

INTELLIGENCE • INSIGHT • INTEGRATION

CONNECTED FUTURE

02 APRIL '26 - VESTAR, ANTWERP





Tripwire Solutions

Oracle AI data platform and the new Oracle Autonomous AI Lakehouse



Tripwire Solutions Services

PRINCIPLES

- Oracle
- Google
- Microsoft Azure
- Databricks
- Data Vault 2.1
- Modern Data Platform
- Reference Architecture

SOLUTIONS

- Data- & Solution Architecture
- Data Platform Automation, DWA, ELT, Zero-ETL
 - Migrate to Data Vault
 - Batch to Real-Time
 - Performance Tuning
- Data Modelling
 - Data Vault 2.1
 - Data Platform Data Modelling
- AI, Data Science & Advanced Analytics
 - Predictive Analysis
 - Machine Learning
 - Deep Learning
 - NLP
- Reporting & Dashboarding
- Data Discovery & Data Visualisation
- Data Governance
- Data Lakehouse
- Cloud and On-Premises

SERVICES

- Assessments
- Data Platform Build & Setup
- Data Platform Managed Operations
- Team Enablement
- Expert Consulting on Demand
- Total Outsourcing





AI Database



AI App Dev



AI Lakehouse



AI Data Platform

Data Platform Emergence for Analytics

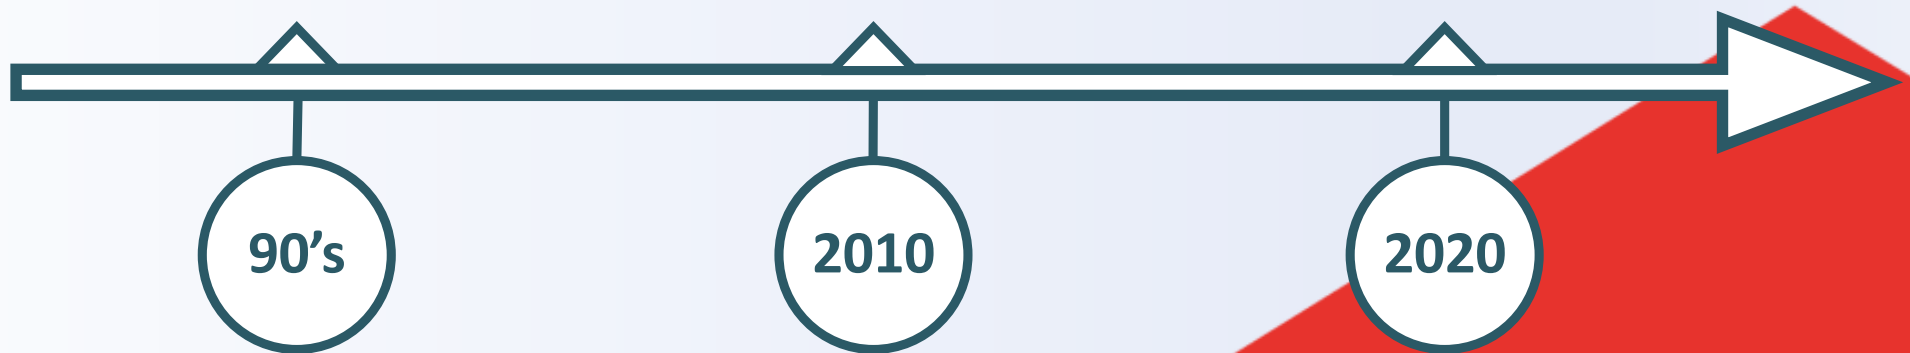
Data
Warehouse



Data Lake



Data
Lakehouse



Data Lakehouse



Data
Warehouse

+

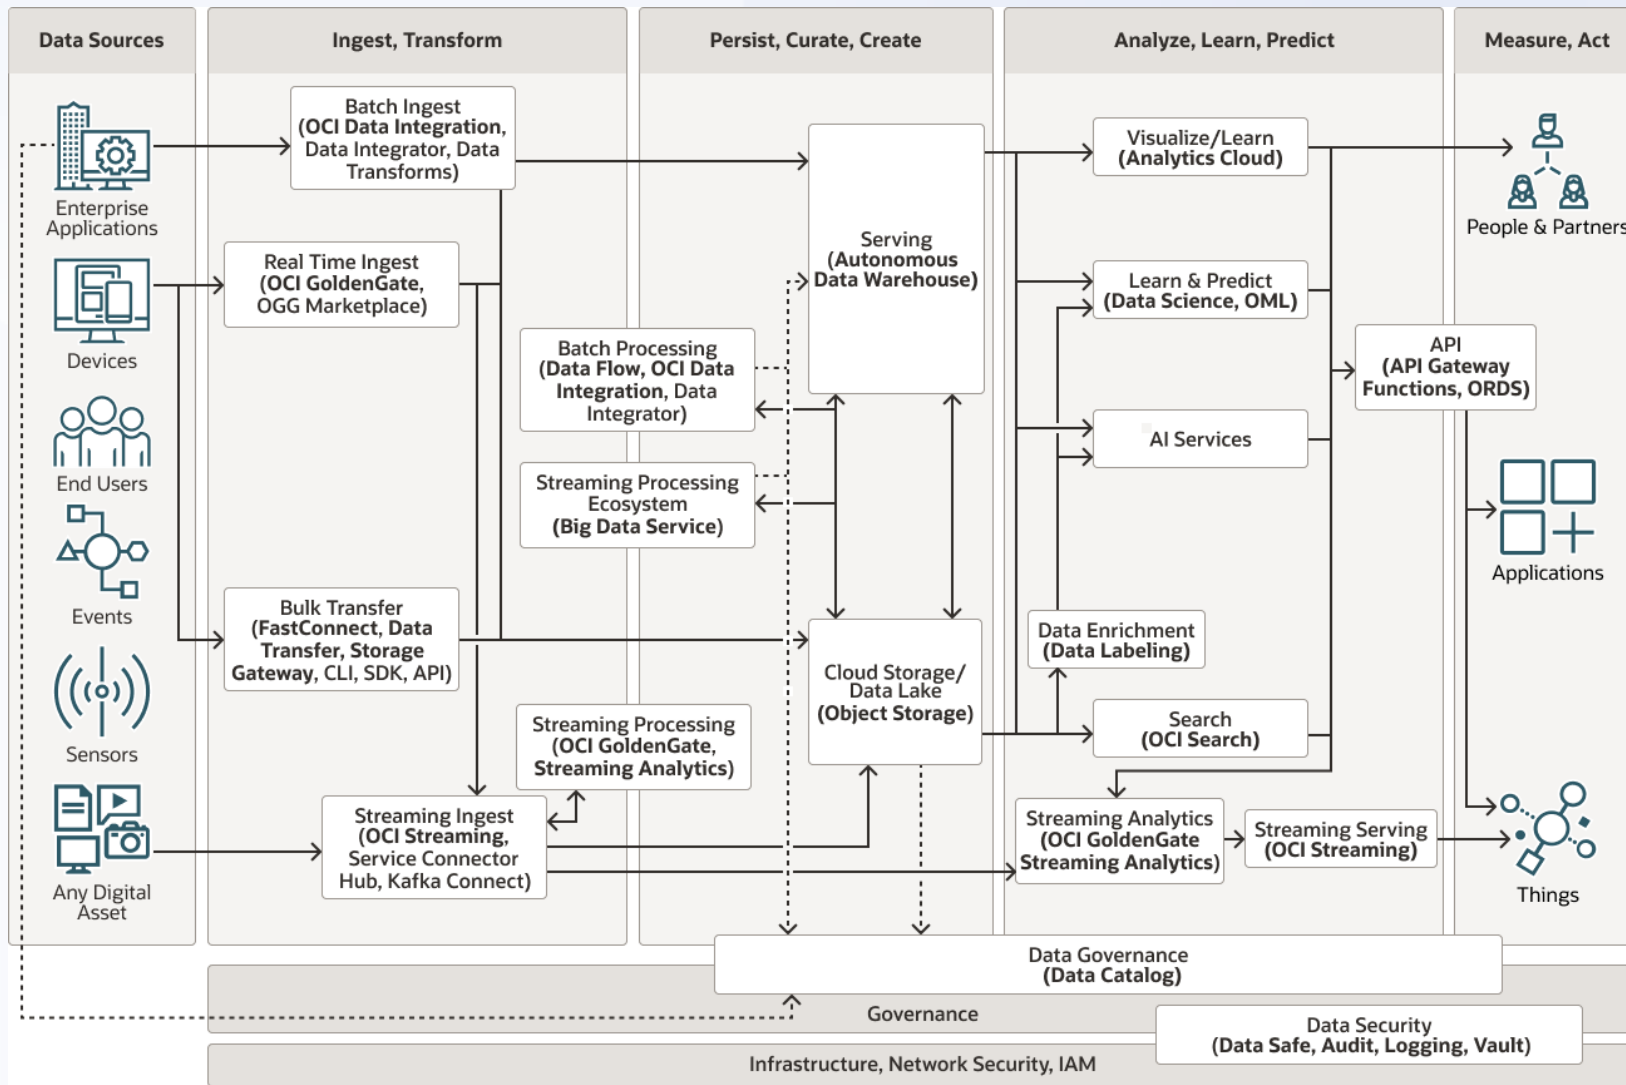


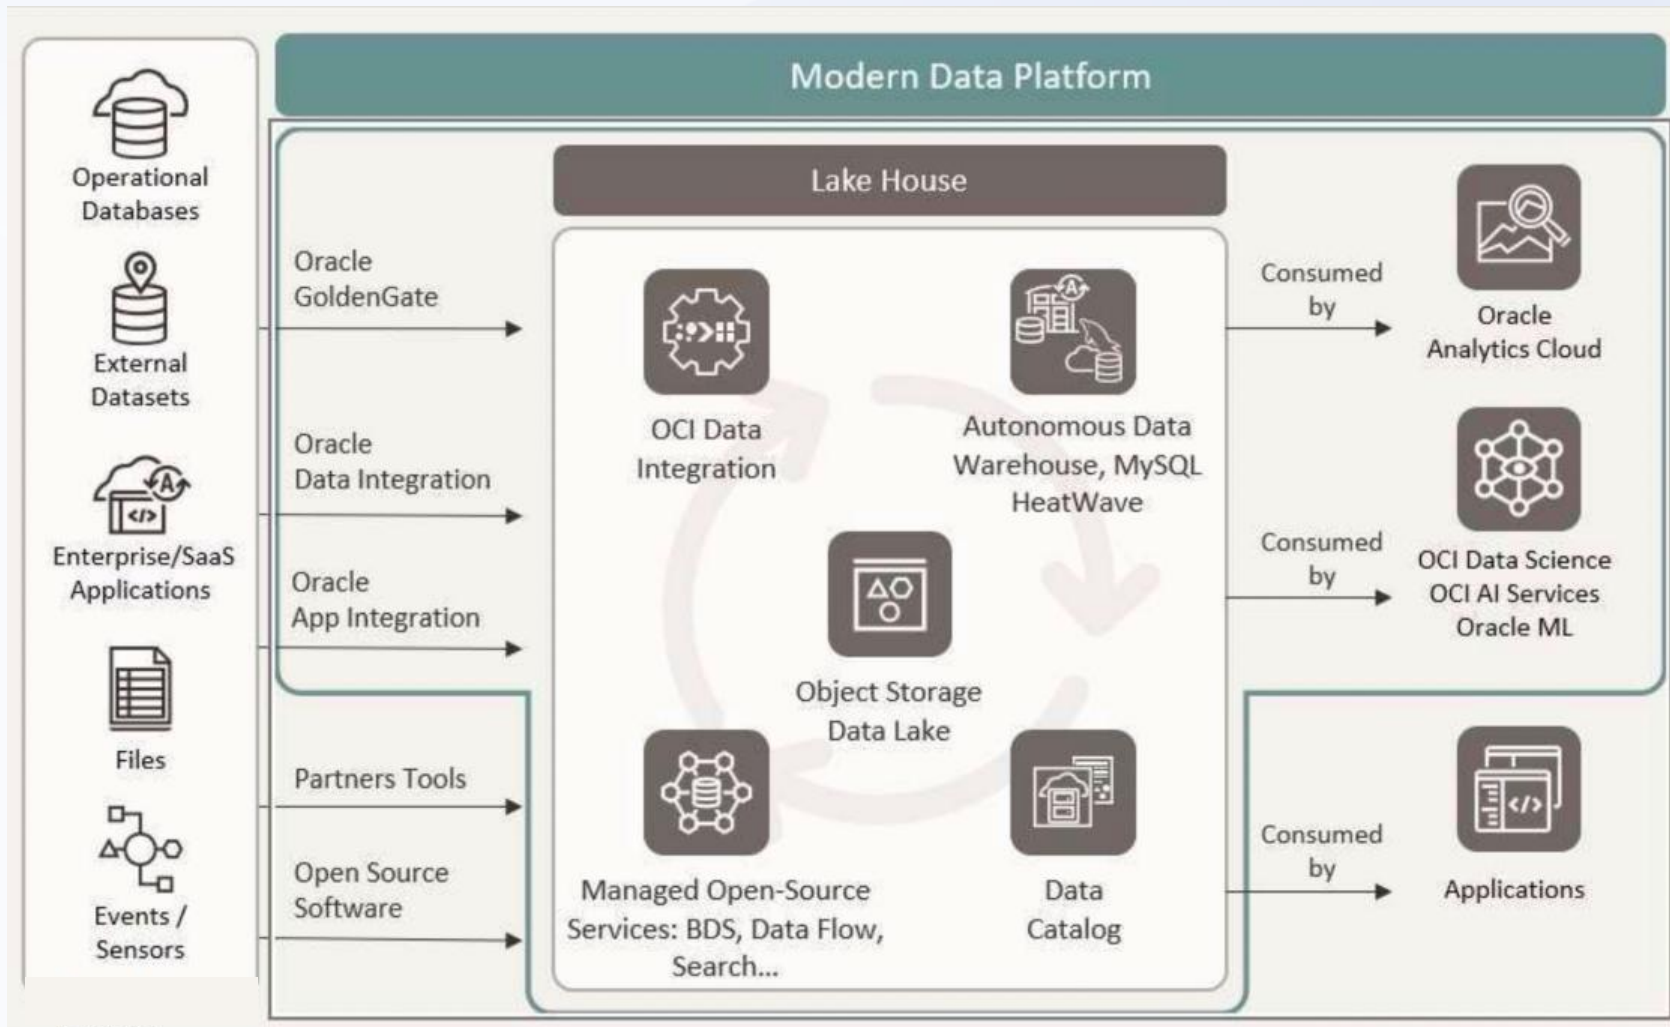
Data Lake

=



Data
Lakehouse



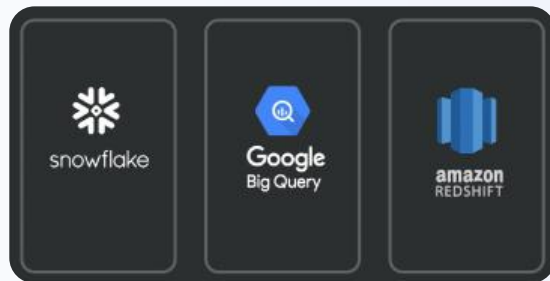


Give your AI & Analytics the enterprise data they deserve

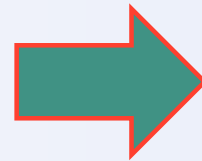
- AI & Analytic initiatives need unified view of data across the enterprise.
- Fragmented data silos hinder AI and analytics initiatives.
- Enterprises Data spread across multiple systems creates:
 - Data drift & Inconsistent records
 - Complex ETL integration pipelines
 - Separate security models
- Lack of unified access across data formats and data location
 - Slows down AI application development
 - Slows time to analytical insights.
