Player Piano - The World of ADW & ATP

Conf42: DevOps 2024

Abstract

Emerging information technology trends for the cloud have the power to transform organizations. In the data management and analytics space, a key cloud service offering has arrived: Oracle Autonomous Data Warehouse (ADW) Cloud and Oracle Autonomous Transaction Processing (ATP)

Rich Niemiec

Chief Innovation Officer @ Viscosity North America

Tweet

Share

January 25 2024 - premiere 5PM GMT

Thu Jan 25 2024 11:00:00 GMT-0600 (Central Standard Time) in America/Chicago

Player Piano: The World of Autonomous Database (ATP/ADW)

Fast - Easy – Elastic

Rich Niemiec

@richniemiec

ORACLE ACE Director

richniemiec@gmail.com

Conf42: DevOps 2024

January 25 2024 - premiere 5PM GMT

Special Thanks:

William Hardie, George Lumpkin, Maria Colgan, Charles Kim, Justin Nugent

VISCOSITY NORTH AMERICA

Agenda – Your Self-Driving Future

- The DBA, Autonomous & the Cloud
- A Robot May Not Look Like One
- Autonomous Transaction Processing (ATP)
- Autonomous Data Warehouse (ADW)
- Machine Learning & Data Visualization Desktop
- Next: **Robots & the Future Ahead**

Quick FREE notes

Send email to (for slides):

hello@viscosityna.com

(richniemiec@gmail.com)

@richniemiec - twitter

Conf42: DevOps 2024

January 25 2024 - premiere 5PM GMT

Viscosity Pillars and Delivery Models

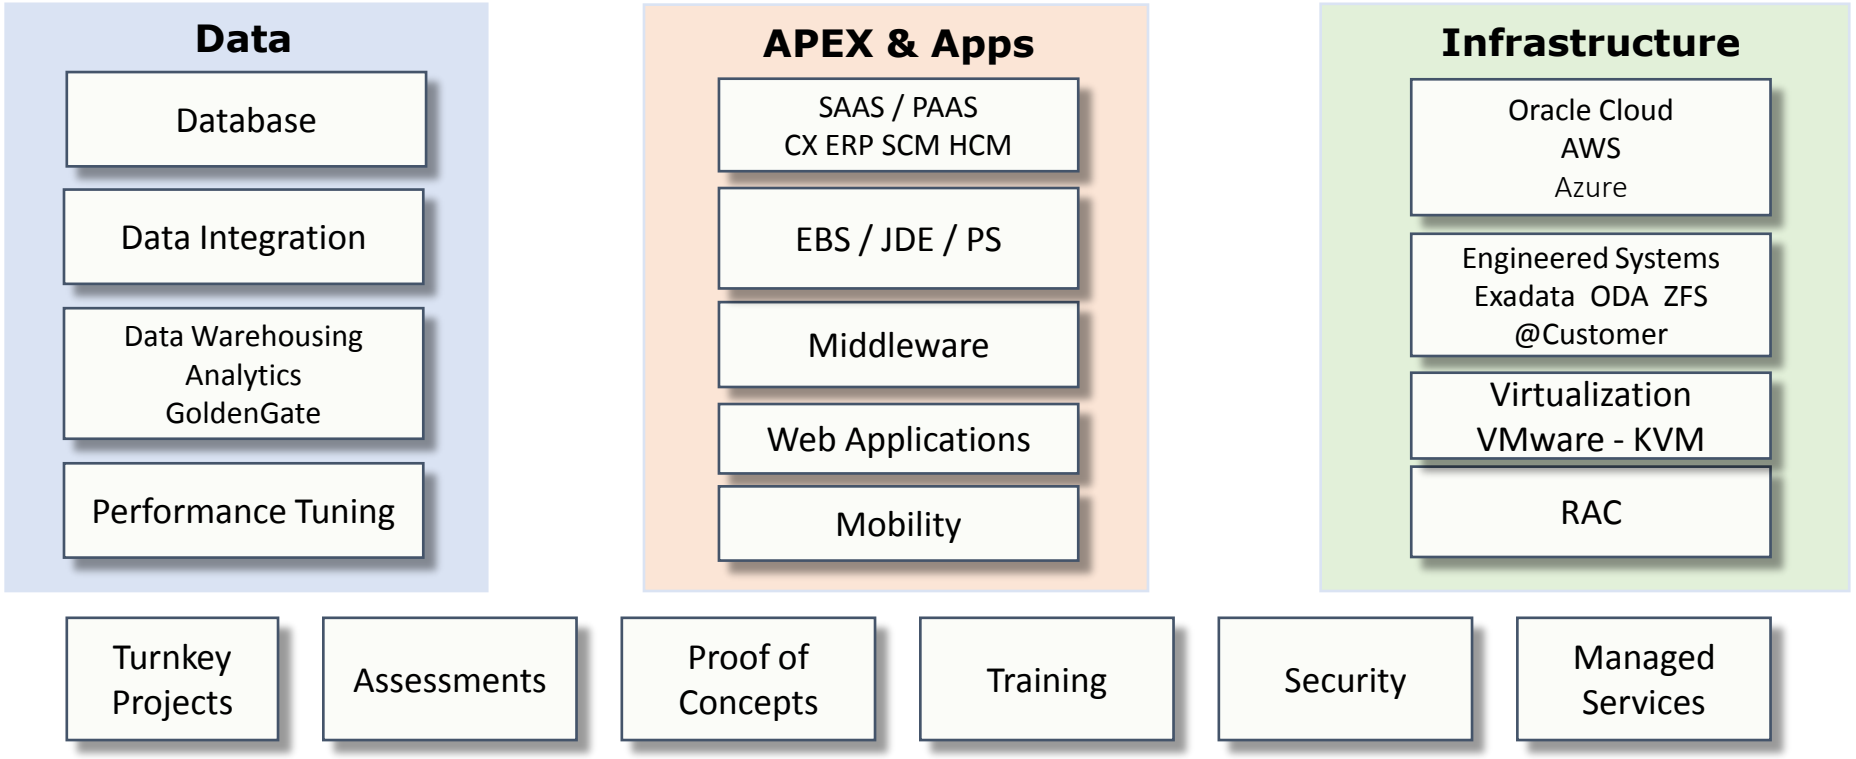

ORACLE Platinum Partner

Viscosity's Oracle ACES

Developer Resource Center /
Oracle ACE Program

Oracle ACE Program

The Oracle ACE Program recognizes and rewards community members for their technical contributions in the Oracle community.

Charles Kim,
CEO & Co-Founder

Twitter: @racdba

ORACLE ACE Director

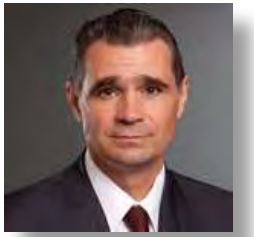

Rich Niemiec,
Chief Innovation Officer

Twitter: @richniemiec

ORACLE ACE Director

Craig Shallahamer
Applied AI Scientist

Twitter: @orapub

ORACLE ACE Director

Sean Scott,
Consultant

Twitter: @oraclesean

ORACLE ACE

Gary Gordhamer,
Consultant

Twitter: @ggordham

ORACLE ACE

Julio Ayapan,
Consultant

ORACLE ACE Associate

We wrote the books - many Experts!

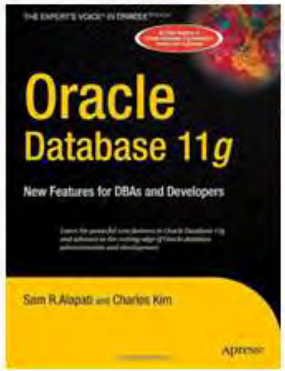

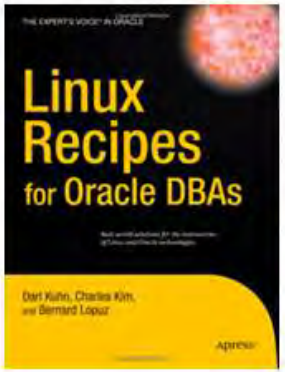

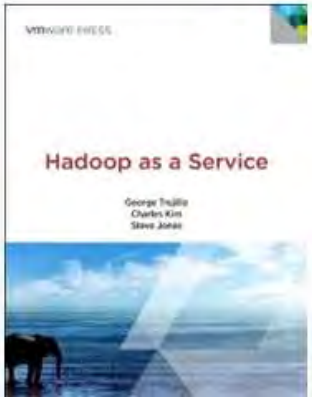

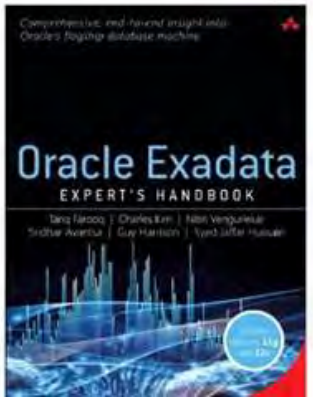

"We Enable Business Transformation at a Time when Companies must Change to Survive"

Go to 19c or 21c (*get to 19c!*)?

Why Upgrade to 19c?

19c

Core Aim : Long Term Stability

Long Term Support Release:
- 5 Years of Premier Support to 2024
- 3 Years Extended Support to 2027

21c

Innovation Release:
- 2 Years of Premier Support to 2023
- No Extended Support

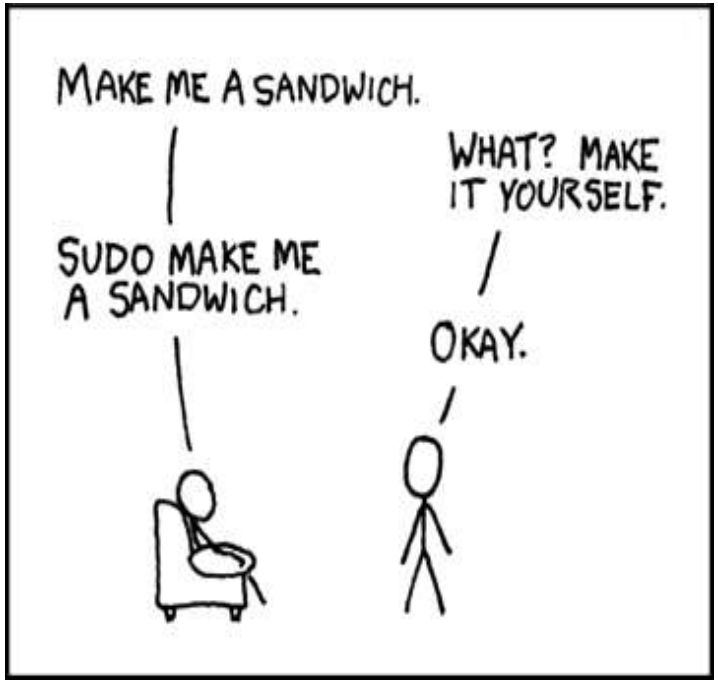

- Need Auto Machine Learning of 21c?
- Want to Export ADW/ATP?
- Leverage Oracle Data Safe. Security Assessment of: Database / Users / Data

Some Jobs gone: Computer Surpassed Them!

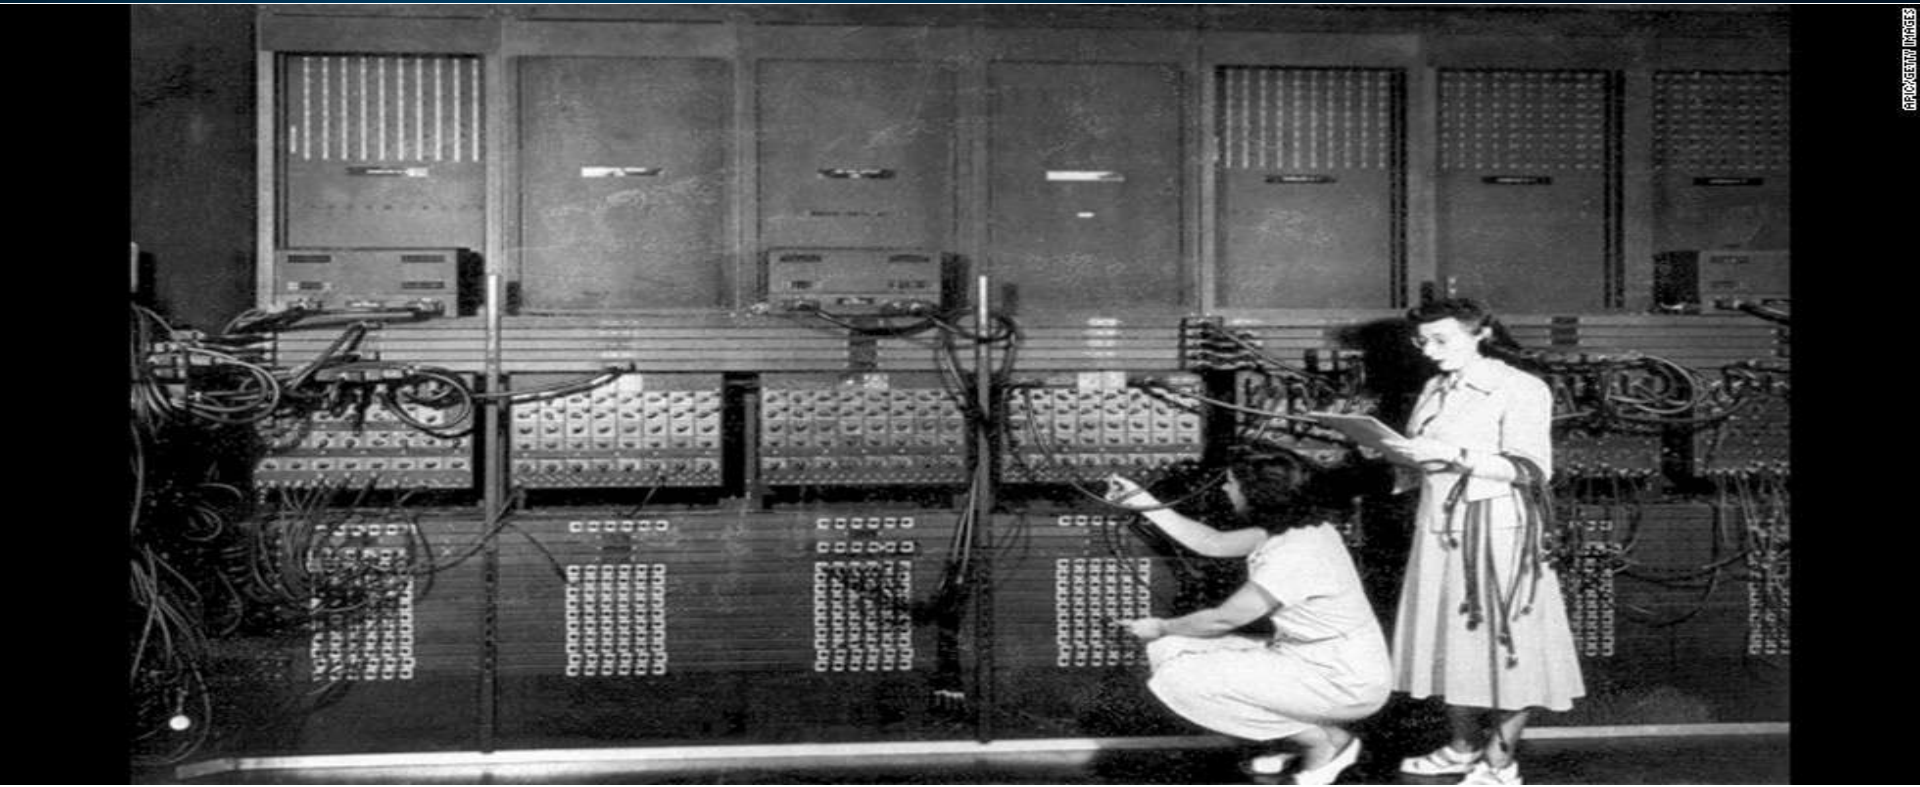

"The competition between Man's Mind vs. the Product of Man's Mind… for this, there is standing room only in the Twilight Zone." - Rod Serling

Autonomous Database – Replacing the DBA?

At OpenWorld

At Modern Cx (Customer Experience)

All 1,000 units of "Pepper" being offered in September sold out within one minute.

AI GUEST

It's time for workers to worry about AI

GARY GROSSMAN, EDELMAN @GARYG02 APRIL 7, 2019 2:22 PM

MOST READ

It's time for workers to worry about AI

Remote AR will make it so we can work —

72% of IT Budget is Spent on Maintenance

DBA & IT Stress by the Numbers

39%

Workloads for DBAs are increasing: **39%** of DBAs handle **50** or more databases

95%

Automation is lacking: **95%** of DBAs create or upgrade databases manually

78%

78% of DBAs will experience unplanned downtime from untested database changes during their careers

2 out of 3

DBAs and IT staff are struggling to provide full protection: **2 out of 3** organizations use multiple tools to backup a single database

Source: Oracle customer survey

ORACLE

The Answer… it's the Self-Driving Database

From Maria Colgan's talk on it…

Oracle's **Vision** for Autonomous Database

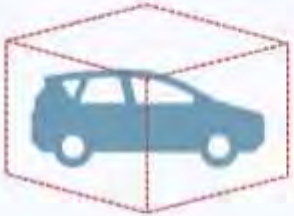

Self-Driving

User defines service levels, database makes them happen

Self-Securing

Protection from both external attacks and malicious internal users

Self-Repairing

Automated protection from all downtime

Concerns with Autonomous (worst DBA article)

Oracle's autonomous database could leave DBAs unemployed

The long-term future of database administrators could be at risk if every enterprise adopts the Oracle 18c autonomous database

Stuart Kennedy
09 May 2018 1:08

Oracle CEO Mark Hurd threw the long-term future of hundreds of thousands of database administrators (DBAs) into question while talking up the benefits of the company's new autonomous database.

The autonomous self-patching, self-healing database, the first version of which is 18c, is a part of a long-term play to help draw the company's customers into Oracle's piece of the cloud – which is increasingly packing itself with cloud-based applications and services.

Latest News

Leading PC makers set to expa
to target existing customers

Artificial intelligence making m
Russian banking

The Autonomous Database Cloud

Vendor’s view of the Cloud

Vendor’s path for you to the Cloud

How Easy to Move to Cloud Quickly

A Few Years Later on the Cloud

Data: Becoming the World's Most Valuable Resource

Data driven organizations are **23 times more** likely to acquire customers, **6x as likely** to retain those customers, and **19x as likely** to be profitable as a result.[4]

$430B Advantage to Data Driven Organizations[2]

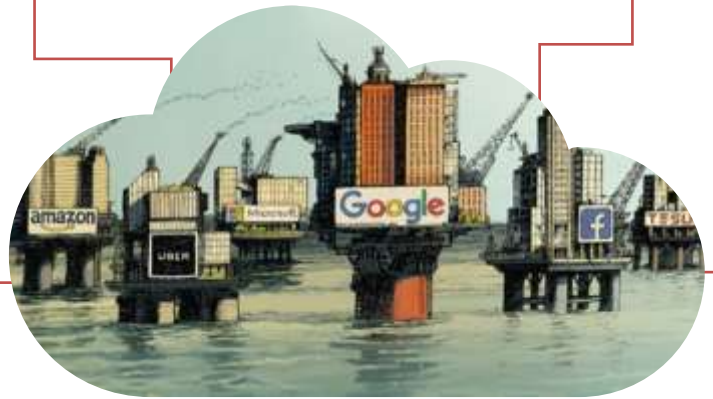

Only about **29%** of Organizations are Using Deep Business Analytics[3]

Tesla is (Data from 25M miles/day) Worth More Than GM Which Shipped **92x** More Vehicles Q1/17[1]

Source:
1. Economist Magazine, May 2017
2. IDC, 2020 Prediction of Value created by Capitalizing on Data
3. Accenture and General Electric Report On Predictive Analytics
4. McKinsey Global Institute

More Data Variety—Better Predictive Models

- Increasing sources of relevant data can boost model accuracy

Model with *"Big Data"* and hundreds -- thousands of input variables including:

- *Demographic data*
- *Purchase POS transactional data*
- *"Unstructured data", text & comments*
- *Spatial location data*
- *Long term vs. recent historical behavior*
- *Web visits*
- *Sensor data*
- *etc.*

Engineered Features – Derived attributes/variable that reflect domain knowledge—key to best models

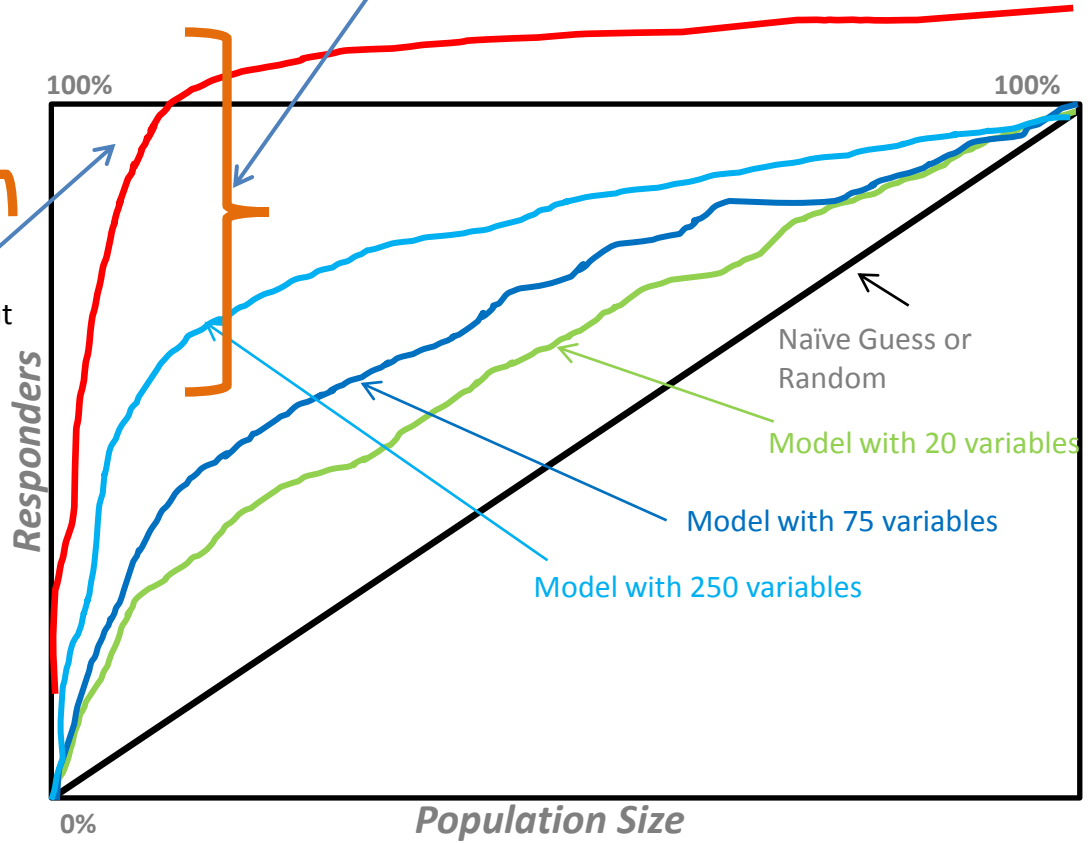

Characteristics of Big Data

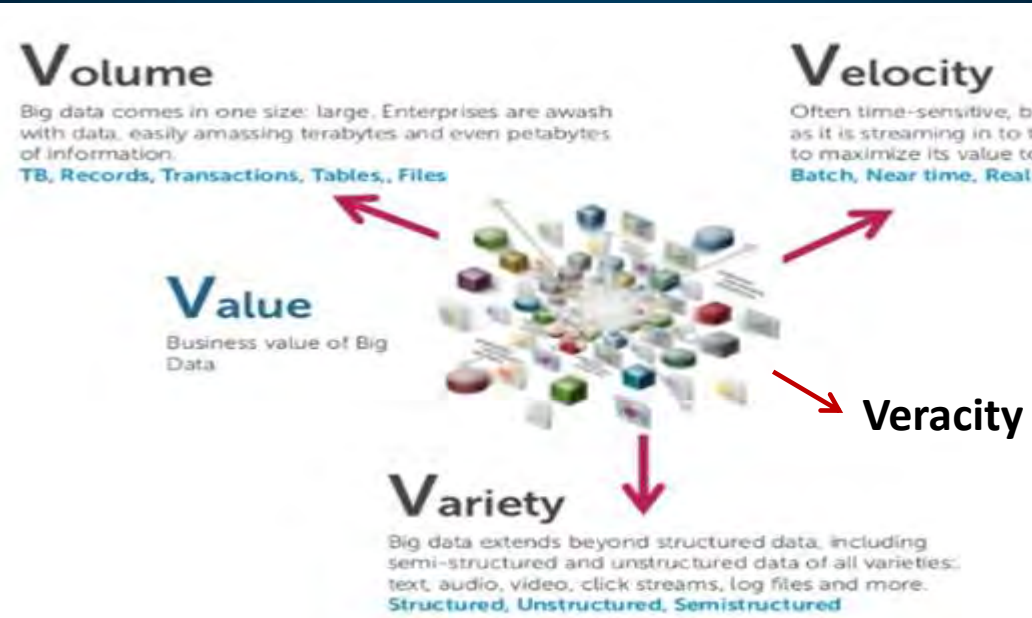

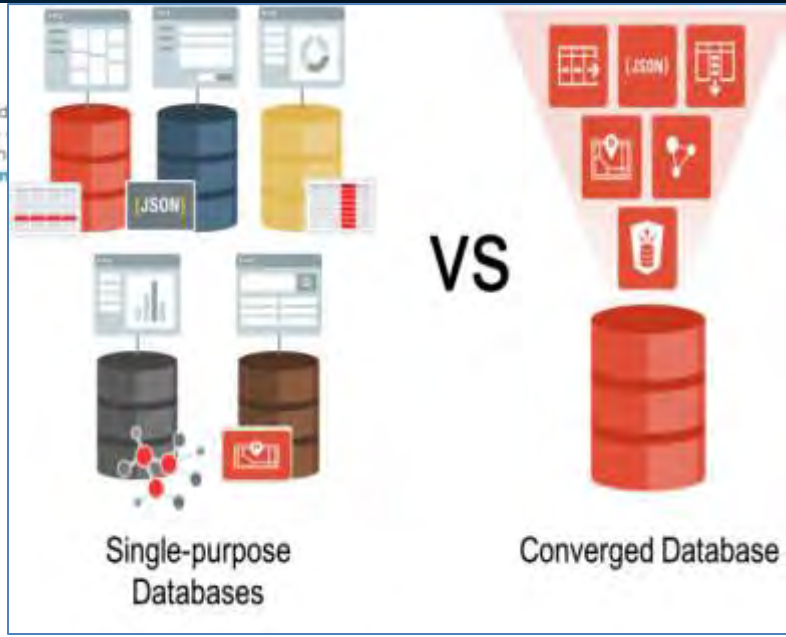

Big Data Themes

- HW & SW technologies for large data volumes
- Focus on Web 2.0 technologies
- Database Scale-out
- Relational & Distributed Data Analytics
- Real Time Analytics

Big Data Domains

- Digital Marketing Optimization
- Data Exploration & Discovery
- Fraud Detection & Prevention
- Social Network & Relationship Analysis
- Machine-generated Data Analytics

Converged Database - Oracle Multi-Model Database*

*Slide from Oracle's, Nitin Vengurlekar

- Benefits of Oracle's Converged Database are broad

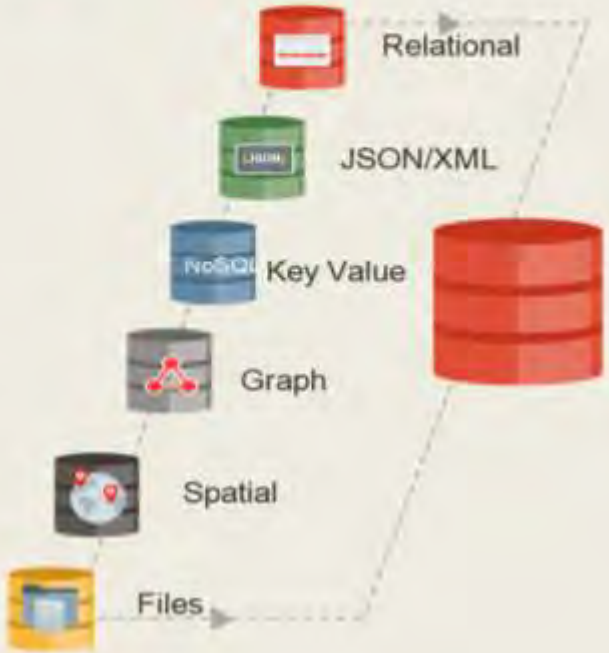

- Integrated development tools (Apex, SQL Dev, Spatial Studio) ✔
- 3rd party and Open Source development tools ✔
- Machine Learning ✔
- Node.js, Python, many others ✔
- In-memory database ✔
- Spatial, Graph support ✔
- NoSQL (JSON, key-value, wide column, XML) ✔
- Containers, microservices, virtualization (Docker, MT) ✔
- Integrated Security ✔
- Deployment choice (on-prem, cloud, hybrid) ✔
- Integrated High Availability and Disaster Recovery

Venturebeat.com – IOT $14.4T this decade

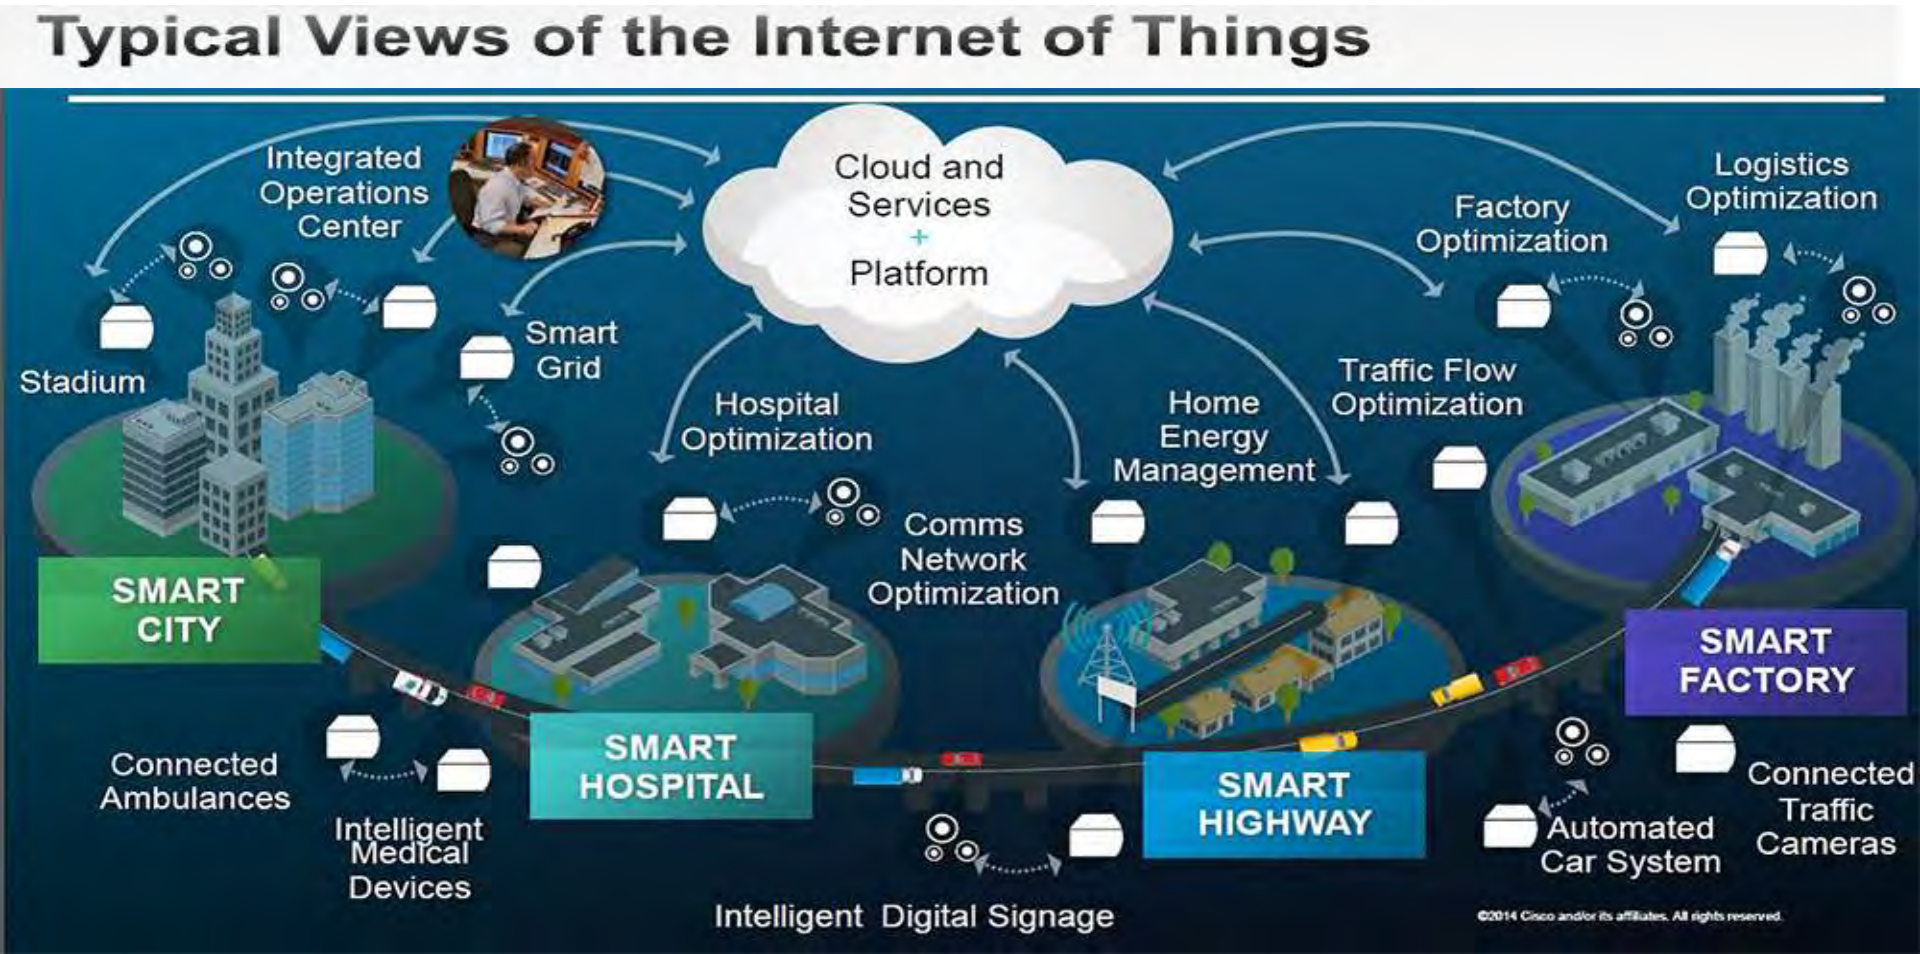

The Consumer IOT Connection

Vessyl smart cup tracks coffee and more for IoT hydration

Big Data / IOT Driving Prescriptive Analytics

Automating the DB gives more time to Impact Business!

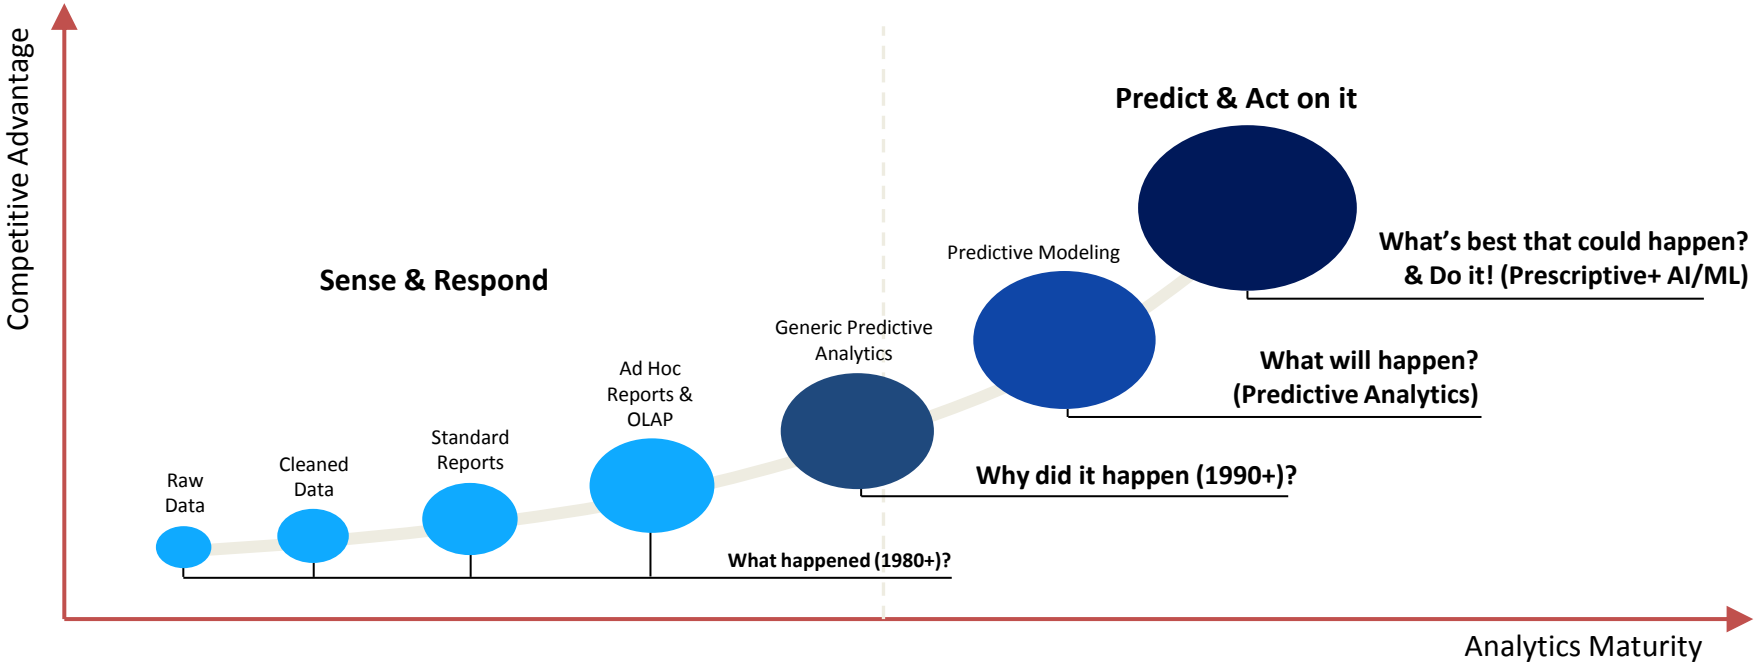

The key is unlocking data to move decision making from sense & respond to predict & act

Biju Thomas at ODTUG - *Emerging* Jobs

#8 Data Engineer

LinkedIn

2020 Emerging Jobs Report

- 33% annual growth
- Skills
  - Apache Spark
  - Hadoop
  - Python
  - ETL
  - AWS
- Industries
  - Computer Software
  - Information Technology
  - Financial Services
  - Healthcare & Hospitals

- Primary job responsibilities involve preparing data for analytical or operational uses.
- Works as part of an analytics team, providing data in a ready-to-use form to data scientists.
- Commonly deal with both structured and unstructured data sets

The DBA

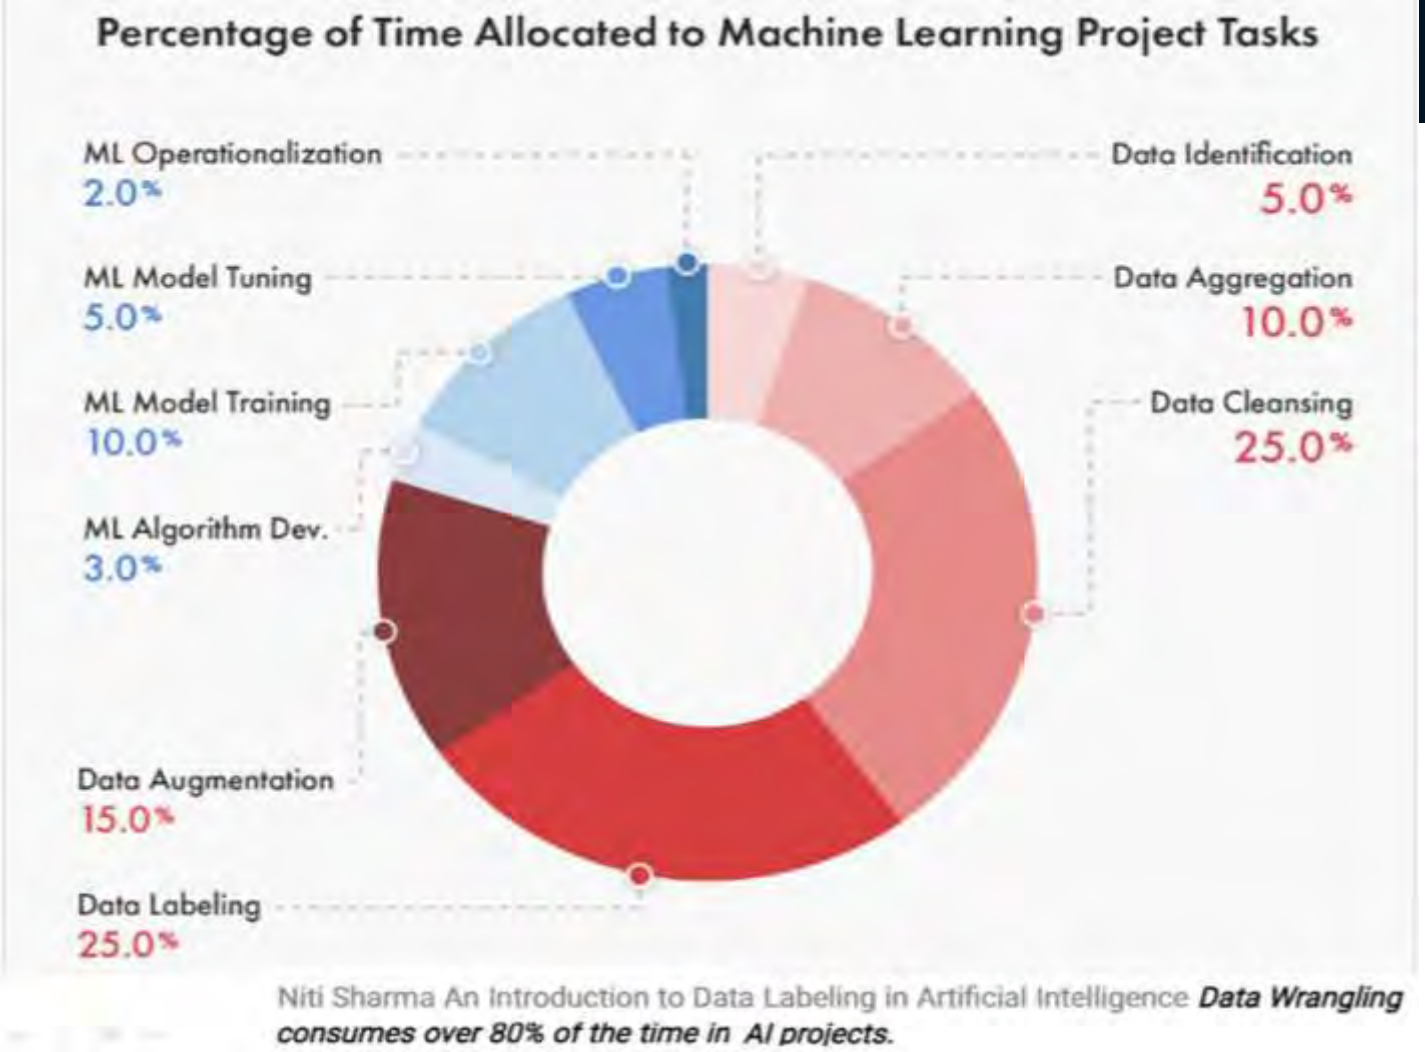

The DBA is <u>most important</u> part of the Machine Learning Process

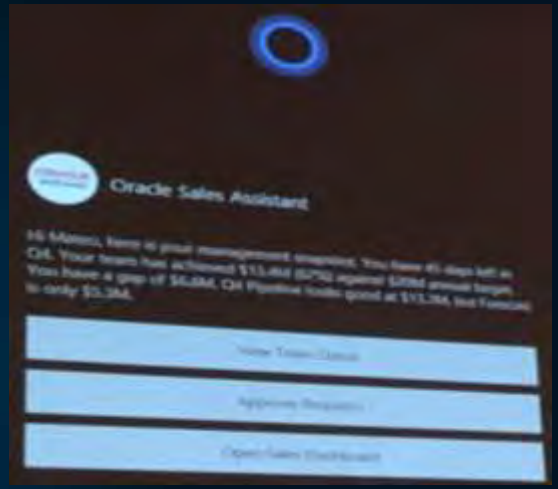

A Robot may not look one!

