



Original Article

# Zero-Trust Multi-Cloud Architecture for Secure Experience Management

Siva Sai Krishna Suryadevara<sup>1</sup> Anjani Kumar Polinati<sup>2</sup>

<sup>1</sup>Sr. AEM Cloud Engineer at Maganti IT Resources , USA.

<sup>2</sup>Senior Software Engineer at Primoris Systems LLC, USA.

*Abstract - Organizations that manage customer and employee experiences across many other cloud platforms now need secure experience management more than ever. As workloads, data, and digital interfaces grow on hyperscalers and SaaS platforms, the user experience is more and more affected by the security & their dependability of service authentication, authorization, and recovery under duress, rather than just by the capabilities of the application. Most multi-cloud security architectures still assume that there are very strong perimeters or trust based on network location, which doesn't work in actual life. They often create identity sprawl, which happens when people, services, and machines have different responsibilities in different cloud environments; policy drift, which happens when access rules differ by their platform and team; and telemetry silos, which make it very hard to see lateral movement and slow down their investigations. This article suggests a zero-trust multi-cloud reference architecture that is specifically made for safe experience management to fill these shortcomings. The architecture assumes that every request is untrustworthy and requires constant verification through strong identity federation, device & workload attestation, and context-sensitive signals (risk, location, behavior, service health). Least-privilege, just-in-time permissions along with micro-segmented service routes limit access. Adaptive policy engines turn intent into consistent controls across cloud environments. A unified observability architecture fully integrates identity, policy decisions, and runtime telemetry, making it easy to find these problems very quickly and respond automatically. A case study shows how to lower risk by lowering the number of over-privileged identities, improve SLA compliance by using strong, policy-driven failover methods & speed up incident resolution by lowering the blast radius and speeding up finding the root cause. The design shows how confidence decisions are constantly available, portable, and completely visible, which may enhance both trustworthy and secure security at the same time. This design accommodates any cloud as well as is based on their general zero-trust concepts. Many other areas, like banking, healthcare, and SaaS, in which seamless and safe multi-cloud experiences are vital, could additionally employ it.*

*Keywords - Zero-Trust, Multi-Cloud Security, Secure Experience Management (SXM), Continuous Authentication, Policy-as-Code, Identity Federation, Micro-segmentation, SASE, Observability, Risk-Adaptive Access Control, Cloud Workload Protection, User Experience Telemetry.*

## 1. Introduction

The expectations of individuals have altered a lot since business enterprises started using computer systems. Applications need to be very fast, versatile, constantly available, and safe, irrespective of where they are kept or who employs them. Because of this notion, companies are now starting to use these distributed cloud designs, which split workloads, pipelines of information, and services that interface with customers over several publicly accessible clouds, private clouds, along with SaaS ecosystems. This approach makes things more flexible & very strong, but it also creates a lot of security problems that can make the digital experience worse for customers, employees, and partners.

Traditional security frameworks were built around central control and fixed boundaries. In a situation where user journeys go across AWS, Azure, GCP, on-premises, and third-party SaaS platforms all in one encounter, these kinds of assumptions don't hold up. As a result, protecting the "experience layer" has become a complex job. Every login, API request, webhook, microservice execution & data movement between clouds could be a security hole.

This is where Secure Experience Management (SXM) comes in. SXM puts the safety, privacy, ease of use & dependability of all these digital interactions first. This includes customers browsing a product site, employees using internal tools, and partners connecting through APIs. In a multi-cloud scenario, making sure that level of safety requires a completely different approach. A single incorrect IAM role in the cloud, an API gateway policy that isn't being watched, or a SaaS integration that isn't being watched can all damage user trust or ruin their experiences.

Security teams often have to deal with visibility silos, unequal control frameworks, and a growing number of "invisible services" that are begun and then forgotten about. At the same time, security precautions can slow down product & platform teams or force them to employ different methods depending on the cloud environment they are using for deployment. The

framework for dealing with this complexity is based on their zero-trust principles, which are based on constantly checking identity, context, and risk.

SXM's Zero-Trust Multi-Cloud Architecture aims to bring together these different parts into a single, always-enforced security framework. Instead of relying on the constraints of any other cloud provider, it requires strict identity verification, consistent policy enforcement, and real-time visibility across all these cloud environments. This introduction explains why this kind of architecture is needed and what problems it solves.

### **1.1. Challenges**

Numerous businesses are looking for flexibility, vendor neutrality, cost-effectiveness & access to better services from more numerous providers, which has made multi-cloud solutions more popular. This distribution gives both apps and duties to several layers of security and operations. As a result, the attack surface gets bigger and more complicated.

A major worry is that identification and identity access management strategies are becoming very less stable. Every cloud has its particular verification plugins, token kinds, authorization models, and role structure. When teams mix these settings together, it makes the handling of identities much less standardized. This makes checks harder and gives attackers additional opportunities to make advantage of mistakes or variances in access privileges.

The network controls are also failing to function right. AWS security groups, Azure NSGs, GCP rules regarding firewalls, and other service mesh arrangements all work in their own ways. Without standardized enforcement, security professionals have to have to manually align restrictions or put up with safeguards that they aren't always the same. Both of those possibilities are exceedingly demanding in massive systems.

Additional issues to worry about are shadow solutions and SaaS drift. Teams quickly start using SaaS apps, managed services, alongside cloud-native capabilities that possibly contravene corporate norms if they aren't kept an eye on. This causes loosely organized identities, secret data flows, as well as integrations which are not allowed.

