What AI Skills Do Security Postings Ask For, and How Do They Differ by Level? AI Asks Went From Almost None to More Than a Third of Postings

AKA Security Research Factory*
AKA Security
akasecurity.io/research/ai-skills-in-security-postings

9 October 2026
AKA-RZ:2610.04v1 · Version 1.1.0

Abstract

US security job postings, spring 2024 against fall 2026: AI expectations went from almost none to more than a third. Senior individual contributors are asked to secure AI, executives to govern it.

*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, 2.8% of 1,038 US security job postings asked anything of the hire about AI. In fall 2026, 36.4% of 4,587 security postings at US-headquartered companies did.

In fall 2026, 20.2% of those postings asked the hire to use AI tools in their own work. Securing AI, LLM or agent systems appeared in 15.5%, building automation or agents with AI in 11.1%, and AI governance, risk or compliance in 5.6%.

The kind of AI asked for differs by level. In fall 2026, 18% of senior individual contributor postings asked the hire to secure AI, against 12% of individual contributor postings and 9% of manager postings. Executive postings asked for AI governance in 20% of cases.

These findings compare two sources: a public dataset of US postings from spring 2024, and AKA's read of company job boards in 2026. They apply best to US-headquartered companies on modern job boards, which lean toward technology, and less to government, staffing and managed-security firms.

2 The answer at a glance

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

Finding	Confidence	Representativeness
Security postings with any AI expectation went from 2.8% in spring 2024 to 36.4% in fall 2026.	**High**. The rise holds at the 74 companies that posted security jobs in both years, so it is not just new companies in the data. It holds with tech companies left out too. That matters because the 2026 data leans toward tech.	**Medium**. Best fit: US employers that lean toward tech. Less certain for government, staffing and managed-security firms.
In 2026, 20.2% of security postings ask the hire to use AI tools in their own work. 15.5% ask them to secure AI systems.	**High**. All five kinds of AI ask rose, each by too much to be chance.	**Medium**. Best fit: US employers that lean toward tech. Less certain for government, staffing and managed-security firms.
Senior individual contributor postings ask the hire to secure AI more often than individual contributor postings: 18% against 12%.	**High**. The gap is too large to be chance.	**Medium**. Level currently from the job title and companies use titles differently. We are recalibrating all roles into equivalent levels, that data will be available in future editions.
Senior individual contributor postings ask the hire to secure AI more often than manager postings: 18% against 9%.	**High**. The gap is too large to be chance.	**Low**. Rests on 310 manager postings, too few to speak for all security managers.
A gap in AI tool use between manager and senior individual contributor postings, 18% against 23%, is too close to call.	**Medium**. Managers may ask for AI tools a little less often, but the gap is too small to be sure.	**Low**. Rests on 310 manager postings, too few to speak for all security managers.
Executive postings ask for AI governance more often than senior individual contributor and manager postings: 20% against 5%.	**Medium**. Executives clearly ask for governance more often. How much more depends on how the postings are read.	**Low**. Rests on 245 executive postings, too few to speak for all security executives.
Privacy and AI security postings ask for AI governance most often. The gap between them is too close to call.	**Low**. Privacy and AI security are too close to tell apart. Needs more postings.	**Medium**. Specialty comes from the job title, which companies use differently.
A gap in AI asks between older and recent open postings is too close to call.	**Low**. Based on postings open on a single day. Needs a second read later.	**Low**. Older postings are the ones that stayed open, so they are not a fair sample.

3 How much AI security postings ask for

Three AI models each read every posting and marked whether it asks something of the hire about AI, and of which kind. The answer used is the one at least two of the three agree on. A posting can carry more than one kind.

Table 2: How much AI security postings ask for.

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

Source: AKA Security analysis of 1,038 US security postings from a public job-postings dataset (listed 5 to 19 April 2024) and 4,587 postings at 831 US-headquartered companies in AKA's job-board data (open on or after 23 September 2026), read by three AI models.

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.

The rise holds without any AI model. A keyword count of AI terms, such as LLM, generative AI, agents, copilot and prompt injection, went from 1.0% of US security postings in spring 2024 to 31.0% in fall 2026. That rise, +30.0 points, is 3.6 points smaller than the three-model rise.

It also holds at the same 74 companies in both years, where any AI expectation went from 4.3% to 37.5%. Outside technology companies, 26.5% of 2026 postings carry an AI expectation, against 3.0% in 2024 postings without staffing firms and government contractors.

In 2024, each of the four kinds of AI expectation sat between 0% and 8% at every level.

4 What AI means differs by level

Table 3: What AI means differs by level.

Level	Postings, 2026 (companies)	Any AI	Secures AI	Uses AI tools	Builds with AI	AI governance
Individual contributor	1,886 (557)	31%	12%	18%	9%	4%
Senior individual contributor	2,146 (616)	41%	18%	23%	13%	5%
Manager	310 (181)	30%	9%	18%	8%	5%
Executive	245 (131)	45%	25%	14%	9%	20%

Source: AKA Security analysis of 4,587 security postings at US-headquartered companies in AKA's job-board data (open on or after 23 September 2026), read by three AI models. Level is read from the title.

4.1 Senior individual contributor postings ask to secure AI

In fall 2026, 18% of senior individual contributor postings at US-headquartered companies asked the hire to secure AI systems. Among individual contributor postings, 12% did. The range of likely values for the gap is +4 to +9 points, and it is clear under each of the three AI models.

Against manager postings, the gap is 18% against 9%, with a range of +5 to +13 points. Executive postings asked to secure AI in 25% of cases. The gap between senior individual contributor and executive postings, −15 to 0 points, is too close to call.

