October 2020, Edition #22
ORAWORLD
e-Magazine for Oracle Users published by the EOUC

Podcasts: A Fascinating Medium

Oracle APEX Series Part 5

US CLOUD Act Series Part 2

Autonomous DB and ML Part 4

EOUC

The US CLOUD Act in light of EU Data Protection Law

page 17

OCI Block Volume: Auto-tune for Detached Volumes

page 45

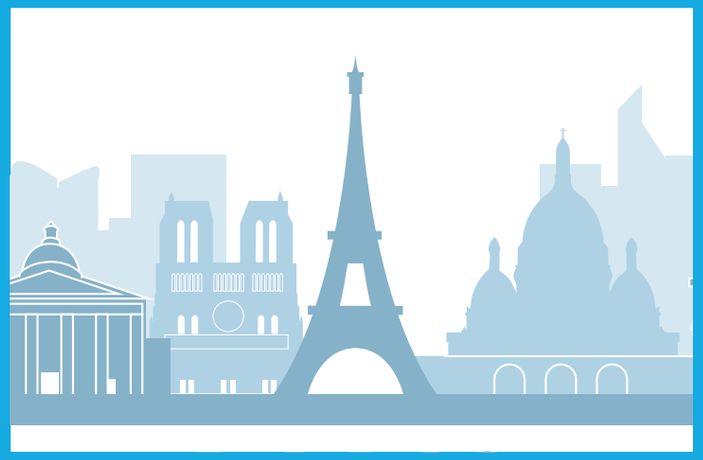

Activities of the French-speaking Oracle User Clubs

page 49

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

Editorial

Dear ORAWORLD Readers,

During the corona pandemic, an e-magazine is certainly a popular publication medium. The 22nd edition of ORAWORLD definitely wants to meet these requirements.

Emmanuel Ruez, President of AUFO (French speaking oracle user group), shows the activities with which the user groups try to get in touch with their members during these special times.

The main topic of this edition are podcasts: ORAWORLD spoke with Bob Rhubart, who had been responsible for the Oracle Groundbreakers Podcast for many years until retiring this summer.

In the second article in our CLOUD Act series, David Bomhard addresses the contradictions between the Regulations of the CLOUD Act and European data protection laws.

In his last article on Autonomous Database and Machine Learning, Jim Czuprynski shows how one can collect data from various sources with APEX and e.g. connect also to Geographic Information Systems.

Carsten Czarski also deals with APEX in his article. But he shows here not the highlights of the new releases, but rather focuses on hidden treasures in APEX.

Andre Luiz Dutra Ontalba is working on a completely new feature. Now in the Oracle Cloud Infrastructure volumes can be detached to low-cost infrastructure and attached again if necessary.

Despite the corona pandemic, I hope that we will meet us again soon face-to-face at an Oracle conference. Stay healthy!

Yours,
Dietmar Neugebauer,
Former President of DOAG

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**.

Not that Simple

www.commitstrip.com

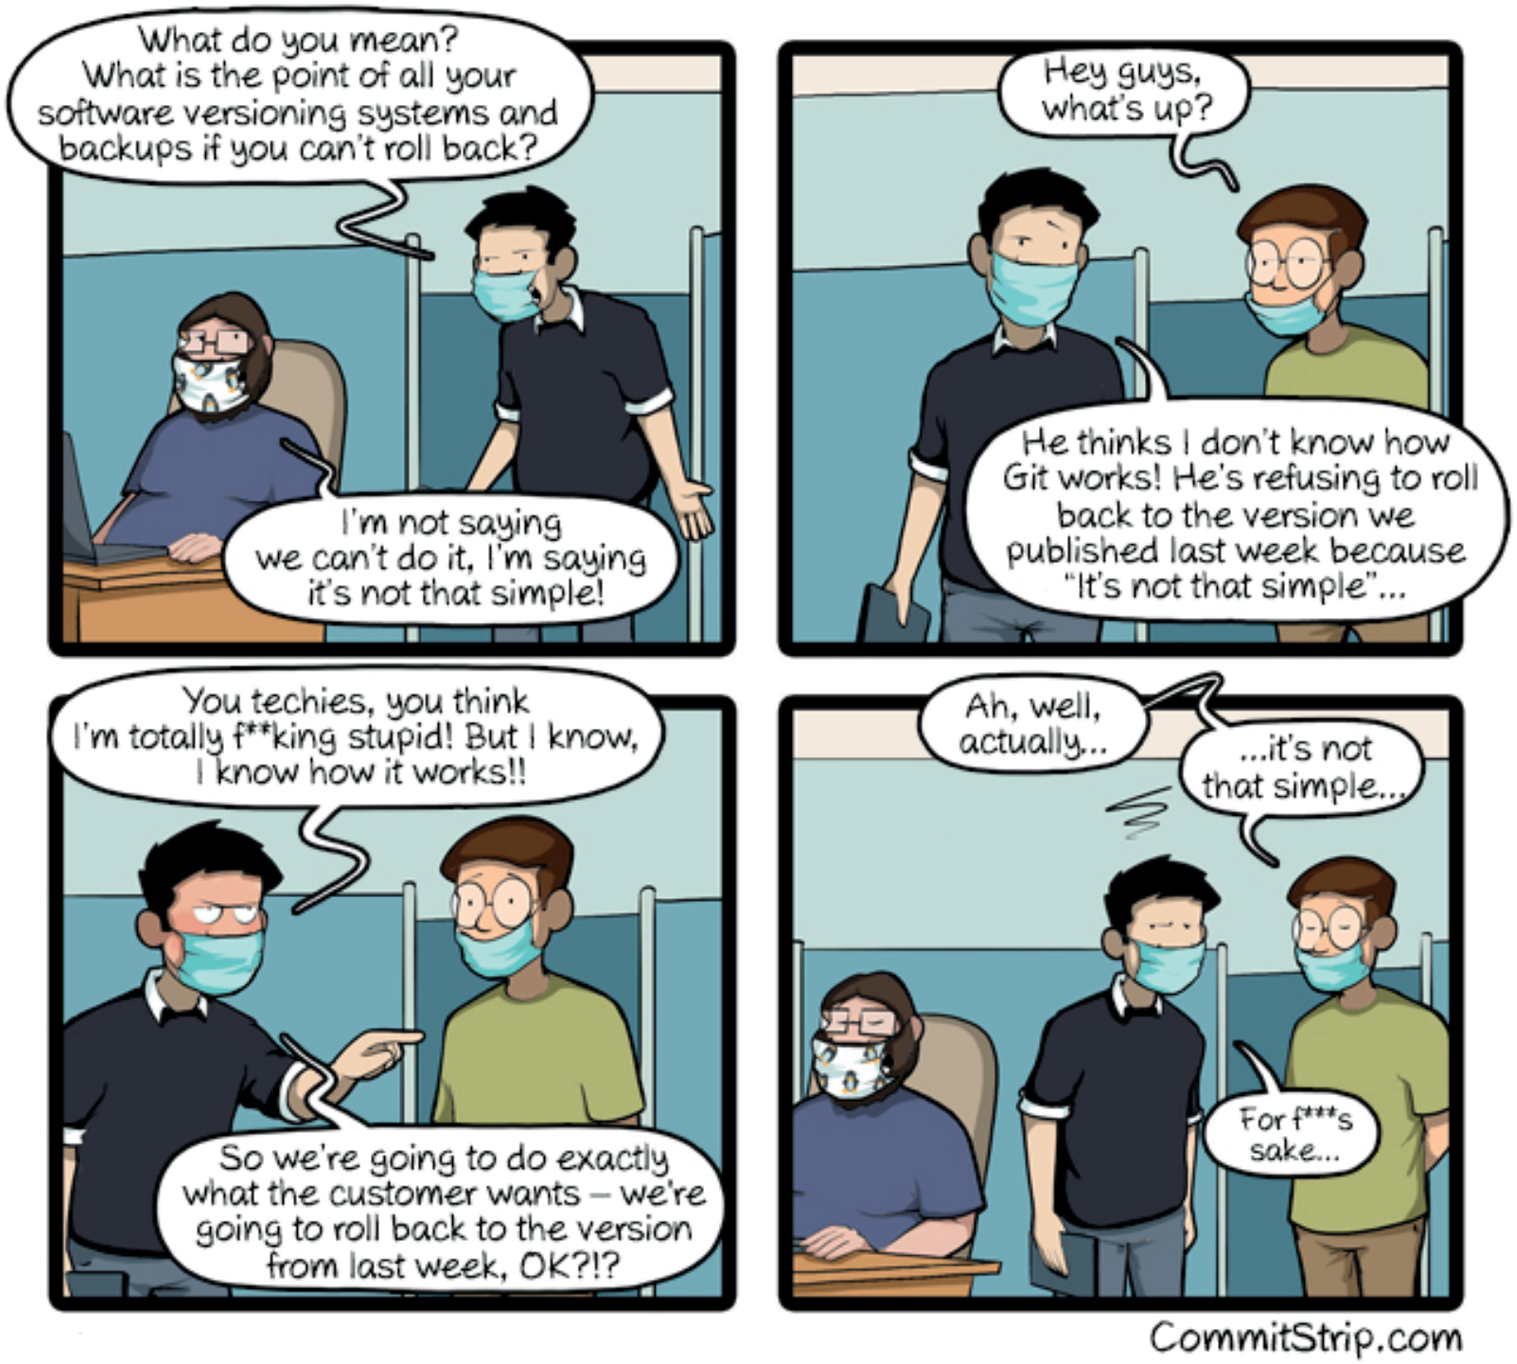

CommitStrip is a daily strip recounting funny anecdotes of life as a coder mixed up with a dash of topical tech news. Find more comics here: ***www.commitstrip.com***

Interview by Christian Luda

Bob Rhubart: “I Love the DIY Aspect of Podcasts, and that the Barriers to Entry are Almost Non-existent”

We all love good podcasts, don’t we? So, as we heard that Bob Rhubart, the man behind the Oracle Groundbreakers Podcast, was retiring at the end of July, we reached out for an interview. During his 12 years at Oracle, Bob has produced hundreds of podcasts, videos, and other content. We spoke to Bob shortly after his retirement and him handing The Groundbreakers Podcast over to community manager Javed Mohammed.

Bob, 12 years ago, you started to work at Oracle. What made you come to Oracle and how was the idea of doing a podcast born?
My career at Oracle began when Oracle acquired BEA Systems in 2008. My job at BEA started just 16 months earlier, when BEA acquired Flashline, a Cleveland-based company where I had worked since 2000. What is now known as the Oracle Groundbreakers Podcast started in March of 2009 as the Oracle Arch2Arch podcast. I had originally recorded a couple of episodes of the Oracle Technology Network TechNet podcast, at the urging of Justin Kestelyn, who ran OTN at that time. Justin then suggested that I launch a separate podcast. We titled it "Arch2Arch," which was a carry-over from the brand attached to the BEA architect community. That title soon gave way to the "OTN ArchBeat Podcast." Then, I think it was around 2017, the OTN brand was changed to "Oracle Developer Community Podcast," and about a year later that became "Oracle Groundbreakers." The podcast title was changed accordingly.

When you started, podcasts were still a relatively new phenomenon. Had you already done podcasts before?
I had done audio interviews in a previous career incarnation. I was a staff editor at a Cleveland-based online bookstore, and did several interviews with authors. This was in the late 90s, long before the term "podcast" was coined. I did a couple of audio interviews at Flashline, but those were one-offs, rather than regularly scheduled podcasts.

Looking back at the hundreds of podcasts you produced: Do you recall special episodes that were highlights for you?
Most of the podcasts I've done were recorded remotely, where I was in my home office and the panelists called in. But there were a few episodes that I recorded in my hotel room during Oracle OpenWorld or other events. I would ask the hotel to send extra chairs to my room. I'd set up a couple of microphones, and then cram five or six panelists into the room. Those programs were always fun. The panelists generally knew each other, so the conversations had a lively vibe.

How much did the production change over the years? What were some of the important things you learned in the process?
For the first couple of years guests would dial in via telephone. I had a device that allowed me to patch the phone line into an audio mixer. The mixer was then plugged into a USB interface which connected to my computer. I recorded that signal with Audacity, an open source audio recording/editing application. Eventually I switched to Skype. All the panelists would connect via Skype, and I used an application called Ecamm Recorder to record the conversation. The audio quality was generally much better than using the phone lines. I would edit those programs using Audacity as well. More recently I've been using Zoom to record the audio, and Adobe Audition for editing and mixing. Audition makes it very easy to edit and mix the various tracks that go into a podcast. Over the last few years, the Groundbreakers Podcast has gotten a bit more sophisticated, with the addition of music tracks, sound effects, and other elements. One of my podcasts typically includes six or more tracks in the final mix.

In recent years podcasts have become more and more popular. What do you think is the reason people love podcasts so much?
One factor is the sheer volume and variety of content. Another is the portability. The closest equivalent is radio, of course. But radio, with very few exceptions, has become a very cookie-cutter medium. Podcasts, on the other hand, offer limitless variety in a medium that you can consume while riding a bicycle – without the risk of decorating the side of a bus with your face.

From your perspective: What are the main ingredients of a good podcast?
Good production values are essential. It goes without saying that in an audio medium the content must *sound* good. But great production values won't save otherwise lousy content. I love the DIY aspect of podcasts, and that the barriers to entry are almost non-existent. But just because anybody can create a podcast doesn't mean everyone should.

Do you listen to other podcasts? What are your favorites, maybe also non-IT-related podcasts?
I'm a news junky, so I listen to a great many news-related programs. My favorite is "**On the Media**" produced by WNYC for National Public Radio. When I want to lighten things up I listen to "**Conan O'Brien Needs a Friend**." I've also gotten into podcast drama. One of the best I've heard is the **Wolverine series from Marvel**. For music I regularly listen to the **Rockabilly and Blues Radio Hour** podcast, and NPR's **All Songs Considered**.

Besides the Groundbreaker Podcast you were involved in a lot of other work at Oracle. What were some of your other career highlights?
Starting in late 2015, I began producing a series of videos for the **Oracle Developers YouTube channel** called **2 Minute Tech Tips**. As the name suggests, these are very short videos in which Oracle ACEs, Groundbreaker Ambassadors, and product team members present useful technical tips. In order to keep things honest, each tip video includes an on-screen countdown timer. These videos have evolved over the years so that the content is crisp and concise. I had a lot of fun developing the series and editing the tips as they were sent in by community members. For several years I also wrote a **regular column for Oracle Magazine** (which was recently replaced by **Oracle Connect**). The columns generally highlighted members of the Oracle developer community.

Now that – after 12 years at Oracle, more than 22 years in IT, and more than 48 years of full-time work in general – you have retired: What are you looking forward to most?
I enjoy the absence of schedules and deadlines. At this point I've been retired for only a few weeks, but I've already noticed that I've completely lost track of what day it is. I'm on a perpetual weekend. I'm also enjoying the tangibly pleasurable sensation associated with the realization that any sense that there is something I should be doing is nothing more than a harmless mental misfire.

Bob, thanks a lot for taking your time.

Bob Rhubart

Liron Amitzi's Top 5 Podcasts

I've been listening to podcasts for years. I started 10 years ago, when my commute to work was about 1 hour each way and I got bored of listening to the radio. In the last 5 years I've been working from home, so no commute, but I still enjoy listening when I'm out for a walk or driving.

I find podcasts to be a very efficient way to learn and expand your knowledge. The fact that it's a sound only format makes it easy to listen to them while you're doing something else such as working in the kitchen or driving. However, I can't listen when doing most work related tasks, as I really like to pay attention to the podcast itself.

When I find an interesting podcast I usually go to the beginning and binge-listen to it. This is why I don't listen to many different ones, but have usually listened to all the episodes of the ones I do listen to.

Here is a list of the top podcasts I listen to today (or have listened to recently). I hope you'll find them interesting too:

Security now

This is one of the best podcasts for geeks. It's been running every week for more than 15 years now and has tons of very valuable and interesting information about news in the computer world. As the title gives away, the topics are mainly security related (new vulnerabilities, new patches, hacks like Stuxnet etc.), but not only. Note that these are long episodes (usually around 2 hours), but they are fun to listen to. The podcast hosts are Steve Gibson and Leo Laporte (the owner of the great TWiT network) and they tend to cover the topics in depth and in a very technical way while still keeping it simple enough to understand over a podcast.

Websites: **https://twit.tv/shows/security-now** (on TWiT) and **https://www.grc.com/securitynow.htm** (on Steve's website)

Manager Tools

This is not a tech podcast, but it's very important for many tech people as most of us work with other people and manage other people. This is a great podcast with lots of content about management, communication, coaching and much more. If you are a manager, want to become a manager or even if you are only interested in these topics, this is for you. Their website has a cool "mind-map" interface to allow you to browse over 1,000 episodes by category to quickly find the ones you are interested in and expand from there. They also have another podcast called "Career Tools" for people who'd like to better manage their career. I haven't listened to it, but I'm sure it's great as well.

Website: **https://www.manager-tools.com/podcasts**

Rebel Entrepreneur

This is the newest podcast on my list with less than 30 episodes so far. Alan Donegan is the owner of “PopUp Business School”, an educational program that made it a goal to help people start their own business with low risk and no debt. Their model was to get sponsorships, fly around the world and provide 2-3 week courses for free. Then Covid-19 happened and these courses were halted. Alan then contacted the ChooseFI team (another successful podcast about financial independence) and started this podcast. Every week he talks to different people about building businesses and coaching them through the PopUp method.

Website: **https://www.popupbusinessschool.co.uk/rebel-podcast.html**

Reply all

Reply All is a podcast from Gimlet media group, hosted by PJ Vogt and Alex Goldman. The purpose of this podcast is to be “about the internet”, but it’s not from the technical point of view (even though some episodes talk about technical stuff). Each episode tells a different story, most are more of an “investigative journalism” type or about solving an internet mystery (for example: a technical issue that someone has or an interesting tweet that needs “deciphering”). Lately the episodes are published less frequently, but when they are out they are still interesting and fun to listen to.

Website: **https://gimletmedia.com/shows/reply-all**

TED TALKS DAILY

TED Talks Daily

We are all familiar with TED talks. This podcast brings some of the interesting talks in the TED network (TED and TEDx events) in podcast format. Each episode is on a completely different topic, presented by very interesting people. As always with TED, the talks are done very well, they are concise and since there are many TED events, there are lots of episodes.

Website: **https://www.ted.com/about/programs-initiatives/ted-talks/ted-talks-daily**

About Liron Amitzi

Liron is an Oracle ACE Director and a senior Oracle DBA consultant with more than 20 years of experience. During these years Liron has worked as a senior consultant with a large number of companies in various fields and has managed Oracle Professional Services teams. He mainly specializes in high availability solutions, performance, backup and recovery, and other infrastructure and application database areas. Liron is the president of BCOUG (British Columbia Oracle User Group), he is a well-known instructor and lectures in Oracle courses, events, and forums and he is the owner of the **gotodba.com website**.

Samuel Nitsche's Top 5 Podcasts

Books about programming transport a deep and well-rounded dive into a topic, online articles and blog posts often provide comprehensive examples, and video tutorials or courses show how something is done in detail. Why in the world do we need yet another information channel, namely "tech podcasts"?

Ever tried to read a book while walking your dog, watch a video while commuting or try out an example from a blog post while washing the dishes?

It might not be for everyone, but for me, podcasts fit a great role in feeding me with interesting content on my trade while I'm doing less mentally exhausting things with my hands. The following are some of my favorite or otherwise special podcasts from the tech universe:

Code Newbie

“Stories from people on their coding journey” - that’s the Code Newbie podcast. Founder and Host Saron Yitbarek interviews people from all kinds of different backgrounds and explores how they got into coding. Very encouraging and entertaining show – not only for Code Newbies.

Website: **https://www.codenewbie.org/podcast**

PawCast with GeePaw Hill

Bite-sized tips and “muses” by GeePaw Hill, who has helped Geeks to produce value for over 40 years. The episodes are usually pretty short, but dive deep into one or two aspects of a topic. The topics are concepts like “Made, Maker and Making”, “Change-Harvesting” or the lack of a “Thick Culture” in the development trade and all of them help developers to become better makers, find more pleasure in making or improve the value of the made.

Website: **https://www.geepawhill.org/tag/podcast/**

Coolest Nerds in the Room

...having the coolest conversations. Stephanie “InfoSteph” and Reggie “MadBlkMan” talk about personal and tech topics, about their challenges at work and in life, get tech-nerdy and philosophical and always spread a lot of encouragement and positive energy.

Website: **https://coolestnerdsintheroom.com/**

Programming Throwdown

The hosts Patrick Wheeler and Jason Gauci dive into different programming topics with an industry expert on the matter. That makes the show a great learning resource. Whether it is a specific programming language, a meta-topic like Design Patterns, Techniques like Functional Programming or even topics at the edges of programming like UI design:
The episodes provide a good overview and sometimes a dive into many aspects of the work of a software developer.

Website: **https://www.programmingthrowdown.com/**

#causeascene POD

#CauseAScene

Tech is not neutral – and therefore it's important that we speak about the real life impact our products and services have and their potential to cause tremendous harm. Podcast host Kim Krayton is a Business Strategist with lots of experience to coach tech leaders and invites a broad variety of incredibly interesting people to her show beyond the well-known (white and male) faces. Her interviews are refreshing, honest, challenging and I never listened to an episode where I didn't learn something new.

Website: **https://hashtagcauseascene.com/podcast/**

About Samuel Nitsche

Samuel is a curiosity-driven software developer with nearly 20 years of experience. He works as Senior Software Developer and trainer at Smart Enterprise Solutions GmbH. His main interest is on modern database development, automated testing and code quality, topics he writes regularly about on different platforms (e.g. his blog **cleandatabase.wordpress.com**, Simple-Talk and several Oracle-related print magazines). He is an Oracle ACE, one of the main contributors and maintainers of **utPLSQL** and loves to share his experience in an entertaining way – gladly in Sith robes – at meetups and conferences.

