AUGUST 2026 EDITION

SkillWeed

THE OMNIBUS IS LAW

THE AGENTS ARE THE RISK

GOVERNING THE MONTH THE AI ACT ACTUALLY LANDED

TABLE OF CONTENTS

EDITOR'S NOTE

Last month we told you the Digital Omnibus was heading toward formal adoption. This month it happened — and then some. Between the last edition and this one, the Council of the EU gave its final green light, the Parliament's endorsement became irreversible, and the finished act was signed on 8 July. It is now sitting in the queue for publication in the Official Journal, days ahead of the original 2 August deadline it was designed to soften.

That is the headline regulatory story. The headline operational story is different, and arguably more urgent: the agents. Multiple independent surveys published in the past few weeks all converge on the same uncomfortable number — somewhere close to half of enterprises have already had a security incident tied directly to an AI agent, not a hypothetical one. Sophos, DigiCert, the Cloud Security Alliance, and OWASP's new Top 10 for Agentic Applications all landed within weeks of each other, and they are not describing the same risk from four different angles. They are describing four separate confirmations of one thing: identity and access management, not model intelligence, is where agentic AI is currently failing organizations.

This edition follows that thread all the way through. We open with what the Omnibus actually changes on the ground. We spend real time inside the agentic AI data, because the numbers are more specific and more alarming than most people realize. We roundup a brutal month of breaches — including one at a Big Four accounting firm and a WordPress flaw large enough to matter to nearly every organization reading this. We look hard at the vendor and OAuth problem the July edition first raised, because it got worse, not better. And because readers told us the beginner playbook was the most-used page in the last issue, we've written a second one — this time for agentic AI policy specifically, since a chatbot policy and an agent policy are not the same document.

We also owe you something on the community front. Readers pushed back on how thin that section ran last time, and they were right to. This edition's community chapter is longer, more specific, and includes real detail on what members are actually building, not just a bulleted list of programs.

Let's get into it.

Akingbade Akinfenwa

CHAPTER 1:
THE DIGITAL OMNIBUS IS NOW LAW — WHAT ACTUALLY CHANGES THE TIMELINE, CONFIRMED

Last edition, the Digital Omnibus on AI was still a provisional agreement waiting on two more institutional steps. Both have now happened. The European Parliament's 16 June endorsement was followed by the Council of the EU's final approval on 29 June, and the finished legislative act was signed on 8 July. It is now awaiting publication in the Official Journal of the European Union, which triggers entry into force three days later — expected to land just ahead of the original 2 August compliance cliff the package was written to soften.

That sequencing matters for one practical reason: until publication actually happens, the original AI Act text and its original August 2026 deadlines remain the technically binding law. Organizations that paused their high-risk compliance work on the strength of "it's basically decided" were operating on a reasonable bet, but a bet all the same. As of this writing, that bet has paid off — the amendments are through every legislative step that was left. Watch for the Official Journal publication date directly if your compliance calendar depends on the exact entry-into-force day.

WHAT'S DEFERRED, RESTATED WITH THE FINAL TEXT

- High-risk Annex III systems (hiring tools, credit-scoring, law enforcement use cases) — deferred from August 2026 to 2 December 2027.
- High-risk Annex I systems embedded in already-regulated products (medical devices, machinery, lifts) — deferred to August 2028.
- Watermarking obligations for AI-generated content under Article 50(2) — a shorter deferral, to 2 December 2026, for systems already on the market before 2 August 2026.
- National AI regulatory sandboxes — member states now have until August 2027 to stand these up.

WHAT'S NEW IN THE FINAL TEXT THAT WASN'T FULLY SETTLED LAST MONTH

- A loosened AI literacy obligation. The amended Article 4 now requires providers and deployers to take measures supporting the development of AI literacy among staff, rather than affirmatively ensuring a sufficient level of it outright — a subtle but real shift from an outcome obligation to an effort obligation.
- Expanded AI Office oversight. The final package hands the AI Office meaningfully broader supervisory powers, including over platforms already regulated under the Digital Services Act — a detail that got less attention than the deadline extensions but may matter more for large platform operators.

- A firm December 2026 date for the new prohibition on AI systems that generate child sexual abuse material or non-consensual intimate imagery ("nudifiers"). Fines for this specific prohibition can reach €15 million or 7% of global annual turnover, whichever is higher — among the steepest penalty exposures in the entire Act.
- The Commission has its own new deadlines to hit: guidance on post-market monitoring plans by September 2027, and guidance to help Annex I providers reconcile the AI Act with sectoral rules by August 2028.

