



An Intent-Aware Zero Trust Identity Architecture for Unifying Human and Machine Access

Badal Bhushan^{1*}, Prassanna R Rajgopal², Kritika Sharma³

¹Cybersecurity Leader, USA

* Corresponding Author Email: Badalbhushan786@gmail.com-ORCID: 0009-0006-6102-7591

²Cybersecurity Researcher/ Member IEEE, ISACA

Email: prassannarr@gmail.com-ORCID: 0009-0009-7461-5220

³Johns Hopkins University

Email: sharma14@outlook.com-ORCID: 0009-0004-8429-0814

Article Info:

DOI: 10.22399/ijcesn.3886

Received : 20 July 2025

Accepted : 26 September 2025

Keywords

Zero Trust Architecture,
Identity Federation,
Context-Aware Access,
Machine Identity,
Intent-Aware Orchestration,
Workload Identity,

Abstract:

Zero Trust is now the de facto standard to secure cloud-native, distributed, and AI-driven enterprise infrastructures. It's not only crucial to address human identities but also to secure non-human entities such as APIs, software agents, RPA bots, and smart city workloads. As hybrid infrastructures become the new normal and agentic AI systems (e.g., self-driving cars) grow more autonomous, identity remains the most stable and trustworthy security control plane. This document proposes an intent-aware Zero Trust Identity Architecture designed to consolidate governance, authentication, and access control for human and non-human entities. The architecture consists of decentralized identity provisioning, policy-as-code enforcement, real-time telemetry ingestion, trust scoring, and AI-powered intent detection to provide inputs for continuous verification and least privilege enforcement. Compliant with standards such as NIST SP 800-207, NIST SP 800-63, CISA Zero Trust Maturity Model, and DoD's Zero Trust Strategy, the architecture also aligns with industry developments from Microsoft Entra ID, AWS IAM Identity Center, Google BeyondCorp, SPIFFE/SPIRE, and W3C DIDs. The whitepaper explores use cases in healthcare, finance, retail, and industrial IoT spaces that are struggling with unique challenges like OT/IT convergence, multi-user devices, and governance of sensitive data access. High-profile attacks such as SolarWinds, MOVEit, and Log4Shell are broken down to highlight weaknesses in legacy IAM architectures and underscore the need for intent-based security. By intersecting behavior, purpose, and identity, this architecture remakes trust in hybrid, edge, and cloud-native settings with a conclusion of actionable paths of mitigation and a vision for intent-based Zero Trust governance.

1. Introduction

The move from perimeter-based security to identity-driven access control represents the modern-day evolution of cybersecurity architecture. Zero Trust, characterized in foundational papers such as the National Institute of Standards and Technology (NIST) Special Publication 800-207, is centered on the principle of "never trust, always verify" with continuous identity assurance, least privilege, and real-time context evaluation regardless of network boundaries [1]. This history has been extended by standards from bodies such as NIST SP 800-63 on digital identity confidence, the Cybersecurity and

Infrastructure Security Agency (CISA) Zero Trust Maturity Model (ZTMM), the Department of Defense (DoD) Zero Trust Strategy, and the emerging guidelines of the IETF Zero Trust Architecture Working Group [2]–[5]. Big cloud providers such as Microsoft, Google, and Amazon Web Services (AWS) have made these ideas reality through the likes of Microsoft Entra ID, Google BeyondCorp Enterprise, and AWS IAM Identity Center delivering fine-grained policy enforcement, secure workload identities, and behavioral telemetry to underlying access control systems [6]–[8]. Despite widespread Zero Trust adoption among human users, enterprise architectures fall behind in

addressing non-human identities such as RPAs, DevOps agents, API-based services, container workloads, and AI agents. IAM systems that use traditional methodologies, despite their success in handling human user passwords, roles, and multi-factor authentication, struggle to scale dynamically among autonomous, decentralized, and ephemeral entities. Even with initiatives such as SPIFFE (Secure Production Identity Framework for Everyone), its reference implementation SPIRE, and W3C's Decentralized Identifiers (DIDs) offering useful primitives of cryptographically bound identity, they remain in their own silos without being integrated into an end-to-end policy enforcement and risk-adaptive trust framework [9]–[11]. This paper presents an intent-aware Zero Trust Identity Architecture intended to bring identity governance together for human and non-human participants in hybrid, cloud-native, and edge-based computing environments. The proposed model features decentralized identity provisioning, dynamic policy enforcement via policy-as-code, behavioral analytics, trust scoring, and AI-driven intent recognition. It also follows closely the NIST, CISA, and CSA security best practices but uses implementation trends from top vendors such as Google, AWS, and Microsoft.

2. Background and Related Work

Legacy Identity and Access Management (IAM) systems were architected with an assumption of trust based on network boundaries and focused primarily on static, human-centric identity credentials. These systems authenticated users through predefined roles and permissions, typically stored in enterprise directories. However, in modern enterprise environments, where workloads are increasingly distributed, ephemeral, and heterogeneous, such identity constructs prove insufficient.

In an effort to enhance the security posture of IAM systems, the industry gradually evolved toward federated identity models, multi-factor authentication (MFA), and single sign-on (SSO) mechanisms. These improvements undeniably increased security resilience by reducing credential reuse and centralizing access control. Yet, they remained fundamentally unscalable and inflexible when applied to non-human identities which now constitute the majority of access requests in cloud-native systems. Entities such as APIs, microservices, containers, serverless functions, and robotic process automations increasingly interact without direct human oversight, demanding a new identity paradigm that supports transparent authentication, context-aware authorization, and real-time decision-making.

To address this gap, industry and open-source communities have introduced short-lived, cryptographically verifiable identities tailored for dynamic workloads. The Secure Production Identity Framework for Everyone (SPIFFE) and its runtime implementation, SPIRE, are leading solutions in this domain. These technologies enable machine identities to be issued and rotated automatically within trusted runtime environments, forming a secure foundation for zero trust workload identity [9].

In parallel, identity innovation is being shaped by the World Wide Web Consortium's (W3C) frameworks such as Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs). These standards promote self-sovereign identity, offering a model where identities are not anchored to centralized authorities but rather exist as portable, cryptographically provable entities. These models are especially relevant in autonomous and cross-domain systems such as IoT networks, multi-cloud orchestrations, and decentralized edge computing [10].

Despite these technical advancements, a major operational limitation persists: many of these identity solutions are siloed and are not natively integrated into broader enterprise governance, risk, and compliance (GRC) frameworks. Enterprises often operate parallel IAM systems for humans and machines, leading to policy fragmentation, audit inconsistencies, and increased attack surfaces. This paper asserts that the next wave of identity innovation must bridge this gap embedding workload identity into enterprise-grade policy engines, continuous monitoring tools, and SOC analytics pipelines.

Simultaneously, the rise of Zero Trust Architecture (ZTA) has catalyzed a philosophical and architectural shift from implicit trust (based on network location or user group membership) to explicit, continuous trust assessment. The foundational principles outlined in NIST Special Publication 800-207, reinforced by initiatives from the Cloud Security Alliance (CSA), Microsoft, and the Cybersecurity and Infrastructure Security Agency (CISA), emphasize least-privilege access, behavioral telemetry, and dynamic policy enforcement [1], [3], [6], [12]. These practices represent a step toward making identity and access contextually aware and resilient to drift or compromise.

However, while considerable progress has been made in the areas of contextual access, policy-as-code engines, and continuous authentication, a critical dimension remains underexplored: the real-time inference of intent for autonomous agents. In human-centric systems, intent is often implied

through behavioral baselines and historic activity patterns. For machine entities, especially those interacting at cloud-scale, intent must be inferred dynamically through a combination of behavioral analytics, threat intelligence correlation, and runtime telemetry.

