

RESEARCH ARTICLE

 **OPEN ACCESS** | Double-Blind Peer Reviewed

Enhancing Cybersecurity in Financial Institutions Through Zero Trust and Cloud Computing Technologies: A Review of Architectures, Standards, and Integration Practices

Jimmy Ben¹, Bowie Dakari^{2*}¹ Department of Computer Science, Harvey Mudd College, USA² Department of Computer Science, Harvey Mudd College, USA

*Corresponding author: Bowie Dakari | Email: bdakari.hmc.edu

ARTICLE INFO

Article History

Received: 20 May 2026**Revised:** 16 June 2026**Accepted:** 14 July 2026**Published:** 05 August 2026

Keywords

zero trust architecture, cloud computing, financial institutions, cybersecurity, identity and access management, micro-segmentation, shared responsibility model, NIST SP 800-207, cloud security posture management, operational resilience

How to Cite

J. Ben, "Enhancing cybersecurity in financial institutions through zero trust and cloud computing technologies: A review of architectures, standards, and integration practices," JATICS, vol. 2, no. 2, pp. 1-14, 2026.

ABSTRACT

At the same time as the old network boundary has vanished due to remote working, open banking, and third-party integration, financial institutions are shifting critical and customer-facing operations from local to public, private, and hybrid cloud environments. As a result of these parallel moves, perimeter based defense is becoming more and more ineffective, and zero trust architecture (ZTA) is becoming an operational must. This review collates peer-reviewed literature, international standards and regulatory and industry guidance, and explores the potential of integrating zero trust and cloud computing technologies to bolster the cybersecurity posture of financial institutions. It describes the threat landscape anytime a cloud layer is added to the environment, as well as the documented cost of breach from both cloud and non-cloud environments, and reviews the logical components of zero trust architecture as documented in NIST SP 800-207 and the core enabling technologies—identity and access management, micro-segmentation, and continuous monitoring; it explores the shared responsibility model across infrastructure, platform, and software-as-a-service environments; and examines regulatory instruments that increasingly presume both technologies. The review suggests that financial institutions adopt an integrated Zero-Trust Cloud Security model, consisting of a combination of identity, device, network, and data controls – all of which are layered across the shared responsibility boundary and continuously checked instead of once at the perimeter. Researchers examined industry breach cost studies that revealed the cost of breaches was materially lower for organizations with mature zero trust deployments and better-governed cloud environments, while also acknowledging significant implementation hurdles such as legacy core-banking integration, identity sprawl and latency sensitive transactions. The review ends with a roadmap of empirical validation, maturity measurement standardization and governance of AI features increasingly integrated into both zero trust and cloud security tooling.

INTRODUCTION

The perimeter approach is a traditional way of securing financial systems which has been used over the years: once users, devices, and applications have been authenticated, they are assumed to be safe within the perimeter of the network. This model has been undermined by two simultaneous changes. First, cloud computing has shifted the center of gravity of core banking, payments, and analytics workloads to cloud infrastructure-, platform-, and software-as-a-service offerings in a shared responsibility model where the institution takes responsibility for certain layers of the stack and the provider for others [1]. Second, employees, partners, and customer applications access systems outside of any fixed network boundary on a regular basis due to remote working and the use of open banking APIs and greater third-party integration. The Basel Committee's analysis of financial-technology developments explicitly identifies cloud reliance and "third-party-dependence" as financial institutions' risks that supervisors and institutions should be aware of and manage as digital transformation continues [2].

In response, zero trust architecture (ZTA) takes a different approach by no longer relying on implicit (based on location) trust and instead continuously verifies the risk of each access attempt, whether it comes from within or outside the traditional boundary of the enterprise perimeter. Kindervag coined the term in his argument for networks to treat all traffic as untrusted until it can be proven otherwise [3] and the approach was later formalized by the U.S. National Institute of Standards and Technology in NIST SP 800-207, which outlines the logical elements, deployment models and threats surrounding zero trust [4]. The essence of zero trust is that there is no single network perimeter to trust, making it an ideal fit with the distributed nature of cloud-based workloads and the fact that the institution does not own or manage the underlying infrastructure.

The risks are highlighted by industry incident reporting. According to the 2024 IBM Cost of a Data Breach Report, the average cost of a breach involving data stored in public cloud storage is USD 5.17 million, which is higher than the global average across all environments (USD 4.88 million) [5]; separately, IBM analysis found that the average cost of a breach is USD 4.83 million less with deployed zero trust architecture than without it [5]. Exploitation of vulnerabilities is also the top initial access vector, as seen in the 2024 Verizon Data Breach Investigations Report, which revealed that the number of attacks on Internet-facing assets and cloud-hosted assets has more than doubled year over year [6].

Though there is clear synergy between zero trust and cloud, the existing literature typically considers them independently; literature on zero trust focuses on enterprise networks [7] [8] without a consistent discussion of identity-centric access control, and literature on cloud security focuses on confidentiality, integrity, and availability without a systematic discussion of identity-centric access control [9] [10]. Even fewer examples are found in the literature specific to financial institutions, in which regulations, outdated core-banking systems and the need to process transactions in real-time pose additional challenges to adoption. This review fills in this void. Its contributions are:

- a synthesis of the threat landscape and loss evidence that motivates the joint adoption of zero trust and cloud security controls in financial institutions;
- a structured account of zero trust architecture, its NIST-defined logical components (Fig. 3), and its principal enabling technologies;
- an analysis of the cloud shared responsibility model and how zero trust principles map onto it across service models (Fig. 2);
- a review of the regulatory and standards landscape that increasingly presupposes both technologies; and
- an integrated Zero-Trust Cloud Security framework for financial institutions (Fig. 5), together with propositions for future research and a practitioner-oriented implementation agenda.

