Complex Systems Informatics and Modeling Quarterly (CSIMQ)
eISSN: 2255-9922
Published online by RTU Press, https://csimq-journals.rtu.lv
Article 271, Issue 48, September/October 2026, Pages 94–119
https://doi.org/10.7250/csimq.2026-48.05

Zero Trust in Heterogeneous Multi-Cloud Environments: A Systematic Literature Review and a Decentralized Federated Trust Framework

Yevhen Darnapuk* and Yevhenii Khudolii

Department of Computer Engineering, Petro Mohyla Black Sea National University,
10 68-Desantnykiv St., 54003, Mykolaiv, Ukraine

yevhen.darnapuk@chmnu.edu.ua, yevhenii.khudoliy@chmnu.edu.ua

Abstract. Implementing Zero Trust across heterogeneous multi-cloud environments is complicated due to incompatible identity models, policy semantics, trust anchors, and credential-lifecycle mechanisms. This article reports a systematic literature review of research published from January 2020 to July 2026. Searches of Scopus, the Web of Science Core Collection, and IEEE Xplore identified 1800 records; 31 distinct studies were retained after deduplication and screening. The synthesis shows that centralized and federated approaches remain operationally mature but continue to depend on identity providers, brokers, and provider-specific policy models. Approaches based on Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs) improve claim portability and local cryptographic verification, yet they do not by themselves resolve issuer governance, credential status, policy mapping, or key compromise. On this basis, the article proposes a decentralized federated trust framework that combines portable credentials, local verification, issuer and credential-lifecycle controls, policy mapping, and resource-local enforcement. The framework is assessed analytically against review-derived requirements and representative architectures; implementation and performance evaluation remain future work.
Keywords: Zero Trust Architecture, Multi-Cloud, Federated Trust, Decentralized Identity, Decentralized Identifiers, Verifiable Credentials.

1 Introduction

Most organizations no longer place all of their critical digital services in a single cloud. The move toward heterogeneous multi-cloud environments is usually driven by practical requirements rather than by architectural preference. Teams distribute workloads across providers to improve resilience, meet regulatory placement requirements, control costs, or gain access to services available only within a specific ecosystem. From an infrastructure perspective, that strategy is often justified. From a security point of view, however, following it may become complicated,

* Corresponding author

© 2026 Yevhen Darnapuk and Yevhenii Khudolii. This is an open-access article licensed under the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0).

Reference: Y. Darnapuk, Y. Khudolii, "Zero Trust in Heterogeneous Multi-Cloud Environments: A Systematic Literature Review and a Decentralized Federated Trust Framework," Complex Systems Informatics and Modeling Quarterly, CSIMQ, no. 48, pp. 94–119, 2026. Available: https://doi.org/10.7250/csimq.2026-48.05

Additional information. Author ORCID iD: Y. Darnapuk – https://orcid.org/0000-0002-7099-5344, and Y. Khudolii – https://orcid.org/0009-0006-0852-9795. PII S225599222600271X. Received: 16 April 2026. Accepted: 19 October 2026. Available online: 31 October 2026.

because identity models, authorization semantics, and policy tooling remain strongly tied to individual vendors [1], [2].

Zero Trust addresses the diminishing usefulness of network location as a basis for security decisions. Cloud-native services, remote administration, application programming interfaces (APIs), mobile access, and machine-to-machine communication require requests to be evaluated using identity, context, policy, device or workload state, and current telemetry rather than inherited perimeter trust [3], [4], [5]. The principle is well established, but its implementation becomes more difficult when the participating services belong to different administrative and cloud domains.

Traditional federation remains useful for single sign-on and delegated authentication. However, an accepted assertion or access token carries a prior trust decision for the lifetime of that artifact, while the identity provider (IdP), federation broker, or authorization server remains a critical coordination point [3], [4], [6]. In heterogeneous multi-cloud environments, this dependency is compounded by incompatible representations of identities, roles, attributes, and access policies. Controls may therefore be mature within each provider but difficult to transfer or interpret consistently across providers [1], [2].

Recent work has explored three related directions: strengthening existing identity and access management (IAM) and policy infrastructures with Zero Trust controls; extending federation across multiple identity providers and computing domains; and using Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs) to support portable claims and local verification [4], [5], [6], [7]. These strands are often examined separately, which obscures their shared limitations concerning trust concentration, revocation, policy interoperability, and key governance. The present review brings them together and uses the resulting evidence to identify the requirements of a cross-cloud trust architecture. It does not assume that decentralized identity is a complete replacement for federation; rather, it examines where credential-based verification addresses weaknesses that remain in centralized and broker-mediated approaches.

The study addresses four research questions (RQs):

RQ1: Which architectural approaches are used to implement Zero Trust across heterogeneous multi-cloud and cross-domain environments?

RQ2: What limitations are reported for centralized federation and provider-specific IAM approaches?

RQ3: How are DIDs, VCs, and related decentralized identity mechanisms used in this context?

RQ4: What architectural requirements follow from the evidence for a more portable federated trust model?

The article makes four contributions. First, it provides a reproducible systematic review of 31 studies concerning Zero Trust, identity federation, IAM interoperability, and decentralized identity in heterogeneous environments. Second, it compares the retained studies by trust coordination, identity mechanism, policy model, governance, revocation, key management, and validation. Third, it derives a set of architectural requirements from the recurring limitations identified in the literature. Fourth, it proposes and analytically evaluates a decentralized federated trust framework based on portable credentials, local verification, and resource-local enforcement.

The article is organized as follows. The main concepts are defined in Section 2. The methodology for the literature review is explained in Section 3. The literature review results are presented in Section 4 and discussed and used for requirements derivation in Section 5. The proposed framework is presented in Section 6. Conclusions and discussion of limitations are available in Section 7.

2 Definitions of Basic Concepts

In this study, Zero Trust Architecture (ZTA) is understood not as a single product or a firm deployment recipe, but as a security paradigm built around the refusal of implicit trust. In this paradigm, access is not granted simply because a subject is located within a network perimeter or has already authenticated. Instead, every request is evaluated in relation to identity, context, policy,

and risk, with the expectation that verification is continuous rather than occasional. This interpretation is consistent with the broader Zero Trust literature, which describes ZTA as a model that replaces inherited trust with repeated, evidence-based decision-making [3], [4], [8]. In other words, Zero Trust is better understood as a shift in security logic than as a narrow technical mechanism.

The term "heterogeneous multi-cloud environment" in this article refers to an infrastructure in which services, data, and workloads are distributed across two or more cloud providers whose technical and administrative models are not natively aligned. Heterogeneity here is not limited to the use of different vendors. It also includes differences in identity and access management semantics, policy languages, monitoring models, resource abstractions, and trust assumptions. In practice, two platforms may both support role-based or attribute-based access control yet differ substantially in how identities are represented, how privileges are inherited, and how authorization decisions are evaluated. This is precisely why multi-cloud security is more sophisticated than simply extending a single-cloud design. Reviews of cloud IAM and multi-cloud security consistently note that interoperability problems often arise not from connectivity itself but from mismatches between provider-specific control models [1], [2]. For this reason, the term should be considered as both an infrastructural and a governance concept. It describes a technical landscape in which uniform enforcement cannot be assumed even when the participating clouds are all individually well secured.

Within such environments, IAM remains the practical core of security enforcement. IAM is generally concerned with authentication, authorization, and the governance of digital access rights across services and resources. In cloud settings, its role expands further because access decisions increasingly involve users, APIs, workloads, devices, and service accounts rather than human operators alone [1], [9]. That broader understanding is important in a Zero Trust context, where identity becomes one of the main anchors for deciding whether a request should proceed at all.

The expression "identity federation" is used here in its conventional sense: a set of mechanisms that allow a subject authenticated in one administrative domain to access resources in another domain through shared trust relationships, token exchanges, or delegated assertions. In practical settings, federation is commonly implemented through identity providers and protocols such as Security Assertion Markup Language (SAML)[†] or OpenID Connect (OIDC)[‡], often combined with OAuth 2.0[§] authorization framework for delegated access. Federation is therefore useful because it reduces duplication of identity processes and lowers the administrative cost of cross-system authentication. At the same time, it usually presumes the existence of one or more entities whose assertions must be accepted by others. The broader notion of federated trust, as used in this article, goes beyond protocol-level interoperability. It refers to the transfer, interpretation, and enforcement of trust decisions across domains, including whether those decisions remain portable when the participating environments lack a centralized identity provider or a single policy plane. This distinction matters because a system may support federated authentication without truly supporting federated trust in an architectural sense. The literature on multi-party authentication and Zero Trust repeatedly shows that cross-domain sign-in is only one part of the larger coordination problem [2], [3], [9].

The phrase "isolated identity domains" describes a condition in which identities, attributes, and access policies remain meaningful largely within the ecosystem that created them. This isolation does not necessarily mean that integration is impossible. More often, this means that interoperability depends on gateways, manual mappings, broker services, or repeated translation between incompatible policy models. As a result, trust is not really portable; it is negotiated anew each time one system attempts to interpret another's claims. In heterogeneous multi-cloud settings, such fragmentation can produce policy drift, privilege inconsistencies, and weak revocation

[†] https://www.cisco.com/site/us/en/learn/topics/security/what-is-saml.html

[‡] https://openid.net/developers/how-connect-works/

[§] https://auth0.com/intro-to-iam/what-is-oauth-2

behavior across administrative boundaries [1], [2], [10]. In this study, the term is used to describe structural fragmentation rather than a complete technical absence of connectivity.

Finally, "decentralized identity" is used in this article to denote identity management approaches that reduce dependence on a permanent centralized intermediary for identifier creation, credential presentation, and trust verification. In this family of approaches, Decentralized Identifiers (DIDs) serve as portable identifiers that can be created and controlled without relying on a single registrar, while Verifiable Credentials (VCs) are cryptographically signed attestations about subjects, attributes, or rights [7], [11], [12]. The underlying idea is not that trust disappears, but that it is expressed and verified differently. Instead of sending every decision back to a central identity provider, a relying party can validate a credential cryptographically, inspect its issuer and proof material, and determine whether the presented claims satisfy policy requirements. This is one reason decentralized identity is increasingly discussed in connection with Zero Trust and multi-domain systems. It promises stronger portability, more selective disclosure of attributes, and less dependence on synchronized IAM silos. Surveys of self-sovereign identity and decentralized identity management, as well as more recent work on zero-trust multi-domain infrastructures, describe DIDs and VCs as key enabling elements for this shift [7], [10], [11], [12]. In this study, decentralized identity is examined as a possible basis for portable cross-domain trust, while its unresolved governance, revocation, policy, and key management requirements are also part of the analysis.

The list of abbreviations and acronyms used in the article is available in Table A1 in Appendix A.

3 Methodology for the Literature Review

The review was conducted as a structured qualitative synthesis of research on Zero Trust in heterogeneous multi-cloud and cross-domain environments, with particular attention to identity federation, IAM interoperability, decentralized identity, policy portability, credential revocation, issuer governance, and key management. Study identification, screening, eligibility assessment, and reporting followed the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 framework [13]. The numerical review corpus included peer-reviewed journal articles, reviews, conference and workshop papers, and peer-reviewed book chapters. Earlier foundational publications, methodological guidance, and normative documents were retained separately to support definitions and framework specification but were not counted as included studies.

The identification process comprised a non-counted query-development stage followed by two formal search rounds, both completed between 22 and 23 July 2026. Round 1 consisted of bibliographic searches in Scopus, the Web of Science Core Collection, and IEEE Xplore and covered publications issued from 1 January 2020 to 23 July 2026. The six finalized query groups were executed once in each database, resulting in 18 database-query executions. Round 2 comprised prior-manuscript bibliographic reconciliation and supplementary backward- and forward-citation searching. Publisher platforms, proceedings archives, institutional repositories, and author-manuscript repositories were used for full-text retrieval and bibliographic verification rather than as additional bibliographic databases.