- Data access & performance challenges
 - Object Storage performance when accessing large data sets
 - Limited support for consistent open table formats

The Data Platform world is changing...

From analytic data silos...



To open and interoperable systems



Oracle is embracing **Open Standards** and **Open Platforms**

- Now Oracle is embracing **Apache Iceberg** to provide an **open AI Lakehouse**
- **Iceberg** standard defines a common **table format** and **open catalog** for tables in object store



Why Iceberg Matters: Open Table Format = Freedom + Flexibility

- What is Iceberg?
 - It is an **open-source table format** designed for cloud data lakes
 - Includes data files presented as a "unified" table
 - Includes Metadata stored in manifests files along with data
- Why does Iceberg matter?
 - Your data is not locked into any vendor
 - Designed for TB to PB scale across object storage
 - Built for multi-engine access
 - Maintains ACID Compliance with multiple engines making updates
 - Enables schema evolution without pain
 - Add Columns / Drop Columns
 - Simplifies building data pipelines to give your AI the data it needs

- Apache Iceberg has many benefits
 - Data versioning
 - Partitioning
 - Basic transactions
 - Basic schema evolution
- And some tradeoffs
 - No indexes
 - Batch-oriented updates
 - Table-level transactions
 - Rudimentary security
 - Slow storage



Iceberg tables are **concurrently readable and writeable** by many data engines



Unified AI & Analytics with Oracle Autonomous AI Lakehouse

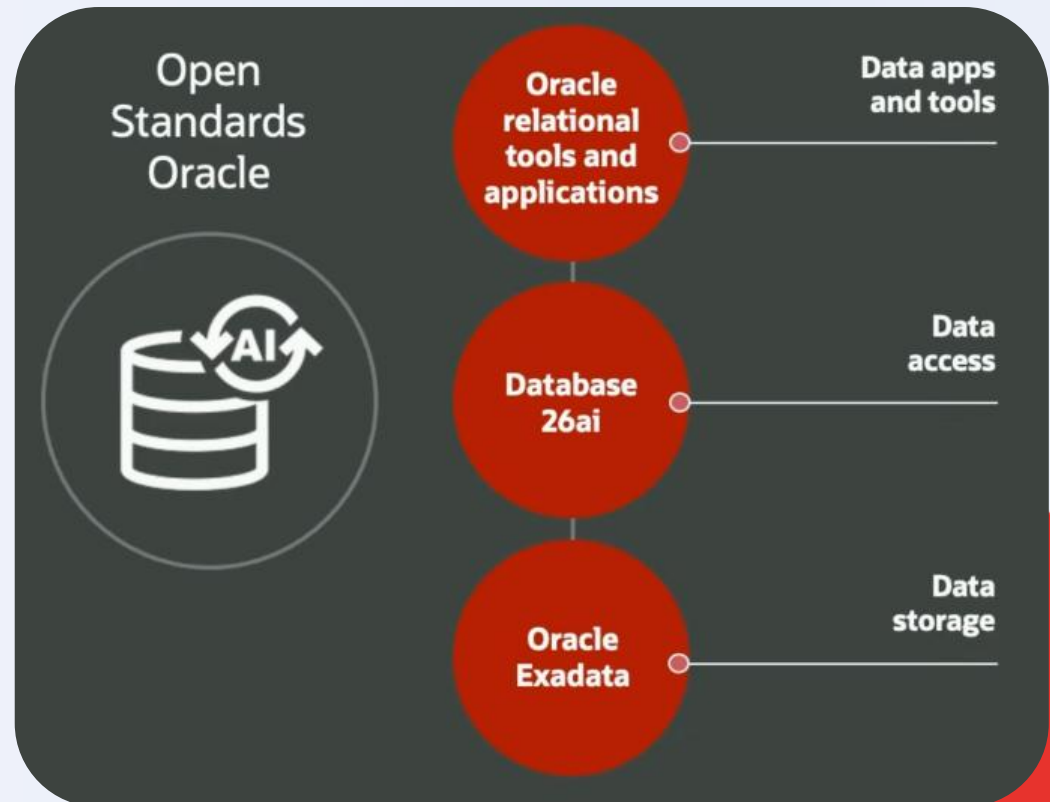
Enterprise-wide AI & analytics with the best of Open Table Iceberg Data Lake
PLUS the best of Oracle



Converged platform that combines the optimized performance and efficiency of the Oracle AI Database with the openness of Apache Iceberg to deliver a unified Lakehouse environment for all of your data, all of your workloads, in all clouds.

