INFORMATION SCIENCE

Critical Media Literacy, Information Ecosystems, and Institutional Cyber-Resilience

Maryam Adamu Muhammed[1], Rabiu Kera[1], Hamza Suleiman[2]

[1]Department of Film and Multimedia Studies, Nuhu Bamalli Polytechnic, Zaria, Nigeria
[2]Department of Strategic Communication and Media Studies, Nuhu Bamalli Polytechnic, Zaria, Nigeria

Published November 2026

Abstract

Democratic institutions today face growing pressure from synthetic information, algorithm-driven polarization, and organized disinformation campaigns. This study examines how critical media literacy, when built into wider information ecosystems, can strengthen the cyber-resilience of institutions. The central problem is that technical perimeter defenses often fail when attackers target human judgment instead of software. To address this problem, the study offers an empirical framework for assessing human-centered threat reduction. It uses a mixed-methods design. The quantitative part analyzes panel data from 52 public and private organizations between 2021 and 2025. The qualitative part uses process tracing to follow cognitive security incidents from start to finish. The econometric results show that investment in critical media literacy and continuous information verification has a significant positive effect on institutional cyber-resilience ($\beta = +0.51$, $p < 0.001$). The same investment is linked to a 38% fall in susceptibility to socio-technical attacks. The findings also show that technical controls, such as firewalls and zero-trust architecture, are necessary but not sufficient. They work best when paired with cognitive security practices that can recognize manipulated information. For policy and practice, the study proposes a roadmap with four parts: mandatory socio-technical risk audits, cognitive security training, decentralized verification of information, and threat intelligence sharing across sectors. By joining human cognitive defense with enterprise cybersecurity, public and private organizations can keep operating even when their information environment is compromised.

Keywords: Critical media literacy; information ecosystems; institutional cyber-resilience; cognitive security; disinformation; socio-technical systems

1 Introduction

1.1 Background of the study

The way organizations protect themselves has changed in a basic way. In the past, institutional cybersecurity focused mainly on hard technical borders, network encryption, and the management of software weaknesses. Today, digital information ecosystems are shaped by algorithmic amplification, generative synthetic media, and real-time social networks. These features have made old security boundaries porous, because false or manipulated content can now enter an organization through the same channels that staff use every day (Starbird *et al.,* 2023; European Union Agency for Cybersecurity [ENISA], 2024).

Threat actors have adapted to this change. Many of them now aim at human thinking rather than software code. They use persuasive disinformation, impersonation, and emotional pressure to damage decision-making and weaken public trust. The evidence from breach investigations supports this concern. The Verizon *Data Breach Investigations Report* found that a non-malicious human ele-

ment was involved in about 68% of breaches in its 2024 edition (Verizon, 2024) and in about 60% of breaches in its 2025 edition (Verizon, 2025). Although the exact share changes from year to year, people remain a central part of how attacks succeed.

In this fast-moving setting, institutional cyber-resilience can no longer depend on technical compliance alone. It must also protect the cognitive side of security. Critical media literacy is the ability to analyze, evaluate, and verify media messages. When it is practiced inside organizational routines, it becomes a first line of defense against cognitive attacks (Hobbs, 2020; Guess *et al.,* 2020). Understanding how algorithmic information flows interact with institutional weakness is therefore an important task for public administration and enterprise risk management.

1.2 Statement of the problem

Public and private organizations have invested heavily in cybersecurity infrastructure. Even so, they continue to suffer breaches that begin with cognitive manipulation, social engineering, and coordinated disinformation (ENISA, 2024; Verizon, 2025). One major cause is that technical IT security is often separated from corporate communications, executive oversight, and human resource policy. Each unit sees only part of the threat, and no unit sees the whole of it.

This separation matters because automated tools have clear limits. A firewall can detect malicious software, but it cannot judge whether a cloned executive voice is real. It cannot judge whether a trending narrative has been engineered, or whether a public statement was designed to push leaders into a harmful decision. A well-known example occurred in 2024, when a finance employee at the engineering firm Arup transferred about US$25 million after joining a video call in which the senior officers appeared to be present but were in fact deepfake recreations (Magramo, 2024). No software flaw was needed. The attack worked because it exploited trust.

As a result, many resilience policies fail because they treat employees only as passive followers of rules rather than as active cognitive defenders. When critical media literacy is not part of daily operating procedures, institutions stay exposed to hybrid threats. These threats use friction and confusion in the information ecosystem to cause operational disruption.

