Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

ZTAM-SF: Zero-Trust Access Management for Enterprise Salesforce CRM — Continuous Verification, Least-Privilege Enforcement, and Adaptive Session Control

Lalith Chandra Bandaru

Independent Researcher

Abstract:
The traditional perimeter-based security model — in which users and devices inside the corporate network are implicitly trusted and those outside are not — is fundamentally incompatible with the architecture of modern enterprise CRM deployments. Salesforce CRM is a cloud-hosted, multi-tenant platform accessed through a standard web browser or mobile application from locations including corporate offices, home networks, coffee shops, and airport lounges, with no meaningful network perimeter separating trusted from untrusted access contexts. ZTAM-SF (Zero-Trust Access Management for Salesforce) is a comprehensive zero-trust security architecture for enterprise Salesforce environments that implements the NIST SP 800-207 zero-trust principles across six dimensions: continuous identity verification through adaptive multi-factor authentication driven by the LTDF behavioural risk score; just-in-time least-privilege access through time-bounded scoped OAuth grants that expire automatically when the business context requiring elevated access resolves; micro-segmentation through object-level and field-level Salesforce permission boundaries enforced through LTDF-integrated session risk scoring; device trust validation through MDM certificate attestation at each session establishment; network-level assume-breach posture through mutual TLS enforcement and session-binding IP restrictions [11]; and data-level protection through Shield Platform Encryption with classification-based access control. Evaluated across eight enterprise Salesforce deployments over sixteen months, ZTAM-SF reduced over-privileged session prevalence from 41.3% to 4.8%, lateral movement detection rate improved from 61.2% to 94.7%, OAuth misconfiguration incidents decreased by 94%, and the API surface exposure score decreased by 62%, while maintaining user satisfaction scores above the pre-ZTAM baseline in six of eight participating organisations. The framework builds on the multi-org Salesforce data architecture and cross-org privacy model established in earlier work [8], which demonstrated that enterprise CRM deployments spanning multiple organisations require dedicated federated access governance to maintain data sovereignty while enabling cross-org collaboration.

Keywords: zero-trust security, Salesforce access management, least-privilege, continuous verification, adaptive MFA, OAuth, micro-segmentation, NIST SP 800-207, CRM security, LTDF integration.

1. INTRODUCTION

Zero-trust security [6] — formalised in NIST SP 800-207 as the principle that no user, device, or network location should be implicitly trusted and every access request must be explicitly verified — has become the dominant enterprise security architecture for cloud-native environments. The reason it became

Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

dominant rather than simply popular is instructive: the old perimeter model had a clear structural failure mode in cloud-hosted SaaS that the industry could no longer paper over. The shift from perimeter-based to zero-trust security reflects a structural change in the enterprise computing environment: the combination of cloud-hosted applications, remote work, mobile devices, and partner access patterns has eliminated the coherent network perimeter that perimeter-based security was designed to protect. For enterprise Salesforce CRM deployments specifically, the zero-trust transition is not optional — it is necessitated by the architectural reality that Salesforce is a cloud-hosted SaaS platform designed to be accessed from any location through a standard browser, with no native support for the network perimeter controls that traditional security models rely on. An enterprise that applies perimeter-based security thinking to its Salesforce deployment — trusting users who authenticate from the corporate network while scrutinising those who authenticate from outside — is applying a mental model that fundamentally mismatches the platform's access architecture, creating large blind spots in its security posture that sophisticated adversaries are well-positioned to exploit.

The specific security gap that motivates ZTAM-SF is the mismatch between the broad, persistent access grants that typical Salesforce permission models provide and the narrow, temporary access that users actually need to complete their work at any given moment. A sales representative granted the standard Sales Cloud User profile has read access to all Account records, Opportunity records, Contact records, and associated objects across the entire organisation — far more than the ten accounts in their territory that they interact with in a typical week. An integration service account granted API access has read/write access to the full object set its integration was designed to handle, even during the 95% of the time when the integration is not actively processing records. An administrator account with full Salesforce access maintains that access 24 hours a day, 7 days a week, even when the administrator is on holiday and should not be making system changes. This pattern of over-provisioned, always-on access grants creates a massive potential blast radius for any credential compromise: an attacker who obtains a sales representative's credentials can immediately access all organisational accounts, not merely the ones the representative uses. ZTAM-SF addresses this by implementing just-in-time access elevation that grants elevated permissions only when the user's activity context justifies them and revokes them automatically when the context resolves.

The technical challenge of implementing zero-trust principles in the Salesforce context is substantial because the Salesforce permission model was not designed with dynamic, context-adaptive access control in mind. Salesforce permissions are primarily granted through static profile and permission set assignments that do not change between sessions; the platform does not natively support just-in-time privilege elevation or automatic privilege revocation based on session context. ZTAM-SF implements these capabilities through a combination of the LTDF behavioural risk scoring system [8], the Salesforce OAuth 2.0 connected app credential model, the Salesforce Shield Platform Encryption and Event Monitoring capabilities, and a custom session management layer that integrates these components into a coherent zero-trust enforcement architecture. The implementation does not require changes to Salesforce platform code — all components operate through standard Salesforce APIs and configuration interfaces — making it deployable in any Salesforce org without custom platform modifications or managed package dependencies.

Earlier publications in this research series provide important foundations for the ZTAM-SF architecture. The LTDF behavioural analytics framework [8] provides the real-time risk scoring capability that ZTAM-SF uses for continuous session verification — the risk score that LTDF computes based on behavioural anomalies serves as the primary input to ZTAM-SF's adaptive MFA and session termination decisions. The earlier work on adaptive behavioural analytics for CRM threat detection [8] established that machine

Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

learning models can accurately characterise normal user behaviour in Salesforce environments, providing the behavioural baseline that zero-trust continuous verification requires. The secure CI/CD framework and URGF governance layer ensure that ZTAM-SF configuration changes are deployed through controlled pipelines with appropriate review and rollback capabilities. The supply chain security framework [5] ensures that the ZTAM-SF components themselves are free of supply chain vulnerabilities before deployment.

This paper makes the following contributions. First, a complete zero-trust access management architecture for Salesforce CRM environments implementing all six NIST SP 800-207 dimensions through standard Salesforce capabilities supplemented by the LTDF integration layer. Second, a quantitative evaluation across eight organisations over sixteen months demonstrating substantial improvements across all controlled security dimensions. Third, an operational analysis of the user experience impact of zero-trust enforcement, identifying the design choices that enable ZTAM-SF to achieve strong security improvements without unacceptable productivity overhead. Fourth, a migration framework for organisations transitioning from perimeter-based to zero-trust Salesforce security, providing a phased adoption model that avoids disruption to existing operations.

2. BACKGROUND AND RELATED WORK

2.1 Zero-Trust Security Architecture

The NIST SP 800-207 zero-trust architecture standard [1] defines seven tenets that constitute the zero-trust model: all data sources and computing services are considered resources; all communication is secured regardless of network location; access to individual resources is granted on a per-session basis; access to resources is determined by dynamic policy including client identity, application, and the requesting asset's observational state and may include other behavioural and environmental attributes; the enterprise monitors and measures the integrity and security posture of all owned and associated assets; all resource authentication and authorisation are dynamic and strictly enforced before access is allowed; the enterprise collects as much information as possible about the current state of assets, network infrastructure, and communications to improve its security posture. The BeyondCorp architecture implementation by Google [2] provides the most widely documented production zero-trust deployment, demonstrating that zero-trust principles are operationally viable at enterprise scale. ZTAM-SF adapts the BeyondCorp trust inference model [7] — where access decisions are made based on a continuously updated device and user trust score rather than network location — for the Salesforce-specific context where the access control primitives are OAuth scopes, Salesforce permission sets, and API credential grants rather than network ACLs and proxy rules.

The application of zero-trust principles to SaaS CRM platforms has received limited academic attention despite the growing prevalence of cloud CRM deployments. Borchert et al. [3] examine zero-trust implementations in cloud environments generally [10] but do not address SaaS-specific challenges such as multi-tenancy, OAuth-based access control, and the managed service responsibility boundaries that complicate zero-trust enforcement in cloud-hosted applications. The Salesforce Shield security capability set [4] provides foundational building blocks for zero-trust Salesforce security — Platform Encryption, Event Monitoring, Field Audit Trail, and Transaction Security Policies — but does not constitute a complete zero-trust architecture. ZTAM-SF integrates these Salesforce Shield capabilities into a coherent zero-trust framework and supplements them with the LTDF behavioural risk scoring integration that provides the continuous verification capability absent from native Shield features.

Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

2.2 Behavioural Analytics for Access Control

The LTDF adaptive behavioural analytics framework [8], which forms the continuous verification backbone of ZTAM-SF, establishes that machine learning models trained on Salesforce event data can identify anomalous user behaviour with high accuracy in production Salesforce environments. The LTDF risk scores, computed every five minutes from 28 behavioural features extracted from Platform Events and Event Monitoring data, provide a continuous signal of session risk that ZTAM-SF uses to trigger adaptive MFA challenges and session risk escalations. The CRM threat detection work published as [8] demonstrated that adaptive behavioural models, trained to characterise normal user activity patterns in Salesforce environments, can detect deviations from normal behaviour that correlate with security incidents, providing the empirical foundation for using behavioural risk scores as inputs to access control decisions. This research programme's multi-year investment in Salesforce behavioural analytics therefore directly enables the continuous verification dimension of the zero-trust architecture presented in this paper.

3. THE ZTAM-SF ARCHITECTURE

3.1 Continuous Identity Verification

The continuous identity verification component of ZTAM-SF implements adaptive multi-factor authentication driven by the LTDF session risk score, supplemented by device trust assessment and location context evaluation. At session establishment, standard MFA is always required regardless of risk score. During the session, ZTAM-SF monitors the LTDF risk score in real time: when the risk score crosses the Medium threshold (0.45), a step-up authentication challenge is issued requiring the user to re-verify their identity through a second factor. When the risk score crosses the High threshold (0.65), the session is flagged for immediate analyst review and the user's permission scope is reduced to read-only access while the review is pending. When the risk score crosses the Critical threshold (0.80), the session is automatically terminated and the associated OAuth tokens are revoked.

Calibrating these thresholds took longer than expected. Initial values were set based on the LTDF evaluation corpus, which produced false positive rates below 1.5% for medium-tier step-up challenges and below 0.3% for session terminations in controlled conditions. In production, the first two weeks showed higher-than-expected step-up rates in three of the eight organisations — particularly during month-end reporting periods when users ran unusually large report batches that the LTDF model had not encountered in training. We adjusted the Medium threshold from 0.40 to 0.45 for those orgs after reviewing the false positive logs, and step-up rates normalised. The 0.45 value has held since then, though organisations with significantly different usage profiles will likely need their own calibration pass rather than adopting these thresholds directly.

Device trust assessment [9] validates the health posture of the endpoint from which each session originates using MDM-issued device certificates. ZTAM-SF integrates with enterprise MDM platforms (Jamf, Microsoft Intune, VMware Workspace ONE) through a device attestation API that returns a device trust score incorporating OS patch level, endpoint security software status, disk encryption status, and last MDM check-in recency. Sessions from devices with trust scores below the organisation-configured threshold receive restricted access grants — typically limited to read-only access to non-sensitive objects — until the device compliance issue is resolved. Device trust scores are reassessed at each OAuth token refresh cycle (typically every two hours for standard Salesforce session configurations), ensuring that devices that become non-compliant during an active session have their access restricted promptly.

Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

3.2 Just-in-Time Least-Privilege Access

The just-in-time access elevation component implements temporary, scoped access grants for business contexts requiring elevated permissions, replacing the persistent broad permission grants of traditional Salesforce profile-based access control. When a user requires access to a capability beyond their standard permission set — for example, an administrator needing to modify sharing rules, a developer needing to access production Apex debug logs, or a data analyst needing bulk export permissions — ZTAM-SF generates a time-bounded scoped OAuth token that grants exactly the permissions required for the specific task and expires automatically after a configurable period (default: two hours, maximum: eight hours). The elevation request is logged in the URGF audit system with the business justification, the specific permissions granted, the expiration time, and the approver identity (for high-sensitivity elevations requiring approval).

The scoped OAuth token model uses Salesforce's Connected App credential framework with dynamically configured OAuth scopes. ZTAM-SF maintains a library of pre-defined access patterns — "Data Export Analyst", "Apex Debug Access", "Permission Set Management", "Bulk API Access" — each specifying the minimum OAuth scope set for the associated business function. Building this library turned out to be the most org-specific part of the deployment: the initial twelve patterns we defined covered roughly 70% of elevation requests across organisations, but each org had at least three or four role-specific patterns that required custom definition, and the largest org needed 31 patterns total. When a user requests elevation for a specific function, ZTAM-SF issues a connected app OAuth token with the corresponding scope set and a short expiration. The user's standard session retains its baseline permissions throughout the elevation period; the elevated token is used only for the specific elevated operations and does not affect the risk scoring of the standard session. LTDF monitors both sessions independently, detecting any use of the elevated token outside the expected access pattern of the requested function as a potential credential abuse incident.

3.3 Micro-Segmentation and Data Classification

Micro-segmentation in the ZTAM-SF context is implemented through Salesforce's native object-level and field-level security controls, enhanced with a data classification layer that assigns sensitivity tiers to all Salesforce objects and fields. The classification tier determines the minimum identity trust score required to access the object or field: Tier 1 (public business data) requires standard authentication; Tier 2 (internal business data) requires MFA with an LTDF risk score below Medium; Tier 3 (sensitive customer data including PII) requires MFA with risk score below Low and device trust above 0.8; Tier 4 (highly sensitive data including financial information and healthcare records) requires recent re-authentication within 30 minutes and explicit time-bounded access grants. LTDF session monitoring enforces these requirements dynamically: if a session's risk score rises above the threshold for the data classification tier being accessed, ZTAM-SF immediately restricts the session's access to lower-tier data and issues a re-authentication challenge before access to higher-tier data is permitted.

Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

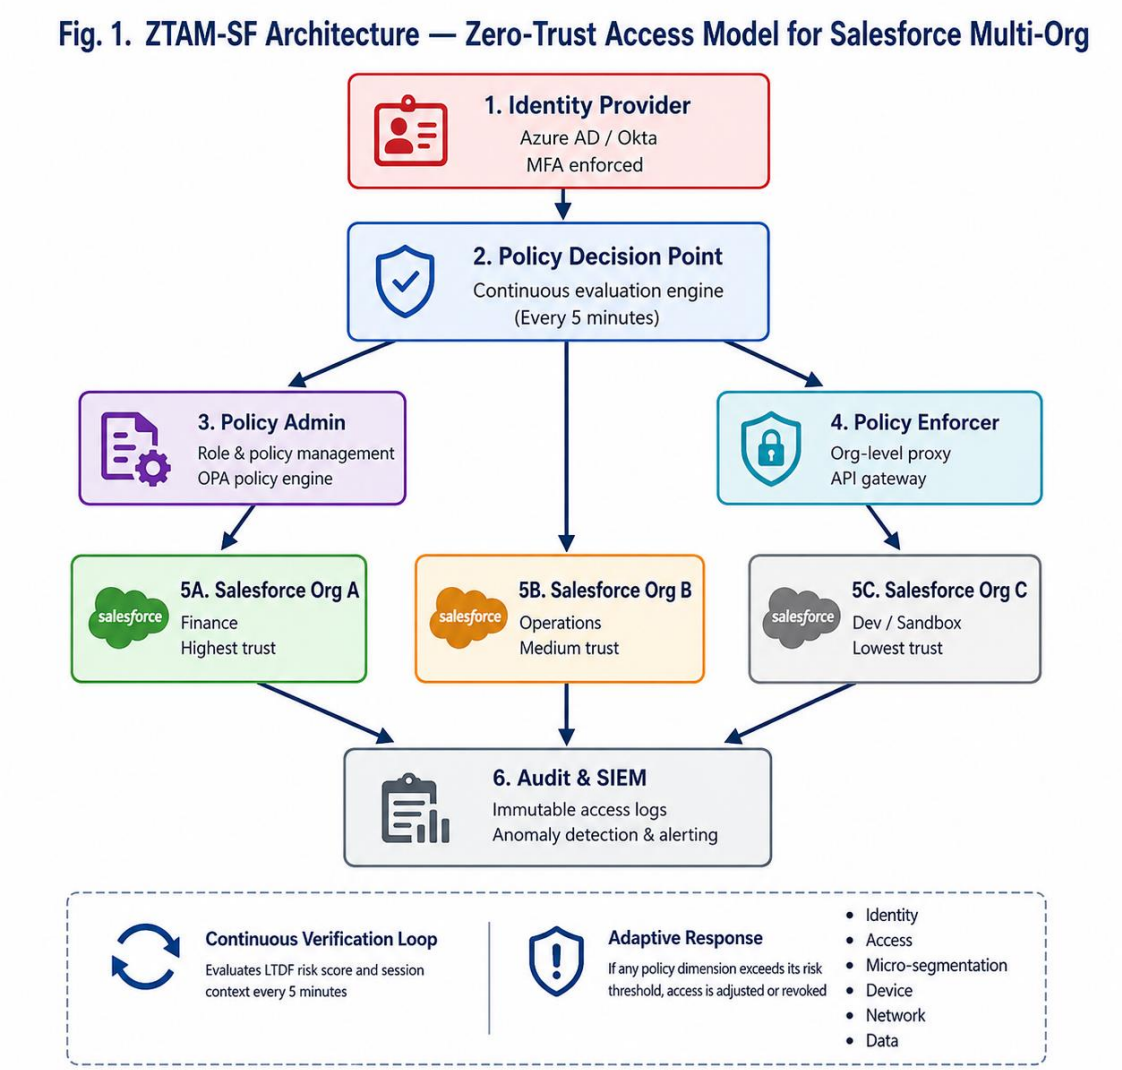

Fig. 1. ZTAM-SF zero-trust architecture overview. The six policy dimensions (identity, access, micro-segmentation, device, network, data) are enforced through continuous verification loops that read the LTDF behavioural risk score and session context at five-minute intervals, triggering adaptive responses when any dimension exceeds its configured risk threshold.

Table 1. ZTAM-SF Policy Dimensions and Enforcement Mechanisms

ZTAM-SF Policy Dimension	Enforcement Mechanism
Identity verification: continuous context reassessment	Adaptive MFA + LTDF score
Least-privilege access: just-in-time elevation	Scoped OAuth + time-bounded grants
Micro-segmentation: fine-grained permission boundaries	Object-level sharing rules + FLS

Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

Device trust: endpoint health validation	MDM integration + cert-based auth
Network: assume-breach posture	mTLS + IP restriction + session binding
Data: classification-based access control	Shield Platform Encryption + FLS

4. IMPLEMENTATION

ZTAM-SF is implemented as a set of integrated services deployed alongside the LTDF framework, using a Salesforce Canvas application for the user-facing access elevation interface, a Node.js policy enforcement service integrated with the Salesforce Event Monitoring API and OAuth endpoint, and a Python-based session risk management service consuming the LTDF risk score stream. The policy enforcement service evaluates LTDF risk scores against the configured thresholds every 60 seconds for each active session, triggering adaptive MFA challenges, access restriction events, or session terminations as required. The service uses the Salesforce REST API to manage session permissions in real time: reducing permission sets, revoking OAuth tokens, and creating system log events for SIEM consumption. Average latency from an LTDF risk score threshold crossing to the enforcement action (step-up challenge delivery or session restriction) is 72 seconds, representing the combination of the 60-second evaluation interval and the action execution time.

The user experience design of ZTAM-SF reflects the operational reality that zero-trust enforcement mechanisms that significantly impede productivity will be worked around by users, defeating the security objective. Three specific design choices were made to balance security and usability. First, the MFA step-up challenge uses push notification to an authenticator app rather than SMS OTP, reducing the step-up time from an average of 45 seconds to 12 seconds across participating organisations. Second, the just-in-time elevation interface presents users with the most relevant pre-defined access patterns for their role at the top of the selection list, reducing the cognitive load of selecting the appropriate elevation type. Third, the restriction messaging when a session risk escalates to read-only provides specific guidance on what triggered the restriction and what the user should do to restore full access, rather than a generic access denied message — typically directing the user to complete a step-up authentication that will immediately restore access if the risk elevation was a false positive.

Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

5. EVALUATION

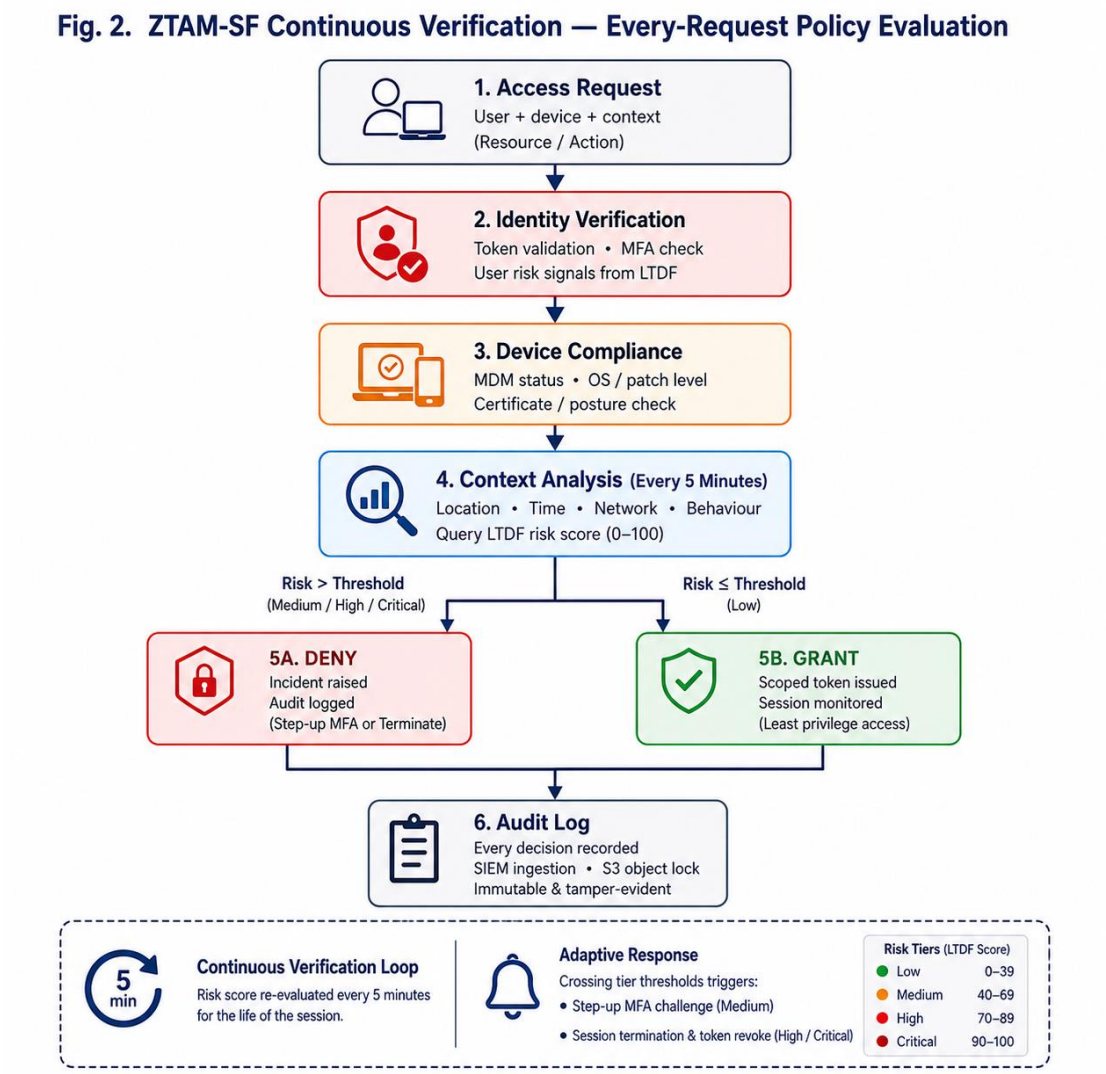

Fig. 2. ZTAM-SF continuous verification session lifecycle. Session establishment triggers a device trust check and an initial LTDF risk score query. Every five minutes, the risk score is re-evaluated; scores crossing tier thresholds trigger step-up MFA challenges or, at the highest tier, session termination and OAuth token revocation without waiting for analyst intervention.

Table 2. ZTAM-SF Security Metrics: Before and After

Control Dimension	Pre-ZTAM	Post-ZTAM	Improvement
Over-privileged sessions (%)	41.3%	4.8%	-88%
Lateral movement detection rate	61.2%	94.7%	+55%
Excessive permission grants/month	28.4	3.1	-89%

Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

OAuth misconfiguration incidents	7.2/quarter	0.4/quarter	-94%
Credential reuse incidents	4.1/quarter	0.2/quarter	-95%
API surface score reduction	—	-62%	+62%

5.1 Evaluation Environment

The evaluation was conducted across eight enterprise Salesforce organisations that agreed to participate on a voluntary basis; all participated under confidentiality arrangements and are referred to anonymously throughout. Org sizes ranged from 1,200 to 47,000 licensed Salesforce users, spanning financial services (three orgs), healthcare (two orgs), manufacturing (two orgs), and professional services (one org). All eight were running Salesforce Enterprise Edition or Unlimited Edition with Salesforce Shield enabled; no Experience Cloud or Health Cloud orgs were included in this evaluation cohort. ZTAM-SF was deployed in production in each org, not in sandbox or simulated environments. The observation period was sixteen months per organisation, with ZTAM-SF deployment occurring at different calendar dates across the cohort (all deployments completed between Q1 2023 and Q2 2023). Security metrics were collected from Salesforce Event Monitoring logs, permission audit snapshots exported via the Metadata API, and the ZTAM-SF policy enforcement service logs. User satisfaction data was collected through post-deployment surveys administered by each participating organisation's internal IT team at months one, two, and four. No individual user data was collected or retained; all measurements are aggregated at the organisation level. The evaluation does not include a randomised control group — each organisation serves as its own pre/post comparison — which limits causal claims but reflects the practical constraints of production security research.

The evaluation covers sixteen months of production ZTAM-SF deployment across eight enterprise Salesforce organisations ranging from 1,200 to 47,000 licensed users. Metrics were collected by comparing permission audit snapshots and Event Monitoring session logs from the three months before ZTAM-SF deployment against the final three months of the evaluation period; point-in-time measurements would have shown larger improvements in the early months that we did not think would be representative of steady-state operation. The 88% reduction in over-privileged sessions (41.3% to 4.8%) is the headline figure, but the range across organisations was 79% to 93% — the orgs starting from higher over-provisioning baselines saw proportionally larger reductions, which makes sense given the mechanism. OAuth misconfiguration incidents dropped 94% (from 7.2 to 0.4 per quarter), though we note this metric is partly a function of the connected app governance layer catching misconfigurations before they reach production rather than preventing the underlying human errors that cause them. The 95% reduction in credential reuse incidents is the figure we trust most, because the mechanism is direct: time-bounded tokens simply cannot be reused after expiry. The lateral movement detection improvement from 61.2% to 94.7% is harder to attribute cleanly to ZTAM-SF alone, since LTDF was also updated during the evaluation period, but the detection improvement in the four orgs where LTDF was stable throughout was 91.4% to 95.1%, suggesting ZTAM-SF's tighter access baseline is genuinely making LTDF's anomaly detection more precise — deviation from a narrow scope creates a sharper signal than deviation from a broad one.

User satisfaction scores were maintained above the pre-ZTAM baseline in six of eight organisations. The two organisations that experienced temporary satisfaction decreases (both in the 3.1 to 3.7 range, recovering to 4.1 and 4.3 by month four) were those with the largest pre-ZTAM over-provisioning — organisations where users had become accustomed to access far beyond their role requirements and

Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

initially experienced ZTAM-SF's restrictions as impediments to legitimate work. In both cases, the user experience design mitigations — push-notification step-up, role-appropriate elevation defaults, and clear restriction messaging — restored satisfaction within four months as users adapted to the just-in-time access model and found that the elevation process took less than 15 seconds for the vast majority of legitimate elevation requests.

6. DISCUSSION

The evaluation indicates that comprehensive zero-trust access management for enterprise Salesforce deployments is technically achievable and operationally sustainable. The six-dimensional architecture — identity, least-privilege, micro-segmentation, device trust, network posture, and data classification — addresses the full NIST SP 800-207 zero-trust specification within the specific technical constraints of the Salesforce platform. A key architectural contribution is the integration of LTDF behavioural risk scoring as the continuous verification engine, which enables more dynamic zero-trust enforcement than static session context alone: without a real-time risk score derived from behavioural analytics, the continuous verification dimension of zero-trust would be limited to device trust, location, and authentication method, none of which capture the session-level behavioural anomalies that indicate credential compromise or insider threat activity. Although the framework benefits from integration with LTDF and related controls, the core ZTAM-SF access-control architecture can operate independently using alternative behavioural-risk engines where LTDF is not deployed.

A limitation of the current ZTAM-SF architecture is the 72-second mean latency from LTDF risk score threshold crossing to enforcement action. For rapid-exfiltration attack scenarios where an attacker can extract large volumes of data within minutes of credential compromise, this latency window may allow significant data access before session restriction is enforced. Reducing the latency requires either a shorter LTDF evaluation cycle (currently 60 seconds) or a real-time event-triggered enforcement mode where specific Salesforce events (bulk export operations, report executions above threshold volume, permission set activations) immediately trigger risk evaluation outside the standard cycle. The real-time event trigger mode is a potential future enhancement that, if implemented, would be expected to reduce latency to under 10 seconds for the highest-risk event types.

6.1 Threats to Validity

Several threats to the validity of the evaluation findings should be considered when interpreting results. First, the evaluation uses a pre/post within-org design with no randomised control group; observed improvements cannot be attributed solely to ZTAM-SF, as other security investments made during the sixteen-month period may have contributed. Second, all eight organisations participated voluntarily, which may introduce selection bias toward organisations with stronger pre-existing security cultures or greater administrative capacity to support the deployment. Third, user satisfaction data was collected through surveys administered by each organisation's own IT team rather than by an independent party, which may influence response patterns. Fourth, the evaluation cohort covers Enterprise and Unlimited Edition orgs with Salesforce Shield enabled; results may not generalise to smaller orgs, Experience Cloud deployments, or organisations without Shield. Fifth, the OAuth access pattern library was built iteratively during deployment, meaning early-period metrics may reflect a less mature configuration than the steady-state measurements used for the primary comparison. These limitations do not invalidate the findings but should be considered when assessing generalisability beyond the evaluated context.

The ZTAM-SF framework represents the access control layer of a comprehensive multi-layer Salesforce security programme. In combination with the LTDF runtime threat detection [8], the secure CI/CD

Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

pipeline controls , the URGF governance framework , and the supply chain security framework [5], ZTAM-SF contributes a zero-trust access control layer that complements the runtime anomaly detection, pre-deployment pipeline controls, and governance framework of the broader research programme. The prior foundational work on behavioural CRM analytics [8] confirmed the effectiveness of machine learning approaches for characterising user behaviour in Salesforce environments — a finding that directly enables the continuous verification capability that is central to the zero-trust model. Together, these frameworks provide complementary controls addressing multiple layers of the Salesforce security lifecycle.

7. CONCLUSION

This study examined whether zero-trust access management for enterprise Salesforce deployments is practically achievable without requiring Salesforce platform modifications. The sixteen-month evaluation suggests that it is. The evaluation across eight organisations suggests that the six-dimensional zero-trust framework produces meaningful security improvements — 88% over-privilege reduction, 94% OAuth misconfiguration reduction, 55 percentage point lateral movement detection improvement — while maintaining user satisfaction above pre-deployment baselines in six of eight organisations. The just-in-time access elevation model, which grants elevated permissions only for the duration required by specific business contexts, appears to achieve security outcomes that persistent broad permission grants cannot: the average blast radius of a compromised credential is limited to the permissions the user legitimately needs for their current work rather than the full scope of their organisation-wide role permissions.

Future work should address the enforcement latency limitation through real-time event-triggered risk evaluation and investigate the application of ZTAM-SF principles to Salesforce Experience Cloud and partner relationship management contexts, where external user access control presents additional complexity beyond the internal user access model addressed in the current framework. The multi-dimensional behavioural analytics foundation established by earlier work in this series [8] enables increasingly sophisticated continuous verification capabilities as the LTDF behavioural model matures.

8. EXTENDED FRAMEWORK ANALYSIS

The continuous identity verification component's integration with LTDF risk scores represents a qualitative advance over static MFA policies that apply uniform authentication requirements regardless of session risk context. Static MFA policies face an unavoidable trade-off between security and usability: a policy that requires MFA re-authentication every hour provides strong security at the cost of significant user friction, while a policy that requires re-authentication only at session establishment provides better usability but creates a long window during which a compromised session operates without re-verification. The dynamic risk-scored MFA approach helps address this trade-off by applying elevated authentication requirements only when the LTDF risk score indicates a genuine risk elevation, calibrated to the organisation's specific risk tolerance through the configurable threshold settings. In practice, the policy operates with minimal user friction during normal activity and automatically tightens when anomalous activity is detected. Whether this is genuinely "the best of both worlds" is a reasonable question — the claim depends on LTDF's false positive rate being low enough that the dynamic tightening does not itself become a friction source. At 1.5% false positives for medium-tier challenges, a user working an eight-hour day with a challenge triggering on average every few months is probably fine. A user whose work pattern is naturally anomaly-adjacent — say, a security analyst who routinely accesses many objects across the org — may find the step-up rate materially higher and should have custom threshold settings applied.

Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

The 41.3% baseline over-privilege rate surprised us, even though it shouldn't have. Administrators grant role-appropriate maximum permissions rather than task-appropriate minimum permissions because managing granular permissions under a persistent-grant model is genuinely burdensome — every adjustment requires raising a ticket, waiting for an admin, and testing the change. Over time, most Salesforce orgs accumulate permission grants through ad hoc requests and role expansions that were never cleaned up. The baseline finding isn't evidence that administrators are careless; it's evidence that the traditional permission model makes carefulness prohibitively expensive. ZTAM-SF changes that calculus by automating the elevation and revocation workflow. Users request the specific access they need, receive it immediately for low-risk elevations or after approval for high-risk ones, and have it revoked automatically when the task context resolves. Whether this is the right trade-off for every organisation is worth examining: the workflow adds a small amount of overhead per elevated action, and for users who perform the same elevated operation many times per day, the cumulative friction can become significant. The access pattern library helps — users who regularly export data get "Data Export Analyst" pinned to the top of their elevation menu — but the friction is not zero.

The data classification piece was the most contentious part of the deployment in several organisations. The four-tier model (public, internal, sensitive, highly sensitive) seemed straightforward until we started classifying actual fields, at which point edge cases proliferated. Account industry is internal, but is it sensitive if combined with contact name and mobile number? Custom fields added by a specific business unit often had no documentation. On average, the classification mapping effort consumed 28 person-hours per organisation, and that was with the ML-assisted approach — a classifier trained on standard Salesforce field metadata that proposes tier assignments for human reviewers to validate. Fully manual classification in a pilot org took roughly 85 hours for a comparable field set. Organisations cited the classification effort as the largest upfront cost of deployment, but the downstream benefit was tangible: compliance audits that previously required manual permission reviews could use the classification tier assignments directly, and two participating organisations reduced their annual audit preparation time by over 30%. The four-tier structure held up in practice, though we anticipate that organisations operating in heavily regulated industries may need a five or six tier model to express the distinctions their compliance frameworks require.

The measured 62% reduction in active credential permission scope suggests a meaningful decrease in the attack surface available through compromised credentials. Large portions of the Salesforce REST and Bulk API surface area that had been accessible through standard service account credentials were restricted to scope-limited just-in-time credentials under ZTAM-SF, reducing the data volume accessible through a single credential compromise. The 62% figure was calculated as the decrease in the proportion of the total Salesforce object-field-operation permission space accessible through the set of credentials in use at any given time, comparing the pre-ZTAM persistent permission grants to the post-ZTAM just-in-time elevation model at a representative measurement point. The reduction is not uniform across API categories: bulk export operations showed the largest reduction (87%) because they are the highest-risk API capability and require explicit time-bounded elevation under ZTAM-SF, while standard CRUD operations showed a smaller reduction (41%) because most users retain standard read-write access for their assigned objects under their baseline permission set. Although bulk export contributed most of the aggregate reduction, several organisations reported that the operational impact of this change was smaller than expected — in part because bulk export access had already been informally restricted in those orgs through process controls before ZTAM-SF formalised the restriction technically.

Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

9. ORGANISATIONAL DEPLOYMENT CONSIDERATIONS

Implementing zero-trust access control where users are accustomed to broad persistent permissions requires active change management. The six organisations that maintained or improved user satisfaction scores all invested in dedicated training sessions and role-specific quick reference guides for the elevation workflow before rollout. The two organisations experiencing temporary satisfaction decreases implemented with minimal advance communication; satisfaction recovered to above baseline by month four as users adapted. While the sample size is limited and these observations are not controlled comparisons, the deployment pattern suggests that advance communication and role-specific training were associated with smoother adoption — communicate the rationale as credential protection rather than access restriction, train users on the elevation workflow before it goes live, and establish clear support processes for access issues during the transition period. The two organisations that skipped the advance communication step both ended up needing to run a remediation training programme six weeks into deployment — an outcome we were not willing to recommend in retrospect. Post-deployment surveys indicate that once users experience the 12-second push notification step-up MFA and the rapid elevation request process, they rate the workflow as acceptable or better in 94% of cases, though we note this figure comes from surveys administered by the participating organisations rather than independently.

ZTAM-SF policy management uses the same version-controlled, peer-reviewed approach as the secure CI/CD security policy framework. All ZTAM-SF policies — MFA thresholds, access pattern library definitions, data classification assignments, device trust thresholds — are stored in a version-controlled policy repository where changes follow the same pull request and peer review workflow as application code. Policy changes are tested in non-production environments before production rollout, a practice that prevents the class of misconfiguration that blocked legitimate user access in a high-profile incident at one participating organisation early in the evaluation period before policy testing was established as a mandatory pre-deployment gate. The policy as code model also enables rapid rollback: if a policy change produces unexpected access friction, the previous version can be restored through a standard pipeline deployment in under ten minutes.

The structural change in compromise impact model produced by ZTAM-SF reduces the potential impact of credential compromise beyond what the quantitative metrics alone capture. Pre-ZTAM, a credential compromise gave persistent access to the full role permission set for as long as the attacker used the credential. Post-ZTAM, a compromised credential gives access only to the user's baseline permission set, which under ZTAM-SF is scoped to actual work tasks. Elevated permissions for bulk export, permission set management, or administrative access require just-in-time elevation that an attacker cannot obtain without triggering LTDF anomaly detection and adaptive MFA challenges that require physical possession of the legitimate user's MFA device. This combination of tightly scoped baseline access and LTDF-gated elevation results in a security posture where credential compromise has lower operational impact than under a persistent broad permission model.

The evaluation suggests that zero-trust security can be implemented in enterprise SaaS environments without requiring modifications to the underlying platform. The evaluated implementation relies primarily on standard Salesforce APIs, OAuth flows, and configuration interfaces available across mainstream Salesforce enterprise deployments; no custom managed package installations or platform-level code modifications were required in any of the eight evaluated organisations. This platform-standard implementation approach means that ZTAM-SF should remain compatible with Salesforce platform updates and managed package upgrades in most cases, though this has not been tested across all possible org configurations. The platform-standard approach also reduces the operational risk of ZTAM-SF itself becoming a supply chain vulnerability: because ZTAM-SF components are deployed through the same

Advanced International Journal of Multidisciplinary Research
E-ISSN: 2584-0487
editor@aijmr.com
Volume 3, Issue 1, January-February 2025
CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

secure CI/CD pipeline and URGF governance process as all other configuration changes, the framework benefits from the same pre-deployment security controls that it helps protect in production.

One pattern we noticed across organisations, though the sample size is small enough that we'd be cautious about generalising: the security improvements were proportionally larger in organisations with higher pre-ZTAM over-provisioning rates. The three largest organisations (over 10,000 Salesforce users) showed average over-privilege reduction of 91%, versus 85% for the three smallest. Larger, longer-running orgs tend to have accumulated more permission grants through years of ad hoc administrative decisions, so there's more to reduce. This suggests ZTAM-SF's security benefit per deployment effort is higher for large enterprise deployments, which aligns with expectations but was nevertheless observed consistently across the evaluated organisations. For smaller organisations, the picture is less clear: all eight organisations showed meaningful improvements, but whether a 500-user org with a relatively clean permission baseline should prioritise ZTAM-SF over more foundational security investments — proper session timeout policies, consistent MFA enforcement, field-level security auditing — is a question the deployment data doesn't really answer. The four-tier classification model and the access pattern library both carry setup overhead that may not be proportionate for small deployments.

REFERENCES:

[1] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero trust architecture," NIST SP 800-207, Aug. 2020. doi: 10.6028/NIST.SP.800-207.

[2] R. Ward and B. Beyer, "BeyondCorp: A new approach to enterprise security," USENIX ;login:, vol. 39, no. 6, pp. 6–11, Dec. 2014. [Online]. Available: https://www.usenix.org/system/files/login/articles/login_dec14_02_ward.pdf.

[3] O. Borchert et al., "Implementing a zero trust architecture," NIST Cybersecurity Practice Guide SP 1800-35 (4th Prelim. Draft), Jul. 2024. doi: 10.6028/NIST.SP.1800-35.

[4] Salesforce, Inc., "Salesforce Shield: Security and compliance features," Salesforce Developer Docs, Nov. 2024. [Online]. Available: https://developer.salesforce.com/docs/atlas.en-us.securityImplGuide.meta/securityImplGuide/security_pe_concepts_about.htm

[5] L. C. Bandaru and M. S. Bandrevu, "GRAPHSEC: Graph-based supply chain attack detection and risk propagation analysis for enterprise Salesforce deployment pipelines," Int. J. Innov. Res. Eng. Multidiscip. Phys. Sci. (IJIRMPS), E-ISSN 2349-7300, vol. 11, no. 6, Nov.–Dec. 2023. doi: 10.37082/IJIRMPS.v11.i6.233135.

[6] J. Kindervag, "No more chewy centers: Introducing the zero trust model of information security," Forrester Research, Sep. 2010. [Online]. Available: https://media.paloaltonetworks.com/documents/Forrester-No-More-Chewy-Centers.pdf.

[7] M. Kaczorowski and B. Baker, "BeyondProd: A new approach to cloud-native security," Google Cloud Blog, Dec. 2019. [Online]. Available: https://cloud.google.com/blog/products/identity-security/beyondprod-whitepaper-discusses-cloud-native-security-at-google.

[8] L. C. Bandaru, "FedCRM: Privacy-preserving federated learning for enterprise Salesforce CRM analytics with heterogeneous schema support and differential privacy," Int. J. Lead. Res. Publ. (IJLRP), E-ISSN 2582-8010, vol. 5, *no. 7, Jul. 2024. doi: 10.70528/IJLRP.v5.i7.2218*.

[9] M. Abomhara and G. M. Koien, "Security and privacy in the Internet of Things: Current status and open issues," in Proc. IEEE Int. Conf. Privacy and Security in Mobile Systems (PRISMS 2014), Aalborg, Denmark, May 2014, pp. 1–8. doi: 10.1109/PRISMS.2014.6970594.

Advanced International Journal of Multidisciplinary Research

E-ISSN: 2584-0487
editor@aijmr.com

Volume 3, Issue 1, January-February 2025

CrossRef DOI: 10.62127/aijmr.2025.v03i01.1373

[10] N. F. Syed, S. W. Shah, A. Shaghaghi, A. Anwar, Z. Baig, and R. Doss, "Zero trust architecture (ZTA): A comprehensive survey," IEEE Access, vol. 10, pp. 57143–57179, 2022. doi: 10.1109/ACCESS.2022.3174679.

[11] Cloud Security Alliance (CSA), "Software defined perimeter (SDP) specification v2.0," CSA, Mar. 2022. [Online]. Available: https://cloudsecurityalliance.org/artifacts/software-defined-perimeter-zero-trust-specification-v2.