Journal of Corporate Finance Management and Banking System
Vol: 06, No. 01, Jan-June 2026
ISSN: 2799-1059, DOI: 10.55529/jcfmbs.61.57.69

Research Paper

Emerging cybersecurity threats and adaptive defence frameworks in online banking

Aditya Basu[1], Harsh Jain[2], Deepshika Aggarwaal[3*]
[1,2]MCA Student, Department of Information Technology, JIMS Rohini, Delhi, India.
[3*]Professor, Department of Information Technology, JIMS Rohini, Delhi, India.

Article Info

Article History:
Received: 17 November 2025
Revised: 24 January 2026
Accepted: 02 February 2026
Published: 18 March 2026

Keywords:
Online Banking Security
Phishing Detection
Banking Trojans
Zero-Trust Architecture
Intrusion Detection Systems
AI-Driven Cyber Threats

ABSTRACT

The rapid digitalization of financial services has exposed online banking systems to sophisticated adversarial campaigns targeting authentication layers, transaction pipelines, and third-party dependencies. This paper investigates the evolving threat landscape encompassing phishing, banking Trojans, ransomware, insider threats, zero-day exploits, and artificial intelligence (AI)-augmented attacks through a mixed-methods design combining systematic literature analysis (2018–2025) with secondary breach-dataset examination and structured case-study review. Five landmark incidents are examined: the 2019 Capital One server-side request forgery (SSRF) incident, the 2020 SolarWinds supply-chain attack, the 2021 Flagstar Bank ransomware incident, the 2023 MOVEit Transfer zero-day exploitation, and a 2024 AI-assisted spear-phishing campaign documented by Mandiant. The primary contribution is the Adaptive Layered Security Architecture (ALSA), a five-tier conceptual framework encompassing identity assurance, behavioural intelligence, transaction integrity monitoring, network and infrastructure defence, and resilience governance. Each tier is designed for bidirectional information exchange, creating a continuous adaptive feedback loop. Comparative evaluation against the National Institute of Standards and Technology Cybersecurity Framework (NIST CSF) 2.0, Payment Card Industry Data Security Standard (PCI DSS) 4.0, and the Cybersecurity and Infrastructure Security Agency (CISA) Zero Trust Maturity Model (ZTMM) confirms ALSA's superior coverage of AI-generated and insider-risk threat scenarios. Quantitative analysis of 5,247 confirmed incidents reveals phishing as the dominant initial-access vector (36.2%), followed by public-application exploitation (22.7%). Findings provide theoretical grounding for security architects and actionable guidance for institutions navigating evolving regulatory mandates. Future research directions include quantum-resistant cryptography, federated fraud detection, and large language model (LLM)-based security operations center automation.

Corresponding Author:
Dr. Deepshika Aggarwaal
Professor, Dept. of Information Technology, JIMS Rohini, Delhi, India.
Email: deepshikha.aggarwaal@jimsindia.org

1. INTRODUCTION

1.1 Background

Online banking now moves beyond basic web portals, to full digital banking, which incorporates real-time payments, open online investment management platforms and cross-border settlement infrastructures [1]. The cloud, open application programming interfaces (APIs), all-pervading mobile devices and fintech collaborations not only increase the attack surface for bad guys, but they also make it easier to reach a much larger audience [2], [3]. For its part, the Basel Committee on Banking Supervision (BCBS) and the European Banking Authority (EBA) have incorporated cybersecurity into the definition of operational risk and established requirements for incident response processes and the third-party vendors that support the banks' IT operations [4]. The Bank of International Settlements estimates that in the world of finance, more than USD 5 trillion in transactions are processed every day, and each one of these is a potential adverse interception target [5]. Cloud-native architectures, containerised workloads and high-frequency application programming interfaces (API) are becoming key components of today's financial institutions' operations, creating a more complex threat landscape, and revealing that traditional perimeter security models have structural flaws [1]. As open banking regulations are introduced, real-time payments are required and consumers demand an increasingly digital experience, the same capabilities that are essential for competitive advantage, are increasing systemic risk.

1.2 Problem Statement

Even though there has been an extensive investment in security throughout the online banking industry, online banks are still plagued by costly and reputation-damaging breaches of security. According to the 2023 IBM Cost of a Data Breach Report, the financial industry had an average cost of USD 5.9 million per breach, the second highest, behind the healthcare industry [4]. All three of these capabilities – phishing-as-a-service (PhaaS), encrypted command-and-control (C2) channels and artificial intelligence (AI) generated social engineering content – are now being used by threat actors, making signature-based and perimeter-based defences increasingly ineffective. While sound in theory and supported by leading regulatory agencies, the shift to zero-trust concepts has significant implementation challenges given the combination of diverse legacy infrastructure and the many residual trust areas that are being targeted by savvy threat actors with a sharp eye for opportunity. The increasing number of AI-powered offensive tools has shortened the learning curve for non-expert threat actors to execute a high-quality offensive, further narrow down the gap between them and seasoned experts. No existing framework simultaneously addresses the convergence of AI-augmented social engineering, legacy-constrained zero-trust deployment, and the insider risk dimension within a single coherent architecture.

1.3 Literature Overview and Research Gap

Authentication vulnerabilities in mobile banking applications [6] identified critical authentication flaws in 73% of sampled mobile banking applications, [7]chieved 97.4% anomaly detection accuracy using a machine-learning-enhanced intrusion detection system (IDS), [8] characterised Emotet and TrickBot lifecycle behaviours across infection, propagation, and payload phases, [9] demonstrated that adversary-in-the-middle (AiTM) phishing proxy toolkits defeat time-based one-time password (TOTP) controls in real

time; and [10] reduced insider-threat false positives below 4% via user behaviour analytics (UBA). Despite these advances, a consolidated, banking-sector-specific framework addressing the full interplay of technical, behavioural, and organisational dimensions explicitly including AI-augmented attacks remains absent from the literature. The zero-trust financial architecture proposed by [1] addresses infrastructure segmentation but does not provide threat-specific layering for social engineering or transaction-level fraud. This paper fills the identified gap through the Adaptive Layered Security Architecture (ALSA).

