

## AGENTIC AI ARCHITECTURE PATTERNS REFERENCE

# Agentic AI Architecture Patterns Reference

Six production patterns with diagrams — ReAct, orchestration, planning, memory, RAG and governance — plus a framework cheat sheet for LangChain, LangGraph, AutoGen and RAG, mapped to the GSDC Agentic AI Expert syllabus.

**6**

PATTERNS + DIAGRAMS

**4**

FRAMEWORKS COMPARED

**12**

SYLLABUS MODULES

**32**

PORTFOLIO ACTIVITIES

PATTERNS & DIAGRAMS · FRAMEWORK CHEAT SHEET · SALARY GUIDE · CAREER ROADMAP · GOVERNANCE ROLES · HIRING TRENDS · SYLLABUS · EXAM

## HOW TO USE THIS REFERENCE

# Six patterns, one mental model

Most production agentic systems are combinations of six core patterns. Each page that follows pairs a diagram with when to use it, trade-offs, and the GSDC modules that teach it.

1

**ReAct**

Reason → act → observe loop with tools. The base agent pattern.

2

**Orchestrator–Worker**

A coordinator routes work to specialised agents and aggregates results.

3

**Planner–Executor**

Decompose a goal into steps, then execute and replan on failure.

4

**Memory-Augmented**

Short-term context plus long-term vector memory for recall.

5

**RAG Pipeline**

Ground answers in retrieved external knowledge.

6

**Governance / Guardrail**

Wrap agents with validation, policy and human-in-the-loop.

**How to read the diagrams**

- **Solid blue arrows** — the main forward flow
- **Dashed amber arrows** — feedback / loops
- **Filled blue boxes** — the LLM / agent core
- **Light boxes** — tools, data & I/O

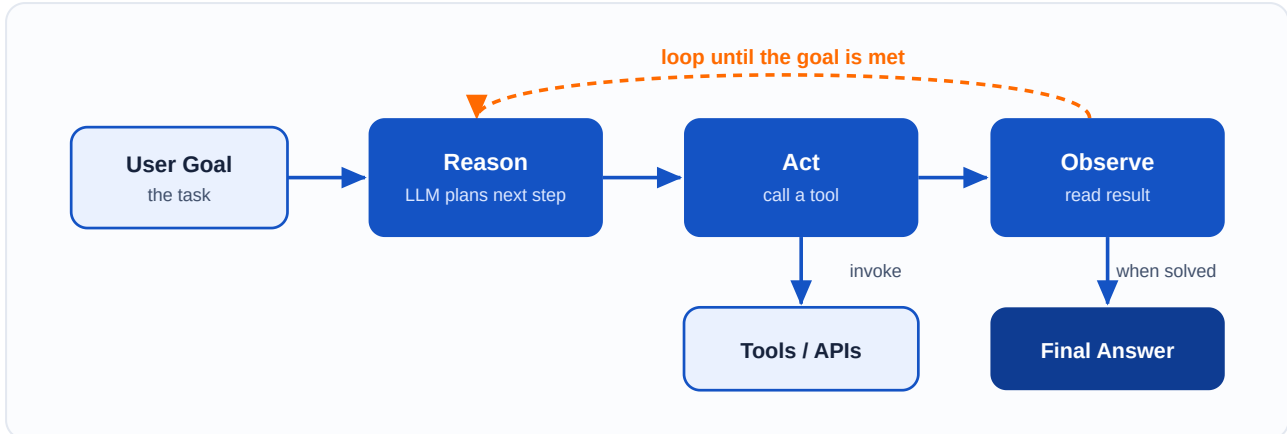
**Patterns compose**

Real systems stack these — e.g. an orchestrator of ReAct agents, each using RAG and long-term memory, all wrapped in a governance layer. Learn them individually, then combine.

## PATTERN 01

# 1 ReAct — Reason + Act

The foundational agent loop: the model reasons about a goal, takes an action (a tool call), observes the result, and repeats until the goal is met.



