pool_size	big integer	2G
olap_page_pool_size	big integer	0
shared_pool_size	big integer	0
streams_pool_size	big integer	0

Listing 1: Activating the MORS before enabling tables for Fast Ingest

```
SQL> ALTER SESSION SET CONTAINER = pdb19;
```

```
Session altered.
```

```
ALTER TABLE simiot.t_meter_readings MEMOPTIMIZE FOR WRITE;
ERROR at line 1: ORA-62172: MEMOPTIMIZE FOR WRITE feature
cannot be enabled on table in encrypted tablespace.
```

Listing 2: Results of initial Fast Ingest attempts

## Preparing an Oracle 19c ATP Instance for Fast Ingest

There is just a single step to prepare an Oracle 19c database to use Fast Ingest features, and that's to reserve memory for the MORS within the instance's Large Pool. I had already loaded all necessary data in an OCI-resident Autonomous Transaction Processing (ATP) instance, so I issued the commands shown in *Listing 1* to size the MORS to 2GB and enable that instance for Fast Ingest. Note that I first needed to set that parameter within my database's **SPFILE** before bouncing the database instance.

MORS is now definitely established within the Large Pool, so now let's attempt to add the **T\_METER\_READINGS** table to the MORS by issuing the **ALTER TABLE ... MEMOPTIMIZE FOR WRITE** command for that table, as *Listing 2* shows.

And here's the first of a few disappointing issues about Fast Ingest's availability: Even though it seems to be perfect for an ATP environment, as of this writing it will not work with a 19c ATP instance because there are some strict limits currently imposed, and one of those limits is that the table's underlying tablespace cannot be encrypted. Since an Autonomous Database doesn't permit any unencrypted data to be stored within it, and since it's not permitted to create a new tablespace for an ADB instance (much less an unencrypted one!), I cannot demonstrate Fast Ingest features there – yet\*.

**\* Note:** I've talked with Oracle Product Managers responsible for Fast Ingest and other in-memory products, and this is a well-known bone of contention. I'll be sure to revisit this article after this limitation has been lifted from ADBs, hopefully in the immediate future.

## Back to the Drawing Board: Simulating an Exadata Environment for 21c on DBCS

Since I needed an unencrypted tablespace for my Fast Ingest examples to work, I shifted my experimentation to my Oracle 21c database environment. After I reloaded my SIMIOT schema's three tables with sufficient data, I again attempted to activate the **T\_METER\_READINGS** table for Fast Ingest processing and received yet another error message.

And that brings forth another limit for Fast Ingest, at least as of this writing: It is only supported for Oracle databases running within an Exadata environment. Fortunately, as I'm simply experimenting here, I can force my 21c database to simulate an Exadata environment by setting the **\_exadata\_feature\_on** initialization parameter to TRUE and then bouncing my database instance to ensure it takes effect, as shown in *Listing 3*.

**WARNING:** This technique should **only** be used for experimentation purposes to simulate an Exadata environment for evaluating features that specifically require Exadata. Even so, it does not always work perfectly in all situations. Most importantly, it should **NEVER** be used within a production environment!

```
SQL> ALTER SYSTEM SET "_exadata_feature_on" = true SCOPE=SPFILE;
SQL> SHUTDOWN IMMEDIATE;
SQL> STARTUP;
. . .
```

Listing 3: Activating simulated Exadata features for experimentation only

```
SQL> ALTER TABLE simiot.t_meter_readings MEMOPTIMIZE FOR WRITE;

ERROR at line 1: ORA-62144: MEMOPTIMIZE FOR WRITE feature not allowed
on table with foreign key constraint.
```

Listing 4: Errors returned when activating Fast Ingest for a table with FK constraints in place

I'll be using the 21c innovation release for this series so we can explore several new features, including the new Native JSON type, in future articles. However, everything I'll show you in this article will also work just fine in your Oracle 19c database, as long as it is running within an Exadata environment.

## Yet Another Unexpected Limitation of Fast Ingest

At long last, I'm ready to activate my target table for Fast Ingest ... and of course, I immediately encountered one last stumbling block, as shown in *Listing 4*.

Another unexpected error! As a responsible DBA who strives to build a good referential data model – by applying a foreign key constraint on all entries in the **T\_METER\_READINGS** table, thus ensuring referential integrity between those rows' Smart Meter ID and those found in its parent **T\_SMARTMETERS** table – I've inadvertently made it impossible to use Fast Ingest without first removing that constraint.

However, this limitation does make sense when we consider the primary benefit of Fast Ingest is *loading data at extreme speeds*. I've captured several other stringent limitations in *Table 1*, but it's by no means complete, and future releases of Oracle are likely to lift some of them. For the most up-to-date information, check



Can I Use Fast Ingest On a Table That ...	Allowed?
Uses <b>REFERENCE, INTERVAL, SYSTEM, or AUTOLIST</b> partitioning?	No
Employs a <b>Foreign Key constraint</b> against another table?	No
Employs <b>triggers</b> ?	No
Is stored within an <b>encrypted</b> tablespace?	No
Uses any <b>storage</b> compression?	No
Leverages <b>IN MEMORY</b> compression?	No
Is <b>EXTERNAL, Index-Organized, TEMPORARY, or NESTED</b> ?	No
Uses a <b>function-based, DOMAIN, BITMAP, or BITMAP JOIN</b> index?	No
Has a column with <b>default</b> values, or that's marked <b>INVISIBLE, VIRTUAL, OR UNUSED</b> ?	No

Table 1: Fast Ingest Table Restrictions

### Top SQL Comparison by Executions

- Ordered by 'Diff' column of '#Executions/sec (DB time)' descending
- '#Plans' column indicates the number of distinct execution plans for the statement in 1st and 2nd periods and in Both periods, combined
- Captured SQL Executions First: 5,384,412, Second: 7,346,324
- Captured SQL Elapsed Time First: 12,472.58 seconds, Second: 12,806.97 seconds
- Captured SQL Elapsed Time as a % of DB time First: 81.06%, Second: 74.32%
- Captured SQL Executions as a % of Total Executions First: 99.99%, Second: 100%
- Common SQL Executions as a % of Total Executions First: .04%, Second: .03%

SQL Id	#Executions/sec (DB time)			#Executions		SQL Text
	1st	2nd	Diff	1st	2nd	
2dnrbaqq0g38		426.19	426.19	7,344,006		INSERT /*+ MEMOPTIMIZE_WRITE *...
34v9c62y8crfv	349.84		-349.84	5,382,604		INSERT INTO T_METER_READINGS (...)
410Zuwrct8shh		0.03	0.03	465		BEGIN :1 := swingbench.storedp...
392zf3aa12zyp	0.02		-0.02	337		BEGIN :1 := swingbench.storedp...
3dbzmtf9ahvzt		0.01	0.01	127		merge /* KSM:OPTIM_DML_INF */...
4wcuxvpykak8n	0.02	0.03	0.01	337	465	SELECT SM_ID FROM SIMIOT.T_SMA...
27q1fj58cnz0k	0.00	0.00	0.00	28	84	select '1' from dual
dhc40pc5aj6p9	0.00	0.00	0.00	28	84	with vs as (select rownum rnum...
6n2qqv1brfhpp	0.01	0.00	0.00	82	69	select ind.type#, ind.property...
97jpmfvjkrck	0.01	0.01	0.00	189	190	SELECT SUBSTR(DECODE(:B2, NUL...
0kkhhb2w93cx0	0.00	0.00	0.00	56	44	update seg\$ set type#=:4, bloc...
5dqz0hqt9fru	0.00	0.00	0.00	54	50	select /*+ connect_by_filterin...
f7xvnp44nukp7	0.01	0.00	0.00	86	86	insert into wrm\$_snapshot_deta...
fdsrx9wqyqtmu	0.00	0.00	0.00	72	72	INSERT INTO wri\$_adv_parameter...
3bcxdbgju78wn	0.00	0.00	0.00	50	50	SELECT U.NAME, NULL, O.NAME, O...

Figure 4: AWR Top SQL Comparison by Execution

out My Oracle Support Note **#2605883.1, Limitations for Using Memoptimized Fast Ingest**, for a complete list of limitations for tables and related objects when implementing Fast Ingest.

### Generating an Intense, Multi-Session Fast Ingest Workload

Now that my 21c instance is ready to accept my simulated workload, I'll initiate 25 user sessions to randomly insert rows into the **T\_METER\_READINGS** table at an extreme rate so that I can demonstrate how to monitor and manage the MORS.

### Fast Ingest Performance Comparisons

Ultimately, of course, what really matters here is how much better Fast Ingest handles a large INSERT-only payload. To determine that, I ran two identical sets of performance tests, both consisting of 25 user sessions hammering away at the **T\_METER\_READINGS** table for 10 minutes, but with one small difference between the tests: enabling the **+MEMOPTIMIZE\_FOR\_WRITE** hint in the INSERT statement. I captured AWR snapshots before and after each test run and then ran an AWR difference report between the two test runs.

The AWR *Top SQL Comparison by Execution* report (**Figure 4**) shows the difference between the number of times the Fast Ingest-enabled INSERT statement versus its non-enabled INSERT statement was executed. Since the workloads were run for an identical length of time, it's clear that enabling Fast Ingest resulted in **7.344M** inserts versus **5.382M** without it enabled – roughly a **37%** performance improvement.

Likewise, the AWR *Top Segments Comparison by Physical Writes* report (**Figure 5**) shows some interesting impacts on the number of physical writes performed during identical execution periods. The Fast Ingest-enabled test showed a lower percentage of physical writes (**12,051**) versus the non-Fast Ingest workload (**83,666**) for the **NEWEST** partition of the target table, even though more physical writes were actually performed (**31,793** vs. **45,190**) because of the corresponding increased number of INSERTs actually executed.

## What's Coming Up Next?

This article demonstrated how to:

- Activate the Memoptimized Row Store (MORS) to enable Fast Ingest features
- Use Fast Ingest to efficiently insert huge volumes of data from numerous sessions
- Measure the performance benefits of Fast Ingest versus traditional INSERT transaction processing

In my next article in this series, I'll demonstrate how to:

- Monitor and manage the MORS using **DBMS\_MEMOPTIMIZE** functions and procedures
- Use Fast Ingest's close cousin, *Fast Lookup*, to take advantage of the MORS for repetitive reads of large "chunks" of data
- Use Oracle's latest Machine Learning algorithms to quickly analyze data from streaming sources for accurate business decision-making

## Top Segments Comparison by Physical Writes

- Ordered by absolute value of 'Diff' column of '% of Total Physical Writes'
- Top 5 Segments from each period are compared
- Total Physical Writes First: 38, Second: 375
- Captured Physical Writes First: 193,085, Second: 293,495
- Captured Physical Writes as % of Total First: 508118.42%, Second: 78265.33%
- When \*\* MISSING \*\* occurs, some of the object attributes may not be available

Owner	Tablespace	Object Name	Subobject Name	Type	Obj#	Dataobj#	% of Total Physical Writes					Physical Writes		
							1st	1st Total	2nd	2nd Total	Diff	1st	2nd	%Diff
SIMIOT	DATA	METER_READINGS_PK_IDX		INDEX	80267	80380	419597.37	419,597.37	65860.00	65,860.00	-353737.37	159,447	246,975	54.89
SIMIOT	DATA	T_METER_READINGS	NEWEST	TABLE PARTITION	80242	80379	83665.79	586,928.95	0.00	65,860.00	-83665.79	31,793	0	-100.00
SIMIOT	DATA	T_METER_READINGS	NEWEST	TABLE PARTITION	80242	80379	83665.79	586,928.95	0.00	65,860.00	-83665.79	31,793	0	-100.00
SIMIOT	DATA	T_METER_READINGS	NEWEST	TABLE PARTITION	80242	80379	83665.79	670,594.74	12050.67	77,910.67	-71615.12	31,793	45,190	42.14
SIMIOT	DATA	T_METER_READINGS	NEWEST	TABLE PARTITION	80242	80301	0.00	670,594.74	12050.67	89,961.33	12050.67	0	45,190	100.00

Figure 5: AWR Top Segments Comparison by Physical Writes



## References

- Architecture of the Memoptimized Rowstore:  
<https://docs.oracle.com/en/database/oracle/oracle-database/19/tgdba/tuning-system-global-area.html#GUID-9752E93D-55A7-4584-B09B-9623B33B5CCF>
- City of Chicago Reference Information:  
<https://data.cityofchicago.org/Community-Economic-Development/Business-Licenses-Current-Active/uupf-x98q/data>

## Some Topical Articles On IoT:

- IoT and the Potential to Save the World  
<https://www.iotforall.com/iot-potential-to-save-the-world>
- What is IIoT? Industrial Internet of Things Explained  
<https://www.hologram.io/blog/what-is-iiot>
- How Smart Devices with Edge Analytics are Helping Business  
<https://www.iotforall.com/how-smart-devices-with-edge-analytics-are-helping-business>
- How 100% Clean Energy Could Power Our Cities And Towns  
<https://grist.org/energy/how-100-clean-energy-could-power-our-cities-and-towns>
- Miami pilots e-cargo bikes to reduce congestion, pollution  
<https://www.smartcitiesdive.com/news/miami-e-cargo-bike-pilot-dhl-city-congestion-pollution/578115/>

*Note: This article has also been published in ODTUG's Database Community portal, TechCeleration, in April 2021.*



## About Jim Czuprynski

**Jim** has over four decades of professional experience in information technology, serving diverse roles at several Fortune 1000 companies before becoming an Oracle DBA in 2001. He was named an Oracle ACE Director in 2014 and is a sought-after public speaker on Oracle Database technology features, presenting at various conferences around the world. He currently serves on the ODTUG Board of Directors.

Jim has authored well over 100 articles focused on facets of Oracle Database administration since 2003 and also co-authored four books on Oracle database technology. Jim's blog, *Generally...It Depends* (<http://jimczuprynski.wordpress.com>), contains regular observations on all things Oracle and the state of the IT industry. He also hosts a biweekly podcast, **Beyond Tech Skills**, with Liron Amitzi that focuses on everything about tech – except tech.

Phil Wilkins

# With Great APIs Comes Great Responsibility ... Making your APIs Safe



*In the last article, I wrote about how GraphQL has been rapid in its adoption and the possibilities GraalVM in the database could offer to run GraphQL from within the database. Making it a powerful solution. Here we're going to pick up the point that with great power comes great responsibility. So, what might that responsibility look like?*

## Responsible Security

You could look at responsibility as what I should know, what I should use, and what I should do. In this day and age, we should all take time to understand data legislation, even if that is so we know and understand what we're agreeing to when it comes to our own data. The next step in developing such understanding is becoming familiar with the Open Web Application Security Project (OWASP) resources. OWASP provides a wealth of useful material around web security, both in terms of causes and mitigations. One of the best-known pieces of work from OWASP is their Top 10. Every couple of years, the experts review the causes of security breaches and identify the top 10 common problems. Until 2019 this was focussed on applications, but now APIs have been separated out. There are plenty of short training courses that can be taken covering each of the Top 10 and how to avoid the pitfalls. **Figure 1** shows the 2019 (current) Top 10.

Capgemini, like many companies, regularly asks staff to complete some e-learning on security, which typically focuses on the more general rules for document handling. I encourage my team to use this as a prompt to see if they need to refresh on the OWASP Top 10 training.

## OCI Services

When you review the Top 10 (see Figure 1), you'll note a number of the issues revolve around configuration and errors with authorization. We're fortunate enough to have both a Web Application Firewall (WAF) and an API Gateway in the Oracle cloud. From a superficial view, it would be easy to assume that you need only one or the other but would be an error. Unlike API Cloud Service, which couldn't accept the availability of a WAF to partner in OCI, this assumption becomes an option. The OCI API Gateway is designed to work with and complement the WAF. We use the WAF to perform the first level of screening. The API Gateway then applies API-specific rules. Today the API Gateway has a strong focus on JSON and REST support, which means we can use some

Not be the cause of a security issue (or better be the person who helped prevent one)

- |  |   |
|--|---|
|  A1:2019- Broken Object Level Authorization   |  A6:2019- Mass Assignment                    |
|  A2:2017- Broken Authentication               |  A7:2019 - Security Misconfiguration         |
|  A3:2019- Excessive Data Exposure             |  A8:2019 - Injection                         |
|  A4:2019 - Lack of Resources & Rate Limiting  |  A9:2019- Improper Assets Management         |
|  A5:2019- Broken Function Level Authorization |  A10:2019- Insufficient Logging & Monitoring |

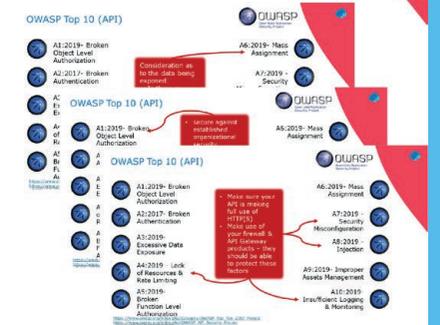


Figure 1: API OWASP Top 10 (previous presentations show factors to consider and differences between Application and API Top 10s, see useful links to get more info).

of the OOTB (Out Of The Box) policies. For example, rate-limiting (even with an elastic backend, rate-limiting protects us from cost spikes for unexpected loads, which could be triggered maliciously or erroneously). Where OCI API Gateway can't help us, we can build small solutions ourselves. For example, using the GraphQL schema to verify the requests in more detail.

