MIS Quarterly Executive

Volume 24 | Issue 4 Article 4

December 2025

Exploring the Agentic Metaverse's Potential for Transforming Cybersecurity Workforce Development

Ersin Dincelli

Haadi Jafarian

MIS Quarterly Executive

Exploring the Agentic Metaverse's Potential for Transforming Cybersecurity Workforce Development

The agentic metaverse (a convergence of immersive metaverse platforms and multi-agent systems) has the potential to transform how organizations approach workforce training and development by enabling scalable, dynamically generated and adaptive learning experiences. This article presents findings from an exploratory qualitative study of an AI-driven metaverse prototype designed for cybersecurity training, which was evaluated by 53 cybersecurity professionals. Drawing on their feedback, we identify challenges and provide six recommendations to guide the operationalization of the agentic metaverse, with particular emphasis on implementation and governance considerations.[1,2]

Ersin Dincelli
University of Colorado Denver (U.S.)

Haadi Jafarian
University of Colorado Denver (U.S.)

Addressing Cybersecurity Workforce Needs through the Agentic Metaverse

Traditional training methods, such as classroom-based instruction and static online courses, often struggle to keep up with rapid technological advancements and the evolving demands of today's job market. To remain competitive, organizations need continuous skill development, hands-on experience and personalized learning pathways. Workforce training and development play a crucial role in organizational success, so companies are seeking innovative solutions that go beyond conventional approaches to improve engagement, retention and practical skill development.[3] Addressing the needs of workforce development, particularly in critical fields like cybersecurity and healthcare, requires immersive and adaptive learning environments that replicate real-world scenarios.[4]

1 Elizabeth White Baker is the senior accepting editor for this article. The authors are grateful for her guidance and recommendations, which significantly strengthened this article.

2 We acknowledge the support of a National Science Foundation grant (Award #2335839) awarded to Ersin Dincelli. We thank the Information Systems Audit and Control Association (ISACA) Denver Chapter for the opportunity to demonstrate our prototype and gather insights from its members during its 2024 fall meeting. We also thank the review team for their constructive feedback throughout the review process..

3 Kapadia, S. *Companies Large and Small Are Using AI for Employee Onboarding. It Can Save HR Days of Time*, Business Insider, March 12, 2025, available at https://www.businessinsider.com/generative-ai-employee-onboarding-human-resources-2025-3/.

4 Galletta, D. F., Moody, G. D., Lowry, P. B., Willison, R., Boss, S., Chen, Y., Luo, X. R., Pienta, D., Polak, P., Schuetz, S. and Thatcher, J. B. "Balancing Fear and Confidence: A Strategic Approach to Mitigating Human Risk in Cybersecurity," *MIS Quarterly Executive* (24:1), March 2025, pp. 1-15.

DOI: 10.17705/2msqe.00122

The cybersecurity workforce faces persistent challenges, including an evolving threat landscape, high burnout rates, a growing workforce gap and a shortage of skilled talent.[5] These issues are further exacerbated by misconceptions about the field, which often discourage younger generations from pursuing careers in cybersecurity.[6] Addressing these challenges is crucial to safeguarding organizational and national security, especially as the demand for skilled cybersecurity professionals continues to grow in response to the increasing frequency and complexity of cyber threats.[7] As a result, there is a growing imperative to explore innovative training solutions that not only attract new talent but also more effectively prepare professionals to navigate the evolving complexities of cybersecurity.

Cybersecurity requires continuous training due to its dynamic threat landscape and the critical importance of preparedness.[8] Given limited organizational resources and significant information asymmetries, cybersecurity training solutions must be both scalable and adaptive to effectively address diverse organizational contexts and individual learning needs. Cyber defense also involves collaboration and time-sensitive decision-making, which are best cultivated through high-fidelity simulations that mirror real-world pressure and complexity.[9] In reality, many cybersecurity practitioners acquire expertise primarily through direct and experiential learning.[10] Though traditional training methods provide valuable foundational knowledge, they often rely on one-size-fits-all approaches[11] and lack the experiential depth necessary to prepare professionals for emerging cyber threats.[12]

The integration of metaverse and AI technologies creates new possibilities for authentic and experiential learning experiences.[13] Specifically, the metaverse enables high-fidelity, immersive and persistent virtual environments where employees can train in interactive, risk-free scenarios that replicate real-world complexity within controlled settings.[14] Complementing this, AI agents can simulate adversaries, guide training efforts and deliver real-time, personalized content tailored to individual learning needs.[15] The convergence of these immersive virtual environments and multi-agent systems, referred to as the *agentic metaverse*, presents a unique opportunity to transform workforce training and development by creating realistic and adaptive scenarios to foster active engagement in unpredictable and high-stakes situations (e.g., nation-state attacks or insider threats) that closely resemble real-world cybersecurity challenges.

In this article, we present insights from an exploratory qualitative study on agentic metaverse-based training. We developed an AI-driven training prototype delivered within a metaverse-like environment to simulate immersive and personalized learning experiences for cybersecurity training. This prototype allowed us to explore the synergy between immersive virtual environments and AI-driven learning for addressing organizations' workforce training and development challenges in the context of cybersecurity. The study involved 53 cybersecurity professionals who participated in a local chapter meeting of the Information

5 Bradley, T. *The Cybersecurity Burnout Risk Crisis Is Reaching The Breaking Point*, Forbes, October 15, 2024, available at https://www.forbes.com/sites/tonybradley/2024/10/15/the-cybersecurity-burnout-crisis-is-reaching-the-breaking-point/.