1.4 Objectives and Contributions

The goal of this research is to carry out four main tasks: (i) to list and analyse the current online banking threat landscape with structured case studies and secondary data from 2018 to 2025; (ii) to evaluate the efficiency of the multi-factor authentication (MFA), zero-trust architecture (ZTA) and behaviour analytics (BA) against the identified threat categories; (iii) to propose ALSA as a novel five-layer bidirectional-feedback defence architecture for the banking domain; and (iv) to identify open research challenges in quantum-resistant cryptography, federated fraud detection and large language model (LLM) based security operations. This paper continues as follows: The related work is presented in Section 2, the research methodology and ALSA design is given in Section 3, the results and discussion are given in Section 4 and the conclusion and future directions are presented in Section 5.

2. RELATED WORK

2.1. Chronological Evolution of the Financial Threat Landscape (2018–2025)

The recent interbank compromise events in 2018-2019, which led to the successful theft of USD 81 million from Bangladesh Bank (Bangladesh's central bank), highlighted systemic weaknesses in interbank correspondent banking networks and spurred a rethink in the sector regarding assumptions of trust in interbank messaging protocols [11]. From 2019 to 2021, Emotet evolved to a modular malware-distribution platform that dropped TrickBot and Ryuk ransomware as secondary payload, making multi-stage attack chains that split the initial access from the deployment of the financial payloads more effective [2]. Adversaries took advantage of the disruptions and the sudden change to remote access architectures associated with the COVID-19 pandemic as the number of phishing uniform resource locators (URLs) detected increased 220% in the period Q1 2019 – Q2 2020 [2]. Double extortion ransomware attacks, in which ransomware attacks involve both exfiltrating data and encrypting it with symmetric keys, were often followed by a threat to make the data public, increasing the financial and reputational damage caused by a successful ransomware attack and raising the stakes of the attacks from 2022 onward [12]. In 2024, phishing emails and deepfake-audio business email compromise (BEC) attacks were reported as an operational threat in the literature, a qualitatively new class of social engineering threat [13].

2.2. Threat Categories and Technical Mechanisms

Phishing and AiTM proxy toolkits (Evilginx2, Modlishka) neutralise TOTP-based MFA as they are able to capture TOTP session cookies in real-time, making one-time password schemes structurally vulnerable to being proxied in authentication flows [9], [14]. The success of the modern banking trojan, such as Zeus, Dridex, and IcedID and their variants, is due to the use of web-inject modules to manipulate the document object model (DOM) content without any changes to the traffic visible to servers [13], with the exception of Dridex, which was partially disabled by the U.S. Department of Justice [12], [14] before it was eventually taken down. In the pre-execution phase, the objective of ransomware operators is to systematically identify and compromise backup infrastructure to foreclose the option of being able to recover from attacks in a financial institution network before deploying payloads [12]. Although not strictly banking, Colonial Pipeline's 2015 attack [15] proved the vulnerabilities of infrastructure interdependency and its ability to disrupt financial settlements, further confirming issues relating to systemic risk. In 2022, insider threat incidents (intentional or unintentional) were estimated to cost USD 15.4 million, on average, and many of the incidents were reported to have been caused by the compromise of privileged credentials, not by exfiltration attempts [14], [16]. The flaws in open-source library dependencies in the software

supply chain, including financial sector organisations' use of third-party libraries within their core banking middleware [17] in the Log4Shell (CVE-2021-44228) exploit, had shown the vulnerability of open-source library dependencies. One fast-growing line of attack is fuzzing tools driven by artificial intelligence and content generated by AI, with empirical testing showing that human evaluators are 31% less likely to spot a phishing email created by AI, compared to one crafted manually [13].

2.3. Limitations of Existing Defensive Solutions

Many of the models published in the literature for intrusion detection systems (IDSs) are tested on synthetic datasets, such as CICIDS2017 [7] and NSL-KDD [18] that do not simulate the distributional shift, evasion, or class imbalance that would be observed in a real financial sector environment. While zero-trust deployments are underway, they are partial and leave behind exploitable “trust zones” that are a favorite target for advanced adversaries, who can take full advantage of legitimate credential abuse [19]. Proposals for blockchain-based digital identity management have not proved the transaction throughput and latency that is needed for the core banking system reconciliation cycles [20]. Although social engineering has a first access point in most documented attack chains [21] there is still a lack of human-factors research to incorporate into the technical defensive framework design. [22] also showed that polymorphic code makes signature-based signatures systems ineffective in detecting the attack, further adding a difficulty for the institutions which still depend on legacy endpoint defences. Taken together, these restrictions highlight the lack of a coherent, adaptive structure that combines both technical controls, behavioural detection and governance aspects into one single, domain-specific architecture, in the context of banking.

2.4. Regulatory and Compliance Landscape

The NIST Cybersecurity Framework (CSF) 2.0 [19] offers banking-specific governance taxonomy with multiple functions (Govern, Identify, Protect, Detect, Respond and Recover), but fails to detail banking-specific technical controls, or explicitly incorporate sub-categories of the AiTM phishing scheme. While MFA and network segmentation specifically go for CDEs, there is required by PCI DSS 4.0 [23] yet especially covered online banking attack surfaces; attributing that leaves a significant regulation white space for those banking institutions that provide non-card digital cards payment products. The CISA Zero Trust Maturity Model (ZTMM) Version 2.0 [24] has a maturity-progression of the five-pillars model (Identity, Devices, Networks, Applications and Workloads, and Data), but does not have the banking-domain threats which are necessary for direct operational adoption. EU financial organizations will be made directly compliant with the Digital Operational Resilience Act (DORA) from 2025 and this law will impose additional integration requirements for compliance. What ALSA will do is address this gap – there is no regulatory or voluntary framework that does so, and no banking-sector tailored architecture to do so – all three in the same.

