Research Synthesis Series • 2025–2026

Cybersecurity & Data Privacy

Research Insights 2025–2026

A synthesis of **97+ peer-reviewed papers**
across 8 cybersecurity and data privacy domains

97+	8	94.3%	22
Papers Reviewed	Research Domains	Federated AI Accuracy	PDF Pages

Atine John Bosco
MBCS · CISA · Google Cloud Digital Leader · ITIL 4 Practitioner
PhD Candidate · ICT Consultant
www.atine.info

May 2026 · CC BY 4.0

Contents

1 Executive Summary **2**

2 The 2025–2026 Cyber Threat Landscape **2**
2.1 AI-Powered Attacks: The New Threat Paradigm . . . 2
2.2 Ransomware Evolution: From Encryption to Extortion . . . 3
2.3 APT Industrialisation . . . 3

3 Zero Trust Architecture **3**
3.1 Core Principles . . . 3
3.2 AI-Augmented Zero Trust . . . 3
3.3 ZTA Implementation Challenges . . . 4

4 AI & Machine Learning for Cybersecurity **4**
4.1 Federated AI: Privacy-Preserving Intrusion Detection . . . 4
4.2 AI Security Operations: SIEM, EDR, and SOAR . . . 4
4.3 Adversarial AI: The Double-Edged Sword . . . 4

5 Data Privacy & Regulatory Compliance **5**
5.1 The Global Regulatory Landscape . . . 5
5.2 Privacy-by-Design for AI Systems . . . 5

6 Cloud Security & Identity Access Management **6**
6.1 AI-Enhanced Federal Cloud Security . . . 6
6.2 IAM Best Practices . . . 6

7 Post-Quantum Cryptography & Quantum-Safe Security **6**
7.1 The Quantum Threat to Current Cryptography . . . 6
7.2 NIST Post-Quantum Standards . . . 6
7.3 Enterprise PQC Migration Roadmap . . . 7

8 Privacy-Enhancing Technologies (PETs) **7**
8.1 PET Taxonomy and Comparative Analysis . . . 7
8.2 PET Integration with Zero Trust . . . 8

9 Strategic Roadmap for ICT Leaders **8**
9.1 Cybersecurity Maturity Model . . . 8
9.2 6-Point Action Framework for ICT Leaders . . . 9

References **9**

Executive Summary

Key Findings at a Glance

- **AI-augmented cyber attacks** are the defining threat of 2025–2026: LLMs powering deepfake phishing, intelligent malware obfuscation, and automated social engineering at industrial scale
- **Zero Trust Architecture (ZTA)** is the consensus defensive posture — replacing perimeter security with continuous verification, microsegmentation, and identity-centric controls
- **Federated AI intrusion detection** achieves **94.3% accuracy** in distributed environments (Mahamad, 2024) while preserving data privacy
- **Post-Quantum Cryptography (PQC)** migration is urgent — CRYSTALS-Kyber and CRYSTALS-Dilithium (NIST FIPS 203/204) are the standardised algorithms for enterprise transitions
- **NIS2 and DORA** are reshaping European cybersecurity compliance; US CISA guidance drives federal cloud security mandates
- **Privacy-Enhancing Technologies** (homomorphic encryption, differential privacy, federated learning) are converging with ZTA for privacy-preserving security operations

The 2025–2026 cybersecurity landscape is defined by a fundamental asymmetry: attackers are harnessing generative AI to industrialise threats at unprecedented scale, while defenders are racing to deploy Zero Trust frameworks, AI-powered detection, and quantum-resistant cryptography. This report synthesises findings from **97+ peer-reviewed papers** across eight domains, providing ICT leaders with evidence-based strategic guidance.

The most consequential finding is the *weaponisation of large language models* for offensive operations — automated spear-phishing, deepfake impersonation, and adaptive malware that evades signature-based detection [**?**, **?**]. Against this backdrop, Zero Trust Architecture emerges not merely as a framework but as a survival imperative for organisations operating in hybrid and multi-cloud environments [**?**, **?**].