This gap is made evident in the analysis of recent high-profile security incidents. Industry compromise studies such as the MOVEit breach (2023), the SolarWinds supply chain attack (2020), and the Log4Shell vulnerability exploitation (2021) have demonstrated systemic failures in identity enforcement, privilege boundaries, and telemetry coverage [13]–[15]. In each case, attackers exploited over-permissioned systems, abused implicit trust relationships, or leveraged vulnerable middleware to traverse environments undetected. These examples underscore the inadequacy of static, role-based IAM systems and the urgent need for adaptive, real-time identity strategies.

To address this evolving threat landscape, a new generation of identity-enabling technologies is emerging. Policy-as-code platforms such as Rego (from the Open Policy Agent project) and Cedar (from Amazon Web Services) are gaining adoption as foundational components of modern authorization systems. These technologies allow access policies to be defined, versioned, and audited like software making them easier to test, scale, and deploy in CI/CD pipelines [16]–[18].

Additionally, vendors such as CrowdStrike and Microsoft Defender are introducing behavioral risk scoring engines, which dynamically calculate trust levels based on endpoint signals, anomaly detection, and peer comparisons. These scores can then be used to make risk-adaptive access decisions for example, denying access to a microservice exhibiting anomalous behavior even if it presents valid credentials.

At the enforcement layer, WebAssembly (WASM) runtimes are being adopted as edge-native enforcement engines that can run policy evaluation logic close to the data or API endpoint. This enables extremely fast, localized decision-making with minimal performance overhead an essential feature for latency-sensitive or high-throughput systems such as trading platforms, industrial control systems, or real-time analytics pipelines.

When combined, these emerging technologies signal a clear transition toward intent-based, identity-first architectures. These architectures rely on dynamic attributes such as process provenance, cryptographic identity, network behavior, and data sensitivity to infer access decisions in real time rather than relying solely on static configurations. The future of IAM is one where identity is ephemeral, contextual, and inherently programmable.

This paper leverages the above technological foundations to propose a multi-layered, governance-aligned identity framework tailored for autonomous workloads in distributed environments. The goal is not only to authenticate and authorize access, but to infer, validate, and continuously reassess trust in an ecosystem where entities human and machine are constantly changing, communicating, and evolving.

3. Core Constructs of Intent-Aware Zero Trust Identity Orchestration

As cybersecurity architecture transitions from perimeter-based controls to Zero Trust principles, the centrality of identity in access decision-making has grown significantly. However, traditional identity systems designed around static role assignments, group memberships, and manually assigned credentials are ill-suited for modern, dynamic, and distributed computing environments. In today's Zero Trust architectures, identity must evolve into an intent-aware orchestration layer a real-time, telemetry-driven control plane that dynamically evaluates access requests based on user behavior, contextual attributes, and inferred purpose rather than solely on predefined rules.

This transformation introduces a series of advanced constructs that redefine identity from a static artifact into an adaptive, policy-aware control signal. These constructs include identity classification, contextual telemetry analysis, behavioral profiling, purpose inference, and real-time policy decisioning. Together, they form the foundational pillars of intent-aware Zero Trust identity orchestration a capability that ensures access decisions reflect not only who is requesting access but also why, when, how, and under what operational context.

3.1 Evolving Identity Taxonomy in Zero Trust Ecosystems

Identity is no longer confined to human actors. A comprehensive identity framework must now account for a diverse taxonomy of entity types:

- Human users: Internal employees, contractors, third-party vendors, and partners.
- Workload identities: Containers, serverless functions, Kubernetes pods, and service accounts.
- Agentic AI: LLM-driven or task-executing AI agents acting on behalf of workflows.
- Robotic Process Automation (RPA): Scripted or adaptive bots interacting with enterprise systems.

- IoT/OT devices: Connected machines, embedded sensors, or industrial control systems.

Each identity type must be provisioned, authenticated, authorized, monitored, and eventually decommissioned. Furthermore, every identity must support dynamic posture assessment and continuous trust validation throughout its lifecycle. This complexity necessitates a flexible and intelligent orchestration engine that can interpret multiple attributes, behaviors, and contextual signals in real time.

3.2 Identity Contextualization and Metadata-Driven Categorization

Intent-aware identity orchestration begins with enriching identities with metadata derived from multiple sources such as role, device posture, geolocation, network health, session behavior, and access history. The addition of metadata enables dynamic categorization of identities not just by their credentials or entitlements, but by their operational purpose and behavioral patterns.

Purpose itself may be explicitly defined, such as through workflow annotations, task metadata, or business process tags embedded in API calls. Alternatively, purpose can be inferred through behavioral analytics and machine learning, which identify activity patterns that correlate with specific intents (e.g., transaction reconciliation, audit preparation, anomaly detection). This distinction allows an orchestrator to understand not only what is being accessed, but why it is being accessed a critical factor for intelligent access control.

Such a model can distinguish between a financial analyst logging into an enterprise dashboard to perform routine monthly closing versus accessing the same system outside normal hours with atypical queries potentially signaling insider threat or compromised credentials. This semantic depth of identity transforms the access control system into a proactive security mechanism.

3.3 Overlaying Intent on Classical IAM Models

In traditional IAM models, authorization is dictated by static attributes like department affiliation, job title, or group membership. Intent-aware orchestration introduces an additional dimension of dynamic context, such as:

- Time of access (e.g., during business hours vs. weekends),
- Device state (e.g., compliant or non-compliant),
- Job function context (e.g., current task vs. anomalous behavior),

- Risk posture (e.g., user risk score, session anomalies).

This enables fine-grained conditional logic such as: “Allow access only if a finance analyst is accessing HR documents during working hours from a compliant corporate device and the action aligns with expected behavior.”

Such policy enforcement becomes crucial in regulated sectors like healthcare or autonomous manufacturing, where misused or misattributed access can have safety or compliance ramifications [16], [17]. Intent-aware policies act as both a proactive guardrail and a forensic anchor enabling analysts to later interpret user behavior within the context of declared or inferred purpose.

3.4 Orchestration Across Multi-Cloud and Cloud-Native Environments

As enterprises adopt multi-cloud architectures, intent-aware identity orchestration must support cross-platform policy coherence. Each cloud provider exposes identity and access primitives differently, but intent-based access patterns can provide a unified abstraction.

For example:

- AWS Verified Access and IAM Identity Center allow ingestion of device telemetry and session trust levels to inform policy logic. These systems enable verification of posture before permitting access, thereby creating a feedback loop between security and identity [18], [19].
- Microsoft Entra ID Protection and Conditional Access policies incorporate user risk scores, sign-in anomalies, and workload sensitivity to enforce access in real-time. These indicators allow enterprises to suspend or deny access dynamically when behavioral risk crosses thresholds [3].
- Google’s BeyondCorp and Access Context Manager provide policy evaluation based on device state, IP reputation, geolocation, and historical behavior, enabling a persistent trust score that adapts per session [5], [20].

This platform-specific telemetry is normalized within the orchestration engine to make intent evaluation cloud-agnostic and resilient across hybrid infrastructures.

3.5 Real-Time Decisioning Engines and Continuous Authorization

A crucial capability in this architecture is the real-time policy decision point (PDP) which must not only evaluate access eligibility but also continuously re-evaluate trust. This necessitates the adoption of policy-as-code frameworks such as Rego and Cedar,

which allow organizations to express complex rules declaratively and link them to behavioral or contextual attributes without hardcoding them in application logic.

These rules can express high-level policies such as:

- “If workload X accesses sensitive data outside maintenance windows, notify and isolate.”
- “If user Y downloads more than 100 files in 15 minutes from HR systems, trigger risk review.”

Such expressions are contextually aware, intent-sensitive, and revocable mid-session, making them ideal for dynamic Zero Trust implementations.

Intent-aware Zero Trust identity orchestration represents a paradigm shift from verifying identities at a single point in time to continuously verifying behavior, context, and purpose. This approach ensures that identity becomes a living signal, capable of adapting to changing risk, usage, and threat conditions in real time.