Another form of limiting that should be applied is for the amount of data requested or allowed to be pushed. If you let users put wild cards into the query (resolver), then your backend could simply end up asking for every record in the system. This might be serviced – and someone has stolen all your data, or more likely, your systems struggle under load. GraphQL supports the concept of pagination, which can ease this problem, but it won't prevent data exfiltration. Conversely, allowing too big a payload for a Create or Update is a form of Denial of Service (DoS) attack as it just swamps your servers with the workload.

Perhaps the most impacting aspect of the OWASP Top 10 is how we implement and apply authentication and authorization

showing in A1, A2, A5, and A7. The bottom line here is, despite the temptation to manage identities and credentials ourselves, we need to manage a lot of user attributes already, so why not add the password. It's easier than integrating an

Identity Management solution directly or federate identity through Identity Cloud Service (IDCS). The reality is we've created long-term headaches. Think of it this way, how many specialist and dedicated engineers are employed by Oracle on IDCS or by Okta. Compare that to how many people you have committed to managing your application credentials, an aspect that is probably not part of the solution's core value proposition (unless you work on IDCS and Okta). Not sold, then let's apply the expression often used with Open Source – many eyes make all bugs shallow – that is to say, the more people looking at code, the more likely they'll spot bugs and fix them. A dedicated team will have more eyes to spot any problems, and even then mistakes can and do happen.

## Good Schemas

Ensuring that our API schema is concisely defined will help our security and reduce our workload. This means we should avoid the classic error of making every attribute a string – strings are the easiest structure to hide injection attacks in. Bounding lengths means you're limiting the chance of getting an injected expression

in – trying putting a SQL statement into 30 characters before you even add the escape sequences to mask the attack. Using numeric values eliminates the risk as well. Well-defined lengths give us several bonuses. We can lean on libraries to validate strings rather than us having to evaluate every field explicitly. Better still, you can avoid data integrity issues because someone had written the word “one” when you wanted the number “1”. The most straightforward rule to use is: big attributes can be big trouble.

## API 1st, the Practical Design Methodology

API 1st practices – a subject that I have written about before, and am a firm believer of. While not a prominent contributor to security, there are benefits to be had in securing your API design. An API 1st model puts the candidate consumer(s) and provider into the domain problem space, not the specific implementation, because you collaborate on the specification. Doing so means we can control the risk of over-sharing (A3 – Excessive Data Exposure) so far as working through the domain will provide insight into what data is needed and what may be required in the future. This may sound like I've contradicted one of the benefits of GraphQL, but the truth be told, we need to walk a balancing act here. GraphQL framework provides a lot of support for security from operations and entity perspective. But there are also open-source utilities that allow you to introduce attribute level controls.

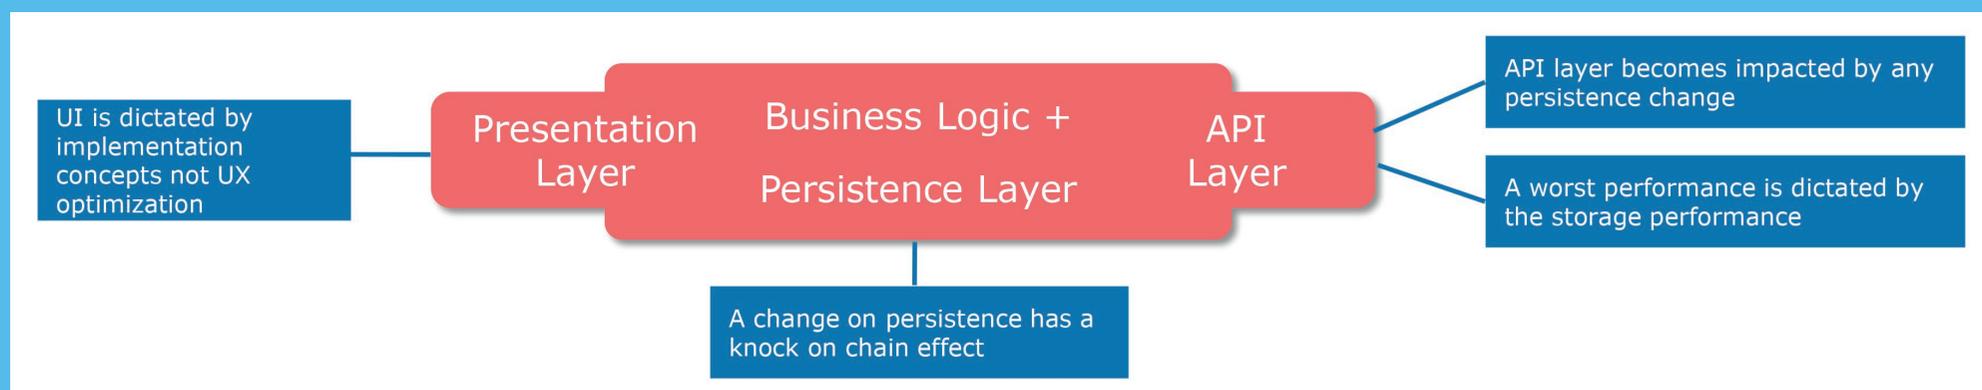


Figure 2: How code led APIs can cause abstraction and application layering to collapse.

Another key benefit of the API 1st approach is that the interface is designed for the problem, not the solution implementation. This means it is possible to apply some levels of abstraction in the code, cushioning the interface from the complex practicalities of data storage. The moment the interface definition becomes informed by implementation, we regularly see a couple of things happening:

- The abstraction layers start getting folded in on themselves – losing the benefits of the abstraction
- The chances of information leaking on how the back end is implemented increase
- While GraphQL provides a good framework for coping with change, it is not impermeable

This is part of why I'm not a huge fan of incorporating APIs directly into databases. It's not that we can't build the abstraction, etc.; it's just too tempting to do the quick and easy, wrong thing when the pressure sooner or later comes down to make a change. We can see in **Figure 2**, how the compression/collapsing of abstraction can impact an API solution.

## Summing Up

The recommendations I've made here aren't unusual, and many of them are even unique to GraphQL. Although how we implement some of these can differ for GraphQL compared with traditional REST APIs. Most of these suggestions aren't even difficult to put into place.

## Useful Reading

The following link to some helpful resources:

- <https://owasp.org/www-project-api-security/>
- [https://cheatsheetseries.owasp.org/cheatsheets/GraphQL\\_Cheat\\_Sheet.html](https://cheatsheetseries.owasp.org/cheatsheets/GraphQL_Cheat_Sheet.html)
- <https://www.howtographql.com/advanced/4-security/>
- [https://owasp.org/www-chapter-vancouver/assets/presentations/2020-06\\_GraphQL\\_Security.pdf](https://owasp.org/www-chapter-vancouver/assets/presentations/2020-06_GraphQL_Security.pdf)

- <https://www.ateam-oracle.com/end-to-end-security-using-oci-waf%2C-oci-api-gateway-and-oci-functions>
- <https://blog.mp3monster.org> – my blog with more content and slides on this and other subjects.



## About Phil Wilkins

**Phil** is an Oracle Ace Director, Technology Evangelist & Snr Consultant Architect for Capgemini. He specializes in PaaS and particularly with API and modern development techniques such as microservices and Integration Cloud. Phil has supported the publication of several Java development books as a technical reviewer; in addition to being a published author himself with several co-authored books to his name, along with regular blog posts (<https://blog.mp3monster.org/>) and magazine articles. His latest book covers Unified Logging. Phil has presented at events around the world from Sweden to California and is a co-organizer to the London Oracle Developer Meetup.

# Oracle Application Express (Part 8): APEX 21.1: What's new?

Carsten Czarski



## Introduction

APEX 21.1, the first APEX release in 2021, has recently been released. As always, the new release contains marquee features like the all-new Map Region, Date Picker, and Application Data Loading. In addition to these, there are a number of smaller features, like native support for the Markdown language, Faceted Search enhancements and the ability to import APEX applications in ZIP format. Even Classic Reports, the component which exists in APEX since "Day One", got some love in 21.1.

This article provides an overview of the key new additions and improvements in APEX 21.1.

## New "Map" Region

One of the marquee features is the new **Map** region, marking the beginning of native Spatial support in Application Express. Displaying coordinates or other spatial data on an APEX page becomes as easy as displaying a Chart: The developer provides a SQL Query (as the data source), some column mapping information (which column contains the spatial data?) and a few details about the type of visualization. Based on this, APEX provides a modern and interactive map visualization.

**Figure 1** shows how an APEX page with both a map and a report region can look like. **Figure 2** shows the same page in Page Designer.

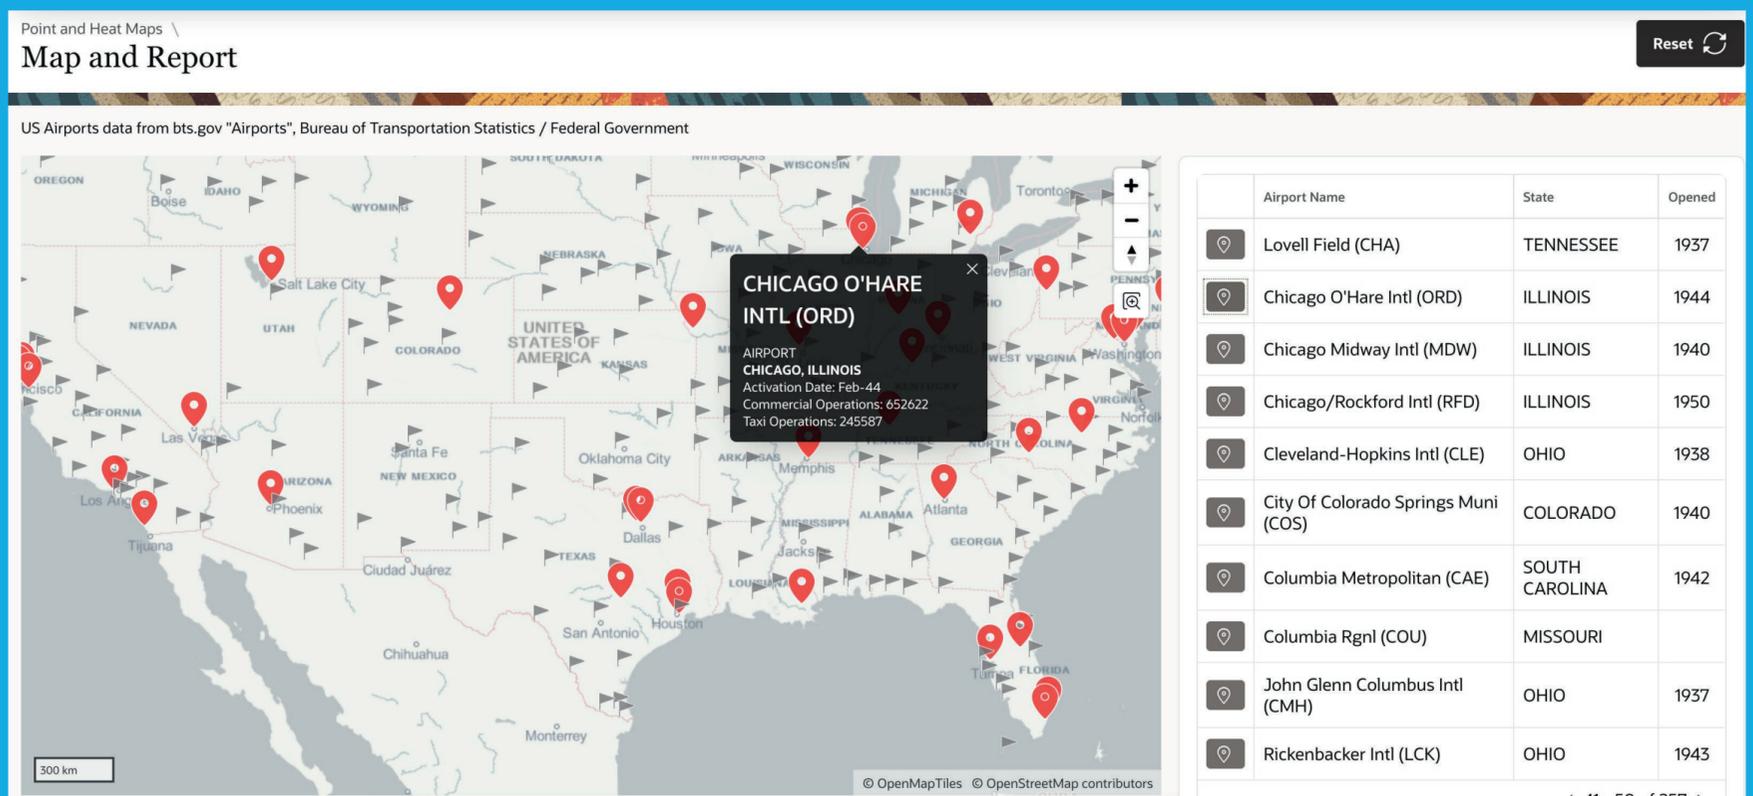


Figure 1: The new Map Region in APEX 21.1

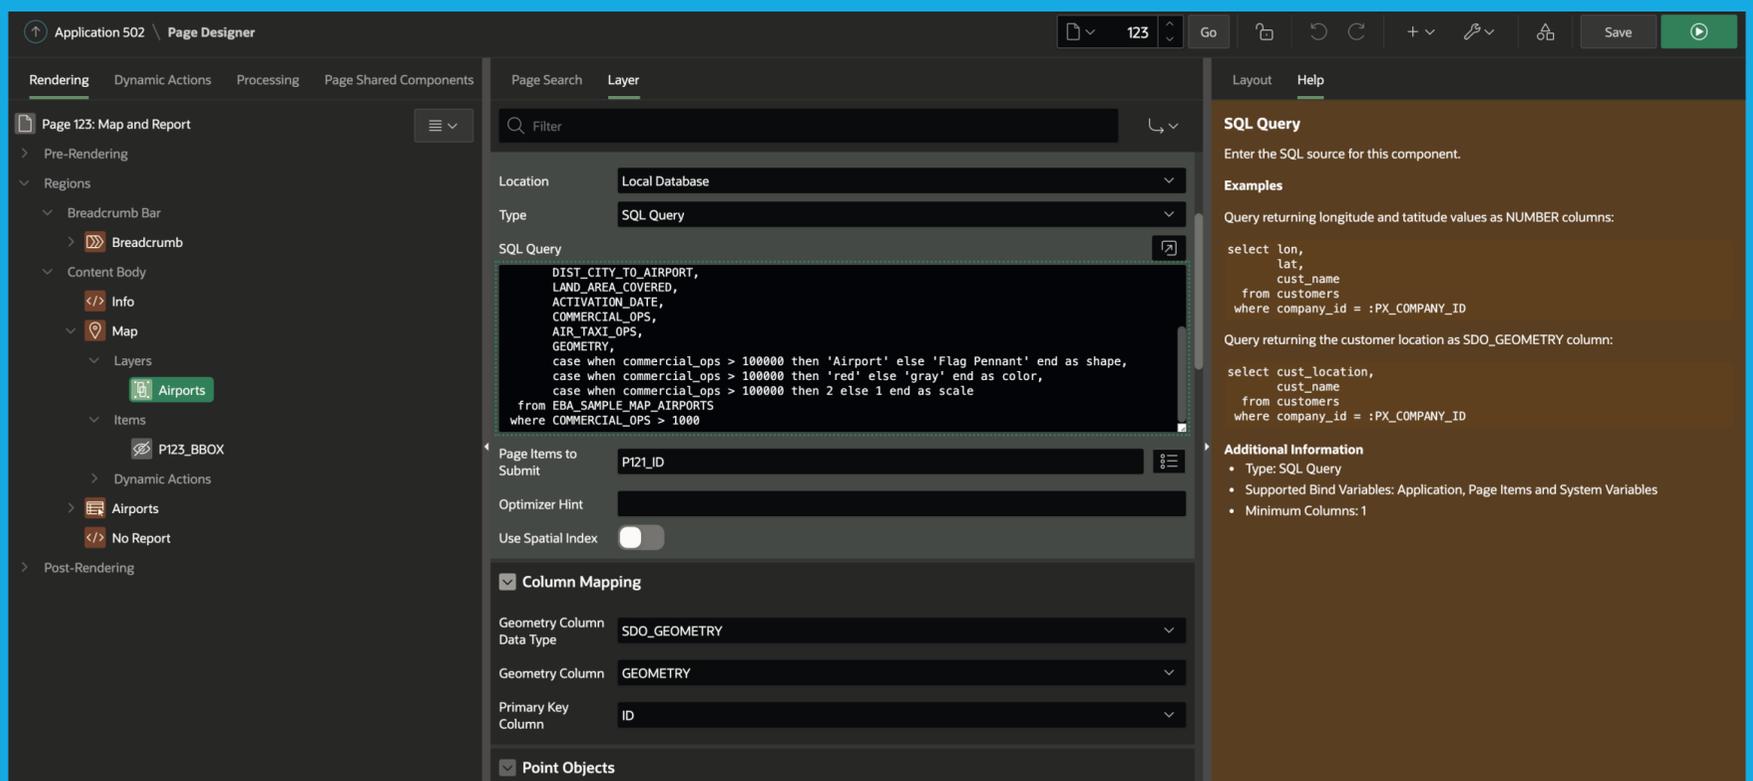


Figure 2: Configuring a map region with Page Designer

As other APEX components, the map region also supports fetching data from a remote database using *REST Enabled SQL*, or from an external REST API, using a *REST Data Source*. The Map region can process coordinate data from two numeric columns (longitude and latitude), from VARCHAR2 or CLOB columns (GeoJSON format), or from a column of the Oracle SDO\_GEOMETRY data type. The latter is of course most recommended, since it can be indexed and allows to use powerful Spatial SQL functions.

