April 2020, Edition #20

# ORAWORLD

e-Magazine for Oracle Users published by the EOUC

# How an App Helps to Heal our Planet:

## Interview with the Founders of fabe

- Oracle and Azure: The New Frontier
- The AskTOM Journey Part 2: Today and the Future
- Tips for Oracle Ace Associates Part 2

EOUC
EMEA ORACLE USERGROUP COMMUNITY

*How an App Helps to Heal our Planet: Interview with the Founders of fabe*

**page 14**

*Oracle and Azure: The New Frontier*

**page 34**


*What's Your Super-Power? Mine is Autonomous Database and Machine Learning (Part 2)*

**page 43**



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

# Editorial

Dear Readers,

For this issue of ORAWORLD magazine, Christian Luda interviewed people behind an amazing community-driven project: fabe (for all a beautiful earth). fabe is an APEX-based progressive web application (PWA) to help its users to slow the rate of planetary warming and to stop human-caused extinctions.

Kellyn Potvin'Gorman from Microsoft is explaining the partnership of Oracle and Microsoft Azure, while Andre Ontalba and Rodrigo Mufalani are showing how to clone an Oracle Autonomous Database, and Mia Urman and Elizabeth Pearl are sharing their experience on Oracle OpenWorld London.

In this issue, you can also read the third part of a very interesting APEX deep dive by Carsten Czarski, as well as the second part of Jim Czuprynski's Machine Learning super power and of course the second part of the AskTOM Journey of Connor McDonald.

For those interested in the Oracle ACE program or in general interested in contributing for the community, there is another interesting article by Phil Wilkins.

**Heli Helskyaho**

I hope you enjoy reading this ORAWORLD issue as much as I did! Please remember to submit your content for the upcoming issue online on our website: www.ORAWORLD.org!

Yours,
Heli Helskyaho

# 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**.

# Unintentional framework

*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***

*Connor McDonald*

# The AskTOM Journey (Part 2): Today and the Future

***In the last issue of ORAWORLD I wrote about the history of AskTOM and how it is intertwined with my own personal history of using Oracle technology. Now we come to the present.***

Currently on AskTOM we have a repository of over 25,000 answered questions and over 130,000 follow-ups and reviews to those questions. As database advocates for Oracle, Chris and I have a number of tasks on our plates at all times, including preparing educational content, preparing conference content, blog posts, tutorials and meeting customers face-to-face, but any spare time that we have we dedicate to answering questions on AskTOM. Similarly to Tom, we endeavour to answer the questions with our own experience first, because I think the best answers are always borne from personal experience. However, we also have a wonderful support network of being able to reach out to the expert product managers inside Oracle Corporation in order to get answers on particular topics that may not be our field of expertise. For example, we will often consult Sergiusz Wolicki for national language support, or Roger Ford for Text.

The ability for the community to get access to experts within the Oracle Corporation whom they might not have get the chance to, is what led to the AskTOM Office Hours program which was an initiative of my manager Steven Feuerstein. Each time Chris and myself reached inside the organisation to get assistance with a particular question, we came to the realisation that there are literally hundreds of experts inside Oracle; people with incredibly detailed knowledge on their particular product areas; people that the customer community perhaps had not only never met, but perhaps never even heard of. This is because some of these experts are “inward facing”, that is, it is not part of their role to be outward facing to the community. This seemed such a lost opportunity because of the incredible wealth of knowledge that they have. The AskTOM Office Hours program is our attempt at bringing that incredible level of knowledge to the community, so that Oracle customers have a much greater likelihood of getting the most out of their Oracle investment and being successful with their applications. The Office Hours program guarantees to make experts from inside Oracle Corporation available to the community for one hour per month. Because of the number of experts that we have, this often equates to three to four sessions every single week, and since its inception last year, we have run nearly 300 sessions across a broad range of topics. This enhancement to AskTOM is one of our team’s proudest achievements because it is a recognition of the focus of the importance of the success of the user community.

I am often asked what the main challenges are when it comes to answering AskTOM questions. Probably the two largest challenges Chris and I face are: breadth and precision. The challenge of breadth is that the Oracle technology platform has grown incredibly over the last 40 years. It is no longer just a database company. It is a solutions company – there is hardware; there is software; there is cloud; and there is the integration of all these things. That myriad of areas naturally leads to questions from people on those areas coming into AskTOM. Sometimes because of the breadth of knowledge required, we have to simply pass those questions on to the various community forums that we provide, but Chris and I always try to bring our years of experience first and foremost. A far greater challenge however is the issue of question precision. In these modern times of fast turnarounds and instantaneous gratification,

typically questions that come in to AskTOM have been very quickly authored and submitted, with not a lot of detail or thought put into them. Chris tweeted a while back that the AskTOM database should not be about answering questions, but more about tutorials on how to intelligently ask a question. I think this an unfortunate reflection of the modern pressures on developers nowadays. If they pause for thought or take the time to really think through the mechanisms via which they should be building their applications, then they are criticised for that, or categorised as being unproductive. “Always be cutting code” seems to be a common mantra. The questions we see on AskTOM are a reflection of that. A developer encounters a problem, or hits a roadblock, and the “resolution” is to jot down some details as quickly as possible and throw that into many question forums as possible. We often see the same question text appearing as an AskTOM submission at the same time it appears on StackOverflow and community.oracle.com. The hope is to get a response as fast as possible to remove the roadblock as quickly as possible. Chris and I always want to help, but it is also in the long-term interest of developers and their customers to be able to craft a question well and build sensible test cases around it.

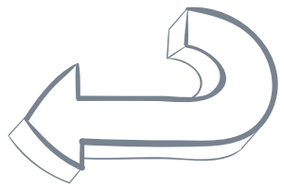

This is one of the reasons we integrated LiveSQL with AskTOM so that people could provide test cases via LiveSQL and link to them from AskTOM. Not only does that make our job much easier, I believe it teaches very good skills to the person asking the question namely

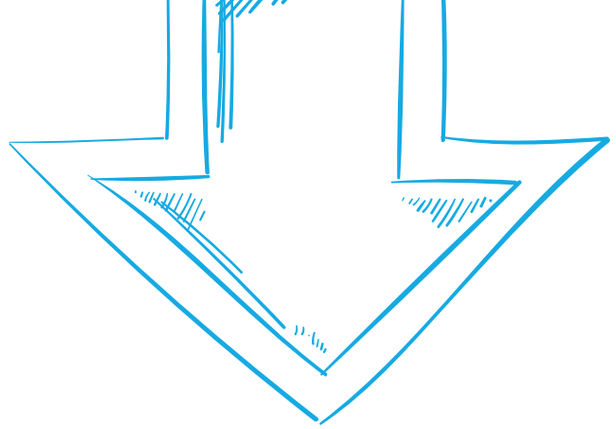

a. the ability to put together a top to bottom testcase that demonstrates an issue or a roadblock, and
b. the ability to remove all unnecessary content from a question, but still includes all the necessary content.

The two AskTOM questions that Chris and I dread the most are the question that has not enough information, or the question that is absolutely flooded with information that is not pertinent to the problem at hand. A developer who can craft a precise question will often be able to solve a problem without assistance and is also more likely to be able to deliver quality code because they understand the benefits of precision.

AskTOM was invented to solve in part the issue of not being able to get quality information to ensure successful database applications. But what of the future? With the internet era of blogs, community forums, webinars etc. is there a future for AskTOM? I don’t see the website going away but we also make no claim to be the only question and answer website that you should use for Oracle questions. AskTOM is not trying to be the next StackOverflow or the next community.oracle.com. AskTOM continues to exist as another means of

continuing to build upon the great Oracle community. Every time Chris and I can reach out to the community via AskTOM, or one of our product managers does so via Office Hours program, I think we are getting more successful customers, which increase our focus on customers to ensure that they get the best return on investment in Oracle technology. The Office Hours program is a good indication of the direction I see AskTOM heading, in that we will do our best to get more experts, more industry experience, and more knowledge from inside the Oracle organisation out to customers to ensure their success with Oracle.

Thanks for using AskTOM!


## About Connor McDonald

Connor is a Database Advocate for Oracle Corporation. Ever being frustrated as a child with the limits imposed by the single kilobyte of RAM in his Sinclair ZX80 computer, he has loved the challenges that come with storing data, which ultimately led to a career in database technology. When the infamous Y2K data issue did not end the world as people thought, he started presenting on his database passions, found that he loved doing it, and has spoken at over 120 conferences around the world since then. Due to his partners predilection for rescuing stray cats, he doesn't just speak the phrase "as hard as herding cats", he lives it every day.

*Phil Wilkins*

# Tips for Oracle Ace Associates and other Advocacy Groups (Part 2: Writing and Presenting)

***In the first part of this article I focussed on blogging. But as blogging becomes easier, then the writing goals can become more advanced by writing an article for a journal, particularly the user-group-based journals.***

## Writing beyond your blog

Many user groups have journals, between the majority and all the content is community sourced, so the editorial teams are always ready for new submissions. Depending on the journal you may need to be a member of the user group. At worst, the review team (all volunteers usually) will feedback to you how they think your submissions needs to be refined – feedback that will also help with your blogging. But sharing articles through major technical websites as well is another good progression, for example DZone, InfoQ and so on.

## Taking the Ultimate Step

The ultimate step when it comes to writing is a book. Many Aces have gone down this route at some point. This means that when it comes to publishers a number of Aces will be able to make suitable introductions and give you some insights into how a particular publisher likes to work. But to be upfront writing a book isn't a trivial undertaking. I've yet to hear of an Ace completing a technical book in less than 9 months, during which time you do need to keep a level of activity up.

The writing of a book is an exercise in commitment to sharing knowledge. It is very rare to make enough money to replace the day job. But, completing a book is deeply satisfying – seeing your work bound and in print, in a bookshop, in Amazon gives a great sense of reward. In some respects, this brings you full circle, people will perceive you as an expert in your field.

## It isn't all about writing

Whilst written media be it a blog or a book is a core way of contributing, it isn't the only one. Many Aces also get involved in presenting. Like writing, such an activity can be daunting to say the least. My personal experience may not work for everyone, but you lose nothing in considering it. I didn't set out to be an active presenter, but this is where I have found myself, whilst I still experience nerves, I do enjoy the interaction and the feedback which reaffirms that my thinking has validity.

My first steps in this direction was simply when working within a team to get up and record on a whiteboard what is being discussed. To say simply standing in front of peers in a room used to make me hot and uncomfortable would have been an understatement. Every presenter that is honest will tell you that they experience some nerves even if they have been doing it for years. The well-liked and respected Oracle Ace Director Tim Hall has spoken about his experiences presenting, and has been very open in talking about nerves and experiencing imposter syndrome. He has posted a number of useful types on his **blog**.

Overtime, you realize actually no one is going to pull you apart. With confidence in your domain, this becomes an opportunity to pitch your ideas, a bit tougher as people will ask questions and challenge your ideas. The next step is potentially more formal meetings where you've had to prepare material as well. You are now essentially a presenter, the only difference between this and being at conferences presenting, is merely numbers.

Several people have said to me (and I agree), the easiest way to deal with the nerves, is to simply stop worrying about what others think. That doesn't mean stop trying to improve or not listen to feedback; but on the day, in the moment of presenting the only thing to focus on is the subject and trying to find some satisfaction in the process – these are the only things you have a chance of influencing.

## Making presenting less intimidating

The only difference between presenting within your office and a conference, is you know the audience at work. But we can do that to an extent in the public context as well, simply attend the meetup, special interest group, conference which you think you'd like to pitch a presentation at. People will get to recognize your face, and you theirs. Smaller user group events and meetups are a great place to start, they're less formal, people haven't 'invested' a lot their time in attending so they are a lot more forgiving and understanding and appreciate you've put in more effort than they have. Once you start feeling comfortable, perhaps aim for bigger events. These will have their sessions recorded to YouTube, so watch videos of previous presenters to get a feel for what is likely to be expected.

## What to present

For me, many of my presentations could be easily characterized as visual/audible blog. Considering what is my message and how to express my idea/solution are exactly the same. I have a couple of additional recommendations:

## 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.

Phil has presented at events around the world from Sweden to California. Phil is a co-organizer to the London Oracle Developer Meetup & Oracle Ace Director.

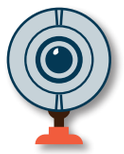

- Be brutally honest with people and set expectations – first time presenting, new presentation deck, say so. People will give you a lot more support and appreciation if you're candid,
- Be comfortable about the subject matter, and the things of the peripheral of your subject – by answering a question does a lot to boost perceptions of expertise,
- Get a question, you can't answer, don't try blagging it, be up front – say you don't know or not 100% but always commit to getting an answer back to the group.

The beauty of presenting, particularly when there is a dialogue: the questions can be the source of inspiration for blogs or new presentations. If you're lucky, the whole process can become a bit of a virtuous cycle.


## Preparation

Preparation for presenting requires a lot more effort whether that is putting slide material or demos together. Demos for me, are the most nerve wracking because there is always something that might go wrong that is beyond your control – poor wi-fi signal or dropouts. In the heat of the moment you miss a key step in the process. Whereas while talking to slides it is easier to reduce the risks – have a copy of the presentation on USB stick, make talking points in your notes etc. etc.

## Conclusion

Writing and presenting are not the only things an aspiring Ace can do to contribute, but these are the most common, and the more you do it, the easier it is. So, if you want to give it a try, then ignore the self-doubts, focus on what would help you, what you can do. The Ace community is an incredible group of people who are willing to help and support you. Asking for an Ace's help is not a weakness, it is gift of giving an Ace the chance to help.

# “We Need to Take Direct Action in our own Lives to Heal our Planet”

*Interview by Christian Luda*


***fabe** (for all a beautiful earth) is an APEX-based PWA launched in 2019. The app helps its users to slow the rate of planetary warming and to stop human-caused extinctions. We talked to the project’s two founders, Steven Feuerstein and Vincent Morneau.*

***Steven and Vincent, you founded fabe in 2018. How did the two of you find together? When and where was the idea born? Were there any previous apps, projects or activists that inspired you?***

**Steven:** Back in 2011, I realized that while I "believed" in evolution, I didn't really know much about it. So, I started reading – and it completely changed my life. For now, acceptance of the reality of evolution by natural selection led directly to a moral challenge: evolution shows that the multitude of species are all different, but no one species is "better" or "smarter" than any other. Different, not better. Which then took a position of abject horror at how humans were killing so many other creatures, many if not all sentient. It hit me hard and from there I asked myself: what is causing all this killing? And my answer was: human consumption – in particular, our addiction to comfort, convenience and entertainment. Every time we buy/consume, we kill our planet a little bit more. However, I do realize that many humans struggle simply to survive – I am not including them in this critique.

*"Extinctions should not happen just so humans can enjoy cat videos"*


All that killing of life to build factories to make stuff, and then deliver that stuff, and stream videos over the internet – all of that is also the main driver of climate change. So, for me, a focus on climate change comes from a deep moral conviction that killing others is wrong, unless it is done for survival. And that belief has informed the fabe project and makes it quite different from other similar apps.

All right, so there were my beliefs. But what to do about it? Well, what I did next was rejoin Oracle Corporation after being away for 22 years. I then decided to think about action to take, but not do anything for a little while. I was very busy building a team of developer advocates. So, I thought and watched and listened and read – and continually asked myself: how can we make change happen fast and how can we get anything done at a time when many of our systems and institutions were so badly degraded? I also experimented with my own life: how could I eat and act in the world to do the least harm and start the healing of our planet? I completely changed my diet to minimize use of plastic and avoid eating animals. I also became obsessed with removing invasive species from natural areas near my home in Chicago, first among them: buckthorn. I spent hour after hour with sharp objects, cutting down and pulling out buckthorn trees, so that native trees would survive and thrive. I got such an intense feeling of purpose and satisfaction from the changes I made in my own life. And I was having an impact.

From these experiences, I saw what could be done if millions of other people followed their own version of my path, which prioritizes taking direct and positive action in one's own life. Then in August 2018, I couldn't take it anymore. I felt the need to act. I remembered that Vincent Morneau had expressed similar views on climate change on Twitter. So, I got in touch, shared my ideas and he said: "Let's get started!" And we did.

**Vincent:** At some point during my career as a developer, I thought my passion in life was tech. I was an avid gamer as a teenager, and gadgets were providing the most pleasure to me.

As I started to make some money, I became a huge consumer of technological goods. I would buy things I'd use once or twice, then return them or throw them away after I got bored. I simply didn't care.

Then some family circumstances happened in my life and I ended up owning a piece of land and a minuscule cabin, so far remote from the city, without access to the power grid or cellular network. Driving up to this cabin became a very important part of my life. Connecting to nature in this way made me realize how much the gadgets I used to love so much were in fact a distraction from my true happiness. From that moment on, I became very sensitive to topics around climate change. Through education and scientific literature, I understood how fragile our ecosystems are, and I started to fear for the future of the natural world and all its inhabitants. Many years went by and I continued my career in tech consultancy. I was in the middle of a cycling trip when Steven emailed me the original fabe idea. I took it as a call to action. Now I know: my passion is nature. And I hope I can use my tech skills to preserve what's left of it.

***You have assembled a team of more than 30 people. How was the process of putting the team together? How did you organize the teamwork?***

Steven: Vincent and I announced fabe in September 2018 to our database/APEX community and asked for volunteers. We were overwhelmed by the response! Dozens of people offered to help. More, we soon realized, than we could really utilize properly as we certainly don't want to waste the time of volunteers. So, we identified a "short list" of talent, reached out, confirmed participation, and got going. As a Canadian non-profit, we were able to take advantage of Google's free offering of its G Suite of tools. We make extensive use of their Drive, Chat, Meet and Calendar apps to coordinate and communicate. We use GitLab for ticket management. And we could use some more help on the project management side. So, if anyone reading this is interested in helping, let us know!

***The project sounds like a full-time job. How did you find the time to work on it besides your regular jobs?***

Vincent: We often say fabe is an all-volunteer effort, but that's a little bit of a lie. For the first year, that statement was true, but in the summer of 2019, we had an idea that could help fabe grow faster. I joined Insum in 2010 as an intern. I was the fifth employee at the time. Fast forward a decade, Insum employs close to a hundred people, most of them APEX consultants and developers. Being in the business of consultancy, you can imagine that it's hard to keep 75 developers busy all the time. It's normal to have a few developers on the bench, because projects come and go, and we need to have them ready to jump in when the business requires them. Steven and I sat down with Insum's CEO Michel St-Amour and CTO Francis Mignault, and we crafted a strategy that would allow fabe to leverage developers on the bench. For one, that would accelerate fabe's development, and it would serve as a great APEX training opportunity, while keeping these developers busy in a fun and engaging project. It's a win-win-win! We are extremely grateful at fabe to be supported by such a great company.

