



// PROGRAM · AGENTIC_AI · 6_MONTHS

Build agents that *plan, reason,* and *act.*



A 6-month, mentor-led, code-first program in LLM reasoning, multi-agent orchestration, and production deployment — for engineers ready to ship autonomous systems.

<code>duration = 6mo</code>	<code>modules = 12</code>	<code>projects ≥ 10</code>	<code>capstone.deployed = true</code>
-----------------------------	---------------------------	----------------------------	---------------------------------------

// PROGRAM THESIS

AI is moving from ~~assistant~~ to *autonomous actor*.

// YESTERDAY

The AI Assistant

Responds to a prompt. Returns a single answer. Stateless. Tool-blind. Stops when the conversation ends.

- ▶ prompt → response
- ▶ no memory between turns
- ▶ cannot execute · cannot persist

// TODAY

The AI *Agent*

Pursues a goal. Plans the steps. Calls tools, runs code, queries memory, delegates to other agents — until the goal is done.

- ▶ goal → plan → act → reflect
- ▶ persistent memory + RAG
- ▶ tool use · multi-agent · deployable

// THE AI IMPERATIVE

Why agentic AI is the *defining skill* of the next decade.

// FIGURE_01

\$1.8T

PROJECTED VALUE

Global AI market by 2030.

src: PwC estimate

// FIGURE_02

70%

OF ENTERPRISES

Plan to deploy AI agents within 2 years.

src: Gartner

// FIGURE_03

4x

FASTER DELIVERY

AI-assisted software teams vs. traditional.

src: GitHub research

// FIGURE_04

Top 3

FASTEST-GROWING ROLE

"AI Engineer" — emerging job category.

src: LinkedIn Jobs Report

Figures are industry-cited estimates and research references for educational context. They are not financial projections or guarantees of any outcome.

// THE OPPORTUNITY

The talent gap is real — and *widening*.