Besides the visualization with **Map Markers** (illustrated in Figure 1) the Map region also supports rendering spatial data as **lines**, **polygons**, **3D polygons** and as a **heat map**. The latter is very useful to visualize the spatial density of data, as shown in **Figure 3**. The color as well as the height for a 3D polygon can be derived from a *value column* returned by the SQL query. APEX also allows to use the Map Region in combination with Faceted Search; enabling developers to easily build powerful, spatial analysis pages with just a few clicks (**Figure 3**).

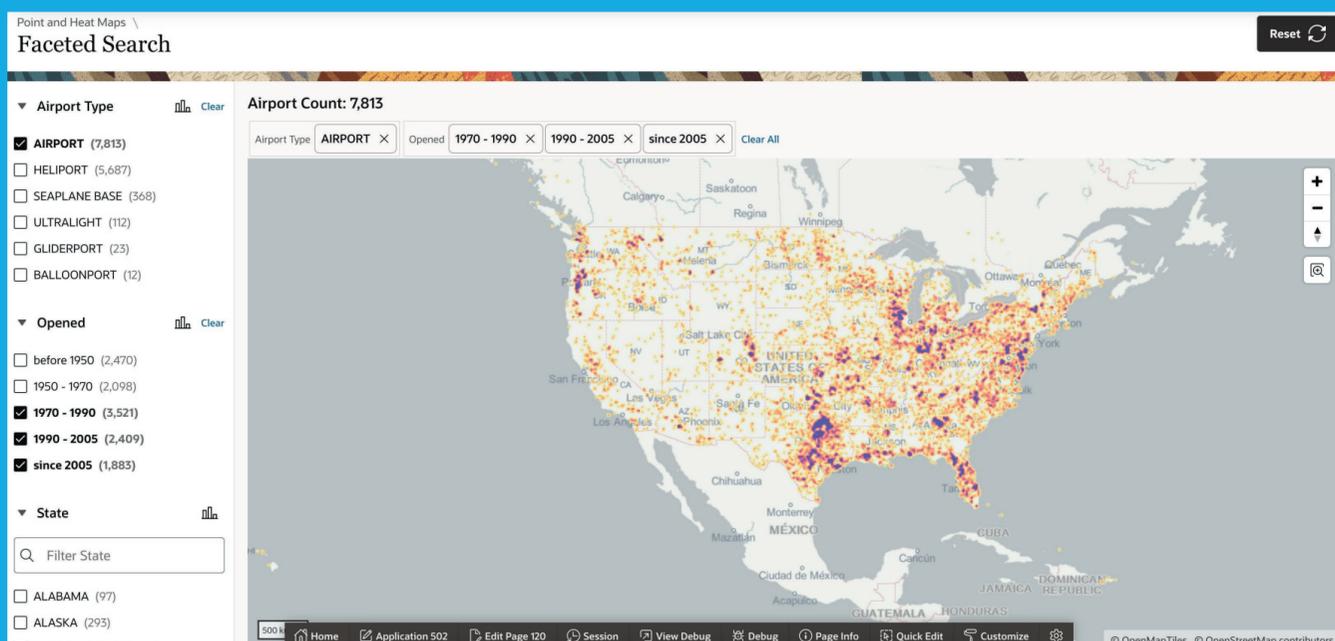


Figure 3: The new Map Region (showing a "Heat Map") in combination with Faceted Search

Support for the Oracle Spatial SDO\_GEOMETRY type has also been added to APEX PL/SQL APIs like **APEX\_EXEC** and **APEX\_JSON**. This allows Plug-In developers to natively work with Spatial data as well.

## New Application Data Loading

APEX 21.1 contains a totally redesigned **Application Data Loading Wizard**, which allows developers to add Data Loading pages to their applications.

Similar to SQL Workshop, the new Application Data Loading supports not only CSV, but also XLSX, JSON or XML files. The functionality is built on top of the **APEX\_DATA\_PARSER** PL/SQL package, which has been available to PL/SQL developers for file parsing since APEX 19.1.

While the old Application Data Loading wizard walks the end user through a 4-step process of data loading (upload a file,

map columns, verify, load), the new Application Data Loading is designed to be much simpler.

1. **Column Mapping** is determined by the Developer *at design time*, and not by the end user *at run time*. Columns in the uploaded file are mapped to table columns either by name, by position or using a regular expression, which provides more flexibility, for instance to map columns based on a prefix. **Figure 4** shows how to maintain column mapping for a Data Load Definition in Shared Components.
2. All Data Loading tasks can be implemented *on one application page*. After uploading the file, the end user can see a **Data Preview**. After clicking the **Load Data** button, data is immediately loaded to the configured target (**Figure 5**).
3. Data Loading not only supports loading to a table, but also to an **APEX Collection**. This allows developers to use the collection as a staging area and to apply some custom PL/SQL logic before finally loading to the target table.

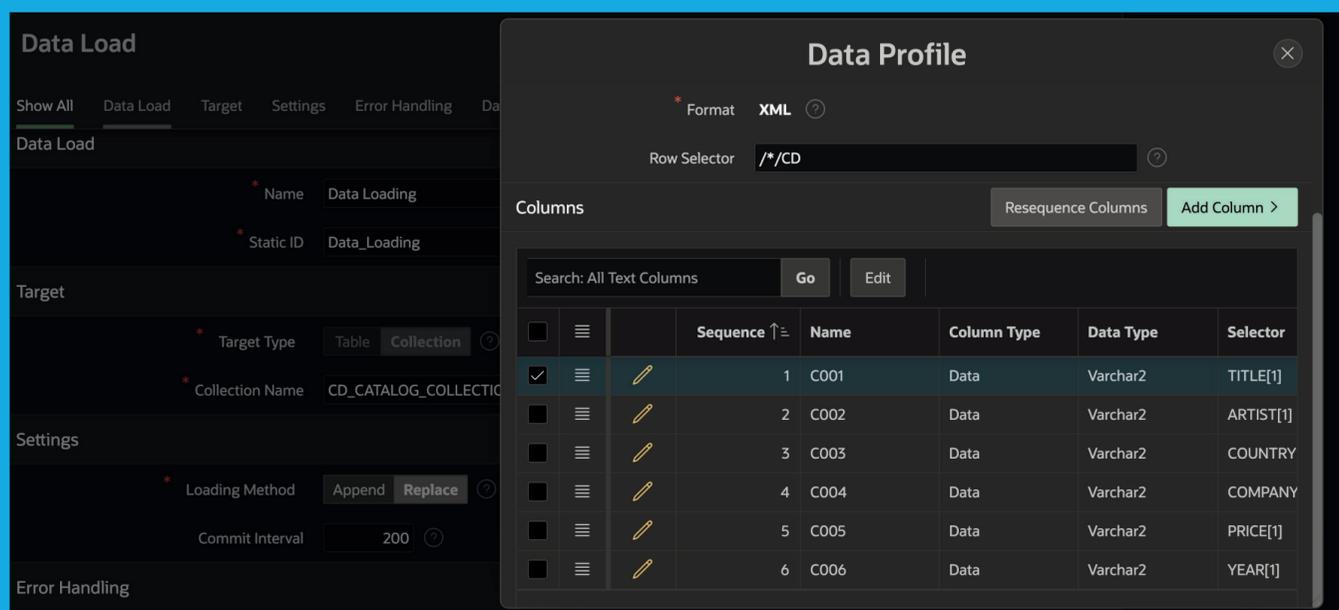


Figure 4: New Data Loading in Shared Components

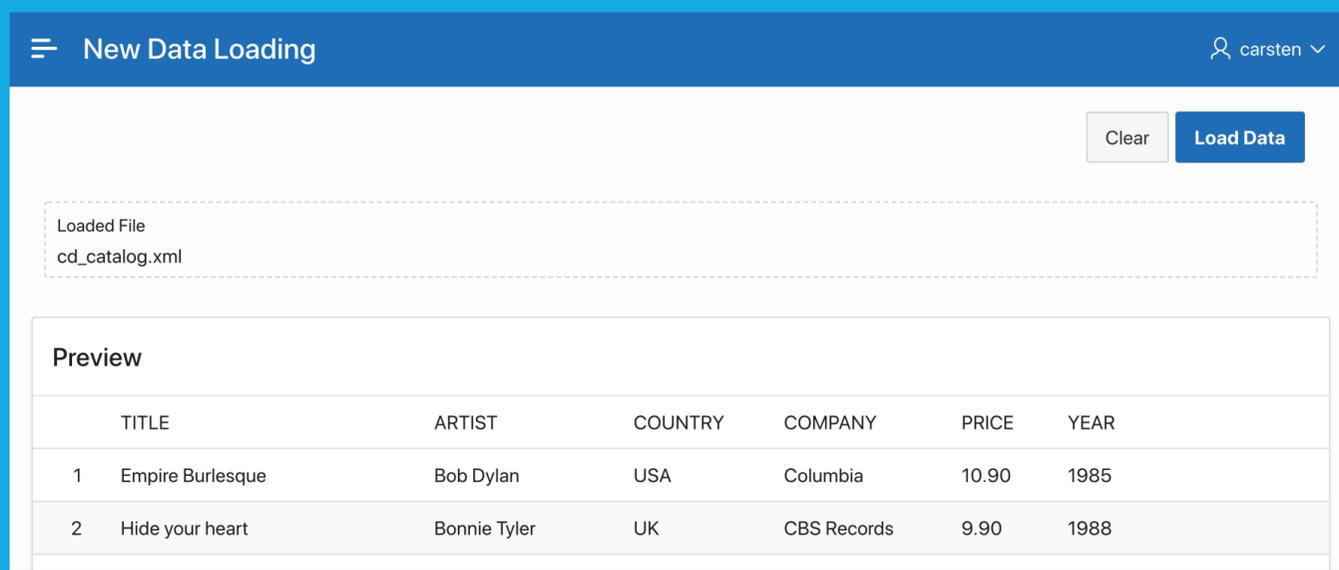


Figure 5: Upload a file, verify the Data Preview and load data with one click.

- Error Handling is configurable: By default, an error aborts the loading process and shows a message. As an alternative, errors can be logged or just silently ignored.
- Finally, all components of the new Application Data Loading functionality can be used by developers without the Create Page wizard: Developers can add the **Data Loading** process to a page using Page Designer. The **APEX\_DATA\_LOADING** PL/SQL package provides the same capability from custom PL/SQL code, for instance within Automations or PL/SQL Page Processes.

The old Data Loading wizard is still available; developers will find it in the **Legacy Page** section in the Create Page wizard. New Application Data Load definitions are created and maintained within Shared Components, Data Load Definitions.

## REST Data Source Enhancements

New Application Data Loading shares internal infrastructure with REST Data Sources. This has allowed the APEX development team to expose some of the new Data Loading functionality to REST Data Sources. REST Data Sources can be created on a CSV file on the web. So, REST Data Sources now support JSON,

XML and CSV formats. XLSX content is not supported (**Figure 6**). Transformation Rules and Lookup Queries can also be used in REST Data Sources.

## Faceted Search

Faceted Search, the flagship feature of APEX, got a few enhancements as well. Facets can be declared as optional, using the **Display Toggling** switch in Page Designer (**Figure 7**).

End users can toggle display of this "optional" facet now. Whether it displays on the first page view, depends on the **Initially Displayed** setting. This feature allows an APEX page to focus on the most important facets and to keep the UI clean. Since APEX memorizes the display status of an optional facet, end users can "configure" their faceted search page, according to individual needs.

## Item Types

APEX 21.1 finally comes with a modernized **Date Picker** item type, which replaces the existing (and now deprecated) Date Picker based on jQuery UI. The new Date Picker is based on the Oracle JET date and time components.

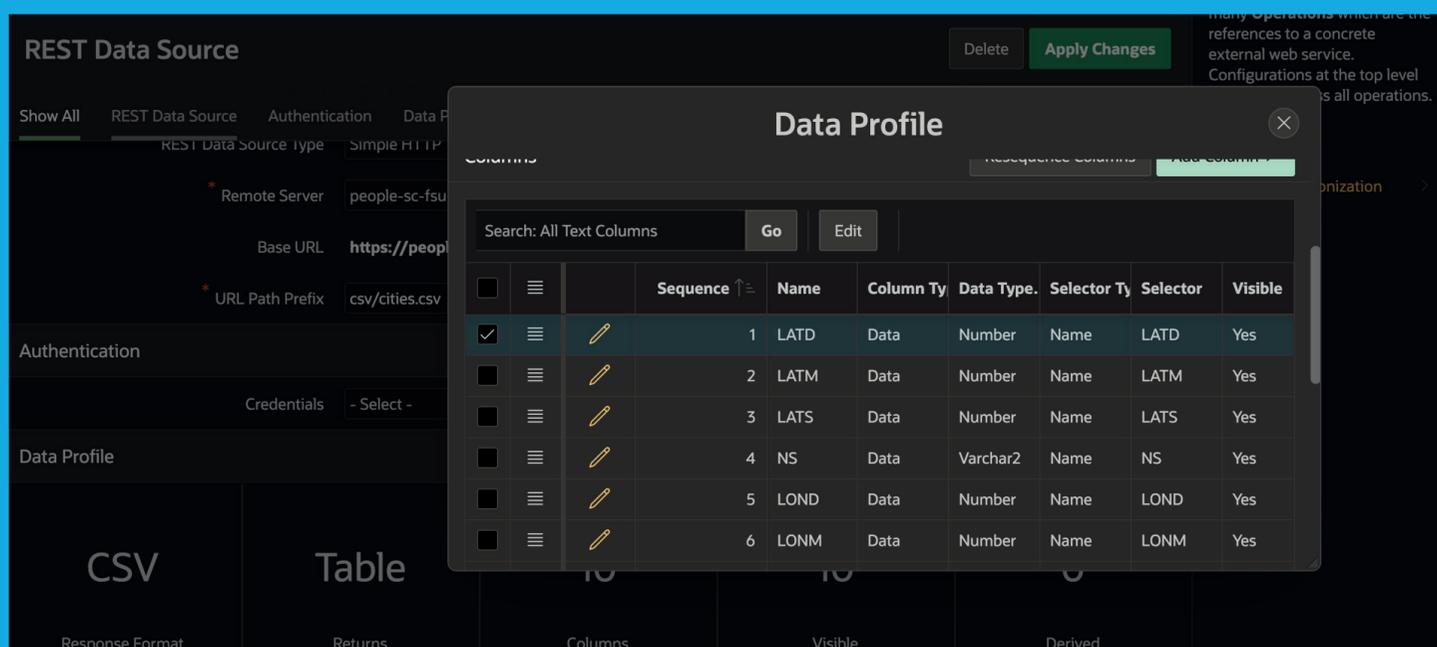


Figure 6: Use a "CSV file on the web" as a REST Source.

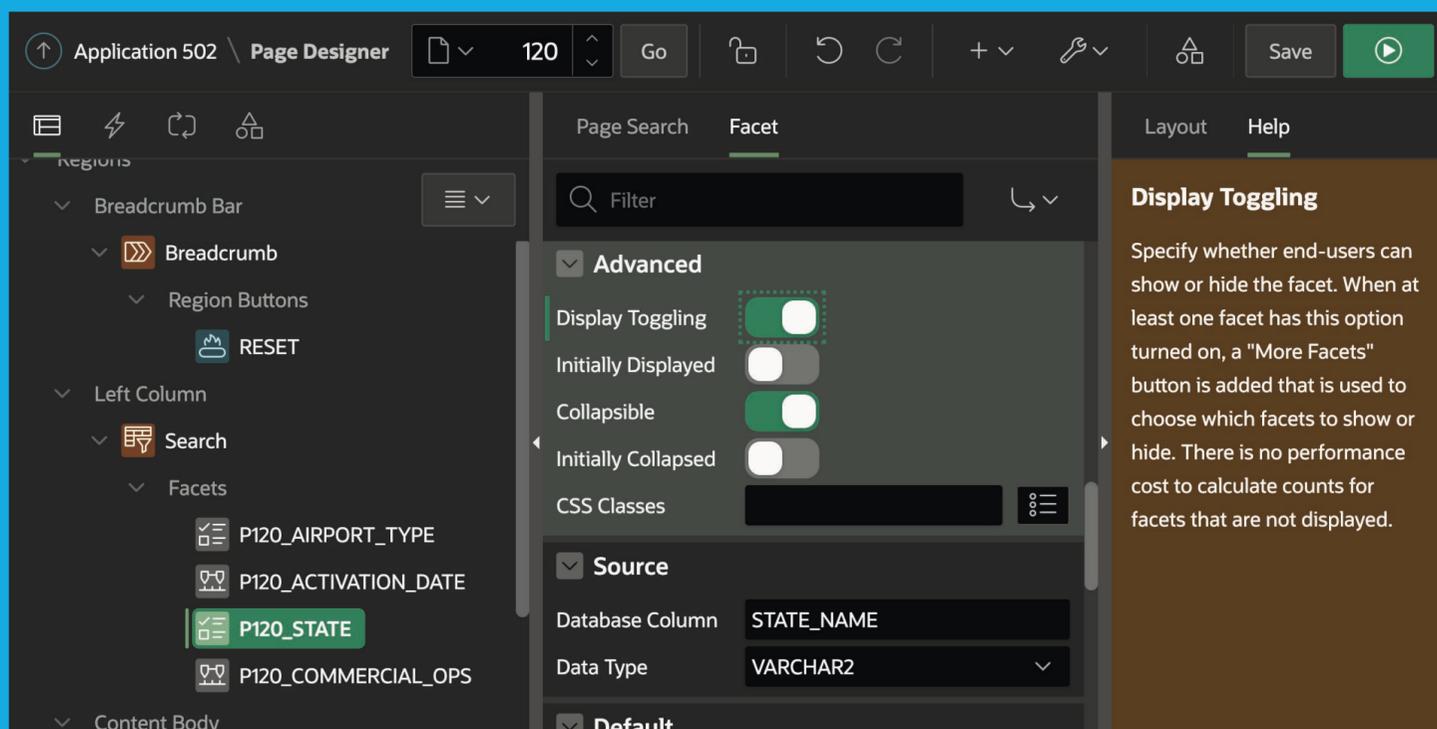


Figure 7: Optional Facets configuration in Page Designer

The new item type provides the following improvements:

- To allow end users entering a time, the item type now has a declarative **Show Time** switch. This is much cleaner than before, when this was dependent on whether the format mask had a time component.
- Developers can declaratively specify minimum and maximum dates, which was often the reason for individual date picker customizations before (*Figure 8*).
- The UI of the new date picker has better accessibility support and also provides improved month and year selection capabilities.
- For improved mobile UX, there is also a native mode, which uses native browser date elements.

