Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

AZ – TAC Adaptive Zero Trust Access Control with Continuous Authentication for Secured Cloud Environments

A Aswini
Vellore Institute of Technology
S Prasanna
sprasanna@vit.ac.in

Vellore Institute of Technology

Article

Keywords: Zero Trust Security, Cloud Access Control, Identity Centric Security, Continuous Authentication, Risk Aware Authorization, Cryptographic Protocol Random Oracle Model

Posted Date: April 1st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8922010/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License. Read Full License

Additional Declarations: No competing interests reported.

AZ – TAC Adaptive Zero Trust Access Control with Continuous Authentication for Secured Cloud Environments

Aswini A [1], Prasanna S [2*]

[1,2]School of Computer Science Engineering and Information Systems, Vellore Institute of Technology, Vellore, India

aswini.2022@vitstudent.ac.in, sprasanna@vit.ac.in*

Abstract

Identity-centric access approaches are becoming fundamental to enable the management of highly dynamic users, services, and workloads in cloud computing environments. However, traditional access control systems are predominantly focused on static tokens, one-time authentication and coarse-grained authorization, which is not ideal to combat modern-day cloud attacks including leaking of credentials, abusing non-human identity, replay attack and lateral movement. In this paper we present an Adaptive Zero -Trust Cloud Access Security Protocol that is based on the concepts of continuous authentication and risk aware authorization applied to human and non-human identities in a cloud environment. The protocol leverages the Zero Trust model to ensure that each access request is cryptographically authenticated, corroborated with contextual security signals, and mathematically verified to present an acceptable risk. In this proposed work, a secure set of techniques for limiting disclosure of personal data so that access and identity are separate, while use of Proof-of-Possession tokens, Mutual Authentication and short-lived credentials mitigate impersonation and token misuse. This framework also ties together dynamic authorization with least-privilege enforcement, automatic credential revocation and optimized key rotation to minimize the effects of identity compromise. For accountability purposes all security-critical events are being recorded into an append-only audit log which is tamper-evident. Formal security of the protocol in the Random Oracle Model is presented, which shows that it can withstand impersonation, replay, man-in-the-middle and credential theft attacks. An informal security analysis also supports its resilience to insider threats and non-human identity violation. The theoretical foundations of strong cryptography combined with flexible risk evaluation and optimization driven approach allows the proposed approach

to deliver improved cloud access security without sacrificing scalability and performance, thus being suitable even in large-scale and dynamic cloud settings.

Keywords: Zero Trust Security, Cloud Access Control, Identity Centric Security, Continuous Authentication, Risk Aware Authorization, Cryptographic Protocol Random Oracle Model

1. Introduction

Cloud Computing has evolved as the prevailing computing model for making a network of shared resources available via Internet, allowing companies to quickly put applications and store huge amounts of data into operation with users distributed worldwide [1]. Today's cloud platforms accommodate highly dynamic workloads, elastic scaling, multi-tenant system architectures and powerful automation systems that greatly enhance operational efficiency [2]. But this has also come with some far more complicated security challenges that traditional network perimeter models are not equipped to handle. With cloud infrastructures transforming into decentralized, service-based and identity-centric designs the concepts of access control and identity security have become focal points for enabling secure cloud operations [3] .

In traditional enterprise settings, security was ensured by network perimeters, firewalls and static trust relations. An authenticator within a trusted network boundary often allowed wide access across internal resources. Such assumptions are fundamentally at odds with cloud, where the users, services, and workloads span multi-networks, shared infrastructures and data centres located in different geographical locations [4]. Furthermore, cloud services are typically interfaced with programmatically rather than via direct human interaction which introduces a huge boon of non-human identities like service accounts, application tokens and workload credentials. They outnumber human users in many cases and have (or at times, even higher) privileges that attract attackers.

Recent security events and empirical studies have demonstrated that identity theft is one of the most common factors contributing to cloud attack. Attackers are now taking advantage of exposed credentials, non-expiring API keys, over-entitled service accounts and inadequate monitoring around access tokens to gain unauthorized access to cloud resources [5]. Once an identity has been compromised, attackers may traverse the network between services, elevate

privileges, exfiltration sensitive data or perform denial of services against important workloads. Previous access control systems, based on static authorizations and single round authentication are not suitable to cope with these kinds of attacks at runtime [6].

The Zero Trust security model is been advocated to overcome these limitations and has gained much interest in both academia and industry [7]. Zero Trust models are about, in a fundamental sense, of ending the respect for implicit trust and affirming that every single access request must be validated at the time it is made. In this model, access decisions are made with explicit checks for identity, context-based attributes, and resource sensitivity [8]. Network location and previous authentication are not grounds for trust, so instead of automatically trusting devices within the network, Zero Trust utilizes a "never trust always verify" mentality. Although the conceptual pillars of Zero Trust are clear, devising the concrete mathematically-based and cloud-native access control protocol that materializes this principle is still an open research challenge [9].

One of the key obstacles to implementing a Zero Trust model in the cloud zone is its context of access that is continually changing. The behaviour of cloud users and the workloads serving them varies significantly, with frequent movement to new execution context (eg VM) and network locations. What does happen is that static policies are quickly obsolete or overly permissive [10]. Moreover, non-human identity exacerbates the issue since they are often idle in authentications or don't participate interactively and continues to rely upon embedded credentials for access. This kind of identity requires services to not just verify requests, but also dynamically check the validity of the running context and access pattern [11].

The dynamic property of cloud is high performance and scalability trade-off to be achieved when strong security guaranteed [12]. Another major obstacle is to achieve Zero Trust in the cloud. Cloud systems have to handle millions of requests per second in low-latency wide-area networks. Cost of crypto – re-authentication/renewal overhead: Re-Authentication or aggressive rotation of credentials etc. can lead to poor system performance and bad user experience [13]. Therefore, an effective CAS-i perspective should contain optimal methodologies that adaptively tune the strength of security controls in reaction to risk while minimizing computational and communication overhead [14].

Cryptography study point of view, there are many cloud access control schemes that relies on bearer tokens in the form of long-term argument that could be re-used if leaked. These

methods lack of guaranteed proof-of-possession and expose vulnerabilities to various types of attacks including replay, token thefts, and credential compromises [15]. Additionally, certain cloud provider audit procedures are not tamper evident in respect to this sort of forensic post-incident analysis [16]. All these shortcomings necessitate for a mathematical framework to ensure strong guarantees on the secrecy, integrity of access requests along with their authenticity, freshness, and accountability [17].

Inspired with these observations, in this paper we propose an Adaptive Zero Trust Cloud Access Security Protocol, which provides the cloud with constant and fine-grained user personal defence. The technique goes beyond traditional role-based or static attribute based access control using cryptographic evidence, context-aware risk assessment and dynamic authorization in a single protocol. Every request for access is a separate security event assessed using mathematically straightforward checks and risks maneuverers. This ensures the visibility updater forces changes dynamically rather than via some static assumptions.

The resulting protocol proposes a privacy-preserving identity representation that separates operational access from real identities and thus limits the potential of revealing an identity. Robust cryptographic primitives are used to derive identity-specific secrets, proof-of-possession tokens and secured communication channels. All user access requests are secure over encryption, integrity assurance and freshness constraint to prevent replay attack and impersonation. Moreover, there are continuous authentication measures in the protocol which re-verify identity trustworthiness during a session so as to detect forging patterns at an early stage.

We present N-HIP, a holistic approach against the ever increasing threat of compromised non-human identities which uses special defence mechanisms like workload binding, honey-credentials and automated secret rotation as part of our proposed architectural structure. These mechanisms protect that credentials can't be used also outside the context since mandated and leaked secrets are identifiable and can be easily revoked. Additional audit-ledger tamper-resistant function is provided to generate auditable logs for all authentication and authorization events, ensuring accountability and compliance.

One of the special characteristics of our work is that, risk-aware optimization is incorporated in the access control decision model. Contrast Information Dance Instead of applying the same security policies to all requests, the protocol actually calculates a continuous risk value

based on various contextual factors such as device health, trustworthiness of the network connection, user behaviour anomalies, signals driven by threat intelligence and sensitivity of resources. Workload Score an access control policy is determined according to this score and the system adapts access scope, token lifetime and enforcement actions accordingly. An optimisation model is used to trade security risk off against the overhead of operations, only applying stronger controls where they are necessary.

The entire protocol is organized in a series of precisely described phases – each related to a particular feature of cloud access security. These stages are system initialization, identity registration, secure channel setup, access request generation, continuous authentication with risk assessment, adaptive authorization and enforcement and revocation handling with optimized key renewal. The properties of each stage are formally specified and accompanied with equations which characterize the security provided by the protocol.

The major contributions of this paper are:

- **Adaptive Identity-Centric Zero Trust Framework**: Presents a new AZ-TAC protocol to provide continuous A&A for both human and non-human identities in the cloud.
- **Risk-Aware Continuous Trust Evaluation:** Presents a risk-based model enabling dynamic access control by combining multi-dimensional risk scores calculated from behaviour analytics, device posture, network trust and threat intelligence.
- **Cryptographically Secure Access Mechanism:** Uses mTLS with certificate pinning and proof-of-possession token binding to prevent credentials leakage, token replay and impersonation.
- **Least-Privilege Dynamic Authorization and Revocation:** This provides us with ABAC-based adaptive authorization and honey token-based compromise detection, short-lived scoped tokens, and efficient secret rotation.
- **Comprehensive Performance Validation:** Shows better latency, detection rate, scalability and revocation response in hybrid dataset-driven and simulation based evaluation than existing IAM and OAuth2 model.

The rest of the paper is organized as follows. Section 2 Related Work Cloud access control, identity security and Zero Trust architectures work are described in Section 2. Section 3 presents the system model and threat models. In section 4, we provide a comprehensive description of the proposed Adaptive Zero Trust Cloud Access Security Protocol along with phase-wise algorithms and mathematical representations. In Section 5 we hand out both formal and informal security analysis. Performance and optimization impact are discussed in Section 6, and the paper is concluded with future research in Section 7.

2. Literature Survey

Various authors have proposed various protocols among them some of the authors survey is discussed in detail. Alazab et al. [18] presented an adaptive cloud security framework that incorporates federated learning–guided anomaly detection, and dynamic communication-protocols customization to preserve virtualized clouds. Local detection models can be trained in a decentralized manner on each hypervisor using their framework, to collaboratively perform threat intelligence without sharing raw traffic data and preserving privacy. Dynamic protocol parameters tuning mechanisms are employed by their system to decrease communication burden and enhance detection accuracy against hyper jacking, side-channel abuse, and VM escape attacks. Experiments showed better anomaly detection performance than the centralized and static security modes, especially in multi-tenant clouds. Their technique is successful in preserving both scalability and privacy. However, their framework focuses on infrastructure level anomaly detection and does not consider identity-centric access control, cryptographic authentication or continuous authorization which reduces its relevance in Zero Trust cloud settings dominated by misused credentials and abuse against non-human identities.

