



Modern Data Platform and Generative AI

Oracle Cloud Infrastructure - The next frontier of enterprise innovation




Safe Harbor Statement

This presentation 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.

Agenda


- Oracle Cloud Infrastructure
- Oracle Modern Data Platform
- AI for Enterprise

Key Challenges for Data and AI Leaders



Strategic & Technology

- Lack of data-driven innovation
- Reduced D&A agility
- Proliferation of stealth D&A
- Multiple points of failure



Organizational

- Loss of influence for D&A team
- Incorrect use of data for decision making
- Non-data-driven approaches to decision making



People

- Employee burnout
- Dependency on dashboard and data silos
- Difficulty sourcing talent for multiple technologies



Financial

- Misallocation of resources
- Higher TCO for D&A projects
- Underused Investments



ORACLE

AI for Enterprises



Generative AI at Enterprise Requires

High quality, clean and well understood Data with known data lineage to base training upon ,while minimize job **failures** due to Infra.

Governance processes will need to be extended to address potential ethical concerns related to the content generated by AI, including bias mitigation and content quality review.

Collaboration between Data Scientists, Business Domain experts, Developers, and Architects.

Generative AI requires a **cultural change and new roles within the organization** (ethics, biases, etc.) with support to users who will interact with generative AI outputs.

Training large generative AI models is very **computationally intensive** (GPUs) and thus rather expensive.

Are you ready to accepts a **results accuracy max up to ~80%**? What does it really mean for your Business? **Predictable Price Performance**

PROCESS

PEOPLE

TECHNOLOGY

AI innovation at Oracle



Address
enterprise
generative AI
requirements

Embed generative
AI across every
layer of the tech
stack

Prioritize data
management,
security and
governance

AI across the Oracle Cloud ecosystem





Applications

- Fusion Applications
- NetSuite
- Fusion Analytics
- Industry Applications
- 3rd Party Applications

AI Services

-  OCI Generative AI
-  OCI Generative AI Agents
-  Digital Assistant
-  Speech
-  Language
-  Vision
-  Document Understanding




ML for data platforms

-  OCI Data Science
-  AI Vector Search in Oracle Database
-  MySQL HeatWave Vector Store
-  OCI Data Labeling

Modern Data Platform

-  No Code Data Integration
-  Real Time Streaming Analytics
-  Data Lake
-  Data Warehouses
-  Data Catalog
-  Self Service Augmented Analytics
-  No Code App Builder

AI infrastructure

-  Compute bare metal instances and VMs
-  OCI Supercluster with RDMA networking
-  Block, object, and file storage; HPC filesystems

Oracle AI Partners

OCI Generative AI

Now Available: OCI Generative AI Service

The screenshot shows the OCI Generative AI overview page. On the left is a sidebar with navigation links: Overview, Playground, Dedicated AI clusters, Custom models, and Endpoints. Below these are sections for Scope and a compartment selector. The main content area is titled 'Generative AI overview' and includes a hero section with a description of the service and a 'Watch service tour' button. Below this is a 'Metrics in my compartment' section with three cards: 'Dedicated AI clusters' with a value of 7, 'Custom models' with a value of 3, and 'Active endpoints' with a value of 12. To the right of these metrics is a 'Resources' section with links to documentation, API reference, workshops, tutorials, and pricing. At the bottom is a 'Get started' section with a 'Go to playground' button and three cards: 'Playground' (describing the visual interface), 'Dedicated AI clusters' (describing hardware units for fine-tuning), 'Custom models' (describing fine-tuning with a dataset), and 'Endpoints' (describing creating and managing endpoints).

High quality pre-built models

We are offering high quality models from Meta and Cohere to meet your business needs with minimal effort

Customize models to meet your needs

Fine tune models with your own data and focus on your most important tasks


Fully hosted inside OCI

All processing and data storage happens inside OCI, no cross-region or cross-cloud communication

Private and secure

Customer-provided training and inference data is secure and cannot be seen by other customers. Oracle does not send customer data to Cohere or Meta.

New models included in the GA release



Llama-2 70B

The Llama-2 70B parameters text generation model developed by Meta that is the leading open source LLM. Free for research and commercial use.



Command

Command is Cohere's highly performant generation model. In two sizes: 6B (Medium) and 52B parameters (XL). Using the Command XL model provides better accuracy while Command M has lower latency and cost.

Summarize

The Summarize model provides high-quality summaries that accurately capture the most important information from your documents.

Embed

The English and multi-lingual Embedding model (V3) that converts text to vector embeddings. A 'light' version of the model exists that is smaller and faster but is slightly less performant (English only).

OCI Gen AI: New Features



New Models

- Meta Llama 2 models
- Cohere Multilingual and English Embed V3



Flexible fine tuning

- Fine-tuning of both Cohere Command 52/6B models
- Vanilla and T-few fine-tuning options along with fine-tuning parameters configuration
- Stacked fine-tuned model serving



Improved cluster UX

- Multi-endpoint support in hosting clusters
- Scaling your clusters by adding/removing units to handle more model requests
- Endpoint Analytics (tokens processed in/out, number of calls, etc.)


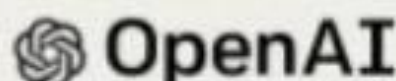
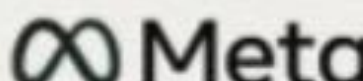








GenAI Ops

- Content moderation controls
- Endpoint model swap with zero-down time.
- Endpoints deactivation/activation
- LangChain Integration



A leader and differentiated relative to the competition

		 cohere	 OpenAI		ANTHROPIC		Google		 Meta	
		Command	GPT-3.5	GPT-4	Claude Instant	Claude 2	Gemini Pro	Gemini Ultra	Llama 2 13B	Llama 2 70B
Customer Focus		Enterprise-focused	←		Consumer and enterprise		→		Research and enterprise	
Accuracy		Accurate MMLU 69.2% → 79% (Q1 projected)	Accurate MMLU 70.0%	Leading MMLU 86.4%	Accurate MMLU 73.4% ¹	Accurate MMLU 78.5% ¹	Accurate MMLU 71.8%	Leading MMLU 83.7%	Specialized MMLU 54.8%	Accurate MMLU 68.9%
TCO		Leading TCO (Bedrock prov. TCO ² : \$1.93 / 1M tokens)	Competitive TCO	Cost Prohibitive	Competitive TCO 1.3x Command ²	Cost Prohibitive 5.7x Command ²	Competitive TCO	Cost Prohibitive	Competitive TCO 1.8x Command ²	Not benchmarked
Deployment		Any cloud or private environment with enterprise support	Limited cloud availability (Azure only) No private options		Limited cloud availability (AWS, Google) No private options		Limited cloud availability (Google only) No private options		Any cloud or private environment w/o enterprise support	
Retrieval-Augmented Generation (RAG) ³		Complete solution (generative, embeddings, rerank, connectors) Avg Lost-in-the-Middle score Command: 78%	Partial solution (no rerank) Avg Lost-in-the-Middle score GPT-3.5: 67%		Incomplete solution (generative model only)		Partial solution (no rerank, no connectors)		Incomplete solution (generative model only) Avg Lost-in-the-Middle score Llama 2 70B: 72%	
Fine-tuning & Support		Fine-tunable Vendor Support Cohere engineers	Fine-tunable	Roadmap	Roadmap		Fine-tunable		Fine-tunable	
			Vendor Support OpenAI engineers (\$2-3m)		Vendor Support Not available		Vendor Support Not available		Vendor Support Not available	

¹ Claude MMLU evaluations are based on 5-shot COT, while standard evals are based on 5-shot. What's this mean? 5-shot COT will have inflated results compared to 5-shot.

² TCO evaluations are benchmarked by Cohere comparing models on Amazon Bedrock controlling for hardware and workloads and based on existing list prices. OpenAI and Google models are not available for benchmarking in similar environments.

³ Lost-in-the-middle eCopyright © 2024, Oracle and/or its affiliates. based on: <https://arxiv.org/abs/2307.03172>.



Retrieval Augmented Generation for OCI Generative AI Agents

Introducing: OCI Generative AI Agents

RAG Agent, the first in a series of Gen AI Agents

Featuring retrieval-augmented generation

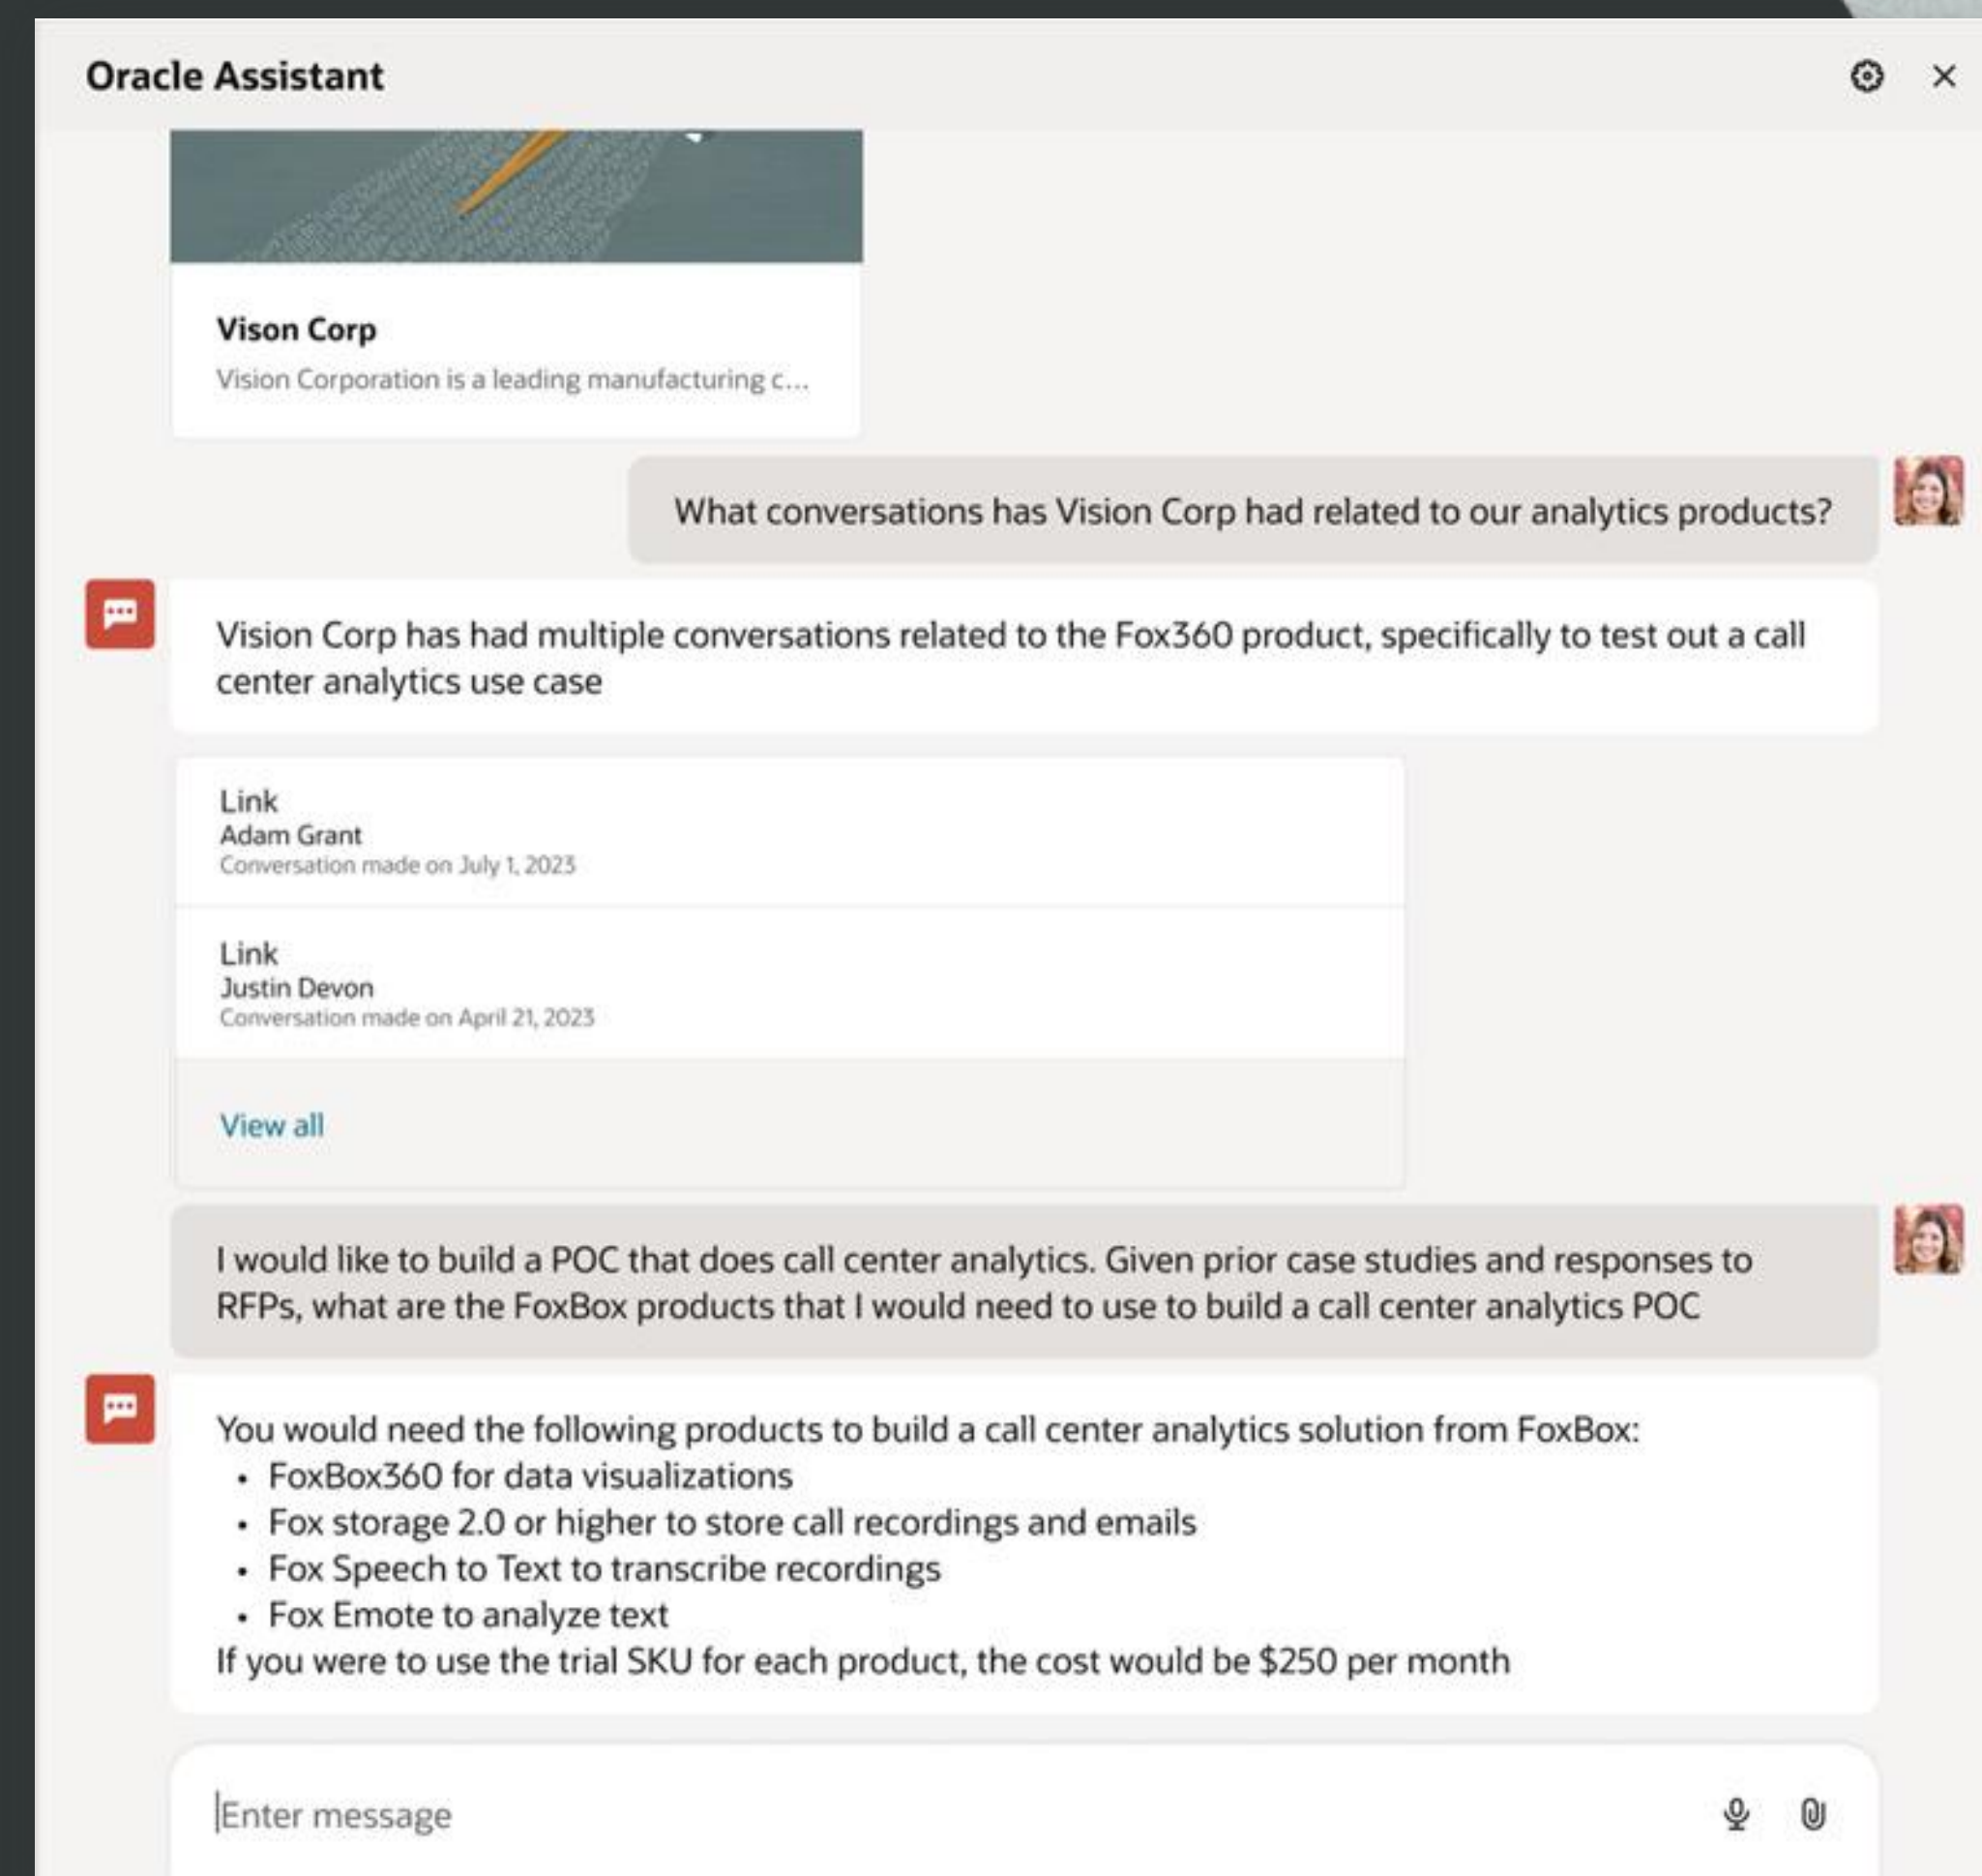
OCI GenAI RAG Agent connected to
Vision Corp's knowledge bases

Analyst asks a natural language question

GenAI RAG Agent responds in humanlike manner
and provides links to relevant source documents

Analyst asks a follow up question

GenAI RAG Agent uses chat history and further
information retrieval to respond



Core capabilities of OCI Generative AI Agents

RAG Agent

Agents act on knowledge bases such as search indexes, for example:

- call transcripts
- internal knowledge sources
- other large corpuses of enterprise proprietary data

