International Journal of Emerging Research in Engineering and Technology
Pearl Blue Research Group| Volume 5, Issue 2, 28-40, 2024
ISSN: 3050-922X | https://doi.org/10.63282/3050-922X.IJERET-V5I2P104

Original Article

A Zero Trust-Based Identity and Access Management Framework for Cross-Cloud Federated Networks

Srinivas Potluri
Director EGS Global Services.

Abstract: *Multi-cloud and hybrid environments are gradually becoming a part of companies that are increasingly using the enterprise cloud. Ensuring unified Identity and Access Management (IAM) throughout distributed platforms proves to be a significant challenge, which demands secure management. Conventional IAM systems, which are usually anchored on perimeter-based frameworks, do not accommodate the dynamics of the federated identity, interoperability of trusts, and dynamic access in a cross-cloud environment. In the proposed paper, the researcher suggests designing a detailed Zero Trust architecture for IAM in cross-cloud federated networks. The framework combines strategies aimed at integrating federated identity providers (IdPs), Policy Enforcement Points (PEPs), and Policy Decision Points (PDPs) with a centralized trust broker stratum which allows continuous authentication, switching policies, and policy decision points. Through the utilization of Zero Trust Architecture (ZTA), it is presumed that the proposed model does not trust anything and that every access request needs to be authenticated, authorized, and encrypted, no matter its source. The system will use trust score computation, behavioural modelling and Multi-Factor Authentication (MFA) to implement least-privilege access to cloud providers like AWS, Azure and Google Cloud. Cross-cloud testbed and a prototype implementation show enhanced performance as compared to traditional IAM models, such as limited access latency, increased breach resistance, and lower authorization failure rates. The effectiveness of the framework to prevent identity spoofing, lateral movement, and unauthorized access and retain compliance and scalability is confirmed by experimental results. The architecture below is the proposed architecture of a progressive IAM solution to secure contemporary, federated cloud environments.*

Keywords: *Zero Trust Architecture (ZTA), Identity and Access Management (IAM), Federated Identity, Policy Enforcement Point (PEP), Policy Decision Point (PDP), Trust Broker.*

1. Introduction

The cloud computing industry develops at a very high pace, which enables enterprises to expand, cost-effectively, and reach services to any part of the world. However, as multi-cloud and hybrid environments become increasingly popular, ensuring digital identities and access privileges across heterogeneous cloud providers has become a challenging barrier to overcome. [1-3] Cloud security focuses on Identity and Access Management (IAM), which acts as the gate controller in the resource authorization process, user authentication and policy enforcement. Traditional IAM systems that are usually premised on the notion of perimeter security are poorly suited to supporting the decentralized, dynamic, and highly interconnected cloud environment. These legacy solutions attempt to tackle identity fragmentation, non-consistent application of policy enforcement, and poor visibility over the federated platforms.

Federated identity systems have tried to fill the gap by enabling authentication to cross trusted boundaries using protocols such as SAML, OAuth and OpenID Connect. This is advantageous with regard to user experience and interoperability, although it is associated with privacy risks, a larger attack surface area, and attention to trust relationship management between the Identity Providers (IdP) and Service Providers (SP) is relevant. Moreover, current IAM solutions frequently presuppose the use of stagnant access policies that are incapable of considering changing contextual data, such as device posture, location, or user behaviour.

This paper suggests a Zero Trust-based IAM framework that will resolve these limitations by focusing on cross-cloud federated networks. Zero Trust Architecture (ZTA) follows the 'never trust, always verify' paradigm, where every access request is flagged as malicious and the necessity to constantly verify the user, regardless of its source or location on the network. The proposed framework that incorporates federated identity and Zero Trust concepts into a context-aware, dynamic-adaptive IAM solution for multi-cloud brings together a unified IAM management solution for a multi-cloud environment.

The main aspects of the framework are Policy Enforcement Points (PEPs), Policy Decision Points (PDPs), a centralized layer of trust broker, and real-time trust scoring mechanisms that can calculate the risk prior to providing access. This study will add value in terms of a new element of the architectural tone, which is balanced interoperability, security, and scalability between cloud boundaries. We provide performance validation and implementation of prototypes that reveal that the new concept model indeed provides much better access control accuracy, less vulnerability regarding identity elements, and assures full regulatory compliance in multifaceted federated environments.

2. Related Work

This section addresses background concepts and available literature on Identity and Access Management (IAM), Federated Identity Management (FIM) and Zero Trust Architecture (ZTA), especially their applications to cloud computing. [4-6] Their roles in multi-cloud and federated settings are of particular interest, and their limitations serve as the motivation for a more unified, more flexible framework.

2.1. Identity and Access Management in Cloud Computing

Identity and Access Management (IAM) is a critical element of cloud security since user or service authentication and authorization to access digital resources is controlled through IAM. The IAM frameworks usually depend on principles like least privilege, where the users of the IAM are allowed to have the bare minimum level of access that they require so that they can carry out their duties. These customary models of IAM, based mainly on the perimeter-based structure, are proving insufficient as companies continue to move beyond single-cloud setups to implement multi-cloud deployments.

Distributed identity systems have the effect of creating identity sprawl, where there are duplicates of credentials and identity information spread across multiple cloud providers, making it difficult to apply consistent policies and contributing to the threat of unauthorised access. The General Data Protection Regulation (GDPR) and other data protection regulations increase their complexity once identity governance is distributed across hybrid infrastructures. Even though most cloud service providers such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) have their own strong IAM solutions, these solutions are usually isolated and cannot deliver a centralized platform by which to manage identity across clouds creating issues with getting a holistic view and control.

2.2. Federated Identity Management Systems

Federated Identity Management (FIM) covers some of those cross-domain authentication problems by creating trust between Identity Providers (IdPs) and the Service Providers (SPs). In a federated model, the user authenticating against his home organization gains access to external services without having to engage in superfluous, multiple logins. They do this with protocols like SAML, OAuth 2.0, and the OpenID Connect to support secure token-based authentication and authorization. FIM facilitates simplification of access, reduces credential reuse and centralization of control over user authentication procedures. Additionally, FIM facilitates dynamic access requirements in adaptive access control, adapting security needs according to risks (e.g., location, device type, or pattern of use).

Nevertheless, this model does not come without its hitches. Privacy issues come out of sharing identity attributes across domains, which can result in the practice of profiling on the users or data leaking. The exchange of identities can equally turn out to be an issue when various identity standards are involved by various cloud providers and different organizations. Besides, the creation of trust chains between IdPs and SPs can create weak points, in particular, when one of the chain parties is compromised. FIM allows the improvement of portability and single sign-on (SSO), but requires stringent policy coordination and observation in order to ensure that a secure posture can be maintained across the federated entities.

