

ORACLE

Data Management Strategy

Çetin Özbütün

Senior Vice President

Oracle Database Server Technology

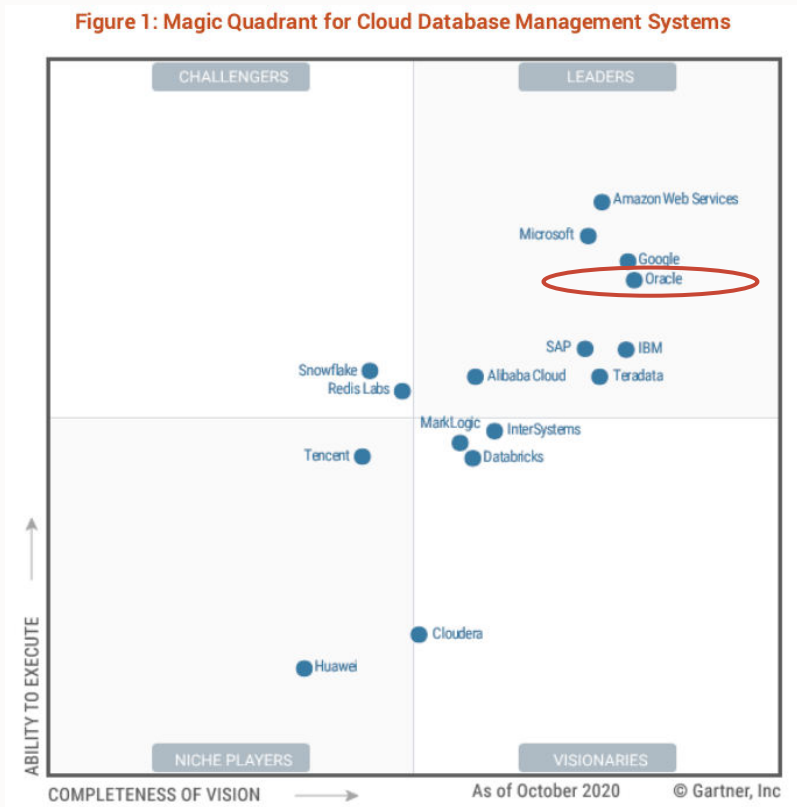
Oracle Cloud Infrastructure Global Footprint

November 2020: 28 Regions Live, 6 Azure Interconnect Regions



Analysts Agree: Oracle *Cloud* DB #1

Gartner: Cloud DBMS MQ



Oracle #1 "Vision"

