

Emerging Trends, Challenges, and Future Directions in Cybersecurity

Vishvesh Revoori¹; Vasudev Karthik Ravindran²; Shubham Agarwal³

¹Senior Software Development Engineer, Amazon Austin, Texas, USA.

²Senior Security Architect, Amazon, Seattle, Washington, USA,

³Information Security Consultant, CIPL, Noida, India

Publication Date: 2026/08/05

Abstract: The dual-use nature of AI, the expanding attack surfaces of cloud computing and the IoT, the evolution of Zero Trust architectures, and the forthcoming disruption of quantum computing to classical public key cryptography are causing rapid, structural shifts in the cybersecurity landscape. This study integrates and builds upon relevant literature of the aforementioned technologies and organizes the AI focused offense and defense, identity-based security, ransomware, cloud-native security, post-quantum cryptography (PQC), supply chain security, IoT/OT security, XDR platform security, and other security frameworks. We review and analyze the gaps that reduce the efficacy of security and defense postures, which include the labor gap in cybersecurity, inconsistent laws and frameworks, aging technologies, the complexity of AI and its governance, and alert fatigue from over-utilization of multiple tools. We identify autonomous security and defense, agile cryptography, AI red team exercises, human-AI defense, and a variety of models using graph theory and network theory to design and implement resilient architectures to mitigate cyber threats, as potential areas for future exploration and research. We aim to provide researchers and practitioners a comprehensive overview of the current cybersecurity environment for the year 2026.

Keywords: Cybersecurity, Artificial Intelligence, Zero Trust, Ransomware, Post-Quantum Cryptography, Cloud Security, IoT Security, Survey.

How to Cite: Vishvesh Revoori; Vasudev Karthik Ravindran; Shubham Agarwal (2026) Emerging Trends, Challenges, and Future Directions in Cybersecurity. *International Journal of Innovative Science and Research Technology*, 11(8), 1-8. <https://doi.org/10.38124/ijisrt/26aug037>

I. INTRODUCTION

Cybersecurity has changed from being a particular IT concern to a cross-industry, board-level priority. Ubiquitous cloud adoption, hybrid and remote work, and the proliferation of connected devices have all created a wider attack surface. The inclusion of generative and agentic AI in the workplace has created new attack and defense tools. Global spending on IT security continues to grow due to the prediction end-user spending on IT security would reach an estimated 240 billion by 2026, driven mainly by the integration of AI in the tools and a platform consolidation [1].

The same tools that have been created to protect against the expanding threat of AI, have now expanded that threat. Security professionals are reporting more and more incidents where AI is used to carry out phishing and other fraudulent social engineering attacks in the form of automated recon-naissance, phishing AI models, rapid variant malware attacks, and deepfake impersonations [2]. The attacks that have been reported are mainly automated AI attacks that carry out attacks with little to no human oversight [3].

This paper surveys the state of the art in cybersecurity as it stands in 2026, around the three central questions detailed below:

- What dominant technological and adversarial 'trends' impacting offense and defense do we see in the near future?
- What are the structural, technical, and organizational 'challenges' to the successful adoption of new defensive paradigms?
- In both theoretical and practical terms, what 'future research directions' are most relevant and promising?

The rest of the paper is detailed below. Section II explains the survey methods and reviews the background and evolution of threat type. Section III provides a framework of current cybersecurity trends. In Section IV the convergence of multiple trends is illustrated by a case study. Section V discusses the principal challenges facing researchers and practitioners. Section VI outlines future research directions. Section VII concludes the paper.

➤ *Survey Methodology*

For this survey we have employed a narrative-synthesis method. Our sources are drawn from two main streams: on the one hand, peer-reviewed work in the fields of post-quantum cryptography, IoT/OT and cloud security, zero-trust models and AI-driven security; on the other, the latest standards and threat-intelligence reports put out by industry bodies such as NIST, Gartner, Kaspersky Securelist and IBM X-Force. The latter are included to account for rapid developments that have not yet made their way into the academic record.

Such a mixed approach is called for in any review of *current* trends given that the pace of cybersecurity threats far outstrips the academic publication cycle, so industry reporting is an indispensable supplement to the literature. In choosing our material we have been guided by its relevance to the taxonomy of trends set out in Section III and, where feasible, have cross-referenced with a minimum of two independent sources.

