

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

AKA Security Research Factory^{*}

AKA Security

akasecurity.io/research/security-hiring-at-ai-first-companies

1 October 2026

AKA-RZ:2610.07v1 · Version 1.0.0

Abstract

Postings at AI-first companies against everyone else: more engineering openings per security engineering opening, more AI work, SOC 2 named twice as 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

At companies with both kinds of opening, AI-first companies post more engineering openings per security engineering opening, and their security postings more often name agent, LLM, MCP or prompt-injection work. Their postings name SOC 2 twice as often as everyone else's, and the EU AI Act four times as often, a likely but borderline gap. Their AI engineering postings less often name a security duty than AI engineering postings elsewhere.

Each figure in this piece comes with a range of likely values, and where that range includes zero we can't tell the gap from no gap. Other measures show no measurable gap. Engineering postings mention coding agents at about the same rate, close enough to count as a tie under a margin set in advance. Reporting lines (who a role reports to), and how long open security openings have been open compared with open engineering openings, show no measurable difference.

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 engineering openings per security engineering opening: a company median of 7, against 5 elsewhere	Medium. Holds as a company median and with each company weighted once. The pooled ratio does not hold.	Medium. 854 companies with both kinds of opening on a complete job board, one day of census data.
More AI engineering openings per security opening: 1.27 against 0.65	High. Holds with the three largest AI-first engineering boards removed.	Medium.
Security postings more often name agent, LLM, MCP or prompt-injection work: 25% of security engineering openings against 17%	Medium. The measure was added after the pre-registered one failed its precision check, and it mixes securing AI with using AI for security work.	Medium.
AI security is a larger share of security postings: 11% against 6% of those with a domain label	Medium. A likely but borderline gap, and with OpenAI, the largest single employer, removed it is too small to tell apart from no gap, but only just.	Medium.
SOC 2 is named twice as often, the EU AI Act four times as often	High for SOC 2, which stays clear on a stricter (99%) range and with the largest employer removed. Medium for the EU AI Act, a likely but borderline gap.	Medium.
AI engineering postings less often name a security duty: 1 in 29, against 1 in 7	High. 115 postings at 64 companies. It passed every check on the 25, 28 and 30 September and 1 October runs.	Low. 115 AI engineering postings at AI-first companies.
Coding-agent mentions in engineering postings show no difference	Medium for a tie. The likely range, -6.6 to $+0.8$ points, sits inside the eight-point tie band but reaches past five points.	Medium.

3 The numbers

The fourth column is the range of likely values for the gap between the two groups. "(99%)" marks a stricter range. "Points" are percentage points.

Table 2: The numbers.

Measure	AI-first	Everyone else	Likely range of the gap (95%)	Base, AI-first against everyone else
Engineering openings per security engineering opening, company median	7.0	5.0	 +0.7 to +3.7	113 and 741 companies
AI engineering postings per security posting	1.27	0.65	 +0.40 to +0.88	196 and 1,397 job boards
Security postings with agent, LLM, MCP or prompt-injection work in the responsibilities or requirements	19.9%	11.3%	 +4.2 to +13.5 points	1,412 and 13,751 postings
The same, security engineering openings only	25.2%	16.6%	 +1.7 to +15.8 points	628 and 5,829 postings
Security postings with a domain label, share in AI security	11.0%	5.9%	+0.3 to +9.6 (99%)	1,151 and 10,869 postings
Security postings naming SOC 2	18.8%	9.1%	+3.1 to +18.8 (99%)	1,412 and 13,751 postings
Security postings naming the EU AI Act	3.9%	1.0%	+0.8 to +5.7 (99%)	1,412 and 13,751 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

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.

Source for every figure in this piece: 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.

4 More engineering openings per security opening

At the typical AI-first company with both kinds of opening, there are 7.0 engineering openings for every security engineering opening. At the typical company elsewhere, 5. Taking a simple average in which every company counts equally gives 11.6 against 9.0.

AI engineering openings show the same direction. AI-first companies post 1.27 AI engineering openings for every security opening, against 0.65 elsewhere. Leaving out the three AI-first job boards with the most engineering openings still gives 1.29 against 0.65.