Kuppinger: Cloud Enterprise DBs



Oracle #1 Overall

Contrasting Database Architecture Strategies

Amazon and Niche DB Vendors

Run **Single-Purpose Proprietary** database for each data type and workload



Amazon Aurora



Amazon DocumentDB



Amazon DynamoDB



Amazon Timestream



Amazon Neptune



Amazon Quantum Ledger Database




Amazon RedShift



Amazon ElastiCache

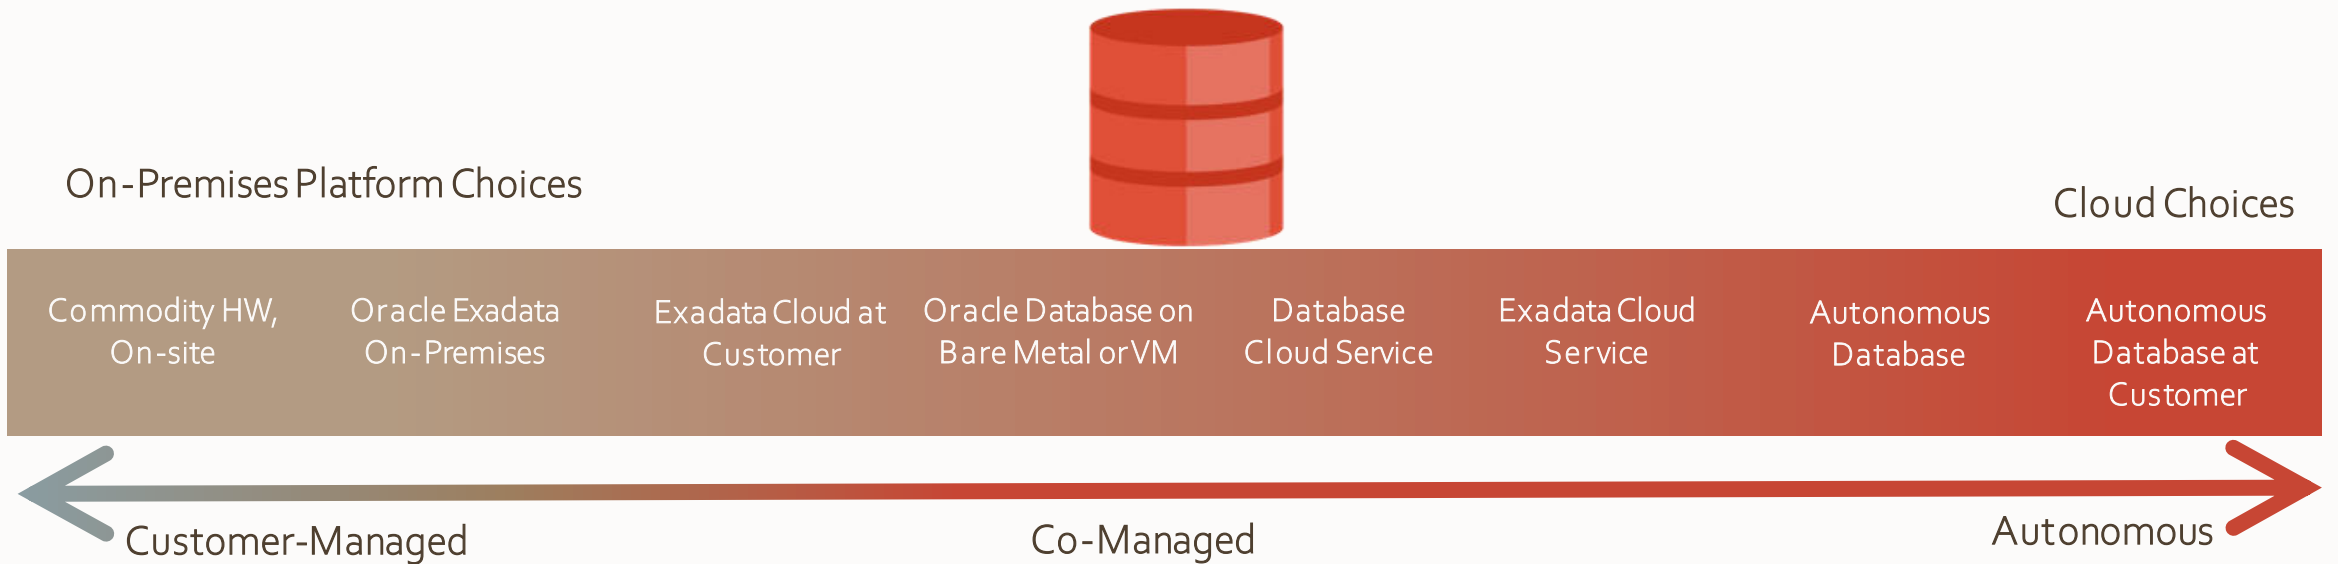
Oracle Strategy

Run **Converged, Open, Oracle** database for multiple data types and workloads




Oracle Database for the Hybrid Enterprise

Same database on-prem, in the cloud or hybrid configurations



Initial Vision of Autonomous Database

GOAL - Remove need for systems/operational admin



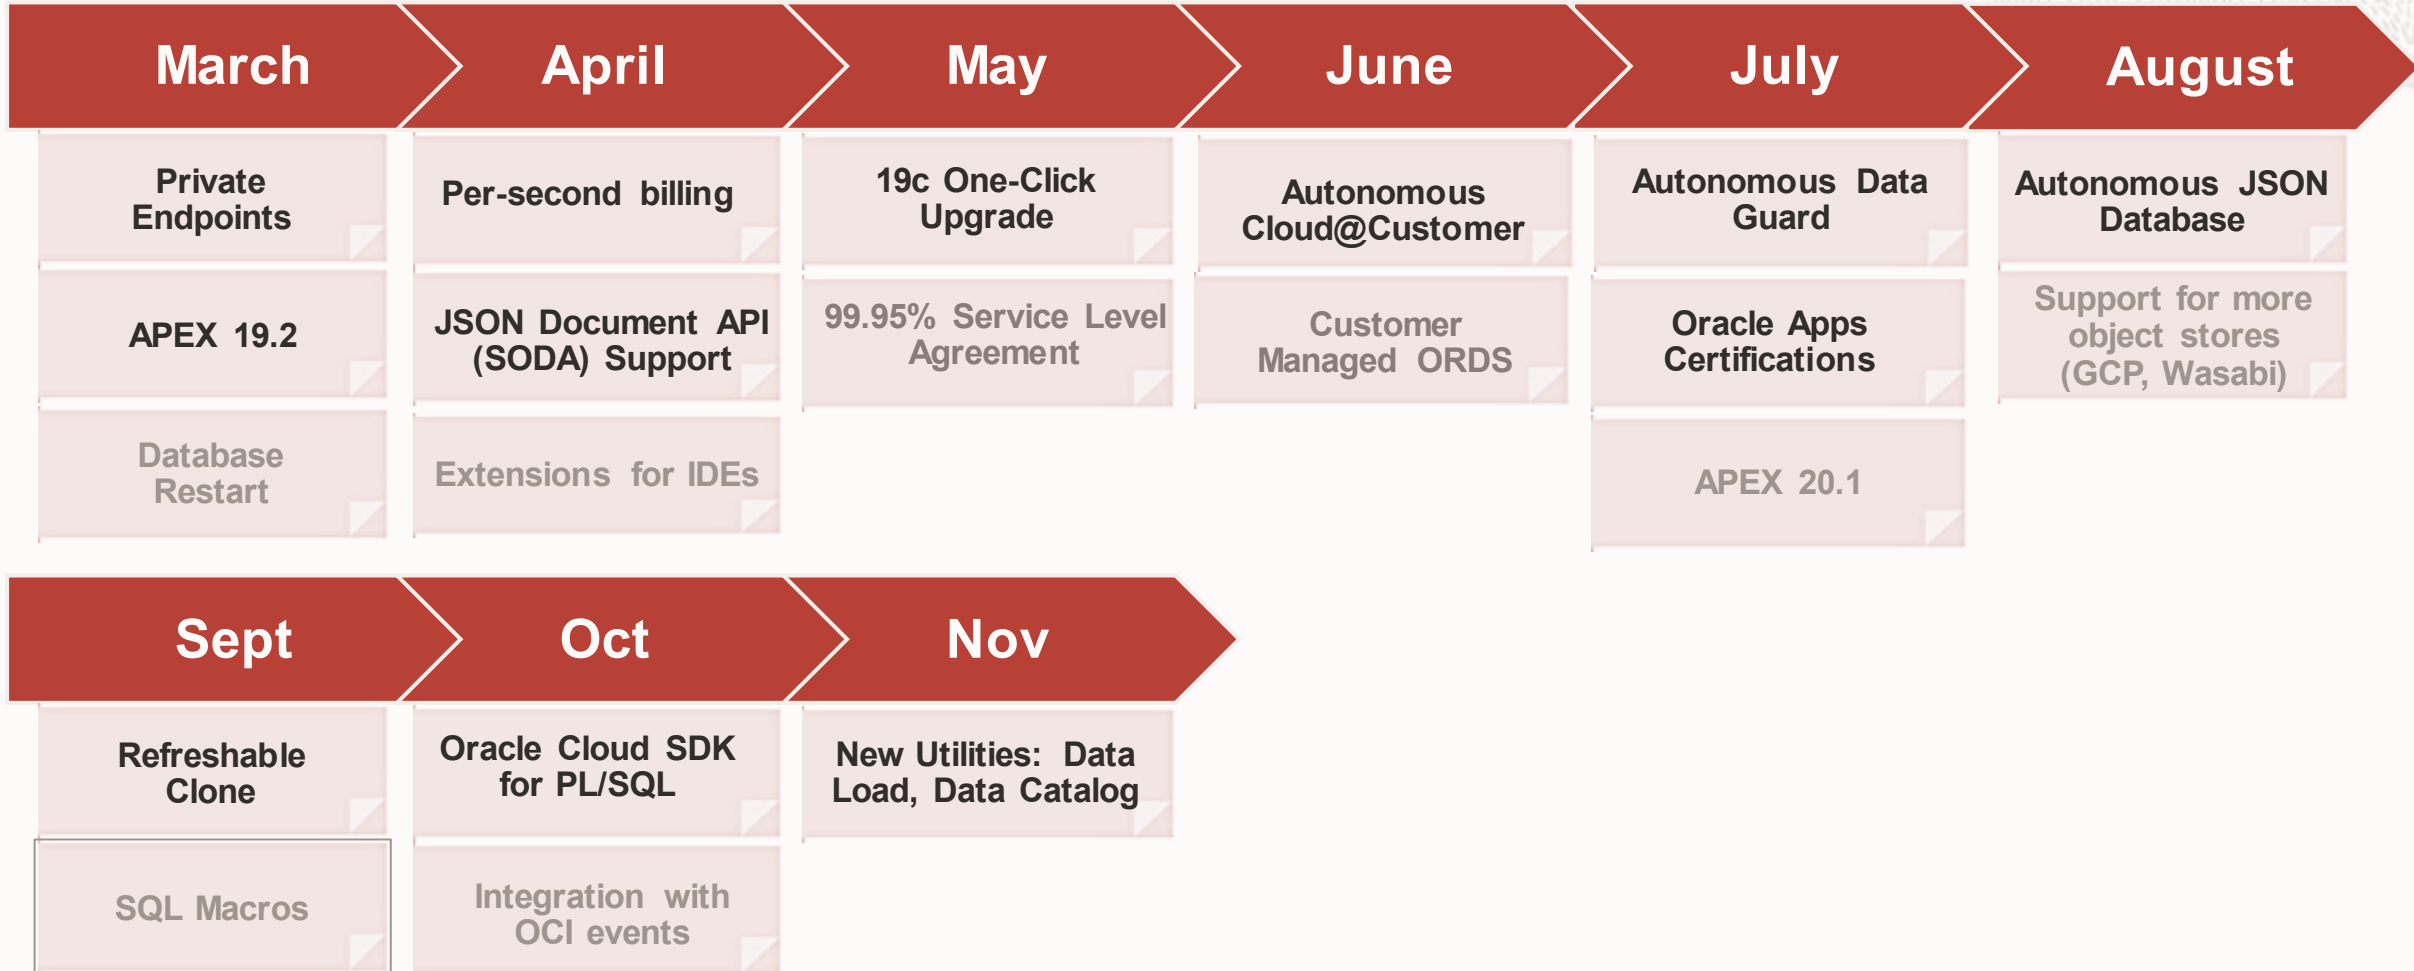
No human labor means
lower cost