Ansari et. al.[19]proposed a Bi-SC-CBALSTM model that integrates both methods into cloud intrusion detection. A pre-processing process and ADASYN-based class balancing pipeline are included in the model to handle unbalanced cloud attack datasets. Through efficient feature subset selection and learning spatial–temporal attacking patterns, their model exhibits strong performance in maintaining high true detection rates, precision and F1-score for various benchmark datasets. Their method performs particularly well at processing high-dimensional data and dynamic attack behaviours. To be sure, the adoption of met heuristic optimization and deep neural networks inevitably introduces more computational cost, training time, and energy consumption than those in previous studies which may prevent real-time measurement over large cloud environments.

Huo et al.[20] introduced DRPE-Cloud a new optical image encryption structure based on the combination of double random phase encoding and thumbnail-preserving encryption designed for secure cloud storage of multimedia contents. Their scheme is based on hyper chaotic key generation and hash function to expand the key space and also to withstand brute force, statistical, and birthday attacks. Controlled visual degradation can be achieved without loss of data using the tunable spectral-phase modulation approach. Experimental tests

showed strong encryption security and high-quality decryption for images with different resolutions. Their method achieves an optimal trade-off between privacy preserving and usability for cloud-based multimedia applications. Nevertheless, the power of practical in-cloud deployment of quantum secure cloud key based synchronization and chaos key management, together with customized encryption thrusts complexity and scalability issues.

Zhang et al. [21] proposed a learning-based cloud security model which has explicit advanced feature representation and intelligent decision models for cyber threat detection. Their high dimensional patterns of cloud users and workloads are captured by their proposed approach, thus adaptive detection on complex and variant attacks can be achieved. Comprehensive experiments on standard datasets demonstrated the higher effectiveness and robustness of the proposed approach in comparison to traditional machine-learning-based intrusion detection systems. Their model is especially suitable for mixed cloud with heterogeneous workloads. However the centralized learning architecture and high computational complexity feature extraction process might raise scalability, latency and resource utilization concerns that make it impractical for real time cloud security enforcement without distributed or light weight enhancements.

Raja et al. Proposed a remote signcryption-based cloud security model, where confidentiality authentication and non-repudiation are provided by the integrated cryptographic process. Their scheme combines encryption with digital signature in a sign-then-encrypt way to spare the computing expenses for information security and lower out data volume. Security analysis as well as the performance tests showed that it was resistant against forgery, replay and chosen message attacks. The lightweight nature of their scheme is conducive for cloud data transmission and users with limited resources. However their architecture has a strong focus on data-level protection and does not include identity-based access control, continuous authentication and adaptive authorization mechanisms necessary to capture modern Zero Trust cloud style architectures.

Singamaneni et al. [22] have suggested an approach of hybrid-cloud secure computing by combining multi-qubit QKD and CP-ABE in order to provide the data confidentiality as well as fine-grained access control. Quantum channel can achieve secure key exchange with the functionality of detecting eavesdropping and CP-ABE provides policy based data access control by using cloud. Their experimental results showed that the rate of generating keys were higher, their distance of transmission was farther and less composite operation than

those of traditional cryptographic schemes. Their framework provides strong post-quantum security assurances. Yet, its dependency on simulated quantum infrastructure and added system complexity challenge its immediate real-world utilization in cost-sensitive and cloud scale-out cases.

Shakor et al. [23] demonstrated a cloud-based data security framework through the combination of AES File encryption and block chain for decentralized key management. They build file-specific encryption keys using hash- based process and save these keys in a private blockchain in an immutably manner, while ECC is utilized for confidential key exchange. Experimental results showed the high entropy, and heavy sensitivity by both key-guessing attacks and statistical attacks. Their method efficiently enhances data-at-rest security and key correctness. But it is primarily focused on encryption and key management, without identity-aware access enforcement, an adaptive authorization or continuous authentication capabilities essential to implementing Zero Trust-based cloud security.

Sasikumar and Nagarajan [24] presented an adaptive cloud security model combining multi-factor with machine learning based dynamic cryptography. Their system utilizes passwords, conditional features, and biometrics for authentication; meanwhile it is supported by a hybrid CNN–Transformer model that helps to predict attack types and choose suitable encryption algorithms at the right moment. Their performance of the proposed approach was evaluated on benchmark datasets and reached high efficiency in attack classification rate along with enhanced security against brute-force, phishing and impersonation attacks. Their adaptive cryptography mechanism provides a hardened environment for new things to consider. Nevertheless the deep learning model components add computational overhead such as time-consuming authorization on video frames, and it does not have on-the-fly identity attestation and risk-based access control in support of Zero Trust properties.

Santos et al. [25] introduced a secure cloud architecture for healthcare systems based on random pattern data division along with the multimodal biometric authentication. It is a frame for the security of privacy-sensitive patient medical data by storing shards of it in NoSQL cloud storage using face and fingerprint fusion for user authentication. Experimental results showed that their proposed scheme would reduce the latency comparing on AES-encryption and to increase authentication reliability. The application is useful for high-latency healthcare applications. But ultimately, it is focused on data fragmentation and initial of what

access, not continuous authentication or identity-centered enforcement or adaptive Zero Trust.

Hussain et al. [26] introduced a TCP. Delta-based cloud security modelling framework, to analyse TCP three-way handshake timing difference for ransom ware in real time. Their method also can circumvent deep packet inspection and neural network learning to perform efficient anomaly detection for high speed cloud traffic by means of statistics analysis and clustering. Their model supports CI/CD pipelines effectively. But it concentrates on network-level signs of behaviour and does not factor in cryptographic authentication, identity validation, and adaptive access control.

Mamidala et al. [27] preferred the hybrid cloud threat detection solution, which integrates Isolation Forest anomaly detection and machine learning driven behaviour analysis. Their method uses unsupervised learning to identify unauthorized access and leakage of the data with scalable and low detection latency. Their results demonstrated better accuracy with fewer false positives than conventional classifiers. Their framework is applicable for big datasets and elastic cloud environments. But it doesn't have crypto identity verification, continuous user authentication and Zero Trust-based access enforcement.

Nauman et al. [28] presented Q-ECS that combines Quantum Key Distribution with attribute-based cryptography, for better cloud security protection. Their model is able to achieve secure key exchange which is resistant to post-quantum attacks and can also provide for fine grained access control. Through simulation-based assessments we found increased efficiency in key generation, decreased latency during encryption as well as efficient handling of noise. The hybrid quantum–classical approach improves security in the long term. However, the practical application for near-term clouds is restricted by its reliance on quantum infrastructure and higher deployment cost.

An approach based on cloud-assisted malicious code defense system for WMN systems is provided in Bie and Ren [29], which used SIR-W epidemic propagation model that contained self-limited variable and Bluetooth characteristics. Infection, warning spreading, and recovery model are established in mathematical manner with equilibrium stability analysis. A better containment of the malware propagation was achieved with respect to classical SIR models on the basis of simulation results. The construction is supported by a rigorous mathematical basis. Nonetheless, the assumption of identical node behaviour or

ideal communication comes with some limitations when applying it to heterogeneous and latency-sensitive real world networks.

Thamil Selvi et al. [30] suggested a post-quantum cloud data security model using Modified EC, SIDH and certificate less searchable encryption. It facilitates secure and efficient keyword search over encrypted cloud data in the presence of Authorized Third Party (ATP) with provable resistance against insiders as well as quantum adversaries. Experimental results indicated less encryption time, memory usage and performance than the schemes based on RSA and Elliptic Curve Cryptography. Their model enjoys strong cryptographic protections. Nevertheless, it does not present identity-based access control, ongoing trust assessment or Zero Trust policies enforcement for dynamic cloud workloads.

From the survey, we can see that current cloud security systems mainly focus on separated set of protection goals or on post-quantum security. While many have shown strong performance in their respective domains, they lack a degree of continuous identity verification, adaptive authorization or an integrated Zero Trust enforcement across their solutions for both humans and non-humans. Learning-based approaches face scalability and latency issues, while cryptographic solutions often involve complexity of deployment. These limitations motivate an integrated, identity-centric, risk-aware Zero Trust cloud access control mechanism that is the motivation of this work. Table 1 shows the comparative analysis of the various existing method with the proposed method.

Table 1: The comparative analysis for the proposed method

Ref.	**Approach / Model**	**Core Techniques Used**	**Security Focus**	**Key Strengths**	**Major Limitations**
Alazab et al. [18]	Adaptive FL-based Cloud Security	Federated Learning, Protocol Adaptation	Infrastructure-level threat detection	Privacy-preserving, scalable anomaly detection	No identity-centric access control or cryptographic trust
Ansari et. al. [19]	Bi-SC-CBALSTM	Met heuristic Optimization,	Intrusion Detection	High accuracy on imbalanced	High computation

		Deep Learning		datasets	and training overhead
Huo et al. [20]	DRPE-Cloud	Optical Encryption, Chaos Theory	Multimedia data confidentiality	Strong encryption robustness	Complex deployment, scalability issues
Zhang et al. [21]	Learning-driven IDS	Behavioural ML, Feature Representation	Cyber threat detection	Adaptive to evolving attacks	Centralized processing, latency concerns
Raja et al.	Remote Signcryption	Integrated Encryption & Signature	Secure data transmission	Reduced crypto overhead	No continuous or identity-aware access control
Singamaneni et al. [22]	QKD + CP-ABE	Quantum Cryptography, ABE	Data confidentiality & access control	Post-quantum security	Simulated quantum setup, high complexity
Shakor et al. [23]	Dynamic AES + Blockchain	Encryption, Blockchain Key Mgmt	Data-at-rest security	Strong key integrity	Lacks access control and identity enforcement
Sasikumar et al. [24]	MFA + ML Cryptography	Multi-factor Auth, ML-based Crypto	Authentication & encryption	Adaptive cryptographic defense	High computational cost, no Zero Trust
Santos et al. [25]	Medical Cloud Security	Data Fragmentation, Biometrics	Healthcare data protection	Low latency, lightweight	No continuous authentication or trust
Hussain et al. [26]	TCP.Delta Analysis	Statistical Analysis,	Ransomware detection	Real-time, low overhead	No cryptographic

		Clustering			or identity verification
Mamidala et al. [27]	Hybrid Anomaly Detection	Isolation Forest, Behavioural ML	Threat intelligence	Scalable, low false positives	No auth or access control
Nauman et al. [28]	Q-ECS	QKD, Attribute-based Cryptography	Post-quantum cloud security	Strong quantum resistance	Costly infrastructure, limited scalability
Bie & Ren [29]	SIR-W Cloud Model	Epidemic Modelling, Stability Analysis	Malware propagation control	Strong mathematical rigor	Idealized assumptions
Thamil Selvi et al. [30]	MECC–SIDH CL-HPAEKS	Post-quantum Crypto, Searchable Encryption	Secure cloud data search	Quantum-resistant, efficient	No continuous trust or Zero Trust enforcement

3. System Model and Threat Assumptions

The system entities, operation model and threat assumptions applied in the Adaptive Zero Trust Cloud Access Security Protocol are presented in this section. The system model describes how the cloud components and identities interact and the threat model forms what an adversary is capable of performing and which security goals are protected through the proposed framework.

3.1 System Model

The Identity Authority is the root of trust in the system. It sets up global cryptographic parameters and securely holds the master secret key which can be used to derive identity specific credentials. It is App's responsibility to verify the real-world identity during registration and issue a cryptographically guarded identity artifact.