Beta:
OCI OpenSearch

Coming Soon:
Oracle Database 23c AI Vector Search
MySQL HeatWave Vector Store

AI Quick Actions for OCI Data Science



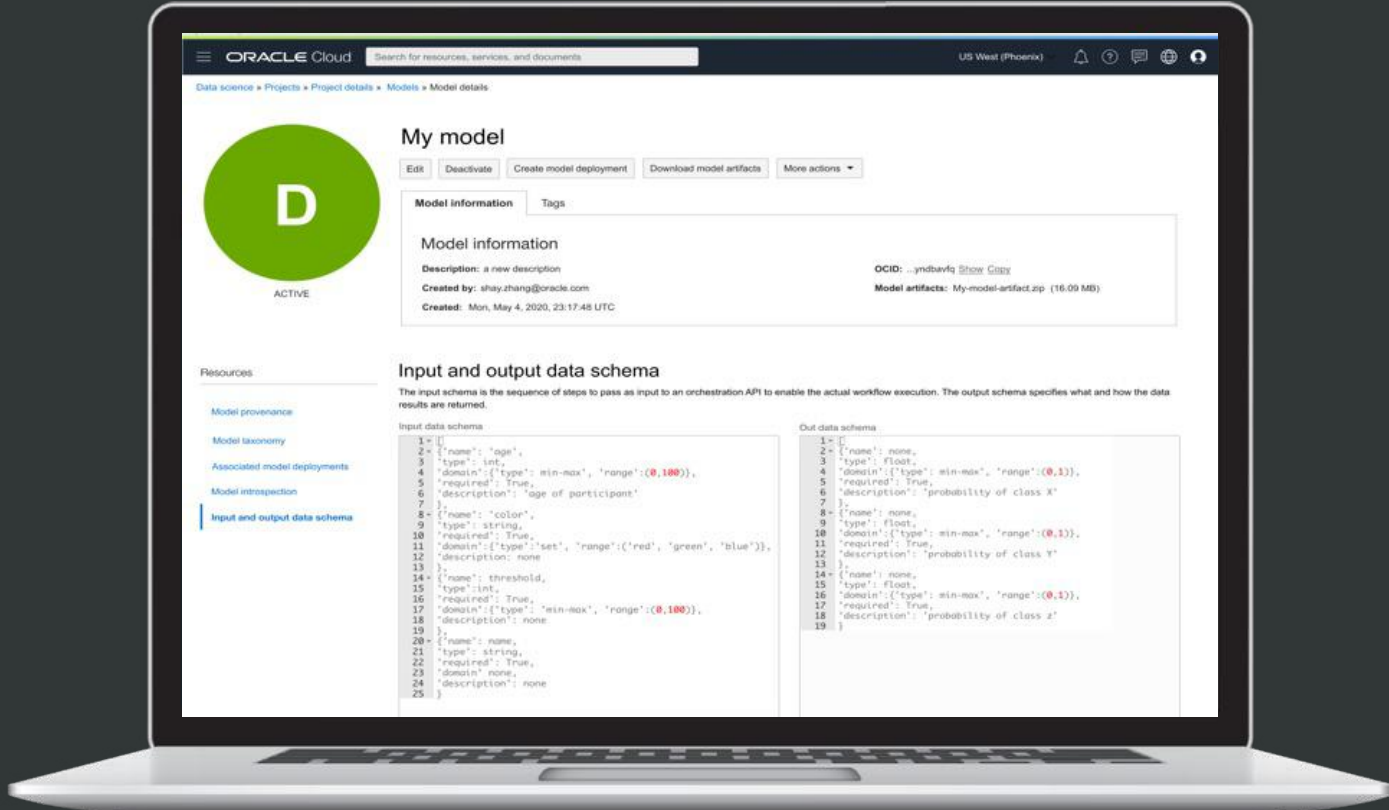
OCI Data Science

Open and extensible

Use your favorite open-source tools and framework

Use your AI anywhere using ONNX

Accelerated Data Science SDK available as an open-source library



Oracle Cloud Infrastructure Data Science

Develop & Experiment

Notebook
Sessions



Local and OCI
Code Editors



Manage, Share, & Reproduce

Model Catalog



Feature Store*



Conda
Environments



Operationalize with MLOps


Jobs



Pipelines



Model
Deployments



Model
Monitoring*



Integrate

ML
Applications*



Data Management

Database – Data Lake – Access – Integration – Preparation

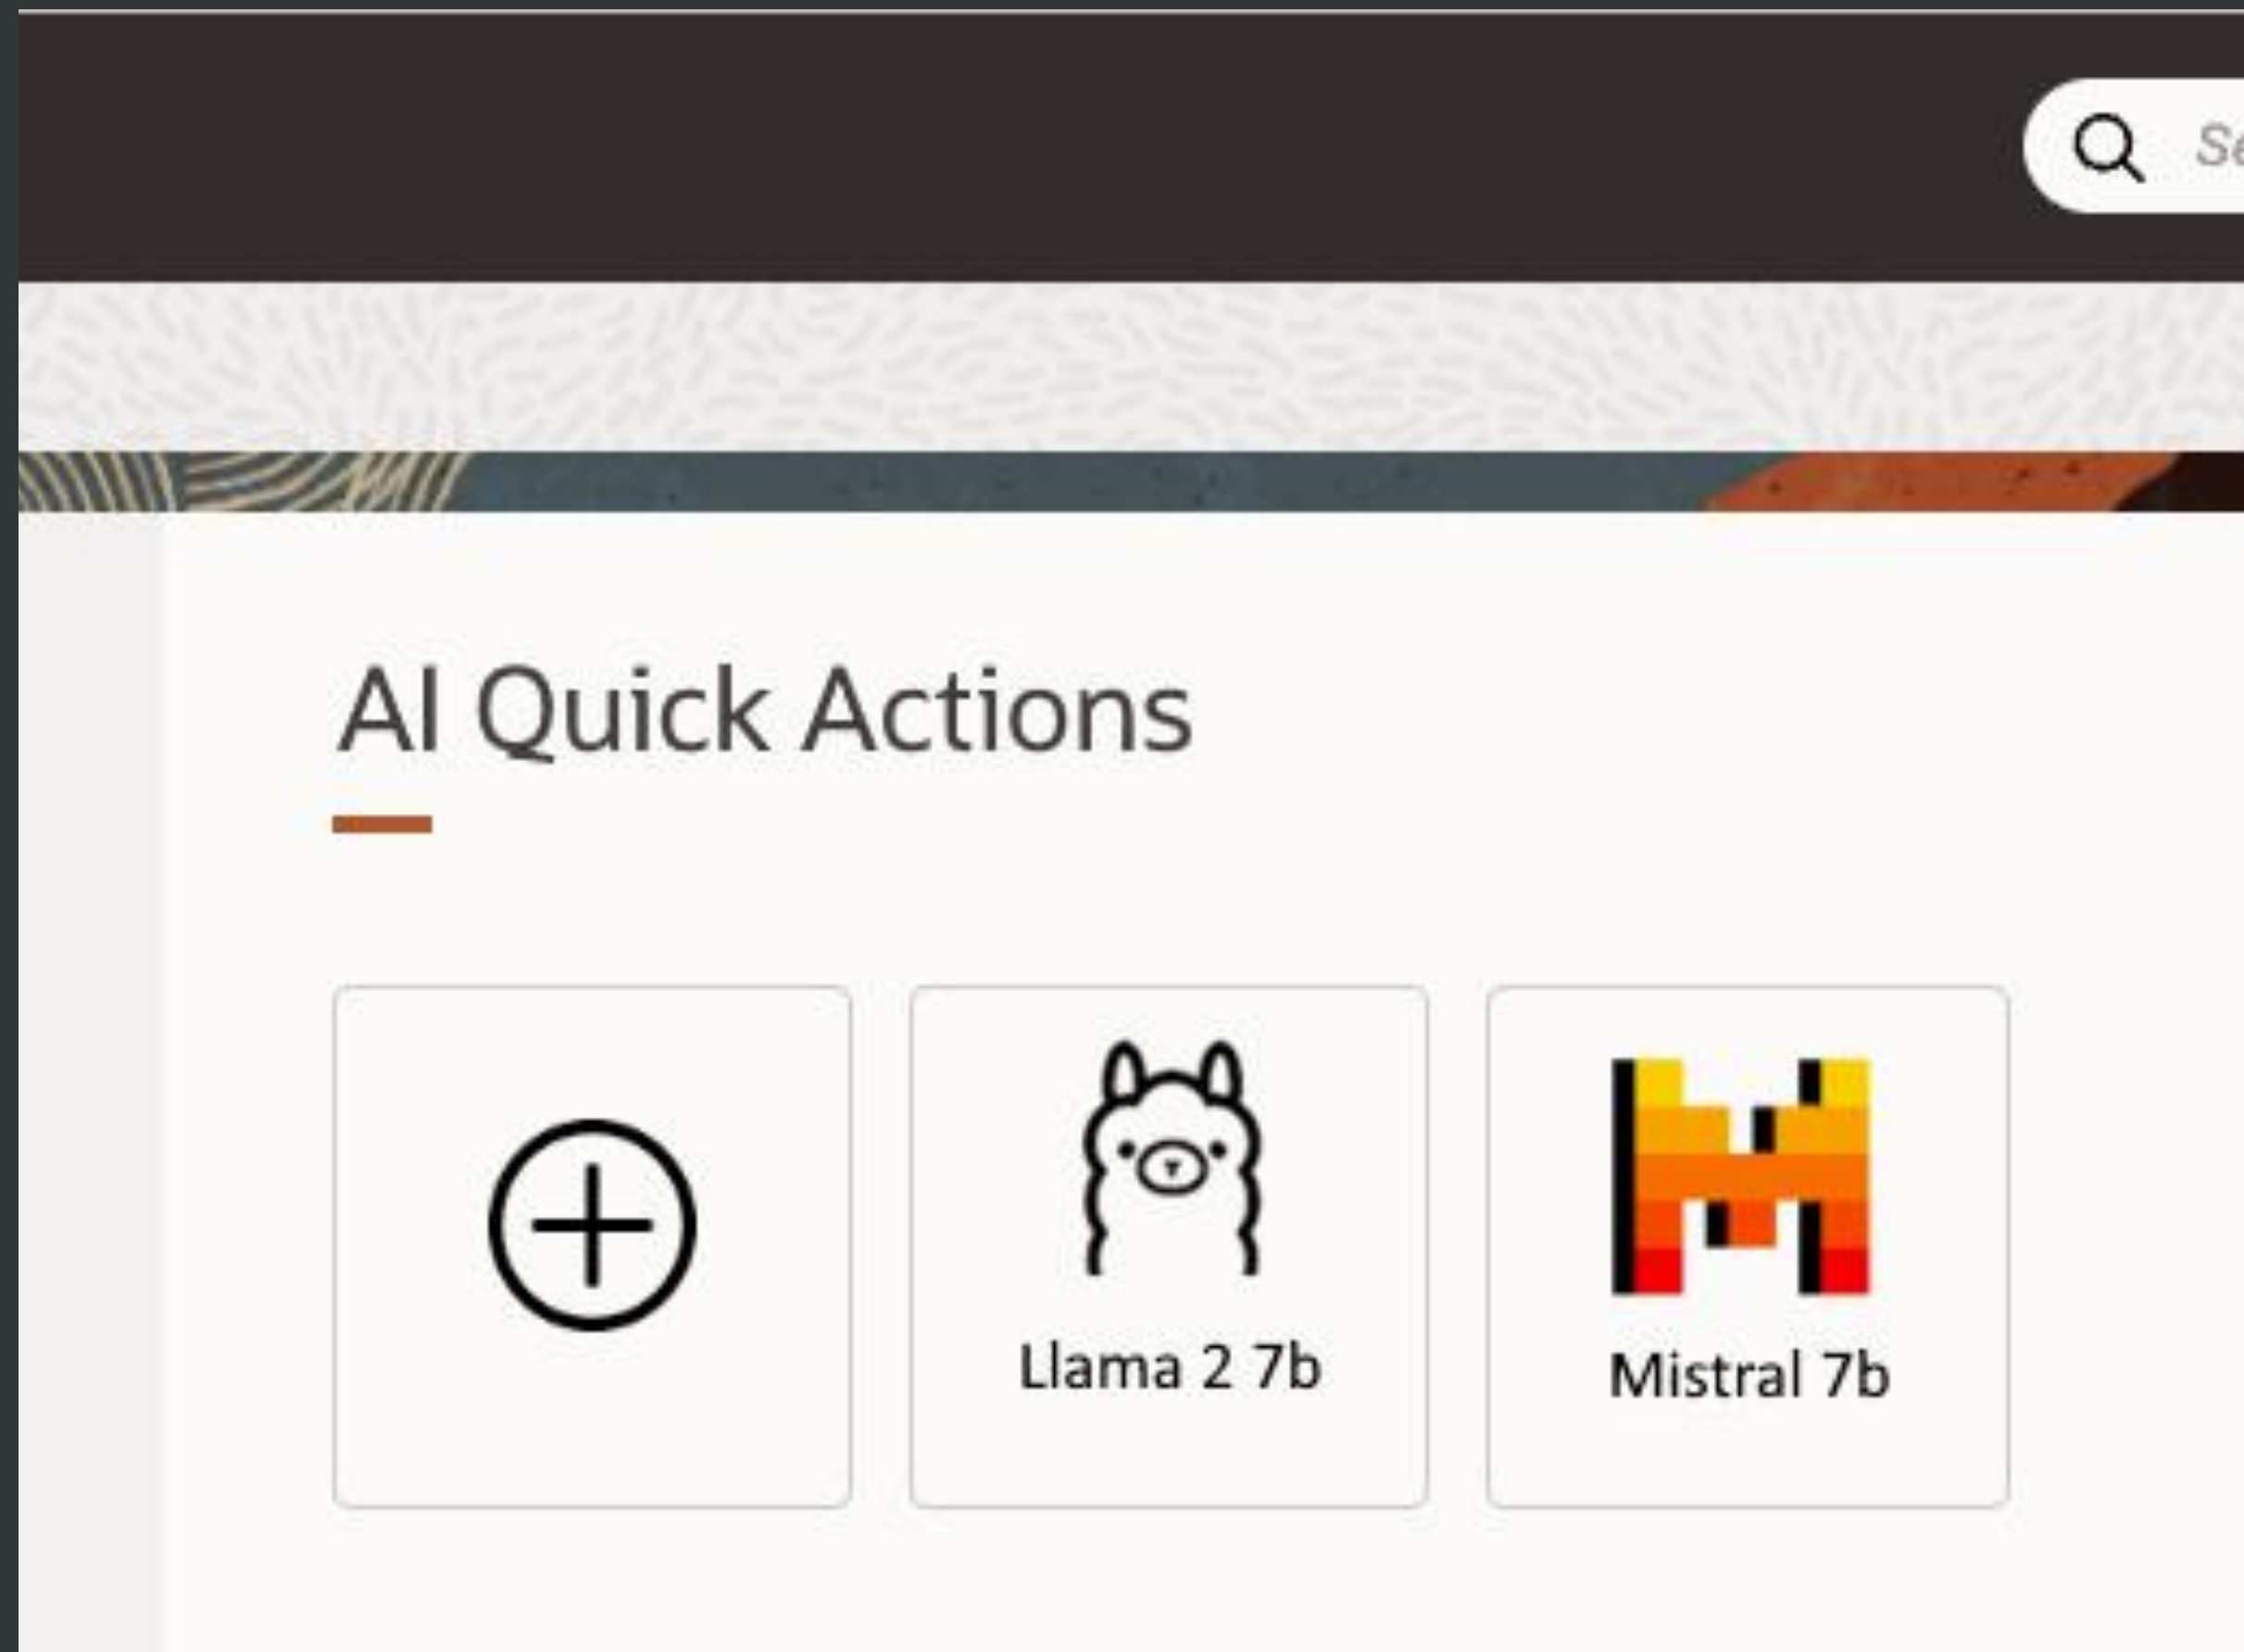
Infrastructure

CPU – GPU – Storage – Network

*Coming soon



Announcing: AI Quick Actions for OCI Data Science



A collection of use cases that can be invoked in a click of a button, all in a user interface inside the Data Science notebook experience.

Get no-code access to LLMs such as Llama2 and Mistral 7B through seamless integration with Data Science notebooks.

Quick actions:

- Deploy
- Fine tune
- Integrate
- Scale

Deploy and Fine Tune

Deployment

- Provides support for model deployment beyond what is available in OCI Gen AI Service with finer grained controls for those who need it
- Support for Text Generation Inference (Hugging Face), vLLM (UC Berkeley) and Nvidia Triton serving with public examples for:
 - Llama2 7b and 13b using A10s
 - Llama2 70b using A100 and A10s via GPTQ Quantization
 - Mistral 7b
 - Jina Embedding Model using A100

Fine Tuning

- Distributed Training with PyTorch, Hugging Face Accelerate and DeepSpeed for Fine-tuning of LLM
- Mount for Object Storage and File System as a Service – enables effortless checkpointing and storage of fine-tuned weights
- Service-provided Condas, eliminates the requirement for custom docker environments

Oracle AI Services


AI Services Benefits

AI for everyone



- An AI starting point, even without data science experience
- State-of-the-art LLMs from Meta and Cohere for generative AI use cases

Prebuilt for enterprise requirements



- Prebuilt models trained on industry-derived data
- Optimized for use cases across finance, manufacturing, and more
- Built-in insight highways into your SaaS applications with Oracle NetSuite, Fusion, and custom apps

Customizable for your needs



- Tailor your AI models without data science experience
- Save costs by training AI models already on OCI
- Fine tune generative AI models for specialized use cases

Best-in class support



- White glove treatment
- Data scientists on staff, dedicated to ensuring your organization's success
- Pricing to support AI experimentation

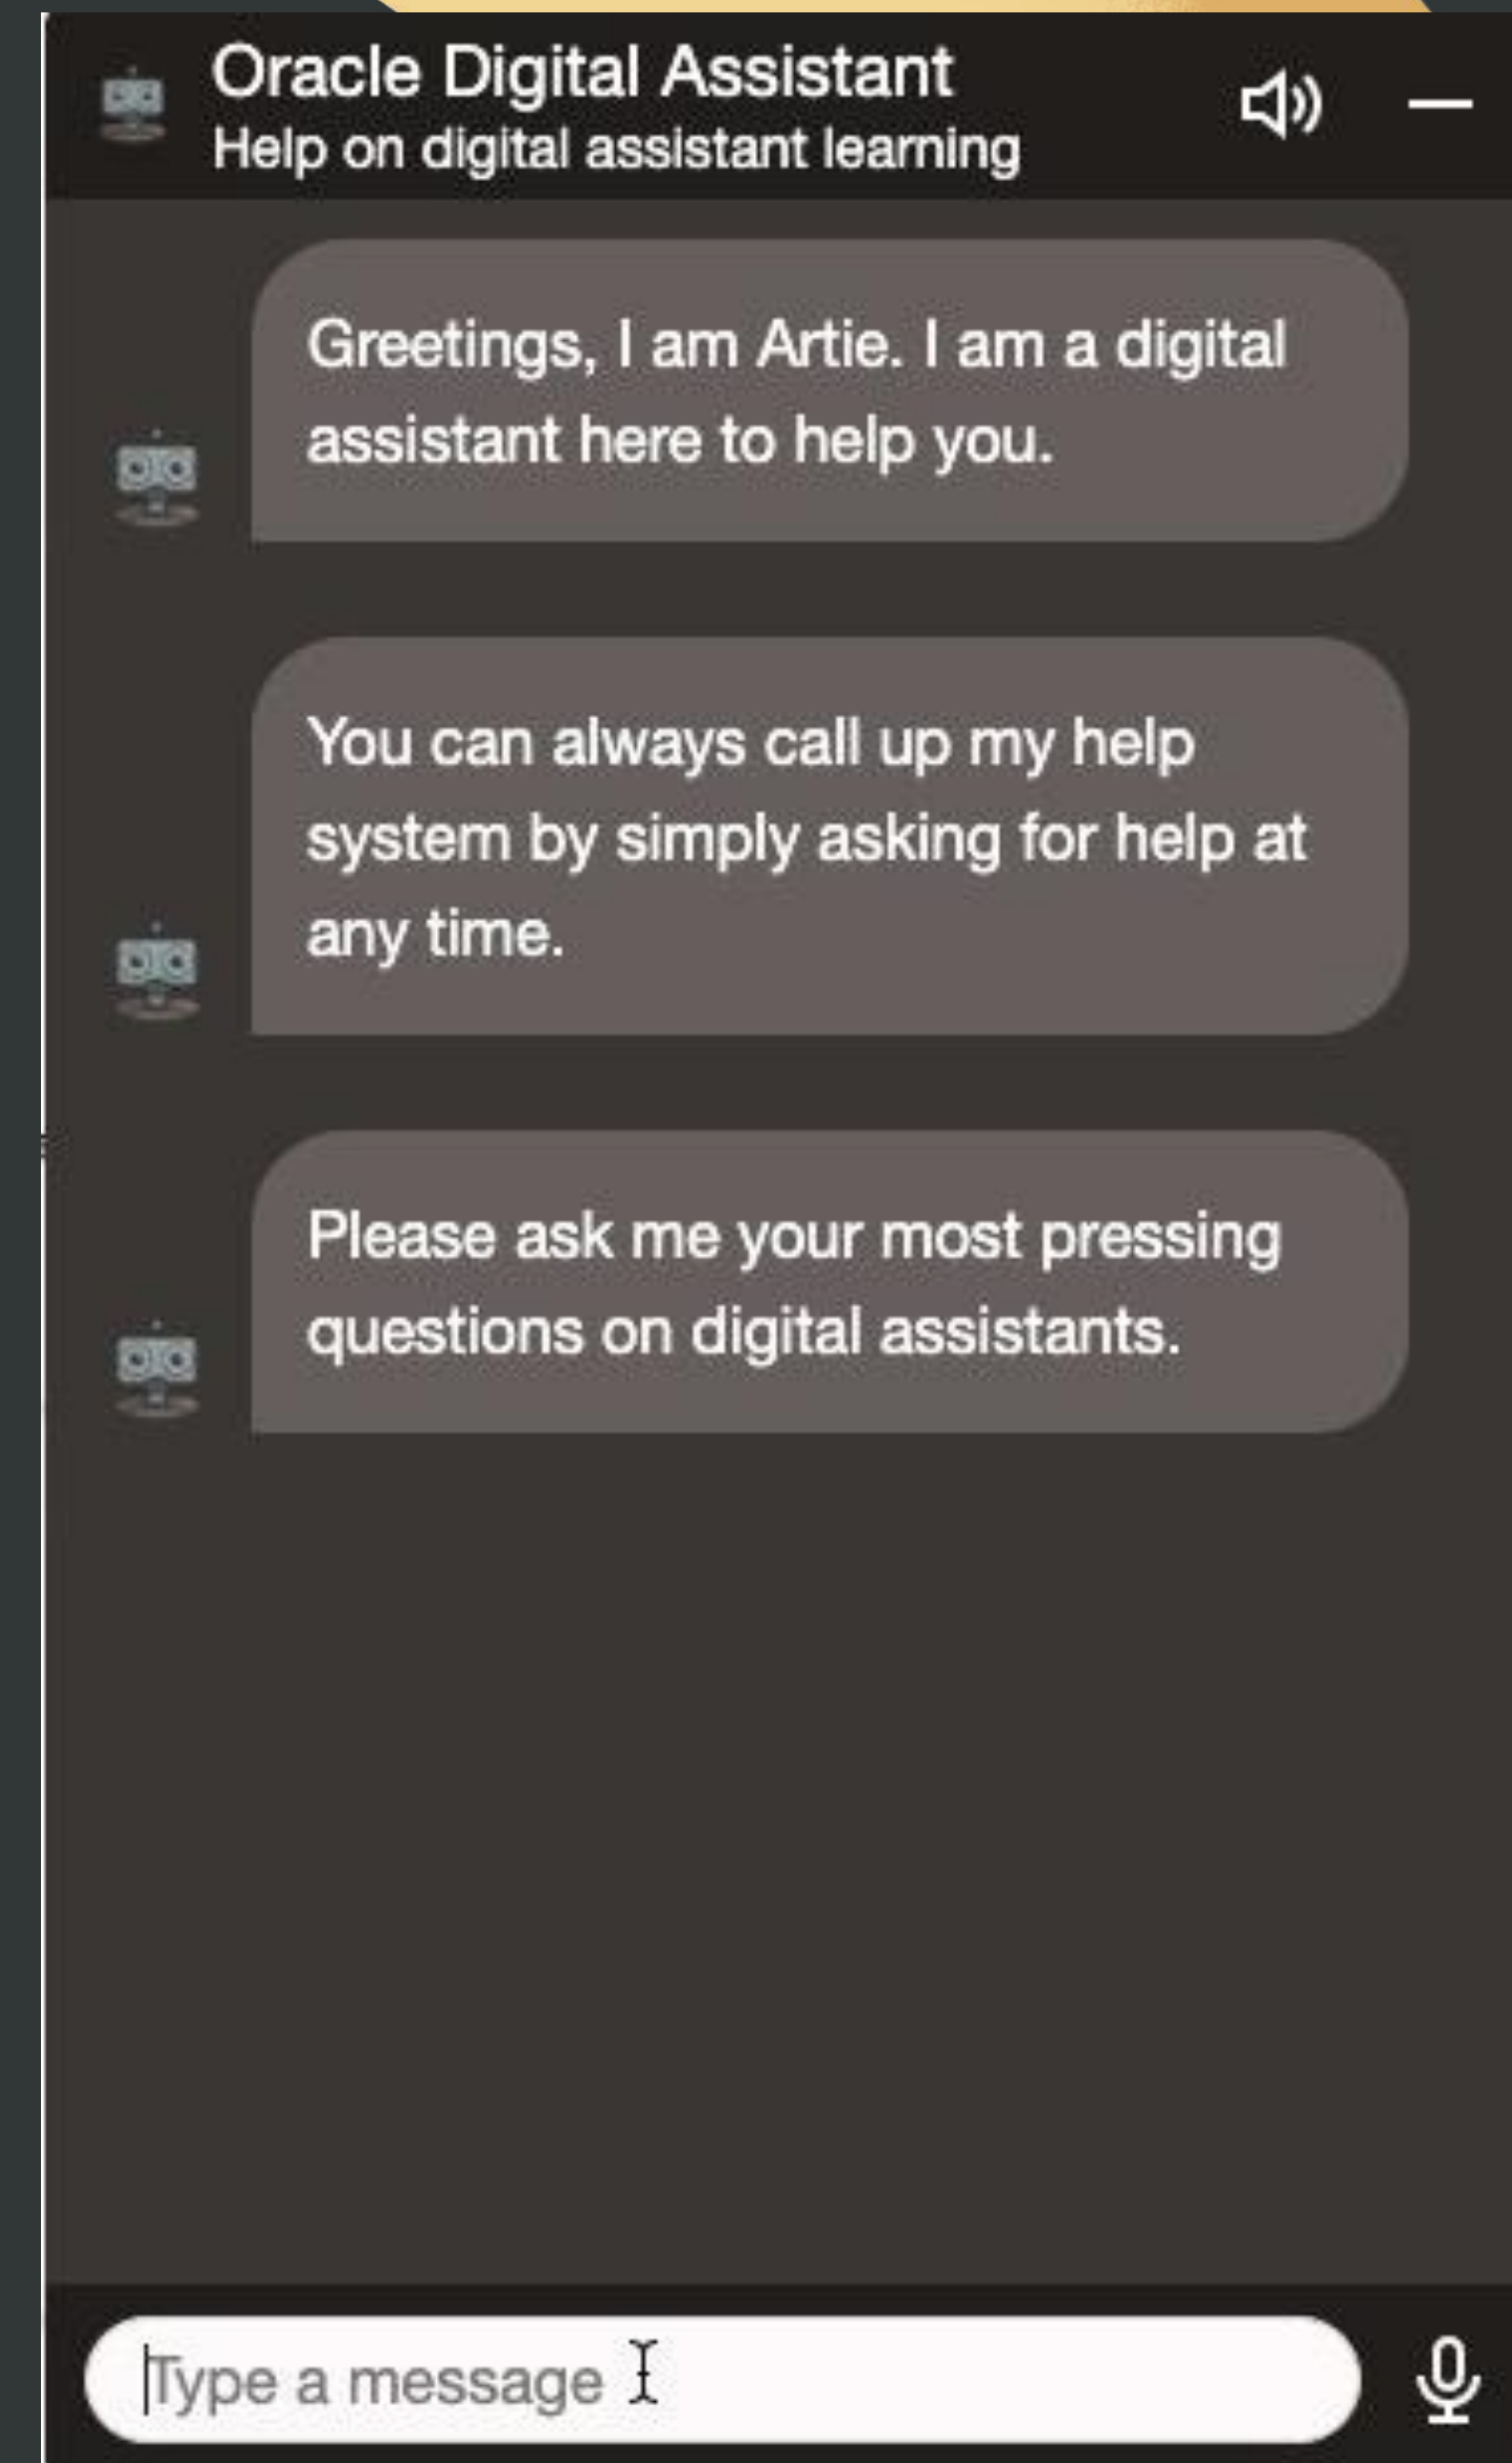
OCI Digital Assistant

Intelligent self service chatbot providing chat, text message, and voice interfaces

Improve self-service experiences with chatbots, voice interfaces, intuitive search

Reach new audiences by creating new application interfaces and dashboards

Create more intelligent customer experiences with added personalization, prediction, analytics, and more



OCI Language

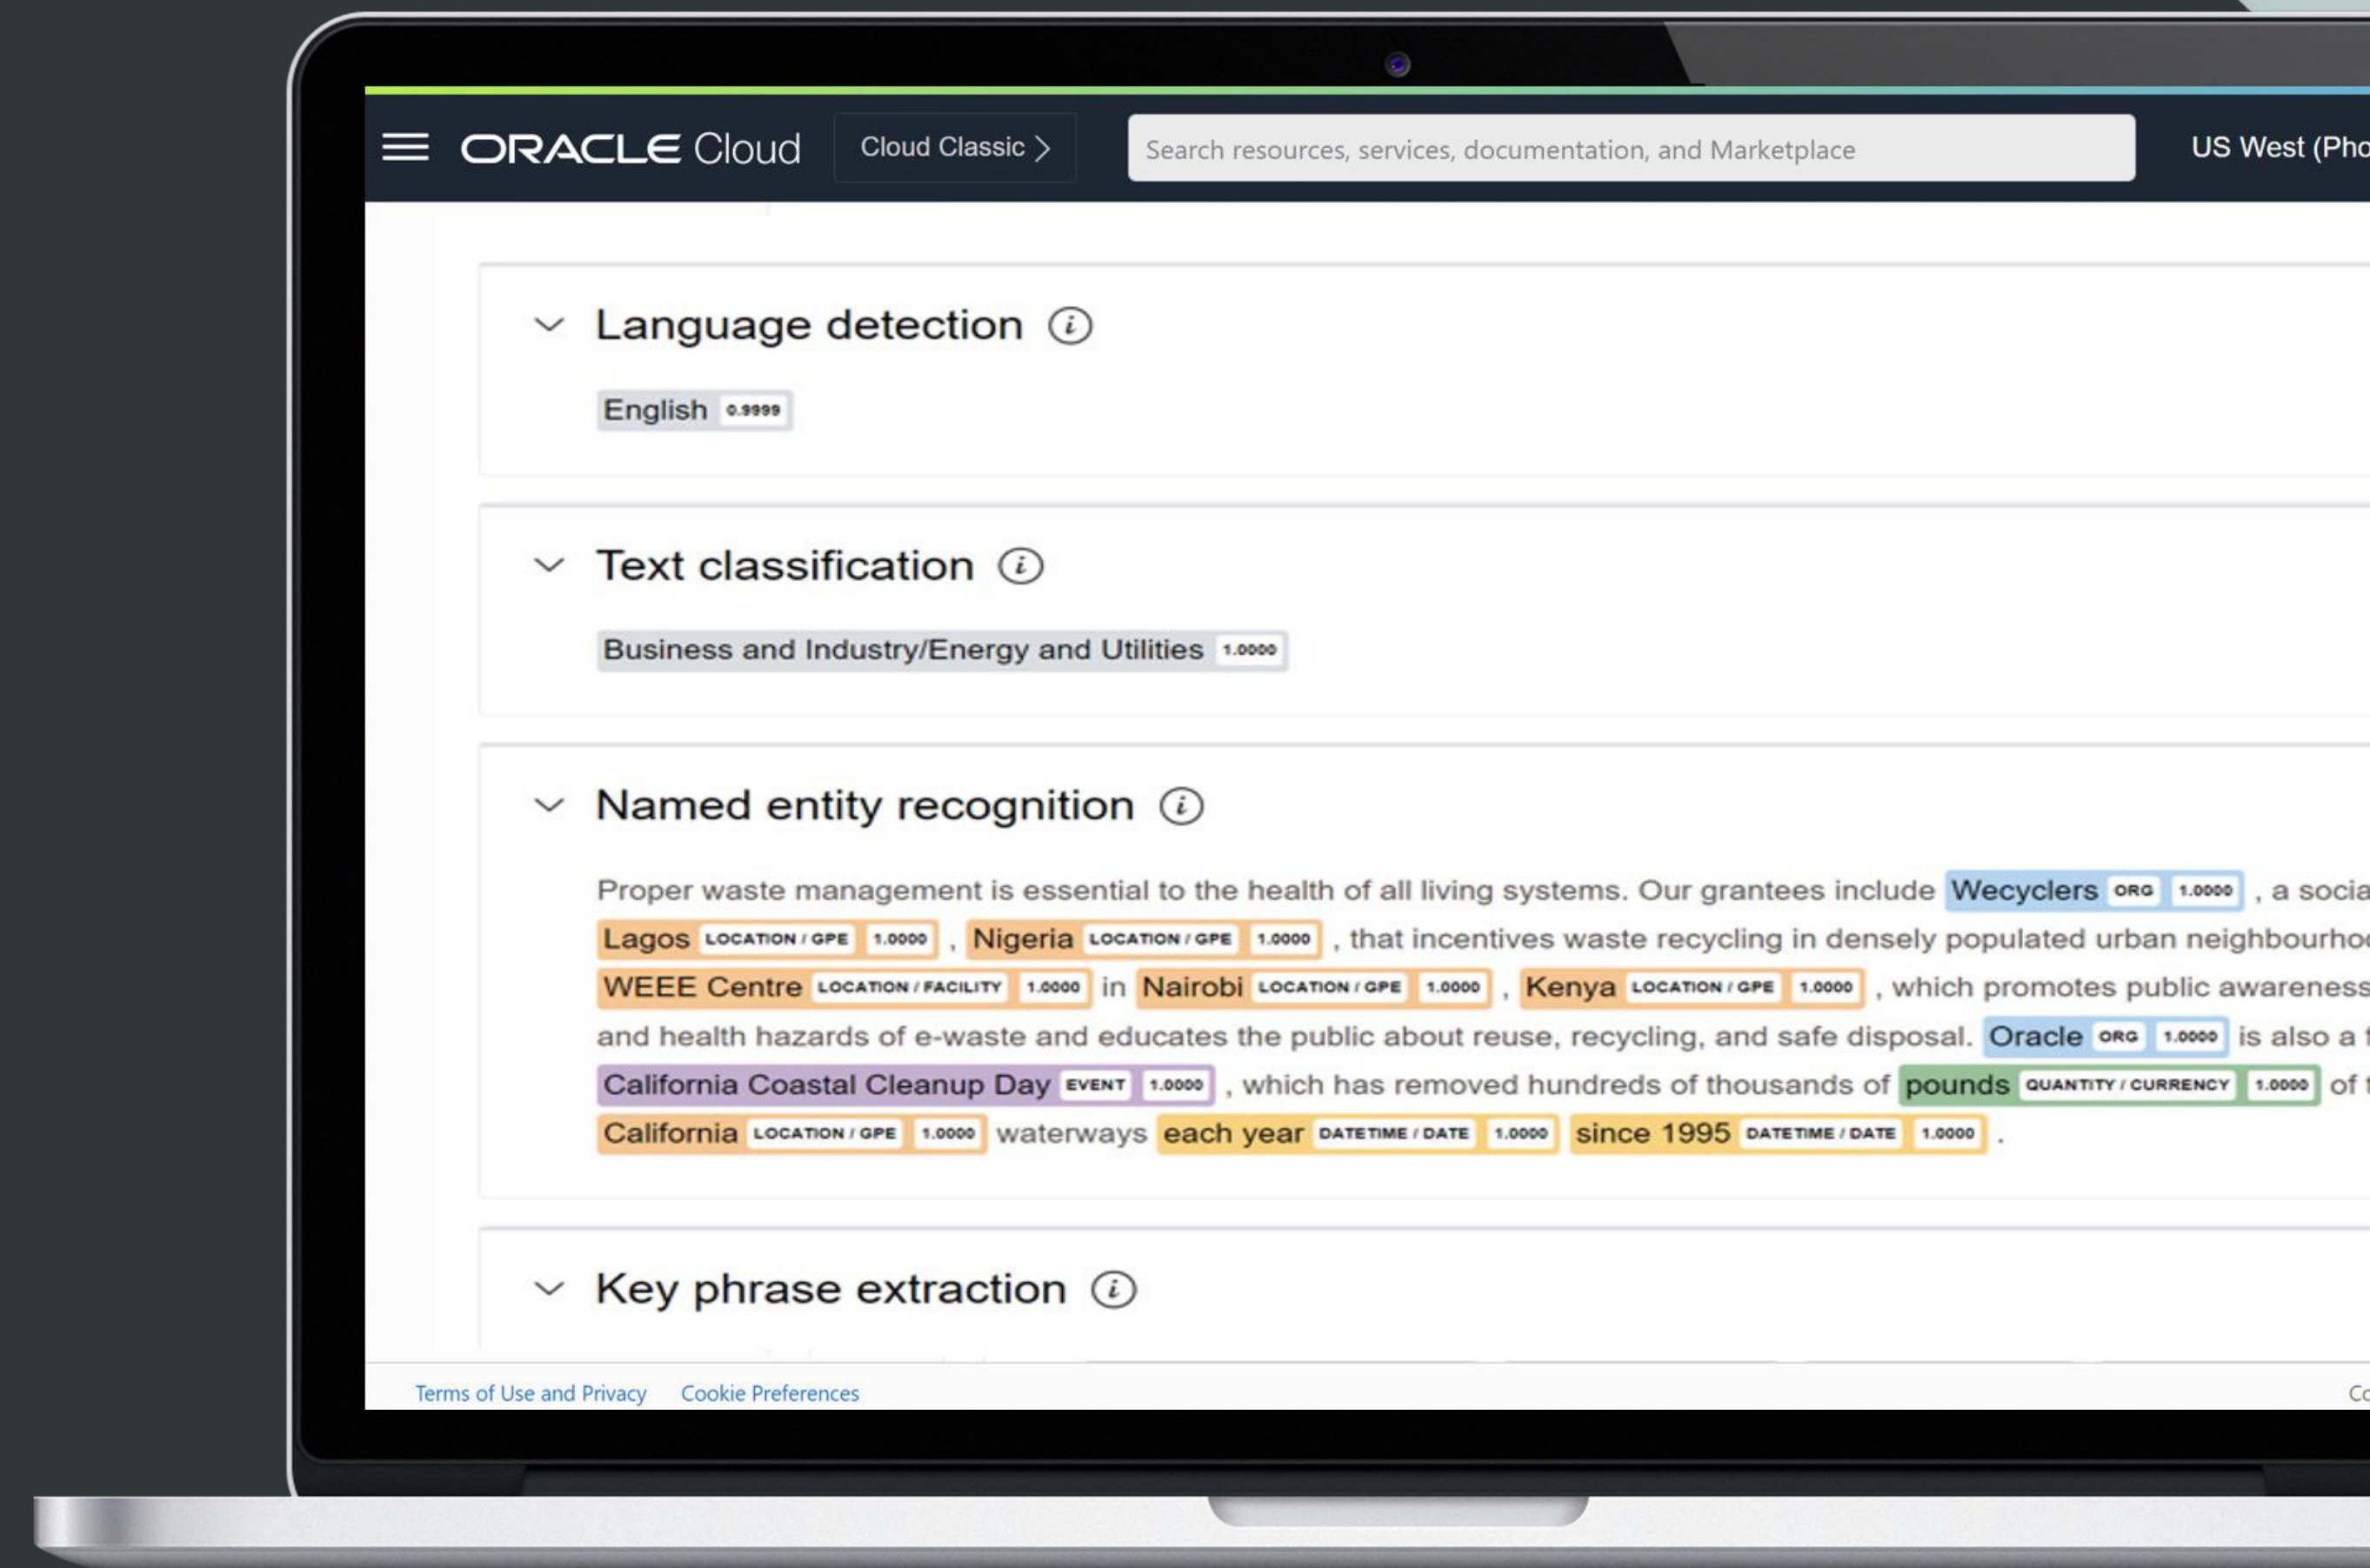
Text AI to classify sentiment and extract important entities