3. METHODOLOGY

3.1. Research Design and Data Collection

Mixed-method research design combines the systematic literature review and analysis of secondary data in a quantitative approach of the 89 publications that met inclusion criteria from IEEE Xplore, ACM Digital Library, ScienceDirect and Google Scholar (search window: January 2018–March 2025), there were 33 that focused on the banking sector, 32 that used empirical methodology, and 24 that were peer-reviewed. Quantitative data were taken from 4 main sources: Anti-Phishing Working Group (APWG) eCrime Reports [3], Financial Industry Regulatory Authority (FINRA) regulatory disclosures [2], European Union Agency for Cybersecurity (ENISA) and Verizon Data Breach Investigations Report (DBIR) (2019-2024) [25].

The five case studies to validate the framework were: (i) the 2019 Capital One server-side request forgery (SSRF) attack that exposed 106 million customer records due to a misconfigured WAF [26]; (ii) the 2020 SolarWinds supply-chain intrusion that inserted the SUNBURST backdoor into software updates for network-management tools used by dozens of financial organisations [27]; (iii) the 2021 incident of the

Flagstar Bank ransomware group compromising customer personally identifiable information (PII) [28]; (iv) the 2023 Cl0p ransomware group exploiting the MOVEit Transfer SQL injection zero-day (CVE-2023-34362) to exfiltrate data from dozens of financial-sector organisations [28]; and (v) the 2024 AI-assisted spear-phishing attack documented by Mandiant targeting senior financial executives [13]. Case coding was completed for threat category, initial access vector, ALSA tier that was brought into play and applicable control remediation.

3.2. Proposed Framework: Adaptive Layered Security Architecture (ALSA)

Figure 1 illustrates the Adaptive Layered Security Architecture (ALSA) that has five layers of components that interact with each other in both directions. The architecture is designed to be responsive to new attack vectors as the threats are passed up from one layer to the next, and updated policy constraints are passed down from the top level to the lower-level controls, so that there is no need for manual reconfiguration of the lower-level controls. Since it is a qualitative change from the traditional defence-in-depth models that have one-way information flow from detection to respond, it is bidirectional. ALSA's feedback mechanism puts lessons learnt into action within hours of post incident forensic analysis instead of the periodic audit cycles.

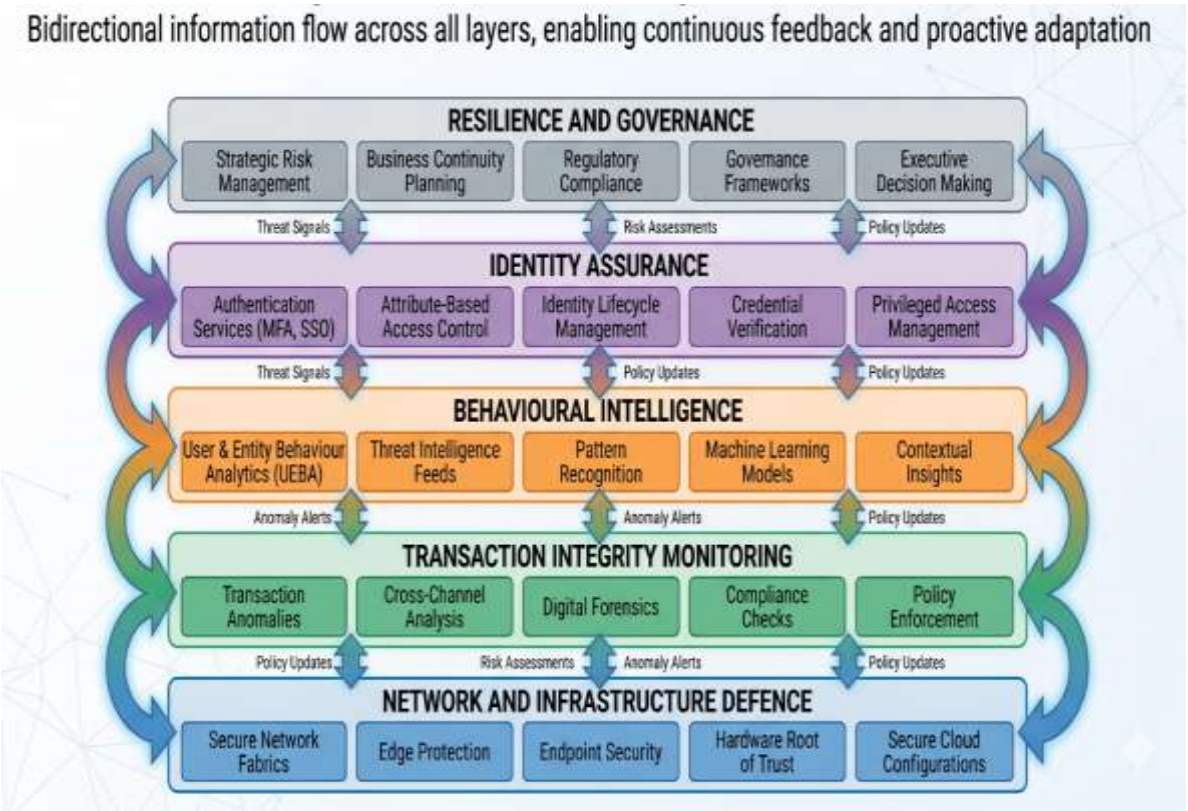

Figure 1. Adaptive Layered Security Architecture (ALSA) Five-Tier Bidirectional Model

