Ali je Oracle Autonomous Database prihodnost razvoja Oracle baze podatkov?

Autonomous Database 2022

What‘s New

Robert Korošec

Oracle

robert.korosec@oracle.com

Oracle Autonomous Database | What's Inside?

Eliminates All Database & Infrastructure Complexity

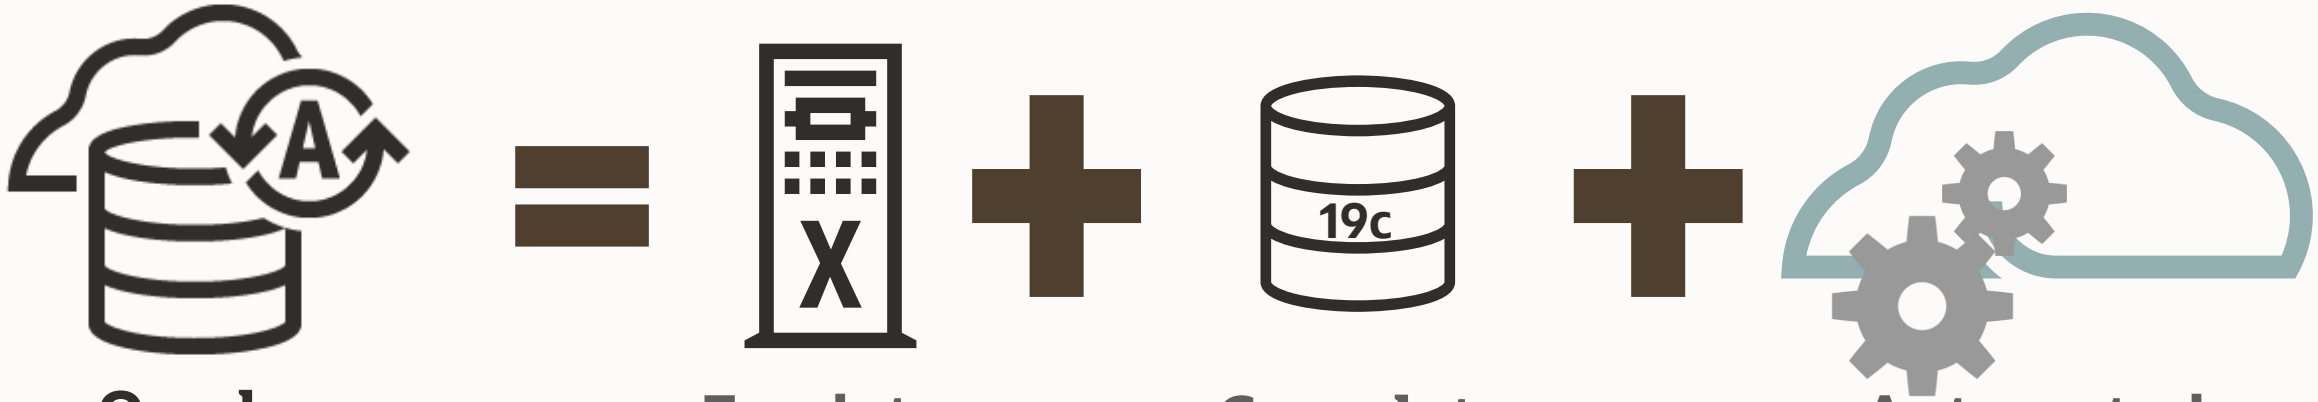

Oracle Autonomouss Database

Exadata Complete Infrastructure Automation

Complete Database Automation

Automated Data Center Operations and Machine Learning

Exadata Vision

Extreme Performance and Availability, Lowest Cost, Available Everywhere

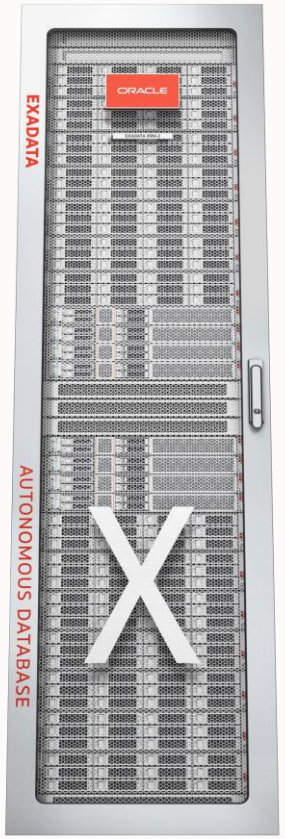

Ideal Database Hardware

Scale-out, database optimized compute, networking, and storage

Database Aware System Software

Unique algorithms vastly improve OLTP, Analytics, and Consolidation

Automated Management

Fully automated and optimized end-to-end

Available

On Premises

Cloud@Customer

Oracle Cloud

Exadata Superior OLTP Architecture - Persistent Memory

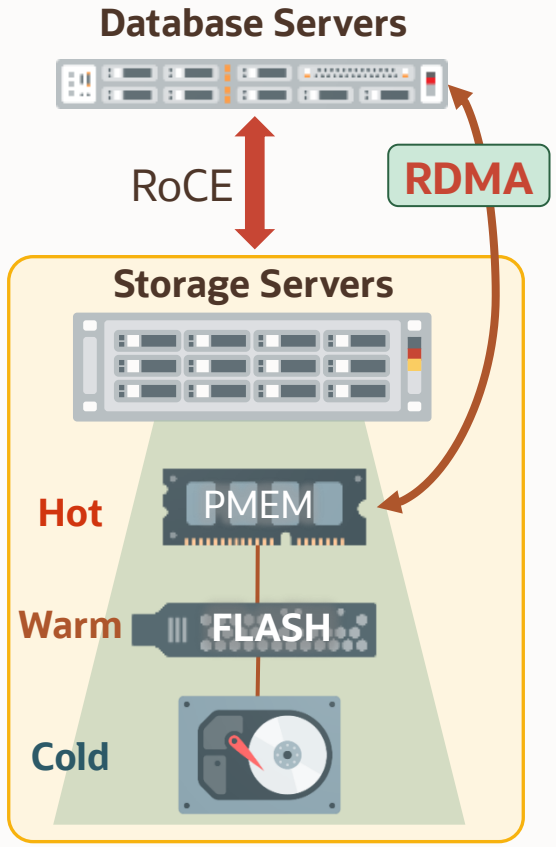

Database uses RDMA instead of I/O to read PMEM in Smart Storage

- Bypasses network and I/O software, interrupts, context switches

Storage servers transparently manage PMEM

- Hottest data automatically cached in PMEM, increasing effective capacity10x
- Automatic PMEM redundancy across multiple storage servers

Speed of RDMA to PMEM for both database reads and commits

Result - 19µs IO latency from Database to PMEM in Storage

- 10X Faster than Flash

World's Only Shared Persistent Memory Optimized for Database

Fastest Cloud Analytics

Smart Scan (SQL Offload)

- Data-intensive processing* runs in massively parallel Exadata Storage, bypassing network bottlenecks and freeing up DB CPUs

Tiered Flash Cache

- Active data is automatically cached on PCI NVMe Flash, inactive data on low cost, high-capacity disks

Storage Indexes

- Eliminates I/O not relevant to a particular query

Hybrid Columnar Compression (HCC)

- Compressed, columnar format in storage, saving space, reducing I/O, speeding analytic queries

In-Memory Columnar (IMC)

- Extends In-Memory database performance to higher capacity Flash memory in storage

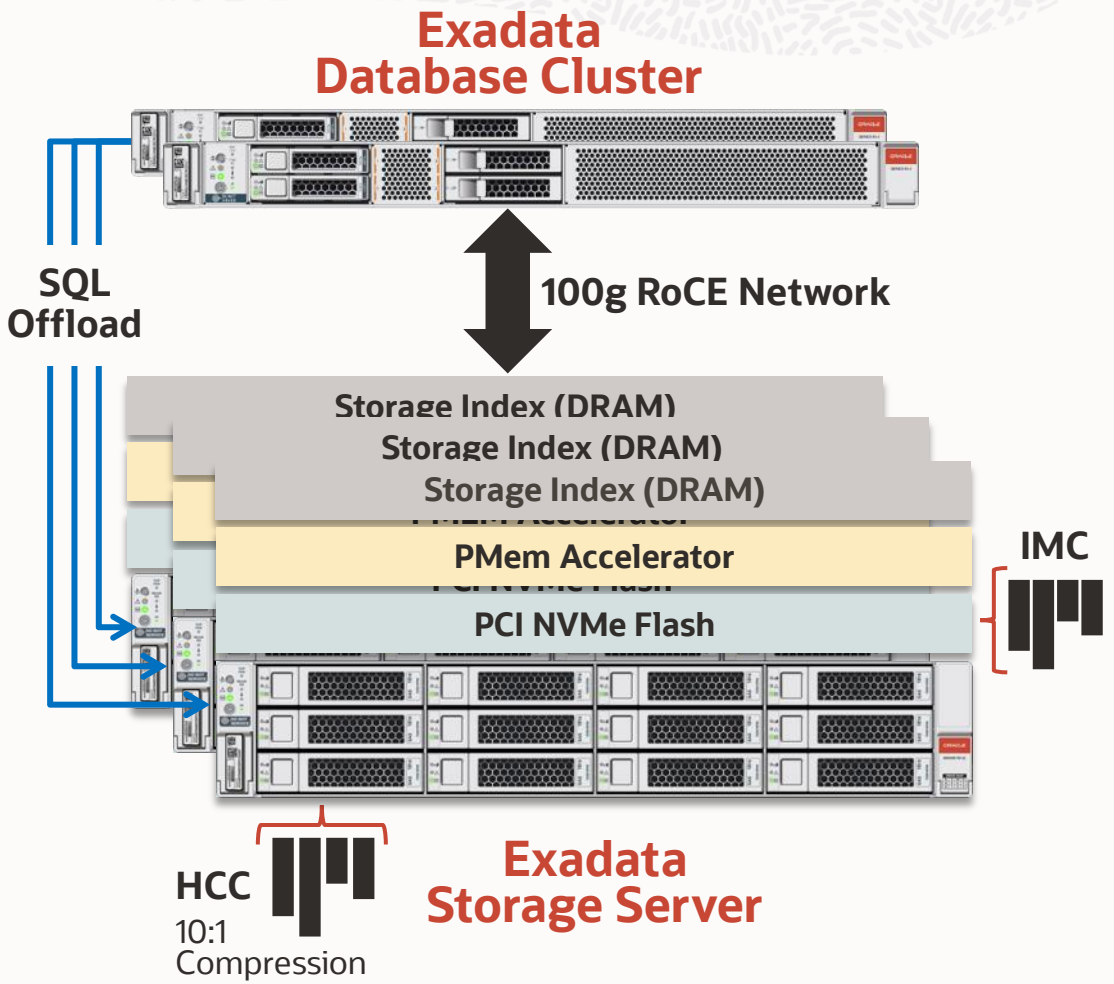

*Includes long-running SQL queries, backups, decryption, aggregation, data mining

What is Autonomous Database?

Family of products based on a common platform, optimized for specific workloads

Autonomous
Data Warehouse

Optimized
for JSON

Autonomous
Transaction Processing

Optimized
for Graph

Shared and Dedicated Infrastructure

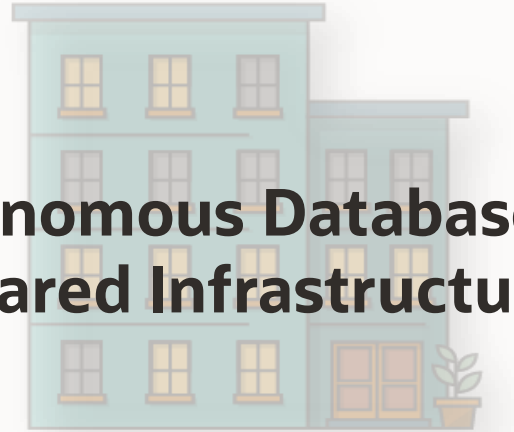

Autonomous Database on Shared Infrastructure

Simplest experience for non-technical users

- Oracle **automates** and **manages everything**

  Deployment, lifecycle, software updates, etc.
- Just choose database resources (compute, storage) & region

Completely **Elastic**

- Low minimum size - 1 OCPU and 1TB
- No fixed sized building blocks
- Scale up and down on-demand with zero downtime
- Automatically scales online **for true pay-per-use**

Autonomous. Database on Dedicated Infrastructure

Provides your own **Database Cloud** running on dedicated Exadata Infrastructure

- Runs all your databases - any size, scale, or criticality

Customizable **Isolation**

- Runs inside Secure Isolation Zone for highest protection from other tenants
- Configure multiple Exadatas or Container Databases for intra-company isolation

Customizable **Operational Policies**

- Control of provisioning, updates, availability, density

Oracle Autonomous Database | Deployment Options

Multiple Deployment Choices

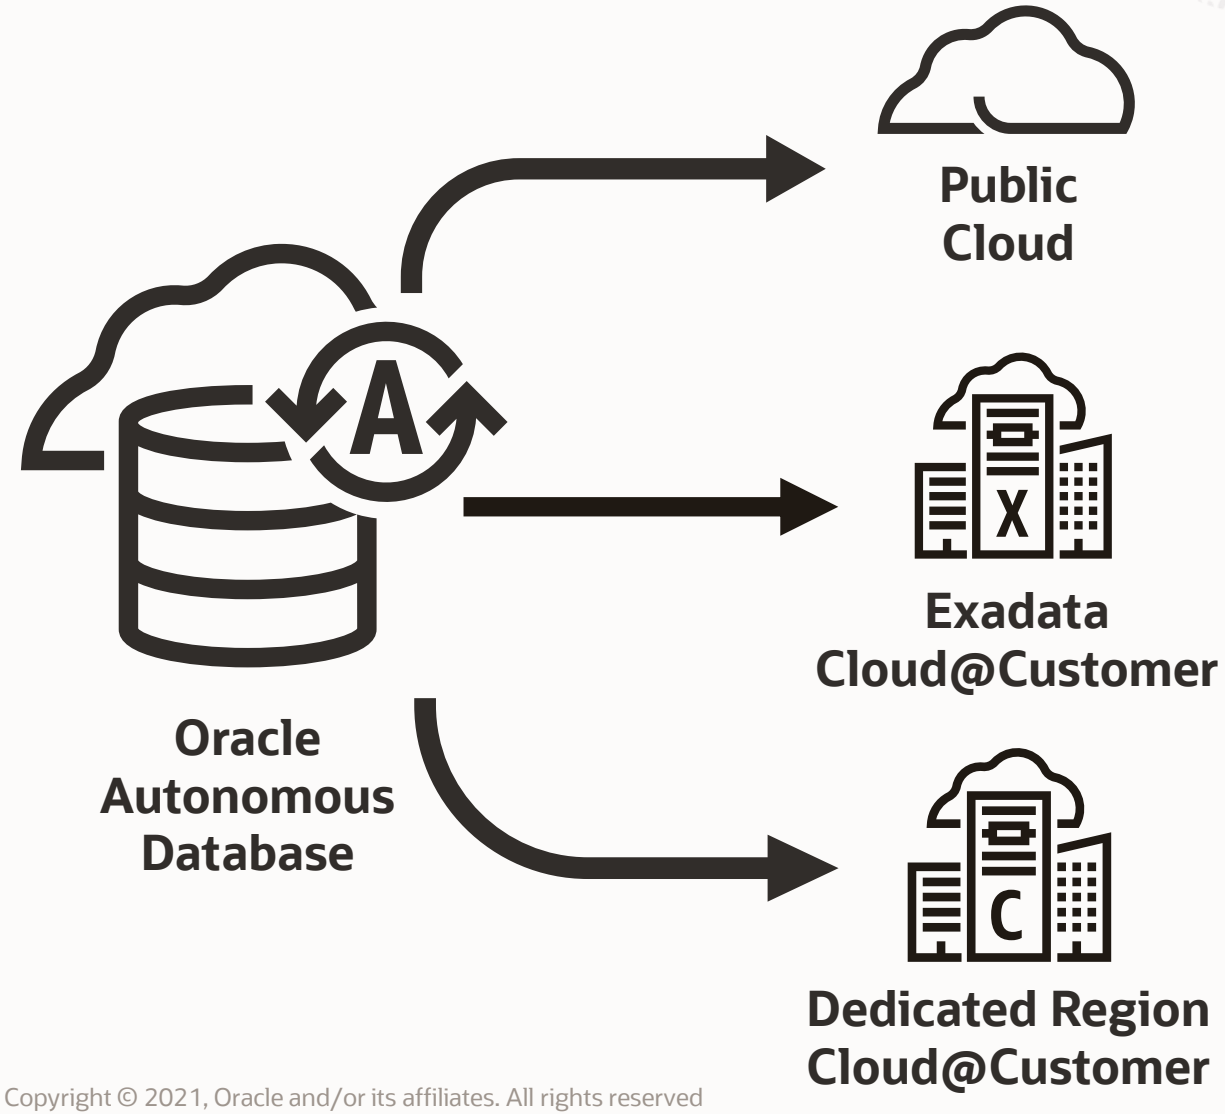

Oracle Cloud Infrastructure Global Footprint

October 2021: 31 Regions Live, 13 planned, 8 Azure Interconnect Regions

100% renewable energy by 2025

Self-Driving | Auto-Provisioning

Instantly deploy a mission critical database

Quickly deploys a RAC Database on Exadata

- Provides a HA configuration
- Ensures seamless scale-out
- Enables online patching

Autonomous Database decides where to place each database during provisioning

- Fewer number of instances preferred
- Databases may be open on only one node
  - Still RAC enabled

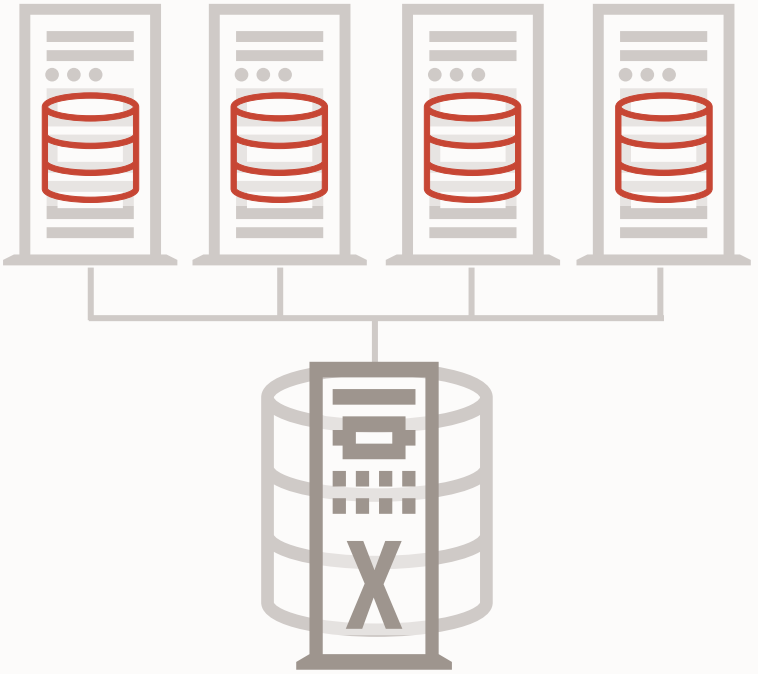

Self-Driving | Auto-Configuration

	Autonomous Data Warehouse	Autonomous Transaction Processing
	Optimizes Complex SQL	Optimizes Response Time
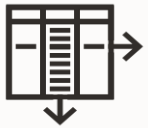	Columnar Format	Row Format
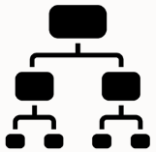	Creates Data Summaries	Creates Indexes
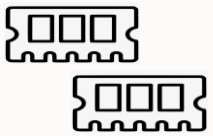	Memory Speeds Joins, Aggs	Memory for Caching, No IO
	Statistics gather as part of bulk load operations	Statistics gather as part of DML operations

Self-Driving |Optimizer Statistics

High-frequency Statistics Collection

- Gathers full optimizer statistics every 15 minutes if statistics are stale
- Controlled using `DBMS_STATS.SET_GLOBAL_PREFS('AUTO_TASK_STATUS','ON');`

Statistics are gathered automatically for direct load operations

- `CREATE TABLE … AS SELECT …;` `INSERT /*+ append */ INTO … SELECT …;`
- Data Pump Import loads
- DBMS_CLOUD loads

Real-Time Statistics are gathered automatically for DML operations*

- Gathers a subset of optimizer statistics for conventional DML operations
- Number of rows, MAX and MIN column values, etc.

