tines

Security

Hype check



© TINES · TINES.COM/SECURITY-HYPE-1

About hype check

Welcome to Hype Check, an original research series from Tines. We ask IT and security leaders — people whose work we admire and whose perspectives we value — what their teams are prioritizing, testing, and struggling with. Then we publish what we hear quickly, while the questions are still being debated, to see where reality clashes with hype.

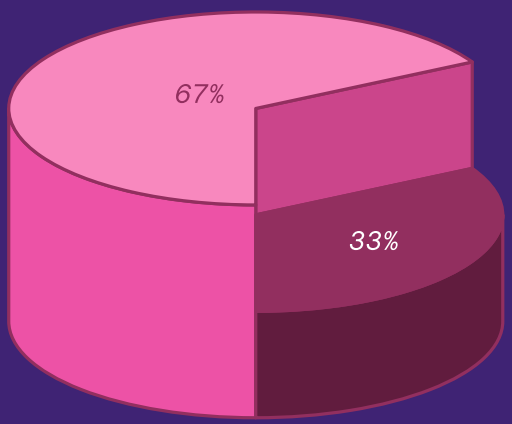
We believe research should start a conversation, not end one. Alongside the survey findings, you'll find excerpts from follow-up interviews that put context behind the numbers. Have a different take? Let us know in the **Tines Community forums.**



About the research

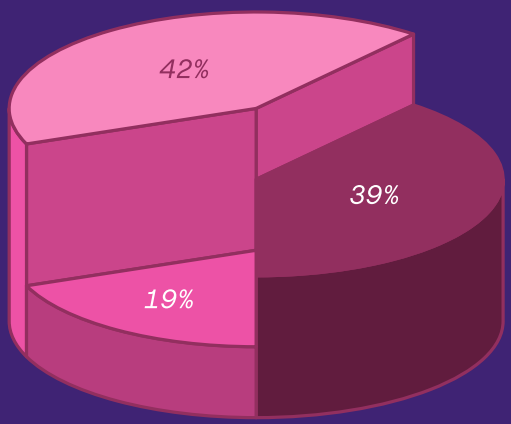
This edition of Hype Check draws on 70 anonymous survey responses from security leaders (67%) and practitioners (33%). Technology and software was the largest industry group (39%), followed by financial services (19%). Respondents' organizations were headquartered in North America, Europe and APAC. Results are self-reported and percentages are rounded.