6 Higgins, K. J. *Millennials Not Pursuing Cybersecurity Careers*, Dark Reading, October 26, 2015, available at https://www.darkreading.com/cybersecurity-operations/millennials-not-pursuing-cybersecurity-careers/.

7 *2024 ISC2 Cybersecurity Workforce Study*, ISC2, October 31, 2024, available at https://www.isc2.org/Insights/2024/10/ISC2-2024-Cybersecurity-Workforce-Study/.

8 Safitra, M. F., Lubis, M. and Fakhrurroja, H. "Counterattacking Cyber Threats: A Framework for the Future of Cybersecurity," *Sustainability* (15:18), September 2023, pp. 1-32.

9 Chowdhury, N., Katsikas, S. and Gkioulos, V. "Modeling Effective Cybersecurity Training Frameworks: A Delphi Method-Based Study," *Computers & Security* (113), February 2022, pp. 1-15.

10 Towhidi, G. and Pridmore, J. "Increasing Cybersecurity Interest and Self-Efficacy through Experiential Labs," *Issues in Information Systems* (23:2), February 2022, pp. 119-131.

11 Bendler, D. and Felderer, M. "Competency Models for Information Security and Cybersecurity Professionals: Analysis of Existing Work and a New Model," *ACM Transactions on Computing Education* (23:2), June 2023, pp. 1-33.

12 Angafor, G. N., Yevseyeva, I. and He, Y. "Game-Based Learning: A Review of Tabletop Exercises for Cybersecurity Incident Response Training," *Security and Privacy* (3:6), July 2020, pp. 1-19.

13 Zhang, Q. "Secure Preschool Education using Machine Learning and Metaverse Technologies," *Applied Artificial Intelligence* (37:1), June 2023.

14 Dincelli, E. and Yayla, A. "Immersive Virtual Reality in the Age of the Metaverse: A Hybrid-Narrative Review Based on the Technology Affordance Perspective," *Journal of Strategic Information Systems* (31:2), June 2022.

15 Molenaar, I. "Towards Hybrid Human-AI Learning Technologies," *European Journal of Education* (57:4), September 2022, pp. 632-645.

Systems Audit and Control Association (ISACA). Participants provided feedback on the prototype through a series of questions. Early feedback was highly positive: 92% believed it would be effective for professional training, 98% believed it would be effective for higher education and 96% agreed that its benefits outweighed potential drawbacks.

Drawing on this exploratory qualitative study, we present practitioner-oriented insights focused on three key areas:

- How the convergence of the metaverse and AI agents, referred to as the agentic metaverse, can transform learning for cybersecurity professionals
- Key challenges associated with implementing agentic metaverse-based training solutions in cybersecurity workforce development
- Practical recommendations to maximize benefits and mitigate challenges in deploying the agentic metaverse for cybersecurity training

Our findings suggest that integrating immersive, experiential training with AI-driven personalization can substantially enhance workforce development efforts. Drawing on feedback from cybersecurity professionals, we identify key challenges and translate them into six practical recommendations. These recommendations serve as a practitioner-oriented playbook for effectively operationalizing the agentic metaverse in organizational training and development contexts.

Understanding the Agentic Metaverse

Before defining what we mean by the agentic metaverse, we first explain the terms metaverse, AI agents and agentic AI.

Metaverse

The metaverse is commonly conceptualized as a persistent, immersive and interactive digital environment that integrates a range of technological ecosystems, such as extended reality (XR),[16] AI, the Internet of Things (IoT), blockchain and spatial computing, to create a shared, decentralized virtual space.[17] Within this space, users can interact, work, learn, socialize and engage with digital assets or other users.[18] Unlike isolated virtual environments, the metaverse allows user identities and digital assets to move seamlessly across different platforms, continues to evolve even when users are offline, and offers users a high level of control and interaction. Together, these capabilities create unique immersive experiences.[19] Within these environments, AI-powered entities can interact dynamically with users, one another and the virtual world itself, offering new opportunities for high-fidelity learning, collaboration and social connection.

AI Agents

An AI agent is an intelligence system that operates autonomously on behalf of users or other systems.[20] Leveraging machine learning, deep learning, reinforcement learning and natural language processing to understand their environment, AI agents learn from interactions, interpret requests and adapt over time to make decisions, perform tasks or respond to inputs.[21] Because they work continuously, AI agents can boost productivity, streamline workflows and create richer user experiences across various industries and platforms,[22] including virtual spaces like the metaverse.[23] The core components of an AI agent include the environment in which it operates, the mechanisms it uses to perceive and process inputs, its internal state or memory, the

16 XR hardware and software includes virtual reality (VR), augmented reality (AR) and mixed reality (MR) technologies, including head-mounted displays, motion controllers, spatial sensors, and software platforms that enable immersive and interactive experiences.

17 Ball, M. *The Metaverse: And How It Will Revolutionize Everything*, Liveright Publishing, 2022.

18 Dolata, M. and Schwabe, G. "What Is the Metaverse and Who Seeks To Define It? Mapping the Site of Social Construction," *Journal of Information Technology* (38:3), February 2023, pp. 239-266.

19 Lacity, M., Mullins, J. K. and Kuai, L. "Evolution of the Metaverse," *MIS Quarterly Executive* (22:2), March 2023, pp. 165-173.

20 Gutowska, A. *What Are AI Agents?* IBM, 2024, available at https://www.ibm.com/think/topics/ai-agents/.

21 Benbya, H., Davenport, T. H. and Pachidi, S. "Special Issue Editorial: Artificial Intelligence in Organizations: Current State and Future Opportunities," *MIS Quarterly Executive* (19:4), December 2020, pp. ix-xxi.