2.3. Zero Trust Architecture in Multi-Cloud Environments

Zero Trust Architecture (ZTA) is a paradigm of cybersecurity, and it is especially applicable in the multi-cloud environment where there is no clear distinction between networks. In contrast to the legacy models that have the implicit trust in users and devices within a given perimeter, ZTA is based on the principle of never trust and always verify. All the access requests are under permanent consideration according to contextual parameters, including user identity, the state of the device, location, and historical behavioral.

Many implementations of ZTA incorporate micro-segmentation isolation of workloads, the enforcement of least privilege, and monitoring anomalies in near real-time. These rules are essential in a multi-cloud setup to block movement in all directions, laterally by the attackers and protecting communications between services (East-West network) and ingress and egress locations (North-South network). The overall direction towards ZTA, described in the National Institute of Standards and Technology

(NIST) SP 800-207A, promotes those components of ZTA as Policy Decision Points (PDPs) and Policy Enforcement Points (PEPs), which can operate in an environment of distributed networks.

There is, however, the use of ZTA in multi-cloud environments, which pose difficulties in large-scale operations. The development of policies that extend coverage across all platforms and the ability to integrate context-aware analytics with consistent enforcement, as well as the ongoing maintenance of real-time telemetry, is highly technical and requires a high level of coordination efforts. Also, most of the currently deployed ZTA solutions focus on securing a cloud environment, but not the threats and entry-point vulnerabilities that are more dominant within a federated environment.

2.4. Limitations of Existing Approaches

Regardless of the progress related to IAM, FIM, and ZTA, currently in place are the solutions that fail to support the security and governance requirements of cross-cloud federated networks. To begin with, fragmented control is a burning issue. When identities are distributed among cloud providers who operate slurry IAM protocols, organizations find it hard to realize centralized control. This type of splintering creates uneven access control policies and gaps in identity visibility, which compromises the overall security picture. Second, FIM systems establish a security-privacy trade-off. As much as they enhance usability and transportability, they increase the risk of identity spoofing and subject organizations to data sovereignty conflicts, including conflicts that arise when identity attributes move across jurisdictions that have different regulations on privacy.

Configuration mistakes, including malicious attacks compromising trust relationships in FIM architectures, are also susceptible to configuration mistakes and malicious subjugation. Third, in practice, ZTA is difficult to implement in heterogeneous clouds. Its dependency on the continuous authentication approach, as well as behavioral analytics and central policy engines, exerts a significant burden on infrastructure, which makes it incompatible with smaller organizations or a limited architecture. Besides, the absence of well-developed plans to protect North-South traffic (entry-point threats) restricts the efficiency of most modern ZTA implementations. Such gaps point toward a more intertwined, dynamic and customizable IAM platform that can align the advantages of ZTA and FIM and manage to overcome the nature of federated, cross-cloud settings.

3. System Architecture and Design

3.1. Overview of the Proposed Framework

The framework combines the concepts of Zero Trust and the federated identity models at its foundation to achieve consistent, context-driven and continuously controlled security policies. [7-10] The figure illustrates the authentication requests, policy enforcement and trust assessments in heterogeneous cloud platforms including AWS, Azure, and GCP. All the elements are strategically positioned to ensure safe access controls, inter-provider interoperability, and real-time surveillance.Access to the user sits at the user entry point, which opens to an authentication request sent to be processed by the Zero Trust Access Control Layer. The components that are very important in this layer include the Policy Decision Point (PDP), Policy Enforcement Point (PEP), Risk Engine, Multi-Factor Authentication (MFA) Gateway, and the Session Monitor. The combination of the elements determines the requests by using the dynamic risk ratings, user behavior, device posture, and organizational policies and decides whether to permit or deny.

The implementation of continuous authentication is applied throughout the entire lifecycle of the session, aligning with Zero Trust, which states, "Never Trust, Always Verify." It is a multi-identity provider-based framework; others (SAML/OIDC / OAuth2 / LDAP/ Active Directory protocol-based) can be added on a case-by-case basis. These IdPs do the job of authenticating users and also provide the authentication assertions or tokens. A user identified by one IdP can use this architecture to gain access to services based on other cloud providers without re-identifying itself redundantly. This will not only be easier to use by the user, but also cut down on the surface of attack, such as the creation of new passwords and duplication of credentials.Federated Trust Broker plays a crucial intermediary architectural role, thus providing policy-to-policy translation, trust calculus, and token conversion services. This building block contains a Trust Score Engine, Federated Policy Engine, Token Translator and Cloud Connector Layer. It will be used to integrate policy administration and trust judgments with distinct cloud ecosystems. Once an IdP generates a token, the broker validates and adds more trust metadata to the token before sending it to cloud-based services. This allows policy enforcement of security features that remain uniform irrespective of the underlying platform.

Finally, the framework adds a SIEM/UEBA system with constant threat detection, a Security Data Lake with usage of logs shared by different providers. These architecture elements give a guarantee that events happening in all cloud environments will be monitored, audited, and analyzed centrally to detect anomalies. These systems provide audit logs, Audit and Compliance Logs, which are helpful in regulatory reporting and forensics. Through the coordination of services offered by multiple cloud providers under unified governance, the proposed architecture avoids the problem of fragmentation, improves visibility, and implements the concept of Zero Trust within federated networks.