No human error means
better reliability and **better security**



Autonomous Database: Recent New Features



Autonomous Database | Optimized by Workload



**Autonomous
Data Warehouse**

Data Warehouse,
Data Mart, Data Lake

**Autonomous
Transaction Processing**

Transactions,
Batch, Reporting

**Autonomous
JSON Database**

Fast, Easy, Low-Cost
JSON Data

New


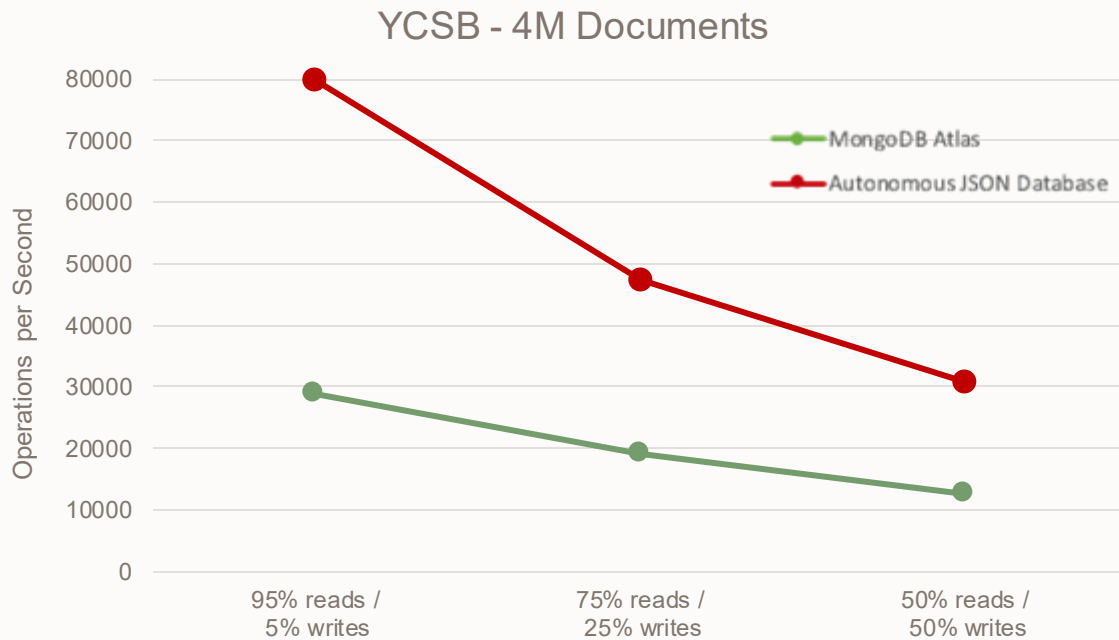
**APEX Application
Development Service**

Fast, Easy, Low-Cost
Low-Code AppDev

Coming soon



Autonomous JSON Database Faster than MongoDB



Autonomous JSON Database with 8 OCPUs compared to MongoDB Atlas on M60
Industry-standard Yahoo Cloud Serving Benchmark (YCSB)

Source of MongoDB results: <https://www.mongodb.com/atlas-vs-amazon-documentdb/performance> as of 8/12/2020

