Implementing Dynamic Confidential Computing for Continuous Cloud Security Posture Monitoring to Develop a Zero Trust-Based Threat Mitigation Model

Olumide Bashiru Abiola[1]; Matthew Onuh Ijiga[2]

[1]Department of Engineering and Computer Science, Colorado Technical University,
Colorado Springs, Colorado, USA
[2]Department of Physics, Joseph Sarwaan Tarkaa University, Makurdi, Benue State. Nigeria

Abstract: The growing complexity of cloud infrastructures and the increasing sophistication of cyber threats necessitate a paradigm shift in cloud security architecture. This review explores the integration of dynamic confidential computing with continuous cloud security posture monitoring (CCSPM) to develop a Zero Trust-based threat mitigation model. Confidential computing, through trusted execution environments (TEEs), ensures data protection during processing, addressing critical gaps in data-in-use security. When combined with CCSPM tools, which provide real-time visibility and risk assessment, organizations can achieve adaptive and proactive defense mechanisms. This paper examines the fundamental principles of confidential computing, the operational mechanisms of CCSPM, and the implementation of Zero Trust frameworks across distributed cloud environments. It further proposes a dynamic model that fuses telemetry from posture monitoring with policy-based access control to enforce continuous verification and threat response. The synergistic approach promises enhanced data integrity, reduced attack surfaces, and scalable threat resilience. Finally, the paper outlines current limitations, standardization challenges, and research opportunities for advancing secure and trustworthy cloud ecosystems.

Keywords: *Confidential Computing, Zero Trust Architecture, Cloud Security Posture Monitoring (CSPM), Trusted Execution Environments (TEEs), Threat Mitigation, Cloud Security Framework.*

How to Cite: Olumide Bashiru Abiola; Matthew Onuh Ijiga (2025). Implementing Dynamic Confidential Computing for Continuous Cloud Security Posture Monitoring to Develop a Zero Trust-Based Threat Mitigation Model. *International Journal of Innovative Science and Research Technology*, 10(5), 69-83. https://doi.org/10.38124/ijisrt/25may587

I. INTRODUCTION

➤ *Background and Significance of Cloud Security Posture Management (CSPM)*

The rapid migration to cloud-native architectures has significantly increased the complexity of managing cybersecurity risks across distributed infrastructures. As enterprises adopt multi-cloud and hybrid-cloud strategies, traditional perimeter-based security models have become inadequate, giving rise to more sophisticated approaches such as Cloud Security Posture Management (CSPM). CSPM systems continuously assess cloud configurations, detect vulnerabilities, and enforce compliance with established security baselines to reduce the attack surface (Enemosah, A., 2024). The significance of CSPM lies in its capability to address configuration drift, mismanaged entitlements, and the lack of visibility across dynamic cloud workloads.

Despite its strengths, conventional CSPM still struggles with securing sensitive workloads due to their exposure during processing, making data vulnerable even in encrypted storage and transit. This is where confidential computing becomes a strategic enhancement. Confidential computing introduces hardware-based Trusted Execution Environments (TEEs) that protect data-in-use by enabling computation on encrypted data, offering a new dimension of security for cloud infrastructures (Choo, K. 2010). By integrating TEEs into CSPM systems, organizations can maintain persistent security observability while preserving data confidentiality, which is a critical requirement in regulated industries such as healthcare, finance, and defense.

Moreover, as threat actors increasingly target cloud orchestration tools, container runtimes, and infrastructure-as-code pipelines, the convergence of dynamic confidential computing and CSPM emerges as a timely evolution. Confidential computing empowers organizations to securely

execute threat analytics, behavioral modeling, and risk scoring within isolated environments, making it a powerful enabler of Zero Trust security frameworks. Consequently, CSPM is no longer just a compliance tool but a dynamic control plane for proactive threat mitigation. The strategic integration of confidential computing into CSPM platforms signals a shift toward data-centric security operations where continuous monitoring and computational privacy form the backbone of Zero Trust enforcement.

➢ *Rise of Confidential Computing in Cloud Environments*

As cloud infrastructures evolve to support distributed workloads and sensitive data exchange, the need for hardware-rooted, runtime data protection has accelerated the adoption of confidential computing technologies. Confidential computing enables data to remain encrypted not only at rest and in transit but also during processing, within isolated hardware-based Trusted Execution Environments (TEEs), significantly reducing the attack surface for cloud-native operations (Wang et al., 2020). These environments operate independently from host operating systems and hypervisors, ensuring that even cloud administrators cannot access the data or code running inside them.

This paradigm shift is particularly relevant in the context of continuous cloud security posture monitoring (CCSPM), where telemetry and threat intelligence must be collected, analyzed, and responded to in near real time. Integrating confidential computing into this workflow ensures that the collection of telemetry does not itself introduce vulnerabilities or expose sensitive operational logic. Furthermore, it aligns with Zero Trust principles by enforcing strict data access control and verifying the integrity of both applications and identities before execution (Mavroeidis & Bromander, 2021).

With the proliferation of advanced persistent threats (APTs) targeting cloud-native applications, confidential computing acts as a foundational enabler for dynamic threat mitigation, allowing cloud systems to maintain resilience while processing sensitive workloads securely. Its emergence is redefining security baselines across sectors, especially in multi-tenant architectures where assurance of runtime confidentiality is paramount to preserving trust and compliance.

➢ *Challenges in Traditional Cloud Security Models*

Traditional cloud security models face inherent limitations in addressing evolving cyber threats, primarily due to their reliance on perimeter-based defenses, static policy enforcement, and centralized trust anchors. These models struggle to provide granular visibility across distributed workloads, especially in multi-cloud and hybrid environments where the attack surface is constantly shifting. A major concern is the lack of real-time, context-aware monitoring, which exposes cloud infrastructures to misconfiguration vulnerabilities, insider threats, and sophisticated lateral movement by adversaries (Spanaki et al., 2019).

Furthermore, the inability to isolate sensitive workloads from the underlying infrastructure limits assurance in data confidentiality and integrity. While encryption at rest and in transit is now standard, data in use remains largely unprotected under traditional architectures. This gap has significant implications for regulatory compliance and secure multi-tenancy, particularly in environments processing sensitive personal or industrial information (Sadeghi et al., 2015). Legacy models also suffer from inconsistent identity and access management controls, complicating the implementation of dynamic, risk-based policies.

The siloed nature of traditional systems impedes unified threat detection and incident response, creating blind spots that attackers exploit. These limitations underscore the urgent need to adopt security paradigms that enforce continuous verification, minimize implicit trust, and dynamically secure data throughout its lifecycle—principles foundational to both confidential computing and Zero Trust Architecture.

➢ *Research Objectives and Scope of the Review*

This review aims to critically examine how dynamic confidential computing can reinforce continuous cloud security posture monitoring (CCSPM) within a Zero Trust Architecture (ZTA) framework. With increasing data breaches and lateral threat movements in multi-cloud environments, it is imperative to explore security models that are resilient, context-aware, and privacy-preserving. The primary objective is to assess how trusted execution environments and adaptive cryptographic mechanisms can dynamically isolate sensitive workloads and enforce granular access policies in real time. Furthermore, the review investigates the role of CCSPM tools in facilitating early detection, autonomous threat mitigation, and risk scoring based on continuous telemetry, thereby enabling the implementation of Zero Trust principles across hybrid architectures. Emphasis is placed on understanding how confidential computing capabilities can operationalize ZTA by integrating secure enclaves into automated threat response loops. The scope also includes identifying existing architectural gaps, performance trade-offs, and compliance risks that hinder effective deployment of such systems in enterprise-scale clouds. By aligning dynamic workload protection with identity-centric access control, this paper contributes to the emerging discourse on cloud-native security resilience.

➢ *Organization of the Paper*

This paper is organized into seven sections to provide a comprehensive exploration of the integration of dynamic confidential computing with continuous cloud security posture monitoring (CCSPM) to develop a Zero Trust-based threat mitigation model. The introduction sets the stage by highlighting the challenges in cloud security and introducing CSPM as a crucial tool for addressing these issues. The conceptual foundations section discusses key concepts, including confidential computing and Zero Trust architecture, and their relevance to cloud security. Dynamic confidential computing models are explored in the third section, focusing on technologies such as Intel SGX and Secure Multi-Party Computation (SMPC). Section four delves into CCSPM,

examining its evolution and role in continuous monitoring and threat remediation. The fifth section outlines the application of Zero Trust principles, enhanced by confidential computing, in building a secure cloud environment. The sixth section addresses the implementation challenges, scalability issues, and future trends, such as AI-driven threat detection and federated Zero Trust models. Finally, the conclusion summarizes the findings and provides recommendations for further research and policy development, emphasizing the transformative potential of integrating confidential computing into CSPM systems for robust cloud security.

II. CONCEPTUAL FOUNDATIONS

➤ *Overview of Confidential Computing: Trusted Execution Environments (TEEs)*

Confidential computing has emerged as a pivotal innovation for enhancing cloud security by enabling computation on encrypted data within hardware-based Trusted Execution Environments (TEEs) as seen in Figure 1. These TEEs are isolated environments within a processor that protect code and data from external access, even by privileged system software or cloud providers. Central to this capability is Intel's Software Guard Extensions (SGX), which provide secure enclaves that restrict access to sensitive data during execution, thus minimizing exposure to attack vectors (Costan & Devadas, 2016). Such a design aligns well with the core tenets of zero trust security, particularly the principle of "never trust, always verify," by enforcing strict runtime protection at the hardware level.