As the threat landscape continues to evolve and identities proliferate across humans, machines, and AI agents, enterprises must architect their IAM systems not as static enforcement layers, but as adaptive, intelligent policy engines that learn, predict, and react with precision and agility.

4. Reference Architecture – Intent-Aware Zero Trust Identity Platform

The foreseen reference architecture for an intent-driven Zero Trust Identity platform is meant to empower aggregated, elastic control of human and non-human identities across modern enterprise use cases. Policy-controlled enforcement, cryptographic roots of trust, real-time telemetry processing, and behavior-enhanced access control are all part of the architecture. Multilayered design enforces least privilege, prevents lateral movement attacks, and promotes resilience in cloud, hybrid, and edge deployments.

At the center of the platform is an Identity Control Plane that is responsible for issuing, managing the lifetime of, and revoking identities. The control plane federates enterprise identity providers such as Microsoft Entra ID, Okta, and PingFederate to govern human identities using protocols such as SAML, OIDC, and SCIM [3], [4], [21]. For workload and machine identities, dynamic provisioning is offered by SPIFFE/SPIRE using X.509 short-term certificates [9], while Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs) offer identity assurance that is distributable to distributed IoT and autonomous systems [10], [22].

Authentication is mandated by hardware-backed credentials—i.e., TPM-bound certificates, FIDO2, and smartcards tied to device trust and posture. Technologies such as Azure AD Certificate-Based Authentication (CBA), Google Workload Identity Federation, and AWS IAM Roles Anywhere provide cryptographically secure, policy-based authentication for human actors and workloads [23], [5], [4]. For privileged identity access, the architecture comes with Just-in-Time (JIT) solutions such as CyberArk Dynamic Privilege and Microsoft Privileged Identity Management (PIM) to eliminate standing permissions [21], [24].

Policy enforcement is collocated in Policy Decision Points (PDPs) and Policy Enforcement Points (PEPs), according to the NIST Zero Trust framework [1]. Policies are defined in policy-as-code frameworks such as Rego (Open Policy Agent) or Cedar (used by AWS Verified Permissions), allowing for version control, testability, and enforcement of rules based on identity attributes, risk posture, and inferred purpose [16], [17], [25]. Policy bundling to WebAssembly (WASM) enables deployment of decision engines to edge locations for low-latency, offline operation [30].

A dedicated telemetry and context aggregation layer gathers signals from MDM (such as Intune, JAMF), EDR (such as Defender for Endpoint, CrowdStrike), and network telemetry sources (such as NAC tools, VPN logs) [31]. They are used to infer device trust, session risk, behavioral anomalies, and purpose inference. Purpose-enriched inference engines integrate employment metadata, behavioral profiling (via UEBA), and self-stated task purpose (e.g., from ServiceNow workflows or CI/CD pipeline contexts) to score and match identities with policy decisions in real time [25], [26].

The architecture also values immutable logging and auditability. All that makes access decisions, issues credentials, or re-checks trust is recorded in tamper-evident storage e.g., CloudTrail (AWS), Azure Monitor, or ledger-integrated stores [35]. Integration with central SIEM systems like Splunk, Sentinel, or Elastic enables correlation of activity across users and workloads, with behavior baselining or outlier-driven alerts [20], [34]. A human-governance and explainability interface is placed on top of the platform, through which stewards and analysts can inspect decision paths, override during emergencies, and glimpse real-time identity behavior. Tools like SHAP (SHapley Additive Explanations), TreeExplainer, and symbolic tracing reveal AI-aided decisions [37].

This NIST SP 800-207 compliant reference architecture combines future-proof concepts of Microsoft, Google, and AWS Zero Trust models to create an integrated, intent-based identity

governance plane. Below, we look at how this architecture matures through industry verticals, and provide case-specific guidance for retail, healthcare, finance, manufacturing, and others.

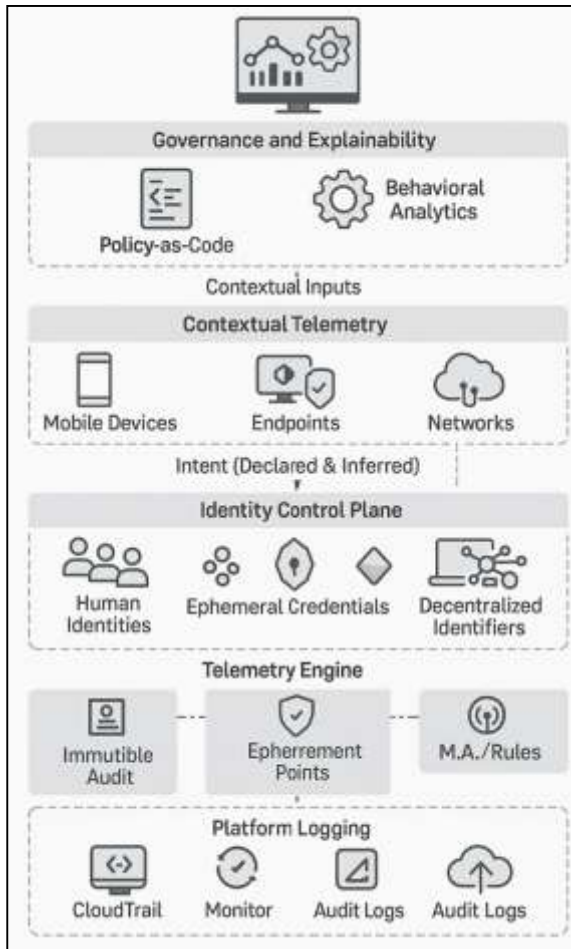


Figure 1 Architecture- Zero Trust Identity Platform

5. Threat Modeling and Security Analysis of the Proposed Solution

This section gives a formal threat model of the proposed Intent-Aware Zero Trust Identity Platform, detailing possible attack vectors and showing how architectural design and its constituents eliminate these threats. One may use a systematic approach, e.g., STRIDE (Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege) or Attack Trees, to analyze threats to the system's sensitive assets and trust boundaries [27].

5.1 Methodology and Scope

This sub-section will outline the chosen threat modeling method (e.g., STRIDE, Attack Trees, or a combination thereof) and scope the analysis. The focus will be placed on the essential components of the Intent-Aware Zero Trust Identity Platform, with particular reference to the Identity Control Plane,

Policy Enforcement, Telemetry and Context Aggregation Layer, and the AI-driven Intent Inference Engine.

5.1.1 Assets Identification

Identify the key assets, i.e., identities (human, machine, AI agents), policy definitions, telemetry data, trust scores, and access decisions [28].

5.1.2 Trust Boundaries

Identify the trust boundaries between architectural components (e.g., between Identity Control Plane and external IdPs, between PDPs/PEPs and target resources, between telemetry sources and the aggregation layer) [29].

5.1.3 Entry Points

Identify how actors (human users, machine identities, attackers) interact with the system [30].

5.2 Core Component Threat Analysis

5.2.1 Identity Control Plane Threats

Spoofing/Tampering with Identities

Threat: An adversary tries to hijack human or machine identities

(e.g., hijacking X.509 certs for workloads, hijacking DID for IoT devices, or hijacking human users using stolen credentials).

Mitigation: Use Strong, hardware-secured credentials (TPM-locked certs, FIDO2) and dynamic provisioning with SPIFFE/SPIRE for transient X.509 certs [9], [23]. Also, use W3C DIDs/VCs for transferable identity assurance and Just-in-Time (JIT) [10].

Repudiation of Identity Issuance/Revocation

Threat: An insider attack or adversary abuse aims to prevent issuance or revocation of an identity.

Mitigation: Immutable logging and auditability of all relevant security events, i.e., credential issuances and revocations, in tamper-proof storages (e.g., CloudTrail, Azure Monitor, ledger-integrated storages) [26], [35].

5.2.2 Policy Enforcement Layer Risks

Policy Tampering (Policy Sprawl/Misconfiguration)

Risk: Unauthorized or inadvertent misconfigurations lead to incorrect or incompatible access policies, enabling unauthorized access or denial of service.

