immersive

WHITE PAPER - GEN AI

GenAI's Impact on Cybersecurity Skills and Training

By James Hadley, Chief Innovation Officer
and Founder, Immersive

Generative AI (GenAI) Has Become the Default Marketing Garnish for Every New Security Launch

Current estimates indicate that over 85% of vendor announcements now claim an LLM or AI component, up from just 18% in 2022[1].

Slide decks herald “self-writing detections” and “AI-generated training labs” that supposedly erase human toil overnight. And yet, the underlying talent crisis continues to deepen.

Like Florida emergency crews tracking a Category-4 hurricane, CISOs refresh zero-day feeds hourly, fully aware that storm landfall is a matter of when, not if.

Beneath the buzz, executives struggle to reconcile vendor optimism with daily reality. Boardrooms ask whether AI is truly improving the breach-cost curve, or merely adds another line item to the tooling budget, primarily to reassure non-technical directors. Such disconnect amplifies the risk of investing in capabilities that look transformative on paper yet leave fundamental skill gaps untouched.

4.2M

ISC² counts 4.2 million unfilled cyber roles globally, a deficit that rose 11% year-over-year despite record enrollment in bootcamps and degree programs.

ISC², Cybersecurity Workforce Study (2024)

62%

Nearly two-thirds (62%) of security leaders feel pressured to purchase AI features they don’t fully understand.

Accenture, Global CISO Survey (2025)

[1] Gartner, Hype Cycle for Cybersecurity (2025)
[2] IBM, Cost of a Data Breach Report (2025)
[3] Proofpoint Threat Labs, LLM-Generated Phishing Campaign Effectiveness Study (2025)
[4] ISACA, State of Cybersecurity (2025)

AI Threats vs. Cyber Skills Gaps

Organizations have invested heavily in AI-enabled SIEM and SOAR pipelines, but these platforms still rely on human expertise to interpret and act on alerts.

And both the skill deficiencies of the security teams involved and the misconfigurations in the solutions used contributed to 58% of successful attacks, even in environments deploying AI-augmented defenses[2]. Meanwhile, attackers are industrializing their own use of GenAI.

A recent proof-of-concept demonstrated that an LLM could generate polymorphic phishing emails that bypassed secure email gateways at a 68% success rate[3].

The paradox is acute: as tooling grows smarter, defenders remain stuck in catch-up mode, creating a widening gap that adversaries exploit with machine-speed precision.

The workforce skills shortage isn’t just a head-count problem; it erodes institutional memory and stretches existing teams to the point of burnout. Over half (54%) of SOC managers cite “staff fatigue” as their top operational risk, surpassing tooling complexity for the first time[4]. No amount of algorithmic triage can replace the calm judgment of an experienced analyst who has weathered multiple incident waves.

Without targeted up-skilling and retention programs, AI deployments risk becoming “Ferraris in the garage”—powerful machines idling for lack of qualified drivers.

IMMERSIVE INSIGHT →

Can GenAI be Used as a Cybersecurity Trainer?

Organizations serious about incident readiness use cyber skills development and training platforms to rehearse real attack scenarios.

Validate that every responder—from SOC analyst to legal counsel—can perform under pressure, and then convert those performance metrics into Resilience Scores the board and C-suite can understand. By treating human capabilities like any other critical control—tested, benchmarked, and reported—these companies gain the evidence they need to justify budgets, satisfy regulators, and sleep at night knowing their teams are truly prepared.

So, can GenAI simply create the training content and close the skills gap?

01

Large language models (LLMs) remain prone to hallucination, context blind spots, and outdated recommendations.

In a recent documented red-team exercise, an LLM was tasked with drafting a lab on container escape; 21% of the generated commands were deprecated or dangerously destructive in production. Worse, the model offered confident but incorrect rationales that novice learners struggled to debunk. Boards can hardly certify readiness on materials whose accuracy even their authors—a black-box neural net—cannot verify.

University of Maryland Cyber Range, Evaluating LLM-Generated Security Labs (2025)

03

AI as a sole contributor is considered dangerous.