Analyze text at scale

Translate to/from 20+ languages to easily reach new audiences and consumers

Improve user experience by understanding sentiment from text

Recognize sensitive information to protect privacy and ensure compliance



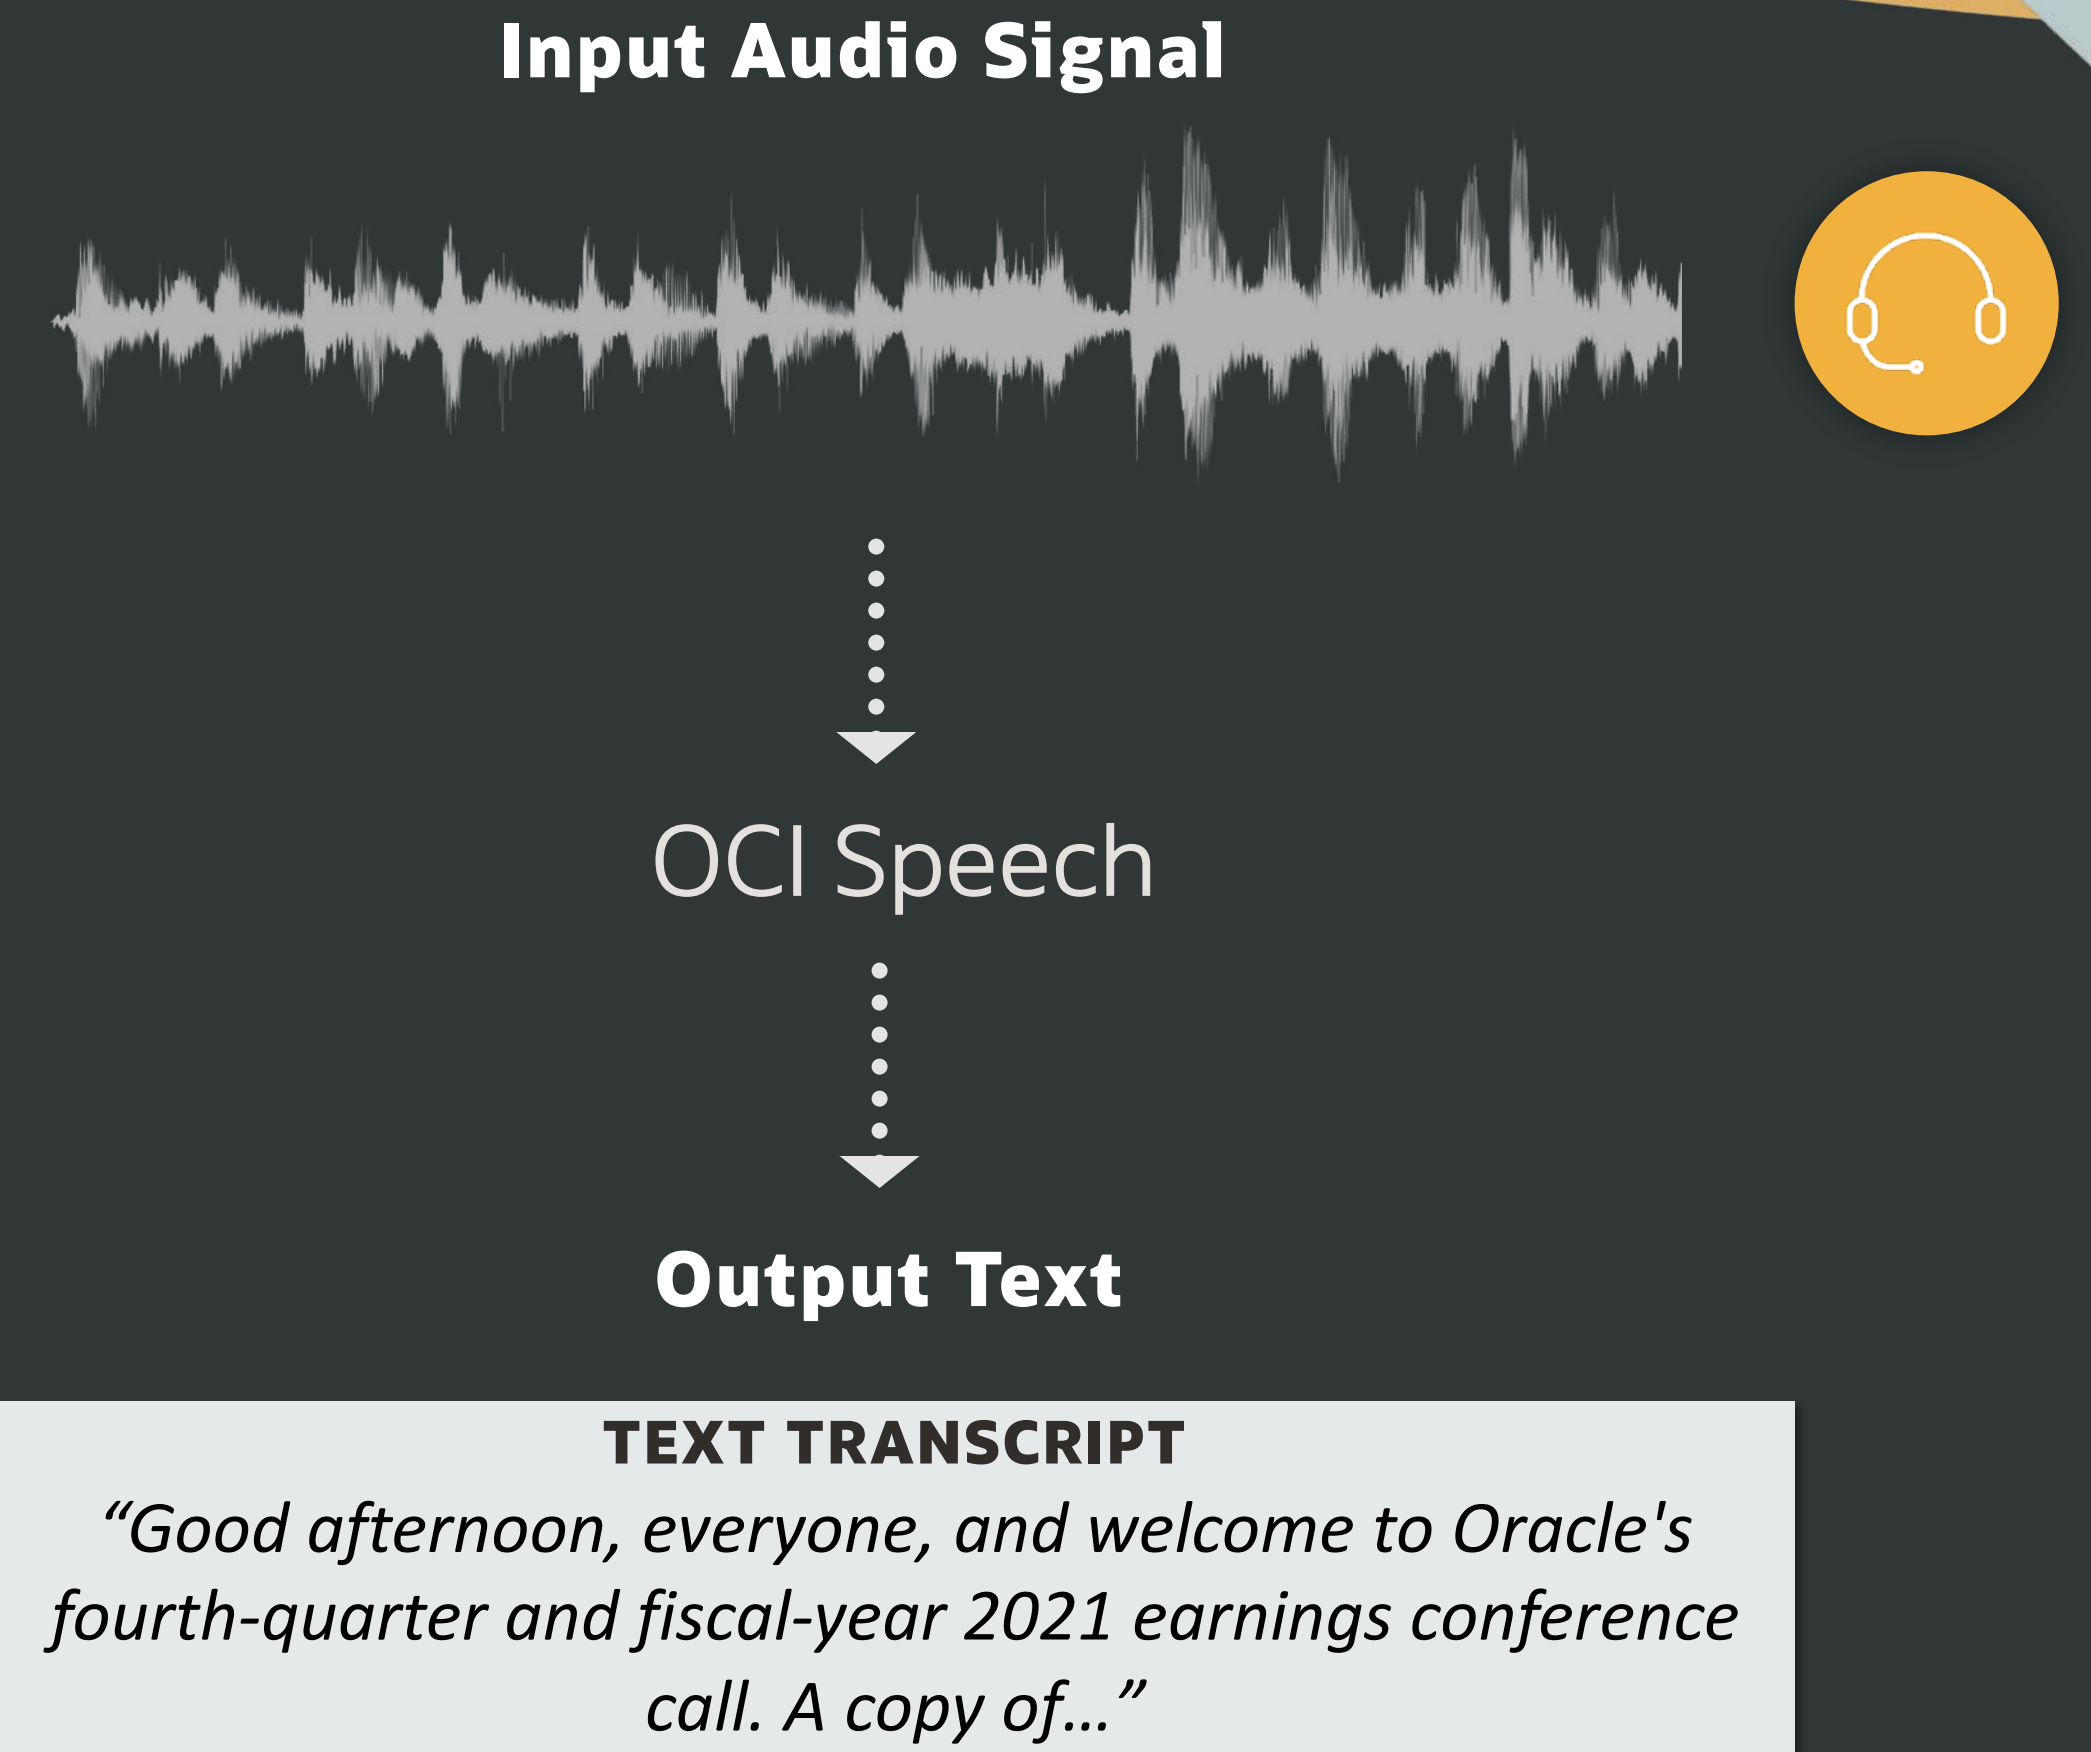
OCI Speech

Transcribe spoken audio with multilanguage support

Create searchable, indexed data by transcribing audio files

Filter profanities and leverage confidence scores

Combine with Language or other AI services to analyze sentiment, translate to other languages, and extract key information





OCI Vision

Understand images and detect objects

Automate workflows and actions based on image analysis

Gain insights from visual data by automatically extracting text and objects

Enrich image-based files with metadata, including objects and colors for better indexing and retrieval in a digital asset management system



OCI Document Understanding

Recognize and extract text, tables, and key data

Extract text from scanned documents, mobile uploads, PDFs, and more

Preserve table formats and key value relationships recognized in documents

Easily automate tedious business processes based on data output in JSON format

Pretrained using enterprise focused data, providing accurate results for enterprise focused use cases

RECEIPT OF SALE

SHOP'S

123 Main St. Any
Tel: 123-

01/02/2023 9:15:00 AM

Transaction ID: 12345678
Type: CREDIT

PURCHASE

Number: XXXXXXXXXXXX1234
Entry Mode: SWIPED
Card Type: VISA

Sub Total: \$8.12
Tax: \$1.12
Tip: \$1.00

Total: \$10.24

Thank you for supporting local businesses!

THANK YOU

JANE DOE
123 MAIN STREET
ANYTOWN, USA 12345

Date 1/1/2023

Pay to the Order of My Apartment Building \$ 1000.00

One thousand and 00/100 Dollars

Memo January Rent

0000000000 0000000000000000 1234

Jane Doe

	Qtr1-13	Qtr2-13	Qtr3-13	Qtr4-13	Total
40000-Revenue	312,993,815	646,566,977	987,076,779	1,388,795,773	3,827,683,913
50000-Materials	30,602,338	61,618,281	93,956,571	127,608,454	570,464,986
60100-Wages	5,254,392	10,665,420	16,704,412	22,786,122	107,367,557
60200-Other Employee Benefits	1,350,031	2,756,745	4,165,028	5,553,025	15,824,829
60500-Travel Expenses	625,122	1,414,180	2,114,809	2,717,731	6,871,842
62000-Freight and Shipping	532,709	1,227,373	1,801,061	2,341,570	5,902,713
64000-Office and Communication	1,205,058	2,409,894	3,614,740	4,819,084	12,048,776
65000-Fees	1,460,054	2,917,672	4,375,301	5,830,490	14,583,517
68000-Depreciation and Amortisation Expense	1,344,096	2,734,982	4,158,649	5,619,781	13,857,508
63100-Other Purchases and Supplies	6,086,510	12,152,621	18,438,759	24,942,737	61,620,627
Total Operating Expenses	39,995,297	81,956,192	129,649,826	185,745,047	499,346,362
77001-Extraordinaries	2,196	12,847	8,662	4,721	28,426
Income Before Income Taxes	2,675,936	6,661,018	15,288,095	30,250,731	54,875,780
79001-Taxes	842,920	2,098,221	4,815,750	9,528,981	17,285,872
Net Income	1,833,016	4,562,797	10,472,345	20,721,750	37,589,902



Streamlining the enterprise AI experience

OCI Generative AI coming in Fusion Applications 24A

Fusion Cloud HCM

- Candidate Experience Summary
- Compensation: Rewards and Recognition
- Survey Questions
- Goal Creation
- Performance Feedback & Development Tips
- Performance Review Summary
- About Me for Connections
- Job Category Landing Pages
- Job Match Explanations
- Feedback Assistance
- Career Summaries for Oracle Grow
- Creating SMART Team Goals

Fusion Cloud CX

- Generate a Knowledge Article from a Source Doc
- Generate a Knowledge Article from an SR
- Assisted Content Creation for Interviews
- Chat Agent Assistance
- Chat Summarization
- Issue/SR Summarization
- SLO and Campaign Text Generation

Fusion Cloud SCM

- Generate Item Descriptions and Attributes
- Sourcing; New Supplier Recommendations for Negotiation Invitation

