ORACLE

Sovereignty, Security, and Compliance for AI

From Risk Management to Real Business Value

Derya Sözen Esen

Director, SaaS Security and Privacy

EMEA Apps, Center of Excellence

May 13, 2026

Derya Sözen Esen

Director, SaaS Security and Privacy

Experience

Computer Engineer with 16 years of sectoral experience

Oracle: 9 years

Expertise

Regulatory compliance, security and privacy domains – help steer customers and Oracle through complex regional regulatory landscape.

Academic

PhD research at Goethe University Frankfurt, Germany

Profound interest in AI regulations and auditing.

Agenda

Let's define:
Trustworthy AI

Preserving Security and Privacy:
Oracle AI

Watch out for:
Regulatory Landscape

Achieving True Sovereignty in the Cloud:
Fusion and AI

Check out:
Useful Resources

The opening poll

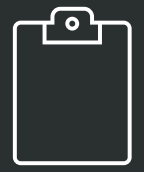

Winning with AI requires control,

not speed alone …

We are part of your
AI Governance !

How to do it responsibly, securely, and compliantly!

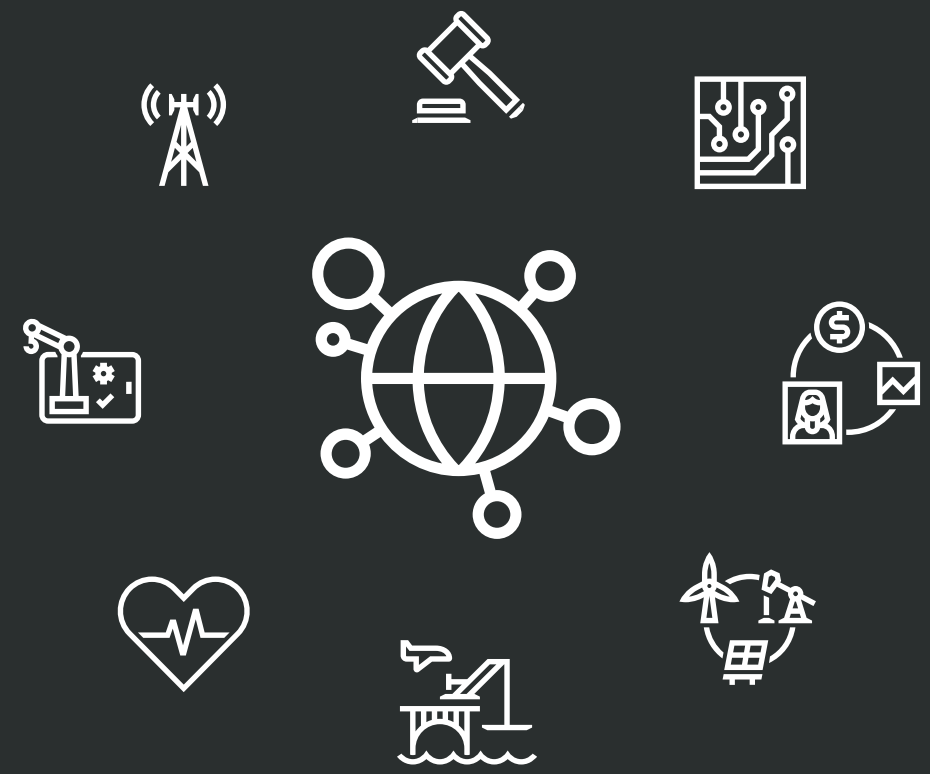

Let’s Define:

Trustworthy AI

Why Trust Is the Real Bottleneck to AI Adoption

AI creates new risk categories:

- Model unpredictability
- Data leakage & privacy violations
- Regulatory non-compliance
- Reputational damage

In regulated industries, risk ≠ theoretical

- One incident can halt deployment entirely

Result:

- Endless pilots
- Shadow AI
- Innovation stalls

Redefining "Trustworthy AI"

Four pillars of AI trust:

Security

- Protecting models, data, prompts, and outputs

Privacy

- Lawful, minimal, purpose-bound data usage

Regulatory Alignment

- Demonstrable compliance, not assumptions

Assurance

- Continuous monitoring, testing, and validation

From Compliance as a Cost → Compliance as an Enabler

Compliance provides:

- Clear boundaries
- Faster approvals
- Confidence to scale

… and it leads to:

Clear risk classification → faster go/no-go decisions

Pre-approved AI controls → teams don't reinvent governance

Audit-ready documentation → leadership confidence

Preserving Security and Privacy

Oracle AI

Oracle AI Cloud Platform → One Oracle Cloud Platform

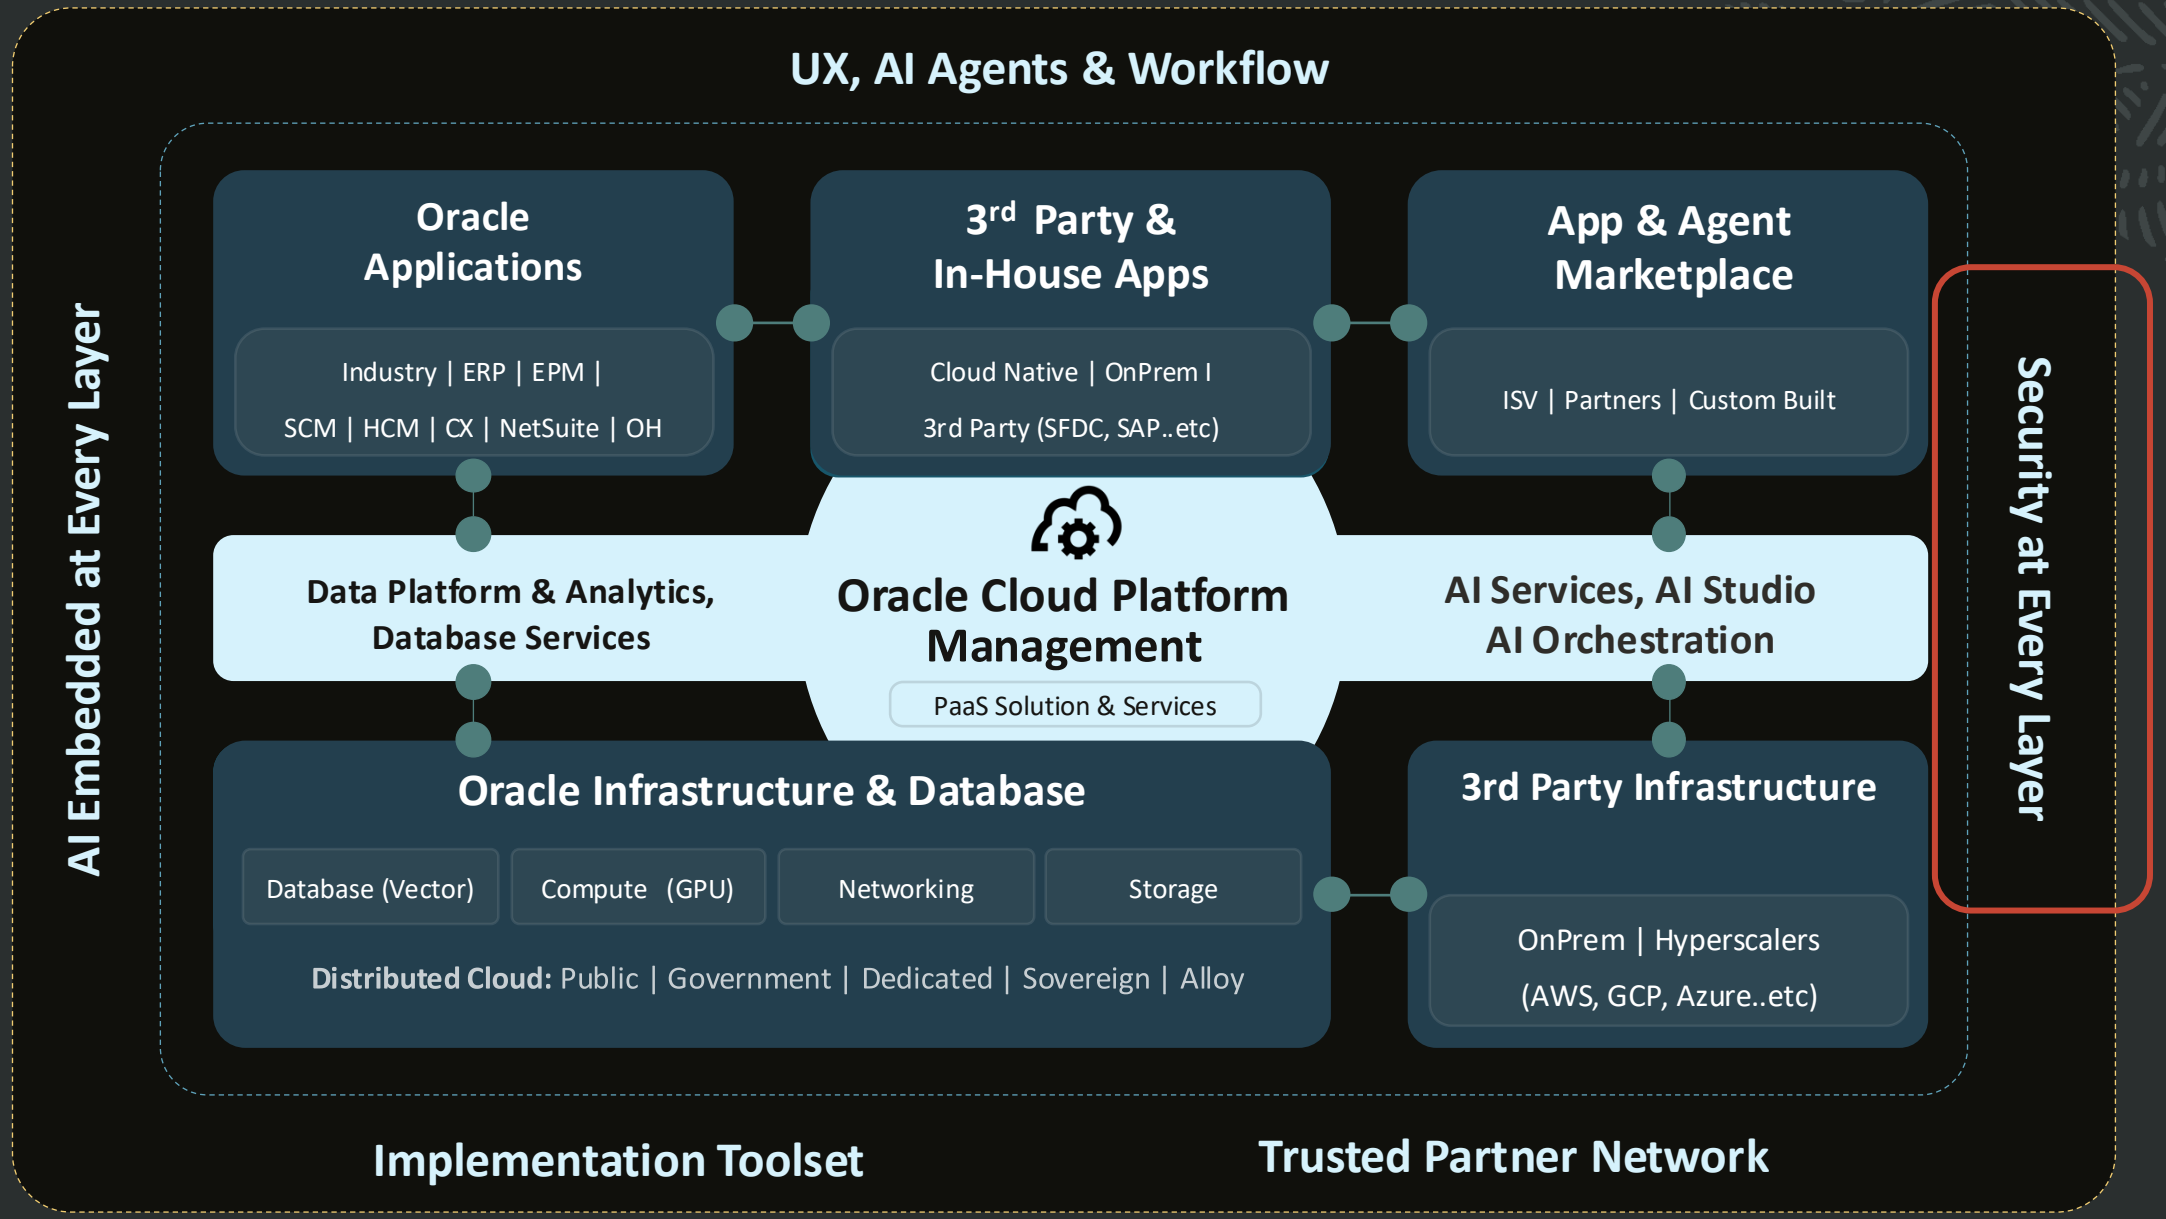

“The models do not get trained on your private data because for some reason, people want to keep their private data private. And that's not going to change.”

By Larry Ellison

Inference: Using a pre-trained AI model to generate insight, predictions, or responses, without retraining the model.

Data is always **secure and is not shared**

Your data and the models we use are hosted and managed all within the same infrastructure (OCI)

None of your data is shared with LLM provider, other customers, or other third parties

Encryption at rest (AES-256) and in transit (TLS).

Same Fusion authentication & authorization model applies.

Users are in the driver seat with **access** and **control** over AI

AI-features are not automatically switched on when delivered but opted into by you

They can also be switched off at any point, for each individual use case

Users remain in control of AI and can augment or override its output

Guardrails protect your business and mitigate risk

We engineer prompts, test results for accuracy and provide outcomes to SaaS users for human approval

Prompts and customer data are not persisted in the generative AI models

Guardrail services part of GenAI, that evaluates the LLM response for compliance to use case-specific requirements.

Prevent prompt injections.

AI is embedded within workflows and augments work

Prompts are engineered and optimized for specific use cases, ensuring accurate and appropriate outcomes for the relevant use case

AI outcomes and responses are presented within applications, ensuring accessibility and ease of use

We monitor legislation and regulatory changes

Oracle is committed to complying with data protection laws and collaborates with data and AI standards organizations

Oracle Responsible AI Development and Deployment

Secure Development Lifecycle

Policies and practices utilized for development and deployment of Oracle AI, are aligned with applicable laws, regulations, and industry standards, including the European Union Artificial Intelligence Act ("EU AI Act"), United States NIST AI Risk Management.

Post-Deployment Monitoring

Oracle monitors deployed AI Systems consistent with our practices and as may be required by applicable law.

