Faysal Aziz

AI for Sales Professionals

GenAI, agents, RAG and infrastructure conversations

FA-0566

AI for Leaders & Business · Agentic & Generative AI

Build clear, credible AI conversations with enterprise customers by connecting use cases to infrastructure requirements, risks and evaluation questions.

Introduction

This course helps sales professionals explain generative AI, tool-using agents and retrieval-augmented generation in accessible business language. It connects customer objectives to high-level compute, storage, networking and deployment considerations.

Participants practise discovery questions and value discussions without promising universal hardware requirements, guaranteed accuracy or business results. The course is conceptual, not AI-system implementation or product-specific certification training.

Learning outcomes

The course teaches participants to:

- Explain generative AI, agentic AI and RAG in business language.
- Connect proposed AI use cases to workload and deployment questions.
- Recognise high-level compute, storage, network and data requirements.
- Identify limitations, oversight and responsible-use issues customers may raise.
- Demonstrate clear prompt/context concepts.
- Discuss tools and infrastructure credibly without overpromising.

Prerequisites

Basic familiarity with the hardware products relevant to the participant's role.

- General familiarity with typical enterprise IT needs.
- Comfort with virtual instructor-led learning.
- No prior AI technical background required.

Training outline

1. 1 – AI in business and core concepts

Business context

- Business adoption drivers and the evidence behind selected examples.
- Why organizations invest in AI capabilities
- Relationship between AI initiatives and IT infrastructure
- Questions customers may ask about data handling and compute capabilities.

Foundations

- Overview of traditional artificial intelligence
- Language models alongside other AI approaches; match the approach to the task.
- Key terminology used in contemporary AI discussions

2. 2 – Generative AI, RAG and agents

Generative AI

- Definition and core principles
- Types of generative models
- General categories of business applications
- Infrastructure considerations at a high level

Retrieval-augmented generation

- Concept and purpose
- Role of organizational data in RAG
- High level architecture components
- Potential relevance and grounding benefits, with retrieval-quality, permission and correctness limitations.

Agentic AI

- Agents as goal-oriented model/tool workflows with configured boundaries.
- How agents differ from simple model queries
- Relevance for workflow automation
- Enterprise implications: integration, evaluation, monitoring and human responsibility.

3. 3 — Customer infrastructure needs

- Workload-dependent compute, storage and networking needs; distinguish training from inference and hosted from local execution.
- Differences between on premises, hybrid, and cloud environments
- Assess when server infrastructure supports the proposed workload; cloud-hosted applications may have different local requirements.
- Align options with customer objectives and validated workload assumptions, not generic AI-readiness claims.

4. 4 — Risks, prompting and tool choices

Responsible use

- Data quality issues
- Bias and fairness concerns
- Security and access control topics
- Reliability and oversight requirements
- Common areas of customer hesitation

Prompt and context foundations

- Structure of input prompts
- Context relevance, quality and limits; extra context does not automatically improve results.
- Importance of precise instructions
- How business users typically interact with AI tools

Tool/platform categories

- General categories of widely used AI tools
- Characteristics that customers evaluate when selecting tools

- Interaction between tools and underlying hardware environments

5. 5 — Sales conversations and future planning

Customer discovery and communication

- Approaching customer discussions about AI objectives
- Identifying potential business drivers
- Linking customer goals to infrastructure implications
- Communicating value without technical overload
- Handling common customer questions

Longer-term considerations

- Evolving directions in enterprise AI
- How workload choices, efficiency and scale may affect compute demand; avoid universal growth claims.
- Longer-term capacity and operational planning using documented assumptions and revisited requirements.

© 2026 Faysal Aziz · Kuala Lumpur