International Journal of Emerging Research in Engineering and Technology
Pearl Blue Research Group| Volume 5, Issue 2, 95-107, 2024
ISSN: 3050-922X | https://doi.org/10.63282/3050-922X.IJERET-V5I2P110

Original Article

A Federated Zero-Trust Security Framework for Multi-Cloud Environments Using Predictive Analytics and AI-Driven Access Control Models

Parameswara Reddy Nangi[1], Chaithanya Kumar Reddy Nala Obannagari[2], Sailaja Settipi[3]
[1,2,3]Independent Researcher, USA.

Abstract *- Companies are moving to multi-cloud more and more to achieve cost-weight, resilience or flexibility of vendors, which leads to identity, policy, and monitoring fragmentation. Such environments cannot support identity sprawl, lateral movement and rapid changing threats with traditional perimeter-based and static controls, which are role-centric. The paper suggests a Federated Zero-Trust Security Framework which aligns security posture between heterogeneous clouds by decentralizing control, using federated identity and ensuring continuous verification. It is a federated identity and credential management layer, a multi-cloud policy orchestration plane, and an AI-driven access control engine which ingests secure telemetry of all providers. Dynamic risk scores are built using predictive analytics and behavior-based models and used to perform threat prediction, and to run adaptive access decisions, such as step-up authentication, just-in-time privilege elevation, or automatic session termination. An experimental implementation on AWS, Azure and GCP highlights small latency overhead, high throughput and better policy consistency whereas experimental analysis illustrates significant improvements on detection accuracy, false positives as well as cross-cloud attack surface reduction compared to RBAC and vendor-native IAM baselines. The findings show that a federated zero-trust predictive, AI-driven access control is a scalable and proactive model of defense in multi-cloud ecosystems of complexities.*

Keywords *- Multi-Cloud Security, Federated Identity Management, Zero-Trust Architecture, AI-Driven Access Control, Predictive Analytics, Policy Orchestration, Continuous Verification, Cloud-Native IAM, Threat Prediction.*

1. Introduction

Multi-cloud strategies have been adopted at a rapid pace reshaping the way enterprises architect and provide digital services and are able to capitalize on the best-of-breed capabilities offered by other cloud providers and achieve optimal cost, performance and resilience. [1-3] but security controls, identity stores and policy management are also disaggregated in this distributed model over heterogeneous platforms. As a result, organizations face an expanded attack surface, inconsistent access policies, increased risk of misconfigurations, and complex compliance obligations. Traditional perimeter-based security architectures, which assume trusted internal networks and static boundaries, are no longer effective in an environment characterized by remote work, API-driven integrations, shadow IT, and sophisticated cyber threats. The threat of identity compromise, lateral movement, and privilege escalation across cloud boundaries continue to pose a threat to attackers and reveal the flaws of traditional access control and monitoring strategies.

The concept of Zero-Trust Security (ZTS) has become a paradigm of solving such issues by presuming that no user, device, and workload are a priori trusted, irrespective of location. However, most zero-trust configurations are centralized, fixed or siloed in a single provider and are more effective in an actual multi-cloud setup. To mitigate this gap, an increasing level of demand exists to have a Federated Zero-Trust Security Framework that can consolidate identity, policy and telemetry across clouds whilst still respecting organizational boundaries and regulatory limitations. Such a framework can become adaptive, risk-aware with the addition of AI-informed access control models and predictive analytics, rather than being reactive as it can enforce rules. In this paper, the authors present federated and AI-enriched zero-trust architecture of the multi-cloud ecosystem that centers on continuous verification, federated identity, and predictive threat prevention as the main design principles.

2. Related Work

2.1. Zero-Trust Security Principles

Zero Trust Architecture (ZTA) is most strictly formalized through the structure of the NIST SP 800-207 where it is described as a security paradigm in which no implicit trust is assigned on the basis of network location or traditional perimeter barriers. [4-6] Rather, access requests should be authenticated authorized and again reviewed using identity, device posture, and context

indications (location, time, workload sensitivity, etc.) on an ongoing basis. The network is regarded as something that cannot be trusted in this perspective, and protection is no longer at the perimeter, but at the resources and data objects themselves.

Subsequent surveys and architectural analysis decomposes ZTA into several essential units, namely high identity and credential management, fine-grained access control (which is often role-, attribute-, or risk-based), continuous verification based on telemetry, and micro-segmentation in order to control lateral movement. They also note that centralized or logically centralized points of policy decision that take in logs, metrics, and security analytics to decide to either permit, reject or step-up authenticate a particular request are essential. Although this conceptual maturity, empirical studies continue to report challenges in deployment in actual hybrid and cloud environments, such as policy inconsistencies, challenges in integrating with legacy systems and limited end-to-end visibility across the heterogeneous infrastructure components challenges that are increased in multi-cloud environments.

2.2. Multi-Cloud Identity and Access Management

The literature on multi-cloud security demonstrates that organizations are moving faster and faster to deploy applications and data on multiple infrastructure and platform vendors, often because of cost-saving, resiliency, or regulatory criteria. The result of this trend is the creation of fragmented silos of identity where individual providers have their own identity and access management (IAM) constructions, including users, roles, policies and trust relations. Research and industry reviews present that this disintegration complicates the imposition of equal access provisions, the least-privilege designing is more hazardous, and the chance of misconfiguration and privilege creep across the accounts and providers is more likely to occur.

Several works suggest unified or centralized identity layers, which are over provider-native IAM, with directory synchronization, common protocols, or abstraction layers to map identities into other clouds. Although these methods can enhance consistency in policy and end user experience, they are quite frequently meticulously orchestrated manually by identity architects in order to map roles, attributes and permissions across various platforms. Case studies have indicated that this manual process cannot increase in scale with the size of the cloud infrastructure number of clouds, tenants, and microservices, and many existing solutions continue to use provider-specific policy artifacts, which restrict their applicability to dynamic and risk-adaptive zero-trust implementation in a multi-cloud setting.

2.3. Federated Identity and Authorization Models

Federated identity management (FIM) has long been proposed as a mechanism for enabling single sign-on and cross-domain access by establishing trust relationships between identity providers (IdPs) and service providers. Classical FIM designs are based on SAML, OAuth 2.0 and OpenID connect standards to provide a layer of identity assertions in tokens that are verifiable across domain or organizational boundaries. This improves usability and reduces the need for duplicate accounts, but introduces dependencies on central IdPs or federation hubs and requires careful management of trust relationships, token lifetimes, and attribute release policies.

More recent work examines federated identity in microservices-heavy and multi-cloud environments, exploring patterns such as centralized IdPs, hub-and-spoke federation brokers, and identity propagation embedded into service meshes or API gateways. Such designs can simplify administration and enable a smooth transition between heterogeneous systems, but bring issues of token translation to heterogeneous systems, privacy and minimal disclosure of attributes, recovery of resilience to IdP failures, and interoperability between multiple providers. FIM surveys have still identified the complexity in integrating, legal and regulatory issues in cross border sharing of identity attributes and the necessity of clear governance structure to establish federation policies, responsibilities, and dispute resolution structures. Notably, the majority of FIM work emphasizes authentication and coarse-grained authorization, as opposed to fine-grained and dynamically risk-adaptive access control combined with a zero-trust and AI-driven decision logic.