Oh Yeah… they never complain, always happy to do more, work anywhere, get smarter as time goes on, leverages AI & ML, works 24x7, doesn't ask for a raise, no union (yet).

A Robot may not look one!

*Robots that Manage a Database (ADW/ATP)!

*Robots that secure a system and use ML & AI

The Autonomous Database Cloud

- Self-Healing
- Self-Driving
- Self-Tuning
- Self-Recovering
- Self-Scaling Administration

Reality of the Autonomous Database

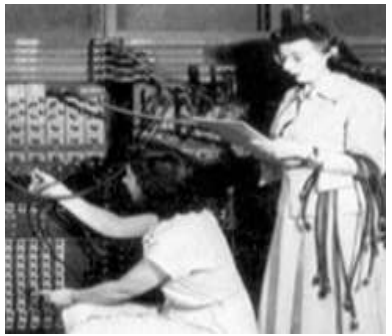

Will my job change?

- Absolutely…sure hope so!
- Hopefully… It already has!
- It has many times in the past…
- It will move **closer to the business** & innovation
- Data Manager instead of DBA & Developer
- Security Expert instead of Security on the DB
- Watching over costs more
  - Cloud Hidden Costs
  - Cloud, Hybrid, or On-Site Decisions
  - Decide which databases should be Autonomous

Oracle Autonomous spreads to OLTP: 8/7/2018

Rich Niemiec @RichNiemiec · Aug 7

Are DBAs up for early retirement? @larryellison says No, but DBA won't do tedious things - managing disk storage & upgrading O/S. DBAs will work with dev on new apps & work with CIO to help business vs. patching. @oracledatabase

Autonomous Optimizations - Specialized by Workload

	ADW	ATP (Creates Indexes) *In 19c
Primary Goal	Fast Complex Analytics	Fast Transaction Processing
Data Formats	Columnar	Row
Data Access Acceleration	Creates Data Summaries	RDMA for messaging and IO
Memory Usage	Parallel Joins and Aggregations	Data Caching to Avoid IO
Statistics	Automatically manages optimizer statistics as data changes	

Automatically Withstands Errors

- Autonomous database must be safer than manually operated database

- **Exceptional SLA Guarantee:**
  - **99.995 NRX%** (No Ridiculous eXclusions)
  - Guaranteed triggered if there is more than 2.5 minutes downtime a month
  - Includes patching and upgrade, database bugs, regional outages
- Applications that are architected for HA using Oracle best practices will usually see pauses, not failures

Outage	Key Feature	Potential Downtime
Server Outage (HA)	**RAC**	Near-Zero
Regional Outage, Disaster Recovery	**ADG**	Seconds
Data Corruption	**ADG**	Zero
Patches (Updates)	**RAC**	Near-Zero
Database Upgrade	**ADG**	Seconds
Table/Index Changes	**Redef**	Zero
User Error	**Flashback**	Time Since Error

The Autonomous Database & the IT Dept.

Autonomous Databases into the future:

- Who ensures database is tuned before it gets to the Cloud?
- Who ensures the cloud vendor is charging correctly?
- Who ensures the backup, security, or recovery is correct?
- Who decides what kind of service the databases will be?
- Who will build the policies for those autonomous databases?
- Who will have the knowledge to decide or estimate the cost of these services?
- Who decides the complex IT Infrastructure when we have more options & vendors?

- The answer is obvious: **A DBA & the rest of IT**, but not a simple DBA; A DBA that has evolved with all this new generation of databases on Cloud (**A Data Manager**).

Get Started FREE: oracle.com/cloud/free

cloud.oracle.com/tryit

oracle.com/cloud/free (Above link takes you here)

- Create your first Autonomous Database Data Warehouse
- Create your first Autonomous Database Transaction Processing
- Try out Analytics & Machine Learning
- Monitor / Clone …etc.

Get Started FREE: oracle.com/cloud/free

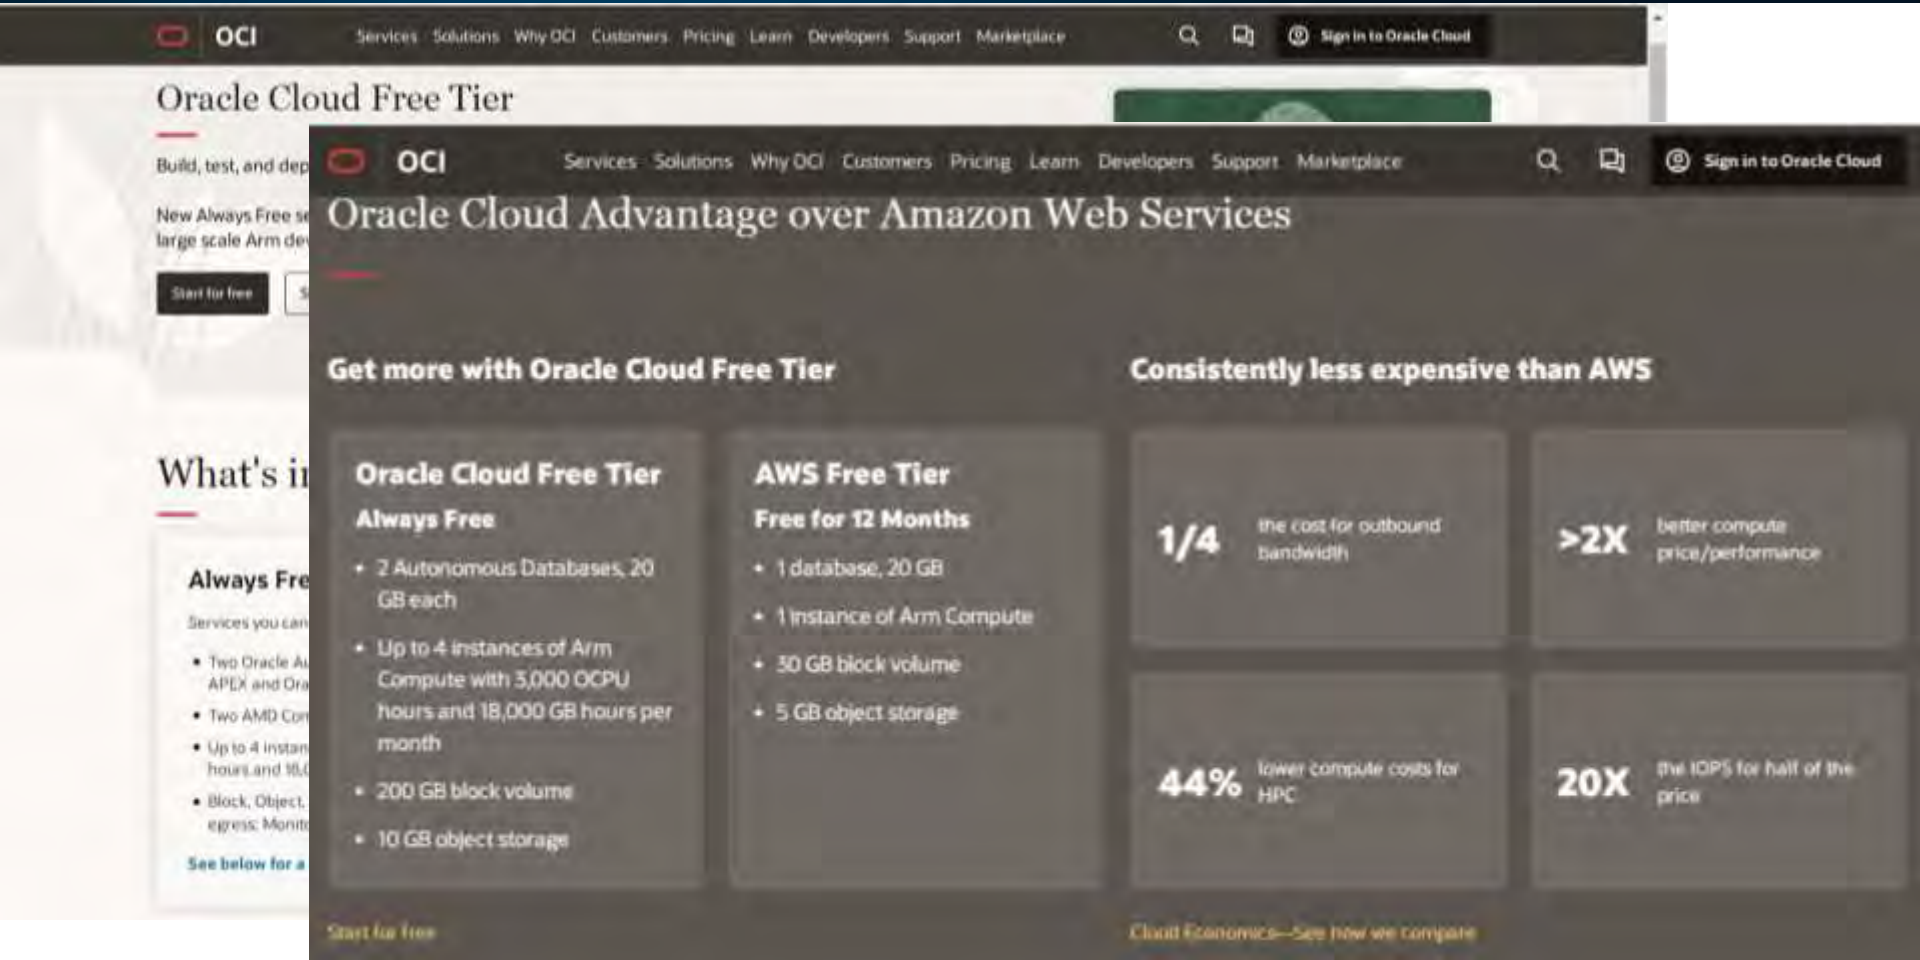

Get Started FREE: oracle.com/cloud/free

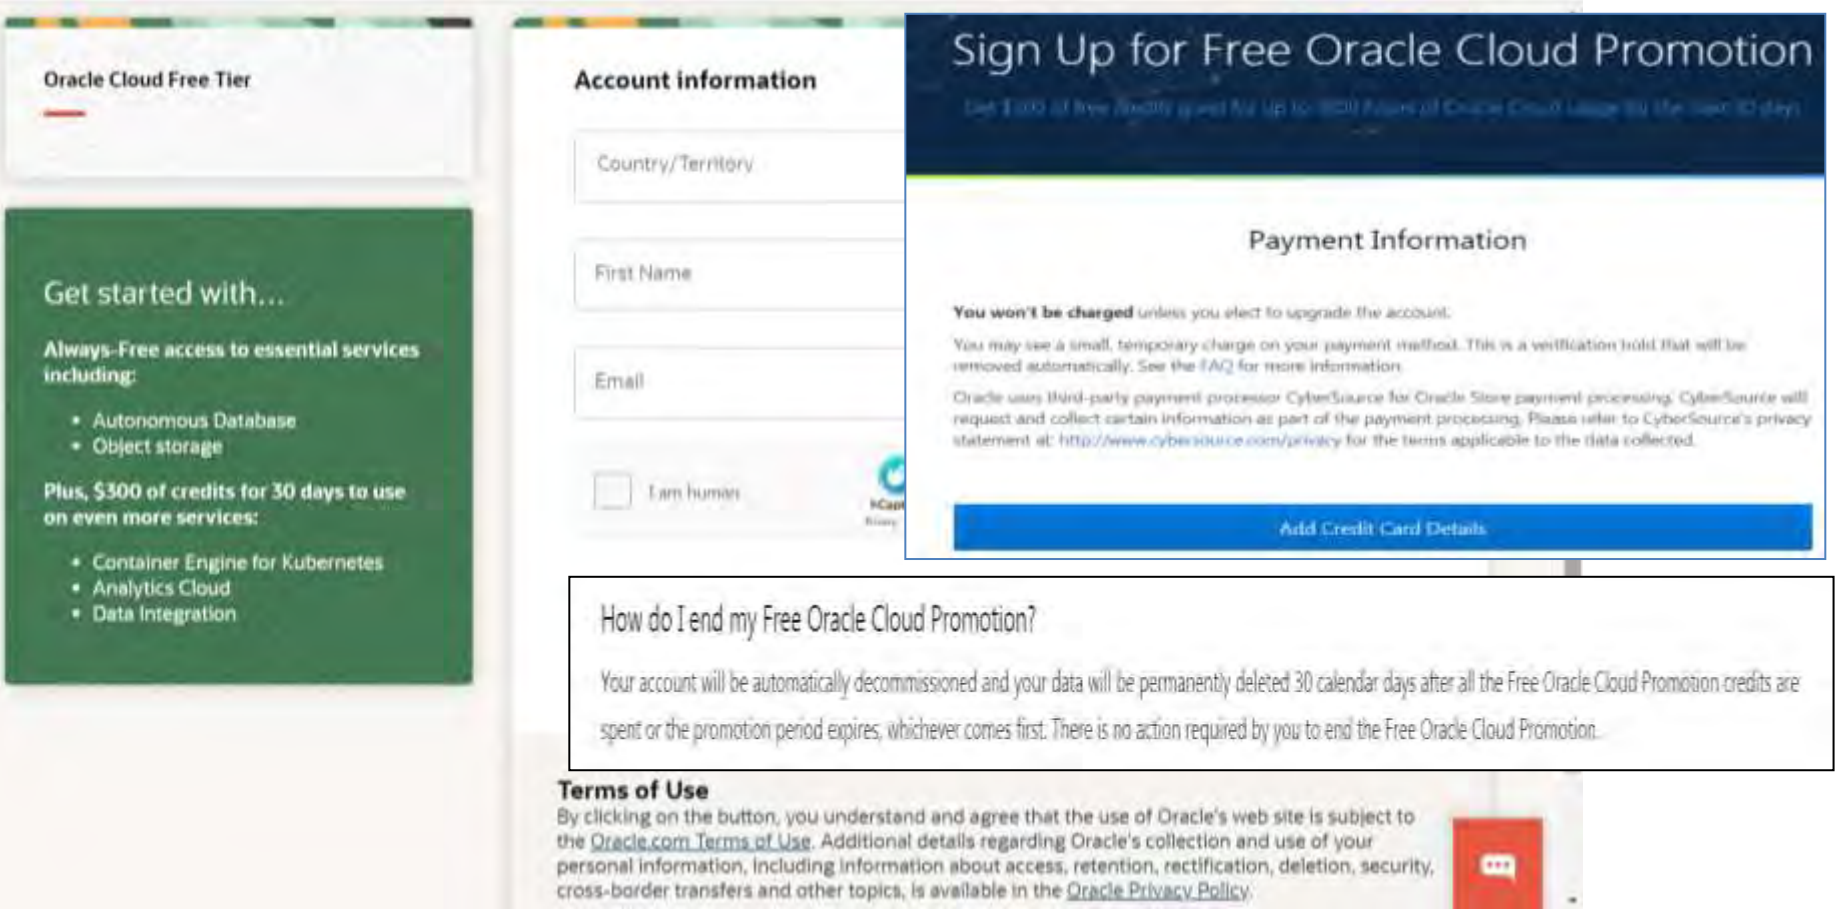

oracle.com/cloud/free

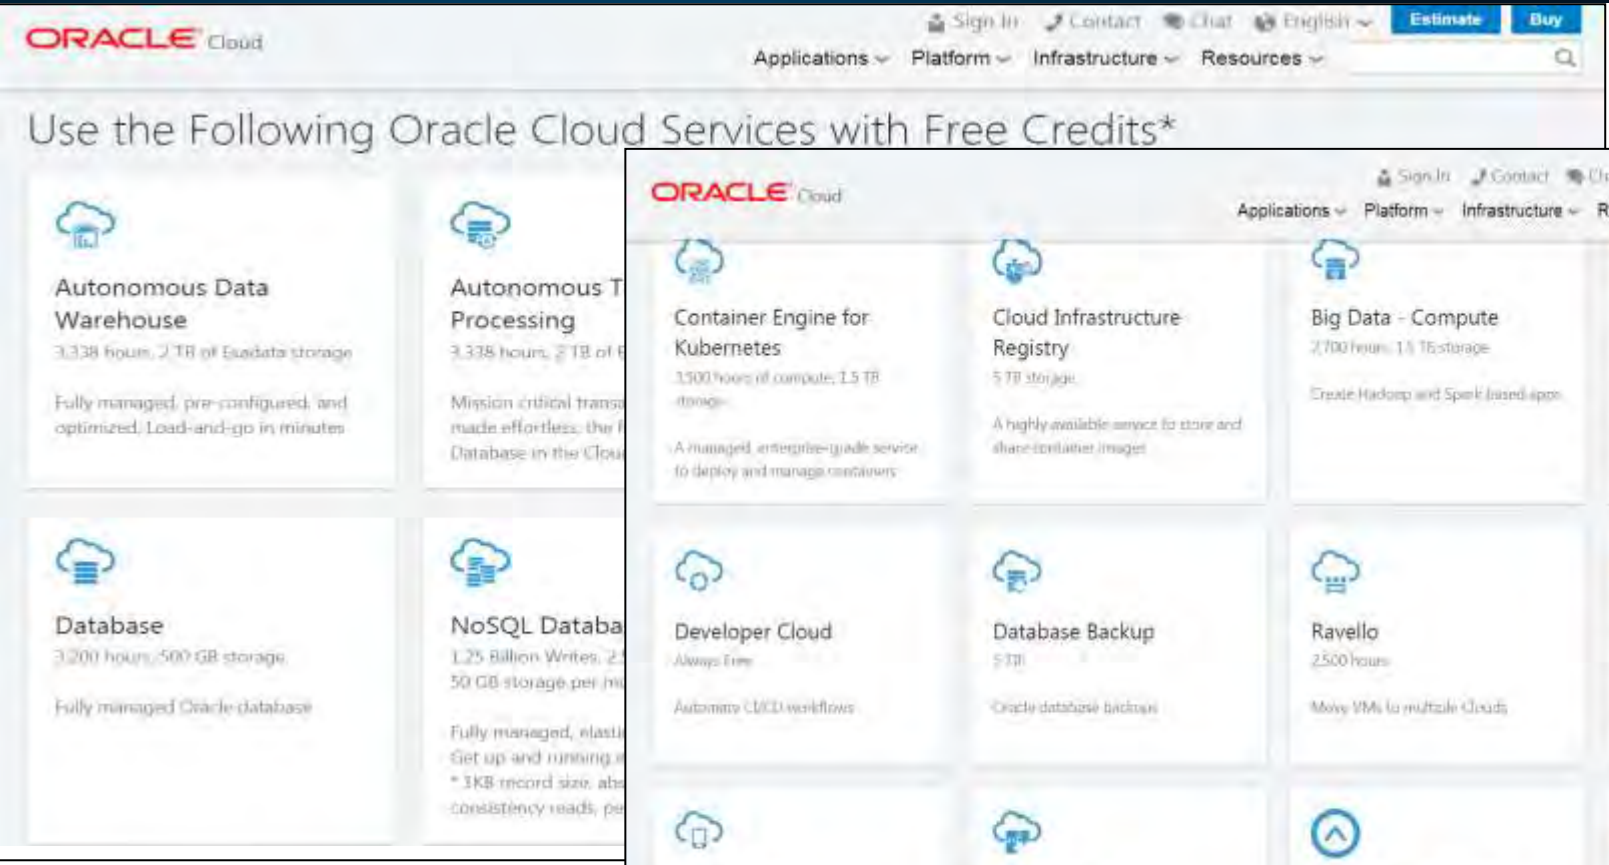

ADW/ATP – Provisioning a Database

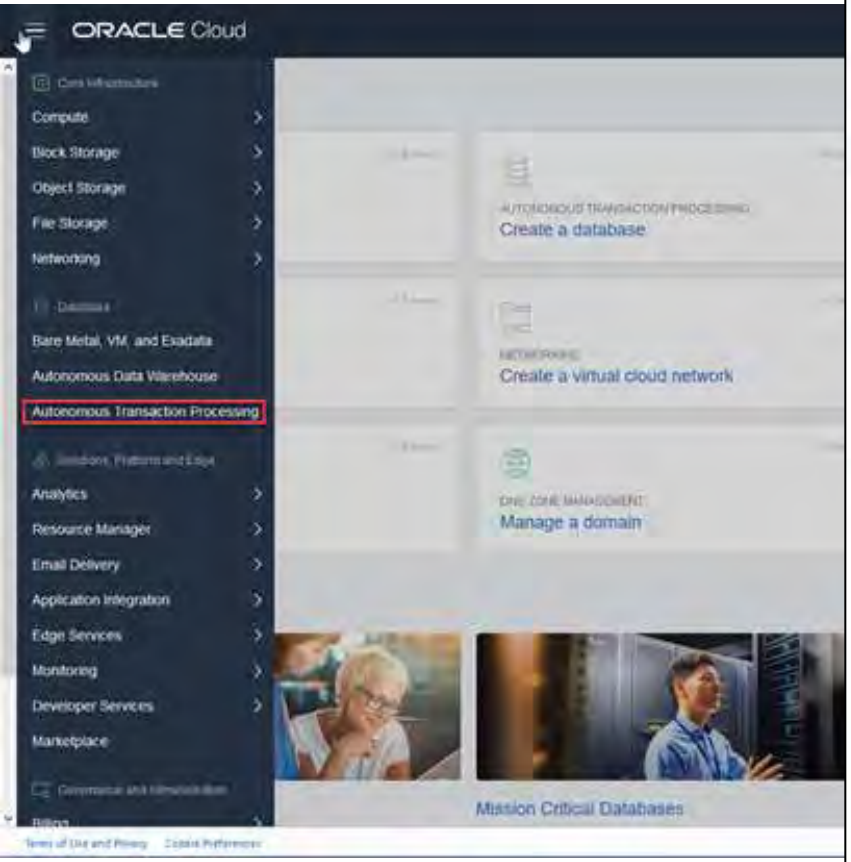

7. The Create Autonomous Database dialog appears. Enter the following information:
   - **Compartment** - Select a compartment for the database from the drop-down list.
   - **Display Name** - Enter a name for the database for display purposes.
   - **Database Name** - Use letters and numbers only, starting with a letter. Maximum length is 14 characters. (Underscores not initially supported.)
   - **Workload Type** - Autonomous Transaction Processing is automatically selected. Alternately, you can choose Data Warehouse as the workload type.
   - **Deployment Type** - Serverless will automatically be selected. This choice runs the autonomous database without provisioning a dedicated infrastructure. Alternately, you can choose the Dedicated Infrastructure deployment type to run the autonomous database on a dedicated Exadata infrastructure.
   - **CPU Core Count** - Number of CPUs for your service.
   - **Storage (TB)** - Select your storage capacity in terabytes. It is the actual space available to your service instance, including system-related space allocations.
   - **Auto Scaling** - If you select the auto scaling option, Autonomous Transaction Processing can use up to three times more CPU and IO resources than specified by the number of OCPUs currently shown in the Scale Up/Down dialog. When auto scaling is enabled, if your workload requires additional CPU and IO resources, the database automatically uses the resources without any manual intervention required.
   - **Administrator Credentials** - Password for ADMIN user of the service instance. The password must meet the following requirements:
     - The password must be between 12 and 30 characters long and must include at least one uppercase letter, one lowercase letter, and one numeric character.
     - The password cannot contain the username.
     - The password cannot contain the double quote (") character.
     - The password must be different from the last 4 passwords used.
     - The password must not be the same password that is set less than 24 hours ago.
   - **License Type** - Select whether you have existing licenses or if you want to subscribe to new database software licenses and the database cloud service.
   - **Tags** - (Optional) Tagging is a metadata system that allows you to organize and track resources within your tenancy. Tags are composed of keys and values which can be attached to resources.

   Click **Create Autonomous Database**

19c - Machine Learning with Automatic Indexes*

Rich Niemiec @RichNiemiec · 7 Aug 2018

Oracle19c #autonomousdatabase will leverage #machinelearning when you upgrade from #Oracle18c by using regression testing to ensure things will be faster in the new version. @oracledatabase @IOUG @OAUG1 @odtug @oracleace @viscosityna

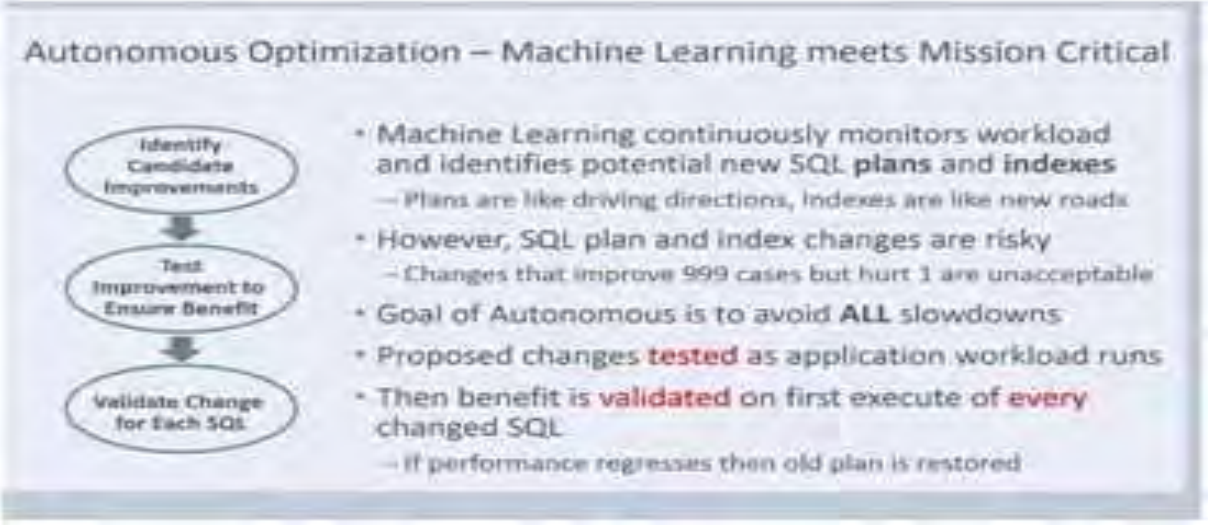

ATP

•Screenshots from Larry Ellison introduction of ATP on 8/7/18 to describe update from Autonomous 18c to Autonomous 19c

19c ORACLE Database

Fast 18c to 19c Upgrade

Autonomous Optimization – Machine Learning meets Mission Critical

- Machine Learning continuously monitors workload and identifies potential new SQL **plans** and **indexes**
  - Plans are like driving directions, Indexes are like new roads
- However, SQL plan and index changes are risky
  - Changes that improve 999 cases but hurt 1 are unacceptable
- Goal of Autonomous is to avoid **ALL** slowdowns
- Proposed changes tested as application workload runs
- Then benefit is validated on first execute of every changed SQL
  - If performance regresses then old plan is restored

ATP – Provisioning a Database (PDB)

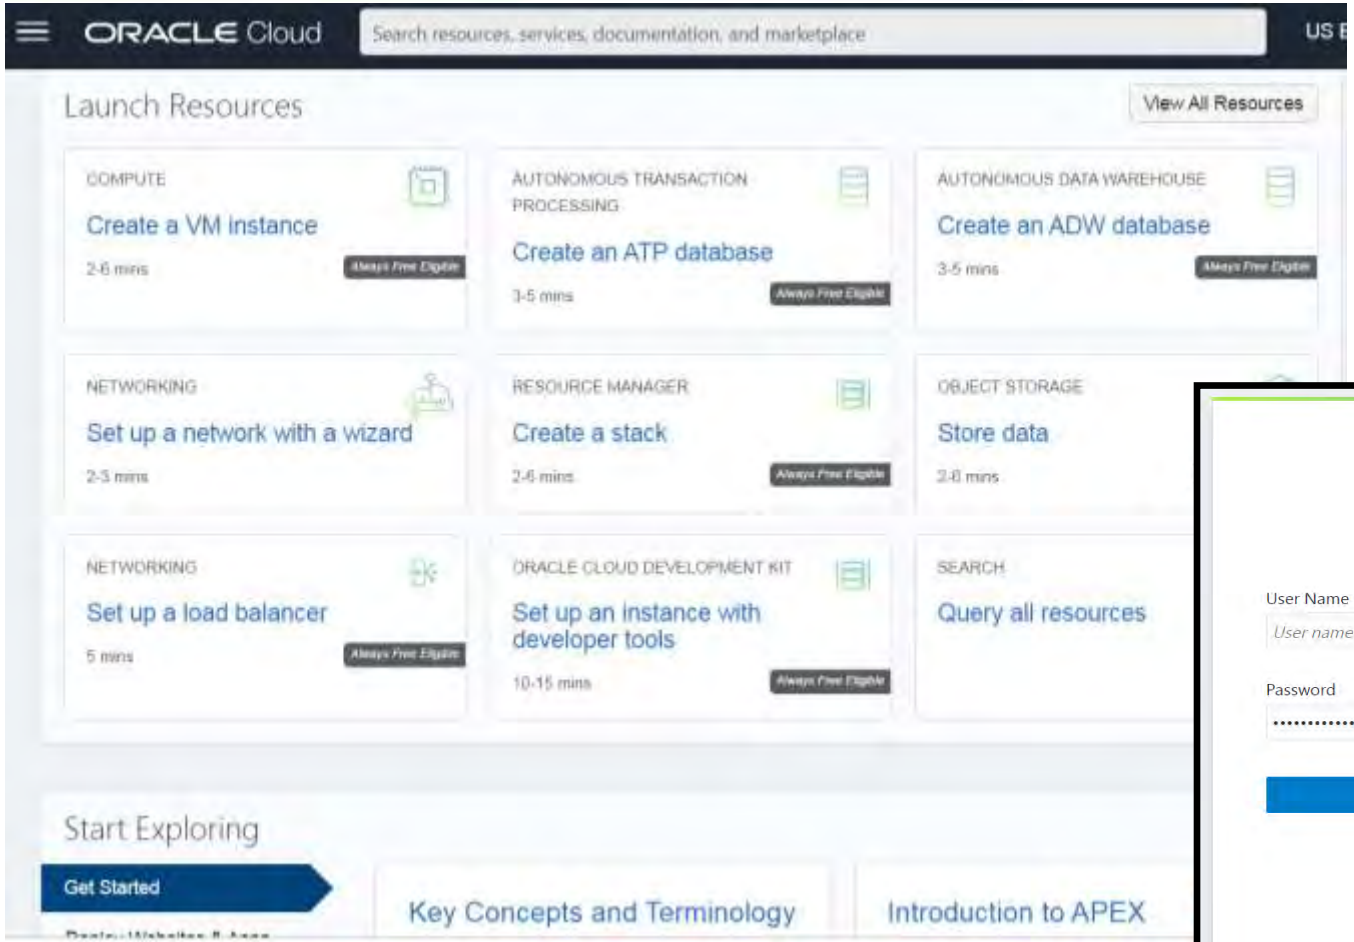

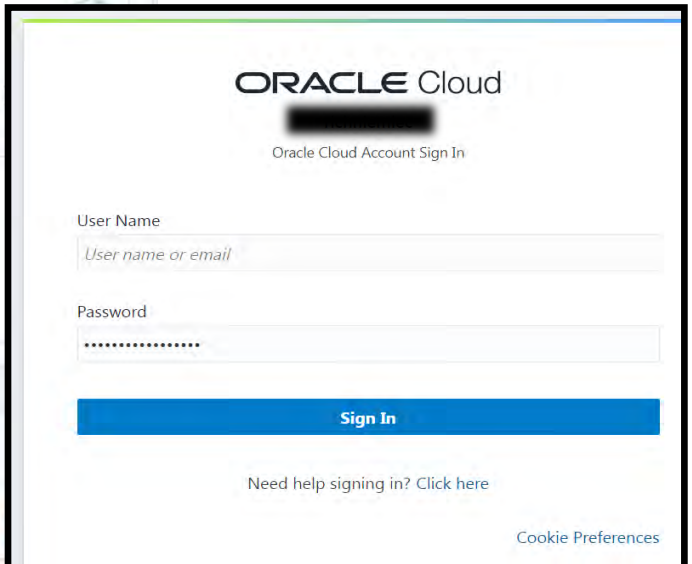

ATP – Provisioning a Database (PDB)

ATP

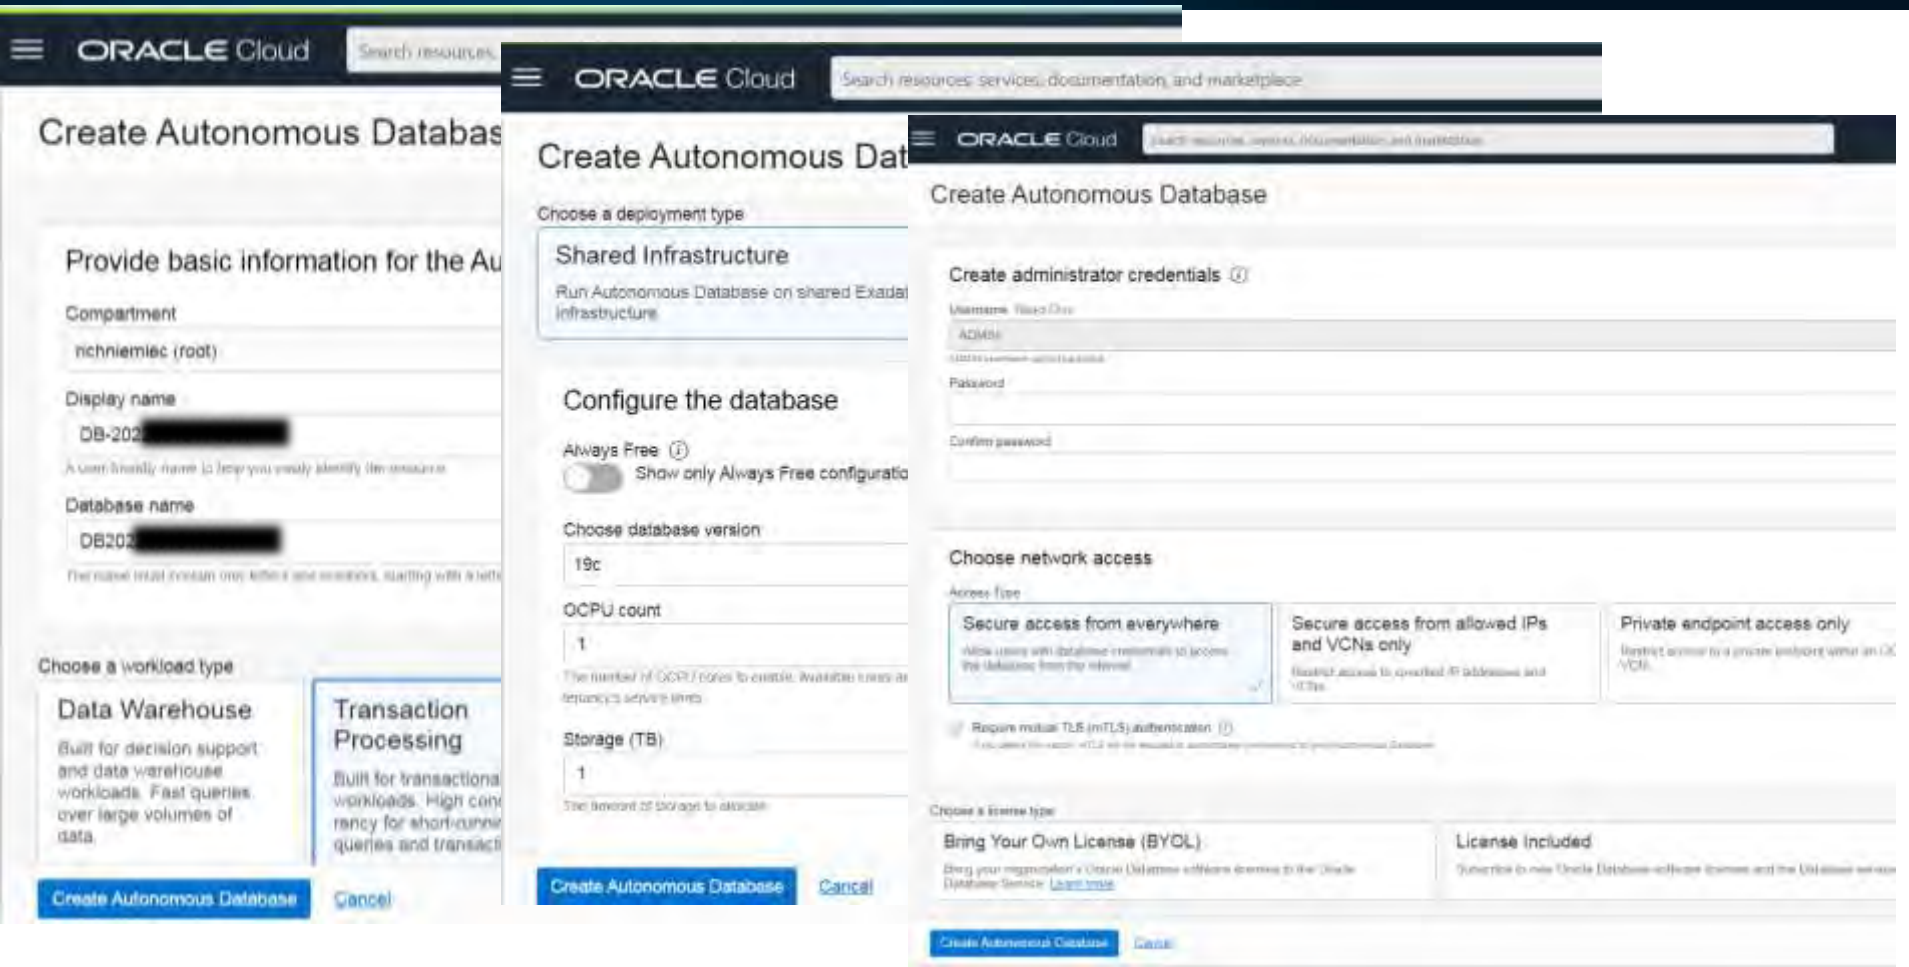

ATP – Provisioning a Database (*2 minutes*)