When an APEX instance is upgraded or applications from older APEX releases are imported into APEX 21.1, existing date picker items will remain unchanged. This is to make sure that date pickers with custom JavaScript extensions still work as before.

However, in the Application Upgrade section, these Date Pickers can be automatically migrated to use the new date picker component.

When using the new date picker in existing applications, developers should make sure to use the latest version of Universal Theme. This is also a general recommendation in order to pick up latest styling enhancements and theme bug fixes.

The **Color Picker** item is also replaced with a new implementation, based on Oracle JET. It provides different display modes (popup vs. inline, with and without an input field) as well as optional contrast checking and different color return formats like RGB, RGBA, HSL, HSLA or CSS color values.

The **Number Field** item type now does automatic formatting and client-side validation based on the item's format mask, which is configured in Page Designer.

APEX 21.1 includes native support for the **Markdown** language. Markdown is very useful for formatted text provided by end users: Markdown contains only a restricted set of the most common text formatting features. Thus, Markdown based editors are simpler than HTML editors, and also more secure,

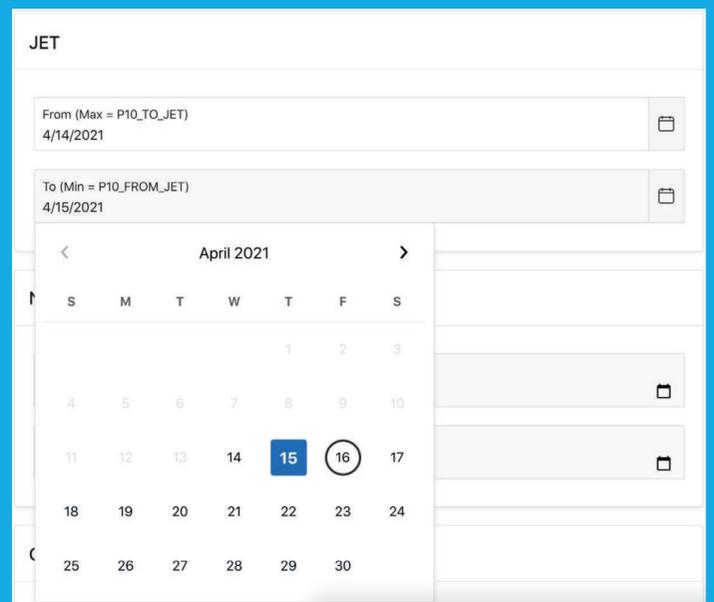


Figure 8: The new Date Picker item type allows declarative date range constraints.

since there is no *Cross Site Scripting* with Markdown. APEX does not allow to embed HTML markup in Markdown content.

The new Markdown support in APEX includes additions to the **Display Only** (display Markdown) and **Rich Text Editor** (generate Markdown, *Figure 9*) item types. **Rich Text** column types were added to Classic and Interactive Reports, and the **APEX\_MARKDOWN** PL/SQL package allows to process markdown programmatically.

### Other Enhancements and Improvements

**Classic Reports** now have a **Lazy Loading** attribute. Once enabled, report contents will be fetched using separate AJAX requests, and not as part of page rendering. This can be used for pages containing expensive report queries to respond faster: the page with the "empty" report renders immediately, the report contents display shortly after query execution completes.

While APEX 20.1 introduced the ability to export an APEX application as a ZIP file, APEX 21.1 closes the loop on this: These ZIP files can now be imported back into APEX (*Figure 10*).

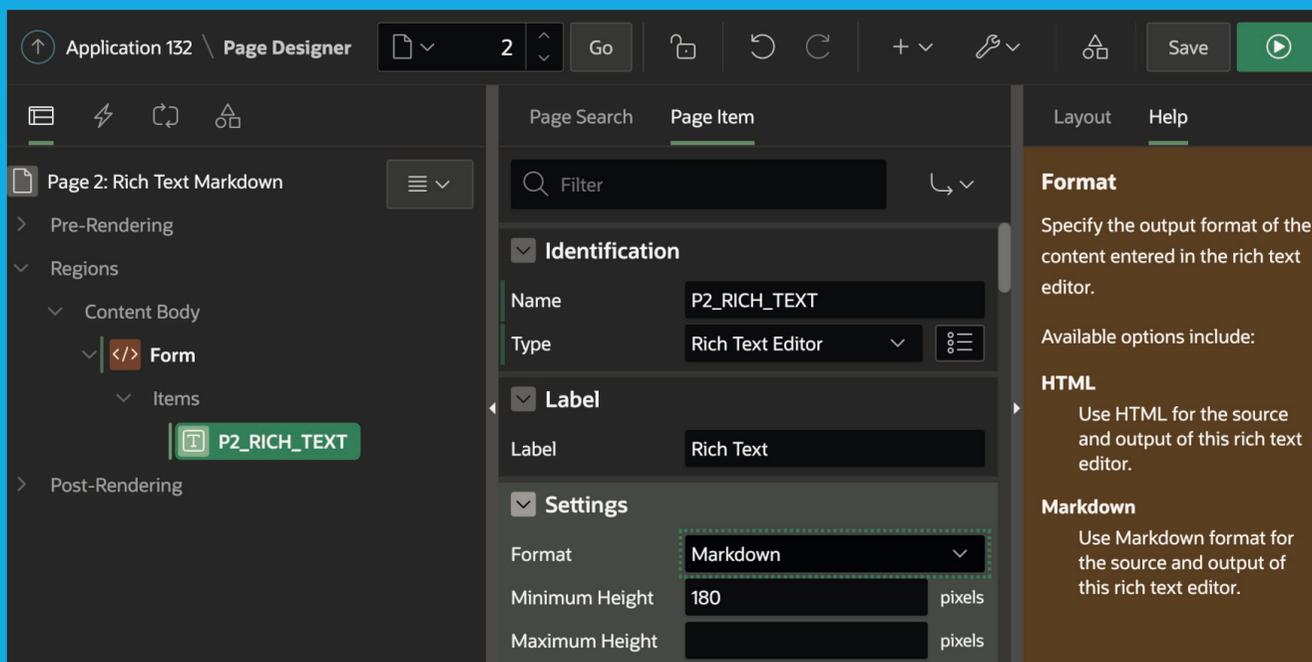


Figure 9: Native Markdown support for the Rich Text Editor

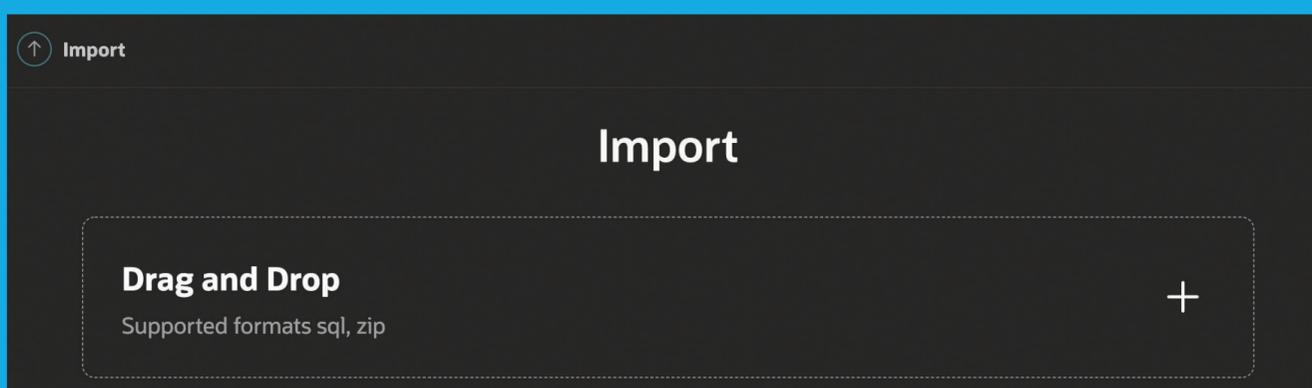


Figure 10: APEX export files in ZIP format can now be directly imported.

**Native Report Printing** was introduced with APEX 20.1 and 20.2. When exporting to PDF, languages requiring *complex text layout*, like Hebrew or Arabic, are now supported.

The **Gallery of Packaged Applications** is no longer part of APEX itself. The Gallery section in the APEX application builder now points to an external landing page on [apex.oracle.com](http://apex.oracle.com); the actual packaged applications are hosted in a GitHub repository. To install a packaged app, simply download it from GitHub and import it to the APEX workspace.

As with every APEX release, underlying JavaScript libraries were upgraded to the latest versions, including:

- Oracle JET 10.0.0
- FullCalendar 5.5.1
- CKEditor 25.0.0
- Monaco Editor 0.22.3

## Summary

APEX 21.1 contains a number of new features and enhancements. The release is particularly interesting for applications dealing with coordinates and Spatial data: Map display can now be implemented with out-of-the-box functionality, additional plug-ins or custom coding is not required any more.

Data Upload pages within applications can now be implemented in a much more flexible manner than before: end users can upload data with just one click and don't have to walk through the lengthy wizard anymore.

And all the APEX developers waiting for a more modern Date Picker item will now find the new Oracle JET based component in APEX 21.1. Besides these three, countless improvements and new features are waiting to be explored.

## More Information

- Information and environment to test-drive APEX:  
<http://apex.oracle.com/en>
- Oracle Application Express Blog:  
<http://blogs.oracle.com/apex>

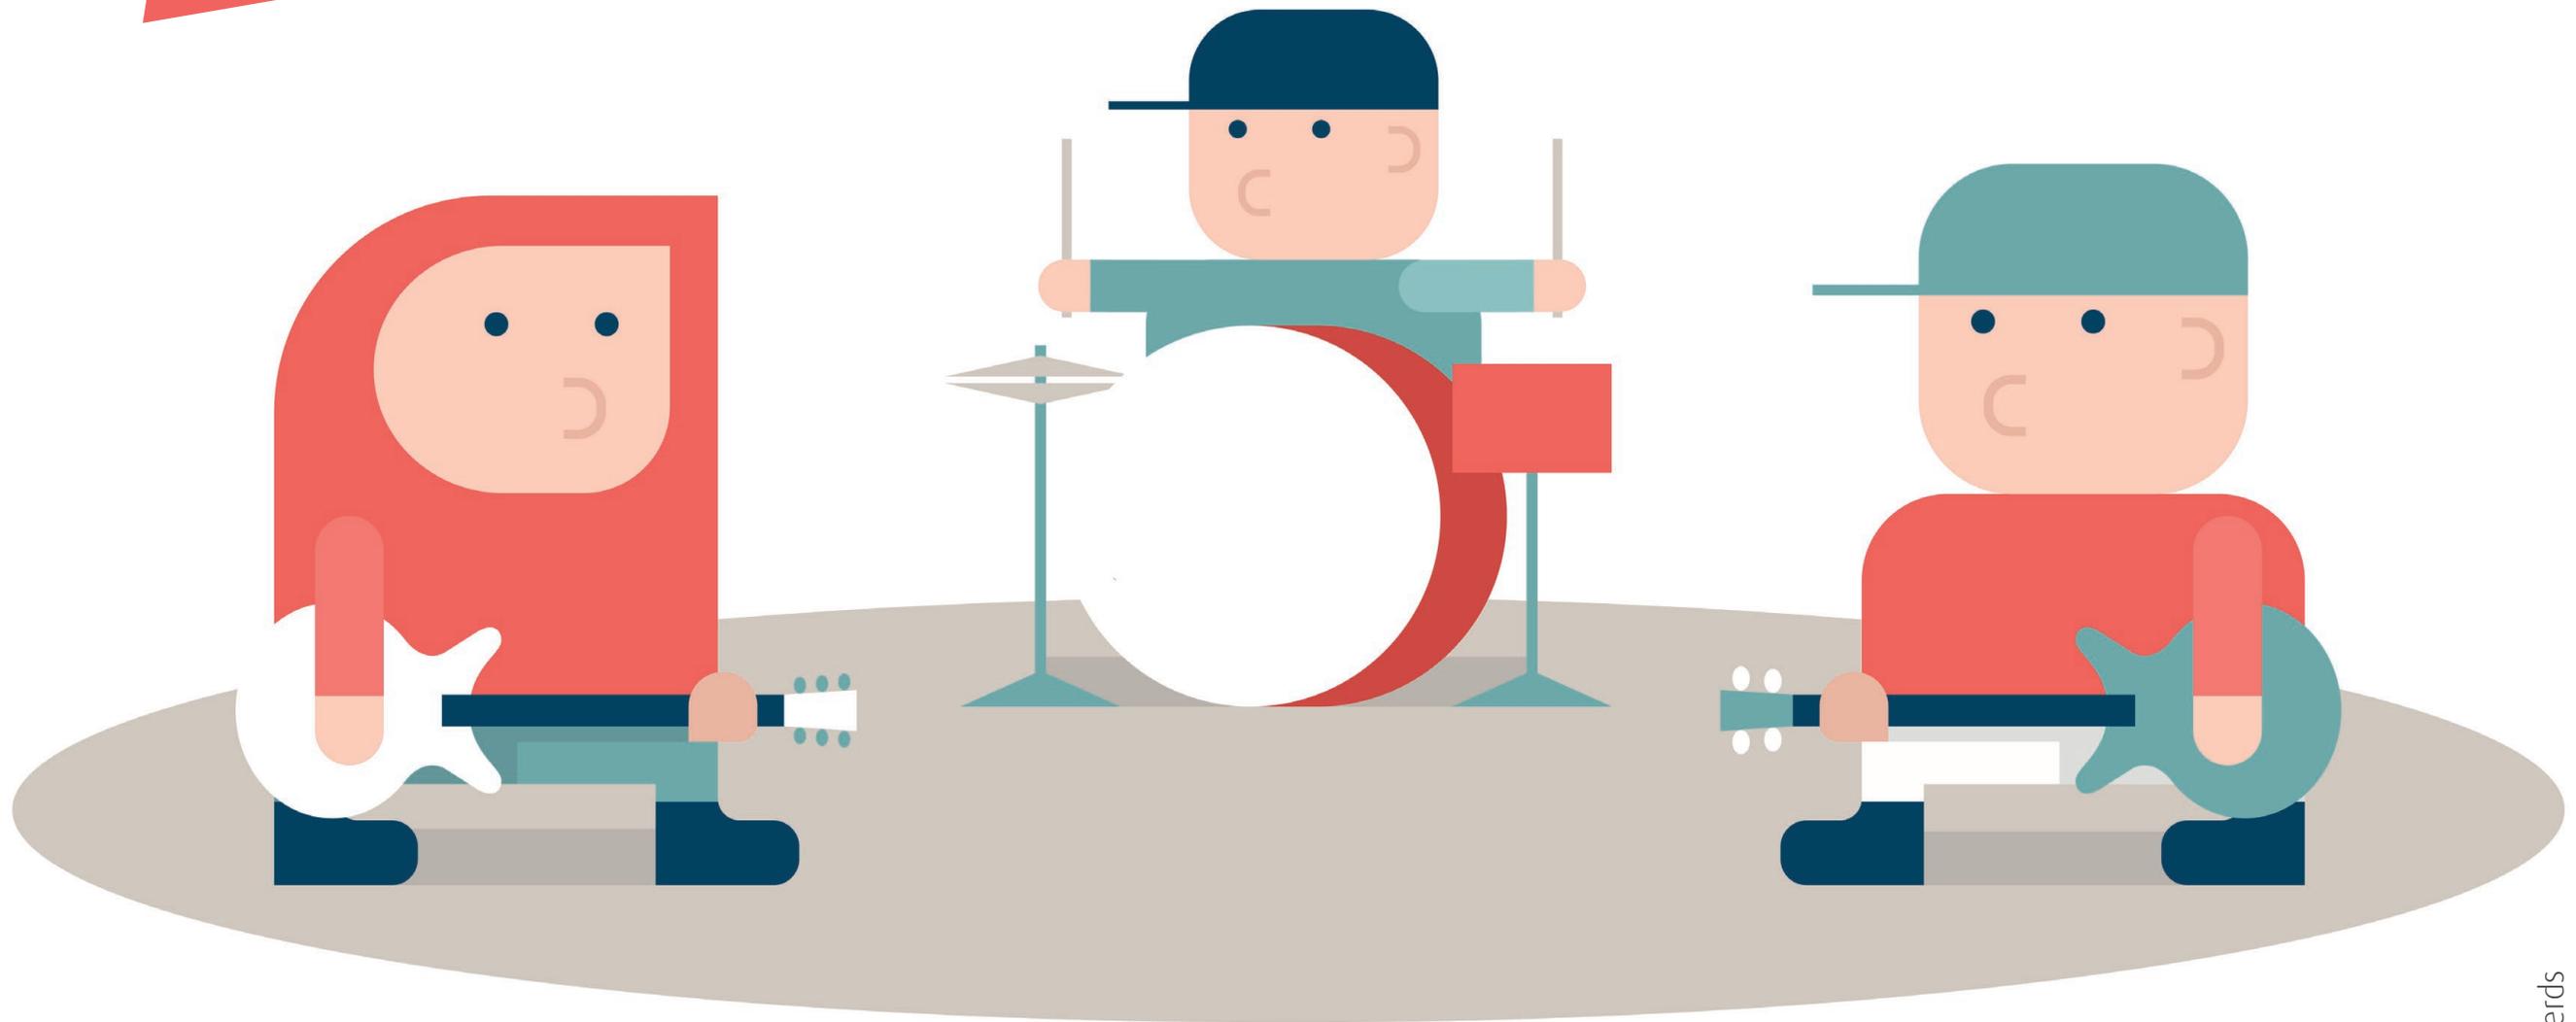


## About Carsten Czarski

**Carsten** works for Oracle in Germany since 2001. He started in the Presales organization helping customers and partners regarding database-centric application development. Since March 2016, Carsten is a member of the Application Express development team. Focus of his work is on the new support for REST services in APEX – beyond that Carsten looks after the Calendar component and the Data Loading facility. He is a frequent speaker at international user group conferences.

