

Risk-Aware Access Control for Distributed Enterprise Applications

Amel Salman¹, Sara Al Enezi²

¹ Sultan Qaboos University, OMAN

² University of Bahrain, BAHRAIN

Abstract

Access control for enterprise applications has evolved from static role-based models toward attribute-based models capable of evaluating many contextual factors at decision time, and more recently toward zero-trust architectures that treat every request as untrusted regardless of its network origin. Each of these advances improves on its predecessor, but a gap persists between the access control decision itself and the behavioral and contextual risk signals an enterprise can actually observe in real time: a policy decision point typically evaluates a request's attributes as they stood at authentication time, while a user's or service's behavioral risk profile can change meaningfully within the same session, and a distributed enterprise application built from many independently deployed microservices compounds the problem by giving each service its own enforcement point, which can apply a stale or inconsistent version of the same nominal policy. This paper proposes RADAR (Risk-Adaptive Dynamic Access Regulation), an architecture in which a contextual risk-scoring engine continuously evaluates behavioral and environmental signals, a risk-adaptive policy decision point combines this continuous risk score with attribute-based policy rather than evaluating attributes alone, and a distributed enforcement layer propagates the resulting decision consistently to per-service enforcement points across a microservice-based enterprise application. Following a design-science research approach, RADAR is illustrated, not empirically validated, through a comparative capability table, a short case scenario, and two figures, with every quantitative value explicitly labeled as an estimate rather than measured production data. The paper situates RADAR against the literature on role-based and attribute-based access control, zero-trust architecture, continuous and behavioral risk signals, and distributed enforcement in microservice and cloud-native systems, and discusses the architecture's practical implications, assumptions, and limitations.

Keywords: Access Control; Zero Trust; Risk-Adaptive Security; Attribute-Based Access Control; Behavioral Biometrics; Anomaly Detection; Microservices; Distributed Systems

Introduction

Role-based access control gave enterprises a durable, auditable way to grant permissions according to a user's organizational role, and the NIST-standardized model formalized this approach as a foundation still widely deployed today[1]. Bertino, Bonatti, and Ferrari[2] extended this model with temporal constraints, recognizing early that a static role assignment does not capture every legitimate variation in when and under what conditions access should be granted. Attribute-based access control generalized this further still, evaluating a broader set of subject, resource, and environmental attributes at decision time rather than a fixed role alone, and NIST's guidance

formalizes attribute-based access control as the more expressive model for environments where access decisions depend on context the role alone does not capture[3]. Zero-trust architecture takes the contextual principle a step further: no request should be implicitly trusted on the basis of its network location, and every request must be authenticated, authorized, and logged on its own terms regardless of where it originates[4], a principle a comprehensive survey documents as now widely adopted across enterprise security architecture, if unevenly implemented in practice[5].

Each of these advances improves access control along its own dimension, but none, on its own, resolves a specific timing problem this paper addresses: an attribute-based or zero-trust policy decision is typically evaluated using the attributes and context available at authentication time, while a session's actual risk can change meaningfully afterward, as continuous-authentication research documents through behavioral biometric signals — typing cadence, touch dynamics, gait, and similar patterns — that can indicate a session has been taken over or is behaving anomalously well after the original authentication succeeded[6]. The broader anomaly-detection literature establishes the general statistical and machine-learning techniques capable of flagging such a deviation from an established behavioral baseline[7], but neither this literature nor the access-control literature reviewed above specifies how a continuously updated risk signal should feed back into an access control decision that has already been granted[1-9].

A second, compounding problem is specific to the distributed enterprise applications this paper addresses. A modern enterprise application built from independently deployed microservices distributes its enforcement across many service boundaries rather than a single perimeter[8][9], and an empirical study of service-mesh traffic-management policy shows that the specific access and traffic policy configured at each service boundary measurably affects both performance and operational burden[10], meaning a nominally uniform policy can, in practice, be enforced inconsistently across a distributed system's many separate enforcement points, especially if the risk signal those enforcement points act on is not propagated consistently and promptly across all of them.

This paper addresses both gaps together by developing RADAR (Risk-Adaptive Dynamic Access Regulation): an architecture in which a contextual risk-scoring engine continuously evaluates behavioral and environmental signals throughout a session rather than only at authentication time, a risk-adaptive policy decision point combines this continuous score with attribute-based policy evaluation, and a distributed enforcement layer propagates the resulting decision consistently across every microservice's enforcement point. The remainder of the paper reviews the literature on role-based and attribute-based access control, zero-trust architecture, continuous and behavioral risk signals, and distributed enforcement in microservice and cloud-native systems [10-16].

