How Does Security Hiring at AI-First Companies Differ From Everyone Else? More AI Engineering Openings per Security Opening

AKA Security Research Factory*
AKA Security
akasecurity.io/research/security-hiring-at-ai-first-companies

20 October 2026
AKA-RZ:2610.07v1 · Version 1.1.1

Abstract

Job boards and security postings, July to October 2026: AI-first companies post more AI engineering openings per security opening, name SOC 2 about twice as often, and their AI engineering postings name a security duty less often.

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

1 Introduction

AI-first companies post more AI engineering openings for every security opening. In a census of complete job boards on 1 October 2026, the 196 AI-first boards carried 1.27 AI engineering openings per security opening. The 1,397 other boards carried 0.65.

Their security postings also read differently. Of security postings collected from 31 July to 1 October 2026, 18.8% at AI-first companies named SOC 2, against 9.1% at everyone else. At AI-first companies, 25.7% of security engineering openings named agent, LLM, MCP or prompt-injection work, against 17.0% elsewhere.

In the other direction, AI engineering postings at AI-first companies name a security duty less often: 3.5%, about 1 in 29, against 13.7%, about 1 in 7, elsewhere.

These findings come from security postings collected from 31 July 2026 and a one-day census of 1,593 complete job boards on 1 October 2026. They lean toward technology companies on modern job boards.

2 The answer at a glance

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

Finding	Confidence	Representativeness
More AI engineering openings per security opening at AI-first companies: 1.27 AI engineering openings for each security opening, against 0.65 at everyone else.	**High**. Still holds when removing the three AI-first companies with the most engineering openings, so a few large companies don't drive it.	**Medium**. Fits tech-leaning employers on modern job boards. Based on one day of job-board counts.
More engineering openings per security engineering opening at AI-first companies: 7.0 at the typical AI-first company, against 5.0 at the typical company elsewhere.	**Medium**. Clear when the two groups' company medians are compared, but not when all openings are pooled together, so it depends on how you count.	**Medium**. Fits tech-leaning employers on modern job boards. Based on one day of job-board counts.
Security engineering openings at AI-first companies name agent, LLM, MCP or prompt-injection work more often: 25.7% against 17.0%.	**Medium**. The gap is clear, but the data hasn't yet separated security-for-AI vs AI-for-security	**Medium**. Fits tech-leaning employers on modern job boards.
SOC 2 is named in 18.8% of security postings at AI-first companies, against 9.1% at everyone else.	**High**. Holds with the largest AI-first employer left out, so one company doesn't drive it.	**Medium**. Fits tech-leaning employers on modern job boards.
AI security is a larger share of security postings at AI-first companies: 11.0% against 5.9%. The gap is likely but borderline.	**Medium**. With the largest AI-first employer left out, the gap is no longer clear, so one company matters.	**Medium**. Fits tech-leaning employers on modern job boards.
The EU AI Act is named in 3.9% of AI-first security postings, against 1.0%. The gap is likely but borderline.	**Medium**. Likely, but only just, even with the largest employer left out.	**Medium**. Fits tech-leaning employers on modern job boards.
AI engineering postings at AI-first companies mention security less often: 3.5%, about 1 in 29. Against other companies where AI engineering roles mention security as a responsibility 1 in 7 or 13.7%	**High**. This is also clear with the largest company outlier taken out.	**Low**. This is out of a pool of 115 AI engineering postings at AI-first companies.
AI-first security postings firmly require a degree about half as often: 8.5% against 16.6%. The gap is likely but borderline. Added after the original run, 20 October 2026.	**Medium**. The gap clears zero at 95% but not at 99%. With the largest AI-first employer left out, the 95% range only just clears zero.	**Low**. Covers the 62% of AI-first and 46% of other postings with an agreed answer, under the entry-level title rules.
A gap in coding-agent asks in engineering postings, 7.7% against 10.9%, is too close to call.	**Medium**. AI-first companies may ask for coding agents less often, but the gap is too small to be sure.	**Medium**. Fits tech-leaning employers on modern job boards.

3 The numbers

Table 2: The numbers.

Measure	AI-first	Everyone else	Likely range of the gap (95%)	Base, AI-first against everyone else
AI engineering postings per security posting	1.27	0.65	+0.40 to +0.88	196 and 1,397 job boards
Engineering openings per security engineering opening, company median	7.0	5.0	+0.7 to +3.7	113 and 741 companies
Security postings naming agent, LLM, MCP or prompt-injection work	20.0%	11.7%	+3.8 to +13.4 points	1,412 and 13,751 postings
The same, security engineering openings only	25.7%	17.0%	+1.5 to +16.2 points	628 and 5,829 postings
Security postings in the AI security domain, 99% range	11.0%	5.9%	+0.3 to +9.6 points	1,151 and 10,869 postings
Security postings firmly requiring a degree, added 20 October 2026	8.5%	16.6%	−13.7 to −1.5 points	873 postings at 88 companies and 6,356 at 766
Security postings naming SOC 2	18.8%	9.1%	+3.1 to +18.8 points	1,412 and 13,751 security postings
Security postings naming the EU AI Act, 99% range	3.9%	1.0%	+0.8 to +5.7 points	1,412 and 13,751 security postings
Security postings naming ISO 42001, 99% range	3.0%	1.3%	+0.1 to +4.0 points	1,412 and 13,751 security postings
Security postings naming ISO 27001, 99% range	17.7%	13.9%	−2.4 to +12.6 points	1,412 and 13,751 security postings
Security postings naming the NIST AI Risk Management Framework, 99% range	2.2%	1.9%	−1.5 to +2.9 points	1,412 and 13,751 security postings
AI engineering postings naming a security duty	3.5%	13.7%	−15.9 to −4.0 points	115 and 439 postings
Engineering postings naming coding agents or AI coding tools	7.7%	10.9%	−6.6 to +0.8 points	573 and 2,313 postings