Before the formal rounds, pilot queries were used to identify relevant terminology and to detect candidate expressions that were either too broad or too restrictive. Pilot results were not exported, screened, or included in the PRISMA counts. The finalized strategy comprised six overlapping thematic query groups:

1. Zero Trust in multi-cloud, hybrid-cloud, cross-cloud, and computing-continuum environments;
2. Zero Trust combined with identity federation, federated authentication, and IAM;
3. Identity, authorization, and trust management in multi-cloud environments;
4. Decentralized identity, DIDs, VCs, and self-sovereign identity (SSI) in Zero Trust or cross-domain systems;
5. Decentralized and cross-domain authorization, policy enforcement, and trust management;
6. Credential revocation, issuer governance, trust registries, and decentralized key management.

Table 1 presents the complete Boolean query bodies Q1-6 used in Round 1. Each query body was executed once in each of the three databases. Scopus applied the TITLE-ABS-KEY field wrapper, the Web of Science Core Collection applied the Topic Search (TS) field wrapper, and IEEE Xplore applied the query to the All Metadata field. The search terms and Boolean operators were unchanged; only the field syntax required by each database differed.

To ensure consistency, explicit inclusion and exclusion criteria were defined before the final full-text review.

The inclusion criteria were as follows:

- Publication in English;
- Publication between 1 January 2020 and 23 July 2026;
- Peer-reviewed journal article, review, conference or workshop paper, or peer-reviewed book chapter;
- Substantive relevance to Zero Trust, heterogeneous multi-cloud trust, IAM interoperability, identity federation, decentralized identity, DIDs, VCs, cross-domain authorization, policy interoperability, credential revocation, issuer governance, trust registries, or key management;
- Sufficient conceptual, architectural, formal, technical, empirical, or systematic-review material for comparative analysis.

Table 1. Thematic Boolean search expressions used to construct database-specific queries

ID	Description	Search query
Q1	Zero Trust in multi-cloud and related distributed environments	(“zero trust” OR “zero-trust” OR “zero trust architecture”) AND (“multi-cloud” OR multicloud OR “multi cloud” OR “hybrid cloud” OR “cross-cloud” OR “cross cloud” OR “cloud federation” OR “federated cloud” OR “multi-domain cloud” OR “computing continuum” OR “cloud-edge continuum” OR “edge-cloud continuum”)
Q2	Zero Trust with federation and IAM	(“zero trust” OR “zero-trust” OR “zero trust architecture”) AND (“identity federation” OR “federated identity” OR “federated authentication” OR “federated trust” OR “identity and access management” OR “cross-domain authentication” OR “cross-domain authorization” OR “cross-domain authorization” OR “policy interoperability” OR “policy portability”) AND (“cloud” OR “cloud computing” OR “multi-cloud” OR multicloud OR “hybrid cloud” OR “computing continuum”)
Q3	Multi-cloud identity and trust management	(“multi-cloud” OR multicloud OR “multi cloud” OR “hybrid cloud” OR “cross-cloud” OR “cross cloud” OR “cloud federation” OR “federated cloud” OR “multi-domain cloud” OR “computing continuum”) AND (“identity and access management” OR “identity federation” OR “federated identity” OR “federated authentication” OR “federated trust” OR “trust management” OR “cross-domain authentication” OR “cross-domain authorization” OR “cross-domain authorization” OR “access control”)
Q4	Decentralized identity in Zero Trust and cross-domain systems	(“decentralized identity” OR “decentralised identity” OR “self-sovereign identity” OR “self sovereign identity” OR “decentralized identifier*” OR “decentralised identifier*” OR “verifiable credential*”) AND (“zero trust” OR “zero-trust” OR “multi-cloud” OR multicloud OR “hybrid cloud” OR “cross-domain” OR “multi-domain” OR “computing continuum” OR “cloud-edge continuum”)
Q5	Decentralized authorization, policy, and trust enforcement	(“decentralized identity” OR “decentralised identity” OR “decentralized identifier*” OR “decentralised identifier*” OR “verifiable credential*”) AND (authentication OR authorization OR authorisation OR “access control” OR “policy enforcement” OR “trust management”) AND (cloud OR “multi-cloud” OR multicloud OR “cross-domain” OR “multi-domain” OR “computing continuum”)
Q6	Revocation, governance, and key management	(“decentralized identity” OR “decentralised identity” OR “decentralized identifier*” OR “decentralised identifier*” OR “verifiable credential*”) AND (revocation OR “credential status” OR “status list” OR “issuer governance” OR “trust registry” OR “trust framework” OR “key management” OR “key rotation” OR “key compromise” OR “issuer compromise”) AND (“zero trust” OR cloud OR “cross-domain” OR “multi-domain” OR federation)

Earlier foundational publications and normative documents were handled outside the numerical corpus. They were retained only where they supplied definitions, established technical primitives, or explained concepts required to interpret the recent studies.

The exclusion criteria were as follows:

- Routine single-cloud administration without a transferable cross-domain contribution;
- Generic cloud security, encryption, anomaly detection, storage, or resource-management research without substantive identity, federation, trust, or policy relevance;
- Domain-specific applications whose mechanisms were not transferable to heterogeneous multi-cloud trust;
- Promotional, editorial, non-technical, or non-peer-reviewed publications;
- Insufficient architectural, methodological, formal, technical, or empirical detail;
- Duplicate or superseded reports of the same study;
- Unavailable full text;
- Failure to satisfy the declared source-selection restriction.

A declared source-selection restriction was applied to publications authored by researchers affiliated with institutions in Russia or Belarus, or issued by publishers based in those countries, in response to Russia's full-scale invasion of Ukraine in February 2022. The same restriction was applied to publications associated with institutions or publishers in Iran, in line with the authors' broader source-selection policy. These exclusions were not based on methodological quality and are acknowledged as a limitation that may have reduced the breadth and generalizability of the synthesis. Institutional affiliation and publisher information were assessed from the publication metadata; nationality was not inferred from authors' names.

This combination of criteria was intended to keep the review focused on architectural and trust-related questions rather than on general cloud security literature.

All database exports were merged into a single record pool. Duplicate records were identified first through normalized Digital Object Identifier (DOI) matching and subsequently through normalized title and publication-year matching where a DOI was absent or inconsistent across databases. Normalization removed differences in capitalization, punctuation, spacing, and database-specific bibliographic formatting. Records were then checked against publication date, language, and document type criteria before title-and-abstract screening.

At the title-and-abstract stage, a publication was retained when it addressed a heterogeneous multi-cloud, hybrid-cloud, cross-domain, multi-domain, cloud-edge, or computing-continuum setting and made a substantive contribution concerning Zero Trust, IAM interoperability, identity federation, decentralized identity, authorization policy, credential lifecycle management, trust governance, or key management. Studies describing mechanisms developed for specialized environments were retained only when their architectural contribution was transferable to cross-domain trust. Generic cloud security, encryption, anomaly detection, storage, resource management, and domain-specific application papers that did not make a substantive contribution to identity or trust were excluded.

Reports passing title-and-abstract screening were sought in full text through publisher platforms, institutional repositories, author-manuscript repositories, and proceedings archives. Full-text reports were assessed against the same eligibility criteria. One primary exclusion reason was recorded for each excluded report. Related conference and journal reports were additionally examined as publication families to prevent multiple reports of the same underlying study from being counted as independent studies.

Round 1 returned 1,800 database records. The exports from the 18 database-query executions were merged into a single record pool. Duplicate removal was performed first by exact and normalized DOI matching, and then by normalized title-and-publication-year matching when a DOI was missing or inconsistent. This procedure removed 866 duplicate records, leaving 934 unique records.

Before title-and-abstract screening, 48 records were removed using objective criteria: three were outside the publication period, 34 were whole books or conference-level records rather than individual studies, six were standards or drafts, and five were not published in English. The remaining 886 records underwent title-and-abstract screening. Of these, 856 were excluded

because they did not satisfy the stated setting and contribution criteria. Round 1 therefore produced 30 reports for attempted full-text retrieval.

Round 2 contributed 15 additional reports. Each was checked against the Round 1 pool to remove duplicates and identify related publication families before the same eligibility criteria were applied. Round 2 therefore produced 15 further reports for attempted full-text retrieval.

The two rounds yielded 45 reports for retrieval: 30 from Round 1 and 15 from Round 2. Seven reports could not be retrieved through publisher, proceedings, repository, or author-manuscript routes, leaving 38 reports for full-text assessment. Six reports were excluded: five did not address a sufficiently relevant heterogeneous multi-cloud or cross-domain setting, and one did not provide a substantive contribution to identity, federation, trust management, or policy interoperability. The qualitative synthesis subsequently included 32 reports describing 31 distinct studies. References [14] and [15] report successive versions of the same multi-cloud identity-management study based on the Solid decentralized-data platform and were therefore assigned a common study-family identifier. Earlier foundational publications, the PRISMA guidance, and normative documents issued by the National Institute of Standards and Technology (NIST) and the World Wide Web Consortium (W3C) were used separately and were not counted in the numerical study corpus.

In Figure 1, the database branch represents Round 1, while the other-methods branch represents Round 2.

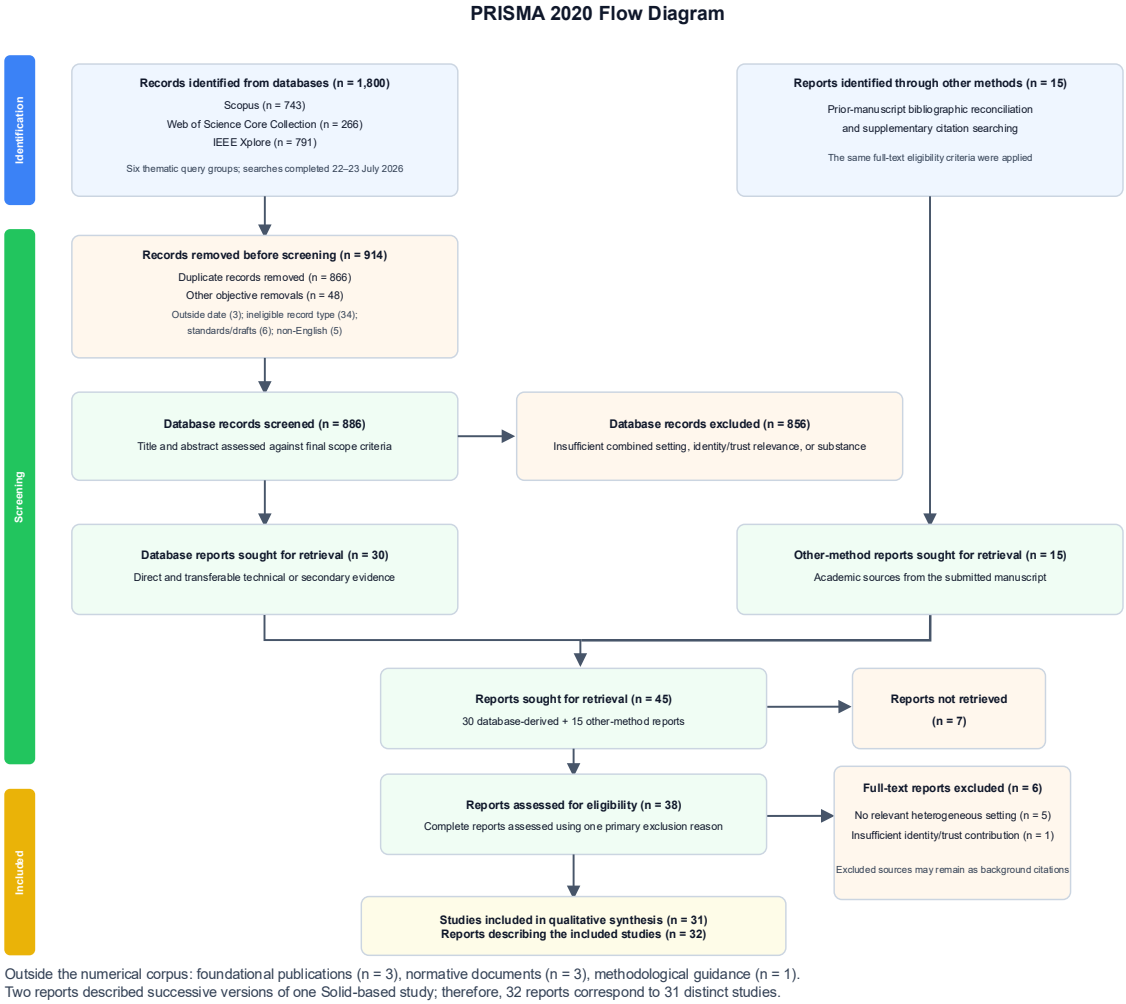

Figure 1. PRISMA 2020 flow diagram for the identification, screening, eligibility assessment, and inclusion of reports and distinct studies.