II. BACKGROUND: EVOLUTION OF THE THREAT LANDSCAPE

Historically, perimeter-based defense was the focus of most cybersecurity strategies, consisting of a trusted internal network protected by firewalls, IDS, and network segmentation.

As networks transitioned to cloud computing and as organizations began adopting remote and hybrid work models, the perimeter-based model began to fade [4]. These types of computing and work environments offer a single defensible perimeter. As the perimeter-based models began to fade, the security community increasingly began to integrate machine learning. Enhancements to security operations made by machine learning included anomaly detection, malware classification, and user-behavior analytics [5], [6].

This began to change with the advent of large language models and agentic artificial intelligence with the ability to reason and complete tasks autonomously. This introduced a truly dual-use system. AI systems that allow security operations to handle alerts at a large scale also allow adversaries to automate reconnaissance and craft social engineering tasks at the same scale [7]. It is estimated that most organizations using security stacks have begun to incorporate generative AI and large language models. A significant number of these same organizations have begun to incorporate agentic AI to perform detection, triage, and containment tasks either autonomously or with minimal oversight [3].

➤ *Threat Actor Landscape*

Regardless of motivation and capability, modern cyber threats involve diverse actors. These threats may come from sophisticated criminals using professional RaaS and initial-access-broker services, advanced persistent threats used by nation-states for espionage and critical infrastructure attacks, or even affiliates making use of available AI tools to per-

form attacks requiring high technical knowledge [8]. Offense AI tools and ransomware have become available and it is making fewer people skilled to perform high-level cyber-attacks. This is known as the “democratization” of cybercrime.

III. LATEST TRENDS IN CYBERSECURITY

The taxonomy in Figure 1 takes a look at eight categories of trends. These are organized into three major themes where multiple trends are present at the same time: most notable, adversarial technology (AI-driven offense, ransomware evolution, deepfakes), and defensive architecture (zero trust, cloud-native security, platform consolidation), and disruptive infrastructures (post-quantum cryptography, IoT/OT and supply-chain security). A prime example of the intersection of multiple trends is AI-based orchestration of Ransomware.

➤ *AI-Driven Offense and Defense*

AI-Enhanced Cybersecurity has become an issue for both Cybersecurity professionals and Cyber adversaries. On the defense side, platforms engineered with AI are employed to streamline analysis for faster detection of irregularities, as well as automating incident response to a level that surpasses the constraints of analysis done by a person [9]. The areas of reported highest influence from AI are the detection of irregularities and new types of threats, as well as the automation of response and containment, and the lessening of the volume of alerts for analysts [3].

On the offense side, AI is adopted to carry out Cyber operations that conduct phishing schemes that are hyper-personalized, exploit chains for automation and mutating malware to escape signature detection, as well as deepfake audio and video for business email scams and executive email fraud [3]. Although not the first conduct of Cyber operations to use AI, a noteworthy advancement was the establishment of almost totally self-sufficient Cyber operations, “agentic” ransomware. Unlike previous Cyber operations that employed AI assistance, this new, more self-sufficient form of ransomware conducts the entire process of a ransomware scheme with very little human involvement.

Researchers have seen a resurgence of interest in adversarial machine learning as a cornerstone of cybersecurity. This subfield is concerned with everything from the evasion and poisoning of ML models employed for detection to the weaponization of generative models for the creation of new attack artifacts [7]. There is an inherent dual-use quality to such capabilities; progress on the defensive side, say in the form of better anomaly-detection architectures, will often be mirrored by an offensive equivalent. In this way, one sees a co-evolutionary dynamic at work between the attacker and the defender that is accelerating and of a different order than the more leisurely arms races of the signature-based malware era [6].

➤ *Zero Trust Architecture and Identity-Centric Security*

As set out in NIST Special Publication 800-207, Zero Trust Architecture (ZTA) does away with the notion of

implicit trust derived from one's position on the network. In its place is a mandate for the ongoing verification of each access request, taking into account identity, the state of the device and any contextual risk [10]. This move to an identity-centric approach to security has been necessitated by the fact that credential abuse and other identity-based vectors have become the go-to method for intruders [9].

There is also a growing preference for Continuous Exposure Management (CEM) programmes that put together adaptive MFA, passkeys and continuous risk scoring. Analysts would have it that those organisations which make the switch to CEM are far less prone to a breach [9]. One need only look at the wider body of zero-

trust work to see the same trend: a departure from static, perimeter-bound models in favour of authorisation that is both context-aware and unrelenting [11].