ROLE



- Security leaders
- Practitioners

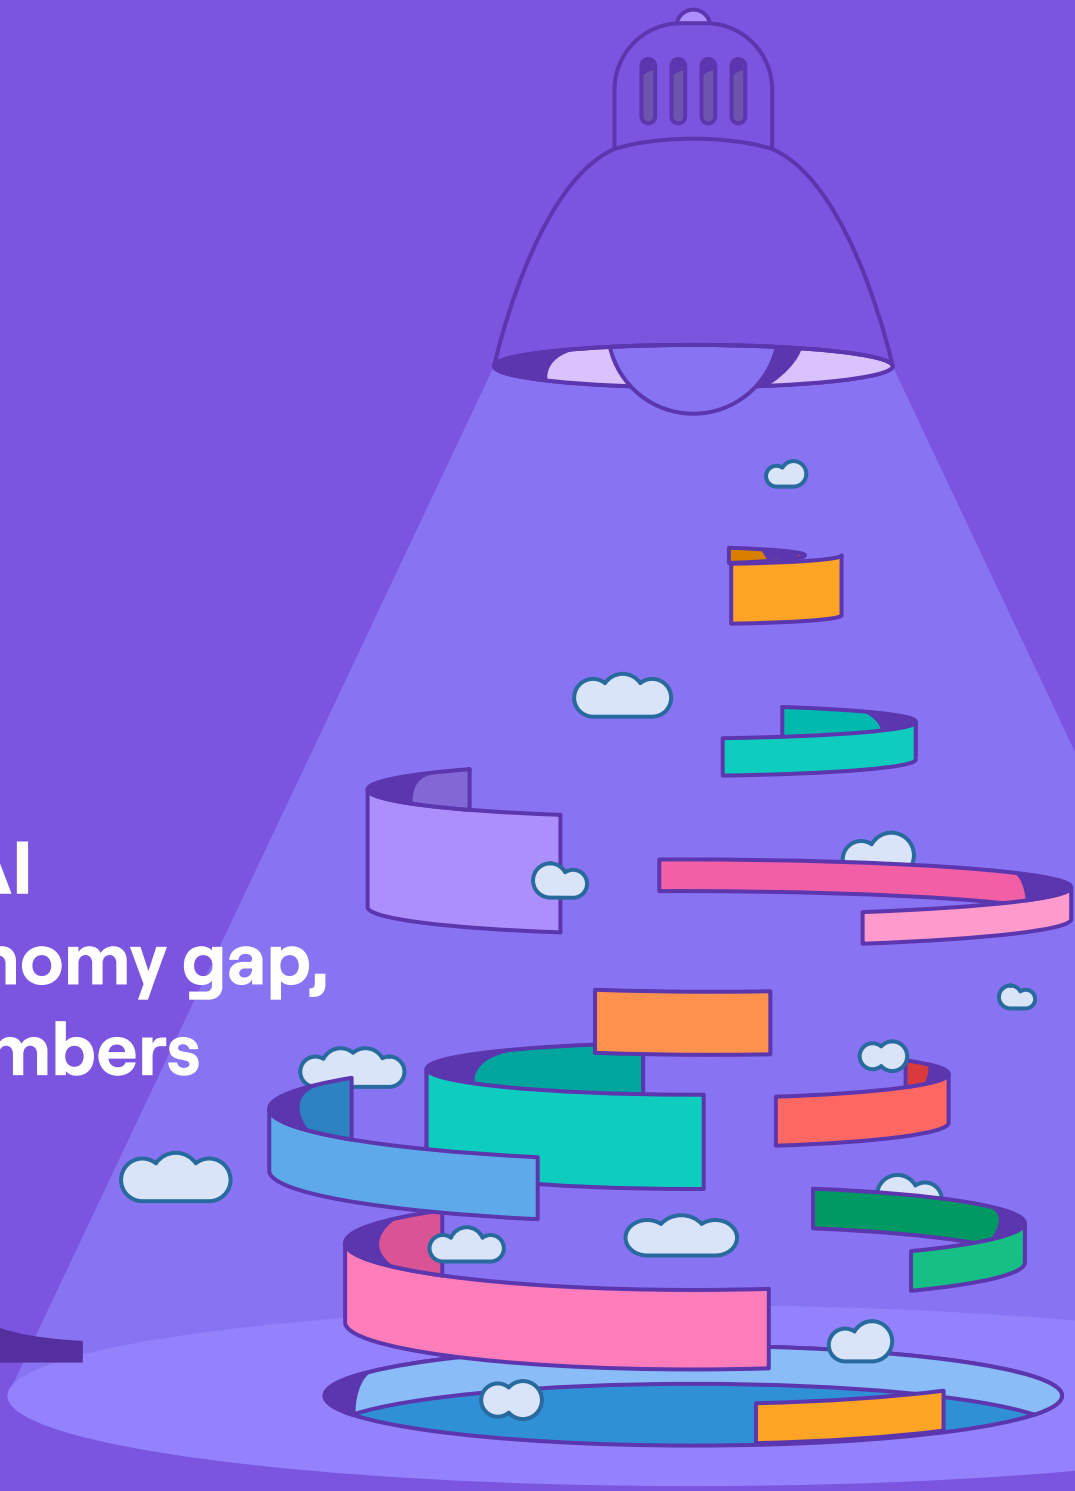
INDUSTRY



- Technology and software
- Financial services
- Other

1

The AI autonomy gap, in numbers





As AI helps attackers move at machine speed, the appetite for "fighting AI with AI" is strong.

For security operations teams, that means looking beyond using AI to generate faster alert summaries. They want AI to help them decide what happens next — and, in some cases, take action without waiting for a human to work through every step. But deciding which actions to permit, and when an analyst should step in, is another matter.