Each included report was analyzed using a structured extraction form. When multiple reports describe the same underlying study, they were assigned a common study-family identifier and

treated as a single study in the numerical synthesis. The extraction was designed both to summarize individual studies and to support comparisons across shared architectural dimensions.

For each source, the following elements were recorded:

- Publication and study type;
- Year;
- Environment and administrative scope;
- Architectural approach;
- Identity mechanism;
- Trust intermediary or trust anchor;
- Authorization and policy model;
- Issuer-governance mechanism;
- Credential-status or revocation mechanism;
- Key-generation, storage, rotation, and recovery approach;
- Validation type;
- Principal findings;
- Author-reported limitations and, where necessary, a review-derived principal limitation based on the study's scope, proposed mechanism, and validation.

The principal limitation was selected in relation to the research questions rather than as a general criticism of the study. Where a source explicitly reported a limitation relevant to heterogeneous multi-cloud trust, that limitation was retained. Otherwise, a review-derived limitation was identified by comparing the study's stated contribution, architectural scope, implemented or modeled functions, evaluation setting, and validation method. Priority was given to the limitations that most directly constrained transferability across independently governed cloud environments, coverage of the identity and credential lifecycle, policy interoperability, or the strength of the available operational evidence. Review-derived limitations were stated as limits on what the reported evidence establishes, rather than as features that every study should have implemented.

The strength of the available evidence was interpreted from the reported validation method rather than reduced to a single numerical quality score. Studies were distinguished as secondary syntheses, conceptual or architectural proposals, formal models, prototypes, simulations, experimental evaluations, case studies, or operational deployments. This classification was used to prevent conceptual proposals and empirically evaluated systems from being treated as equivalent evidence.

In addition to these general characteristics, special attention was paid to several issues that are central to the present study:

- Dependence on centralized identity providers;
- Persistence of trust throughout a token's lifetime;
- Revocation handling;
- Compatibility of claims and policies across different cloud-specific IAM systems;
- Support for decentralized identity mechanisms;
- Evidence of portability across domains.

The final analysis combines two complementary layers. The first layer was descriptive and focused on the types of studies in the selected corpus: which publication types were most common, which technical settings were discussed most often, and which architectural approaches dominated the literature. The second layer was interpretive and concentrated on recurring weaknesses that appeared across otherwise different studies. These included trust centralization, fragmented IAM semantics, limited portability of access-control logic, weak revocation across administrative boundaries, and the still-emerging role of decentralized identity in addressing such problems.

The descriptive and interpretive analyses provided the basis for the requirements derived in Section 5 and the framework proposed in Section 6.

4 Literature Review Results

4.1 Overview of the Retained Literature

The qualitative synthesis comprised 31 distinct studies described in 32 reports. Twenty-five studies were primary technical, applied, empirical, formal, or architectural investigations; five were secondary reviews or syntheses; and one developed and assessed a Zero Trust maturity framework. The reported validation methods ranged from conceptual and formal analyses to prototypes, performance experiments, case studies, and applied demonstrations. Earlier foundational publications, the PRISMA reporting guidance, and the NIST and W3C documents were used separately and were not counted as included studies.

The evidence base was uneven in both scope and validation. Several studies directly examined multi-cloud or hybrid-cloud identity, Zero Trust workload authentication, policy conflicts, role mapping, multi-cloud access control, and cross-cloud trust management. Other studies investigated transferable mechanisms in cloud-edge federations, computing continua, distributed research infrastructures, telecommunications environments, or service-based architectures. DID- and VC-based studies generally provided stronger evidence for portable identifiers and local verification than for issuer governance, cross-domain credential status, or semantic policy interoperability.

Centralized and federation-based approaches were generally more operationally mature but retained dependencies on identity providers, brokers, authorization servers, or synchronized policy services. Decentralized approaches reduce some of these dependencies but introduce separate requirements for trusted issuers, credential lifecycle management, resolver availability, key protection, and local policy interpretation. Appendix B summarizes the characteristics of the 31 distinct studies included in the qualitative synthesis, with one row for each study. References [14] and [15] are treated as one study family because they describe successive versions of the same Solid-based architecture. The final column notes the limitation that most directly affects each study's contribution to this review and explains how that limitation relates to the study's scope, method, or validation. An omitted mechanism is identified as a limitation only when its absence restricts transferability to heterogeneous multi-cloud trust or limits the evidence relevant to the research questions. Appendix B documents the study-level findings and limitations compared in Sections 4.2–4.6. This comparison provides the basis for the requirements derived in Section 5. The framework proposed in Section 6 brings together mechanisms intended to address those requirements.

4.2 Zero Trust in Cloud and Multi-Cloud Research

The secondary literature establishes a broadly consistent Zero Trust baseline. The retained reviews describe a move from network-location assumptions toward identity-, context-, telemetry-, and policy-based decisions, while also showing that implementation maturity varies substantially across organizations and technical environments [3], [4], [5]. NIST Special Publication (SP) 800-207 [16] provides the normative architectural vocabulary used in this article, but it is not included in the numerical review corpus.

The multi-cloud studies provide more concrete evidence of the differences between general Zero Trust principles and their cross-provider implementation. The performance study in [17] evaluates the overhead introduced by Zero Trust controls in a multi-cloud setting. The workload-authentication framework in [18] applies short-lived and federated identity mechanisms to workloads operating across clouds. The hybrid-cloud IAM framework in [19] examines identity administration and access to multi-tenant resources, while the single sign-on (SSO) and multifactor authentication (MFA) implementation in [20] provides an applied centralized or federated identity baseline.

These studies show that multi-cloud Zero Trust is not a single architectural pattern. Some approaches add stronger authentication and request controls to existing IAM infrastructures; others introduce workload federation, distributed policy components, or independent verification mechanisms. The empirical studies demonstrate that these controls are implementable, but they also expose performance overhead, provider-specific integration requirements, and continued dependence on identity and policy coordination services [17], [18], [19], [20]. The evidence therefore supports Zero Trust as a cross-cloud design objective, but not the existence of a settled, interoperable implementation model.

4.3 Federation, IAM Interoperability, and Centralized Trust Dependencies

Federation remains the most established mechanism for extending authentication beyond a single administrative domain, but the retained implementations rely on different coordination structures. The multi-IdP portal in [6] accepts assertions through a trusted third party and demonstrates practical authentication among independent institutions. The cloud-edge architecture in [21] achieves interoperability through a transparent proxy placed between otherwise incompatible authentication systems. The artificial intelligence/high-performance computing (AI/HPC) deployment in [22] combines federated single sign-on, multifactor authentication, and time-limited roles within a coordinated research infrastructure. The cross-domain Industrial Internet of Things (IIoT) architecture in [23] instead proposes decentralized federation to reduce dependence on a global identity provider, although it does not provide equivalent implementation or performance evidence.

Workload and hybrid-cloud studies extend federation beyond conventional user sign-in. The framework in [18] uses short-lived workload identities and removes long-lived service-account secrets, but permissions remain governed by the participating cloud providers. The mirror-identity system in [19] coordinates on-premises and cloud identities through duplication and synchronization rather than making them portable. The SSO/MFA implementation in [20] strengthens authentication through federated sign-in and risk-sensitive controls but does not evaluate whether authorization semantics remain equivalent across clouds.

Policy-oriented studies show why successful authentication is not sufficient. MC-ABAC [24] supports attribute-based collaboration among cloud tenants, whereas the formal model in [25] detects conflicts among authorization rules from Amazon Web Services, Microsoft Azure, and Google Cloud. The CAPODAZ architecture [26] separates policy decision and enforcement services, and the review in [27] identifies requirements for mapping organizational roles to provider-specific entitlements. These studies address different stages of authorization, but each requires an explicit interpretation of claims, roles, resources, or policies in the relying environment.

Federation therefore provides operational cross-domain authentication but does not, by itself, provide portable authorization. The retained implementations depend on a trusted third party [6], an online proxy [21], institutional identity and role governance [22], provider-managed workload federation [18], or identity synchronization [19]. These mechanisms remain useful compatibility solutions, but their trust dependencies and policy assumptions must be made explicit when they are applied across independently governed clouds.

4.4 Decentralized Identity as an Emerging Architectural Direction

The included decentralized-identity studies move beyond general descriptions of DIDs and VCs and examine how independently verifiable identity evidence can be used across administrative domains. The computing-continuum architecture in [7] and the IAM implementation in [28] combine decentralized identity with heterogeneous cloud-edge environments. The multi-cloud service-management framework in [29] integrates self-sovereign identity into a multi-cloud architecture, while the POLARIS cross-domain access-control architecture [30] connects verifiable identity with policy-based cross-domain authorization.

Several studies evaluate more specialized uses of decentralized identity. The blockchain-enabled identity management (BIdM) system in [31] implements a blockchain-supported cross-domain identity lifecycle. The 6G access-control architecture in [32] and the cross-domain communication protocol in [33] use DIDs and VCs across separate service or network domains. Although these environments are not conventional enterprise multi-cloud deployments, their mechanisms are transferable because they evaluate identifier resolution, credential verification, cross-domain communication, and relying-party enforcement.

Credential-based systems also provide evidence of privacy and repeated authorization. The dynamic-attribute Verifiable Credential model in [34] addresses changes to credential attributes, while the continuous-authorization system in [35] combines VCs with request-time authorization. The two Solid-based architecture reports [14], [15] describe successive versions of one privacy-preserving multi-cloud identity-management study. The IEEE Big Data study in [36] provides a broader synthesis of privacy-preserving IAM models across multiple cloud environments.

The W3C DID and VC documents [37], [38] define the relevant data models and verification primitives but do not prescribe a multi-cloud trust architecture. The included systems demonstrate that portable verification is technically feasible; they do not show that issuer admission, credential status, key compromise, or cloud-specific policy semantics disappear. Decentralized identity therefore changes how trust evidence is carried and verified, but it does not remove the requirement to govern and interpret that evidence.

4.5 Governance, Revocation, Policy Interoperability, and Key Management

The included studies address governance and credential lifecycle management less consistently than authentication and verification. Policy interoperability receives the most direct treatment. MC-ABAC [24] defines attribute-based collaboration across clouds; the formal model in [25] identifies conflicting authorization policies; the role-mapping review in [27] examines entitlement translation; the CAPODAZ architecture [26] separates decision and enforcement services; and the leakage-free ABAC model in [39] addresses the protection and evaluation of cross-cloud authorization rules. These studies show that portable credentials cannot substitute for policy interpretation. A relying domain must still map claims to its own resources, actions, attributes, and authorization semantics.