➤ Ransomware Evolution and Encryptionless Extortion

In the realm of cyberattacks, ransomware is still among the most financially ruinous and hard to put behind one [12]. Yet the way it is run is changing. With more victim organizations refusing to part with their money and the proportion of ransoms that are paid dropping to well below a third of all cases, threat actors have been forced to adapt. They are turning to what is known as encryptionless extortion: exfiltrating data for leverage without the forensic trail or overhead that comes with encrypting files [13].




Fig 1 Taxonomy of Latest Cybersecurity Trends, Grouped into Adversarial-Technology, Defensive-Architecture, and Infrastructural-Transition Themes.

There is also evidence of adversaries being just as quick as defenders to make use of quantum-resistant cryptography. A number of ransomware families have started employing post-quantum-secure key-exchange primitives like Kyber (NIST's ML-KEM) in their key-negotiation and command-and-control protocols [13], [14]. On the supply chain side, initial access brokers are as vital an intermediary as ever, with remote-access services serving as the entry point of choice [13].

➤ Cloud-Native and Multi-Cloud Security

As workloads make their way into hybrid and multi-cloud settings, the attack surface has grown in direct proportion. Misconfigured storage and the theft of credentials are still the usual means of gaining unauthorized entry [1], [15]. To counter this, enterprise security architecture now relies heavily on cloud-native application protection platforms (CNAPPs) to bring together identity governance, workload protection and posture management from various providers. It is not an easy task; the very nature of these heterogeneous environments, along with the lack of uniform configuration standards, makes it difficult to keep a unified view of things across providers [15].

➤ Post-Quantum Cryptography Migration

Though quantum computers are still not available to break popular public-key cryptography algorithms like RSA and ECC, the transition schedule for replacing the cryptographic infrastructure across the globe is already scheduled for years and even decades [16]. It is precisely this kind of imbalance that explains the concept of "harvest now, decrypt later," meaning that adversaries, including state-sponsored ones, currently harvest and keep encrypted data so that it can be decrypted when CRQs become available [17]. The NIST published its first three post-quantum standards in the form of FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), and FIPS 205 (SLH-DSA), and as a result of further policy and regulation, migration deadlines were set starting from 2027, with full compliance to be reached in the early 2030s [14]. Hybrid approaches utilizing both classical and post-quantum primitives such as combining X25519 with Kyber/ML-KEM have already been adopted in TLS 1.3 and QUIC by technology vendors [13].

➤ IoT, OT, and Supply-Chain Security

Manufacturing, energy, healthcare and the transportation sector are all seeing a steady increase in the deployment of Internet-of-Things (IoT) and operational-technology (OT).

This growth has put devices on the field that lack the computational heft for robust cryptography and are subject to uneven patching, rendering them convenient pivot points for lateral movement [18].

At the same time there has been a sharp rise in software supply-chain attacks. By way of a compromised vendor or upstream build pipeline, an attacker can make his way to numerous downstream victims; it is estimated that such third-party and supply-chain breaches have quadrupled in the last five years alone [8], [19]. In response, the industry is turning its attention to mitigations like reproducible builds, dependency provenance checks and software bills of materials (SBOMs). The situation in OT is of particular concern. The push for predictive-maintenance analytics and remote monitoring has brought enterprise IT and what were once siloed industrial control systems (ICS) into closer contact. An initial breach on the IT side now has new avenues to spread to physical processes where safety is paramount. Dealing with the realities of these environments – long lifecycles, hard-to-patch equipment and limited processing power – is not the same as standard enterprise IT work. As a result, there is a fresh emphasis on network segmentation, anomaly detection for industrial protocol traffic and the kind of lightweight cryptography that a resource-starved device can handle [18].

➤ *Deepfakes and AI-Enabled Social Engineering*

The ability of Generative AI to produce authentic-seeming synthetic audio, video and text, markedly at a lower cost than before, creates new challenges for security professionals. These technologies enable deepfake voice and video fraud, executive impersonation, and highly customized phishing attacks. These attacks can now be performed to a level of fidelity and on a scale that was not possible previously [3], [20]. Highly customized phishing attacks are now considered the leading concern for security professionals. These attacks are now more concerning than adaptive malware and automated exploit chaining [3].