Advanced implementations like SCONE (Secure Computing Nodes) demonstrate how TEEs can be leveraged to run containerized applications with SGX, thereby preserving confidentiality in cloud-native architectures without substantial performance trade-offs (Arnautov et al., 2018). This development is particularly relevant for continuous cloud security posture monitoring, where telemetry data may contain sensitive operational or user-specific metadata. By integrating confidential computing into continuous monitoring frameworks, organizations can uphold privacy and regulatory compliance while maintaining real-time visibility. These capabilities significantly reinforce the Zero Trust paradigm by ensuring that even within shared or compromised infrastructure, sensitive computations remain protected. The use of TEEs thus forms a foundational element in building dynamic, resilient, and secure cloud infrastructures under zero trust-based threat mitigation models.
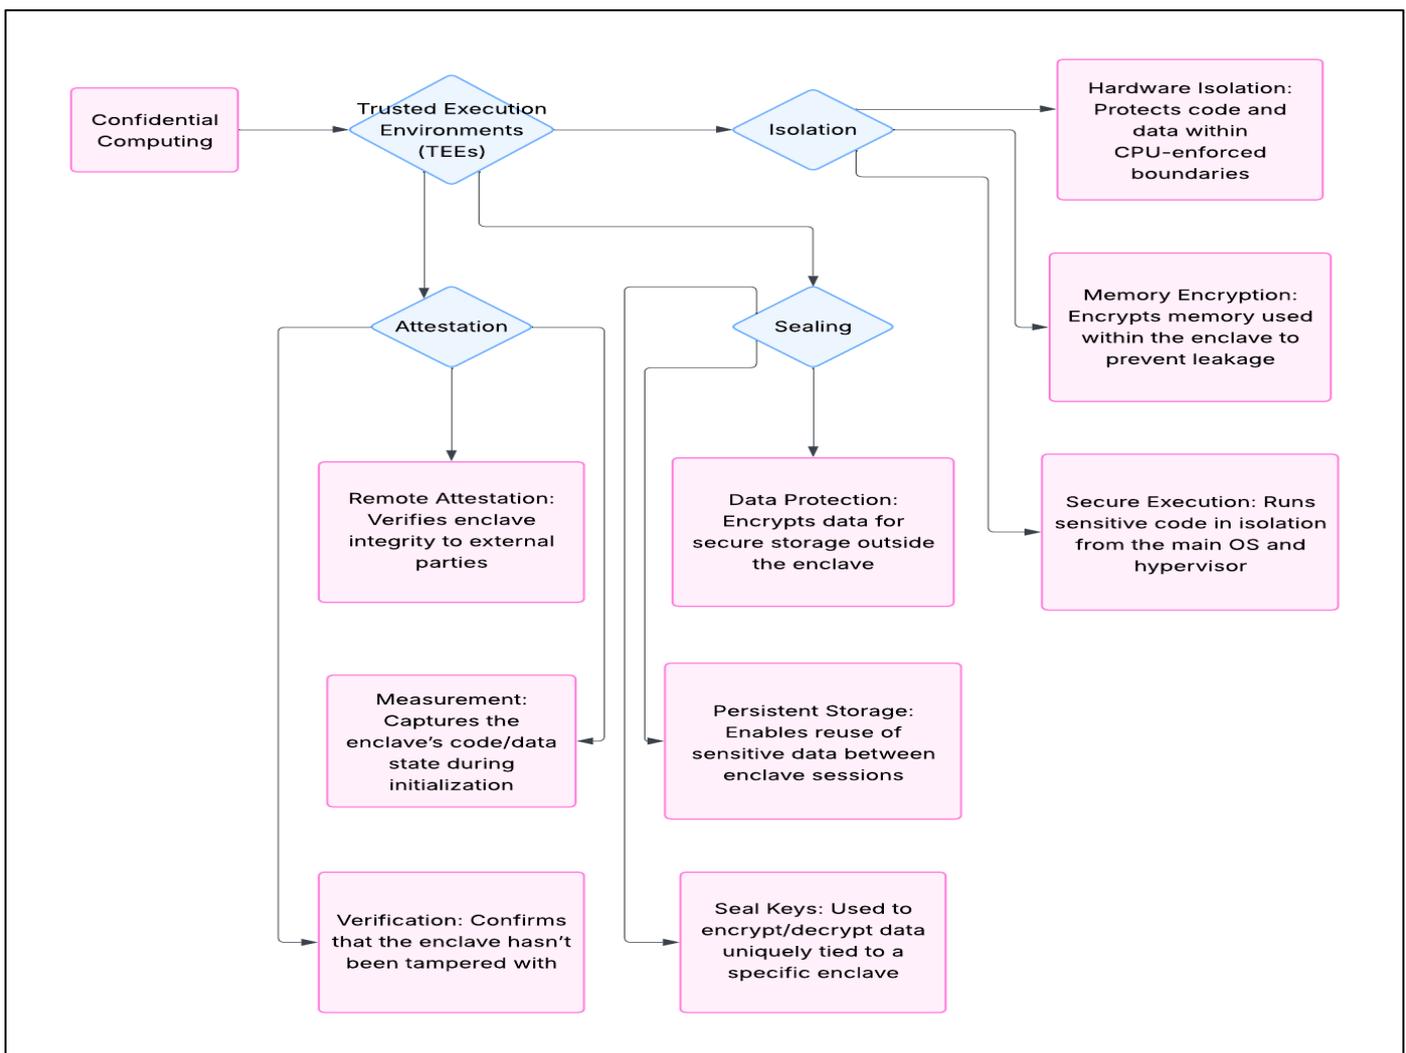

Fig 1 Diagram Showing the Architecture of Confidential Computing via Trusted Execution Environments (TEEs)

Figure 1 represents the architectural framework of Trusted Execution Environments (TEEs) under the broader scope of confidential computing, emphasizing its foundational security capabilities. At the core, TEEs are classified into three functional pillars—Isolation, Attestation, and Sealing—each addressing a unique security objective. Isolation ensures that sensitive data and computations are shielded from external interference through hardware-enforced separation, memory encryption, and secure execution environments. Attestation provides mechanisms to verify the integrity and authenticity of the TEE via remote attestation, measurement, and verification processes, thereby building trust with external systems. Sealing secures data that must persist outside the enclave by encrypting it using enclave-specific seal keys, thus enabling data protection and safe storage continuity. Together, these elements establish a secure foundation for processing confidential workloads in untrusted cloud environments, aligning with modern confidentiality and integrity demands.

➢ *Key Components of Continuous Cloud Security Monitoring*

Continuous Cloud Security Posture Monitoring (CCSPM) has emerged as a fundamental pillar of cloud-native defense mechanisms, enabling real-time visibility into misconfigurations, threats, and compliance drifts in complex hybrid environments. At its core, CCSPM operates through an integrated ecosystem of telemetry, automation, and dynamic assessment tools that provide continuous insight into an organization's cloud security health as seen in Table 1. Unlike static assessments, CCSPM tools ingest log streams, audit data, and network flows to evaluate cloud configurations against evolving security baselines and regulatory requirements (Torkura et al., 2021). This persistent assessment model enhances threat detection accuracy and fosters proactive remediation before adversaries can exploit exposed assets.

A robust CCSPM framework typically encompasses identity and access misconfiguration monitoring, encryption policy validation, behavioral analytics, and automated response orchestration. These components collectively support the implementation of zero trust security principles by ensuring that all access and behaviors within the cloud are verified, contextualized, and risk-scored continuously (Alsadie, D. 2024). Furthermore, the integration of artificial intelligence and machine learning into CCSPM platforms allows for adaptive learning and predictive anomaly detection, significantly improving the system's ability to detect emerging threats. As cloud infrastructures scale and diversify, these intelligent and dynamic monitoring systems are essential to maintaining compliance, enforcing security policies, and ultimately building a resilient zero trust-based security model.

Table 1 Key Components of Continuous Cloud Security Monitoring (CCSM)

Component	Description	Function	Examples/Tools
Telemetry Collection	Captures real-time and historical data from cloud workloads and infrastructure.	Enables threat detection, baselining, and behavioral analysis.	AWS CloudTrail, Azure Monitor, GCP Operations Suite
Security Configuration Assessment	Evaluates cloud services against security best practices and compliance policies.	Identifies misconfigurations and policy violations.	Prisma Cloud, AWS Config, Microsoft Defender CSPM
Threat Detection Engine	Uses rules and behavioral analytics to detect anomalies and potential attacks.	Alerts security teams to suspicious or malicious activity.	Amazon GuardDuty, Azure Sentinel, Splunk
Automated Remediation	Applies predefined actions to mitigate or isolate detected threats.	Reduces response time and limits attack impact.	SOAR platforms, Lambda functions, Terraform

➢ *Zero Trust Architecture (ZTA): Principles and Pillars*

Zero Trust Architecture (ZTA) has emerged as a critical paradigm for securing modern cloud environments by eliminating implicit trust and continuously validating every access request, irrespective of origin. The foundational principles of ZTA—least privilege, continuous authentication, and segmentation—are now being redefined in conjunction with confidential computing to ensure secure computation even in potentially compromised cloud infrastructures. This integration introduces a new layer of hardware-based isolation that strengthens ZTA implementations by mitigating insider threats and external breaches, especially in multi-tenant and distributed environments (Confidential Computing Consortium. 2020).