Credential lifecycle mechanisms are addressed through several technical approaches. The dynamic-attribute credential model in [34] supports changes to credential attributes, the CP-ABE multi-cloud model in [40] includes revocation and key-version handling, and the continuous-authorization system in [35] re-evaluates credential evidence during protected-resource access. These mechanisms support short-lived, status-aware, and dynamically updated trust evidence rather than indefinite reliance on a previously accepted assertion.

Key management receives more focused treatment in [41] and [42]. The distributed organizational-key approach in [41] reduces dependence on one signing key or administrator. Study [42] covers key generation, protected storage, isolation, rotation, delegation, recovery, and scalable management for decentralized identity. These studies provide useful mechanisms for issuer and holder keys but do not by themselves specify a complete multi-cloud issuer-governance model.

The evidence is weakest for issuer admission and suspension. The reviewed architectures generally identify trusted issuers or consortium participants, but fewer studies define which authority admits an issuer, which schemas it may use, how assurance levels are assigned, or how relying domains respond to an issuer-key compromise. This gap directly motivates the explicit governance and compromise-response functions introduced in Section 6.

4.6 Synthesis of the Review Results

The synthesis below is based on the study-level evidence summarized in Appendix B, and brings together the recurring findings that cut across the individual studies. Federation has the strongest

implementation record in the retained literature, although the participating systems rely on different coordination mechanisms. The portal in [6] accepts assertions through a trusted third party, while the cloud-edge architecture in [21] uses an online proxy to connect otherwise incompatible authentication systems. The workload framework in [18] replaces long-lived service-account secrets with provider-managed short-lived identities. The AI/HPC deployment in [22] operates through institutional identity providers and coordinated role governance. The multi-cloud SSO/MFA implementation in [20] combines federated sign-in with multifactor and risk-sensitive authentication but does not extend the evaluation to cross-provider authorization semantics. The IIoT architecture in [23] proposes decentralized federation to reduce dependence on a global identity provider, although it is supported by architectural analysis rather than an implemented evaluation. Taken together, these studies show that cross-domain authentication is feasible, but they do not establish portable authorization across independently governed clouds.

The decentralized-identity studies move more of the verification process to the relying domain. The computing-continuum implementation in [7] evaluates DID and VC verification with selective disclosure. The framework in [29] applies self-sovereign identity to multi-cloud service orchestration, and POLARIS [30] combines verifiable presentations with policy-based authorization. The Solid-based study reported in [14] and [15] focuses on user-controlled identity data and semantic verification. Dynamic credential attributes are treated in [34], blockchain-supported cross-domain authentication in [31], and repeated authorization and revocation in [35]. The telecommunications studies in [32] and [33] provide further evidence from cross-domain access and communication. None of these systems evaluates issuer admission, credential status, key compromise, and cloud-policy mapping as one complete lifecycle.

Policy and lifecycle controls are distributed across a separate group of studies. MC-ABAC [24] addresses the interpretation of attributes across cloud tenants, while [25] detects conflicts after provider policies have been translated into a common formal model. CAPODAZ [26] separates policy decision from enforcement, [27] examines role mapping, and [39] addresses leakage constraints in policy generation. Key custody and recovery receive more direct treatment in [41] and [42]. Credential updating or revocation is evaluated in [34], [40], and [35]. Each study covers part of the required path, but the included implementations do not combine issuer governance, current credential status, policy interpretation, key recovery, and resource-local enforcement in a heterogeneous multi-cloud deployment. The framework proposed in Section 6 combines these functions to address the limitations summarized in Appendix B. This combination is derived from the reviewed mechanisms and limitations; it is not presented as an architecture already validated by the included studies.

5 Discussion and Requirements Derived from the Review

The corpus contains three recurring architectural arrangements. Existing IAM is extended through identity-provider federation, proxies, or provider-managed workload identities in [6], [21], and [18]. DID- and VC-based systems in [7], [29], and [30] move cryptographic verification toward the relying domain. Other studies concentrate on supporting functions, including distributed key management [41]; policy analysis [24], [25]; and separation of policy decision and enforcement [26]. These arrangements are not mutually exclusive; an implementation may combine federation with short-lived credentials, local policy enforcement, or independently verifiable claims.

The first three research questions (RQ1–3) presented in Section 1 lead to the governance, status, policy-mapping, and failure-handling requirements in Table 2.

For RQ1, the most mature implementations remain federation-based. Reference [6] reports an operational multi-IdP portal, [22] describes federation in an AI/HPC research infrastructure, and [21] evaluates a cloud-edge authentication proxy. Enterprise workload federation is implemented in [18], while [20] evaluates SSO, MFA, and risk-sensitive authentication in a multi-cloud setting. The decentralized systems have narrower but more concrete validation. Reference

[7] evaluates DID/VC verification in a computing continuum, [29] measures credential-based service orchestration, and [30] evaluates VC-based policy authorization.

Table 2. Review-derived requirements and their realization in the proposed framework

ID and requirement	Corpus basis	Realization in the proposed framework (Section 6)
R1. Portable identity for users, workloads, and devices	The Solid-based study [14], [15] addresses user-controlled identity across clouds; [28] applies DIDs and VCs to identities in a computing continuum; [18] implements short-lived workload identity across cloud providers.	A user, workload, service, or device presents a DID and VC evidence package that can be carried across administrative domains. Cryptographic and DID-method verification of that evidence is addressed under R2.
R2. Local cryptographic verification of identity evidence	Local DID/VC verification is evaluated in [7]; POLARIS [30] verifies issuer-bound presentations under policy; [35] repeats credential verification during protected-resource access.	The relying domain verifies signatures, proof material, and DID verification methods through V_{sig} and V_{did}, without relying on remote token introspection for every request.
R3. Explicit issuer admission, permitted scope, suspension, and removal	The service-management framework in [29] identifies credential-ecosystem participants, POLARIS [30] evaluates issuer-bound claims, and [43] treats trust governance as a separate multi-cloud concern. None defines the complete issuer-state process required here.	The governance policy records accepted issuer DIDs, credential schemas, permitted scopes, assurance levels, and active, suspended, or revoked states. The request-time test is T_{issuer}.
R4. Credential expiration, status checking, suspension, and revocation	Dynamic credential attributes are evaluated in [34], key-version revocation in [40], and credential revocation during continuous authorization in [35].	S_{cred} checks expiration, suspension, revocation, and the permitted freshness of status information.
R5. Semantic mapping between credential claims and local authorization policies	MC-ABAC [24] requires attributes to be interpreted across cloud tenants; [25] detects conflicts only after provider policies have been translated; [27] derives requirements for mapping organizational roles to cloud entitlements.	M_{claims} maps verified claims to local attributes, roles, resources, and actions. Missing or conflicting mappings produce denial.
R6. Protected key storage, rotation, delegation, compromise response, and recovery	Distributed organizational signing is evaluated in [41]; [42] covers protected storage, isolation, rotation, delegation, and recovery for decentralized-identity keys.	Issuer and holder keys are versioned assets subject to protected storage, rotation, emergency suspension, recovery, and invalidation of affected credentials or presentations.
R7. Context-sensitive authorization and enforcement close to the protected resource	POLARIS [30] evaluates policy-based authorization, CAPODAZ [26] separates decision from enforcement, and [35] performs request-level re-evaluation.	P_{local} evaluates local authorization and C_{context} evaluates resource, action, posture, risk, and other request conditions. The PEP enforces the result.
R8. Defined behavior when resolution, registry, or credential-status information is unavailable or stale	The proxy in [21], DID resolution in [33], and current credential evidence in [35] expose availability and freshness dependencies, but do not provide one common failure rule.	High-risk requests fail closed. Lower-risk requests may use previously verified information only within an explicit freshness interval. Conflicting registry information suspends the affected trust entry.
R9. Auditability of decisions, mappings, credential status, and policy versions	Policy-conflict analysis in [25], separated decision and enforcement in [26], and policy generation in [39] depend on identifiable policy inputs and outcomes. Explicit decision-record requirements are derived from this need for traceability.	Each decision records the issuer, credential-status reference, verification-method version, mapping and policy versions, context, outcome, and decision time.

For RQ2, the difficulty is not an inability to authenticate across domains. In [6] and [21], the trusted third party or proxy remains on the request path. Reference [18] removes long-lived workload secrets but leaves permissions under provider control. The hybrid-cloud system in [19] depends on identity duplication and synchronization. At the authorization layer, MC-ABAC [24] requires attributes to be interpreted across tenants, the conflict analysis in [25] assumes that

provider policies have first been translated, and [27] identifies unresolved role-mapping requirements. Authentication may therefore succeed while authorization remains local or semantically inconsistent.

For RQ3, DID- and VC-based systems show that claims can be presented and checked without consulting one central IdP for every request. This behavior is demonstrated directly in [7], [29], and [30]. Reference [34] supports changes to credential attributes, while [35] re-evaluates credential evidence during authorization. These mechanisms do not determine whether an issuer is acceptable, whether status information is sufficiently current, or how a verified claim maps to a local permission.

RQ4 concerns the functions that a more portable federated trust model must contain. A requirement was retained when an included study demonstrated the relevant function or when a recurring omission limited cross-domain transferability, lifecycle coverage, policy interpretation, or local enforcement. Table 2 distinguishes the corpus basis of each requirement from its realization in the proposed framework. “Framework realization” denotes architectural coverage; it does not imply that the requirement has already been implemented or experimentally achieved.

The requirements enter the framework in different ways. R1–R5 and R7 are represented directly in the request and decision path. R6 supports signature and issuer verification through the key lifecycle rather than through a separate Boolean condition. R8 governs the treatment of unavailable, stale, or conflicting evidence, and R9 governs the decision record. This mapping permits an assessment of architectural coverage, but not of latency, scalability, or operational security, which are the properties requiring implementation and controlled experiments under representative multi-cloud workloads.

6 Reference Framework for Decentralized Federated Trust

The federated systems in the corpus show why successful cross-domain authentication is not equivalent to portable trust. The portal in [6] relies on a trusted third party, the cloud-edge design in [21] relies on an online proxy, and the workload framework in [18] retains provider-managed federation. Decentralized systems move more of the verification process to the relying domain: [7] evaluates local DID/VC verification, [29] applies SSI to multi-cloud service management, POLARIS [30] combines verifiable presentations with policy authorization, and [35] rechecks credential evidence during resource access.

The proposed framework brings these mechanisms together with explicit issuer governance, credential-status evaluation, claim mapping, key controls, and defined failure behavior. It is a review-derived reference architecture rather than an implemented system. NIST SP 800-207 [16] supplies the decision and enforcement terminology, while the W3C specifications [37], [38] define the VC and DID primitives.

The framework changes the principal artifact through which cross-domain trust is conveyed. Instead of relying only on a cloud-specific account, role, or broker-issued token, it uses DIDs and VCs to carry verifiable claims between administrative domains. The relying environment verifies the evidence and applies its own policy before granting access. A central identity provider is therefore not required to mediate every request, although credential issuers, resolvers, registries, and status services still form part of the trust environment.

Figure 2 presents the principal components and request path. The Requesting Subject may be a user, workload, service, or device, since Zero Trust decisions in cloud systems apply to machine identities as well as human users [5], [16], [28]. The Policy Enforcement Point (PEP) intercepts the request and obtains a decision from the Trust Engine before allowing access to the protected API or data resource. The decentralized registry in Figure 2 is represented as a distributed ledger technology (DLT) layer that stores or anchors the public material required for DID resolution and credential-schema verification.

