

SPOTLIGHT REPORT · 2026

AI agents and the impact on cybersecurity

Industry leaders and cybersecurity experts on how AI agents are reshaping the threat landscape, security operations, the skills the field now demands and how we adapt teaching and learning.

Version final · 19 August 2026

AI Workforce Consortium

Members: Accenture · Cisco · Cornerstone · Eightfold · Google · Honeywell · IBM · Indeed · Intel · Microsoft · Pearson · SAP

Advisors: College Board, DIGITALEUROPE, the Business-Higher Education Forum (BHEF), the European Vocational Training Association and the National Applied AI Consortium.

Contents

Executive summary

01 TECHNOLOGY

- 1.1 Growing AI capabilities accelerate cybercrime
- 1.2 More than a third of security leaders plan investments in AI-powered defense a priority
- 1.3 Use cases of agentic AI deployed for defense
- 1.4 Governance drives AI success
- 1.5 Recommendations

02 WORKFORCE

- 2.1 Cybersecurity job demand has grown 9.5% in the last six months
- 2.2 AI skill integration doubled year over year
- 2.3 How cybersecurity roles are evolving
- 2.4 Recommendations

03 SKILLS

- 3.1 A unified skills framework for the AI-enabled knowledge worker
- 3.2 Technical Domain skills

- 3.3 Lead AI Agents skills

- 3.4 Interpersonal skills

- 3.5 Business Context skills

- 3.6 Intrapersonal skills

- 3.7 Recommendations

04 LEARNING

- 4.1 The increasing demand for experience in cybersecurity

- 4.2 The learning transformation to respond to new needs

- 4.3 Recommendations

05 APPENDIX

- 5.1 About the AI Workforce Consortium

- 5.2 Methodology

- 5.3 Developing the skills framework

- 5.4 Key definitions

- 5.5 Career Signal Radar Prompt

- 5.6 Resources

- 5.7 Terms of use and disclaimer

Figures are numbered 1 to 13 and appear at the point they are discussed.

Executive summary

The cybersecurity world is changing fast, and agentic AI is the main reason why. This technology provides defenders with unprecedented speed and scale, but it also empowers attackers to find vulnerabilities and bypass controls using AI agents that outperform most human experts

Across the cybersecurity workforce, the work is shifting from manual execution to strategic orchestration: people set the guardrails, the objectives and the measures of success, and agents execute inside them. In the [World Economic Forum's Global Cybersecurity Outlook 2026](#), 94% of respondents expect AI to be the most significant driver of change in cybersecurity in the year ahead, up from 66% in the 2025 edition.

The last eight months of incident reports involving AI agents show why. In November 2025 [Anthropic](#) disclosed a cyber-espionage campaign in which AI executed 80–90% of the operation. In July 2026 [Sysdig](#) documented what it called the first case of 100% agentic ransomware, a complete extortion operation driven end to end by a large language model (LLM), which gained entry through an unpatched internet-facing service. Later that month, Hugging Face disclosed an intrusion "driven, entirely, by an autonomous AI agent system." OpenAI considered this an "unprecedented cyber incident" after confirming it involved one of its unpublished models running with reduced cyber refusals during an internal evaluation. On 22 June 2026 the [Five Eyes cyber security agencies](#) urged a collective response: "The timeline is not years, it is months."

This report explores that transformation in four parts: how technology is reshaping the threat landscape, how this shift is changing job roles, which skills matter most now, and how learning must change to prepare people for what comes next.

Note: This report serves as a companion to the webinar "AI Agents and the Impact on Cybersecurity," recorded live on May 21, 2026.

Technology

How growing AI capabilities are reshaping the cyber threat landscape, and how security leaders, AI labs and governments are responding.

1.1 Growing AI capabilities accelerate cybercrime

As generative AI evolves from content generation toward reasoning and agentic execution, it has progressively reshaped the cyber threat landscape, increasing speed, scale, and sophistication. This follows a well-documented pattern known as “scaling laws”: AI model general capability improves predictably as compute, data, and model size increase. Capability on specific tasks, including offensive-security tasks, emerges unevenly and must be measured directly. If scaling holds, capabilities relevant to attackers should be expected to keep improving. Organization readiness cannot wait.

How modern AI became a threat vector

The initial wave of generative AI, roughly spanning 2022 to 2024, primarily functioned as a catalyst for less experienced threat actors. During this period, frontier model chatbots, when jailbroken, lowered the barrier to entry for cybercrime, simplifying the discovery of exploitation techniques and enabling novice attackers to operate above their skill level. This early phase was defined by knowledge assistance rather than autonomous action.

The agentic shift

The cybersecurity threat changed materially in 2025 with the arrival of reasoning models and agentic AI: frontier systems could now plan and execute complex, self-directed workflows rather than merely retrieve knowledge. In November 2025, Anthropic reported suspicious activity in the use of its Claude agent, which it later identified as a highly sophisticated cyber-espionage campaign that they attributed, with high confidence, to a state-sponsored group it designates GTG-1002. The AI agent ran 80–90% of the operation autonomously, with limited human direction.¹

¹Anthropic, “Disrupting the first reported AI-orchestrated cyber espionage campaign,” November 2025.

anthropic.com/news/disrupting-AI-espionage The report also describes human operators intervening at four to six decision points; this report leads with the 80–90% figure throughout for consistency.

The 2026 IBM X-Force Threat Intelligence Index reported a 44% year-over-year rise in attacks exploiting public-facing applications, which IBM attributes largely to missing authentication controls alongside AI-enabled vulnerability discovery.² In early July 2026, Sysdig's Threat Research Team documented the first fully autonomous ransomware operation,³ two weeks later also Hugging Face disclosed an incident "driven, end to end, by an autonomous AI agent system."⁴ Soon after that, OpenAI disclosed that the incident involved one of its unpublished models running with reduced cyber refusals during an internal evaluation, describing it as an "unprecedented cyber incident".

The Mythos moment (2026)

Frontier agents such as Claude Mythos Preview are now demonstrating capabilities that Anthropic describes as "better than all but the most skilled humans" in vulnerability discovery, and in chaining flaws to navigate defensive perimeters and bypass sophisticated sandboxes without human direction.

This new wave of capabilities compounds on prior shifts in knowledge dissemination and agentic automation. As frontier models evolve from passive assistance to autonomous offensive execution matching expert human proficiency, they fundamentally reshape the cyber threat landscape.

These events mark an inflection point. AI's ability to coordinate multi-stage, multi-target intrusions at machine speed is now an observed reality rather than a theoretical concern.

²IBM X-Force, *X-Force Threat Intelligence Index 2026*, 25 February 2026. ibm.com/reports/threat-intelligence


³Sysdig Threat Research Team, "JadePuffer: agentic ransomware for automated database extortion," July 2026. sysdig.com/blog/jadepuffer-agentic-ransomware-for-automated-database-extortion

⁴OpenAI statement of 21 July 2026: the incident involved GPT-5.6 Sol "and an even more capable pre-release model, all with reduced cyber refusals for evaluation purposes," during internal ExploitGym testing. openai.com/index/hugging-face-model-evaluation-security-incident

FIGURE 1

The three compounding waves of AI-augmented cybercrime

The three compounding waves of AI-augmented cybercrime



Source: Consortium synthesis. Qualitative interpretation of the compounding effect.

WAVE	WHAT IT DOES	IMPACT
First wave, 2022–24: chatbots	Extend the reach	Jailbroken chatbots put existing attack techniques in the hands of people who can then operate above their skill level. Diffuses attack knowledge and lowers the barrier to entry. Knowledge assistance, not autonomous action: a human still runs every step.
Second wave, 2025–26: agentic automation	Accelerate execution	Runs existing attack techniques end to end, without a human in the loop. Brings scale and sophistication to groups that previously lacked both. Human involvement falls from a supervised campaign (Anthropic, November 2025, 80–90% AI-executed) to none in execution (Sysdig, July 2026, 100%).
Third wave, 2026: autonomous discovery	Enable discovery	Finds vulnerabilities and attack paths that were not previously known. Puts zero-day capability, once nation-state only, within reach. Restricted to vetted partners today, though Epoch AI measures open-weight models trailing the frontier by about four months.

A responsible path forward

In response, frontier AI labs are prioritizing safety by restricting access to their most advanced capabilities. This is taking two broad forms.

- **Targeted collaborations.** Initiatives such as Anthropic's Project Glasswing aim to "secure the world's most critical software and provide defenders with a durable, strategic advantage." OpenAI's Daybreak initiative takes a similar approach, offering tools such as Codex Security and GPT-5.5-Cyber to help organizations "find, validate, and patch vulnerabilities at scale."
- **Guarded public releases.** For public models, AI labs are building-in guardrails to prevent misuse in high-stakes domains such as cyber and biology (for example, Anthropic's Fable 5). In some cases, they limit early preview access to trusted partners before a wider release, as OpenAI did with GPT-5.6 Sol, Terra, and Luna.

These measures are responsible, but they do not eliminate risk. Two factors warrant close attention.

- **Jailbreak susceptibility.** In principle, every guardrail remains vulnerable to bypass, potentially exposing the restricted, high-risk capabilities it was meant to contain.
- **Narrowing lead times.** Expert consensus, supported by data from Epoch AI, suggests that the advantage closed labs currently hold over the open-source community is shrinking, with the gap now measured in a few months.⁵ That raises the possibility that these advanced capabilities will soon reach open-weight models.⁶

Governments are responding. On 2 June 2026 the United States administration issued an executive order titled Promoting Advanced Artificial Intelligence Innovation and Security, establishing a voluntary "covered frontier model" designation under which developers may grant up to 30 days of pre-release government access, with no mandatory licensing requirement.⁷ Weeks later, on 22 June 2026, the Five Eyes cyber

⁵Epoch AI, "Open models lag state-of-the-art closed models by 4 months" (May 2026), measured on the Epoch Capabilities Index. epoch.ai/data-insights/open-closed-eci-gap

⁶On 16 July 2026 Moonshot AI announced Kimi K3, a 2.8-trillion-parameter model, and published its open weights on 27 July 2026. <https://www.interconnects.ai/p/kimi-k3-the-open-weights-escalation>.