Adrian Png

# NEW KID ON THE BLOCK: OCI LANGUAGE SERVICE



The Oracle Cloud Infrastructure (OCI) Language service is the newest kid on the block for Oracle's long list of AI/ML oriented offerings, possibly the first of many to come. These services are designed for utility and ready to use without any knowledge about machine learning necessary. Pretrained models like these and a complementary REST API will help increase adoption by developers to include AI/ML technology in the applications they build.

The Language service currently includes five language processing operations, and they are:

- Aspect-based Sentiment Analysis
- Named Entity Recognition
- Key Phrase Extraction
- Language Detection
- Text Classification

**Aspect-based Sentiment Analysis** uses Natural Language Processing (NLP) to analyze a body of text, extract aspects, and then classify them into three sentiments: positive, neutral, or negative. The results also include a confidence score between 0 and 1 for each sentiment, with 1 representing a higher amount of confidence.

There are a few limitations that must be considered. The API currently only supports the English language, and the model performance is affected by grammatical and spelling mistakes. If the text contains other languages or non-ascii symbols, only the English text will be processed. The submitted text must contain between 1 and 1,000 characters.

**Named Entity Recognition** as the name implies, detects entities in text. The results return contains all the entities found, the type of entity, position of the text where it was found, and a

confidence score. It supports a wide variety of entity types like organizations, events, persons etc. Brings back memories of [OpenCalais](#).

Again, sentence structure is important, and it will affect the performance of the processor. The body of text should contain as much contextual information as possible, or you might not get the results expected.

**Key Phrase Extraction** analyzes the text document and extracts keywords and key phrases along with confidence scores that range between 0 and 1.

**Language Detection** is a feature that evaluates a body of text and infers the language. According to the documentation, it currently supports about 75 languages. The results also include the language code and a confidence score. Unfortunately, the processor can only deal with one language at a time, and the text should preferably contain 10 or more words.

**Text Classification** uses Deep Learning techniques to analyze the submitted text and suggests a text category with a confidence score. It appears to only support the English language currently, and it requires at least 10 words for a useful determination of the text category. Only one category result is returned even if the classifier finds multiple suggestions.

The fastest way to test drive the Language service is to use the OCI Command Line Interface (CLI). For example, perform the sentiment detection using this command:

```
oci ai language detect-sentiments \  
  --text "I feel great with the launch of OCI AI Language Services!"
```



It should produce the following JSON output:

```
{
  "data": {
    "aspects": [
      {
        "length": 6,
        "offset": 22,
        "scores": {
          "Negative": 0.0004437448224052787,
          "Neutral": 0.0,
          "Positive": 0.9995562434196472
        },
        "sentiment": "Positive",
        "text": "launch"
      },
      {
        "length": 24,
        "offset": 32,
        "scores": {
          "Negative": 0.00011140753485960886,
          "Neutral": 0.0,
          "Positive": 0.99988853931427
        },
        "sentiment": "Positive",
        "text": "OCI AI Language Services"
      }
    ]
  }
}
```

Oracle is also rolling out support for the new Language service in their OCI Software Development Kits (SDKs). However, please note that you will need to configure the CLI or SDKs to run as a user with the necessary privileges and APIs. We will discuss some of the necessary steps shortly.

In the next and final segment, we will examine how you can integrate the OCI Language Service in Oracle Application Express (APEX). APEX is a low-code programming platform with extensive support for working with web services. More recent versions of APEX include support for authenticating with *OCI REST APIs*. Please note that it is assumed that you, the reader, are somewhat familiar with OCI concepts.

## Required OCI Resources

As a best practice, we should create an OCI *Identity and Access Management* (IAM) user specifically for interacting with the OCI REST API on behalf of the APEX application. The user should have limited OCI capabilities and should only be allowed to have API access and nothing more. Also, create an IAM Group, e.g., *AiLanguageUsers*, and assign the user to the group.

To work with the OCI REST API, you will need to generate the API keys for signing the HTTP requests. This can be done easily from the user's details page in the OCI console (*Figure 1*).

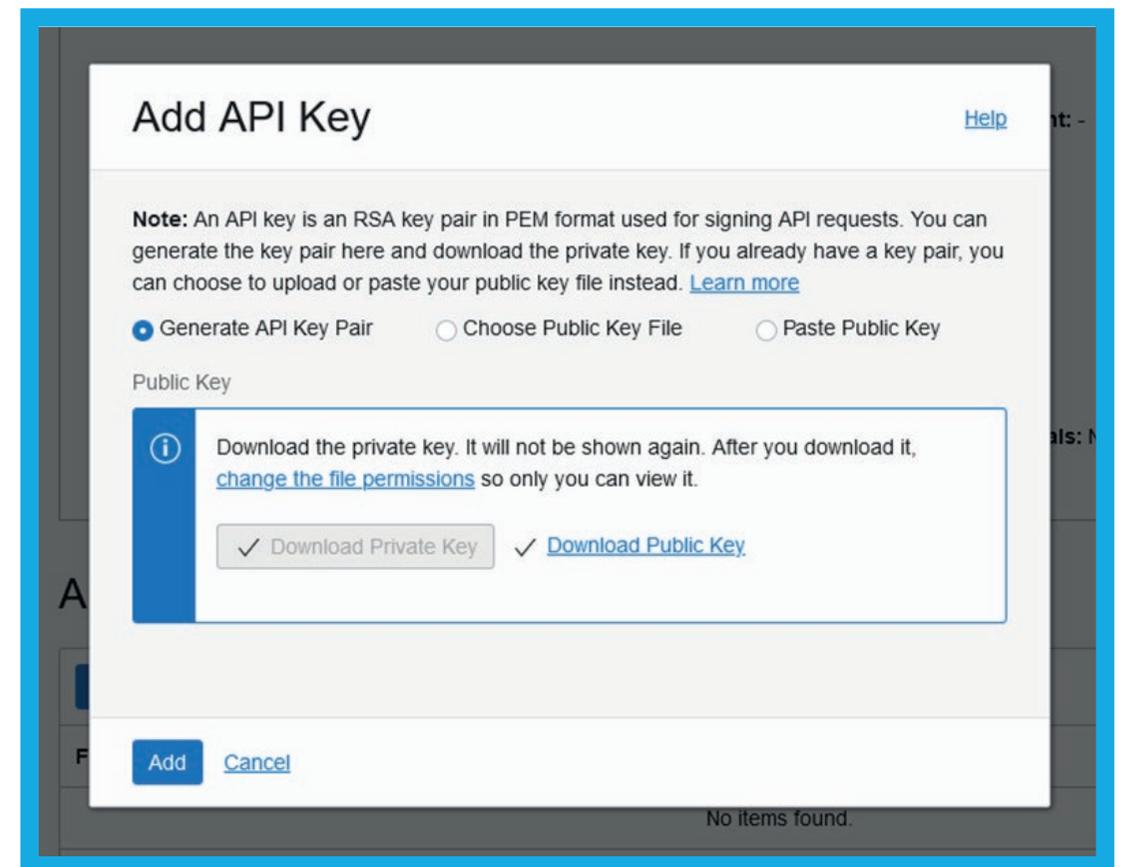


Figure 1: Generate the API key pair using the OCI console

You should also download the keys in PEM format. Also, gather the information (click the *View Configuration File* link on the

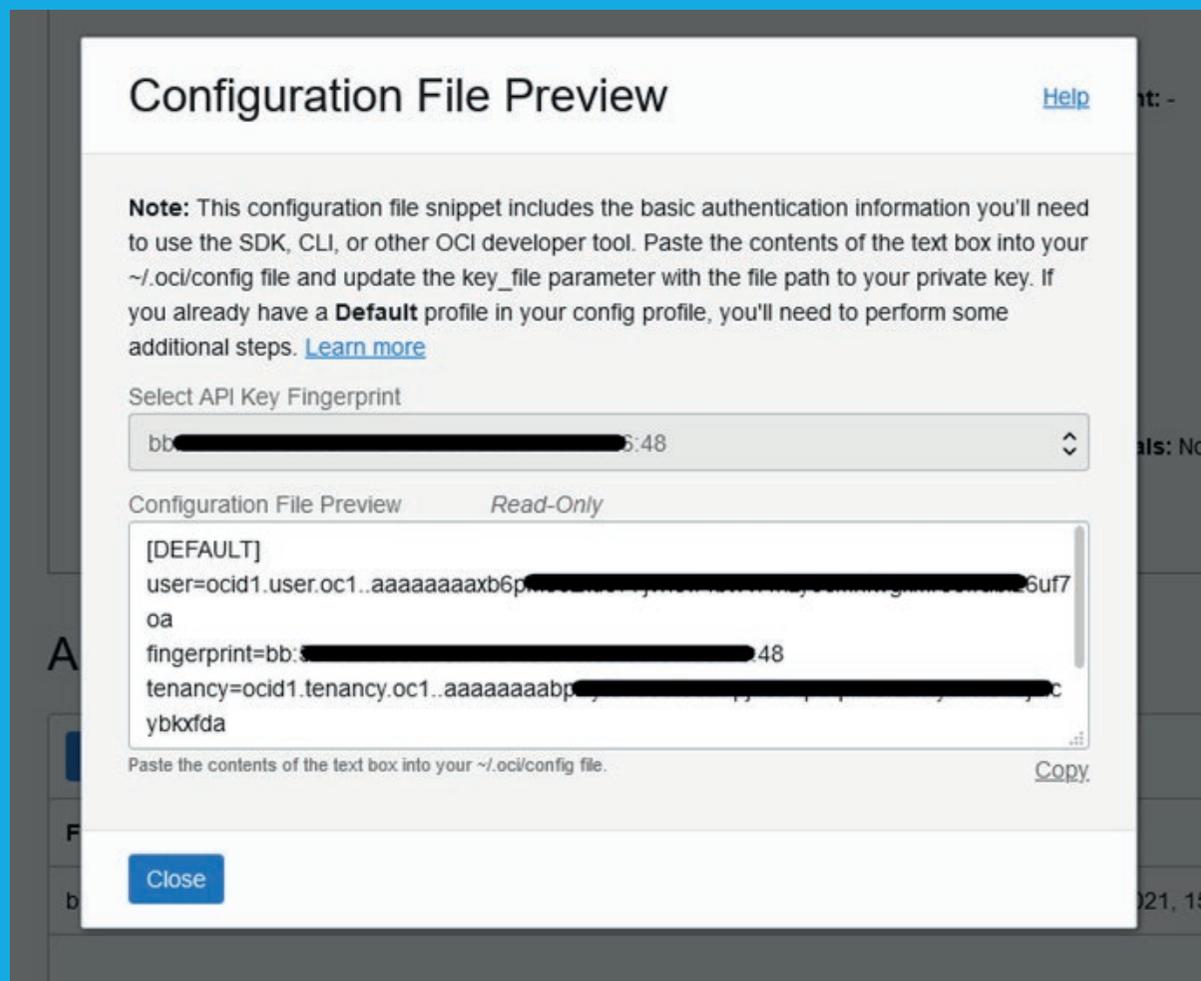


Figure 2: Collate information to be used to setup the APEX Web Credential.

user's details page) that you will need to setup the *APEX Web Credentials* for OCI (**Figure 2**), specifically:

- User OCID
- Tenant OCID; and
- The API key's fingerprint

Open the private key that you had downloaded from the OCI console and remove (a) the enclosing tags

```
-----BEGIN PRIVATE KEY-----
```

and

```
-----END PRIVATE KEY-----;
```

and (b) all newline characters. The result should be a long string of characters. This is the OCI *Private Key* that you will need to provide in the APEX Web Credentials. Copy the text, but **do not** save the PEM file.

Next, to use the Language Services, we will need to create an IAM Policy to allow users access. Create a policy with the following statement:

```
allow group AiLanguageUsers to use ai-service-language-family in tenancy
```

With these OCI resources created, you can setup the CLI, SDKs and APEX to call the Language service REST APIs.



## The APEX Application

The first task is to create the *Web Credentials* in APEX with the required information about your user and tenancy (*Figure 3*).

Create an APEX application with the following components (*Figure 4*):

- Page 1
  - Static Content region – Input
    - Text area page item P1\_INPUT
    - Button that triggers a dynamic action with the following actions:
      - Execute PL/SQL server-side code (*Listing 1*)
      - Refresh Analysis Report region
  - Classic Report region – Analysis Report
    - Code for query provided in *Listing 2*

When you are done, save and run the page. Enter some text into the Input field and then click the Analyze button. The response from the Language service will be parsed and displayed in the table below (*Figure 5*).

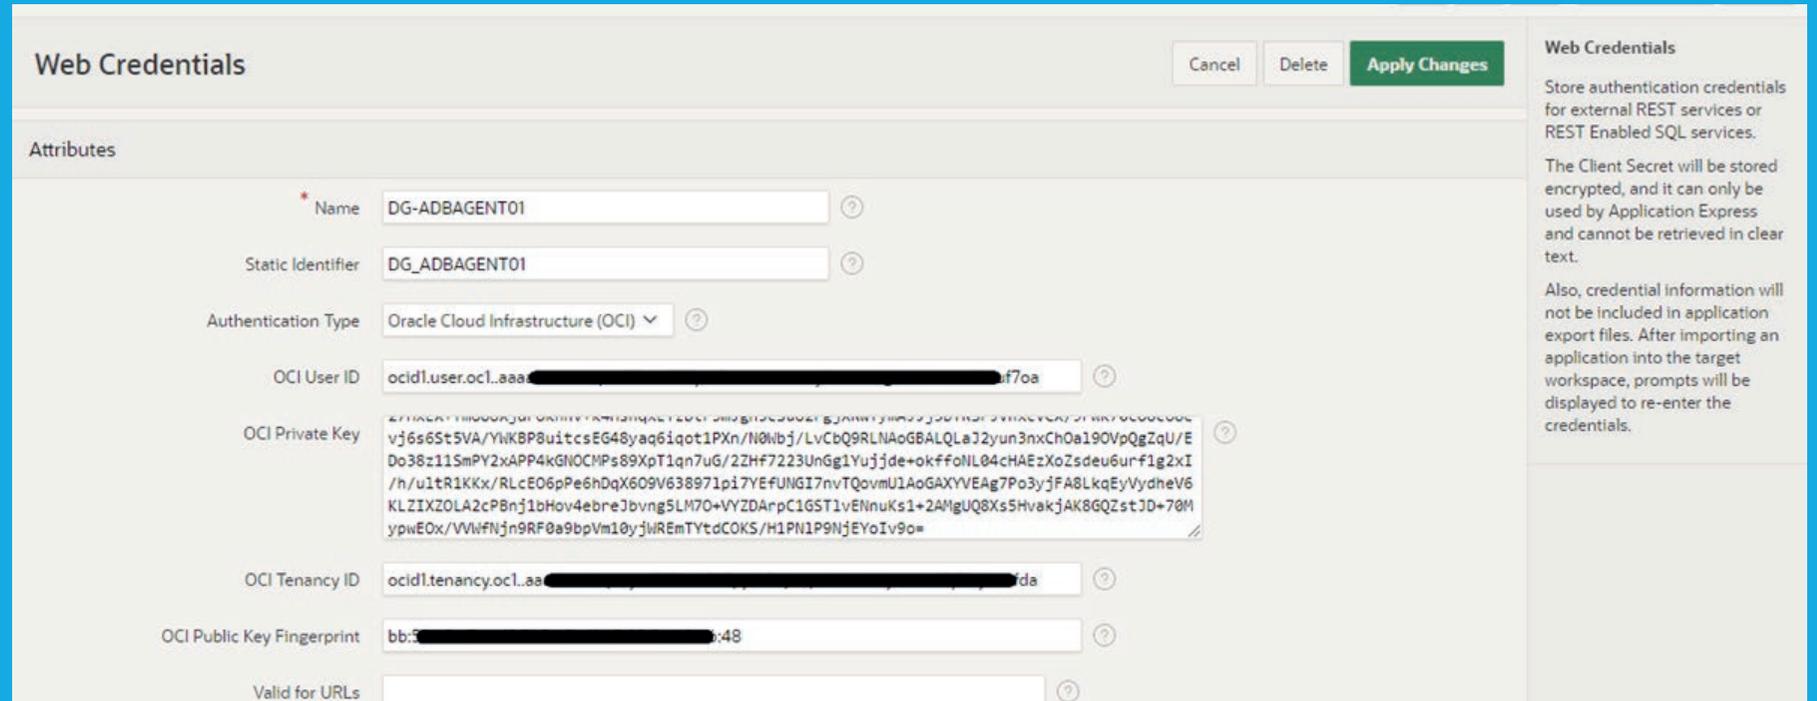


Figure 3: Create the Web Credential.