* Controlled via DBMS_STATS – Not on by default

Self-Driving | Database Parameters

Ability to customize the service for a given region or time zone

All parameters set to optimal values based on workload type

- May be different than regular database defaults

Users can only change a limited number of parameters

Allowed Parameters

APPROX_FOR_AGGREGATION
APPROX_FOR_COUNT_DISTINCT
APPROX_FOR_PERCENTILE
AWR_PDB_AUTOFLUSH_ENABLED
CURRENT_SCHEMA (Session only)
CURSOR_SHARING
DB_BLOCK_CHECKING
DDL_LOCK_TIMEOUT
FIXED_DATE
GLOBAL_NAMES
HEAT_MAP
MAX_IDLE_TIME
NLS_CALENDAR
NLS_COMP
NLS_CURRENCY
NLS_DATE_FORMAT
NLS_DATE_LANGUAGE
NLS_DUAL_CURRENCY
NLS_ISO_CURRENCY
NLS_LANGUAGE
NLS_LENGTH_SEMANTICS
NLS_NCHAR_CONV_EXCP
NLS_NUMERIC_CHARACTERS
NLS_SORT
NLS_TERRITORY
NLS_TIME_FORMAT
NLS_TIME_TZ_FORMAT
NLS_TIMESTAMP_FORMAT
NLS_TIMESTAMP_TZ_FORMAT
OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES (Session only)
OPTIMIZER_IGNORE_HINTS
OPTIMIZER_IGNORE_PARALLEL_HINTS
OPTIMIZER_MODE
PLSCOPE_SETTINGS
PLSQL_CCFLAGS
PLSQL_DEBUG
PLSQL_OPTIMIZE_LEVEL
PLSQL_WARNINGS
RECYCLEBIN
STATISTICS_LEVEL (Session only)
TIME_ZONE

Self-Driving | Automatic Indexing

Machine Learning auto-tunes your workload

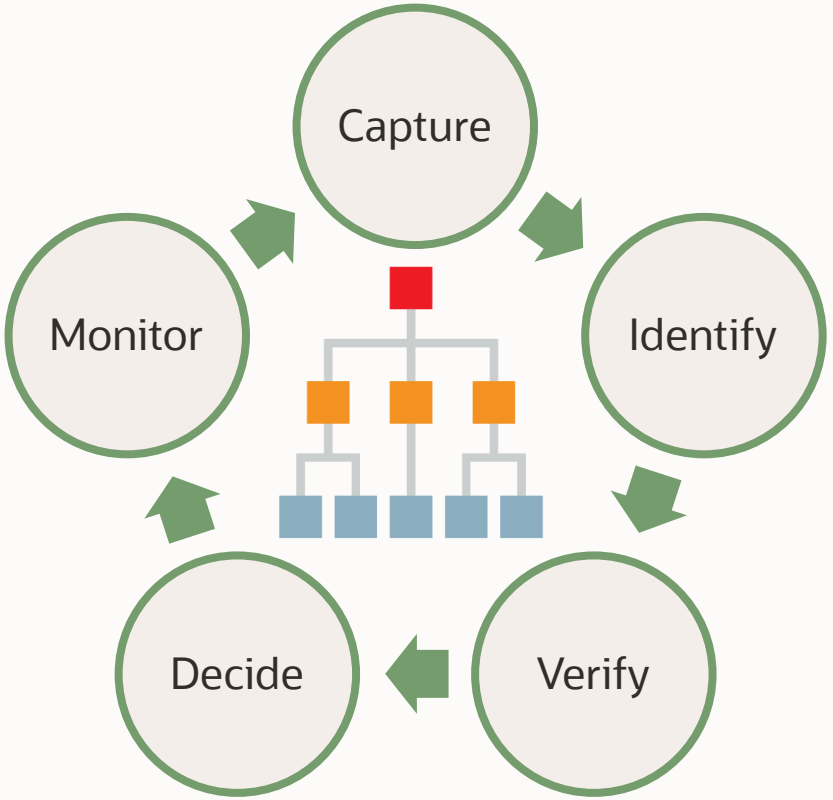

- An expert system that implements indexes based on what a skilled performance engineer would do
- Reinforcement Learning allows it to learn from its own actions as all candidate indexes are validated before being implementing
- The entire process is fully automatic
- Transparency is equally important as sophisticated automation
  - All tuning activities are auditable

Self-Driving | Automatic Partitioning

Applies Partitioning to Tables and Indexes to Improve Performance on Shared Infrastructure

Identifies candidate tables by analysing the captured workload

- Tables must be 64GB or larger and have up to date statistics

Uses one of the following strategies with a single column key:

- INTERVAL AUTOMATIC: best suited for ranges of partition key values
- LIST AUTOMATIC: applies to distinct partition key values
- HASH: partitioning on the partition key's hash values

Reinforcement Learning allows it to learn from its own actions as all candidate strategies are validated before being used

- Selects strategy with the highest estimated IO reduction

The entire process is fully automatic, but transparency is equally important as sophisticated automation

- Controlled via new `DBMS_AUTO_PARTITIONING` package

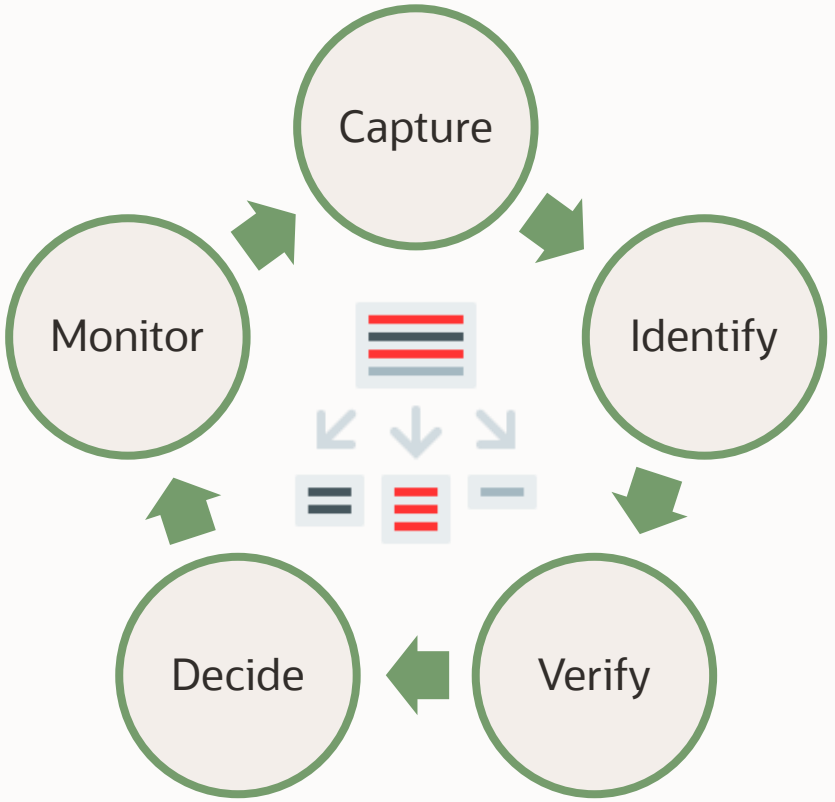

Self-Driving | Auto Stop/Start

- Scheduling automatic database shutdowns and restarts for times when the system is not in use to stop charges
  - E.g., weekends & holidays
- Available via the SDK, REST API or cloud UI
- Only pay for storage when Database is stopped

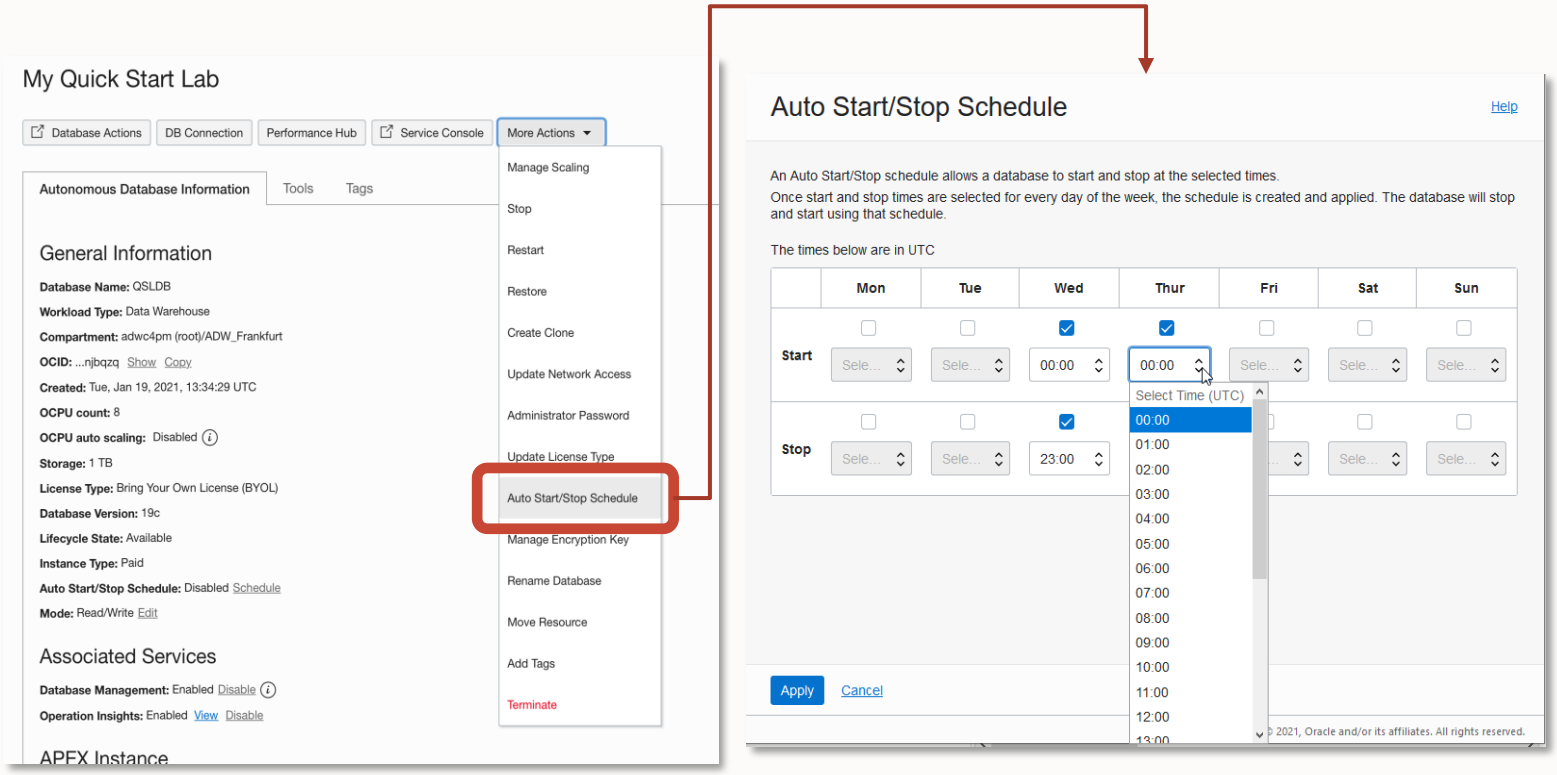

Elastic: Scaling Reduces Costs

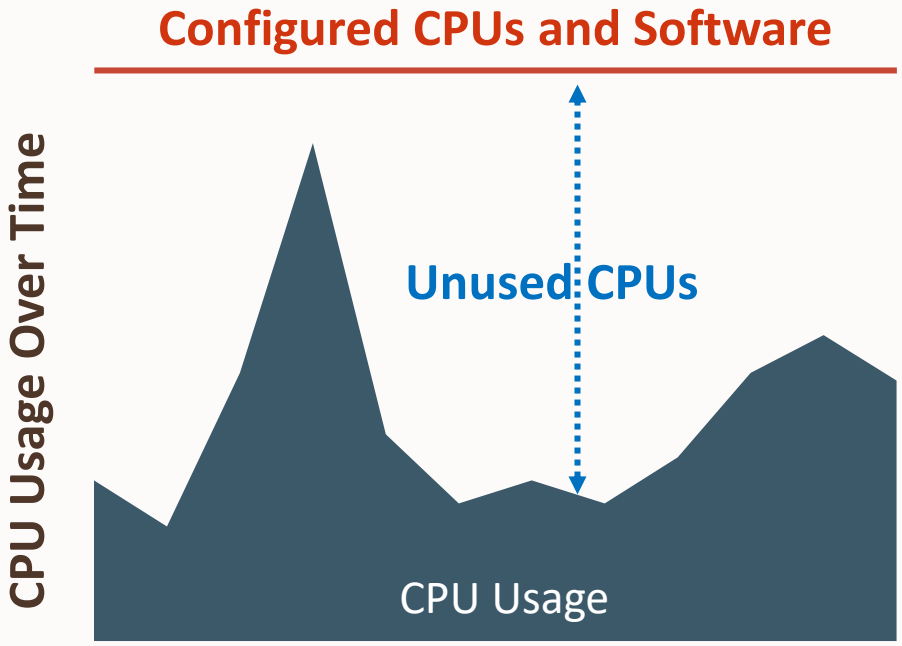

On-Premises - Static

Purchase CPUs and software licenses
for highest projected peak load

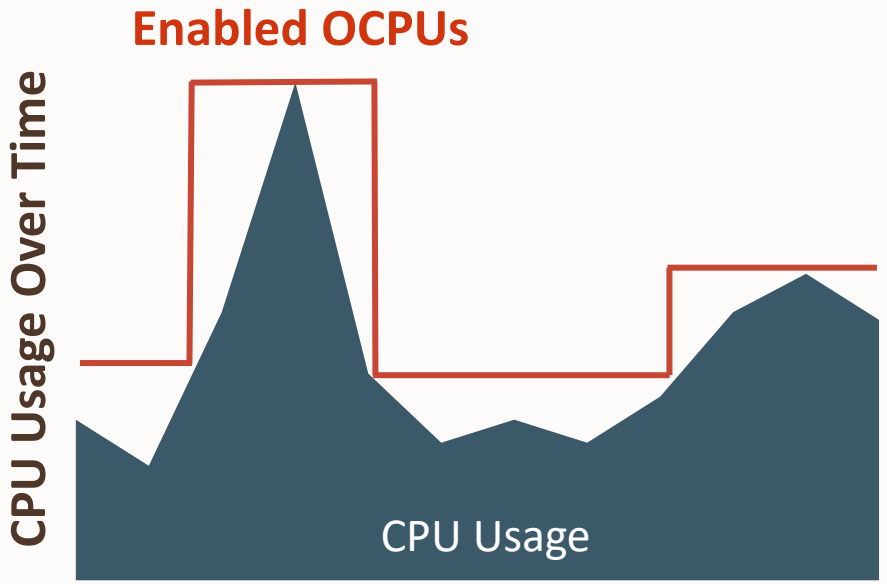

Cloud - Elastic

Adjust enabled OCPUs to match actual load
OCPUs are charged hourly
Zero OCPU Cost when stopped

Auto Scaling Delivers New Database Paradigm – True Pay per Use

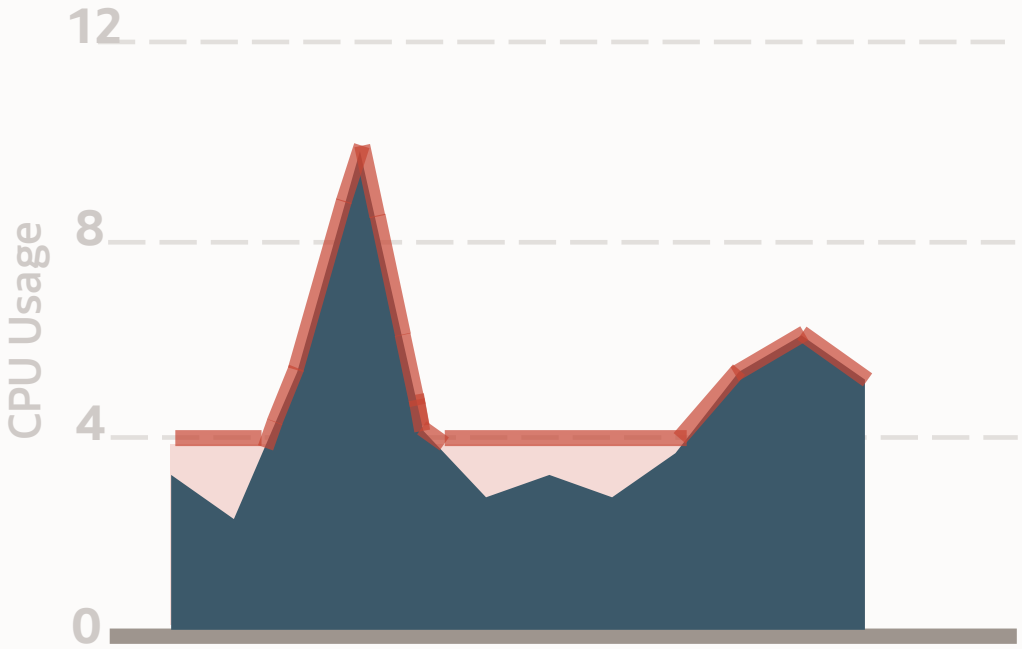

Dynamic Auto-Scale
Automatically scale up to 3X
Scaling with zero downtime

Autonomous Database

Automatically scales compute resources when needed by workload.

All scaling occurs online, while the application continuously runs.

Enables true pay per use.

Self-Driving | Automatic scaling of storage

Works in a similar way to auto scaling for OCPUs….

System automatically (transparently) expands storage on-demand up to 3x base amount.

Expands in 1TB increments

Simplifies storage management but leaves you in control!

Helps you manage ADB costs more effectively.

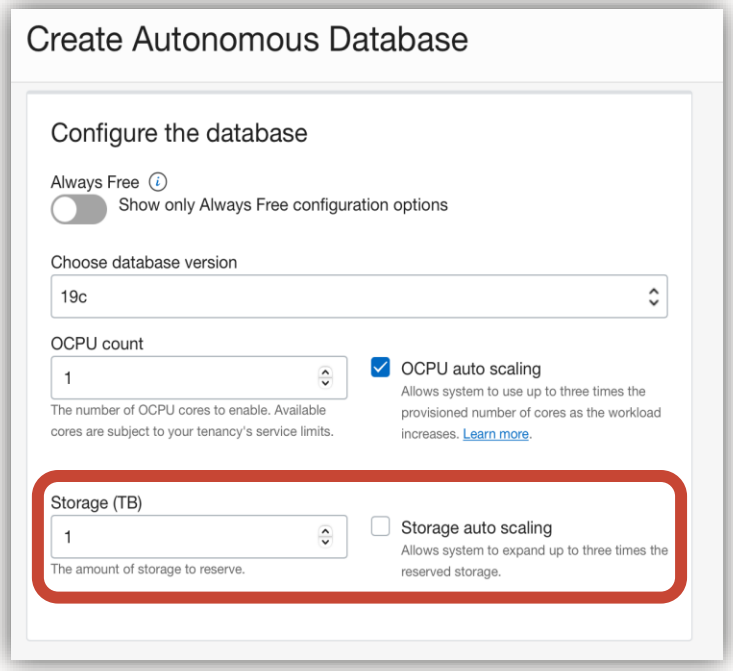

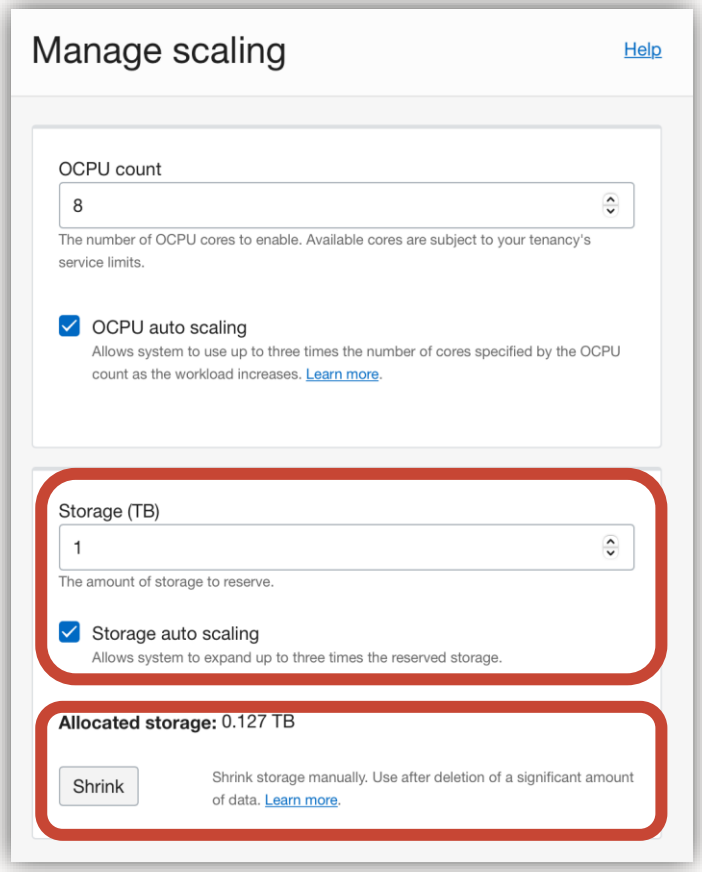

Self-Driving | Cloning

Quickly provision a point-in-time copy

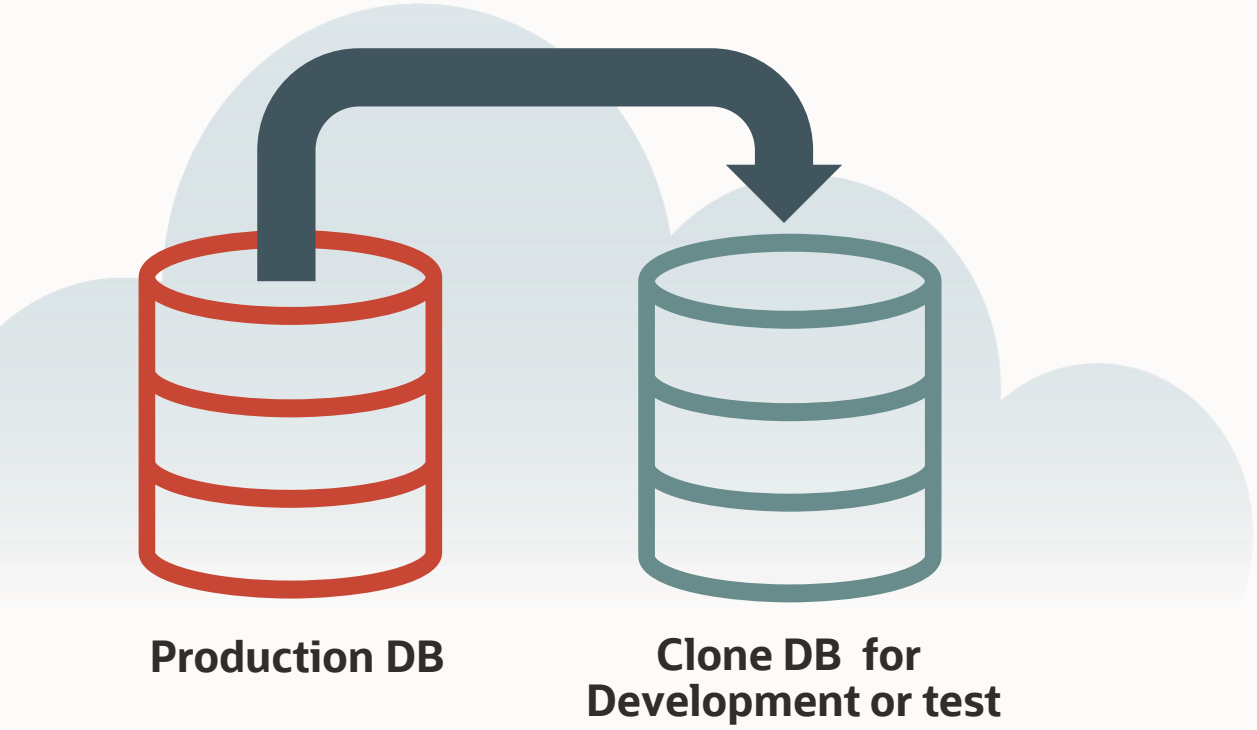

- Cloning creates a point-in-time copy of
  - An Autonomous Database
  - A backup of an Autonomous Database

