International Journal of Computer Applications (0975 - 8887)
Volume 187 - No.142, September 2026

Comparative Analysis of Blockchain Platforms for Decentralized Applications

Kyrylo Sotnykov
Independent Researcher
Tampa, FL, USA

ABSTRACT

Blockchain technology has experienced rapid growth and has led to the creation of many different blockchain platforms that enable developers to develop decentralized applications on them. Selecting an appropriate blockchain platform to build a real-world system on is difficult due to many variables that affect the choice of a platform, including architecture, scalability, transaction cost of using the platform, and developer ecosystem.
This paper reviews, compares, and analyzes the various blockchain and Distributed Ledger technology platforms available, namely Ethereum, Polygon, Arbitrum, Base, and Hedera, to aid developers in determining the appropriate platform for their project.
The platforms are evaluated based on multiple criteria, including transaction cost, throughput, finality, consensus mechanisms, and the accessibility for developers to work on them.
The analysis also provides contextualization for how the differences in architecture of platforms determine how feasible it will be to build scalable decentralized applications on these platforms, using the area of credential systems as a workload requiring secure, tamper-resistant record management and efficient transaction processing.
A weighted multi-criteria decision model is applied to examine platform suitability under high-volume credential, security-oriented, and developer-oriented application scenarios. The analysis demonstrates that platform rankings change substantially when the relative importance of transaction cost, throughput, finality, decentralization, and ecosystem maturity is modified.
The results of the study indicate that different blockchains provide different trade-offs in terms of decentralization, performance, and operational cost; therefore, this study will be beneficial to researchers, developers, and organizations as they seek to determine the best blockchain platform for decentralized systems.

Keywords

Blockchain Platforms, Distributed Ledger Technology, Decentralized Systems, Blockchain Scalability, Consensus Mechanisms, Credential Verification, Smart Contract Platforms, Decentralized Infrastructure

1. INTRODUCTION

The past ten years have seen an expansive evolution of blockchain technology from the original underpinning of cryptocurrency to an expansive platform for decentralized applications [12, 21]. A large number of today's blockchain networks offer a programmable environment for building distributed systems of various types, including but not limited to: financial services, digital identity services, and frameworks for data verification. Although many blockchain networks have been developed in the past several years, selecting the best blockchain platform for a specific application continues to present a challenge to both developers and organizations.
Different blockchain platforms have significant architectural differences [19, 5]. While some prioritize decentralization and security, others emphasize scalability, transaction speed, or usability by developers. As such, practical performance characteristics for different blockchain platforms may vary greatly from one another. Factors such as transaction costs, network throughput, latency for confirmations, and tooling support will all determine if a blockchain platform has the appropriate use case.
In recent years, the blockchain ecosystem has expanded beyond traditional Layer-1 blockchains to include Layer-2 scaling solutions and other distributed ledger (DLT) architectures that do not conform to the original Layer-1 model of development and operation. For example, Ethereum has been the original foundation for the Layer-1 model [20], while Arbitrum and Base support rollup-type scaling solutions for Layer-2 blockchains [11]. Other examples such as Polygon provide sidechain-type architectures [17]; they seek to increase performance and interoperability with existing Layer-1 blockchain ecosystems.
Additionally, distributed ledger platforms such as Hedera implement alternative consensus mechanisms designed to provide high throughput and predictable transaction costs [1, 4].
Academic research is lacking in the area of academic understanding about the differences between various forms of blockchains, since the diversity in type has continued to develop. The bulk of existing literature has been focused on theoretical analyses of individual protocols or general analyses of protocols that would work well for specific purposes, as opposed to looking at multiple platforms that provide similar capabilities and/or performance under similar applications.
This paper attempts to fill this void by comparing the various blockchains and distributed ledgers currently available in the marketplace that are commonly used. The evaluation will look strictly at core system characteristics that have an effect on the development of decentralized applications. Examples of these core system characteristics include cost efficiency, scalability, transaction finality, consensus model and maturity of the developer ecosystem.
In order to better frame the research being undertaken, this study uses credential management systems as a representative use case

from which the research work can be referenced. Credential management systems require secure and tamper-resistance recording of data, efficient processing of transactions and a method by which to ultimately scale up to a lot of users [18]. This work will help inform the reader how the various platforms evaluated will meet the aforementioned credential management system requirements and provide practical insight into the strengths and weaknesses of the various blockchain infrastructures examined.