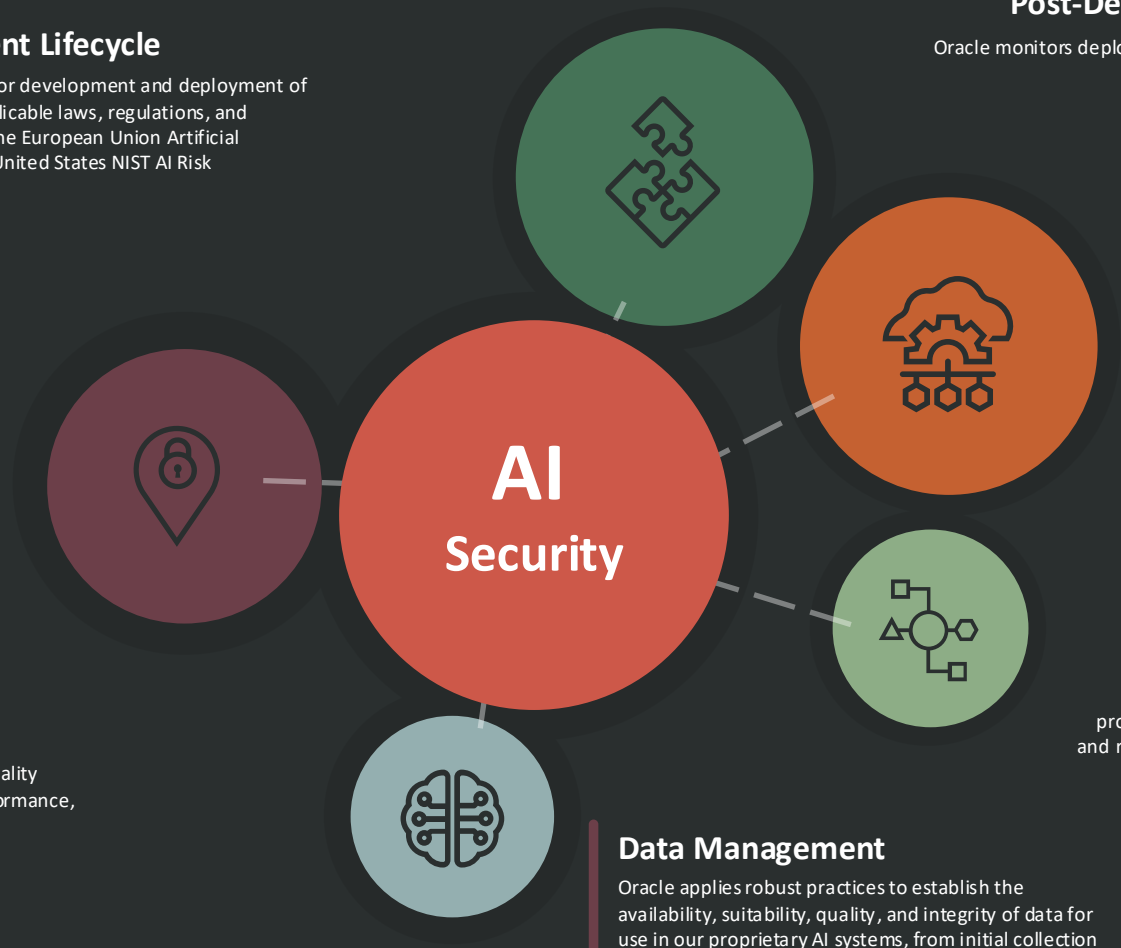

Risk Management

Oracle identifies, analyse, manage, and mitigate AI system risks throughout the development lifecycle.

Technical Documentation

Creating, maintaining, and making documentation available is another core element of our AI governance practices.

Quality Assurance

Oracle AI systems are subject to a rigorous quality assurance process designed to establish performance, reliability, and alignment with intended use.

Compliance with Law

Oracle continually monitor developments around the world to proactively update AI policies and practices to comply with new laws and regulations (and the implementation of existing laws like the EU AI Act), and to remain aligned with emerging or evolving industry standards.

Data Management

Oracle applies robust practices to establish the availability, suitability, quality, and integrity of data for use in our proprietary AI systems, from initial collection to use in training and testing.

AI Threats and Risks

THREAT VECTOR	BUSINESS IMPACT EXAMPLE	STRATEGIC MITIGATION
Data Privacy and Leakage	Staff pasting PII/IP into public models.	Deploy private instances; DLP, Guardrails on prompts.
Data Security	Unencrypted training sets exposed in breaches.	Encryption at rest/transit; rigorous key management.
AI API Security	Unsecured endpoints allowing unauthorized usage.	Strict AuthN/AuthZ; API Gateway rate limiting.
Model Bias and Fairness	Discriminatory hiring or lending decisions.	Mandate fairness audits; diverse training data.
Supply Chain Risk	Compromised third-party models or libraries.	Vulnerability scanning; SBOMs for AI artifacts.
Hallucinations	AI generating confident but false legal facts.	Implement RAG to ground answers in trusted data.
Adversarial Attacks	"Jailbreaking" to bypass safety filters.	Red Teaming exercises; adversarial training.
Intellectual Property	Output infringing on copyrighted material.	Verify data provenance; vendor IP indemnification.
Shadow AI	Unvetted tools bypassing IT security.	Centralized AI Governance; approved Service Catalog.

Just a snapshot of some security and privacy preserving capabilities in Oracle AI

