RUNNING ON EMPTY

Melissa Bischoping
identifies a new form of fatigue in cyber security

Picture credit: iStock

Security teams are no strangers to fatigue. Over the years, they've dealt with a constant stream of alerts, managed increasingly complex toolsets, and defended IT environments that change by the hour. It's demanding work and even the best-prepared teams can feel the strain of keeping pace with today's threat landscape. Recent research underscores this pressure: the 2025 Pulse of the AI SOC report found that 66 percent of security analysts describe their workload as 'unsustainable,' and 73 percent have experienced alert-related burnout in the past year.

Now, artificial intelligence (AI) is reshaping that landscape once again. As AI systems take agentic workflows for defence, a new version of this same challenge is emerging: agent fatigue. Where teams previously struggled with alert overload, they are now grappling with the cognitive strain of overseeing and approving AI-driven actions at speed and scale. Budgets continue to get tighter, yet expectations are higher. In response, many organisations are embracing AI to detect, classify and remediate threats in real-time. These systems promise faster, smarter, more autonomous security. But without clarity into how AI reaches its decisions – and the confidence to question them – fatigue can quickly resurface in a new form, with default acceptance of what it's suggesting.

We're entering a transitional period and the next phase of cyber security maturity depends on trust: trust in automation, in human judgment, and in the visibility that connects the two. When teams can clearly see what AI is doing, why and where, fatigue can be swiftly replaced with confidence.

The push towards agentic security tools and AI-powered capabilities has changed the rhythm of security teams. Where analysts once fought to keep up with endless notifications, they're now managing a flow of automated decisions – each demanding context, confidence and approval. The noise hasn't disappeared; it's simply evolved. AI agents can now isolate devices, deploy patches or adjust policies in seconds. Analysts remain accountable for the outcome, but often without the full picture: which systems are affected, what dependencies exist or what the business impact might be. Over time, those gaps turn confident oversight into routine 'rubber stamp' approval.

Where analysts once fought to keep up with endless notifications, they're now having to manage a flow of automated decisions

Agent fatigue doesn't mean automation has failed; it means that our analysts and engineers must hone their critical thinking and technical prowess in order to trust – but still verify – as the ecosystem continues to scale in complexity. Today's challenge is keeping pace with decisions made faster than humans can assess them.

AI's strength is speed, but decisions made in milliseconds may not provide the context that humans rely on. An isolation request, a patch push or a policy change may protect one system while disrupting another. Without real-time visibility into dependencies and impact enriched by environmental nuance, even simple approvals can feel uncertain.

This is the confidence gap – the point where automation's pace exceeds human understanding. You've still got control, but you've lost clarity. When analysts can't see how an agent's recommendation was reached, trust becomes fragile and oversight turns reactive. This doesn't mean that we need to limit automation. It just needs to be built on clarity and context. Teams need to know what every endpoint is running, what's changed and why an action is being proposed. That transparency turns fatigue into focus. With complete, real-time data, analysts can verify AI recommendations quickly and confidently - allowing automation to empower rather than overwhelm. Contextual understanding of the business and technology environment is something that, at least today, is best and solely provided by the human experts in the process.

Technology alone can't close the trust gap; culture plays an equal part. For years, cyber security has rewarded speed with metrics like mean time to detect and mean time to respond, driving teams to act fast, sometimes at the expense of reflection. AI accelerates that tempo even further, introducing new layers of cognitive pressure and decision complexity.

A sustainable security culture recognises that pressure and creates space for context. Cyber security experts should feel confident questioning recommendations, asking for clarification and taking the time to understand why a system has made a decision. That confidence grows when leaders actively encourage it. Managers and CISOs set the tone – showing that caution isn't hesitation, it's professionalism. This is how junior analysts become trusted senior experts with wisdom, not just technical knowledge. When leaders treat AI as a co-pilot rather than a replacement for expertise, it reinforces the value of human judgment at every level.

Empowering people to think critically also strengthens team wellbeing. Many analysts enter the field because they value curiosity and problem solving, not because they want to rubber stamp machine outputs. By giving them the autonomy to challenge, interpret and learn from AI decisions, security leaders preserve that sense of purpose among their team members. The result is a healthier relationship between people and technology; one built on collaboration and informed confidence.

Culture lays the foundation for trust, but design determines whether or not it lasts. The right systems make human oversight intuitive, not optional. Re-establishing confidence requires transparency and feedback to be embedded into every stage of the decision-making cycle.

The best way to achieve this is through a "human-in-the-loop" approach. It ensures teams remain active participants with full visibility into what AI is doing and why. When experts understand the reasoning behind a recommendation – and the context of the systems it affects – they can validate its logic, apply judgment and act with confidence.