22 Rai, A., Constantinides, P. and Sarker, S. "Next Generation Digital Platforms: Toward Human-AI Hybrids," *MIS Quarterly* (43:1), March 2019, pp. iii-ix.

23 Huynh-The, T., Pham, Q. V., Pham, X. Q., Nguyen, T. T., Han, Z. and Kim, D. S. "Artificial Intelligence for the Metaverse: A Survey," *Engineering Applications of Artificial Intelligence* (117), January 2023.

decision-making or policy mechanisms that guide its behavior and the actions it takes in response.

Agentic AI

While traditional AI agents typically operate based on predefined rules or static learning models, the emergence of agentic AI introduces a new level of autonomy and sophistication. Self-learning agentic AI agents extend beyond conventional AI agents, which often lack self-directed intelligence.[24] Agentic AI not only responds to instructions but also adds self-improving capabilities, perceiving context, planning, coordinating with other agents and revising strategies to pursue evolving objectives under human-defined constraints. This advancement is particularly important for training, as it enables adaptive scenarios that respond dynamically to learner behavior rather than following fixed scripts.[25]

Definition of the Agentic Metaverse

We define the agentic metaverse as a new class of immersive and persistent virtual ecosystem in which governing logic is delegated to goal-driven AI agents capable of autonomously and collaboratively perceiving, reasoning and acting based on real-time contextual cues and evolving objectives. Unlike conventional metaverse applications, typically limited to pre-scripted and human-authored scenarios, the agentic metaverse distributes decision-making authority across autonomous agents and human participants. It dynamically generates narratives, defines and adapts rules, and creates virtual objects and environments in direct response to user behavior and system-level objectives, thus allowing the virtual world to coevolve without continuous human intervention.

The agentic metaverse pairs the adaptive agent infrastructure with the high sensory fidelity and embodiment capabilities of immersive platforms to enhance learning effectiveness. Features such as spatial audio, haptics, and natural hand, eye and voice input contribute to a strong sense of presence, helping to anchor user attention, engage spatial memory and facilitate skill acquisition in environments that closely resemble real-world operational conditions. This combination of deep immersion and agent-led adaptability enables realistic and open-ended training experiences that traditional pre-scripted simulations are often unable to deliver or sustain.

The Agentic Metaverse and the Future of Workforce Training and Development

In contrast to traditional training methods that often struggle to accommodate individual learning styles and evolving industry demands,[26] the agentic metaverse can continuously assess performance, adjust content in real time and personalize learning pathways, while still aligning with instructional objectives. Agentic metaverse-based training therefore allows skill verification to emerge through demonstrated actions rather than self-reported knowledge or retrospective testing. Beyond individual learning, the agentic metaverse can foster collaboration. Agentic metaverse-based training simulates real-world, team-based training scenarios where agents take on specialized roles, such as attackers, defenders, mentors and analysts in different scenarios, injects situational complexity on demand, provides just-in-time feedback and continuously recalibrates learning pathways based on evolving organizational goals and individual performance. This supports not only immediate skill acquisition but also long-term workforce development and role progression. Agentic metaverse-based training, therefore, provides unique human-AI multi-agent teaming experiences that are unattainable in traditional training sessions.

Additionally, the agentic metaverse can enable proactive skill-gap analysis, as it leverages aggregated, privacy-preserving industry data and performance trends to anticipate future workforce needs and recommend targeted upskilling opportunities. By integrating large language models (LLMs) and visual-language

24 For additional details on agentic AI, see *What Is Agentic AI*, Amazon Web Services, available at https://aws.amazon.com/what-is/agentic-ai/.

25 Stryker, C. *Agentic AI: 4 Reasons Why It's the Next Big Thing in AI Research*, IBM, 2025, available at https://www.ibm.com/think/insights/agentic-ai/.

26 Dincelli, E. and Chengalur-Smith, I. "Choose Your Own Training Adventure: Designing a Gamified SETA Artefact for Improving Information Security and Privacy through Interactive Storytelling," *European Journal of Information Systems* (29:6), July 2020.

models (VLMs),[27] the agentic metaverse can dynamically generate natural language dialogue and coordinate with specialized text-to-image or text-to-3D generative models to create rich visual content in real time. This multimodal integration enhances realism and immersion, provided that content moderation and governance frameworks are in place.[28]

AI-Driven Metaverse Training Prototype as a Precursor to the Agentic Metaverse

We conducted an exploratory qualitative study with cybersecurity professionals to gather insights about an educational metaverse platform enhanced with various AI features, designed to support cybersecurity workforce training and development. This prototype serves as a foundational step toward the development of the agentic metaverse by enabling early evaluation of immersive and AI-supported learning experiences that mirror the envisioned capabilities of the agentic metaverse. Specifically, we demonstrated a functional prototype during the ISACA Denver Chapter's fall 2024 meeting. ISACA is an international professional association focused on IT governance, security and assurance. With more than 200 local chapters globally, it offers education and networking opportunities for IT professionals.[29] The meeting was conducted virtually via Zoom, involving 53 cybersecurity professionals from various local companies. Prior to the meeting, we collaborated with organizers to incorporate polling and open-ended questions into the session's agenda. Participation was required for meeting attendance credit, which facilitated a high response rate among attendees.

Description of the Training Prototype

The prototype showcased key elements related to the metaverse, including an interactive 3D environment, a dynamic virtual world, event-based quests, customizable avatars, non-player characters (NPCs), a virtual marketplace and a digital assistant. It also incorporated gamification features such as leaderboards, achievements and rewards to promote engagement. Additionally, the prototype featured several AI-driven components, including a chatbot powered by OpenAI GPT-4o, AI-generated avatars and quests developed through dynamic storytelling, and a journal enhanced by generative AI for content follow-up and note-taking.