—a comparative architectural analysis of five blockchain and distributed ledger platforms;

—a unified evaluation framework covering cost, throughput, finality, security, and developer accessibility;

—a multi-criteria decision model for application-specific platform selection;

—scenario-based analysis for high-volume, security-oriented, and developer-oriented decentralized applications;

—sensitivity and trade-off analysis demonstrating how platform rankings change under different operational priorities.

2. OVERVIEW OF BLOCKCHAIN PLATFORM ARCHITECTURES

This section outlines the major architectural classifications of contemporary blockchain platforms at a high level. The foundational understanding of these architectural models is necessary when evaluating how a diverse range of networks enable the implementation of decentralized applications and large-scale systems. For the purposes of this paper, the following types of blockchain platforms are analyzed in broad categories: Layer-1 blockchain platforms; layer 2 scaling platforms; sidechains; and alternate Distributed Ledger Technologies (DLTs).

2.1 Layer-1 Blockchains

Layer-1 blockchain platforms are the fundamental protocol layer that generates accordance, approves transactions and secures the network. Each layer-1 blockchain platform works independently of one another and has its own validator/nodes to process transactions on behalf of users. Layer-1 architectures have high efficiency and low costs associated with developing decentralized applications, but they may be limited in their capability to process transactions quickly due to scalability restrictions.
A good example of a layer 1 blockchain platform is Ethereum [20]. Ethereum uses ”proof of stake” (PoS) to generate consensus and enables the execution of programmer statements through smart contracts [20]. The majority of development activity on decentralized applications takes place on layer 1 examples due to their maturity and regulatory framework. However, high transaction fees and the limited capability to process transaction quickly may make layer 1 blockchains inefficient when used for applications that have frequent or significant transactional activity.

2.2 Layer-2 Scaling Solutions

Layer 2 solutions enhance scalability by processing transactions off the base blockchain, but still provides security guarantees through existing or prior Layer 1 based networks. This not only relieves congestion from the base layer but allows for much lower transaction costs.
Many Layer 2 solutions implement a rollup architecture where the individual transactions occur off the blockchain and are processed back to the main blockchain in a batch aggregation manner [11]. Optimistic rollups allow transactions to be assumed authentic unless contested, allowing for greater transaction throughput and decreasing costs. While Layer 2 networks greatly address scalability, they can add additional challenges such as bridging assets and maintaining synchronization with the base layer.

2.3 Sidechain Architectures

Sidechains are independent blockchains that run parallel to a primary blockchain. Sidechains usually have interfaces that enable interaction with one another — transferring either assets or data between chains. Sidechains operate as means of providing superior performance and lower transaction costs than the primary blockchain.
While Layer 2 solutions use the consensus mechanism and validator sets of the primary blockchain, sidechains have their own independent consensus mechanisms and validator sets. Consequently, although this provides greater scalability, sidechains introduce a different security assumption than the primary blockchain. Sidechain architectures are frequently used for high transaction throughput and minimal fees while being compatible with existing blockchain ecosystems [17].

2.4 Alternative Distributed Ledger Architectures

Beyond classic blockchains, various platforms use distinct architectures of distributed ledger technology (DLT) that do not necessarily follow a linear structure composed of consecutive blocks of data. These types of systems are often engineered using different consensus techniques (e.g., DAG) for the purpose of attaining high speeds and low latency finality of transactions.
Examples of these types of architectures are Hashgraph-based consensus systems [1]. Their method of communication employs a gossip model between nodes, combined with virtual voting to validate the order in which transactions occur within the DLT [4]. These types of architectures are designed to provide optimal levels of throughput with minimal transaction costs and to achieve rapid finality. These characteristics make alternative DLT systems particularly well suited to high-volume transactional needs (many transactions), where reliable event ordering is critical.
The examination of each of these various types of DLT architecture offers an ability to evaluate how well each model addresses their inherent trade-offs between decentralization, scalability, and operational effectiveness.