The 2025–2026 Cyber Threat Landscape

AI-Powered Attacks: The New Threat Paradigm

The convergence of generative AI with cyber offensive operations marks the most significant threat evolution since ransomware. Sharma (2025) documents the comprehensive 2024–2025 threat landscape, identifying three primary AI-enabled attack vectors [1]:

Three AI-Enabled Attack Vectors

1. **Deepfake Social Engineering:** LLMs generate hyper-personalised phishing content; voice/video deepfakes impersonate executives for business email compromise (BEC) fraud
2. **Intelligent Malware Obfuscation:** AI-generated polymorphic code that mutates to evade antivirus signatures and sandboxed analysis environments
3. **Automated Reconnaissance:** AI-driven OSINT tools that map attack surfaces, identify vulnerabilities, and prioritise targets at machine speed

Ransomware Evolution: From Encryption to Extortion

Modern ransomware groups have pivoted from encryption-first to *extortion-first* strategies. Key trends identified in the literature [**?**, **?**]:

- **"Harvest Now, Decrypt Later" (HNDL):** Adversaries exfiltrate encrypted data today, planning to decrypt once quantum computers mature — making current data breaches a long-term risk
- **Double/Triple Extortion:** Data encryption + public leak threat + DDoS attack combined for maximum leverage
- **Ransomware-as-a-Service (RaaS):** Industrialised affiliate models enabling non-technical actors to deploy sophisticated ransomware campaigns

APT Industrialisation

Advanced Persistent Threats (APTs) are operating with commercial-grade scalability. Sharma et al. (2024) describe adversaries functioning as *commercial syndicates* — with dedicated R&D, QA testing of attack tools, and customer-service-style support for RaaS affiliates [3]. This industrialisation demands equivalent organisational maturity in defence.

Zero Trust Architecture

Core Principles

Zero Trust Architecture (ZTA) replaces the implicit trust of perimeter-based security with the principle: *"Never trust, always verify."* The NIST ZTA framework (SP 800-207) defines three core tenets [**?**, **?**]:

NIST Zero Trust Architecture — Core Tenets

1. **Verify Explicitly:** Authenticate and authorise every request based on all available data points (identity, location, device health, service/workload, data classification, anomalies)
2. **Use Least-Privilege Access:** Limit user access with just-in-time (JIT) and just-enough-access (JEA), risk-based adaptive policies, and data protection
3. **Assume Breach:** Minimise blast radius, segment access, verify end-to-end encryption, use analytics to detect and respond to threats

AI-Augmented Zero Trust

Chokkanathan et al. (2024) demonstrate that AI/ML significantly enhances ZTA effectiveness through dynamic policy decisions and continuous behavioural authentication [10]:

AI Enhancements to Zero Trust

- **Behavioural Biometrics:** Continuous authentication via keystroke dynamics, mouse patterns, and usage behaviour — invisible to users
- **Adaptive Access Policies:** ML models dynamically adjusting access permissions based on real-time risk scores
- **Micro-segmentation Automation:** AI-driven network segmentation that adapts to workload patterns and detected threats
- **Anomaly Detection:** Statistical models flagging deviations from established user and entity baselines (UEBA)

ZTA Implementation Challenges

Table 1: Zero Trust Implementation Challenges and Mitigations

Challenge	Impact	Mitigation Strategy
Legacy System Integration	Cannot support modern identity protocols	Wrap legacy systems with ZTA proxies/gateways
Policy Orchestration	Complexity of managing thousands of micro-policies	Automate with AI policy engines; use policy-as-code
User Experience	Friction from continuous verification	Invisible authentication (biometrics, device trust)
Telemetry Privacy	Behavioural monitoring raises privacy concerns	Anonymise telemetry; implement GDPR-compliant logging
Cost and Expertise	Significant investment in tools and skills	Phased implementation; start with identity and IAM

AI & Machine Learning for Cybersecurity

Federated AI: Privacy-Preserving Intrusion Detection

The landmark empirical result in AI-powered cybersecurity is Mahamad's (2024) federated CNN architecture for distributed intrusion detection [4]:

Federated AI Intrusion Detection — Key Results

- **Architecture:** Federated CNN combining edge node models with cloud aggregation — no raw data leaves local environments
- **Accuracy:** 94.3% intrusion detection accuracy in proof-of-concept evaluation
- **Privacy Benefit:** Local model training preserves data sovereignty; only model gradients are shared
- **Scalability:** Demonstrated across heterogeneous cloud and edge nodes
- **Application:** Particularly valuable for healthcare, finance, and government where data cannot be centralised

AI Security Operations: SIEM, EDR, and SOAR

Folorunso et al. (2024) document the transformative impact of AI across security operations [13]:

- **AI-Powered SIEM:** Correlates millions of security events per second; ML models reduce false positives by 60–80% compared to rule-based systems
- **Next-Generation EDR:** Behavioural detection of fileless malware, living-off-the-land (LotL) attacks, and zero-day exploits that evade signatures
- **SOAR Automation:** AI-orchestrated playbooks for automated threat containment, IOC enrichment, and incident response
- **Threat Intelligence:** NLP-powered analysis of threat feeds, dark web monitoring, and vulnerability correlation