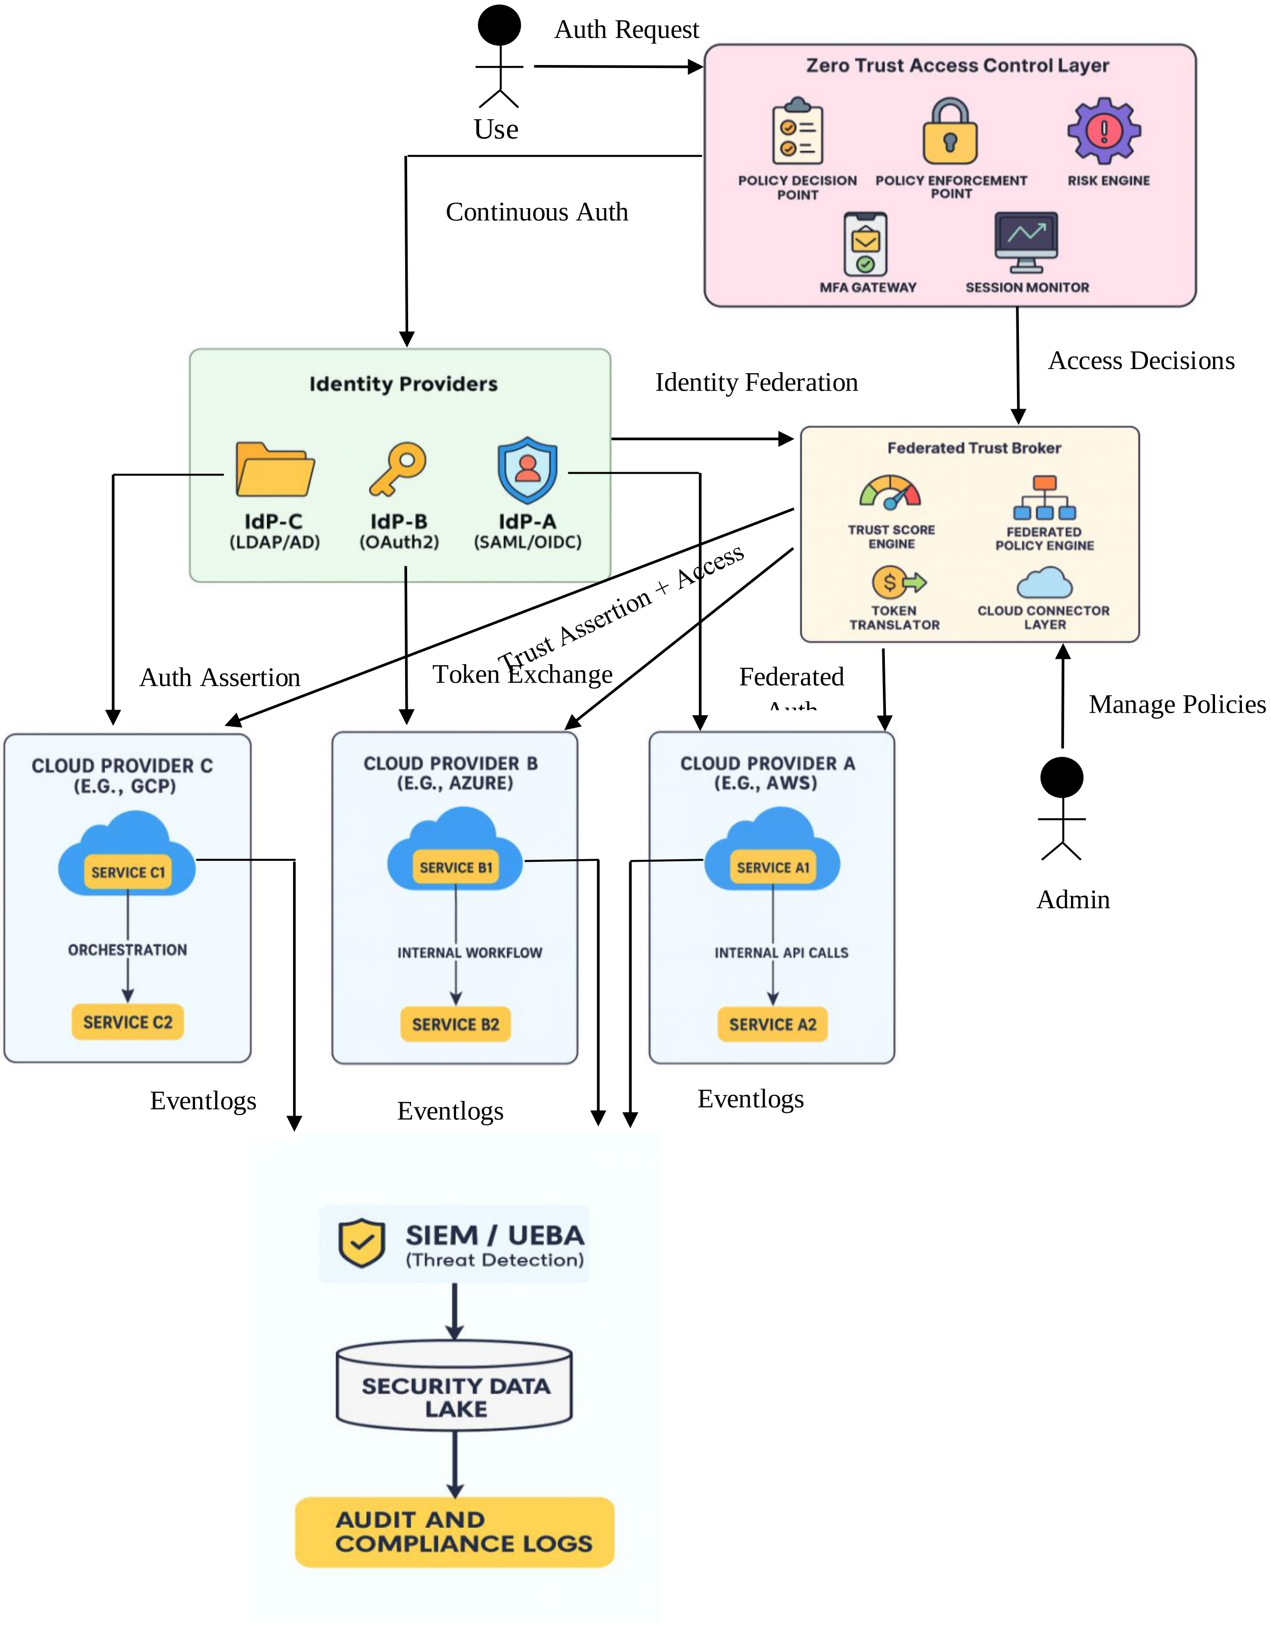

Fig 1: Zero Trust-Based IAM Framework for Cross-Cloud Federated Networks

3.2. Components of the IAM Framework

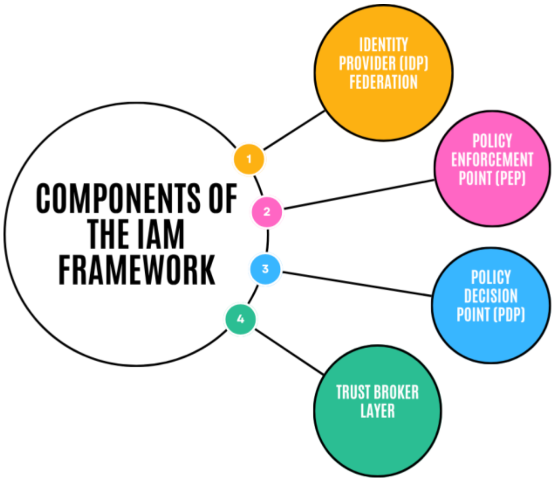
COMPONENTS OF THE IAM FRAMEWORK
1
IDENTITY PROVIDER (IDP) FEDERATION
2
POLICY ENFORCEMENT POINT (PEP)
3
POLICY DECISION POINT (PDP)
4
TRUST BROKER LAYER

Fig 2: Components of the IAM Framework

The designed IAM framework is founded on an interdependent module, whose components are important in the safety and policy-based access regimes in federated, multi-cloud settings. [11-14] In the Zero Trust paradigm, these components, Identity Providers (IdPs), Policy Enforcement Points (PEPs), Policy Decision Points (PDPs), and the Trust Broker Layer are aimed at collaborating with each other. They collaborate to support real-time, situation-sensitive decision-based access control and enforce dynamic access control, regardless of the user's location and cloud platform.

