

ORACLE

Data Management Strategy

Çetin Özbütün

Senior Vice President

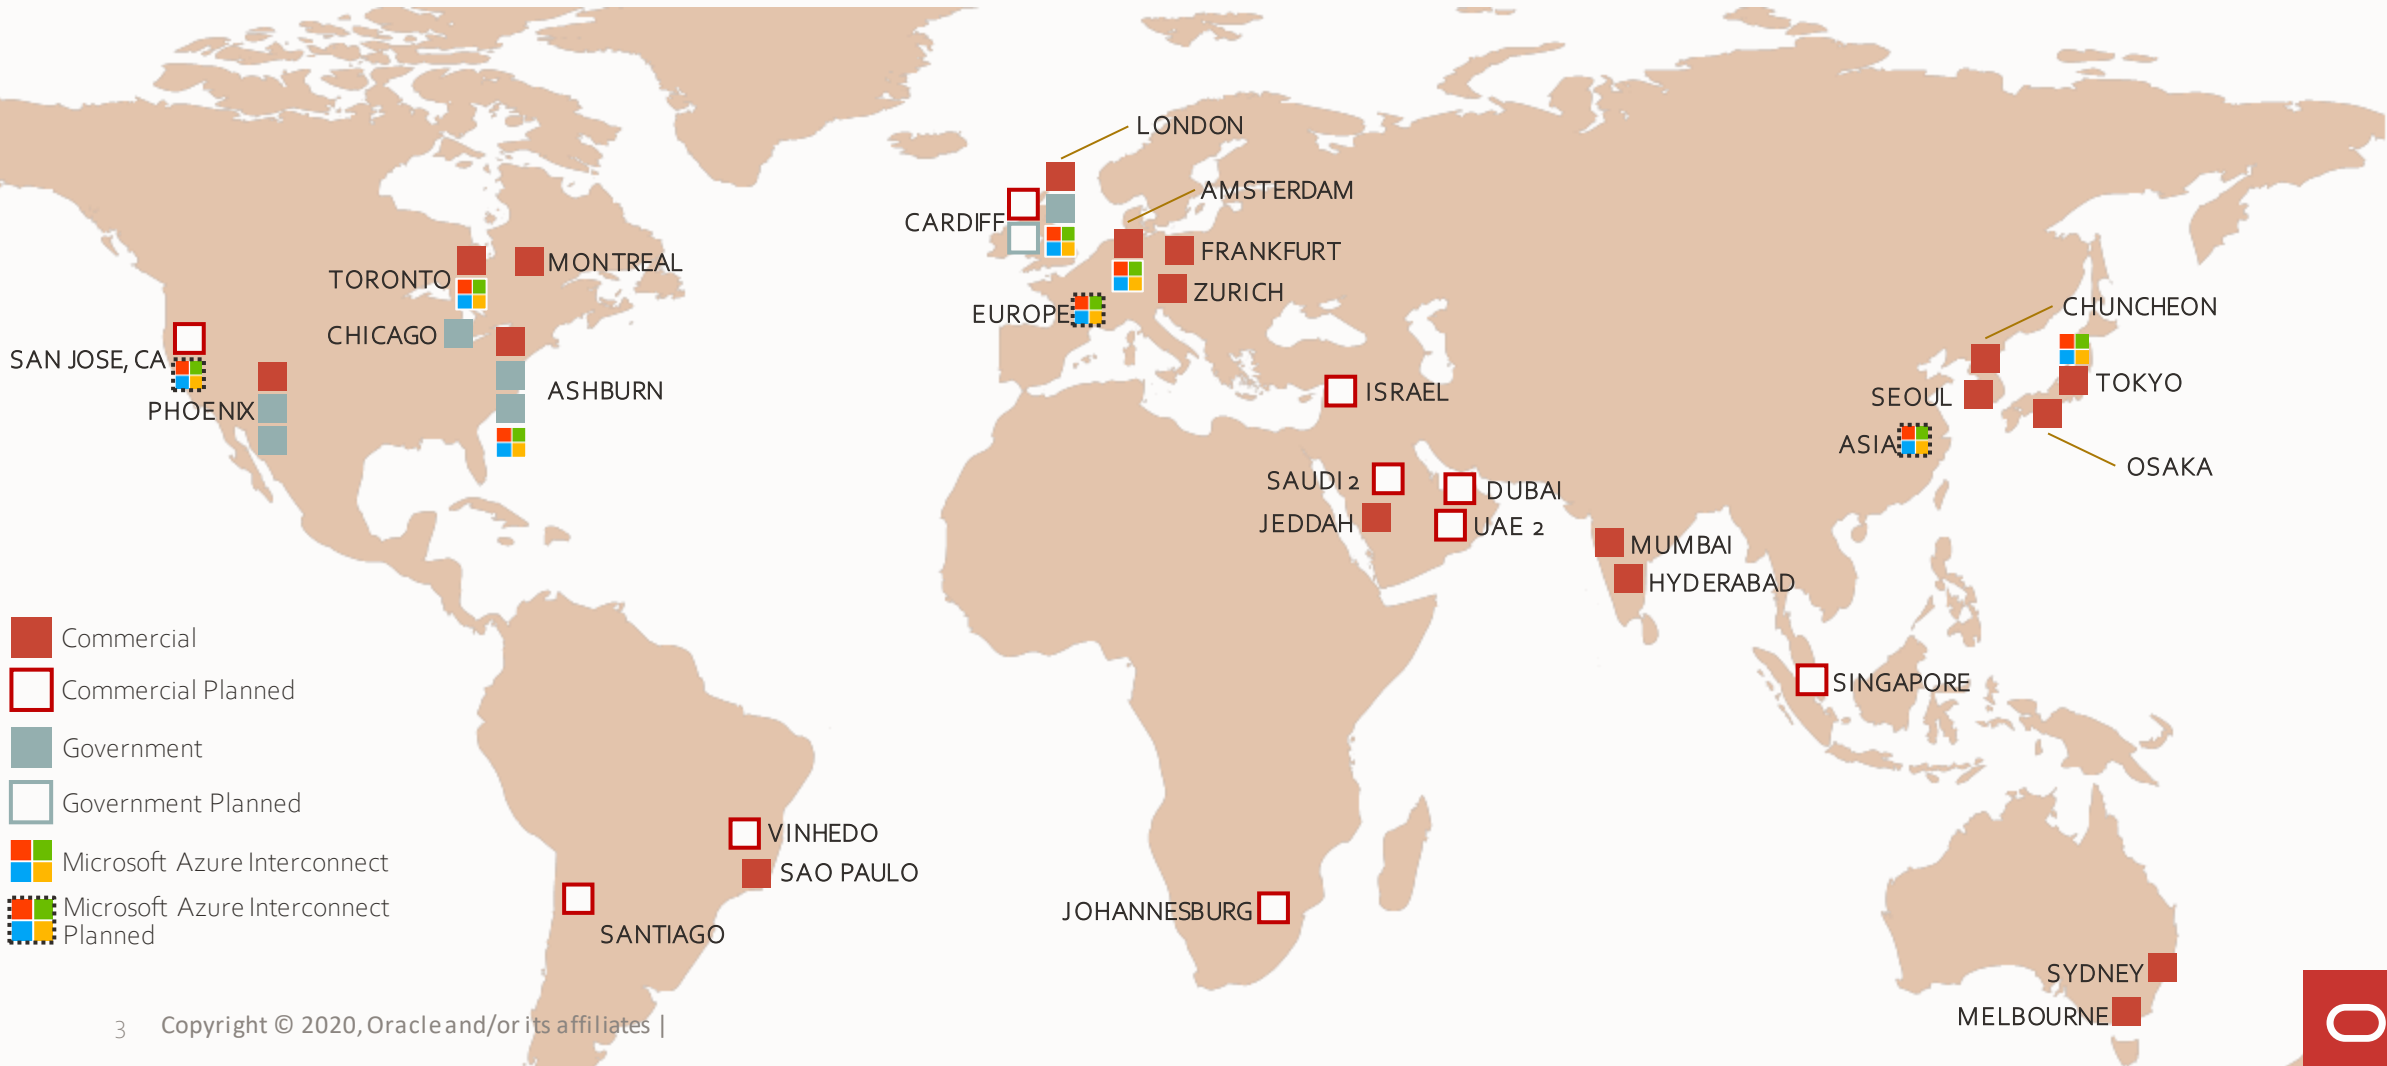
Oracle Database Server Technology

Safe harbor statement