2.4. Predictive Analytics and AI in Access Control

Literature in security concerns is more actively exploring the effectiveness of predictive analytics in security, including machine learning (ML) and user and entity behavior analytics (UEBA) in improving authorization decisions. The common methods will use the past records and telemetry of logins, resources accessed, device attributes, and network paths and will compute real-time scores of risk or forecast anomalies when new access requests do not follow anticipated patterns. These mechanisms are frequently combined with attribute based or risk adaptive access control models where context-dependent decisions, including step up authentications, temporary elevation of privileges or automatic session revocation can be made.

Within zero-trust and cloud security, AI and ML techniques have been applied to threat detection, insider risk analysis, micro-segmentation tuning, and prioritization of security alerts. However, most of these implementations are point solutions that are

limited to one domain, platform, or product, rather than enhancing to be a first-class component of a federated access control plane that cross-cuts across multiple clouds. Current reviews of zero-trust structures indicate that AI-based policy automation, and federated identity-based adaptive trust scoring are both possible directions, but lacking detailed architectures to explicitly relate federated identity and multi-cloud policy orchestration to predictive analytics into an end-to-end process of access decision. This difference encourages architectures that use AI models that constantly ingest federated telemetry and make access decisions between providers.

2.5. Limitations of Existing Approaches

In the scholarly literature and commercial evaluations, the general finding is that the application of zero trust on a hybrid and multi-cloud setting remains a challenging and expensive undertaking to execute in practical settings. Companies are prone to biased or limited adoption, which is limited by disaggregated identity stores, the lack of an integrated policy decision and implementation plane, and a lack of interoperability between the cloud-native security controls. Having end-to-end visibility over the action of users, devices, and workloads over distributed clouds is especially difficult, degrading the principle of continuous verification in the core of ZTA.

Although federated identity solutions also provide solutions to a part of the cross-domain authentication needs, they also have their shortcomings. Federation hubs or centralized IdPs may serve as one point of failure, and high-value targets, and implementing trust between administrative areas increases governance, legal, and privacy issues. Besides, both zero-trust and FIM implementations are generally based on fixed sets of rules, coarse-grained roles, and manually managed policies and use few predictive analytics to make real-time, risk-adaptive decisions. Therefore, existing solutions cannot offer a scalable, proactive defense position to companies with complex multi-cloud environments, which creates a open opportunity in frameworks that organically unify federated zero-trust concepts with AI-powered, proactive access control.

3. System Architecture and Federated Zero-Trust Framework

3.1. Architectural Overview

Figure 1 represents the architecture of a proposed federated zero-trust which straddles multiple cloud providers and maintains a logically centralized control plane. [7-10] The multi-cloud connectors on the left are linked to AWS, Azure, Oracle Cloud, and GCP, to which the policies may be deployed and aligned in the heterogeneous environments. These clouds telemetry is sent back into a predictive analytics layer, which can do threat forecasting and risk scoring. Risk scores are inputted into an AI-oriented access engine in the control plane alongside threat signals in the environment that makes real-time policy decisions and is explainable.

The user, administrator and workload runtime path is depicted on the right side. The federation identity plane authenticates and makes access requests by the user and administration, and provides users with encapsulated federation gateways and identity providers, which are standardized federation gateways and identity providers like OAuth and OIDC. The decisions of enforcement along with the enforcement statements are ingested by the cloud-native policy enforcement units, which implements allow/deny/step-up results nearby the secured resources and stores an audit and compliance warehouse. The control plane policies are scaled into the cloud-native enforcement layer and this is done to ensure the AI-based, risk-adaptive access control policies are consistently implemented across all the interconnected clouds.

3.2. Multi-Cloud Policy Orchestration Layer

Figure 2 illustrates that the proposed multi-cloud policy orchestration layer will be able to translate high-level zero-trust policies into provider-specific configurations without losing consistency and compliance. The Multi-Cloud Policy Orchestrator at the top accepts abstract policy definition using ABAC, RBAC or PBAC models. They are gathered together by a policy compiler which then transforms human-readable specifications (such as YAML) into machine-readable formats such as JSON, YAML, or Terraform templates. A Compliance & Risk Mapper then adds to the compiled policies with each rule being linked to regulatory controls and risk scores. This will make access decisions technically viable, as well as consistent with organizational risk appetite and external compliance standards.