US CLOUD Act (Part 2): CLOUD Act in light of EU Data Protection Law

Dr. David Bomhard

Introduction

The general data protection regulation (hereafter "GDPR"), a milestone in European data protection law, is applicable since 25 May 2018. Two months earlier, on 23 March 2018, another legal act came into force, which also covers the handling of data: The Clarifying Lawful Overseas Use of Data Act (hereafter "CLOUD Act"). Even though a good two years have passed by now, the relationship between the two laws is still highly controversial today and has considerable practical implications.

The CLOUD Act applies to companies around the world with at least one branch or business activity in the United States (for more information on the CLOUD Act, **see Part 1 of this series**). The CLOUD Act aims to enable U.S. law enforcement agencies to access users' communications data without recourse to international mutual legal assistance agreements. This applies in particular in the event that data is stored by a U.S. company on a server within the EU. This is highly problematic from the perspective of European data protection law. While the CLOUD Act requires the disclosure of certain data to U.S. criminal prosecution authorities, such disclosure is in principle prohibited under the GDPR. This leads to a massive tension between the CLOUD Act and European data protection law. This article would like to shed more light on this dilemma and show possible ways how this dilemma could be resolved in practice.

Possible grounds for justification under GDPR

According to Art. 44 sentence 1 GDPR, any transfer of personal data to a third country shall take place only if a justification under Articles 45 to 49 GDPR is applicable and the general provisions and principles of the GDPR are also complied with. The underlying legal mechanism: data export is prohibited in principle and requires an explicit justification.

1. Compliance with a legal obligation?

According to Art. 6 para. 1 sentence 1 lit. c GDPR, the processing of personal data can be justified if this is necessary for compliance with a legal obligation to which the respective controller is subject. In this context, legal obligation means only EU obligations, i.e. the obligation must result from a norm of EU law or that of an EU member state. A request for information based on the CLOUD Act is not sufficient for this purpose. In general, the high EU data protection standard shall not be easily undermined by legal acts of third countries.

2. Privacy Shield?

Furthermore, the so-called Privacy Shield, a data protection agreement between the United States and the European Union, on the basis of which the EU Commission had issued an adequacy decision, did not help here. At best, the Privacy Shield served as a basis of justification for the transfer of data from private individuals, but not for the purposes of criminal prosecution (which is what the CLOUD Act is aimed at). In any case, the Privacy Shield was successfully challenged in the European Courts of Justice (ECJ) and as a result declared invalid on 16 July 2020 (Court of Justice of the European Union Judgment in Case C-311/18). For this reason, the Privacy Shield cannot be considered as a basis of justification for the transfer of data to U.S. criminal prosecution authorities.

3. Appropriate safeguards?
In the absence of an adequacy decision, personal data may only be transferred to a third if the controller or processor has provided appropriate safeguards.

In particular, standard data protection clauses adopted by the EU Commission can serve as appropriate safeguards. In such case, the party exporting data from the EU would have to conclude the standard data protection clause with the data importer in the third country. The aim is to ensure that the data importer actually complies with the obligations arising for him from the standard data protection clauses so that an adequate level of data protection is maintained. Also, the ECJ has recently recognized them as an effective means of establishing an adequate level of data protection. However, standard data protection clauses are not an appropriate means of ensuring that data is transferred to U.S. law enforcement authorities. The same applied to binding corporate rules (BCR), which until recently were used in connection with the Privacy Shield to establish an adequate level of data protection. Since binding corporate rules only apply within the corporate group, they could not be considered as a basis for justifying the transfer of data to U.S. criminal prosecution authorities. Nonetheless, they will no longer play a role in the future after the Privacy Shield has been declared ineffective by the ECJ (see above).

4. Mutual Legal Assistance Treaty?
The GDPR aims to prevent legal acts of third countries from undermining the EU data protection standard. According to Art. 48 GDPR, personal data may be transferred to third countries if a mutual legal assistance treaty exists with the respective EU member state. Since the CLOUD Act is a national US law, it does not meet these requirements – there is no appropriate EU-US mutual legal assistance agreement in place.

5. Solution in sight?
After the Privacy Shield was declared ineffective (see above), the U.S. government and the EU Commission have started discussion about a potential successor regulation. The current aim is to assess the prospects for an agreement that complies with the requirements of Art. 44 sentence 1 GDPR and thus stands up to a potential review by the ECJ, i.e. a regulation that may possibly resolve the tension between the Cloud Act and the GDPR.

Interim Conclusion

From the perspective of the GDPR, the transfer of data to US law enforcement agencies cannot be easily justified. This can lead to a difficult dilemma in which the company has to make a decision: If it does not comply with the U.S. criminal prosecution authorities' request for information, it risks the consequences of US law, in particular, timely increasing fines. If it complies with the request for information, it risks violating European data protection law by transferring data from the EU to the United States. In case of data protection breaches the GDPR provides for administrative fines of up to 20,000,000 EUR, or 4% of the total worldwide annual turnover, whichever is higher.

Practical solutions

In order to solve the conflict between the GDPR and the CLOUD Act, the following options are conceivable in practice:

1. Spin-off of European subsidiaries

A feasible way to resolve the dilemma could in practice be to split off a European entity or to set up a European subsidiary, thus avoiding the applicability of the CLOUD Act. In case the new spin-off is no longer an addressee of the CLOUD Act, it would have better prospects to refuse the disclosure of personal data to the U.S. criminal prosecution authorities.

2. Anonymization

A further considerable option may be to anonymize the respective data to be disclosed, insofar as this is compatible with the requirements of the CLOUD Act. Anonymization leads to the fact that the data no longer fall within the applicability of the GDPR. However, the mere encryption of data (since this is reversible, unlike anonymization) does not in principle remove the personal reference and thus the applicability of the GDPR.

3. Use of data trustees

The use of so-called data trustees could also be conceivable. This is the case if a service provider stores the user data inside the EU and has access to it. US law requires the provider to own and control the data in order to be able to access the data directly, which might not be the case with a service provider as a data trustee. However, most cloud trustee providers have a representative office in the United States, which is why US law applies to them anyway and they would have to hand over the data if a U.S. search warrant were available.

4. Appeal under the CLOUD Act
Appeals to the US law enforcement authorities requesting information may also be considered (for the course of the appeal proceedings, **see Part 1 of this series**). However, for practical reasons, it is difficult for non-U.S. citizens to take legal action against measures adopted on the basis of the CLOUD Act. This is only possible if the country of the entity concerned has concluded an agreement on data exchange with the United States. There is no such agreement between EU member states and the United States. In such a case, the company would be dependent on its IT provider taking legal action against U.S. authorities. But even for a U.S.-located IT provider such a possibility is only available within narrow limits.

5. Risk allocation towards IT provider
Finally, a contractual allocation of risk to the respective IT provider should be considered, flanked by corresponding information, indemnity and recourse obligations.

Summary
Many Companies may be exposed to the above-mentioned dilemma between the CLOUD Act and the GDPR. At present, personal data can only be transferred to the United States in conformity with data protection laws within very narrow limits. The European Data Protection Board (EDPB) recently emphasized the need for an international data protection agreement with regard to the export of data to the United States. It should become clear by then which of the outlined solutions will prevail in practice as “legal work around”.

About Dr. David Bomhard

Dr. David Bomhard is a physicist and IT lawyer at Noerr LLP in Munich specialising in legal advice on complex IT projects, especially IT outsourcing, cloud computing, agile software development, automation of corporate processes, use of artificial intelligence, data usage contracts, software as a service (SaaS), cyber security, IT transactions and software copyright. Dr. David Bomhard is also a lecturer for negotiation management and legal informatics at the University of Passau.

NUMBER OF THE MONTH:

88 Million *Followers*

The much talked about deal between Oracle and TikTok is still not fully approved at the time of our editorial deadline. While we continue to follow the news surrounding this deal, let's have a look at the biggest star of the popular video-sharing app: Charli D'Amelio.

The 16-year-old social media star was a competitive dancer for over 10 years before she started posting content on TikTok, in late 2019. Her videos that show her dancing to trending songs quickly earned her a large following, making her the most popular creator on TikTok with an impressing number of over 88 million followers.

This success quickly led to other ventures outside of social media which include a nail polish collection, a makeup line, and a sweatshirt collection, as well as a voice role in the 2020 animated movie "StarDog and TurboCat". As a result, D'Amelio has been placed on Fortune's "40 Under 40" list in 2020.

But what really is TikTok? While being a guest on Jimmy Fallon's Tonight Show, Charli D'Amelio herself breaks down how the app works. The TikTok star also talks about becoming one of the biggest influencers on the platform and using her fame to raise awareness about cyberbullying.

What's Your Super-Power? Mine is Autonomous Database and Machine Learning (Part 4)

Jim Czuprynski

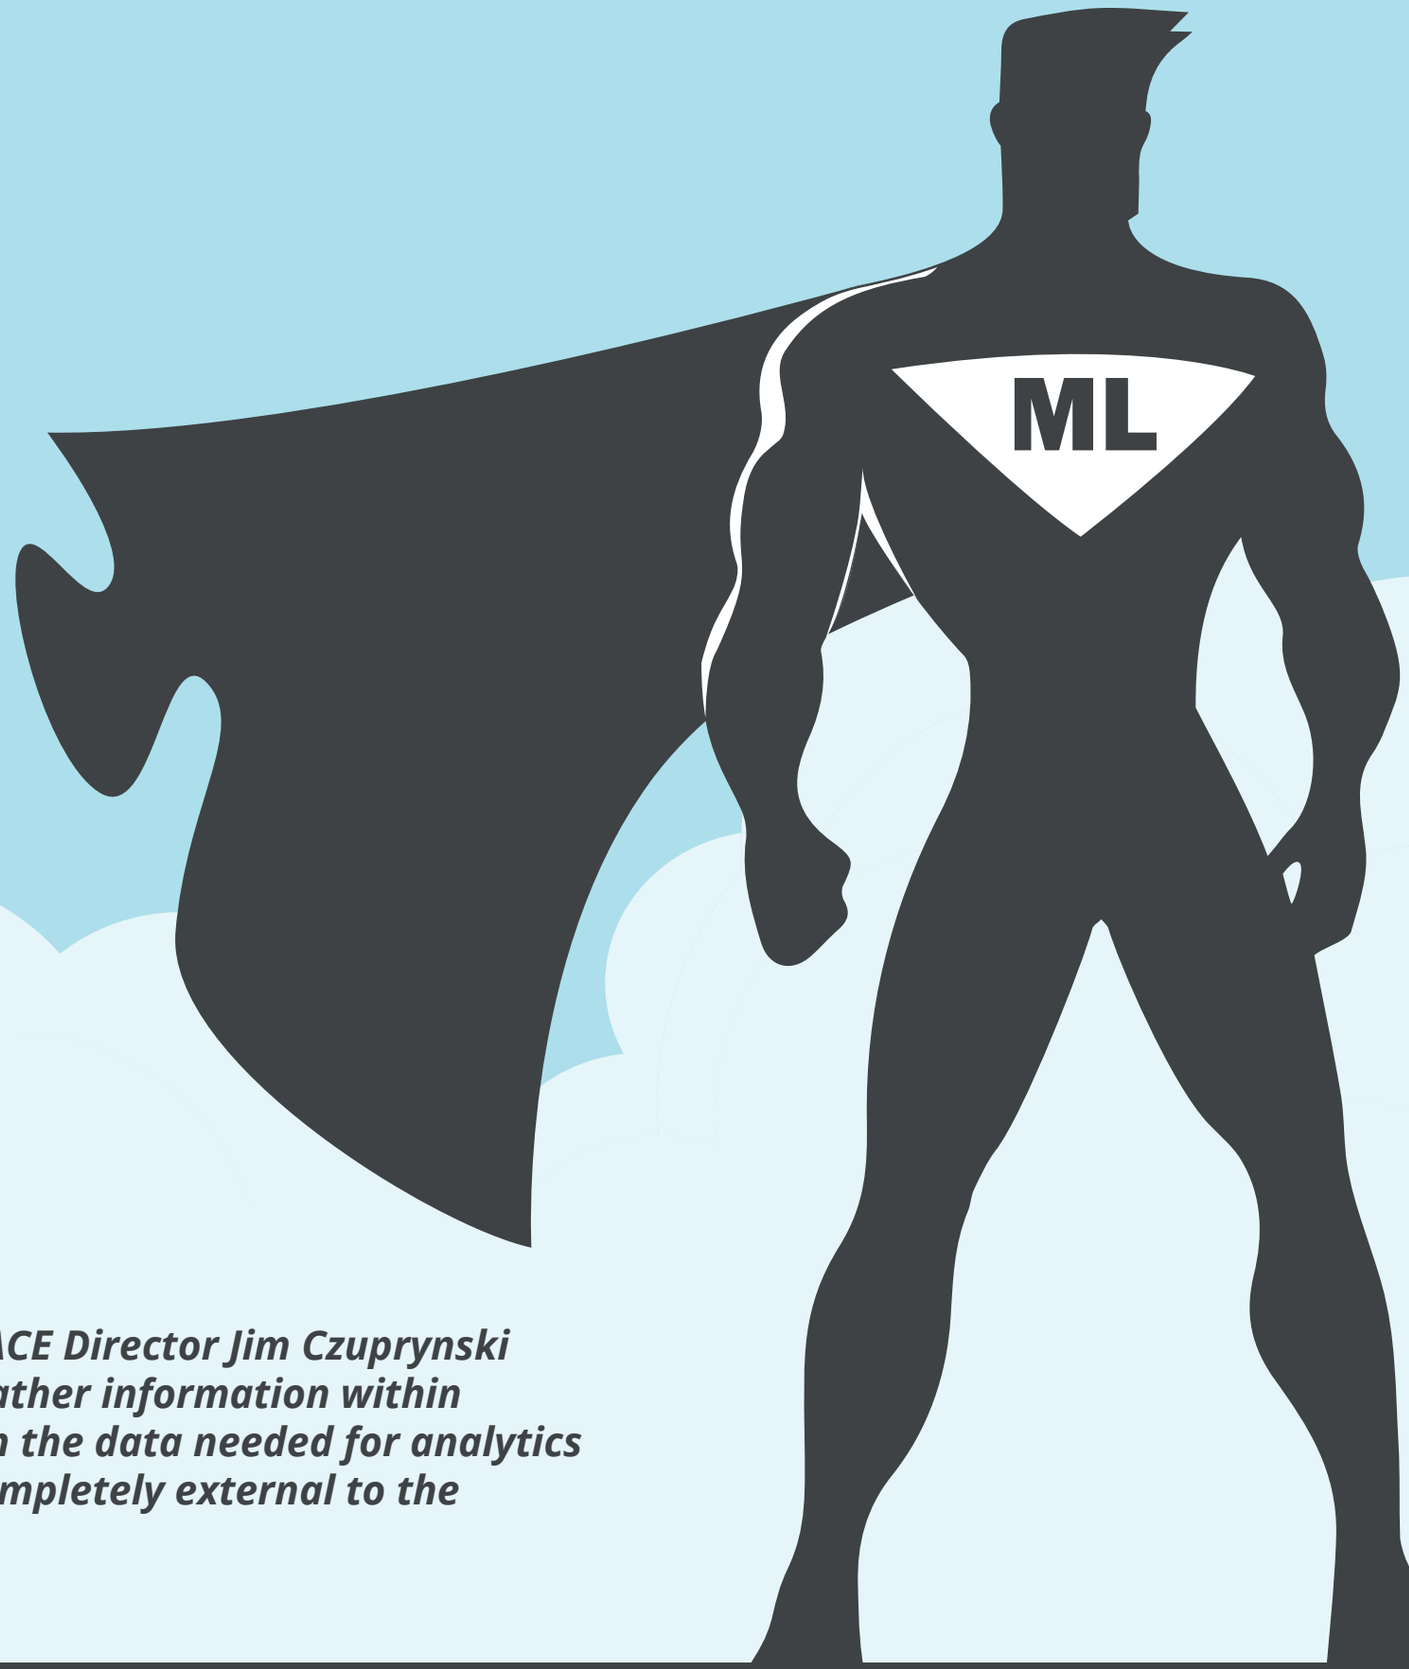

Building on the third article in this series, Oracle ACE Director Jim Czuprynski wraps up this topic with a look at other ways to gather information within Oracle Application Express (APEX), especially when the data needed for analytics and fulfillment of our busines units' requests is completely external to the application's Autonomous Database environment.

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

The **previous article** in this series took a close look at how to leverage Oracle's machine learning capabilities as I:

- Adjusted my initial Decision Tree model for deeper insights into "flippable" voter patterns
- Reviewed methods to determine the accuracy of the model's results via Confusion Matrix and ROC reporting
- Used *Application Express* (APEX) to build new variations of the model and display those variations through APEX's visualization toolsets

This final article in the series demonstrates how to leverage *Oracle Application Express* (APEX) to capture data from diverse sources, even when they're external to a traditional Oracle database, including web sites that offer APIs to capture data. And then I'll close with some real-world examples of how to capture and use *geographic information system* (GIS) information within APEX to assist in identifying voter locations and even help build efficient routing methods using APEX *plug-ins*.

Shifting Priorities, New Imperatives

As I described in the prior article in this series, I'd spent quite a bit of time preparing analyses of "flippable" voters – folks that might be persuaded to consider voting for our candidate based on their prior party declaration history during primary elections. However, as is not uncommon within campaigns, I soon found myself faced with some new imperatives from our volunteer leadership specific to procuring and using geolocation data to assist with delivery of campaign merchandise to interested voters, as well as distributing small and large campaign signs throughout our far-flung congressional district.

Again With the Excel.

During the last election cycle, we'd been hard-pressed to tackle these types of requests within the last few weeks of the campaign. Our efforts had included numerous exchanges of requestors' data, often via Excel spreadsheets and Google sheets, and my colleagues also decided to use a nearly-free solution to obtain geo-coded addresses. This time around, however, my colleagues realized that we really needed a method that would allow us to gather *multiple* sets of latitude and longitude attributes in batches instead of singular requests.

I knew it would have taken only a matter of hours to construct an APEX application that would have handled capture and processing of multiple merchandise and sign requests, but our team had decided to stay with what our volunteer team was already most comfortable with: exchanging these data via Excel spreadsheets. Thankfully, they had decided to adopt a standard format of a multi-sheet Excel workbook, one for each subset of geographic *regions* our volunteers would be delivering signs and merchandise within.

The good news here is that there's a simple way to handle capturing these data into a more secure environment for processing; even better, there's a great example already included within the sample applications that are part of APEX 19.2. The *Sample Data Loading* sample application provides several excellent examples of how to capture, validate, process, and eventually load data directly from various external data sources.

Figure 1 shows the end result of adapting *Page 31* from the sample application by copying that page directly into my **VEVO2** application as a new page. After a few simple modifications to match my data loading requirements – for example, limiting the list of input files to only those with an XLSX extension – the page was ready for immediate use.

A really neat feature of this page is that once the appropriate XLSX Excel workbook is selected, its corresponding worksheets are automatically displayed as part of a drop-down list populated automatically as a List of Values (LOV), as shown in ***Figure 2***. Selecting the desired worksheet and clicking the **Preview** button populates a preview report of the first ten rows available for loading, so that the end user can visually validate they've chosen the right data for import and processing (I've obscured the requestor's address information to protect their PII; the names shown are actually random-generated for obscurity as well).