Adding up every opening across companies, instead of going company by company, there are 5.99 engineering openings per security opening at AI-first companies against 4.41 elsewhere. The range of likely values for that gap runs from +0.14 to +3.40, so the gap is likely but borderline. Per security engineering opening, added up the same way, it is 11.3 against 9.7, a gap too small to tell apart from no gap.

5 Security postings name more AI work

Counting only the responsibilities and requirements sections, 20% of AI-first security postings ask for work with AI agents, large language models (LLMs), MCP (the Model Context Protocol, which connects AI agents to tools) or prompt injection, against 11% elsewhere. Among security engineering openings it is 25% against 17%.

That count mixes two kinds of work. Mistral AI's CyberSecurity, Offensive Security Engineer will "proactively hunt for vulnerabilities in the interactions between our agentic applications, cloud in-

frastructure, and foundational models, with a focus on realistic, high-impact attack vectors." Vercel's Product Security Engineer will "build and operate LLM/agent-based reasoning that can validate business logic, auth, and design-level findings." The first secures AI. The second uses it to do security work.

AI security as its own domain (a model sorts each security posting into one specialty area) is 11% of AI-first security postings with a domain label, against 6% elsewhere. On a stricter (99%) range of likely values the gap is likely but borderline.

6 Compliance vocabulary leans toward SOC 2

SOC 2 appears in 18.8% of AI-first security postings and 9.1% elsewhere. The EU AI Act appears in 3.9% against 1.0%. The SOC 2 gap is clear, even on a stricter (99%) range of likely values and with OpenAI, the largest AI-first contributor, left out. The EU AI Act gap is likely but borderline on that stricter range, and still borderline with OpenAI left out. For ISO 42001 the gap is too small to tell apart from no gap, but only just. For ISO 27001 (17.7% against 13.9%) and the NIST AI Risk Management Framework, the gaps are too small to tell apart from no gap.

7 AI engineering postings less often name a security duty

At AI-first companies, 3.5% of AI engineering postings name a security duty such as threat modelling, secure design or access control. At other companies, 13.7% of AI engineering postings do. The sample is small: 115 AI-first postings at 64 companies.

A broader comparison only just holds. Across all companies, AI engineering postings name a security duty in 11.6% of cases against 18.4% for other engineering postings, and on the latest data that gap is likely but borderline.

Coding-agent mentions show no difference. 7.7% of AI-first engineering postings name coding agents or AI coding tools, against 10.9% elsewhere. The range of likely values for the gap, -6.6 to $+0.8$ percentage points, sits inside the eight-point margin that the questions fixed in advance as a tie.

8 What did not hold

- **Whether AI-first companies post security openings at all.** 64.8% of AI-first companies have a security posting open, against 71.7% of others. On boards with 44 or more open postings the gap falls to 0.2 percentage points (89.1% against 89.3%), and most comparison companies were added to the data because they posted a security role. The way companies were collected cannot support this comparison.
- **A security engineering opening.** 59.2% against 57.5%, no measurable difference.
- **Reporting lines.** Only 5% of postings state one. Among those, 69.6% of AI-first postings report into security against 65.3%, with a range of likely values for the gap from -9.8 to $+17.5$ percentage points.
- **Security leader postings that report into engineering or the CEO.** 43.8% against 22.4%, on 16 AI-first postings, a gap too small to tell apart from no gap.
- **The age of open security openings against open engineering openings.** Dividing how long security openings have been open by how long engineering openings have been open gives 0.9 at AI-first companies and 0.9 elsewhere, and the range of likely values for the gap runs from -0.44 to $+0.09$.
- **Reposted openings.** AI-first security openings were reposted less often (7.6% against 11.8%, a likely but borderline gap), but when every company counts equally the gap shrinks to 1.0 percentage points.

9 Reference figures across the data

- Across 1,593 complete job boards there are 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.
- Open security openings show no measurable age difference from open engineering openings at the same company. The median ratio of their ages is 0.94 across 247 companies, and security openings are older at 40.9% of them.
- AI security postings of any title that state a minimum experience ask for seven years or more in 52% of cases. Among security engineering postings outside AI security that state one, 34% do.

10 Where the data comes from

AI-first companies are a size-ranked list of engineering- and AI-driven companies, capped at 200, plus every company that has appeared on a Forbes AI 50 list. 204 of them matched to job boards in the data. The list is not published. **Everyone else** is every other company in the data.