ORACLE

OCI - Modern Data Platform




Enterprises Transitioning to Democratize Data Across Many Stakeholders

Customer Personas

Visualization



Data Architect



Notebook



ELTL Tools



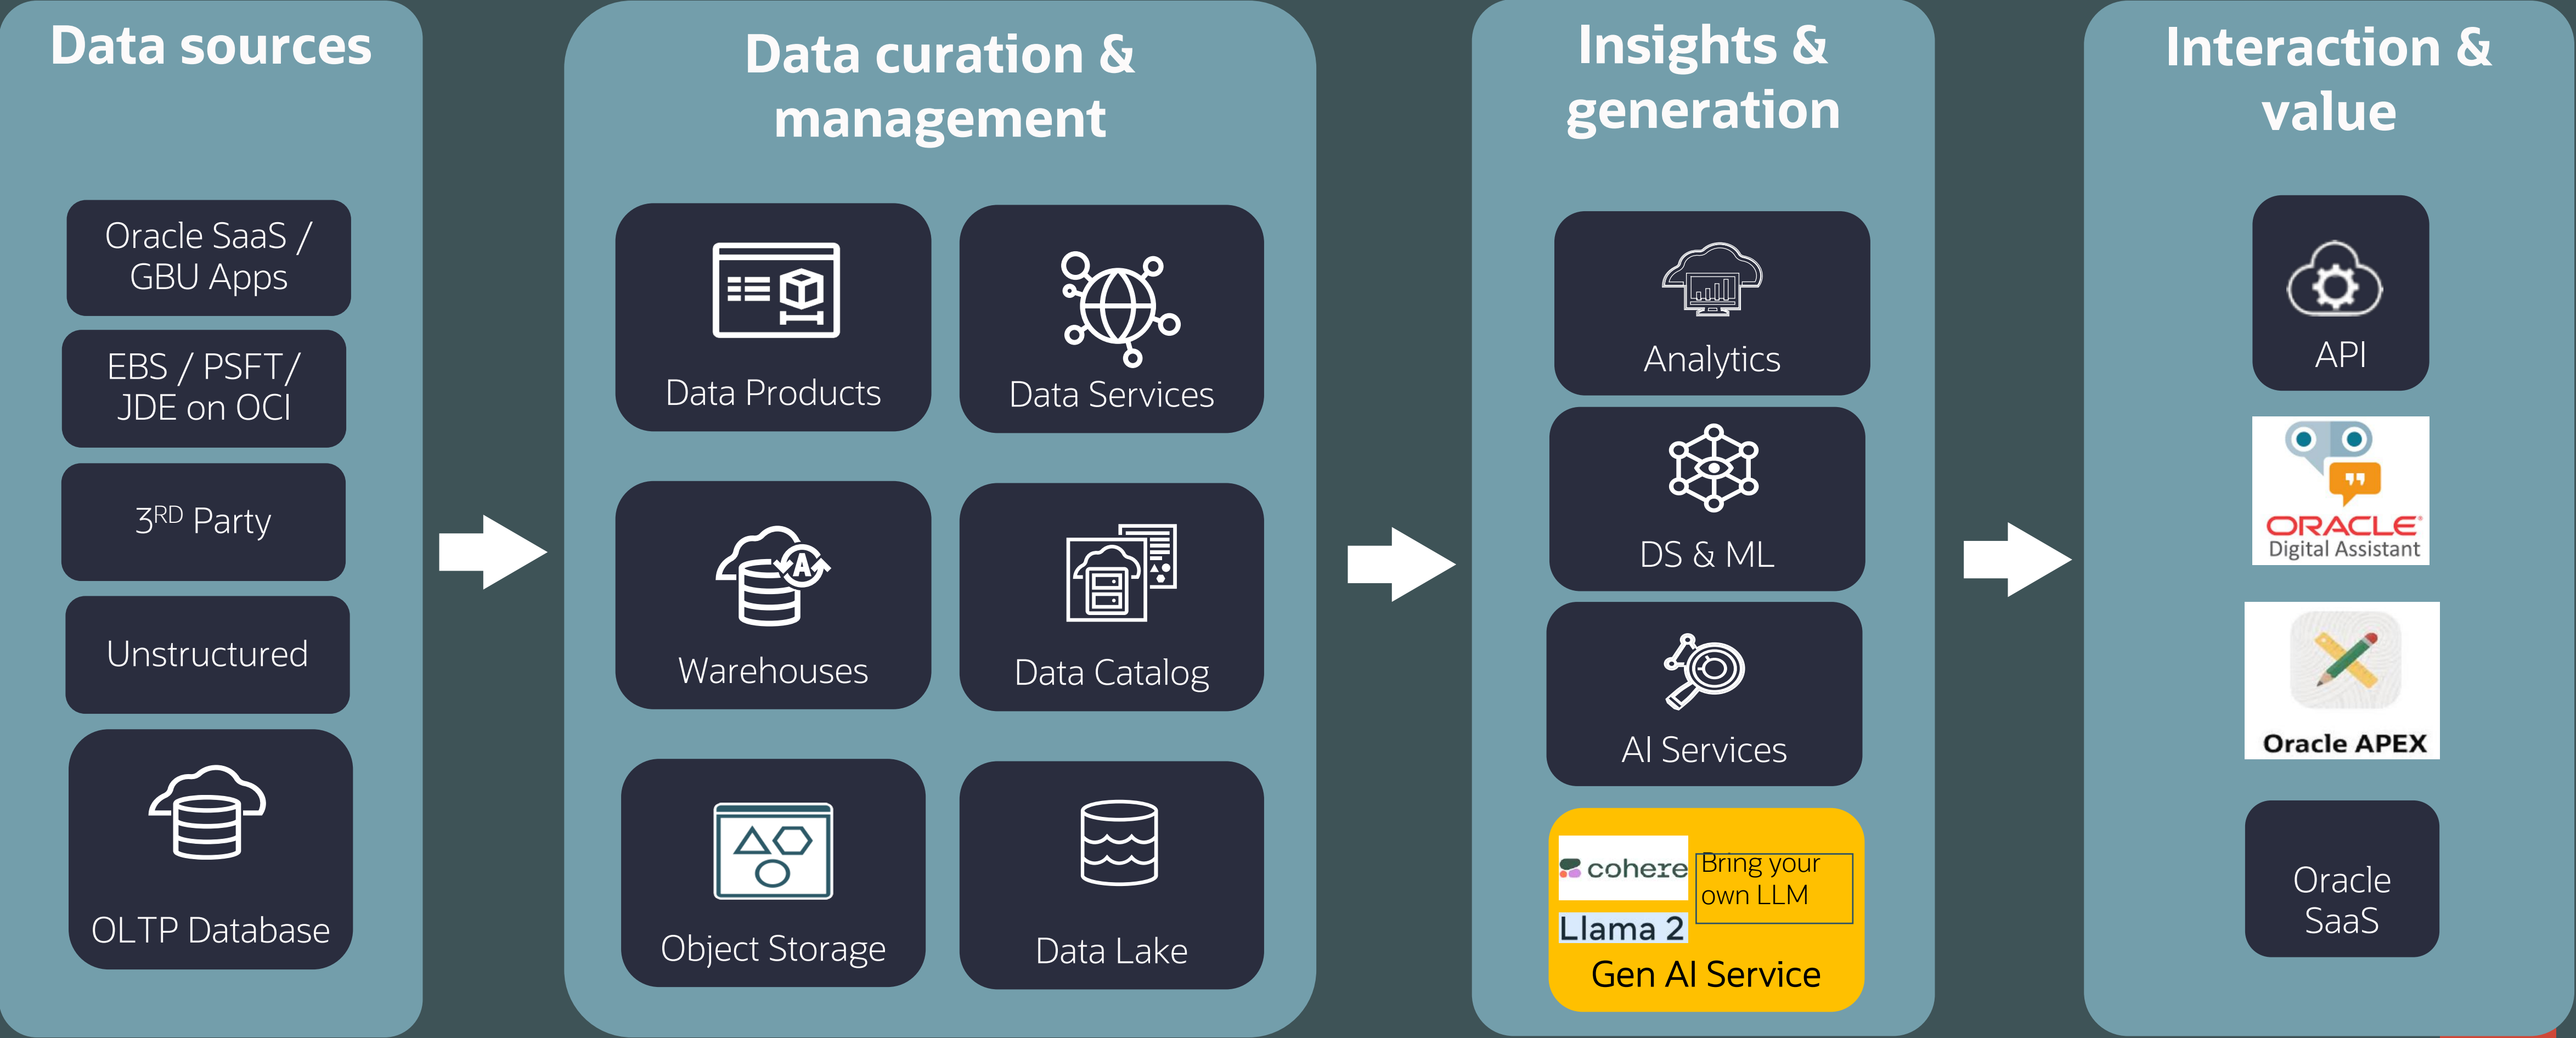
Low Code AppBuilder



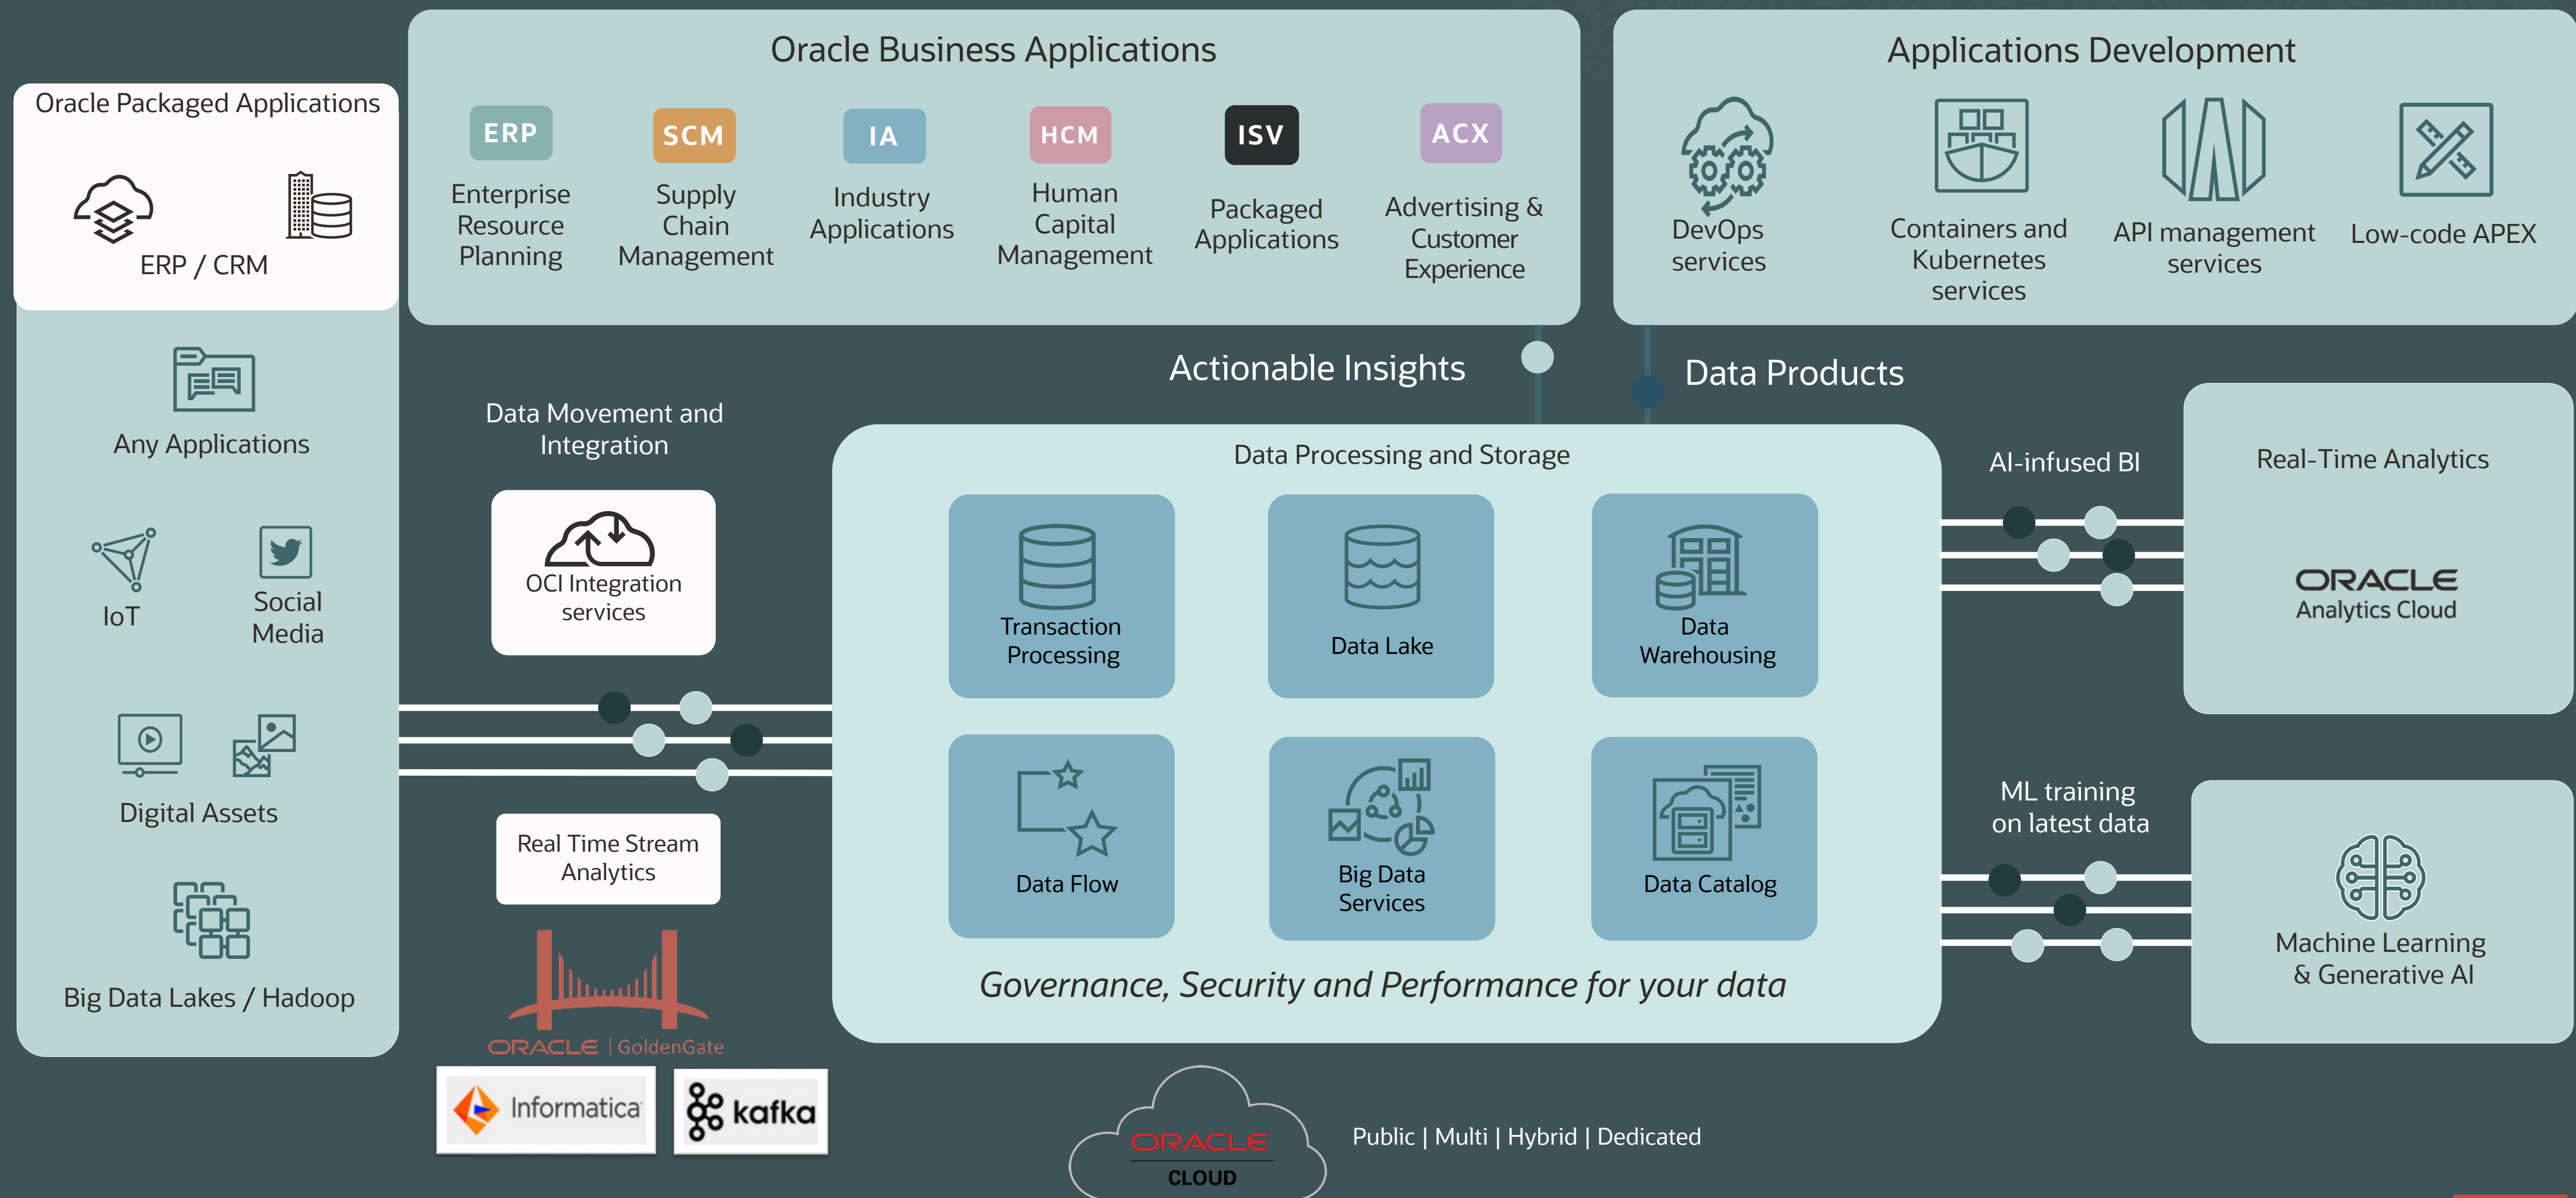
- Cloud helps bring silos of data together
- Make compliance and compartmentalization of data easy for various personas
- Democratize access to data securely using tools like Slack, Teams etc...

Generative AI - Needs a Modern Data Platform

Gen-AI and Data Platform fuels Growth. Data must be centralized, governed, processed and managed at scale.



Oracle Modern Data Platform



OCI Modern Data Platform

Data Sources and discovery



Oracle Fusion Modules



SaaS apps - ADP, SFDC, WD, and many others



On-premises Application



Third Party Applications



Unstructured Data



Any Other Database

Data Integration

Data Flow (Spark)

Data Integration (Spark)

Golden Gate (Realtime)

Streaming Service (Near-Realtime)

Oracle Integration Cloud (App Connector)

Oracle Data Integrator (ETL Tool)

Data Store



Data Safe



Autonomous Data Warehouse or ExaCS



Data Flow



Big Data Service



Data Catalog



Object Storage



Raw Layer



Enriched Layer



Archive Layer

Analytics



Oracle Analytics Cloud



Data Science



AI and Gen-AI Services



Machine Learning

Measure, Act



Enterprise reporting



Transaction Investigation



Fraud tracking



Dashboards and reports



Machine learning models



Web-based apps



Other consumers

Low and Predictable Cost
Granular control to match workload needs

Support Architecture Pattern
(Data Fabric, Data Mesh, Lakehouse)

Comprehensive Data Protection
Security assessment, sensitive data masking

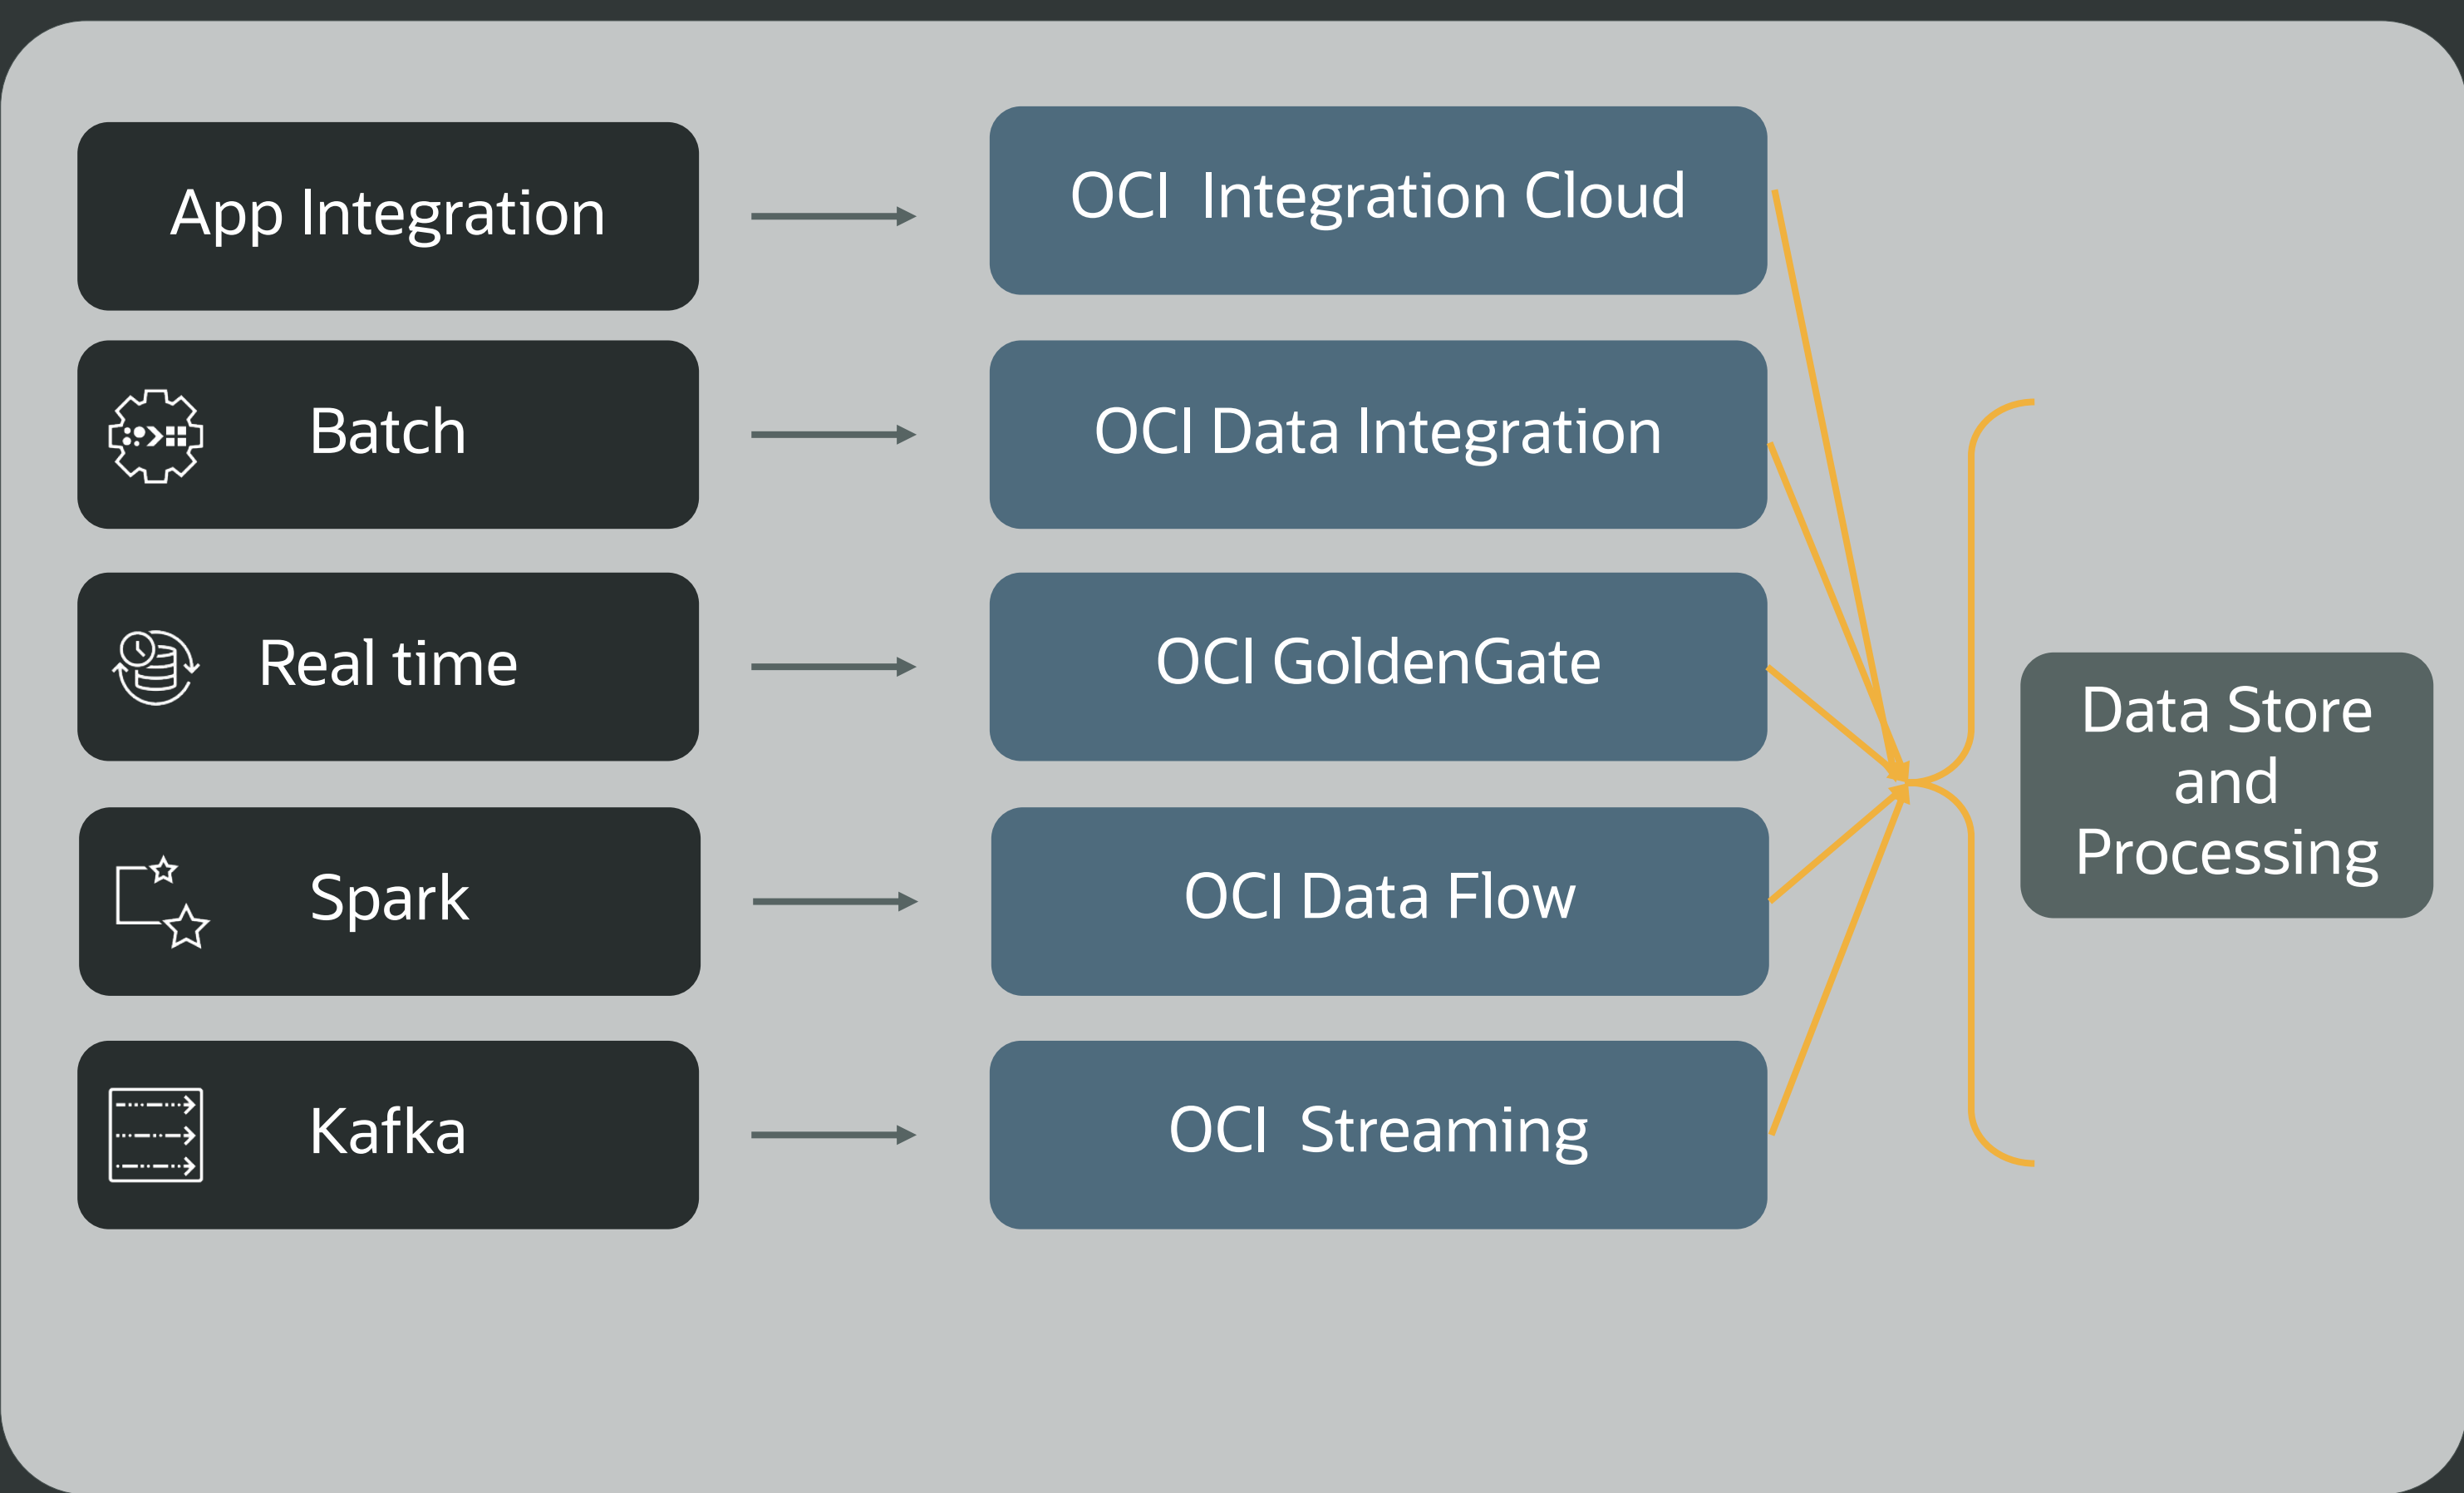
Modern Bigdata platform
Along with Serverless Spark data integration

Analytics, AI / Data Science Services
with prebuilt machine learning models



Data Integration comprehensive portfolio

Oracle offers multiple ways to load your data



OCI Integration Cloud
Apps and API integration

OCI Data Integration
Batch data integration

OCI GoldenGate
Real-time data integration and streaming

OCI Data Flow
Managed Apache Spark service

OCI Streaming
Apache Kafka-compatible event streaming Service



Oracle Modern Data Platform

We meet you where you are!

