Generative AI Bootcamp

Beginner Level | 1-Day Program

Core Design Principle

The bootcamp should feel like building one real product over two days. Every concept taught should connect back to a single evolving example that participants can take home as a working artifact.

Program Overview

The program spans 6 modules for one day. The table below summarises the complete flow with recommended durations and delivery methods.

Day	#	Module Title	Level	Duration	Delivery
1	1	Introduction to Generative AI & LLMs	Beginner	45 min	Concept + Demo
1	2	Prompt Engineering Essentials	Beginner	60 min	Hands-On
1	3	From Chatbots to Agents	Beginner	45 min	Concept + Demo
1	4	Building & Connecting Agents with Zapier	Beginner	60 min	Hands-On Build
1	5	Knowledge Agents & RAG: Introduction	Beginner	45 min	Trainer Demo
1	6	Responsible AI & Day 1 Wrap-Up	Beginner	45 min	Discussion + Quiz

DAY 1

BEGINNER WORKSHOP | 5 HOURS | Microsoft Copilot Studio + Zapier

Module 1

Introduction to Generative AI & LLMs

BEGINNER — DAY 1
45 minutes

What This Module Must Achieve

Before any hands-on work, participants need a working mental model of what they are actually using. This module should leave every participant (regardless of technical background) with a clear, non-mystified understanding of how LLMs work and why they sometimes fail. The tone should be conversational and demystifying, not academic.

Topics to Cover

- What Generative AI is and what it is not. Distinguish from traditional ML (classification, prediction) vs. generation (creating new content)
- How LLMs work at a conceptual level: tokenisation, next-token prediction, temperature, and context windows. No equations needed. Use plain language and a whiteboard-style visual
- The training process in brief: pre-training on internet-scale data, instruction fine-tuning, RLHF. Participants should understand why LLMs know so much and why they hallucinate.
- Open vs. proprietary models: what the difference means in practice (privacy, cost, customisability, performance)

Expected Learning Outcomes

- Participants can explain in plain English what an LLM is and how it generates output
- Participants understand the concept of hallucination and why it is a structural property of LLMs, not a bug
- Participants can name at least three real business applications of GenAI and articulate the risk in each
- Participants feel oriented and curious, not overwhelmed

Suggested Tweak / Recommendation

The existing curriculum lists this as a standalone module with a demo but no hands-on element. We recommend the trainer spend the final 10 minutes having each participant type their first prompt into Copilot and note the quality of the output. This creates an immediate emotional investment in the 'why does this work?' question that the rest of Day 1 answers.

Module 2

Prompt Engineering Essentials

BEGINNER — DAY 1
60 minutes

What This Module Must Achieve

Prompt engineering is the first skill participants can immediately apply at work after Day 1. This module should be practical and iterative. Participants should leave with a reusable template and the confidence to improve any bad prompt they encounter.

Topics to Cover

- Anatomy of a good prompt: Role, Context, Task, Format, Constraints. Walk through each component with examples from the HR domain
- System prompts vs. user prompts: why the distinction matters for building products vs. one-off queries
- Zero-shot, few-shot, and chain-of-thought prompting: when and why to use each
- Prompt anti-patterns: vague instructions, missing context, conflicting constraints, over-reliance on the model's judgment
- Iterative refinement. Treat a prompt like code: version it, test it, improve it.
- Hands-on exercise: participants receive three poorly written HR prompts and must rewrite them using the template. Trainer reviews live

Expected Learning Outcomes

- Participants can write a structured prompt using the Role-Context-Task-Format-Constraints template
- Participants can diagnose why a given prompt produces poor output and articulate specific improvements

Suggested Tweak / Recommendation

The existing curriculum names this module 'Prompt Engineering Essentials' but the hands-on component, 'Crafting prompts for business scenarios', is underspecified. We recommend providing three written prompt challenges with evaluation criteria, so participants know what 'good' looks like.

Recommended addition: include one prompt injection example so participants understand the security implications of prompts in deployed systems. This takes five minutes and meaningfully raises awareness.

Module 3

From Chatbots to Agents

BEGINNER — DAY 1
45 minutes

What This Module Must Achieve

Participants need to understand why a chatbot and an agent are fundamentally different before they build one. Participants should be able to draw the agent architecture on a whiteboard.

Topics to Cover

- The evolution from rule-based chatbots to conversational AI to autonomous agents: what changed and why
- The three pillars of an agent: Brain (the LLM that reasons), Memory (short-term context and long-term knowledge), Tools (APIs and external systems the agent can call)
- The agent loop: Perceive → Plan → Act → Observe → Repeat. Walk through this with a concrete HR scenario
- When to use an agent vs. a simple LLM call: cost, latency, and complexity tradeoffs
- Demo: trainer shows HireIQ attempting to answer Sarah's question about her probation end date: first as a chatbot (guesses), then as an agent that calls the HR system (correct, personalised answer)

Expected Learning Outcomes

- Participants can articulate the difference between a chatbot and an agent to a non-technical stakeholder
- Participants can describe the agent loop and map it to a real business task
- Participants can identify the Brain, Memory, and Tools components in any given agent architecture

Suggested Tweak / Recommendation

The existing curriculum covers this well conceptually.