Participants were briefly introduced to key agentic metaverse concepts, including metaverse, AI agents, multi-agent systems and agentic AI. We also shared our vision in detail to foster meaningful discussions among participants. While developing a fully functional agentic metaverse was beyond the scope of our study, demonstrating a prototype as a precursor to the agentic metaverse allowed participants to experience how such a system could be applied in real-world training scenarios and envision its potential role in shaping the future of workforce development. The discussion and demonstration, in turn, provided us with valuable insights from cybersecurity professionals. Figure 1 presents a screenshot of the prototype.

Analysis of Prototype Feedback

Discussions during the ISACA meeting revealed a consistent pattern across participant reactions. On the one hand, there was strong enthusiasm for the concept's novelty and training potential. Participants highlighted how the agentic metaverse could make abstract cybersecurity challenges more tangible and engaging. On the other hand, participants moved beyond their initial enthusiasm to raise practical concerns. They asked questions about how the platform would be used in practice (usability), how its development and content would be governed responsibly (governance) and how it would be integrated into existing organizational workflows (rollout). These themes of excitement and pragmatism framed the discussion and revealed both positive sentiments and concerns, as well as challenges that any real-world deployment of the agentic metaverse would need to address.

27 For additional details on visual language models, see *What Are Vision Language Models*, NVIDIA, available at https://www.nvidia.com/en-us/glossary/vision-language-models/.

28 Liu, X., Tai, Y. W. and Tang, C. K. "Agentic 3D Scene Generation with Spatially Contextualized VLMs," *arXiv preprint arXiv:2505.20129*. May 2025.

29 For more details on Information Systems Audit and Control Association (ISACA), see *We Are ISACA*, available at https://www.isaca.org/about-us/who-we-are/.

Figure 1: Screenshot from the Prototype

Positive Sentiments and Concerns

From the discussion, four themes emerged, one reflecting a positive sentiment and three highlighting areas of concern:

1. Engagement and realism: Many participants reacted to the excitement of the experience: "Impressive and helps visualization," "Ambitious, well done!" and "Pretty cool!" were some of the comments. Some participants noted that the scenarios felt authentic and closely aligned with real-world challenges, reinforcing the prototype's value as a training tool. Typical comments were: "This seems like a great learning tool" and "Great tool for education and training."

2. Learning curve and usability: A recurring concern involved onboarding and the potentially steep learning curve, particularly for participants without gaming experience. As one participant put it: "Difficult to understand the tool." Another added: "… might be quite a learning curve to use it for someone who does not do gaming."

3. Governance, safety and content quality: Participants also underscored the importance of risk management, content curation, data protection and instructor oversight. Concerns voiced included the "risk of [the] game consuming confidential information" and that "controlling the quality of open-source additions could be challenging; [they] may not be based on solid 'ethical hacking' course material."

4. Time-to-value and operational fit: Participants emphasized that successful adoption would depend on how quickly the platform could deliver tangible impact: "Time to value is the primary factor preventing long-term development." A forward-looking vision also emerged: "If you could make the SOC [Security Operations Center] work like a video game, that would be something valuable!" In addition, concerns were raised about distribution and delivery models: "Platform access … web-based, PC game … or console?"

In summary, participants described the platform as an "interesting concept" with "a lot of potential," while emphasizing the importance of "onboarding for non-gamers," "modular scenario authoring," "classroom controls," "content governance," "privacy safeguards," "clear delivery channels" and "rapid time-to-value." These practitioner priorities directly informed the core challenges and recommendations that follow.

Five Core Challenges to Realizing the Agentic Metaverse

Based on the recurring themes in the prototype feedback, we identified five interdependent challenges to realizing the agentic metaverse. These challenges highlight that success will depend not only on technical capability but also on robust governance, trust and sustained operational support.

Challenge 1. Addressing Infrastructure and Real-Time Performance Issues

The agentic metaverse requires substantial computing resources to render immersive 3D environments while simultaneously running and training multiple AI agents. The resource demand includes continuous graphical processing, high-bandwidth network connectivity and reliable cloud or edge resources to maintain low-latency interaction. One participant summarized the infrastructure challenge as follows: "You need to support [these systems] with hardware allocation, distribution and device management, as well as software support, debugging, content development … and comprehensive user support with reporting and analytics. Achieving this requires a strong infrastructure and dedicated organizational support."

Real-time synchronization and data management add another layer of difficulty. Agentic metaverse environments generate large volumes of live data, such as user interactions, agent decisions and environment updates, that must remain consistent across distributed instances. Participants also asked for more realism tied to live threat signals, which raises the bar on ingestion and synchronization: "Make the game 'real;' consider using a honeypot that is taking on real intrusions … it would add randomness to the intrusions and themes of attacks."

Device heterogeneity further complicates delivery. The agentic metaverse must perform across virtual reality headsets, augmented reality glasses, mobile devices and desktops. Input modalities, rendering capabilities and sensor stacks differ widely across these devices. Without careful abstraction, the experience would degrade or become costly to maintain. Cross-platform support often forces trade-offs between fidelity and reach, as emphasized by one participant: "Platform access: will it be web-based, PC game (Steam) or console (Sony, Xbox)?"

Finally, enterprise integration and interoperability are persistent issues. Organizations rely on diverse legacy systems and third-party tools with inconsistent application programming interfaces (APIs) and separate identity stores. Instructors also require operational controls that align with enterprise learning workflows, including instructor consoles for customizing learning and learning management system (LMS) integration: "A proctored [moderated] quest where you can start and stop the class and discuss progress would be helpful in a classroom setting."