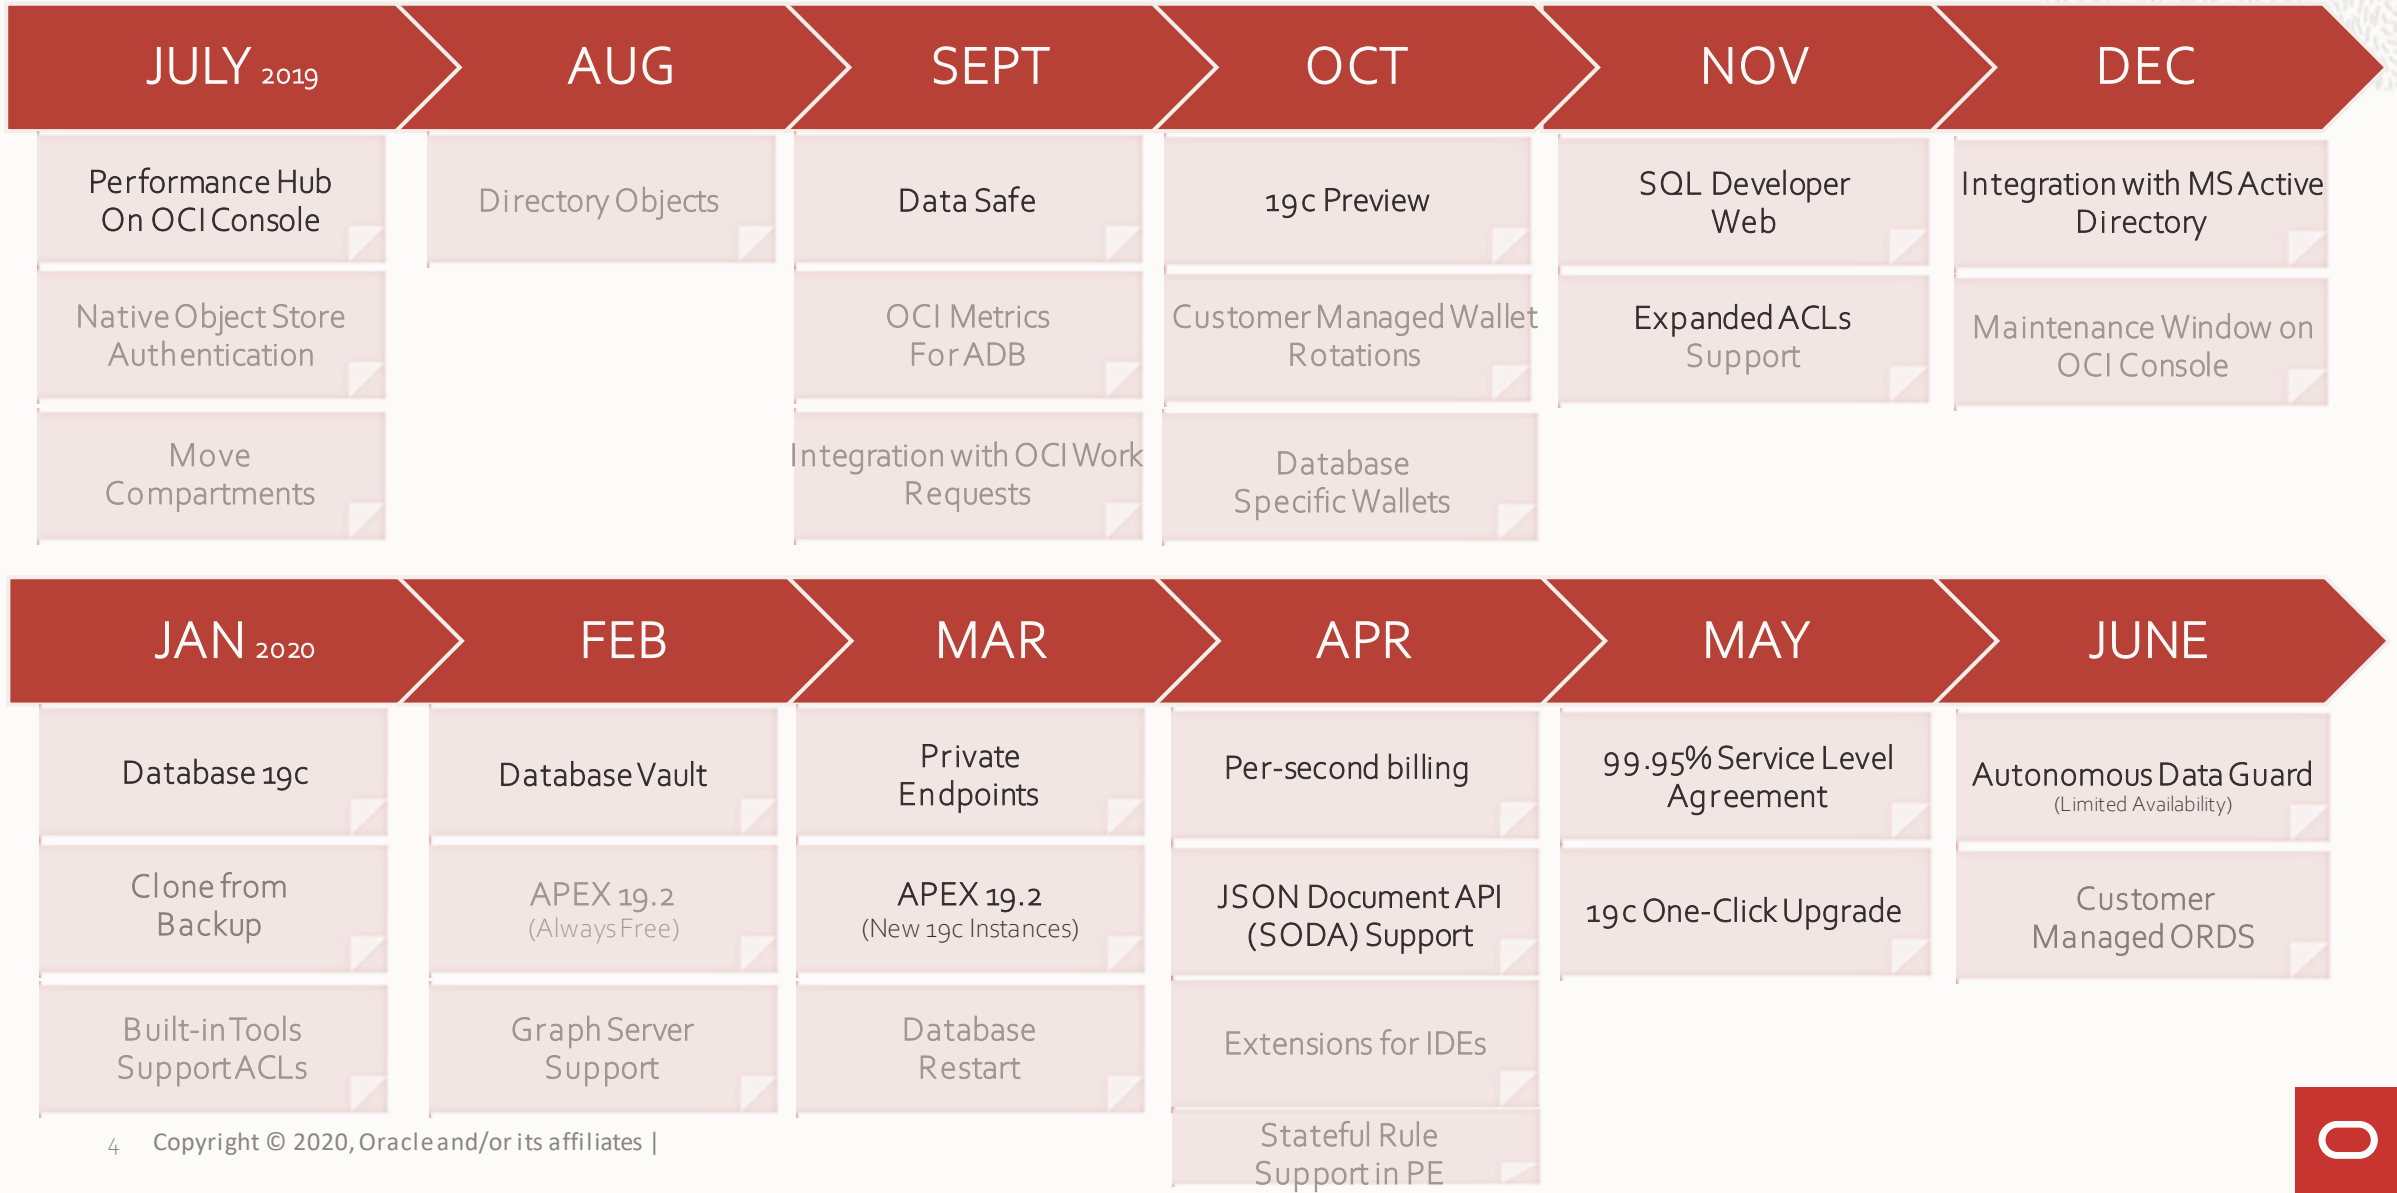
The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