ATP

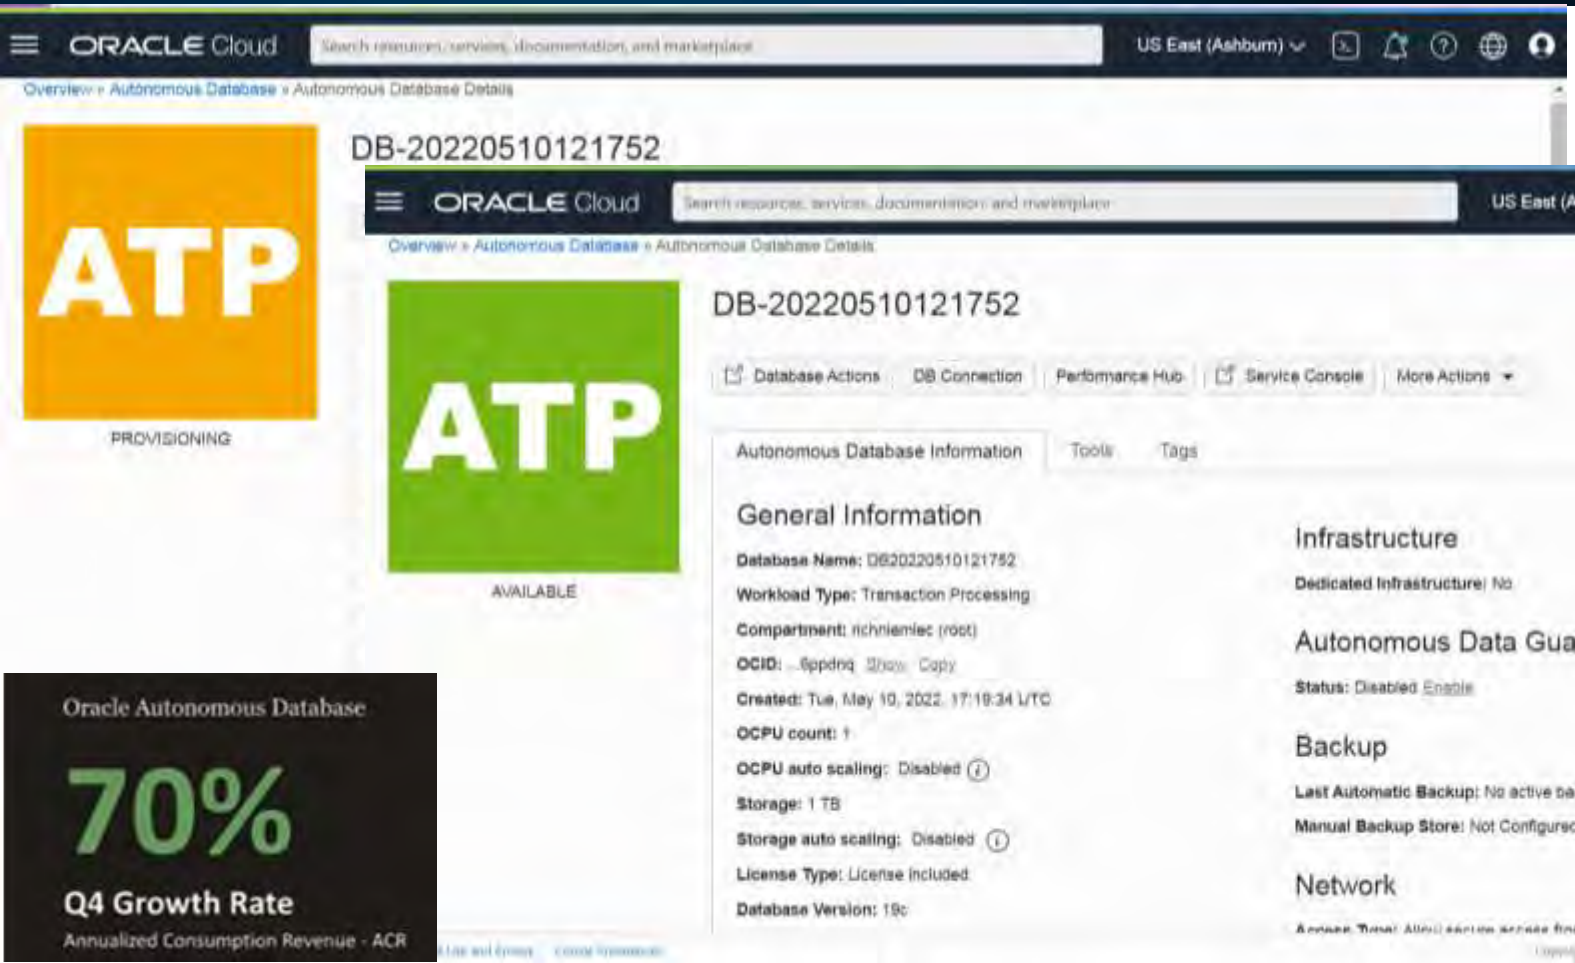

ATP – Stop Database *(25 sec.)*

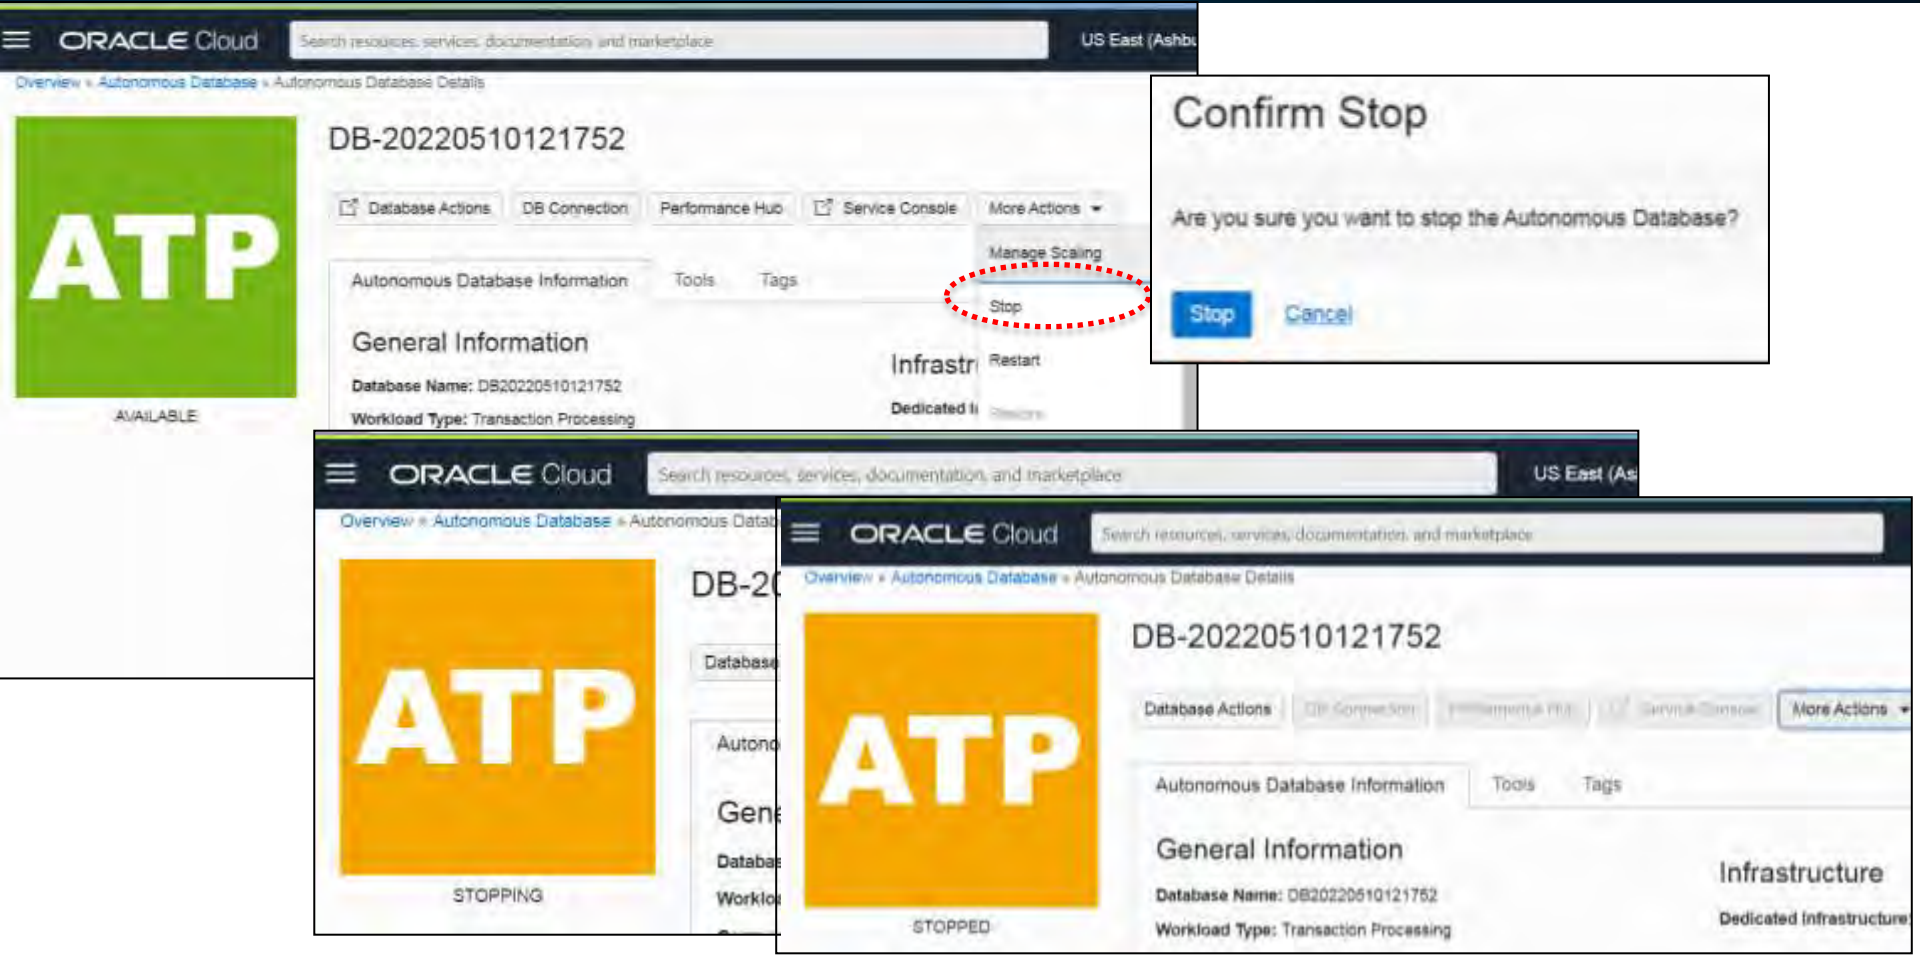

ATP – Start Database *(30 sec.)*

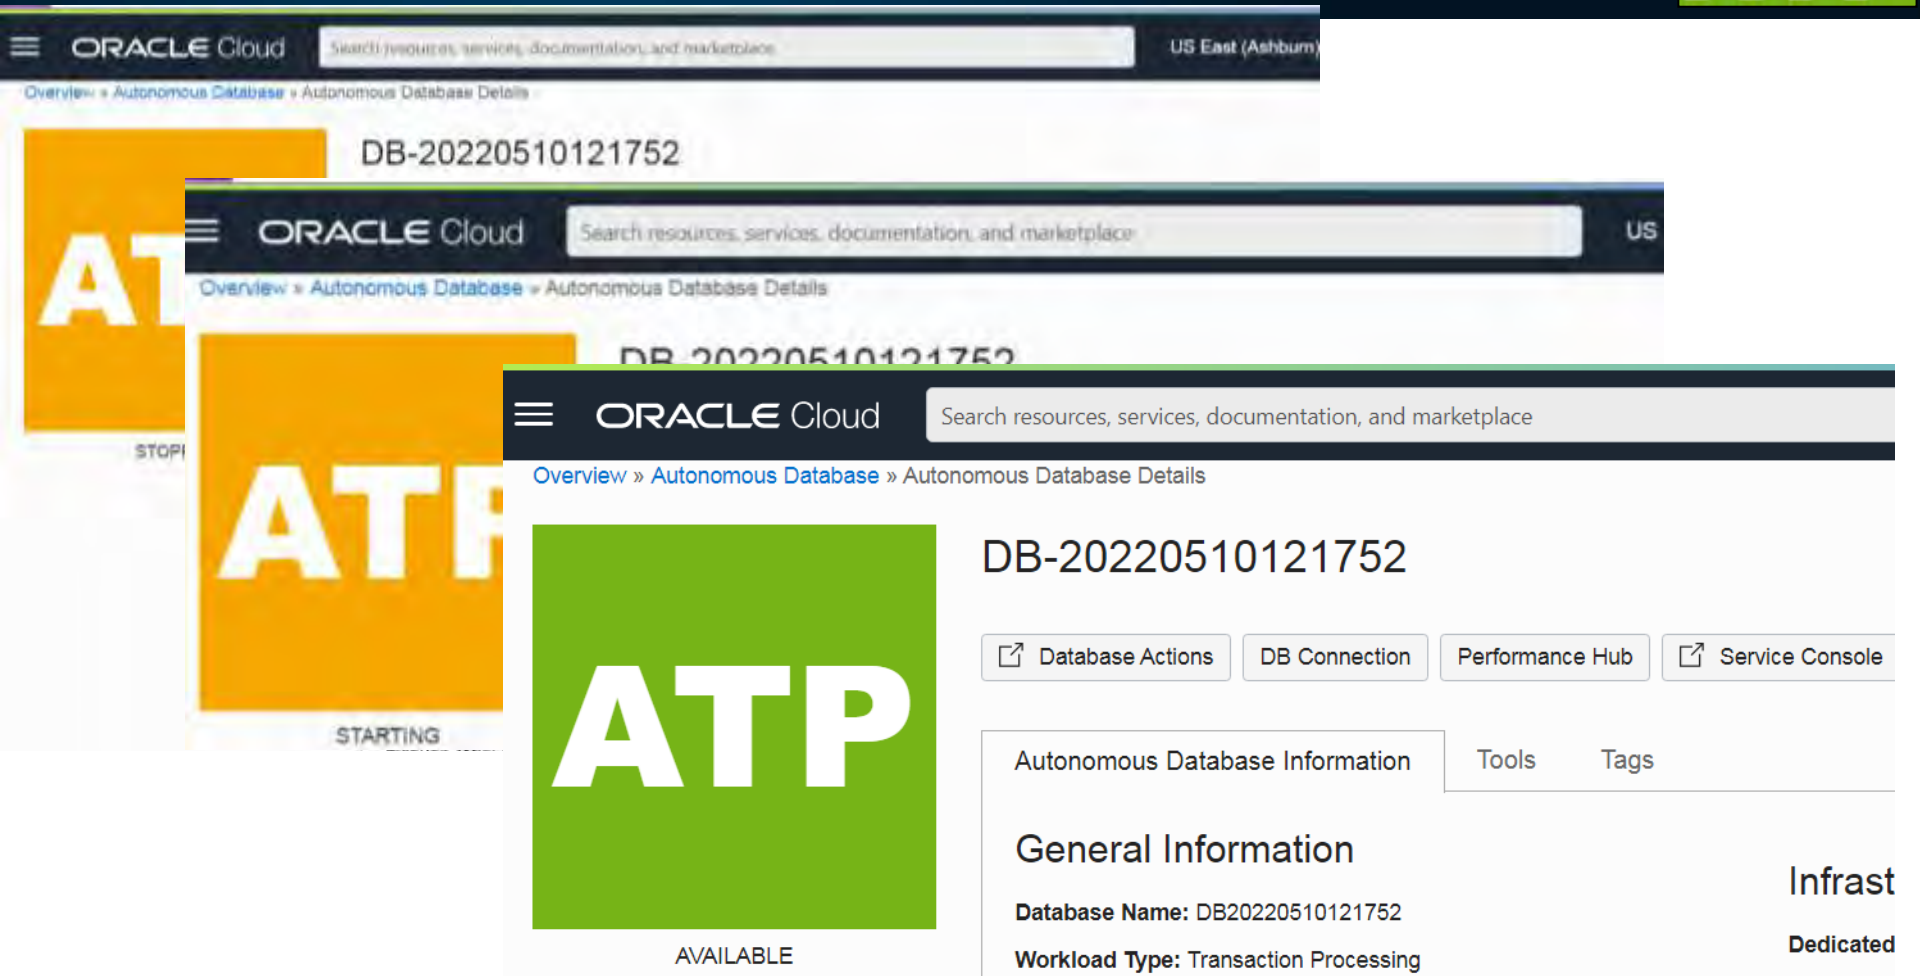

ATP – Scaling Database *(50 sec.)*

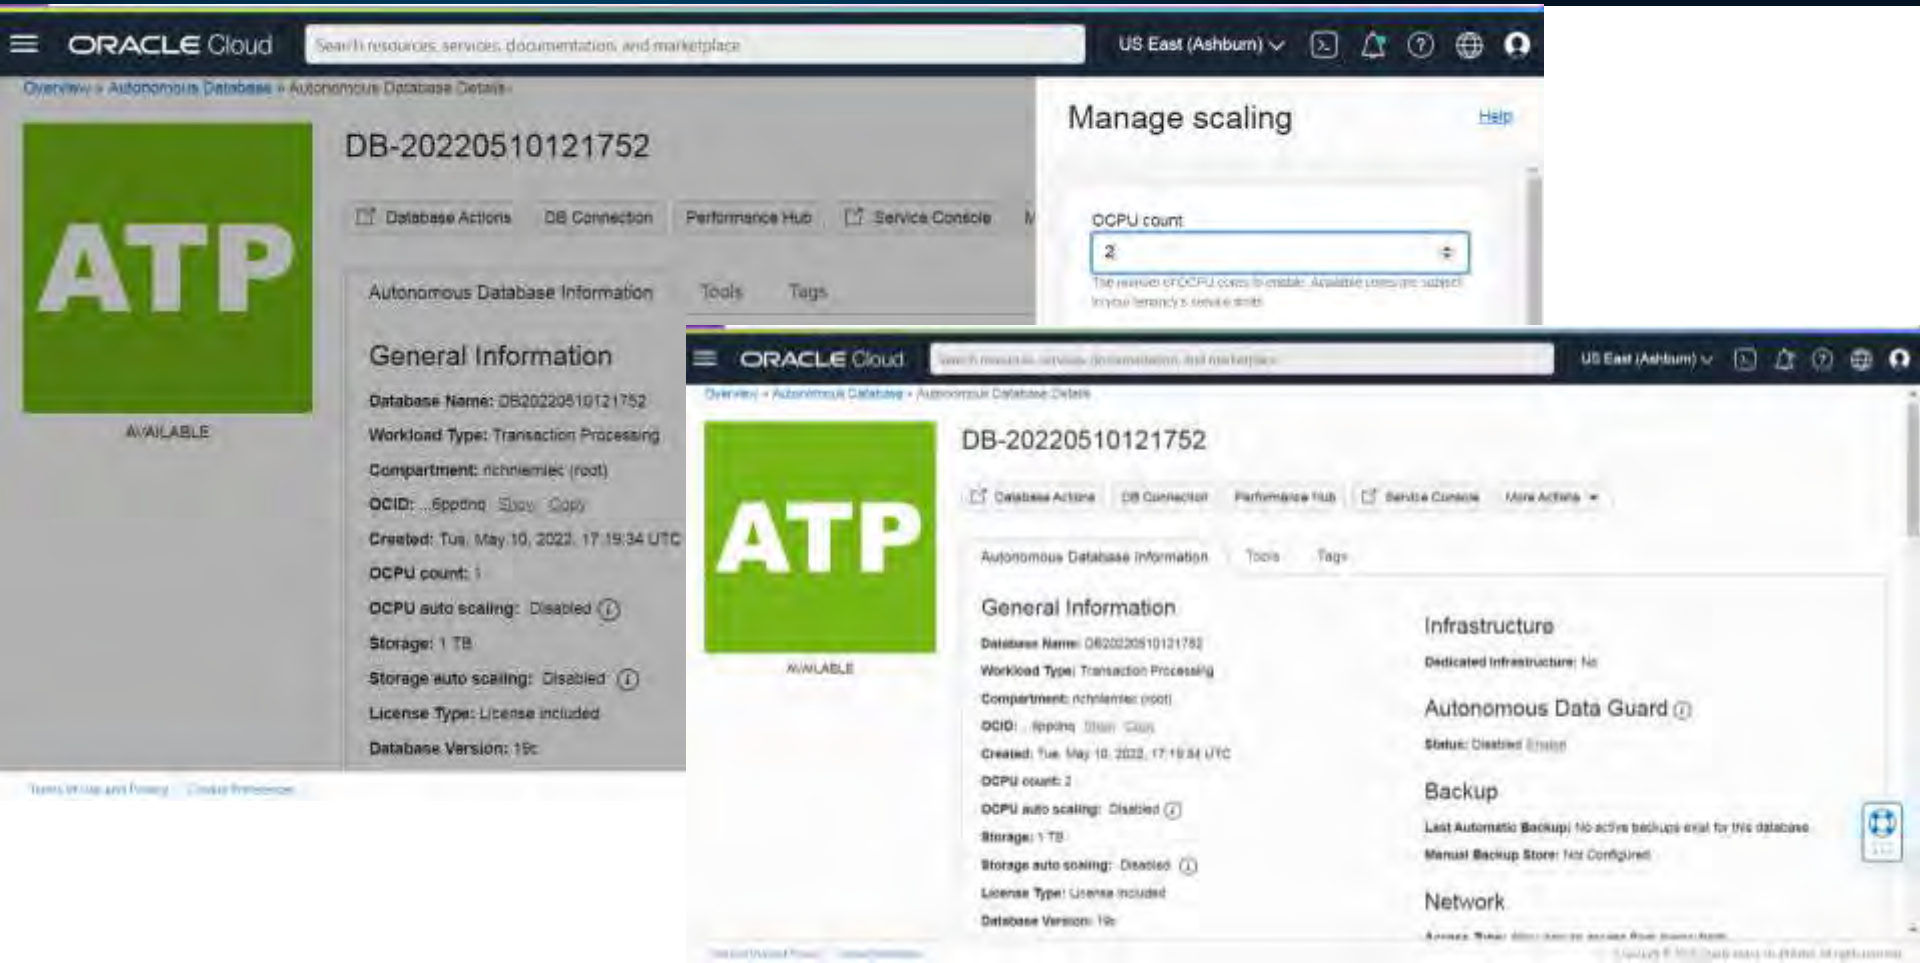

ATP – Restart Database (Stop/Start) (40 sec.)

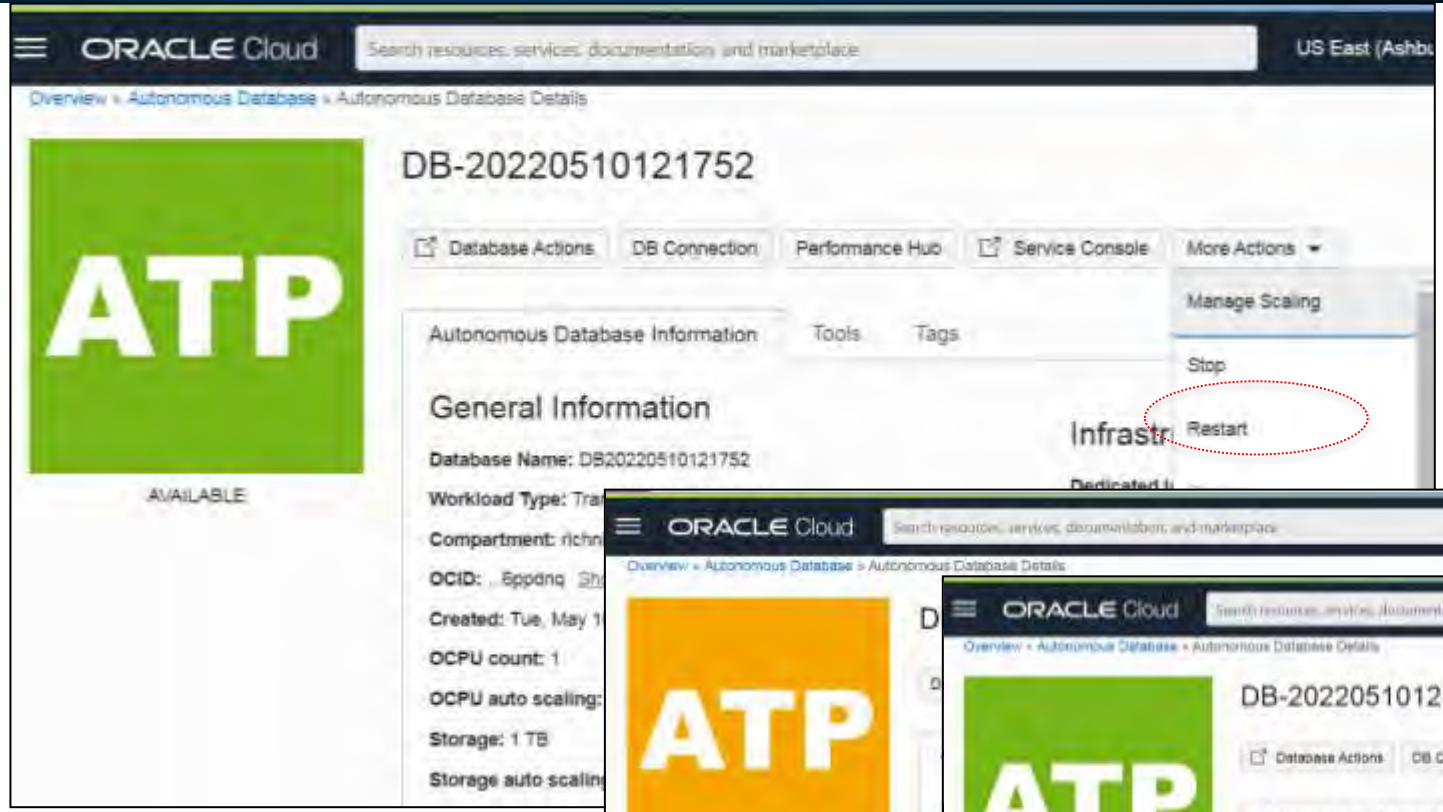

Automatic Scheduling to Start/Stop DB!

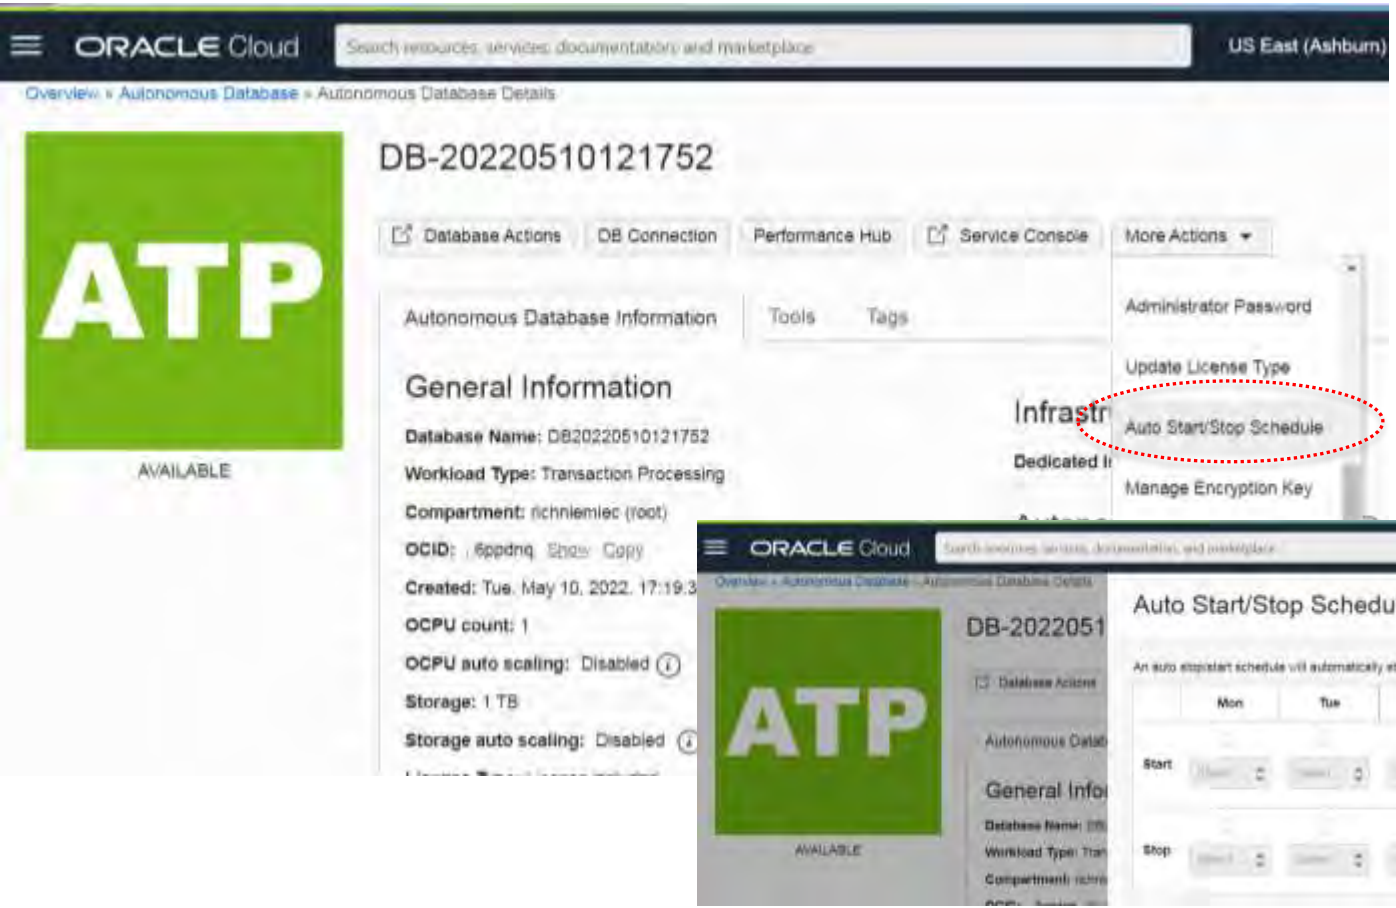

ATP - Create a Clone

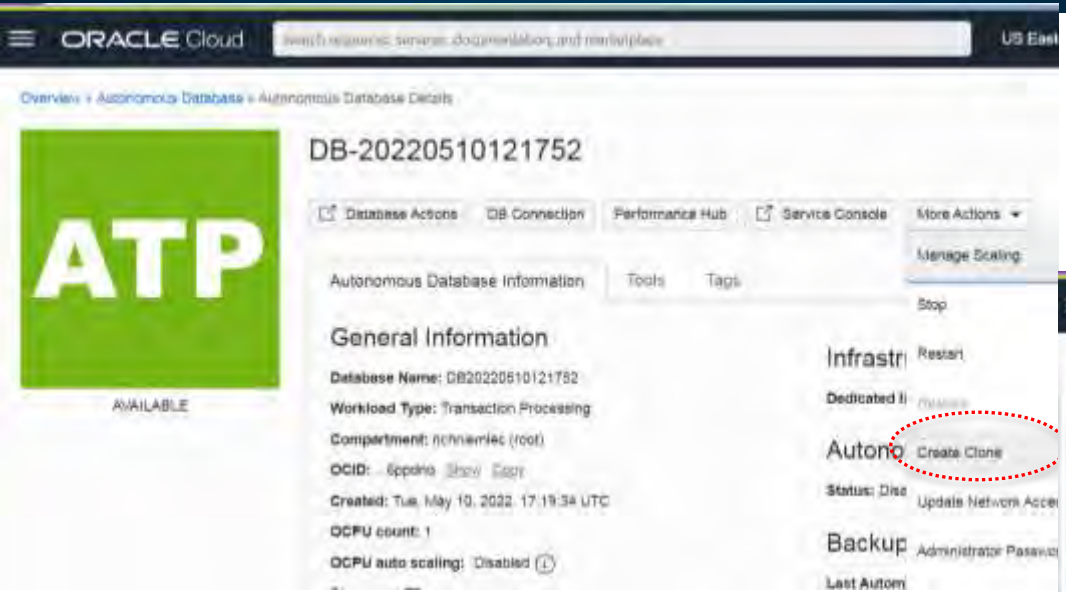

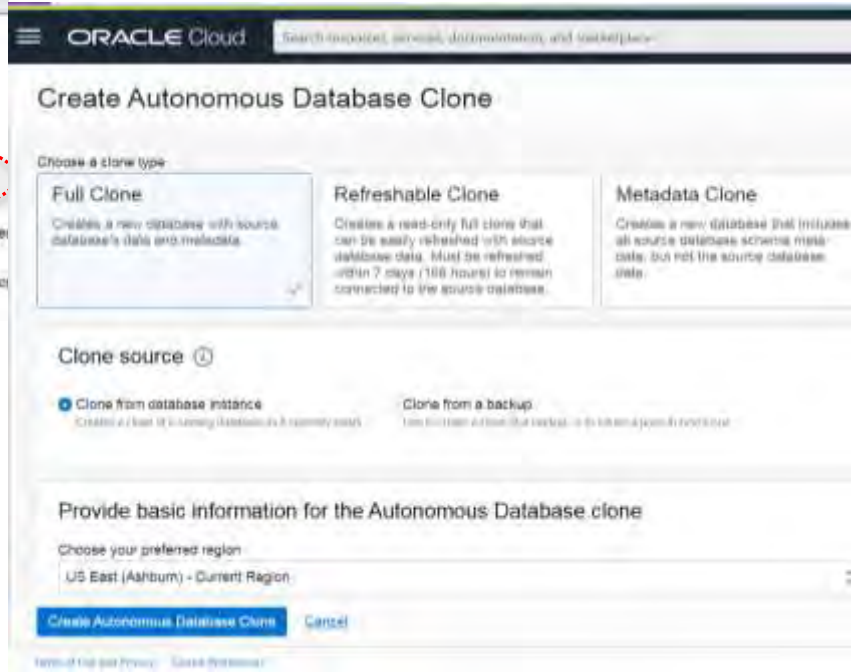

Database Actions are Numerous!

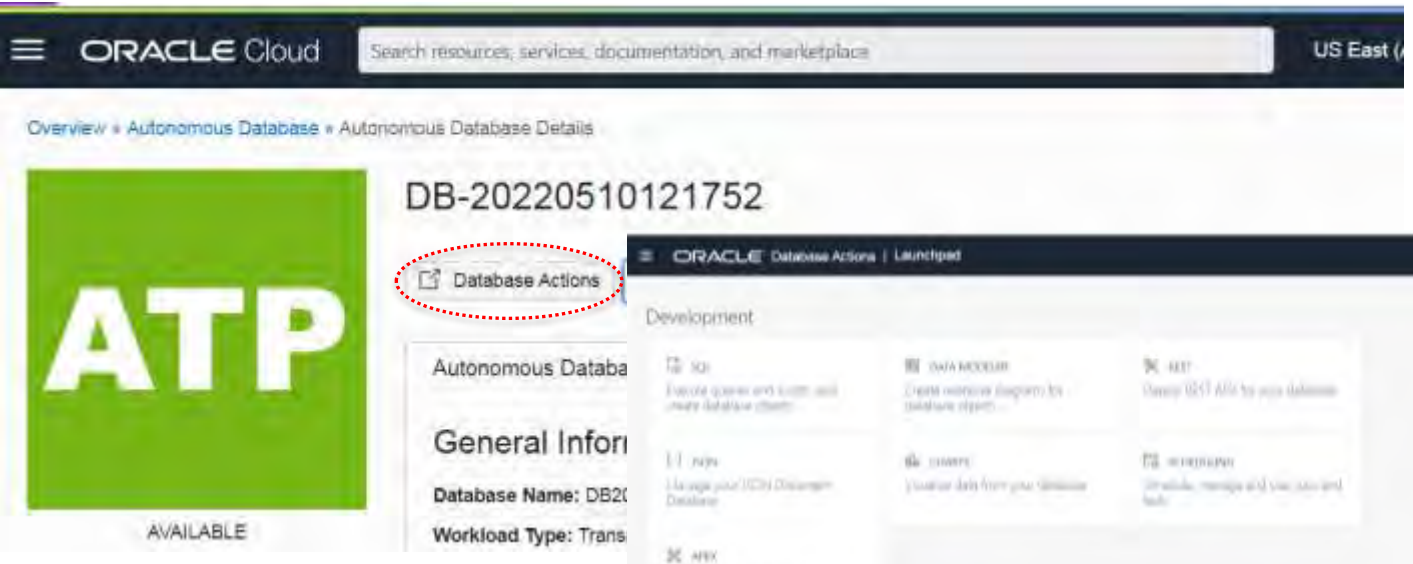

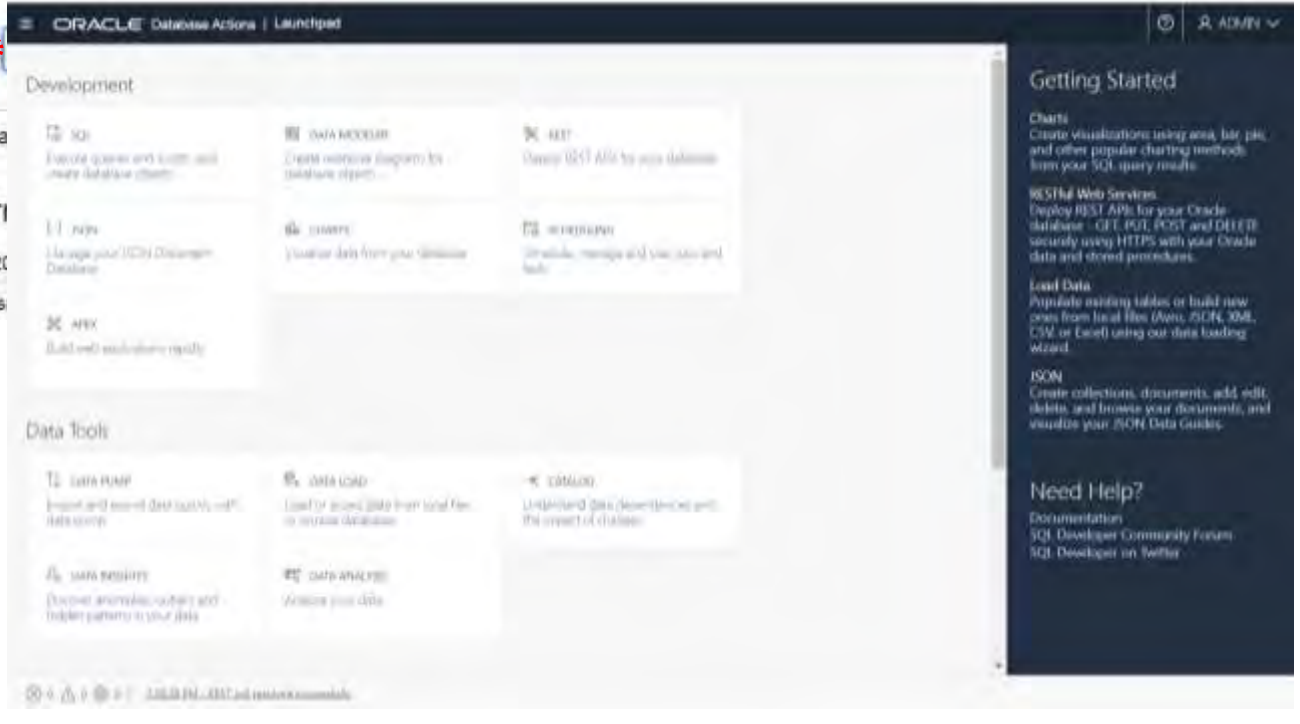

ATP - Performance Hub

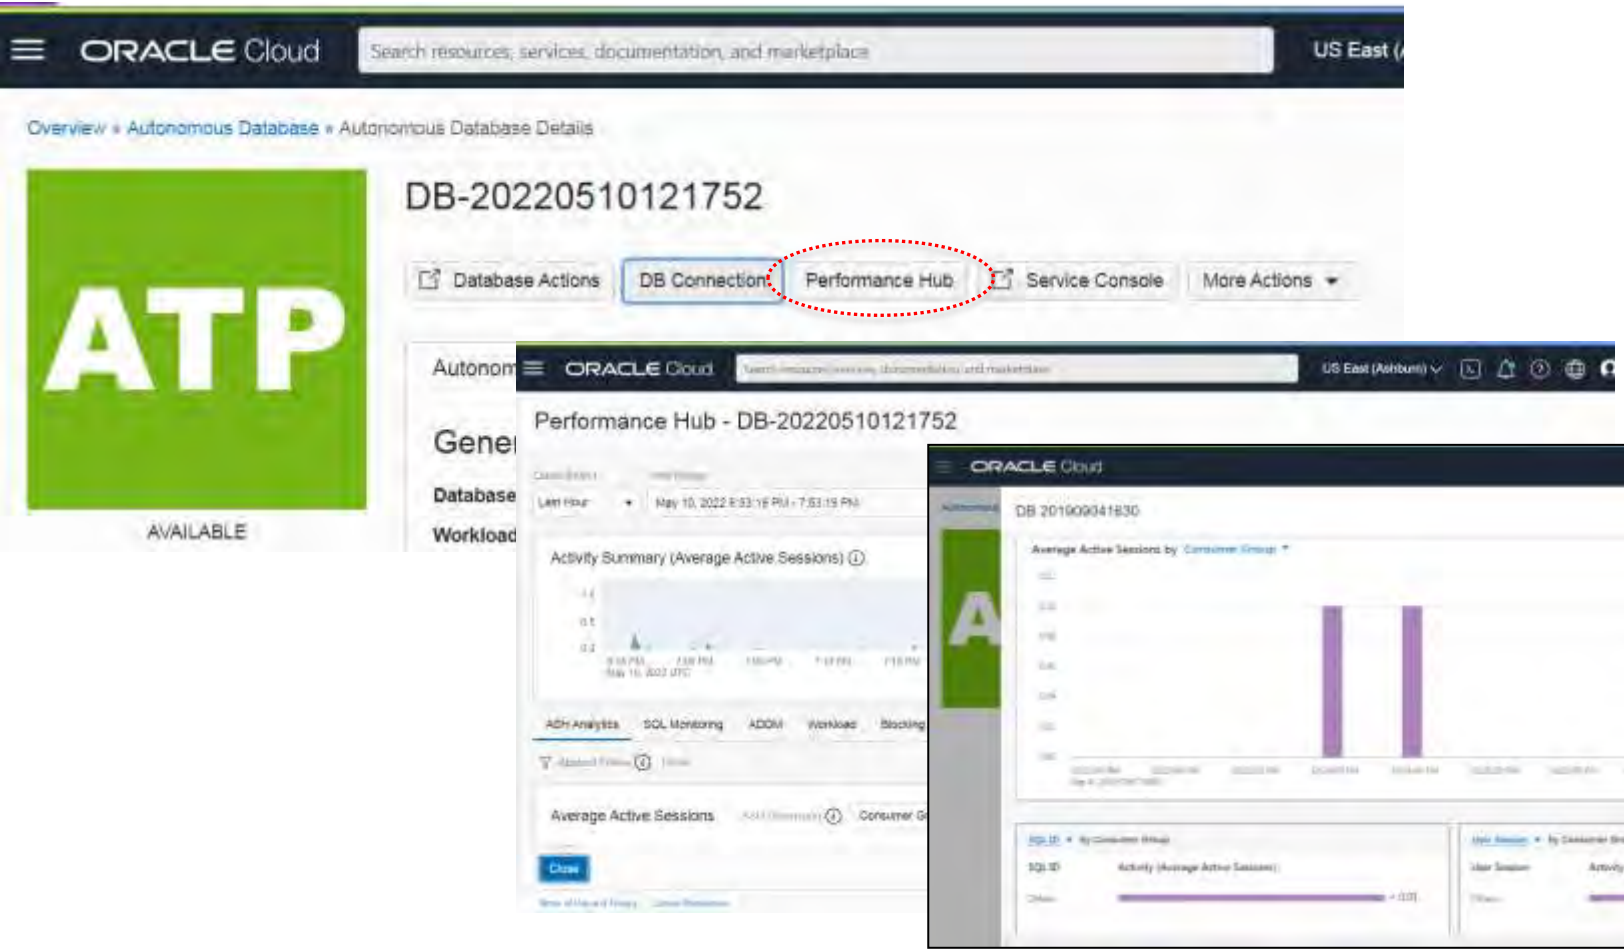

ATP - Service Console (Overview/Activity)

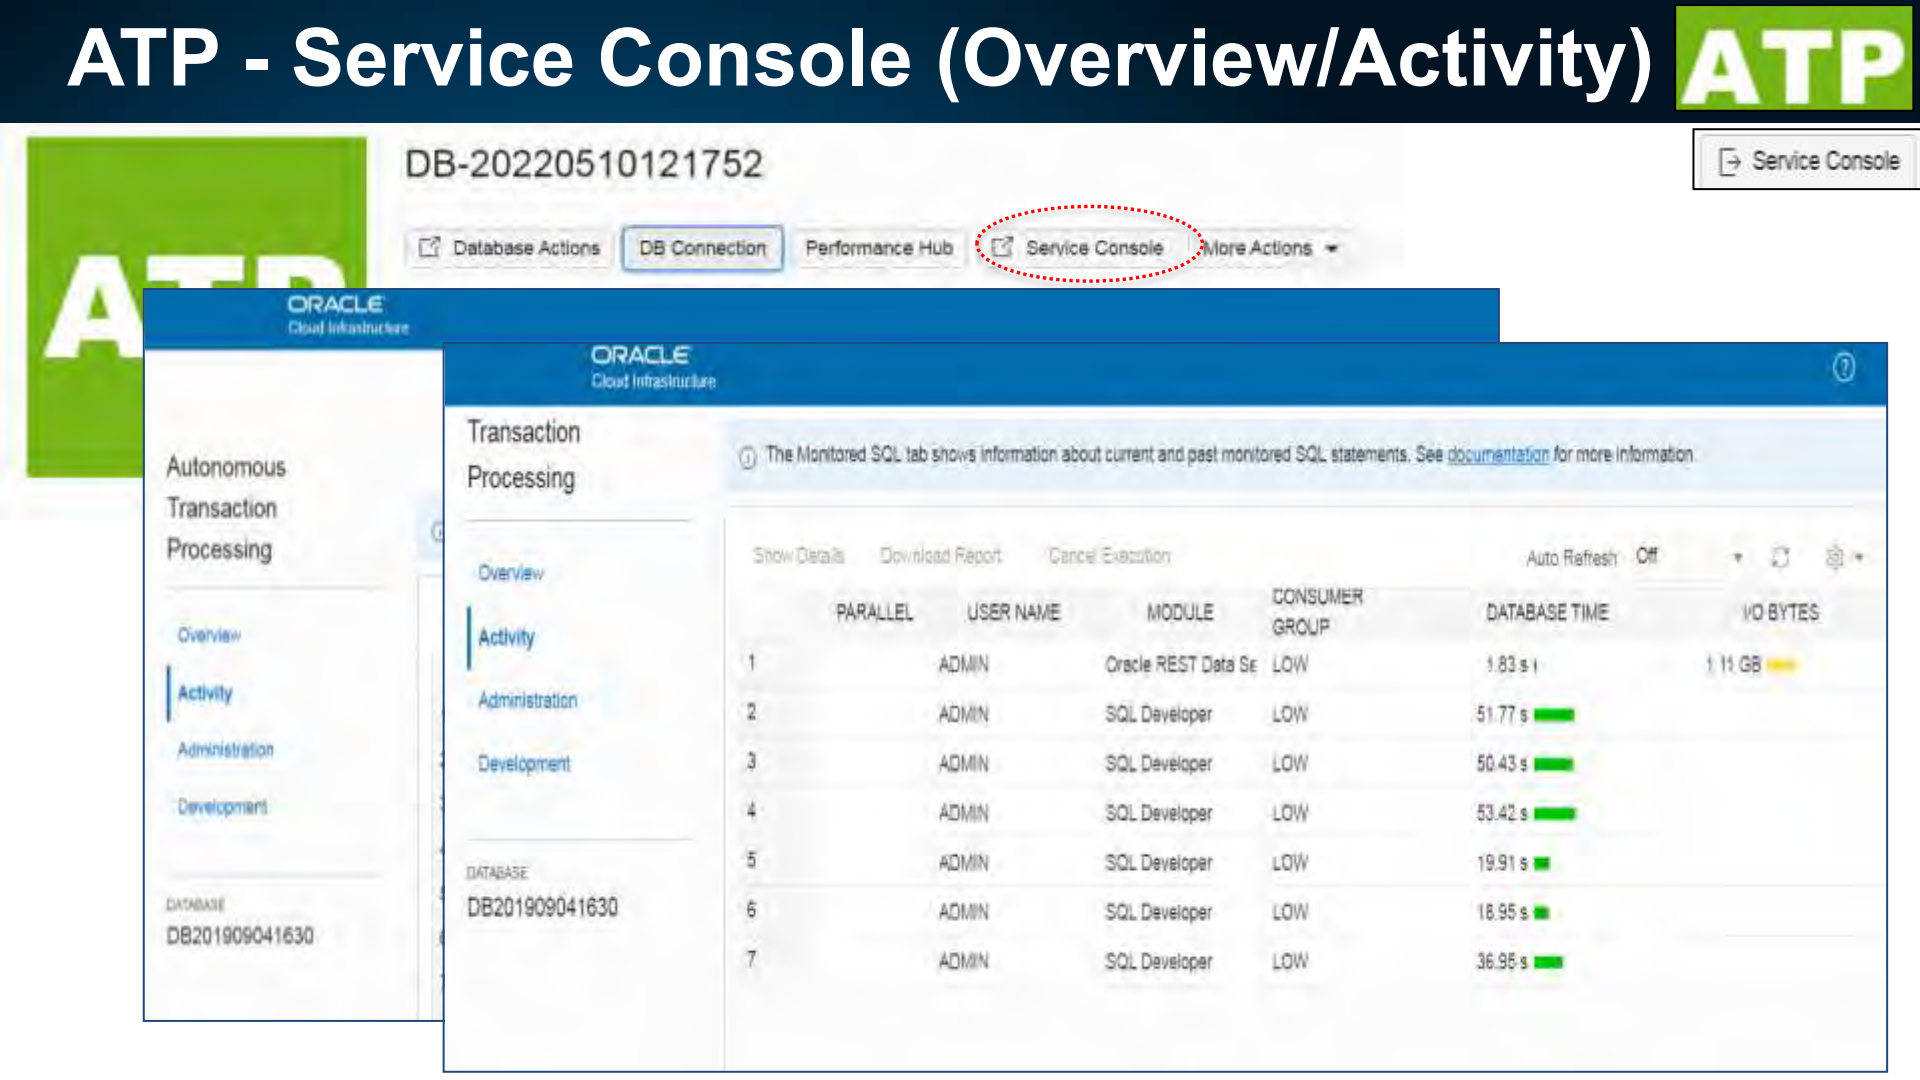

ADW - Provision Database *(1 min. 20 sec.)*

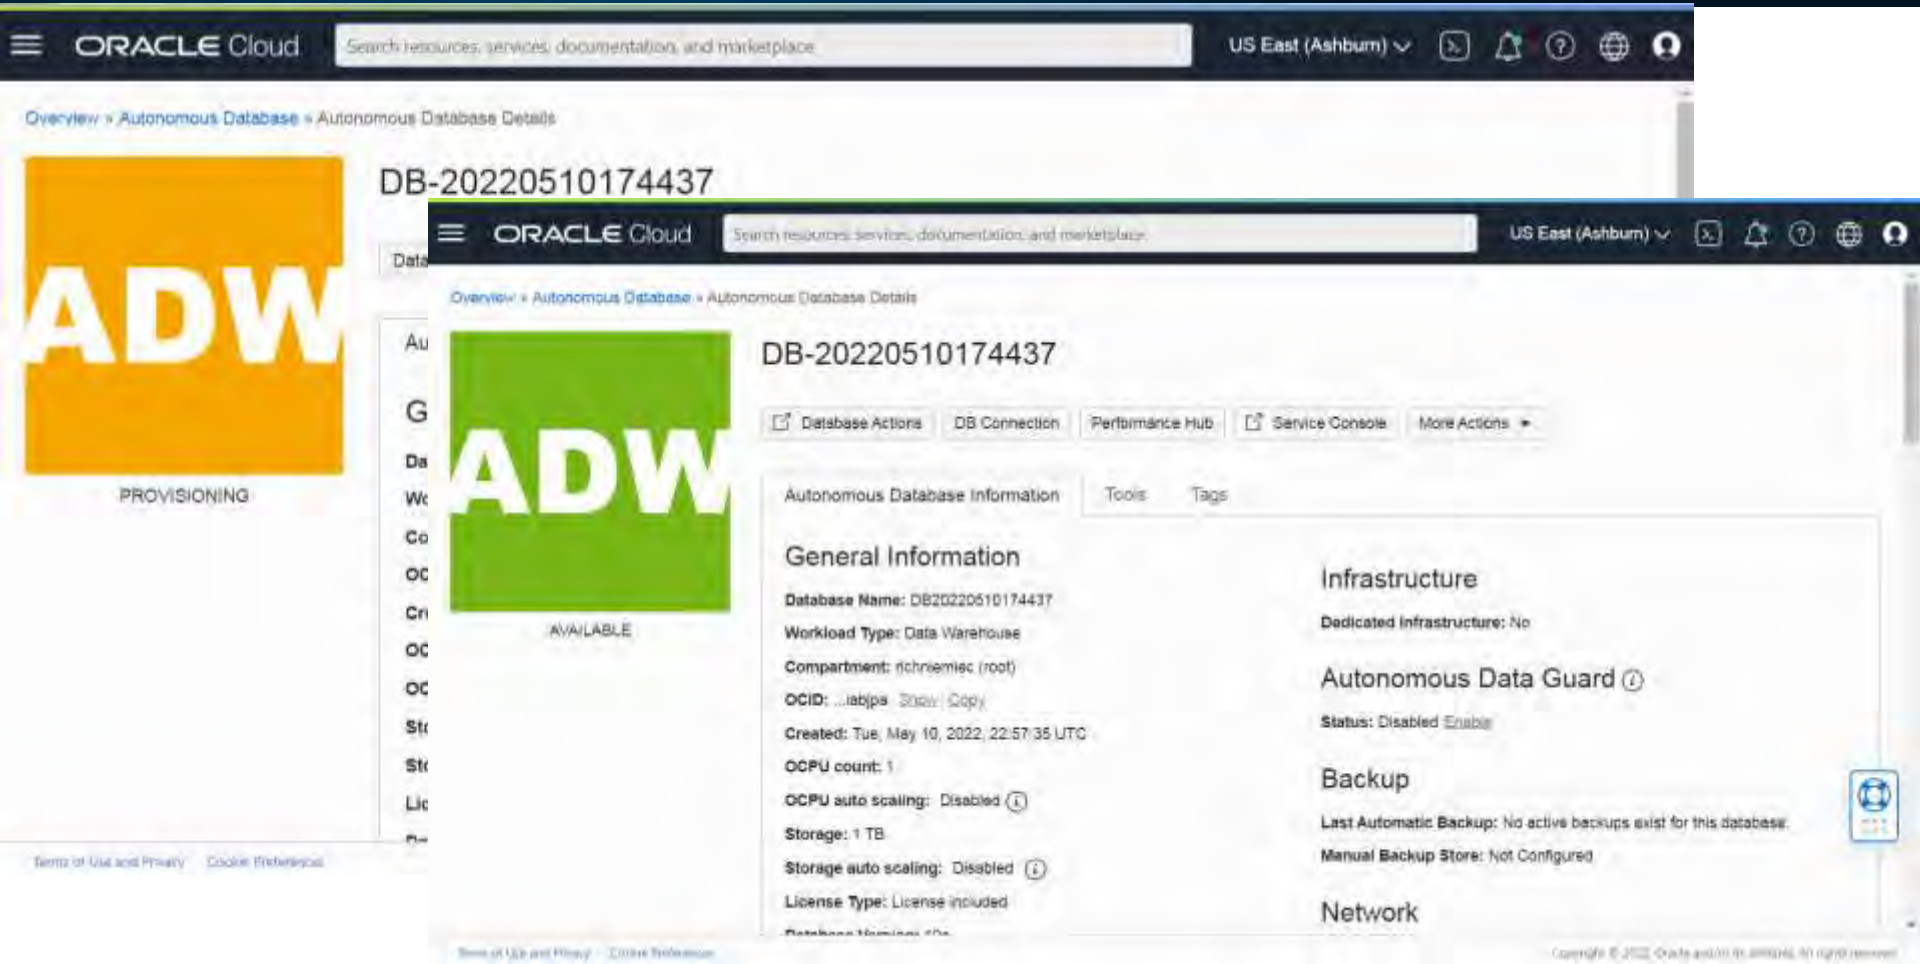

Optimized for Autonomous Database

- Managed service for data warehouse workloads
- **Easy to import data** from various source: existing Oracle installation, external tables...
- ADB Originally configured for **best DW (ADW) performance**: no parameter settings required. **New ATP** with auto-indexes (19c+ only)!
- Optimizer **statistics are automatically gathered and maintained**; pre-configured for accuracy and efficiency
- **No SQL tuning or hints needed**
- DOP for SQL is automatically chosen for best performance

