



- New APEX Series – Part 1: “Low Code with the Database”
- Running Your JuPyTeR Notebooks on the Cloud
- What’s New after Oracle OpenWorld 2019



Self-Service Integration, What?, Why? and How?

page 9



Running Your JuPyTer Notebooks on the Cloud

page 22



The Groundbreakers EMEA Tour 2019

page 41

Editorial

Submit Your Article!

## I. Work & Life

Comic: Discount

Number of the Month

## II. Techs & Nerds

What's New after  
Oracle OpenWorld 2019

3	Self-Service Integration, What?, Why? and How?	9
4	Oracle Application Express (Part 1): Low Code with the Database	15
5	Running Your JuPyTer Notebooks on the Cloud	22
6	Video: Amazon Says „Bye-Bye Oracle”	30
	Ubuntu® and Its Circle of Friends	32
8	Oracle Exadata Database Machine: Decade of Innovation	34

## III. Users & Groups

The Groundbreakers  
EMEA Tour 2019 41

Ambassador's Corner 45

## IV. Past & Future

Call for Papers 46

Events 47

Contact us 49

Legal notice 49

# Editorial

Dear Oracle User Group Members,

Information Systems is a “world” which seems to run like a pendulum between development on one hand and integration architecture on the other hand.

The oldest of us remember the great old days of mainframes and development languages, then came the software packages ERPs. That was the time when our roles seemed to be integrators and maybe to develop a few (the less possible) custom extensions.

Then it shifted again to the other side with the web wave by the beginning of this century.

I feel that we are now back again to the other extreme, as we can see in this ORAWORLD edition and in other sources, the main words being SaaS (looks a bit like Software packages promises of before?), Autonomous DB and everything, low code (see the article about APEX), Architecture - part of it being “how do we integrate all that?” Does that ring a bell?

What will be the next evolution triggering the movement: Digital transformation is the current buzzword, interconnection of organisations to improve for example supply chain, API, .. could be it.

This always changing “world” is certainly why we love it so much, and this continuous change is one of the reasons we benefit from exchanging experiences and brainstorming together in our user groups.

Keep being surprised, and agile!

Jean-Jacques Camps



**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](http://www.oraworld.org).








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

# Number of the Month: Martin Meyer

Digital Transformation: only 11% of Europeans trust that their national governments are best placed to take responsibility for ensuring, that new technologies improve society.





According to a **representative survey** conducted by the opinion polling institute Ipsos on behalf of the Vodafone Institute, an average of 44 % of Europeans say that they feel very or somewhat enthusiastic about the growing influence of digital technologies in many areas of life. The study represents the views of people aged over 18 years from the five countries Bulgaria, Italy, Germany, Sweden and United Kingdom.

The general feelings towards the influence of digital technologies vary from 35 % (Germany), 38 % (United Kingdom), 47 % (Italy), 45 % (Sweden) to the most enthusiastic Bulgarians (57 %).

The question of who should ensure the safe and sustainable use of new technologies to improve society presented a complex picture. On average, 25 % of Europeans would leave the responsibility to consumers themselves. 23 % believed that the government is the most likely to do so, with 19 % each of the tech companies and consumer protection agencies and only 14 % of non-governmental organizations.

When it comes to the matter of trust, only an average of 11 % of the surveyed Europeans believe that their national governments are the best actors to take responsibility for ensuring, that new technologies improve society. Only in Sweden the government enjoys with 26 % the highest rating value among the potential actors.

Among the questions concerning aspects of digital transformation were also topics like personal robots, virtual


reality technology in schools, tiny bio-robots implemented in people's bodies, automated administrative procedures, autonomous cars, distribution of goods by drones and doctor decisions based on data stored on huge databases.

Considering the trust in technology the survey included also the question, whether people were willing to trust a humanoid robot as a caregiver, that has been programmed or even only controlled remotely by experienced caregivers.

Respondents clearly preferred human experienced caregivers to 56 % on average, followed by family members assisted by caregivers (43 %) or friends assisted by caregivers (21 %). Humanoid robots, which could take over care tasks in the future, only had the confidence of an average of about 10 % of the survey participants.

According to another poll, the representative study **"The Tech Divide: Policy"** – which revealed contrasting attitudes towards digitization in Europe, Asia and the USA – the majority of Europeans do not trust their respective governments to advance digital transformation. In European states, only 40 % of those surveyed see their respective governments as willing to advance digitization; and only 34 % believe that they have the adequate skills and capacity to do so.

"The government symbolizes the opposite of what digitization stands for. It symbolizes safety and continuation, while digitalization stands for transformation and change." said an industry expert and CEO of Start-Up Hub, Sweden.



Follow us on



@EOUC

@ORAWORLD\_Mag

#ORAWORLD



@emeaoracleusergroups

@ORAWORLDMag



www.oraworld.org





### Autonomous Linux

Oracle expands its Gen 2 Cloud services and introduces the first autonomous OS: Oracle Autonomous Linux patches, tunes, and updates itself while running, removing human errors, cutting downtime and improving security. It comes together with the new OS Management Service, included in Oracle's Premium Support package.

### Cloud Applications

Oracle's cloud applications have received updates and new features. For CX, advertising and marketing have been unified and new B2B sales capabilities have been added. HCM has new links to LinkedIn that allow importing profiles and connecting with recruiters. ERP has received an intelligent document recognition function, and SCM comes with the new service Oracle Business Network.



### Oracle Cloud Free Tier

The new Oracle Cloud Free Tier includes Always Free services for anyone to try the self-driving database and OCI for an unlimited time. The program offers two autonomous databases, two compute VMs, two block volumes, object and archive storage, a load balancer, and substantial allowances for outbound data transfer, monitoring services, and notifications.



### Exadata

The new Oracle Exadata Database Machine X8M integrates new technologies, including direct remote memory access via Converged Ethernet (RoCE) as a high-bandwidth, low-latency network structure and persistent memory accelerator with Intel® Optane™ DC persistent memory modules.



### Microsoft partnership

An integration between Oracle Digital Assistant and Microsoft Teams allows enterprise customers to query various Oracle Cloud Applications through text chat. The attendees were also reminded that Oracle's partnership with Microsoft makes it easier to connect workloads in both Azure and the Oracle Cloud.

### Oracle Cloud Marketplace

Shield IO, Couchbase and Fortinet are now among the software providers selling their products through Oracle Cloud Marketplace. From now on, customers can rent software by the hour without having to make long-term licensing commitments. Payments can be made with Universal Credits, the same system with which customers buy Oracle Cloud Services.



### VMware partnership

Oracle has announced a strategic partnership with VMware and will officially support Oracle products running on VMware virtualized environments. VMware will help Oracle run VMware workloads in the Oracle cloud. It may also improve support for the many databases running on VMware.



# What's New after Oracle OpenWorld 2019



### Oracle Database

A preview of the most important software innovations contained in Oracle Database 20c and Oracle Autonomous Database included a native persistent storage database (PMEM), Automated Machine Learning (AutoML), native blockchain tables for secure transactions, and a binary JSON data type for better performance.



### Digital Assistants


Digital Assistants is now available across all Cloud Applications. In addition to text chat they can now help employees and customers by voice. Compared to the well-known assistants for consumers like Siri or Alexa, Oracle's Digital Assistants can better distinguish enterprise-specific vocabulary.



Arturo Viveros

# Self-Service Integration, What?, Why? and How?





## We're all becoming citizen-integrators...

Integration has long been a key enabler when it comes to realizing almost any kind of business-driven IT initiative. In order to get the most out of their investments in multiple technologies, organizations need systems, services and applications to be able to easily communicate and share data with each other through secure and reliable mechanisms.

Moreover, innovation is constantly disrupting the IT ecosystem and influencing the way systems are designed and implemented. The never ending stream of requirements for modern solutions constantly challenges and stretches the traditional boundaries of software development, fostering the adoption of new paradigms such as Cloud-Native, which aims at building and deploying systems in a way which ensures that the benefits of the Cloud Computing delivery model will be vastly exploited.

This is where the concept of self-service comes in. It's well documented that the three core flavors in which Cloud Computing offerings are commonly delivered are:

- Integration as a Service (IaaS)
- Platform as a Service (PaaS)
- Software as a Service (SaaS)

Each one of these models has its own characteristics and benefits, but one of the most important things they have in common is the notion of the consumer of the services being able to procure them and customize them to a sufficient extent without the intervention of a third party.

Self-service of course requires an important level of automation as well as built-in support for multi-tenancy, auto-scaling and a highly cohesive user experience in order to function properly, allowing consumers to effectively and rapidly leverage their Cloud subscriptions.

Even though the latter applies equally to all of the three aforementioned styles, it is SaaS which usually commands the most visibility from a business perspective and may even become part of an organization's value chain while abstracting the intricacies of foundational IaaS and PaaS services.

## So, what is self-service integration?

For a handful of years now we've been hearing about the rise of the "Citizen Developer", which can be any person with no professional software development training who nonetheless is able to ideate and build applications for the use of others by exploiting an IT-provided toolkit consisting of elements such as:

- Low-code frameworks
- Responsive UX / UI
- Natural language recognition
- Built-in assistance
- Etc...

It seems quite logical then, under the same line of thinking, to discuss as well the emergence of a "Citizen Integrator". This human consumer of technology would have the ability, provided the right set of tools, to purposefully and seamlessly connect the dots and pull together a diverse set of intrinsically interoperable Cloud-based services as a means of realizing and/or enhancing business-driven scenarios requiring integration. All of this of course, without having to write a single line of code.

Let's consider possible use cases for self-service integration: We have already mentioned the impact and appeal SaaS

can have towards the business. Taking a closer look, we'll see that organizations with meaningful investments in technology tend to have an extensive portfolio of services and applications, ranging from utilitarian apps such as Slack, Dropbox, OneNote to mission critical solutions like JIRA, Oracle Sales Cloud or Eloqua.

Such a scenario means that integration may be needed at any moment in order to provide key capabilities stemming from the combination of two or more of these tools. Moreover, it needs to happen as soon as possible and without much hassle; so take for example:

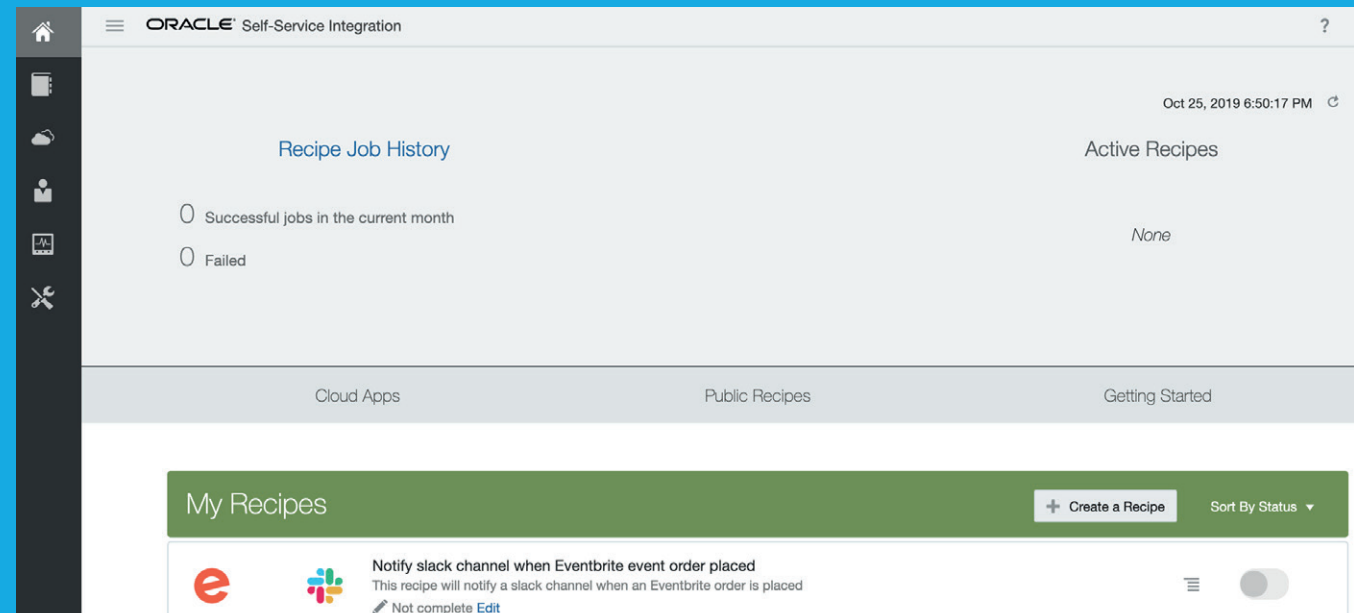
- A project manager who wants to enhance communication by automatically broadcasting the status and outcomes of certain JIRA tasks to a Slack channel.
- A sales VP who has identified that Eventbrite registrations for a targeted marketing event should trigger the creation of leads on Oracle Sales Cloud, thus facilitating her team's work and leveraging the marketing team's efforts.
- A marketing campaign leader tasked with following through on multiple opportunities who needs a constant flow of updated information between Oracle Sales Cloud and multiple other recipients such as Eloqua, Google Drive, Asana, Email etc. in order to be effective.