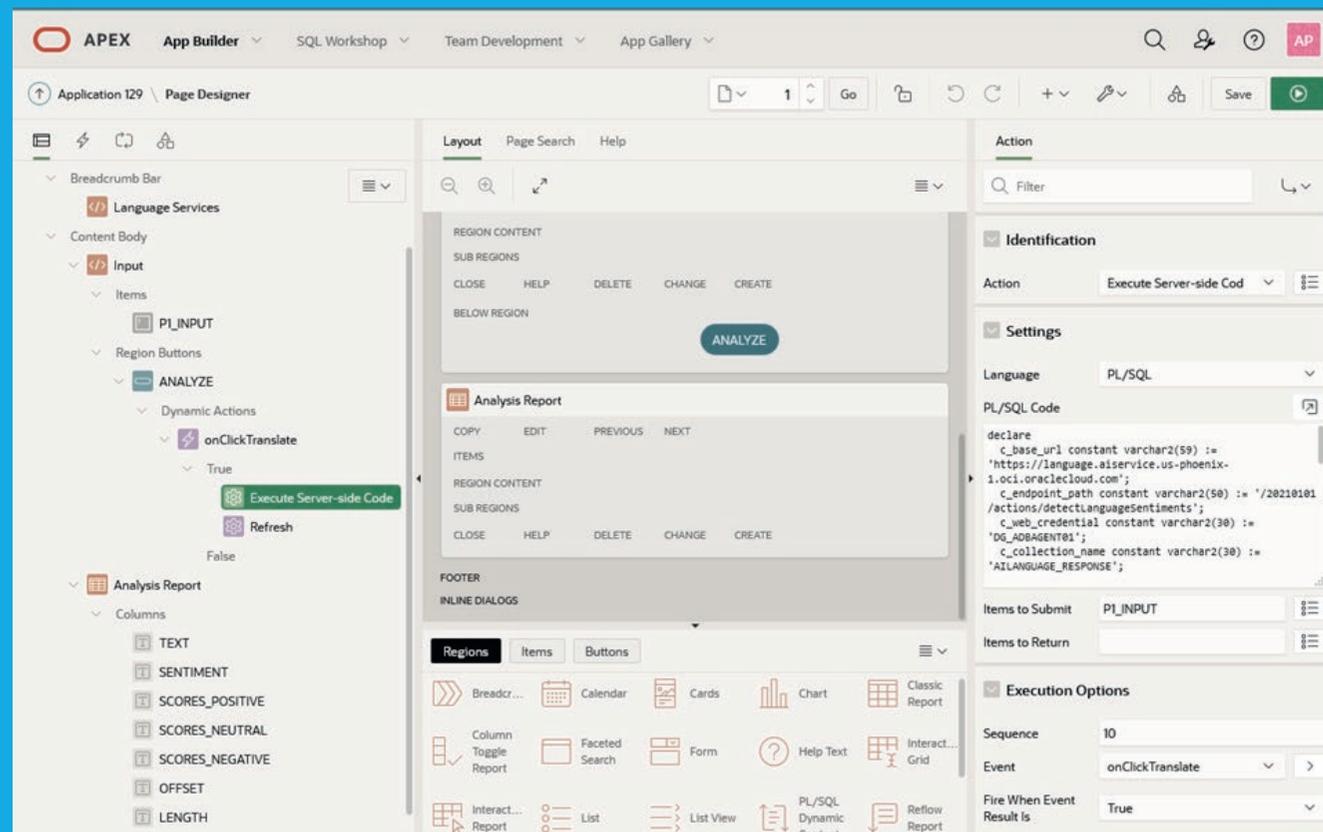


Figure 4: APEX Page and components.



```

declare
c_base_url constant varchar2(59) :=
'https://language.aiservice.us-phoenix-1.oci.oraclecloud.com';
c_endpoint_path constant varchar2(50) :=
'/20210101/actions/detectLanguageSentiments';
c_web_credential constant varchar2(30) := 'DG_ADBAGENT01';
c_collection_name constant varchar2(30) := 'AILANGUAGE_RESPONSE';

l_response clob;
begin
apex_collection.create_or_truncate_collection(c_collection_name);

apex_web_service.g_request_headers(1).name := 'Content-Type';
apex_web_service.g_request_headers(1).value := 'application/json';

l_response := apex_web_service.make_rest_request(
p_url => c_base_url || c_endpoint_path
, p_http_method => 'POST'
, p_body => '{ "text": "' || apex_escape.json(:P1_INPUT)
|| '"' }'
, p_credential_static_id => c_web_credential
);

apex_collection.add_member(
p_collection_name => c_collection_name
, p_clob001 => l_response
);
end;

```

Listing 1

```

with d as (
select clob001 as json_data
from apex_collections
where collection_name = 'AILANGUAGE_RESPONSE'
and seq_id = 1
)
select
x.text
, x.sentiment
, x.scores_positive
, x.scores_neutral
, x.scores_negative
, x.offset
, x.length
from d
, json_table(
json_data
, '$.aspects[*]'
columns(
offset number path '$.offset'
, length number path '$.length'
, text varchar2(32767) path '$.text'
, sentiment varchar2(32767) path '$.sentiment'
, scores_positive number path '$.scores.Positive'
, scores_negative number path '$.scores.Negative'
, scores_neutral number path '$.scores.Neutral'
)
) as x

```

Listing 2

Language Services

Input

OCI recently added new services to existing compliance program including SOC, HIPAA, and ISO to enable our customers to solve their use cases. We also released new white papers and guidance documents related to Object Storage, the Australian Prudential Regulation Authority (APRA), and the Central Bank of Brazil. These resources help regulated customers better understand how OCI supports their regional and industry-specific compliance requirements. Not only are we expanding our number of compliance offerings and regulatory alignments, we continue to add regions and services at a faster clip.

Analyze

Analysis Report

Text	Sentiment	Scores Positive	Scores Neutral	Scores Negative	Offset	Length
OCI	Positive	0.99	0.01	0.00	377	3
regions	Positive	0.99	0.01	0.00	559	7
resources	Positive	1.00	0.00	0.00	320	9
services	Positive	1.00	0.00	0.00	23	8

1 - 4

Release 1.0

Figure 5: Analyze the text and detect aspects and their sentiments.

## Costs

Finally, let us discuss the elephant in the room. How much will this cost? At the time of writing, I have not found a product page that provides pricing information. However, from the Oracle Cloud Cost Estimator (<https://www.oracle.com/cloud/cost-estimator.html>), it appears that the first 5,000 transactions are free (**Figure 6**), after which, it will cost USD 1 for every 5,000 transactions.

This information is provide solely based on observation. Please **do not** make any purchasing decisions based on the estimates in this article. For accurate pricing information, please check with your Oracle representative.

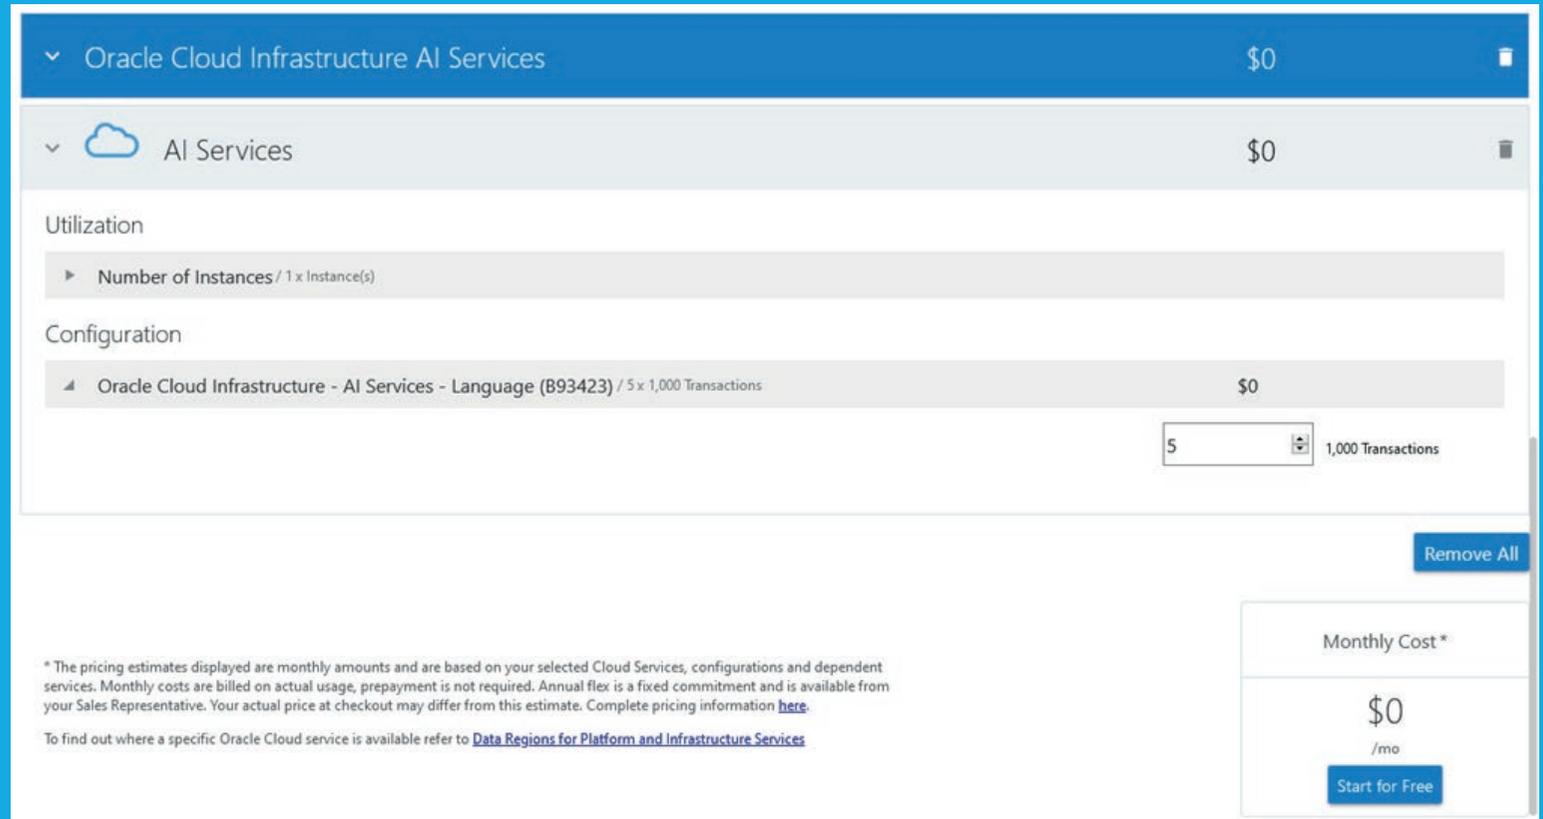


Figure 6: Language service pricing found on Oracle Cloud Cost Estimator.

## About Adrian Png

**Adrian** is a seasoned solutions architect with over 20 years of experience working with clients to design and implement state-of-the-art infrastructure and applications. He earned a Master of Technology (Knowledge Engineering) degree from the National University of Singapore and has applied his knowledge and skills in Artificial Intelligence and Machine Learning in practice, notably in his paper *Primer Design for Whole Genome Amplification using genetic algorithms* (In Silico Biol. 2006; 6(6): 505-14). Adrian is also trained and certified in several Oracle technologies including Oracle Cloud Infrastructure, Oracle Autonomous Database, Oracle cloud native services, Oracle Database and Oracle APEX. He is an Oracle ACE and is a recognized contributor to the Oracle community. Most recently, he co-authored the book *Getting Started with the Oracle Cloud Free Tier*, an indispensable reference for anyone just starting out with the Oracle Cloud and wished to get the most out of Oracle's cloud offerings.



A decorative graphic featuring a white cloud icon on the left. To its right, the word "CLOUD" is written in large, bold, white capital letters, with "Object Storage" written below it in a smaller, white, sans-serif font. The background consists of several parallel, slanted lines in shades of gray, blue, orange, and purple, creating a sense of depth and motion.

**CLOUD**  
Object Storage

*André Luiz Dutra Ontalba*

**Mounting an Oracle Cloud Object Storage**

**Bucket as a File System on Linux**

Object Storage allows you to securely store any type of data in its native format. The Object Storage is ideal for creating applications or provider solutions of storage, that need to have scale and flexibility, since we can consolidate several data sources for different purposes such as backup or archiving.

The Block Volume service enables users to dynamically provision and manage block storage volumes. You can create, attach, connect and move volumes, change the volume's performance, according to your needs.

The idea for this article arose from the need of presenting an Object Storage on Linux to perform a backup of an environment.

## Prerequisites

First, ensure that the bucket you're trying to mount is in the compartment listed for S3 compatibility; by default, it's a root compartment of the tenancy. If you need to change that, you will find the settings located under Administration->Tenancy Details->Edit Object Storage Setting.

Create an Object Storage Bucket from the OCI Console by clicking in Object Storage (see Figure 1), then enter a name for the bucket and click on Create Bucket (Figure 2).

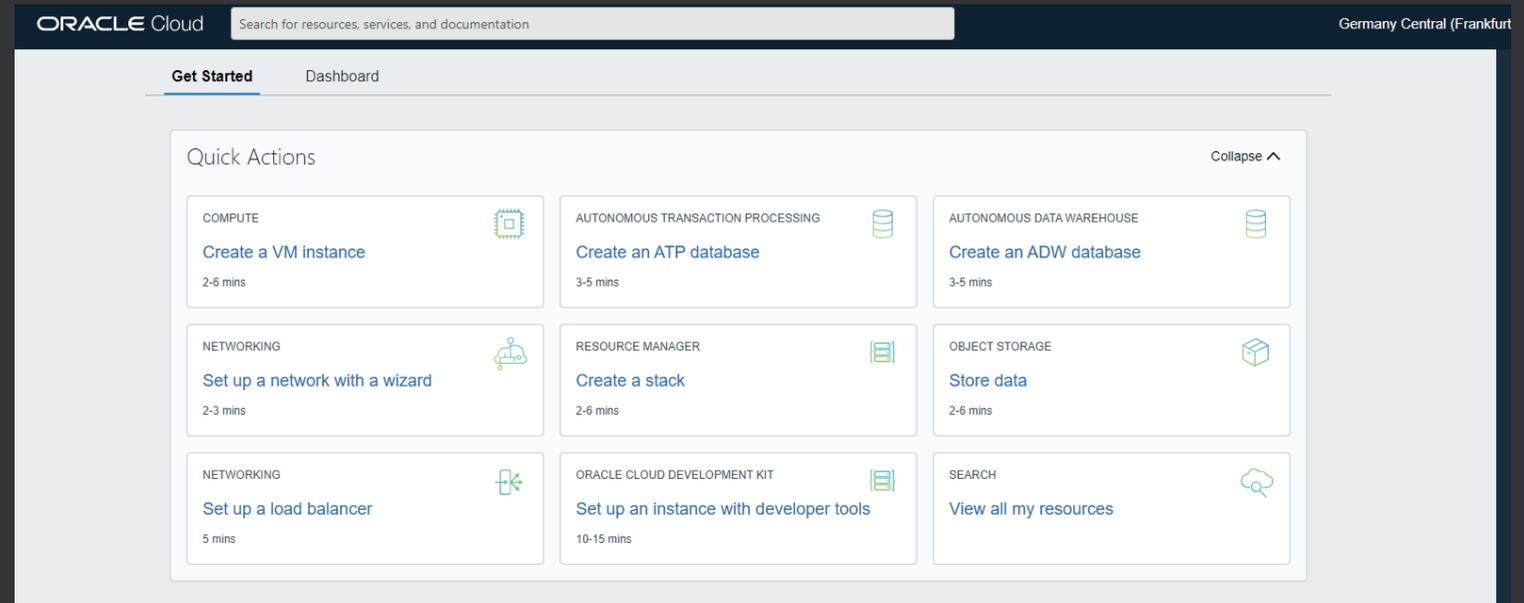


Figure 1: Click in Object Storage.