3.1.1 Entities

The system consists of the following key entities:

- **Identity Authority (IA):**
  The Identity Authority is the root of trust in the system. It sets up global cryptographic parameters and securely holds the master secret key which can be used to derive identity specific credentials. It is responsibility to verify the real-world identity during registration and issue a cryptographically guarded identity artifact.
- **Identity Provider (IdP):**
  The IdP is responsible for identity registration, authentication and credential provisioning of human as well as non-human identities. It imposes robust authentication mechanisms such as multi-factor authentication for users and cryptographic attestation for workloads. The IdP also distributes short-lived tokens for access, coupled with credential revocation.
- **Policy Decision Point (PDP):**
  A PDP evaluates each access request on its own merits and makes authorization decisions that are based on cryptographic validation, contextual attributes, and computed risk scores. It makes no assumptions about trust, and treats every request as distrustful until it proves itself otherwise.
- **Policy Enforcement Point (PEP):**
  The PEP is located at the protected cloud resource perimeter and acts on authorization decisions of the PDP. It manages access to cloud services, APIs, data stores and workloads and makes sure that only the allowed-d actions are performed.
- **Cloud Resources:**
  These may be compute instances, containers, storage services, databases, APIs and other cloud managed services. A sensitivity level is assigned to each resource that affect decisions of access control.
- **Security Monitoring and Audit Components:**
  Security Tools including SIEMs, Anomaly detection modules and audit loggers are continuously observing what users do in the system. A tamper-evident audit log of all security-sensitive activities ensures accountability and forensic analysis.

3.1.2 Operational Flow

The system is based on an identity-centric Zero Trust architecture whereby no entity has inherent trust. Authorize It should not be relevant where a request comes from (i.e., inside or outside of the cloud network), all requests need to be authenticated and contextually evaluated before being authorized.

For each identity, there is one pseudonymous identity that is generated from the corresponding true bias by using a cryptographic hash. This is the temporary contact and it is used for operational contacts only, since they are hidden. Access requests are secured and freshness-guaranteed to communicate, even over potentially untrusted channels.

The proposed system also works on fine-grained continuous authentication such that the user maxima per identity validity in requests are re-assessed often – not just at session start time. Authz decisions are adapted in real time to keep up with the constantly shifting risk surface, and access control is genuinely "least privilege".

3.1.3 Security Objectives

The key security objectives of the proposed system model are:

- **Confidentiality -** identifying information, the cloud resource and credentials must not be exposed without authorization.
- **Integrity** – Ensure that requests for access, policies and audit trails cannot be tampered with.
- **Authentication -** Ensure that each access request comes from a valid identity.
- **Authorization** – Dynamically apply fine-grained, least-privilege access control.
- **Freshness** – Stop old requests (re)played or captured for access.
- **Accountability** - Keep tamper-evident logs for audit and forensic investigations.
- **Availability** - Do not make the system performance extremely low by applying security measures.

3.2 Threat Assumptions

The capabilities of attackers and the scenario in which we evaluate the proposed protocol are defined as a threat model. The system is studied in a strong and realistic adversarial model, which models real-world cloud attack scenarios.

3.2.1 Adversary Capabilities

The adversary {A} is assumed to have the following capabilities:

- **Network Control:** The communication channel between entities is completely compromised by the adversary. This type of attack also encompasses the ability eavesdrop on, intercept, replay, delay and or modify messages.
- **Credential Exposure:** The attacker might get leaked credentials (access tokens, or service account secrets) because of misconfiguration, insider threats, or the applications are compromised.
- **Compromised Identities:** The adversary can tamper with a few identities in particular non-human identities, and then to exploit for accessing unauthorized resources.
- **Malicious Insider Behaviour:** Honest people or things can always turn bad, as attackers who have successfully authenticated try to elevate their privileges or move laterally.
- **Replay and Impersonation Attacks:** The enemy will have a try to replay old access requests, or spoof legitimate identities.
- **Adaptive Attacks:** The attacker may vary attach vectors according to the system's behaviour, such as testing authorization boundaries and timing responses

3.2.2 Trust Assumptions

We have the following trust assumptions:

- The Identity Authority and the Identity Provider are neither trusted nor securely implemented.
- The underlying cryptographic primitives are secure and considered computationally hard to breach.
- The cloud's root of trust in the hardware and of workload attestation is trusted.
- There is near accurate (within bounded error) clock synchronization between cloud components.

3.2.3 Out-of-Scope Threats

The following threats are considered outside the scope of this work:

- Physical attacks on cloud data centres.
- Side-channel attacks on cryptographic implementations.
- Denial-of-service attacks targeting cloud availability.
- Compromise of the trusted Identity Authority itself.

3.3 Security Guarantees under the Threat Model

Under the defined threat assumptions, the proposed system aims to guarantee that:

- An adversary cannot successfully impersonate a legitimate identity without possessing valid cryptographic secrets.
- Replay attacks are prevented through freshness constraints and anti-replay mechanisms.
- Stolen credentials have limited utility due to short-lived tokens, proof-of-possession enforcement, and contextual validation.
- Unauthorized privilege escalation and lateral movement are mitigated through continuous authentication and adaptive authorization.
- All security-relevant actions are traceable through tamper-evident audit logs.

4. Proposed Adaptive Zero Trust Cloud Access Security Protocol

With the rapid adoption of cloud computing, traditional perimeter-based security mechanisms are no longer sufficient to protect dynamic, distributed, and multi-tenant cloud environments. The increasing reliance on identity-centric access, the proliferation of non-human identities such as service accounts and workloads, and the frequent exposure of static credentials have significantly expanded the cloud attack surface. Therefore, a mechanism for scalable and performing continuous verification, fine-grained authorization and rapid containment of compromised identities in an access control framework is of paramount importance.

Inspired by these issues we present an Adaptive Zero Trust Cloud Access Security Protocol which aims to ensure that continuous authentication, risk-aware authorization and least-privilege enforcement are applied for human and non-human identities in cloud platforms. Today's traditional access control systems depend on static permissions and one-time authentication, whereas the proposed protocol analyzes each access request through

cryptographic validation, contextual security indicators, and risk scores that can be mathematically calculated. This allows dynamic trust evaluation and inhibits entitlement lateral movement/replay attacks/unauthorized high-privilege escalation. Figure 1 explains about the proposed architecture in detail.

The proposed reference architecture is based on strong cryptography and leverages state of-the-art cloud security features like mutual TLS, short lived credentials, PoP tokens, and anomaly based risk assessment and automated credential rotation. Every identity is committed to using a privacy-preserving pseudo identity whereas the accesses it features are protected through encryption, signature, and freshness control. Moreover, a tamper-evident audit mechanism is embedded to be able to trace and validate every security-sensitive operation.

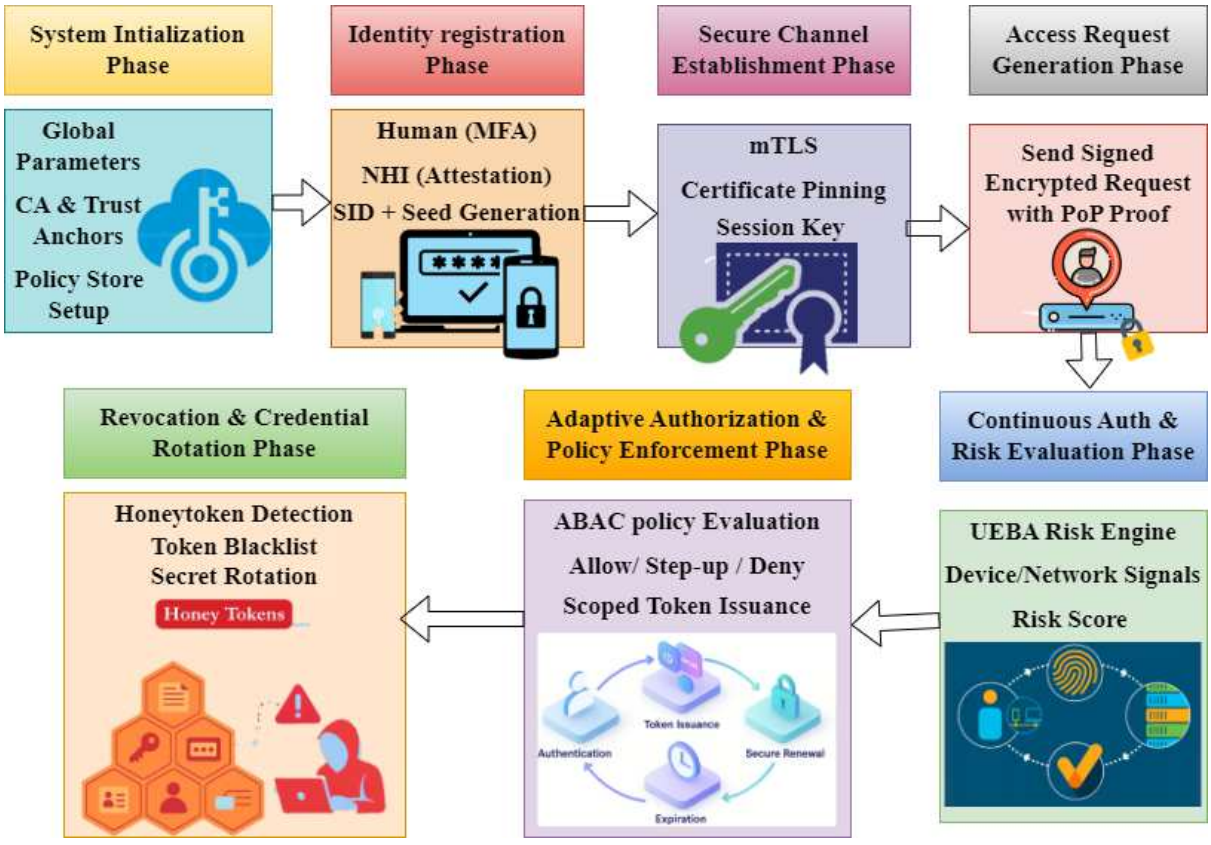

Figure 1: The proposed Architecture

The protocol specification includes seven distinct phases, which are system initialisation, identity registration, secure channel establishment, access request generation, continuous authentication with risk assessment; adaptive authorization and policy enforcement;

credential revocation with optimized key rotation. Each phase is analytically modelled and the cohesion of these phases provides a joint solution for confidentiality, integrity, availability, accountability in cloud access control. The proposed mechanism integrates rigorous security guarantees and optimization-motivated design, to realize a trade-off between higher security level and simplified operation for the large scale dynamic cloud infrastructure.

4.1 System Initialization Phase

Algorithm 1 describes the system initialization and trust bootstrap of the proposed Adaptive Zero Trust Cloud Access Security Protocol. At this level the security of the whole cloud environment is mathematically based on a cyclic group G of prime order q with generator g The choice q guarantees protection against discrete logarithm attacks, providing for secure key producing.