Confidential computing, enabled through Trusted Execution Environments (TEEs), aligns closely with Zero Trust by enforcing runtime encryption and limiting data visibility to only authorized and verified execution processes.

This ensures that sensitive operations remain protected even if the underlying host is untrusted, further reinforcing the "never trust, always verify" mantra of ZTA (Ahmadi, S. 2024). Additionally, these technologies support dynamic policy enforcement and real-time integrity checks, enhancing the fidelity of identity, device, and workload assessments integral to Zero Trust strategies. In cloud ecosystems where lateral movement by threat actors remains a pressing concern, confidential computing provides critical runtime assurances that bolster Zero Trust microsegmentation practices. Consequently, the convergence of these two models—confidential computing and ZTA—offers a promising trajectory toward achieving resilient, scalable, and secure cloud infrastructures.

➢ *Synergies Between Confidential Computing and Zero Trust*

The convergence of confidential computing and Zero Trust Architecture (ZTA) marks a transformative shift in securing cloud environments, particularly in dynamic and adversarial threat landscapes. Confidential computing enables the encryption of data in use through trusted execution environments (TEEs), offering hardware-level isolation for sensitive workloads. When integrated with ZTA, which enforces continuous verification of users, devices, and workloads regardless of network location, this synergy fortifies cloud systems against lateral movement, insider threats, and data exfiltration (Sulochana et al., 2021). In traditional cloud deployments, trust boundaries often rely on static perimeters, which are increasingly ineffective in multi-tenant and distributed systems. Confidential computing overcomes this by ensuring that data remains protected even during computation, aligning with the ZTA principle of "never trust, always verify."

Moreover, Zero Trust models benefit from the attestation and verification mechanisms embedded within TEEs, enabling adaptive access decisions based on real-time workload integrity (Wu et al., 2020). This creates a feedback loop where security posture assessments feed into the policy enforcement engine, enhancing resilience. The integration also supports secure policy orchestration in cloud-native applications by isolating microservices and enforcing least privilege at the granular level. This dynamic alignment offers a robust framework for continuous cloud security posture monitoring, mitigating sophisticated threats while preserving data confidentiality. The synergy is especially critical in regulated industries and sensitive operations, where trust minimization and verifiable computation are foundational to operational integrity.

III. DYNAMIC CONFIDENTIAL COMPUTING MODELS

➢ *Hardware-Enforced Confidentiality: Intel SGX, AMD SEV, and Arm TrustZone*

Hardware-enforced confidentiality represents a pivotal mechanism in enabling secure execution within untrusted cloud environments. Intel Software Guard Extensions (SGX), AMD Secure Encrypted Virtualization (SEV), and Arm TrustZone are foundational technologies designed to isolate sensitive workloads using Trusted Execution Environments (TEEs), thereby supporting the broader objectives of confidential computing. These TEEs create isolated memory regions, shielding computation and data from both cloud providers and potential malicious insiders, which is vital for the implementation of Zero Trust architectures (Costan & Devadas, 2016).

Intel SGX offers enclave-based computation, facilitating minimal Trusted Computing Bases (TCBs) and remote attestation, enabling verifiable execution of code in encrypted memory. This capability becomes instrumental in continuous cloud security posture monitoring (CCSPM), where runtime integrity and confidentiality must be maintained across distributed systems. AMD SEV, on the other hand, encrypts entire virtual machines with a per-VM key without requiring software modification, thus offering scalable protection for dynamic multi-tenant environments.

Meanwhile, Arm TrustZone as shown in Table 2, enables system-on-chip partitioning, separating secure and non-secure execution states. Its application in edge-to-cloud workflows enhances device-level attestation, which complements threat mitigation in Zero Trust-based models. As modern cloud environments increasingly adopt decentralized and hybrid configurations, these hardware features empower secure telemetry collection, access governance, and encrypted policy enforcement (Azab et al., 2016). Consequently, hardware-enforced confidentiality is not only a technical prerequisite but also a strategic enabler for implementing resilient, adaptive, and verifiable Zero Trust models in confidential cloud computing infrastructures.

Table 2 Summary of Hardware-Enforced Confidentiality Technologies in Cloud Security

Technology	Key Features	Security Mechanism	Use Cases in Cloud Security
Intel SGX	Enables creation of secure enclaves for code and data	Memory encryption and enclave-based isolation	Protecting sensitive computations, secure multi-party analytics
AMD SEV	Encrypts entire virtual machine memory with per-VM keys	Hardware-based full memory encryption	Isolating tenant workloads in virtualized cloud environments
Arm TrustZone	Creates a secure world alongside a normal execution environment	Trusted execution environment with separate privilege levels	Secure boot, cryptographic key management, and trusted mobile services

➢ *Secure Multi-Party Computation (SMPC) and Homomorphic Encryption*

Secure Multi-Party Computation (SMPC) and homomorphic encryption are two foundational cryptographic techniques essential for extending the principles of confidential computing within dynamic and multi-tenant cloud environments. These techniques empower cloud providers to process encrypted data without accessing its plaintext form, making them pivotal in environments where continuous monitoring and strict data privacy co-exist. SMPC enables collaborative computation between distrusting parties while ensuring that none of them can infer the private data of others—supporting Zero Trust principles by design (Sardar et al., 2023). Within the scope of continuous cloud security posture monitoring, SMPC can facilitate distributed anomaly

detection, secure threat correlation, and compliance validation without exposing sensitive telemetry.

Homomorphic encryption further advances this paradigm by allowing encrypted datasets to be computed over without decryption, aligning well with hardware-based Trusted Execution Environments (TEEs). By applying homomorphic models to runtime security metrics, organizations can audit and enforce compliance transparently, even in outsourced or federated infrastructures (Elrabaa, et al., 2019). These tools reinforce policy enforcement engines by enabling risk-adaptive responses that preserve confidentiality and integrity simultaneously. As threat surfaces evolve, the integration of SMPC and homomorphic encryption into confidential computing enhances not only trust boundaries but also operational resilience. Thus, these techniques form a critical layer of the proposed dynamic threat mitigation model, embedding security deeply into data lifecycle processes while ensuring privacy-preserving cloud operations.

➤ *Adaptive Confidentiality in Multi-Tenant Cloud*
Adaptive confidentiality in multi-tenant cloud environments addresses the complex challenge of isolating sensitive workloads while maintaining performance and scalability across dynamically shifting cloud resources. This dynamic isolation is essential for supporting secure, continuous posture monitoring across diverse workloads without sacrificing the integrity or privacy of any individual tenant's data. In multi-tenant cloud platforms, where resources are pooled and shared among users, static security boundaries often fail to prevent lateral threat propagation. Adaptive trusted execution environments (TEEs) play a vital role in this context by enabling real-time adjustment of security parameters based on contextual risk assessments and policy enforcement mechanisms (Pires, R. P. 2020). These TEEs facilitate confidential computing paradigms that can dynamically scale and respond to threat intelligence signals, thereby enabling selective isolation of risky workloads and enforcing zero trust segmentation at the hardware level.

Hashim et al. (2024) highlight that traditional TEEs fall short when deployed in multi-tenant settings, primarily due to their static nature and inability to manage complex access control across users with varying privileges. The incorporation of dynamic TEEs resolves these limitations by allowing workload migration, encrypted computation, and real-time auditability without disrupting operational continuity. Consequently, multi-tenant cloud infrastructures benefit from enhanced resilience, where trust is no longer assumed but continuously verified. This implementation aligns with the broader zero trust-based threat mitigation model, supporting continuous risk evaluation and access enforcement while preserving user data confidentiality in shared execution environments.

➤ *Case Studies and Industry Implementations*
Recent implementations of confidential computing within enterprise cloud systems demonstrate its potential to enforce secure enclaves for continuous security posture monitoring. For instance, cloud providers like Microsoft

Azure and Google Cloud have incorporated Trusted Execution Environments (TEEs) such as Intel Software Guard Extensions (SGX) into their infrastructure to ensure that sensitive data remains confidential even during processing. This hardware-based isolation mechanism supports critical security operations without exposing plaintext data to system administrators or unauthorized applications. (Brenner, et al, 2017) examined an Intel SGX-based deployment scenario where virtual secure mode (VSM) was used to protect customer data during multi-tenant processing. The study highlighted reduced attack surfaces and enhanced resilience in scenarios with untrusted hosts, reinforcing the viability of dynamic confidential computing as a pillar in Zero Trust enforcement.