3. EVALUATION METHODOLOGY

This study applies a combined comparative, decision-analytical, and simulation-based methodology to evaluate the suitability of blockchain and distributed ledger platforms for decentralized applications. The methodological framework consists of three complementary components: (1) comparative architectural analysis, (2) multi-criteria decision analysis (MCDA), and (3) simulation-based experimental evaluation.
The comparative architectural analysis identifies the principal technical and operational characteristics of each platform and establishes a common basis for comparison. The MCDA subsequently converts these characteristics into normalized ordinal scores and examines how platform suitability changes when different application priorities are applied. Finally, the simulation-based evaluation complements the architectural and decision analysis by applying a consistent modeled transaction workload to all evaluated platforms under controlled workload and congestion conditions. This combi-

nation is intended to avoid relying on either a single performance metric or a purely qualitative architectural comparison.
The methodology is therefore structured so that Sections 4 and 5 establish the comparative evidence, Section 6 evaluates application-specific trade-offs through weighted decision analysis, and Section 7 examines selected operational characteristics through controlled simulation.

3.1 Platform Selection and Scope

Five platforms are included in the study: Ethereum, Polygon, Arbitrum, Base, and Hedera. These platforms were selected to provide representation of several important architectural approaches used for contemporary decentralized applications.
Ethereum represents a mature Layer-1 smart contract blockchain. Polygon represents a sidechain-based architecture with compatibility with the Ethereum ecosystem. Arbitrum and Base represent Ethereum Layer-2 rollup architectures in which execution is moved away from the base layer while maintaining a relationship with Ethereum for settlement and security. Hedera represents an alternative distributed ledger architecture based on hashgraph consensus rather than a conventional sequential blockchain structure.
The purpose of this selection is not to provide an exhaustive survey of all available blockchain networks. Instead, the five platforms provide a bounded comparison across different architectural models while maintaining a practical focus on platforms capable of supporting decentralized applications and transaction-based systems.
The scope of the comparison is limited primarily to technical and operational characteristics that directly influence application development and deployment. These include transaction cost, throughput, transaction finality, security and decentralization characteristics, and developer ecosystem maturity. Other potentially important deployment considerations, including regulatory compliance, organizational governance requirements, privacy requirements, infrastructure-provider availability, and long-term network sustainability, are recognized but are not treated as primary decision criteria in the comparative model.
Credential management is used as a representative application context because such systems require tamper-resistant recording, repeated transaction processing, reliable confirmation of state changes, and the ability to scale as the number of issued or verified credentials increases. The credential workload provides a practical reference point rather than restricting the conclusions exclusively to credential applications.

3.2 Data and Evidence Sources

The comparative evaluation is based on information obtained from official platform documentation, published platform specifications and whitepapers, and academic or technical literature cited throughout the study. Official documentation is used primarily to characterize platform-specific properties such as consensus mechanisms, fee structures, transaction-processing models, finality behavior, developer tools, and reported network capabilities. Academic and technical sources provide additional context regarding blockchain architecture, scalability, consensus, security, and distributed ledger design.
The methodology distinguishes between documented platform characteristics and modeled assumptions. Values reported directly by platform documentation or technical sources are treated as descriptive evidence of network design or reported capabilities. Values introduced for the simulation, including representative transaction costs, effective throughput values, congestion multipliers, and modeled confirmation delays, are treated as controlled simulation inputs rather than direct measurements of contemporaneous mainnet performance.
This distinction is important because publicly reported platform metrics are not necessarily collected under identical workloads, network conditions, transaction types, or time periods. Consequently, documented characteristics support the architectural and comparative analysis, whereas the simulation standardizes selected operational parameters to enable comparison under equivalent modeled conditions.

