September 2021, Edition #26

ORAWORLD

e-Magazine for Oracle Users published by the EOUC

Machine Learning
For Beginners

What's your Passion with Jože Senegačnik

page 6

Capitalizing on Oracle 19c Fast Ingest & Fast Lookup Features

page 20

Location Matters – Introduction to the APEX Map Region

page 30

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

Editorial

Dear ORAWORLD Readers,

It is a great pleasure and honor for me to welcome you to the first edition of ORAWORLD for which I have acted as editor-in-chief.

Firstly, thanks are due to my predecessor Dr. Dietmar Neugebauer, former DOAG President, who has done a fantastic job over the past 5 years. Secondly, thanks also go to Dr. Dietmar Neugebauer, former President of DOAG, for his confidence in me. I will do my best to continue the amazing work started by him: strengthening the community.

We are living in strange times. Our personal life has changed during the last one and a half year. For many of us, our professional environment has changed. The changes we faced are not limited to the individual level, they also reach a macro-dimension: the Oracle ACE Program is changing, we don't know exactly how, we don't know exactly when. The Oracle User Group Community program does not exist anymore right now, we don’t know if it’s only in a changing phase or if it’s gone.

But none of this stops us from keeping the community alive. User groups are refocusing, redefining themselves and continuing their mission to facilitate information sharing among community members.

Five very brave user groups – POUG, nlOUG, HROUG, SPOUG and UKOUG – are re-organizing in-person/hybrid conferences this fall. Beside of this, October will bring an 11-days webinar-marathon, coordinated by EOUC: the OGBEMEA 2021 virtual tour.

November will bring the DOAG conference and the EOUC leaders’ online summit. In between these major events, there are the local online meetings and there is our latest edition of ORAWORLD, both of which facilitate collaboration between members and knowledge sharing.

Please share ORAWORLD Magazine and our website and invite your colleagues to subscribe for free.

Thank you, our EOUC readers, for your ongoing contribution, for sharing your articles, thoughts and ideas with the user group community. We all hope you will find this edition of our eMagazine useful and enjoyable and we encourage you to send us your thoughts or any other comments you might have.

I am looking forward for your submissions and feedback!

Yours,
Mirela Ardelean,
RoOUG Vice president
EOUC Board member

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**.

„It haunts us“

www.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***

Interview by Christian Luda

What's Your Passion?
Jože Senegačnik about his Love for Baking and Flying

In our series "What's Your Passion?" we interview known people from the Oracle community about what drives them outside of work. For the third part of the series, we had the pleasure to talk to Jože Senegačnik from Slovenia who – besides being an Oracle ACE Director and OakTable member as well as a recognized conference speaker – is a very passionate baker and pilot.

Jože, when did your passion for baking start?
Around 18 years ago when we moved to a new house, and I bought a baking machine. At the same time my neighbor got a pizza oven in his cellar. So, I wanted to have one as well and then actually built it myself – of course not from scratch: I used some pre-fabricated parts for the inside and did the outside and the construction myself. Then, I met somebody who told me that his mother had been baking for almost 60 years and was making bread every Saturday in a big tiled stove. So, I asked him to arrange a training with his mother. It was funny, she was quite nervous at first, and me too, but more excited to finally learn it. It took her a bit too feel comfortable but then we became good friends. That was my first baking session. It was less about preparing the dough but rather the baking techniques, like timing and temperature.

From this initial spark, how did your passion for baking evolve over the years?
In 2007, I bought a sour dough culture while I was at Oracle OpenWorld. Back at home, I started baking with it but unfortunately the recipes included in the package weren't very fine. I still believe that the instructions were somehow wrong, so I didn't have a big success with that and decided to put it away. Then, in 2016, I started to develop my own cultures and bake with them. At that time there was much more information available on the internet. That was the real starting point for me. Like with the database, I wanted to know what was going on behind. I wasn't happy with "do this, this, this, and then, here's the success". Of course, you can just say "use this command", but sometimes it's good to know what is going on in the background, otherwise it is just a black box. So, I read a lot of papers trying to understand the fermentation process which was a whole new area for me as my knowledge of biology and chemistry had ended 46 years ago in high school. It wasn't so easy to read these scientific papers but most of the times I was just looking at the summary (laughs). I wanted to know how I could impact the quality of the bread with factors like water, time, temperature, or the decision between white or full grain flour. For example, more water results in more lactic acid in the bread. There is a lot of science behind it.

Is it true that you also grow your own grain?
Actually, this is the first year. I have a friend – we are singing together in the church choir – who is a farmer. Last year, I got some special kind of durum wheat from Austria. Durum wheat has a yellow color and is typical for South Italy where it's used for pasta. I like the bread very much that is made from it in the Apulia region, between Bari and Naples. Of course, they have ideal conditions which I don't have in Slovenia. Currently, we have good weather but I'm not sure yet, we will see in one month what will come out of it. In that region, there is the town Altamura which is also known as "the town of bread". They have been using the same recipe practically for more than 2,000 years and their bread has even been granted PDO status. The Roman poet Horace described it in his writings and recommended travelers to buy the bread because it stays fresh for two weeks. So, this the first time for me that and we just made an arrangement to put a bit of the durum wheat on his field. We'll see. And there is another small field with it in the North-Eastern part of Slovenia, close to Maribor which has the same environment as in Austria, around Vienna – a much drier climate which is probably very good for this bread. I'm always happy if I can get a grain from a local farmer, and many times I know exactly on which field it was cultivated. I have a very small mill at home, so I can mill it and get a fresh flour. Although the fresh flour is not very good for normal baking – it has to be about three to four weeks old, so that an oxidation already took place. However, it is much healthier if you mill it and make a bread immediately after that, because you keep all the vitamins and the good stuff inside. And I also prefer it because I love

the freshness of it. Also, the outcome is always dependent on the input. If you put in good ingredients, the outcome will likely be very good, but of course you can screw it (laughs).

Do you develop your own recipes?
Yes, but you start with some simple recipes. There are a lot of different recipes but basically, my recipes are the result of what I like to eat, simple as that. I'm not a friend of white bread. It doesn't have any good things inside, just the starch, ok, all carbohydrates, but nothing else. I like to put a whole grain flour inside and mix it with different types of grains, like barley, rye, or durum wheat, also corn and buckwheat – which is actually not a grain so you can make a gluten-free bread from it. But I'm not using it very much, just occasionally for my daughter.

Do you try different breads on your travels, and do you have a favorite country for bread?
I would say, in Germany the rye rules (laughs). I've had very good rye bread in Germany, also in Austria, but especially when I was in Hamburg I had a wonderful bread at the hotel, it was amazing. I also had some good bread in Scandinavia, but most of my travels was on Oracle conferences where you don't have much time to visit local bakeries. You often only get the bread at the hotel where it's mostly normal bread – bread that fits the taste of the majority. But what I usually did was that in every country I bought some flour in the grocery store. And I did something else, which was very unique and connected to my other hobby, the flying: At the conferences we often had welcome receptions for the foreign speakers a day before and I usually made an agreement with the organizer that I could bring my bread for the dinner. So, I was baking the bread, delivering it to the conference by flying myself and then was a speaker at the conference. The days before I always asked about their menu and then made a bread that fit it. I flew to the DOAG conference a couple of times, and also to conferences or Oracle classes in Denmark, Finland, Bulgaria, Scotland, Latvia, Estonia, Poland, a lot of countries. This way, I could always combine my hobby and the business. Actually, my hobby was rewarded by the business with their money which I would have otherwise received for driving by car or buying an airline ticket. But the major obstacle was always the weather. I changed my plane for a better one but it's still old, it's 56 years now, eleven years younger than me and we share the same birthday, December 28. After I changed the plane, I also did the training for flying by instruments, which is the same training that pilots of big airlines have. So now, I can also fly in bad weather. To receive the required instrument rating, you must have 200 hours of flying and then you have another 50 hours of training where you fly only by reference to the instruments without looking out of the cockpit. Then, I had to pass seven exams on theory, like meteorology. I did those exams on Gatwick in England where the CAA has their residence. On the final practical exam, you are flying with an examiner, and I was wearing some special kind of glasses so didn't see anything outside but just the instruments.

If you're doing it the first time, isn't it hard to trust the instruments?
Yes, it is. Because your body is telling you something, which is not reflected by the instruments. And it requires a lot of concentration, if you're flying one hour under instrument

conditions you get tired. For instance, when I was flying to Scotland the last time, I think the conference was in Glasgow, I made a middle stop to refuel in Lille. The weather was not really bad, but it was raining with low visibility. When I started to descend, approaching Lille, I entered the clouds at 9,000 feet, around 2700 meters, and exited just before landing. It was 12 o'clock and the visibility was low, and if they hadn't put on the lights on the runway, I wouldn't have seen it. I landed, refueled, and then departed to Scotland. I entered the clouds immediately after takeoff and exited the clouds well over Dover in England. So, all the time in the clouds I was just referencing the instruments.

How long have you been flying and was there ever a risky situation?

I started in 2007, before that I had been playing flight simulator on my computer. When I decided to really start flying, I got my private license within five months. I actually bought my airplane before, I bought it in October and then got the license in November. It was a Piper Aircraft which I bought in Mülheim, Germany. Every year I had to fly to that airport for the annual check, because I had the maintenance shop there. Until recently, when I slowed down a bit, I used to fly about 150 hours per year which is quite a lot for a private pilot. And I never had a serious situation. I was very cautious on deciding whether to fly or not. Several times I had to buy a airline ticket at the very last moment because of the weather. When I was teaching meteorology for pilots I always said: "It's much better being on the ground and thinking how nice it would be to fly now than the opposite way – to be in the sky and wishing to be on the ground." I had those moments myself.

So, you have these two passions, baking and flying, if you compare both, are there similarities?

People think the bread is not alive, but it is. You never know the outcome. You can follow instructions and be very precise, but doing the same thing on two different days,

the outcome will not be the same. There are so many small variables, especially if you're a private baker and don't have special devices for keeping constant temperature or humidity. It's a challenge. So, you never know the result until you cut the bread. The flying is almost the same. It's always a challenge with the weather. For example, flying to Hamburg, I have to cross Austria and the whole Germany, it's very likely that on my route I enter different kinds of weather conditions. So. you never know how your flight will be until you land.

Is this part of not knowing the outcome part of the fun for you?

I like the challenge. Like when a customer calls and says: we have a problem with the database. The bigger the problem, the more adrenaline you have. I like that. And the flying was a dream I had all my life. It already started when I was at the end of grammar school and built my first radio control for boats and also for airplanes. However, I was never thinking that one day I would be a pilot. I started flying so late because I didn't think

I was capable of it. But when I got my medical certificate, the doctor said: “Well, we are not looking for the superman.” So, I had a completely wrong impression. And it was funny, when I started at the flight school, I asked: “How much time do I need? When can I expect to get a pilot license?” They said: “Well, you’re 53 years old, your reflections are slow, you’ll probably need more than a year.” At the end I needed less than half of that. I think everybody has some kind of goal in their life. But they often think it’s undoable, it’s too far, that they cannot achieve it. So, when I actually started flying, that was a really good feeling.

Jože, thanks a lot for taking your time.

About Jože Senegačnik

Jože has more than 33 years of experience in working with Oracle products. He began in 1988 with Oracle Database version 4 while working for the City of Ljubljana, Slovenia, where he had charge over the city’s municipal and geographic information systems. From 1993 to 2003, he worked in developing GIS systems for the Surveying and Mapping Authority of the Republic of Slovenia, and in the development of applications for other governmental institutions, all based on the Oracle database. More recently, he has specialized in performance optimization, having developed his own toolset for monitoring performance and analyzing trace files. Jože is an internationally recognized speaker, a member of the OakTable Network and has become an Oracle ACE Director for his long record of positive contributions to the Oracle community.

Ann-Sofie Vikström Often

KEEP CALMS AND DEVOPS

Many of us are focusing on DevOps as a 'how to' in projects, but do we really do it as it should be done to achieve the values we want?
Since the early 2000s, agile approach has been transformed from agile development to continuous integration, then to continuous delivery and finally to DevOps.
There has been a long tradition to separate development from operations, and I must admit that Ops seems to be the hardest part to transform. Some companies even outsourced the operations to one 3rd party, doing the development themselves or outsourced development to another 3rd party.
As coming from a consultant company, I find that project contract is not handling DevOps at all, even if the customer is asking for it, or they think they ask for it but mean continuous integration or ditto delivery. And is the customer's IT organization at all prepared?
In this article I will focus on the big change that is needed. It all starts with the change of culture.

Going from Silos and Complex Ticketing Systems to a Team-Based Collaboration

The enabler for this transformation is when development and operations talk about the:

- product
- production
- build metrics

And throughout the lifecycle discussing

- requirements
- features
- schedules
- resources
- whatever else might come up

And how hard can that be? If we look at different terminologies, the confusion is complete. Many of them focus only on the process. But Gartner says: "Rather than being a market per se, DevOps is a philosophy, a cultural shift that merges operations with development and demands a linked toolchain of technologies to facilitate collaborative change."

DevOps is not a process, or a different approach to development – the successful transformation is definitely culture change.

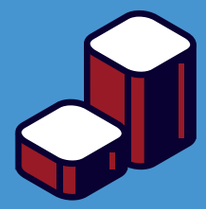

The framework **CALMS** has become popular for assessment of the readiness to adopt DevOps, but also measuring success during the transformation. Introduced in "The DevOps Handbook" by authors Gene Kim, Patrick Debois, John Willis and Jez Humble, **CALMS** is an acronym for:

Culture – there is a culture of shared responsibility where collaboration is the key.
Automation – team members seek out ways to automate as many tasks as possible and are comfortable with the idea behind continuous delivery.

Lean – team members are enabled to use lean principles for software development, and the way they work.

Measure everything (applications, business values, infrastructure) for continuous improvement and refining of cycles

Sharing experiences, both bad and successful, to enable others to learn.

Culture and Collaboration

Culture and collaboration are the hardest challenge and comes with no quick fix:

- The developer wants to change and the operations wants stability.
- IT contracts are almost always about delivery, and the aspect of involving the customer's IT people is not included.
- IT operations are often steered by the ITIL framework, audit, risk management etc. The solution for this challenge could be to use ITIL Agile together with the CALMS framework.
- How to deal with 3rd party operations in a DevOps environment?
  - Contract with SLAs is not providing a culture of DevOps

The Solution:

- Establish cross-functional teams of development and operations
- Share common goals
- The change of mindset will take time
- Mind the gap:
  - When you operate, you are risk-averse and are naturally inclined to limit changes to systems as you are rewarded for stability.
  - When you develop, you are rewarded for innovation and delivering at speed.

Team configuration examples:

Daily deployment
➔ DevOps team with shared skillsets. In this team you will find the DevOps Engineer.
Several Deployments per week?
➔ Dev and Ops in the same team.
Strict release management and/or fixed releases?
➔ Dev and Ops team in close collaboration.

Automation – Speeding up the process

- Start with Build, Test, Deploy and provisioning automation if not in place already
  - Developers commit to the version control system
  - Software is built on every commit
  - Automatic testing on every commit
  - Automatic build
- Feedback from the automatic build and test is given to the developer quickly
- When getting more mature include Infrastructure as a code

Lean Principles

Use lean principles as the enabler for continuous improvement. Sources can be:

- Retrospective
- Input from the metrics monitored
- A/B-tests (user experience research/experiment)
- Interviews
- Surveys
- Web analytics
- Customer journey mapping
- Root cause analysis

Measure – Closing the DevOps Loop

Collect events from everything, everywhere

- application logs, database logs, network logs, interactions with external systems, configuration files, performance data, usage data etc.

The format of log events should be in JSON to be consumed by products like

- Splunk
- Elasticsearch Logstash Kibana (ELK)

What to measure? Examples:

- How long did it take to go from development to deployment?
- How often do recurring bugs or failures happen?
- How long does it take to recover after a system failure?
- How many new customers did we gain/loose last week?
- How long did the “Customer journey” take?

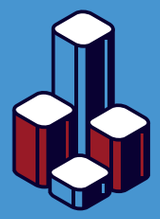

Sharing

The last part of CALMS encourages to share, share and share through:

- Retrospectives
- Developer and operator pairing with each other, if not already in the same team, throughout the application’s lifecycle.
- Developer also working with operational stuff from time to time.
- Anything else that will help break down the silos.

Interested in more reading?

- **https://www.slideshare.net/JanJoostBouwman/itil-and-devops-can-be-friends**
- **https://devblog.axway.com/dev-insights/keep-calms-and-devops/**
- **https://www.atlassian.com/devops/frameworks/calms-framework**

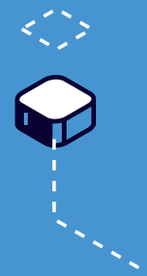

AutoML: ORACLE

Machine Learning for Beginners

Heli Helskyaho

Oracle offers several tools for machine learning. You can, for example, use the in-database machine learning with models built in SQL, R, or Python. Or you can connect to Oracle Database with different libraries, such as cx_Oracle, and use the data from the Oracle Database with different IDEs for machine learning. Or you can use Oracle Data Science Cloud that has an environment for Python machine learning including special Oracle libraries, and the possibility to pip install any Python libraries.

The in-database machine learning means that the data will remain where it is: there is no need to move large amounts of data outside the database. Also, typically a database is the best place to perform operations to data. For a long time, Oracle has offered in-database machine learning for SQL (OML4SQL) and R (OML4R), but now also Python (OML4Py) is supported both in the Autonomous Database and the Oracle Database 21c on-premises. Together with OML4Py Oracle released also the Oracle Machine Learning (OML) Auto ML functionality. OML Auto ML is based on special Oracle libraries for Python that perform several machine learning steps automatically. All these options of in-database machine learning are cost-free options of the Oracle Database and the easiest start would be an Oracle Cloud Free Tier with an Autonomous Oracle Database that automatically includes in-database OML.

OML AutoML User Interface in the Oracle Autonomous Database is a great option for somebody who is new to machine learning and/or Python. It creates a whole script set needed for a machine learning process and the user can save the process as a Notebook. That Notebook can be opened, investigated, and edited. The machine learning model can also be deployed using AutoML.

To start with the OML AutoML, sign in using the OML credentials. If you do not have a machine learning user created, you need to create one. In the Oracle Autonomous Database, it can be done using the Manage Oracle ML Users functionality under Autonomous Database Administration.
The Oracle Machine Learning Notebooks have a lot of different functionalities, but in this article, we only talk about OML AutoML that can be found under Quick Actions.
If you select OML AutoML, the AutoML Experiment is opened. Here you can create a new AutoML Experiment, or Edit/Delete/Duplicate existing ones. To create an AutoML Experiment, press Create. Then the Create Experiment opens as shown in ***Figure 1***.

Define a name for the experiment, and select the data source. Select the feature (column) you want to predict and the prediction type (Classification or Regression). You can also optionally select the Case ID as being the identifier of each row in the dataset.

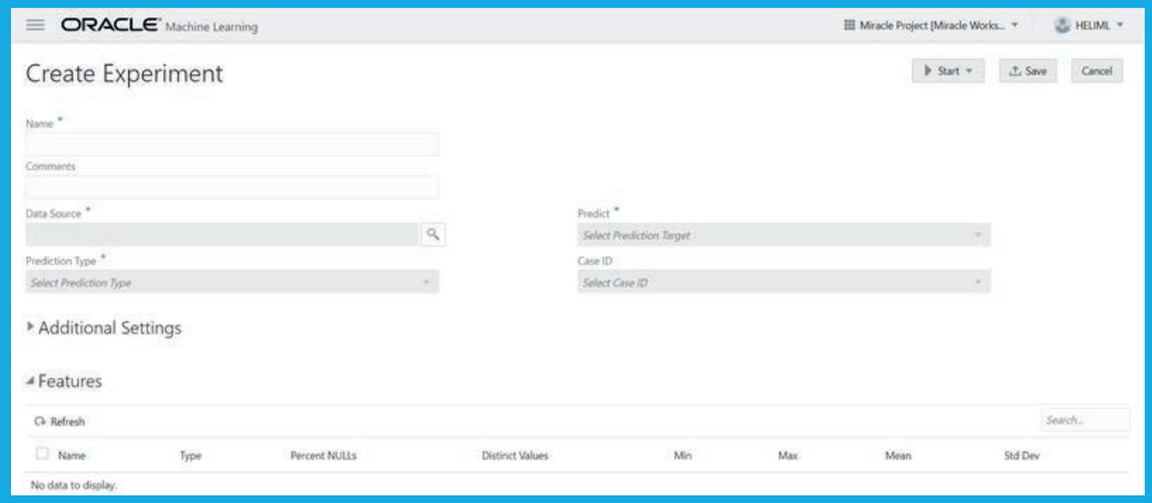

Figure 1: Create Experiment.