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

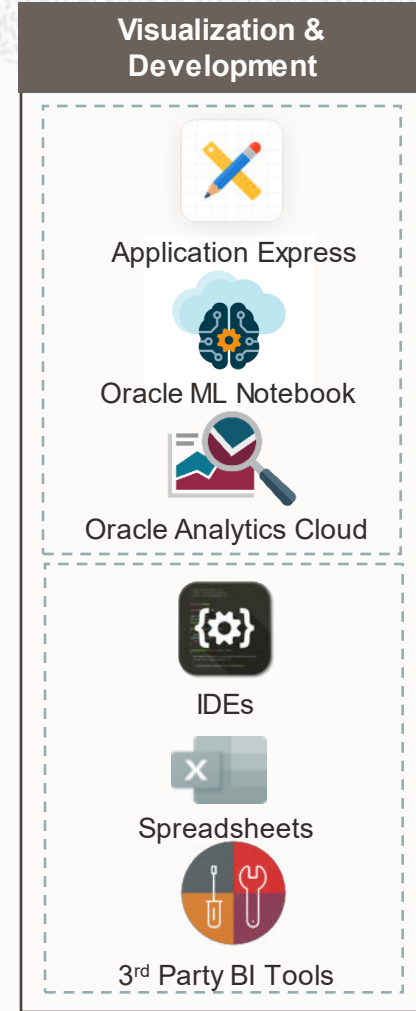
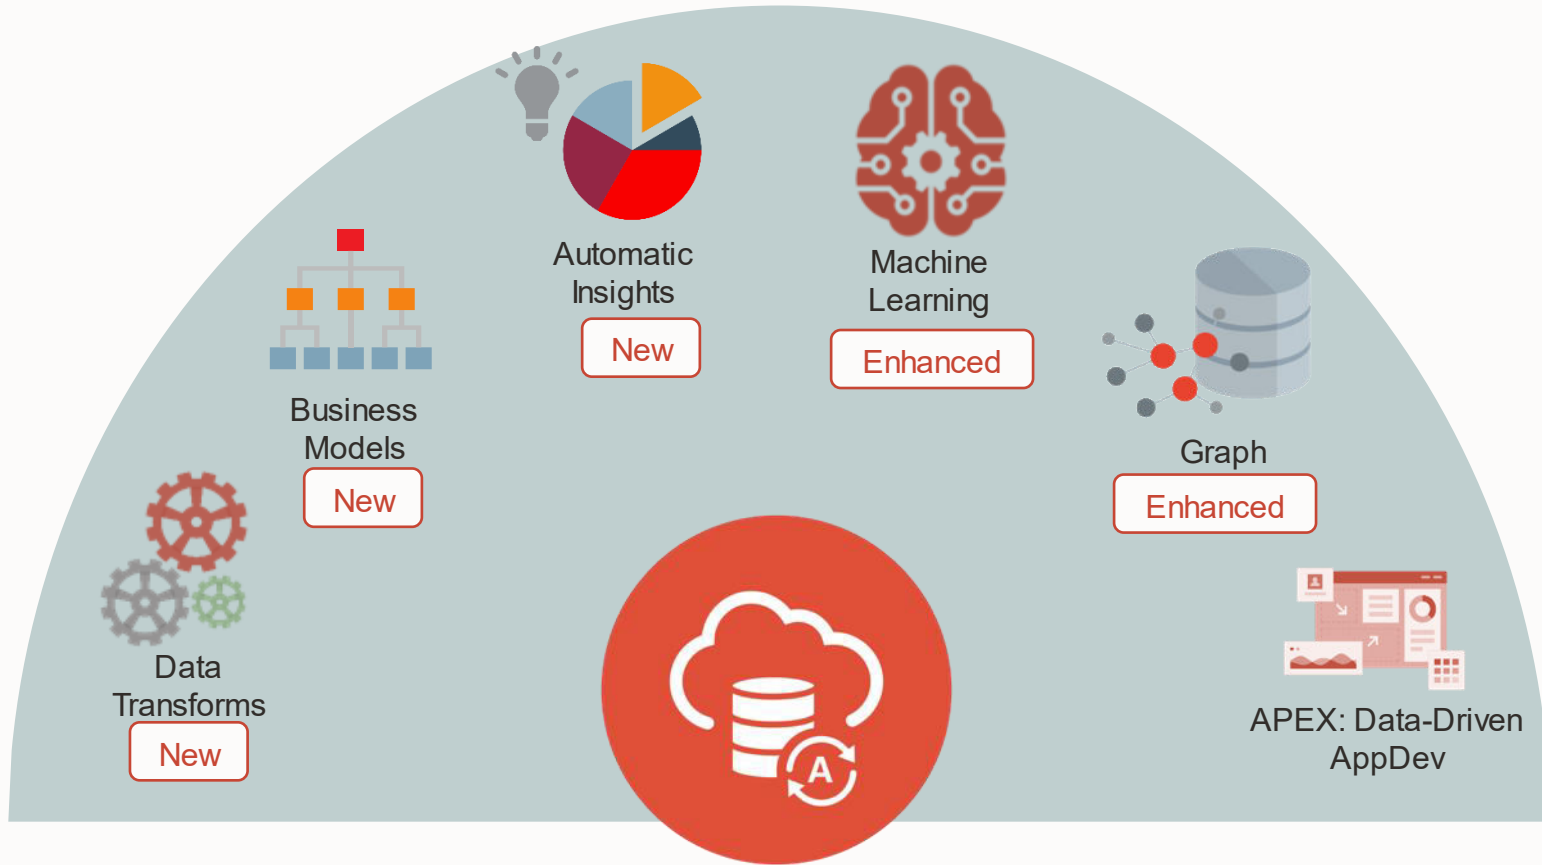
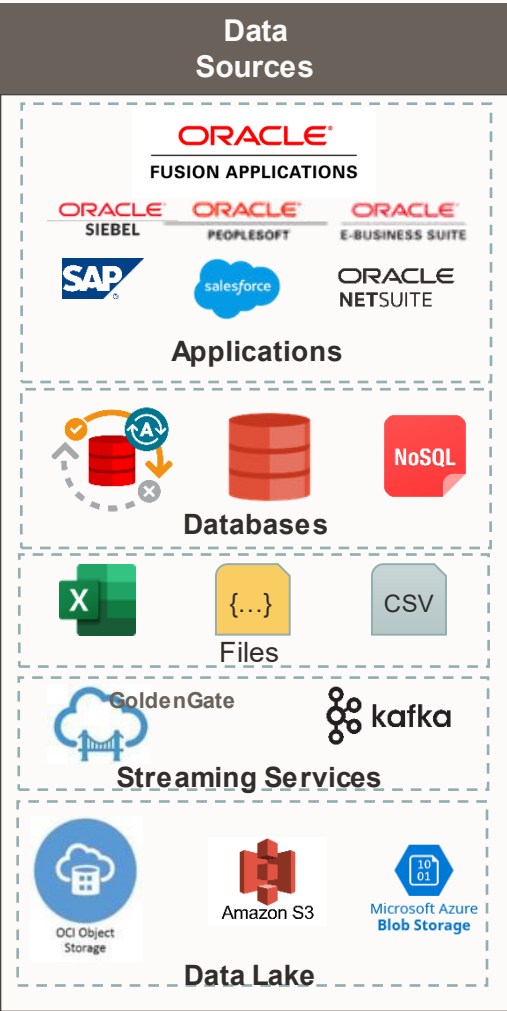
Goal:

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

Solution:

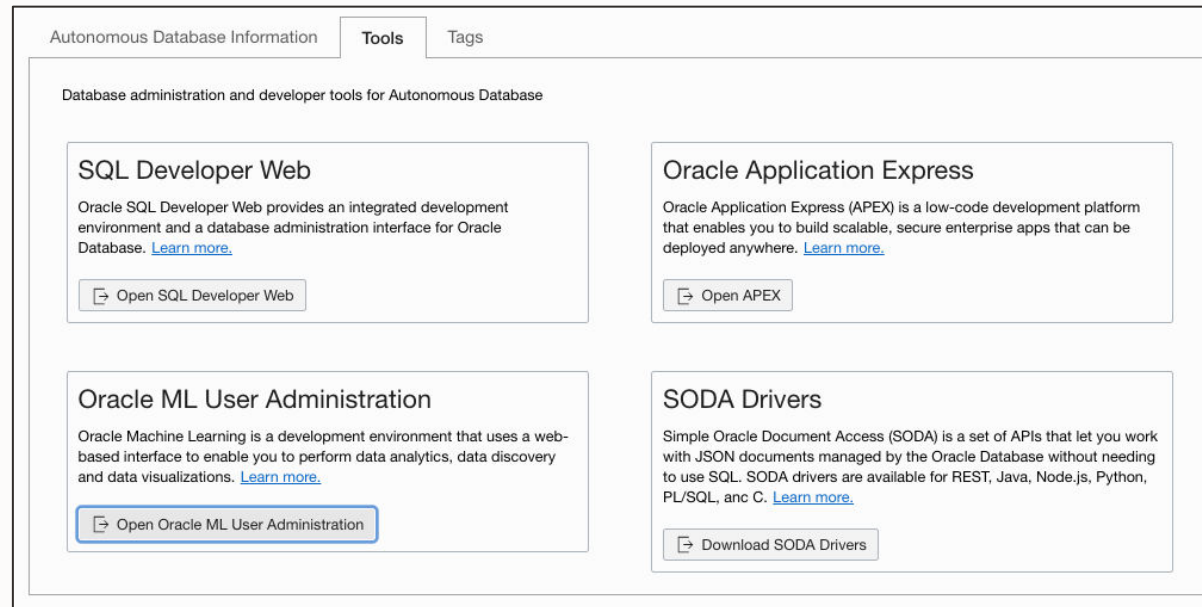
- Extend Autonomous Database for:
 - Data ingestion and transformations
 - Business modelling and analysis
 - Machine learning and automatic insights