WHAT THIS MEANS FOR YOU THIS MONTH

If your organization built a compliance roadmap around the original August 2026 date, this is the moment to formally re-baseline it — not quietly let it slip. Document the change, note the new dates, and use the recovered runway deliberately: real risk assessments and control testing, not a longer runway to eventually rush the same thin version of the work. Transparency obligations under most of Article 50 still land on schedule in August regardless of the Omnibus, so don't let the high-risk deferral create a false sense that everything moved.

THE BOTTOM LINE

The Omnibus stopped being a forecast this month and became a fact. That doesn't reduce the workload — it resets the clock on it. Organizations that treat 2 December 2027 the way they were treating 2 August 2026 a year out will be fine. Organizations that treat this as the regulation going away will be the ones scrambling in eighteen months.

CHAPTER 2:
INSIDE 2026'S AGENTIC AI RECKONING THE NUMBERS THAT MATTER

Last month's issue argued that agentic AI changes the risk calculus for GRC teams. This month, four separate research efforts — from a security vendor survey, an industry summit, a cloud security association, and the OWASP Top 10 project — converged to make that argument with hard data instead of theory.

A DigiCert-commissioned survey of just over 1,000 IT and security leaders across the US, UK, and Australia found that 78% had experienced some kind of AI-related security issue in the prior six months. Of that group, roughly 28% caught a vulnerability before it became an incident — but the remaining 50% reported an actual breach or operational disruption tied to an unauthorized or misconfigured AI agent. Only about half of surveyed organizations had a formal AI governance program in place at all. Science and technology firms reported the highest incident rates, with banking, financial services, telecom, and retail close behind.

Separately, research compiled by security firm Shattered found that among enterprises that have actually deployed AI agents, 88% reported at least one agent-related security incident. The same body of research traced most failures to two root causes: agents granted more access than their task required, and agents acting on data they should never have been able to reach in the first place. Over-permissioned credentials alone accounted for 61% of incidents; prompt injection, the most basic form of agentic attack, was implicated in roughly a third.

A Cloud Security Alliance research note went further, finding that 98% of assessed production AI agents combine three risk factors simultaneously: access to private data, exposure to untrusted content, and the ability to take outbound action. The CSA's blunt framing was that an agent's capability and its defensibility are inversely correlated — the more useful an agent is, the harder it typically is to fully secure.

WHAT'S NEW SINCE LAST MONTH'S CHAPTER

- OWASP formalized its Top 10 for Agentic Applications, the first peer-reviewed framework built specifically around autonomous AI rather than chat-style models. It reframes the entire threat model around identity and tool use instead of the traditional network perimeter — a genuinely different lens for GRC teams used to thinking in terms of firewalls and endpoints.
- Published July 22, the Sophos AI Security 2026 Report concluded that AI agent identities are now the fastest-growing exposed attack surface inside enterprises, driven largely by employees adopting coding agents and assistants that are quietly granted privileged access to core systems.

- Security researchers have already built and demonstrated proof-of-concept adaptive AI "worms" — malicious code that uses an agent's own reasoning loop to discover vulnerabilities and self-propagate, rather than relying on a fixed exploit chain. This remains a research demonstration, not a confirmed in-the-wild incident, but it signals where the threat model is heading.
- The dominant defensive theme researchers converged on this month is "Agent Zero Trust" — frameworks from major AI labs now explicitly argue that autonomous agents should be treated as potential insider threats by default, with cryptographically scoped identities and continuous runtime monitoring, rather than as trusted extensions of the employee who deployed them.

WHY THIS ISN'T THE SAME STORY AS LAST MONTH

July's issue used the Vercel/Context.ai incident as its central case study — an employee's overly broad OAuth grant to a productivity tool that, once compromised, gave attackers a path into a company that wasn't even that tool's customer. That pattern hasn't gone away; if anything, this month's data confirms it's the dominant failure mode industry-wide, not an isolated event. What's changed is the scale of confirmation: this is no longer one incident illustrating a risk, it's a majority of surveyed enterprises reporting the same category of failure.

