

When a Security Posting Expects AI, Are Managers Asked to Stay Hands- On? More Often: 78% Against 43%

AKA Security Research Factory^{*}

AKA Security

akasecurity.io/research/ai-and-hands-on-security-managers

1 October 2026

AKA-RZ:2610.06v1 · Version 1.0.0

Abstract

149 security people-manager postings read by three AI models: manager postings that expect AI ask for hands-on work in 78% of 32 cases, against 43% of 47 without.

^{*}Written entirely by agents in the Research Factory, within guardrails and controls set by our team. Panel figures as of 23 September 2026. Second-read figures as of 25 September 2026.

1 Introduction

Three AI models each read 149 security people-manager postings, and we use the answer at least two agree on. Where the posting expects the hire to work with AI, 78% of manager postings ask the leader to stay hands-on, against 43% where it does not. Executive postings with an AI expectation centre on deciding in 81% of cases, against 65% without. Each figure comes with a range of likely values. If that range includes zero, we can't tell the gap from no gap. With the three models together, the executive gap is too small to tell apart from none, though two of the models each find a clear gap on their own.

Whether manager postings with no AI expectation differ from security manager postings in a spring 2024 sample of US LinkedIn postings, read the same way, is not measurable with this sample. Manager postings that ask for hands-on work are concentrated where AI shows up. A second read two days later, by newer versions of two of the models on a separate set of US postings, reproduced the manager gap.

2 The answer at a glance

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

Finding	Confidence	Representativeness
Manager postings that expect AI ask for hands-on work more often: 78% against 43%	Medium. A clear gap, also clear for two of the three models alone, without AI and security-vendor companies, and on the second read. For the other model alone it is likely but borderline. The size moves by up to 14 percentage points between methods, and a check where no single model judges both sides shows no clear gap.	Low. 149 managers, mostly security operations and GRC, mostly at enterprises.
Executive postings that expect AI centre on deciding 81% of the time, against 65%, a gap that is not measurable	Low. No clear gap with the three models together (likely range -4 to +36). A clear gap for two of the three models alone, and on the second read for one of its two models but not the other.	Low. 69 executives.
Coordinating as the core of manager postings: 25% with an AI expectation, 40% without	Low. No clear gap, and the models disagree most on this label.	Low. Manager postings only.
AI-exposed manager postings in 2026 ask for hands-on work more often than 2024 manager postings: 78% against 52%	Medium. A clear difference with the three models together. On the second read it was clear for one of the two models and not the other.	Low. Two one-time snapshots, collected two different ways, US only for 2024.
49% of the 149 security people-manager postings expect hands-on work	High as a level for these postings.	Low. Manager postings only.

3 The numbers

Each figure is the answer at least two of the three models agree on. "Must stay hands-on" means the posting says the leader personally does technical work: player-coach, works real cases, writes detections, does the compliance work. "Core" is the activity the responsibilities weight most heavily.

Table 2: The numbers.

Level	AI expectation in the posting	Postings	Must stay hands-on	Core is deciding	Core is coordinating
Manager	No	47	43%	30%	40%
Manager	Yes	32	78%	19%	25%
Executive	No	37	41%	65%	11%
Executive	Yes	32	38%	81%	3%
All people managers		149	49%	47%	21%

Source: AKA Security analysis of 149 security people-manager postings in AKA's job-board data, read by three AI models on 23 September 2026. The level rows sum to 148. The remaining posting, a Lead IAM Engineer who manages people by all three models' reading, carries a senior individual-contributor title and counts in the total only.

4 Manager postings that expect AI ask for hands-on work

The manager gap is 36 percentage points, with a likely range of +13 to +56 (allowing for several postings from one company). Two of the three models each find it on their own, the second at 78% against 45% and the third at 90% against 44%. The first, at 71% against 47%, finds a likely but borderline gap (+1 to +45).

The postings describe a working lead. A SOC Manager posting listed by Artemis Search: "our AI-native platform investigates and resolves cases for customers today... This is a player-coach leadership role... while still working real cases alongside your team." Mattermost's GRC Manager: "You will do the hands-on compliance work while coordinating across internal stakeholders in engineering, infrastructure, and IT." Forward Financing's Manager, Security Operations, will "lead the advancement of our AI security strategy, evolving our detection tooling and response processes... while owning and tuning our detection/response funnel."