Mitigation: Code-written policy with policy-as-code tools (Rego to OPA, Cedar to AWS Verified Permissions) permits versioning, testing, and dynamic enforcement [16], [17], [31].

Privilege Elevation through Policy Bypass

Threat: A malicious actor finds a policy logic bug or uses a bug in a distributed Policy Decision Point

(PDP) or Policy Enforcement Point (PEP) to gain unauthorized elevated privileges.

Mitigation: Distributed PDPs and PEPs according to NIST Zero Trust architecture. Policy bundling to WebAssembly (WASM) for secure, latency-sensitive deployment to edge points of presence to reduce reliance on central points of failure [30]. Continuous verification and least privilege enforcement are foundational principles [1].

5.2.3 Telemetry and Context Aggregation Layer Vulnerabilities

Information Disclosure/Tampering of Telemetry Data

Threat: An attacker is hijacking or modifying telemetry signals (e.g., MDM, EDR, network logs) to influence trust scores or intent inference, leading to false access decisions.

Mitigation: The architecture favors gathering live signals from heterogeneous, trusted sources. Data validity checks and secure communication processes would be presumed but necessary in this case. Integration with SIEM systems to cross-correlate is helpful for identifying telemetry anomalies [25], [31], [32].

Denial of Service (DoS) to Telemetry Ingestion

Threat: An attacker causes a flood attack on the telemetry aggregation layer, rendering real-time context updates unavailable and potentially leading to stale or incorrect trust determination.

Mitigation: Strong, scalable ingestion pipelines and the decentralization of telemetry sources and aggregation points would be needed. The paper indicates resilience in its distributed form [33].

5.2.4 AI-Driven Intent Inference Engine Threats Adversarial AI Attacks (Spoofing/Tampering of Intent)

Threat: An attacker constructs baiting inputs or manipulates behavioral patterns to mislead the AI into a spurious "intent" conclusion, thereby illegitimately granting access.

Mitigation: The architecture integrates employment metadata, User and Entity Behavior Analytics (UEBA), and self-reported purpose of tasks for determining intent. Future work includes "Autonomy Trust Anchoring" through cryptographic attestation methods on AI models [25], [36].

Information Disclosure (AI Model Inversion/Extraction)

Threat: The attacker attempts to reverse-engineer the AI model for sensitive information regarding its decision-making process or its training data.

Mitigation: The "Human Governance and Explainability Interface" with tools like SHAP and TreeExplainer provides controlled visibility of AI

decisions without necessarily exposing the complete complexity of the underlying model [28], [37].

5.3 Overall Security Posture and Resilience

Continuous Verification: The "never trust, always verify" strategy, applied continuously, greatly reduces the attack surface by eliminating implicit trust [1].

Least Privilege: Dynamic policy enforcement and JIT access patterns ensure that identities have only the least number of permissions they need, limiting the impact of a breach [22].

Unified Governance: By managing human and machine identities under a single governance fabric, the architecture closes critical gaps exploited by attackers who exploit the "human-machine identity blur" [4], [7].

Auditability and Explainability: Immutable logging and AI explainability technology support transparency and accountability, critical for detection, response, and compliance [25], [28], [37].

5.4 Threat Model Limitations

- **Complexity of AI/ML:** The "black box" nature of some AI/ML models may render comprehensive threat modeling challenging, requiring ongoing research in adversarial AI [36].
- **Changing Threat Landscape:** There may be novel attack forms on the horizon that are not countered by the current model.
- **Implementation-Specific Vulnerabilities:** This mathematical model is structural; implementation-specific idiosyncrasies may create new vulnerabilities.

6. Implementation Layers and Identity Control Plane

The intent-aware Zero Trust Identity Platform includes multiple architectural layers that collectively enable policy-enforced, context-sensitive security at scale. These layers allow organizations to orchestrate identity across diverse systems—from cloud-native workloads to industrial edge devices while enforcing granular control and enabling real-time risk evaluation.

At the core of this infrastructure is the Identity Control Plane, responsible for issuing, managing, and revoking both human and non-human identities. For workload identities, dynamic issuance is achieved through solutions like SPIFFE/SPIRE, which use ephemeral X.509 certificates to apply workload identity to runtime assertions [9]. For device and agentic identities, Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs) provide cryptographic authenticity, privacy-preserving claims, and portability across different

trust domains [10]. Human identity management also extends to enterprise identity systems such as Microsoft Entra ID, AWS IAM Identity Center, and Ping Identity. These systems support trust federation using SAML, OIDC, and SCIM protocols [3]–[5], and act as Identity Providers (IdPs) in Zero Trust architectures, offering multi-factor authentication, adaptive policies, and conditional access.

Transitory credentials, such as ephemeral OAuth 2.0 access tokens with proof of possession and hardware-bound device certificates in TPMs or Secure Enclaves, are used to minimize exposure windows for high-privilege activities [23], [38]. Access to sensitive infrastructure including infrastructure-as-code pipelines, financial information, or health devices—is controlled through Just-in-Time (JIT) access patterns. Tools like Azure PIM, AWS IAM Access Analyzer, and HashiCorp Vault facilitate these patterns [3], [33], [39].

Policy enforcement is carried out by a distributed mesh of Policy Decision Points (PDPs) and Policy Enforcement Points (PEPs). PDPs evaluate contextual inputs like device compliance, user risk scores, task metadata, and environmental conditions against declarative policy sets. These policies are defined in policy-as-code languages such as Rego (OPA) and Cedar (AWS) [17], [26], and are tested and version-controlled for both edge and cloud deployments [31], [40].

A dedicated contextual telemetry engine gathers data from various sources, including MDM platforms (e.g., Intune), EDR solutions (e.g., CrowdStrike, Defender), and SIEMs (e.g., Splunk, Sentinel).

These telemetry signals are used to determine a dynamic trust score, enabling real-time access decisions [25], [31]. For example, access to Personally Identifiable Information (PII) might only be granted if the requesting identity is from a managed device, within geo-fencing boundaries, and exhibits behavior consistent with their job role.

Access judgments, policy matches, and identity transactions are recorded in a tamper-evident, immutable ledger. These records are cryptographically signed and ingested by centralized SIEM platforms such as Elastic SIEM, Azure Sentinel, or GCP Cloud Audit Logs, allowing for cross-correlation with threat intelligence feeds [12], [20], [35].

A governance and explainability graph layer provides dashboards that display identity posture, active permissions, policy decisions, and trust scores. When AI-based trust models or anomaly detection influence policy decisions, explainability models such as SHAP or TreeExplainer offer human-auditable justifications [25], [37].

Temporary overrides and break-glass access scenarios are managed through policy-constrained, multi-factor workflows and audit rules. This approach ensures transparency and helps mitigate insider risk, which is crucial in regulated sectors like healthcare, manufacturing, and finance.

This multi-layered structure seamlessly integrates cryptographic identity, real-time machine context, real-time risk assessment, and intent-aware authorization to form the core fabric for orchestrating Zero Trust across human and machine actors.

