

› APEX 20.2:
What's new?

› EMEA Virtual Tour 2020:
#BackToBasics

› CLOUD Act:
Protection of
Business Secrets

At Face Crypto Value

A Look behind the Hype of Bitcoin and Co.





CLOUD Act Series Part 3: Protection of Business Secrets

page 13



APEX Series Part 6: APEX 20.2 – What's new?

page 18



Smart Sizing of Fast Recovery Area in Cloud Environments

page 27

Editorial

Submit Your Article!

I. Work & Life

Comic: Another Day, Another Daily

At Face Crypto Value

Number of the Month

US CLOUD Act (Part 3):
CLOUD Act and the Protection
of Business Secrets

3

4

5

6

12

13

II. Techs & Nerds

Oracle Application Express (Part 6):
APEX 20.2: What's new? 18

Smart Sizing of Fast Recovery
Area in Cloud Environments 27

New Feature: Bring Your Own IP 32

III. Users & Groups

EMEA Virtual Tour 2020
#BackToBasics
15 Tips, 15 Speakers, 15 Days 36

EOUC Leaders' Virtual Summit 41

Tech Cycle "User Word"
Migrating Mission-Critical
Applications to the Oracle Cloud 44

IV. Past & Future

Events 47

Contact us 48

Legal notice 48

Editorial

Dear ORAWORLD Readers,
2020 what a wonderful year!

Let's see the bright side of this sometimes-gloomy year. The weather was extraordinary sunny and mild, on the professional side, a formidable acceleration of ongoing trends, cloud adoption, collaborative work, remote working, on the personal side, an opportunity to be with our families, a great leap forward for new mobility like bicycle usage helped by so-called "Corona cycling lanes".

Resilience was a 2020 keyword, and thanks to our expertise as IT professionals, our organizations kept on delivering services to our customers/users. Would we have thought 1 year ago that government agencies could switch to remote working (except for E-stonia of course)? That virtual classrooms could be so generalized? That our user groups would organize so many events, some even with a wider attendance!

We certainly faced difficulties, remote admin, security risks increase, but as the Chinese character for crisis, concatenation of Risk and Opportunity shows, there is a bright side: we made a huge step forward, creating new practices, new ways of working. We are building a "new normal". Some requirements will appear for better traceability (think blockchain), more versatile systems will be the norm, office spaces will evolve to a more user-friendly design; a lot of work and changes to support!



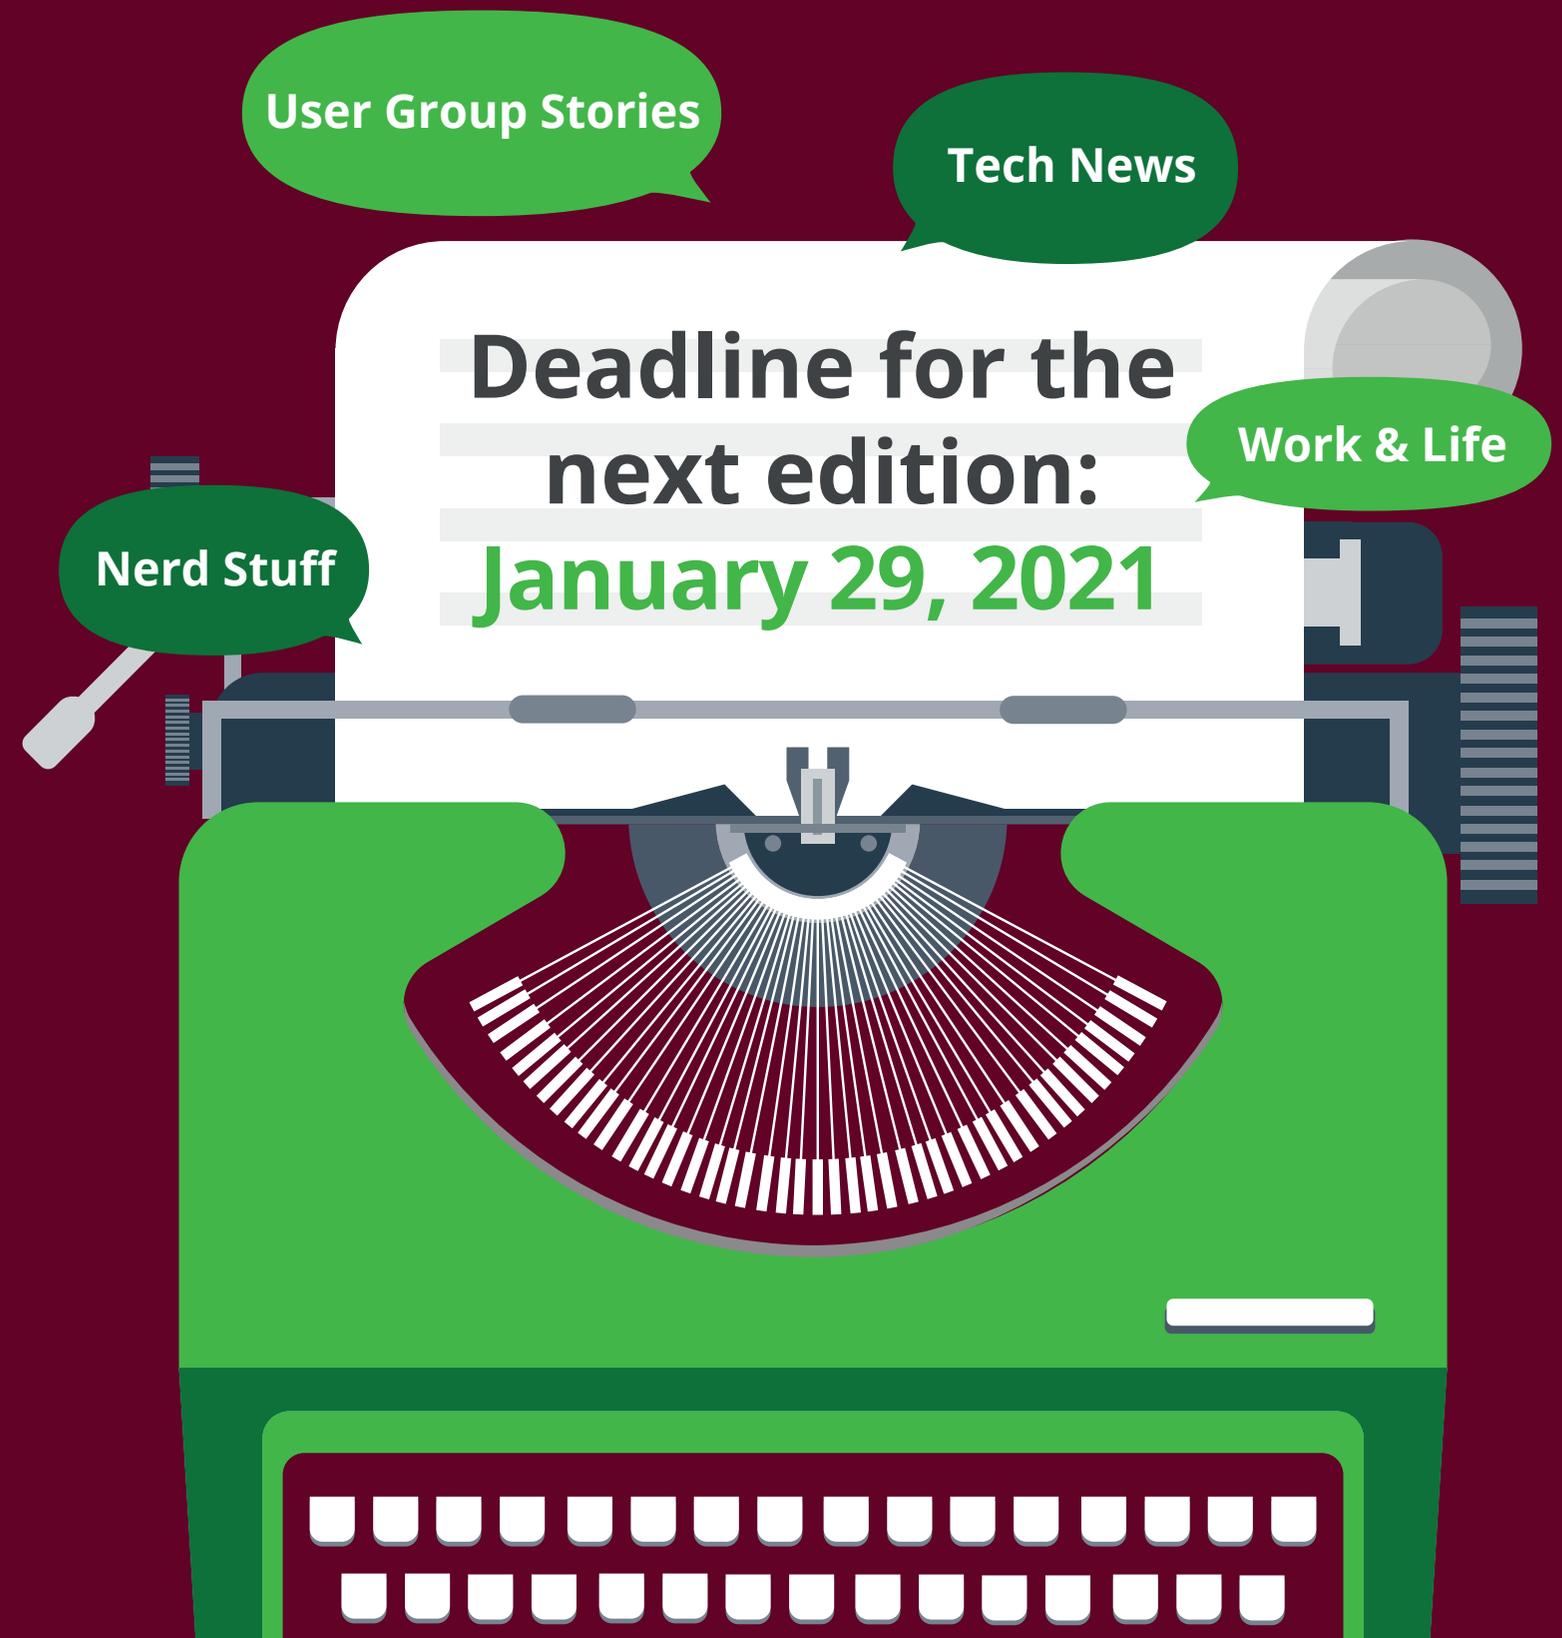
Enjoy reading our last ORAWORLD issue in 2020, and let us wish for a formidable 2021, and keep our creativity and know-how to deliver a new and better world.

Yours,
Jean-Jacques Camps
President AUFO (France)

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.