➤ *Security Platform Consolidation and XDR*

The challenges of tool sprawl and alert fatigue mean that more organizations are consolidating purchases to integrated security solutions. This trend is expected to continue into 2026 [3]. Extended Detection and Response (XDR) solutions are intended to address this trend. These solutions integrate security telemetry from email, network, endpoint, cloud and identity. XDR solutions attempt to address the limitations of other security solutions by employing integrated security telemetry to identify attack patterns that cross multiple security domains [3], [9].

➤ *Regulatory and Policy Drivers*

Cybersecurity trends have increasingly become influenced by regulatory and policy changes rather than technological or market considerations. Examples of regulations that will drive cybersecurity practices include the Network and Information Security Directive (NIS2) and Cyber Resilience Act (CRA) within the EU, specialized frameworks such as DORA (Digital Operational Resilience

Act) for financial institutions, and changing breach notification rules, which will result in stricter supply chain due diligence and incident reporting, and security by design for connected products. On the other hand, post-quantum readiness at the national level, with mandatory migration deadlines for government and national security systems, is increasingly influencing the private sector via supplier and vendor compliance requirements [17].

IV. CASE STUDY: CONVERGENCE OF TRENDS IN AGENTIC RANSOMWARE

To highlight how the trends outlined in Section III inter-act in the real world, consider the development of agentic ransomware campaigns in 2026, where an autonomous AI agent carried out a full attack lifecycle from gaining initial access all the way through file encryption with minimal human intervention [3]. This example demonstrates convergent developments across three of the branches in our taxonomy illustrated in Fig. 1:

- **Adversarial technology:** The agent used automated reconnaissance and exploit chaining along with adaptive and self-modifying malware behavior, as part of the larger trend of AI-enabled attack chains [3].
- **Infrastructural transition:** The malware's built-in key exchange protocol reportedly employed a lattice-based key encapsulation scheme similar to the NIST's ML-KEM protocol; this is a clear demonstration of the threat actors' early adoption of post-quantum secure cryptography [13], [14].
- **Defensive architecture:** Ultimately, the detection required using cloud workload telemetry and behavioral analysis, which highlights the importance of cloud native and XDR architectures outlined in Section III-D and III-G.

This convergence is representative of a broader research methodology issue in general: studying individual trends in isolation greatly understates the danger of their interaction, and defense research should take into account compound attack patterns rather than addressing each trend independently.

V. CHALLENGES

In spite of considerable expenditure of money and development of new techniques, various fundamental difficulties have remained barriers to good practice in cyber-security. Table 2 provides an overview of the key problems together with a basic cause and possible solutions. These are discussed below.

➤ *Workforce and Skills Gap*

The chronic insufficiency of skilled protection experts limits the ability of organizations to use, improve, and understand more complex tools. The problem is aggravated by the complexity of multi-cloud and hybrid environments that require deep knowledge in areas like networking, identity, cloud design, and possibly even AI/ML systems.

➤ *Regulatory Fragmentation and Compliance Burden*

Organizations that operate in many jurisdictions must comply with a growing number of complicated and unclear laws related to data protection, data breach notification, and specific sectors. Hence, the process of complying with the rules takes time, especially for small firms that do not have a compliance department.

➤ *Legacy Systems and Technical Debt*

Critical infrastructure, financial institutions, and other enterprises use a lot of technology that cannot be easily up-dated, replaced, or combined with modern zero-trust and post-quantum technologies, due to the requirement for the complete change of the architecture rather than just installation of new software. Standardization of all processes falls behind the time, which leads to embedding new algorithms into legacy software becoming even more difficult. [17].

➤ *AI Governance and Dual-Use Risk*

The same generative and agentic AI powers used to enhance defense also reduce the barrier to entry for

sophisticated attackers. While “cryptographic agility” — the ability to swap out crypto primitives without massive amounts of redesign— is something of Holy Grail for most companies, NIST’s “10+ years” to transition are actually quite ambitious, given that there is a (non-zero but currently unknowable) chance of timely relevant breakthrough discoveries [3, 5]. There remains significant room between vendor marketing and real, cross-domain AI capability to hinder any effective vendor selection process [3].

➤ *Post-Quantum Migration Difficulty*

While NIST has released its initial PQC standards, migrating away from today’s crypto is hampered by the enormous size of crypto dependencies throughout enterprise software, hardware and protocol standards. “Cryptographic Agility,” which is defined as the capability to change out crypto primitives without rewriting large portions of system design, remains elusive for many organizations, and NIST’s 10+ year migration period is likely to be shorter due to the non-zero, but unknown likelihood of relevant quantum breakthroughs [16], [17].