Telemetry problems make the situation worse. Different clouds have quite different logs, traces, and metrics, which makes it very hard to put them all together into a single tale. Without cross-cloud observability, threats might not be found, and it gets harder to figure out what's wrong with the user experience.

In the end, modern digital services depend heavily on links to third parties, such as payment gateways and analytics systems. These integrations often bypass centralized security safeguards and create unverified trust links, which makes things more vulnerable.

In this way, Secure Experience Management means keeping every interaction in the user or customer journey safe while making sure that privacy, reliability & performance are all maintained. SXM knows that security is an important part of the whole experience, and any other failure, whether it's an outage or a breach, makes people less trustworthy.

### **1.2. Problem Statement**

Organizations need a security architecture that is more consistent and works well in a variety of multi-cloud environments to give users the best digital experiences. Businesses have adopted cloud services to become more flexible, but they often have trouble making sure that their security protocols are the same across AWS, Azure, GCP, and the growing number of SaaS platforms in their ecosystem. As a result, security teams face operational blind spots, inconsistent enforcement procedures & rules that don't work together.

This inconsistency has a direct effect on measurable outcomes. Policy latency, or the time between when a security rule is made and when it is put into action across all these cloud settings, creates security holes. User friction happens when each cloud has its own rules for logging in, multi-factor authentication, or session management. This makes it harder for businesses to provide a smooth experience. When visibility is broken up, breach dwell time usually goes up, which means that incident responders have to manually link information from different platforms. SLA variation happens when workloads spread across clouds show different patterns of reliability, making it harder to guarantee that these users will always have the same experience. Compliance risk rises as regional rules need clear data lineage and consistent controls across different areas.

To sum up, organizations work in these ecosystems where the quality of experience and security posture depend on each other. Without a zero-trust framework that includes identity, access, telemetry & policy enforcement, organizations can't ensure that digital experiences will be reliable & resilient. This difference means that a Zero-Trust Multi-Cloud Architecture needs to be put in place that is specifically made for Secure Experience Management.

### 1.3. Motivation

Zero-trust is no longer just a way to protect yourself; it's a must for every business that offers digital experiences in several other cloud environments. As services are spread out among many providers, the line between them gets less, and identification becomes the main way to govern them. It is important to always check all requests, API calls, and data transactions based on the context rather than assumptions.

This need becomes much stronger in the field of Secure Experience Management. Outages cause problems with the user experience, which hurts trust in the brand & the impression it leaves. Breaches lead to trust problems directly, especially in fields that deal with personal or financial information. Product teams have to slow down or get around security measures because of inconsistent rules across cloud environments. This makes it harder for the latest ideas to come about.

The industry additionally places hardships on them. As rules concerning privacy, managing data regardless of borders, and the supply chain risk grow stricter, it's becoming increasingly essential for enterprises to have identical and verifiable controls, irrespective of which cloud service supplier is hosting the computing power. Zero-trust assists you attain these types of goals while simultaneously being sure that they operate well and have been reliable.

In the end, enterprises can give prospects safe, pleasant interactions with SXM without bogging things down or causing them to be less flexible.

## 2. Literature Review

### 2.1. Zero-Trust Foundations

Zero-Trust Architecture (ZTA) has evolved from a niche security notion to a widespread design goal, mostly due to the failure of conventional perimeter defenses to correspond with their actual system usage. The main idea behind Zero Trust Architecture (ZTA) is very simple: "never trust, always verify," use the least amount of privilege, and assume that a breach is the normal state of affairs. Zero Trust Architecture (ZTA) doesn't assume that everything inside a network perimeter is safe. Instead, it treats any request—whether it's from a person, device, service, or workload—as potentially very harmful until proven otherwise.

Three realistic variations keep coming up during the research. Identity-centric zero trust puts authentication & ongoing authorization at the top of the list. It focuses on who is asking for access (human users, service accounts, API clients) & uses things like device posture, location, behavioral baselines, and risk assessment to decide who can get in. This profession says that identification is the only consistent control point in their networks that change and workloads that change.

Network-centric zero trust takes a very old but still useful view: separate all these parts, make trust zones as small as possible, and check traffic at every point. It happens a lot when network security teams already have the tools for enterprise retrofits. In modern cloud systems, it doesn't function since these networks alone don't define purpose; a "clean" network channel doesn't always mean a safe activity.

Workload zero trust is the third part. It extends Zero Trust Architecture (ZTA) to service-to-service communication, such as containers, serverless computing, and microservices. The main goals are mutual authentication, service identities with the least amount of access, and enforcing policies at runtime. The literature on workload ZT often crosses with service mesh research, as meshes provide the infrastructure for identity, encryption & policy among services.

A key part of comparing these types is the part that is seen as the basis of trust. Identity-centric Zero Trust is based on the identities of users and services. Network-centric Zero Trust is based on traffic paths, and workload-centric Zero Trust is based on software components. Most writers agree that actual ZTA systems have all three parts, although the balance changes depending on how mature the environment is and what threats are most important.

**Table 1. Zero-Trust Variants and Primary Enforcement Points**

Zero-Trust Variant	Primary Trust Anchor	Typical Enforcement	Relevance to This Work
Identity-centric ZT	User/service identity	IdP, MFA, conditional access, session risk	Core for consistent cross-cloud access
Network-centric ZT	Traffic path/zone	Micro-segmentation, ZTNA, firewall policy	Limits lateral movement between clouds
Workload ZT	Service/workload identity	mTLS, service mesh policy, runtime attestation	Secures east-west & service-to-service calls

## **2.2. Security Frameworks for Multiple Clouds**

People usually choose to use many other clouds because they are more reliable, cheaper, or because they believe in using the "best tool for the job." As each cloud provider sees itself as the main one, security gets harder and harder. The research shows that native security stacks, especially Identity and Access Management (IAM), are consistent on the inside but not so much on the outside. AWS IAM policies, Azure RBAC and Conditional Access models, and Google Cloud IAM roles all make different assumptions about how identity hierarchy, resource scope & policy evaluation work.