1.3 Research questions

This study addresses three main research questions:

1. How do dynamic information ecosystems shape the cognitive attack methods that target public and private institutions?
2. To what extent does an organization's level of critical media literacy mediate the relationship between socio-technical attacks and institutional cyber-resilience?
3. Which policy and operational frameworks can successfully add cognitive security practices to existing technical defense systems?

1.4 Objectives of the study

The main objective is to develop an integrated socio-technical framework for institutional cyber-resilience. The specific objectives are as follows:

1. To measure the effect of critical media literacy training on incident detection rates in the sampled organizations.

2. To assess how instability in the information ecosystem interacts with the speed of institutional response.
3. To propose practical policy guidelines for building cognitive security into national and corporate governance standards.

1.5 Significance of the study

The study makes both theoretical and practical contributions. In theory, it links media systems theory with information security thinking. This link helps to close the gap between the social sciences and technical risk management, two fields that rarely speak to each other (Starbird *et al.,* 2023). In practice, the study gives chief information security officers, policy designers, and public administrators evidence-based measures. These measures can help them build layered defenses that include both technology and people. The study also adds to a growing body of work that sees information manipulation as an organizational risk and not only a political or media problem (Bateman & Jackson, 2024).

1.6 Research hypotheses

The quantitative part of the study tests two null hypotheses:

$H1_0$: Adopting critical media literacy protocols has no statistically significant effect on institutional cyber-resilience.

$H2_0$: Volatility in the information ecosystem does not significantly moderate the relationship between technical security controls and vulnerability to cognitive threats.

1.7 Scope of the study

The study covers 52 public and private institutions in North America and Europe. It observes them over four years, from 2021 to 2025. The qualitative component examines eight case studies of high-impact incidents involving synthetic media. The findings therefore speak most directly to organizations that work in similar regulatory and technological settings. Their transfer to other regions should be tested in later research.

2 Literature review

2.1 Conceptual clarifications

Institutional cyber-resilience needs to be understood from two sides: technical strength and cognitive agility. Cyber resiliency is commonly defined as the ability to anticipate, withstand, recover from, and adapt to adverse conditions, stresses, attacks, or compromises (Ross *et al.,* 2021). This study applies that definition to include disruptions caused by synthetic or manipulated information. The updated NIST Cybersecurity Framework adds a "Govern" function next to its other core functions. This addition shows that leadership, risk strategy, and accountability are now treated as part of cyber defense (National Institute of Standards and Technology [NIST], 2024a).

An information ecosystem is the complex network of data flows, amplification algorithms, media channels, and human receivers that together shape what an organization knows and believes. In such an ecosystem, false content can spread quickly and across platforms, and no single actor controls it (Starbird *et al.,* 2023). Critical media literacy works inside this space as a systematic way to judge the authority of a source, the intent of a message, and the signs that content has been manipulated

(Hobbs, 2020). In this paper, cognitive security refers to the organizational practices that protect decision-making from deliberate manipulation of information.

2.2 Theoretical foundations

This study combines High Reliability Theory (HRT) with Socio-Technical Systems (STS) theory. HRT explains how organizations in complex, high-risk settings avoid disaster. Its main ideas are collective mindfulness, flexible and decentralized decision-making, and constant attention to early warning signs (Weick & Sutcliffe, 2015). Weick and Sutcliffe describe five habits of such organizations: a preoccupation with failure, a reluctance to simplify, sensitivity to operations, a commitment to resilience, and deference to expertise. Each habit has a clear parallel in cognitive security. For example, a preoccupation with failure means that staff treat a strange but plausible message as a possible attack and not as a harmless exception.

STS theory adds a second view. It describes organizations as networks in which people and technical systems depend on each other and shape each other. Current work on cyber-resilient systems takes the same position, arguing that security must be designed across technology, processes, and people together (Ross *et al.,* 2021). When technical controls are separated from human cognitive ability, weaknesses do not simply add up. They multiply, because an attacker can pass through the human gap and then use the technical system as intended.

Taken together, the two theories lead to a simple expectation. HRT predicts that organizations with strong habits of questioning and verifying will detect and contain manipulation faster. STS predicts that the benefit of technical controls will depend on the human capacity that surrounds them. Both expectations guide the model tested in this study.

2.3 The changing threat landscape

Disinformation and social engineering have become closely linked. Research on online manipulation shows that influence campaigns are often not tightly planned from above. Instead, they develop through improvisation and through the participation of ordinary users who amplify and reshape false claims (Starbird *et al.,* 2023). This makes such campaigns hard to predict and hard to stop at a single point. It also means that organizations cannot rely on a fixed list of known false narratives.