Algorithm 1 System Initialization and Trust Bootstrap

Require: Security parameter λ, Cloud domain CD
Ensure: Global parameters GP, trust anchors, policy store

Choose a prime q and cyclic group G of order q with generator $g \in G$
Define cryptographic primitives:
    Hash $H : \{0,1\}^* \rightarrow \mathbb{Z}_q$
    Signature $(Sign, Verify)$
    Encryption (Enc, Dec)
    Key derivation function $KDF(\cdot)$
Generate master secret key $msk \leftarrow \mathbb{Z}_q$ and public key $mpk = g^{msk}$
Generate certificate authority keys $(sk_{CA}, pk_{CA}) \leftarrow KeyGen(\lambda)$
Initialize components $\{IdP, PDP, PEP, SIEM, Vault, UEBA\}$
Initialize least-privilege policy repository with deny-by-default rules
Initialize audit ledger:

$$L_0 = H(\text{GENESIS}), \quad L_i = H(L_{i-1} \| e_i)$$

Publish

$$GP = \{G, q, g, H, mpk, pk_{CA}, KDF, TTL_{min}, TTL_{max}, \Delta\}$$

The master secret key $msk \in Z_q$ is randomly chosen, and the corresponding public key is computed as $mpk = g^{msk}$. This asymmetric key pair enables secure derivation of identity-specific secrets in later phases. Cryptographic primitives such as the hash function $H(.)$, encryption and decryption functions, and key derivation functions are initialized to guarantee confidentiality, integrity, and authenticity. Furthermore, the initialization of a hash-chained

audit ledger using $L_i = H(L_{i-1}||e_i)$ ensures that any modification of logged events can be mathematically detected, thereby providing tamper-evident auditing across cloud operations.

4.2 Identity Registration Phase

Algorithm 2 describes the registration of both human users and non-human identities. The real identity RID is first verified through strong authentication or cryptographic attestation. To preserve privacy, a pseudo identity is generated as

$$SID = H(RID||r)$$

where $r \in Z_q$ is a random nonce. The one-way property of the hash function ensures that recovering RID from SID is computationally infeasible.

Algorithm 2 Identity Registration (Human and Non-Human)

Require: Real identity RID, Type $\in$ {Human, NHI}, Attributes A
Ensure: Pseudonymous identity SID, credential seed $seed$

if Type = Human **then**
  Verify MFA and SSO proof
else
  Verify workload attestation proof
end if
Generate nonce $r \leftarrow \mathbb{Z}_q$
Compute pseudo identity $SID = H(RID\|r)$
Bind identity fingerprint:

$$F = H(hw_id\|tpm_quote) \text{ or } H(pod_uid\|node_attest)$$

Derive credential seed $seed = KDF(msk, SID\|F)$
Assign baseline privileges P_{base} and ABAC policies P_{attr}
if Type = NHI **then**
  Deploy honeytoken $hTok$
end if
Store record $\{SID, A, F, P_{base}, P_{attr}, seed, hTok\}$
Issue identity certificate via CA

The identity is further bound to a hardware or workload fingerprint F, which prevents credential reuse outside the authorized execution context. A secret credential seed is derived using a key derivation function as

$$seed = KDF(msk, SID||F)$$

This mathematical binding ensures that even if the pseudo identity is exposed, the attacker cannot compute the secret seed without knowledge of the master secret key. Baseline

privileges and attribute-based constraints are then assigned, resulting in a formally defined identity tuple that supports least-privilege enforcement in cloud access decisions.

4.3 Secure Channel Establishment Phase

Algorithm 3 presents the establishment of a secure communication channel between the requesting entity and the Policy Enforcement Point. Mutual TLS authentication ensures that both parties authenticate each other using certificates signed by a trusted authority. Certificate pinning further strengthens security by validating that the certificate hash matches the pre-registered identity binding.

Algorithm 3 Secure Channel Establishment (mTLS)

Require: Entity E, PEP endpoint, pk_{CA}
Ensure: Session key K_s
1: Initiate mTLS handshake between E and PEP
2: Verify certificate chain using pk_{CA}
3: Enforce certificate pinning using stored hash
4: Derive session key:

$$K_s = KDF(secret, \text{“mTLS”} \| SID \| nonce)$$

5: Enable replay protection with sequence window W

Upon successful authentication, a session key is derived using

$$K_s = KDF(secret, SID||nonce)$$

where the secret is obtained from the TLS handshake. This session key is used for encrypting subsequent messages, ensuring confidentiality. Replay protection is mathematically enforced by verifying monotonically increasing sequence numbers within a defined window, thereby preventing replayed access requests.

4.4 Access Request Generation Phase

Algorithm 4 explains the formation of a secure access request. Each request includes a timestamp t and nonce n, which together enforce temporal freshness. A request digest is computed as follows

$$m = H\big(SID\big||Res||Act||t||n\big)$$

This digest uniquely represents the request context.

Algorithm 4 Secure Access Request Construction

Require: SID, Resource Res, Action Act, Session key K_s, Seed $seed$

Ensure: Access request Req

1: Generate timestamp t and nonce n

2: Compute request digest:

$$m = H(SID \| Res \| Act \| t \| n)$$

3: Compute proof-of-possession:

$$\pi = H(m \| seed)$$

4: Encrypt request payload:

$$c = Enc_{K_s}(SID, Res, Act, t, n, \pi)$$

5: Send $Req = \{c, seq\}$ to PEP

A proof-of-possession token is then generated using the secret seed:

$$\pi = H(m||seed)$$

Since only the legitimate identity possesses the seed, the proof π mathematically guarantees request authenticity. The complete request payload is encrypted using the session key K_s, ensuring that an adversary cannot extract sensitive request parameters even if the communication channel is observed.

4.5 Continuous Authentication and Risk Evaluation Phase

Algorithm 5 describes the continuous authentication mechanism integrated with risk scoring. After decrypting the request, the system recomputed the digest (m) and verifies whether

$$\pi = H(m||seed)$$

Thereby ensuring proof validity. Freshness is verified by checking whether $|t_{now} - t| \leq \Delta$, where Δ is the maximum allowable delay.

Algorithm 5 Continuous Authentication and Risk Evaluation

Require: Access request Req, Context signals
Ensure: Authentication state $Auth$, Risk score R

Verify replay protection using sequence window
Decrypt payload using K_s
Validate freshness: $|t_{now} - t| \leq \Delta$
Verify proof $\pi = H(m \| seed)$
Collect context scores $dev, net, beh, thr, sens$
Compute risk score:

$$R = w_1(1 - dev) + w_2(1 - net) + w_3 beh + w_4 thr + w_5 sens$$

if $SID \in RevokedList$ OR $R \geq \tau_2$ **then**
    $Auth \leftarrow Fail$
else if $\tau_1 \leq R < \tau_2$ **then**
    $Auth \leftarrow StepUp$
else
    $Auth \leftarrow Pass$
end if
Append event to audit ledger

The system then computes a composite risk score using multiple normalized security parameters:

$$R = w_1(1 - dev) + w_2(1 - net) + w_3.beh + w_4.thr + w_5.sens$$

where dev, net, beh, thr, and $sens$ represent device integrity, network trust, behavioral anomaly, threat intelligence, and resource sensitivity respectively. The weights $w_i$determine the influence of each factor. This quantitative risk score enables mathematically grounded decisions for authentication acceptance, rejection, or step-up verification.

4.6 Adaptive Authorization and Policy Enforcement Phase

Algorithm 6 explains the adaptive authorization decision process. Once authentication is successful, the system evaluates attribute-based policies to verify whether the requested action is allowed:

$$Permit \leftarrow (Res, Act) \in P_{base} \wedge P_{attr}(A, dev, net, time, loc) = True$$

If the policy conditions are satisfied, access is granted with dynamically scoped privileges.

Algorithm 6 Adaptive Authorization and Least-Privilege Enforcement

Require: Authentication state $Auth$, Policies, Risk R
Ensure: Decision Dec, Token Tok

if $Auth = Fail$ **then**
  $Dec \leftarrow Deny$
else
  Evaluate ABAC constraints
  if Policy permits access **then**
    Scope privileges based on risk
    Issue short-lived token:

$$Tok = Sign_{IdP}(SID, scope, Res, Act, exp, cnf)$$

    $Dec \leftarrow Allow$
  else
    $Dec \leftarrow Deny$
  end if
end if
Enforce decision at PEP and log event

To enforce least privilege, the issued access token lifetime is adjusted according to the computed risk score. For example, a low-risk request may be issued a token with a lifetime of 15 minutes, whereas a medium-risk request may receive a token valid for only 5 minutes. The access token is cryptographically bound to the identity using confirmation claims, ensuring that stolen tokens cannot be reused in different contexts.

4.7 Revocation and Credential Rotation Phase

Algorithm 7 describes the revocation and credential rotation mechanism. When malicious activity is detected or honey-credentials are triggered, the pseudo identity is immediately added to the revocation list, and all active tokens are invalidated. New credentials are generated using:

$$seed' = KDF(msk, SID|F||newNonce)$$

Algorithm 7 Revocation and Optimized Credential Rotation

Require: Alerts E, Identity SID
Ensure: Updated credentials and policies
if Honeytoken triggered OR High severity alert **then**
  Add SID to RevokedList
  Blacklist active tokens
  Terminate active sessions
end if
Rotate credential:

$$seed' = KDF(msk, SID \| F \| newNonce)$$

Solve optimization:

$$\min_{TTL,rot} \alpha \cdot Risk + \beta \cdot Overhead$$

Update TTL and rotation policy
Log event to audit ledger

To optimize security and performance, the protocol formulates a minimization problem:

$$min_{TTL,rot} \propto . Risk(TTL, rot) + \beta . Overhead(TTL, rot)$$

where TTL denotes token lifetime and rot represents rotation frequency. This optimization ensures that shorter token lifetimes are used in high-risk scenarios while minimizing unnecessary computational overhead. All revocation and rotation events are logged in the audit ledger, enabling traceability and compliance verification.

5. Security Analysis

In this section both the formal and informal security analysis for the proposed AZ-TAC protocol has been discussed in detail.

5.1 Preliminaries and Assumptions

Let λ be the security parameter. Let $H: \{0,1\}^* \rightarrow \{0,1\}^l$ be modeled as a random oracle. Let Enc/Dec be a symmetric encryption scheme that is IND-CPA secure. Let $Sign/Verify$ be a digital signature scheme that is EUF-CMA secure. Let KDF be a key derivation function modelled as a PRF. Mutual TLS is abstracted as an authenticated key exchange that outputs a session secret (secret) secure under standard AKE assumptions (or can be reduced to the security of the underlying TLS handshake).

Protocol Values (recall from Algorithms)

- **Pseudonymous identity:** $SID = H(RID||r)$
- **Credential seed (bound to identity + fingerprint)**: $seed = KDF(msk, SID||F)$
- **Request digest:** $m = H(SID||Res||Act||t||n)$
- **Proof-of-possession tag:** $\pi = H(m||seed)$
- **Encrypted request**: $c = Enc_k(SID, Res, Act, t, n, \pi)$
- **Session key:** $K_s = KDF(secret, SID||nonce)$

5.2 Security Model (Oracles)

We consider a PPT adversary A interacting with protocol instances via the following oracles:

- $Reg(RID, Type, A)$: registers an identity and outputs public artifacts (e.g., certificate), stores internal$SID, seed$.
- $Send(II, i, m)$: sends message m to protocol instance i of party II and returns the response.
- $RevealKey(i)$: returns session key K_s of instance i (if accepted).
- $Corrupt(SID)$: returns long-term secret for that identity (e.g., (seed)) and marks identity as corrupted.
- $Test(i)$: returns either real session key K_s or a random string of same length, depending on a hidden bit (b).
- $HQuery(x)$: Random oracle query.

