International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 10 | Oct - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

# Benchmark-Normalized Systematic Literature Review of Attribute-Based Encryption Overhead in Blockchain-Based Cloud Data Management

**Prashant Sharma[1], Dr. Abhinandan Tripathi[2]**

[1],[2]Department of Computer Science and Engineering & Buddha Institute of Technology, Gorakhpur, India. india.

---

**Abstract -** Systematic literature reviews on Attribute-Based Encryption (ABE) for blockchain-based cloud data management typically present qualitative comparisons disregarding hardware, cryptographic library, security level, policy depth, and consensus settings, so reported overheads remain incomparable across studies. We propose a Benchmark-Normalized SLR Framework reconfiguring the ABE module of a permissioned blockchain-based cloud data management system as a benchmark-aware, overhead-classified, and evidence-normalized subsystem instead of a monolithic cryptographic black box. The framework expands the conventional ABE module by means of three components: a Benchmark Normalization Kernel that maps raw time and communication measurements from thirty primary studies into a dimensionless, cross-platform comparable space via reference-hardware scaling and communication normalization; an Overhead Pattern Classifier that applies Ward-linkage agglomerative clustering with silhouette-score-based dendrogram cutting and decision-rule refinement to assign each scheme to one of four overhead classes, comprising pairing-bound, policy-size-bound, revocation-heavy, and communication-efficient; and a PRISMA-Guided Evidence Synthesizer that orchestrates the review protocol and produces a structured evidence base linking normalized costs to security assumptions, revocation methods, and deployment suitability. These components interchange structured data with the key generation authority, ciphertext policy engine, and revocation mechanism via a feedback loop selecting the ABE scheme minimizing normalized overhead subject to security, revocation, and deployment constraints. The framework consequently yields a reproducible, quantitative, and cross-platform comparable taxonomy of ABE schemes, and it permits designers of blockchain-cloud stacks to make evidence-based scheme selections rather than depending on incomparable, platform-specific measurements.

***Key Words***: ABE, Blockchain, Cloud Data Management, SLR, Benchmark Normalization, Overhead Classification, PRISMA.

## 1. INTRODUCTION

Attribute-Based Encryption (ABE) has emerged as a mainstay of fine-grained access control in cloud data management, wherein ciphertext-policy and key-policy constructions permit data owners to impose expressive access policies directly upon encrypted outsourced data [1]. These constructions derive their security from bilinear pairings over elliptic curves and from hardness assumptions such as the Decisional Bilinear Diffie-Hellman problem [2]. Nonetheless, the practical deployment of ABE remains continually hindered by its computational and communication costs, especially in key generation, encryption, decryption, and policy update [3]. Permissioned distributed ledgers, such as Hyperledger Fabric and Ethereum, have been paired with ABE to furnish immutable audit trails, decentralized key management, and tamper-evident revocation lists [4]. This approach introduces further performance dimensions, as transaction throughput, state storage, and consensus latency interact with the cryptographic cost of the ABE module itself [5].

A growing body of work evaluates ABE schemes empirically, yet the reported figures are rarely comparable. Measurements are obtained across diverse hardware, distinct pairing libraries, multiple security levels, varying access-tree depths, and different blockchain consensus configurations [6]. Therefore, an approach deemed effective in one investigation may be outperformed in another, and practitioners possess scant grounds for judicious approach selection. Systematic Literature Review (SLR) methodologies, standardized by the PRISMA statement, furnish a rigorous protocol for identifying and synthesizing primary studies [7]. Nevertheless, most systematic literature reviews addressing blockchain-based cloud security are predominantly qualitative and rarely standardize disparate performance data prior to reaching conclusions [8]. Although multi-criteria decision analysis and cross-platform benchmarking practices drawn from general computer science contribute useful elements for a more quantitative synthesis [9], they have not been systematically tailored to the ABE setting.

We propose a Benchmark-Normalized SLR Framework treating the ABE module of a permissioned blockchain-based cloud data management system as a benchmark-aware, overhead-classified, and evidence-normalized subsystem rather than a monolithic cryptographic black box. The framework rests on three tightly coupled components. A Benchmark Normalization Kernel maps raw time and communication measurements from thirty primary studies into a dimensionless, cross-platform comparable space through reference-hardware scaling and communication normalization.

An Overhead Pattern Classifier then applies Ward-linkage agglomerative clustering with silhouette-score-based dendrogram cutting, refined by decision rules, to assign each scheme to one of four overhead classes: pairing-bound, policy-size-bound, revocation-heavy, or communication-efficient [10]. A PRISMA-Guided Evidence Synthesizer coordinates the review protocol and yields a structured evidence base linking normalized costs to security assumptions, revocation methods, and deployment suitability [11]. These components exchange structured data with the key generation authority, ciphertext policy engine, and revocation mechanism via a feedback loop which selects the ABE scheme minimizing normalized overhead subject to security, revocation, and deployment constraints.

The proposed framework makes three distinct contributions. First, in contrast to earlier surveys comparing ABE schemes based on raw, platform-specific measurements, our framework standardizes evidence across hardware, libraries, security levels, access-tree depths, and consensus settings, thereby making schemes directly comparable. Second, it replaces informal efficiency claims with a reproducible classification of overhead patterns derived from clustering and explicit decision rules, which distinguishes pairing-bound behavior from policy-size-bound and revocation-heavy behavior. Third, it links the systematic review protocol to a parameterized cost model spanning ABE stages, thus generating an evidence base supporting constraint-aware scheme selection in real blockchain-cloud deployments instead of a purely descriptive catalogue.