GOVERNING AGENTIC SYSTEMS: WHAT TO ADD THIS MONTH

- Treat every agent as a distinct identity with its own credential lifecycle — not a feature bolted onto a human employee's account.
- Assume prompt injection will be attempted against any agent that reads untrusted content (emails, web pages, uploaded files, code comments). Build detection for it rather than hoping it doesn't happen.
- Map which of your production agents combine all three CSA risk factors — private data access, untrusted content exposure, and outbound action — since that combination is now the norm, not the exception, and each one deserves a dedicated review.
- Revisit your incident response runbook specifically for cascading multi-agent failures, where one compromised agent's bad output becomes another agent's trusted input.

THE BOTTOM LINE

The theoretical case for agentic AI governance is closed. The industry now has the incident data to prove the risk isn't hypothetical, and multiple independent research bodies reached the same conclusion within weeks of each other. The professionals who move fastest from "we should probably look into this" to "here's our agent identity inventory and our zero-trust rollout plan" will be the ones setting the standard other GRC teams get measured against.

CHAPTER 3:
BREACH ROUNDUP — WHAT BROKE IN JULY
THE MONTH IN ONE LINE

If June's breach roundup was about AI tools becoming an attack surface, July was the month that lesson showed up everywhere from automotive e-commerce to Big Four accounting.

NOTABLE DISCLOSURES

REVOLUTIONPARTS (DISCLOSED JULY 2026).

An e-commerce platform serving automotive dealerships disclosed a breach exposing more than 5.1 million records, underscoring that mid-market SaaS vendors serving entire industries remain high-value targets precisely because a single breach cascades across every dealership relying on the platform.

ERNST & YOUNG (DISCLOSED JULY 2026).

One of the Big Four professional services firms disclosed a breach affecting client tax and investment data — a reminder that even organizations built around audit and assurance are not exempt from the same identity and access failures they advise clients on.

NAIC / SHINYHUNTERS (ONGOING).

The National Association of Insurance Commissioners confirmed that the ShinyHunters extortion group's breach of its PeopleSoft environment exposed only publicly available data, outdated logs, and configuration files — a case worth studying specifically because the organization's own disclosure pushed back on the attacker's inflated claims, a useful template for how to communicate proportionately during an active extortion attempt.

KDDI (DISCLOSED LATE JUNE, STILL RESOLVING IN JULY).

The Japanese telecom operator disclosed a breach of an email system shared by five other Japanese ISPs, exposing email addresses and passwords for as many as 14.2 million customers. Reporting noted that only some passwords appeared to have been hashed, with no clear explanation yet of why any were stored in plaintext at all — a basic control failure at enterprise scale.

COCA-COLA SUBSIDIARY / ANUBIS RANSOMWARE GROUP.

The Anubis group claimed to have stolen roughly 1 TB of confidential data from a Coca-Cola subsidiary, part of a wider pattern of ransomware groups targeting subsidiaries of large multinational brands rather than parent companies directly.

SYNOPSYS AND BOSCH / D1R GROUP.

The D1R cybercrime group claimed to have exfiltrated data from both engineering firms, threatening to leak it absent payment — another entry in the growing list of extortion-first, encryption-optional ransomware operations.

NIKE / WORLDLEAKS.

Nike confirmed it is investigating a potential incident after the WorldLeaks extortion gang claimed to have stolen and posted samples from 1.4 TB of internal files tied to design and product development.

PAYPAL WORKING CAPITAL (ONGOING DISCLOSURE).

Reporting confirmed the earlier-disclosed breach exposed names, emails, phone numbers, business addresses, Social Security numbers, and dates of birth spanning July 2025 through December 2025, with some reports of unauthorized transactions, refunds, and password resets tied to affected accounts.

INFRASTRUCTURE AND PATCH NEWS THAT MATTERS TO GRC TEAMS, NOT JUST ENGINEERING

Microsoft's July Patch Tuesday addressed roughly 570 vulnerabilities, including three actively exploited zero-days — one of the largest single patch cycles on record, and a useful data point for anyone building a patch-management control narrative for an upcoming audit.

A remote-code-execution flaw dubbed wp2shell was disclosed affecting the WordPress core, putting an estimated 500 million-plus WordPress sites at risk of unauthenticated takeover. Given how much of the web still runs on WordPress, this belongs in any third-party and website-risk inventory, even for organizations that don't consider themselves "WordPress shops" — marketing sites, blogs, and vendor microsites are common blind spots.

Oracle's July Critical Patch Update drew specific attention because researchers and Oracle itself indicated many of the fixed vulnerabilities were likely discovered with the help of

AI-assisted vulnerability research — a small but telling signal that AI is reshaping the defender side of the equation as much as the attacker side.