All of these roles seem like great candidates for becoming a “Citizen Integrator”, if they could only realize their intentions with no help or intervention at all from the IT department.

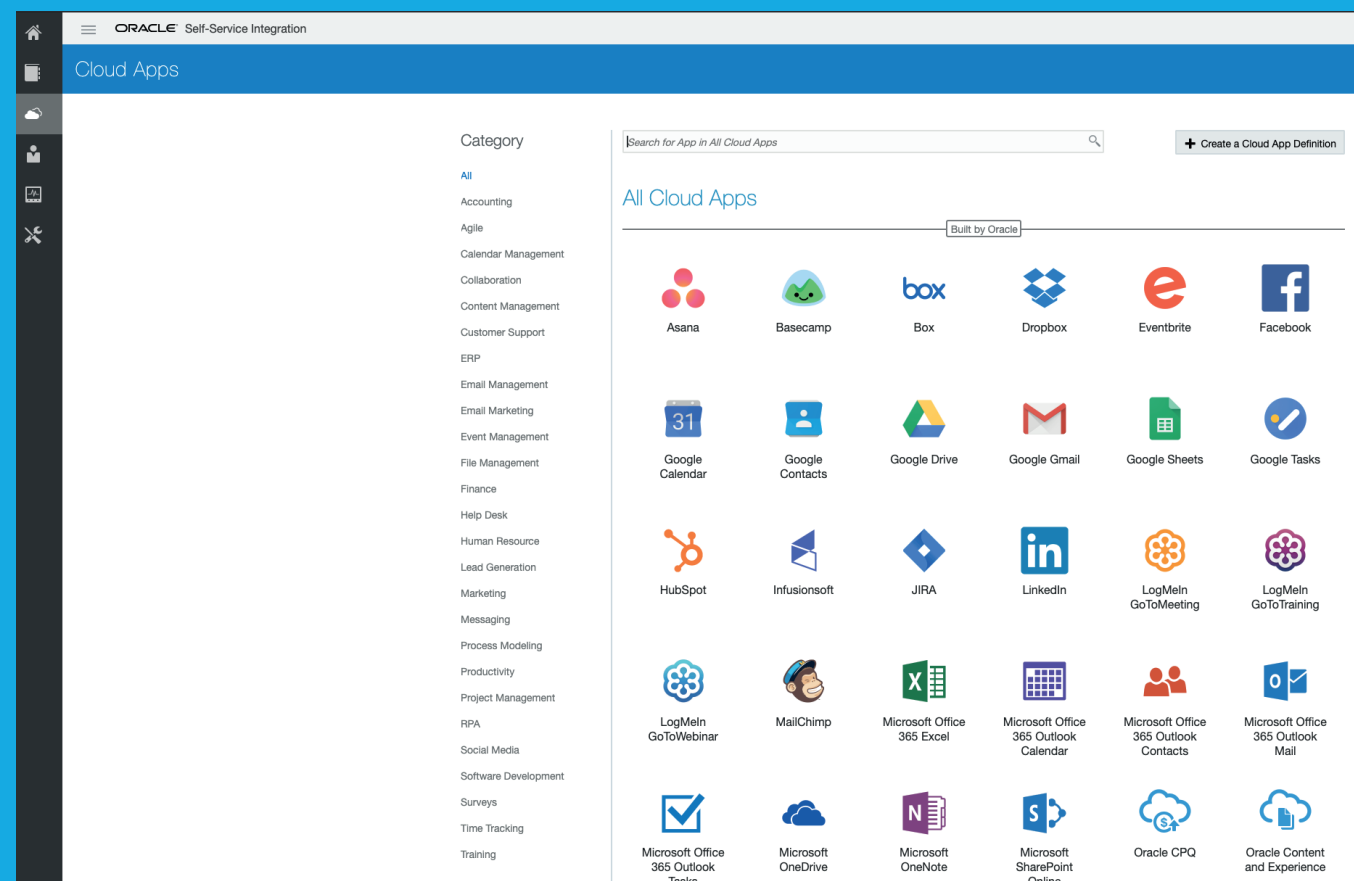
So is there any option out there yet with the ability to provide this kind of experience? The answer is definitely yes, as self-service integration has surged strongly into the landscape with several cloud providers supporting it by means of their own offerings (e.g. Oracle SSI, IFTTT, Jitterbit, Mulesoft).

Let’s take a look at the Oracle Self-Service Integration cloud service, which works with a mix of Oracle SaaS and common third-party apps.

The self-service portal (see *screenshots 1 and 2*) has several facets to it, as users can potentially assume different roles ranging from citizen integrator to experienced developer depending on the task at hand.

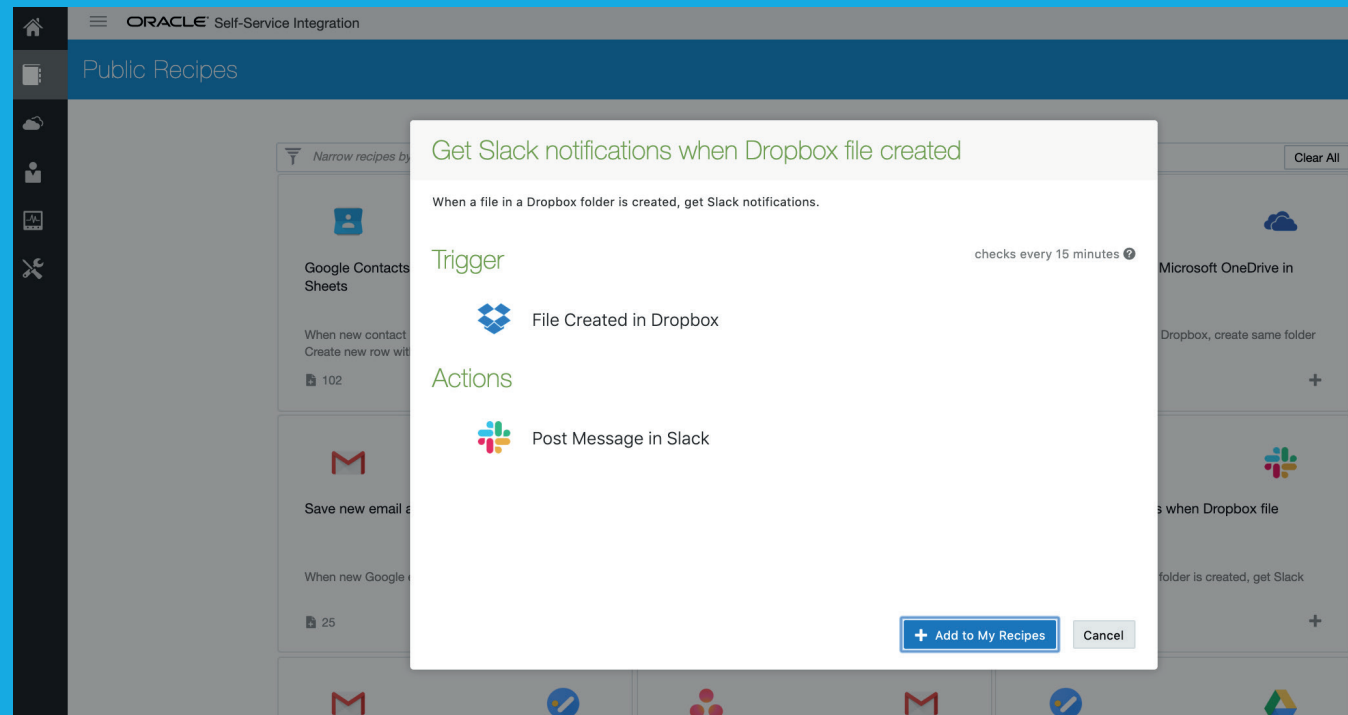


Screenshot 1



Screenshot 2



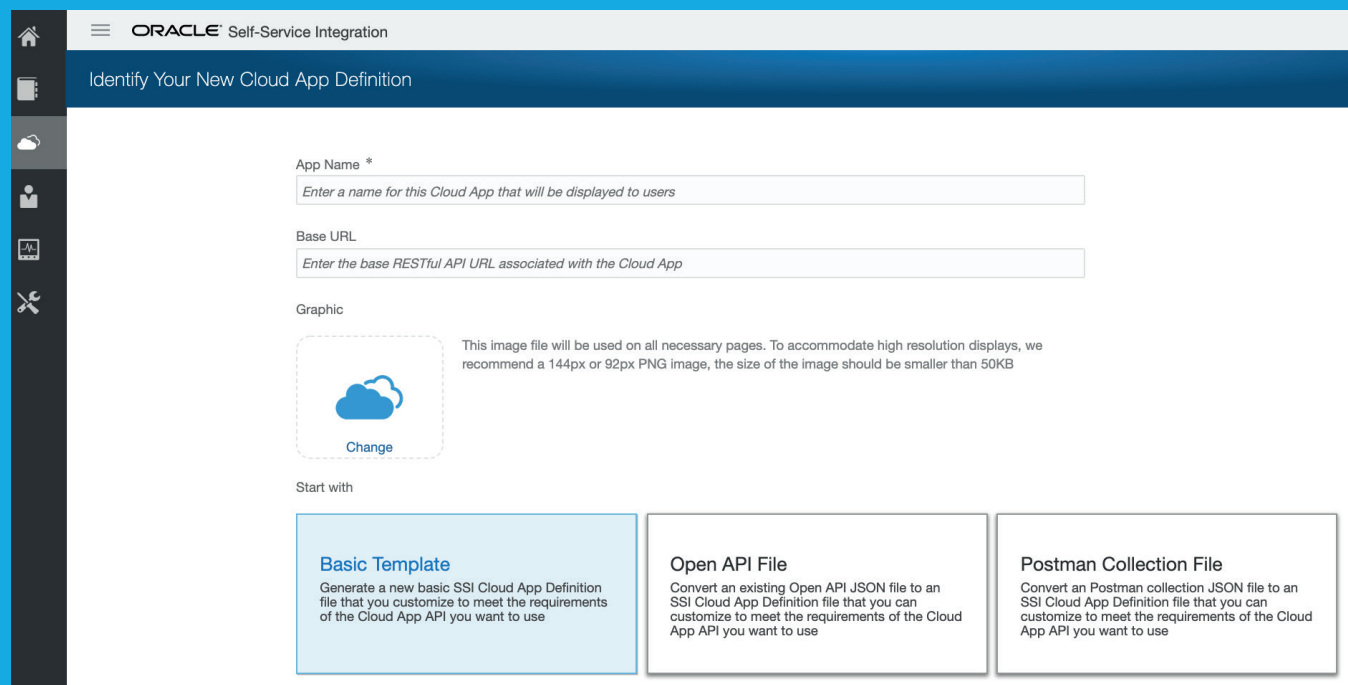


Screenshot 3

Basically, the idea is to have an ever growing catalog of Cloud App definitions which can then be freely combined into ready-made recipes.

For the citizen integrator it should be very simple to leverage existing public or private recipes by filling only a couple of blanks which don't require technical knowledge at all; for example, customizing the message that will be posted to Slack every time a file is uploaded into a given Dropbox folder (see *screenshot 3*).

Developers on the other hand can easily add new Cloud App definitions or update existing ones in order to extend the recipe book (see *screenshot 4*).



Screenshot 4

Other roles with varied levels of technical acumen such as curators, business analysts and IT managers, might also play a part within the self-service integration ecosystem.

It is important to point out that this kind of cloud feature is not really meant to be used for Enterprise Integration purposes (as would be an ESB, SOA, iPaaS for example). The resulting workflows normally have a very clear-cut shelf life in support of specific initiatives, projects or business requirements, which fosters flexibility and continuous improvement by making them quite disposable.

As a conclusion, we can say that use cases related to self-service integration definitely have a place in most organizations, particularly in those that are heavily invested in a diverse portfolio of SaaS solutions. Cloud vendors have also identified this niche and are actively looking to provide the needed features within their offering. The paradigm of empowering a “Citizen Integrator” by leveraging comprehensive recipe catalogs seems like the right one, as long as connecting the dots without the help or direct intervention of IT remains as simple as possible, while providing robust enough mechanisms for developers to extend the list of recipes and Cloud app definitions.



## About Arturo Viveros

Arturo Viveros is a Mexican IT Professional currently based in Oslo, Norway. He is a certified Cloud Integration Architect with over 14 years of experience in the design, development and delivery of software for a variety of customers and industries. Arturo is also a Groundbreaker Ambassador, Oracle ACE and a published technology writer both in English and in Spanish; he strives constantly to be involved with and support Developer Communities / User Groups focusing on technologies such as Oracle, Cloud, Java, DevOps, Software Architecture, Open Source, Blockchain etc. In his spare time Arturo enjoys family life, reading, hiking, travelling, playing guitar and practicing sports such as tennis, football and skiing.