Researchers assert that this inconsistency obstructs appropriate composition. Two clouds may have similar primitives (roles, tokens, groups, policies), but their meanings are very different in subtle ways. The meanings of "resource boundary," how inheritance works, and how exceptions are handled may all be different. So, translation layers or manual equivalence are often used in multi-cloud setups, and both of them are very risky.

"Best-effort parity" is a constant concern since teams strive to set up the same rules on numerous clouds, thinking that symmetry means security. But parity is usually not exact. Features come and go at different times, logs are not always reliable, and provider-specific features make teams want to include localized exceptions. The literature warns that thinking in terms of parity might lead to blind spots because attackers only need the weakest cloud to get an advantage.

Single-cloud security methods, on the other hand, can assume that telemetry is the same and that rules are always followed. Multi-cloud architectures need to take into account that things are different and change all the time. This is why many studies suggest that these policies should be higher-level and not favor any one provider, with local enforcement mechanisms instead of trying to make everything the same at the most basic level.

## **2.3. Research on Secure Experience Management**

Secure Experience Management (SXM) combines security assurance with monitoring of digital experiences. Digital Experience Monitoring (DEM) is where the idea for SXM comes from. DEM looks at how users interact with apps, including latency, failures, transaction paths & actual user monitoring. SXM improves this framework by adding a security-conscious context. The goal is not just to find a broken user journey, but also to figure out if the problem is caused by an attack, a misconfiguration, or dangerous conduct.

Many frameworks link SXM to ITSM and ITOM workstreams. Experience signals can begin incident procedures, automate processes for fixing problems, or change how change management works. The "closed loop" is important for scalability since it doesn't just notice experiencing difficulties; it also fixes them.

The current SXM research faces challenges in multi-cloud threat modeling. Most suggested SXM pipelines assume that user journeys happen in a single provider or a carefully controlled hybrid environment. They are great at connecting application performance to identity & endpoint risk, but not so good at looking at cross-cloud lateral movement, policy drift between providers, or inconsistent telemetry. A user's experience may include different clouds, SaaS edges & service meshes. In a single cloud-focused SXM model, the weak link may not be easy to see.

SXM solutions capture "user sentiment," but they often don't take into their account the real movement of risk when trips are spread out over multiple clouds.

## **2.4. Related Architectures and Shortcomings**

Several adjacent designs try to solve parts of this problem. SASE/SSE stresses moving security limits to the edge of the network and making sure that users & devices can only access certain areas of the network based on their identity. They believe in zero trust, but they often prefer human access channels to trusting service-to-service or cross-cloud workloads.

CNAPP platforms try to bring together cloud posture, workload protection, and managing access rights. They do a better job of managing multi-cloud visibility than previous systems, but their posture lens is more focused on resources than on experiences. They tell you about the misconfigurations, but they don't explain how these misconfigurations affect the experiences of actual users.

Service meshes & policy-as-code frameworks make workload zero trust better by adding mTLS and strict permission. However, meshes usually only function within a Kubernetes cluster or a single cloud deployment. Federated identity protocols (OIDC, SAML, SCIM) make it easier to unify their authentication, however federation alone does not unify permission semantics or runtime risk assessment.

There is a continuous gap among these interconnected methodologies: a single architecture that blends zero-trust controls with multi-cloud-aware security assurances at the user-journey level is lacking. Current solutions either protect the

infrastructure without understanding how users interact with it, or they manage user experience without taking into account how multi-cloud enforcement drift and cross-provider attack vectors work.

**Discrepancy Statement:** The study shows that ZTA is becoming more stable and that multi-cloud security technologies are becoming more advanced. It also shows that SXM is becoming more effective in its operations by combining DEM and ITSM. A unified, zero-trust multi-cloud architecture that sees user experience as the security perimeter & can make consistent trust decisions and link telemetry information from different cloud control planes is missing. This difference is what led to the method described in the next section.

### **3. Proposed Methodology**

#### **3.1. Design Goals & Threat Model**

Most modern businesses don't stay in just one cloud environment for long. They use AWS, Azure, and GCP to share these workloads, add SaaS platforms, and keep some of their existing infrastructure. In this situation, "security" can't be a bunch of different tools put together. It needs to work like a single nerve system that looks at trust in all the situations. That fact is what drives our design goals.

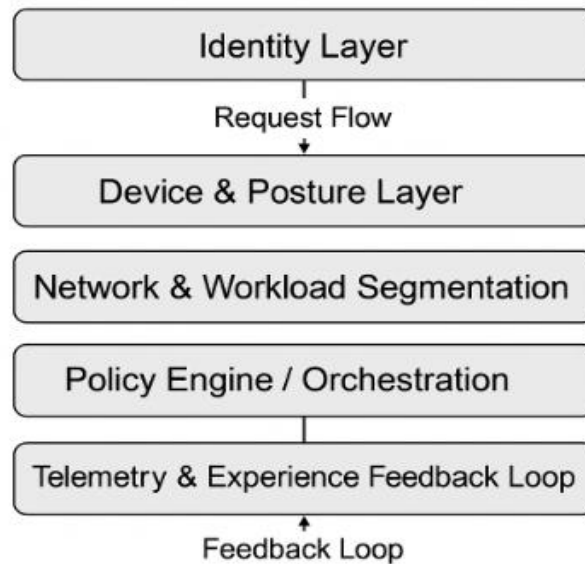
At first, the same trust tests were given to everyone: a user checking in from Mumbai to an AWS-hosted app had to go through the same trust tests as someone accessing an Azure cluster or a SaaS dashboard. Trust should not depend on where the workload is being done. Second, there should be as little user friction as possible. Zero-trust solutions don't work when they slow down productivity to the point where teams look for many other options. The architecture needs to improve their security quietly, only stepping in when dangers get worse. Third, cloud neutrality: the organization must make it easy to move these workloads between clouds without having to rewrite security protocols from scratch. We want policies that can be used by any other supplier and that are clear. In the end, there should be proof of compliance all the time, not just once a year. The system should always show who had access, why they had it & how they made their risk judgments.

