Graduate Journal of Interdisciplinary Research, Reports & Reviews
(Grad.J.InteR^3)
ark:/85846/JPR3.2026AR.04010047

ISSN(E): 2584-2919
Vol.04, No.01 (2026)
pp.47-53

Protecting Identities and Secrets at Enterprise Scale: A Zero Trust Framework for the Age of AI-Augmented Threats

Sudheer Reddy Koppula
Microsoft Corporation Redmond, WA, USA

*Corresponding author: sudheerreddy.koppula@gmail.com

(Received:23-05-2026; Accepted: 15-07-2026; Published Online:25-07-2026)

ABSTRACT: Background: As contemporary cloud-native systems evolve, identity has quickly become the foremost surface where attackers can compromise accounts. In addition to the 600 million identity-based attacks seen on a daily basis, 97% of these attacks leverage credentials i.e., passwords to gain access to systems.
Purpose: Rapid cloud adoption, new workloads created through cloud services, and emerging methods of attack (e.g., AI-assisted phishing or AiTM, which is often paired with credential theft) have made off-the-shelf authentication models obsolete. The objective of this paper is to present a comprehensive framework to protect an enterprise's identity and secret at an enterprise-level and is based on Zero Trust principles as well as the experience of the author who serves as an SRE Platform manager for a large public cloud company.
Method: We will also address the evolving landscape of potential identity threats; for example: credential stuffing, token replay, OAuth consent phishing and attacks against non-person identities; and provide a multi-layered defense architecture to include: phishing-resistant multi-factor authentication (MFA), passwordless migration, workload identity governance, secrets lifecycle management, and AI-based anomaly detection.
Results: As part of this paper, we will also share results from an enterprise-scale implementation of MFA where we achieved more than 99.6% MFA adoption across 200,000+ employees and discuss the operational, cultural, and architectural lessons learned from this implementation.
Conclusion: We will conclude this paper with action items for those organizations that hope to move from password-based authentication to a Zero Trust identity model.

Keywords: identity protection, Zero Trust, passwordless authentication, FIDO2, MFA, secrets management, workload identity, credential attacks, phishing-resistant authentication, AI-driven security

1. Introduction

The concept of identity has taken over as the new 'perimeter' and is changing the way that businesses approach security. With organizations moving to the cloud, eliminating traditional network boundaries, and developing hybrid work models, verifying who is attempting to gain access is the most important part of the security stack. The data supports this change: today, the largest cloud platforms block over 7000 attempted password attacks every second and see more than 600 million identity-based attacks each day [1,2]. In fact, as of June 2022 there was an increase of over 32% in identity-based attacks of which 97% were targeting password-based authentication [2]. Identity compromise can have disastrous effects that are severe and long-lasting. Stolen credentials are the most common way attackers gain access to enterprise systems and 90% of organizations have had at least one incident related to identity in the previous year [3]. The extreme costs associated with breaches such as the UnitedHealth and Change Healthcare breach of $786M show identity failures are not only technical problems, but they are also existential risks to the enterprise [2]. Despite these risks, 65% of enterprises still use only passwords for some of their systems, and only 21% have deployed phishing-resistant FIDO2 (Fast IDentity Online) authenticators [4]. This disparity between

© 2026 Author(s). This article is published under the CC-BY license at http://jpr.vyomhansjournals.com.

threat severity and maturity of defensive controls is a pressing challenge for enterprises today. In this paper, we provide a complete framework for protecting identities and secrets at enterprise-level scale by extending the Protect Identities & Secrets pillar of the Enterprise Security Transformation Framework [5]. The material is built upon the author's experience as a Platform SRE Manager with a leading cloud provider, in addition to conversations with over Five Hundred (500) security practitioners at the Global Security Tour Seattle in 2026, and includes both a reference architecture and practical guidance for organizations at any point on the identity protection continuum.

2. The Identity Threat Landscape

The identity threat landscape has evolved dramatically, driven by the convergence of cloud adoption, AI-augmented attacks, and the proliferation of non-human identities.

2.1. A. Password and Credential Attacks at Scale:

Identity-based attacks utilizing passwords are still the number one target today, with password-based attacks representing over 97% of all identity attacks seen on the major Cloud providers [2]. These attacks fall into a variety of categories:

- ***Password Spraying:*** An automated set of attempts against many accounts, using a small set of commonly used passwords, while still remaining under lockout thresholds.
- ***Credential Stuffing:*** Utilizing username and password combinations from previous breaches to gain access to other services where users have reused their passwords.
- ***Brute Force Attacks:*** An automated high volume of attempts to guess at credentials; Typically done against service accounts and older systems that have poor password policies.

Number wise, this is unreal; For all the Cloud infrastructure globally, there is an average of 7,000 password attacks blocked every second; This number has continued to grow each year as attackers are automating and trying to exploit credentials at an unprecedented level [1].

2.2. B. Adversary-in-the-Middle and Token Replay:

When organizations utilize conventional MFA (SMS codes, push notifications), hackers have learned to implement adversary-in-the-middle (AiTM) methods, allowing them to capture the authentication tokens during the authentication process in real-time [2]. In the AiTM attack, the attacker controls the proxy server that the victim uses to authenticate; once authenticated via the proxy, the attacker captures the session token and is able to leverage it against the victim to gain access to the victim's account. Therefore, with the aid of AiTM attacks, MFA methods that are not resistant to phishing become nearly useless when it comes to stopping highly-skilled/advanced adversaries. The rate of success for token replay attacks targeting cloud session credentials to bypass MFA has risen substantially. To identify a token replay attack, continuous anomalous identity activity monitoring must take place particularly with regard to unusual activity such as impossible travel, sudden privilege escalation, and concurrent session activity occurring from geographically-separated locations.

2.3. C. Auth Consent Phishing and Application Abuse:

The OAuth Consent Model is increasingly being used as a means of launching Identity Attacks. In this type of identity attack, malicious applications are created by attackers in order to entice potential victims to provide high levels of access via what appear to be legitimate consent prompts [2]. Once the victim has provided consent, the attacker has access even after performing password resets or enrolling in MFA services. Access is maintained in a manner that is very difficult to detect and remediate.

2.4. D. Attacks on Workload and Non-Human Identities:

As user-focused protections improve through MFA and conditional access policies, attackers are pivoting to target workload identities applications, service accounts, scripts, and API keys that often have elevated

© 2026 Author(s). This article is published under the CC-BY license at http://jpr.vyomhansjournals.com.

permissions but lack the governance controls applied to human identities [2]. In many enterprise environments, workload identities outnumber human identities by a factor of 10:1 or more.

2.5. E. AI-Augmented Identity Threats:

AI has drastically changed the economics of identity breaches. Phishing emails now made with AI are much more believable, almost impossible to discover, and can be created at scale [5]. In addition, things like the use of deepfake technology to create a voice and video impersonation make impersonating someone's identity for social engineering purposes much easier for attackers. Conversely, AI will also revolutionise the way organisations defend against identity breaches by providing AI-driven identity protection systems with the ability to recognise unusual authentication behaviour and automate response actions in seconds.

3. A Framework For Identity And Secrets Protection

From our experience implementing identity protection solutions at large corporations and guided by the Zero Trust concepts found in NIST SP 800-207 [6], we have developed a multi-pronged approach to protect identities that includes five levels of defense.

3.1. A. Layer 1: Phishing-Resistant Authentication:

To protect identities, we must eliminate all methods of phishable authentication factors. Phishing-resistant MFA such as: FIDO2 security keys; platform authenticators (Windows Hello, Apple's Touch ID); and certificate-based authentication will stop 99% of identity attacks [2, 7]. Default enforcement for all users. Every user should have phishing resistant MFA as their default. This cannot be an option for them. Progressive rollout will occur by starting with privileged accounts, then those who are higher risk, and finally all users. Example data shows that with over 200,000 employees using MFA, that it is very feasible to have 99.6% user adoption [5]. We must systematically eliminate legacy protocols such as NTLM; Basic Authentication; and legacy SMTP so as to not bypass new, modern controls.

3.2. B. Layer 2:

Passwordless Migration: In addition to Multi-Factor Authentication (MFA), eliminating passwords entirely from authentication will be a goal of the strategy. Passwords are inherently risky and can be either guessed, stolen, phished, sprayed or stuffed. By eliminating passwords as an authentication method through a passwordless authentication system you will eliminate this entire type of attack against your organization.

3.3. C. Layer 3: Workload Identity Governance

- Replace stored credentials with a platform-managing identity: When a platform has the capability for using managed identities, it is recommended to use that functionality rather than using stored credentials.
- Minimize permissions or access to least privilege: Workload identities should be assigned as small a set of permissions as necessary based on an access review process.
- Use conditional access for workloads: Workloads that require authentication should leverage IP range restrictions or verified/compliant device requirements and risk-based access evaluation.
- Use administrative workflow for app consent management: Users should only provide consent for low-risk and verified applications; all other applications should use the admin workflow to provide consent.

3.4. D. Layer 4: Secrets Lifecycle Management

- Centralized vault: All secrets reside in a dedicated secrets management infrastructure with identity-based access controls.
- Automated rotation: Secrets are rotated on predefined schedules (e.g., lower frequency for higher sensitivity credentials).
- CI/CD secrets scans: Credential scanning is part of every build & deploy pipeline in an automated fashion.

© 2026 Author(s). This article is published under the CC-BY license at http://jpr.vyomhansjournals.com.

Table 1: Comparison of Passwordless Authentication Methods

Method	Mechanism	Strengths	Phishing Resistant
FIDO2 Security Keys	Hardware-bound cryptographic key pair	Strongest phishing resistance; works offline	Yes
Platform Authenticators	Biometric + device-bound key	Frictionless UX; no separate device	Yes
Certificate-Based Auth	X.509 certificates on managed devices	Strong device and identity binding	Yes
Passkeys	Synced FIDO2 credentials across devices	Cross-device convenience	Yes
Authenticator App	Push with number matching	Familiar UX; broad device support	Partial

- Just-in-time (JIT) access: JIT provides time-limited access to secrets with complete audit logging, rather than continuous access.
- No standing privileges: You cannot perform administrative functions unless you have elevated privileges, which can only be done through approval and time-limited access.

3.5. E. Layer 5: AI-Driven Identity Threat Detection

- Behavioral analytics: Using machine learning models to create baseline behaviors and flagging anomalies, such as impossible travel, unusual access patterns, and privilege escalation attempts.
- Graph-based threat modeling: Identity relationship graphs are used to model lateral movement and attack chains [5].
- Automated response: When a high-confidence compromise occurs, the compromised session should be revoked, the credentials reset, and enforcement of conditional access to be initiated.
- Continuous access evaluation: Provides real-time trust reassessment according to the changing risk signals.

4. Case Study: Implementation At Enterprise Scale

4.1. A. Quantitative Outcomes

- MFA Implementation: 99.6% of 200,000+ users have adopted MFA; methods that protect user accounts from phishing attacks were prioritized for all management accounts. [5]
- Tougher to compromise passwords: Following an enforcement of MFA and discontinuing legacy protocols there was a stark decrease in successful credential-based attacks.
- Secrets Hygiene: All CI/CD pipeline secrets are automatically scanned.
- Workload Identity Governance: Default use of managed identities reducing stored credentials by more than 50% in target areas.

4.2. B. Operational Lessons

- Implementing phased rollouts is crucial as trying to implement phishing resistant MFA at once creates too much of a support load overall throughout the organization. During the pilot phase of implementation, a phased rollout approach (i.e. 1. Privileged accounts, 2. High risk population and 3. General users) worked very well.
- Long-standing protocol deprecation is the unspoken obstacle as many lines of business applications were built on basic authentication - this required individual application remediation for every system affected.
- Workload identities required separate governance (meaning, separate tooling and ownership)

© 2026 Author(s). This article is published under the CC-BY license at http://jpr.vyomhansjournals.com.

in order to build a workload identity lifecycle (i.e., creation of the identity, assigning permissions to the identity, monitoring the identity, rotating the identity, and ultimately decommissioning the identity).

- Leadership informing engineers that security took priority over all other processes required a cultural shift from the engineers' standard way of designing identity on a daily basis [5].

4.3. C. Practitioner Insights from the Security Community:

While conducting the Global Security Tour Seattle (2026), over 500 of the security professional's provided feedback to common themes shared among their experiences across the industry. The trends are as follows:

- Organizations that implement phishing resistant multifactor authentication (MFA), have experienced a disproportionately high number of identity incidents versus organizations that implemented either SMS or push-based MFA.

- Every security professional identified that managing identities will be the upcoming major initiative within the security profession.

- Although most security professionals feel AI will significantly enhance threat detection related to identity security; at present, many organizations are still early in their implementation/usage of AI.

- Concern regarding post-quantum cryptography's impact on an organization's identity infrastructure, continues to increase

5. Recommendations

5.1. A. Eliminate Passwords as a Primary Authentication Factor:

Organizations should establish a concrete roadmap for passwordless migration, beginning with administrative accounts. FIDO2 security keys and platform authenticators should be the target state for all human authentication [5].

5.2. B. Govern Workload Identities with the Same Rigor as Human Identities:

Every workload identity should have an owner, a defined purpose, least-privilege permissions, and an expiration or review date. Managed identities should replace stored credentials wherever platform support exists.

5.3. C. Implement Continuous Access Evaluation:

Static, session-based trust models are insufficient. Organizations should implement continuous access evaluation that reassesses trust based on real-time risk signals and can revoke or step up authentication dynamically.

5.4. D. Automate Secrets Lifecycle Management:

Secrets scanning, rotation, and revocation should be fully automated. Integration of secrets scanning into CI/CD pipelines should be a non-negotiable requirement for all engineering teams.

5.5. E. Invest in AI-Driven Identity Threat Detection:

Organizations should invest in identity threat detection capabilities that leverage behavioral analytics, graph-based modeling, and automated response. Defenders who do not adopt AI-driven identity protection will be systematically disadvantaged.

5.6. F. Prepare Identity Infrastructure for Post-Quantum Cryptography:

Organizations should begin inventorying cryptographic dependencies in their identity infrastructure. Migration planning toward quantum-resistant algorithms should begin now [5].

6. Conclusion

The foundation of Zero Trust security is built upon identity protection. This need stems from the incredible threat landscape over 600 million identity attacks daily, with the vast majority of those (97%) targeting passwords; an alarming rise in the incidence of Account Take Over (AiT) attacks via token replay; and ungoverned workload identity at increasing risk of exploitation. Therefore, there must be a fundamental shift away from traditional methods of password-based authentication to newer, more secure methods that

© 2026 Author(s). This article is published under the CC-BY license at http://jpr.vyomhansjournals.com.

are resistant to phishing attacks and do not rely on passwords, with identity verification needing to continuously evolve relative to emerging threats. This paper provides a structured framework (that encompasses phishing-resistant authentication, passwordless migration, governance of workload identities, secrets lifecycle management, and AI-enabled detection of threats) to address these challenges across the enterprise. There are three key imperatives that security leaders need to understand: 1) Passwords need to be completely removed as a factor of authentication, not just augmented; 2) Workload identities represent the next frontier in terms of risk, and require specific governance; and 3) AI-enabled identity protection is essential - it is the only method that can keep pace with the scale and speed of a modern attacker.

References

[1] Microsoft, "microsoft Digital Defense Report 2024, " Microsoft Corporation, Tech. Rep., Oct. 2024.

[2] Microsoft, "microsoft Digital Defense Report 2025, " Microsoft Corporation, Tech. Rep., Jan. 2026.

[3] Beyondtrust, "the State Of Identity Security 2024, " Beyondtrust, Tech. Rep., 2024.

[4] BusinessWire, "New Report Reveals Persistent Enterprise Reliance on Passwords Despite Rising Identity Threats," Mar. 2025.

[5] S. R. Koppula, "Enterprise Security Transformation: Insights from Global Security Tour Seattle 2026," LinkedIn Pulse, Apr. 2026.

[6] S. Rose et al., "Zero Trust Architecture," NIST SP 800-207, Aug. 2020.

[7] FIDO Alliance, "FIDO2: Web Authentication (WebAuthn)," FIDO Alliance Specifications, 2019.

[8] Executive Office of the President, "Executive Order 14028: Improving the Nation's Cybersecurity," Federal Register, May 2021.

[9] European Parliament and Council, "Directive (EU) 2022/2555 (NIS2 Directive)," Official Journal of the EU, Dec. 2022.

[10] Identity Defined Security Alliance, "2024 Trends In Securing Digital Identities, " Idsa, Tech. Rep., 2024.

© 2026 Author(s). This article is published under the CC-BY license at http://jpr.vyomhansjournals.com.

Copyright and License

[© 2026 **Sudheer Reddy Koppula**] This is an Open Access article published in "Graduate Journal of Interdisciplinary Research, Reports & Reviews" (Grad.J.InteR[3]), an **Open Access Journal** by Vyom Hans Publications. It is published with a Creative Commons Attribution - CC-BY4.0 International License. This license permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

How to Cite?

Sudheer Reddy Koppula. (2026). Protecting Identities and Secrets at Enterprise Scale: A Zero Trust Framework for the Age of AI-Augmented Threats. *Graduate Journal of Interdisciplinary Research, Reports and Reviews*, 4(01), 47-53. Retrieved from https://jpr.vyomhansjournals.com/index.php/gjir/article/view/61

© 2026 Author(s). This article is published under the CC-BY license at http://jpr.vyomhansjournals.com.