Just 9% of those surveyed describe their current AI use as heavily integrated, with autonomy to act on some decisions. Why is ambition so far ahead of practice?

91% plan to increase AI autonomy

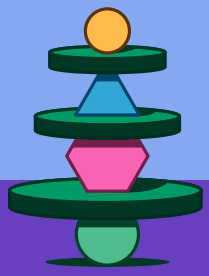


9% say AI can act on decisions today



> In our organization, we're building response automation from the ground up. For us, giving AI responsibility for deciding which response actions to take — not just carrying out predefined steps — is still a few years down the track.

Senior Director of
Cybersecurity Operations





AI can spot the threat. Should it also contain it?



The first clue is where teams use AI most — and where they're more cautious.

64% report using AI for threat detection or anomaly identification, while just 26% use it for containment actions, such as isolating a host or disabling an account. That makes AI use for detection 2.5 times more common than containment.

The promise of the AI SOC goes beyond spotting threats faster: it's helping teams decide what to do next. But before AI acts, leaders need answers. What evidence would trigger it? How far could the consequences spread? Could we reverse the action — and explain why it happened? For some teams, that's a level of uncertainty they're just not willing to accept.

64%



USE AI FOR DETECTION

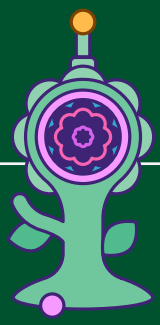
26%

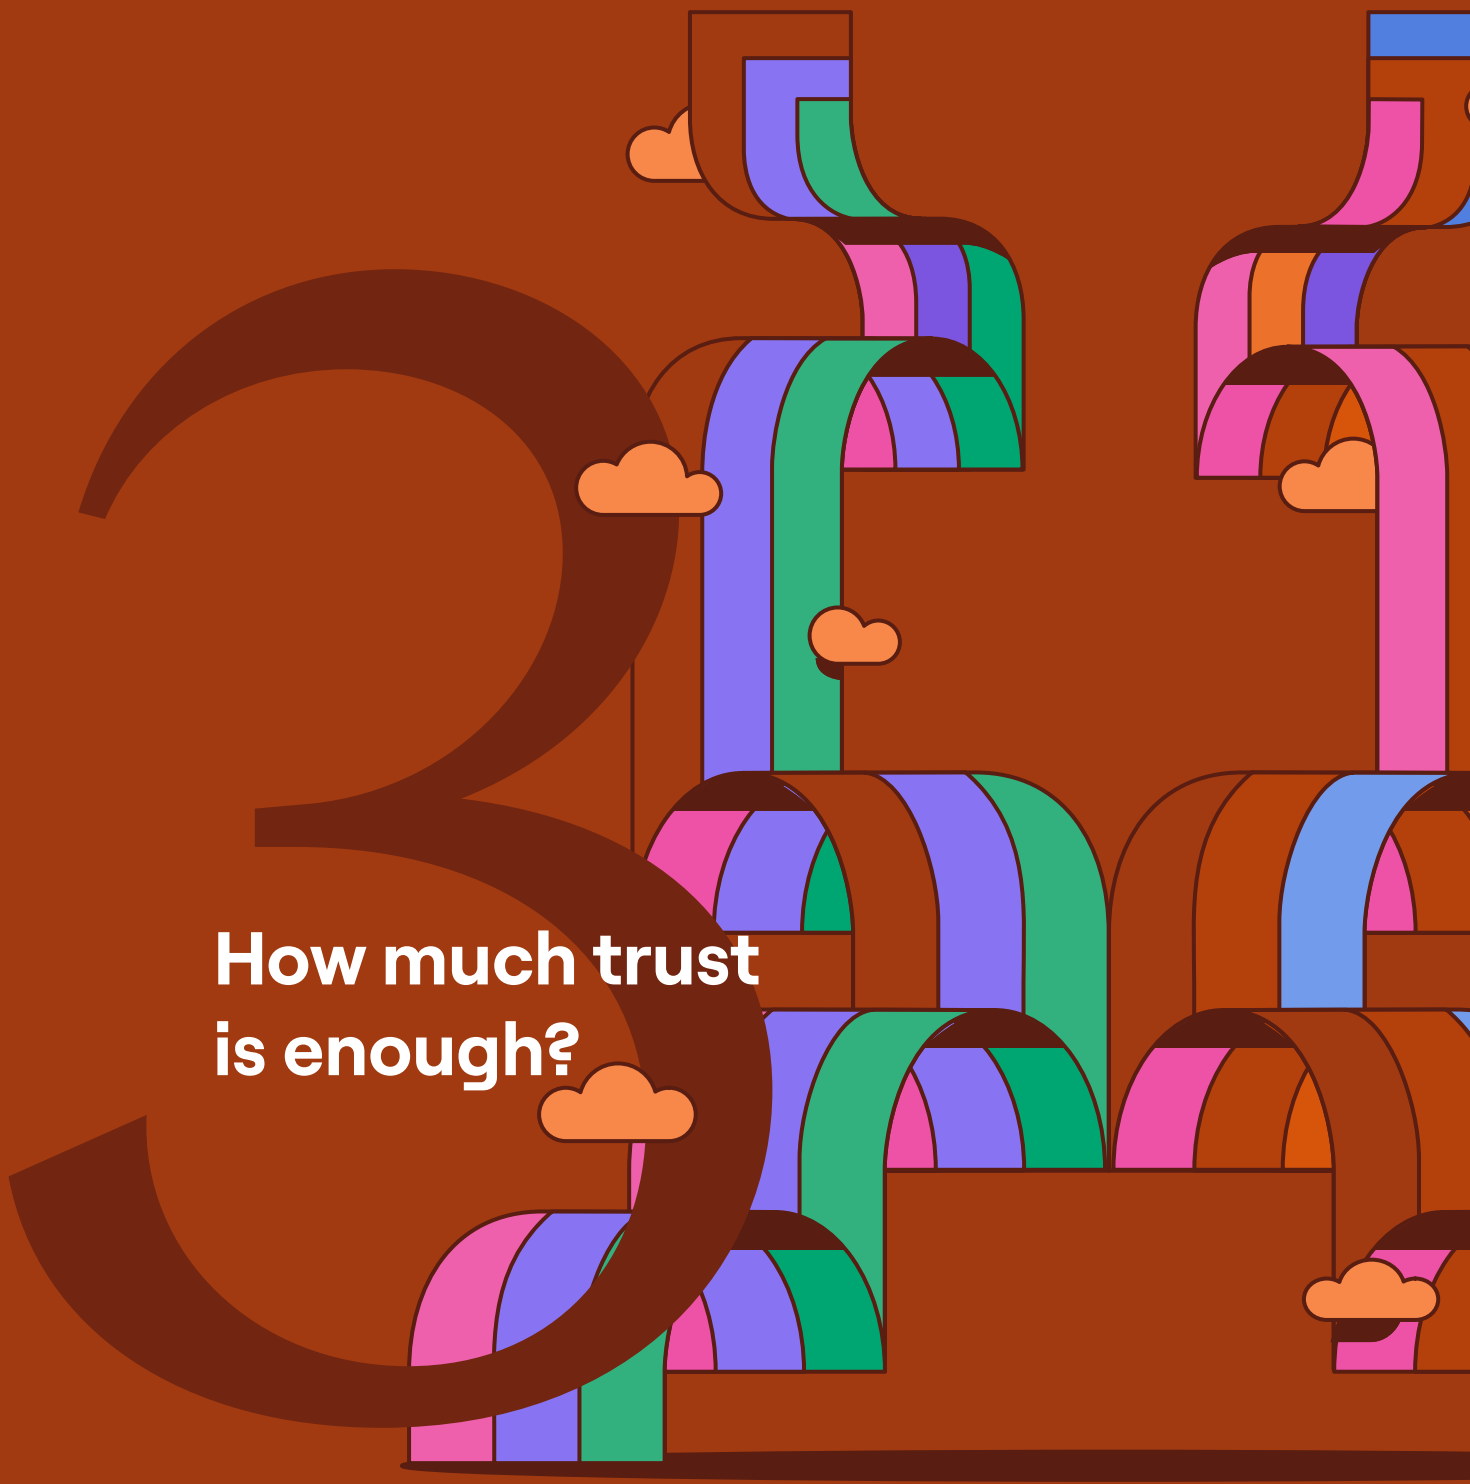


USE AI FOR CONTAINMENT

"Do you have a good data security posture? Do you have control over your APIs in the environment? Do you have an orchestrated process to onboard and offboard APIs? Do you have role-based access control? This is some of the basic security hygiene that needs to be checked off before you even jump into an AI-enabled environment."