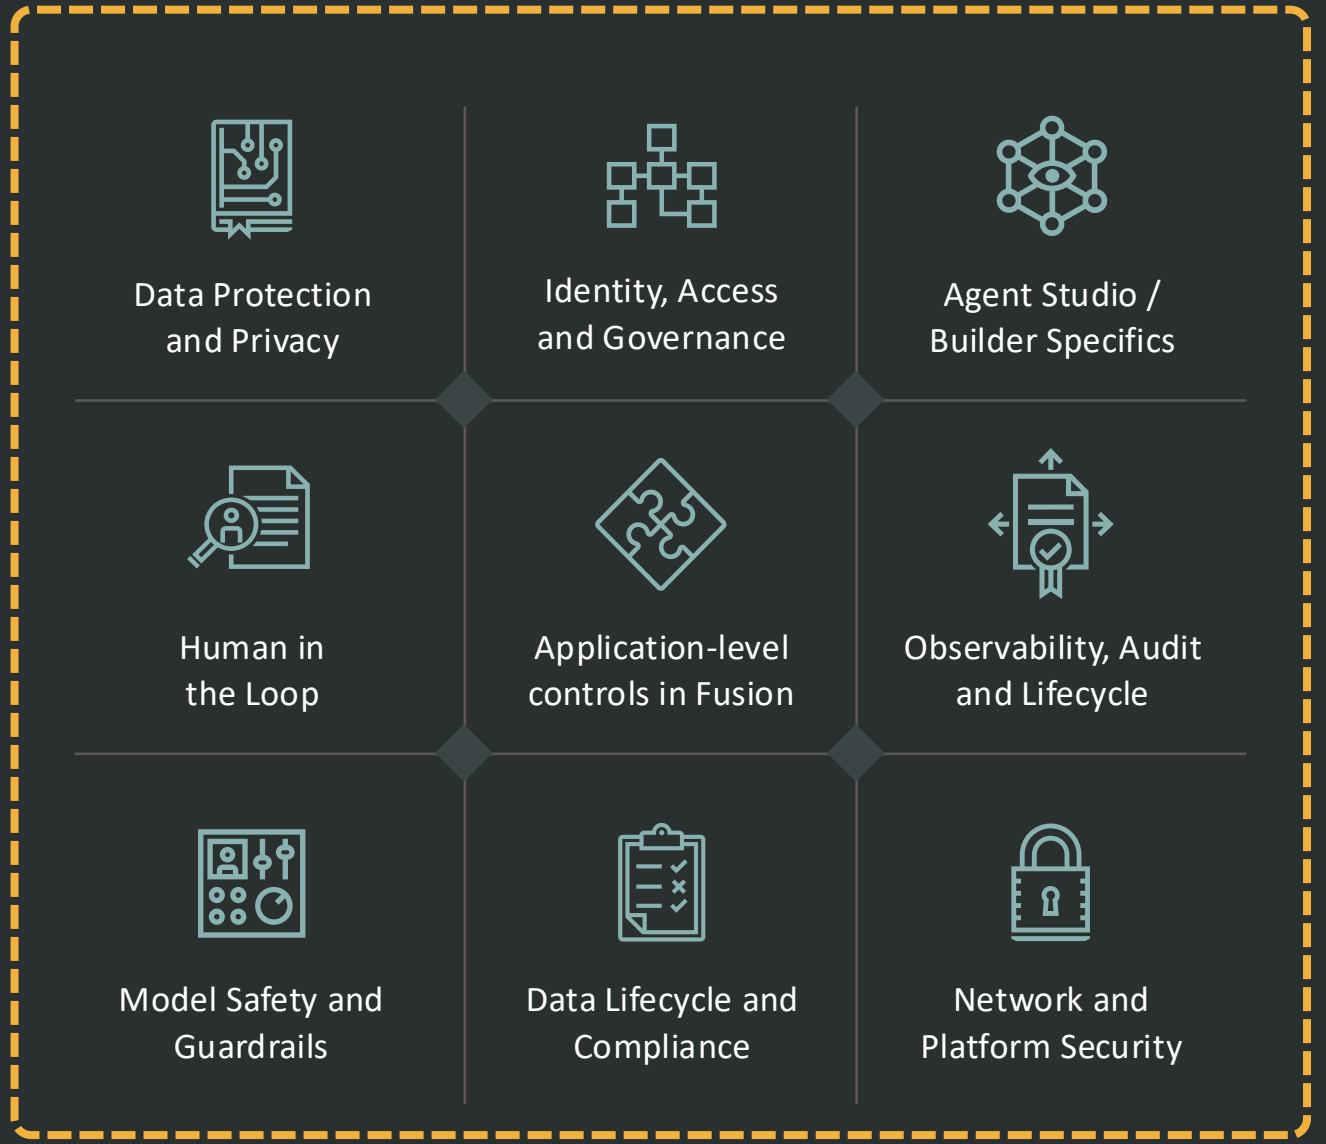

Delivering Trustworthy Agents

Defense-in-depth through a combination of technical, contractual and policy guardrails