Furthermore, leading cybersecurity operations have started to integrate confidential computing into continuous cloud security monitoring pipelines. Conti et al. (2020) presented a detailed analysis of SGX deployment in various industrial environments, including fintech and healthtech, where enclave-based analytics helped maintain compliance while supporting real-time telemetry collection. These secure enclaves facilitated the enforcement of continuous risk scoring, adaptive access decisions, and policy automation based on protected analytics engines. The operational feedback loop created by these implementations supports Zero Trust principles by authenticating every workload and verifying security posture dynamically, without reliance on static perimeters.

Together, these cases reveal the emerging standard of integrating confidential computing into production-level systems for secure cloud workload execution. They validate the argument that hardware-enforced isolation and encrypted computation are essential to operationalizing Zero Trust architectures in dynamic, threat-prone cloud ecosystems.

IV. CONTINUOUS CLOUD SECURITY POSTURE MONITORING (CCSPM)

➤ *Definition and Evolution of CCSPM Tools*
Continuous Cloud Security Posture Monitoring (CCSPM) has emerged as a proactive approach to address the dynamic security risks of cloud environments. Unlike static security assessments, CCSPM enables organizations to perform continuous evaluations of their configurations, permissions, and security controls to maintain compliance and mitigate threats in real time. It is built upon principles of visibility, automation, and policy-driven governance that align with the zero trust paradigm. Modern CCSPM tools offer multi-cloud integration, automated remediation, and real-time policy enforcement, bridging the operational gap between development and security teams in cloud-native environments (Lang et al., 2011).

The evolution of CCSPM has been significantly influenced by the shift toward infrastructure-as-code (IaC) and DevSecOps practices, which demand continuous feedback loops for security validation as seen in figure 2. Contemporary tools are not only reactive but also predictive, employing rule-based engines and machine learning to

identify misconfigurations and non-compliant assets across distributed systems. As cloud complexity increases, CCSPM tools now provide advanced telemetry and behavior analytics that support scalable threat detection without violating the confidentiality of user workloads (Alzubaidi et al., 2020). These advancements have expanded the scope of posture monitoring from compliance auditing to real-time risk governance.

In the context of confidential computing, CCSPM tools play a crucial role in enforcing encrypted policy enforcement across trusted execution environments (TEEs), enabling privacy-preserving visibility while maintaining security oversight. This synergy reinforces the zero trust model by ensuring continuous validation of every asset, identity, and interaction across the cloud lifecycle.
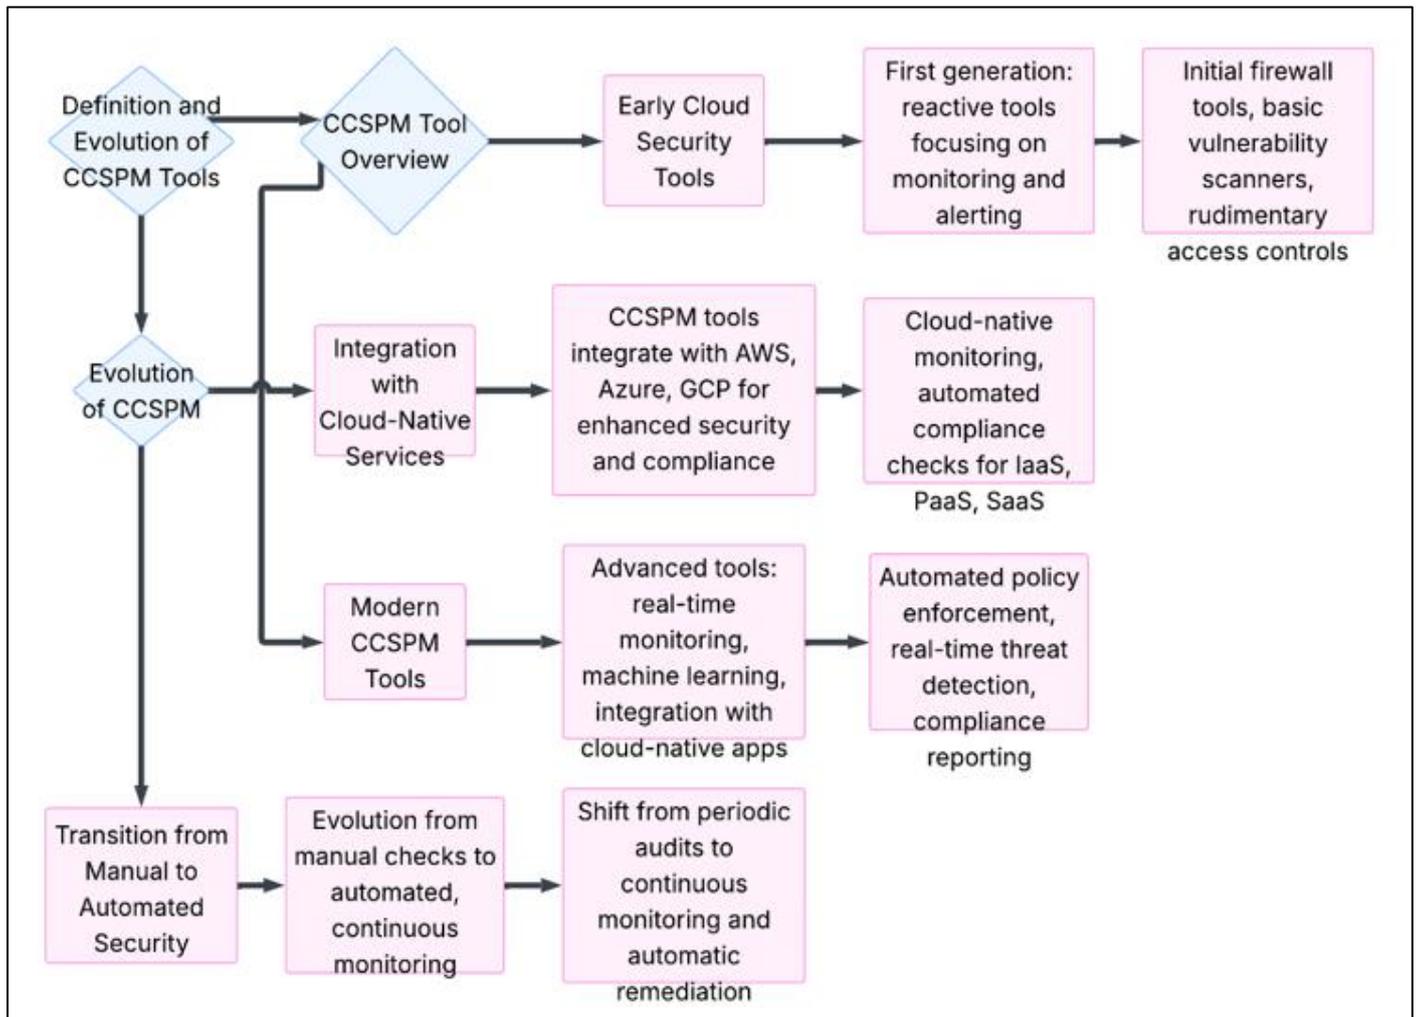

Fig 2 Evolution and Overview of CCSPM Tools

Figure 2 illustrates the development and key components of Cloud Security Posture Management (CCSPM) tools, tracing their evolution from early reactive tools to modern, advanced systems integrated with cloud-native services. The central node, "Definition and Evolution of CCSPM Tools," branches into two main sections: CCSPM Tool Overview and Evolution of CCSPM. The first branch, CCSPM Tool Overview, highlights the early cloud security tools, which were primarily reactive, focusing on monitoring and alerting, and outlines the transition to more advanced systems integrated with cloud services like AWS, Azure, and GCP, enhancing security and compliance. It then shows how modern CCSPM tools evolved, offering real-time monitoring, threat detection, and automated policy enforcement. The second branch, Evolution of CCSPM, shows the progression from manual security checks to automated, continuous monitoring, emphasizing the shift

from periodic audits to dynamic, automatic remediation. This comprehensive layout visually represents the ongoing advancements in CCSPM tools and their increasing integration with cloud-native architectures to enhance security management.

➤ *Telemetry, Logging, and Risk Scoring*

Telemetry and logging are foundational to any dynamic and continuous cloud security posture monitoring framework, particularly when integrated with confidential computing principles. As enterprise workloads scale across hybrid and multi-cloud environments, the volume and diversity of telemetry data—such as access logs, API calls, and system metrics—require sophisticated mechanisms for privacy-preserving collection and contextual interpretation. Dynamic telemetry not only enables real-time visibility into system states but also supports enriched threat analytics through

behavioral baselining and anomaly detection (Kodakandla, N. 2024). However, this process becomes significantly more complex when sensitive data must be analyzed without exposure, underscoring the relevance of confidential computing. Trusted Execution Environments (TEEs) facilitate this secure processing by encrypting telemetry at rest, in transit, and even during computation.

Risk scoring further operationalizes telemetry insights by assigning quantifiable threat levels based on contextual behavior, system vulnerabilities, and trust metrics as shown in Table 3. In Zero Trust architectures, these scores drive automated access decisions and remediation policies. By coupling logging pipelines with confidential computing enclaves, organizations can ensure the integrity and authenticity of event streams, reducing the risk of tampering or misattribution (Shiraz & Gani, 2020). Integrating machine learning within secure telemetry workflows can further enhance the precision of risk modeling, allowing adaptive security controls that align with evolving threat landscapes. This seamless fusion of telemetry, logging, and risk scoring within a confidential framework not only elevates cloud observability but also fortifies continuous trust evaluation in dynamic Zero Trust environments.