The rest of the article follows. In section II, the review method is introduced. In Section III, the characteristics of the cloud-driven threat landscape and loss evidence are discussed. Section IV provides an overview of cloud computing service and deployment models, and a shared responsibility model. In Section V, we're going to look at zero trust architecture and the logical elements of it and then in Section VI, we're going to look at the technologies that enable it. Section VII reviews the regulatory context, Section

VIII outlines an integrated framework and implementation guidance, Section IX discusses open challenges, and Section X concludes.

The concept of zero trust and cloud is that they complement each other because they solve each other's problems and they have different owners and cost lines. Cloud computing provides elasticity, patching managed by another party, and built-in security services, all of which most institutions could not replicate (at the same cost) on premises, but it does so by moving portions of this control boundary away from the institution and into the cloud. But zero trust provides a model for securing access through just this kind of shared boundary that is distributed, and it relies on a deep and detailed understanding of identity, device, and behavior that's often more readily available in modern cloud-native platforms than from older on-premises systems. When paired up, the two technologies will provide a more complete solution; when used independently, an institution can find itself with a strong cloud infrastructure behind a perimeter model that the cloud infrastructure already makes redundant, or a zero trust policy engine with no reliable telemetry it can reason over.

REVIEW APPROACH

The narrative design of this review is structured according to the lines of an art of research that is the field of a security research, with its line of evidence extending from peer-reviewed security research, to formal architecture specifications, to international standards, to regulatory instruments. These sources included peer-reviewed journal articles and conference papers on zero trust and cloud security, foundational and current technical standards and specifications, financial-sector regulatory and supervisory guidance, and industry breach-cost and threat studies that provided disclosed methodology. The scope and selection criteria applied are outlined in Table I.

Criterion	Description
Topical scope	Zero trust architecture, cloud computing security, and their combined application in financial institutions.
Evidence types	Peer-reviewed literature, standards and technical specifications, regulatory and supervisory instruments, and industry breach or incident studies.
Inclusion	Direct relevance to identity-centric access control, cloud shared-responsibility security, or financial-sector cyber regulation; transparent methodology or authoritative issuing body.
Exclusion	Vendor material without disclosed methodology; opinion content without supporting evidence; superseded editions of a cited standard.
Analytical lens	Thematic synthesis across four themes: threat landscape, cloud architecture and shared responsibility, zero trust architecture and enablers, and regulation.
Temporal focus	Emphasis on 2010–2025, spanning the origin of zero trust to the most recent applicable EU and NIST instruments.

Table I: Scope And Selection Criteria Of The Review

As with any synthesis that draws on industry breach studies, the loss figures cited here reflect self-reported and vendor-collected samples rather than a census of all incidents, and should be read as indicative of direction and order of magnitude rather than as precise population parameters [5], [6].

THE CLOUD-DRIVEN THREAT LANDSCAPE IN FINANCIAL SERVICES

A. WHY THE PERIMETER MODEL FAILS IN CLOUD-NATIVE FINANCE

Financial institutions are operating more workloads on top of multiple clouds (public cloud, private cloud and on-premises infrastructure), multiple applications, and even different parts of the same applications on top of different clouds. This distribution invalidates two premises for perimeter security, that resources are located within a perimeter and that a user or device within a perimeter can be trusted during the course of the session. This is not the case if core banking is hosted on non-institutional owned infrastructure or if employees log in from their home networks or if APIs present account functionality directly to third-party applications [1], [2].

B. EVIDENCE ON BREACH COST BY ENVIRONMENT

In Fig. 1(a), the cost of breaches with public cloud data is plotted against the average reported cost for breaches across environments for 2024. An average breach of data in public cloud storage cost USD 5.17

million, a 6% increase from the global average breach cost of USD 4.88 million across all environments, and shows a positive trend that has been observed in breaches of data in the cloud space for several years [5]. A complementary comparison from the longitudinal study of zero trust deployment at IBM is presented in Fig. 1(b): Organizations with no zero trust architecture had an average cost of breaches of USD 5.04 million, while organizations with a deployed zero trust architecture had an average cost of breaches of USD 3.28 million, representing a difference of approximately 35% [5]. These are association-based industry findings rather than controlled experiments, and the maturity of an organization in overall security is likely to vary by organization that implements zero trust, but the numbers do align with the rest of the literature that suggests that identity-centric, continually verified access limits the blast radius and dwell time of an intrusion [4], [11].




Fig. 1: Breach cost evidence relevant to cloud and zero trust adoption. (a) Average breach cost by primary data environment, 2024. (b) Average breach cost by zero-trust deployment status. Data: IBM Cost of a Data Breach Report [5].

One additional challenge that is particular to financial institutions is that cloud migration itself is often regulator-visible: many jurisdictions mandate that a systemically important function be notified or non-objected prior to migration to a public cloud provider, and more and more supervisors are requiring financial institutions to prove that the same or greater security is provided after migration, rather than simply asserting it. This provides a real benefit, beyond the merits of the security case, to be able to develop an architecture that has controls that can be easily proven to a third party examiner. The need for evidence can be met through the emphasis on explicit, logged, per-session access decisions, which is well suited to the need for evidence of the security offered by zero trust, as opposed to perimeter security, where the assurance is largely the absence of an observed intrusion [4], [12].