⁷The White House, Executive Order, "Promoting Advanced Artificial Intelligence Innovation and Security," June 2026; published in the Federal Register 5 June 2026. <https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/>

security agencies issued⁸ a more urgent note on frontier AI models “anticipated to exceed current industry expectations, fundamentally transforming both offensive and defensive cyber capabilities. The timeline is not years, it is months.”

1.2 Leaders prioritize AI security investment.

Threat actors' increasing use of AI has rendered traditional operating models inadequate: today's threats operate at a velocity human analysts cannot match. In response, organizations worldwide are shifting to AI-powered defense strategies to keep pace with the new threat environment.

According to the WEF Global Cybersecurity Outlook 2026⁹, 94% of survey respondents anticipate that AI will be the most significant driver of change in cybersecurity in the year ahead, up from 66% in the 2025 edition.

Cisco's May 2026 survey of 8,000 security leaders across 30 markets and 14 industries shows the response taking shape across five priorities. Asked to select up to three investment areas for the next two years, leaders put investment in AI-powered defense first (36%), followed by improving detection and response time (32%), improving application security (27%), with investment in threat intelligence and investment in people tied at 25%.¹⁰

Investment in people ties for fourth place, well behind AI-powered defense. As set out in section 4.1, due to the increasing demand for experience in cybersecurity, leaders are investing heavily in advanced tools while reporting a widespread inability to find entry-level talent with the skills required to operate them.

Together these point toward a proactive, agentic defense model, in which humans define the guardrails and objectives while AI agents execute at speed and scale.

⁸ NCSC (UK), CISA, NSA, ACSC, CCCS and NCSC-NZ, joint statement on the AI shift in cyber threat, 22 June 2026. [ncsc.gov.uk — Five Eyes cyber security agencies statement](https://ncsc.gov.uk/five-eyes-cyber-security-agencies-statement).


⁹ World Economic Forum in collaboration with Accenture, *Global Cybersecurity Outlook 2026*, 12 January 2026. Fieldwork 25 August – 1 October 2025; 804 qualified respondents across 92 countries, including 316 CISOs. weforum.org/publications/global-cybersecurity-outlook-2026.

¹⁰ Cisco survey of 8,000 security leaders across 30 markets and 14 industries, fieldwork conducted in May 2026. Respondents could select up to three competencies, so percentages sum to more than 100%. See 5.2, Methodology. [Relentless Defense Report](#), September 2026.

FIGURE 2

Key areas for security investment, next two years

Key areas for security investment, next two years



Source: Cisco survey of 8,000 security leaders across 30 markets and 14 industries, fieldwork conducted in May 2026. Respondents could select up to three competencies, so percentages sum to more than 100%. [Relentless Defense Report](#), September 2026.

1.3 Use cases of agentic AI deployed for defense

Agentic AI is already deployed in defense: automating trust validation, unifying threat visibility, and closing the gap between offensive and defensive operations. The examples below are a small subset to exemplify use cases.

ORGANIZATION	PROGRAM	DESCRIPTION
Accenture	Agent Oliver	An AI capability for external attack-surface management, unifying enterprise-wide asset discovery and exposure analysis.
Cisco	The Foundry Security Spec	An open-source specification developed by Cisco that defines an architectural framework for evaluating software security using agentic AI. A model-agnostic blueprint—built around 8 core agent roles (such as Orchestrator, Detector, and Validator), ~130 functional requirements, and strict guardrails—to systematically detect, verify, and report vulnerabilities with traceable, auditable proof.


ORGANIZATION	PROGRAM	DESCRIPTION
European Commission	Action Plan on Cybersecurity and AI, COM(2026) 577	Commits €200 million of EU funding, drawn from existing programmes, to secure AI defense, simulation platforms and automated exploit healing. The funding is committed rather than disbursed.
DARPA	AI Cyber Challenge (AIxCC)	A two-year competition in which seven finalist teams built AI systems to find and patch vulnerabilities in real open-source code. All seven cyber reasoning systems were open-sourced.
Google	CodeMender	Google uses AI agents like CodeMender to automatically detect and patch software vulnerabilities at scale.
IBM	AI-powered purple teaming	Agentic AI embedded in researchers' daily workflows to accelerate deep exploration of complex security problems and close the gap between red and blue teams.

Source: Descriptions supplied by the named organizations, July 2026.

1.4 Governance drives AI success

AI systems and agents are also an attack surface, and most organizations are deploying them faster than they can govern them. The oversight gap has a measurable cost.

IBM's Cost of a Data Breach Report 2025 found that 63% of breached organizations had no AI governance policy or were still developing one; one in five organizations reported a breach involving "shadow AI" — tools adopted outside official channels; organizations with high levels of shadow AI carried an average of \$670,000 in additional breach cost; and 97% of organizations that suffered an AI-related security incident lacked basic AI access controls.¹¹



¹¹IBM Security and Ponemon Institute, *Cost of a Data Breach Report 2025*, 30 July 2025. ibm.com/reports/data-breach. Summary of the AI findings: newsroom.ibm.com

1 in 5	Suffered a breach involving “shadow AI” — tools adopted outside official channels
+\$670K	Average extra breach cost for organizations with high levels of shadow AI
97%	Of victims of direct AI-model compromise lacked basic AI access controls

Source: IBM’s Cost of a Data Breach Report 2025.

These are not failures of detection technology. They are failures of governance: policy, access control, and visibility.

Why does governance lag?

The standards themselves are not missing. The [NIST AI Risk Management Framework](#) and [ISO/IEC 42001](#) supply the management structure; [OWASP’s Top 10 for Agentic Applications](#) and [MITRE ATLAS](#) supply the threat model, covering prompt injection, tool poisoning and excessive agency.

The Consortium’s ICT in Motion report (2025) found that demand for AI governance skills across the 50 ICT roles analyzed is up 150%, while demand for AI ethics skills is up 125%. Both skills sit among the critical skill gaps, with fewer than 30% of roles able to meet current demand.¹²

So, organizations are not avoiding AI oversight for lack of a standard. They are struggling to implement one, because they have not built the workforce capacity to carry it.

FROM THE PANEL

Two practitioners on the same point: authority over an agent must be narrow, named and governed.

“The organizations that succeed won’t be the ones with the most automation — they’ll be the ones with the strongest governance around autonomous systems.”

Kyu Kwak · CISO, Pearson

¹²AI Workforce Consortium, *ICT in Motion* (2025), §4.2, Technical Skill Gap.

"[In implementing your agentic strategy for cybersecurity] don't boil the ocean. Pick one workflow and give an agent real authority over it... The formula from my side is: narrow scope, real authority, and named owner."

Omar Santos · Distinguished Engineer, Cisco

Note: Panel views, see 5.2 section.

1.5 Recommendations

FOR EVERYONE — LEADERS, WORKFORCE AND EDUCATORS

- **Treat AI safety as joint accountability, as cybersecurity already is.** Builders must embed safety and transparency into design and training. Deployers must enforce governance, access controls, and monitoring before go-live. Users must be equipped to use AI critically and flag anomalies. No single actor can carry it.

FOR LEADERS

- **Track where AI capability is heading.** Treat frontier-level AI capability as present and diffusing, and assess risk and readiness on an ongoing basis rather than at a single point in time.
- **Stay informed from primary sources** — frontier labs' model cards and independent research, including the [UK AI Security Institute](#).
- **Accelerate the foundations before the AI-specific work.** Assess risk, evaluate readiness, establish clear accountability and confirm controls are in place: asset and identity visibility, zero-trust architecture, phishing-resistant multi-factor authentication (MFA), patch and vulnerability management. Layer AI-specific risk governance on top of that.

FOR THE WORKFORCE

- **Prioritize upskilling in agentic AI** and build a method for staying current rather than treating it as a one-off course.¹³
- **Insist on hands-on lab practice, not theory alone.** Working with these systems directly is what builds a concrete sense of what they can do today and how they are improving.

¹³See appendix 5.5 for a method for running weekly agentic research for continuous learning and updates.

FOR EDUCATORS

- **Prepare students for entry-level roles that already assume agentic AI proficiency** and effective human-AI collaboration. Update curricula to cover agentic vulnerability discovery and the use of AI agents in defense.¹⁴
- **Teach agentic secure development against a named framework** in software development courses — the [Foundry Security Spec](#), [NIST SP 800-218A](#) for secure software development lifecycle (SDLC) practices, and the [CSA AI Controls Matrix](#) where controls need to map to ISO/IEC 42001.
- **Anchor course outcomes to a workforce framework**, so graduates' skills are legible to employers — the [NICE Workforce Framework for Cybersecurity](#) and its [AI Security competency area](#) in the US, [ENISA's European Cybersecurity Skills Framework \(ECSF\)](#) in Europe.
- **Teach and assess responsible AI in context.** Practice setting agent boundaries, reviewing output, escalating exceptions and documenting decisions inside a specific cybersecurity function — not as general AI principles taught in the abstract.
- **Pair curriculum design with continuous teacher professional development.** Use [UNESCO's AI Competency Framework for Teachers](#) to structure that development, so educators are equipped to teach with AI tools as well as about them.

¹⁴Consider small language models fine-tuned for a specific task, such as Antares — a pair of open-weight models (350M and 1B parameters) fine-tuned for vulnerability localization, released by Cisco Foundation AI on 21 July 2026. blogs.cisco.com/ai/introducing-antares

Workforce

How cybersecurity job demand is evolving across G7 economies, and how agentic AI is redefining existing roles and creating new ones.