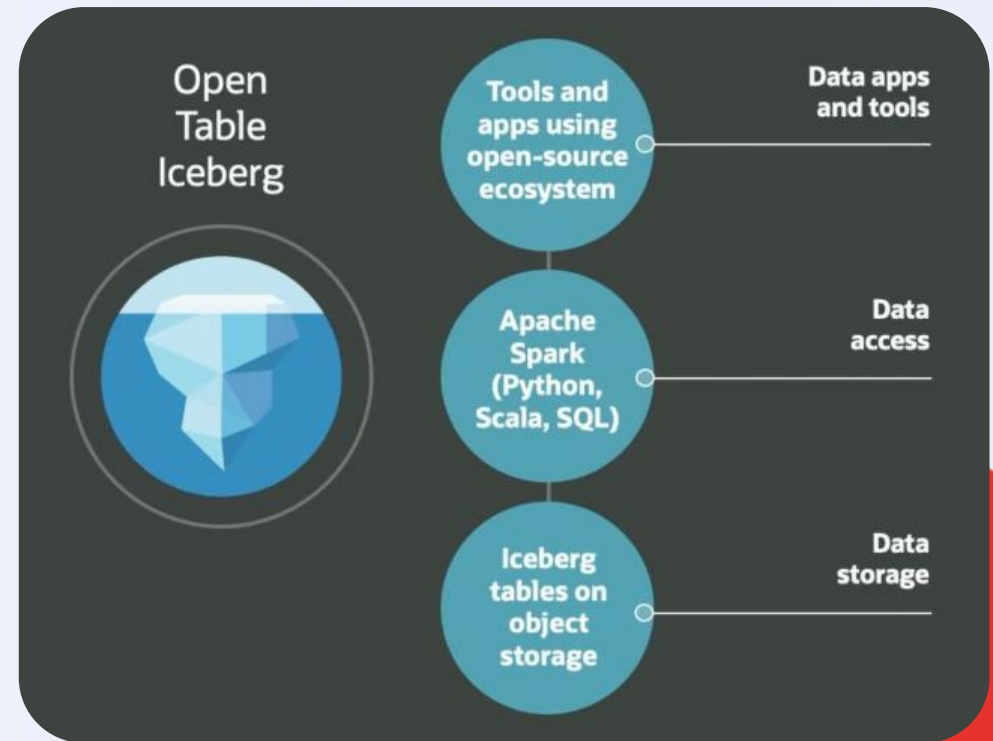
The Oracle data platform is arguably the most capable

- Converged Data Architecture supports all leading data types
- Sophisticated secure analytic SQL perfected over 45 years
- Translytical real-time, enterprise workloads
- Broadest ecosystem of tools and apps
- Extreme-performance Exadata storage for same price as object store + no cost per I/O

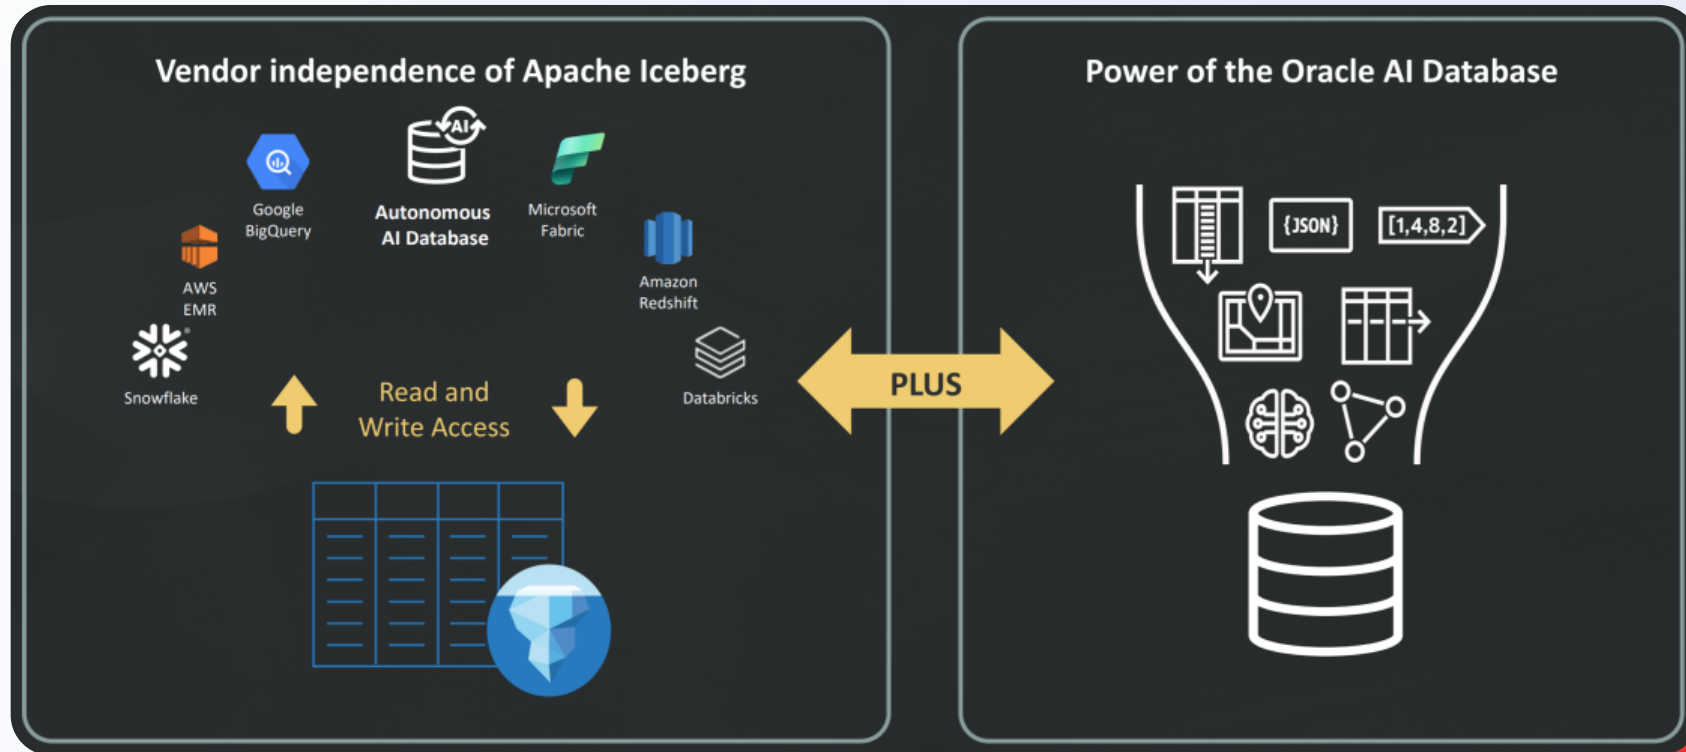


An Iceberg-based data platform is arguably the most open

- Popular with data science and AI communities
- Open-source analytics engines such as Spark and Flink, to process data
- Open data formats: Parquet, CSV, JSON, Avro, Delta UniForm, Iceberg
- Highly scalable and low-cost object storage