4.2 A gap in AI tool use for managers is too close to call

In fall 2026, 18% of security manager postings at US-headquartered companies asked the hire to use AI tools, against 23% of senior individual contributor postings. The range for the gap, −11 to 0 points, reaches zero, so the gap is too close to call.

Manager postings carry any AI expectation in 30% of cases, building with AI in 8% and AI governance in 5%.

4.3 Executive postings ask to govern AI

In fall 2026, 20% of security executive postings at US-headquartered companies asked for AI governance, risk or compliance work. Senior individual contributor and manager postings did in 5% of cases each. The gap against senior individual contributors is +9 to +23 points, and against managers +8 to +24 points.

Each AI model agrees on the direction, but the models disagree on the size. The two hosted models each put the executive governance rate at 21%, and the local model at 11%. That is why this finding is rated Medium.

Executive postings carry any AI expectation in 45% of cases, against 41% for senior individual contributors. That gap, −3 to +13 points, is too close to call.

5 Where each kind of AI lands by specialty

Table 4: Where each kind of AI lands by specialty.

Specialty	Postings, 2026 (companies)	Any AI	Secures AI	Uses AI tools	Builds with AI	AI governance
AI security	333 (200)	95%	70%	28%	36%	27%
Privacy	101 (70)	46%	10%	22%	7%	26%
Application security	730 (367)	42%	21%	25%	12%	2%
Offensive security	135 (93)	38%	19%	25%	13%	1%
Detection and response	484 (244)	35%	5%	26%	13%	2%
GRC	309 (187)	30%	4%	21%	8%	14%

Source: AKA Security analysis of 4,587 security postings at US-headquartered companies (fall 2026), read by three AI models. Specialty is read from the title.

AI governance appears most in AI security postings (27%), privacy postings (26%) and GRC postings (14%). The gap between privacy and AI security, −13 to +12 points, is too close to call. Privacy leads all other specialties on governance by a range of +10 to +33 points.

6 Where the data comes from

The 2024 postings come from a public job-postings dataset: 1,038 US-located security postings listed 5 to 19 April 2024. The 2026 postings come from AKA's daily read of 2,194 company job boards on 16 applicant tracking systems: 4,587 security postings at 831 US-headquartered companies, any job location, open on or after 23 September 2026. Collection began on 31 July 2026.

Both years use the same security title rules. A local model rules whether each posting is a security role, one posting is kept per company and title, and job aggregators are left out. 2024 is US by job location and 2026 is US by company headquarters, so the comparison mixes change over time with differences in source, location and season.

Until the evening of 21 September 2026, the collector skipped titles containing analyst, auditor, consultant, CISO, director, head of, specialist and "manager, security". The 2026 side is limited to postings open

on or after 23 September 2026.

The two hosted models agree with each other 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 expectation, and 0.68 to 0.77 on each kind. Ranges are found by resampling whole companies, not single postings. Postings state what a company says it wants. They are not headcount, and not proof of hiring.

7 How this answer was stress-tested

The main result is a rise in any AI expectation from 2.8% to 36.4%, with a range of +30.4 to +36.9 points.

Table 5: How this answer was stress-tested.

Test	What it found
Each model alone.	Hosted model A puts 2026 at 38.2%, hosted model B at 36.1% and the local model at 32.7%. Each range clears zero.
Keyword count, no model.	1.0% to 31.0%, range +27.1 to +33.1. Same direction, smaller by 3.6 points.
Same companies. The 74 companies with postings in both years.	4.3% to 37.5%, range +22.8 to +43.3. Same direction, smaller by 0.4 points.
Company mix. 2026 postings outside technology companies, against 2024 without staffing firms and government contractors.	3.0% to 26.5%, range +19.8 to +27.5. Same direction, smaller by 10.1 points.
Newly published postings. 2026 postings first published in the last 30 days, local-model count.	31.3% for new postings against 32.7% for all open postings. The change clears zero.
Postings of similar length. 2026 descriptions run 60% longer, so the test was repeated on postings of 3,000 to 6,000 characters.	3.6% to 29.2%, range +21.4 to +29.7. Same direction, smaller by 8.0 points.
Scope check of every posting.	A local model ruled 26% of 1,405 postings in 2024 and 12% of 5,196 in 2026 not security roles. They were removed before any AI measure.
Calibration of the local model. A fourth, larger model labelled 180 postings sampled across groups.	Of postings with no AI expectation, the local model marked 96% to 100% as none. It caught 40% to 67% of the AI expectations. An estimate built on those rates puts any AI expectation at 45%, with a range of 39% to 52%.
Posting age. 2026 postings open on 1 October, all countries, by when they were first published.	28% to 31% for each group published in the last six months, and 23% for postings open longer. The gap, −12 to 0 points, is too close to call.

8 What would strengthen this answer

Table 6: What would strengthen this answer.

Test	What it would settle	Why it has not run
The fourth, larger model as a further reader on all 5,625 postings	Whether the calibration-based estimate of 45% or the three-model result of 36.4% is closer	
Larger manager and executive samples	The level gaps, which rest on 310 manager and 245 executive postings	
A later re-read	A direct trend under the 2026 collection method	
Postings from government, managed-security firms and employers that post only on one large professional network	The employers the company list under-covers	

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

Cite as

```
@techreport{aka-rz-2026-04,
  title       = {What AI Skills Do Security Postings Ask For, and How Do They Differ by Level? AI Asks 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     = {1.1.0},
  year        = {2026},
  month       = oct,
  url         = {https://akasecurity.io/research/ai-skills-in-security-postings}
}
```