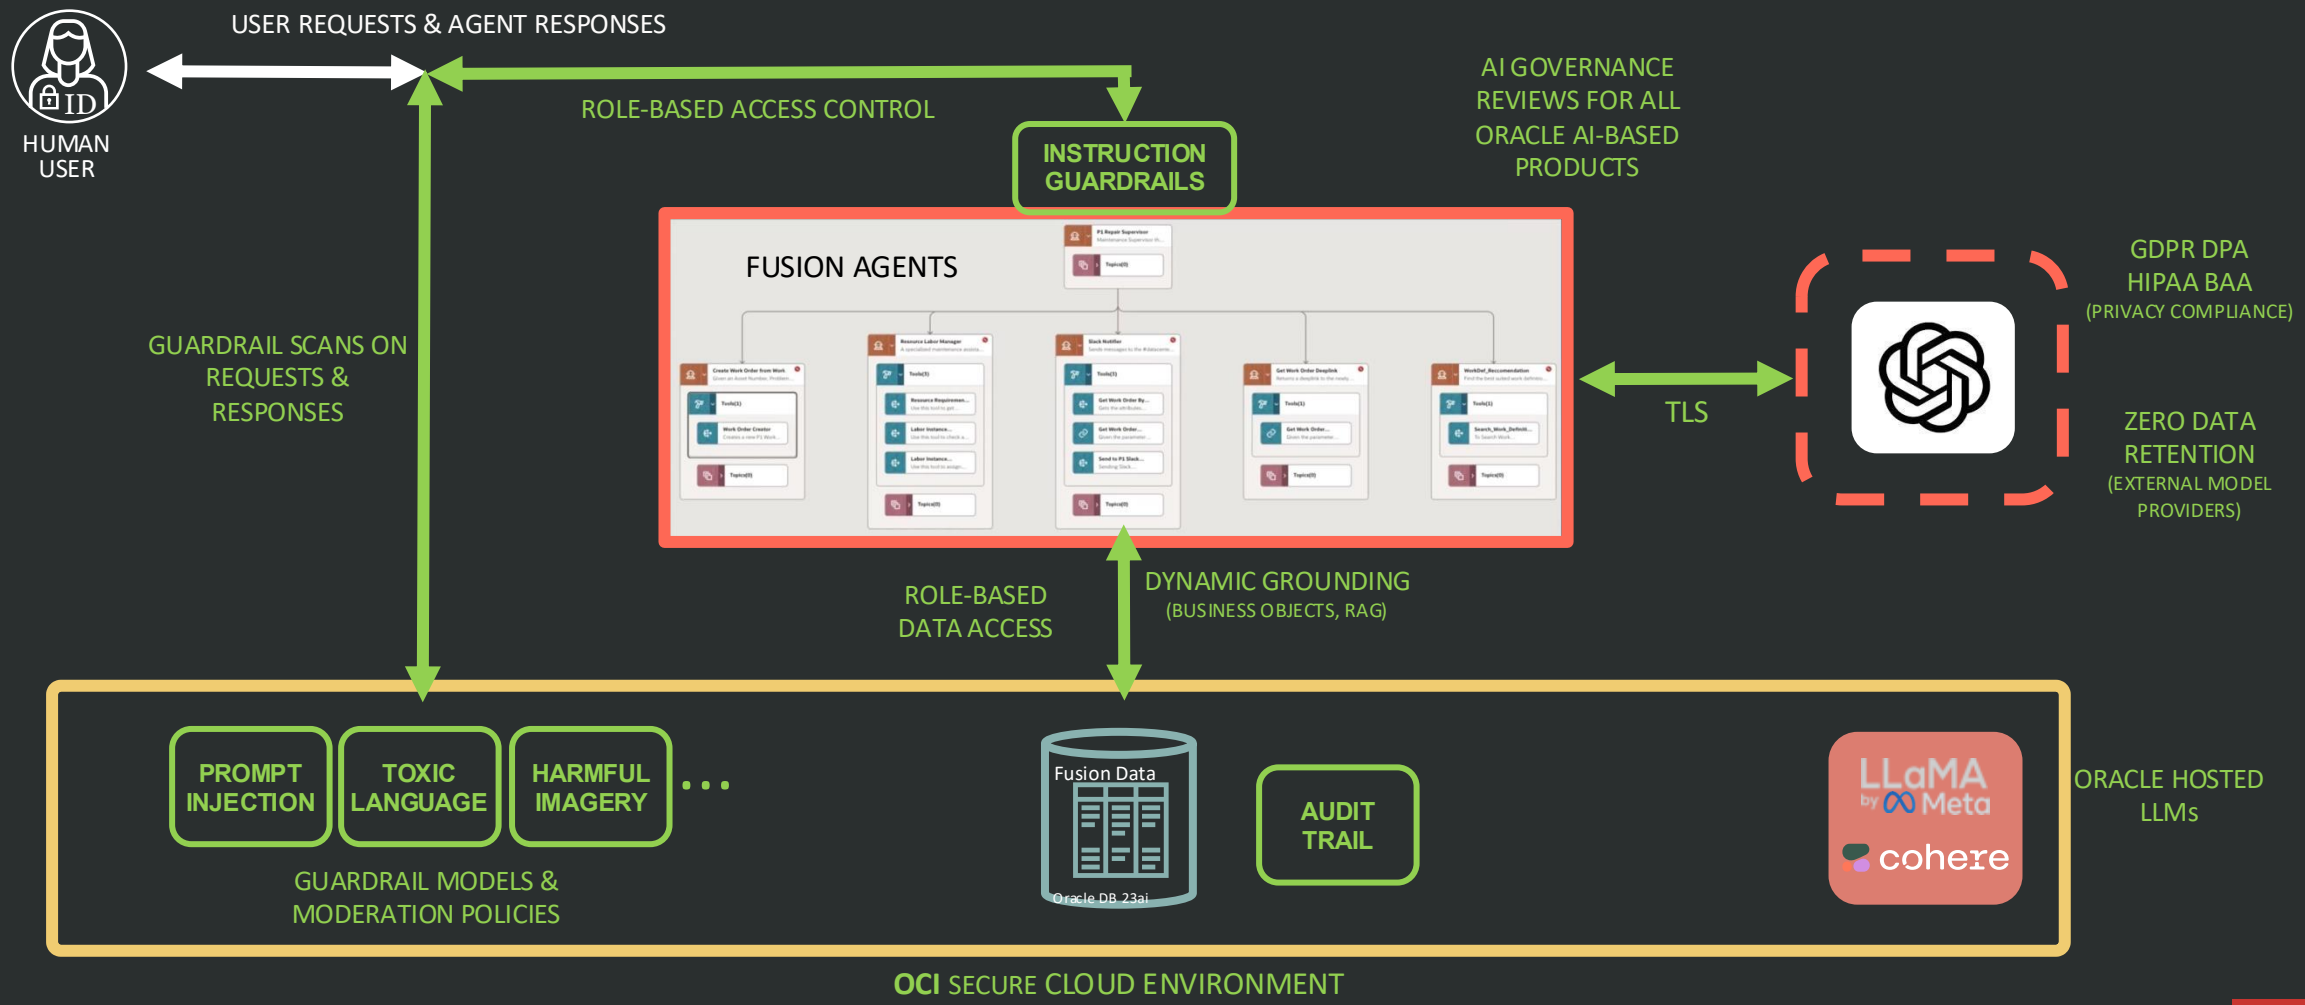

Watch out for:

Regulatory Landscape

Good Compliance is Good Business!

Aug 2024 - AI Act entered into Force (EU): The multi-year compliance countdown officially began.

Jan 2025 - Data (Use & Access) Act Enacted (UK): Reforming data rules to support AI innovation and automated decision-making.

Feb 2, 2025 - Prohibited AI Systems (EU): Bans on "unacceptable risk" AI systems take effect, including social scoring and cognitive manipulation.

Aug 2025 - GPAI Obligations (EU): General Purpose AI models must meet transparency and copyright requirements.

Oct 2025 - AI Growth Lab Launch (UAE): Regulatory sandbox testing in healthcare and transportation sectors.

Jan 2026 - ADM Codes of Practice (UK): New statutory codes for automated decision-making go live.

Feb 2026 - Post-Market Monitoring (EU): Deadline for implementation guidelines for high-risk systems.

Mar 2026 - Copyright & AI Reports (UK): Statutory deadline outlining protections for AI-generated outputs.

Aug 2026 - High-Risk AI Deadline (EU): Rules for Annex III high-risk systems become generally applicable.

Late 2026 - Global AI Hub Law (Saudi Arabia): Expected transition to binding federal AI statute.

Aug 2027 - Product Integration (EU): Compliance deadline for AI used as safety components in regulated products.

Dec 2030 - Public IT Systems (EU): Final deadline for AI components in large-scale public infrastructure.

Trust becomes the passport that allows AI to cross borders.