Table 3 Key Components and Functions of Telemetry, Logging, and Risk Scoring in Continuous Cloud Security Posture Monitoring

Component	Description	Function in CCSPM	Relevance to Zero Trust Model
Telemetry	Real-time data collection from cloud assets, workloads, and network traffic	Enables continuous visibility, behavioral analytics, and threat detection	Supports dynamic trust decisions and contextual access validation
Logging	Structured records of events, user actions, and system states	Facilitates audit trails, forensic analysis, and compliance reporting	Ensures accountability and traceability of all access and actions
Risk Scoring	Quantitative evaluation of asset, user, or activity risk levels	Prioritizes threats, supports policy enforcement, and triggers remediation	Enables adaptive access control and policy refinement based on behavior
Correlation Engine	Integration layer that aggregates and contextualizes telemetry and logs	Detects patterns, anomalies, and attack indicators across environments	Central to enforcing microsegmentation and dynamic threat response

➤ *Integrating Threat Intelligence with Confidential Data Workloads*

The integration of threat intelligence into confidential computing environments marks a pivotal advancement in continuous cloud security posture monitoring. As cloud infrastructure increasingly supports dynamic, multi-tenant workloads, the ability to securely ingest and analyze threat intelligence within trusted execution environments (TEEs) becomes essential. TEEs enable confidential computing systems to perform sensitive analytics on encrypted threat data without exposing it to the host infrastructure, preserving data confidentiality and operational integrity. This aligns with Zero Trust principles, where no component—internal or external—is inherently trusted. By embedding threat detection models within TEEs, organizations can proactively identify anomalous behaviors and apply risk-weighted mitigation policies in real time.

Moreover, threat intelligence feeds—when combined with telemetry from cloud assets—enhance the contextual relevance of detected threats, especially in hybrid or edge-cloud architectures. Yang, et al, 2023) demonstrated how AI-driven, context-aware threat detection mechanisms can adaptively respond to evolving attack surfaces across distributed systems. This flexibility is especially critical for workloads running in confidential containers or enclaves where operational transparency is limited by design. Similarly, Han et al. (2020) emphasize the role of secure data channels and real-time analytics for threat signal correlation in sensitive healthcare applications, underscoring the growing relevance of confidential data processing.

As confidential computing continues to evolve, embedding threat intelligence directly into encrypted data processing pipelines can amplify both the accuracy and speed of threat response—transforming passive detection into an active, Zero Trust-driven defense mechanism.

➤ *Real-Time Remediation and Policy Enforcement*

Real-time remediation and policy enforcement are critical pillars of continuous cloud security posture monitoring, especially when integrating dynamic confidential computing into Zero Trust architectures. As organizations increasingly adopt cloud-native and containerized infrastructures, enforcing dynamic security policies in real-time ensures resilience against lateral movement and privilege escalation attacks. Confidential computing—through trusted execution environments (TEEs)—allows security controls to be enforced directly within the enclave, eliminating the risk of exposure during monitoring or policy application (Abwnawar, N. 2020). This approach enables runtime inspection of encrypted workloads and autonomous remediation workflows without compromising data confidentiality or operational latency.

Advanced orchestration platforms now support automated responses to risk indicators, such as anomalous access attempts or configuration drifts, by deploying encrypted policy agents that validate and enforce predefined behavioral baselines. This aligns with Zero Trust principles, where policy enforcement is continuous and adaptive to user context, device state, and workload sensitivity (Jim, M., 2024). Moreover, integration with telemetry feeds and policy

engines allows for granular enforcement—automatically restricting access, quarantining workloads, or invoking re-authentication in response to threat signals. These real-time capabilities significantly enhance situational awareness and containment precision within cloud environments. Thus, confidential computing not only protects sensitive processes but also transforms them into autonomous security actors capable of enforcing organizational policy at the data and compute layers in real-time.

V. ZERO TRUST-BASED THREAT MITIGATION MODEL

➢ *Building Blocks: Identity, Devices, Network, and Applications*

Implementing a Zero Trust-based threat mitigation model within cloud infrastructure requires a precise focus on its foundational components: identity, devices, network, and applications. The Zero Trust Architecture (ZTA) eliminates implicit trust by continuously validating users and devices before granting access to cloud resources. Identity becomes the control plane, necessitating strong authentication, dynamic authorization, and behavioral analytics to enforce least privilege access across federated domains (Rose et al., 2020). Device integrity is equally critical, where endpoint posture assessments ensure only secure, policy-compliant

devices interact with sensitive workloads. Confidential computing strengthens this assurance by executing code within secure enclaves, isolating data even from privileged cloud providers.

Network segmentation—an essential Zero Trust pillar—undergoes refinement through microsegmentation and adaptive routing strategies as presented in Table 4. These ensure east-west traffic is verified and encrypted, reducing lateral movement risks. Moreover, dynamic policy enforcement, grounded in real-time telemetry and behavioral baselines, enables rapid anomaly detection and containment within distributed cloud environments (Hubbard et al., 2021). Applications, especially those handling confidential data, are embedded within secure runtime environments, where continuous risk evaluation governs access, execution, and integration with third-party services.

This layered approach—underpinned by identity-centric controls, device health telemetry, network microsegmentation, and confidential computing—constructs a proactive, context-aware threat mitigation model. It aligns seamlessly with continuous security posture monitoring and supports dynamic reconfiguration in response to evolving risk landscapes.

Table 4 Summary of Building Blocks in a Zero Trust-Based Threat Mitigation Model

Building Block	Core Function	Zero Trust Implementation	Confidential Computing Integration
Identity	Verifies user and entity authenticity	Enforces continuous identity verification using multi-factor authentication (MFA), least privilege, and behavioral analytics	Protects identity attributes during authentication and authorization processes using secure enclaves
Devices	Monitors endpoint posture and compliance	Validates device health, ensures endpoint compliance, and manages access based on device trustworthiness	Enables encrypted processing of device telemetry data for risk scoring
Network	Controls traffic flow and segmentation	Applies microsegmentation, encryption-in-transit, and adaptive routing policies	Secures network policies and traffic metadata within isolated trusted execution environments
Applications	Protects workload execution and inter-app communication	Limits access to sensitive application components through role-based controls and just-in-time access	Shields in-memory application logic and data processing from external threats

➢ *Incorporating Confidential Computing into Microsegmentation and Access Control*

Incorporating confidential computing into microsegmentation and access control mechanisms enables granular enforcement of Zero Trust principles in cloud-native architectures. Traditional microsegmentation relies on predefined trust zones, but these are increasingly inadequate due to mutable workloads and dynamic access demands in multi-tenant environments. Confidential computing enhances this model by securing workloads within Trusted Execution Environments (TEEs) as seen in figure 3, thereby isolating sensitive operations from both unauthorized users and compromised host systems. This hardware-based confidentiality reinforces least-privilege principles and eliminates implicit trust, even within the same subnet or virtual machine cluster (Feng et al., 2024).