## Use when

- A task needs tools (search, code, APIs)
- Steps aren't known in advance
- You want transparent, traceable reasoning

## Watch-outs

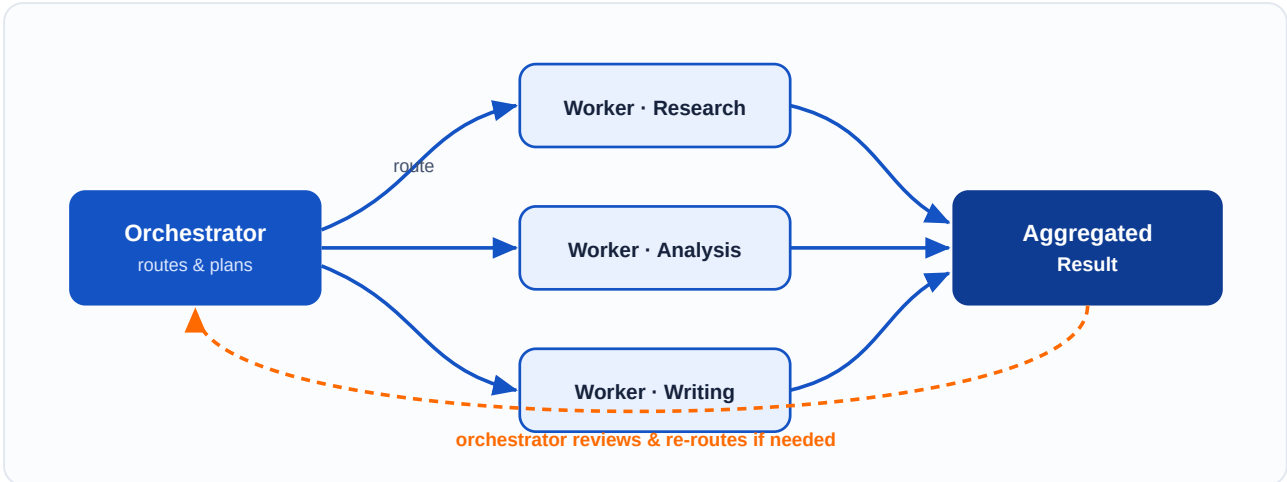
- Can loop indefinitely — cap iterations
- Tool errors need handling
- Token cost grows with each cycle

Taught in Modules 1 & 4 · implement with LangChain / LangGraph.

PATTERN 02

## 2 Orchestrator-Worker

A coordinator agent decomposes a request, routes sub-tasks to specialised worker agents in parallel, then aggregates their outputs into one result.



**Use when**

- Work splits into specialised sub-tasks
- Parallelism speeds things up
- Different tools/prompts per sub-task

**Watch-outs**

- Coordination overhead & latency
- Conflicting worker outputs to reconcile
- Cost scales with worker count

50% OFF

**Build multi-agent systems — at half off**

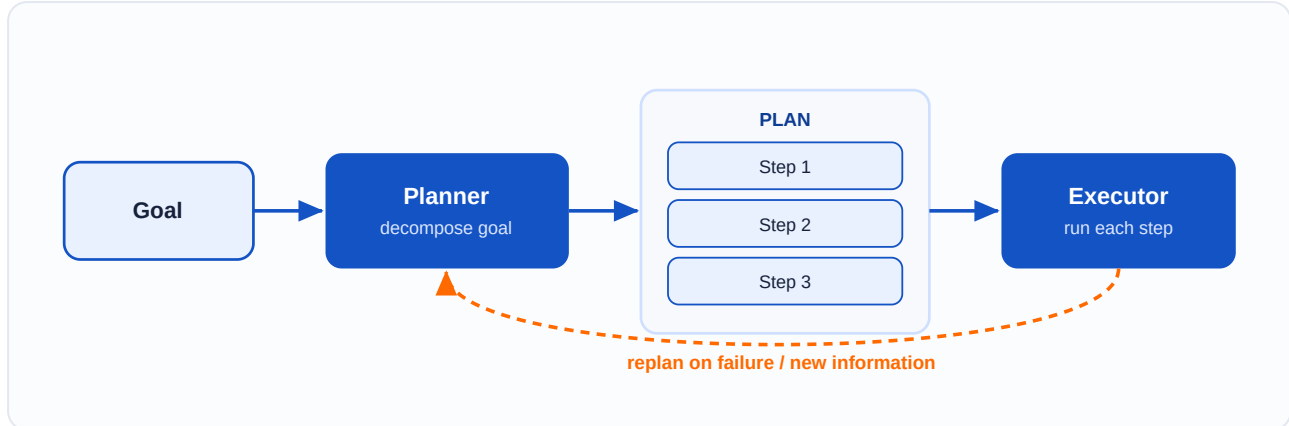
Module 5 teaches orchestration end-to-end. Enroll now at 50% off and ship coordinated multi-agent systems with confidence.