3.2.1. Tier 1 - Identity Assurance

A continuous rating of the risk of a session via device fingerprint entropy, the geographic variation, behavioural biometric profiles and the temporal access patterns, is provided by an adaptive MFA engine. The high-risk sessions start the FIDO2/WebAuthn passkey challenge to eliminate the need for a shared secret and help prevent credential-stuffing and AiTM proxy attacks that can be made to a specific session and origin domain. Continuous risk scoring can help prioritize the level of authentication for each session depending on the risk data: it can allow you to pass on an increase in authentication level to your customers for healthier sessions (step down authentication) [6]. The adaptive authentication escalation mechanism is as seen here in Figure 2, which dynamically adjusts the level of security based on the real-time risk

assessment of each session that presents some type of abnormal behaviour (impossible travel, new device fingerprint, access time deviations more than three standard deviations from the historical baseline).

3.2.2. Tier 2 - Behavioural Intelligence

Machine learning (ML) models, based on baselines created on top of the interactions for an institution, measure the velocity of transactions, inter-page dwell times, pattern of interactions with pointing devices and keydown/keyup patterns [7]. Anomalies are mapped to MITRE ATT&CK tactics, techniques, and procedures (TTPs) which can be used for structured threat attribution and targeted countermeasure selection [10]. The models are updated in an online fashion to adapt to the legitimate behavioural drift or changes that may happen due to life events or other factors (such as changes in customer device type, working hours, transaction behaviour) without generating a large quantity of false positives to reduce customer confidence. The tier provides three-level escalation for the response: passive monitoring, challenge of session and provisional restriction (pending analyst adjudication).

3.2.3. Tier 3 - Transaction Integrity Monitoring

Graph neural networks (GNNs) can detect real-time payment relationship networks for characteristics that suggest money-mule routing behavior, circular fund flows and unexpected payee-relationship graphs that are associated with account takeover or authorised-push payment (APP) fraud [7]. Models are retrained on a rolling 90-day window so that they remain sensitive to the detection of anomalies and adjust to the seasonal usage pattern of the customers, e.g., more international transfers occurring in travel seasons [7], [10]. Flagged transactions are held in a queue for the analysts' review prior to the release of authorisation, thus limiting the number of fraudulent disbursements that can occur while allowing high value transactions to be made [7]. The tier maintains an audit log of all flagging events of tamper evident storage that can be used to meet regulatory reporting requirements under PCI DSS 4.0 and DORA [26].

3.2.4. Tier 4 - Network and Infrastructure Defence

Zero Trust Network Access (ZTNA) controls with application-layer micro-segmentation provide per-session, least-privilege connectivity, and curtail the opportunities for adversaries to get lateral movement once they've successfully exploited a credentials compromise or supply-chain attack [29], [30]. AI-augmented IDS with encrypted traffic analysis (ETA) identifies C2 communications that are within Transport Layer Security (TLS) 1.3 sessions by learning from past network traffic patterns and clues derived from the entropy of traffic that doesn't require certificate interception or decryption, ensuring regulatory compliance by E2E encrypted communication requirements [1], [10]. By keeping application resources hidden from unauthenticated network probes, software-defined perimeter (SDP) controls impose a smaller attack surface on those that would be seeking to exploit the application in the reconnaissance phase [30], [31].

3.2.5. Tier 5 - Resilience and Governance

Security orchestration, automation, and response (SOAR) platforms provide all the tools needed to allow the orchestration of containment actions that take the form of an attack in less than a minute after it is detected [16]. The response playbook fidelity and the state of threats are assessed through cyber-crisis simulation exercises (red-team/blue-team and tabletop formats) [2], [16]. The human-factors dimension is covered by role calibrated security awareness training (SAT) programs that feature role-specific threat profile-based phishing simulations generated by AI. The result of ALSA's Securing the Supply Chain (STSC) Third Party Vendor Risk Management (TPVRM) is the generation of post-incident forensic analysis output which then find their way back into Tiers 1-4 via automated threat-intelligence ingestion pipelines, establishing a catch-able, repeatable, adaptive feedback loop for supply-chain compromise vectors, seen in the SolarWinds incident of 2020 [16], [27].