3.3 Evaluation Criteria

All five platforms are evaluated using a common set of five criteria: transaction cost efficiency, transaction throughput, transaction finality, security and decentralization, and developer ecosystem maturity. The criteria were selected because together they capture major technical and operational trade-offs that influence whether a decentralized application can be deployed and operated effectively.
Transaction cost efficiency represents the financial cost associated with recording and processing transactions. This criterion is particularly important for applications that generate large numbers of relatively low-value transactions because small differences in per-transaction cost can become substantial at scale. The comparison focuses primarily on relative cost characteristics rather than treating a single observed transaction price as permanently representative, because network fees may change with congestion, protocol updates, and market conditions.
Transaction throughput represents the platform's ability to process transaction demand over time. Throughput is relevant because applications with increasing user or event volumes require sufficient processing capacity to avoid excessive delays or congestion. Reported or representative throughput values are interpreted in the context of each platform's architecture rather than as perfectly equivalent measurements across different execution models.
Transaction finality represents the time and mechanism through which a transaction becomes sufficiently confirmed or irreversible for application use. The evaluation considers both the approximate delay associated with confirmation or finality and the type of assurance provided. This distinction is necessary because deterministic finality, probabilistic finality, Layer-2 execution confirmation, and Layer-1 settlement do not represent identical stages of transaction processing.
Security and decentralization represent the trust and validation characteristics of the platform. This criterion considers factors such as the consensus model, validator participation, dependence on additional infrastructure or execution components, and the relative distribution of network validation or governance. The criterion is evaluated comparatively because different architectures provide different security assumptions and cannot always be reduced to a single directly measurable variable.
Developer ecosystem maturity represents the practical accessibility of a platform for application development and maintenance. The assessment considers documentation, software development kits, frameworks, libraries, tooling compatibility, community support, and overall ecosystem maturity. This criterion is included because platform suitability depends not only on transaction performance but also on the engineering effort required to develop, integrate, deploy, and maintain production applications.

3.4 Ordinal Normalization and Comparative Scoring

To enable the criteria to be evaluated within a common decision framework, the comparative evidence is translated into a five-point ordinal scoring scale used in the MCDA presented in Section 6.

Each platform receives a score from 1 to 5 for each evaluation criterion.
A score of 1 represents relatively weak suitability among the platforms examined for the corresponding criterion, while a score of 5 represents relatively strong suitability. Intermediate values represent progressively stronger comparative positions. The scale is ordinal and comparative; therefore, the difference between scores of 4 and 5 should not be interpreted as an exact quantitative interval equivalent to the difference between scores of 2 and 3.
Scores are assigned by comparing the architectural and operational evidence summarized in Sections 4 and 5. Quantitative characteristics, such as transaction cost and throughput, provide direct comparative evidence where representative values are available. More qualitative characteristics, particularly security, decentralization, and ecosystem maturity, require interpretation of factors such as consensus architecture, validator participation, tooling availability, documentation, ecosystem maturity, and operational dependencies.
Accordingly, the scoring process includes analytical judgment and should not be interpreted as a purely mechanical measurement procedure. Its purpose is to convert heterogeneous technical characteristics into a common ordinal representation that can be used to examine decision trade-offs. The complete scores applied in the analysis are reported in Section 6.