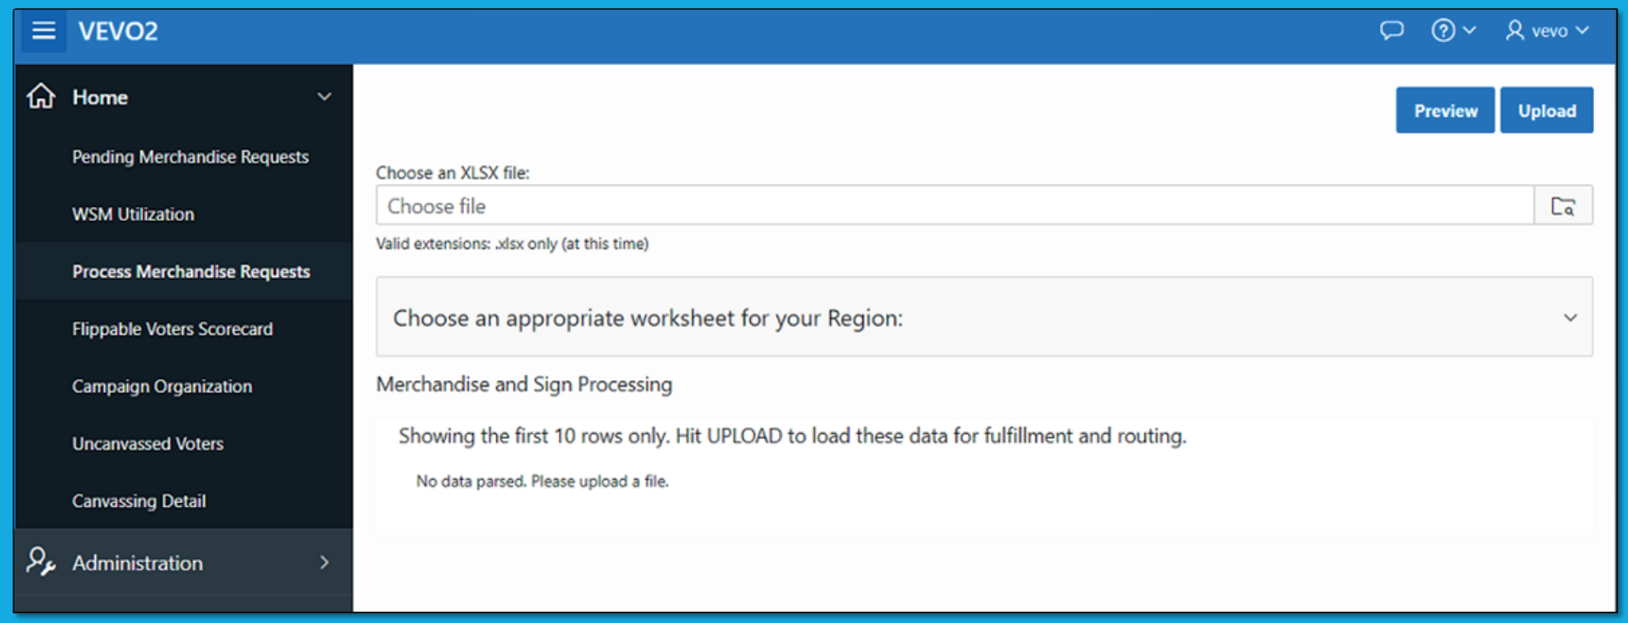

Figure 1. Specifying an Excel Workbook for Ingestion

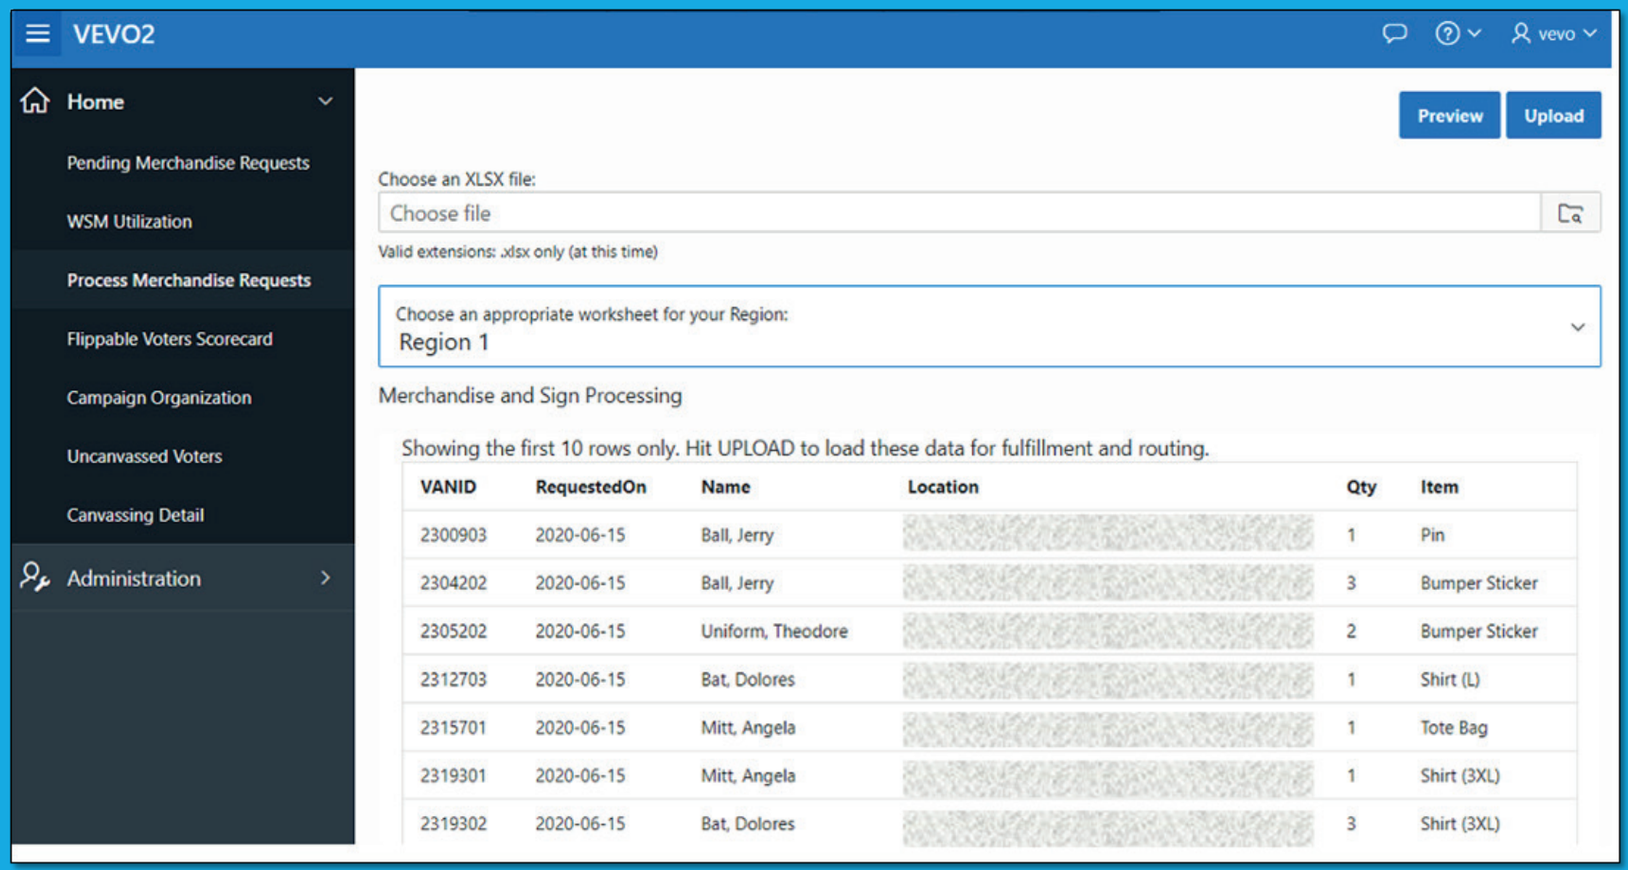

Figure 2. Choosing a Single Worksheet and Previewing Its Contents

After reviewing the chosen data, the end user simply clicks the **Upload** button to begin processing the merchandise and campaign sign requests. The data will be loaded into a new table named **VEVO**. **T_VOTER_MERCH_REQUESTS**; *Listing 1* shows the DDL for that table.

```
-----
-- Build table containing limited geographic data and SDO_GEOMETRY
object
-----
DROP TABLE vevo.t_voter_merch_requests PURGE;
CREATE TABLE vevo.t_voter_merch_requests (
    vm_id             NUMBER        NOT NULL
   ,vm_rqst_dt        DATE          NOT NULL
   ,vm_rqstr_name     VARCHAR2(100) NOT NULL
   ,vm_rqstr_addr     VARCHAR2(100) NOT NULL
   ,vm_qty            NUMBER(3,0)   NOT NULL
   ,vm_type           VARCHAR2(20)  NOT NULL
   ,vm_dlvr_id        NUMBER
   ,vm_dlvr_dt        DATE
   ,vm_lat            NUMBER
   ,vm_lng            NUMBER
   ,vm_geopoint       SDO_GEOMETRY
);

-----
-- Create PK IDX
-----
 ALTER TABLE vevo.t_voter_merch_requests
    ADD CONSTRAINT voter_merch_requests_pk
    PRIMARY KEY (vm_id, vm_rqst_dt)
    USING INDEX (
       CREATE UNIQUE INDEX vevo.voter_merch_requests_pk_idx
           ON vevo.t_voter_merch_requests(vm_id, vm_rqst_dt)
            TABLESPACE vevo_idx
        );

-----
-- Create a corresponding VIEW for the table
-----
CREATE OR REPLACE VIEW vevo.voter_merch_requests AS
  SELECT * FROM vevo.t_voter_merch_requests;
```

Listing 1. Creating Table VEVO. T_VOTER_MERCH_REQUESTS

Note: I can already hear the sound from a thousand face-palms about the potential security risks of gathering data via Excel worksheets, Google sheets, CSV files, or indeed any other externally-sourced data. I'm keenly aware that it's possible to imbed malware into *any* external source of data, especially if the end users passing files back and forth haven't completely secured their at-home networked devices, implemented a robust antivirus protection scheme, and followed common-sense security best practices to avoid incursions by bad actors. Suffice it to say it is a continuous battle I'm always waging on many fronts, and a great topic for a future article series.

Leveraging Oracle Spatial and Graph Features

My new table is quite standard, but with one significant difference: The **VM_GEOPOINT** column's datatype of **SDO_GEOMETRY** enables me to take advantage of Oracle Database's Spatial and Graph features, which are available **at no additional licensing cost since December 2019**. I'll be exploring Spatial and Graph features in much greater depth in an upcoming series of videos and webinars, but for now let's look at how to enable this table for those features.

Listing 2 shows how I activated the **T_VOTER_MERCH_REQUESTS** table for Spatial and Graph capabilities by adding an entry for it into the **USER_SDO_GEOM_METADATA** table. This new entry tells the database that the **VM_GEOPOINT** column *contains geolocation data*, the *permissible ranges* for Latitude and Longitude values, and the *spatial reference system* we will be using (**8307**, which specifies the WGS 84 latitude and longitude system of coordinates and is most typically used for navigation via GPS devices).

```
-----
-- Update SDO Geometrics metadata to reflect Longitude and Lati-
tude is being
-- applied within the table, and which column contains the SDO_GE-
OMETRY obj
-----
DELETE FROM user_sdo_geom_metadata
 WHERE table_name = 'T_VOTER_MERCH_REQUESTS'
   AND column_name = 'VM_GEOPOINT';

INSERT INTO user_sdo_geom_metadata
VALUES (
    'T_VOTER_MERCH_REQUESTS'
   ,'VM_GEOPOINT'
   ,SDO_DIM_ARRAY(
       SDO_DIM_ELEMENT('Longitude', -180, 180, 0.5)
      ,SDO_DIM_ELEMENT('Latitude',  -90,   90, 0.5)
   )
  ,8307);

COMMIT;
```

Listing 2. Enabling Table VEVO. T_VOTER_MERCH_REQUESTS for Spatial and Graph usage

Note: The broad scope of geolocation capabilities precludes an in-depth discussion within this article series, but my recent **2-Minute Tech Tip** gives you a brief introduction of how easy it is to use them with just a few lines of code. Also, here's an **excellent introduction** to the **SDO_GEOMETRY** datatype and general spatial concepts to begin your journey into Oracle Spatial and Graph's rich feature set.

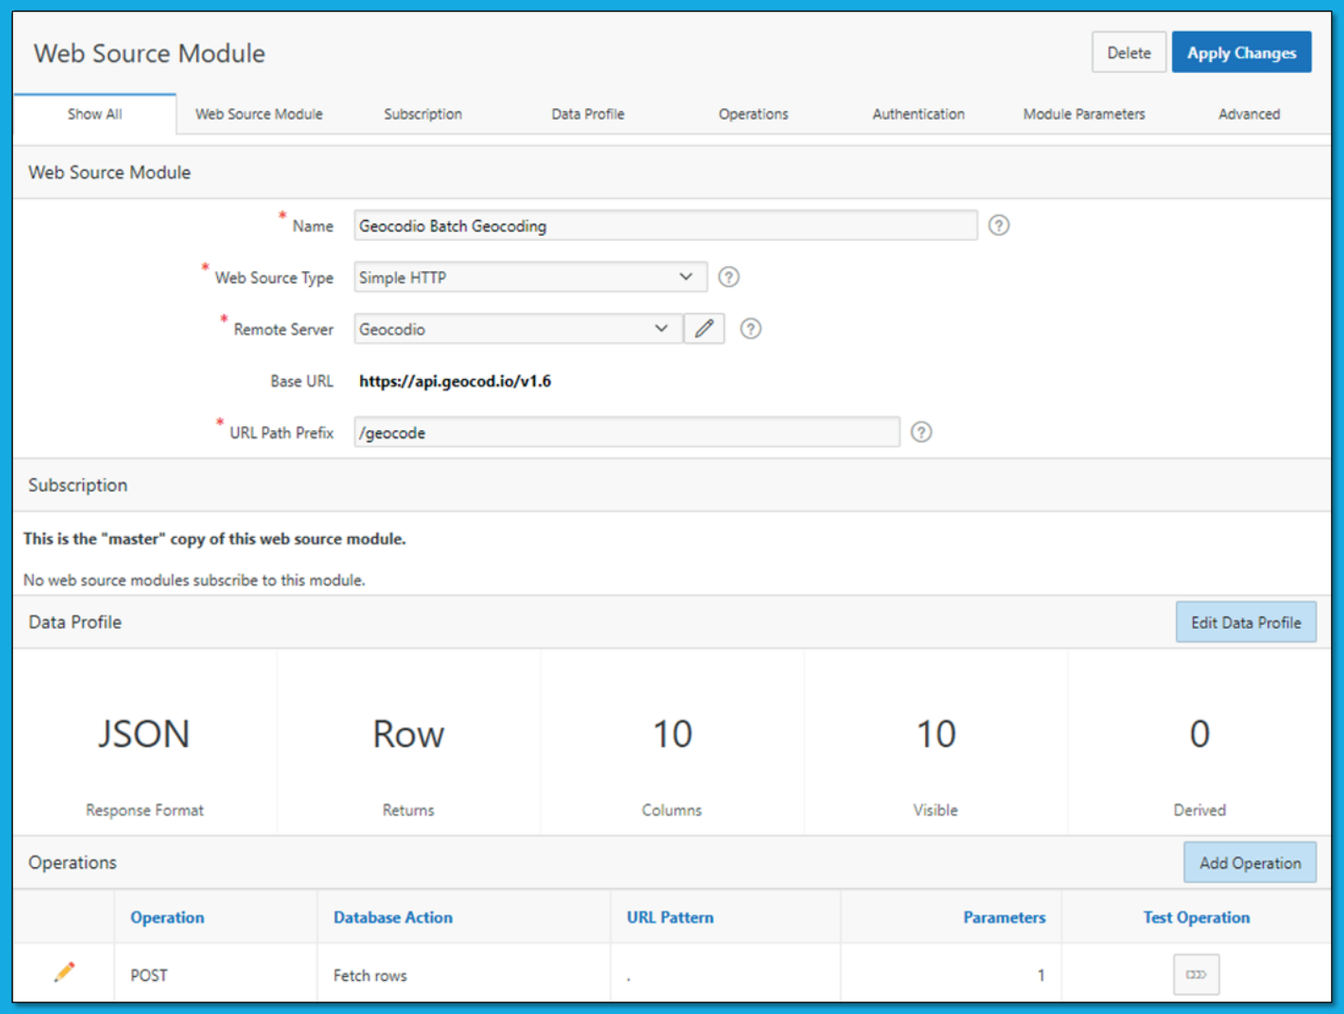

Figure 3. Web Source Module: REST API Sources and Actions

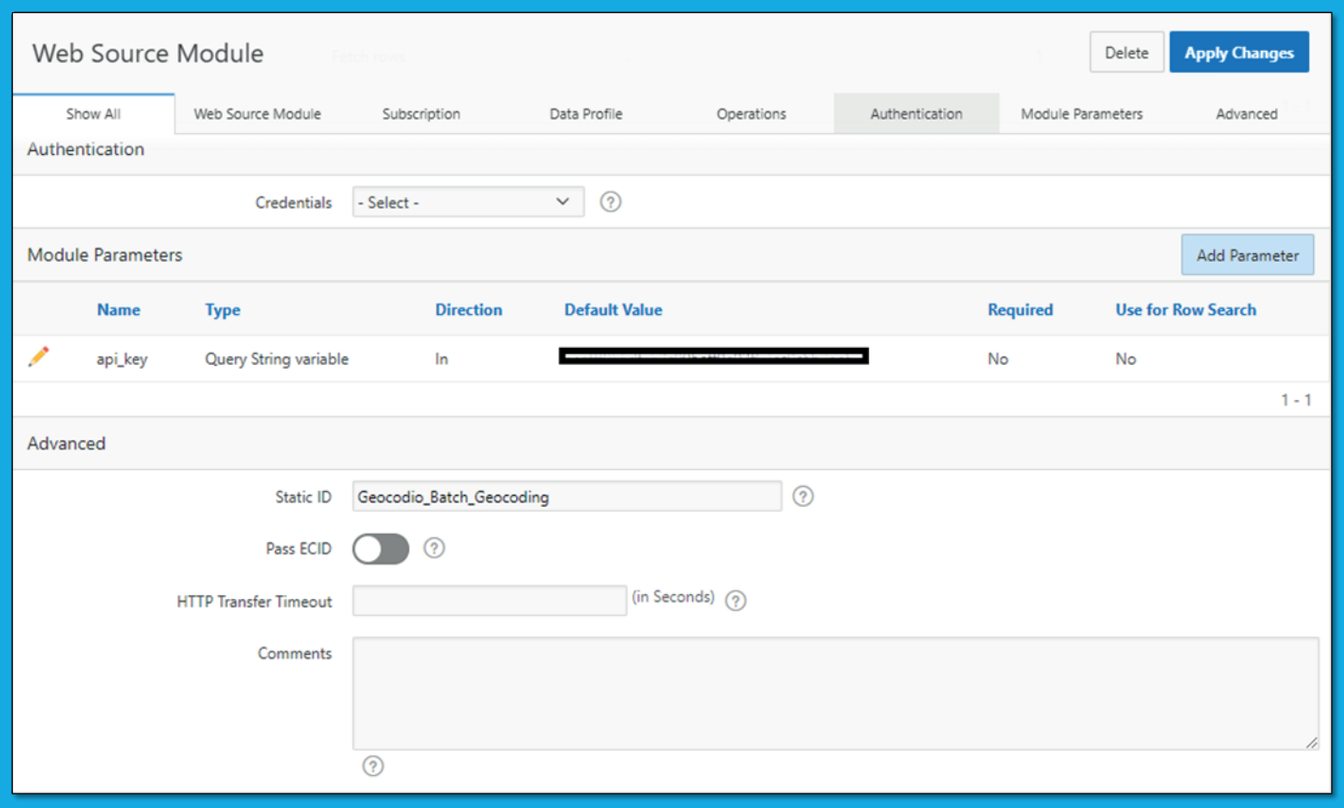

Figure 4. Web Source Module: Input Parameters List

Delivering on (Campaign) Promises

Now that I've loaded the requests for merchandise and signs into my new table, I'll next need to plot the precise geolocation of each address so that campaign organizers can figure out exactly where volunteers will deliver the items. Of course, exactly which volunteer we should assign to each set of deliveries is contingent on how close that volunteer is to the requestor. And we also need to insure we're not overburdening volunteers in our district's rural area with excessively long routes beyond the boundaries each volunteer has specified they are willing to travel.

Batch Population of Latitude and Longitude Values

As I mentioned earlier, during the last campaign cycle my colleagues had decided to handle all their geolocation needs on an address-by-address basis via calls via the free Google Maps API layer. While that technique would certainly work fine for small numbers of geolocation requests, Google assesses $5 per 1000 address requests. While each Google developer account comes with $200 in free credits per month, it's not unlikely that **credit limit** could be exceeded depending on which mapping, routing, and geolocation features are used and at what intensity.

Since I anticipated there might be times near the end of the campaign's efforts when we'd likely exceed that limit while building potential routes for our volunteers, I realized I'd need to find an alternative. I soon discovered **another address geocoding service**, *Geocod.io*, that permitted a maximum of 2,500 requests per API key per day for free. Geocod.io provided a relatively robust REST API for its geolocation services, and that meant I could take advantage of APEX's ability to directly access external data via calls through a *Web Source Module* (WSM).

I quickly built a corresponding WSM for the Geocod.io REST API calls I'd need to make, as shown in ***Figure 3***. Note that the API call requires just an API key for an input parameter ***(Figure 4)***.

There are just few wrinkles to using this WSM effectively, however:

- The WSM actually requires a **POST** REST API call, and I need to *pass in a specially-formatted JSON document* of all the addresses for which I needed to retrieve geolocation information.
- I also need to *parse the returned JSON output* upon receipt of the newly-geocoded addresses and then apply the latitude and longitude to the corresponding merchandise request entries in the **T_VOTER_MERCH_REQUESTS** table.
- Finally, I need to *update those geocoded attributes* in the **VM_GEOPOINT** object so I can use that information effectively for future routing calculations.

