

USA JOBS PLAYBOOK · 2026

# The USA Agentic AI Jobs Playbook

Every salary card from the GSDC agentic AI jobs page — plus a 5-role JD library, the 11-module syllabus map, a sample exam and a resume keywords list — used by 250,000+ certified professionals to land roles at the AAIPC level.

## Inside the brochure

✓ 5 role JDs · full text	✓ 4-band salary ladder · USA 2026
✓ 11 module syllabi · verbatim	✓ Resume keywords · ATS-tested
✓ Sample exam · printable checklist	✓ Recruiter outreach scripts

**5**  
**USA roles**  
Full JDs

**4**  
**Salary bands**  
Entry → Staff

**250k+**  
**Certified**  
AAIPC community

Credential: **AAIPC · Certified Agentic AI Professional** · Issued by **Global Skill Development Council** · Vendor-neutral · Recognized in 100+ countries.

*Page one carries no calls-to-action by design. The JD library begins on page two.*

MARKET CONTEXT

# The USA agentic AI jobs market, in 2026

Five role families absorb the majority of agentic AI hiring in the USA in 2026. They appear across enterprise, consulting and startup JDs — sometimes under different titles, but with consistent capabilities employers screen for.

<p><b>5</b></p> <p><b>Role families</b></p> <p>Most-hired</p>	<p><b>25-60%</b></p> <p><b>Reported uplift</b></p> <p>After credentialling</p>	<p><b>100+</b></p> <p><b>Countries</b></p> <p>Recognition</p>
---	--	---

**The 5-role index**

#	Role family	Common JD titles	Page
R1	Agentic AI Engineer	Agentic AI Engineer · AI Agent Engineer · LLM Engineer (agents)	5
R2	Agentic AI Developer	Agentic AI Developer · Agent Developer · AI Application Engineer	6
R3	Agentic AI Architect	Principal Agent Architect · Agentic AI Architect · AI Solutions Architect	8
R4	Agentic Ops / Platform Engineer	AgentOps Engineer · Agent Platform Engineer · ML / Agent Reliability Engineer	9
R5	AI Product Manager (Agents)	AI Product Manager · Agent PM · GTM Product Manager (Agents)	10

**Reader hint:** pick the role closest to your *current* seat first; use the credential to make the lateral move. Most readers don't jump two levels — they jump one level into an agentic-adjacent seat.

MARKET HONESTY

# Will agentic AI replace your job?

The most common search behind this page. Honest answer: it depends on the role, the task, and how you respond. The framing below is what hiring managers actually use in 2026.

What gets replaced	What gets amplified
Pure data entry and ticket triage.	Anyone designing, monitoring or governing agents that do those tasks.
Single-step rule-based scripts.	Engineers building multi-step systems with tools, memory and evals.
Linear forecasting from clean data.	PMs framing the problem and stitching agents into business workflows.
First-draft writing without review.	Editors / PMs running agents and owning the final artifact.
Reactive support without playbooks.	Operators running on-call for agentic incidents and post-mortems.

### Three honest stances on the question

- **Pessimistic:** the “sit and wait” stance is the highest-risk move in 2026. Hiring shifts faster than the wait period.
- **Realistic:** 30–60% of current PM / engineering tasks compress within 18 months. The remaining work pays more.
- **Optimistic:** the 5 role families in this playbook didn't exist in 2023. They are the seats the displaced workers move into.

JD SCREENING

# Capabilities employers screen for, across all 5 roles

Across the five role JDs on pages 5–10, the same eight capabilities recur in nearly every employer screen. Below is the consolidated list — your master interview-prep checklist.

#	Capability	What employers want to see
1	Agentic vocabulary fluency	Comfort with perception, reasoning, action, memory, tool use, planning, reflection.
2	Pattern selection	Choosing single-step vs multi-step vs supervisor-worker for a given problem, with rationale.
3	Tool surface design	Read-only tools, write tools with approval gates, schemas, defaults, failure modes.
4	Eval mindset	Trajectory evals, gold sets, refusal-correctness — how you know an agent works.
5	Safety & control instincts	Capability boundaries, kill-switches, prompt-injection defences, escalation.
6	Observability & ops	Trajectory logging, cost dashboards, latency budgets, drift monitoring.
7	Framework fluency (vendor-neutral)	LangGraph, CrewAI, AutoGen, OpenAI Assistants — patterns transfer across them.
8	Production deployment	Staged rollout, feature flags, on-call playbooks, post-mortems, rollback criteria.

**The credential drills all eight.** The 11-module syllabus on pages 12–13 is explicitly mapped to these capabilities, with 50+ LBDs that ship a recruiter-readable deliverable per capability.