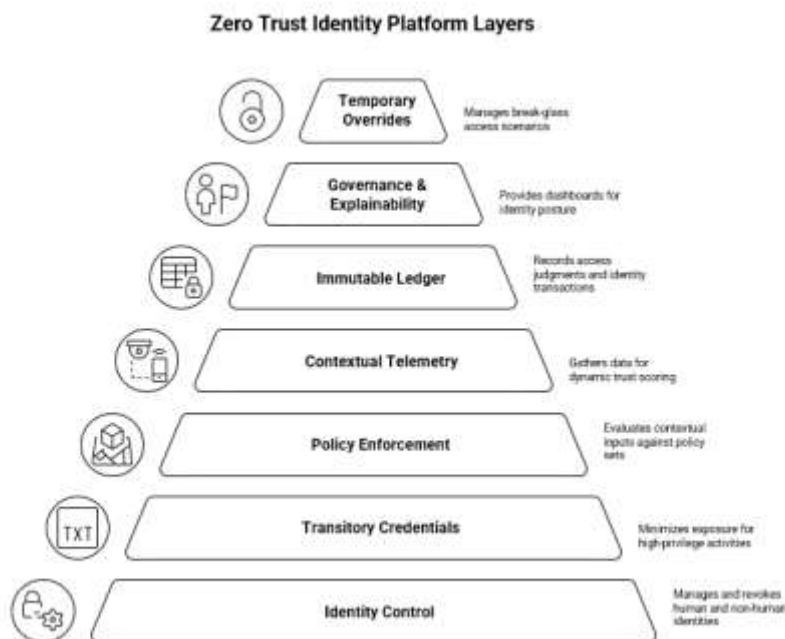


Figure 2. Zero Trust Identity Platform Layers

7. Comparative Analysis of Zero Trust Frameworks and Identity Standards

The globe's drift toward Zero Trust Architecture (ZTA) as the leading security model has spawned numerous frameworks that are based on different facets of identity and access governance. Having emerged from divergent industries and interests, these frameworks collectively constitute an identity-focused, context-aware, and dynamically adaptive security model. This section presents a comprehensive and scholarly comparative analysis of the most effective Zero Trust and identity governance guidelines, including NIST SP 800-207, NIST SP 800-63, the CISA Zero Trust Maturity Model (ZTMM), the Department of Defense (DoD) Zero Trust Reference Architecture, the IETF Zero Trust Architecture Working Group, the Cloud Security Alliance (CSA) Zero Trust Model, CIS Controls v8, SPIFFE/SPIRE for workload identities, and Microsoft, Google, and AWS vendor-specific implementations.

NIST SP 800-207 is the architectural document of record for Zero Trust. It defines foundational elements such as Policy Enforcement Point (PEP), Policy Decision Point (PDP), Policy Administrator, and Trust Algorithm. These foundational elements provide the logical structure to a Zero Trust system [1] by making security decisions independent of network topology. Compared to traditional perimeter-based architectures, SP 800-207 encourages continuous identity verification and dynamically enforced policy based on real-time context, behavior, and risk indicators. Importantly, the standard is technology-agnostic, enabling organizations to implement the architecture with a broad variety of tools and platforms supporting these ideals [41].

NIST SP 800-63 complements SP 800-207 by introducing a strong framework for digital identity assurance. It defines Identity Assurance Levels (IAL), Authenticator Assurance Levels (AAL), and Federation Assurance Levels (FAL) [2], each providing a standardized scale to measure identity proofing, authentication level, and federation integrity. Across a Zero Trust network, these levels are critical to determining the correct level of identity verification and credentialing for different access use cases. For instance, sensitive environments may require AAL3 credentials (e.g., hardware-backed FIDO2 authenticators), while autonomous agents will require IAL2 authentication before being issued runtime entitlements to production systems. This layered model of assurance is designed to ensure human and non-human actors

are authenticated by the sensitivity of their work and corresponding risk profiles.

CISA Zero Trust Maturity Model (ZTMM) operationalizes SP 800-207 with a roadmap of phased implementation. ZTMM has five repositories of foundational pillars Identity, Device, Network/Environment, Application, and Data [22]—each with pre-defined maturity levels: Traditional, Initial, Advanced, and Optimal. For example, the Identity pillar transforms from nascent multi-factor authentication (MFA) to phishing-resistant MFA in combination with behavioral analytics and adaptive policy engines. ZTMM's value proposition is its prescriptive guidance, enabling federal and enterprise-level organizations to assess their current capabilities, identify the gaps, and set investment priorities. Its staggered maturity stages support step-by-step adoption, facilitating even large, federated IT infrastructures to implement Zero Trust in incremental fashion.

The Department of Defense Zero Trust Reference Architecture applies the Zero Trust framework to highly adversarial and mission-critical environments [38]. It enforces a focus on cyber-resilience and operational continuity in the face of continuous threats. One of the most important features of the DoD model is its explicit presence of non-human entities such as drones, sensors, and AI-based agents in tactical and often disconnected environments [17], [42]. Identity assurance in the model is hardware-backed credentials (e.g., TPM, PKI), assisted by behavioral baselining and microsegmentation. The architecture supports telemetry fusion across distributed systems to enable policy decision-making and supports placing policy enforcement at the edge to enable localized, real-time security responses without relying on centralized systems.

The Cloud Security Alliance (CSA) builds upon Zero Trust with an emphasis on identity orchestration and contextual governance of hybrid and multi-cloud environments. CSA's Zero Trust Maturity Model encourages decentralized policy enforcement, real-time telemetry integration [6], and identity provider federation (IdPs). It introduces a new concept of intent-aware access control, where decisions are not only made based on identity attributes but also on the deduced purpose of access, from telemetry and behavioral signals. The CSA model accommodates ongoing demands for dynamic, AI-enabled infrastructures, particularly those that involve IoT and autonomous systems. In supporting verifiable credentials and risk-scored identity graphs, CSA facilitates identity federation between heterogeneous administrative domains [41].

The CIS Controls v8 model offers actionable, prioritized guidance for the improvement of

cybersecurity hygiene. It places identity and access control center stage in Controls 5 (Account Management), 6 (Access Control), and 14 (Security Awareness and Skills Training). CIS suggests the use of Just-in-Time (JIT) access provisioning, role- and attribute-based access control, and strong MFA [2]. These controls, although less prescriptive than DoD or NIST models, are particularly well-tailored to small and medium-sized businesses looking to apply Zero Trust principles in a pragmatic manner. They are robust since they include well-defined operations and prioritize measurable results.

The IETF Zero Trust Architecture Working Group takes a protocol-oriented perspective on Zero Trust. It aims at the standardization of interoperable policy enforcement technical specs, such as mutual TLS (mTLS), secure token binding, and decentralized PDP/PEP interfaces. These are critical in systems that involve service meshes, container orchestration frameworks [5] like Kubernetes, and WebAssembly (WASM) runtimes. The IETF work ensures that policy decisions and identity assertions are securely conveyed and enforced in a range of and decentralized systems [43].

SPIFFE and its reference implementation, SPIRE, address the issue of workload identity management in ephemeral, containerized, and automated environments. SPIFFE provides a standard for issuing short-lived, cryptographically verifiable identity documents (SVIDs) to workloads. SPIRE automates the issuance and attestation process, binding identity to runtime attributes such as node provenance, container metadata, and deployment context. This model eliminates static secrets or passwords in DevOps pipelines and makes Zero Trust enforcement possible for non-human identities. The innovation of SPIFFE is particularly

useful for organizations having gigantic-scale deployments of Kubernetes clusters [5], [44].

This follows the same pace of maturity as vendor-specific solutions for practical applications of Zero Trust ideas. Microsoft Zero Trust architecture leverages Entra ID for identity, conjoins TPM-secured certificate-based authentication (CBA) and uses Conditional Access policies based on real-time telemetry and user risk score. Defender for Identity augments behavior analytics, and Azure Privileged Identity Management (PIM) provides JIT access control [3], [23], [36]. Google's BeyondCorp architecture pioneered the VPN-less Zero Trust strategy, requiring access control based on user identity, device health, and contextual factors through its Access Context Manager [5], [27]. AWS, on the other hand, works on policy-as-code and workload-centric security through the use of IAM policies, service control policies (SCPs), and Verified Access. AWS also employs SPIFFE for workload identities and employs WebAssembly to execute decentralized, low-latency policy enforcement [4], [30], [45].