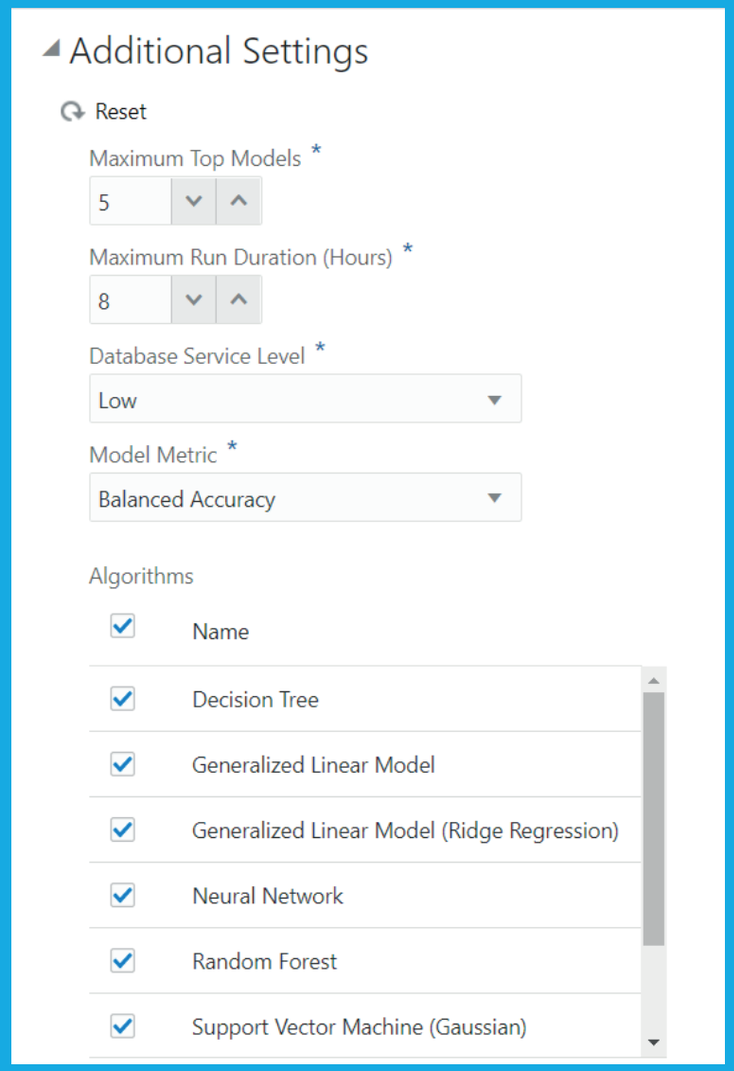

Figure 2: Advanced Settings for Classification.

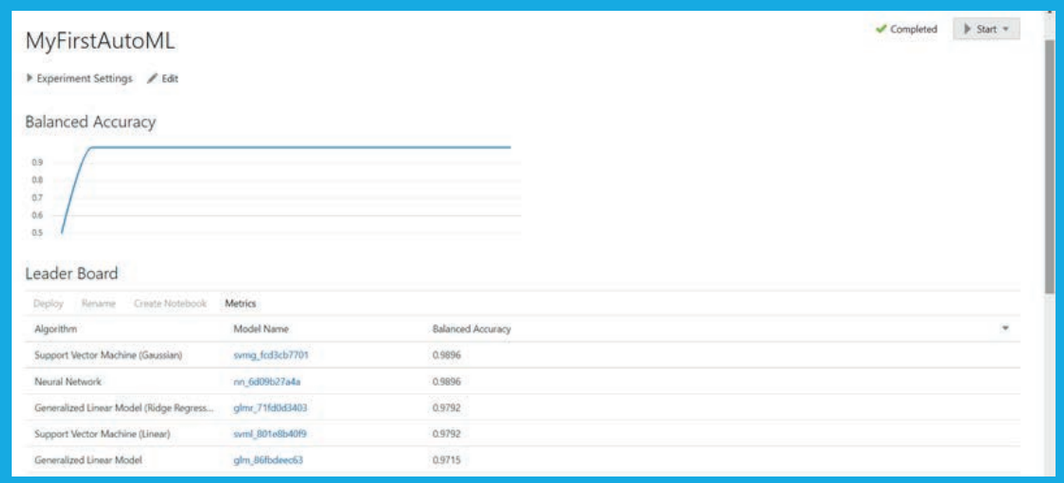

Figure 3: The Machine Learning Experiment has completed.

Under Additional Settings you can define relevant hyperparameters, model metric, and select the algorithms that will be used to train the models. The possible selections depend on the chosen prediction type. In ***Figure 2*** you can see examples of Addition Settings for Classification.

In the Features section you can observe detailed information about the features, including the name, type, percent of null values, distinct values, minimum and maximum values, mean, and standard deviation. You can select or deselect these attributes to be included in the model training.

When you are ready with your Machine Learning Experiment, just press Start on to top right corner as seen in ***Figure 1***. There are two possibilities for running the experiment: Faster Results or Better Accuracy. Choose the preferred option. As the name implies, these are experiments that you can run as many times as you want. When the experiment has completed you can see a Completed sign on the top right corner (***Figure 3***). If you are not happy with the outcome, you can press Edit, change the settings and run it again.

The Leader Board (***Figure 3***) shows how the different algorithms performed for the metric chosen in settings. In this example we chose Balanced Accuracy. Both Support Vector Machine and Neural Network performed equally in our experiment. If you scroll down that page, you can also see a section called Features. Features show details of each feature and the importance of it for the prediction.

If you want, you can also add additional metrics to compare the models. From the Leader Board select Metrics and the list of possible additional metrics is shown. Select the ones you want and close the pane.

If you select one of the algorithms, you can rename the model. By clicking the model's name, you can see more details about

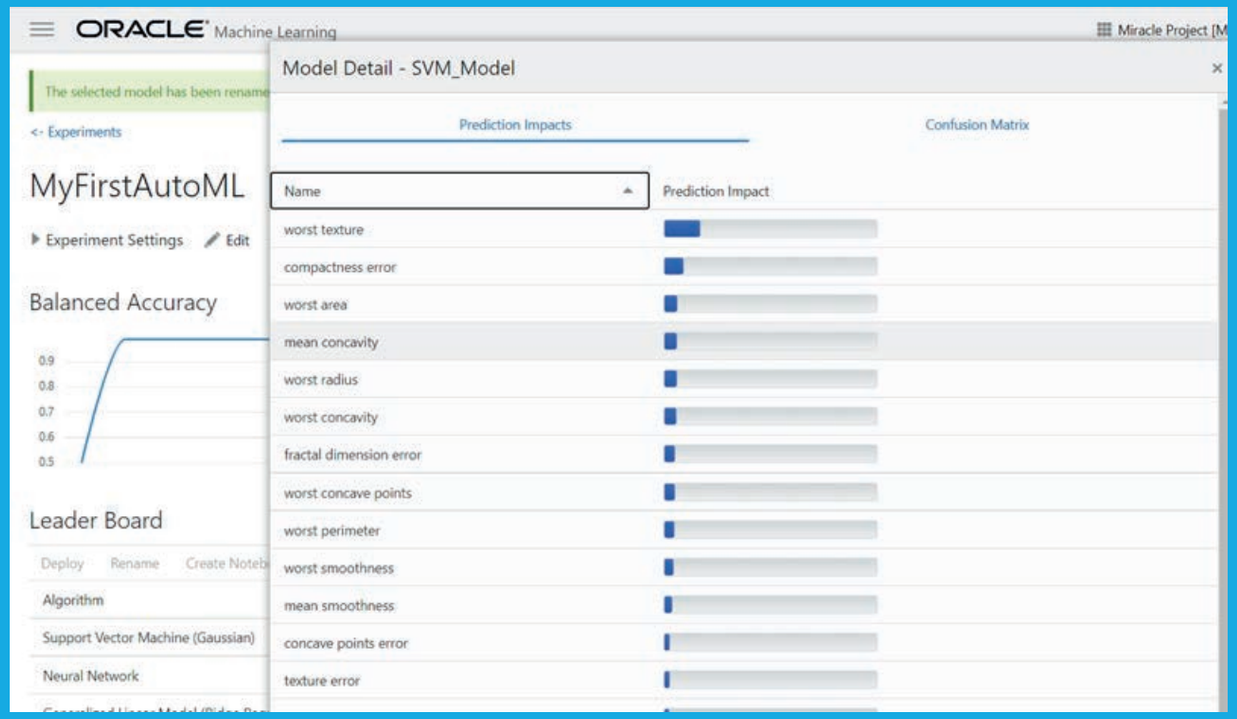

Figure 4: Prediction Impact of each feature to the selected model.

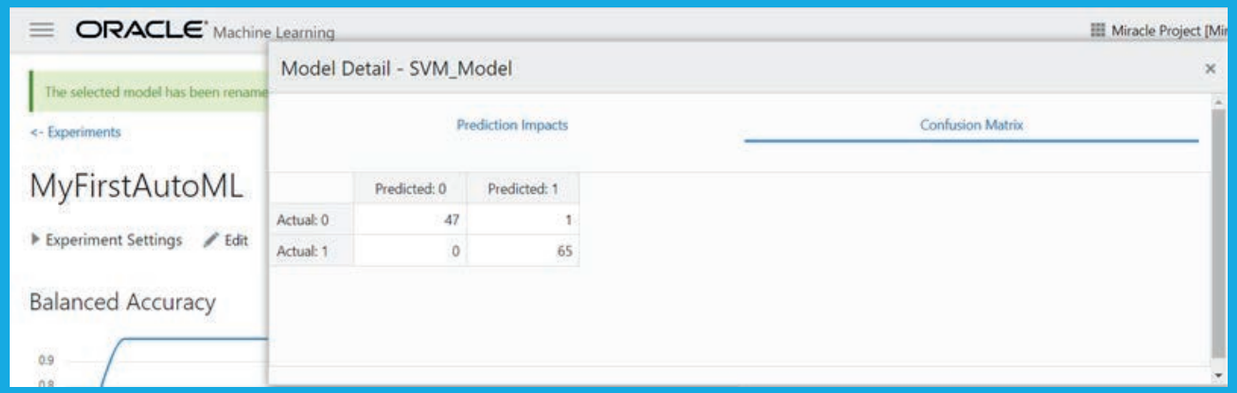

Figure 5: Confusion Matrix.

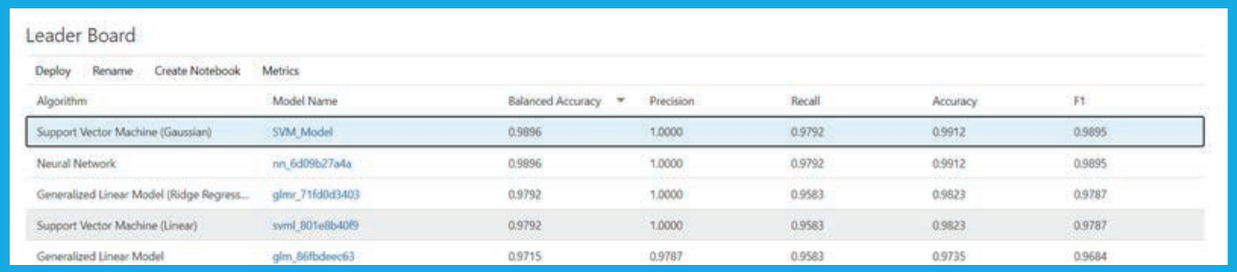

Figure 6: Deploy, Rename, or Create Notebook from the model chosen.

that particular model. There are two tabs: Prediction Impact and Confusion Matrix. The Prediction Impact (***Figure 4***) shows what impact each feature had for the model, while as Confusion Matrix (***Figure 5***) demonstrates how well the model predicted.

When you have decided what algorithm created the best model, select that model and either deploy it by pressing Deploy or create a new Notebook based on it by selecting Create Notebook. You can see these selections in ***Figure 6***.

If you select deploy, just insert the data needed (***Figure 7***) and press OK.

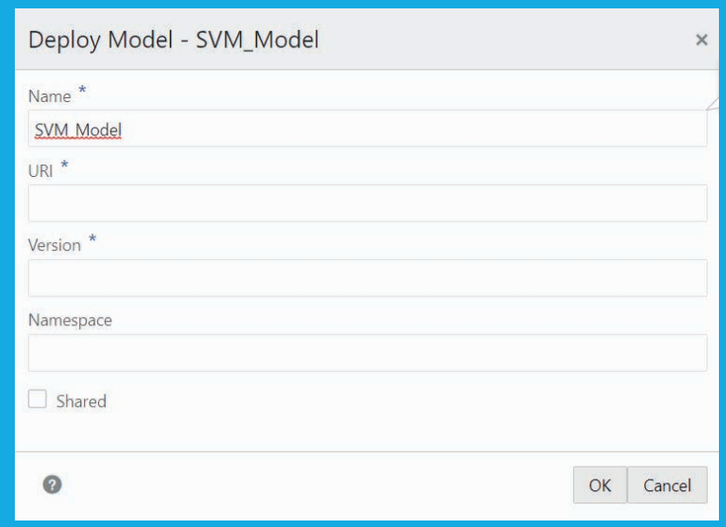

Figure 7: Deploy model called SVM_Model.

If you select Create Notebook, you are asked for the name of the new notebook and then the notebook is created. You will find the notebook from the list of notebooks from Quick Actions or from the hamburger menu on top left.

Now you are able to view and edit the new notebook as shown in ***Figure 8***.

This new OML AutoML functionality is an easy way of learning both machine learning and Python. It has its limitations but for a beginner that is probably a good thing since machine learning is a very broad topic to learn. It is easier to start with something limited. Enjoy machine learning and OML AutoML!

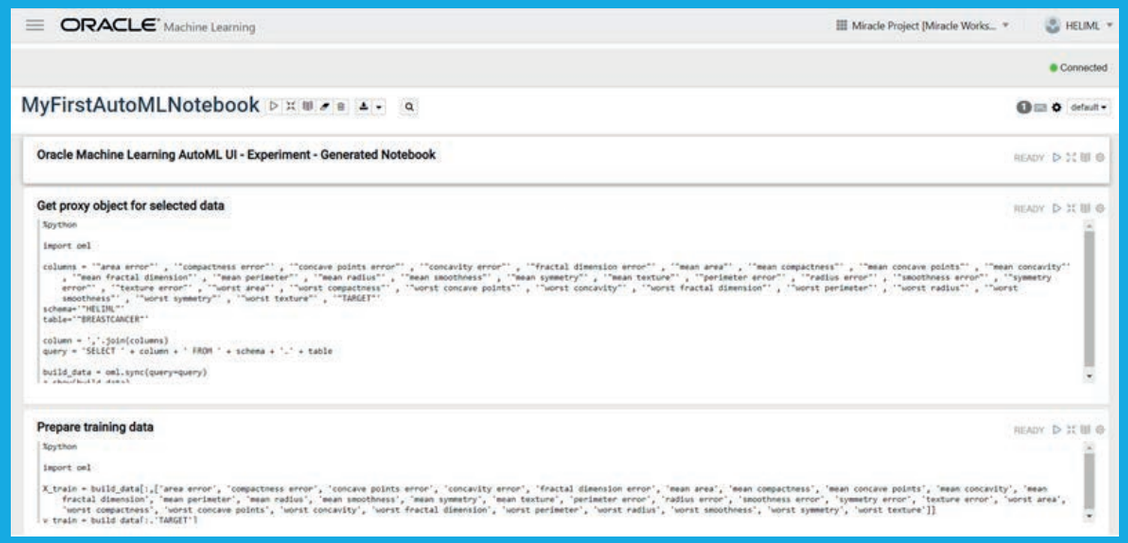

Figure 8: A notebook.

About Heli Helskyaho

Heli is an Oracle ACE Director, Oracle Groundbreaker Ambassador, and CEO of Miracle Finland Oy. Heli holds a Master's degree (Computer Science) at the University of Helsinki and she is specialized on databases. At the moment she is working on her doctoral studies, researching and teaching at the University of Helsinki.
Heli believes that understanding your data makes using the data much easier. She is the author of Machine Learning for Oracle Database Professionals: Deploying Model-Driven Applications and Automation Pipelines (Apress, 2021), Oracle SQL Developer Data Modeler for Database Design Mastery (Oracle Press 2015) and a co-author of Real World SQL and PL/SQL: Advice from the Experts (Oracle Press 2016).
Heli was listed as one of the TOP 100 influences on IT sector in Finland in 2015, in 2016, in 2017, in 2018, in 2019 and again in 2020. Heli is also a Data Vault 2.0 Certified Practitioner.