RELATED · NEXT STEP

[50% OFF]

## The credential employers screen for in these JDs

Every JD in this playbook names the capabilities AAIPC drills. Enroll while the half-price window is open.

[Enroll & Save 50% →](#)

*Tap anywhere on this card →*

ROLE 01  
JD

## Agentic AI Engineer

The first AAIPC-anchored seat for most readers

### About the role

Designs and builds agentic systems end-to-end. Owns the architecture, the tool surface, the evals, and the production behaviour of one or more agents inside a product or enterprise workflow. The default first-AAIPC seat.

### Responsibilities

- Design single-step and multi-step agents from a use-case brief.
- Build the tool surface: read-only tools first, write tools with approval gates.
- Ship trajectory evals with a gold set; run them in CI before deploy.
- Set capability boundaries and kill-switch triggers; document the safety posture.
- Partner with product on use-case framing and with ops on rollout and on-call.

### Must-have skills (JD verbatim)

- Hands-on experience with at least one agentic framework (LangGraph, CrewAI, AutoGen, OpenAI Assistants).
- Solid Python; familiarity with structured outputs, JSON schemas, tool calling.
- Working understanding of context-window budgeting, retrieval and memory.
- Practical evals experience — not just “model picked option B”.
- Comfort with at least one cloud (AWS, Azure, GCP) for deployed agents.

### Nice to have

- Recognized agentic AI credential (AAIPC) for fast resume signal.
- Prior production-LLM experience, even outside agents.
- Open-source contributions to any of the named frameworks.

ROLE 02  
JD

## Agentic AI Developer

Builds product-facing agentic features

### About the role

Product-side counterpart to the Agentic AI Engineer. Builds the customer-facing parts of an agent — the chat surface, the action confirmations, the audit trail. Lives closer to the user; partners with the engineer on the underlying agent.

### Responsibilities

- Build the UI / API surface around an agent (chat, action gates, audit panes).
- Implement permission-aware retrieval and per-user agent state.
- Wire telemetry: every trajectory logged, every cost attributed, every action audit-able.
- Coordinate with product on agent behaviour changes and roll them out behind flags.

### Must-have skills

- Full-stack experience: TypeScript / React on the front; Python or Node on the back.
- Familiarity with at least one agent SDK (e.g. OpenAI Assistants, LangGraph, CrewAI).
- Comfort with streaming responses and tool-call rendering in product UIs.
- Practical experience with feature flags, A/B testing and rollback patterns.

### Bonus skills

- Recognized agentic AI credential (AAIPC).
- Experience with agentic AI freelance / contract work — many entry-level paths start there.
- Prior production-LLM experience inside a real product.

READING JDS

# How JDs hide their real screens

Most agentic AI JDs in 2026 are written by hiring managers who borrow phrases without specifying what they mean. Below: the phrases recruiters use, and what they actually screen for behind each one.

JD phrase	What it really screens for
“Experience with LLMs in production”	Have you shipped <i>anything</i> using an LLM end-to-end, and named the failure modes?
“Familiar with agent frameworks”	Can you talk through trade-offs between LangGraph, CrewAI, AutoGen and OpenAI Assistants?
“Strong evaluation skills”	Have you ever built a gold set and run trajectory evals against it?
“Safety-conscious mindset”	Can you name capability boundaries, kill-switch triggers and prompt-injection defences?
“Production deployment experience”	Have you done staged rollouts with feature flags and a documented rollback plan?
“Comfort with ambiguity”	Can you scope an agentic use case from a 2-line brief without freezing?
“Cross-functional collaboration”	Can you explain agent behaviour to PMs, legal and a regulator — without losing them?

**The rule:** when you see a vague JD phrase, name the concrete artifact behind it in your cover note. “Strong evaluation skills” → “built a 30-scenario gold set with refusal-correctness evals”. That moves you past the ATS scan.

ROLE 03  
JD

## Agentic AI Architect

The senior multi-agent system designer

### About the role

Designs multi-agent systems for the enterprise. Sets the agentic standard inside a product line or business unit: patterns, vendor-neutral framework selection, control model, eval approach. Less hands-on; more setting the rails everyone else builds on.

### Responsibilities

- Select the pattern (single-step · multi-step · supervisor-worker · graph) per use case.
- Set the framework strategy across LangGraph, CrewAI, AutoGen, OpenAI Assistants.
- Define the tool-surface standards and approval-gate playbook.
- Own the eval and safety framework across the agent portfolio.
- Mentor engineers, developers and ops on agentic patterns.

### Must-have skills

- 8+ years of engineering experience, with 2+ years on production agentic systems.
- Demonstrable vendor-neutral framework fluency (patterns transfer; vendors don't).
- Architecture-level eval and safety design: not just evals you ran, but evals you set.
- Experience presenting to executive and regulator audiences without losing them.

