

# What AI Skills Do Security Job Postings Ask For? AI Expectations Went From Almost None to More Than a Third of Postings

AKA Security Research Factory<sup>\*</sup>

AKA Security

[akasecurity.io/research/ai-skills-in-security-postings](https://akasecurity.io/research/ai-skills-in-security-postings)

1 October 2026

AKA-RZ:2610.04v1 · Version 1.0.0

## Abstract

AI asks in security postings went from almost none in 2024 to more than a third in 2026. Senior individual contributors are asked to secure AI more often than peers and managers, executives to govern it.

---

<sup>\*</sup>Written entirely by agents in the Research Factory, within guardrails and controls set by our team. Figures as of 1 October 2026.

## 1 Introduction

In spring 2024, almost no security job posting located in the US asked anything of the hire about AI. In September 2026, 36% of security postings at US-headquartered companies do. What the posting asks for differs by level. Senior individual contributor postings ask the hire to secure AI systems more often than other individual contributor and manager postings do: 18% of them ask for it. Manager postings ask for AI tool use in 18%, and for securing AI in 9%. Executive postings ask for AI governance in 20%, against 5% of senior individual contributor postings.

## 2 The answer at a glance

Table 1: Findings with their confidence and representativeness ratings, on the scales defined in [1].

| Finding                                                                                                                                                                | Confidence                                                                                                                                                                                                                                                                                                             | Representativeness                                                                                    |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| AI expectations went from almost none to 36% of security postings                                                                                                      | <b>High.</b> Three AI models each read every posting, and the rise holds for each of them on its own, for a simple count of AI keywords, at companies hiring in both years, outside technology companies and on newly published postings. The keyword count puts the rise 3.6 percentage points lower than the models. | <b>Medium.</b> The company list leans toward technology companies, the segment that asks for AI most. |
| Senior individual contributor postings ask to secure AI more often than other individual contributor and manager postings (18%)                                        | <b>High.</b> Holds for each model separately. Executive postings ask about as often (25%): that gap is too small to tell apart from no change, and in one model's answers alone executive postings ask more.                                                                                                           | <b>Medium.</b>                                                                                        |
| Manager postings ask for AI tool use in 18%, against 23% of senior postings, and for securing AI in 9%                                                                 | <b>High</b> for the securing gap, which holds for each model. <b>Medium</b> for tool use, where the gap is likely but borderline (−11 to 0 percentage points).                                                                                                                                                         | <b>Low.</b> Manager postings only.                                                                    |
| Executive postings ask for AI governance more often than senior individual contributor postings (20% against 5%)                                                       | <b>Medium.</b> The gap holds for each model, and against manager postings too, where the local model's gap is likely but borderline. The models differ on its size. 245 executive postings.                                                                                                                            | <b>Low.</b> Executive postings only.                                                                  |
| Privacy and AI security roles have the highest AI governance rates (26% and 27%), and the gap between them is too small to tell apart from no change                   | <b>Low.</b> No clear gap. 101 privacy postings.                                                                                                                                                                                                                                                                        | <b>Medium.</b>                                                                                        |
| Among postings open on 1 October, grouped by when they were first published over the prior six months, the local model put AI expectations at 28% to 31% in each group | <b>Low.</b> A snapshot of the postings open on one day, in all countries, not a trend over time.                                                                                                                                                                                                                       | <b>Low.</b>                                                                                           |

## 3 How much AI security postings ask for

Three AI models each read every security posting in both years and marked five kinds of AI expectation on the hire. Two are large hosted models and one is a local model. The table uses the answer at least two of the three agree on.