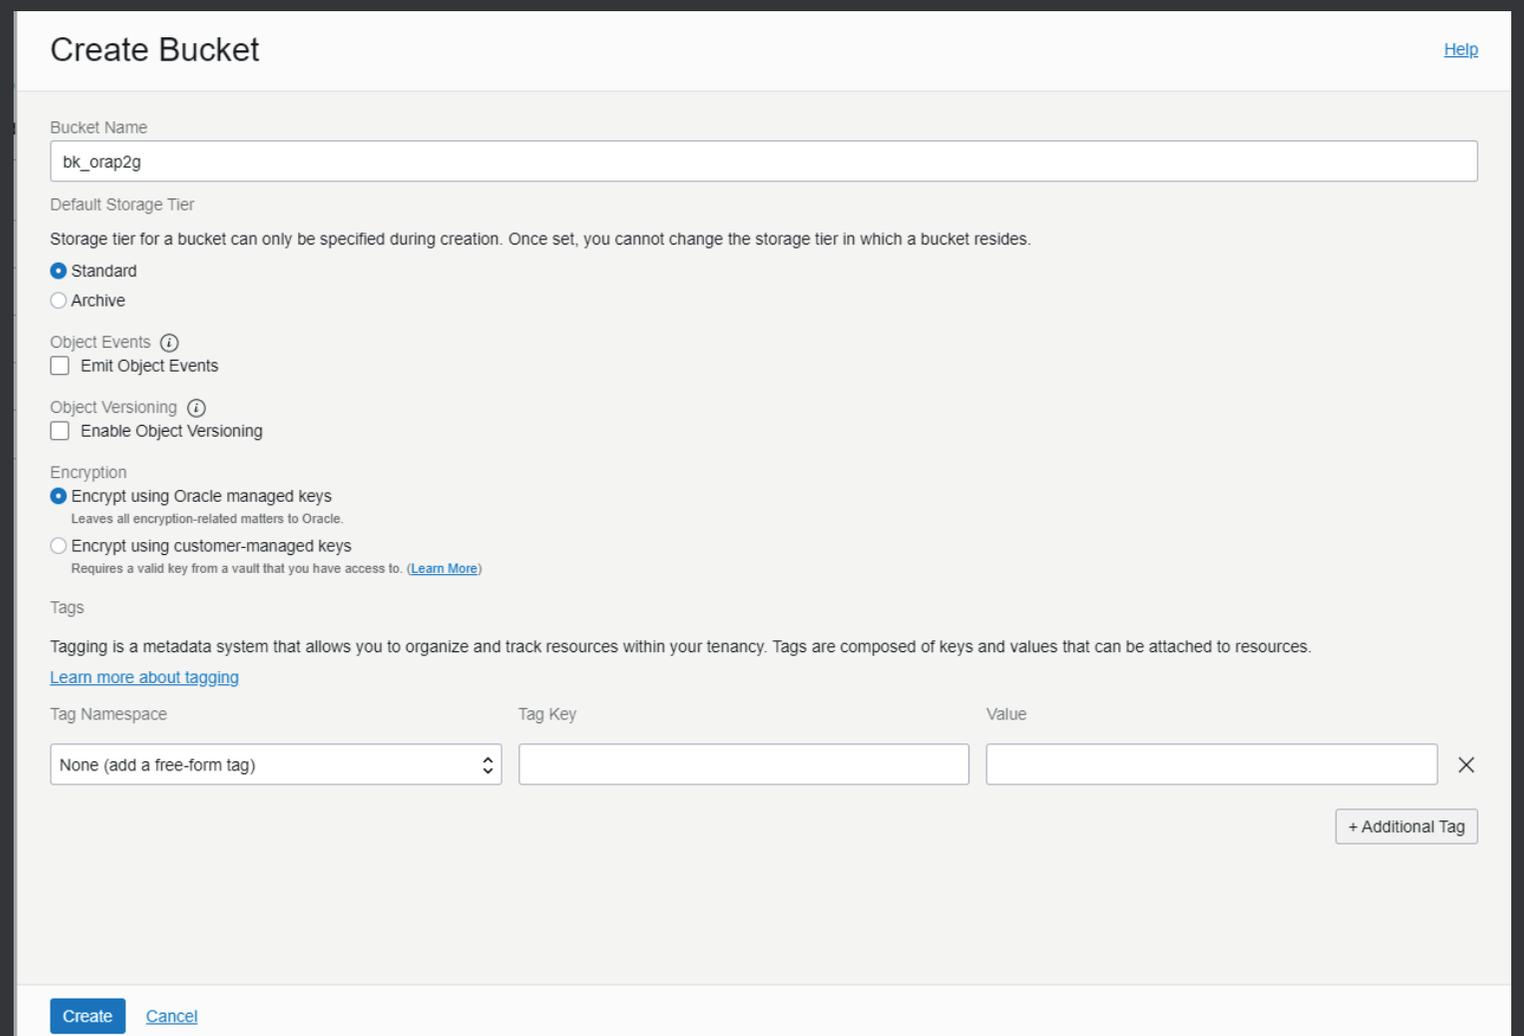
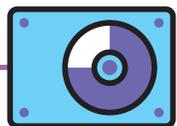
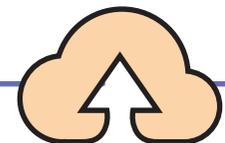


Figure 2: Create a Bucket.





## Install s3fs-fuse

You can install s3fs-fuse by using a prebuilt package from the Oracle Linux EPEL repository. For this example, I'm using the binary RPM.

Connect with opc user. Then execute this command to install the s3fs-fuse:

```
sudo yum install s3fs-fuse
```

## Configure Credentials

In the Oracle Cloud Infrastructure Console, click the Profile icon in the top-right corner, and select **User Settings**. Then click **Customer Secret Keys**, and after that click **Generate Secret Key** and give the key a meaningful name, for example **bk\_orap2g** (*Figure 3*).

Copy and save the secret key because it won't be shown again. The S3 credentials are created by using an access key and the secret key. The access key is displayed in the Customer Secret Keys area of the Console.

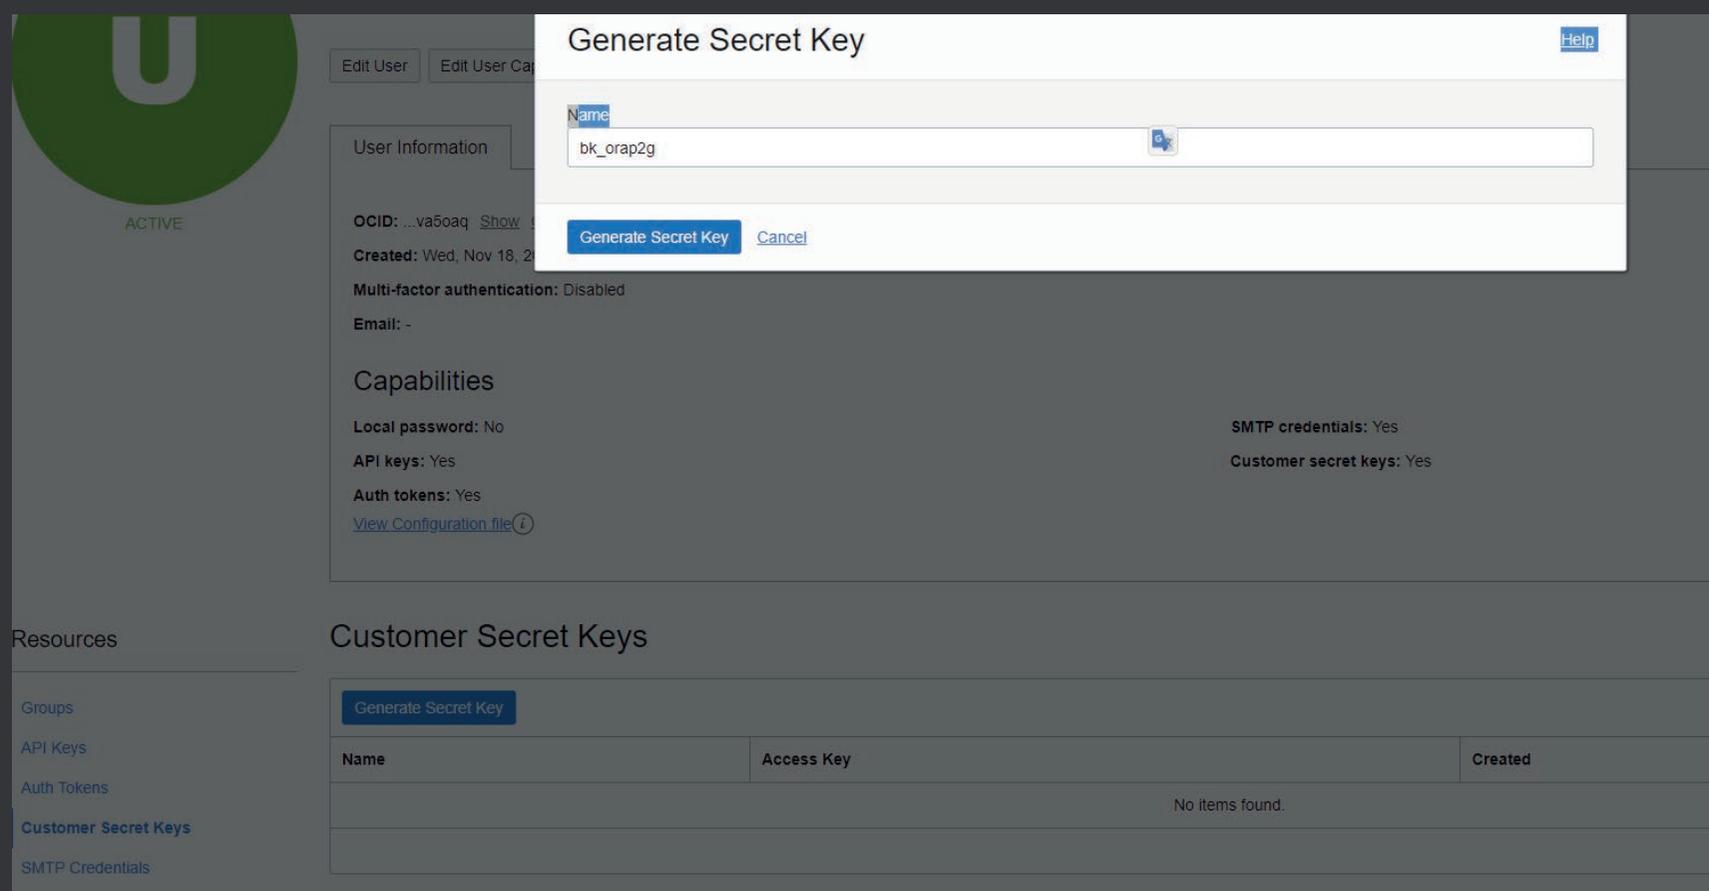


Figure 3: Generating the Secret Key.



```
[root@ora2pg ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        16G   0    16G   0% /dev
tmpfs           16G   8.0K  16G   1% /dev/shm
tmpfs           16G  681M  15G   5% /run
tmpfs           16G   0    16G   0% /sys/fs/cgroup
/dev/sda3       39G   11G   28G  29% /
/dev/sda1       200M  9.7M  191M   5% /boot/efi
/dev/sda4       201G  133G   58G  70% /u01
tmpfs           3.1G   0    3.1G   0% /run/user/0
tmpfs           3.1G   0    3.1G   0% /run/user/993
tmpfs           3.1G   0    3.1G   0% /run/user/1000
s3fs            39G   11G   28G  29% /home/oracle/bkp_db
[root@ora2pg ~]#
```

Figure 4: Mount point validation

**Generate Secret Key:** bk\_orap2g

**Generated Key:** xxxxxxxxxxxxN5ujFRYICdG5RQbiuyxi07+PWrLdE=

**Access key:** cc3e970b898676154xxxxxxxxxxxxxxxxxxxxxxxx

**Create the directory and assign ownership:**

```
# mkdir /home/oracle/bkp_db
# chown oracle:oinstall /home/oracle/bkp_db
```

**Mount the File System:**

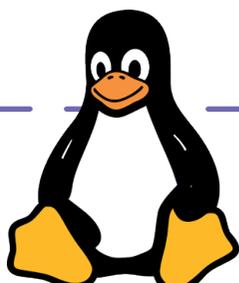
```
$ s3fs bucket2fs /home/oracle/bkp_db -o passwd_file==${HOME}/.pass-
word-bk -o url=https://mynamespace.compat.objectstorage..oracle-
cloud.com/ -o nomultipart -o use_path_request_style
```

**Add an entry in the /etc/fstab file (see also Figure 4):**

```
#s3fs
bk_orap2g /home/oracle/bkp_db fuse.s3fs _netdev,allow_other,no-
multipart,use_path_request_style,passwd_file=/home/oracle/.pass-
word-bk,url=https://mynamespace.bej2.objectstorage.eu-frankfurt-1.
oraclecloud.com/ 0 0
```

**If you get the error “fuse: failed to exec fusermount: Permission denied”, change the permission of /usr/bin/fusermount:**

```
[oracle@ora2pg ~]$ ls -l /usr/bin/fusermount
-rwsr-x---. 1 root fuse 32584 Feb 25 2021 /usr/bin/fusermount
```



Run the following command, and then try again:

```
[oracle@ora2pg ~]# sudo chmod +x /usr/bin/fusermount  
[oracle@ora2pg ~]$ ls -l /usr/bin/fusermount  
-rwsr-x--x. 1 root fuse 32584 Feb 25 2021 /usr/bin/fusermount
```

If there's an issue, add the Debug parameter to help troubleshoot:

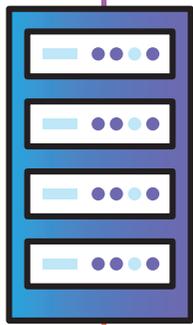
```
[oracle@ora2pg ~]$ s3fs bucket-1 /test_bucket -o passwd_  
file=${HOME}/.passwd-s3fs -o dbglevel=info -f -o curldb=
```

I hope this helps you!



## About André Luiz Dutra Ontalba

**André**, graduated in Computer Science, specializes in Oracle Database with solid knowledge in Engineered Systems, Performance & Tuning, RAC, Oracle Cloud and Oracle ERP's System; he has worked with Oracle for 20 years, Oracle OCP 10g/11/12g/2019 Oracle Cloud Professional, AWS Professional, Azure Expert, and has more than 40 other Oracle product certifications. He currently works as a Cloud Architect Team Lead at Digora. André is also a founder of the Luxembourg Oracle User Group.





# MASH PROGRAM

**MASH -**

*Kim Berg Hansen*

**Mentor and Speaker Hub**



***Being a speaker in the tech community can be a very rewarding experience. The first step can be a giant one, though, and can often seem overwhelming. A mentor can help break the start down into manageable pieces.***

Sharing what we know makes us all much better at what we do than if we all sit alone. Sharing knowledge goes in all directions, not just from expert to novice. You – yes, you – have unique knowledge gained from your unique experiences. Even an expert may gain insight from that.

We're a **small team** of enthusiastic community members that love to share, being active in the community as speakers, bloggers, user group members and whatever else we can think of.

## **MASH Program**

To help others starting out on their sharing journey as tech speakers, we have founded the MASH Program.

MASH stands for **Mentor And Speaker Hub**.

We:

- facilitate contact between new tech speakers and experienced mentors.
- guide the mentoring process with our 10-step how-to program.
- provide a stepping stone into the global community of speakers.
- help you grow into your new role as speaker.



We'd love to welcome more speakers to the rewarding experience of community involvement, and we think we can help you getting started by putting you in touch with a mentor that can guide your first steps.

## **The rewarding experience**

Facilitation of knowledge sharing in our tech world is often in communities – formal or informal. In the Oracle world many countries and states have Oracle User Groups (OUGs) dedicated to serving the community.

OUGs and other networking groups arrange conferences and webinars where some of us do presentations about what we have learned from working with our tech, so that others may learn more than what they could have learned alone. Everybody has something they can share to the benefit of others – there's only a difference in degree that the most experienced have most to share, but we all learn from all.

You can be an active part in your community without speaking, but becoming a speaker adds an extra dimension to your network. Sharing your knowledge will add a lot to your own knowledge – researching your chosen topic makes you dive deeper into the topic, but also attendees to your session will often chime in with details you haven't found yourself.



There are a lot of good arguments for you to become a tech speaker – even if you think you are not an expert, you most certainly are in some topic that will help others. And nothing beats having someone tell you after the session that it has been very helpful and educational.

## **An overwhelming first step**

Perhaps you're considering doing a presentation, but it just seems like there's so much to take into consideration. Where to



speaking, what to speak about, how to build a good presentation, how to present, and so on and so on.

It's just too easy to become a victim of imposter syndrome and feel you are not nearly good enough to overcome these obstacles. And you think that the audience wants to hear acknowledged experts and not your inadequate (according to yourself) knowledge.

On top of that it can be scary to think about presenting to an audience of strangers. Even actors can get stage fright, so it's not a surprise that we (people mostly used to sitting by our trusted computer) also get nervous on stage.

## Many smaller first steps

But in our daily life, most of us are used to breaking out a task into smaller sub-tasks. For developers it's called modularization, but it applies to a lot of what we do. You can apply the principle to speaking as well, for example splitting it up in these smaller steps:

- Choosing a Topic
- Choosing a Title
- Writing an Abstract
- Submitting the Abstract
- Building a Session
- Creating the Presentation
- Practicing the Session
- How to Present
- Promoting the Session
- What Else To Do



Taken one at a time, it is no longer so daunting. When you gain practice, it becomes easier of course – but until you gain your own practice, it can help to utilize the practice of someone else.



## Mentoring

When you start out as a speaker, getting some good tips, helpful advice, and constructive criticism can make a real difference for achieving sufficient success at your first attempt to make you want to do it again.



This is where you can make good use of a mentor with speaking experience. The mentor can assist your entry on the speaking stage, helping you to avoid some beginner mistakes (don't worry too much, we all make mistakes) and gain some confidence before you meet an audience.

## Your local User Group

An OUG near you can quite possibly be a good place for a first presentation. (Depending on topics, specialized events might be more suitable – discuss it with your mentor.)



The user groups are a big part of the sharing community – community members involve themselves in the OUGs and the OUGs support the sharing by community members.