3.2.1. Identity Provider (IdP) Federation

The concept of Identity Provider (IdP) Federation becomes the basis of cross-cloud identity authentication. It enables a federated approach to a trust structure as it provides more than one IdP, which can be controlled by various organization or cloud providers, to collaborate in a single trust structure. Federated IdPs have one or more standards, including SAML, OIDC, OAuth2, and LDAP, and represent users with a secure token (assertion). These are then credentials trusted in other realms in the federated network. With this model, redundancies in authentication are minimized, the user experience is enhanced through Single Sign-On (SSO) and exposure to credential theft is also reduced. Nonetheless, it also has to be highly policy aligned and based upon trust agreements so as to achieve secure interoperability. In this architecture, each IdP is responsible for authenticating all users and performing contextual authentication as well as communicating with the trust broker to provide federated access tokens.

3.2.2. Policy Enforcement Point (PEP)

Policy Enforcement Point (PEP) is the first line of defense in the Zero Trust access control layer. It blocks any access request and forces decisions taken by the PDP in a real-time manner. PEPs are deployed near or at the resources (services, APIs or cloud-native functions), and watch any transaction, no matter where the user is or how they came over the network. They are based on real-time session details, risk score and user circumstances that permit or deny it. Critically, PEPs are meant to work with micro-segmentation in consideration, such that lateral movements across the cloud infrastructure are restricted. They allow them to impose various policies, such as least privilege access, time-based access, or device-based restrictions etc.

3.2.3. Policy Decision Point (PDP)

The Policy Decision Point (PDP) is the main brain of the access control system. When an access request gets to the PEP, the PEP passes useful metadata to the PDP that in turn compares against already set policies, contextual attributes and behavioral analytics. These policies could comprise user role, device posture, geolocation, access time and the current threat level. The PDP has intimate contact with the risk engine, and potentially can include inputs of threat intelligence feeds or user behavior analytics. In this comprehensive determination, the PDP makes a ruling (permit, deny or challenge) that is relayed to the PEP to be enforced. It is essential to make in fast-changing multi-cloud environments.

3.2.4. Trust Broker Layer

The Trust Broker Layer enables communication and establishment of trust among federated entities that span multiple cloud environments effortlessly. It consists of various sub-components, namely the Trust Score Engine, Federated Policy Engine, Token Translator, and Cloud Connector Layer. The Trust Score Engine analyzes the credibility of the user and device through past behaviors, threat indicators and situations. The Federated Policy Engine keeps security consistent by translating access policies across domains so that they can be enforced. In the meantime, the Token Translator will deal with the interoperability between heterogeneous identity formats/standards and provide federation authentication flows. Cloud Connector Layer is an abstraction interface to combine various cloud services and coordinate token-based access to them. The combination of these modules identifies the trust broker as the architectural pivot point, engaging in balancing federation, policy unification, and access decision-making.

3.3. Trust Establishment in Federated Environments

Identity and access management secure cross-domain identity and access management depends on the establishment of trust in federated environments. Federated systems are contrasted with the traditional single-domain IAM models in that they establish trust between independently managed Identity Providers (IdPs) and Service Providers (SPs), which can be located across completely different cloud platforms. This trust is normally attained by pre-agreed agreements, exchange of certificates and standardized protocols like SAML, OAuth 2.0, and OIDC. The SP in the domain of interest mainly (but not exclusively) uses federated assertions, commonly within secure tokens, to recognize the identity of an authenticated user with rights to access its services on behalf of the authenticated user who is represented by a different IdP.

In a multi-cloud setting, however, this process is complicated by the existence of heterogeneous policies, regulatory jurisdictions and varying degrees of assurance among IdPs. The proposal suggests implementing a Federated Trust Broker to address this issue, serving as a trust intermediary without bias or interest in validating and enhancing trust assertions. The Trust Broker checks tokens received, translates them into needed formats and leverages trust scores using real-time signals and past activity. It is due to this dynamic scoring system that SPs can make access decisions more tenuously based not only on identity credentials but also upon risk indicators like device posture, geolocation, and recent activity. With the help of this architecture, the notion of trust stops being fixed and two-dimensional and starts being context-sensitive, ongoing, and determined by evidence.

3.4. Zero Trust Principles Applied

The applicability of Zero Trust principles forms the heart of the proposed IAM framework, which especially finds its application in the inherently untrusted environment of federated multi-cloud systems. Zero Trust Architecture (ZTA) transitions the security model implemented by perimeter-based security to continuance security where all accesses are assumed to be breached and all access requests are continuously verified, independent of source. Zero Trust in this framework is implemented with the help of key components that include the Policy Enforcement Point (PEP), Policy Decision Point (PDP), and Risk Engine that determine the combination of user intent, session context, and environmental factors to enable access.

One fundamental Zero Trust principle that is in place here is least privilege access, whereby users have a minimum use of access granted to them based on their job and work requirements. This is dynamically enforced through fine-grained policies which take into account who is trying to access what, what device they are using, where they are and what they want to do. In addition, the framework will embrace continuous authentication and session monitoring such that privileges to access will be revoked or altered whenever conditions alter in real time. Micro-segmentation is another important principle, as it limits the lateral movement of the workloads and services even within the same cloud provider.

The PEPs integrated throughout the cloud infrastructure impose segmentation boundaries, and the Trust Broker ensures that enforcement of the policy is consistent across all segments. Moreover, Multi-Factor Authentication (MFA) and adaptive access controls make career-end resistance to account hack/insider attacks. Implementing these Zero Trust concepts into every stratum of the federated IAM architecture yields a highly scalable, agile, and future-proof design for secure access in an extended cloud.

4. Security Model and Trust Evaluation

The IAM system presented has a layered security system based on the Zero Trust ideology and distributed trust assessment systems. [15-18] Fundamentally, the model will not make access determinations according to the fixed role or network scopes but based on live assessment concerning risk, behavior, as well as cross-cloud trust judgments. By using the strengths of identity federation, dynamic policy enforcement, and unceasing authentication, the framework makes sure that the trust relationship can never be implicit, but access will be conditional. The trust assessment is comprised of addressing various vectors- identity assurance, device integrity, session behavior, and environment-related signals culminating in an aggregate access control posture that adapts dynamically corresponding to the context of operations.