The EU AI Act classifies AI systems used in “management and operation of critical infrastructure” as high-risk, demanding tight human oversight and audit trails.

European Parliament, Artificial Intelligence Act (2025)

The takeaway: GenAI can add value to improving an organization's cybersecurity resilience, but it can’t simply be put in charge of it.

The current state of GenAI indicates we’re far from it:

02

AI struggles with domain nuance.

A separate pilot involving financial-sector simulations showed that the LLM failed to incorporate PCI-DSS and SOX controls, leaving critical regulatory requirements unmentioned. Subject-matter experts spent more time correcting misaligned content than they would have spent writing from scratch.

SANS Institute, Financial Sector AI Training Pilot Findings (2025)

04

Cyber insurers don’t trust AI.

Meanwhile, a leading cyber-insurer announced a 7% surcharge on policies where AI training content lacks documented human review

Marsh McLennan, Cyber Insurance Underwriting Trends Q2 2025 (2025)

Then, what exactly is GenAI’s role in cyber skills development and training?

What GenAI is and is Not in Cybersecurity Training

LLMs predict the next token in a sequence, not the next best defensive action on a live network. They excel at summarizing threat-intel feeds, drafting lab scaffolds, and translating jargon into plain English, but they lack the situational awareness to guarantee that a newly generated exercise is accurate, current, and safe.

When unreviewed content feeds directly into a learning portal, outdated commands or hallucinated controls can mislead learners and even seed dangerous habits. In short, GenAI is better suited to be a high-speed drafting assistant, and definitely not a domain-qualified instructor.

The value of GenAI in cybersecurity training is, therefore, selective, not universal. It excels at automating repetitive or data-heavy tasks—sketching initial lab outlines, clustering emerging TTPs, or offering context-sensitive hints—but falters when nuanced judgment, regulatory nuance, or ethical trade-offs come into play.

In some critical areas, such as certifying that a scenario meets compliance mandates or balancing business-impact decisions during a live drill, AI’s role is effectively non-existent without explicit SME validation. Organizations must therefore deploy GenAI where it truly augments human expertise, while enforcing strict guardrails everywhere else.

The value of GenAI in cybersecurity training is, therefore, selective, not universal

GenAI does have a role, but it's one that needs to be honed and utilized to improve the outcomes of cyber readiness efforts.

There are three key tenets to how organizations should see GenAI's role in the cybersecurity skills and training market.

GenAI Can't Replace All Cybersecurity Experience and Expertise

Real incidents rarely unfold like textbook examples. Responders must weigh business constraints, compliance deadlines, and ambiguous evidence under intense time pressure.

LLMs cannot replicate the contextual judgment of practitioners who have navigated multiple breaches, balanced legal-disclosure timelines, and briefed executives in boardrooms that double as war rooms. A recent blue-team workshop pitted an AI-generated incident-response playbook against a human-curated one.

Analysts following the AI script achieved containment in 56 minutes, while the human script averaged 34 minutes—an intelligence gap that doubled the potential possible data exfiltration in that scenario[5].

GenAI also struggles with socio-technical nuance. When a hospital drills a ransomware scenario, triage doctors must decide whether to divert ambulances even before the IT team restores EHR access. Gen AI cannot evaluate ethical trade-offs that weigh human life against data confidentiality. Nor can it grasp sector-specific legal mandates: under PCI DSS 4.0, cardholder environments must be segmented within six hours of compromise, but under HIPAA, a covered entity has 60 days to disclose. Seasoned responders internalize these subtleties; a model fine-tuned on broad internet text does not.

