ORACLE

Agentic AI in the Enterprise

OCI Enterprise Agents

Pank Sharma

Director AI COE

May, 2026

Confidential - Oracle Restricted \Employees Only

The enterprise shift: AI moves from answering to doing

Agentic AI turns intelligence into accountable execution across systems.

Passive AI

Answers questions
Drafts content
Summarizes knowledge
Leaves handoffs to people

Agentic AI

Plans toward a goal
Uses tools and memory
Executes across systems
Verifies and reports outcomes

Enterprise value = trusted work completion with visible control.

Challenges of building production-ready AI Agents

Go beyond traditional software or workflow automation

CHALLENGES:

- x Coordinating action
- x Confirming success
- x Handling errors
- x Maintaining audit trails

RESULT IN:

- x Brittle integrations
- x Rigid workflows
- x Opaque governance
- x Inflexible orchestration

Confidential - Oracle Restricted \Employees Only

What makes an AI system agentic?

A goal-driven loop: reason, choose tools, act, verify, and improve.

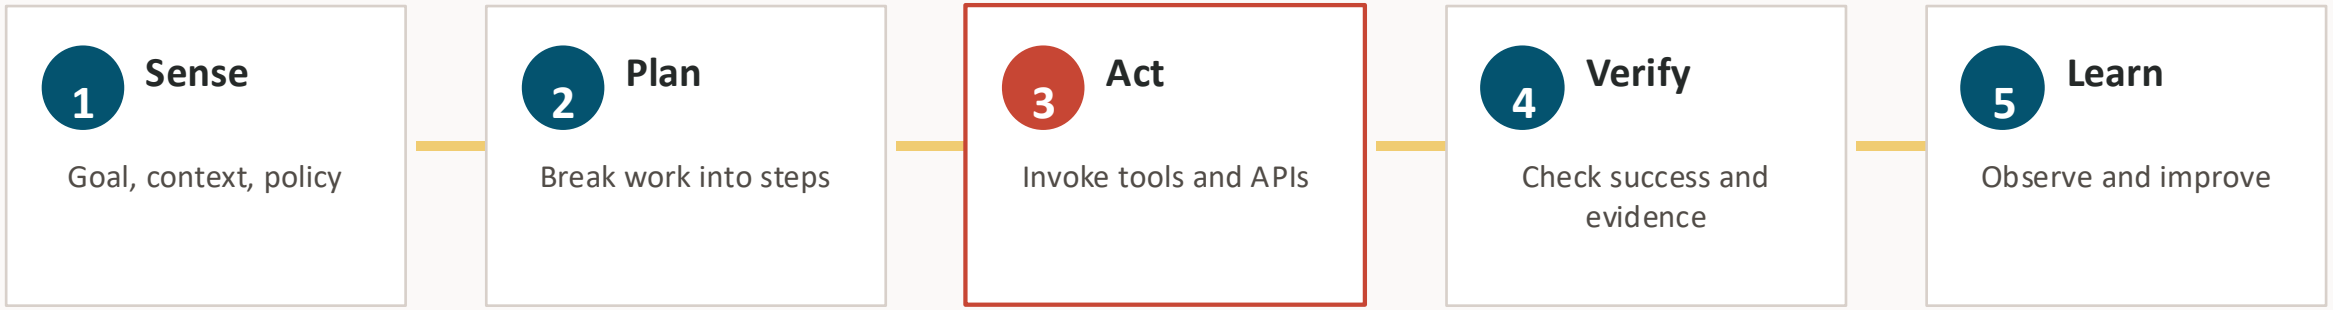

Enterprise agents are productive only when the loop is bounded by policy, identity, data access, and observable outcomes.

Adoption path: move autonomy one boundary at a time

Scale ambition by making controls, ownership, and evidence stronger at each step.

Assist

Human asks;
AI prepares

Drafts and summaries

Automate

AI completes bounded tasks with approval

Triage and retrieval

Orchestrate

AI coordinates multi-step workflows

Tools, memory, human checkpoints

Operate

AI monitors and acts by exception

Runbooks, recovery, continuous optimization

Autonomy increases only as guardrails, evaluation, and observability become explicit.