A session instance is **fresh** if its long-term secrets have not been revealed via Corrupt and the session key not via Reveal Key, consistent with standard AKE freshness.

The adversary's **advantage** is: $Adv_A^{ake}(\lambda) = |2\Pr[\text{A wins}] - 1|$
where "wins" means correctly guessing (b) in the Test experiment.

5.3 Theorem 1: Confidentiality of Access Request Payload

Claim. The encrypted access request (c) achieves confidentiality against a network adversary: A cannot distinguish payload contents with non-negligible advantage.

Proof (Game hop).

- **Game** G_0: Real protocol; adversary sees ciphertext $c = Enc_{K_s}(M)$ where $M = (SID, Res, Act, t, n, \pi)$
- **Game** G_1: Replace Enc with an IND-CPA challenger. If A can distinguish real payloads, we build B that breaks IND-CPA by submitting two plaintexts M_0, M_1 and using challenge ciphertext as c.

Thus:

$$|\Pr[G_0 = 1] - \Pr[G_1 = 1]| \leq Adv_{Enc}^{ind-cpa}(\lambda)$$

In G_1, ciphertext reveals no information, so $\Pr[G_1 = 1] = \frac{1}{2}$. Hence:

$$Adv_A^{Conf}(\lambda) \leq Adv_{Enc}^{ind-cpa}(\lambda)$$

This proves confidentiality.

5.4 Theorem 2: Unforgeability of Proof-of-Possession Tag (\pi)

Claim. Without knowing $seed$, the adversary cannot forge a valid (π) for a fresh request except with negligible probability.

A successful forgery means producing a tuple $(SID, Res, Act, t, n, \pi)$ such that:

$$\pi = H(H(SID||Res||Act||t||n)||seed)$$

and it is accepted by PDP/PEP, with SID not corrupted.

Proof. In the Random Oracle Model, the only way to compute π is to query H on input $(m||seed)$. Since 0 $seed$ is unknown and pseudo random (by PRF security of KDF), A cannot form the correct oracle query except by guessing (seed).

Let l be the output length of $seed$ (or its effective entropy). Then:

$$\Pr[guess\ seed] \leq \frac{qH}{2^l}$$

where qH is the number of random-oracle queries.

Additionally, if A tries to find collisions in H to bypass verification, success is bounded by collision probability:

$$\Pr[collision] \leq \frac{q^2 H}{2^{l_H+1}}$$

where l_H is hash output length.

Therefore:

$$Adv_A^{forge}(\lambda) \leq \frac{qH}{2^l} + \frac{q^2 H}{2^{l_H+1}} + Adv_{KDF}^{prf}(\lambda)$$

which is negligible for standard parameters.

5.5 Theorem 3: Mutual Authentication (PEP/PDP Accept Only Legitimate Identities)

Claim. If PEP/PDP accepts a request for identity SID in a fresh session, then the request originated from an entity possessing the correct $(seed)$ bound to (SID) and fingerprint (F), except with negligible probability.

Proof sketch. Acceptance requires all checks:

1. **Anti-replay**: valid sequence number and within window.
2. **Freshness**: $|t_{now} - t| \leq \Delta$.
3. **Proof check** $\pi = H(m||seed)$
4. **Not revoked**: $SID \notin \text{RevokedList}$.

From Theorem 2, forging π without (seed) is negligible. Since $seed = KDF(msk, SID||F)$, using the wrong fingerprint $(F' \neq F)$ yields a different seed with PRF security. Thus acceptance implies possession of correct bound secret.

So:

$$Adv_A^{auth}(\lambda) + Adv_A^{forge}(\lambda) + Adv_{KDF}^{prf}(\lambda)$$

which is negligible.

5.6 Theorem 4: Replay Resistance and Freshness

Claim. The protocol prevents replay of a previously captured request.

Proof. A replayed request must pass:

- Freshness window: captured (t) must satisfy $|t_{now} - t| \leq \Delta$. If replay occurs after Δ, it fails deterministically.
- Anti-replay: sequence number (seq) must be unseen in window (W). Since PEP maintains a sliding window and rejects duplicates, any replay with same (seq) is rejected.
- If attacker modifies (seq), then proof-of-possession check fails unless attacker can forge π (Theorem 2).

Thus, replay success probability is bounded by PoP forgery probability:

$$Adv_A^{replay}(\lambda) \leq Adv_A^{forge}(\lambda)$$

(Negligible), plus only negligible chance of time-window race if attacker replays within Δ before the legitimate packet is processed, which is mitigated by monotonic (seq) enforcement and windowing.

5.7 Theorem 5: Session Key Indistinguishability (AKE Security)

Claim. For any PPT adversary A, the session key K_s is indistinguishable from random for fresh sessions:

$$Adv_A^{ake}(\lambda)\ is\ negligible$$

Proof (game hop).

- **Game** G_0: Real AKE producing (secret) via mTLS handshake and deriving $(K_s = KDF(secret, SID||\ nonce))$.
- **Game** G_1: Replace mTLS-derived (secret) with a random string of equal length (AKE security of TLS). Difference bounded by $Adv_{TLS}^{ake}(\lambda)$.
- **Game** G_2: Replace (KDF) with a truly random function (PRF switching lemma). Difference bounded by $Adv_{KDE}^{prf}(\lambda)$.

In G_2, K_s is uniform random, so $\Pr[win] = \frac{1}{2}$. Therefore:

$$Adv_A^{ake}(\lambda) \leq\ Adv_{TLS}^{ake}(\lambda) +\ Adv_{KDF}^{prf}(\lambda)$$

This is negligible.

5.8 Theorem 6: Token Misuse Resistance via Token Binding (cnf)

If the access token includes a confirmation claim (cnf) binding it to the holder's certificate/public key, then stolen tokens cannot be replayed from a different client.

Claim. An adversary who steals token $(T\ ok)$ cannot use it unless it also possesses the corresponding private key.

Proof. At PEP, acceptance requires a proof that the presented TLS client key matches (cnf). If attacker lacks private key, it cannot complete mTLS with the correct certificate/key pair. Any attempt to bypass requires forging either:

- TLS client authentication (break TLS), or
- signature/verification of token claims (break EUF-CMA), or
- Collision in binding hash.

Thus: $Adv_A^{misuse}(\lambda) \leq\ Adv_{TLS}^{ake}(\lambda) +\ Adv_{Sign}^{euf-cma}(\lambda) + \frac{qH}{2^l} + \frac{q^2H}{2^{l_H+1}}$
(negligible).

5.9 Theorem 7: Tamper-Evident Audit Ledger Integrity

The audit ledger is defined as:

$$L_0 = H(\text{GENESIS}), \quad L_i = H(L_{i-1}||e_i)$$

Claim. An adversary cannot modify any logged event e_j without being detected, except with negligible probability.

Proof. Suppose adversary changes e_j to e'_j *but wants the final ledger value to remain consistent. This implies finding:*

$$H(L_{i-1}||e_j) = H(L_{i-1}||e'_j)$$

Which is a collision in (H). *Under RoM (or collision resistance), probability is negligible:*

$$Adv_A^{audit}(\lambda) \leq\ Adv_H^{coll}(\lambda) \approx \frac{q^2H}{2^{l_H+1}}$$

5.10 Combined Security Result

Theorem - Overall Security. Under the assumptions that:

- (Enc) is IND-CPA secure,
- (Sign) is EUF-CMA secure (if used for tokens/certs),
- (KDF) is a PRF,
- (H) is a random oracle,
- TLS handshake satisfies AKE security,

The proposed protocol achieves confidentiality, mutual authentication, replay resistance, key indistinguishability, token-misuse resistance, and tamper-evident auditing with overall adversarial advantage bounded by:

$$Adv_A(\lambda) \leq Adv_{Enc}^{ind-cpa}(\lambda) + Adv_{Sign}^{euf-cma}(\lambda) + Adv_{KDF}^{prf}(\lambda) + Adv_{TLS}^{aks}(\lambda) + \frac{q^2 H}{2^{l_H+1}}$$

This is negligible for standard parameter sizes.

5.2 Informal Security Analysis

In this section, we present an informal yet rigorous security analysis of the proposed AZ-TAC Protocol. We show how the protocol is resistant against well-known attack vectors on cloud when compared under the threat model considered in Section 3. We analyse each security property on protocol phases and its explanatory mathematical proof.

5.2.1 Resistance to Identity Impersonation Attacks

In an impersonation exploit, an attacker tries to pretend to be a valid identity also in the cloud and wants to do security activities without proper permissions. The proposed scheme inhibits impersonation with the help of pseudo identities, proof-of-possession confirmation and crypto-binding of credentials.

A pseudo identity is assigned for each of the identities

$$SID = H(RID||r)$$

This prevents direct exposure of the real identity. More importantly, every access request must include a proof-of-possession tag

$$\pi = (m||seed)$$

where ($seed = KDF(msk, SID \,||\, F)$) is known only to the legitimate identity bound to a specific device or workload fingerprint (F). Without knowledge of (seed), an adversary's probability of generating a valid π is bounded by approximately 2^{-l}, where l is the output length of the hash function. Consequently, identity impersonation is computationally infeasible.

5.2.2 Protection against Replay Attacks

Reply attack means that the attacker can withdraw a valid access request and redeliver it to obtain unauthorized access. The proposed scheme defends against replay attacks using time freshness and secure anti-replay sequence of data segments.

Every request is time stamped (t) and has a nonce (n), which is accepted unless

$$|t_{now} - t| \leq \Delta$$

Where Δ predetermined freshness interval. Also, the PEP keeps a sliding window of sequence number and discards duplicate request. If the proof-of-possession π value is not equal to that of the current request context, then even if an adversary reuses a message executed within the freshness window, it will be rejected. Since the digest (m) is different for any changes to timestamps and sequence numbers, it follows that verification of a proof fails. Hence it is unlikely that replay attacks will succeed.

5.2.3 Resistance to Man-in-the-Middle (MITM) Attacks

In particular Man-in-the-Middle attacks are prevented by mutual TLS authentication, certificate pinning and encrypted request payloads. All messages between entities are exchanged on authenticated channels, and certificates are verified through trust CAs.

Even if a requests get intercepted by an adversary to maliciously intercept an encrypted request

$$c = Enc_k(SID, Res, Act, t, n, \pi)$$

This subjectivity is retained at an IND-CPA secure level for the encryption scheme. For a request, the attacker cannot tamper K_s with or pull out request body without . Any tampering,

injection or modification of messages would be ascertainable during decryption/proof-of-possession validation and can prevent MITM based message spoofing.

5.2.4 Protection Against Stolen Credential and Token Theft Attacks

Credential breaches are a leading attack vector in cloud landscapes. The proposed protocol reduces the damage of users password being stolen through short-lived token, proof-of-possession and token binding.