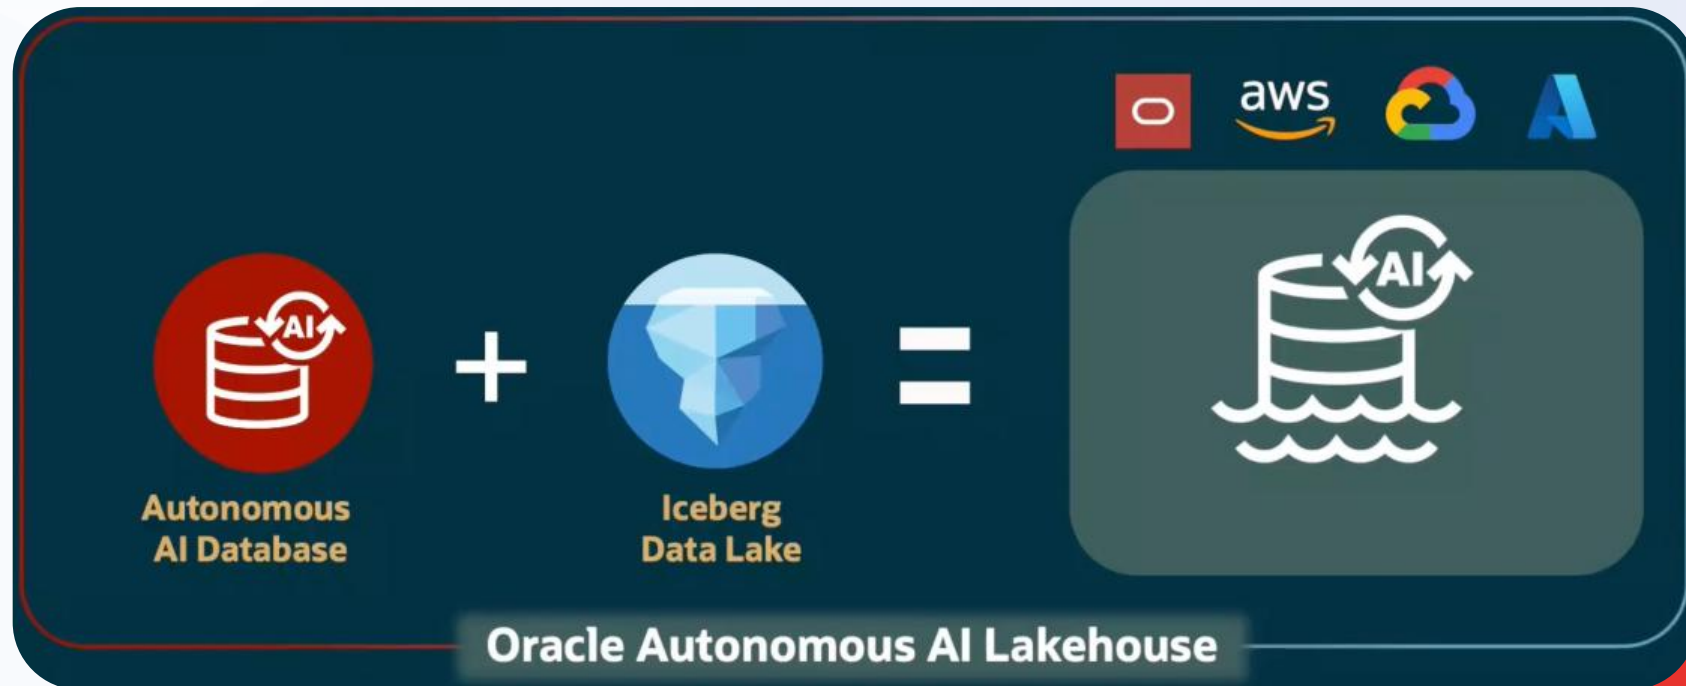


Oracle's new Autonomous AI Lakehouse: The best of both worlds

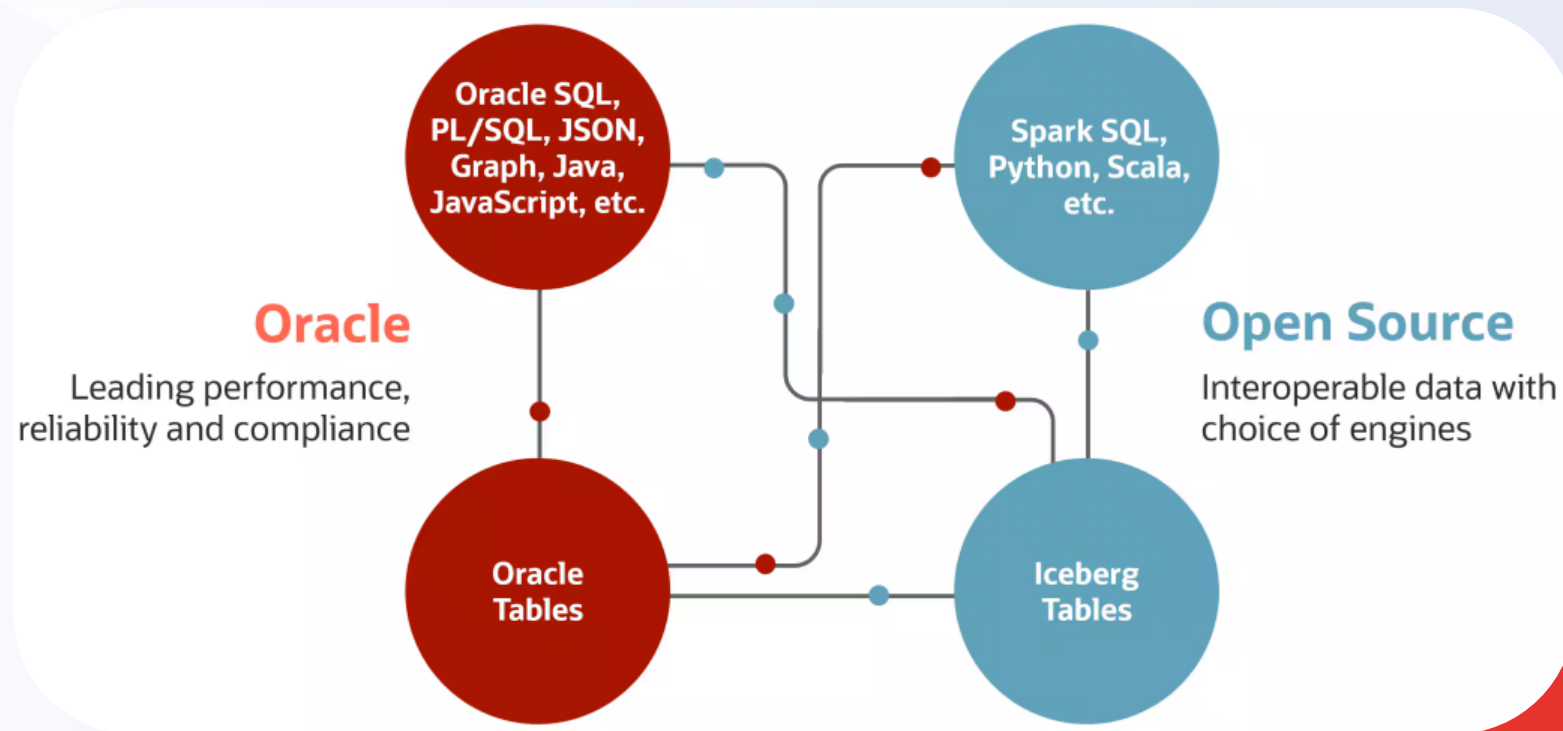


Oracle Autonomous Data Warehouses have **automatically** become **Autonomous AI Lakehouses**

A true multi-cloud Lakehouse



Instead of having to choose between most capable and most open platform, you can now get the best of both with the **Oracle Autonomous AI Lakehouse**

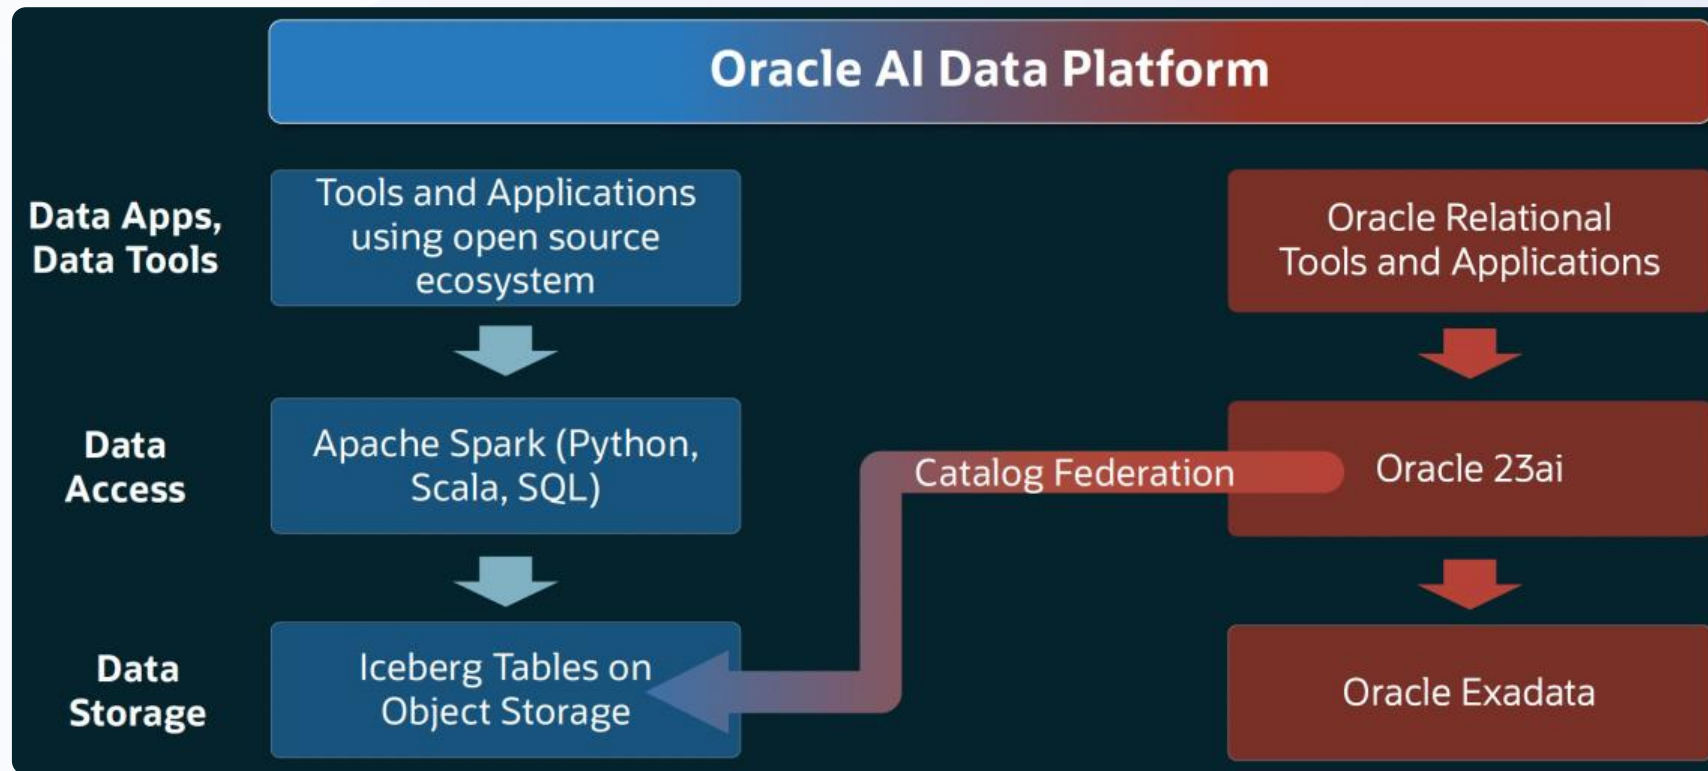


Oracle Autonomous AI Lakehouse

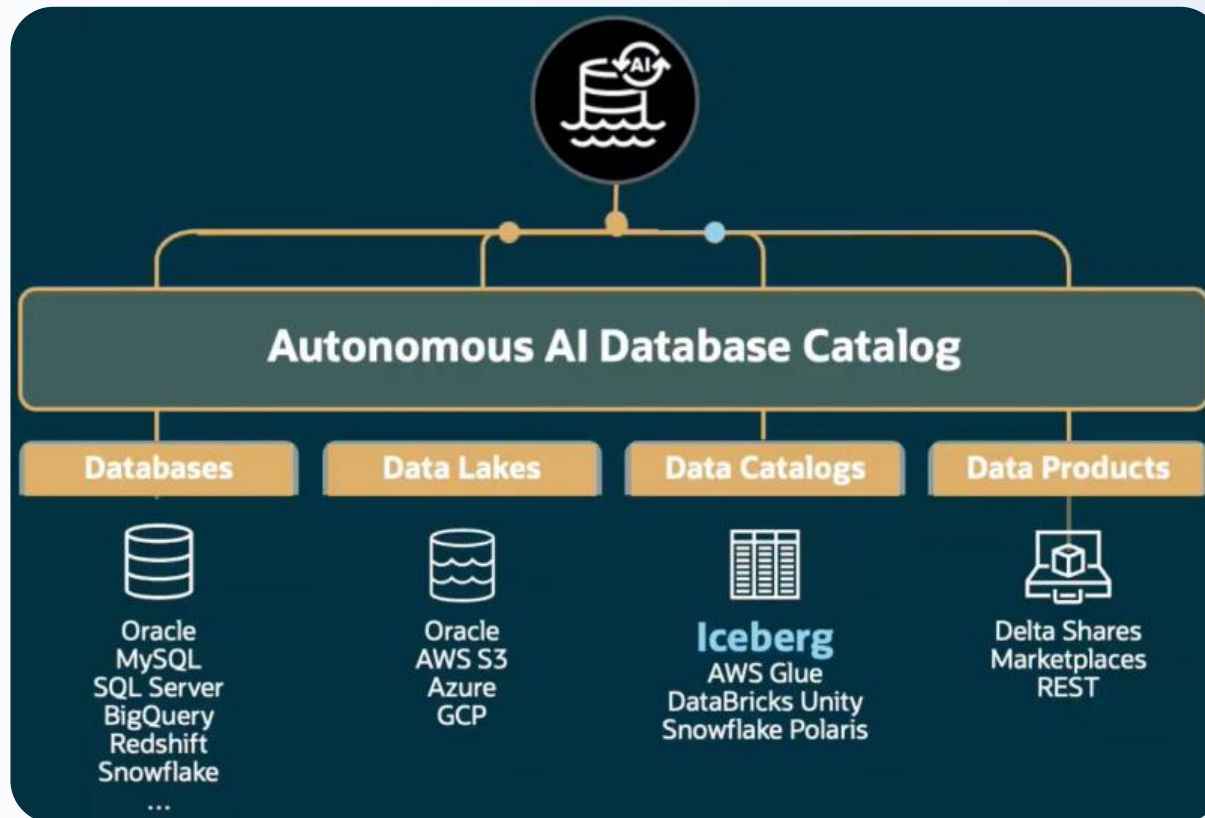
- Delivers the best of both worlds-powerful plus open
- Oracle now brings its **performance, reliability, and AI** to Iceberg with zero data movement
- A built-in data catalog provides a **unified view** of all enterprise data
- Oracle Data Lake Accelerator provides **fast and scalable queries** of Iceberg tables
- Fully managed, **autonomous operation** eliminates administrative overhead



Run any Oracle SQL workload transparently against Iceberg tables



Unified Enterprise View of Data with Autonomous AI Database Catalog



- All the capabilities of the AI Lakehouse on all your data
- "Catalog of catalogs" architecture
- Enables Oracle AI Lakehouse to discover, access, optimize, and analyze data from other catalogs and sources

Autonomous Database Catalog provides Plug-and-Play Access to Iceberg

- Simple SQL access to any Iceberg table
 - Browse from any Iceberg catalog, or one-command setup for Iceberg Tables in object storage
- Full read-only access for analytics and AI

The image displays three sequential screenshots of a SQL interface, illustrating the process of accessing Iceberg tables through an Autonomous Database Catalog.