OCI AI Platform Agents Lifecycle

01 **IDEATE & DESIGN:**
Identify key business problems and define how AI agents can address them, ensuring objectives align with organizational needs.

02 **BUILD:**
Develop the AI agent's architecture, core logic, and integrations, focusing on scalability, security, and compliance from the outset.

03 **TEST:**
Validate agent performance, accuracy, and reliability in real-world scenarios using robust test data and evaluation criteria.

04 **DEPLOY:**
Implement the AI agent within the chosen environment, following best practices around access, security, and compliance.

05 **ORCHESTRATE:**
Coordinate how multiple agents or services work together, optimizing workflows for seamless automation and communication.

06 **ACT:**
The agent executes tasks by invoking tools, querying memory, and taking real-world actions.

07 **OBSERVE:**
Continuously track agent performance for ongoing improvement, troubleshooting, and risk mitigation.

OCI AI Platform Agents Lifecycle

Agent Management
Agents are built, deployed, and observed.

- Authoring
- Testing & evaluation
- Observability
- Hosted deployment

Agent Orchestration
Agents plan how to execute the task given.

- Human-in-the-loop
- Open-source Frameworks
- Protocols (MCP, A2A)

Agent Runtime
Agents act. Where the work happens

- Response API
- Vector Stores
- Tools
- Memory

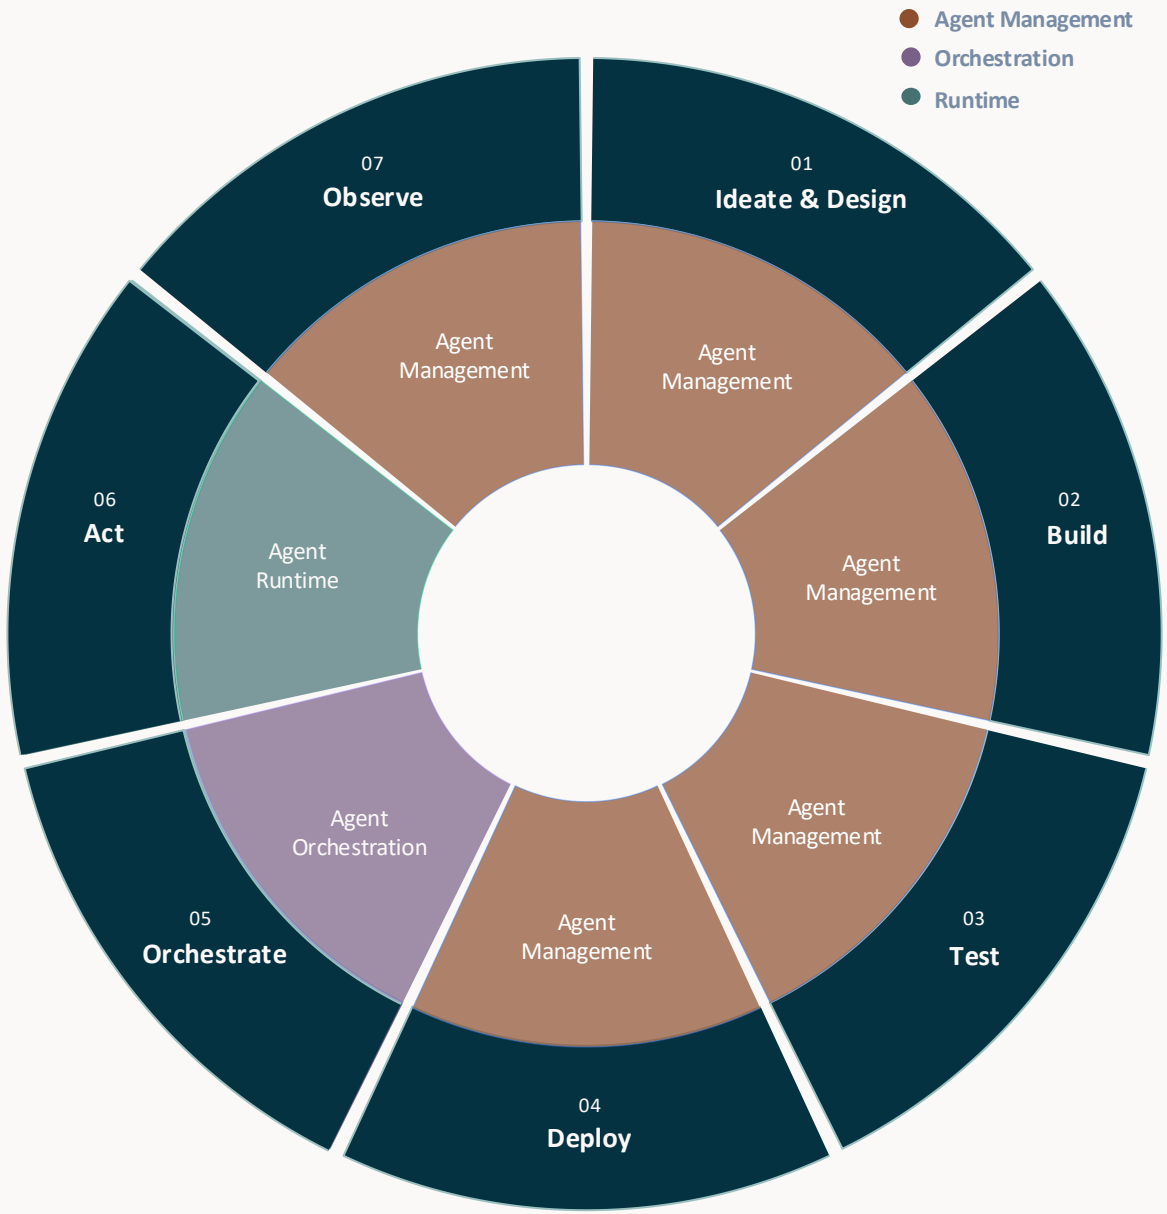

Confidential - Oracle Restricted \Employees Only

Building agentic workflows with OCI Enterprise AI Agents

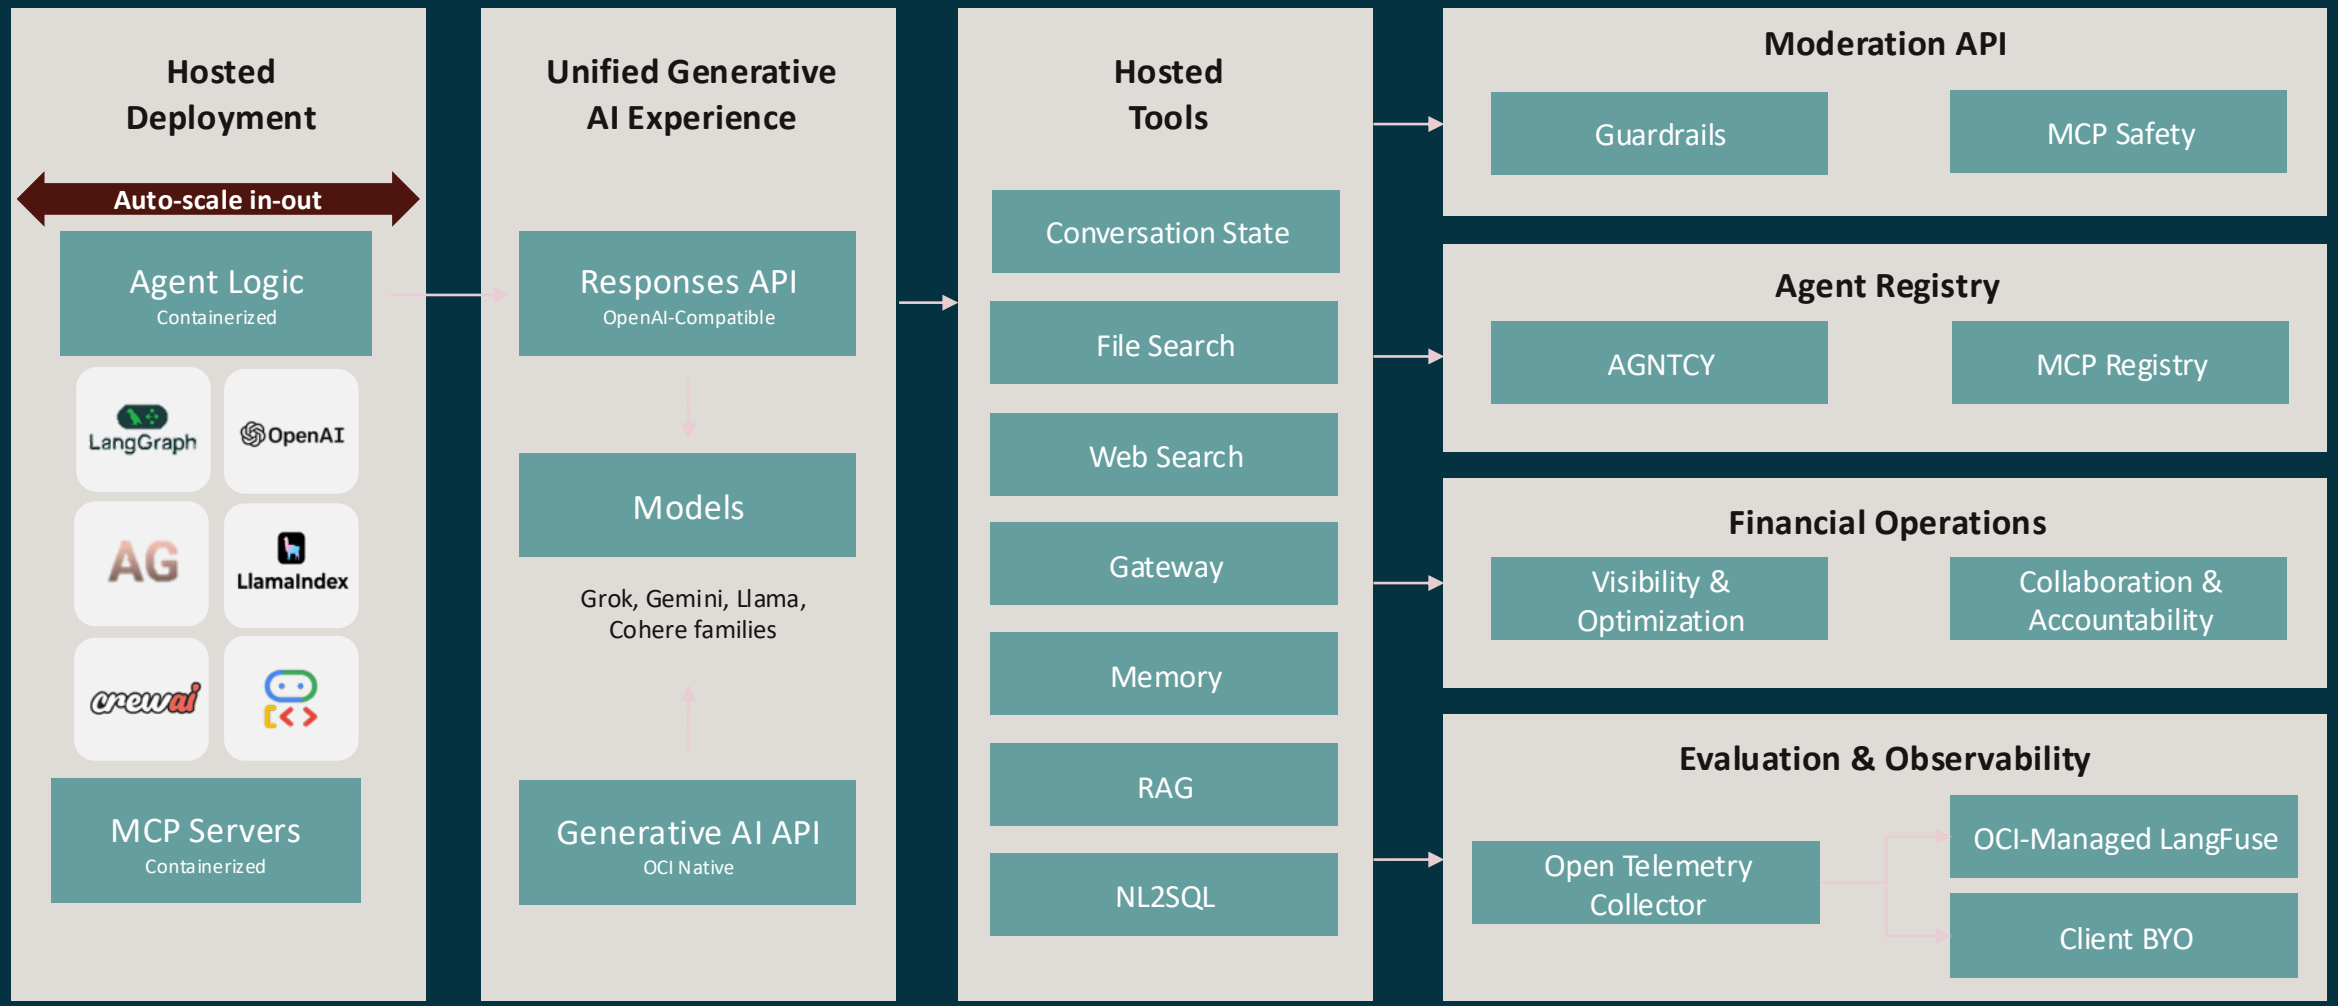

Confidential - Oracle Restricted \Employees Only

Three decisions determine whether agents scale

The platform story is easier when management, orchestration, and runtime have separate owners.

Manage agents

How do we author, test, deploy, observe, and version agents?

1

Orchestrate work

How do agents coordinate frameworks, protocols, tools, and people?

2

Run actions

Where do tools, memory, APIs, routing, and execution actually happen?

3

OCI Enterprise AI Agents: Agent Management

Hosted deployment options

Build & Self-Host

