ANOMALI

Executive Brief

SECURITY EFFICIENCY THAT DOESN'T INCREASE RISK

An executive brief for CISO and CFO alignment

In 2025, the average breach took 241 days to identify and contain, the lowest figure in nine years of IBM research. Organizations that used AI and automation extensively across their security operations cut that lifecycle by about 80 days and spent roughly $1.9 million less per breach. The number that moves both risk and cost is the same one: how long it takes to reach a confident decision.

That single measure is why security productivity and risk reduction are not competing budget lines. They are the same initiative, if the program is built correctly.

THE PROBLEM IS NOT A SHORTAGE OF TOOLS

Most security programs are short on decisions, not tooling. New platforms get deployed, dashboards multiply, and activity rises, while analysts stay overloaded and leadership struggles to show that risk actually fell.

AI often speeds up individual tasks without reducing the total work that still reaches a human analyst, because investigations require the same validation and alerts still get double-checked. The result is motion without progress: faster activity, unchanged outcomes.

This means that the meaning of "productivity" needs to change for a security organization. Speed is only worth paying for when it produces earlier, better-founded decisions. When a program removes noise and resolves uncertainty before it reaches an analyst, two things improve at once.

1. Cost falls, because less low-value work runs through expensive people and infrastructure.
2. Risk falls, because real intrusions surface earlier, while there is still margin to act. Cost discipline and risk reduction stop pulling against each other.

GOVERNANCE GUARDRAILS: WHY SPEED HERE DOES NOT ADD RISK

Speed raises risk only when it removes the ability to explain a decision. The exposure most boards now worry about is real: IBM found that 97% of organizations that suffered an AI-related breach lacked basic AI access controls, and 63% had no AI governance policy at all. The danger comes from deploying AI faster than the controls around it. Governance is what converts speed into something a CISO can defend.

THREE CONTROLS MAKE FASTER OPERATIONS DEFENSIBLE

Traceability. Every detection and every decision traces back to its source. Telemetry is normalized to a common schema, de-duplicated, and enriched with intelligence at ingestion, so an analyst or an auditor can follow any conclusion to the specific data and threat intelligence that produced it. Reasoning that cannot be traced cannot be trusted, and the platform is built so that it always can be.

Auditability. Assisted and autonomous actions are logged, explainable, and reversible, with human oversight built into every level of the autonomy model. Compliance evidence is produced from operational data rather than reconstructed in a separate, parallel exercise, which shortens audit cycles and removes a recurring source of manual effort.

Quality thresholds. Not every signal earns an analyst's attention, and not every action clears the bar to run without a human. Detections carry confidence scoring, analysts supervise and decide, and the degree of autonomy is calibrated to the risk of the action.

HOW LEADERSHIP SHOULD MEASURE IT

The outcomes worth tracking are the ones that show work has been removed from the human queue, not just moved around it. Leadership should look at four things:

- Fewer alerts requiring human review
- Shorter investigations
- Reduced escalation to senior engineers
- Faster time to decision

Time to decision deserves the most weight. Faster alerting means little if it still takes hours to confirm whether an incident is real, and IBM's data is consistent on the point that longer detection and containment directly raise breach cost.

There's a second signal worth paying attention to: whether the team trusts the system. Trust becomes visible when analysts stop re-checking automated output, when established responses run without hesitation, and when senior engineers are no longer pulled in to validate routine decisions. Where that trust is missing, the same tooling adds cognitive load instead of relieving it, which is exactly how "productivity" investments quietly fail.

Most security leaders have already decided to adopt AI. The harder question is what form of it they can stand behind in front of a board, an auditor, or a regulator. Doing more with the same team is achievable today. Keeping full traceability, auditability, and control while doing it is what makes the investment defensible, and what separates a productivity story leadership can sign off on from one it cannot.