[Claim 50% Off →](#)

## PATTERN 03

### 3 Planner–Executor

Separate the thinking from the doing: a planner decomposes the goal into an ordered plan, an executor runs each step, and the planner revises when a step fails.



#### Use when

- Goals are multi-step & long-horizon
- You want an inspectable plan up front
- Steps may need re-ordering mid-run

#### Watch-outs

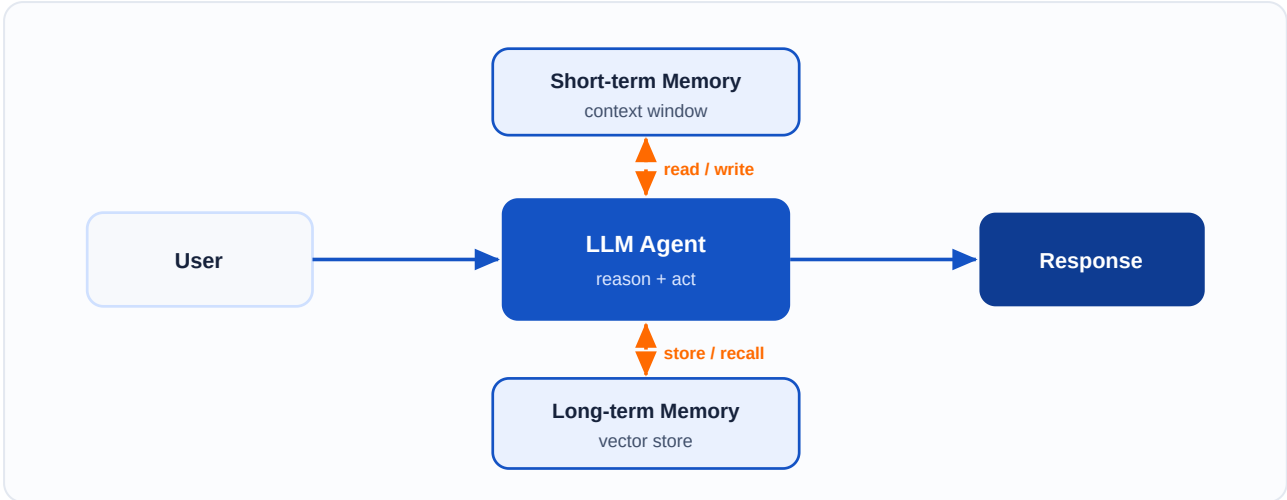
- Bad initial plans waste cycles
- Needs clear step success criteria
- Replans can cascade — bound them

*Taught in Modules 1, 2 & 10 (capstone) · implement with LangGraph state machines.*

PATTERN 04

# 4 Memory-Augmented Agent

Give the agent recall: a short-term working memory (the context window) plus a long-term store (a vector database) it can write to and retrieve from across sessions.