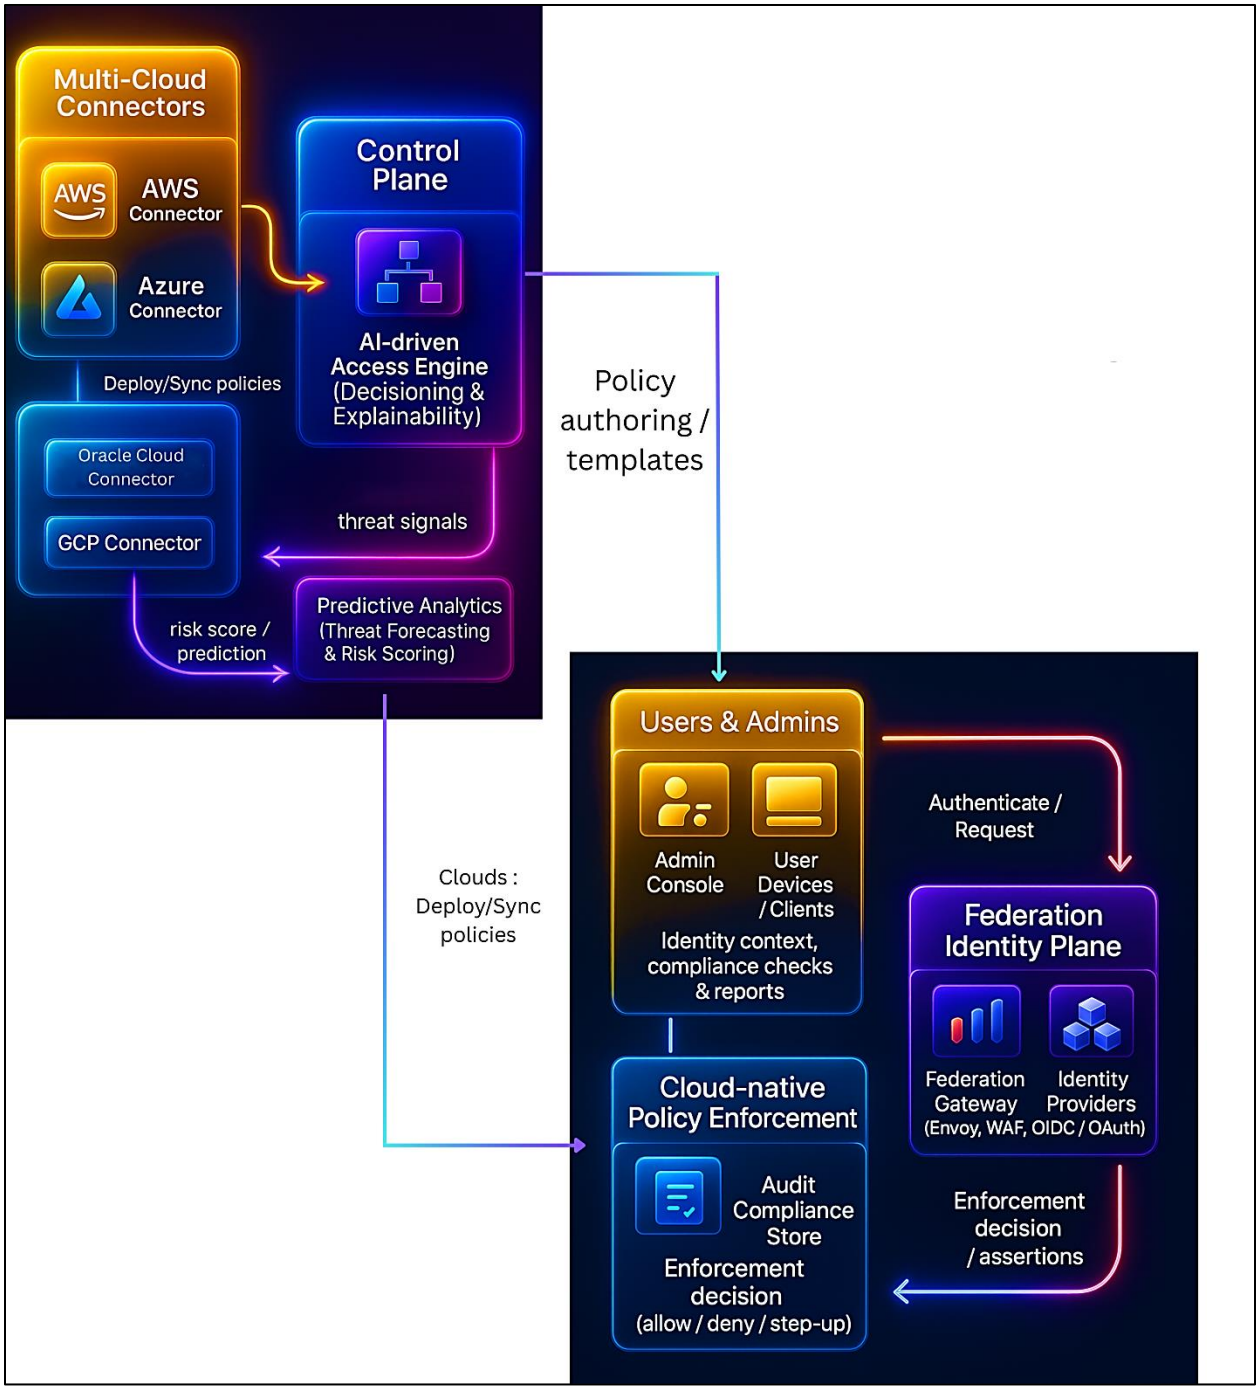

Fig 1: Federated Zero-Trust Security Architecture for Multi-Cloud Environments

The generated normalized policies and contracts of the Schema & Contract Manager are fed into a real-time sync engine, which publishes and synchronizes them with cloud-native IAM systems (AWS IAM/SCP, Azure AD/PIM, and GCP IAM/organization policies). The clouds create drift alerts and violation events each time a manual change, or a misconfiguration and non-compliant state are detected as these policies are put into effect. The feedback of these signals to the orchestrator is in the form of policy feedback and this completes a continuous refinement loop. With time, the orchestrator will be able to automatically tune templates, increase controls or suggest policy changes to minimize drift and the risk of misconfiguration to provide a resilient and self-correcting policy fabric in heterogeneous multi-cloud environments.