Self-Driving | Cloning

Quickly provision a point-in-time copy

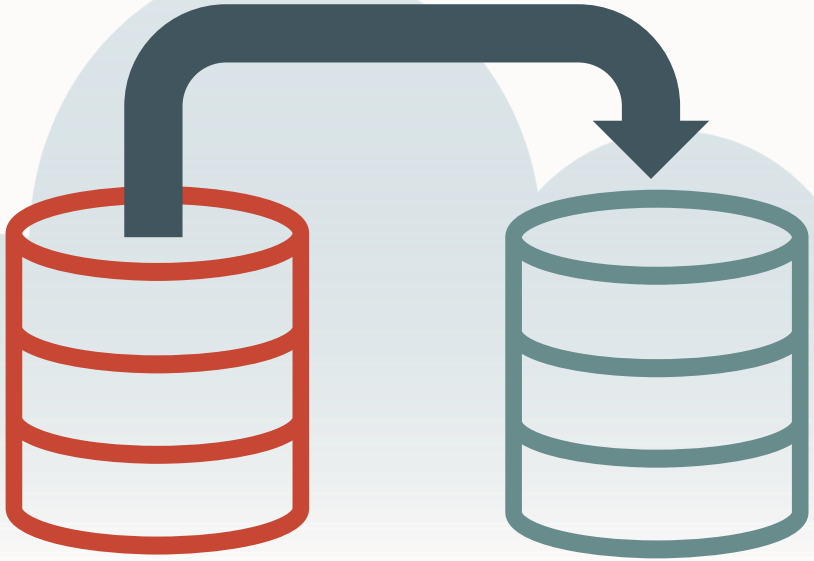

- Cloning creates a point-in-time copy of
  - An Autonomous Database
  - A backup of an Autonomous Database
- Three types of clone can be created:
  - A full database clone
  - A metadata clone (Schema but no data)
  - Refreshable clone (read-only full clone)

Note: it is possible to temporarily discount a refreshable clone for up to 24 hours to make local changes before reconnecting to the primary

Self-Securing

Only Database is exposed to users – SQL access only

- No highly privileged access – no SYSDBA access
- No login allowed to CDB - only login to PDB
- No callouts to OS allowed

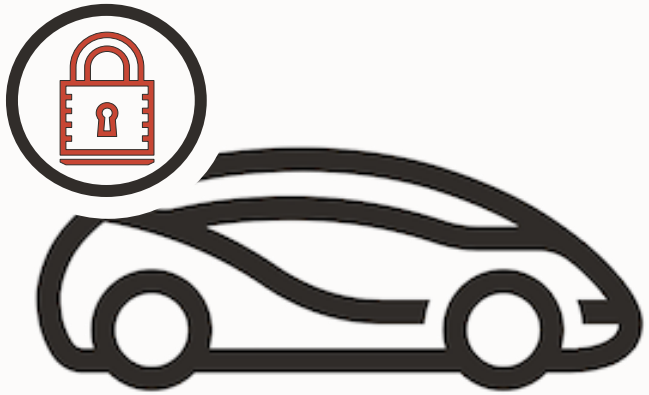

Database Vault's Automatic protects customer data from Oracle operations staff

Oracle automatically applies security updates for the entire stack

Self-Securing | Auto Patching

Automatic patching without downtime

Automatic Patching of all components (on-demand for critical security issue)

Firmware, OS, Hypervisor, Clusterware, Database

Patches applied in a rolling fashion across RAC nodes and Exadata storage servers

Database is continuously available to application

Applications using Application Continuity best practices, run without interruption

Patching is automatically scheduled

Customer can adjust patching window within a time range on Dedicated deployments

Next patching windows shown on console

Note: Early access to patches now possible on both Shared & Dedicated Infrastructure

Self-Securing | Separation of Duty

Security is a **shared** responsibility

Oracle automatically takes care of

- Data encryption by default
- Network security and monitoring
- OS and platform security
- Database patches and upgrades
- Administrative separation of duties

However, there are still areas of security that need to be managed by the customer

- Ongoing security assessments
- Users & Privileges
- Sensitive data discovery
- Data protection
- Activity auditing

Self-Securing | Oracle Data Safe

Automated Data Protection

Unified database security control center

- Security configuration assessment
- User risk assessment
- User activity auditing
- Sensitive data discovery
- Data masking

Defense in depth for all customers

- Saves time and mitigates security risks
- No special security expertise needed

Free with all Oracle Cloud Databases

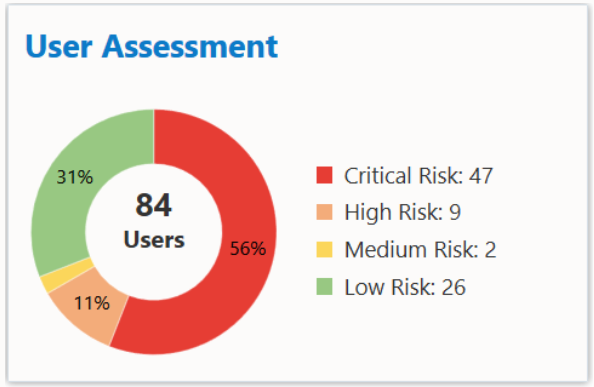

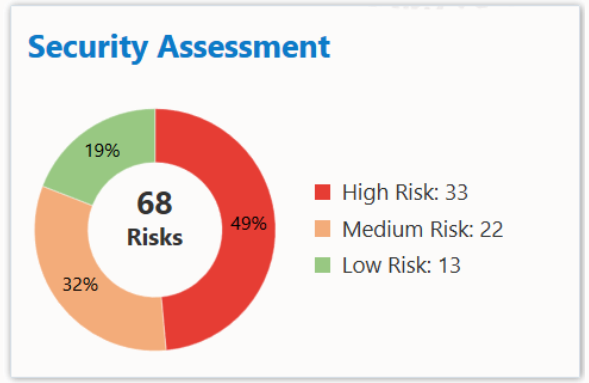

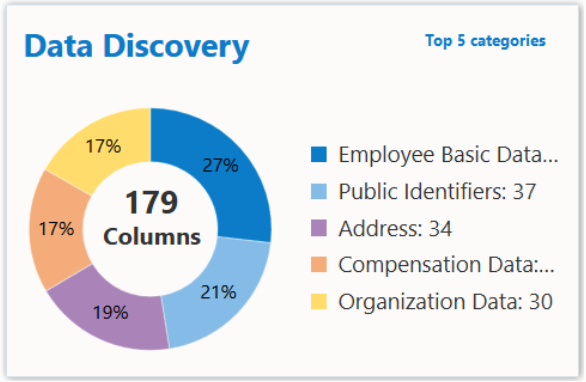

Self-Repairing | Autonomous Data Guard

Zero data loss

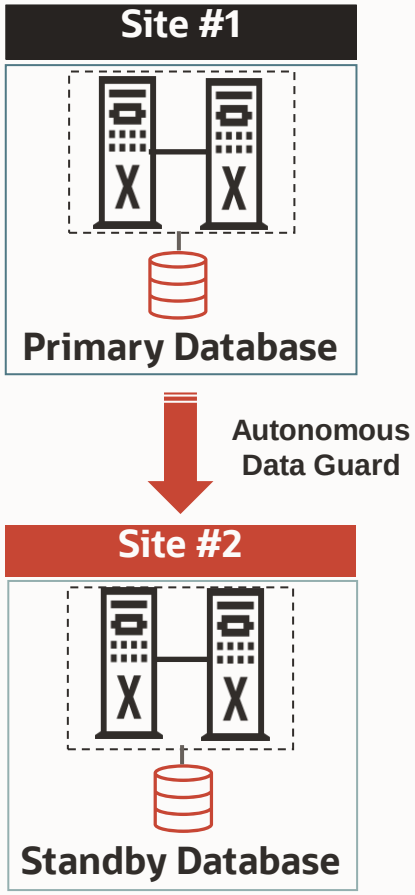

Maintains a real-time remote copy of a production database

- Protects from physical disasters, network outages, database corruptions
- Completely transparent to customer applications
- Standby can be in the same (RPO=0) or across regions

Automates everything

- Creation, operation, patching, backup, recovery, and monitoring
- Automatically handles fail over to the remote copy
- Automates resynchronization after the primary site is restored

Maintains copy by applying logical changes, not physical changes

- Prevents propagating data corruptions to standby
- Validates data consistency as changes are applied

Note: ADB on Shared Infra does not provide access to standby database(s)

Self-Repairing | Autonomous Data Guard

Zero data loss

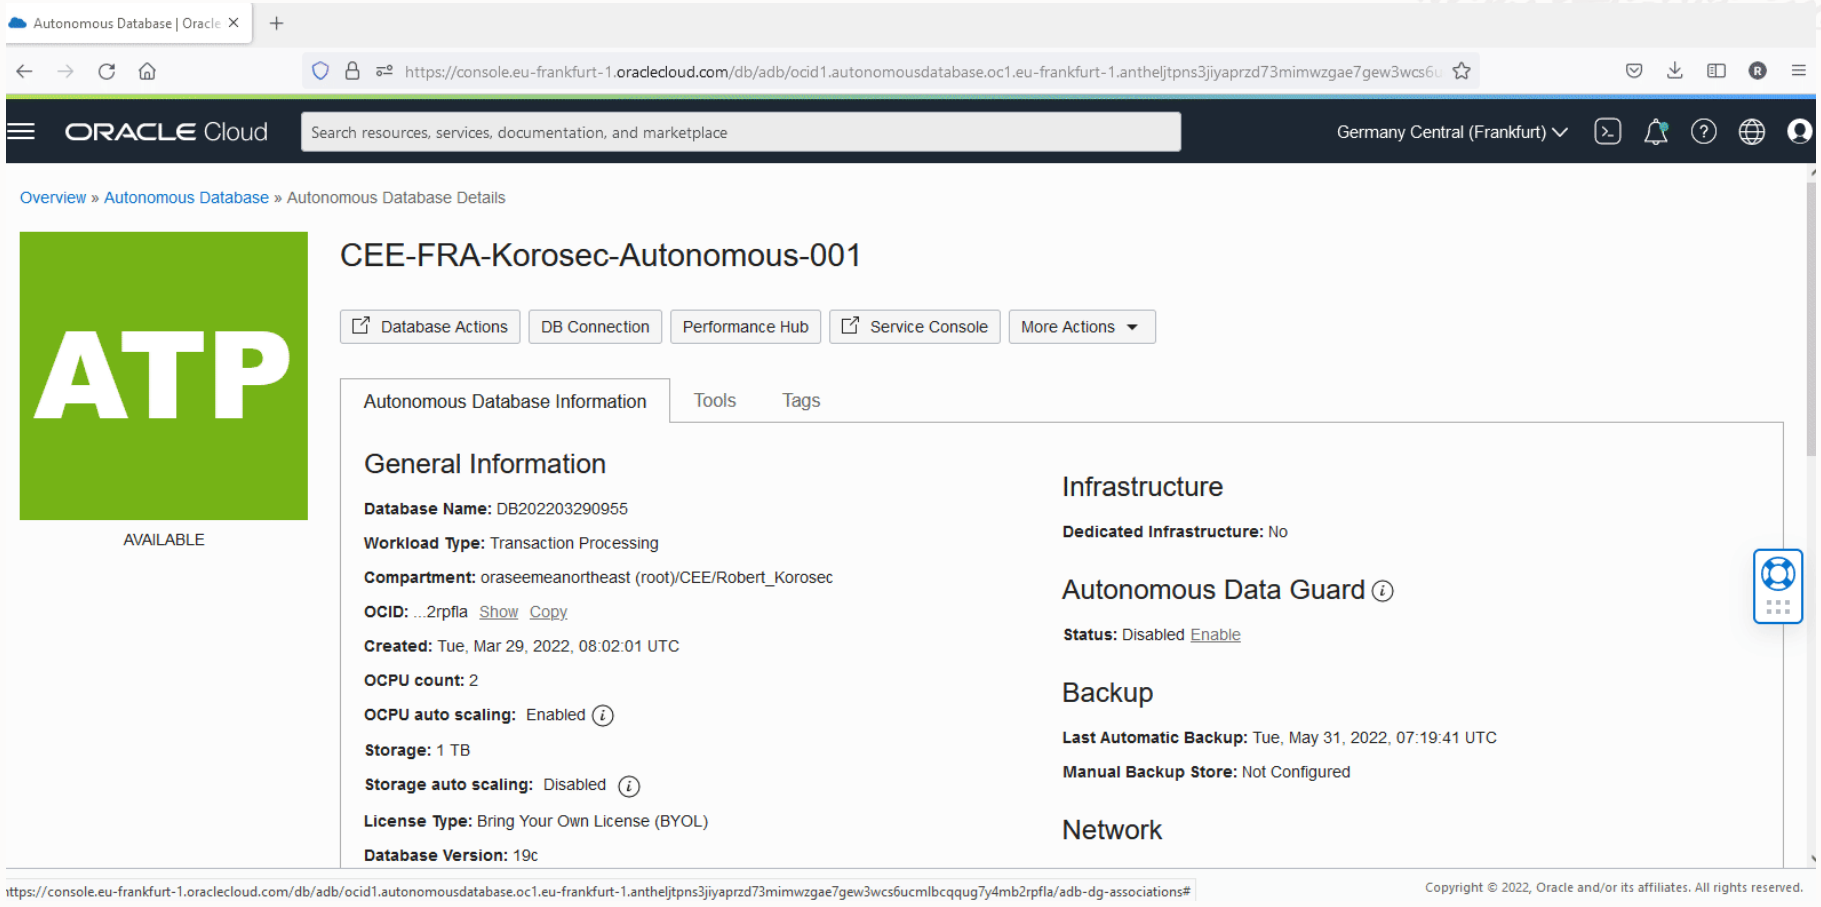

Cloud automation for

- Create
- Delete
- Switchover
- Failover
- Reinstate

CLI-support to configure additional Data Guard attributes

Enable Disaster Protection With a Single Click

Self-Repairing | Automated cloud operations

Pro-active detection and resolution of incidents

Autonomous Database detects over 85% of issues before customer

Continuous monitoring for each database: 8000+ metrics and 1500+ alarms

- Much broader than many on-premises customers
- Consolidated monitoring of entire stack: infrastructure, load balancer, connection manager, database, ORDS, APEX, OML

Automatic service requests (SRs) are generated for each deviation

- Immediate investigation and resolution by cloud ops
- Root cause analysis for every issue
- Zero customer actions required

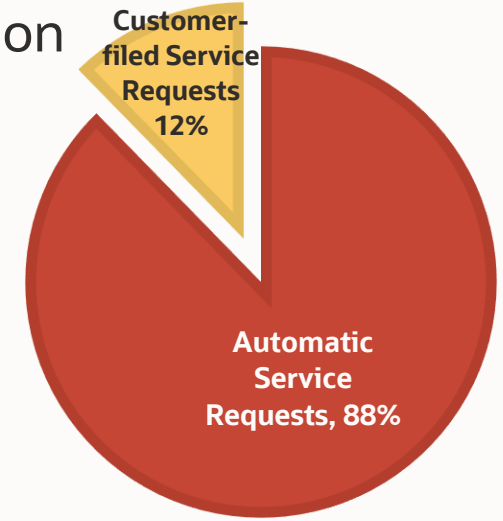

7 out of every 8 issues are automatically resolved with zero customer actions required

Significant benefits: Five-year ROI of 417%

The real-world business value of Oracle Autonomous Data Warehouse

Business value highlights

- **417%** five-year ROI
- **63%** reduced total cost of operations
- **Five months** to payback
- **68%** more efficient database administrators
- **84%** more efficient IT infrastructure management
- **45%** reduction in IT infrastructure costs
- **94%** reduction in unplanned downtime
- **27%** more productive

IDC
ANALYZE THE FUTURE

TABLE 8

Five-Year ROI Analysis

	Per Organization	Per Database
Benefit (discounted)	$7.42M	$536.9K
Investment (discounted)	$1.44M	$103.8K
Net Present Value	$5.99M	$433.9K
ROI (NPV/Investment)	417%	417%
Payback (Months)	5 months	5 months
Discount Factor	12%	12%

n = 7, Source: IDC In-depth Interviews, December 2020

Oracle Autonomous Database Service
Under the hood – A converged database

- Multitenant for Efficient, Agile Database Clouds
- AutoML for simple integrated Machine Learning
- In-Memory for Database Acceleration
- Native JSON for Document Data
- In-Memory Ingest for Fastest IoT
- Cloud SQL for integrating Object Store Data Lake
- Persistent Memory Store for Lowest Latency
- Spatial and Graph for Mapping and Social Networks
- Licence
  - Subscription (Licence Included) – Elastic, €1,14679 €/OCPU/hour
  - BYOL, €0,27524/OCPU/hour -> €204/OCPU/mo.
    - DB EE + Multitenant -> 2 OCPUs (<16 OCPUs; RAC otherwise)
    - Active DataGuard BYOL for Autonomous Dataguard
    - DB SE2 -> 4 OCPUs (max 8 OCPUs), no ADG, no Multitenant, no RAC

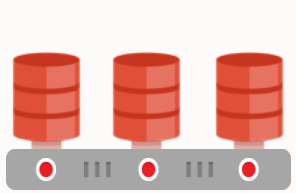
Multitenant


Machine Learning


In-Memory Analytics


JSON

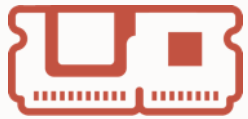
Persistent Memory


In-Memory IoT


Spatial


Graph


Cloud Integration

Exadata Cloud@Customer (on-premise installation)

Multi Customer Use Cases

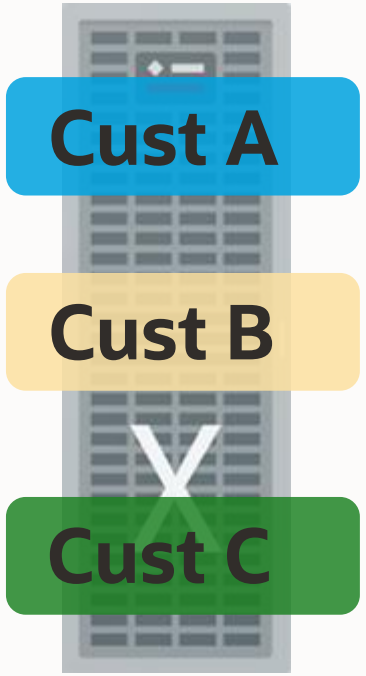

Multi-Customer Model Available To

- **Government Owned Clouds**
- **Hosting Service Providers (HSP)**

License Types Permitted Under Multi-Customer Model

- **BYOL**
- **License Included**

Service Available:

- Exadata Database Service
- Autonomous Database service

“… we want all our Oracle Database customers running Autonomous Database …”

Larry Ellison

Oracle Applications Certified with Oracle ATP, Autonomous Transaction Processing

Oracle Applications Certified with Oracle ATP, Autonomous Transaction Processing

Oracle Product	Version Number	Notes/Links
Oracle Peoplesoft	PeopleSoft 9.2, PeopleTools 8.57, 8.58, 8.59 and above	Certified with Autonomous Database-Dedicated and Shared
PeopleTools	8.57.16 and 8.58.05 and above	Certified with Autonomous Database-Dedicated and Shared
Siebel CRM	21.8	Certified with Oracle Autonomous Transaction Processing through Shared (Autonomous Database-S) or Dedicated (Autonomous Database-D) Infrastructure
Oracle JD Edwards	Applications Release 9.2	Certified with Oracle Autonomous Database on both Dedicated and Shared Exadata Infrastructure on Oracle Cloud Infrastructure - FAQ (PDF)
Oracle JD Edwards	Tools Release 9.2.4.3	Certified with Oracle Autonomous Database on both Dedicated and Shared Exadata Infrastructure on Oracle Cloud Infrastructure - FAQ (PDF)

References: Oracle internal Autonomous Deployments

https://www.oracle.com/autonomous-database/powered-by-autonomous-database/#accelerate

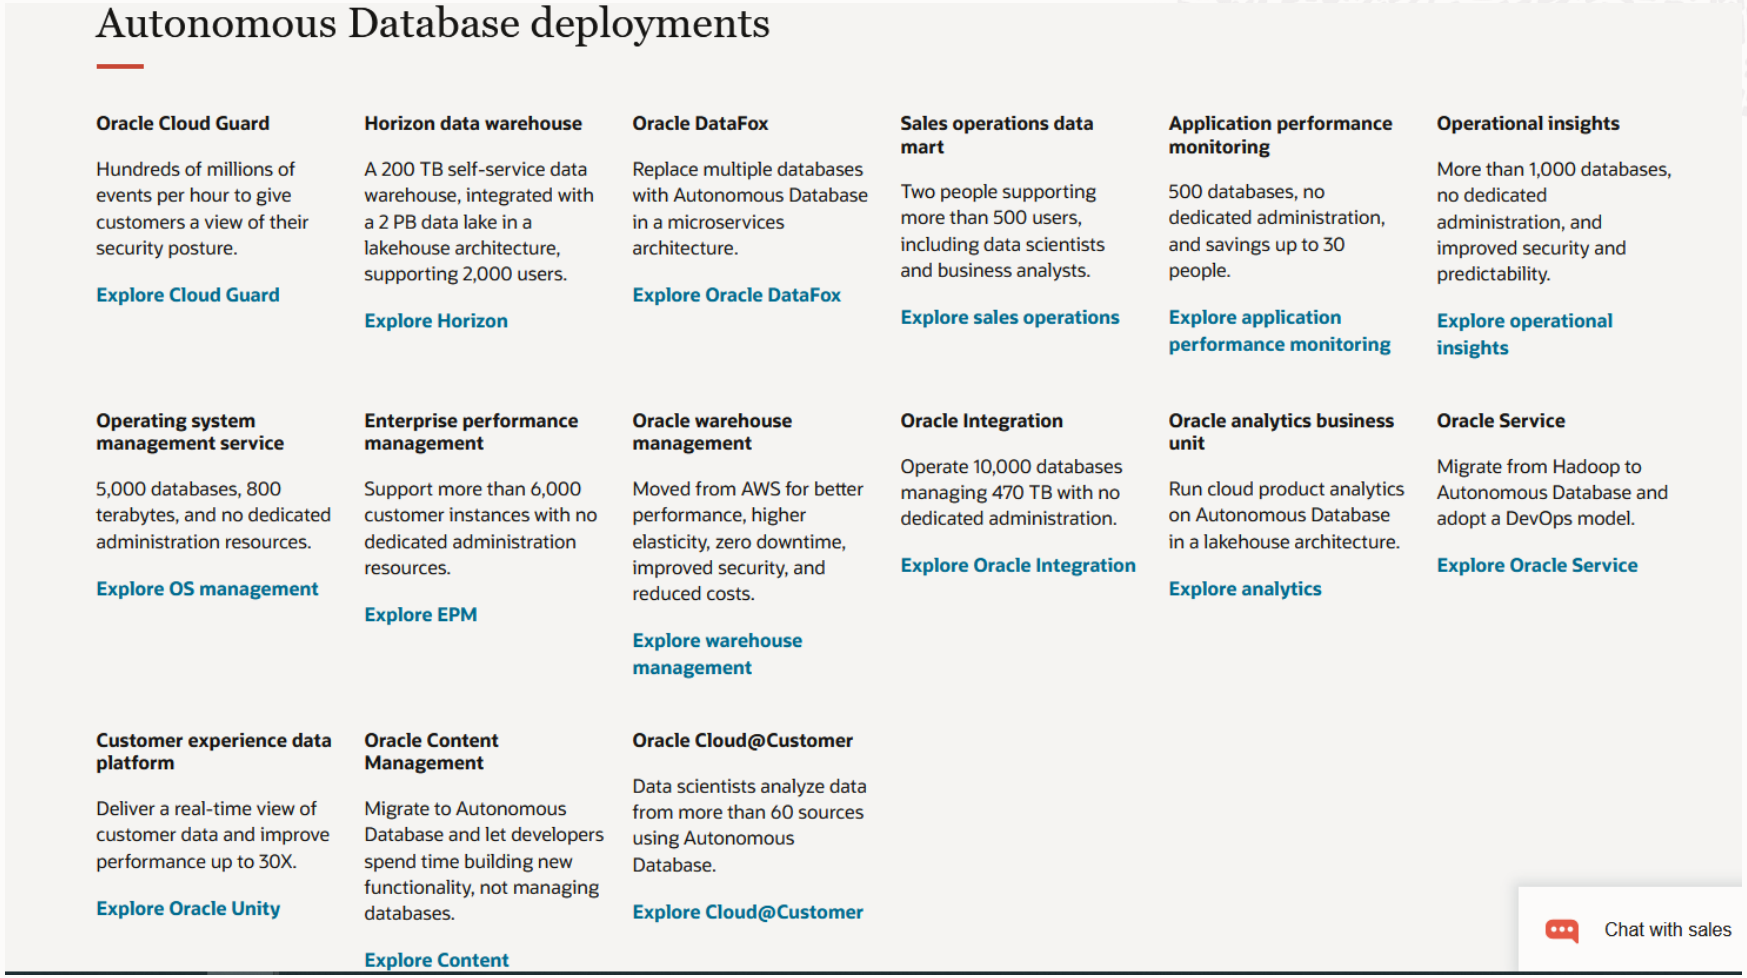

Autonomous Database deployments

Oracle Cloud Guard

Hundreds of millions of events per hour to give customers a view of their security posture.