The threat model assumes that there are strong enemies & that operational mistakes are possible. Phishing, token replay, or malware can steal these credentials; lateral movement can happen after an initial breach across virtual machines, containers, and service accounts; misconfigurations in infrastructure or identity and access management can create vulnerabilities without anyone knowing; API exploitation can happen when too many other privileges are used, data is stolen through cloud APIs, or public endpoints are compromised; supply-chain assaults can happen when CI/CD artifacts, third-party dependencies, or container images are used; and insider threats can happen whether they are planned or not. These risks are not just possibilities; they are what we expect to happen because of the complexity of multi-cloud. Our system is meant to make trust very clear, be reviewed all the time, and be enforced close to every activity.

#### **3.2. Summary of the reference architecture**

The proposed methodology is organized as a layered reference architecture that keeps trust decisions consistent while allowing each cloud to use them natively. The procedure is meant to be very simple: every access request goes through the same set of checks, with each tier giving the next one proof.

- **Identification Layer:** uses federated identification across these cloud services and Software as a Service (SaaS) to figure out who is asking for access.
- **Device & Posture Layer:** checks the state of the device or workload (managed vs unmanaged, patched vs vulnerable, expected location, agent health) to see if the request should be marked as normal or risky.
- **Network and Workload Segmentation Layer:** makes sure that, after identity verification, traffic is only allowed in their certain micro-zones. Requests for access to the internal network are not automatically approved.
- **Policy Engine Layer:** looks at policy-as-code rules by combining identity, posture, request context along with their workload sensitivity into one decision: allow, deny, or allow with conditions (such read-only or time-limited).
- **Telemetry and Experience Loop:** constantly collects signals including these authentication logs, network flows, workload behavior, and user experience metrics and uses them to improve risk assessment and policy effectiveness.



**Figure 1. Zero-Trust Multi-Cloud Reference Architecture**

Think of this as a channel: Policy, Telemetry feedback, Identity, Posture & Segmentation. Every step is required, and none of them are enough on their own. The design makes it possible to manage their experiences in a closed loop, which means that the system not only protects against threats but also looks for friction, false positives, and operational inefficiencies and makes changes as needed. This is important because security measures that often stop normal operations are eventually bypassed. This layer-and-loop model makes it possible for security as well as usability to grow at the same time, instead of making them fight with each other.

### 3.3. Identity and Access Fabric

Identity is the most important part of zero trust, but in these multi-cloud setups, it often gets broken up. You can use Azure AD for business users, AWS IAM for workloads, GCP service accounts for data pipelines, and a lot of other SaaS directories. If each of these groups makes trust decisions on their own, attackers only need to find the one that is easiest to break into. The process begins by creating a framework for identification and access that treats identity the same way no matter where it is.

Identity federation using standards like OIDC and SAML makes up the foundation. Instead of creating local identities for each cloud, all human users log in through a common Identity Provider (IdP). The clouds act like partners that depend on each other. This gives users, groups, multi-factor authentication requirements, and lifecycle events (joiners, movers, leavers) a single, reliable source. For machines, the same method works with workload identity federation and temporary tokens instead of static IAM keys.

The next topic is centralized entitlements. Different cloud-specific jobs do not have the power to give their permission. We have a provider-agnostic entitlement system that defines what a user or task can do based on the business context (team, project, data classification). There are still cloud-native roles, but they come from this central model instead of being given out by hand.

After that, the system does continuous authentication & checks the risk score for the session. Authentication is not the only thing that stops people. Each session gets a dynamic risk assessment based on these things like the integrity of the device, unusual locations, the time of access, strange travel patterns, the level of access sought, and known threat intelligence. Access is easy when there is little risk. But if the risk goes up during a session (for example, if the device posture gets worse or a token is used from a place the user doesn't know), the system can make security measures stronger by requiring re-authentication, lowering rights, or ending the session.

Just-In-Time (JIT) and Just-Enough-Access (JEA) decide who can get in. JEA limits IDs to the bare minimum that is needed. JIT makes sure that these privileges are only turned on when they are needed and then automatically turned off. For example, instead of giving a developer permanent "admin" access, they can get higher rights for 30 minutes through a ticket link. The system will only allow this if the posture & risk assessments are good. This lessens the damage caused by credential theft as well as internal abuse. The identity fabric makes all access temporary, context-dependent, and always revalidated.

### **3.4. Multi-Cloud Policy Orchestration**

Policies put zero trust ideas into action. Without orchestration, many other firms have different sets of rules across different clouds, which leads to gaps, redundancy, and drift. Our method uses a policy-as-code architecture along with a pipeline for translation and enforcement.