Challenge 2. Providing Multi-Agent Scenario Design and Scalability

Designing realistic AI agents is a complex and multifaceted task. Scenario design must blend domain expertise with pedagogy to anticipate diverse learner profiles. Without joint modeling, simulations risk becoming "gamey," failing to capture the authenticity of real incidents and team dynamics. Participants highlighted the steep learning curve and the risk of over-gamification: "Complex [difficult] to understand the tool," and "Learners are likely already familiar with video games … progress and achievements encourage people to keep playing. Drawback: Some concepts could be difficult to teach with a 'game.'"

Several participants advocated moving beyond scripted narratives. They also emphasized the need for modularity to support flexibility and adaptability in content design: "I would like it to be more modular; make it something that practitioners can add their own scenario to, more of a platform."

Scale magnifies these issues. As the number of learners and AI agents increases, computing and coordination demand grow non-linearly. Traditional server architectures struggle when dozens of autonomous agents, such as attackers, defenders, mentors and analysts, must share state, negotiate goals and adapt to humans in real time. Assessment at this scale is equally complex. Conventional measures (e.g., quizzes) fail to capture the multi-agent dynamics that matter in cybersecurity, including the quality

of interactions, decision-making latency under pressure and a learner's ability to adapt to evolving threats. Robust evaluation requires mixed-method analytics that combine real-time agent and user telemetry with instructor observation and structured post-session debriefing.

Challenge 3. Ensuring Security and Privacy

The agentic metaverse introduces new attack surfaces and vulnerabilities that can expose sensitive assets at risk. Training materials, proprietary resources and personally identifiable information may be exposed through data breaches, identity spoofing or adversarial AI exploits. In addition, the agentic metaverse will require the continuous collection of highly sensitive data, including performance metrics, decision pathways, behavioral patterns and, in some cases, biometric cues. The combination of rich agent telemetry and real-time analytics raises concerns about unauthorized profiling and the secondary use of data. A participant voiced these concerns explicitly: "[The system has] high engagement and strong simulation potential. But there is a risk of such an AI-driven game consuming confidential information."

Challenge 4. Managing the Safety and Governance of Social Dynamics

The agentic metaverse depends on lifelike interactions, which can create risks, such as harassment, discrimination and unethical behavior. Managing real-time interactions and potential misconduct requires clear codes of conduct, active content moderation, and live and post hoc monitoring for reporting, reviewing and resolving incidents. In the words of one participant: "Who is going to monitor social interactions in these immersive worlds?"

Beyond social risks, content safety and ethics remain key challenges. AI agents trained on biased data or shaped by narrow developer perspectives can unintentionally reinforce systemic bias, unfairness in assessments or unequal learning experiences. Accessibility is another key concern. Some learners may experience motion sickness and eye strain, while neurodiverse users may be overwhelmed by high-stimuli environments or fast-paced agent interactions. Realism must be balanced with usability. High-fidelity simulations, rich visuals and continuously adaptive agents enhance authenticity, but they can also increase cognitive load and induce mental and even physiological fatigue, especially for non-gamers, as emphasized by participants: "[The tool could be] challenging to navigate" and "[It] might be quite a learning curve to use it for someone who does not do gaming."

Because these social, ethical and ergonomic issues span technology, policy and people, they require governance. In the agentic metaverse, governance must extend beyond regulatory compliance to function as a dynamic control system for immersive, high-stakes training. Organizations need to define who owns and maintains the environment, establish monitoring protocols for user and agent interactions, set health and safety protocols, and ensure transparency and explainability in AI decision-making. Deployments must also comply with regional and sector-specific regulations and standards, including the European General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA) and ISO 27001, and emerging metaverse-specific standards (e.g., IEEE P2048, IEEE P7016), to ensure lawful data collection, secure storage and ethical use of data across on-premises and cloud environments.

Challenge 5. Navigating Adoption and Change Management

The agentic metaverse, including XR hardware and software as well as AI capacity-building, requires substantial investment, often without immediate returns. Long-term sustainability adds further pressure, as ongoing platform maintenance, content updates and system expansion are essential. As indicated by one participant, pilots should be designed to demonstrate early impact and provide justification for continued investment: "Time to value is the primary factor preventing long-term development."

Organizations are also likely to face cross-disciplinary expertise gaps. Effective agentic metaverse programs will require coordination among AI developers, XR designers, cybersecurity professionals, 3D artists, pedagogy experts and instructional designers. In particular, small

and medium-sized enterprises may struggle to recruit or retain all of these roles, making it challenging to develop, integrate and sustain high-quality training. Rather than depending on external developers, organizations should pursue approaches such as partnered delivery models, reusable scenario frameworks, and internal capability building to maintain and extend content after launch. One participant stated that the extensibility of the platform may help address these challenges: "Make it something that practitioners can add their own scenario to, more of a platform."

Even with adequate resources, cultural resistance and lack of employee buy-in can stall progress. Concerns about automation, skepticism about immersive learning, generational differences in technology comfort and technostress can all reduce participation. Without clear communication emphasizing augmentation rather than replacement of human mentorship, employees may perceive the agentic metaverse environment as impersonal or risky. Addressing such adoption challenges requires internal champions. As one participant advised: "You need a technology champion who is going to carry most of the burden ... and documentation should be ready to help your champion convince leadership and employees." Effective change management would pair these champions with clearly defined roles and responsibilities, structured communication strategies and practical playbooks, such as onboarding, user interface and troubleshooting guides.