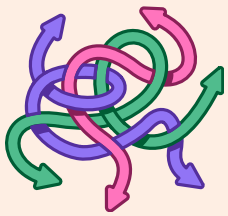
Senior Cybersecurity Operations Specialist





How much trust is enough?

Respondents largely agree on the direction of travel — greater AI autonomy. They differ on what needs to happen before they take the next step.



61% plan to increase AI autonomy even with some trust concerns remaining. 30% say they'll increase it only if those concerns are resolved first. Interestingly, responses were similar across all regions, despite differences in their regulatory environments.

That's not a split between leaders who see value in AI for security operations and people who don't. It's a difference in how they approach uncertainty and risk. Some are prepared to expand autonomy while managing remaining concerns. Others want greater confidence before they do so.

Should security teams move ahead with some trust concerns unresolved, or wait until they're comfortable giving AI more control? Only a leader with influence over their organization's risk appetite can make this call.



61% will increase autonomy despite trust concerns



30% will wait until concerns are resolved

> I think security teams are heading toward being too confident with autonomy. We need to have checks and balances in place, and putting those checks and balances in at the scale and speed that AI does things is challenging.

Sam Lalli,
ICAM Security Lead, Miro



4

Barriers to AI autonomy



Wanting to give AI more autonomy is one thing. Building the conditions to do it is another.

Asked to name their biggest barrier, respondents most often selected difficulty explaining an AI decision during an audit or review (23%). Budget or resourcing constraints (including human resources) and regulatory or compliance restrictions followed at 22% each.

These problems call for very different solutions. And together, they show that granting AI more autonomy takes more than a capable model.

Other responses included lack of trust in the underlying AI models (14%) and insufficient internal validation of AI performance (9%). Elsewhere in the survey, one respondent pointed to a "lack of trusted and verified data to feed agents." It's a reminder that the decision to let AI act depends not only on trusting the model, but also on trusting the information it uses.



23%



EXPLAINING DECISIONS IN AN AUDIT OR REVIEW

22%



REGULATORY OR COMPLIANCE RESTRICTIONS

22%



BUDGET OR HUMAN RESOURCE CONSTRAINTS

"Security has traditionally been seen as a blocker. I'm challenging my team to change that. I say to them: we're an enablement team, but equally we can't enable everything. We can't just open the floodgates. We need to be pragmatic about things."

Sam Lalli,
ICAM Security Lead, Miro





Faster defense needs a baseline

Giving AI more autonomy has its challenges, but teams are already seeing measurable gains from the ways they use it today. For example, 29% report that AI has cut their mean time to detect by at least half.

That reported gain makes another finding harder to ignore: 34% say their organization has no formal measurement of median threat detection time. 20% report no formal measurement of median time to resolve or contain an incident.

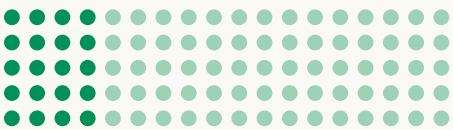
If a team gives AI more responsibility, how will it know whether that change improved its defense?





