What's possible with AI in IT operations and security

Evidence current to mid-2026.

A clear breakdown of where AI does and does not fit across the core tasks of an IT operations and security function. A practical starting point, not the last word. Evidence current to mid-2026.

Where AI is mature: strong fit today

Proven and available today. AI does the bulk of the work and the IT team reviews it.

Function	The job today	With AI
Service desk ticket triage and routing	Staff submit requests by email, Teams or phone; IT person reads, classifies and routes each one, with back-and-forth to clarify what the user actually needs	AI classifies incoming tickets, suggests a priority and routes to the right queue; IT person handles escalations, edge cases and anything requiring judgement
Vulnerability scan report summarisation	Run the scanner, read through hundreds of Common Vulnerabilities and Exposures (CVE) entries, work out which ones apply to the actual environment and write up the action items	AI reads the scanner output and produces a prioritised summary of findings mapped to the real estate; the IT person confirms which apply and owns every remediation decision
Patch compliance gap reporting	Manually check patching status across devices and systems, build a list of what is patched and what is not, chase down the gaps	AI monitors patching status and generates gap reports by device, OS and severity level; the IT person reviews and drives remediation
Security content and policy drafting	Write phishing simulation emails, security bulletins, acceptable use policies and awareness training content from scratch	AI drafts from a brief; IT manager reviews against Essential Eight controls and the Information Security Manual (ISM) baseline, adapts to the business context and owns sign-off
Incident documentation and post-incident reports	Piece together what happened from memory, system logs and chat threads after an incident, then write up the timeline and findings	AI drafts the incident timeline from log summaries, alert data and notes; the IT person checks accuracy, fills in the context and owns the official record

Where AI is emerging: consider piloting with a human gate

Promising but not yet proven at this scale. AI assists and IT staff stay in the loop, so trial it on a contained scope first.

Function	The job today	With AI
Security operations centre (SOC) alert triage	Analyst works through the alert queue in sequence; each alert needs context lookup, judgement and a disposition; alert volume means real threats get buried in noise	AI scores and prioritises alerts, suppresses known false positives and surfaces the most likely true positives; analyst investigates escalations and owns every containment call
Phishing email triage	Staff report suspicious emails; analyst reviews each one manually, looks up URLs and senders and decides whether it is malicious	AI pre-scores and classifies each submission and flags high-confidence indicators; analyst confirms the verdict and handles edge cases. AI does not block autonomously
Vulnerability prioritisation	Team prioritises remediation by Common Vulnerability Scoring System (CVSS) severity score alone, making it hard to separate actually-exploitable findings from high-severity-but-never-attacked ones	AI combines CVSS severity, Exploit Prediction Scoring System (EPSS) exploit-probability and Known Exploited Vulnerabilities data into a ranked shortlist; the IT person owns the remediation schedule

Function	The job today	With AI
Identity and access review	IT manager manually reviews access rights in batches from spreadsheets or the directory, working out who has what and whether it still makes sense	AI surfaces anomalous permission patterns, overprivileged accounts and unused access that has accumulated over time; IT manager owns and executes all access changes
AI for IT operations (AIOps) event correlation and noise reduction	On-call engineer manages a flood of monitoring alerts during an incident, manually linking related events and trying to find the root cause in the noise	AI correlates and deduplicates events across monitoring sources, surfaces root-cause candidates and suppresses noise; engineer investigates and decides
Incident response playbook drafting	Security lead writes incident response (IR) playbooks from scratch, working from Australian Cyber Security Centre (ACSC) or US National Institute of Standards and Technology (NIST) guidance and adapting to the specific environment	AI drafts the playbook from a brief and applicable frameworks; security lead validates the steps, adapts to the environment and owns the final document
Endpoint hardening scripts and cloud provisioning	IT generalist writes scripts to apply hardening controls or configures cloud infrastructure manually in the console	AI drafts scripts and configuration from a brief; the IT person reviews every line, tests in a non-production environment and owns deployment. Never apply untested AI-generated config to live systems

Where AI is not ready or suitable today: keep with people

These stay with people, either because the tooling is not reliable enough or because controls and compliance rule it out.

- **Access revocation decisions.** AI may surface candidates for review. The IT team must own and execute every removal: access left live after a departure or role change is one of the most common precursors to data theft.
- **Incident response judgement calls.** Containment scope, escalation path, breach notification timing and crisis communications under pressure. The IT manager or incident lead leads the response; AI supports with log summaries and draft comms at most.
- **Security-risk sign-off and ransom-payment decisions.** The business owner or head of IT accepts a residual security risk, and decides whether to pay a ransom under extortion. AI can summarise the options and consequences but must not decide. A ransom payment also triggers a mandatory report under the Cyber Security Act 2024.
- **Auto-remediation of production systems.** AI agents can reverse manual emergency patches not yet committed to the repository. The IT team approves every change to live systems. The human approval gate is the control.

Worth weighing: for any of these, the upfront work (connecting tooling, defining rules and response criteria, cleaning the asset inventory) is a one-off setup; the ongoing job on the strong-fit items is mostly reviewing AI output and handling escalations, not the original grind. That time goes back into the hardening, the judgement calls and the threats that actually matter.

One practical note on the threat side: AI is also the attacker's tool. AI-written phishing emails and deepfake voice and video have made business email compromise (BEC) and executive-impersonation fraud cheaper and more convincing. That is not a reason to avoid AI tools; it is a reason the human gate on payments, access changes and unusual requests matters more, not less.