OCI Data Flow with Apache Spark

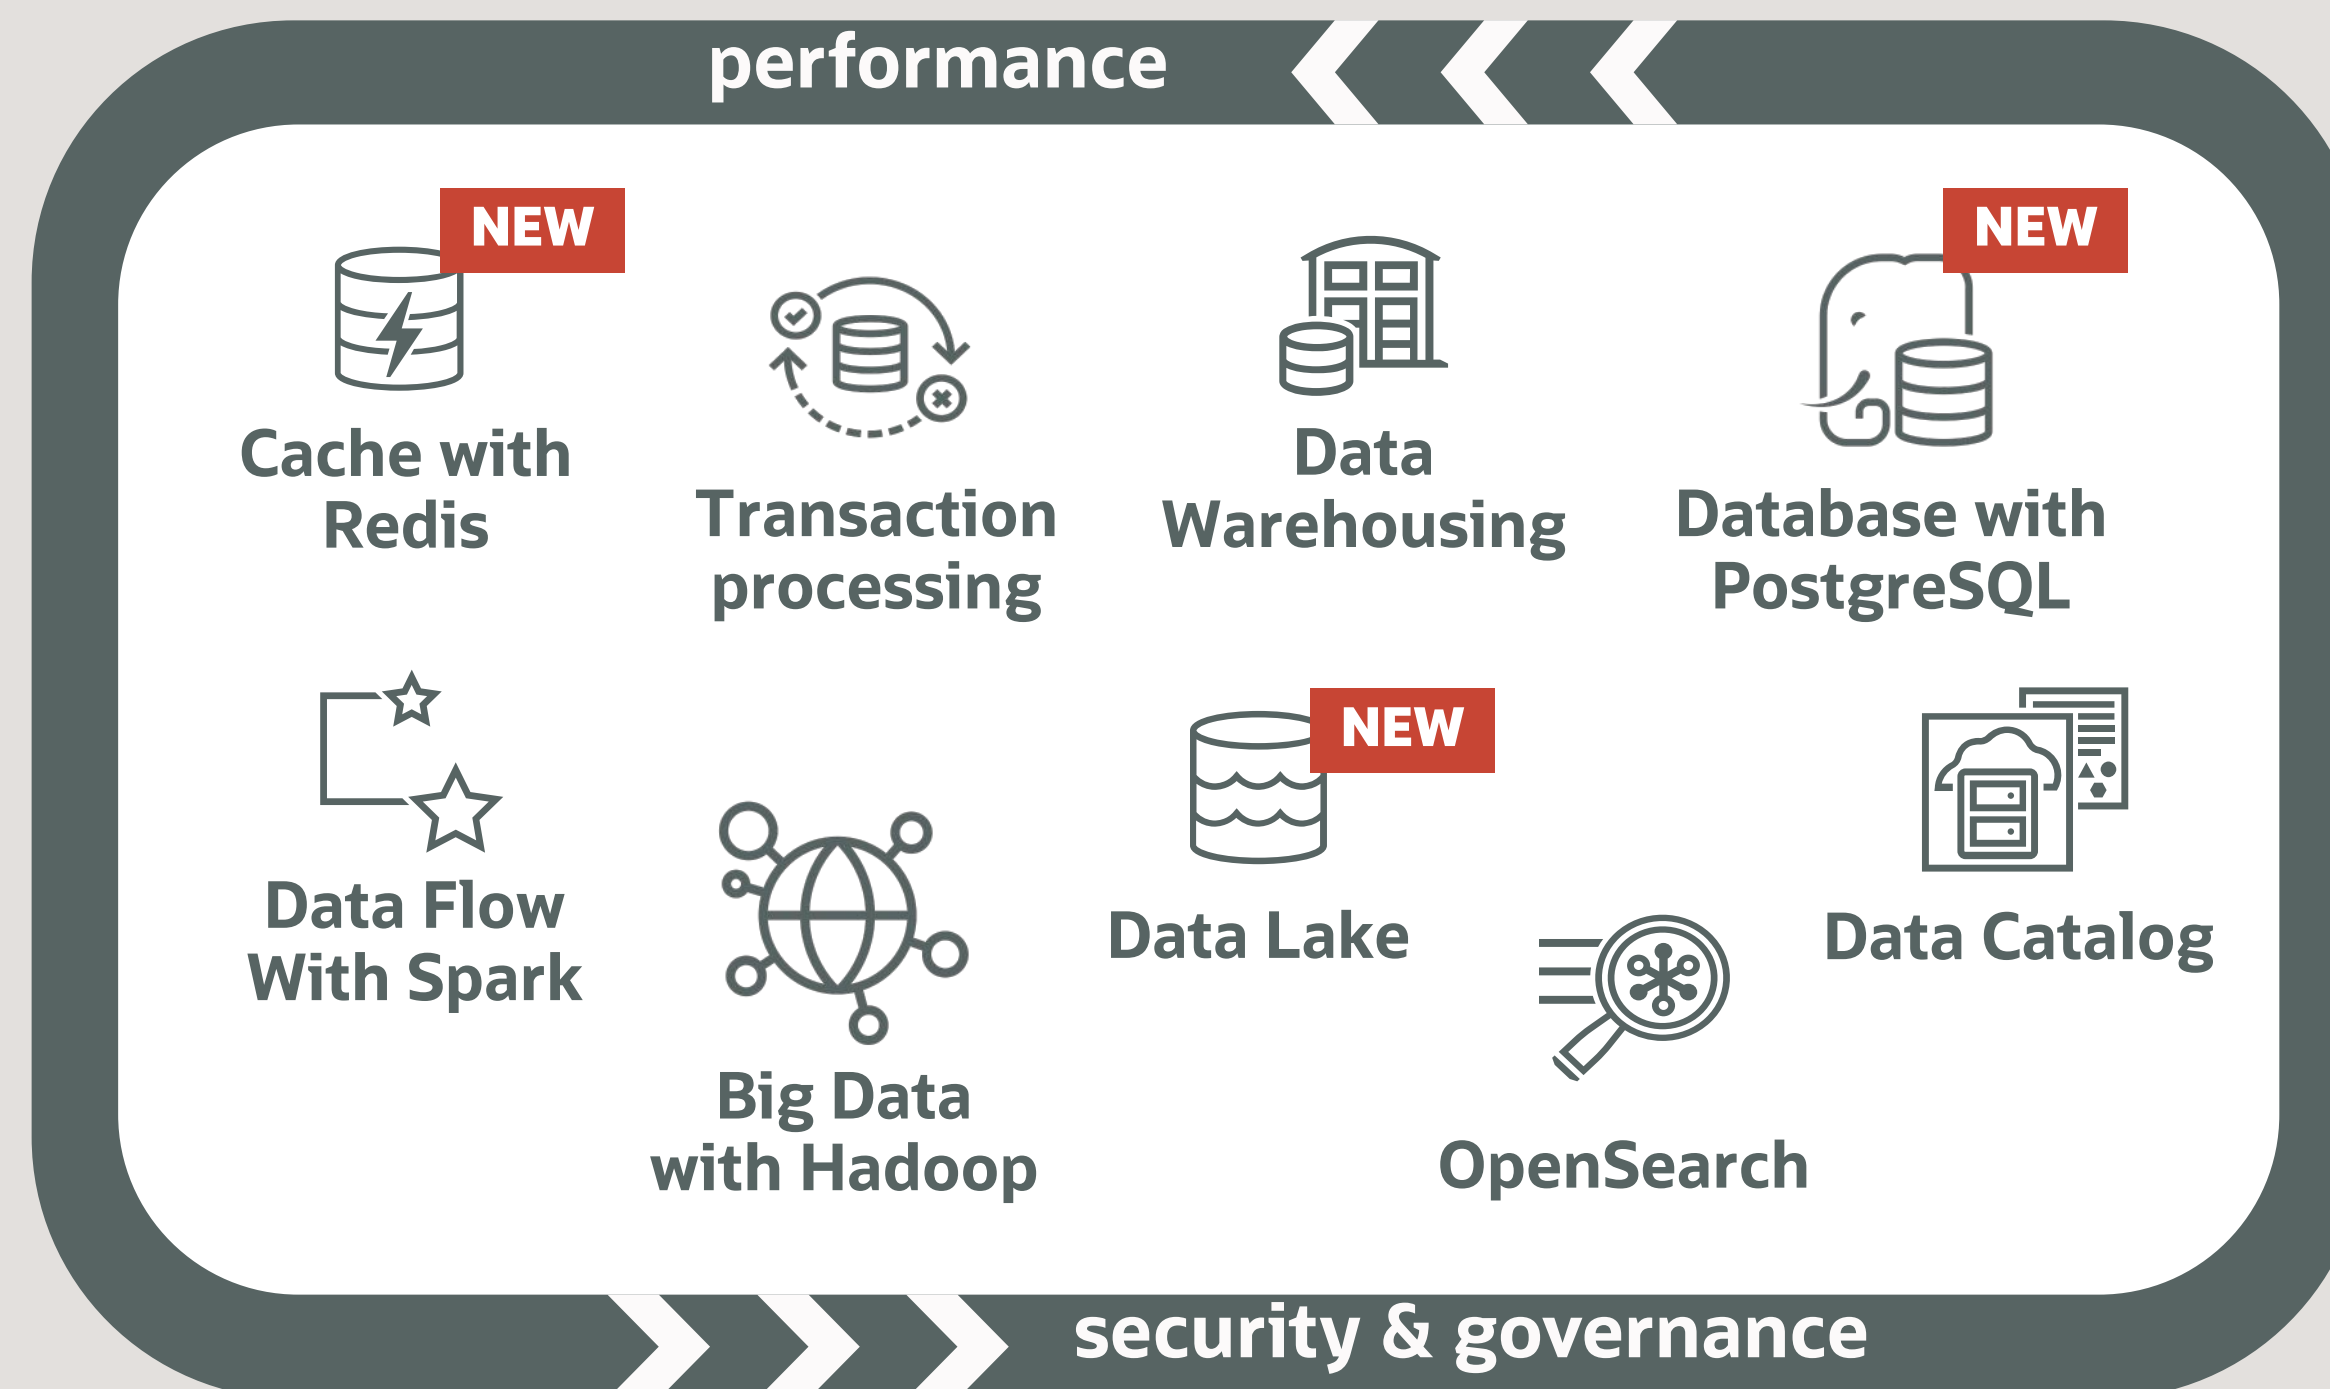
Fully managed Apache Spark service for large datasets

Big Data with Apache Hadoop

Oracle Distribution Including Apache Hadoop, Apache Ambari, Apache HBase, Apache Hive, Apache Spark, more

Data Catalog

Metadata management to help discover data and support data governance on OCI



OpenSearch


Fully managed open source search service for you application search and log analysis

Redis

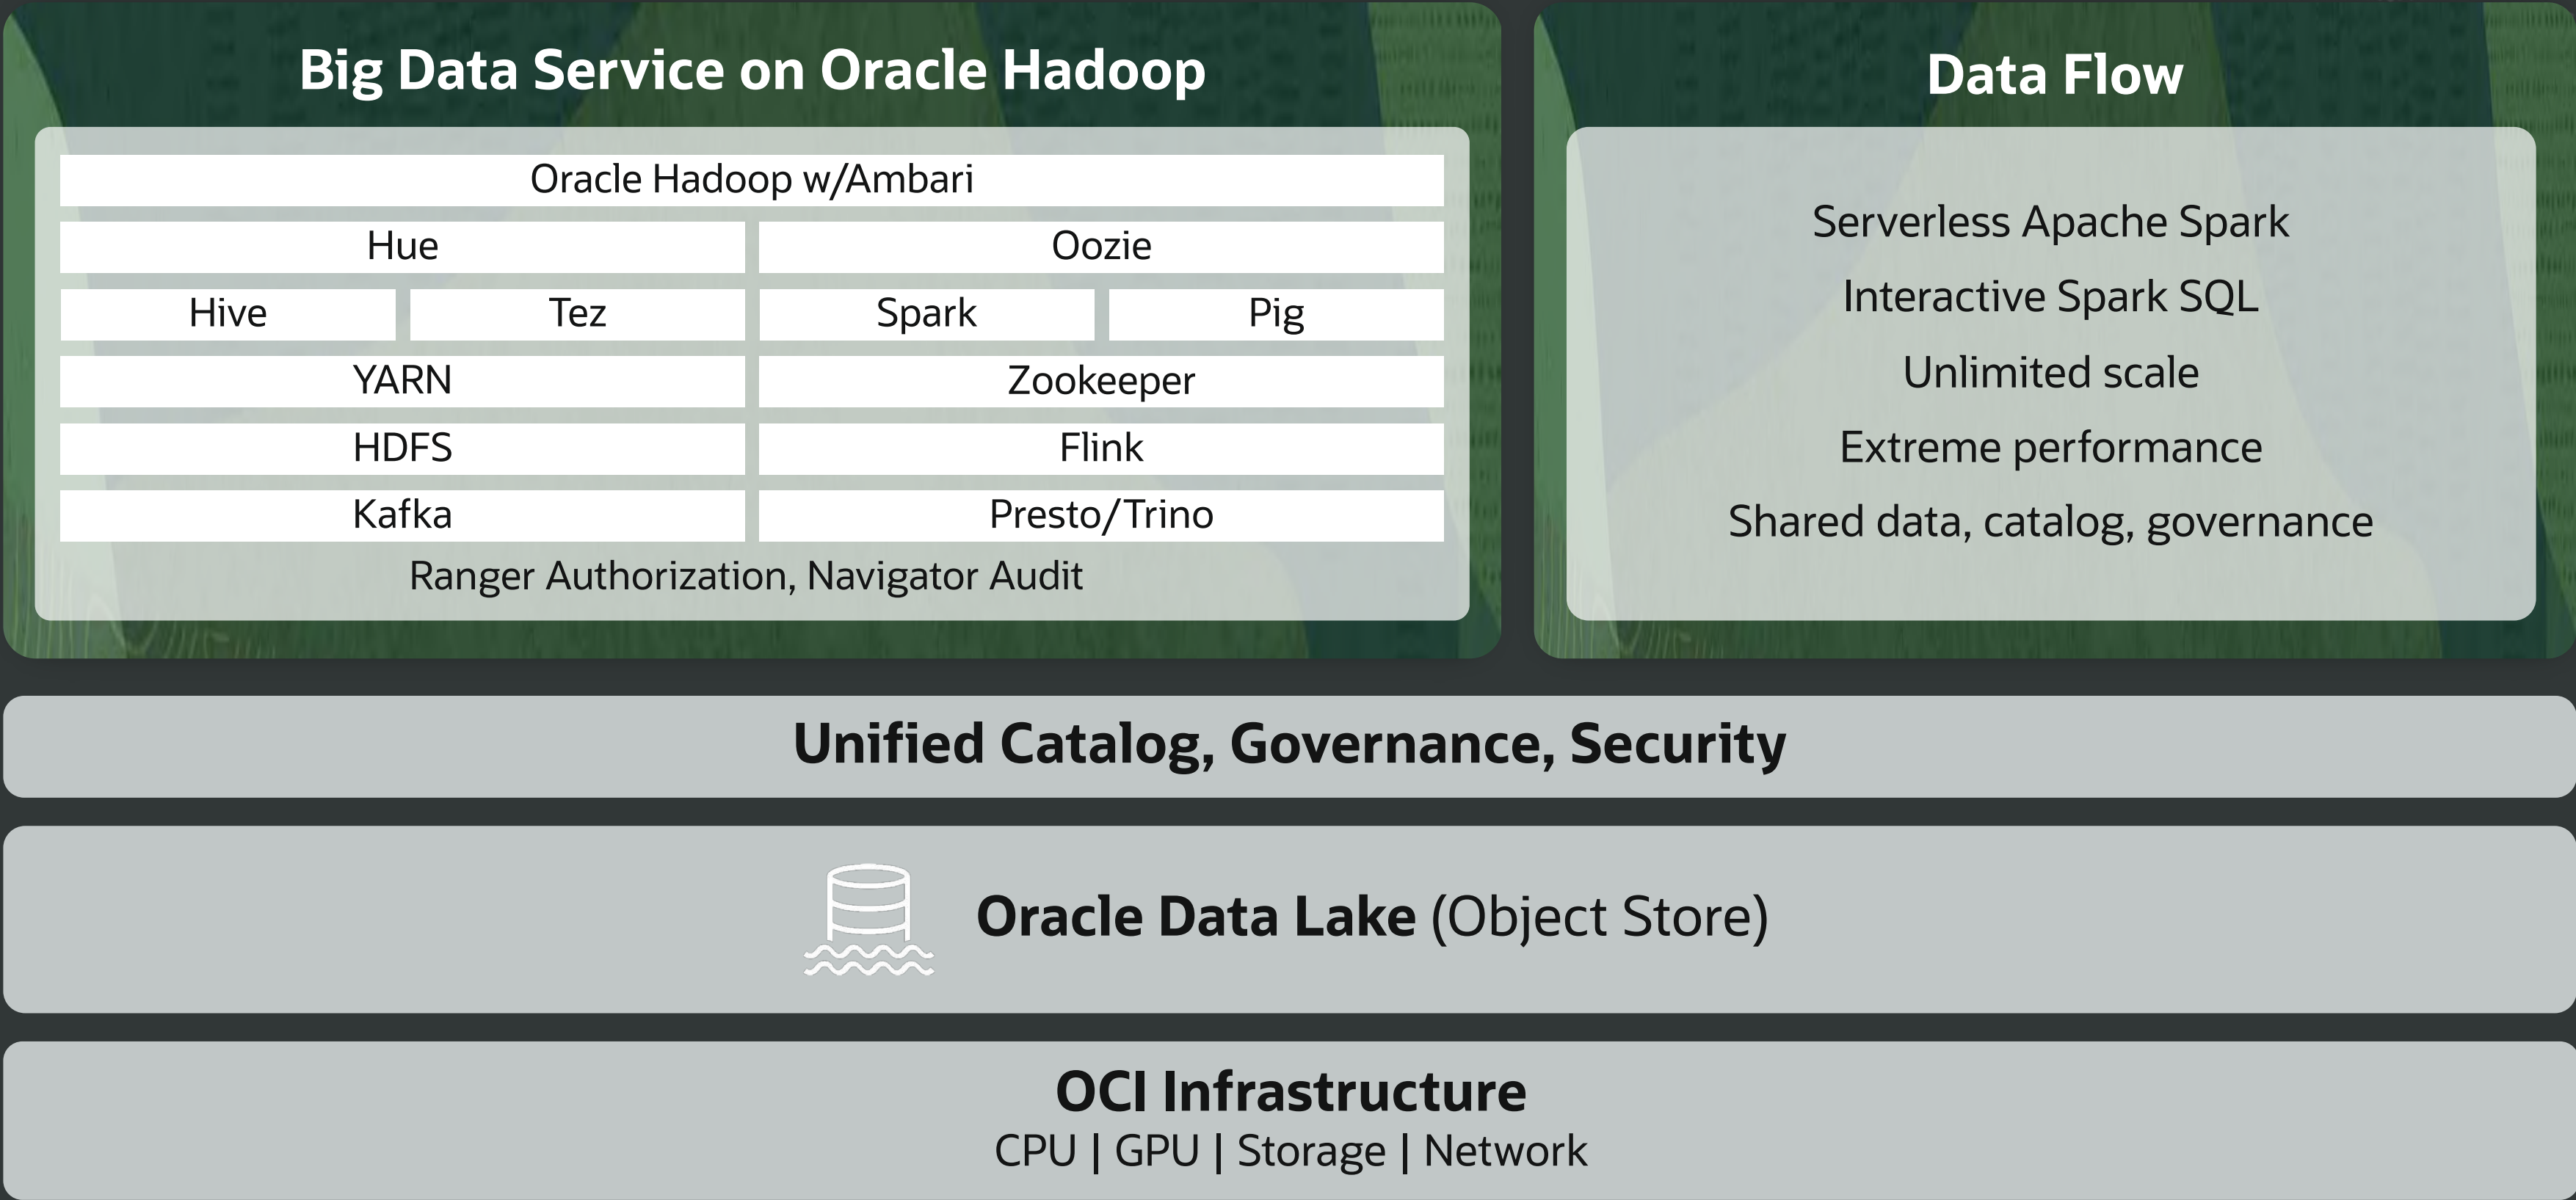
Managed in-memory caching solution to improve your application performance and customer experience

PostgreSQL

Managed PostgreSQL database with optimized performance, dynamically scale storage



OCI Big Data Service and OCI Data Flow managed open-source services on shared storage and catalog



Data Flow simplifies the Spark ecosystem


Fully managed Apache Spark ETL, SQL and Streaming

Data Engineering Jobs



Serverless Spark for
Big Data and ML at any scale
with no administration.

Interactive Spark SQL



Petabyte-scale data discovery, preparation
and interactive query on OCI Object Storage
from BI tools using ODBC or JDBC.

Spark Streaming



Serverless Spark Streaming with
deep OCI integration.

Data Engineering with Data Flow

Spark on Demand

Ephemeral Spark jobs on demand in about a minute.



Simple Operations

Built-in monitoring, alerting and diagnostic UIs.




Enterprise Security

RBAC and Impersonation for fine-grained data security.



Lowest Cost

No charge for the service, only pay for IaaS while it runs.