2.1 Cybersecurity job demand has grown 9.5% in the last six months


Cybersecurity job demand across G7 economies grew 9.5% in the period October 2025 – March 2026 against the same six months a year earlier. In the preceding half-year, April – September 2025, growth was 6.8%. The trend is accelerating, and it reflects a consistent need for specialized talent. Security Engineer and Cybersecurity Analyst remain the most frequently listed roles across G7 cybersecurity postings.¹⁵

FIGURE 3

Cybersecurity job demand growth, G7 economies

Cybersecurity job demand growth, G7 economies

YEAR-OVER-YEAR GROWTH IN ALL G7 CYBERSECURITY POSTINGS




Source: Indeed demand-side data, G7 economies. Each half-year is compared with the same half-year a year earlier.

¹⁵Source: Indeed demand-side data, G7 economies. Each half-year is compared with the same half-year a year earlier.

FIGURE 4

Distribution of roles requested in the overall demand (G7 Economies)

Distribution of roles requested in the overall demand (G7 Economies)



Source: Cornerstone & Indeed demand-side data, trailing 12 months. Represents the share of job postings across G7 economies.

2.2 AI skill integration doubled year over year

AI skill integration is the share of job postings for a role that require AI skills. It tracks the rise of AI-related skills within professional roles as expressed in job postings.

The share of G7 cybersecurity postings requiring AI skills in Oct 2025 – Mar 2026 was twice the share a year earlier: a six-month average of 28.5%, against 14.2% in Oct 2024 – Mar 2025, with the latest month at 29.7%. That moves cybersecurity from initial integration to significant integration, using the bands introduced in ICT in Motion Report (AI Workforce Consortium, 2025) and defined in section 5.4. (Figure 5)

Alongside this, an “agentic skill stack” is emerging as a baseline for high-volume cybersecurity roles such as security engineering, cloud security, and detection and response engineering. Five pillars recur in AI-integrated cybersecurity postings: Python, Prompt and context engineering, AI security, agent orchestration and machine learning operations (MLOps). (Figure 6)

x2

AI skills appearing in G7 cybersecurity postings, Oct 2025 – Mar 2026 against a year earlier (28.5% against 14.2%)

>25%	Cybersecurity moved from initial to significant AI integration
5	Pillars of the agentic skill stack: Python, Prompt and context engineering, AI security, agent orchestration and MLOps

Lightcast advertised-salary data covering the six-month window, October 2025 – March 2026, puts the median advertised salary in cybersecurity postings that require AI skills 14.9% above the median for all cybersecurity postings in the US.


The premium is real: employers do pay more when a cybersecurity posting asks for AI skills.

14.9%	Median advertised-salary premium for AI-skill postings at the end of the window
--------------	---

Source: Lightcast Data, 2026 Cybersecurity Sector United States 2024-2026.

FIGURE 5

Average AI skills integration for cybersecurity roles, G7 economies




Source: Indeed, based on keyword analysis of job postings following the methodology of the Indeed Hiring Lab AI Tracker. G7 data, April 2024 to March 2026. Shaded columns are the two six-month comparison windows. The Oct 2025 – Mar 2026 average is 2.0× the Oct 2024 – Mar 2025 average.

2025 average, and the series crossed the significant-integration threshold in August 2025. The chart highlights two of the five integration bands; all five are defined in section 5.4.

FIGURE 6

Top in-demand skills for agentic AI security roles

Top in-demand skills for agentic AI security roles



Source: Cornerstone & Indeed demand-side data, latest 12 months for G7 economies. The chart shows the share of agent-tagged cyber job postings requiring each skill, illustrating the agentic skill stack that employers are seeking. Percentages represent the proportion of agent-tagged postings that reference each skill and should not be interpreted as mutually exclusive, as a single posting may require multiple skills.

2.3 How cybersecurity roles are evolving

AI skills are not new to cybersecurity. What has changed in the last few years with the emergence of agents built on large language models (LLMs), is how pervasive they have become, and how demand has shifted from building models to directing and verifying them.

The two job profiles below illustrate this transformation (Figure 7 and 8).

The SOC analyst role is being rapidly redefined as agentic AI systems assume high-volume, repetitive tasks such as triaging alerts, correlating threat intelligence feeds, and executing standard decision-tree workflows. This shifts the analyst’s primary function from manual processing to orchestration and oversight. Technical requirements are shifting from SIEM platform proficiency and query writing toward agent supervision and output validation, while durable skills remain more essential than ever.

Agents do not replace the analyst's technical foundation. Nobody can check a machine's reasoning about a packet capture without being able to read one. The work shifts from producing findings to directing the systems that produce them and answering for what they get wrong.

FROM THE PANEL

Two practitioners describing, from different angles, the same emerging mandate: Cybersecurity roles are evolving

“One of the roles that is getting reshaped the fastest is the tier 1 SOC analyst. That job involves triaging alerts, pulling information from tools like Splunk, checking threat intelligence feeds, and correlating data. Those repetitive tasks and decision-tree workflows are where agentic systems are now generally effective. Everybody is an orchestrator now, so it is reshaping that work.”

Omar Santos · Distinguished Engineer, Cisco

“The rise of AI-driven threats — such as deep fakes and unusual activity within traditional threat vectors — highlights the need for new roles like AI forensics experts or agent security experts who can dissect these complex, modern attacks.”

Kyu Kwak · CISO, Pearson

Note: Panel views, see 5.2 section.

FIGURE 7

The SOC operator: from manual analyst to orchestrator

Security operations center (SOC) analyst — monitors, investigates and responds to threats against organizational networks and data.

PRE-AGENTIC ERA		AGENTIC ERA
Technical skills	Threat detection and analysis · SIEM platform proficiency · Search Processing Language (SPL) query writing · Incident investigation and response · Networking and operating system knowledge · Scripting · Customer service · Cyber threat intelligence · Endpoint detection and response · Firewall · Incident management	New skills: Investigation using AI-layer telemetry · AI agent supervision · AI governance and compliance · AI-assisted threat hunting · Risk-based decision making · AI output validation · Security workflow orchestration · Human-AI collaborative investigation · Automation exception handling · Prompt injection

	PRE-AGENTIC ERA	AGENTIC ERA
	- Incident response - Log analysis - Penetration testing - Triage - Vulnerability management	- recognition and response - Detection quality validation - Identity-centric security analysis - Cloud-native security monitoring
Durable skills	- Critical thinking - Problem solving - Communication - Continuous learning	- Critical thinking - Problem solving - Communication - Continuous learning - Collaboration across teams - Adaptability - Evaluation and judgment - Attention to detail - Working independently - Evidence synthesis - Business-context judgment - Risk evaluation - Assessment and issue analysis

Source: Cornerstone & Indeed demand-side data, latest 12 months for G7 economies. The chart shows the share of agent-tagged cyber job postings requiring each skill, illustrating the agentic skill stack that employers are seeking. Percentages represent the proportion of agent-tagged postings that reference each skill and should not be interpreted as mutually exclusive, as a single posting may require multiple skills

FIGURE 8

Elsewhere the shift creates roles outright, such as the AI forensics expert.

AI forensics expert: an emerging role in cybersecurity

AI forensics expert—AI forensic experts investigate artificial intelligence systems after they fail, cause harm, or come under legal challenge. The evidence is unfamiliar: training data, model weights, prompt histories and inference logs rather than a disk image or a packet capture, read for bias, poisoning, tampering and adversarial manipulation. Regulation is starting to guarantee it exists. The EU AI Act obliges high-risk systems to log events automatically across their lifetime so incidents can be reconstructed later (Article 12).

TECHNICAL SKILLS	DURABLE SKILLS
Digital forensics and forensic investigation · Artificial intelligence platforms and applied AI · Python and SQL · Data science and data analytics · Incident remediation · Cyber threat intelligence · Machine learning · Malware analysis · Electronic discovery · Legal and regulatory analysis · Regulatory compliance	Communication · Assessment and issue analysis · Collaboration across teams · Adaptability · Team orientation · Evaluation and judgment · Evidence synthesis · Problem solving · Attention to detail · Working independently

Note: Cornerstone posting analysis, G7. The table above presents the principal skills identified in AI-skilled forensics job postings during the 12-month period ending March 2026 (n = 197 postings). The digital forensics analyst title alone grew 453% in the analysis (May 2026). This is the fastest relative growth of any cybersecurity role. The absolute base is small, on the order of 150 to 200 AI-skilled forensics postings per six-month window in G7, or 4.2% of all forensics postings, so the finding is presented as an emerging-role signal.

2.4 Recommendations

FOR LEADERS

- **Implement role-specific learning** to empower workers to delegate to agents, setting guardrails, managing escalations, and verifying AI outputs.
- **Redesign the roles the automation changes.** Start with the profiles: name which tasks move to agents, what the human role becomes, and what supervision the new shape requires.
- **Publish the career path through the change.** If tier-one triage shrinks, say what the route to senior judgment now is: agent supervision, exception investigation, detection engineering.

FOR THE WORKFORCE

- **Use AI agents to scan emerging trends** and curate a personal knowledge base, shifting from manual research to automated expertise.¹⁶
- **Gain general and role-specific hands-on experience with agents** in secure, isolated environments to develop governance skills and a solid understanding of evolving capabilities.
- **Share workflows and lessons learned with colleagues** to accelerate collective AI fluency.

FOR EDUCATORS

- **Evolve the technical curriculum** so labs require students to audit, validate, and fact-check agent-generated security outputs rather than only performing manual tasks.
- **Prioritize durable skills** such as ethical reasoning, systems and critical thinking, empathy and linguistic precision.
- **Use experiential learning** to prepare students to navigate open-ended, ambiguous challenges, supporting independent inquiry and complex, real-world problem solving.
- **Partner with industry** to integrate real-world agentic workflows into project-based learning, giving direct visibility into evolving jobs as technology adoption progresses.

¹⁶See appendix 5.5 for a method for running weekly agentic research for continuous learning and updates.

Skills

A unified skills framework for the AI-enabled knowledge worker, built on five dimensions: Technical Domain, Lead AI Agents, Business Context, Interpersonal and Intrapersonal.

