

ORACLE
Data Deep Dive

Agenda

10 April, 2026 | Jacob K. Javits Convention Center

New York City



Oracle Data Deep Dive Overview

Oracle Data Deep Dive takes place on the second day of the AI World Tour. This document provides a catalog of sessions and hands on experience topics that you will be able to select from once your registration is approved. You will work directly with Oracle AI Database features that bring AI to the data to reduce latency and complexity across multi-cloud and on-premises.

The Developer Workshop track teaches foundational concepts for building AI applications using the database's integrated AI features. You will learn how vector search works inside the database, how to design effective retrieval and context pipelines, and how to use agent patterns that persist memory and apply business rules. Hands on labs guide you through data ingestion, indexing, retrieval, and response generation so you can turn prototypes into production ready solutions.

Oracle experts and peers will share best practices, reference architectures, deployment patterns, and performance techniques. Learn how to run workloads where they fit best using Oracle's distributed cloud model on-premises, in OCI, or across other hyperscalers while maintaining consistency, cost efficiency, and compliance. You will walk away with patterns you can apply immediately.

How to register: RSVP for "Oracle Data Deep Dive" in your AI World Tour New York City registration to secure your seat. [Click here to Register](#)

New York City Agenda

ORACLE
Data Deep Dive

9:00–10:00AM	The “AI for Data” Revolution is Here—How to Survive and Thrive			
10:00-10:45AM	Build with Oracle Autonomous AI Lakehouse: Petabyte-Scale Analytics, Zero ETL			
10:45-11:00AM	Break			
	Track 1: Learning Sessions	Track 2: Developer Workshop	Track 3: HOLs	Track 4: HOLs
11:00AM-11:45AM	Building Enterprise-Grade AI Agents: Inside Oracle’s Private AI Agent Factory	Data fundamentals for AI application development	Getting Started with Oracle AI Vector Search	From Compliance to Confidence: Supercharging Database Security with Oracle Data Safe
11:45AM–12:30PM	AI Vector Search: Under the Hood			
12:30PM–1:15PM	Lunch Break			
1:15PM –2:00 PM	Create and Implement multi-cloud Applications for Any Runtime and Hyperscaler	Building RAG and Agentic applications	Turn Forgetful AI into Consistent Decision-Making Agents Using Oracle AI Database	The Autonomous AI Lakehouse Data Engineer: The Superhero Enabling Generative AI
2:00PM –2:45PM	Uncompromising Security for your multi-cloud Database			
2:45PM –3:15PM	Break			
3:15PM – 4:00PM	GoldenGate 26ai Everywhere: Real-time Multicloud Streaming for AI	Practical guide to Agent Memory, Context, and Memory Engineering	Build an Innovative Q&A Interface Powered by Generative AI with Oracle APEX	Best Practices to Upgrade to Oracle AI Database 26ai
4:00PM-4:45PM	Build multi-cloud Apps That Never Go Down with Oracle's Globally Distributed DB			
	Closing Remarks			

■ Keynote ■ Breakouts ■ Hands On Lab ■ Developer Workshop - Hands On ■ Ask Anything ■ Breaks



Keynote

“AI for Data” Revolution is Here

How to Survive and Thrive

AI is rewriting the playbook for data, applications, and enterprise innovation. Discover what an AI-driven future looks like at your organization: build faster with a converged data architecture; scale smarter with flexible multi-cloud deployment; protect better with advanced ransomware defense and autonomous recovery. It's all designed for today's demands and delivered wherever you need, from public clouds of your choice to your own data center.



Juan Loaiza

Executive Vice President

Oracle Database Technologies

Track 1 | Learning Sessions

ORACLE
Data Deep Dive

Build with Oracle Autonomous AI Lakehouse: Petabyte-Scale Analytics, Zero ETL

Enterprise data is growing fast, and traditional lakehouses often face performance issues at scale and rely on complex, costly ETL pipelines. These challenges slow insights and hinder AI projects. In this session, discover how to build a high-performance, petabyte-scale lakehouse with Oracle AI Database, eliminating the need for expensive ETL. Explore real-world architectures, performance benchmarks, and customer success stories that showcase Oracle's efficient data management at scale.

