

Contextual Risk-Based Verification for Strengthening JWT-Based Authentication

Arben Kola

Department of Emerging Engineering and Technology, Albanian Institute of Applied Engineering, Albania

Elira Dervishi

Faculty of Advanced Engineering Technologies, Tirana Institute of Technology, Albania

Received: 01 July 2026; **Accepted:** 16 August 2026; **Published:** 03 September 2026

Abstract: JSON Web Token (JWT)-based authentication provides a scalable mechanism for representing authenticated identity and authorization claims across distributed applications, APIs, and microservices. However, conventional JWT validation generally emphasizes cryptographic validity, issuer, audience, expiration, and token structure while providing limited consideration of the contextual circumstances in which a valid token is presented. This creates an important security limitation: possession of a correctly signed and unexpired token may be treated as sufficient evidence of legitimacy even when the surrounding request context is anomalous. This paper proposes a contextual risk-based verification framework that supplements conventional JWT validation with adaptive assessment of contextual signals, including device characteristics, request origin, temporal behavior, access patterns, session consistency, and resource sensitivity. The methodology conceptualizes authentication as a continuous risk evaluation process rather than a one-time token-validation event. A multi-stage architecture is developed consisting of cryptographic validation, contextual feature extraction, risk scoring, policy evaluation, and adaptive response. The analysis indicates that contextual verification can strengthen JWT-based authentication by distinguishing technically valid tokens from potentially suspicious token usage. The approach also introduces trade-offs involving latency, privacy, false positives, implementation complexity, and policy calibration. The proposed framework therefore positions JWT as one component of a broader adaptive authentication architecture rather than as an independently sufficient security control.

Keywords: JSON Web Token, JWT Authentication, Contextual Verification, Risk-Based Authentication, Adaptive Security, Token Security, Identity Verification, Access Control, Anomaly Detection

1. INTRODUCTION

Modern web applications increasingly depend on distributed authentication mechanisms that allow identity information to move securely between clients, APIs, microservices, and authorization services. JSON Web Token (JWT) has become an important mechanism for representing authentication and authorization claims because tokens can be validated without maintaining extensive server-side session state. This characteristic is particularly valuable in distributed architectures in which multiple services must independently evaluate

authentication information.

Despite these advantages, the possession of a valid JWT does not necessarily establish that every subsequent request made with that token is legitimate. A token may remain cryptographically valid after being copied, exposed, or replayed from an unexpected environment. Conventional validation can therefore establish token integrity without fully establishing contextual legitimacy. The distinction is fundamental: cryptographic validity answers whether

a token is authentic, whereas contextual verification evaluates whether its current use is consistent with expected behavior.

The security problem addressed in this paper is consequently the gap between static token validation and dynamic authentication assurance. A conventional JWT workflow may validate the signature, expiration time, issuer, audience, and relevant claims before granting access. However, an attacker possessing a legitimate token may satisfy those conditions without being the legitimate user or trusted device. A stronger architecture must therefore evaluate both the token and the circumstances surrounding its presentation.

Token-binding and contextual verification provide a conceptual foundation for addressing this limitation. Ganapathy (2025) positions contextual verification as a mechanism for strengthening JWT authentication by considering information beyond conventional token validation. This perspective supports the central argument of the present study: authentication assurance should depend not only on what a token contains but also on where, when, how, and under what behavioral conditions the token is being used (Ganapathy, 2025).

The objective of this paper is to develop a research-oriented framework for contextual risk-based verification of JWT authentication. The proposed approach integrates conventional JWT validation with contextual feature analysis and adaptive risk decisions. The study specifically examines how contextual information can improve resistance to token replay, credential misuse, anomalous access, and compromised-session scenarios while preserving the scalability advantages of token-based authentication.

The scope of the paper is architectural and analytical rather than experimental. It proposes a framework, evaluates its expected security behavior, and identifies implementation limitations. No fabricated benchmark results are presented because the supplied reference set does not contain an experimental dataset concerning JWT authentication.