Interactive SQL with Data Flow SQL Endpoint

Fully-Managed Apache SparkSQL

Deploy SparkSQL endpoints in minutes with nothing to maintain



Discover your Data Lake

Query structured and unstructured data in object storage with SQL



Fast and Scalable

Fast in-memory performance
Scale to thousands of cores



Curate, Share and Secure

Publish data to Oracle Catalog

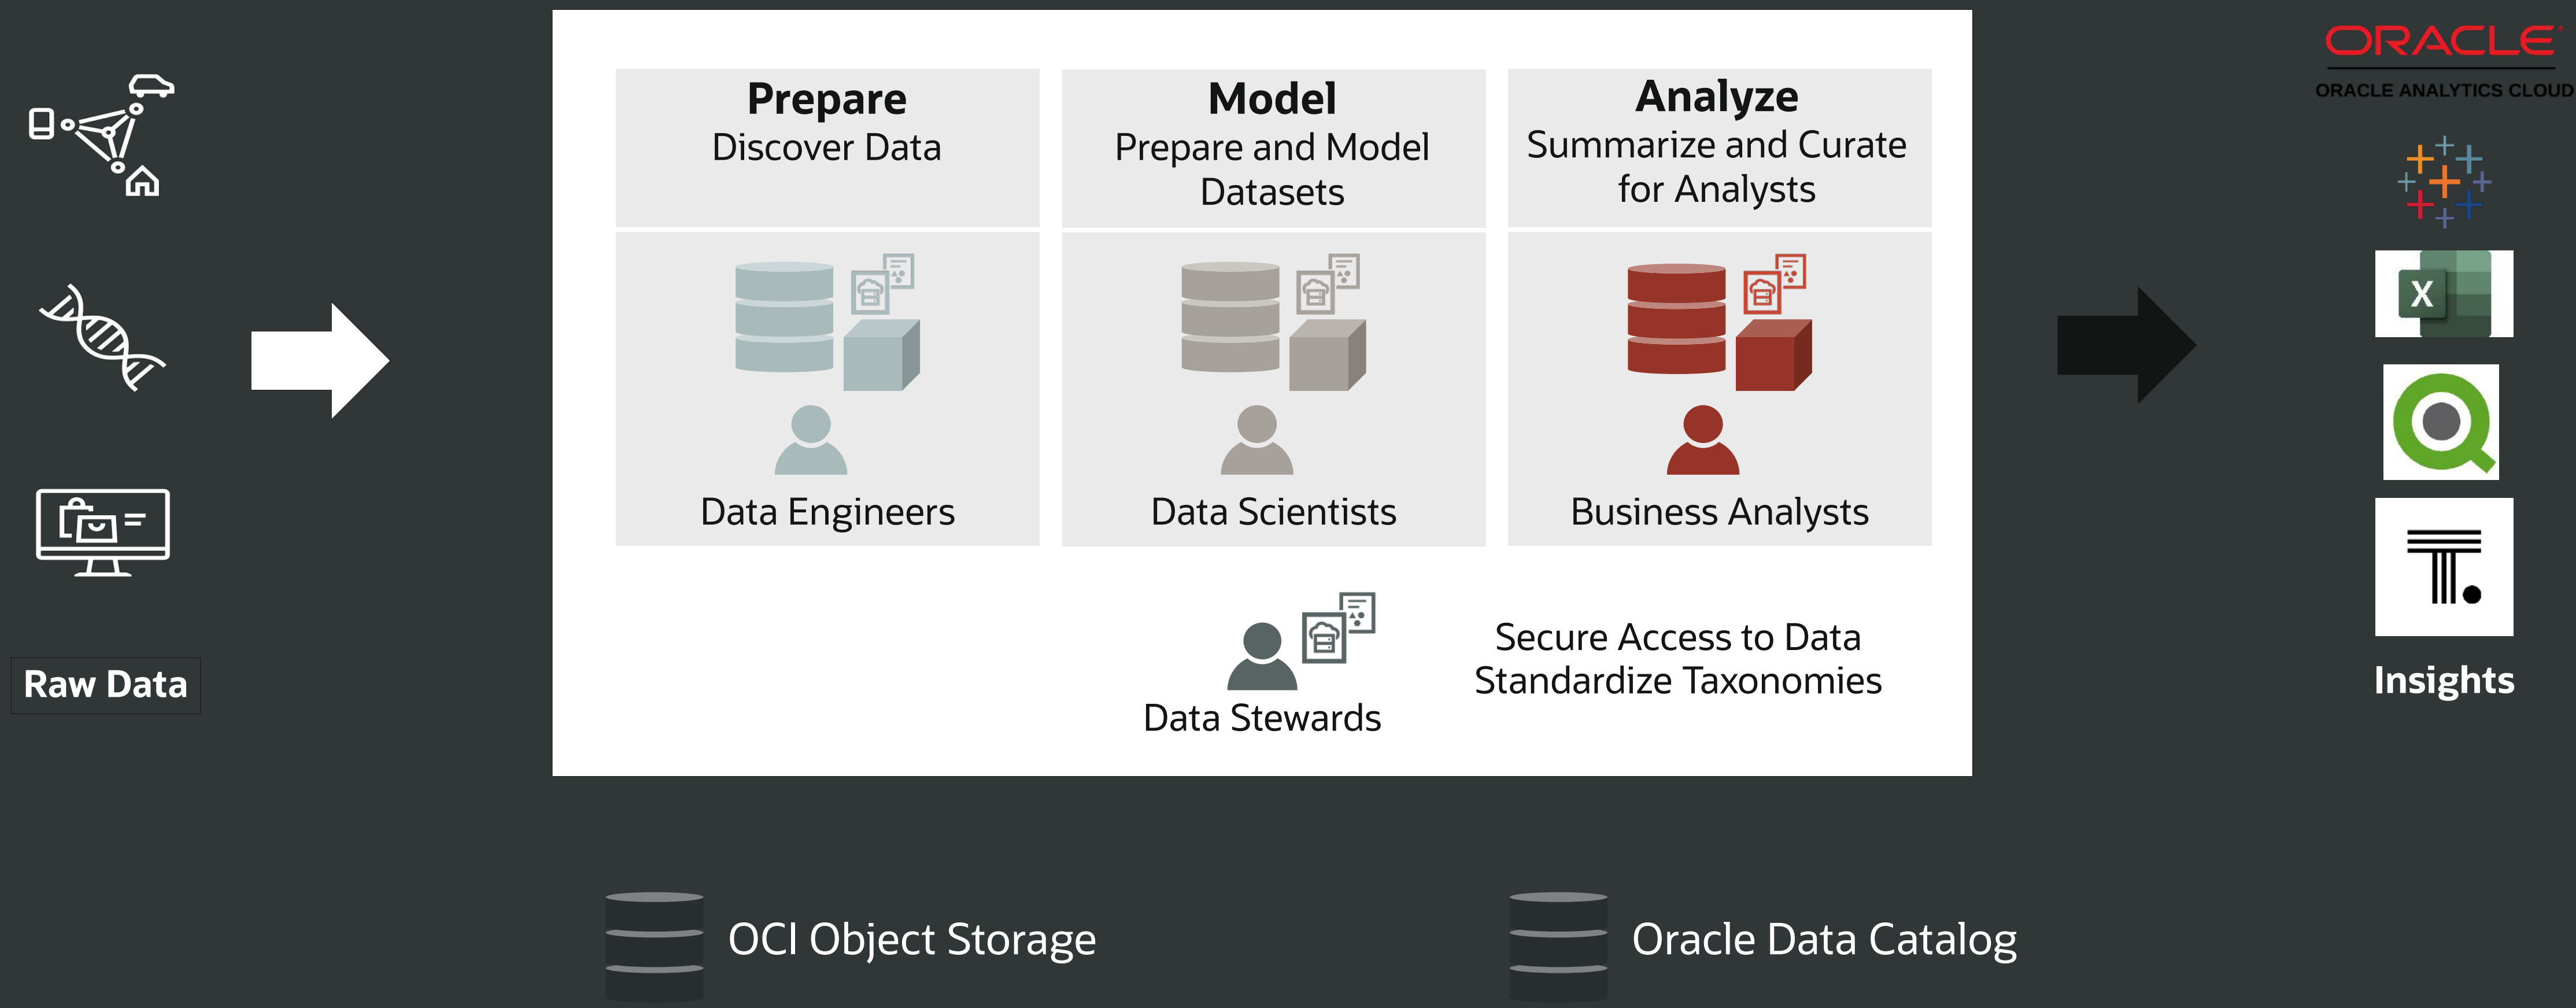


Integrated

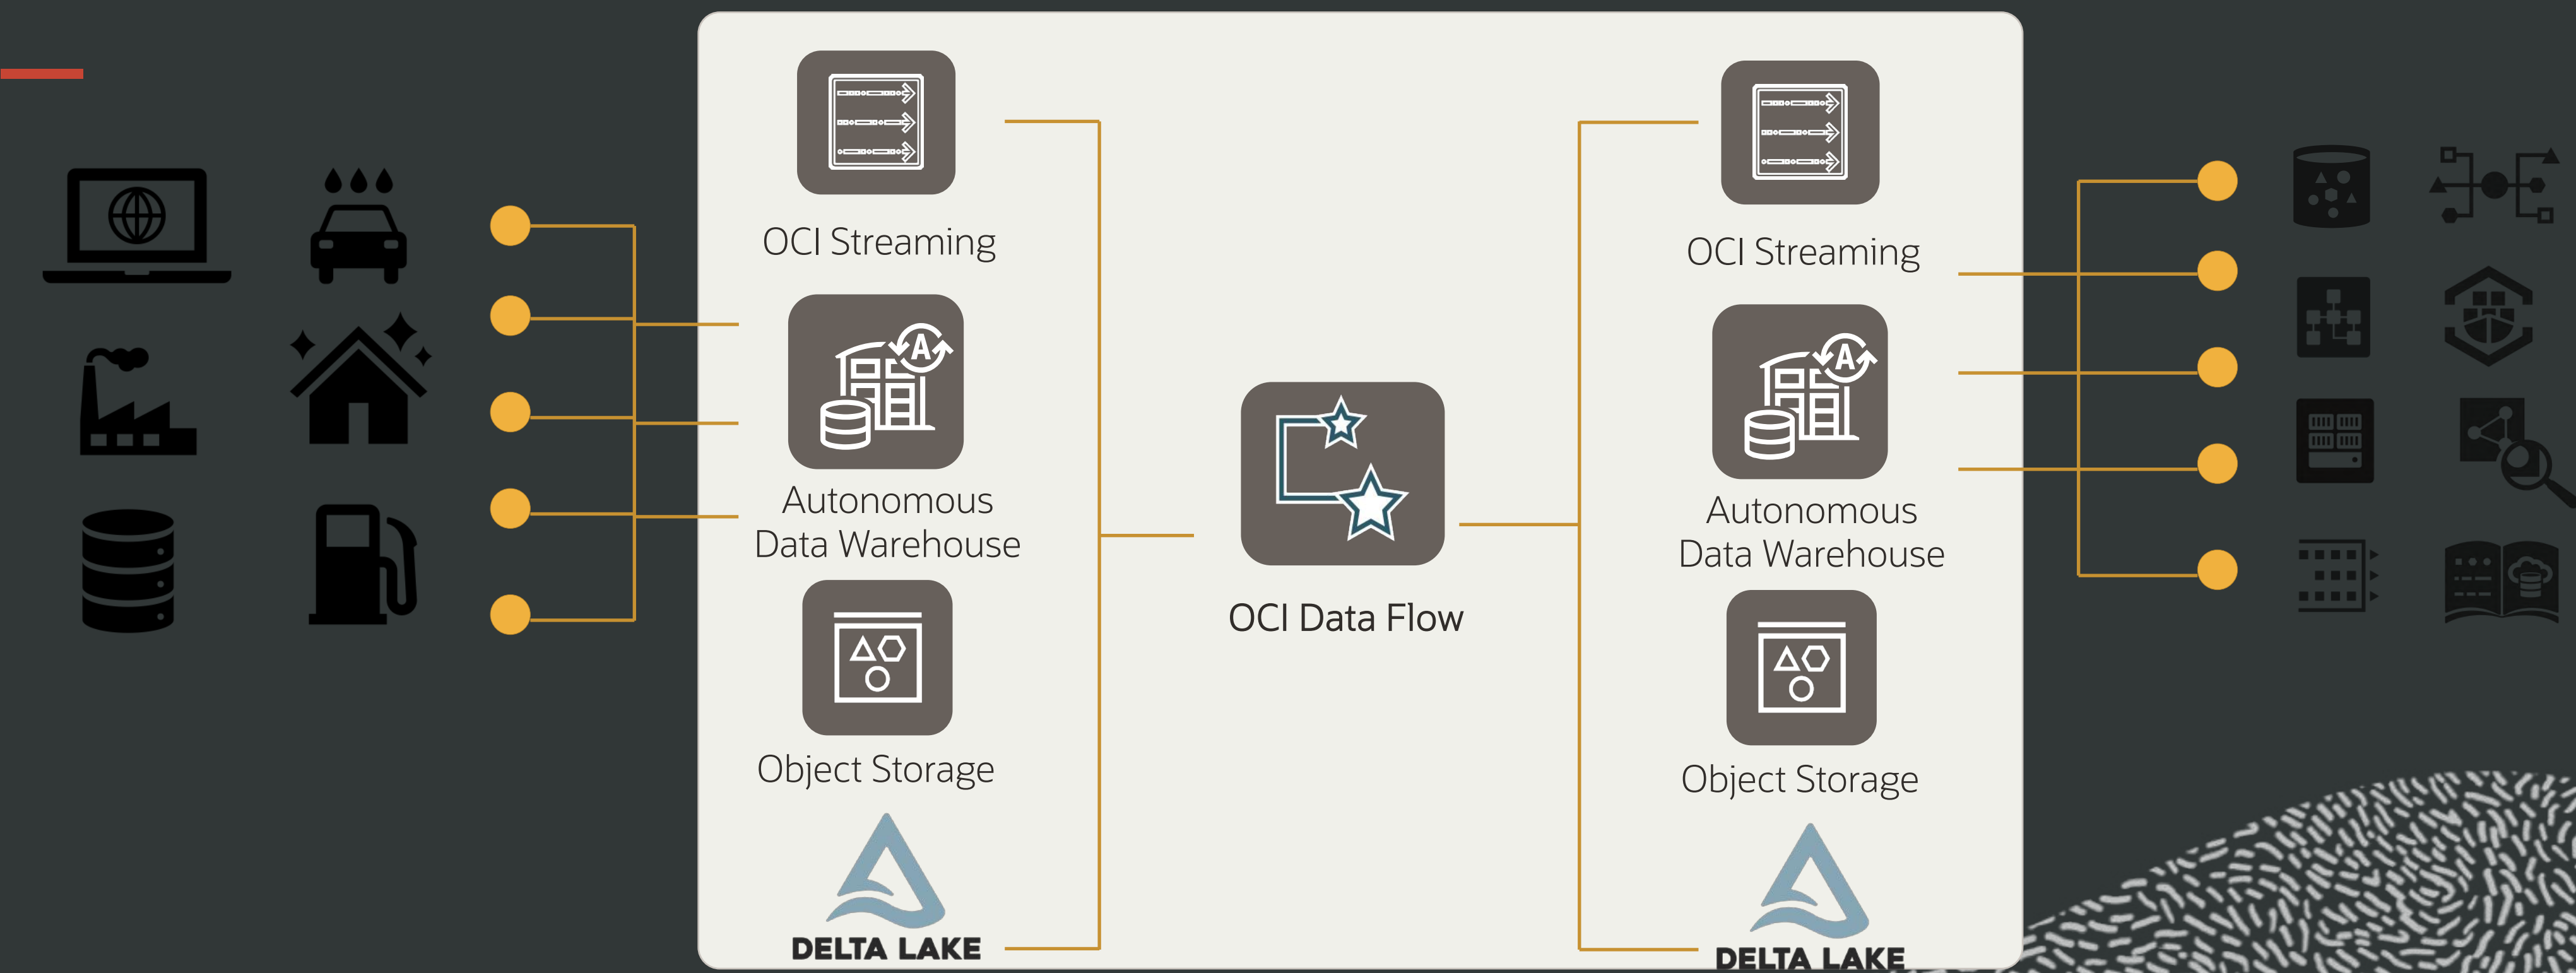
Query with Oracle Analytics Cloud
or bring your favorite BI Tool



Data Flow Spark SQL endpoints can be connected over JDBC/ ODBC for ad-hoc queries and analytics at petabyte scale



Data Flow now pre-includes Delta Lake libraries with Spark



						
ACID Transactions	Scalable Metadata	Time Travel	Unified Batch/Streaming	Schema Evolution / Enforcement	Audit History	DML Operations
Protect your data with serializability, the strongest level of isolation	Handle petabyte-scale tables with billions of partitions and files with ease	Access/revert to earlier versions of data for audits, rollbacks, or reproduce	Exactly once semantics ingestion to backfill to interactive queries	Prevent bad data from causing data corruption	Delta Lake log all change details providing a full audit trail	SQL, Scala/Java and Python APIs to merge, update and delete datasets




Introducing OCI Cache with Redis


OCI Cache with Redis is a fully managed distributed caching service that uses open-source **Redis 7.0.5**.

The service automates the complex and routine tasks associated with deploying and managing a distributed environment so you can focus on building great applications.

We offer easy cluster creation, automated HA, patching, security updates, and resizing.



Key features



Fully Managed

- Automated cluster creation
- OS-level patching
- Monitoring
- Failover
- SLA 99.95%



Unparalleled Simplicity

- Specify the amount of memory and number of nodes and we take care of everything else
- Choose exact amount of memory for your workload. No need to overprovision



Scalable

- Scale your cache on the fly with no downtime
- Scale from 2GB up to 500GB per node and 5 nodes per cluster to meet your caching needs



Powerful Insights

- Out of box metrics to monitor cluster health and performance
- Access CPU and Memory utilization
- See transmitted and received network bytes



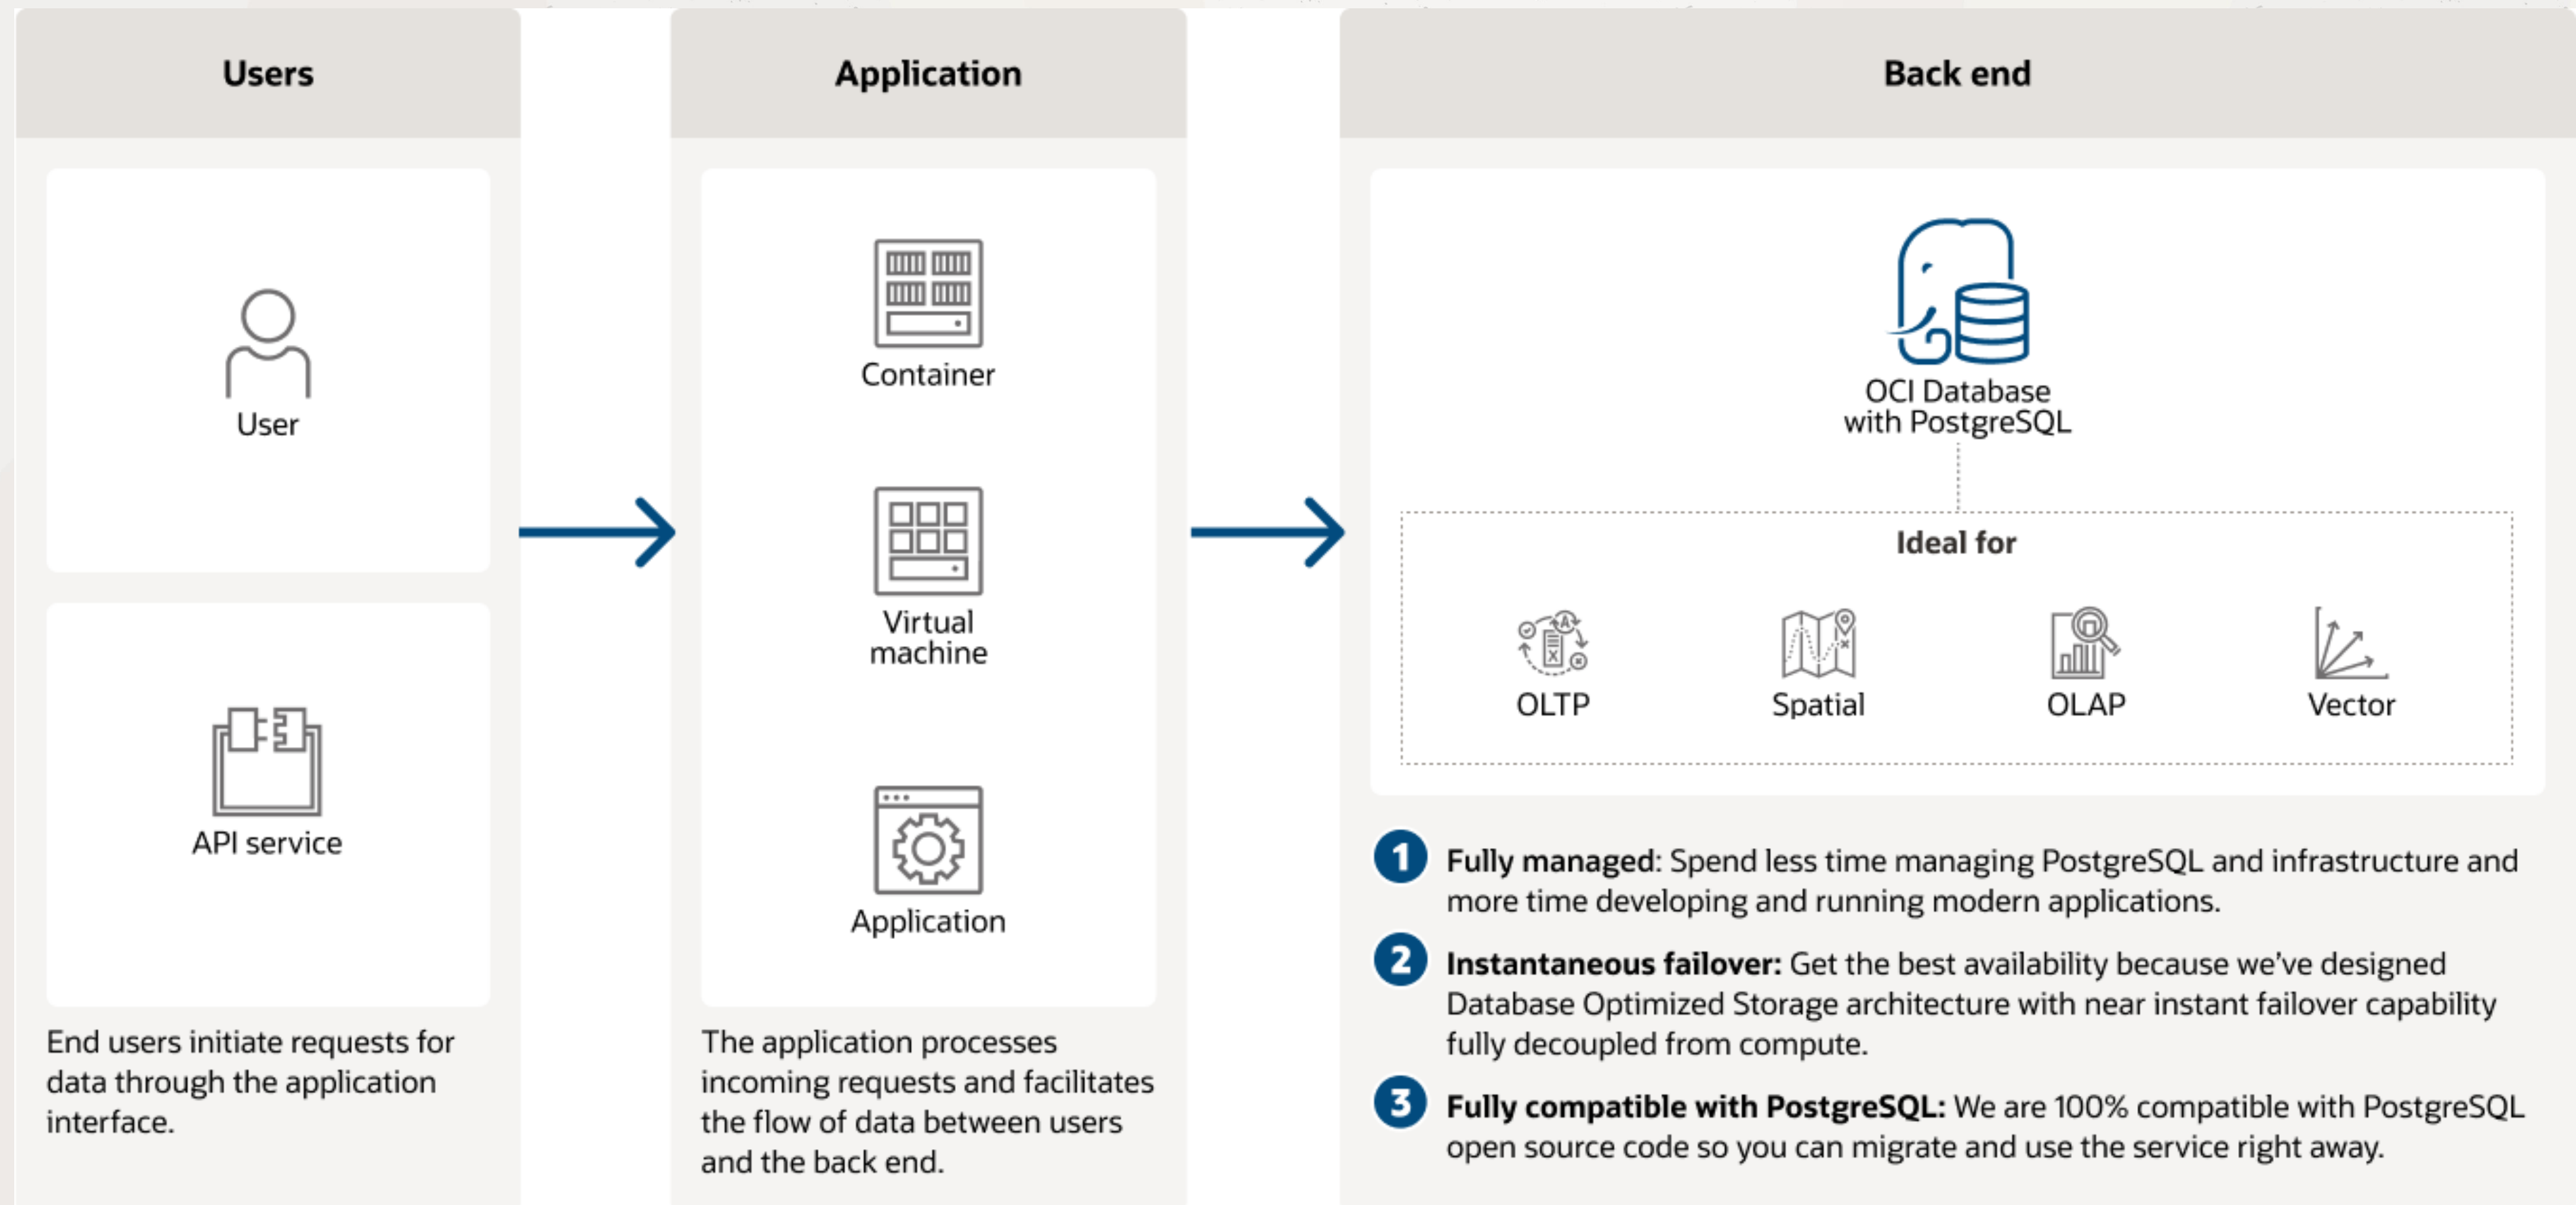
Automated HA

- Deploy 2 or more nodes and OCI automatically spreads them across multiple Availability or Fault Domains



OCI Database with PostgreSQL

A fully managed service



OCI Database with PostgreSQL

A fully managed service

60%

Less costly than Amazon Aurora with PostgreSQL¹

3X

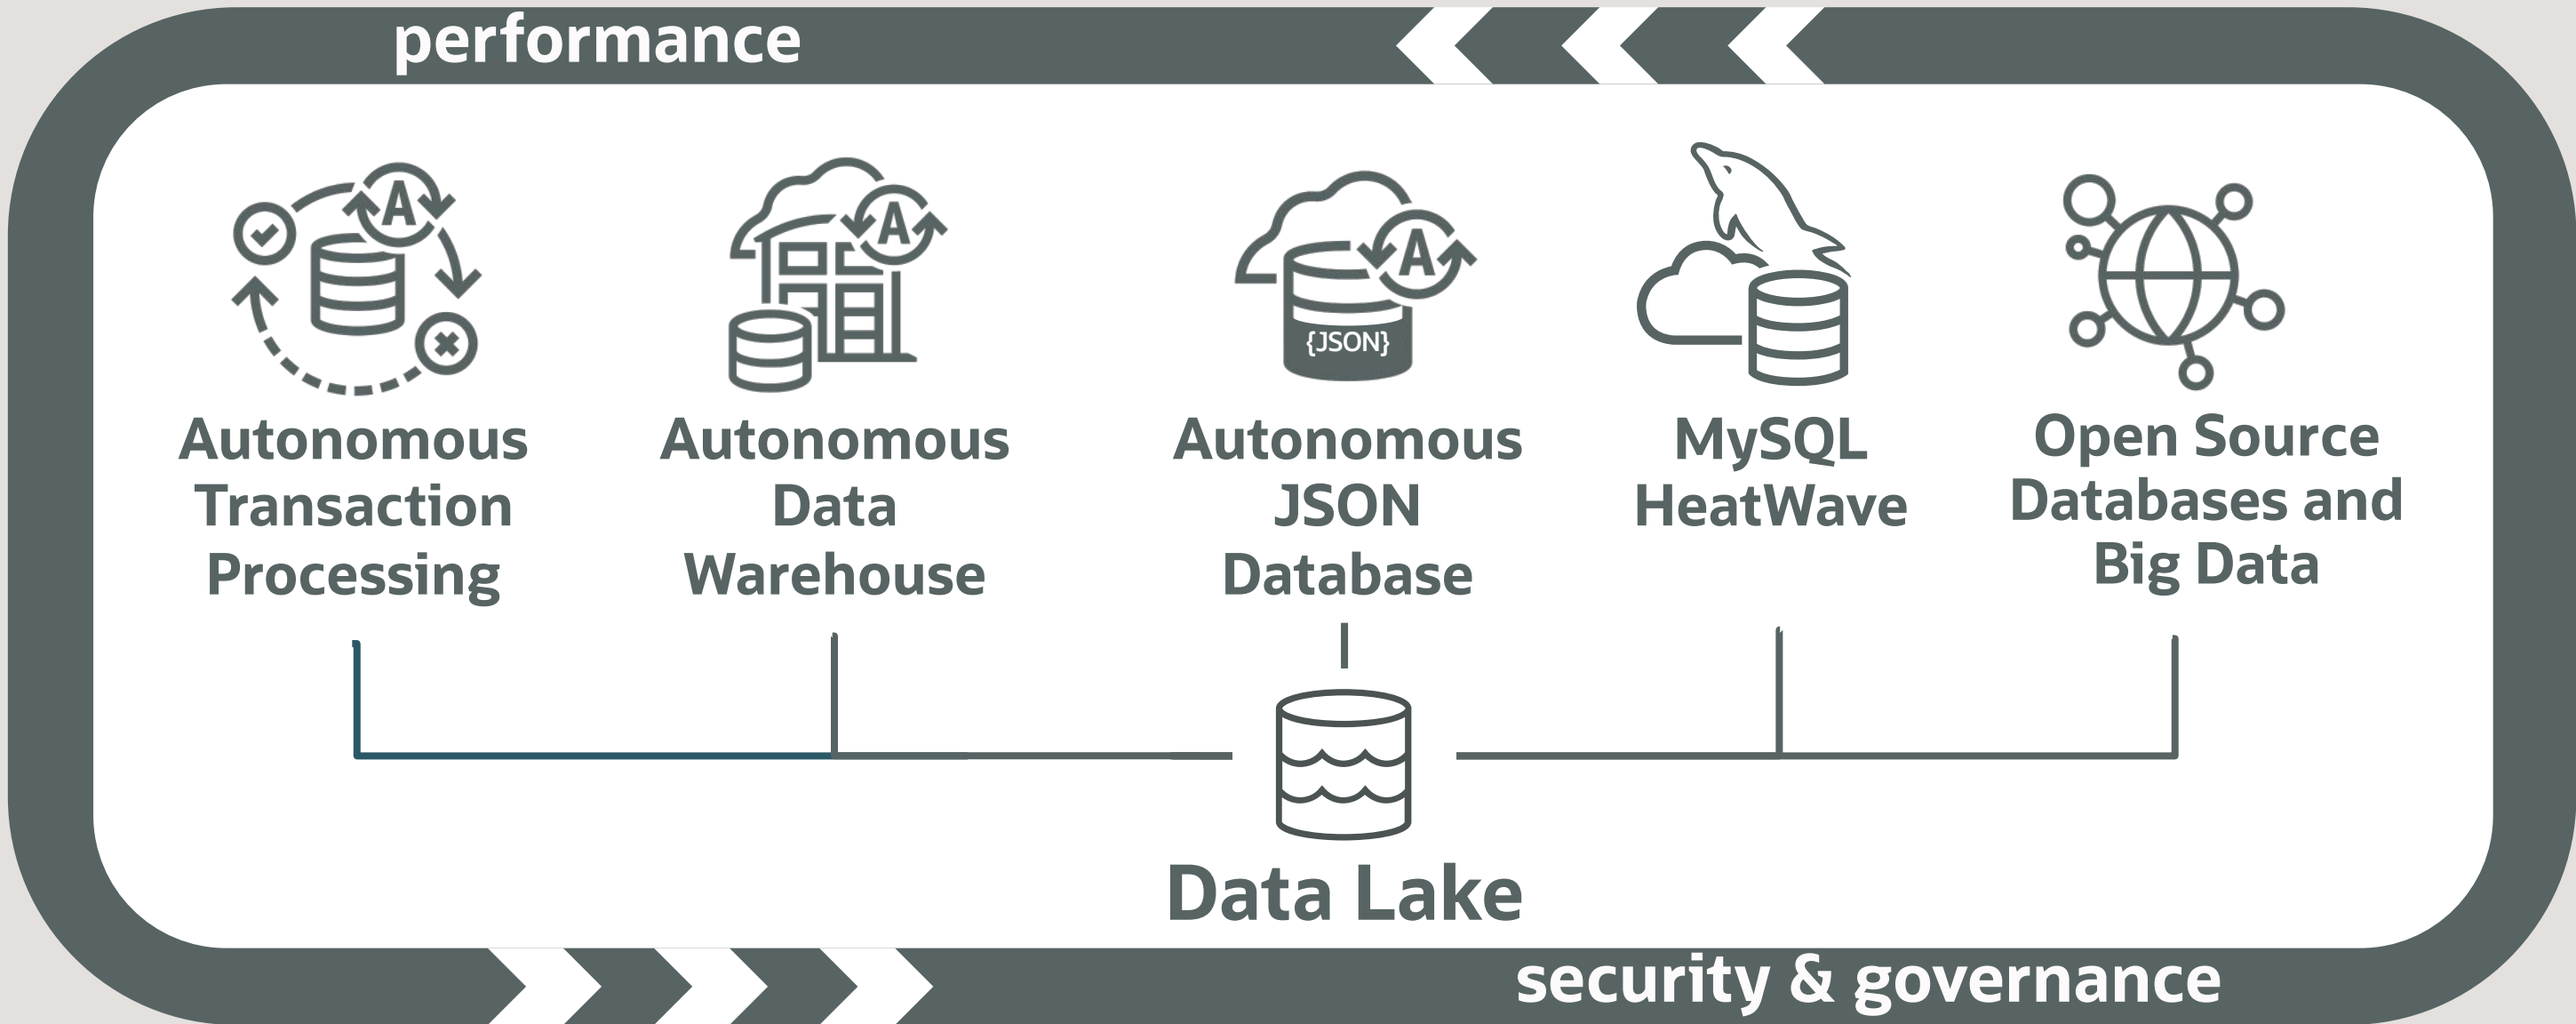
Faster than a self-managed cluster²

99.99%

Availability SLA

Complete and Simple Platform for All Data Management Needs

Choose your Oracle journey



Autonomous Transaction Processing
Pre-configured for row format, indexes, and data caching to accelerate transaction processing and mixed workloads

Autonomous Data Warehouse
Pre-configured for columnar format, partitioning, and large joins to accelerate analytics and data warehouse

Autonomous JSON
Transactions and analytics on JSON data, and includes MongoDB compatible API