Jim Czuprynski

Fast and Furious: Capitalizing on Oracle 19c Fast Ingest & Fast Lookup Features

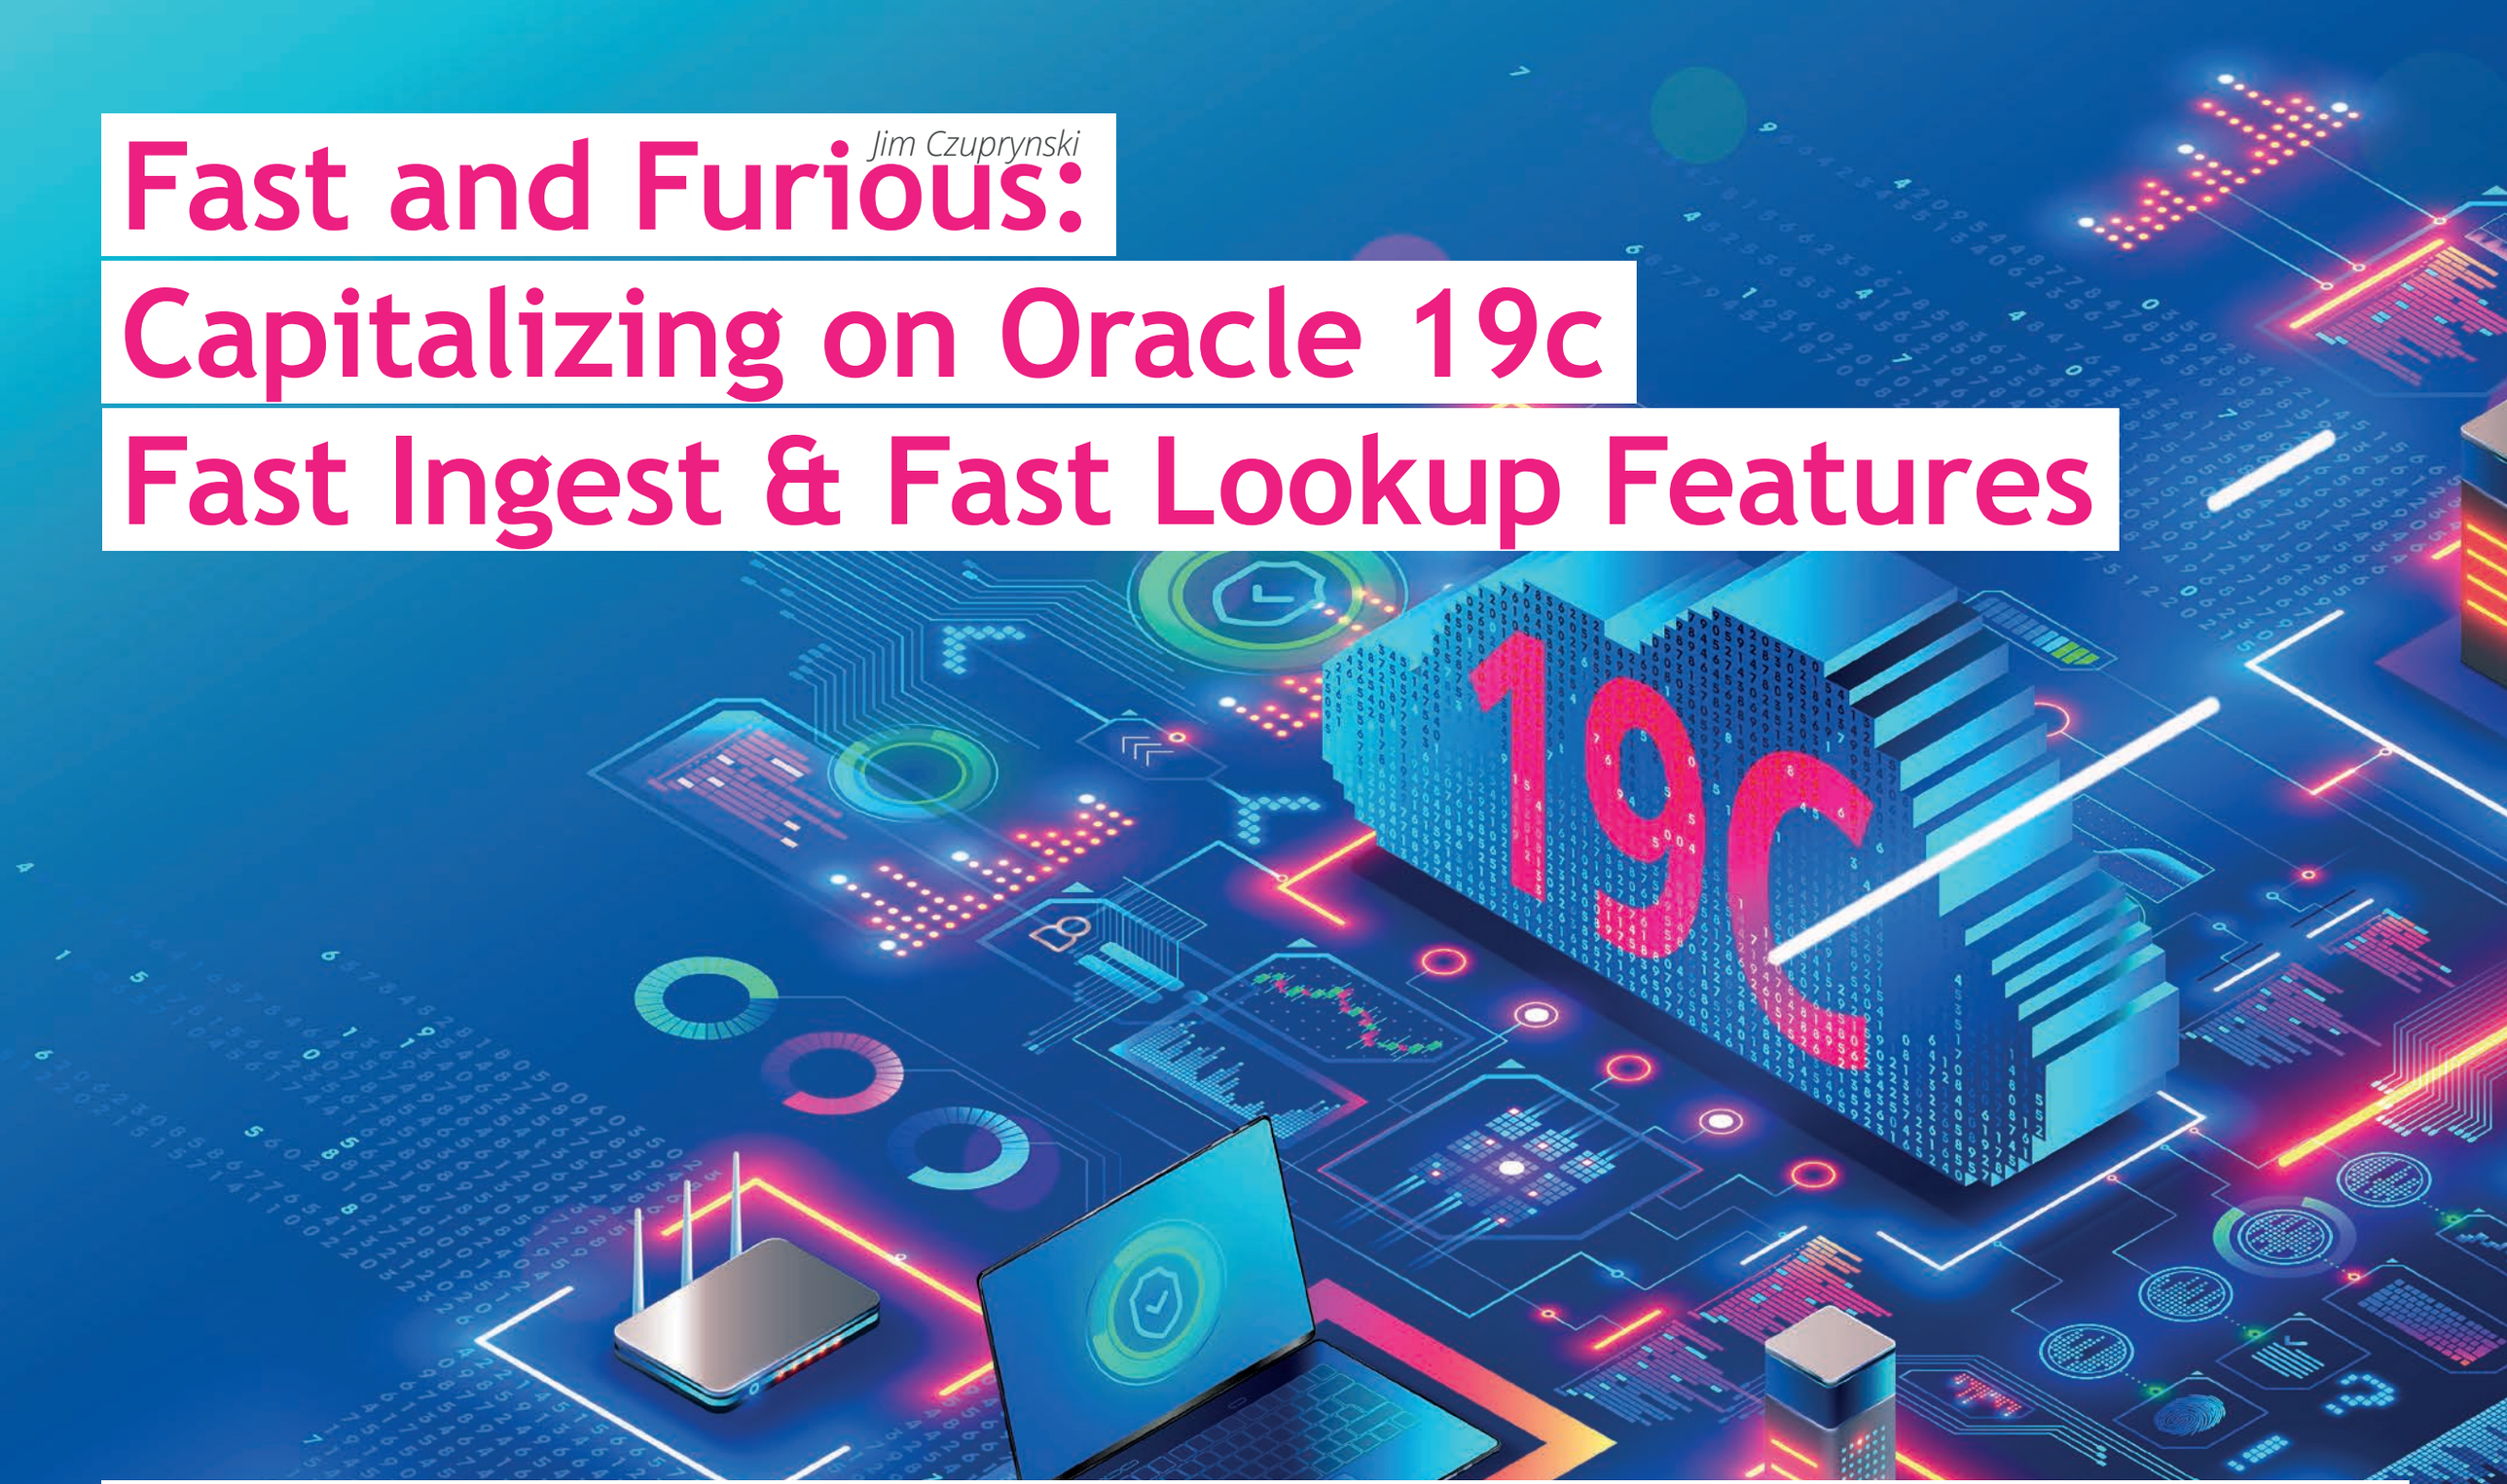

The Internet of Things (IoT) has penetrated every corner of our modern world, from the smart meters that monitor our home's electricity usage to the smart home technology that lets us turn on lights or watch our favorite shows with a simple voice command. All these devices generate enormous quantities 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.

Continuing the previous article in this series, Oracle ACE Director Jim Czuprynski wraps up his investigations into Oracle Database 19c's Fast Ingest features and turns his attention to the features of Fast Ingest's close cousin, *Fast Lookup*.

Monitoring and Managing Fast Ingest Workloads

In the **first part** of this series, I showed you how to activate the *Memoptimized Row Store* (MORS) within an Oracle database instance to enable Fast Ingest features, and then I used those features to insert huge volumes of data from numerous sessions via an intense SwingBench-generated workload to simulate capturing data from streaming sources. I also measured the performance benefits of Fast Ingest versus traditional INSERT transaction processing via AWR difference reports.

In this last part of the series, we'll take a look at 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

Monitoring MORS Using V$MEMOPTIMIZE_WRITE_AREA

Let's briefly return to the streaming simulation I demonstrated in the last article, which consisted of 25 user sessions that randomly inserted rows into the **T_METER_READINGS** table at an extreme rate to demonstrate the performance advantages of Fast Ingest.

To take a look at how well my database instance is handling the Fast Ingest workload, I can query the V$MEMOPTIMIZE_WRITE_AREA dynamic view to see how much space is available for Fast Ingest INSERT transactions, how much has been used, and how many cumulative writes have occurred so far. ***Listing 1*** shows a sample query and an example of its output as that query was executed over time while Fast Ingest workloads were being generated.

What's interesting here is that even as the number of transactions written and the number of writer processes gradually *increase* over time – they are apparently a cumulative count - the free space allocated to the MORS actually *decreases*

```
SQL> SELECT
       (total_size / (1024*1024)) total_size_mb
     ,(used_space / (1024*1024)) used_space_mb
     ,(free_space / (1024*1024)) free_space_mb
     ,num_writes
     ,num_writers
     ,con_id
  FROM v$memoptimize_write_area;


TOTAL_SIZE_MB USED_SPACE_MB FREE_SPACE_MB NUM_WRITES NUM_WRITERS CON_ID
------------- ------------- ------------- ---------- ----------- ------
         2055    .156784058    2054.84322          0         106      3

. . .

TOTAL_SIZE_MB USED_SPACE_MB FREE_SPACE_MB NUM_WRITES NUM_WRITERS CON_ID
------------- ------------- ------------- ---------- ----------- ------
         2055    1.15670776    2053.84329       8830         107      3

. . .

TOTAL_SIZE_MB USED_SPACE_MB FREE_SPACE_MB NUM_WRITES NUM_WRITERS CON_ID
------------- ------------- ------------- ---------- ----------- ------
         2055    3.15655518    2051.84344      24681         109      3
```

Listing 1: Query against V$MEMOPTIMIZE_WRITE_AREA and resulting output

over time as the data stream is consumed and INSERTs are performed via Fast Ingest outside of standard Oracle transaction processing.

Monitoring Fast Ingest Worker Processes

Listing 2 shows another way to monitor what's happening within my database instance as the sample workload generation

```
SELECT
    P.program "Executor"
    ,N.name "Statistic"
    ,S.value "Value"
  FROM
    v$sesstat S
    ,v$statname N
    ,v$session P
  WHERE S.statistic#  = N.statistic#
    AND S.sid = P.sid
    AND N.name LIKE '%memopt%'
    AND S.value > 0
  ORDER BY P.program, N.name;
```

Listing 2: Query to observe MORS activity within database instance

```
Executor             Statistic                          Value
-------------------- ---------------------------------- ----------
JDBC Thin Client     memopt w buffer gets                       10
JDBC Thin Client     memopt w buffer gets                        9
JDBC Thin Client     memopt w buffer gets                        9
. . .
JDBC Thin Client     memopt w buffer gotcur                     10
JDBC Thin Client     memopt w buffer gotcur                     10
JDBC Thin Client     memopt w buffer gotcur                      9
. . .
JDBC Thin Client     memopt w buffer hit bucket 0                9
JDBC Thin Client     memopt w buffer hit bucket 0                9
JDBC Thin Client     memopt w buffer hit bucket 0                9
. . .
JDBC Thin Client     memopt w rows written                  230993
JDBC Thin Client     memopt w rows written                  231090
JDBC Thin Client     memopt w rows written                  228340
. . .
oracle@host21c (W001) memopt w drain sleep                    554
oracle@host21c (W001) memopt w drain sleep wake post          224
oracle@host21c (W001) memopt w drain sleep work                 3
```

Figure 1: Results of querying dynamic views in Listing 2 to observe MORS activity

unfolds. Any session-level statistic that's prefixed with **memopt** is pertinent to activity that's occurring within the MORS.

Figure 1 shows the results of running the query in ***Listing 2*** against my 21c instance using the SYS account. (I've edited the output to remove repetitive content per each session.) Note that the last few lines in this abbreviated list include statistics from the actual "worker" processes – in this case, **W001** – that are responsible for handling the processing of the rows being inserted via Fast Ingest.

Controlling the In-Memory Row Store Via DBMS_MEMOPTIMIZE

I mentioned earlier that Fast Ingest essentially bypasses normal transaction processing. Instead, it maintains an instance-level low water mark (LWM) as well as session-level high water marks (HWM) for determining whether all INSERTs have successfully completed.

Fortunately, the **DBMS_MEMOPTIMIZE** package supplies us with functions and procedures we can use to monitor the LWM and HWMs. It's also possible to use **DBMS_MEMOPTIMIZE** to expedite Fast Ingest processing when necessary – for example, when terminating a still-connected application session, or just prior to shutting down a database instance for maintenance, patching, or other administrative tasks.

Monitoring Fast Ingest LWM and HWMs

Listing 3 shows how to use the **GET_APPLY_HWM_SEQID** and **GET_WRITE_HWM_SEQID** functions to determine the exact status of Fast Ingest processing at the session and instance level.

To get a better understanding of how these two watermarks can be used to determine if it's safe to proceed with administrative operations, let's take a look at what happens within a single

session. First, I'll run the code in ***Listing 3*** to display the LWM and HWMs before starting a new Fast Ingest transaction:

```
Low-Water-Mark for All Rows Applied (Written Globally):498413730056
High Water Mark for Rows Written For This Session:      0
```

At this point, there's no activity for this session, so the HWM is zero. However, note that after I generated a Fast Ingest transaction within this same session, the session's HWM has advanced beyond the instance's LWM, which indicates the data for that session hasn't yet been written to the database:

```
Low-Water-Mark for All Rows Applied (Written Globally):498536313235
High Water Mark for Rows Written For This Session:     498536328525
```

Finally, after just a few more seconds have passed, note that global LWM has surpassed the session's HWM, thus indicating that all data has been successfully written to the database, so it's now safe to end this session.

```
Low-Water-Mark for All Rows Applied (Written Globally):499537736037
High Water Mark for Rows Written For This Session:     498536328525
```

Obviously, I would need to check that all application sessions using Fast Ingest showed a similar status before terminating the database instance or performing any administrative tasks that could disrupt the application's integrity – for example, modifying any of the application's base tables that are currently enabled for Fast Ingest.

Flushing Unwritten Fast Ingest Data

Instead of waiting for all sessions to complete their transaction activities, I can also flush any unwritten data to the database at either the session or system level via two **DBMS_MEMOPTIMIZE** procedures.

Listing 4 shows how the **WRITE_END** procedure flushes *all* Fast Ingest data from the Large Pool for just the current session, while ***Listing 5*** illustrates how procedure **WRITES_FLUSH** will force *all* unwritten data for *all* sessions to the database. (Thankfully, pluralizing the global procedure helps me to remember which one does what.)

Fast Lookup: Concepts and Strategies for Use

Hopefully I've demonstrated that Fast Ingest works quite well for

```
DECLARE
  applied_hwm NUMBER(15,0);
  written_hwm NUMBER(15,0);
BEGIN
  applied_hwm := DBMS_MEMOPTIMIZE.GET_APPLY_HWM_SEQID;
  written_hwm := DBMS_MEMOPTIMIZE.GET_WRITE_HWM_SEQID;
  DBMS_OUTPUT.PUT_LINE(
     'Low-Water-Mark for All Rows Applied (Written Globally): ' || applied_hwm);
  DBMS_OUTPUT.PUT_LINE(
       'High Water Mark for Rows Written For This Session:    ' || written_hwm);
EXCEPTION
  WHEN OTHERS THEN
     DBMS_OUTPUT.PUT_LINE(
       'Fatal unexpected error: ' ||
       SQLCODE || ' - ' || SQLERRM);
END;
```

Listing 3: Monitoring Fast Ingest low and high water marks for written transactions

```
BEGIN
  DBMS_MEMOPTIMIZE.WRITE_END;
EXCEPTION
  WHEN OTHERS THEN
     DBMS_OUTPUT.PUT_LINE(
     'Fatal unexpected error: ‘ ||
       SQLCODE || ' - ' || SQLERRM);
END;
/
```

Listing 4: Flushing all unwritten Fast Ingest data for just this session with WRITE_END

```
BEGIN
  DBMS_MEMOPTIMIZE_ADMIN.WRITES_FLUSH;
EXCEPTION
  WHEN OTHERS THEN
     DBMS_OUTPUT.PUT_LINE(
     'Fatal unexpected error: ' ||
     SQLCODE || ' - ' || SQLERRM);
END;
/
```

Listing 5: Flushing all unwritten Fast Ingest data for all sessions with WRITES_FLUSH

adding streaming data quickly to a database. However, there's another facet of the MORS that allows us to *access* selected data extremely quickly as well: *Fast Lookup*. Interestingly, this feature set is focused at retrieving much smaller "chunks" of data from a database table, essentially at a single-row level. ***Figure 2*** illustrates how it's similar to Fast Ingest and how its command set overlaps with that set of features.

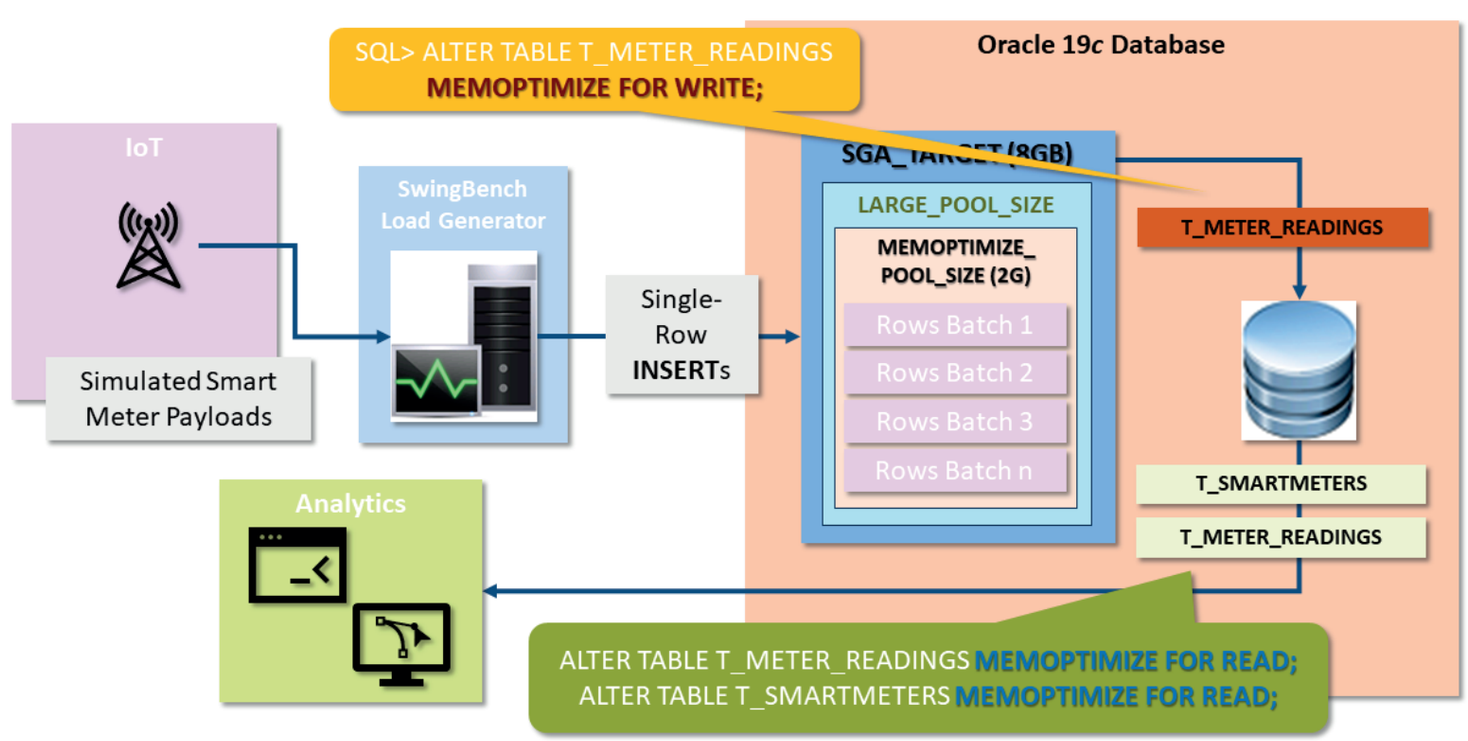

Figure 2: Fast Lookup concepts

A good metaphor to explain the difference? if Fast Ingest is like my database is *drinking from a firehose*, Fast Lookup is its microscopic counterpart *enabled with a pair of tweezers* to retrieve and retain for read-only access only the data needed to answer a SQL query. In short, here's how it works:

- After a table is placed in **MEMOPTIMIZE FOR READ** mode, any Fast Lookup request against its data builds a hash table of its primary key (PK) values.
- Any rows so retrieved are added to separate buffers in the **MEMOPTIMIZEd** pool and remain pinned there until they are removed.
- The hash table is thus a direct shortcut to reading a row directly from its corresponding buffer. This is actually much faster and efficient than reading it from a block in the traditional database buffer cache, since only the rows so selected are retained in the MORS.