**Use when**

- Continuity across turns or sessions
- Personalisation from past interactions
- Knowledge exceeds the context window

**Watch-outs**

- Stale or conflicting memories
- Retrieval relevance & privacy
- Memory growth & pruning strategy

**LIMITED-TIME OFFER**

## Master memory & retrieval design

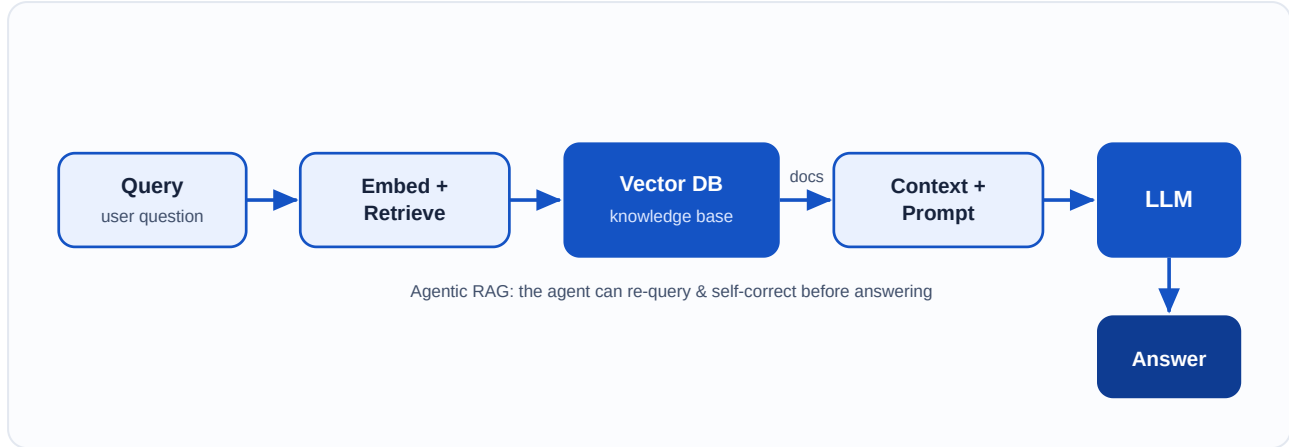
Enrollment is open for a limited time. Learn to architect short- and long-term memory the way production teams do.

[Enroll Now →](#)

## PATTERN 05

## 5 Retrieval-Augmented Generation

Ground answers in external knowledge: embed the query, retrieve relevant chunks from a vector store, add them to the prompt, and let the LLM answer from grounded context.



### Use when

- Answers must cite current/private data
- Reducing hallucination matters
- Knowledge changes faster than retraining

### Watch-outs

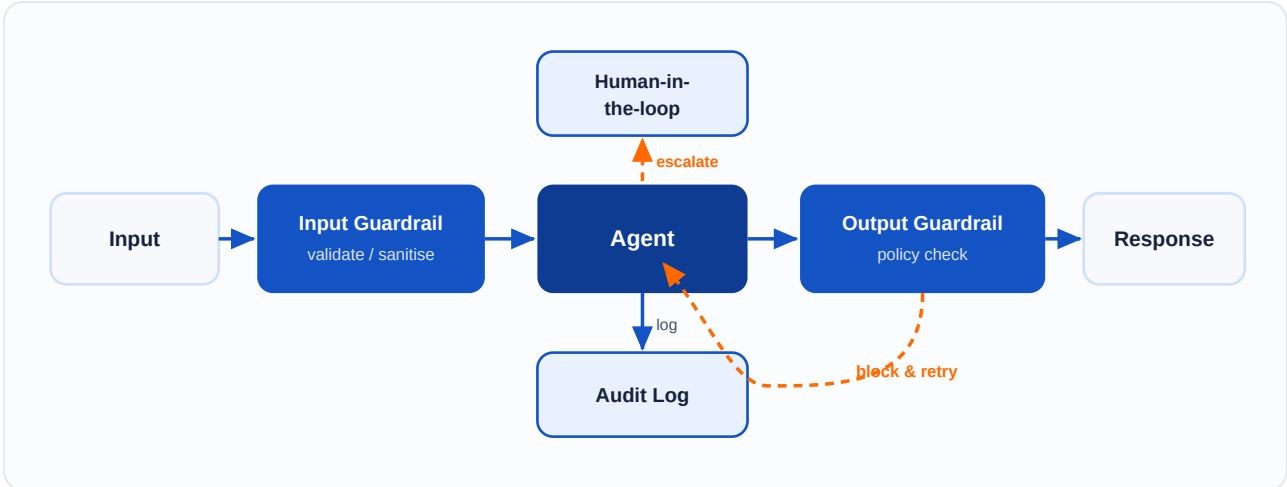
- Retrieval quality caps answer quality
- Chunking & embedding choices matter
- Keep the index fresh