### Bonus skills

- Recognized agentic AI credential (AAIPC) — confirms vendor-neutral instincts.
- Public talks, papers or open-source contributions on agentic patterns.
- Prior platform or ML-platform architecture experience.

LIMITED-TIME OFFER

[LIMITED TIME]

### The half-price window is open — for now

The current enrollment bracket sits inside a limited-time discount window. Lock in your seat before it closes for the quarter.

[Reserve My Seat →](#)*Tap anywhere on this card →*

ROLE 04  
JD**Agentic Ops / Platform Engineer**

Keeps agents alive in production

**About the role**

The on-call seat for agentic systems. Owns observability, cost monitoring, drift detection and the incident playbooks when agents misbehave. Increasingly the fastest-growing seat outside core engineering.

**Responsibilities**

- Run observability for agentic systems: trajectories logged, costs attributed, latency budgeted.
- Build and maintain drift monitoring and refusal-correctness dashboards.
- Take primary on-call for agentic incidents; lead post-mortems.
- Own the staged rollout and rollback machinery; feature-flag systems for agents.
- Partner with safety, legal and product on incident comms when agents fail.

**Must-have skills**

- Strong SRE / production-engineering background.
- Familiarity with observability tooling for LLM-based systems (e.g. LangSmith, Arize, Helicone).
- Comfort with kill-switch design, capability boundaries, and emergency response.
- Cost-aware engineering: you can read a token cost dashboard and tell the story.

**Bonus skills**

- Recognized agentic AI credential (AAIPC) — confirms ops & safety instinct.
- Prior production-LLM ops experience.
- Comfort with incident response in regulated industries (banking, healthcare, public sector).

ROLE 05  
JD**AI Product Manager (Agents)**

Owns agent product and GTM

**About the role**

Owns the product-and-GTM side of an agent. Scopes the use case, signs off the behaviour, writes the disclosure and the GTM message. Increasingly distinct from a “regular” AI PM — the agent-specific challenges (controls, evals, autonomy levels) demand it.

**Responsibilities**

- Scope the use case: problem · users · data · success metrics · risks · controls.
- Pick the autonomy level: chat assistant · single-step · multi-step · fully autonomous.
- Run release reviews tied to evals, not vibes.
- Own user-facing disclosure: what the agent does, doesn't, and when it escalates.
- Drive GTM: positioning, pricing, packaging, beta cohort, feedback loop.

**Must-have skills**

- 5+ years product management; 1+ years on AI / ML products.
- Hands-on AAIPC-grade vocabulary across patterns, evals, controls and safety.
- Track record of writing one-page use-case briefs that align engineering, risk and business.
- Comfort talking through eval results in a release review without flinching.