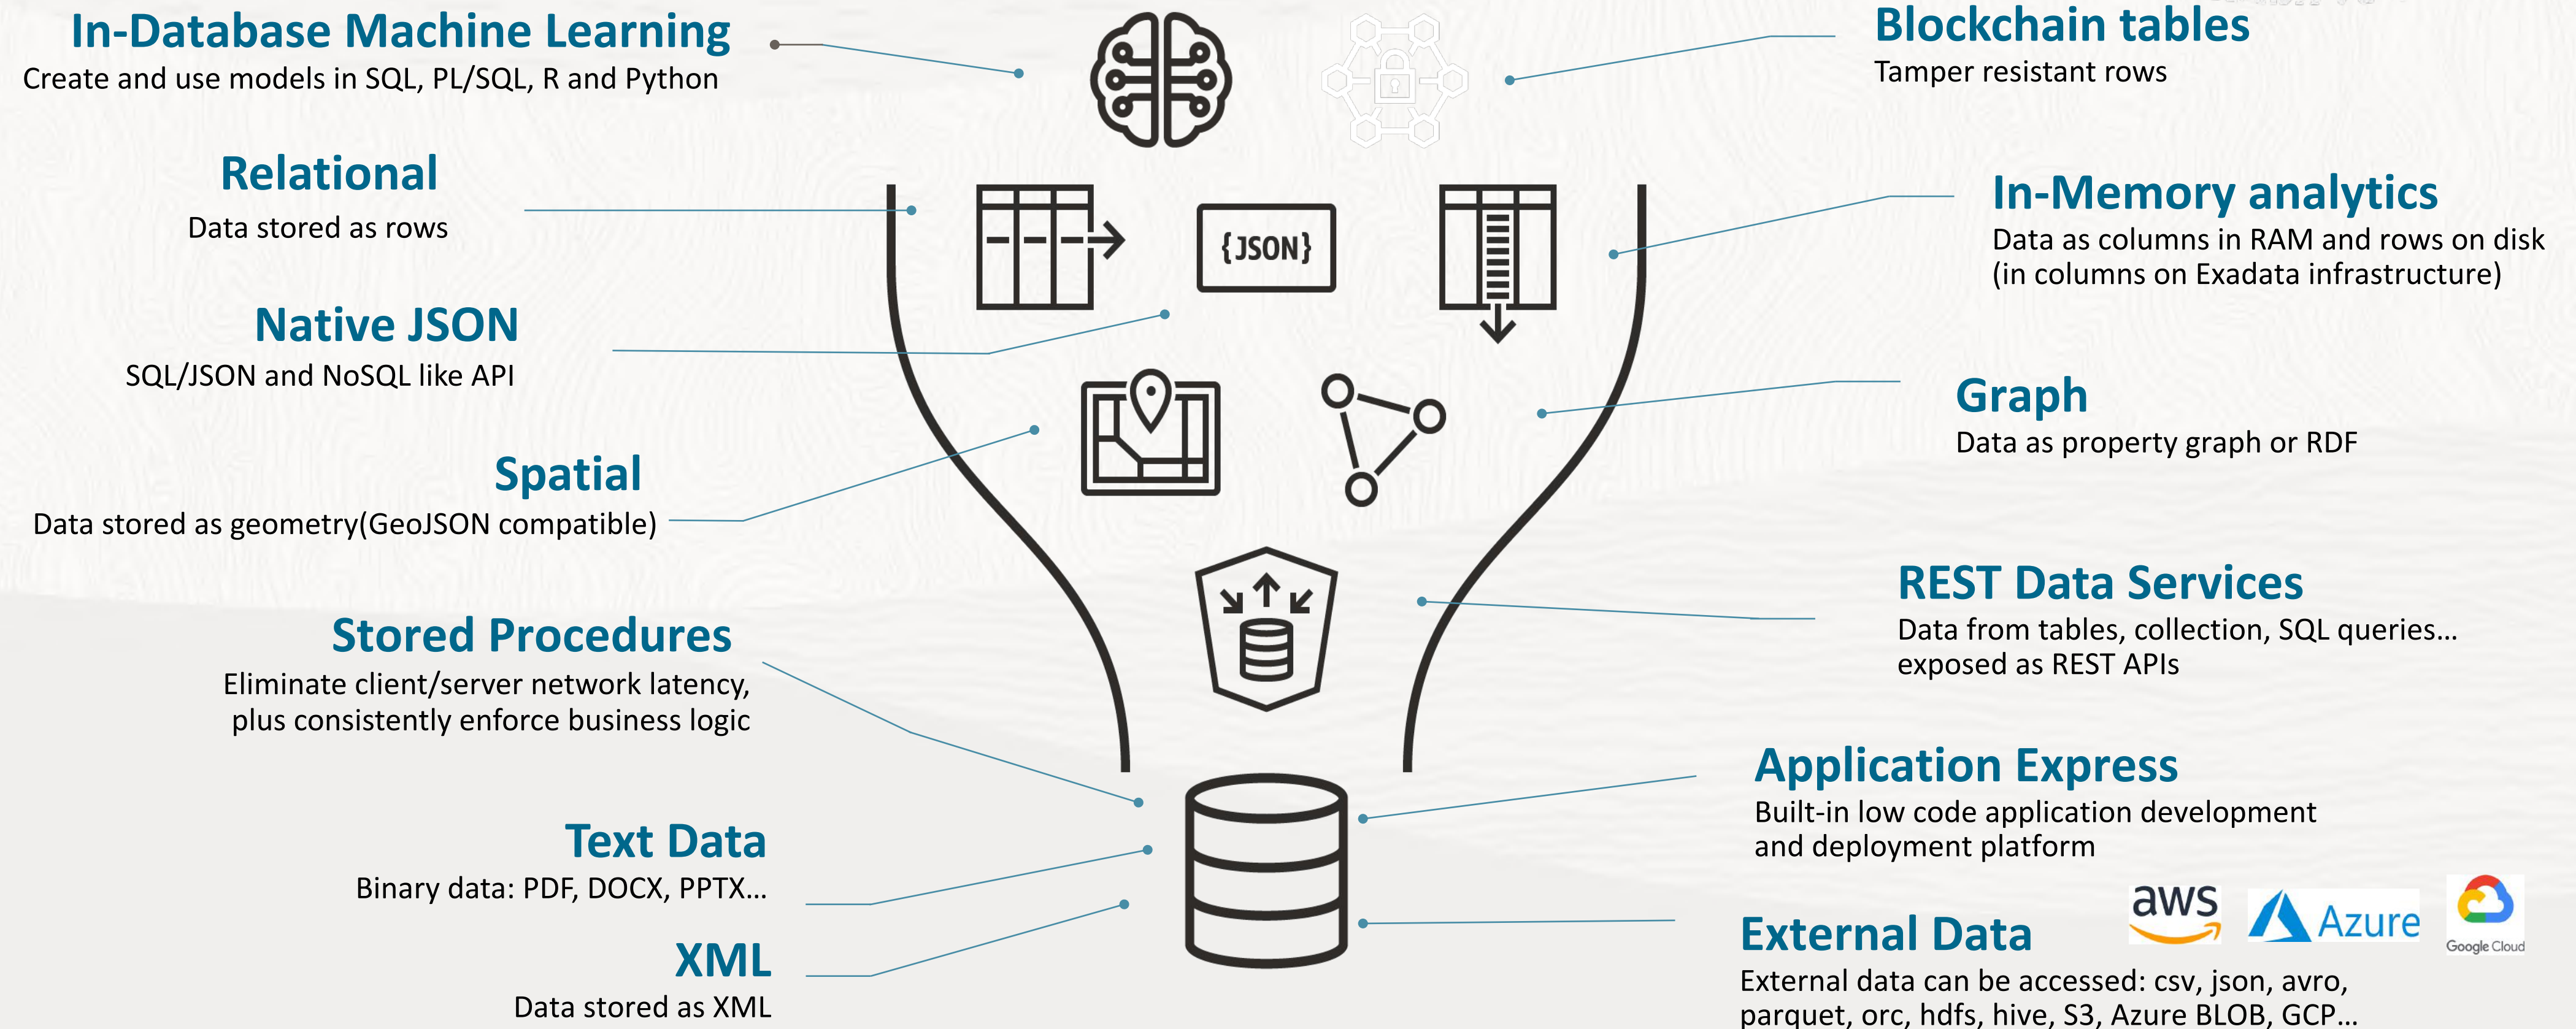
MySQL HeatWave
One MySQL cloud database service for transactions, real-time analytics across data warehouses and data lakes, and machine learning (ML)—without ETL

Open Source Databases and Big Data
A portfolio of managed open source databases and big data services




Autonomous Database Powered by Oracle Converged Database Engine

Runs on Exadata: Move Compute to the Data; Dramatically Simplifies Application Development



Multi-model & Multi-workload

Oracle Database 23^c – The next Long Term Support Release



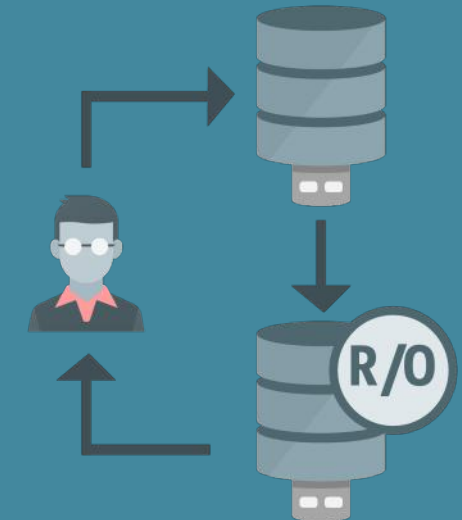
SQL Domains

Schema Privileges

Oracle Database

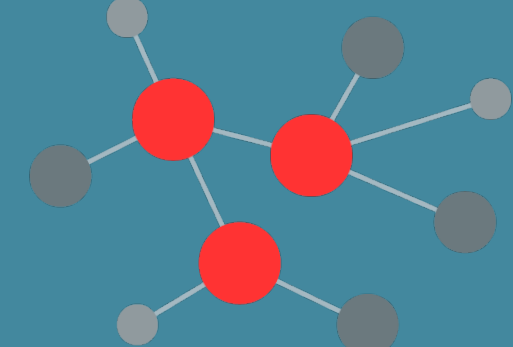
23^c

App Simple




Read-Only Per-PDB Standby

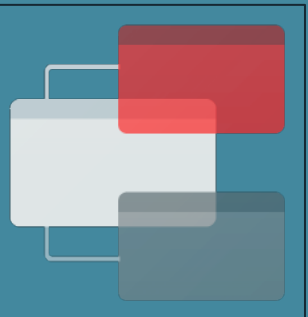
Operational Property Graphs

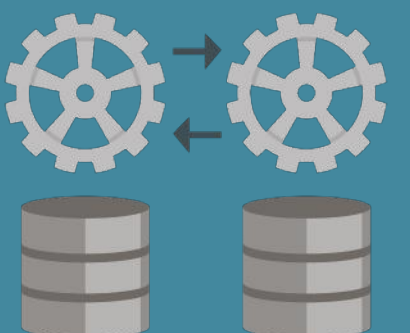


Real-time SQL Plan Management



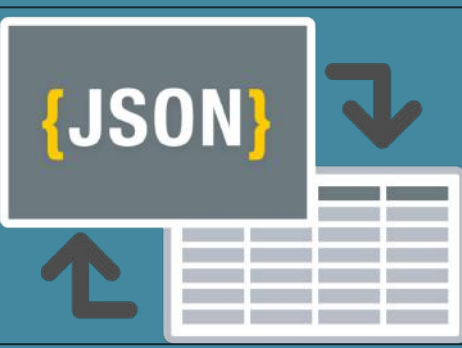
JSON Schema

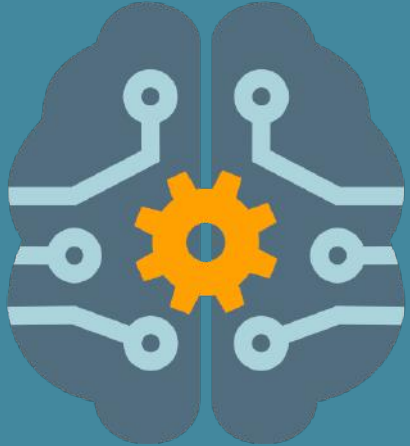




Microservice Support

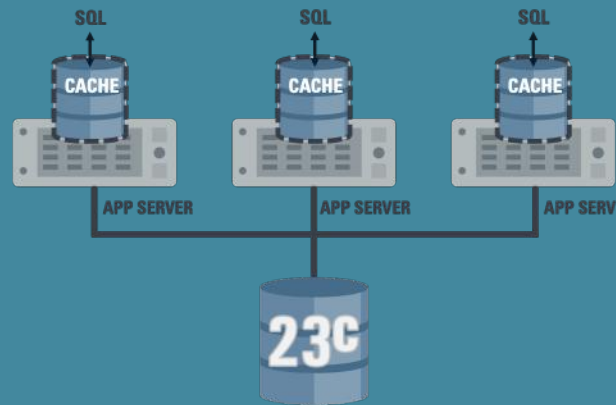
JSON Relational Duality






AI Vector Search

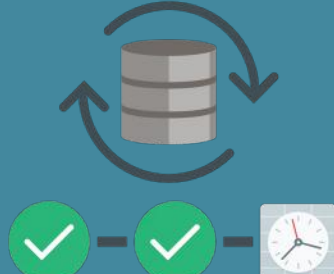
True Cache






SQL Firewall


Priority Transactions



Rolling Patching

JS Stored Procedures






Developer Role

Shrink Tablespace

Boolean Datatype

Globally Distributed Database





MySQL HeatWave

Transactions, real-time analytics across data warehouses and data lakes, and machine learning in one cloud database service

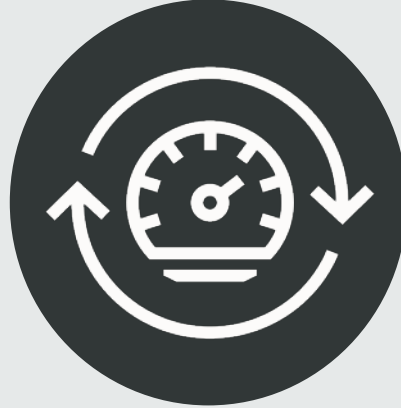

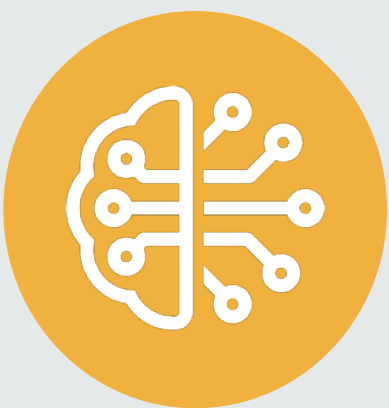



Social, eCommerce, gaming, healthcare, fintech, IoT apps. Analytics and ML tools.



Queries

Results

MySQL HeatWave




OLTP Analytics Vector Store Generative AI In-Database ML Autopilot

Object Store



Parquet

Database exports



For both MySQL and non-MySQL workloads

Data remains in object store, processing is done in HeatWave



Oracle Analytics

Converged analytics platform for all personas, workloads and data



Governed Analytics

LOB/Self-Service Analytics

Augmented Analytics

Dashboards

Distributed Pixel-Perfect Reports

Data Visualization

Self-Service Data Preparation

Voice & Chatbot

Natural Language

Semantic Models

Query Federation

Storytelling

Direct Connectivity

Data Profiling & Enrichment

Gen AI

Briefing Books

Data Export

Collaboration

Mobile


Automated Insights

Machine Learning

Open Connectivity
Oracle Cloud | Third Party Cloud | On-Premises



- Search
- OECD FLOOD HISTORY v2
- ID
 - Row Count
 - Flood Type
 - Geo Type
 - IPAC
 - Geo Code
 - Geo Name
 - Parent Geo Code
 - Country
 - Period
 - Year
 - Flood Recover
 - Pop. Exp %
 - Total Population
 - Pop. Exposed
 - Built Area Exp %
 - Total Built Area
 - Built Area Exposed
 - Crop Land Exp %
 - Land Exp %
 - Total Land Surface
 - Land Exposed
 - Gdp Per Capita
 - Gdp Exposed
 - My Calculations
 - Value Labels



Insights Chat

Other countries with significant flood prevention projects include the **United Kingdom, Germany, France, and Australia**. The projects typically involve a combination of structural measures, such as levees and dams, and non-structural measures, such as flood forecasting and warning system.

What is the population impact of flooding in these countries?

Among the 4 countries with the biggest flood prevention projects, The Netherlands has the biggest population exposure with 57% affected by flooding. 51% is due to Coastal flooding while 16% is due to River flooding. China has the next highest population exposure with 21% affected by flooding. 3% is due to Coastal flooding and 18% is due to River flooding.

Add Flood Type

China, Japan, and the United States has a higher population exposure % from River flooding. The Netherlands has a higher population exposure % from Coastal flooding, which is primarily due to over 50% of the land being exposed to Coastal flooding.

Change Chart Type Add field Replace field Remove field Filter

Create a calculation for GDP Exp ?

Low-code Application development with APEX



World's most popular enterprise low-code platform

Build scalable and secure apps 20X faster, with 100X less code

Delivers the most productive way to develop and deploy mobile, web, and desktop apps everywhere – cloud and on-premises

19M+ apps created

2M+ live APEX apps

3K+ APEX apps created daily

700K+ APEX developers

75% of Fortune 500 use Oracle APEX

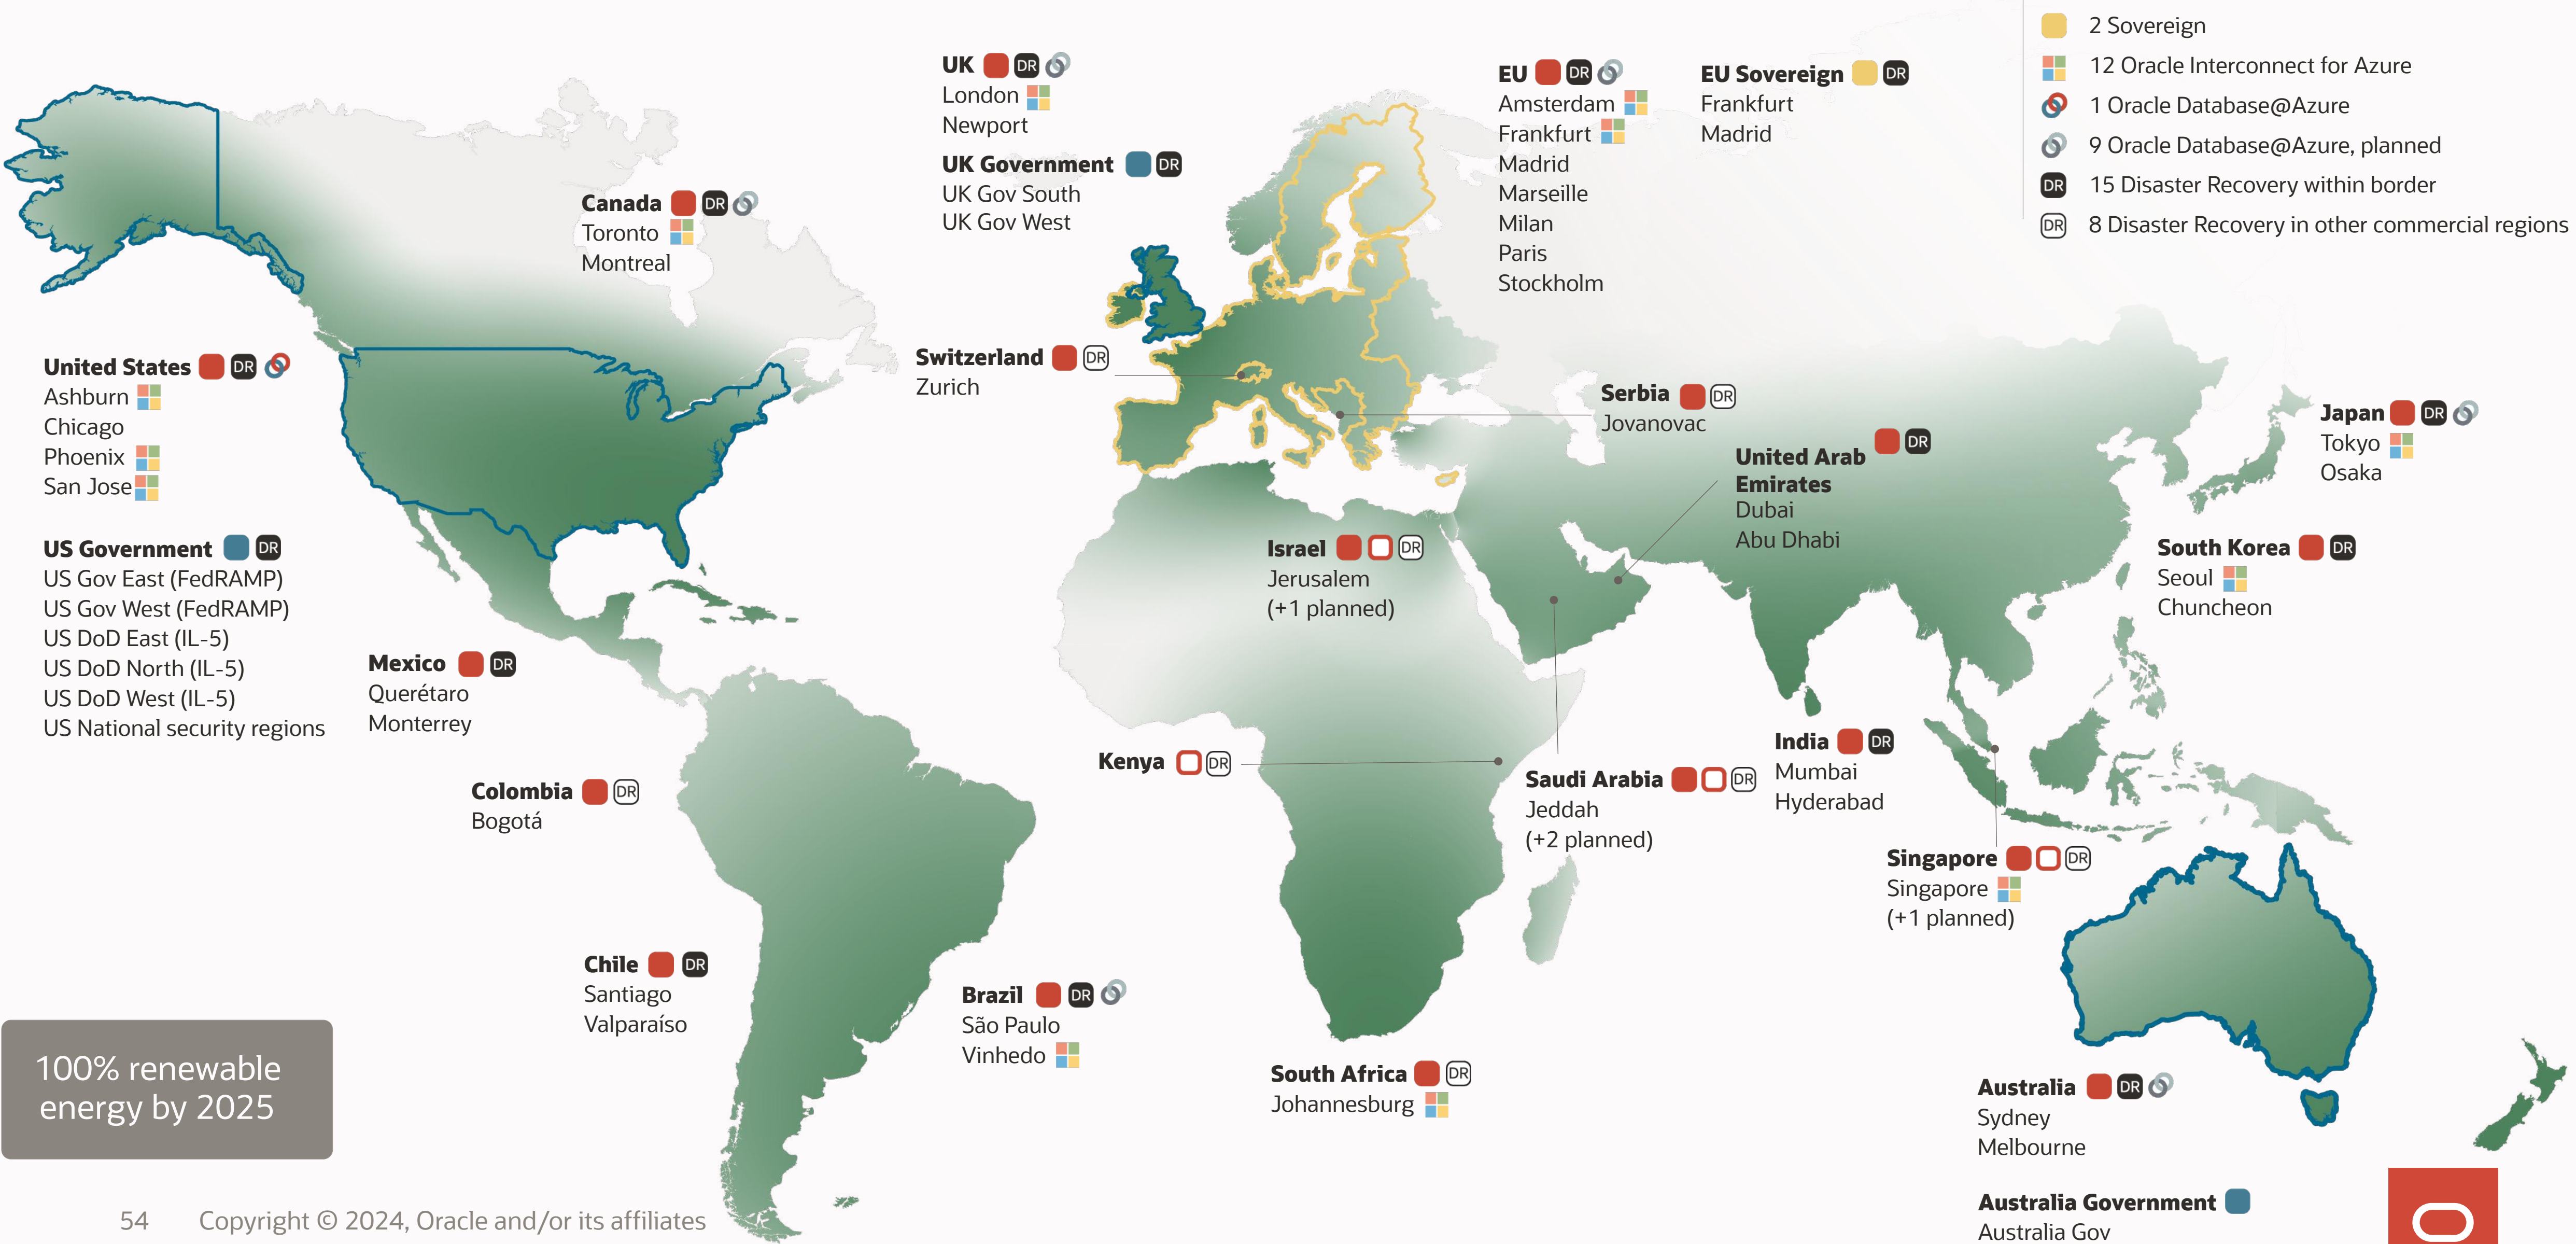
ORACLE

Oracle Cloud Infrastructure


Generation 2 – Internet of Cloud



Deploy in commercial, government, or sovereign regions



OCI's unique architecture




Flex infrastructure

Precise resource fit



Off-box virtualization

Isolation for security, performance, and economics.




Nonblocking networks



L2 network virtualization

Network optimized for performance and service reliability



RDMA cluster networking

Powers New *and* Existing Workloads

Traditional apps

Performance, elastic benefits without re-architecture, modernize incrementally

HPC / Machine Learning

Scale up to 20,000 CPUs or or 512 GPUs in a cluster

Compute-intensive apps

Bare metal instances for CPU and GPUs

Network-intensive apps

No network contention, lowest cost outbound in the market

Critical databases

Native database clustering, scale up to PBs and thousands of cores



Annual cost of a virtual machine
(95% usage over all 12 months)


Flexible sizing vs savings plans

We're (still) noticeably cheaper

OCI and Google Cloud offer flexible CPU virtual machines

AWS and Azure offer fixed sizes, forcing you to “upsize” even if you just need “a little bit more”

Needed vCPUs	OCI	AWS	Azure	Google Cloud
2	2 (1 OCPU)	2 .large	2	2
4	4 (2 OCPU)	4 .xlarge	4	4
6	6 (3 OCPU)	8 .2xlarge	8	6
8	8 (4 OCPU)	8 .2xlarge	8	8
10	10 (5 OCPU)	16 .4xlarge	16	10
12	12 (6 OCPU)	16 .4xlarge	16	12
14	14 (7 OCPU)	16 .4xlarge	16	14
16	16 (8 OCPU)	16 .4xlarge	16	16
18	18 (9 OCPU)	32 .8xlarge	32	18
20	20 (10 OCPU)	32 .8xlarge	32	20
22	22 (11 OCPU)	32 .8xlarge	32	22
24	24 (12 OCPU)	32 .8xlarge	32	24
26	26 (13 OCPU)	32 .8xlarge	32	26
28	28 (14 OCPU)	32 .8xlarge	32	28
30	30 (15 OCPU)	32 .8xlarge	32	30
32	32 (16 OCPU)	32 .8xlarge	32	32



Savings plans are used when cheaper than on-demand pricing.
AWS and Azure lines overlap (with only Azure showing).



