
AI Business Solution Professional – 120 Hours

Course Overview

The AI Leader Program is designed for professionals who want to bridge the gap between business strategy and AI implementation. While many organizations struggle to translate AI potential into real business value, this program equips participants with the knowledge and practical skills needed to identify AI opportunities, define business requirements, lead AI initiatives, and collaborate effectively with technical teams.

Unlike traditional AI programs that focus solely on technology, the AI Leader Program combines business discovery, AI strategy, project definition, and hands-on technical understanding.

Participants learn how to evaluate AI use cases, conduct stakeholder workshops, build business cases, define Statements of Work (SOWs), assess data readiness, and oversee the development of AI-powered solutions using Large Language Models (LLMs), Retrieval-Augmented Generation (RAG), AI Agents, and modern AI Platforms.

Throughout the program, participants gain a comprehensive understanding of the entire AI project lifecycle—from identifying an opportunity and securing stakeholder buy-in to designing, building, validating, and presenting an AI solution that delivers measurable business impact.

The program culminates in a capstone project where participants design and build an end-to-end AI solution based on a real-world business scenario, demonstrating their ability to lead AI initiatives from concept to implementation.

Target Audience

- **BI Professionals, Data Analysts, and Data Specialists** seeking to transition into AI leadership and AI-driven business roles.
- **Software Developers and IT Professionals** who want to complement their technical expertise with strong business and strategic AI capabilities.
- **Technical Product Managers and AI Pre-Sales Professionals** looking to lead AI initiatives and effectively bridge business requirements with technical solutions.

Learning Outcomes

- Identify AI opportunities within an organization and create a comprehensive Statement of Work (SOW), including project scope, KPIs, success metrics, and milestones.
- Lead discovery processes with business stakeholders and translate business requirements into clear technical specifications.
- Develop Python applications for document processing, information extraction, interaction with Large Language Models (LLMs), and response handling.

-
- Design and implement end-to-end Knowledge Bases using Retrieval-Augmented Generation (RAG) and LangChain on real-world organizational documents.
 - Build AI Agents with custom tools and integrations using FastMCP.
 - Integrate AI components into a unified AI Platform that can be demonstrated and delivered to business stakeholders.
 - Validate the feasibility and business value of AI initiatives through the development, deployment, and presentation of a complete pilot solution.

Course Syllabus:

Part 1: Business & Discovery

Module 1 — AI Foundations

Purpose: Build an accurate mental model of the AI landscape — what is possible, what is not, and where this role sits.

Topics

- The AI landscape: traditional ML vs. Generative AI vs. Agentic AI
- How LLMs work — tokens, context window, inference
- Capabilities and limits: hallucination, knowledge cutoff, reasoning limits
- The cost / latency / quality triangle
- The AI Leader role — responsibilities, interfaces, and boundaries with Data Science and IT
- The AI market: providers, models, open vs. closed

Module 2: AI Strategy & Use Case Prioritization

Purpose: Tools for prioritizing AI opportunities and building a business case that reaches management.

Topics

- Prioritization frameworks: value vs. feasibility, effort vs. impact
- Building a business case: ROI, TCO, build vs. buy
- Organizational AI maturity models
- Change management and AI adoption across the organization
- Governance and risk: bias, security, privacy, IP
- Communicating AI to leadership — framing and managing expectations

Module 3: Business Discovery & Workflow Definition

Purpose: How to interview stakeholders, map processes, and define an AI workflow in the field.

Topics

- Stakeholder interview techniques — the right questions, surfacing pain points
- Process mapping: identifying automation and augmentation candidates
- Distinguishing good AI use cases from poor ones
- Requirements gathering for AI projects — how it differs from traditional IT
- Workflow design: task decomposition, human-in-the-loop
- Defining success metrics and acceptance criteria

Module 4: Project Definition & Statement of Work (SOW)

Purpose: What must be written before anyone starts building.

Topics

- SOW structure for an AI project: scope, deliverables, assumptions, exclusions
- Defining measurable KPIs and success metrics
- Defining MVP vs. pilot vs. production
- Milestones and decision points in an AI project
- Managing uncertainty and experimentation in AI projects
- Cost structure: tokens, compute, development, maintenance
- Interfacing with technical teams — a shared language

Module 5: Data Readiness for AI

Purpose: Assessing data readiness — the most underrated stage, and often the cause of failure.