3.1 A unified skills framework for the AI-enabled knowledge worker

AI is changing how work gets done: workflows, team structures, what leaders spend their time on, and the required skills.

Cybersecurity skills are mapped in depth by [NICE \(NIST SP 800-181r1\)](#), [SFIA 9](#), the [European e-Competence Framework \(e-CF, EN 16234-1\)](#), [ESCO](#), and the World Economic Forum's [Future of Jobs Report](#). Each describes the knowledge and competences a person needs to perform the tasks of a defined job role.

That is the assumption now under pressure. When the work is done with AI, people cross the boundaries of their roles. The person moves from directly performing the task to delegating it to an agent while remaining accountable for the outcome. Workflows and teams organized around scarcity, such as hunting for a new vulnerability, become work organized around abundance, where triaging many findings is the constraint rather than producing one.

This report introduces a unified skills framework to help workers, leaders and educators frame the conversation about skills as work moves to agentic AI. It is a new lens over the existing frameworks, not a replacement for them. Instead of a static list of skills that are in or out, it proposes areas of skill and sets out why the mix within them is changing. It is published as a starting point for discussion and testing rather than a settled answer and will be updated as the technology and the market move.

FIGURE 9

A unified skills framework for the AI-enabled knowledge worker

A unified skills framework for the AI-enabled knowledge worker



Source: AI Workforce Consortium. The framework is built on five dimensions: Technical Domain, Lead AI Agents, Business Context, Interpersonal, and Intrapersonal. The dimensions are areas of emphasis, not a partition. If a competency such as critical thinking appears in more than one — it is named once in the dimension where the report argues it matters most and referenced from the others. For derivation see the appendix (5.3).

3.2 Technical Domain skills

AI agents are approaching human-level performance in some domains and surpassing it in others. Human proficiency in the domain nevertheless remains essential. The frontier of AI capability is jagged: strong on one task, unreliable on an adjacent one that looks no harder.¹⁷ Output is stochastic. Task delegation must therefore be risk-based.

¹⁷Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraye, L., Candelon, F., and Lakhani, K. R. (2023), "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality," Harvard Business School Working Paper 24-013. [hbs.edu — working paper](https://hbs.edu/working-paper). Published version: [Organization Science \(2025\)](#)

Execution can be delegated. Judgment cannot, and judgment requires understanding. Deciding what to hand over takes enough understanding of the work to weigh what failure would cost and to check what comes back. The judgment you never delegate is the judgment about what is delegable.

The risk is blind delegation, which could cost twice: skill erosion over time, and high-stakes mistakes in the moment.

As AI lets professionals work in adjacent roles, a broad first-principles grasp of the wider technical domain becomes more important. Without it, people make naive decisions and undervalue the expertise the work actually requires. Judging AI output in a domain that is not your own depends on knowing its foundations well enough to see what is missing. That leaves the practitioner with a call to make case by case: whether this particular output needs verification by a qualified expert.

KEY SKILLS

01 Technical breadth. Including institutional and industry tacit knowledge. Learn how your organization and sector actual work.

02 Foundational depth. Master the key principles and big ideas: identity, protocols, operating systems, data, cryptography and threat modeling.

03 First-principles reasoning. Hold a model of the system strong enough to contest an answer you did not produce.

AVOID

"Vibe coding" a business application, ignoring the secure software development lifecycle and delegating to the agent without reaching an internal expert for guidance or validation.

GOOD PRACTICE

A solid appreciation of software engineering principles — development lifecycle, version control, testing — lets a non-technical product manager build a prototype for idea socialization and validation.

FROM THE PANEL

Three practitioners on the same risk: expertise is what makes an agent's output checkable, and it decays when unused.

"Someone without domain expertise may be more willing to trust what an AI agent is saying and continuously approve its recommendations. One of the challenges will be ensuring people apply appropriate judgment rather than accepting outputs at face value."

Ian Molloy · Principal Research Staff Member and Department Head, Security Research, IBM

"Success in an AI-enabled landscape requires a dual foundation: deep technical domain expertise paired with rigorous critical thinking to deconstruct and solve fundamental business problems."

Omar Santos · Distinguished Engineer, Cisco

"Cultivating and maintaining deep human domain expertise is essential to preventing skill atrophy and mitigating long-term talent risks in an increasingly automated environment."

Hannah Calhoon · VP of AI, Indeed

Note: Panel views, see 5.2 section.

3.3 Lead AI Agents skills

Agent use is spreading quickly through cybersecurity work: AI skill integration in G7 cybersecurity postings has doubled year over year and crossed 25% of postings (section 2.2). For security professionals, agentic proficiency is more than a productivity skill; it is the defense capability: it is essential to deliver AI-enabled defense at the speed and scale required. Someone in the organization must define what a Model Context Protocol (MCP) server — the interface between agents and company tools and data — may expose, decide where human checkpoints sit in an autonomous workflow, and reason about the blast radius of an agent's actions when it hallucinates. That is the security professional's role.

This mandate calls for workflow re-engineering combined with judgment, across three layers.

01 AI literacy and tool fluency. Adapting prompts and configuration across tools to reach a reliable outcome.

02 Agent orchestration. Knowing what agents can do today and anticipating the next release cycle, scoping what they may touch, placing human checkpoints at irreversible actions, designing the workflow they run inside.

03 Agent evaluation. Defining what good looks like and measuring against it continuously, because agent behavior drifts as models and tools change.

AVOID

Deploying an autonomous SOC triage agent with broad API permissions and vendor-default guardrails. Nobody maps out what the agent can do versus what it should do. The first real test of its blast radius is the day people see the impact of a true positive the agent auto-closed and nobody caught.

GOOD PRACTICE

A security analyst applies orchestration principles: scoping tool exposure, setting human checkpoints at irreversible actions and defining explicit success criteria. They pilot an investigation agent on a bounded alert category, measure it against the human baseline, then argue for expanded autonomy with evidence.

The professional is being asked to operate at a higher order: orchestrator, governor, reasoner about consequences. Yet the traditional route to that judgment, the tier-one apprenticeship of triage and escalation, is exactly the layer AI is automating away. The layer AI automates first is the layer where judgment was learned.

FROM THE PANEL

Three practitioners describing, from different angles, the same emerging mandate: evaluate the tool and govern their boundaries and recognize their failures.

"The ability to critically evaluate tool efficacy and execute rapid experimentation is now a core organizational competency."

Henry Deng · Employer and Public Sector Partnerships Lead, Grow with Google

"By establishing prescriptive standards for MCP server exposure and agent operations, we ensure that our AI ecosystem remains governed, transparent, and inherently trustworthy."

Kyu Kwak · CISO, Pearson

"You need enough domain expertise to understand what an AI agent is trying to do and to recognize the consequences if its recommendations are based on a hallucination."

Ian Molloy · Principal Research Staff Member and Department Head, Security Research, IBM

Note: Panel views, see 5.2 section.

3.4 Interpersonal skills

When agents absorb the routine technical work, what remains for the human is disproportionately interpersonal: explaining an agent's decision to an auditor, negotiating guardrails with a product team, telling an executive why the fastest path is the wrong one.

Members on agent-supported teams will also work closer to clients and more across functions than their job descriptions anticipated. Three capabilities carry that load.

01 Ethical reasoning and empathy. Someone must break the tie when the agent's answer is fast, cheap and wrong. The hardest conversations in security are not with machines.

02 Stakeholder engagement and influence. The ability to influence and collaborate across teams is moving from useful to decisive.

03 Audience translation and precision of language. Instructing an agent is now an engineering act, and the same instructions need to be distilled and contextualized for an executive.

AVOID

The insider-threat agent could flag more if it also read personal messages and after-hours activity. Detection improves in testing. Nobody asks whether it should, because nobody wants to be the person slowing down a tool that works.

GOOD PRACTICE

Someone asks what the team would be willing to explain to the employees being monitored. The agent's scope stays on work systems, the exception path is written down, and the decision is recorded with its reasoning, not just its configuration.

FROM THE PANEL

Two practitioners describing, from different angles, the same emerging mandate: as AI takes on more of the execution, human value shifts to the judgment, precision, and adaptability required to direct it effectively.

"Soft skills, specifically judgment, adaptability, and conceptual flexibility, are becoming increasingly critical, as they allow professionals to effectively harness AI to solve complex challenges."

Kyu Kwak · CISO, Pearson

"Being able to tell the agent and inform the agent what you actually wanted to do requires being very, very precise with your language."

Ian Molloy · Principal Research Staff Member and Department Head, Security Research, IBM

Note: Panel views, see 5.2 section.

FIGURE 10

Human skills in cyber roles: year-over-year growth in demand

GROWTH	SKILL
+533%	YoY growth in demand for ethical reasoning
+125%	YoY growth in demand for stakeholder engagement

Source: Cornerstone and Indeed demand-side year-over-year growth data across the G7 countries.¹⁸

3.5 Business Context skills

AI-native companies demonstrate how AI can compress organizational layers, bringing technical workers closer to core business operations. Treated as a hypothesis for all organization, it has a clear implication for skills: workers become “bilingual” translators between the technical and the business, able to explain how the technology they run creates and sustains value for a customer, whether that customer is external or the team down the corridor.

Thriving in a flatter organization takes four competencies.

01 Strategic and systems thinking. Understanding the big picture and how interconnected parts interact.

02 Customer orientation and business acumen. If doing is no longer the bottleneck, everyone should pay more attention to outcomes, and outcomes mean understanding the value created for a customer, the often-cited importance of “taste” for developers.

03 Critical thinking and problem framing. Intelligent systems automate more work while performing unevenly, failing in unfamiliar ways and carrying evolving risks. Professionals have to evaluate outputs, spot the gaps and decide soundly when AI falls short.