2. LITERATURE REVIEW

The provided literature contains two distinct knowledge groups. The first group directly concerns JWT authentication through the work of Ganapathy (2025), while the remaining references investigate corrosion protection, coatings, electrospinning,

nanocomposites, and surface modification. Although these domains are technically different, several methodological principles can be used cautiously as conceptual analogies for layered protection and surface-level security enhancement. Such analogies should not be interpreted as evidence that corrosion studies experimentally validate JWT security.

Ganapathy (2025) provides the most direct theoretical positioning for contextual JWT verification. The underlying principle is that authentication can be strengthened when token validation is combined with additional contextual or binding information. Rather than considering a token as an isolated bearer credential, contextual verification introduces an additional relationship between the token and the environment in which authentication occurs. This creates a stronger authentication condition because successful possession of the token alone becomes insufficient under selected security policies.

The corrosion-oriented studies illustrate another useful conceptual principle: a base protection mechanism can be strengthened through additional layers or treatments. Behzadnasab et al. (2011), for example, examine epoxy coatings containing silane-treated ZrO₂ nanoparticles for corrosion protection, while Bagherzadeh et al. (2016) investigate magnetite/PANI nanocomposite protection. Bakhsheshi-Rad et al. (2019) examine electrospun coating structures for enhanced corrosion resistance. These studies are not authentication studies, but they demonstrate a broader engineering concept in which a primary protective layer can be supplemented with additional mechanisms to improve resistance under adverse conditions.

Similarly, Li and Xia (2004) discuss electrospinning as a technique for producing nanofibers, while Panahi et al. (2020) and Karimi et al. (2020) investigate electrospun polymeric structures and their functional properties. From a methodological perspective, these works reinforce the idea that protection can be engineered through structural layering and functional modification. In the JWT domain, contextual verification applies an analogous systems principle: cryptographic validation constitutes the foundational layer, while contextual risk assessment functions as an additional adaptive layer.

Other supplied references examine specialized coating and surface-enhancement mechanisms. Wan et al. (2017) investigate PVD-CrN coatings with ALD sealing layers, and Zhao et al. (2020) examine

composite coatings with core-shell nanofibers and their mechanical and anticorrosive performance. Again, these studies cannot be directly translated into authentication evidence. Their relevance is limited to the conceptual observation that a primary protective mechanism may be reinforced by secondary controls designed for specific failure conditions.

The major research gap in the supplied literature is therefore clear. There is insufficient directly relevant literature in the provided reference set to support an empirical review of contextual JWT risk scoring, token replay detection, device binding, or adaptive authorization. Consequently, the present study does not claim that the corrosion literature establishes these security mechanisms. Instead, it uses the directly relevant JWT work as the principal foundation and develops a new conceptual framework from established authentication principles. Ganapathy (2025) specifically supports the proposition that contextual verification can strengthen JWT authentication, providing the central theoretical basis for the framework developed here.

3. METHODOLOGY

3.1 Research Design

This study adopts a conceptual framework-development methodology. The research process begins by identifying the principal weakness of conventional JWT authentication: the difference between token validity and request legitimacy. The framework then introduces contextual signals as supplementary evidence and establishes a risk-based decision process.

The methodology is structured into five sequential layers: token validation, contextual feature extraction, risk calculation, policy evaluation, and adaptive response. This architecture preserves the conventional JWT mechanism while introducing an additional decision layer capable of responding to contextual anomalies.

The fundamental authentication condition can be represented as:

Authentication Assurance = Token Validity + Contextual Consistency + Risk-Based Policy Decision

This expression does not imply that these factors are mathematically interchangeable. Instead, it represents a layered decision model in which cryptographic verification establishes a baseline

condition and contextual information determines whether that baseline is sufficient for the requested operation.