1. List available catalogs

```
1 SELECT
2   CATALOG_NAME,
3   CATALOG_TYPE
4 FROM
5   USER_MOUNTED_CATALOGS;
```

	CATALOG_NAME	CATALOG_TYPE
1	GLUE	DATA_CATALOG
2	ADW_DEV	DB_LINK
3	DATABRICKS	SHARE

2. List the tables in a catalog

```
1 SELECT
2   OWNER,
3   TABLE_NAME
4 FROM
5   ALL_TABLES@DATABRICKS;
```

	OWNER	TABLE_NAME
1	MOVIES	CUSTOMER_EXTENSION
2	MOVIES	CUSTOMER_SEGMENT

3. Query a catalog table

```
1 SELECT
2   CUST_ID,
3   AGE,
4   CREDIT_BALANCE
5 FROM
6   MOVIES.CUSTOMER_EXTENSION@DATABRICKS
7 WHERE
8   AGE BETWEEN 19 AND 29;
```

	CUST_ID	AGE	CREDIT_BALANCE
1	1000007	23	1643
2	1000022	28	877

Plug-and-play access to Iceberg data from Oracle SQL

```
SELECT employee_name
FROM   hr.employee_event@Databricks
WHERE  period_begin = '01-July-2025';
```

- Simple syntax to use Iceberg tables
schema.table@**catalog**
- Use any Iceberg catalog
 - Databricks Unity
 - Snowflake Horizon
 - AWS Glue
 - And more

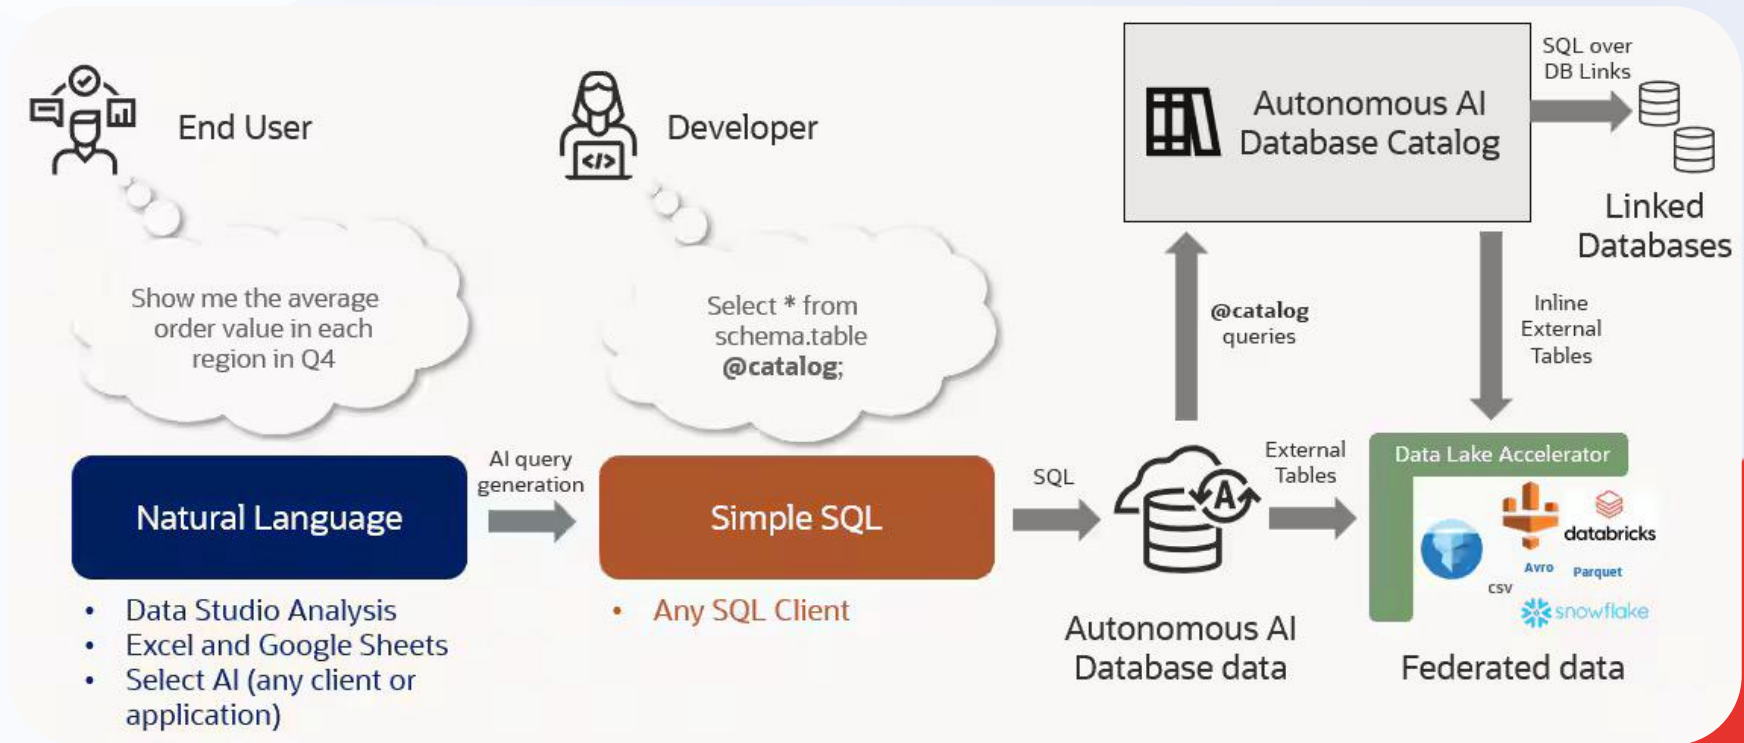
Self Service AI Capabilities: Data Science Agent

Use a pre-built AI assistant to search data catalogs, prepare and explore data, uncover key insights, and transform them into actionable insights.

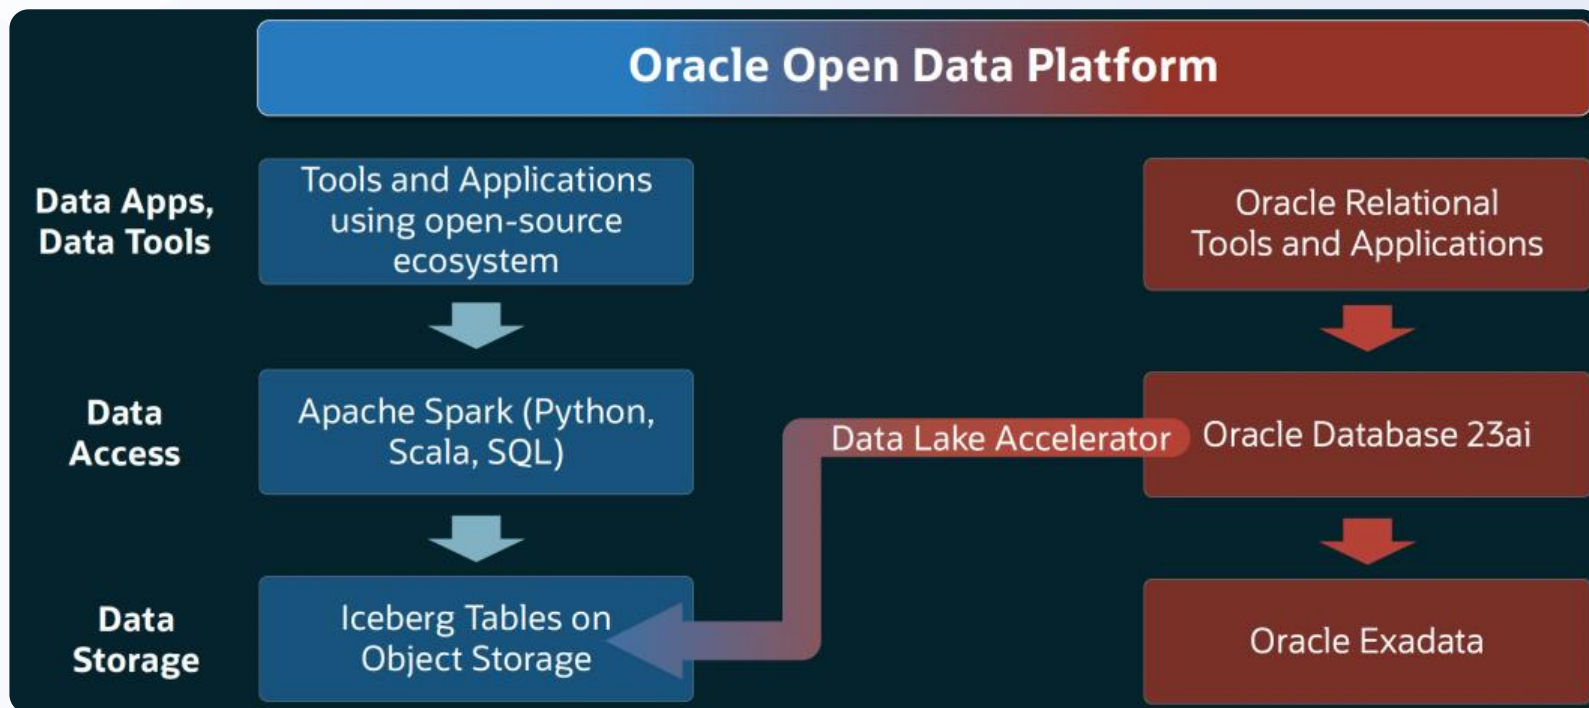


This agent is your digital teammate that speaks natural language.
You ask questions. It gets the answers.