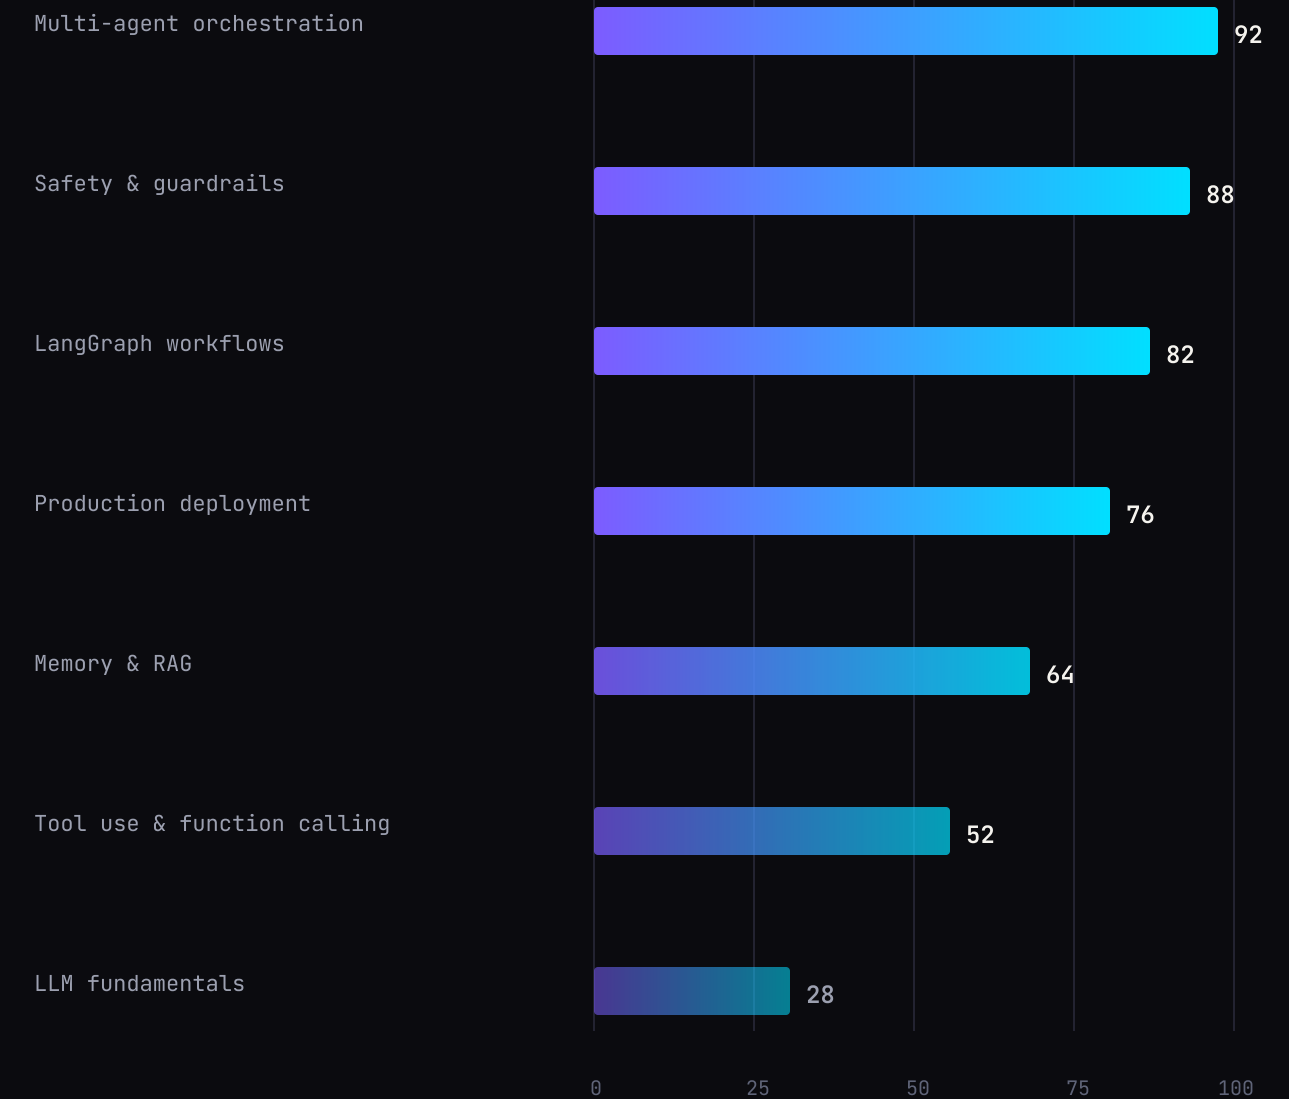
Enterprises are racing to deploy autonomous AI systems, but the pool of practitioners who can architect and safely deploy them remains thin.

- 01 LangChain, LangGraph, and agent skills are the fastest-rising keywords in AI job postings.
- 02 FAANG, consulting, and enterprise firms are creating "AI Engineer" and "Agent Engineer" roles.
- 03 Agentic workflows are replacing bespoke automation in data, support, and dev tooling.
- 04 Responsible deployment skills — safety, evals, guardrails — are the hardest to hire.

// CHART · SKILL_GAP_INDEX

Agentic AI skill demand vs. supply

Talent gap index by skill area · higher = greater shortage (illustrative)



// ANATOMY OF AN AGENT

Six components. *One autonomous system.*

[01] PERCEPTION

Parse input · understand state

Structured outputs · context windows · ReAct framing

[03] MEMORY

Persist context · retrieve facts

Vector stores · RAG · short / long / episodic

[05] ACTION

Execute · observe · iterate

Tool dispatch · result parsing · loop control

[02] PLANNING

Decompose goals · sequence steps

Chain-of-thought · tree-of-thought · self-critique

[04] TOOL USE

Call functions · run code

APIs · sandboxed exec · DB queries · function calling

[06] SAFETY

Guardrails · evals · constraints

Prompt injection defence · output filtering



// BUILT DIFFERENTLY – BECAUSE THE FIELD DEMANDS IT

Six principles that shape *every module*.