APEX_JSON and APEX_WEB_SERVICE Procedures To the Rescue

The code in ***Listing 3*** shows one solution to these three requirements. Let's break down at a high level what this code is actually accomplishing:

- First, I declare a cursor that extracts just 40 addresses at a time for which no latitude and longitude is yet populated. The limit of 40 addresses also fits nicely within the maximum 32KB size of the input CLOB; as I'll share in a bit, I'll also use that limit to insure that I never exceed the 2,500 address limit in a single day.
- I then call **APEX_JSON** to create a new JSON document that contains the unique identifier for each requestor's address and the address itself. When all of the pending requests are processed, I close the JSON document and pass it to the **SEND_CLOB** variable for processing.
- Next, I call **APEX_WEB_SERVICE** to first prepare the WSM for the incoming data's format, and then pass the contents of the **SEND_CLOB** variable as payload to the Geocod.io REST API via a POST call.
- The results of the POST REST API call are captured in the **RECV_CLOB** variable.
- I then call **APEX_JSON** to parse the returned geocoded addresses into their component key value pairs, locating the unique request ID, latitude, and longitude values and applying them to their corresponding columns via an **UPDATE** statement.
- Finally, once all individual merchandise request records are updated, I apply the newly-derived latitude and longitude values *en masse* to the **VM_GEOPOINT** object as well, and commit all changes.

```
-----
-- Leverage APEX_JSON and APEX_WEB_SERVICE to call a predefined APEX
-- Web Source Module (WSM) via POSTing JSON output to the GeoCodio API
-----

DECLARE
  -- Processing variables:
  SQLERRNUM     INTEGER := 0;
  SQLERRMSG     VARCHAR2(255);

  -- CLOBs for input and output:
  sent_clob     CLOB;
  recv_clob     CLOB;
  -- JSON parsing variables:
  recv_values   APEX_JSON.T_VALUES;
  mbr_count     PLS_INTEGER;
  vid           VARCHAR2(4000);
  lat           NUMBER(9,6);
  lng           NUMBER(9,6);

  -- Process each set of next 40 voters at one time
  CURSOR curNeedLatLng IS
    SELECT
         TO_CHAR(vm_id) AS van_id
        ,vm_rqstr_addr  AS formatted_address
      FROM vevo.t_voter_merch_requests
     WHERE vm_lat IS NULL
       AND vm_lng IS NULL
       AND rownum <= 40;

BEGIN
  -----
  -- Create a CLOB containing the necessary address elements
  -----
  APEX_JSON.INITIALIZE_CLOB_OUTPUT;
  APEX_JSON.OPEN_OBJECT;

  FOR i IN curNeedLatLng
    LOOP
      APEX_JSON.WRITE(i.van_id, i.formatted_address);
    END LOOP;

  APEX_JSON.CLOSE_OBJECT;

  sent_clob := APEX_JSON.GET_CLOB_OUTPUT;
  APEX_JSON.FREE_OUTPUT;

  -----
  -- Set header values so that incoming input is recognized as JSON
  -----
  APEX_WEB_SERVICE.G_REQUEST_HEADERS.DELETE();
  APEX_WEB_SERVICE.G_REQUEST_HEADERS(1).name := 'Content-Type';
  APEX_WEB_SERVICE.G_REQUEST_HEADERS(1).value := 'application/json';

  recv_clob :=
    APEX_WEB_SERVICE.MAKE_REST_REQUEST(
      p_url =>          'https://api.geocod.io/v1.6/geocode?api_key={YourAPIKeyHere}'
     ,p_http_method =>  'POST'
     ,p_body =>         sent_clob
     );

  -----
  -- Populate the received JSON output into a CLOB, and then process the results
  -----
  APEX_JSON.PARSE(
     p_values => recv_values
    ,p_source => recv_clob
    );

  mbr_count := APEX_JSON.GET_COUNT(p_path => 'results', p_values => recv_values);

  FOR i IN 1 .. mbr_count
    LOOP
      vid := (APEX_JSON.GET_MEMBERS(p_path => 'results', p_values => recv_values)(i));
      lat := APEX_JSON.GET_VARCHAR2(p_path => 'results.'|| vid || '.response.results[1].location.lat', p_values => recv_values);
      lng := APEX_JSON.GET_VARCHAR2(p_path => 'results.'|| vid || '.response.results[1].location.lng', p_values => recv_values);

      UPDATE vevo.t_voter_merch_requests
        SET
          vm_lat = lat
         ,vm_lng = lng
       WHERE vm_id = TO_NUMBER(vid);

    END LOOP;

  COMMIT;

  -----
  -- Finally, apply all updated Lat/Long data in SDO_GEOMETRY type
  -----
  UPDATE vevo.t_voter_merch_requests
     SET vm_geopoint =
       SDO_GEOMETRY(
           2001
          ,8307
          ,SDO_POINT_TYPE(vm_lng, vm_lat, NULL)
          ,NULL
          ,NULL
       );

  COMMIT;

EXCEPTION
  WHEN OTHERS THEN
    SQLERRNUM := SQLCODE;
    SQLERRMSG := SQLERRM;
    DBMS_OUTPUT.PUT_LINE('Unexpected error: ' || SQLERRNUM || ' - ' || SQLERRMSG);

END;
/
```

Listing 3. Requesting, Capturing, and Applying GeoLocation Attributes

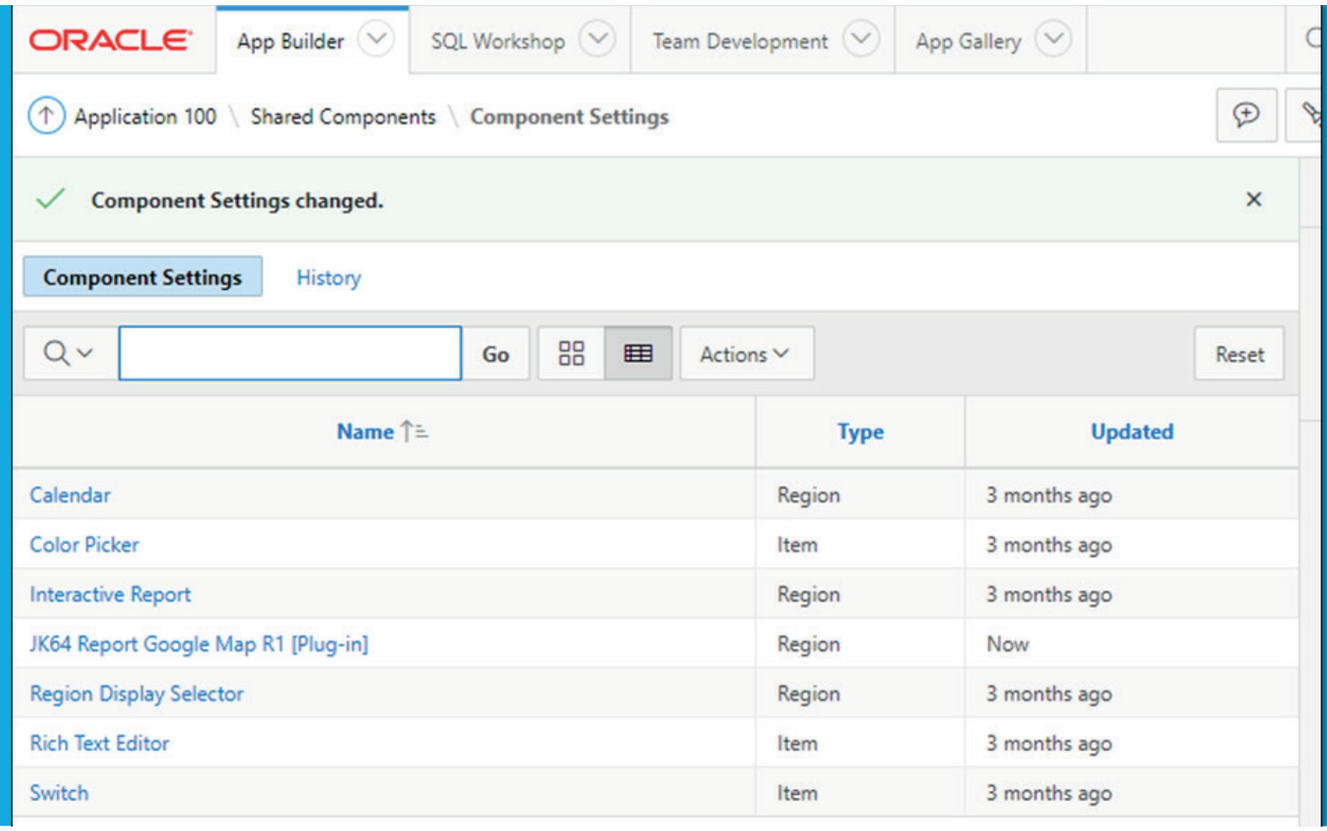

Figure 5. List of Plug-Ins After Installing Google Map Plug-In

I then scheduled this anonymous PL/SQL block to run as a scheduled task once every minute from 1:00 AM – 2:00 AM every morning within the **VEVO** schema of my Oracle database. Since each execution will only process 40 addresses at a time, that means I'd complete geocoding for up to 2,400 addresses per night, thus forestalling any possibility of exceeding the generous 2,500 per day limit that Geocod.io provides at no charge.

So … Where Is Everybody?

Now that established mechanisms to capture geolocation attributes for the campaign's merchandise requests, my next step is to figure out how to effectively dispense campaign volunteers to distribute those items. Since campaign signs are larger than coffee mugs, we may need to cache supplies to facilitate faster dispersal and display of those signs to increase potential support for our candidate. To do that, I need to display these requests via some sort of mapping software, and I also need to provide the ability to provide routing for volunteers to deliver the requested items.

Visualizing Geolocations With Google Maps: A Most Excellent Plug-In

Fortunately, my fellow Oracle ACE **Jeffrey Kemp** has provided an elegant solution: an excellent *Google Maps Plug-In* for Oracle APEX solutions. I downloaded his Plug-In from his **GitHub repository** and followed the detailed instructions to install it as part of the Shared Components for my **VEVO** application. Once installed ***(Figure 5)***, I was able to take advantage of its myriad features to display specific locations within a setting that's familiar to anyone who has ever used Google Maps on their mobile phone or within their internet browser.

This APEX plug-in makes short work of displaying specific geolocations through a simple query, as shown in ***Figure 6***. The latitude / longitude pair identifies the exact location of each request, but I also included information like the *unique ID* of each requesting voter and their *physical address* so an end user could display this information during a mouse-over of the location. Also note that I differentiated between Large and Small campaign signs versus "normal" merchandise by specifying different icons for each item type.

Note: It's important to follow the precise **instructions** as to how columns must be specified in the query you're building to display points within the Google Map plug-in. The extra 5-10 minutes spent reading Jeff's excellent documentation are definitely worth it. (I wasted almost two hours of frustrating experimentation before I realized my query's issues were simply a failure to RTFM.)

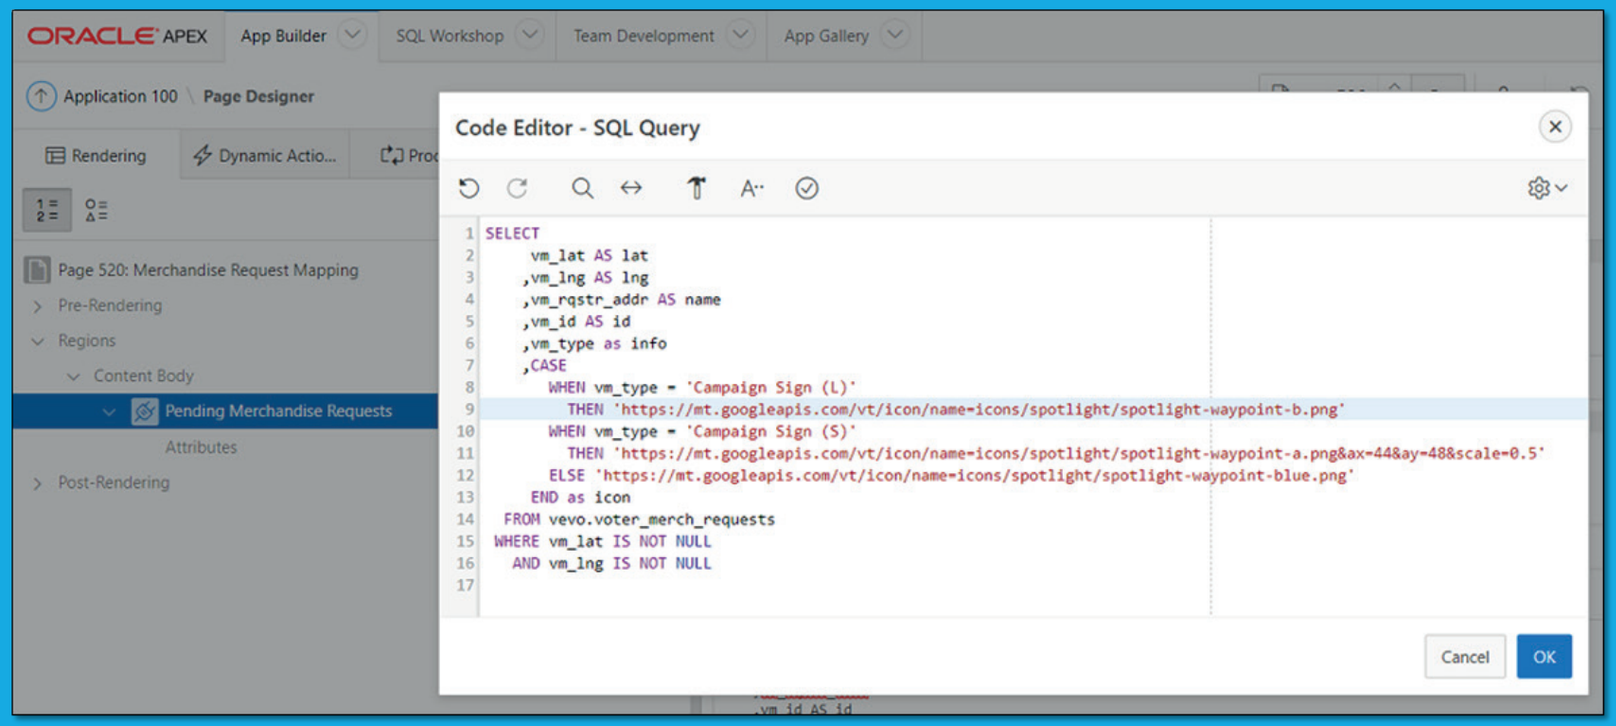

Figure 6. Providing Data for Plug-In Consumption

This plug-in also gives me quite a bit of control over how the map area itself displays initially, what mapping features are activated (for example, whether GeoJSON documents can be drag-and-dropped directly into the Map area), and several other settings. ***Figure 7*** lists just a few of the settings available, including the ability to place the initial center of my Map area based on a latitude / longitude pairing of 42.5° North and 87.5° East – in other words, just to the northwest of Chicago, Illinois.

Finally, here's the end result ***(Figure 8)***: a completely functional Google Map served up within my APEX application. Note how I the color-coded icons I specified in the query in Figure 6 to differentiate the large and small campaign signs from other merchandise to provide immediate visual feedback for the relative demand between those three categories.