Taught in Module 4 - implement with LangChain (linear RAG) or LangGraph (agentic RAG).

PATTERN 06

# 6 Governance & Guardrail Wrapper

Wrap the agent in oversight: validate inputs, check outputs against policy, log every decision, and escalate to a human when confidence or risk thresholds are crossed.



**Use when**

- Agents act with real-world impact
- Compliance & auditability are required
- Safety / brand risk must be contained

**Watch-outs**

- Over-blocking hurts usefulness
- Guardrails add latency
- Define clear escalation thresholds

**48 HOURS ONLY**

**Learn to govern autonomous agents**

Module 7 covers guardrails, policy and human-in-the-loop. This offer is valid for the next 48 hours — secure your seat.

[Secure My Seat →](#)

## FRAMEWORK CHEAT SHEET

# LangChain · LangGraph · AutoGen · RAG

Choose the right tool for the pattern. Current as of 2026 — note AutoGen's status change.

Framework	Best for	Strengths	Watch-outs
<b>LangChain</b>	Linear pipelines & RAG; tool/API integration	Modular, model-agnostic, huge ecosystem	Less ideal for complex stateful loops
<b>LangGraph</b>	Stateful agents with loops, branching, human pauses	Graph model; explicit state; persistence & rollback	Steeper curve; overkill for simple RAG
<b>AutoGen</b>	Multi-agent conversation patterns	Pioneered multi-agent chat; strong capabilities	In maintenance mode; merging into MS Agent Framework
<b>RAG</b>	Knowledge-intensive, grounded answers	Cuts hallucination; uses live/private data	A pattern, not a framework; retrieval quality is key

## When to choose what

- Simple Q&A / RAG bot → **LangChain**
- Loops, retries, human-in-loop → **LangGraph**
- Dynamic agent-to-agent dialogue → **AutoGen**
- Answers grounded in your data → **RAG**

## 2026 notes

- LangGraph is the LangChain team's evolution for agents
- AutoGen + Semantic Kernel → Microsoft Agent Framework
- Security & guardrails are your responsibility, not the framework's
- Agentic RAG lets agents re-query and self-correct

Framework status per 2026 industry sources (LangChain/LangGraph guidance; Microsoft's consolidation of AutoGen). Verify the latest before committing.

## PUTTING IT TOGETHER

# Pattern → framework → role

How each pattern maps to a framework choice, a syllabus module, and the role that owns it day to day.

Pattern	Typical framework	Module	Owned by
ReAct	LangChain / LangGraph	1, 4	Agentic AI Engineer
Orchestrator-Worker	LangGraph / AutoGen	5	Multi-Agent Systems Engineer
Planner-Executor	LangGraph	1, 2, 10	AI Solution Architect
Memory-Augmented	LangChain + vector DB	4	Agentic AI Engineer
RAG Pipeline	LangChain / LlamaIndex	4	AI / ML Engineer
Governance / Guardrail	Framework-agnostic	7, 8	AI Governance Lead

### The capstone ties it together

Module 10's capstone has you combine several patterns into one end-to-end, governed solution — the portfolio piece that proves architect- and principal-level readiness.