WHY THIS BELONGS IN A GRC PUBLICATION

None of these incidents required a novel technique. Weak credential hygiene, shared infrastructure between vendors, overly broad platform access, and patch backlogs did the damage — the same fundamentals GRC programs exist to catch, executed against organizations of every size and sector in a single month.

THE BOTTOM LINE

A single bad month like this is not an argument for panic. It's an argument for treating the risk register as a living document rather than a quarterly ritual — because the assumptions it was built on can go stale inside thirty days.

CHAPTER 4:
VENDOR RISK PULSE — THE MCP AND OAUTH PROBLEM THE STORY THAT WON'T GO AWAY

July's issue used the Vercel/Context.ai incident to make a simple point: attackers increasingly log in rather than break in, riding an employee's own OAuth grant into systems the breached vendor never had permission to touch. This month, the same pattern showed up at industry scale, with a new and more specific culprit — the Model Context Protocol (MCP), the connective layer increasingly used to let AI agents call tools and reach external systems.

WHAT'S NEW THIS MONTH

Security researchers identified 492 MCP servers exposed directly to the internet with no authentication at all — meaning any of the tools, data sources, or internal systems those servers connect an AI agent to were, in effect, reachable by anyone who found them.

Separately, malicious-skill scanning uncovered more than 1,100 malicious skills distributed across ClawHub, a marketplace for a popular open-source agent framework, illustrating that the same supply-chain risks long familiar from npm and PyPI packages have arrived in the AI agent ecosystem largely unchanged in shape.

A related supply-chain compromise affecting a widely used AI plugin ecosystem reportedly harvested compromised agent credentials from dozens of enterprise deployments, giving attackers standing access to customer data, financial records, and proprietary code for months before discovery — the agentic-era equivalent of the OAuth-token pattern from last month's Vercel case, but with a longer dwell time.

Vendor concentration risk also showed up in a more conventional form: a supply-chain breach affecting Salesforce-adjacent data pipelines exposed customer names, business emails, phone numbers, job titles, sales notes, CRM records, and pricing information across multiple downstream companies, including well-known security and productivity vendors — a reminder that even security-focused vendors are not immune to being the "someone the target chose to trust" in another organization's breach story.

WHY THIS IS A GOVERNANCE PROBLEM, NOT JUST AN ENGINEERING ONE

Most Google Workspace and Microsoft 365 tenants still allow any employee to grant a third-party application access to the corporate account by default. MCP servers compound this because they often sit between an AI agent and multiple internal systems at once — meaning a single unauthenticated MCP server can expose far more than a single OAuth grant would. The Pentagon reportedly designated a major AI vendor a "supply chain risk" this year, a first-of-its-kind classification that signals regulators and large customers are starting to treat AI infrastructure providers with the same scrutiny historically reserved for hardware and cloud vendors.

WHAT TO DO ABOUT IT — UPDATED FOR MCP

- Extend your vendor risk questionnaire to explicitly ask whether a vendor's product uses MCP or a similar agent-tool-calling protocol, and if so, how server authentication is enforced.
- Inventory every MCP server or agent-tool integration your organization runs, whether centrally deployed or adopted by individual teams, the same way you would inventory SaaS applications with OAuth access.
- Move OAuth and MCP consent to admin review by default — the single highest-leverage control change available to most organizations this year, and still not the default configuration at most companies.
- Ask AI and agent-platform vendors directly what scopes their integrations request and why, and decline anything broader than the function requires.

THE BOTTOM LINE

Vendor risk didn't get a new category this year so much as a new door. The rigor doesn't need reinventing — it needs to be pointed, deliberately, at MCP servers and agent-tool integrations with the same seriousness your program already applies to any vendor touching customer data.