Building Enterprise-Grade AI Agents: Inside Oracle's Private AI Agent Factory

As enterprises move from AI experimentation to operational impact, the need for secure, governed, and scalable AI agents has never been greater. Oracle's Private AI Agent Factory enables organizations to design, deploy, and manage AI agents entirely within controlled environments using standard REST interfaces, MCP-based agent coordination, and zero-trust security principles. This session explores how enterprises can accelerate AI adoption while maintaining data sovereignty, regulatory compliance, and operational resilience, turning AI agents into trusted digital coworkers rather than isolated experiments.

AI Vector Search: Under the Hood

Discover how Oracle AI Database 26ai streamlines AI app development with built-in AI Vector Search. Learn about the new vector data type, approximate indexes, enhanced SQL, and performance optimizations that make integration fast and seamless. See how vectors power Large Language Models via RAG and Agentic RAG, using your business context for advanced solutions. Learn how Vector Search works with core enterprise features like transactions, RAC, Exadata, security, partitioning, and sharding. See how Oracle delivers scalable, enterprise-grade AI search for real-world needs.

Build multi-cloud Apps That Never Go Down with Oracle's Globally Distributed DB

Enterprise data and transaction volumes are surging, demanding 24/7 availability and potential compliance with data sovereignty regulations. To meet these needs, organizations are distributing applications and data across availability zones, regions, countries, and clouds. Learn how customers are building global-scale applications that process millions of transactions per second, manage petabytes of data, eliminate downtime, and address compliance. Discover how Oracle Globally Distributed Autonomous AI Database seamlessly delivers strong data consistency, full SQL capabilities, and support for structured, unstructured, and generative AI-driven data.

GoldenGate 26ai Everywhere: Real-time Multicloud Streaming for AI

Discover how Oracle GoldenGate 26ai enables real-time data replication and streaming across Oracle Cloud Infrastructure, Azure and Google Cloud. Explore architecture patterns that focus on real-time AI use cases for Agentic, RAG and GraphRAG use cases. Learn to deploy GoldenGate in hybrid and multicloud environments, optimize latency and scale, and integrate with native services like AI Data Platform, AI Lakehouse, BigQuery and Microsoft Fabric. Plus, get an exclusive preview of upcoming features including enhanced automation, observability, and deeper cloud-native integration - making multicloud smarter than ever.

Uncompromising Security for your multi-cloud Database

Maintain full control of your data by leveraging the built-in automation, controls, and security features of Oracle Database services in partner clouds. Join this session to discover how Oracle Database Security delivers 360-degree visibility and enforces robust controls at the source to prevent unauthorized access. Oracle provides comprehensive security at enterprise scale for your entire database fleet, no matter where they reside. Learn how Oracle AI Database 26ai offers multi-layered, real-time defense and quantum-resistant encryption, delivering a comprehensive framework to secure organizations against Agentic AI-powered attacks.

Create and Implement multi-cloud Applications for Any Runtime and Hyperscaler

In today's multi-cloud environment, simplicity is key for developers and administrators. In this session, discover how centralized configuration providers streamline management of Oracle Database client settings, identity, secrets, and monitoring—reducing complexity and overhead. Oracle's JDBC, Python, .NET, and C drivers support these providers across on-premises, Oracle Cloud, Azure, and now Google Cloud Platform and AWS, enabling consistent and efficient database connectivity everywhere.



Data fundamentals for AI application development	This workshop introduces the data modeling and database concepts that underpin modern AI applications, with specific attention to Unified Modeling Theory (UMT), vector data, understanding similarity metrics, vector search, and more. The hands-on labs tie these concepts directly to higher-level patterns such as Retrieval-Augmented Generation (RAG) and agentic systems, making the session ideal for developers building production-ready applications.
Building RAG and Agentic applications	Building on the first session, this workshop focuses on implementing a production-style RAG pipeline and evolving it into an agentic application. Using LangChain and Python, developers work through hands-on labs covering data ingestion, vector search, retrieval, and response generation, then extend into agent workflows, all backed by Oracle AI Database as the system of record and memory core. All implementation is done in Python, giving you practical, runnable code throughout. The emphasis is on practical patterns you'll actually use.
Practical guide to Agent Memory, Context, and Memory Engineering	Building on the previous session, this hands-on workshop shows how to implement agent memory as a first-class capability end to end on Oracle AI Database. Explore episodic, semantic, and procedural memory architectures. With LangChain and Python, you will build agents with persistent state, cross-session recall, and reusable patterns that improve over time. Learn context engineering techniques such as offloading and compaction, plus memory engineering for retrieval and indexing. See how memory fits into the agent loop from ingestion and recall to reasoning and action, with runnable Python code and production-ready patterns.