Adversarial AI: The Double-Edged Sword

Zhang (2025) provides a comprehensive analysis of AI security risks [14]:

AI Security Risks — Classification

Risk Type	Description
Data Poisoning	Adversaries inject malicious training samples to corrupt model behaviour
Model Evasion	Adversarial examples crafted to fool ML classifiers (e.g., malware that mimics benign traffic)
Model Inversion	Extracting training data from model outputs — privacy violation risk
Membership Inference	Determining whether a specific record was in the training dataset
Prompt Injection	Manipulating LLM-based security tools via adversarial inputs

Data Privacy & Regulatory Compliance

The Global Regulatory Landscape

The regulatory environment for data privacy has intensified dramatically in 2024–2026, with new obligations spanning AI governance, operational resilience, and cross-border data flows [**?**, **?**]:

Key Regulatory Frameworks 2025–2026

EU AI Act (2024): First comprehensive AI regulation; classifies AI systems by risk level. High-risk AI (biometric systems, critical infrastructure) requires conformity assessments, human oversight, and transparency obligations

NIS2 Directive: Expands scope to 18 sectors; mandates 24-hour incident reporting, supply-chain security, and board-level accountability for cybersecurity

DORA (Digital Operational Resilience Act): Financial sector mandate for ICT risk management, incident classification, resilience testing, and third-party oversight

GDPR (Ongoing): Article 22 right-to-explanation drives demand for interpretable AI; DPAs increasingly scrutinising AI decision-making systems

CISA Cloud Security Guidance: US federal mandates for Zero Trust adoption, cloud security controls, and threat intelligence sharing

Privacy-by-Design for AI Systems

Mbah & Evelyn (2024) provide strategic guidance for embedding privacy into AI-powered cybersecurity systems [15]:

- Implement *data minimisation* — collect only what is necessary for security operations; avoid retaining raw user behavioural data
- Apply *purpose limitation* — security telemetry must not be repurposed for commercial profiling or employee monitoring
- Conduct *Data Protection Impact Assessments (DPIAs)* for all AI security deployments processing personal data at scale
- Establish *retention schedules* for security logs aligned with GDPR requirements and operational necessity

Cloud Security & Identity Access Management

AI-Enhanced Federal Cloud Security

Ofili et al. (2025) present a comprehensive framework for AI-enhanced federal cloud security aligned with CISA mandates [6]:

AI-Enhanced Federal Cloud Security Framework

1. **AI Threat Intelligence Integration:** Real-time threat feeds correlated with AI models for predictive threat detection in cloud workloads
2. **Zero Trust Cloud Architecture:** Identity-centric access control with continuous verification for all cloud service access
3. **CISA Compliance Automation:** Automated compliance checking against CISA Binding Operational Directives and Emergency Directives
4. **Federated Identity Management:** Cross-agency identity federation with SAML 2.0 / OIDC and risk-based MFA
5. **Cloud Security Posture Management (CSPM):** Continuous misconfiguration detection and automated remediation

IAM Best Practices

Identity Access Management — 2025–2026 Best Practices

- **Privileged Access Workstations (PAWs):** Dedicated, hardened devices for administrative tasks — isolated from general-purpose browsing
- **Just-in-Time (JIT) Access:** Temporary privilege elevation approved through automated workflows; no standing privileged access
- **Passwordless Authentication:** FIDO2/WebAuthn passkeys replacing passwords; phishing-resistant by design
- **Conditional Access Policies:** Risk-based access decisions incorporating device compliance, location, and user behaviour signals
- **Non-Human Identity (NHI) Management:** Securing service accounts, API keys, and machine identities — often the weakest link in cloud environments

Post-Quantum Cryptography & Quantum-Safe Security

The Quantum Threat to Current Cryptography

Quantum computers running Shor's algorithm will break RSA, ECC, and Diffie-Hellman — the cryptographic foundations of TLS, VPNs, and digital signatures. The "Harvest Now, Decrypt Later" (HNDL) strategy means adversaries are already collecting encrypted data for future quantum decryption [**?**, **?**].