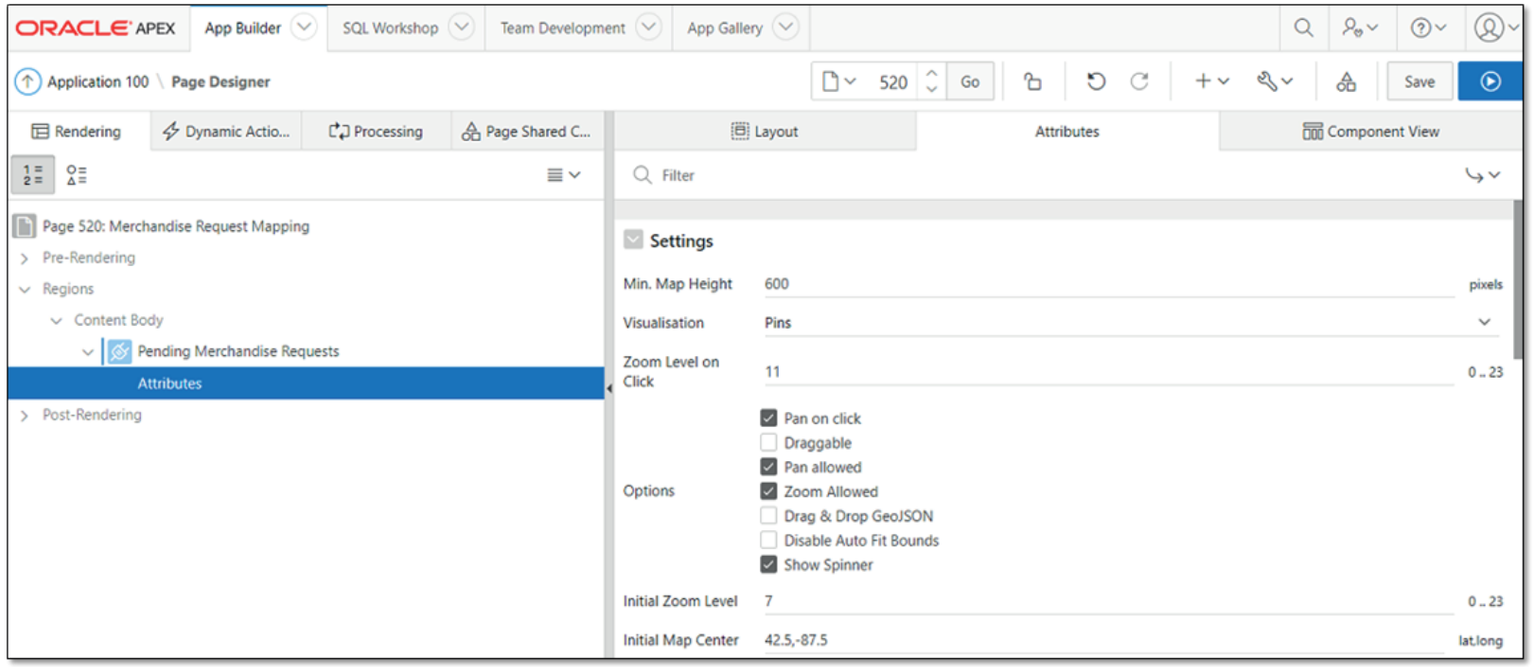

Figure 7. Controlling Google Maps API Plug-In Behaviors

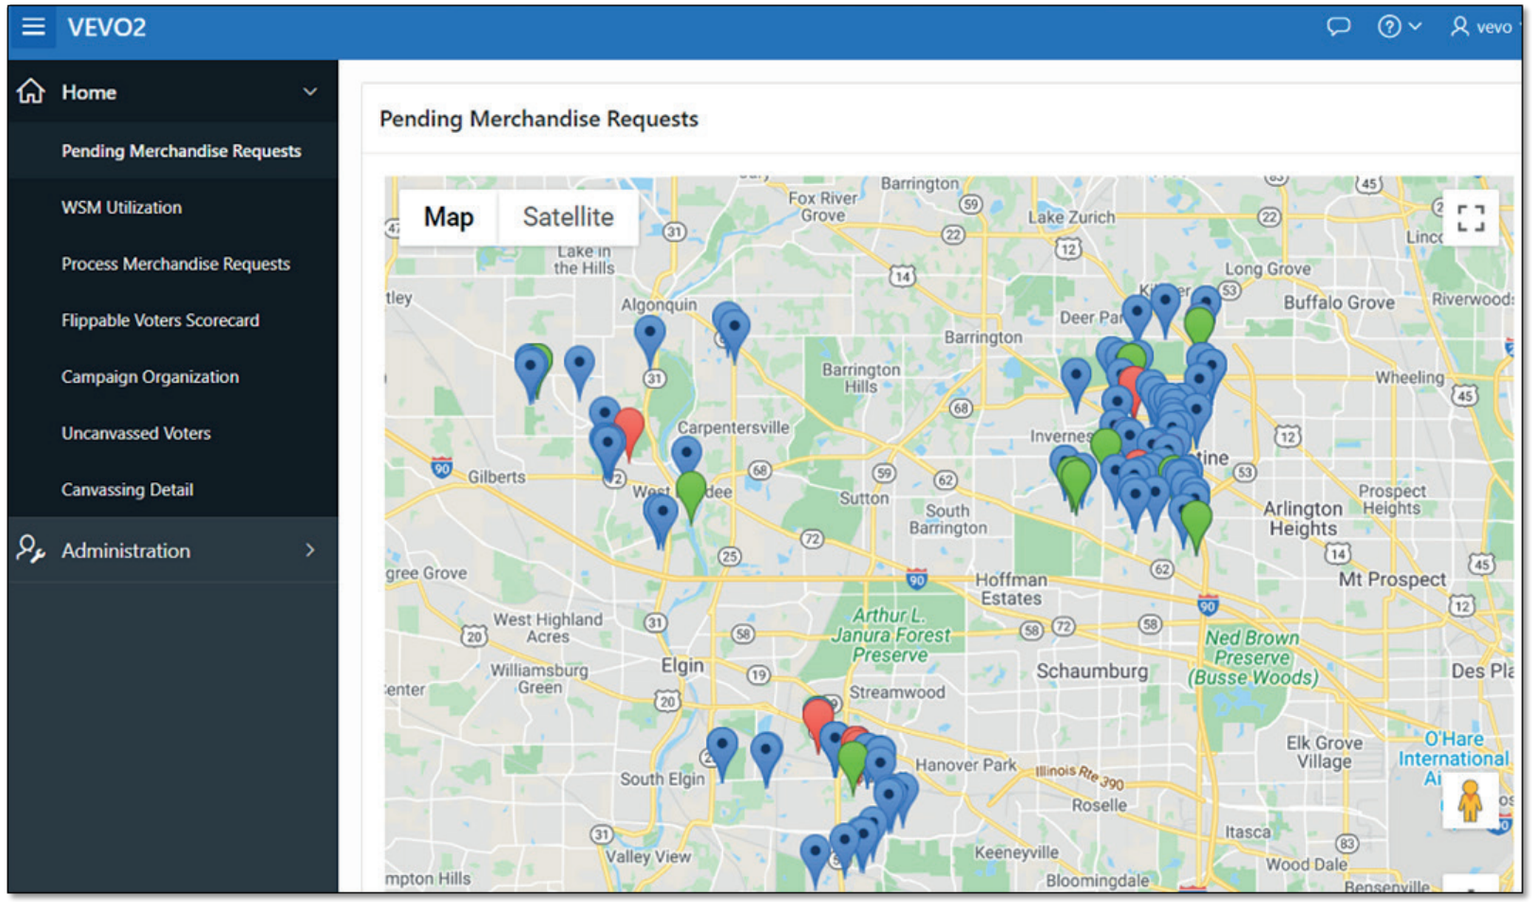

Figure 8. Displaying Color-Coded Points Within the Google Maps Plug-In

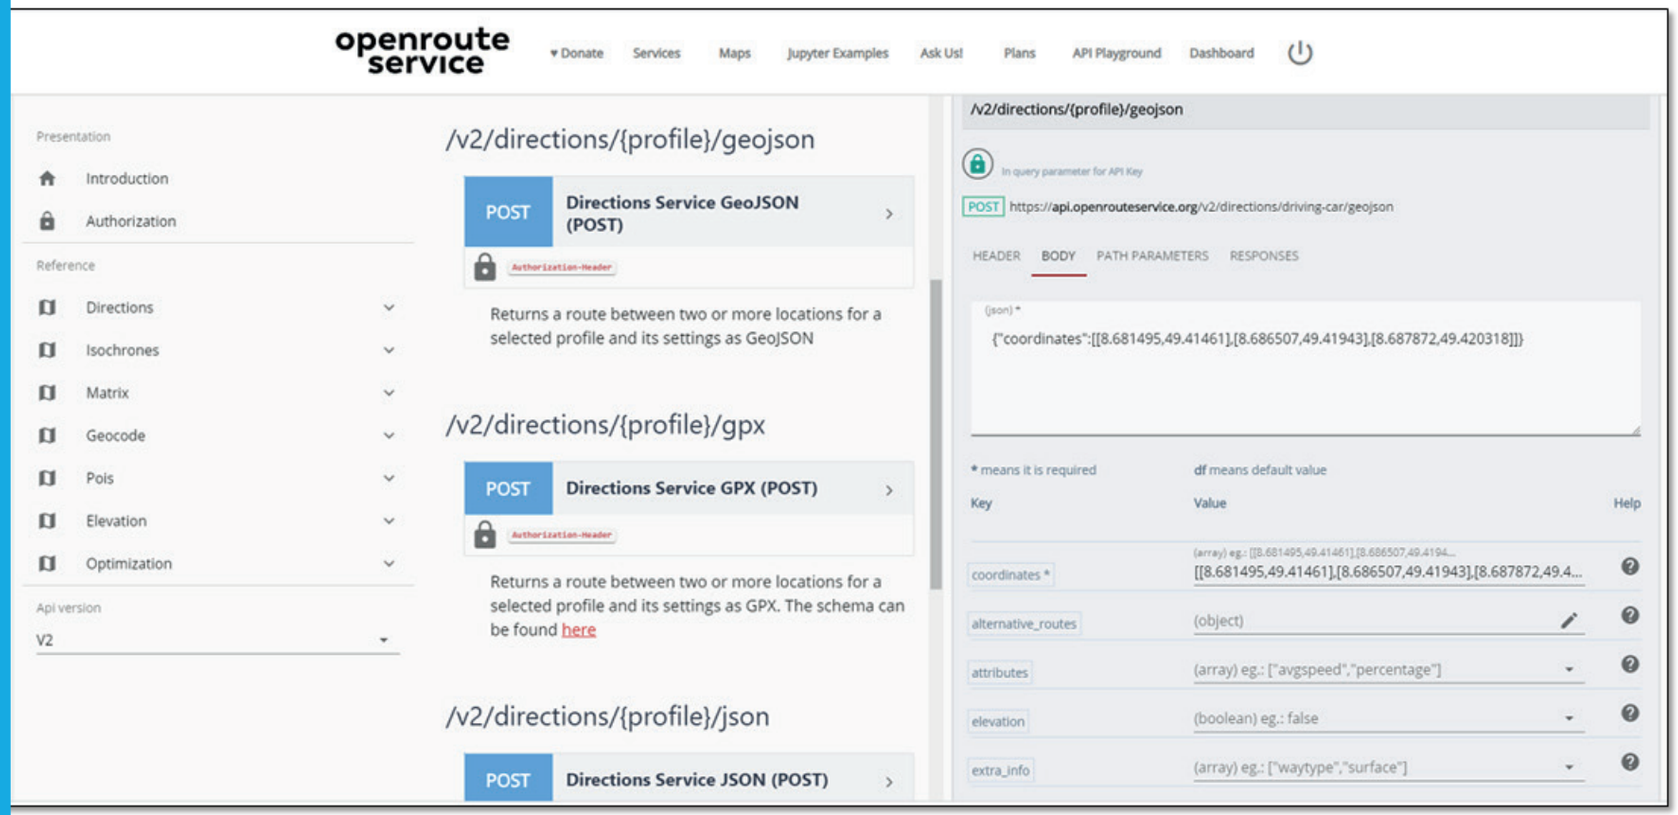

Figure 9. OpenRouteService: Testing Coordinates in JSON Format

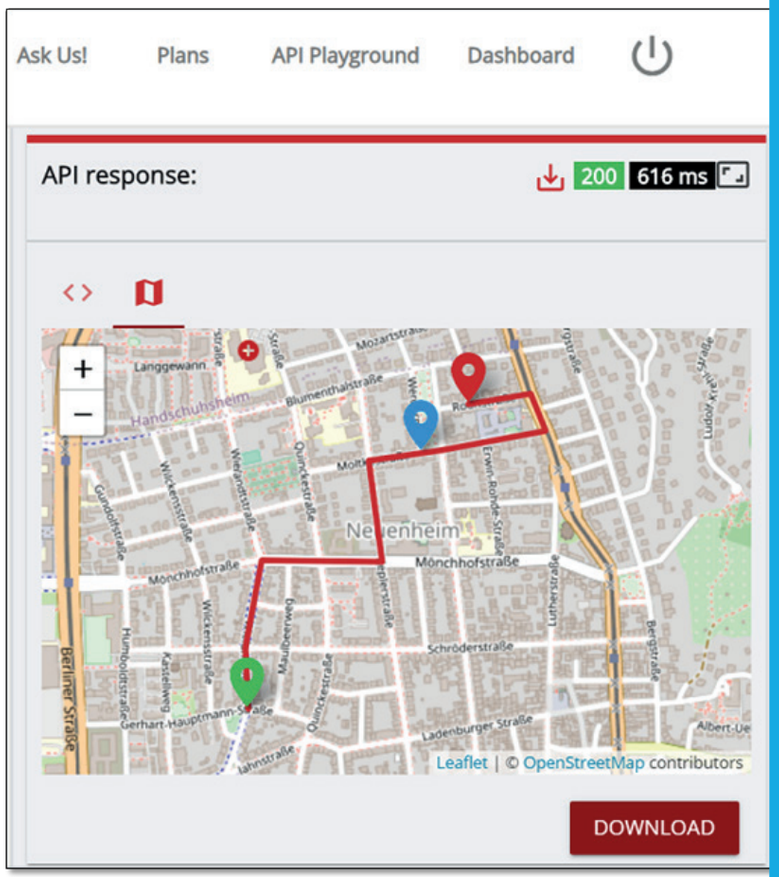

Figure 10. OpenRouteService: Displaying Mapping Results

Routing Solutions: Google Ain't the Only Game in Town

The Google Map API does provide methods to obtain a route based on a set of map points, and Jeffrey's plug-in does provide access to that **method** as well. However, as already mentioned, ramping up that routing solution could likely also exceed **the $200 monthly credit** for each Google developer license limit, as even simple route requests cost $5 per each 1000 requests.

The good news is that there are several lower-cost alternatives for developing routes, including a completely open-source solution, **OpenRouteService**, that's already in use by numerous public and private institutions. It provides a robust REST API that's quite simple to leverage, and its UI offers the ability to verify the formatting of input geolocation coordinates in JSON format ***(Figure 9)***. OpenRouteService even provides its own mapping capability to view and confirm the route it produces ***(Figure 10)***.

While these mapping results may not appear to be immediately useful for my APEX application's Google Map plug-in solution, OpenRouteService's REST API also allows me to download a JSON document that contains the results of its proposed route. That's something I can take advantage of, if I'm willing to delve a bit deeper into one of the plug-in's brilliant features: the **capability to drag and drop a GeoJSON shape into the Map area**.

As a simple demonstration, I uploaded a sample set of coordinates to OpenRouteService's routing engine and then clicked the **Download** button from the route map it produced. This produced a GeoJSON document – essentially, nothing more than the individual stops on the map, and a colored line that represented the most effective traffic route between those points. Then on the page that's using the Google Maps API plug-in, I added a new button to the same region as the Map area, and added a dynamic action for the button. Finally, I copied the GeoJSON output from OpenRouteService into the *Code* section of the *True* condition for that dynamic action, which completes a call to plug-in's **loadGeoJsonString** JavaScript method ***(Figure 11)***.

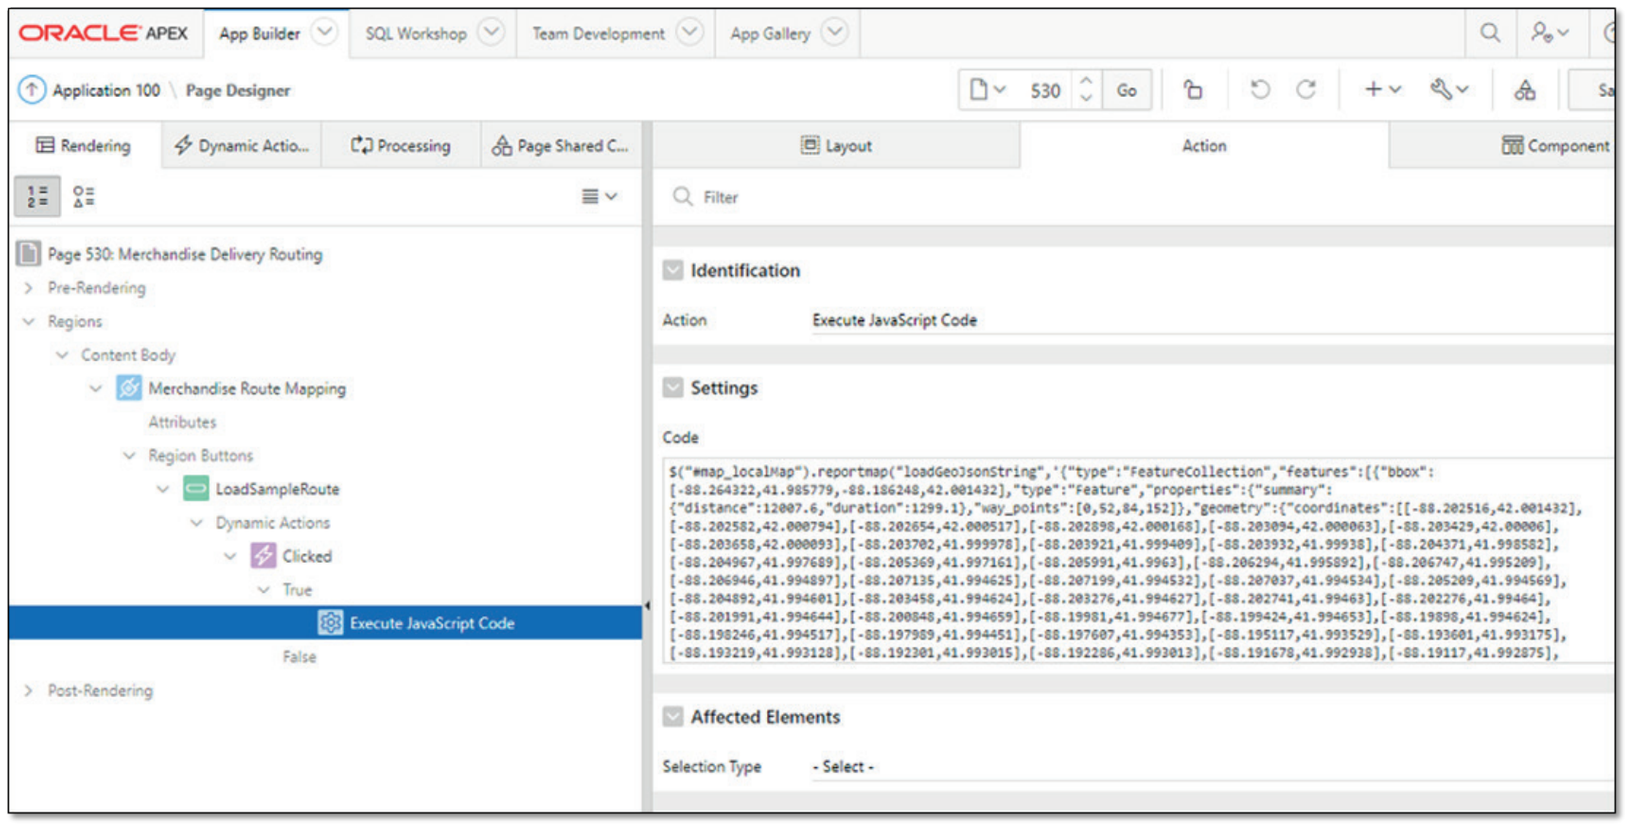

Figure 11. Adding GeoJSON for Display in Google Maps Plug-in

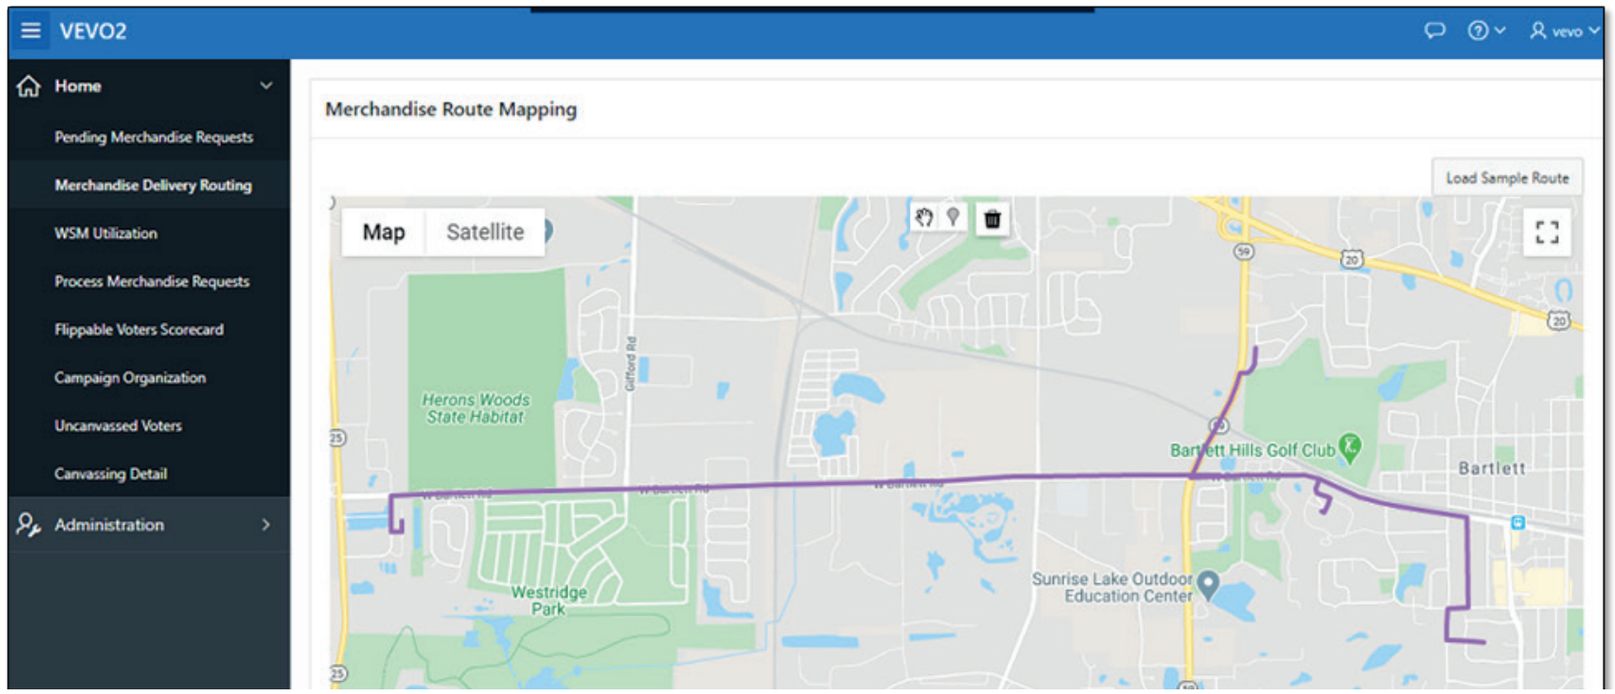

Figure 12. Resulting Route Map From GeoJSON Input

The resulting route map is shown in ***Figure 12***. While I'm sure you'd agree that the map this approach produces may not be as sophisticated as the one the Google Maps API methods would generate, that's not the point of my demonstration:

- First, this shows the inherent power of Jeffrey Kemp's plug-in to provide just about any mapping solution via its GeoJSON drag-and-drop method.
- This method also provides an effective way to use other open-source alternatives to the Google Maps API. That's especially important in a situation that requires close attention to the often-overlooked cost of accessing pricier APIs.

At Last, My Work Here Is Done.

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

- Use **APEX_DATA_PARSER** to load data directly into a table from external sources
- Apply features of *Web Source Modules* (WSMs) to obtain geocoding, mapping, and routing information
- Use **APEX_JSON** and **APEX_WEB_SERVICE** to construct input and process returned data in JSON format in concert with WSMs
- Incorporate an *APEX plug-in* to display voter merchandise requests on Google Maps from within my application

Thanks for coming along on the journey! I hope these articles and their numerous examples of how to deploy Machine Learning algorithms in various Oracle habitats – including OML, OAC, and APEX – have sparked your curiosity enough to continue learning about what the future holds for our industry and how you can take advantage of these trends to expand your professional accomplishments.

References

These references and links provide additional information and insight into the techniques used within this article. And be sure to pay particular attention to the examples and practices in Jeffrey Kemp's Google Maps API Plug-In – a little careful reading will save you an immense amount of wasted time and effort!

- APEX_DATA_PARSER: **https://docs.oracle.com/en/database/oracle/application-express/19.2/aeapi/APEX_DATA_PARSER.html**
- APEX_JSON: **https://docs.oracle.com/en/database/oracle/application-express/19.2/aeapi/APEX_JSON.html**
- APEX_WEB_SERVICE: **https://docs.oracle.com/en/database/oracle/application-express/19.2/aeapi/APEX_WEB_SERVICE.html**
- APEX Web Source Modules: **https://docs.oracle.com/en/database/oracle/application-express/19.2/htmdb/managing-web-source-modules.html**
- GeoCodio Geolocation Services: **https://geocod.io/**
- Jeffrey Kemp's Google Maps APEX Plug-In: **https://jeffkemponoracle.com/2016/02/google-map-apex-plugins/**
- OpenRouteServices Direction and Routing Services: **https://openrouteservice.org/**

This article was first published in ODTUG TechCeleration in July 2020.

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 blog **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.

Carsten Czarski

Oracle Application Express (Part 5): A Few Hidden Treasures

Introduction

Every release of Oracle Application Express comes with a number of new features. The marquee features like *Faceted Search*, *Interactive Grid* or *Web Source Modules* are often talked about and are thus known to almost every APEX developer.

However, there are also the smaller new features; which are hardly the topic of conference talks. And after some time, these are often forgotten. This article is about such smaller APEX features, or features which are available for years, but which people hardly talk about these days.

APEX is full of "Hidden Treasures" – and it's impossible for one article to list them all. Thus, this article talks about a choice of smaller APEX features which should be valuable for almost every APEX developer.

Classic Report Templates

When an APEX developer creates a report on a page, the default layout is as displayed in ***Figure 1***: a typical, tabular layout.

But Classic Reports belong to the most versatile components in Application Express. The reason is that these are *template-driven*. Developers can create own report templates and visualize data however wanted. And APEX even provides a few powerful templates out of the box. With these, developers can provide data visualizations, which make it hard to believe that there is an APEX classic report behind this.

Figure 2 shows an example for the **Cards**, and ***Figure 3*** is an example for the **Comments** report template, both provided out-of-the-box by APEX since APEX 5.0 (which was released back in 2015).