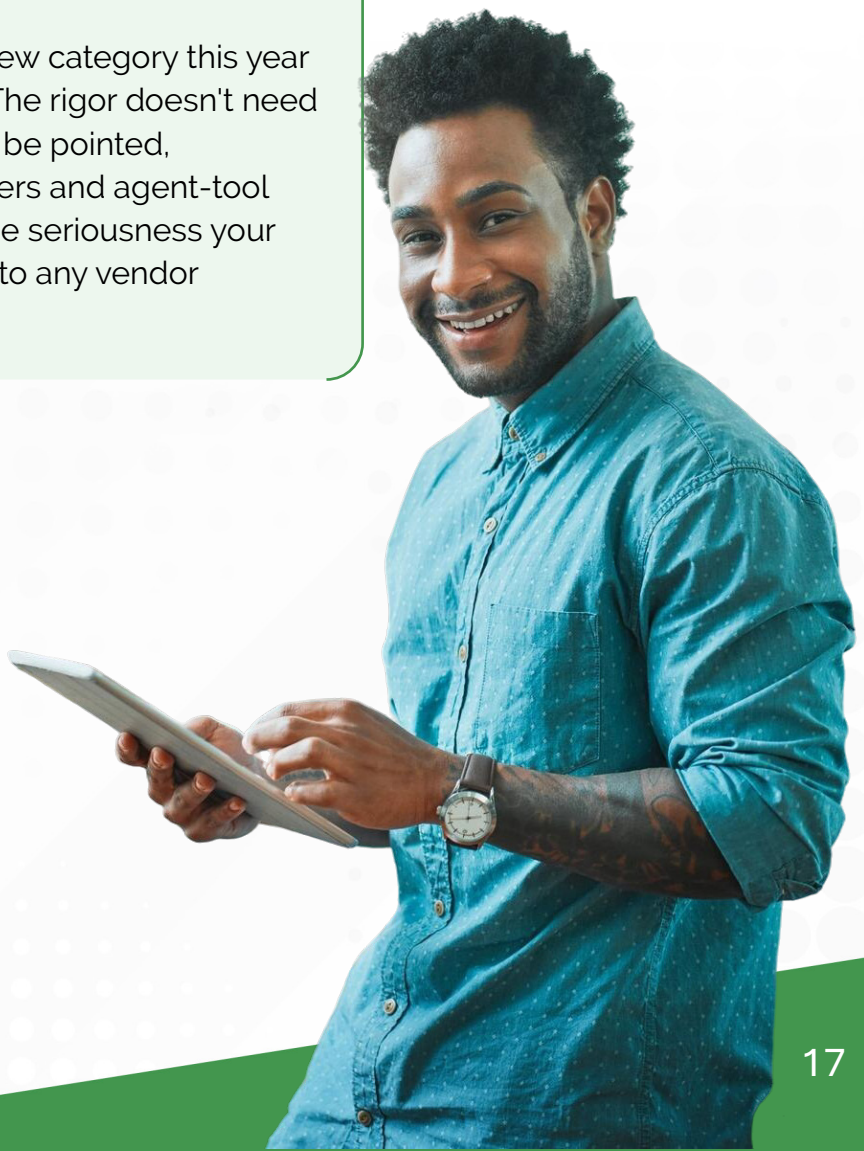

CHAPTER 5:
BEGINNER'S PLAYBOOK — WRITING YOUR FIRST AGENTIC AI POLICY WHY THIS ISN'T THE SAME PLAYBOOK AS LAST MONTH

July's beginner playbook covered writing a first AI usage policy — a document mostly concerned with what tools employees may use and what data they may enter into them. That playbook still matters, but it assumes a human asks a question and a human decides what to do with the answer. An agentic AI policy has to cover something that playbook never touched: a system that acts on its own, using credentials, without a human approving each step.

STEP 1 — INVENTORY AGENTS SEPARATELY FROM TOOLS

Before writing anything, distinguish between AI tools your organization uses (a chatbot, a writing assistant) and AI agents that can take action (an agent that reads email and replies, an agent with access to your CRM or codebase, a coding agent that can commit changes). Most organizations already did the Step 1 inventory from last month's playbook for tools. Agents need their own separate inventory, because the risk profile is categorically different.

STEP 2 — GIVE EVERY AGENT ITS OWN IDENTITY

An agent should never simply inherit the permissions of the employee who set it up. Give each agent a distinct, scoped identity with the minimum access its specific job requires. If an agent only needs to read a calendar, it should not also be able to send email on the organization's behalf.

STEP 3 — DEFINE WHICH ACTIONS REQUIRE A HUMAN CHECKPOINT

Pick the handful of actions that would cause real damage if an agent got them wrong — sending money, deleting data, messaging customers externally, pushing code to production — and require a named person to approve each one before it executes, regardless of how autonomous the rest of the workflow is.