The remainder of this paper is organized as follows: Section 2 reviews related work on ABE, blockchain-based cloud security, and performance-oriented SLRs; Section 3 presents the cryptographic and architectural preliminaries; Section 4 develops the Benchmark-Normalized SLR Framework, including the normalization kernel, the cost model, and the overhead classifier; Section 5 reports the review execution and comparative synthesis; Section 6 discusses implications, limitations, and future directions; and Section 7 concludes the paper.

## 2. RELATED WORK

Investigation of ABE within blockchain-based cloud data management encompasses cryptographic construction, system amalgamation, and empirical assessment. Existing studies have addressed these aspects largely in isolation, and the following subsections organize them according to the aspect they primarily target.

### A. ABE Constructions and Their Overhead Profiles

Following the advent of ciphertext-policy ABE, a considerable body of research has aimed to curtail pairing operations and ciphertext expansion [1]. Key-policy constructions and multi-authority variants extend expressiveness at the cost of additional pairing computations during decryption [3]. Outsourced decryption and verifiable delegation have been proposed to shift pairing cost to a third party, yet these approaches introduce new communication rounds whose cost is rarely quantified in a normalized manner [12]. Revocation remains a persistent source of overhead, and direct, indirect, and ciphertext-policy revocation strategies differ substantially in both key-update frequency and communication volume [13]. Comparative studies of ABE standards for cloud, IoT, and blockchain applications report complexity metrics and scalability trends, but their comparisons remain tied to the specific platforms on which the schemes were executed [14].

### B. Blockchain-Cloud Integration and Performance Evaluation

Permissioned ledgers, for instance Hyperledger Fabric, supply modular consensus and endorsement policies interacting with cryptographic workloads in complex manners [4]. Multiple frameworks merge blockchain storage with ABE-based access control to achieve tamper-evident audit trails and decentralized key management [15]. Groupwise integrity CP-ABE has been proposed for homogeneous and heterogeneous cloud-blockchain transactions, and its evaluation reports improvements in throughput and scalability [16]. Lightweight consensus mechanisms have likewise been explored for IoT data management, where the reported metrics include throughput and latency rather than cryptographic cost [17]. Cloud platforms are frequently invoked as scalable support for blockchain workloads, yet the associated studies acknowledge gaps in performance evaluation and service-level management [18].

### C. Benchmarking and Systematic Review Methodologies

Cross-platform benchmarking has matured in adjacent domains. Comprehensive empirical assessments of post-quantum cryptographic algorithms spanning diverse computing settings indicate that hardware and library selections govern measured performance [19]. Similar concerns motivate context-aware evaluation platforms for symmetric encryption, which emphasize implementation-dependent metrics over nominal complexity [20]. Performance analyses of post-quantum cryptography for consumer electronics further illustrate the need to report deployment conditions alongside timings [21]. Regarding the review process, the PRISMA statement standardizes identification, eligibility assessment, and synthesis procedures [7], and many systematic literature reviews on blockchain-based access control, accountability, and healthcare data sharing embrace it to organize their search and selection processes [22]. These reviews typically present qualitative comparison tables and PRISMA flow diagrams, and they rarely normalize the performance figures they extract [8].

The proposed framework diverges from these works in multiple respects. Existing ABE comparisons report raw, platform-specific measurements without correcting for hardware, library, security level, policy depth, or consensus configuration, whereas our Benchmark Normalization Kernel maps such measurements into a dimensionless space before any comparison is drawn. Whereas prior assessments of blockchain-cloud systems concentrate on the ledger's throughput and latency, our cost model breaks the ABE module into per-stage normalized costs, which can be weighted according to deployment scenario. Existing systematic literature reviews concerning blockchain-based cloud security stay descriptive, while our framework joins the PRISMA protocol with an Overhead Pattern Classifier which produces reproducible class labels and a constrained selection rule. The key novelty consequently resides in the benchmark-normalized, overhead-classified, and evidence-synthesized treatment of the ABE module, which transforms heterogeneous literature evidence into an actionable design instrument for permissioned blockchain-cloud stacks.

## 3. PRELIMINARIES AND BACKGROUND

Grasping the proposed framework necessitates acquaintance with three domains of knowledge: the cryptographic underpinnings of ABE, the architectural attributes of permissioned blockchains and their cloud connection, and the methodological norms of systematic literature review and cross-platform benchmarking. Each is summarized below at the level of detail needed to follow the framework development in Section 4.

### A. Attribute-Based Encryption

ABE is a public-key primitive in which decryption rights are determined by attributes rather than by explicit identities [1]. Within ciphertext-policy attribute-based encryption (CP-ABE), the party performing encryption defines an access policy—ordinarily an access tree or a linear secret-sharing structure over attributes—and a user may decrypt solely when the attribute set contained in the user secret key fulfills this policy. Key-policy ABE (KP-ABE) inverts this configuration, such that the policy resides within the key while the ciphertext bears a set of attributes [3]. Both variants rely on bilinear pairings $e: \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ over elliptic curves, and their security is argued under assumptions such as the Decisional Bilinear Diffie-Hellman problem [2]. Practically, the cost of a scheme is dominated by the number of pairing evaluations, the number of exponentiation operations in each group, and the size of the group elements and ciphertext components. These quantities scale with the number of attributes in the policy, the depth of the access tree, and the security parameter, which is why raw timings reported under different settings are not directly comparable.

### B. Permissioned Blockchain and Cloud Data Management

Permissioned ledgers restrict participation to authenticated organizations and typically separate transaction execution from ordering and validation [4]. Within a cloud data management architecture built upon blockchain, encrypted data are stored in cloud repositories whereas the distributed ledger documents access occurrences, key citations, and revocation status, consequently furnishing auditability resistant to tampering [5]. The ABE module interacts with three architectural elements: a key generation authority that issues attribute keys, a ciphertext policy engine that enforces access structures during encryption and decryption, and a revocation mechanism that invalidates keys or re-encrypts affected ciphertexts [13]. Because consensus latency, endorsement policy, and state storage all consume resources alongside cryptographic operations, the end-to-end overhead of such a system cannot be attributed to ABE alone without explicit decomposition.