Security postings are every security posting collected since 31 July: 15,163, including 1,412 at 149 AI-first companies. Postings by job aggregators, which are not employers, are left out of every count in this piece. Questions about security engineering openings use only that narrower group of postings. The census is a full count: it reads every open posting on 1,593 job boards where the read was complete, 196 of them AI-first, and counts engineering, AI engineering and security openings. The engineering sample is a fixed one-in-ten draw of non-security engineering postings from those boards, 2,886 postings. Security domains come from a model's label on 12,895 postings. Domain shares leave out the 453 the model read as not security work.

The questions and pass rules were written down before any numbers existed. Results were appended to that record after the run, and the questions above them were not edited. Two measures were added after the first run and are marked as such: the agent-work count limited to the responsibilities and requirements sections, and the reporting lines confirmed by both the model and a keyword rule.

11 How this answer was stress-tested

Table 3: How this answer was stress-tested.

Test	What it found
Pre-registration. Twelve questions, their measures and their pass rules fixed before the first run.	One question reversed its own premise: open security openings showed no measurable age difference from open engineering openings at the same company.
Hand-read precision. 20 matched sentences read by hand for every keyword measure. An instrument under 80% is not quoted.	The pre-registered agent, LLM and MCP keyword count read 13 genuine matches in 20, mostly company boilerplate and one MCP certification. It is not quoted. The section-scoped version used here read 19 in 20, and is marked as added after the run. Security duties read 17 in 20, coding agents 19 in 20, each compliance term 20 in 20.
Domain labels. The model's domain label against a keyword rule, then against a blind read of 100 postings by a second model.	Model against rule: kappa 0.52, below the 0.6 bar. Model against the second model: 0.75, rule against the second model 0.51. Domain shares come from the model. The second-model check was added after the run, once the pre-registered kappa check failed.
Reporting lines. The model's reporting-line reads against the second model on 65 postings.	The model found lines that were not there in 11 of 20 cases where it alone saw one. Only lines both the model and a keyword rule find are counted.
Redrawing whole companies, weighting each company once and removing the largest contributor for every comparison.	With OpenAI removed, the AI security share gap is too small to tell apart from no gap, but only just. Removing Anthropic kept the AI engineering security-duty gap.
Size control for whether companies post security openings at all.	The gap ran -29, -14 and -0 points across small, medium and large boards, and the question was ruled unreadable.
Stricter (99%) ranges for the five compliance terms.	SOC 2 stays clear. The EU AI Act is likely but borderline. ISO 42001 is too small to tell apart from no gap, but only just. ISO 27001 and the NIST AI Risk Management Framework show no clear gap.
Fresh runs on 24, 25 and 28 September. The full battery rebuilt from each day's data.	One finding came down. On 23 September, AI engineering postings named a security duty in 9.9% of cases against 19.7% for other engineering (likely range -19 to -2). On 24 September it was 11.4% against 18.6%, and the likely range reached +0.1. On 25 September it was 11.4% against 18.2%, and the likely range reached +0.2. On 28 September it was 11.3% against 18.1%, and the likely range reached +0.1. It was replaced with the pre-registered AI-first comparison. On 23 September that comparison's likely range reached +0.1. On 24, 25 and 28 September it passed every check, and again on 30 September and 1 October.

12 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 of these gaps is growing	The census began on 23 September.
A comparison group drawn at random	Whether AI-first companies post security openings at all at a different rate	Most comparison companies were found through their security postings. Fixing that needs new collection.
Separating securing AI from using AI in the agent-work measure	How much of the 25% against 17% is each	Needs a model label on every matching posting, with agreement testing.
More AI engineering postings at AI-first companies	A tighter likely range on the 1 in 29 against 1 in 7	115 postings today. The engineering sample grows as boards are read.
Reporting lines from more postings	Who security reports to at AI-first companies	Only 5% of postings state a line.

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 Differ at AI-First Companies? More Engineering Openings per Security Engineering Opening},
  author     = {{AKA Security Research Factory}},
  institution = {AKA Security},
  type       = {Research report},
  number     = {AKA-RZ:2610.07v1},
  version    = {1.0.0},
  year       = {2026},
  month      = oct,
  url        = {https://akasecurity.io/research/security-hiring-at-ai-first-companies}
}
```