Computational Systems & Infrastructure
SYSTEMS PERSPECTIVE | COMMISSIONED SYNTHESIS

Zero-Trust Service Identity in Cloud-Native Infrastructure

Robert Wilson - Corresponding author

Abstract

This systems perspective examines zero-trust service identity in cloud-native infrastructure. The organizing question is how identity, authorization, and workload provenance should survive rapid scaling and service churn. Ten related scholarly sources are synthesized through a decision-centered framework spanning problem definition, mechanism, measurement, evaluation, implementation, and governance. The review does not invent experiments, pooled estimates, or unreported quantitative results. It instead evaluates the strength and transferability of the available evidence, with particular attention to equating network segmentation with verified workload identity. The resulting framework links technical or empirical performance to explicit use conditions and identifies tests that should precede wider adoption in multi-tenant cloud and edge service meshes.

Keywords: zero trust, cloud security, service identity, microservices, authorization

Synthesis lens	Principal question
Mechanism	the chain linking workload attestation, credential issuance, policy evaluation, network enforcement, and logging
Evaluation	credential compromise, stale policy, confused deputies, revocation latency, and control-plane failure
Primary risk	equating network segmentation with verified workload identity

1. Introduction

Research on zero-trust service identity in cloud-native infrastructure sits at the boundary between a promising component and a consequential decision. The core question is how identity, authorization, and workload provenance should survive rapid scaling and service churn. Answering it requires more than assembling favorable results. It requires a chain of evidence that makes the problem boundary, mechanism, measurement, comparator, uncertainty, and accountable decision visible to the reader.

This article presents a structured narrative review of ten related publications. Sources were selected for topical relevance, traceable publication metadata, and complementary coverage of technical, empirical, and governance questions. No meta-analytic pooling is attempted because the included studies differ in designs, populations, system boundaries, and outcomes. The purpose is decision-oriented synthesis: to identify what each source can support, where transfer remains uncertain, and which evidence would justify the next stage of use.

The central risk is equating network segmentation with verified workload identity. A top-line metric or broad policy label can appear decisive while concealing the conditions that produced it. Accordingly, the synthesis separates observation from interpretation and implementation from validation. This separation is especially important when the same system creates benefits for one user or institution while transferring cost, risk, or administrative burden to another.

2. Evidence Base and Problem Boundary

A defensible review begins by defining the unit of analysis. For the present topic, the relevant boundary includes inputs, institutional or technical context, the mechanism producing an output, the person or system acting on that output, and the downstream outcome. Evidence falls outside the boundary when it measures only an intermediate proxy yet is used to claim an end-to-end benefit.

Syed et al. (2022) addresses "Zero Trust Architecture (ZTA): A Comprehensive Survey." In this synthesis, that work is used as a source on problem definition within zero-trust service identity in cloud-native infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for multi-tenant cloud and edge service meshes. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Polese et al. (2023) addresses "Understanding O-RAN: Architecture, Interfaces, Algorithms, Security, and Research Challenges." In this synthesis, that work is used as a source on conceptual boundary within zero-trust service identity in cloud-native infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for multi-tenant cloud and edge service meshes. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Porambage et al. (2021) addresses "The Roadmap to 6G Security and Privacy." In this synthesis, that work is used as a source on measurement within zero-trust service identity in cloud-native infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for multi-tenant cloud and edge service meshes. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Together, these works provide complementary entry points, but their conclusions should not be flattened into a single ranking. Differences in data generation, setting, temporal horizon, and outcome definition may be substantively meaningful. A review should therefore preserve those differences and state where analogy, rather than direct replication, connects one domain to another.

3. Mechanism and System Design

The operative mechanism is the chain linking workload attestation, credential issuance, policy evaluation, network enforcement, and logging. This formulation discourages component-only reasoning. A model, platform, policy, assay, or infrastructure service acquires practical meaning only when its interfaces and use conditions are specified. The evidence chain should describe what information is available, what transformation occurs, which assumptions make that transformation credible, and who is authorized to act.

Berardi et al. (2022) addresses "Microservice security: a systematic literature review." In this synthesis, that work is used as a source on system mechanism within zero-trust service identity in cloud-native infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for multi-tenant cloud and edge service meshes. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Rivera et al. (2024) addresses "Securing Digital Identity in the Zero Trust Architecture: A Blockchain Approach to Privacy-Focused Multi-Factor Authentication." In this synthesis, that work is used as a source on representation choice within zero-trust service identity in cloud-native infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for multi-tenant cloud and edge service meshes. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Zanasi et al. (2024) addresses "Flexible zero trust architecture for the cybersecurity of industrial IoT infrastructures." In this synthesis, that work is used as a source on failure analysis within zero-trust service identity in cloud-native infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for multi-tenant cloud and edge service meshes. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

These studies also show why modular claims require interface tests. A component may perform well while the complete pathway fails because inputs drift, users interpret outputs differently, recovery semantics are ambiguous, or institutional incentives change. Design documentation should therefore include not only the intended pathway but also foreseeable deviations, degraded modes, and conditions under which the system should stop or defer.

4. Evaluation and Uncertainty

Evaluation should center on credential compromise, stale policy, confused deputies, revocation latency, and control-plane failure. Average performance remains useful, but it should be accompanied by distributions, error categories, relevant subgroups or operating regimes, and uncertainty at the point where a decision changes. Comparators must isolate the value of the proposed addition rather than compare a mature intervention with an intentionally weak baseline.