ADW vs. ATP - *sqlmaria.com blog*

	Autonomous Data Warehouse	Autonomous Transaction Processing
	Optimizes Complex SQL	Optimizes Response Time
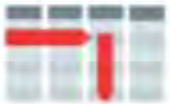	Columnar Format	Row Format
	Creates Data Summaries	Creates Indexes **(Auto-Indexing on 19c)**
	Memory Speeds Joins, Aggs **(Columnar In-Memory Usage)**	Memory for Caching, No IO
	Statistics gather as part of bulk load operations	Statistics gather as part of DML operations

Autonomous Database: Performance vs. AWS

Rich Niemiec @RichNiemiec · Oct 1

Replying to @RichNiemiec

Demo by #LarryEllison of #oracle vs #Amazon is 15x faster "but you have to be willing to pay less." #ioug #viscosityna #oracleace

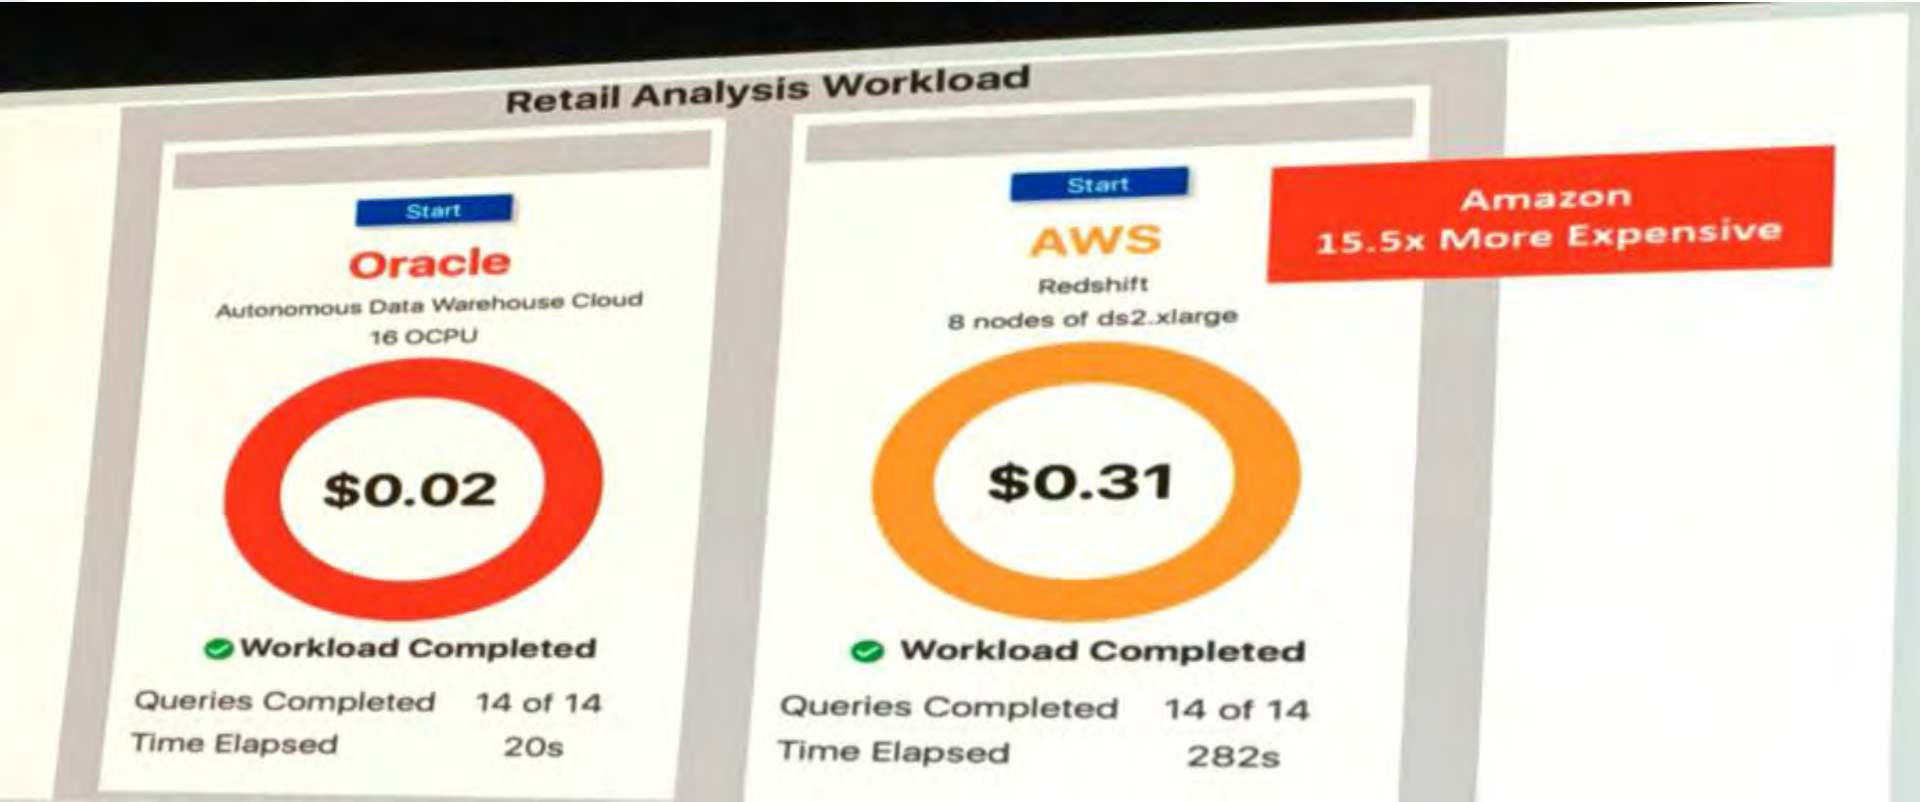

Autonomous Database Summary

MY ADW/ATP Benefits:

- Oracle19c PDB –**19c features** (Server, Storage, VM, RAC)
- 3rd Party **BI and Integration tools** as well as Oracle's
- Pay/Scale what you need only! **Shut down idle instances**
- Runs on Oracle Cloud or YOUR Cloud at Customer
- Auto Full stack **patching** once/quarter with pre-checks prior
- **Automatic Patching** rolls across nodes for availability
- Automatic Quarterly patches or off-cycle Security Patches
- **ML Security**, Fully Encrypted, with Data Vault & Masking
- Oracle "break glass" access when approved by client
- Monitors CPU, Disk, Memory with Alerts & Resource Mgr.
- Cluster Health Monitor & Error Handling use ML algorithms

Results in Less Admin, More Uptime & Better Security!

Autonomous Database Summary

Database management simplified…

* Sean Stacey Session

1 Auto-Provisioning
Automatically deploys mission-critical databases (RAC on Exadata infrastructure) which are fault-tolerant and highly available. Enables seamless scale-out, protection in case of a server failure and allows updates to be applied in a rolling fashion, while apps continue to run.

2 Auto-Configuration
Automatically configures the database to optimize for specific workloads. Everything from the memory configuration, the data formats, and access structures are optimized to improve performance. Customers can simply load data and go.

3 Auto-Indexing
Automatically monitors workload and detects missing indexes that could accelerate applications. It validates each index to ensure its benefit, before implementing it and uses machine learning to learn from its own mistakes.

4 Auto-Scaling
Automatically scales compute resources when needed by workload. All scaling occurs online, while the application continuously runs. Enables true pay per use.

5 Automated Data Protection
Automatically protect sensitive and regulated data in the database, all via a unified management console. Assess the security of your configuration, users, sensitive data, and unusual database activities.

6 Automated Security
Automatic encryption for the entire database, backups and all network connections. No access to OS or admin privileges prevents phishing attacks. Protects the system from both cloud operations and any malicious internal users.

7 Auto-Backups
Automatic daily backup of database or on-demand. Restore or recover a database to any point-in-time you specify in the last 60 days.

8 Auto-Patching
Automatically patch or upgrade with zero downtime. Applications continue to run as patching occurs in a round-robin fashion across RAC nodes or servers.

9 Automated Detection and Resolution
Using pattern recognition, hardware failures are automatically predicted without long timeouts. IOs are immediately redirected around unhealthy devices to avoid database hangs. Continuous monitoring for each database automatically generates service requests for any deviation.

10 Automatic Failover
Automatic failover with zero-data loss to standby. It's completely transparent to end-user applications.

Autonomous Database - Cost Analysis

ATP ADW

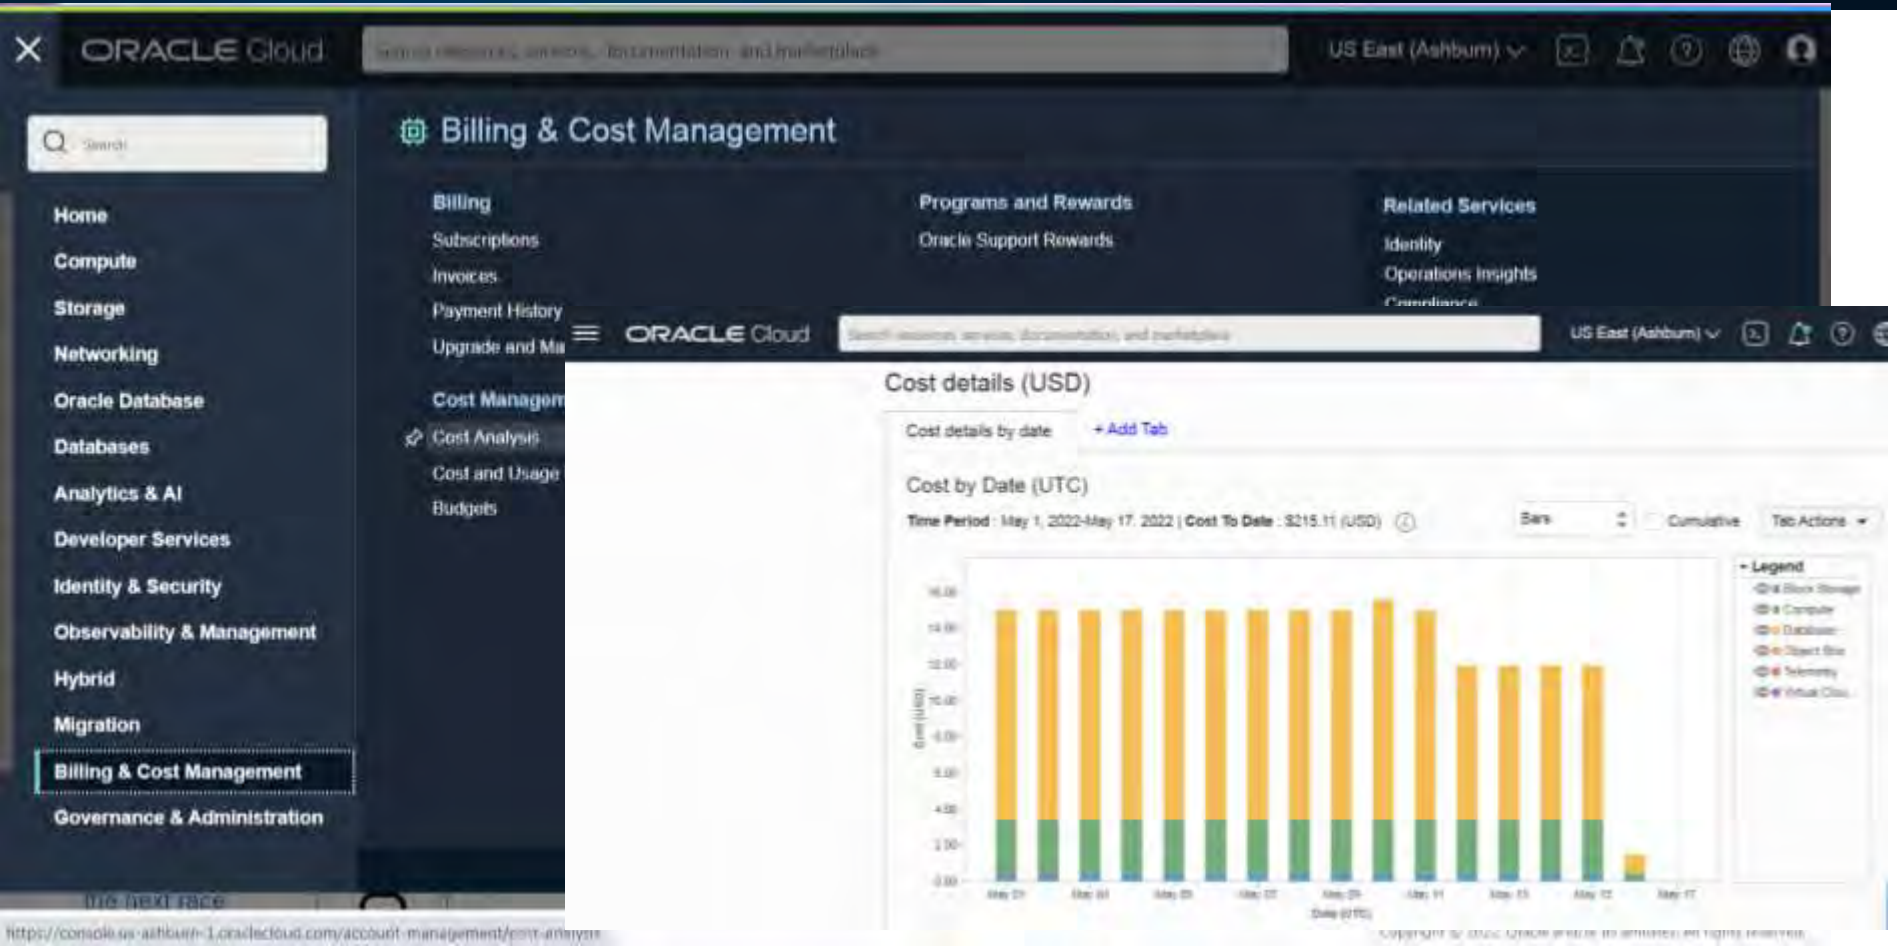

Autonomous DB: ATP-ADW-AJD-APEX *all in root*

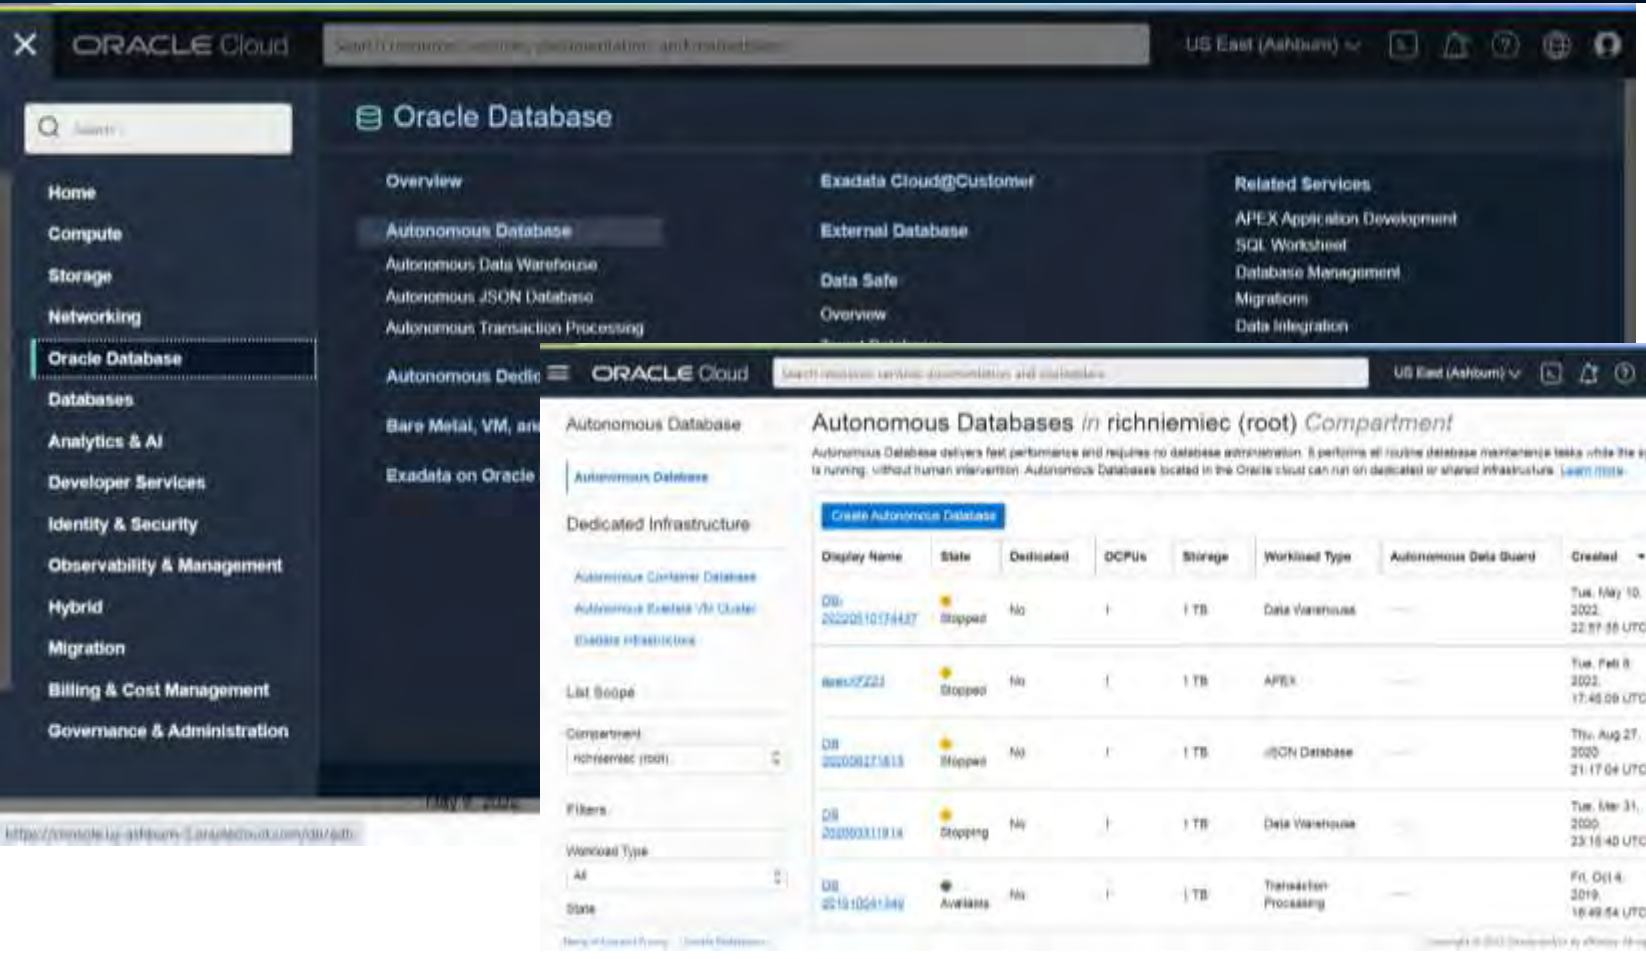

Create/Run Notebooks in Oracle ML

- **Quick Actions**

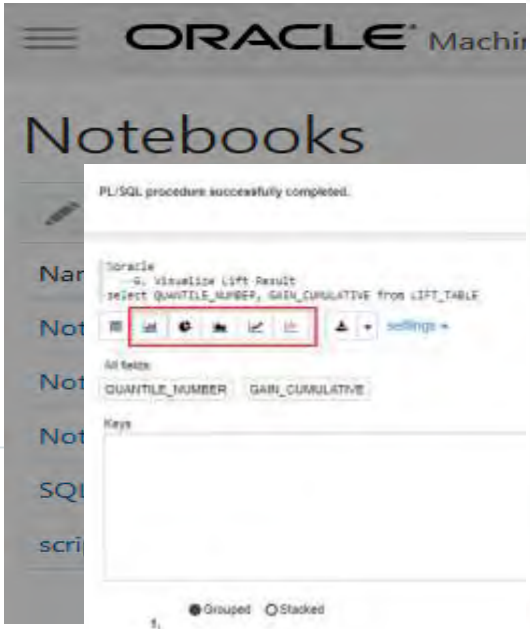

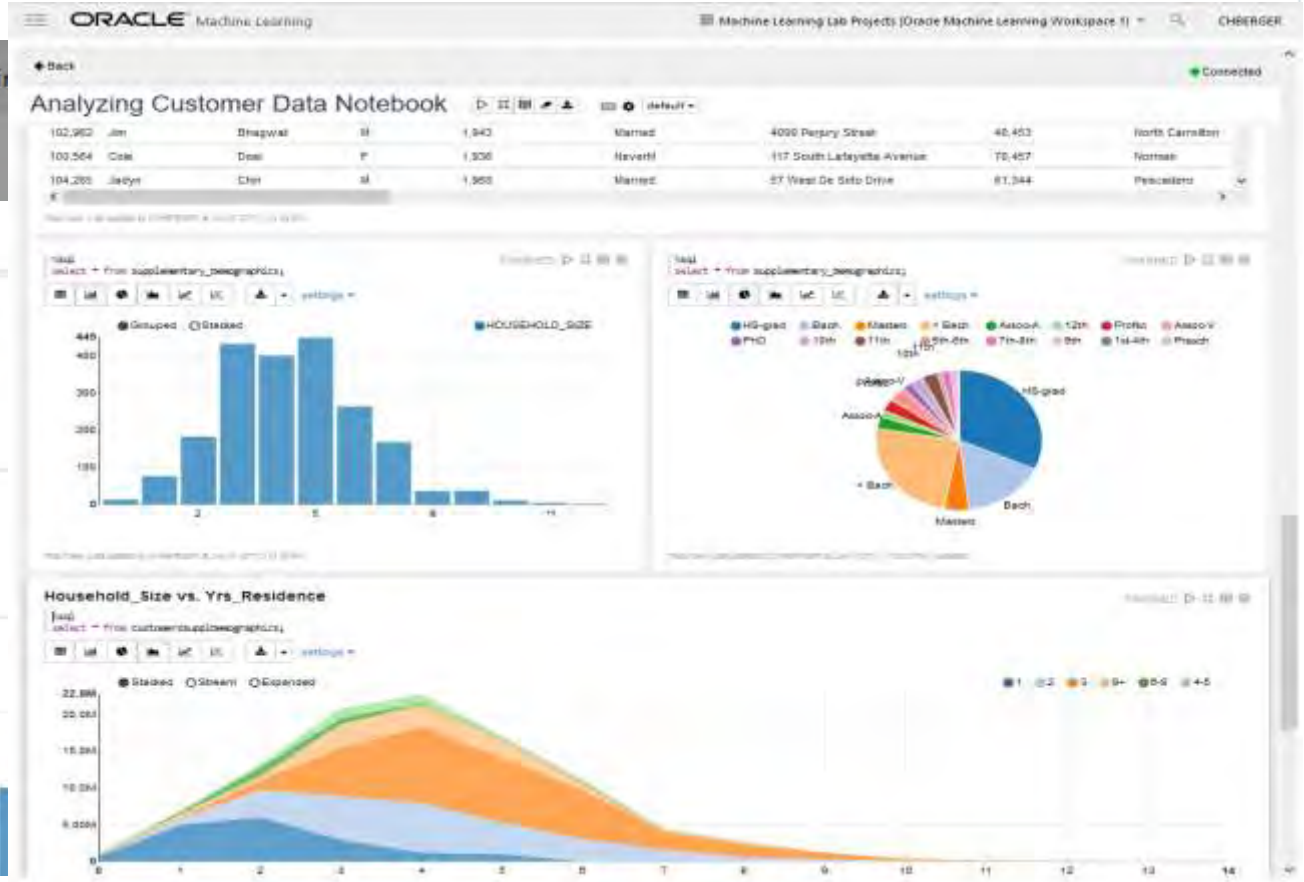

What You Need; Nick of Time *(Predictive Analytics)*

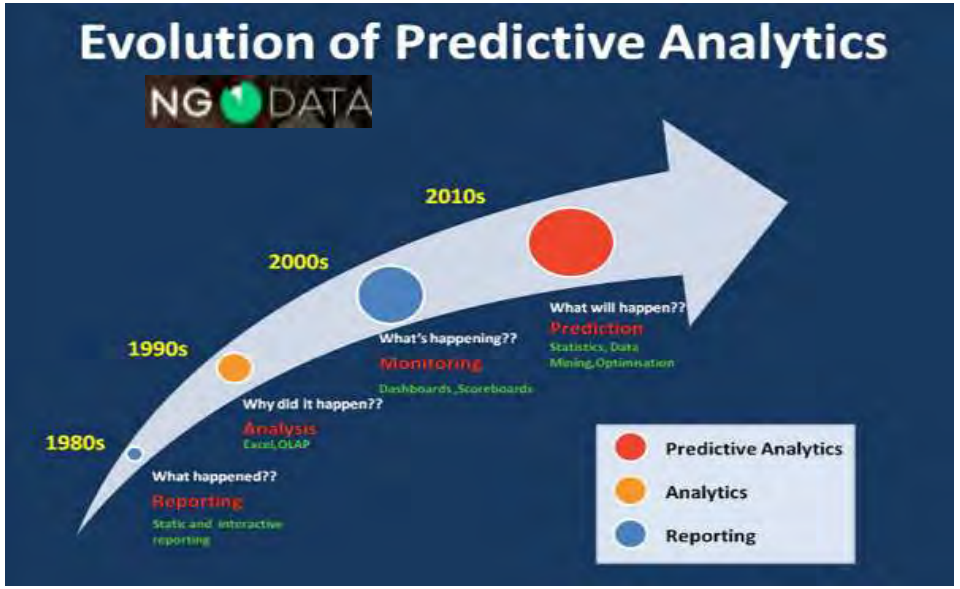

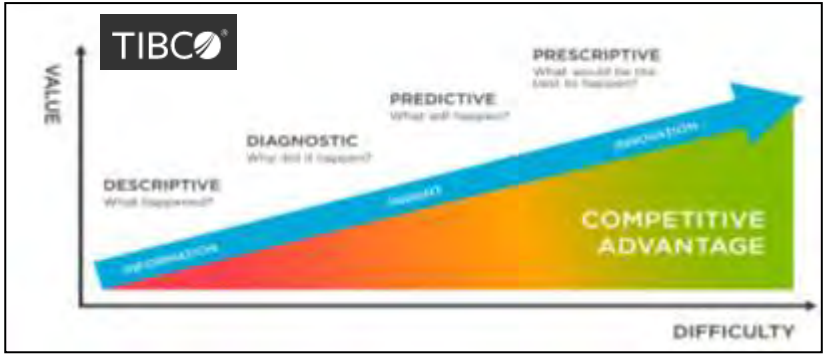

ATP - Service Console (Development)

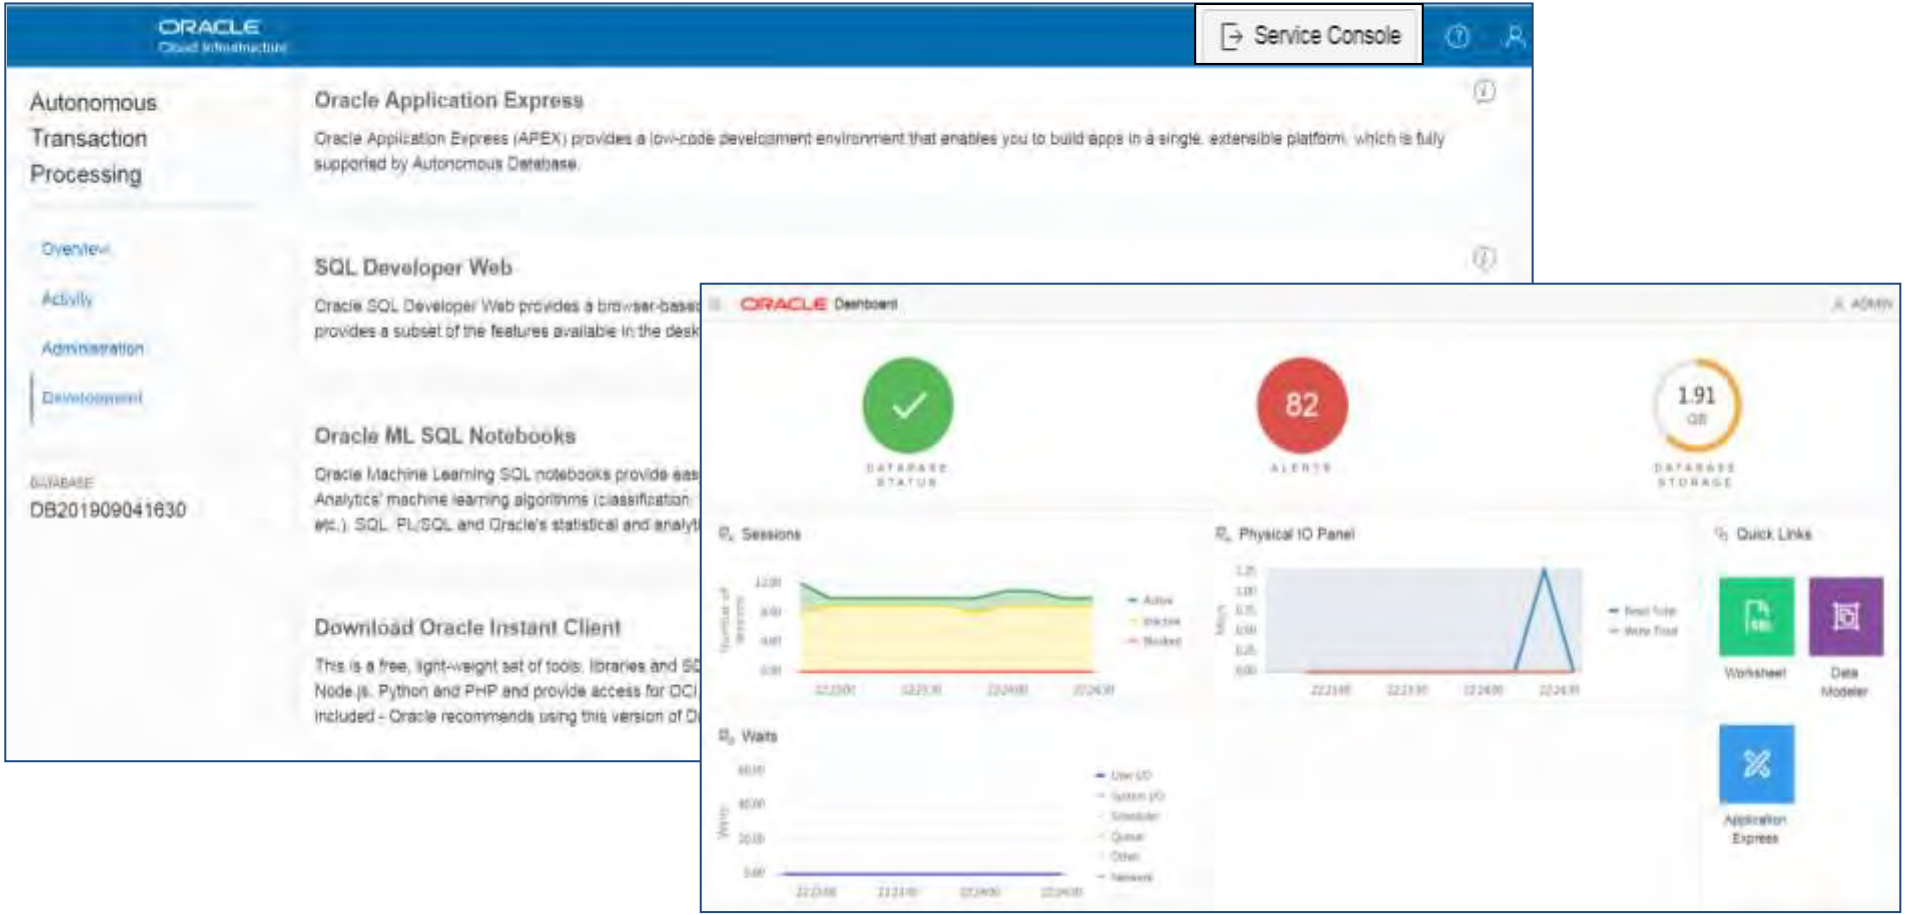

ATP - Service Console (Development/ML)

ORACLE Machine Learning

ADMIN_RICH2 Project

How Do I?

Get Started
Get started with Oracle Machine Learning

Create Notebooks
How to create a notebook

Create Jobs
How to create a job

Manage Permissions
How to manage collaborative permissions in workspaces

Try It
Follow along with a hands on workshop

Quick Actions

Run SQL Statements
Enter and run SQL statements

Run SQL Scripts
Enter and run SQL scripts

Notebooks
The place for data discovery and analytics

Jobs
Schedule notebooks to run at certain times

Examples
Check out some examples

Machine Learning connection to ADW/ATP - SQL

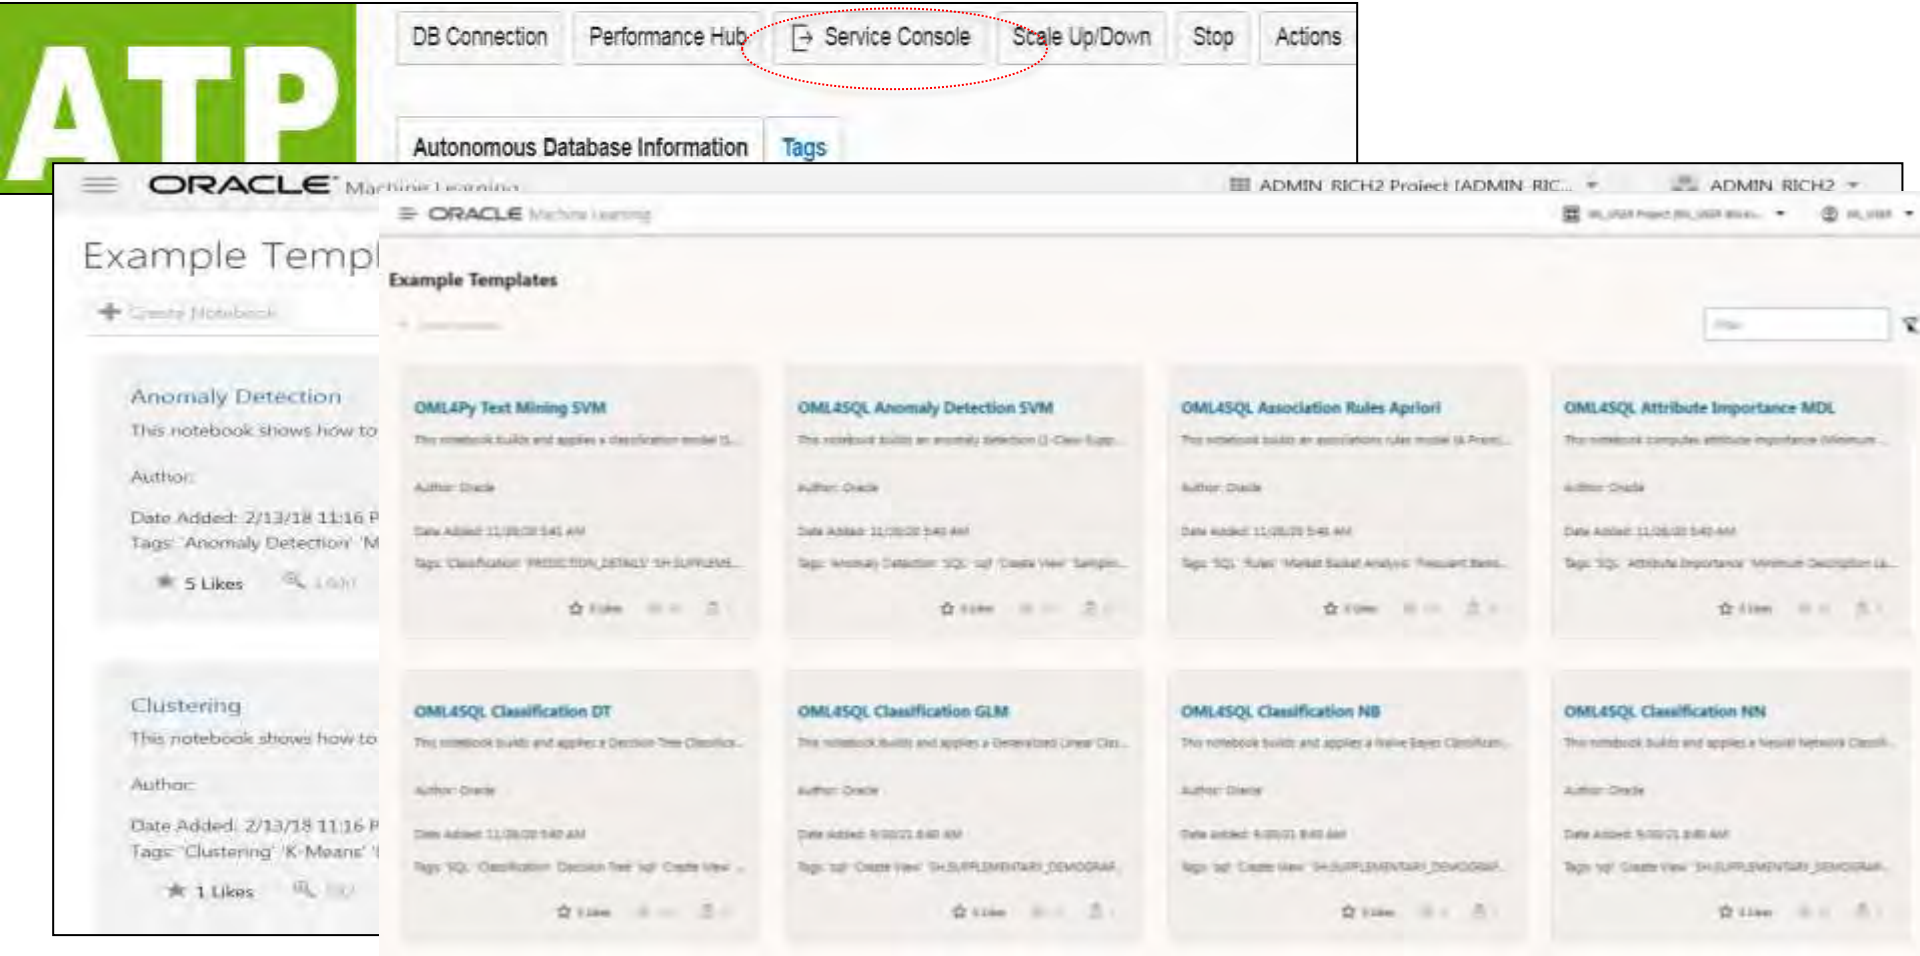

Machine Learning connection ADW/ATP - Python

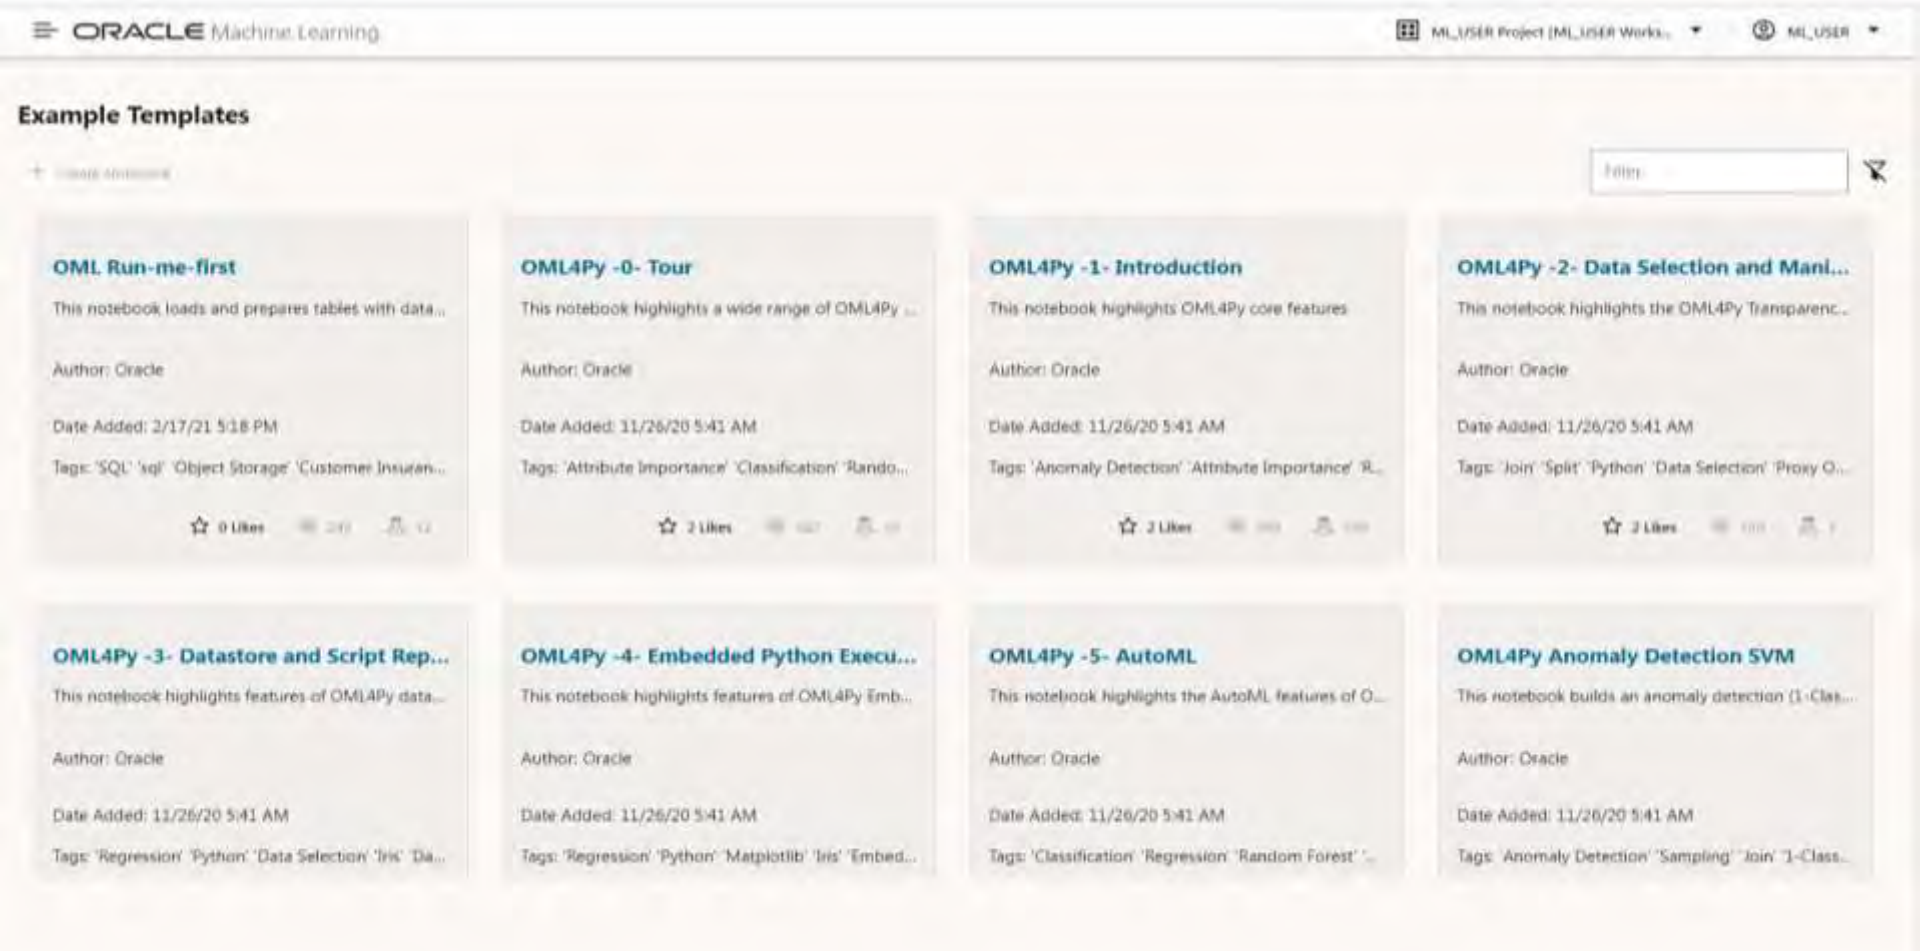

Machine Learning Functions* - Oracle Docs.