NIST Post-Quantum Standards

The NIST PQC standardisation process concluded in 2024 with three primary standards [16]:

Table 2: NIST Post-Quantum Cryptography Standards

Algorithm	NIST Standard	Primary Use Case
CRYSTALS-Kyber	FIPS 203	Key Encapsulation Mechanism (KEM) — replaces RSA/ECC for key exchange
CRYSTALS-Dilithium	FIPS 204	Digital signatures — replaces RSA/ECDSA
SPHINCS+	FIPS 205	Hash-based signatures — conservative, stateless alternative
FALCON	(Draft)	Compact lattice-based signatures for constrained environments

Enterprise PQC Migration Roadmap

5-Step Post-Quantum Migration Framework

1. **Cryptographic Inventory:** Catalogue all cryptographic assets — certificates, keys, algorithms, protocols — using automated discovery tools
2. **Risk Prioritisation:** Identify long-lived secrets (classified data, financial records) most vulnerable to HNDL attacks; prioritise these for immediate migration
3. **Hybrid Deployment:** Implement classical + PQC hybrid schemes (e.g., X25519 + CRYSTALS-Kyber) to maintain backward compatibility during transition
4. **PKI Modernisation:** Upgrade certificate lifecycle management for PQC certificate sizes (larger than classical); update HSMs and key management systems
5. **Protocol Updates:** Migrate TLS, SSH, VPN, and code-signing to PQC-capable versions; test interoperability across all vendor stacks

Privacy-Enhancing Technologies (PETs)

PET Taxonomy and Comparative Analysis

Murtaza (2025) provides a comprehensive analysis of Privacy-Enhancing Technologies for network security [12]:

Table 3: Privacy-Enhancing Technologies — Comparative Analysis

PET	Primary Use Case	Key Limitation	Maturity
Homomorphic Encryption	Compute on encrypted cloud data	High performance cost; engineering complexity	Research/Early Commercial
Differential Privacy	Privacy-preserving aggregate analytics	Utility vs. privacy tradeoff tuning	Production (Apple, Google)
Federated Learning	Collaborative ML without raw data sharing	Model update leakage; orchestration complexity	Production
Secure Multi-Party Computation	Joint computation without revealing inputs	Communication overhead; complex protocols	Early Commercial
Trusted Execution Environments (TEE)	Isolated secure enclaves for sensitive computation	Hardware dependency; side-channel attacks	Production (Intel SGX, AMD SEV)
Zero-Knowledge Proofs	Verify claims without revealing underlying data	Computational cost; complex implementation	Emerging

PET Integration with Zero Trust

Yelkoti (2025) demonstrates that combining PETs with Zero Trust creates a comprehensive privacy-preserving security architecture [17]:

PET + Zero Trust Integration Architecture

- **Federated Learning + ZTA:** Distributed threat model training across organisational boundaries without sharing sensitive security telemetry
- **Differential Privacy + SIEM:** Aggregate security metrics with provable privacy guarantees — enabling cross-sector threat intelligence sharing
- **TEE + Cloud Security:** Confidential computing for sensitive workloads in untrusted cloud environments (Intel SGX, AMD SEV, AWS Nitro Enclaves)
- **Homomorphic Encryption + Analytics:** Encrypted search and analysis of security logs without decryption — preserving data sovereignty in multi-cloud deployments

Strategic Roadmap for ICT Leaders

Cybersecurity Maturity Model

Five-Stage Cybersecurity Maturity Model

1. **Stage 1 — Reactive:** Incident response only; no proactive monitoring; basic antivirus and firewall. *Risk: Very High*
2. **Stage 2 — Compliant:** Meeting minimum regulatory requirements; basic SIEM; documented policies. *Risk: High*
3. **Stage 3 — Proactive:** Threat hunting; vulnerability management programme; security awareness training. *Risk: Medium*

4. **Stage 4 — Predictive:** AI-powered threat detection; Zero Trust implementation; automated response. *Risk: Low*
5. **Stage 5 — Adaptive:** Continuous improvement; quantum-ready cryptography; federated security operations; AI governance. *Risk: Minimal*

6-Point Action Framework for ICT Leaders

Strategic Action Framework 2025–2026

1. **Adopt Zero Trust Architecture:** Begin with identity and IAM hardening (passwordless auth, JIT access, PAWs); extend to network microsegmentation and workload protection over 18–24 months
2. **Deploy AI-Powered Security Operations:** Implement AI-enhanced SIEM/SOAR with behavioural analytics; establish a Security Operations Centre (SOC) with automated playbooks
3. **Initiate PQC Migration:** Conduct cryptographic inventory now; implement hybrid PQC schemes for highest-risk data; plan full migration by 2027–2030
4. **Achieve Regulatory Compliance:** Map controls to NIS2, DORA, GDPR, and EU AI Act; appoint a dedicated Data Protection Officer (DPO) and AI compliance function
5. **Invest in Human Capital:** Cybersecurity skills gap is critical — invest in training, certifications (CISSP, CISM, CEH), and security awareness programmes for all staff
6. **Build Supply Chain Resilience:** Implement vendor risk assessments, software bill of materials (SBOM), and third-party security audits as part of procurement processes