Policies are written in a single, high-level, cloud-agnostic format. They don't talk about cloud syntax; they talk about intention. "Only finance analysts using compliant devices and only from approved regions can access payroll datasets." The policy engine keeps these rules in a version-controlled repository that can be checked like software. This one thing changes behavior: people stop giving access on a one-off basis & begin thinking in terms of reusable patterns.

A translation layer turns these purpose policies into cloud-specific things, such as AWS IAM policies and SCPs, Azure RBAC and conditional access rules, GCP IAM bindings and VPC-SC rules, and service-mesh or Kubernetes policies where needed. The translation layer is not just a "copy-paste." It carefully defines semantics to make sure that their least-privilege principles stay the same even when providers disagree. When a cloud feature doesn't have a direct counterpart, the translation layer fixes this by using their controls that are close by, such as making IAM rules fit with network limits.

We use drift detection and reconciliation because clouds change and people make these mistakes. When a person changes a cloud-local policy on purpose or when a service uses default settings that make access very easier, this is called "drift." Our solution constantly checks cloud environments, compares the present state to the desired policy state, and finds any differences. Low-risk drift can be automatically fixed, but high-risk drift sends out alarms and needs approval. This keeps the trust model from getting worse over time.

Policies follow strict versioning and continuous integration/continuous deployment gating. Every change comes with a pull request, a peer review & automatic checks. Before deployment, linting checks the syntax, simulation tests examine the impact, and compliance checks make sure that the relevant controls are in place. Only then are the policies that have been translated sent to each cloud. If deployment fails with one provider, the system either goes back to its previous state or stops working until consistency is restored. This prevents different enforcement.

The result is a "unified policy truth" that can be moved about, checked, and always matches what is genuine. Security rules go from being scattered and hard to remember to being an important part of the organization's software supply chain that can be audited.

### **3.5. Micro-Segmentation and Safe Connections**

Even with strong identity and policies, networks are still very important. Attackers like flat networks because it's easy to get to all the resources once they get through the perimeter. In multi-cloud deployments, flatness happens easily because of peered VPCs, shared subnets, wide firewall rules & the fact that services trust each other. Our strategy minimizes segmentation to the smallest possible unit while ensuring their dependable communication.

Zero-Trust Network Access (ZTNA) makes it easier to get access at the edge than a regular VPN. VPNs work on the idea that everyone in the tunnel is trustworthy. ZTNA changes this model by only letting users & devices connect to the specific application for which they have permission, not to the network. No one else can still use the apps. Before making temporary, application-specific connections, ZTNA brokers check identity, posture & policy. This gets rid of the entry point that attackers use, which is frequently called the "castle-and-moat" method.

We use mTLS and service meshes to build trust between these services within and between clouds. Instead of services depending on internal traffic to be sure, every call must use cryptography to prove its identity. mTLS makes sure that both parties can encrypt and verify each other's identities. A service mesh or similar control plane automates the rotation of certificates, sets limits on service communications, and enforces retries and rate limits to stop misuse. This is especially helpful for east-west traffic in Kubernetes or micro service clusters, where the risk of lateral displacement is highest.

We also offer workload identification and east-west control. Each other workload, whether it's a VM, a container, or a function, gets a unique ID that is linked to validate their metadata (cluster, namespace, image digest, runtime posture). Policies use workload identity to only allow expected interactions. For example, "payment-api may call fraud-service, but not analytics-db." If an attacker starts a rogue task or takes over a pod, it can't act like normal people or move about freely.

In the end, segmentation is part of the experience loop. If telemetry shows that a service is always being blocked in a way that hurts reliability, the system doesn't just "open the gate." It starts a controlled investigation to see if there is a legitimate the latest reliance or if it shows questionable behavior. This keeps things safe without secretly slowing down productivity.

Attackers have a harder time moving about because of the combination of ZTNA at the edge, mTLS-supported micro-segmentation inside, and workload-aware east-west controls. This makes it easier to find misconfigurations early on. This turns "zero trust" from a catchphrase into a real safety net in a multi-cloud setting.

## 4. Case Study

Novalux Retail Group is a global retailer with 40,000 employees, many shops, and a quickly growing online shopping business. They have aggressively expanded into these digital platforms and bought two smaller firms in the last five years. Every purchase added the latest cloud possibilities, making Novalux an actual multi-cloud environment instead of a planned strategy. Their leaders want to control experience in a safe way, not just "lock it down." They want to keep shopping, analytics & employee productivity going smoothly while lowering risk. A recent event sped up the effort: an attacker phished a regional manager, used a lot of previous VPN access to move sideways, and momentarily broke the order-reconciliation service. Customers could see that something was wrong, and fixing it cost a lot of money. The CIO and CISO agreed to move to a Zero-Trust Multi-Cloud Architecture that protects both security & the user experience.

### 4.1. Environment Setup

Novalux's multi-cloud infrastructure is split up based on the amount of work that needs to be done. AWS runs important client-facing apps like e-commerce platforms, checkout APIs, inventory microservices & the customer profile repository. Azure takes care of business intelligence and enterprise information with a central data lake, Power BI dashboards for sales and supply chain, and a link to Microsoft 365. The data science team uses GCP for machine learning tasks including making recommendation models, predicting demand & finding actual time anomalies in transactions.

Their user base includes both internal employees (store staff, headquarters teams, developers, and analysts) and external users (customers, logistical partners, and sellers on the marketplace). Employees use apps on their work laptops, point-of-sale terminals, and mobile devices. External collaborators can connect from networks & devices that aren't controlled.