34%

HAVE NO FORMAL MEASUREMENT OF MEDIAN RESOLUTION OR CONTAINMENT TIME



20%

HAVE NO FORMAL MEASUREMENT OF MEDIAN THREAT DETECTION TIME

"

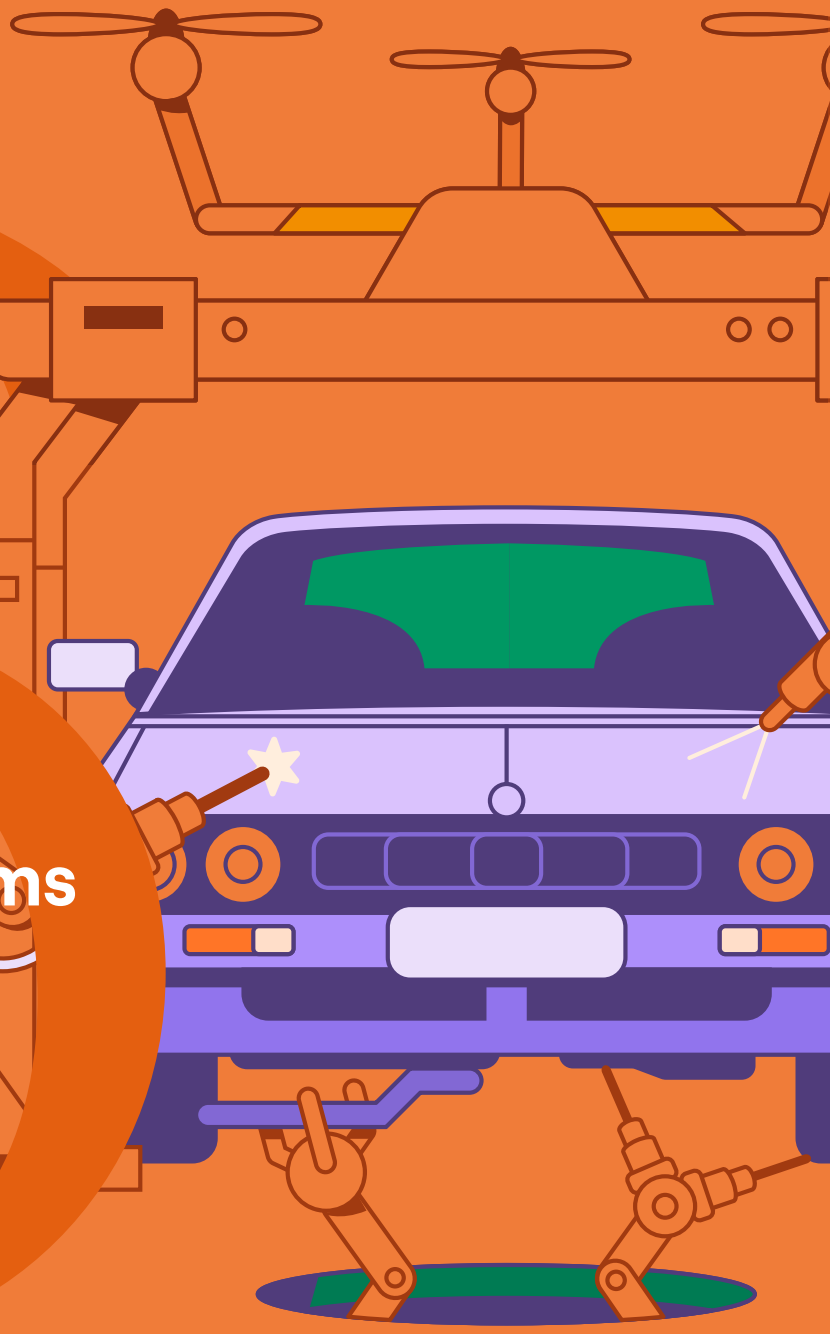
AI promises to save time, but it's hard to trust a time-saved figure without understanding how it's calculated and whether it reflects the work the workflow is actually doing.

Senior Cybersecurity Operations Specialist



6

How fast are security teams moving?





Super fast detection and containment times remain uncommon — but not unheard of.

7% report a current median detection time in the 1–5 minute band, while 6% report a median resolution or containment time under 15 minutes. By contrast, 16% report a median detection time of one hour or more, showing how widely detection times vary.

Meanwhile, some teams are setting ambitious goals for the end of 2027: 4% target detection in under a minute, while 13% target resolution or containment in under 15 minutes.