# Oracle Application Express (Part 1): Low Code with the Database

*Carsten Czarski*





## Introduction

The Low Code approach, which is there to speed up application development, becomes more and more popular among IT landscapes. As the name indicates, low code applications are built by putting together standardized application components, like charts, forms or tabular reports. Low code does not mean “no code”: Some coding, to adopt to specific requirements, is needed in almost every case. However, the developer can concentrate on the actual business logic, while the low code framework takes care of low-level implementations.

Application Express (APEX) is the low code platform delivered with the Oracle Database. As such, it's strongly focused on integrating with the RDBMS. Data for application components typically comes from executing a SQL query – and since APEX runs directly in the database, this comes at almost no additional cost: APEX provides **zero-latency data access**. However, since version 18.1 APEX can consume data from external sources, like REST services, as easy as local database tables.

APEX is part of Oracle's Autonomous Database Cloud service; after provisioning an **Autonomous Data Warehouse** or **Autonomous Transaction Processing** instance, APEX is immediately available – all required installation and setup is already done. These cloud services are **managed databases**, so database administration is done by the Oracle Cloud – developers can focus on their APEX application development.

## Architecture

APEX is implemented in Oracle PL/SQL and runs completely within the Oracle Database. The database is front-ended by Oracle REST Data Services (ORDS), which maintains a pool of database connections and dispatches browser requests to the

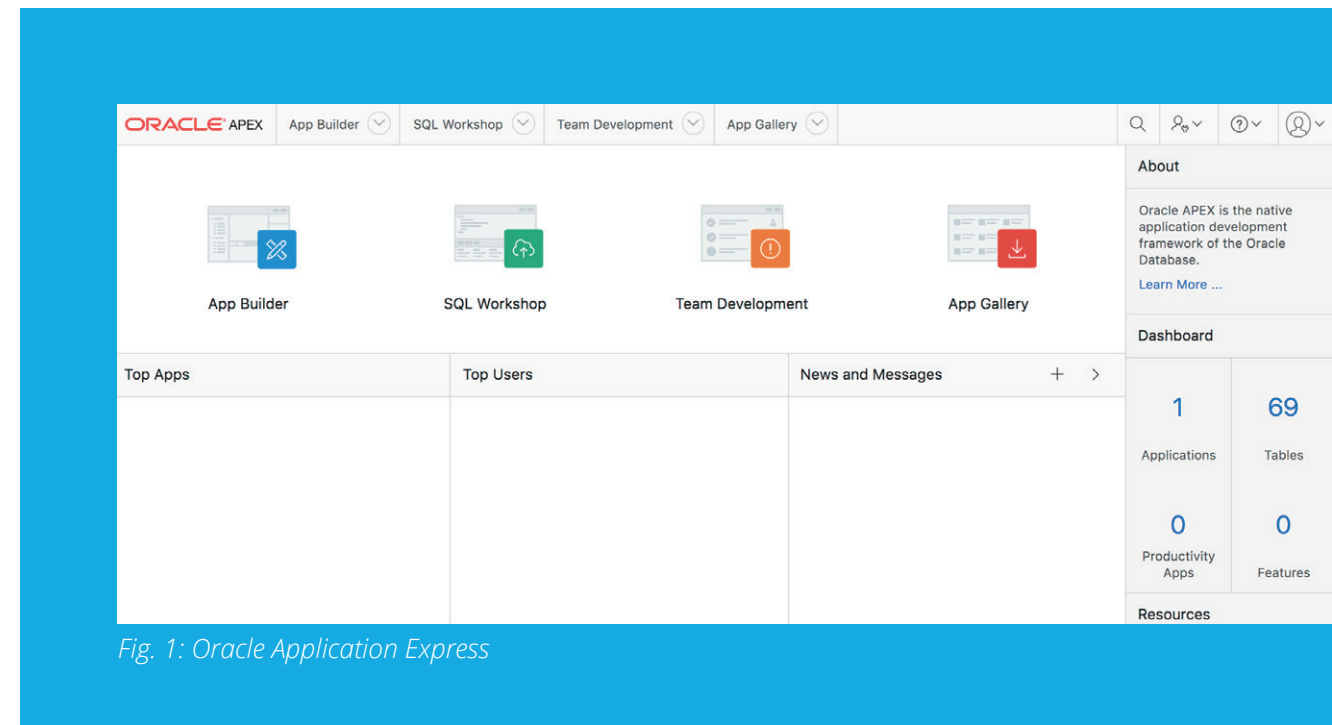


Fig. 1: Oracle Application Express

database. All request processing and page rendering happens completely in the database.

APEX comes with a multitenant architecture out of the box. A **workspace** is a logical area in APEX; assigned to one or multiple database schemas and isolated from all other workspaces.

Each workspace has its own user and application management; and workspace administrators can maintain their workspaces without needing the central instance administrator.

One of the most important aspects of APEX is, that it is completely meta data driven: All parts of an APEX application (pages, regions, processes etc.) and their attributes are stored



Rows

10

▼

?

Clear Command

Find Tables

Save

Run

```
select application_id, application_name, page_id, page_name
  from apex_application_pages
 order by 1,3
```

Results

Explain

Describe

Saved SQL

History

APPLICATION_ID	APPLICATION_NAME	PAGE_ID	PAGE_NAME
10683	Sample Trees	1	Home
10683	Sample Trees	2	Manage Sample Data
10683	Sample Trees	3	Project Tracking

Fig. 2: Querying the APEX Data Dictionary

as table rows in the database. When a page is rendered, the APEX engine (written in PL/SQL) reads this meta data and generates HTML markup or JSON output to render application pages. When developers build or maintain their applications, all changes are processed as changes to the rows in APEX meta data tables. **APEX never generates code.**

This approach has several advantages: One is, that meta data does not store **how** something is implemented, but rather **what** is implemented. APEX demonstrated this greatly for chart regions on application pages: For instance, an APEX application contains a **Bar Chart** on some data returned by a SQL query. When the first version of APEX came out back in 2004, these charts were either pure HTML or SVG charts (in these times, a separate browser plugin was required to display SVG).

As charting technologies evolved over time, APEX provided different ways to display this chart: The AnyChart library rendering Flash charts was introduced – application meta data remained unchanged. AnyChart started to support HTML5 – no change to the meta data. Finally, APEX moved to Oracle's own JET (JavaScript Extension Toolkit) for chart display. Application meta data can still remain unchanged, or can be upgraded automatically. Code-based implementations would likely have required new implementations.

Another great advantage is that APEX provides access to all application meta data in form of data dictionary views. Just by using SQL, a developer can query APEX for all sorts of information about their applications, which can be a great information source for QA purposes (*fig. 2*).

Fig. 3: Create a new application in APEX: Create Application Wizard

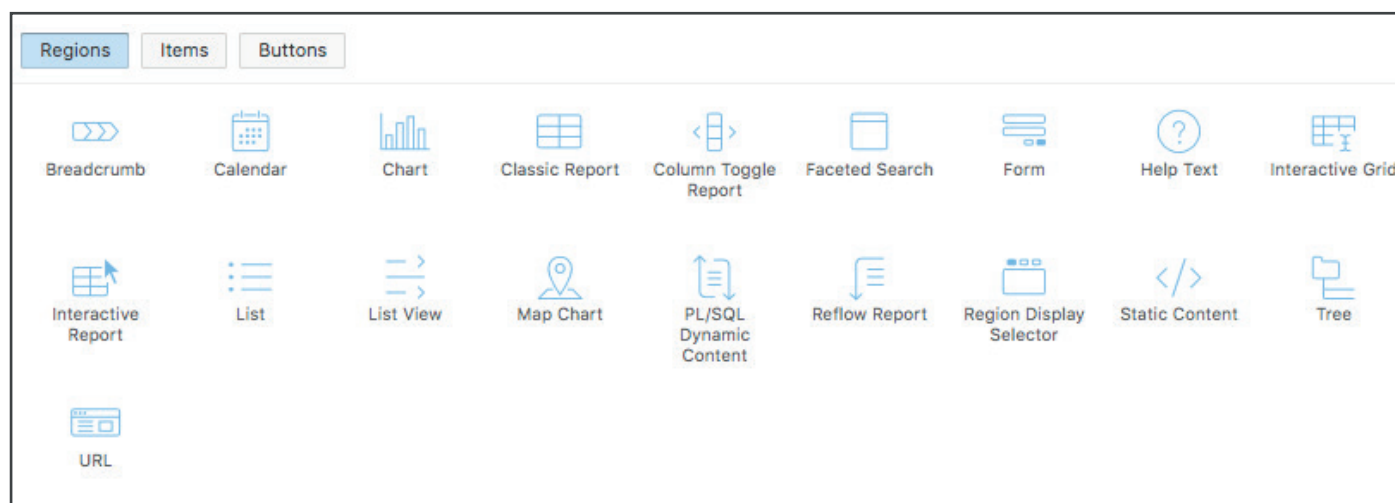


Fig. 4: Available page region types in APEX Page Designer

## Building applications

A new application is built using the Create Application Wizard (fig. 3). After a name has been provided and a set of initial pages and application features has been configured, the application is being created – and is ready to run immediately.

As already stated, the application was “created” by inserting meta data rows into APEX engine tables. Since pages are rendered based on the meta data at runtime, the app can be started immediately, without any compilation process. Also, every attribute of the application can be changed any time – unlike for code generators, there is never a “point of no return”, from which the generated application cannot be changed any more.

APEX provides a variety of components to use on a page. Examples include:

- **Classic Report:** visualize data in tabular form or based on a template
- **Interactive Report:** tabular display of data, end users can interactively filter
- **Interactive Grid:** an editable data grid
- **Forms:** a form page with form items and buttons to save changes
- **Calendar**
- **Tree**
- **Charts** (bar, line, pie, and many other visualization options)

Since version 4.0 in 2010, APEX components are extensible. Developers can create new region, item or other component types by implementing a *Plug-In interface*. Until today, the developer community provides a large collection of plug-ins on the community website [apex.world](https://apex.world).

The latest Application Express release 19.2 introduces the new *Faceted Search* Component (shown in fig. 5). Faceted Search is typically known from sales sites on the internet: on the left, a list of *Facets* is displayed – the end user can pick a brand, color, price or other attribute. On the right, or below, the page shows relevant data, matching the filter criteria from facets. APEX 19.2 allows to create faceted search pages simply by picking a database table. APEX will analyze the data and propose columns to create facets for (fig. 6).

Faceted Search again demonstrates the value of a meta data driven framework. Meta data stores information about faceted search: the table to use, the facets with their assigned columns and Lists of Values, and details on visualization of the actual data. APEX executes the required SQL queries to fetch data and item count information automatically – and makes sure that page regions are properly refreshed once the end user changed their query criteria. Application developers can focus on their business requirements.

### Other Application Components

An application consists of more than just page regions, of course. It requires business logic, authorizations, authentication, validation of end user input and much more. APEX provides all these things as components.