„Another Day, Another Daily“

www.commitstrip.com



CommitStrip.com

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

Follow us on  @EOUC @ORAWORLD_Mag #ORAWORLD

 @emeaoracleusergroups @ORAWORLDMag

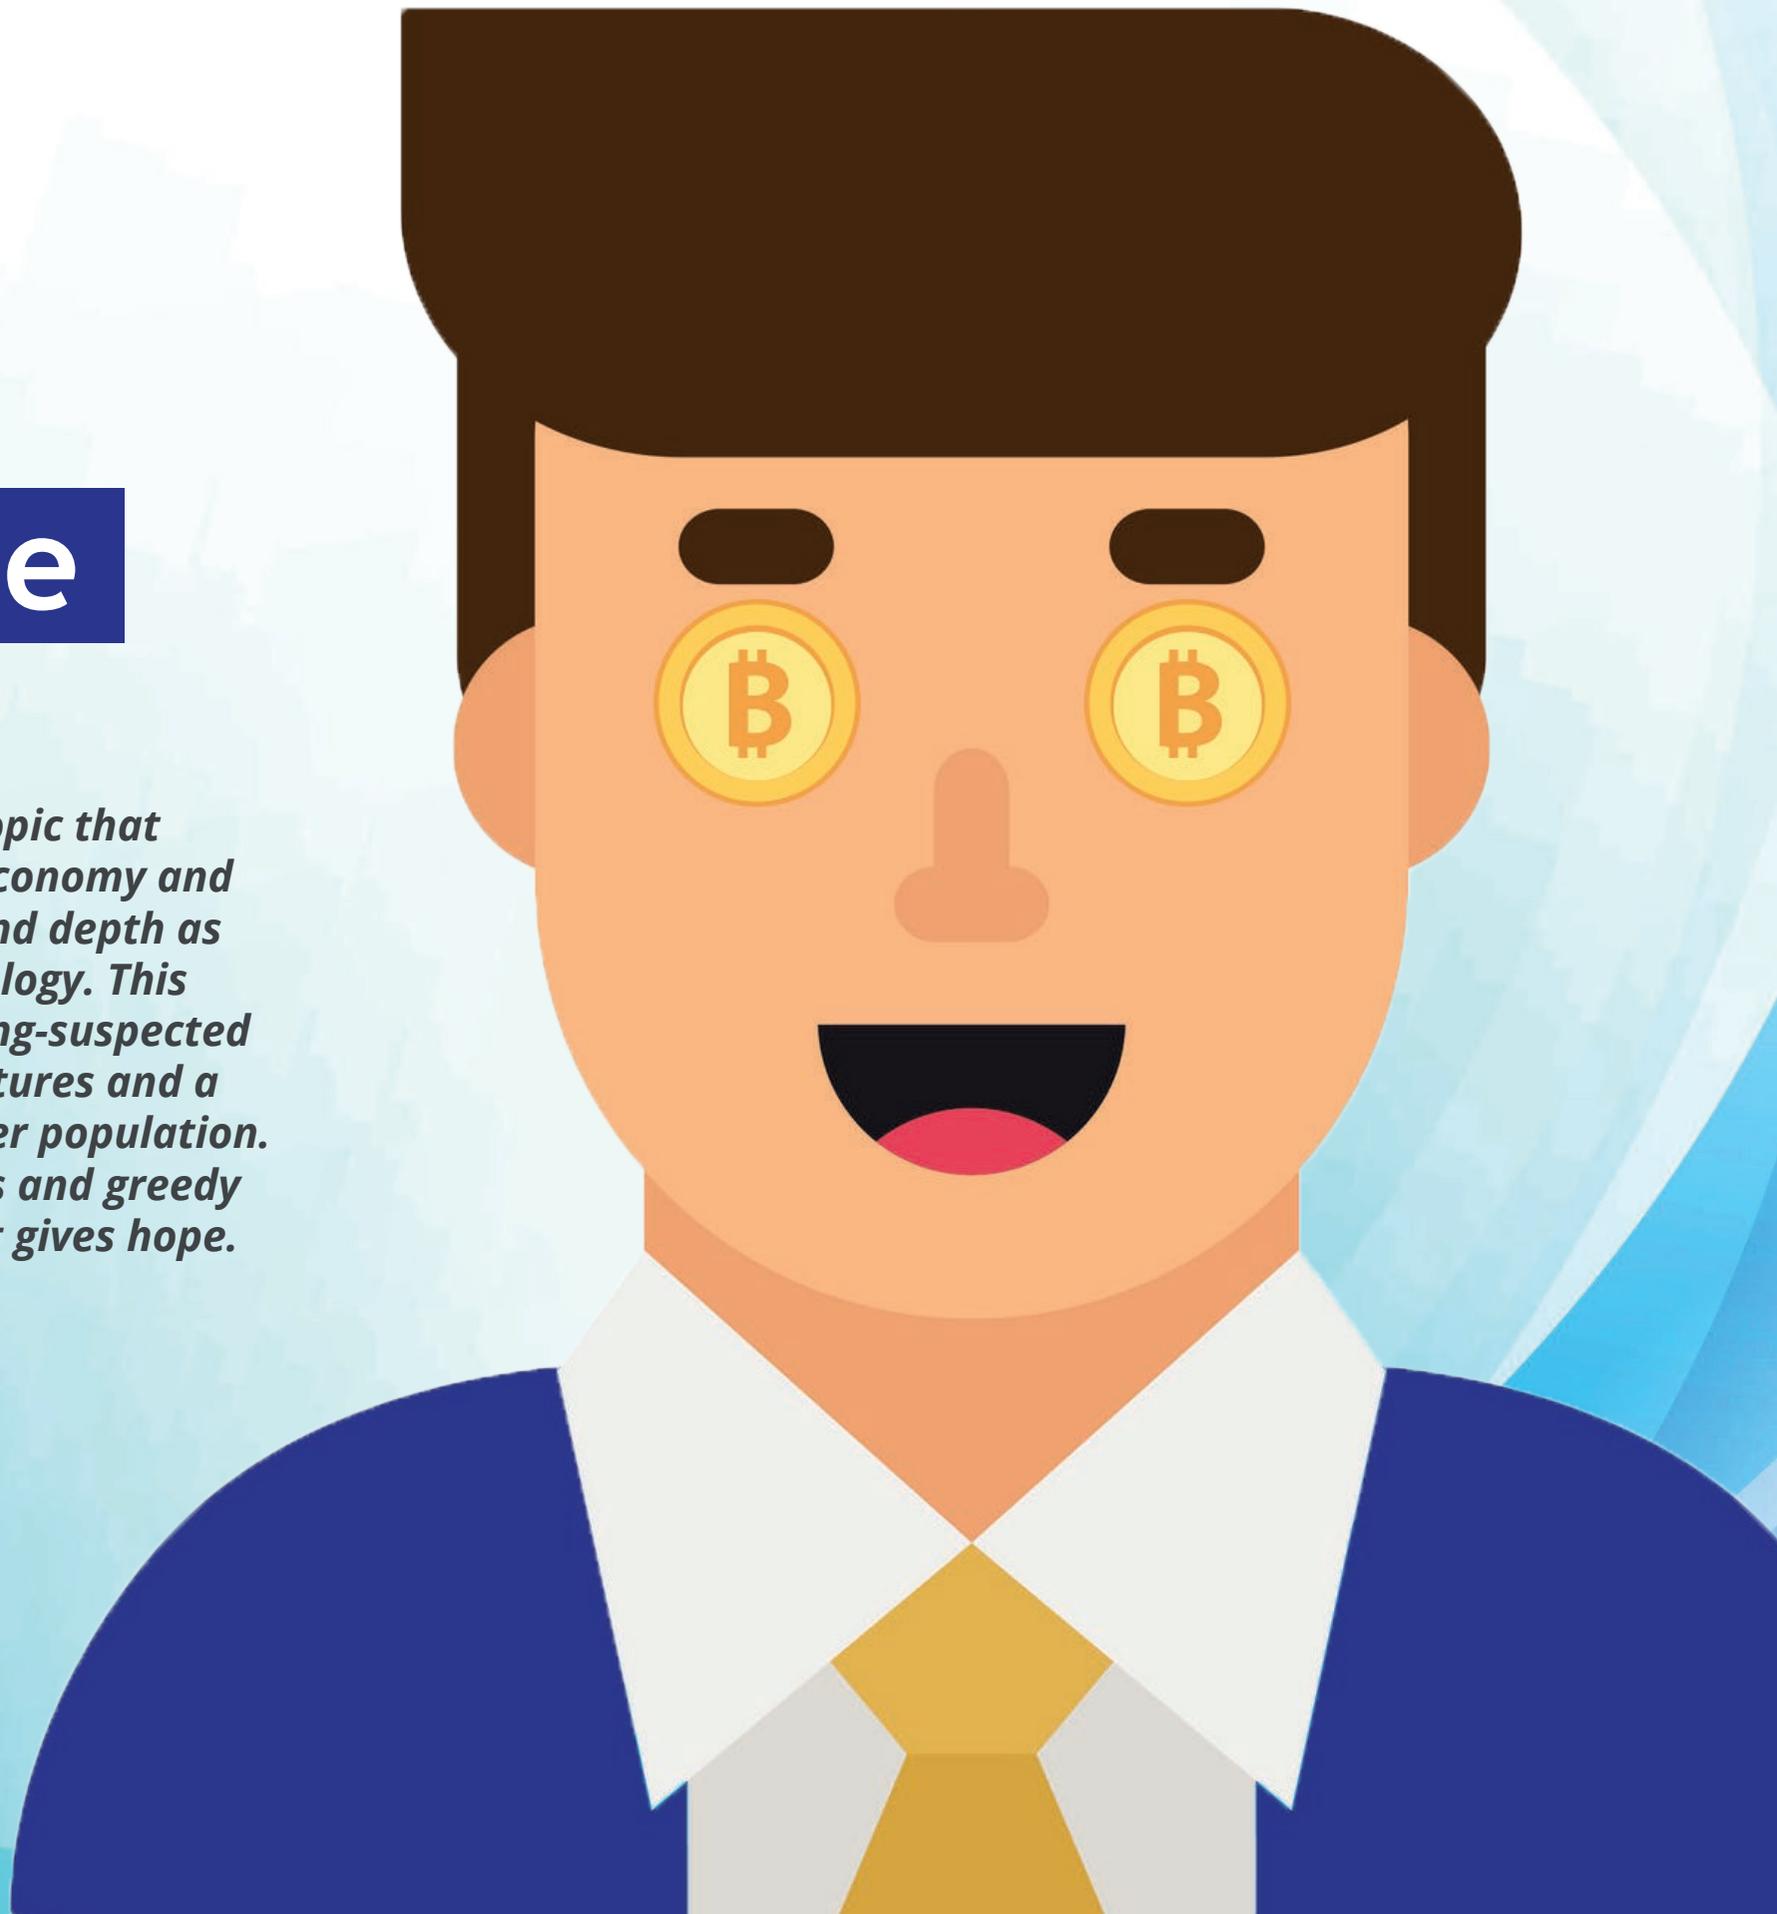


www.oraworld.org

At Face

Nikita Fuchs
Crypto Value

Before Corona, there was hardly a topic that penetrated such large parts of the economy and various social groups at the speed and depth as different forms of Blockchain technology. This relentlessly reveals to developers long-suspected reliability deficits in corporate structures and a financial illiteracy among the broader population. However, the tango with the clueless and greedy reveals a social glimmer of light that gives hope.



“If you’d bought those Bitcoins back then. For one Euro! Then you would be a multimillionaire now, these things are now worth over € 10,000, the piece!” he says and leans back in his seat, visibly proud, having conjured a mixture of confusion and amazement on the faces of his counterpart and the rest of the seatmates in his 4 seater seat. This in itself – the counterparts were three women in their mid-50s – is still nothing special when it comes to topics related to the Internet in my home country Germany: Everyone with a specialist IT profession knows this from the inquiries made by relatives at annual family get-togethers, to please explain what exactly you are doing again. As a developer for blockchain applications, I had my share.

What was actually amazing about this situation, which has happened over the years one way or another in the most unpredictable places with people who are most unsuspecting for this topic, is the person speaking: It was not a thirteen-year-old budding Internet crack of grandchildren. Rather, a man in his late 50s, more of a type where system administrators would bet their lives without hesitation on the existence of an Internet Explorer full of toolbars.

I should have put the naivety of my subconsciously prevailing attitude at the beginning of the blockchain hype out of my mind, that people first dealt dryly and rationally with its peculiarities before the enthusiasm for a piece of technology, against the background of the absurd, primeval market and behavioral models, which I and other offspring of the “decision-makers of tomorrow” had been taught years earlier at a private business school. Admittedly, a pinch of the deeper understanding of finances conveyed there would have saved some people from losing a lot of money “investing” in the 243rd Bitcoin clone, where even the original has even fewer uses than diamonds (but is also valued, due to a certain scarcity, hype and crime).



The American late-night comedian John Oliver excellently distilled the situation around Bitcoin & Co: “It’s everything you don’t understand about money combined with everything you don’t understand about computers”.

Therefore, let’s jump back to technology for a moment and summarize again: The first generation of blockchains consists of well-ordered data blocks with signed transactions, the validity of which can be checked by all participants in the network through the openness of the data. So, to make sure that not just any blocks are spread through the network, the creation of a block – the so-called “mining” – is technically linked to criteria that are difficult to achieve, but easy to validate (keyword “target hash difficulty”), which should not play a role here.

The decentralized nature of Blockchain technology means that everyone can “mine” a block in a race with others and receive a reward for this. This reward is public, including all inherited transactions. The Bitcoin protocol also belongs to this generation. The second generation of blockchain protocols can



also send instructions with every transaction changing a state machine for all participants in the network, i.e., execute simple code, so-called “smart contracts”. This includes protocols such as Ethereum and Aeternity, on which, for example, decentralized trading venues or contracts can be represented.

In order to keep the amount of operations to be carried out and accumulating data for the individual participants small (each usually synchronizes a copy of all transaction data), **transaction volumes in Bitcoin are limited to around 7 per second**, in Ethereum to about 6-16 smart contract executions per second – depending on their **complexity**. And this applies to all participants, worldwide. Other protocols are not much better off, which unfortunately makes this technology only attractive for certain niches for the foreseeable future.

I congratulate you, your knowledge of blockchain technology now exceeds that of the PR departments of large (DAX) corporations and their superiors:



- In 2017, Daimler AG trumpeted the “Launch of a 1-year-corporate Schuldschein [sic!; “Schuldschein” being the German word for IOU] with a volume of € 100 million” in the digital media, which “was digitally carried out via blockchain technology” (source: [media.daimler.com](https://www.media.daimler.com)). On which protocol and on the basis of how this process should have been designed remains a mystery to this day, because, despite expert inquiries, those involved still refuse to share the transaction ID to validate their technically vaguely explained process to verify the transactions that are open per se in real blockchains and can therefore be checked.
- The aspect of the openness of the data by design is something Porsche’s CFO Lutz Meschke did not get a full explanation of as well when he said the following to German business newspaper Handelsblatt in the years of **PR-heavy blockchain announcements** on a prototype project with a Berlin start-up: “With the blockchain, the car owner remains in control of the data and can determine who is allowed to use his data”. And in the slow blockchain system, in which each participant inefficiently always verifies all transactions himself, when accessing the car, “the authorization check should also run six times faster than with today’s technology”. This technology, which is limited to a throughput of a few kilobytes per second for technical reasons (cf. “Size”: etherscan.io/block/11365773), “also ensures that the communication between vehicles, which is essential for autonomous driving, is reliable and, above all, secure against hackers”. And because it couldn’t get any smaller: “It is somewhat surprising how long the importance of blockchain for the car was not recognized in its dimensions”. Nothing more followed the prototype; the start-up has now completely banned the term “blockchain” from its website.



- The economical data throughput rate is apparently also enough for DB AG [German railway company Deutsche Bahn] to allow “trains to communicate with each other and with other elements, such as signals, switches, and routes via the blockchain” (source: blockchain.deutschebahn.com). Because: “The blockchain creates the necessary transparency in real time, e.g. which section is occupied by which train at what time”. And: “The previous [revenue] distribution system is [...] non-transparent and inflexible – and again leads to a delay in the distribution of revenue. Based on blockchain technology, [DB’s sub-company] DB Regio has therefore [...] found a fair and transparent solution for the revenue sharing procedure” (source: digitalspirit.dbsystem.de). Experienced IT staff recognizes the recurring pattern of management, as recently with “the cloud”, trying to magic away problems with the latest IT hype, which have cemented themselves in the company’s infrastructure over decades in the form of technical debt due to inability and ignorance. This impression is also involuntarily and cynically underlined at the end of the same PR hymn of praise for blockchain technology: “A classic database requires an administrator or an institution to take care of data management [...]. The [blockchain] system controls itself.” (Ibid).

There is hardly a company with a public focus that has not announced anything about blockchain in the past three years. Many ambitious employees in middle positions will have been amazed at how long they vainly competed for budgets for large-scale future projects on the subject of digitization, while entire departments are set up overnight for the latest technology craze.

According to its own information, Deutsche Bahn had 30 developers at the beginning of 2019 (ibid) to research 20 possible application scenarios for blockchain technology, while passengers were still allowed to try their luck at spending 200 MB of high-speed volume during their ICE train trip through Germany’s infamous dead spots.

But: Is it really the case with every new technological craze? After all, the Internet was still uncharted territory long after the turn of the millennium. Then where did this wave of attention for blockchain come from? Can it perhaps have something to do with the sensational capital gains of Bitcoin and co.? Or the so-called initial coin offerings (ICOs), effective IPOs carried out on blockchains, with which projects raised billions in capital, especially in 2017 and 2018? Both of these have been the tireless topic of every financial magazine in print and TV over the years, and demographically, the Old Boys Club of German decision-making levels would definitely fit the target group.

The buzz was there, and the curiosity and the greed must have been great, otherwise the inquiries from financial to the criminal community from home and abroad cannot be explained. But what exactly is it that makes innovation-resistant digitization dinosaurs rush headlong into a technology that they have not understood? A friend, owner of an advertising agency for corporate communications, once revealed his secret of success to me: “Switch off your rational IT logic and think like the common German Mittelstand-Manfred”.

Measured against the numbers of German medium-sized companies, this thinking may not be the worst, but if you tried succinctly, the result should be something like: “Now everyone is talking about it. And even Bernd recently raved in the Lions Club about how much money his brother-in-law made recently with bitcoins.” So there must be something to this Bitcoin blockchain thing that half a dozen consulting firms are now always happy to tell you something about for good money.

What remains less uncertain is an uneasy feeling as a private person: If even well-known corporations have no scruples about presenting blockchain technology to the public as brazenly as an all-in-one solution for selfish purposes, how often is this the case with other innovations, for which news





platforms and consumers understandably still lack the know-how (or simply desire) to unmask them as obvious flops? Many developers might now think of one of the Potemkin villages he was forced to work on himself when the management had to prove something to the board again.

The way stock prices reacted to “blockchain” (as the case of an ice tea company shows, whose shares have multiplied after being renamed to “Long Blockchain Corp” (source: vanityfair.com), at least for listed companies, § 256 of the German Criminal Code should be checked, which criminalizes the creation of financial advantages under false pretenses and the maintenance of falsehoods. And, unfortunately, the list of the above examples is only intentionally this short, there would potentially be no lack of material for prosecutors.

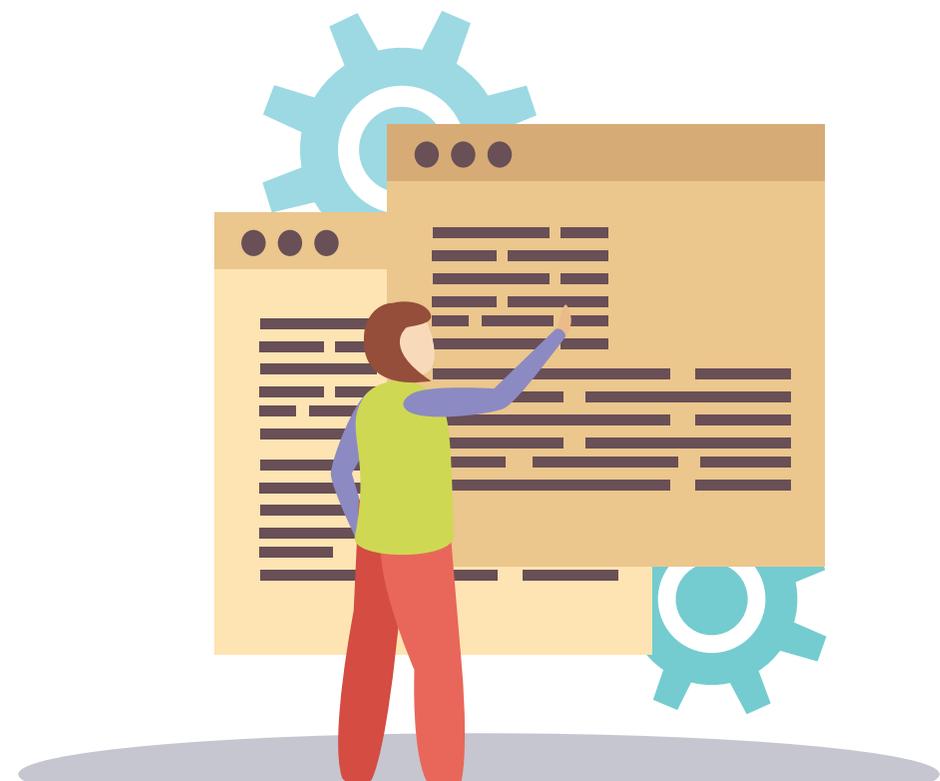
It is tragic and frustrating to watch the technology field I passionately work in having been misused as a job creation niche where one-eyed people tell the blind tall tales for short-term profit. After the hype has flattened out for the first time, many have long since become “experts” in other areas, while others continue to use their contacts in companies and authorities to put senseless projects in the ears of decision-makers: There is no other way to explain, how even technology companies like Samsung jump on the bandwagon of the use case of supply chain tracking by means of blockchain, which is nowhere near feasible in the long term (just think of the data and the complicated private key handling for signing transactions at each loading station).

Or how it can be that the headhunter advertises on the phone for a position with a high salary at the Federal Office for Migration and Refugees to work on a blockchain project to record refugees? The inability to set up a uniform data management solution, which has been known there for years, is apparently to be concealed with the miracle cure blockchain, after decision-makers there were evidently given a similar

picture of the technology as Porsche and DB: In blockchains “data cannot be lost” (in contrast to documents), and they “manage themselves” (which sounds cheaper than database experts and Oracle licenses).

Incidentally, the process of such projects is always the same technologically: If you realize the capacity limits of real blockchains like the Ethereum protocol, you quickly switch to IBM’s Hyperledger: It’s called blockchain, but it is neither decentralized nor transparent, and thus a kind of distributable ACID database with some application logic. Has nothing to do with the original idea of blockchains but is a helper in need.

Distributed ledger technology, as blockchains are often called, has tremendous strengths, provided you know and respect its weaknesses. With transaction costs of (still) dozens of cents up to a few euros, banks cannot be replaced. However, when customers from Israel try to settle a bill without success for a week due to bureaucratic nonsense in our financial system, crypto currencies have come just in time. Other use cases that rarely require cost-intensive writing on the chain but read more often (always free of charge) from their current state are





absolutely attractive: The so-called tokenization, i.e. the breaking down and/or mapping of valuable goods into digital parts, for example, which can be managed decentrally via smart contracts are attractive for many projects. Anyone who owns 20 percent of all tokens (shares) of a legally correct and bindingly tokenized painting can easily transfer these tokens to a wallet address (blockchain account) of a bank even in the Philippines and thus deposit it as security for a loan. The transaction fees of one to three Euros – at the worst – don't hurt.

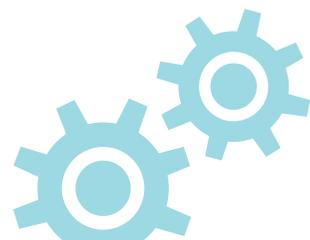
As a consolation for the misuse and misunderstanding of this revolutionary technology, there is, in my opinion, at least a social added value that has nothing to do with technology. Because the low interest rate policy of the central banks turned countless small savers into hobby investors, countless billions in the form of small amounts of a few hundred or thousand euros have flowed into the crypto exchanges. The bursting price bubbles made some people wonder whether the ECB's money printers might not sometimes make sense as a price stabilization mechanism. The interest in the overall topic has made an incredibly broad social spectrum aware on the subject of how our monetary system works per se that Henry Ford said about, not entirely unfounded, "If people understood the monetary system, we would have a revolution tomorrow morning." In the comment columns of today's Internet it is not uncommon to even see a kind of people reasoning with the vigor of a German TV stock market commentator about the the ECB's monetary policies, whereas you would rather expect to find reality TV on button 1 of that folk's TV remote. And this is a delightful situation.