Table 2: How much AI security postings ask for.

| Kind of AI expectation              | Spring 2024 | September 2026 | Range of likely values for the change (95%), points                                                              | Range across the three models, 2026 |
|-------------------------------------|-------------|----------------|------------------------------------------------------------------------------------------------------------------|-------------------------------------|
| Any AI expectation                  | 2.8%        | 36.4%          |  + <b>30.4</b> to + <b>36.9</b> | 32.7% to 38.2%                      |
| Uses AI tools in their own work     | 0.5%        | 20.2%          |  + <b>17.2</b> to + <b>22.3</b> | 16.7% to 22.3%                      |
| Secures AI, LLM or agent systems    | 1.0%        | 15.5%          |  + <b>12.9</b> to + <b>16.3</b> | 15.1% to 18.1%                      |
| Builds automation or agents with AI | 0.6%        | 11.1%          |  + <b>9.0</b> to + <b>12.0</b>  | 8.3% to 12.0%                       |
| AI governance, risk or compliance   | 0.4%        | 5.6%           |  + <b>4.2</b> to + <b>6.3</b>   | 4.3% to 6.6%                        |

Source: AKA Security analysis of 1,038 US-located security postings from the public LinkedIn Job Postings dataset (5 to 19 April 2024) and 4,587 security postings at US-headquartered companies, any job location, in AKA's job-board data (1 October 2026). A posting can carry more than one kind. A change in bold is a clear change.

Bars show each range of likely values against zero (dashed). Solid: the whole range is above or below zero. Grey: it ends at zero. Hollow: it includes zero.

A simple keyword count of the same descriptions, which does not rely on the AI models, finds terms such as LLM, generative AI, agents, copilot and prompt injection in 1.0% of 2024 postings and 31.0% of 2026 postings.

### 3.1 A direct read replaced an estimate

An earlier version of this analysis counted AI expectations with the local model alone, then corrected for the postings it missed using a fourth, larger model as the reference. That correction put the share near 45%, with a likely range of 39% to 52%. Reading every posting directly with the two hosted models gives 36% to 38%. The local model's count of 33% sits close to both, and agrees with them at a kappa of 0.81 to 0.82 on whether a posting carries any AI expectation. Kappa is a standard agreement score: one is perfect agreement, zero is no better than chance.

## 4 What AI means differs by level

Share of September 2026 security postings at US-headquartered companies, by the level in the title. The answer at least two of the three models agree on.

Table 3: What AI means differs by level.

| Level                         | Postings | Any AI | Secures AI | Uses AI tools | Builds with AI | AI governance |
|-------------------------------|----------|--------|------------|---------------|----------------|---------------|
| Individual contributor        | 1,886    | 31%    | 12%        | 18%           | 9%             | 4%            |
| Senior individual contributor | 2,146    | 41%    | 18%        | 23%           | 13%            | 5%            |
| Manager                       | 310      | 30%    | 9%         | 18%           | 8%             | 5%            |
| Executive                     | 245      | 45%    | 25%        | 14%           | 9%             | 20%           |

Source: AKA Security analysis of 4,587 security postings at US-headquartered companies (1 October 2026), labelled by the three models. In the 1,038 US-located postings from spring 2024, every cell in the five AI columns sat between 0% and 8%.

### 4.1 Senior individual contributor postings ask to secure AI

Senior individual contributor postings ask to secure AI more often than other individual contributor and manager postings: 18%, against 12% and 9%. Executive postings ask as often or more, at 25%: that gap is too small to tell apart from no change, and in one hosted model's answers alone executive postings ask more (−17 to −2 points). Fifth Third Bank's Principal Information Security Architect posting asks for "es-

tablishing enterprise architectural patterns and standards for prompt-layer protections, tool-execution guardrails, agent identity, observability, and monitoring of agent planning and tool usage." A founding security engineer posting at Withpace names "an emerging category we care deeply about: designing secure AI and agentic systems."

#### 4.2 Manager postings ask to use it

Managers are asked to use AI tools in their own work in 18% of postings, against 23% for senior individual contributors. The difference, -11 to 0 points, is likely but borderline. They are asked to secure AI less often: 9% against 18%, a gap of 5 to 13 points that holds for each of the three models separately. Vanta's Manager, Security Operations posting describes "a view of security that is AI-first, human-centric, and trust-based" and asks the hire to "leverage AI to improve efficiency."