Oracle Autonomous AI Database Catalog

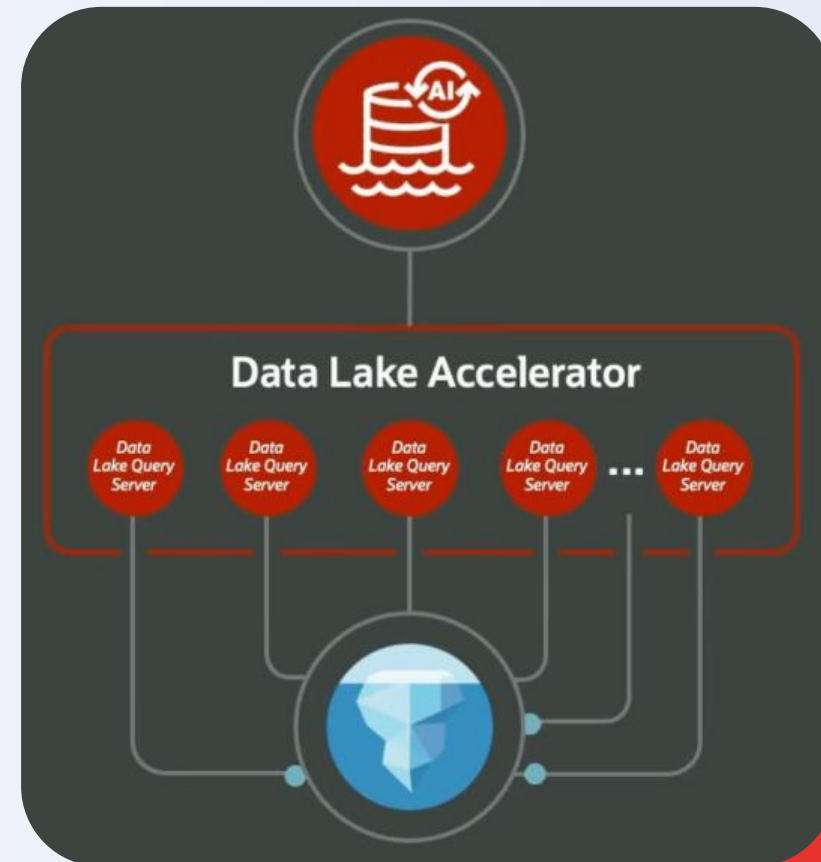


What about performance? Scale-out access using the Oracle Data Lake Accelerator



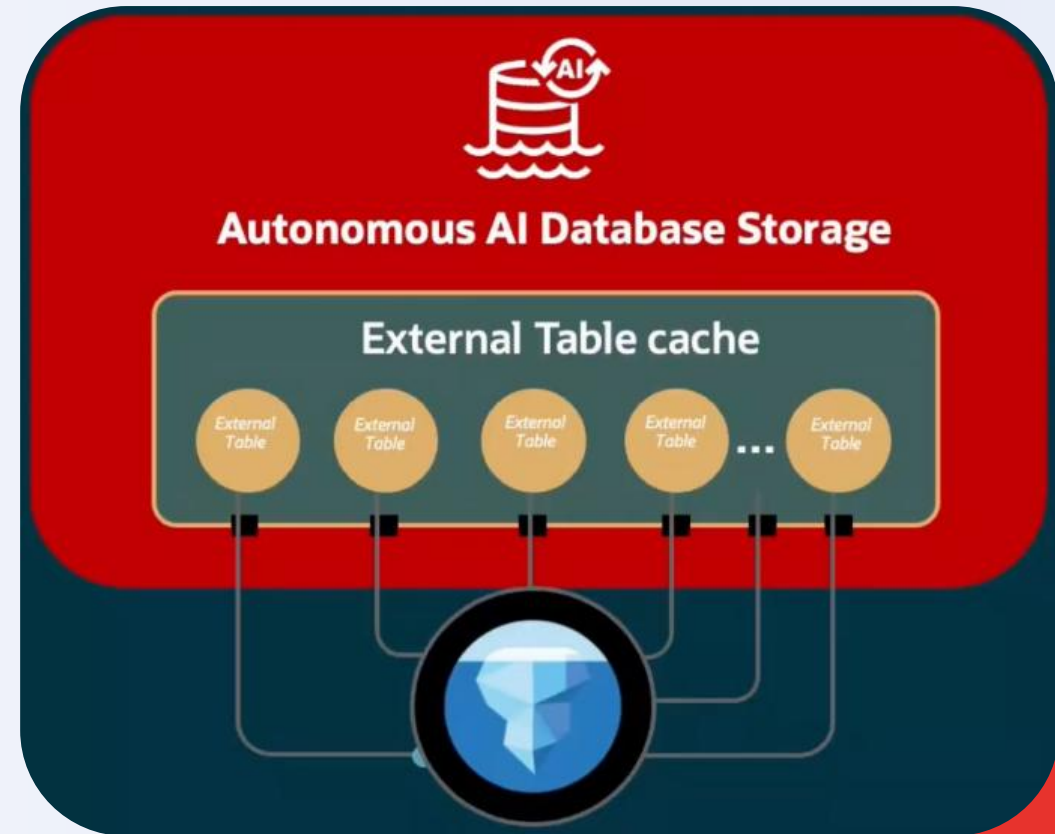
Oracle Data Lake Accelerator

- Scale out network bandwidth dynamically for scanning very large Iceberg tables
- Automatic and transparent: accelerator allocates additional network and compute resources to optimize speed
- Pay-as-you-go: activated during query execution and deactivated afterwards