Maximum flexibility

- Define agent behavior locally
- Use LangChain, LangGraph, AutoGen, OpenAI Agent SDK
- Maintain local tool definitions and context
- OCI as inference and orchestration backend
- No proprietary abstractions required

Seamless Deployment

Operational simplicity

- CLI, API, and UI-driven deployment
- No infrastructure setup or tuning
- Clear dev/staging/production separation
- Fast iteration across environments
- For teams focused on business problems

Fully Hosted Agents

Zero operational overhead

- Operated, scaled, and secured by OCI
- Enterprise governance by default
- No customer-side operational ownership
- Predictable cost, managed availability
- Rapid time-to-value

Confidential - Oracle Restricted \Employees Only

OCI Enterprise AI Agent Agents: Agent Orchestration

Flexible options for frameworks and protocols

Frameworks

LangGraph

LangChain

crewai

AG
AutoGen

Semantic Kernel

LlamaIndex

Protocols

- Model Context Protocol (MCP)
- Agent-to-Agent (A2A)

Resulting in:

- Build with the frameworks your team already knows. No migration. No retraining
- Accelerate time-to-value
- Avoid vendor lock-in
- Keep your orchestration layer open, composable, and future-proof
- Insert human checkpoints where your workflows demand it.

Confidential - Oracle Restricted \Employees Only

OCI Enterprise/Generative AI Agents

Confidential - Oracle Restricted \Employees Only

OCI Enterprise AI Agents: Agent Runtime

Tools for agents

File Search

Vector store-backed retrieval with semantic search and reranking

Code Interpreter

Sandboxed code execution for data analysis and computation

Containers

Bring your own code as hosted tools with managed lifecycle

SQL

Schema-aware SQL generation and execution against Oracle databases

RAG

Managed retrieval-augmented generation with vector stores

Memory

Short-term compression and long-term memory for agent continuity

Gateway

Connect to external APIs and services with managed auth

Supporting capabilities

Vector Stores

Fully managed, auto-scaling knowledge bases with semantic search, reranking, metadata filtering.