As one might guess, Fast Lookup is thus most useful for data that will be accessed frequently based on exact primary key values. ***Table 1*** highlights a few use cases that would be of particular value in this scenario.

Enabling a Table for Fast Lookup

Let's explore Fast Lookup's usefulness against some of the tables in my application's **SIMIOT** schema.

First, I attempted to activate the **T_SMARTMETERS** table for Fast Lookup by issuing the **ALTER TABLE** … MEMOPTIMIZE FOR READ command for the desired table ***(Listing 6)***. I was a bit surprised to receive an unexpected error because I had forgotten to add a valid primary key constraint for the table.

In retrospect, this error message made perfect sense because, as I explained earlier, one requirement for Fast Lookup is an active PK constraint enforced by a unique index on the table. After I added an appropriate index for that table, it was successfully added to the MORS for Fast Lookup.

Next, I tried to enable the **T_METER_READINGS** table for Fast Lookup, and received yet another unexpected error, as shown in ***Listing 7***.

This one was a bit more puzzling because I had already verified that I'd loaded data into that table. However, I'd forgotten that the **CURRENT** partition as yet contained no data, as that partition is essentially reserved for population through SwingBench workload generation. After inserting a few rows into that partition and then reissuing the command, the table was activated successfully for Fast Lookup.

Populating Tables and Partitions for Fast Lookup

Enabling a table for Fast Lookup is just the first step – I also need to populate it into the MORS. ***Listing 8*** shows how to use

Scenario	Example
Often-accessed rows	Rows containing values most likely to be accessed most frequently – for example, movies **most recently added to a streaming service**
Most recently loaded rows	Users tend to access data about **most recent store** sales than yesterday's / last week's sales
Repetitive reads on the same larger "chunks" of data	Rows containing a nearly-fully-populated **VARCHAR2(2000)** column

Table 1: Fast Lookup – Appropriate Use Case Examples

```
ALTER TABLE t_smartmeters MEMOPTIMIZE FOR READ;

ERROR at line 1: ORA-62142: MEMOPTIMIZE FOR READ feature requires NOT
DEFERRABLE PRIMARY KEY constraint on the table
```

Listing 6: Enabling Fast Lookup: An expected error

```
ALTER TABLE t_meter_readings MEMOPTIMIZE FOR READ;

ERROR at line 1: ORA-62156: MEMOPTIMIZE FOR READ feature not allowed on
segment with deferred storage
```

Listing 7: Enabling Fast Lookup: A surprising error

NOTE: During my search for Fast Lookup's limitations, I was unable to uncover any My Oracle Support (MOS) notes that describes them in detail as of this writing. For now, the best advice I have is to review the **Oracle Error Messages** guide for your current database release for all error messages in the range of **ORA-62100** thru **ORA-62200**.

```
BEGIN
  DBMS_MEMOPTIMIZE.POPULATE(
     schema_name => 'SIMIOT'
     ,table_name => 'T_METER_READINGS'
     ,partition_name => NULL
  );
EXCEPTION
  WHEN OTHERS THEN
     DBMS_OUTPUT.PUT_LINE(
     'Fatal unexpected error: ' ||
       SQLCODE || ' - ' || SQLERRM);
END;
/
```

Listing 8: Populating an object into the Fast Lookup pool

```
BEGIN
  DBMS_MEMOPTIMIZE.DROP_OBJECT(
     schema_name => 'SIMIOT'
     ,table_name => 'T_METER_READINGS'
     ,partition_name => NULL
  );
EXCEPTION
  WHEN OTHERS THEN
     DBMS_OUTPUT.PUT_LINE(
     'Fatal unexpected error: ' ||
       SQLCODE || ' - ' || SQLERRM);
END;
/
```

Listing 9: Removing an object from the Fast Lookup pool

```
SELECT *
  FROM t_smartcos
WHERE sm_id BETWEEN 2688079 AND 2688081;
```

Listing 10: Searching for a small range of Smart Meters

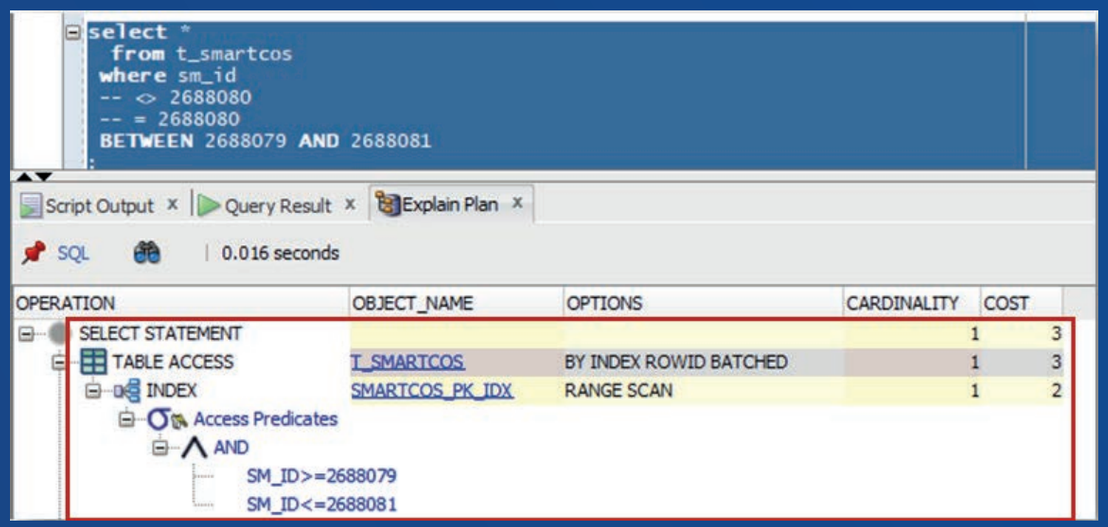

Figure 3: Fast Lookup: Range scan failure

the **POPULATE** procedure of **DBMS_MEMOPTIMIZE** to add the table – or, if it's partitioned, just a single partition – into the Fast Lookup pool. As shown, setting the **partition_name** parameter to **NULL** will populate all of the table's partitions into the MORS.

Likewise, I can also remove an entire table or even a single partition of the selected partitioned table from the Fast Lookup pool via the **DROP_OBJECT** procedure, as shown in ***Listing 9***.

Fast Lookup: Indications of Successful Use

To demonstrate this simply and quickly, I'll build a small sample table named **T_SMARTCOS** using the code that I've posted to this **GitHub repository**. This table contains about 250K rows that I've built using the **CREATE TABLE AS SELECT** method, but also contains a column (**SM_DESCRIPTION**) that stores a large amount of randomized text about each customer that my application will need to display whenever I retrieve a customer's record via my customer relationship management application. As you'll recall, this is a perfect use case for using Fast Lookup to store these data directly in memory for ultra-fast retrieval.

To demonstrate that Fast Lookup was employed successfully, consider the query in ***Listing 10***, which attempts to retrieve information from **T_SMARTCOS** for just a few customers.

Figure 3 shows the execution plan for this statement. Note that even though it's retrieving a maximum of three rows based on the possible range of values for the **SM_ID** column – its primary key - the optimizer still chooses to use the primary key index for a range scan to quickly retrieve the rows.

However, note that when accessing the same table using an *equality predicate* – in this case, setting the selection criteria for a single specific Smart Meter ID, as shown in ***Listing 11*** – the optimizer can take advantage of Fast Lookup features via two

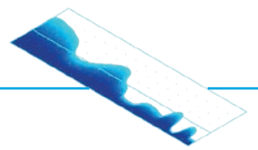
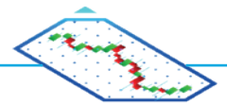
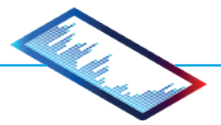

new operations: **BY INDEX ROWID READ OPTIM** and **UNIQUE SCAN READ OPTIM** (*see Figure 4*). These operations confirm that the optimizer is indeed using Fast Lookup features to quickly access the larger data "chunk."

Monitoring Use of Fast Lookup Across Multiple Sessions

The query in ***Listing 12*** may look familiar, as I introduced an earlier version of it in the previous article. We're isolating Fast Lookup activity within the database by selecting only those statistics that contain the **memopt r** string.

Figure 5 shows the result from that query, which I ran during numerous experiments of Fast Lookup capabilities. Note that Fast Lookup deploys several worker processes (**W00n**) to copy individual rows into the MORS whenever I executed the **POPULATE** procedure of **DBMS_OPTIMIZE** as indicated by the **memopt r rows populated** statistic. Likewise, the **memopt r**

```
SELECT *
  FROM t_smartcos
WHERE sm_id = 2688080;
```

Listing 11: Searching for just one Smart Meter

SQL | 0.017 seconds

OPERATION	OBJECT_NAME	OPTIONS	CARDINALITY	COST
SELECT STATEMENT			1	2
TABLE ACCESS	T_SMARTCOS	BY INDEX ROWID READ OPTIM	1	2
INDEX	SMARTCOS_PK_IDX	UNIQUE SCAN READ OPTIM	1	1
Access Predicates				
SM_ID=2688080				

Figure 4: Fast Lookup: Equality predicate success

Executor	Statistic	Value
SQL Developer	memopt r NO IM tasks accepted	4
SQL Developer	memopt r populate tasks accepted	2
oracle@localhost.localdomain (W000)	memopt r blocks populated	457
oracle@localhost.localdomain (W000)	memopt r populate	1
oracle@localhost.localdomain (W000)	memopt r puts	74314
oracle@localhost.localdomain (W000)	memopt r rows populated	74314
oracle@localhost.localdomain (W000)	memopt r successful puts	74314
oracle@localhost.localdomain (W001)	memopt r cleanup	1
oracle@localhost.localdomain (W001)	memopt r entries deleted	32817
oracle@localhost.localdomain (W002)	memopt r blocks populated	34734
oracle@localhost.localdomain (W002)	memopt r puts	5660881
oracle@localhost.localdomain (W002)	memopt r puts:buckets full	681453
oracle@localhost.localdomain (W002)	memopt r rows populated	5660881
oracle@localhost.localdomain (W002)	memopt r successful puts	5660881
oracle@localhost.localdomain (W002)	memopt r successful puts:with evictions	681453
oracle@localhost.localdomain (W002)	memopt r tag collisions	104
oracle@localhost.localdomain (W007)	memopt r cleanup	1
oracle@localhost.localdomain (W007)	memopt r entries deleted	2800857

Figure 5: Results of query for MEMOPT R statistics

```
SELECT
    P.program "Executor"
    ,N.name "Statistic"
    ,S.value "Value"
 FROM
    v$sesstat S
    ,v$statname N
    ,v$session P
 WHERE S.statistic#  = N.statistic#
    AND S.sid = P.sid
    AND N.name LIKE '%memopt r%'
    AND S.value > 0
 ORDER BY P.program, N.name;
```

Listing 12: Querying against dynamic views to discover MEMOPT R statistics

entries deleted statistic indicates rows were removed from the MORS whenever I executed the **DROP_OBJECT** procedure.

Wrapping Up: Considerations for Using Fast Ingest and Fast Lookup

Hopefully the business use case and solutions I've provided in this article series will help you to decide if Fast Ingest and Fast Lookup offer appropriate strategies for improving the performance of your application's workloads. Since Fast Ingest is definitely the most convoluted of these two feature sets, I've summarized what I see as its benefits and drawbacks in ***Table 2*** below. Of course, please be aware that as Fast Ingest is undoubtedly improved in the future, many of the drawbacks I've listed may diminish or even disappear.

NOTE: This article was previously published in ODTUG TechCeleration in July 2021.

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**
- Understanding Fast Lookup: **https://docs.oracle.com/en/database/oracle/oracle-database/19/tgdba/tuning-system-global-area.html#GUID-E46EF11C-E999-4277-950F-E78EEC895ABB**
- MOS Note on Fast Ingest Limitations: **https://support.oracle.com/epmos/faces/DocContentDisplay?&id=2605883.1**

Feature	Advantages	Drawbacks
Ingested data is captured in batches within the Large Pool, but not immediately written to the database	Ingesting data is quite fast, and huge volumes of data from numerous sessions can be captured with extreme efficiency because the database isn't processing individual rows	Should the database instance crash before all ingested data is written to the database, *it is possible to lose data*
Since Fast Ingest is not a transaction in the traditional Oracle sense, COMMITs don't occur within its context	"Normal" Oracle transaction mechanisms are bypassed to enable rapid data capture	• No COMMITs mean no ROLLBACKs, either! • Parent-child transactions must be coordinated to ensure against data loss • Data cannot be queried until it's actually been flushed to disk from Fast Ingest buffers
Index operations and constraint checking only happens when data is finally written from the Large Pool Fast Ingest area to disk	Not necessarily a bad thing!	Should a primary key violation occur while the "drainer" background processes are writing data from the Fast Ingest buffers to disk, the violating rows won't be INSERTed, but *no exception will be raised*
Unless the insert fails for valid reasons, the application itself must ensure that all valid data have actually been inserted	**DBMS_MEMOPTIMIZE** procedures can be called to verify all data has been written successfully to the database	The application itself must now ensure all data has indeed been written to the database

Table 2: Fast Ingest: Benefits and Drawbacks

About Jim Czuprynski

Jim has over four decades of professional experience in information technology throughout his career, 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 often at Oracle OpenWorld, Oracle CODE, COLLABORATE / INSYNC, KSCOPE, Oracle Development Community tours, and Oracle User Group 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 published on databasejournal.com, IOUG SELECT, ODTUG TechCeleration, and OracleWorld. He has also co-authored four books on Oracle database technology. He also hosts a biweekly podcast, **Beyond Tech Skills**, with his colleague Liron Amitzi that focuses on everything about tech – except tech.

Carsten Czarski

Oracle Application Express (Part 9): Location Matters - Introduction to the APEX Map Region

Location data, visualized on a map, is a common requirement for application developers today. And of course, this is also true for APEX applications. With APEX 21.1, released in May 2021, the *Map Region* was introduced.

This new component allows developers to add maps to their applications, without coding, in a fully declarative fashion: After providing a data source, as a SQL Query, Table or REST Data Source, and configuring some visualization attributes, the map region displays a modern and interactive map within the APEX application.

This article provides an overview about the Map Region ***(Figure 1)*** and its capabilities, and particularly highlights its usage in combination with the *Spatial SQL functionality of the Oracle Database*.

Map Regions and Layers

In a nutshell, the map region displays spatial data (coordinates) on a background map. While the background map is served *to the browser* from the **Oracle Maps** cloud service **(maps.oracle.com)**, the spatial data comes from the data source configured in map region settings, i.e., from the APEX database or an external REST Source. All rendering happens in the browser.

Like other APEX components, the map region is fully declarative; the most straightforward way to create the first map is the Create Page wizard ***(Figure 2)***.

The Map region allows to choose from five pre-configured background maps ***(Figure 3)***. Support for custom background maps is planned for a future release.