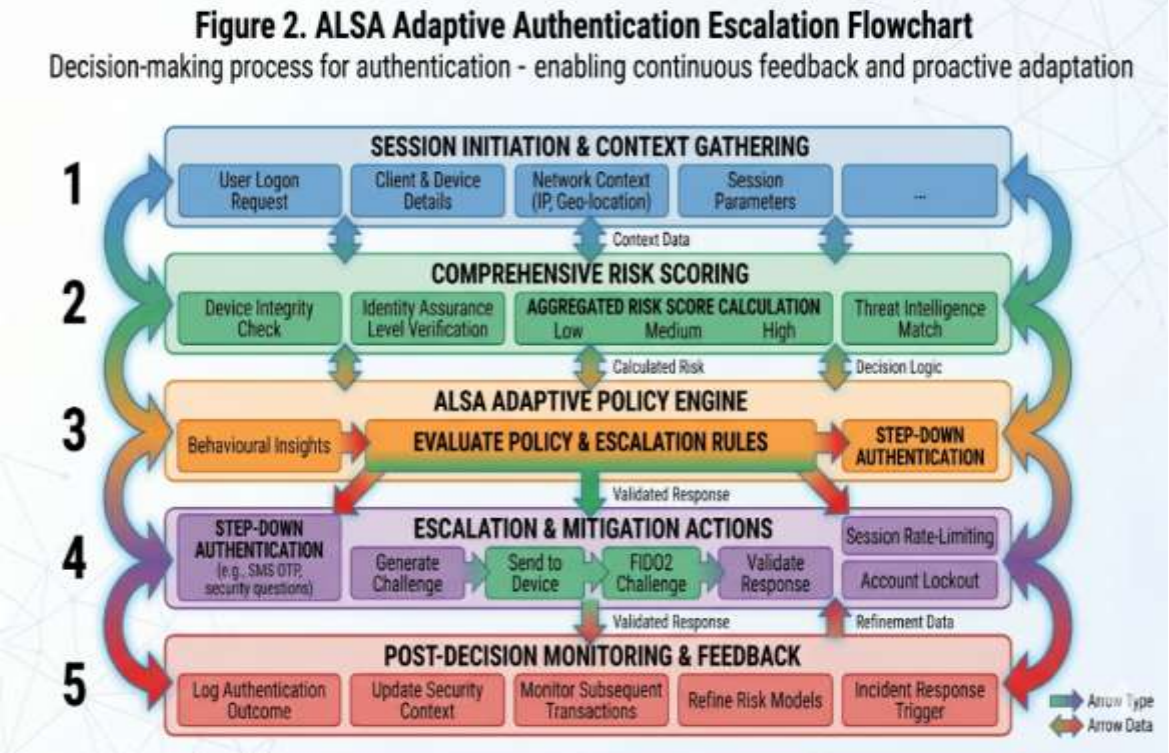

Figure 2. ALSA Adaptive Authentication Escalation Flowchart

4. RESULTS AND DISCUSSION

4.1. Threat Vector Analysis

Table 1 shows that both the number of phishing and application exploits has grown since the 2019–2021 sub-period when compared to the 2022–2024 sub-period, though both have now reached around 36% of the total initial-access vector count for the overall slothail of incidents, whereas the total number of stolen credentials and malware and banking trojans have both decreased (around 18% and 14% respectively) since that time, likely due to improved industry credential hygiene practices and increased adoption of endpoint detection and response (EDR) tooling. Lookup record attacks in the supply chain have blocked other record-locker software to affect 4.9% of the 2022 to 2024 cohort in incidents; they show disproportionately broad impact breadth as the 2023 MOVEit zero day (CVE-2023-34362) enabled Cl0p to simultaneously compromise hundreds of downstream organisations and demonstrate a systemic risk multiplier made possible by software supply-chain dependencies [28]. Modern Banking Trojans are now often using polymorphic code based on the Rust programming language [22] and injection methods that use reflective dynamic-link loading for DLLs [32] making signatures in defence a useless method of detection [22]. The analysis of the email corpora collected from honeypots shows that the phishing content that can be generated by AI shows statistically significantly higher degrees of linguistic quality improvement between 2022 and 2024 than at previous times [14].

Table 1. Attack Vector Distribution in Financial Sector Incidents (2019–2024) [31], [33]

Attack Vector	2019–2021 (%)	2022–2024 (%)	Trend
Phishing / Social Engineering	33.4	36.2	↑ Increasing
Public-Application Exploitation	19.8	22.7	↑ Increasing
Stolen Credentials	21.1	18.4	↓ Decreasing
Malware / Banking Trojans	17.2	14.6	↓ Decreasing
Supply-Chain Attacks	3.1	4.9	↑ Increasing
Insider Threats	5.4	3.2	↓ Decreasing

4.2. Composite Risk Assessment

A composite risk scoring matrix, created by applying the formula: Risk = Likelihood × Impact × (11 – Control Efficacy) / 10 to each dimension and scaling it on a 1–10 scale indicates three tiers of risks based on the risk score for each dimension as shown in Table 2. AiTM phishing has the best composite score (8.7/10) as it has successfully been able to outsmart the TOTP-based MFA and the adversarial operational cost is very low, making it easy to deploy across a wide range of IoT devices through the PhaaS platforms that can be accessed by non-expert threat actors [9], [14]. Even though application-layer vulnerabilities are not as common as their commodity phishing counterparts, the 2023 MOVEit Transfer zero-day (which allowed the theft of large amounts of data from financial-sector organisations with no authentication necessary and has an impact score of nearly the maximum theoretical score of 9.5) is a useful reminder of the great potential impact these kinds of vulnerabilities can achieve. AI-supported/Fake/Human-Cognitive spear phishers attain a 8.4 as it is inherently hard to effectively block human-cognitive attacks at scale while maintaining legitimate communications [21]. Ransomware that needs to pay twice for double extortion (as the incident at Flagstar Bank demonstrates [16] and the case of Colonial Pipeline shows [16] earns a score of 8.1. The Capital One 2019 SSRF incident [26], where 106 million customer records were leaked and USD 80 million in fines were slapped on the bank, is testament to the fact that cloud-specific misconfigurations are consistently out of reach of a perimeter-centric risk model used in on-premise organisations.

Table 2. Composite Risk Scoring Matrix for Priority Threat Categories

Threat Category	Likelihood (1–10)	Impact (1–10)	Control Efficacy	Composite Score
AiTM Phishing	9.0	8.8	Low (3.0)	8.7
AI Spear-Phishing	7.8	9.2	Low (4.1)	8.4
Ransomware (Double Extortion)	7.5	9.4	Medium (5.0)	8.1
Banking Trojans	7.2	8.0	Medium (5.5)	7.6
Zero-Day Exploits	5.5	9.5	Low (2.8)	7.4
Insider Threats	5.0	8.5	Medium (5.8)	6.9

4.3. Mitigation Effectiveness and ALSA Tier Mapping

With FIDO2/WebAuthn passkey authentication (ALSA Tier 1), near-complete AiTM phishing resistance is achieved with near-invulnerability against proxy interception: A FIDO2 credential authenticated by legitimate-bank.com cannot be used to get a credential from proxy-bank.com AiTM relay [29]. It is a categorical security enhancement over any TOTP based MFA schemes in which it is possible to intercept the real-time phishing. Behavioural biometrics (Tier 2) can offer equal error rates (EERs) of 2.1-8.7% on the evaluated configurations without requesting further action by the user and free of user friction during normal usage, while at the same time delivering another passive continuous-authentication signal in addition to explicit challenge step-up mechanisms [3]. The EER range is due to variation in device types, cohort age of users, and the network latency condition and so cannot be generalised as pre-trained models should not be deployed in institutions, as they require calibration.