Badii et al. (2020) addresses "Smart City IoT Platform Respecting GDPR Privacy and Security Aspects." In this synthesis, that work is used as a source on comparative evaluation within zero-trust service identity in cloud-native infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by

itself establish readiness for multi-tenant cloud and edge service meshes. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Patel (2024) addresses "Zero Trust and DevSecOps in Cloud-Native Environments with Security Frameworks and Best Practices." In this synthesis, that work is used as a source on uncertainty within zero-trust service identity in cloud-native infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for multi-tenant cloud and edge service meshes. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

External validation should introduce meaningful shifts in setting, time, population, workload, market structure, or environmental conditions. A nominally independent sample can still be too similar to test transfer. Conversely, a failed transfer is scientifically informative when the changed conditions are measured and the mechanism of failure is examined rather than hidden in an aggregate score.

5. Translation and Governance

Translation to multi-tenant cloud and edge service meshes depends on more than technical readiness. Decision rights, documentation, maintenance, monitoring, appeal, and recovery are part of the intervention. The governance requirement is short-lived credentials, policy testing, separation of duties, and observable exception handling. Each control should be connected to a named failure mode and should identify an owner, evidence source, review interval, and escalation path.

Ajish (2024) addresses "The significance of artificial intelligence in zero trust technologies: a comprehensive review." In this synthesis, that work is used as a source on implementation within zero-trust service identity in cloud-native infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for multi-tenant cloud and edge service meshes. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Alboqmi and Gamble (2025) addresses "Enhancing Microservice Security Through Vulnerability-Driven Trust in the Service Mesh Architecture." In this synthesis, that work is used as a source on governance within zero-trust service identity in cloud-native infrastructure. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for multi-tenant cloud and edge service meshes. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

A credible implementation plan also defines sunset and revalidation conditions. New data, software updates, institutional restructuring, population change, or shifts in incentives can invalidate previously reasonable assumptions. Monitoring should therefore test the validity of the evidence chain, not merely whether a service remains online or a headline metric remains stable.

6. Research Agenda

Future studies should preregister or otherwise predefine the intended decision, principal outcomes, exclusions, and analysis plan when feasible. They should report missingness, sensitivity analyses, negative or null findings, and the cost of errors to differently situated stakeholders. Reproducible artifacts should include sufficient metadata to reconstruct data provenance, model or analytical versions, parameter choices, and the sequence from evidence to conclusion.

Cross-disciplinary work needs a shared object of explanation rather than a collection of adjacent vocabularies. For zero-trust service identity in cloud-native infrastructure, that shared object is the complete decision pathway. Technical, clinical, social, environmental, or economic expertise should each change the design or interpretation of the study. When evidence remains incomplete, the useful output is a bounded hypothesis, a transparent uncertainty statement, or a request for additional measurement.

7. Conclusion

The literature supports a disciplined approach to zero-trust service identity in cloud-native infrastructure: define the decision, preserve system boundaries, test consequential failure modes, connect uncertainty to action, and assign accountable control. The strongest conclusion is not that one method is universally superior. It is that credible use in multi-tenant cloud and edge service meshes requires an auditable link from source evidence to design choice, evaluation result, and governed decision.

References

Ajish, D. (2024). The significance of artificial intelligence in zero trust technologies: a comprehensive review. Journal of Electrical Systems and Information Technology, 11(1). https://doi.org/10.1186/s43067-024-00155-z

Alboqmi, R., & Gamble, R. F. (2025). Enhancing Microservice Security Through Vulnerability-Driven Trust in the Service Mesh Architecture. Sensors, 25(3), 914. https://doi.org/10.3390/s25030914

Badii, C., Bellini, P., Difino, A., & Nesi, P. (2020). Smart City IoT Platform Respecting GDPR Privacy and Security Aspects. IEEE Access, 8, 23601-23623. https://doi.org/10.1109/access.2020.2968741

Berardi, D., Giallorenzo, S., Mauro, J., Melis, A., Montesi, F., & Prandini, M. (2022). Microservice security: a systematic literature review. PeerJ Computer Science, 7, e779. https://doi.org/10.7717/peerj-cs.779

Patel, D. (2024). Zero Trust and DevSecOps in Cloud-Native Environments with Security Frameworks and Best Practices. International Journal of Advanced Research in Science Communication and Technology, 454-464. https://doi.org/10.48175/ijarsct-11900d

Polese, M., Bonati, L., D'Oro, S., Basagni, S., & Melodia, T. (2023). Understanding O-RAN: Architecture, Interfaces, Algorithms, Security, and Research Challenges. IEEE Communications Surveys & Tutorials, 25(2), 1376-1411. https://doi.org/10.1109/comst.2023.3239220

Porambage, P., Gur, G., Osorio, D. P. M., Liyanage, M., Gurtov, A., & Ylianttila, M. (2021). The Roadmap to 6G Security and Privacy. IEEE Open Journal of the Communications Society, 2, 1094-1122. https://doi.org/10.1109/ojcoms.2021.3078081

Rivera, J. J. D., Muhammad, A., & Song, W. C. (2024). Securing Digital Identity in the Zero Trust Architecture: A Blockchain Approach to Privacy-Focused Multi-Factor Authentication. IEEE Open Journal of the Communications Society, 5, 2792-2814. https://doi.org/10.1109/ojcoms.2024.3391728

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. https://doi.org/10.1109/access.2022.3174679

Zanasi, C., Russo, S., & Colajanni, M. (2024). Flexible zero trust architecture for the cybersecurity of industrial IoT infrastructures. Ad Hoc Networks, 156, 103414. https://doi.org/10.1016/j.adhoc.2024.103414