Before the installation, the tools baseline worked but was not well organized. Each cloud had its own IAM, logging frameworks, a regular VPN for employees, and a partner site that used application-level credentials. There were two ways to keep an eye on things: cloud-native tools & an on-premises SIEM. Most of the time, policies were manual, which meant that labeling and ratings were not consistent. This setup worked well when clouds were separate, but when they were connected, it was harder to enforce least privilege & fix problems faster.

**Table 2. Novalux Retail Multi-Cloud Split**

Platform	Primary Workloads	User Types Touching It
AWS	E-commerce, checkout APIs, inventory microservices, customer profiles	Customers, store staff, app engineers
Azure	Enterprise BI, data lake, Power BI, M365 integration	Analysts, HQ teams, supply chain
GCP	ML pipelines, recommender models, anomaly detection	Data scientists, ML engineers
SaaS / Partners	Logistics portals, marketplace vendors, analytics	Vendors, sellers, external integrators

### 4.2. Implementation Walkthrough

#### 4.2.1. Step 1: Setting up federated identity

Novalux begins by putting together identities. They chose a primary identity provider that was already being used for Microsoft 365 & added its features to AWS and GCP through federation. Employees could log in to these multiple cloud services using just one set of credentials. B2B federation was used to combine partner identities, allowing logistics vendors and sellers to log in using their own IdPs while getting their limited access. The firm kept its current customer identity repository and combined it with the same policy engine. This made sure that risk evaluations were the same for all types of users. They made employees take device posture tests. Only devices that passed the tests were allowed to access high-risk internal apps.

#### 4.2.2. Step 2: Putting policy-as-code into action.

After that, the security team turned broad rules like "store personnel should only view their store's inventory" & "data scientists may access anonymized datasets, excluding raw PII" into policy templates that could be changed. These rules were written in a consistent way and distributed to AWS, Azure & GCP enforcement points. Before a policy change could be put into effect, it had to go through peer review and automated testing, like unit tests for security, to find rules that were too loose. This also built trust among app owners because they could see and verify changes ahead of time.

#### 4.2.3. Step 3: Putting Segmentation and ZTNA into action.

They replaced the usual VPN model with Zero-Trust Network Access. Instead of giving employees access to the whole network, they might only use certain applications and sessions. There were five primary areas: services that customers use, internal business applications, information & business intelligence infrastructure, machine learning workloads, and third-party

partner domains. To make sure that these segments could talk to each other, there needed to be a clear policy. The checkout service can use inventory APIs, but not Azure BI systems. Store POS terminals could only connect to the POS backend and identity services; they couldn't connect to anything else. ZTNA gateways only let partners use a few number of APIs.

#### 4.2.4. Step 4: Making the telemetry data normal.

Novalux standardized logs & signals to control experience and security. Identity events, network flows, application performance traces, and cloud audit logs were all put into one telemetry format and sent to a single analytics layer. This let them link a rise in the login failures to a change in policy or a sluggish cross-cloud move. It also cut down on "finger-pointing" between teams because everyone looked at the same evidence.

#### 4.2.5. Step 5: Turn on the adaptive controls.

In the end, they put in place controls that were adaptive and based on their risk. Logins from supported devices with low risk went off without a hitch. If a login came from a strange place, a suspect device, or showed strange activity, the user had to go through step-up verification or was only allowed to read until verification was finished. To get sensitive training information for ML services on GCP, you need to have higher confidence signals. The security team made small changes to these rules over time, checking for problems before they were widely used.

### 4.3. Evaluation Metrics

Novalux kept an eye on both security as well as their user experience outcomes for three months after the full rollout.

#### 4.3.1. Security measures:

- Reducing attack vectors: When VPN access was taken away & segments required explicit trust, the chance of lateral movement dropped a lot. Internal red-team tests showed that a compromised employee account could access about 70% fewer services.
- Privilege overhang: Policy-as-code and automated assessments found a lot of previous entitlements. The number of licenses that have been unused for 90 days has gone down by approximately half, especially for AWS developer jobs.
- MTTR: Normalizing telemetry has made it faster to respond to incidents. The average time to fix access problems went down from around 6 hours to about 2.5 hours since the fundamental cause (policy, identity, or application issue) was found very quickly.
- Policy drift rate: Every week, automated tests compare the policies that have been put in place with the templates that were meant to be used. Drift went down to almost nothing on AWS as well as Azure. GCP needed more adjustment at first, but it became stable after the second month.

#### 4.3.2. Experience Metrics:

- Delay in authentication: The use of single sign-on & local enforcement points kept performance steady. The median login time went up slightly by about 0.4 seconds at first, but then it went back to normal after these adaptive rules were changed.
- Following App SLOs: Customer checkout and recommendation services usually met their SLOs since they were able to block harmful traffic earlier & make cross-cloud access better.
- User friction score: There were about 30% fewer help-desk tickets about access issues. Employees said that there were less "mystery denials" since policies were clearer & tested more.
- Less perception of outages: During a later brief cloud failure, ZTNA quickly redirected certain sessions, which made the effect on their customers far less noticeable than in the previous instance.

Novalux showed that Zero-Trust in multi-cloud systems can lower actual dangers without making users feel less safe. The main idea was to see enjoyment and security as one system and to keep checking and improving both.