STEP 4 — WRITE THE ONE-PAGE POLICY

Your first agentic AI policy should answer five questions in plain language: which agents exist and what can they touch, what identity and access scope each one has, which actions require human approval, who can authorize a new agent or a scope change, and what the kill-switch procedure is if an agent misbehaves. That's the whole document to start.

STEP 5 — BUILD THE KILL SWITCH BEFORE YOU NEED IT

Know, in writing, exactly how to immediately revoke a specific agent's access — not "disable AI tools generally," but cut off one agent's credentials without disrupting every other system it might share infrastructure with. Test this procedure once, on a low-stakes agent, before you ever need it on a high-stakes one.

STEP 6 — REVIEW MONTHLY, NOT QUARTERLY

Agentic AI tooling is changing faster than the quarterly review cycle recommended for general AI policy last month. For agent-specific policies, put a monthly check on the calendar for at least the next two quarters.

A STARTER TEMPLATE TO ADAPT

"[Organization] operates the following AI agents for [approved workflows]. Each agent is provisioned its own scoped identity with the minimum access required for its task, never a shared or inherited credential. [High-impact action list] require review and approval by [named role] before execution. Any new agent or expansion of an existing agent's access must be requested through [process] and reviewed by [approving role]. In the event an agent behaves unexpectedly, [named role] can revoke its access immediately via [kill-switch process]. This policy will be reviewed monthly."

You will not get this exactly right on the first draft, and neither did any of the organizations in this month's incident data. Write it, test the kill switch, and improve it before an actual incident forces the improvement.

CHAPTER 6:
TOOL OF THE MONTH — VANTA'S AGENT FOR RISK WHY WE'RE FEATURING VANTA THIS MONTH

Last edition covered Drata. This month, Vanta made the more directly relevant move for the theme of this issue: a purpose-built agent for managing the exact category of risk this magazine has spent two editions unpacking.

WHAT IS VANTA'S AGENT FOR RISK

Vanta describes itself as an Agentic Trust Platform, and its newest release — the Agent for Risk is built to bring internal risk and third-party/vendor risk into a single, continuously updated view rather than two separate workflows. It sits on top of what Vanta calls its Trust Graph, a unified data layer spanning more than 400 integrations and over 1,400 continuous automated tests, creating what the company describes as a living, connected map of an organization's controls, vendor relationships, assets, and compliance obligations.

WHAT STANDS OUT

- Unified internal and third-party risk. Rather than treating vendor risk and internal control risk as separate programs with separate tooling, the Agent for Risk is built to give a single continuously updated posture across both.
- An AI Risk Library. A purpose-built knowledge base specifically addressing AI-related risk scenarios, reflecting the same shift this magazine has been tracking toward agentic and AI-specific risk categories needing dedicated tooling rather than a bolt-on to existing GRC platforms.
- Context-aware agent guidance. Building on Vanta's earlier AI Agent release, the platform proactively flags inconsistencies, suggests fixes, and can execute some remediation steps directly, rather than only surfacing a dashboard for a human to interpret.
- Industry recognition. Vanta was named a Leader in its first-ever appearance in a major analyst evaluation of GRC platforms this year, a marker of how quickly automation-first, AI-native GRC tooling has moved from niche to mainstream in analyst frameworks.

WHAT THIS MEANS FOR GRC PROFESSIONALS

The tools featured in this magazine over the past two editions all point toward the same shift: platforms are no longer just automating evidence collection, they're increasingly automating the risk-scoring and remediation-suggestion work that used to require a human analyst's judgment call. That doesn't remove the analyst from the loop — it changes what the analyst spends their time on, from gathering evidence to reviewing what an agent already gathered and deciding what it got wrong.

THE CAVEAT

A unified risk view is only as good as the integrations and scoring logic underneath it. Teams adopting tools like this still need to understand what "good" looks like independent of the dashboard — the same caveat that applied to Drata last month applies here, arguably more so, given how much more the platform is now doing autonomously.

CHAPTER 7:
CYBER GRC ROLE SPOTLIGHT — THE 2026 JOB MARKET, BY THE NUMBERS WHY WE'RE REVISITING THIS

Last edition's role spotlight covered the basics of the Cyber GRC Analyst role. This month, several labor-market reports landed with hard numbers on exactly how that role is changing — worth a dedicated look rather than folding into a general update.

THE HEADLINE STATS

An analysis of 100 real cybersecurity and AI job listings in 2026 found that half now require some form of AI governance competency. Separately, machine learning skills now appear in roughly 57% of cybersecurity-plus-AI job listings, with AI governance knowledge specifically showing up in about half. ISC2 identifies AI/ML as the single most in-demand cybersecurity skill for 2026, cited by 41% of security teams as their top requirement, ahead of cloud security in second place.