3.2 Layer 1: Cryptographic JWT Validation

The first layer performs conventional JWT validation. The receiving service verifies the token's cryptographic signature and evaluates security-relevant claims such as issuer, audience, expiration, and other application-defined constraints.

This layer remains essential because contextual verification cannot compensate for an invalid or forged token. A token that fails cryptographic validation should be rejected before contextual processing. The architecture therefore follows a defense-in-depth principle: contextual verification supplements rather than replaces cryptographic authentication.

For example, if a token has expired, the request should be rejected immediately rather than subjected to further contextual analysis. Similarly, an invalid signature should result in denial because contextual consistency cannot establish the authenticity of an invalid token.

3.3 Layer 2: Contextual Feature Extraction

After successful token validation, the system evaluates contextual characteristics associated with the request. Relevant signals may include device identity, network origin, approximate geographic consistency, access time, user-agent characteristics, session continuity, request frequency, resource sensitivity, and historical behavioral patterns.

The purpose is not to collect every available attribute. Excessive contextual collection can increase privacy risks and operational complexity. Instead, features should be selected according to their security relevance.

A simple contextual representation can be expressed as:

$C = \{D, N, T, B, S, R\}$

where D represents device context, N network context, T temporal context, B behavioral context, S session consistency, and R resource sensitivity.

Consider a user who normally accesses an application from a recognized device during working hours. A token request originating from the same device and

consistent network environment would generate comparatively low contextual risk. Conversely, a token suddenly used from a previously unseen device with unusually high request frequency could produce elevated risk even though the JWT itself remains valid.

3.4 Layer 3: Risk Scoring

The extracted contextual features are transformed into a risk score. A conceptual model may be represented as:

$$\text{Risk Score} = \sum (w_i \times r_i)$$

where r_i represents the risk contribution of contextual feature i , and w_i represents its policy-defined weight.

The purpose of the model is not necessarily to produce a universally standardized score. Different applications require different risk thresholds. A banking API, for example, may assign greater weight to device inconsistency and unusual transaction behavior than a low-sensitivity information portal.

Three broad risk levels can be established. Low-risk requests may proceed normally. Medium-risk requests may trigger additional verification. High-risk requests may be rejected or require administrative intervention.

This adaptive structure is particularly important because binary authentication can be insufficient for modern distributed systems. Ganapathy (2025) emphasizes contextual verification as a strengthening mechanism for JWT authentication; the present framework operationalizes that concept by transforming contextual evidence into explicit risk decisions.

3.5 Layer 4: Policy Evaluation

Risk scores must be interpreted through security policies rather than treated as autonomous authorization decisions. The policy engine maps risk levels to appropriate responses.

For low-risk access, the system can permit the request without additional friction. For medium-risk access, step-up authentication or additional verification can be required. For high-risk access, the system can deny the request, revoke the associated session, or require reauthentication.

Resource sensitivity should also influence policy. A

medium-risk request for a public profile may not justify the same response as a medium-risk request for confidential financial or administrative information.

This creates a two-dimensional decision structure involving risk level and resource sensitivity. Such a design reduces the likelihood that contextual anomalies automatically create unnecessary authentication barriers for low-value operations.

3.6 Layer 5: Adaptive Response

The final layer applies the selected response. Possible actions include allowing access, requesting additional verification, reducing privileges, limiting request frequency, terminating the session, or requiring issuance of a new token.

An important advantage of this architecture is that the system does not need to treat authentication as a permanent state. Instead, authentication assurance can change as contextual conditions change.

For example, a token initially used under normal circumstances may later appear in an abnormal environment. A contextual risk engine can detect this transition and respond accordingly. This is stronger than relying exclusively on the original authentication event because it recognizes that security conditions may change after token issuance.

3.7 Hypothetical Operational Scenario

Consider a microservice application in which a user authenticates successfully and receives a JWT. The token contains valid identity and authorization claims. During normal operation, requests originate from the user's recognized device and follow established behavioral patterns. The risk engine therefore assigns a low contextual risk.