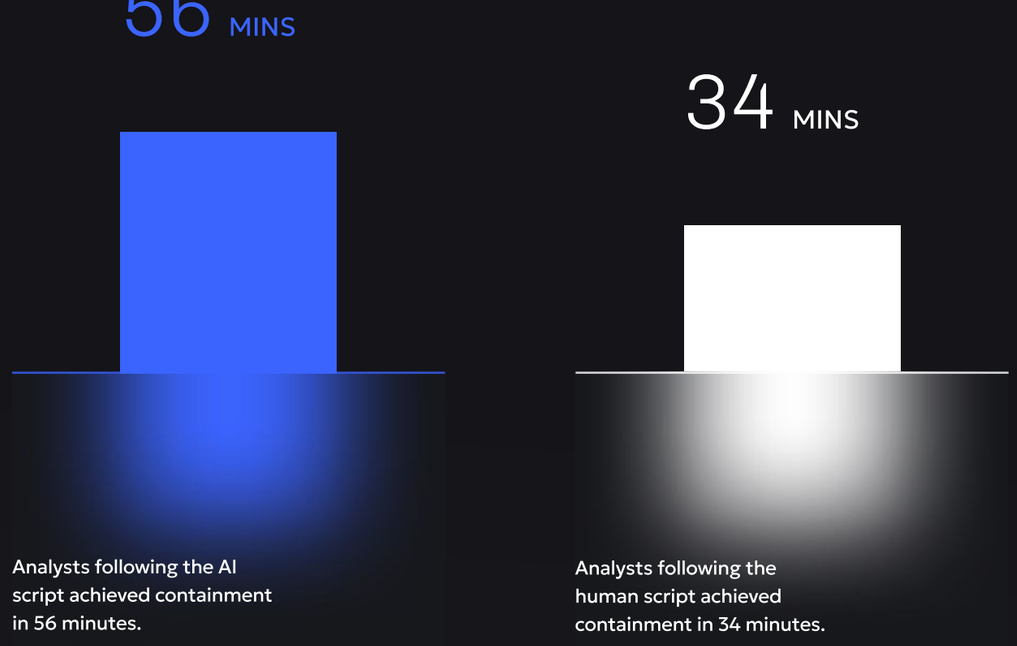

Analysts following the AI script achieved containment in 56 minutes.

Analysts following the human script achieved containment in 34 minutes.

[5] SANS Institute, AI vs. Human IR Playbooks Workshop Findings (2025)

IMMERSIVE INSIGHT →

GenAI Can be Used as a Force-Multiplier

While GenAI can't replace experts - it can dramatically expand their reach with effective prompt engineering.

LLMs can be used to draft objectives, code snippets, and assessment questions, improving cybersecurity SME lab-development throughput. This creates a second-order benefit: SMEs spend less time on rote editing and more on injecting high-impact context—sector-specific regulations, edge-case exploit chains, and executive-communication drills. The projections are that such AI-assisted workflows could save the average training vendor 4,400 developer hours annually, freeing budget for deeper scenario research[6].

GenAI also handles multilingual localization at near-real time, enabling a global workforce to upskill simultaneously instead of in staggered release cycles. Threat-intel analysts can leverage LLM summarization to condense 60-page CVE write-ups into three-paragraph briefs in single-digit minutes, accelerating the turnaround from disclosure to lab publication.

Yet, note that every AI draft routes through a human-in-the-loop review that checks for hallucinations, outdated commands, and compliance gaps—maintaining quality without sacrificing velocity.

[6] IDC, AI Productivity Benefits in Cyber Training Market (2025)

[7] Coursera, AI Tutoring Impact (2024)

GenAI Can Enhance the Learner Experience

For learners, GenAI excels as a just-in-time tutor and personalized navigator.

Adaptive engines analyze drill performance, career objectives, and even preferred learning styles to queue the next best lab, reducing cognitive overload and cutting time-to-competence.

GenAI-guided sequencing improves completion rates by 22% across technical subjects[7], and lowers lab abandonment after integrating chat-assist hints to help students that are stuck.

In-lab chat assistants can translate error logs, suggest alternative commands, or link to bite-sized videos that clarify complex exploits. Accessibility improves too: voice-enabled prompts aid visually impaired learners, and real-time translation opens specialized labs to non-English speakers.

At the organizational level, a no-code simulation builder lets risk managers describe a "supply-chain ransomware scenario," which the LLM expands into a drill complete with role injects, scoring rubrics, and post-mortem templates. This feature slashes custom-exercise development time from weeks to hours, democratizing simulation design beyond the security team.