This relationship depends on transparency in both directions. Humans need insight into how models operate, while AI systems must be designed to adapt from human input. Each adjustment or override should improve the model's future performance, creating a continuous feedback loop that strengthens trust on both sides. When oversight is structured this way, AI and human judgment reinforce one another. Analysts gain confidence through clarity, while automation gains precision through human experience. Over time, this partnership builds resilience – a balanced relationship where technology enhances decision making rather than overwhelming it.

HOW DO ORGANISATIONS EVOLVE THEIR PEOPLE, PROCESSES AND PRINCIPLES TO KEEP PACE?

Strong culture and thoughtful design can keep human oversight alive, but leadership and ongoing training determines whether it scales. As AI becomes more deeply embedded in security workflows, governance must evolve to ensure that efficiency never comes at the cost of accountability. The goal is to define how automation operates and how its actions are understood, reviewed and improved over time.

The UK Government's 2025 Cyber Security Skills Report found that only 42 percent of organisations deploying AI tools have provided formal AI training to staff – leaving most experts ill-equipped to interpret or challenge automated decisions. This skills gap highlights a growing governance challenge: ensuring that human capability evolves alongside technological capability.

For CISOs and security leaders, that means establishing clear frameworks for when and how AI is allowed to act. Some tasks can be executed autonomously, others should trigger approval and all must be auditable. Every automated decision – from a simple patch deployment to a containment action – should leave a transparent record of its reasoning, outcome and any subsequent human intervention. That auditability is what turns automation from a black box into a source of confidence. This also becomes a source of data an organisation can learn from to optimise how and where it implements these solutions.

Governance also extends to evaluation. Leaders should regularly review whether automation is truly reducing risk or simply shifting it elsewhere, such as creating blind spots in oversight or over reliance on machine judgment. Treating AI outputs as living data rather than static truth encourages this kind of ongoing review and improvement. Much like we are never truly 'done' with security work, we will never be 'done' with AI - it is a part of our ecosystems that must adapt and evolve alongside the business.

Perhaps most importantly, governance sets the tone for trust. When leaders demonstrate that they can explain and justify how automated systems make decisions and where the humans retain ultimate authority, it reassures both internal teams and external stakeholders. Boards, regulators and customers increasingly expect this level of clarity. Building those accountability mechanisms early creates a culture where innovation and responsibility advance together.

EMPOWERING PEOPLE TO THINK CRITICALLY ALSO STRENGTHENS TEAM WELLBEING

When governance implements accountability mechanisms, every AI-driven decision, whether successful or flawed, provides insight into how systems and teams can adapt together. Mature organisations treat these moments as learning exercises rather than failures, using them to strengthen both processes and people.

When an automated action produces an unexpected result, the focus should shift quickly to understanding why. Was key context missing? Was the model's data incomplete or hallucinated? Did the approval process give analysts enough accurate, relevant information to intervene effectively? Asking such questions turns isolated incidents into collective progress. This continuous learning mindset translates to focus. Analysts see feedback influences how systems behave as AI becomes better aligned with human judgment. Over time, both sides adapt to create a cycle of trust built on shared experience and knowledge.

Fatigue caused by AI agents in this light becomes a signal that something needs adjusting. It points to gaps in visibility, communication or governance that can be corrected before they undermine resilience. Perhaps it points to an over zealous adoption - too much, too soon - and warrants revisiting adoption strategy to take more manageable steps. By responding to those signals early, organisations build a more confident partnership between human and machine.

The rise of agentic AI and automation has moved beyond tech just being a tool. It's becoming a collaborator, capable of acting and adapting alongside its human counterparts. The question for the coming years as AI becomes a permanently interwoven part of our technology ecosystem is how organisations will evolve their people, processes and principles to keep pace.

True resilience will depend less on reacting to threats and more on managing relationships between data and decision, automation and oversight, and trust and transparency. The ability to understand why a system acts and to proactively change what it does next time, not just see what it does, will become the new measure of maturity.

As security grows more autonomous, the human role won't disappear; it will continue to shift. Analysts will become curators of context - guiding, interpreting and refining the choices made by machines. Leaders will move to defining principles, ensuring that AI systems reflect the organisation's values as much as its risk tolerance. The organisations that succeed will be those that pair automation's speed with human judgment, data with discernment and insight with integrity. When that balance is achieved, the fatigue that once drained security teams can become something else entirely: a catalyst for confidence, clarity and enduring trust ●

Melissa Bischoping
is Director of Endpoint Security Research at Tanium.

Today's challenge is keeping pace with decisions made faster than humans can properly assess them

Picture credit: iStock