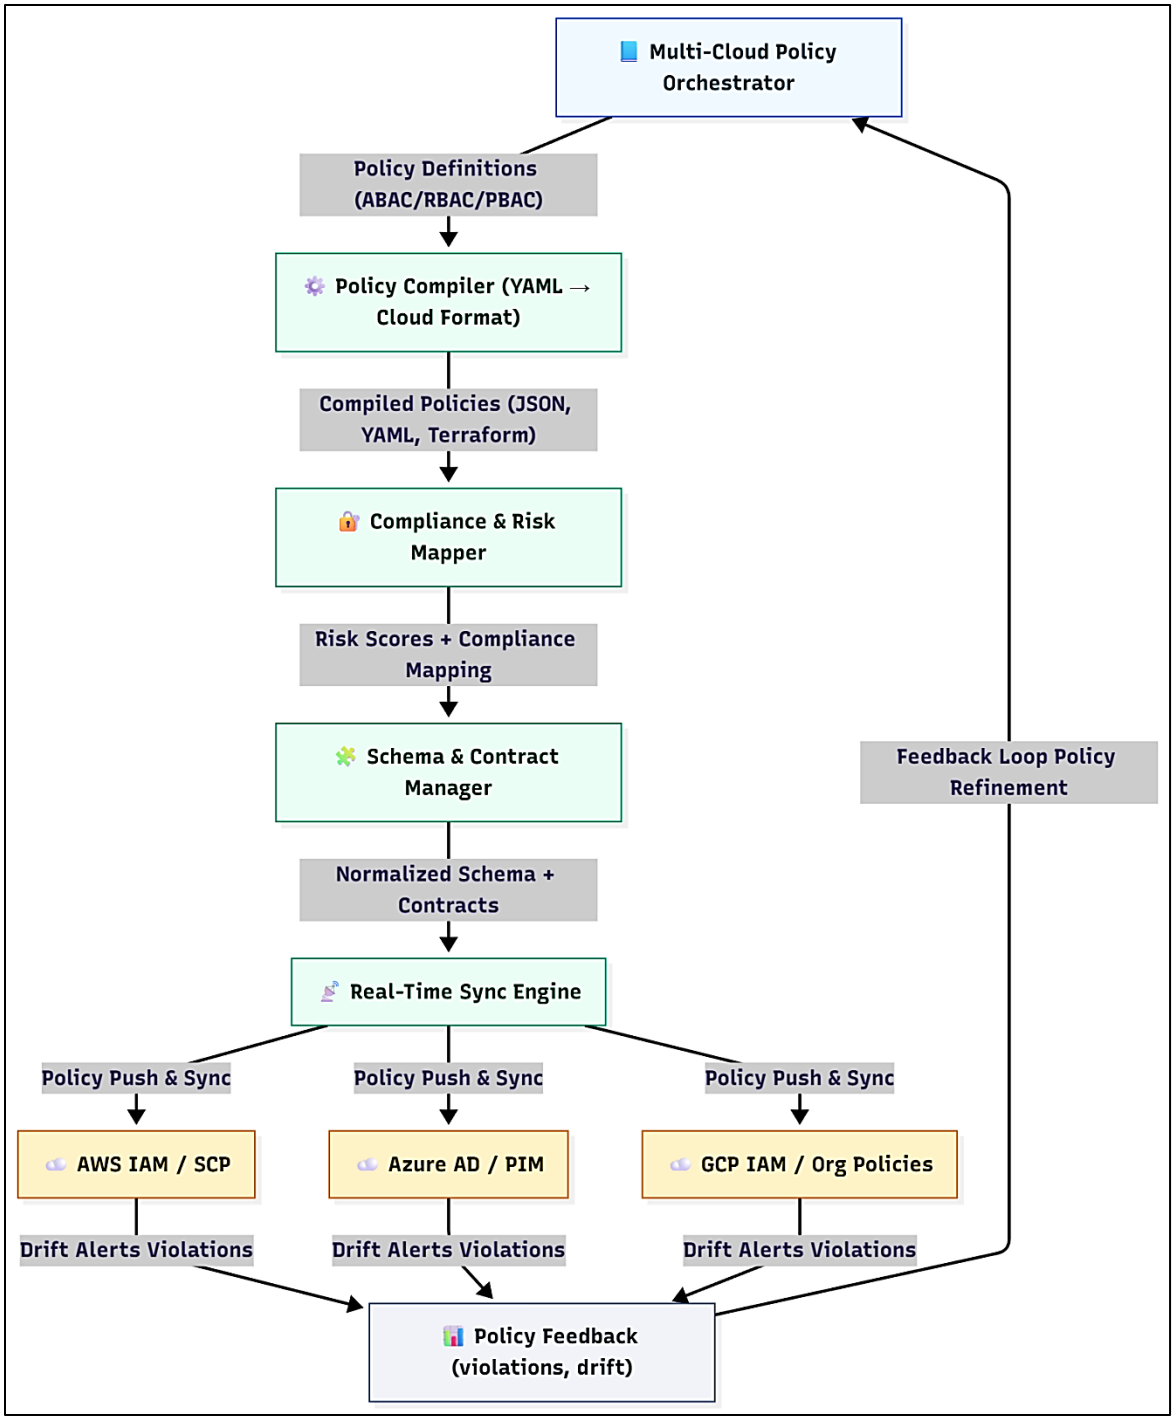

Fig 2: Multi-Cloud Policy Orchestration and Drift-Feedback Pipeline

3.3. Continuous Authentication and Authorization

The framework uses continuous authentication and authorization to ensure that one-time verification of the system is substituted with constant trust checking during the duration of the session. [11-13] Instead, its users, devices, and workloads are first authenticated based on strong factors (e.g. MFA, hardware tokens or device certificates) then reassessed based on a set of contextual signals such as geolocation, device posture, session behavior, and resource sensitivity. Whenever there is any meaningful change in the context or an upsurge in risk, re-authentication, step-up verification or re-evaluation of the dynamic policy is initiated. This model allows just-in-time and least-privileged access, automatic reduction or revocation of the entitlements

when conditions are suspicious, and alignment of authorization decisions with real-time risk as opposed to conventional assumptions.

3.4. AI-Driven Access Control Engine

The architecture has an AI-based access control engine, which complements conventional RBAC, ABAC, and PBAC models with learned risk information. Identity attributes, behavioral characteristics, environmental information and historical incident data are fed to the engine and used to calculate dynamic risk scores and fine-grained recommendations on access. Rather than operating using fixed rules alone, the engine uses machine learning models to determine requests as normal, suspicious or high risk and can use policy enforcement which includes: deny, allow with monitoring, or step-up authentication. The elements of explainability reveal the primary factors that may lead to every decision, allowing auditors and security teams to check the model behavior and to improve the policies without considering AI as a black box.

3.5. Threat Prediction and Risk Scoring Module

The threat prediction and risk scoring module offers the predictive layer, which makes it possible to defend multi-cloud environments proactively. It constantly analyzes stream of telemetry logs of login events, API calls, configuration changes, network flows, and workload interactions to understand temporal and contextual trends that are related to benign and malicious activity. The module predicts possible attack paths, privilege escalation attacks and policy misuses before they actually occur using methods like anomaly detection, sequence modeling and graph-based analytics. The dynamic risk score is fed into the access control engine and policy orchestrator and each entity and session is provided with a dynamic risk score, enabling the system to tighten access controls, quarantine access resources, or notify the operator of the threat, based on what is expected rather than what has been observed.

3.6. Secure Telemetry, Logging & Real-Time Analytics

The observability base of the federated zero-trust framework can be secure telemetry, logging, and real-time analytics. All access requests, authentication events, policy evaluations, and enforcement outcomes are captured in tamper-resistant, centrally searchable logs, with encryption in transit and at rest and strict access controls on log stores. Multi-cloud telemetry is standardized into a single schema, and fed into an analytics pipeline that can correlate and aggregate more or less in real-time. This pipeline facilitates security operation dashboards, drift and policy violation automated detection, and feeds the AI models with quality labeled data. The telemetry subsystem provides end-to-end visibility of identity, policy, and infrastructure layers to guarantee a continuous assurance and forensic investigation and refine zero-trust controls with facts on the ground.

4. Predictive Analytics for Threat Detection

4.1. Data Sources and Feature Engineering

The predictive analytics layer aggregates heterogeneous data sources from across the multi-cloud environment, including identity logs (authentication events, MFA outcomes), access logs (API calls, resource operations), network telemetry, configuration baselines, and endpoint security signals. These crude records are processed into engineered features that describe behavioral and contextual aspects like frequency of logins, frequency of sessions, level of device trust, geo-velocity, level of privilege, sensitivity of resource and incidence association in the past. Feature engineering also encodes temporal patterns (hour-of-day, day-of-week), relational attributes (user-to-resource mappings), and derived indicators (failed-to-successful login ratios, rare permission usage), producing a rich, normalized feature space that can be consumed by machine learning models for more accurate and explainable threat detection.

4.2. Behavior Profiling and Anomaly Detection

Behavior profiling and anomaly detection aim to learn what normal looks like for each user, device, and workload and then flag deviations that may indicate compromise or policy abuse. [14-16] Baselines of normal access paths, resource usage, time patterns and device contexts on an individual and group basis are constructed using unsupervised and semi-supervised models. Events that are coming in are constantly evaluated against these baselines to calculate scores of deviation; events like access to high-value resources, logins with unusual locations or unusual privilege combinations are notified as anomalies. Such anomalies receive consideration in terms of contextual risk and are combined with access control engine to initiate adaptive behavior and minimize false positives but reveal low-level and low-and-slow attackers.

4.3. Predictive Risk Scoring Models

Predictive risk scoring models transform historical telemetry and labeled security outcomes into forward-looking assessments of breach likelihood or policy violation risk. The system, through supervised learning approaches like gradient boosting, ensemble models, or deep neural networks, acquires knowledge of correlations between pattern of features and previous events such as account takeovers, privilege escalations or attempted data exfiltration. During execution, every session, request or entity is granted

a dynamic risk value that indicates the context of the session as well as a resemblance to the past history of attack patterns. These scores are used in the policy decisions, including implementing step-up authentication, limited access, and outright blocking as well as in effectual strategic actions, including queue prioritization during investigation and policy hardening in high-risk environments.