The GRC market specifically is projected to grow at roughly 13% annually through 2027, driven by expanding global privacy regulation and tightening cyber insurance requirements — a growth rate that has held up even through a year that saw layoffs at large enterprises undergoing restructuring.

WHAT'S ACTUALLY CHANGING ABOUT THE ROLE

- Quantitative risk thinking is displacing gut-feeling risk ratings. The 2026 expectation is increasingly the ability to translate "we are at high risk" into a specific dollar-value estimate of potential loss, using tools built for that purpose.
- ISO 42001 Lead Auditor has emerged as one of the most in-demand certifications specifically for anyone moving into AI governance work, alongside the CRISC and CISA certifications already covered last edition.
- Board-level reporting expectations are rising. Coverage this year has made the case that AI-driven cyber risk is now a leadership issue, not only a technical one — boards increasingly need oversight of AI risk alongside traditional audit responsibilities, which raises the bar for GRC professionals expected to translate technical AI risk into business language executives can act on.

WHAT HASN'T CHANGED

The entry path remains genuinely open. GRC hiring managers continue to prioritize the ability to think critically about risk and communicate clearly over a computer science background or coding experience — the same accessible-entry-point case made last edition still holds, even as the skill bar for AI-specific competency rises within the role.

WHY THIS MATTERS FOR ANYONE CONSIDERING THE FIELD

The gap between demand and supply is not closing on its own. Roughly a third of hiring managers report struggling specifically to find candidates with AI/ML security skills, even as GRC and AI governance postings continue to climb. For someone building this skill set now, that gap is the opportunity.

THE BOTTOM LINE

The Cyber GRC Analyst role hasn't been replaced by AI — it's been redefined by it. The professionals building AI governance literacy into their existing risk, compliance, and communication skills right now are positioning themselves for the roles the market is actively struggling to fill.

CHAPTER 8:
VENDOR COMPARISON — AI GOVERNANCE PLATFORMS, REVISITED WHY WE'RE COMING BACK TO THIS TABLE

Last edition's four-platform comparison (Credo AI, Holistic AI, OneTrust AI Governance, and IBM watsonx.governance) is still a fair map of the landscape's starting points. This month's tool-of-the-month chapter adds a fifth consideration worth naming explicitly: platforms like Vanta that started in compliance automation and are now building dedicated AI-risk capabilities on top of an existing GRC foundation, rather than being purpose-built for AI governance from day one.

THE UPDATED FRAME

Tool	Best For	What This Month's Data Adds
Credo AI	Regulation-mapped compliance documentation	Remains the strongest fit for EU AI Act and NIST AI RMF-specific documentation needs, especially with the Omnibus's final text now settled
Holistic AI	Bias and fairness auditing	Still the deepest technical option for regulated, high-stakes use cases like hiring and lending
OneTrust AI Governance	Extending an existing GRC/privacy program	Its 2026 additions of real-time monitoring and AI agent detection are increasingly relevant given this month's agent-identity findings
IBM watsonx.governance	Model lifecycle governance at enterprise scale	Best fit unchanged — strongest inside an existing IBM ecosystem
Vanta (Agent for Risk)	Compliance-automation platforms extending into unified AI and vendor risk	New this month — best fit for teams that want AI and third-party risk in one continuously updated view rather than a dedicated point solution

HOW TO CHOOSE, STILL AS A BEGINNER

The rule of thumb from last edition holds: match the tool to where your organization already is, not where you hope to be in two years. What's new is that "where your organization already is" increasingly includes whether you're managing agent identities at all yet — if you're not, that gap may matter more than which vendor logo ends up on the dashboard.

THE BOTTOM LINE

No new platform unseated last month's four leaders — but the market's center of gravity moved visibly toward unified risk views that treat AI, vendor, and internal control risk as one continuously monitored surface rather than three separate spreadsheets. That shift is worth tracking even if you don't buy anything new this quarter.

CHAPTER 9:
COMMUNITY HIGHLIGHT
WE HEARD YOU

Readers told us last edition's community section read like a bulleted press release instead of an actual look inside what's happening at Skillweed. Fair criticism — here's the fuller version.