Table 1 Summary of Latest Cybersecurity Trends

Trend	Description	Representative Driver
AI-driven offense/defense	Dual-use AI for automated attack chains and AI-augmented SOC operations	Agentic ransomware, AI-SOC platforms [3]
Zero Trust / identity-centric security	Continuous verification of every access request; no implicit network trust	Credential abuse, CEM adoption [9], [10]
Ransomware evolution	Shift toward encryptionless extortion; adoption of PQC in malware	Falling ransom payment rates [13]
Cloud-native security	Unified posture and workload protection across multi-cloud environments	Misconfiguration, credential theft [15]
Post-quantum cryptography	Migration to ML-KEM, ML-DSA, SLH-DSA and hybrid schemes	Harvest-now-decrypt-later, NIST FIPS 203–205 [14]
IoT/OT and supply chain	Securing constrained devices and upstream software dependencies	Quadrupling of supply-chain breaches [8]
Deepfakes / social engineering	Synthetic audio/video used for impersonation and fraud	Generative AI accessibility [20]
Platform consolidation (XDR)	Unifying telemetry across domains to reduce alert fatigue	Tool sprawl, cross-domain visibility gaps [3]

Table 2 Principal Challenges, Underlying Causes, and Potential Mitigations

Challenge	Underlying Cause	Potential Mitigation Direction
Workforce/skills gap	Rapid growth in tooling complexity outpacing training pipelines	AI-assisted analyst augmentation; curriculum modernization
Regulatory fragmentation	Inconsistent multi-jurisdictional data-protection regimes	Harmonized compliance frameworks; automated compliance tooling
Legacy technical debt	Long-lived critical infrastructure incompatible with modern controls	Incremental cryptographic agility; hybrid legacy/modern gateways
AI governance & dual-use risk	Same AI capabilities usable for attack and defense	AI red-teaming; secure model pipelines; usage policy
PQC migration complexity	Deep embedding of classical cryptography across stacks	Crypto-agile architecture; phased hybrid deployment
Privacy–security tension	Monitoring required for detection conflicts with data minimization	Privacy-preserving analytics; federated learning
Metrics/evaluation gaps	Inconsistent, vendor-driven effectiveness reporting	Open, reproducible security-AI benchmarks
Alert fatigue/tool sprawl	Fragmented, siloed point solutions	Platform consolidation (XDR); alert-prioritization ML

➤ *The Privacy-Security Conflict*

More sophisticated behavioral analysis, identification/validation tools, and surveillance techniques necessitated by robust security often collide with privacy demands and data minimization goals, requiring sensitive architectural and policy choices.

➤ *Metrics and Evaluation Challenges*

Metrics and Evaluation Challenges. A common, persistent technical challenge throughout all the trend areas discussed in this review is the absence of consistent and comparable measures for evaluating security posture and the effectiveness of tools. Detection rates, MTDD, and MTTR provided by vendors are typically calculated using highly varying parameters and testing methodologies, hindering inter-vendor and inter-organization comparison.

This challenge is especially significant with the advent of security-AI solutions in which platform offerings often claim cross-domain coverage that doesn't necessarily materialise in practice [3].

➤ *Alert Fatigue and Tool Sprawl*

The development of openly available, reproducible benchmarking datasets that will support the evaluation of security AI will be an important next step for the research community. Alert Fatigue and Tool Sprawl. The trend towards increasing numbers of siloed, point security products has resulted in the poor security posture visibility, high number of low fidelity alerts, leading to analyst alert-fatigue and missed attacks. Some consolidation efforts, which are addressed in this review, offer a potential path to addressing alert-fatigue and improving security visibility [3].

VI. FUTURE SCOPE AND RESEARCH DIRECTIONS

➤ *Autonomous and Human-Supervised Security Operations*

We expect future operations centres to adopt a hybrid model: leveraging autonomous agents for triage, correlation, and containment, with humans reserved for high-consequence decision-making. Challenges of this future model include defining appropriate degrees of autonomy, producing verifiable audit trails of AI-assisted security operations, and engineering secure and meaningful human control and escalation pathways.

➤ *Cryptographic Agility and Post-Quantum Deployment*

Embedding cryptographic agility in system design to permit easy substitution of cryptographic primitives is an extensive, and largely unresolved, engineering research problem, especially in the context of resource-constrained IoT/OT devices and legacy systems used by governments and the financial sector [16]. We also need to investigate hybrid classical-post quantum performance and interoperability at scale through the internet.