5 Executive postings that expect AI: A gap in decisions that is not measurable

The executive gap is 16 percentage points, with a likely range of -4 to +36, too small to tell apart from no difference. The first model finds 81% against 51%, the second 80% against 64%, the third 87% against 63%. Coordinating is rarely the core of an executive posting in either group (3% and 11%).

KLA's Director, AI Governance & Privacy Attorney, will "drive its global AI governance program" and "lead the development of policies, controls, and governance frameworks for responsible AI use across products, internal tools, and enterprise operations."

6 The coordination gap is not measurable

Managers whose postings expect AI are coordination-centred 25% of the time, against 40% without. The likely range runs from -36 to +6 percentage points, so the gap is too small to tell apart from none. The models also disagree most on this label: across all 149 postings, the second model assigns coordinating to 12%, the first to 26% and the third to 23%. Delinea's Senior Manager, Cloud Engineering Governance (Fe-dRAMP), with no AI expectation, "will own day-to-day team operations, backlog execution, and cross-functional coordination, enabling the Director to focus on strategy and stakeholder engagement."

7 Against spring 2024

The same three models and the same instructions read the 2024 US postings with manager or executive titles from a public LinkedIn dataset (5 to 19 April 2024). The first model ruled 100 of 245 not security roles, including financial audit, model risk and physical security. After the series' own check of which postings are security roles, and the removal of job aggregators (sites that repost other employers' jobs), 72 were security people managers. AI expectations were almost absent (3%).

Table 3: Against spring 2024.

At least two of three models agree	2024	2026, no AI	2026, AI	2024 → 2026 AI, likely range
Manager must stay hands-on	52% (44)	43% (47)	78% (32)	 +4 to +46
Manager core is deciding	11%	30%	19%	
Manager core is coordinating	41%	40%	25%	 -36 to +7
Executive core is deciding	68% (28)	65% (37)	81% (32)	 -9 to +36

Source: AKA Security analysis, answers at least two models agree on, postings in brackets. 2024: 72 in-scope people-manager postings from a public LinkedIn dataset of US postings (5 to 19 April 2024). 2026: the 23 September panel. Likely ranges come from re-drawing the sample 2,000 times, a company at a time. 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.

Whether managers without AI differ from 2024 managers is not measurable with this sample: hands-on 43% against 52% (-32 to +13), coordinating 40% against 41% (-20 to +20). The AI-exposed manager postings are the ones that differ from 2024. Whether managers' deciding share differs by AI in 2026 is not

measurable (30% without, 19% with, -28 to +6). It is 11% of the 44 managers in 2024 and 25% of the 79 in 2026. Each model alone gives 14% in 2024 and between 27% and 29% in 2026. With the three models together, the rise is likely but borderline (+1 to +26).

8 A second read, two days later

On 25 September, the second and third models labelled every US security posting in both years for a separate study, using newer settings than the panel. That set includes postings with manager and executive titles, at US-headquartered companies only. Where both models agree the hire manages people, the manager gap reproduces.

Table 4: A second read, two days later.

Second read, US postings	Postings	Second model	Third model
2026 managers must stay hands-on: no AI against AI	32 against 24	38% against 79% (+16 to +64)	34% against 75% (+15 to +64)
2026 executives, core is deciding: no AI against AI	20 against 22	80% against 91% (-9 to +33)	75% against 100% (+10 to +43)
2024 managers must stay hands-on	44	59%	48%

Source: AKA Security analysis of the second and third models' labels from 25 September 2026. Counts are postings where both models agree the hire manages people and agree on the AI expectation. Ranges in brackets are likely ranges, allowing for several postings from one company.

The comparison of 2024 managers with 2026 AI-exposed managers shows a clear rise for the third model (+3 to +49) and no clear change for the second (-2 to +40).

9 Where the data comes from

The 2026 postings come from AKA's daily read of company job boards. 585 postings carried a leadership title (manager, director, head, chief, VP or lead) after removing program, product, project, account and risk managers. The first model read 432 of them. It also read 18 with "Management" in the title, which a corrected title rule no longer counts as leadership, and an earlier calibration sample across all levels. Only 166 turned out to be security people managers with reports. Many "Lead" and "Director" titles have none, and about a quarter of the rest were not security roles.

The local model the series uses to read every posting in both years then checked which of the 166 are security roles, and ruled 14 out. The series' rules on which postings count then took out 3 more: postings by job aggregators, which are not employers, and postings that closed before 23 September. That leaves 149, 15 of them from the calibration sample.