The labels “DID document + VC presentation” and “Verifiable Credential (VC)” in Figure 2 are compact descriptions of the exchanged identity evidence. Operationally, the subject submits

an access request and a verifiable presentation. The DID Resolver obtains the relevant DID documents and verification methods, rather than accepting them as trusted documents supplied by the subject. The VC Verifier then checks the credential structure, issuer signature, and proof material according to the applicable W3C specifications [37], [38].

The implementations in [7] and [30] show that selective disclosure can be combined with credential verification and policy evaluation. In the proposed framework, the Trust Engine acts as a policy decision point (PDP) and may accept zero-knowledge proofs (ZKPs) when a claim must be demonstrated without revealing its complete value. The framework applies the conjunctive decision rule (Expression 1) to a request r:

$$\text{Permit}(r) = V_{\text{sig}} \wedge V_{\text{did}} \wedge T_{\text{issuer}} \wedge S_{\text{cred}} \wedge M_{\text{claims}} \wedge P_{\text{local}} \wedge C_{\text{context}}. \quad (1)$$

V_{sig} verifies the presentation and credential signatures, and V_{did} confirms that the verification methods were obtained from acceptable DID documents. T_{issuer} establishes whether the issuer is trusted for the credential type, scope, and assurance level. S_{cred} checks expiration, suspension, revocation, and status freshness. M_{claims} confirms that the claims map unambiguously to the relying domain's semantic model. P_{local} evaluates the mapped claims under local policy, while C_{context} evaluates the resource, action, workload or device state, risk level, and other request conditions. If any mandatory condition is false or cannot be established from current and unambiguous evidence, the request is denied.

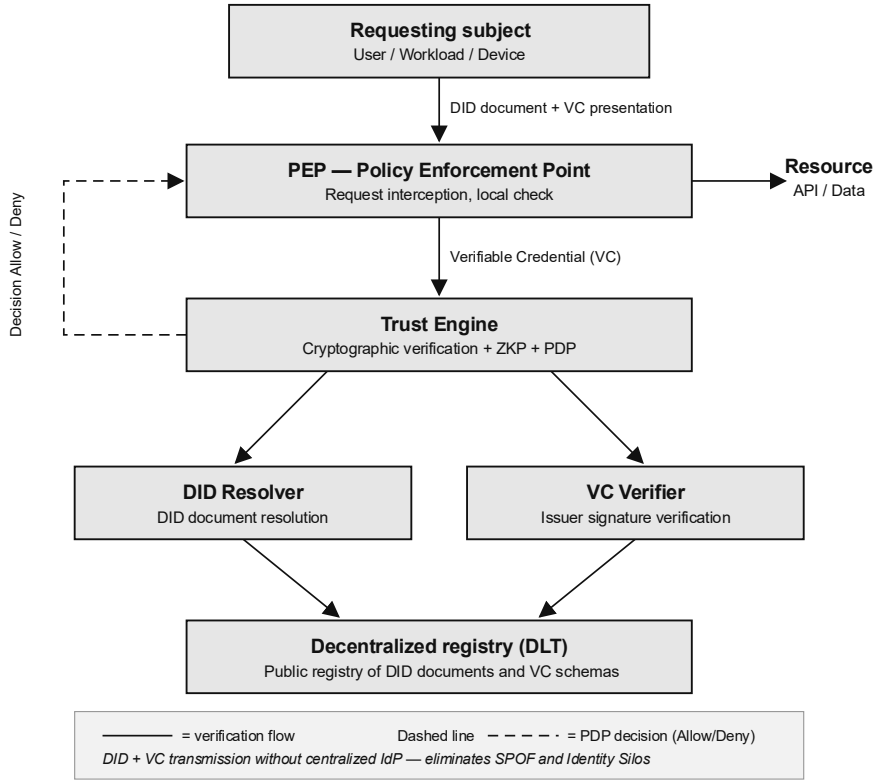

Figure 2. Reference framework for decentralized federated trust based on DIDs and VCs

The relationship between the decision rule (Expression 1) and Table 2 is as follows. R1 is realized by the DID and VC evidence package presented by the subject and accepted across administrative domains. R2 is implemented through V_{sig} and V_{did}; R3 through T_{issuer}; R4 through S_{cred}; R5 through M_{claims}; and R7 through P_{local}, C_{context}, and enforcement by the PEP. R6 is not represented by a separate request-time predicate because it governs the issuer and holder keys on which V_{sig} and T_{issuer} depend. R8 determines how V_{did}, T_{issuer}, and S_{cred} are handled when

resolution or status information is unavailable, inconsistent, or stale. R9 determines which evidence, mappings, policy versions, and outcomes are written to the decision record.

The studies cited in Table 2 address different parts of the credential lifecycle. The service-management framework in [29] identifies participants in a multi-party credential ecosystem, while POLARIS [30] evaluates issuer-bound claims. Dynamic credential updates are examined in [34], and the implementation in [35] checks revocation during repeated authorization. Distributed signing-key custody is evaluated in [41], whereas [42] covers protected storage, rotation, delegation, and recovery. The review in [43] treats governance as a separate multi-cloud coordination problem. None of these studies defines a cross-cloud procedure for issuer admission, scope restriction, suspension, revocation, and recovery after key compromise.

In the proposed framework, an issuer is admitted by a governance authority accepted by the relying domain. Publication of a DID does not make the issuer trusted. The admission record identifies the issuer DID, permitted credential schemas, authorized claim scope, assurance level, validity period, current status, and the policy version under which admission was granted. Trust management is rule-based rather than reputation-based: an issuer is trusted only for the schemas, claims, assurance levels, and validity period recorded in that policy. The issuer states are active, suspended, and revoked. Suspension blocks new authorization decisions while an incident, policy violation, or key change is being investigated. Revocation removes the issuer from the accepted trust set.

Credential status is evaluated under S_{cred}. The check covers expiration, suspension, revocation, the issuer-key version, the holder-key version, and the permitted age of the status information. The credential status reference may identify a signed status list or another status service accepted by local policy. A revoked credential is rejected. A suspended credential is rejected until its status is restored. If the status source cannot be reached, the request is processed under the R8 freshness rule rather than being accepted solely because its signature remains valid.

Holder-key compromise invalidates presentations made with the affected key version. Issuer-key compromise suspends the issuer, prevents new credentials from being accepted under the compromised key, and triggers rotation and reissuance or invalidation of affected credentials. Historical key material may be retained for auditing earlier decisions, but it cannot authorize new requests after its suspension or revocation time.

Credential portability does not make provider permissions semantically equivalent. MC-ABAC [24] requires attributes to be interpreted across cloud tenants. The formal model in [25] detects conflicts only after provider policies have been translated into a common representation, while [27] identifies a similar problem for organizational roles and cloud-specific entitlements. CAPODAZ [26] shows how policy decision-making can be separated from enforcement. These studies support explicit mapping and local enforcement, but do not prescribe one common authorization language.

In the proposed framework, verified claims are mapped to the relying domain's attributes, roles, resources, and actions before policy evaluation. Each mapping rule identifies the accepted credential schema and issuer scope, the source claim and expected value, the corresponding local attribute or role, the permitted resource and action, the required contextual conditions, and the policy version. Missing, expired, or conflicting mappings result in denial.

For instance, a relying domain may define that an active issuer authorized for the *WorkloadOperator* schema can map the claim *project = alpha* to the local role deployer for resources belonging to project alpha. The mapping is valid only when credential status is current, the workload posture is compliant, and the risk score is below the locally defined threshold. The same credential may map to a different role or to no role in another cloud because final authorization remains under the relying domain's policy.

Several reviewed systems depend on information that must remain available during authorization. The proxy in [21] remains on the authentication path, the communication protocol in [33] depends on DID resolution, and the system in [35] requires current credential evidence for

repeated authorization. These dependencies make unavailable or stale trust information an architectural issue rather than only an implementation fault.

The proposed framework handles such failures according to the sensitivity of the requested operation. High-risk requests are denied when mandatory resolution or status information is unavailable, stale, or inconsistent. Conflicting registry responses or a failed integrity check suspend the affected registry or DID-method trust entry until a trusted resolution path is restored. A lower-risk policy may use previously verified information within a defined freshness interval, provided that its version and age are recorded. The decision record identifies the issuer, credential-status reference, verification-method version, policy-mapping version, relevant context, outcome, and decision time. A reference or digest may be retained instead of the complete credential presentation where that is sufficient for later audit.

A successful cross-cloud workload request illustrates the complete path. A workload in one cloud presents a *WorkloadOperator* credential to a resource in another cloud. The signature and DID verification method satisfy V_{sig} and V_{did}; the issuer is active and authorized for the schema and project scope under T_{issuer}; the credential is current under S_{cred}; and the project claim maps to the local deployment role under M_{claims}. If local policy and contextual checks also succeed, the PEP permits the deployment request and records the evidence and policy versions used in the decision.

Revocation or compromise produces a different result. If the status service marks the credential as revoked, S_{cred} is false and the request is denied. If the resolver returns conflicting verification methods, V_{did} cannot be established. If the governance service reports an issuer-key compromise, T_{issuer} fails and the issuer is suspended. High-risk requests remain denied until a trusted resolution path is restored, the issuer key is rotated, and the affected credentials have been invalidated or reissued.

Once these checks are complete, the Trust Engine returns a decision to the PEP, which enforces it against the protected resource. CAPODAZ [26] provides an example of separating decision from enforcement. MC-ABAC [24], the conflict-analysis model in [25], and the role-mapping review in [27] retain local interpretation of provider permissions. The mirror-identity approach in [19] illustrates the alternative cost of reproducing and synchronizing identities across domains.

The proposed framework therefore keeps the final authorization decision and enforcement within the relying environment. DIDs and VCs reduce the need for participating domains to share one native identity system, but they do not remove differences among AWS IAM, Azure IAM, Google Cloud IAM, or application-specific policy models.

Table 3 and the condition mapping above show how R1–R9 are represented in the framework. The comparison uses six studies selected to represent the principal coordination models found in the selected literature review corpus: multi-IdP federation, proxy federation, workload federation, decentralized computing-continuum identity, multi-cloud SSI, and VC-based policy authorization. Each selected study provides enough architectural and validation details for a source-specific comparison. The selection is intended to cover distinct architectural patterns rather than to rank every included study.

The criteria are derived directly from R1–R9. Identity evidence and decision path cover R1, R2, and R7 and describe what identity artifact is presented, where it is verified, and where authorization is enforced. Governance and lifecycle cover R3, R4, and R6 and describe issuer admission, credential status and revocation, and key protection and recovery. Policy, failure behavior, and auditability cover R5, R8, and R9 and describe claim mapping, behavior when trust information is unavailable or stale, and the information retained about a decision. Validation is reported with the approach name because it indicates evidence maturity rather than architectural coverage.

The comparison shows two distinct trade-offs. The portal in [6], the proxy architecture in [21], and the workload system in [18] provide stronger implementation evidence and integrate with established IAM infrastructures, but their trust paths remain tied to a third party, proxy, or provider-managed federation. The computing-continuum system in [7], the multi-cloud SSI framework in [29], and POLARIS [30] support portable evidence and local cryptographic

verification, but none makes issuer governance, credential status, semantic policy mapping, key-compromise recovery, failure behavior, and auditability equally explicit.