Track 3 | Hands-on Labs

Getting Started with Oracle AI Vector Search

Discover how Oracle AI Database 26ai streamlines AI app development with built-in AI Vector Search. Learn about the new vector data type, approximate indexes, enhanced SQL, and performance optimizations that make integration fast and seamless. See how vectors power Large Language Models via RAG and Agentic RAG, using your business context for advanced solutions. Learn how Vector Search works with core enterprise features like transactions, RAC, Exadata, security, partitioning, and sharding. See how Oracle delivers scalable, enterprise-grade AI search for real-world needs.

Turn Forgetful AI into Consistent Decision-Making Agents Using Oracle AI Database

Most AI assistants forget everything once a conversation ends. Real businesses cannot afford that. In this hands-on lab, you will follow Seer Equity as it uses AI to support loan approvals. You will see what breaks when AI cannot remember prior decisions, customer context, or internal policies, leading to inconsistent outcomes and lost trust. Through guided exercises, you will turn a forgetful assistant into a reliable agent running in the database that remembers what matters, applies business rules consistently, and knows when to act. You will get hands-on with Oracle AI Database and Select AI Agent, focusing on concepts and outcomes rather than code.

Build an Innovative Q&A Interface Powered by Generative AI with Oracle APEX

From precisely summarizing intricate data to crafting context-aware responses, Generative AI stands at the forefront of a technological shift, promising to redefine how we navigate and leverage knowledge in our day-to-day interactions. Enabling your APEX applications with generative AI capabilities has never been easier. In this workshop, you learn how to bring generative AI capabilities based on large language models (LLMs) to your applications built using Oracle APEX, your favourite low-code Platform. The Generative AI service can be accessed through REST APIs, and by using the powerful REST Data Source capabilities of APEX, you can effortlessly incorporate this advanced technology into your applications using a low-code approach.

Track 4 | Hands-on Labs

ORACLE
Data Deep Dive

From Compliance to Confidence: Supercharging Database Security with Oracle Data Safe

Seven years of audit retention. Developers need realistic test data. New releases to validate. Orphaned accounts to clean up. You are not alone. Oracle Data Safe is a cloud-native service that turns these security chores into an integrated control center for your Oracle Databases. Assess configurations for drift, discover and mask sensitive data for non-production, continuously monitor user activity, and centralize audit retention at scale. This session covers real scenarios to prepare compliant datasets, detect risky privileges, prove audit readiness, and help teams reduce risk while accelerating delivery. You will leave with actionable patterns and a clear path to security confidence.

Best Practices to Upgrade to Oracle AI Database 26ai

Oracle AI Database 26ai is available on-prem finally. Now is the time to start your upgrade testing. With AutoUpgrade and refreshable clone PDBs, you can significantly reduce downtime and streamline the upgrade process. Learn from our real world customer examples as we are going to demonstrate how to upgrade to Oracle AI Database 26ai with minimal disruption, covering key pre- and post-upgrade best practices for a smooth and efficient transition. Don't miss this opportunity to optimize your upgrade strategy and maximize database availability.

The Autonomous AI Lakehouse Data Engineer: The Superhero Enabling Generative AI

AI's impact depends on trusted data. Explore how Oracle Autonomous AI Lakehouse data engineers deliver reliable data from lakes and multi-cloud sources. In this hands-on lab, use Data Studio, an intuitive embedded toolset, to integrate, transform, and share data efficiently, enabling AI models to use high-quality data for better insights. By the end, you will integrate internal and external sources, create and publish data products, and securely share insights across teams without duplication or delay. Gain practical skills to design scalable data workflows that drive collaboration and innovation in finance.