Either way, blockchain technology has the potential to move societies forward and bring the world closer together. Even if it's just for boastful conversations on the ICE train.



About Nikita Fuchs

Nikita is a graduate economist and business informatics specialist who works as an independent consultant and developer for various projects. As a developer of smart contracts, he has been helping projects raise capital through tokenization of their products since 2017. He is currently active at the digital fashion start-up LUKSO, which makes digital fashion tradable using blockchain technology, as well as the blockchain protocol Aeternity, which aims to overcome the current technical limitations of blockchains.





NUMBER OF THE MONTH:

18.5 Million

In the third quarter of 2020, **18.5 million Bitcoins were circulating in the world**. In 2008, Bitcoin was invented by an unknown person or group calling itself "Satoshi Nakamoto" and then released one year later as an open-source software.

Bitcoin was the first cryptocurrency and has become the de facto standard, inspiring a constantly growing number of other cryptocurrencies. Some of the most well-known

other virtual currencies on the market are Tether, Ethereum, Litecoin, Ripple or Libra.

Maybe you are now thinking: "Well, of course I've heard about Bitcoin but what really is a cryptocurrency and how does it work?" Then click the play button above to view the 190-seconds video by The Guardian showing the idea and concepts behind cryptocurrencies and, in particular, Bitcoin.



US CLOUD Act (Part 3):
CLOUD Act and the Protection
Dr. Michael Isler
of Business Secrets

Introduction

The Clarifying Lawful Overseas Use of Data Act (hereafter “CLOUD Act”) is a piece of federal legislation of the United States that aims at enabling U.S. law enforcement agencies to access electronic records stored with “a provider of electronic communication service or remote computing service”. The CLOUD Act applies irrespective if such data are stored in the USA or abroad. All providers that are subject to U.S. jurisdiction due to a certain (even loose) nexus to the United States must abide by the disclosure obligation under the CLOUD Act (for more detailed information, see [Part 1 of this article series](#)).

The CLOUD Act targets “wire or electronic communication and any record or other information pertaining to a customer or subscriber within such provider’s possession, custody, or control.” Apparently, whenever such information relates to an individual, it qualifies as “personal data” pursuant to the General Data Protection Regulation (“GDPR”) and other privacy legislation. The transfer of personal data to the United States on the basis of an official access request is likely to expose the provider to a dilemma to comply with the GDPR and violate the CLOUD Act, or vice versa (for a more detailed analysis, see [Part 2 of this article series](#)).

However, the implications of the CLOUD Act go even further. Data that is susceptible to a disclosure obligation may be subject to protection as professional secret, trade secret, or confidential information by virtue of a non-disclosure agreement (hereafter “business secrets”). For instance, certain information is protected under special professional or sectorial statutory secrecy obligations, such

as patient data processed by a healthcare provider (medical secrecy), information entrusted to an attorney (attorney-client privilege), or voice and data transmissions carried by a telecom provider (telecommunications secrecy). Also, government officials are generally bound to keep information secret that was entrusted on or collected by them (official secrecy). Common to all these examples is that secret information of others is kept by a custodian. In contrast, another example of trade secrets are intangible assets, such as information about innovative technologies, which the owner himself has an interest in keeping secret. These should not be allowed to enter the public domain, as this would destroy their value for the owner of the assets.

A structured view of the different kinds of business secrets is provided in Image 1. This third part of our article series about the CLOUD Act deals with the handling of business secrets.



Image 1: Different kinds of business secrets.



When does the “CLOUD Act risk” materialize?

In the normal course of events, a U.S. law enforcement agency interested in accessing information that constitutes business secrets would first approach the owner of the information (hereafter the “principal”). Should the principal decline to comply with this request for legal reasons, the authority may take recourse to a known third party who stores the information on behalf of the principal, e.g. an outsourcing or cloud service provider (hereafter the “service provider”). If such service provider is qualified as a “US company” in the sense of the CLOUD Act, the probability that the access request will be successful increases. The ensuing risk needs to be validated, both with respect to business secrets of others entrusted to the principal and own business secrets of the principal.

Protection of business secrets of others

In general, the violation of statutory secrecy obligations constitutes a criminal offence. The criminal liability is not only imposed on the persons acting on behalf of the principal, who are subject to the statutory secrecy obligation in the first place, but also on those working for the service provider on instruction of the principal.

The normal procedure in case of a lawful access request would be that it is enforced by way of judicial aid with the help of the authorities in the principal’s home country. Alternatively, the person whose protected data is sought may consent to the access request, which is however rarely the case in a criminal investigation. Accordingly, the service provider confronted with an access request by a U.S. law enforcement agency under the CLOUD Act has a choice between the devil and the deep blue sea: Either he or she will commit a crime for secrecy violation – or disobey the order of the U.S. authority. To make things worse, in some

jurisdictions it is forbidden to release information to foreign authorities without following the path of the official judicial aid system.

In case the access request is directed to information of a third party protected under a mere contractual confidentiality obligation, the dilemma for the service provider is less virulent, because it may be excused from liability when forced to comply with a binding order under the CLOUD Act. It has also become standard in non-disclosure agreements that confidential information may be disclosed if so mandated by a legal obligation.

Thorough risk assessment required

Any person entrusted with the custody of business secrets of third parties needs to perform a vendor risk assessment of the service provider to whom it wishes to transfer data.





Such transfer of data to a trusted third party is a permitted disclosure and does not constitute a violation of the secrecy obligation. However, if the service provider later is ordered to disclose secret information in its custody due to an access request under the CLOUD Act, the principal will have to prove that it has diligently assessed the probability and impact of the risks resulting from the engagement of that third-party provider. In fact, the exposure of the provider under the CLOUD Act is one of several aspects that needs to be taken into consideration during a vendor risk assessment. Failure to properly verify the legal implications of third-party vendor engagement may indeed yield criminal sanctions.

In this respect, it is fair to state that the relevance of the CLOUD Act has so far rather been overestimated when it comes to its practical impact. According to [Microsoft's law enforcement requests report](#), out of approximately 5,500 access requests in the first half year of 2020, 115 were extraterritorial, i.e. related to content stored outside of the territory of the United States. Among these, 113 affected consumer services such as Hotmail/Outlook.com, OneDrive, Xbox Live, and Skype. Only two access requests (0.03%) were related to enterprise content. In earlier periods, the figures are similar.

This should not downplay the concerns raised around the CLOUD Act, but put them into right proportion. The potential applicability of the CLOUD Act does not rule out the engagement of a U.S. affiliated service provider per se, but it is one of the factors that must be considered. In this respect, the approach that the service provider takes towards access requests by foreign authorities must be laid out in a transparent manner and underpinned by a contractual obligation to exhaust legal means against any disclosure obligation of so directed by the principal. Encryption key management and organizational procedures for granting plain text data access to provider personnel in case of support events are also important factors. Further, the principal should determine how likely it is that the information entrusted to the service provider could become the object of a U.S. criminal prosecution.

Protection of own business secrets

When it comes to the protection of own business secrets, such as trade secrets or other valuable know-how of the principal, the disclosure resulting from an access request by a U.S. law enforcement authority under the CLOUD Act will not infringe upon the right of others, but may impair the legal protection and value of the business secrets. According to the

EU trade secrets directive, which had to be transposed into national law of the EU member states by June 2018, a trade secret is only legally protected if it is not publicly known and reasonably shielded against inadvertent disclosure.

Hence, each holder of a trade secret needs to assess if the measures to keep the information secret are sufficient. Whilst we have seen that access requests under the CLOUD Act are not very frequent, the risk must still be weighed in the light of the overall efforts of the secret holder to keep the valuable information confidential. Hence, placing trade secrets in the hands of providers that are exposed to the long arm of U.S. law enforcement may be delicate, although the abstract risk of a successful access request will surely not deprive information of its status as business secret.

Summary

The CLOUD Act does not only cause concerns when it comes to compliance with data protection legislation, but also in relation to business secrets entrusted to a third-party provider who may be compelled to disclose data to U.S. law enforcement agencies. Such business secrets may be privileged information by virtue of a statutory obligation of secrecy, information subject to a contractual obligation of confidentiality, or trade secrets. Whilst the transfer of business secrets to a third-party provider with a likely exposure to the CLOUD Act is not prohibited, it needs to be scrutinized as part of a vendor risk assessment exercise. In this respect, the CLOUD Act is one risk factor among others. In terms of practical relevance, the probability of the risk of inadvertent disclosure is generally overestimated, but given the potential impact it must be taken seriously.



About Dr. Michael Isler

Dr. Michael Isler is managing partner at Walder Wyss Ltd., a business law firm based in Zurich, Switzerland. He advises in complex outsourcing, technology transfer and platform projects from the conceptual and negotiation phase to dispute settlement. Moreover, he regularly publishes and lectures in his practice areas and takes an active role in several professional organizations.

Carsten Czarski

Oracle Application Express (Part 6): APEX 20.2: What's new?



Introduction

The second APEX release in 2020, **APEX 20.2**, has been released on October 21, 2020. As every release, it comes with a number of new features, which this article will provide an overview on: Besides the marquee features, like the new Cards Region, Automations or the new printing capabilities, there is a number of small enhancements which make APEX applications richer and developer lives easier.

New Cards Region

Cards is a new region type in APEX 20.2, which allows to have a declarative “Cards” display for data from a SQL Query or a REST Data Source (**Figure 1**). The card layout is a very common visualization type for web applications – it provides much more flexibility than classic tabular layouts. Unlike the already existing “Cards” Classic Report template, the new Cards region does not require a specially crafted SQL query.

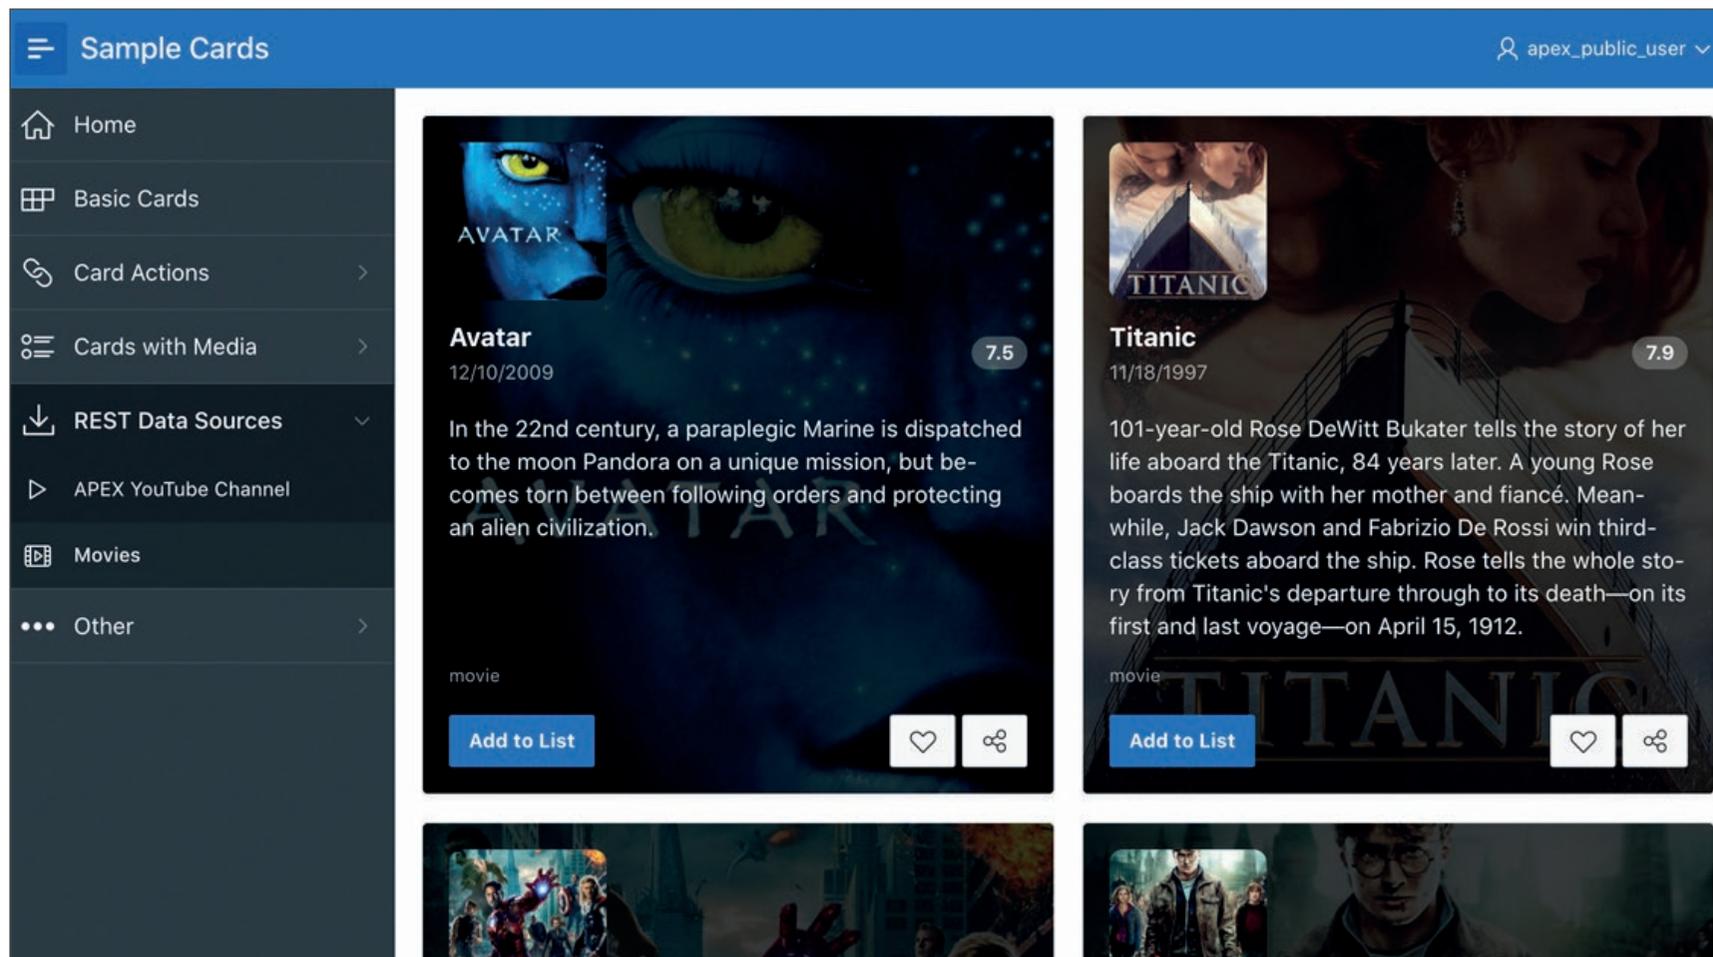


Figure 1: The new Cards Region in APEX 20.2

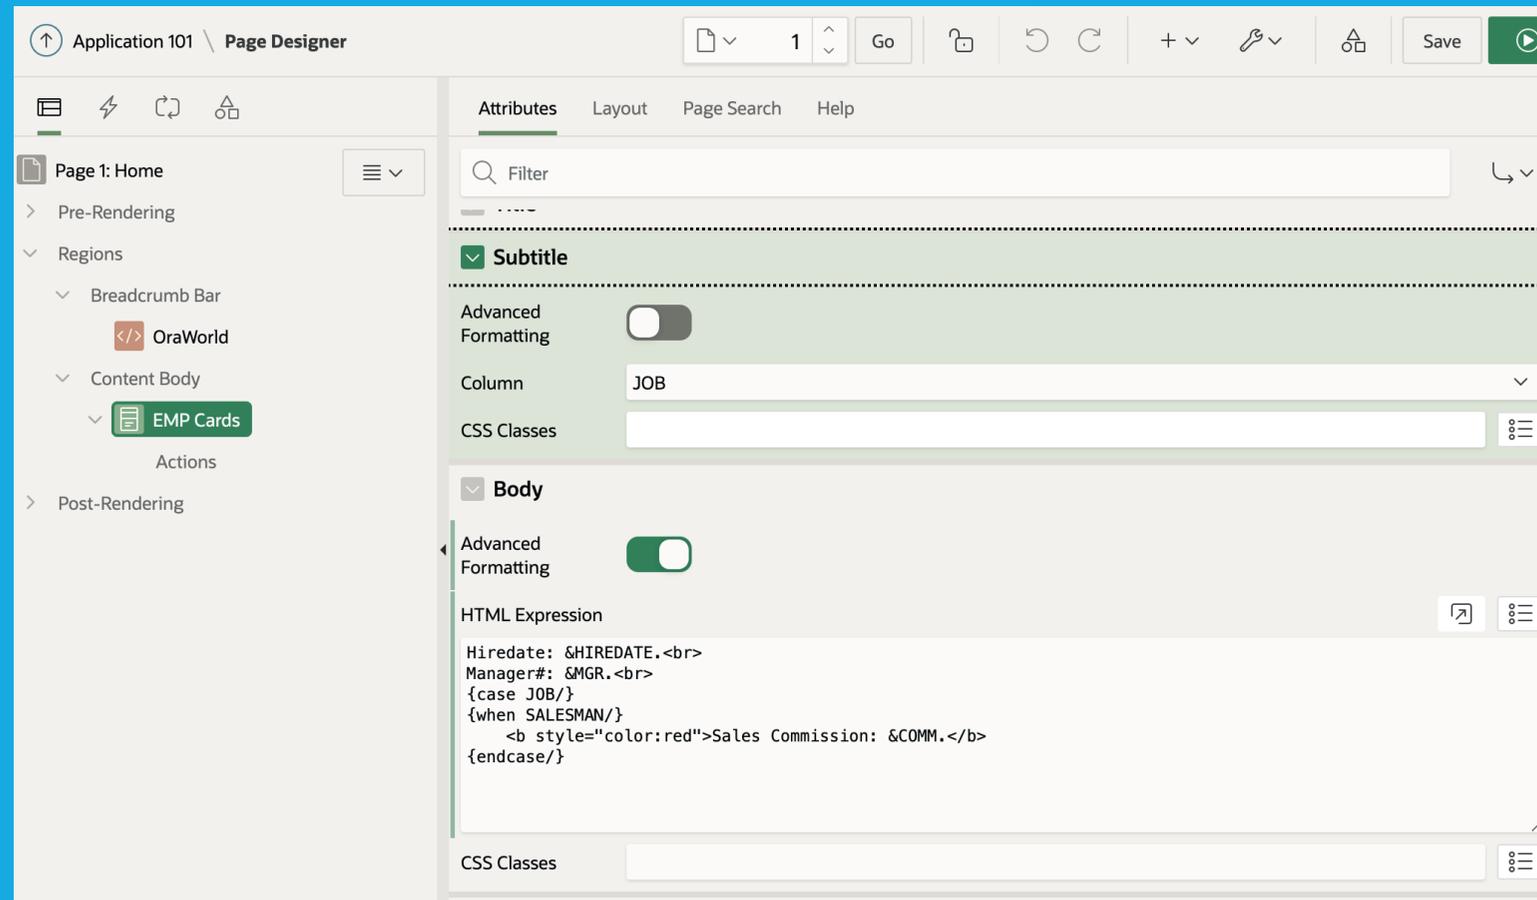


Figure 2: Using HTML Expressions with Template directives for advanced formatting