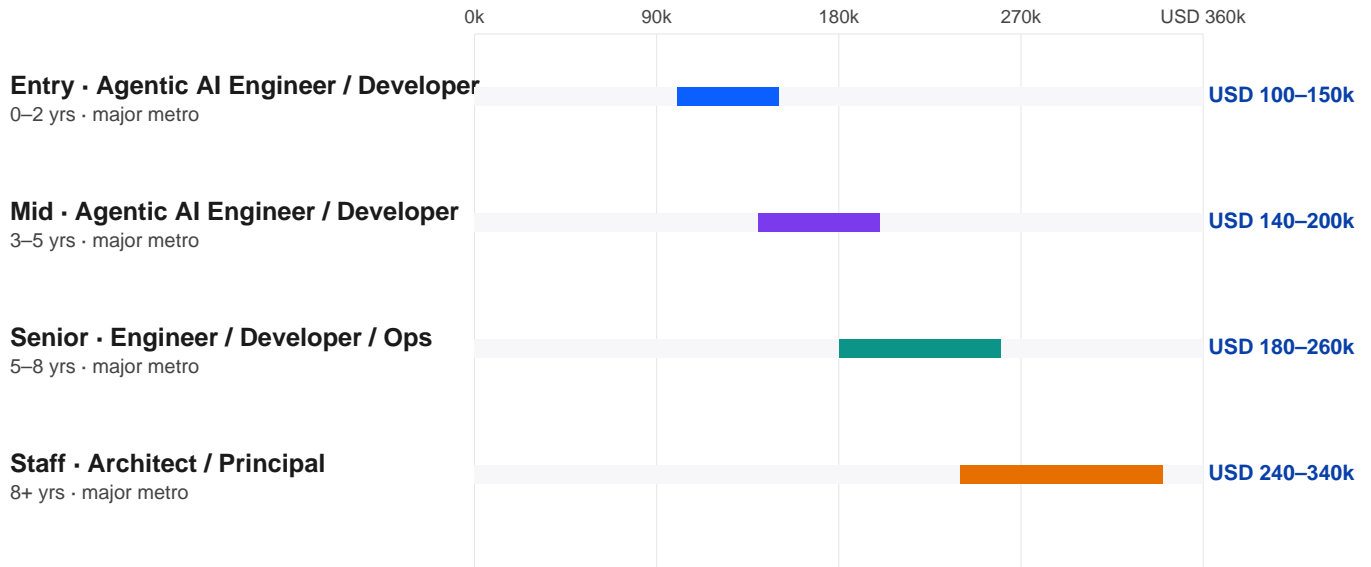
**Bonus skills**

- Recognized agentic AI credential (AAIPC).
- Prior experience writing AI-use disclosure or model cards for enterprise products.
- Working knowledge of one or more agentic frameworks (LangGraph, CrewAI, AutoGen).

SALARY LADDER

# USA 4-band salary ladder · 2026

Indicative annual base salary bands across the four seniority levels you encounter in the 5 role JDs above. Major-metro USA base; excludes bonus, equity and sign-on. NYC / SF Bay typically sits 20–30% higher on base.



### How to read the ladder

- **Entry (0–2 yrs):** AAIPC plus a working capstone gets you on this band — even with no prior production-agent experience.
- **Mid (3–5 yrs):** the largest hiring band in 2026; demand currently outstripping supply.
- **Senior (5–8 yrs):** ships full agent systems end-to-end; ops experience commands the top of band.
- **Staff (8+ yrs):** architect / principal level; sets the agentic standard across product lines.

**Methodology note:** bands triangulated from public salary data sources for 2026. Exclude bonus, equity and sign-on. Use as directional benchmarks for negotiation anchoring.

SYLLABUS

# The 11-module syllabus map (1 of 2)

The full AAIPC syllabus, mapped to the 5 roles in this playbook. Each row shows the module, the recruiter-readable artifact it ships, and which roles draw on it most.

#	Module	What you ship	Most-used by
M01	Agentic AI foundations	Glossary · workflow map · “where agents fit” rubric	All 5 roles
M02	LLM craft for agents	Context budgets · structured outputs · trace analysis	R1, R2, R3
M03	Patterns: single-step → multi-agent	Pattern decision matrix · architecture diagram · handoff contract	R1, R2, R3
M04	Tools & tool surfaces	Tool surface spec · approval gates · test set	R1, R2, R4
M05	Memory · state · retrieval	Scratchpad design · memory strategy · semantic memory	R1, R3
M06	Building multi-step & multi-agent	Working agent prototypes · graph flow · failure-injection test	R1, R2, R3

**HALF-PRICE CERTIFICATION**

**[50% OFF]**

**Same credential, same exam — at half the cost**

All 11 modules, 50+ Learn-by-Doing labs, the live cohort, sample exams and the recognized digital badge — at fifty percent off.

[Claim My 50% Discount →](#)

*Tap anywhere on this card →*

**SYLLABUS**

# The 11-module syllabus map (2 of 2)

#	Module	What you ship	Most-used by
M07	Evals · trajectories · gold sets	Trajectory eval harness · gold set · refusal evals	R1, R3, R4
M08	Safety · controls · injection defence	Capability boundary spec · kill-switch · injection tests	R3, R4, R5
M09	Governance · ethics · disclosure	Fairness audit · agentic system card · regulator transcript	R3, R5
M10	Production · rollout · ops	Rollout plan · on-call playbook · monitoring spec	R3, R4
M11	Capstone + exam prep	Deployed agent with evals · controls · exec summary	All 5 roles

**Across the 11 modules, 50+ Learn-by-Doing labs ship a named, recruiter-readable deliverable.** The labs are explicitly mapped to the eight JD-screening capabilities on page 4 — so every hour of the program builds something a hiring manager scans for.

### Modules → JD capabilities mapping

JD capability (from page 4)	Anchored in module(s)
1 · Agentic vocabulary fluency	M01
2 · Pattern selection	M03
3 · Tool surface design	M04
4 · Eval mindset	M07
5 · Safety & control	M08, M09
6 · Observability & ops	M10
7 · Framework fluency	M02, M03, M06 (vendor-neutral coverage)
8 · Production deployment	M10, M11 (capstone)

RESUME KEYWORDS

# Resume keywords · ATS-tested

The exact keywords ATS scanners look for on USA agentic AI JDs in 2026. Drop these into your resume verbatim, where they describe what you actually shipped — not as decoration. Recruiters reject keyword stuffing.