¹⁸¹⁸ Source: Cornerstone and Indeed demand-side year-over-year growth data across the G7 countries. Growth rates are measured from a relatively small baseline and should be interpreted as a directional signal of emerging demand rather than the magnitude of market growth. Ethical reasoning is discussed in 3.4 and 3.6; stakeholder engagement in 3.4 and 3.5; this figure serves both.

04 Cost-per-outcome economics. Knowing what an AI-assisted result actually costs and being able to say so in the language of the people who approve budgets.

AVOID

The team reports how many alerts the agent handled and how many tokens it consumed. Nobody can say what a resolved incident now costs, or whether the cheap model running at twice the volume is actually cheaper once its mistakes are reworked.

GOOD PRACTICE

The team measures cost per resolved incident by model tier. Routine triage routes to a small model, the hard tail goes to a frontier one, and the routing rule is reviewed each quarter against outcomes. When finance asks what the agent is worth, there is a number.

FROM THE PANEL

Two practitioners describe the same movement, one at the level of the individual and one at the level of the career.

"We are actively fostering a product-management mindset, encouraging individuals to analyze their own roles and tasks to identify high-impact opportunities for AI-driven efficiency."

Henry Deng · Employer and Public Sector Partnerships Lead, Grow with Google

"We are observing a significant shift as professionals transition from traditional operational roles, such as SOC analysts, into product-leadership functions, a move that is substantially accelerating our organizational capabilities."

Kyu Kwak · CISO, Pearson

Note: Panel views, see 5.2 section.

3.6 Intrapersonal skills

AI is changing how professional value is defined. As agentic systems take on more of the task-based work, value shifts from completing tasks to driving outcomes.

That is a challenge to identity. When professional worth is tied to specific tasks, automation forces a redefinition. Four personal qualities, plus an overarching mindset, sustain effectiveness and wellbeing through it.

01 Self-awareness and purpose. The foundational anchor: decoupling professional identity from specific tasks and re-aligning it with broader, strategic goals.

02 Resilience and growth mindset. The mechanism for processing the discomfort of constant change, turning the stress of uncertainty into iterative improvement.

03 Curiosity and learning agility. The engine of continuous evolution, keeping knowledge current against fast-moving AI capability.

04 Adaptability and agency. The operational output. Combining the first three, individuals gain the capacity to navigate complex environments proactively rather than react to them.

In a word, an entrepreneurial mindset — now relevant to every professional, not only founders: viewing problems as opportunities, keeping a strong customer orientation, learning fast through experimentation, and making confident, data-driven decisions.

AVOID

An analyst tries the agent, it confidently gets something wrong, and they quietly stop using it¹⁹. Nobody notices, because nothing broke. Six months later they are the person on the team who cannot work the way everyone else now works.

GOOD PRACTICE

The same analyst notes what went wrong and tries again on a smaller task where the answer can be checked. The tool is still imperfect. They now know which parts of it to trust, which is the only useful thing to know about any tool.

FROM THE PANEL

Three practitioners describe the same disposition, from courage through mindset to discomfort.

"We are operating in uncharted territory; it requires both the courage to explore the unknown and the conviction to actively define and shape the future of work."

Marci Paino · Chief Learning Officer, Cisco

"The most successful professionals are those who not only dedicate time to mastering new tools but also maintain a growth mindset, embracing the reality of constant, rapid technological evolution."

Hannah Calhoon · VP of AI, Indeed

¹⁹ On human-AI collaboration in security operations generally, see Singh, R., Tariq, S., Jalalvand, F., Baruwat Chhetri, M., Nepal, S., Paris, C., and Lochner, M. (2025), *LLMs in the SOC: An Empirical Study of Human-AI Collaboration in Security Operations Centres*, arXiv:2508.18947. arxiv.org/abs/2508.18947.

"Embracing the feeling of being slightly out of your depth is far more productive than settling into complacency; it signals a growth mindset and a vital sense of urgency to keep learning."

Henry Deng · Employer and Public Sector Partnerships Lead, Grow with Google

Note: Panel views, see 5.2 section.

3.7 Recommendations

Technical Domain

FOR LEADERS

- **Make the delegation threshold explicit.** Define guidance by risk tier: which technical work an agent may perform with review, and which requires verification by a qualified expert, with named human accountability on every agent-assisted work product.
- **Protect the talent pipeline.** Require junior staff to complete foundational tasks manually through dedicated rotations and position them as primary auditors of AI-generated work. Take inspiration from airline pilots, who still rehearse engine failures in simulators at intervals set by regulation, because aviation learned through painful experience what happens when automation lets critical skills atrophy.

FOR THE WORKFORCE

- **Review the agent's output yourself and explain it to a colleague.** Skill erosion is not an automatic consequence of using AI; it is a consequence of using it without engaging with what comes back. Lira et al. (2025) found that people who practiced with AI exerted less effort and still performed better in later tests without it.²⁰ Reviewing and explaining is what keeps tool use in the category of practice rather than substitution.
- **Get familiar with the work beyond your own role.** Understand the principles other roles rely on and the ideas they treat as important; the NICE Workforce Framework for Cybersecurity is a map of them. Challenge the assumptions inside common practice, reason from first principles, and ask where AI could produce a better outcome.

²⁰Lira, B., Rogers, T., Goldstein, D. G., Ungar, L., and Duckworth, A. L. (2025), "Coach not crutch: Evidence that AI can improve writing skill despite reducing effort," arXiv:2502.02880. arxiv.org/abs/2502.02880

FOR EDUCATORS

- **Adopt a T-shaped model of technical proficiency:** broad knowledge across the domain, grounded in first principles, alongside deep expertise in the focus area of the specific role.
- **Design instruction and assessment across the full range of AI involvement,** from manual execution with no AI, through AI-augmented workflows, to directing AI agents end to end. The AI Assessment Scale offers a ready structure.²¹ At every point, test whether the learner can review, explain and validate the work product.

Lead AI Agents

FOR LEADERS

- **Own agent governance as a security function.** No agent goes into production without a named owner who can answer three questions: what it can access, which of its actions are irreversible, and who reviews it when it gets something wrong.
- **Be straight about what is changing.** If agent adoption means the triage role shrinks, say so, and name what people become instead: agent supervisors, tuners, exception investigators.

FOR THE WORKFORCE

- **Experiment in a properly isolated sandbox,** and write the agent's permission scope and escalation rules before you run it. The environment needs three properties: a disposable filesystem, controlled network egress with an allow-list, and scoped throwaway credentials that exist only for the exercise. Reverting the host is the easy part; constraining what the agent can reach is what matters.
- **Automate your technology trend scanning.** Use AI agents to keep pace with emerging agentic technology, building targeted prompts into a personal knowledge base that informs your expertise. Three approaches, in order of sophistication:
 - **Basic, ad hoc.** Once a week, run a search prompt tailored to your area in any frontier AI chatbot to surface developments you need to know about. A meta-prompt for building your own is in the appendix at 5.5.

²¹Perkins, M., Furze, L., Roe, J., and MacVaugh, J. (2024), "The Artificial Intelligence Assessment Scale (AIAS): A Framework for Ethical Integration of Generative AI in Educational Assessment," *Journal of University Teaching and Learning Practice*, 21(6). aiassessmentscale.com — see also the 2025 revision.

- **Routine, continuous.** Use native platform features to automate recurring scans. Background agents sweep technical forums on a schedule, filter noise, and deliver high-signal summaries as markdown files.
- **Advanced, knowledge base.** Move from finding to reasoning. Following the “LLM wiki” pattern, structure those markdown reports into a connected wiki you can browse in a note-taking application such as Obsidian and query through a chatbot, then periodically lint it for contradictions and stale claims.²²

FOR EDUCATORS

- **Integrate agentic AI into the cyber curriculum** through robust partnerships with industry, taught through active learning: experiential, project-based and problem-based work that builds curiosity, advanced problem-solving and metacognition, and leaves students comfortable sitting with open-ended questions and identifying their own knowledge gaps.
- **Simulate the internship and apprenticeship** with structured triage exercises, incident retrospectives and agent-supervision scenarios that build the judgment entry-level roles once taught on the job.

Interpersonal

FOR LEADERS

- **Make dissent cheap.** Give people a way to record disagreement with an agent’s output in under five minutes, and thank the ones who use it whether or not they turn out to be right. Disagreement that costs status only surfaces when someone is certain, which is too late.
- **Check what your scorecards are actually measuring.** If agents absorb the throughput, a metric built on volume is now measuring the agent rather than the person. Ask where the human contribution shows up in how work is assessed, and whether anything currently captures it.

FOR THE WORKFORCE

- **Write specs an agent cannot misread, then translate the other way.** Vague instructions produce vague agents. Practice both directions: a task specification precise enough that a colleague could

²²Andrej Karpathy, “llm-wiki.md” — a pattern in which immutable raw sources are kept alongside an LLM-generated, interlinked markdown wiki with ingest, query and lint operations. gist.github.com/karpathy. Obsidian is a markdown knowledge-base application. Tooling in this area changes quickly; product names are given as of July 2026.

execute it without a clarifying question, and agent output compressed into two sentences for an executive.

- **Volunteer for the conversations the agent cannot have:** Audit walkthroughs, tabletop exercises, briefing the business after an incident. That is where your value is moving.

FOR EDUCATORS

- **Build critical thinking and communication together.** Set scenarios where an agent is confidently wrong and the student must push back in writing, to someone senior. Grade the reasoning and the phrasing, not just the verdict.
- **Use the Feynman technique as standard practice.** Students explain the same technical work to three audiences: a fellow engineer, an executive, and a non-technical listener. The technique — explaining a concept in plain language until the gaps in your own understanding become visible — is defined in section 5.4.

Business Context

FOR LEADERS

- **Buy intelligence like a metered utility.** Through 2025 and the first half of 2026, per-unit model prices repriced roughly every quarter; treat the rate as something to re-check rather than a fixed assumption. Measure cost per outcome rather than per token: the cheap model that is wrong twice as often is the expensive one. Route routine work to small models and reserve frontier intelligence for the hard tail.
- **Give your team a number the business already understands:** cost per resolved incident, time to restore service, downtime avoided. Make the security team accountable for one of them alongside its technical metrics and report it where the business reports everything else.