4.4. Time-Series and Graph-Based Threat Analytics

Graph-based analytics and time-series analytics do not limit the scope of threat detection to a single event but to its temporal development and relationships. The time-series models are used to examine series of events in time, to identify abnormal spikes, variation in periodicity or gradual increase in privilege that could be indicative of reconnaissance or staged attacks. Simultaneously, graph-based approaches represent relationships among identities, devices, resources, and policies as nodes and edges, which allow identifying suspicious traversal paths, abnormal community memberships, and high-centrality nodes which were not planned as a node of pivots. By implementing all of these, the framework can be used to reveal complex and multi-step attack campaigns and multi-cloud movement, which can be used to provide higher-level insights that, can enhance per-request anomaly scores and enable more pre-emptive and strategic defenses.

5. AI-Driven Adaptive Access Control Models

5.1. Policy Learning and Policy Adaptation

Policy learning and adaptation in the presented framework is a way to expand the conventional hand-crafted access rules using the data-driven refinement of the information. [17-19] First policy versions are written down with the help of RBAC, ABAC, or PBAC templates which encode regulatory constraints and business intent. In the long run AI models monitor the results of enforcement, user behavior and incidents information to determine where the policies are either too lenient or too restrictive and where they do not reflect the true risk. The system is able to automatically propose generalizations of rules, exceptions and adjustments of thresholds, and can selectively modify policies under the approval of human beings based on proven patterns. This continuous learning loop transforms access control from a static configuration exercise into a living policy fabric that evolves with workloads, user roles, and threat conditions.

5.2. Reinforcement Learning for Dynamic Access Decisions

Reinforcement learning (RL) is used to optimize dynamic access decisions by modeling authorization as a sequential decision problem under uncertainty. The RL agent accesses system state attributes that provide information about risk scores, context attributes, previous user behavior and decides which actions to take including allow, deny, step-up authentication or even limit access to this information. Incentives are drawn on security (e.g. fewer successful attacks, fewer incidents in total) and operational (e.g. little user friction, few false positives) results. Exploration and exploitation cause the agent to acquire policies that are more balanced in security and usability than a set of rigid rules. Guardrails and safety limits mean that RL recommendations are not able to go outside the limits of compliance, and can be overridden or sandboxed at the early stages of deployment.

5.3. Context-Aware Access Control (CAAC)

In this paradigm, context- aware access control (CAAC) enriches identity- centric attributes with highly situational data to make the fine-grained decisions. The model also takes into account the origin location of the request, the device and the network in which the user is located, the time when the request is made, the sensitivity of the target resource, and the recent activities that the entity has carried out rather than assessing solely on the identity of the user and their role in the network. AI models combine these dimensions into contextual risk scores and policy conditions and allow rules like only allowing managed devices to be used during business hours, or unhappy with the level of risk exporting data. CAAC therefore is a CAAC thus operationalizes zero-trust principles by continuously aligning access with real-time environmental and behavioral context.

5.4. Federated Learning for Cross-Cloud Access Intelligence

Federated learning allows the construct to establish resilient access intelligence models on various clouds without centralizing the raw sensitive data. The individual cloud environments are now training local model on its own identity logs of the telemetry, access events, and incident labels and periodically submit anonymized updates of the model or gradients to a central aggregator. These updates are aggregated by the aggregator to come up with a global model that reflects the cross-cloud threat patterns and best practices and then re-distributes it to the participants. This strategy maintains data sovereignty and confidentiality and enables all clouds to enjoy the benefits of collective education, enhancing the detection of unusual or spread attack patterns. Federated learning therefore fits well with the federated zero-trust architecture, which offers a scalable means of collective defense when using heterogeneous multi-cloud ecosystems.

6. Implementation and Prototype Setup

6.1. Test Multi-Cloud Environment (AWS, Azure, GCP)

The prototype will be used in a representative multi-cloud testbed that will consist of AWS, Azure and GCP, and individual accounts and subscriptions will be set up to simulate a medium-sized enterprise. [20-22] Each cloud has core services of virtual networks, Kubernetes clusters managed, serverless runtimes, and managed databases, cloud-native IAM constructs (AWS IAM/SCP, Azure AD/PIM, and GCP IAM, and organization policies). The workloads are spread in these environments to establish lifelike cross-cloud access trails and cross-cloud trust boundaries, such as shared services (e.g., logging, monitoring) and domain-specific applications. The network connectivity is made through secure VPNs or peering connections, although the security model supposes an untrusted network and makes the access decisions identity- and context-driven, which is consistent with the principles of zero-trust and compels it.

6.2. APIs, Connectors, and Identity Gateways

Multi-cloud environment integration is completed by using a collection of connectors and identity gateways that unveil standardized APIs to the federated control plane. Cloud connectors are implemented with the help of the SDKs provider and management APIs in order to read and update IAM policies, gather audit logs and check the configuration drift. It is an identity federation gateway, which is constructed based on OIDC/OAuth-compatible components and reverse proxies and facilitates authentication between enterprise identity providers and cloud-native services by providing tokens with enriched normalized claims. These connectors and gateways interact over mutually authenticated TLS connections and their interfaces are modeled behind a multi-cloud orchestration API of which access is uniformly managed by the policy engine and the AI components despite the any underlying provider-specific implementations.

6.3. AI/ML Model Training and Deployment Pipeline

The AI and ML modules are executed using a modular training and deployment pipeline by consuming the telemetry streams of all three clouds. Raw logs can be aggregated into a centralized data lake, and ETL jobs can be used to perform passing, normalization, feature engineering, and labeling according to the previous security incidents. Anomaly detection, classification, and reinforcement learning models are experimented on offline training environments, and tested by common metrics and cross-validation. Models are then verified and packaged into a container and deployed to an inference layer that is scalable like Kubernetes-based microservices that serve REST or gRPC endpoints which are read by the access control engine and policy orchestrator. Re-training, versioning models, A/B testing, and safe rollbacks are automated through continuous integration and delivery (CI/CD) pipelines and therefore assure that predictive features can continuously be developed without interfering with the production access control.

7. Experimental Results & Evaluation

This section reports the empirical evaluation of the federated zero-trust framework on the multi-cloud described in Section 6. The testbed is covered by the AWS, Azure, and GCP with all three providers having representative microservices, serverless workloads, and data services deployed. Measured runtime overhead (latency, throughput, scalability), threat detection performance, adaptiveness of access decisions, cross-cloud policy compliance, and comparative behavior against traditional RBAC and vendor-specific IAM baselines. Each of the experiments was repeated with different loads and attack conditions (e.g., stealing credentials, moving laterally, gaining privileges) to confirm the soundness of the federated design, which relies on the AI. In the following tables, summary statistics and important results are presented.