***fabe stands for "for all a beautiful earth". What is your vision of a beautiful earth for all?***

Steven: The best way to answer this question is to repeat the four objectives of fabe when we first announced it and called for volunteers. Our first priority is to slow and stop as many human-caused extinctions as possible, especially of humans and other known sentients. Extinction means gone forever. Extinctions happen in the normal course of events but surely, they should not happen just so humans can enjoy cat videos and the like. Our second priority is to ensure that future generations are able to lead sustainable, healthy lives. Climate

change and environmental degradation are having a terrible effect on people around the world, with worse yet to come. We must take action to ensure our children and grandchildren are able to survive and thrive. Our third priority is to end the mass brutalization and killing of domesticated sentients, such as pigs and cows. Yes, pigs and cows are sentient creatures, solving problems and caring for their families, just like us. It is flat out wrong to treat fellow sentients the way we do. The least we can do is treat well and with respect the animals we are going to eat. Our fourth priority is quite different. It is to saturate our lives with purpose and meaning. We want our lives to be about more than simply getting through the day. We feel the greatest joy and satisfaction from helping others.


***A key contribution of fabe is the introduction of three new R's – reduce, rescue and reconnect, the latter two replacing recycle and reuse. The removal of recycle and reuse comes with a plea for minimum consumption. What is the hardest thing for yourselves to reduce?***

Steven: For me, dairy and eggs. I made the shift to vegetarian several years ago. I will not eat the flesh of sentients if I can avoid it. But it's been harder for me to give up dairy (milk and cheese) and eggs. I only buy milk and cheese from grass-fed, pasture-raised cows, so I am trying to minimize my collaboration in (and benefiting from) animal abuse. But still: eggs I buy from a nearby friend. I know her chickens are happy and healthy and so as long as I can do that, I see no reason to give up eggs. But I am working to integrate other protein sources into my diet and may at some point be able to give up dairy entirely. And – breaking news! – I have decided to take a break from making my own yogurt (the main reason I buy milk). I will instead try oat milk with my granola, nuts, raisins and chopped fruit. Maybe this won't be so hard after all!

Vincent: My plan to reduce consumption has been drastic, but honestly also fairly easy. I actually get a lot of pleasure out of it. I cook more than ever before, out of whole foods. I seek less business travel than before, which diminishes unnecessary stress. A lot of good things happen when you reduce. However, what has been more difficult is the social effect around me. I can get a little grumpy when friends or family order take-out food with more plastic and styrofoam than actual food. Even close relatives still don't believe me when I tell them I prefer a no-gift birthday. I need to get better at managing expectations at certain events. There is something I reduced but did not eliminate from my life: my desire to travel for pilgrimage. There is so much beauty to see in this world. There are more sustainable ways to travel than others, but a plane is a plane unfortunately.

Steven: Also, we need to provide you with an update: we are now thinking and organizing our app around four Rs, adding the new R "reject", meaning reject the status quo: it is now clear that the systems humans created to manage modern civilization (economic, political, national, religious etc.) are not adequate to the tax of making radical change in face of an existential threat. We are, collectively, failing at this very moment. This means we have to either change the people running those systems or get

the systems to change faster and with more urgency. Probably we need to do both. We each need to remind ourselves every day that the status quo leads straight to catastrophe and death to hundreds of billions of living creatures, including humans. Then we need to make sure we apply direct pressure for changing the status quo.

***Have you encountered certain personal habits in the course of the project that you were not aware of before, how harmful they are?***

Steven: Throwing away – and that includes, sadly, "recycling" plastic. I go to great lengths to minimize the amount of plastic in my life. One way I do that is to minimize the amount of, well, anything new in my life. I try really hard to never buy anything unless I actually need it. Also, I try not to buy things with single-use plastic if it is at all possible. Still I end up with plastic. And here's the thing: plastic kills. Plastic kills as it is being made and it kills when we throw it away or "recycle" it. Why do I put the word "recycle" in quotes? Because – and I hate to say this – worldwide, recycling has largely broken down. There's a good chance that any plastic you put in the recycling bin will end up in a landfill. So, I decided that the safest thing I could do with my plastic is keep it. Sequester it. Don't throw it away. That way, it cannot hurt other living creatures. We'll see how it goes. I have room to store it without causing great distress to my wife or neighbors. And it's a big disincentive to consume any new plastic when you know you will have to put it somewhere!

Vincent: Most of my life, I considered "energy" as a monetary concern, not an ecological one. Electricity is pretty cheap where I live, so I could go to work all day and leave the AC turned on because, hell, it would cost only a few cents and I'd come back to a cool apartment. On the other hand, I would not keep my car engine running idle, only because that would be too expensive on gas. I was hardly thinking about the footprint I was leaving on this planet. Looking back, I was a very selfish human.

***You have team members from all over the world. Did you come across different priorities and new perspectives? How did you benefit from having such an international team?***

Steven: The bigger the team, the more widely it is distributed, the bigger the variation in priorities and perspectives, for sure! One of the biggest challenges for me, as the team grows, as the project evolves, has been to realize that just because I have one idea about how things should be done, it doesn't mean that it is a good idea, the best idea, or an idea shared by others on the team. We have benefited so much from the different areas of passion and energy from members of the team. Some of us come to fabe with a strong, almost exclusive, emphasis on reducing emissions (a focus on planetary warming). Others, like Steven, get very worked up over extinctions, animal abuses, and how fabe can help reduce them. The end result is an app and messaging that is inclusive and welcoming to all interests. Another big benefit from an international team that we are just beginning to explore is producing social media content in multiple languages. We have some very interesting and unique ways of presenting the challenge of responding to the climate emergency. We will explore how to present those in different languages, to different cultures.

***You have chosen a PWA with APEX for fabe. What was the reason for this decision?***

Vincent: This is an interesting question because we didn't choose PWA. PWA chose us. When Steven and I first announced we were looking for volunteer developers on social media, it reached the broad Oracle development community, but more specifically the vibrant APEX community. We ended up putting together a team of 12+ Oracle developers, all of them being proficient in Oracle APEX, including myself. Then we were left with a challenge: how do we build an APEX application

that can look and feel like a modern mobile app? Since APEX essentially produces web applications, we had to use web technologies. Today, PWA is the best web technology to create a web application and put them into app stores. We decided to use APEX for many reasons. First, remember that fabe is a volunteer-based initiative, so that means development time is highly variable for our developers and one of the absolute strengths of APEX is its efficiency to navigate through the application builder allowing us to add new functionality in relatively little time. We also believe the maturity of the framework and the commitment by Oracle to support APEX will help fabe grow sustainably in the future. We are already looking forward to having many known elements of the APEX Statement of Direction in 2020, and we can't wait to see what's in store next. Another great productivity helper is that our application is hosted on the Oracle Cloud using the Oracle Database Cloud Service. Our infrastructure efforts are minimal, and we can focus on what matters most: creating a secure, fast and modern mobile application that scales worldwide.

***What were the biggest technical challenges during the implementation of the app?***

Vincent: As much as we love APEX, some important mobile features users expect were not yet baked into the framework, so we had to get our hands dirty a few times. For example, when I first downloaded Twitter on my phone, I was asked to create my account, or login to my existing account. After that initial process, every time I open Twitter, I see my feed and my notifications. I don't have to redo the login every single time I open Twitter. This is common practice for most mobile apps. APEX apps don't do this out of the box, so we had to create an automatic login authentication scheme that stores an encrypted fingerprint on the mobile device which is used to identify your user and get you in the app seamlessly. Also, because fabe is a content heavy platform eventually dealing with thousands of users posting images and videos, we had to architect a system where regular file uploads are parsed through automatic compression. We obviously don't want to store 15 MB of PNG pictures on our servers, so we made sure any image uploaded with APEX runs first through our compression service. Facebook does that, so I think fabe can too. Lastly, we quickly found out one navigation pattern where APEX falls short: It's common practice to use modal dialogs in APEX, and fabe does too. One problem with modal dialogs is that the user gets the impression that you are one page "ahead" of the parent, and a normal reflex a user might have in order to come back to the parent page would be press the back button. Now of course, the browser does not know a modal dialog is opened, so pressing the back button will take the user back before the parent page was opened. As this navigation pattern is very awkward, we created a new navigation mechanism in fabe where modal dialog actually counts as a browser state change. This results in a natural use of the back button.

"We are, collectively, failing at this very moment"


***Who are you trying to reach with your app? Are you aiming primarily at people that already have a certain level of awareness and want to do something? What about people that deny that the environmental crisis is man-made?***

Steven: We are not focused on people who deny human-caused climate change. There are literally hundreds of millions of people who feel the crisis urgently and are taking action or

want to get active. We want to give them a way to internalize climate action into their busy lives and maximize their response to the climate emergency. We are particularly interested in getting fabe in the hands of school strikers, Extinction Rebellion members, and Sunrise Movement activists. They are doing incredibly important work and if they complement their activism with changes in their consumption patterns, we believe things can start changing must faster.

***You released the app at the end of August 2019. How has the feedback been thus far?***

Steven: We got some very mixed feedback, from "OMG we've totally changed how we buy things now." to "How do I get started? What am I supposed to do?" It was, in other words, very much a 1.0 release, in that almost immediately on getting into the hands of users, we realized that: First, it was a significant improvement on other climate action tracking apps, with its focus on building and carrying out action plans and not just a variety of individual actions. Second, it needed some major work to make it as sticky and impactful as is needed. We are now working on a substantial reworking of fabe that we believe will make it compelling and very popular.

***The app proposes various actions. What are some of your user's favorite actions that you did not expect to be so popular? Which ones do you think should be more popular?***

Steven: Actions more popular than expected? Two stand out: "Don't (always) flush for #1." and "Bring your own reusable dishware and utensils." In other words: don't use plastic, throwaway forks and knives and spoons when you are traveling. We are hoping to get more people to add rescue actions to their daily plans, such as volunteering at a nature preserve and removing invasives. Consumption reduction is, of course, key. But we also need to take direct action – each and every one of us to the extent possible – in our own lives to heal our planet.

*"Looking back, I was a very selfish human"*

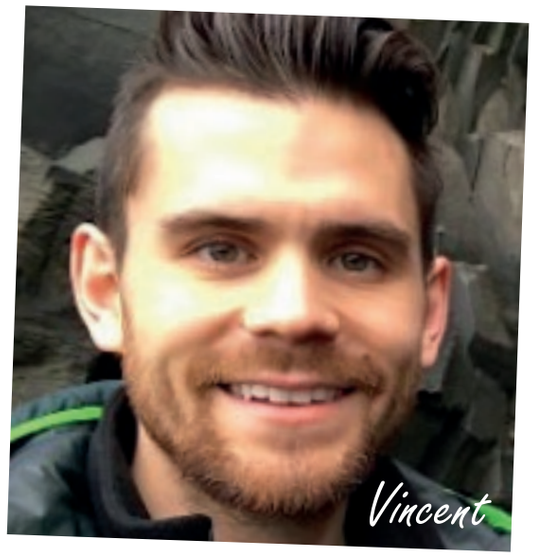


***Can you give us an outlook of the planned gamification features?***

Vincent: With our first version of fabe, users collect life points for each action taken. For instance, you use your bike to work and you gain three life points. You skip meat for a day and you gain five life points. As you start taking more and more actions, you accumulate those points as part of your journey down the fabe path. It was a great first step at gamifying our action library, but if we are being honest with ourselves, it was missing some depth. In the next version of fabe, we are looking at implementing levels, achievements, badges and many forms of reward elements that will make the fabe experience more fun and more addictive.

***Has there been any criticism from other activists?***

Steven: We have not much criticism of our app specifically, but there is a disturbing trend among a number of climate action leaders to be dismissive about personal change. They think the idea of changing individual consumption is a good thing to do, but to really make a difference you have to organize for system change – change in government, in laws, in corporate policies. We totally agree about the need for systemic change. But we

also believe that we must also change our own lives in order to (a) reduce the burning of fossil fuel, (b) convince politicians and CEOs of our seriousness, (c) start healing our planet now, and (d) be energized to take on collective action with others. In other words, the fabe team believes that the only response to the climate emergency that has any chance of success is: Everyone taking action every day in every possible way – except violence or coercion. Reduce your consumption, volunteer at a nature preserve, protest bank loans for fossil fuel projects. Buy only what you need, remove invasive species in your community, join a local climate action group. Avoid single-use plastic, plant trees and tend them to make sure they survive, ask to meet your political representative to talk about saving life on our planet. Everyone, every day in every way. Our next version of fabe will help all of us do this.

***If you look at Oracle and other big tech companies: Do you think they do enough to fight climate change and environmental destruction? What are the biggest challenges for the tech industry?***

Steven: No one is doing enough to slow down planetary heating and human-caused extinctions. But many people and some organizations are doing a lot. For the most part, tech companies are a big net negative right now. They consume enormous amounts of energy and concentrate enormous amounts of wealth. Microsoft recently made a big and very positive announcement about their climate plans. We hope they meet their aggressive goals. Oracle is taking some steps, particularly around reducing the environmental impact of its cloud, but is certainly not a leader at this time. The biggest challenge for the tech industry is that it is a major enabler of the Internet, automation, and global scale-up of human activity: all things that are used overwhelmingly to increase consumption and further devastate life on our planet. How does that stop without compromising the bottom line for many of these companies? That's a bit unclear right now. On the positive side, these companies have so much money, they could right now pour money into existing, proven, on-the-ground initiatives to save species at risk of extinction. Imagine if Oracle and Larry Ellison decided that they would save orangutans – because, of course, alliteration is key: Oracle for Orangutans!

***Steven and Vincent, thanks for taking your time.***

## Further links:

Read more about three of fabe's initial 3 Rs (reduce, reuse, reconnect) on their website: **https://forallabeautifulearth.org**
They will be adding more about the 4th R, reject, soon.

Find out more about the breakdown of "recycling":
**https://www.theguardian.com/environment/2019/aug/17/plastic-recycling-myth-what-really-happens-your-rubbish**

Read more about Steven's plastic challenge:
**https://forallabeautifulearth.org/taking-on-my-next-plastic-challenge/**

Check on Oracle's plans when it comes to sustainability of operations here: **https://www.oracle.com/corporate/citizenship/sustainability/operations.html**

*Christian Luda*

# The Most Famous Home Worker Returns

In these times of Corona most of us are working from home. It is a great technical achievement that allows us to work remotely and thus helping to slow down the spreading of the virus. You can only imagine how we could have dealt with such a crisis before the internet. However, especially for people with young kids, this situation means a big challenge.

Who knows better than Professor Kelly and his wife Jung-a Kim who, well, went viral in 2017 when their children Marion and James **crashed a live interview**, Kelly was giving to the British broadcaster BBC.

Thanks to Corona and working from home becoming such a big topic, the three-year-old clip recently became a YouTube hit again. So, Professor Kelly and his family returned to BBC giving some insights about the lockdown in their hometown of Busan, South Korea.

Before the interview, Kelly had spoken out for workers with kids. On Twitter he posted a picture of him and his son, stating: “This is what happens when I sit down at my desk now to try to work. It is basically impossible for me to work now. Be kind to your employees with kids. After two weeks penned up in the house, those kids are gonna be climbing the walls.”

# Number of the Month


# 95% of IT Leaders say Open Source is Strategically Important

More and more economy players are relying on open source, while proprietary software is losing importance. This is the result of the Red Hat study "**The State of Enterprise Open Source 2020**", for which market researchers surveyed 950 IT executives from the United States, Great Britain, Latin America and Southeast Asia. The respondents were unaware that Red Hat was behind the research.

According to the survey, 95% of enterprise experts consider freely licensed, easily redistributable and customizable programs to be an important component of an infrastructure software strategy.

The share of open source software in use is currently 36% but expected to rise to 44% in the next two years. Free software is no longer seen primarily as an option to reduce costs. When asked about the driving factors for use, 33% of the participants named higher software quality. Lower total cost of ownership (TCO) was cited by 30%, better security by 29%.

63% of the respondents have a hybrid cloud infrastructure. Of the remaining 37% who don't, 54% are also planning to introduce one within the next two years. 86% think that enterprise open source software is used by the most innovative companies while 83% see it instrumental in their organization's ability to take advantage of the cloud architectures.

Red Hat first took the survey a year ago. At that time, 55% of the software used by the companies surveyed was proprietary. According to new survey this proportion has dropped to 42%. In two years, the respondents expect a share of only 32%.

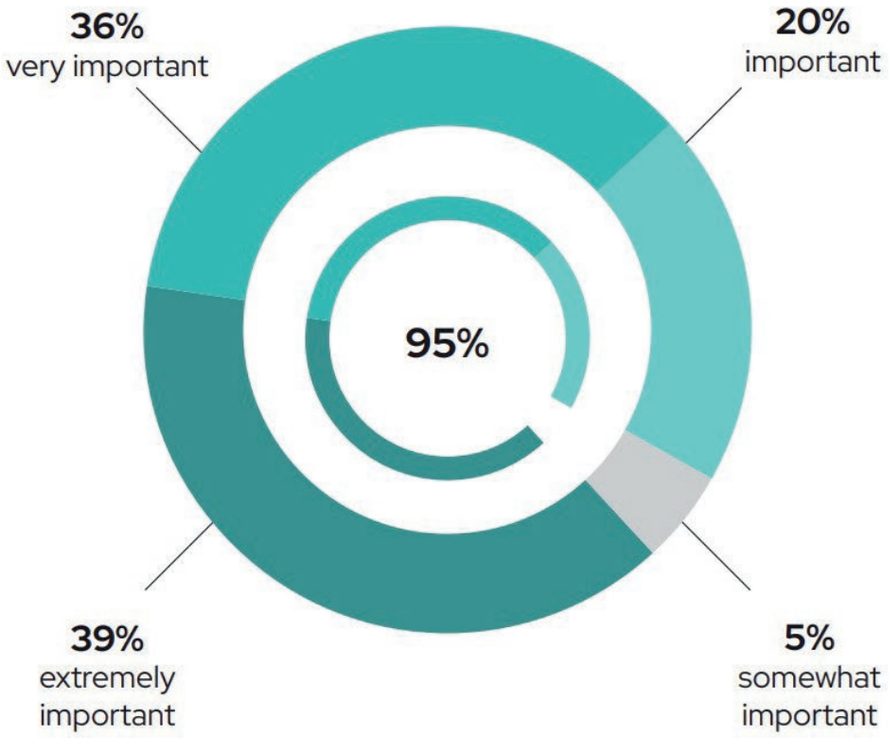


*The importance of enterprise open source according to the Red Hat survey respondents.*

# Oracle Application Express (Part 3): APEX and the REST of the World – Integrating REST Services into APEX

*Carsten Czarski*

## Introduction

Integrating other IT systems with REST services is daily developer work nowadays. Especially Cloud services and applications expose their functionality as REST interfaces. The actual development task is rather straightforward: An HTTP client package and a JSON or XML parser is all it needs.