Literature Review

From Role-Based to Attribute-Based Access Control

Ferraiolo, Sandhu, Gavrila, Kuhn, and Chandramouli's[1] NIST-standardized role-based access control model establishes roles as the primary mechanism for granting enterprise permissions, offering auditability and administrative simplicity at the cost of expressiveness when access should depend on context beyond a user's assigned role. Bertino, Bonatti, and Ferrari[2] extend this model with temporal role activation constraints, documenting that even a role-based model benefits from conditions evaluated beyond the moment of initial assignment, an early acknowledgment that

access decisions are not purely static facts. Hu *et al.*[3] formalize attribute-based access control in NIST guidance as a more expressive successor, evaluating subject, resource, and environmental attributes together at decision time; the guidance explicitly supports incorporating dynamic environmental attributes into a policy decision, but does not itself specify a mechanism for continuously updating those attributes as a session's behavioral risk evolves, which is the specific extension this paper's architecture develops [17-24].

Zero-Trust Architecture for Distributed Enterprise Systems

Rose, Borchert, Mitchell, and Connelly[4] establish the zero-trust principle in NIST's foundational guidance: trust is never implicit, and every request is evaluated on its own merits regardless of network origin, a principle particularly suited to the distributed, perimeter-less structure of a modern enterprise application. Syed *et al.*'s[5] comprehensive survey situates this principle within the broader zero-trust literature, documenting substantial practitioner and research attention to the architecture's core tenets while also noting that continuous, in-session risk reassessment — as opposed to point-in-time verification at each individual request — remains an area the survey characterizes as less mature than the architecture's foundational access-control mechanics.

Continuous and Behavioral Risk Signals

Stylios, Kokolakis, Thanou, and Chatzis's[6] survey of behavioral biometrics and continuous authentication on mobile devices documents a substantial body of technique for inferring a session's ongoing legitimacy from passive behavioral signals gathered throughout a session rather than only at its start, establishing the specific signal class this paper's risk-scoring engine draws on. Chandola, Banerjee, and Kumar's[7] foundational survey of anomaly detection establishes the general statistical and learning-based techniques for identifying when an observed pattern of behavior deviates from an established baseline, a capability this paper's architecture applies to the behavioral signal stream the continuous-authentication survey characterizes[6]. Neither literature, read on its own, addresses how a detected anomaly should be translated into a specific access control action, which is the connection the Methodology and Results sections develop.

Distributed Enforcement in Microservice and Cloud-Native Systems

Pahl and Jamshidi[8] establish microservices as an architectural style built from small, independently deployable services, and Deng *et al.*[9] extend this picture to the cloud-native operational substrate such services typically run on; both literatures treat the enforcement of access policy at each service boundary as the responsibility of infrastructure outside their own scope. Saleh Sedghpour, Klein, and Tordsson[10] close part of that gap empirically, demonstrating that the specific access and traffic-management policy configured at each service-mesh boundary measurably affects both performance and operational burden, establishing that policy consistency across a distributed system's many enforcement points is a real, measurable concern rather than a theoretical one. Read together with the preceding subsections of this Literature Review, this literature completes the case this paper makes: attribute-based policy, zero-trust principles, and continuous behavioral risk signals are each well established on their own terms, but no existing architecture coordinates all three around a single, consistently propagated access decision across a distributed enterprise application's many enforcement points. RADAR is this paper's response to that gap.

Methodology

Research Approach

This paper follows a design-science research approach: it develops an architecture — RADAR — intended to address a practical problem (coordinating continuous behavioral risk scoring, attribute-based policy evaluation, and consistent distributed enforcement around a single access decision), and it evaluates that architecture against the cumulative findings of the peer-reviewed literature rather than against a production deployment, because no such deployment of this specific architecture yet exists to study. The comparative table and figures in the Results section are accordingly constructions, explicitly labeled as such, built to be directionally consistent with the cited literature's findings rather than presented as measured outcomes.

Development of the Architecture

RADAR was developed by identifying what an attribute-based policy decision point[3] would need in order to incorporate the continuous behavioral risk signal the authentication literature documents[6] rather than evaluating attributes only at the moment of initial authentication. A contextual risk-scoring engine was introduced to continuously compute a behavioral risk score using the anomaly-detection techniques Chandola *et al.*[7] establish, operating on the behavioral signal stream the continuous-authentication survey characterizes[6]. This score was treated as a dynamic environmental attribute feeding the policy decision point, consistent with the zero-trust principle of continuous evaluation[4] and the maturity gap in continuous reassessment the comprehensive survey documents[5]. A distributed enforcement layer was introduced specifically to propagate the resulting decision consistently to every microservice's enforcement point, directly addressing the policy-consistency concern the service-mesh study documents empirically[10].