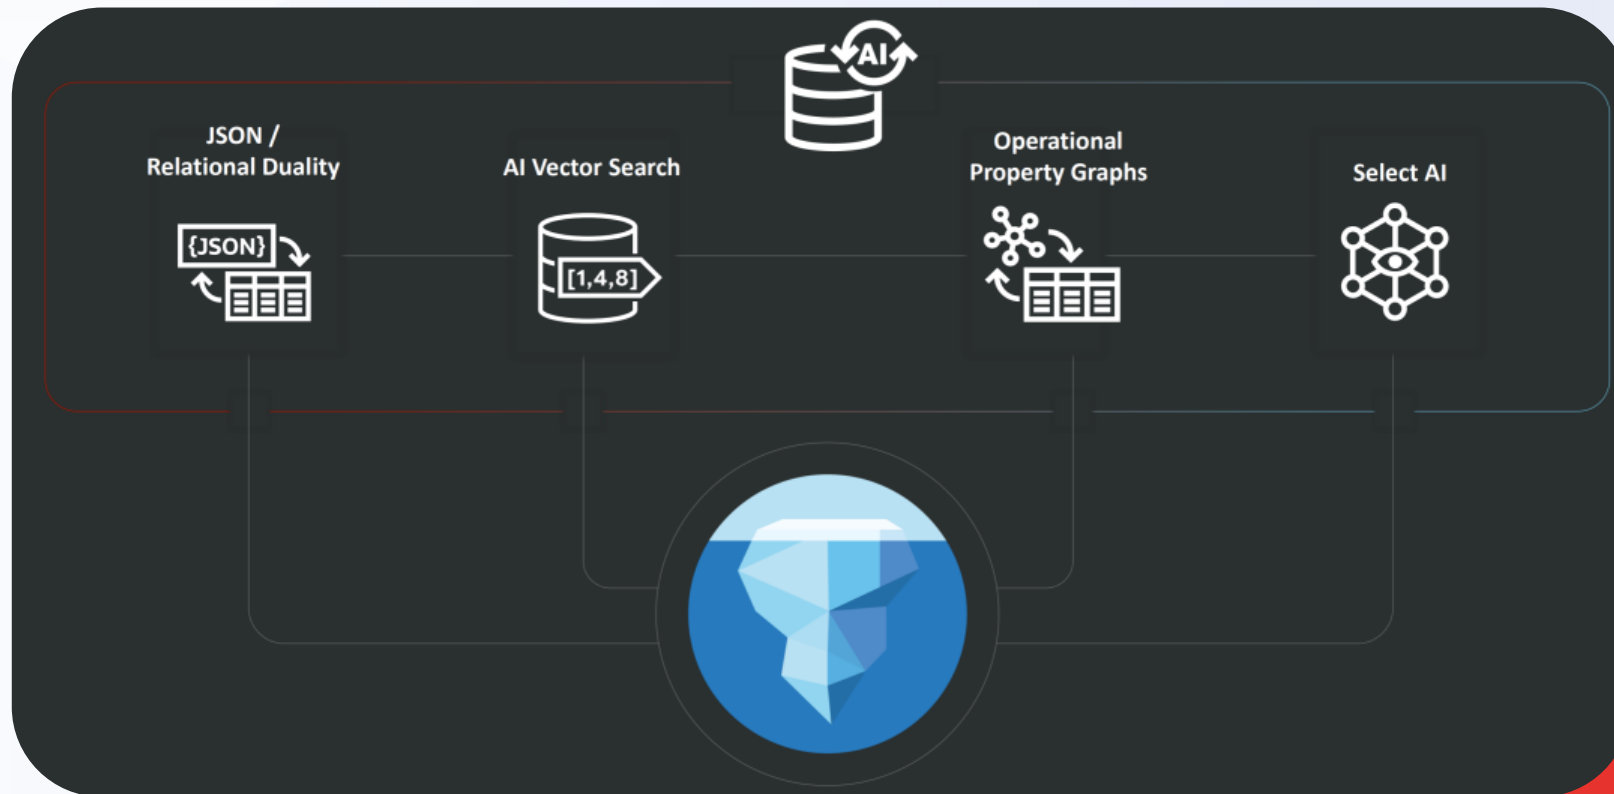


External Table Cache

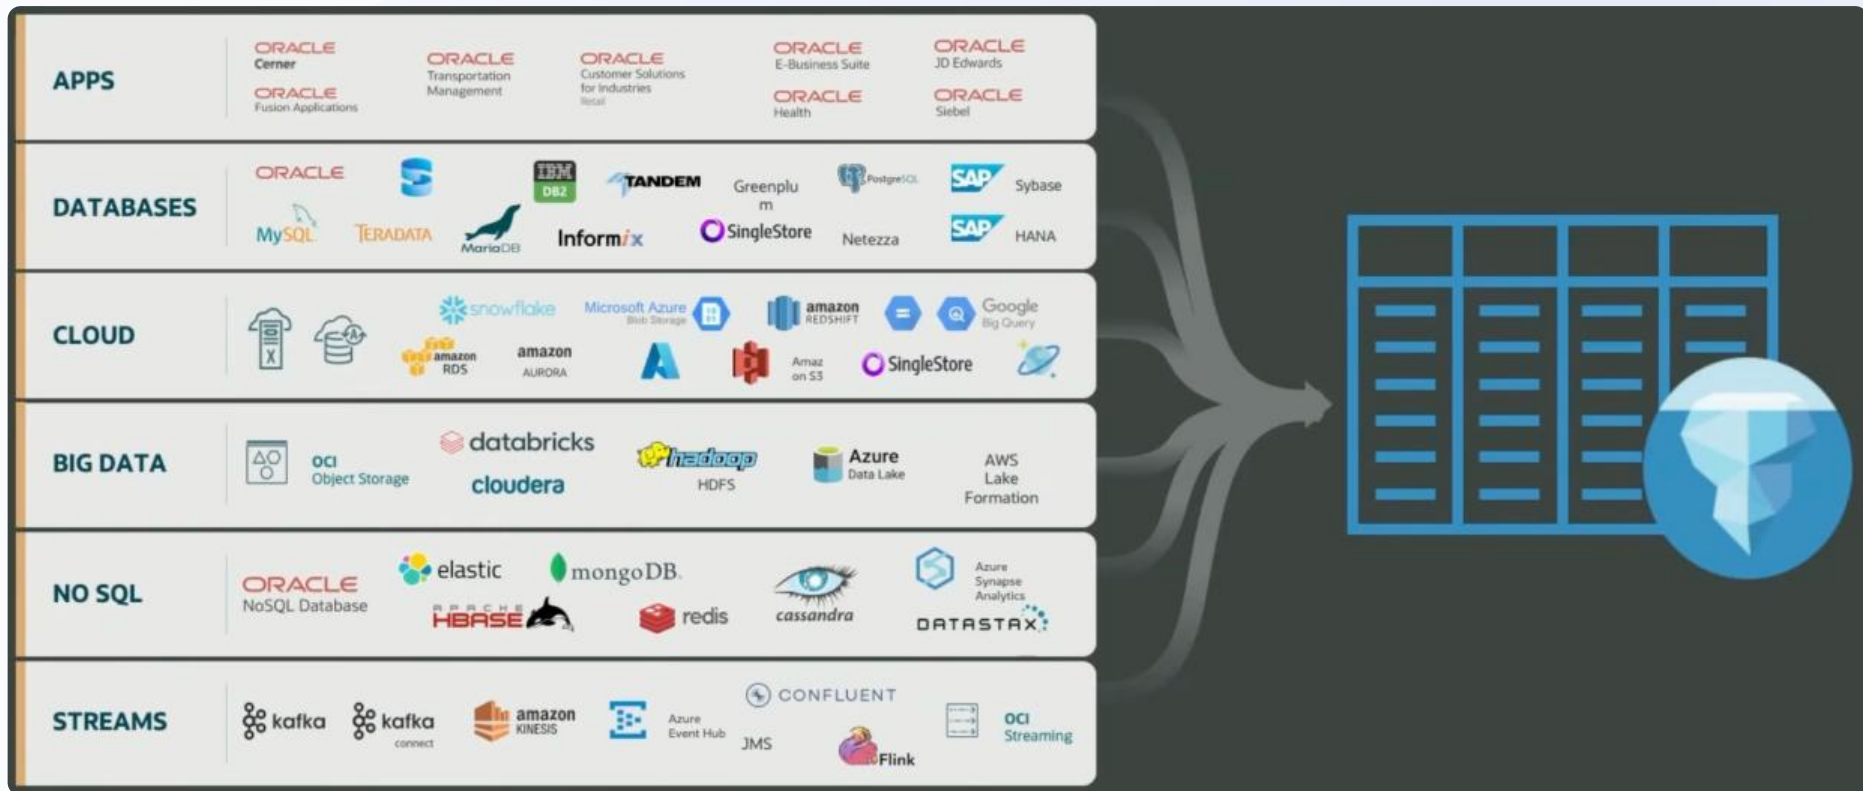
- Eliminate latency and repeated cloud fetches
- Keep your existing ETL pipelines, governance models & controls
- Use policy driven automatic caching to optimize cost & performance
- Exadata table cache:
Boost Iceberg data performance by caching frequently accessed tables in Oracle Exadata flash storage, delivering the performance of native Exadata tables.








Power of Oracle AI Database on Iceberg Tables with zero data movement



GoldenGate and Oracle ETL provide real-time data pipelines from hundreds of data sources



Improves open table Iceberg in every dimension

 <p>100% of Oracle SQL analytics</p>	 <p>100% of Oracle Workloads and Apps</p>	 <p>100% of Oracle AI capabilities</p>	 <p>100% of data in your Enterprise in one catalog</p>	 <p>100% of Clouds and On-premises</p>
powered by Exadata performance	including real- time processing	including Oracle AI Vector Search	including on- prem databases	with leading public clouds and cloud@customer
Oracle Autonomous AI Lakehouse				

Run AI and analytics securely on all your data without lock-in

Open and interoperable



- Open by design with Apache Iceberg
- Runs everywhere

Lakehouse without compromises



- The openness of Iceberg with the performance, reliability and trust of Oracle
- Extend existing warehouses to lakehouse architecture

AI-powered operations



- Operationalize AI insights directly in business operations
- Built-in AI, analytics, and compliance across **all** data

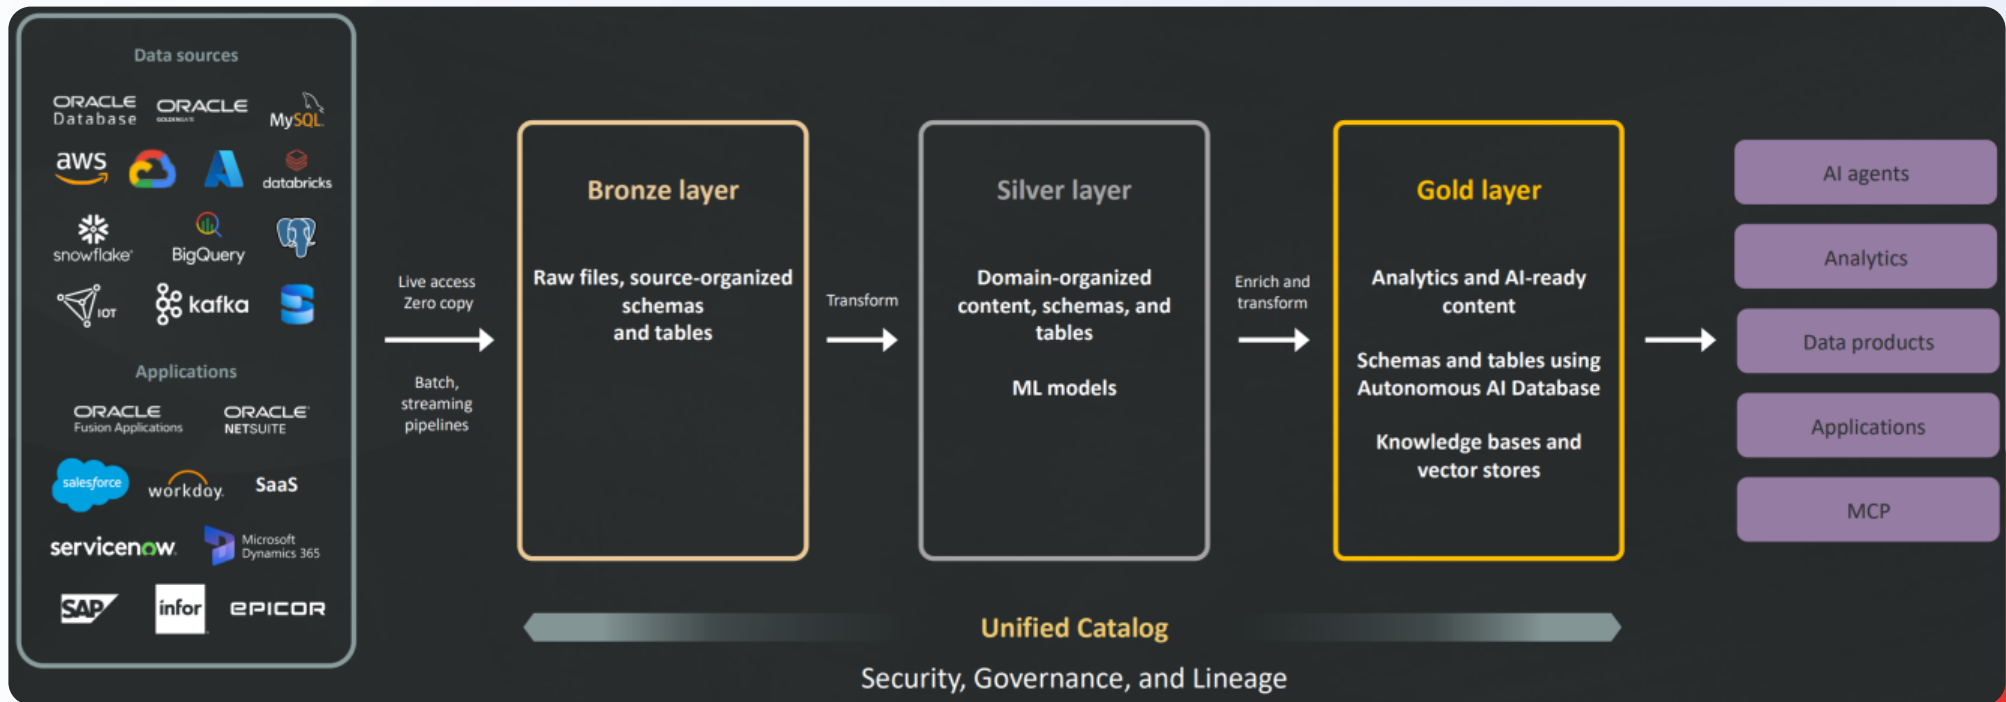
What makes the Autonomous AI Lakehouse Awesome?

○ Key Benefits:

- Unified view of your data via the catalog, so AI can actually find it
- Works the same in every cloud (OCI, AWS, Azure, GCP)
- Takes care of tuning and optimization for you
- Leverages open formats like Iceberg and Parquet so you keep control
- Super secure with built-in guardrails
- Plays nicely with other tools you already use
- Fast analytics on huge object storage datasets
- Saves money by using efficient storage and serverless tech

Data for AI

Enterprise-grade lakehouse foundation



Oracle AI Data Platform (AIDP)

- Oracle AI Data Platform brings together data, analytics, and AI in a single, governed environment. It unifies the entire data-to-AI lifecycle—from ingestion and preparation to model development and deployment—on a secure, scalable foundation.
- With integrated services like Autonomous AI Database, Oracle Analytics Cloud, Apache Spark and OCI Generative AI, AI Data Platform helps teams prepare, govern, and apply AI to their data efficiently.
- Whether you're engineering data, developing models, or building intelligent applications, AI Data Platform provides the speed, governance, and integration needed to operationalize AI across the enterprise.