Classic Report

Empno	Ename	Job	Mgr	Hiredate	Sal	Comm	Deptno
7369	SMITH	CLERK	7902	12/17/1980	800		20
7499	ALLEN	SALESMAN	7698	2/20/1981	1600	300	30
7521	WARD	SALESMAN	7698	2/22/1981	1250	500	30
7566	JONES	MANAGER	7839	4/2/1981	2975		20
7654	MARTIN	SALESMAN	7698	9/28/1981	1250	1400	30
7698	BLAKE	MANAGER	7839	5/1/1981	2850		30
7782	CLARK	MANAGER	7839	6/9/1981	2450		10
7788	SCOTT	ANALYST	7566	12/9/1982	3000		20
7839	KING	PRESIDENT		11/17/1981	5000		10

Figure 1: By default, APEX reports display data in a tabular layout

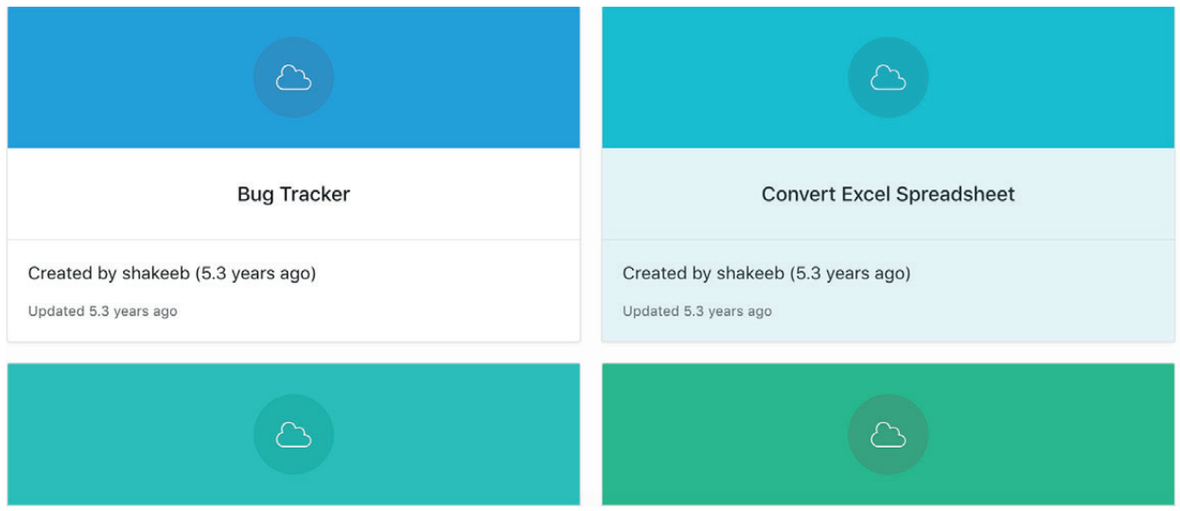

Figure 2: Cards Report Template

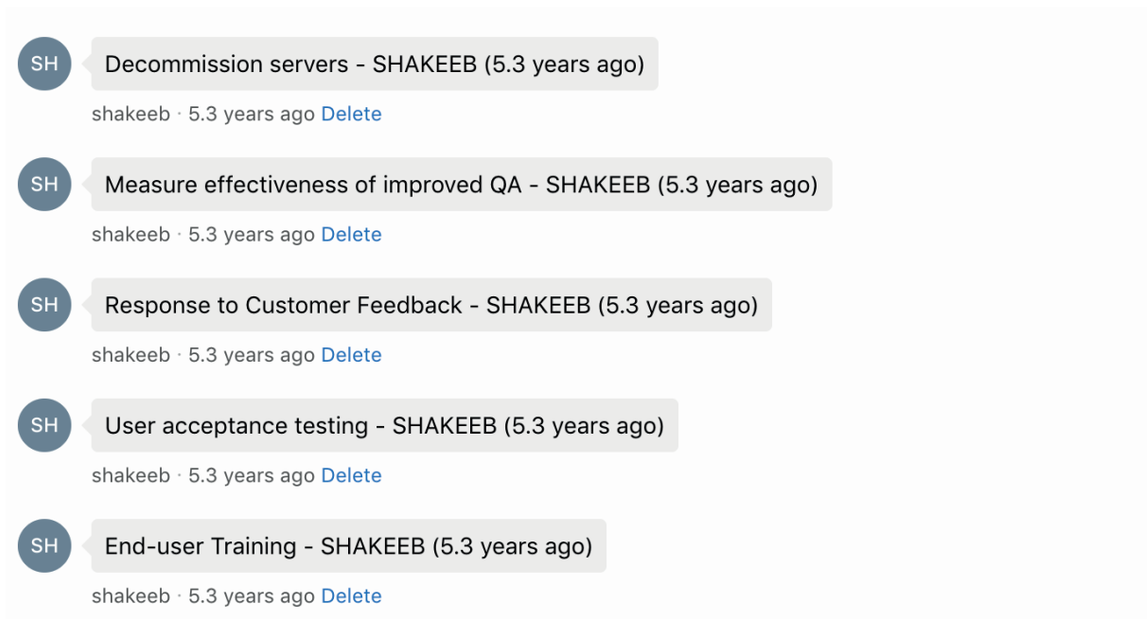

Figure 3: Comments Report template

Each of these report templates expects the SQL query to return *specific column names*: For instance, the **Comments** report template expects **COMMENT_DATE**, **USER_NAME**, **COMMENT_TEXT** and others. ***Listing 1*** shows an example query.

```
Select
  user_icon, -- use apex_string.get_initials(str) to get initials
  comment_date,
  user_name,
  comment_text,
  ' ' comment_modifiers,
  'u-color-'||ora_hash(user_name,45) icon_modifier,
  'Edit' actions,
  ' ' attribute_1,
:
from …
```

Listing 1: Example query for the "Comments" report template

Universal Theme Sample Application

This sample application is available in the App Gallery (Packaged Applications) and also publicly on **apex.oracle.com/ut** and it is more than just a demonstration app. It acts as a full reference for the Universal Theme. ***Figure 4*** shows the reference page for colors and their CSS classes in Universal Theme.

The **Button** and **Icon** builders are special gems within the Sample Universal Theme app. After configuring a button or a *Font APEX* icon, the app returns the HTML or CSS markup to be directly copied and pasted into the target APEX application ***(Figure 5)***.

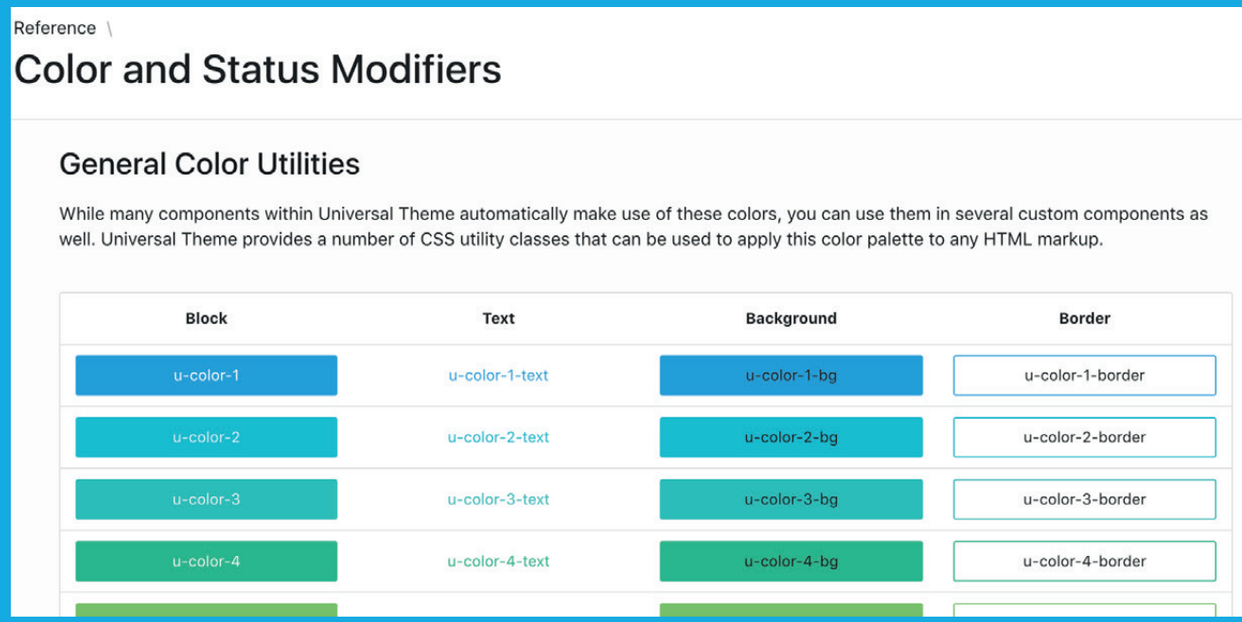

Figure 4: Reference Information about Universal Theme Colors.

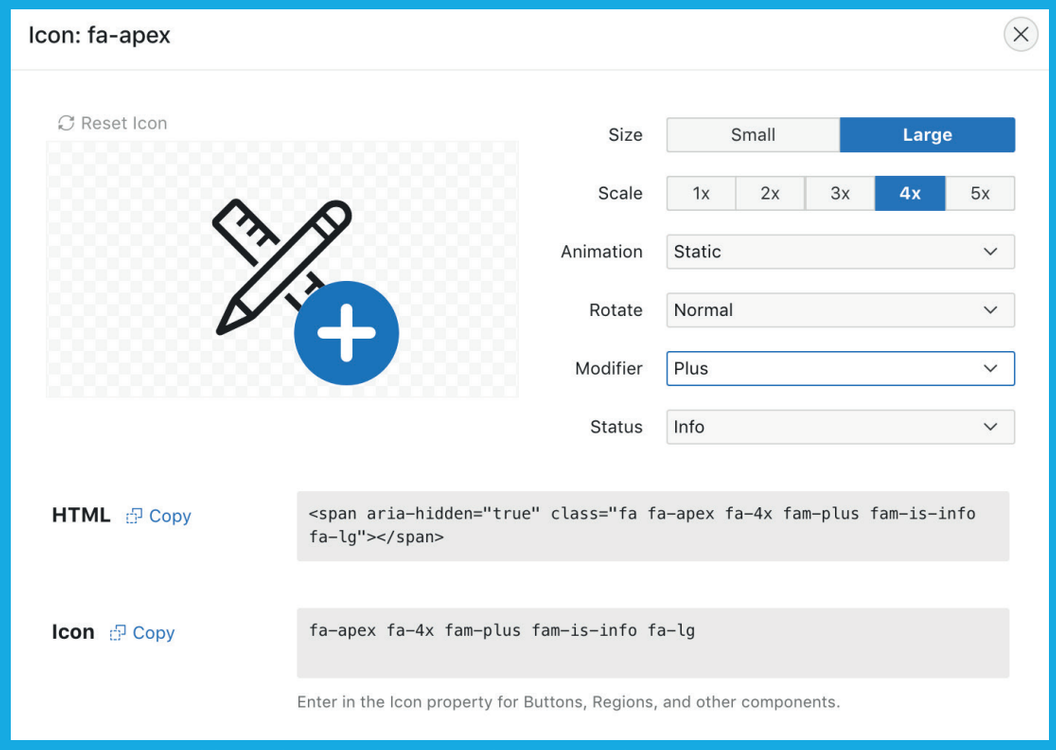

Figure 5: Icon Builder in the Sample Universal Theme Application

APEX PL/SQL Packages

APEX comes with a number of PL/SQL packages for application developers. These packages have the **APEX_** prefix and many of them provide very useful functionality.

APEX_STRING

APEX_STRING was introduced with APEX 5.1 and it contains, as the name indicates, useful utilities when working with strings. This article will highlight only two of them: **SPLIT** and **GREP**. However, the package contains even more functions, so developers should check the documentation out.
APEX_STRING.SPLIT tokenizes a string, based on a delimiter character ***(Listing 2)***.

```
select column_value as word
  from table( apex_string.split( 'This is a sentence with words',
' ' ));

WORD
---------
This
is
a
sentence
:
```

Listing 2: APEX_STRING.SPLIT tokenizes a string

The **GREP** function can be used to extract substrings based on regular expressions. Unlike the REGEXP_SUBSTR SQL function, **APEX_STRING.GREP** is able to extract *multiple* matches. ***Listing 3*** illustrates this by extracting all email addresses from a string.

```
select * from table( apex_string.grep(
  p_str =>      'This is a text with some mail adresses. The first
one' ||
                    ' is a.a@oracle.com, we have noname@company.
com'   ||
                    ' and thesupport@yetanothercompany.com',
  p_pattern  => '[A-Z0-9._%+-]*@[A-Z0-9.-]*\.[A-Z]*',
  p_modifier => 'i'
));

COLUMN_VALUE
--------------------------
a.a@oracle.com
noname@company.com
thesupport@ yetanothercompany.com
```

Listing 3: APEX_STRING.GREP extracts based on regular expressions

In APEX 20.1, the **APEX_STRING_UTIL** package was added, containing more specialized string utilities like **FIND_TAGS**, **GET_EMAIL_ADDRESSES** and others. The **TO_SLUG** function is often useful, as it converts a string to a "-" separated string, with special characters removed. The result can be made unique by providing a value for the **p_hash_length** parameter ***(Listing 4)***.

```
select apex_string_util.to_slug('This is a #question# !title!',4) slug
  from dual;

SLUG
--------------------------------------------------------------------------------
this-is-a-question-title-4253
```

Listing 4: TO_SLUG can simplify and "uniquify" a string

APEX_DATA_PARSER

The **APEX_DATA_PARSER** package (introduced in APEX 19.1) is the programmer's interface to the file parser in **SQL Workshop Data Loading**. As SQL Workshop, APEX_DATA_PARSER.PARSE is able to parse CSV, XLSX, XML or JSON files. The example in ***Listing 5*** downloads a sample XLSX (Spreadsheet) file from the internet, parses it and returns the contents as rows and columns.

APEX_DATA_PARSER.PARSE also contains a very special gem: Try to parse an iCal file (with the **.ics** extension). When downloading calendar information like a collection of public holidays from the internet, these are often provided as .ics files.

```
select col001, col002, col003, col004
  from table( apex_data_parser.parse(
                 p_file_name => 'file.xlsx',
                 p_content   => apex_web_service.make_rest_request_b(
                                  'https://examples.../XLSX_100.xlsx',
                                  'GET' ) ) );

COL001   COL002       COL003       COL004
-------- ------------ ------------ -------
0        First Name   Last Name    Gender
1        Dulce        Abril        Female
2        Mara         Hashimoto    Female
3        Philip       Gent         Male
:        :            :            :
```

Listing 5: Parse an XLSX file from the internet with APEX_DATA_PARSER

APEX_ZIP

APEX_ZIP is there since APEX 5.0 and allows to create or process ZIP archives directly in APEX. When end users upload ZIP files to an APEX application, contents can be extracted and processed with APEX_ZIP. Also, multiple files can be provided as a single ZIP download, instead of the end user having to downloads all files one after the other.

Listing 6 shows, how an uploaded file (e.g. from a File Browse item) is unzipped with **APEX_ZIP.GET_FILES**. All extracted files are then stored into the **MY_UNZIPPED_FILES** table.

```
declare
    l_uploaded_file blob;
    l_unzipped_file blob;
    l_zip_files     apex_zip.t_files;
begin
    select blob_content
      into l_uploaded_file
      from apex_application_temp_files
     where name = :PX_FILE_BROWSE;

    -- get the list of file names within the ZIP archive
    l_zip_files := apex_zip.get_files( l_uploaded_file );

    for f in l_zip_files.first..l_zip_files.last loop
        -- extract the contents of the file from the ZIP archive
        l_unzipped_file := apex_zip.get_file_content( l_uploaded_file, l_zip_files( f ) );

        insert into my_unzipped_files ( file_name, file_content )
        values ( l_zip_files( f ), l_unzipped_file );
    end loop;
end;
```

Listing 6: Unzipping an uploaded ZIP archive

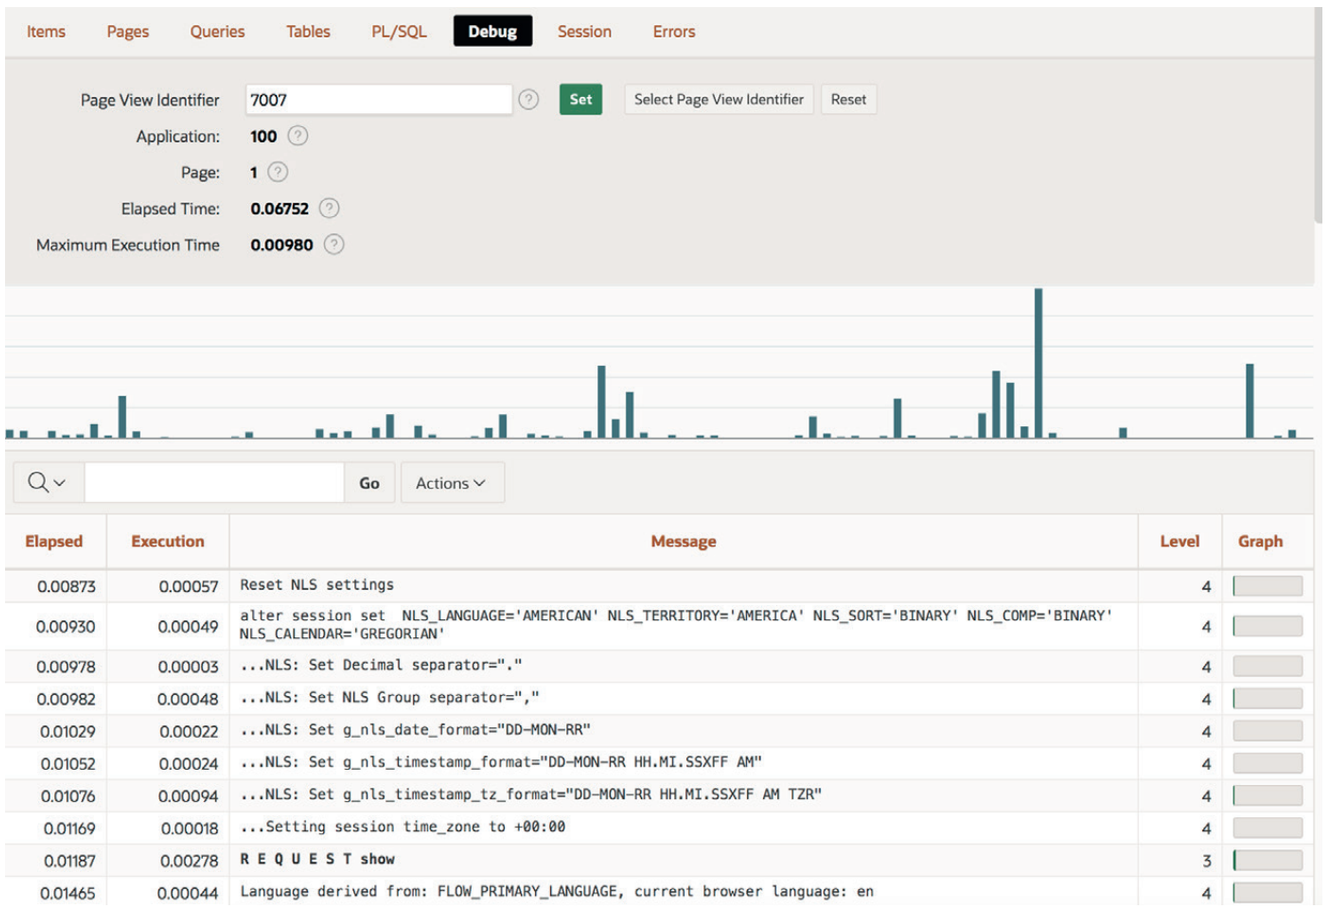

Elapsed	Execution	Message	Level	Graph
0.00873	0.00057	Reset NLS settings	4	
0.00930	0.00049	alter session set NLS_LANGUAGE='AMERICAN' NLS_TERRITORY='AMERICA' NLS_SORT='BINARY' NLS_COMP='BINARY' NLS_CALENDAR='GREGORIAN'	4	
0.00978	0.00003	...NLS: Set Decimal separator="."	4	
0.00982	0.00048	...NLS: Set NLS Group separator=","	4	
0.01029	0.00022	...NLS: Set g_nls_date_format="DD-MON-RR"	4	
0.01052	0.00024	...NLS: Set g_nls_timestamp_format="DD-MON-RR HH.MI.SSXFF AM"	4	
0.01076	0.00094	...NLS: Set g_nls_timestamp_tz_format="DD-MON-RR HH.MI.SSXFF AM TZR"	4	
0.01169	0.00018	...Setting session time_zone to +00:00	4	
0.01187	0.00278	**R E Q U E S T show**	3	
0.01465	0.00044	Language derived from: FLOW_PRIMARY_LANGUAGE, current browser language: en	4	

Figure 6: Reviewing Debug Output for an APEX page

```
declare
    l_number number;
begin
    apex_debug.info('This is an page process on page %s', :APP_PAGE_ID );

    l_number := 10 / 0;
exception
    when others then
        apex_debug.error( 'An error occurred: %s', sqlerrm);
end;
```

Listing 7: PL/SQL Page Process with APEX_DEBUG instrumentation

0.01229	0.00002	**> APEX_APPLICATION_PAGE_PROCESS So something**	3	
0.01230	0.00085	...Process "So something" - Type: NATIVE_PLSQL	4	

0.01315	0.00083	begin declare
		l_number number;
		begin
		apex_debug.info('This is an page process on page %s', :APP_PAGE_ID);
		
		l_number := 10 / 0;
		exception
		when others then
		apex_debug.error('An error occuured: %s', sqlerrm);
		end;
		end;
			4	

0.01398	0.00042	This is an page process on page 1	4	
0.01440	0.00012	**An error occuured: ORA-01476: divisor is equal to zero**	1	
0.01452	0.00002	**< APEX_APPLICATION_PAGE_PROCESS So something**	3	
0.01454	0.00036	...close http header	4	
0.01490	0.00324	Show page template header	4	

Figure 7: Custom messages logged with APEX_DEBUG, will be part of debug output

Debug Mode

At the first glance, working with APEX *Debug Mode* is rather straightforward: It is enabled by clicking **Debug** in the *Developer Toolbar*. The page will be reloaded, while debug information is being recorded. This debug information can then be reviewed by clicking **View Debug** in the Developer Toolbar. Debug output contains the individual actions executed by APEX, together with elapsed time information ***(Figure 6)***.

Developers can also write their own messages to the debug log, by using the **APEX_DEBUG** package. The example in ***Listing 7***; executed as a page process in the **Pre-Rendering** section, writes a few messages to the debug log when the page is loaded (and also has some exception processing).

If debug is enabled, these debug messages will be part of debug output ***(Figure 7)***. Unhandled PL/SQL exceptions and messages logged with APEX_DEBUG.ERROR are displayed in red.