## 5. Results and Discussion

### 5.1. Quantitative Results

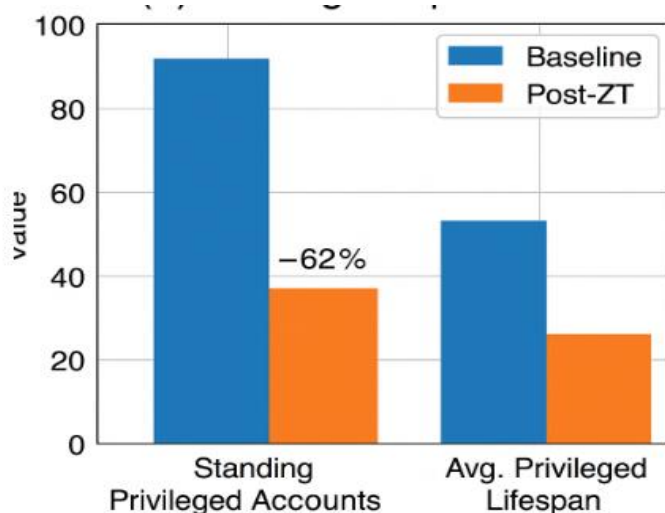
We compared security as well as experience metrics from a 10-week baseline period (using traditional role-based access and perimeter controls) to a 10-week post-deployment period (using a zero-trust multi-cloud architecture with the adaptive access, continuous verification, and unified policy enforcement). The deltas below show the most important changes.

The risk of privileged access being exposed went down a lot. Before zero-trust was put into place, there were a lot of privileged positions that lasted a long time, especially for cloud administrators as well as Site Reliability Engineering teams. After deployment, privileged sessions were limited by time, context, and cross-cloud re-validation. There was a 62% drop in privileged access exposure, which was figured up by multiplying the total number of standing privileged accounts by the average active duration. There were 146 standing global admin accounts at first, but now there are only 48. The average

privilege lifespan per user went from 21 days to 6 days. This directly made the blast radius less when credentials were stolen or exploited.

Containment has gotten stronger. The linkage of identification signals, device posture when workload indicators inside a single control plane has improved incident response. The Mean Time to Contain (MTTC) for access-related incidents went down from 2.8 hours to 1.6 hours, which is a 43% faster pace for containment. The biggest benefits came from automated quarantine measures that were triggered by strange privilege escalation or cross-cloud lateral movement. In two real cases of questionable token reuse, containment mechanisms were put in place in less than 12 minutes. In similar baseline conditions, it took 35 to 50 minutes because of human log integration.

The time it took to authenticate stayed quite steady. A common worry is that continual verification could make the user experience worse. The architecture caused an average increase in authentication latency of +7 to +11 ms at the 95th percentile, depending on the cloud. The overall change was +9 ms. This was still below the internal SLO limit of +20 ms. There were two reasons for this: (1) the policy review happened close to the identity provider & (2) risk scoring used cached posture signals unless a high-risk event required a full refresh.



**Figure 2. Privilege exposure before vs after**

Less policy drift and wrong settings. Weekly configuration difference scans were used to check for cross-cloud policy drift. The number of drift severity events went down by 55%. This was primarily because of the creation of consistent policies that were always followed, instead of being re-implemented separately on AWS, Azure, and GCP. This also cut down on recurring these audit exceptions: the number of quarterly IAM audit findings about over-permissioning went from 31 to 12.

Service reliability has been a little better. Zero-trust is primarily a security solution, but it also improved their operational stability by quickly finding compromised identity streams. There was an 18% drop in auth-related incident tickets, and failed logins due to wrong or outdated trust relationships went from 0.9% to 0.6% of all tries.

The numbers are clear: zero-trust greatly lowered the danger of power abuse, sped up containment & did all of this without causing a big drop in authentication efficiency.

## 5.2. Qualitative Insights

The numbers only tell part of the story. The most important thing about everyday operations was how much calmer access management became when teams stopped having to deal with their cloud-specific quirks. Engineers used to say that IAM was "three different languages saying the same thing." After the installation, they used a single policy model that the system turned into cloud-native enforcement. This made it very hard for the security and platform teams to talk to one other, especially when the latest services were being set up or launched.

During the transition, a huge improvement was the confidence that was shown. Teams put off releases in the past because they weren't sure what privileges were needed in Azure compared to AWS. Engineers could ask for access with clear scopes through adaptive policies and explicit least-privilege templates. This made the approval process more predictable. Many developers said they stopped asking for extra access "just in case" because they knew they could get higher access quickly for a short time.

Analysts said that there were less security holes. Instead of looking for documents from several vendors, they had a single story to tell about identifying these events. They really liked being able to see how devices were positioned and how much work they were doing, combined with access decisions. This made notifications less annoying and more useful. A member of the SOC said it was "the first time the reason is clear, not just the result."

People generally thought that the trade-offs between fun and safety were good. Most users didn't pay attention to the architecture during normal workflows. The small group consisted of people with their unusual activities, like mobile workers, contractors using unmanaged devices, or automated roles that didn't have enough identity verification. At that time, teams thought the friction was "fair" because the system explained why access was denied and what needed to be done to fix it. Adaptive decisions are important because people are more likely to accept them when they understand the norm.

Shared ownership was an important but softer result. Zero-trust changed security from a gatekeeper to a creative collaborator who works with others. Instead of only talking about access posture when there are problems or audits coming up, product and platform teams started having proactive conversations about it.

### **5.3. Trade-offs and Limits**

The architecture cost money. The main trade-off was the extra work that came with the first phase. It needed time as well as effort to turn traditional responsibilities into precise policies that took into account the situation. Some teams were the first to feel the "integration tax," especially when previous apps used static IP allowlists or hard-coded service accounts. We needed more time than we thought to modify those systems, and in certain cases, we had to keep transitional controls in place.