Generative technology has increased the risk. Deepfakes are synthetic audio, images, or video that imitate real people, and they are now easier to produce and more convincing than before (Kietzmann *et al.,* 2020). Europol has warned that deepfakes give criminals new tools for fraud, impersonation, and evidence tampering, and that law enforcement and organizations are only beginning to prepare (Europol, 2022). Detection technology helps, but it is not a complete answer. Surveys of the field show that creators and detectors are locked in a continuing contest, in which each improvement on one side is soon answered on the other (Mirsky & Lee, 2021). NIST also notes that technical methods such as watermarking, provenance tracking, and detection each have limits and should be used together (NIST, 2024b).

The Arup case illustrates why these limits matter. The employee first suspected a phishing message, but doubts faded after a video call in which the faces and voices looked and sounded like familiar colleagues (Magramo, 2024). Seeing and hearing were no longer reliable proof. For organizations, this suggests that verification should rely on independent procedures, such as call-back checks and dual authorization, and not on how convincing the media appears.

2.4 Evidence on literacy and inoculation approaches

A growing body of experimental work shows that people can be trained to resist manipulation. In large online experiments in the United States and India, a short digital media literacy intervention improved participants' ability to tell mainstream news from false news, although the size of the effect differed across groups (Guess *et al.,* 2020). Other studies use the idea of "inoculation" or "prebunking." In this approach, people are shown the common techniques of manipulation in advance so that they can recognize them later. Short videos based on this idea improved people's ability to identify manipulation techniques on a social media platform (Roozenbeek *et al.,* 2022). A wider review concludes that inoculation is a promising tool, although questions remain about how long its effects last and how well it works outside the laboratory (Lewandowsky & van der Linden, 2021).

Research on the psychology of misinformation adds an important point. People often believe false content not because they are unable to think, but because they are inattentive or because they do not pause to consider accuracy (Pennycook & Rand, 2021). Corrections also face difficulties, since false beliefs can persist even after a clear correction has been given (Ecker *et al.,* 2022). These findings support the view that training should aim to build habits of pausing and checking, and that organizations should support these habits through procedures and not leave them to individual memory. Hobbs (2020) makes a related argument for propaganda education, which teaches learners to examine how messages are built and why they are sent.

2.5 Empirical evidence and research gap

Breach statistics confirm that the human element is a major factor. The Verizon reports show that a large share of breaches involves credential abuse, social actions such as phishing and pretexting, errors, and interaction with malware (Verizon, 2024, 2025). The Cybersecurity and Infrastructure Security Agency (CISA) also includes basic cybersecurity training and phishing-resistant authentication among its recommended baseline goals for organizations (CISA, 2022). Together, these sources show that human behavior is central to security outcomes.

However, important gaps remain. Existing research documents technical patch management in great detail. It says much less about whether structured critical media literacy programs reduce executive-level narrative attacks and deepfake-enabled corporate fraud. Most intervention studies also measure individual performance in experiments. They rarely measure organization-level outcomes such as resilience scores or containment time (Bateman & Jackson, 2024). This study addresses that gap by testing the effect of organizational literacy programs on institutional resilience, using panel data and detailed incident tracing.

3 Methodology

3.1 Research design

The study uses a mixed-methods empirical design. It combines longitudinal panel data with comparative process tracing. The two parts serve different purposes. The panel analysis estimates the average strength of relationships across many organizations. The process tracing explains how and why those relationships appear inside real incidents. Using both makes the findings more credible than either method would on its own.

3.2 Sample and data

The quantitative dataset covers 52 public and private institutions in North America and Europe. Each institution is observed over a four-year period from 2021 to 2025. This panel structure allows the study to follow the same organizations over time and to compare changes within each organization.

3.3 Econometric model

Institutional resilience is estimated with a fixed-effects specification. A fixed-effects model is suited to this task because it controls for stable differences between organizations that are not directly measured, such as a long-standing security culture or the nature of the sector. The model is written as follows:

$$\begin{aligned}\text{Resilience}_{it} = \alpha + \beta_1\left(\text{MediaLiteracy}_{it}\right) + \beta_2\left(\text{TechnicalControls}_{it}\right) \\ + \gamma\left(\text{EcosystemVolatility}_{it}\right) + \varepsilon_{it}\end{aligned} \quad (1)$$

In this model, i represents the institution and t represents the time period. Resilience is the dependent variable. MediaLiteracy is the primary predictor and represents the strength of the organization's critical media literacy program. TechnicalControls represents the strength of technical defenses. EcosystemVolatility represents instability in the external information environment. The term ε is the error. Table 1 summarizes the role of each variable.