#### RECOMMENDED FOR YOU

### Go from patterns to a shipped system

Related to this reference: the certification turns these patterns into a built, governed capstone. Start your enrollment.

[Get Certified →](#)

## FULL SYLLABUS

# The 12-module curriculum

36+ hours of learning · 2 practice exams · daily live sessions · job-support program.

**1 Advanced Agentic AI Concepts & Foundations**

- Multi-agent systems & emergent behavior
- Meta-learning & self-adaptive agents
- Goal-directed decision-making

**2 Designing Scalable Agentic AI Architectures**

- Microservices for AI agents
- Distributed orchestration & fault tolerance
- High availability & scalability

**3 Integration with Emerging Technologies**

- IoT for real-time sensing & action
- Edge computing for latency & privacy
- Blockchain & quantum applications

**4 Advanced LLMs & Framework Customization**

- Fine-tuning LLMs for domain tasks
- Prompt engineering & RLHF
- LangChain, LangGraph, AutoGen & RAG

**5 Multi-Agent Coordination & Protocols**

- Collaboration & consensus building
- Secure communication protocols
- Conflict resolution & negotiation

**6 Agentic AI in Enterprise Use Cases**

- Finance, healthcare, manufacturing
- Government & public services
- Sector-specific deployment patterns

Patterns 1–6 are taught across Modules 1, 2, 4, 5, 7 & 10.

## FULL SYLLABUS

# Governance, scale, capstone

**7 Governance, Compliance & Ethical AI Leadership**

- Governance frameworks & policy
- Bias detection & mitigation
- Regulation, compliance & accountability

**10 Capstone: End-to-End Agentic AI Solution**

- Requirement analysis & design
- Build & integrate multi-agent systems
- Deploy, validate, document & present

**8 Operationalizing Agentic AI**

- Monitoring & alerting systems
- Anomaly detection & fault diagnosis
- Retraining, incident response & rollback

**11 Building AI Agents — No Coding Required**

- Drag-and-drop workflow design
- Pretrained components & business rules
- One-click deployment & monitoring

**9 Preparing Organizations at Scale**

- AI adoption & transformation roadmaps
- Closing talent gaps & upskilling
- Change management & stakeholders

**12 Personalized 1-1 & Daily Live Sessions**

- 1-on-1 Trainer / SME connect
- Daily expert-led live sessions
- Global peer community & lifelong access

**EXCLUSIVE OFFER****Master the full curriculum**

Unlock an exclusive enrollment offer and learn every pattern, framework and governance practice in one structured program.

**Enroll Now** →

EXAM BLUEPRINT

# How you earn the credential

40

Questions

90

Minutes

65%

To Pass

5 yr

Validity

Parameter	Detail
Exam questions	40 questions
Format	Multiple choice
Duration	90 minutes
Passing score	65%
Language	English
Open book	No
Certification validity	5 years
Complimentary retake	Yes — one free retake included

### Indicative weighting

- Foundations & architecture patterns — ~30%
- LLMs, frameworks & coordination — ~30%
- Enterprise use cases & operations — ~20%
- Governance, ethics & scale — ~20%

### How to prepare

- Know all six patterns & when to use them
- Complete all 12 modules & 32 labs
- Sit both included practice exams
- Use the GSDC Book of Knowledge

Weightings indicative for study planning. Refer to the official exam guide for the authoritative blueprint.

## SALARY GUIDE

# What these skills pay

Indicative 2026 US base by role, with median hourly. Pattern fluency and multi-agent specialisation pull pay toward the upper percentiles.

Role	Median base	90th pct	\$/hr*
Agentic AI Engineer	\$165k	\$235k	\$79
Multi-Agent Systems Engineer	\$195k	\$300k	\$94
AI Solution Architect	\$225k	\$315k	\$108
Principal AI Engineer	\$320k	\$450k	\$154
AI Governance Lead	\$200k	\$260k	\$96

\*Median ÷ 2,080 hrs. Sources: ZipRecruiter (2026) & 2026 industry salary guides; PwC 2025 reports a 56% AI-skills wage premium. Indicative; not a guarantee.