Explore Cloud Guard

Horizon data warehouse

A 200 TB self-service data warehouse, integrated with a 2 PB data lake in a lakehouse architecture, supporting 2,000 users.

Explore Horizon

Oracle DataFox

Replace multiple databases with Autonomous Database in a microservices architecture.

Explore Oracle DataFox

Sales operations data mart

Two people supporting more than 500 users, including data scientists and business analysts.

Explore sales operations

Application performance monitoring

500 databases, no dedicated administration, and savings up to 30 people.

Explore application performance monitoring

Operational insights

More than 1,000 databases, no dedicated administration, and improved security and predictability.

Explore operational insights

Operating system management service

5,000 databases, 800 terabytes, and no dedicated administration resources.

Explore OS management

Enterprise performance management

Support more than 6,000 customer instances with no dedicated administration resources.

Explore EPM

Oracle warehouse management

Moved from AWS for better performance, higher elasticity, zero downtime, improved security, and reduced costs.

Explore warehouse management

Oracle Integration

Operate 10,000 databases managing 470 TB with no dedicated administration.

Explore Oracle Integration

Oracle analytics business unit

Run cloud product analytics on Autonomous Database in a lakehouse architecture.

Explore analytics

Oracle Service

Migrate from Hadoop to Autonomous Database and adopt a DevOps model.

Explore Oracle Service

Customer experience data platform

Deliver a real-time view of customer data and improve performance up to 30X.

Explore Oracle Unity

Oracle Content Management

Migrate to Autonomous Database and let developers spend time building new functionality, not managing databases.

Explore Content

Oracle Cloud@Customer

Data scientists analyze data from more than 60 sources using Autonomous Database.

Explore Cloud@Customer

Chat with sales

Major new features for ADB-Shared over last 12 months

https://blogs.oracle.com/datawarehousing/category/dwh-newsletter

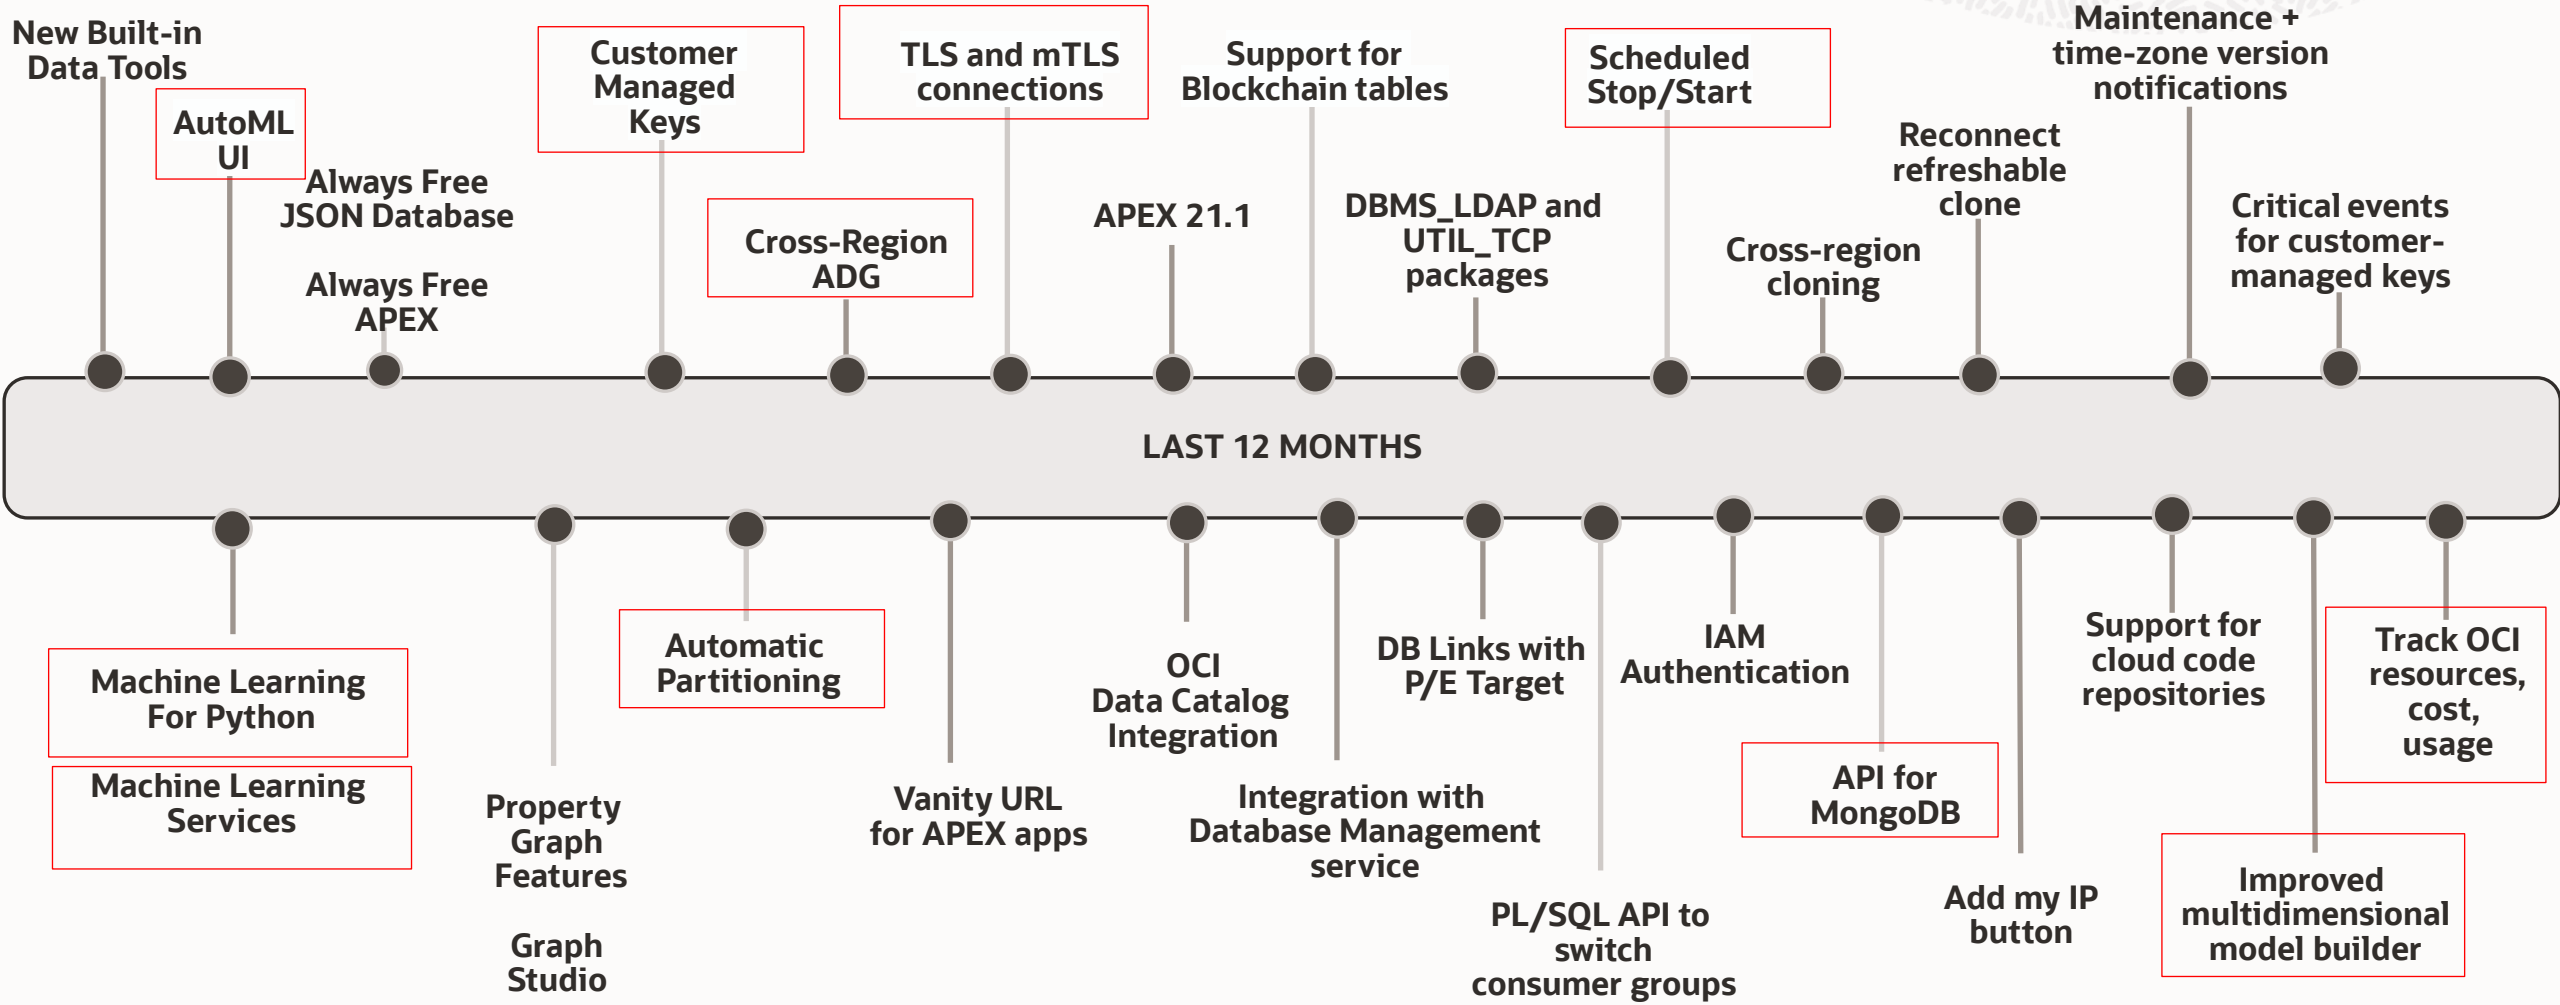

Create Database Links to Non-Oracle Databases

Free

db_type Value	Database Type
AWSREDSHIFT	Amazon Redshift
AZURE	Microsoft SQL Server Azure SQL Azure Synapse Analytics
MYSQL	MySQL
POSTGRES	PostgreSQL
SNOWFLAKE	Snowflake

Innovate More = Built-In, Self-Service Tools Expand Autonomous Vision

An Complete Ecosystem Empowering Business Users To Do More With Their Data

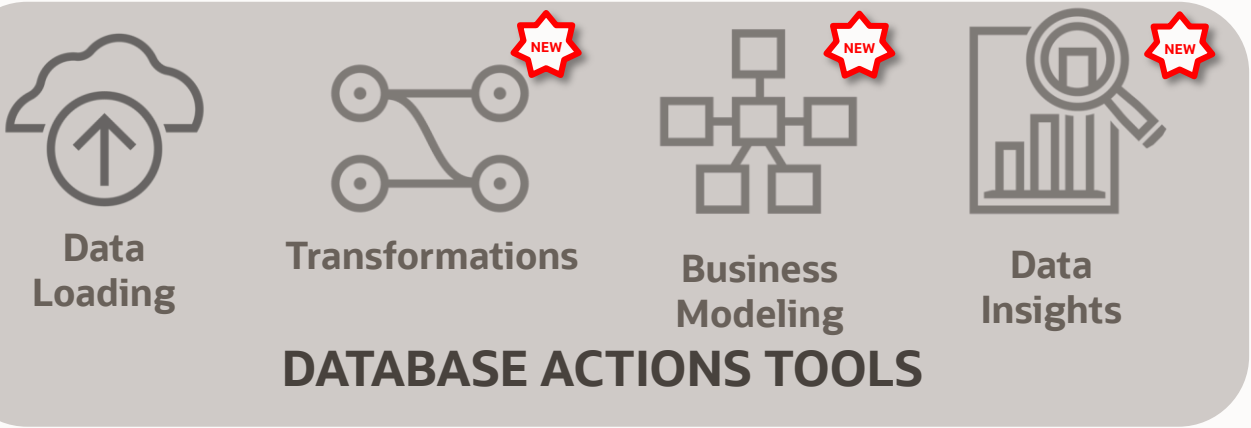

Oracle + 3rd Party Applications	Oracle + 3rd Party Databases	Files	Oracle + 3rd Party Streaming Services	Data Lakes over Object Stores

Database Actions – new suite of self-service tools for non-IT users

Load

Simple drag & drop loading (DEMO)

Transforms

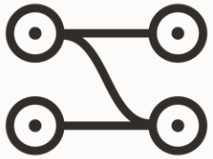

Declarative transformations and data cleansing (DEMO)

Data Analysis

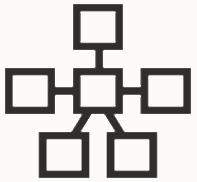

Automatically create powerful business models (DEMO)

Guided discovery of hidden patterns and anomalies (DEMO)

Graph

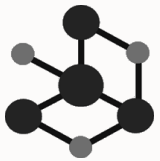

Prepare and explore graph data (DEMO)

ML Model

Easily create models with AutoML (DEMO)

Deploy

Integrate ML models into apps via REST or SQL (DEMO)

Build

Low-Code app-dev platform (APEX)

OML AutoML UI

No-code AutoML-based user interface supporting automated machine learning

Powerful, easy to use UI

Automates model building, tuning, and deployment

- Supports model management
- Enhance data scientist productivity
- Empower data professionals who are not ML experts

Featuring

- Minimal user input: data, target
- Model leaderboard
- Model deployment via REST endpoints
- Generate OML4Py notebooks from models

Multiple algorithms compared and best is selected

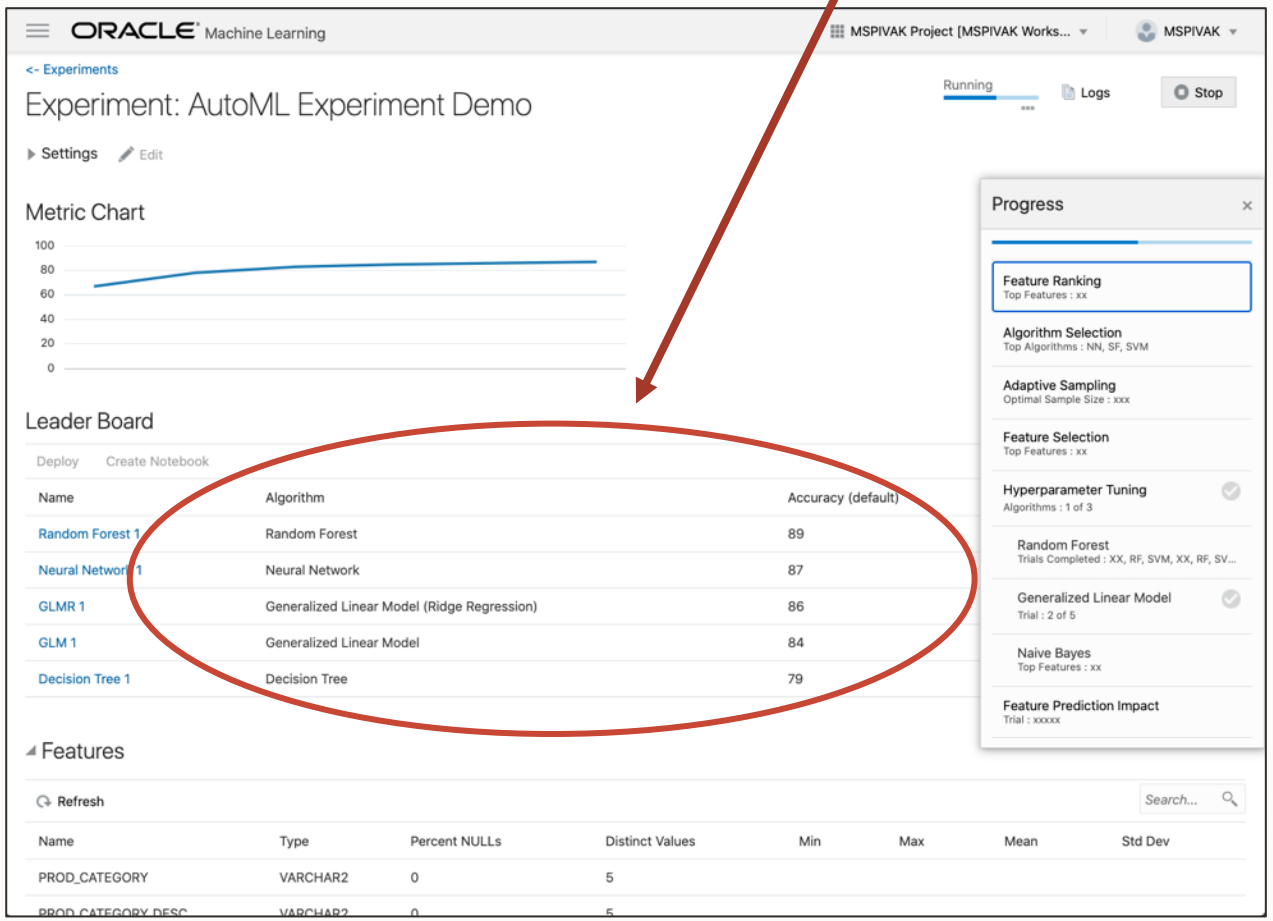

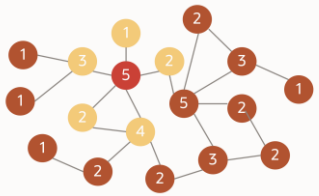

Graph Studio Now Fully GA

One-click start to analyzing with graphs in Oracle Autonomous Database

Graph Studio provides comprehensive set of features:

- Graph modeling tool to map relational data to graphs
- Launched directly from OCI Console

- Browser-based notebooks for interactive analysis and collaboration
- Integrated graph visualization
- PGQL: SQL-like property graph query language
- Nearly 60 pre-built property graph algorithms
  - PageRank, Community Detection, Shortest path, etc.