Practical Recommendations for Implementing the Agentic Metaverse for Workforce Development

Drawing from the recurring themes and the five challenges identified above, we provide six practical recommendations to guide practitioners in implementing the agentic metaverse for training and workforce development.

1. Build a Scalable and Interoperable Infrastructure

Organizations should prioritize the development of a scalable and interoperable infrastructure that blends edge computing with cloud platforms to meet resource-intensive rendering and multi-agent inference with low latency. They should containerize services and use orchestration (e.g., Kubernetes) alongside event-driven pipelines, edge caching and real-time observability (e.g., metrics, logs with clear service-level objectives [SLOs] and tracing) to keep experiences responsive under load. To reduce integration friction, organizations should publish interoperability guidelines and adopt standardized APIs and middleware so that the platform plugs cleanly into LMS, ERP and customer relationship management (CRM) software, and runs consistently across different platforms, such as VR headsets, desktop and mobile computers, through capability negotiation. Organizations should define latency budgets and device profiles up front to guide feature tiering and performance targets across heterogeneous hardware. Finally, common data and content schemas for events, assessments and telemetry should be agreed internally to support reliable synchronization, analytics and long-term maintainability.

2. Develop Realistic and Adaptive Multi-Agent Training Scenarios

Effective scenario design is essential for translating abstract knowledge into applied competencies, particularly in complex domains such as cybersecurity. Scenarios should reflect the sociotechnical dynamics, decision-making constraints and ethical dilemmas employees face in real-world workflows. Achieving this requires cross-functional teams to collaboratively design exercises that are both realistic and pedagogically sound. AI agents can simulate stakeholders (e.g., IT managers, HR) to increase authenticity or emulate adversarial actors using various tactics, such as techniques and procedures from frameworks like MITRE ATT&CK,[30] and enable dynamic branching based on learner decisions. Each scenario should be explicitly mapped to target competencies, ranging from procedural knowledge (e.g., incident reporting protocols) to leadership under uncertainty.

Scenario complexity should match learner roles and expertise levels. For example, an entry-level employee might engage with guided decision trees to recognize suspicious

30 For details, see *ATT&CK Matrix for Enterprise*, MITRE, available at https://attack.mitre.org/.

Table 1: Design Principles for Adaptive Cybersecurity Training Scenarios

Design Principle	Example Scenario	AI Capabilities
Realism via threat-based modeling	Simulated ransomware attack that begins with a phishing email and lateral movement	AI agents emulate attacker tactics, techniques and procedures from MITRE ATT&CK and adapt behavior based on learners' input
Competency-aligned progression	Tiered incident response scenarios (e.g., beginner: detect phishing; advanced: remediate breach)	AI adjusts difficulty in real time based on learner proficiency and decision patterns
Ethical and legal ambiguity	Whistleblower suspects data misuse and must choose escalation path	AI simulates HR, legal and managerial roles with differing perspectives
Real-time feedback and debriefing	Live alerts, logging of actions, automated post-incident reports	AI generates performance summaries and suggests best practices
Inclusivity and role adaptability	User plays both attacker and defender roles with varying technical depth	AI tailors content complexity and provides role-specific hints
Adaptive branching and narrative consequences	Choosing to ignore a vulnerability leads to simulated breach in later scenarios	AI tracks decisions and triggers downstream scenario consequences
Psychological safety	Trainees can opt out for sensitive content or request AI pause, rewind or restart scenario without penalty	AI provides checkpoint saves, non-judgmental feedback and coaching mode
Replayability and scenario comparison	Trainees can replay the same scenario to explore alternate paths (e.g., breach vs. containment)	AI logs and compares multiple runs and offers adaptive commentary and progress tracking

cybersecurity activity, whereas a manager might navigate multi-agent negotiations during a simulated crisis under time constraints. Scenario design must be role-sensitive, offering alternate perspectives (e.g., attacker vs. defender) to cultivate a more comprehensive understanding of cybersecurity dynamics. Table 1 presents a set of principles for designing cybersecurity training scenarios, along with illustrative examples and the role of AI in supporting scenario realism and adaptivity.

3. Embed Security and Privacy by Design

Organizations should adopt a security-by-design posture tailored to the agentic metaverse, XR technologies and multi-agent systems from day one. This typically means conducting threat modeling that includes agent behaviors and prompt-injection vectors and involves segmenting networks and isolating training datasets from production systems, enforcing encryption in transit and at rest, and aligning identity and access management with single sign-on (SSO) and least-privilege role-based or attribute-based access control (RBAC/ABAC) for both users and agents. Continuous system-wide monitoring should be embedded early using security information and event management (SIEM) systems. This should be supported by documented model and data provenance, comprehensive audit trails of agent and user actions, prompt interactions and regular red-teaming of agents. Adversarial machine learning testing should be employed to identify vulnerabilities, and human-in-the-loop safeguards, such as pause, rollback and override mechanisms, should be in place to mitigate risk during consequential AI-driven decisions.

Participants specifically warned about scenarios where "an AI-driven game consumes confidential information." In such scenarios, data loss prevention measures, strict environment sandboxing, and content-ingress and egress controls are essential to prevent the leakage of proprietary assets and personally identifiable

information. Privacy-by-design should guide data handling throughout the lifecycle. Organizations should minimize and pseudonymize telemetry, specify purposes up front, obtain informed consent where appropriate, set clear retention schedules and provide user controls (access, correction, deletion) as well as explanations for consequential AI feedback.