C. PRINCIPAL THREAT VECTORS

The following four threat categories are common and relevant to cloud-hosted financial workloads in the sources examined: And about 2/3 of breaches were caused by a non-malicious human element, including error or social engineering, highlighting the importance of identity in the 2024 Verizon report [6]. In the same study, the use of Internet-facing and cloud-based vulnerabilities as an initial attack point has increased significantly year-over-year, almost tripling the number of attacks [6]. Beyond any attack on the underlying cloud platform, storage access policies that are too open and unrestricted management-plane access are cited in the cloud security literature as among the most frequent causes of exposure resulting from cloud misconfiguration [9, 10]. For institutions that offer open banking application programming interfaces (APIs) and/or delegated authorizations, vulnerabilities in these interfaces are also important; a formal analysis of the popular OAuth 2.0 authorization protocol identified previously unknown attacks in actual implementations [13].

CLOUD COMPUTING IN FINANCIAL INSTITUTIONS

A. SERVICE AND DEPLOYMENT MODELS

In the NIST definition of cloud computing, there are four cloud deployment models (private, community, public and hybrid cloud) and three services (software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS)) [1]. Financial institutions typically take a hybrid approach with key core ledgers and regulated customer data kept on private infrastructure or a private cloud, and the less sensitive or highly elastic workloads, like digital channels, analytics, and disaster recovery, running in the public cloud IaaS or PaaS. SaaS is frequently applied to non-regulated core business areas, such as human-resources and customer-relationship-management systems.

B. THE SHARED RESPONSIBILITY MODEL

One of the hallmarks of cloud computing is the sharing of security duties between the cloud service provider and the customer institution with the separation line changing depending on the service model. At the IaaS level, the provider has the responsibility to provide physical equipment and virtualization layer while the customer is responsible for operating systems, network controls, identity and data. With PaaS, the provider also takes care of the run-time and middleware, leaving the customer with only the application logic, data, and access configuration. SaaS requires the provider to take care of almost the entire stack, and the customer to be responsible for identity governance, data classification, and configuration of the controls that the provider exposes. ISO/IEC 27017 builds upon ISO/IEC 27002 and adds guidance specific to cloud that clarifies this division of responsibility for a specific subset of controls [14] and ISO/IEC 27018 does the same for the protection of personally identifiable information (PII) processed by a public cloud provider [15]. The Cloud Controls Matrix by the Cloud Security Alliance takes the shared responsibility model one step further by correlating controls with the domain area that they apply to and the party or parties that are responsible for each control [16].

C. MAPPING ZERO TRUST CONTROLS TO THE SHARED RESPONSIBILITY MODEL

In a cloud deployment, the customer institution has no control of the entire stack, so understanding which control families reside with the institution, which are shared with the provider, and which the provider helps to substantially implement – or provides as a native capability – is key to effective security. Fig. 2 shows a qualitative synthesis of this mapping for the 3 main service models.




Fig. 2: Qualitative synthesis of the primary locus of control responsibility for representative zero-trust control families across cloud service models. Ratings are the authors’ interpretive synthesis of the shared-responsibility literature [1], [14], [16], not a certification or vendor-specific mapping. One practical implication of the pattern in Fig. 2 is that institutions can’t expect that shifting a workload to a higher abstraction service model (IaaS to SaaS) automatically decreases their security load. In SaaS, IAM is still very much a customer responsibility, as the provider has no way of knowing which users of the institution should be granted which access. In contrast, IaaS brings network segmentation and encryption key management closer to the provider when the institution uses the provider’s own software-defined networking and key-management systems, but does not configure them to their default settings.

D. THIRD-PARTY AND SUPPLY-CHAIN EXPOSURE IN CLOUD-HOSTED FINANCE

Cloud adoption focuses an increasing amount of financial infrastructure into the hands of a few hyperscale firms, transforming the nature of third-party risk from bilateral vendor relationship to a potential single point of correlated failure for many institutions. This is clearly reflected in the development of an oversight framework for critical ICT third-party providers by DORA [17] and in the Basel Committee’s guidance on the mapping of the dependencies of critical operations, including cloud operations [18]. But, from a zero trust point of view, the risk of concentration is only partially reduced by rigorous access controls at an institution’s boundary, because the risk under the shared responsibility model mentioned above in Section IV [1] is the availability and integrity of the provider’s own infrastructure, which is not subject to the institution’s control. The reviewed regulatory literature highlights two key mitigations for this type of risk: exit planning and multi-region or multi-provider resilience arrangements, as tested, and not simply documented.

E. ZERO TRUST ARCHITECTURE: PRINCIPLES AND LOGICAL COMPONENTS

I. Core Tenets

NIST SP 800-207 defines zero trust as a set of tenets: All data sources and computing services are considered resources; All communication is secured regardless of network location; Access to individual resources is granted on a per-session basis; Access is granted via the dynamic policy that considers who is requesting access, the device’s state, and other behavioral and environmental attributes; and The

enterprise monitors and measures the integrity and security posture of all owned and associated assets [4]. A policy-based architecture is created to make these principles generalizable for standardization and audit, extending Kindervag's original prescription never trust and always verify [3] to a policy-driven architecture.

II. Logical Components

The logical architecture defined in NIST SP 800-207 is shown in Fig. 3. A policy engine, based on ongoing diagnostics and mitigation information, identity provider information, threat intelligence, activity logs, and compliance policies, calculates a trust score and access determination for every request. A policy administrator then implements that decision by creating or removing communication between a subject and a resource by using a point at which the subject can access the resource—a policy enforcement point (PEP). All subjects on the subject side of the PEP are seen as an untrusted zone; all subjects on the other side of the PEP, including cloud-based workloads, are seen as an implicit trust zone that the architecture continues to monitor, but is not considered trusted [4].




Fig. 3: Logical components of a zero trust architecture as defined in NIST SP 800-207 [4], showing the control-plane policy decision point and the data-plane policy enforcement point that mediates every subject-to-resource request.