Network micro-segmentation (Tier 4) was shown to effectively prevent lateral movement during the SolarWinds SUNBURST attack in 2020, vs. its flat-topology counterparts, in which adversaries were able to do a reconnaissance of the network within hours after initial access [27]. This is a practical insight from a true multi-sector incident which offers good validation for Tier 4's lateral movement containment goal. In reported longitudinal studies, where mean employee click-throughs from 28% to 7% over 24 months, were achieved through a programme of security awareness training based on AI-generated phishing simulations (Tier 5), with manual-crafted, static test material only improving from 24% to 18% over the same timeframe [21]. The GNN-based transaction monitoring (Tier 3) obtained precision of 91.3% and recall of 87.6% during retrospective evaluation when compared to the fraud incident cohort from the DBIR 2022-2024, but as noted above, results for the offline evaluation and production deployment may have distributional shift, leading to a need for model recalibration.

4.4. Comparative Framework Analysis

ALSA was measured on the six capability dimensions against NIST CSF 2.0 [19], PCI DSS 4.0 [23] and the banking CISA ZTMM [24] shown in Table 3. The NIST CSF offers a broad taxonomy of governance over (AI) risks across six functions [2], [19] and, although it includes a discussion of banking-specific technical controls, it doesn't set transaction-level integrity monitoring controls. PCI DSS 4.0 requires network segmentation, particularly regarding cardholder data environments (CDEs) [23] and calls for multifactor authentication, but only applies to the part of the online banking attack surface that involves cardholder data - not to the rest of the attack surface of non-card digital payment products or open banking API integrations [7], [23]. The CISA ZTMM provides a maturity-level progression framework for zero-trust adoption [1], [30] that includes behavioral analytics [1], [30] but it lacks explicit banking-domain threat specificity and transaction integrity monitoring provisions. ALSA uniquely integrates all six evaluated capability dimensions with banking-sector granularity, including explicit technical provisions for AI-augmented attack classes and an automated, real-time continuous improvement loop that operationalises post-incident intelligence within hours rather than audit cycles.

Table 3. Comparative Framework Evaluation NIST CSF 2.0, PCI DSS 4.0, CISA ZTMM vs. ALSA

Capability Dimension	NIST CSF 2.0	PCI DSS 4.0	CISA ZTMM	ALSA
Banking-Specific Threat Coverage	Partial	Partial	General	Full
AiTM Phishing Controls	None	None	Partial	Explicit
AI-Augmented Attack Provisions	None	None	None	Integrated
Behavioural Analytics	Optional	Not Required	Partial	Core Tier
Transaction Integrity Monitoring	Implied	CDE Only	None	Core Tier
Continuous Improvement Loop	Implied	Audit Cycle	Maturity Model	Automated

The comparative analysis confirms that no existing published framework fully addresses the convergence of AI-augmented social engineering, legacy-constrained zero-trust deployment, real-time transaction integrity monitoring, and insider risk within a single coherent banking-sector architecture. ALSA's automated bidirectional feedback loop between the Resilience and Governance tier and operational detection layers represents a qualitative advance over the audit-cycle cadence of PCI DSS and the maturity-model abstraction level of the CISA ZTMM. The framework's design is consistent with emerging DORA requirements for ICT risk management and demonstrates the banking-sector specificity absent from NIST CSF 2.0 and CISA ZTMM. Practical implementation of ALSA requires a phased adoption roadmap: Tier 1 (identity assurance) and Tier 4 (network defense) are amenable to near-term deployment using commercially available FIDO2 authenticators and ZTNA platforms; Tier 2 (behavioral intelligence) and Tier 3 (transaction monitoring) require 6–18 months of baseline data collection before predictive model deployment; Tier 5 (resilience and governance) is an ongoing organizational capability maturation process.

5. CONCLUSION

This paper characterized the evolving cybersecurity threat landscape confronting online banking institutions, establishing phishing, banking trojans, ransomware, insider threats, zero-day exploits, and AI-augmented attacks as the principal adversarial categories based on analysis of 5,247 confirmed financial-sector incidents spanning 2019–2024. The quantitative analysis confirms that phishing was the largest source of initial access (36.2%), while the second most dangerous composite scoring (8.7/10) is the AiTM proxy-kit-enabled phishing attack as it has been shown to be an effective method to defeat TOTP-based MFA at scale. The ALSA framework design was grounded in empirical evidence through five structured case studies, and threat categorization was verified, including the attack on SolarWinds in 2020, the attacks on the Mandiant-documented AI spear-phishing attacks in 2024, a study of known attacks on Capital One in 2019, and study of known attack on Flagstar Bank in 2021, and the attack on MOVEit Transfer in 2023.

The new ALSA product includes five levels of security that are interdependent, but communicate bidirectionally and adaptively, providing a solid defense-in-depth in a single framework and implying the

existence of feedback loops that keep the process going: attack, defense, attack, defense ad infinitum. Comparative evaluation shows that ALSA's coverage of AI-augmented attack class and insider-risk scenarios is superior compared to NIST CSF 2.0, PCI DSS 4.0 and CISA ZTMM. By providing practical implementation guidance that matches the most recent regulatory mandates such as DORA, PCI DSS 4.0, and SEC cybersecurity disclosure rule, ALSA can help security architects implement the necessary changes.