For APEX developers, the APEX_WEB_SERVICE PL/SQL package acts as the HTTP client, and the JSON or XML parser is provided by database SQL functions. However, in recent versions, APEX also provides a meta data driven, declarative approach to integrate REST services: **Web Source Modules**. This article will provide an overview on using Web Source Modules versus pure PL/SQL coding with APEX_WEB_SERVICE.

## PL/SQL API: APEX_WEB_SERVICE

The **APEX_WEB_SERVICE** PL/SQL package was introduced with APEX 4.0 and allows to invoke REST APIs using the MAKE_REST_REQUEST function. The example in Listing 1 invokes the Github REST API (api.github.com) in order to get all repositories belonging to a specific Github user (e.g. “oracle”).

```
begin
  apex_web_service.g_request_headers.delete;
  apex_web_service.g_request_headers( 1 ).name  := 'User-Agent';
  apex_web_service.g_request_headers( 1 ).value := 'APEX';
end;
/

select apex_web_service.make_rest_request(
    p_url          => 'https://api.github.com/users/oracle/repos',
    p_http_method  => 'GET' )
from dual
/

[{"id":33195015,"node_id":"MDEwOlJlcG9zaXRvcnkzMzE5NTAxNQ==","name":"accel
erators","full_name":"oracle/accelerators","private":false,"owner":{"login
":"oracle","id":4430336,"node_id":"MDEyOk9yZ2FuaXphdGlvbjQ0MzAzMzY=","avat
ar_url":"https://avatars1.githubusercontent.com/u/4430336?v=4","gravatar_i
:
```

*Listing 1: Access the Github REST API with APEX_WEB_SERVICE*

```
with github as (
  select apex_web_service.make_rest_request(
    p_url          => 'https://api.github.com/users/oracle/repos',
    p_http_method  => 'GET' ) as json
  from dual )
select name,
    language,
    forks_cnt,
    owner_id,
    owner_name
  from github, json_table(
    json,
    '$[*]'
    columns(
      name       varchar2(50)  path '$.name',
      language   varchar2(50)  path '$.language',
      forks_cnt  number        path '$.forks_count',
      owner_id   varchar2(15)  path '$.owner.id',
      owner_name varchar2(30) path '$.owner.login' ) )
/

NAME                           LANGUAGE   FORKS_CNT  OWNER_ID.   OWNER_NA
------------------------------ ---------  ---------  ----------  --------
accelerators                   C#         54         4430336     oracle
accs-caching-java-sdk          Java       6          4430336     oracle
adf-samples                    Java       18         4430336     oracle
analytical-sql-examples        PLSQL      66         4430336     oracle
apiplatform-mgmtscripts-sample Python     2          4430336     oracle
:
```

*Listing 2: Invoke the Github REST API and immediately parse JSON results*

APEX_WEB_SERVICE can be used within APEX applications, but also outside of an APEX session context.

As required by the Github API, the code in Listing 1 first sets the **User-Agent** request header. Then it invokes the REST endpoint with the MAKE_REST_REQUEST function. The REST service will respond with a JSON document containing the requested data. Parsing the JSON, and processing its data, would be the next step. The **JSON_TABLE** SQL function, introduced in version 12.1.0.2 of the Oracle database, is the most efficient way to do this (Listing 2).

The **P_HTTP_METHOD** parameter of MAKE_REST_REQUEST allows to use other HTTP methods, like **POST**, **PUT** or **DELETE**. MAKE_REST_REQUEST_B returns a BLOB and can be used for REST APIs returning binary data.

With APEX_WEB_SERVICE, developers can use all sorts of REST APIs, either within the company network or on the internet. All details of the HTTP request, like HTTP headers, cookies or the request body can be configured using parameters or global variables of the APEX_WEB_SERVICE package. The response comes as a BLOB or CLOB and can be processed however the developer likes.

However, when accessing REST APIs programmatically, developers find themselves in writing the same low-level code, over and over again. For instance, the JSON_TABLE query will always have the same structure and semantics, just the JSON attribute names and data types will change.

Imagine, we had all details of our REST API stored as meta data. Then, a generic engine could do all the low-level tasks, like executing the HTTP request and parsing JSON, automatically. Results would be provided as rows and columns, as it was data from a local table.

## APEX 18.1: Web Source Modules

APEX 18.1, released in April 2018, introduced a new approach to support external REST APIs. **Web Source Modules** store all details of an external REST API as meta data. Having all this information, APEX is able to execute the HTTP request and to parse the response automatically. An APEX component, which consumes the Web Source Module, just receives rows and columns. The REST service becomes as easy to consume as a table or a SQL query; no low-level coding is required any more. Web Source Modules are found in **Shared Components**, in the **Data Sources** section.

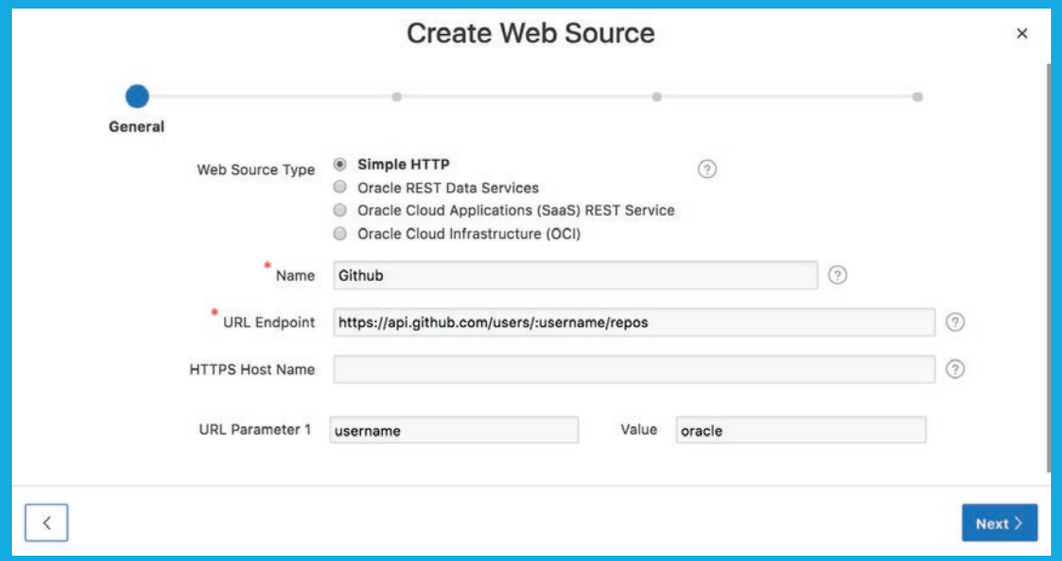

Fig. 1: Wizard to create a new Web Source Module

The wizard to create a new Web Source Module starts with providing general information like the HTTP method and the endpoint URL. URL endpoints can contain placeholders, which are replaced with values from session state, when the HTTP request is executed. For instance, "**oracle**" within "**https://api.github.com/users/oracle/repos**" denotes the Github user name, thus we will use a placeholder (:**username** or **{username}**) and assign a default value for it (Fig. 1).

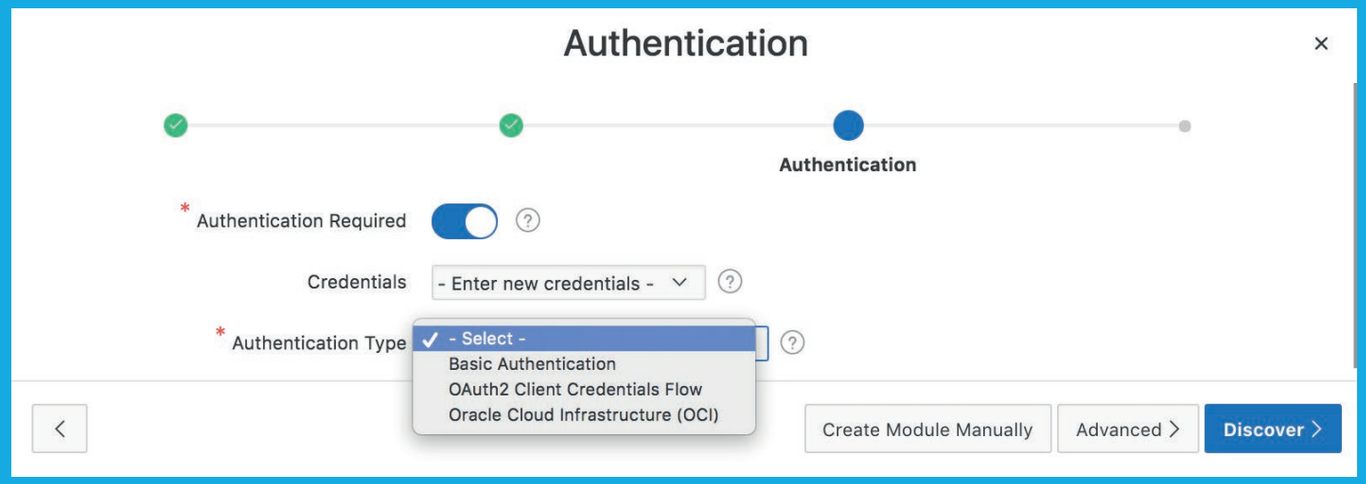


*Fig. 2: Provide authentication details for the Web Source Module*

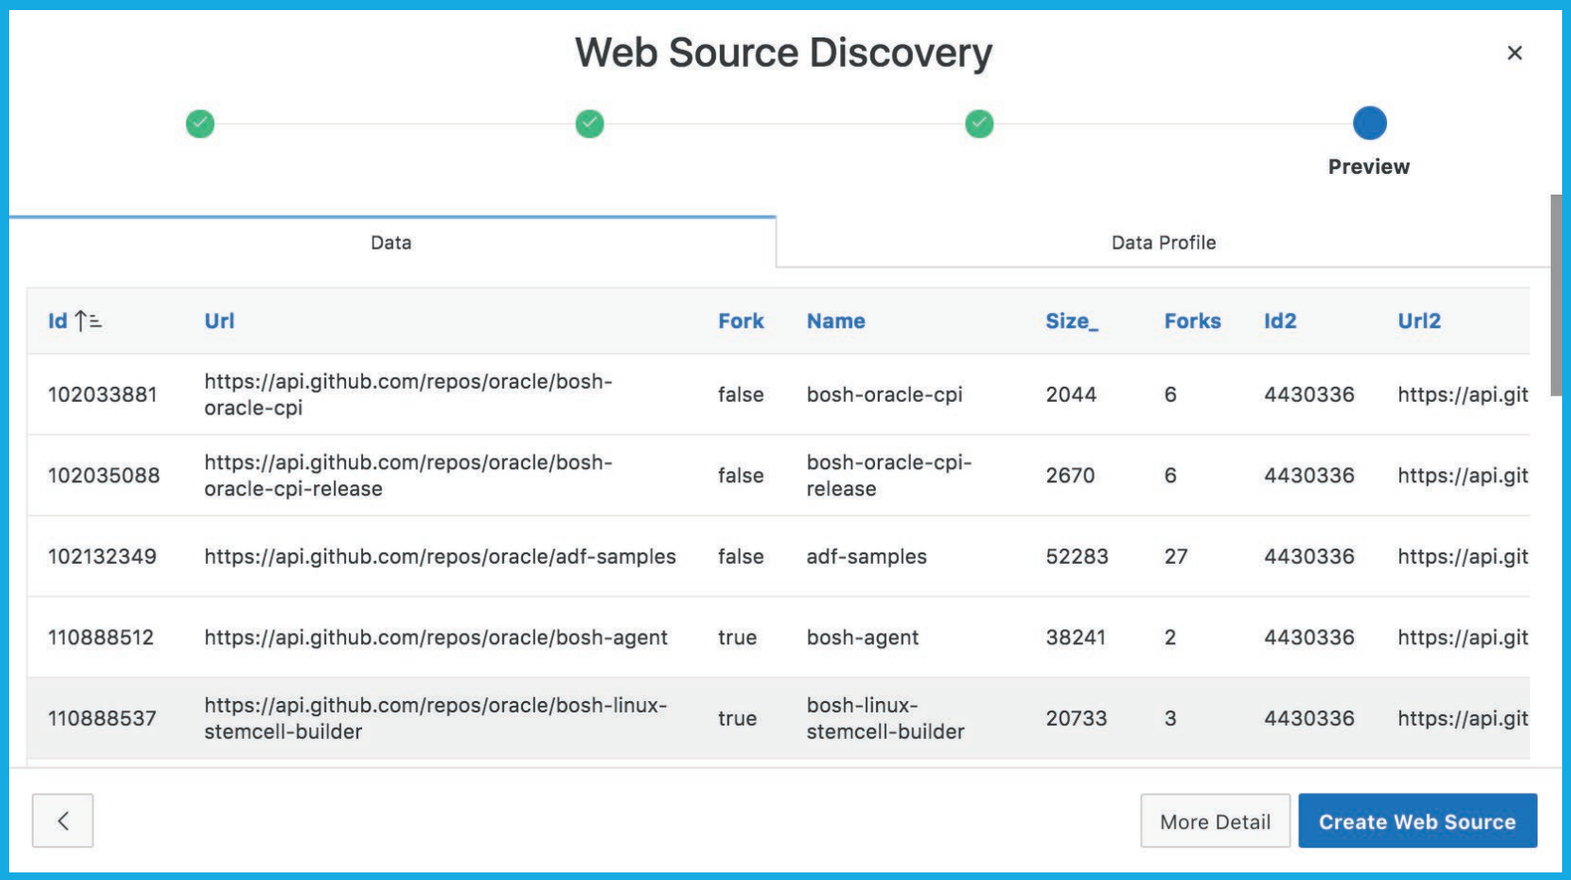


*Fig. 3: Web Source Module discovery result*

APEX also allows to provide authentication details like user names and passwords. Those will be stored encrypted, cannot be returned in clear text, and only the APEX engine is able to use them. APEX supports *Basic Authentication*, the *OAuth2 Client Credentials* flow and native Authentication for the *Oracle Cloud Infrastructure (OCI)*. The Github API, used in this example, is public and does not require authentication.

In most cases, clicking the **Discover** button is the final step of creating a new Web Source Module (the **Advanced** section allows to provide additional details like HTTP headers or JSON parsing hints). APEX will invoke the REST endpoint, receive and parse the response in order to determine a Data Profile. The *data profile* contains all required information to treat the REST results like a table, i.e. all details, which were manually passed to the JSON_TABLE function in Listing 2.

Clicking the **Create Web Source** button will store the new Web Source Module. All details can be reviewed or changed later on. As Fig. 4 shows, we now

have a declarative **Web Source Module Parameter** for the ":username" placeholder in the endpoint URL.

The **Data Profile** section (fig. 5) contains all information on how to parse the JSON or XML response. With this information, APEX is able to correctly parse the REST service response and to provide rows and columns to consuming components.

This part of the meta data can also be edited later on: APEX allows to add new columns or to change the definition of an existing column. However, once a Web Source Module is in use by an APEX component, existing columns can only be hidden, but not deleted any more.

Now it's time to actually use the new Web Source Module.

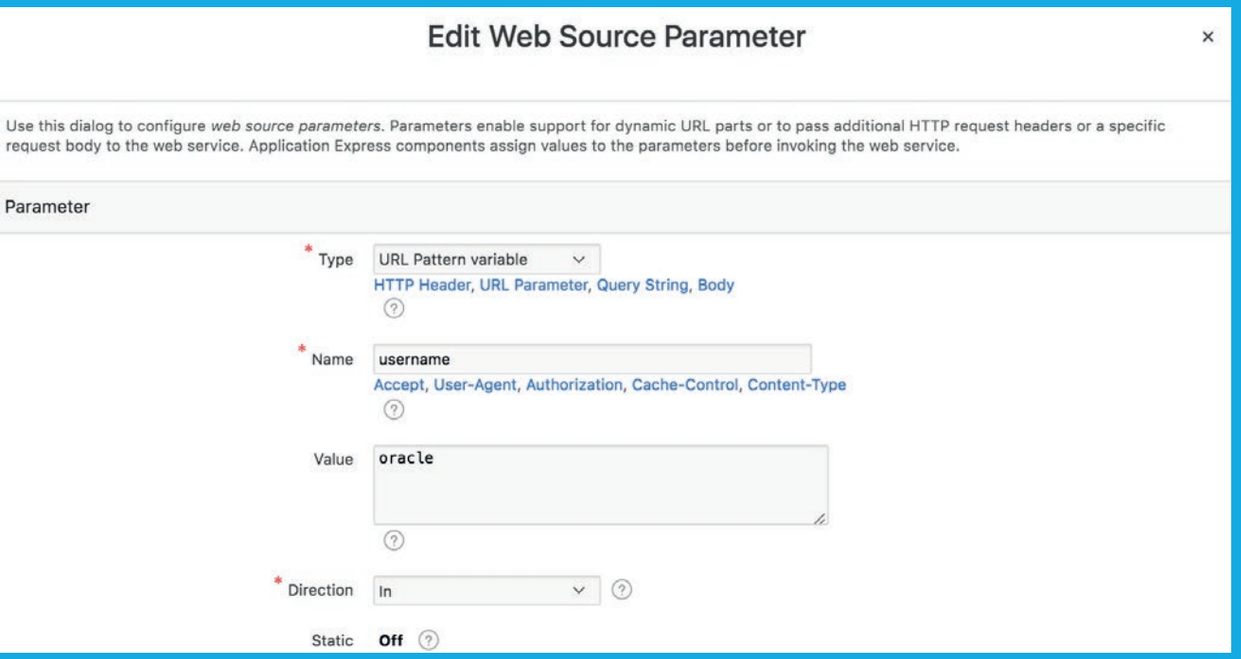


Fig. 4: The "username" part of the Github URL is dynamic

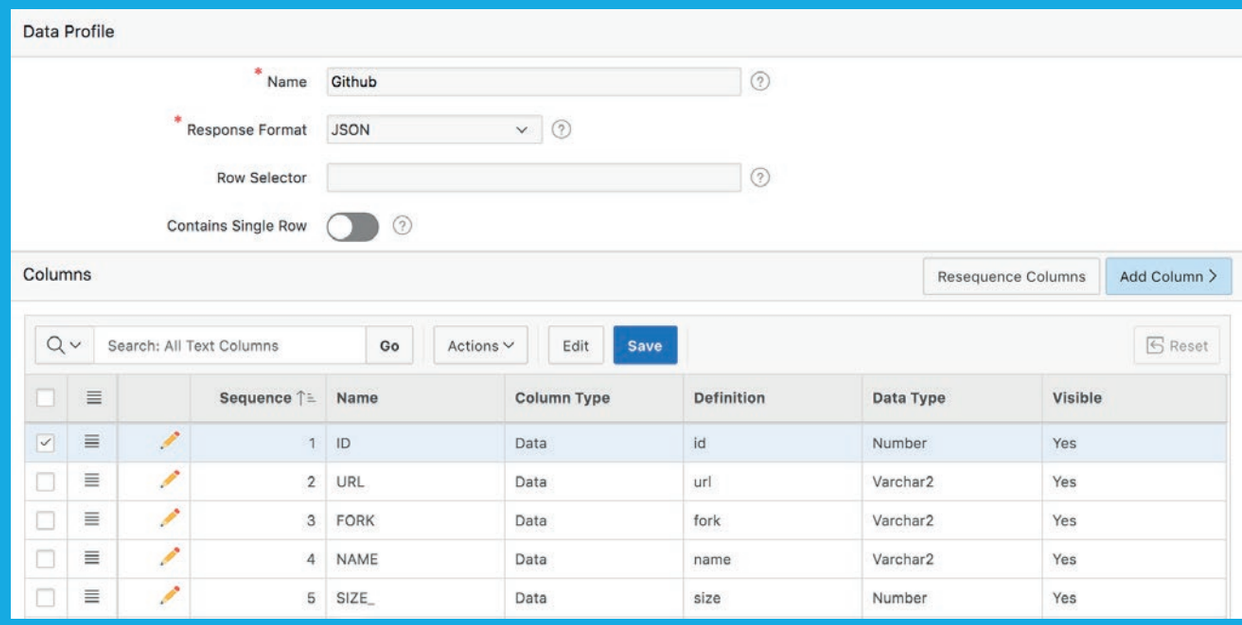


Fig.5: The Data Profile contains all information about the JSON response

If an APEX component supports Web Source Modules, the Create Page wizard will contain the **Data Source** step shown in Fig. 6. Instead of executing a SQL Query in the **Local Database**, the developer chooses the **Web Source Module** location, and one of the existing Web Source Modules. The columns in the shuttle item are directly derived from the **Data Profile** definition. In many cases it's advisable to select only a subset of these columns. For instance, the Github API returns 95 columns. Many of these contain URLs to other Github API endpoints, so displaying in a report makes not too much sense.