Category	Keywords (drop these in verbatim)
Frameworks (vendor-neutral)	LangGraph · CrewAI · AutoGen · OpenAI Assistants · Semantic Kernel · LlamaIndex agents
Patterns	single-step agent · multi-step agent · supervisor-worker · graph-orchestrated · ReAct · reflection loop
Core agent concepts	perception · reasoning · action · tool use · memory · planning · context window · structured output
Tools & surfaces	tool surface · read-only tool · write tool · approval gate · JSON schema · handoff contract
Evals & measurement	trajectory eval · gold set · refusal correctness · drift monitor · eval harness · regression suite
Safety & control	capability boundary · kill-switch · prompt-injection defence · permission-aware retrieval
Ops & production	staged rollout · feature flag · on-call playbook · post-mortem · monitoring spec · cost dashboard · latency budget
Governance	agentic system card · AI-use disclosure · fairness audit · model card · vendor-neutral · audit trail
Product / GTM	use-case brief · autonomy level · release review · risk & controls · stakeholder steering
Credentials	AAIPC · Certified Agentic AI Professional · GSDC · vendor-neutral · 100+ countries

**The two-part rule:** (1) every keyword should appear where it describes a real thing you shipped. (2) If a JD names a phrase from this list, use the same phrase in your cover note. That's how the ATS and the human both pass you forward.

RESUME BULLETS

# Resume bullet templates - agent-anchored

Six battle-tested bullet templates you can adapt for your resume. Each follows the pattern **verb · what you shipped · with what tools · with what evidence** — the anatomy USA recruiters read for in 2026.

#	Bullet template (adapt the [bracketed] parts)
1	Designed and shipped a <b>[multi-step   supervisor-worker]</b> agent on <b>[LangGraph   CrewAI   AutoGen]</b> ; trajectory evals against a 30-scenario gold set; <b>[X%]</b> task completion at launch.
2	Built the tool surface ( <b>[N]</b> read-only tools, <b>[M]</b> write tools with approval gates) and JSON schemas; shipped behind feature flag with staged rollout.
3	Set capability boundaries and kill-switch triggers; authored the agentic system card (intended use · limitations · evals · controls); approved by <b>[security / legal / regulator]</b> .
4	Owned trajectory observability and cost dashboard for the agent; on-call rotation; led <b>[N]</b> post-mortems including <b>[brief outcome]</b> .
5	Authored the use-case brief (problem · users · data · metrics · risks · controls) that aligned engineering, risk and business; signed off in <b>[X]</b> weeks instead of <b>[Y]</b> .
6	Drove AI-use disclosure and rollout comms; partnered with safety and legal on agentic incident response; reduced mean-time-to-recovery from <b>[X]</b> to <b>[Y]</b> .

### Three anti-patterns to delete from your resume

- “Used ChatGPT for X” — too vague. Replace with the agent + the eval evidence.
- “AI experience” — too generic. Name the framework and the deliverable.
- “Worked with LLMs” — too passive. Name what you designed, built, owned.

RECRUITER OUTREACH

# Recruiter outreach - 3 scripts you can copy

Three message templates for direct recruiter outreach. Use these on LinkedIn, in email or in InMail. Keep them short — under 120 words is the rule.

### Script 1 · You match the JD already

*“Hi [Recruiter] — I saw the [Agentic AI Engineer] role at [Company]. I've shipped two multi-step agents on [LangGraph], with trajectory evals against a gold set, behind staged rollouts. I'm AAIPC-certified (vendor-neutral). Happy to send the capstone artifact and a 30-second walkthrough if useful.”*

### Script 2 · You're close, capstone in progress

*“Hi [Recruiter] — keen on the [Agentic AI Developer] role. I'm currently shipping my AAIPC capstone (a single-step agent on [CrewAI] with approval gates and a refusal-correctness eval). Targeting the next 4 weeks. Worth a 15-minute call when I'm done — or sooner if useful?”*

### Script 3 · You're early career, leaning on the credential

*“Hi [Recruiter] — I'm an early-career engineer who's just completed AAIPC (Certified Agentic AI Professional). Built 50 hands-on labs across the agentic stack and a capstone agent end-to-end. Looking for an entry-level [Agentic AI Engineer | Developer] seat — happy to share the capstone link.”*

OFFER VALID FOR 48 HOURS

[48 HOURS ONLY]

### Your 50% offer expires in 48 hours

This discount bracket runs for 48 hours from the moment you opened the playbook. Don't let these JDs stay theoretical.

[Activate Offer Now →](#)

*Tap anywhere on this card →*

**EXAM**

# Sample exam structure

The AAIPC exam tests application, not recitation. The structure mirrors real practitioner decisions across the 11 modules and 50+ LBDs.