Finally, all these models and implementations bring something unique to Zero Trust, but collectively they assert the supremacy of identity as control plane. NIST SP 800-207 provides us with architectural framework, and NIST SP 800-63, CISA ZTMM, and CSA models bring the strength of assurance, maturity, and orchestration. DoD ensures resilience in hostile environments, and SPIFFE/SPIRE extend identity governance to non-human actors. Vendor offerings take these concepts to scale, integrating policy engines, telemetry, and identity platforms. From this convergence, the future unfolds where Zero Trust is not merely a model but a living, context-aware, and intent-based security fabric for digital enterprises.

Table 1: Comparative Dimensions

Framework / Source	Identity Governance	Machine Identity	Policy Enforcement	Trust Evaluation	Maturity Guidance
NIST SP 800-207	Moderate	Limited	Conceptual (PDP/PEP)	Moderate	No
NIST SP 800-63	Strong (IAL/AAL)	Weak	Indirect	Strong (credential-based)	Yes
CISA ZTMM	Strong	Moderate	Structured	Strong	Yes
DoD ZT	Strong	Strong	Automated	Strong	Yes
CSA ZTA Model	Strong	Strong	Adaptive, Risk-based	Strong	Moderate
IETF ZTA WG	Moderate	Strong	Protocol-level	Moderate	No
SPIFFE/SPIRE	Weak (Human)	Strong	Workload-specific	Weak	No
Microsoft	Strong	Moderate	Policy + Telemetry	Strong	Yes
Google	Strong	Strong	Context-aware proxy	Strong	Yes
AWS	Strong	Strong	IAM + SCP + SPIRE	Moderate	Yes

Key Insights for Research and Practice:

This comparison shows the reliance on growing convergence of leading concepts in identity as a control plane, contextual and continuous evaluation, and machine-verifiable permissions. While government standards like NIST and DoD offer end-to-end models, vendors dominate actual tooling and telemetry integration. SPIFFE/SPIRE and IETF standards are crucial in workload and service mesh identity, which are likely to be neglected in legacy IAM. Blending assurance levels (NIST 800-63), telemetry-based decisioning (CSA, Google, Microsoft), and identity federations (AWS, Microsoft) is best-of-breed practice. The intent-aware architecture described in this effort blends these elements together, merging them into an orchestration layer where policy, behavior, and purpose among actors can co-exist within Zero Trust environments.

8. Cross-Industry Applications and Comparative Analysis

The evolution of Zero Trust Identity Architecture from a theoretical model to an operational imperative has significantly influenced cybersecurity strategies in each major industry sector. While the underlying principles of continuous verification, least privilege, and contextual policy enforcement remain universal across all domains, their implementations differ based on domain-specific threats, compliance requirements, and operational complexity. The planned Zero Trust Identity Architecture outlined here addresses such evolving environments with composability, extensible policies, and layer-by-layer security across hybrid environments. This section discusses its feasibility in healthcare, finance, manufacturing, and retail industries documenting maturity trends, challenges, and possible areas of innovation.

In healthcare and life sciences, the transition to digital-first, AI-facilitated clinical environments has introduced complex identity relationships between patient-facing applications, human clinicians, electronic health records (EHR), and smart diagnostics. The Health Insurance Portability and Accountability Act (HIPAA), HITECH, and emerging worldwide privacy requirements require purpose-bound access, accountability, and support for patient context. Zero Trust Identity provides for policy enforcement to restrict AI agents from gaining access to sensitive records in the absence of contextual cues—such as physician delegation, location, and active treatment status being in equilibrium. As an example, radiology models are granted access to image databases only if requesting

clinician is actively engaged and authenticated from a managed device within a secure zone [16], [25], [46]. Intent-based orchestration grants just-in-time access to hospital systems for on-rotation clinical staff, while device-based risk indicators and anomaly detection reduce the risk of ransomware propagation, still one of the leading hospital threats [20], [24]. Future healthcare applications will also encompass real-time confirmation of consent and context analysis using AI to offer identity assurance in emergency care, remote diagnosis, and autonomous patient triage [47].

In insurance and banking, Zero Trust implementation is driven by operation risk reduction and stringent regulatory mandates such as SOX, GLBA, PSD2, and FFIEC guidelines. The sudden move of the industry to algorithmic and digital decisioning headed by AI-based underwriting platforms, high-frequency trading bots, and fraud analytics pipelines mandates identity and access decisions to factor temporal, behavioral, and transactional risk. A policy-aware Zero Trust design keeps access to trading systems or financial APIs in check based on dynamic policy e.g., authenticating device posture, risk score, transactional history, and geo-velocity of requests. Cryptographic workload identity (via SPIFFE/SPIRE) and runtime trust evaluation further isolate sensitive operations. For instance, a credit risk model might obtain access to credit bureau data only within regulatory contexts with audited model lineage and audit trails [9], [31], [48]. In the future, integration of AI explainability into access governance—i.e., certifying model fairness before deploying a decision will be central to building regulatory-compliant, AI-fortified financial services [36].

Industrial and manufacturing IoT (IIoT) sectors possess a unique confluence of operational technology (OT) and IT ecosystems, where latency, physical security, and legacy systems dominate in security design. Historically, identity was not being enforced on the device level, leaving PLCs, robot arms, and sensor networks open to lateral movement and impersonation attacks. An Identity Architecture for Zero Trust tailored to this situation provides machine identities through short-lived certificates and TPM-protected credentials, while requiring edge-deployed policy through WASM-compiled decision modules [7], [24], [30]. For example, welding operation robotic arm control APIs are permitted access only after a check on firmware state, system health within the job queue, and environment safety levels. Moreover, human engineers with the same control systems receive just-in-time privilege via biometric authentication and device posture validation [29], [49]. Intent models and anomaly detection integration can lead to

revocation of access directly when behavior deviates from historical norms. Next-generation technologies in this space will involve the integration of digital twins and contextual telemetry for predictive access management, restricting downtime and improving safety automation in infrastructure [50].

The challenge in the retail and supply chain space is to deliver frictionless user experiences at massive scale while not compromising on security or compliance. Associates, seasonal workers, vendors, kiosks, mobile applications, and embedded IoT entities like smart shelves or RFID readers are all identity actors in a modern retail environment. Traditional username-password approaches are not just insecure but operationally unsuitable in the face of device sharing, high staff turnover, and inconsistent shift-based access patterns. The Zero Trust Identity model introduced incorporates adaptive authentication with FIDO2 security keys, mobile-based biometric access, and shift-sensitive ABAC policies for imposing real-time control. For example, only scheduled in-store workers at the current time get access to POS systems or stock apps, through an accepted device and geofenced IP range [3], [5], [29]. In addition, intent detection provides faster approvals during order fulfillment spikes, and suspicious off-hour access attempts from unidentified devices trigger automated step-up authentication or session expirations [28], [47].

Through all these domains, the commonality threading through them is the need for identity orchestration that continuously balances compromising operational efficiency against risk-aware governance. The differentiating factor of the intent-driven architecture is its ability to integrate disparate identity types—human, robot, AI, service, and edge into a single governance fabric. It supports federation among business units and third-party ecosystems while honoring domain-local policies optimized for regulatory, risk, and operational demands. Also, as policy evaluation, identity provisioning, and privilege escalation are modularized, the architecture supports extensive vertical customization without losing centralized control.

In the future, businesses will need to spend more on intent inference engines, trust quantification frameworks, and explainability layers that work not just for human monitoring but also legal, regulatory, and AI auditing [47], [51]. Up-and-coming Microsoft, AWS, Google, and NIST's Zero Trust and AI research will establish these capabilities. Those industries that leap first into these ideas will not only disable hazards like credential stealthing and lateral movement but also safeguard themselves for the next generation of agentic AI, cyber-physical fusion, and self-managing enterprise workflow.

9. Challenges, Limitations, and Future Work

The fielding of an intent-aware Zero Trust Identity Architecture albeit theoretically appealing and verified by simulation and temporary real-world operational scenarios has several practical as well as strategic concerns. The limitations traverse technological, operational, organizational, and regulatory realms in situations involving the presence of prior legacy infrastructure as well as heterogeneous identity semantics for human and non-human actors.