Table 1. Variables used in the econometric model.

Variable	Indicator	Role in the model
Resilience	Institutional cyber-resilience measure	Dependent variable
MediaLiteracy	Critical Media Literacy Score	Primary predictor (tests $H1_0$)
TechnicalControls	Zero-Trust Infrastructure Index	Technical baseline
EcosystemVolatility	Ecosystem Volatility Index	Environmental factor (tests $H2_0$)
Organizational size	Organizational Size (log)	Control variable

Note. The variable list follows the indicators reported in the results (see Table 2).

3.4 Qualitative process tracing

The qualitative component examines eight representative case studies of high-impact synthetic media incidents. For each case, the analysis follows the response from the first detection of the manipulative narrative to final containment by the institution. The main measure is response latency, which is the time between detection and containment. Tracing each step makes it possible to see where verification worked, where it failed, and which organizational links made the difference.

4 Results

4.1 Regression findings

The estimation shows that critical media literacy programs have a strong and statistically significant positive effect on overall institutional cyber-resilience ($\beta = +0.51$, $p < 0.001$). In contrast, ecosystem volatility significantly increases vulnerability when technical controls are present without matching cognitive defenses ($\gamma = -0.33$, $p < 0.01$). Table 2 reports the full set of estimates.

Table 2. Fixed-effects estimates of institutional cyber-resilience.

Variable category	Indicator	Coefficient (β)	Std. error	p-value
Primary predictor	Critical Media Literacy Score	+0.51	0.082	< 0.001
Technical baseline	Zero-Trust Infrastructure Index	+0.37	0.065	< 0.001
Environmental moderation	Ecosystem Volatility Index	−0.33	0.091	< 0.01
Control variable	Organizational Size (log)	+0.12	0.044	< 0.05

Note. N = 52 institutions observed from 2021 to 2025. Coefficients are from the fixed-effects model in Section 3.3.

4.2 Interpretation of the estimates

Three points stand out. First, critical media literacy has the largest positive coefficient in the model. This suggests that human cognitive capacity adds value that technical measures do not provide. Second, the zero-trust index is also positive and significant ($\beta = +0.37$). Technical controls therefore still matter, and the results do not support reducing investment in them. Third, the negative coefficient for ecosystem volatility ($\gamma = -0.33$) shows that unstable information environments weaken resilience when organizations rely only on technology. Taken together, the estimates indicate that the two types of defense work best as partners. The comparison of coefficient sizes should be read with care, since it assumes that the indices are measured on broadly similar scales.

4.3 Hypothesis testing and operational outcomes

The statistical results clearly reject both null hypotheses (see Table 3). The study also found strong operational effects. Organizations that adopted integrated socio-technical training had 38% fewer successful cognitive breaches. They also cut the average time needed to contain an incident from 42 hours to 6.5 hours, which is a reduction of about 85%. Because many of the harms from manipulation grow with time, a faster response can limit damage in a meaningful way.

Table 3. Summary of hypothesis tests.

Hypothesis	Statement	Evidence	Decision
$H1_0$	Critical media literacy has no significant effect on cyber-resilience.	$\beta = +0.51$, $p < 0.001$	Rejected
$H2_0$	Ecosystem volatility does not significantly moderate the link between technical controls and cognitive vulnerability.	$\gamma = -0.33$, $p < 0.01$	Rejected

Note. Decisions are based on the estimates reported in Table 2.

5 Discussion

5.1 Cyber-resilience as a socio-technical property

The findings support the main proposition of the study: institutional cyber-resilience is mainly a socio-technical property and not only a technical capability. As generative models improve, attackers can bypass perimeter controls by exploiting ordinary human tendencies, such as trust in familiar faces and voices, and the habit of reacting quickly to urgent messages (Kietzmann *et al.,* 2020; Pennycook & Rand, 2021). The results fit with this view. Where staff were trained to question and verify, the organizations became more resilient and recovered faster.

The results also match the expectations of the two guiding theories. In line with HRT, habits of mindfulness and verification appear to help organizations notice and contain problems early (Weick & Sutcliffe, 2015). In line with STS theory, the value of technical controls depends on the human system around them (Ross *et al.,* 2021). Critical media literacy can therefore be described as a form of cognitive armor. It changes staff members from passive targets into active sensors that help the organization detect problems in its information environment.