C. Deployment Approaches

The approaches outlined in NIST SP 800-207 for implementing zero trust are not mutually exclusive, but include several different strategies: Enhanced identity governance, in which policy is focused mainly on the identity of the requesting subject; Micro-segmentation, in which the network is broken up into small, individual, secured zones; and Software-defined perimeter, in which network infrastructure is dynamically provisioned in accordance with identity [4]. Elements of this approach were anticipated in the Cloud Security Alliance's software-defined perimeter specification, which set out a framework in which network resources are kept hidden until a device and its user have been authenticated and authorized [19]. For financial institutions that have a large legacy footprint, the most achievable place to begin is enhanced identity governance, as this can be added to existing networks without an immediate redesign of the network; a full micro-segmentation and software-defined perimeter approach would generally entail a broader network redesign.

Approach	Core mechanism	Suitability for legacy core banking	Typical first target
Enhanced identity governance [4]	Policy centered on subject identity and entitlement	High; layers onto existing network without re-architecture	Privileged and administrative accounts
Micro-segmentation [11]	Network divided into small, independently secured zones	Moderate; requires network redesign but can be phased	Cardholder-data and payment-processing segments

Approach	Core mechanism	Suitability for legacy core banking	Typical first target
Software-defined perimeter [19]	Resources hidden until device and user are authenticated	Lower; best suited to greenfield or cloud-native segments	Remote-access and partner API gateways
Cloud-native posture management [16]	Continuous configuration assessment against a control baseline	High for cloud-hosted workloads; not applicable on-premises	Public and hybrid cloud digital-channel platforms

Table II: Comparison of Zero Trust Deployment Approaches for Financial Institutions

D. THREATS CONSIDERED IN THE NIST MODEL

NIST SP 800-207 also lists threats related to zero trust deployments, such as the ability to subvert the policy engine or administrator, DoS or network disruption on the control plane, credential theft to impersonate a valid subject, and the lack of visibility when using encrypted traffic or third-party network segments [4]. These threats are directly applicable to financial institutions where the policy engine and administrator are high-value targets because compromising these components can give the attacker access to much of the financial institution's traffic, and where third-party payment traffic is prevalent and encrypted.

TECHNOLOGIES ENABLING ZERO TRUST IN THE CLOUD

A. IDENTITY AND ACCESS MANAGEMENT

In a zero trust architecture, identity is the controlling factor and in cloud environments, location on the network has little security significance. The foundation is strong authentication, preferably multi-factor authentication and risk-adaptive, along with least-privilege authorization. When using open banking APIs, it is important to the correctness of the underlying authorization protocol: Fett et al. pointed out that seemingly minor implementation decisions could introduce vulnerabilities that the protocol had been designed to prevent, and proposed corrected configurations [13]. The larger NIST cybersecurity control ecosystem that is organized at the governance level by CSF 2.0, and referenced by many zero trust implementations, is used in many regulated industries for identification, authentication and access-control controls [20].

B. MICRO-SEGMENTATION AND SOFTWARE-DEFINED NETWORKING

Micro-segmentation breaks up a network, whether on-prem, cloud or hybrid into small segments that have access policy that can be enforced separately, meaning that a breach of one segment does not automatically give a user lateral access to another segment. Sheikh et al. present practical strategies for implementing network micro-segmentation as a zero trust enabler and analysis of the operational trade-offs between segmentation granularity and management overhead [11]. In cloud environments, micro-segmentation is generally applied on top of provider-native (SDN) software-defined networking (SDN) primitives, security groups, and policies in a service-mesh, enabling institutions to define boundaries without managing the physical network, in keeping with the shared-responsibility paradigm outlined in Section IV.

C. CLOUD SECURITY POSTURE AND ACCESS BROKERING

Cloud access security brokers and cloud security posture management tools deliver visibility and policy enforcement between an institution's customers and its cloud services, and add a layer of zero trust policy enforcement to SaaS and PaaS environments where an institution can't deploy a traditional network-based PEP. These tools continuously evaluate the configuration against a baseline, such as the ISO/IEC 27017 cloud control set [14] or the CSA Cloud Controls Matrix [16] and mark deviations from a baseline, or drift, because misconfiguration is the more frequently cited reason for cloud exposure mentioned in the surveyed literature [9] or [10] rather than platform vulnerability.

D. CONTINUOUS MONITORING AND ADAPTIVE POLICY

The NIST model demands reevaluation of trust constantly instead of during the course of a session, which in practice relies on the telemetry received from endpoint detection, network traffic analysis, and cloud workload monitoring that are returned to the policy engine [4]. This continuous evaluation cycle, also referred to as adaptive or risk-based access control, is one of the more technically complex concepts to implement at scale, especially when policy decisions need to be made within the financial transaction's latency envelope [8].

E. BALANCING SECURITY INVESTMENT AND OPERATIONAL FRICTION

From an operational cost perspective, zero trust and cloud security controls are not free: risk-adaptive authentication requires additional steps for legitimate access and micro-segmentation introduces additional administrative burden with network changes. He et al. point out that the technical challenge of continuous, risk-based evaluation is one of the tougher parts of zero trust to operationalize to the desired

scale and latency for production systems [8], and Buck et al. review the literature on zero trust and find that organizational and cultural resistance—not only technical—often hampers adoption [21]. To financial institutions, this means that a more stringent policy should be applied to workloads that are more sensitive and/or latencies are greater, and less stringent for others, as explained in more detail in Section VIII.