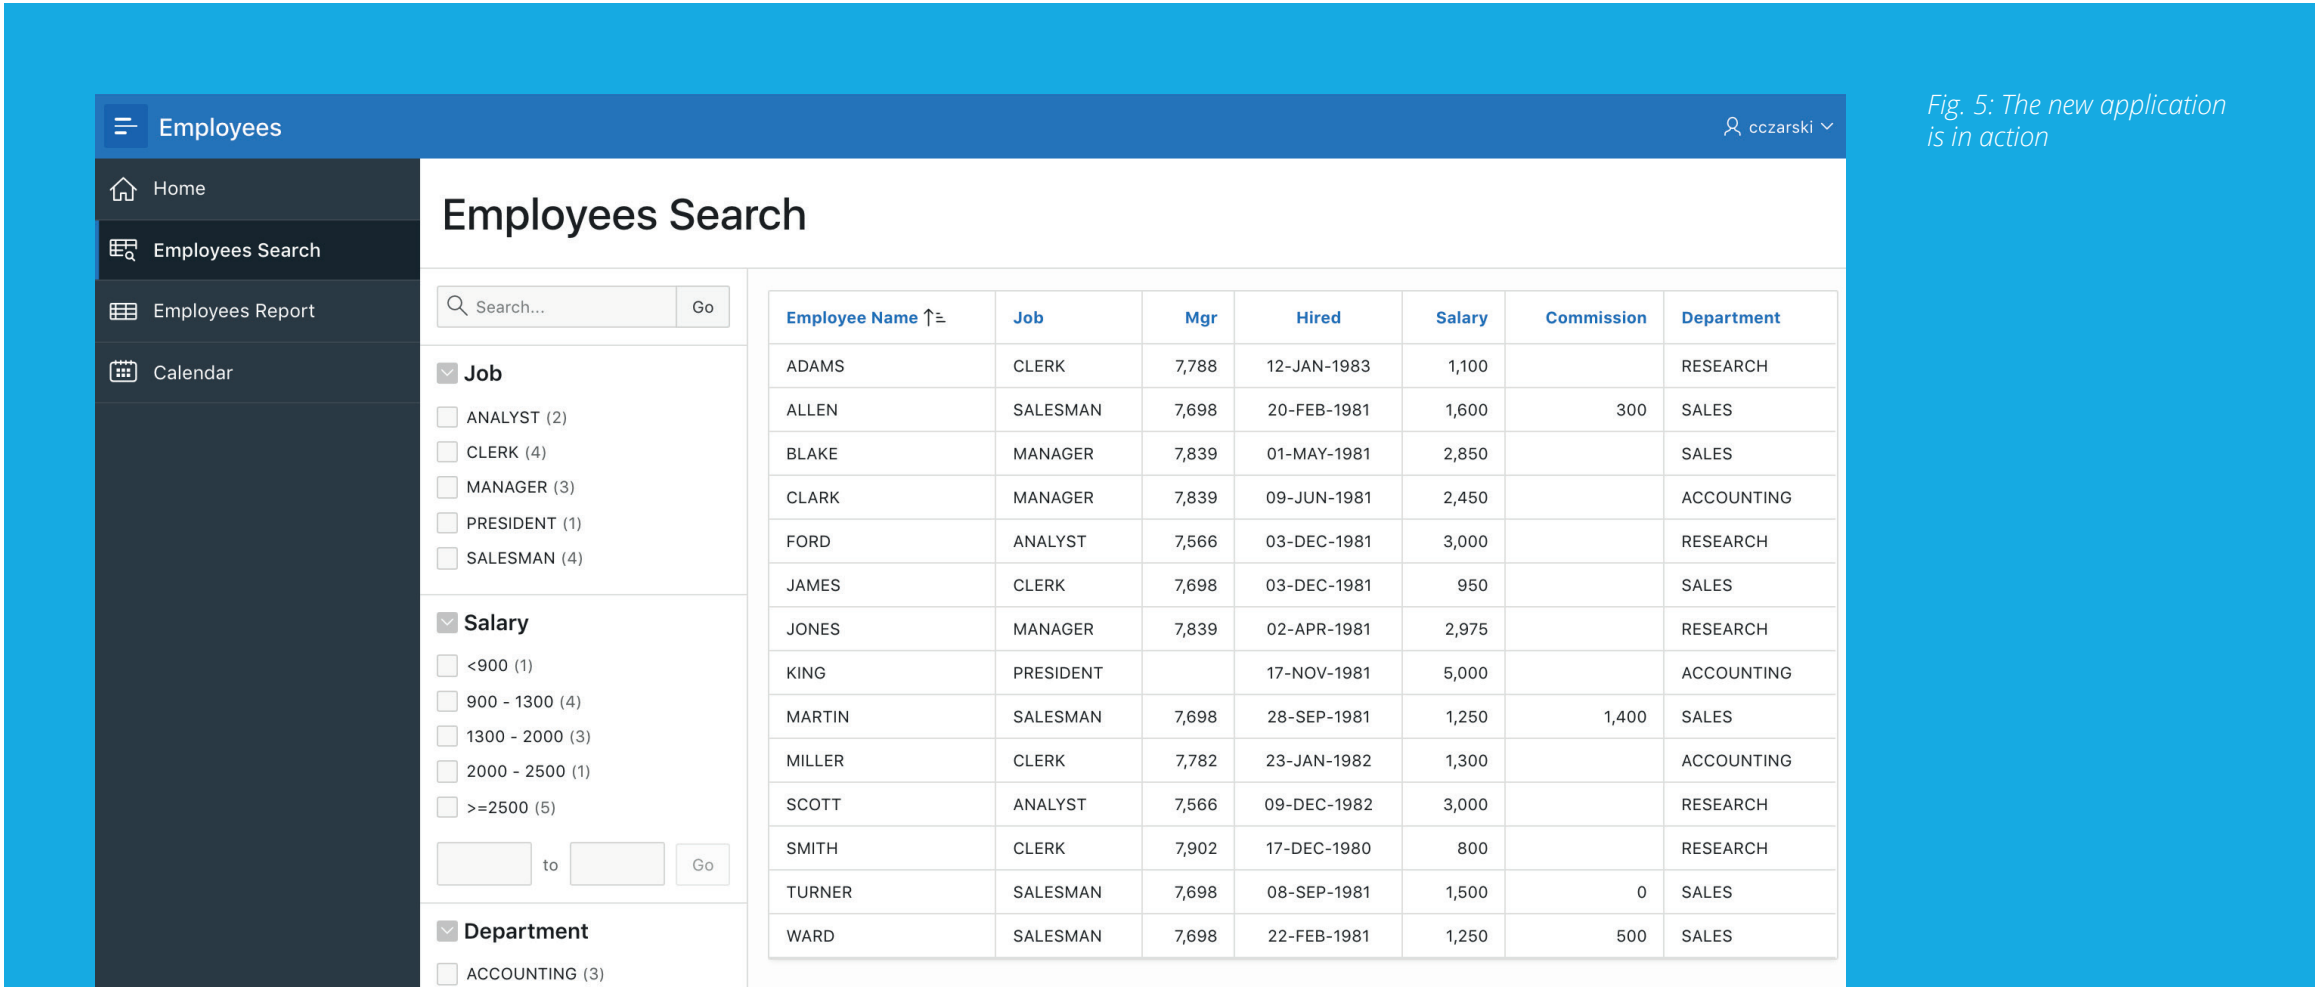


Fig. 5: The new application is in action

### Create Faceted Search

Local Database REST Enabled SQL Service Web Source

Table SQL Query

Table / View Owner: CCZARSKI

Table / View Name: AIRBNB (table)

Report Cards

Select the facets you want to include in this report. Note that facet recommendations are based on statistics of data in your table. You can refresh the column statistics to provide better recommendations.

	Column	Facet	Searchable	Nulls	Distinct	Average Length
	ID (number)		No	0	7,833	6
	HOST_NAME (varchar2)		Yes	1	2,132	8
<input checked="" type="checkbox"/>	NEIGHBORHOOD (varchar2)	Checkbox	Yes	0	36	15
<input checked="" type="checkbox"/>	ROOM_TYPE (varchar2)	Checkbox	Yes	0	3	15
<input checked="" type="checkbox"/>	PRICE (number)	Range	No	0	517	4

Fig. 6: Create a Faceted Search page in APEX

The following list enumerates a few examples for such components – APEX offers many more.

- **Authentication Schemes** control *how* end users log into the application. APEX provides LDAP, Single Sign-On, Database Users and others by default.
- **Authorization Schemes** control *what* the authenticated user is allowed to do. These are implemented similar to a

“Is the User allowed to ...?” question and then assigned to the individual component. At runtime, these will be evaluated based on the current user – and components show up or not.

- **Processes** can execute all sort of procedural business logic. As APEX runs in the database, PL/SQL is the first choice for implementing processes, especially when interacting with tables, SQL and data.
- **Themes and Templates** control page layout and how components are rendered.
- **Lists of Values** are named entities which provide data in form of Display Value and Return Value. Typical use cases are select lists, radio groups or checkbox items in forms.

## APEX Releases

The APEX development team plans to release new APEX versions twice a year: APEX 18.1 came out in May, 18.2 in September 2018. APEX 19.1 was released in March 2019 and the next APEX release 19.2 is scheduled to be released around the time this article is published. The following new features are planned:

- Faceted Search
- Modern Popup List of Values (LOV):  
A modern, feature-rich UI for Popup LOVs
- List of Values on top of external data sources:  
LOV can use REST Services as their data source
- New built-in Team Development application:  
APEX provides an application to track issues in small development teams
- New form input items:  
Markdown Editor and Star Rating
- Editable Interactive Grid support external data sources



## APEX Community and developer resources

APEX has a very active developer community. Watching the **#orclapex** hashtag on Twitter, is a common way for APEX developers to stay up to date. The apex.world website is maintained by neither Oracle nor an Oracle User Group, but purely by the developer community. apex.world provides access to a repository of community-driven APEX plug-ins, a Slack discussion channel, a job exchange and many more. Other useful resources for APEX are:

- Oracle OTN discussion forum for questions and deeper discussions  
[https://community.oracle.com/community/database/developer-tools/application\\_express/content](https://community.oracle.com/community/database/developer-tools/application_express/content)
- apex.oracle.com evaluation server. Everybody can request a free workspace to play with APEX on this instance. Also, the home page of apex.oracle.com contains a huge collection of tutorials, presentations and other collateral.  
<https://apex.oracle.com>
- The tryapexnow.com evaluation instance is launched a few weeks before a new APEX version is released. People can try out new features and provide feedback to the development team.  
<https://tryapexnow.com>



## 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.



# Running Your JuPyTeR Notebooks on the Cloud

*Mani Sarkar*





## Introduction

**JuPyteR notebooks** gives us the facility to write notebooks using Java, Scala and other JVM and non-JVM languages in addition to Python, R and Julia. We'll try to do this on **Oracle Cloud Infrastructure (OCI)**. If you are new to Oracle Cloud, I would suggest getting familiar with the **docs** and **Getting Started** sections of the docs.

I found the **tutorial to setup a VM instance** simple and useful – I recommend having a glance and following the steps. Take note of the pre-requisites before actually creating a VM instance and **ssh**-ing into it – it will involve Creating Compartments, Subnets, Security Lists, among other things before you can create a VM.

## Signing up

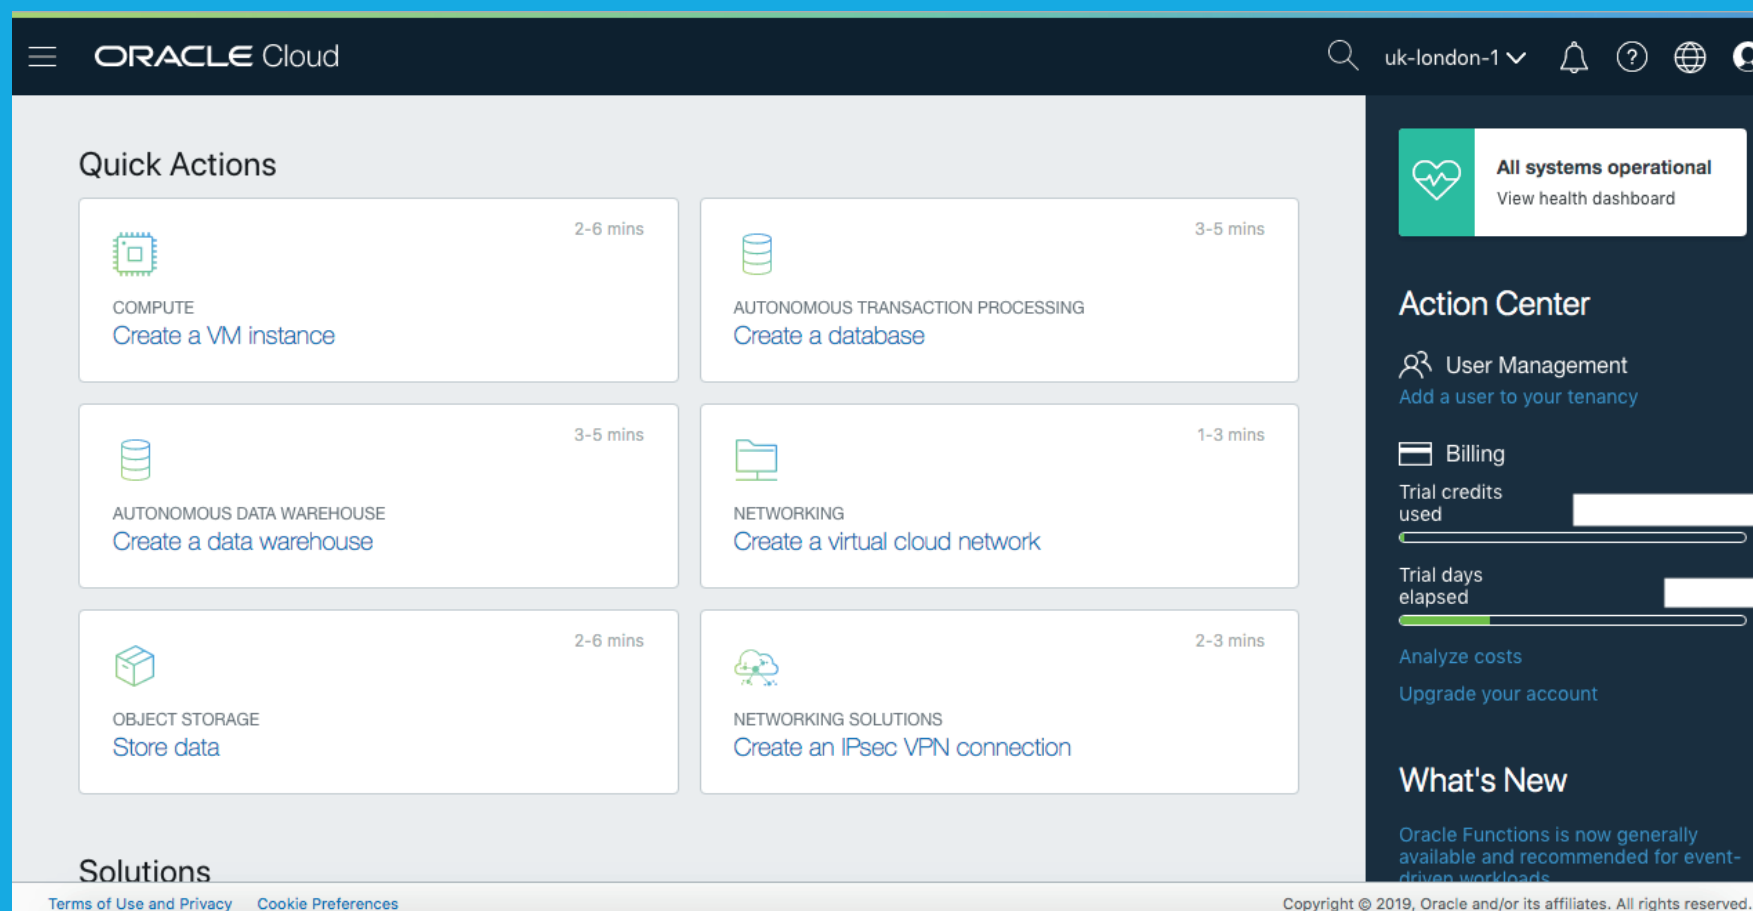
You will have noticed you have to have an account to be able to get access to the Cloud Dashboard and proceed.