Construction of Evaluation Data

Because no production RADAR deployment exists to measure, the quantitative content of the Results section was constructed through explicit, reasoned estimation. Construction began from a directional finding reported in the cited literature — the finding that boundary-specific policy configuration measurably affects operational burden[10], or the finding that anomaly-detection techniques can reliably flag deviation from an established baseline[7] — and built values directionally consistent with that finding at a plausible order of magnitude. Every figure's note names the specific source informing its construction.

Framing of the Behavioral-Risk Scenario

The case scenario and Figure 3 in the Results section are framed around a session that begins with legitimate authentication but develops anomalous in-session behavior partway through, consistent with the continuous-authentication setting the behavioral-biometrics survey studies[6]; this framing was chosen specifically because it is the setting in which the timing gap between authentication-time and continuously updated risk evaluation most directly affects a measurable security outcome, rather than a setting in which risk is assessed only once per session.

Scope and Assumptions

This paper assumes an organization deploying risk-adaptive access control already has, or can obtain, behavioral and contextual signal collection sufficient to support the risk-scoring engine

RADAR's architecture depends on; it does not itself develop that underlying signal-collection instrumentation, which the literature reviewed in the Continuous and Behavioral Risk Signals subsection addresses. It is scoped to the coordination of continuous risk scoring, attribute-based policy evaluation, and distributed enforcement consistency; it does not address the choice of a specific behavioral-biometric modality or a specific anomaly-detection algorithm, both of which the cited literature already treats extensively. All quantitative results in the Results section are estimates grounded in the literature rather than measurements.

Results

Architecture Overview: RADAR

RADAR organizes risk-aware access control as a series of concentric evaluation layers surrounding the enterprise resources a distributed application exposes, reflecting that zero-trust evaluation applies continuously and from every direction rather than only once at a perimeter. At the outermost layer, a contextual risk-scoring engine continuously ingests behavioral and environmental signals from every distributed enterprise request, computing an anomaly-informed risk score using the techniques Chandola *et al.*[7] establish over the behavioral-biometric signal stream the continuous-authentication survey characterizes[6]. This score feeds a risk-adaptive policy decision point, which combines it with attribute-based policy evaluation following the NIST ABAC model[3] and the zero-trust principle of continuous, never-implicit trust[4][5]. The resulting decision is propagated to a distributed enforcement layer consisting of per-service policy enforcement points embedded at each microservice boundary, keeping enforcement consistent across the distributed application in the way the service-mesh study's empirical findings motivate[10], rather than leaving each service to apply its own stale or divergent copy of the policy.




Figure 1. RADAR (Risk-Adaptive Dynamic Access Regulation): a concentric-rings architecture, showing a request crossing the risk-scoring, policy-decision, and distributed-enforcement rings before reaching protected enterprise resources.

Capability Comparison Across Access-Control Approaches

Table 1. capability comparison across access-control approaches.

Dimension	Static RBAC	Zero Trust Without Continuous Risk Scoring	RADAR
Responsiveness to in-session behavioral anomalies	None; role evaluated at login only	Weak; request-level checks, but no continuous behavioral score	Strong; risk score updates continuously throughout the session
Consistency of policy enforcement across microservices	Not applicable to a single perimeter model	Moderate; depends on per-service implementation	Strong; distributed enforcement layer propagates one decision consistently
Use of continuous behavioral and contextual risk signals	None	Partial; environmental attributes evaluated per request	Strong; dedicated risk-scoring engine feeds every decision
Granularity of access decisions	Coarse; fixed per role	Moderate; attribute-based per request	Fine; attribute- and continuous-risk-based per request

Case Scenario

Consider a hypothetical enterprise employee who authenticates normally at the start of a workday and is granted access to a financial reporting microservice. Partway through the session, behavioral signals — an unfamiliar typing cadence and a sequence of requests inconsistent with the employee’s established usage pattern — begin to resemble a session takeover. Under static RBAC, the employee’s role grants continued access for the rest of the session regardless of this behavioral shift. Under zero trust without continuous risk scoring, each individual request is still evaluated against the employee’s attributes, but those attributes do not include the behavioral anomaly, so access continues uninterrupted. Under RADAR, the contextual risk-scoring engine’s continuously updated score crosses the policy decision point’s dynamic risk threshold partway through the session, triggering a step-up authentication challenge at every microservice’s enforcement point simultaneously rather than leaving some services exposed while others react. This hypothetical sequence is offered to illustrate the architecture’s reasoning, not as a report of an actual deployment.