FOR THE WORKFORCE

- **Audit your role like an outside analyst.** List your recurring tasks, mark which an agent could own within a year, and write the case for what you would do with the time freed.
- **Learn what your work costs.** Find out what one unit of your team's output costs to produce, including model spend and the rework when the agent gets it wrong. If you cannot say the number, you cannot argue for the budget.

FOR EDUCATORS

- **Add a business case to the capstone.** The technical solution is only part of the grade. The other part is the economics: what the incident cost, what prevention costs, and whether the student can defend the trade-off to a non-technical audience.
- **Put students in front of someone who pushes back on budget.** Brief a real internal stakeholder where you can and a role-played one where you cannot, then grade how well students defend the solution in business terms rather than technical ones.

Intrapersonal

FOR LEADERS

- **Make it safe to feel lost.** If admitting “I don’t know this tool yet” costs people status on your team, they will hide the gap and it will compound. Go first: learn in public and share what you learned, even imperfectly.
- **Watch for the quiet resisters.** The senior person avoiding the agent may be protecting an identity built over years. Help them build a new one, with agency and support.

FOR THE WORKFORCE

- **Treat feeling lost as evidence you are learning.** Feeling settled in this field means something is quietly expiring. Ask periodically: what am I avoiding because it makes me feel like a beginner?
- **Anchor identity to outcomes, not tasks.** If your worth is “I am great at analyzing packet captures,” an agent can take it. If it is “I am great at investigation that keeps this organization safe,” the agent is a new tool for the same mission.

FOR EDUCATORS

- **Assess the disposition, not only the answer.** Build coursework where the correct move is to say “I do not know this yet, and here is how I would find out.” Give credit for a documented, reasoned change of mind — the habit that keeps a practitioner effective when the tooling turns over.
- **Rebuild the apprenticeship the automation removed.** If tier-one work is where judgment used to be learned, replace it deliberately: reviewed casework, deliberately un-automated exercises, and reflection on what the student delegated and why. On this dimension educators have the most to do, not the least.

Learning

How education and learning must evolve to prepare the workforce for agentic cybersecurity work.

4.1 The increasing demand for experience in cybersecurity

The previous sections traced AI's impact on cybersecurity, how the workforce is adapting, and how to frame the new skills required. This section completes the skilling and education analysis. Upskilling those already in the field matters, but the harder question is how future workers are developed, and that requires industry and education alignment.

Entry-level ICT roles are among the most exposed to AI-driven task change. Research published in June 2026 by the World Economic Forum in collaboration with PwC finds that “globally, more than one in three young workers are employed in occupations with medium to high exposure to AI-driven task change.”²³

PwC's 2026 Global AI Jobs Barometer — a separate publication from the study above — finds that entry-level roles in the most AI-exposed occupations are seven times more likely to require traditionally senior, human-intensive skills, such as leadership, creativity and face-to-face interaction, than equivalent roles in the least exposed occupations.²⁴

The Consortium's ICT in Motion report (2025) identified a growing preference among ICT employers for more experienced candidates, which contributed to weaker demand for entry-level talent.

For this report, cybersecurity job demand was analyzed across G7 economies, comparing each half-year against the same half-year a year earlier and grouping postings by whether they carry a senior or a junior title. The split is clear. Senior-titled postings grew 46.1% in Apr – Sep 2025 and 65.0% in Oct 2025 – Mar 2026. Junior-titled postings moved from –0.6% to +5.9% across the same windows. As noted in section

²³World Economic Forum in collaboration with PwC, *Artificial Intelligence and the Future of Entry-Level Work: A Framework for Safeguarding and Reinventing Early Career Pathways*, 22 June 2026. [weforum.org — AI and the future of entry-level work](https://weforum.org/ai-and-the-future-of-entry-level-work)


²⁴PwC, 2026 Global AI Jobs Barometer. pwc.com/gx/en/services/ai/ai-jobs-barometer.html

2.1, total cybersecurity job demand grew 9.5% in the most recent six-month window, compared to 6.8% in the previous one.²⁵

The experience paradox. Employers are trending toward more senior titles and more senior skills, even in roles that are entry level.

FIGURE 11

Year-over-year growth in G7 cybersecurity postings, by seniority of title



Source: Indeed. Cybersecurity job postings, G7 economies, each half-year against the same half-year a year earlier. Titles were classified as senior or junior through job title analysis.

What security leaders say they cannot find

The expectations of security leaders help explain the paradox. In May 2026, Cisco surveyed 8,000 security leaders across 30 markets and 14 industries. Asked which competencies are hardest to find in entry-level cybersecurity candidates, three cluster tightly at the top: hands-on experience with AI agents (49%), technical cybersecurity depth (48%) and human-centric professional skills (45%). Four percentage points separate the top and bottom of that list, so the three should be read as a cluster rather than a ranking.

²⁵A note on method. Required years of experience is the more direct signal of seniority, but only a subset of postings state it. Titles are always present, so postings are clustered on them instead, accepting that a title is partly a marketing choice. Source: Indeed, cybersecurity job postings, G7, Oct 2025 – Mar 2026 against the same period a year earlier, and the preceding half-year on the same basis.

FIGURE 12

Competencies hardest to find in entry-level cybersecurity candidates

Competencies hardest to find in entry-level cybersecurity candidates



Source: Cisco survey of 8,000 security leaders across 30 markets and 14 industries, conducted in May 2026. Respondents could select up to three competencies, so percentages sum to more than 100%. See 5.2, Methodology. [Relentless Defense Report](#), September 2026.

Asked about the most important skills for the future, the same security leaders put AI-system security first: 53% rank securing complex, autonomous AI systems as the top entry-level competency for 2030, and 50% name a strong technical foundation paired with the ability to keep learning. These findings highlight an urgent need to evolve higher education and professional learning programs to meet these emerging requirements.

4.2 The learning transformation to respond to new needs

The technology shift and the new employer expectations have asked the learning world to respond. If employers expect AI-agent proficiency, technical depth and human skills on day one, the way those skills are built must change too. A wave of recent skilling programs is replacing lecture-based instruction with immersive environments where learners solve realistic, open-ended problems before their first job. The programs map closely onto the three gaps security leaders identified.

- [Cisco's Networking Academy Capture the Flag](#) is an interactive, gamified cybersecurity competition and learning platform. Instead of step-by-step lab guides, students solve real-world, story-driven challenges to find hidden "flags," building hands-on experience with ethical hacking and networking technologies.
- [The DC Tech Hub](#) addresses the technical-depth gap through a work-based route: a paid 12-month cybersecurity apprenticeship run by the University of the District of Columbia with Accenture, PeopleShores and Per Scholas, combining 15 weeks of intensive technical training with nine months of hands-on experience on Accenture delivery teams.

- [Google’s Vantage project](#) — a research protocol published by Google Research with NYU and OpenMic in April 2026 — is an AI simulation platform that assesses critical thinking, creativity and collaboration through scenario-based conversations with AI avatars.
- [IBM’s AI Builders Challenge with IBM Bob](#), launched in June 2026, is a global student competition built on its Bob development environment, now available to 20,000 post-secondary institutions, where participants focus on building working AI solutions rather than completing coursework.
- [Pearson’s AI Skills for the Future playbook](#) targets the human-skills gap, the hardest to teach from a textbook. Its D.E.E.P. framework — Diagnose the task-augmentation plan, Embed learning in the flow of work, Evaluate skill progress, and Prioritize learning as a strategic investment — is an operating model for closing skill gaps, aimed at the strategic human capabilities that AI productivity gains depend on but do not replace.

Note: The five programs above are illustrative examples supplied by Consortium members, not a survey of the field and not an evaluation.

FROM THE PANEL

One practitioner describes the need for human-guided, motivating learning experiences

“There’s an element of how we as instructors or those building learning experiences can work to motivate individuals and guide their learning. While AI can provide skills assessments and experiential feedback, the human in the loop, the human driving the process, remains essential to the learning experience.”

Henry Deng · [Employer and Public Sector Partnerships Lead, Grow with Google](#)

Note: Panel views, see 5.2 section

Organizations are rethinking how they train the people they already have

Static, one-size-fits-all curricula are giving way to closed-loop adaptive systems that treat learning the way engineers treat software: measure, adjust, release, repeat. The pressure is well documented and current.

63%	Of employers name skill gaps as the single biggest barrier to business transformation	WEF Future of Jobs Report 2025
85%	Of leaders call building an adaptable workforce critical	Deloitte 2026

7%

Say they are leading in helping their people grow continuously

Deloitte 2026

Sources: World Economic Forum, *The Future of Jobs Report 2025* ([weforum.org](https://www.weforum.org)); Deloitte, *2026 Global Human Capital Trends*, 4 March 2026 ([deloitte.com](https://www.deloitte.com)).

FROM THE PANEL

One practitioner describes the need for personalized, flow-of-work learning delivery

"To achieve true workforce readiness, we must move toward a holistic approach where learning is delivered in the flow of work and in a personalized way. As role requirements evolve, we must close the loop by validating competencies and processes to ensure they deliver measurable value to organizational performance."

Guna Jayaraman · Chief AI Officer, Cornerstone

Note: Panel views, see 5.2 section

4.3 Recommendations

FOR LEADERS

- **Adopt a continuous learning model** that builds personalized skill development, assessment and validation into daily workflows, so professional growth keeps pace with the technology it serves.
- **Invest in early-career talent and create more entry points** to avoid long-term skill gaps — internships, apprenticeships and projects, and connections with mentors and peers in the field.