Oracle Cloud Infrastructure Global Footprint

June 2020: 24 Regions Live, 12 Planned




Autonomous Database: 12-Months of Key New Features



Coming Soon

Autonomous Database on Exadata Cloud@Customer

- Running on Gen2 Exadata Cloud@Customer
- Autonomous Database in your data center
 - Autonomously managed via Oracle Cloud Infrastructure (OCI) control plane
 - Secure websockets tunnel between on-prem firewall and OCI firewall
 - On-prem Exadata hardware, software, interfaces
 - Same as cloud pay-per-use subscription model



Extend Autonomous Database to help more people get insights into their data

Autonomous Database automates almost all database administration

Autonomous Database does not automate tasks for Data Engineers, Data Analysts, or Data Scientists

Goal:

- Help Data Analysts and Data Scientists to use Autonomous Database to more easily gain insights into their data

Solution:

- Extend Autonomous Database for:
 - Data ingestion and transformations
 - Business modelling and analysis
 - Machine learning and automatic insights



New Autonomous Database Architecture

Data Sources

ORACLE®
FUSION APPLICATIONS

ORACLE SIEBEL ORACLE PEOPLESOFT ORACLE E-BUSINESS SUITE

SAP salesforce ORACLE NETSUITE

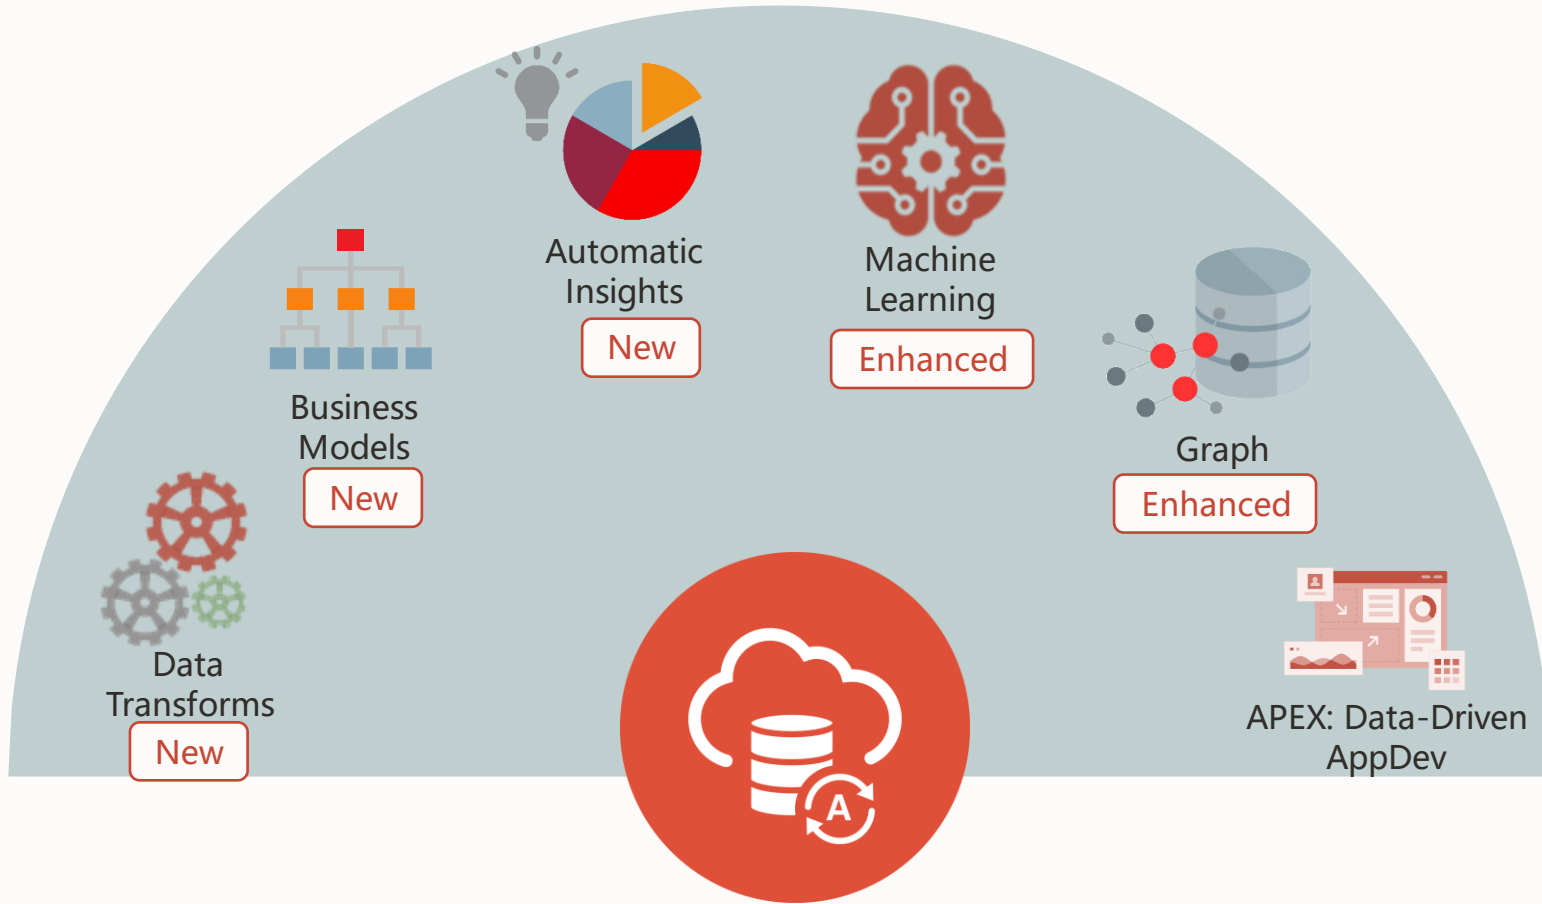
Applications

Databases: NoSQL

Files: X, {...}, CSV

Streaming Services: GoldenGate, kafka

Data Lake: OCI Object Storage, Amazon S3, Microsoft Azure Blob Storage



Visualization & Development

Application Express

Oracle ML Notebook

Oracle Analytics Cloud

IDEs

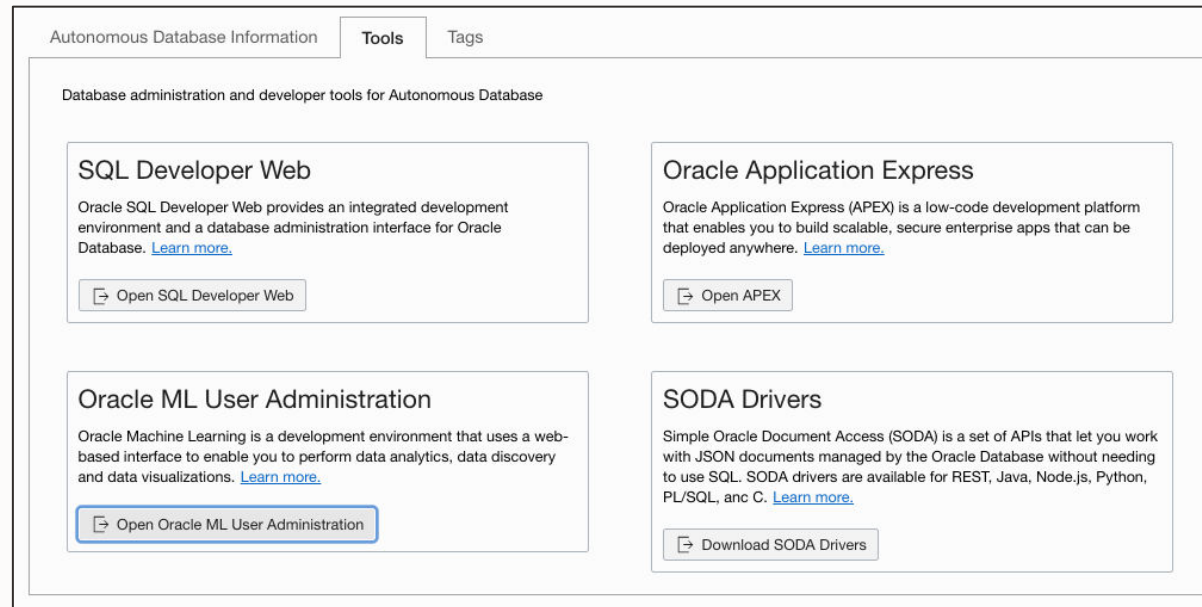
Spreadsheets

3rd Party BI Tools



Autonomous Database today

A handful of integrated tools