Table 3. Requirement-based comparison of representative architectural approaches

Approach and validation basis	Identity evidence and decision path (R1, R2, R7)	Governance and lifecycle (R3, R4, R6)	Policy, failure behavior, and auditability (R5, R8, R9)
Multi-IdP federation [6] – operational portal implementation	Institutional IdP assertions are accepted through a trusted third party. Authorization by the resource domain is outside the evaluated federation path.	Issuer trust and token handling remain within the institutional federation relationships. No portable issuer-state, credential-status, or cross-domain key-recovery process is defined.	Authorization remains provider-local. Loss of the trusted third-party interrupts authentication. Cross-cloud mapping, stale-information handling, and a common decision record are not evaluated.
Proxy federation [21] – prototype and comparative delay evaluation	A proxy translates authentication between cloud and edge domains. The target domain still performs its own authorization.	Trust and credential handling are administered through the proxy and participating providers. Portable credential status and key recovery are not part of the design.	The proxy is an online dependency, and provider policies remain local. The study does not define semantic mapping, bounded use of stale evidence, or a shared audit format.
Workload federation [18] – enterprise implementation	Short-lived OIDC workload identities replace long-lived service-account secrets. The receiving provider verifies the external workload identity and applies its own permissions.	Issuer configuration and token lifetime remain provider-managed. The study does not cover user or device credentials or an issuer-suspension and recovery process.	Permissions remain cloud-specific. Cross-provider claim mapping, registry or status-service failure, and a common decision-record structure are not addressed.
Computing-continuum identity [7] – prototype and performance experiments	DIDs and VCs support local verification, selective disclosure, and decentralized access control.	Credential issuance and blockchain anchoring are present, but issuer admission, status freshness, and compromise recovery are not evaluated as one lifecycle.	Policy enforcement is demonstrated, but cloud-specific semantic mapping, resolver or status-service failure, and audit-record requirements remain incomplete.
Multi-cloud SSI [29] – prototype scenario and latency measurements	SSI credentials support cross-cloud service management evidence and relying-party verification.	The framework defines ecosystem participants and blockchain-supported trust, but revocation and compromise recovery remain partially specified.	Credential-based orchestration is evaluated, but general mapping to cloud-native IAM, defined failure behavior, and decision-record requirements remain incomplete.
POLARIS [30] – working prototype and experiments	VCs, selective disclosure, and policy-based authorization are evaluated together.	Issuer-bound claims are supported, but credential status, issuer suspension, and key recovery are not defined as a complete lifecycle.	Policy evaluation is implemented, but mapping to several cloud-native IAM models, resolver or status failure, and a common audit record are not demonstrated.
Proposed framework – requirements-based analytical evaluation	DID and VC evidence is evaluated through V_{sig}, V_{did}, T_{issuer}, S_{cred}, M_{claims}, P_{local}, and C_{context}. Final enforcement remains with the PEP in the relying domain.	The framework defines issuer admission and scope; active, suspended, and revoked states; credential-status freshness; versioned issuer and holder keys; rotation; invalidation; and recovery.	Mapping rules are versioned, and ambiguous mappings are denied. High-risk requests fail closed; lower-risk use of earlier evidence is bound by freshness rules. Each decision records evidence, mappings, policy versions, context, outcome, and time.

The proposed framework shares portable credential verification with [7], [29], and [30], and retains local authorization and enforcement in the same way that the federated systems retain provider control. Its distinguishing feature is not a new credential format or cryptographic primitive. It places issuer state, credential status, claim mapping, context, key lifecycle, failure handling, and decision recording in one traceable request path. This is broader architectural coverage, not proof of operational superiority. Unlike the six compared systems, the proposed framework has not been implemented or evaluated under a representative multi-cloud workload.

7 Conclusions and Limitations

This review included 31 distinct studies described in 32 reports and examined how Zero Trust, identity federation, IAM interoperability, and decentralized identity have been applied across multi-cloud and cross-domain environments. Centralized and federated mechanisms remain the most operationally mature, but their use across cloud boundaries commonly depends on identity providers, proxies, provider-managed federation, and cloud-specific policy semantics. DID- and VC-based approaches make identity evidence more portable and support local cryptographic verification, but they do not remove the need to govern issuers, determine credential status, map claims to local policies, and protect signing and holder keys. The principal unresolved problem is therefore not authentication alone, but the consistent interpretation and enforcement of trust evidence across independent domains.

The proposed framework maps derived requirements R1–R9 (Section 5) to the request path, supporting controls, and the decision rule in Expression 1 (Section 6). A request is permitted only when the credential and presentation signatures are valid, the DID verification method is acceptable, the issuer is trusted for the relevant credential type and scope, the credential is current, its claims map unambiguously to the relying domain, local policy permits the requested action, and the current context is acceptable. Key-lifecycle controls support signature and issuer verification, while the failure rule governs unavailable, inconsistent, or stale evidence. A valid credential is therefore necessary but not sufficient for access.

The framework is intended for cloud-security architects and IAM engineers designing cross-cloud access paths, identity-governance teams responsible for issuer and credential lifecycles, platform operators integrating local enforcement, and auditors who need to reconstruct authorization decisions. It provides these stakeholders with a design and assessment baseline rather than deployable software. The framework identifies where trust decisions are made, which conditions require denial, which governance records must be maintained, and which evidence should remain available for audit. It also provides researchers with a common structure for implementing and comparing heterogeneous multi-cloud trust systems.

Its expected practical value follows from the combination of portable identity evidence, local verification, issuer-state management, current credential-status checks, explicit claim mapping, protected key lifecycles, resource-local enforcement, and versioned decision records. Under a deployment that provides these functions, the framework may reduce dependence on an online federation intermediary, shorten the period during which revoked or compromised credentials remain usable, and improve the reproducibility of authorization decisions across providers. These are testable expectations rather than effects demonstrated by the present study.

Several limitations qualify these conclusions. The evidence base combines reviews, conceptual and formal models, conference studies, prototypes, and applied implementations, which provide different levels of empirical support. Many technical evaluations were conducted in specialized settings, including computing continua, telecommunications infrastructures, research-computing platforms, and particular cloud services; their findings cannot be transferred directly to every multi-cloud deployment. The restriction concerning affiliations and publishers in Russia, Belarus, and Iran also reduced the breadth of the corpus and may affect the generalizability of the synthesis. In addition, the framework has not been implemented, experimentally tested, formally verified, or assessed by independent experts. Its latency, scalability, policy-maintenance cost, and behavior

during resolver failure, registry compromise, credential theft, and issuer-key compromise remain to be measured.

Evaluation should begin with a prototype spanning at least two cloud providers and should include both a workload credential and a human or device credential. The minimum measures should include median and tail decision latency, throughput, the proportion of requests requiring an online identity-provider or proxy call, the interval between publication of a revocation and enforced denial, and the consistency of claim-to-policy mappings across providers. A predefined authorization test set should also be used to measure false-permit and false-deny rates. Recovery tests should record the time required to suspend an affected issuer, rotate keys, invalidate credentials, and restore trusted operation after issuer- or registry-key compromise. Auditability can be measured as the proportion of decisions containing the required issuer, status, verification-method, mapping, policy-version, context, outcome, and timestamp fields.

These measurements should be collected during normal operation and under status-service outages, stale status information, conflicting registry responses, credential theft, and issuer-key compromise. Expert assessments by cloud-security architects, IAM engineers, identity-governance specialists, and auditors should examine whether the governance records, mapping rules, failure behavior, and decision evidence are sufficiently clear to implement and operate.

References

[1] I. Indu, P. M. R. Anand, and V. Bhaskar, "Identity and access management in cloud environment: Mechanisms and challenges," *Engineering Science and Technology, an International Journal*, vol. 21, no. 4, pp. 574–588, 2018. Available: https://doi.org/10.1016/j.jestch.2018.05.010

[2] S. Ali et al., "Security and privacy in multi-cloud and hybrid cloud environments: Challenges, strategies, and future directions," *Computers & Security*, vol. 157, article 104599, 2025. Available: https://doi.org/10.1016/j.cose.2025.104599

[3] C. Buck, C. Olenberger, A. Schweizer, F. Völter, and T. Eymann, "Never trust, always verify: A multivocal literature review on current knowledge and research gaps of zero-trust," *Computers & Security*, vol. 110, article 102436, 2021. Available: https://doi.org/10.1016/j.cose.2021.102436

[4] W. Yeoh, M. Liu, M. Shore, and F. Jiang, "Zero trust cybersecurity: Critical success factors and A maturity assessment framework," *Computers & Security*, vol. 133, article 103412, 2023. Available: https://doi.org/10.1016/j.cose.2023.103412

[5] S. Sarkar, G. Choudhary, S. K. Shandilya, A. Hussain, and H. Kim, "Security of Zero Trust Networks in Cloud Computing: A Comparative Review," *Sustainability*, vol. 14, no. 18, article 11213, 2022. Available: https://doi.org/10.3390/su141811213

[6] J. R. Almeida, A. Zúquete, A. Pazos, and J. L. Oliveira, "A federated authentication schema among multiple identity providers," *Heliyon*, vol. 10, no. 7, article e28560, 2024. Available: https://doi.org/10.1016/j.heliyon.2024.e28560

[7] J. M. Bernabé Murcia, E. Cánovas, J. García-Rodríguez, A. M. Zarca, and A. Skarmeta, "Decentralised Identity Management solution for zero-trust multi-domain Computing Continuum frameworks," *Future Generation Computer Systems*, vol. 162, article 107479, 2025. Available: https://doi.org/10.1016/j.future.2024.08.003

[8] M. A. Azad, S. Abdullah, J. Arshad, H. Lallie, and Y. H. Ahmed, "Verify and trust: A multidimensional survey of zero-trust security in the age of IoT," *Internet of Things*, vol. 27, article 101227, 2024. Available: https://doi.org/10.1016/j.iot.2024.101227

[9] S. Aboukadri, A. Ouaddah, and A. Mezrioui, "Machine learning in identity and access management systems: Survey and deep dive," *Computers & Security*, vol. 139, article 103729, 2024. Available: https://doi.org/10.1016/j.cose.2024.103729

[10] A. Mühle, A. Grüner, T. Gayvoronskaya, and C. Meinel, "A survey on essential components of a self-sovereign identity," *Computer Science Review*, vol. 30, pp. 80–86, 2018. Available: https://doi.org/10.1016/j.cosrev.2018.10.002

[11] J. Sedlmeir, R. Smethurst, A. Rieger, and G. Fridgen, "Digital Identities and Verifiable Credentials," *Business & Information Systems Engineering*, vol. 63, pp. 603–613, 2021. Available: https://doi.org/10.1007/s12599-021-00722-y

[12] J. Xian, L. You, Q. Yi, J. Wang, and G. Hu, “A survey on decentralized identity management systems,” *Computer Science Review*, vol. 58, article 100811, 2025. Available: https://doi.org/10.1016/j.cosrev.2025.100811

[13] M. J. Page et al., “The PRISMA 2020 statement: an updated guideline for reporting systematic reviews,” *BMJ*, vol. 372, article n71, 2021. Available: https://doi.org/10.1136/bmj.n71

[14] A. Cuzzocrea and I. Belmerabet, “Towards Privacy-Preserving Multi-Cloud Identity Management Using SOLID,” in *Proceedings of the 21st International Conference on Security and Cryptography, SCITEPRESS – Science and Technology Publications*, 2024, pp. 649–654. Available: https://doi.org/10.5220/0012719900003767

[15] A. Cuzzocrea and I. Belmerabet, “A SOLID-Based Framework for Supporting Privacy-Preserving Identity Management in Multi-Cloud Environments,” in *2024 IEEE 48th Annual Computers, Software, and Applications Conference (COMPSAC)*, IEEE, 2024, pp. 2334–2340. Available: https://doi.org/10.1109/COMPSAC61105.2024.00375

[16] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, “Zero Trust Architecture,” *NIST Special Publication 800–207*, 2020. Available: https://doi.org/10.6028/NIST.SP.800-207

[17] S. Rodigari, D. O’Shea, P. McCarthy, M. McCarry, and S. McSweeney, “Performance Analysis of Zero-Trust multi-cloud,” in *2021 IEEE 14th International Conference on Cloud Computing (CLOUD)*, IEEE, 2021, pp. 730–732. Available: https://doi.org/10.1109/CLOUD53861.2021.00097

[18] S. Deochake, R. Murphy, and J. Gearheart, “A Multi-Cloud Framework for Zero Trust Workload Authentication,” in *2025 Annual Computer Security Applications Conference Workshops (ACSAC Workshops)*, IEEE, 2025, pp. 523–530. Available: https://doi.org/10.1109/ACSACW69556.2025.00061

[19] S. Deochake and V. Channapattan, “Identity and Access Management Framework for Multi-tenant Resources in Hybrid Cloud Computing,” in *Proceedings of the 17th International Conference on Availability, Reliability and Security*, ACM, 2022, pp. 1–8. Available: https://doi.org/10.1145/3538969.3544896

[20] M. I. Hussain et al., “AAAA: SSO and MFA Implementation in Multi-Cloud to Mitigate Rising Threats and Concerns Related to User Metadata,” *Applied Sciences*, vol. 11, no. 7, article 3012, 2021. Available: https://doi.org/10.3390/app11073012

[21] Y.-D. Lin, D.-T. Truong, A. Ali, C.-Y. Li, Y.-C. Lai, and T.-M. T. Dinh, “Proxy-Based Federated Authentication: A Transparent Third-Party Solution for Cloud-Edge Federation,” *IEEE Network*, vol. 34, no. 6, pp. 220–227, 2020. Available: https://doi.org/10.1109/MNET.011.2000136

[22] S. R. Alam et al., “Federated Single Sign-On and Zero Trust Co-design for AI and HPC Digital Research Infrastructures,” in *SC24-W: Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis*, IEEE, 2024, pp. 1756–1764. Available: https://doi.org/10.1109/SCW63240.2024.00220

[23] E. Kim et al., “Can We Create a Cross-Domain Federated Identity for the Industrial Internet of Things without Google?” *IEEE Internet of Things Magazine*, vol. 3, no. 4, pp. 82–87, 2020. Available: https://doi.org/10.1109/IOTM.0001.2000050

[24] M. A. Madani, A. Kerkri, and M. Aissaoui, “MC-ABAC: An ABAC-based Model for Collaboration in Multi-Cloud Environment,” *International Journal of Advanced Computer Science and Applications*, vol. 14, no. 6, 2023. Available: https://doi.org/10.14569/IJACSA.2023.01406126

[25] E. Zahoor, A. Ikram, S. Akhtar, and O. Perrin, “A Formal Approach for the Identification of Authorization Policy Conflicts within Multi-Cloud Environments,” *Journal of Grid Computing*, vol. 20, no. 2, article 18, 2022. Available: https://doi.org/10.1007/s10723-022-09606-1

[26] D. Kallergis, Z. Garofalaki, G. Katsikogiannis, and C. Douligeris, “CAPODAZ: A containerised authorisation and policy-driven architecture using microservices,” *Ad Hoc Networks*, vol. 104, article 102153, 2020. Available: https://doi.org/10.1016/j.adhoc.2020.102153

[27] M. Niedermeier and H. Wittges, “Towards Role Mappings in Hybrid Cloud Environments: A Systematic Literature Review,” in *Enterprise Design, Operations, and Computing, Lecture Notes in Computer Science*, vol. 15409, 2025, pp. 119–139. Available: https://doi.org/10.1007/978-3-031-78338-8_7

[28] C. D. N. Kyriakidou et al., “Identity and Access Management for the Computing Continuum,” in *Proceedings of the 2nd International Workshop on MetaOS for the Cloud-Edge-IoT Continuum*, ACM, 2025, pp. 33–39. Available: https://doi.org/10.1145/3721889.3721926

[29] E. Zeydan, J. Baranda, J. Mangues-Bafalluy, S. S. Arslan, and Y. Turk, “A Trustworthy Framework for Multi-Cloud Service Management: Self-Sovereign Identity Integration,” *IEEE Transactions on Network Science and Engineering*, vol. 11, no. 3, pp. 3135–3147, 2024. Available: https://doi.org/10.1109/TNSE.2024.3361809

[30] A. Zhang, X. Lee, Z. Zhuang, J. Wei, Y. Fu, and B. Peng, “POLARIS: Cross-Domain Access Control via Verifiable Identity and Policy-Based Authorization,” in *2025 IEEE 24th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, IEEE, 2025, pp. 1489–1498. Available: https://doi.org/10.1109/Trustcom66490.2025.00173

[31] R. Chen et al., “BIdM: A Blockchain-Enabled Cross-Domain Identity Management System,” *Journal of Communications and Information Networks*, vol. 6, no. 1, pp. 44–58, 2021. Available: https://doi.org/10.23919/JCIN.2021.9387704

[32] S. R. Garzon, H. Dinh-Tuan, M. M. Martinez, A. Küpper, H. J. Einsiedler, and D. Schneider, “Beyond Certificates: 6G-ready Access Control for the Service-Based Architecture with Decentralized Identifiers and Verifiable Credentials,” in *2024 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, IEEE, 2024, pp. 830–835. Available: https://doi.org/10.1109/EuCNC/6GSummit60053.2024.10597085

[33] H. Dinh-Tuan, S. R. Garzon, and J. Fu, “Secure and Trustful Cross-Domain Communication with Decentralized Identifiers in 5G and Beyond,” in *2024 27th Conference on Innovation in Clouds, Internet and Networks (ICIN)*, IEEE, 2024, pp. 32–36. Available: https://doi.org/10.1109/ICIN60470.2024.10494437

[34] F. Wang, S. Yang, M. Zhang, Y. Liu, Q. Zhang, and F. Tian, “DAVC: A Verifiable Credentials Model for Dynamic Attribute Management,” in *2025 IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA)*, IEEE, 2025, pp. 1431–1438. Available: https://doi.org/10.1109/ISPA67752.2025.00197

[35] N. Fotiou et al., “Continuous authorization over HTTP using Verifiable Credentials and OAuth 2.0,” in *Open Identity Summit*, 2022, pp. 39–50. Available: https://doi.org/10.18420/OID2022_03

[36] A. Cuzzocrea and I. Belmerabet, “Privacy-Preserving Identity and Access Management in Multiple Cloud Environments: Models, Issues, and Solutions,” in *2024 IEEE International Conference on Big Data (BigData)*, IEEE, 2024, pp. 6223–6229. Available: https://doi.org/10.1109/BigData62323.2024.10826004

[37] M. Sporny, D. Longley, D. Chadwick, and I. Herman, “Verifiable Credentials Data Model v2.0.” Available: https://www.w3.org/TR/vc-data-model-2.0/. Accessed on Apr. 1, 2026.

[38] M. Sporny, D. Longley, M. Sabadello, D. Reed, O. Steele, and C. Allen, “Decentralized Identifiers (DIDs) v1.0: Core architecture, data model, and representations.” Available: https://www.w3.org/TR/did-1.0/. Accessed on Apr. 4, 2026.

[39] J. C. John, A. Gupta, and S. Sural, “Secure multi-cloud collaboration using data leakage free attribute-based access control policies,” *Computers & Security*, vol. 161, article 104736, 2026. Available: https://doi.org/10.1016/j.cose.2025.104736

[40] J. Tian, “Zero trust anonymous access algorithm for multi cloud storage system based on CP-ABE,” *Egyptian Informatics Journal*, vol. 30, article 100681, 2025. Available: https://doi.org/10.1016/j.eij.2025.100681

[41] M. Bakhtina, J. Kvapil, P. Švenda, and R. Matulevičius, “The Power of Many: Securing Organisational Identity Through Distributed Key Management,” in *Advanced Information Systems Engineering. CAiSE 2024. Lecture Notes in Computer Science*, vol. 14663, 2024, pp. 475–491. Available: https://doi.org/10.1007/978-3-031-61057-8_28

[42] M. Yıldız and Ş. Bahtiyar, “A Novel Key Management Framework for Secure and Scalable Decentralized Identity Systems,” in *2024 17th International Conference on Security of Information and Networks (SIN)*, IEEE, 2024, pp. 1–8. Available: https://doi.org/10.1109/SIN63213.2024.10871880

[43] A. Alajroush, R. Latif, and T. Saba, “Trust Management Frameworks In Multi-Cloud Environment: A Review,” in *2023 Sixth International Conference of Women in Data Science at Prince Sultan University (WiDS PSU)*, IEEE, 2023, pp. 117–120. Available: https://doi.org/10.1109/WiDS-PSU57071.2023.00034

Appendix A

Table A1. List of abbreviations and acronyms

Abbreviation	Meaning	Abbreviation	Meaning
5G	fifth-generation mobile communications	6G	sixth-generation mobile communications
ABAC	attribute-based access control	AI	artificial intelligence
API	application programming interface	AWS	Amazon Web Services
CP-ABE	ciphertext-policy attribute-based encryption	DID	Decentralized Identifier
DLT	distributed ledger technology	DOI	Digital Object Identifier
EHDEN	European Health Data & Evidence Network	HPC	high-performance computing
HSM	hardware security module	HTTP	Hypertext Transfer Protocol
IAM	identity and access management	IdP	identity provider
IIoT	Industrial Internet of Things	IoT	Internet of Things
MC-ABAC	multi-cloud attribute-based access control	MFA	multifactor authentication
NIST	National Institute of Standards and Technology	OIDC	OpenID Connect
PDP	policy decision point	PEP	policy enforcement point
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses	RBAC	role-based access control
RQ	research question	SAML	Security Assertion Markup Language
SP	Special Publication	SSI	self-sovereign identity
SSO	single sign-on	TCP	Transmission Control Protocol
TLS	Transport Layer Security	TS	Topic Search
UML	Unified Modeling Language	VC	Verifiable Credential
W3C	World Wide Web Consortium	ZKP	zero-knowledge proof
ZTA	Zero Trust Architecture		

Appendix B

Table B1. Characteristics of the studies included in the qualitative synthesis. The reference numbers in Table B1 correspond to the numbers in the list of references of the main body of the article.