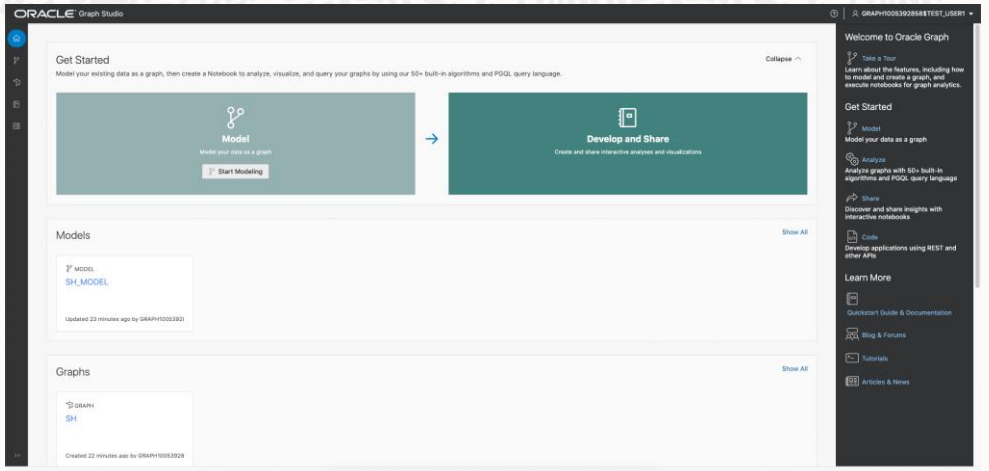

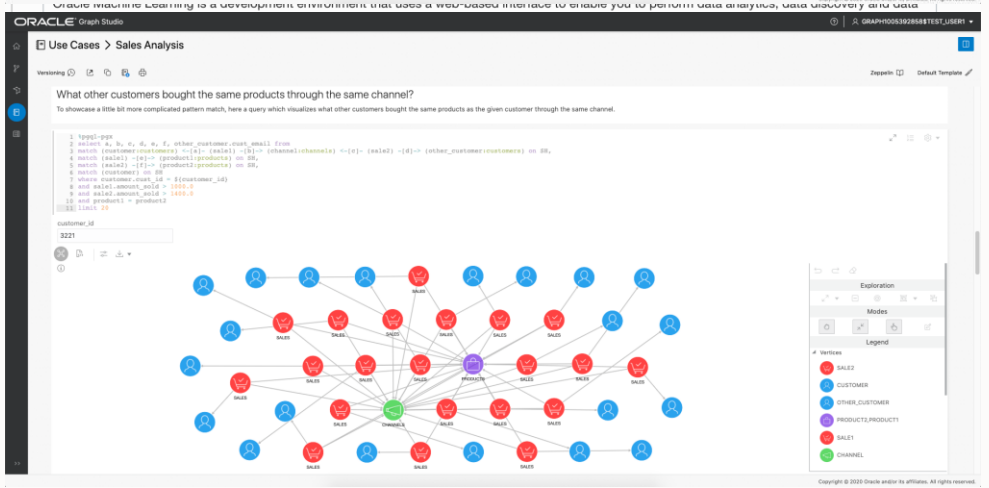

NEW Self-Service Tools for Data Analysts

From data to insights with built-in self-service data tools

Load

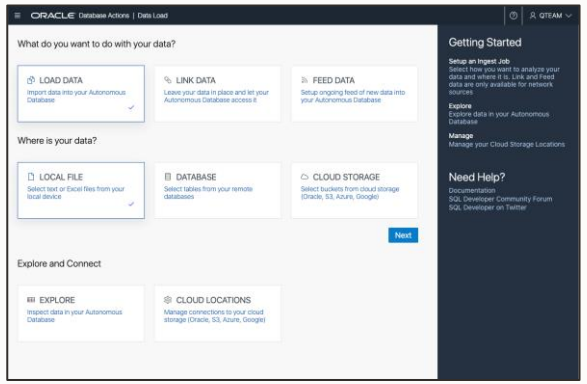

Simple drag & drop loading

NEW Transforms

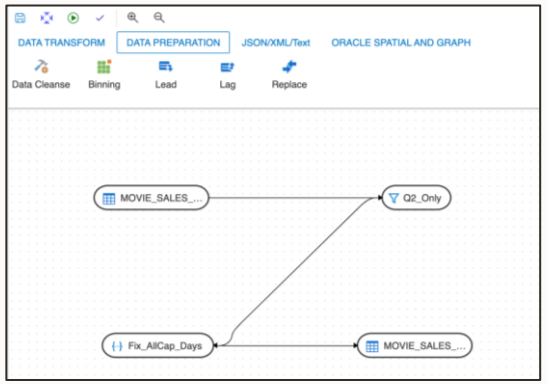

Declarative transformations and data cleansing

NEW Business Model

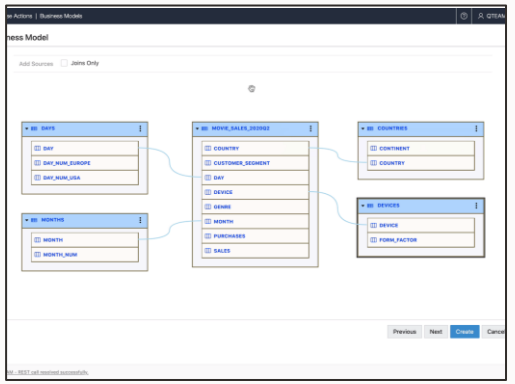

Automatically create business models

NEW Data Insights

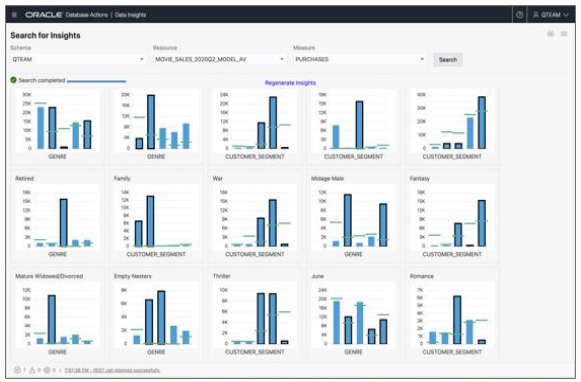

Automatically discover hidden patterns and anomalies

Built-in zero-code declarative data transformation tool

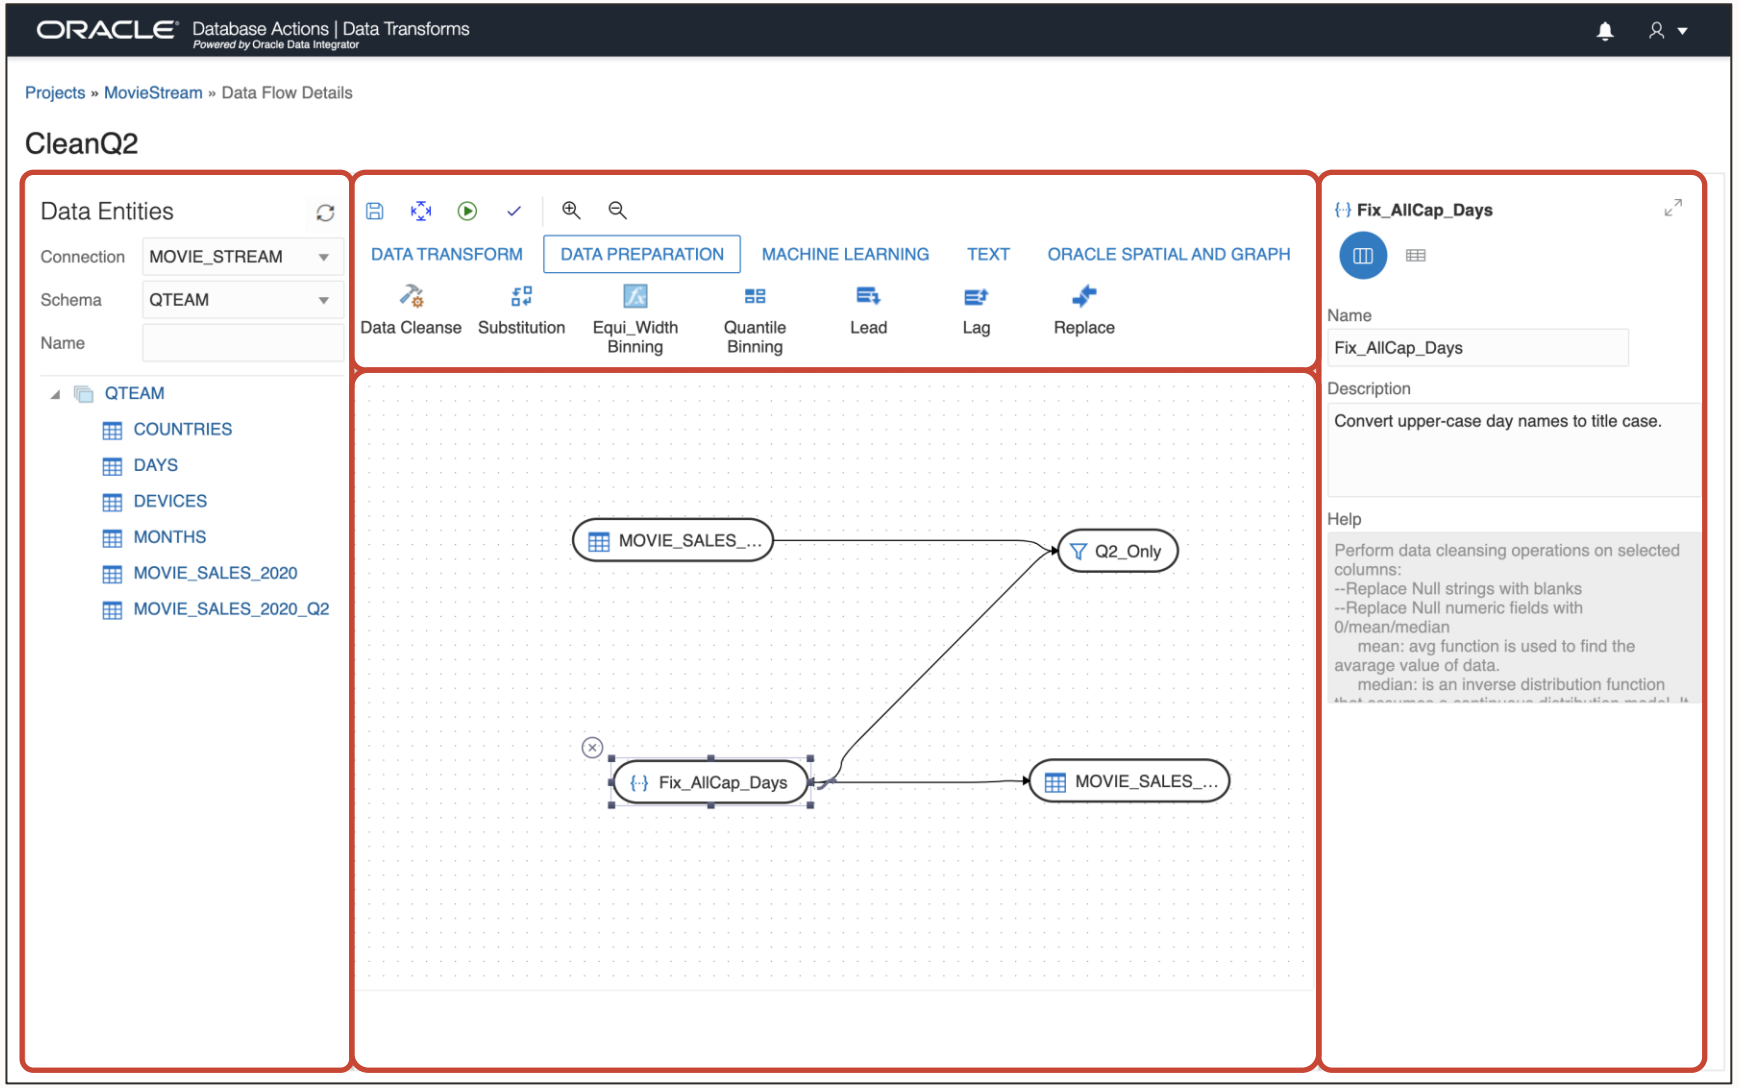

Simple “drag-and-drop” data transformations

- Auto code generation for all ODI sources and targets including Fusion, NetSuite and Salesforce
- Built-in data quality

Based on Oracle Data Integrator

- New, easy-to-use web UI
- Simple migration to ADB for ODI customers
- Exposes all built-in database operators

Data Analysis tool:
Machine learning driven anomaly and outlier detection

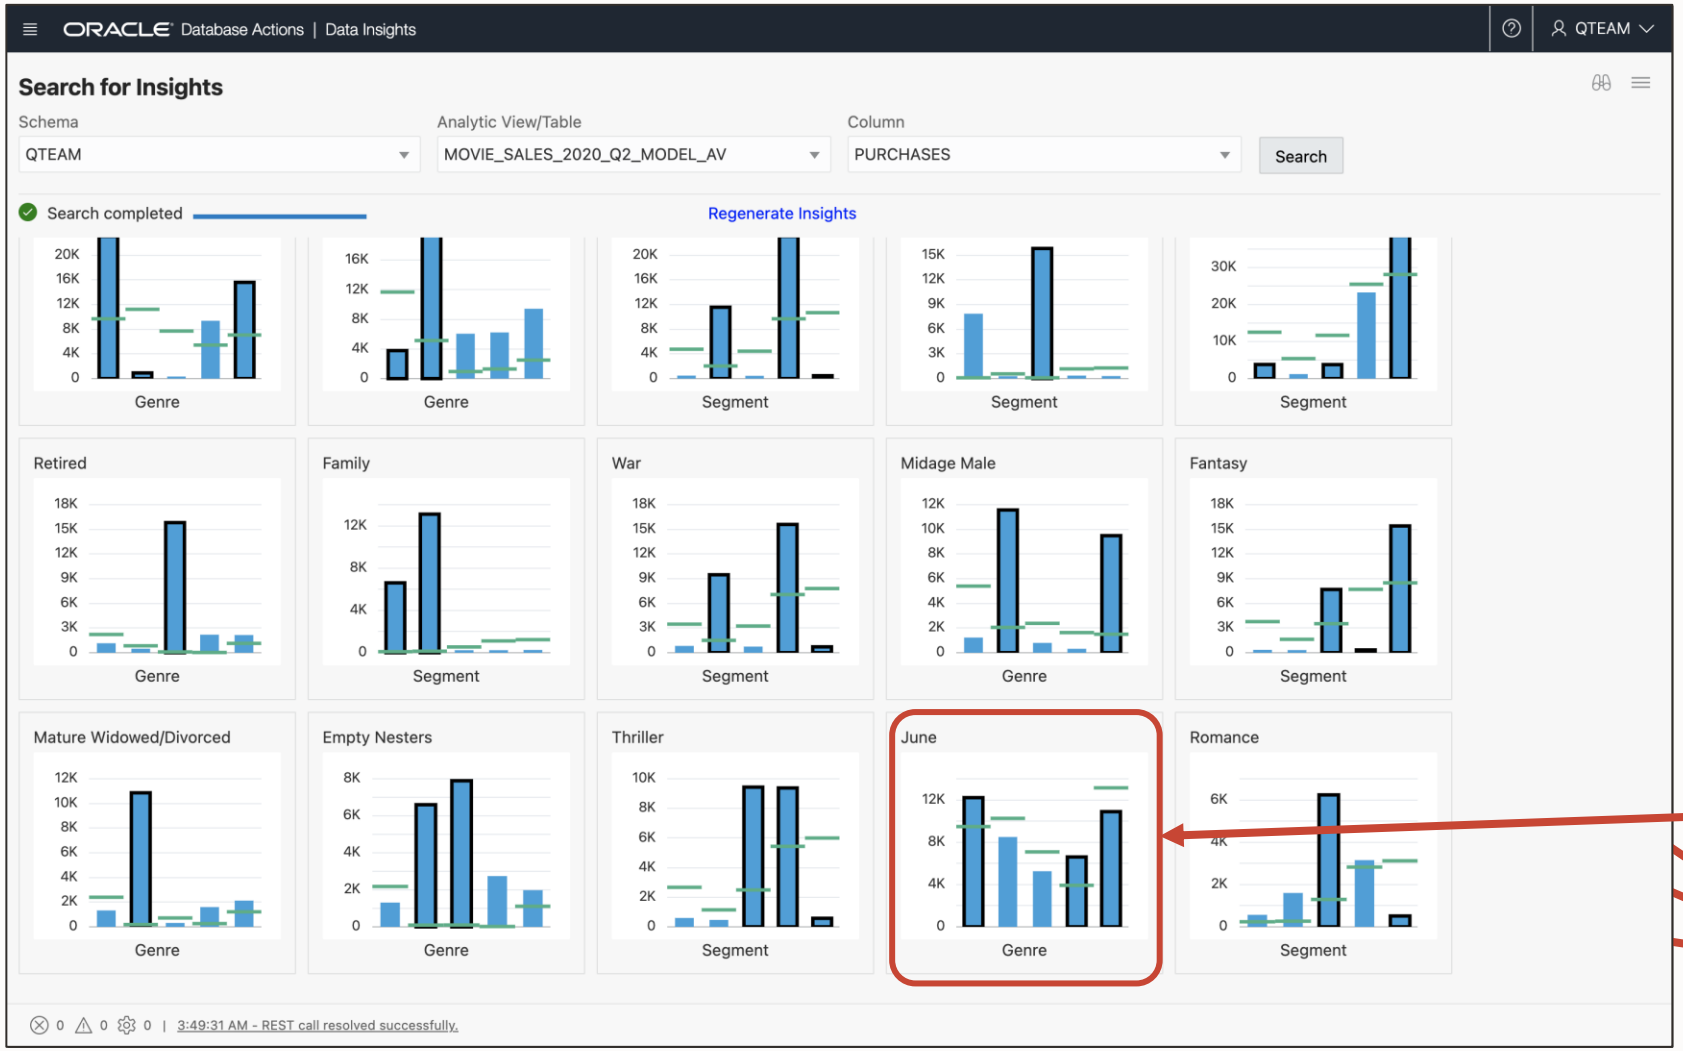

Automatic insight discovery

- Crawls over business model, running as background process
- Discovers hidden patterns, anomalies and outliers
- Variety of algorithms including regression slope

Usage

- Drill down on specific insight
- Significant differences between predicted and actual values highlighted