0.55711	0.00115	execute	8
0.55826	0.45972	determining execution plan	9

```
SQL_ID  asy9dzbhh808h, child number 1
-------------------------------------
select * from(select a.*,row_number() over (order by null) apx$rownum
from(select i.*  from (select "EMPNO","ENAME","JOB","MGR","HIREDATE","SA
L","COMM","DEPTNO" from ((select /*+ qb_name(apex$inner)
*/d."EMPNO",d."ENAME",d."JOB",d."MGR",d."HIREDATE",d."SAL",d."COMM",d."D
EPTNO" from (select x.* from "EMP" x  ) d  )) i  ) i where 1=1  order
by "EMPNO" asc nulls last )a )where apx$rownum<=:p$_max_rows

Plan hash value: 1478970339

-----------------------------------------------------------------------------------------
| Id  | Operation                     | Name     | Rows  | Bytes | Cost (CPU)| Time     |
-----------------------------------------------------------------------------------------
|   0 | SELECT STATEMENT              |          |       |       |     2 (100)|          |
|*  1 |  VIEW                         |          |    14 |  1400 |     2   (0)| 00:00:01 |
|*  2 |   WINDOW NOSORT STOPKEY       |          |    14 |  1218 |     2   (0)| 00:00:01 |
|   3 |    VIEW                       |          |    14 |  1218 |     2   (0)| 00:00:01 |
|   4 |     TABLE ACCESS BY INDEX ROWID| EMP      |    14 |   532 |     2   (0)| 00:00:01 |
|   5 |      INDEX FULL SCAN          | PK_EMPNO |    14 |       |     1   (0)| 00:00:01 |
-----------------------------------------------------------------------------------------

Predicate Information (identified by operation id):
---------------------------------------------------

   1 - filter("APX$ROWNUM"<=:P$_MAX_ROWS)
   2 - filter(ROW_NUMBER() OVER ( ORDER BY  NULL )<=:P$_MAX_ROWS)
```

1.01798	0.00067		9
1.01865	0.00035	populate_buffer p_columns_values_row=>,p_context.query_result.has_more=>true	8
1.01900	0.00172	fetch_rows	8

Figure 8: Level 9 Debug will record information about the SQL execution plan

Enabling debug mode and reviewing debug output is very helpful in diagnosing all sorts of issues with an APEX page or application. If an APEX component (e.g. report, chart or others) fails during execution, there is always valuable information available in debug output.

But debug mode can do even more. If enabled at the highest level (**Level 9**), APEX will log information about the actual *execution plan* being used when the SQL statement was executed ***(Figure 8)***. This is super-helpful for improving query performance.

When debug is enabled, APEX uses Level 4 *(Info)* by default. Switching to higher levels gets easier with the latest APEX release 20.1: Simply pick the desired level in the developer toolbar.

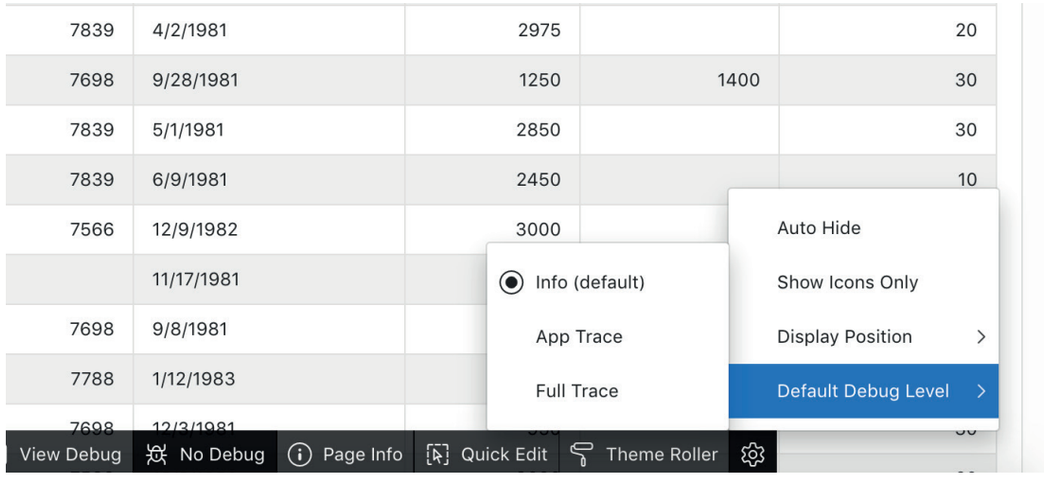

Fig. 9: Choosing another Default Debug Level

Page Designer: Two Tips

Page Designer includes a number of configuration options to increase developer productivity. Many developers think that Page Designer wastes too much space with the *Layout View* in the center of the screen. These can simply switch from to **Two Pane Mode** ***(Figure 10)***. After that, all tabs can be rearranged by just dragging them around.

The **Pin Filter** functionality is very useful, when the same property has to be changed for many different components on the page. In this situation, look up the property by using the **Filter** functionality on the top of the property pane. Then click the "pin" icon in order to keep that filter active (and the desired property focused) when another page component is selected ***(Figure 11)***.

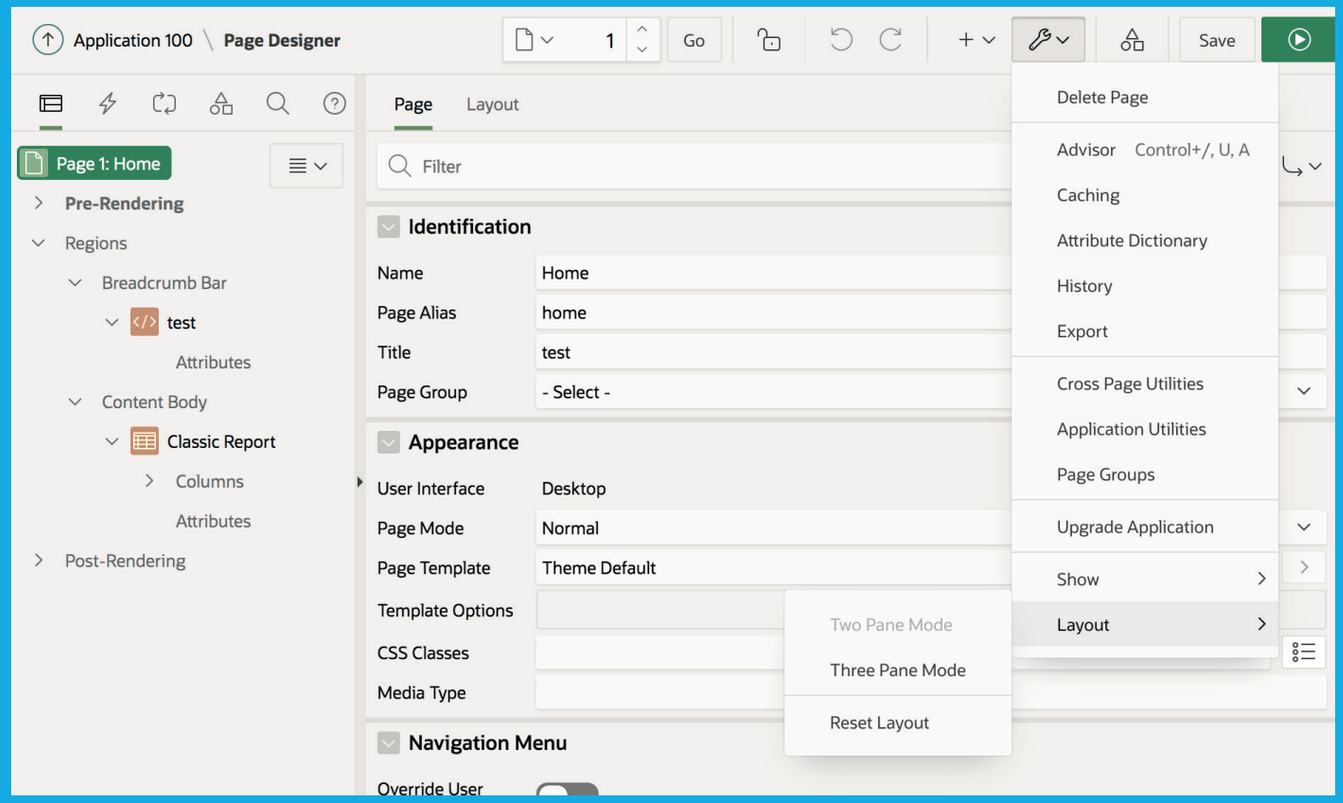

Figure 10: Page Designer in Two Pane Mode

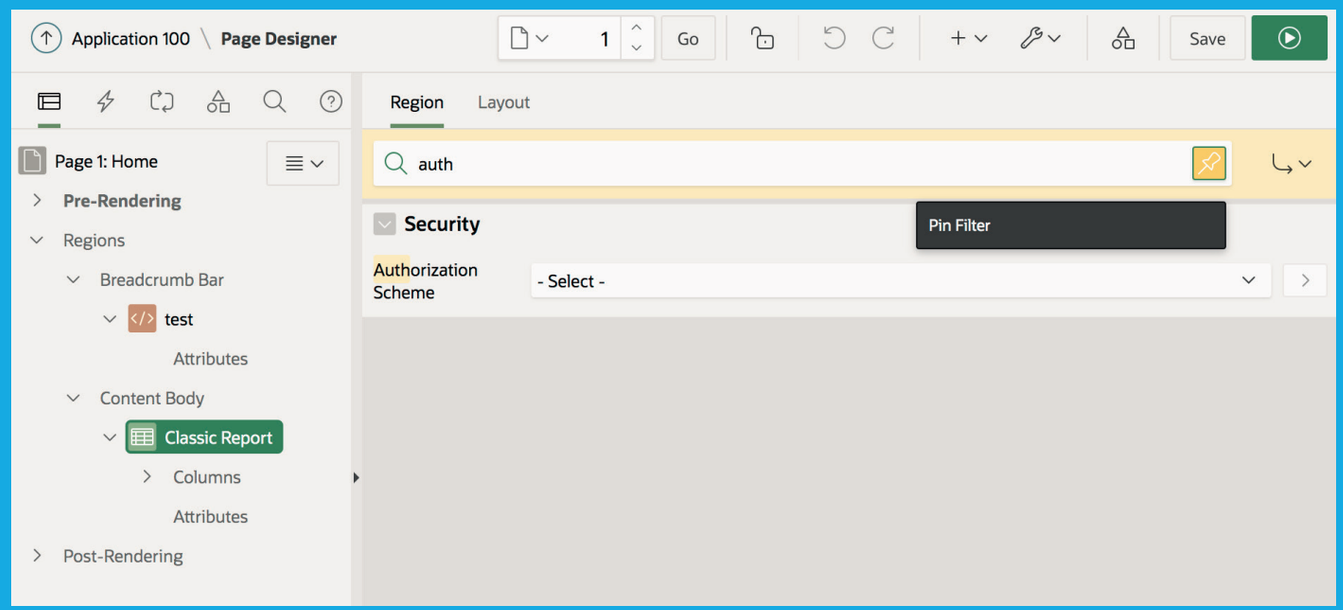

Figure 11: Pin Filter functionality in Page Designer

More Information

- Universal Theme Sample Application
  http://apex.oracle.com/ut
- Blog Posting on APEX Classic Report Templates
  https://blogs.oracle.com/apex/thats-a-classic-report-really
- Blog Posting on the APEX_DATA_PARSER package
  https://blogs.oracle.com/apex/super-easy-csv-xlsx-json-or-xml-parsing-about-the-apex_data_parser-package
- APEX 20.1 API Reference
  https://apex.oracle.com/api
- Information and environment to test-drive APEX
  http://apex.oracle.com/en
- Oracle Application Express Blog
  http://blogs.oracle.com/apex

About Carsten Czarski

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

OCI Block Volume: Auto-tune for Detached Volumes

André Luiz Dutra Ontalba

This summer, Oracle launched a great feature that helps reducing your costs for the Oracle Cloud Infrastructure Block Volume service. Let's take a look at this feature!

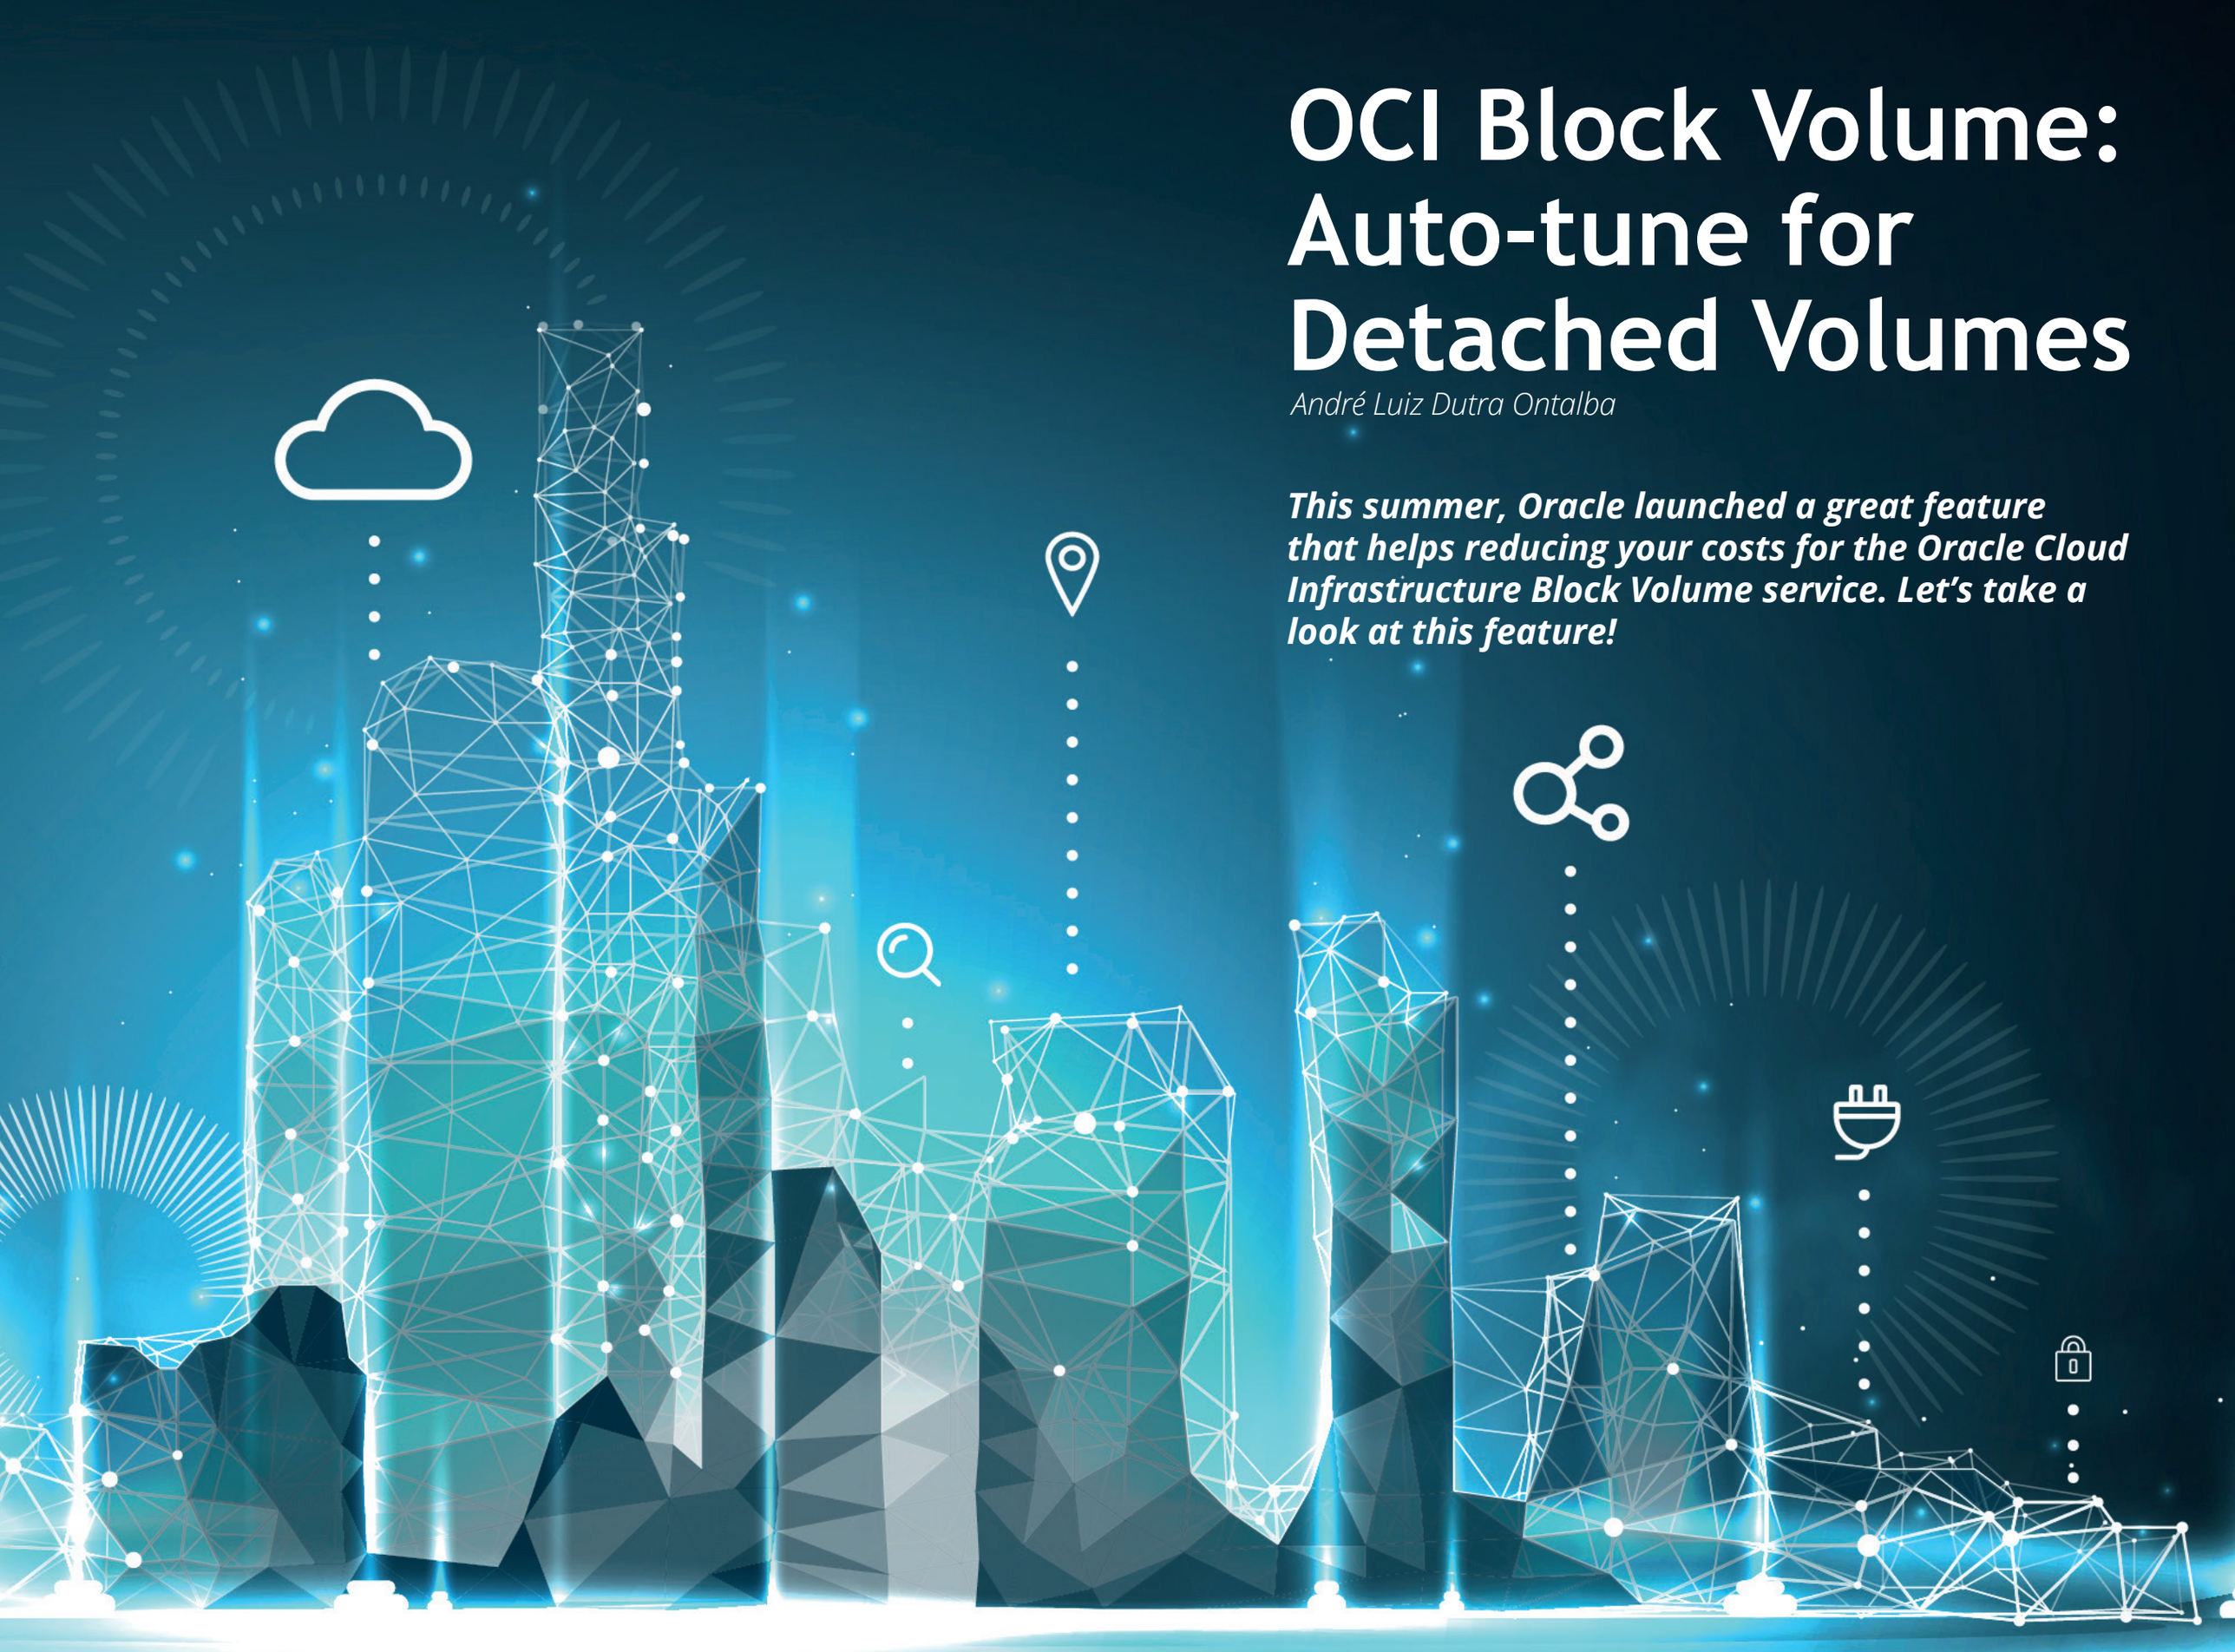

You can now tune the performance of your detached volumes to the Lower Cost setting automatically. With this new capability, you can achieve significant cost savings while your volumes stay in a detached state. When you're ready to use them for your workloads, simply attach them, and their performance and cost are automatically and instantaneously adjusted to the performance setting you originally configured.