FOR THE WORKFORCE

- **Build a standing habit of skilled practice.** Cyber ranges, learning platforms, competitions and real projects all produce evidence of competence an employer can see. Skill development can no longer be an event.
- **Spend time with AI agents now**, both as tools you direct and as systems you will eventually have to secure. And keep working on judgment, communication and collaboration; leaders put these among their top three priorities, and they are what carries over when roles change.

FOR EDUCATORS

- **Implement hands-on sandbox environments.** Learning providers need to recreate the experience layer. Sandbox environments should let students practice real vulnerability assessment and defense against shifting threats on top of solid security fundamentals.
- **Expand practical experience** by integrating internships, apprenticeships and industry projects to give students real-world exposure. Credentialing outcomes from live environments lets learners communicate their applied skill experience to employers.

Appendix

Who produced this report, how the evidence was gathered, the reasoning behind the skills framework, and the definitions the report depends on.

5.1 About the AI Workforce Consortium

The AI Workforce Consortium is a group of eleven leading companies — Accenture, Cisco, Cornerstone, Eightfold, Google, Honeywell, IBM, Indeed, Intel, Microsoft, Pearson and SAP — in partnership with global advisors including College Board, DIGITALEUROPE, the Business-Higher Education Forum (BHEF), the European Vocational Training Association and the National Applied AI Consortium.

The shared mission is to prepare today’s and tomorrow’s workforce with actionable insights and scalable frameworks to leverage the transformational opportunity of AI in information and communications technology (ICT) jobs across all industries.

Consortium members are pursuing large-scale initiatives with the stated goal of upskilling 130 million individuals globally.

Figure 13:

AI Upskilling commitments among AI Workforce Consortium members

CONSORTIUM MEMBER (Commitments by 2030)	LEARNERS
Accenture	10M
Cisco	25M
Cornerstone Foundation	3M
IBM	30M
Intel	30M

CONSORTIUM MEMBER (Commitments by 2030)	LEARNERS
<u>Microsoft*</u>	20M
<u>SAP</u>	12M

(*) Microsoft: Commitment period 2025-2027

(**) Google: To help ensure AI's benefits are widely shared, Google.org is launching a \$75 million dollar AI Opportunity Fund

5.2 Methodology

This report draws on two bodies of evidence, used for different purposes. Quantitative sources support all claims about workforce trends; qualitative sources inform how those trends are interpreted and how the skills framework was designed.

Quantitative

- **Job posting analytics (Cornerstone and Indeed).** Job posting data covering 24 months, April 2024 to March 2026, across G7 markets. This underpins the report's observations on shifting demand for cybersecurity skills, roles and qualifications. Coverage is limited to G7 economies, so findings should not be read as global.
- **Cisco survey of security leaders.** A survey of 8,000 security leaders across 30 markets and 14 industries, fieldwork conducted in May 2026. Respondents could select up to three competencies, so percentages sum to more than 100%. Data from this survey informs Figures 2 and 12, as well as several claims within the report. [Relentless Defense Report](#), September 2026.
- Job posting data captures employer demand as expressed in the market; the survey captures the stated priorities of security leaders.

Qualitative

- **Practitioner interviews,** conducted during the webinar "AI Agents and the Impact on Cybersecurity," recorded live on 21 May 2026. Seven practitioners from six Consortium member organizations took part, including leaders from Cisco alongside representatives of Cornerstone, Google, IBM, Indeed and Pearson.
- **Secondary research:** published reports from Consortium members and the wider industry, cited individually throughout.

- **Consortium contributions:** written input and review from member organizations.

Quotations

Quotations come from the 21 May 2026 discussion, selected to illustrate recurring themes and edited only for length and verbal filler; none has been altered in substance. This is a bounded panel, not a survey. The views are those of the individual speakers, are not attributable to their employers, and are not offered as representative of the profession. Claims about workforce trends rest on the quantitative sources and cited research, not on the panel.

Limitations

- **Scope.** Job posting analysis covers G7 markets only; the Cisco survey spans 30 markets. The two datasets do not cover identical populations.
- **Panel size.** Seven practitioners from eleven consortium members.
- **Pace of change.** Job posting data is a lagging indicator of capabilities advancing faster than hiring can register. The framework is a forward-looking hypothesis and should be revisited as evidence accumulates.

Authorship and review

AI tools were used to review and format the text. Expert members of the AI Workforce Consortium reviewed the draft and provided feedback and contributions. The analysis, findings, framework and recommendations remain those of the authors.

5.3 Developing the skills framework

The framework moves beyond narrow technical silos to address the shift toward AI-native organizations, and centers on the five dimensions used throughout this report: Technical Domain, Lead AI Agents, Business Context, Interpersonal and Intrapersonal. The thesis was built in three steps.

- 01 Evidence of AI capabilities.** A factual baseline of current AI capabilities and limitations, derived from documented evidence (section 5.3.1)
- 02 Hypotheses on the future of work.** An analysis of how work will transform, using AI-native organizations as leading indicators for traditional enterprises. This section is anchored in webinar transcripts, expert interviews, and secondary research (section 5.3.2).

03 The skills thesis. A definition of the mindsets and capabilities required to thrive, grounded in observed signals. (section 5.3.3 and section 3).

5.3.1 Evidence of AI capabilities

E1 · AI elevates and expands human capability

AI elevates by turning intent into finished artifacts and removing repetitive effort, so people spend more time on judgment and less on mechanics. It expands by letting people reach upstream and downstream in a workflow, doing credible work beyond their own core expertise. People do more, and do more kinds of things.

E2 · AI automates workflows under human supervision

AI can now execute not only discrete tasks but complete workflows on its own, within a frame set by people. Humans define the objectives, monitor performance, and keep the authority to intervene, correct or approve. Autonomy in execution; human authority over intent and outcome.

E3 · AI enables new operating models

The impact is not limited to individual productivity. AI changes how organizations are structured to deliver work, enabling leaner teams, flatter hierarchies and new combinations of human and machine effort. It reshapes the organization chart, not only the task list.

E4 · AI capabilities are accelerating but remain jagged

The pace itself is a defining condition. The capability frontier is expanding quickly, yet it is uneven: superhuman on some tasks, weak on others that look no harder, and stochastic by nature. Dell'Acqua et al. (2023) named and measured this jagged frontier in a field experiment with knowledge workers. Any framing of skills or operating models should therefore be treated as provisional and built to be revisited as the technology moves.

5.3.2 Five hypotheses on the future of work

If AI capabilities are the mechanisms, these five hypotheses are the outcomes. They describe how the nature of work, the structure of teams and the definition of value shift as agentic AI matures, moving the human worker from a doer of tasks to a director of outcomes.

H1 • The new shape of AI-human collaboration

Agentic AI is increasingly taking over operational tasks, positioning humans as the primary architects of intent and objectives. By defining workflows, establishing guardrails, and providing the necessary discernment and "taste," humans transition from active execution to being "in the loop" or "on the loop"—setting success criteria and acting as the final authority.

H2 • Everyone is an agent manager and orchestrator

The professional landscape is shifting toward an orchestration model. Every worker, not only technical leads, will manage a hierarchy of agents to increase productivity and solve operational challenges. Roles blur as individuals use agents to work both upstream in strategy and downstream in execution, beyond their traditional perimeters.

H3 • Small teams with agent leverage

Teams get smaller. By automating routine workflows, teams of two or three can execute at a scale that previously required entire departments. The focus shifts from headcount to high-leverage talent acting as internal startups.

H4 • Everyone operates closer to the business

Organizations are flattening toward the top. As transactional tasks are automated, every worker moves closer to the business context. Employees once focused on execution transition into strategic roles defined by critical judgment, adopting a product-management perspective regardless of their function.

H5 • Adaptability as the primary skill (VUCA)

In a world of accelerating change, the ability to learn, unlearn and adapt is the only certainty. As AI assumes responsibility for technical execution, human value shifts toward higher-order strategic contributions and human skills. Success is defined by comfort with volatility and with the speed of technological advancement.

H1 and H2 redefine the individual's day-to-day role from worker to manager. H3 and H4 redefine how those individuals cluster into lean, business-centric organizations. H5 is the overarching condition: on a jagged frontier, the sustainable advantage is the speed at which an organization and its people can adapt.

5.3.3 From capability to skill: the logical chain

The following table maps the capability evidence (section 5.3.1) to the hypotheses (section 5.3.2) and the skill dimensions outlined in section 3.

AI CAPABILITY — THE MECHANISM	FUTURE OF WORK HYPOTHESIS — THE OUTCOME	MAIN SKILL DIMENSION	THE LOGICAL CONNECTION — THE “WHY”
E2: Automate workflows — AI executes; humans supervise	H1: New collaboration shape — humans own intent; AI owns operative work	Technical Domain	To supervise automation, humans must retain first-principles expertise to exercise judgment and discernment where AI is stochastic.
E1: Elevate and expand — intent to artifact; blurring role perimeters	H2: Everyone is an orchestrator — individuals manage hierarchies of agents	Lead AI Agents	As work expands upstream and downstream, workers shift from doing to governing, prompting, and reasoning about agent actions.
E3: New operating models — leaner teams; flatter hierarchies	H3: Small teams with agent leverage — teams of two or three execute at scale	Business Context	Small teams require systems thinking and customer orientation to act as internal startups and drive direct business impact.
E3: New operating models — flatter structures; proximity to business	H4: Closer to the business — workers move from transactional to strategic roles	Interpersonal	In flatter organizations, value is created through influence, stakeholder engagement and empathy — work agents cannot perform.
E4: Accelerating and jagged frontier — fast, uneven, provisional	H5: Increasing VUCA — constant change; rapid skill half-life	Intrapersonal	A jagged frontier requires resilience, curiosity and learning agility to decouple identity from tasks and embrace constant re-skilling.