Suppose the same token is subsequently presented from an unfamiliar device and begins generating requests at a substantially higher frequency. Cryptographic validation still succeeds, but contextual consistency deteriorates. The system therefore increases the risk score. If the requested resource is highly sensitive, the policy engine may require step-up verification rather than allowing direct access.

This scenario demonstrates the principal value of contextual verification: the system does not declare the token invalid merely because the environment

changed. Instead, it recognizes uncertainty and applies an adaptive control proportional to the risk.

3.8 Evaluation Criteria

The framework should be evaluated using five principal criteria: security effectiveness, false-positive rate, authentication latency, operational scalability, and privacy impact.

Security effectiveness concerns the ability to identify suspicious token use that would pass ordinary JWT validation. False-positive rate measures the frequency with which legitimate users are challenged because of unusual but harmless contextual changes. Latency determines whether contextual analysis creates unacceptable delays in API processing. Scalability examines whether the architecture remains practical across large distributed systems. Privacy impact considers whether contextual data collection introduces unnecessary surveillance or data-retention risks.

4. RESULTS

The conceptual analysis indicates that contextual risk-based verification can strengthen JWT authentication primarily by addressing a limitation that cryptographic validation alone cannot resolve: the possibility that a legitimate token is used under illegitimate circumstances. A correctly signed JWT establishes token authenticity, but it does not independently prove that the current requester, device, environment, and behavioral context remain trustworthy.

The first finding is that authentication assurance becomes multidimensional when contextual evidence is incorporated. Instead of treating authentication as a binary valid/invalid condition, the proposed framework establishes a continuum of confidence. This permits the system to distinguish normal token use from anomalous token use without immediately invalidating every unusual request. Such contextual strengthening is consistent with the central proposition of Ganapathy (2025).

The second finding concerns token replay resistance. A stolen JWT can potentially satisfy ordinary signature and expiration checks. Contextual verification introduces additional conditions that an attacker may be unable to reproduce consistently. A token used from an unfamiliar device, abnormal network environment, or inconsistent behavioral pattern can therefore be subjected to increased scrutiny.

Contextual verification does not eliminate replay attacks completely, but it can reduce the likelihood that a replayed token receives unrestricted access.

The third finding is the importance of adaptive response rather than universal rejection. Not every contextual anomaly indicates compromise. Legitimate users may change networks, devices, locations, or access times. A rigid rule would therefore create excessive false positives. The proposed risk-based model addresses this issue by allowing different responses according to risk severity and resource sensitivity.

The fourth finding is that resource sensitivity should be integrated into authentication decisions. A contextual anomaly associated with a low-impact operation may justify monitoring rather than immediate denial. The same anomaly associated with privileged administrative functionality should produce a stronger response. This makes contextual authentication more proportional and operationally practical.

The fifth finding is that layered protection improves architectural resilience. The supplied coating literature, while unrelated to cybersecurity, repeatedly demonstrates the broader engineering principle of strengthening a primary protective mechanism through additional functional layers (Behzadnasab et al., 2011; Bakhsheshi-Rad et al., 2019; Wan et al., 2017). In the present framework, this principle translates conceptually into placing contextual verification above foundational JWT validation rather than attempting to replace cryptographic controls.

Finally, the analysis identifies significant implementation trade-offs. Contextual verification requires additional processing, feature management, policy configuration, and potentially stateful behavioral information. Therefore, the security benefit must be balanced against latency, privacy, operational complexity, and false-positive risk. The framework is consequently best understood as an adaptive security layer whose effectiveness depends heavily on appropriate feature selection and policy calibration.

5. DISCUSSION

The findings support a theoretical shift from static credential validation toward continuous authentication assurance. Conventional JWT authentication is highly effective at validating the

structural and cryptographic properties of a token, but contextual verification changes the security question from “Is this token valid?” to “Is this valid token being used in a context that remains trustworthy?” This distinction is central to strengthening bearer-token architectures.