Interestingly, the sub-minute target isn't limited to teams already detecting threats within minutes. Some teams aiming for sub-minute detection currently report times of 15–30 minutes. Reaching that goal would cut their detection time by more than 93% — a substantial leap.

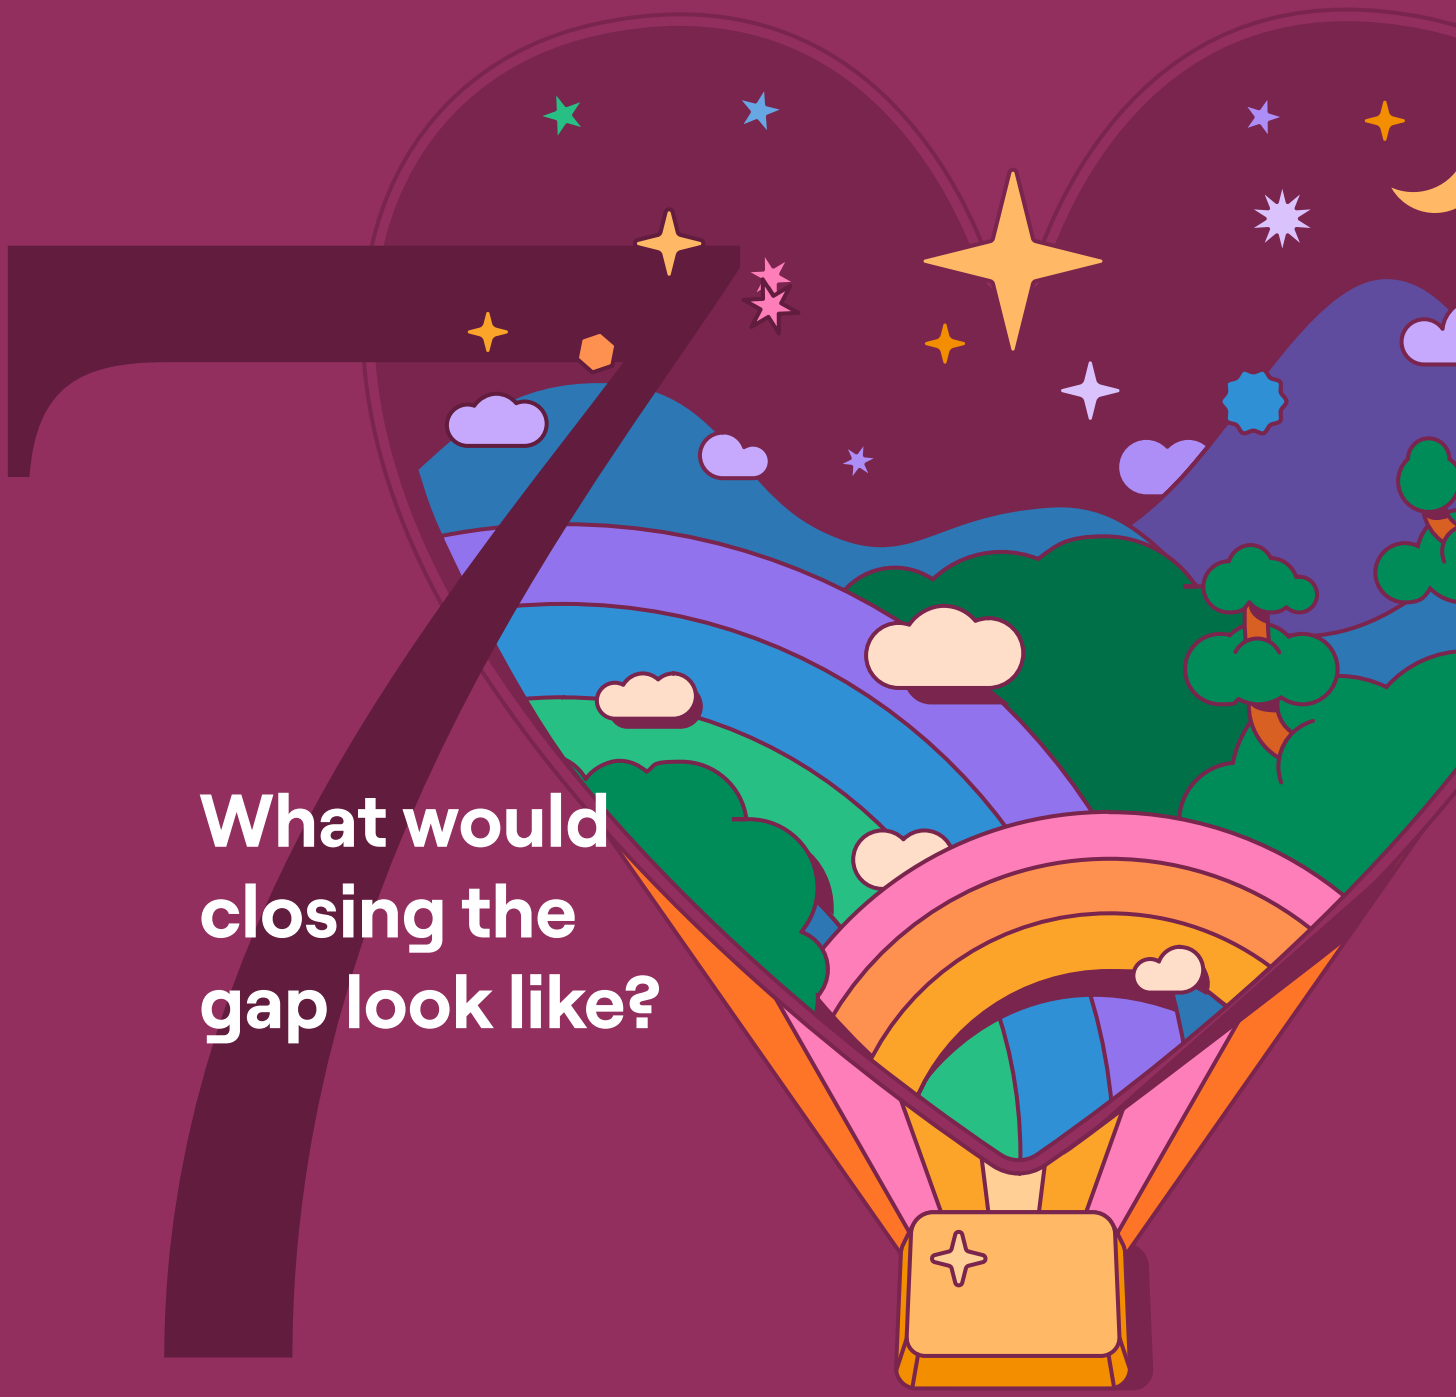
These are goals, not forecasts. Fast detection matters most when teams can also investigate and respond effectively, and measure the difference.