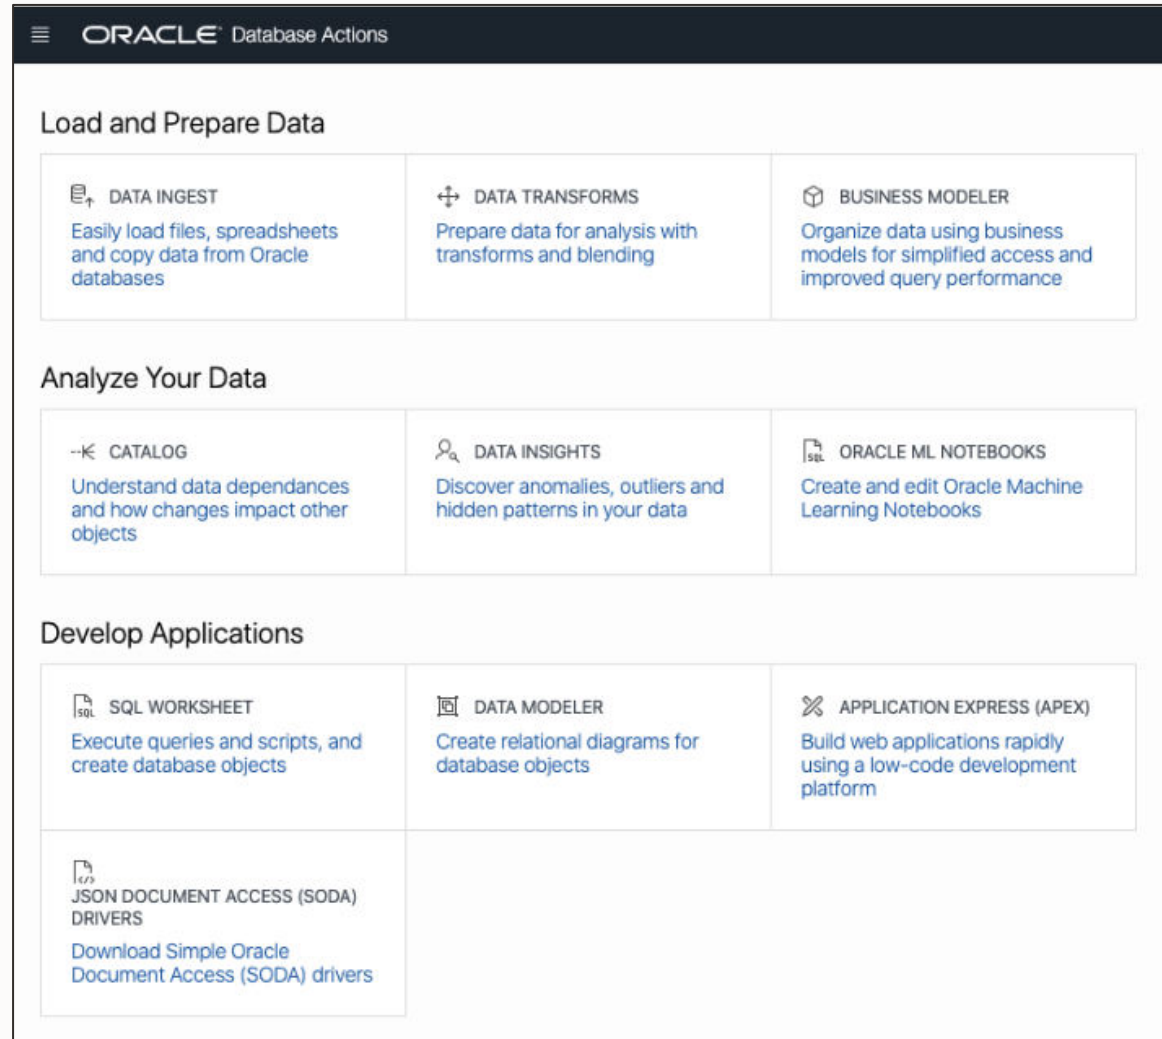


The screenshot shows a web interface for Oracle Autonomous Database tools. At the top, there are three tabs: 'Autonomous Database Information', 'Tools' (which is selected), and 'Tags'. Below the tabs, the main heading reads 'Database administration and developer tools for Autonomous Database'. The page is divided into four tool cards, each with a title, a brief description, and an 'Open' button with an external link icon.

- SQL Developer Web**: Oracle SQL Developer Web provides an integrated development environment and a database administration interface for Oracle Database. [Learn more.](#) Button: Open SQL Developer Web
- Oracle Application Express**: Oracle Application Express (APEX) is a low-code development platform that enables you to build scalable, secure enterprise apps that can be deployed anywhere. [Learn more.](#) Button: Open APEX
- Oracle ML User Administration**: Oracle Machine Learning is a development environment that uses a web-based interface to enable you to perform data analytics, data discovery and data visualizations. [Learn more.](#) Button: Open Oracle ML User Administration
- SODA Drivers**: Simple Oracle Document Access (SODA) is a set of APIs that let you work with JSON documents managed by the Oracle Database without needing to use SQL. SODA drivers are available for REST, Java, Node.js, Python, PL/SQL, and C. [Learn more.](#) Button: Download SODA Drivers

Autonomous Database – coming soon

Broad set of integrated tools and utilities



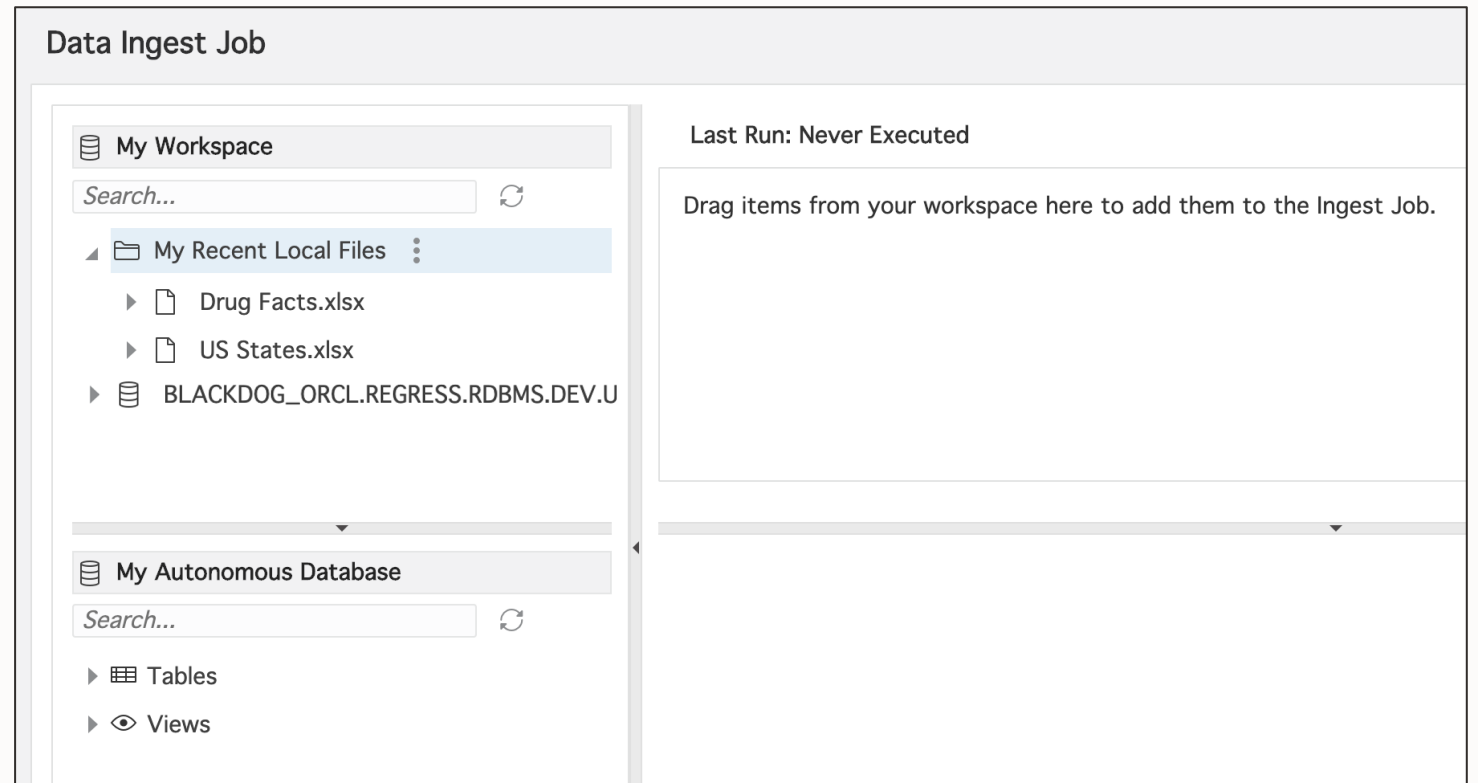
The screenshot displays the Oracle Database Actions interface, which is organized into three main sections:

- Load and Prepare Data**
 - DATA INGEST**: Easily load files, spreadsheets and copy data from Oracle databases.
 - DATA TRANSFORMS**: Prepare data for analysis with transforms and blending.
 - BUSINESS MODELER**: Organize data using business models for simplified access and improved query performance.
- Analyze Your Data**
 - CATALOG**: Understand data dependences and how changes impact other objects.
 - DATA INSIGHTS**: Discover anomalies, outliers and hidden patterns in your data.
 - ORACLE ML NOTEBOOKS**: Create and edit Oracle Machine Learning Notebooks.
- Develop Applications**
 - SQL WORKSHEET**: Execute queries and scripts, and create database objects.
 - DATA MODELER**: Create relational diagrams for database objects.
 - APPLICATION EXPRESS (APEX)**: Build web applications rapidly using a low-code development platform.
 - JSON DOCUMENT ACCESS (SODA) DRIVERS**: Download Simple Oracle Document Access (SODA) drivers.

Data Ingest

Simple 'Drag and Drop' Data Loading

- Files on local computer
- Files in Object Storage (incl AWS S3, Azure Blob Storage)
- Oracle Databases (on-prem and cloud)



Data Transforms

Based upon Oracle Data Integrator

- All ODI connectors (incl. Fusion, SFDC)
- Simple migration to ADB for ODI customers

Declarative, no-code development

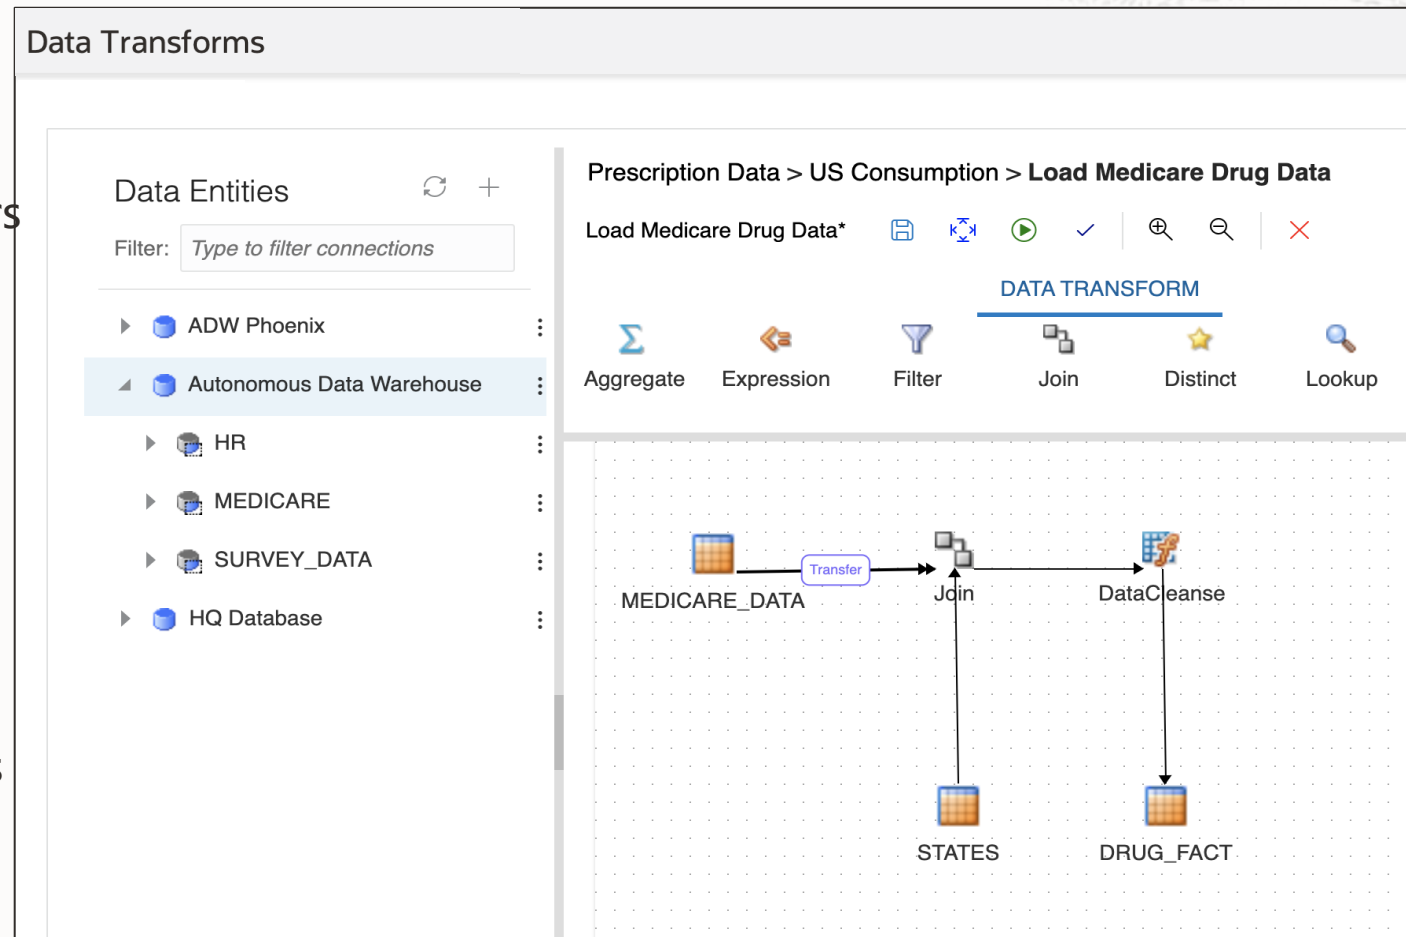
- New, easy-to-use cloud UI
- 'Drag and Drop' to create Maps

Rich set of Operators

- Transform, Quality, Analytic, Spatial, ML
- All DB Operators

Autonomous

- Discover relationships, recommend actions
- Auto code generation



Business Models (Analytic Views)