Don’t Forget Risk and Governance of AI in Your Training

Velocity without governance invites catastrophe. The NIST AI Risk Management Framework advises explicit human review, data-provenance tracking, and bias testing for any GenAI used in critical-infrastructure training. Executive sponsors should require a “GenAI bill of materials” listing model version, training-data window, validation status, and copyright provenance for each generated asset. The EU AI Act, previously mentioned, goes further, mandating “post-market monitoring” of high-risk AI systems to ensure they don’t drift into non-compliance over time.

Put simply, GenAI’s use in cybersecurity training requires governance.

Organizations must also establish rollback procedures: if a model update introduces faulty commands, content should be quarantined quickly via version control, and not linger in a catalog for weeks. Audit logs should capture which SME approved each AI-generated item, preserving a chain of accountability for regulators and insurers.

Used carelessly, GenAI hallucinates commands, obscures accountability, and inflates the very risk it aims to mitigate. But when balanced correctly, it accelerates content velocity, personalizes learner journeys, and frees experts to focus on high-value judgment calls—shifting GenAI from a potential liability into a strategic asset.

But how can you practically and responsibly put GenAI to use in the context of cyber readiness?

Practical Ways to Utilize Generative AI Inside a Cyber Skills Development and Training Program

Before any dashboards light up or labs go live, a program needs guardrails that spell out how — and how not — to apply GenAI. These principles anchor automation to human judgment, establish transparency rules, and hard-wire ethical boundaries so machine speed never outruns professional rigor. These guardrails should include:

01 Human-in-the-Loop Validation

Every AI-generated lab, hint, or simulation passes SME review before release; experts verify accuracy, regulatory fit, and ethical soundness.

02 Transparency and Explainability

Prompts, model versions, and validation steps are logged so auditors can trace how content was produced and why it’s trustworthy.

03 Test that Content is Still Relevant (if Not, Retire it)

New materials face quarterly drills; if they don’t raise Resilience-Score deltas, they’re revised or retired.

04 Framework Mapping

Assets are tagged to MITRE ATT&CK techniques and NIST NICE tasks, giving boards a common language for coverage.

05 Ethical Guardrails

Bias scans, copyright checks, and privacy filters run automatically; anything that fails routes to legal review before publication.

These guardrails should already be a part of a cyber skills development and training platform, but are listed to ensure that they be put in place via people and process if the technology isn’t doing it already.

The following sections highlight concrete examples of where, why, and how GenAI can accelerate the impact of cyber skills development and training efforts without sacrificing human oversight.

AI-Augmented Content Lifecycle

PRACTICAL EXAMPLE: TURNING A PROMPT INTO A VALIDATED LAB

GenAI can draft a lab outline in seconds, but that draft is only the starting point. The lifecycle below shows how each AI artifact moves through review, safety checks, pilot testing, and continuous evaluation so every asset remains accurate, compliant, and measurable from day one to retirement.

Stage	Key Action	AI Contribution	Human Gatekeeper	Evidence Artifact
Prompt & Draft	Generate initial outline	LLM + threat-feed context	Senior SME	Draft doc
Review & Edit	Add sector regs, fix hallucinations	Grammar, localization	Compliance SME	Red-lined diff
Safety & Bias Scan	Detect unsafe or biased commands	Automated scanners	AI Ethics Lead	Scan log
Pilot Drill	Run lab in staging range	Hint-bot support	SOC testers	Pilot AAR
Publish	Release to catalog	Auto-tag frameworks	Content Ops	Version record
Evaluate	Measure RS delta after use	Analytics LLM	CISO	Quarterly dashboard
Retire/Update	Deprecate stale content	Model suggests fixes	SME approves	Change log

GenAI-Enhanced User Experience

PRACTICAL EXAMPLE: ADAPTIVE COACHING THAT SPEEDS UP COMPETENCE

Learners judge training by how quickly it helps them solve real problems, not by flashy widgets. The use-cases below illustrate how Gen AI personalizes paths, answers questions in real-time, and even builds custom drills—transforming a static catalog into an adaptive coach that shortens the journey from “I don’t know” to “I’ve got this.”