The principal theoretical implication is that authentication should be treated as an evidence-aggregation process. Token validity represents one category of evidence, whereas device, behavioral, temporal, network, and session characteristics represent additional evidence. Ganapathy (2025) provides the directly relevant foundation for this contextual strengthening model. The present framework extends that concept by organizing contextual evidence into a formal risk-evaluation pipeline.

A major practical implication is improved resilience against token theft. If an attacker obtains a legitimate token, conventional JWT validation may provide little distinction between the legitimate owner and the attacker. Contextual verification introduces additional environmental conditions. The attacker may possess the token but lack the expected device characteristics or behavioral consistency. Consequently, token possession becomes necessary but potentially insufficient for sensitive operations.

However, contextual verification introduces an important contradiction: stronger security can increase authentication friction. Users legitimately change devices, networks, and locations. A system that treats every contextual deviation as malicious will generate excessive step-up challenges and may reduce usability. Therefore, contextual risk should not be interpreted as a deterministic proof of compromise. It should be treated as an uncertainty indicator.

Another trade-off involves privacy. Device fingerprints, behavioral patterns, and location-related information can improve risk assessment but may also constitute sensitive contextual data. A responsible implementation should therefore minimize feature collection, restrict retention, apply access controls to risk data, and avoid collecting information that does not materially improve security.

Latency is another limitation. JWT validation is comparatively efficient because services can often validate tokens locally. Contextual verification may require additional computations, policy lookups,

behavioral comparisons, or risk-service communication. Excessive dependence on remote risk services could undermine the scalability advantages associated with stateless token validation. An efficient implementation should therefore distinguish between lightweight contextual checks that can occur locally and higher-cost analysis reserved for suspicious or sensitive requests.

The provided corrosion literature offers only indirect methodological support for the framework. Studies involving layered coatings and functional modifications demonstrate a broad engineering pattern in which a foundational protective mechanism is reinforced by additional mechanisms (Bagherzadeh et al., 2016; Zhao et al., 2020). However, these studies should not be interpreted as empirical evidence for cybersecurity performance. Their contribution to this paper is conceptual rather than domain-specific.

The most significant limitation of the present study is the absence of a JWT-specific experimental dataset in the supplied references. Consequently, the proposed risk model has not been empirically calibrated against real-world authentication traffic. Future experimental research should evaluate detection accuracy, false-positive rates, response latency, token-replay detection capability, and scalability under realistic workloads.

Another limitation is that risk weights cannot be assumed to be universal. A contextual feature that is highly predictive in one environment may be relatively insignificant in another. Enterprise applications, consumer platforms, financial systems, and internal administrative services may require different risk models. Therefore, contextual verification should be designed as a policy-driven system rather than a fixed universal algorithm.

Overall, the framework indicates that JWT authentication can be strengthened without abandoning the efficiency of token-based architectures. The most effective approach is not to replace JWT but to position it within a broader layered authentication model in which cryptographic validation provides the foundation and contextual risk assessment provides adaptive assurance.

6. CONCLUSION

This paper proposed a contextual risk-based verification framework for strengthening JWT-based authentication. The central argument is that

cryptographically valid tokens should not automatically be treated as permanently trustworthy credentials because token validity does not necessarily establish the legitimacy of the environment in which the token is being used.

The proposed architecture combines five layers: JWT cryptographic validation, contextual feature extraction, risk scoring, policy evaluation, and adaptive response. This structure enables authentication decisions to incorporate device, network, temporal, behavioral, session, and resource-sensitivity information. The resulting model supports graduated responses ranging from normal access to step-up verification and denial.

The principal contribution is the conceptual integration of JWT validation with contextual risk assessment. Rather than treating authentication as a single event, the framework supports continuously evaluated authentication assurance. This approach can improve resilience against token replay and misuse while avoiding the unnecessary replacement of efficient JWT mechanisms.