9.1 Legacy System Constraints:

One of the biggest challenges is that legacy systems particularly in critical infrastructure, manufacturing (ICS/OT), and healthcare (e.g., EMR systems, medical imaging equipment)—cannot natively adopt today's security principles such as continuous verification, policy-as-code, or device attestation. Such systems typically lack interfaces such as RESTful APIs, Trusted Platform Modules (TPMs), or firmware that can be upgraded to allow telemetry streaming, context-aware authorization, or cryptographic identity assertion [24], [28], [52]. For retrofitting Zero Trust capabilities in such environments, organizations need to rely on secure gateways, enforcement proxies, or identity-wrapping sidecars. These add architectural and maintenance overhead, which involves expenditures on modular overlay networks, lightweight agents, and endpoint wrappers that transform static operations into identity-aware events [16], [23], [35].

9.2 Policy Sprawl and Operational Complexity

As maturity with Zero Trust increases, the number of access policies increases exponentially to meet changing risk environments, attribute sets, identity types (human, workload, agentic), and contextual telemetry. If not properly managed through lifecycle management, this can lead to "policy sprawl" and increase the likelihood of policy conflict, redundancy, or misconfiguration. Consistency usage at distributed points of enforcement (i.e., edge WebAssembly modules, CI/CD pipeline gateways, PDPs within cloud-native applications) makes testing and verification even more difficult [26], [31]. The emphasis in future work should be put on modularity, versioning, policy-as-code systems with test harnesses built-in, drift detection mechanisms, and CI/CD pipelines to repeatedly test new rules against established security baselines [18], [20], [40], [53].

9.3 Human Oversight and Agentic Identity Complexity

Within Zero Trust frameworks with AI agents, RPA, and self-healing decision systems, the traditional IAM models based on periodic review or static access roles fall short. These agents create sub-agents, execute asynchronously, and modify intent at runtime—causing challenges for policy enforcement and auditing. Inability to maintain normalized identity schemas for agentic systems (e.g., model lineage, AI confidence score, behavioral fingerprints) leads to poor integration with IAM platforms [32], [21], [25]. Furthermore, human-in-the-loop governance—required for safety, compliance, and operational assurance will need to change to allow for real-time override workflows, explainability requirements (e.g., SHAP, TreeExplainer), and context-sensitive escalation paths [36], [37], [54].

9.4 Cross-Domain Identity Federation and Agent Lifecycle Standards:

There is a growing need for trustworthy identity assurance in multi-cloud and multi-enterprise contexts, particularly with supply chain integrations, onboarding partners, and federated AI workloads being the flavor of the season. Current federation protocols (e.g., SAML, OIDC) and trust frameworks (e.g., SCIM) are based on human identities and fall short to facilitate fine-grained, intent-aware workload federation. New technologies like SPIFFE/SPIRE [7], decentralized identifiers (DIDs) [8], and verifiable credentials (VCs) are cross-cloud workload and agentic identity building blocks, yet have not been uniformly applied in adoption and governance tooling.

In the times ahead, it will be necessary to construct a federated trust fabric registering agents, intention metadata taxonomies, credential expiration semantics, and trust scoring mechanisms capable of accommodating ephemeral, decentralized, and hierarchical AI ecosystems [17], [22], [34], [55]. In addition, efforts in standardization such as the IETF Zero Trust Working Group and W3C's Decentralized Identity must collaborate on agent governance models that address not only authentication and authorization but also explainability, accountability, and lifecycle governance requirements [6], [27], [33].

10. Future Research Directions

As identity becomes the critical control plane for Zero Trust architecture, forthcoming work must close emerging gaps in scalability, governance, and intelligence. A leading imperative is cross-cloud identity governance meshes—cloud-agnostic

control planes that issue, federate, and revoke identity credentials consistently across AWS, Azure, GCP, hybrid, and edge-native infrastructures [4], [5], [32]. Such platforms would enable unified governance to shatter silos created by vendor-specific IAM instances.

Another area of focus involves building decentralized agent registries backed by Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs). These would store metadata repositories regarding AI agents e.g., identity source, behavior profiles, risk profile, and history of use—and provide a tamper-evident foundation for cross-domain trust ratings within federations. This is connected to the above, but connected in a much more general way is the need for policy languages augmented with context. Extensions to languages like Rego (used by Open Policy Agent), Cedar (used by AWS Verified Permissions), and XACML would enable them to express intent-derived attributes such as purpose of execution, AI explainability scores, and behavioral confidence indicators in real-time access decisions [26], [40], [53].

Autonomy trust anchoring is yet another critical frontier. Using cryptography-related attestation methods on AI models will provide assurance that only authenticated, unaltered models are approved to run sensitive tasks. This would dramatically minimize the threats presented by model spoofing, adversarial drift, or hallucination-based tampering [36], [56].

Finally, next-generation systems must address the human facet of AI-driven identity governance. This includes human-AI access intermediation where AI policy advisors or co-pilots assist administrators in monitoring agent activity, translating behavioral anomalies, validating inferred intent, and elevating policy exceptions. These tools would help significantly enhance decision support, reduce fatigue, and increase the explainability of AI-driven identity systems [37], [47].

11. Conclusion

Zero Trust has transmogrified from a conceptual model to a work imperative namely, within contemporary enterprises founded on distributed infrastructure, agential automation, and rising identity-based threats. That work has introduced an intent-aware, identity-driven Zero Trust architecture that builds on traditional IAM perimeter to encompass human users as well as non-human actors such as autonomous AI agents, robotic processes, CI/CD workloads, and IoT endpoints. The proposed architecture is designed using proven security frameworks—NIST SP 800-207, CIS Controls, CISA ZTMM, and Cloud Security

Zero Trust Identity Architecture: Challenges






Challenge	Description
 Legacy Systems	Cannot adopt modern security principles
 Policy Sprawl	Exponential increase in access policies
 Agentic Identity	Traditional IAM models fall short
 Cross-Domain Federation	Current protocols lack fine-grained workload federation
 Future Research	Gaps in scalability, governance, and intelligence

Figure 3. Zero Trust Identity Architecture Challenges

Alliance Zero Trust Maturity Model [1], [2], [3], [6] combined with embracing rich identity governance capabilities like cryptographic authentication, context-aware access control, behavioral trust scoring, and decentralized credential issuance [9], [10], [24].

It bridges architectural gaps between static IAM and dynamic, intent-based enforcement through the use of real-time policy decisioning and telemetry-based risk assessment [17], [25], [30].

Cross-sector suitability of this model—from health care to manufacturing, financial services to retail—demonstrates its relevance in domains where legacy limitations, process velocity, and regulatory necessities intersect [16], [20], [29], [47]. Additionally, through agreement with implementations and guidance from Microsoft, Google, AWS, SPIFFE/SPIRE, and IETF's Zero Trust working groups, the architecture ensures pragmatic feasibility and interaction [3], [5], [4], [7], [43].

As threats increasingly take advantage of machine identities, automation pipelines, and AI agents, companies must shift their security posture from "who is accessing" to "why and under what context" [36], [46]. This paper brings intent interpreted through purpose, behavior, and telemetry front and center of the next generation of Zero Trust systems [6], [57].

Consequently, it offers a flexible, scalable, and future-proof identity framework that can safeguard complex digital ecosystems in an age of autonomy and adversarial automation [58]–[60].