After defining the data source, which can be a table, SQL query or an external REST Data Source, the developer configures the different attributes of a Cards region: Data source columns can be directly mapped to card attributes like **Title**, **Subtitle** or others. One or many buttons can be added to a card as **Card Actions**. The **Advanced Formatting** option allows to specify simple HTML expressions or even to use *advanced template syntax* (**Figure 2**).

The advanced template syntax allows much better separation between display logic and the actual data: It's no longer required to generate HTML snippets in the SQL query, in order to conditionally display a HTML block or not – the template directives take care of that. Detailed documentation on the new syntax is available in the APEX JavaScript API Reference, and here within the **apex.util.applyTemplates** function.

In APEX 20.2, the advanced template syntax can be used for the Cards region and for Interactive Grid. Support in other region types will be added in future releases.

Automations

Executing business logic, based on end user actions (for instance, clicking a button or saving a row), is the normal case within an APEX application, and developers are doing that every day. However, pretty often the requirement is to perform a query and execute some actions – without anybody clicking anything: It just has to run on a defined schedule.

Of course, that can be done manually, with PL/SQL code and with configuring the *Oracle Database Scheduler*: just build a stored procedure to perform a query and to execute some actions. Then use the `DBMS_SCHEDULER` package to configure the job which runs on regular schedule.

However, that requires a lot of low-level and technical code. The new **Automations** feature, which is configured in Shared Components, makes this task much easier. An Automation consists of:

- A data source, which can be a table, SQL Query or a REST Data Source.
- An execution schedule definition.
- A set of *Actions* which are executed for each result row, or once.

All of this is highly configurable: each Action can be conditional, to only execute it in certain situations. In addition to using plain PL/SQL code, Automations also support using *Process Type Plug-Ins*. The automation will run on the defined schedule (**Figure 3**), but developers can also invoke it “on demand”, using the **APEX_AUTOMATION** package.

Within the action code, column data from the automation query is available as *bind variables* (:EMPNO, :ENAME and so on). Action execution is logged (**Figure 4**), so that developers can verify which actions have been executed and whether error messages were raised or not.

The screenshot shows the configuration for an automation. In the 'Settings' section, the 'Type' is 'Scheduled', the 'Schedule Expression' is 'FREQ=HOURLY; INTERVAL=6; BYMINUTE=0', the 'Schedule Status' is 'Disabled', and 'Actions Initiated On' is 'Query'. In the 'Source' section, the 'Data Source' is 'Local Database' and the 'Source Type' is 'SQL Query'. The SQL statement entered is 'select * from emp'.

Figure 3: Query the EMP table every 6 hours, and execute Actions

The screenshot shows the 'Execution Log' for an automation named 'EMP Automation'. The log shows that the automation was aborted on 11/26/2020 at 04:08:07 PM. The log table contains five rows of execution details, including timestamps, messages, message types, and primary key values.

Automation	EMP Automation		
Execution Start Timestamp	11/26/2020 04:08:07 PM		
Status	Aborted		
Timestamp	Message	Message Type	Primary Key Value
11/26/2020 04:08:07 PM	SALESMAN row WARD processed.	Information	
11/26/2020 04:08:07 PM	SALESMAN row TURNER processed.	Information	
11/26/2020 04:08:07 PM	SALESMAN row MARTIN processed.	Information	
11/26/2020 04:08:07 PM	SALESMAN row ALLEN processed.	Information	
11/26/2020 04:08:07 PM	ORA-20000: Cannot process row for empno 7839	Error	

Figure 4: Review the Automation Execution Log

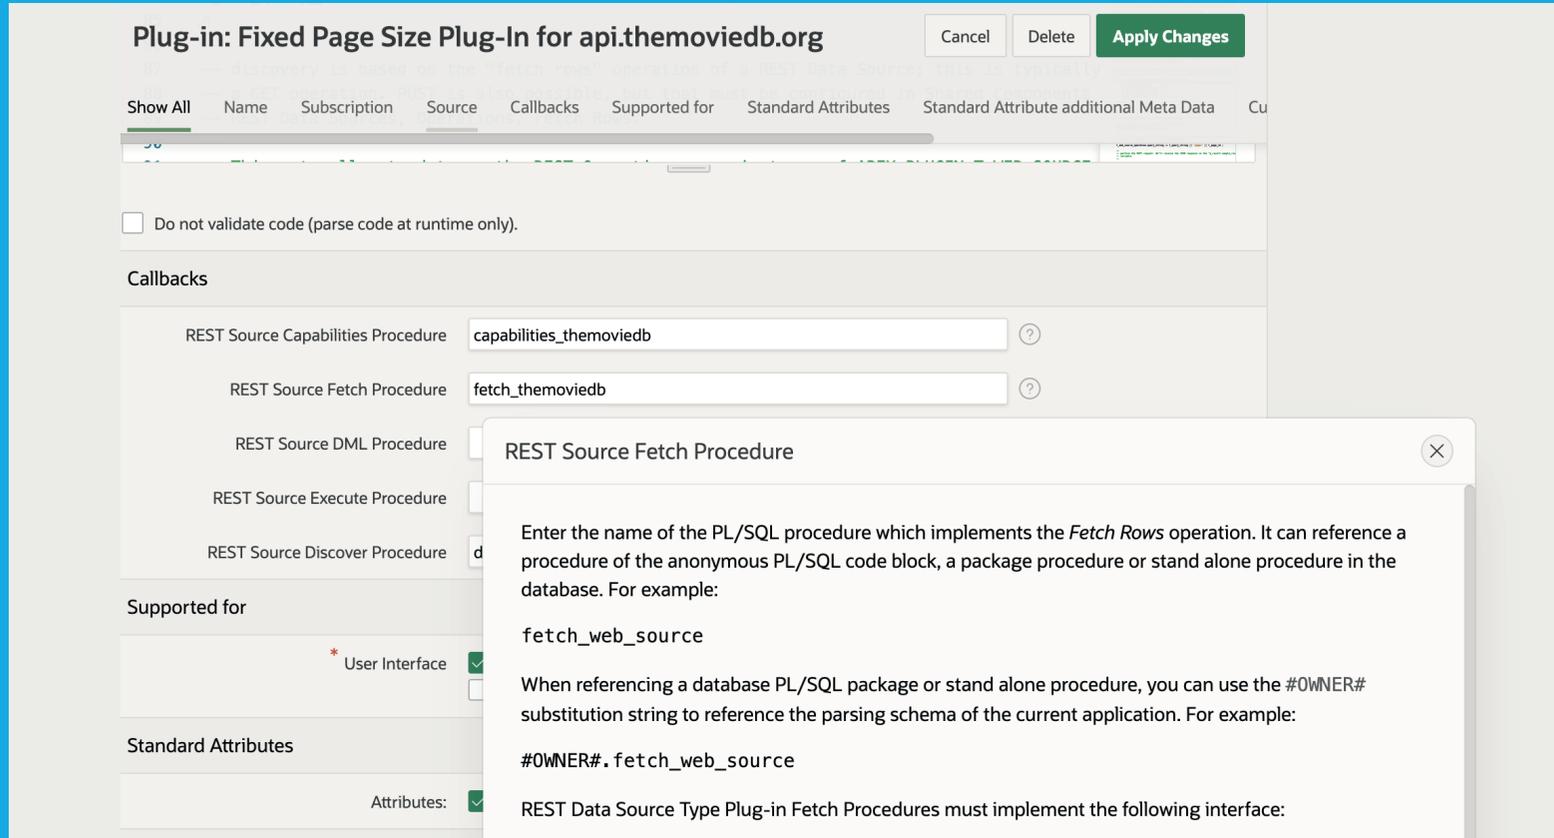


Figure 5: REST Data Source Plug-Ins

With the APEX_AUTOMATION package, developers can amend these logs with own messages, using **Information**, **Warning** or **Error** message types.

REST Data Sources

In APEX 20.2, Web Source Modules have been renamed to **REST Data Sources**. Besides this simple name change, there are a lot of new functionality.

REST Data Source Plug-Ins (Figure 5) enable APEX to fully leverage REST API features, like *result pagination or server-side filtering*. Since REST is an architecture, and not a standard, there is a variety of dialects on how to invoke a REST service and only request a specific result page. It also varies from REST service to REST service, how *filters* have to be passed to the server.

With APEX 20.2., developers can implement support for specific REST APIs as a REST Data Source Plug-In. The Plug-In contains

PL/SQL code, which is called by the APEX engine when APEX interacts with the REST API. APEX passes in all required context information; among others this is:

- rows requested by the APEX component
- whether all rows are needed, when it, for instance, computes an aggregation
- component filters, for instance Interactive Report filters
- new values for DML operations (POST, PUT, DELETE)

The plug-in implementation then crafts its specific HTTP request, passing the required HTTP headers, URL parameters and request body, and returns the received response back to the APEX engine. APEX will proceed with JSON or XML parsing and with providing rows and columns to the requesting component, which can be a report, a chart or another component.

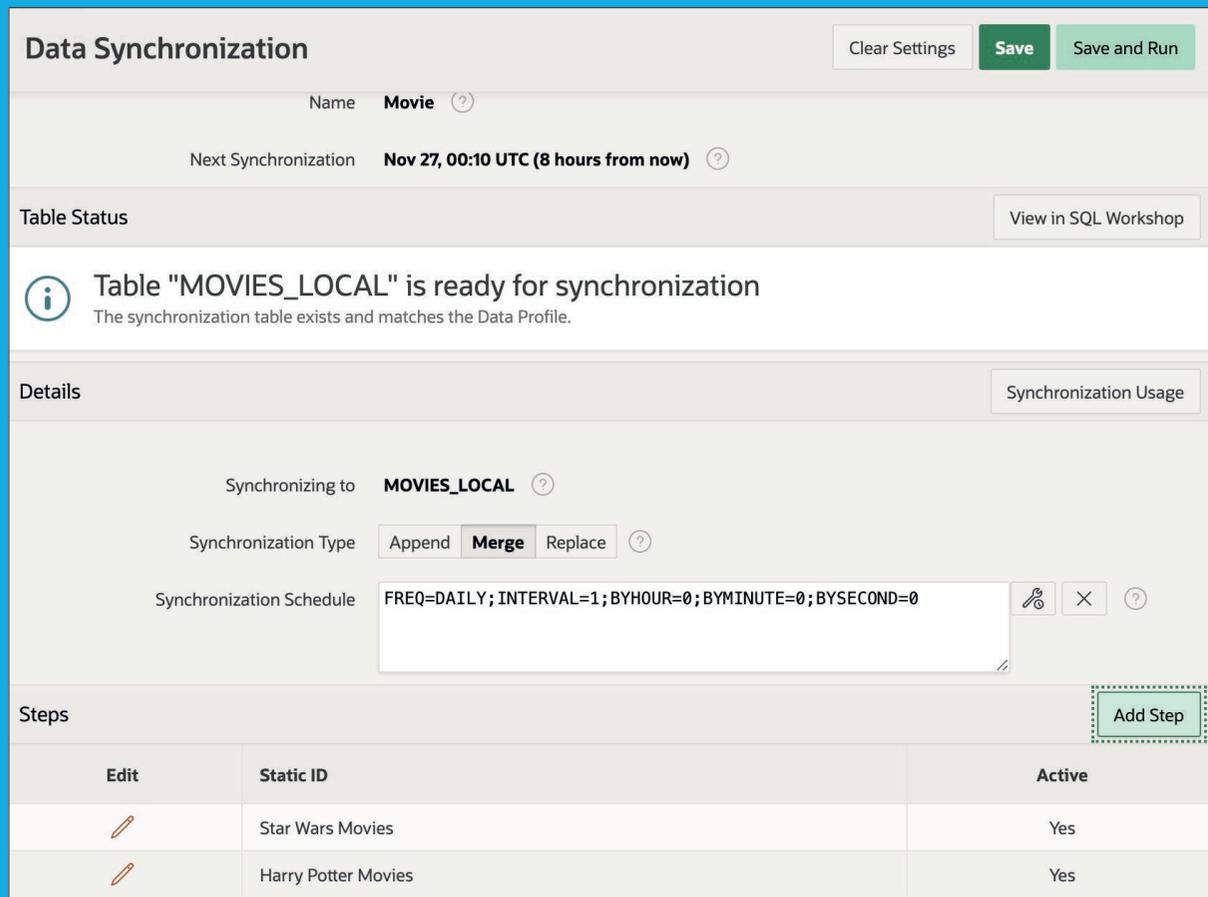


Figure 6: Daily synchronization from a REST API to the MOVIES_LOCAL table

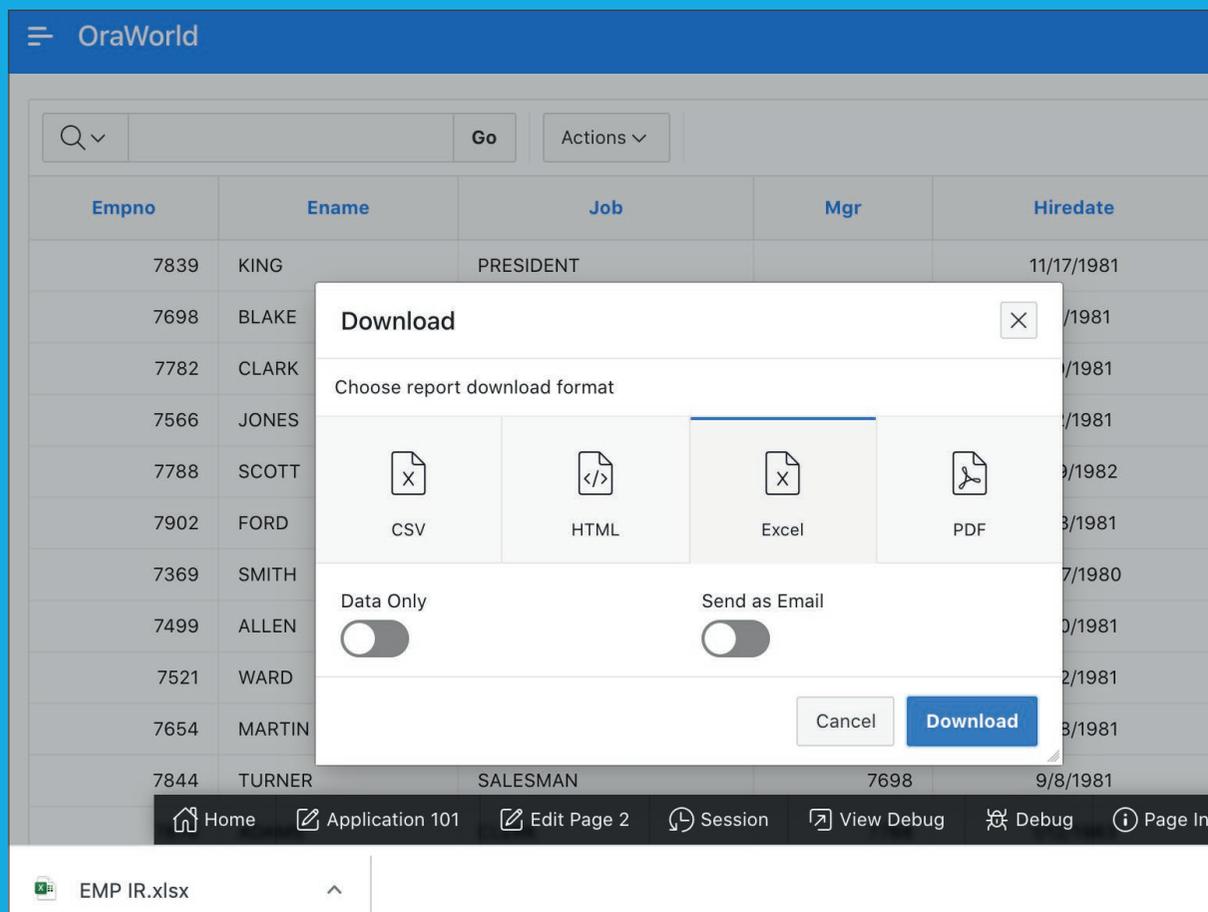


Figure 7: Downloading an Interactive Report as XLSX with APEX 20.2

REST Synchronizations make it super-easy to copy data from a REST API to a local table: All is done declaratively, and no PL/SQL coding is required. Synchronizations can be defined for all REST Data Sources, which have a **Fetch Rows** database operation defined.