➤ *AI Red-Teaming and Secure Model Pipelines*

The use of private AI models in organisations' security-sensitive work-flows means that red-teaming of

models, evaluation of adversarial robustness, and secure model pipelines are a quickly expanding field, including such concerns as prompt injection, data poisoning and model extraction risks [1].

Graph- and Network-Theoretic Approaches to Threat Propagation Attacks and other disruptions to an enterprise IT network such as supply-chain disruptions, compromise, and lateral movement are likely to be propagated through interconnected nodes of system. Graph neural networks and network-science approaches, including centrality-based threat prioritization, analysis of network structure (e.g., small-world and scale-free networks), and dynamic modeling of communication between nodes on networks may be used to model, predict, and mitigate distributed intrusions and malware propagation across enterprise networks and IoT networks in a manner similar to epidemic modeling on networks.

Federated and Privacy-Preserving Security Analytics Federated machine learning and other privacy-preserving security analytic approaches have the potential to enable collaborative development of threat detection and prediction systems without the need to pool sensitive log data, which can be a critical component of enhanced security posture with an respect to data sovereignty and privacy [21]. Another important research thread is to develop mechanisms that effectively meld the use of pattern recognition capabilities of machine intelligence with the judgment of human security professionals for situations with new threats, those with considerable ambiguity and for which autonomous systems may carry excessive risks.

➤ *Explainable and Auditable Security AI*

With Increasing AI autonomy in security functions, explainability shifts from being a "nice to have" feature to a required "must have" in security operations from trust, regulatory compliance, to incident forensics. We need further research into the explainability of anomaly detection; verifiable logging for automated containment action; and the use of standard formats for auditing of AI decisions in the security domain, ensuring increased automation doesn't reduce accountability.

➤ *Quantum-Safe Transition Roadmapping*

In addition to the cryptographic agility research discussed previously, formal quantitative road-mapping methods - going beyond Mosca's risk-timeline framework [16] - will be required to enable organisations to make rational decisions regarding system prioritisation for quantum-safe migration. The prioritisation will depend on the sensitivity of and data life time confidentiality required for, the information residing in the organisation's systems, practical concerns about the feasibility of migration in the short term and organisations' exposure to harvest-now-decrypt-later threats.

There remains significant scope to conduct empirical benchmarks of hybrid post-quantum schemes on live networks.

➤ Resilience Engineering for Compound Attacks

As shown by the case-study detailed within Section IV, attacks in the future are likely to leverage multiple threat trend vectors in unison (for instance, AI driven scanning, post-quantum authenticated communications, and supply chain driven initial compromise).

Resilience engineering to model and assess defences to compound attacks rather than individually isolated controls is another key area for future, and to some extent, as-yet underdeveloped research.

VII. CONCLUSION

The security landscape in 2026 features an increasing dual-use trajectory for AI, an ongoing architectural transition to identity-based, zero-trust security, a nascent but increasingly serious transition to post-quantum cryptography, and the ever-growing attack surface enabled by interdependencies within cloud, IoT, and software supply chains. Although significant advances in security defenses, from AI in security operations to platform convergence and the initial deployment of post-quantum defenses, are occurring, security organizations remain hampered by pervasive structural constraints in terms of skills shortages, fractured regulation, technical legacy, and the novel AI-related governance challenges. These challenges will need to be addressed through sustained R&D into, among other areas, autonomous security operations, cryptographic agility, AI red-teaming, and network-theoretic modeling of threat propagation, combined with ongoing collaboration among industry, standards organizations, and the academic research community.

This case study highlights an important lesson of more general methodological significance: increasingly impactful emerging threats are inherently compound in character, blending AI automation, evolving cryptographic abilities, and widening infrastructural attack surfaces in ways that can be effectively countered only by tightly integrated, systems-wide defenses, not by siloed controls.

Future work will need to conceive of cybersecurity as a socio-technical system rather than as an array of technical issues. This understanding, both data-driven and based on formal modeling, may prove to be the critical long-term project of cybersecurity research.