At the same time, contextual verification introduces challenges involving privacy, latency, false positives, feature reliability, and policy calibration. These limitations demonstrate that contextual authentication should be implemented carefully and proportionally. Future research should develop empirical datasets, compare alternative risk-scoring techniques, evaluate adaptive thresholds, and measure the framework under realistic distributed API workloads.

The broader conclusion is that JWT should be regarded as a foundational authentication mechanism rather than a complete security decision by itself. Contextual risk-based verification provides a practical pathway for transforming static token validation into a more adaptive and security-aware authentication architecture.

REFERENCES

1. Abd El-Lateef, H, M, Mohamed, I,M,A, Zhu, J, H, Khalaf, M,M(2020) An efficient synthesis of electrospun TiO₂-nanofibers/Schiff base phenylalanine composite and its inhibition behavior for C-steel corrosion in acidic chloride environments. J. Taiwan Inst. Chem. Eng.
2. Bagherzadeh, M, Ghahfarokhi, Z,S, Yazdi, E,G (2016) Electrochemical and surface evaluation of the anti-corrosion properties of reduced graphene oxide. RSC Adv. 6,
3. Bagherzadeh, M, Haddadi H, Iranpour M (2016) Electrochemical evaluation and surface study of magnetite/PANI nanocomposite for carbon steel protection in 3.5% NaCl. Prog. Org. Coatings.
4. Bakhsheshi-Rad, H,R, Akbari, M, Ismail, A, F, Aziz M, Hadisi, Z, Pagan, M, Daroonparvar, Chen, X (2019) Coating biodegradable magnesium alloys with electrospun poly-L-lactic acid- α -kermanite-doxycycline nanofibers for enhanced biocompatibility, antibacterial activity, and corrosion resistance. Surf. Coatings Technol.
5. Behzadnasab, M, Mirabedini, S,M, Kabiri K, Jamali, S (2011) Corrosion performance of epoxy coatings containing silane treated ZrO₂nanoparticles on mild steel in 3.5 % NaCl solution.
6. "Comparative Corrosion Inhibition Test Study in Phosphate Buffer Saline for 316L Stainless Steel with ElectrospunPCL Coating with Different Additives"5304Ersin Kamberli1, ETJ Volume 09 Issue 10 October 2024Corros. Sci.
7. Ganapathy, S. K. (2025). Strengthening JWT Authentication Through Token Binding and Contextual Verification. *Frontiers in Emerging Engineering & Technologies*, 2(05), 15–23. Retrieved from <https://irjernet.com/index.php/feet/article/view/jwt-authentication-security>
8. Karimi, Kichi, M, Torkaman, R, Mohammadi, H, Toutounchi, A, Kharaziha, M,Alihosseini, F (2020) Electrochemical and in vitro bioactivity behavior of poly (ϵ -caprolactone) (PCL)-gelatin-forsterite nano coating on titanium for biomedical application. Mater. Today Commun.
9. Li, D, Xia, Y (2004)Electrospinning of nanofibers: Reinventing the wheel?Adv. Mater.
10. Panahi Z, Tamjid, E, Rezaei, M (2020) Surface modification of biodegradable AZ91 magnesium alloy by electrospun polymer nanocomposite: Evaluation of in vitro degradation and cytocompatibility. Surf. Coatings Technol.,
11. Wan T, F, Zhang J, C,M, Kim S,W, Park, Y, Yang, Kim K, H, Kwon S,H (2017)Enhanced Corrosion Resistance of PVD-CrN Coatings by ALD Sealing

Layers Nanoscale Res. Lett.

- 12.** Zhao, X, Yuan, S, Jin, Z, Zhu, Q, Zheng, Q, Jiang, Q, Song, H, Duan, J (2020) Fabrication of composite coatings with core-shell nanofibers and their mechanical properties, anticorrosive performance, and mechanism in seawater.