This is useful when data, returned by the REST service, does not change too often, but the APEX pages are frequently accessed. Instead of reaching out to the REST API for every page view, the page only accesses the local table. Based on a regular schedule, the local table is refreshed from the REST API (**Figure 6**) by either appending, merging or replacing rows.

For **Web Credentials**, APEX 20.2 introduces the **URL Query String** and **HTTP Header** types. This allows developers to use the secure and encrypted credential storage also for consumer or API keys, which are often passed as part of the URL or as a HTTP header. APEX makes sure that such sensitive parts are not written to debug or web service logs.

To protect the credential from accidentally being sent to the wrong REST API, developers can provide a **Valid URL Pattern**: APEX will only use the Web Credential for URLs starting with the given pattern - if that is not the case, an error will be raised. To change the Valid URL Pattern, the secret part of the Web Credential must be entered again.

Native Printing

For Interactive Grid, native printing capabilities were already introduced with APEX 20.1, released back in April 2020. When downloading an Interactive Grid in PDF format, the file is generated by the APEX engine itself, without reaching out to an external printing engine.

APEX 20.2 extends this native printing capability not only to classic and interactive reports, it also introduces native XLSX (Spreadsheet) generation (**Figure 7**).

All data export and printing functionality is available to the PL/SQL developer in the **APEX_DATA_EXPORT** package. That allows developers to generate CSV, XLSX or PDF files from custom SQL queries, as illustrated in *Listing 1*.

Faceted Search

For a **Faceted Search** page, APEX is now able to display item counts as a **Chart (Figure 8)**, without any additional code. For new Faceted Search pages, charts are enabled by default; for existing applications, the feature is enabled by toggling a switch in Page Designer.

```

declare
    l_context apex_exec.t_context;
    l_export apex_data_export.t_export;
begin
    l_context := apex_exec.open_query_context(
        p_location => apex_exec.c_location_local_db,
        p_sql_query => 'select * from emp' );

    l_export := apex_data_export.export (
        p_context => l_context,
        p_format => apex_data_export.c_format_csv,
        p_file_name => 'employees' );

    apex_data_export.download(
        p_export => l_export );
end;

```

Listing 1: Convert query results to CSV using APEX_DATA_EXPORT

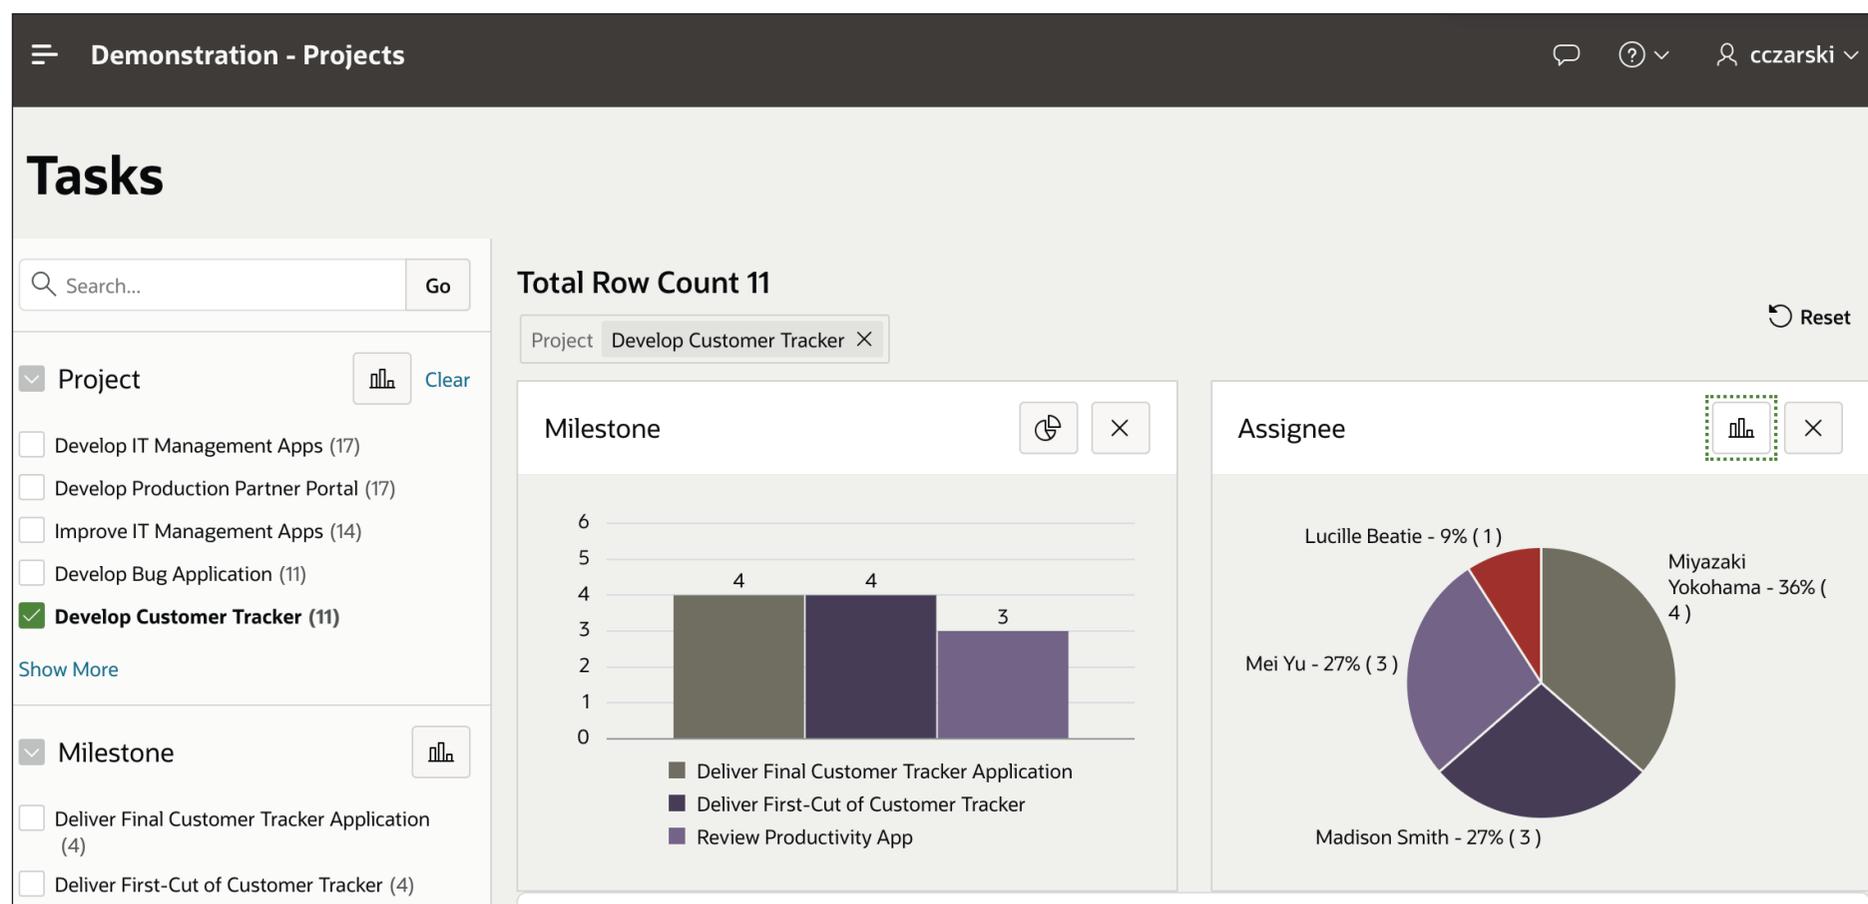


Figure 8: Facet item counts, displayed as a chart



New facet types have been introduced as well. For the **Input Field** facet, the developer specifies the filter operator and the column to apply the filter on. End users simply type in values and APEX applies the filter accordingly. This is useful for facets filtering with, for instance, less-than or greater-than comparisons.

Facet Groups allow to better deal with multiple flag columns containing only values like "Y" or "N". Faceted Search allows to create simple **Checkbox** facets, and to combine those into a Facet Group. End users experience the facet group like a normal checkbox facet.

Other New Features

APEX 20.1 introduced "Redwood", Oracle's look and feel, to the Application Builder. Starting with 20.2, the **Redwood Light** theme style is available for APEX applications (Figure 9). For new applications, the new style is available immediately; existing applications need to first upgrade Universal Theme (in Shared Components, Themes).

The **Embedded Code** utility provides easy access to all PL/SQL, SQL and JavaScript code within the application. A faceted search page allows to navigate; developers can review the code directly or export it as a ZIP file (Figure 10).

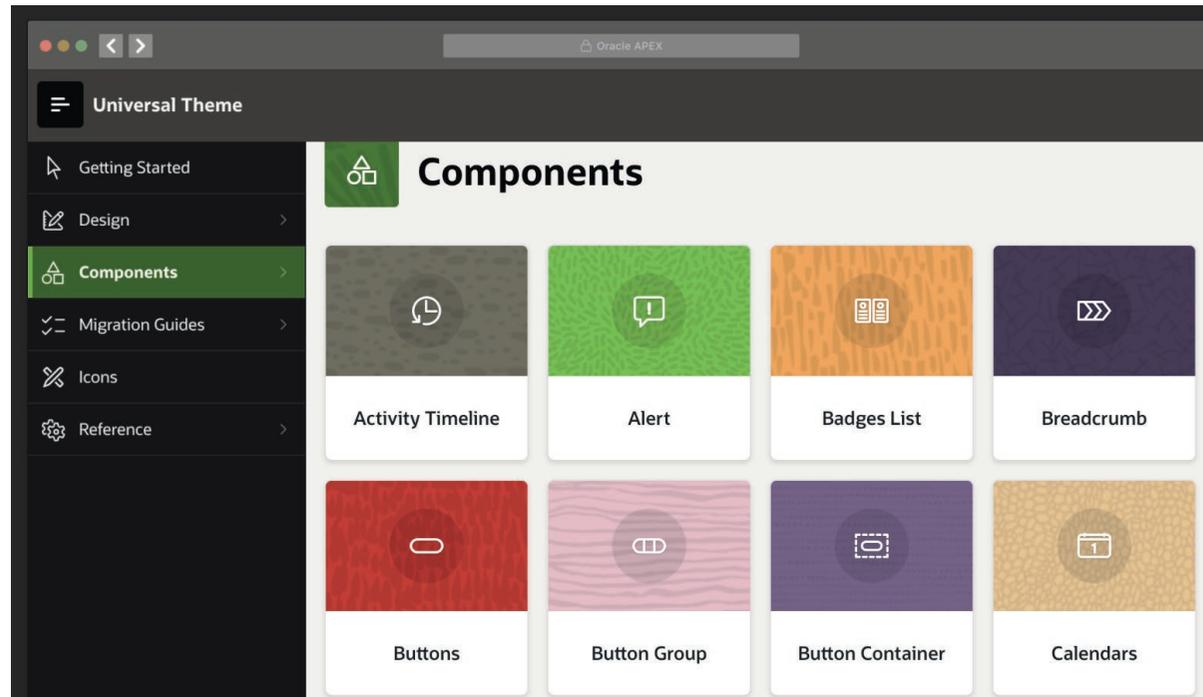


Figure 9: Redwood Light Theme Style for APEX applications



Figure 10: Embedded Code: All SQL, PLSQL or JavaScript snippets within the application



All code editors within the Application Builder or within SQL Workshop have been changed to use the popular **Monaco Editor** library, which provides in-context code completion, syntax highlighting, and better accessibility.

The new **Single Checkbox** item type offers an alternative to the switch type for Boolean columns. This also works in Interactive Grid, even when not in edit mode. The previous Checkbox type has been renamed to **Checkbox Group**.

The **File Browse** item type has been improved as well: It now supports rendering as a *drop zone*, supporting drag & drop of a file to be uploaded. The Rich Text Editor item type is now based on the **CKEditor 5** library.

Finally, the plain **Text Field** item type has a new **Text Case** setting to optionally transform the user-entered text to upper or lower case. The **Trim Spaces** and **Text Case** settings are applied on the client as well as the server.

More Information

- Cards Sample Application:
http://apex.oracle.com/pls/apex/apex_pm/r/sample-cards
- Example REST Data Source Plug-In on GitHub:
<https://github.com/oracle/oracle-db-examples/tree/master/apex/plugins/rest-source/fixed-page-size>
- 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.



Szymon Skorupinski

Smart Sizing of Fast Recovery Area in Cloud Environments



One of the most important aspects of keeping databases up and running is a capacity planning. Done properly at the project phase can save us from many critical problems later on. But that work doesn't end after the project phase is completed. During operational phase, we need to constantly monitor the requirements and react appropriately. There are different types of resources one has to consider and plan accordingly like processing power, memory or networking. What I want to focus on today is a database storage, more precisely, the special part of Oracle Database storage aimed to centralize and keep backup and recovery related files – Oracle Fast Recovery Area (formerly known as Flash Recovery Area).

Every Database Engineer knows how crucial the sizing of this area is. Oracle manages its contents quite well, but only if there is enough space available. Without proper sizing we can quickly render our database to be blocked for the users or jobs, which is of course something we prefer to avoid... The challenge with this specific area is that the requirements can change rapidly and in both directions. Sizing the storage for database files is usually easier, as the growth happens slower, most of the time in a constant rate and is rather growing, than decreasing. Bigger spikes caused for example by significant data loads or introducing new pluggable databases can be typically planned in advance. All those operations will increase the demand on FRA space too, but it can also significantly increase without a big change in overall database size.

Consider for example the situation with a workload which generates high amounts of redo, by processing huge amounts of data, but overwriting instead of adding it. Real-life example? Something went wrong and the processing needs to be re-executed, several times. Is flashback logging enabled? Amount of space required is basically doubled... That would be about keeping FRA big enough to avoid free-space issues. However,

nobody wants to waste the storage, as this can be quite expensive, especially if we start to speak about big numbers. So, we need to decrease the FRA space allocation at some point, for example after the busiest PDBs are being removed from our CDB or after cleaning up/archiving big parts of our data. That is inevitable, as Oracle tries to use as much FRA space as possible, in order to have the biggest number of files available to speed-up and ease potential recovery operations. Makes sense, but it means that transient files will be deleted only if FRA is nearly full and of course only if they are already marked as reclaimable or eligible for deletion (e.g. orphaned, redundant or obsolete due to whatever reason).

Now think about it in the context of increasingly popular cloud deployments, where one database serves multiple clients, using for example Oracle Multitenant option. For private cloud deployments we can still try to group PDBs in a way to easier manage what's going on or maybe even try to stay in touch with the owners to at least try to plan in advance. But let's face it – level of complication increases exponentially with the number of databases. At some point it is just too difficult to plan or forecast the periods with higher redo generation rate, new PDB creation (especially if it is done in automated way triggered for example by the client through some web interface), huge data loads etc. So, it is basically impossible to prepare proper sizing of FRA beforehand. The case of public clouds where any client can come with its own requirements at any point in time is even worse.

What would be the correct solution to those problems? The answer is quite short: proper and smart enough automation. It doesn't matter much how the resizing will be scheduled and run. It can be done using even good old Cron instead of for example more modern Ansible. The most important is how often it should run and on what basis the decision to act or not should be made.

Automating FRA resizing – when and how often

Let's discuss first on when and how often to run the automation. It depends on:

1. How big is the safety buffer we want/need to keep?
2. How fast is the database able to fill the safety buffer?
3. How fast is it possible to add/extend underlying FRA storage?

Answering above questions will let you decide when and how often to run your resizing jobs. Let's assume for example that it takes around 15 minutes to resize the underlying storage (e.g. Exadata with a lot of cells having big number of cell/grid disks) and the database is able to generate 1 TB of redo per hour. Not to waste too much space on safety buffer, you should probably run the resizing job every half an hour, always targeting to keep at least 625 GB of free space (as in the worst case you have to survive about 45 minutes before the new storage space will become available). It should be basically doubled in case you have flashback database enabled.

It is worth adding that the way on how to manage the safety buffer is also connected with how FRA behaves – by default, it tries to delete files after the threshold of more than 85% of used space is reached. However, depending on the environment and other variables, we may want to lower (in case of very big FRA) or increase this threshold (in case of small FRA). It is possible by setting special event, with the following command (number specified for level is the percentage after FRA starts to delete files):

```
SQL> alter system set events '19823 trace name context forever, level 70';
```

As you see, there are a plenty of options on how to implement the safety buffer – playing with thresholds allows you to do it inside FRA, but it is also possible to keep at least part of it outside, by having FRA always smaller than underlying storage. In the case of it being completely outside, we could even avoid “losing” 15 minutes of time on underlying storage expansion and do it in parallel with just increasing FRA by changing the `DB_RECOVERY_FILE_DEST_SIZE` parameter. Nevertheless, my suggestion is to keep it as simple as possible.