#### 4.3 Executive postings ask to govern it

AI governance appears in 20% of executive postings against 5% of senior individual contributor postings. The gap is likely 9 to 23 points. Over manager postings the lead is likely 8 to 24 points. That lead is clear for each hosted model taken separately, and likely but borderline for the local model. The two hosted models each put executive governance at 21%. Executive postings carry any AI expectation in 45%, against 41% for senior individual contributor postings, a gap too small to tell apart from no change. The local model read 33% of executive postings as carrying one. DeepHealth's Director of AI Governance posting describes "managing DeepHealth's AI Governance program" for "safe, legal, and ethical use of AI across the organization."

### 5 Where each kind of AI lands by specialty

The same split shows up across specialties in 2026. Dedicated AI security roles carry the most AI of any specialty, 95% of 333 postings, and 70% ask the hire to secure AI. Privacy postings carry an AI expectation in 46% of 101 postings, against 42% in application security (AppSec), a gap too small to tell apart from no change.

Privacy's AI governance rate, 26%, sits next to 27% for AI security roles, with no clear gap between them. Both lead every other specialty, and AI security leads the third-highest specialty by 6 to 20 points. Securing AI reaches 21% in AppSec (730 postings) and 19% in offensive security (135). Detection and response postings ask to use AI tools (26% of 484) and rarely to secure AI (5%). Governance, risk and compliance (GRC) postings split between using tools (21% of 309) and governance (14%).

Source: AKA Security analysis of 4,587 security postings at US-headquartered companies (1 October 2026). Specialty comes from the title. The answer at least two of the three models agree on.

### 6 Where the data comes from

The 2024 postings come from the LinkedIn Job Postings dataset on Hugging Face ([datastax/linkedin-job-listings](#) [2]): postings located in the US, listed between 5 and 19 April 2024. The 2026 postings come from AKA's daily read of company job boards, limited to US-headquartered companies, with any job location. Postings in both years are counted the same way: the series' security title rules, a local model's read of whether each posting is a security role, and one posting kept per company and job title. Job aggregators, which repost other employers' jobs, are left out in both years. The 2026 postings are those open on or after 23 September 2026.

Level comes from the title (executive: chief, CISO, VP, head of, director; manager: manager or supervisor, excluding program, project, product and account managers; senior: senior, staff, principal, lead). The five kinds of AI expectation come from the same instructions given to all three models. The likely ranges (95%) come from resampling the data many times, a whole company at a time, because one company's postings tend to resemble each other. They reflect which postings happened to be collected, not the differences in source, location and season between the two years.

These figures measure what the security postings in this data ask of the hire, at every level. They are different from the homepage figure on agent, LLM, MCP (the protocol agents use to connect to tools) or prompt-injection work in security engineering openings at AI-first companies. That figure is a keyword count limited to the responsibilities and requirements sections and to engineering openings, and it compares groups of companies.

## **7 How this answer was stress-tested**

Tests marked 23 September were run on that day's data with the local model's labels. The rest were run on the 1 October data in this piece.

Table 4: How this answer was stress-tested.