Better Cloud Economics

- Private network connectivity that costs **74% less**
- **More than 3X** better price-performance for compute
- **Up to 44% less** expensive infrastructure with local solid-state disks, twice the RAM, RDMA networking, and a performance SLA
- **20X the input/output** operations per second for less than half the cost

		ORACLE	aws	Azure	Google Cloud
COMPUTE	Virtual Machine ¹ (AMD, 4 vCPU, 16 GB RAM Monthly)	\$54	+ 132%	+ 132%	+ 83%
	Dense IO Virtual Machine Instances (\$/OCPU/Hour)	\$0.1275	+ 22%	+ 35%	+ 3% ²
	Bare Metal Standard (\$/OCPU/Hour)	\$0.0638	+ 50%	N/A ³	N/A ⁴
	Kubernetes Cluster (100 vCPU, 750 GB RAM, Monthly)	\$2,373	+ 101%	+ 80%	+ 45%
STORAGE	Block Storage ⁵ (400 GB, 20K IOPS, Monthly)	\$17	69X	25X	69X
	Object Storage ⁶ (1 TB, standard access, Monthly)	\$26	- 6%	- 29%	- 22%
NETWORK	Internet Data Egress (50 TB, Monthly)	\$340	12X	9X	12X
	VPN (10 site to site VPN with 24 hrs connection)	Always Free	Paid	Paid	Paid
	Network Load Balancer	Always Free	Paid	Paid	Paid
	NAT Gateway	Always Free	Paid	Paid	Paid
DATABASE	Private Line Network (1 Gbps, 100 TB Data, Monthly)	\$155	14X	18X	13X
SECURITY	MySQL Database (16 vCPU, 64 GB RAM, 500 GB, Monthly)	\$345	+ 206%	+ 98%	+ 159%
MONITORING	Vulnerability & Threat Scanning	Always Free	Paid	Paid	Paid
	Log Storage	Free 10 GB	Paid	Paid	Paid
	CloudWatch - Detailed and custom Metrics	Free 500Mn Datapoints	Paid	Paid	Paid

Growth Tax Contributors
~ 25%
=> Impact grows with Scale

Gen1 Legacy Tax
~ 15%
Always free services on OCI

¹ Leveraging OCI flexible shape (AMD processors for all cloud vendors)
² No local storage is included. This will be an additional cost
³ Microsoft does not publicly disclose its bare metal pricing
⁴ Google does not publicly disclose its bare metal pricing

⁵ AWS EBS io2, Azure LRS P30, GCP Persistent Disk
⁶ AWS S3 Intelligent Tiering, Azure Block Blob, Hot (LRS), GCP Cloud Storage, Standard. Read and write requests are additional costs (not included here)

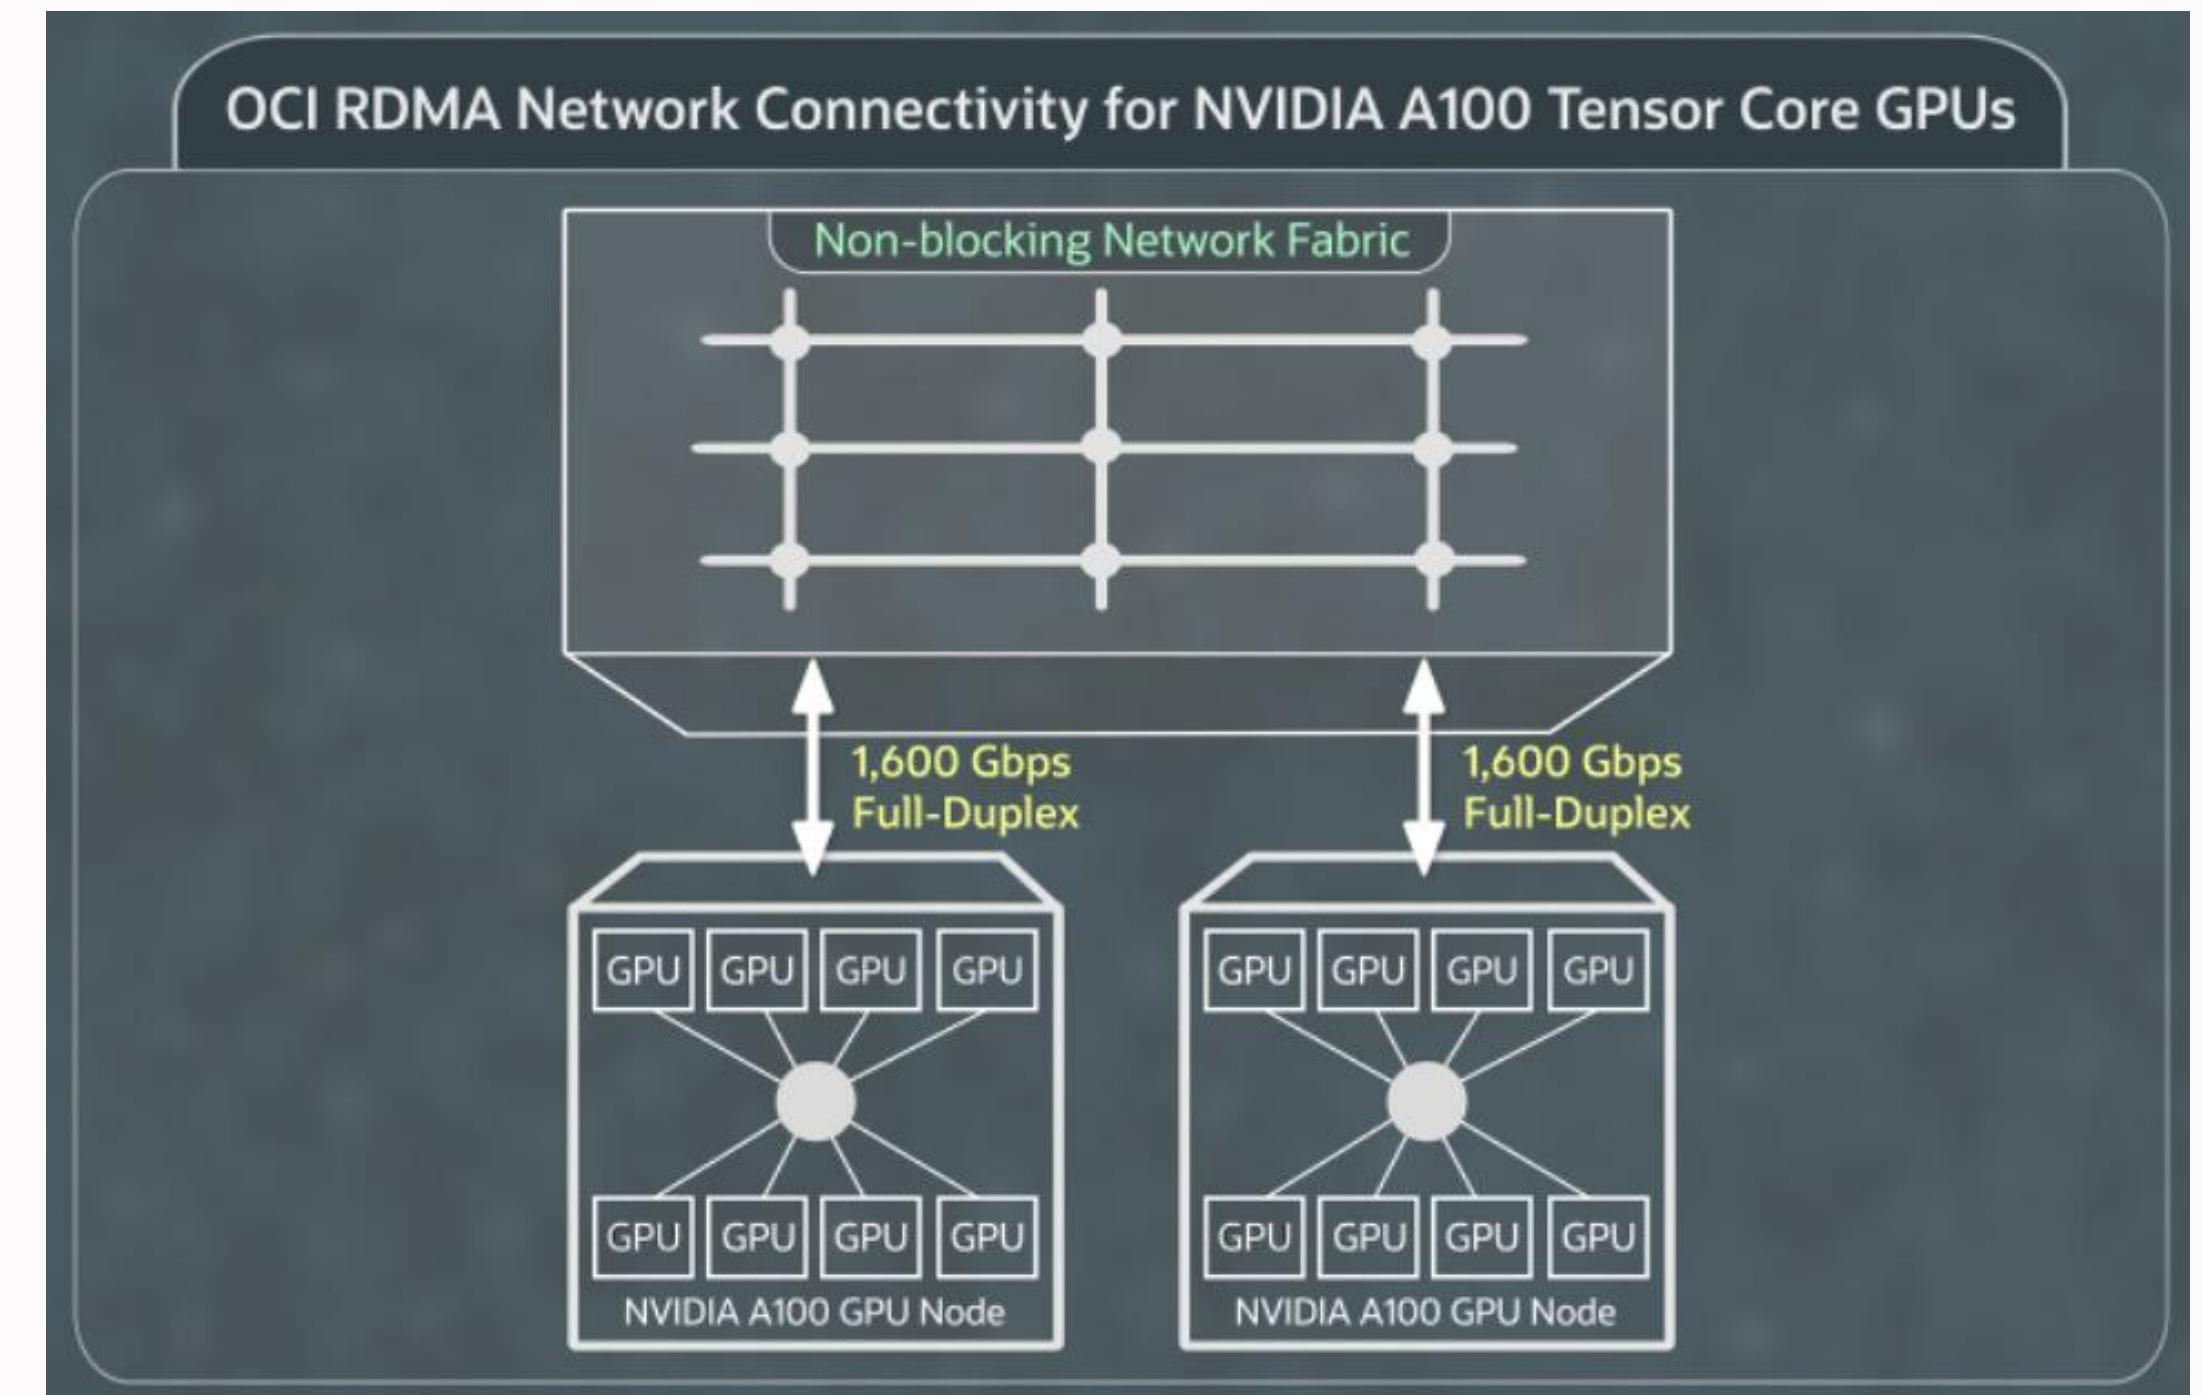


Cluster networking


- Workload parallelization for AI training and Cloud Native Application to reduces time and cost for each run and performance
- Achieved with RDMA, using decades of prior experience in Oracle Database
- Purpose-built, dedicated network based on RoCE v2

Future


- Qualifying next generation Ethernet and Infiniband technologies



OCI Supercluster enables AI training at massive scale



Native support for bare metal instances




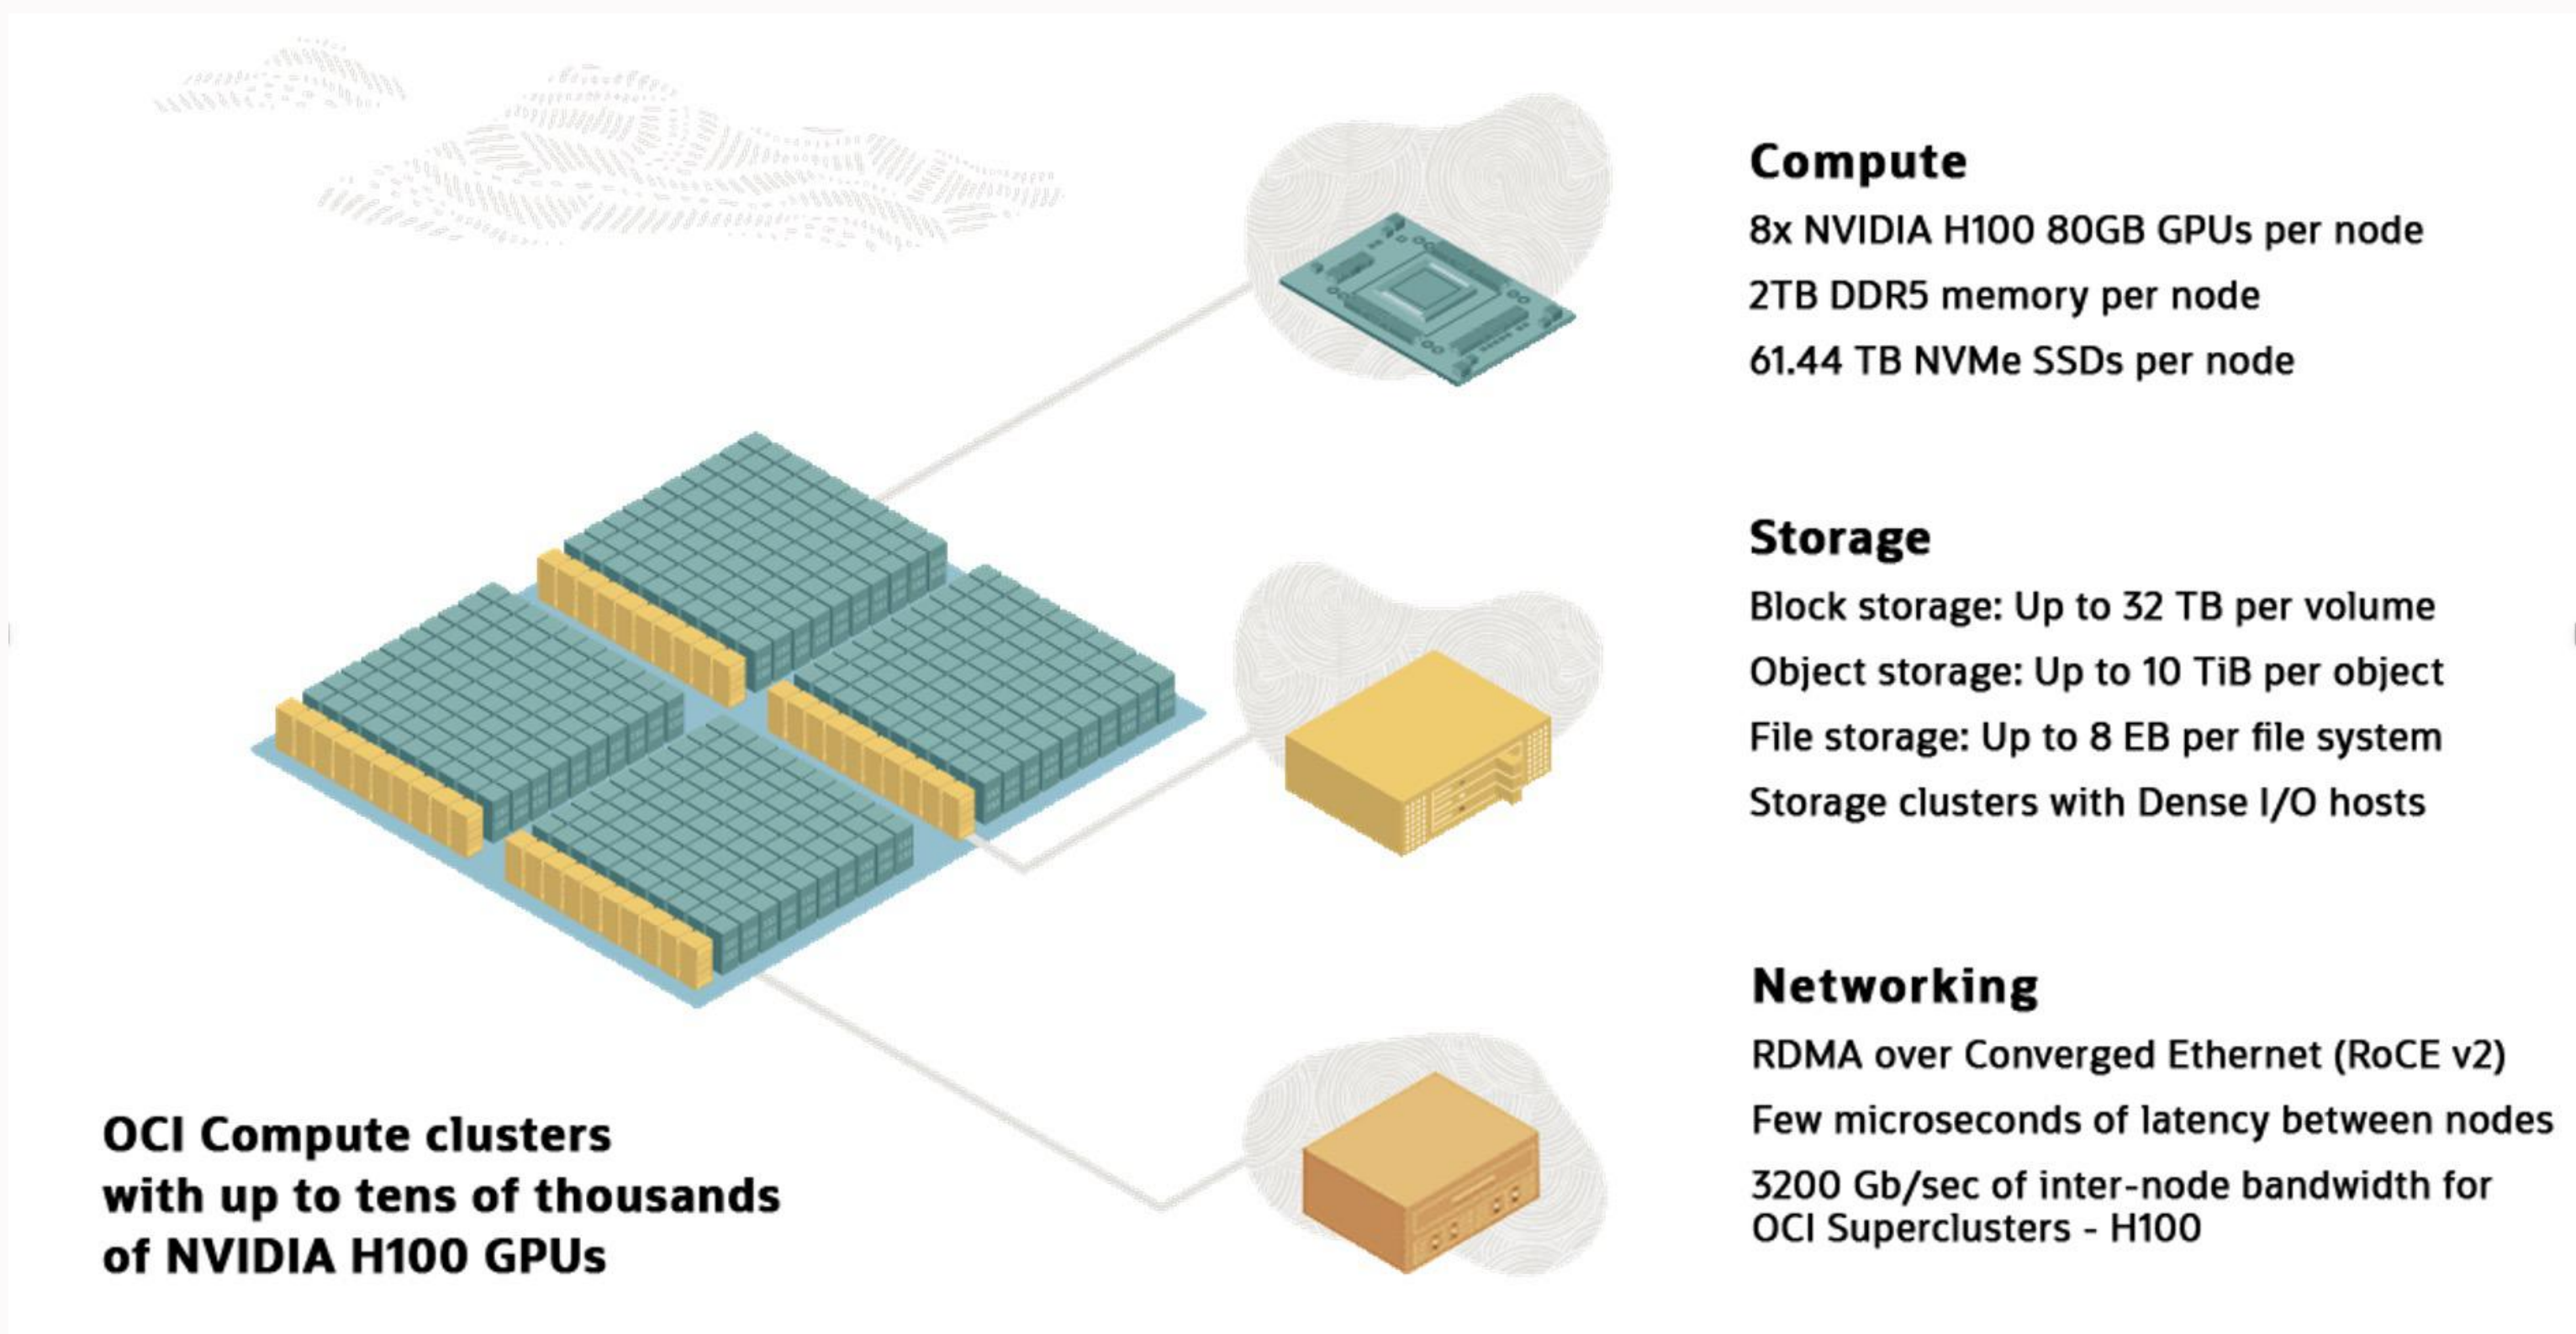
NVIDIA A100 and H100 GPUs



Scalability to tens of thousands of GPUs



Local NVMe Storage



OCI AI Infrastructure (GPU)

OCI

H100-80

8 GPU, 640GB GPU Memory

- ✓ 112 Cores
- ✓ 2048 GB Memory
- ✓ 3200 Gbps Networking
- ✓ 61.4 TB Local NVMe
- ✓ \$80/node

AZURE

- 👎 96 Cores
- 👎 1900 GB Memory
- ✓ 3200 Gbps Networking
- 👎 1 TB Local Storage
- \$/node

AWS

- 96 Cores
- 2048 GB Memory
- ✓ 3200 Gbps Networking
- 30.7 TB Local Storage
- 👎 \$98.3/node

GCP

- 👎 ? Cores
- 2048 GB Memory
- ? Gbps Networking
- ? TB Local NVMe
- ?/node

A100-80

8 GPU, 640GB GPU Memory

- ✓ 128 Cores
- ✓ 2048 GB Memory
- ✓ 1600 Gbps Networking
- ✓ 27.2 TB Local NVMe
- ✓ \$32/node

- 96 Cores (96 VCPU)
- 1900 GB Memory
- ✓ 1600 Gbps Networking
- 6.4 TB Local SSD
- \$37.2/node

- 👎 48 Cores (96 VCPU)
- 👎 1152 GB Memory
- 400 Gbps Networking
- 8 TB Local NVMe
- 👎 \$41.0/node

- 👎 48 Cores (96 VCPU)
- 1360 GB Memory
- 👎 100-200 Gbps Networking
- 👎 3 TB Local Storage
- \$40.6/node

A100-40

8 GPU, 320GB GPU Memory

- 64 Cores
- ✓ 2048 GB Memory
- ✓ 1600 Gbps Networking
- ✓ 27.2 TB Local NVMe
- ✓ \$24.4/node

- ✓ 96 Cores (96 VCPU)
- 900 GB Memory
- ✓ 1600 Gbps Networking
- 6.4 TB Local SSD
- \$27.2/node

- 👎 48 Cores (96 VCPU)
- 1152 GB Memory
- 400 Gbps Networking
- 8 TB Local NVMe
- 👎 \$32.8/node

- 👎 48 Cores (96 VCPU)
- 👎 680 GB Memory
- 👎 100-200 Gbps Networking
- 👎 0
- \$29.4/node

A10-24

4 GPU, 96GB GPU Memory

- 64 Cores
- ✓ 1024 GB Memory
- ✓ n/a
- ✓ 7.68 TB Local NVMe
- ✓ \$8/node

- ✓ 72 Cores (72 VCPU)
- 880 GB Memory
- n/a
- 👎 1.4 TB Local SSD
- 👎 \$13/node

- 👎 48 Cores (96 VCPU)
- 👎 384 GB Memory
- n/a
- 3.8 TB Local NVMe
- \$8.2/node

- 👎 Not Available

✓ best configuration

👎 worst configuration

Details please refer to excel file



Best-of-breed processors for training and inference



OCI Compute

NVIDIA Grace Hopper/Blackwell

Bare metal

NVIDIA H200, B200

Bare metal

VMs for A100, H100, L40S



OCI Compute

AMD MI300X

Bare metal




OCI Compute

Ampere A2

Flexible VMs

Lowest Data Egress Charges by an Order of Magnitude

Outbound Traffic

	First GB	1GB	1TB	10TB	100TB	150TB
Amazon Web Services*	Free – up to 1GB		-	0.085	0.070	0.050
Google Cloud			0.11	0.08	-	-
IBM Cloud	Free – up to 10TB , <u>select</u> regions			0.083	-	-
Microsoft Azure*	Free – up to 5TB , any region			0.083	-	0.05
 Oracle Cloud	Free – up to 10TB, any region			0.0085	-	-






Source: <https://projector.cloud-mercato.com/projects/aws-azure-google-ibm-oracle-vms-q1-2021>


* Only free for Tier-1 customers leaving Cloud

Uniquely, Oracle doesn't charge for data transfer between availability domains in the same region



Recent notable Oracle Cloud Infrastructure customer wins

				
Ridesharing platform modernization	AI training platform	ERP modernization 180TB SAP deployment	Data center exit VMware	Microsoft Bing Conversational Search



Forward-looking statements

This presentation 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.

Statements in this presentation relating to Oracle's future plans, expectations, beliefs, intentions, and prospects are "forward-looking statements" and are subject to material risks and uncertainties. A detailed discussion of these factors and other risks that affect our business is contained in Oracle's Securities and Exchange Commission (SEC) filings, including our most recent reports on Form 10-K and Form 10-Q under the heading "Risk Factors." These filings are available on the SEC's website or on Oracle's website at <http://www.oracle.com/investor>. All information in this presentation is current as of January 2024 and Oracle undertakes no duty to update any statement in light of new information or future events.

Thank you