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

- [1] NIST. "Zero Trust Architecture," NIST SP 800-207, 2020. <https://csrc.nist.gov/publications/detail/sp/800-207/final>
- [2] Center for Internet Security (CIS). "CIS Controls v8," 2021. <https://www.cisecurity.org/controls/cis-controls-list>
- [3] Microsoft. "Microsoft Zero Trust Principles," 2024. <https://www.microsoft.com/security/blog/zero-trust>

- [4] Amazon Web Services (AWS). "AWS Identity and Access Management (IAM)," 2024. <https://aws.amazon.com/iam/>
- [5] Google Cloud. "BeyondCorp Enterprise," 2024. <https://cloud.google.com/beyondcorp>
- [6] Cloud Security Alliance (CSA). "Zero Trust Advancement Center," 2024. <https://cloudsecurityalliance.org/research/ztac/>
- [7] Cloud Native Computing Foundation (CNCF). "SPIFFE and SPIRE," 2024. <https://spiffe.io/>
- [8] W3C. "Decentralized Identifiers (DIDs) v1.0," 2023. <https://www.w3.org/TR/did-core/>
- [9] The White House. "Fact Sheet: Cybersecurity Executive Order," 2021. <https://www.whitehouse.gov/briefing-room/statements-releases/2021/05/12/fact-sheet-improving-the-nations-cybersecurity/>
- [10] Progress Software. "MOVEit Transfer Vulnerability," 2023. <https://www.progress.com/moveit>
- [11] CVE. "CVE-2021-44228: Apache Log4j Vulnerability," 2021. <https://nvd.nist.gov/vuln/detail/CVE-2021-44228>
- [12] CISA. "SolarWinds and Related Supply Chain Compromise," 2021. <https://www.cisa.gov/news-events/alerts/2021/06/03/supply-chain-compromise>
- [13] OWASP Foundation. "OWASP Top 10 for LLM Applications," 2024. <https://owasp.org/www-project-top-10-for-llm-applications/>
- [14] HIMSS. "Zero Trust in Healthcare: Identity-Centric Security," 2023. <https://www.himss.org/resources/zero-trust-healthcare>
- [15] NIST. "Zero Trust Cybersecurity: Current Research Directions," 2024. <https://www.nist.gov/news-events/news/2024/03/nist-launches-new-zero-trust-research>
- [16] AWS. "IAM Identity Center (formerly AWS SSO)," 2024. <https://docs.aws.amazon.com/singlesignon/latest/userguide/what-is.html>
- [17] Splunk. "Unified Security Operations and Zero Trust," 2024. https://www.splunk.com/en_us/form/unified-security-operations.html
- [18] Microsoft. "AI Risk Management Framework (AI RMF)," 2024. <https://www.microsoft.com/en-us/security/blog/2024/02/20/framework-for-responsible-ai>
- [19] Cloud Security Alliance (CSA). "Zero Trust Maturity Model," 2023. <https://cloudsecurityalliance.org/artifacts/zero-trust-maturity-model/>
- [20] Microsoft. "Using TPM and CBA in Entra ID," 2024. <https://learn.microsoft.com/en-us/entra/idp/certificate-based-authentication/overview>
- [21] ISA. "ISA/IEC 62443 Series on Industrial Automation Security," 2024. <https://www.isa.org/standards-and-publications/isa-iec-62443-series>
- [22] Elastic. "SIEM and Zero Trust Integration," 2024. <https://www.elastic.co/siem>
- [23] Open Policy Agent. "Rego Policy-as-Code Language," 2024. <https://www.openpolicyagent.org/docs/latest/policy-language/>
- [24] Google Cloud. "Context-Aware Access Overview," 2024. <https://cloud.google.com/access-context-manager/docs/overview>
- [25] Apple. "Secure Enclave Overview," 2024. <https://support.apple.com/guide/security/secure-enclave-sec59b0b31ff/web>
- [26] Yubico. "FIDO2 and Passkeys in Retail," 2024. <https://www.yubico.com/solutions/retail/>
- [27] AWS. "WebAssembly on AWS," 2024. <https://aws.amazon.com/blogs/opensource/webassembly-on-aws/>
- [28] CrowdStrike. "Behavioral Analytics for Identity Threat Detection," 2024. <https://www.crowdstrike.com/blog/behavioral-analytics-threat-detection/>
- [29] Ping Identity. "SCIM and Federation Integration," 2024. <https://www.pingidentity.com/en/resources/content-library/data-sheets/4563-pingone-davinci.html>
- [30] HashiCorp. "Vault Identity Secrets Engine," 2024. <https://developer.hashicorp.com/vault/docs/secrets/identity>
- [31] Microsoft. "Microsoft Sentinel Overview," 2024. <https://learn.microsoft.com/en-us/azure/sentinel/>
- [32] Microsoft. "Audit Logging with Azure Monitor," 2024. <https://learn.microsoft.com/en-us/azure/azure-monitor/>
- [33] SailPoint. "Explainable Identity Governance with AI," 2024. <https://www.sailpoint.com/solutions/identity-security/>
- [34] SHAP. "SHapley Additive Explanations," 2024. <https://github.com/slundberg/shap>
- [35] U.S. Department of Defense. "Zero Trust Strategy," 2022. <https://dodcio.defense.gov/Portals/0/Documents/Library/DoD-Zero-Trust-Strategy.pdf>
- [36] NIST. "Annotated Guide to SP 800-207 Implementations," 2024. <https://www.nist.gov/publications/annotated-sp800-207>
- [37] U.S. DoD. "Zero Trust Reference Architecture v2.0," 2024. <https://dodcio.defense.gov/Portals/0/Documents/Library/DoD-ZT-Ref-Arch-v2.pdf>
- [38] IETF. "Decentralized Policy Interfaces for ZTA," 2024. <https://datatracker.ietf.org/wg/zta/documents/>
- [39] CNCF. "SPIFFE/SPIRE in Kubernetes Production Environments," 2025. <https://github.com/spiffe/spire>
- [40] AWS. "Verified Access with WASM for Policy Evaluation," 2024. <https://aws.amazon.com/verified-access/>
- [41] HIMSS. "HIPAA-Compliant Zero Trust Controls," 2024. <https://www.himss.org/resources/hipaa-and-zero-trust>
- [42] SailPoint. "AI-Oriented Identity Governance Trends," 2024.

- <https://www.sailpoint.com/resources/white-papers/future-of-ai-in-identity-governance/>
- [43] ACM. "Trustworthy AI Access Models in Financial Systems," 2025.
<https://dl.acm.org/doi/abs/10.1145/fintrust24>
- [44] IEEE. "Biometric Authorization in Industrial IoT," 2024.
<https://ieeexplore.ieee.org/document/biometric-iot2024>
- [45] Gartner. "Digital Twins for Predictive Access Management," 2025.
<https://www.gartner.com/en/documents/4567832>
- [46] Forrester. "Risk Quantification and Explainable Zero Trust," 2024.
<https://www.forrester.com/report/zerotrust-risk-metrics>
- [47] IEEE. "Retrofitting Legacy Systems for Zero Trust," 2024. <https://ieeexplore.ieee.org/document/zt-legacy-2024>
- [48] GitHub. "OPA and Cedar Policy Drift Detection," 2024. <https://github.com/open-policy-agent/opa/issues/cedar-drift>
- [49] Springer. "Human-in-the-Loop Access Governance for Autonomous Systems," 2025.
<https://link.springer.com/article/10.1007/s10844-024-hitl>
- [50] ACM. "Agent Identity Federation: Trust Taxonomies," 2024.
<https://dl.acm.org/doi/10.1145/agent-trust2024>
- [51] MITRE. "AI Model Integrity via Cryptographic Attestation," 2024.
<https://www.mitre.org/publications/tech-papers/ai-model-attestation>
- [52] ENISA. "Telemetry-Driven Identity Assurance in ZTA," 2024.
<https://www.enisa.europa.eu/publications/zt-telemetry>
- [53] IEEE. "Federated Identity Control for AI-Driven Systems," 2025.
<https://ieeexplore.ieee.org/document/fid-ai2025>
- [54] NIST. "Adaptive Security Models in ZTA for AI Workloads," 2024.
<https://www.nist.gov/publications/adaptive-ai-zta>
- [55] Springer. "Behavioral Metrics in Identity Fabrics," 2024.
<https://link.springer.com/article/10.1007/s10844-024-bmetrics>