REFERENCES

- [1]. Fortinet Inc., “Cybersecurity trends 2026: Defending against agentic & ai threats,” <https://www.fortinet.com/resources/cyberglossary/cybersecurity-trends-2026>, 2026.
- [2]. TierPoint LLC, “Cybersecurity trends in 2026: Rising threats strategies,” <https://www.tierpoint.com/blog/cybersecurity/cybersecurity-trends/>, 2026.
- [3]. Kiteworks, “Ai cybersecurity in 2026: Key trends threats,” <https://www.kiteworks.com/cybersecurity-risk-management/ai-cybersecurity-2026-trends-report/>, 2026.
- [4]. D. C. Diningrat, B. Rahardjo *et al.*, “Security issues in multi-cloud: A systematic literature review,” *IEEE Access*, 2025.
- [5]. K. Shaukat, S. Luo, V. Varadharajan, I. A. Hameed, and M. Xu, “A survey on machine learning techniques for cyber security in the last decade,” *IEEE Access*, vol. 8, pp. 222 310–222 354, 2020.
- [6]. I. H. Sarker, M. H. Furhad, and R. Nowrozy, “Ai-driven cybersecurity: An overview, security intelligence modeling and research directions,” *SN Computer Science*, vol. 2, no. 3, pp. 1–18, 2021.
- [7]. N. Kaloudi and J. Li, “The ai-based cyber threat landscape: A survey,” *ACM Computing Surveys*, vol. 53, no. 1, pp. 1–34, 2020.
- [8]. IBM Security, “X-Force Threat Intelligence Index 2026,” IBM Corporation, Tech. Rep., 2022.
- [9]. Seceon Inc., “Cybersecurity trends 2026: Top security challenges, ai threats,” <https://seceon.com/cybersecurity-trends-2026/>, 2026.
- [10]. S. Rose, O. Borchert, S. Mitchell, and S. Connelly, “Zero trust architecture,” NIST Special Publication 800-207, Tech. Rep., 2020.
- [11]. X. Yan and H. Wang, “Survey on zero-trust network security,” in *International Conference on Artificial Intelligence and Security*. Springer, 2020, pp. 50–60.
- [12]. M. Humayun, M. Niazi, N. Z. Jhanjhi, M. Alshayeb, and S. Mahmood, “Cyber security threats and vulnerabilities: A systematic mapping study,” *Arabian Journal for Science and Engineering*, vol. 45, no. 4, pp. 3171–3189, 2020.
- [13]. Kaspersky Securelist, “Reviewing the trends in ransomware attacks in 2026,” <https://securelist.com/state-of-ransomware-in-2026/119761/>, 2026.
- [14]. G. Alagic, Q. Dang, D. Moody, A. Robinson, H. Silberg, D. Smith-Tone *et al.*, “Module-lattice-based key-encapsulation mechanism standard,” 2024.
- [15]. B. Alouffi, M. Hasnain, A. Alharbi, W. Alosaimi, H. Alyami, and M. Ayaz, “A survey of multi-cloud security issues, challenges, and solutions,” *IEEE Access*, vol. 9, pp. 57 792–57 807, 2021.
- [16]. M. Mosca, “Cybersecurity in an era with quantum computers: Will we be ready?” *IEEE Security & Privacy*, vol. 16, no. 5, pp. 38–41, 2018.
- [17]. R Street Institute, “Post-quantum cryptography migration in the united states: Managing risk and advancing cyber readiness in critical infrastructure,” <https://www.rstreet.org>, 2026.
- [18]. F. Alwahedi, A. Aldhaheri, M. A. Ferrag, A. Battah, and N. Tihanyi, “Machine learning techniques for iot security: Current research and future vision with generative ai and large language models,” *Internet of Things and Cyber-Physical Systems*, vol. 4, pp. 167–185, 2025.
- [19]. C. Okafor, T. R. Schorlemmer, S. Torres-Arias, and J. C. Davis, “Sok: Analysis of software supply chain security by establishing secure design properties,” in *Proceedings of the 2022 ACM Workshop on Software*

Supply Chain Offensive Research and Ecosystem Defenses, 2022, pp. 15–24.

- [20]. P. Korshunov and S. Marcel, “Deepfakes: A new threat to face recognition? assessment and detection,” *arXiv preprint arXiv:1812.08685*, 2018.
- [21]. T. Z. Sana, S. Abdulla, A. Das, A. Nag, M. M. Hassan, Z. Z. Fiza, Karim, and S. R. R. Kabir, “Advancing federated learning: A systematic literature review of methods, challenges, and applications,” *IEEE Access*, 2025.