MEMBER SPOTLIGHT: MR K, GRC SPECIALIST

The class ends. The support doesn't. Every Skillweed graduate who lands a role gets access to our closed WhatsApp group — a private space exclusively for graduates who've landed a role to keep getting support after the class ends, not just during it. This month's spotlight comes straight out of that group.

One graduate, who we'll refer to here as Mr K, landed a job as a GRC Specialist and has stayed active in the group ever since — not to show off the win, but to keep asking the exact questions a new GRC hire actually runs into on the job.

This means real, on-the-job questions, answered in real time by people who've already sat in his seat. No searching forums, no guessing alone through a first review or a tool decision. Just direct access to people who've done it.

That's the actual value of the closed group: not a rehash of class material, but real questions from someone doing the job, answered by people who've already been where he is.

That's the part of Skillweed most people don't find out about until after they've already gotten the job — and it's often the reason they keep recommending us to the next person.

ONE MORE THING: THE HIRING TRUTH SESSION

If you know us, you know we don't stop at teaching the frameworks — a real part of the mission is making sure the class turns into an actual job. This month, that showed up directly: we hosted a session called The Hiring Truth, taught by Coach Queenette of PJY Group, whose business has a track record of successfully placing job seekers into roles. Coach Queenette extended a 15% discount on her placement services specifically to Skillweed community members.

The session was genuinely insightful, and it's exactly the kind of support this community exists to provide — the class gets you the skills, sessions like this one help make sure that translates into an offer. If you'd like the meeting link and recording, email info@skillweed.com.

AVAILABLE RESOURCES

- Free e-books: the resource library has grown past 60 downloadable guides, with a new addition this month specifically on agentic AI policy basics, built around the same framework in Chapter 5 of this issue. Browse the full library at skillweed.com/assets/resources.
- Paid books: the existing Amazon bestselling titles — Third-Party Risk Management, Risk Assessment, Guide to IT General Controls, and Cyber Security Operational Technology Best Practice — remain available at academy.skillweed.com/pages/resource-library.
- YouTube: the channel added a new short-form series this month specifically breaking down agentic AI incidents as they're reported, rather than only covering evergreen framework explainers. Subscribe at youtube.com/@skillweedAcademy.

WHAT MEMBERS ARE SAYING

> "The CRISC track didn't just teach me the framework — it taught me how to defend a risk rating in a room full of people who outrank me."

> "I used the AI Governance Internship capstone as an actual interview talking point, and it's the thing every interviewer asked follow-up questions about."

> "The community chat answered a question about scoping an agent's permissions faster than I could have found it in any vendor's documentation."

READY TO BUILD THE SKILLS THIS MAGAZINE JUST WALKED YOU THROUGH?

Join the October 3rd Cyber GRC & AI Governance Class

No coding. No IT background needed. Six weeks, live instructor-led classes, real-world labs, and a curriculum built around exactly what this issue covered — the Digital Omnibus, agentic AI governance, vendor risk, and the frameworks employers are actively asking for.

Reserve your spot: www.skillweed.com

MORE FROM SKILLWEED

Geotela — Geotela is launching soon and we couldn't be more excited! Geotela is Skillweed's geointelligence navigation tool, built for more than getting from A to B — it helps investors and explorers uncover real opportunity in a region, place by place. Already a community of 500+ and growing, with a waitlist open now.

Join the waitlist by visiting www.geotela.com

Landzille — If you've ever wanted something more in terms of investments especially in building a legacy, you should check out Landzille. Landzille goes beyond giving you great land deals in North Texas to providing after sales services that ensures you are tax compliant without breaking the bank. Visit Landzille.com to find available projects.

CHAPTER 10:
FINAL MESSAGE: TURNING KNOWLEDGE INTO ACTION

Understanding the trends in this edition is only useful if it leads somewhere. Here's where Skillweed can take you next.

Until Next Month, that's a wrap on the August edition of the Skillweed Cyber Magazine.

If there's one thing this issue should leave you with, it's this: the gap between "we should probably govern our AI agents" and "here's our agent identity inventory and our kill-switch procedure" is exactly where the professionals who win the next five years of this industry will separate themselves from everyone still writing policy for the chatbot era.

That's exactly the skill set we build at Skillweed.

We'll be back next month with more frameworks, more real-world incidents, and more of the practical, no-fluff guidance that's helped thousands of people break into this industry from zero.

Until then — stay curious, stay compliant, and stay ahead. The Skillweed Team


SkillWeed