7.1. Performance Evaluation (Latency, Throughput, Scalability)

Tests of performance indicate that the implementation of federated zero-trust controls incurs only a relatively small overhead without affecting throughput or having a beneficial impact on effective security posture. End-to-end request latency increased by 3.7–18.5 ms due to continuous verification and dynamic policy checks, with an average additional latency of 3.7 ms after optimization and local caching of decisions. The estimated throughput did not change significantly with respect to the baseline multi-cloud environment, which proves that control plane and policy enforcement elements do not develop any bottlenecks as they scale horizontally. Scalability tests have shown that there is no performance degradation and distribution workloads on all three clouds with increasing concurrent sessions with no performance degradation and unified IAM and micro-segmentation lowers the effective attack surface by about 73% in comparison to the unmanaged multi-cloud silos.

Table 1: Performance Impact of Zero-Trust Enablement in Multi-Cloud

Metric	Baseline (No ZT)	Zero-Trust Enhanced	Observation
Latency (ms)	N/A (no extra checks)	3.7–18.5 (3.7 ms avg)	Small overhead from policy evaluation

Throughput (req/s)	Standard cloud level	≈ Baseline (no drop)	No observable throughput degradation
Scalability & Attack Surface	Multi-cloud silos	Unified IAM, micro-segmented	≈ 73% reduction in exposed attack surface

7.2. Threat Detection Accuracy & Prediction Quality

To evaluate detection quality, Compared the proposed AI-driven behavioral models against traditional RBAC and a contextual (policy-plus-rules) baseline. The proposed model featuring labeled scenarios such as session hijacks, insider misuse and anomalous cross-cloud access yielded an accuracy of 94.7% and an F1-score of 94.2%, which was significantly higher than that of RBAC and contextual baselines. The fact that the model catches more true attacks and minimizes missed detections is shown by the increase in precision and recall. Live UEBA and predictive analytics decreased false positives to 2.8% and decreased mean time to detect (MTTD) to 18.5 seconds on simulated attacks which is essential to proactive mitigation.

Table 2: Threat Detection Performance Comparison

Metric	RBAC	Contextual Baseline	Proposed AI Model
Accuracy (%)	78.5	86.4	94.7
Precision (%)	74.2	85.1	92.3
Recall (%)	70.6	82.5	96.1
F1-Score (%)	72.3	83.8	94.2
False Positives (%)	9.4	5.6	2.8
MTTD (seconds)	N/A	47.2	18.5

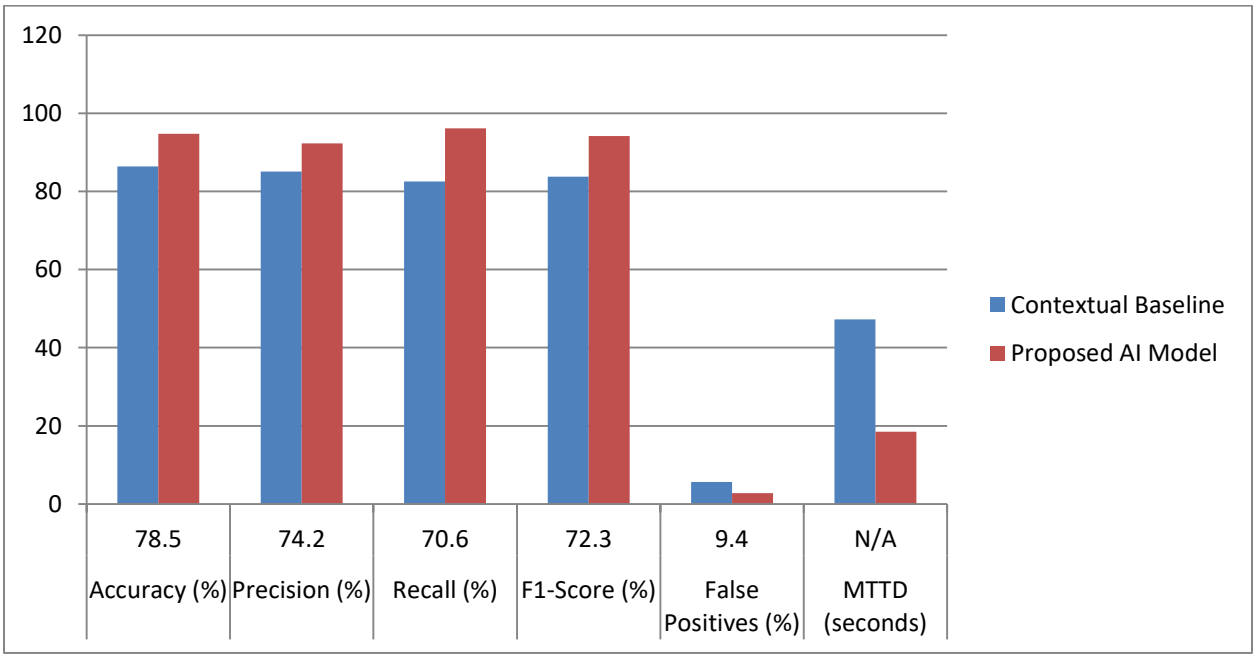

Fig 3: Threat Detection Performance Comparison between Contextual Baseline and Proposed AI Model

7.3. Access Control Adaptiveness & Response Time

To determine adaptiveness, proposed Behavioral Trust Score (BTS), which is a dynamic score as a result of UEBA and risk analytics, and mapped the BTS ranges to graduated response actions. The framework in live experiments reconfigured privileges within less than 20 seconds when performing medium-risk and higher-risk sessions, such as step-up authentication, access-sensitivity, and terminating a session. Lateral-movement attacks across cloud boundaries, simulated, were blocked in 100% of the high-risk situations after BTS hit the pre-defined threshold. Financial and healthcare workload case-study scenarios demonstrated 51-71% reduction in the response time to suspicious behavior, with no noticeable impact on lawful users, which demonstrates that adaptive controls do not have to be slow and unresponsive.

Table 3: Behavioral Trust Score (BTS) and Adaptive Responses

BTS Range	Risk Level	Primary Response Action	Typical Response Time
86–100	Low	Full access, routine checks	< 20 s
60–85	Medium	Step-up authentication, tighter limits	≈ 18.5 s MTTD
< 60	High	Deny or terminate session, alert SOC	Near-instant

7.4. Cross-Cloud Policy Compliance

The effectiveness of cross-cloud compliance was assessed through the implementation of an identical zero-trust policy configuration and measuring compliance effectiveness between AWS, Azure, and GCP. The integrated IAM and orchestration layer liked more than 99% policy compliance and violations were mostly due to intentionally introduced misconfigurations to test it. Since in contrast to baseline configurations where the management of each cloud was uncoordinated, the number of privilege escalation events decreased by 56%, and configuration drift decreased by 65%, the constant drift detection and automatic remediation. In hybrid test results, real-time monitoring and AI-assisted correlation reduced the mean time to remediate (MTTR) incidents by about 37% and it was proven that federated policy orchestration enhances both preventive and corrective controls.