New Autonomous Database Architecture



Autonomous Database today

A handful of integrated tools

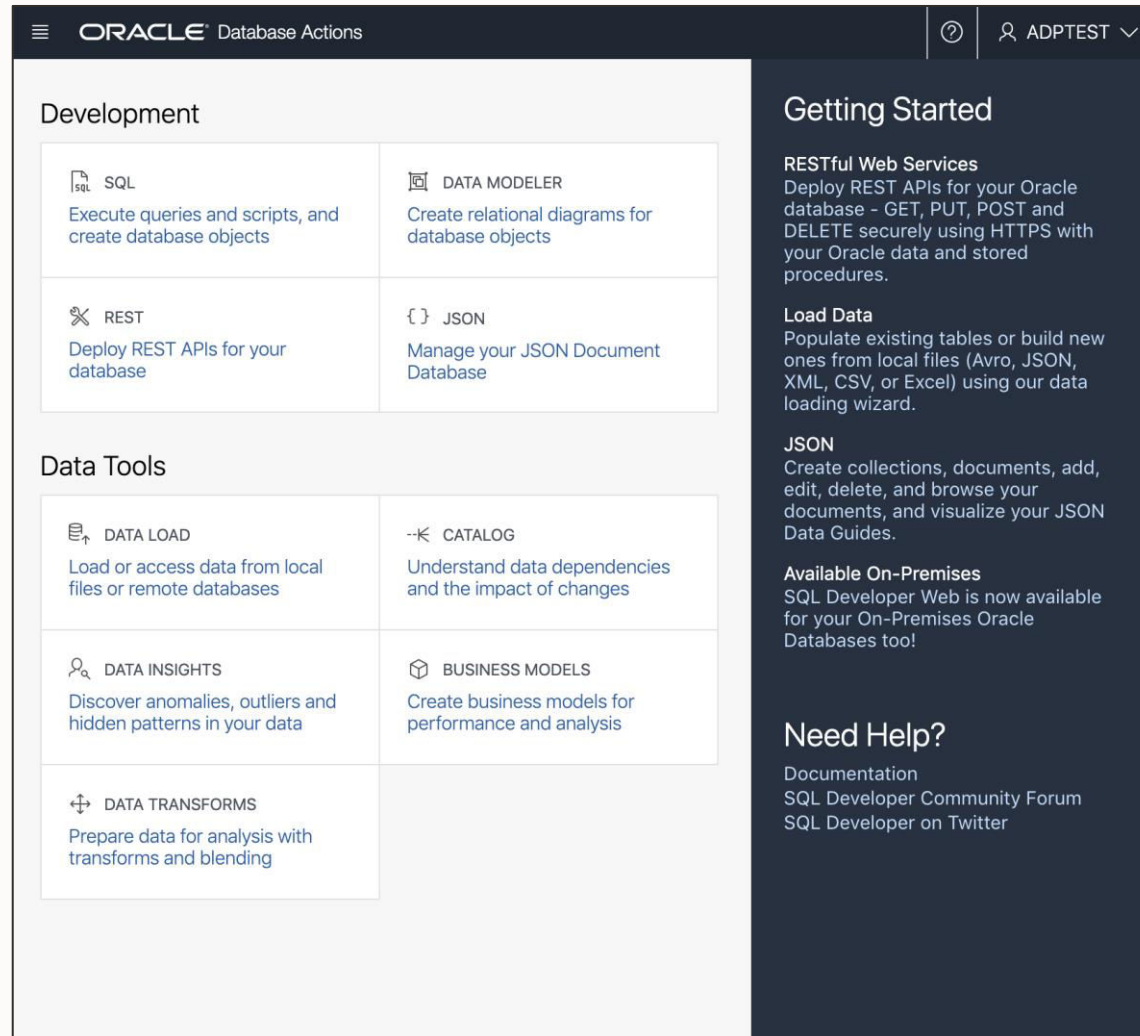


The screenshot shows a web interface for Oracle Autonomous Database tools. At the top, there are three tabs: "Autonomous Database Information", "Tools" (which is selected), and "Tags". Below the tabs, the main heading reads "Database administration and developer tools for Autonomous Database". The page is divided into four tool cards:

- SQL Developer Web**: Oracle SQL Developer Web provides an integrated development environment and a database administration interface for Oracle Database. [Learn more.](#) Below the text is a button with a right-pointing arrow and the text "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.](#) Below the text is a button with a right-pointing arrow and the text "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.](#) Below the text is a button with a right-pointing arrow and the text "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.](#) Below the text is a button with a right-pointing arrow and the text "Download SODA Drivers".

Autonomous Database – coming soon

Broad set of integrated tools and utilities



The screenshot displays the Oracle Database Actions interface. The top navigation bar includes the Oracle logo, the text 'Database Actions', a help icon, and a user profile icon labeled 'ADPTEST'. The main content area is divided into two columns. The left column features two sections: 'Development' and 'Data Tools'. 'Development' contains four tool cards: 'SQL' (Execute queries and scripts, and create database objects), 'DATA MODELER' (Create relational diagrams for database objects), 'REST' (Deploy REST APIs for your database), and 'JSON' (Manage your JSON Document Database). 'Data Tools' contains four tool cards: 'DATA LOAD' (Load or access data from local files or remote databases), 'CATALOG' (Understand data dependencies and the impact of changes), 'DATA INSIGHTS' (Discover anomalies, outliers and hidden patterns in your data), and 'BUSINESS MODELS' (Create business models for performance and analysis). 'DATA TRANSFORMS' (Prepare data for analysis with transforms and blending) is listed below the 'DATA INSIGHTS' card. The right column is titled 'Getting Started' and contains three sections: 'RESTful Web Services' (Deploy REST APIs for your Oracle database - GET, PUT, POST and DELETE securely using HTTPS with your Oracle data and stored procedures), 'Load Data' (Populate existing tables or build new ones from local files (Avro, JSON, XML, CSV, or Excel) using our data loading wizard), and 'JSON' (Create collections, documents, add, edit, delete, and browse your JSON documents, and visualize your JSON Data Guides). Below these are 'Available On-Premises' (SQL Developer Web is now available for your On-Premises Oracle Databases too!) and 'Need Help?' (Documentation, SQL Developer Community Forum, SQL Developer on Twitter).

