

AGENTIC AI USE-CASE &amp; TOOLKIT PACK

# Agentic AI Use-Case & Toolkit Pack

Four vertical playbooks with illustrative case studies, an enterprise readiness & ROI checklist, the 32 hands-on builds indexed by category, and the full GSDC Agentic AI Expert syllabus — everything you need to turn agentic AI into real outcomes.

4

VERTICAL PLAYBOOKS

32

BUILDS BY CATEGORY

12

SYLLABUS MODULES

2

READINESS CHECKLISTS

VERTICAL PLAYBOOKS · CASE STUDIES · READINESS & ROI · 32 BUILDS BY CATEGORY · SALARY GUIDE · CAREER ROADMAP · GOVERNANCE · SYLLABUS · EXAM

## HOW TO USE THIS PACK

# From use case to capability

This pack pairs real-world agentic-AI use cases with the skills to build them. Start with the vertical that matches your industry, sanity-check readiness and ROI, then use the build index and syllabus as your learning plan.

1

**Finance**

Trading, fraud detection, risk.

2

**Healthcare**

Diagnostics, monitoring, care plans.

3

**Manufacturing**

Maintenance, supply chain, QC.

4

**Government**

Service automation, crisis mgmt.

**What's in the pack**

- Four vertical playbooks + illustrative case studies
- Enterprise readiness checklist (printable)
- ROI checklist & value model
- 32 Learn-by-Doing builds, indexed by category
- Full 12-module syllabus & exam blueprint

**How to use it**

- Pick your vertical playbook (pages 3–6)
- Run the readiness & ROI checklists (7–8)
- Plan skills from the build index (9–10)
- Study the modules & sit the exam

*Case studies are illustrative, composite scenarios for planning — not accounts of specific named organisations. Outcome figures are indicative.*

VERTICAL PLAYBOOK 01

# 1 Finance

Agentic AI in finance targets high-volume, time-sensitive decisions where speed, accuracy and auditability all matter.

<b>Automated trading</b> Agents monitor signals and execute within policy limits.	<b>Fraud detection</b> Real-time anomaly detection across transactions.	<b>Risk management</b> Continuous exposure modelling & alerts.
--	--	---

**Illustrative case study — real-time fraud detection**

A mid-size bank deploys a fraud-detection agent that scores transactions in real time, retrieves customer and merchant context (RAG), and escalates high-risk cases to a human analyst with an explanation.

<b>↓ false +</b> Fewer false positives vs. rules-only	<b>Real-time</b> Sub-second scoring	<b>Audit</b> Every decision logged
--	--	---------------------------------------

**Patterns & skills used**

- RAG
- ReAct
- Governance wrapper
- Human-in-the-loop
- Monitoring

Maps to Modules 6 & 7 · builds 16, 21, 25, 28. Illustrative composite; outcomes indicative.

VERTICAL PLAYBOOK 02

# 2 Healthcare

In healthcare, agents augment clinicians — never replace them — with strict guardrails, traceability and human sign-off.

<b>Personalized care plans</b> Synthesise history & guidelines into draft plans.	<b>Diagnostics support</b> Surface relevant evidence for clinician review.	<b>Patient monitoring</b> Watch signals & flag deterioration early.
---	---	--

**Illustrative case study — clinician decision support**

A hospital network pilots a monitoring agent that watches vitals, retrieves the patient's record and guidelines, and alerts the care team with a summarised rationale — always routing final decisions to a clinician.

<b>Earlier</b> Faster deterioration alerts	<b>Human</b> Clinician sign-off required	<b>Traceable</b> Evidence shown for each alert
---	---	---

**Patterns & skills used**

- Memory-augmented
- RAG
- Governance wrapper
- Human-in-the-loop
- Bias auditing

**50% OFF**

**Build safe, governed healthcare agents**

Modules 6 & 7 cover sector use cases and governance. Enroll now at 50% off and learn to ship agents responsibly.

[Claim 50% Off →](#)

VERTICAL PLAYBOOK 03

# 3 Manufacturing

Manufacturing pairs agents with sensors and operational systems to keep lines running and supply chains resilient.

### Predictive maintenance

Predict failures before they stop the line.

### Supply-chain automation

Agents re-plan around disruptions.

### Quality control

Detect defects and trigger interventions.

### Illustrative case study — predictive maintenance

A manufacturer connects edge agents to machine sensors. Agents detect anomalies, diagnose likely causes, schedule maintenance, and escalate when confidence is low — reducing unplanned downtime.

#### ↓ downtime

Fewer unplanned stoppages

#### Edge

Low-latency on-site agents

#### Plan

Auto-scheduled maintenance

### Patterns & skills used

Edge deployment

Planner-Executor

Anomaly detection

Orchestration

Monitoring

Maps to Modules 3, 6 & 8 · builds 10, 23, 28, 29. Illustrative composite; outcomes indicative.