Trust = demonstrable compliance

Oracle AI System Development Practices are aligned with applicable regulations and industry standards, including without limitation the **EU AI Act, ISO 42001** and the **NIST AI RMF**

https://www.oracle.com/corporate/cloud-compliance

https://cloudsecurityalliance.org/star/registry/oracle-corporation

Artificial Intelligence

Oracle participates in the development of artificial intelligence (AI) standards. The projects in these standards bodies address a variety of aspects of AI systems – bias/fairness, human oversight, explainability/transparency, risk management, privacy, and security. Such standards advance technical harmonization and support market access.

ISO/IEC JTC 1 SC 42	This is a joint body under the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) established AI standardization.
CEN/CENELEC JTC 21	This is a joint body under the European Committee for Standardization (CEN) and the European Committee for Electrotechnical Standardization (CENELEC). Joint Technical Committee 21 develops and adopts standards for AI and related data. This body has also been tasked with establishing harmonized standards (European Norms) for the AI Act.
ETSI ISG SAI	The European Telecommunications Standards Institute ("ETSI") established an Industry Specifications Group on Securing Artificial Intelligence (ISG SAI). This group addresses security of AI and the use of AI to support security.

We are also engaging in the AI activities of:

- The Organisation for Economic Co-operation and Development (OECD), including the Working Party on Artificial Intelligence Governance (AIGO).
- The EU-US Trade and Technology Council (TTC), including Working Group 1 on Technology Standards and its AI subgroup as well as the Joint Roadmap on Trustworthy AI and Risk Management
- The Linux Foundation: https://www.linuxfoundation.org/press/linux-foundation-welcomes-the-agntcy-project-to-standardize-open-multi-agent-system-infrastructure-and-break-down-ai-agent-silos
- Cloud Security Alliance: https://cloudsecurityalliance.org/research/working-groups/ai-controls
- AI Alliance: https://thealliance.ai/
- OECD AI: http://oecd.ai/
- OWASP – Review Board of Agentic Security Initiatives: https://genai.owasp.org/

4. COMPLIANCE WITH AI LAWS

4.1 Oracle shall comply with all laws related to AI to the extent that such laws, by their terms, are expressly applicable to Oracle's provision of AI Functionality and Output (as defined below) and impose obligations directly upon Oracle in its role as an information technology services provider with respect to AI Functionality and Output. You shall comply with all laws related to AI to the extent that such laws, by their terms, are applicable to Your use and receipt of AI Functionality and Output and impose obligations directly upon You with respect to AI Functionality and Output.

4.2 Neither Oracle nor You shall provide or use AI Functionality or Output in a jurisdiction in a manner that is prohibited by applicable law in such jurisdiction (including under Article 5 of the European Union Artificial Intelligence Act).

Achieving True Sovereignty in the Cloud

Oracle AI Stack

Applications

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

AI Services

- GenAI
- GenAI Agents
- Code Assist
- Digital Assistant
- Speech
- Language
- Vision
- Document Understanding

GenAI and ML for Data Platforms

- Database AI Vector Search
- Autonomous Database Select AI
- HeatWave GenAI with Vector Store
- Data Science
- ML in Oracle Database
- Language
- HeatWave AutoML
- Data Labeling

AI Infrastructure

- Supercluster with RDMA networking
- Compute GPU: Bare metal, VMs, Kubernetes
- Block, Object, File Storage, HPC filesystems

Distributed Cloud

- Public
- Dedicated
- Sovereign
- Alloy
- Government

AI Partners and ISVs

Oracle’s sovereign-by-design philosophy

Localization

Isolation

Personnel Requirements

Encryption

Access Management

Data Access Requests

Core Principles

No compromise on security | No compromise on functionality | No compromise on innovation

Sovereign AI for EU and UK

Sovereign Cloud Infrastructure

Unique Value Proposition

Pure sovereign GenAI availability in public cloud.

Facilitate Compliance & Mitigate Disclosure Requests

Simplify compliance with current and upcoming regulations like GDPR, DORA, PS Requirements, etc.

Separated legal entities, encryption and isolation mitigate extraterritorial disclosure requests.

Increase Autonomy and Digital Sovereignty

Ensuring that AI operations and customer data remain within EU/UK borders.

Separate legal entities for owning hardware & data centre leases.

Similar Fusion features as Commercial Cloud

LLMs available: Cohere & Llama

Enterprise AI at every level of the stack to help you create value and drive results in **EU and UK Sovereign Cloud** solutions