F. PRIVILEGED ACCESS AND INSIDER RISK

Financial institutions have a higher concentration of privileged accounts – from database administrators to cloud-platform administrators and automated deployment pipelines – that can circumvent many of the controls mentioned above if compromised or misused. The principle of zero trust architecture holds that privilege is granted for each session as opposed to having a blanket right, making the case for just in time elevation, session recording, and automatic termination of admin access privileges without a persistent admin credential [4]. In cloud environments, this includes the provider's own management plane: Cloud security posture management tools generally treat access to the management plane as a separate control family and one that is monitored closely, often in conjunction with normal user access, rather than as an integral part of that user access [9] [16]. Insider misuse and human error are two of the most common drivers of breach cost across the wider evidence base [5] [6] and privileged-access governance is a zero trust control, not simply an extension of ordinary identity and access management.

REGULATORY AND STANDARDS LANDSCAPE

Cloud computing and zero trust are increasingly a given for financial institutions, and a comprehensive set of regulatory and standards guidance is emerging that takes both of those technologies for granted. The milestones from the zero trust, cloud standards and financial-sector regulatory initiatives are now aligned along a common timeline as shown in Fig. 4.




Fig. 4: Milestones in the zero trust, cloud computing standards, and financial-sector regulatory landscape relevant to this review [3], [1], [19], [14], [2], [15], [12], [4], [18], [22], [17].

A. CLOUD-SPECIFIC FINANCIAL-SECTOR GUIDANCE

The Federal Financial Institutions Examination Council (FFIEC) released guidance focused on security in cloud computing environments, noting that outsource providers do not relieve institutions of their responsibility for security, and that banks must conduct thorough due diligence on their vendors, configuration management, and identity governance of their cloud providers [12]. This guidance aligns with the shared responsibility model discussed in Section IV and reiterates that identity and configuration controls continue to be the institution's responsibility, whether it is using a service model or not, and are essential to the zero trust model.

B. OPERATIONAL RESILIENCE AND DIGITAL OPERATIONAL RESILIENCE ACT

The Basel Committee's operational resilience principles involve identifying all dependencies, such as those of cloud and other third parties, that support critical operations and to simulate how well banks can sustain within the tolerable limits of disruption [18]. The EU's Digital Operational Resilience Act (DORA) introduces tightly related requirements as enforceable obligations on financial entities within the EU, such as the management of their ICT risk, testing of their ICT resilience and ensuring ICT third-party providers are reliable and resilient [17]. Combined with NIST CSF 2.0, which explicitly includes supply-chain and third-party risk-management, these documents move towards continuous assurance, one for which both zero trust's continuous verification and cloud security posture management's continuous configuration assessment are appropriate.

C. PAYMENT AND DATA PROTECTION REQUIREMENTS

The 2022 PCI DSS update actively reinforced authentication and monitoring protocols in a manner that closely mirrors the zero trust framework, such as a bespoke option in which an institution can prove that its data is equally secure using other controls [22]. The GDPR calls for appropriate technical and organizational measures, as well as the requirement to notify of a breach within 72 hours, to be taken with regard to personal data, regardless of whether it is processed on premises or on the cloud, and ISO/IEC 27018 offers cloud-specific guidance on how to meet comparable obligations when a public cloud provider acts as a processor [15], [23].

D. CERTIFIABLE ASSURANCE

Based on the ISO/IEC 27017 cloud control extension and external assurance of the information security management system to counterparties, auditors and regulators, ISO/IEC 27001 is a qualification that provides third-party confidence that an institution's information security management system conforms to a recognized standard [14], [24]. Like any certification, the currency and the clarity of the scope of the certification are the key factors that make it useful; a certification that addresses only a portion of an institution's cloud estate should not be interpreted to address the entire estate.

E. COMPARATIVE VIEW OF REGULATORY AND STANDARDS INSTRUMENTS

The instruments reviewed in this section are summarized in Table III, which places each instrument in terms of binding and its relevance to the zero trust and cloud posture of a financial institution.

Instrument	Nature	Primary relevance to ZT/cloud
FFIEC cloud guidance [12]	Supervisory guidance (US)	Accountability persists under shared responsibility; vendor due diligence
BCBS operational resilience [18]	Supervisory principles (international)	Mapping of cloud and third-party dependencies for critical operations
DORA [17]	Binding EU regulation	ICT risk management, resilience testing, oversight of critical cloud providers
NIST CSF 2.0 [20]	Voluntary framework (US)	Govern function incorporates supply-chain and third-party risk
PCI DSS v4.0 [22]	Contractual industry standard	Authentication and monitoring requirements aligned with ZT principles
GDPR [23]	Binding EU regulation	Technical/organizational measures and breach notification irrespective of hosting location
ISO/IEC 27001 and 27017 [14], [24]	Certifiable standard and cloud extension	External assurance of ISMS scope including cloud-hosted systems

Table III: COMPARISON OF REGULATORY AND STANDARDS INSTRUMENTS RELEVANT TO ZERO TRUST AND CLOUD ADOPTION

INTEGRATED ZERO-TRUST CLOUD SECURITY FRAMEWORK

A. FRAMEWORK DESCRIPTION

The material discussed in Sections III-VII are brought together in a financial institution-specific framework in Fig. 5. The left-hand side of the table lists the exposure drivers that motivate the framework: dissolution of the “network perimeter” due to SaaS and remote access, multi-cloud and hybrid infrastructure, legacy core-banking systems that were not built for identity-centric access control, third-party and FinTech API integration, and the rapid growth of privileged and machine (non-human) identities. These drivers are working to support a control core around the NIST SP 800-207 logic [4] that defines who is setting policy and making risk-based access decisions, who and what is requesting access, in what state, and how lateral movement is being constrained once access is granted, and how security posture management is being extended into the IaaS, PaaS, and SaaS environments as part of the shared responsibility model [1], [16]. Monitoring and analytics connect the layers, feeding back to the policy engine, and the resiliency outcomes, i.e. continuous assurance, contained blast radius, elastic and resilient operations, regulatory confidence align with DORA concerns, Basel principles, and FFIEC guidance [26], [27], [28].