Source: AKA Security analysis of 15,163 security postings (1,412 at 149 AI-first companies), a census of 1,593 complete job boards (196 AI-first) and a sample of 2,886 engineering postings, 1 October 2026. The SOC 2 range is at 99%.

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.

4 More engineering openings per security opening

On 1 October 2026, the 196 complete job boards of AI-first companies carried 1.27 AI engineering openings for every security opening. The 1,397 other complete boards carried 0.65. The range of likely values for the gap is +0.40 to +0.88. With the three AI-first boards that have the most engineering openings removed, the figure is 1.29 against 0.65.

The same census counts engineering openings per security engineering opening. At the median AI-first company it was 7.0, against 5.0 at the median of everyone else. That comparison covers 113 AI-first and 741 other companies with both kinds of opening on a complete board. The range for the gap is +0.7 to +3.7.

That finding is rated Medium. Pooled across all openings rather than taken per company, the figure is 11.3 against 9.7, and the range, −1.01 to +4.80, includes zero. With each company weighted once in a simple average, it is 11.6 against 9.0, the same direction.

5 Security postings name more AI work

Of security postings collected from 31 July to 1 October 2026, 20.0% at AI-first companies named agent, LLM, MCP or prompt-injection work in their responsibilities or requirements. Elsewhere, 11.7% did. Among security engineering openings only, the shares are 25.7% at AI-first companies and 17.0% elsewhere.

The count mixes two kinds of work: securing AI, and using AI to do security work. They cannot be separated in this measure. It was added after the planned keyword count failed its precision check, with 13 of 20 matches genuine. For both reasons, the finding is rated Medium, one step below the scale.

AI security, as its own domain, is 11.0% of labelled security postings at AI-first companies, against 5.9% elsewhere. The 99% range for that gap is +0.3 to +9.6 points, which is likely but borderline. With the largest AI-first employer removed, the range reaches zero.

6 Degrees are firmly required less often

Added after the original run, 20 October 2026.

AI-first security postings firmly require a degree about half as often: 8.5% against 16.6%, a likely but borderline gap. The difference is −8.2 points, with a 95% range of −13.7 to −1.5. At 99% the range, −14.9 to +1.1, includes zero. With the largest AI-first employer removed, the 95% range is −14.0 to −1.0, so the upper end sits within a point of zero.

"Firmly required" is the answer at least two of the three models agree on as "required with no alternative", the same measure the entry-level piece uses. Only postings with an agreed answer count: 873 postings at 88 AI-first companies, 62% of the 1,412, and 6,356 postings at 766 other companies, 46% of the 13,751. The measure follows the entry-level title rules, so it describes that subset and not all 15,163 postings. Each model on its own gives the same direction: 9.6%, 8.5% and 8.6% at AI-first companies against 15.5%, 16.3% and 17.3% elsewhere.

7 Compliance vocabulary leans toward SOC 2

Of security postings collected from 31 July to 1 October 2026, 18.8% at AI-first companies named SOC 2, against 9.1% at everyone else. That is about twice as often. The 99% range for the gap is +3.1 to +18.8 points, and it holds with the largest AI-first employer removed.

AI-first security postings also name the EU AI Act more often, 3.9% against 1.0%, and ISO 42001 more often, 3.0% against 1.3%. Both gaps are likely but borderline at 99%. For ISO 27001 and the NIST AI Risk Management Framework, a gap is too close to call.

8 AI engineering postings less often name a security duty

Of AI engineering postings at AI-first companies, 3.5% named a security duty, about 1 in 29. At other companies, 13.7% did, about 1 in 7. The range for the gap is −15.9 to −4.0 points. The AI-first figure rests on 115 postings at 64 companies.

The comparison cleared zero in each run from 25 September to 1 October 2026. It failed on 23 September.

Across all companies, AI engineering postings name a security duty in 11.6% of cases, against 18.4% for other engineering postings. That gap is likely but borderline.

9 Measures with no clear gap

- **Coding agents.** Engineering postings naming coding agents or AI coding tools: 7.7% at AI-first companies, 10.9% elsewhere. The range, −6.6 to +0.8 points, includes zero.