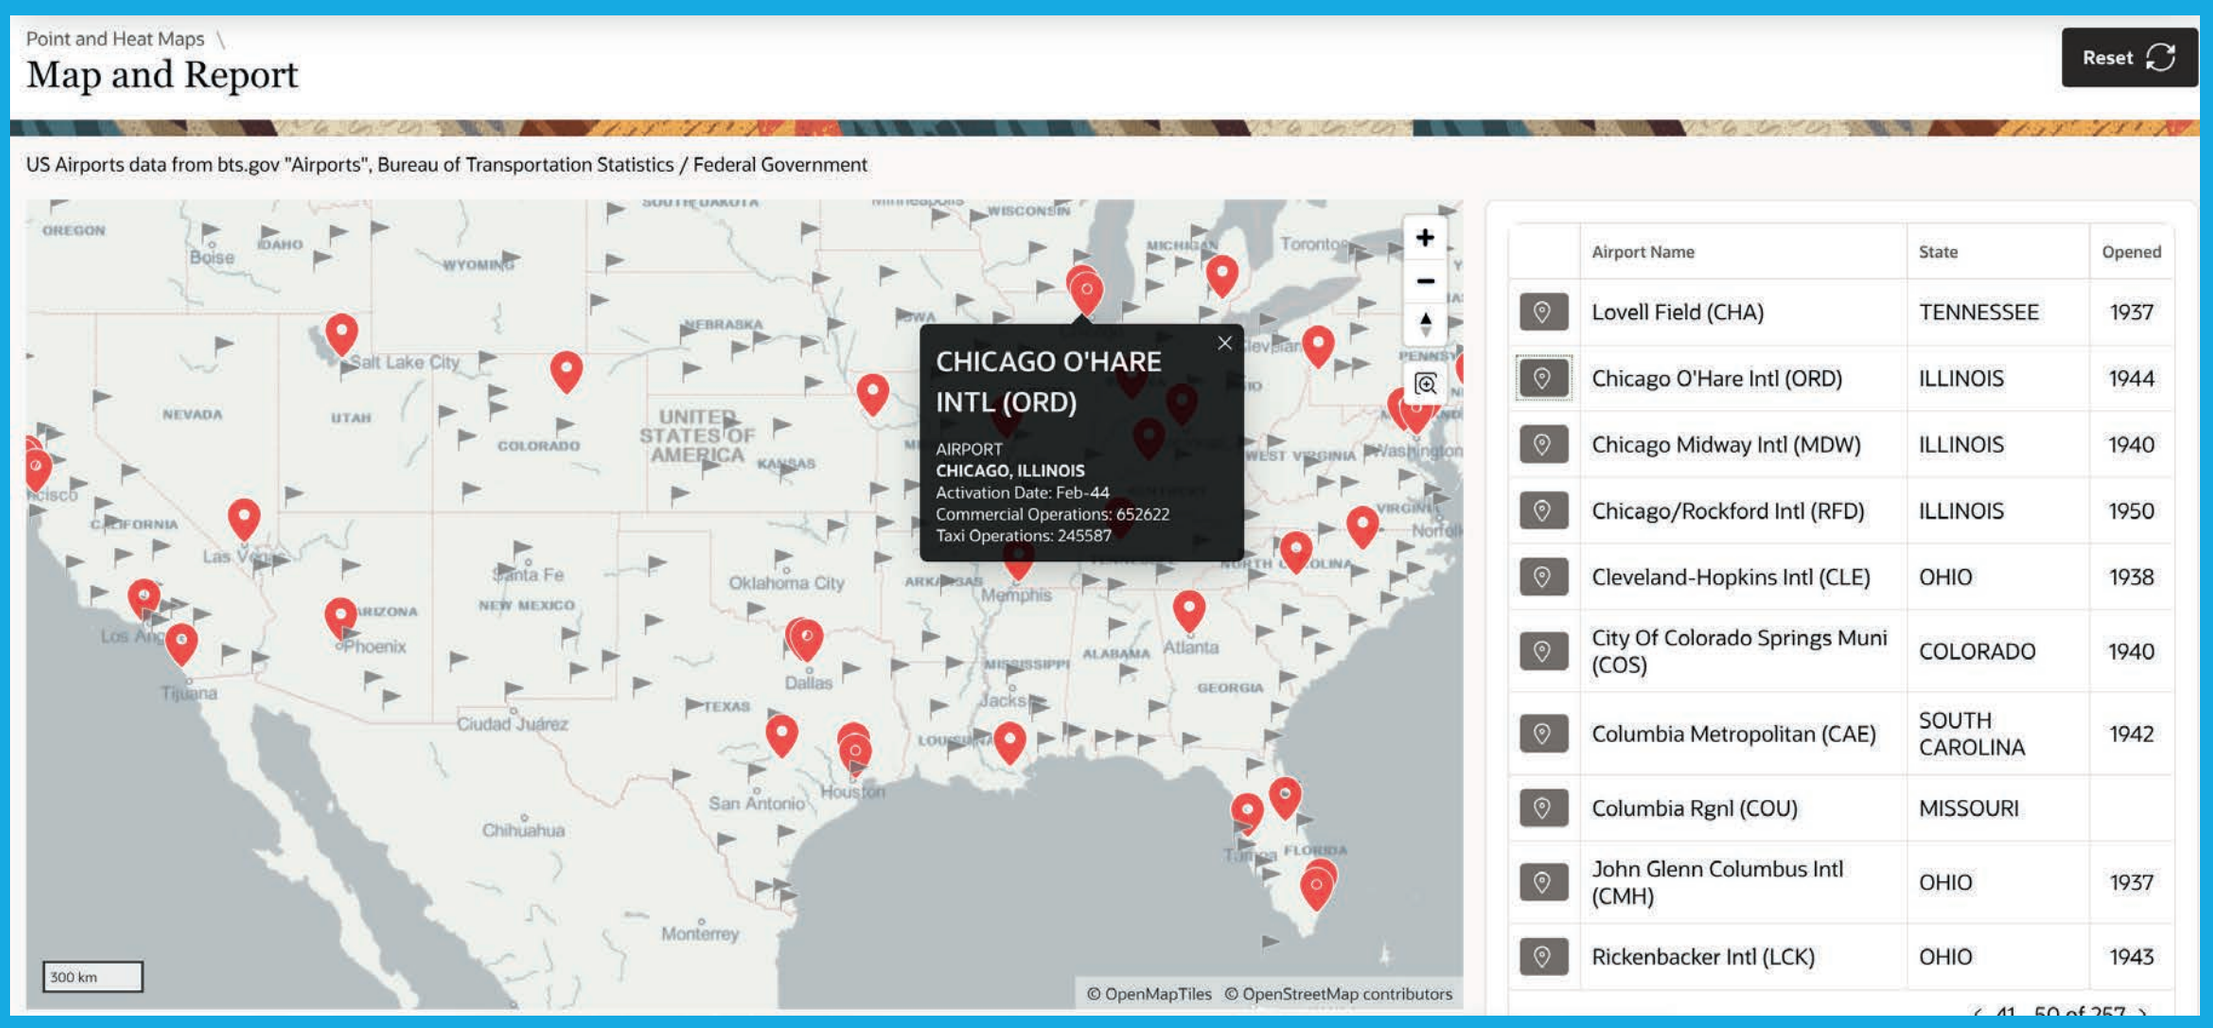

Figure 1: The new Map Region in APEX 21.1.

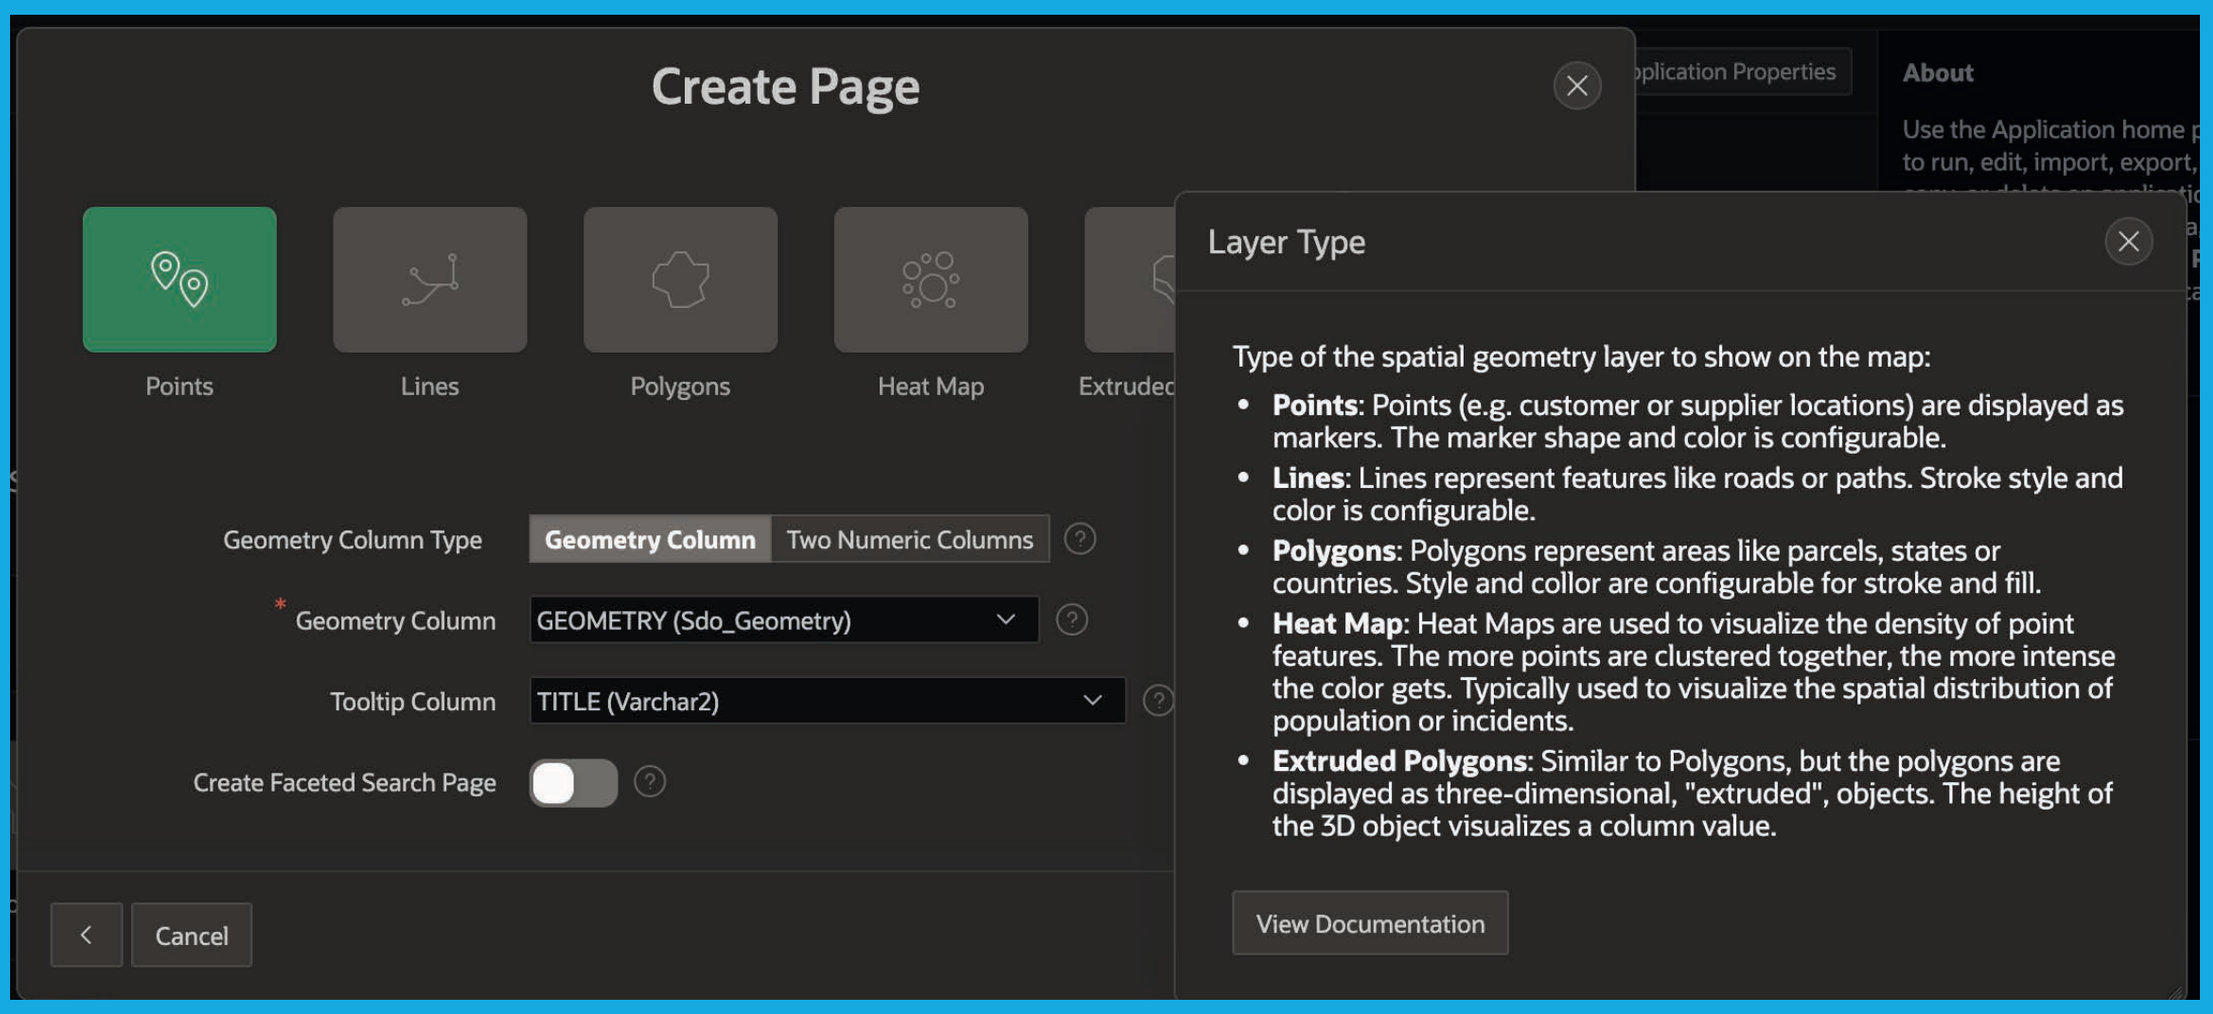

Figure 2: Create a new Map with the APEX Create Page wizard.

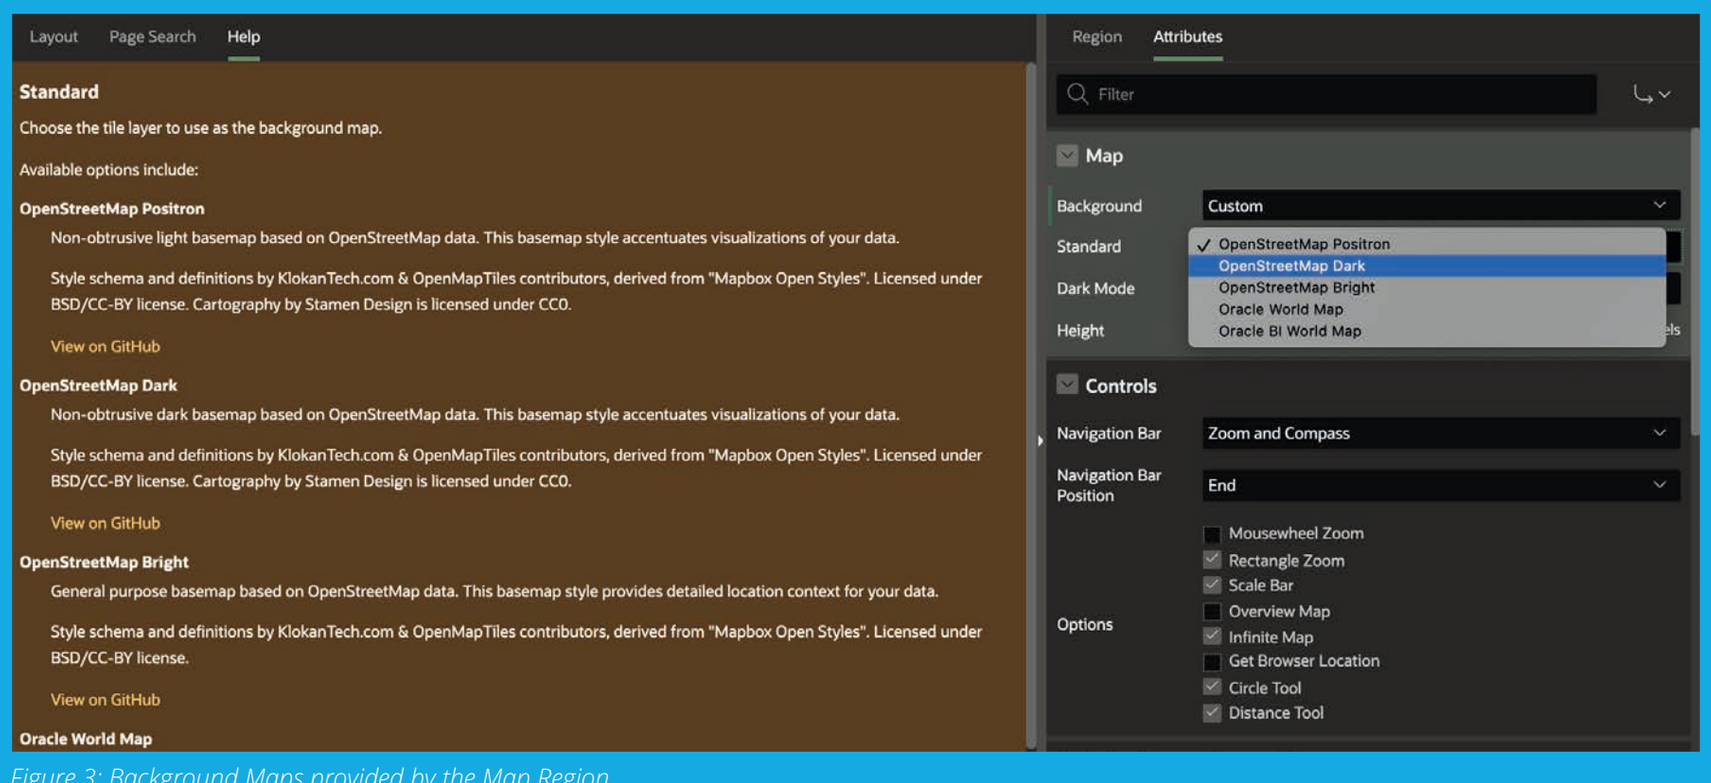

Figure 3: Background Maps provided by the Map Region.

Spatial data, from the map region data source, is displayed as *layers* on top of the background map. The map region supports the following visualization types:

- Points
- Lines
- Polygons
- Extruded (3D) Polygons
- Heat Maps

For Point and Heat Map layers, the map region can work with two NUMBER columns, containing longitude and latitude data. A single column, containing GeoJSON or the Oracle **SDO_GEOMETRY** data type, can be used as an alternative.

Using SDO_GEOMETRY is strongly recommended, primarily because it allows to create a *spatial index* and to use powerful "Spatial SQL" queries. For GeoJSON, it's important to know that the map region only expects the "geometry" part of a GeoJSON object; all other attributes must come from normal query columns.

The best visualization type depends on the actual spatial data and the use-case. For line and polygon layers, the spatial data must represent actual lines or polygons, of course. As data consists of simple longitude and latitude coordinate pairs in most cases, the first examples will be *Point* ***(Figure 4)*** and *Heat Map* layers.

Point layers are most common and visualize spatial objects ("Feature") using *markers*. The APEX map region contains 21 built-in markers but also allows to provide SVG definitions for additional custom markers. *Font APEX icons* can also be used, as an alternative to SVG markers.

Marker colors and sizes are configured declaratively but can also be derived from query result columns, using standard APEX substitution syntax, for instance **&COLOR_COLUMN**..

Information Windows, which display when clicking a marker, and *Tool Tips* (display when hovering over a marker) can be used to show additional information about the spatial feature. These popups are fully customizable with custom HTML markup, column substitution syntax and CSS classes. *Template Directives*, which are known from the Cards Region and from Interactive Grid, allow even more customized layouts.

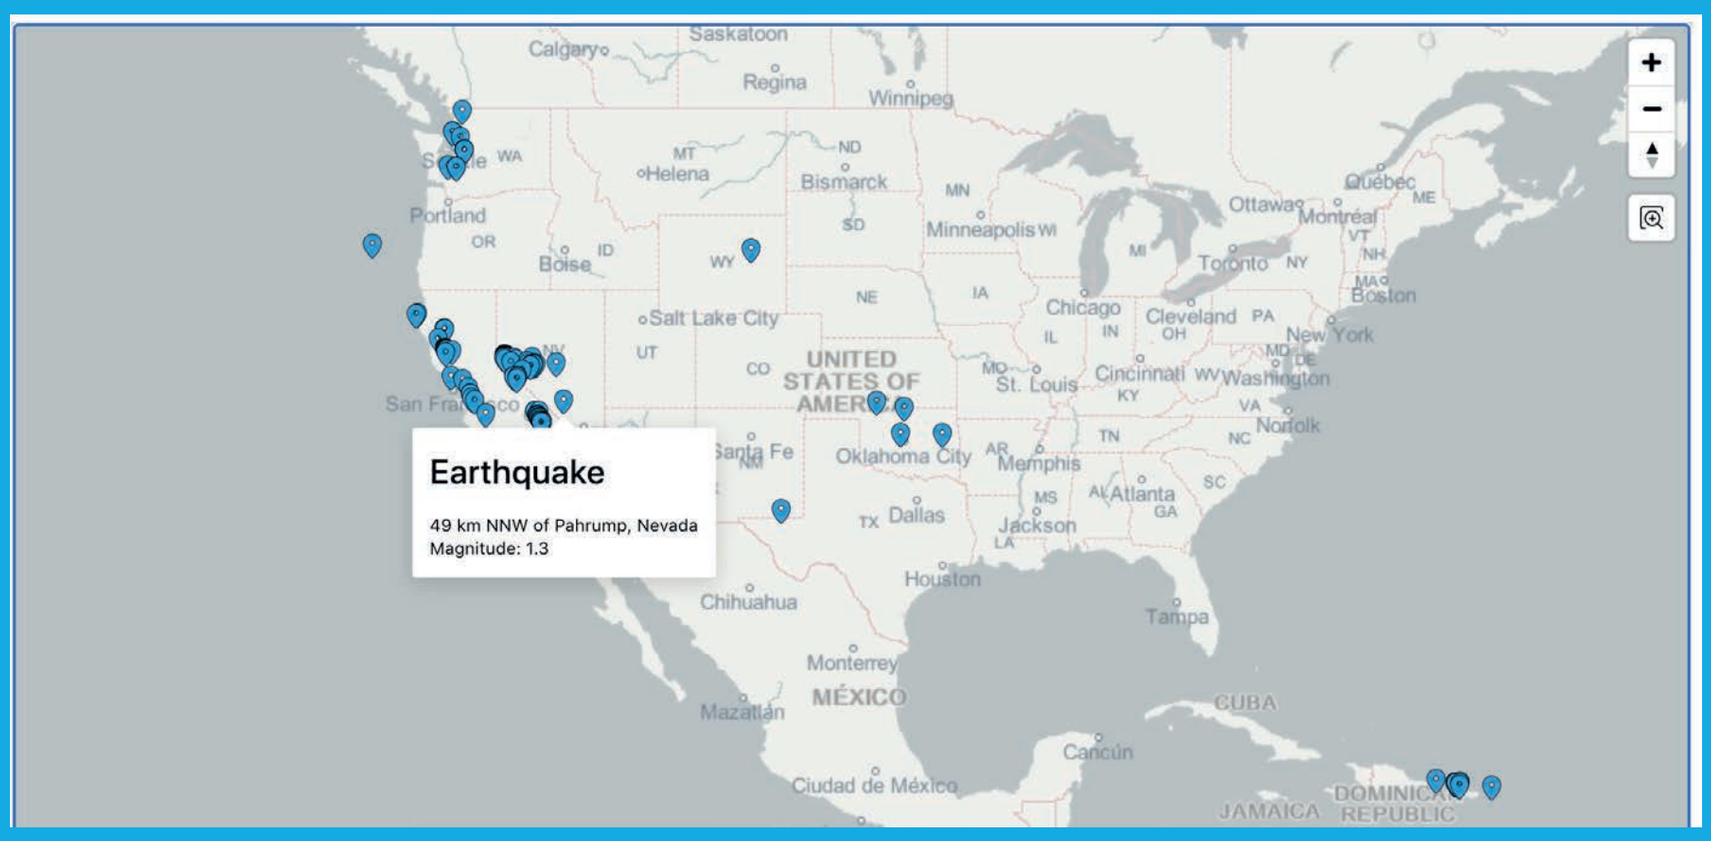

Figure 4: A Point layer allows to interact with individual features.