Complementing these security- and privacy-by-design controls, a compliance-by-design approach should tie them to applicable legal and assurance frameworks. This would include aligning with the GDPR, CCPA, ISO 27001 and the System and Organization Controls 2 (SOC 2) framework;[31] performing accessibility reviews; documenting cross-border data flow; and conducting vendor risk assessments before pilots. Teams should maintain a living control matrix that maps organizational policies and technical safeguards to specific compliance requirements. These controls should be validated at each stage, and reports, including incidents, audit findings and remediation actions, should be regularly reviewed to sustain trust as the agentic metaverse scales.

4. Embed Safety and Accessibility by Design

Accessibility and safety must be integral to the design process and treated as first-order constraints, with universal design, adjustable immersion and pacing, and non-immersive alternatives built in. To ensure inclusive participation, agentic metaverse implementations should align with the Americans with Disabilities Act (ADA), the European Accessibility Act (EAA), Web Content Accessibility Guidelines (WCAG) and universal-design principles. This means that implementations should provide multimodal access, such as screen-based interfaces and text or audio alternatives, together with adjustable immersion levels, flexible pacing and comfort settings tailored for motion-sensitive or neurodiverse learners.

Ergonomic safeguards, including motion-sickness thresholds, session-length limits with scheduled breaks and device-calibration guidance, should be implemented to promote well-being and support sustained engagement. These features should be explicitly defined and enforced as part of governance to reduce sensory loads, help prevent fatigue and improve learning outcomes across diverse user groups. Moreover, the accessibility and safety policies should be clearly documented and communicated.

Facilitators must be trained to take account of accessibility issues, and immersive environments should be instrumented with real-time monitoring and feedback loops. These practices will not only mitigate social and cognitive risks but also foster broader participation and build trust in the training experience.

5. Govern Social Dynamics and Oversight

Organizations must treat social interactions within immersive training environments as a governed space, not merely a configurable feature. Establishing clear codes of conduct, role-based moderation protocols and documented escalation procedures, including conflict mediation, incident reporting and decision review, should be defined before any pilot deployment. Governance mechanisms must encompass both real-time and retrospective oversight, supported by comprehensive audit trails of user and agent interactions, explainable AI feedback for consequential decisions, and human-in-the-loop controls to pause, override or revise scenarios as needed. To mitigate bias, organizations should combine content curation safeguards, such as diverse datasets and personas, with technical checks, including bias testing, red-teaming of agents and prompts, and procedural reviews, such as appeals for moderation or AI decisions. Clear ownership and accountability for these controls must be established across key organizational stakeholders, including those working in learning and development, information security, HR and compliance.

Governance should not be an afterthought or limited to compliance. Rather, it should function as a dynamic control layer that actively shapes the integrity, safety and inclusivity of immersive learning environments. Key governance decisions include: whether the environment will be closed (e.g., employees only) or open to mixed audiences; who owns and maintains the infrastructure and AI models; how interactions

31 A compliance framework that evaluates how service organizations manage and protect customer data.

Table 2: Recommendations, Challenges Addressed, and Key Actions

Recommendation	Challenges	Key Actions
Build a Scalable and Interoperable Infrastructure	C1, C5	• Define device profiles and latency budgets up front • Containerize services and orchestrate with Kubernetes • Standardize APIs/middleware with capability negotiation across platforms • Implement SSO/LMS/ERP/CRM integration • Set observability SLOs and real-time dashboards • Blend edge and cloud for low-latency rendering and agent inference
Develop Realistic and Adaptive Multi-Agent Training Scenarios	C2, C4	• Form a cross-functional authoring team • Map each scenario to a competency model • Offer role-sensitive variants and perspectives • Develop branching storytelling based on learner decisions • Use AI agents to simulate stakeholders and adversaries
Embed Security and Privacy by Design	C3, C4	• Run threat modeling, including agent behaviors and prompt-injection vectors • Encrypt data in transit and at rest; provide provenance tracking and audit trails • Align identity with SSO and least-privilege RBAC/ABAC • Employ SIEM-backed monitoring; adversarial machine learning tests; red-team agents/prompts • Implement human-in-the-loop fail-safes
Embed Safety and Accessibility by Design	C2, C4	• Provide accessibility options and non-immersive modes • Align with ADA/EAA/WCAG and universal-design principles • Enforce ergonomic safeguards; adjustable pacing, calibration guidance • Provide comfort/sensory controls for neurodiverse learners • Train facilitators on accessibility and safety protocols • Conduct real-time monitoring and feedback loops
Govern Social Dynamics and Oversight	C2, C4	• Publish code of conduct; define role-based moderation • Enable real-time and post hoc oversight (audit trails, explainability, human-in-the-loop controls) • Perform bias testing; curate diverse datasets and personas • Provide explainable AI for consequential decisions plus human-review/appeal mechanisms and override controls. • Make governance choices explicit (closed vs. mixed audiences, ownership)
Drive Adoption with Staged Rollout, Champions and Enablement	C1, C5	• Conduct pilots with measurable objectives and time-to-value milestones • Appoint internal champions; deliver train-the-trainer sessions • Provide guidelines (e.g., onboarding, troubleshooting) • Build a community of practice; capture feedback in backlog • Run change communications; showcase impact metrics

between users and agents are monitored and audited; how user-generated data is collected, stored and used; and how explainability and appeal mechanisms are provided for AI-generated outputs.

6. Drive Adoption with Staged Rollout, Champions and Enablement

Organizations should approach adoption of the agentic metaverse as a phased change management initiative rather than a one-time deployment, beginning with narrowly scoped pilots focused on specific roles or high-priority use cases. The pilots should be instrumented for before-and-after comparisons and aligned with measurable, high-value outcomes, such as skill acquisition, reduced onboarding time and improved incident response coordination. Communicating clear time-to-value milestones will help build credibility and secure stakeholder buy-in for broader scaling.