Data Analysis tool:
Built-in wizard to help build multidimensional models

Coming soon

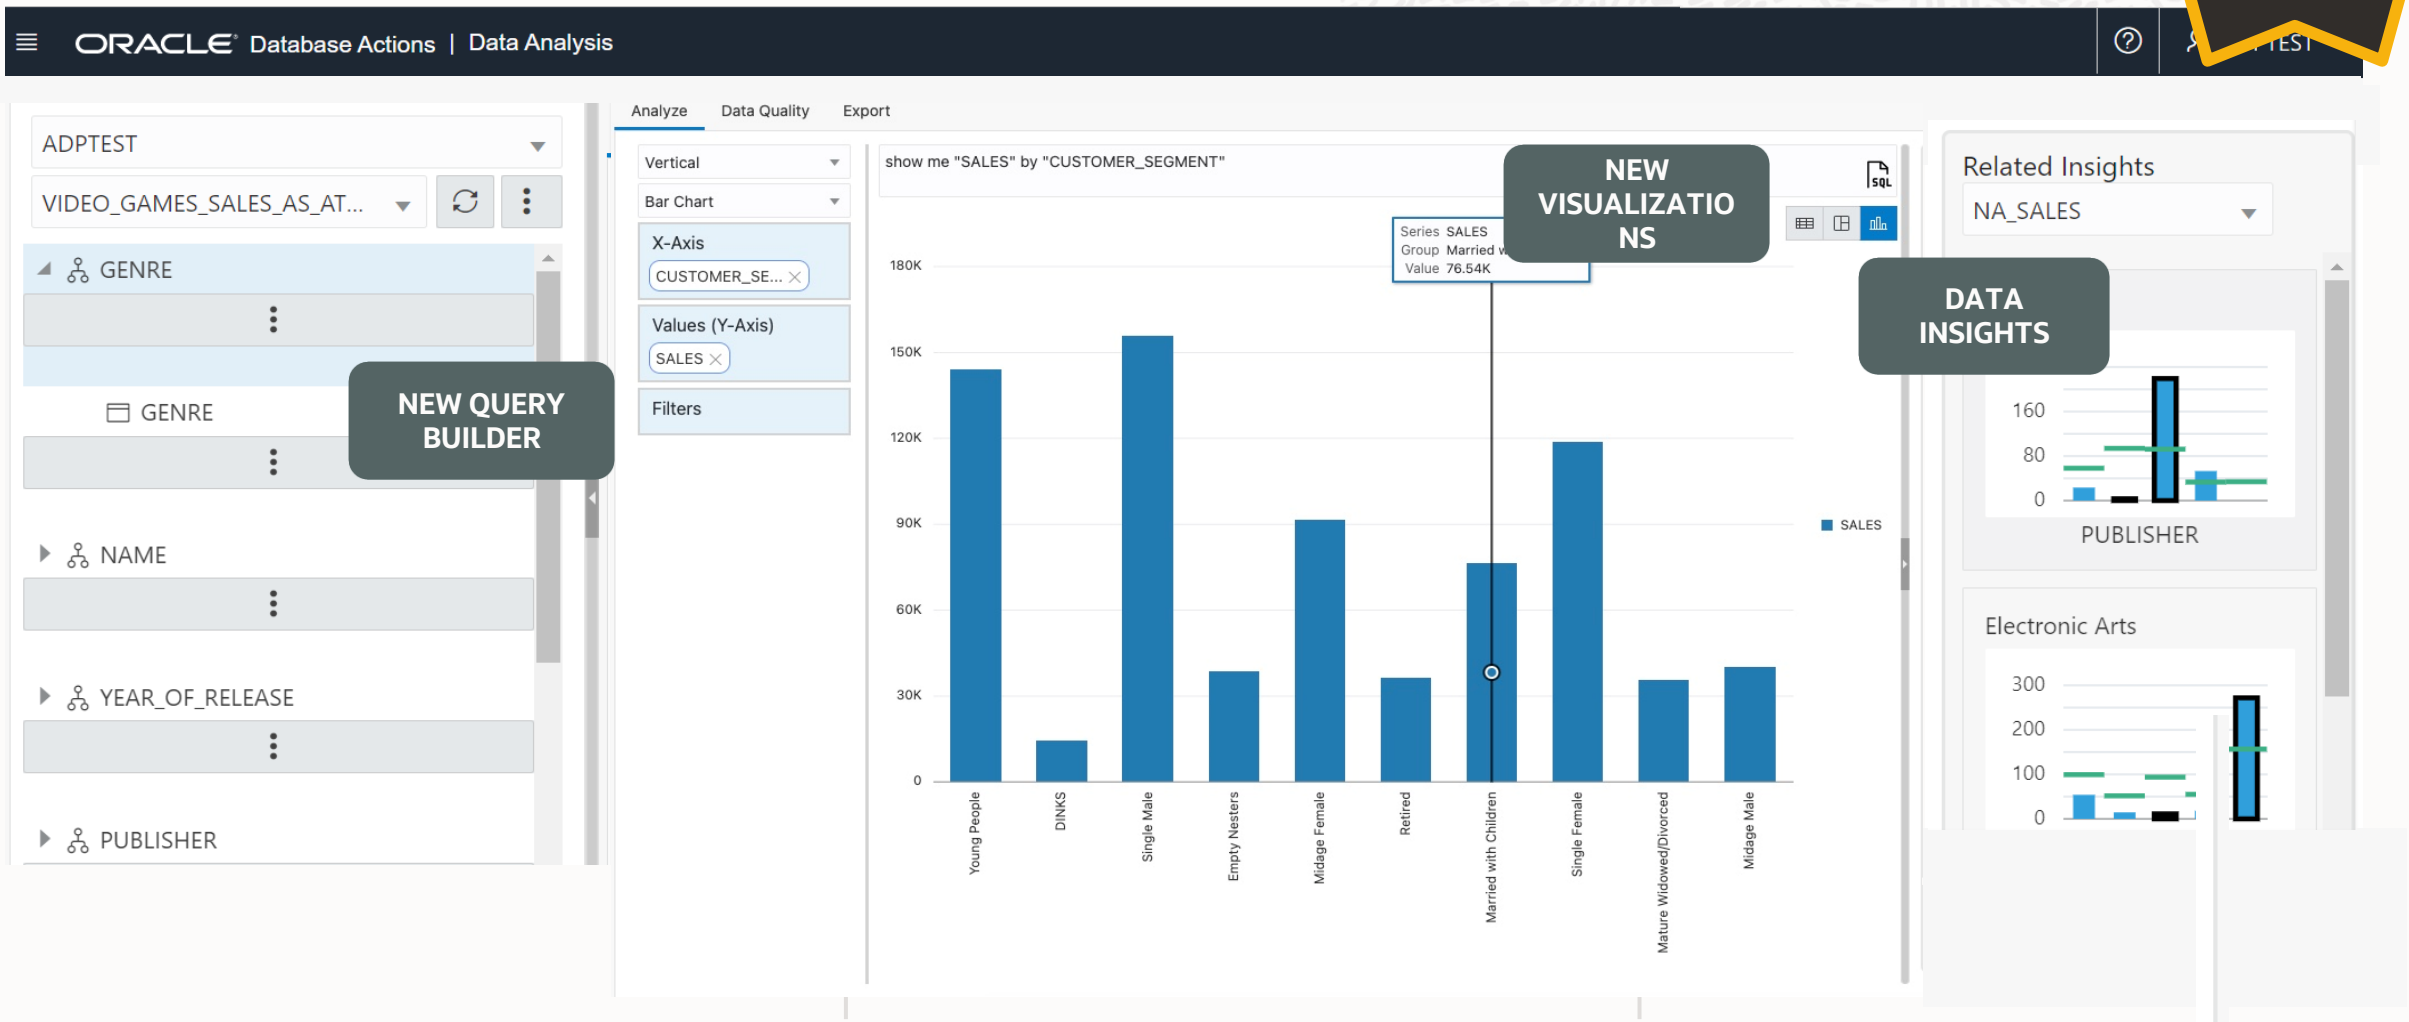

Data Analysis tool:
Directly test/query the multidimensional model with SQL

Coming soon

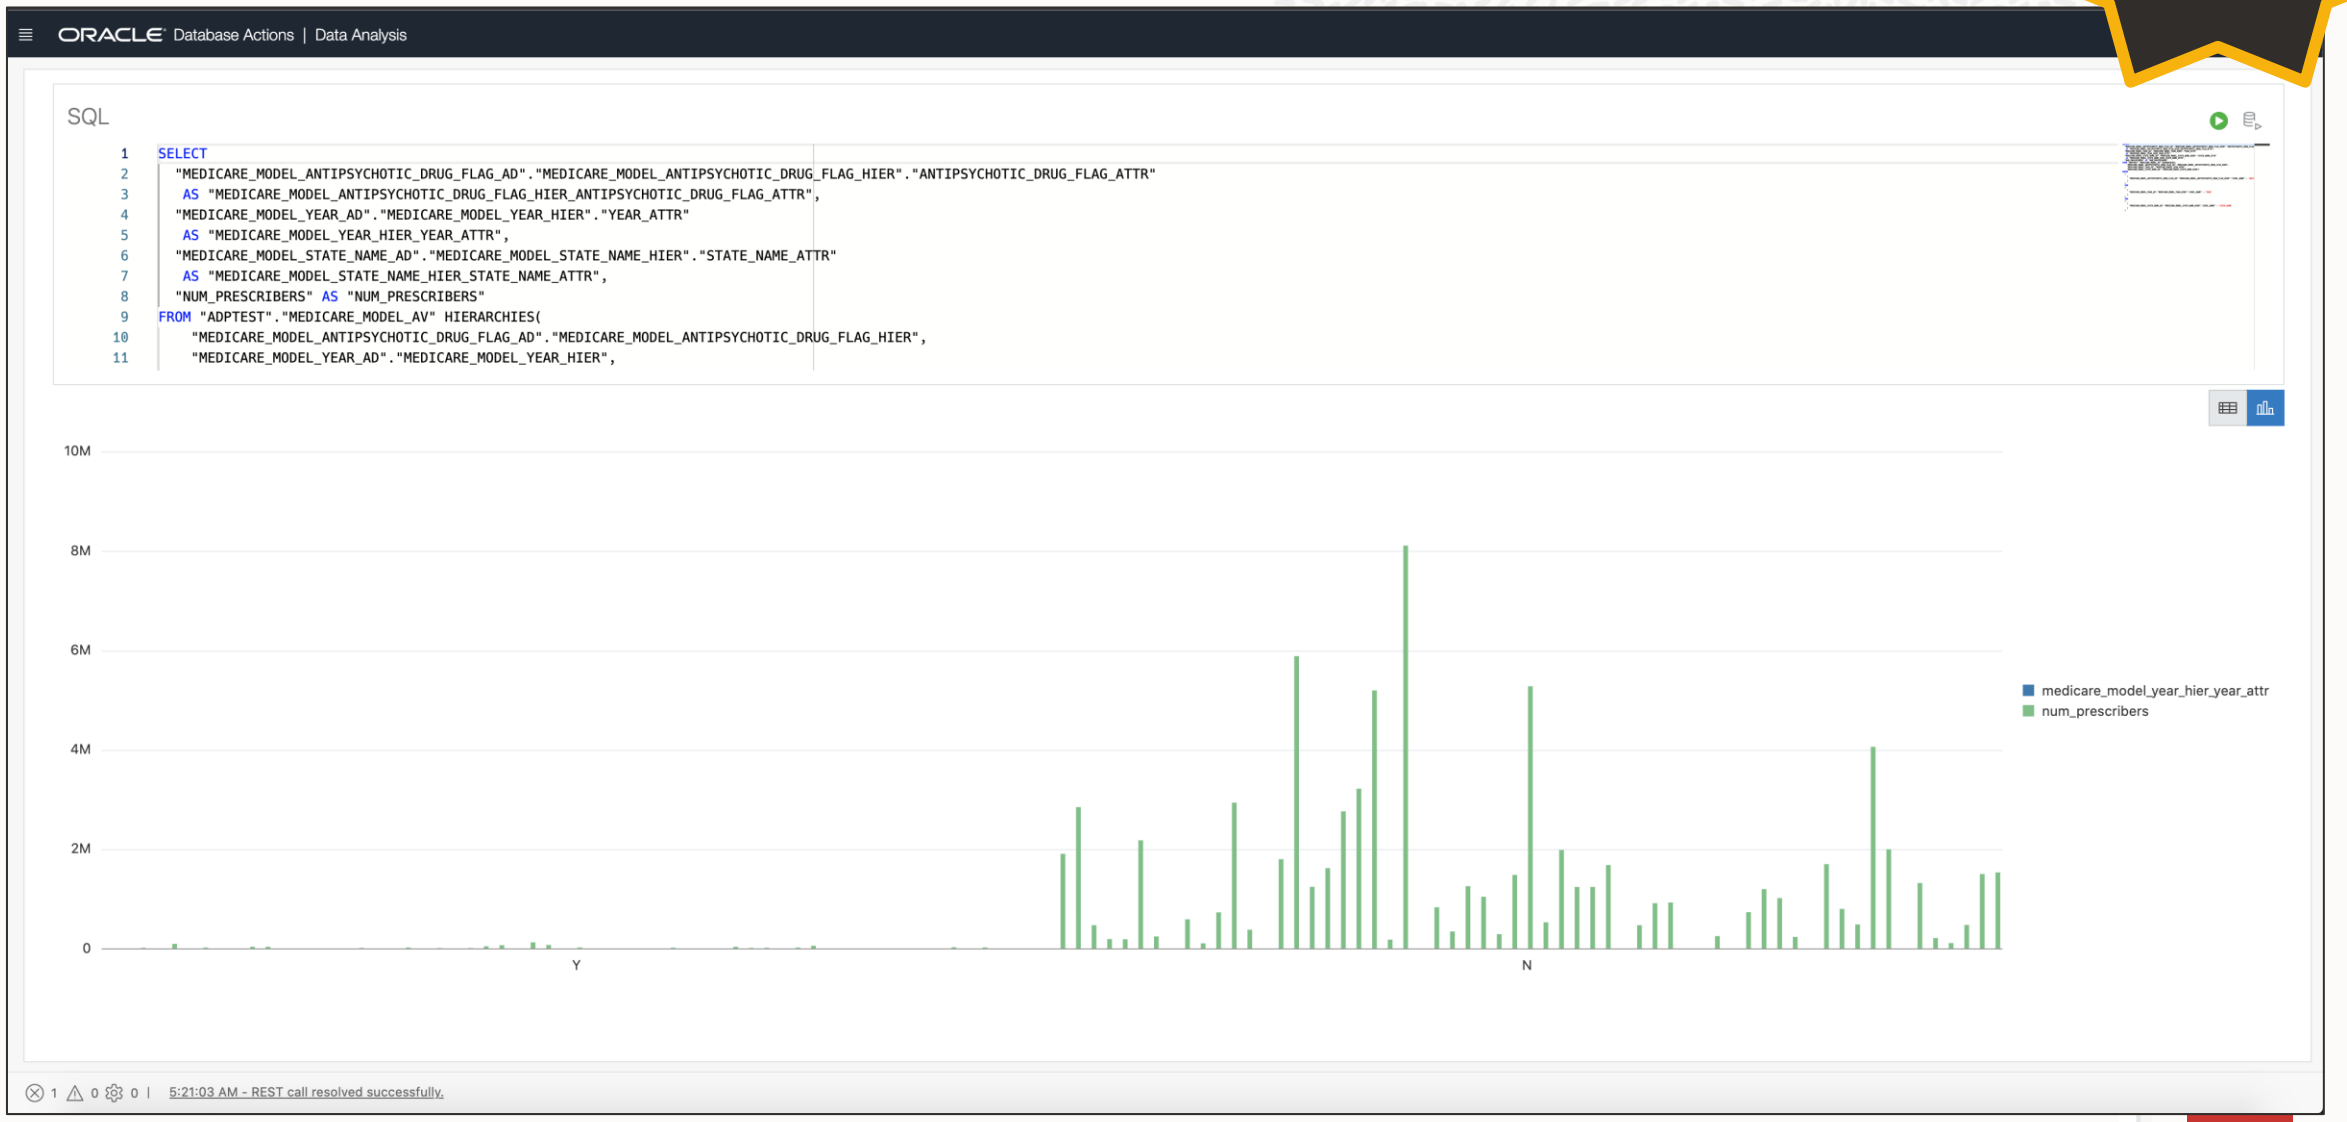

Data Analysis tool:
Export metadata for multidimensional model to Tableau

Coming soon

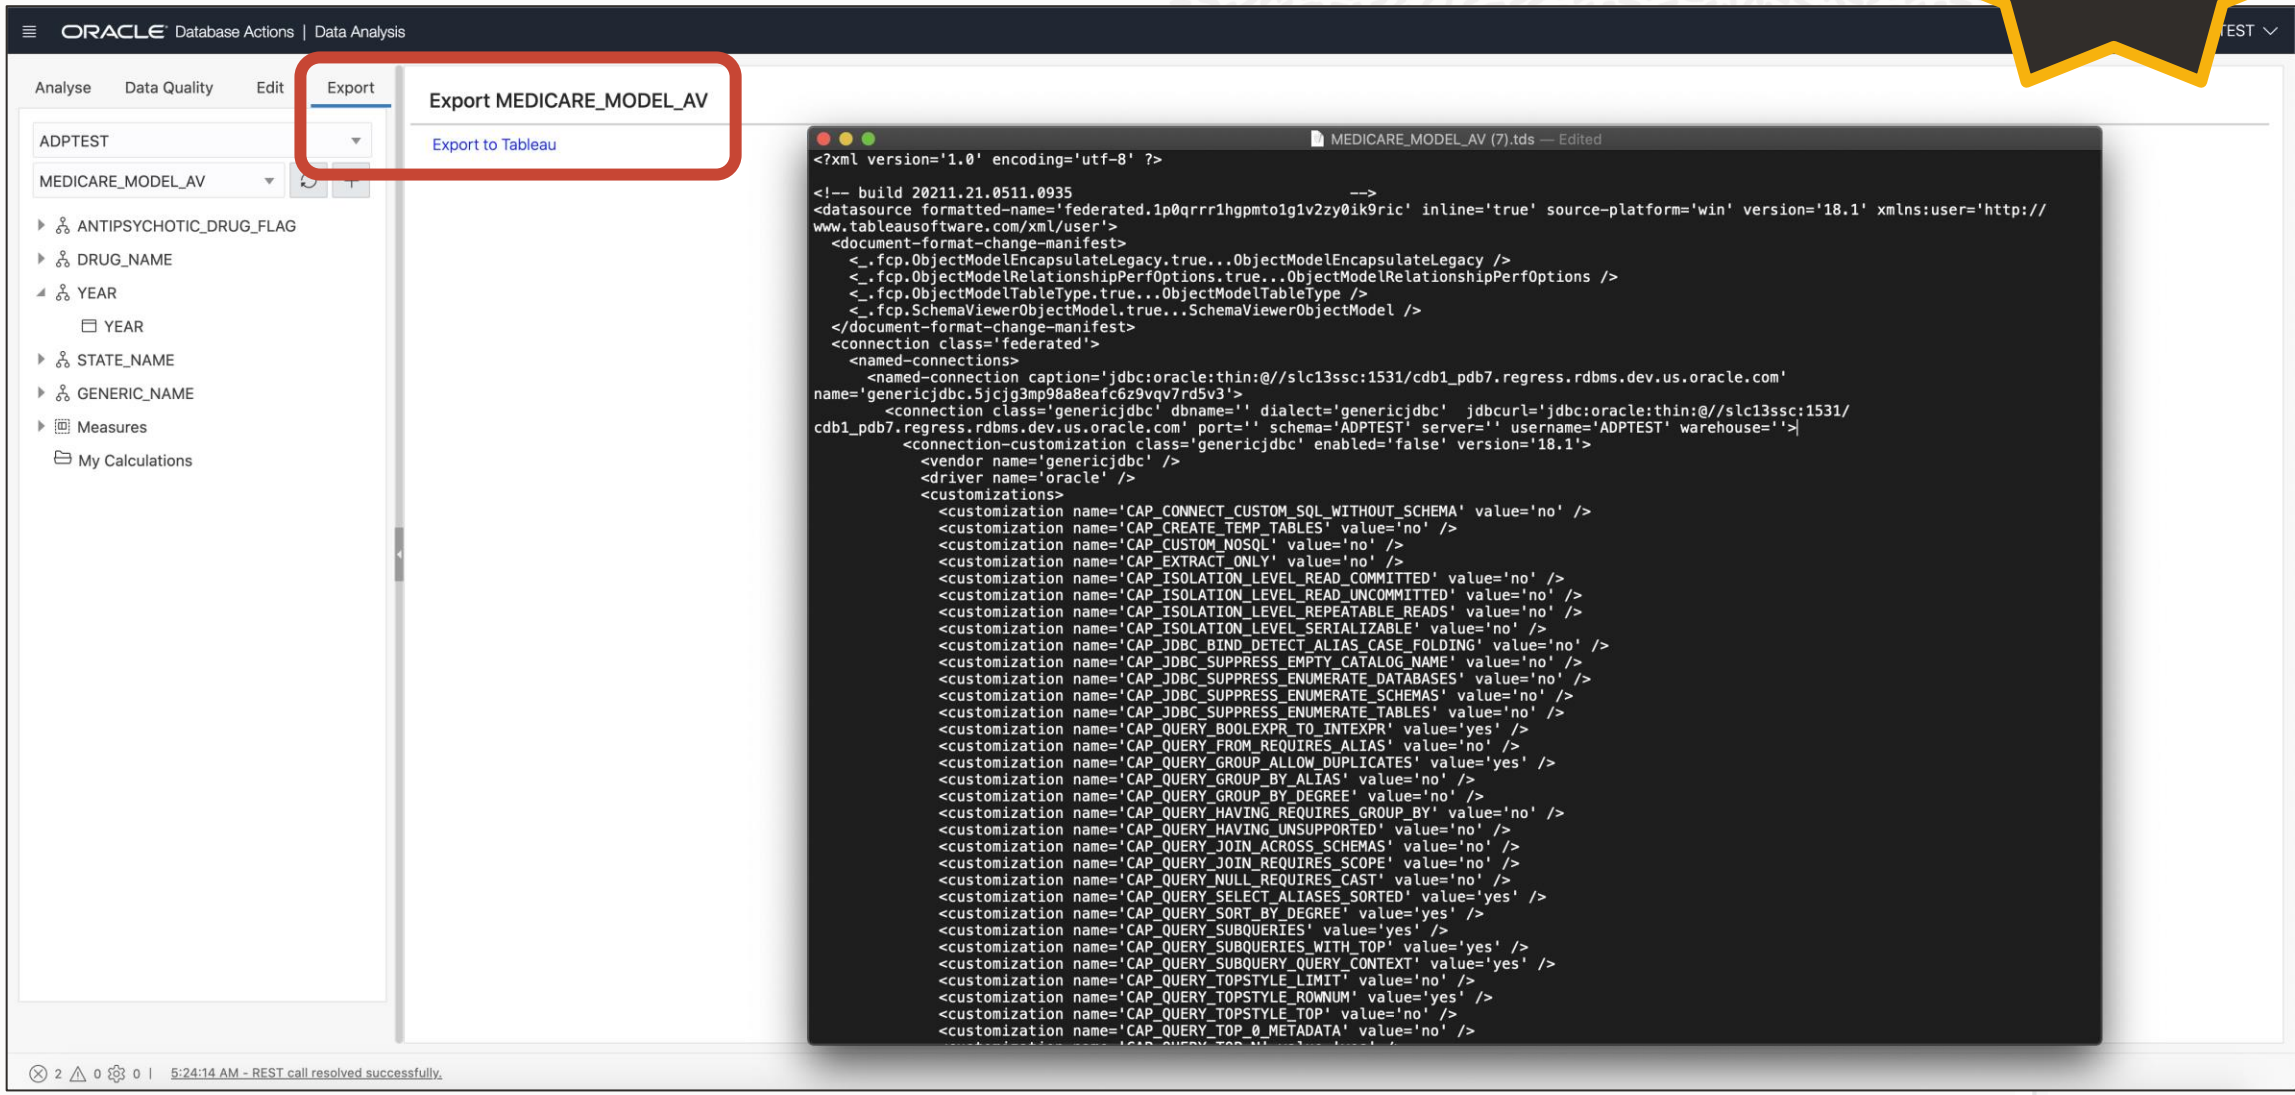

Import metadata and query model in Tableau

Coming soon

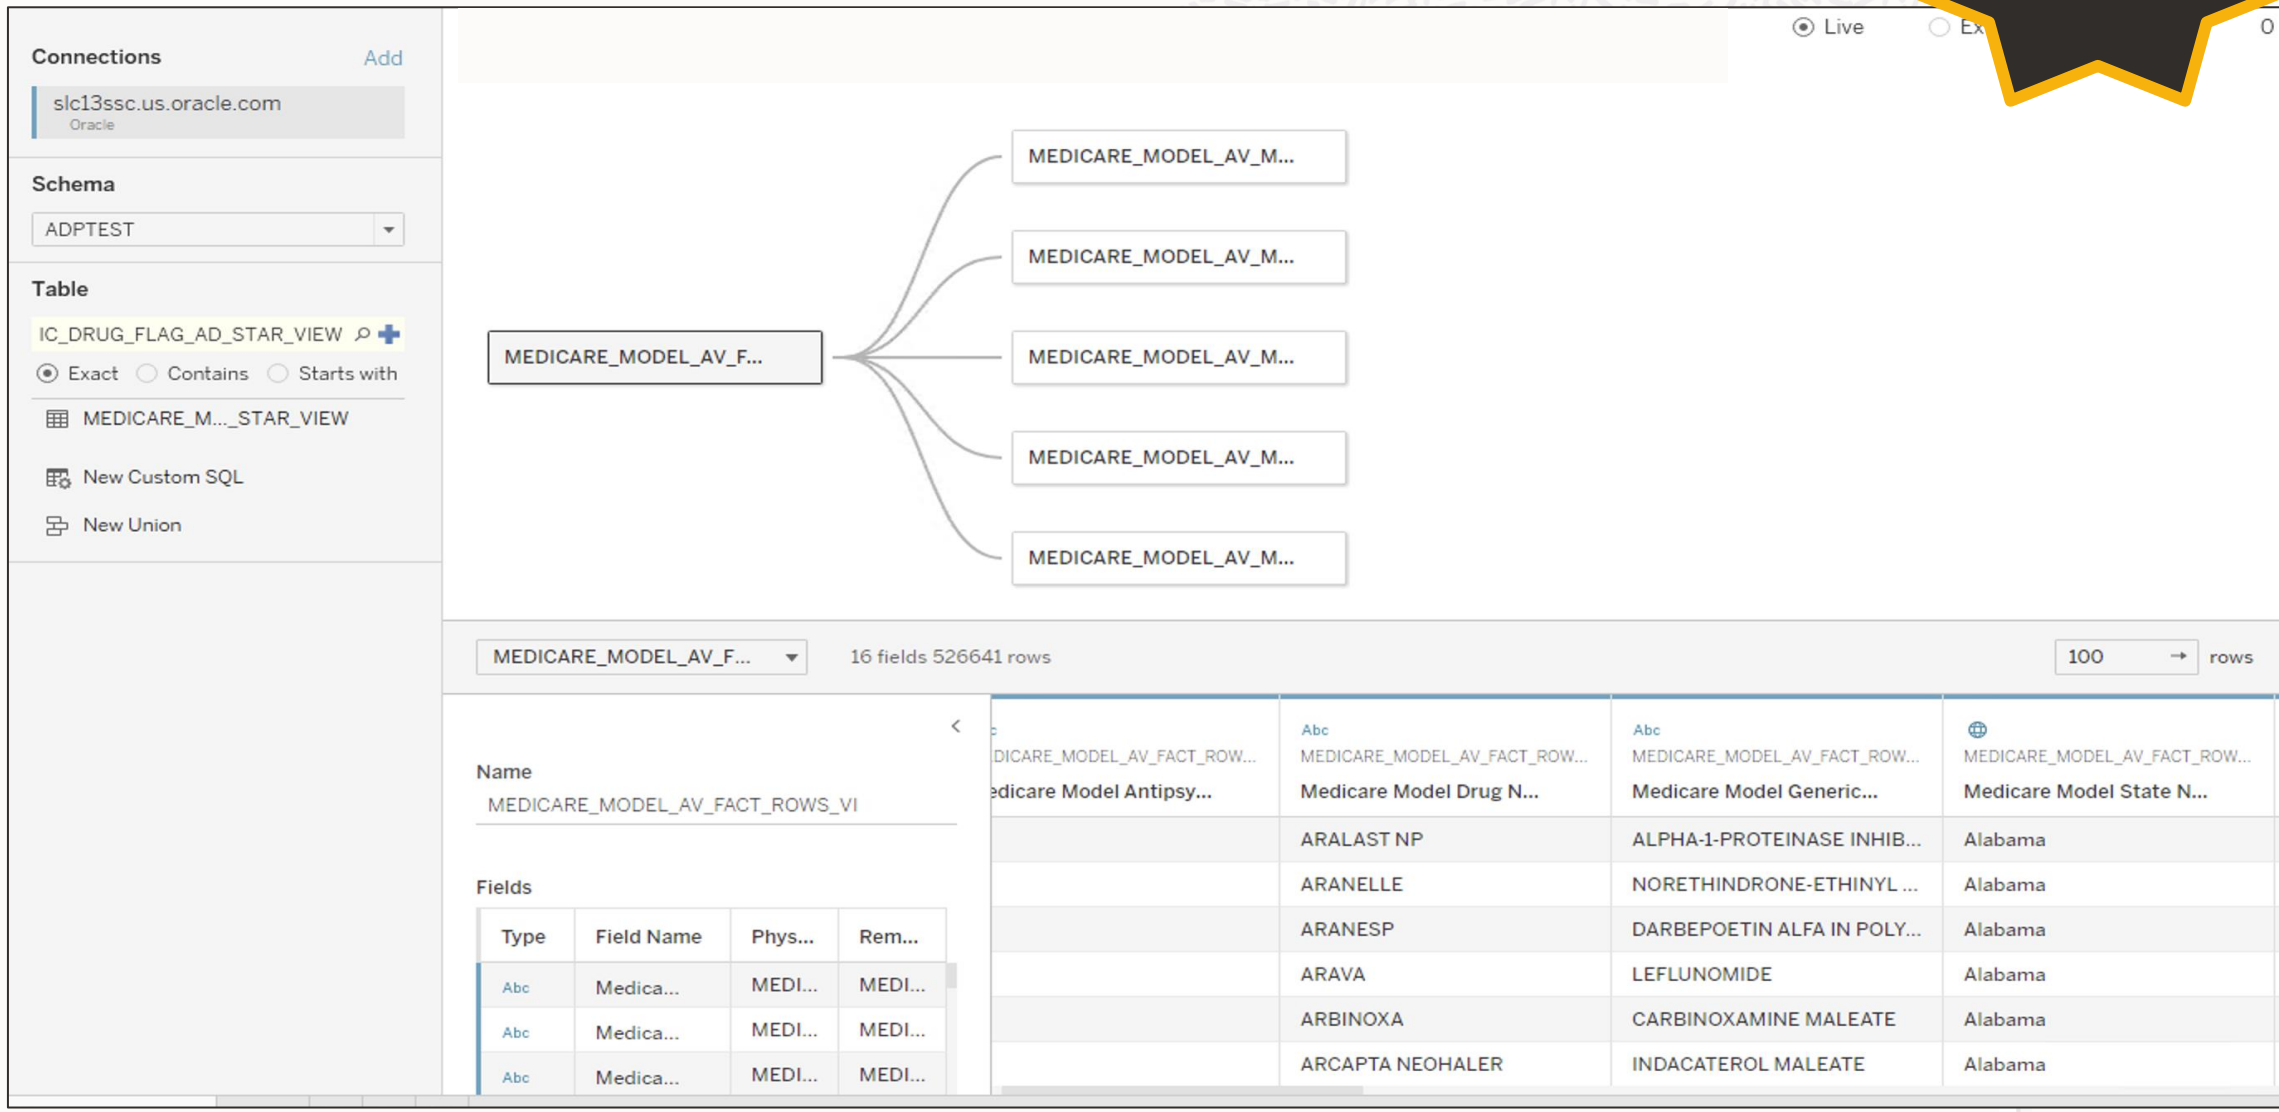

Directly read metadata and query model in OAC

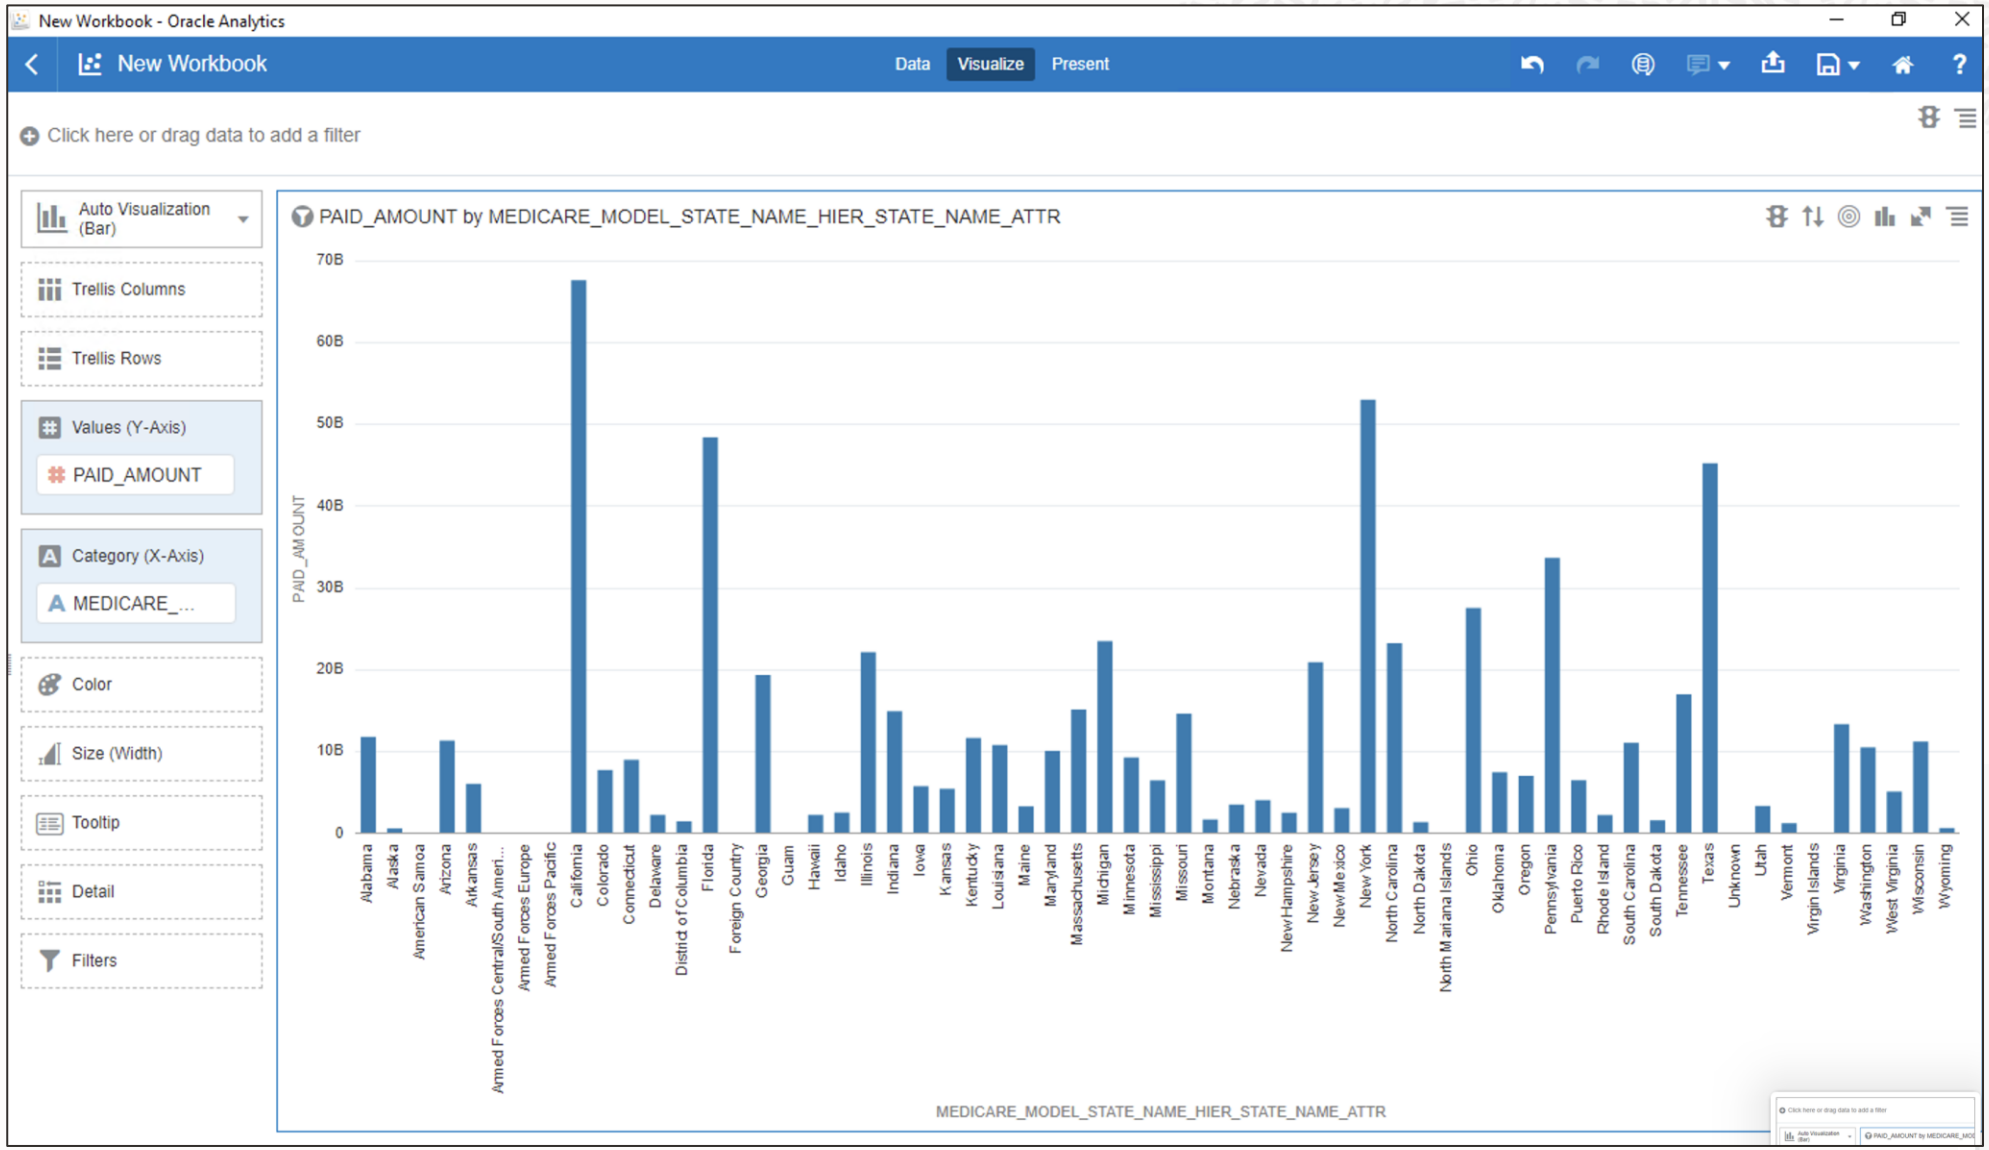

Data Lake Accelerator: Analyze Data At Scale

Object Store queries just get better / smarter

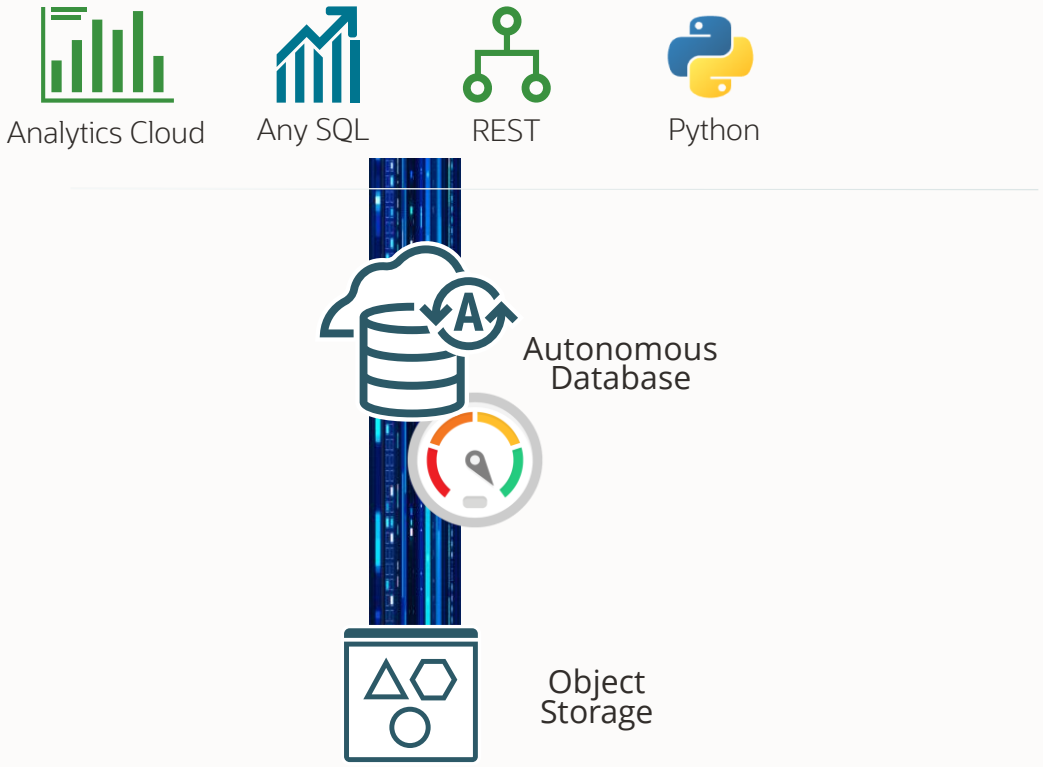