**HALF PRICE**

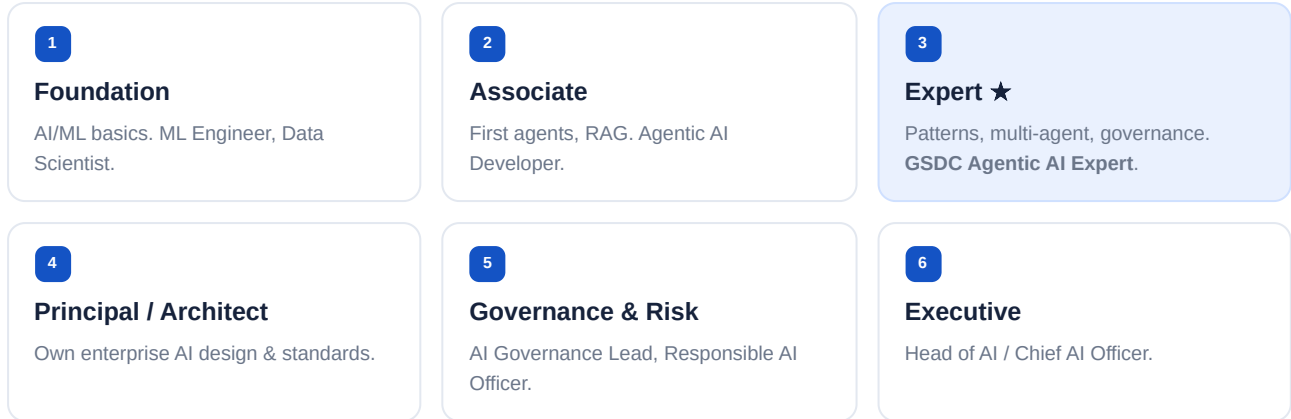
## Skills that move you up the pay curve

Certify the exact pattern and framework skills these roles screen for  
— currently available at half price for new enrollees.

[Start Prep →](#)


CAREER ROADMAP

# From practitioner to AI leader



HIRING TRENDS

## What 2026 employers screen for



AI-skill job postings rose 7.5% YoY even as total postings fell 11.3% (PwC 2025 Global AI Jobs Barometer).

## AI GOVERNANCE ROLES

# The oversight career lane

Pattern 6 (governance) maps to a fast-growing, well-paid set of oversight roles — taught directly in Module 7.

Role	Median base	Core mandate
AI Governance Lead	\$200k	Frameworks, policy & oversight
Responsible AI Officer	\$210k	Ethics & accountability in delivery
AI Risk & Compliance Mgr	\$185k	Regulatory mapping & audit
AI Ethics Specialist	\$175k	Bias detection & mitigation
Chief AI Officer	\$350k+	Enterprise AI strategy & assurance

### Why it's booming

- Agents act without a human in the loop
- AI regulation is tightening globally
- Boards now demand AI assurance

### What Module 7 covers

- Governance frameworks & policy
- Bias detection & mitigation
- Compliance & accountability

**RISK-FREE OFFER**

## Move into governance — with a 7-day guarantee

Try the program risk-free. If it isn't right within 7 days, get a full refund. Build the oversight skills enterprises need.

[Enroll Risk-Free →](#)

## 32 LEARN-BY-DOING PORTFOLIO ACTIVITIES

# Implement the patterns — 1–16

Each lesson ships with a built-in, AI-powered challenge. These 32 artifacts become your hiring portfolio. Activities 1–16 below.