Analytic Views

- Enables high-performance multidimensional analysis over relational data
- Essbase, ERP, EPM and RPAS teams have reviewed their usage of Analytic Views with you

Automatic creation

- Business Models embedded in ADW

Automatic multidimensional cache

- For both local and remote data

Use with any analytics tool

- Standard SQL queries
- No application changes required

The screenshot shows the Oracle Database Actions - Analytic View Designer interface. The main title is "Create Business Model". The interface is divided into sections for configuration:

- Add Details:** Name: Health Insurance Coverage Data; Caption: Yearly Heath Insurance Coverage D;
- Add Dimensions:** Description: Insured populations by year, state/i
- Set Options:** Schema: ADPTEST; Show recommended fact tables only; Fact Table: SAHIE_FACT

Below the configuration, there are tabs for "Preview" and "Statistics". The "Preview" tab is active, showing a table with the following data:

	YEAR	COUNTY_FIPS_CODE	AGE_CODE	GENDER_CODE
1	2008	01001	1	
2	2008	01001	1	
3	2008	01001	1	
4	2008	01001	1	
5	2008	01001	1	
6	2008	01001	1	
7	2008	01001	1	
8	2008	01001	1	
9	2008	01001	1	

At the bottom left, there are "Next" and "Cancel" buttons.



Machine Learning

Extensive collection of in-database ML algorithms

Built-in Notebook for Data Scientists

- Based on Apache Zeppelin

New: support for Python and R

- Use Python, R, and SQL side-by-side
- Leverage database-native algorithms from Python and R

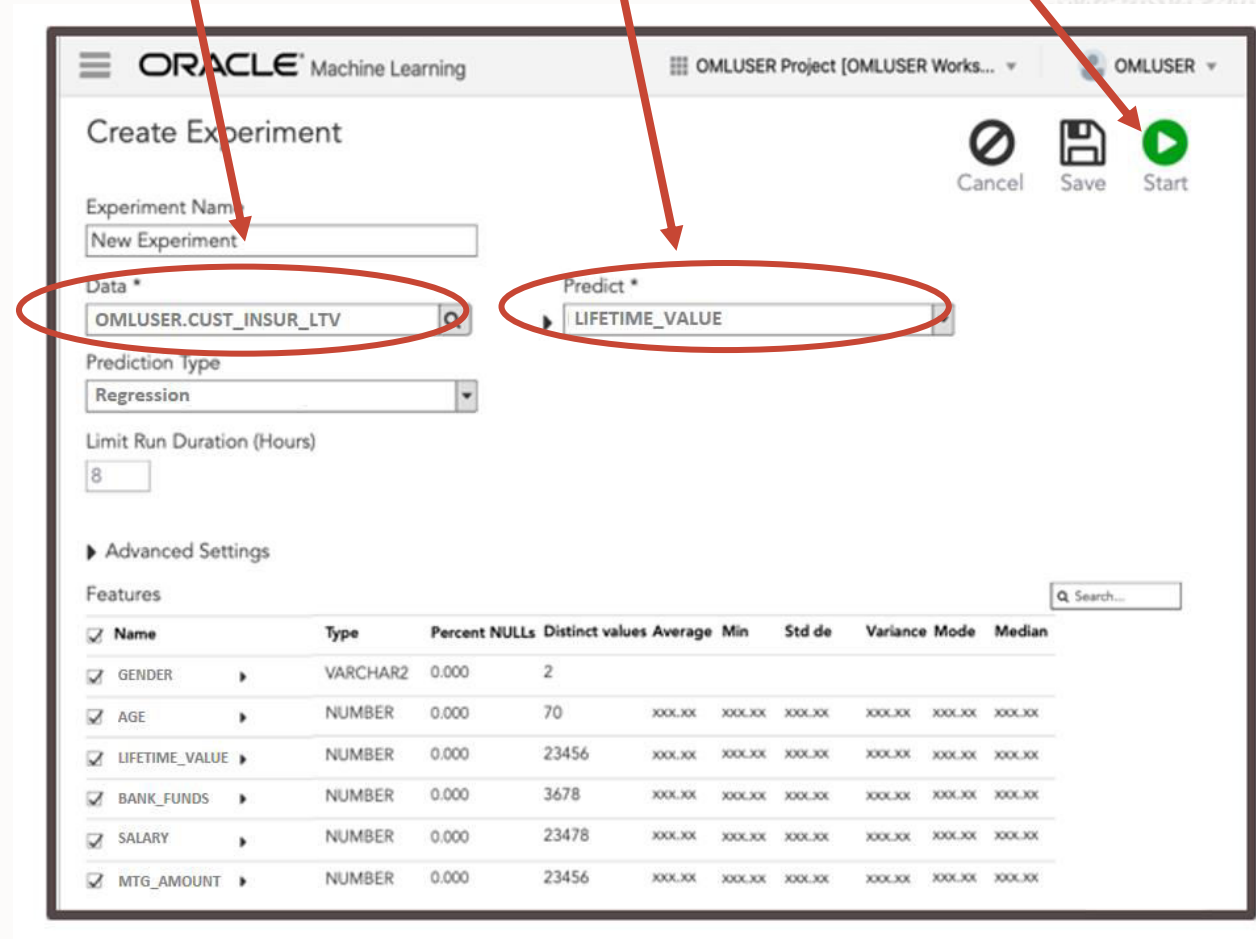
New: automatic machine learning (AutoML)

- Declarative user interface for creating, managing, and deploying ML models
- Based on Oracle Labs AutoML

1. Choose data set

2. Choose data for prediction

3. Press start



Machine Learning

Extensive collection of in-database ML algorithms

Built-in Notebook for Data Scientists

- Based on Apache Zeppelin

New: support for Python and R

- Use Python, R, and SQL side-by-side
- Leverage database-native algorithms from Python and R

New: automatic machine learning (AutoML)

- No-code user interface for creating, managing, and deploying ML models
- Based on Oracle Labs AutoML

Multiple algorithms compared and best is selected

The screenshot displays the Oracle Machine Learning interface for an 'AutoML Experiment Demo'. The 'Leader Board' section is circled in red, showing a table of results:

Name	Algorithm	Accuracy (default)
Random Forest 1	Random Forest	89
Neural Network 1	Neural Network	87
GLMR 1	Generalized Linear Model (Ridge Regression)	86
GLM 1	Generalized Linear Model	84
Decision Tree 1	Decision Tree	79