VERTICAL PLAYBOOK 04

# 4 Government & Public Sector

Public-sector agents prioritise transparency, equity and accountability — with governance front and centre.

<b>Service automation</b> Faster, fairer handling of citizen requests.	<b>Policy simulation</b> Model outcomes before rollout.	<b>Crisis management</b> Coordinate response across agencies.
---	--	--

**Illustrative case study — citizen service automation**

An agency deploys an orchestrated set of agents to triage requests, retrieve the right policy, draft responses, and route edge cases to staff — with full audit trails and bias checks built in.

<b>Faster</b> Shorter response times	<b>Fair</b> Bias audited & logged	<b>Open</b> Transparent decision trails
---	--------------------------------------	--

**Patterns & skills used**

- Orchestration
- RAG
- Governance wrapper
- Compliance mapping
- Audit logging

**LIMITED-TIME OFFER**

## Deliver agents the public sector can trust

Enrollment is open for a limited time. Learn the governance and orchestration skills these projects demand.

[Enroll Now →](#)

## ENTERPRISE READINESS CHECKLIST

# Are you ready to deploy agents?

Print and score yourself across six dimensions before you build. Gaps here are the top reasons pilots stall.

**Strategy & scope**

- A high-value workflow is clearly identified
- Success metrics & guardrails are defined
- Executive sponsor & budget secured

**Data & knowledge**

- Relevant data is accessible & permissioned
- A retrieval/knowledge source is in place
- Data quality & freshness are managed

**Platform & infrastructure**

- LLM access & tool integrations available
- Environments for test & production
- Monitoring & logging in place

**Governance & risk**

- Input/output guardrails defined
- Human-in-the-loop & escalation paths
- Audit, bias & compliance review

**Talent & skills**

- Team can build & operate agents
- Upskilling plan for gaps
- Clear ownership & roles

**Operations & scale**

- Incident response & rollback ready
- Cost & latency budgets agreed
- A path from pilot to production

Score 0–3 per dimension. Below 2 in governance or data? Close that gap before scaling.

ROI CHECKLIST & VALUE MODEL

# Make the business case

Four value levers, what to measure, and a simple way to frame return — so your pilot earns its next round of funding.

Value lever	What to measure	Example metric
Time saved	Hours of manual work removed	Cycle time per task
Cost reduced	Lower cost to serve / operate	Cost per transaction
Revenue / throughput	More output or faster delivery	Volume handled per week
Risk reduced	Fewer errors, incidents, fines	Error / incident rate

### ROI in one line

ROI = (annual value created – annual run cost) ÷ total investment. Count build, model/usage, oversight and maintenance in run cost — not just the licence.

### Costs to remember

- Model/token & infra usage
- Human oversight & review time
- Monitoring & maintenance
- Governance & audit overhead

### Quick ROI checklist

- Baseline measured before launch
- Full run cost (not just licence) counted
- Value levers picked & quantified
- Review cadence to track actuals set

48 HOURS ONLY

## Learn to build the business case & the agents

Module 9 covers scaling & ROI; the capstone proves it. This offer is valid for the next 48 hours.

Secure My Seat →

## 32 LEARN-BY-DOING BUILDS · BY CATEGORY

# The full build index — part 1

All 32 hands-on builds, grouped so you can pick by skill area. Each becomes a portfolio artifact.

**A Foundations & Concepts**

- 1 Map a multi-agent system & emergent behaviors.
- 2 Design a self-adaptive agent (meta-learning).
- 3 Model goal-directed autonomy for a decision agent.
- 4 Build a ReAct agent with tools & iteration cap.

**B Architecture & Scale**

- 5 Architect microservices for an agent ecosystem.
- 6 Add fault tolerance to a distributed cluster.
- 7 Load-balance agents & optimise resources.
- 8 Design for high availability & horizontal scale.

**C Integration & Emerging Tech**

- 9 Connect IoT sensors for real-time action.
- 10 Deploy an edge agent for latency & privacy.
- 11 Use blockchain for verifiable interactions.
- 12 Prototype a quantum-inspired optimisation.

**D LLMs & Frameworks**

- 13 Fine-tune an LLM for a domain task.
- 14 Engineer prompts with context management.
- 15 Run an RLHF loop to align responses.
- 16 Build a RAG workflow in LangChain / LangGraph.

**E Coordination & Multi-Agent**

- 17 Build an orchestrator routing to workers.
- 18 Implement consensus across agents.
- 19 Build a negotiation protocol for agents.
- 20 Secure agent messaging with a protocol.

32 LEARN-BY-DOING BUILDS · BY CATEGORY

# The full build index — part 2

## F Enterprise Use Cases

- 21 Build a finance agent (fraud / risk).
- 22 Prototype a healthcare diagnostics support agent.
- 23 Automate a supply chain with predictive agents.
- 24 Design a public-service automation workflow.