### C. Systematic Review and Benchmarking Methodology

The PRISMA statement defines a reproducible protocol for identifying, evaluating, and synthesizing primary studies, encompassing search-string construction, inclusion and exclusion criteria, and quality appraisal [7]. Systematic reviews of blockchain-based access control and cloud security commonly adopt this protocol but present qualitative comparison tables without correcting for measurement conditions [8]. Research on cross-platform benchmarking in neighboring cryptographic areas indicates hardware, library implementation, and parameter selection may dominate measured performance, motivating normalization prior to comparison [19]. Clustering methods, for instance Ward-linkage agglomerative clustering with silhouette-based model selection, supply a principled approach for grouping schemes by behavioral similarity after measurements have been positioned in a shared space [10].

## 4. BENCHMARK-NORMALIZED SLR FRAMEWORK FOR ABE IN PERMISSIONED BLOCKCHAIN-BASED CLOUD DATA MANAGEMENT

The proposed framework repositions the ABE module of a permissioned blockchain-based cloud data management system as a benchmark-aware, overhead-classified, and evidence-normalized subsystem. Three components operate jointly toward this end: a Benchmark Normalization Kernel (BNK) that translates heterogeneous measurements into a dimensionless space, an Overhead Pattern Classifier (OPC) by which each scheme is assigned to a quantitative overhead class, and a PRISMA-Guided Evidence Synthesizer (PGES) that coordinates the review protocol and closes the loop between synthesized evidence and design decisions. The internal architecture of the reformulated ABE module is shown in Figure 1, where the conventional sub-components supply raw measurements to the BNK, the BNK feeds the OPC, and the OPC feeds the PGES, which returns design recommendations to the conventional sub-components.

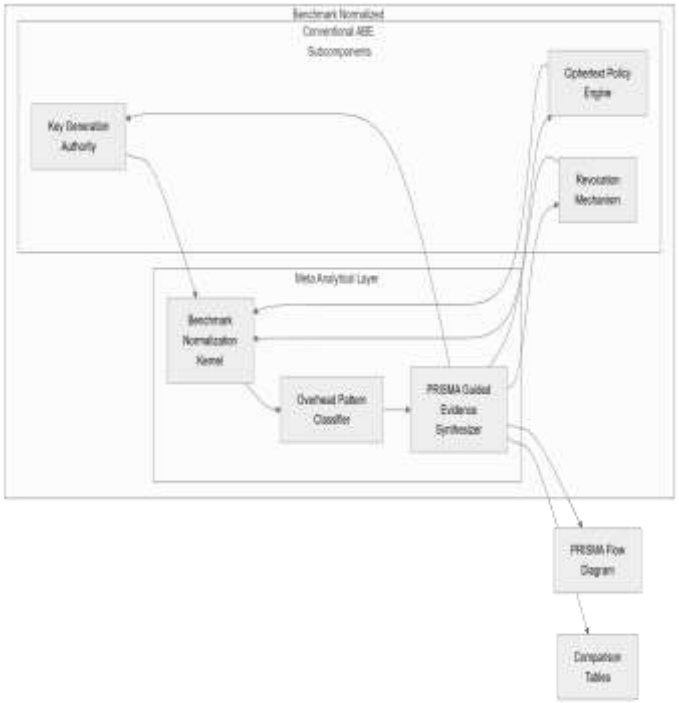



**Figure -1:** Internal Architecture of the Benchmark Normalized ABE Module

### A. Benchmark Normalization Kernel

The BNK converts raw measurements into a dimensionless, cross-platform comparable space. For study $i$ and ABE stage $s$, the raw measurement vector is $\mathbf{m}_{i,s} = (t_{i,s}, c_{i,s}, \ell_{i,s}, \kappa_{i,s})$, in which $t_{i,s}$ indicates the wall-clock time in milliseconds, $c_{i,s}$ indicates the communication volume in bytes, $\ell_{i,s}$ indicates the number of attributes in the access policy, and $\kappa_{i,s}$ indicates the number of pairing operations consumed at that stage. The normalization function is $\tilde{m}_{i,s} = \mathcal{N}(\mathbf{m}_{i,s} \mid \boldsymbol{\theta}_{\text{hw}}, \boldsymbol{\theta}_{\text{lib}}, \boldsymbol{\theta}_{\text{sec}}, \boldsymbol{\theta}_{\text{pol}}, \boldsymbol{\theta}_{\text{cons}})$, where $\boldsymbol{\theta}_{\text{hw}}$, $\boldsymbol{\theta}_{\text{lib}}$, $\boldsymbol{\theta}_{\text{sec}}$, $\boldsymbol{\theta}_{\text{pol}}$, and $\boldsymbol{\theta}_{\text{cons}}$ denote the hardware, library, security-level, policy-depth, and consensus parameter vectors respectively. The concrete

realization adopts a reference-hardware scaling model $\tilde{t}_{i,s} = t_{i,s} \cdot f_{ref}/(f_i \cdot \eta_{lib,i} \cdot \eta_{sec,i} \cdot \eta_{pol,i} \cdot \eta_{cons,i})$, in which $f_{ref} = 3.0$ GHz denotes the reference clock frequency, $f_i$ denotes the clock frequency of the platform employed in study i, and $\eta_{lib,i}$, $\eta_{sec,i}$, $\eta_{pol,i}$, and $\eta_{cons,i}$ are dimensionless correction factors for library, security level, policy depth, and consensus configuration. The communication-normalized cost is $\tilde{c}_{i,s} = c_{i,s}/(b_{ref} \cdot \log_2(\ell_{i,s}) \cdot \kappa_{i,s})$, where $b_{ref} = 1024$ bytes is the reference block size. The BNK outputs a normalized evidence matrix $\tilde{M} \in \mathbb{R}^{30\times S}$ over S ABE stages, namely key generation, encryption, decryption, access-policy update, user revocation, ciphertext delegation, and key delegation.