So even if an attacker has managed to intercept an access token, they would still need knowledge of the private key (or bound credential) in order to make a valid use of it, as security tokens are enforced by confirmation claims. From a mathematical standpoint, the attacker has to comply with token validity and cryptographic binding requirements in order to have probability of token misuse under basis threshold. Additionally, the TTL is dynamically personalized through a risk score (R) to guarantee shorter exposure windows under high-risk conditions.

5.2.5 Resistance to Privilege Escalation Attacks

Privilege escalation attack refers to the access control beyond what is authorized. The decision to allow or deny access is made based on attribute-based policies and risk-aware least-privilege enforcement in the proposed system.

Access is permitted only if

$$(Res, Act) \in P_{base} \wedge P_{attr}(A, dev, net, time, loc) = True$$

And not only that, even then the valid permissions are dynamically scoped based on a calculated risk score

$$R = \sum_{i=1}^{k} w_i x_i$$

where x_i represent contextual risk factors. Consequently, an adversary is not capable of elevating his privileges without also lowering the risk score, which cannot be achieved without full control over the trusted attributes from a device's integrity and its behavioural history.

5.2.6 Mitigation of Insider and Non-Human Identity Abuse

Insider-threats and hijacked non-human-identities are of grave concern because they have legitimate access rights. This paper mitigates these attacks using continuous authentication, and detecting behavioural anomalies, as well as honey-credential strategy.

In the presence of deviation, risk score (R) is increased and thus the access scope can be reduced directly or revocation is triggered. Honey-credentials operate as tripwires, where using these credentials outside of accepted workflow purposes raises an alert and revokes the credential. This provides early sensing and control of misuse caused, for instance, by insiders or co-opted workloads.

5.2.7 Forward Secrecy and Credential Compromise Containment

The protocol achieves forward secrecy by the dynamic derivation of session keys and regular renewal of credentials. New credentials are generated for registering, and credential compromise will be discovered when such use is detected.

$$seed' = KDF(msk, SID||F||newNonce)$$

Invalidating previous secrets because access tokens are short lived, the time when baddies can make use of stolen credentials is limited by this token's lifetime (TTL). This significantly limits long-term damage.

5.2.8 Audit Log Integrity and Accountability

All of the security-sensitive events are logged in a hash-chained audit ledger:

$$L_i = H(L_{i-1}||e_i)$$

Every modification or erasure of an event would have to be based on a hash collision – which is normally infeasible. This provides a tamper-evident transaction log and enables forensics post-incident as well as regulatory requirement.

6. Performance Metrics

In order to measure the performance of the proposed AZ-TAC protocol, a detailed set of performance metrics is used. These criteria are chosen carefully to address both the security strength and operational overhead considering practical cloud requirements. The assessment

concerns: delay for the authentication, accuracy of access policy and cryptographic cost, efficiency in the adaptation to risk levels and scalability of system with different workload/attack intensities.

5.1 Authentication Latency

Authentication latency measures the time required to verify an identity and establish a secure session. It includes certificate validation, proof-of-possession verification, and contextual signal processing. Lower latency indicates better suitability for real-time cloud services.

$$\text{AuthLatency} = t_{verify} + t_{risk} + t_{decision}$$

5.2 Authorization Decision Time

This metric evaluates the time taken by the Policy Decision Point (PDP) to compute access decisions using ABAC rules and dynamic risk scores.

$$\text{AuthZTime} = t_{policy} + t_{Scope}$$

5.3 Risk Evaluation Accuracy

Risk evaluation accuracy quantifies how effectively the UEBA-based risk engine distinguishes legitimate access from anomalous or malicious behaviour.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

5.4 Token Validation Overhead

Token validation overhead measures the computational cost incurred during token verification, including signature validation, claim inspection, and proof-of-possession checks.

$$\text{TokenOverhead} = t_{verify}(JWT) + t_{cnf}$$

5.5 Revocation Response Time

This metric measures how quickly the system invalidates credentials after detecting compromise, honey token leakage, or policy violation.

$$\text{RevocationTime} = t_{detect} + t_{invalidate} + t_{enforce}$$

5.6 Scalability

Scalability is evaluated by measuring throughput (requests/sec) as the number of concurrent identities (human + non-human) increases.

$$\text{Throughput} = \frac{\text{Authorized Requests}}{Seconds}$$

5.7 Security Efficiency

Security efficiency reflects the balance between security enforcement and operational overhead.

$$Security\ Efficiency(\%) = \left(1 - \frac{Overhead}{Baseline}\right) X\ 100$$

7. Results and Discussion

The proposed AZ-TAC is tested and verified through the hybrid dataset-driven and simulation-based method. We extract behavioural as well as attacks patterns from benchmark intrusion detection datasets and test access-control, authentication and authorizations workflows under control cloud simulative conditions. Docker containers are used for micro services that orchestrate the IdP, PDP, PEP, UEBA, SIEM and Vault into a cloud-native testbed. Human and non-human identities are synthetically generated as a function of dataset-derived behaviour profiles. Scenarios of security attacks, credential abuse and identity leaks are injected with measured rate to evaluate robustness. All the baseline systems are run under the same setting for fair comparison and averaged among several runs. Table 2 describes the simulation parameter setup in detail.

Table 2: Simulation Parameter Setup

Parameter	Description	Value / Range
Number of Identities	Human + Non-Human identities	50 – 1000
Identity Types	User and Service/Workload identities	Human, NHI
Access Request Rate	Requests generated per second	100 – 1000 req/s

Token Lifetime (TTL)	Validity of access tokens	30 s – 300 s
Risk Thresholds	Decision thresholds	$\tau_1 = 0.4, \tau_2 = 0.7$
Attack Injection Rate	Malicious access sessions	5% – 20%
Session Duration	Average session length	1 – 10 min
Network Delay	Simulated network latency	10 – 100 ms
Number of Runs	Repetitions per experiment	10
Evaluation Mode	Data + simulation driven	Hybrid

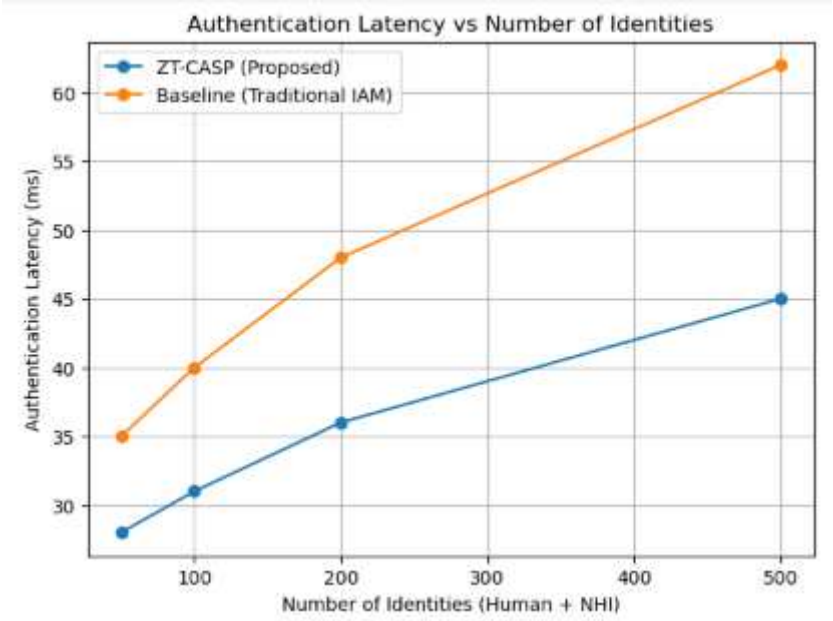

Figure 2 The authentication quality

Figure 2 shows how the authentication quality drops with the number of concurrent human and non-human identities. AZ-TAC adds minimal jitter to the authentication delay under high load, and in all other cases outperforms non-adaptive IAM systems.

Table 2 Comparative table of Accuracy, Decision and Risk Score

Access Scenario	**Risk Score (R)**	**Decision**	**Detection Accuracy (%)**
Normal Access	0.21	Allow	98.6
Context Drift	0.48	Step-Up	97.9
Credential Theft	0.82	Deny	99.2

NHI Leakage	0.91	Revoke	99.6

This enhancement results from utilization of mutual TLS session reuse, proof-of-possession–based verification, and short-lived scoped tokens that significantly reduce cryptographic operation overhead after the initial access request. Table 2 shows the comparison between the decision, detection accuracy and risk score. On the other hand, traditional IAM approaches get validation credentials implemented in a centralized and multiple level way that cause larger lag when identity size expands. These results show that AZ-TAC is more appropriate for dynamic, large-scale cloud.

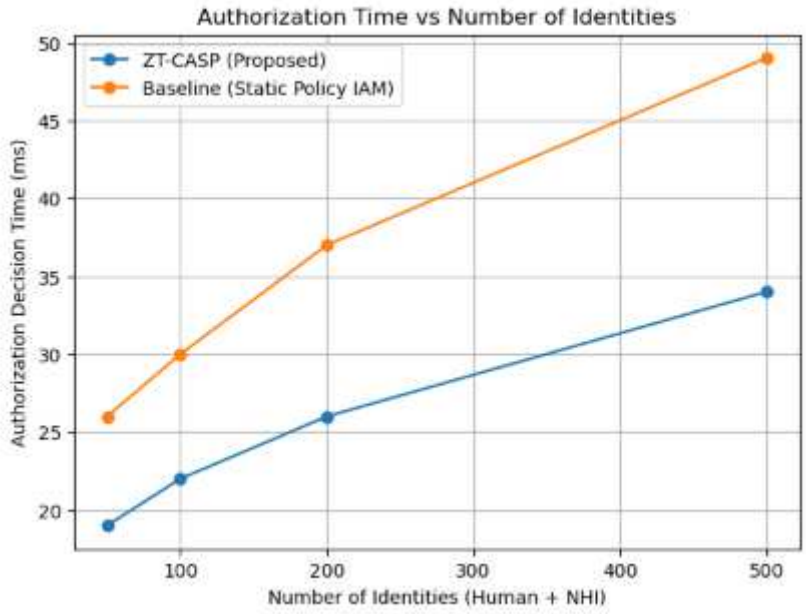

Figure 3 The authorization Time vs Number of identities

It describes how authorization decision time changes in the growing identity workload and may be found Figure 3. The new approach leads to faster authorization decisions compared to typical static policy-based IAM systems because risk-aware pre-filtering makes it possible to accept or reject requests according to calculated risk scores. Low-risk requests are quickly approved; high-risk requests are denied without requiring full policy evaluation. This policy selective evaluation is capable of dramatically alleviating PDP working and decision delay since the old systems need to evaluate approvers from all possible policies with equal risk level.