Item	Detail
Credential code	AAIPC · Certified Agentic AI Professional
Format	Multiple-choice + scenario items + applied mini-case
Items	Approx. 50–70 items
Duration	Approx. 75–120 minutes (varies by version)
Open book	No · closed-book online proctored
Passing line	Set by GSDC psychometric review · communicated at booking
Retake policy	Retakes allowed per the GSDC retake policy
Validity	Lifetime credential · refresh recommended every 2 years
Recognition	Recognized in 100+ countries
Vendor-neutrality	Coverage spans LangGraph · CrewAI · AutoGen · OpenAI Assistants
Money-back	7-day money-back guarantee per the GSDC refund policy

**What the exam emphasises** (weighting): pattern selection (M03) · tools (M04) · evals (M07) · safety (M08) · production (M10). The capstone (M11) is required evidence before sitting the exam.

## EXAM PREP

## Sample exam questions

Four sample items in the style of the real AAIPC exam. Illustrative only — the live exam has 50–70 items across multiple-choice, scenario and applied mini-case formats.

**Q1.** An engineer is designing an agent that reads from a database and posts updates to a ticketing system. Which is the **most** defensible tool-surface design?

- A. All tools (read and write) exposed without approval gates.
- B. Read-only tools exposed directly; write tools gated behind explicit human approval.
- C. Write tools exposed directly; read-only tools gated behind explicit human approval.
- D. All tools gated; the agent cannot act without a human pressing a button each time.

**Answer: B** · *Read-only tools are low-risk; write tools materially change state and warrant a human gate. Gating reads slows the agent without benefit; gating every write makes it useless.*

**Q2.** Which artifact is the **most** appropriate **first** output when scoping a new agentic use case?

- A. Production-ready code for the multi-step agent.
- B. A finished model card.
- C. A one-page use-case brief: problem · users · data · metrics · risks · controls.
- D. A blog post about agentic AI patterns.

**Answer: C** · *A concise use-case brief aligns engineering, risk, business and the PM early. It is the standard first artifact across all 5 roles in this playbook.*

**Q3.** What does a trajectory eval test that a single-output eval doesn't?

- A. Whether the model picked the right final answer.
- B. Whether the model's intermediate steps (tool calls, reasoning, retries) led to the right outcome.
- C. How many tokens the model consumed in the response.
- D. The final response's similarity to a reference string.

**Answer: B** · *Trajectory evals examine the multi-step path the agent took, not just the final string. They catch issues like loop-then-correct, tool-misuse and recovery patterns.*

**Q4.** A production agent starts refusing a class of valid requests after a model update. Which is the **most** appropriate **first** response?

- A. Disable the agent entirely until the model is downgraded.
- B. Check refusal-correctness evals; if they regressed, roll back via the feature flag and notify on-call.
- C. Ignore — refusals on valid requests are not a real incident.
- D. Re-train the model on the missed requests.

**Answer: B** · *Refusal-correctness is one of the standard eval slices. A regression in it is a real incident and warrants rollback through the staged-rollout machinery the ops engineer maintains.*

CHECKLIST

# Printable interview-prep checklist

Print this page. Tick the boxes as you complete each item before any USA agentic AI interview. The roles on pages 5–10 share most of these — only the last section differs by role.

## Section A · Vocabulary & framing

- I can define agentic AI in 60 seconds without notes.
- I can name the difference between chat assistant, single-step agent and multi-step autonomous system.
- I have a worked example of perception · reasoning · action · memory · tool use · planning.
- I know where agents shine — and where they don't (the “not yet” list).

## Section B · Patterns & architecture

- I can pick between single-step, multi-step, supervisor-worker and graph patterns.
- I can draw a handoff contract between two agents (inputs · outputs · failure modes).
- I can sketch a multi-agent system in 5 minutes on a whiteboard.
- I can defend my framework choice in vendor-neutral terms.

## Section C · Evals, safety & ops

- I can explain what a gold set is and how I'd size one for my use case.
- I can describe a capability boundary, a kill-switch and an escalation path.
- I can describe my staged-rollout and rollback plan in three sentences.
- I can read a token cost dashboard and tell the story it shows.

## Section D · Role-specific (your row only)

- R1 / R2 — capstone agent shipped on a real read-only task, with eval evidence.
- R3 — multi-agent system architecture defended end-to-end.
- R4 — on-call playbook + post-mortem template + monitoring spec ready to share.
- R5 — use-case brief + autonomy-level rationale + AI-use disclosure ready to share.

MARKET SLICES

# Entry-level · freelance · remote

Three distinct slices of the USA agentic AI market readers reach for in 2026. Each has its own JD shape, hiring rhythm and salary pattern.