Fig. 5: Integrated Zero-Trust Cloud Security framework for financial institutions. Exposure drivers shape a control core organized on the NIST SP 800-207 logic [4]; monitoring and analytics close the loop back into policy.

The following caution applies to the framework as presented: It is not a tested reference architecture, but a summary of the literature and standards that have been reviewed; it is intended to structure governance discussions and control mapping, not to provide a guarantee of immunity to compromise once a governance system is followed. Its strength is in communicating the relationships between the various exposure drivers, controls, and assurance evidence that regulators and counterparties are now increasingly demanding, and that this discipline should always be used to expose the weaknesses in one area and prevent them from being covered by the strengths in another.

B. PROPOSITIONS FOR FUTURE RESEARCH

The framework yields propositions intended to guide empirical validation rather than findings already established by the reviewed literature.

- P1 (Layered defense). Zero trust identity controls and cloud-native network segmentation are complementary rather than substitutable; institutions that deploy only one layer retain materially higher exposure than those that deploy both [29], [11].
- P2 (Shared responsibility clarity). Breach likelihood is higher where an institution's security team cannot correctly state which controls the cloud provider owns for a given service, independent of the technical adequacy of the controls actually in place [1], [30].
- P3 (Identity as the binding constraint). As institutions adopt more SaaS, the marginal security return on identity and access management investment rises relative to network-layer investment, because network segmentation is decreasingly available to the customer in higher-abstraction service models (Fig. 2).
- P4 (Latency-security trade-off). Continuous, risk-based policy evaluation improves security outcomes but introduces latency that is more consequential for real-time payment processing than for batch or analytical workloads, implying that uniform zero trust policy stringency across all financial workloads is not optimal [8].
- P5 (Assurance transparency). Certifications and attestations that disclose their precise cloud scope contribute more to counterparty and regulator confidence than broader but less specific claims of compliance [12], [24].

C. ILLUSTRATIVE APPLICATION

Take, as an example, a retail bank that has built a digital-channel platform in the public cloud on a PaaS, has kept its core ledger on-premises, and provides payment-initiation APIs to third-party FinTech apps as part of an open banking initiative. The framework calls for identity and device trust to be based on strong, risk-adaptive authentication for both customers and partner applications, and the use of narrow, short-lived scope and duration of tokens, recognizing the weaknesses of the protocol introduced by the OAuth literature [13]. By limiting the access of the digital-channel platform to the on-premises core to a small

number of service accounts and API calls, a compromise of the cloud-hosted channel does not have direct access to the ledger [11]. Cloud security posture management constantly scans the configuration of the PaaS environment against the institution's baseline configuration, which is based on ISO/IEC 27017 [14]; it monitors and analyzes the correlation of all identity, network and cloud telemetry to identify any unusual pattern, such as a partner application requesting an unexpected amount of account data, which can be flagged for review. The confidence in the regulatory posture is embedded in the capacity to provide evidence of each of these controls as needed, a posture for which DORA and FFIEC guidance are increasingly expecting [12], [17]. This vignette is representative and not the presentation of any particular institution or product.

D. IMPLEMENTATION GUIDANCE

- Start from identity, not network topology. Inventory human and machine identities and their entitlements before redesigning network segmentation; identity gaps undermine segmentation regardless of its technical quality [4].
- Map the shared responsibility boundary explicitly, service by service. A single enterprise-wide assumption about what the cloud provider secures is rarely accurate across a mixed IaaS, PaaS, and SaaS estate (Fig. 2).
- Sequence migration by criticality and coupling. Begin zero trust and cloud posture hardening with systems that are both highly exposed (internet-facing, third-party-integrated) and highly consequential (payment-critical), rather than migrating by convenience.
- Instrument before enforcing. Deploy monitoring and logging ahead of strict policy enforcement so that legitimate but unusual access patterns, common in institutions with complex legacy integration, can be identified before they are blocked.
- Treat certification scope as a living artifact. Keep ISO/IEC 27001 and related certifications current with the institution's actual cloud footprint, and disclose scope precisely to counterparties and examiners [12], [24].

E. INDICATIVE METRICS FOR TRACKING MATURITY

However, there is no agreed upon maturity metric for zero trust in the literature that was reviewed [21] and institutions tend to create their own set of indicators based on surrounding, well-established measurement practices. Table IV shows a sample of the items in the control families referred to above, and is only meant as an example of what might be present and is not a validated instrument.

Control family	Illustrative metric	Data source
Identity and access	Share of privileged accounts under MFA and just-in-time access	Identity provider and PAM logs
Micro-segmentation	Share of critical workloads in a dedicated, independently enforced segment	Network policy inventory
Cloud posture	Mean time to remediate a configuration drift from baseline	CSPM / CASB alerts
Monitoring	Share of access decisions informed by real-time device or behavioral signal	Policy engine decision logs
Third-party oversight	Share of critical ICT providers with a tested exit or substitution plan	Vendor risk register

Table IV: Illustrative Metrics for Tracking Zero-Trust and Cloud Security Posture

OPEN CHALLENGES AND RESEARCH AGENDA

Several challenges recur across the literature reviewed and define a research and practice agenda.

- Legacy core-banking integration. Many institutions have core systems that were built before the advent of IAM solutions focused on identity-based access control and lack the ability to show the granular access control mechanisms which a zero trust policy engine needs, resulting in compensating controls being added at the integration layer [2].
- Latency-sensitive transaction processing. There is limited research quantifying the latency that can be introduced to real-time financial workloads due to continuous risk-based policy evaluation, which may be incompatible with the throughput and response-time requirements of real-time payments [8].
- Identity sprawl and machine identities. In most institutions, the number of non-human identities (service accounts, APIs, automated agents) is increasing rapidly, and governance of non-human identities is less developed than for human identities [7], [21].
- Measuring zero trust maturity. Unlike breach cost where there is a consensus metric, there is no consensus metric for zero trust maturity, most models are self-assessed, and not independently verifiable [21], [25].