Smart sizing query

What I want to focus on is to give you an idea on how to get some numbers which will allow you to work on resizing algorithm appropriate for your environments. My approach is to avoid any manual deletions and let the database manage the files by itself. There are several things required to do it in the smart way:

1. Knowledge on how FRA manages the deletion of the files (e.g. deleting only when there is space pressure).
2. What is the impact and possibility to change the thresholds on the above algorithm?
3. What is the space usage of different types of the files inside FRA, as well as their status (reclaimable or not)?
4. What is the historical rate of generating those files?

The last point is crucial to properly size FRA considering current workload. Combined with well calculated and maintained safety buffer should allow us to size FRA appropriately, which means not to oversize, while keeping the database up and running if unexpected workload hits it.



The challenge is how to get the data described in the last two points. There are some summary views like [V\\$RECOVERY_FILE_DEST](#) (used and reclaimable space totals) or [V\\$RECOVERY_AREA_USAGE](#) which shows the division and status of different types of files, but not the raw numbers, just percentages. That's why I decided to use the definition of this view to get the values needed for doing the calculations. Unfortunately, its definition is bigger than 4,000 characters, which means that we won't be able to get it fully from [V\\$FIXED_VIEW_DEFINITION](#) view. However, this can be retrieved by [strings](#) and [grep](#) commands executed on [\\$ORACLE_HOME/lib/libserver12.a](#) file (please change the number to your database version accordingly).

For the purpose of the article, I want to show you the query focusing only on 2 types of files in FRA (archived logs and flashback logs) where underlying storage is based on ASM (*see Listing 1*). The query was tested on Exadata environment running both Oracle 12c and 19c.

```

define no_of_days = 14
define rounding = 2
define reco_dg = 'RECO%'
with archlog_daily_rate as (
  select trunc(completion_time) day,
         round(sum(
           case when ceilasm = 1 and name like '+%'
                then ceil(((blocks * block_size)+1)/1048576)*1048576)/power(1024,3)
                else blocks*block_size/power(1024,3)
           end),&rounding) gb
  from v$archived_log,
       (select /*+ no_merge */ ceilasm
        from x$krasga)
  where standby_dest != 'YES'
     and archived = 'YES'
     and fal = 'NO'
  group by trunc(completion_time)
)
select
  (select db_unique_name || ' - ' || decode(database_role,
    'PHYSICAL STANDBY', 'STANDBY', database_role) from v$database) cdb,
  (select round(case type
    when 'HIGH' then total_mb/1024/3
    when 'NORMAL' then total_mb/1024/2
    else total_mb/1024
  end,&rounding)
  from v$asm_diskgroup where name like '&reco_dg') recodg_gb,
  (select round(space_limit/power(1024,3),&rounding)
  from v$recovery_file_dest) fra_gb,
  (select round(space_used/power(1024,3),&rounding)
  from v$recovery_file_dest) fra_used_gb,
  (select round(space_reclaimable/power(1024,3),&rounding)
  from v$recovery_file_dest) fra_reclaimable_gb,
  fra_flashback.*,
  fra_archlogs.*,
  generated_archlogs.*
from (
  select
    max(gb_2days) max2d, round(avg(gb_2days),&rounding) avg2d,
    round(median(gb_2days),&rounding) med2d, max(gb_3days) max3d,
    round(avg(gb_3days),&rounding) avg3d,
    round(median(gb_3days),&rounding) med3d
  from
    (select
      to_char(day,'DD.MM.YYYY'),
      sum(gb) over (order by day range interval '1' day preceding)
      as gb_2days,
      sum(gb) over (order by day range interval '2' day preceding)
      as gb_3days
    from archlog_daily_rate
    where day >= sysdate - &no_of_days
       order by day)) generated_archlogs,
  (select
    round(nvl(fl.space_used, 0)/power(1024,3),&rounding) fb_used_gb,
    round(nvl(fb.reclsiz, 0)/power(1024,3),&rounding) fb_reclaimable_gb
  from (select
    round(sum(case when ceilasm = 1 and name like '+%'
      then ceil(( fl.bytes + 1 ) / 1048576) * 1048576
      else bytes
    end),2) space_used
  from
    v$flashback_database_logfile fl,
    (select /*+ no_merge */ ceilasm
     from x$krasga) fl,
    (select sum(to_number(fblogreclsiz)) reclsiz
     from x$krfblog) fb) fra_flashback,
  (select
    round(sum(al.file_size)/power(1024,3),&rounding) archlogs_used_gb,
    round(sum(case when dl.rectype = 11
      then al.file_size
      else 0
    end)/power(1024,3),&rounding) archlogs_reclaimable_gb
  from
    (select
      recid,
      round(case when ceilasm = 1 and name like '+%'
        then ceil(((blocks * block_size)+1)/1048576)*1048576
        else blocks * block_size
      end,&rounding) file_size
    from v$archived_log,
         (select /*+ no_merge */ ceilasm
          from x$krasga)
    where is_recovery_dest_file = 'YES'
       and name is not null) al,
         x$kcagf dl
    where al.recid = dl.recid(+)
       and dl.rectype(+) = 11) fra_archlogs;

```

Listing 1

Thanks to subquery factoring and a few subqueries, all the information is retrieved in one line, which allows to use it for getting the report from all the databases and import it for example to Excel for further analyzing or visualizing. The query gives all the data points described above and required to calculate the target size of FRA. Apart from the CDB name and type, we have also: the sizes of underlying disk group and FRA, the sizes of both used and reclaimable FRA space in total, as well as only for flashback and archived logs. Additionally, thanks to the first subquery, we can calculate maximum, average and median redo generation rate throughout 2 and 3 consecutive days (more periods can be added) during the last 14 days. How much we would like to go into the past is parametrized and of course it affects how quickly we will react to changing circumstances (e.g. using shorter periods will lead to more frequent changes of target FRA size). Having those results, we can finally decide what we would like to size FRA for (e.g. be able to keep at least 2 days of redo logs), please just ensure to add the safety buffer to this number.

I hope that provided information would allow everyone to design efficient FRA sizing algorithms taking into consideration each environment specifics, so it would behave correctly in all circumstances. Having above query as an example should be sufficient to adjust the results according to the requirements, for example by adding other types of the files stored in FRA to the equation.



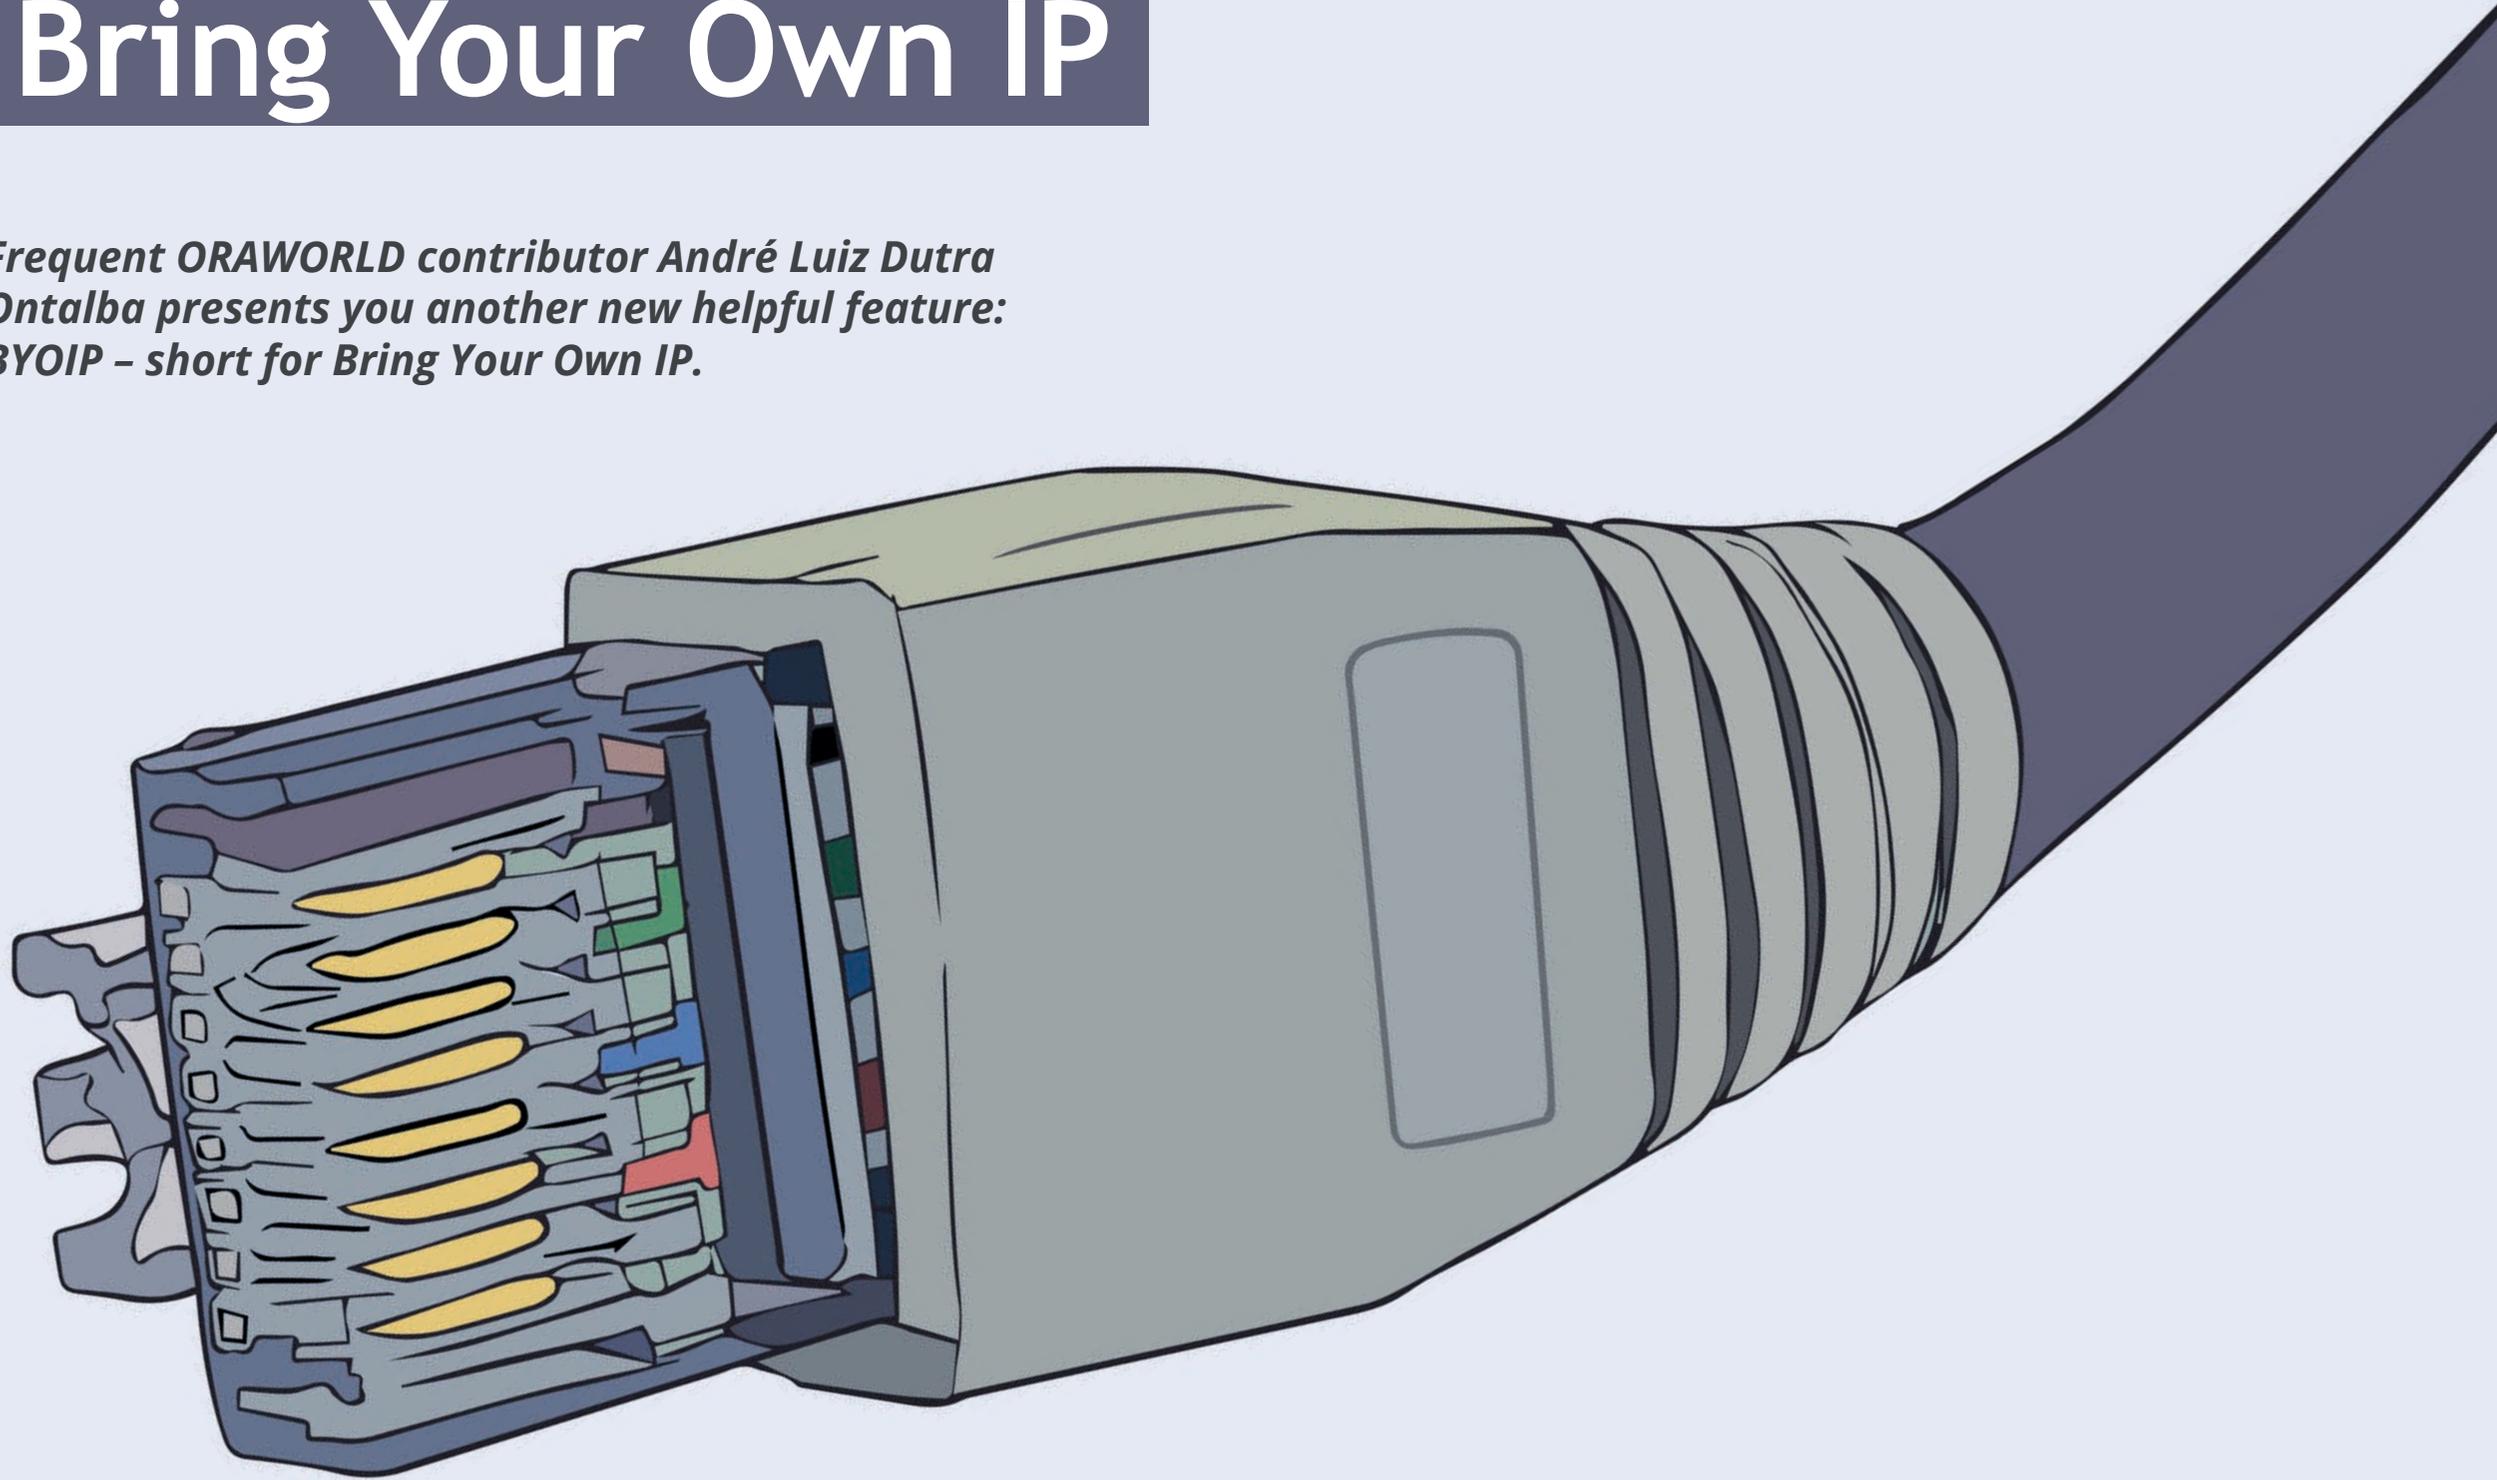
About Szymon Skorupinski

Szymon is a Principal Database Engineer at Trivadis AG and Oracle ACE Associate, with over 16 years of experience working with complex IT environments at every stage of their lifecycle. He is an Oracle OCP for 10g, 11g and 12c, as well as an Oracle Cloud Infrastructure 2019 Certified Architect Associate and Oracle Autonomous Database Cloud 2019 Certified Specialist.

