When a Posting Expects AI Use, Are Security Managers Asked to Be More Hands-On? Manager Postings That Expect AI Ask for Hands-On Work More Often

AKA Security Research Factory*
AKA Security
akasecurity.io/research/ai-and-hands-on-security-managers

9 October 2026
AKA-RZ:2610.06v1 · Version 1.1.0

Abstract

149 security people-manager postings, fall 2026: manager-level postings that expect AI ask the hire to stay hands-on in 78% of cases, against 43% without. A shift in executive decision work is not measurable.

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

1 Introduction

More often, at the manager level. In fall 2026, 78% of 32 security manager postings that expected AI use asked the hire to stay hands-on. Among 47 manager postings with no AI expectation, 43% did.

At the executive level, a difference is too close to call. Executive postings that expect AI make deciding the core of the job in 81% of 32 cases, against 65% of 37 without, but the range of likely values includes zero.

These figures come from a fixed panel of 149 security people-manager postings at 120 companies. The postings cannot show what produces a difference. AI-exposed postings come from different teams and companies.

These findings come from one panel of 149 security people-manager postings at 120 companies, mostly in security operations and GRC, and 90 of them at enterprises.

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-level postings that call out AI ask the hire to be hands-on more often than those that don't: 78% of the 32 that call out AI, against 43% of the 47 that don't.	**Medium**. The gap is too large to be chance but the size of the gap differs based on how the postings are read.	**Low**. With only 79 manager postings, there are too few to speak for all security managers. The panel as a whole is mostly security operations and GRC.
About half of security people-manager postings ask the hire to stay hands-on: 49% of 149.	**High**. Each of the three AI models that read the postings puts the share close to half.	**Low**. Covers only postings for managers and executives, from one fixed set of 149.
2026 manager postings that expect AI ask for hands-on work more often than 2024 manager postings: 78% against 52%.	**Medium**. There is a clear rise in the main reading but the second model doesn't always agree.	**Low**. Two one-time snapshots, collected in different ways.
Executive postings may center on making decisions more often when they expect AI: 81% of those that do, against 65% of those that don't. The gap is too close to call.	**Low**. Too close to be sure. Needs a larger executive sample.	**Low**. Only 69 executive postings.
A gap in coordination work between manager postings with and without AI, 25% against 40%, is also too close to call.	**Low**. Too close to be sure, and readings of the main activity often disagree. Needs a larger manager sample.	**Low**. Manager postings only.

3 The numbers

Three AI models each read all 149 postings. The answer used is the one at least two of the three agree on.

Table 2: The numbers.

Level	AI expectation in the posting	Postings	Stays 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 at 120 companies in AKA's job-board data (open on or after 23 September 2026), read by three AI models.

An 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. Stays hands-on means the posting says the hire mainly does the technical work personally, such as a player-coach who works real cases, builds detections or does the compliance work. The core is the activity the responsibilities weight most heavily.

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

In security manager postings open in fall 2026, 78% of the 32 that expected AI asked the hire to stay hands-on. Of the 47 with no AI expectation, 43% did. The gap is +36 points, with a range of likely values of +13 to +56 points. The 32 postings come from 28 companies and the 47 from 35.

The gap holds for each AI model alone, at +24, +33 and +46 points. It also holds with AI-native companies, AI infrastructure companies and security vendors removed, at +37 points.

The finding is rated Medium. In a split-model check, one model judged the AI expectation and another judged the work mix, so no single model judged both sides. That gap is +22 points, with a range of −1 to +43 that includes zero. The three models also differ on size by up to 14 points.

Across all 149 people-manager postings, 49% ask the hire to stay hands-on. The range is 40% to 58%, and each model alone puts the share between 49% and 51%.

5 Against spring 2024

The 2024 comparison uses 72 people-manager postings from a public US job-postings dataset, listed 5 to 19 April 2024. Of those, 44 are at manager level, from 40 companies. AI expectations appear in 3% of the 72.

In 2024, 52% of US security manager postings asked the hire to stay hands-on. In 2026, 78% of AI-exposed manager postings did. The difference is +26 points, with a range of +4 to +46.

Table 3: Against spring 2024.

Comparison	2024	2026	Range of likely values for the change, points
Stays hands-on: all 2024 managers, 2026 managers with an AI expectation	52%	78%	+4 to +46
Stays hands-on: all 2024 managers, 2026 managers with no AI expectation	52%	43%	−32 to +13
Coordinating: all 2024 managers, 2026 managers with no AI expectation	41%	40%	−20 to +20
Coordinating: all 2024 managers, 2026 managers with an AI expectation	41%	25%	−36 to +7
Executive, AI expectation, deciding	68%	81%	−9 to +36
Manager, deciding, all postings	11%	25%	+1 to +26

Source: AKA Security analysis of 72 US people-manager postings from a public job-postings dataset (April 2024) and 149 people-manager postings in AKA's job-board data (fall 2026).

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.

Manager postings without an AI expectation show no clear change from 2024. Deciding as the core of manager postings rose from 11% of 44 in 2024 to 25% of 79 in 2026, a likely but borderline rise.

6 Executive postings that expect AI: A gap in decisions is too close to call

In security executive postings open in fall 2026, 81% of the 32 that expected AI made deciding the core of the job. Of the 37 with no AI expectation, 65% did. The gap, +16 points, has a range of −4 to +36, so it is too close to call.