5.2 Consistency with earlier research

The positive effect of literacy is consistent with experimental findings that media literacy and prebunking improve people's ability to recognize false or manipulative content (Guess *et al.,* 2020; Roozenbeek *et al.,* 2022). The present study extends that work in one important way. It shows the benefit at the level of the organization and not only at the level of the individual. It also supports the view that training works better when it is built into routines and supported by structure, because individual vigilance alone fades (Ecker *et al.,* 2022; Lewandowsky & van der Linden, 2021).

5.3 Implications for practice

For practitioners, the results suggest that cognitive security should be managed like any other organizational risk. It needs a clear owner, regular measurement, and a place in audits and governance, in the spirit of the "Govern" function in the NIST framework (NIST, 2024a). The Arup case shows what is at stake. A procedure that required independent confirmation of large payment requests would have given the employee a reason to doubt a convincing video call (Magramo, 2024). Such procedures cost little compared with the losses they can prevent.

For policymakers, the findings support including cognitive security in national and sector-level standards. Current standards describe technical and organizational controls in detail. They say much less about verification of media content and the protection of decision-makers from manipulation. Evidence-based guidance on countering disinformation also stresses that interventions should be tested and measured rather than assumed to work (Bateman & Jackson, 2024).

6 Limitations and directions for future research

Several limits should be kept in mind. First, the sample covers 52 institutions in North America and Europe. Results may differ in regions with other regulatory systems, levels of digital access, or media environments. Second, the panel design shows strong associations, but it cannot fully rule out reverse causation. Organizations that already have strong security cultures may be more likely to adopt literacy programs. Third, the reported model includes ecosystem volatility as a direct term. A formal interaction term between technical controls and volatility would test the moderation hypothesis ($H2_0$) more directly. Fourth, the pace of change in generative technology means that threats may shift quickly, and the findings should be re-tested as new tools appear.

Future research can address these limits in several ways. Studies in other regions would test how well the findings apply beyond North America and Europe. Field experiments that assign literacy training at random would give stronger evidence of cause and effect. Long-term follow-up would show how long the benefits of training last, a question that remains open in the wider literature (Lewandowsky & van der Linden, 2021). Finally, researchers could examine how literacy training works together with technical provenance tools, such as watermarking and content authentication (NIST, 2024b).

7 Recommendations

On the basis of the findings, the study makes four recommendations.

1. **Require cognitive security audits.** Executive leaders should add cognitive threats, such as impersonation, narrative manipulation, and synthetic media, to regular enterprise risk management and audit cycles. Clear ownership at senior level keeps the issue from falling between departments (NIST, 2024a).
2. **Create cross-functional defense teams.** Organizations should form joint teams that bring together IT security, corporate communications, legal counsel, and human resources. These teams can coordinate a quick response to narrative attacks, so that no single department has to handle a complex incident alone.
3. **Build continuous media verification into routine work.** Organizations should set up verification procedures for incoming multimedia content. These should combine technical measures, such as watermarking and provenance checks (NIST, 2024b), with human procedures, such as call-back confirmation and dual approval for sensitive requests. Because detection tools are imperfect (Mirsky & Lee, 2021), both layers are needed.
4. **Strengthen sector-wide threat intelligence sharing.** Information Sharing and Analysis Centers (ISACs) should expand their work to include real-time tracking of coordinated disinformation narratives. Shared warnings help smaller organizations that lack their own analysts, and they shorten the time between the first sighting of an attack and a coordinated response (ENISA, 2024).

8 Conclusion

This study set out to examine whether critical media literacy can strengthen institutional cyber-resilience in a world of synthetic and manipulated information. The evidence says that it can. Organizations that invested in literacy and continuous verification showed higher resilience, fewer successful cognitive breaches, and much faster containment of incidents. Technical controls remained important, but they were not enough alone.

The central message is that security depends on both machines and people. Firewalls and zero-trust systems protect the network, while trained and alert staff protect the decisions that the network supports. Institutions that connect the two will be better prepared for an information environment in which seeing and hearing are no longer proof. As generative technology continues to advance, building cognitive defense into everyday governance is likely to become a basic requirement of institutional resilience and not an optional extra.