Learner Chat-Assist

An embedded LLM explains why a PowerShell command failed, cites MITRE T1055, and suggests a corrective flag—reducing frustration and boosting retention.

Adaptive Pathing

Drill scores and career goals feed a recommendation engine that queues the lab most likely to raise the learner’s weakest RS sub-score.

Real-Time Hints

During red-team ranges, an AI coach offers escalating clues increasing completion rates without SME intervention.

No-Code Sim Builder

Risk managers type “B2B supplier compromise,” and the LLM outputs a full crisis scenario with injects, scoring weights, and legal-notice timelines.

Using AI to Measure it's Own Impact

AI moves from experimental spend to a permanent budget line only when it proves its worth. The metric set below converts faster lab production, higher completion rates, and RS gains into hard numbers a CFO or insurer can weigh against cost.

PRACTICAL EXAMPLE: TRANSLATING AUTOMATION GAINS INTO BUSINESS VALUE

+10

Increase to resilience score as a target using AI

Metric	Baseline	Target w/AI	Evidence Source
Resilience-Score Delta	+0	+10 in 12 mo	RS dashboard
Lab Cycle Time	14 days	9 days	Content tracker
SME Review Hours Saved	—	30% reduction	Timesheets
Learner NPS	46	60	Post-lab survey
Insurance Discount	0%	≥ 10%	Broker statement

Governance and Roles

Role	Accountability	Cadence
CISO	AI policy, budget	Quarterly board
Head of L&D	Content QA	Monthly
AI Ethics Lead	Bias, privacy	Per release
SOC Manager	Pilot drills	Sprint demo
Compliance Counsel	Framework mapping	Quarterly
Board Risk Committee	Oversight	Quarterly

Even perfect tools misfire when accountability is vague. Governance assigns named owners to every AI touch-point—from policy writing and bias scanning to pilot-drill sign-off—ensuring metrics stay trustworthy and decisions traceable. The table below outlines an example set of governance roles needed to oversee AI's use, accuracy, validity, and value.

What to Look For in an AI-Powered Cyber Skills Development and Training Platform

GenAI branding is everywhere, but "AI-ready" slides don't guarantee the platform delivers real value—or even uses machine learning beyond a marketing blurb.

When evaluating solutions, focus on concrete functionality that proves AI is embedded in ways that accelerate learning, strengthen measurement, and preserve auditability. The checklist below outlines essential capability areas and the specific features that separate substance from hype.

Capability Area	Why it Matters	Specific Features to Verify
AI Transparency & Auditability	Boards and regulators must trust any AI that shapes readiness evidence.	▪ Immutable logs of prompts, model versions, and outputs ▪ Exportable audit trail in JSON/CSV ▪ Explainability dashboard that traces how an answer was formed
Human-in-the-Loop Controls	Prevents hallucinations and compliance mis-steps.	▪ Mandatory SME approval workflow for every AI-generated lab or drill ▪ Role-based permissions to accept, edit, or reject AI drafts
Adaptive Learning Engine	Personalizes pathways so time-pressed staff close the biggest gaps first.	▪ Real-time skill analysis tied to Resilience Score or similar metric ▪ Auto-recommendations for "next best lab/drill" with rationale
AI-Assisted Content Creation	Cuts lead time without sacrificing rigor.	▪ Draft lab/scenario generator seeded by ATT&CK® technique or custom prompt ▪ Side-by-side diff showing AI draft vs. SME final ▪ Time-to-publish metrics for proof of velocity gain
Real-Time Learner Assistance	Reduces friction and abandonment.	▪ In-lab chat coach that references live environment logs—not canned FAQs ▪ Context-aware hints that cite frameworks or MITRE IDs
Framework Auto-Mapping	Speeds audit prep and highlights control gaps.	▪ One-click tagging of labs and drills to MITRE, NIST CSF, PCI-DSS, or DORA ▪ Heat map that updates automatically as new content is added
Safety & Bias Filters	Protects IP and ensures inclusive training data.	▪ Automated scans for deprecated commands, malicious payloads, and toxic language ▪ Copyright detection and data-provenance tracking
Performance Evidence	Converts AI claims into measurable ROI.	▪ Dashboards correlating AI-generated content to Resilience Score deltas ▪ Reports showing reduced lab cycle time or increased learner completion
Integration & Export Options	Keeps AI outputs from becoming data silos.	▪ REST or GraphQL API for pushing Resilience metrics into GRC tools ▪ Webhook support for insurance questionnaires or board portals
Data Privacy & Governance	Avoids accidental disclosure of sensitive logs or PII.	▪ On-prem or region-pinned model-hosting options ▪ Prompt-firewall to prevent leakage of regulated data