You can sign up by going to **oracle.com** and also to **cloud.oracle.com** – I recommend signing up via these portals. You might even be eligible for **FREE** credits once you do that (enough to spend your weekend running your favorite instances).

## Setup

### Dashboard – sign-in

Once you are signed up, you sign-in via **cloud.oracle.com/sign-in** which will take you to a page like this:



Follow the instructions as described in the [tutorial to setup a VM instance](#) and give your VM and other resources names (use initials as a prefix) you can identify easily. This will kick off the request to create the VM (if all your entries are valid) – and in less than 15 seconds you should have a VM ready to be used.

Once the VM instance is created, make a note of the **Public IP Address** of the instance. All running VMs can be found by going to the **Compute > Instance** on the navigation menu on the left. Select the running VM by clicking it which will take you to the VM details page, where you can spot the **Public IP Address**:

The screenshot shows the Oracle Cloud console interface. At the top, the header includes the Oracle Cloud logo, a search bar, and the region 'UK South (London)'. The main content area displays the details for a VM instance named 'instance-20190912-2037'. On the left, there is a green square icon with a white 'I' and the word 'RUNNING' below it. To the right of the icon are buttons for 'Start', 'Stop', 'Reboot', 'Move Resource', 'Apply Tag(s)', and an 'Actions' dropdown menu. Below these buttons are two tabs: 'Instance Information' (selected) and 'Tags'. The 'Instance Information' section contains the following details:

- Availability Domain:** QsiE:UK-LONDON-1-AD-1
- Fault Domain:** FAULT-DOMAIN-1
- Region:** uk-london-1
- Shape:** VM.Standard2.1
- Virtual Cloud Network:** [my-vcn](#)
- Maintenance Reboot:** -
- Image:** [Canonical-Ubuntu-16.04-2019.08.14-0](#)
- OCID:** ...q6xjdq [Show](#) [Copy](#)
- Launched:** Thu, 12 Sep 2019 19:38:39 UTC
- Compartment:** neomatrix369gba (root)/Sandbox
- Launch Mode:** NATIVE

Below the instance information is the 'Primary VNIC Information' section, which includes:

- Private IP Address:** 10.0.0.2
- Public IP Address:** 132.145.78.136 (highlighted with a red box)
- Internal FQDN:** instance-20190912-2037... [Show](#) [Copy](#)
- Subnet:** [Public Subnet QsiE:UK-LONDON-1-AD-1](#)

At the bottom of the page, there are links for 'Terms of Use and Privacy' and 'Cookie Preferences', and a copyright notice: 'Copyright © 2019, Oracle and/or its affiliates. All rights reserved.'

*Note: the Public IP Address will be different for every VM created, the above is temporary.*



## CLI

OCI can be accessed with a command-line tool called **oci-cli** which can be [installed by following instructions](#) mention in the [CLI docs](#). Once installed the command to invoke it is called **oci** and you can invoke it by doing the below:

## Actions to get on the cloud

```
$ oci -help

Usage: oci [OPTIONS] COMMAND [ARGS]...

Oracle Cloud Infrastructure command line interface, with support
for Audit, Block Volume, Compute, Database, IAM, Load Balancing,
Networking, DNS, File Storage, Email Delivery and Object Storage
Services.

Most commands must specify a service, followed by a resource type
and then an action. For example, to list users (where $T contains
the OCID of the current tenant):

oci iam user list --compartment-id $T

Output is in JSON format.

For information on configuration, see
https://docs.cloud.oracle.com/Content/API/Concepts/sdkconfig.htm.
Options:
<-- snipped -->

Commands:
<-- snipped -->
```

As such we won't need the dashboard for the most part here onwards. We will also NOT be covering the use of the CLI tool in this post.

## Logging into the VM instance

You can then **ssh** into the box (see [docs on connecting via ssh](#)) and proceed with rest of the actions below:

```
#### Oracle Linux and CentOS images, user name: opc
#### the Ubuntu image, user name: ubuntu

$ ssh -i ~/.ssh/id_rsa ubuntu@132.145.78.136

or

$ ssh ubuntu@132.145.78.136
```

## Installing git

We selected the **Canonical Ubuntu Linux** (*Canonical-Ubuntu-16.04-2019.08.14-0*) as our OS image, which comes with **apt-get** and **git** installed so we don't need to do anything there.

## Running JuPyteR Notebooks

### Cloning our repo

We can clone our repo and perform the rest of the steps:

```
$ git clone https://github.com/neomatrix369/awesome-ai-ml-dl
$ cd awesome-ai-ml-dl/examples/JuPyteR
```

## Installing Docker

The Docker docs for installing Docker on Ubuntu can be found on [the Docker site](#). A bash-script has also been provided to quicken the process, although the target OS here is Ubuntu 16.04 or higher:

```
$ cd build-docker-image
$ ./installDocker.sh
```

*Note: in case you choose another OS image during VM creation, you will have to install Docker manually with the docs from Docker or modify the above script to make it work for the target OS.*



## Building the JuPyteR Docker Image

```
$ cd build-docker-image
$ sudo ./buildDockerImage.sh
```

In this specific environment, we need to pass the **sudo** keyword before every docker command. You may not have to do that in your local environment or elsewhere.

## Running the JuPyteR notebook as a Docker container

```
$ cd [back into the project root folder]
$ sudo ./runDockerContainer.sh
```

This will show you a console like this:

```
<--- snipped --->

OpenJDK Runtime Environment (build 9.0.4+11)
OpenJDK 64-Bit Server VM (build 9.0.4+11, mixed mode)
PATH=/home/jupyter/.local/bin:/opt/java/openjdk/bin:/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin

~~~ JDK9, Linux only: We are enabling JVMCI flags (enabling Graal as Tier-2 compiler) ~~~
~~~ Graal setting: please check docs for higher versions of Java and for other platforms ~~~
JAVA_OPTS=-XX:+UnlockExperimentalVMOptions -XX:+EnableJVMCI -XX:+UseJVMCICompiler
JAVA_TOOL_OPTIONS=-XX:+UnlockExperimentalVMOptions -XX:+UseGroupMemoryLimitForHeap -XX:+UnlockExperimentalVMOptions -
XX:+EnableJVMCI -XX:+UseJVMCICompiler

Available kernels:
  python2    /home/jupyter/.local/share/jupyter/kernels/python2
  java       /usr/share/jupyter/kernels/java

[I 13:39:35.993 NotebookApp] Writing notebook server cookie secret to /home/jupyter/.local/share/jupyter/runtime/notebook_cookie_secret
[I 13:39:36.293 NotebookApp] Serving notebooks from local directory: /home/jupyter
[I 13:39:36.294 NotebookApp] The Jupyter Notebook is running at:
[I 13:39:36.295 NotebookApp] http://(81dde8675279 or 127.0.0.1):8888/?
token=bb0c81ef7e9f3932355b953163702aa2d9f75e18005e6e48
[I 13:39:36.297 NotebookApp] Use Control-C to stop this server and shut down all kernels (twice to skip confirmation).
[C 13:39:36.310 NotebookApp]

To access the notebook, open this file in a browser:
    file:///home/jupyter/.local/share/jupyter/runtime/nbserver-28-open.html
Or copy and paste one of these URLs:
    http://(81dde8675279 or 127.0.0.1):8888/?
token=bb0c81ef7e9f3932355b953163702aa2d9f75e18005e6e48
```

Make a note of the URL, and replace the `127.0.0.1` with your *Public IP Address* i.e. `132.145.78.136`.

You can also see from the above logs we are using Java 9 (built on the [AdoptOpenJDK](#) farm) and enabling the [GraalVM compiler](#) as HotSpot's [C2 compiler](#) (see [Switches to enable the GraalVM compiler in Java 9](#) to enabled the GraalVM compiler). It's also because the [Java extension for JuPyteR](#) requires Java 9 or higher to work.

**Opening the JuPyteR notebook in your browser**

Go to the browser and try to open this: <http://132.145.78.136:8888/?token=bb0c81ef7e9f3932355b953163702aa2d9f75e18005e6e48>

Aargh! It does NOT work! That is because we haven't opened up the port 8888 from within our cloud network (via *Ingress Rules*, read more about it [here](#)) to the outside world (public):

ORACLE Cloud

uk-london-1

Resources

Ingress Rules (4)

Egress Rules (1)

Ingress Rules

Add Ingress Rules

Edit

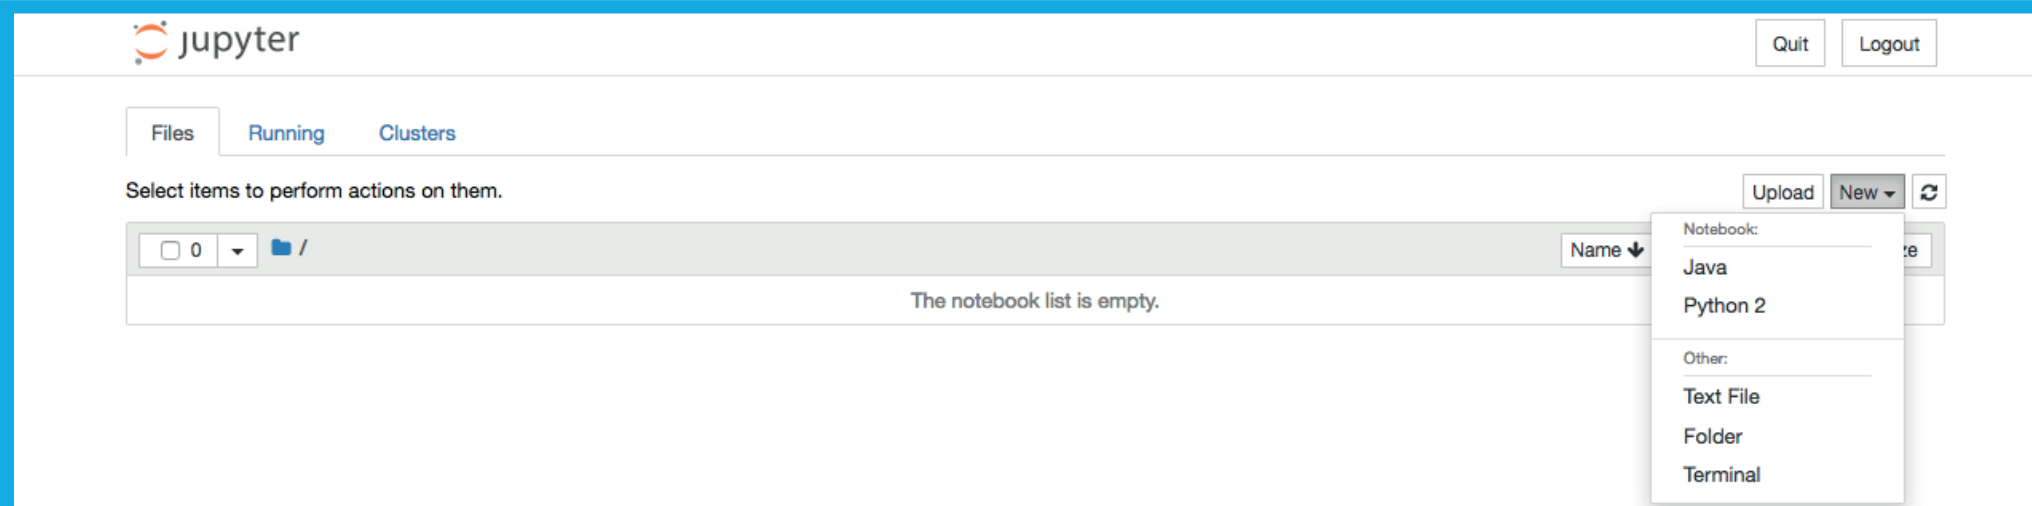
Remove

	Stateless	Source	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows
<input type="checkbox"/>	No	0.0.0.0/0	TCP	All	22		TCP traffic for ports: 22 SSH Remote Logi n Protocol
<input type="checkbox"/>	No	0.0.0.0/0	ICMP			3, 4	ICMP traffic for: 3, 4 Destination Unreach able: Fragmentation Needed and Don't Fr agment was Set
<input type="checkbox"/>	No	10.0.0.0/16	ICMP			3	ICMP traffic for: 3 De stination Unreachabl e
<input type="checkbox"/>	No	0.0.0.0/0	TCP	All	8888		TCP traffic for ports: 8888

We would need to add the above entry to the **Ingress Rules** section, you can get to **Ingress Rules** page via the navigation menu: **Networking > Virtual Cloud Networks > Virtual Cloud Network Details > Security Lists**, which brings you to the page with the **Default Security Lists**. On clicking the Security List that corresponds to your **Virtual Cloud Network (VCN)** you will land on the above **Ingress Rules** page.