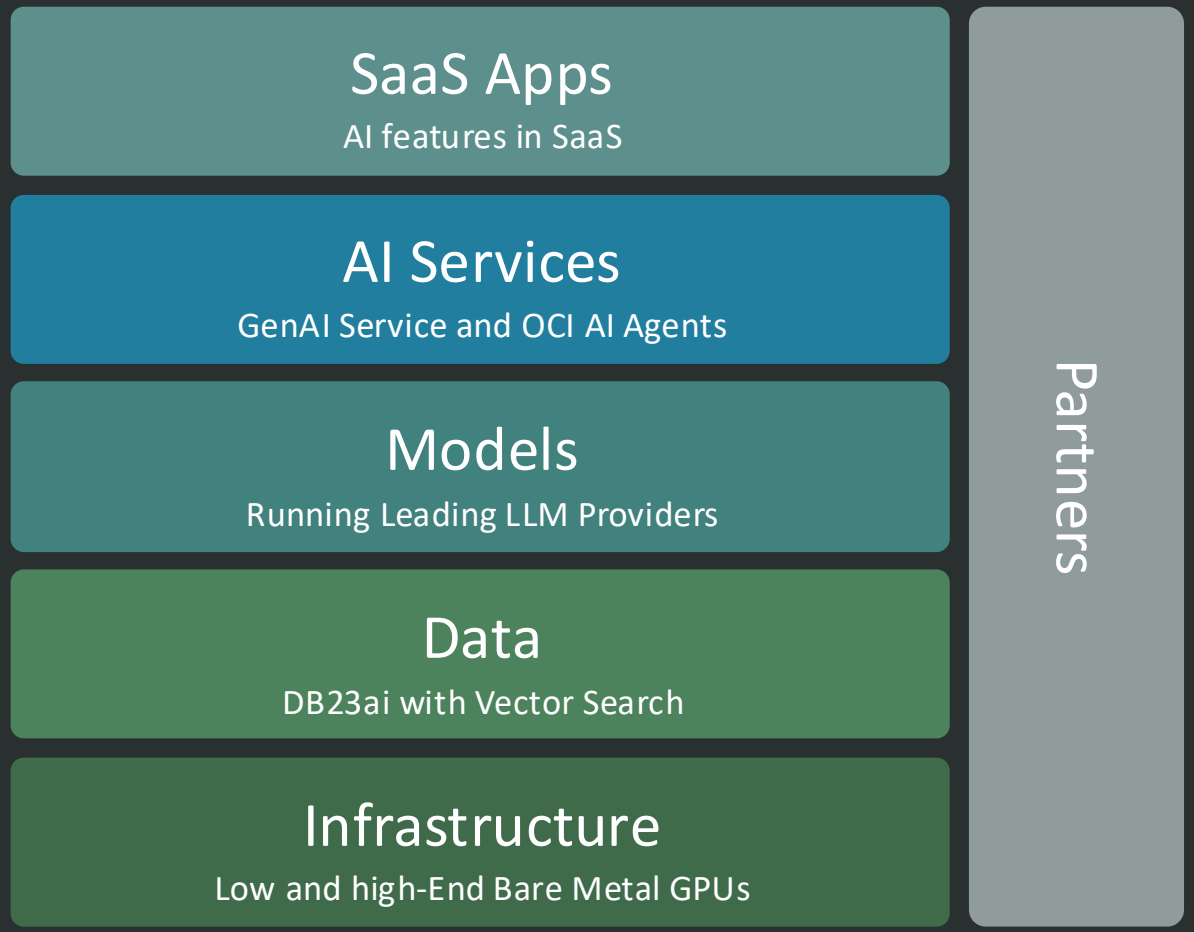

In Summary … Oracle AI as a Key Differentiator

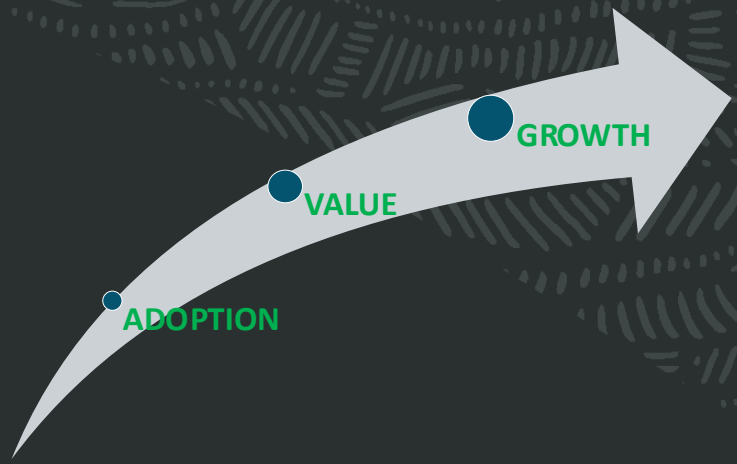

Full stack

- GenAI is delivered on the same cloud platform and under the same security model.
- Available across multiple Oracle Cloud regions and can also be deployed on-premises through the OCI Dedicated Region (DRCC) platform.

Sovereign

- Available in **EU Sovereign Cloud**, Oracle **UK Sovereign Cloud** & DRCCs.
- Enabling public-sector agencies and regulated industries to build and operate AI applications while maintaining full control of their data, operations, and infrastructure.

Natively Embedded

- Oracle AI works directly with transactional data, unlike other vendors that require copying transactional data into a separate AI data cloud.
- By embedding AI within the applications, Oracle ensures that transactional relationships are preserved and that existing security and access controls are consistently enforced.

Check out:

Useful Resources

Oracle AI Useful Resources

- AI Consensus Assessment Initiative Questionnaire for Oracle SaaS Cloud Applications
- Oracle Artificial Intelligence Terms
- AI agent observability and evaluation
- Sovereign AI
- Oracle AI for Fusion Applications Q&A
- Oracle's AI standards participation
- For more, check out AI security blogs on Oracle Blogs

Learn more at Sovereign AI

Unlock real business value
.. without compromise !
AND WIN

Thank you

Any questions?

derya.soezen.esen@oracle.com

What’s next? and our Closing Poll

Oracle AI Live

8-14 June | London

Join us 8-14 June at Future Stores on London's Oxford Street to explore what's next in AI and digital innovation—and what these trends mean for leaders shaping the future of business, government, and society. Each day of the event will focus on a different sector.

Join us for deep conversations on strategies and tactics you can take back to your organization.

AI Changes Everything - Agenda

Plenary Session

09:30 – 10:15

Neil Sholay

VP of AI

Agentic AI in the Enterprise

10:30 – 11:00

Pankaj Sharma

Director AI, EMEA Technology Engineering

AI Agent Demo

11:15 – 11:45

Guillaume Voisin

Principal Solution Engineer

AI Driven Customer Service

11:45 – 12:30

Jose Cruz

Senior Director, Data Strategy & Architecture

Sovereignty, Security, and Compliance for AI

12:30 – 13:00

Derya Soezen Esen

Director, SaaS Security & Privacy

Real Ai for real value in construction & engineering

13:30 – 14:00

Josh Kanner

Senior Director, Product Management

Fast-Track AI Adoption

15:00 – 15:30

Blair Bozada

Senior Product Marketing Manager, AI


ORACLE