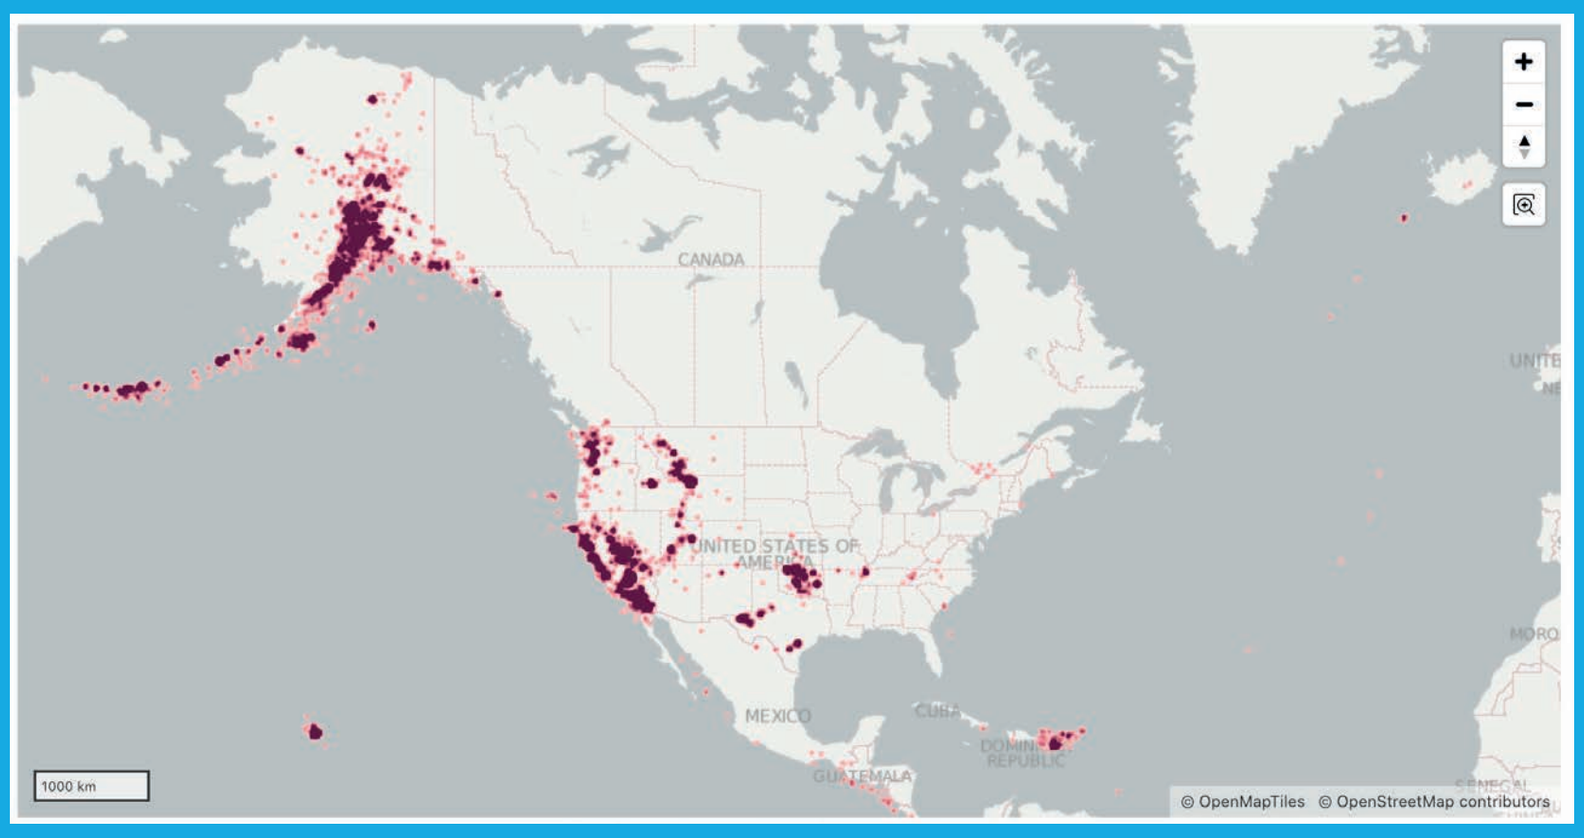

Figure 5: A Heat Map Layer is useful to visualize spatial density.

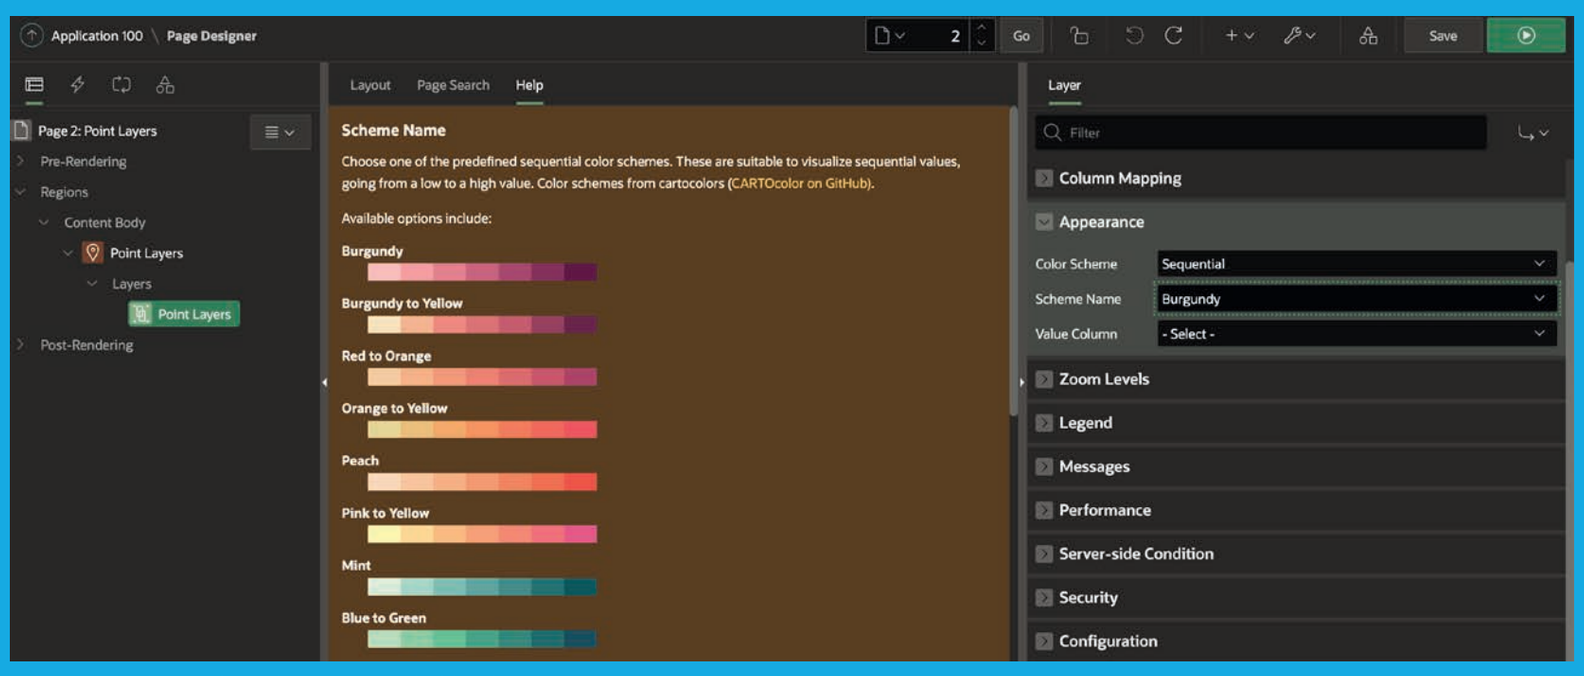

Figure 6: Built-In color schemes provided by the APEX Map Region.

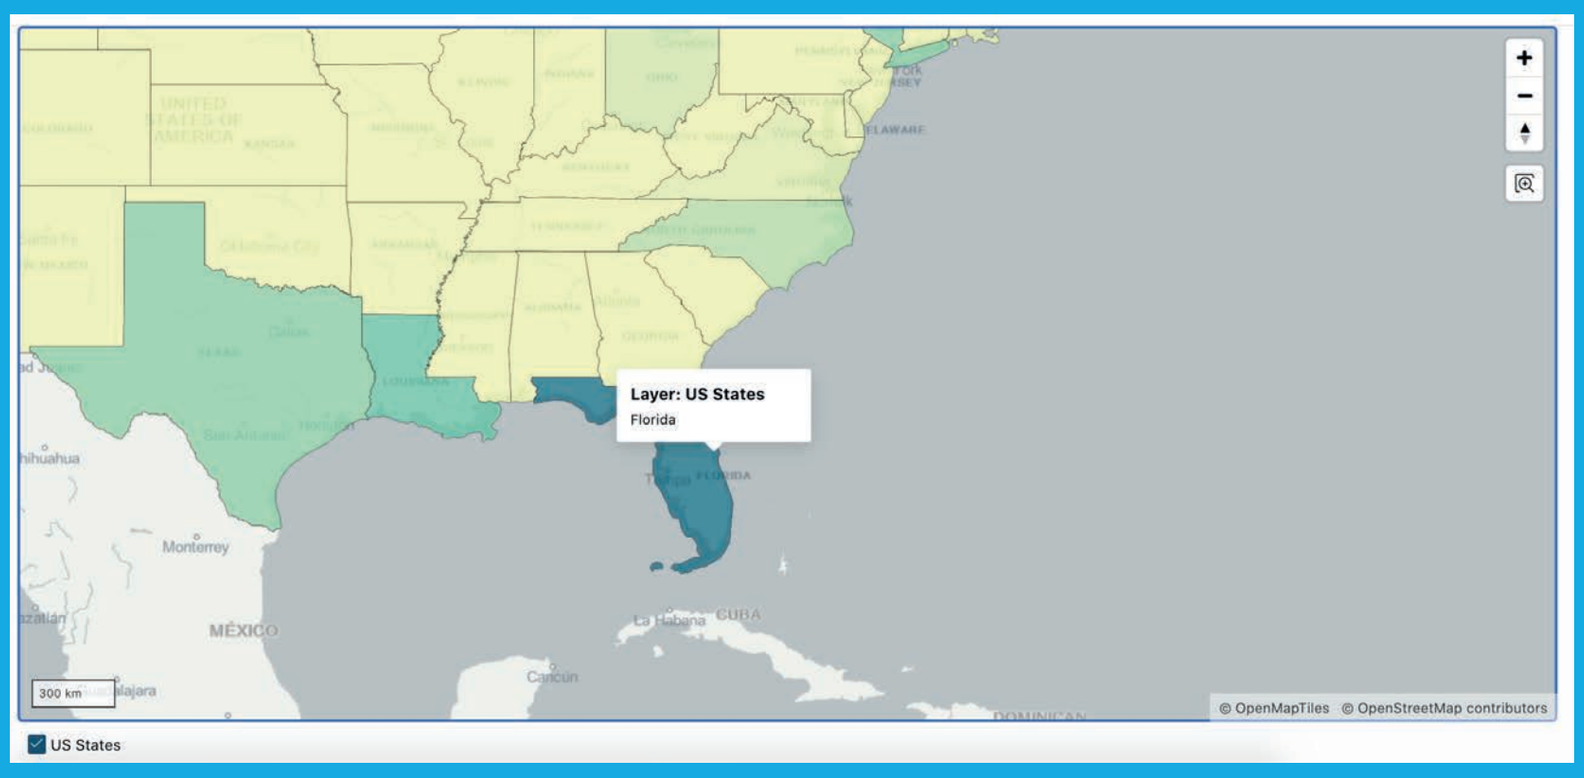

Figure 7: Thematic Maps with APEX: Visualize US States by "Water Area".

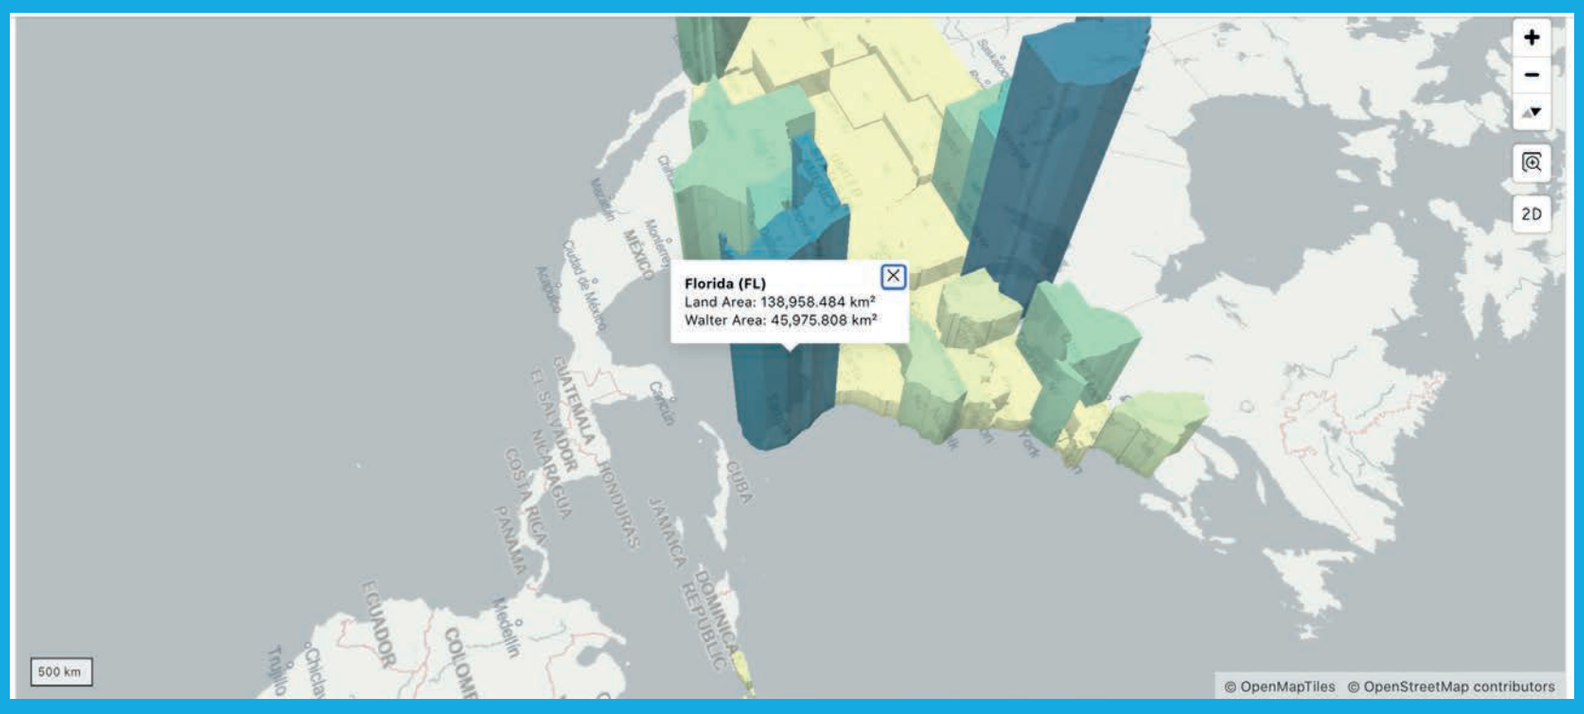

Figure 8: Extruded Polygons use the height of an object to visualize a value.

For each layer, the map region allows to specify a minimum and maximum zoom level. Zoom levels range from 1 (fully zoomed out) to 18 (zoom to street level). A layer is only displayed when the minimum zoom level has been reached and does not display any more once beyond the maximum zoom level. This is very useful to turn off layers when the current map zoom level is not appropriate for the layer to convey useful information.

Heat Map ***(Figure 5)*** layers are great for visualizing the *density* of spatial features. Unlike with Point Layers, end users cannot interact with individual features here: information windows or tool tips are not available.

However, based on the used color scheme, the user gets an impression about which areas a very densely "populated" with spatial features, and which (other) areas contain only a few sparsely distributed features. The Map Region contains a set of built-in color schemes, but also allows to provide a custom list of colors ***(Figure 6)***.

Polygon layers are used to visualize areas and their attributes. The polygon color can be set declaratively in Page Designer but can also be derived from a result column (using substitution syntax), or from a numeric attribute and a *color scheme*. ***Figure 7*** illustrates such a thematic map: it shows US states, with the area color derived from the amount of "water area" in the given state. Information Windows and Tool Tips are available for Polygon layers as well.

A special kind of polygon layer is the *Extruded (3D) Polygon* type. In addition to displaying the colored polygon, it allows to *extrude the polygon to a height*, derived from another query result value. ***Figure 8*** illustrates how such a layer can look like.

Of course, a map region can have multiple layers, which are displayed on top of each other ***(Figure 9)***. Each layer can even use a totally different data source: For instance, a layer based on a REST Data Source can sit on top of a layer based on a SQL Query. If the developer allows with the **Enable to Hide** attribute, end users can show and hide layers by clicking on the name in the legend.

Spatial SQL with the Map Region

The Map Region becomes even more powerful when combined with Spatial functionality in the database. Since Oracle Database 18c, there is no longer a *Spatial & Graph option*, so once the Oracle Database is properly licensed, Spatial SQL functionality can be used without having to worry about additional licenses.

Assumed, we have tables containing Earthquake and US States data, it's just a matter of a SQL Query to find out the earthquakes happening in a specific US state. And map region can nicely visualize these query results.