- 1 **Map a multi-agent system** and its emergent behaviors.
- 2 **Design a self-adaptive agent** with a meta-learning loop.
- 3 **Model goal-directed autonomy** for a decision agent.
- 4 **Extract an agentic pattern** from current research.
- 5 **Architect microservices** for an agent ecosystem.
- 6 **Add fault tolerance** to a distributed agent cluster.
- 7 **Load-balance agents** and optimise resources.
- 8 **Design for high availability** and horizontal scale.
- 9 **Connect IoT sensors** for real-time agent action.
- 10 **Deploy an edge agent** for latency & privacy.
- 11 **Use blockchain** for verifiable agent interactions.
- 12 **Prototype a quantum-inspired** optimisation routine.
- 13 **Fine-tune an LLM** for a domain-specific task.
- 14 **Build a ReAct agent** with tools & an iteration cap.
- 15 **Run an RLHF loop** to align agent responses.
- 16 **Build a RAG workflow** in LangChain / LangGraph.

## 32 LEARN-BY-DOING PORTFOLIO ACTIVITIES

# Implement the patterns — 17–32

- 17 **Build an orchestrator** routing to worker agents.
- 18 **Implement consensus** across collaborating agents.
- 19 **Build a negotiation protocol** for competing agents.
- 20 **Secure agent messaging** with an efficient protocol.
- 21 **Build a finance agent** for fraud or risk detection.
- 22 **Prototype a healthcare** diagnostics support agent.
- 23 **Automate a supply chain** with predictive agents.
- 24 **Design a public-service** automation workflow.
- 25 **Add a governance wrapper** with input/output guardrails.
- 26 **Run a bias audit** and apply mitigation.
- 27 **Map a compliance checklist** to a regulation.
- 28 **Add a planner–executor** with replan-on-failure.
- 29 **Detect anomalies** and diagnose agent faults.
- 30 **Wire long-term memory** with a vector store.
- 31 **Build a no-code agent** with drag-and-drop tooling.
- 32 **Ship the capstone** — combine patterns end-to-end.

## Your portfolio outcome

By completion you'll have implemented all six patterns plus a combined capstone — concrete, demonstrable proof of architecture capability for hiring managers and clients.

## WHAT YOU GET

# Everything in your enrollment

**Self-paced expert-led videos**

Full tutorial library across all 12 modules, lifetime access.

**GSDC Studio — daily live sessions**

100+ live monthly sessions with global AI experts.

**1-on-1 SME connect**

Personal trainer / SME sessions to resolve any query.

**GSDC Book of Knowledge**

Structured study material aligned to the exam.

**Exam + free retake & practice**

Certification exam with a complimentary retake.

**Capstone + job support**

Portfolio capstone plus a job-support program.

**7-DAY  
BACK****Globally recognized · 5-year validity · 7-day money-back guarantee**

A vendor-neutral credential trusted by 2,50,000+ professionals. If within 7 days you haven't gained the skills you're looking for, you're covered by a full 100% money-back guarantee. (See refund policy for terms.)

Oracle

Amazon

Microsoft

Cisco

HP

Deloitte

Bank of America

Volkswagen

Logos denote organisations where GSDC-certified professionals work; not an endorsement.

## YOUR NEXT MOVE

# From reference to real systems

You now have the six patterns, the framework cheat sheet, and the role and pay context. The next step is building them for real — with mentorship, a capstone, and a recognized credential.

**6 patterns**

ReAct to governance, with diagrams.

**32 labs + capstone**

Implement every pattern, then combine them.

**Job support**

Capstone, mentorship & a job-support program.

## FINAL CALL

**Claim your seat in the current cohort**

This is your last prompt in this reference — enroll now to lock in your place, your mentorship and your offer before the cohort fills.

[Enroll Now →](#)**Get in touch**

Web

[gsdcouncil.org](https://gsdcouncil.org)

Email

[accreditations@gsdcouncil.org](mailto:accreditations@gsdcouncil.org)

WhatsApp

[+1 251-333-1717](https://wa.me/12513331717)

Offices: Lewes, Delaware (USA) · Glattbrugg (Switzerland) · Singapore. © 2025 GSDCouncil. All rights reserved. Framework status & salary data indicative; sources cited within.