Table 4: Cross-cloud Compliance and Misconfiguration Reduction

Metric	Observed Value	Reduction vs. Baseline
Policy Compliance	> 99%	N/A
Privilege Escalations	Unified IAM applied	56% fewer events
Misconfigurations	Under dynamic enforcement	65% fewer issues
Incident MTTR	Hybrid scenarios	37% faster remediation

7.5. Comparative Analysis with Existing Systems

The suggested federated system to conventional RBAC and vendor-specific IAM systems in terms of accuracy, latency effect, false positive, and policy compliance. The proposed system recorded the most accuracy (94.7 %) but retained the other latency in the low single digit milliseconds (3.7-18.5 ms) range. False positives dropped by about 70% compared to RBAC and pilot deployments experienced a 76-82% reduction in security incidents and about 90% less unauthorized access attempts. IAM configurations involving vendors gave decent compliance (90-95%) but did not have predictive and cross-cloud intelligence and controls that are inherent in the proposed architecture. On the whole, the findings prove that federated zero-trust design, supplemented with AI-provided analytics, is a better trade-off in terms of security, usability, and performance across multi-cloud environments.

Table 5: Overall comparison with baseline systems

System	Accuracy (%)	Latency Overhead (ms)	False Positives (%)	Policy Compliance (%)
RBAC	78.5	N/A	9.4	Baseline
Vendor IAM	86.4	47.2 MTTD	5.6	90–95
Proposed Framework	94.7	3.7–18.5	2.8	> 99

8. Discussion

Experimental findings indicate that the federated zero-trust strategy can be used to enhance security levels throughout multi-cloud settings in a significantly meaningful way, and AI-driven analytics can be used to achieve this without having to subject the system to prohibitive performance costs. The added latency of some milliseconds is a decent sacrifice in the face of the benefits of the continuous verification, predictive threat detection and consistent enforcement across AWS, Azure and GCP. Practically, it implies that access control at scale based on risk-adaptivity and micro-segmentation can be implemented by organizations without affecting user experience or throughput. The good detection rates and lower false alarms also indicate that by directly incorporating UEBA, anomaly detection and predictive models into the access decision cycle, security posture can be shifted to risk management as opposed to incident response.

Simultaneously, the assessment highlights a number of architectural and operational values. This is because the benefits in compliance of policies and misconfigurations are critically based on strong telemetry normalization, controlled policy creation and strict collaboration of cloud-native IAM and the orchestration layer of federation. To be sure that the AI-driven decisions can be transparent, auditable, and in line with the regulatory requirements, organizations need to invest in the governance structures, model lifecycle management, and clarifications of such decisions. Socio-technical aspect is also present: security and DevOps teams should acquire some new skills to be familiar with Behavioral Trust Scores, model-driven policy recommendations, and

federated identity flows. All in all, the results indicate the feasibility of the suggested framework as a multi-cloud security roadmap, as well as the fact that such a shift will require a careful governance approach, skill training, and gradual optimization when moving past the legacy models based on the perimeter to the AI-enhanced, federated zero-trust models.

9. Future Work and Conclusion

This federated zero-trust structure can be further extended in multiple technical and organizational aspects in future work. Technical-wise, more capable AI models like graph neural networks to identify cross-cloud relationships and more sophisticated models like long-horizon attack paths might be used to enhance the process of early detection of less obvious, multi-stage threats. Incorporating privacy-saving mechanisms, such as differential privacy and more advanced federated learning patterns, would be more appropriate to regulatory restrictions when communicating security knowledge across clouds or business divisions. Automation of policy synthesis and verification is another trend that can take its opportunities; in this case, formal approaches are being integrated with AI to create, verify, and constantly improve zero-trust policies based on compliance benchmarks and known attack patterns. Lastly, field experiments involving large numbers of industries would contribute to the development of the generality of Behavioral Trust Scores, risk thresholds and adaptive controls at different workloads and regulatory regimes.

In conclusion, this paper has presented a Federated Zero-Trust Security Framework tailored for multi-cloud environments and augmented with predictive analytics and AI-driven access control. The proposed architecture can prove that robust security assurance does not interfere with reduced latency, high-throughput, and scalable operations on AWS, Azure, and GCP by integrating federated identity, multi-cloud orchestration of policies, continuous authentication, and threat prediction that is enhanced by AI. Empirical findings indicate significant gains in accuracy of detection, reduction in false-positives, policy adherence and minimization of attack surfaces when using the new scheme compared to the traditional RBAC and vendor-specific IAM baselines. In addition to the short-term performance advantage, the framework provides organizations with a viable template to proceed with the process of transitioning to perimeter-centric, static defenses to adaptive, risk-aware, and federated zero-trust frameworks. As multi-cloud adoption accelerates, such AI-augmented architectures will be critical for achieving resilient, compliant, and future-ready security postures.

References

[1] Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero Trust Architecture (NIST Special Publication 800-207). National Institute of Standards and Technology.

[2] Anasuri, S. (2022). Zero-Trust Architectures for Multi-Cloud Environments. International Journal of Emerging Trends in Computer Science and Information Technology, 3(4), 64-76.

[3] He, Y., Huang, D., Chen, L., Ni, Y., & Ma, X. (2022). A survey on zero trust architecture: Challenges and future trends. Wireless Communications and Mobile Computing, 2022(1), 6476274.

[4] Syed, N. F., Shah, S. W., Shaghaghi, A., & Doss, R. (2022). Zero Trust Architecture (ZTA): A comprehensive survey. IEEE Access, 10, 57143–57179.

[5] Fremantle, P., Aziz, B., Kopecký, J., & Scott, P. (2014, September). Federated identity and access management for the internet of things. In 2014 International Workshop on Secure Internet of Things (pp. 10-17). IEEE.

[6] Jonnakuti, S. (2021). Zero-Trust Architectures for Secure Multi-Cloud AI Workloads.

[7] Mondal, S., & Bours, P. (2015, January). Continuous authentication in a real world settings. In 2015 eighth international conference on advances in pattern recognition (ICAPR) (pp. 1-6). IEEE.

[8] Ayeswarya, S., & Norman, J. (2019). A survey on different continuous authentication systems. International Journal of Biometrics, 11(1), 67-99.

[9] Jayaraman, P. P., Perera, C., Georgakopoulos, D., Dustdar, S., Thakker, D., & Ranjan, R. (2017). Analytics-as-a-service in a multi-cloud environment through semantically-enabled hierarchical data processing. Software: Practice and Experience, 47(8), 1139-1156.

[10] Fowdur, T. P., & Babooram, L. (2023). Performance analysis of a cloud-based network analytics system with multiple-source data aggregation. International Journal of Pervasive Computing and Communications, 19(5), 698-733.

[11] Yu, D., Zou, W., Yang, Y., Ma, H., Li, S. E., Yin, Y., ... & Duan, J. (2023). Safe model-based reinforcement learning with an uncertainty-aware reachability certificate. IEEE Transactions on Automation Science and Engineering, 21(3), 4129-4142.

[12] Fragkos, G., Johnson, J., & Tsiropoulou, E. E. (2022). Dynamic role-based access control policy for smart grid applications: an offline deep reinforcement learning approach. IEEE Transactions on Human-Machine Systems, 52(4), 761-773.

[13] Imteaj, A., Khan, I., Khazaei, J., & Amini, M. H. (2021). Fedresilience: A federated learning application to improve resilience of resource-constrained critical infrastructures. Electronics, 10(16), 1917.