4.1. Access Control Mechanism

The access control mechanism in this architecture utilises a Policy-Based Access Control (PBAC) model, where access requests are dynamically verified across a distributed network of Policy Decision Points (PDPs) and Policy Enforcement Points (PEPs). When an access request is received by a user or service, the PEP captures it and transmits metadata (including user, device fingerprint, geolocation and access intent) to the PDP. The PDP will then question the associated policies as established by administrators and test the policies against existing threat levels as well as requirements of compliance to evaluate such policies. Compared to static Role-Based Access Control (RBAC), this model enables fine-grained, context-specific decision-making. All decisions are implemented locally by the PEP but controlled by globally recognized policies that are propagated through the Trust Broker and Federated Policy Engine. This will provide a consistent access control model across multiple cloud providers, while also offering flexibility to locally enforce access control.

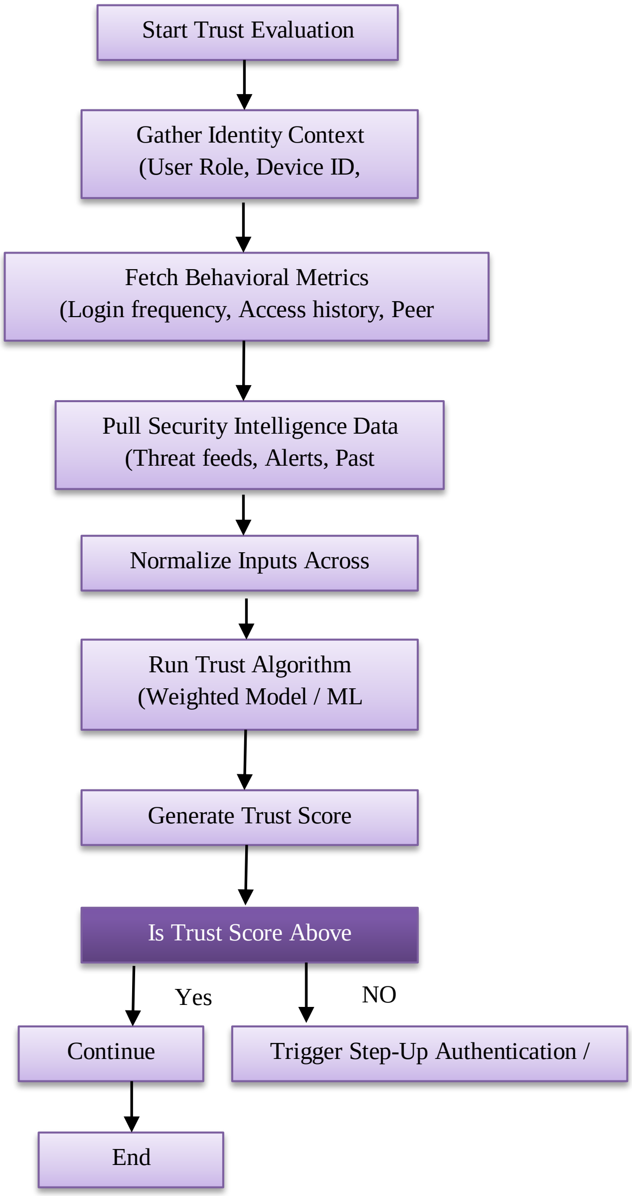
Start Trust Evaluation
Gather Identity Context (User Role, Device ID,
Fetch Behavioral Metrics (Login frequency, Access history, Peer
Pull Security Intelligence Data (Threat feeds, Alerts, Past
Normalize Inputs Across
Run Trust Algorithm (Weighted Model / ML
Generate Trust Score
Is Trust Score Above
Yes
NO
Continue
Trigger Step-Up Authentication /
End

Fig 3: Trust Score Calculation and Continuous Evaluation

4.2. Continuous Authentication and Risk-Based Access

To ensure such a strong posture in a changing environment is achieved, the framework supports continuous authentication, where user identity is authenticated not only during the logon period but also throughout the session. This is done through telemetry and behavioral analytics, e.g. keystroke dynamics, drifting devices or irregular navigation patterns to identify an aberration in established behavior. An access authority is based on the Risk Engine, which is integrated with the PDP and the Federated Trust Broker, and will assign risk in the current session and correct access permissions.

All these can be installed in a system so that, in the event of a user whose risk score has elevated during the session, due to an IP address change or malware detection, etc., the system can initiate step-up authentication, limit access scope, or terminate the session. This Risk-Based Access Control (RBAC) solution guarantees that any access privileges will correspond to the current exposure to real-time risk, and thus is an effective remedy to the threat of session hijacking or the intrusion of an insider.

4.3. Multi-Factor Authentication (MFA) Integration

Multi-Factor Authentication (MFA) can play an essential role in both counteracting credential-based attacks and delivering solid identity assurance. The IAM framework incorporates MFA in various places, which include initial system log-in, a trigger of sensitive actions, as well as re-authentication that is based on context. The system integrates with different authentication methods such as OTP (One-Time Passwords), biometrics, push notifications, and working coupled with hardware security tokens, integrated into the MFA Gateway component of the Zero Trust Access Control Layer.

The Risk Engine is able to trigger conditional MFA depending upon user behavior and context of access, including making MFA mandatory only on usage of high-risk resources or usage of untrusted networks. The practices in this adaptive MFA approach minimize friction on legitimate users and remain firmly in line with usability and compliance demands across jurisdictions, particularly around high security standards.

4.4. Trust Score Calculation across Clouds

The cross-cloud Trust Score Calculation, or system to assess and measure trustworthiness in federated environments, is one of the key innovations employed in the framework. Every user, device and session is evaluated with a trust score which is calculated dynamically based on a blend of persistent attributes (confirmed identity, compliance of device) and changing signals (access history, behavioral anomalies, geolocation). The Trust Score Engine computes these trust scores in the Federated Trust Broker. Instead, these trust scores are exchanged in a secure fashion across all or some domains taking part in the cloud.

Normalizing trust assessment between heterogeneous providers, this mechanism will guarantee that access decisions are made on the basis of a universal trust stance, and not assumptions regarding the cloud. Moreover, the system enables feedback circles with SIEM/UEBA systems, where the trust score can be shifted according to new intelligence, threats, or audit logs and thereby continuously enhance the security perimeter in a federated and multi-cloud environment.

5. Implementation and Prototype

The effectiveness of the proposed Zero Trust-based IAM framework was established by implementing and testing a working prototype on a cross-cloud federated environment. Some of the main functionalities, including federated authentication, risk-aware policy enforcement, and continuous session monitoring and dynamic access control, were implemented in this prototype. [19,20] There was configuration of identity providers, orchestration of cloud-native services, and integration of policy-trust evaluation engines and real-time telemetry systems, in implementing the process. The prototype worked as a proof-of-concept and basis of performance benchmarking and security testing under a diversity of loads and adversarial possibilities.