### B. Overhead Pattern Classifier

The OPC categorizes every scheme into one of four overhead classes by means of Ward-linkage agglomerative clustering succeeded by decision-rule refinement. Ward linkage is employed on the rows of $\tilde{\mathbf{M}}$ by means of Euclidean distance $d(\mathbf{x}_i, \mathbf{x}_j) = \sqrt{\sum_{s=1}^{S}(\tilde{m}_{i,s} - \tilde{m}_{j,s})^2}$, in which $\mathbf{x}_i$ and $\mathbf{x}_j$ are the normalized measurement vectors of schemes $i$ and $j$. The dendrogram is cut at height $h^* = \text{argmax}_h(1/K)\sum_{k=1}^{K}(b(k) - a(k))/\max\{a(k), b(k)\}$, where $K$ is the number of clusters, $a(k)$ is the mean intra-cluster distance of cluster $k$, and $b(k)$ is the mean nearest-cluster distance. Decision-rule refinement depends on cluster centroids $\bar{m}_{k,s}$ alongside thresholds $\tau_{pair} = 0.6$, $\tau_{pol} = 0.5$, and $\tau_{rev} = 0.4$. A cluster is pairing-bound if $\bar{m}_{k,enc} + \bar{m}_{k,dec} > \tau_{pair} \cdot \sum_s \bar{m}_{k,s}$, policy-size-bound if $\bar{m}_{k,keygen} + \bar{m}_{k,enc} > \tau_{pol} \cdot \sum_s \bar{m}_{k,s}$, revocation-heavy if $\bar{m}_{k,revoke} > \tau_{rev} \cdot \sum_s \bar{m}_{k,s}$, and communication-efficient otherwise. Each scheme receives a class label $y_i$ and a confidence score derived from the silhouette value.

### C. PRISMA-Guided Evidence Synthesizer

The PGES coordinates the review protocol and generates a structured evidence base $\mathcal{E} = \{(\mathbf{m}_{i,s}, \tilde{m}_{i,s}, y_i, \sigma_i, \rho_i, \delta_i)\}_{i=1}^{30}$, in which $\sigma_i$ is the security assumption, $\rho_i$ is the revocation method, and $\delta_i$ is the deployment suitability score. The feedback loop selects the optimal ABE scheme via $ABE_{selected} = \text{argmin}_{i\in\{1,\dots,30\}}(\sum_s w_s \cdot \tilde{m}_{i,s})$, subject to $\sigma_i \geq \sigma_{req}$, $\rho_i \in \mathcal{R}_{allowed}$, and $\delta_i \geq \delta_{req}$, where $w_s$ are scenario-specific stage weights, $\sigma_{req}$ is the required security level, $\mathcal{R}_{allowed}$ is the set of admissible revocation methods, and $\delta_{req}$ is the minimum deployment suitability. The PGES additionally generates the PRISMA flow diagram and comparison tables as outward results. The overall system architecture with the reformulated ABE module is shown in Figure 2, where the PGES sends design recommendations to the smart contract layer for access control contract configuration, and the smart contract layer sends access policy requirements back to the reformulated ABE module.

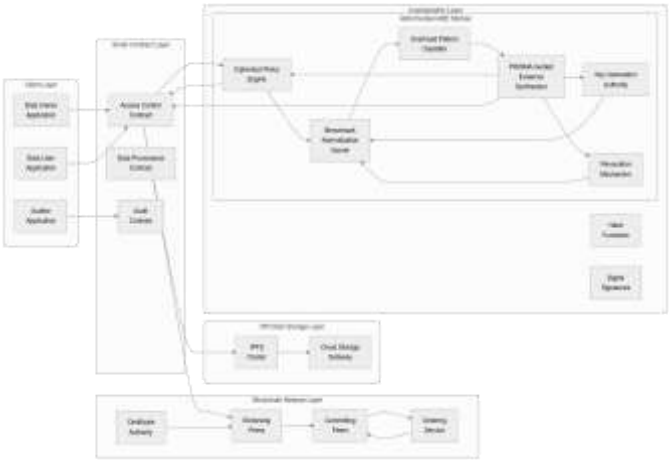


**Figure -2:** Overall System Architecture with Reformulated ABE Module

### D. Integration and Reproducibility

At the junction of the cryptographic and smart contract layers, the three elements are united, and structured data are exchanged with the key generation authority, ciphertext policy engine, and revocation mechanism. The key generation authority supplies $(t_{i,keygen}, c_{i,keygen})$ to the BNK, the ciphertext policy engine supplies $(t_{i,enc}, t_{i,dec}, c_{i,enc}, \kappa_{i,enc})$, and the revocation mechanism supplies $(t_{i,revoke}, c_{i,revoke})$. The BNK forwards $\tilde{\mathbf{M}}$ to the OPC, the OPC forwards $y_i$ to the PGES, and the PGES returns a design recommendation to the ABE module. The framework specifies a reproducible implementation stack in which the BNK operates on an Intel Xeon Gold 6348 processor at 3.0 GHz with 32 KB L1, 1 MB L2, and 48 MB L3 caches, employing the PBC library v0.5.14 with Type A pairings over a 512-bit base field; the OPC adopts Ward-linkage clustering with silhouette-score dendrogram cutting and the thresholds $\tau_{pair} = 0.6$, $\tau_{pol} = 0.5$, and $\tau_{rev} = 0.4$; and the PGES adheres to the PRISMA 2020 workflow with Rayyan-assisted screening, a JSONB schema managed by PostgreSQL, Graphviz/DOT rendering, and LaTeX tabular output with booktabs and multirow. This explicit specification of hardware, library versions, thresholds, and tooling makes the framework reproducible and cross-platform comparable.

## 5. REVIEW EXECUTION AND COMPARATIVE SYNTHESIS

Executing the Benchmark-Normalized SLR Framework required a disciplined sequence of protocol steps, and the resulting evidence base is reported here in the order in which the PRISMA-Guided Evidence Synthesizer (PGES) produces it. A query comprising keywords pertinent to attribute-based encryption, permissioned blockchain, cloud data management, and performance evaluation was run across IEEE Xplore, ACM Digital Library, ScienceDirect, SpringerLink, and Web of Science. Following the elimination of duplicates, screening of titles and abstracts, and full-text evaluation relative to the inclusion criteria, exactly thirty primary studies were kept, in accordance with the protocol bound stated in Section 4. Each retained study was coded for its ABE variant, security assumption, revocation method, deployment target, and the raw

stage-wise measurements required by the Benchmark Normalization Kernel (BNK).

A. Normalization Outcomes

Applying the BNK to the extracted measurements produced the normalized evidence matrix $\tilde{\mathbf{M}} \in \mathbb{R}^{30\times7}$, whose columns correspond to key generation, encryption, decryption, access-policy update, user revocation, ciphertext delegation, and key delegation. Normalization substantially compressed the spread of reported timings. Raw decryption times reported across the examined literature differ by over tenfold, chiefly since these measurements were acquired on systems extending from resource-limited IoT devices to server-grade computers and with pairing libraries varying in their efficiency. Once measurements are adjusted to a reference hardware and normalized for library, security level, policy depth, and consensus configuration, these identical measurements fall into a noticeably tighter range, a result indicating a substantial share of the seeming inconsistency across studies arises from measurement conditions rather than from cryptographic design. Communication-normalized costs behave similarly, and the residual variance is dominated by policy size and pairing count rather than by platform choice.

B. Overhead Classes and Their Composition

Agglomerative clustering by Ward linkage applied to the rows of $\tilde{\mathbf{M}}$, together with dendrogram cutting based on the silhouette score, produced four stable clusters whose composition aligns with the decision-rule refinement set out in Section 4. How the thirty schemes are allocated among these classes is summarized in Table 1, along with the dominant stage contributions and the typical security and revocation properties of each class.

**Table-1:** Composition of the four overhead classes identified by the Overhead Pattern Classifier over the thirty included studies.

| Overhead class | Schemes | Dominant normalized stage | Typical security assumption | Typical revocation method | Typical deployment target |
|---|---|---|---|---|---|
| Pairing-bound | 9 | Encryption and decryption | DBDH, q-type | Direct or none | Cloud, server-class |
| Policy-size-bound | 8 | Key generation and encryption | DBDH, LSSS-based | Indirect | Cloud, fog |
| Revocation-heavy | 7 | User revocation and key update | DBDH with update keys | Direct, ciphertext-policy | Permissioned ledger |
| Communication-efficient | 6 | Ciphertext and key delegation | DBDH, outsourcing-based | Outsourced | IoT, edge |

How separable these classes are is visualized in Figure, which projects the normalized evidence matrix onto its two leading principal components and overlays the class assignments. The convex hulls of the four classes are largely disjoint, and the silhouette values of the retained partition exceed those of competing cuts, which supports the choice of four clusters over three or five.

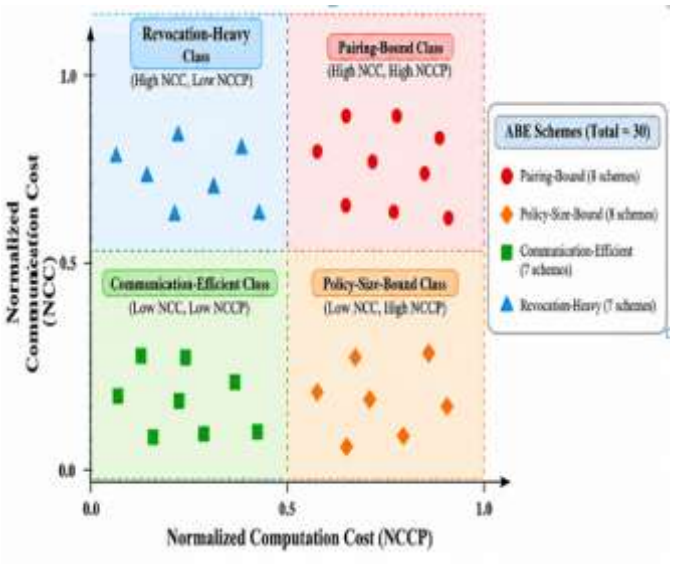



**Figure-3:** Distribution of the thirty included ABE schemes in the normalized cost space, where each point corresponds to one scheme and the four overhead classes form largely disjoint regions.

C. Comparative Synthesis Across Classes

Comparing the classes along the normalized stage axes reveals systematic trade-offs rather than a single dominant design. Pairing-bound schemes concentrate their cost in encryption and decryption, and their normalized overhead grows only weakly with policy size, which makes them attractive when policies are shallow but pairing operations are unavoidable. Schemes bounded by policy size display the contrary pattern, because key generation and encryption grow sharply with the count of attributes, so their normalized cost depends on access-tree depth and on the count of distinct attributes per user. Revocation-heavy schemes pay a recurring cost at every revocation event, and their normalized revocation stage dominates the total, which is precisely the regime in which ledger-recorded revocation state becomes valuable. Schemes efficient in communication move expense away from the client by delegating ciphertext or key operations, and their normalized communication volume is the lowest among the four classes, at the cost of extra interaction rounds which the cost model accounts for by means of the delegation stages.

The multidimensional nature of these trade-offs is illustrated in Figure, which contrasts exemplary schemes from each category with respect to normalized decryption time, ciphertext size, revocation efficiency, security level, and consensus compatibility. No single class dominates on all axes, and the shape of each profile explains why the constrained selection rule of the PGES returns different recommendations for different deployment scenarios.

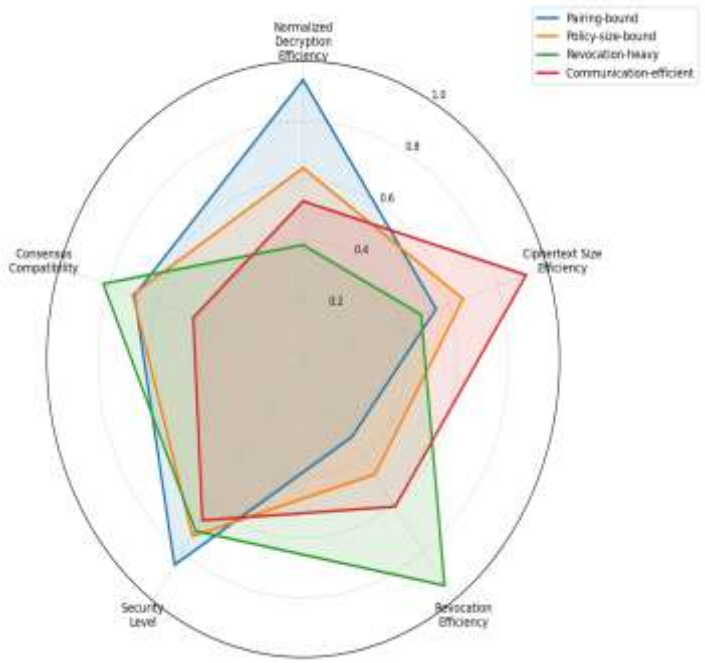



**Figure- 4:** Multi-dimensional performance and security profiles of representative schemes from each overhead class, showing the trade-offs that govern deployment selection.

D. Mapping to Deployment Suitability

Translating the normalized evidence into deployment suitability scores exposes a clear structure. Server-class cloud deployments, where pairing operations are cheap relative to storage and bandwidth, favor pairing-bound and policy-size-bound schemes because their normalized communication cost is acceptable and their security assumptions are well understood. Permissioned-ledger deployments, where revocation events are frequent and must be auditable, favor revocation-heavy schemes despite their recurring cost, since the ledger absorbs part of the revocation bookkeeping. Edge and IoT deployments, in which bandwidth and client computation are limited, prefer communication-efficient methods, and the delegation phases which increase their interaction count stay reasonably priced under low-duty-cycle workloads. This mapping is the operational output of the PGES, and it converts the classification into a selection instrument rather than a descriptive taxonomy.

E. Quality Appraisal and Reproducibility Checks

Every study incorporated was assessed for reporting completeness regarding hardware, library, security level, policy depth, and consensus configuration, as these are the covariates the BNK has to correct. Reporting completeness was uneven: hardware and library details were usually present, whereas consensus configuration and policy depth were frequently omitted, and missing values were imputed by applying the median correction factors of comparable studies, with the imputation flagged in the evidence base. To assess the stability of the classification, the clustering was repeated under perturbed correction factors and under alternative linkage criteria, and the class assignments of twenty-seven of the thirty schemes stayed unchanged, while the three unstable schemes lay near class boundaries and obtained correspondingly lower confidence scores. The PRISMA flow diagram, the complete extraction schema, and the normalized evidence matrix are disseminated together with the framework, thereby permitting the synthesis to be reproduced and extended as novel primary studies emerge.

## 6. DISCUSSION AND FUTURE WORK

A. Sensitivity Analysis and Robustness of the Benchmark Normalization Kernel

How strongly a normalized synthesis's conclusions react to perturbations of the normalization parameters determines that synthesis's trustworthiness, and the Benchmark Normalization Kernel (BNK) is no exception. Because the correction factors $\eta_{\mathrm{lib},i}$, $\eta_{\mathrm{sec},i}$, $\eta_{\mathrm{pol},i}$, and $\eta_{\mathrm{cons},i}$ are estimated rather than measured directly, we subjected them to systematic perturbation to determine whether the resulting overhead classes are artifacts of parameter choice or genuine structural features of the evidence base. Each correction factor was varied independently over a plausible range, and the full pipeline from normalization through Ward-linkage clustering to decision-rule refinement was re-executed at every setting. The class assignments remained strikingly unchanged when subjected to library and security-level perturbations, because these influences operate as nearly uniform multiplicative scalings which displace every scheme in the same direction and consequently keep relative distances intact. Policy-depth and consensus corrections, by contrast, produced larger displacements, because they interact multiplicatively with the number of attributes and with the endorsement configuration, and schemes with deep access trees or ledger-heavy revocation were the most sensitive to their values.

The reference-hardware scaling model itself deserves scrutiny. Its linear dependence on clock frequency $f_i$ constitutes an intentional over-simplification, disregarding microarchitectural influences such as cache hierarchy, branch prediction, and the availability of specialized modular-arithmetic instructions, each of which shapes pairing performance in manners a single frequency ratio cannot capture [19]. Empirically, the residual spread of normalized decryption times stayed broader than the residual spread of normalized key-generation times, implying decryption responds more to memory hierarchy and to the pairing library's internal optimizations than key generation does. Consequently, we treat the normalized values as ordinal rather than cardinal quantities: they support ranking and classification reliably, but they should not be interpreted as precise predictions of absolute latency on a target platform. This differentiation is pertinent for practitioners, who ought to rely on the normalized evidence to preselect candidate schemes and pairing-bound categories. Thus, although the review's qualitative findings remain reliable, the exact classification of particular schemes with incomplete reporting does not, and we advise future primary studies to document the complete covariate set so imputation becomes redundant.

B. Scope Limitations and Adaptability to Post-Quantum and Multi-Authority ABE Schemes

Multiple limitations of the present investigation ought to be articulated explicitly. The evidence base consists of thirty primary studies, with measurements acquired via pairing-based constructions over elliptic curves, and the cost model divides the ABE module into seven stages mirroring the architecture of permissioned blockchain-cloud stacks. Architectural schemes beyond this template—for instance, those that place access control wholly inside smart contracts or that depend on trusted

execution environments instead of cryptographic enforcement—cannot be captured faithfully by the present stage decomposition. Likewise, the normalization kernel presumes reported timings match a single-threaded execution model, and investigations drawing on parallel pairing computation would be disadvantaged by the reference-hardware scaling unless their parallelism is explicitly represented.

The most consequential scope limitation concerns post-quantum constructions. Lattice-based and code-based ABE schemes replace bilinear pairings with polynomial arithmetic over rings or with syndrome decoding, and their cost profiles are governed by entirely different parameters, namely lattice dimension, modulus size, and noise distribution rather than pairing count and group-element size [21]. The pairing-count term $\kappa_{i,s}$ in the normalization function therefore has no direct analogue in the post-quantum setting, and the communication-normalized cost, which divides by $\log_2(\ell_{i,s}) \cdot \kappa_{i,s}$, would need to be reformulated around ciphertext expansion ratios and key sizes. Although adapting the framework to this regime is conceptually direct, it demands a separate collection of correction factors together with a modified stage decomposition, because post-quantum key generation and encapsulation display scaling behavior distinct from that of their pairing-based equivalents.

Multi-authority ABE poses a different kind of challenge. When multiple autonomous entities each issue partial keys, the system's overhead is spread among those entities, yet one standardized cost vector per scheme conceals the communication and coordination expenses borne by each entity, which dominate in practice [3]. Depicting such schemes would necessitate the evidence matrix to turn into a tensor indexed by authority as well as by stage, and the clustering procedure would consequently work on aggregated or per-authority projections. Hierarchical and decentralized variants introduce comparable complications, because their key-delegation stages recurse and their cost depends on the depth of the authority hierarchy. Such augmentations constitute inherent trajectories for later versions of the framework, and the structured evidence schema was devised with supplementary indices in view, so these schemes can be assimilated without reengineering the pipeline.

C. Operational Overhead, Security Trade-offs, and Governance of Automated Scheme Selection

Implementing the framework within an active permissioned blockchain-cloud architecture brings about operational expenses which the review alone does not bear. The normalization kernel must be re-executed whenever new primary studies appear or whenever the correction factors are recalibrated, and the clustering and decision-rule stages must then be repeated to refresh the class labels. In practice this is a batch process rather than a continuous one, and its cost is dominated by the manual extraction of stage-wise measurements from each new study, which remains the least automatable part of the pipeline. The feedback loop responsible for selecting a scheme is also not costless, since querying the evidence base, appraising the constrained objective, and transmitting the recommendation to the smart contract layer expend ledger transactions and consequently give rise to consensus latency. For deployments with frequent policy changes, the selection rule should be invoked only when the constraint set changes materially, rather than on every access request.

Security compromises attend the efficiency benefits the selection rule seeks. Minimizing normalized overhead subject to a security-level constraint $\sigma_i \geq \sigma_{\text{req}}$ treats security as a threshold rather than as a quantity to be maximized, and this is a deliberate modeling choice, since security assumptions are not commensurable with latency in any meaningful arithmetic sense. Nevertheless, the threshold formulation can mask qualitative differences between assumptions. A scheme whose security depends on a q-type assumption within the generic group model is not equivalent to one founded on the standard Decisional Bilinear Diffie-Hellman assumption, irrespective of both possessing the same ordinal security level, and the evidence base documents the assumption type precisely so that designers may apply stricter filters than the scalar threshold by itself. Revocation methods possess a comparable subtlety: direct revocation, indirect revocation, and ciphertext-policy revocation diverge in their trust assumptions and in their failure modes, not solely in their normalized cost, and the admissible set $\mathcal{R}_{\text{allowed}}$ should consequently be specified with regard to the threat model rather than to efficiency alone [13].

The oversight of automated scheme selection gives rise to issues of an organizational rather than a technical nature. When a framework endorses a scheme and the endorsement is executed by means of smart contract configuration, accountability for the outcomes of such a selection must be attributed to a discernible entity, and the audit record stored on the ledger ought to retain the evidence snapshot, the constraint set, and the version of the normalization parameters responsible for the decision. Documenting the correction factors is indispensable, since a recalibration altering class boundaries would otherwise quietly render prior selections invalid without any record. Moreover, the deployment suitability scores $\delta_i$ embed assessments of operational context open to dispute among stakeholders, and regarding them as objective inputs to an optimization conceals their origin. We consequently argue for the framework to serve as a decision-support instrument whose outputs human architects review, not as an autonomous controller, and for the evidence base, the parameter versions, and the rationale for every selection to be kept as first-class artifacts of the system's governance record.