When integrated with context-aware access control systems, confidential computing dynamically validates both user identity and workload integrity before permitting access, enhancing the decision-making process beyond static credential checks. This layered defense model is vital for defending against lateral movement attacks, especially in hybrid or federated cloud setups. Moreover, embedding TEEs within access policy enforcement points allows organizations to implement runtime verification of workloads, ensuring that malicious code or altered configurations cannot influence access outcomes (Sarkar, et al., 2022). This convergence of secure enclaves with zero trust microsegmentation thus enables continuous enforcement of conditional access, resilience against insider threats, and scalable policy automation aligned with Zero Trust principles.
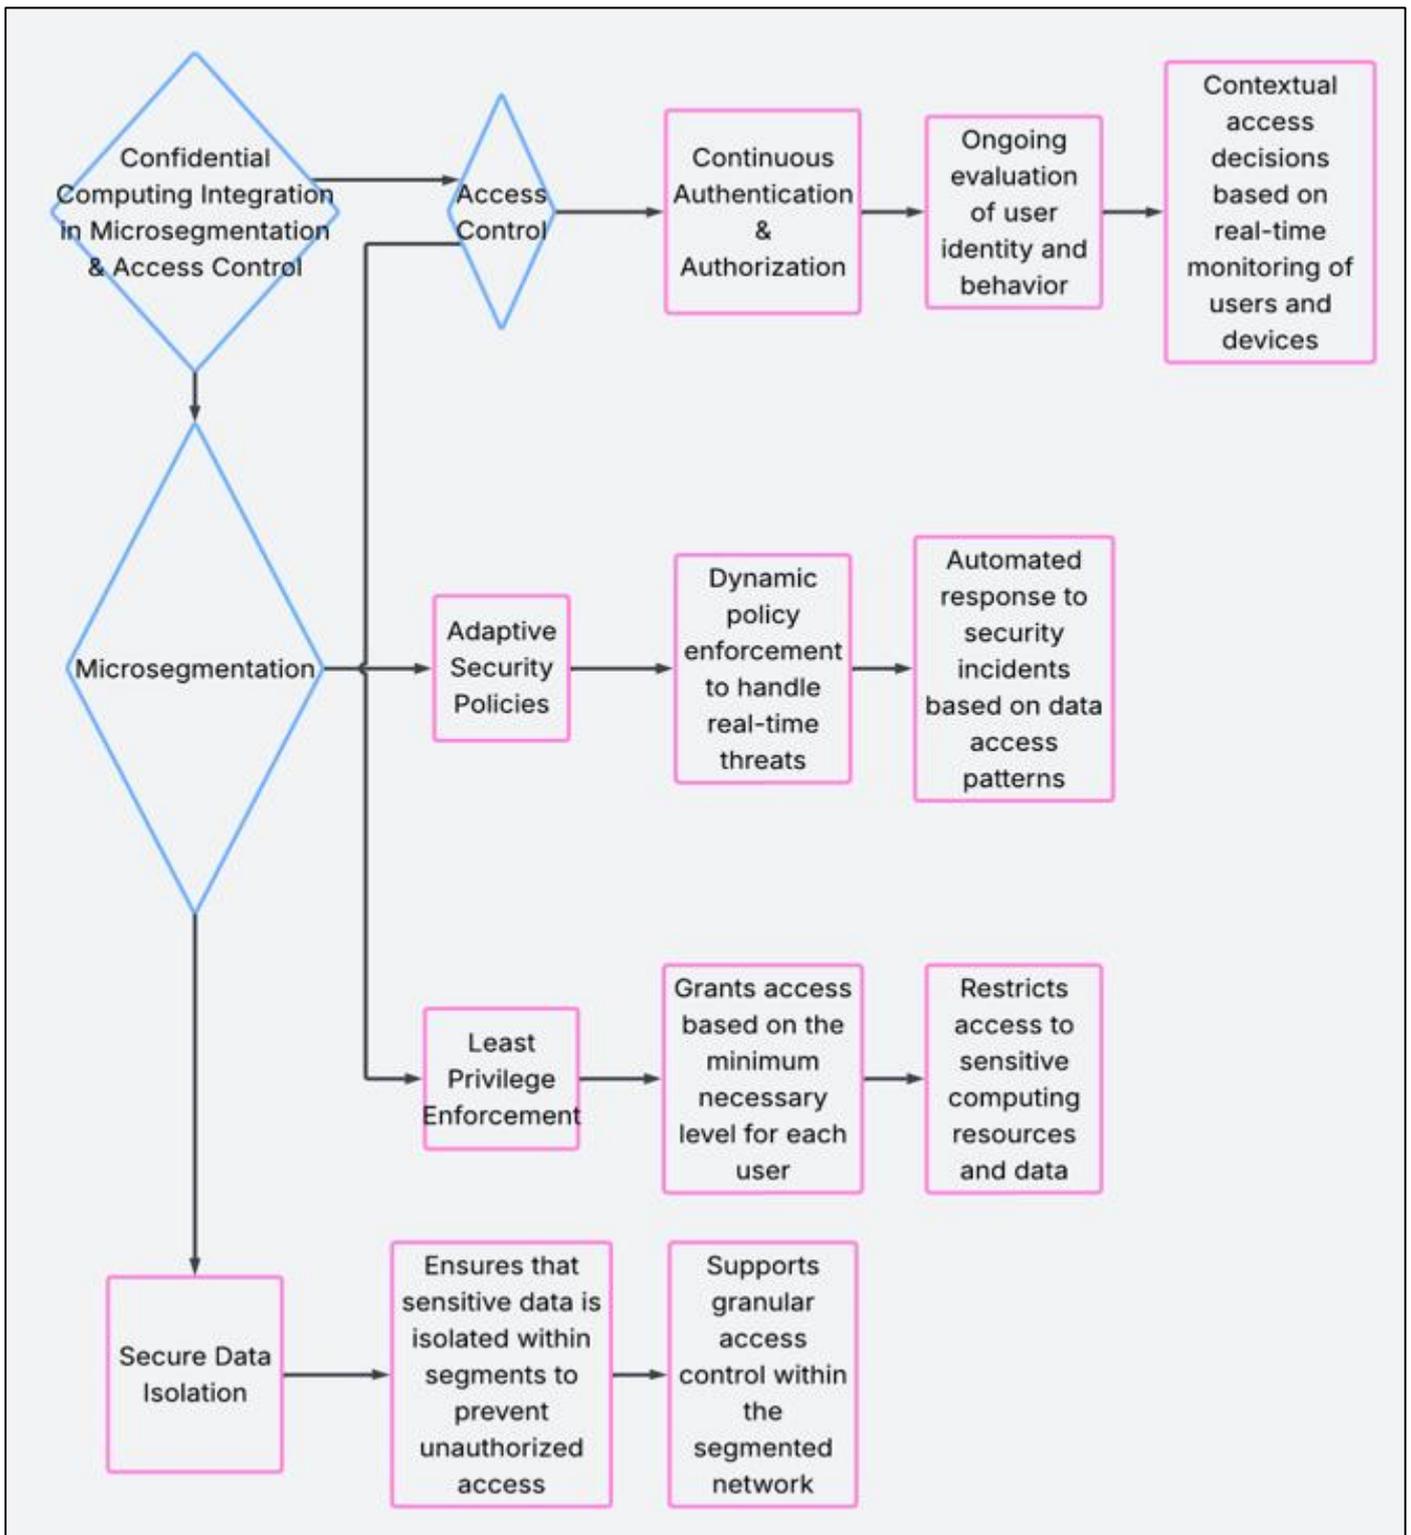
Fig 3 Incorporating Confidential Computing into Microsegmentation and Access Control

Figure 3. Illustrates the integration of confidential computing into microsegmentation and access control within a cloud security framework. At the center, the flowchart begins with Confidential Computing Integration in Microsegmentation & Access Control, leading into two main branches: Microsegmentation and Access Control. The Microsegmentation branch details practices like Adaptive Security Policies, ensuring dynamic enforcement to handle real-time threats, and Secure Data Isolation, which isolates

sensitive data within specific segments to prevent unauthorized access. The Access Control branch focuses on Continuous Authentication & Authorization, which involves ongoing user identity evaluation and behavior monitoring, and Least Privilege Enforcement, ensuring that access to resources is granted based on the minimum required level for each user. The diagram visually captures how confidential computing enhances both security segmentation and access management through robust policies that enforce data

protection, mitigate threats, and provide granular control within a segmented network environment.

➢ Risk-Adaptive Access Management

Risk-adaptive access management (RAAM) plays a pivotal role in actualizing zero trust principles by dynamically adjusting access permissions based on real-time contextual risk evaluations. Unlike static access control models that rely solely on predefined roles or attributes, RAAM leverages behavioral analytics, threat intelligence, and continuous monitoring to assess risk and grant or restrict access accordingly (Ahmed, M. et al., 2020). In cloud environments where confidentiality and data sovereignty are critical, RAAM frameworks integrated with confidential computing paradigms can enforce micro-policies that operate within trusted execution environments (TEEs), ensuring that decisions are both context-aware and privacy-preserving.

Confidential computing enhances RAAM by securing the execution of access control logic and identity verification processes within hardware-enforced enclaves. This capability mitigates the risk of privilege escalation and insider threats that exploit vulnerabilities outside of enclave boundaries (Xiao et al., 2022). Furthermore, RAAM supports granular policy enforcement through real-time telemetry and anomaly detection, aligning with zero trust mandates that "never trust, always verify." As cloud workloads increase in complexity and threat landscapes evolve, adaptive mechanisms that respond proportionately to risk signals are essential to preserving security without compromising operational agility. This section underscores the significance of coupling risk-adaptive strategies with confidential computing infrastructure to support zero trust-based threat mitigation models at scale.

➢ Detection and Response in Isolated Enclaves

In the context of Zero Trust Architecture (ZTA), the integration of confidential computing technologies, particularly through the use of isolated enclaves, has emerged as a pivotal advancement in enhancing threat detection and response mechanisms within cloud environments. These enclaves, often implemented via Trusted Execution Environments (TEEs), provide hardware-enforced isolation, ensuring that sensitive data and processes remain protected even in the presence of potential system compromises.

The deployment of TEEs facilitates secure enclaves that are instrumental in executing critical security functions such as real-time threat detection, behavioral analytics, and automated incident response. By isolating these processes, organizations can mitigate the risks associated with lateral movement attacks and unauthorized data access, which are prevalent in traditional security models. Moreover, the ability of TEEs to perform remote attestation ensures that only verified and trusted code is executed within these enclaves, aligning with the ZTA principle of "never trust, always verify" (InformationWeek, 2025).

Recent advancements have demonstrated the efficacy of combining TEEs with other security measures to bolster cloud security. For instance, the integration of confidential computing with identity and access management systems enhances the enforcement of least privilege access controls, a core tenet of ZTA. Additionally, the use of secure enclaves in processing sensitive workloads allows for continuous monitoring and rapid response to threats without exposing critical data to potential breaches (Hasan, 2024).

Furthermore, the application of confidential computing extends to supporting compliance with regulatory standards by providing verifiable assurances of data integrity and confidentiality. This is particularly significant in sectors handling sensitive information, where maintaining data privacy is paramount. The adoption of these technologies signifies a shift towards a more resilient and proactive security posture, essential for addressing the evolving threat landscape in cloud computing.