Many OUGs arrange events like meetups and conferences. The best such events have a mix of well-known speakers and new speakers with fresh perspectives.

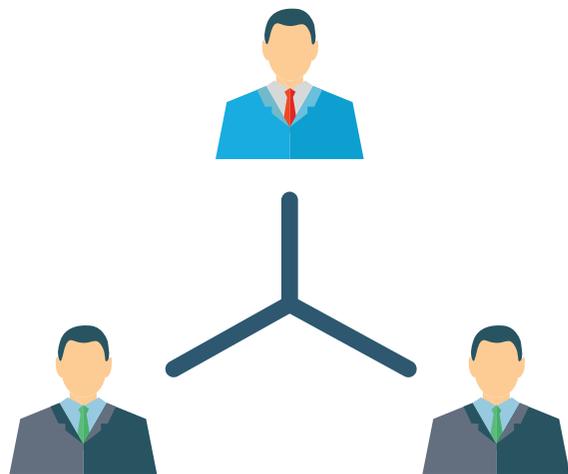
Your OUG can help this process by referring to the MASH website in their call-for-papers and newsletters, showing potential new speakers where to find a mentor to help make a quality submission to their CFP.

Also the OUG can reserve a slot in their conference agenda for a new speaker, giving them a steady influx of fresh and interesting insights to attract the audience.

### Go for it!

Ready to give it try? Get in touch with MASH and we can help you start on your journey as a conference speaker.

Read more on: <https://mashprogram.wordpress.com/>



## About Kim Berg Hansen

**Kim** is a database developer from Middelfart in Denmark. Originally wanting to work with electronics, he almost coincidentally tried computer programming and discovered where his talent lay. He has worked extensively with Oracle SQL and PL/SQL since 2000, utilizing the SQL language to the fullest.

Kim is an Oracle ACE Director, the author of the Practical Oracle SQL book, shares knowledge by blogging, presenting at conferences, and being the SQL quizmaster at the Oracle Dev Gym. Outside the coding world, Kim is married, loves to cook, and is a card-carrying member of Danish Beer Enthusiasts Association.



Lisa Damerow

# A Look Back at APEX connect 2021 [online]

The logo features a large, stylized 'A' shape composed of several overlapping geometric elements. A prominent blue chevron points upwards from the bottom left towards the top right. To its right, a grey chevron points downwards. A horizontal grey bar intersects the center of the 'A'. The text 'APEX connect' is written in white, with 'APEX' in a bold, sans-serif font and 'connect' in a lighter, lowercase sans-serif font. Below 'connect', the text 'by DOAG' is written in a smaller, lowercase sans-serif font. The background of the entire graphic is light grey, decorated with various geometric patterns: a dotted triangle at the top left, a striped triangle at the bottom right, and several smaller triangles and lines scattered throughout.

**APEX** connect  
by DOAG

*APEX connect 2021 welcomed around 320 participants from all over the world for the second time as a virtual event on May 4 and 5, 2021.*

Last year, we as organizers from DOAG gained some experience and knowledge regarding online events: In addition to the web conference platform, the tool Gather.town was available to visitors this year. It allowed participants to walk through the virtual image of Phantasialand with their own avatar and exchange with others using microphone and webcam. The two sponsors of the event, MT AG and Database Consult, each had a small virtual booth in Gather, where staff welcomed visitors to exchange information.

Gather.town was utilized thoroughly throughout **APEX connect 2021**! On the evening of the first conference day, a moderated quiz show provided great entertainment. Here, the participants could put their knowledge to the test and compete against each other. The moderator asked, among other things, how much a full-grown blue whale weighs or what “Twitter” translates to. From matching musicians to given song lyrics to Jeopardy and useless knowledge, everyone was involved.

Afterwards, the avatars of the visitors and speakers danced polonaise through the virtual Phantasialand and formed the name of their favorite tool together: APEX. The night owls among the participants concluded the evening at the virtual APEX bar.

To kick off the event, Niels de Bruijn, DOAG Board Member and Head of the DOAG Development Community, officially welcomed the participants to APEX connect 2021. Immediately afterwards, the participants received information directly from the source! Thus, Carsten Czarski and Patrick Wolf from the Oracle APEX Development Team presented the current version 21.1 of Application Express in detail in their keynote and gave an outlook into the future. More than 30 presentations on the

topics APEX, JavaScript and PL/SQL awaited the participants on the two conference days. Part of the presentation program for the first time was a “Good Practices Series” aimed at beginners and advanced users. In these virtual workshops, speakers provided interested participants with basics as well as advanced knowledge on a wide range of topics. In the second day’s keynote, Steven Feuerstein took participants behind the scenes of the “Evry” app developed with APEX. It allows its users to make one or more personal contributions to the environment and sustainability every day.

So that APEX fans would not have to split into two, all presentations were recorded and made available to participants on our website afterwards. This allows people to review what they have already learned and catch up on presentations they may have missed.

A big thank you to all participants, speakers, sponsors, and contributors!



# Call for Papers

## MakeIT 2021

until June 27, 2021  
Portorož, Slovenia + online  
<https://www.makeit.si/index.php/en/>

## DOAG 2021 Conference + Exhibition

until July 7, 2021  
online  
<https://2021.doag.org/en/home/>

## HrOUG2021

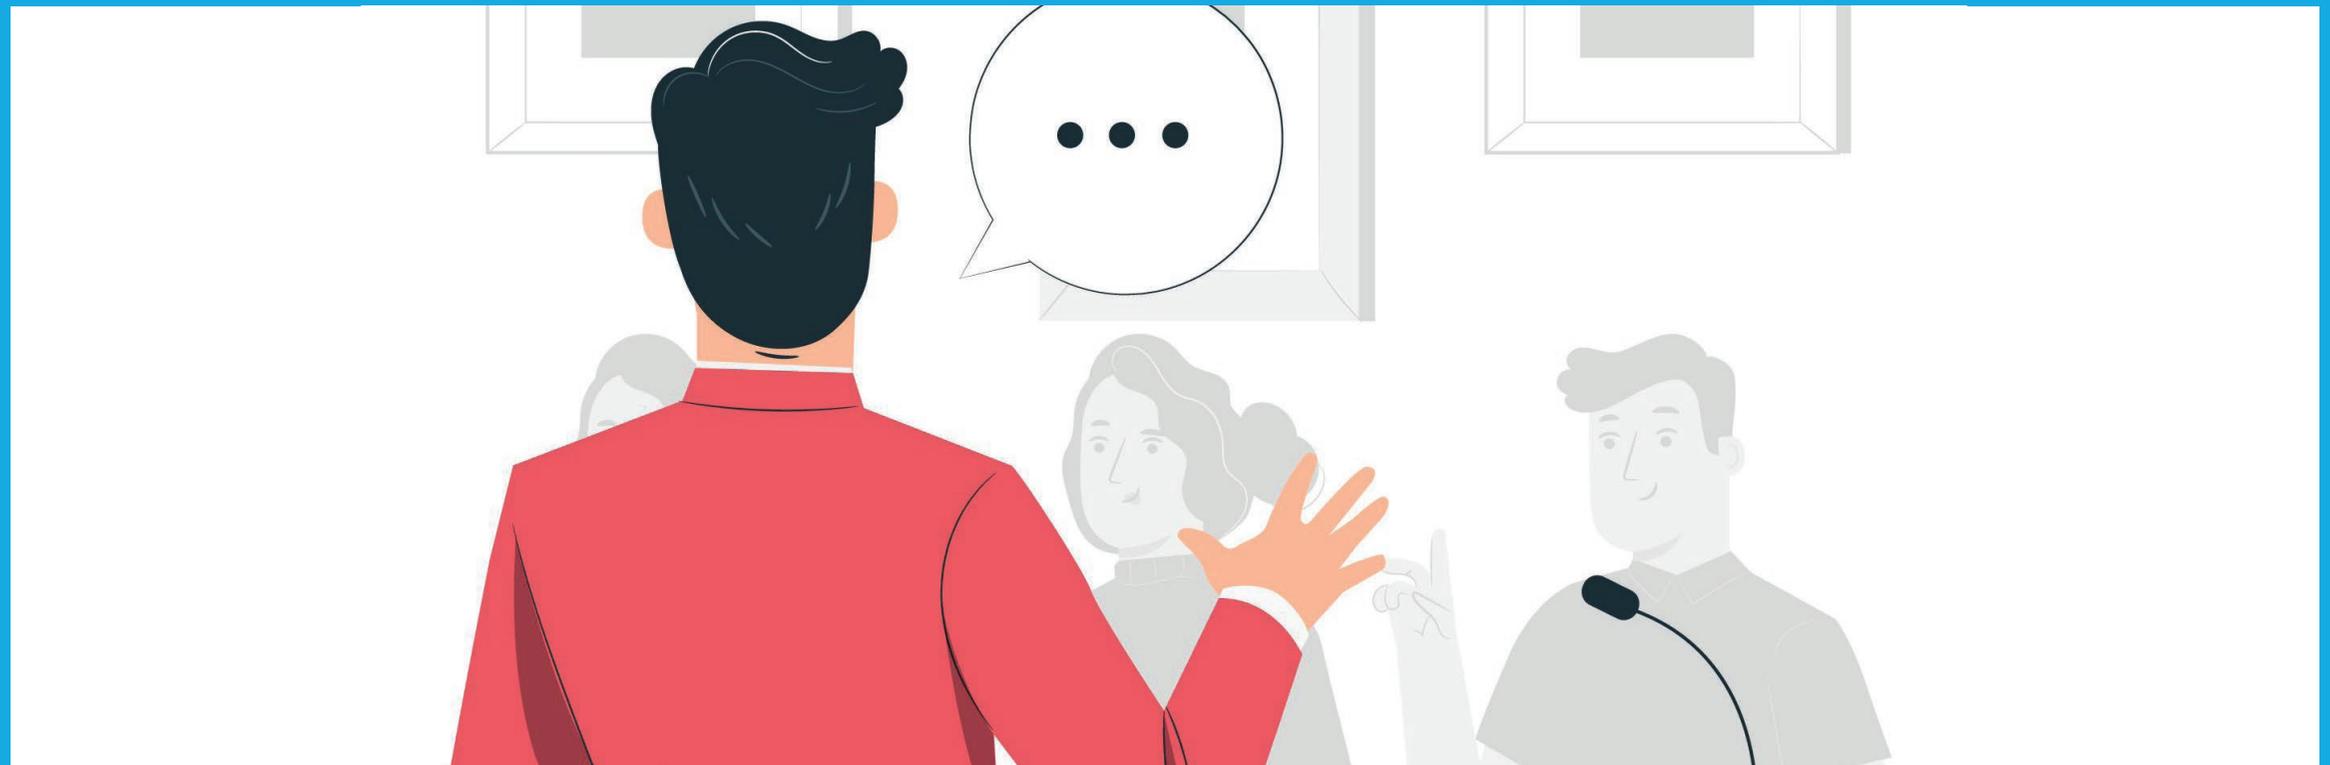
until July 15, 2021  
Rovinj, Croatia  
<https://2021.hroug.hr/eng>

## Oracle Groundbreakers Tour EMEA 2021

until August 15, 2021  
online  
<http://ogbemea.com/>

## UKOUG 2021 | Together

until August 18, 2021  
The Oval, London, UK  
<https://ukoug.org/page/UKOUG2021>



# Events

## Ascend 2021

August 15 - 18, 2021  
Hollywood, Florida, US + online  
<https://ascendusersconference.com/>

## POUG 2021

September 10 + 11, 2021  
Warsaw, Poland + online  
<https://poug.org/en/>

## APEX World 2021

September 29 + 30, 2021  
Zeist, Netherlands  
<https://www.nloug.nl/events/apex-world-2021/>

## MakeIT 2021

October 11 + 12, 2021  
Portorož, Slovenia + online  
<https://www.makeit.si/index.php/en/>

## HrOUG2021

October 12 - 16, 2021  
Rovinj, Croatia  
<https://2021.hroug.hr/eng>

## Oracle Groundbreakers Tour EMEA 2021 (online)

October 18 - 29, 2021  
online  
<http://ogbemea.com/>

## DOAG 2021 Conference + Exhibition

November 16 - 18, 2021  
online  
<https://2021.doag.org/en/home/>

## UKOUG 2021 | Together

November 29 + 30, 2021  
The Oval, London, UK  
<https://ukoug.org/page/UKOUG2021>



## ORAWORLD is a publication of the EOUC — EMEA ORACLE USERGROUP COMMUNITY

### The following user groups belong to EOUC:

Angola Oracle User Group, Oracle User Group Armenia, Austrian Oracle User Group, Azerbaijan Oracle User Group, Bulgarian Association of Software Developer, Bulgarian Oracle User Group, Hrvatska udruuga Oracle korisnika, Czech Oracle Applications User Group, Danish Oracle User Group, Egypt Oracle Users Group, Oracle User Group Estonia, Oracle User Group Finland, Club Français des Utilisateurs JD Edwards, Association des Utilisateurs Francophones d'Oracle, Club des Utilisateurs PeopleSoft, Oracle User Group Georgia, Deutsche Oracle Anwendergruppe, PeopleSoft Germany, Hungarian Oracle User Group, Israel Oracle User Group, Taranta Valley Oracle User Group, Italian Oracle User Group, Jordan Amman Oracle User Group, Latvian Oracle Users Group, Lithuanian Oracle Users Group, Mauritius Oracle User Group, Oracle Gebruikersclub Holland, Oracle Benelux User Group, Oracle User Group Norway, Polish Oracle Users Group, Oracle Users Group Portugal, Romanian Oracle User Group, Russian Oracle User Group, EBS Finance Special Interest Group Russia, Arab Oracle User Group, Serbia and Montenegro, Slovenian Oracle User Group, South African Oracle User Group, Spanish Oracle User Group, Swedish Oracle User Group, Swiss Oracle User Group, Tajikistan Oracle User Group, Turkey Oracle Users Group, Ukraine Oracle User Group, Middle East Oracle User Group, United Kingdom Oracle User Group, Zimbabwe Oracle User Group.

### Editorial board:

**Registered office:** DOAG Dienstleistungen GmbH  
Tempelhofer Weg 64, 12347 Berlin, Germany

**www.doag.org,**

Director Fried Saacke,  
AG Berlin Charlottenburg HRB 95694B,  
VAT ID DE240700058

**Contact:** redaktion@doag.org

**Editor-in-chief (ViSdP):** Dr. Dietmar Neugebauer

### Editorial team:

Mirela Ardelean, Jean-Jacques Camps, Heli Helskyaho, Dr. Dietmar Neugebauer,  
Ann-Sofie Vikström Often, Andrejs Vorobjovs

Other editors: Lisa Damerow, Marcos López,  
Christian Luda, Martin Meyer

Other authors are credited by name with their article. They are subject to the terms and conditions for authors: [www.oraworld.org/terms-and-conditions-for-authors](http://www.oraworld.org/terms-and-conditions-for-authors)

### Graphic design:

Alexander Kermas, DOAG Dienstleistungen GmbH,  
Tempelhofer Weg 64, 12347 Berlin, Germany

### Disclaimer:

All rights reserved. Duplication or retransmission in whatever form or whatever medium either in whole or in part requires written permission, to the extent the content is not made available for duplication or retransmission.

The information in this publication has been duly researched and is correct to the best of our knowledge and belief. The use of this information is at your own risk. No liability for the accuracy of the information is accepted and, in particular, for its practical application in individual cases. Opinions represent the views of the individual author and do not necessarily represent the view of the publisher.

The ORAWORLD e-magazine reports on events in the Oracle and IT world. The publication covers current topics in the international user group network, as well as news items on products and technologies and their use. The purpose of the magazine is to foster the sharing of knowledge and experience among readers. ORAWORLD is independent of Oracle and does not represent its commercial interests either directly or indirectly.

ORAWORLD is published by DOAG Dienstleistungen GmbH, Tempelhofer Weg 64, 12347 Berlin, Germany, legally represented by director Fried Saacke, the nature and purpose of whose business is management of the group, organising events and publishing.

DOAG Deutsche Oracle Anwendergruppe e.V. holds 100 percent of the capital invested in DOAG Dienstleistungen GmbH. DOAG Deutsche Oracle Anwendergruppe e.V. is legally represented by the managing committee; Chair: Björn Bröhl.

### Article submission:

If you are interested in submitting an article, please e-mail us your article via the online form at [www.oraworld.org](http://www.oraworld.org). Deadline for the next edition: **July 28, 2021**.

### Photo credits:

Title: © Aleksandr Baidin | <https://123rf.com>

P. 4: © floralset | <https://123rf.com>

P. 5: © Comic | [www.commitstrip.com](http://www.commitstrip.com)

P. 6: © freepik | <https://freepik.com>

P. 7: © macrovector | <https://freepik.com>

P. 8 - 11: © created by freepik | <https://freepik.com>

P. 22: © kpdmedia | <https://123rf.com>

P. 26: © Anastasiia Nevestenko | <https://123rf.com>

P. 32: © emuemu | <https://stock.adobe.com>

P. 39: © prabowo | <https://stock.adobe.com>

P. 40 - 43: © sentavio | <https://freepik.com>

P. 42, Linux: | <https://wikipedia.com>

P. 44: © Spencer Garner | <https://pixabay.com>

P. 45 - 74: © macrovector | <https://freepik.com>

P. 48 - 49: © phochi | <https://istockphoto.com>

P. 50: © freepik | <https://freepik.com>

P. 51: © DOAG | [www.doag.org/en/home](http://www.doag.org/en/home)