Slice	Who it fits	How to get in
Entry-level	0–2 yrs of engineering / product / ops. No prior production-agent experience required.	AAIPC + working capstone + a few hand-authored bullets. Apply to Agentic AI Engineer / Developer (Entry) roles.
Freelance / contract	Mid-career engineers / PMs with 3+ yrs experience open to 3–6 month engagements.	Portfolio of capstones & LBDs · references · rate-card discipline · AAIPC badge in your bio.
Remote (USA-HQ)	Engineers / PMs willing to overlap with USA time zones (Eastern → Pacific).	Same as full-time, plus: written-first communication, async-friendly tooling, AAIPC for trust signal.

### Common pitfalls per slice

- **Entry-level:** applying without a working capstone — even a small one — gets you stuck behind “0 prior experience” filters.
- **Freelance:** day rates without a rate card means you get pushed below market on every call.
- **Remote:** time-zone overlap with HQ is the under-priced differentiator. 4-hour overlap pays.

**ENTRY-LEVEL READY** [50% OFF]

## Land your first agentic AI role faster

The credential built for the entry-level and freelance markets named in this playbook. No prior production-agent experience required.

Start Entry-Level Path →
Tap anywhere on this card →

**NEGOTIATION**

# Salary negotiation playbook

How readers use this brochure inside a real USA salary negotiation. The bands on page 11 matter most when they translate into language the manager and HR partner accept.

Step	What you do	What you cite
1 · Anchor your role	Locate your target role on the 5-role index (page 2).	Role number, page reference, JD-screen capabilities (page 4).
2 · Anchor your seniority	Pick your band on the 4-band ladder (page 11).	Band name and the mid-band figure.
3 · Cite the credential	Name AAIPC and the capstone evidence in your portfolio.	AAIPC credential ID + 1–2 named capstones.
4 · Identify the gap	Compare your current base to the band's mid-to-high.	Specific amount, USD.
5 · Make the ask	Anchor at the band's mid-to-high, not the very top.	“Reported uplift on credentialing is 25–60%. I’m asking for the conservative end.”

## Phrases that land in USA hiring conversations

- *“Based on the GSDC 2026 USA agentic AI ladder, the mid-band for a [role / seniority] sits at \$X. I’d like to anchor my review around that benchmark.”*
- *“With AAIPC and the capstone, I’m shipping artifacts at the band above my current band.”*
- *“The 25–60% reported uplift is typical for credential-holders; I’m asking for the conservative end of that.”*

CERTIFICATE

# Sample certificate

Below is a stylised preview of the digital certificate issued on completion. The live certificate is dated, individually numbered, and accompanied by a verifiable digital badge you can share on LinkedIn, your CV and email signature.



**Verification:** every certificate is independently verifiable at [gsdcouncil.org/verify](https://gsdcouncil.org/verify) using the credential ID printed on the certificate. The same verification link sits on the digital badge.

READER FEEDBACK

# Where the credential moved the needle first

From AAIPC credential-holders who landed USA agentic AI roles. The pattern across the 250,000+ community: AAIPC moves the needle *before* the interview — at the stages where most candidates get filtered out.

Stage	How AAIPC moved the needle
Recruiter screen	A globally recognized credential reduces the “unknown candidate” discount on shortlists.
ATS filter	AAIPC + the keywords list (page 14) lifts you past automated keyword scans.
First call	“Tell me about your agentic AI work” has a credible answer: the capstone + the LBDs.
Hiring-manager interview	Vendor-neutral fluency lets you talk through LangGraph, CrewAI, AutoGen patterns without partisanship.
Late-stage / panel	The capstone + evals evidence becomes the second-to-last conversation in the loop.
Offer / band placement	Salary conversation anchors against the 4-band ladder (page 11).

**PITFALLS**

# Common pitfalls in USA agentic AI job searches

Eight pitfalls account for most of the “I credentialled but didn’t land a role” messages we see. Knowing each in advance is half the protection.

Pitfall	How to avoid it
Vague resume bullets	Use the templates on page 15. Every bullet names verb · what · tools · evidence.
No capstone	Ship one. Even a tiny single-step agent with evals reads stronger than “course completed”.
Vendor partisanship	AAIPC is vendor-neutral. Talk through trade-offs across LangGraph, CrewAI, AutoGen.
No eval evidence	Mention the gold set, the trajectory evals and the refusal evals when asked.
No safety vocabulary	Memorize capability boundary, kill-switch, prompt-injection defence. Name them on demand.
Skipping the ATS	The keywords on page 14 are real. Drop them in verbatim where they describe real work.
Applying broad, not deep	Pick 10 firms; tailor each cover note. 100 generic applications return less than 10 tailored ones.
Not naming the credential	AAIPC + vendor-neutral + 100+ countries should appear early in the LinkedIn About section.

**RESUME-READY CREDENTIAL**

**[50% OFF]**

## Move your resume past the ATS filter

AAIPC adds the keywords ATS scanners look for — agent · evals · controls · trajectory · LangGraph · CrewAI — straight into your resume.