Scale out queries against Object Storage

- Specialized, object storage processing
- Scans, filters and aggregates data

Automatic and transparent

- Engages only when necessary
- Uses auto-scale to augment database compute for the life of the query

Reduce impact on your database workload

- Object store processing is isolated from database cores

Autonomous JSON Database

Better and cheaper than MongoDB Atlas

Oracle Converged Database

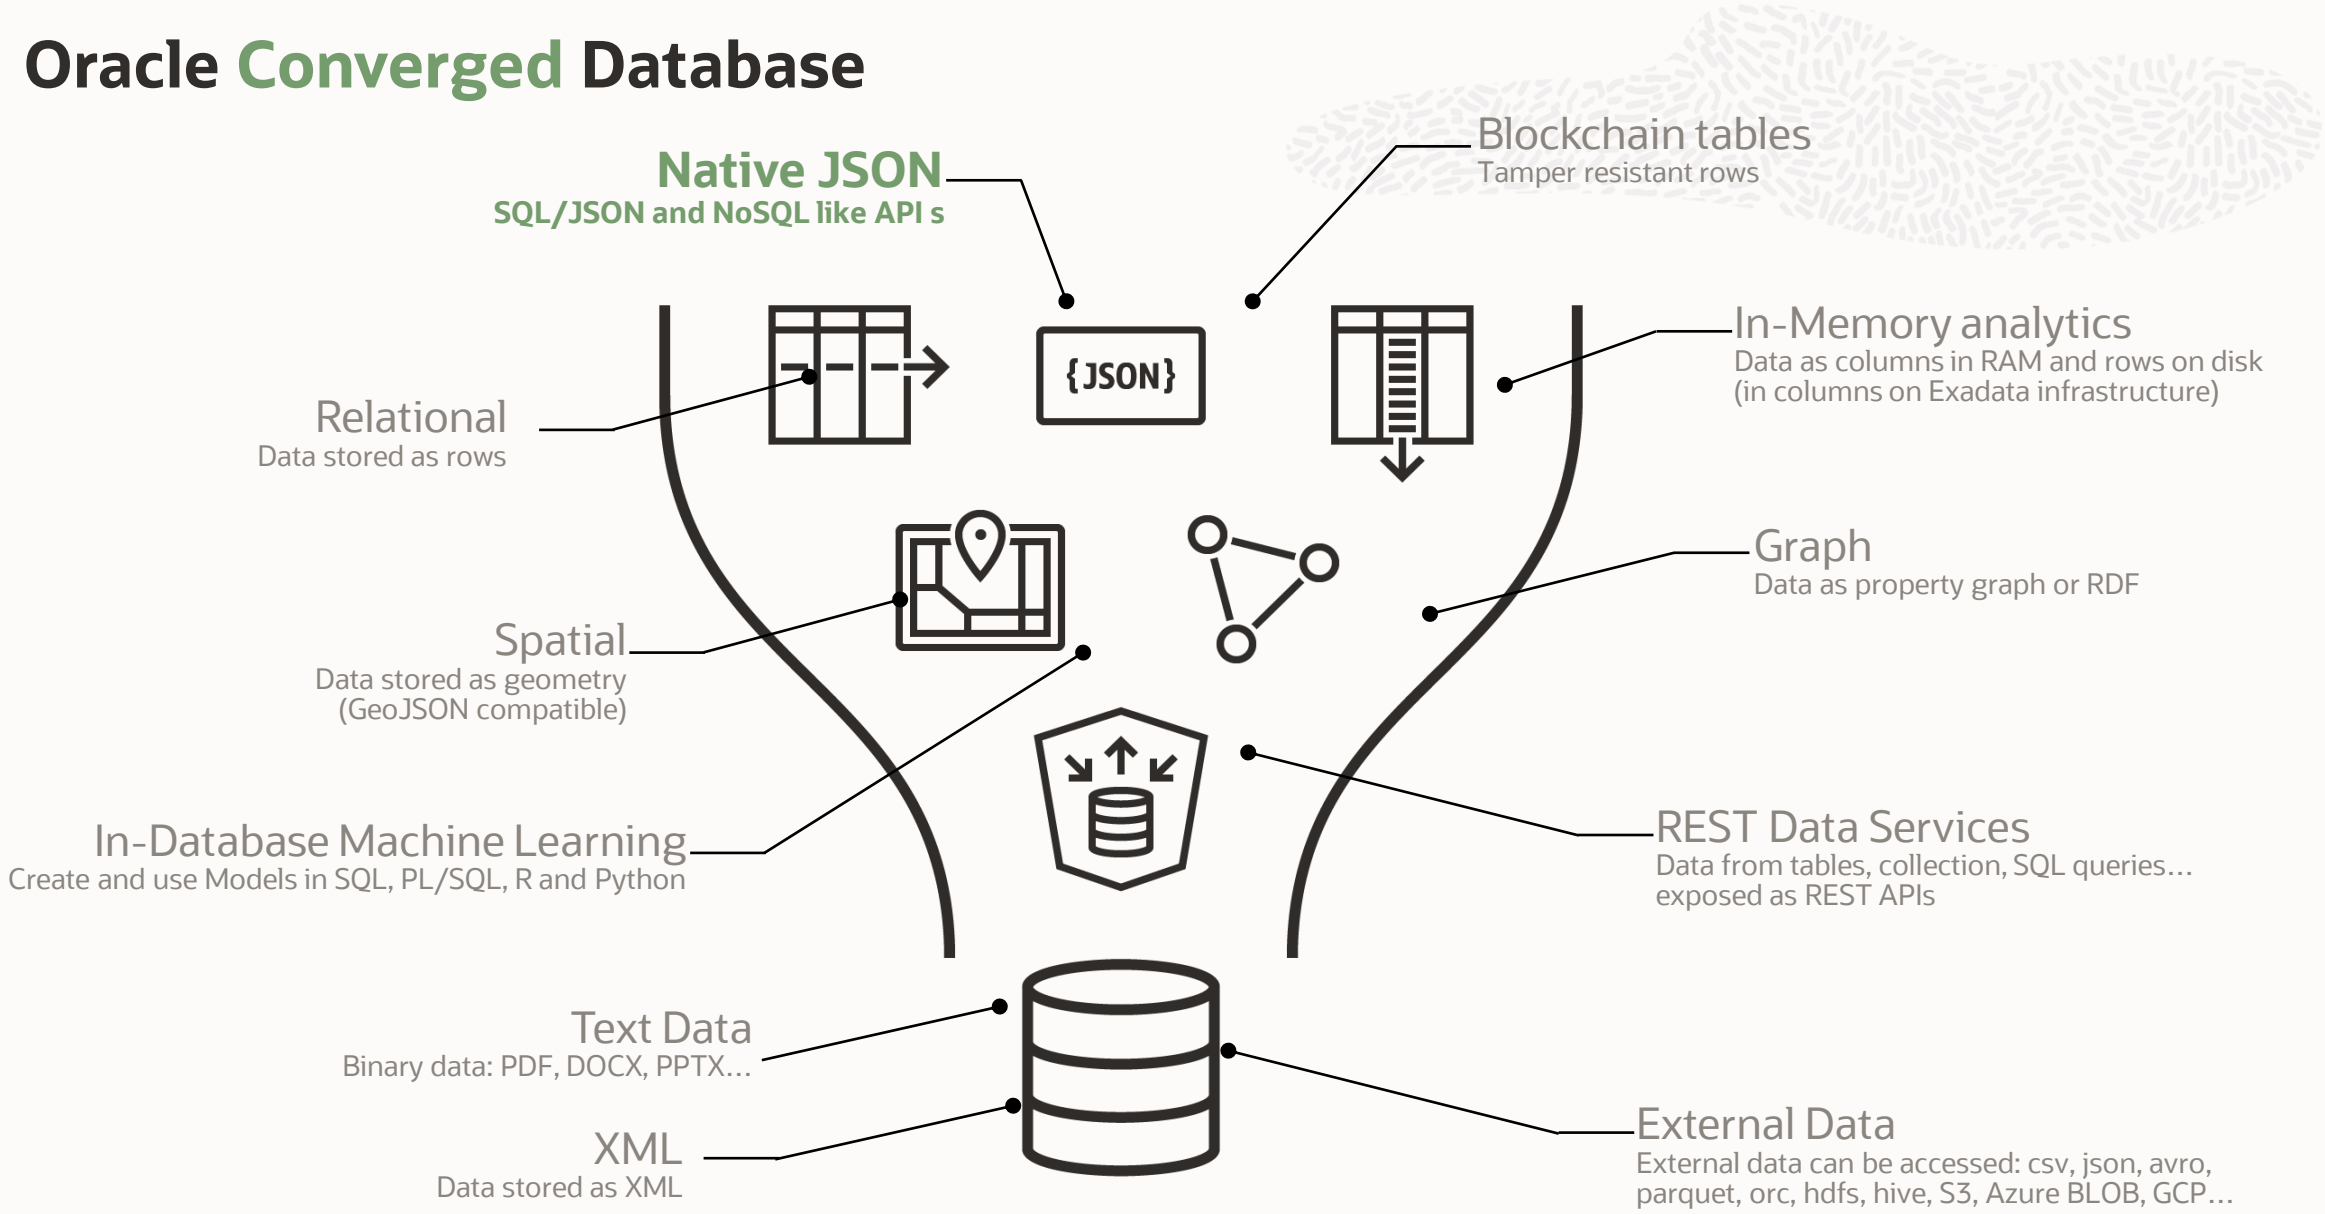

Oracle Autonomous JSON Database

A managed cloud service for JSON storage

- Cloud service for **JSON-centric** development
- Flexible and fast at scale
- **Native** JSON storage
- Simple **document APIs**
  - Language drivers, command-line, and REST
  - SQL not required