When reviewing the resulting page in Page Designer (fig. 7), we can see how Web Source Modules are integrated into the APEX engine: There is no generated SQL query or PL/SQL code. Instead we see the declarative information, that this report gets its data from the Web Source Module defined in Shared Components. All HTTP request execution and response parsing happens inside the APEX engine.

At runtime, the classic report behaves like every other classic report: features like column Format Masks or HTML Expressions will work as usual. The end user will not be able to recognize differences to a report based on a SQL query.

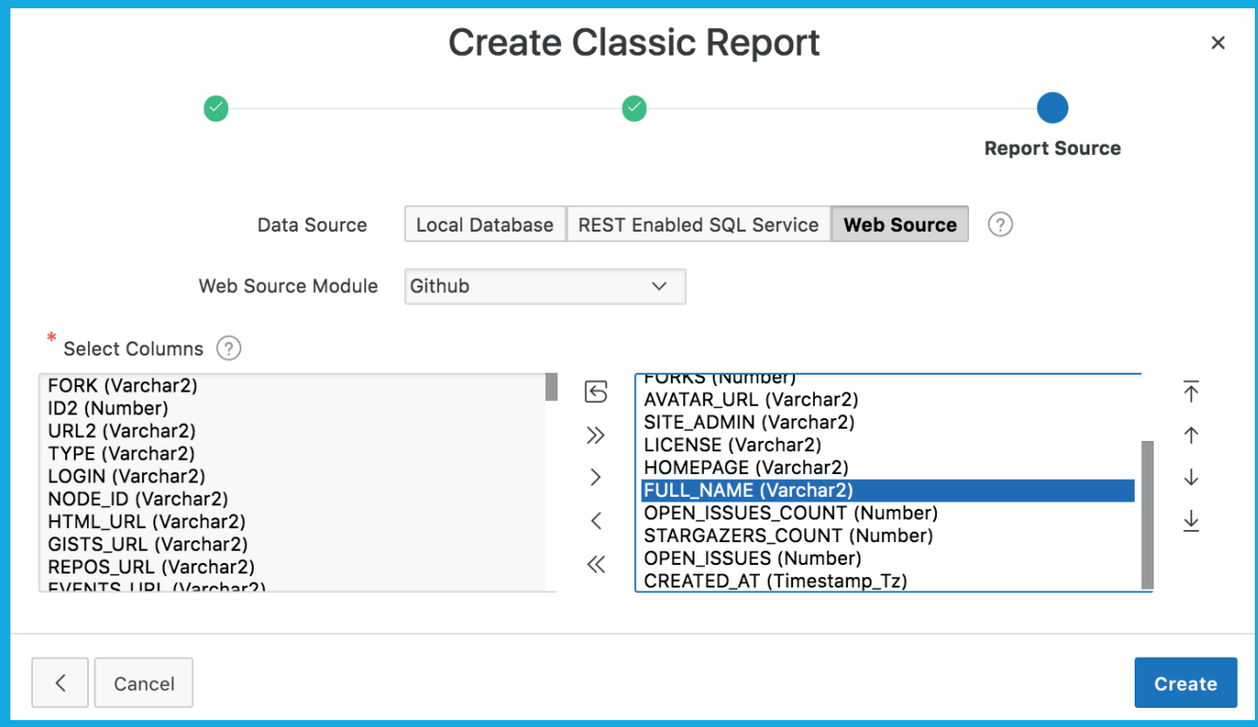


Fig. 6: Create Page wizard: Simply pick the Web Source Module as Data Source

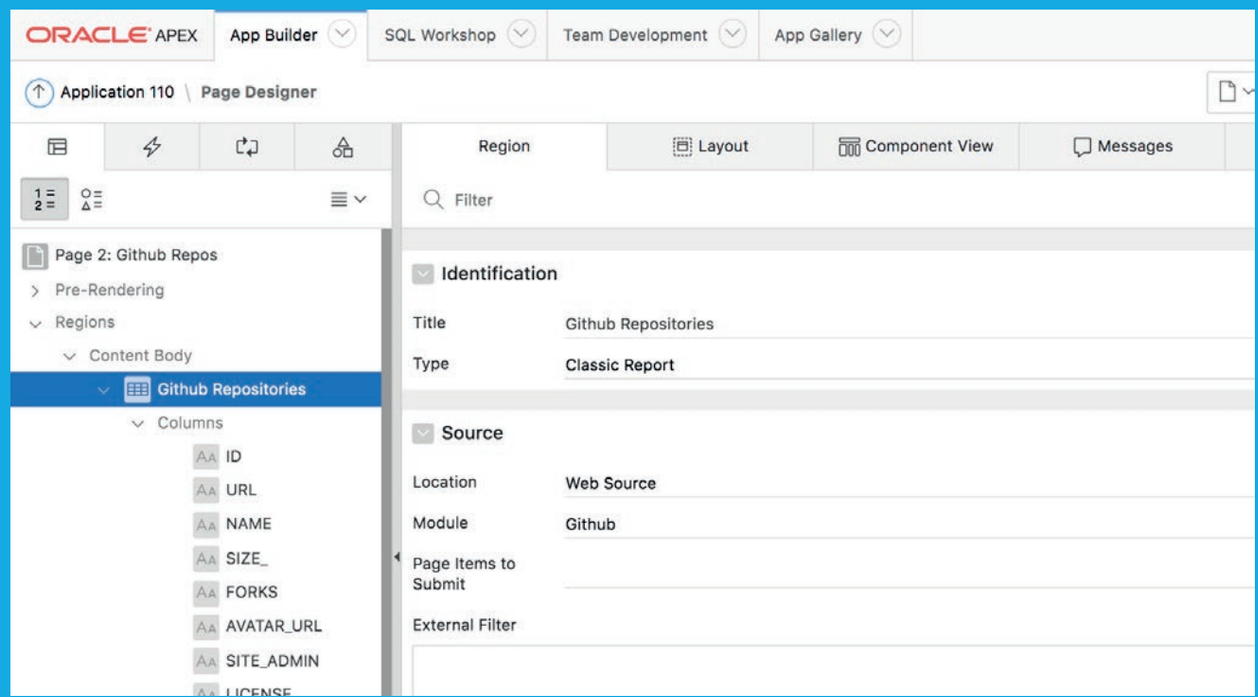


Fig. 7: Classic Report on a Web Source Module in Page Designer

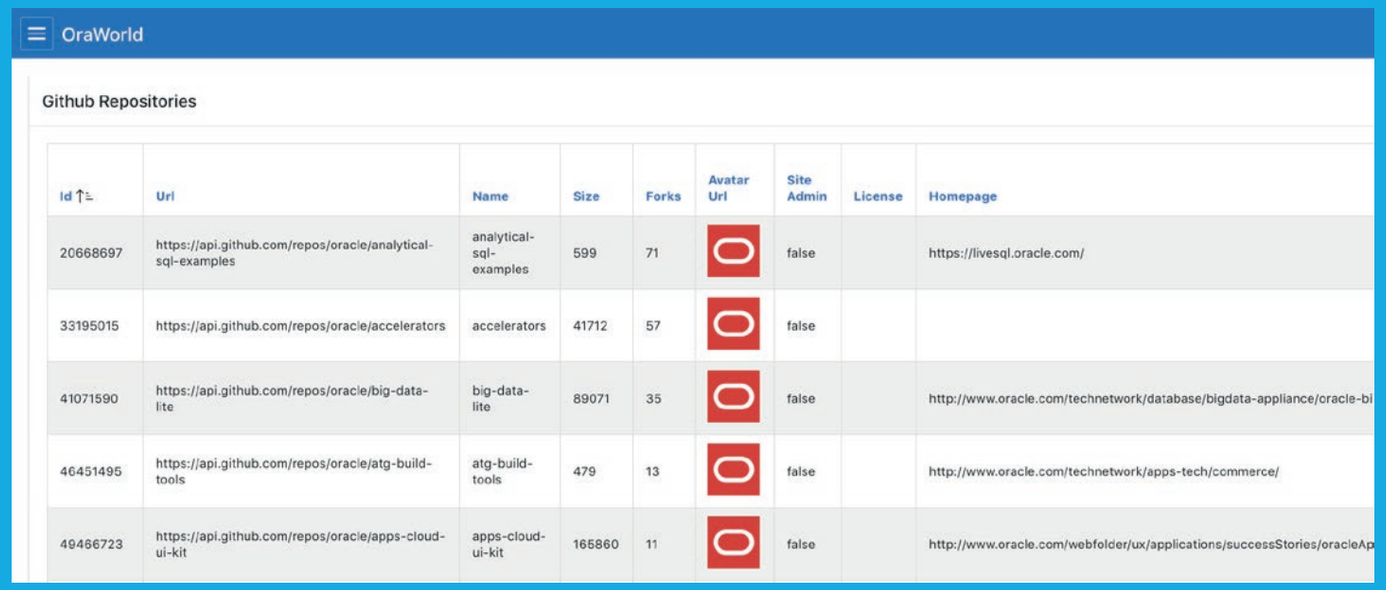


Fig. 8: Classic Report on a Web Source Module in Action

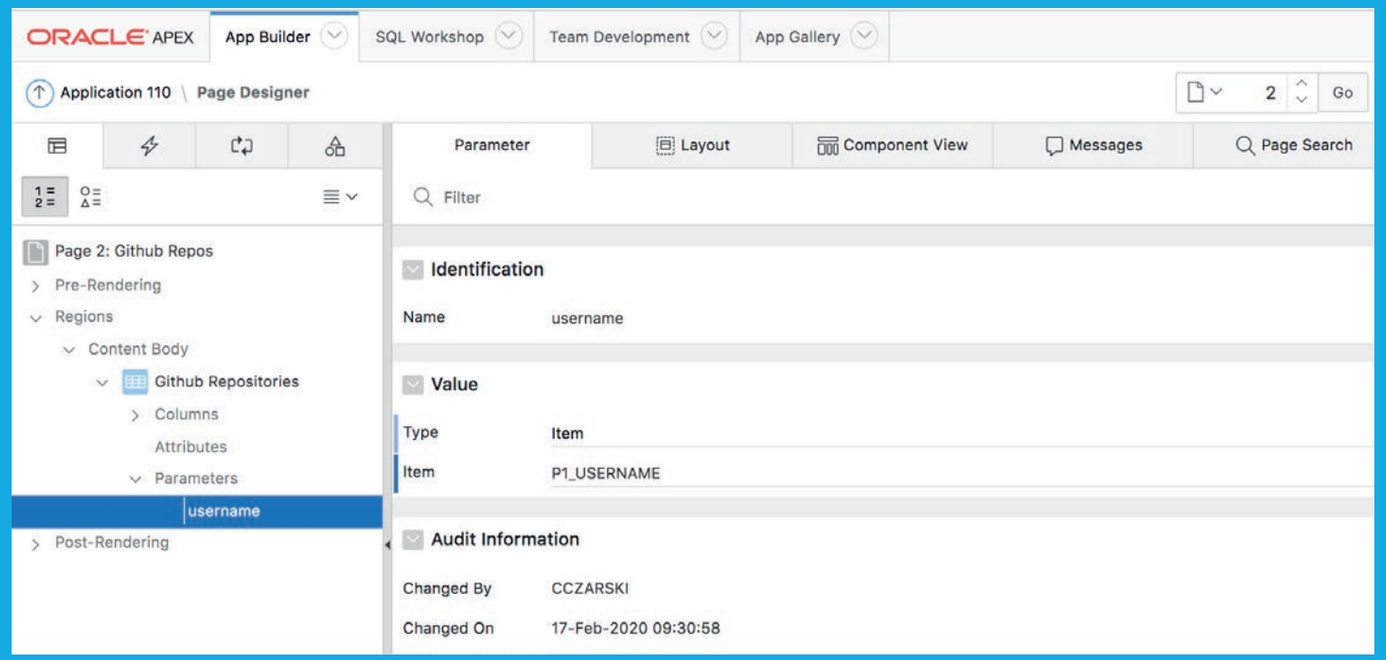


Fig. 9: Link a page item to the Web Source Module parameter

Remember that the Web Source Module's URL endpoint contains a placeholder (":username"), together with a declarative **Web Source Module Parameter**? Fig. 9 shows, how this **username** parameter is linked to a page item, so that the value of **P1_USERNAME** will be used for REST service invocation.

Parameters illustrate very nicely, how Web Source Modules act as an abstraction between the APEX component and the actual REST endpoint. For the APEX component, there is just a parameter named "username" (e.g. to be linked to a page item). The classic report does not care about whether this will be

passed as a HTTP Header, a cookie or as a part of the URL; these details are all encapsulated within the Web Source Module.

## Web Source Modules and PL/SQL

Web Source Modules are not limited to declarative usage in APEX: PL/SQL developers can use them with the **APEX_EXEC** package. But why should a PL/SQL developer use a Web Source Module and not the APEX_WEB_SERVICE package?

The *abstraction* (again) of Web Source Modules is a good reason to use them even in custom PL/SQL code. A Web Source Module acts as a *unit*: It takes in parameters and it returns result data in rows and columns. *How* the input parameters are processed and *how* a JSON response is turned into rows and columns, is the "implementation" of the Web Source Module. A consuming PL/SQL unit is only dependent on the interface of the Web Source Module, and not on details of the REST invocation. Listing 3 illustrates how to access the Github Web Source Module with PL/SQL and the APEX_EXEC package.

```
declare
  l_context apex_exec.t_context;
  l_columns apex_exec.t_columns;
begin
  -- We need an APEX session in order to access the WSM.
  apex_session.create_session( 110, 2, 'CCZARSKI' );

  apex_exec.add_column( l_columns, 'ID' );
  apex_exec.add_column( l_columns, 'NAME' );
  apex_exec.add_column( l_columns, 'SIZE_' );
  apex_exec.add_column( l_columns, 'LANGUAGE' );

  l_context := apex_exec.open_web_source_query(
    p_module_static_id => 'Github',
    p_columns          => l_columns );

  while apex_exec.next_row( p_context => l_context ) loop
    dbms_output.put_line(
      '- ' || apex_exec.get_varchar2( l_context, 'NAME' )
           || ' ( ID '
           || apex_exec.get_varchar2( l_context, 'ID' ) || ')' );
      dbms_output.put_line(
        '  ' || apex_exec.get_varchar2( l_context, 'LANGUAGE' ) );
      dbms_output.put_line(
        '  ' || apex_exec.get_varchar2( l_context, 'SIZE_' ) );
      dbms_output.put_line( '  ' );
  end loop;

  apex_exec.close( l_context );
exception when others then
    apex_exec.close( l_context );
    raise;
end;
```

*Listing 3: Accessing a Web Source Module with PL/SQL*

## Summary

Web Source Modules, introduced in APEX 18.1, significantly improve the integration of external REST APIs with APEX applications. Details about the REST API are stored as meta data, so that APEX components like reports, forms or charts can use them declaratively.

APEX cares about all low-level technical details, like sending HTTP headers, building the URL or parsing the response. Developers work with rows and columns and can thus focus on their application and business logic.

- APEX 18.1 supports Web Source Modules for *read only* components. This includes Classic and Interactive Reports, Charts, the Calendar, Tree region and Toggle Column and

Reflow report regions. Plug-In developers can enable their Plug-Ins for Web Source Modules.

- APEX 19.1 enables Web Source Modules for form pages. If the REST API provides PUT, POST and DELETE handlers, and the Web Source Module contains meta data for these, then a form page on a REST API can be created like a form on a table.
- APEX 19.2 enables Web Source Modules for the Interactive Grid and shared lists of values.

## More Information

- Information and environment to test-drive APEX
  **http://apex.oracle.com/en**
- APEX tutorials (lookup the “REST Lab”)
  **https://apex.oracle.com/en/learn/tutorials/**
- 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.

# Oracle and Azure: The New Frontier

Kellyn Gorman

***I like learning about new technology, and I went to Microsoft to learn just that:  a new analytics platform and discoveries in AI. After six months in my position, one thing became very clear – technology is an ever-connected web and data is at the center of it.  My customers wanted to move forward with a massive push to analytics, AI and machine learning, but to do so, they had to interweave their previous and existing data systems into the new. Doing this in the cloud posed an interesting challenge and many of those existing data sources were on Oracle, an expertise I possessed that wasn't common at Microsoft.***

## Keep it Simple

With the new opportunity to move Oracle into the Azure cloud, I approached it as I'd always had, by breaking down the complex into the simplest of terms: a block is a block and data is data, so I took the time to verify that I could use these blocks and build out the data, no matter what the platform was. In this case, it was Infrastructure as a Service in Azure, which I quickly realized wasn't as challenging as I first thought. Why hinder my customers by limiting where data could reside? The network was more likely to be our bottleneck, so the closer I could get the data to their solutions, the better off any technology communicating with Oracle was. I began to build out Oracle on Azure Infrastructure as a Service (IaaS) in January 2019 as a solution for a few customers and it quickly grew.