- **First**: Clear *Business Problem* to Solve
- **Second**: *Function* to Perform
- **Third**: *Algorithm* to use

- **First**: Build/*Train the Model* when you build it using about 60% of data.
- **Second**: *Test/Score Model* for accuracy/precision using about 40% of data.

Compare Algorithms!

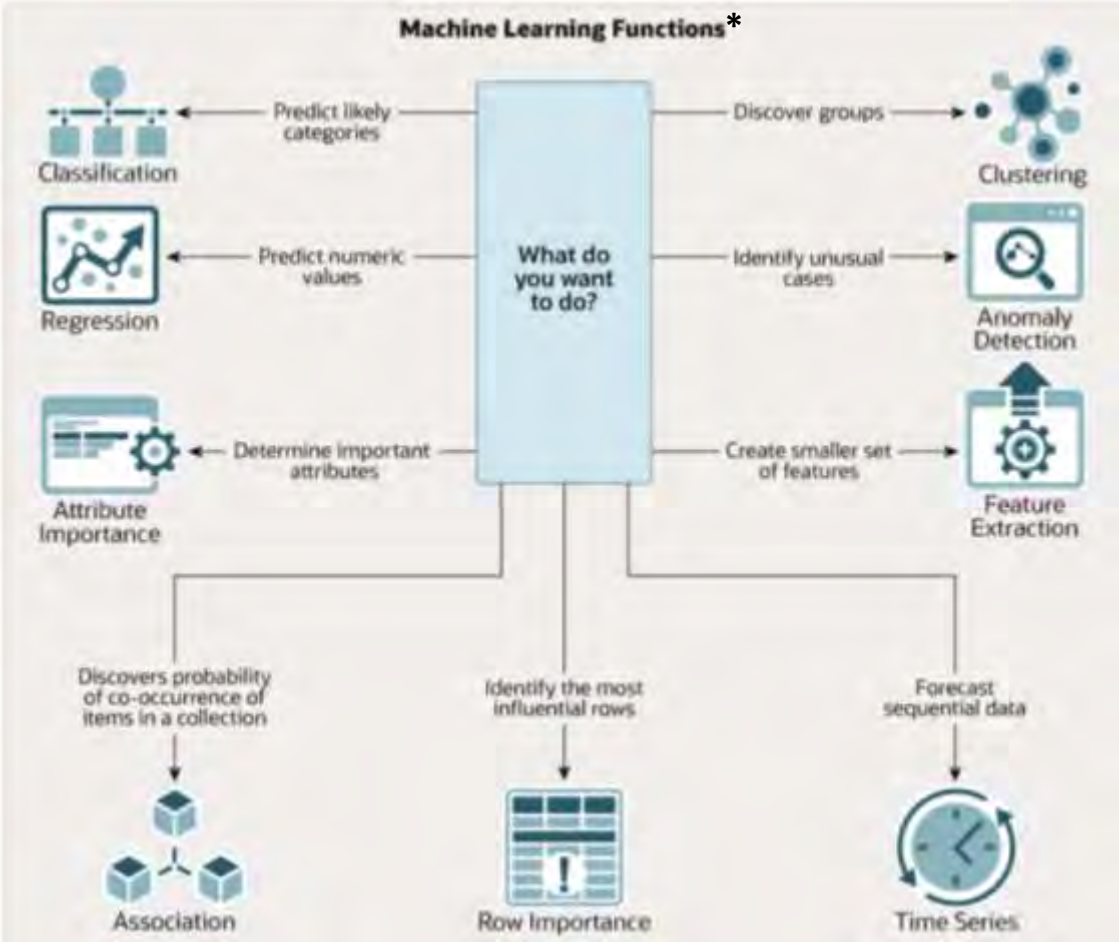

Business Understanding

Be Extremely Specific in Problem Statement:

Poorly Defined	Better	Data Mining Technique
Predict employees that leave	• Based on past employees that voluntarily left: • Create New Attribute EmplTurnover → O/1	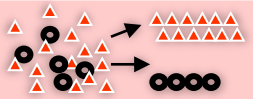
Predict customers that churn	• Based on past customers that left (churn): • Create New Attribute Churn → YES/NO	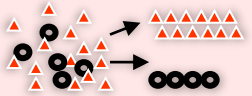
Target "best" customers	• Recency, Frequency Monetary (RFM) Analysis • Specific Dollar Amount over Time Window: • Who has spent $500+ in most recent 18 months	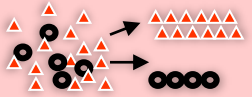
How can I make more $$?	• What helps me sell soft drinks & coffee?	
Which customers are likely to buy?	• How much is each customer likely to spend?	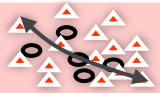
Who are my "best customers"?	• What descriptive "rules" describe "best customers"?	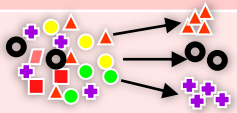
How can I combat fraud?	• Which transactions are the most anomalous? • Then roll-up to physician, claimant, employee, etc.	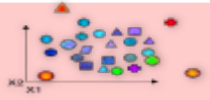

One-Class SVM (ML Anomaly Detection)*

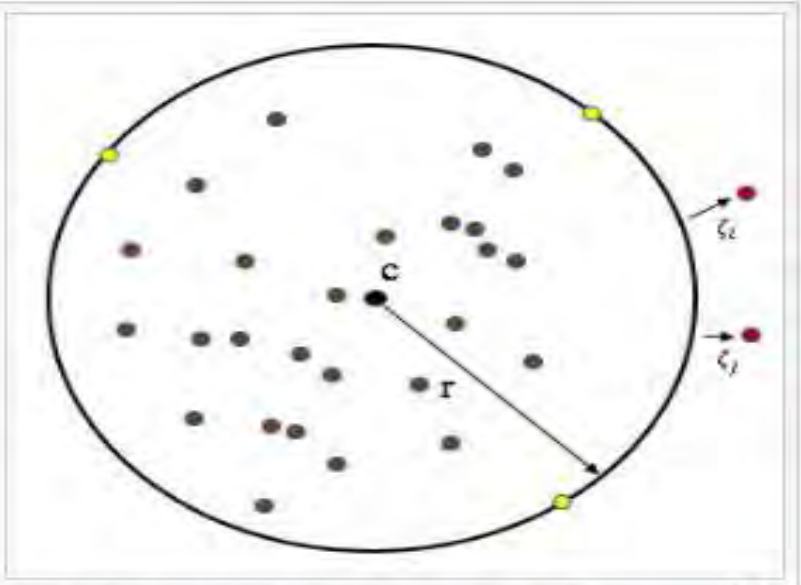

The hypersphere containing the target data having center a and radius R. Objects on the boundary are support vectors, and two objects lie outside the boundary having slack greater than 0.

Support Vector Data Description (SVDD): Find the smallest hypersphere containing all data points (use supervised training to get it)

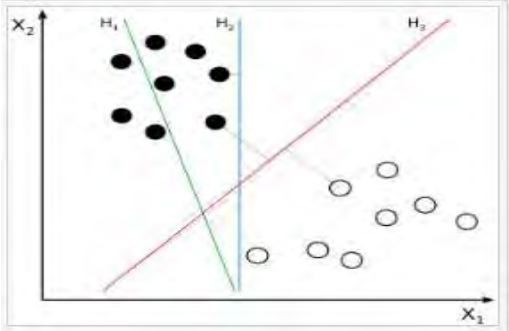

H_1 does not separate the classes. H_2 does, but only with a small margin. H_3 separates them with the maximal margin.

Linear SVM

*Wikipedia

One-Class SVM (ML Anomaly Detection) - FYI Only

- Support Vector Machine (**SVM**)
- One-Class Classification (**OCC**)
- Used to Classify Data
- **Identify specific objects of a class** based on supervised learning based on objects of that class.
- **Supervised Learning** using **Algorithms**
- SVM maps training data into separated areas
- **Unsupervised Learning** does data **clustering**
- They analyze the data used for **regression**
- SVM as binary non-probabilistic linear classifier
- **Classify images, text, & even hand written text**.

Machine Learning connection to ADW/ATP

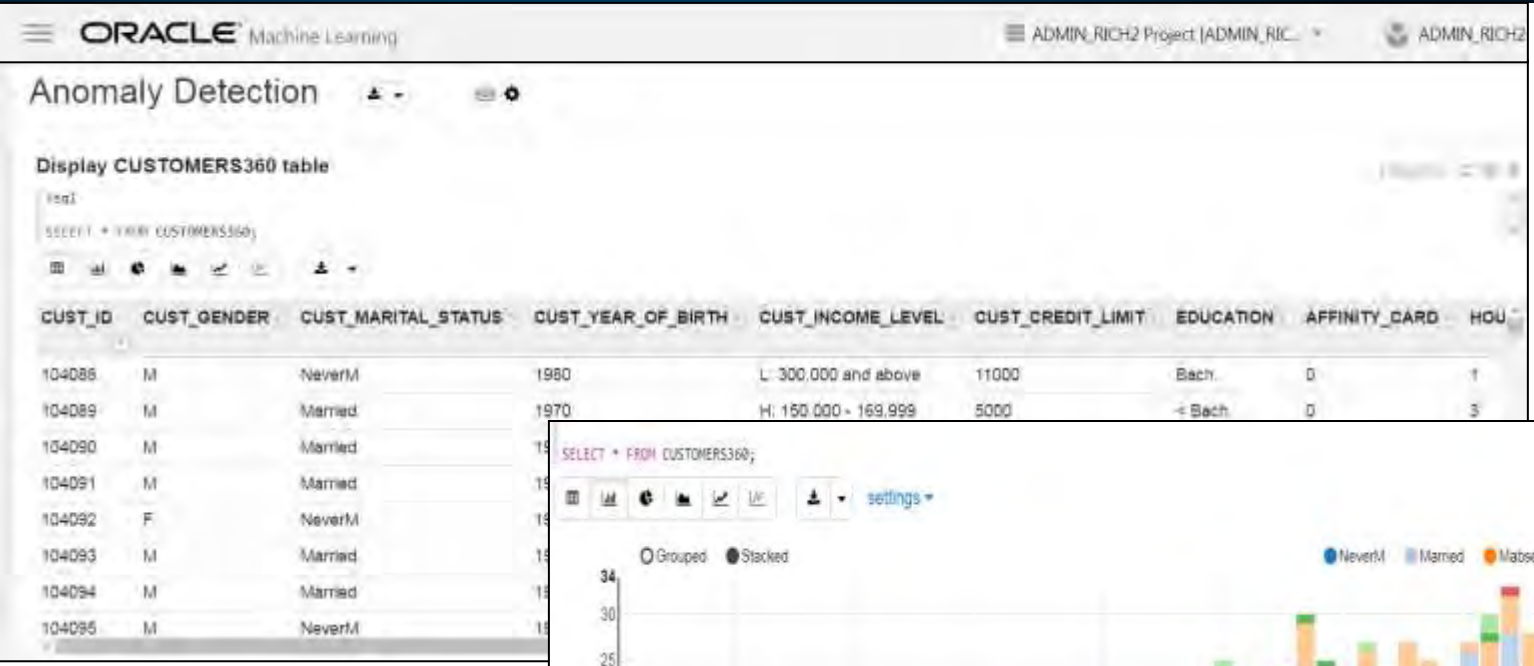

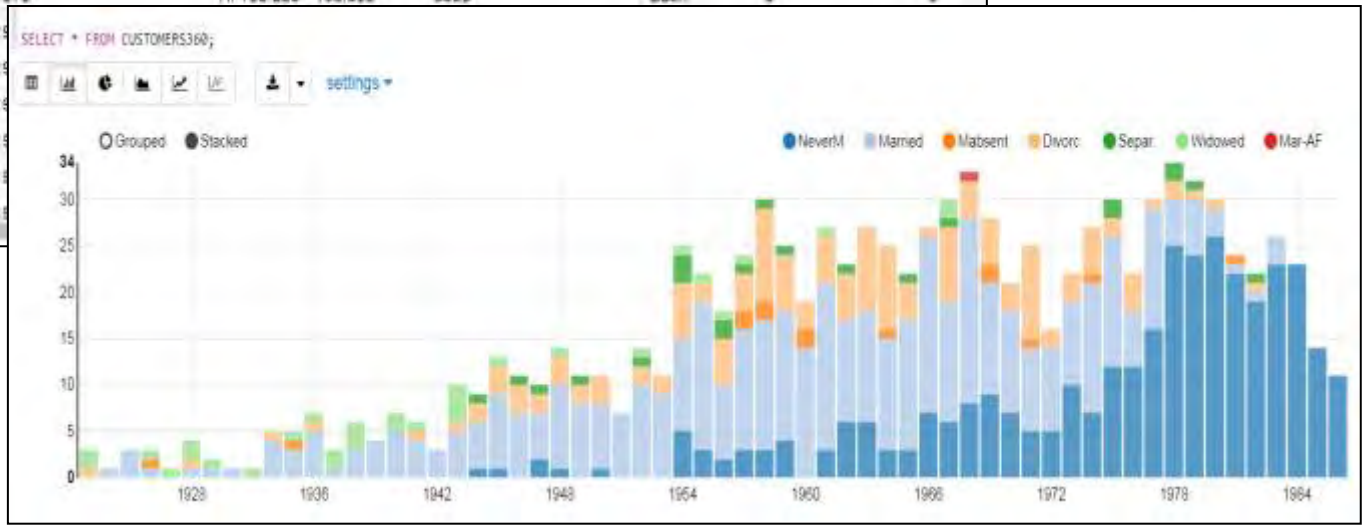

Machine Learning connection to ADW/ATP

Build Anomaly Detection model (1-Class Support Vector Machine)

```
%script
--Build Anomaly Detection Model (1-Class SVM) on CUSTOMERS360 data

DECLARE
v_sql varchar2(100);

BEGIN

--Create a Build Setting table for Model Build

  EXECUTE IMMEDIATE 'CREATE TABLE CUSTOMERS360_SET (setting_name VARCHAR2(30),setting_value VARCHAR2(4000))';
  EXECUTE IMMEDIATE 'INSERT INTO CUSTOMERS360_SET (setting_name, setting_value) VALUES (''ALGO_NAME'', ''ALGO_SUPPORT_VECTOR_MACHINES'')';
  EXECUTE IMMEDIATE 'INSERT INTO CUSTOMERS360_SET (setting_name, setting_value) VALUES (''PREP_AUTO'', ''ON'')';
  DBMS_OUTPUT.PUT_LINE ('Created model build settings table: CUSTOMERS360_SET ');

--Build the 1-Class SVM model.

  EXECUTE IMMEDIATE 'CALL DBMS_DATA_MINING.CREATE_MODEL(''CUSTOMERS360MODEL'', ''CLASSIFICATION'', ''CUSTOMERS360'', ''CUST_ID    '', null, ''CUSTOMERS360_SET'')';
  DBMS_OUTPUT.PUT_LINE ('Created model: CUSTOMERS360_MODEL ');

END;
```

```
Created model build settings table: CUSTOMERS360_SET
Created model: CUSTOMERS360_MODEL
PL/SQL procedure successfully completed.
```

Graph all Customers's PROBABILITY_ANOMALOUS vs. YRS_RESIDENCE grouped by CUST_MARITAL_STATUS

```
%sql

-- Click on the bar chart icon and expand settings.   Drag PROBABILITY_ANOMALOUS
-- HINT: Select items in the legend to focus on certain categories.

SELECT *
FROM (SELECT CUST_ID, round(probability_anomalous,2) probability_anomalous, HOUSEHOLD_SIZE, YRS_RESIDENCE, CUST_GENDER, CUST_MARITAL_STATUS,
      rank() over (ORDER BY probability_anomalous DESC) rnk
      FROM (SELECT CUST_ID, HOUSEHOLD_SIZE, YRS_RESIDENCE, CUST_GENDER, CUST_MARITAL_STATUS, prediction_probability(CUSTOMERS360MODEL, '0' using *) probability_anomalous
            FROM CUSTOMERS360))
ORDER BY probability_anomalous DESC;
```

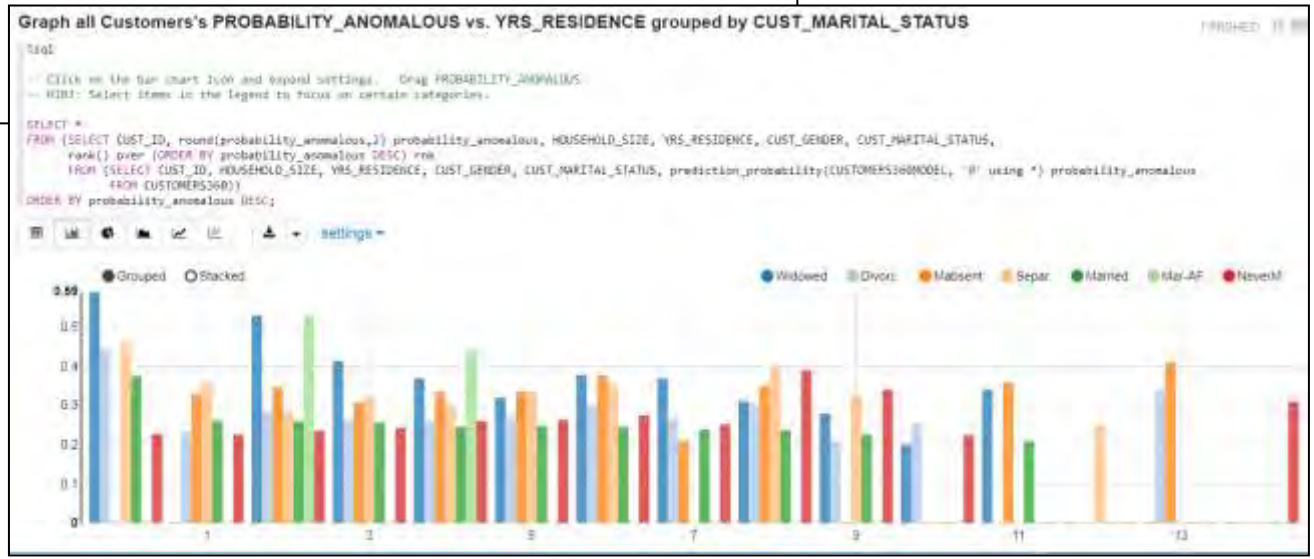

Machine Learning connection to ADW/ATP

View Prediction_Details that explain why the record was selected as anomalous

```
%sql
-- Select customers with OCCUPATION of 'TechSup' and more than 55% probability of being anomalous

SELECT CUST_ID,
       RTRIM(TRIM(SUBSTR(OUTPRED."Attribute1",17,100)),'rank="1"/>') FIRST_ATTRIBUTE,
       RTRIM(TRIM(SUBSTR(OUTPRED."Attribute2",17,100)),'rank="2"/>') SECOND_ATTRIBUTE,
       RTRIM(TRIM(SUBSTR(OUTPRED."Attribute3",17,100)),'rank="3"/>') THIRD_ATTRIBUTE,
       RTRIM(TRIM(SUBSTR(OUTPRED."Attribute4",17,100)),'rank="4"/>') FOURTH_ATTRIBUTE,
       RTRIM(TRIM(SUBSTR(OUTPRED."Attribute5",17,100)),'rank="5"/>') FIFTH_ATTRIBUTE
       FROM (SELECT CUST_ID, PREDICTION_DETAILS(CUSTOMERS360MODEL, '0' USING *) PREDICTION_DETAILS FROM CUSTOMERS360
           WHERE PREDICTION_PROBABILITY(CUSTOMERS360MODEL, '0' USING *) > 0.50
           AND OCCUPATION = 'TechSup'
           ORDER BY CUST_ID) OUT,
           XMLTABLE('/Details'
                PASSING OUT.PREDICTION_DETAILS
                COLUMNS
                   "Attribute1" XMLType PATH 'Attribute[1]',
                   "Attribute2" XMLType PATH 'Attribute[2]',
                   "Attribute3" XMLType PATH 'Attribute[3]',
                   "Attribute4" XMLType PATH 'Attribute[4]',
                   "Attribute5" XMLType PATH 'Attribute[5]')
       OUTPRED;
```

CUST_ID	FIRST_ATTRIBUTE	SECOND_ATTRIBUTE	THIRD_ATTRIBUTE
100646	"CUS... weig		
102922	"CUS		

CUST_ID	FIRST_ATTRIBUTE	SECOND_ATTRIBUTE	THIRD_ATTRIBUTE
100646	"CUST_MARITAL_STATUS" actualValue="Widowed" weight=".226"	"CUST_YEAR_OF_BIRTH" actualValue="1941" weight=".118"	"CUST_CREDIT_LIMIT" actualValue="1500" weight=".
102922	"CUST_MARITAL_STATUS" actualValue="Widowed" weight=".222"	"CUST_YEAR_OF_BIRTH" actualValue="1931" weight=".169"	"CUST_CREDIT_LIMIT" actualValue="1500" weight=".
103441	"CUST_MARITAL_STATUS" actualValue="Widowed" weight=".222"	"CUST_YEAR_OF_BIRTH" actualValue="1941" weight=".117"	"EDUCATION" actualValue="Bach." weight=".076"
104286	"EDUCATION" actualValue="9th" weight=".165"	"HOUSEHOLD_SIZE" actualValue="4-5" weight=".146"	"CUST_CREDIT_LIMIT" actualValue="1500" weight=".

Decision Tree Algorithm (ML Classifier)

A decision tree consists of three types of nodes:[1]

1. Decision nodes – typically represented by squares
2. Chance nodes – typically represented by circles
3. End nodes – typically represented by triangles

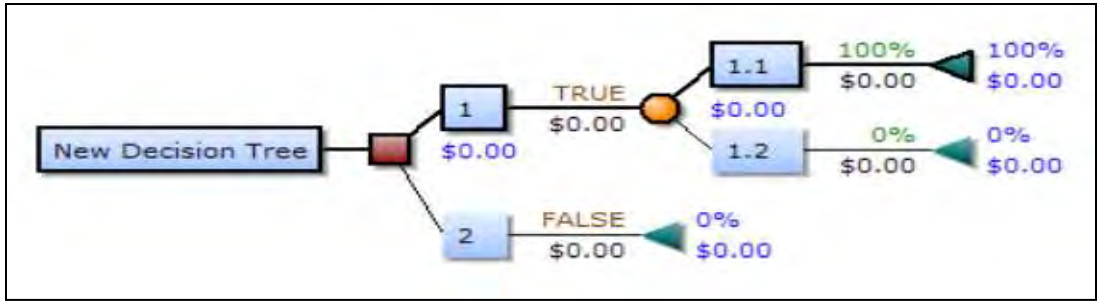

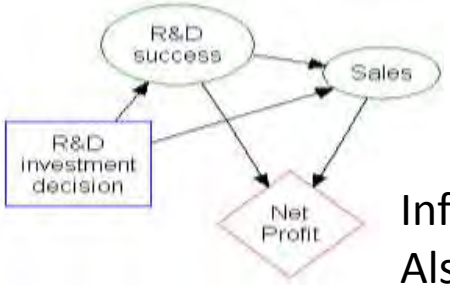

Influence Diagram is Also a Decision Tree

Flow Chart calculates whether to settle a case or not based on costs/probabilities

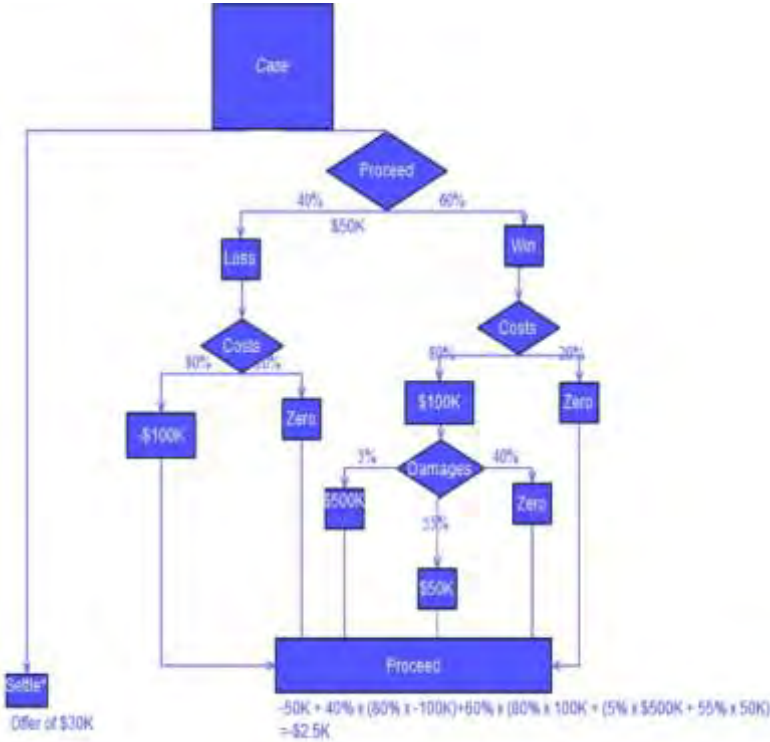

OML (OAA) Oracle Data Mining SQL Sample (PARTIAL)

dmdtdemo.sql – DBMS_DATA_MINING package – Decision Tree

```
-- Given demographic data about a set of customers, predict the
-- customer response to an affinity card program using a classifier
-- based on Decision Trees algorithm.

 INSERT INTO dt_sh_sample_settings VALUES
   (dbms_data_mining.algo_name, dbms_data_mining.algo_decision_tree);
--
-- CREATE A NEW MODEL
-- Build a DT model
```

```
BEGIN
  DBMS_DATA_MINING.CREATE_MODEL(
    model_name          => 'DT_SH_Clas_sample',
    mining_function     => dbms_data_mining.classification,
    data_table_name     => 'mining_data_build_v',
    case_id_column_name => 'cust_id',
    target_column_name  => 'affinity_card',
    settings_table_name => 'dt_sh_sample_settings');
END;
/
```


```
SELECT T.cust_id, S.prediction, S.probability, S.cost
  FROM (SELECT cust_id,
               PREDICTION_SET(dt_sh_clas_sample COST MODEL USING *) pset
               FROM    mining_data_apply_v
               WHERE cust_id < 100011) T,
            TABLE(T.pset) S
ORDER BY cust_id, S.prediction;
```

Script Output

Task completed in 3.075 seconds

CUST_ID	PREDICTION	PROBABILITY	COST
100001	0	.966183575	.270531401
100001	1	.033816425	.966183575
100002	0	.740384615	2.076923077
100002	1	.259615385	.740384615
100003	0	.909090909	.727272727
100003	1	.090909091	.909090909
100004	0	.909090909	.727272727
100004	1	.090909091	.909090909
100005	0	.272357724	5.821138211
100005	1	.727642276	.272357724
100006	0	1.000000000	.000000000
100006	1	.000000000	1.000000000
100007	0	.909090909	.727272727
100007	1	.090909091	.909090909
100008	0	.909090909	.727272727
100008	1	.090909091	.909090909
100009	0	.272357724	5.821138211
100009	1	.727642276	.272357724
100010	0	.675965665	2.592274678
100010	1	.324034335	.675965665

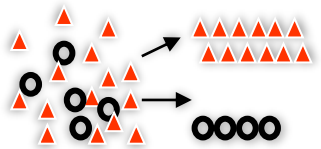

Seasonal, Irregular & Missing Data: Time Series Algorithm

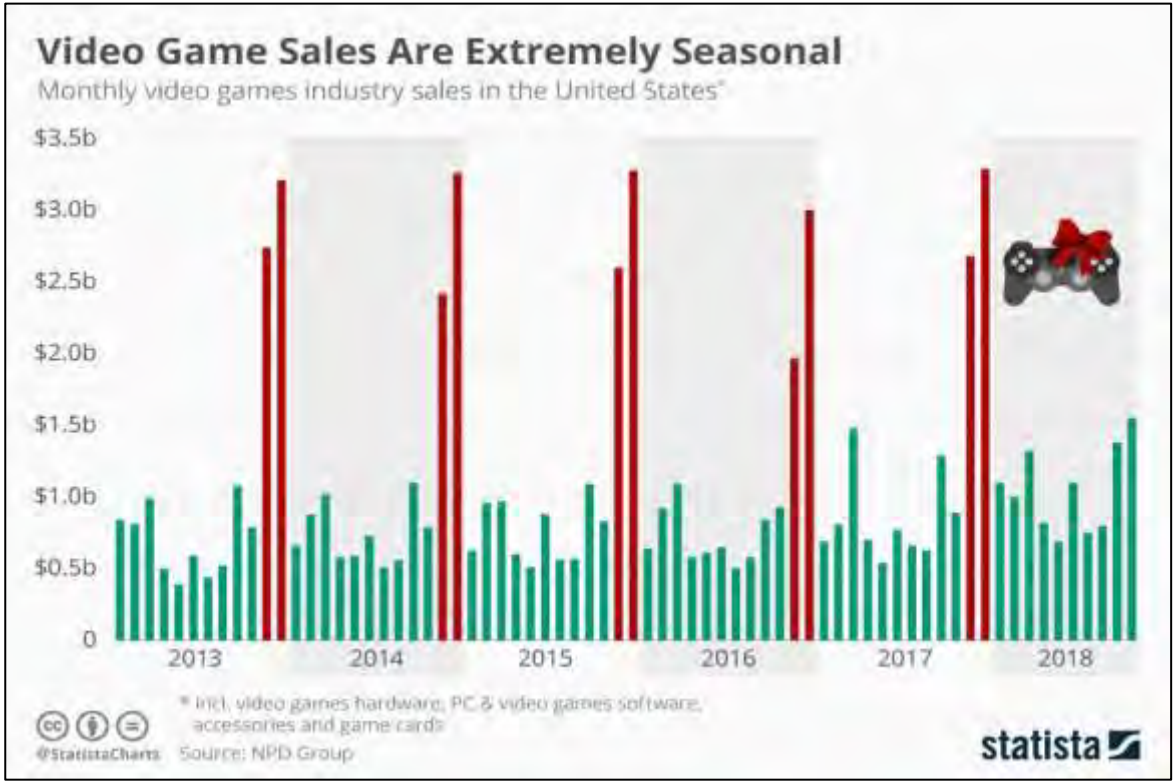

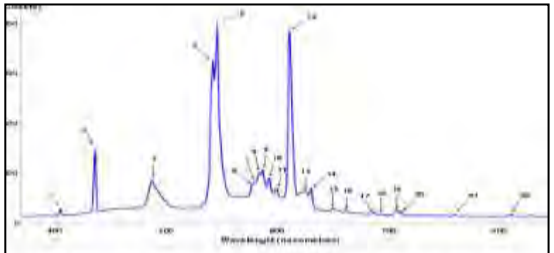

The power spectrum $S_{xx}(f)$ of a time series $x(t)$ describes the distribution of power into frequency components composing that signal. According to Fourier analysis, any physical signal can be decomposed into a number of discrete frequencies, or a spectrum of frequencies over a continuous range.

Exponential Smoothing Algorithm (Time Series)

Stock Application (Left)
A **Time Series** (Below)

Use **Exponential Smoothing** to Weight older data to have less effect than newer data

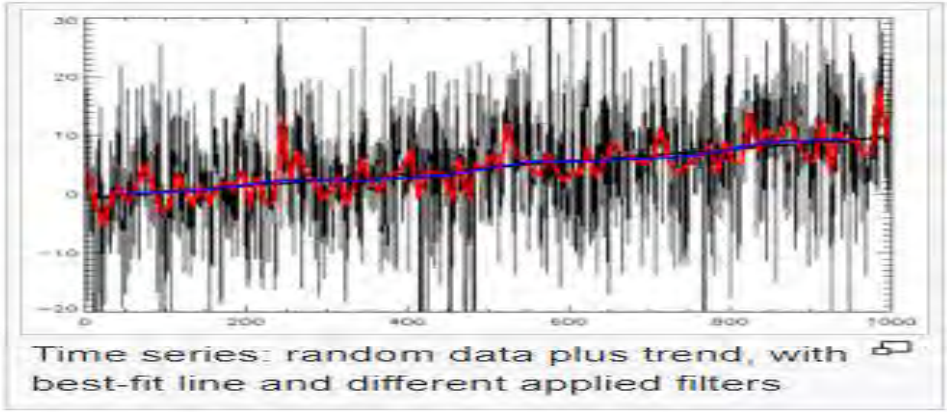

Machine Learning & AI - Oracle's Built-In Algorithms

Oracle's Machine Learning & Adv. Analytics Algorithms

Classification

- Naïve Bayes
- Logistic Regression (GLM)
- Decision Tree
- Random Forest
- Neural Network
- Support Vector Machine
- Explicit Semantic Analysis

Clustering

- Hierarchical K-Means
- Hierarchical O-Cluster
- Expectation Maximization (EM)

Anomaly Detection

- One-Class SVM

Time Series

- State of the art forecasting using Exponential Smoothing.
- Includes all popular models e.g. Holt-Winters with trends, seasons, irregularity, missing data

Regression

- Linear Model
- Generalized Linear Model
- Support Vector Machine (SVM)
- Stepwise Linear regression
- Neural Network
- LASSO *

Attribute Importance
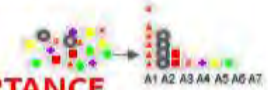

- Minimum Description Length
- Principal Comp Analysis (PCA)
- Unsupervised Pair-wise KL Div
- CUR decomposition for row & AI

Association Rules

- A priori/ market basket

Predictive Queries
- Predict, cluster, detect, features

SQL Analytics


- SQL Windows, SQL Patterns, SQL Aggregates

Feature Extraction
- Principal Comp Analysis (PCA)
- Non-negative Matrix Factorization
- Singular Value Decomposition (SVD)
- Explicit Semantic Analysis (ESA)

Text Mining Support

- Algorithms support text type
- Tokenization and theme extraction
- Explicit Semantic Analysis (ESA) for document similarity

Statistical Functions

- Basic statistics: min, max, median, stdev, t-test, F-test, Pearson's, Chi-Sq, ANOVA, etc.