Oracle AI Data Platform

Autonomous AI Database

Analytics

Open-Source Engines

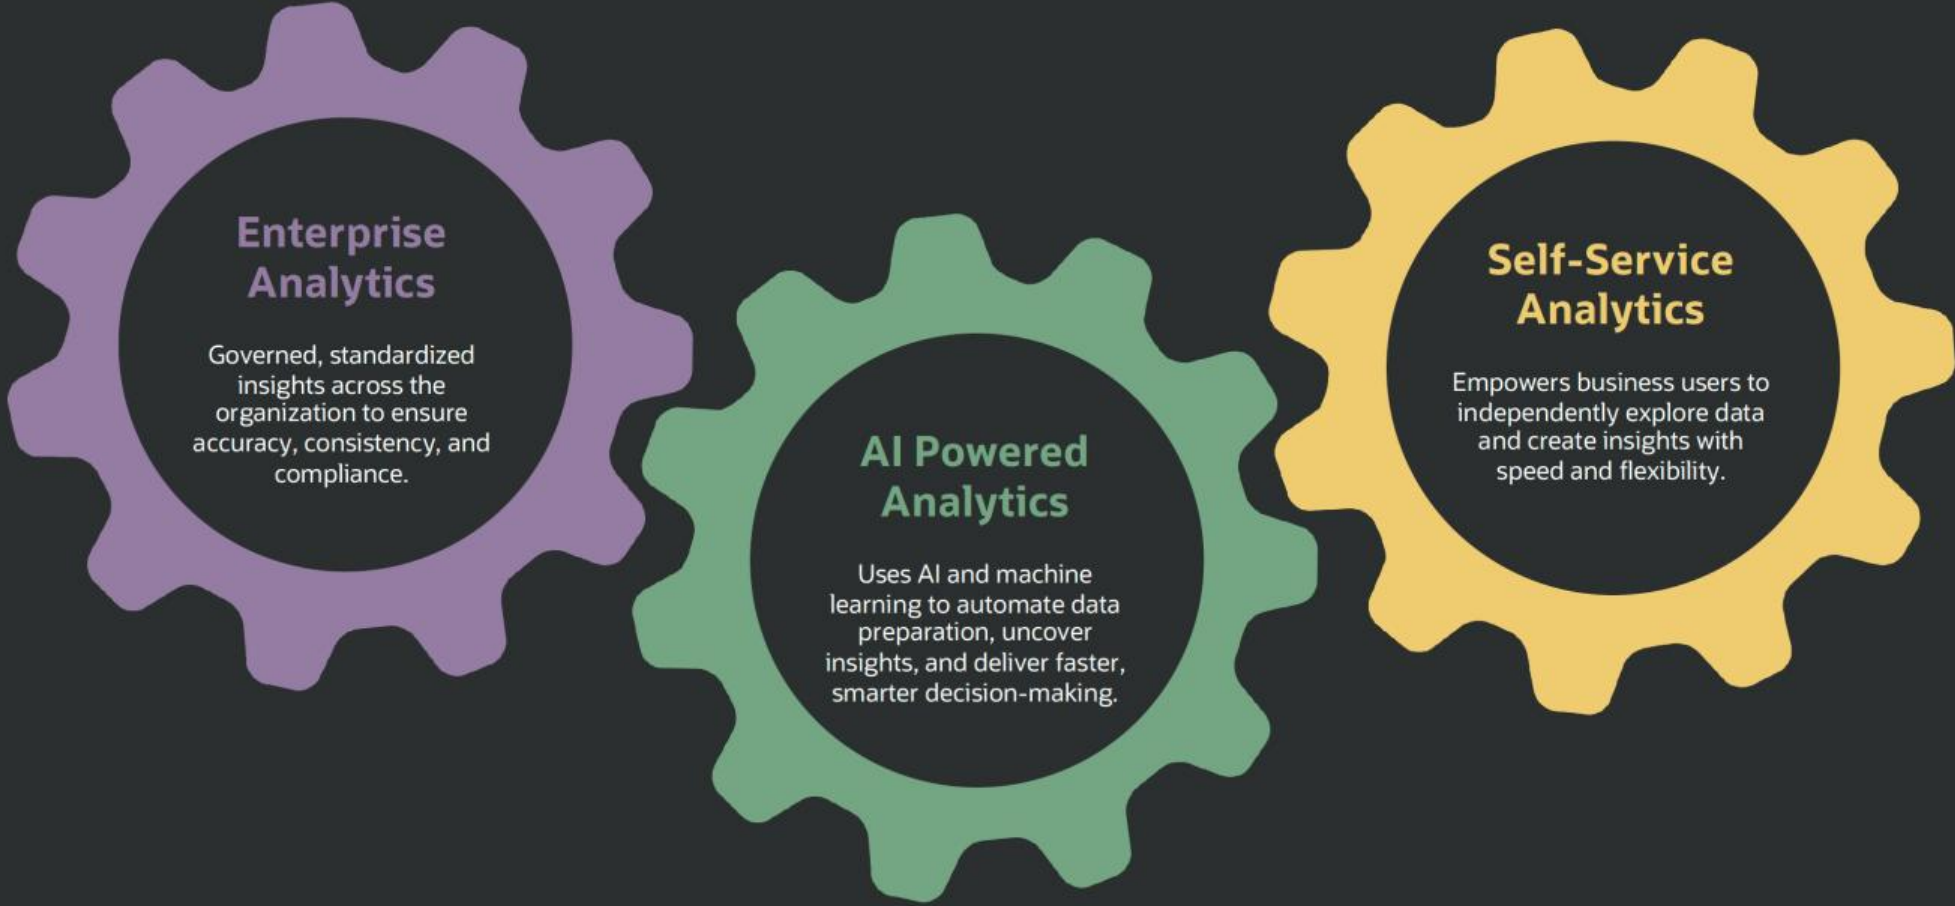
Infrastructure

Generative AI

Seamlessly integrates multiple OCI services including compute, storage, database, analytics, and AI into a cohesive experience that enables organizations to build, deploy, and manage AI solutions at enterprise scale

Oracle Analytics

Converged Analytics Platform
for all personas, workloads
and data



Enterprise Analytics

Governed, standardized insights across the organization to ensure accuracy, consistency, and compliance.

AI Powered Analytics

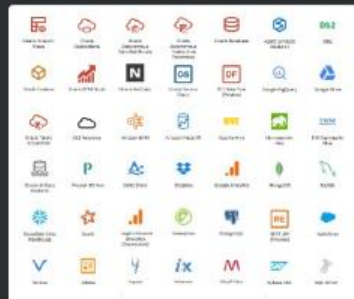
Uses AI and machine learning to automate data preparation, uncover insights, and deliver faster, smarter decision-making.

Self-Service Analytics

Empowers business users to independently explore data and create insights with speed and flexibility.

Self-Service Analytics

Open Connectivity



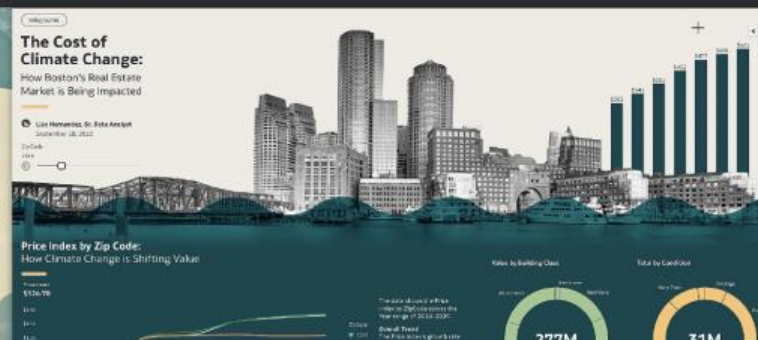
Data Preparation



Data Modeling



Data Visualization



AI Powered Analytics



Data Enrichment



AI-Assistants



Contextual Insights



Advanced Analytics



AI Storytelling



No-code Custom ML



Explain



Natural Language Interactions



Natural Language Generation



Auto Insights

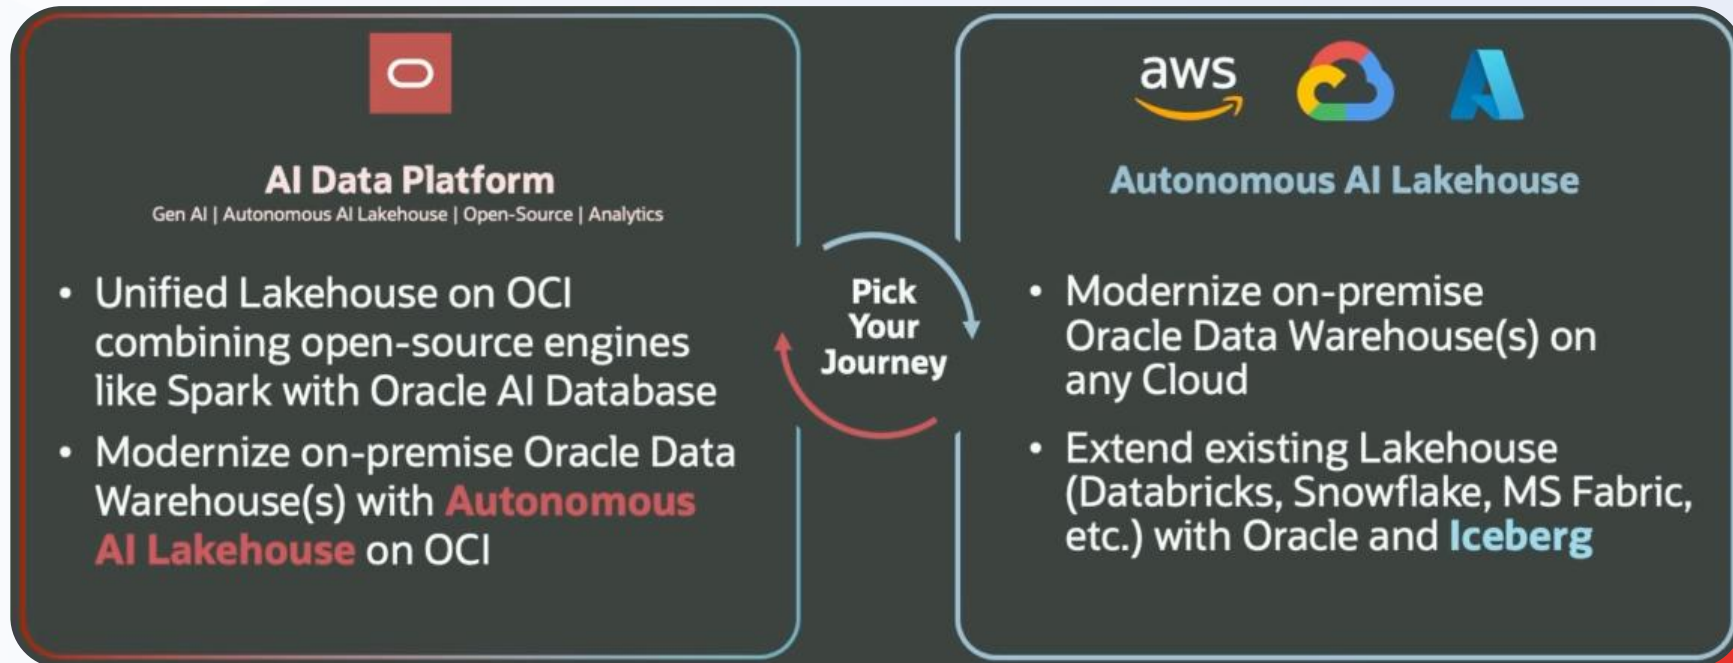


OCI AI Services Integration



Database ML Integration

Complete and Comprehensive on OCI or multicloud



Tripwire Solutions: How can we help you?

- Assessments
- Setup your Data Platform environments
- Team enablement and Training
- Provide Experts for point solutions

Q & A