The 149 postings come from 120 companies, and 90 are at enterprises. The functions are uneven: security operations 53, GRC 51, privacy 21, offensive 11, the rest in single digits.

AI expectation means the posting asks the hire to use AI tools, build AI automation, secure AI systems, govern AI or bring machine-learning skills. The panel was read on 23 September 2026 and is a fixed sample. It does not refresh with the daily data.

These postings cannot show what produces the difference. AI-exposed postings come from different teams and companies. One check leaves out AI-native companies, AI infrastructure companies and security vendors, but other differences between teams remain.

10 How this answer was stress-tested

Table 5: How this answer was stress-tested.

Test	What it found
Three models separately. Each model labelled all 149 postings.	Manager hands-on gap, first to third model: 71 against 47, 78 against 45, 90 against 44. Executive deciding gap, in the same order: 81 against 51, 80 against 64, 87 against 63. The manager gap is clear for the second and third models and likely but borderline for the first (+1 to +45). The executive gap is clear for the first and third and not clear for the second (-5 to +36).
Rater agreement. Cohen's kappa for each pair of models (agreement beyond chance: zero is chance, one is perfect).	Hands-on 0.74 to 0.84. AI expectation 0.86 to 0.92. Deciding 0.49 to 0.69. Coordinating 0.39 to 0.78, with the second model the outlier. The tables report the answer at least two models agree on.
Split-model check. The first model's AI label paired with the second model's work-mix label, so no single model judges both sides.	Manager hands-on 71 against 49. Executive deciding 78 against 65. Neither gap is clear: both likely ranges include zero (-1 to +43 and -7 to +34).
Company-type exclusion. AI-native companies, AI infrastructure companies and security vendors removed.	Manager hands-on 75 against 38 (+8 to +63). Executive deciding 83 against 60 (+1 to +44). The manager gap is clear outside the companies most likely to build with AI, and the executive gap is likely but borderline.
Redrawing whole companies: the sample redrawn 2,000 times, a company at a time, to find each likely range.	Manager hands-on gap +13 to +56. Executive deciding gap -4 to +36. Coordination gap -36 to +6.
Title rule audit. The corpus's own seniority tag compared with a stricter title rule.	Found and fixed an error: titles naming a domain called "management" counted as managers, including "Security Engineer, Vulnerability Management". The panel moved by at most a point.
Scope check of the 2024 set. The first model read all 245 leadership-titled 2024 postings.	100 were not security roles. The 2024 comparison uses the 72 in-scope people managers, after the series' own scope check and the removal of job aggregators as well. Hands-on agreement on 2024 was weaker (kappa 0.47 to 0.64) than on 2026.
Scope check of the 2026 panel. The series' local model, which reads every posting in both years, read the 166 panel postings.	It ruled 14 not security roles, and they leave the panel, as do 3 postings by job aggregators or closed before 23 September. The manager gap and the 2024 comparison held. The executive gap moved from clear to too small to tell apart from none.
Second read on 25 September. Newer settings of the second and third models on US-headquartered postings in both years.	The manager gap reproduced under both models. The executive gap was clear for the third model and not for the second. The 2024 comparison was clear for the third model and not for the second.

11 What would strengthen this answer

Table 6: What would strengthen this answer.

Test	What it would settle	Why it has not run
A larger manager sample	Narrower likely ranges, and room to split by function beyond operations and GRC	149 managers is what 54 days of collection produced, and 124 of them were first seen after collection widened to more titles on 21 September. The earliest re-read with three months of the wider collection is late December 2026.
The same teams over time	Whether a team's manager posting became more hands-on after the team added AI expectations	Needs repeated postings from the same team, which are rare in two months of data.
Matching on company and function	Separates AI from the kind of team that tends to post AI expectations	Needs several times the sample to match within company and function.
What managers actually do	Whether stated expectations match the job as worked	Postings show stated requirements. Time-use data comes from surveys or licensed sources.
A practitioner read of 50 postings	Whether "must stay hands-on" matches how a security leader reads the role	Needs practitioner time.

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-2026-06,
  title      = {When a Security Posting Expects AI, Are Managers Asked to Stay Hands-On? More Often: 78% Against 43%},
  author     = {{AKA Security Research Factory}},
  institution = {AKA Security},
  type       = {Research report},
  number     = {AKA-RZ:2610.06v1},
  version    = {1.0.0},
  year       = {2026},
  month      = {oct},
  url        = {https://akasecurity.io/research/ai-and-hands-on-security-managers}
}
```