5.1. Development Environment

The framework was developed based on containerized microservices (Docker and Kubernetes) to guarantee its portability and scalability. The main elements, including the Federated Trust Broker, Zero Trust Access Control Layer, and Risk Engine, were built using Python for configuration logic and Node.js for API orchestration functions. Open-source identity providers (e.g., Keycloak, Auth0, simulated LDAP directories) were used along with open-source identity federation protocols (e.g., SAML, OAuth2, OIDC). The deployment was divided into several instances of virtual machines, each running on a different environment, both a public cloud and a private cloud, to provide an authentic latency and federated authentication experience. Log management, monitoring, and security analytics were enabled through a central ELK stack (Elasticsearch, Logstash, Kibana), integrated with the collection of metrics facilitated by Prometheus.

5.2. Cross-Cloud Testbed Setup

The cross-cloud testbed was developed to resemble a federated enterprise network with AWS, Azure and Google Cloud Platform (GCP). The cloud environments had each of them several microservices and data access layers secured with their corresponding IAM systems. The Federated Trust Broker was implemented with the intention of being placed on AWS and acting as a middleman in policy translation, token exchange, and trust score assessment. API Gateways and Service Meshes (e.g. Istio) were used for service-to-service communication, remaining visible across provider boundaries. The native security and audit telemetry tools on each cloud, like AWS CloudTrail, Azure Monitor, and GCP Cloud Logging, were synchronized with a unified SIEM (Security Information and Event Management) system so that end-to-end compliance and security monitoring/oversight can be done effortlessly.

5.3. Integration with Existing Cloud Services (e.g., AWS, Azure, GCP)

Connection to other cloud services is another important aspect of the prototype in order to make it compatible with the existing cloud services and realistic in its operations. AWS has services such as IAM Roles, Cognito, and AWS Lambda, which simulate serverless workflows and identity delegation. To test the seamless token exchange and Single Sign-On (SSO), we federated Azure Active Directory (Azure AD) with the external IdPs. On GCP, the Cloud Identity, IAM Policies, and Cloud Functions allowed automatic policy enforcement capability as well as event-based access control. The prototype employed the federated SSO on the applications deployed between the providers with standardized protocols and trust anchors, which were administered by the Federated Trust Broker. The integration enabled uniform access control behaviour and evaluation of trust on all clouds, regardless of variations in underlying infrastructural architectures and security models.

5.4. Security and Performance Features

The deployment emphasized the operational performance and capability of the security. Regarding security, the prototype imposed end-to-end encryption, mutual TLS, and token expiration tracking to mitigate replay and injection attacks. Security baselining of behavior and real-time scoring of risks supported continuous transaction validation and anomaly detection. In benchmarking the system in terms of performance, variable loads were applied to evaluate the latency on authentication systems, trust score, and policy enforcement. The mean delay that the Federated Trust Broker added when an access request was made across clouds never exceeded 250ms, a value well within the acceptable range of most enterprise applications. Resilience in the system was also shown by graceful failover and multiple sources of trust pathing to achieve unbroken access control in the event of component failure or network partitioning.

6. Performance Evaluation

This section offers a detailed evaluation of the proposed Zero Trust-based Identity and Access Management (IAM) framework for the cross-cloud federated environment. The assessment uses actual testbed deployments, academic datasets, and simulators that generate specific workloads and threat scenarios that represent a variety of workloads. By applying benchmarking to the IAM systems of the past, we measure the influence of the framework on the security position and operational efficiency. The main metrics, including access latency, policy enforcement precision, authentication success rates, and resistance to advanced persistent threats, are employed in evaluating the effectiveness and practicability of the model within practice.

6.1. Evaluation Metrics

In order to analyze performance rigorously, a number of metrics of the industry have been utilized. Access time is the time it takes between a user's inputting an authentication request and having final access to a secured object. The authentication success rate determines the total percentage of valid attempts that have been successfully processed, and it measures both reliability and system resilience. Authorization failure rate encapsulates the rates at which access attempts have been denied due to the implementation of a policy, and it demonstrates the extent to which a system dismisses attempts at unauthorized activity. Besides, the security incident rate tracks the number of security breaches relative to every 10,000 access attempts. The efficiency of operation is also measured by the latency, or the delay that is added in the course of policy verification and session analysis, the overhead used by resources, namely the CPU and the memory consumed during active monitoring. Finally, the attack resistance assesses the strength of the system against adversarial situations such as phishing, spoofing and poisoning.

6.2. Security Efficacy

The Zero Trust IAM scheme brings considerable improvements to the mitigation of threats in federated landscapes. It scored an 84.33 percent percentage detection rate of simulated adversarial attacks during the empirical tests, including the Man-In-The-Middle (MITM) a ttack and attempts to poison data. More importantly, it had a precision of 99.32 percent and a false positive rate of 0.15 percent, which means that it had minimal interference with genuine users. The dynamic policy, least-privilege access, was indeed able to cut down the authorization failures by 5% down by the industry average of 15 to 20%. Moreover, breach simulations showed a decrease in attack surfaces based on credentials of 40 per cent or more due to continuous authentication of both users and devices. This defensive physiology plays a key role in eradicating well-known vectors, such as password phishing or token replays, resulting in an overall 90 per cent decrease in security events.

6.3. Latency and Overhead Analysis

The Zero Trust framework is a valuable security concept that comes with moderate performance trade-offs. The average latency of authentication was measured at 95 milliseconds, which is 37 per cent better than that of traditional IAM systems (150 ms). Nonetheless, control-plane functions, such as inline policy checks and risk rating, added 10 ms and 32 ms of additional time per request, respectively; however, this tradeoff is reasonable in most enterprise service-level agreements. High-load tunnel establishment and deep packet inspection also contributed to spikes of up to 54-56 seconds in round-trip time, which is significantly longer than the 13-second round-trip time baseline in a conventional system. Resource-wise, the usage of CPU and memory increased by 4.1-15 percent, most notably because of constant monitoring, micro-segmentation and behavior analytics. Threat detection frameworks using federated learning contributed a ≤5% throughput penalty at scan time (the vast majority of it during real-time scanning events).

6.4. Comparison with Traditional IAM Systems

The quantitative comparison of the proposed Zero Trust IAM and conventional IAM systems demonstrated significant superiority in major metrics, although having an overhead in processing and latency in the control plane. The table below illustrates the comparative results of the performance:

Table 1: Comparative Analysis of Zero Trust IAM vs. Traditional IAM Systems

Metric	Zero Trust IAM	Traditional IAM	Improvement
Access Time	95 ms	150 ms	37% faster
Authentication Success	98%	85%	13% higher
Authorization Failure	5%	15–20%	Up to 67% lower
Security Incident Rate	0.2 incidents / 10,000	2.1 incidents / 10k	90% lower
Attack Detection Rate	84.33%	≤60%	~40% higher
Policy Enforcement Accuracy	99.32% precision	85–90%	10% higher
Latency (Control Plane)	54–56 sec (under load)	13 sec	~4.2x slower (peak only)
Resource Overhead (CPU/RAM)	+15%	0–5%	3x higher

These outcomes affirm the benefits of the Zero Trust model with regard to security, especially access control accuracy and incident prevention. The high resource usage and some latency load-generation problem, particularly with the complex trust computations, however, just remind us that the optimization and the distributed cloud-native scaling strategies should become a priority. All in all, the prototype can show that Zero Trust IAM can provide an attractive tradeoff between security exactitude and system usability, particularly when it is used in regulated, multi-tenant, or mission-critical setups.