VI. IMPLEMENTATION CHALLENGES AND FUTURE DIRECTIONS

➢ Scalability and Interoperability Issues

Despite its promise, implementing dynamic confidential computing for continuous cloud security posture monitoring poses significant scalability and interoperability challenges. As cloud-native infrastructures evolve into more decentralized and multi-cloud architectures, the orchestration of Trusted Execution Environments (TEEs) across heterogeneous hardware and cloud service providers becomes complex and inefficient. TEEs such as Intel SGX and AMD SEV are designed with limited enclave memory and rely heavily on specific processor capabilities, which constrain their scalability in distributed workloads (Sardar et al., 2023). These limitations hinder seamless workload migration and horizontal scaling, critical for maintaining real-time visibility and policy enforcement in large-scale Zero Trust deployments.

Table 5 Scalability and Interoperability Issues in Confidential Computing Integration

Issue	Description	Technical Challenges	Potential Solutions
Scalability of Confidential Computing	Confidential computing systems must handle large-scale, multi-tenant cloud environments, where resources are distributed across numerous virtual machines.	High computational overhead, latency in encrypted data processing, and difficulty in scaling secure enclaves across various cloud environments.	Use of hardware accelerators like GPUs and specialized cloud infrastructure that optimizes enclave scalability.
Interoperability with Legacy Systems	Integrating confidential computing with existing infrastructure and legacy systems often results in	Difficulty in integrating with older systems that were not designed with confidentiality or	Adopting hybrid cloud models and standardized APIs for seamless

	compatibility challenges due to differing protocols.	zero trust in mind, causing operational bottlenecks.	communication between old and new systems.
Data Movement Between Segments	Moving sensitive data between microsegmented zones or secure enclaves while maintaining confidentiality can cause delays and security vulnerabilities.	Network latency, risk of data leaks during transitions, and potential inefficiencies in cross-segment data flow.	Implementing optimized data routing protocols with secure tunneling techniques for better data protection.
Policy Enforcement Across Multiple Domains	Enforcing consistent policies across diverse cloud and on-premises environments, particularly with varying regulations and security protocols.	Complexities in maintaining uniform security policies across various cloud platforms, third-party services, and hybrid environments.	Deployment of centralized policy management platforms to enforce consistent access control and security policies.

Interoperability also presents a fundamental barrier when integrating confidential computing technologies into existing Continuous Cloud Security Posture Monitoring (CCSPM) systems. The fragmented support across cloud vendors and the lack of standardized APIs or protocols for enclave attestation, policy synchronization, and telemetry collection impair consistent security enforcement (Zhang et al., 2021). Moreover, ensuring policy coherence across enclaves hosted in different cloud environments adds to the architectural and operational overhead as seen in table 5. These constraints challenge the agility and reliability of threat mitigation models that rely on real-time, context-aware data processing.

Without a unified framework for scalable confidential computing, the adoption of Zero Trust principles across hybrid and multi-cloud environments risks becoming fragmented. To realize the full potential of dynamic confidential computing in cloud-native security, a coordinated push toward hardware-agnostic architectures, open standards, and federated attestation frameworks is essential.

➢ *Performance Overhead of Encrypted Execution*
Confidential computing introduces significant security advantages by enabling encrypted data processing within Trusted Execution Environments (TEEs); however, these benefits often come with measurable performance trade-offs. Execution within TEEs such as Intel SGX can lead to increased memory access latency, cache misses, and I/O overhead, particularly when large-scale cloud-native applications are containerized for secure processing (Arnautov et al., 2018). This performance degradation becomes more pronounced in continuous cloud security posture monitoring (CCSPM) environments, where telemetry analysis, policy enforcement, and threat response require real-time or near-real-time execution. The overhead not only affects latency-sensitive workloads but also limits the scalability of Zero Trust-based models deployed across hybrid and multi-cloud systems.

Further complicating performance optimization is the constraint that enclaves are restricted in memory size and do not inherently support multithreading efficiently. Studies have shown that computational sandboxing frameworks, such as Ryoan, can mitigate some of these limitations through distributed enclave design, but at the cost of increased architectural complexity and reduced throughput under high workload concurrency (Hunt et al., 2018). These factors impose critical trade-offs in operationalizing threat mitigation strategies that depend on encrypted runtime analytics. For Zero Trust to be dynamically enforced within confidential environments, cloud systems must adapt by balancing security guarantees with processing efficiency. As cloud security monitoring becomes increasingly continuous and adaptive, reconciling the tension between execution integrity and operational performance remains central to sustainable deployment models.

➢ *Compliance and Regulatory Considerations*
The integration of confidential computing within cloud security posture monitoring frameworks introduces significant compliance and regulatory considerations. Confidential computing technologies, such as Trusted Execution Environments (TEEs), aim to protect data in use, thereby enhancing data confidentiality and integrity. However, their implementation must align with existing regulatory frameworks like the General Data Protection Regulation (GDPR) and the Health Insurance Portability and Accountability Act (HIPAA), which mandate stringent data protection measures.

Eiselt et al. (2025) critically examine the deployment of Confidential Virtual Machines (CVMs) in public cloud infrastructures, highlighting that while CVMs offer enhanced security, they often fall short in fully isolating workloads from cloud providers. This partial isolation raises concerns about compliance, as regulatory standards require clear boundaries to prevent unauthorized data access. The study underscores the necessity for transparent attestation processes and verifiable trust mechanisms to meet compliance obligations.

Similarly, Korada, L. (2024). discuss the challenges of implementing confidential computing in multi-tenant cloud environments. They emphasize that while TEEs can bolster data security, their effectiveness in ensuring compliance depends on proper key management and the ability to provide audit trails. The authors advocate for standardized protocols and certifications to validate the security claims of confidential computing solutions, which are essential for regulatory adherence.

Incorporating confidential computing into cloud security posture monitoring necessitates a comprehensive understanding of regulatory requirements and the

development of mechanisms to demonstrate compliance. Organizations must ensure that their confidential computing implementations not only enhance security but also provide the necessary transparency and accountability demanded by regulatory bodies.

➢ *Future Trends: AI-Augmented Confidential Computing, Federated ZT Models*

The future of continuous cloud security posture monitoring lies in the integration of artificial intelligence (AI) with confidential computing to enable intelligent, self-adaptive threat detection and mitigation. AI-augmented confidential computing environments can dynamically interpret encrypted telemetry data within trusted execution environments (TEEs), enabling real-time anomaly detection without compromising data confidentiality (Ghorbanian et al., 2019). Additionally, the advancement of federated zero trust (ZT) models across hybrid and multi-cloud ecosystems represents a pivotal direction. These models decentralize authentication, authorization, and policy enforcement, ensuring that no implicit trust exists across any node or service—irrespective of its location. As edge computing proliferates, the demand for lightweight, hardware-based confidential computing mechanisms such as Intel SGX or ARM TrustZone increases to support distributed risk analysis and compliance monitoring (Hua et al., 2020). The convergence of federated learning, AI, and confidential computing can further allow for collaborative anomaly modeling across organizations without exposing raw data, thereby promoting resilience against advanced persistent threats. These emerging paradigms not only extend zero trust beyond traditional perimeters but also build dynamic, context-aware security operations that scale with evolving cloud-native infrastructures.

VII. SUMMARY AND CONCLUSION

This review highlights the convergence of dynamic confidential computing and Zero Trust frameworks as a transformative paradigm in enhancing continuous cloud security posture monitoring (CCSPM). Confidential computing enables secure data processing in isolated enclaves, ensuring data remains encrypted during computation, while Zero Trust enforces strict identity verification and microsegmentation (Chandramouli & Kuhn, 2020). Together, these technologies mitigate internal and external threats by eliminating implicit trust and ensuring cryptographic protection for sensitive workloads across dynamic, multi-cloud environments.

The strategic advantage of integrating these technologies into cloud operations lies in their capacity to continuously monitor, detect, and contain threats without compromising data privacy or performance. By embedding confidential computing into real-time posture management workflows, organizations can establish verifiable trust anchors and dynamic access controls, fostering resilience against sophisticated cyberattacks (Conti et al., 2018). This layered approach aligns with evolving compliance mandates and enhances response times to security anomalies, thereby reducing the risk window for data breaches.

Future research should explore lightweight secure enclaves for edge devices, scalable attestation models, and AI-driven anomaly detection within Zero Trust boundaries. Policymakers should also prioritize standardizing secure workload isolation and adaptive trust evaluation to support confidential computing adoption in regulated sectors. These directions will be pivotal in advancing a verifiable, zero-trust-aligned cloud security architecture.

The integration of dynamic confidential computing with continuous cloud security posture monitoring offers a groundbreaking path toward operationalizing Zero Trust principles in modern cloud environments. By leveraging trusted execution environments and real-time telemetry, organizations can ensure that sensitive data remains secure throughout its lifecycle while maintaining visibility into evolving threat landscapes. This fusion not only strengthens data confidentiality and access control but also enhances the agility and precision of threat mitigation strategies in distributed and multi-tenant architectures. As cyber threats become more sophisticated, this Zero Trust-aligned model emerges as a necessary evolution in securing digital assets and maintaining regulatory compliance. Moving forward, the continued development of scalable, interoperable, and performance-efficient solutions will be essential to realizing the full potential of this architecture across cloud-native ecosystems.

