
Generative AI for Business Executives Programme

26 – 30 September 2027
Manama





Generative AI for Business Executives Programme

Ref: 1703_5250 **Date:** 26 – 30 September 2027 **Location:** Manama **Fees:** 4300 GBP

Programme Summary

Workplaces face an influx of synthetic media and foundation models operating directly within departmental workflows, requiring structured governance rather than informal adoption. Generative AI for Business Executives addresses the architectural mechanics, deployment pathways, and oversight standards governing foundation models and autonomous systems. Participants evaluate agentic AI architectures, institutional prompt engineering, workflow automation, and operational guardrails to build an actionable capability roadmap. This masterclass is delivered by Mars Training Centre UK.

The Business Case

Uncoordinated deployment of foundation models exposes enterprises to proprietary data leakage, intellectual property infringement, hallucinated operational guidance, and unverified algorithmic decisions. Establishing systematic operational controls converts fragmented experimentation into audited workplace productivity integration while securing institutional information assets.

- **Productivity Velocity:** Accelerates information synthesis, code synthesis, and structured documentation across commercial operations.
- **Intellectual Property Protection:** Prevents unapproved exposure of confidential records to external public training corpora.
- **Operational Output Reliability:** Reduces hallucination rates through retrieval-augmented generation and systematic prompt engineering standards.
- **Policy Alignment:** Replaces informal employee adoption with enforceable acceptable-use policies and transparent auditing frameworks.

Principles and Evidence Base

- **ISO/IEC 42001:2023 Artificial Intelligence Management System:** Establishes enterprise requirements for establishing, implementing, and maintaining governed AI systems.
- **NIST AI Risk Management Framework:** Guides structured protocols for governing, mapping, measuring, and managing foundation model exposures.
- **Enterprise Architecture Framework for Generative Workflows:** Informs retrieval-augmented generation architectures, model orchestration layers, and agentic workflows.



Intended Participants

- Departmental heads managing the deployment of language models and autonomous agents within core functional pipelines.
- Operational leaders establishing enterprise guidelines for workplace productivity integration and synthetic content generation.
- Corporate counsel, information security managers, and compliance leads formulating acceptable-use policies.
- Enterprise architects and transformation specialists evaluating generative model APIs, commercial tools, and agentic workflows.

The curriculum supports professionals overseeing knowledge-intensive operations, customer engagement hubs, software delivery teams, and administrative workflows.

Learning Outcomes

1. Participants will be able to evaluate foundation models and agentic AI architectures using technical and operational criteria.
2. Participants will be able to construct structured prompt engineering protocols to extract deterministic outputs from conversational models.
3. Participants will be able to design workplace productivity integration pathways that link proprietary enterprise datasets to generation engines.
4. Participants will be able to draft institutional acceptable-use policies addressing data security, intellectual property, and user accountability.
5. Participants will be able to assess vendor roadmaps, model fine-tuning requirements, and infrastructure costs for departmental workflows.
6. Participants will be able to establish supervisory audit loops and monitoring mechanisms for autonomous task-planning agents.

Programme Structure

Day 1: Foundation Models and Technical Mechanics

How do large language models, transformer mechanisms, and synthetic media engines process unstructured enterprise data?

- Transformer Architecture Fundamentals and Attention Mechanisms
- Tokenisation Parameters, Context Windows, and Compute Scaling
- Pre-training Foundations, Fine-Tuning Pathways, and Distillation
- Modalities Across Text, Code, Tabular Data, and Multimodal Synthesis
- Commercial Application Programming Interfaces Versus Open-Weights Model Hosting



Day 2: Institutional Prompt Engineering and Retrieval Architectures

How do operational teams configure deterministic prompt logic and integrate private institutional knowledge?

- Structured Prompt Engineering Patterns and Zero-Shot and Few-Shot In-Context Learning
- System Prompt Guardrails and Contextual Output Formatting
- Retrieval-Augmented Generation Architecture and Vector Embedding Databases
- Semantic Search Indexes and Enterprise Knowledge Grounding
- Evaluation Harnesses for Hallucination Detection and Factual Accuracy

Day 3: Agentic AI Architectures and Workflow Automation

How do autonomous agents chain actions, invoke corporate tooling, and execute multi-step operational tasks?

- Agentic AI Architectures and Autonomous Goal Decomposition
- ReAct Reasoning Frameworks and External Tool Calling Protocols
- Multi-Agent Coordination and Task Handoff Workflows
- Human-in-the-Loop Approval Checkpoints and Safety Escalation Gates
- Autonomous Code Generation and Enterprise Process Orchestration

Day 4: Governance, Acceptable-Use Policies, and Risk Controls

How are algorithmic exposures, privacy requirements, and intellectual property liabilities systematically governed?

- ISO/IEC 42001:2023 Controls for Foundation Model Deployments
- Acceptable-Use Policies for Workforce Interaction and Code Generation
- Data Protection, Zero-Retention Agreements, and Data Leakage Prevention
- Intellectual Property Provenance, Copyright Exposure, and Synthetic Media Watermarking
- Red-Teaming Routines, Jailbreak Defences, and Adversarial Vulnerabilities

Day 5: Workplace Productivity Integration and Implementation Roadmaps

How are generative initiatives prioritised, budgeted, measured, and scaled across operational departments?

- Workplace Productivity Integration Playbooks Across Functional Operations
- Total Cost of Ownership Calculation: Tokens, Compute, Licensing, and Staff Enablement
- Vendor Evaluation Matrix and API Service Level Agreement Auditing
- Capability Upskilling Frameworks and Workforce Transformation Metrics
- Capstone Enterprise Generative Architecture and Governance Blueprint



Methodology

The course employs an applied framework analysis model combining technical demonstrations, policy structuring laboratories, and operational scenario testing. Participants interrogate architectural patterns and formulate enforceable policy documentation.

- Suggested activity: Designing and stress-testing a structured prompt engineering template to extract structured data from unstructured enterprise documents.
- Suggested activity: Drafting a departmental acceptable-use policy governing employee interaction with public and private language models.
- Suggested activity: Mapping an agentic workflow architecture that incorporates human-in-the-loop controls for administrative approvals.

Capstone Output

Each participant produces an enterprise operational blueprint consisting of an acceptable-use policy draft, a retrieval-augmented generation deployment assessment, and a capability implementation roadmap. This documentation serves as a direct operational specification for institutionalising generative systems within sponsoring departments.

Frequently Asked Questions

What technical infrastructure is required to participate?

The sessions explore conceptual mechanics, system architectures, and operational frameworks through standard cloud interfaces; no specialised hardware or local software installation is required.

How does this programme differ from general technical coding courses?

The syllabus concentrates on architectural evaluation, operational integration, and governance controls, targeting decision-makers directing capability rollouts rather than software engineers training custom weights.

Are the prompt engineering techniques applicable across different model providers?

Yes. The prompt engineering patterns and retrieval-augmented generation concepts taught reflect model-agnostic principles applicable to proprietary commercial services and self-hosted models.

Can the policy frameworks be integrated into existing information governance systems?

Yes. All acceptable-use policies, risk mitigation checklists, and governance structures follow international benchmarks such as ISO/IEC 42001:2023 and NIST guidelines for direct corporate adoption.