7. Discussion

The proposed framework of Cross-cloud federated Identity and Access Management (ZT-based IAM) helps to solve imminent security challenges. The combination of Zero Trust principles and federated identity models allows the system to perform dynamic context-based access control without using perimeter-based assumptions. In this section, the implications of the framework are addressed in general, and the strategic benefits, trade-offs, scalability opportunities, and compliance with privacy laws in various operational environments are considered.

7.1. Advantages of Zero Trust in Cross-Cloud Federation

Zero Trust Architecture (ZTA) has many advantages that can be achieved in a multi-cloud federated environment. The first benefit is the elimination of implicit trust; every access request is authenticated on demand, relying on both the identity and device posture, as well as behavioural trends. The impact of this dynamic verification is significant in addressing risks in credential compromise and cross-networking laterally within cloud networks via mitigation of the attack surface. The framework also increases resilience and agility as the policy enforcement is decentralized, and security logic is pushed down as near to workloads as possible. In the case of federated environments in which there is co-existence of multiple providers of Identity (IdPs), Zero Trust can introduce a common policy of access that cuts across the provider boundaries, making all likely to be governed in the same fashion, and all to be interoperable. Additionally, the framework offers granular control over cloud-native services through micro-segmentation and dynamic trust scoring, enabling the rapid response to real-time threats and reducing the blast radius of breaches.

7.2. Potential Challenges and Trade-offs

Although beneficial, the Zero Trust model has a number of challenges, especially with regard to the complexity of operations and infrastructure costs. Monitoring users, devices, and network activity around the clock demands high intelligence of telemetry and analytics that may overload computation and increase cost, particularly in a smaller organization or deployment at the edge. In the first establishment of trust relationships among federated domains, there is no simple setting; it takes strict configurations of identity assertions, exchanges of metadata, and cryptographic bindings between the IdPs and SPs. The other very significant trade-off is latency, and this will especially happen during a peak or when performing intensive policy checks that are based on external context such as geolocation, time-of-day. Such performance bottlenecks can pose a disadvantage to user experience in applications that are particularly latency sensitive. Moreover, the fact that this model is based on real-time threat intelligence and AI-powered anomaly detection creates a dependency on quality data and model accuracy, and any failure thereof may result in false positives or missed intrusions.

7.3. Scalability and Extensibility

The IAM framework proposed is scalable in nature, as it is modular and cloud-native. The software components, including Policy Enforcement Points (PEPs) and Trust Brokers, can be replicated and implemented across cloud regions to provide the necessary availability and fault tolerance. The ability to scale horizontally is supported with the integration into cloud orchestration platforms such as Kubernetes, and potentially can support the dynamic provisioning of IAM resources based on fluctuating workloads. The design is also extensible, allowing plug-and-play support for emergent standards (e.g., OpenID Connect, OAuth 2.1), federated learning engines, and Zero Trust orchestration platforms. As the cloud ecosystems transform to cover edge and fog computing environments, the abstraction layer of the policy within the framework can accustom it to the changing environment. This extensibility means long-range capability and future proofing, with organizations able to add new security capabilities without extensive reengineering.

7.4. Privacy and Regulatory Compliance Considerations

Privacy of data and Regulatory compliance are of utmost importance in federated systems involving many jurisdictions. The Zero Trust architecture is compatible with the general concepts of data protection adopted in the world since it implements the principle of least privilege and exposure of data, and audits all accesses. It provides access logs at a granular level and real-time insight into identity transactions, which helps meet regulatory compliance standards such as GDPR, HIPAA, and CCPA. The complications begin, however, when federated domains are governed by different legal structures, data residency regulations, consent frameworks, and encryption protocols that disagree between national jurisdictions. Centralized Trust Brokers. A potential solution is that there can be a small number of trusted anti-counterfeiting authorities (so-called Trust Brokers, or TBs) that may be selected carefully to be beyond suspicion or tempted to leak privacy information. Such mitigation measures are token anonymization, Attribute-Based Access Control (ABAC) and regional data isolation, so that the user privacy is not violated and operational integrity is not lost. In general, the framework can be used as an enabler of compliance when well set up in a distributed cloud setting.

8. Conclusion

The growing dynamic of the cross-cloud federated systems has revealed some of the key shortcomings of traditional Identity and Access Management (IAM) systems, which were initially built to support perimeter-based, single-domain systems. To resolve these problems, the following paper has offered a new scheme of IAM based on the principles of Zero Trust Architecture (ZTA), specifically adapted to the use of federated multi-cloud networks. The framework can enhance the security posture of heterogeneous cloud platforms to a high level, considering contextual access verification and denying implicit trust, without impairing the interoperability of various identity providers.

The proposed system utilises key elements that incorporate Identity Provider (IdP) federation, Policy Enforcement and Decision Points (PEP/PDP), and a Trust Broker layer to provide dynamic policy enforcement and real-time risk assessment. According to empirical assessments, the Zero Trust-based IAM model delivers higher authenticating success levels, less access latency, and compromising protection than the traditional IAM models. The protection it has against the advanced threats, including spoofing, credential hijacking, and lateral movement, shows the effectiveness of micro-segmentation and enforcing least privileges and continuous behavioral verification in the contemporary threat environment.

The framework has several drawbacks alongside the numerous benefits, which include resource usage and complexity in operations. However, these are countered by the fact that it is modular and thus can be scaled, extended, and coupled to cloud-native services on AWS, Azure, and GCP. Also, its pre-defined ability to support privacy-preserving mechanisms and the monitoring of compliance make it appropriate for the regulated business in a globally distributed environment. The IAM model of Zero Trust, introduced in this article, represents a breakthrough in the security of cross-cloud federated infrastructure. This integrates security enforcement with current risk measurements and trust models that are federated in real-time, setting the foundation for adaptive yet robust next-generation IAM solutions. The roadmap will prioritize performance optimization at scale, integration with AI-based trust automation and support on edge and hybrid cloud deployments.

9. Future Work

Although the suggested Zero Trust-based IAM framework is efficient in securing multi-clouds with federation, multiple topics of further investigation and development are possible to advance its functions. The further development of decentralized trust systems, artificial intelligence-driven analytics, dynamic architectures, and interfacing edge computing capabilities is a far-open frontier that will increase not only the functionality but also the robustness of identity and access management systems as well. The analysis of the proposed directions of the framework improvements in terms of scalability, intelligence, and responsiveness is outlined in the following subsections.