The models split on this one. Two of the three models alone find a clear gap, and one does not. Agreement on the deciding label runs from 0.49 to 0.69 across the model pairs. Coordinating was the core of 3% of executive postings with AI and 11% without.

7 The coordination gap is too close to call

Among security manager postings, coordinating was the core in 25% of the 32 that expected AI and 40% of the 47 without. The gap, −15 points, has a range of −36 to +6.

Agreement on the coordinating label is low. One hosted model is the outlier, and no model pair that includes it agrees above a kappa of 0.39. Deciding as the core of manager postings went from 30% without AI to 19% with AI, which is also too close to call.

8 A second read, two days later

On 25 September 2026, two hosted models read the US-headquartered postings again with newer settings. Only postings where both agree the hire manages people, and agree on the AI expectation, count.

Table 4: A second read, two days later.

Second read, US postings	Postings, no AI and AI	Hosted model A	Hosted model B
Manager hands-on, AI minus no AI	32 and 24	+42 (+16 to +64)	+41 (+15 to +64)
Executive deciding, AI minus no AI	20 and 22	+11 (−9 to +33)	+25 (+10 to +43)
2026 AI-exposed managers against 2024 managers, hands-on	44 (2024) and 24	+20 (−2 to +40)	+27 (+3 to +49)

Source: AKA Security analysis of US-headquartered people-manager postings in AKA's job-board data (read 25 September 2026) and 44 US manager postings from a public job-postings dataset (April 2024). Figures are gaps in points, with the range of likely values.

The manager hands-on gap holds in both second reads. The executive decision gap and the 2024 comparison hold in one and not the other.

9 Where the data comes from

The 2026 panel is 149 security people-manager postings at 120 companies, 90 of them at enterprises. The postings were open on or after 23 September 2026, and the panel does not refresh with daily data. By level, 79 are manager postings and 69 executive postings, with 1 at another level. By function, 53 are in security operations, 51 in GRC, 21 in privacy and 11 in offensive security.

The panel was built over 54 days. 124 of the 149 postings were first seen after the collection widened to more titles on the evening of 21 September 2026.

It started from 585 postings with a leadership title: manager, director, head, chief, VP or lead, after removing program, product, project, account and risk managers. Of those, 166 were security people managers with reports. A local model ruled 14 of them not security roles, and 3 more were removed as job-aggregator postings or postings closed before 23 September.

The 2024 set starts from 245 leadership-titled US postings. A scope check ruled 100 of them not security roles, such as financial audit, model risk and physical security, leaving 72 people managers after aggregators were removed.

Three hosted AI models each read all 149 postings. Agreement across the model pairs is 0.74 to 0.84 on hands-on work, 0.86 to 0.92 on AI expectation, 0.49 to 0.69 on deciding and 0.39 to 0.78 on coordinating. Each range of likely values (95%) is found by redrawing the sample 2,000 times, a company at a time.

10 How this answer was stress-tested

Main results, AI minus no AI, in points: manager hands-on +36, executive deciding +16, manager coordinating −15.

Table 5: How this answer was stress-tested.

Test	What it found
Each model alone, manager hands-on.	+24 (+1 to +45), +33 (+12 to +53) and +46 (+26 to +64). Same direction, each clear of zero.
Split-model check, manager hands-on. The AI label from one model, the work-mix label from another.	+22 (−1 to +43). No clear change, smaller by 14 points.
Without AI-native companies, AI infrastructure companies and security vendors, manager hands-on.	+37 (+8 to +63). Same direction, about the same size.
Each model alone, executive deciding.	+30 (+10 to +50) and +24 (+5 to +42) are clear. +16 (−5 to +36) is not.
Split-model check, executive deciding.	+13 (−7 to +34). No clear change.
Without AI-native companies, AI infrastructure companies and security vendors, executive deciding.	+23 (+1 to +44). Likely but borderline.
Second read, 25 September.	Manager hands-on +42 and +41, both clear. Executive deciding +11, not clear, and +25, clear.
Second read against 2024.	+20 (−2 to +40), not clear, and +27 (+3 to +49), clear.
Stricter title rule. Titles naming a domain called "management" not counted as managers.	The panel moved by 1 point at most.
Scope check of the 2024 set.	100 of 245 postings were not security roles. Hands-on agreement is lower in 2024: 0.47 to 0.64, against 0.74 to 0.84 in 2026.
Scope check of the 2026 panel. A local model read the 166 postings.	14 ruled out and 3 more removed, leaving 149.

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 ranges, and a split by function beyond operations and GRC	
The same teams over time	Whether a team's manager posting became more hands-on after the team added AI expectations	
Matching on company and function	Separates AI from the kind of team that tends to post AI expectations	
What managers actually do	Whether stated expectations match the job as worked	
A practitioner read of 50 postings	Whether "stays hands-on" matches how a security leader reads the role	

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 Posting Expects AI Use, Are Security Managers Asked to Be More Hands-On? Manager
Postings That Expect AI Ask for Hands-On Work More Often},
  author      = {{AKA Security Research Factory}},
  institution = {AKA Security},
  type        = {Research report},
  number      = {AKA-RZ:2610.06v1},
  version     = {1.1.0},
  year        = {2026},
  month       = oct,
  url         = {https://akasecurity.io/research/ai-and-hands-on-security-managers}
}
```