References

[1] Sharma, P. (2025). A Comprehensive Study on Modern Cyber Security Threats and Defense Mechanisms: The 2024–2025 Landscape. *IJSREM*. doi:10.55041/ijsrem55286

[2] Arslan, R., Özseven, M., & Aydın, M. M. (2025). Transforming European Cybersecurity: AI-Powered Threat Analysis, Quantum Age, Blockchain/Crypto Risks, and Regulatory Strategies. *IJEAS*. doi:10.24107/ijeas.1619600

[3] Sharma, S., Agrawal, S., & Kumar, S. (2024). Unlocking Cybersecurity Horizons: Exploring Cutting-Edge Technologies, Strategies, and Trends. *IEEE ICEC*. doi:10.1109/icec59683.2024.10837210

[4] Mahamad, H. (2024). Guardians of the Data Galaxy: A Federated AI and Cloud Synergy for Zero-Trust Cybersecurity Models. *IJEHSS*. doi:10.54922/ijehss.2024.0746

[5] Kolawole, I. (2025). Leveraging Cloud-based AI and Zero Trust Architecture to Enhance US Cybersecurity and Counteract Foreign Threats. *ResearchGate*. ResearchGate

[6] Ofili, B. et al. (2025). Enhancing Federal Cloud Security with AI: Zero Trust, Threat Intelligence, and CISA Compliance. *WJARR*. doi:10.30574/wjarr.2025.25.2.0620

[7] Anonymous. (2024). Zero-Trust Security Model for Enhanced Cloud Security and Data Privacy. *SSRN*. doi:10.2139/ssrn.4904958

[8] Joshi, H. (2024). Emerging Technologies Driving Zero Trust Maturity Across Industries. *IEEE OJCS*. doi:10.1109/ojcs.2024.3505056

[9] Bartakke, J. & Kashyap, R. (2024). The Usage of Clouds in Zero-Trust Security Strategy. *JIOS*. doi:10.31341/jios.48.1.8

[10] Chokkanathan, K. et al. (2024). AI-Driven Zero Trust Architecture: Enhancing Cyber-Security Resilience. *IEEE CSITSS*. doi:10.1109/csitss64042.2024.10816746

[11] Gurram, A. (2025). Generative AI for Enhanced Cybersecurity: Building a Zero-Trust Architecture with Agentic AI. *WJAETS*. doi:10.30574/wjaets.2025.15.1.0504

[12] Murtaza, M. (2025). Emerging Trends in Network Security: Next-Generation Firewalls, AI-Driven Threat Detection, and Privacy-Preserving Technologies. *Zenodo*. doi:10.5281/zenodo.16813761

[13] Folorunso, A. et al. (2024). Impact of AI on Cybersecurity and Security Compliance. *GJETA*. doi:10.30574/gjeta.2024.21.1.0193

[14] Zhang, T. (2025). Research on Information Security Issues in the Field of Artificial Intelligence. *ACE*. doi:10.54254/2755-2721/2025.bj27253

[15] Mbah, G. O. & Evelyn, A. N. (2024). AI-Powered Cybersecurity: Strategic Approaches to Mitigate Risk and Safeguard Data Privacy. *WJARR*. doi:10.30574/wjarr.2024.24.3.3695

[16] Anonymous. (2025). Post-Quantum Cryptography-Safe Network Architectures: Design Frameworks and Implementation Strategies for Enterprise Zero-Trust Environments. *Zenodo*. doi:10.5281/zenodo.1693

[17] Yelkoti, N. K. K. R. (2025). Next-Generation Cryptographic Security in Multi-Cloud Enterprises: AI-Enhanced Data Privacy, Protection, and Threat-Resilient Automation. *IJCESEN*. doi:10.22399/ijcesen.4247

[18] Anonymous. (2024). Privacy-Preserving Technique in Cybersecurity: Balancing Data Protection and User Rights. *Al-Kindi Publishers*. URL

About the Author

Atine John Bosco is a Chartered Information Scientist (MBCS), Certified Information Systems Auditor (CISA), Google Cloud Digital Leader, and ITIL 4 Practitioner.

With 15+ years in ICT leadership across Africa and international organisations, he specialises in digital transformation, cybersecurity governance, and AI-driven ICT strategy.

www.atine.info

Research Insights Series · CC BY 4.0 · May 2026