- **Any security posting open.** 64.8% of the 196 complete AI-first boards had a security posting open on 1 October, against 71.7% of other boards. The gap is too close to call, but only just, and it changes with board size. On boards with 44 or more open postings, the shares are 89.1% and 89.3%.
- **A security engineering opening open.** 59.2% of AI-first companies against 57.5% of others.
- **Reporting lines.** Where a line is confirmed by both methods, 69.6% of AI-first security postings report into security, against 65.3% elsewhere. Only 5% of labelled security postings have a confirmed line.
- **Security leaders reporting into engineering or the CEO.** 43.8% at AI-first companies against 22.4%, from 16 AI-first postings. The range is −8.8 to +51.1 points.
- **Age of open security openings.** Against open engineering openings at the same company, the ratio of ages is 0.9 at AI-first companies and 0.9 elsewhere.
- **Reposted openings.** 7.6% at AI-first companies against 11.8%, a likely but borderline gap. With each company counted once, the difference is 1.0 points.

10 Reference figures across the data

- Across all 1,593 complete job boards on 1 October 2026, there were 4.7 engineering openings per security opening and 10 per security engineering opening.
- Where a posting states a reporting line, 65.8% report into security and 15.1% into engineering.
- Across 247 companies, the median age of open security openings against engineering openings is 0.94, and security openings are older at 40.9% of them.
- Experience asks: 52% for AI security postings of any title, against 34% for security engineering postings outside AI security.

11 Where the data comes from

AI-first companies are 204 companies matched to job boards in the data. The list is not published. Everyone else is every other company in the data. Job aggregators, which repost other employers' openings, are left out of every count.

Security postings are every security posting collected from 31 July to 1 October 2026: 15,163 in all, 1,412 at 149 AI-first companies and 13,751 at 2,388 others. Security engineering openings are a narrower group within them. The engineering sample is 2,886 postings.

The census is a full count of every open posting on the 1,593 job boards where the read was complete on 1 October 2026: 196 AI-first and 1,397 others. It counts engineering, AI engineering and security openings, on one day.

A local model labels each security posting with a domain. 12,895 postings carry a label, and 453 were read as not security work. Domain shares use 1,151 AI-first and 10,869 other postings. Reporting lines are counted only when both the local model and a keyword rule confirm them.

Ranges of likely values are 95% unless marked 99%. The AI security share and the five compliance terms are computed at 99% only.

12 How this answer was stress-tested

Table 3: How this answer was stress-tested.

Test	What it found
Three AI-first boards with the most engineering openings removed. AI engineering per security opening.	1.29 against 0.65, range +0.39 to +0.93. Same direction and about the same size.
Pooled ratio. Engineering per security engineering opening across all openings.	11.3 against 9.7, range −1.01 to +4.80. No clear gap.
Largest AI-first employer removed. Agent-work count.	All security postings +4.1 to +14.2 points, and security engineering openings +1.4 to +16.5. Both above zero.
Largest AI-first employer removed. SOC 2.	+4.4 to +19.6 points at 99%. Above zero.
Largest AI-first employer removed. EU AI Act.	+0.8 to +6.0 points at 99%. Borderline, as in the main result.
Largest AI-first employer removed. AI security share.	−0.1 to +9.2 points at 99%. Reaches zero.
Largest contributing company removed. Security duty in AI engineering postings.	−15.6 to −3.6 points. Below zero.
Stricter 99% ranges. Five compliance terms.	SOC 2 clear of zero. EU AI Act and ISO 42001 borderline. ISO 27001 and the NIST AI Risk Management Framework show no clear gap.
Size control. Whether companies post security openings at all.	The gap changes with board size: −29, −14 and 0 points across small, medium and large boards.
Repeat runs. Security duty comparison.	Clear of zero on 25, 28 and 30 September and 1 October. Failed on 23 September.
Domain labels, local model against a keyword rule.	Kappa 0.52, below the 0.6 bar.
Domain labels, local model against a second AI model's blind read of 100 postings.	Kappa 0.75. The keyword rule against the second model is 0.51, so domain shares come from the local model.
Reporting lines, local model against a second AI model on 65 postings.	Of 20 cases where the local model alone saw a line, 11 were not real. Only lines both the local model and a keyword rule find are counted.

13 What would strengthen this answer

Table 4: What would strengthen this answer.

Test	What it would settle	Why it has not run
A census series	Whether any gap is growing	
A comparison group drawn at random	Whether AI-first companies post security openings at all at a different rate	
Separating securing AI from using AI in the agent-work measure	How much of 26% against 17% is each	
More AI engineering postings at AI-first companies	A tighter range on 1 in 29 against 1 in 7, from 115 postings today	
Reporting lines from more postings	Who security reports to at AI-first companies. 5% of labelled security postings have a line both methods confirm	

Acknowledgements

Question asked by William Lin.

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-2026-07,
  title       = {How Does Security Hiring at AI-First Companies Differ From Everyone Else? More AI
Engineering Openings per Security Opening},
  author      = {{AKA Security Research Factory}},
  institution = {AKA Security},
  type        = {Research report},
  number      = {AKA-RZ:2610.07v1},
  version     = {1.1.1},
  year        = {2026},
  month       = oct,
  url         = {https://akasecurity.io/research/security-hiring-at-ai-first-companies}
}
```