01 Code-first from day one

Every concept is implemented, not explained. You write agents, build pipelines, deploy systems.

// 12 implemented projects

02 The full agent stack

LLM reasoning → tool use → memory → orchestration
→ safety → deployment. End-to-end.

// end-to-end progression

03 Framework-native training

LangChain, LangGraph, CrewAI, AutoGen — as used in production, not toy demos.

// industry frameworks

04 Safety is not an afterthought

Dedicated module: evals, guardrails, prompt injection defence, responsible deployment.

// production-grade safety

05 Deployment & observability

Containerise, serve, monitor, cost-manage. LangSmith · FastAPI · Docker · cloud endpoints.

// production deployment

06 A capstone worth showing

Multi-agent system deployed to a cloud endpoint with docs, architecture review & demo.

// portfolio-ready

From LLM fundamentals to *memory systems*.

[MODULE]

01

Foundations of Large Language Models

Transformers, tokenization, prompting, and the LLM API landscape. What LLMs do well — and where agents are required.

▶ prompt evaluation lab · structured output exercises

[MODULE]

02

Introduction to AI Agents

Agent architecture: perception, reasoning, action, memory. ReAct, chain-of-thought, and tree-of-thought reasoning patterns.

▶ ReAct agent implementation from scratch

[MODULE]

03

Tool Use & Function Calling

Equipping agents with tools: web search, code execution, database lookup, third-party APIs. OpenAI function calling and tool schemas.

▶ tool-equipped research agent

[MODULE]

04

Memory Systems for Agents

Short-term, long-term, semantic, episodic memory. Vector stores (FAISS, Chroma, Pinecone), RAG, memory strategies.

▶ RAG pipeline + benchmark

LangChain → LangGraph → *multi-agent*.

[MODULE]

05

LangChain for Agent Development

Chains, LCEL, agents, runnables. Integrating LLMs, tools, and memory into coherent applications.

▶ multi-tool LangChain agent with memory

[MODULE]

06

LangGraph · Stateful Workflows

Graph-based orchestration with state, cycles, branching. Building complex, stateful reasoning pipelines.

▶ stateful customer support agent

[MODULE]

07

Multi-Agent Systems · CrewAI & AutoGen

Crews of specialised agents. Role assignment, delegation, inter-agent communication, task decomposition.

▶ multi-agent content generation pipeline

[MODULE]

08

Planning & Task Decomposition

Goal decomposition, sub-task planning, self-critique, reflection loops. Recovering from failure.

▶ autonomous planner with self-correction

Code · safety · ship — *to production.*

[MODULE]

09

Code Generation & Execution Agents

Agents that write, execute, and debug code. Sandboxed environments, interpreter agents, data-analysis workflows.

▸ data analysis agent with Python execution

[MODULE · critical]

10

Evaluation, Safety & Guardrails

Agent reliability & output quality. Prompt-injection defence, output filtering, constitutional AI, responsible deployment.

▸ evaluation framework + safety audit

[MODULE]

11

Production Deployment

Containerise (Docker), serve (FastAPI), observe (LangSmith, Arize), manage cost & latency.

▸ deploy agent to cloud endpoint w/ monitoring

[CAPSTONE]

12

Capstone · Autonomous Agent

End-to-end design & deployment of a multi-tool, multi-agent system on a real business problem. Mentor-reviewed.

▸ architecture doc · code · deployment · demo

// YOUR 6-MONTH PATH

Four phases · one *deployed agent*.

M1-2

M2-4

M4-5

M5-6

// PHASE 01

Foundations

// STUDY

LLM fundamentals · agent architecture · tool use · memory

// BUILD

prompt_lab · react_agent · rag_pipeline

// PHASE 02

Frameworks

// STUDY

LangChain · LangGraph · CrewAI · AutoGen orchestration

// BUILD

langchain_agent · langgraph_flow · agent_crew

// PHASE 03

Advanced Agents

// STUDY

Planning · code execution · evals · safety · guardrails

// BUILD

task_planner · code_agent · safety_audit

// PHASE 04

Deployment & Capstone

// STUDY

Production deployment · end-to-end capstone agent

// BUILD

cloud_endpoint + capstone_portfolio

// THE AGENT TECHNOLOGY STACK

Hands-on with the *production stack*.