As a prerequisite, make sure that tables use the **SDO_GEOMETRY** data type, and that there is a spatial index created for the SDO_GEOMETRY column. Starting with APEX 21.1, SQL Workshop Object Browser contains a

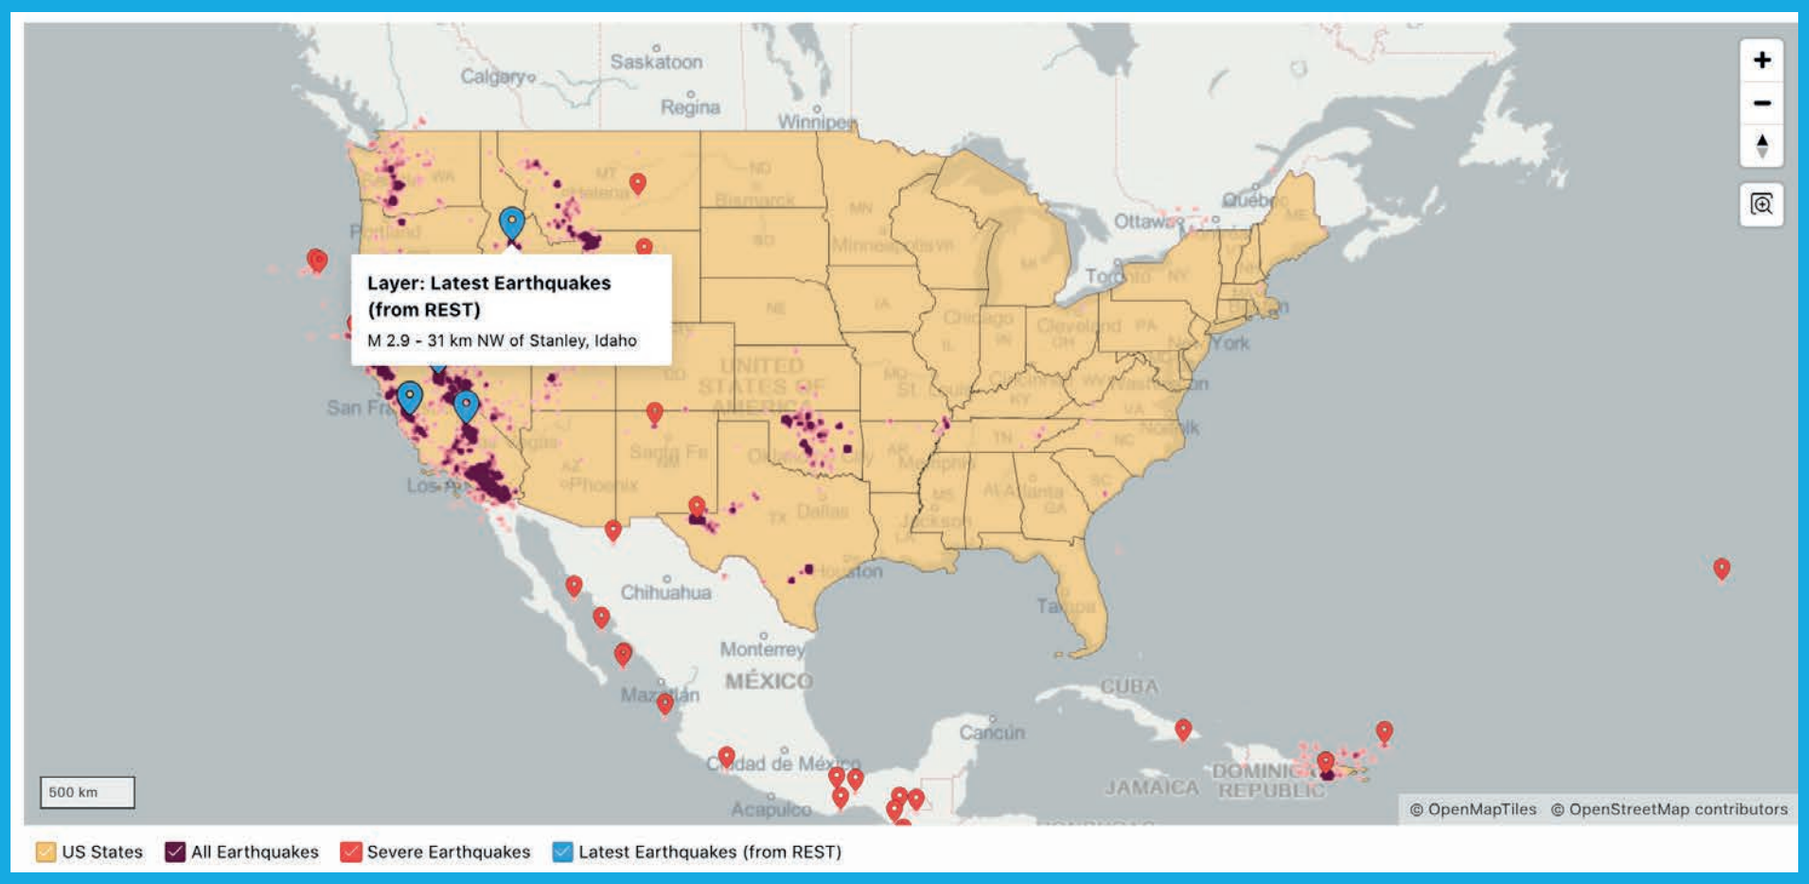

Figure 9: Multiple Layers in a Map Region.

new Create Index wizard, which helps with this task ***(Figure 10)***.

Once the spatial index is created, we can craft a SQL query ***(Listing 1)*** to find all earthquakes, which affected the state of California: The **SDO_ANYINTERACT** SQL function does the required spatial computations for us (and uses the spatial index for query performance). The Map Region allows easy visual inspection as shown in ***Figure 11***.

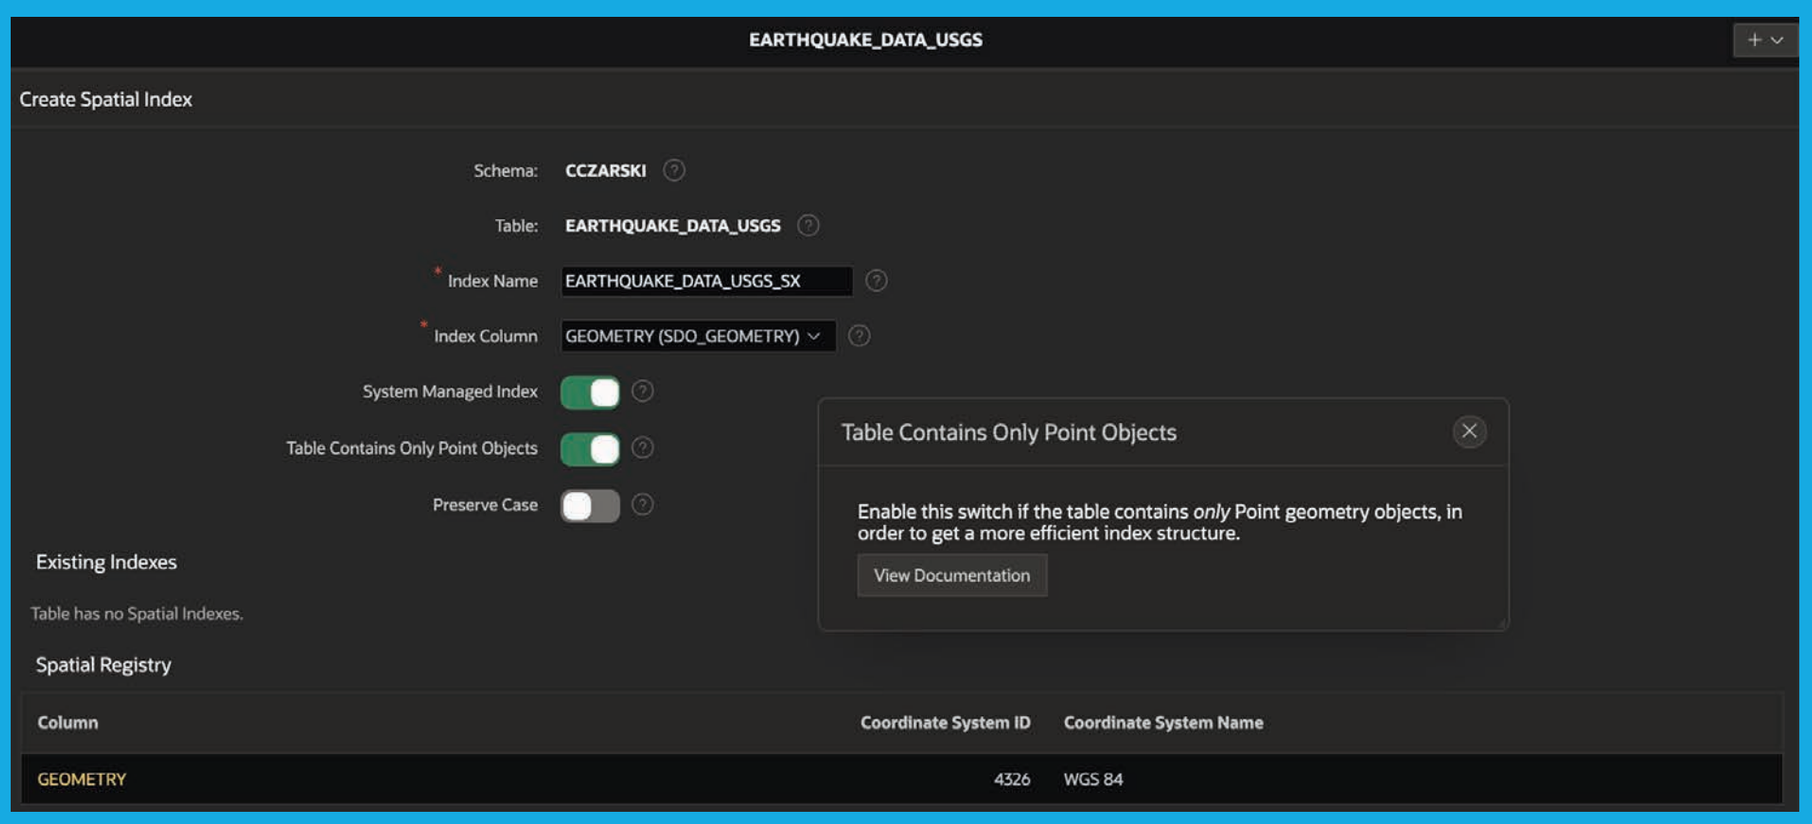

Figure 10: Create a Spatial Index with SQL Workshop.

```
select e.id,
       e.geometry,
       e.title,
       s.name
    from earthquake_data e,
       us_states_data  s
 where sdo_anyinteract( e.geometry, s.geometry ) = 'TRUE'
    and s.name = 'California'
```

Listing 1: SQL Query finding all Earthquakes in California.

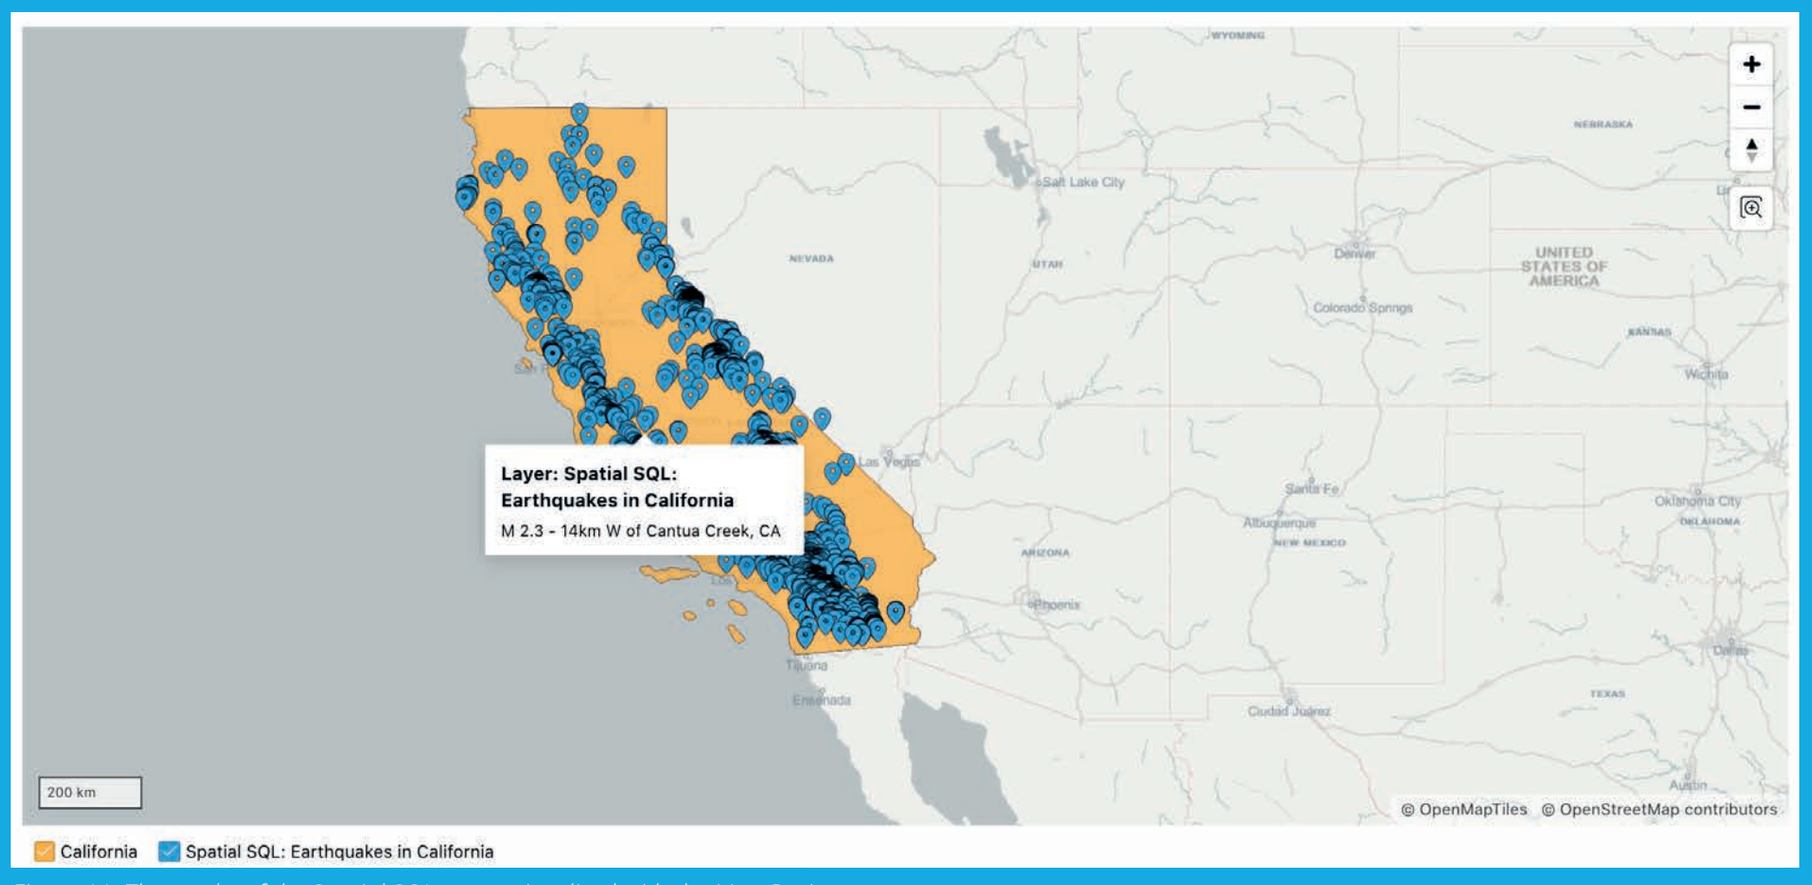

Figure 11: The results of the Spatial SQL query, visualized with the Map Region.

SDO_WITHIN_DISTANCE is another example for powerful Spatial SQL: As the name indicates, it finds objects within a given distance to a specified point. ***Listing 2*** contains an example, and ***Figure 12*** shows how the results look like.

The SDO_NN function can find Nearest Neighbors, based on a given coordinate, and based on additional query filters in the SQL WHERE clause. The more restrictive the filters become, the farther away the returned neighbors are.

Figures 13 and 14, which were taken from the Sample Maps packaged application, show how results of a Nearest Neighbor search can be displayed with the Map Region.

```
select a.id,
       a.geometry,
       a.title,
       s.name
    from airport_data   a,
       us_states_data s
  where sdo_within_distance(
         a.geometry,
         sdo_geometry( 2001, 4326,
              sdo_point_type( {lon}, {lat}, null ),
              null, null ),
         'distance=50 unit=km' ) = 'TRUE'
```

Listing 2: Spatial SQL using SDO_WITHIN_DISTANCE.

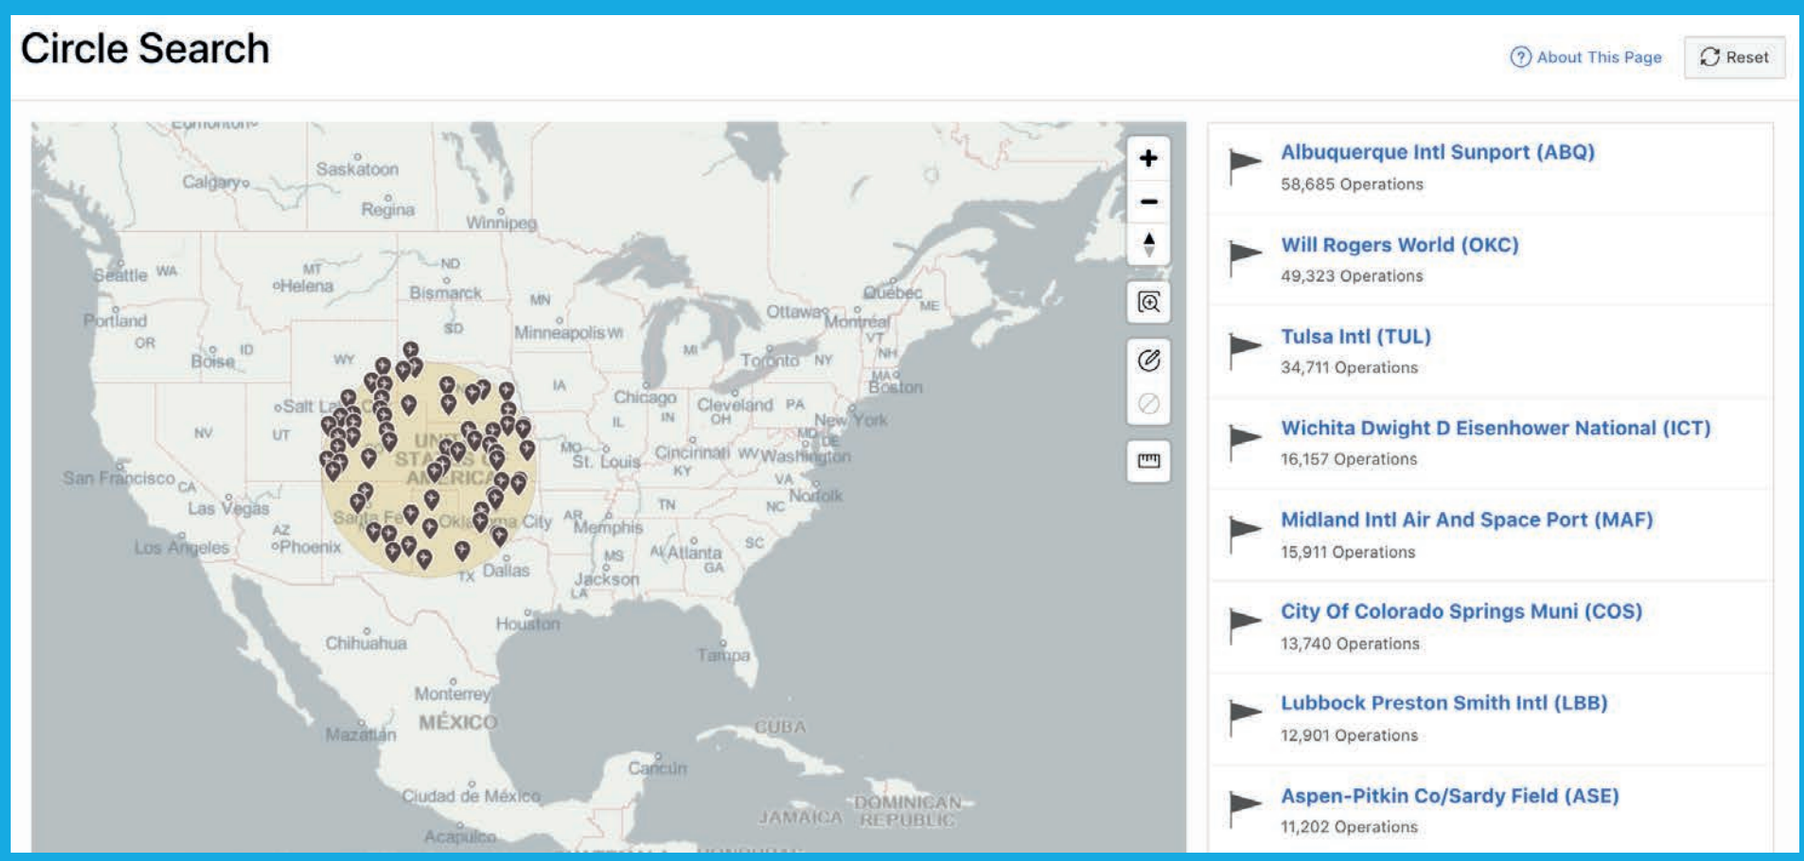

Figure 12: Circle Search with the SDO_WITHIN_DISTANCE function.