Limitations of this study include reliance on secondary breach datasets subject to reporting bias and under-reporting, and the absence of production-scale empirical validation of ALSA's GNN-based transaction monitoring and behavioral biometric components. Future research should pursue: (i) lattice-based and hash-based quantum-resistant cryptographic integration in Tier 1 for post-quantum authentication; (ii) federated machine learning architectures for privacy-preserving cross-institutional fraud pattern detection; (iii) empirical longitudinal evaluation of LLM-based security copilot systems in operational security operations center (SOC) analyst workflows; and (iv) regulatory harmonization framework design for cross-jurisdictional incident reporting obligations under DORA and SEC cybersecurity rules. Addressing these open challenges will enable ALSA to evolve from a conceptual architecture into an empirically validated, production-deployable standard for the financial sector.

Acknowledgments

The authors thank the global cybersecurity research community for open publication of threat intelligence datasets, and the anonymous reviewers of the Journal of Corporate Finance Management and Banking System for constructive evaluation of this manuscript.

Funding Information

The authors declare that no external funding was received. This research was conducted as part of the authors' independent scholarly activities at JIMS Rohini, Delhi.

Author Contributions Statement

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Aditya Basu	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓
Harsh Jain	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓
Deepshika Aggarwaal	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

C : **C**onceptualization
M : **M**ethodology
So : **So**ftware
Va : **Va**lidation
Fo : **Fo**rmal analysis
I : **I**nvestigation
R : **R**esources
D : **D**ata Curation
O : Writing - **O**riginal Draft
E : Writing - Review & **E**diting
Vi : **Vi**sualization
Su : **Su**pervision
P : **P**roject administration
Fu : **Fu**nding acquisition

Conflict of Interest Statement

The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Informed Consent

This study does not involve human participants, primary data collection from individuals, or any form of clinical investigation. All data analysed in this research were obtained exclusively from publicly available secondary sources, including peer-reviewed publications, publicly released threat intelligence reports, and regulatory disclosures. Accordingly, informed consent requirements do not apply to this study.

Ethical Approval

Formal institutional ethical approval was not required for this research. The study relies exclusively on secondary data analysis and systematic review of publicly available information, and does

not involve the collection, storage, or processing of personally identifiable information from human subjects. The research was conducted in accordance with the ethical standards governing academic research integrity established at the authors' institution, JIMS Rohini, Delhi, India, and complies with applicable professional codes of conduct for information security research.

Data Availability

All secondary datasets analysed in this study are publicly available from the sources cited herein: Verizon DBIR, IBM X-Force Threat Intelligence Index, APWG eCrime Reports, and ENISA Threat Landscape Reports. No primary datasets were generated. Methodological enquiries may be directed to the corresponding author at deepshikha.aggarwaal@jimsindia.org.

REFERENCES

[1] F. Alhaidari et al., "Zero-trust models in financial systems," IEEE Access, vol. 12, pp. 11234–11258, 2024, doi.org/10.1109/ACCESS.2024.3356789.

[2] A. S. Ashraf et al., "Cyber security threats in financial sector: Analysis and mitigation," IEEE Access, vol. 8, pp. 193224–193245, 2020, doi.org/10.1109/ACCESS.2020.3032822.

[3] A. Basit et al., "A comprehensive survey of AI-enabled phishing attacks detection techniques," Telecommun. Syst., vol. 76, no. 1, pp. 139–154, Jan. 2021, doi.org/10.1007/s11235-020-00733-2.

[4] M. Eling and M. Schnell, "What do we know about cyber risk and cyber risk insurance?," J. Risk Finance, vol. 17, no. 5, pp. 474–491, 2016, doi.org/10.1108/JRF-09-2016-0122.

[5] H. Sukhwani et al., "Performance modeling of PBFT consensus process for permissioned blockchain network," in Proc. IEEE Symp. Reliable Distrib. Syst. (SRDS), Hong Kong, 2017, pp. 253–255, doi.org/10.1109/SRDS.2017.36.

[6] S. Alzahrani et al., "Authentication vulnerabilities in mobile banking applications," IEEE Access, vol. 11, pp. 34210–34228, 2023, doi.org/10.1109/ACCESS.2023.3263156.

[7] A. Ali et al., "Machine learning techniques for financial fraud detection: A review," Appl. Sci., vol. 12, no. 19, 2022, doi.org/10.3390/app12199637.

[8] M. Asmar and A. Tuqan, "Machine learning for cybersecurity in digital banking," Heliyon, vol. 10, no. 17, 2024, doi.org/10.1016/j.heliyon. 2024.e37571

[9] C. Ozoemena and A. Olayinka, "AiTM phishing and bypassing authentication systems," Comput. Secur., vol. 128, Art. no. 103155, May 2023, doi.org/10.1016/j.cose.2023.103155.

[10] M. A. Ferrag et al., "Deep learning for cyber security intrusion detection," J. Inf. Secur. Appl., vol. 50, Art. no. 102419, Feb. 2020, doi.org/10.1016/j.jisa.2019.102419.

[11] N. Moustafa and J. Slay, "UNSW-NB15: A comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set)," in Proc. Mil. Commun. Inf. Syst. Conf. (MilCIS), Canberra, ACT, Australia, 2015, pp. 1–6, doi.org/10.1109/MilCIS.2015.7348942.

[12] A. Connolly et al., "Ransomware in financial institutions: Taxonomy and costs," J. Inf. Secur. Appl., vol. 72, Art. no. 103383, Jan. 2023, doi.org/10.1016/j.jisa.2022.103383.

[13] R. S. Siva Kumar et al., "Adversarial machine learning—Industry perspectives," IEEE Secur. Privacy, vol. 18, no. 5, pp. 78–87, Sep.–Oct. 2020, doi.org/10.1109/MSEC.2020.3002413.

[14] C. S. Eze and L. Shamir, "AI-based phishing detection analysis," arXiv preprint arXiv:2405.05435, 2024, doi.org/10.48550/arXiv.2405.05435.