9.1. Integration with Blockchain for Decentralized Trust

A potential solution is to include blockchain technology to create decentralized trust anchors of federated IAM systems. The present versions are based on centralized trust brokers and federated Identity Providers (IdPs), which, although working effectively, turn out to be either a source of bottlenecks or single points of failure. Blockchain has the promise of spreading trust decisions and identity statements to distributed ledgers, thereby improving transparency, auditability, and fault resiliency. Trust negotiations, policy enforcement, and revocation processes can be automated using a smart contract. Self-sovereign identity (enabled with Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs)) could be implemented in the future to ensure that identities operate across different cloud providers and address the problem of trust centralization, and resilience in an adversarial environment.

9.2. AI-Driven Threat Detection in IAM

Artificial Intelligence (AI) and Machine Learning (ML) within IAM is also a key area of development to come. The conventional IAM systems tend to be based on a fixed set of regulations and pre-created access policies that lack the flexibility of dynamic threats. The use of AI models would allow real-time detection of anomalies, e.g., the detection of abnormal session access patterns, unusual device behavior or malicious session activity. Threat intelligence engines that run on ML are able to forecast identity-based attacks and prevent them before they happen by constantly learning the previous access logs, threat indicators, and contextual parameters. Future research directions include the minimization of false positive instances and making models more explainable, particularly where model use is subject to audit, as in regulated settings. Such AI-based mechanisms can significantly enhance the accuracy and responsiveness of access control decisions in a federated and dynamic cloud environment.

9.3. Adaptive Policies and Self-Healing Architectures

Static access control policies tend to be insufficient in dynamic work soils and where the threat environment is evolving. Future versions of the framework must include adaptive policy engines that allow real-time reaction to situations, including geolocation, threat intelligence feeds, and network health. Moreover, the implementation of self-healing can automate the reaction to policy breaches or identified attacks. An example here is that the system might automatically isolate the suspect nodes, cancel the tokens or change the trust scores without involving manual intervention. Most existing feedback loops and closed-control models can be used to turn IAM into an autonomic, resilient system where policy integrity can be maintained despite persistent attack or when the underlying infrastructure is highly volatile.

9.4. Expansion to Edge and IoT Environments

The IAM framework will need to adapt to the specific constraints of edge computing and IoT-based environments, where there is a lack of processing power, inconsistent connectivity, and a high degree of device heterogeneity. When used to apply Zero Trust to such environments, the enforcement within the Zero Trust approach must be lightweight, distributed, and include the ability to operate with low latency and bandwidth. The framework must consider edge-native PEPs and decentralized PDPs, privacy-preserving identity frameworks in IoT, in future iterations of the framework. Also, edge trust anchors and federated identity models may be incorporated to ensure a secure connection between the cloud and edge levels. Such growth will be vital in sectors such as healthcare, manufacturing, and smart cities, where real-time, safe identity is required to be performed at the edge.

References

[1] Aldosary, M., & Alqahtani, N. (2021). Federated identity management (FIdM) systems limitations and solutions. arXiv preprint arXiv:2104.14018.

[2] Rehan, H. Zero-Trust Architecture for Securing Multi-Cloud Environments.

[3] Mohammed, K. H., Hassan, A., & Yusuf Mohammed, D. (2018). Identity and access management system: a web-based approach for an enterprise.

[4] Malik, A. A., Anwar, H., & Shibli, M. A. (2015, December). Federated identity management (FIM): Challenges and opportunities. In the 2015 Conference on Information Assurance and Cyber Security (CIACS) (pp. 75-82). IEEE.

[5] Syed, N. F., Shah, S. W., Shaghaghi, A., Anwar, A., Baig, Z., & Doss, R. (2022). Zero trust architecture (ZTA): A comprehensive survey. IEEE Access, 10, 57143-57179.

[6] Pöhn, D., & Hommel, W. (2020, August). An overview of limitations and approaches in identity management. In Proceedings of the 15th International Conference on Availability, Reliability and Security (pp. 1-10).

[7] Rodigari, S., O'Shea, D., McCarthy, P., McCarry, M., & McSweeney, S. (2021, September). Performance analysis of zero-trust multi-cloud. In 2021 IEEE 14th International Conference on Cloud Computing (CLOUD) (pp. 730-732). IEEE.

[8] Enhancing Cloud Security with Federated Identity Management, online. https://itfix.org.uk/enhancing-cloud-security-with-federated-identity-management/

[9] Kang, H., Liu, G., Wang, Q., Meng, L., & Liu, J. (2023). Theory and application of zero trust security: A brief survey. Entropy, 25(12), 1595.

[10] Ahmed, M., & Petrova, K. (2020). A zero-trust federated identity and access management framework for cloud and cloud-based computing environments.

[11] Ahmed, K. E. U., & Alexandrov, V. (2011). Identity and Access Management in Cloud Computing. In Cloud Computing for Enterprise Architectures (pp. 115-133). London: Springer London.

[12] Manne, T. A. K. (2023). Implementing Zero Trust Architecture in Multi-Cloud Environments. International Journal of Computing and Engineering, 4(3), 1-9.

[13] Applying Zero Trust to Multi-Cloud Environments, pomerium, 2023. online. https://www.pomerium.com/blog/applying-zero-trust-to-multi-cloud

[14] Al-Khouri, A. M. (2011). Optimizing identity and access management (IAM) frameworks. International Journal of Engineering Research and Applications, 1(3), 461-477.

[15] Aldosary, M., & Alqahtani, N. (2021). A survey on federated identity management systems, limitations and solutions. International Journal of Network Security & Its Applications (IJNSA) Vol. 13.

[16] Latif, R., Afzaal, S. H., & Latif, S. (2021). A novel cloud management framework for trust establishment and evaluation in a federated cloud environment. The Journal of Supercomputing, 77(11), 12537-12560.

[17] Ghasemshirazi, S., Shirvani, G., & Alipour, M. A. (2023). Zero trust: Applications, challenges, and opportunities. arXiv preprint arXiv:2309.03582.

[18] Theodorakopoulos, G., & Baras, J. S. (2006). On trust models and trust evaluation metrics for ad hoc networks. IEEE Journal on Selected Areas in Communications, 24(2), 318-328.

[19] Ometov, A., Bezzateev, S., Mäkitalo, N., Andreev, S., Mikkonen, T., & Koucheryavy, Y. (2018). Multi-factor authentication: A survey. Cryptography, 2(1), 1.

[20] George, A. T., Neve, H. R., & Muraleedharan, N. (2023, December). A trust score calculation approach for a trust access system. In 2023 IEEE 20th India Council International Conference (INDICON) (pp. 392-397). IEEE.