In case, you are still not able to find it, search for the term **security** using the search facility on any page in the Cloud Console (see the magnifying glass 🔍 at the top of the page). This will show you all the **Default Security Lists** and clicking on it will bring you to the **Ingress Rules** page above (you might have just one Security List entry). **Note: Ingress means traffic coming into the network/VM instance.**

Why port 8888, that's because we set it up like that in the docker scripts, have a look at the **sources** to find out why and how. Having done all of the above: voila! We see the JuPyteR startup page in the browser:



And you can see a Java-based notebook is available to play with! Try out <https://github.com/ligee/kotlin-jupyter/raw/master/samples/ScreenShotInJupyter.png> by creating a new Java notebook in the browser. You are free to create Python notebooks as well, not just Java ones – this is the beauty of JuPyteR notebooks.

## Create custom image for reuse

As we have been able to successfully run JuPyteR Notebook from inside a VM instance, we can save this image for future re-use share with others. Creating an image of the VM instance can be done via **Compute > Instances > Instance Details** in the navigation menu, and Create Custom Image from the Actions drop-down menu. Whilst in the process of creating a custom image, your



original VM instance is shut-down. This can take under a couple of minutes to complete depending on the size of the original VM instance. When successfully created, it becomes available among the list of Custom Images to choose from, the next time we go to create a new VM instance.

## Power-user

If all of this was piece of cake for you or you have survived without much hassle, then try out all the deep-dive stuff mentioned in the [README page here](#).

To be able to code in other languages in the JuPyteR environment all you need is a JuPyteR extension – it's only a matter of installing and configuring. You can learn all about this [here](#).

## Signing off

```
[I 12:02:58.743 NotebookApp] 302 GET /?token=c5c8814e7e6c7e69c4e9ebe8e40f028bded87b7fee3e8afa (2.24.105.122) 0.80ms
[W 12:04:11.908 NotebookApp] Forbidden
[W 12:04:11.908 NotebookApp] 403 GET /api/config/tree?_=1567944251210 (2.24.105.122) 1.22ms referer=http://132.145.79.76:8888/tree
[W 12:04:11.910 NotebookApp] Forbidden
[W 12:04:11.911 NotebookApp] 403 GET /api/config/common?_=1567944251211 (2.24.105.122) 1.11ms referer=http://132.145.79.76:8888/tree
[W 12:04:11.924 NotebookApp] Forbidden
[W 12:04:11.925 NotebookApp] 403 GET /api/terminals?_=1567944251212 (2.24.105.122) 1.22ms referer=http://132.145.79.76:8888/tree
[W 12:04:11.952 NotebookApp] Forbidden
[W 12:04:11.952 NotebookApp] 403 GET /api/kernelspecs (2.24.105.122) 1.12ms referer=http://132.145.79.76:8888/tree
[W 12:04:11.983 NotebookApp] Forbidden
[W 12:04:11.984 NotebookApp] 403 GET /api/sessions?_=1567944251213 (2.24.105.122) 1.18ms referer=http://132.145.79.76:8888/tree
[W 12:04:11.986 NotebookApp] Forbidden
[W 12:04:11.986 NotebookApp] 403 GET /api/terminals?_=1567944251214 (2.24.105.122) 0.98ms referer=http://132.145.79.76:8888/tree
^C[I 12:36:41.114 NotebookApp] interrupted
Serving notebooks from local directory: /jupyter-notebooks
0 active kernels
The Jupyter Notebook is running at:
http://(ad312d6886be or 127.0.0.1):8888/?token=c5c8814e7e6c7e69c4e9ebe8e40f028bded87b7fee3e8afa
Shutdown this notebook server (y/[n])? y
[C 12:36:44.300 NotebookApp] Shutdown confirmed
[I 12:36:44.301 NotebookApp] Shutting down 0 kernels
ubuntu@instance-20190908-1229:~/awesome-ai-ml-dl/examples/JuPyteR$
```

In case you have created a notebook, it gets saved in the sub-directory called jupyter/notebooks, you can retrieve this using *scp* from your local machine (see [here](#) on how to do that).

Make sure you have *signed out of both* the [oracle.com](#) and [cloud.oracle.com](#) login sessions, it's easy to forget one or the other. But before doing that please also have a look at the [Cleaning up of resources](#) page in the docs – you don't want your instance running forever while you are not looking at it!

## Conclusion

A good set of scripts (including Docker) and an easy-to-use cloud environment can help in many ways. In this case, enabling us to run a JuPyteR notebook instance that can be shared publicly or privately depending on your network security settings. The JuPyteR environment is flexible and allows extending functionality via configurations and extensions.

We didn't cover things like cloud security and partitioning of user instances. Please look into this further, if they are important to you. Please do ensure it meets the necessary levels of security for your application or use-case. Check out the docs on [Security on the OCI docs page](#) to learn more.

*This article is a condensed version of [Mani Sarkar's original blog post](#).*

```

SQL>
System altered.

SQL>
System altered.

SQL> oracle-control [rman2na on db-rman2na-a-r1] Killing off all non local dedicated servers, shared servers and dispatchers
Database closed.
Database dismounted.
oracle-control [rman2na on db-rman2na-a-r1] Oracle server is running
oracle-control [rman2na on db-rman2na-a-r1] the database is still up after 1 seconds
oracle-control [rman2na on db-rman2na-a-r1] Killing off all non local dedicated servers, shared servers and dispatchers
oracle-control [rman2na on db-rman2na-a-r1] Oracle server is running
oracle-control [rman2na on db-rman2na-a-r1] the database is still up after 3 seconds
oracle-control [rman2na on db-rman2na-a-r1] Killing off all non local dedicated servers, shared servers and dispatchers
oracle-control [rman2na on db-rman2na-a-r1] Oracle server is not running
oracle-control [rman2na on db-rman2na-a-r1] the database is DOWN! after 5 seconds
ORACLE instance shut down.
SQL> Disconnected from Oracle Database 11g Enterprise Edition Release 11.2.0.4.0 - 64bit Production
With the Partitioning, OLAP, Data Mining and Real Application Testing options
oracle-control [rman2na on db-rman2na-a-r1] Running ipcs -m to wait for shared memory segment de-allocation
oracle-control [rman2na on db-rman2na-a-r1] ## THIS CAN TAKE 10 Seconds Per 1 GBytes of SGA ##

--- Shared Segments -----
key          owner      perms      bytes      natch      status
0x00000000   root       0          2000       0
oracle-control [rman2na on db-rman2na-a-r1] stopped oracle
db-rman2na-a-r1.iad6$ █

```

Christian Luda

# Amazon Says „Bye-Bye Oracle”

*What you see in video above is the shutdown of the last Oracle database of Amazon’s Consumer Business and the subsequent celebration of the Amazon team.*

These visuals were part of a [blog post](#) from October 15, in which Amazon Web Services (AWS) announced the turn off. The Consumer Business division includes the services Amazon Prime, Alexa, and Kindle. According to the post, 75 petabytes of data from almost 7,500 Oracle databases had been migrated to various AWS database services like Amazon DynamoDB, Amazon Aurora, Amazon Relational Database Service (RDS), and Amazon Redshift.

Jeff Barr, AWS's Chief Evangelist and author of the blog post, had the following explanation for Amazon's decision: "Over the years we realized that we were spending too much time managing and scaling thousands of legacy Oracle databases. Instead of focusing on high-value differentiated work, our database administrators (DBAs) spent a lot of time simply keeping the lights on while transaction rates climbed and the overall amount of stored data mounted. This included time spent dealing with complex and inefficient hardware provisioning, license management, and many other issues that are now best handled by modern, managed database services."

There had been a lot of back and forth between the two companies in recent years. In December 2018, Oracle CTO

Larry Ellison had declared that "in terms of technology, there is no way that any normal person would move from an Oracle database to an Amazon database." At this year's OpenWorld, Ellison revealed Oracle's plans to take on AWS which includes partnerships with Microsoft and VMware. During his speech he highlighted Oracle's autonomous technology in comparison to Amazon: "AWS offers databases based on the type of application. AI has a different database, IoT has a different database. They are really not elastic and server-less; they are expensive. Oracle's database, with its autonomous operating system, ensures that organizations don't need manual processes to configure their systems."

Against this backdrop, it comes as no surprise that Amazon's blog post cannot do without malice: a graphic showing the advantages of the migration bears the file name "bye\_bye\_oracle\_1", while Jeff Barr elsewhere refers to Oracle as a "hand-managed database".

Especially since Oracle shifts its focus more and more towards the cloud, these will certainly not have been the last words exchanged between the two big companies.

# Ubuntu<sup>®</sup> Lisa Damerow and Its Circle of Friends





# ubuntu

Ubuntu® (also known as Ubuntu Linux) is a popular, free and open source Linux operating system. Since its first launch in October 2004, Ubuntu® gradually became more popular and is one of the most used Linux distributions nowadays. Let's take a look at the logo and where this exotic name actually comes from.

In Zulu, the word “ubuntu” is an African philosophy and stands for “humanity”. It is often translated to “I am because we are” or “humanity towards others” but is often used in the sense of “the belief in a universal bond of sharing that

connects all of humanity”. Which, in our opinion, makes it a very clever thing to name an open source project. They embraced this deeper meaning further by incorporating it into the logo as well. At first glance you might only see an orange circle with an intermitted white ring and some dots around it. Think again! This logo actually symbolizes three people holding their arms and creating a circle with it. According to the Ubuntu® website, this symbol, named “Circle of Friends”, represents freedom, collaboration, precision and reliability. The wordmark itself is a unique, custom designed font.

---

<https://de.wikipedia.org/wiki/Ubuntu>  
[https://en.wikipedia.org/wiki/Ubuntu\\_philosophy](https://en.wikipedia.org/wiki/Ubuntu_philosophy)  
<https://askubuntu.com/questions/42724/whats-the-meaning-of-the-ubuntu-logo-where-does-it-come-from>  
<https://design.ubuntu.com/brand/ubuntu-logo/>

# Oracle Exadata Database Machine: Decade of Innovation

*Venkata Ravi Kumar Yenugula and Mari Kupatadze*



# Introduction










Oracle Exadata Database Machine, or simply Exadata is an Oracle computer hardware and software optimized for running Oracle Database(s). The study of history of this product is important. It can tell you what decisions were made to improve its performance, reliability, availability and what made it the most popular Engineered System.

Historically, it was marketed exclusively as a Data Warehouse platform and was poor for OLTP. But Exadata made

optimizations for mixed workloads. OLTP workloads benefit from Flash memory, Smart Flash Cache and Smart Flash Log optimize response time for heavy writes.