LAYER 04

**Agent
Orchestration****LangChain**

Chains · runnables · agent executors

LangGraph

Stateful graph-based orchestration

CrewAI

Multi-agent role-based collaboration

AutoGen

Conversational multi-agent framework

LAYER 03

**Runtime &
Deployment****Python**

primary implementation language

FastAPI

Agent REST API serving

Docker

Containerisation & portability

LangSmith

Observability · tracing · evals

LAYER 02

**Memory &
Retrieval****FAISS**

Local vector similarity search

Chroma

Embedded vector store

Pinecone

Managed vector retrieval at scale

RAG patterns

Retrieval-augmented generation

LAYER 01

**Models &
Cloud****OpenAI APIs**

GPT-4o · function calling

Hugging Face

Open-source LLMs & tools

Cloud Endpoints

GCP · AWS · Azure deployment

Arize

LLM observability & eval

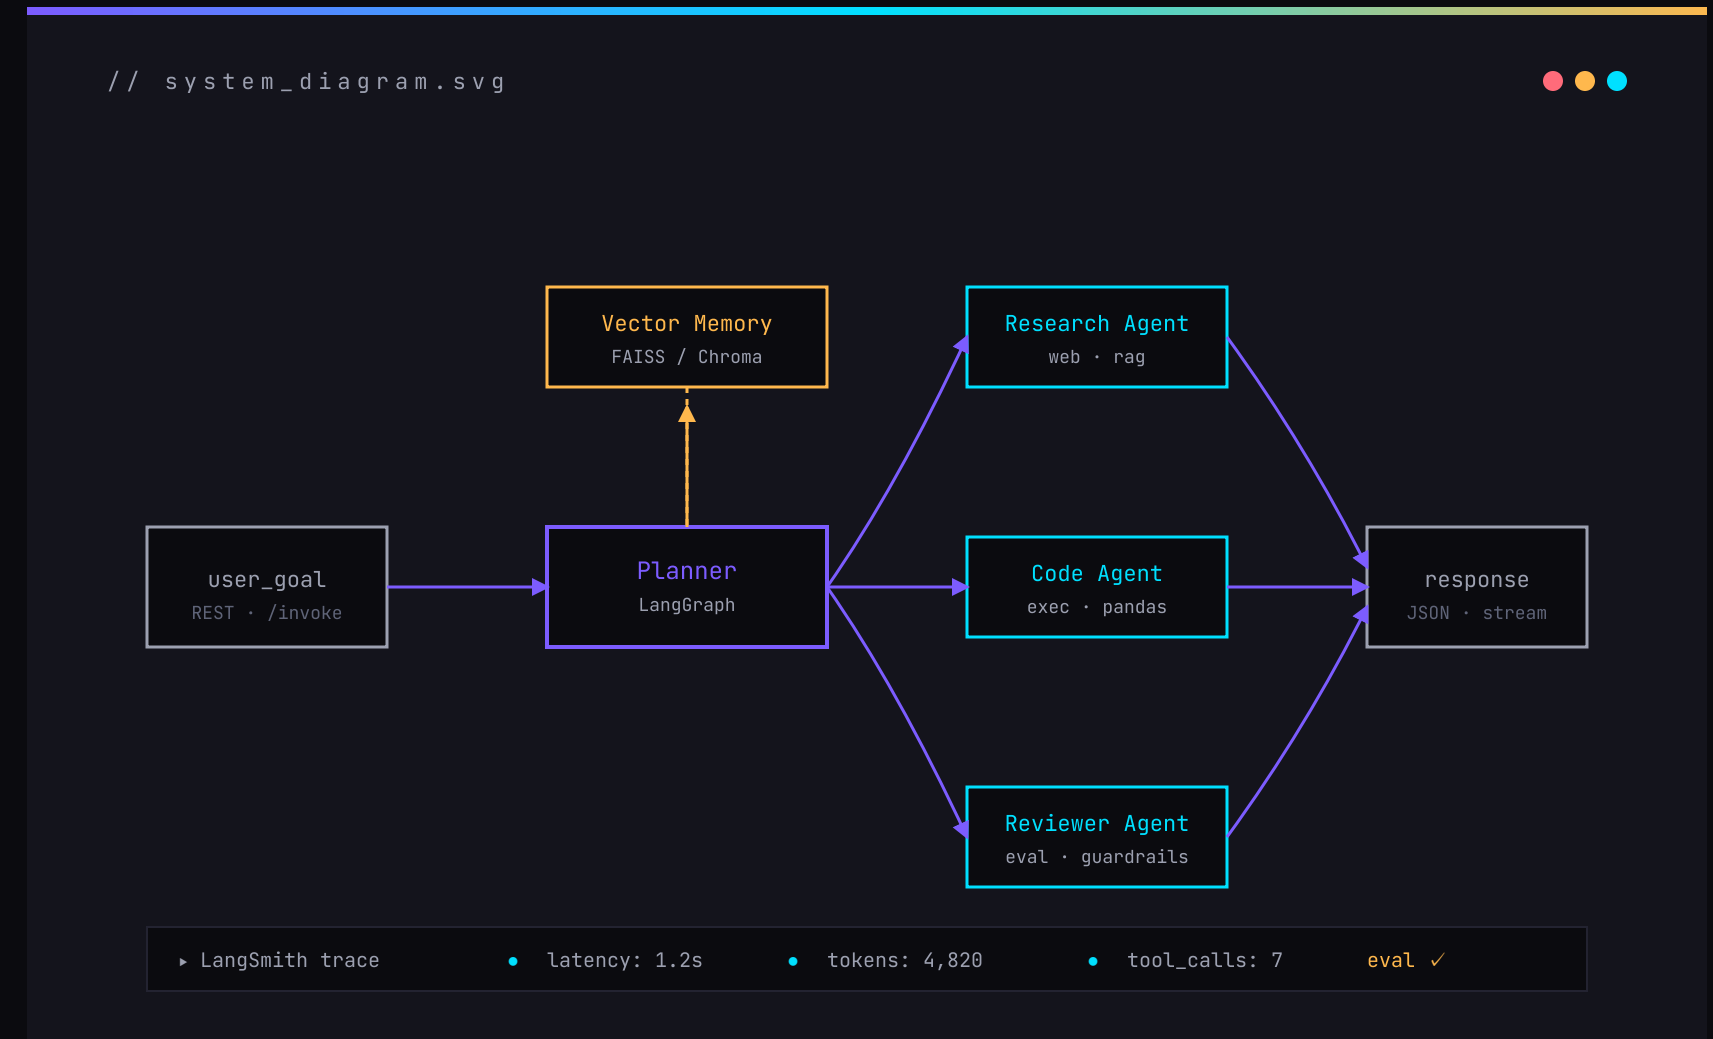
// THE CAPSTONE

An end-to-end autonomous AI *system*.

Design and deploy a multi-tool, multi-agent system addressing a real business problem of your choice — architecture, implementation, cloud deployment, mentor-reviewed demo, and a professional README.

// DELIVERABLES_LIST

- ✓ `architecture_design.pdf`
- ✓ `implementation/` // LangChain · LangGraph · CrewAI
- ✓ `deploy.sh` // cloud endpoint + monitoring
- ✓ `demo.mp4` // live or recorded · mentor-reviewed
- ✓ `README.md` // recruiter-ready documentation



// DESIGNED FOR PRACTITIONERS

Not an intro survey. *A technical program.*

▶ **PREREQ** Python required · basic ML helpful · prior LLM experience optional



Software Engineers

Extend your skills to agentic workflows, tool-using LLMs, and autonomous pipelines.



Data & ML Engineers

Move from model training to deploying AI systems that act on predictions.