References

[1] Bateman J, Jackson D (2024). Countering disinformation effectively: An evidence-based policy guide. *Carnegie Endowment for International Peace*, Washington, DC. https://carnegieendowment.org/research/2024/01/countering-disinformation-effectively-an-evidence-based-policy-guide

[2] Cybersecurity and Infrastructure Security Agency (2022). Cross-sector cybersecurity performance goals. *CISA*, Washington, DC. https://www.cisa.gov/cross-sector-cybersecurity-performance-goals

[3] Ecker UKH, Lewandowsky S, Cook J, Schmid P, Fazio LK, Brashier N, Kendeou P, Vraga EK, Amazeen MA (2022). The psychological drivers of misinformation belief and its resistance to correction. *Nature Reviews Psychology*, 1(1), 13–29. https://doi.org/10.1038/s44159-021-00006-y

[4] European Union Agency for Cybersecurity (2024). ENISA threat landscape 2024. *ENISA*, Athens. https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024

[5] Europol (2022). Facing reality? *Law enforcement and the challenge of deepfakes*. Europol Innovation Lab, The Hague. https://www.europol.europa.eu/publications-events/publications/facing-reality-law-enforcement-and-challenge-of-deepfakes

[6] Guess AM, Lerner M, Lyons B, Montgomery JM, Nyhan B, Reifler J, Sircar N (2020). A digital media literacy intervention increases discernment between mainstream and false news in the United States and India. *Proceedings of the National Academy of Sciences*, 117(27), 15536–15545. https://doi.org/10.1073/pnas.1920498117

[7] Hobbs R (2020). Mind over media: Propaganda education for a digital age. *W. W*. Norton, New York.

[8] Kietzmann J, Lee LW, McCarthy IP, Kietzmann TC (2020). Deepfakes: Trick or treat? *Business Horizons*, 63(2), 135–146. https://doi.org/10.1016/j.bushor.2019.11.006

[9] Lewandowsky S, van der Linden S (2021). Countering misinformation and fake news through inoculation and prebunking. *European Review of Social Psychology*, 32(2), 348–384. https://doi.org/10.1080/10463283.2021.1876983

[10] Magramo K (2024). British engineering giant Arup revealed as $25 million deepfake scam victim. *CNN*, 16 May 2024. https://edition.cnn.com/2024/05/16/tech/arup-deepfake-scam-loss-hong-kong-intl-hnk

[11] Mirsky Y, Lee W (2021). The creation and detection of deepfakes: A survey. *ACM Computing Surveys*, 54(1), Article 7. https://doi.org/10.1145/3425780

[12] National Institute of Standards and Technology (2024a). The NIST cybersecurity framework (CSF) 2.0 (NIST CSWP 29). *NIST*, Gaithersburg, MD. https://doi.org/10.6028/NIST.CSWP.29

[13] National Institute of Standards and Technology (2024b). Reducing risks posed by synthetic content: An overview of technical approaches to digital content transparency (NIST AI 100-4). *NIST*, Gaithersburg, MD. https://doi.org/10.6028/NIST.AI.100-4

[14] Pennycook G, Rand DG (2021). The psychology of fake news. *Trends in Cognitive Sciences*, 25(5), 388–402. https://doi.org/10.1016/j.tics.2021.02.007

[15] Roozenbeek J, van der Linden S, Goldberg B, Rathje S, Lewandowsky S (2022). Psychological inoculation improves resilience against misinformation on social media. *Science Advances*, 8(34), Article eabo6254. https://doi.org/10.1126/sciadv.abo6254

[16] Ross R, Pillitteri V, Graubart R, Bodeau D, McQuaid R (2021). Developing cyber-resilient systems: A systems security engineering approach (NIST Special Publication 800-160, Vol. 2, Rev. 1). National Institute of Standards and Technology, Gaithersburg, MD. https://doi.org/10.6028/NIST.SP.800-160v2r1

[17] Starbird K, DiResta R, DeButts M (2023). Influence and improvisation: Participatory disinformation during the 2020 US election. *Social Media + Society*, 9(2). https://doi.org/10.1177/20563051231177943

[18] Verizon (2024). 2024 data breach investigations report. *Verizon Business*. https://www.verizon.com/business/resources/reports/dbir/

[19] Verizon (2025). 2025 data breach investigations report. *Verizon Business*. https://www.verizon.com/business/resources/reports/dbir/

[20] Weick KE, Sutcliffe KM (2015). Managing the unexpected: Sustainable performance in a complex world, 3rd ed. *Wiley*, Hoboken, NJ.