Topics

- Data readiness assessment: is the data present, accessible, and good enough?
- Working with data and IT teams — asking the right questions
- Types of data sources: files, DBs, APIs, unstructured documents
- Data quality, lineage, and governance
- Embeddings and vector databases — concept and when to use them
- Preparing data for RAG: chunking, indexing, metadata
- Privacy, security, and access to sensitive data

Module 6: Discovery to SOW Workshop

Purpose: A capstone exercise for Part A — from an undefined business use case to a complete, presented SOW.

Topics

- Consolidating and weaving together all Part A tools
- Structuring a business-case presentation to a panel
- Preparing for the technical part: what the SOW tells us to build
- Hands-on practice

Part 2: Technical Build

Module 7 — LLMs & Prompt Engineering

Purpose: From understanding LLMs in theory to writing Python code that calls a model and processes its responses.

Topics

- High-level LLM architecture — tokens, context, temperature
- Model types: closed vs. open, and the major providers
- Working with an LLM API in Python: request, parameters, error handling
- Prompt design: instruction, context, examples, output format
- Structured outputs and JSON mode
- Cost management and token budgeting
- Evaluating response quality: building test sets

Module 8 — Knowledge Bases & RAG Implementation

Purpose: Build Agent 1 — Policy Assistant: an agent that answers questions from organizational documents.

Topics

- Why AI systems need a Knowledge Base
- Extracting text from PDF and Word with a dedicated Python library
- Text processing: chunking strategies and best practices
- Embedding models — what they are and how to use them
- Vector databases — concept and implementation
- Building a RAG pipeline with LangChain
- Connecting the KB to the LLM for document-grounded answers
- Reducing hallucinations and grounding responses

Module 9 — Tools & MCP Integration

Purpose: Build Agent 2 — Data Analyst Agent: an agent that writes SQL, runs it, and returns an answer.

Topics

- What MCP (Model Context Protocol) is and why it exists
- Installing and configuring a FastMCP environment
- Writing a simple tool in Python with FastMCP
- How an LLM decides when to call a tool and what to send
- The agent loop: plan → tool call → observe → respond
- Agent-as-a-tool: using one agent as a tool of another
- Connecting to the AI Platform: exposing an MCP server to the interface

Module 10: Agentic Workflows & Multi-Agent Systems

Purpose: Build Agent 3 — Intake & Routing Agent: an agent with several tools that classifies and routes requests.

Topics

- The difference between a bot and an agent
- Agent architectures: ReAct, Plan-and-Execute
- Memory and state management in agents
- Multi-agent systems: orchestration and specialization
- Reliability: guardrails, error recovery, human approval
- Testing and debugging agents: tracing, observing decisions
- Security: prompt injection, tool misuse, scope limitation
- Code review of agents — common pitfalls

Module 11 — AI Flow Builders & Automation

Purpose: Connect all the agents into one complete business workflow that can be run and presented.

Topics

- AI data flows and pipelines — AI-oriented ETL
- Designing conversational workflows (Chat-as-a-Flow)
- AI Platform: building a visual flow, connecting nodes
- AI Platform: automating business processes
- Connecting multiple agents in a single flow
- Error handling and fallback paths in a flow
- Automating business processes with AI

Module 12 — Analysis, Debugging & Optimization

Purpose: How to make sure the solution works well — before presenting it to the client.

Topics

- Testing and validation of AI systems
- Identifying bottlenecks — latency, cost, quality
- Improving accuracy and response quality
- Performance optimization: caching, batching
- Monitoring: logging, tracing, dashboards
- An iterative improvement process
- When the solution is ready for a pilot

Module 13 — Pilot Deployment & Stakeholder Presentation

Purpose: The move from POC to pilot — how to present a solution to a business stakeholder and get approval to continue.

Topics

- Deployment patterns: from POC to pilot to production
- Managing stakeholder expectations at the pilot stage
- Defining success criteria for a pilot
- Feedback loops: collection, analysis, and iteration
- Preparing a demo for a non-technical audience
- Governance in a pilot: who is accountable, what is measured
- What happens after the pilot — a continuation plan

Module 14 — Final Project — End-to-End AI Solution

Purpose: Student takes the SOW written in Part A and builds the solution it defined in Part B.

Deliverables

- A complete, updated SOW — based on Module 4 and refined after Part B
- Knowledge Base — relevant documents for the use case, indexed
- A working agent — at least 2 functional tools, connected to the KB
- AI Platform — a flow that presents the use case end-to-end
- Demo — a presentation to a panel: business context, live demo, Q&A