5.4 Key definitions

TERM	DEFINITION
Agentic AI	Systems that combine the intelligence of advanced AI models with access to digital tools, empowering the agents to take actions on behalf of users, under their control. Throughout this report, “AI agents” and “agentic AI” are used interchangeably.
AI forensics expert	An expert who investigates artificial intelligence systems after they fail, cause harm, or come under legal challenge. The evidence is unfamiliar: training data, model weights, prompt histories and inference logs rather than a disk image or a packet capture, read for bias, poisoning, tampering and adversarial manipulation. Regulation is starting to guarantee it exists. The EU AI Act obliges high-risk systems to log events automatically across their lifetime so incidents can be reconstructed later (Article 12).
AI skill integration	A metric that quantifies the prevalence of AI-related skills within job roles, based on requirements specified in job postings. Roles are classified into five bands by the proportion of postings requiring AI skills: immaterial ($\leq 10\%$), initial integration (10–25%), significant integration (25–50%), established integration (50–70%) and core ($> 70\%$).
Blast radius	The scope of systems, data and actions an agent can affect if it behaves incorrectly — the practical measure of how much a mistake can cost.
Codex Security	OpenAI’s agentic application security tool, designed to act as an automated security researcher.
Cyber range	A controlled virtual environment used for cybersecurity training, software testing and threat simulation — a digital flight simulator in which practitioners can rehearse against attacks safely.
Cybersecurity tabletop exercise	A collaborative, discussion-based session where teams simulate a cyber incident to evaluate their response plans, communication strategies and decision-making processes without disrupting actual operations.
ENISA ECSF	The European Cybersecurity Skills Framework, published by the EU Agency for Cybersecurity: twelve role profiles with associated skills, knowledge and competences.

TERM	DEFINITION
Feynman technique	A learning method in which a concept is explained in plain language, as if to a non-expert, until the gaps in one's own understanding become visible.
GPT-5.5-Cyber	It is a specialized cybersecurity-focused AI model designed to assist with advanced security research, threat analysis, vulnerability discovery, secure coding, and defensive security operations.
GRC	Governance, risk and compliance — the combined disciplines that set policy, assess risk against it and evidence conformance.
IAM	Identity and access management — the controls that establish who or what a principal is and what it is permitted to do.
ISO/IEC 42001	The international management-system standard for artificial intelligence, published in December 2023.
Jagged frontier	The observation that AI capability is uneven — strong on one task and unreliable on an adjacent one of apparently similar difficulty — so capability must be measured task by task rather than assumed.
Jailbreak incident	A security event in which an adversarial user successfully bypasses, evades or deactivates the safety filters, ethical guidelines and operational guardrails of an AI model.
LLM	Large language model — a model trained on very large text corpora that generates and reasons over language, and that underpins current agentic systems.
MCP server	A Model Context Protocol server: provides AI models with standardized, secure access to external data sources and tools.
MITRE ATLAS	A knowledge base of adversary tactics and techniques against AI-enabled systems, structured like MITRE ATT&CK.
MLOps	Machine learning operations — the practices for deploying, monitoring and maintaining machine learning models in production.

TERM	DEFINITION
Multi-factor authentication (MFA)	A security process requiring multiple forms of verification for access.
Mythos	Anthropic's frontier model, released to vetted partners as Claude Mythos Preview and described in a published system card; it is the core intelligence behind Project Glasswing. The release status and date should be stated consistently wherever the model is referenced.
NICE Framework	The NIST Workforce Framework for Cybersecurity (SP 800-181r1), which describes cybersecurity work in terms of task, knowledge and skill statements, and its competency areas, including AI Security.
NIST AI RMF	The NIST Artificial Intelligence Risk Management Framework (AI 100-1, 2023): a voluntary structure for governing, mapping, measuring and managing AI risk.
OpenAI Daybreak	An enterprise cybersecurity initiative launched by OpenAI on 12 May 2026, intended to tip the scaling advantage back toward defenders.
Project Glasswing	A global cybersecurity coalition launched by Anthropic.
SDLC	Software development lifecycle — the stages through which software is specified, built, tested, released and maintained.
Security Operations Center (SOC) analyst	An analyst who monitors, investigates and responds to threats to protect critical organizational networks and data.
SIEM	Security information and event management — the platform that aggregates logs and events from across an estate for detection, correlation and investigation.
SPL	Search Processing Language — the query language used in Splunk to search and analyze machine data.
VUCA	Volatility, uncertainty, complexity and ambiguity. It describes difficult, fast-changing and unpredictable environments.
Zero-trust architecture	A security framework requiring strict identity verification for every person and device, with no implicit trust based on network location.

5.5 Career Signal Radar Prompt

A portable metaprompt that builds a personalised weekly "how is AI changing my job" research prompt for any cybersecurity practitioner.

Works with the latest frontier models such as ChatGPT, Claude, Gemini and any other good reasoner model with searching tools. No vendor-specific syntax, no plugins, no code.

How to use the Career Signal Radar

1. Open any AI chat that can search the web.
2. Paste the Prompt.
3. Answer the questions to personalize your radar.
4. Check the results.

Recommendation - High-Stakes Validation (Council of Models)

For significant decisions—such as a role change, a certification spend, or a team restructure—run your weekly prompt through two different AI models simultaneously. Where the models agree, the finding is likely robust. Where they diverge, the discrepancy highlights exactly where you should invest twenty minutes of your own focused reading to verify the truth.

Meta prompt

The metaprompt text is not reproduced in this bundle.

Check online on the AI Workforce Consortium website: <https://www.cisco.com/site/m/ai-workforce-consortium/index.html>

5.6 Resources

Incidents and capability evidence

— [Anthropic — Disrupting the first reported AI-orchestrated cyber espionage campaign \(November 2025\)](#)

- [IBM X-Force Threat Intelligence Index 2026 \(25 February 2026\)](#)
- [Sysdig Threat Research Team — JadePuffer: agentic ransomware for automated database extortion \(July 2026\)](#)
- [OpenAI — statement on the Hugging Face model-evaluation security incident \(21 July 2026\)](#)
- [Epoch AI — open-weight models lag the frontier by about four months \(May 2026\)](#)
- [Anthropic — *System Card: Claude Mythos Preview* \(2026\).](#)
- [Moonshot AI — Kimi K3 announcement \(16 July 2026\) and open weights \(27 July 2026\).](#)

Policy and government

- [Five Eyes cyber security agencies — joint statement on the AI shift in cyber threat \(22 June 2026\)](#)
- [US Executive Order — Promoting Advanced Artificial Intelligence Innovation and Security \(Federal Register, 5 June 2026\)](#)
- [European Commission — Action Plan on Cybersecurity and Artificial Intelligence, COM\(2026\) 577 \(7 July 2026\)](#)

Standards and frameworks

- [NIST AI Risk Management Framework \(AI 100-1\)](#)
- [NIST Cybersecurity Framework 2.0](#)
- [NIST SP 800-181r1 — Workforce Framework for Cybersecurity \(NICE\) and the AI Security competency area](#)
- [NIST SP 800-218A — Secure Software Development Practices for Generative AI](#)
- [ISO/IEC 42001 — AI management system](#)
- [OWASP Top 10 for Agentic Applications \(2026\)](#)
- [MITRE ATLAS](#)
- [Cloud Security Alliance — AI Controls Matrix](#)
- [ENISA — European Cybersecurity Skills Framework](#)
- [SFIA 9 · European e-Competence Framework \(e-CF\) · ESCO](#)
- [UNESCO — AI Competency Framework for Teachers \(2024\)](#)

Research and surveys

- [WEF — Global Cybersecurity Outlook 2026 \(12 January 2026\)](#)

- [WEF — The Future of Jobs Report 2025](#)
- [WEF with PwC — Artificial Intelligence and the Future of Entry-Level Work \(22 June 2026\)](#)
- [PwC — 2026 Global AI Jobs Barometer](#)
- [Deloitte — 2026 Global Human Capital Trends \(4 March 2026\)](#)
- [IBM — Cost of a Data Breach Report 2025](#)
- Dell’Acqua et al. (2023) — [Navigating the Jagged Technological Frontier](#)
- Singh et al. (2025) — [LLMs in the SOC \(arXiv:2508.18947\)](#)
- Lira et al. (2025) — [Coach not crutch \(arXiv:2502.02880\)](#)
- Perkins, Furze, Roe and MacVaugh (2024) — [The AI Assessment Scale](#)
- [AI Workforce Consortium](#) — [ICT in Motion \(2025\)](#).
- [TRUST Framework Google, a governance blueprint for AI agents \(2026\)](#)
- Cisco survey of 8,000 security leaders across 30 markets and 14 industries, fieldwork conducted in May 2026. [Relentless Defense Report](#), September 2026.

Programs and tools

- [Cisco — Foundry Security Spec \(12 May 2026\)](#)
- [Cisco Foundation AI — Antares \(21 July 2026\)](#)
- [Cisco Networking Academy — Capture the Flag](#)
- [Google DeepMind — CodeMender](#)
- [DARPA — AI Cyber Challenge \(AIxCC\) and its results](#)
- [Indeed Hiring Lab — data portal](#)
- [Karpathy — llm-wiki pattern](#)

Webinars

- Watch the [“Agentic AI and the Impact on Cybersecurity”](#) Webinar
- Register [here](#) for our upcoming AI Workforce Consortium webinar on Governance and Responsible AI.

5.7 Terms of use and disclaimer

The insights presented in this report are provided solely for informational purposes and are presented “as-is.” While every effort has been made to ensure the accuracy and relevance of the information, the AI Workforce Consortium (the “Consortium”) does not assume responsibility for any decisions made based on the data included herein. It is recommended that organizations and individuals conduct their own research and due diligence to inform their decision-making processes. The Consortium expressly disclaims any responsibility and shall not be liable for any damages, losses, injuries, or liabilities arising from reliance on the information contained in this report. Users bear the sole responsibility for evaluating the accuracy and usefulness of the information obtained.

Copyright © August 2026, AI Workforce Consortium. All rights reserved.

AI Workforce Consortium Spotlight Report, July 2026. A companion to the webinar “AI Agents and the Impact on Cybersecurity,” recorded on 21 May 2026.