André Luiz Dutra Ontalba

New Feature: Bring Your Own IP

*Frequent ORAWORLD contributor André Luiz Dutra Ontalba presents you another new helpful feature: **BYOIP** – short for **Bring Your Own IP**.*



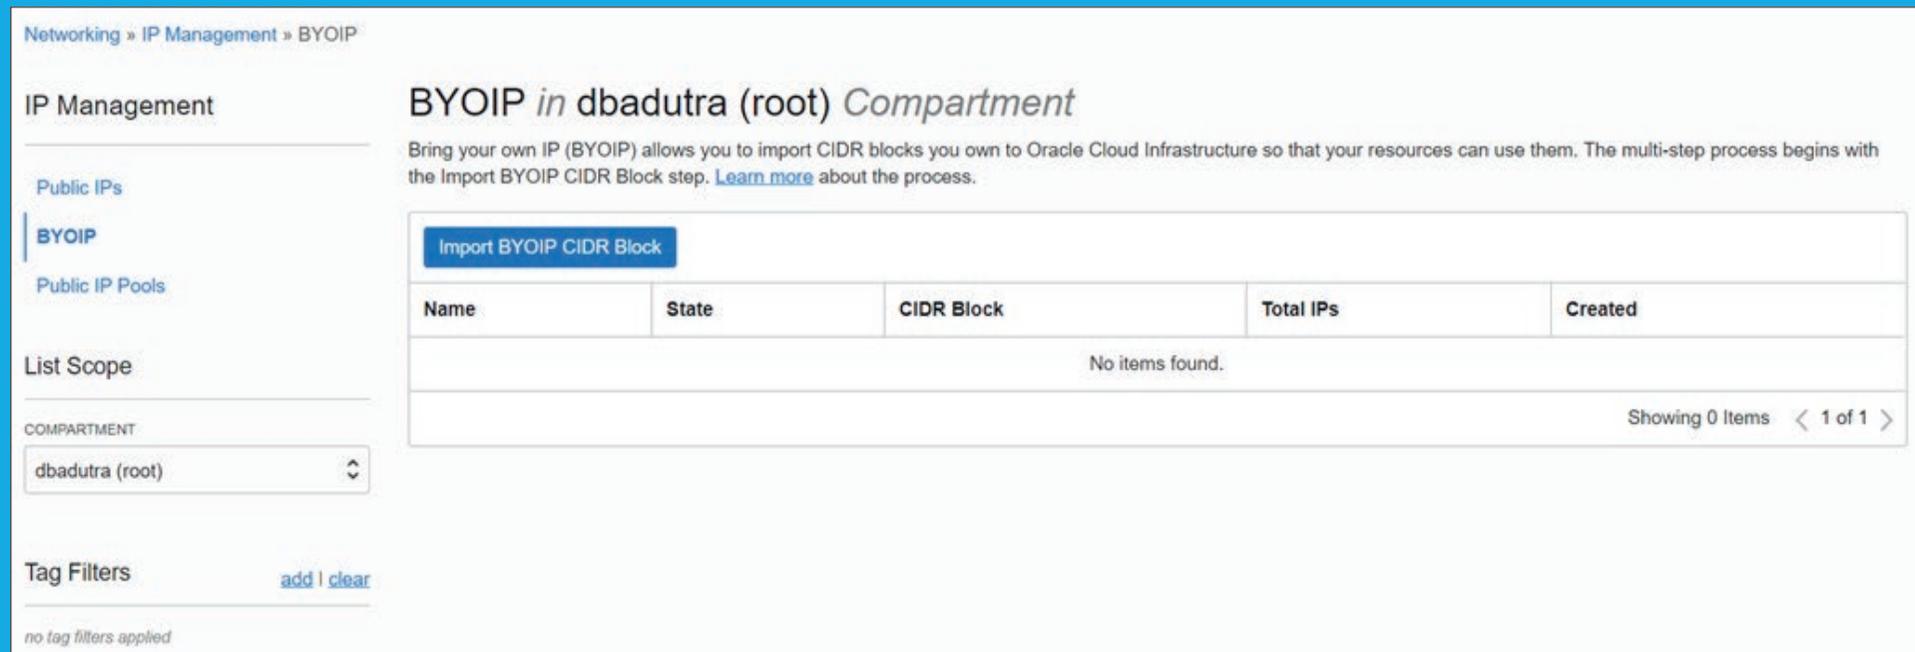


Image 1: The new BYOIP feature

Oracle Cloud Infrastructure (OCI) now comes with an exciting new feature: It allows you to Bring Your Own IP (BYOIP) address space to use with resources in OCI. BYOIP (*see Image 1*) lets you manage your IPv4 CIDR blocks to align with your existing deployment policies, security, management and achieve the following:

- Solution continuity and hardcoded dependencies
- IP pool management
- IP reputation

After a validation process on imported CIDR blocks you are notified that the CIDR block is available for advertisement.

Limits and Quotas

It is important to note that the feature comes with the following restrictions:

- A CIDR block that is /24 or smaller is required for addresses to be advertised to the Internet.
- You can't bring more than 10 IPv4 address ranges to your OCI account.
- You can create up to 10 pools.
- BYOIP is not available to Oracle Cloud Infrastructure Free Tier or Pay As You Go services.
- Your addresses can only be imported to a specific Oracle region.
- You can use BYOIP with a CIDR block that is a minimum of /24 and a maximum of /8.
- You can't bring the same address range to more than one compartment at a time.

Import BYOIP CIDR Block [Help](#)

The Import BYOIP CIDR Block workflow follows these steps:

- Enter the BYOIP CIDR block on this form. You will get a validation token in the next step.
- Make changes to the validation token string. See the [documentation](#).
- Provide the modified validation token to your regional internet registry service. Wait for the registry to update the information for that CIDR block.
- Return to the Console and click **Finish Import** so Oracle can complete the process. The process can take up to 10 business days. You are notified by email when the process is complete. You can also check the progress in your work requests.

BYOIP CIDR BLOCK NAME

COMPARTMENT ⓘ

dbadutra (root) ↕

CIDR BLOCK

Maximum size /8, minimum size /24

[Cancel](#)

Image 2: Importing the BYOIP CIDR Block

Requirements and Preparation

You must have ownership of the public IPv4 CIDR block (*see Image 2*) you want to import into Oracle Cloud Infrastructure, and the ownership must be registered with a supported Regional Internet Registry (RIR). The Oracle validates ownership of your addresses. Only the following registries are supported, and the addresses must have a specified type or status:

American Registry for Internet Numbers (ARIN) -
“Direct Allocation” and “Direct Assignment” network types

Réseaux IP Européens Network Coordination Centre (RIPE NCC) - “ALLOCATED PA,” “LEGACY,” “ASSIGNED PI,” and “ALLOCATED-BY-RIR” allocation statuses

Asia-Pacific Network Information Centre (APNIC) –
“ALLOCATED PORTABLE” and “ASSIGNED PORTABLE” allocation statuses

The addresses in the IP address range must have a clear history. Oracle can investigate the reputation of the IP address range and may reject a range of IP addresses that contains an IP address associated with malicious behavior.

The steps required for BYOIP on Oracle Cloud Infrastructure require significant time, so I recommend planning beforehand to perform such an import. I hope this article 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.

7 1 - 15 October



67 Live Sessions

3574 Attendees

7 Host Countries

6 Tracks

70+ Hours

99 Countries represented

61 Speakers

6 Ground Breakers Ambassadors

13 Oracle advocates

18 ACEDs

12 Days

19 ACEs

1 ACE Associate



EMEA Virtual Tour 2020

#BackToBasics

15 Tips, 15 Speakers, 15 Days Mirela Ardelean

For several years now, I have been following Oracle tech conferences stage and I have noticed that autumn is always full of events. The year 2020 was no exception, on the contrary...

Follow us on



@EOUC

@ORAWORLD_Mag

#ORAWORLD



@emeaoracleusergroups

@ORAWORLDMag



www.oraworld.org

Among the hundreds of hours of webinars running during this fall, a virtual tour managed to slip in: the Oracle Groundbreakers EMEA Virtual Tour. User Groups from Azerbaijan, Belgium, Croatia, Luxembourg, Romania, Saudi Arabia, and Tajikistan worked together under EOUC guidance to build this series of online events.

For two weeks, from October 1 to October 15, more than 60 speakers from all over the world did a virtual trip, from Central Asia all the way to Western Europe, sharing their knowledge with us, the community.

This year was different in all its aspects. For the Oracle Groundbreakers EMEA Virtual Tour 2020 was about #diversity and was about #BackToBasics.

But how did the idea of #BacktoBasics come about? In a Twitter group I belong to, a question was asked early this year: "is a presentation about <topic> too basic?" The common opinion was: "no, it isn't". There are people at the beginning of their career or people who already have a few years of experience but want to change their technical field of expertise; everyone will learn something from a 101-type presentation, regardless how many years of experience they have or how "fresh and new" they are.

The discussion, but especially the answer to the question, gave me the idea to set #BackToBasics as the motto for Oracle Groundbreakers EMEA Virtual Tour. At the moment of actually planning the virtual tour I was asking myself: how can we make the #BackToBasics messages to be heard by as many people as possible? The obvious answer was: by using social media!

So I addressed the idea in the Twitter group (almost everyone in this group is a rock-star speaker in the Oracle world): "I am looking for 15 volunteers to post on Twitter one technical tip #BackToBasics, one on each day of the Oracle Groundbreakers EMEA Virtual Tour."

I was surprised how quickly the 15 volunteers responded, but I shouldn't have been surprised: they are recognized as eager to share their knowledge with the community. Below, are the famous 15 speakers and their tips. For anyone who wants to read the full threads, here is the collection on Twitter: <https://twitter.com/i/events/1316681343804338176>

Will 101 presentations become trendy (again)? Hard to say yet. Considering the number of Twitter impressions, which is high for each of the above tweets, and the fact that some future events included 101 presentations in their agenda, I would say:

Let's go #BackToBasics!





 **Heli Helskyaho @ Home** @HeliFromFinland · Oct 1

Data quality is the most important requirement for a database and can only be achieved by designing the database. Designing database should not be done with a paper&pen but with a tool that supports the process. Free and easy to use: oracle.com/database/techn... #BackToBasics

1 14 31

 **Jeff Smith** 🍷 @thatjeffsmith · Oct 3

My Tweet for #OGBEMEA #BackToBasics:

If it ends up in production, it should be source controlled.

This includes your DATABASE!

Object DDL, Stored Procedures, Upgrade scripts, lookup fields, etc.

For extra credit, source control your data models.

6 20 49

 **Erik van Roon** @evrocs_nl · Oct 2

My Tweet for #OGBEMEA #BackToBasics:

When you execute a PLSQL function in a SQL query, please always be aware that any SQL inside that function will not return the truth as it was when your main query started, but when that query inside the function started #ReadConsistency

5 12 55

 **Samuel Nitsche** @Der_Pesse · Oct 5

#OGBEMEA #Back2Basics Daily Tip Nr. 5:

EVERY program functionality is ALWAYS a transformation from INPUT to OUTPUT.
Often in DB, the input is indirect (e.g. values in a table).

If you identify input and output, your functionality will become much clearer - and more testable.

7 18

 **Kim Berg Hansen** ♠️ @kibeha · Oct 4

The #OGBEMEA #BackToBasics Daily Tip Tweet no 4:

Getting TOP-N per group is easy in #SQL with analytic functions (aka window functions).

The hard part is deciding which function is applicable to your task:

- dense_rank()
- rank()
- row_number()

#AnalyticsRockAnalyticsRoll

```
partition by s.grp
order by s.qty desc, s.item
) rnum
from item_sales s
where drnk <= 3 or rnk <= 3 or rnum <= 3
order by grp, qty desc, item;
```

Script Output x Query Result x

SQL | All Rows Fetched: 8 in 0,03 seconds

GRP	ITEM	NAME	QTY	DRNK	RNK	RNUM
1	AUTO	101010 Brake disc	33	1	1	1

3 23 87





Liron Amitzi @amitzil · Oct 6

#OGBEMEA #Back2Basics Daily Tip #6:
Constraints are your friends. They are the only real way to enforce data integrity and consistency across all of your applications, and they can even lead to performance improvements.

1 12 37



Martin Berger @martinberx · Oct 7

My Tweet for #OGBEMEA #BackToBasics:
The best way to improve performance is to avoid all unnecessary activities.

Consider this on all levels,
might it be the avoidable business process,
report, code/SQL or simple I/O

1 12 24



RogerMacNicol @RogerMacNicol · Oct 8

#backtobasics #OGBEMEA Table free space is your friend. Calculate avg row size when rows are first inserted and avg row size when they are fully populated by business processes. The difference between the two is how much initial free space your tables need.



4 12 26



Franck Pachot @FranckPachot · Oct 10

#OGBEMEA #Back2Basics Daily Tip Nr. 10:
There are two SQL languages

One for data scientists, ad-hoc queries, DSS/DWH/BI, with literal values and maybe SELECT *

And one programmatic, declaring prepared cursors, with bind :variables and explicit list of columns in select clause

1 4 17



Piet de Visser (Товарищ Пит - в спячке) @pdevisser · Oct 9

Indexes must match the Question.
If a PK, UK or FK is used to look up records,
And if your queries often contain a "status=Active",
Then you Should add that "status" field to the index.

#KnowYourData
#Design
#OverloadedIndexes
#fatindex
#compositeIndex
#BacktoBasics

7 14





Ron Ekins @RonEkins · Oct 14

The #OGBEMEA #BackToBasics Daily Tip Tweet

Whenever you are troubleshooting or tuning remember to: 1)establish a baseline 2)agree end goals 3)understand impact of change 4)limit number of changes being applied at once 5)do representative testing 6)document and analyse results



Martin Widlake @MDWidlake · Oct 13

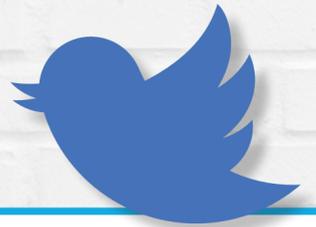
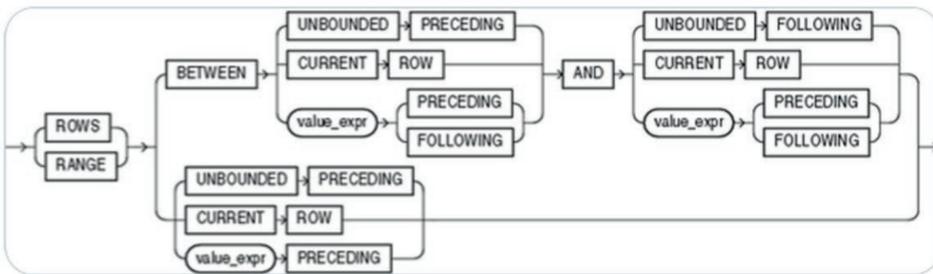
#OGBEMEA #BackToBasics

SQL is designed from the ground up to process huge volumes of data:
Joining it
grouping it
min, max, average, sum.

But it does so much more:

Windows functions.
statistical analysis
pattern matching

The better you learn SQL the more it will do for you.



Sabine Heimsath @home @oraesque · Oct 14

#OGBEMEA #backtobasics tip for #PLSQL:

Always comment on exceptions to the rules

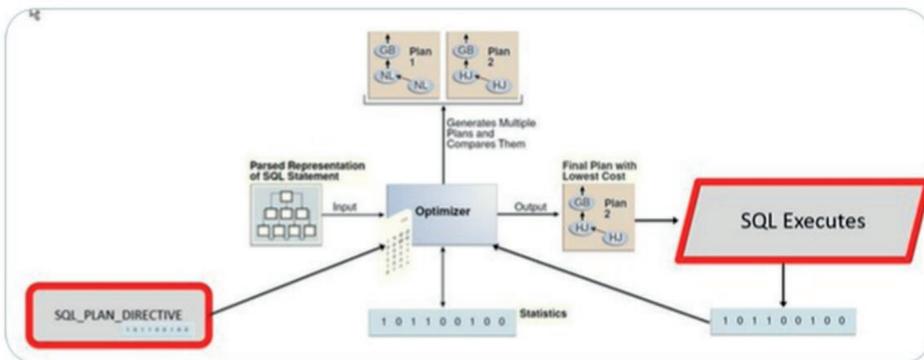
If you write 'when others then null', or use 'union' instead of 'union all', put the reason right next to it.
The next person to read the code will know immediately that you did it on purpose.



Neil Chandler @ChandlerDBA · Oct 11

#OGBEMEA #BackToBasics

The oracle optimizer is only as clever as you let it be. Having good statistics and lots of constraints allows it to transform SQL into more efficient code, avoiding work and picking the quickest path to your data. Help it, and it will help you!