Moderation & Guardrails

Responsible AI enforced consistently across inference, agent actions, and tool invocations. Configurable per deployment.

Observability

Request/response logs, scaling events, auth logs, correlation IDs. Compatible with LangSmith, LangFuse, and more without code changes.

Confidential - Oracle Restricted \Employees Only

Oracle's AI Capabilities

Enerprise AI Platform

Enerprise AI Services

Fully Managed Multimodal platform, secure private access to LLM

Responses API
OpenAI-Compatible

Enterprise AI Agents

Rapidly build, manage, and deploy agents. RAG, SQL or Custom

Generative AI API
OCI Native

Oracle Applications

Embedded Agents

Custom Agents

OpenAI · Llama · Gemini · Grok · cohere · Hugging Face

AI Data Platform

Multi-Cloud
AI lakehouse

Agentic UI & AI tools

The agent UI interface where agents manage user interaction

Chat | Insight | Workflows | Agent

AI Lakehouse

Integrates advanced data engineering, AI models and tooling, Analytics

Analytics | Data Science

Data and AI Foundation

Industry-leading AI models and frameworks, running on optimized OCI infrastructure

Models | Open Source | Data & AI Catalog

PyTorch · TensorFlow

Data Science

Model training · **ML ops** · **Inference**

Model building · **Data preparation** · **Fine tuning**

AI Accelerator Packs

NVIDIA NIMs

AI.Q · **Route Optimizer** · VSS

Outcome based Ready-made solutions

Confidential - Oracle Restricted \Employees Only

OCI AI Supported Models for the Enterprise

xAI	OpenAI	Meta	cohere	Google
NEW Grok 4.1	NEW gpt-oss	Llama 3.1	NEW Command A	NEW Gemini 3 pro
Grok 4		Llama 3.2	Command R+	Gemini 3 flash
Grok 4 Fast		Llama 3.3	Command R	Gemini 2.5 pro
Grok Code Fast 1		Llama 4 Scout	NEW Embed 4	Gemini 2.5 flash
Grok 3 models		Llama 4 Maverick	Embed 3	Gemini 2.5 flash-lite
			Rerank 3.5	

BYOM

Confidential - Oracle Restricted \Employees Only

Common use cases

Business Operations
As a business operations lead, I want multi-agent systems to orchestrate complex cross-functional processes (e.g., end-to-end order fulfillment), so that I can achieve proactive automation with minimal oversight.

Compliance Officer
As a compliance officer, I want an agent to monitor real-time data, flag risks, and generate reports with citations, so that I can ensure adherence in high-stakes environments.

DevOps Engineer
As a DevOps engineer, I want AI agents to handle incident response and deployment workflows autonomously, so that I can minimize downtime and improve system reliability.

Financial Analysis
As a financial analyst, I want an agent to combine market data searches, portfolio file analysis, and risk modeling functions, so that I can make informed trading or investment decisions in volatile markets.

HR – Talent Manager
As an HR professional, I want an agent to assess candidates, generate competency frameworks, and handle interview scheduling/guidance, so that I can streamline talent acquisition.

Legal
As a legal professional, I want an agent to search and summarize case files or documents from vector stores, so that I can quickly reference precedents during case preparation.

Market Research
As a market researcher, I want an agent to perform web searches, synthesize data from multiple sources, and cite evidence, so that I can complete weeks of research in minutes.

Confidential - Oracle Restricted \Employees Only

Next steps: turn agentic AI into an enterprise program

Use a bounded workflow to prove value, controls, and operating ownership.

1	**Select one workflow**	High handoff pain, clear owner, visible outcome.
2	**Map controls**	Systems, data, approvals, audit, and risk boundaries.
3	**Run a constrained POC**	Evaluation criteria, tool access, and human checkpoints.
4	**Build production blueprint**	Ownership, monitoring, cost model, runbook, governance.

Pick the first workflow where actions, evidence, and ownership are already clear.

Q&A

Confidential - Oracle Restricted \Employees Only