Development	
SQL Execute queries and scripts, and create database objects	DATA MODELER Create relational diagrams for database objects
REST Deploy REST APIs for your database	JSON Manage your JSON Document Database

Data Tools	
DATA LOAD Load or access data from local files or remote databases	CATALOG Understand data dependencies and the impact of changes
DATA INSIGHTS Discover anomalies, outliers and hidden patterns in your data	BUSINESS MODELS Create business models for performance and analysis
DATA TRANSFORMS Prepare data for analysis with transforms and blending	

Getting Started

RESTful Web Services
Deploy REST APIs for your Oracle database - GET, PUT, POST and DELETE securely using HTTPS with your Oracle data and stored procedures.

Load Data
Populate existing tables or build new ones from local files (Avro, JSON, XML, CSV, or Excel) using our data loading wizard.

JSON
Create collections, documents, add, edit, delete, and browse your JSON documents, and visualize your JSON Data Guides.

Available On-Premises
SQL Developer Web is now available for your On-Premises Oracle Databases too!

Need Help?

Documentation
SQL Developer Community Forum
SQL Developer on Twitter

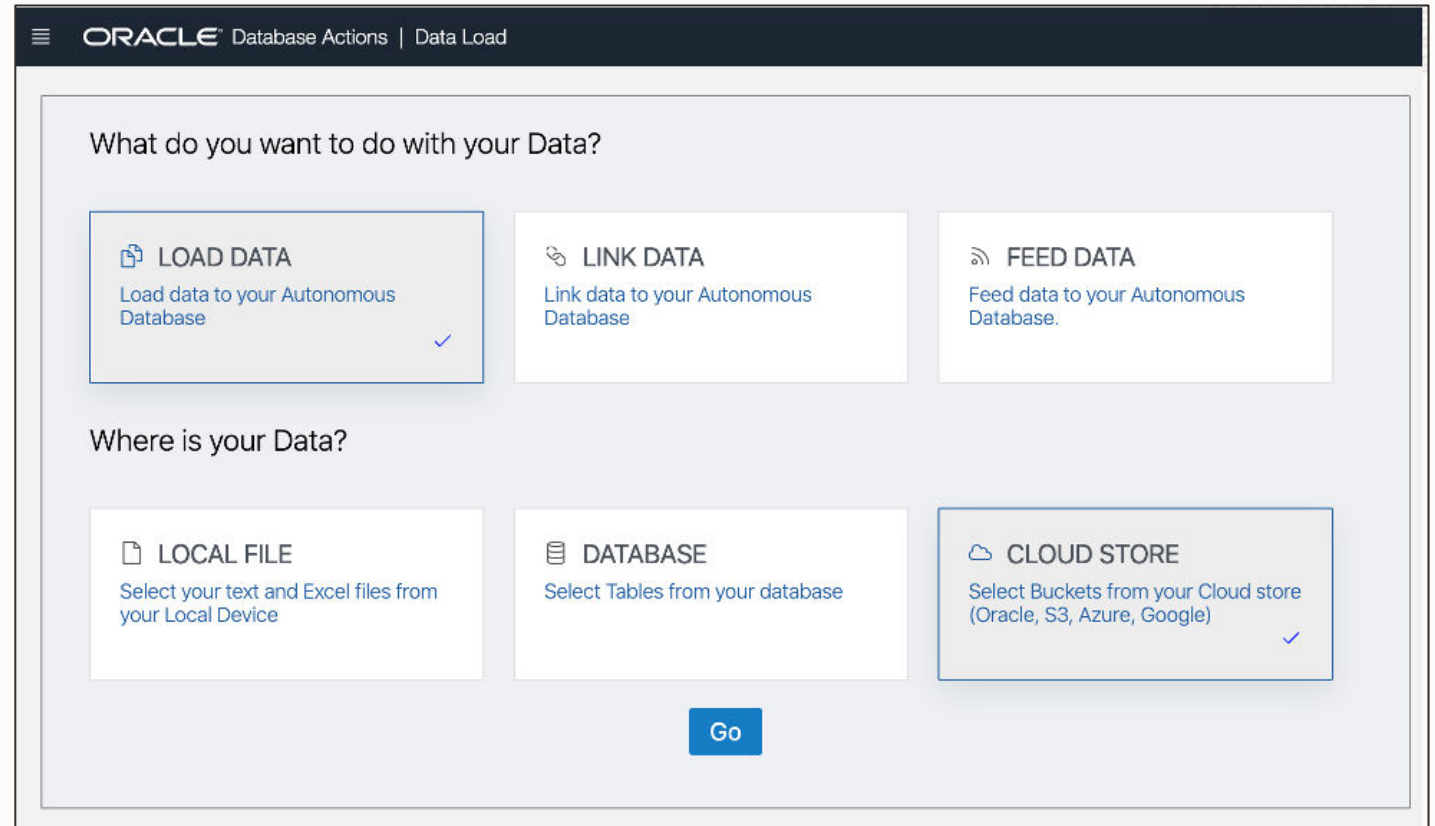


Available Today!

Data Load

Simple 'Drag and Drop' Data Loading

- Data Sources:
 - Files on local computer
 - Oracle and non-Oracle Databases (on-prem and cloud)
 - Cloud Storage (incl AWS S3, Azure Blob Storage, GCP)
- Continuous data loading
 - New files loaded as soon as they arrive in cloud storage



Available Today!
(via ODI on Marketplace)

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

The screenshot displays the Oracle Data Integrator (ODI) interface. The main window shows a Data Flow Map for a project named 'Covid-19'. The map consists of four nodes: 'COVID19_Case_...', 'NOV_ONLY', 'LAG_DAILY_NEW', and 'DAILY_NOV_20'. The 'LAG_DAILY_NEW' node is selected, and its properties are displayed in the right-hand pane. The properties include:

- Name:** LAG_DAILY_NEW
- Function Name:** Lag
- Description:** Need to convert cumulative totals to daily new. Use lag, with window on State, County, order by date.
- Help:** Lag
- Function Signature:** LAG(expression, offset, default) OVER (PARTITION BY partition ORDER BY order)



Business Models (Analytic Views)