- Multi-cloud consistency. Where institutions rely on multiple cloud providers, they face challenges of dissonant native identities, segmentation, and monitoring primitives; and the literature is sparse of empirical advice on how to deal with such diversity on an institution-wide scale [10].
- Governance of AI-enabled security tooling. Cloud security posture management and zero trust policy engines are frequently adding in machine learning for anomaly detection and risk scoring, and so are beginning to cause questions about explainability and model-risk in the answers generated, which current standards are only partially addressing [8, 20].
- Cost and skills constraints for smaller institutions. Compared to large banks, community and regional financial institutions (CRFIs) are more likely to have disproportionately high implementation cost and more likely to be severely affected by a shortage of security skills, and the literature has a sparse collection of validated, less-costly approaches to achieving zero trust and cloud security maturity that can be applied to CRFIs [12], [20].

These obstacles are not mutually exclusive. Legacy integration constraints often lead towards coarser grained, higher latency policy checking, which in turn accentuates the latency-security trade off; and the resulting compensating controls lead to an increase in machine identities and service accounts, deepening the identity-sprawl problem. A research and practice program that considers them as a coordinated group of constraints, rather than as individual pieces of work to be performed by different groups/teams, is more likely to lead to workable guidance.

CONCLUSION

In this review, the authors have explored the potential application of ZTA and cloud computing technologies to boost cybersecurity in financial institutions. The evidence shows that these two technologies complement one another, with one technology replacing a fixed network boundary that served as the foundation for traditional security, and the other providing the identity-centric, always verified approach that a boundary-less environment demands. Although it is observational, the evidence shows that both misconfiguration of cloud resources and lack of zero trust is correlated with increased losses, while mature deployment of zero trust is correlated with significantly reduced costs. Under the shared responsibility model, the cost of adoption of cloud infrastructure does not diminish an institution's responsibilities for security, but rather shifts responsibilities, and identity and access governance still largely remains the responsibility of the institution regardless of the service model.

The Zero-Trust Cloud Security framework proposed by the Cloud Security Alliance provides financial institutions with a structure for applying identity, device, network and data controls to the shared responsibility model in light of the expectations of the regulatory frameworks outlined in DORA, the Basel operational resilience principles and FFIEC cloud guidance. Actual deployment is limited by legacy integration, latency sensitivity, identity sprawl and lack of maturity metrics. Aimed at tackling these limitations empirically, and managing the artificial intelligence elements that are becoming part of the fabric of both zero trust and cloud security tools, these are the next directions of research and practice.

REFERENCES

- [1] U. A. Kithran, S. Ilyas, K. Amber, H. Abbas, and S. Zahid, "AI-based collision prediction and prevention systems in urban traffic," *Spectrum of Engineering Sciences*, vol. 3, no. 12, pp. 1907–1920, 2025.
- [2] M. R. Islam, I. A. Badhan, M. H. Rahman, and M. N. Hasnain, "The economics of global renewable energy supply chains: Opportunities and vulnerabilities for the U.S.," *International Journal of AI, Engineering and Management Studies (IJAIEMS)*, vol. 1, no. 1, pp. 108–136, 2026.
- [3] U. Imtiaz, "Ghost signals: Ethical RF adversarial testing of consumer alarm ecosystems," in *Data Science and Security*, S. Shukla, H. Sayama, K. Tiwari, J. P. George, and J. V. Kureethara, Eds. Cham, Switzerland: Springer, 2026, Lecture Notes in Networks and Systems, vol. 1945, doi: 10.1007/978-3-032-24075-0_21.
- [4] S. M. Ali, M. Hadi, F. Usama, M. S. Khan, M. Iqbal, and S. Ahmad, "From artificial to quantum intelligence: The next frontier of low-carbon energy systems," in *Proc. 2026 Int. Conf. Integrated Intelligence and Cognitive Engineering (ICIICE)*, Dubai, United Arab Emirates, 2026, pp. 1–6, doi: 10.1109/ICIICE69672.2026.11565487.
- [5] A. Badhan, M. N. Hasnain, M. H. Rahman, I. Chowdhury, and M. A. Sayem, "Strategic deployment of advance surveillance ecosystems: An analytical study on mitigating unauthorized U.S. border entry," *Inverge Journal of Social Sciences*, vol. 3, no. 4, pp. 82–94, 2024.
- [6] M. Delshadi, N. Ahmad, M. D. Rasheed, M. S. Sardar, and U. Aslam, "Deep learning-based network traffic analysis for cyber threat detection," *Spectrum of Engineering Sciences*, vol. 3, no. 12, pp. 607–614, 2025.
- [7] K. Khaliq, M. S. Khan, S. Sajid, I. Shahbaz, M. W. Iqbal, K. S. Siddiqui, *et al.*, "A study on the application of deep learning and machine learning in adaptive intelligent tutoring," *Spectrum of Engineering Sciences*, vol. 4, no. 3, pp. 818–835, 2026.