R Packages


- CRAN R Algorithm Packages through Embedded R Execution
- Spark MLlib algorithm integration

Exportable ML Models
- REST APIs for deployment

ORACLE

- OAA (Oracle Data Mining + Oracle R Enterprise) and ORAAH combined
- OAA includes support for Partitioned Models, Transactional, Unstructured, Geo-spatial, Graph data, etc.
- OAA is Oracle Advanced Analytics; ORAAH is Oracle R Advanced Analytics for Hadoop

Types of ML Algorithms in Healthcare on PubMed*

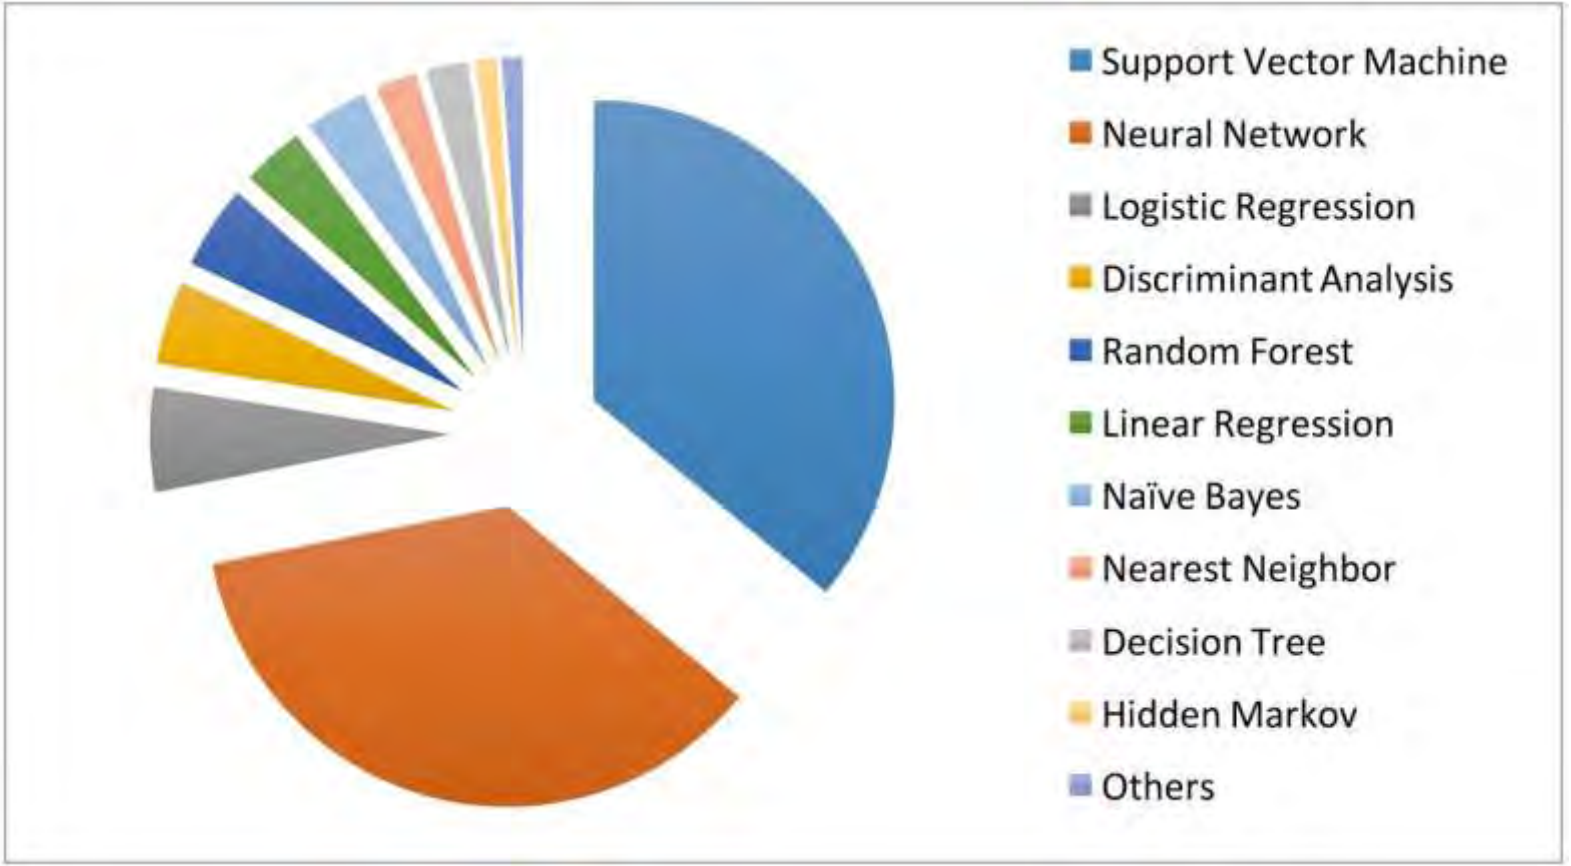

***Stroke and vascular neurology (snj.bmj.com)**

Note: Markov is Bayesian, NN is k-NN, DA close to PCA

AutoML for Autonomous Database- 3/18/2021

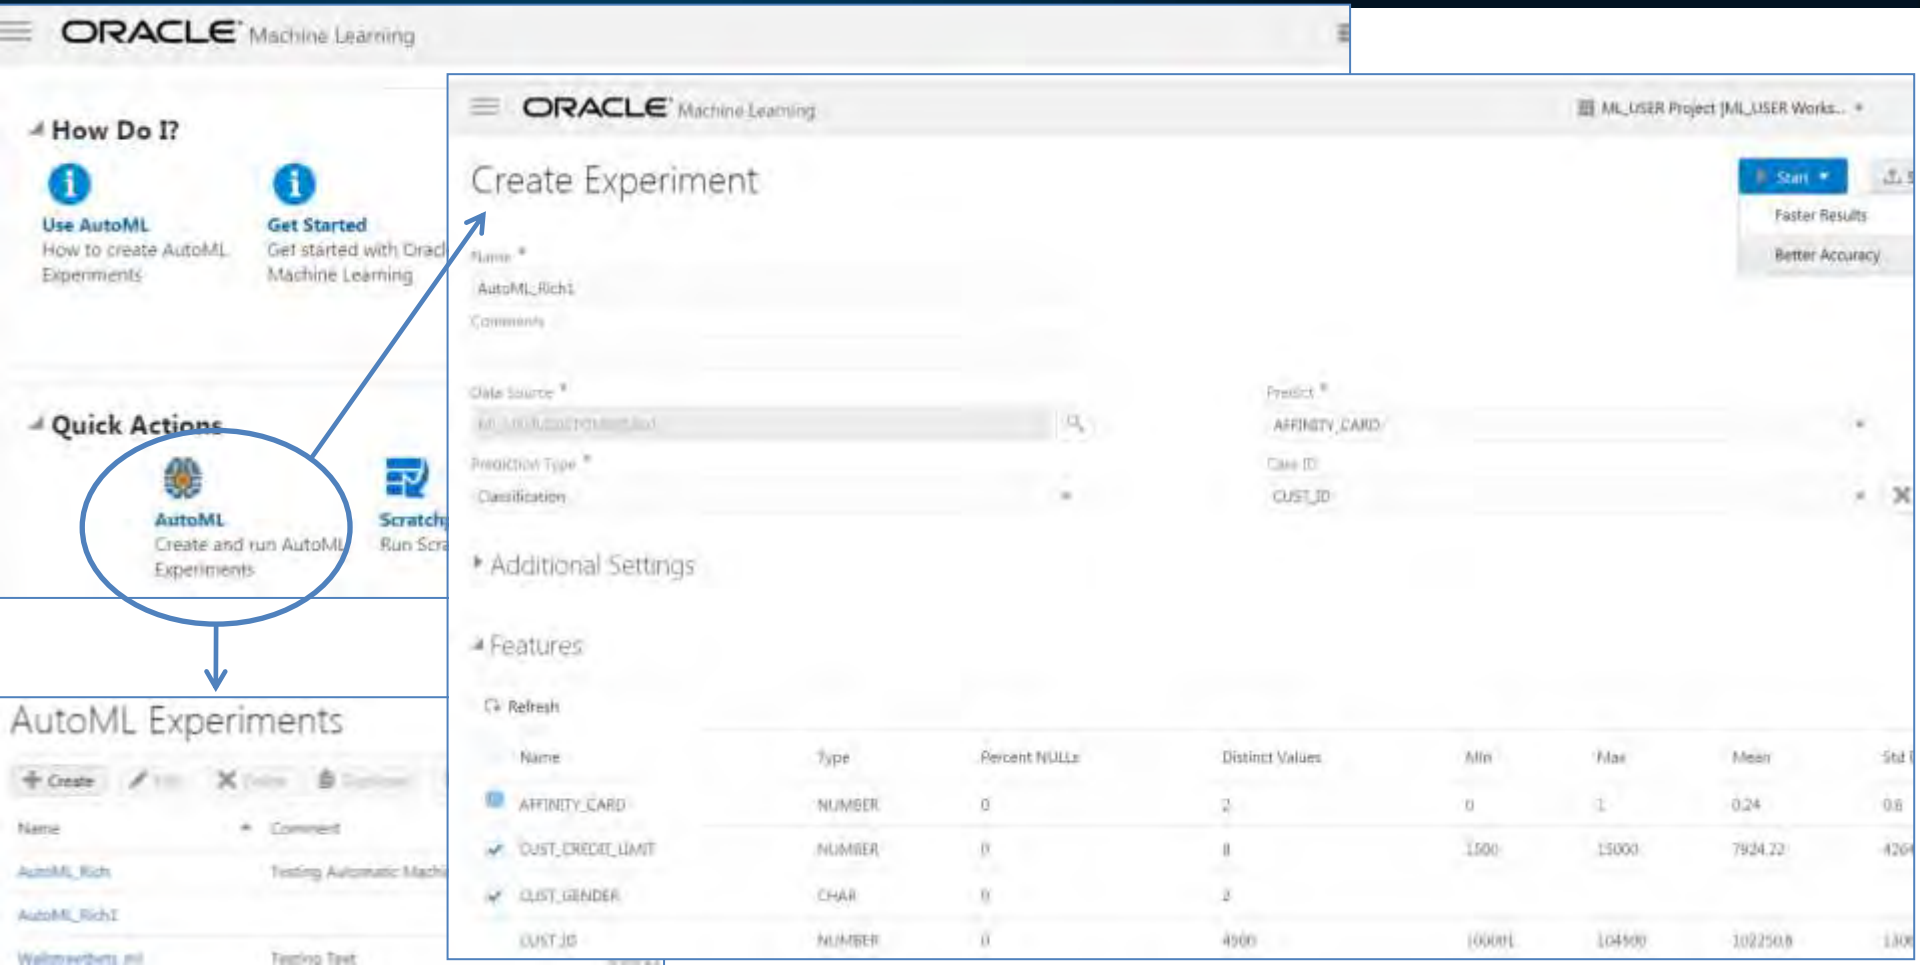

AutoML is here for Autonomous Database

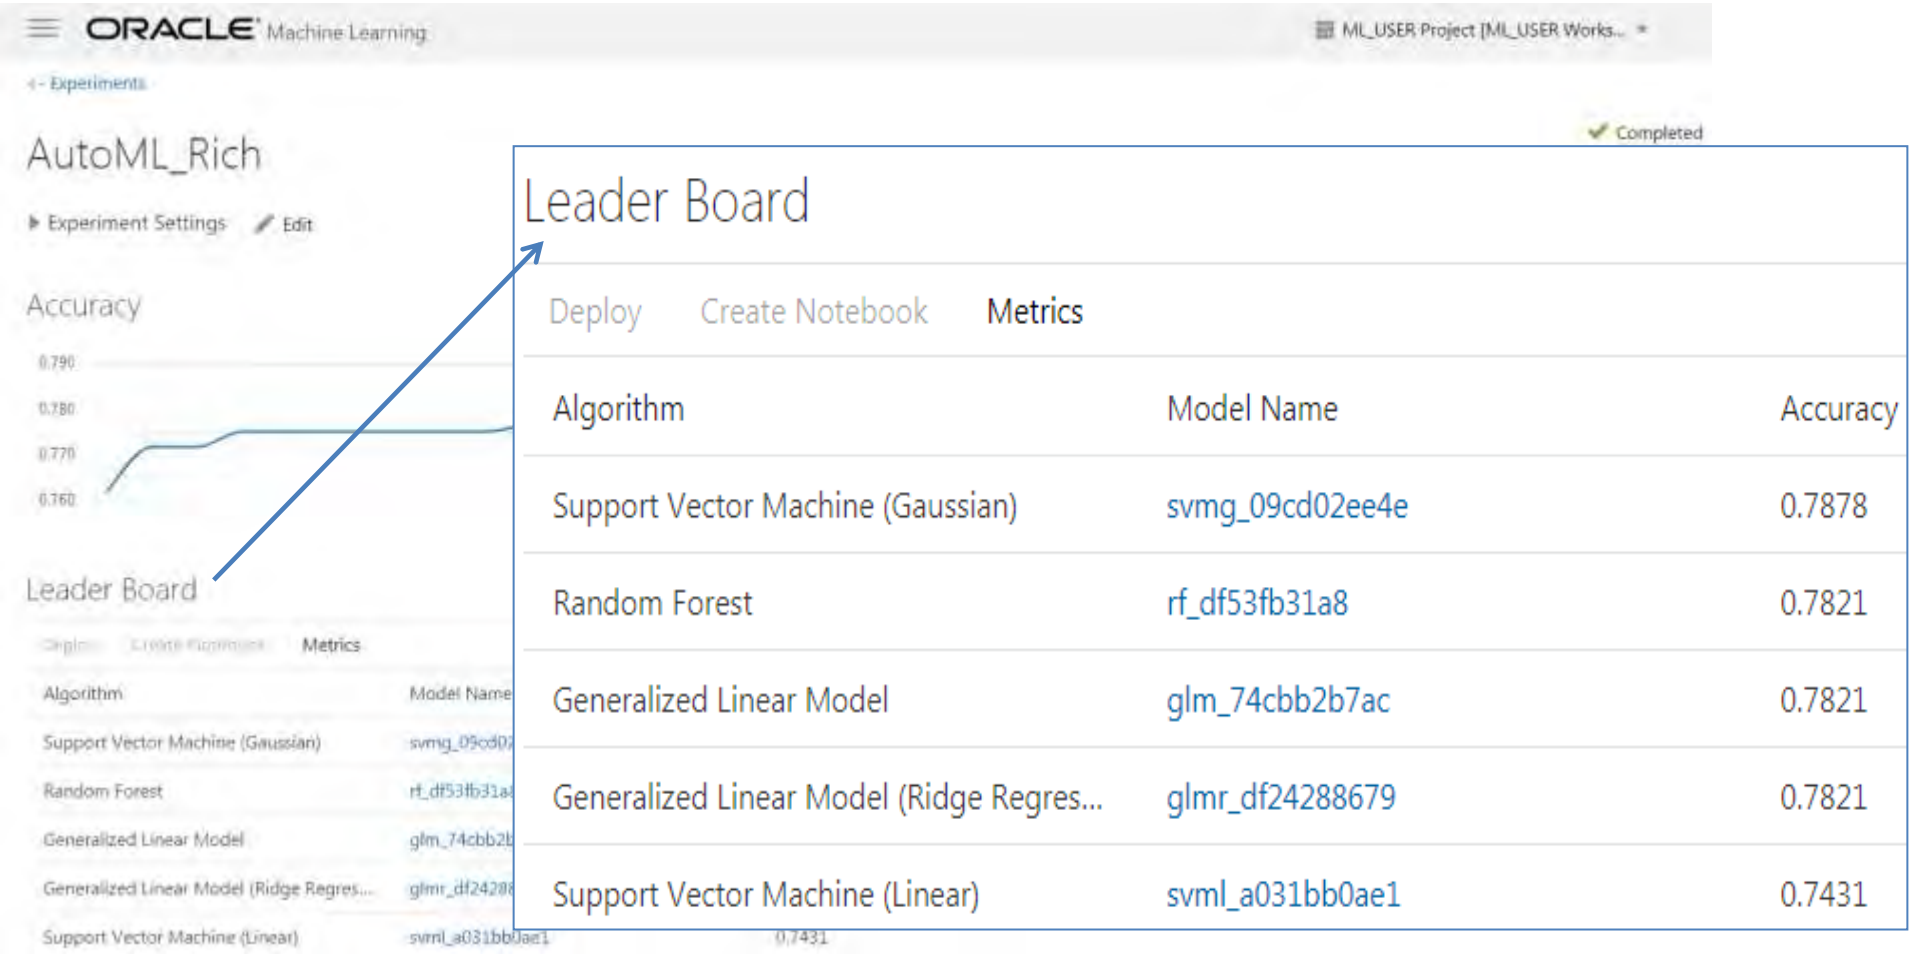

AutoML is here for Autonomous Database

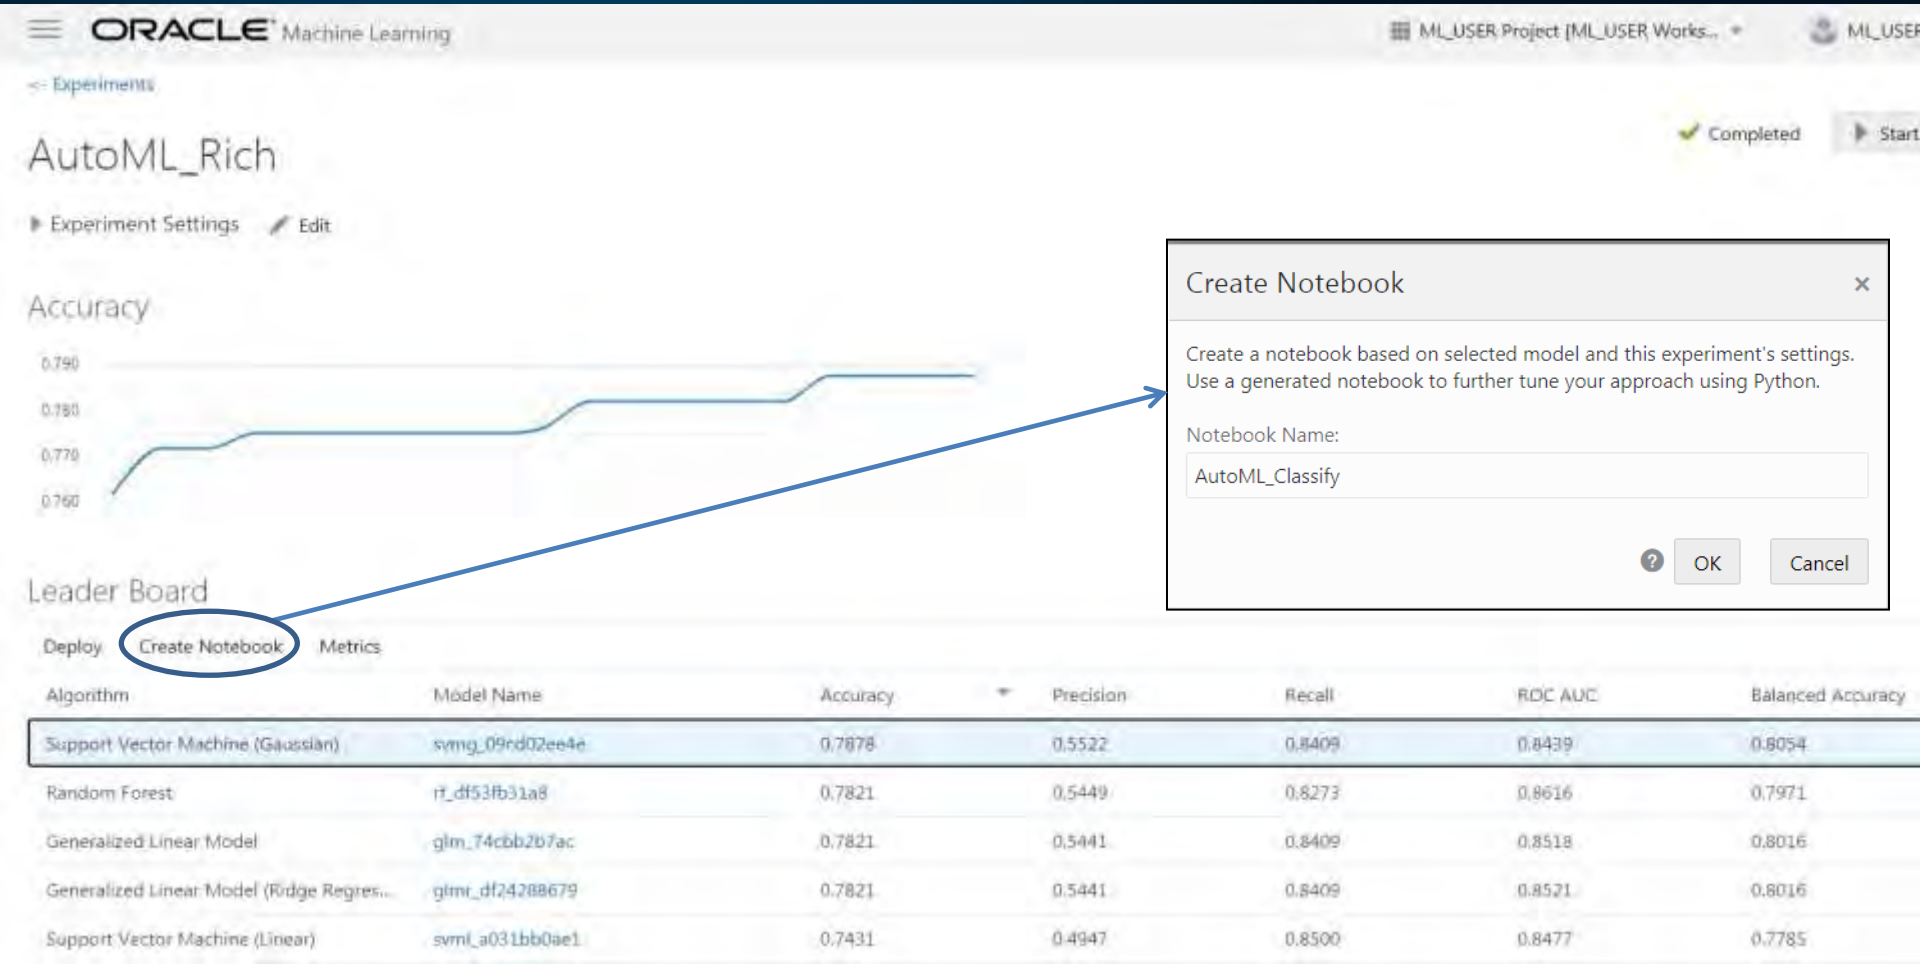

ORACLE Machine Learning

AutoML_Classify

Build Data

```
%python

import oml

columns = 'CUST_ID' , 'CUST_CREDIT_LIMIT' , 'CUST_GENDER' , 'CUST_INCOME_LEVEL' , 'CUST_MARITAL_STATUS' , 'CUST_YEAR_OF_BIRTH' , 'EDUCATION' , 'HOUSEHOLD_SIZE' , 'OCCUPATION' , 'YRS_RESIDENCE' ,
    'Y_BOX_GAMES' , 'AFFINITY_CARD'
schema='HL_USER'
table='CUSTOMERS300'

column = ','.join(columns)
query = 'SELECT ' + column + ' FROM ' + schema + '.' + table

data_build = oml.sync(query=query)
print('data_build')
```

Create Train Data

READY

```
%python

import oml

X_train = data_build[:,['CUST_ID', 'CUST_CREDIT_LIMIT', 'CUST_GENDER', 'CUST_INCOME_LEVEL', 'CUST_MARITAL_STATUS', 'CUST_YEAR_OF_BIRTH', 'EDUCATION', 'HOUSEHOLD_SIZE', 'OCCUPATION', 'YRS_RESIDENCE',
    'Y_BOX_GAMES']]
y_train = data_build[:,'AFFINITY_CARD']
```

Build 'SUPPORT_VECTOR_MACHINES' Model

READY

```
%python

import oml

svm_settings = {
  'ODMS_SAMPLE_SIZE' : '1294' , 'SVMS_COMPLEXITY_FACTOR' : '66.70000000000002' , 'SVMS_KERNEL_FUNCTION' : 'SVMS_GAUSSIAN' , 'SVMS_STD_DEV' : '2.2360679774997' , 'CLAS_WEIGHTS_BALANCED' : 'OFF' ,
    'SVMS_NUM_PIVOTS' : '200' , 'ODMS_DETAILS' : 'ODMS_DISABLE' , 'ODMS_SAMPLING' : 'ODMS_SAMPLING_ENABLE' , 'ALGO_NAME' : 'ALGO_SUPPORT_VECTOR_MACHINES'
  }

svm_mod = oml.svm(**svm_settings)

svm_mod = svm_mod.fit(X_train, y_train)
```

Machine Learning & AI - Oracle's Algorithms

21c

Oracle Machine Learning Algorithms

Classification
Naïve Bayes
Logistic Regression (GLM)
Decision Tree
Random Forest
Neural Network
Support Vector Machine
Explicit Semantic Analysis
XGBoost*

Clustering
Hierarchical K-Means
Hierarchical O-Cluster
Expectation Maximization (EM)

Anomaly Detection
One-Class SVM
MSET-SPRT*

Time Series
Forecasting - Exponential Smoothing
Includes popular models
e.g. Holt-Winters with trends, seasonality, irregularity, missing data

Regression
Linear Model
Generalized Linear Model
Support Vector Machine (SVM)
Stepwise Linear regression
Neural Network
XGBoost*

Attribute Importance
Minimum Description Length
Principal Comp Analysis (PCA)
Unsupervised Pair-wise KL Div
CUR decomposition for row & Al

Association Rules
A priori/ market basket

Predictive Queries
Predict, cluster, detect, features

SQL Analytics
SQL Windows
SQL Patterns
SQL Aggregates

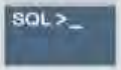

Feature Extraction
Principal Comp Analysis (PCA)
Non-negative Matrix Factorization
Singular Value Decomposition (SVD)
Explicit Semantic Analysis (ESA)

Text Mining Support
Algorithms support text
Tokenization and theme extraction
Explicit Semantic Analysis (ESA) for document similarity

Statistical Functions
Basic statistics: min, max, median, stdev, t-test, F-test, Pearson's, Chi-Sq, ANOVA, etc.

R & Python *Coming soon*
Third-party R & Python Packages through Embedded Execution
Spark MLlib algorithm integration

Model Deployment & Monitoring
SQL—1st Class Objects
Oracle RESTful API (ORDS)
OML Web Services (for Apps)

Includes support for Partitioned Models, Transactional data and aggregations, Unstructured data, Geo-spatial data, Graph data, etc,

* New in 21c

• ORAAH is Oracle R Advanced Analytics for Hadoop

Machine Learning has many parts

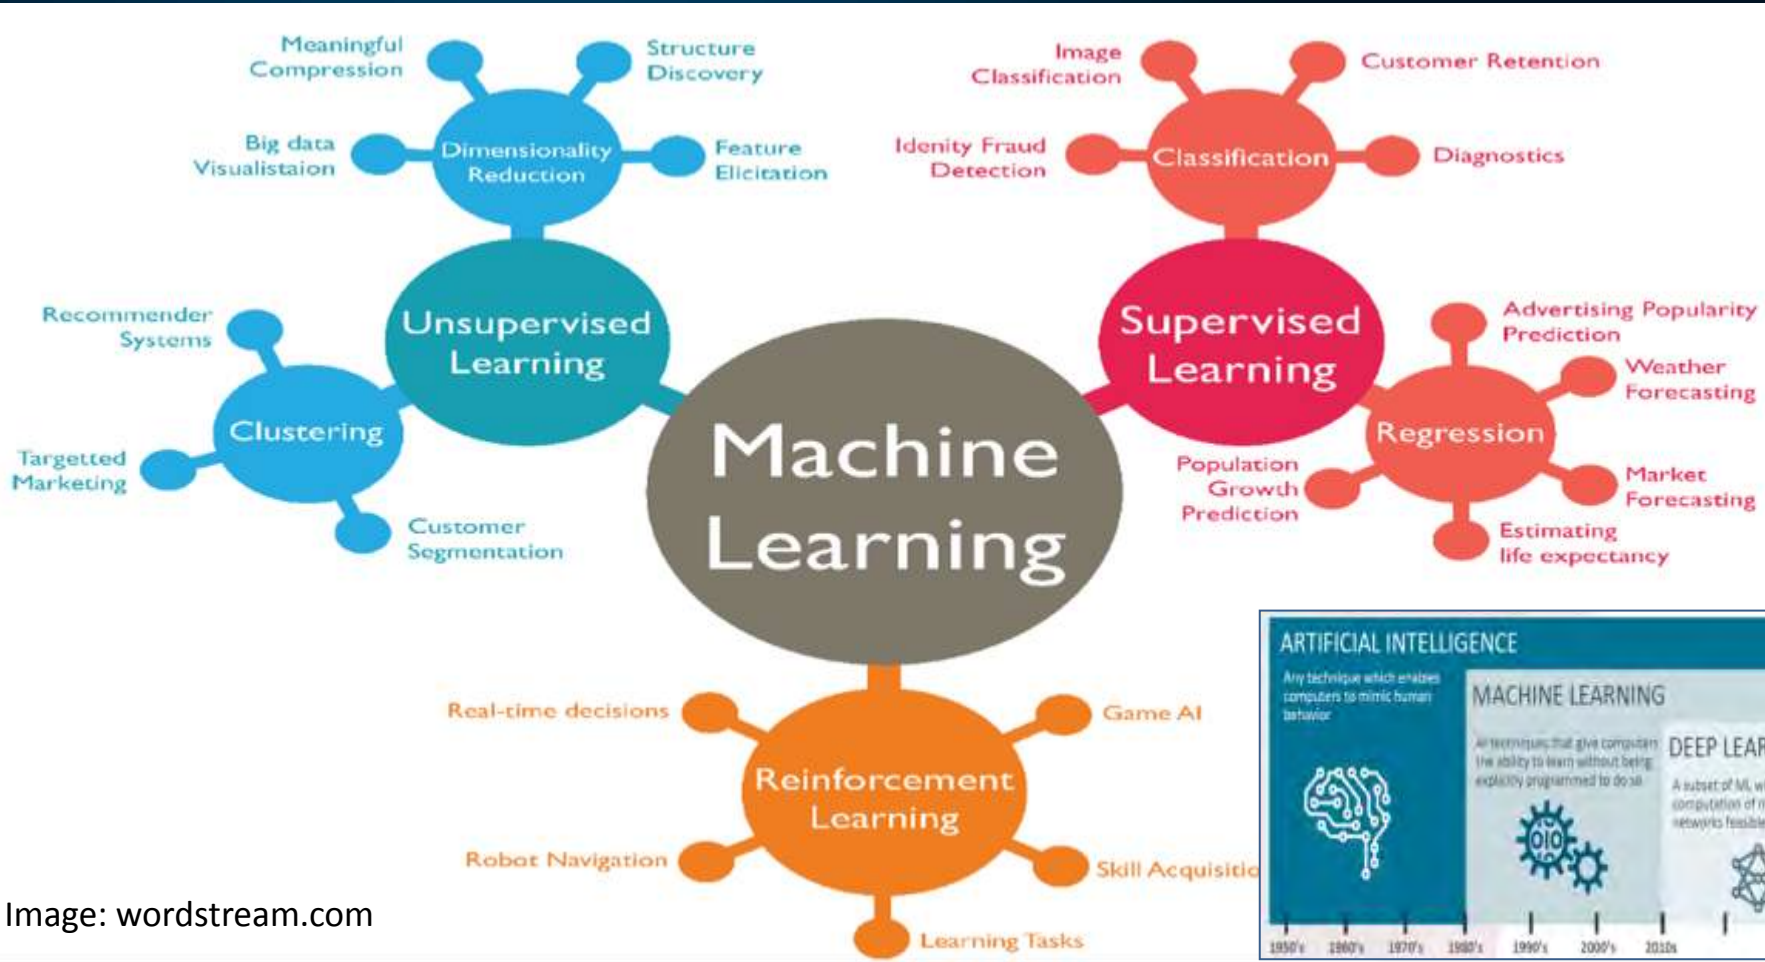

Image: wordstream.com

Exadata Cloud Machine: ALL Features

(Oracle's Juan Loaiza presentation on Exadata Cloud)

Exadata Cloud: Compatible, Scalable, Available, Secure

Decades of Database Innovation Proven at Millions of Mission-Critical Deployments

	Multitenant
	In-Memory DB
	Real Application Clusters
	Active Data Guard
	Partitioning
	Advanced Compression
	Advanced Security, Label Security, DB Vault
	Real Application Testing
	Advanced Analytics, Spatial and Graph
	Management Packs for Oracle Database

All Oracle Database Innovations

All Exadata DB Machine Innovations

Offload SQL to Storage	
InfiniBand Fabric	
Smart Flash Cache, Log	PCI Flash
Storage Indexes	
Columnar Flash Cache	
Hybrid Columnar Compression	10:1
I/O Resource Management	
Network Resource Management	
In-Memory Fault Tolerance	
Exafusion Direct-to-Wire Protocol	

George Shows How to Mine ADW with Oracle Analytics

Many Graph Choices

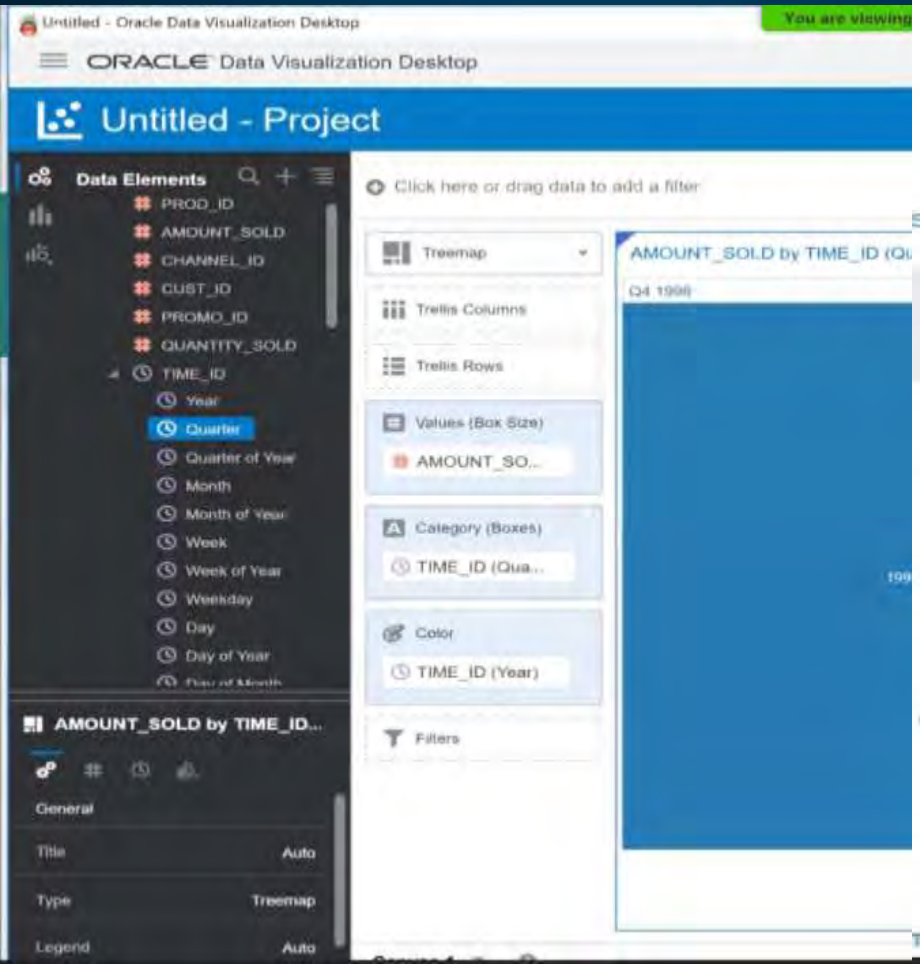

More Complex Treemap

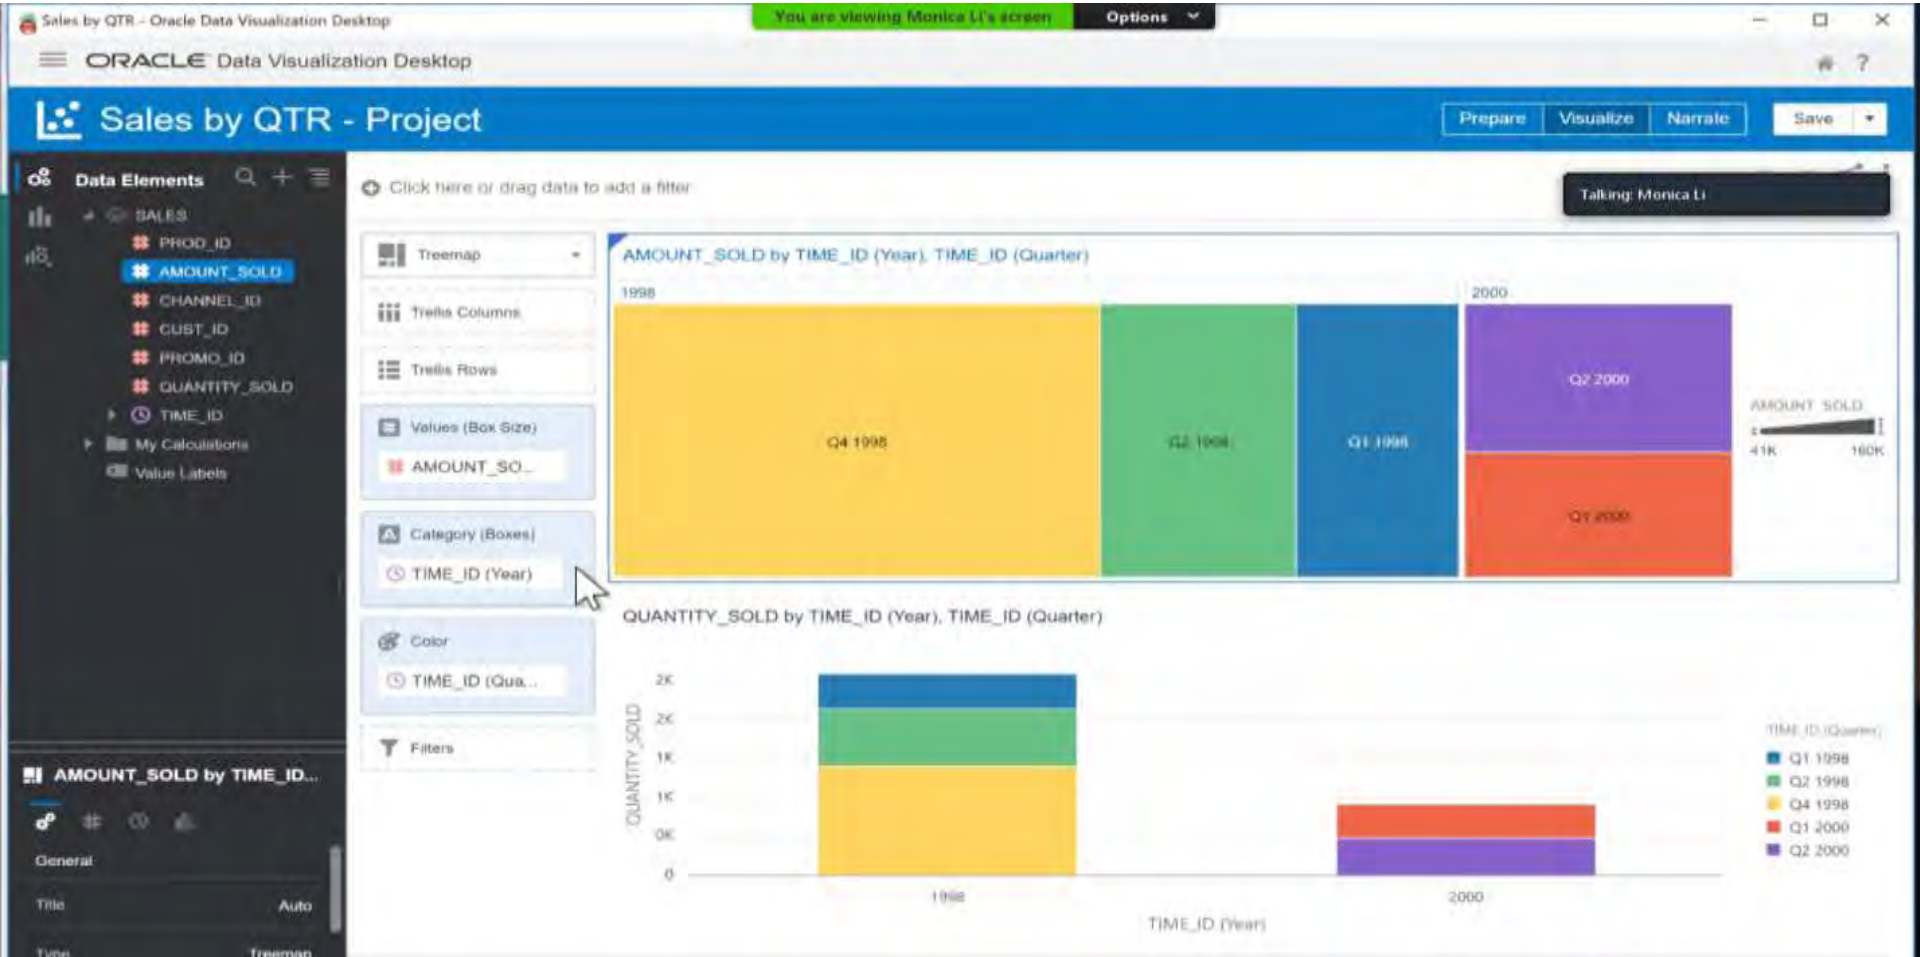

Much More Complex Available

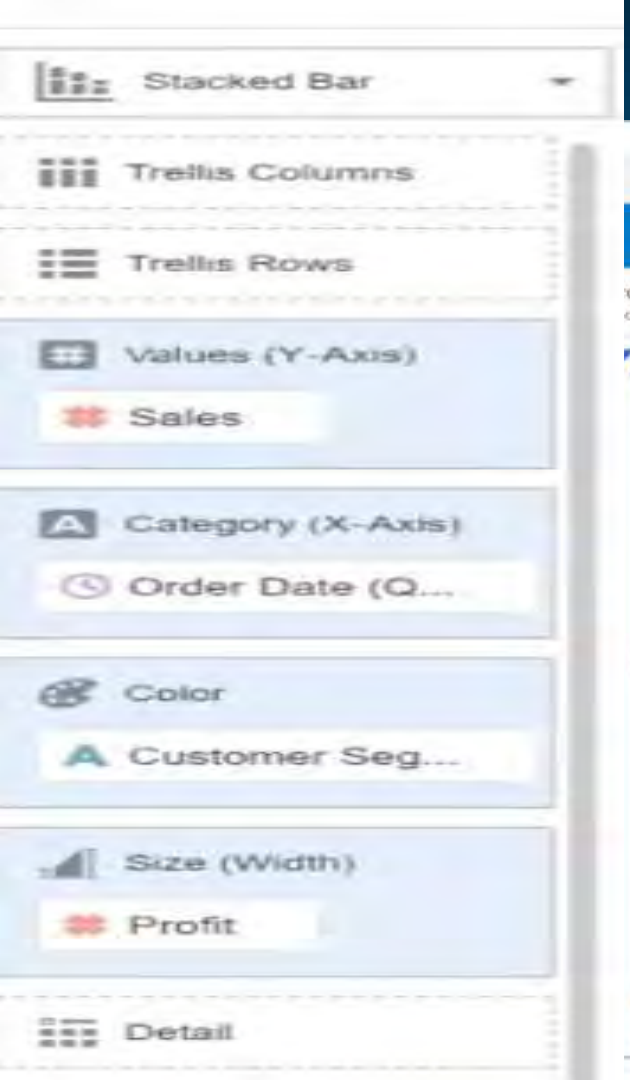

Oracle Analytics Cloud (OAC) to Cluster Data

Easy Models

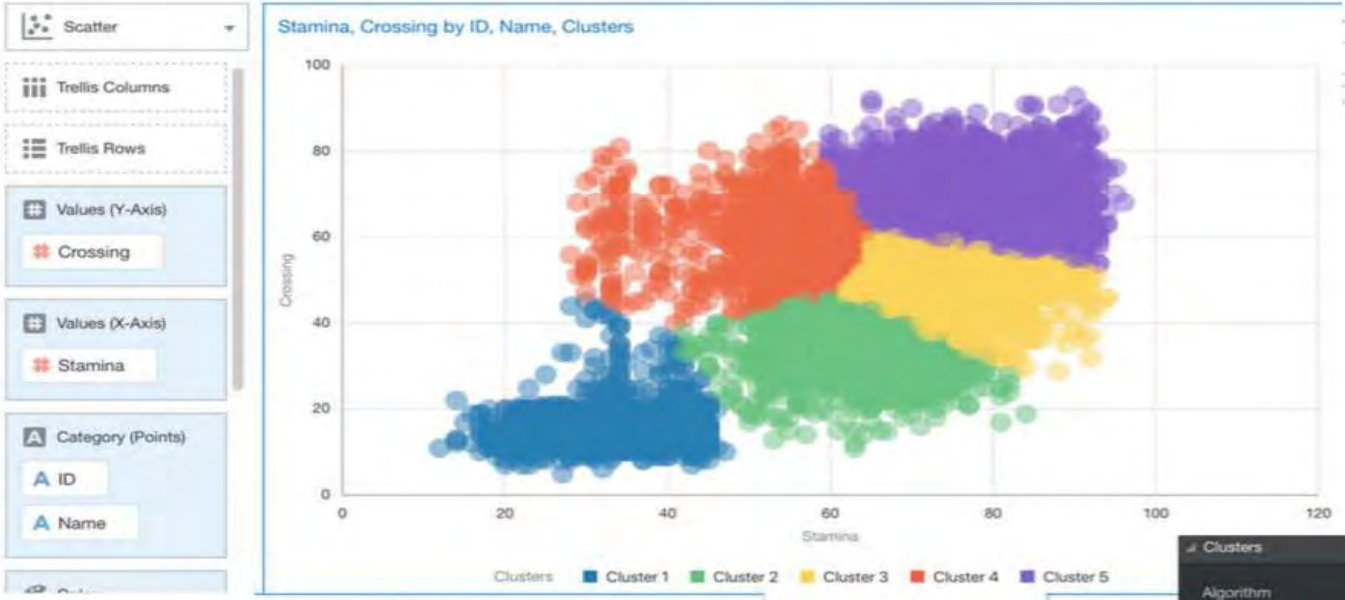

Oracle Analytic Summit 2020

Advanced Innovations to Leverage from Oracle:

Built-In Innovations to Drive Faster Transformation

AI / ML Embedded into Apps

IoT Apps

Human Interface

ORACLE

Future World – Enabling Innovation!

64-Bit advancement of Directly addressable memory

	Address Direct	Indirect/Extended
4 Bit:	16	(640)
8 Bit:	256	(65,536)
16 Bit:	65,536	(1,048,576)
32 Bit:	4,294,967,296	
64 Bit:	18,446,744,073,709,551,616	

As the hardware physically implements the theoretical possibilities of 64-Bit, things will dramatically change…. …moving from 32 bit to 64 bit will be like moving from 4 bit to 32 bit or like moving from 1971 to 2000 overnight.

96

Addressable Memory Growth

Shift	Increase (Result)
• 8-bit to 16-bit	**1 mph (Windows)**
• 16-bit to 32-bit	65K mph (Internet)
• 32-bit to 64-bit	300T mph (Robotics/3D…etc.)
• 64-bit to 128-bit	5T*T*B mph (Implants/4D)

Here Come the Robots!