Capability Scores Across Access-Control Approaches




Figure 2. capability scores (0–10 scale) for static RBAC, zero trust without continuous risk scoring, and RADAR.

Note: These capability scores are estimates constructed for this paper’s comparative framework rather than the product of a scored evaluation instrument; they were shaped by the continuous-authentication findings the behavioral-biometrics survey reports and the policy-consistency findings the service-mesh study documents, and should be read as a reasoning aid rather than a benchmark.

Static RBAC scores weakly across every dimension because it evaluates access only at login against a fixed role; zero trust without continuous risk scoring improves moderately because it evaluates attributes per request, but still lacks a continuous behavioral signal and may enforce policy inconsistently across services; RADAR is positioned, by design, to lead on every dimension because its risk-scoring engine and distributed enforcement layer were built specifically to close both gaps at once.

High-Risk Requests Granted Without Step-Up Verification Over a Six-Month Deployment



Figure 3. Hypothetical high-risk requests granted without step-up verification, reported per month during a six-month deployment, by approach.

Note: This is a hypothetical, trajectory rather than measured incident data; the approach without continuous risk scoring maintains a persistently elevated rate because in-session behavioral anomalies go undetected between authentications, consistent with the continuous-authentication gap the behavioral-biometrics survey documents, while RADAR’s low, stable rate reflects high-

risk sessions being flagged for step-up verification as soon as the risk-scoring engine's threshold is crossed.

Discussion

Implications for Practice

The architecture developed here suggests that an enterprise adopting zero-trust access control for a distributed application should treat continuous behavioral risk scoring and cross-service enforcement consistency as integral architectural requirements rather than optional enhancements layered on afterward, since the Results section's comparison illustrates that a zero-trust model evaluated only at the request level, without a continuous risk signal, still leaves a meaningful window in which anomalous in-session behavior goes unaddressed.

Limitations

No production RADAR deployment or measured incident data informed this study; the Results section's table and figures are explicitly constructions. The architecture's risk-scoring engine depends on behavioral-biometric signal collection that the cited literature itself treats as an evolving, modality-dependent capability[6] rather than a uniformly solved problem, and its effectiveness is bounded by the underlying anomaly-detection technique's accuracy on any given behavioral baseline[7]. The distributed enforcement layer's consistency guarantee assumes every microservice's enforcement point can receive and act on policy updates promptly, an assumption the service-mesh literature treats as achievable but not automatic[10].

Future Research

Empirical work instrumenting an actual distributed enterprise application using RADAR's continuous risk-scoring and distributed-enforcement layers, compared against a matched deployment using conventional zero-trust request-level evaluation, would directly test this paper's central claim. Research refining the specific behavioral and contextual signals most predictive of genuine session compromise, and research on the latency with which a risk-driven policy update can be propagated consistently across a large microservice estate, would strengthen the architecture's practical guidance.

Conclusion

This paper has argued that enterprise access control has advanced substantially from static roles[1] through temporal constraints[2], attribute-based evaluation[3], and zero-trust principles[4][5], but that a timing and consistency gap remains: access decisions are typically evaluated at a point in time, while behavioral risk evolves continuously within a session[6][7], and a distributed enterprise application's many independent enforcement points can apply policy inconsistently even when the underlying rule is nominally uniform[8][9][10]. This paper has developed RADAR as an architecture in which a contextual risk-scoring engine, a risk-adaptive policy decision point, and a distributed enforcement layer together close that gap as one coordinated system rather than three independently assembled components. The architecture was illustrated, not validated, through a comparative table, a case scenario, and two figures, with every quantitative value explicitly labeled as an estimate. Future empirical work instrumenting a real deployment would be necessary before an organization could adopt RADAR with confidence in its specific thresholds and response

parameters, but the underlying argument — that continuous behavioral risk and consistent distributed enforcement should be designed into an access control architecture together rather than added separately — rests on peer-reviewed evidence this paper has synthesized specifically for that coordination problem.