Ref.	Environment	Proposed solution	Method, validation, and principal contribution	Main limitation and justification
[3]	General enterprise and network Zero Trust literature	Multivocal review and research-gap framework for Zero Trust	Synthesizes academic and practitioner evidence and identifies recurring architectural and organizational research gaps	The review combines heterogeneous academic and practitioner evidence without a common empirical benchmark. It identifies broad research gaps but cannot establish which mechanisms operate effectively across real multi-cloud deployments.
[4]	Organizational Zero Trust adoption across industries	Critical-success-factor and maturity-assessment framework	Three-round Delphi study with 12 cybersecurity experts; defines organizational factors and maturity dimensions for Zero Trust adoption	The Delphi study assesses organizational adoption factors rather than a technical cross-cloud architecture. It therefore informs readiness and governance but does not validate identity portability, credential lifecycle, or cross-provider policy enforcement.
[5]	Cloud, distributed, and multi-cloud networks	Comparative review and taxonomy of Zero Trust models and implementations	Compares architectures, enabling technologies, enforcement mechanisms, and proofs of concept in cloud environments	The review compares architectures and proofs of concept but contains limited operational multi-cloud evidence. It therefore cannot establish that any reviewed approach provides interoperable trust across independently governed provider IAM systems.
[6]	Multi-institution web ecosystem and health-data portal	Federated authentication among multiple institutional identity providers through a trusted third party	Implemented in the European Health Data & Evidence Network (EHDEN) portal; demonstrates practical authentication across independent institutional IdPs	Authentication is mediated by a trusted third party, while authorization remains outside the evaluated federation path. Central availability and policy portability therefore remain unresolved, limiting the architecture as a complete cross-domain trust model.
[7]	Multi-domain cloud-edge computing continuum	Decentralized IAM integrated with Zero Trust using DIDs, VCs, SSI, and blockchain anchoring	Prototype and performance experiments demonstrate portable identity evidence, local verification, and selective disclosure	The prototype validates portable identity evidence, selective disclosure, and local verification, but does not fully define issuer admission, status-service failure, or key-compromise recovery. The evidence therefore does not cover the complete credential lifecycle.
[14], [15]	Privacy-preserving identity management across multiple clouds	Solid-based architecture using user-controlled data stores, semantic matching, zero-knowledge proofs, and a software-component blueprint	Architectural design study. The first report introduces the architecture; the second extends it with software components and defines Unified Modeling Language models for user-controlled identity and semantic verification	The two reports progressively specify the architecture, but neither provides an implemented multi-cloud evaluation or operational issuer-governance, revocation, and key-management procedures. The evidence therefore remains architectural rather than operational.
[28]	Dynamic cloud-edge–Internet of Things (IoT) computing continuum	Zero Trust access control using DIDs, VCs, and relationship-based authorization	Proof-of-concept impl. demonstrates decentralized identification and access control across continuum domains	The proof of concept is limited to a computing-continuum scenario and does not evaluate a complete credential lifecycle or semantic policy mapping. Transferability to independently governed cloud providers is therefore not established.
[22]	Artificial intelligence (AI) and high-performance computing (HPC) digital research infrastructures	Federated single sign-on (SSO) combined with multifactor authentication (MFA), time-limited roles, monitoring, and layered Zero Trust controls	Deployed architecture and case-study scenarios demonstrating operational federation in research infrastructures	The deployed design is operational but tied to institutional identity providers and coordinated role governance in one research infrastructure. It therefore does not demonstrate portable trust across providers with independent IAM and policy models.
[41]	Cross-organizational data-exchange systems	Distributed organizational key management using threshold or distributed signing	Design-science evaluation and proof of concept show reduced dependence on one administrator or signing key	The study addresses organizational signing-key custody rather than a complete multi-cloud identity and authorization workflow. Its contribution is therefore directly relevant to key governance but does not validate end-to-end cross-cloud trust.
[23]	Cross-domain Industrial Internet of Things	Blockchain-oriented decentralized federated-identity architecture	Architectural analysis shows how decentralized identities may reduce dependence on a global identity provider	The decentralized federation architecture is not accompanied by a complete implementation or measured evaluation. Issuer governance and revocation remain unspecified, so operational feasibility and credential-lifecycle coverage are not demonstrated.
[17]	Cloud-native applications distributed across multiple clouds	Service-mesh-based Zero Trust enforcement using service identities and per-request controls	Experimental evaluation measures performance and resource overhead introduced by Zero Trust controls	The experiments measure the performance overhead of service-mesh Zero Trust controls but do not examine identity portability, issuer governance, revocation, or cross-provider policy mapping. Performance evidence alone cannot establish a complete multi-cloud trust solution.

Table B1. Continued

Ref.	Environment	Proposed solution	Method, validation, and principal contribution	Main limitation and justification
[29]	Telecommunications service management across multiple clouds	Multi-party service-management framework integrating SSI and blockchain-based trust	Prototype scenario and latency measurements demonstrate credential-based cross-cloud service orchestration	The prototype is evaluated in telecommunications service orchestration and only partly specifies revocation and compromise recovery. Its lifecycle controls and transferability to general enterprise multi-cloud environments therefore remain uncertain.
[21]	Federated cloud-edge computing	Transparent proxy-based authentication federation between cloud and edge domains	Prototype and comparative delay evaluation demonstrate interoperability between otherwise incompatible authentication systems	The proxy enables authentication interoperability but remains a mandatory online trust dependency, while authorization stays provider-local. The solution therefore relocates rather than removes centralized availability and policy-portability constraints.
[18]	Enterprise Kubernetes workloads across multiple clouds	Secretless workload authentication using OpenID Connect (OIDC) and Workload Identity Federation	Enterprise implementation demonstrates short-lived workload identities without long-lived service-account secrets	The implementation eliminates long-lived workload secrets but remains provider-managed and excludes human and device identities and semantic policy mapping. The demonstrated portability is therefore limited to workload authentication.
[30]	Heterogeneous cross-domain applications and organizations	POLARIS architecture using VCs, selective disclosure, and policy-based authorization	Working prototype and experiments demonstrating issuer-bound attributes, verifiable presentations, and cross-domain policy evaluation	The prototype evaluates VC-based policy authorization, but issuer admission, credential status, key recovery, and cloud-specific policy mappings remain incomplete. It therefore does not demonstrate an end-to-end multi-cloud trust lifecycle.
[42]	Large-scale decentralized digital-identity systems	Hierarchical key-management framework using hardware security modules (HSMs), isolation, rotation, delegation, and recovery	Security analysis and framework assessment provide comprehensive decentralized-identity key-lifecycle functions	The framework covers key-lifecycle controls comprehensively but is not evaluated in a heterogeneous multi-cloud deployment and retains hierarchical administration. Its cross-domain applicability and decentralization benefits therefore remain unverified.
[34]	Decentralized Web3 identity ecosystems	Dynamic Attribute-Oriented Verifiable Credential model with Merkle structures and zero-knowledge proofs	Prototype experiments demonstrate dynamic attribute updates, selective disclosure, and efficient attribute verification	The experiments validate dynamic attributes and selective disclosure in a Web3-oriented design, but issuer governance and cloud-policy mapping are outside the evaluated scope. The contribution is therefore limited mainly to credential representation and update.
[40]	Multi-cloud storage	Zero Trust anonymous-access algorithm combining hierarchical identity-based encryption and ciphertext-policy attribute-based encryption (CP-ABE)	Algorithmic and experimental evaluation demonstrates fine-grained encrypted access and key-version-based revocation	The scheme validates fine-grained encrypted storage access and key-version revocation, but its storage-specific model does not address general federation, issuer governance, or policy interoperability. It cannot therefore represent a complete cross-cloud trust architecture.
[31]	Cross-domain network services	Blockchain-enabled decentralized identity-management system using DIDs and accumulators	Theoretical analysis and prototype evaluation demonstrate cross-domain identity authentication and distributed trust establishment	The system evaluates cross-domain identity authentication but gives limited treatment to authorization semantics, policy mapping, and recovery. This consequently establishes identity portability more strongly than authorization portability.
[32]	Multi-stakeholder fifth-generation/sixth-generation (5G/6G) service-based architecture	DID- and VC-based access control replacing conventional certificates and centralized authorization tokens	Prototype experiments demonstrate credential-based authentication and authorization across administrative domains	The prototype demonstrates DID- and VC-based access control in a telecommunications architecture but lacks complete issuer, revocation, and key-recovery processes. Transferability to general multi-cloud governance therefore remains untested.
[33]	Cross-domain 5G and future mobile networks	DID-based communication agent and application-layer secure transport protocol	Implementation and comparison of the Transmission Control Protocol and Transport Layer Security (TCP/TLS) demonstrating decentralized identifier resolution and secure channel establishment	This study evaluates DID resolution and secure cross-domain communication rather than fine-grained resource authorization, which introduces protocol overhead. Its contribution is therefore confined mainly to the communication layer.
[24]	Collaborating tenants across multiple cloud providers	Multi-cloud attribute-based access control (MC-ABAC) model for cross-tenant resource sharing	Python implementation and feasibility evaluation demonstrate attribute-based collaboration among independent cloud tenants	The model and Python evaluation demonstrate cross-tenant ABAC feasibility but do not cover realistic deployment, credential lifecycle, or decentralized identity integration. Operational applicability to a complete multi-cloud trust environment therefore remains unverified.

Table B1. Continued

Ref.	Environment	Proposed solution	Method, validation, and principal contribution	Main limitation and justification
[25]	Amazon Web Services (AWS), Microsoft Azure, and Google Cloud authorization environments	Formal policy aggregation and conflict detection using Event Calculus	Automated tool and performance evaluation demonstrate detection of conflicting authorization rules across three providers	Conflict detection depends on the prior translation of provider policies and does not itself establish semantic equivalence. The reliability of the detected conflicts is therefore constrained by the accuracy of the mapping supplied to the formal model.
[19]	Multi-tenant hybrid cloud implemented on Google Cloud Platform	Mirror-Identity IAM framework linking on-premises and cloud identities	Prototype and scalability assessment demonstrate coordinated identity administration for shared hybrid-cloud resources	The prototype coordinates on-premises and cloud identities through duplication and synchronization but does not demonstrate portability across different provider IAM models. It therefore preserves the synchronization problem identified in RQ2.
[20]	Consumer and provider domains in a multi-cloud environment	Federated trust combining SSO, MFA, and risk-sensitive authentication	Implementation and comparative case evaluation demonstrate strengthened authentication across participating cloud domains	The implementation strengthens authentication through SSO, MFA, and risk-sensitive controls but relies on centralized federation and does not evaluate cross-cloud authorization semantics or credential lifecycle. Its contribution therefore remains authentication-centered.
[26]	Hybrid Cloud-IoT machine-to-machine services	Containerized microservice architecture separating policy decisioning and enforcement	Prototype evaluation measures latency, throughput, and successful authorization requests	The prototype evaluates authorization-service performance in a specialized Cloud-IoT scenario but continues to depend on central policy services. Availability and cross-domain decision autonomy therefore remain limited.
[27]	Hybrid-cloud role-based access control (RBAC) within one organization	Systematic review and requirements framework for mapping central roles to cloud-specific entitlements	Synthesizes role-mapping methods and derives requirements for inter-domain RBAC interoperability	The review derives role-mapping requirements but does not implement or experimentally evaluate a mapping solution. Semantic equivalence and the maintenance of provider-specific mappings therefore remain unresolved.
[43]	Trust management in multi-cloud environments	Review and classification of existing multi-cloud trust-management frameworks	Compares trust assessment, coordination, compliance, and governance mechanisms across published approaches	The review compares trust-management frameworks descriptively without a common empirical benchmark or implemented reference architecture. It therefore cannot determine the relative operational effectiveness of the reviewed governance mechanisms.
[39]	Collaborative data sharing across multiple clouds	Data-leakage-free attribute-based access control (ABAC) policy-generation model based on trust and access constraints	Formal analysis and experiments on large datasets demonstrate policy generation that permits collaboration while limiting leakage paths	The model validates leakage-free policy generation under assumed trust and access data but does not address identity federation or credential lifecycle. Its evidence therefore covers policy construction rather than complete cross-domain trust management.
[35]	Hypertext Transfer Protocol (HTTP) resources accessed through web browsers	Continuous authorization using VCs, the OAuth 2.0 authorization framework, proof of possession, and request-level verification	System implementation and experiments demonstrate repeated authorization and credential revocation for HTTP requests	The implementation demonstrates repeated HTTP authorization and credential revocation but is browser- and HTTP-specific and provides limited treatment to issuer governance, recovery, and multi-cloud policy mapping. Broader applicability therefore remains untested.
[36]	Privacy-preserving IAM across multiple-cloud environments	Taxonomy and comparative synthesis of privacy-preserving identity and access-management models	Organizes existing approaches by privacy, interoperability, data minimization, and access-control properties	The study is a secondary taxonomy without a prototype, and the compared models use heterogeneous assumptions and evaluation criteria. It supports the classification of the literature but not conclusions about operational effectiveness.

Note. References [14] and [15] describe successive versions of the same Solid-based multi-cloud identity-management study and are summarized in one row. They were treated as a single study in the numerical synthesis.