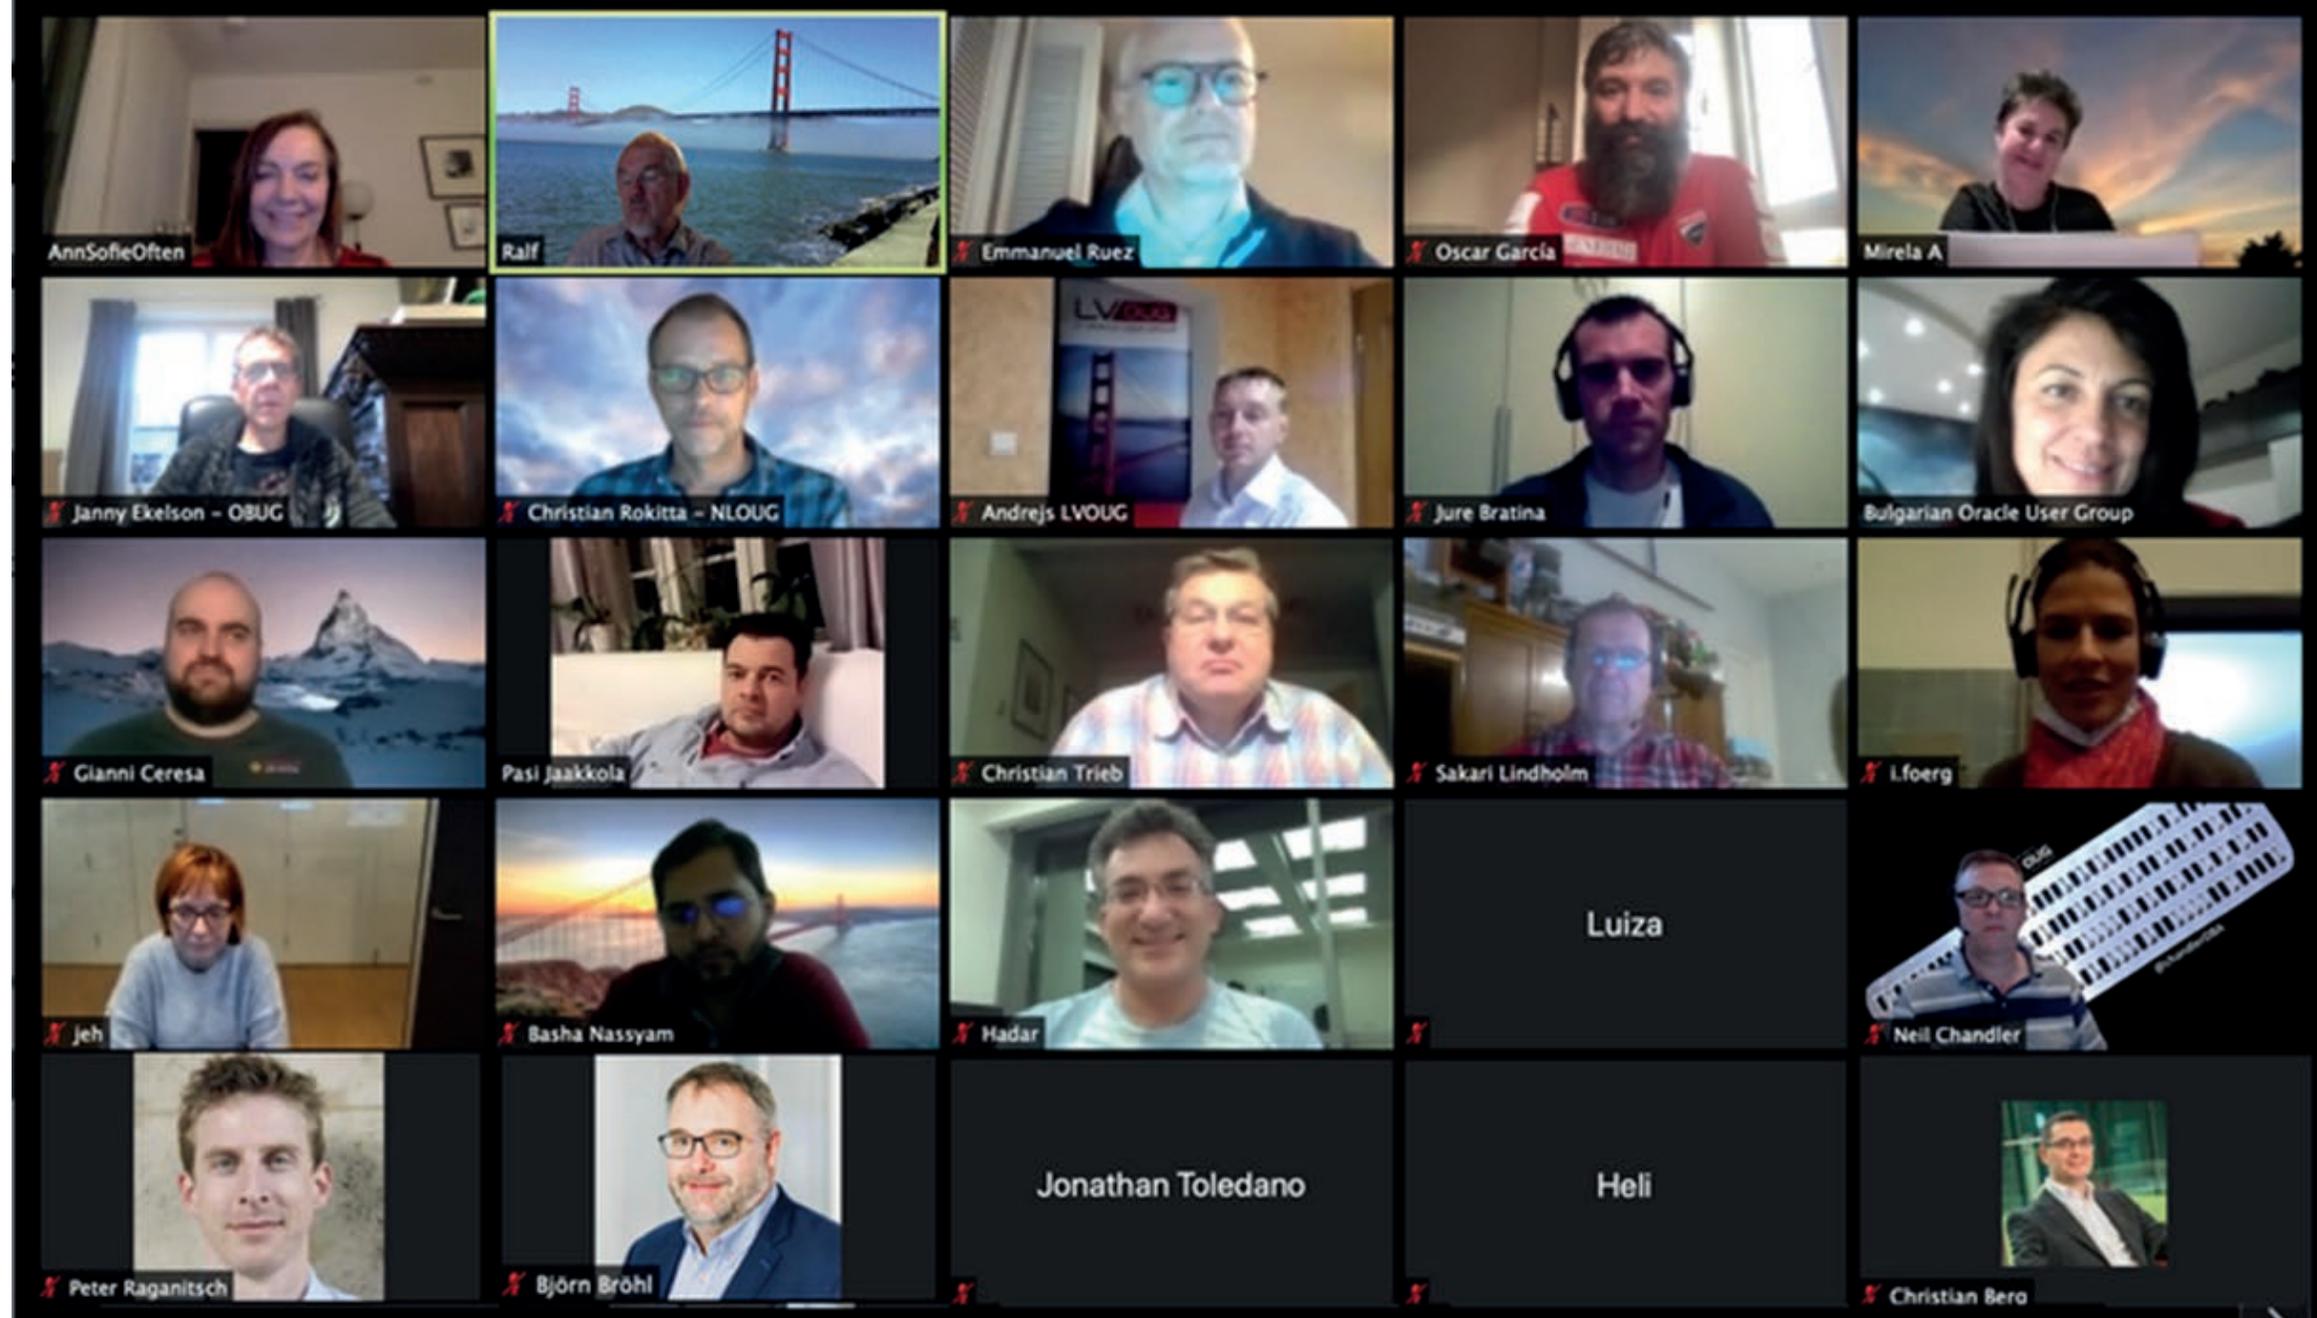
Flora Barriele @floo_bar · Oct 15

#OGBEMEA #BackToBasics:

To prevent accidental/intentional damage in a system, maintain a healthy privilege management by enforcing principle of least privilege + avoiding privilege creep.

In #database words: grant the bare minimum to users/roles + revoke unnecessary privilege!





EOUC Leaders' Virtual Summit Mirela Ardelean

*For many ORAWORLD readers the month of November is equivalent with **DOAG Conference**. In addition, for Oracle User Group (OUG) leaders, November means also that the traditional EMEA OUGs Leaders' Summit takes place, organized by Oracle User Groups for Oracle User Groups.*

Preparation

In recent years, two annual meetings of OUGs Leaders have been organized: one in May-June, organized by Oracle (let's call it Spring Summit), and one in November, organized by OUGs (the Fall Summit). At the end of 2019 we already knew that Oracle would not organize the Spring Summit in 2020: following OUGs leaders vote, EOUC had already started the process of transformation from an informal community to a legal entity. In that particular context, EOUC decided to take over the organization of both summits. At that time, it was not suspected that the year 2020 would be different.

January 2020: we had already started the discussions for the organization of the spring Summit. First option put on table: April, one day before APEX Alpe Adria. When things started to look bad in Europe, due to Covid-19, we decided to move the Summit closer to summer, hoping that the situation would be better and that physical meetings would be allowed. Second option: May, the day before Riga Dev Days 2020. Then travel restrictions were imposed, entering as a tourist in certain countries was banned. We gave up, resigned ourselves to the organization of the Spring Summit.

But still, it was hope for the Fall Summit; being a little cautious, we decided to consider the hybrid format for it. During the discussions over the summer, two options seemed valid: October, before HROUG Conference, or November, as usual, before DOAG Conference. But as soon as HROUG was cancelled, we concluded: let's do a virtual summit, during DOAG Conference, and that's it! At least, we keep our tradition.

The Summit

On the second day of DOAG Conference, the EMEA OUG Leaders' Summit took place. For reasons strictly related to the availability of participants and their possible overexposure to online meetings, we reserved only two hours for discussions.

We were very happy when we received the confirmations that over 30 user group leaders from more than 25 countries across EMEA, plus the representative of ODTUG, were joining our virtual event. The main theme of this year's summit was "The Future of the Oracle User Groups" with a short agenda:

- round table: OUGs in pandemic time
- community events: plans for 2021

For the round table, each representative was invited to describe how their user group managed their activities during pandemic, the impact of the 2020's restrictions in their user group activities, challenges when swapping from live-events to online-events, sponsorship issues or success stories.

We heard stories about successful events, about struggling with board members leaving, about finding the right format for an online event, about adopting diversity in topics and technology areas, about financial loss, about wondering "what's the correct price for an online event when there are so many free events?"

We also got some hints which pitfalls to be avoided when planning or conducting such events. The fact that most of our user groups are still active, organizing different kinds of events and trying to make the best out of the current situation, is very positive. We, all, are looking forward to a hopefully more 'normal' next year to physically meet each other again. The common attitude towards planning the events in 2021 is: "online events for the first half of the year, then ... we'll see".

A topic raised during the discussion was the need to intensify networking of the user groups for event planning, especially for the time when on-site conferences will be allowed again:



events should be coordinated efficiently, to lead to attractive scheduling for the speakers. A common place to check the calendar of future events is needed.

Regarding the community plans for 2021, the EOUC board announced the launch of the MASH program: Mentor and Speaker Hub. The main goal of MASH is to provide help to anyone who wants to start their way as a speaker and would like to present a session in tech conferences. Also, the EOUC board presented the intention to continue organizing the Oracle Groundbreakers Tour 2021 (date and format will be discussed during January 2021).

Last, but not least important, the EOUC board will organize the 2021 Spring and Fall Summits: the spring one will be virtual, that is for sure. For the autumn one we all hope it will be on-site, somewhere in Europe, close to a user group conference.

The discussion ended too soon but we plan to continue it somehow during the next summit: user groups who shared their success stories have been kindly asked to be prepared to present a more detailed version of their achievements.

Conclusion

Whether a user group is small or big, organizing big or small events, we think it is useful to know how others are doing, what struggles they have and how they try to deal with the daily challenges. We are looking forward (e-)meeting everyone again.

Stay safe!



About Mirela Ardelean

Mirela is a Oracle Database Developer with more than 20 years of experience, an APEX enthusiast who feels some nostalgies for Oracle Forms. She was the co-founder of RoOUG (Romanian Oracle User Group) in 2012 and has been the Vice President of RoOUG since the beginning. Mirela is very much involved in Oracle User Groups communities, trying to build a strong collaboration between OUGs. When EOUC was founded, in November 2019, Mirela became one of the board members. She tweets as [@Mirela_RoOUG](#) about what's happening in the OUGs' world.

Tech Cycle

“User Word”

Migrating

Mission-Critical

Applications

to the

Oracle Cloud

Emmanuel Ruez





In the previous ORAWORLD issue #22 we presented the **activities of the French-speaking Oracle User Clubs during the pandemic** which included a large number of webinars. The 3 French User Clubs have organized many more webinars since then.

Among them was a webinar as part of the Tech Cycle “User Word” series organized together with Oracle France. For **the first session** on September 17, the client was Essilor, an “historic client” of Oracle and an “historic member” of the Board of AUFO (l’Association des Utilisateurs Francophones d’Oracle).

The topic of the session was “sharing experience of moving to cloud”. In particular, a “Lift & Shift” migration of mission-critical applications from an on-premises E-Business Suite Oracle Cloud Infrastructure (OCI) – in total a migration of 7 instances in Europe, Asia, and the United States.

After a presentation of the client and its organization, the two speakers explained the nature of the project and the reasons why the choice was made to use Oracle Cloud. They shared the successes and the specific challenges they faced. They concluded the session by talking about the benefits that this project has generated for their organization.

In introduction of the session, Nicolas Villette, Cloud Director France at Oracle, presented the worldwide deployment of Oracle Cloud Infrastructure and how this meets user security requirements.

The participants of the webinar particularly appreciated the quality of the information provided and the transparency with which the speakers answered the questions put to them.

All of our past and upcoming webinars can be found on the [clubutilisateursoracle.org website](https://clubutilisateursoracle.org).



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.

Events



ILOUG TechDays 2021

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

JavaLand 2021

March 16 - 18, 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/>



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

The following user groups belong to EOUC:

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

Editorial board:

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

www.doag.org,

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

Contact: redaktion@doag.org

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

Editorial team:

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

Disclaimer:

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

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

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

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

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

Article submission:

If you are interested in submitting an article, please e-mail us your article via the online form at **www.oraworld.org**. Deadline for the next edition: **January 29, 2021**.

Photo credits:

Title: © Abdul Qaiyoom | <https://stock.adobe.com>

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

P. 5: © Comic | www.commitstrip.com

P. 6, Person: © Pavlo Plakhotia | <https://123rf.com>

P. 6, Background: © antishock | <https://123rf.com>

P. 7 - 11, Objects: © marcovector | <https://freepik.com>

P. 12: © vectorjuice | <https://freepik.com>

P. 13: © Alexander | <https://stock.adobe.com>

P. 14 - 17, Clouds: © Freepik | <https://freepik.com>

P. 18: © Anastasiia Nevestenko | <https://123rf.com>

P. 27: © Graf Vishenka | <https://stock.adobe.com>

P. 32: © Yada Maneesri | <https://123rf.com>

P. 37, Background: © kues1 | <https://freepik.com>

P. 38 - 39, Background: © kues1 | <https://freepik.com>

P. 38, Person with Birds: © stories | <https://freepik.com>

P. 44: © CoreDESIGN | <https://www.shutterstock.com>

P. 47: © DOAG | www.doag.org/en/home