Analytic Views

- Enable high-performance multidimensional analysis over relational data

Automatic creation

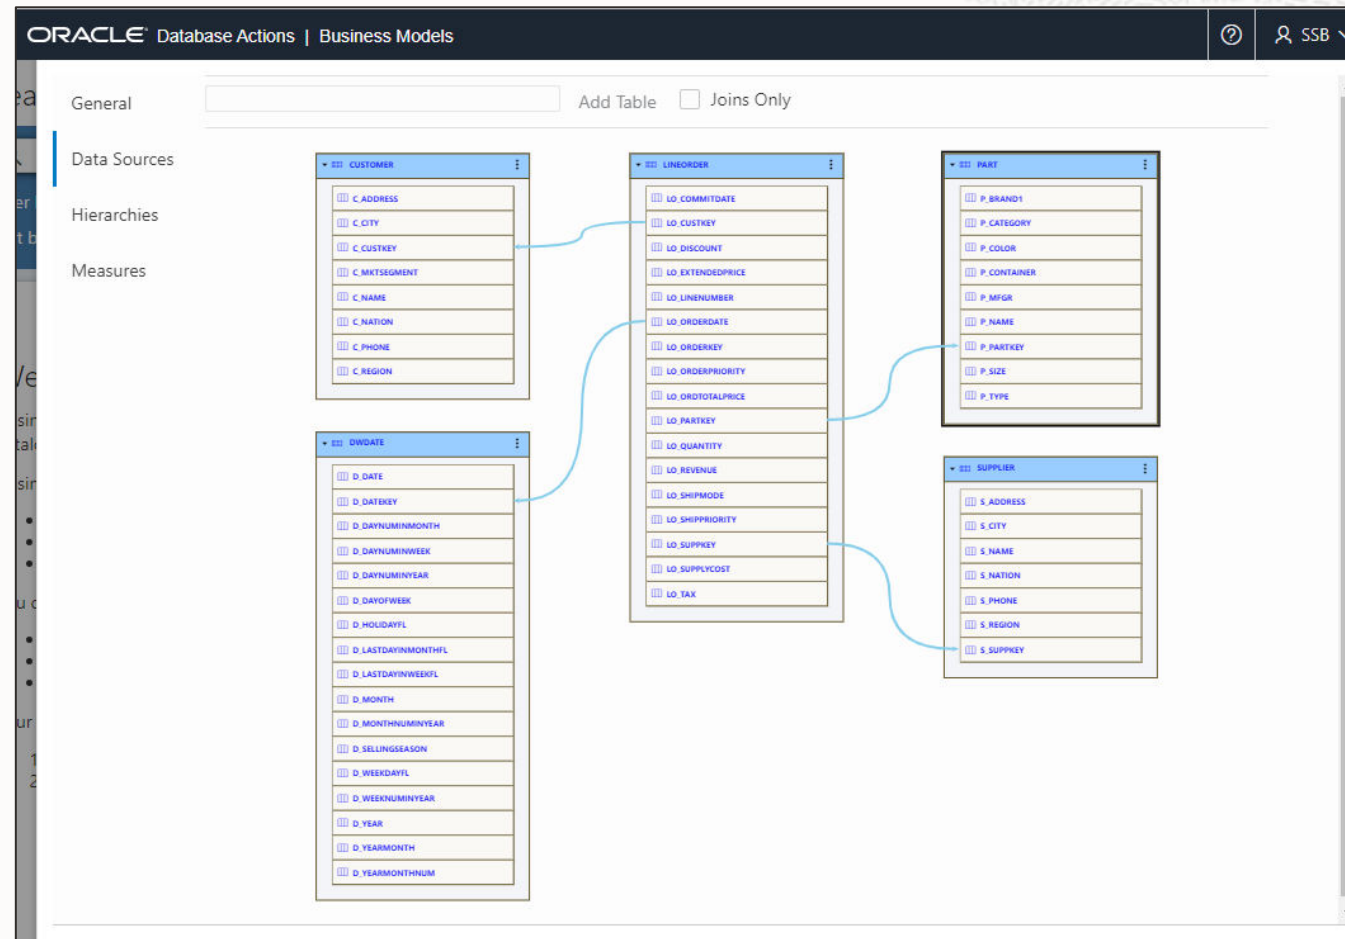
- Business Models embedded in ADW

Automatic multidimensional cache

- For both local and remote data

Use with any analytics tool

- Standard SQL queries
- No application changes required



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

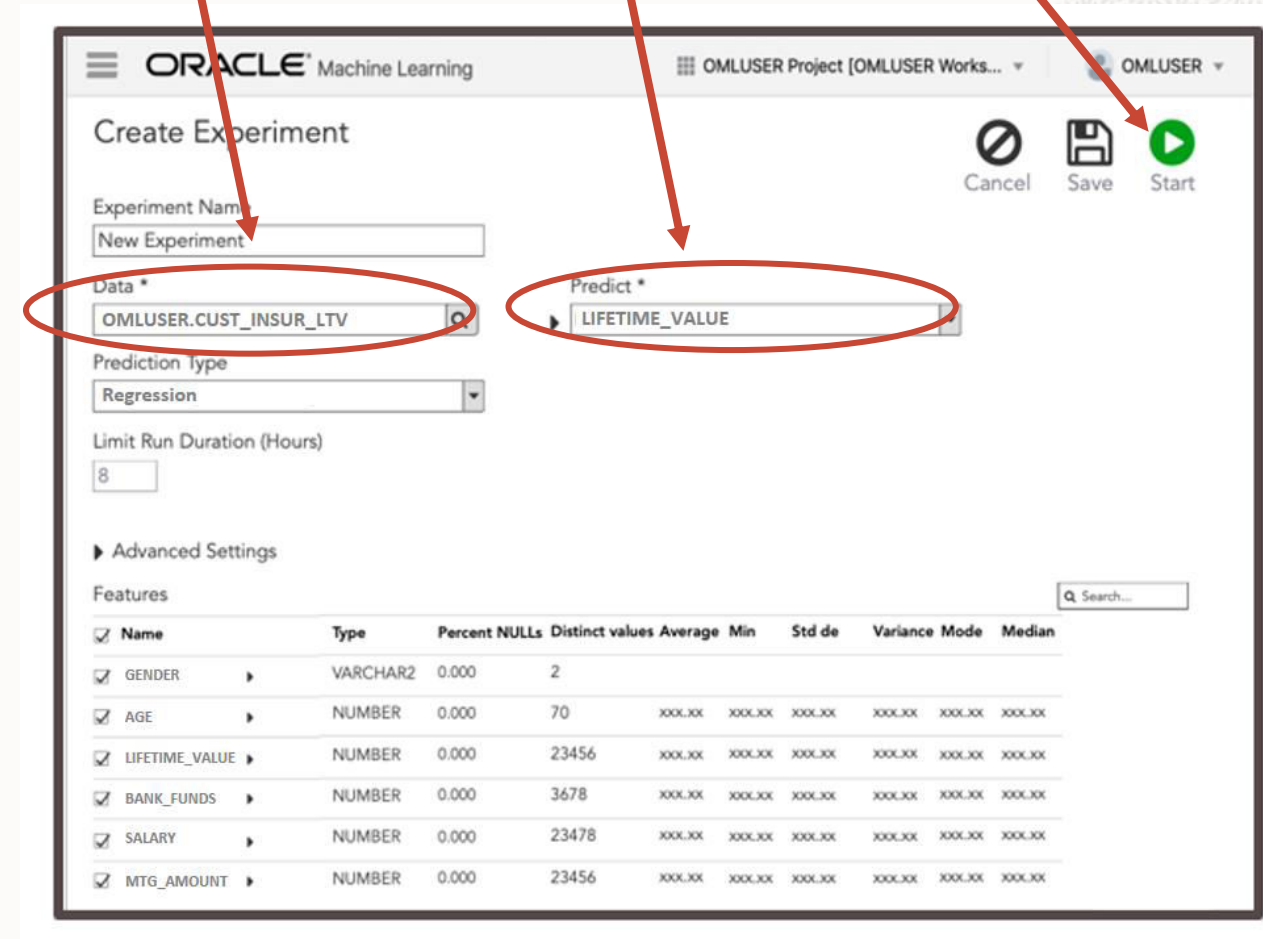
New: automatic machine learning (AutoML)