No way was I the first person to do this type of work inside or outside of Microsoft, but I knew I had a unique perspective on how to move Oracle workloads efficiently. I didn't want to focus on Azure for the solution, but on what Oracle products could partner with Azure to provide success.

Oracle Enterprise Edition with diagnostic and tuning pack, provides incredible value in the way of the Automatic Workload Repository (AWR) that could be used as the cornerstone for how we moved these workloads to the cloud. I didn't focus on the hardware in any way, as hardware was to be left behind in the end. The goal was to migrate the workload and to do this, AWR was the key.

The reasons for the migrations were varied – many were migrating to Azure for the simple reason their applications or databases were already in Azure. The goal was to have these valuable and vital data sources near their other Azure resources. There wasn't dissatisfaction with the product, but just a need to centralize and simplify their cloud journey. Surprisingly, not having a PaaS or SaaS for Oracle didn't hinder my customers from taking their Oracle journey to Azure.

For Oracle to be migrated to Azure, an Infrastructure as a Service (IaaS) solution is the only option currently for Oracle database and customers simply wanted to know we had the expertise to offer them success. Back in June 2019, if customers still desired a Platform as a Service (PaaS) solution, we currently engage our Oracle Cloud Infrastructure (OCI)/Azure partnership, which adds another layer of offerings on top of what we already possess. Will we have a PaaS solution for Oracle on Azure? We continually discuss this with our customers and are surprised that they discover IaaS suits their needs more often.

The Oracle Cloud and Azure partnership also serves for most of my customers that are dedicated to Real Application Clusters (RAC). Although we do have two partner solutions for RAC: FlashGrid and Cloud Simple from VMWare, I request a full review to verify a real need for RAC before I will choose this route. We've verified most existing RAC databases can be architected with Active DataGuard and robust Azure VMs to both scale up and scale out the workload successfully.

## Scaling Out with a Cloud Solution

Understanding what is available is important for success. There are always cheap deployments from cloud vendors to try out their services, but for Oracle, rarely are these offerings going to be satisfactory for a relational database deployment. The **Azure Calculator** can give you some insight on what is available, but a few tips can go a long way to ensuring greater success:

- Different **VM series** are available in different regions. Choose the region that is closest to your location but note the different offerings that are available per region.
- Although many VM Images are available via the portal, there are more available from the command line. If you don't see what you need, check the catalog with an Azure CLI command. You can use a number of different filters and add Linux utilities to assist:

```
az vm image list --offer Oracle --all --publisher
Oracle --output table
```

The output will look similar to this, but list everything for Oracle databases. The following won't count all the Linux VMs etc.:

```
Offer                  Publisher    Sku
---------------------  -----------  -------------
Oracle-Database-Ee         Oracle   12.1.0.2
Oracle-Database-Ee         Oracle   12.2.0.1
Oracle-Database-Ee         Oracle   18.3.0.0
Oracle-Database-Se         Oracle   12.1.0.2
Oracle-Database-Se         Oracle   12.2.0.1
Oracle-Database-Se         Oracle   18.3.0.0
```

- Make sure you use the correct OS version and if a built in HA solution is desired by the customer.
- Although D series and E series can look attractive to the customer from the price, they often don't have the power to keep up with the demands of an Oracle database. Use the AWR report to view CPU and Memory usage to size accordingly. You'll most often find the **M, F and G series** meets the requirements needed for different Oracle workloads.
- Use an **Automatic Storage Management** (ASM) instance to help simplify the storage management. It will also assist in TEMP management in the database.
- As **storage** is completely separate from the VM choice, allocate P40-P80 disks to match the IO requirements, (again, gathered from the AWR report) and use ultra-disk for redo logs, which are a known constraint when migrating to the cloud.
- Consider the impressive savings that can be had by using the 1- or 3-year reserved pricing over pay-as-you-go.

There are additional specifics that my team looks for when migrating, but these are an excellent starting point for migration and more information can be found form Microsoft on best

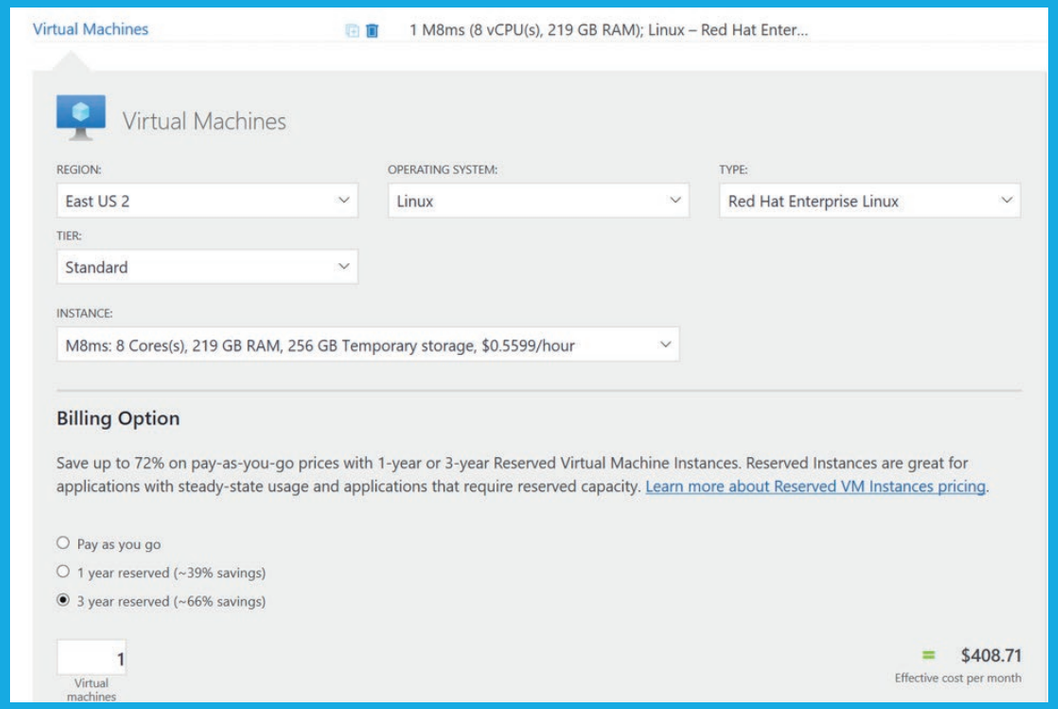


*Fig. 1, Azure Pricing Calculator for an Oracle VM, sans storage.*

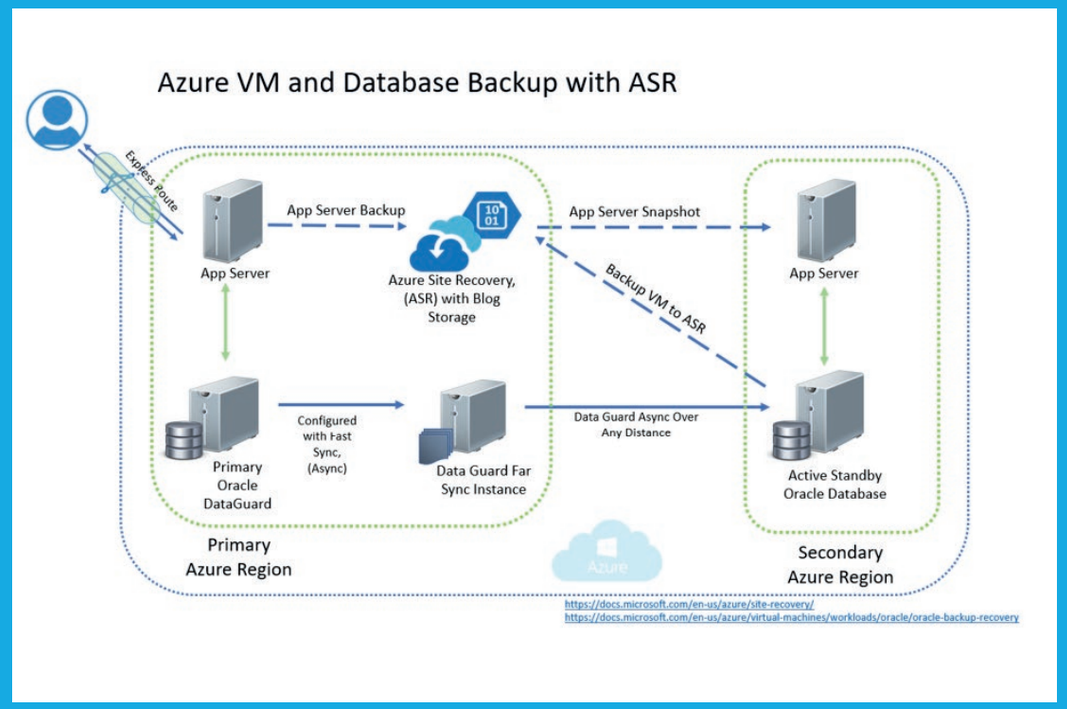


*Fig. 2, Azure VM and Database Backup with ASR.*

practices when migrating Oracle to Azure.

## Disaster Recovery and High Availability

One of the most common conversations I have revolves around how we may architect an environment on-prem isn't how the layer we're responsible for in the cloud may be architected. There is considerable redundancy built into a cloud infrastructure and services available that can replicate on-prem products and are no longer required once the migration to the cloud is done.

With architecture like the above diagram in place, there are significant changes in requirements for hardware and resources vs. the older environment.

Oracle DataGuard secondaries can be in several geo regions, safeguarding if a regional data center were to suffer an outage. **Oracle DataGuard Far Sync** can be used to transport logs to DataGuard secondaries over vast distances and protect transactional data from the primary to secondaries.

We configure **Oracle DataGuard** with automatic failover to secondaries, even multiple failover configurations for customers who need it. An application that can be made "RAC aware" to failover in a cluster situation can just as easily be enhanced to failover to an Oracle DataGuard database.

Azure Site Recovery (ASR) can be used to take **snapshots** and create replicas for full VMs with Oracle databases residing on them. Although some configurations can impact the ability to create a consistent snapshot, a final step of recovering the database to a point in time can be performed, which saves time from a full RMAN recovery.

An **RMAN repository catalog** can be built on an inexpensive, Azure VM to support the Oracle environment built in Azure. This offers a small performance improvement from having control-file history retained inside the database and offers an Oracle DBA to have the same trusty backup repository in Azure they had had on-prem without difficulty and backups can be performed inexpensively to **Azure Blob Storage**.

## BASH Is All

When I joined Microsoft, I assumed I would need to learn how to write PowerShell better and my existing BASH skills would gather dust. In a matter of a few months, I was writing deployments all in BASH scripts, adding in **Azure CLI** commands to the mix.  The amount of Linux environments in Azure surprises most. I use Linux VMs and BASH almost exclusively in my day-to-day work. I've never been one to use graphical user interfaces, so chose to adopt the Azure cloud shell for most of my work. By adding **Azure cloud storage**, I've created a location for my suite of scripts that I use in my day-to-day job at Microsoft. I deploy environments, perform administration tasks an automate work that I need to perform more than twice.

With the **Azure Cloud Shell**, I'm able to configure my environment with everything I need to work just as I did with Putty or other terminal emulators. This includes personalizing my interaction with the command line, including aliases, environment variables and profiles. All of this is configured to work with multiple accounts and be completely interactive with the Azure cloud **(fig. 3)**. As much as I love my BASH scripts, if I have a solution in PowerShell that I want

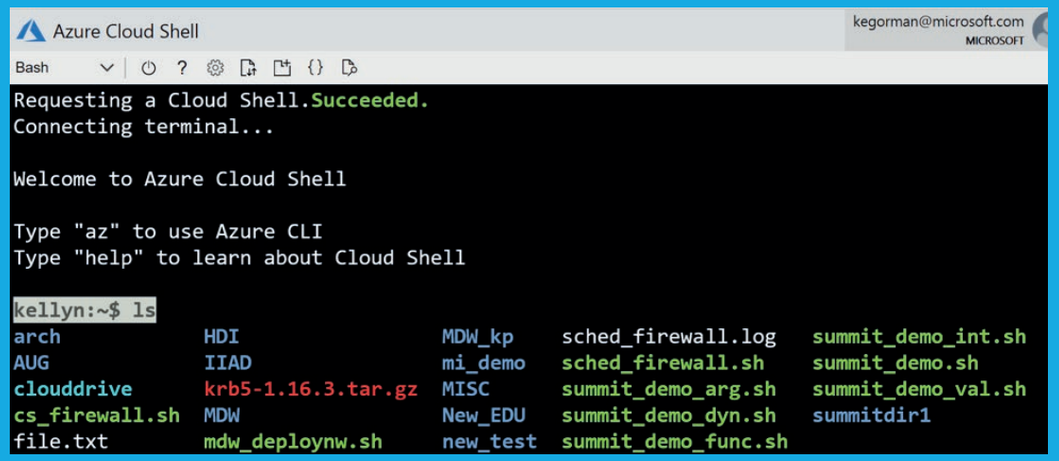


Fig. 3, The Azure Cloud Shell with color identified interface and personalized prompts.

to use with a BASH command or script, I can from the Azure Cloud Shell, it just requires me to call the script from the command line like I would any other script, application or utility:

```
PWSH <insert script name>.ps1
```

I'm able to version control all **my scripts in Github** and share them with team members or even the community. As other members of the community and Microsoft are doing the same, I'm given the opportunity to collaborate with folks I might never had the chance before the Azure cloud.

## Patch Automation

For many companies, the move to the cloud is to automate patching and upgrading of the OS and software layer. Upon investigation, my customers were pleased to discover Azure can automate Linux VM patching by enabling/scheduling the **Update Management** feature in the portal. This only leaves the database and application patch automation and once discussion are underway, we discovered, outside of new projects, automation of patching is riddled with hurdles that render it unfeasible for many. For those that still want patch automation, I quickly let them know, along with the full infrastructure tools available in **Oracle Cloud Control**, also built in the Azure cloud IaaS VMs, we can add the Lifecycle and Cloud management packs, implementing the **patch management feature** from Oracle for their new cloud environment.

With all these products and features, from both Oracle and Microsoft Azure, we're able to create a powerful and satisfying solution for customers who wish to incorporate their Oracle data sources in Azure IaaS VMs as part of their cloud journey.

## About Kellyn Pot'Vin-Gorman

Kellyn is a Customer Success Engineer at Microsoft specializing on Oracle and data platforms on Azure. An alumnus of both Microsoft's Idera ACE and Oracle ACE Director programs, a Friend of Redgate, she has been recognized with numerous awards over the years for her technical contributions and community volunteerism. She is one of only six women part of the Oak Table, a network for the Oracle scientist. She has extensive experience in environment migrations, optimization, automation and architecture. Kellyn is well known for her technical content and thought leadership through her presentations, keynotes, webinars, publications and engaging with her on social media presence as DBAKevlar or her blog, **dbakevlar.com**.