AI / ML Practitioners

Deepen your LLM knowledge into agent architectures, RAG, and memory systems.




Automation & DevOps

Build AI-powered automation that orchestrates tools, APIs, and workflows.



AI Product Managers

Understand what agents can and cannot do — and scope agentic products credibly.



Career Switchers

Structured path from Python & LLM foundations to production-grade agent dev.

// FROM PROGRAM TO PLACEMENT

Roles this program *prepares you for.*

// TARGET_ROLES[]

- AI Engineer
- Agentic AI Developer
- LLM Application Engineer**
- ML Engineer (NLP/LLM)
- AI Automation Engineer
- AI Product Engineer**
- Conversational AI Engineer
- AI Research Engineer
- Generative AI Specialist**
- AI Solutions Architect

// service_01

AI Role Readiness

Resume + LinkedIn + GitHub portfolio for AI Engineer / LLM roles.

// service_02

Technical Interview Prep

LLM & agent architecture drills · mock interviews · system design.

// service_03

Portfolio & Capstone Review

Capstone mentor polish · architecture doc · presentation coaching.

// service_04

Job Search Strategy

Role mapping · recruiter outreach · offer evaluation guidance.

// CHART · AI_ENGINEER_COMPENSATION

AI Engineer bands (USD, illustrative)



Illustrative global reference ranges; actual compensation varies by region, employer, and experience. Job placement assistance is provided; employment outcomes are not guaranteed.

// WHY THIS - NOT A GENERIC AI COURSE

A program built for the *production agent stack*.

FEATURE	GENERIC AI/ML	LLM SURVEY VIDEOS	LEARNING SAINT · AGENTIC AI
Agentic architecture focus	⊖ Rarely	⊖ Rarely	✓ Core curriculum
LangGraph · CrewAI · AutoGen	⊖ Partial	⊖ Variable	✓ Dedicated modules
Memory · RAG · Tool use	⊖ Limited	⊖ Limited	✓ In every project
Production deployment	⊖ Rarely	× No	✓ Module 11 + capstone
Safety & guardrails	⊖ Brief mention	× No	✓ Dedicated module
Mentor feedback	⊖ Limited	× No	✓ Structured
Career support	× No	× No	✓ Included



// AWARDED BY LEARNING SAINT

Certificate of Completion


Agentic AI & Autonomous Systems

- ✓ 12 modules complete
- ✓ project assignments submitted
- ✓ capstone deployed & reviewed

Recognises structured training for professional development. Not a regulatory credential or employment guarantee.

// ENROLL IN FIVE STEPS

Initialize your *agent journey*.



PROGRAM_CONFIG.JSON

DURATION 6 months	MODE Live online + recordings
LEVEL Intermediate → Advanced	INSTRUCTION English
INSTALLMENTS Available	PROGRAM FEE Talk to admissions

// PREREQ
Python required · basic ML helpful · prior LLM experience optional.

// next_action

Start building autonomous AI systems.

Mentor-led instruction · hands-on projects · capstone mentoring · career support.

TALK TO AN ADVISOR


// REACH US · THREE REGIONS


Talk to an advisor in your *region*.


Three offices, one program — speak to admissions in the region nearest to you.

 // region_01
USA Office

 COMPANY
Learning Saint Inc.


 ADDRESS
3201 Leith Ln, Louisville, KY 40218
16192 Coastal Highway, Lewes, Delaware 19958

 PHONE
+1 (302) 592-6724


 EMAIL
info@learningsaint.com

 // region_02 · headquarters HQ
United Kingdom


 COMPANY
Learning Saint Ltd.


 ADDRESS
182-184 High Street North,
East Ham, London E6 2JA

 PHONE
+44 20 4621 6235


 EMAIL
support@learningsaint.com

 // region_03
India Office

 COMPANY
Learning Saint Pvt. Ltd.

 ADDRESS
H-70, H Block, Sector 63,
Noida, Uttar Pradesh 201309, India

 PHONE
+91 96959 13043

 EMAIL
info@learningsaint.com