7%

REPORT A MEDIAN DETECTION TIME UNDER 5 MINUTES

6%

REPORT A MEDIAN RESOLUTION OR CONTAINMENT TIME UNDER 15 MINUTES



What would closing the gap look like?

This survey can't tell us exactly when teams will close the AI autonomy gap — **only that 91% plan to narrow it in 2027.**

The primary challenge for security teams in 2027 will not simply be getting AI to spot threats faster, but managing the overwhelming volume of alerts that rapid triage generates. As Tines' Field CISO Tim Bandos writes, this acceleration creates a **"decision gap"** where more information reaches analysts, but they are still left to determine what happens next.

To close this gap and turn rapid detection into reliable action, teams must move beyond out-of-the-box solutions to build custom automations tailored to their specific enterprise processes. Ultimately, success requires carefully defining the boundaries of AI by determining which decisions it can support, where it can act autonomously, and where human operators must maintain control.

That's the approach security teams are taking with **Tines 3B**, building AI-assisted apps, agents, and automations around their own processes in a secure environment governed by their enterprise. Rather than accepting a fixed definition of an "autonomous SOC," they can decide where AI helps, where it acts, and how its actions are overseen and explained.

The question for 2027 isn't whether AI belongs in security operations. It's whether teams can turn faster detection into action they trust.