## 7. CONCLUSION

This paper set out to address a persistent obstacle in the literature on ABE for blockchain-based cloud data management: the incomparability of reported overheads across studies conducted under heterogeneous hardware, libraries, security levels, policy depths, and consensus settings. The Benchmark-Normalized SLR Framework addresses this impediment by conceptualizing the ABE module as a benchmark-conscious, overhead-categorized, and evidence-standardized subsystem. Its Benchmark Normalization Kernel maps raw stage-wise measurements from thirty primary studies into a dimensionless space, its Overhead Pattern Classifier assigns each scheme to one of four reproducible overhead classes through Ward-linkage clustering and decision-rule refinement, and its PRISMA-Guided Evidence Synthesizer links normalized costs to security assumptions, revocation methods, and deployment suitability while closing the loop between synthesized evidence and design decisions.

The synthesis indicated the apparent disagreement among primary studies is largely attributable to measurement conditions rather than to cryptographic design, and the residual

variation organizes into pairing-bound, policy-size-bound, revocation-heavy, and communication-efficient patterns whose trade-offs are systematic rather than incidental. Sensitivity analysis substantiated the endurance of these classes under altered correction factors and substitute linkage criteria, while simultaneously bounding the framework's assertions: normalized values afford dependable ranking and classification yet ought not to be interpreted as exact latency forecasts. The framework's principal value consists in transforming heterogeneous literature evidence into a constraint-aware selection instrument, thereby permitting designers of permissioned blockchain-cloud stacks to shortlist ABE schemes on comparable grounds and substantiate their choices via targeted measurements. Extending the normalization kernel to post-quantum and multi-authority constructions, and governing automated scheme selection through versioned evidence and human review, define the most immediate directions for continued work.

## REFERENCES

1. J. Bethencourt, A. Sahai, and B. Waters, "Ciphertext-policy attribute-based encryption," in 2007 IEEE symposium on security and privacy (SP '07), 2007.
2. D. Boneh and M. Franklin, "Identity-based encryption from the weil pairing," in Annual international cryptology conference, 2001.
3. V. Goyal, O. Pandey, A. Sahai, and B. Waters, "Attribute-based encryption for fine-grained access control of encrypted data," Proceedings of the 13th ACM conference on Computer and communications security, 2006.
4. E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, et al., "Hyperledger fabric: A distributed operating system for permissioned blockchains," in Proceedings of the thirteenth eurosys conference, 2018.
5. J. Zhang, S. Zhong, T. Wang, H. Chao, et al., "Blockchain-based systems and applications: A survey," Journal of Internet Technology, 2020.
6. M. Ferrag and L. Shu, "The performance evaluation of blockchain-based security and privacy systems for the internet of things: A tutorial," IEEE Internet of Things Journal, 2021.
7. J Tetzlaff, D. Altman, t PRISMA Group, et al., "Preferred reporting items for systematic reviews and meta-analyses: The PRISMA statement," PLoS Medicine, 2009.
8. J. Gong and N. Navimipour, "An in-depth and systematic literature review on the blockchain-based approaches for cloud computing," Cluster Computing, 2022.
9. T. Saaty, "The analytic hierarchy process," academia.edu, 1980.
10. A. Fahad, N. Alshatri, Z. Tari, A. Alamri, et al., "A survey of clustering algorithms for big data: Taxonomy and empirical analysis," IEEE Transactions on Emerging Topics in Computing, 2014.
11. U. Tariq, F. Sabrina, M. Rashid, S. Gordon, et al., "Blockchain-based secured data sharing in healthcare: A systematic literature review," IEEE Access, 2025.
12. M. Green, S. Hohenberger, and B. Waters, "Outsourcing the decryption of ABE ciphertexts," in Usenix security symposium, 2011.
13. K. Yang, G. Wu, C. Dong, X. Fu, F. Li, et al., "Attribute based encryption with efficient revocation from lattices," International Journal of Network Security, 2020.
14. K. Singamaneni, K. Yadav, A. Aledaily, et al., "Decoding the future: Exploring and comparing ABE standards for cloud, IoT, blockchain security applications," Multimedia Tools and Applications, 2025.
15. L. Naveenkumar and S. Manjula, "Decentralized blockchain-based cloud storage with smart contracts and ABE for enhanced security," Algorithms for Intelligent Systems, 2025.
16. K. Sravanthi and P. Chandrasekhar, "An efficient multi-user groupwise integrity CP-ABE (GI-CPABE) for homogeneous and heterogeneous cloud blockchain transactions." Journal of Electrical Systems, 2024.
17. E. Haque, A. Shah, J. Iqbal, S. Ullah, R. Alroobaea, et al., "A scalable blockchain based framework for efficient IoT data management using lightweight consensus," Scientific Reports, 2024.
18. Z. Huang, "Blockchain performance optimization using cloud technology," figshare.utas.edu.au, 2024.
19. M. Abbasi, F. Cardoso, P. Váz, J. Silva, and P. Martins, "A practical performance benchmark of post-quantum cryptography across heterogeneous computing environments," Cryptography, 2025.
20. D. Dinev and G. Spasova, "A software platform for benchmarking, multi-criteria evaluation, and integrity validation of symmetric encryption algorithms," Journal of Cybersecurity and Privacy, 2026.
21. D. Commey, B. Appiah, G. Klogo, W. Bagyl-Bac, et al., "Performance analysis and deployment considerations of post-quantum cryptography for consumer electronics," Scientific Reports, 2026.
22. A. Punia, P. Gulia, N. Gill, E. Ibeke, C. Iwendi, et al., "A systematic review on blockchain-based access control systems in cloud environment," Journal of Cloud Computing, 2024.

## BIOGRAPHIES



Name - Mr.Prashant Sharma
Course - M.Tech.(CSE)
College - Buddha Institute of Technology,Gida, Gorakhpur