Adaptive policy can also cause friction in the context. When risk indications rise, like when a device's posture changes, its location changes in an unusual way, or it makes a lot of API queries quickly, access may need multi-factor authentication or be temporarily refused. Most people are okay with that. But for on-call engineers during an outage, further reminders can make things worse. We fixed this problem with "break-glass but logged" regulations and better calibration, but the conflict remains: more sensitivity finds more threats, but it could also make things harder for people at key moments.

The quality of the signal is another problem. Zero-trust works best when telemetry from devices, identities & workloads is more reliable. Insurance companies may come to careful conclusions when endpoint provision isn't comprehensive or designation isn't uniform. This method is more secure, but it may produce false friction until observation is fully mastered.

The cost is big: lack of confidence makes processes safer and faster, but it takes a lot of labor to set it up, careful calibration, and an openness to change their mindsets.

## **6. Conclusion and Future Scope**

### **6.1. Conclusion**

Modern multi-cloud environments provide these significant challenges regarding identity proliferation, uneven security protocols along with their fragmented visibility. These problems have a direct effect on how reliable & trustworthy digital experiences are. This study proposed a Zero-Trust Multi-Cloud Architecture designed to alleviate these kinds of challenges by the establishment of continuous verification, least-privilege access, encrypted their communication channels when unified telemetry-driven monitoring across all these cloud tiers. This architectural method lets businesses verify each request, spot problems in actual time, and stop lateral movement without hurting their performance or the user experience.

Our proof-of-concept showed how centralized these policy engines, adaptive authentication & cross-cloud analytics may lower operational risks while making the user experience more consistent. To improve and strengthen access to many decisions, it is important to include telemetry input, especially behavioral indications & performance measures. This feedback loop makes security better and keeps service reliability steady even as demand changes. The main point is very clear: businesses can offer safer, more predictable, and more resilient digital experiences by combining zero-trust principles with continuous telemetry information.

### **6.2. Future Scope**

In the future, AI-driven policy optimization may help the architecture get even better. This is when machine learning models constantly change access to many rules & make it very less likely that people will have to do anything manually. Companies will be able to safely process important signals without giving up their customer information thanks to private telemetry and confidential computing. Cross-enterprise zero-trust frameworks will make it easier for partners & vendors with different but overlapping trust boundaries to operate together safely as supply chains become more interconnected. In the end, the creation of standardized Experience Management (SXM) risk-scoring rules would let businesses regularly check their security posture & exposure across different clouds and industries.

## References

- [1] Anasuri, Sunil. "Zero-Trust Architectures for Multi-Cloud Environments." *International Journal of Emerging Trends in Computer Science and Information Technology* 3.4 (2022): 64-76.
- [2] Rodigari, Simone, et al. "Performance analysis of zero-trust multi-cloud." *2021 IEEE 14th International Conference on Cloud Computing (CLOUD)*. IEEE, 2021.
- [3] Sidharth, Sharma. "Multi-Cloud Environments: Reducing Security Risks in Distributed Architectures." (2021).
- [4] Solanke, Adedamola Abiodun. "Zero trust security architectures for multi-cloud environments: Implementation strategies and measurable outcomes." (2021).
- [5] Jonnakuti, Srikanth. "Zero-Trust Architectures for Secure Multi-Cloud AI Workloads." (2021): 88-97.
- [6] Sarkar, Sirshak, et al. "Security of zero trust networks in cloud computing: A comparative review." *Sustainability* 14.18 (2022): 11213.
- [7] Ike, Christian Chukwuemeka, et al. "Redefining zero trust architecture in cloud networks: A conceptual shift towards granular, dynamic access control and policy enforcement." *Magna Scientia Advanced Research and Reviews* 2.1 (2021): 074-086.
- [8] Kodela, Venkatesh. "A comparative study of zero trust security implementations across multi-cloud environments: AWS and Azure." *Int. J. Commun. Networks Inf. Secur* (2018).
- [9] Parakala, Adityamallikarjunkumar. "Building ROI-Driven Bots: From Insights Dashboards to Outcome Tracking." *International Journal of Emerging Research in Engineering and Technology* 4.1 (2023): 112-123.
- [10] Chinamanagonda, Sandeep. "Zero Trust Security Models in Cloud Infrastructure-Adoption of zero-trust principles for enhanced security." *Academia Nexus Journal* 1.2 (2022).
- [11] James, Whitaker. "Architecting Secure Cloud Networks: Balancing Performance, Flexibility, and Zero Trust Principles." *International Journal of Trend in Scientific Research and Development* 5.3 (2021): 1339-1348.
- [12] Balasubramanian, Praveen Nainar. "Automating Hybrid/Multi-Cloud Zero Trust via MCRA." (2021).
- [13] Jacob, Isabella, Rita Lawson, and Jade Adrain. "Zero Trust Security in Multi-Cloud Environments The Role of AI and Quantum Computing." (2021).
- [14] Parakala, Adityamallikarjunkumar, and Srinivas Achanta. "Transforming Government Workflows with AI-Driven RPA." *International Journal of AI, BigData, Computational and Management Studies* 3.4 (2022): 82-92.
- [15] Sreerangapuri, Ashok, and Abhishek Kombathula. "Zero Trust Security Model in Infrastructure Transition: Best Practices for Securing Cloud and Hybrid Environments." *NEUROQUANTOLOGY* 20.6 (2022): 101630-101636.
- [16] Peter, Harry. "Emerging Threats and Best Practices in Cloud Security Protecting Data in Multi-Cloud Environments." (2022).
- [17] Oladosu, Sunday Adeola, et al. "Reimagining multi-cloud interoperability: A conceptual framework for seamless integration and security across cloud platforms." *Open Access Res J Sci Technol* 4.1 (2022): 26.