# How to Clone your Autonomous Database

André Ontalba & Rodrigo Mufalani

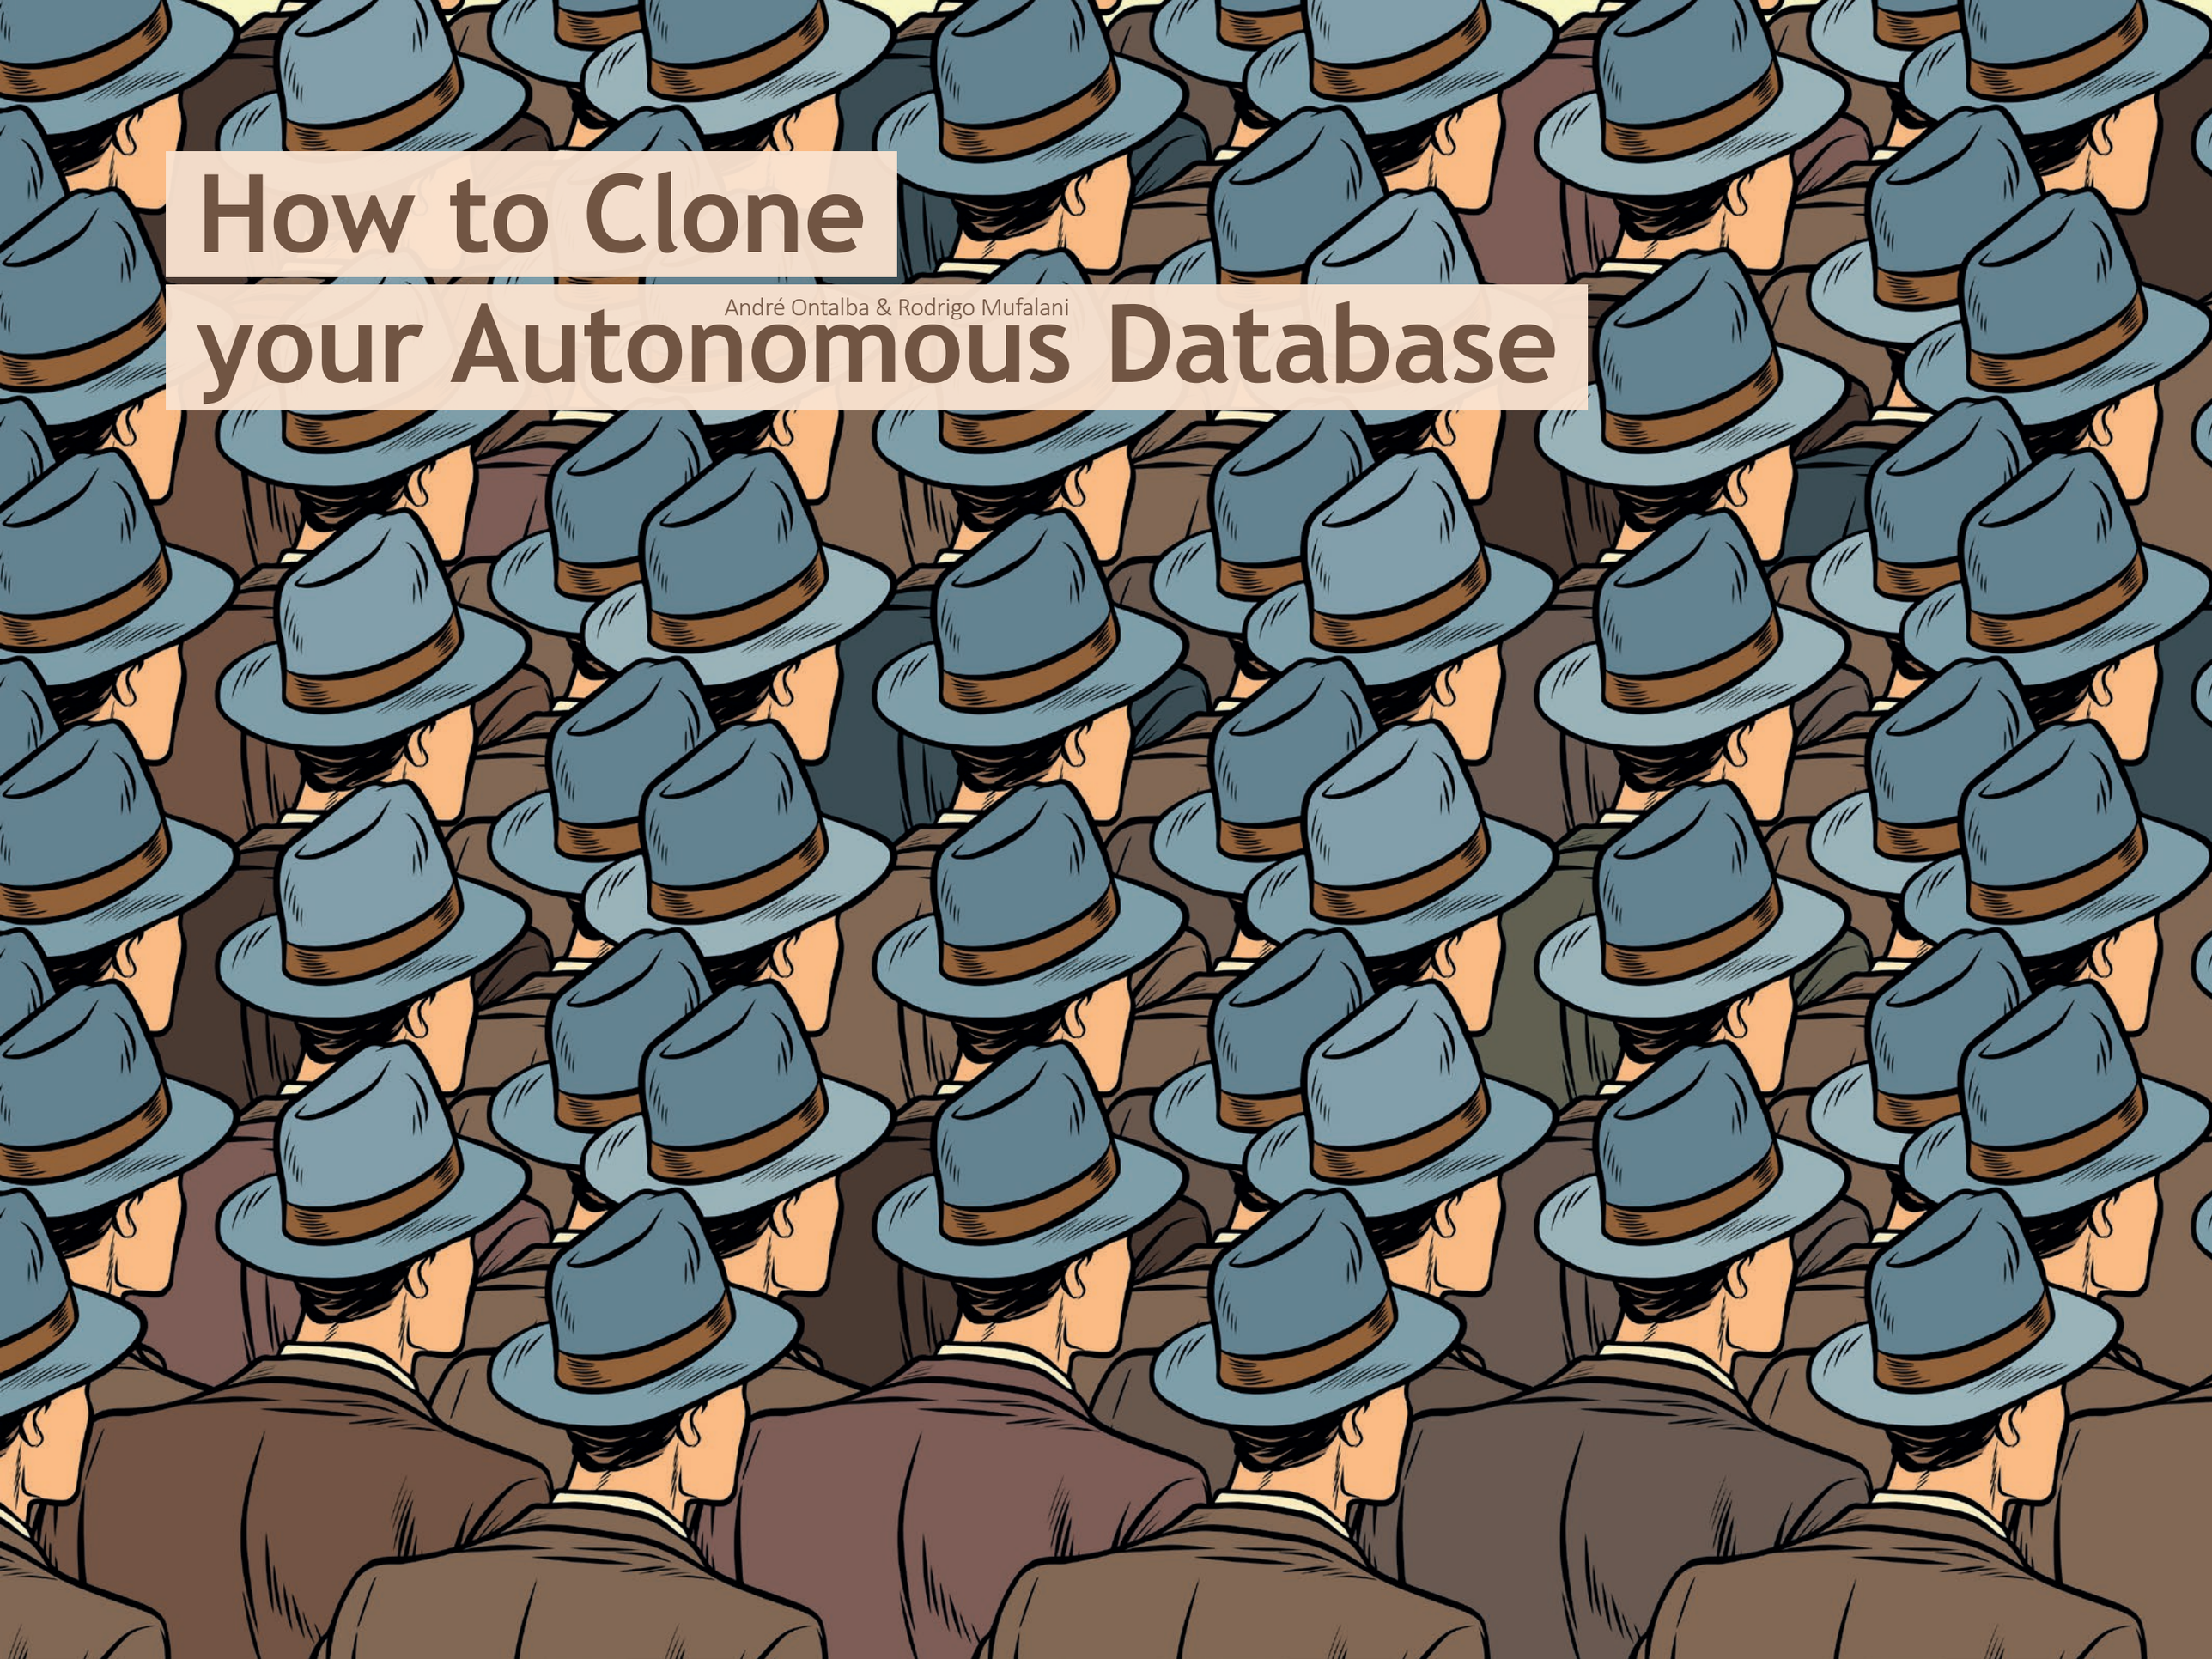

The Oracle Cloud Infrastructure Console allows you to clone an existing Autonomous Database. You may wish to use this cloning feature to create a point-in-time copy of your Autonomous Database for purposes such as testing, development or analytics. If you need to clone only the database schema of your source database, the «metadata clone» option is a quick and easy way to accomplish this task.

To get started, log into Oracle Cloud and click the "Create an ATP database" link **(fig. 1)**.

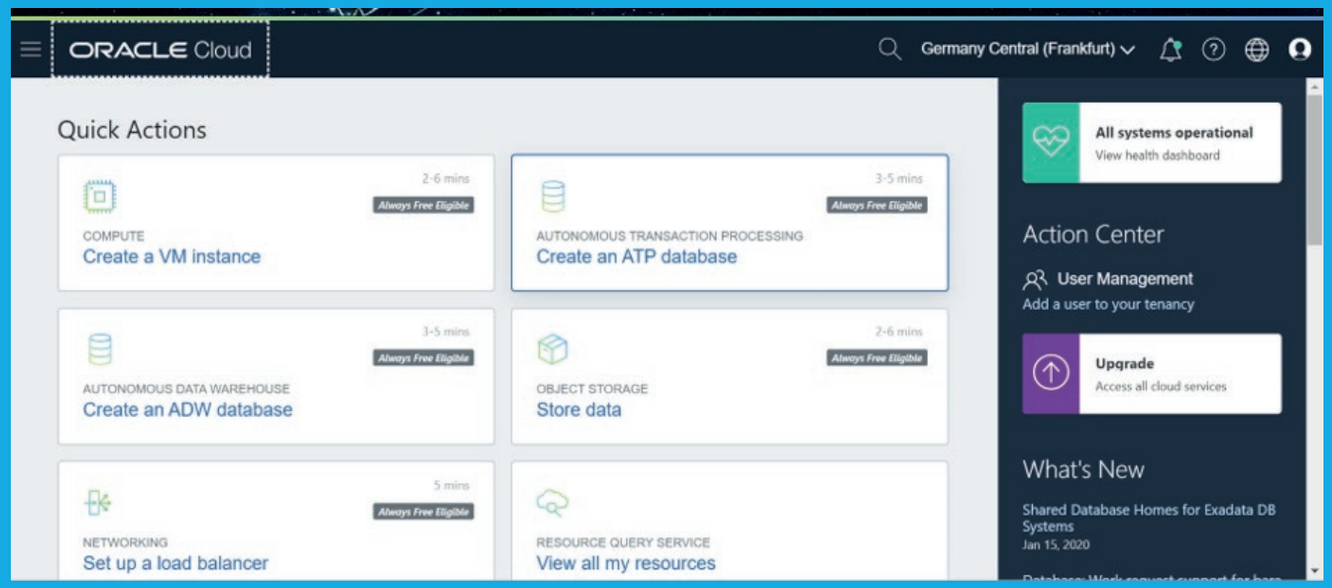

Fig. 1

Fill in the information and select the options to create the ATP instance.

*Choose a compartment:* dbadutra(root)

*Display name:* DBCLONE

*Database name:* DBCLONE

*Choose a workload type:* Transaction Processing

*Choose a deployment type:* Shared Infrastructure

*Configure the database:* In this case, as we use Always Free we can't change the CPU and storage size **(see fig. 2)**.

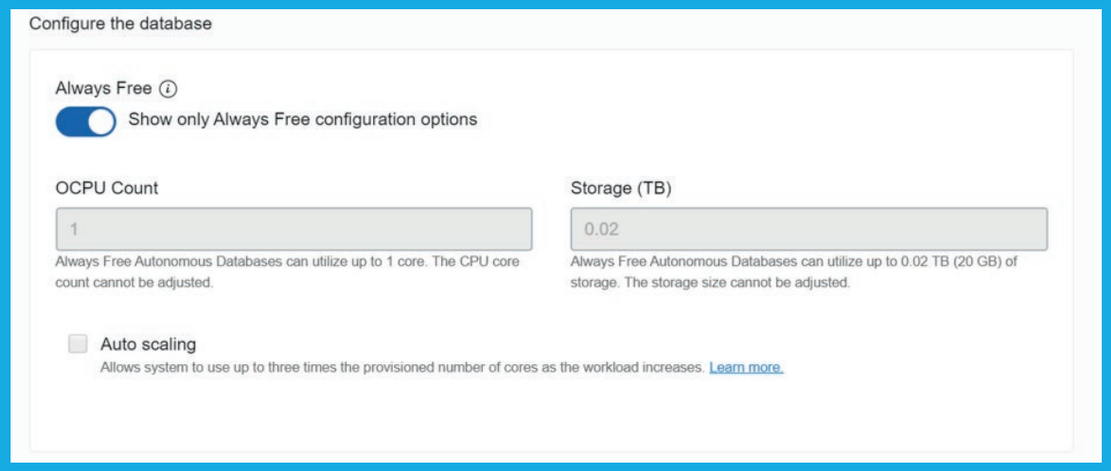

Fig. 2

*Create administrator credentials:* OracleATP2020

*Choose a license type:* In this case select License Included

After you have filled everything in, click "Create Autonomous Database" and wait a few minutes for the ATP being created **(see fig. 3)**.

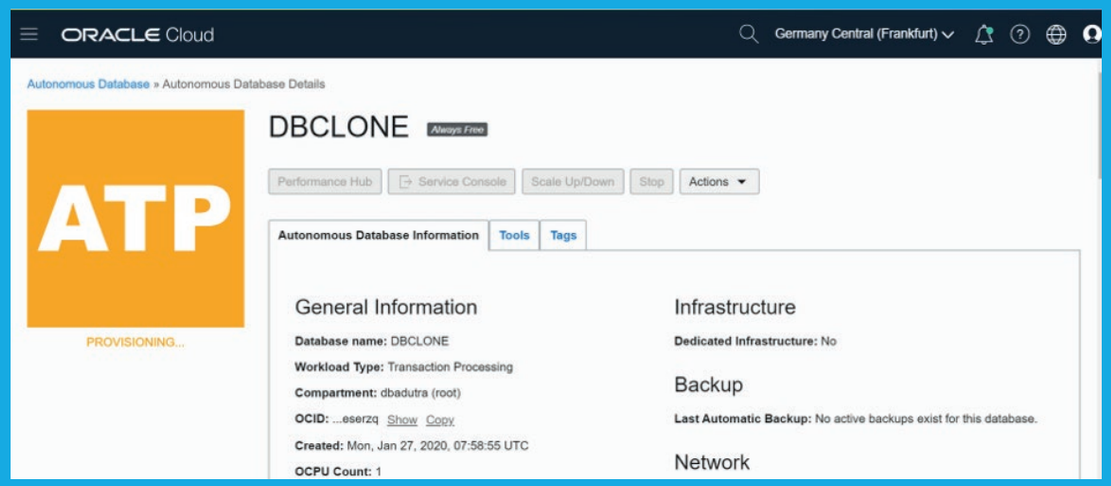

Fig. 3

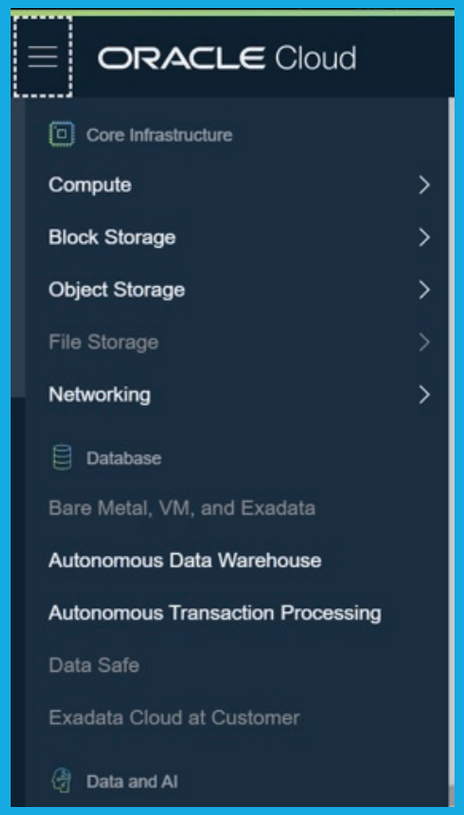


Fig. 4

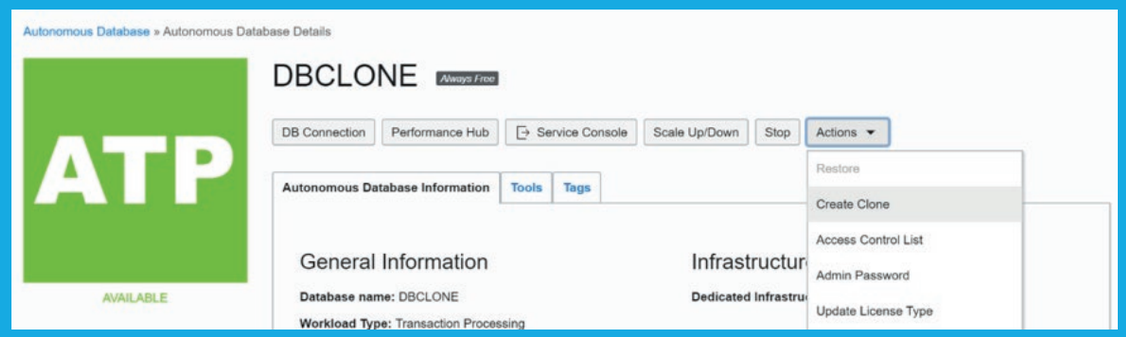


Fig. 5

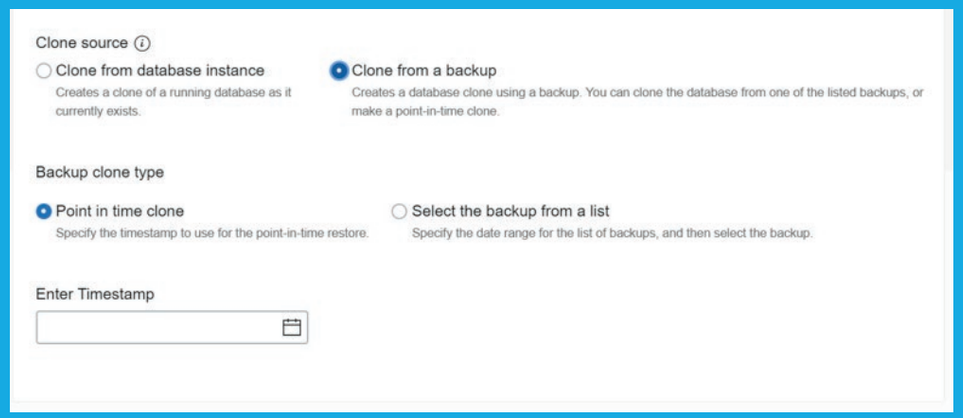


Fig. 6

When the ATP is available, select “Autonomous Transaction Processing” in the menu **(see fig. 4)**, click in “DBCLONE” and then select “Create Clone” in the “Actions” menu **(see fig. 5)**.

Now let’s create a clone of our ATP, filling in the information below:

*Clone Type:* In this option, we can choose a complete clone or just the structure without data. In this case we will select “Full Clone”.

In this option we can choose a clone from the Database Instance or use a backup to perform this clone and then perform a point in time recovery **(see fig. 6)**.