## The Evolution of Exadata and Database platform leader since 2008

Oracle releases Exadata every twelve to eighteen months. The following is a brief overview of each Exadata generation:

	V1	V2	X2	X3	X4	X5	X6	X7	X8	V1 – X8 Growth
										
	Sep 2008 Xeon E5430 Harpertown	Sep 2009 Xeon E5540 Nehalem	Sep 2010 Xeon X5670 Westmere	Sep 2012 Xeon E5-2690 Sandy Bridge	Nov 2013 E5-2697 v2 Ivy Bridge	Dec 2014 E5-2699 v3 Haswell	Apr 2016 E5-2699 v4 Broadwell	Oct 2017 Xeon 8160 Skylake	Apr 2019 Xeon 8260 Cascade Lake	
Storage (TB)	168	336	504	504	672	1344	1344	1680	2.35 PB	14 X
Flash Cache (TB)	0	5.3	5.3	22.4	44.8	89.6	179.2	358	358 TB	64 X
CPU (cores)	64	64	96	128	192	288	352	384	384 cores	6 X
Max Mem (GB)	256	576	1152	2048	4096	6144	12288	12288	12 TB	48 X
Ethernet (Gb/s)	8	24	184	400	400	400	400	800	800 Gb/s	100 X
Scan Rate (GB/s)	14	50	75	100	100	263	301	350	560 GB/s	40 X
Max Read IOPS (M)	.05	1	1.5	1.5	2.66	4.14	5.6	5.97	6.57 M	132 X

The Only Engineered System provides Full-Stack Security at two levels:

- **Exadata Database Machine Security**

- Industry policing: Banks, Government, Retail, Telcos
- Advanced Intrusion Detection Environment (AIDE)
- Regular security scans
- FIPS 140-2 certification
- PCI-DSS compliance
- Data and network encryption
- Linux minimal distribution
- Secure erase
- System lockdown
- Live kernel patching

- **Oracle Database Maximum Security Architecture**

- Identity Management
- Transparent Data Encryption
- Network Encryption
- Database Vault
- Audit Vault
- Key Vault
- Database Firewall
- Virtual Private Database
- Label Security
- Data Redaction
- Data Masking & Subsetting

## Oracle Exadata Advantages Increase Every Year

Oracle Exadata Smart Software	Oracle Exadata Smart Hardware
<ul style="list-style-type: none"> <li>• Smart Scan</li> <li>• InfiniBand Scale-Out</li> <li>• Database Aware Flash Cache</li> <li>• Storage Indexes</li> <li>• Columnar Compression</li> <li>• IO Priorities</li> <li>• Data Mining Offload</li> <li>• Offload Decrypt on Scans</li> <li>• Network Resource Management</li> <li>• Multitenant Aware Resource Mgmt</li> <li>• Prioritized File Recovery</li> <li>• In-Memory Fault Tolerance</li> <li>• Direct-to-wire Protocol</li> <li>• JSON and XML offload</li> <li>• Instant failure detection</li> <li>• In-Memory Columnar in Flash</li> <li>• Exadata Cloud Service</li> <li>• Smart Fusion Block Transfer</li> <li>• Exadata Cloud at Customer</li> <li>• In-Memory OLTP Acceleration</li> <li>• Automatic Indexing</li> <li>• Autonomous Database</li> <li>• KVM Virtualization</li> </ul>	<ul style="list-style-type: none"> <li>• Scale-Out Servers</li> <li>• Scale-Out Storage</li> <li>• DB Processors in Storage</li> <li>• Unified InfiniBand</li> <li>• PCIe NVMe Flash</li> <li>• Tiered Disk/ Flash</li> <li>• Software-in-Silicon</li> <li>• 3D V-NAND Flash</li> <li>• 25 GigE Client Network</li> <li>• Hot Swappable Flash</li> <li>• Persistent Memory</li> <li>• RoCE Networking</li> </ul>



**V1:** The first Exadata debuted in October 2008 focusing on Data Warehouse workloads. It supported IORM (IO Resource Manager) for allocating IO bandwidth between several databases and database filtering operations were moved into cells. Flash storage was not included yet. Exadata V1 was available in Full and Half Rack sizes. The hardware vendor was HP.

**V2:** The second version of Exadata was released in 2009. The hardware vendor changed to Sun. By the time, Oracle announced that it was in process of attempting to acquire Sun Microsystems. Exadata was available in Full, Half and Quarter Rack sizes. This version added OLTP workload support using Flash storage. Hybrid Columnar Compression and Storage Indexes features were first introduced in V2.

**X2:** The third generation was introduced in 2010 and was labeled as X2-2 and X2-8. The Flash Logging feature and 10 GB/sec Ethernet connectivity were added in this edition. The Storage Expansion racks (racks that contain only storage servers) were introduced in 2011.

**X3:** The next editions of Exadata, X3-2 and X3-8 were released in 2012 with new Eight Rack support. The new Write-Back Flash Cache feature increased OLTP throughput by 20x.

**X4:** Oracle released X4-2 in 2013 and another model X4-8 in 2014. The well-known Capacity on Demand first was introduced in this edition.

**X5:** X5-2 and X5-8 were released in 2015. Adding one server at a time became possible. Virtualization, Database snapshots on Exadata storage, Fault Tolerant redundancy with In-Memory were introduced in this version. The High Performance Exadata storage servers were replaced

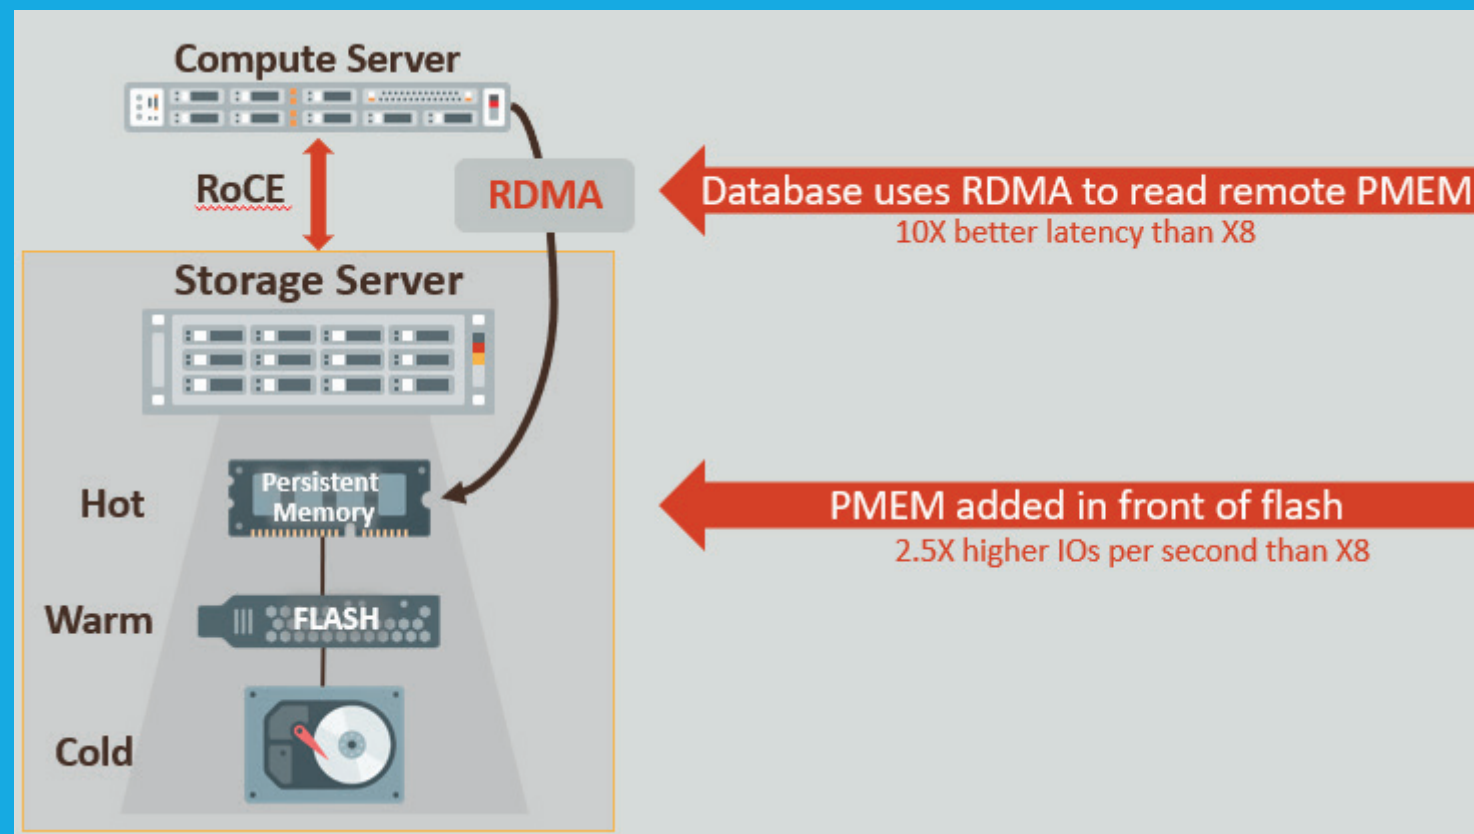
with Extreme Performance storage servers containing all NVMe disks. Oracle also introduced Exadata Cloud Service. Organizations were able to run Oracle database on Exadata without having a data center, cooling, network, power and so on via **Oracle Public Cloud**.

**X6:** In 2016, Oracle released X6-2 and X6-8. The new feature Exafusion Direct-to-Wire protocol was introduced. The disk controller batteries were replaced with super capacitors.

**X7:** In 2017, the X7-2 and X7-8 Exadata models were released. Flash cards became Hot pluggable. 25 GB/sec Ethernet connectivity was added and In-Memory processing was extended into flash storage. M.2 drives house operating system and cell software were also introduced. The Exadata storage server hard disks and flash drives contained application data only. M.2 drives were protected with Intel RSTe RAID.

**X8:** Oracle has released X8 in 2019. Cascade Lake Intel Xeon Platinum Processor 8200 Family comes with Increase CPU speed and Storage CPU count increased. X8 comes with Larger Capacity Drives, Larger local drives and X8 will ship with Cisco 9348.

**X8M:** Oracle has released X8M in 2019. Enhanced Internal Fabric with 100 GB RDMA over Converged Ethernet (RoCE), 1.5 TB Persistent Memory per storage server and Enhanced consolidation using Linux KVM. Storage Servers (HC/EF) includes Persistent Memory, therefore no memory expansion is required for Extreme Flash or High Capacity X8M-2 Storage Servers. RDMA over Converged Ethernet is a protocol that runs InfiniBand RDMA software on top of Ethernet. Oracle Exadata X8M Persistent Memory Data Accelerator delivers Extreme Performance and No Application Changes Required.



## Unique Smart Database Software

### Exadata Smart Analytics

- Move queries to Storage server, not Storage server to queries
- Automatically offload and parallelize queries across all Storage servers
- Extend In-Memory DB with flash
- Run In-Memory DB on standby

### Exadata Smart Storage

- Exadata Hybrid Columnar Compression (EHCC) reduces space usage by 10x
- Database-aware Flash Caching gives speed of flash with capacity of disk
- Exadata Storage Indexes eliminate unnecessary I/O

### Exadata Smart OLTP

- Special InfiniBand protocol enables 3x faster OLTP messaging
- Ultra-fast DB-optimized flash logging

### Exadata Smart Consolidation

- Critical DB messages jump to head of queue for ultra-fast latency
- CPU, I/O and network resources prioritized for end-to-end quality of service

## Monitoring complete system using Oracle Enterprise Manager Cloud Control

Using OEMCC you can discover compute nodes, cell nodes and InfiniBand switches. Monitor critical hardware metrics and view aggregated alerts from all components.

All subcomponent targets, hardware or software can be monitored from the central location – including CRS, ASM, database, hosts and IB network. Instead of adding each target individually, you can view the whole Exadata as a single target and drill down to individual targets whenever necessary. OEMCC also provides Automated Patching, EXAchk integration, Database consolidation workbench like Planning, Migration and Validation and Lifecycle management of Exadata in virtualized environment like Create/Delete RAC databases & Virtual Machines.

## Using ZFS and Zero Data Loss Recovery Appliance (ZDLRA) for backups

Exadata has everything for MAA. But it cannot protect you from Human Errors. If the company has not implemented the principle of least privileges and Junior DBA was able to drop important table with purge option, the only option is to restore table from backup. If the backup is located on slow disks or even on tapes, then meeting company's Service Level Agreement (SLA) is a tough process.

The capability of ZDLRA is continuous database-aware validation ensuring data integrity and database recovery to any point in time and Incremental forever backups eliminate backup windows and database server overhead.

Protecting the mission-critical data that resides on an Oracle Exadata is the top priority. The Oracle ZFS is an ideal solution for saving backups due to its performance, reliability and powerful features. The high bandwidth interconnects reduce backup and recovery time.

Another backup and recovery monster is ZDLRA. It is a completely different solution, which uses real-time redo transport to create full backups at any point in time.

Connection between Oracle Exadata and ZDLRA is either through 10GigE or InfiniBand.

Single Zero Data Loss Recovery Appliance (ZDLRA) scales to protect an entire Data Center.

## Exadata Maximum Availability Architecture

My Oracle Support Note 757552.1 is frequently updated to provide the latest information about MAA. Additional best practice papers are published at [www.oracle.com/goto/maa](http://www.oracle.com/goto/maa)

Protection from planned and unplanned outages are HA challenges. Exadata has hardware level fault tolerance:

- Database nodes, Cell nodes, DB Server disk drives, Cell disk drives, Flash drives are all redundant. ASM file system interface tolerates several types of failures using HIGH or NORMAL redundancy
- Redundant InfiniBand network connectivity using dual-ported QDR Host Channel Adapters and redundant switches
- Linux channel bonding configurable during deployment
- Redundant Power Distribution Units (PDUs)
- Hot-swappable power supplier and fans
- Redundant I/O paths
- Fastest Real Application Cluster (RAC) Instance and Node Failure Recovery
- Fastest Backup – Recovery Manager (RMAN) Offload to Storage Servers
- Deep Automatic Storage Management (ASM) Mirroring Integration
- Fastest Oracle Data Guard Redo Apply
- Complete Failure Testing with Lowest Brownouts

## Virtualization in Oracle Exadata

Virtualization provides CPU, memory and storage isolation for consolidated workloads. Exadata VMs have almost raw hardware performance by delivering I/O through ultra-fast InfiniBand network. Implementing IB network isolation across the clusters using IB partitioning provides better security. The traffic of one RAC cluster is not accessible to another RAC cluster.