- [8] U. A. Kithran, A. Khanam, S. Ilyas, R. Yaqoob, and M. Hannan, "Demand forecasting in transportation networks using deep learning," *Spectrum of Engineering Sciences*, vol. 3, no. 11, pp. 687–700, 2025.
- [9] U. Iqbal and Y. Bhutto, "Digital transformation through artificial intelligence and advance business analytic in American operational management," *Journal of Theoretical and Applied Econometrics*, vol. 3, no. 1, pp. 37–50, 2026.
- [10] M. M. Islam, "AI-driven fraud detection and prevention using human behavior analysis to enhance U.S. social and financial security," *International Journal of Applied Mathematics*, vol. 38, no. 8S, pp. 861–871, 2025.
- [11] D. Safiullah, A. A. Khan, N. Ahmad, M. Danish, and S. Rana, "Integrating artificial intelligence and soil ecology to enhance plant resilience under environmental stress," *ASSAJ*, vol. 4, no. 2, pp. 3508–3516, 2025.
- [12] R. Singh, "Brønsted-acid-guided esterification: Mechanistic and kinetic insights for long-chain fatty acid to alkyl ester conversion on a heterogenized porous acid platform," *International Journal of Clinical Research and Medical Sciences*, vol. 1, no. 1, 2026, doi: 10.67231/rzakra50.
- [13] U. Iqbal, "AI-enhanced network optimization for electric vehicle charging infrastructure expansion in the United States using graph theory and demand analytics," *Journal of Engineering and Computational Intelligence Review*, vol. 2, no. 2, pp. 112–129, 2024.
- [14] M. Chaudhari, "Hydrophilic nanofiber barrier coatings for aluminum: Waterborne crosslinking, microstructure–property links, and EIS-based protection," *Genetics and Molecular Research*, vol. 25, no. 3, Art. no. 3017, Jul. 2026, doi: 10.4238/rhd2k915.
- [15] R. Singh, "Smart prediction of molecular behavior in liquid chromatography for better compound detection," *International Journal of Clinical Research and Medical Sciences*, vol. 1, no. 1, 2026, doi: 10.67231/v3jyx227.
- [16] R. Parveen, "An explainable AI framework for ethical fraud prevention in U.S. federal welfare programs," *International Journal of Applied Mathematics*, vol. 38, no. 11S, pp. 1425–1435, 2025.
- [17] M. N. Amjad, "Machine learning for renewable energy forecasting and smart grid optimization," *Journal of Engineering and Computational Intelligence Review*, vol. 2, no. 2, pp. 130–145, 2024.
- [18] U. Imtiaz, "Dynamic security certification framework for evolving distributed architectures," in *Data Science and Security*, S. Shukla, H. Sayama, K. Tiwari, J. P. George, and J. V. Kureethara, Eds. Cham, Switzerland: Springer, 2026, Lecture Notes in Networks and Systems, vol. 1947, doi: 10.1007/978-3-032-24360-7_32.
- [19] M. Chaudhari, "Epoxy barrier coatings reinforced with solidified-waste microfillers for aggressive acids," *Genetics and Molecular Research*, vol. 25, no. 3, Art. no. 3020, Jul. 2026, doi: 10.4238/96y1s945.
- [20] Kumar and K. K. S. Nair, "The impact of cloud-based business analytics on operational efficiency in U.S. small and medium-sized firms," *Spanish Journal of Innovation and Integrity*, vol. 49, pp. 303–318, 2025.
- [21] S. T. Hasan, "Labor market intelligence using natural language processing for the U.S. direct care workforce: Predicting skill gaps, recruitment needs, and emerging healthcare workforce demand," *American Journal of Multidisciplinary Knowledge Insights*, vol. 3, no. 1, pp. 1–16, 2024.
- [22] S. Mirza, S. Butt, A. Arif, and N. Jahangir, "COVID and post-COVID e-learning reforms in higher education institutions: A systematic review of literature," *Journal of Business Insight and Innovation*, vol. 5, no. 8, pp. 364–383, 2026.
- [23] M. N. Amjad, "Artificial intelligence-based voltage stability assessment in modern power systems: Edge intelligence for real-time monitoring of smart electrical networks," *Journal of Business Insight and Innovation*, vol. 3, no. 1, pp. 102–115, 2024.
- [24] M. A. Jawed, "Uncertainty-aware geostatistical reconstruction of regional groundwater hydraulics under sparse monitoring conditions: A case study of the Evergreen Underground Water Conservation District (EUWCD), Texas, USA," *Bishop International Journal of Mathematics and Computer Science*, vol. 1, no. 1, pp. 57–71, 2025.
- [25] W. Reese and D. Sinclair, "Bridging the digital divide: Digital inequality and socioeconomic outcomes in the United States," *American Journal of Multidisciplinary Knowledge Insights*, vol. 5, no. 1, pp. 1–8, 2026.
- [26] A. Kumar, "The role of banking credit access in healthcare analytics-driven practice growth: Implications for regional economic development in the U.S.," *Journal of Computational Systems & Engineering Insights*, vol. 2, no. 2, pp. 15–28, 2024.
- [27] N. J. Mercer, "Enhancing organizational performance through strategic cloud computing in management," *International Journal of Business Integrity and Technology Advancements*, vol. 1, no. 1, pp. 11–17, 2025.
- [28] K. Hamdan, "AI-powered cyber resilience strategies for modern organizations," *Journal of Computational Systems & Engineering Insights*, vol. 4, no. 1, pp. 1–10, 2026.

- [29] R. M. Cachart, F. Ahmed, S. Pochincharla, V. Gummadavelli, R. S. Ganapathy, and A. Devineni, "Adaptive program governance platform using reinforcement learning and policy optimization," in *Proc. 2026 Int. Conf. Intelligent and Sustainable AI Systems (ICOSAAS)*, Bali, Indonesia, 2026, pp. 1563–1570, doi: 10.1109/ICOSAAS68663.2026.11648756.
- [30] M. S. Uddin, A. Khan, I. Alim, and D. Khan, "The evolution of cybersecurity in the digital age: Contemporary threats, technologies, and future directions," *Journal of Computational Systems & Engineering Insights*, vol. 4, no. 2, pp. 12–29, 2026.