In this case, we select “Clone from the database instance”.

*Create In Compartment:* dbadutra (root)

*Display name:* Clone of DBCLONE

*Database name:* DBCLONE2

*Configure the database:* In this case, as I use Always Free, we can’t change the CPU and storage size.

*Create administrator credentials:* OracleATP2020

*Choose a license type:* In this case select License Included

Now, we click on “Create Autonomous Database Clone” and wait a few minutes for Clone ATP creation.

Then, our clone is created and available for use **(see fig. 7)**.

We hope we helped you with this great new feature. To learn more on the Autonomous Database Clone check here: **https://docs.cloud.oracle.com/en-us/iaas/Content/Database/Tasks/adbcloning.htm**


## About André Luiz Dutra Ontalba

André is an Oracle ACE member who graduated in Computer Science and specializes in Oracle Database with solid knowledge in Engineered Systems, Performance & Tuning, RAC, Oracle Cloud and Oracle ERP's System; He has been working with Oracle for 17 years, certified OCP Oracle 11 / 12g / Cloud and has more than 27 other certifications in Oracle products. He currently works as a Senior Database Architect at Sogeti Luxembourg, a Capgemini Group company. André is also the founder of the Luxembourg Oracle Users Group (LUXOUG) and a writer for OTN, GPO (Oracle Brazil User Group) and LUXOUG.

Twitter: **@aontalba** / blog: **www.dbadutra.com**

## About Rodrigo Mufalani

Rodrigo is an Oracle ACE member and Oracle Certified Master (OCM) with 15 years of experience. He started with Oracle 8i, but had the opportunity to support Oracle 7.3.4 onwards. He is an Oracle database specialist with a primary focus on Engineered Systems, Performance & Tuning and RAC. He is the founder and president and also a speaker for the Luxembourg Oracle User Group. He has spoken at OTN LAD TOUR, OTN EMEA TOUR and other Oracle events. Rodrigo currently works as Principal DB Architect at eProseed Europe. He was the third Oracle ACE to be nominated in Brazil.

Twitter: **@mufalani** / blog: **Mufalani.worpress.com**

# What's Your Super-Power? Mine is Autonomous Database and Machine Learning (Part 2)

*Jim Czuprynski*

So, you've poked around in the Oracle Machine Learning (OML) tools that are part of your Always Free Oracle Cloud subscription, and you're no longer afraid that the robots are coming for your job anytime soon. But now it's time to apply what you've learned to a real-world situation. Where should you start?

The answer is simple: *Everywhere you look.*

Continuing the **previous article** in this series, Oracle ACE Director and Senior Enterprise Data Architect Jim Czuprynski from Viscosity NA builds upon the basics of AI and ML he explored earlier to experiment with a practical example of leveraging machine learning algorithms: analyzing voter election data to determine the potential of "flipping" a voter to another political party during an upcoming general election. Jim also shows how to verify the accuracy of the machine learning model obtained, and concludes with an exploration of applying the powerful analytic and visualization capabilities of Oracle Analytics Cloud (OAX) to voters' demographic data to understand where the candidate's re-election campaign team should focus its resources to accomplish its goals.

## Developing Our Super-Powers: Where Are We So Far?

In the **previous article** in this series, I demonstrated several aspects of the potential power of Oracle Machine Learning—especially in the context of Oracle Autonomous Database – including:

- Creating an OML user, configuring it for use against both an ATP and ADW schema in Autonomous Database
- Leveraging Zeppelin Notebook features for tabular reporting and data visualization tools
- Exploiting Oracle's **DBMS_DATA_MINING** package to create a simple ML Time Series analysis and visualize the results from applying that algorithm

This article will delve deeper into using ML techniques to discover previously-unseen patterns hiding in plain sight within my data by leveraging the sophisticated data mining tools already built into Oracle Database. This time around, I'm going to apply ML to some simple use cases that every one of us can appreciate because most of us will have had at least some exposure – painful or otherwise! – since our early days in secondary school. And no, I'm not going to discuss my personally embarrassing experiences about the class I was least successful at and most loved to hate: physical education.

## Let's Talk Politics.

One of the more challenging parts of discussing ML in a context we can all appreciate is finding a wide-enough topic with sufficient variables and data for building sufficiently-complex models based on real-world business cases while simultaneously discussing the often-complex underlying questions we're actually trying to answer without making those quandaries so obscure that no one understands why a data scientist might be interested in solving them.

Fortunately, I've hit upon a plethora of use cases that anyone within our global civilization can immediately appreciate: elections.

I have actively volunteered for a local campaign for a US Congressional seat since early 2018, so I've had direct exposure to some fascinating challenges that a modern political campaign has to endure to accomplish its goal of electing a candidate. Participating at this level has given me a virtual "catbird's seat" that lets me observe the trials and tribulations of a modern political organization.

To my great surprise, I've discovered that a political campaign's requirements aren't really that much different from those

of any modern business: Finite goals must be achieved with often-limited resources, and while making decisions on where, when, and how to deploy those resources can be fraught with uncertainty, data-driven decision-making adds quite a bit of confidence that the right choices have been made.

**Disclaimer:** The use cases I will be discussing are only meant to be exemplary, and they reflect no particular political orientation. In fact, any of these cases can be completely reversed so that the opposite party or candidate wins the election. That's the neat part of political science, and why it yields such excellent and comprehensible examples for data science as well.

## A Typical Challenge: Who Is Likely To Vote, and How?

Modern political campaigns in the USA are focused on identifying intimately the attributes of voters within each locality, township, county or parish, and state. In many ways, this is not unlike a sales organization that's trying to identify its best customers that are likely to continue to purchase their products for the long term, or who might be convinced to switch from a competitor's product to theirs instead. However, unlike a retail store, a campaign really does not know how likely an individual voter is to vote for their particular candidate.

For example, my resident state of Illinois, like 15 others in the USA, conducts semi-closed primaries. In other words, the voter must declare her intent to vote for a specific political party's slate of candidates when voting in that primary. Moreover, that declaration is recorded as a matter of public record, and any valid electoral campaign is permitted and encouraged to access that information when plotting strategies for canvassing via phone or in person in what are typically called Get Out The Vote (GOTV) efforts.

Tightly-secured data portals such as **VoteBuilder** and **i360** retain these public voting records. Every campaigns' data scientists thus have access to an astonishing array of information, including answers to survey questions about which issues most motivated voters during past and current campaigns. Of course – and this is important to stress here! – since our elections are conducted by observers from all political parties and closely monitored by trained election judges, the *actual ballots cast for candidates during elections are kept completely secret*.

## The Goal: Discover "Flippable" Voters

During my volunteer involvement with the electoral campaign I assisted in during the 2018 election, I was able to learn some valuable lessons about which voters the campaign most wanted to focus on to vote for its candidate. While it was relatively easy to identify which voters are extremely likely to cast their vote based on past Democratic party affiliation, the campaign was also interested in voters who might be "on the edge" and could be possibly shifted to vote for their candidate.

My volunteer activities in the 2020 campaign has given me several ideas for the campaign for locating these "flippable" voters, and I decided to deploy Oracle's data mining and modeling tools to attempt to identify them. My model is based on how many times a voter has voted for the party opposite of the candidate's party – either Republican or Independent - over the last several primaries as well as the 2018 campaign. Here's the assumptions I'll be using in these experiments:

- The voter voted for the opposite party at least three times in the 2010, 2012, and 2014 primaries.

- The voter then voted either Independent or Democrat in the 2016 and 2018 primaries.
- The voter did vote in the 2018 general election.

## Gathering Data Sources

To capture all necessary data for my analyses, I downloaded voting information from votebuilder.com for the US congressional district campaign for all voters and their corresponding voting history as a CSV file. (Obviously, the confidentiality of my fellow voters' balloting information prevents me from providing a representative sample of these data.)

I then loaded that data into an Autonomous Transaction Processing (ATP) Database I created specially for these purposes. Since ATP databases retain data only in encrypted format using the AES-256 encryption algorithm, my fellow campaign volunteer colleagues were assured that all voter data was sufficiently protected.

I created a new database schema named **VEVO** and then created two tables, **T_VOTERS** and **T_VOTING_RESULTS**, within that schema. I then loaded the data I had previously downloaded into those tables, created all necessary PK indexes, and gathered optimizer statistics. The code I used to complete all of these tasks is available in this **personal Github repository**.

## Selecting the Appropriate Data Mining Model and Algorithm

A key decision I need to make is which data mining model algorithm to apply to my source data. Fortunately, Oracle provides a brief but excellent **cheat sheet** that can be quite helpful in choosing the right model, or at least eliminating the models that are inappropriate to the use case at hand. It's also a great primer for discovering all the different models that the **DBMS_DATA_MINING** package does offer.

Based on that cheat sheet, it appears that a classification model is probably most appropriate because it allows me to determine the flippability of a group of voters based on one or more voter attributes, and then rank the attributes that are most significant to identifying which sets of voters the campaign should focus its limited resources to flip. And while for now I'm planning to predict voters' behaviors based on a single criterion – whether they've shown a tendency to vote for my candidate's party now after having voted consistently for the opposite party *in the past* – the *Decision Tree* algorithm offers a plethora of capabilities to classify targets in several dimensions.

**Note:** I'll revisit the power of this algorithm in the next article as well because not only does it allow me to apply a cost-benefit matrix to its discoveries, but it also lets me observe the rules it generates to make its classification decisions. I'll demonstrate how best to explore those rules via Oracle SQL Developer's Data Mining tools, but feel free to **take a look at those capabilities** in advance.

```
-----
-- Locate and score all voters who:
-- 1.) Voted Democratic in 2018 primary elections; and
-- 2.) Voted at least once for a non-Democratic candidate in the prior four
--     primary elections; and
-- 3.) Voted at least once in the last five general elections.
-----
CREATE OR REPLACE VIEW vevo.likely_converts
AS
SELECT
    v_id
    ,v_county_name
    ,v_situs_city
    ,v_situs_zip5
    ,v_precinct_name
    ,v_race_designation
    ,v_gender
    ,v_age
    ,likely_dem
 FROM
   vevo.voters
 ,(SELECT
         vr_dems.vr_v_id
        ,vr_dems.dem_count
        ,vr_nondems.nondem_count
        ,vr_consistent_voter.gen_election_count
        ,CASE
            WHEN (vr_consistent_voter.gen_election_count > 1)
             AND (vr_dems.dem_count > 0)
             AND (vr_nondems.nondem_count > 2) THEN '1'
            WHEN (vr_consistent_voter.gen_election_count > 1)
             AND (vr_dems.dem_count > 0)
             AND (vr_nondems.nondem_count > 0) THEN '1'
            ELSE '0'
        END AS likely_dem
    FROM
     -----
     -- Which voters have voted Democrat in the 2018 primary?
     -----
     (SELECT
            vr_v_id
           ,DECODE(vr_party_abbr, 'D', 1, 0) dem_count
      FROM vevo.voting_results
     WHERE vr_election_abbr = 'PRI2018'
     ) vr_dems
      -----
      -- Which voters have declared themselves as either Independent or
      -- at least twice in the prior four primaries?
      -----
    ,(SELECT
            vr_v_id
           ,COUNT(vr_election_abbr) nondem_count
       FROM vevo.voting_results
      WHERE vr_election_abbr IN ('PRI2010','PRI2012','PRI2014','PRI2016')
        AND vr_party_abbr IN ('I', 'R')
      GROUP BY vr_v_id
      ) vr_nondems
      -----
      -- Which voters have voted at least once in the last five
-- general elections?
      -----
    ,(SELECT
            vr_v_id
           ,COUNT(vr_election_abbr) gen_election_count
       FROM vevo.voting_results
      WHERE vr_election_abbr
          IN('GEN2010','GEN2012','GEN2014','GEN2016','GEN2018')
      GROUP BY vr_v_id
      ) vr_consistent_voter
  WHERE vr_dems.vr_v_id = vr_nondems.vr_v_id
    AND vr_dems.vr_v_id = vr_consistent_voter.vr_v_id
  ) vr_scoring
WHERE v_id = vr_scoring.vr_v_id;
```

*Listing 1. Constructing View VEVO.LIKELY_CONVERTS*

To implement the flippable voter model, I built a new view, **VEVO.LIKELY_CONVERTS**, that incorporates the criteria that I described above for their identification. **Listing 1** shows that new view.

## Building the ML Model

What I find fascinating about the Decision Tree algorithm is its ability to prove that it's actually producing reasonably accurate

results by determining just how trustworthy those results are. Here's a brief sketch of how I'm planning to accomplish this:

- First, I'll need to prepare two sets of input data:
  - The *training* set is a randomized collection of individual observations drawn from well over 50% of the complete sample set.
  - The *testing* set is built from the remainder of the original sample set.
- Hopefully, if I've done my best to provide sufficiently clean data to the model, I should see *lift* – in other words, evidence of the desired positive influence – upon a significant number of use cases in the training subset.
- To verify that the model is accurately calculating lift, I will also apply the model to the remaining randomly-selected data in the testing subset to determine if there are any anomalies like *false positives* and *false negatives*.

**Listing 2** shows how I prepared representative samples of the training versus testing data subsets. I used a simple randomizing method – calculating the modulus of each voter's unique id (**v_id**) and then selecting only those with a remainder between one and six – to randomly select 60% of all voters likely to be flippable – to capture my training sample into view **VEVO.CONVERTS_TRAINING_SAMPLE**. Likewise, I reserved the remaining 40% of potentially flippable voters for later evaluation in view **VEVO.CONVERTS_TESTING_SAMPLE**.

```
-----
-- Prepare training and evaluation views, splitting the data into
-- CONVERTS_TRAINING_SAMPLE (randomly selected 60% of the data) and
-- CONVERTS_TESTING_SAMPLE (the remaining 40%)
-----
CREATE OR REPLACE VIEW vevo.converts_training_sample
AS
SELECT *
    FROM vevo.likely_converts
  WHERE MOD(v_id,10) BETWEEN 1 AND 6;

CREATE OR REPLACE VIEW vevo.converts_testing_sample
AS
SELECT *
    FROM vevo.likely_converts
  MINUS
  (SELECT *
  FROM converts_training_sample
  );
```

*Listing 2. Constructing Views for Training and Testing the Decision Tree Model*

At last, it's time to build the Decision Tree ML model. I'll first drop the model (just in case it already exists) and then create a new one using procedure CREATE_MODEL2 of DBMS_DATA. MINING, as shown in **Listing 3**.

```
-----
-- Drop and rebuild Decision Tree Model
-----
BEGIN
  DBMS_DATA_MINING.DROP_MODEL('DT_VEVO_SAMPLE');
EXCEPTION
  WHEN OTHERS THEN NULL;
END;
/

DECLARE
  v_setlist DBMS_DATA_MINING.SETTING_LIST;
BEGIN
  -- Add settings
  v_setlist('PREP_AUTO') := 'ON';
  v_setlist('ALGO_NAME') := 'ALGO_DECISION_TREE';

  DBMS_DATA_MINING.CREATE_MODEL2(
     model_name           => 'DT_VEVO_SAMPLE'
    ,mining_function      => DBMS_DATA_MINING.CLASSIFICATION
    ,data_query           => 'SELECT * FROM converts_training_sample'
    ,set_list             => v_setlist
    ,case_id_column_name  => 'v_id'
    ,target_column_name   => 'likely_dem'
    );
END;
/
```

*Listing 3. Creating the Decision Tree Model*

```
-----
-- Train the newly-created model against the training sample in
-- CONVERTS_TRAINING_SAMPLE, calculating the “lift” of the model
-- as well
-----
DROP TABLE vevo.converts_applied_result PURGE;
DROP TABLE vevo.converts_lifted PURGE;

BEGIN
  -- Apply the testing sample against the testing sample
  DBMS_DATA_MINING.APPLY(
      model_name           => ‘DT_VEVO_SAMPLE’
     ,data_table_name      => ‘CONVERTS_TESTING_SAMPLE’
     ,case_id_column_name  => ‘v_id’
     ,result_table_name    => ‘CONVERTS_APPLIED_RESULT’
     ,data_schema_name     => ‘VEVO’
     );

END;
/

BEGIN
  -- Calculate and retain “lift” (i.e. the positive difference by
  -- applying the testing data)
  DBMS_DATA_MINING.COMPUTE_LIFT(
      apply_result_table_name     => ‘CONVERTS_APPLIED_RESULT’
     ,target_table_name           => ‘CONVERTS_TESTING_SAMPLE’
     ,case_id_column_name         => ‘V_ID’
     ,target_column_name          => ‘LIKELY_DEM’
     ,lift_table_name             => ‘CONVERTS_LIFTED’
     ,positive_target_value       => ‘1’
     ,score_column_name           => ‘PREDICTION’
     ,score_criterion_column_name => ‘PROBABILITY’
     ,num_quantiles               => 25
     ,cost_matrix_table_name      => NULL
     ,apply_result_schema_name    => ‘VEVO’
     ,target_schema_name          => ,VEVO‘
     ,cost_matrix_schema_name     => NULL
     ,score_criterion_type        => ,PROBABILITY‘
  );
END;
/
```

*Listing 4. Testing the Decision Tree Model*

## Trust, But Verify: Testing the Decision Tree Algorithm

Now that the model is built from the training sample, I’ll capture the results into a new table named

| Parameter | Value | Purpose / Description |
|---|---|---|
| apply_result_table_name | CONVERTS_APPLIED_RESULT | The table in which the results of the ML model’s training exercise will be recorded |
| target_table_name | CONVERTS_TRAINING_SAMPLE | The table in which the training data for the ML model exists |
| case_id_column_name | V_ID | The column that represents a unique case number for each data point – here, the voter’s unique ID |
| target_column_name | LIKELY_DEM | The column that defines the measurement |
| lift_table_name | CONVERTS_LIFTED | The table in which to store the results of the COMPUTE_LIFT calculations |
| positive_target_value | 1 | The column value in the target_column_name column identifies a positive result |
| score_column_name | PREDICTION | The column in the apply_results_table_name table contains the predictions against which lift will be computed |
| score_criterion_column_name | PROBABILITY | The column in the apply_results_table_name table that contains the scoring criterion |
| num_quantiles | 25 | How many quantiles across which to distribute the resulting computed lift factors (default: 10) |
| cost_matrix_table_name | NULL | An optional table in which the applicable costs associated with any misclassifications are stored |
| apply_result_schema_name | VEVO | The schema that owns the apply_results_table_name table |
| target_schema_name | VEVO | The schema that owns the target_table_name table |
| cost_matrix_schema_name | NULL | The schema in which the optional cost_matrix_table_name table resides |
| score_criterion_type | PROBABILITY | Determines whether *probabilities* (the default) or *defined costs* are used to score the computed lift |

*Table 1. Decision Tree ML Model Parameters and Controls*

**Note:** Those of you who have deep experience using DBMS_DATA_MINING realize that I've chosen mostly default values for these settings. In the next article in this series, I will explore these parameters and other potential settings for them in greater detail. I'll also illustrate how to manipulate some of these parameters from within an APEX application to fine-tune the resulting model in real time.