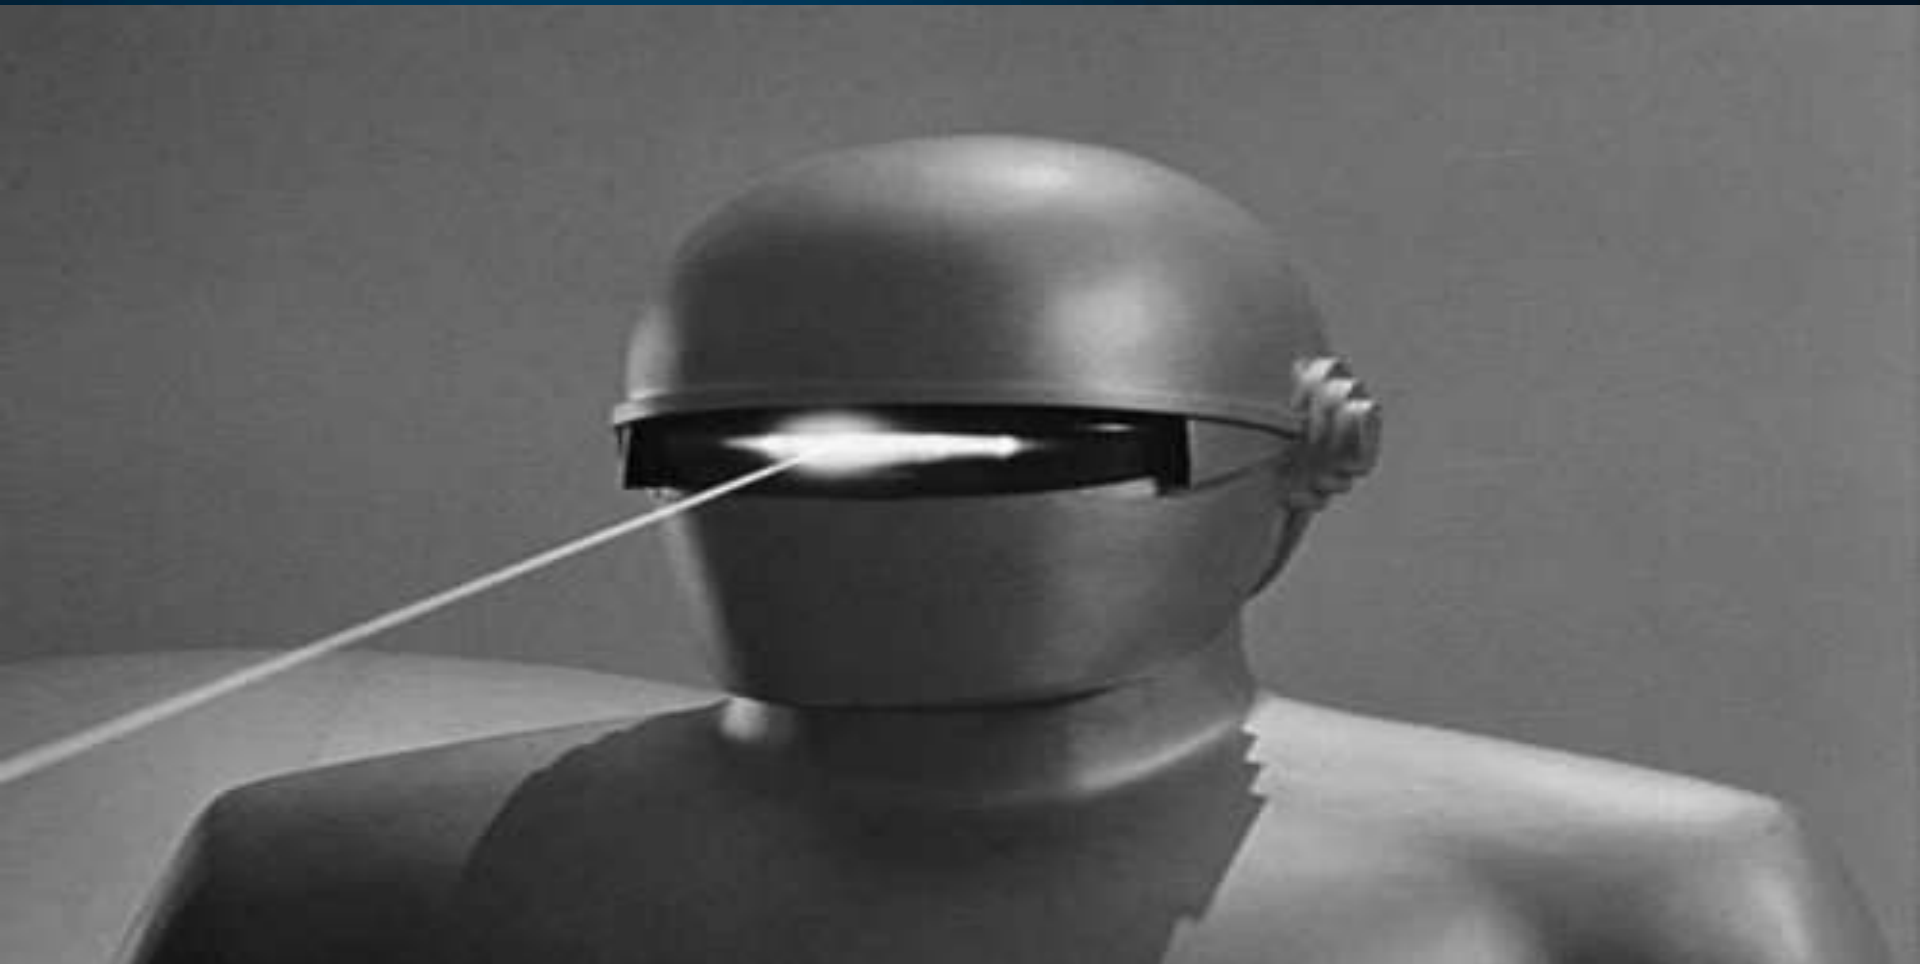

Robots We Grew up With…

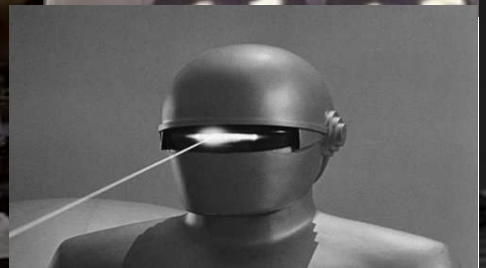

Your Robotic Backup DBA … *may be Your Clone!*

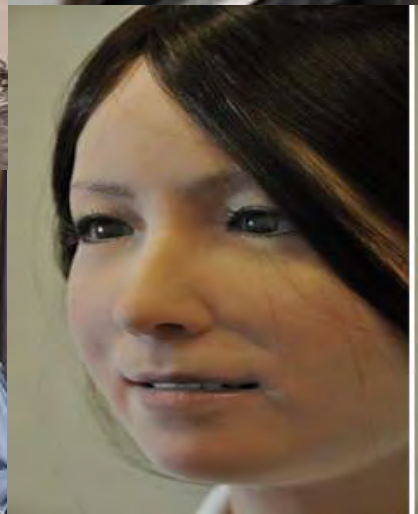

From Agnes with Love *(Computers try to help)*

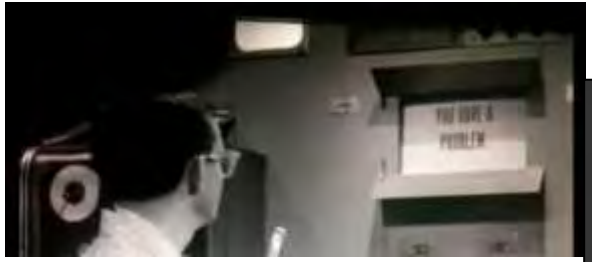

Service Robots at Work… 24x7

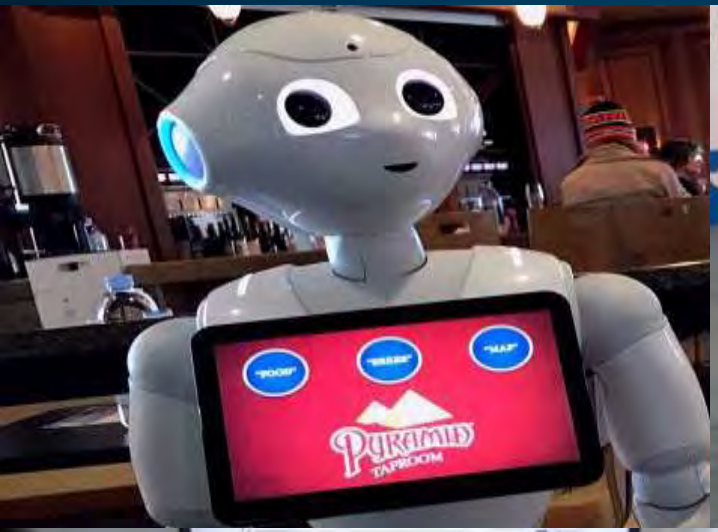

How to Connect Pepper to Oracle (so far)

How Pepper + Oracle are Connected

ORACLE MOBILE

- ORACLE POLICY AUTOMATION
- ORACLE Commerce Cloud
- ORACLE Service Cloud
- ORACLE CONFIGURE, PRICE, AND QUOTE
- ORACLE MARKETING CLOUD
- ORACLE SALES CLOUD

Meet Pepper

- Humanoid robot designed by SoftBank Robotics
- Able to recognize principal human emotions and tailor interactions
- Facial / object recognition
- Spatial awareness

SoftBank Robotics

Use Oracle Virtual Assistant with Robots - FYI Only!

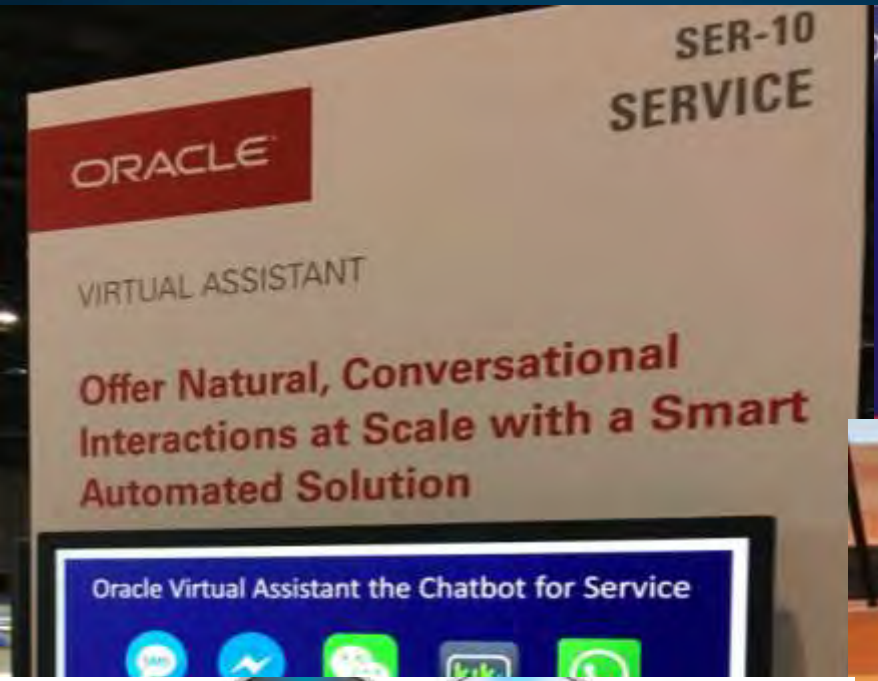

Oracle Virtual Assistant Interface

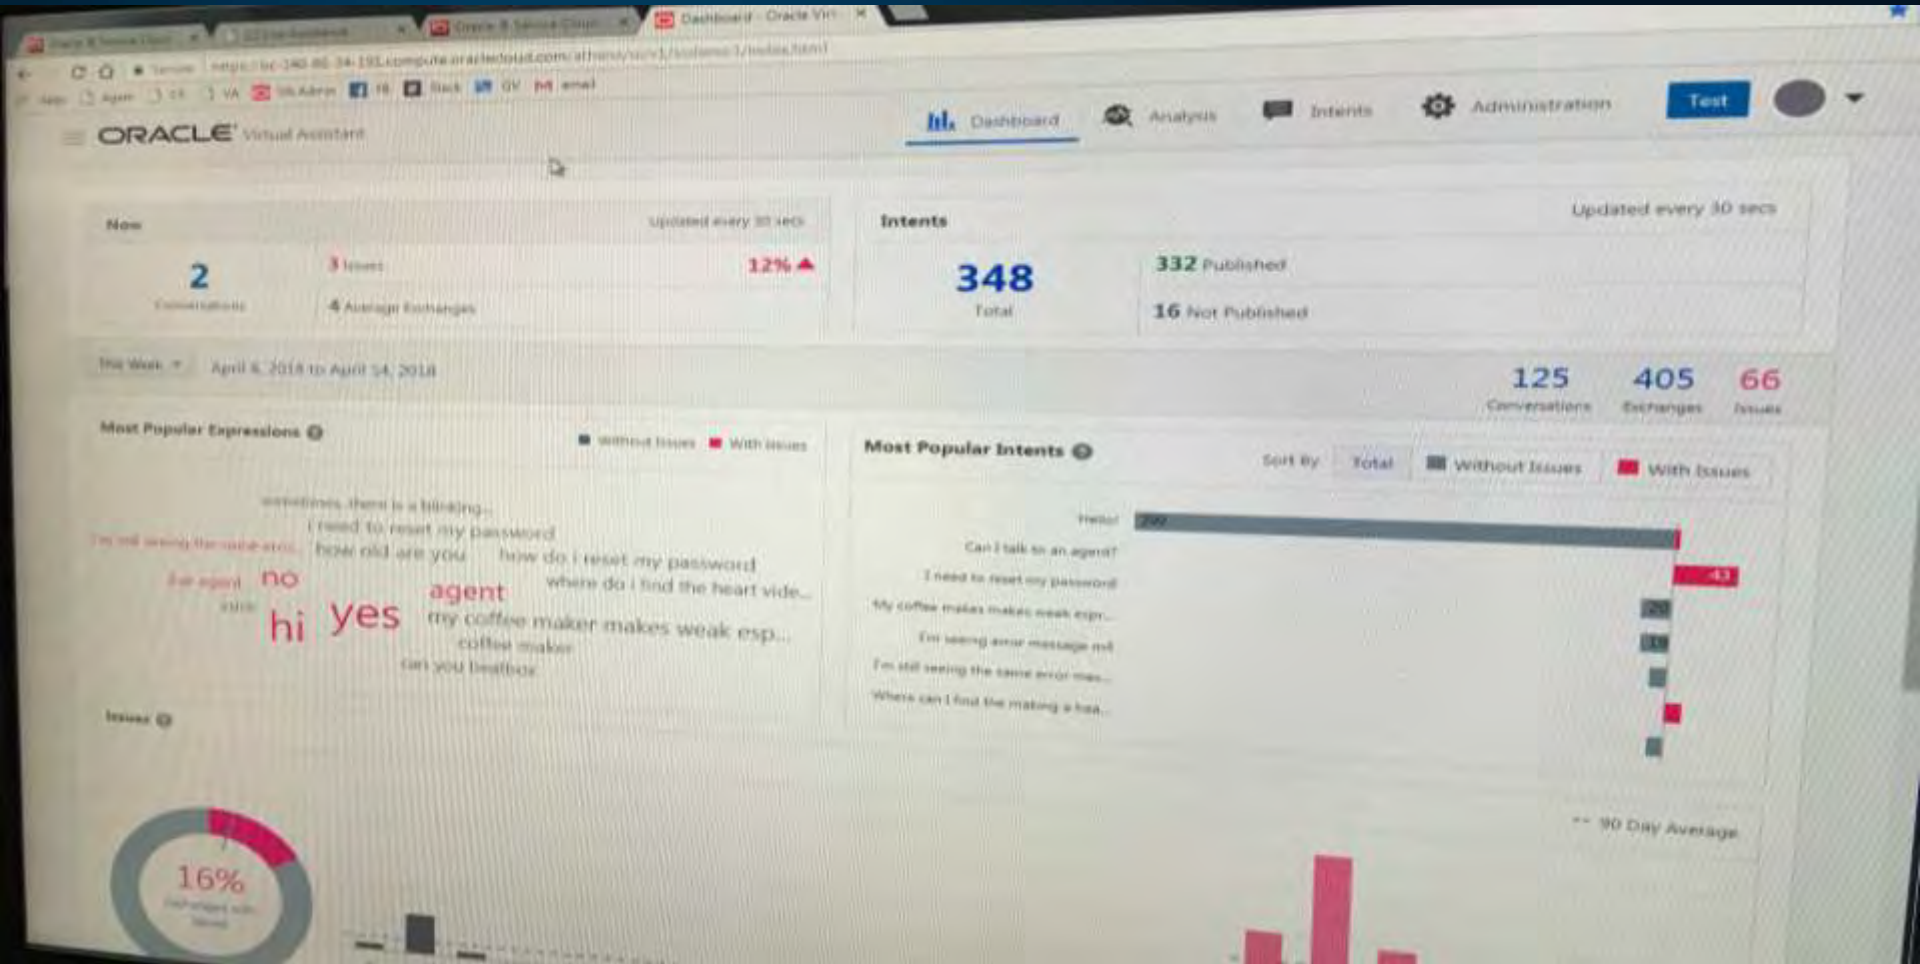

The Brain Center at Whipple's - *Robots Coming*

Robotics/Automation Impact to Jobs

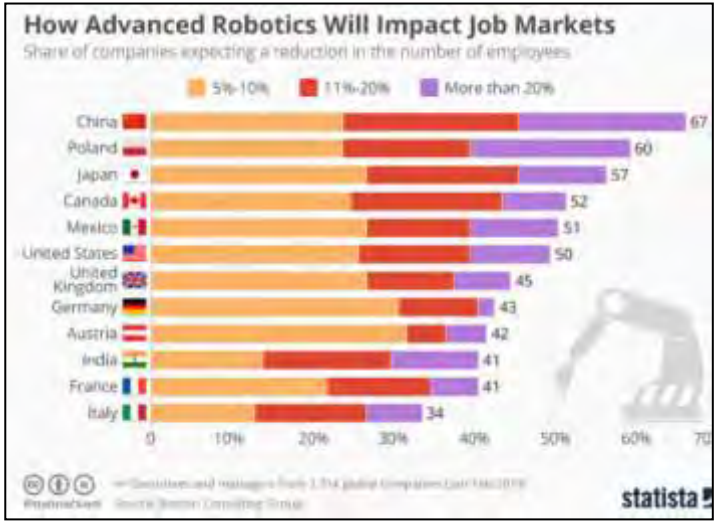

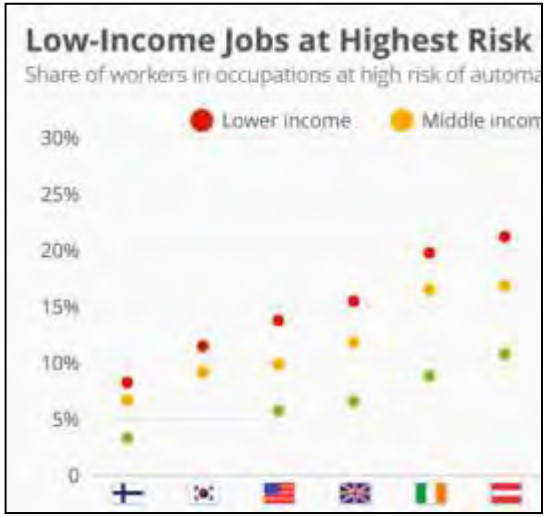

Robotics/Automation Impact to Jobs

The Obsolete Man (Twilight Zone)

Autonomous Database – Replacing the DBA?

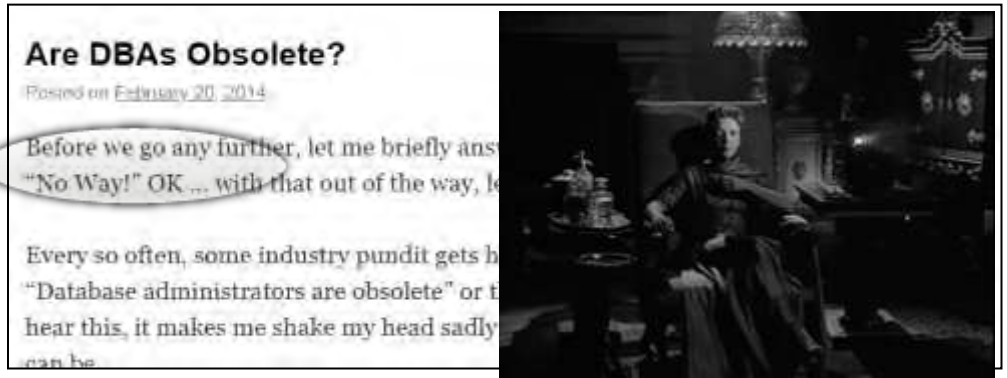

Good News: DBA +11% Increase (2016-2026)

Leveraging – DB, AI & Virtual Reality!

Virtual Reality

- Immersion in virtual worlds
- Total interaction with virtual
- E.g. Oculus Rift

Mixed Reality

- Virtual World integrated to reality
- Interaction between reality and virtual
- E.g. Microsoft HoloLens

Augmented Reality

- Virtual on top of reality
- Limited interaction with the virtual
- E.g. Smartphones & tablets

actimage | digital intelligence

Apple – Tech Innovator!

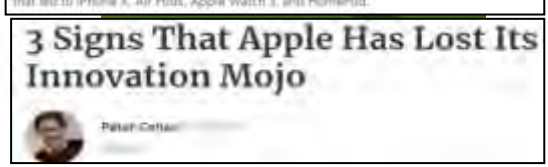

Amazon - Retail Innovator (since 1994)

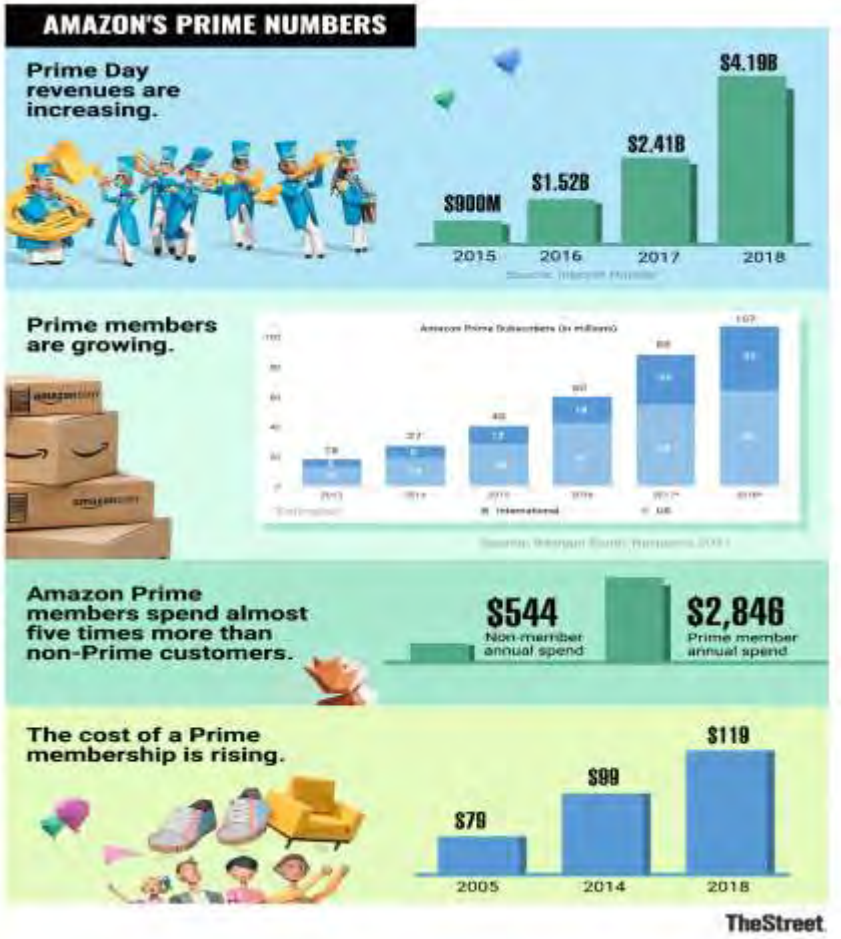

Google - Marketing Innovator (& Leveraging Data)!

Oracle Focus: Cloud, Data, Apps, Security, Self-Service, ML, & AI!

Innovation Components for Business!

Machine Learning: Oracle ML

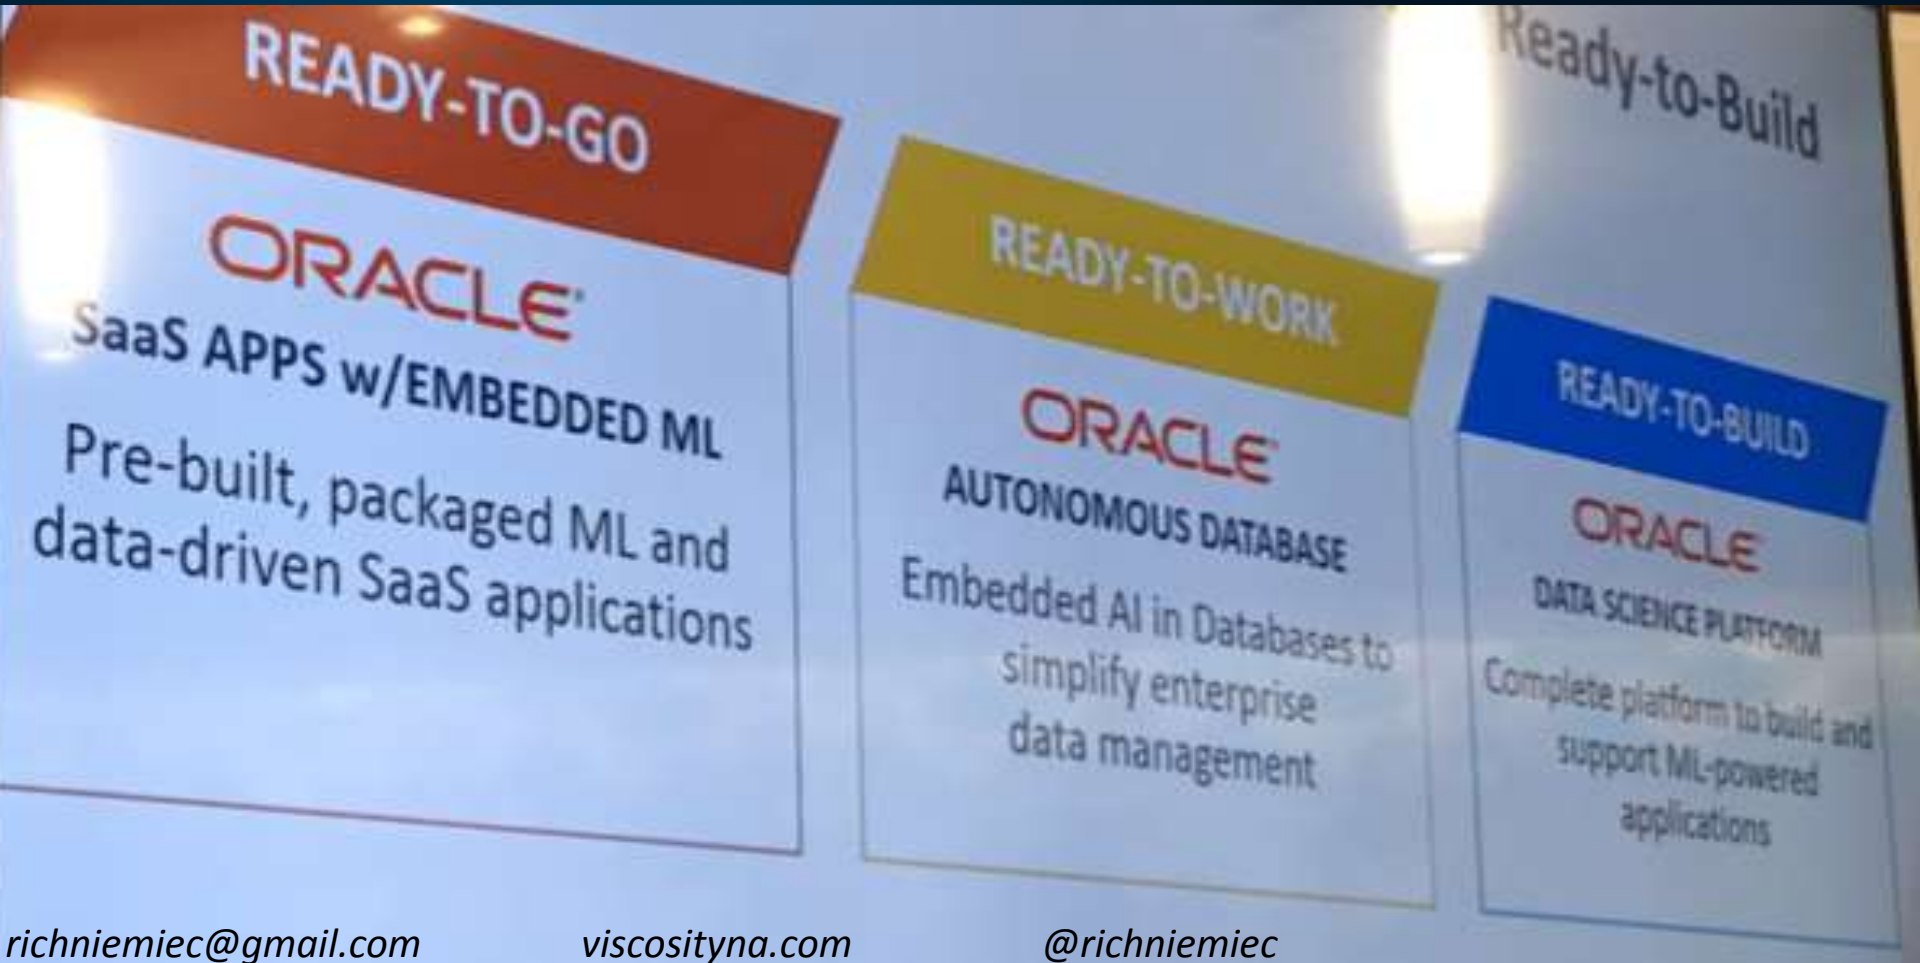

Oracle Chatbots & AI

Viscosity helps to Leverage NOW!

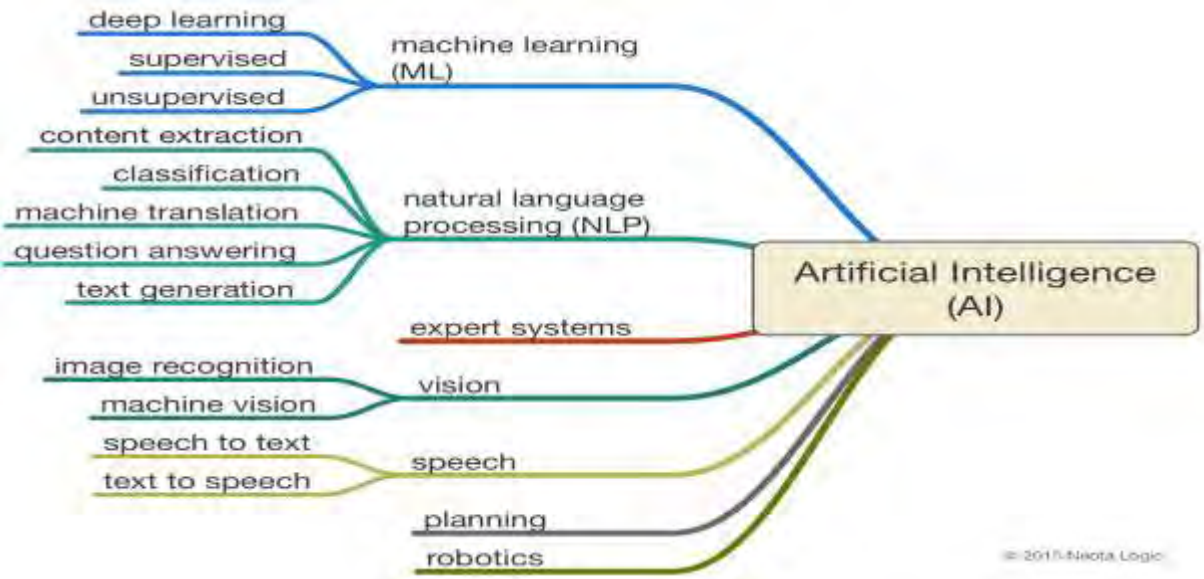

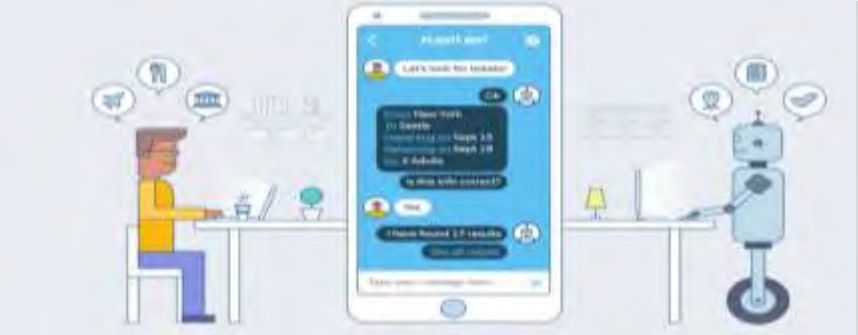

OpenAI - Ten Years Later… the Baby Talks!

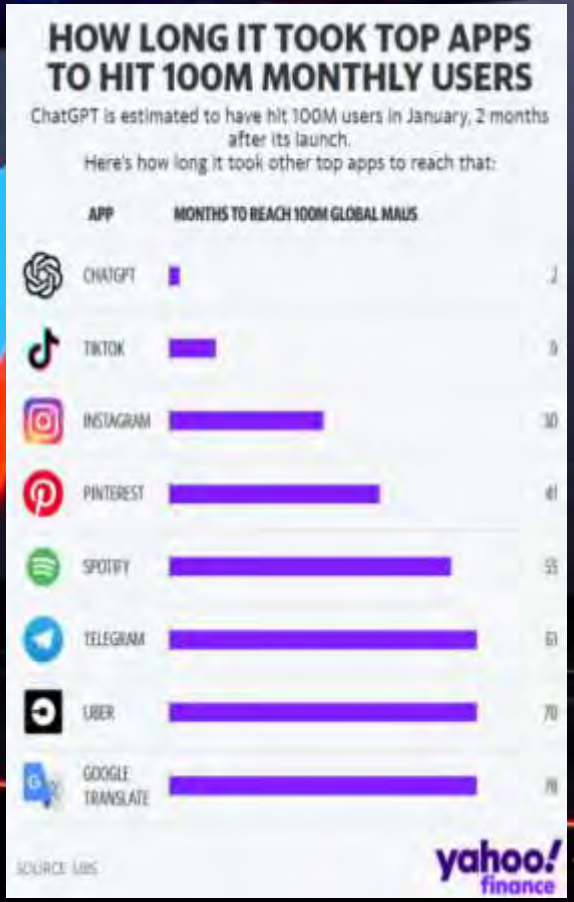

About One Year Ago: ChatGPT 3.5 was Born … the Baby Talks

- Developed by OpenAI: A Company Founded by Elon Musk 10 Years Ago
- Huge Artificial Neural Network: Multi-Billion Parameter Large Language Model (LLM)
- Trained on Language: Trillions of Data Elements … Wikipedia + Public Internet
- ChatGPT 3.5 Talks: Generates Answers to Questions … Engages in Conversations

A Big Surprise

ChatGPT's Scale Enabled it to Understand and Generate Language

Main Difference Between GPT-4 and ChatGPT

Unlike ChatGPT, GPT-4 can see and understand images, whereas the former is limited to text.

While ChatGPT had several issues upon its arrival, Microsoft claims that GPT-4 has been trained to avoid several malicious prompts.

ChatGPT

GPT-4 has arrived. It will blow ChatGPT out of the water The Washington Post 3/14/2023

ChatGPT was launched on November 30, 2022, by San Francisco–based OpenAI (the creator of the GPT series of large language models; DALL·E 2, a diffusion model used to generate images; and Whisper, a speech transcription model).

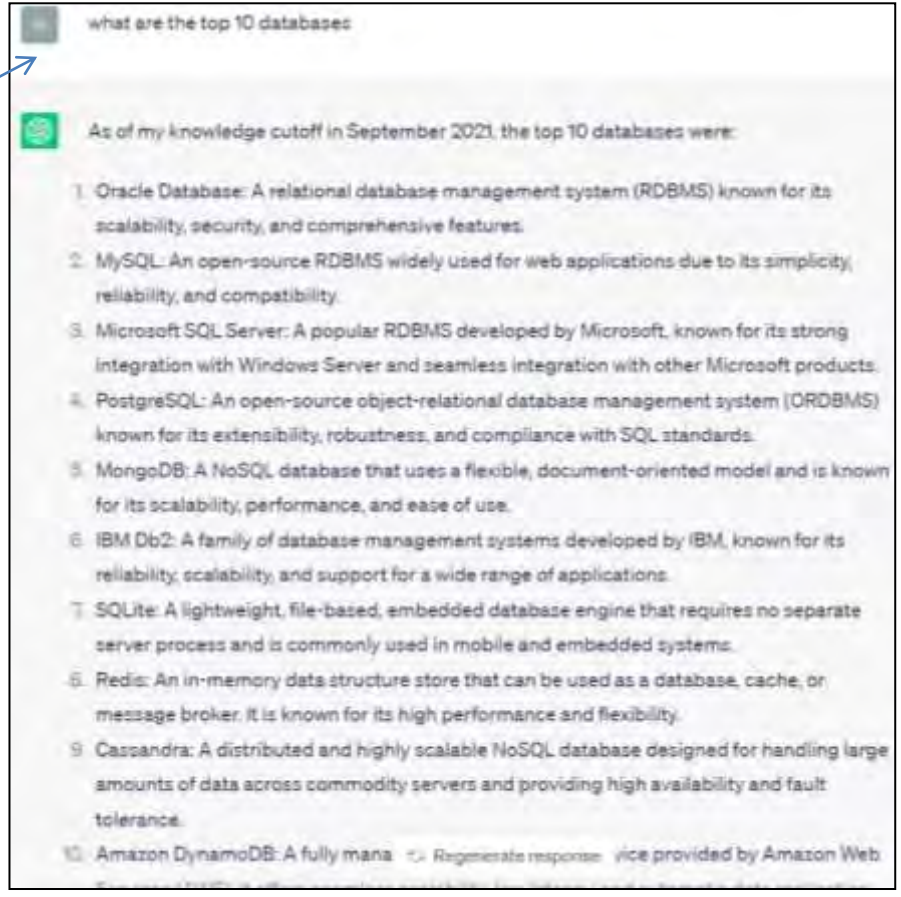

Google's Bard

Bard

Reset chat

Bard Activity

FAQ

Updates

Help

R What are the top 10 databases

Here is a brief overview of each of the top 10 databases:

- **Oracle** is a commercial database that is known for its scalability and performance. It is often used by large enterprises that need to store and manage large amounts of data.
- **MySQL** is an open-source database that is known for its ease of use and flexibility. It is often used for web applications and other applications that need to be scalable and easy to maintain.
- **Microsoft SQL Server** is a commercial database that is known for its integration with other Microsoft products. It is often used by businesses that use Microsoft's Windows operating system and other Microsoft products.
- **PostgreSQL** is an open-source database that is known for its power and flexibility. It is often used for applications that need to store complex data or that need to be highly scalable.
- **MongoDB** is a NoSQL database that is known for its flexibility and scalability. It is often used for applications that need to store large amounts of unstructured data.
- **Redis** is a NoSQL database that is known for its speed and performance. It is often used for applications that need to store and manage large amounts of data in real time.
- **IBM DB2** is a commercial database that is known for its scalability and performance. It is often used by businesses that need to store and manage large amounts of data.
- **Elasticsearch** is a search engine database that is known for its speed and scalability. It is often used for applications that need to search and index large amounts of data.
- **Cassandra** is a NoSQL database that is known for its scalability and fault tolerance. It is often used for applications that need to store and manage large amounts of data in a distributed environment.
- **MariaDB** is a fork of MySQL that is known for its stability and performance. It is often used as a replacement for MySQL when MySQL is not available or when a more stable database is needed.

Enter a prompt here

Bard may display inaccurate or offensive information that doesn't represent Google's views.

Transformers - Google, 2017

See: Illustrated Guide to Transformers Neural Network: YouTube step by step

FYI

Attention Is All You Need

Ashish Vaswani* Google Brain avaswani@google.com
Noam Shazeer* Google Brain noam@google.com
Niki Parmar* Google Research nikip@google.com
Jakob Uszkoreit* Google Research usz@google.com
Llion Jones* Google Research llion@google.com
Aidan N. Gomez* † University of Toronto aidan@cs.toronto.edu
Łukasz Kaiser* Google Brain lukaszkaiser@google.com
Illia Polosukhin* ‡ illia.polosukhin@gmail.com

Abstract

The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English-to-German translation task, improving over the existing best results, including ensembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.8 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature. We show that the Transformer generalizes well to other tasks by applying it successfully to English constituency parsing both with large and limited training data.

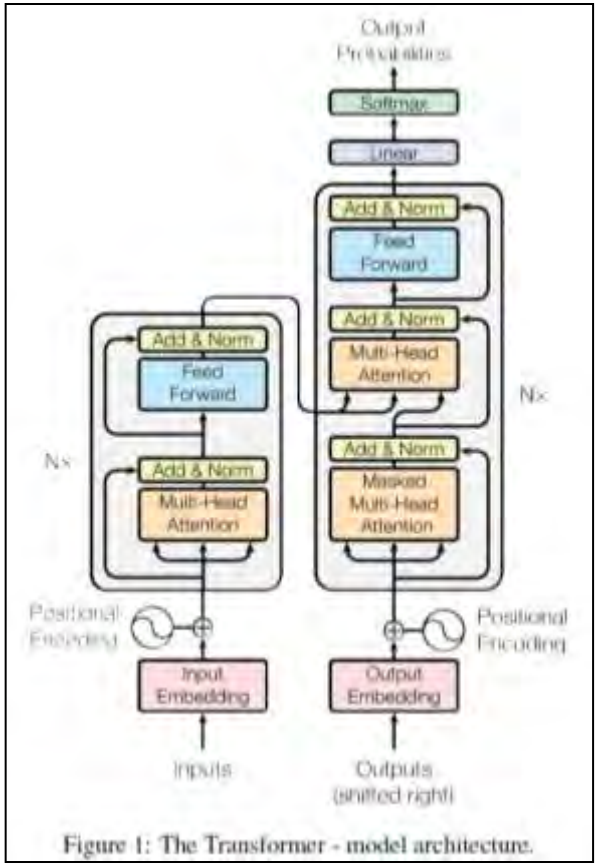

Figure 1: The Transformer - model architecture.

Outperformed Recurrent NN, GRU, & LSTM

Transformers Have Longer Term Memory

- Google BERT
- OpenAI GPT

Next word learned tend to much earlier words thru Backprop

Leverages Parallelism to train model

Stanford (113 authors / 215 pgs.) on Foundation Models 2021/2022

On the Opportunities and Risks of Foundation Models

Rishi Bommasani* Drew A. Hudson Ehsan Adeli Russ Altman Simran Arora
Sydney von Arx Michael S. Bernstein Jeannette Bohg Antoine Bosselut Emma Brunskill
Erik Brynjolfsson Shyamal Buch Dallas Card Rodrigo Castellon Niladri Chatterji
Annie Chen Kathleen Creel Jared Quincy Davis Dorottya Demszky Chris Donahue
Moussa Doumbouya Esin Durmus Stefano Ermon John Etchemendy Kawin Ethayarajh
Li Fei-Fei Chelsea Finn Trevor Gale Lauren Gillespie Karan Goel Noah Goodman
Shelby Grossman Neel Guha Tatsunori Hashimoto Peter Henderson John Hewitt
Daniel E. Ho Jenny Hong Kyle Hsu Jing Huang Thomas Icard Saahil Jain
Dan Jurafsky Pratyusha Kalluri Siddharth Karamcheti Geoff Keeling Fereshte Khani
Omar Khattab Pang Wei Koh Mark Krass Ranjay Krishna Rohith Kuditipudi
Ananya Kumar Faisal Ladhak Mina Lee Tony Lee Jure Leskovec Isabelle Levent
Xiang Lisa Li Xuechen Li Tengyu Ma Ali Malik Christopher D. Manning
Suvir Mirchandani Eric Mitchell Zanele Munyikwa Suraj Nair Avanika Narayan
Deepak Narayanan Ben Newman Allen Nie Juan Carlos Niebles Hamed Nilforoshan
Julian Nyarko Giray Ogut Laurel Orr Isabel Papadimitriou Joon Sung Park Chris Piech
Eva Portelance Christopher Potts Aditi Raghunathan Rob Reich Hongyu Ren
Frieda Rong Yusuf Roohani Camilo Ruiz Jack Ryan Christopher Ré Dorsa Sadigh
Shiori Sagawa Keshav Santhanam Andy Shih Krishnan Srinivasan Alex Tamkin
Rohan Taori Armin W. Thomas Florian Tramèr Rose E. Wang William Wang Bohan Wu
Jiajun Wu Yuhuai Wu Sang Michael Xie Michihiro Yasunaga Jiaxuan You Matei Zaharia
Michael Zhang Tianyi Zhang Xikun Zhang Yuhui Zhang Lucia Zheng Kaitlyn Zhou
Percy Liang*[1]

Center for Research on Foundation Models (CRFM)
Stanford Institute for Human-Centered Artificial Intelligence (HAI)
Stanford University

AI is undergoing a paradigm shift with the rise of models (e.g., BERT, DALL-E, GPT-3) trained on broad data (generally using self-supervision at scale) that can be adapted to a wide range of downstream tasks. We call these models foundation models *to underscore their critically central yet incomplete character. This report provides a thorough account of the opportunities and risks of foundation models, ranging*

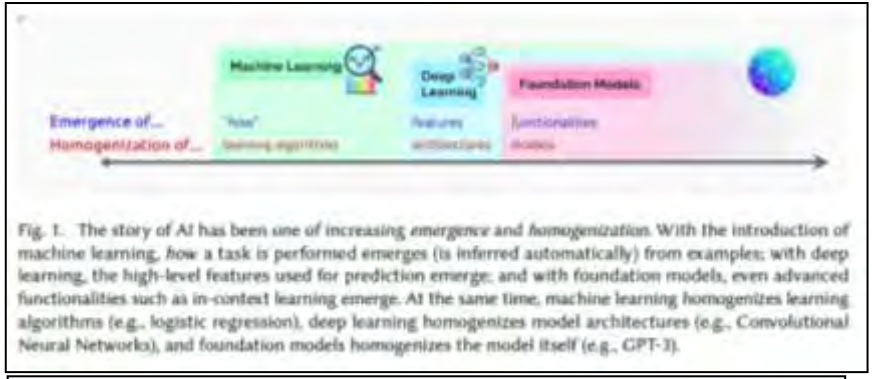

Fig. 1. The story of AI has been one of increasing *emergence* and *homogenization*. With the introduction of machine learning, *how* a task is performed emerges (is inferred automatically) from examples; with deep learning, the high-level features used for prediction emerge; and with foundation models, even advanced functionalities such as in-context learning emerge. At the same time, machine learning homogenizes learning algorithms (e.g., logistic regression), deep learning homogenizes model architectures (e.g., Convolutional Neural Networks), and foundation models homogenizes the model itself (e.g., GPT-3).