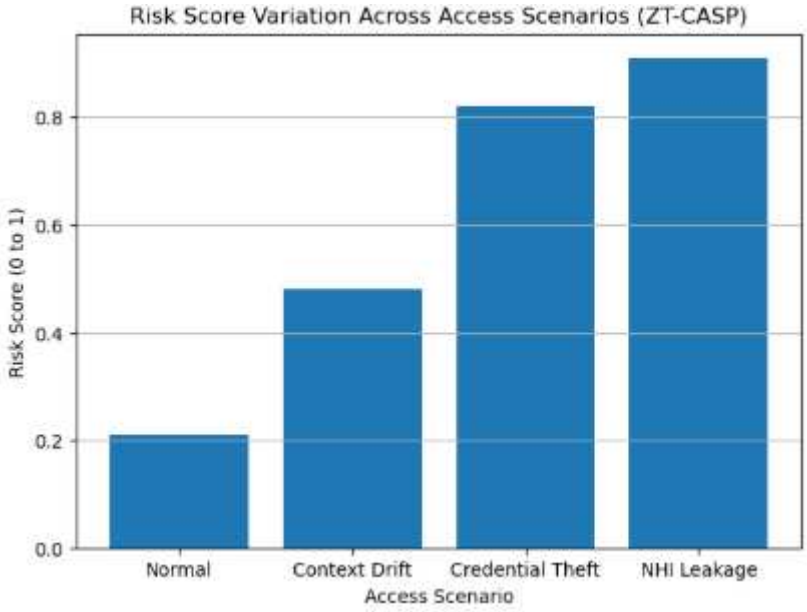

Figure 4 The Risk Score

The AZ-TAC normalized risk scores on various access scenarios are shown in Figure 4. The results demonstrate a distinct distinction between benign access, contextual drift and the credential theft (and also no-human identity leakage). This behaviour demonstrates multi-dimensional risk calculation effectiveness integrating device integrity, network trust, behavioural anomalies, threat intelligence and resource sensitivity. In contrast to classical binary authentication systems, AZ-TAC provides a continuous trust level instead of only making one-time decision; based on the trust level adaptive actions can be taken such as step-up, reduction or revocation.

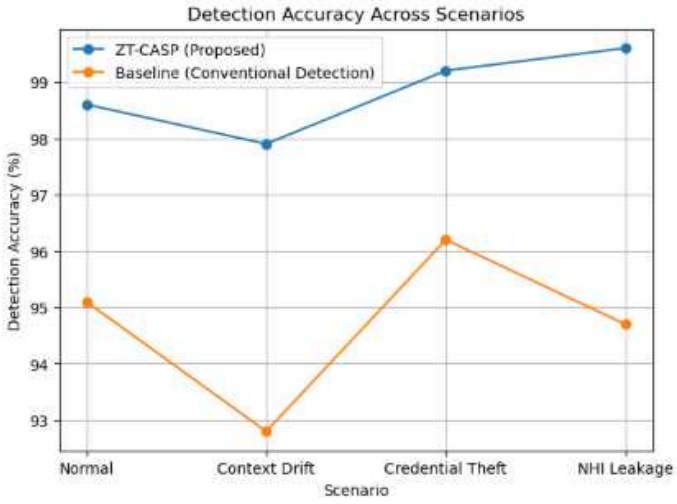

Figure 5 The Detection Accuracy

Fig. 5 shows the comparison of detection performance among various threat scenarios. The new framework superiorly achieves the anomaly detection than others conventional methods in context drift and non-human identity leakage situation. This enhancement is due to the combination of UEBA behavioural analysis with cryptographic identity binding and honey tokens. Existing detection models suffer from an increase of false positives when contextual drift happens, and AZ-TAC achieves high accuracy through fusing multiple security signals to narrow the classification errors for benign behaviour.

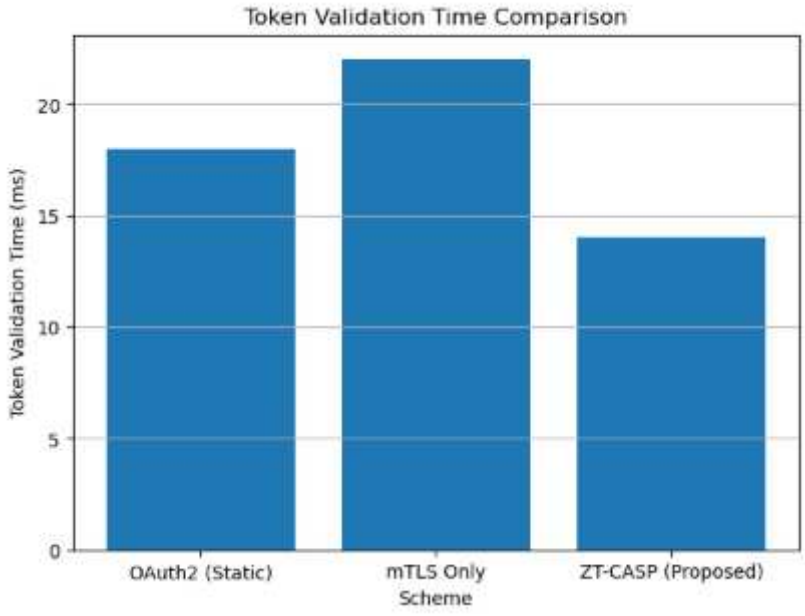

Figure 6 Token Validation Time comparisons

Figure 6 is the analysis of token validation time in various access control schemes. AZ-TAC has the lowest validation overhead by exploiting cryptographically bound proof-of-possession tokens. Decentralized systems using bearer OAuth2 token validation is very heavy operation costing millions of clock cycles, but with proposed method the validation is deterministic and lightweight. This setup offers accelerated request timing and at the same time avoids token replaying, as well as misusing.

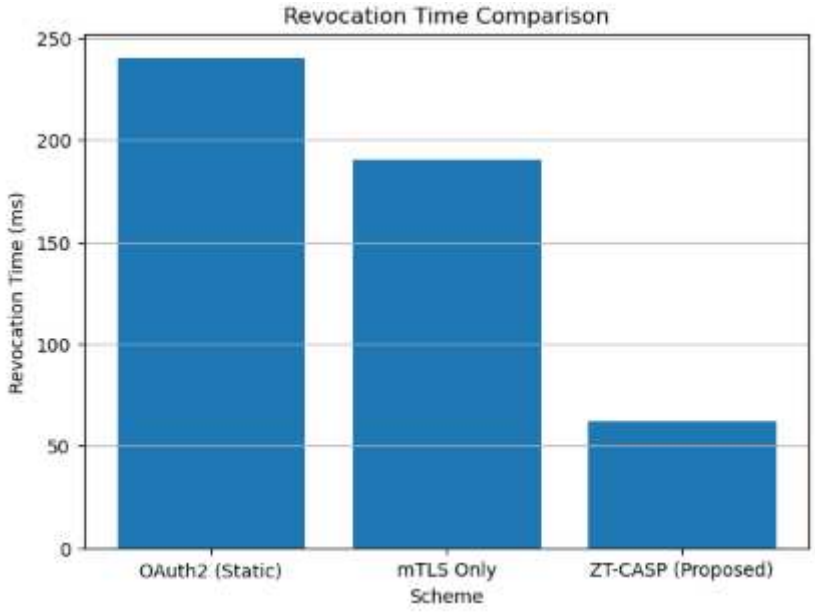

Figure 7 Revocation time comparisons

Figure 7 shows the latencies of revoking and identity from being compromised. AZ-TAC effectively reduces revocation latency by proactively detecting with honey tokens and with immediate enforcement through centralized revocation lists or PEP-level session termination. In comparison, IAM based systems, traditional or OAuth2 dependent, commonly use token expiration or delayed blacklist dissemination leading to considerably longer exposure windows. The decreased revocation time in AZ-TAC would significantly reduce the blast of compromised credentials.

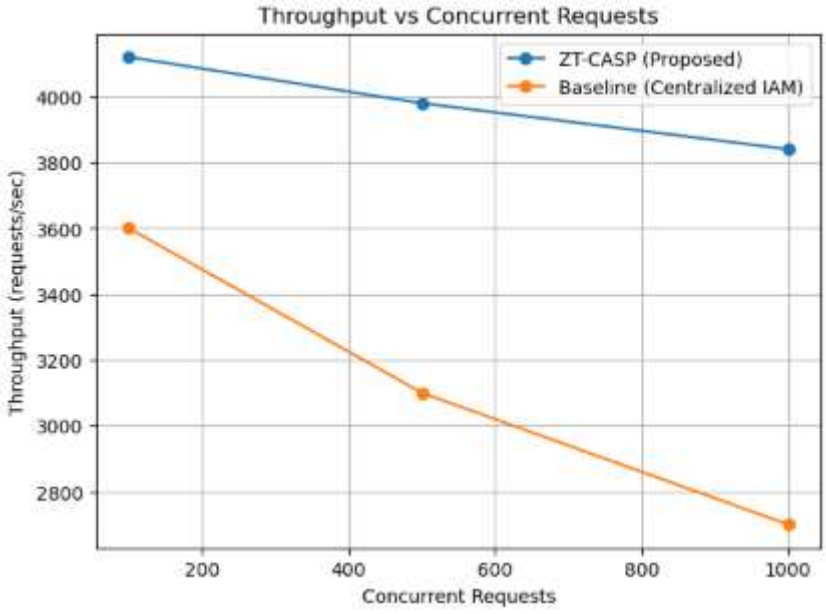

Figure 8 Throughput vs concurrent request

In Figure 8 we measure the system throughput under a growing load of concurrent requests. The proposed approach demonstrates a stable throughput as an increasing number of queries are executed, so scalability is impressive. It is accomplished by sharing enforcement to PEPs, and the reduction of centralized PDP processing enforcing risk-aware decision shortcuts. The current centralised IAM systems have challenges of lower throughput as the system is bottlenecked by policy evaluation and over synchronization.

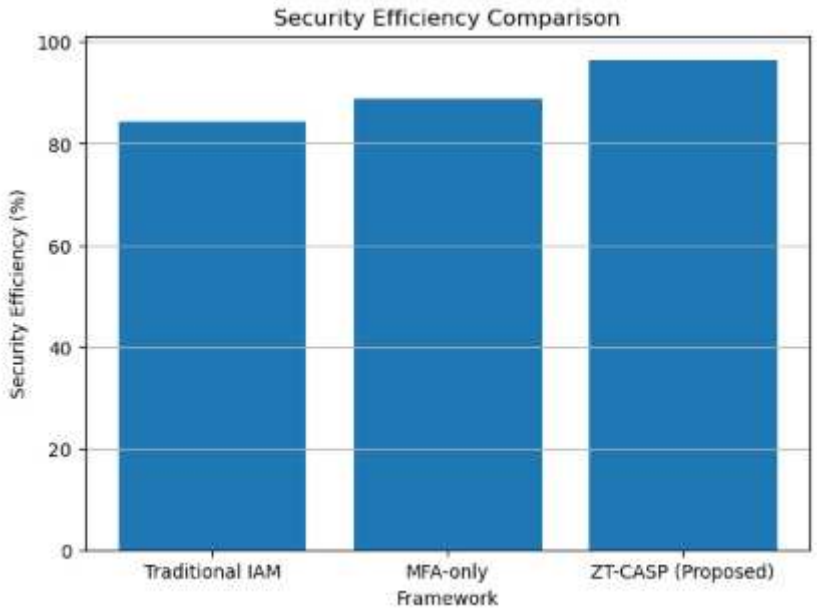

Figure 9 Security Efficiency Comparisons

Last but not least, Fig. 9 contrasts the overall security effectiveness with regard to various security frameworks. AZ-TAC has the best security efficiency as it dynamically scales up protection strength while controlling operational overhead. In contrast, MFA-only or static IAM systems over-protect or under-enforce depending on too broad threat context: the proposed architectural framework adjusts token lifetime, access scope and enforcement intensity according to risk in real-time. The proposed adaptive optimization provides stronger security guarantees with negligible performance overheads, proving that AZ-TAC is a very practical solution to be deployed in real-world clouds.