- 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'. It includes a 'Metric Chart' showing a line graph, a 'Leader Board' table, and a 'Features' table. A red circle highlights the Leader Board table, and a red arrow points from the text 'Multiple algorithms compared and best is selected' to it.

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

Name	Type	Percent NULLs	Distinct Values	Min	Max	Mean	Std Dev
PROD_CATEGORY	VARCHAR2	0	5				
PROD_CATEGORY_DESC	VARCHAR2	0	5				

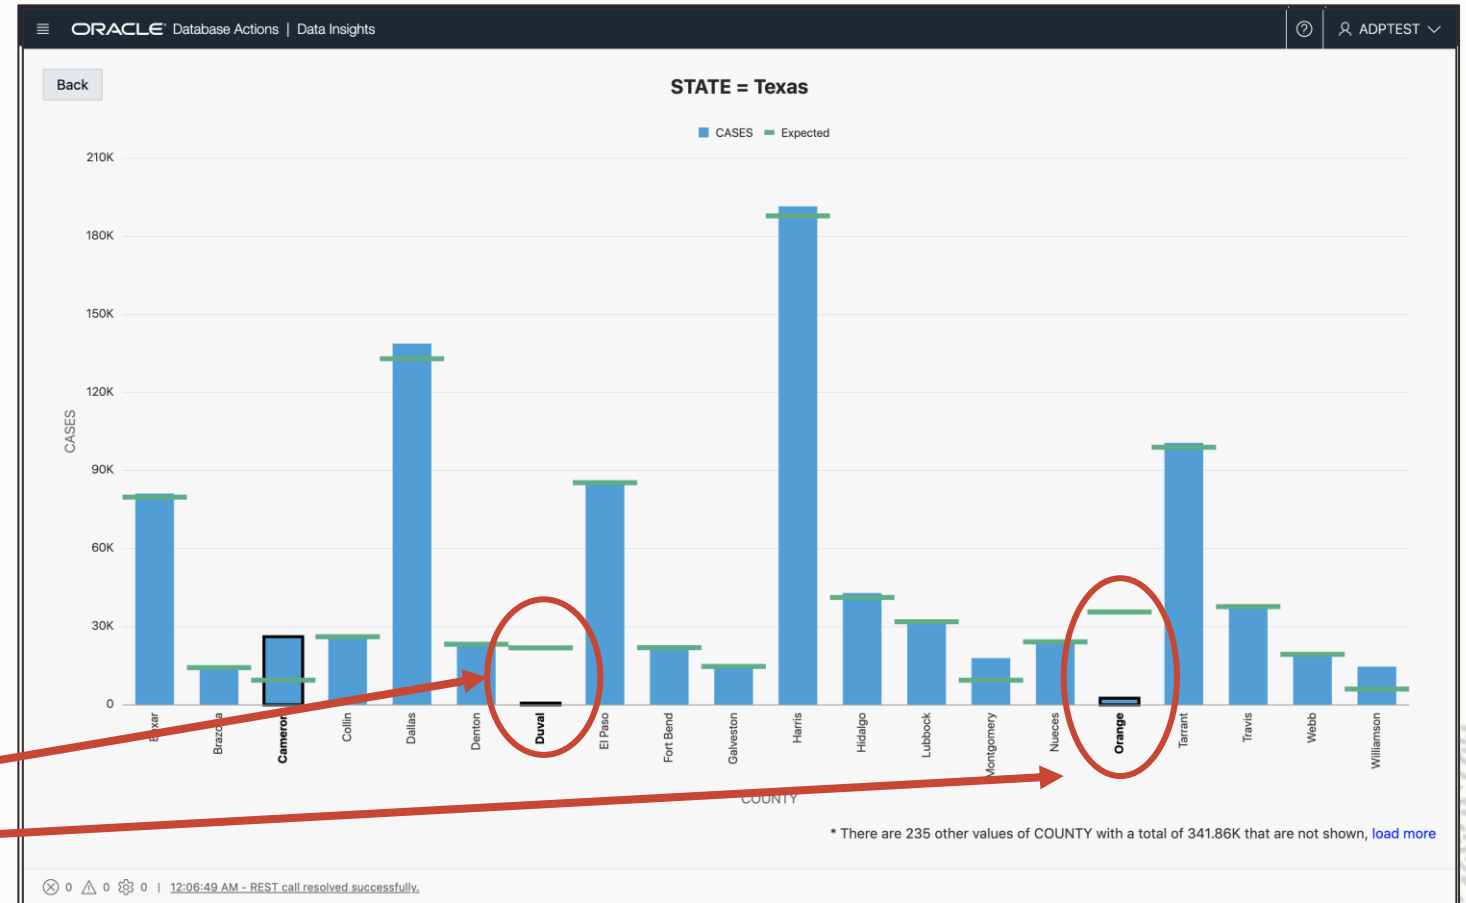


Auto-Insights – Driven by Machine Learning

Automatic Insight Discovery

- Discovers **hidden patterns** and outliers
- **Crawl over business model**, running as background process
- **Variety of algorithms** including singular value decomposition

Automatically identify anomalies and predicted values



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 Actions | Catalog interface. The main window shows the 'GEOG_DIM' entity with a 'Lineage' view selected. The lineage diagram illustrates the flow of data from the 'GEOG_DIM' entity (containing 'COUNTY', 'FIPS', 'REGION', and 'STATE') to the 'ADP_DATALOAD_LOCALFI...' table, which then feeds into a 'Geog_Dim.csv' file. The left sidebar contains filters for Schema (ADPTEST), Entity type (ANALYTIC_VIEW), and various other options like 'Entity name', 'Entity type', 'Created on', 'Filters', 'Schema', 'Entity type', and 'Applications'. The top navigation bar shows 'ORACLE Database Actions | Catalog' and the user 'ADPTEST'.





ORACLE