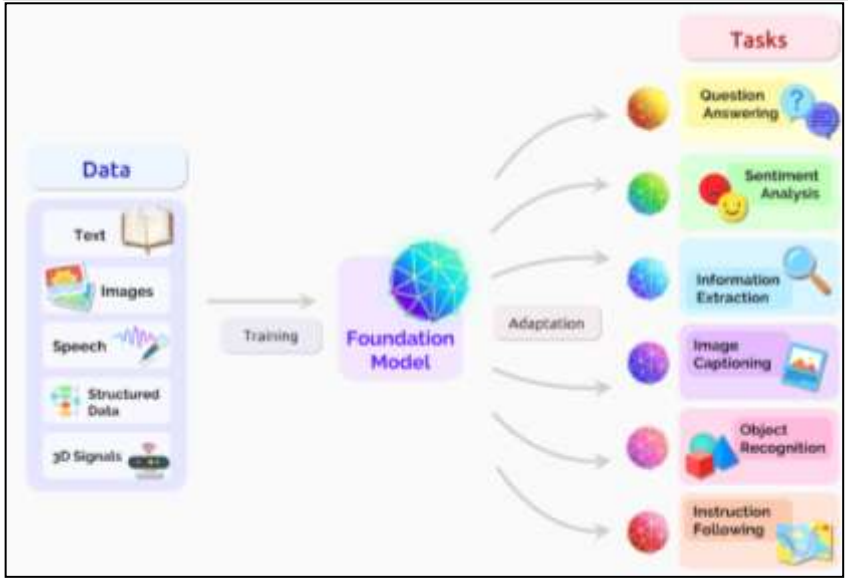

From Stanford Paper

FYI

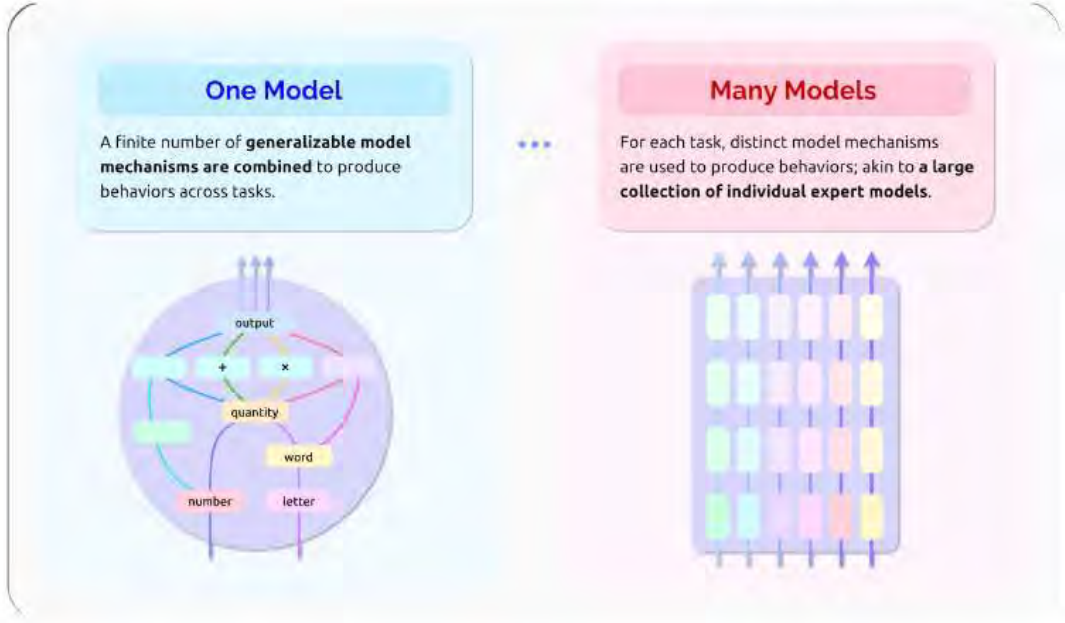

Fig. 23. The one model–many model nature of foundation models: A central interpretability question is to understand where a foundation model lies on the spectrum between *one model* and *many models*. As one model, behavior can be made interpretable by identifying and characterising the finite number of generalizable model mechanisms used to produce behaviors across tasks (e.g., mechanisms that assign meaning to words, compare quantities, and perform arithmetic). As many models, explanations of model behavior in one task are not necessarily informative about behavior in other tasks, thus requiring the independent study of behavior in each task.

4.7 Security and privacy

Authors: Florian Tramèr, Rohith Kuditipudi*, Xuechen Li**

Fig. 20. Risks and opportunities raised by foundation models for security and privacy of ML systems.

Will Businesses of the Future just be a Single Foundation Model to do all Tasks? *(my own question)*

From Stanford Paper (Example Application)

FYI

Section on Healthcare & Biomedicine

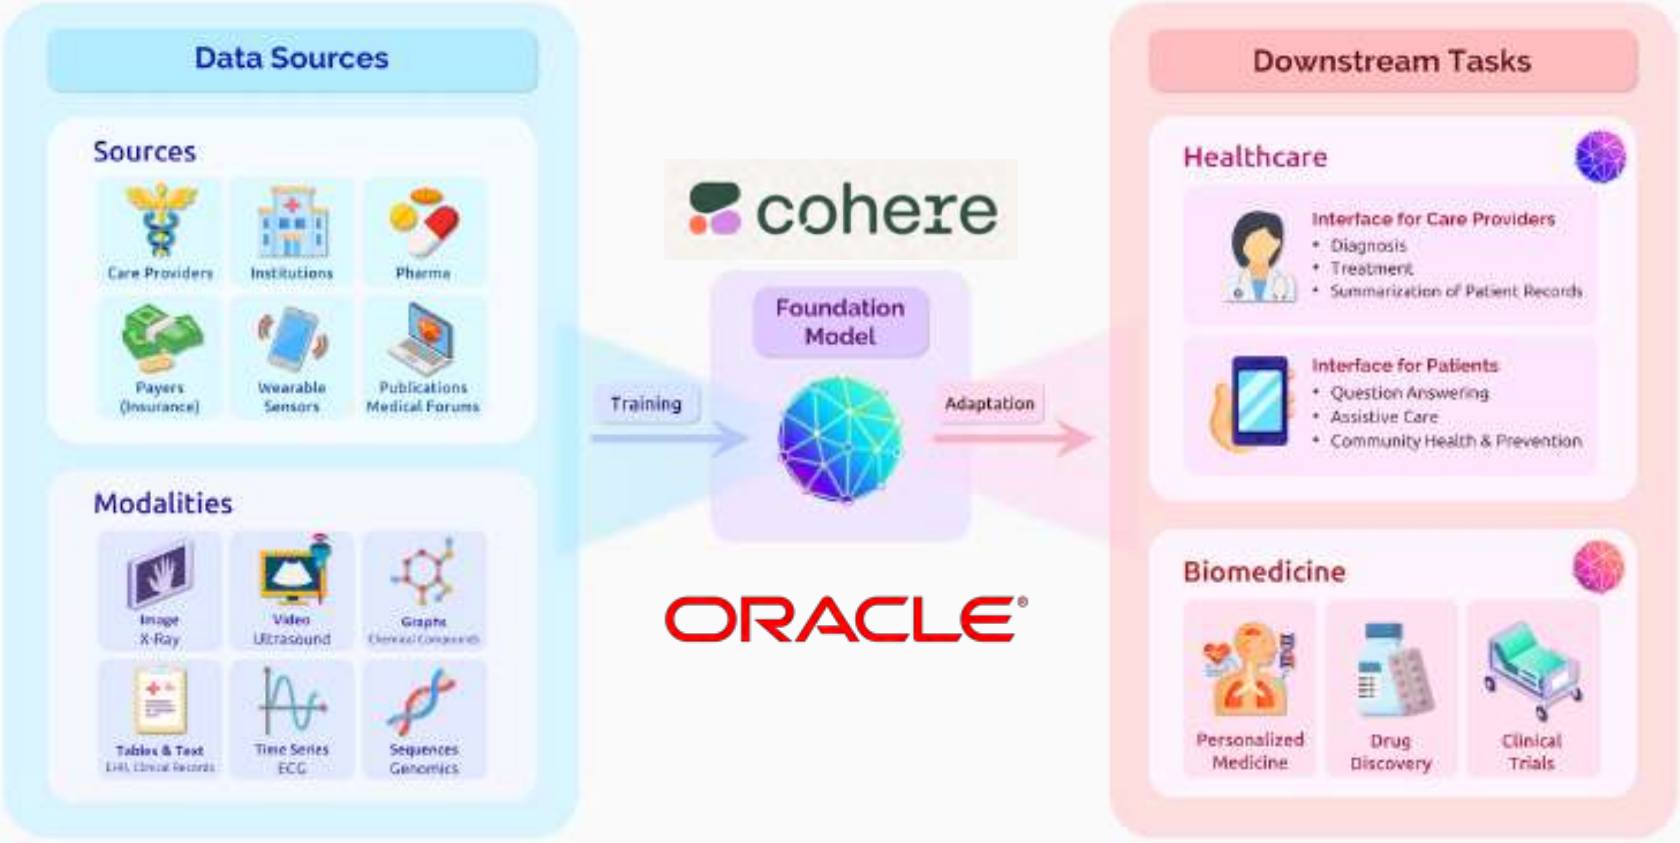

Vector Search

Imagine a house-hunting app that helps customers find houses for sale that are similar to a picture the customer uploads

The distance between the vectors is proportional to their semantic similarity

Create table with Vector Data Type & Blob

```
CREATE TABLE house_for_sale (house_id      number,
                             price         number,
                             city          varchar2(400),
                             house_photo   blob,
                             house_vector  vector
);
```

Find houses that are similar to this picture and match the customer's preferred city and budget

```
SELECT …
FROM   house_for_sale
WHERE  price <= (SELECT budget      FROM customer …)
AND    city  in (SELECT search_city FROM customer …)
ORDER BY vector_distance(house_vector, :input_vector);
```

AI Vector search Preview!

Announcing:
AI Vector Search in
Oracle Database 23c

Sign up
for Preview Now

The Oracle AI Stack

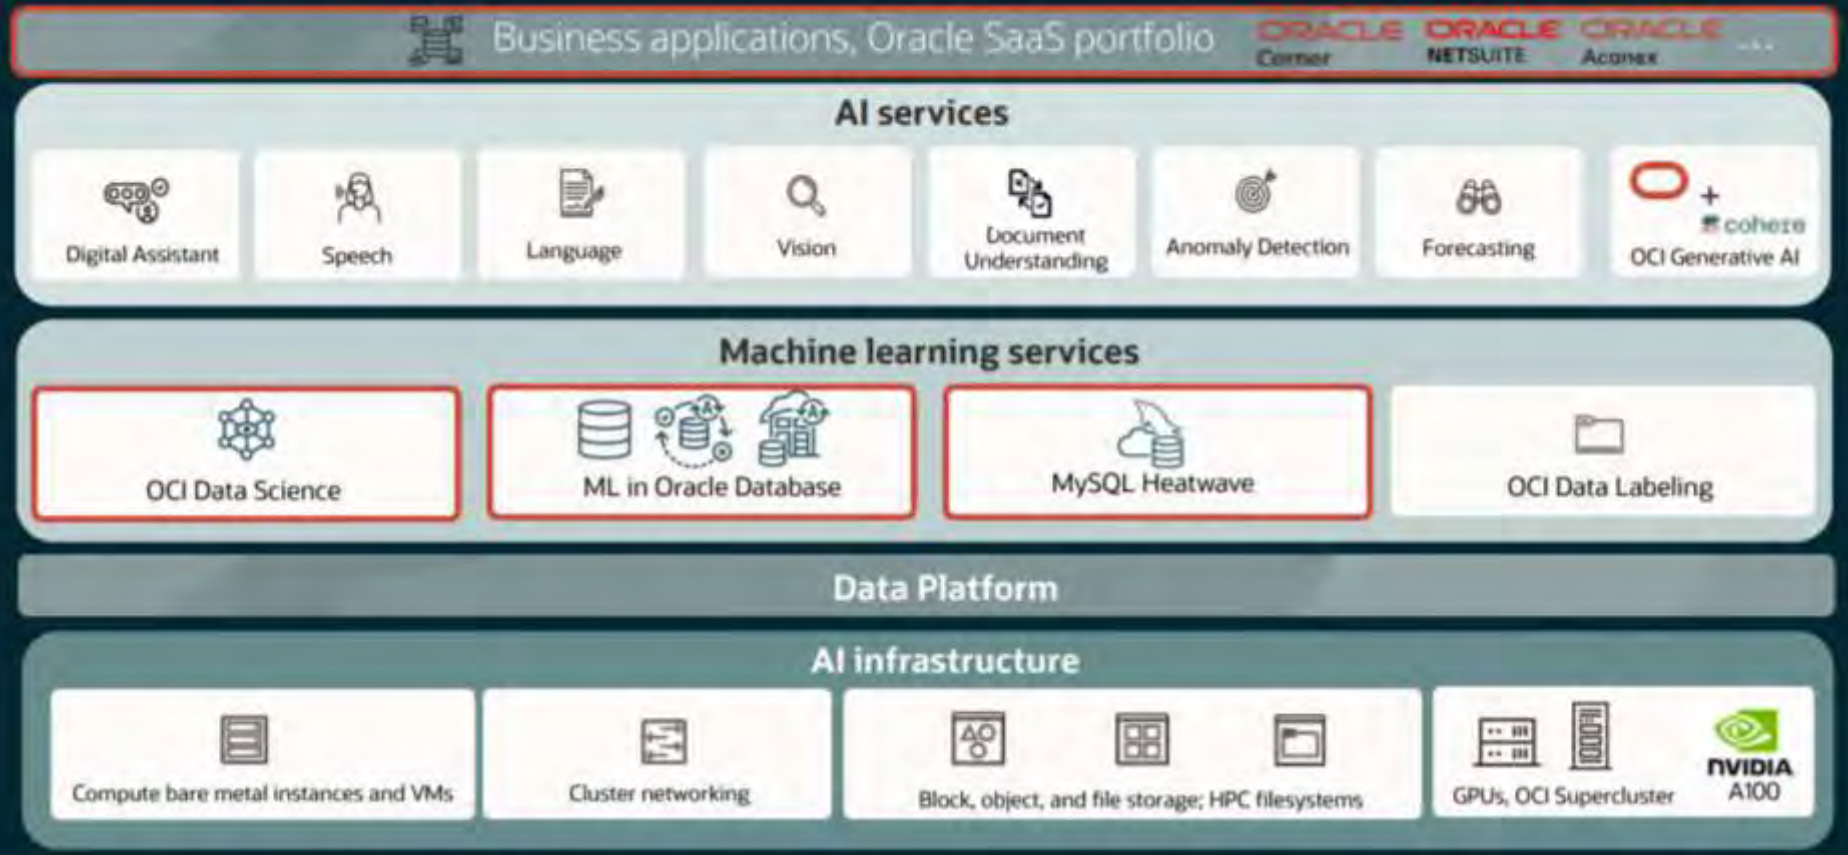

Tech Trends - Gartner Hype Cycle

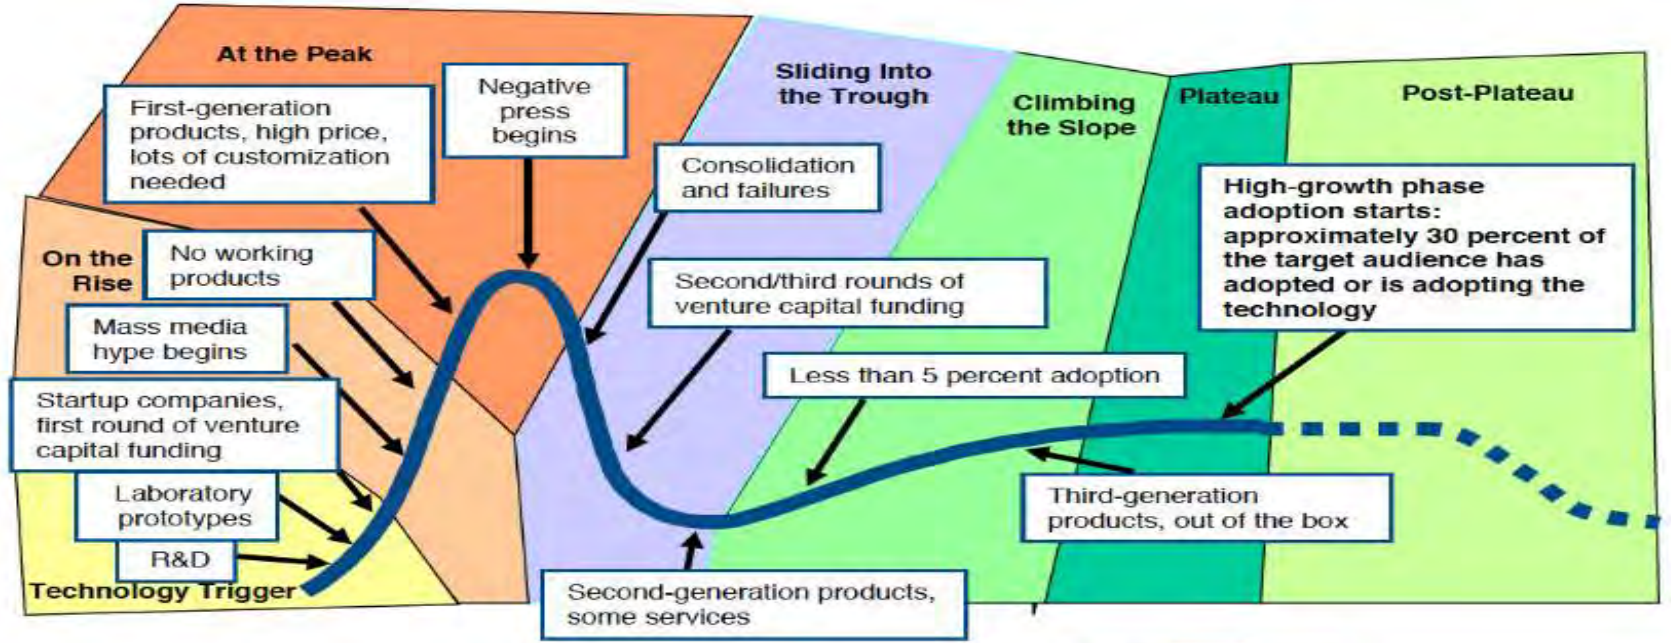

Source: Gartner Research (May 2003)

Tech Trends - Gartner Hype Cycle 2013: All about Tech …

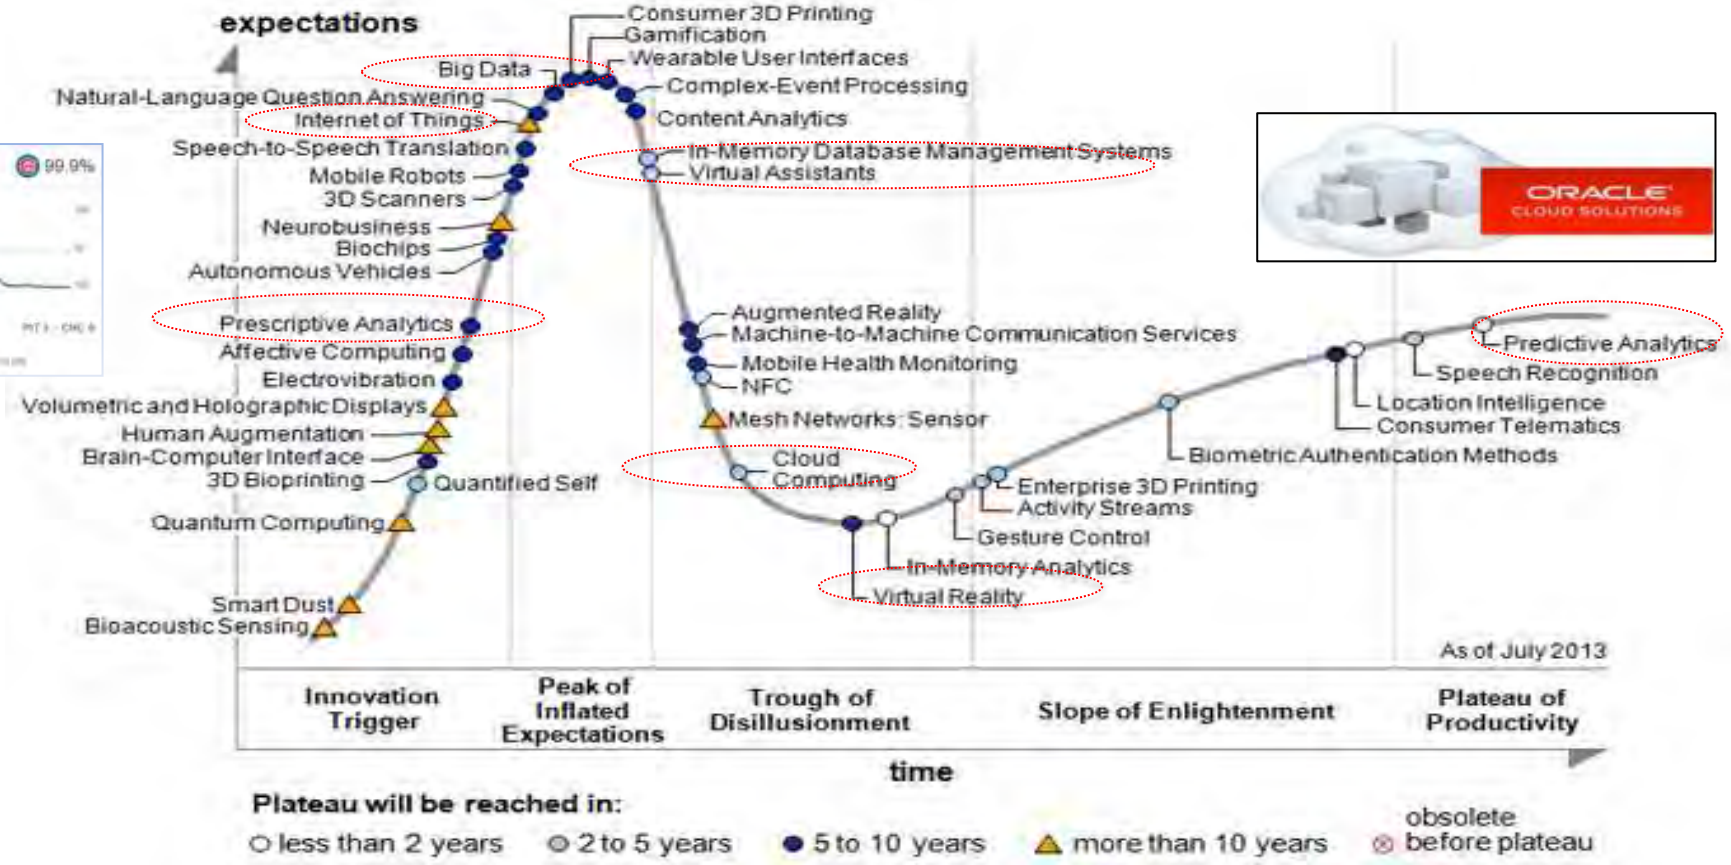

Tech Trends - Gartner Hype Cycle 2015: All about Robotics …

Gartner Hype Cycle July, 2016

All about Implants & the Twilight Zone

Source: Gartner (July 2016)

Gartner Hype Cycle August 2018 - All about Robotics …
All about Tech Creating a New Reality (just 5 years later!)

Digital Twin
Biochips
Smart Workspace
Brain-Computer Interface
Autonomous Mobile Robots
Smart Robots
Deep Neural Network ASICs
AI PaaS
Quantum Computing
5G
Volumetric Displays
Self-Healing System Technology
Conversational AI Platform
Autonomous Driving Level 5
Edge AI
Exoskeleton
Blockchain for Data Security
Knowledge Graphs
4D Printing
eral Intelligence
Smart Dust
Flying Autonomous Vehicles
Biotech — Cultured or Artificial Tissue
Neuromorphic Hardware
Deep Neural Nets (Deep Learning)
Carbon Nanotube
IoT Platform
Virtual Assistants
Silicon Anode Batteries
Blockchain
Connected Home
Autonomous Driving Level 4
Mixed Reality
Smart Fabrics
Augmented Reality

Innovation Trigger | Peak of Inflated Expectations | Trough of Disillusionment | Slope of Enlightenment | Plateau of Productivity

time

Plateau will be reached:

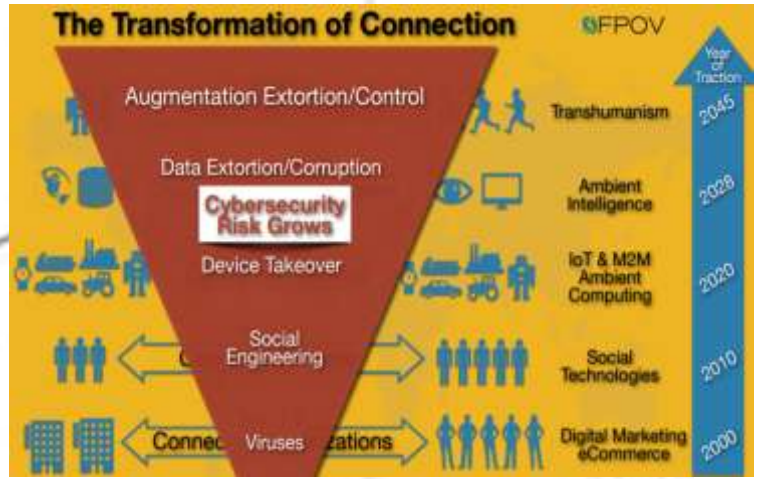

Gartner 2020 hype cycle

Gartner 2021 hype cycle for Emerging Tech

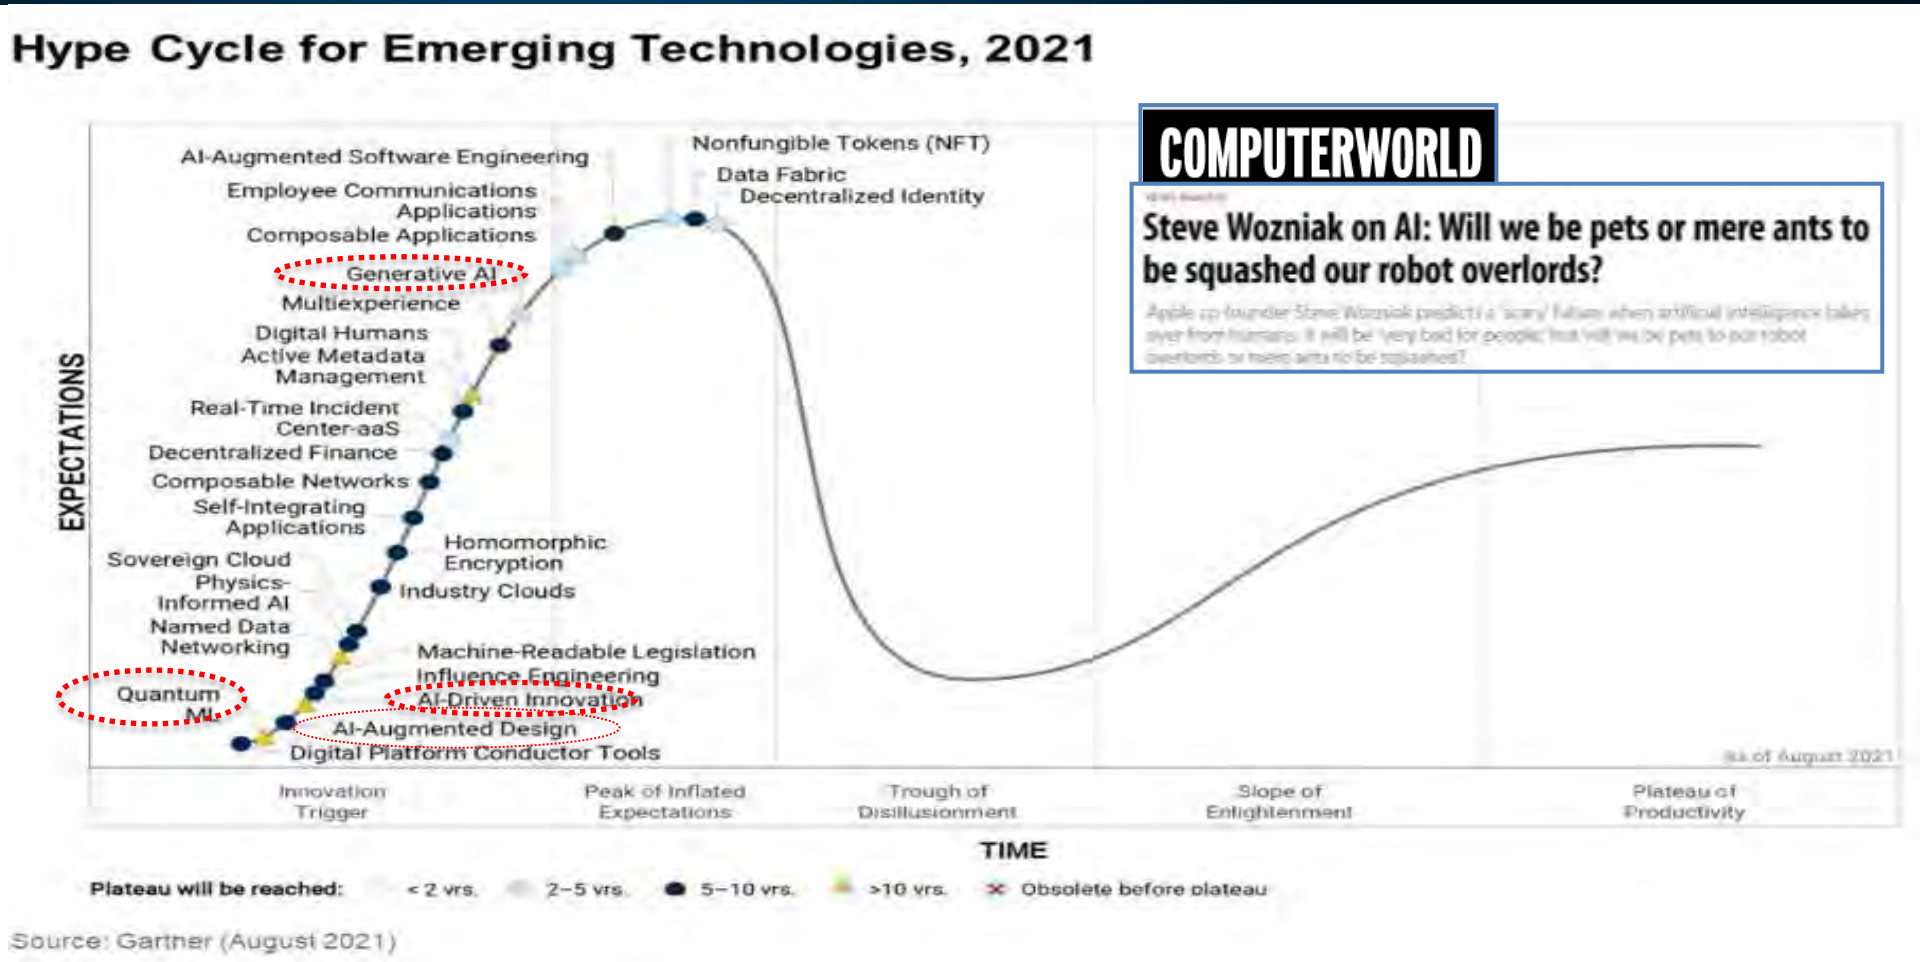

Source: Gartner (August 2021)

Final Thoughts… world changing fast!

"Those who use things of the world should not become attached to them. For the world in its present form is passing away."

1 Corinthians 7:31

Star Trek

- Communicator - Motorola Flip Phone
- Phaser - EPM or Stun Gun / Taser
- Tablets (Medical) - Tablet Computers
- Tricorders - Many Medical Devices (below)
- Translators - Google Translate (others)
- Tractor Beam - MIT has it in concept
- Telepresence - Zoom Virtual Conferences
- Geordi's Visor - Robotic Eyes / Implants
- Communicator Badges - Many Security Badges
- Food Replicator - 3D Printer
- Holodeck - VR
- Teleportation - Quantum Entanglement
- Big Screen TV - Everyone has it
- Hands Free Phone - Bluetooth headset / Airpods

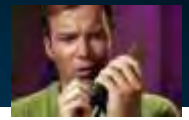

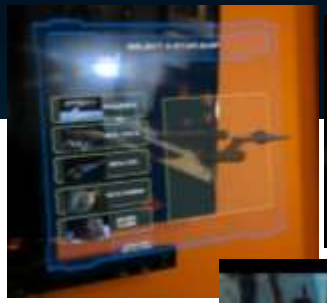

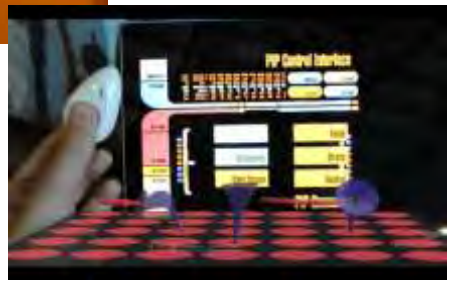

Hololens VR Apps Above

Microsoft working on Holodeck

MouthLab

Final Thoughts… Catch your Ride!

"Things may come to those who wait, but only the things left by those who hustle."

— Abraham Lincoln

The Digital Transformation Ahead

Digital Transformation 2000 to 2050

A historically significant change in humanity...

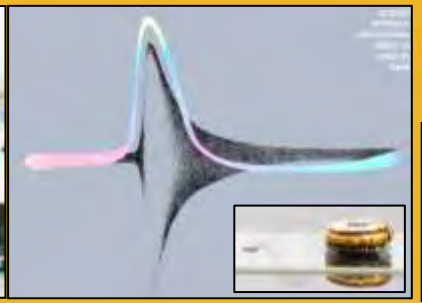

Gerd

Disconnected *The Dark Ages* | Using Digital | Wearing Digital | Implanting Digital | The Hive Mind

ADB Gives You Time Enough at Last

Summary – Prepare now for the Future!

- Innovation, Autonomous & the Cloud
- A Robot May Not Look Like One…
- Autonomous Transaction Processing (ATP)
- Autonomous Data Warehouse (ADW)
- Machine Learning & Data Visualization Desktop
- Next: **Robots & the Future Ahead**

"We make a Living by what we get. We make a Life by what we give."

-Sir Winston Churchill

Conf42: DevOps 2024

January 25 2024 - premiere 5PM GMT

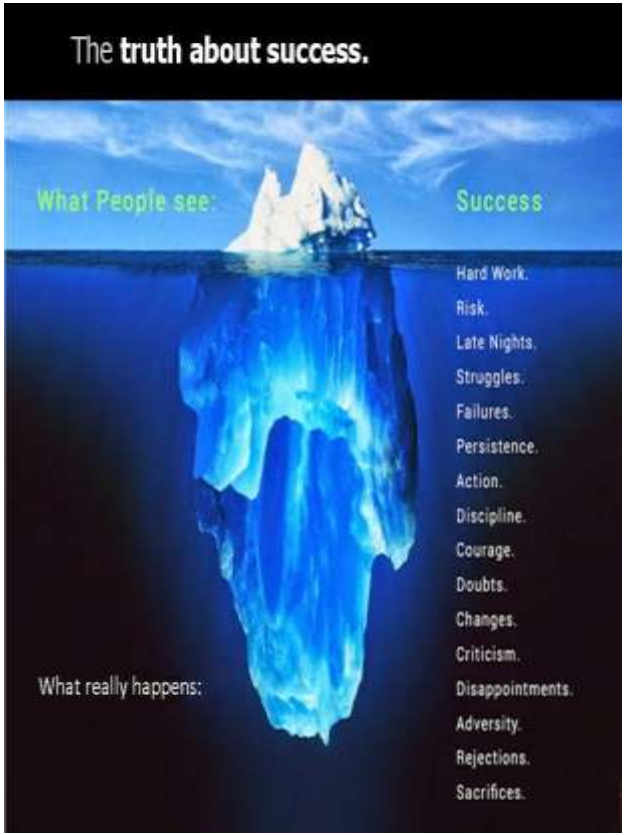

Smartsalessolutions.net

For More Information

- *Oracle 12c Release 2 Performance Tuning Tips & Techniques; Richard J. Niemiec; Oracle Press (Available now – 3/14/2017)*

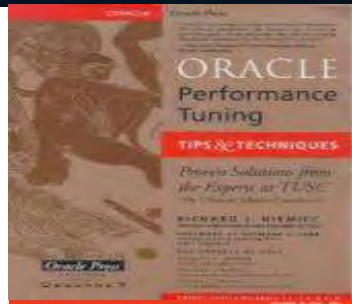

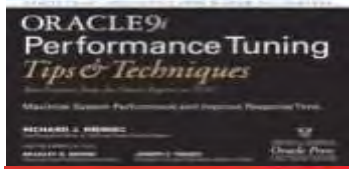

"If you are going through hell, keep going" - Churchill

12c R2 Book – Available Now!

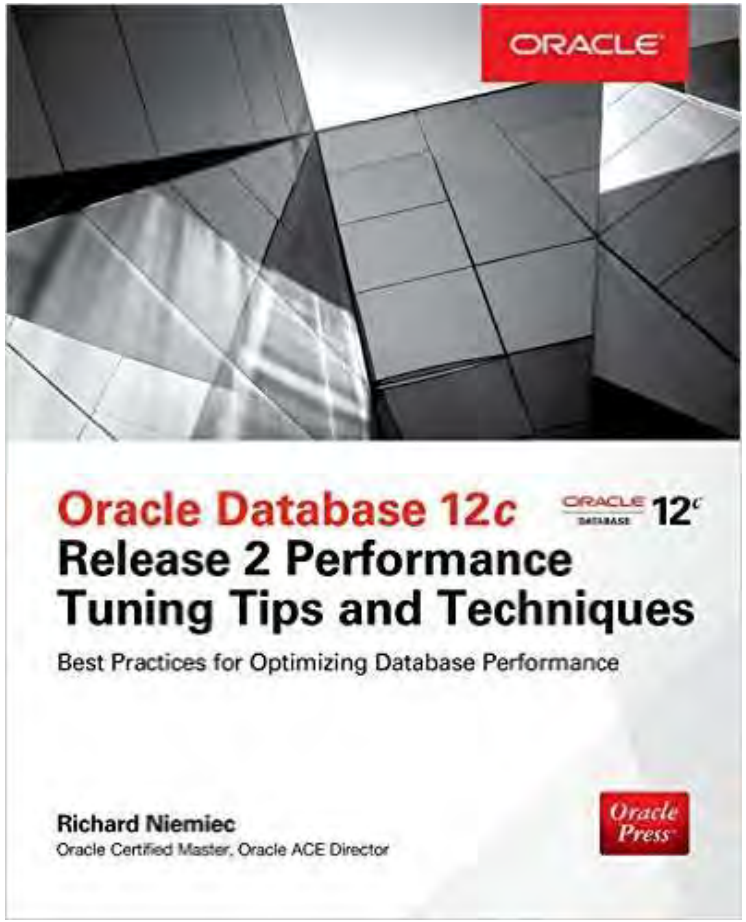

Top New Release

References

- The Emerging Technology Roadmap, Scott Klososky
- Futurist Gerd Leonhard The Futures Agency, Technology vs. Humanity, Gerd Leonhard, oracle.com & Juan Loaiza presentations, amazon.com, smartcitiescouncil.com, youtube.com, business coach, libelium.com, monetate, en.wikipedia.org, Netflix, Black Mirror, cooking-hacks.com, hometoys.com, FPOV, huffingtonpost.com, thegardian.com, nationalgeographic.com, newscientist.com, enswmu.blogspot.com, dailymail.co.uk, FutureRobot, theguardian.com, thinkhealthwireless.blogspot.com, ge.com, cmswire.com, runningsupplement.co.uk, quickmeme.com, nike.com, thisiswhyimbroke.com, businessinsider.com, slideshare.com, forrester.com, spiceworks.com, mwaintel.com, humancapitalist.com, wired.com and Wired Magazine, shodanhq.com, developer.nokia.com, extremetech.com, Getty Images, & any other company products are the property of their respective companies.

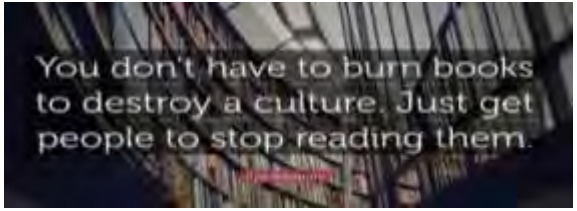

Copyright Information

- Neither the Viscosity nor the author guarantee this document to be error-free. Please provide comments/questions to rich.niemiec@viscosityna.com & richniemiec@gmail.com. I'm always looking to improve!

- Rich Niemiec ©2022. This document cannot be reproduced without expressed written consent from Rich Niemiec, but may be reproduced or copied for presentation and conference use.

- References include Rich Niemiec's Exadata Presentation & Oracle 12cR2 Database Performance Tuning Tips & Techniques book, www.oracle.com, en.wikipedia.org, slashgear.com, gifsoup.com, www.amazon.com, Tech Crunch, www.rolta.com, The Twilight Zone, Information Week, Gartner, Computerworld, & Oracle OpenWorld

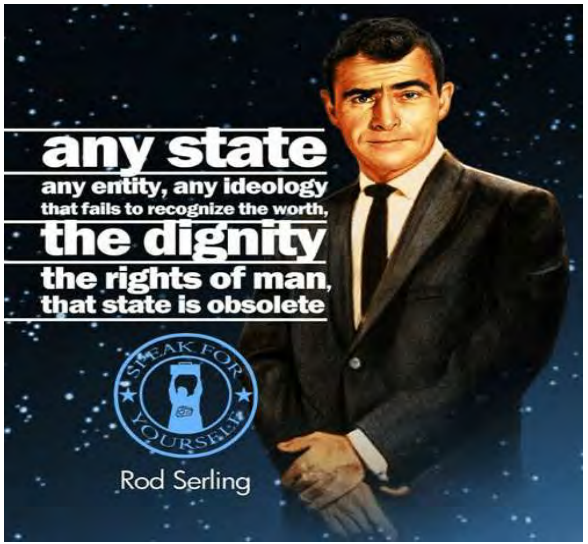

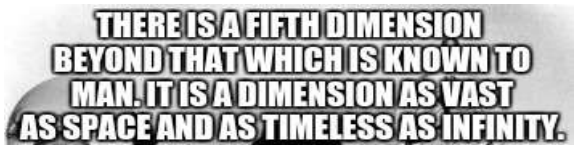

Contact Information: richniemiec@gmail.com

Rich's Overview @richniemiec

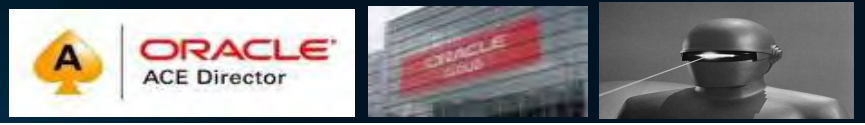

- Chief Innovation Officer, Viscosity North America
- Board Member – TEC, Entrigna, Ask DB Experts
- Former President of TUSC
  - Inc. 500 Company (Fastest Growing 500 Private Companies)
  - 10 Offices in the United States (U.S.); Based in Chicago
  - Oracle Advantage Partner in Tech & Applications
- Former President Rolta TUSC & President Rolta EICT International
- Author (5 Oracle Best Sellers – #1 Oracle Tuning Book for over a Decade):
  - Oracle Performing Tips & Techniques (Covers Oracle7 & 8i)
  - Oracle9i Performance Tips & Techniques
  - Oracle Database 10g Performance Tips & Techniques
  - Oracle Database 11g Performance Tips & Techniques
  - Quick Start Guide to Oracle Query Tuning (**2015**)
  - Oracle Database 12c Performance Tips & Techniques
- Former President of the International Oracle Users Group
- IOUG Top Speaker in 1991, 1994, 1997, 2001, 2006, 2007
- MOUG Current President & Top Speaker Twelve Times
- National Trio Achiever award - 2006
- Oracle Certified Master & Oracle Ace Director
- Purdue Outstanding Electrical & Computer and Engineer – 2007
- Honorary Senior Technical Advisor to Oracle China - 2014
- Chris Wooldridge Award – 1998, 2012
- Chicago Entrepreneur Hall of Fame - 1998
- E&Y Entrepreneur of Year & National Hall of Fame - 2001

Join Thousands Of Happy Customers And End The Frustration In Tuning Oracle Databases

Our services give you the training you need so your Oracle database runs faster and more efficiently. And, you get the credit!

MEMBERSHIPS

TRAINING

FREE TOOLS

BOOKS

Follow Us Online!


Facebook.com/ViscosityNA


Linkedin.com/company/Viscosity-North-America


@ViscosityNA

Viscosity North America


Facebook.com/ViscosityNA


@Viscosity_NA

Quick FREE notes

Send email to (for slides):
hello@viscosityna.com
(richniemiec@gmail.com)

@richniemiec - twitter

Conf42: DevOps 2024
January 25 2024 - premiere 5PM GMT