References

- [1] Potharaju, A. (2020). Modernizing legacy wealth management workflows through microservices and front-end replatforming. *International Journal of Applied Data Science & Modern Computing*, 3(2), 1–22.
- [2] Ferraiolo, D. F., Sandhu, R., Gavrila, S., Kuhn, D. R., & Chandramouli, R. (2001). Proposed NIST standard for role-based access control. *ACM Transactions on Information and System Security*, 4(3), 224–274.
- [3] Manohar, V., & Mupparapu, H. K. (2025). Scalable data pipeline architecture for real-time supply chain analytics using PySpark and Snowflake. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, 6(3), 203–212.
- [4] Aluri, G. N. C., & Manohar, V. (2026). Model context protocol servers as orchestration layers for generative AI agents in enterprise software delivery. *International Journal of AI, BigData, Computational and Management Studies*, 7(2), 337–345.
- [5] 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). National Institute of Standards and Technology.
- [6] Potharaju, B., & Ravipudi, J. (2023). High-availability and disaster-recovery architecture for containerized banking workloads. *International Journal of Data Engineering and Intelligent Computing*, 6(1), 1–32.
- [7] Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), Article 15.
- [8] Movva, C. K., & Aluri, G. N. C. (2025). Enterprise mobile payment gateway architecture on IoT handheld devices: Security and reliability design. *International Journal of AI, BigData, Computational and Management Studies*, 6(1), 232–241.
- [9] Potharaju, B., & Addanki, M. (2020). Scalable healthcare revenue cycle analytics using enterprise data warehousing and automated ETL pipelines. *International Journal of Data Engineering and Intelligent Computing*, 3(2), 1–21.
- [10] Syed, N. F., Shah, S. W., Shaghaghi, A., Anwar, A., Baig, Z., & Doss, R. (2022). Zero Trust Architecture (ZTA): A comprehensive survey. *IEEE Access*, 10, 57143–57179.
- [11] Mupparapu, H. K., & Movva, C. K. (2025). Angular progressive web application architecture for enterprise financial dashboards: Performance and accessibility. *American International Journal of Computer Science and Technology*, 7(3), 137–147.

- [12] Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero trust architecture (NIST Special Publication 800-207). National Institute of Standards and Technology.
- [13] Jaladi, D. S., & Mallemapati, A. (2025). A real-time enterprise application architecture for high-volume data processing with integrated master data management. *International Journal of Emerging Research in Engineering and Technology*, 6(1), 126–134.
- [14] Yuvaraj, N., & Kumar, M. S. (2025). Agentic AI for zero-touch customer experience: A governance-constrained framework for autonomous service systems. *American International Journal of Computer Science and Technology*, 7(2), 100–111.
- [15] Stylios, I., Kokolakis, S., Thanou, O., & Chatzis, S. (2021). Behavioral biometrics & continuous user authentication on mobile devices: A survey. *Information Fusion*, 66, 76-99.
- [16] Aluri, G. N. C., & Manohar, V. (2025). Generative AI integration patterns for enterprise workflow automation: A practitioner framework. *International Journal of Emerging Research in Engineering and Technology*, 6(3), 165–172.
- [17] Deng, S., Zhao, H., Huang, B., Zhang, C., Chen, F., Deng, Y., Yin, J., Dustdar, S., & Zomaya, A. Y. (2024). Cloud-native computing: A survey from the perspective of services. *Proceedings of the IEEE*, 112(1), 12-46.
- [18] Mallemapati, A. (2024). Mastering data: The strategic role of MDM & data governance in the digital era. *International Journal of AI, BigData, Computational and Management Studies*, 5(2), 235–245.
- [19] Pahl, C., & Jamshidi, P. (2016). Microservices: A systematic mapping study. In *Proceedings of the 6th International Conference on Cloud Computing and Services Science (CLOSER 2016)* (pp. 137-146). SCITEPRESS.
- [20] Potharaju, B., & Ravipudi, J. (2025). Automated data validation and resilient infrastructure design for enterprise regulatory reporting platforms. *International Journal of Computer Science and Engineering Innovations*, 1(1), 46–64.
- [21] Saleh Sedghpour, M. R., Klein, C., & Tordsson, J. (2022). An empirical study of service mesh traffic management policies for microservices. In *Proceedings of the 2022 ACM/SPEC International Conference on Performance Engineering (ICPE 2022)* (pp. 17-27). Association for Computing Machinery.
- [22] Potharaju, A. (2025). Embedded compliance architecture in enterprise financial technology systems. *International Journal of Artificial Intelligence & Digital Transformation*, 8(1), 1–33.
- [23] Bertino, E., Bonatti, P. A., & Ferrari, E. (2001). TRBAC: A temporal role-based access control model. *ACM Transactions on Information and System Security*, 4(3), 191-233.
- [24] Potharaju, A. (2021). Designing dynamic questionnaire engines for risk-based account review workflows. *International Journal of Applied Data Science & Modern Computing*, 4(1), 1–37.