## G Governance & Compliance

- 25 Add a governance wrapper with guardrails.
- 26 Run a bias audit & apply mitigation.
- 27 Map a compliance checklist to a regulation.

## H Operations & Reliability

- 28 Stand up monitoring & alerting for live agents.
- 29 Detect anomalies & diagnose agent faults.
- 30 Wire long-term memory with a vector store.

## I No-Code & Capstone

- 31 Build a no-code agent with drag-and-drop tooling.
- 32 Ship the capstone — combine patterns end-to-end.

### Use the index as a menu

Working a finance project? Prioritise categories D, E, F & G. Operating in production? Add H. The capstone (I) pulls it all together into one portfolio piece.

**RECOMMENDED FOR YOU**

### Build the ones your use case needs

Related to your vertical: enroll to access every build, the AI-powered practice platform and mentor support.

[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

The four vertical playbooks in this pack come straight from Module 6.

## 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 & ROI

**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 every module behind these use cases**

Unlock an exclusive enrollment offer and learn to design, govern and ship agentic systems end-to-end.

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

## How to prepare

- Work through the build index by category
- Complete all 12 modules & 32 builds
- 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. Delivering real use cases pulls pay toward the upper percentiles.

**\$136.8K**

US median (≈\$66/hr)

**\$208K**

75th percentile

**\$233.5K**

90th percentile

**56%**

AI-skills premium (PwC)

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 AI-skills wage premium. Indicative; not a guarantee of earnings.

HALF PRICE

## Turn use-case delivery into higher pay

Certify the skills behind these case studies — currently available at half price for new enrollees.

Start Prep →

## CAREER ROADMAP

# From practitioner to AI leader

Delivering use cases like these is how you climb. The Expert certification is the inflection point into higher-paid ownership.

1

**Foundation** · ~\$110k–\$160k

AI/ML basics. ML Engineer, Data Scientist.

2

**Associate** · ~\$160k–\$200k

First agents, RAG. Agentic AI Developer.

3

**Expert** · ~\$200k–\$280k ★Use cases, multi-agent, governance. **GSDC Agentic AI Expert**.

4

**Principal / Architect** · ~\$250k–\$450k

Own enterprise AI design &amp; standards.

5

**Governance & Risk** · ~\$185k–\$275k

AI Governance Lead, Responsible AI Officer.

6

**Executive** · \$350k+

Head of AI / Chief AI Officer.

Pay figures indicative US base; total comp higher with equity. For planning only.

## AI GOVERNANCE ROLES

# The oversight career lane

Every use case in this pack needs governance. Module 7 maps directly to a fast-growing, well-paid set of oversight roles.

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**

## Own 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 every project needs.

[Enroll Risk-Free →](#)

## HIRING TRENDS

# What 2026 employers screen for

Even as overall postings fell, AI-skill demand rose. Employers want proof you can deliver real use cases — exactly what this pack helps you build.

**+7.5%**

AI-skill postings YoY

**-11.3%**

Total postings YoY

**3.5×**

Faster AI-job growth

**Remote**

Global hiring common

**01****Shipped use cases**

Delivered, measurable projects beat generic résumés.

**02****Domain + AI**

Pair industry depth with agent-design fluency.

**03****Governance literacy**

Safety & compliance now table-stakes for senior roles.

Posting trends per PwC 2025 Global AI Jobs Barometer; remote/hiring patterns per 2026 job-market reporting.

## FROM USE CASE TO CAPABILITY

# Map each vertical to what you learn

A quick lookup tying each vertical to the builds, modules and roles that deliver it.

Vertical	Key builds	Modules	Role
Finance	16, 21, 25, 28	4, 6, 7	Agentic AI Engineer
Healthcare	22, 25, 26, 30	4, 6, 7	AI Solution Architect
Manufacturing	10, 23, 28, 29	3, 6, 8	Multi-Agent Systems Eng.
Government	17, 24, 25, 27	5, 6, 7	AI Governance Lead

### Read it as a learning plan

Pick your vertical, do its key builds first, study the listed modules, and target the matching role. The capstone then combines them into proof you can show.

### Common thread

Every vertical needs governance (Module 7) and operations (Module 8). Strong candidates show they can not only build agents but run them safely.

WHAT YOU GET

# Everything to deliver real use cases

**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.



**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

# Turn use cases into outcomes

You have the playbooks, the readiness and ROI checklists, the build index and the syllabus. The next step is building these for real — with mentorship, a capstone, and a recognized credential.

**4 verticals**

Finance, healthcare, manufacturing, government.

**32 builds + capstone**

A portfolio that proves real capability.

**Job support**

Capstone, mentorship & a job-support program.

## FINAL CALL

**Claim your seat in the current cohort**

This is your last prompt in this pack — 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) · Glatbrugg (Switzerland) · Singapore. © 2025 GSDCouncil. All rights reserved. Case studies illustrative; salary & market data indicative; sources cited within.