REFERENCES

[1]. Ahmadi, S. (2024). Zero trust architecture in cloud networks: Application, challenges and future opportunities. *Journal of Engineering Research and Reports*, *26*(2), 215-228.

[2]. Akinade, A. O., Adepoju, P. A., Ige, A. B., & Afolabi, A. I. (2024). Cloud security challenges and solutions: A review of current best practices. *Int. J. Multidiscip. Res. Growth Eval*, *6*, 26-35.

[3]. Alasmary, W., Mehmood, R., & Katib, I. (2021). Intelligent cloud-native security posture management for adaptive threat mitigation. *Future Generation Computer Systems*, *125*, 503–514. https://doi.org/10.1016/j.future.2021.06.016

[4]. Ali, W., & Awad, A. I. (2021). A trust-aware cloud security posture assessment framework based on continuous monitoring. *Future Generation Computer Systems*, 124, 178–190. https://doi.org/10.1016/j.future.2021.05.014

[5]. Alim, M. A., Eshete, B., & Liu, Z. (2021). Adaptive security posture monitoring in cloud systems using machine learning and real-time analytics. *Journal of Cloud Computing: Advances, Systems and Applications*, 10(1), 1–18. https://doi.org/10.1186/s13677-021-00239-2

[6]. Alshamrani, A., Myneni, S., Chowdhary, A., & Huang, D. (2019). A defense-in-depth framework for cloud security using automated policy enforcement and anomaly detection. *Future Generation Computer Systems,* 99, 605–614. https://doi.org/10.1016/j.future.2019.05.031

[7]. Chandramouli, R., & Dang, Q. H. (2020). *Zero Trust Architecture*. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207

[8]. Chandrasekaran, S., Gupta, G., & Shanthini, A. (2021). A zero trust-based adaptive threat detection model for securing hybrid cloud infrastructures. *Journal of Cloud Computing*, *10*(1), 1–19. https://doi.org/10.1186/s13677-021-00247-5

[9]. Cheng, R., Zhang, F., Kos, J., He, W., Hynes, N., Johnson, N., ... & Song, D. (2018). Ekiden: A platform for confidentiality-preserving, trustworthy, and performant smart contract execution. *arXiv preprint arXiv:1804.05141*.

[10]. Costan, V., & Devadas, S. (2016). Intel SGX explained. *IACR Cryptology ePrint Archive*, 2016, 86. https://eprint.iacr.org/2016/086

[11]. Costan, V., & Devadas, S. (2016). Intel SGX explained. *Proceedings of the IEEE Symposium on Security and Privacy*, 1(1), 1–27. https://scholar.google.com/scholar_lookup?title=Intel%20SGX%20explained&author=Costan&publication_year=2016

[12]. Costan, V., & Devadas, S. (2016). Intel SGX Explained. *Proceedings of the 2016 IEEE Symposium on Security and Privacy*, 1–18. https://doi.org/10.1109/SP.2016.35

[13]. Costan, V., Leblanc, S., & Devadas, S. (2016). Sanctum: Minimal hardware extensions for strong software isolation. *Proceedings of the 25th USENIX Security Symposium*, 857–874. https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/costan

[14]. Damaraju, A. (2022). Integrating Zero Trust with Cloud Security: A Comprehensive Approach. *Journal Environmental Sciences and Technology*, *1*(1), 279-291.

[15]. Gade, K. R. (2022). Cloud-Native Architecture: Security Challenges and Best Practices in Cloud-Native Environments. *Journal of Computing and Information Technology*, *2*(1).

[16]. Goltzsche, D., Rüsch, S., Nieke, M., Vaucher, S., Weichbrodt, N., Schiavoni, V., & Kapitza, R. (2018, June). Endbox: Scalable middlebox functions using client-side trusted execution. In *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)* (pp. 386-397). IEEE.

[17]. Jimmy, F. N. U. (2023). Cloud security posture management: tools and techniques. *Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online)*, *2*(3).

[18]. Jimmy, F. N. U. (2023). Cloud security posture management: tools and techniques. *Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online)*, *2*(3).

[19]. Kodakandla, N. (2024). Securing Cloud-Native Infrastructure with Zero Trust Architecture. *Journal of Current Science and Research Review*, *2*(02), 18-28.

[20]. Kodakandla, N. (2024). Securing Cloud-Native Infrastructure with Zero Trust Architecture. *Journal of Current Science and Research Review*, *2*(02), 18-28.

[21]. Korada, L. (2024). Use Confidential Computing to Secure Your Critical Services in Cloud. *Machine Intelligence Research*, *18*(2), 290-307.

[22]. Mavroeidis, V., Vishi, K., Zych, M. D., & Jøsang, A. (2020). The security and privacy of smart environments: A critical review. *Journal of Information Security and Applications*, 52, 102500. https://doi.org/10.1016/j.jisa.2020.102500

[23]. Mohammadi, M., Al-Fuqaha, A., Guizani, M., & Oh, J. (2018). Semi-supervised deep reinforcement learning in support of IoT and smart city services. *IEEE Internet of Things Journal, 5*(2), 624–635. https://doi.org/10.1109/JIOT.2017.2744138

[24]. Raj, H., Le T., Saroiu, S., Wolman, A., & England, P. (2016). *Flicker: A flexible platform for secure system extensions*. ACM SIGOPS Operating Systems Review, 40(4), 315–328. https://doi.org/10.1145/1168857.1168901

[25]. Ristenpart, T., Tromer, E., Shacham, H., & Savage, S. (2009). Hey, you, get off of my cloud: Exploring information leakage in third-party compute clouds. *Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS)*, 199–212. https://doi.org/10.1145/1653662.1653687

[26]. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero Trust Architecture*. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207

[27]. Russinovich, M., Costa, M., Fournet, C., Chisnall, D., & Delignat-Lavaud, A. (2023). *Confidential computing: Elevating cloud security and privacy*. Queue, 21(4), 44–48. https://doi.org/10.1145/3623461(ResearchGate)

[28]. Sabt, M., Achemlal, M., & Bouabdallah, A. (2015). Trusted execution environment: What it is, and what it is not. *IEEE Trustcom/BigDataSE/ISPA*, 1, 57–64. https://doi.org/10.1109/Trustcom.2015.357

[29]. Shafagh, H., Burkhalter, L., Hithnawi, A., & Hubaux, J. P. (2017). Towards blockchain-based auditable storage and sharing of IoT data. *Proceedings of the 2017 on Cloud Computing Security Workshop (CCSW)*, 45–50. https://doi.org/10.1145/3140649.3140656

[30]. Shaukat, K., Luo, S., Varadharajan, V., Hameed, I. A., Chen, S., & Li, J. (2021). A review on security challenges in cloud computing: Issues, solutions, and future directions. *Journal of Network and Computer Applications*, 179, 102983. https://doi.org/10.1016/j.jnca.2020.102983

[31]. Shih, M. W., Wang, J., Dautenhahn, N., & Lee, R. B. (2021). Enforcing zero trust security policies using enclave-based isolation in cloud environments. *IEEE Transactions on Dependable and Secure Computing*, 18(3), 1073–1087. https://doi.org/10.1109/TDSC.2019.2906006

[32]. Shinde, S., Shrestha, B., Shanbhogue, V., Pei, F., Xu, Z., Wolff, F., & Seshia, S. A. (2017). PANOPLY: Low-TCB Linux Applications with SGX Enclaves. *Proceedings of the 2017 USENIX Annual Technical Conference (USENIX ATC 17)*, 1–14.

https://www.usenix.org/conference/atc17/technical-sessions/presentation/shinde

[33]. Taherkordi, A., Zahid, F., Verginadis, Y., & Horn, G. (2018). Future cloud systems design: challenges and research directions. *IEEE Access*, *6*, 74120-74150.

[34]. Tuyishime, E., Balan, T. C., Cotfas, P. A., Cotfas, D. T., & Rekeraho, A. (2023). Enhancing cloud security—proactive threat monitoring and detection using a siem-based approach. *Applied Sciences*, *13*(22), 12359.

[35]. Ullah, I., Ahmad, F., & Kim, D. (2021). Security and performance evaluation of cloud-native security posture management systems. *Future Generation Computer Systems, 124*, 131–145. https://doi.org/10.1016/j.future.2021.05.006

[36]. Zhang, Y., Chen, X., Zhang, L., & Liu, Y. (2021). Towards secure and efficient data sharing in cloud computing using trusted execution environments. *IEEE Transactions on Cloud Computing*, 9(3), 1247–1259. https://doi.org/10.1109/TCC.2019.2904461

[37]. Zhang, Y., Qiu, M., Tsai, C. W., Hassan, M. M., & Alamri, A. (2017). Health-CPS: Healthcare Cyber-Physical System Assisted by Cloud and Big Data. *IEEE Systems Journal, 11*(1), 88–95. https://doi.org/10.1109/JSYST.2015.2460747