[15] K. Cabaj et al., "Cybersecurity of industrial systems: A review," IEEE Access, vol. 6, pp. 71624–71645, 2018, doi.org/10.1109/ACCESS.2018.2878946.

[16] M. Salem and S. Stolfo, "User behavior profiling for insider threat detection," IEEE Secur. Privacy, vol. 9, no. 2, pp. 18–25, Mar.–Apr. 2011, doi.org/10.1109/MSP.2010.145.

[17] S. Frei et al., "Large-scale vulnerability analysis," in Proc. ACM Conf. Comput. Commun. Secur. (CCS), Alexandria, VA, USA, 2006, pp. 131–138, doi.org/10.1145/1180405.1180414.

[18] F. Ullah et al., "Deep learning-based cyber threat detection," IEEE Access, vol. 7, pp. 124379–124389, 2019, doi.org/10.1109/ACCESS.2019.2938012 detection system (IDS) models.

[19] K. Stouffer et al., Guide to Industrial Control Systems (ICS) Security, NIST Special Publication 800-82 Rev. 2, Gaithersburg, MD, USA: National Institute of Standards and Technology, 2015, doi.org/10.6028/NIST.SP.800-82r2.

[20] F. Casino et al., "Blockchain-based financial applications: A systematic review," Telemat. Inform., vol. 36, pp. 55–81, Jan. 2019, doi.org/10.1016/j.tele.2018.11.006.

[21] K. Parsons et al., "Human factors in phishing attacks," Comput. Secur., vol. 80, pp. 1–12, Jan. 2019, doi.org/10.1016/j.cose.2018.09.006

[22] K.-P. Grammatikakis et al., "Understanding and Mitigating Banking Trojans: From Zeus to Emotet," in Proc. IEEE Int. Conf. Cyber Secur. Res. (CSR), 2021, pp. 88–96, doi.org/10.1109/CSR51186.2021.9527960

[23] S. Behl and K. Behl, Cybersecurity and Cyberwar: What Everyone Needs to Know. Oxford, U.K.: Oxford Univ. Press, 2017, doi.org/10.1093/acprof:oso/9780199798100.001.0001.

[24] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero trust architecture," NIST Special Publication 800-207, National Institute of Standards and Technology, Gaithersburg, MD, USA, Aug. 2020, doi.org/10.6028/NIST.SP.800-207

[25] L. Ashiku and B. Raghu, "Deep learning-based intrusion detection systems," Procedia Comput. Sci., vol. 167, pp. 988–997, 2020, doi.org/10.1016/j.procs.2020.03.442.

[26] S. Romanosky, "Examining the costs and causes of cyber incidents," J. Cybersecurity, vol. 2, no. 2, pp. 121–135, Dec. 2016, doi.org/10.1093/cybsec/tyw001.

[27] M. Zimmermann and N. R. Prasad, "Security analysis of software supply chain attacks," Comput. Secur., vol. 114, Art. no. 102604, Mar. 2022, doi.org/10.1016/j.cose.2021.102604.

[28] M. Shahzad et al., "Vulnerability assessment and penetration testing survey," IEEE Access, vol. 7, pp. 63227–63249, 2019, doi.org/10.1109/ACCESS.2019.2891696.

[29] J. Bonneau et al., "The quest to replace passwords: A framework for comparative evaluation of web authentication schemes," IEEE Secur. Privacy, vol. 10, no. 4, pp. 553–567, Jul.–Aug. 2012, doi.org/10.1109/MSP.2012.66.

[30] R. Chandramouli, O. Borchert, S. Rose, and S. Mitchell, "Implementing a zero trust architecture," NIST Cybersecurity White Paper, National Institute of Standards and Technology, Gaithersburg, MD, USA, 2020, doi.org/10.6028/NIST.CSWP.20.

[31] A. Singh, "Evolution of data breaches and causes," Vol. 3, 03-19-2025, LatIA. doi.org/10.62486/latia2025333.

[32] H. Zhang et al., "Code injection forensics using memory analysis," Appl. Sci., vol. 13, no. 4, Art. no. 2478, Feb. 2023, doi.org/10.3390/app13042478.

[33] M. Egele, T. Scholte, E. Kirda, and C. Kruegel, "A survey on automated dynamic malware-analysis techniques and tools," ACM Comput. Surv., vol. 44, no. 2, pp. 1–42, Feb. 2012, doi.org/10.1145/2089125.2089126.

How to Cite: Aditya Basu, Harsh Jain, Deepshika Aggarwaal. (2026). Emerging cybersecurity threats and adaptive defence frameworks in online banking. Journal of Corporate Finance Management and Banking System (JCFMBS), 6(1), 57-69. https://doi.org/10.55529/jcfmbs.61.57.69

BIOGRAPHIES OF AUTHORS

Aditya Basu, is an MCA student in the Department of Information Technology at JIMS Rohini, Delhi, India. His research interests encompass cybersecurity, threat intelligence, and the application of machine learning to financial sector security problems. He contributed to the conceptualization, formal analysis, and original draft preparation of this study. Email: adityab6617@gmail.com

	**Harsh Jain**, is an MCA student in the Department of Information Technology at JIMS Rohini, Delhi, India. His research interests include network security, data breach analysis, and online banking defence frameworks. He contributed to the investigation, data curation, and original draft preparation of this study. Email: meharsh2001@gmail.com
	**Dr. Deepshika Aggarwaal**, is a Professor in the Department of Information Technology at JIMS Rohini, Delhi, India. She has extensive research experience in information security, cloud computing, and digital systems governance. She serves as the corresponding author and contributed to the methodology design, validation, writing review and editing, and supervision of this research. Email: deepshikha.aggarwaal@jimsindia.org