Oracle Autonomous JSON Database

More than a simple document store

- ✓ Autonomous
- ✓ Full SQL support
- ✓ ACID transactions
- ✓ Advanced security
- ✓ APEX low-code development
- ✓ One-click instant expansion to ATP

Java APIs Comparison SODA versus MongoDB

MongoDB 4.1.X API

```
MongoClient mongoClient = MongoClients.create(connString);
MongoDatabase database = mongoClient.getDatabase("finance");
MongoCollection<Document> coll = database.getCollection("pos"
);

Document po = Document.parse(json);
coll.insertOne(po);

Bson filter = eq("Requestor", "Alexis Bull");
MongoCursor<Document> cursor = coll.find(filter).cursor();
Document doc = cursor.next();

coll.createIndex(Indexes.ascending("Requestor"));
```

SODA for Java 1.1.X API

```
OracleClient sodaClient = new OracleRDBMSClient();
OracleDatabase database = sodaClient.getDatabase(jdbcConn);
OracleCollection coll = database.openCollection("pos");

OracleDocument po = database.createDocumentFromString(json);
coll.insert(po);

OracleCursor cursor = coll.find().filter("""
        {"Requestor": "Alexis Bull"}""").getCursor();
OracleDocument doc = cursor.next();

coll.admin().createIndex(database.createDocumentFromString(""
"
          {"name" : "REQUESTOR_IDX",
           "fields" : [{"path" : "Requestor"}]}"""));
```

The left-hand part of the x-axis shows the YCSB results for a small 4 million document database. The right-hand part of the x-axis shows the results for a much larger database with 81 million documents. Within each part are three benchmarks with different IO configurations (95% find/5% write, 50% find/50% write, 5% find/95% write).

The Oracle Autonomous JSON Database is shown in red, MongoDB is in green, and Amazon DocumentDB is in blue.

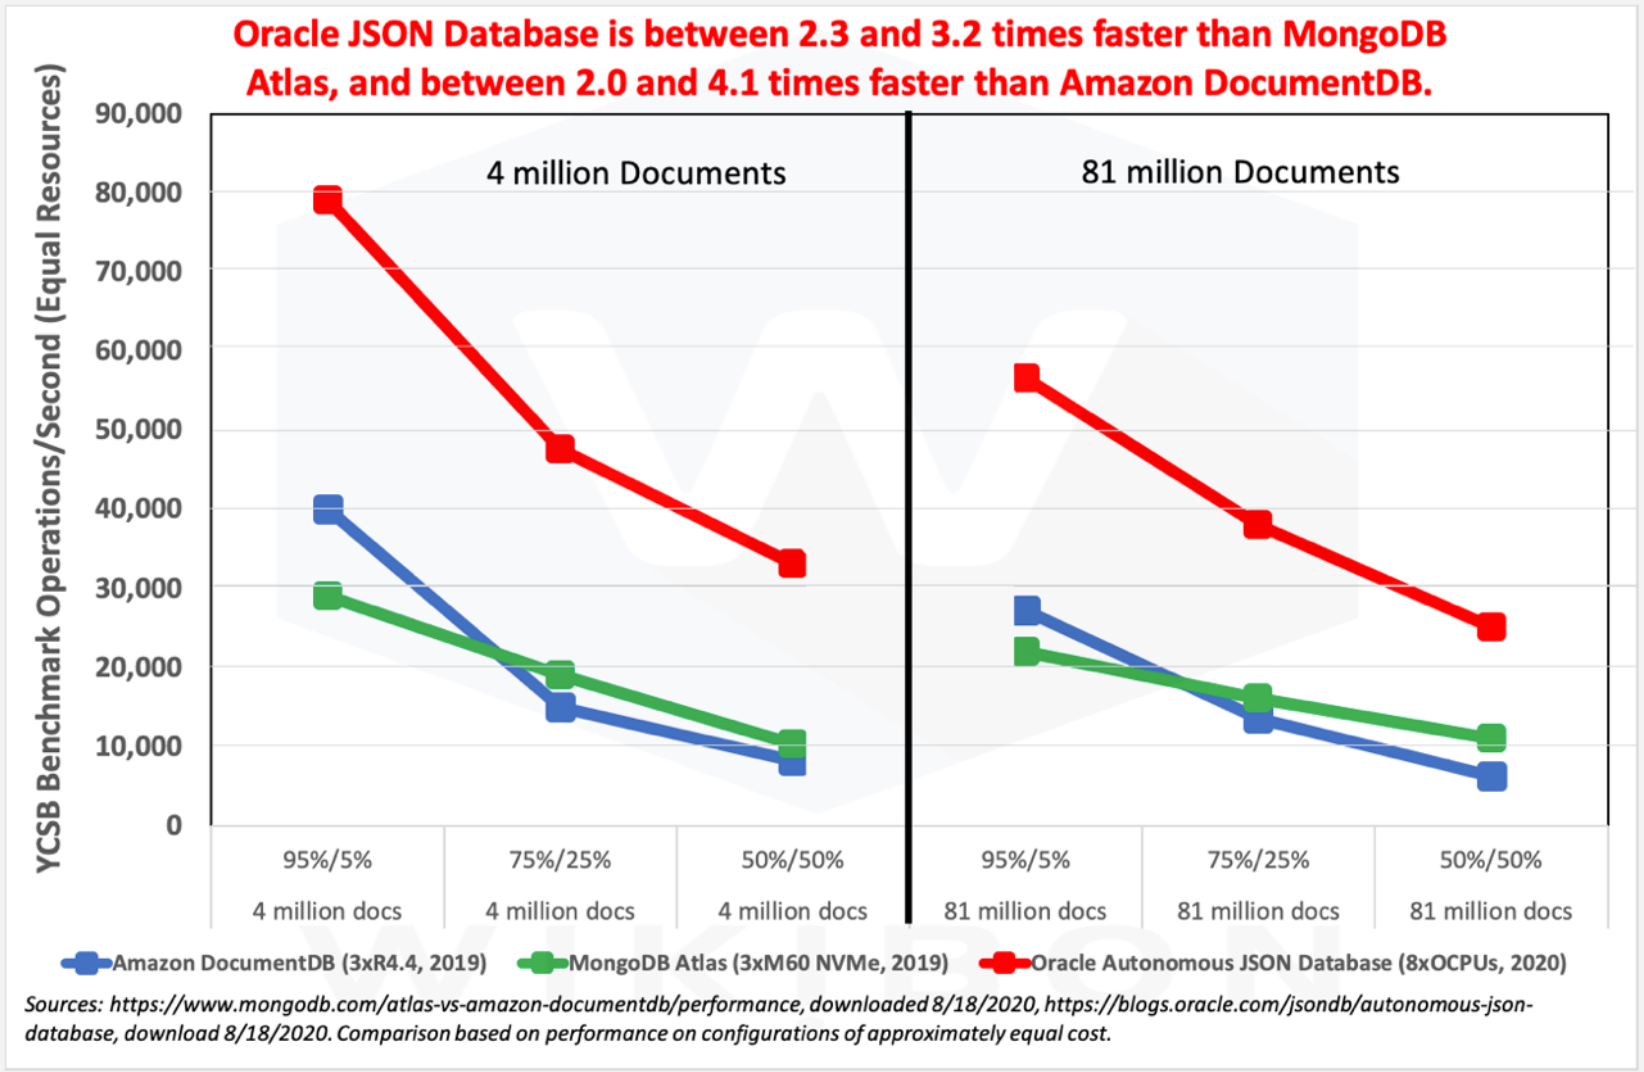

Figure 1: Document Database Performance Comparisons between Amazon DocumentDB, MongoDB Atlas, and Oracle Autonomous JSON Database.

Source: YCSB Benchmark between MongoDB Atlas & Amazon DocumentDB, & YCSB Benchmark between MongoDB Atlas and Oracle Autonomous JSON Database. Comparison based on Operations/second on configurations of approximately equal cloud hardware and costs.

Database API for MongoDB API

Demo

https://blogs.oracle.com/database/post/mongodb-api

Oracle Database API for MongoDB

Direct connections from MongoDB drivers to Oracle Database

- Developers keep their skills and continue to use MongoDB's tools, drivers, etc
- Easy migrations of MongoDB workloads to Oracle
  - Often minimal code changes required (exception aggregation pipeline)
- Enable SQL and Oracle ecosystem for JSON collections
  - More and faster analytical capabilities, machine learning
  - Query JSON alongside other data models: relational, XML, spatial, etc. using SQL
  - Expose relational data, reports, query results as MongoDB collections

Now available in Autonomous Database (shared only)

- Autonomous JSON Database, ATP, ADW

Future support is investigated for ADB- dedicated, DBCS, ExaCC, On-premises

Autonomous JSON Database (AJD) versus Autonomous Transaction Processing (ATP)

- All features of AJD also available in ATP
- AJD: 20GB limit on non-JSON data
- AJD costs approx. 25% of ATP
- One-click upgrade from AJD to ATP
- *Always Free* versions available for both

Autonomous „Killing“ Feature

ORDS – Auto Rest

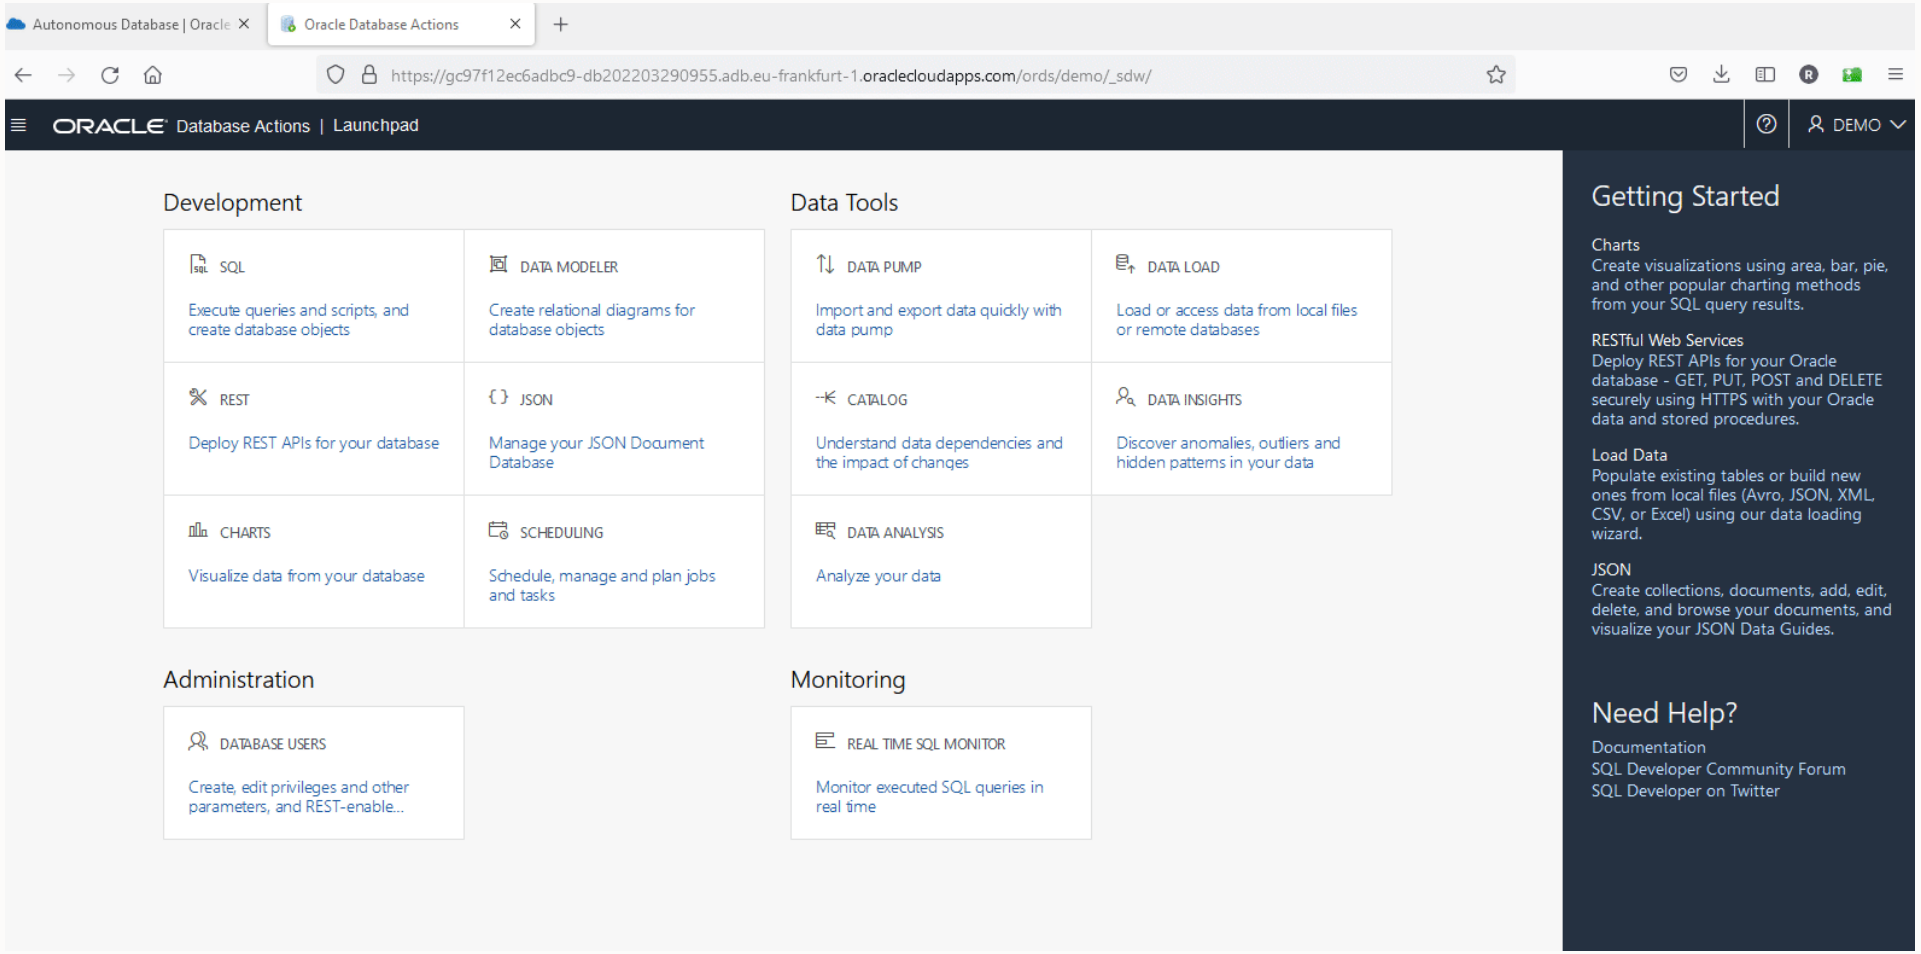

Try Now for Free

https://cloud.oracle.com/try-autonomous-database

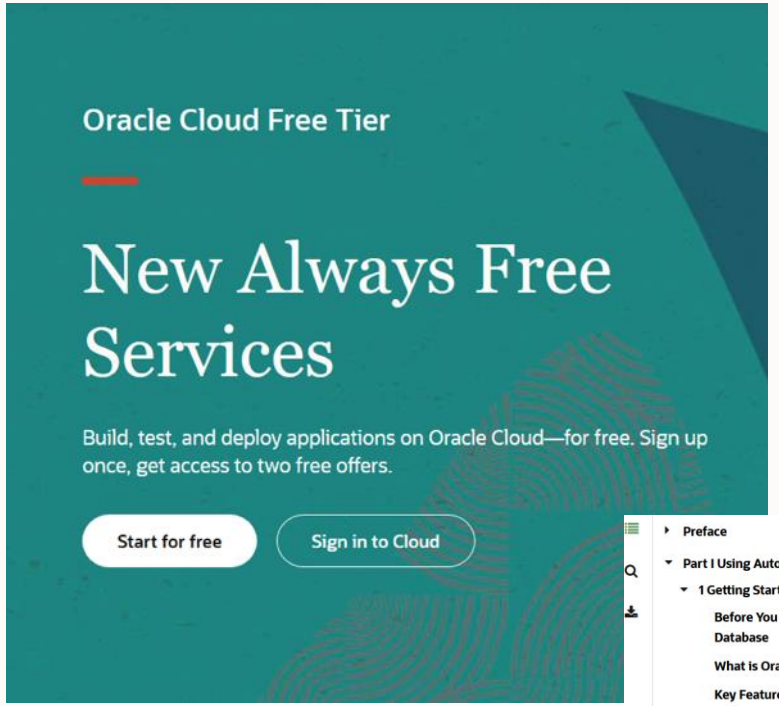

Resources:

- Reference Architectures
- Hands-on-Labs
- Demos
- Cloud Navigator
- Documentation all-in-one
- 875 pages (and growing)

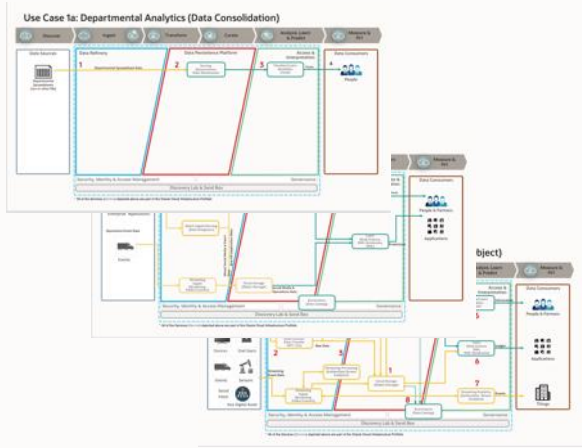

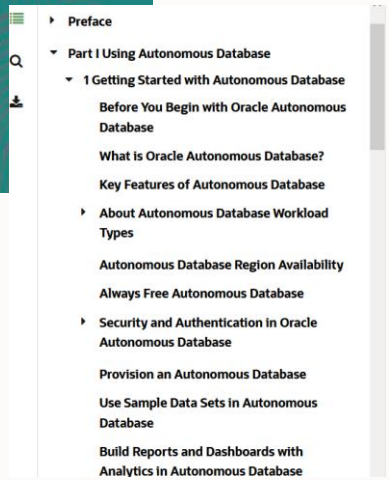

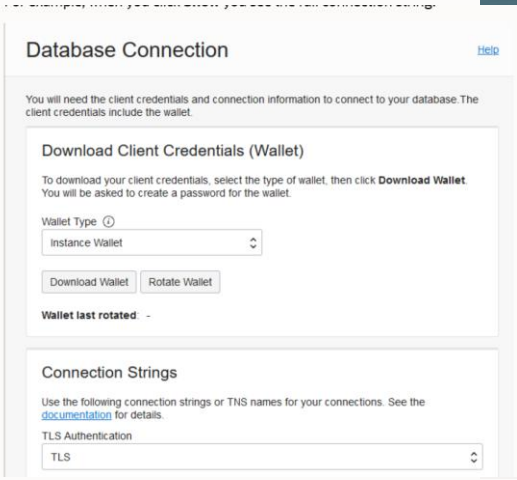

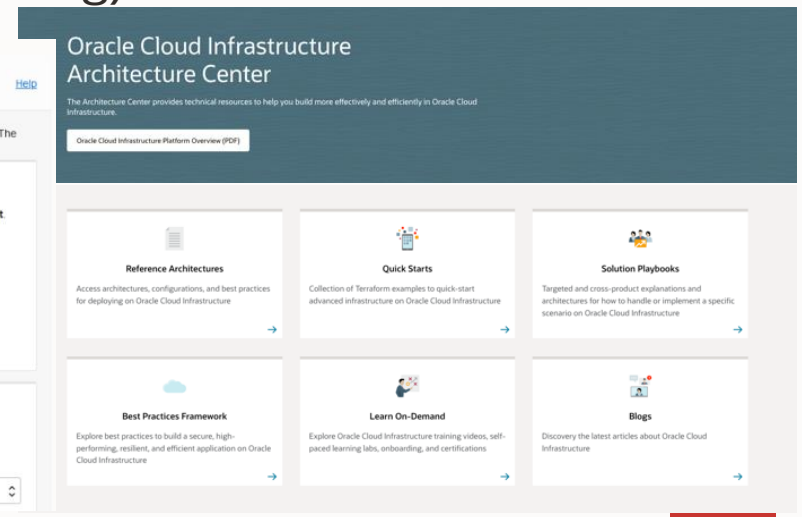

LiveLabs – Deep Dive With Autonomous Database

https://bit.ly/get-started-with-adb

Workshop Outline

- Provision ADB
- Load Data
- Query and Visualize Data
- Wallets
- Manage and Monitor
- Scale
- Machine Learning Notebooks
- Build a Machine Learning Algorithm


ORACLE