Exadata virtualization can only run certified Oracle Linux versions. Windows, RedHat, and other guest OS are not supported. Trusted Partitions allow licensing by virtual machine.

## Choice of Oracle Exadata Deployment Models in Cloud Platform

- **On-Premises:** Customer Data Center – Purchased – Customer Managed
- **Cloud at Customer:** Customer Data Center – Subscription Model – Oracle Managed
- **Public Cloud Service:** Oracle Cloud - Subscription Model – Oracle Managed

## Benefits from Oracle Cloud Deployment Models

- Flexible Subscription Model
- Database PaaS Services
- Software Defined Networking
- Cloud Security and Hardening
- Oracle Managed Exadata InfraStructure

## Summary

Oracle database is a powerful software but installing it on slow performance servers may cause its inefficiency and customers may mistakenly think that Oracle is a “slow database”. Oracle recognized the need for a powerful server that could efficiently run database and released Exadata. Almost all hardware and software in Exadata is “Database-Aware” and is tuned for running Oracle database efficiently.



### Mari Kupatadze

Oracle Certified Master,  
Senior Solutions Architect at  
FlashGrid Inc.



### Venkata Ravi Kumar Yenugula

Oracle ACE Director &  
Oracle Certified Master,  
Oracle Engineered Systems  
Architect at Infoblox Solutions Inc.





The six cities of the EMEA Tour 2019 (credit: TajOUG)

# Heli Helskyaho The Groundbreakers EMEA Tour 2019

*During October we conducted the Groundbreakers EMEA Tour event. I was lucky to attend five of its total six legs.*





TajOUG president Rustam Khodjaev cutting the anniversary cake (credit: TajOUG)



Group photo from the TajOUG (credit: TajOUG)

The journey started in Dushanbe, Tajikistan on October 10. A full room of eager learners was waiting for the event to start at 8 o'clock in the morning. The speakers on this leg of the tour were I, Kamran Aghayev, and Rustam Khodjaev. I was talking about database designing and machine learning, Kamran was explaining how to move a data warehouse to Oracle Big Data Cloud Service and how to change the career from DBA to Data Engineer. Rustam was explaining Oracle

Academy Program being used at Educational Institutions of Tajikistan. At the end we had an "ask an expert" panel where the audience was able to pose questions about whatever they wanted to know. I ended up explaining a lot about reinforcement learning that happens to be one of my favorite research areas at the moment. Kamran was giving both technical and career related advice. The day ended with a big cake and a celebration of TajOUG's fifth anniversary.



Very early the next morning we flew to Istanbul for the next leg of the tour. In Istanbul, Sandesh Rao joined the speakers' team. On Saturday morning we arrived at the local university where the event took place. I talked about machine learning, Sandesh continued about the topic from the point I finished and shared some tips and tricks for troubleshooting Database 19c. Kamran gave the same interesting talks he gave in Dushanbe. I was amazed by the level of knowledge people had and the eagerness they had for learning. I was also very happy to meet an old friend of mine, Gurcan Orhan, whom I had not seen for a long time. After a very busy day we took a taxi to the airport to be in Baku, Azerbaijan for the next day's event.

On Sunday morning we all rushed to the beautiful venue in the center of Baku. Now also Ludovico Caldara and Rodrigo Mufalani joined our speakers' team for this part of the tour. Sandesh and I were talking about machine learning and I also explained how to get started with the Always Free Cloud Service and an autonomous database. Ludovico was explaining Oracle Drivers configuration for High Availability and made the audience laugh explaining how a pomegranate almost won him and how he went to a hammam. Rodrigo taught everybody how to be a patching hero while Rustam explained how the Oracle Academy Program was used at Educational Institutions of Tajikistan.

In Baku, I also visited a new Azerbaijan Finland school to meet their principal and the headteacher from Finland. This school really amazed me! They are using the Finnish methods of teaching but in an improved way. All the equipment in the school was the best one can buy, the school building was official but cozy at the same time. I wish them a great success with the school!



Our travelling group waiting for a taxi in Istanbul (credit: Sandesh Rao)

On Monday it was time to fly again. This time we flew to Ljubljana, Slovenia via Istanbul. Now, Oren Nakdimon joined the team in Istanbul. In Ljubljana we had a rental car to take us to Portoroz on the Slovenian coast where the event took place. Portoroz is such a beautiful place! We arrived there late in the evening and needed to rush to my favorite ice cream place (Cocoa) before closing time. Healthy dinner. 😊 The next morning it started. Oren explained how to upgrade an application with no downtime, Ludovico talked about



Me and Davor Ranković, president of the Croatian user group HrOUG  
(credit: Sandesh Rao)

CMAN, Kamran encouraged people to career change from DBA to Data Engineer, and I talked about machine learning and autonomous database. When we were done with all the presentations, we jumped into the rental car again to get to Rovinj, Croatia.

Tuesday was the tour day for Croatia. The event was on a beautiful island in Rovinj. It was so calming to listen to the

waves and to watch the sea. In Croatia, Oren showed how to upgrade an application with no downtime and how to do constraint optimization. Kamran explained how to move a data warehouse to Oracle Big Data Cloud Service and how to change the career from DBA to Data Engineer. Sandesh talked about machine learning and shared some tips and tricks for troubleshooting Database 19c. Ludovico gave a talk about REST for DBAs. I talked about in-database machine learning.

The next day the rest of the team continued to Bucharest, Romania. I stayed and gave my keynote presentation about the basics of machine learning.

In Bucharest they had an amazing venue full of excited audience. Which does not surprise me at all since Maria Colgan now joined the speakers' team. Maria talked about explain plan and showed what is under the hood of an autonomous database. Kamran explained how to move a data warehouse to Oracle Big Data Cloud Service and how to change the career from DBA to Data Engineer. Oren described how to upgrade an application with no downtime and how to do constraint optimization. Sandesh was talking about machine learning and shared his troubleshooting tips and tricks for Database 19c.

The tour was exhausting since we were always travelling or talking without any breaks, but I think it was a great tour and I am happy I was part of it. All the user groups were amazing, their leaders are fantastic and all the speakers on the tour were true superheroes! Thank you for letting me be a part of the tour! Thank you, Javed Mohammed for being there the whole time and making great videos and interviews. Thank you, Jennifer Nicholson to make this possible and thank you Kamran for organizing it!



# Ambassador's Corner

Heli Helskyaho

Dear user group leaders,  
I hope you and your user groups are doing well!

## EOUC Leaders' Meeting

This is the last Ambassadors' corner, since there are no more Ambassadors! The EOUC leaders' face-to-face meeting was held November 18 in Nuremberg, Germany. This meeting decided to establish a new legal entity for EOUC, EOUC e.V., and elected the new board for it. The new board members are: Mirela Ardelean (RoOUG, Romania), Björn Bröhl (DOAG, Germany), Ralf Kölling (DOAG, Germany), Wolfgang Scherrer (SAOUG, South Africa), Ann-Sofie Vikström Often (OUGN, Norway), and Andrejs Vorobjovs (LVOUG, Latvia). I want to thank EOUC for all the years I was your ambassador and also want to wish all the best to the new EOUC e.V. and its new board.

## Groundbreakers' Tour EMEA

The EMEA version of the Groundbreakers' Tour took place in October. The tour (<http://ogbemea.com>) started in Dushanbe, Tajikistan October 10, then moved to Istanbul, Turkey October 12, Baku, Azerbaijan October 13, Portoroz, Slovenia October 15, Rovinj, Croatia October 16 and finally to Bucharest, Romania October 18. Based on my own experience and all the feedback I heard the tour was a great success!

I want to take the opportunity to thank especially Kamran Aghayev A. and all the user groups involved for organizing it. Thank you to the Groundbreaker programs at Oracle for making this possible (especially Jennifer Nicholson and Javed Mohammed). Thank you for all the speakers that were part

## Your Ambassadors:

If you have anything we can help with,  
please do not hesitate to contact us!



**Ami Aharonovich**  
**iIOUG (Israel)**  
[ami@dbaces.com](mailto:ami@dbaces.com)



**Heli Helskyaho**  
**OUGF (Finland)**  
[heli@miracleoy.fi](mailto:heli@miracleoy.fi)

of this tour! It was a very hard trip to travel but you all did it with a smile on your face! And thank you very much for all the attendees! Without you this tour would be nothing. The plans for the Groundbreakers' Tour EMEA 2020 have already started.

All the best to you and your user group!  
Best regards,  
Heli



# Call for Papers

## ILOUG Techdays

Petah Tikva, Israel  
<https://www.iloug.org>

## APEX World 2020 'Autonomous APEX'

SS Rotterdam, Rotterdam, The Netherlands  
<https://www.nloug.nl/page.aspx?event=1476>

## Kscope 20

until December 5, 2019  
Boston, Massachusetts, USA  
<https://kscope20.odtug.com>

## RigaDevDays 2020


until December 9, 2019  
Riga, Latvija  
<https://papercall.io/rdd20>

## OUGN Spring Conference 2020

until December 15, 2019  
Color Fantasy (cruise ship), Oslo - Kiel - Oslo  
[c4p.ougn.no](https://c4p.ougn.no)

## DOAG 2020 Datenbank

until January 15, 2020  
Düsseldorf, Germany  
<https://datenbank.doag.org/en/speakers/speakers>





# Events

## Sangam19

December 6 - 7, 2019  
Hyderabad, India  
<http://sangam19.info>

## DBA dag

December 13, 2019  
Oracle Headquarters Utrecht, Utrecht, The Netherlands  
<https://www.nloug.nl/page.aspx?event=1480>

## ItOUG Tech Days 2020

January 29 + 31, 2020  
Milan + Rome, Italy  
[www.itoug.it](http://www.itoug.it)

## ILOUG Techdays

February 3 - 4, 2020  
Petah Tikva, Israel  
<https://www.iloug.org/program>

## OUG Ireland 2020

March 12 - 13, 2020  
The Gresham, Dublin, Ireland  
<https://ukoug.org/page/ougireland2020>

## APEX World 2020 'Autonomous APEX'

March 16 - 17, 2020  
SS Rotterdam, Rotterdam, The Netherlands  
<https://www.nloug.nl/page.aspx?event=1476>

## JavaLand 2020

March 17 - 19, 2020  
Phantasialand, Brühl, Germany  
<https://www.javaland.eu/en/home>

## OUGN Spring Conference 2020

March 19 - 21, 2020  
Color Fantasy (cruise ship), Oslo - Kiel - Oslo  
<https://ougn2020.com>

## POUG Workshop

March 27, 2020  
Warsaw, Poland  
<https://poug.org>

## COLLABORATE 20

April 19 - 23, 2020  
Mandalay Bay Resort and Casino, Las Vegas, Nevada, USA  
<https://questoraclecommunity.org/events>





# Events

## APEX connect 2020

May 5 - 7, 2020  
Phantasialand, Brühl, Germany  
<https://apex.doag.org/de/home/>

## OUG Scotland

May 20, 2020  
The Studio, Glasgow, UK  
<https://ukoug.org/page/ougscotland2020>

## DOAG 2020 Datenbank

May 25 - 26, 2020  
Düsseldorf, Germany  
<https://datenbank.doag.org/de/home>

## RigaDevDays 2020

May 26 - 28, 2020  
Riga, Latvia  
<https://rigadevdays.lv>

## UKOUG Technology Summit

June 3, 2020  
Birmingham, UK  
<https://ukoug.org/page/techsummituno>

## UKOUG Business Applications Exchange

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

## Kscope 20

June 28 - July 2, 2020  
Boston, Massachusetts, USA  
<https://kscope20.odtug.com/>

## RECONNECT 20

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

## INFOCUS 20

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





## 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: Stefan Kinnen.

### 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 #20 edition: **February 10, 2020**.

### Photo credits:

Title: © Dejan Bozic | <https://123rf.com>  
P. 4: © floralset | <https://123rf.com>  
P. 5: © Comic: [www.commitstrip.com](http://www.commitstrip.com)  
P. 6: © aimage | <https://123rf.com>  
P. 9: © Maxim Evseev | <https://123rf.com>  
P. 15: © Anastasiia Nevestenko | <https://123rf.com>  
P. 22: © Natalia Merzlyakova | <https://123rf.com>  
P. 32: © Freepik | <https://freepik.com>  
P. 32: Logo © Ubuntu | <https://ubuntu.com>  
P. 34: © fullvector | <https://freepik.com>  
P. 46 - 48: © Schuchrat Kurbanov | DOAG