Should Companies Use GenAI to Downsize?

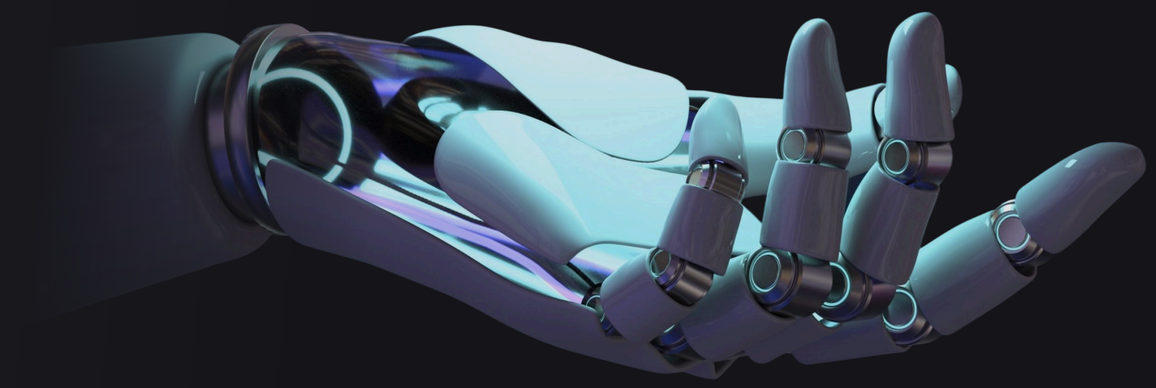

IMMERSIVE INSIGHT →

There are more than a few prognosticators warning that GenAI advancements will result in massive job loss.

While it may be tempting for leaders to reduce headcount in anticipation of the automation new AI tools will bring, it would be premature to cut roles without a proper strategy and strong cybersecurity hygiene.

In other words, if you have gaps in cyber readiness and lack the ability to prove and improve critical skills across the organization, cutting staff will not make any of these problems better.

The truth is, while GenAI tools can automate and accelerate vulnerability detection and other tasks, they also introduce new risks, such as prompt injection attacks, that require a sophisticated human defense - a cyber readiness workforce. Savvy leaders will need to ensure their teams have the capabilities needed to defend against new threats by implementing a cyber readiness program that continuously assesses, builds, and proves cyber resilience.

Leveraging GenAI for a Better Cyber Readiness Outcome

GenAI is reshaping security, but its real power emerges only when paired with human judgment, rigorous governance, and clear evidence of impact.

GenAI can and does play a role in accelerating value in cybersecurity skills and training—from content drafting to adaptive coaching—as long as the quality of its work product is safeguarded with human-in-the-loop controls, transparency, and ethics.

Organizations that leverage GenAI as part of their cybersecurity skills and training platforms within a broader Cyber Readiness Program can transform raw machine intelligence into a strategic advantage—closing skill gaps faster, proving readiness to leadership, and adapting at the pace of attackers, and positioning every role to respond with greater confidence and efficacy when the next incident strikes.

To learn more about Immersive's stance and thought leadership on GenAI, visit us here.

Resilience

Immersive is trusted by the world’s largest organizations and governments, including Citi, Pfizer, Humana, HSBC, the UK Ministry of Defence, and the UK National Health Service. We are backed by Goldman Sachs Asset Management, Ten Eleven Ventures, Menlo Ventures, Summit Partners, Insight Partners and Citi Ventures.

www.immersivelabs.com | sales@immersivelabs.com