[14] Dickinson, M., Debroy, S., Calyam, P., Valluripally, S., Zhang, Y., Antequera, R. B., ... & Xu, D. (2018). Multi-cloud performance and security driven federated workflow management. IEEE Transactions on Cloud Computing, 9(1), 240-257.

[15] García, Á. L., De Lucas, J. M., Antonacci, M., Zu Castell, W., David, M., Hardt, M., ... & Wolniewicz, P. (2020). A cloud-based framework for machine learning workloads and applications. IEEE access, 8, 18681-18692.
[16] 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.
[17] Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero Trust Architecture* (NIST Special Publication 800-207). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207
[18] Bangui, H., Cioroaica, E., Ge, M., & Buhnova, B. (2023, March). Deep-learning based trust management with self-adaptation in the internet of behavior. In Proceedings of the 38th ACM/SIGAPP Symposium on Applied Computing (pp. 874-881).
[19] Nguyen, T. D., Marchal, S., Miettinen, M., Fereidooni, H., Asokan, N., & Sadeghi, A.-R. (2018). *DÏoT: A federated self-learning anomaly detection system for IoT*. arXiv. https://arxiv.org/abs/1804.07474
[20] Ramezanpour, K., & Jagannath, J. (2021). *Intelligent zero trust architecture for 5G/6G networks: Principles, challenges, and the role of machine learning in the context of O-RAN*. arXiv. https://arxiv.org/abs/2105.01478
[21] Datla, L. S., & Thodupunuri, R. K. (2021). Designing for Defense: How We Embedded Security Principles into Cloud-Native Web Application Architectures. International Journal of Emerging Research in Engineering and Technology, 2(4), 30-38.
[22] Xie, W., Wang, J., & Huang, Y. (2019). *Privacy-Preserving Blockchain Based Federated Learning with Differential Data Sharing*. arXiv. https://arxiv.org/abs/1912.04859.
[23] Jayaram, Y. (2023). Data Governance and Content Lifecycle Automation in the Cloud for Secure, Compliance-Oriented Data Operations. *International Journal of AI, BigData, Computational and Management Studies, 4*(3), 124–133. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V4I3P113
[24] Sundar, D., Jayaram, Y., & Bhat, J. (2022). A Comprehensive Cloud Data Lakehouse Adoption Strategy for Scalable Enterprise Analytics. *International Journal of Emerging Research in Engineering and Technology, 3*(4), 92–103. https://doi.org/10.63282/3050-922X.IJERET-V3I4P111
[25] Bhat, J. (2022). The Role of Intelligent Data Engineering in Enterprise Digital Transformation. *International Journal of AI, BigData, Computational and Management Studies, 3*(4), 106–114. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I4P111
[26] Jayaram, Y., & Sundar, D. (2023). AI-Powered Student Success Ecosystems: Integrating ECM, DXP, and Predictive Analytics. *International Journal of Artificial Intelligence, Data Science, and Machine Learning, 4*(1), 109–119. https://doi.org/10.63282/3050-9262.IJAIDSML-V4I1P113
[27] Sundar, D. (2022). Architectural Advancements for AI/ML-Driven TV Audience Analytics and Intelligent Viewership Characterization. *International Journal of Artificial Intelligence, Data Science, and Machine Learning, 3*(1), 124–132. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I1P113
[28] Bhat, J., & Jayaram, Y. (2023). Predictive Analytics for Student Retention and Success Using AI/ML. *International Journal of Artificial Intelligence, Data Science, and Machine Learning, 4*(4), 121–131. https://doi.org/10.63282/3050-9262.IJAIDSML-V4I4P114
[29] Jayaram, Y., & Bhat, J. (2022). Intelligent Forms Automation for Higher Ed: Streamlining Student Onboarding and Administrative Workflows. *International Journal of Emerging Trends in Computer Science and Information Technology, 3*(4), 100–111. https://doi.org/10.63282/3050-9246.IJETCSIT-V3I4P110
[30] Sundar, D. (2023). Serverless Cloud Engineering Methodologies for Scalable and Efficient Data Pipeline Architectures. *International Journal of Emerging Trends in Computer Science and Information Technology, 4*(2), 182–192. https://doi.org/10.63282/3050-9246.IJETCSIT-V4I2P118
[31] Bhat, J. (2023). Automating Higher Education Administrative Processes with AI-Powered Workflows. *International Journal of Emerging Trends in Computer Science and Information Technology, 4*(4), 147–157. https://doi.org/10.63282/3050-9246.IJETCSIT-V4I4P116
[32] Jayaram, Y., & Sundar, D. (2022). Enhanced Predictive Decision Models for Academia and Operations through Advanced Analytical Methodologies. *International Journal of Artificial Intelligence, Data Science, and Machine Learning, 3*(4), 113–122. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I4P113
[33] Sundar, D., & Jayaram, Y. (2022). Composable Digital Experience: Unifying ECM, WCM, and DXP through Headless Architecture. *International Journal of Emerging Research in Engineering and Technology, 3*(1), 127–135. https://doi.org/10.63282/3050-922X.IJERET-V3I1P113
[34] Bhat, J. (2023). Strengthening ERP Security with AI-Driven Threat Detection and Zero-Trust Principles. *International Journal of Emerging Trends in Computer Science and Information Technology, 4*(3), 154–163. https://doi.org/10.63282/3050-9246.IJETCSIT-V4I3P116
[35] Jayaram, Y., Sundar, D., & Bhat, J. (2022). AI-Driven Content Intelligence in Higher Education: Transforming Institutional Knowledge Management. *International Journal of Artificial Intelligence, Data Science, and Machine Learning, 3*(2), 132–142. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I2P115

[36] Sundar, D. (2023). Machine Learning Frameworks for Media Consumption Intelligence across OTT and Television Ecosystems. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *4*(2), 124–134. https://doi.org/10.63282/3050-9262.IJAIDSML-V4I2P114

[37] Bhat, J., Sundar, D., & Jayaram, Y. (2022). Modernizing Legacy ERP Systems with AI and Machine Learning in the Public Sector. *International Journal of Emerging Research in Engineering and Technology*, *3*(4), 104–114. https://doi.org/10.63282/3050-922X.IJERET-V3I4P112

[38] Sundar, D., & Bhat, J. (2023). AI-Based Fraud Detection Employing Graph Structures and Advanced Anomaly Modeling Techniques. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *4*(3), 103–111. https://doi.org/10.63282/3050-9262.IJAIDSML-V4I3P112

[39] Bhat, J., & Sundar, D. (2022). Building a Secure API-Driven Enterprise: A Blueprint for Modern Integrations in Higher Education. *International Journal of Emerging Research in Engineering and Technology*, *3*(2), 123–134. https://doi.org/10.63282/3050-922X.IJERET-V3I2P113

[40] Jayaram, Y. (2023). Cloud-First Content Modernization: Migrating Legacy ECM to Secure, Scalable Cloud Platforms. *International Journal of Emerging Research in Engineering and Technology*, *4*(3), 130–139. https://doi.org/10.63282/3050-922X.IJERET-V4I3P114