Successful adoption will also require visible executive sponsorship and robust enablement at multiple levels. Organizations should identify cross-functional champions within information security, learning and development and IT and equip them with practical resources, such as onboarding playbooks, facilitation guides, troubleshooting workflows and demo scripts. "Train-the-trainer" sessions can further build internal capacity and momentum. This enablement requires organization-wide support in the form of open office hours, dedicated help desk support, targeted change communications and thoughtful user experience accommodations to lower the barrier for non-gamers or less tech-savvy users. Early pilot participants should be treated as a community of practice whose feedback is used to shape product backlogs, refine success metrics and update governance protocols.

Mapping the Recommendations to the Identified Challenges

Table 2 maps our six recommendations to the five core challenges: (C1) Addressing infrastructure and real-time performance issues, (C2) providing multi-agent scenario design and scalability, (C3) ensuring security and privacy, (C4) managing the safety and governance of social dynamics, and (C5) navigating adoption and change management. It also summarizes the key actions required to effectively implement each recommendation.

Concluding Comments

Researchers and practitioners are continuing to explore the role and readiness of AI for deployment within the metaverse.[32] While the realization of a fully functional agentic metaverse will still require advancements in infrastructure, interoperability and governance frameworks, early prototypes are demonstrating its transformative potential. As these prototypes reveal the benefits of AI-integrated virtual environments, it is becoming increasingly clear how the agentic metaverse differs from traditional approaches.[33] Unlike conventional virtual environments that rely on pre-scripted scenarios, the agentic metaverse introduces autonomous agents capable of authentic role-playing, adaptive task complexity and real-time interaction with learners. This integration fundamentally enhances immersive learning by enabling personalized and context-sensitive experiences that are increasingly critical for modern workforce training and development.

The agentic metaverse is especially valuable in cybersecurity training because it can incorporate cooperative and competitive elements simultaneously. One AI agent may act as a defender, identifying and patching vulnerabilities, while another emulates an advanced persistent threat seeking to infiltrate systems. As trainees engage with these agents, the system provides immediate feedback and analytics on metrics such as escalation times, detection accuracy and decision quality. These data-driven insights allow trainees to identify their weaknesses, develop more efficient strategies and refine their skills through continuous iteration. Early adopters experimenting with metaverse-based "cyber ranges" that simulate real-world threat landscapes[34] report improved engagement and knowledge retention compared to traditional methods. By diversifying training experiences,

32 Dolata, M. and Schwabe, G., op. cit., February 2023.

33 Kamalov, F., Calonge, D. S., Smail, L., Azizov, D., Thadani, D. R., Kwong, T. and Atif, A. "Evolution of AI in Education: Agentic Workflows," *arXiv preprint arXiv:2504.20082*, April 2025.

34 Torro, O., Jalo, H. and Pirkkalainen, H. "Six Reasons Why Virtual Reality Is a Game-Changing Computing and Communication Platform for Organizations," *Communications of the ACM* (64:10), October 2021, pp. 48-55.

these initiatives can also help address burnout in cybersecurity roles, shifting practice away from repetitive drills toward engaging and skill-reinforcing exercises.

Looking ahead, the agentic metaverse may evolve into digital twins of real-world infrastructures, where trainees interact with AI agents that simulate real-time threats informed by live intelligence feeds and recent organizational events. However, as with any emerging technology, organizations should exercise caution before making significant investments in the agentic metaverse. Current enthusiasm surrounding the metaverse and AI may generate inflated expectations that may not align with practical realities or technological limitations.[35] Moreover, the convergence of AI and immersive environments introduces complex sociotechnical dynamics that can result in unexpected consequences, including social and ethical risks, bias, workplace inequalities, surveillance and user profiling. The challenges and recommendations presented in this article will support practitioners in navigating these complexities and inform the responsible deployment of agentic metaverse technologies for workforce training and development in an increasingly dynamic digital landscape.

About the Authors

Ersin Dincelli

Ersin Dincelli (ersin.dincelli@ucdenver.edu) is an associate professor of information systems at the University of Colorado Denver Business School. He holds a Ph.D. in informatics from the University at Albany, State University of New York. His research concerns the behavioral aspects of information security and human-computer interaction, focusing on decision-making processes in cybersecurity and design for immersive systems. His work has been funded by the National Science Foundation and the National Security Agency. He has received several prestigious honors, most notably the AIS Outstanding Contribution to Education Award, the Distinguished Dissertation Award and recognition as a TIAA Chancellor's Urban Engaged Scholar.

Haadi Jafarian

Haadi Jafarian (haadi.jafarian@ucdenver.edu) is an associate professor of computer science at the University of Colorado Denver's College of Engineering, Design and Computing. He holds a Ph.D. in computing and information systems from the University of North Carolina Charlotte. His work, supported by the National Science Foundation, Department of Defense, Department of Education and Colorado Department of Transportation, integrates formal methods, adversarial modeling and AI-driven analytics to develop proactive security solutions for critical cyber infrastructures. He is also deeply involved in cybersecurity education, exemplified by his role in developing the university's bachelor of science cybersecurity program.

35 Krüger, K., Weking, J., Fielt, E., Böttcher, T., Kowalkiewicz, M. and Krcmar, H. "Value Drivers for Metaverse Business Models: A Complementor Perspective," *Journal of Management Information Systems* (42:1), February 2025, pp. 143-173.