Module 4

Building & Connecting Agents with Zapier

BEGINNER — DAY 1
60 minutes

What This Module Must Achieve

This is the primary hands-on module of Day 1. Participants should leave having built and tested a real, running automation. The experience should feel tangible, not a simulated exercise. The Zapier integration should work. If it does not, the session fails.

Topics to Cover

- Defining an agent persona: name, role, tone, permitted scope, escalation rules: documented as a specification, not just a system prompt
- Task decomposition: taking a multi-step business workflow and breaking it into the discrete actions an agent must perform in sequence
- Introduction to Zapier: triggers, actions, Zaps, and the AI step that connects to an LLM
- Connecting Copilot Studio to Microsoft Outlook and Microsoft Teams via Zapier: walking through authentication, mapping fields, and testing
- Hands-on build: participants construct a Daily HR Digest: a scheduled Zap that pulls open onboarding tasks from SharePoint, summarises them with an LLM, and emails the HR manager each morning

Expected Learning Outcomes

- Participants have a live, tested Zapier automation that runs the Daily HR Digest
- Participants can write an agent persona specification document
- Participants understand the architectural pattern: Trigger → LLM reasoning → Action → Output
- Participants can connect Copilot Studio to at least one external Microsoft application

Suggested Tweak / Recommendation

The existing curriculum refers to this as 'Hands-on: Scheduled mail on daily task', which is accurate but undersells the learning.

Critical operational requirement: Zapier trial accounts must be set up before the session begins, not during it. Microsoft should include Zapier account creation in the pre-work instructions distributed before Day 1.

Module 5

Knowledge Agents & RAG: Introduction

BEGINNER — DAY 1
45 minutes

What This Module Must Achieve

Participants should leave understanding why hallucination happens and how RAG addresses it, conceptually and through a live demonstration in Copilot Studio. They do not need to implement RAG themselves at this stage; that happens in the intermediate version of the bootcamp and requires Python. But they must leave with a clear mental model.

Topics to Cover

- The hallucination problem revisited: why LLMs generate confident but false information, and why fine-tuning alone is not the right solution for dynamic, company-specific knowledge
- Two approaches: fine-tuning (baking knowledge into model weights) vs. RAG (retrieving knowledge at query time), and why RAG is preferred for most business use cases
- The RAG pipeline at a conceptual level: Ingest documents → Chunk → Embed → Store in vector database → Retrieve relevant chunks → Augment the prompt → Generate grounded response
- What a vector database is and why it is needed, explained without mathematics using a library analogy
- Trainer demo: upload the Contoso HR Policy PDF into Copilot Studio and watch HireIQ answer Sarah's benefits question with a cited, accurate response

Expected Learning Outcomes

- Participants can explain RAG to a business stakeholder using plain language
- Participants understand why RAG is preferred over fine-tuning for knowledge that changes regularly
- Participants can identify which documents in their own organisation would improve an agent's RAG knowledge base

Suggested Tweak / Recommendation

The existing curriculum marks this as a 'Trainer demo', which is appropriate for Day 1. However, participants should be invited to suggest which documents HireIQ should be given. This five-minute participatory exercise makes the concept concrete and personal, and it sets up the Day 2 discussion about document quality and chunking strategy.

Module 6

Responsible AI & Day 1 Wrap-Up

BEGINNER — DAY 1
45 minutes

What This Module Must Achieve

Responsible AI is not a compliance checkbox, it is a design practice. This module should feel practical, not preachy. Participants should leave with specific design decisions they would make differently because of what they learned here.

Topics to Cover

- Microsoft's six Responsible AI principles and what they mean in practice: Fairness, Reliability and Safety, Privacy and Security, Inclusiveness, Transparency, Accountability
- Real-world AI failures and their root causes like selected examples from HR, healthcare, and finance that resonate with a business audience
- Human-in-the-loop design: when agents must escalate to a human, how to design that escalation gracefully, and how to avoid the failure mode of agents that attempt to answer everything
- Governance requirements for deployed AI: audit logs, rate limits, content filters, data retention policies
- Prompt injection: a brief, practical explanation of how adversarial users can hijack an agent and what mitigations exist
- Day 1 recap quiz: 15 questions covering key concepts from all five prior modules. Group discussion on two or three questions where participants disagree

Expected Learning Outcomes

- Participants can name and explain the six Microsoft Responsible AI principles without referring to notes
- Participants can design a human-in-the-loop escalation pathway for a given scenario
- Participants score at least 70% on the Day 1 recap quiz

Suggested Tweak / Recommendation

The existing curriculum covers this well conceptually.