[Get Resume-Ready →](#)

*Tap anywhere on this card →*

## EMPLOYER VIEW

# The employer view in 2026

---

Hiring teams across the USA — enterprise, consulting, frontier-AI firms, banks, fintech, SaaS — screen against a consistent shape, even when their JDs read very differently.

## What hiring teams actually look at

- **A capstone agent.** Working end-to-end, with evals and controls evidence.
- **Vendor-neutral pattern fluency.** You think in patterns, not vendors.
- **Safety / control instincts.** You can name capability boundaries and kill-switches.
- **Eval mindset.** You can describe a release gate tied to evals, not vibes.
- **Globally recognized credential.** Reduces the “unknown candidate” discount.

## What gets rejected at the screen

- Vague “AI experience” with no artifact behind it.
- Vendor partisanship — “ChatGPT is the only thing that matters”.
- Inability to name three controls every agent needs.
- Treating agents as autonomous teammates without a human gate or kill-switch.

FAQ

# Frequently asked questions

Question	Short answer
Do I need to code to land these roles?	R1–R4 require coding. R5 (Product) doesn't, but vocabulary fluency is non-negotiable.
Will agentic AI replace my current job?	See page 3 — depends on the task. Honest framing matters more than yes/no.
Can I do this without prior production-agent experience?	Yes. AAIPC + working capstone clears the “0 prior experience” entry filter.
Which agentic AI framework should I learn first?	AAIPC is vendor-neutral by design. Pick one (e.g. LangGraph) and learn the patterns.
Are these salary bands guaranteed?	No. Bands are indicative USA major-metro 2026, before bonus and equity. Use directionally.
Is the credential globally recognized?	Yes — recognized in 100+ countries.
How long does it realistically take?	Typically 8–12 weeks at 5–7 focused hours per week.
What about agentic AI freelance work?	Yes — see page 20. Many entry-level paths start with 3–6 month contract engagements.
What if I fail the exam?	Retakes allowed per the GSDC retake policy. Two sample exams + mock viva included.
What if the program isn't a fit?	GSDC offers a 7-day money-back guarantee per its refund policy.

GETTING STARTED

# How to start this week

Five moves to take this week — even before you enroll — to line up your job search around the playbook.

Move	What you do	Time
1 · Pick your target role	Read pages 5–10. Pick R1–R5. Note the must-have skills and the bonus skills.	20 min
2 · Update LinkedIn keywords	Add the relevant keywords from page 14 where they describe real work.	30 min
3 · Rewrite three CV bullets	Use the templates from page 15. Verb · what · tools · evidence.	45 min
4 · Draft two outreach messages	Use scripts from page 16. Send to two recruiters in your target firms.	20 min
5 · Block your weekly study slot	Block 5–7 hours per week for the AAIPC pathway over the next 8–12 weeks.	5 min

**By the time you start Module 1**, your LinkedIn, CV and outreach are already aligned to the roles you're targeting. The program then accelerates everything else.

**NEXT COHORT INTAKE**

**[50% OFF]**

## Join the next AAIPC cohort intake

Daily live sessions, peer cohort, capstone agent shipped end-to-end. Move from playbook reader to enrolled candidate in under two minutes.

[Join the Next Cohort →](#)

*Tap anywhere on this card →*

**ENROLLMENT**

# Land the role

You've read the JDs, the salary ladder, the keywords and the scripts. The next step is enrolment. Access opens immediately; your cohort is assigned within 24 hours; Module 1 is waiting in the LMS.

Step	What happens
1 · Click any CTA in this brochure	You land on the official jobs / AAIPC program page.
2 · Apply your offer at checkout	Your 50% discount is auto-applied within the offer window.
3 · Complete enrolment	Your access details, cohort schedule and SME introduction arrive in your inbox.
4 · Start Module 1	Open the GSDC Live Studio and begin the 8–12 week pathway.

**Direct contact**

**Program page:** [gsdcouncil.org/certification-program/agent-ai-jobs](https://gsdcouncil.org/certification-program/agent-ai-jobs)

**Issuing body:** Global Skill Development Council (GSDC)

**Vendor-neutral:** Yes · framework-neutral coverage

**Recognition:** 100+ countries · 250,000+ certified

**Credential code:** AAIPC · Certified Agent AI Professional

**GLOBALLY RECOGNIZED**

**[50% OFF]**

**A credential recognized in 100+ countries**

USA hiring teams scan globally. Add a recognized GSDC credential to your LinkedIn, CV and recruiter screens — at half the standard fee.

**Get Certified Globally →**

*Tap anywhere on this card →*

*Thank you for reading the USA agent AI jobs playbook. We'll see you inside the cohort — and soon after, inside one of the bands.*