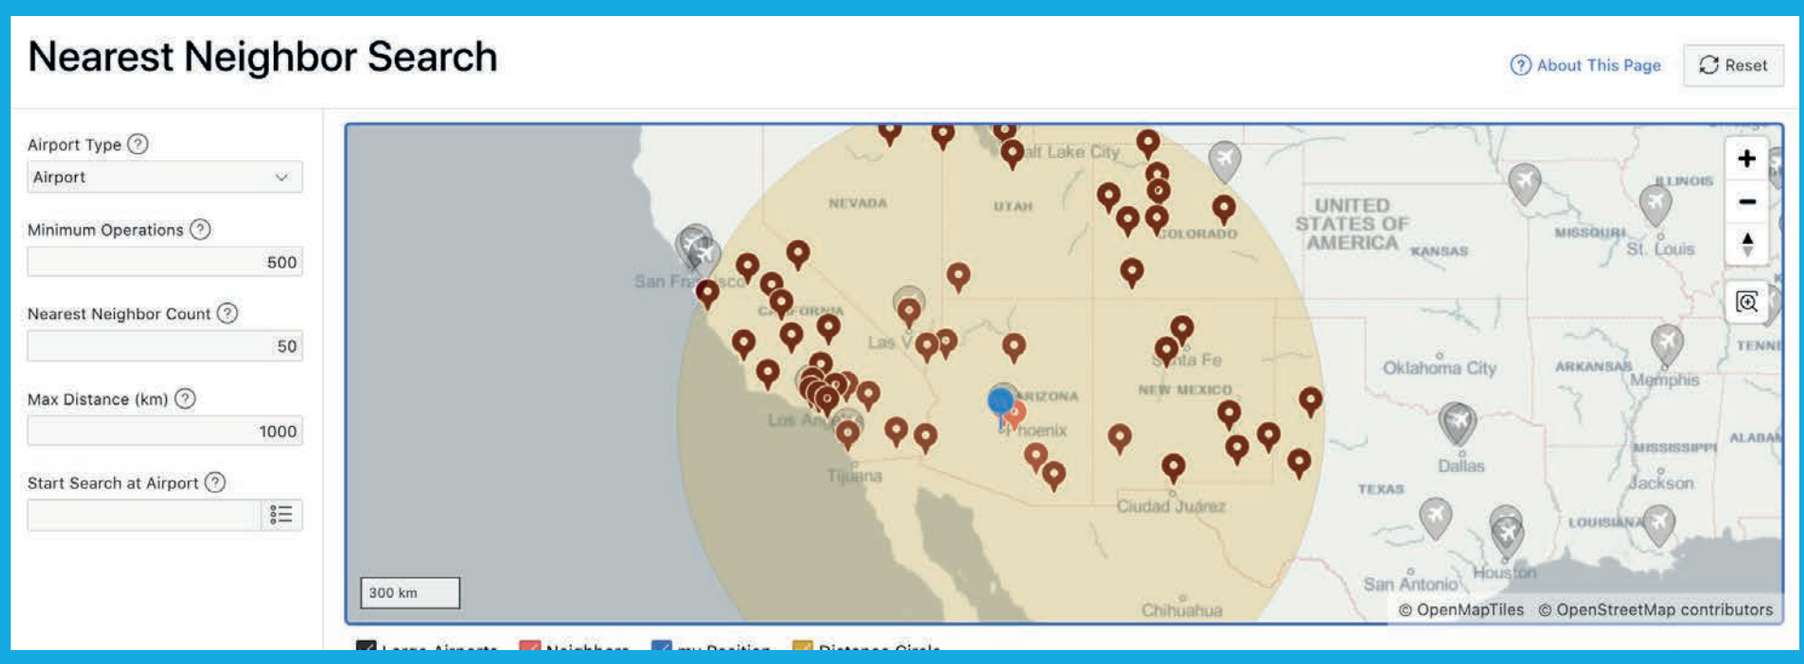

Figure 13: SDO_NN with a lax filter finds many nearest neighbors.

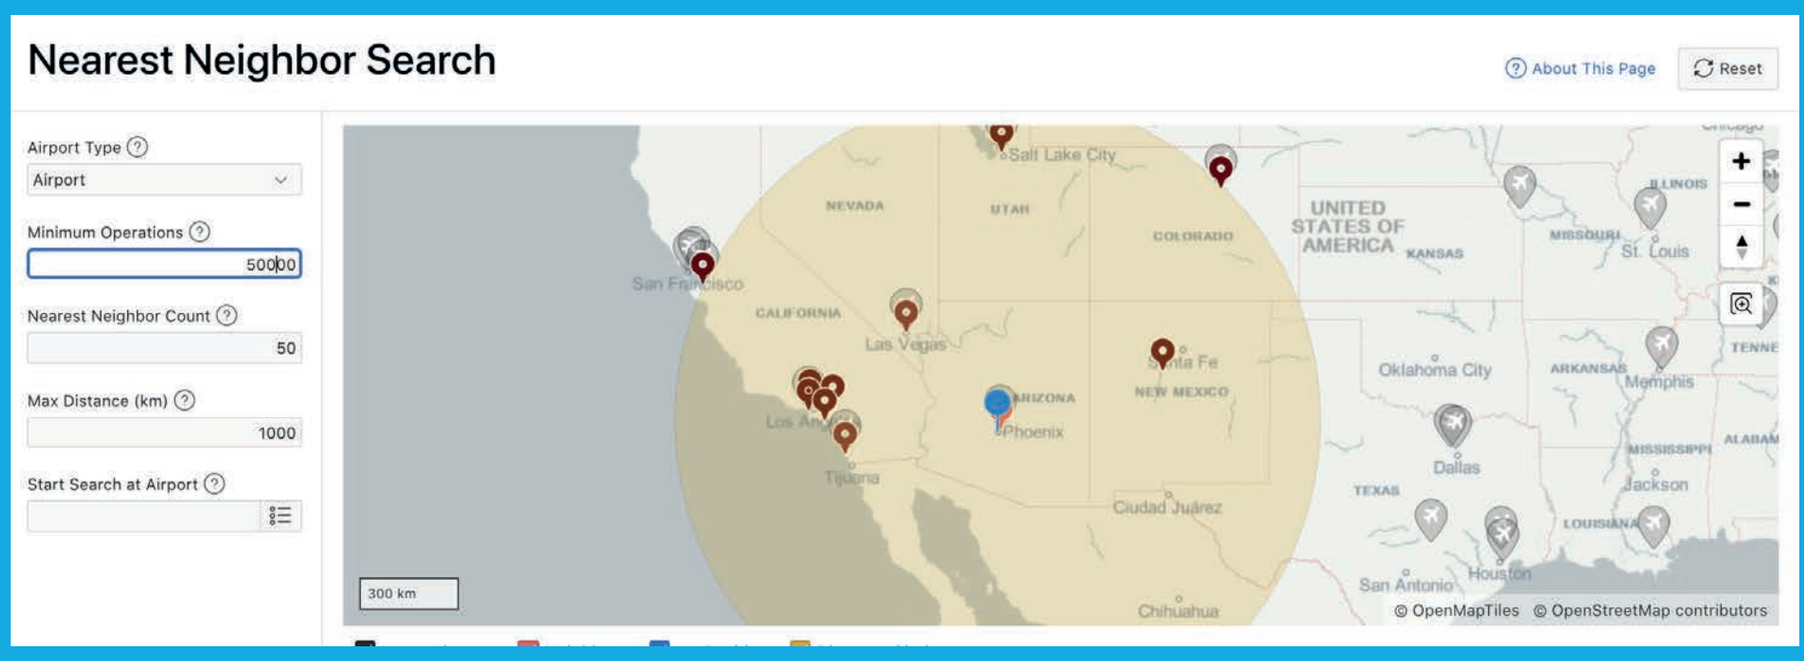

Figure 14: With a stricter filter, nearest neighbors are farther away.

Summary

Using Maps in an APEX application has never been easier than with APEX 21.1. It only requires a SQL Query or REST Source to get spatial data from - the rest is done by the new Map Region. Map display attributes like colors, markers, information windows or tool tips are declaratively configured in Page Designer. Maps are added to an APEX application in minutes.

But most value comes from the combination with spatial capabilities of the Oracle Database. Developers can do more than simply selecting all rows from a table - powerful spatial functionality like Within Distance or Nearest Neighbor Search are available as SQL functions, ready to use.

For most readers, the Sample Maps packaged application, which is available on GitHub, should be the best starting point to explore the capabilities of the Map Region. Simply download the application from GitHub and import into the APEX 21.1 environment. The application not only contains simply visualizations, but also examples for usage of Spatial SQL.

More Information

- Information and environment to test-drive APEX: **http://apex.oracle.com/en**
- Oracle Application Express Blog: **http://blogs.oracle.com/apex**
- Packaged Application "Sample Maps" on GitHub **https://github.com/oracle/apex/tree/21.1/sample-apps/sample-maps**

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.

Another Year, Another OGBEMEA Virtual Tour

Mirela Ardelean

That time of the year will arrive. No, I am not talking about Christmas (yet 😉). I am talking about the Fall, the "Conference Season", when it's time to hit the road and … and … Unfortunately, no, not this year.
But the season, the "Conference Season"? Travel restrictions…
So, what are we going to do? Simple: We are organizing the Oracle Groundbreakers EMEA Tour (OGBEMEA), as we did for the last 6 years. And we'll do the tour virtual … again.

We, the organizers team, had this kind of discussion some months ago. Now, we are in the middle of selecting topics for OGBEMEA Virtual Tour 2021 and, by the time you will read this article, the agenda will be finished and published on **ogbemea.com**.

How will the OGBEMEA Virtual Tour 2021 look like?

The "recipe" used last year was successful, so we are following it again: 10 Oracle User Groups (OUGs) are collaborating under EOUC coordination to build an interesting agenda and to deliver a marathon of more than 10 days with webinars and labs. Everything for free, everything as a result of voluntary work of organizers, OUGs and speakers.

Who are the 2021 participants? Have a look at the impressive list of OUGs and countries:

Azerbaijan - AzerOUG
BeNeLux - OBUG
Croatia - HROUG
Kingdom of Saudi Arabia - KSAOUG
Kyrgyzstan - CASOUG and TajOUG
Luxembourg - LuxOUG
Norway - OUGN
Switzerland - Swiss OUC
Tajikistan - TajOUG
Uzbekistan - CASOUG and TajOUG

Imagine someone should have to travel to all these 10 countries in 2 weeks. It would have been almost impossible! But in a virtual tour, this imaginary trip is doable.

What to expect from OGBEMEA Virtual Tour 2021?

We are doing our best to build an agenda covering a large area of topics related to Oracle technologies. Some very known names will be recognized on the agenda.

But this is not all: The OGBEMEA Tour wants to welcome new presenters on its virtual scene. This is a result of collaboration between OGBEMEA Tour organizers and the MASH program (MASH stands for Mentor and Speaker Hub; read more about the program in our previous ORAWORLD issue and on mashprogram.wordpress.com). We are keeping fingers crossed for every newbie speaker who will present a topic on this virtual tour.

This year's agenda includes also a live tour (how does this sound: a live tour within a virtual tour).

The live tour will take you to the Oracle Industries Innovation Lab and, from what I have seen already, it will be like a trip into the future.

This is only a short highlight of what you can expect to happen between October 15 and 29. Check the agenda and register to as many sessions as you can. It's free and it is online.

OGBEMEA Virtual Tour is built by the community for the community. Don't miss it!

Last but not least, I would like to thank my organizer teammates (Kamran Aghayev, Janny Ekelson, Jure Bratina, Rodrigo Mufalani, Rustam Khodjaev) and all the 10 OUG boards working together to set up this huge virtual event!

Interview by Christian Luda

Tech Pub Quiz: “It's Fantastic Fun, and That's the Only Goal”

During the pandemic user groups have come up with interesting online formats – one that sticks out is the Tech Pub Quiz started by the Swiss Oracle User Community. During the summer break ORAWORLD had the pleasure to talk to the quiz' organizers Julian Frey, Flora Barriele, Jasmin Fluri, and Christian Berg.

How did you come up with the idea of doing the "Tech Pub Quiz"?

Julian: Before, I was doing meetups with the Oracle Beer Bern group. And last year, every one of us had online conferences, and we stopped enjoying attending online presentations. Before Covid, I was a frequent attendee of pub quizzes – in pubs and in-person. Somehow, I missed it and thought: "Is it possible to hold something like this virtually on a tech topic?" We then started in December with planning the first one. In February, we did the "Oracle Beer Pub Quiz Edition", but we also had in mind opening it up for all tech, which was the outcome of the actual first "Tech Pub Quiz" in March.

For people who haven't joined yet, how does the "Tech Pub Quiz" work?

Julian: We decided that the fun part should be the main focus. It should be a fun evening, having some beers with some friends, and probably learn something – maybe something useless, but learn something. I think that's almost the only goal we have. We use Zoom and Kahoot as the quiz platform.

How about the others: Did you also have experiences in real-life pub quizzes?

Jasmin: I've never attended a pub quiz before. So, for me, it was a first. It was interesting to see how the people interacted because there's always a discussion besides the questions. And yeah, I had an entertaining impression when I first joined.

Christian: I did some before on different topics, so when this malarkey started last year, I also did some Kahoot sessions with friends. It's just a way to pass the time, have some fun and, as Jasmin said, talk between the questions. There can be some questions that raise, let's say, discussion potential, to put it mildly (laughs). It's geeks talking about tech questions, and depending on who writes the questions and how he writes the questions, you get reactions like, "Come on, my answer should count!" It's fantastic fun, and that's the only goal. To keep the community alive and keep people together, at least virtually seeing each other and talking to each other.

Flora: On my side, I'm a big fan of pubs and a big fan of quizzes, and I have attended a few pub quizzes before. What impressed me with the "Tech Pub Quiz" online is that you have almost the same feeling and atmosphere as in regular pub quizzes. It's always interesting to see which questions raise discussions: Even the tiniest topics can lead to something huge, and conversations continue even afterwards.

Do you have an example of such a question?

Flora: I have a silly example and a serious example. The silly one – I don't know how we started with Pokémon questions, but it became crazy (everybody laughs). We had a group of people totally against Pokémon questions and another group of people defending them. And the serious example: Sometimes we have a question about a feature, for example, an Oracle feature, but we forget to precise the version. So, an extensive discussion starts depending on the release number, the Oracle edition, or a hidden parameter that doesn't work the same way. So, it's technically interesting.

Jasmin: We also discovered that it's tough to phrase precise questions that don't raise this kind of uncertainty with the answers. We are limited with the length of the questions and the answers. You cannot put endless texts into those fields. Also, participants need to have time to read the questions before they answer and usually, there's a timer that goes on to limit the time you have to answer.

Christian: If you give them too much time to answer, the whole thing becomes just tedious; it goes on and on and on. If you don't provide them with enough time, they sometimes have trouble reading the answers. Especially questions that go very much into detail, like really going into operating system versions. Pokémon questions are a lot easier, but as Florence said, there are two camps. We did them only in the first edition, and we stopped because people kept complaining. And now there's a whole group of people saying, "Hey, we want them back!". But we haven't done it. There's endless potential with that.

So, the questions come partly from you, and then other people can also submit questions?

Julian: We have this submit form where we got over 90 questions, and the rest of the questions are provided by us and by friends of us.

Are there some people who have more fun submitting questions while others enjoy more answering them?

Julian: I think almost everyone who sent in questions complained how hard it is to write questions.

Christian: Whenever we have our internal question writing sessions, at least for me, I tend to do that late at night with the TV screen on in the background, just browsing around thinking, "What could we come up with, what did I see in projects?" or just completely random tech stuff: "Oh, I read this recently, and throw this in to throw people completely off track". But usually, it works well. We do have a reserved cache of questions that we haven't used yet.

Flora: I admire Jasmin and Julian because they are always very productive and have a lot of imagination on different topics. On my side, I need to search for inspiration. I sometimes open Oracle documentation and start browsing topics to get inspired. Otherwise, writing questions and especially writing funny questions is very difficult.

During the quiz, how much of a factor is the level of alcohol?

Christian: You can't deduce a correlation between the time passed in the quiz and the presumably amount someone has drunk. It's bizarre, we had a couple of statistical outliers where it was very late in the evening, and suddenly, somebody had like twice the amount of points as the best one in the round before. It's just random also with the mix of questions. Again, it's just about fun. And the good thing is people do tend to stick around for the entire length. That's the one thing that we have noticed. There's hardly any dropouts. We usually do six rounds with tech questions and one round with just geek questions, so basically being silly. And people tend to stick around, which is cool, which shows you they are having fun and are enjoying it.

Jasmin: What's also impressive is that some people only want to watch the quiz; they don't answer questions. They only watch other people answering questions and then see the results.

Julian: We should start streaming it on Twitch.

Christian: That's a good idea! But, actually thinking about this: No, it's not! I know that some people tend to get a question wrong and are like, "Oh my god, that's my area of specialty". It's super funny when you have one of the Oracle product managers sitting in there, and after the answer comes up, he's like, "Oh my god, what did I click", and you're like ", Nobody saw it, don't worry!" (laughs).

How many attendees did you have the first time, and how did it evolve? Did it also get more international?

Julian: On the first meetup, we were around 35 players, and then it evolved up to almost 50. International: Yes, we've had people from nearly all over the planet.

Christian: Some people are using a very long lunch break to join.

Julian: We had attendees from California, from Canada, Israel, all over Europe and also one from APAC.

So, are you meeting people you didn't know before?

Yes!

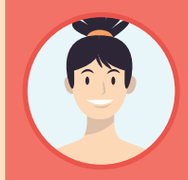

Jasmin: And people talking to each other that didn't speak before. They get to know each other.

Julian: We even grow outside of our circle. We called it "Tech Pub Quiz" and not "Oracle Pub Quiz" because we also wanted to target communities outside our Oracle bubble. For example, we had Google employees as well as SQL server specialists attending.

Julian: One story that comes to mind: One Google employee had participated in the third quiz for the first time, and then for the fourth quiz, he was in some meeting but was eager to attend the geek section at the end. While on the train, he already joined the Zoom call, rushed at home, and then I got a message from a friend who knows him quite well saying, "He's now on; we can start with the geek section".

Christian: That was cool, mainly because he's not from the hardcore Oracle community.

Do you have funny stories you would like to share, maybe someone getting drunk?

Christian: Honestly, the quiz is quite civil. People have fun, but at some point, their eagerness kicks in, and they think, "Actually, I'm trying to win here". Also, we've created our beer glasses which the winners are getting. So, at some point, people become competitive.

Flora: And, what I think is fair, when the same person wins several rounds, we don't send the glasses to them again, which gives space for other people to get them too. Whatever the results may be.

What are the plans for the "Tech Pub Quiz"?

Christian: To do it again! We all love it. It just shows how desperately we miss meeting each other, hanging out together, and having fun together.

Thanks everybody for taking your time.

Jasmin Fluri
Working at the development side of the database, Jasmin's focus is automation and how to build robust database applications. She's an Oracle ACE and besides working in consulting at Schaltstelle GmbH, teaches at the university of applied sciences.

Christian Berg
Working with Analytics since the day of nQuire's acquisition by Siebel Systems in 2001, Christian has 20+ years of business analytics and data engineering experience. He's an ACE Director Alumnus and currently works for Oracle as an Analytics Domain Specialist.

Julian Frey
Julian works as an Expert Oracle Database Consultant for Edorex AG in Ostermundigen (Switzerland). During his more than 10 years of consulting he has gained a lot of experience in many areas, especially in the areas of HA, Engineered Systems and Monitoring. He is a frequent speaker and founder of the Oracle Beer Bern Meetup Group. Julian is a member of the Oracle ACE Program as an Oracle ACE. Also is he a founding member of the SwissOUC User Group.

Flora Barriele
Flora is an Oracle Database Engineer and Oracle ACE. She's an open-minded technology enthusiast, constantly learning, who likes to connect with people and loves sharing knowledge experience, and challenges.

Events

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

October 15 - 29, 2021
online
http://ogbemea.com/

SPOUG 21

November 3, 2021
Madrid, Spain + online
https://spoug.es/spoug-21-english/

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 udruga 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): Mirela Ardelean

Editorial team:

Mirela Ardelean, Jean-Jacques Camps, Heli Helskyaho, 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**

Article submission:

If you are interested in submitting an article, please e-mail us your article via the online form at **www.oraworld.org**. Deadline for the next edition: **October 27, 2021**.

Photo credits:

Title: © created by freepik | https://freepik.com
P. 4: © floralset | https://123rf.com
P. 5: © Comic | www.commitstrip.com
P. 6: © created by freepik | https://freepik.com
P. 11 Source: https://devblog.axway.com/dev-insights/keep-calms-and-devops/
P. 11 - 14: © upklyak | https://freepik.com
P. 15: © rawpixel.com | https://freepik.com
P. 20 - 29: © AndSus | https://123rf.com
P. 30: © Anastasiia Nevestenko | https://123rf.com
P. 38: © macrovector_official | https://freepik.com
P. 39: © pch.vector | https://freepik.com
P. 40-42: © pikisuperstar | https://freepik.com
P. 43: © DOAG | www.doag.org/en/home

Graphic design:

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

Disclaimer: