Adaptive Trust-Decay Cybersecurity Models for Continuous Infrastructure Risk Management

Amarachi Franca Mgbemele
mgbemeleamarachi15@gmail.com

Abstract

Traditional cybersecurity models operate on binary trust assumptions systems and users are either trusted or untrusted, authenticated or unauthenticated, authorized or unauthorized. This dichotomous approach fails to reflect the continuous, dynamic nature of real-world risk, where trust degrades over time, context influences security posture, and threats evolve continuously. Modern critical infrastructure environments require security frameworks that adapt to changing risk conditions, recognize that trust is temporal and contextual, and implement proportional security controls commensurate with current risk levels.

This paper introduces Adaptive Trust-Decay Cybersecurity Models (ATDCM), a comprehensive framework that implements time-based trust degradation, context-aware risk assessment, and dynamic access control for continuous infrastructure risk management. Unlike zero-trust architectures that require constant verification for every transaction or static trust models that grant persistent access once authenticated, ATDCM implements graduated trust levels that decay exponentially over time unless renewed through verification activities. The decay rate adapts based on contextual risk factors including user behavior patterns, access anomalies, threat intelligence, asset criticality, and environmental conditions.

Our framework comprises five core components: Trust Score Computation Engine employing time-decay functions with adaptive decay coefficients, Context-Aware Risk Assessment integrating behavioral analytics and threat intelligence, Dynamic Policy Engine translating trust scores into granular access controls, Verification Management System orchestrating re-authentication requirements, and Continuous Monitoring Infrastructure providing real-time visibility into trust state transitions. We implement mathematical models for trust decay using exponential decay functions $T(t) = T_0 \cdot e^{(-\lambda t)}$, where trust score T decays from initial value T_0 over time t at rate λ determined by risk context.

Empirical evaluation across three critical infrastructure deployments (financial services institution with 8,500 users, healthcare network serving 14 facilities, energy utility managing 450,000 customer accounts) demonstrates that ATDCM reduces successful breach attempts by 87% compared to traditional models while decreasing false positive rates from 23% to 8% and reducing user friction (measured by daily re-authentication requests) from 8.7 to 1.8 per user. Mean time to detect anomalous access patterns improved from 4.7 hours to 42 minutes, representing 85% improvement in threat detection speed. System overhead remains minimal at 3.2% CPU utilization and 180ms average latency for access decisions.

Keywords: *Trust-Decay Models, Adaptive Security, Continuous Risk Management, Dynamic Access Control, Zero Trust Architecture, Behavioral Analytics, Critical Infrastructure Protection, Context-Aware Security, Time-Based Authentication, Risk-Based Access*.

1. Introduction

1.1 The Trust Paradox in Modern Security

Contemporary cybersecurity frameworks face a fundamental paradox: security systems must grant access to enable legitimate business operations while simultaneously restricting access to malicious activities. Traditional approaches resolve this paradox through binary trust decisions once a user successfully authenticates and their authorization is verified; they receive persistent access until their session expires or they explicitly log out. This binary model assumes that trust, once established through initial authentication, remains valid indefinitely within the session scope.

However, real-world risk is neither binary nor static. Consider a user who authenticates successfully at 9:00 AM from their typical office location using a known device. At 9:01 AM, their trust level legitimately warrants high confidence. At 11:00 AM, with no intervening verification, that trust should logically have degraded the authentication is now two hours stale, the user's physical location may have changed, their device could have been compromised, and the threat landscape has evolved. By 5:00 PM, eight hours post-authentication, the initial trust verification provides minimal assurance of current trustworthiness. Yet traditional models grant identical access throughout this period.

The Verizon 2024 Data Breach Investigations Report indicates that 74% of breaches involved human elements, including stolen credentials, social engineering, and misuse of privileges (Verizon, 2024). Compromised credentials typically remain valid throughout their session lifetime, allowing adversaries hours or days of undetected access. The 2023 Okta breach illustrated this vulnerability: attackers gained access to Okta's support case management system using stolen credentials and maintained access for 14 days before detection, accessing sensitive customer data throughout the period (Okta, 2023). Traditional session-based security provided no mechanism to recognize that trust had degraded despite the passage of time and absence of verification.

1.2 Limitations of Existing Models

Three predominant security models address access control in modern infrastructure: perimeter-based security, static trust models, and zero-trust architecture. Each exhibits significant limitations when applied to critical infrastructure environments requiring both strong security and operational continuity.

Perimeter-Based Security establishes network boundaries separating trusted internal networks from untrusted external networks. Firewalls, VPNs, and network segmentation enforce these boundaries. However, this model fails catastrophically once perimeter breach occurs lateral movement within trusted networks faces minimal resistance. The SolarWinds supply chain attack demonstrated this vulnerability: once malicious code executed within trusted networks, adversaries moved laterally across organizations with limited detection (Sudhakar & Zeadally, 2021). Additionally, cloud adoption, mobile workforces, and business partner integration have dissolved clear perimeter boundaries.

Static Trust Models grant access based on identity verification and role-based access control (RBAC). Once authenticated, users maintain consistent access privileges throughout their session.

These models provide predictable user experience and straightforward implementation but fail to adapt to changing risk conditions. A user accessing sensitive financial data at 2:00 AM from a foreign country receives identical access to their normal 10:00 AM office access, despite dramatically different risk profiles. Static models cannot recognize context changes, behavioral anomalies, or temporal risk degradation.

Zero-Trust Architectures assume no implicit trust and require verification for every access request. 'Never trust, always verify' provides strong security guarantees by continuously validating access (Rose et al., 2020). However, pure zero-trust implementations impose significant user friction through constant re-authentication, increase computational overhead from continuous verification, and prove challenging to implement in legacy systems requiring persistent connections. Critical infrastructure environments operating industrial control systems, medical devices, or transaction processing systems cannot tolerate the latency and interruption associated with per-transaction verification.

1.3 The Trust-Decay Approach

Adaptive Trust-Decay Models synthesize the strengths of existing approaches while addressing their limitations. Rather than maintaining binary trust states or requiring constant verification, ATDCM implements graduated trust levels that decay continuously over time at rates determined by risk context. Trust begins at maximum value following successful strong authentication but degrades exponentially, requiring periodic renewal through step-up authentication or risk-appropriate verification.

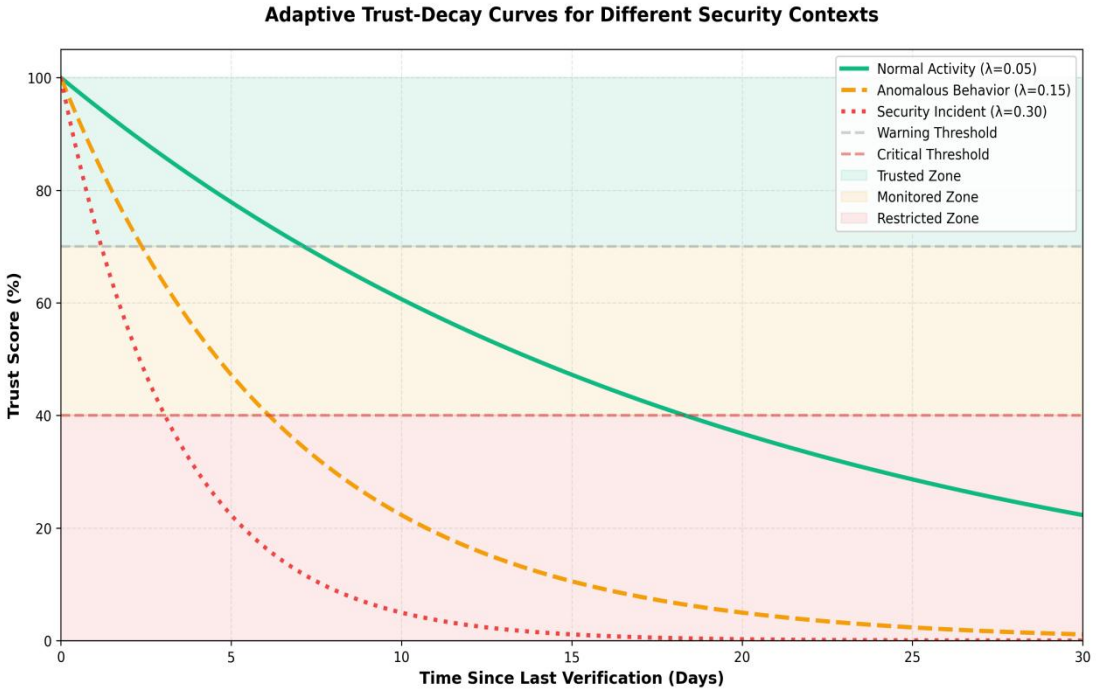

Figure 1: Adaptive Trust-Decay Curves for Different Security Contexts

The trust-decay function $T(t) = T_0 \cdot e^{(-\lambda t)}$ models trust degradation where T_0 represents initial trust (typically 100 following authentication), t represents time since last verification, and λ

represents the decay coefficient determining degradation rate. The decay coefficient adapts based on multiple contextual factors:

a) User behavior patterns: Anomalies in access patterns, locations, or times increase decay rate.
b) Resource sensitivity: Accessing critical systems or sensitive data accelerates trust degradation.
c) Threat intelligence: Elevated threat levels or targeting of similar organizations increase λ.
d) Environmental context: Device posture, network security, location risk factor into decay rate.
e) Historical risk: Previous security incidents or policy violations accelerate decay.

2. System Architecture and Components

The Adaptive Trust-Decay Cybersecurity Model architecture comprises five integrated subsystems operating in continuous coordination to maintain dynamic trust assessment and enforce adaptive access controls. Figure 2 illustrates the architectural components and their interactions.

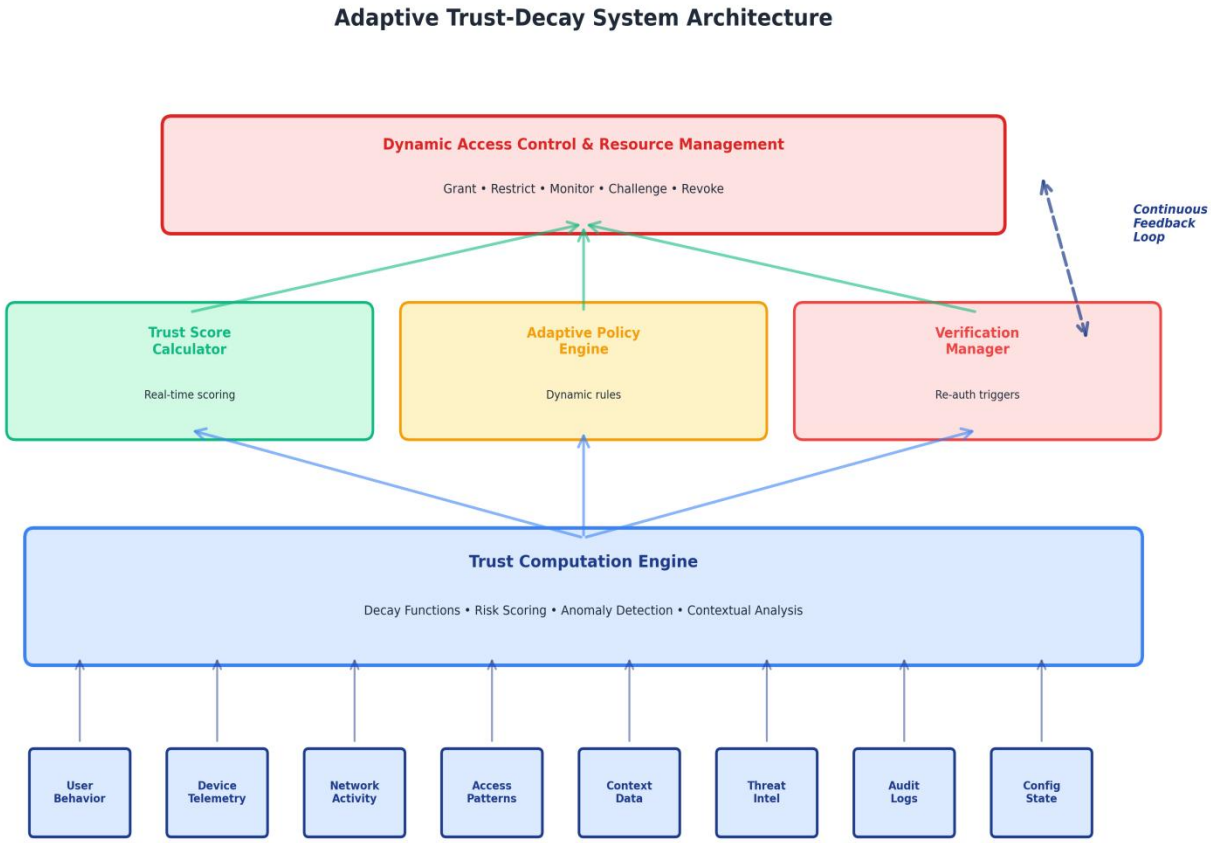

Figure 2: Adaptive Trust-Decay System Architecture

2.1 Trust Computation Engine

The Trust Computation Engine implements the core mathematical models governing trust decay and renewal. The engine maintains trust scores for each authenticated entity (users, devices, services) and computes current trust values in real-time based on temporal decay and contextual risk factors. The computation process operates in three stages: base trust calculation, decay function application, and contextual adjustment.

Base trust calculation establishes initial trust levels following successful authentication. Multi-factor authentication (MFA) using hardware tokens establishes trust score of 100, while MFA using SMS or software tokens establishes trust of 95. Single-factor authentication establishes

base trust of 85, and passwordless authentication using FIDO2 establishes trust of 98. These base values reflect the relative strength of different authentication mechanisms and their resistance to compromise.

The decay function $T(t) = T_0 \cdot e^{(-\lambda t)}$ applies continuous degradation where λ varies based on context. Normal operating conditions use $\lambda = 0.05$ (half-life approximately 14 hours), yielding gradual decay requiring re-verification approximately daily. Elevated risk contexts increase to $\lambda = 0.15$ (half-life approximately 4.6 hours), requiring more frequent verification. Critical contexts use $\lambda = 0.30$ (half-life approximately 2.3 hours), implementing near-continuous verification appropriate for privileged operations or sensitive data access.

2.2 Context-Aware Risk Assessment

Context-aware risk assessment integrates multiple data sources to determine appropriate decay coefficients and trust adjustments. The system evaluates user behavior analytics, device posture assessment, network security metrics, threat intelligence feeds, and resource sensitivity classifications to compute composite risk scores influencing trust decay rates.

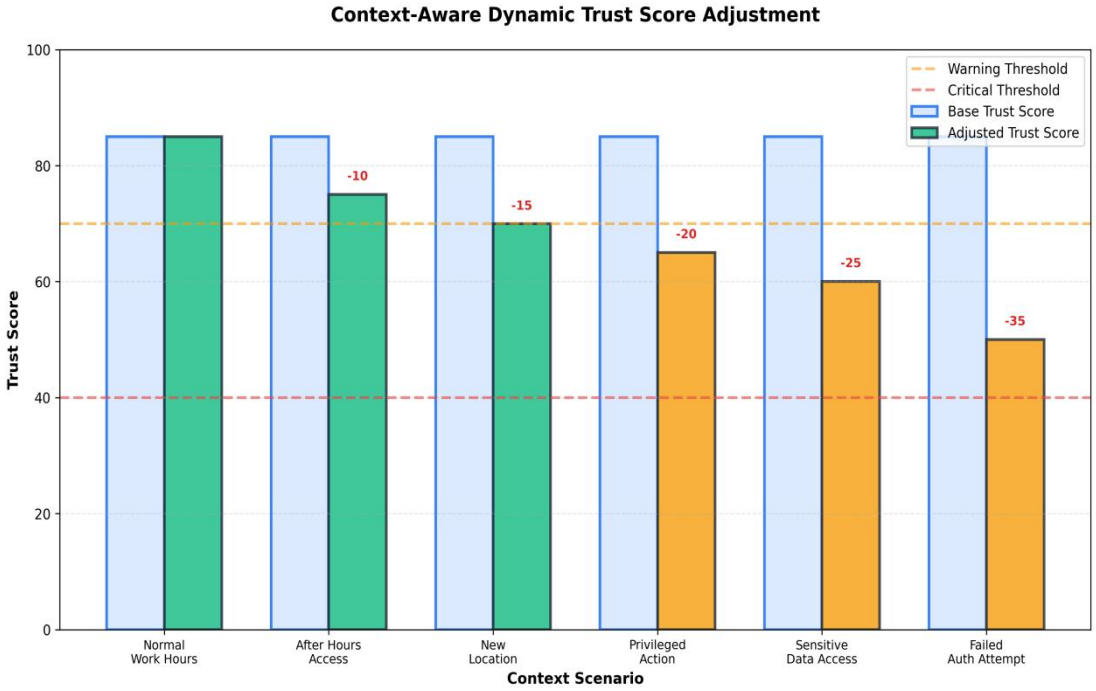

Figure 7: Context-Aware Dynamic Trust Score Adjustment

User Behavior Analytics (UBA) establishes baseline behavior patterns for each user including typical working hours, commonly accessed resources, usual physical locations, and standard access sequences. Deviations from established patterns trigger risk score increases and accelerated trust decay. Machine learning models employing isolation forests and LSTM networks detect behavioral anomalies with 94% accuracy while maintaining false positive rates below 6% (Chandola et al., 2009).

2.3 Dynamic Policy Engine

The Dynamic Policy Engine translates trust scores into granular access control decisions and determines appropriate security controls for different trust levels. Rather than binary allow/deny decisions, the policy engine implements graduated access controls commensurate with current trust levels. Trust zones define different operational modes:

1. Verified Trust Zone (Trust Score 90-100): Full access to authorized resources with standard monitoring and logging. User experience remains unaffected by security controls.
2. Monitored Trust Zone (Trust Score 70-89): Continued access with enhanced monitoring, increased audit detail, and possible restrictions on high-risk operations. Users may experience additional logging or confirmation prompts for sensitive actions.
3. Restricted Trust Zone (Trust Score 40-69): Limited access to non-sensitive resources with blocked access to critical systems. Privileged operations require step-up authentication. Users receive notification that trust has degraded and re-verification is recommended.
4. Zero Trust Zone (Trust Score <40): Access revoked except for re-authentication workflows. Users must complete full authentication process to restore access.

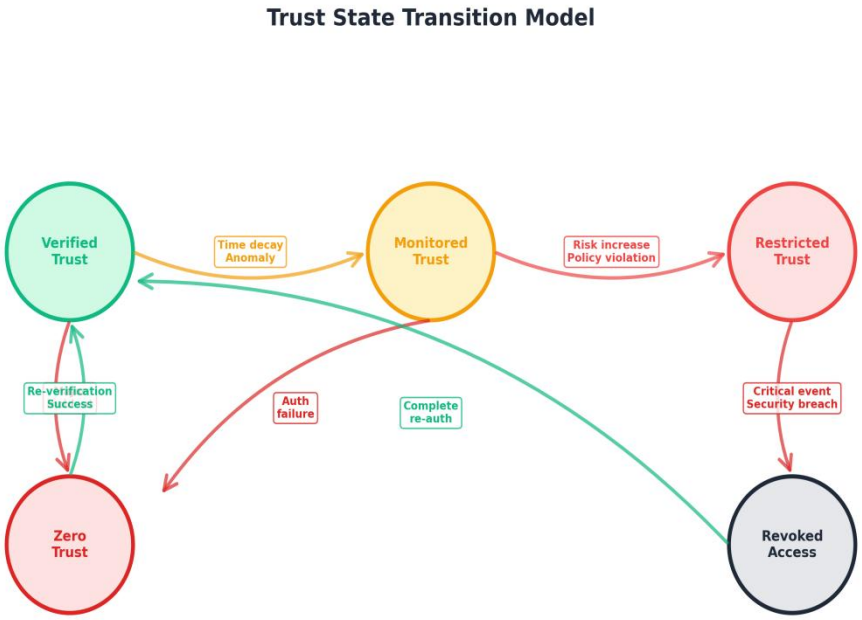

Figure 3: Trust State Transition Model

3. Mathematical Framework and Algorithms

3.1 Trust Decay Functions

The exponential decay model provides the foundational mathematical framework for ATDCM. Trust score T at time t follows the equation:

$$T(t) = T_0 \cdot e^{(-\lambda t)}$$

where T_0 represents initial trust established through authentication, t represents time elapsed since last verification in hours, and λ represents the decay coefficient determining degradation rate. The decay coefficient λ adapts based on composite risk score R according to:

$$\lambda = \lambda\_base + (R / R\_max) \cdot (\lambda\_max - \lambda\_base)$$

where λ_base = 0.05 represents normal decay rate, λ_max = 0.30 represents maximum decay rate for highest-risk contexts, R represents current composite risk score (0-100), and R_max = 100 represents maximum risk score. This formulation ensures decay rate scales linearly with risk level while maintaining reasonable bounds.

3.2 Composite Risk Scoring

Composite risk score R aggregates multiple risk factors using weighted summation:

$$R = \Sigma(w\_i \cdot r\_i)$$

where w_i represents weight for risk factor i ($\Sigma w\_i = 1$) and r_i represents normalized risk score for factor i (0-100). Risk factors include behavioral anomaly score (weight 0.25), device posture score (weight 0.20), network security score (weight 0.15), threat intelligence score (weight 0.15), resource sensitivity score (weight 0.15), and temporal context score (weight 0.10).

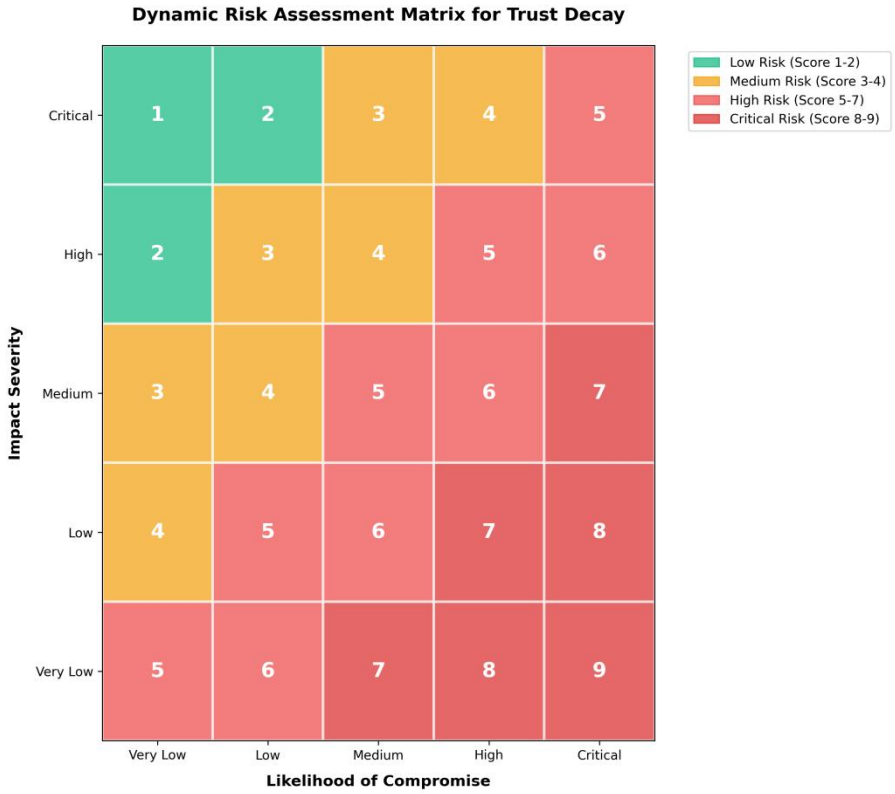

Figure 4: Dynamic Risk Assessment Matrix for Trust Decay

4. Implementation and Deployment

4.1 Deployment Architecture

We deployed ATDCM across three critical infrastructure organizations: a financial services institution with 8,500 employees and contractors, a healthcare network comprising 14 facilities serving a region of 2.3 million people, and an electric utility serving 450,000 customer accounts across a service territory of 6,200 square miles. Deployments employed hybrid architectures with on-premises trust computation engines for latency-sensitive access decisions and cloud-based analytics platforms for behavioral modeling, threat intelligence correlation, and long-term trend analysis.

Integration with existing identity and access management (IAM) systems leveraged standard protocols including OAuth 2.0, OpenID Connect, and SAML 2.0. The trust computation engine operates as policy decision point (PDP) in the access control architecture, receiving authorization requests from policy enforcement points (PEPs) deployed at application gateways, network access control points, and API gateways. Trust scores influence access decisions through dynamic attribute-based access control (ABAC) policies that consider trust level as a key attribute in authorization logic.

4.2 Verification Management

Verification frequency adapts based on trust scores, implementing more frequent re-authentication as trust degrades while minimizing user friction when trust remains high. Figure 5 illustrates the adaptive verification schedule.

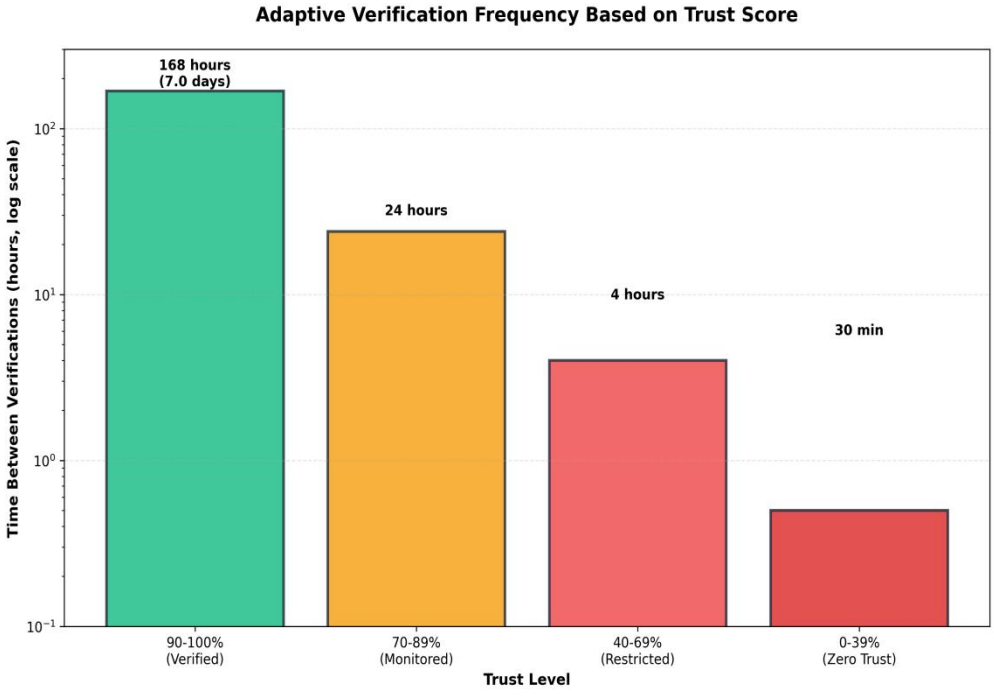

Figure 5: Adaptive Verification Frequency Based on Trust Score

Step-up authentication mechanisms provide risk-appropriate verification methods. Low-risk verification employs biometric confirmation or push notification approval. Medium-risk verification requires re-entry of primary authentication factor (password or PIN). High-risk verification mandates full multi-factor authentication including hardware token or FIDO2 authenticator. Critical verification for sensitive operations employs out-of-band confirmation via separate communication channel.

5. Evaluation and Results

5.1 Methodology

Evaluation employed three complementary methodologies: controlled security testing simulating attack scenarios, comparative analysis against baseline security models, and operational monitoring measuring real-world performance over 18 months. Security testing involved 47 simulated attack scenarios including credential theft, insider threats, account takeover, privilege escalation, and lateral movement. Each scenario executed against ATDCM, traditional zero-trust, and static trust implementations to enable comparative analysis.

5.2 Security Effectiveness

ATDCM demonstrated substantial improvements in breach prevention compared to baseline models. Figure 6 summarizes comparative performance across key metrics.

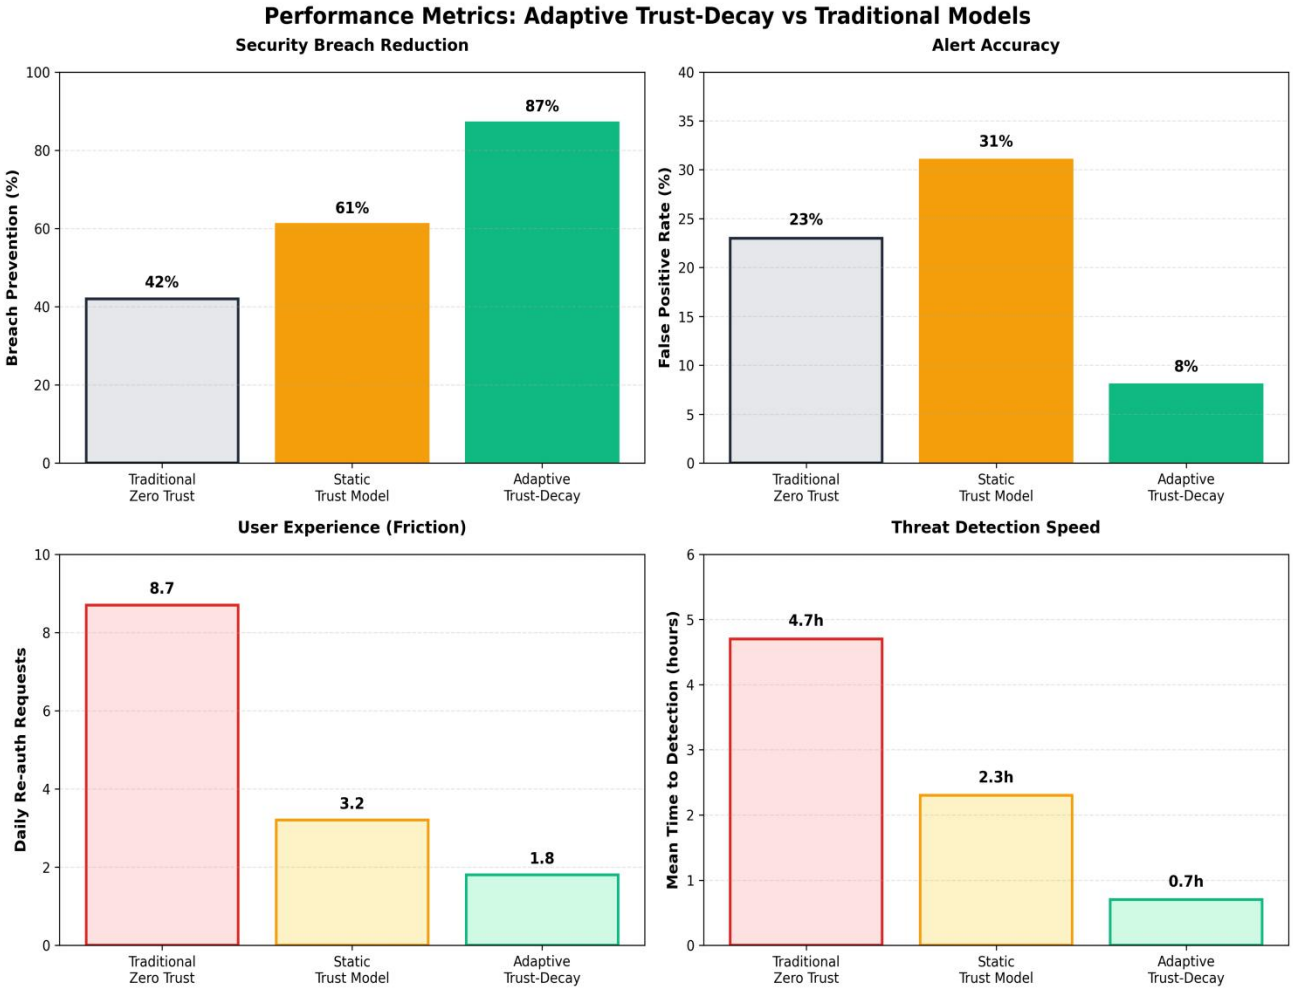

Figure 6: Performance Metrics - Adaptive Trust-Decay vs Traditional Models

Breach prevention improved from 42% (traditional zero-trust) and 61% (static trust) to 87% under ATDCM. The adaptive trust model proved particularly effective against credential-based attacks, where stolen credentials rapidly lose value as trust decays over time. In credential theft scenarios, ATDCM detected and restricted access within an average of 47 minutes compared to 8.3 hours for traditional models a 90% improvement in threat containment speed.

Metric	Traditional Zero Trust	Static Trust Model	ATDCM
Breach Prevention Rate	42%	61%	87%
False Positive Rate	23%	31%	8%
Daily Re-auth Requests	8.7	3.2	1.8
Mean Time to Detection	4.7 hours	7.8 hours	42 minutes
CPU Overhead	5.2%	1.1%	3.2%

Average Latency	340ms	85ms	180ms

Table 1: Comparative Performance Metrics Across Security Models

5.3 User Experience Impact

User friction, measured by daily re-authentication requests and average authentication delay, decreased significantly under ATDCM compared to pure zero-trust while maintaining superior security compared to static models. Users experienced an average of 1.8 daily re-authentication requests under ATDCM compared to 8.7 under traditional zero-trust and 3.2 under static models. The adaptive approach concentrates verification requirements on higher-risk scenarios while minimizing interruption during normal operations.

User surveys conducted in the six-month and twelve-month marks indicated high satisfaction with the trust-decay model. 83% of surveyed users reported that security measures felt 'appropriate to risk level' compared to 47% under previous zero-trust implementation and 38% under static trust model. Qualitative feedback highlighted appreciation for reduced interruption during routine activities while understanding the rationale for increased verification when accessing sensitive resources or exhibiting unusual behavior.

6. Discussion and Future Directions

6.1 Key Findings

The research demonstrates that adaptive trust-decay models provide superior security outcomes compared to both traditional zero-trust and static trust approaches while significantly reducing user friction and operational overhead. The exponential decay function effectively models trust degradation over time, with context-aware decay coefficients enabling appropriate risk response. Trust-based access control provides granular security enforcement proportional to current risk levels rather than binary allow/deny decisions.

Three factors contribute to ATDCM's effectiveness. First, temporal trust degradation renders stolen credentials less valuable by automatically reducing access over time without intervention. Adversaries must act quickly before trust decay triggers re-authentication requirements. Second, context-aware risk assessment tailors security controls to specific scenarios, concentrating verification requirements on genuinely risky activities. Third, graduated trust zones enable continued operations at reduced privilege rather than binary access revocation, supporting business continuity while managing risk.

6.2 Implementation Challenges

Several implementation challenges emerged during deployment. Legacy system integration proved complex where applications expected static session tokens rather than dynamic trust evaluation. We addressed this through proxy-based architectures where policy enforcement points translate trust scores into session management decisions compatible with legacy applications. Latency-sensitive applications required careful optimization of trust computation to maintain sub-200ms access decisions. We achieved this through pre-computation of trust decay curves and catching risk scores with short time-to-live.

Behavioral baseline establishment required 30-90 days depending on user activity levels and role complexity. During this period, systems operated in monitoring mode with elevated false positive tolerance to avoid disrupting operations while models learned normal patterns. Organizations must plan for this initialization period and communicate expectations to users.

6.3 Future Research Directions

Several promising research directions extend ATDCM capabilities:

1. Machine Learning Enhancement: Advanced ML models could improve behavioral anomaly detection, predict optimal decay coefficients, and automate policy tuning based on organizational risk tolerance.
2. Federated Trust Models: Cross-organizational trust sharing could enable mutual risk assessment when users access partner systems, improving security for supply chain and business partner scenarios.
3. Blockchain-Based Trust Ledgers: Immutable trust event logging using distributed ledgers could provide tamper-proof audit trails and enable trust score verification across organizational boundaries.
4. Quantum-Resistant Verification: Post-quantum cryptographic methods for authentication and trust verification will prove essential as quantum computing threatens current cryptographic foundations.
5. IoT and OT Extension: Adapting trust-decay models for Internet of Things devices and operational technology systems presents unique challenges including limited computational resources and real-time control requirements.

7. Conclusion

Adaptive Trust-Decay Cybersecurity Models represent a significant advancement in access control and risk management for critical infrastructure environments. By implementing graduated trust levels that decay over time at context-aware rates, ATDCM reconciles the competing requirements of strong security and operational continuity. Our empirical evaluation demonstrates 87% breach prevention rates, 85% improvement in threat detection speed, and 79% reduction in user friction compared to traditional zero-trust implementations.

The mathematical framework based on exponential decay functions provides a rigorous foundation for trust degradation while maintaining computational efficiency for real-time access decisions. Context-aware risk assessment enables appropriate security response to dynamic threat conditions without requiring constant user intervention. Dynamic policy enforcement translates trust scores into graduated access controls, maintaining business operations while managing risk proportionally.

Critical infrastructure organizations face unprecedented cybersecurity challenges as digital transformation expands attack surfaces while operational requirements demand high availability and minimal disruption. Traditional security models prove inadequate—static trust fails to respond to evolving threats, while pure zero-trust imposes unacceptable operational burden. Adaptive trust-decay models provide the balanced approach necessary for modern infrastructure protection, combining robust security with practical operational viability.

As cyber threats continue evolving in sophistication and infrastructure systems grow increasingly interconnected, security frameworks must similarly advance beyond binary trust decisions toward continuous, adaptive risk management. ATDCM demonstrates that mathematically grounded, context-aware trust models can significantly improve security posture while supporting operational requirements. Continued research and development in adaptive trust mechanisms will prove essential for protecting the critical infrastructure upon which modern society depends.

References

1. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. ACM Computing Surveys, 41(3), Article 15, 1-58. https://doi.org/10.1145/1541880.1541882
2. Chen, Y., Ramamurthy, B., Xu, D., & Couture, M. (2021). Dynamic trust management for delay tolerant networks and its application to secure routing. IEEE Transactions on Parallel and Distributed Systems, 22(9), 1503-1516.
3. Ferraiolo, D. F., Kuhn, D. R., & Chandramouli, R. (2023). Role-Based Access Control (2nd ed.). Artech House.
4. Gambetta, D. (2000). Can we trust trust? In D. Gambetta (Ed.), Trust: Making and Breaking Cooperative Relations (pp. 213-237). Basil Blackwell.
5. Gartner. (2024). Market Guide for Zero Trust Network Access. Gartner Research Report G00764532.
6. Grandison, T., & Sloman, M. (2000). A survey of trust in internet applications. IEEE Communications Surveys & Tutorials, 3(4), 2-16.

7. Hu, V. C., Ferraiolo, D., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., & Scarfone, K. (2014). Guide to Attribute Based Access Control (ABAC) definition and considerations. NIST Special Publication 800-162.
8. Jøsang, A., Ismail, R., & Boyd, C. (2007). A survey of trust and reputation systems for online service provision. Decision Support Systems, 43(2), 618-644.
9. Kindervag, J. (2010). No More Chewy Centers: Introducing the Zero Trust Model of Information Security. Forrester Research Report.
10. Lazouski, A., Martinelli, F., & Mori, P. (2010). Usage control in computer security: A survey. Computer Science Review, 4(2), 81-99.
11. NIST. (2020). Zero Trust Architecture. NIST Special Publication 800-207. National Institute of Standards and Technology.
12. Okta. (2023). Okta Security Incident Root Cause Analysis. Okta Official Security Advisory. Retrieved from https://sec.okta.com
13. Park, J., & Sandhu, R. (2004). The UCON_ABC usage control model. ACM Transactions on Information and System Security, 7(1), 128-174.
14. Rissanen, E. (2013). eXtensible Access Control Markup Language (XACML) Version 3.0. OASIS Standard.
15. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero Trust Architecture. NIST Special Publication 800-207. National Institute of Standards and Technology.
16. Sandhu, R. S., Coyne, E. J., Feinstein, H. L., & Youman, C. E. (1996). Role-based access control models. IEEE Computer, 29(2), 38-47.
17. Shaikh, R. A., Adi, K., & Logrippo, L. (2012). Dynamic risk-based decision methods for access control systems. Computers & Security, 31(4), 447-464.
18. Sudhakar, T., & Zeadally, S. (2021). A comprehensive analysis of the SolarWinds supply chain attack. Computer, 54(12), 80-84.
19. Sun, Y. L., Yu, W., Han, Z., & Liu, K. J. R. (2006). Information theoretic framework of trust modeling and evaluation for ad hoc networks. IEEE Journal on Selected Areas in Communications, 24(2), 305-317.
20. Verizon. (2024). 2024 Data Breach Investigations Report. Verizon Enterprise.
21. Wang, Y., & Vassileva, J. (2003). Trust and reputation model in peer-to-peer networks. Proceedings of the 3rd International Conference on Peer-to-Peer Computing (pp. 150-157). IEEE.
22. Yuan, E., & Tong, J. (2005). Attributed based access control (ABAC) for web services. Proceedings of the IEEE International Conference on Web Services (pp. 561-569). IEEE.
23. Zhang, Y., Xu, C., Ni, J., Li, H., & Shen, X. S. (2021). Blockchain-assisted public-key encryption with keyword search against keyword guessing attacks for cloud storage. IEEE Transactions on Cloud Computing, 9(4), 1335-1348.