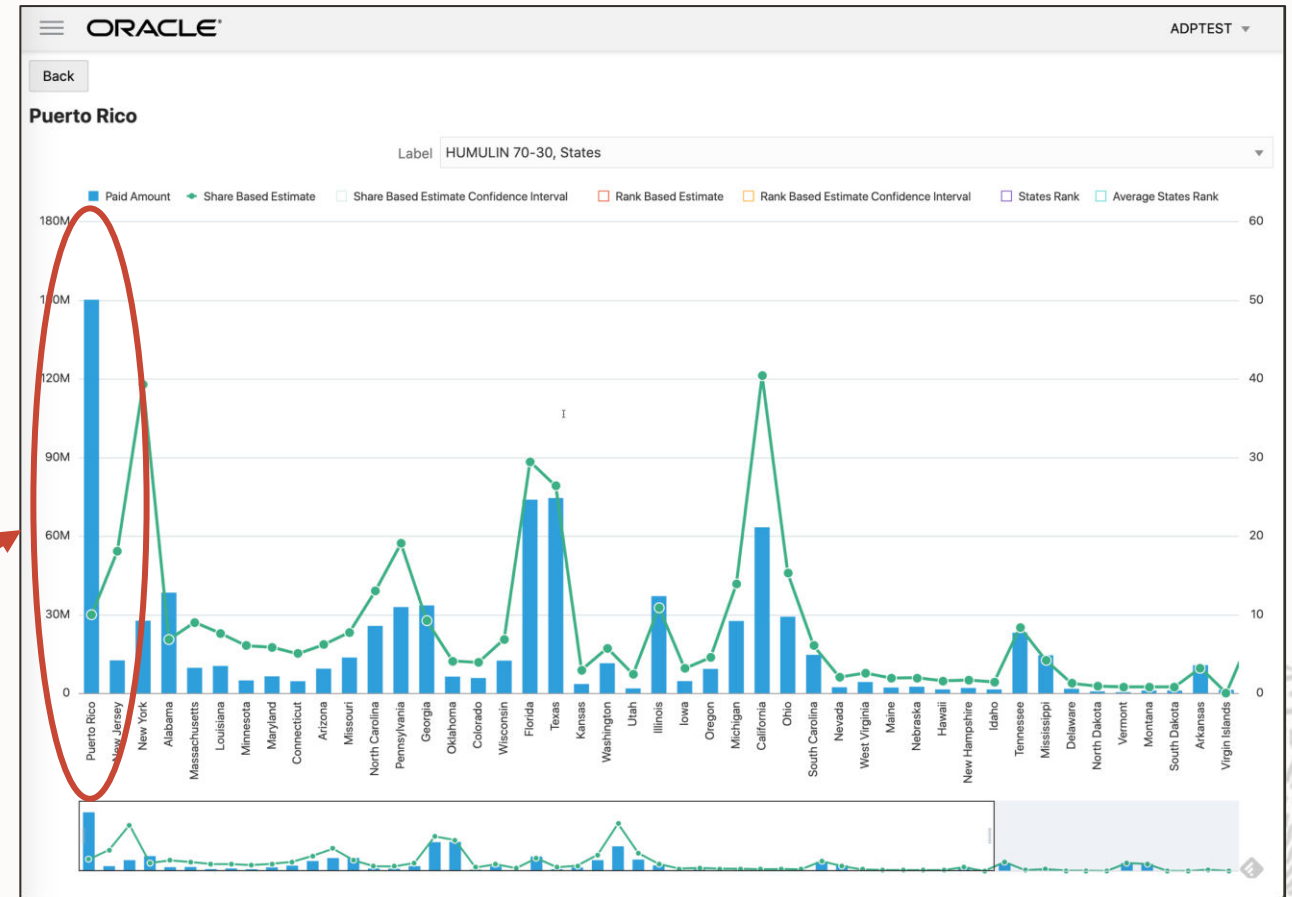
The 'Progress' sidebar on the right shows various stages of the experiment, including Feature Ranking, Algorithm Selection (Top Algorithms: NN, SF, SVM), Adaptive Sampling, Feature Selection, Hyperparameter Tuning, and Feature Prediction Impact.



Auto-Insights – Driven by Machine Learning

Automatic Insight Discovery

- Discovers **hidden patterns** and outliers
 - **Crawl over business model**, running as background process
 - **Variety of algorithms** including singular value decomposition
- Example:**
- Puerto Rico payments for a given drug are far higher than expected. Needs to be investigated for fraud or improper prescriptions.



Graph

Graph Studio UI

- Graph Modeler
- Notebooks and Visualization (based on Oracle Labs Data Studio)

Graph API's

- Property Graph Query Language (PGQL)
 - Declaration graph query language
- Graph Analytics
 - Based on Oracle Labs PGX
- RDF Graphs: SPARQL and REST

SQL for Reporting and Analytics

The screenshot displays the Oracle Graph Studio interface. At the top, the title bar reads "ORACLE Graph Studio" with a "PRODUCT" dropdown menu. Below the title bar, the breadcrumb navigation shows "Use Cases > My Sales Analysis". The main workspace is divided into three sections:

- Code Editor (Top):** Contains a Java snippet for creating an analyst:

```
var new_analyst = session.createAnalyst()
sales
```
- Query Editor (Middle):** Contains a PGQL query:

```
%pgql
select e, e2 from match (s:customer) --[e] -> (p:product) <- [e2] - (c2:customer) on REC4_HVI
where s.cust_id = 3221
limit 500;
```
- Graph Visualization (Bottom):** Shows a network graph with nodes and edges. A red box highlights a specific node with its properties:

id: 1227004981436686681	
id	1227004981436686681
CUST_FIRST_NAME	Randall
CUST_GENDER	M
CUST_ID	2048.0
CUST_LAST_NAME	Oakfield
CUST_YEAR_OF_BIR	1944.0
TH	
PROD_ID	0.0
PROD_LIST_PRICE	0.0
label	CUSTOMER

At the bottom of the interface, there is a Java snippet for calculating PageRank:

```
%pgx-java
var usales = sales.undirect("usales")
VertexProperty<Integer, Double> pagerank = usales.getOrCreateVertexProperty(PropertyType.DOUBLE, "pagerank")
new_analyst.pagerank(usales)
```



Data Lineage and Impact Analysis

Built-in metadata repository

- Shared across all components
- Future integration with OCI Data Catalog

UI for data lineage and impact analysis

- Visible to all users
- SQL and PL/SQL interfaces

The screenshot displays the Oracle Database Cloud Service Lineage Explorer interface. The main window shows the details for the entity **MEDICARE_AV**, which is an **ANALYTIC_VIEW**. The details include the namespace (DB), owner (ADPUSER1), application (DATABASE), and creation/updated timestamps. Below the details, the **Lineage** tab is active, showing a lineage graph. The graph illustrates the data flow from **MEDICARE_AV** to various data sources and dimensions, including **MEDICARE_DATA**, **MEDICARE_DRUGS**, **MEDICARE_DRUG_DIM**, **MEDICARE_DRUG_HIER**, **MEDICARE_STATES**, **MEDICARE_STATE_DIM**, **MEDICARE_STATE_HIER**, **MEDICARE_YEARS**, and **MEDICARE_YEAR_DIM**. The interface also includes a search bar, a browse catalog sidebar, and a status bar at the bottom indicating a successful REST call.





ORACLE