When you enable this feature for your volumes, the volume is monitored and changed to the Lower Cost performance option automatically when it is disconnected and maintained in this setting until you reconnect it. This feature now comes integrated with the storage platform. You can take advantage of this with a click on the console or by using a command line option in the CLI for each of your volumes.

Auto-tune for detached volumes capability is generally available for all existing and new boot and block volumes in all global regions, on the Console, and through CLI, SDK, API, and Terraform.

Enabling and Managing Auto-tune for Detached Volumes

Enabling the Auto-tune feature for detached volumes is straightforward with a click in the Oracle Cloud Infrastructure Console. To Enable auto-tune for a volume, on the Block Volume Details screen of the Console, click **Edit** and slide the Auto-tune Performance setting from **Off** to **On** for the volume ***(Figure 1)***. The Edit dialogue window is also revised as part of this feature update.

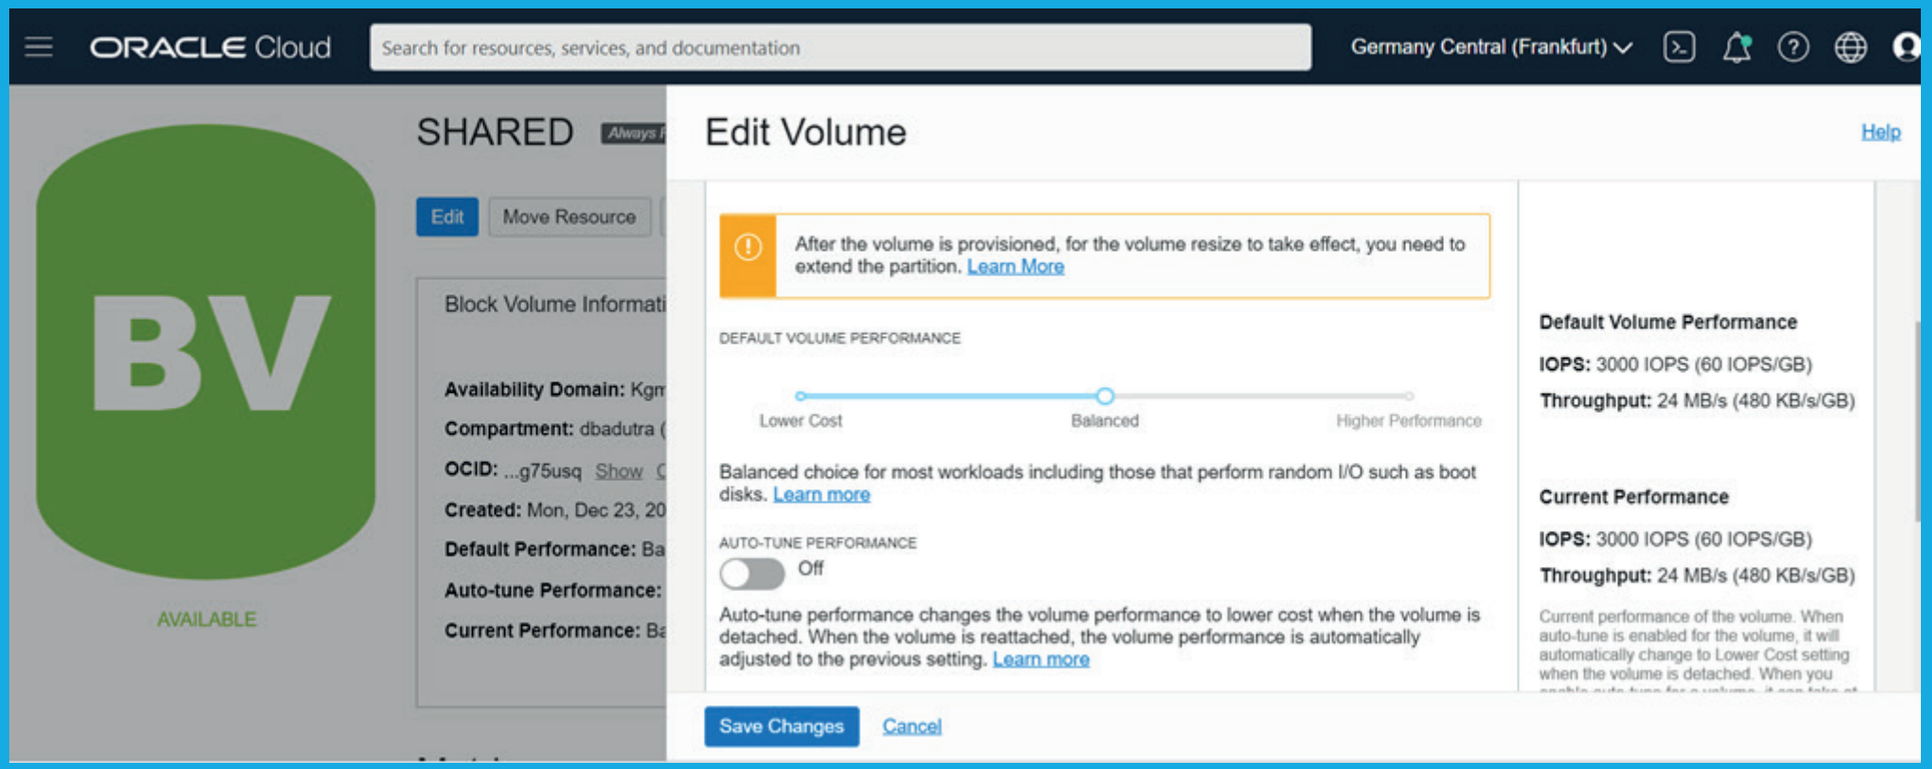

Figure 1

When the Auto-tune Performance is set to **On** for a detached volume, the auto-tuning takes effect after 24 hours ***(Figure 2)***. After that, if the volume is still detached, its performance and cost is lowered to the Lower Cost setting automatically ***(Figure 3)***.

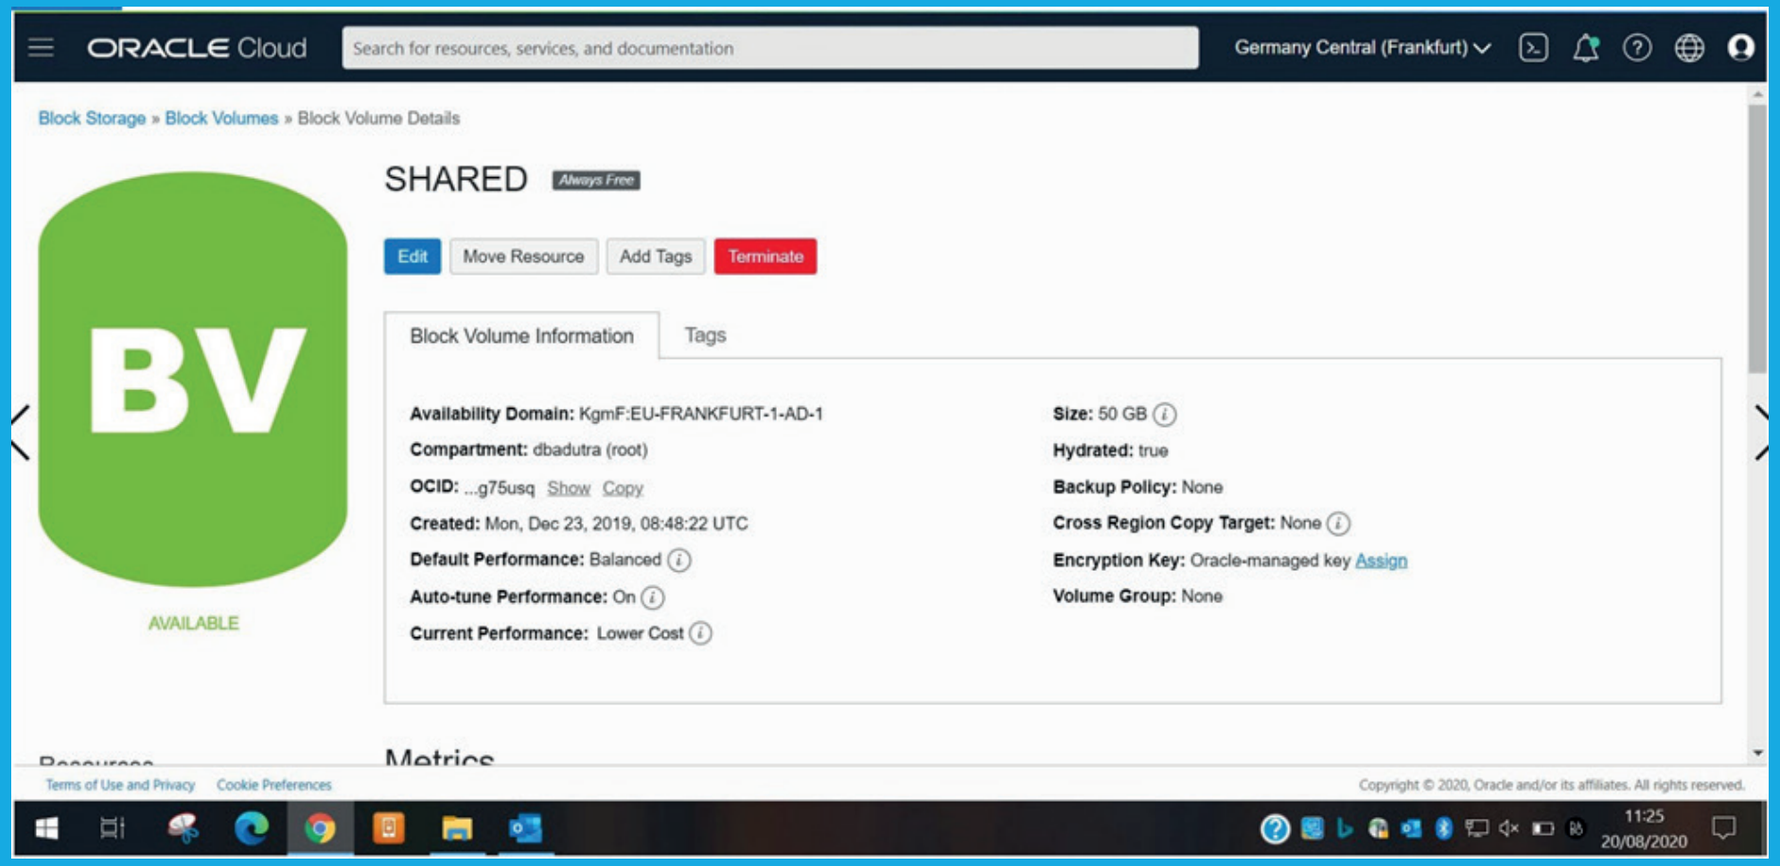

Figure 2

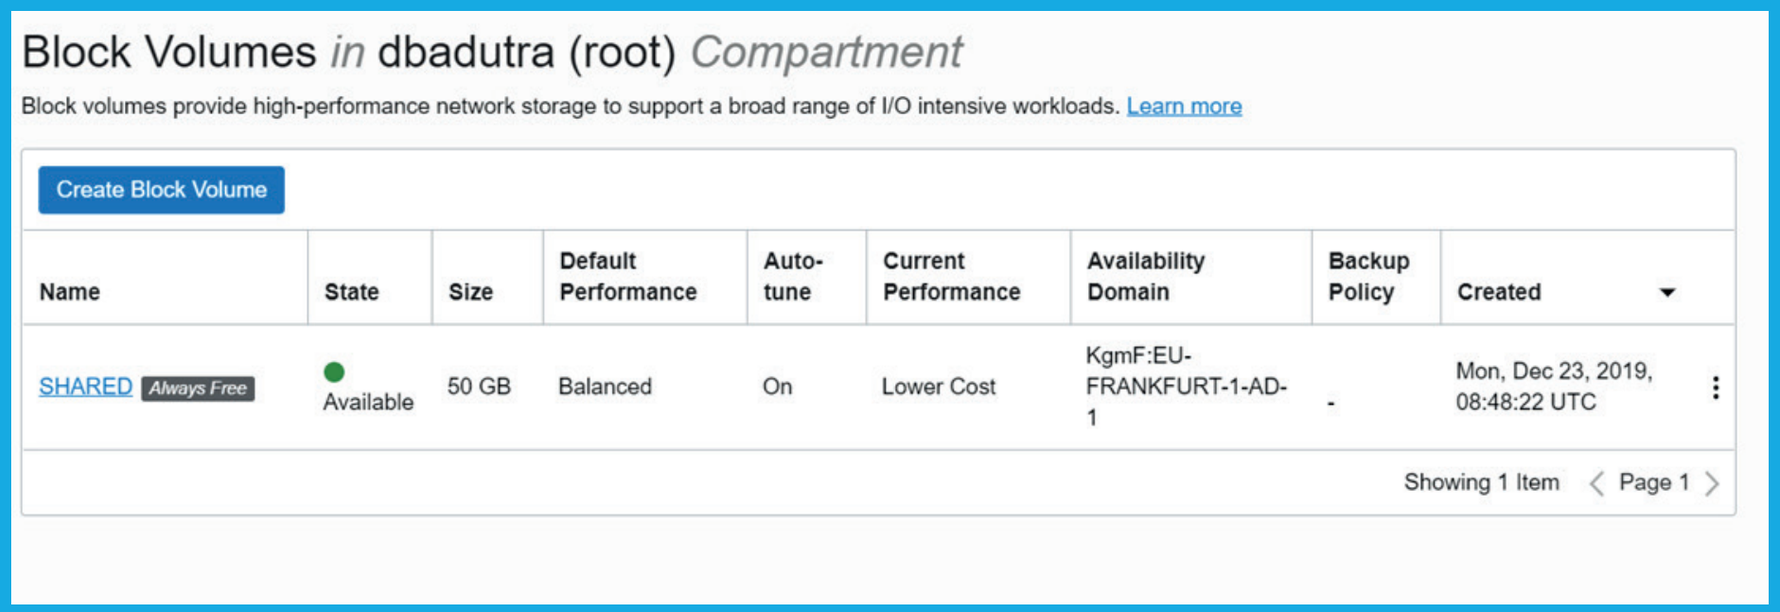

Figure 3

When the volume is attached again, its performance is set to the Default Performance setting immediately and automatically.

Here is an overview of the highlights that come with this new feature:

- Build-in feature – new volumes or existing ones
- Applicable to Block & Boot volumes
- Available in all regions
- Reduces operational expenses
- Once enabled, Auto-tune takes 24 hours to be effective
- Done via Console, CLI, SDK, API, Terraform
- Switches automatically to lower cost performance option when detached from instance
- When attached again it gets to previous defined settings

I hope this helps you!

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.

Emmanuel Ruez

Activities of the French-speaking Oracle User Clubs during the Pandemic

During the health crisis and the containment imposed by the French Government, the three Oracle User Clubs, AUFO (Association des Utilisateurs Francophones d'Oracle), the Groupe Francophone des Utilisateurs J.D. Edwards and the PeopleSoft User Club remained very close to their members.

The 3 clubs have been particularly dynamic by offering their members numerous webinars. Nearly 20 virtual meetings were organized on various subjects close to the users' concerns, in relation to Oracle news and the health context. Oracle France made its professional tool Zoom available to the clubs. Some webinars were co-organized with Oracle partners, which also made it possible to maintain links with the publisher's ecosystem. A few examples of webinars organized (the full program can be found **on our website**):

AUFO Webinar | April 23, 2020
Cloud Migration Commission | Core Model in an International Context | SMCP & Auchan Feedback

J.D. Edwards Webinar | April, 28, 2020
Dematerialization of supplier invoices | Eranove feedback

PeopleSoft Webinar | March 31, 2020
Migrating your PeopleSoft ERP to the Cloud with the SOAR method

In general, although the workload for many of our members as well as for the Club administrators was particularly heavy during this period, we had more subscribers than usual to these webinars and fewer no-shows. All webinars, after agreement of the speakers and participants, were recorded. We were thus able to make the presentations available to the users (as we usually do) but also the replay of the webinar, accessible on the YouTube channel of the Oracle User Clubs. We have had very good feedback from members on the organized webinars and even new members!

We already have a nice program of meetings planned for the new school year. All meetings are scheduled in webinar format for the moment, according to the schedule **on our website**. Some of them will be organized as face-to-face meetings if possible.

As some of the topics are transversal to the three clubs, members of all three clubs will be invited, which will allow for an even wider exchange between users.

In the last months, we were proud to see that our members' interest in the clubs remained intact and even grew thanks to the exchanges proposed. We appreciated the webinars organized by Oracle France which gave the floor to users: Mutuelle Générale, Louvre Hotels.

Finally, we took the time to take part in webinars on very interesting topics to help us better understand IT evolution, such as those organized by Hub Institute (one of whose partners will be speaking at our next Users' Day).

We were able to see that companies that had already implemented digitization solutions were able to continue their business more easily. Remote access was facilitated, and work could be carried out efficiently. This was discussed during the various webinars organized and the feedback was very eloquent. For example, the Eureden testimonial given on June 30th during the Mobility webinar organized by the J.D. Edwards Club demonstrated the value of the solutions implemented. Similarly, companies that had already implemented Business Continuity Plans with their partners were little affected by the crisis.

We would like to return to face-to-face meetings as soon as possible. Indeed, these exchanges are much richer. Speaking up is simpler. Speech is freer.

However, we will certainly propose more systematically the possibility of remote participation.

In France, we were confined from 17 March to 11 May, few French companies welcome all their employees back to their premises.

There are two trends: a partial reception (maximum 50% of the workforce at the same time) or a complete maintenance of teleworking for all employees. Concerning the possibility of participating in an external event, the tendency is rather to discourage or even prohibit it.

We hope that the health situation will improve and that we will be able to organize our annual event on November 25. To be continued...

About Emmanuel Ruez

Emmanuel, 56, is a graduate of the University of Grenoble. He spent the first ten years of his career working for consulting firms and software publishers, including Oracle, and has worked with a number of major French groups to assist them in implementing IT solutions. From 1998, Emmanuel was recruited by Gaumont where he was in charge of the eBs solution. In 2001, he joined the Accounting Department of the Orange Group and worked on a number of IT solution creation and implementation projects. Emmanuel has been a member of AUFO since 1998 and became its president this year.

Call for Papers

Kscope 21

November 17, 2020
Nashville, TN, USA
https://kscope21.odtug.com/

APEX World 2021

November 30, 2020
Netherlands + online
https://www.nloug.nl/events/apex-world-2021

Events

DOAG 2020 Conference + Exhibition

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

Oracle Groundbreakers EMEA Virtual Tour

October 1 - 15, 2020
online
http://ogbemea.com/

ILOUG TechDays 2021

February 3 - 5, 2021
Tel Aviv, Israel
https://www.iloug.org/

JavaLand 2021

March 16 + 17, 2021
Phantasialand Brühl, Germany + online
https://www.javaland.eu/en/hybrid/

APEX World 2021

March 18, 2021
Netherlands + online
https://www.nloug.nl/events/apex-world-2021

Kscope21

June 20 - 24, 2021
Nashville, TN, USA
https://kscope21.odtug.com/

HrOUG 2020 (postponed)

Postponed to 2021
Rovinj, Croatia
https://2020.hroug.hr/eng/

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:

Caroline Sengpiel, 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 the next edition: **October 29, 2020**.

Photo credits:

Title: © pch.vector | https://freepik.com
P. 4: © floralset | https://123rf.com
P. 5: © Comic | www.commitstrip.com
P. 6: © Freepik | https://freepik.com
P. 7 + 8, Objects: © marcovector | https://freepik.com
P. 9 + 13 © Feodora Chiosea | https://123rf.com
P. 17: © Alexander | https://stock.adobe.com
P. 18 - 21, Clouds: © Freepik | https://freepik.com
P. 23: © Freepik | https://freepik.com
P. 36: © Anastasiia Nevestenko | https://123rf.com
P. 45: © O_du_van | https://123rf.com
P. 49: © Freepik | https://freepik.com
P. 52 + 53: © DOAG | www.doag.org/en/home