**VEVO.CONVERTS_APPLIED_RESULT** using the **APPLY** procedure of **DBMS_DATA_MINING**. I'll then use the results in **CONVERTS_APPLIED_RESULT** as input into the **COMPUTE_LIFT** procedure to capture the corresponding lift for each voter in the testing set so that I can evaluate how well the model is performing based on the sample input data. These steps are shown in **Listing 4**.

**Table 1** breaks down how each of the **COMPUTE_LIFT** procedure's parameters influence the ML model's training process.

## Reviewing the Model's Results

At last it's time to review the results of this exercise and the several levels of data it returned. Let's take a look at the resulting lift for potentially flippable voters across the 25 quantiles that the ML model has gathered. As the simple query in **Listing 5** shows, the **30,176** voters identified within the testing set are distributed equally across those buckets – about **1,200** per quantile - with increasing lift levels.

## Don't Tell Me, Show Me: Visualizing ML Model Results With Oracle Analytic Cloud (OAC)

Experienced data scientists realize the value of visualization techniques to determine quickly if the results of their data

```
-----
-- Show the resulting "lift" provided through ML statistics:
-----

COL quantile_number         FORMAT A08       HEADING "Quantile"
COL quantile_total_count    FORMAT 999,999   HEADING "Voter|Count"
COL cumulative_lift         FORMAT 9.9999    HEADING "Cumulative|Lift"

SELECT
     quantile_number Quantile
    ,quantile_total_count Qty
    ,ROUND(gain_cumulative,4) Cum_Lift
  FROM vevo.converts_lifted;

              Voter        Cumulative
Quantile      Count              Lift
-------- ---------- ---------------
       1      1,208            .0555
       2      1,207            .1100
       3      1,207            .1644
       4      1,207            .2189
       5      1,207            .2734
       6      1,207            .3257
       7      1,207            .3715
       8      1,207            .4172
       9      1,207            .4561
      10      1,207            .4950
      11      1,207            .5339
      12      1,207            .5728
      13      1,207            .6117
      14      1,207            .6499
      15      1,207            .6882
      16      1,207            .7265
      17      1,207            .7639
      18      1,207            .7901
      19      1,207            .8202
      20      1,207            .8502
      21      1,207            .8803
      22      1,207            .9103
      23      1,207            .9404
      24      1,207            .9705
      25      1,207           1.0000

25 rows selected.
```

*Listing 5. Decision Tree Model: Querying The Results*

modeling are yielding the results they expected, including recognition of missing data points, outliers, and trends. The **prior article** in this series demonstrated how to leverage *Oracle Machine Learning* (OML) that's already built into ADB

to do that, but there's another even more valuable set of data visualization tools available within OCI: *Oracle Analytic Cloud* (OAC).

> **Note:** Unlike OML, its smaller and less powerful sibling, OAC does require separate licensing. I'm only demonstrating it here to show how its data visualization tools compare to OML so that you can decide if it's worth the additional cost. I'll also delve much deeper into OAC's analytic capabilities in the next article in this series so that you can make an intelligent decision about including it into your IT organization's future data science tool plans.

After I followed the OCI wizard's steps to create a new OAC instance, the instance was available for access within a few minutes.

Since I've already constructed my data model's tables containing the results of the "lift", it's a pretty simple matter to leverage OAC's powerful data visualization capabilities to show those results. I've connected to my OAC's URL and selected the **RECLONED** ATP database as the connection I'll be using to gather data, chosen the **VEVO** schema, and finally pointed to the **CONVERTS_LIFTED** table for my data source.

From the **CONVERTS_LIFTED** table, I'll select just two columns – **QUANTILE_NUMBER** and **LIFT_CUMULATIVE** – and then display their contents before marking them as an attribute and a measure, respectively. I'll then save my work as a new OAC project named Voter *Flippability Analytics* before heading on to the visualization phase.


Figure 1. OAC: Graphing the Results with a Bar Graph

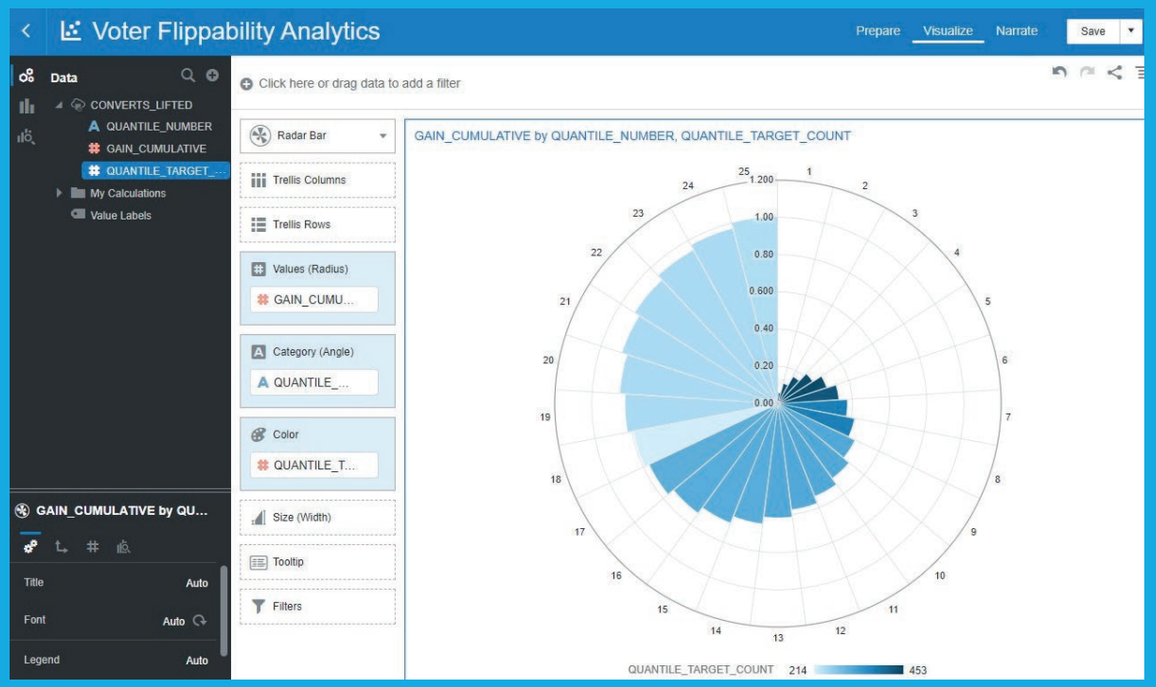


Figure 2. OAC: Graphing the Results with a Radar Bar Graph

Almost done! In my saved project, I clicked on the *Visualize* link and dropped the appropriate attribute and measure into a simple *bar graph* **(Figure 1)**. I'm using a variation of the query in **Listing 5** above as the source of this data visualization. It's immediately apparent from this bar graph that the ML model has indeed observed some significant lift among the test cases of voters identified as flippable for the campaign.

For a slightly different view of these results, here's something quite different from the traditional visualization methods to present the cumulative lift within quantile, something I can't do with Oracle Machine Learning (OML) and its Zeppelin notebook: *a Radar Bar* graph. Note that I used the values retained in the *quantile target* count column to drive the intensity of the color **(Figure 2)**.

This is just a small demonstration of OAC's power for data visualization, of course, and in the next article in this series, I'll explore its myriad features for more complex data visualization as well as its capabilities to analyze source data for completeness and accuracy before I've even applied any ML modeling.

## What's Coming Up Next?

We've come a long way since the first part of this series. This article demonstrated how to:

- Identify a specific real-world use case – the need to identify "flippable" voters for an election campaign – to focus volunteer resources more effectively
- Construct and validate a Decision Tree ML model to predict the outcome of a simple strategy to identify those "flippable" voters
- Leverage *Oracle Analytic Cloud* to report against and visualize the ML model's results

But we're not quite done yet – not in the least! In the next article in this series, I'm planning to:

- Explore how to leverage *Oracle Application Express* (APEX) to develop robust applications to assist with data management - especially when those data sources are external to a traditional Oracle database
- Leverage APEX to refine and execute ML models, as well as visualize their results
- Explore OAC to demonstrate much more of its built-in analytic engine and data cleansing capabilities.
- Demonstrate how SQL Developer's built-in *Data Mining* toolsets further aid in constructing, managing, and visualizing ML models

## References

These reference guides are invaluable to understanding more about how to leverage OML within Oracle Autonomous Database and gain an initial foothold to understanding the extensive Data Mining capabilities built into Oracle Database 19c.


- Oracle 19c Data Mining Methods "Cheat Sheet": **https://www.oracle.com/a/tech/docs/oml4sql-algorithm-cheat-sheet.pdf**
- Oracle Analytic Cloud (OAC) Visualization and Reporting Techniques: **https://docs.oracle.com/en/cloud/paas/analytics-cloud/acubi/index.html**
- SQL Developer Data Mining Features: **https://www.oracle.com/database/technologies/datawarehouse-bigdata/dataminer.html**

- Oracle 19c Data Mining User Guide: **https://docs.oracle.com/en/database/oracle/oracle-database/19/dmprg/index.html**
- Oracle 19c DBMS_DATA_MINING Package Documentation: **https://docs.oracle.com/en/database/oracle/oracle-**

## About Jim Czuprynski

Jim Czuprynski has nearly 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, IOUG COLLABORATE, ODTUG Kscope, Oracle Development Community tours, and Oracle User Group conferences around the world.

Jim has authored over 100 articles on facets of Oracle DB administration since 2003 at databasejournal.com and IOUG SELECT. His **Generally … It Depends** contains regular observations on all things Oracle and the state of the IT industry. He is currently the Senior Enterprise Data Architect for Viscosity North America.


*Mia Urman & Elizabeth Pearl*

# Top Take-Aways from Oracle OpenWorld London

Oracle OpenWorld London took place on February 12. Although much smaller than the grand affair that was Oracle OpenWorld in San Francisco, it had the big star lineup of the usual conference. Safra Katz, Oracle's CEO, was a highlight as well as other visionary keynotes that marked the agenda.

Let's start at the top, shall we: We loved the intimate format of the event. It was an all in one, a major conference hosting the top stars in the Oracle world and at the same time a more intimate and casual setting. This was assisted by the fact that at OOW London, the sessions took place within the exhibit hall. It was a nice balance of learning technology and meeting incredible vendors that had the next generation cutting edge solutions. Our technologies of choices at our company AuraPlayer are Oracle digital assistants, mobility and integration, so we were thrilled to see so many great sessions and vendors in this space.

## Oracle's Vision

We were happy to see that Oracle still has a bit up their sleeve. We were most excited that Oracle seems to be taking the lead in innovations with their digital assistants. Product management for the tool reviewed new features in the area of usability, intuitive chats design (using the conversation designer – a no-code low code way of developing your chat dialogue), to the incredible innovations in voice technology and security. This is definitely something to keep your eye out for!


*Suhas Uliyar talking about Integration*

## Oracle and Integration

The conference also revealed to us that Oracle has taken a great leap in the area of integration. Suhas Uliyar, VP of Digital Assistant and Integration, has taken over the Integration position recently and seems to be going at it with the same drive and passion as he has in the digital assistant realm.

Also, a huge plus is the fact that the integration technologies already support built-in plugins to the Oracle Cloud and Fusion based applications. Additionally, they provide a flexible framework to add your own plugins which is similar to what we have done with traditional Oracle EBS automation and Oracle RPA.

The conference layout was a great place for us to meet members of the EBS community. It helped spark conversations around system UI and look & feel, usability, and general legacy modernizations. We were happy to see that the Oracle community members are headed in the right direction for 2020. In addition to the meetups, we thought it is important to mention how beautiful the location was, take a look on the right side.

Thank you, Oracle, for a wonderful conference! We look forward to seeing everyone in London next year or Vegas in September.

## About Mia Urman

Mia is an Oracle ACE Director and an expert in Oracle development tools and middleware. Mia is a seasoned presenter on Oracle technologies and often presents at OOW, Kscope, UKOUG, DOAG and Collaborate. While providing services to customers, Mia became intimately familiar with some of the leading challenges facing Oracle Forms/ERP clients and co-founded AuraPlayer (formerly OraPlayer), a development house providing solutions in the mobile, SOA, integration, auditing, and RPA fields. She currently serves as AuraPlayer's CEO.

## About Elizabeth Pearl

Elizabeth is the Business Development Manager at AuraPlayer. With only a few years of experience in the Oracle world, she has already amassed vast experience providing customers with digital solutions to modernize their Oracle on-premises and legacy systems. Her attendance at many Oracle events has bolstered her knowledge in Oracle Forms and Oracle E-Business Suite. With this knowledge, Elizabeth increases the value of current customers while attracting new ones.

# Our User Groups in Times of Corona

***The Corona crisis has caused events all over the world to be either canceled or postponed. We asked some of the User Group representatives about how the Coronavirus has affected their activities.***

**Luiza Nowak of POUG**

We had to postpone our March event and all the activities are now frozen. Our POUG Workshop is postponed to July 3 and POUG2020 is set to September 11 to 12 (no changes at the moment, but the situation is very dynamic). We want to organize a webinar dedicated to POUG Workshop attendees in April with one of our partners, Microsoft. Stay safe!

**Francesco Tisiot of ITOUG**

ITOUG has been „lucky" on this aspect: Our main event, the ITOUG Tech Days in Milan and Rome, happened just before the virus spread in our country. In Rome, the first two cases were discovered just a couple of hours before the speakers arrived in the city and the two people affected were staying in a hotel 500 meters from where the speakers were staying. Our event wasn't impacted by that, the situation started becoming serious a few weeks later. We don't have any physical planned event until the ITOUG Tech Days in 2021 for which we don't have any date yet. From March 30 to April 1, we organized the „ACEs @ home": a series of 3 afternoons with 4 sessions each day over 3 topics: Analytics, APEX and Database.

**Ciprian Onofreiciuc of RoOUG**

In Romania, we decided to cancel several events including the conference scheduled for June 3, RoOUG Contech 2020. In of our meetups we managed to reschedule online, and we will try to do the same for the next one, in May. We hope we will come back in the community with live events in autumn. Stay safe!

**Klaus-Michael Hatzinger of AOUG**

Due to the corona crisis, the AOUG has completely cancelled this year's user conference in June. There are still three months to go, but we do not believe that all restrictions will be lifted by then. In addition, many speakers would come from other countries where it is not clear whether the border closures and flight connections will work again by then. Apart from the user conference we have no remote events planned at the moment. The customers are far too busy with themselves, so we don't believe it would make sense right now.

**Ann-Sofie Vikström Often of OUGN**

Our meetup activities stopped temporarily. The board meetings have been reduced to a minimum. We have postponed the Nordic Oracle APEX Day as a tour, going to Iceland, Denmark, Norway, Finland and Sweden, and also postponed the OUGN2020 conference to October 21 to 23 and hope that this is enough time. We need to reschedule with the speakers but for now we have no plans to start a new process for the content. We are informing our members about the different initiatives going in ACE@ Home, APEX@Home.

**Andrejs Vorobjovs of LVOUG**

We postponed our annual conference, Riga Dev Days. This means a small loss because we had agreements with the venue etc. The EOUC meeting is also postponed as we had planned to do it during Riga Dev Days. We have no plans for new events, until there will be clear information about COVID-19. Maybe we will do a small remote meetup in April.

**Irene Förg of SOUG:**

We had planned a few events in Switzerland and will now do them all remotely. The general assembly of the association has already taken place, we even did the voting via Zoom. The feedback was positive throughout. Now we are organizing two more events where we offer 2 to 3 streams with lectures in the morning.

# Call for Papers

## Oracle Groundbreakers EMEA Tour 2020

March 16 - May 3, 2020
EMEA
http://ogbemea.com/

## BGOUG

April 15, 2020
Pamporovo, Bulgaria
https://bgoug.org/en/become-a-speaker

## Oracle Open World

May 8, 2020
Las Vegas, Nevada, USA
https://www.oracle.com/code-one

## HrOUG 2020

May 31, 2020
Rovinj, Croatia
https://2020.hroug.hr/eng/Call-for-papers

# Events

## COLLABORATE 20

April 20 - 23, 2020
online
https://questoraclecommunity.org/collaborate/

## UKOUG Business Applications Exchange

June 15 - 16, 2020
The Oval, London, UK
https://ukoug.org/page/bax20

## RECONNECT 20

July 21 - 23, 2020
Hyatt Regency St. Louis At The Arch, St. Louis, Missouri, USA
https://questoraclecommunity.org/reconnect

## INFOCUS 20

August 25 – 27, 2020
Sheraton Downtown Denver, Denver, Colorado, USA
https://questoraclecommunity.org/infocus/

## UKOUG Technology Summit

September 8, 2020
Birmingham, UK
https://ukoug.org/page/techsummitdos

## POUG2020

September 11 - 12, 2020
Wroclaw, Poland
https://poug.org/

# Events

## Oracle Open World / Code One

September 21 - 24, 2020
Las Vegas, Nevada, USA
https://www.oracle.com/code-one/

## Oracle Groundbreakers EMEA Tour 2020

October, 2020
EMEA

## HrOUG 2020

October 13 - 16, 2020
Rovinj, Croatia
https://2020.hroug.hr/eng/

## OUGN “Spring” Conference (postponed)

October 21 - 23, 2020
Cruise Ship (Oslo, Norway - Kiel, Germany)
https://ougn2020.com/

## SAOUG Connect 2020

November 9 - 10, 2020
Cape Town, South Africa
http://www.saoug.co.za/

## DOAG 2020 Conference & Exhibition

November 17 - 20, 2020
Nuremberg, Germany
https://www.doag.org/en/home/

## 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):** Dr. Dietmar Neugebauer

## Editorial team:

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

Other editors: Lisa Damerow, Mylène Diacquenod, 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**

## Graphic design:

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

## 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 #22 edition: **Juni 6, 2020**.

## Photo credits:

Title: © Kowit Paikhamnam | https://123rf.com
P. 4: © floralset | https://123rf.com
P. 5: © Comic | www.commitstrip.com
P. 6: © fotogestoeber | https://stock.adobe.com
P. 6 - 9 Pfeile: © MicroOne | https://stock.adobe.com
P. 10 links - 13: © Chalermsuk Bootvises | https://123rf.com
P. 10 rechts - 13: © Andrey Ikryannikov | https://123rf.com
P. 25: © Anastasiia Nevestenko | https://123rf.com
P. 34: © Sergey Nivens | https://123rf.com
P. 39: © ??????? ?????? | https://123rf.com
P. 43: © Freepik | https://freepik.com
P. 53 - 55: © Mia Urman & Elizabeth Pearl
P. 58 - 60: DOAG