| Test                                                                                                                                                                                                           | What it found                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hosted-model labels on both years.</b> The two hosted models each read all 5,625 postings for the five kinds of AI expectation.                                                                             | The two hosted models agree at a kappa of 0.91 on any AI expectation and 0.85 to 0.87 on each kind. The local model agrees with them at 0.81 to 0.82 on any AI and 0.68 to 0.77 on the kinds. The level gaps the text states hold for each of the three models separately. On securing AI, the gap between senior and executive postings is too small to tell apart from no change for the local model and likely but borderline for one hosted model, and in the other hosted model's answers executive postings ask more. On executives, the local model read 33% against 45% for the answer at least two models agree on. |
| <b>A second instrument.</b> AI keywords counted in every description.                                                                                                                                          | 1.0% of 2024 postings, 31.0% of 2026 postings.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Same companies.</b> 74 companies appear in both years. Their 2024 postings are located in the US. Their 2026 postings can be anywhere.                                                                      | Any AI expectation 4.3% → 37.5%.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Company mix.</b> 2026 postings outside technology companies. 2024 without staffing firms and government contractors. The 2024 data have no industry field, so the 2024 cut removes those employers by name. | 26.5% of 2026 non-technology postings carry an AI expectation. 2024 without staffing and contractors: 3.0%.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Open postings against new ones.</b> 2026 postings first published in the last 30 days.                                                                                                                      | 31.3%, against 32.7% for all open postings in the same local-model count.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description length.</b> 2026 descriptions run 60% longer. Measure repeated on postings of 3,000 to 6,000 characters in both years.                                                                          | 3.6% → 29.2%.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Scope check of every posting, both years.</b> A local model read all 1,405 2024 postings and 5,196 2026 postings that pass the title rules.                                                                 | It ruled 26% of 2024 postings and 12% of 2026 postings not security roles, and they were removed before any AI measure.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Seniority rule audit.</b> The corpus's own seniority tag compared with a stricter title rule.                                                                                                               | The tag counted any title containing "Management" as a manager. 184 of 495 "manager" postings were engineers ("Security Engineer, Vulnerability Management"). Fixed before this piece. The manager rate of securing AI fell from 4.6% to 3.9% in the local-model count. (23 September)                                                                                                                                                                                                                                                                                                                                       |
| <b>Posting age.</b> 2026 postings in all countries, grouped by first-published date.                                                                                                                           | 28% to 31% in each group from the last six months. Postings open longer than six months carry less (23%), a likely but borderline difference, but those are the postings that stayed open, and cannot be read as a trend.                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Calibration against a fourth model.</b> A fourth, larger model labelled 180 postings, sampled across groups, to score the local model.                                                                      | Of the postings that model marked with no AI expectation, the local model agreed on 96% to 100%, and it caught 40% to 67% of the AI expectations that model marked. The correction built on those rates produced the 45% estimate. The direct read by the two hosted models above replaced it. (23 September)                                                                                                                                                                                                                                                                                                                |
| <b>A fresh run on 28 September.</b> The full pipeline rebuilt from that day's data.                                                                                                                            | Local-model figures against 25 September: any AI expectation unchanged at 33.9%, manager 33.2% → 33.6%.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

8 What would strengthen this answer

Table 5: What would strengthen this answer.

| Test                                                                         | What it would settle                                                                                                                       | Why it has not run                                                                            |
|------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| The calibration model as a further reader on all 5,625 postings              | Whether that model reads AI expectations more broadly than the two hosted models, the gap between the 45% estimate and the 36% direct read | Possible now. Deferred: the level and specialty findings already hold for both hosted models. |
| Larger manager and executive samples                                         | Narrower likely ranges for the manager and executive rows (310 and 245 postings)                                                           | Collection widened to leadership titles on 18 and 22 September. The samples grow daily.       |
| A re-read of our own data in three to six months                             | A direct trend under stable collection                                                                                                     | The earliest clean read is late December 2026.                                                |
| Licensed longitudinal postings (Lightcast, Revelio Labs, Indeed Hiring Lab)  | Year-by-year AI expectations from 2022, with seasonality                                                                                   | Paid data.                                                                                    |
| Postings from government, managed-security firms and LinkedIn-only employers | How far the 36% holds beyond technology-leaning employers                                                                                  | Needs collection the current data does not cover.                                             |

Acknowledgements

Question asked by [Lenny Zeltser](#).

References

[1] AKA Security Research Factory. *Security Hiring Research: What Job Postings Say About Security Work*. AKA Security, 2026. [akasecurity.io/research/security-hiring-research](https://akasecurity.io/research/security-hiring-research)

[2] DataStax. *LinkedIn Job Postings* (datastax/linkedin\_job\_listings). Hugging Face dataset, 2024. [huggingface.co/datasets/datastax/linkedin\\_job\\_listings](https://huggingface.co/datasets/datastax/linkedin_job_listings)

Cite as

```
@techreport{aka-rz-2026-04,
  title      = {What AI Skills Do Security Job Postings Ask For? AI Expectations Went From Almost None to More Than a Third of Postings},
  author     = {{AKA Security Research Factory}},
  institution = {AKA Security},
  type       = {Research report},
  number     = {AKA-RZ:2610.04v1},
  version    = {1.0.0},
  year       = {2026},
  month      = oct,
  url        = {https://akasecurity.io/research/ai-skills-in-security-postings}
}
```