8. Conclusion

In this paper an Adaptive Zero Trust Access Cloud Security Protocol (AZ-TAC) was proposed to mitigate the lack of security in classic cloud access control mechanisms within dynamic and multi-tenant deployments. The architecture imposes continual authentication and risk-adjusted authorization with least privilege for both human and non-human identities using cryptographic validation and contextual risk determination. By adding proof-of-possession tokens, UEBA-driven risk scoring and adaptive policy enforcement features, AZ-TAC can effectively address the most damaging forms of credential compromise, token misuse and lateral movement attacks. A hybrid dataset and simulation-driven assessment established that the new ABG scheme can deliver significantly reduced authentication delay, revocation response time, improved detection accuracy and more scalability than in current IAM or OAuth2-aligned schemes. The findings also reiterate that adaptive security enforcement does contribute significantly toward protection efficacy with minimum operational overhead. In conclusion, AZ-TAC offers a practical and scalable Zero Trust solution for the modern cloud fabric environment. In the future, we plan to extend our model to consider autonomous policy learning with reinforcement learning for optimizing risk thresholds and access scopes dynamically, post-quantum cryptographic schemes integration, testing on real-world cloud platforms for advancing long-term security and deployment readiness.

Declarations

Data Availability: The datasets used during the current study available from the corresponding author on reasonable request

Funding Statement: Not Applicable.

Ethics Approval: The author provides ethical approval for the given manuscript.

Consent to Publish: All the authors gave permission and consent to publish.

Author Contributions: All the authors have contributed in an equal manner.

Competing Interests: The authors declare that there are no competing interests among the authors.

References:

[1] S. Ali *et al.*, “Computers & Security Security and privacy in multi-cloud and hybrid cloud environments : Challenges , strategies , and future directions,” *Comput. Secur.*,

vol. 157, no. July, p. 104599, 2025, doi: 10.1016/j.cose.2025.104599.

[2] S. Mittal, C. Monga, K. Upreti, N. Kumar, R. D. Raut, and M. S. Alam, "Light Weight Cryptography for Cloud-Based E-Health Records," *2022 7th Int. Conf. Commun. Electron. Syst.*, no. Icces, pp. 690–696, 2022, doi: 10.1109/ICCES54183.2022.9835827.

[3] D. Rayappan, "Lightweight Feistel structure based hybrid-crypto model for multimedia data security over uncertain cloud environment media communication and allied information exchange has," *Wirel. Networks*, vol. 27, no. 2, pp. 981–999, 2021, doi: 10.1007/s11276-020-02486-x.

[4] F. Thabit, O. Can, S. Alhomdy, G. H. Al-gaphari, and S. Jagtap, "International Journal of Intelligent Networks A Novel Effective Lightweight Homomorphic Cryptographic Algorithm for data security in cloud computing," *Int. J. Intell. Networks*, vol. 3, no. April, pp. 16–30, 2022, doi: 10.1016/j.ijin.2022.04.001.

[5] A. L. A. A. H. Ali, E. K. Gbashi, H. Alaskar, A. J. Hussain, and S. Member, "A Lightweight Image Encryption Algorithm Based on Secure Key Generation," *IEEE Access*, vol. 12, no. July, pp. 95871–95883, 2024, doi: 10.1109/ACCESS.2024.3414334.

[6] Q. Huang, Y. Yang, W. Yue, and Y. He, "Secure Data Group Sharing and Conditional Dissemination with Multi-Owner in Cloud Computing," *IEEE Trans. Cloud Comput.*, vol. 9, no. 4, pp. 1607–1618, 2021, doi: 10.1109/TCC.2019.2908163.

[7] S. Mohammed, S. Nanthini, N. B. Krishna, I. V Srinivas, M. Rajagopal, and M. A. Kumar, "Measurement : Sensors A new lightweight data security system for data security in the cloud computing," *Meas. Sensors*, vol. 29, no. March, p. 100856, 2023, doi: 10.1016/j.measen.2023.100856.

[8] J. Zhou, S. Chen, K. R. Choo, and S. Member, "EPNS : Efficient Privacy-Preserving Intelligent Traffic Navigation From Multiparty Delegated Computation in Cloud-Assisted VANETs," *IEEE Trans. Mob. Comput.*, vol. 22, no. 3, pp. 1491–1506, 2023, doi: 10.1109/TMC.2021.3110718.

[9] V. Lamba, J. Goyal, R. N. Koppar, and K. Puttaswamy, "Deep learning - based steganography framework to enhance ad - hoc cloud security," *Int. J. Inf. Technol.*, no.

0123456789, 2025, doi: 10.1007/s41870-025-02648-0.

[10] C. Ghazel, I. Merdassi, and L. Saidane, “An efficient mobile cloud security method based on strong multi - factor authentication and smart card technology,” *Ann. Telecommun.*, no. 0123456789, 2025, doi: 10.1007/s12243-025-01076-2.

[11] E. Science, “A Enhancing Data Security Solutions for Smart Energy Systems in IoT-Enabled Cloud Computing Environments through Lightweight Cryptographic Techniques A Enhancing Data Security Solutions for Smart Energy Systems in IoT-Enabled Cloud Computing Environments through Lightweight Cryptographic Techniques”, doi: 10.1088/1755-1315/1375/1/012003.

[12] A. Sharma and U. Kumar, “Cloud computing security assurance modelling through risk analysis using machine learning,” *Int. J. Syst. Assur. Eng. Manag.*, vol. 16, no. 3, pp. 1287–1300, 2025, doi: 10.1007/s13198-025-02705-8.

[13] S. K. B. Sangeetha, C. Selvarathi, and S. V. Easwaramoorthy, “Secure Healthcare Access Control System (SHACS) for Anomaly Detection and Enhanced Security in Cloud-Based Healthcare Applications,” *IEEE Access*, vol. 12, no. November, pp. 164543–164559, 2024, doi: 10.1109/ACCESS.2024.3492024.

[14] P. Yan, H. U. A. Peng, and F. Qi, “Enhancing medical cloud security : a multi-level approach with twofish encryption and fingerprint authentication,” vol. 0123456789, 2025.

[15] S. Makkar, J. Sidhu, T. Zaidi, R. Batra, P. Garg, and J. Shekhawat, “Advanced model for maximizing multi - cloud security through job scheduling,” *Int. J. Syst. Assur. Eng. Manag.*, no. 0123456789, 2024, doi: 10.1007/s13198-024-02567-6.

[16] N. Indira, S. Rukmanidevi, and A. V Kalpana, “Light Weight Proactive Padding Based Crypto Security System in Distributed Cloud Environment,” vol. 13, no. 1, pp. 36–43, 2020.

[17] J. A. K. Sashi, R. R. Anto, and A. Rosaline, “Enhancing security in cloud computing systems using hybrid feature selection and ensemble - based machine learning for intrusion detection,” *Evol. Syst.*, vol. 16, no. 3, pp. 1–27, 2025, doi: 10.1007/s12530-025-09725-6.

[18] M. Alazab, A. Awajan, A. Obeidat, N. Faruqui, and A. Bere, “Engineering Applications of Artificial Intelligence Adaptive protocols for hypervisor security in cloud infrastructure using federated learning-based anomaly detection,” *Eng. Appl. Artif. Intell.*, vol. 152, no. April, p. 110750, 2025, doi: 10.1016/j.engappai.2025.110750.

[19] L. Ansari, “Enhanced cloud security with Bi-Optimized Sand Cat Swarm and Conv-Bi-ALSTM deep learning models,” *Expert Syst. Appl.*, vol. 286, no. April, p. 128128, 2025, doi: 10.1016/j.eswa.2025.128128.

[20] D. Huo *et al.*, “Double random phase encoding optical cryptosystem for cloud-optimized image security,” *Opt. Lasers Eng.*, vol. 196, no. September 2025, p. 109388, 2026, doi: 10.1016/j.optlaseng.2025.109388.

[21] M. Ali, “A Novel Framework in Cloud Security : Remote Signcryption,” *IEEE Internet Things J.*, vol. 11, no. 24, pp. 41207–41222, 2024, doi: 10.1109/JIOT.2024.3457988.

[22] K. K. Singamaneni, G. Muhammad, S. Member, and Z. Ali, “A Novel Multi-Qubit Quantum Key Distribution Ciphertext-Policy Attribute-Based Encryption Model to Improve Cloud Security for Consumers,” *IEEE Trans. Consum. Electron.*, vol. 70, no. 1, pp. 1092–1101, 2024, doi: 10.1109/TCE.2023.3331306.

[23] M. Y. Shakor, M. I. Khaleel, M. Safran, S. Alfarhood, and M. Zhu, “Dynamic AES Encryption and Blockchain Key Management : A Novel Solution for Cloud Data Security,” vol. 12, no. November 2023, 2024.

[24] K. Sasikumar and S. Nagarajan, “Enhancing Cloud Security : A Multi-Factor Authentication and Adaptive Cryptography,” *IEEE Open J. Comput. Soc.*, vol. 6, no. January, pp. 392–402, 2025, doi: 10.1109/OJCS.2025.3538557.

[25] M. B. Santos, A. C. Gomes, A. N. Pinto, and P. Mateus, “Private Computation of Phylogenetic Trees Based on Quantum Technologies,” *IEEE Access*, vol. 10, pp. 38065–38088, 2022, doi: 10.1109/ACCESS.2022.3158416.

[26] M. Fazil, H. Salwa, S. Hasan, and H. Rauf, “Cloud security modeling : using TCP deltas with data analytics and machine learning techniques,” *Cluster Comput.*, vol. 28, no. 4, pp. 1–19, 2025, doi: 10.1007/s10586-024-04884-8.

[27] V. Mamidala, A. Raj, G. Yallamelli, and M. Veerappermal, “Hybrid anomaly-based cloud security and machine learning-based behavioral analysis for threat intelligence detection in cloud environment,” *Cluster Comput.*, vol. 28, no. 8, pp. 1–20, 2025, doi: 10.1007/s10586-024-05025-x.

[28] U. Nauman, M. Deng, U. Salman, Y. Zhang, and S. Khan, “Q-ECS : quantum-enhanced cloud security with attribute-based cryptography and quantum key distribution,” *Quantum Inf. Process.*, vol. 24, no. 6, pp. 1–32, 2025, doi: 10.1007/s11128-025-04786-7.

[29] Y. Bie and J. Ren, “A model for malicious code propagation in wireless mesh networks based on cloud security defense,” *Telecommun. Syst.*, vol. 88, no. 4, pp. 1–15, 2025, doi: 10.1007/s11235-025-01355-2.

[30] C. P. T. Selvi and R. L. Kumar, “Enhancing data security in the cloud using MECC - SIDH enhanced CL - HPAEKS scheme,” *Wirel. Networks*, vol. 31, no. 3, pp. 2477–2494, 2025, doi: 10.1007/s11276-024-03891-2.