3.5 Scenario-Based Weighting and Multi-Criteria Decision Analysis

A single universal ranking is not used because different decentralized applications assign different levels of importance to cost, performance, security, and development accessibility. Instead, Section 6 applies different criterion weights under three representative application scenarios.
The **high-volume credential scenario** assigns greater importance to transaction cost, throughput, and finality because issuing, updating, revoking, and verifying large numbers of credentials can generate substantial transaction volume. Under this type of workload, relatively small differences in transaction cost or processing capacity may have a significant effect on operational feasibility.
The **security-oriented scenario** assigns greater importance to security and decentralization. This scenario represents applications in which trust assumptions, validator participation, consensus characteristics, and the security of high-value records may be more important than minimizing transaction cost.
The **developer-oriented scenario** assigns greater importance to ecosystem maturity and developer accessibility. This reflects projects in which implementation time, availability of libraries and frameworks, documentation quality, community support, and compatibility with established development workflows substantially influence platform selection.
The precise criterion weights and resulting weighted suitability scores are presented in Section 6. Scenario-based weighting is used because it allows the same underlying platform evidence to be examined under different operational priorities. This approach provides a more informative comparison than a single fixed ranking and makes explicit the trade-offs that cause a platform to become more or less suitable as application requirements change.

3.6 Simulation-Based Experimental Evaluation

The simulation-based evaluation presented in Section 7 provides a quantitative complement to the comparative and multi-criteria analyses. Whereas the architectural analysis examines documented platform characteristics and the MCDA evaluates their relative importance, the simulation applies a standardized credential-management workload to all five platforms using the same workload structure and defined network-condition scenarios.
The simulation examines transaction cost, effective throughput, processing and completion time, finality delay, and cost stability as workload size and modeled congestion change. It does not attempt to reproduce every internal mechanism of the evaluated networks or represent direct contemporaneous mainnet benchmarking. Instead, representative inputs derived from documented platform characteristics are used to create a controlled cross-platform comparison.
Detailed workload distributions, baseline input values, congestion multipliers, equations, experimental objectives, and evaluation metrics are provided in Section 7 and are therefore not repeated in this section.

3.7 Reproducibility and Methodological Limitations

Several methodological controls are used to improve consistency and reproducibility. The same evaluation criteria are applied to all platforms, the MCDA uses a common ordinal scoring scale, and the simulation applies an equivalent workload structure and explicitly defined congestion conditions across the five platforms.
For the simulation, representative baseline parameters are held constant within the defined scenarios and are modified only through the specified congestion assumptions. Each platform, workload-size, and congestion-level combination is evaluated through repeated simulation runs, and a fixed random seed is used to support repeatability. The detailed experimental configuration is reported in Section 7.
Nevertheless, the methodology has limitations. Blockchain transaction fees, network utilization, effective throughput, protocol behavior, and ecosystem maturity can change over time. Reported throughput values may also describe theoretical capacity, observed operation, or architecture-specific transaction definitions and therefore should not always be interpreted as directly equivalent measurements.
Similarly, the ordinal MCDA scores summarize heterogeneous quantitative and qualitative evidence and consequently involve analytical judgment, particularly for security, decentralization, and ecosystem maturity. The resulting scores and rankings should therefore be interpreted as comparative decision-support indicators rather than absolute measurements of platform quality.
Finally, the simulation relies on representative parameters and explicitly defined assumptions rather than simultaneous live-network benchmarking of all five platforms. The experimental results should therefore be interpreted as controlled estimates of relative operational behavior under the modeled conditions. This separation between documented characteristics, comparative analytical judgments, and modeled simulation assumptions is maintained throughout the study to provide a transparent basis for interpreting the results.

4. PLATFORM ANALYSIS

This part of the study presents a summary of the different blockchain and distributed ledger methods that were examined. Each of the various platforms highlights a diverse way to approach development of a decentralized infrastructure i.e. Layer-1, Layer-2 scaling solutions, sidechains, and an array of other distributed ledger technologies. The information included in this section aims to provide the reader with an overview of each of the platforms so that when performing a comparative evaluation at a later point in

time all relevant aspects are available for consideration including; technical features, consensus algorithms, and ecosystem maturity.

4.1 Ethereum

Ethereum is one of the most popular blockchain platforms available for creating decentralized applications and operates as a Layer-1 network that provides developers with programmable capabilities via smart contracts [20]. At the beginning of 2020, Ethereum made an important transition from using a proof-of-work (PoW) consensus algorithm to proof-of-stake (PoS), which has greatly decreased its overall power consumption as well as maintained the creation and functionality of a decentralized pool of validators.
Ethereum's overall ecosystem has matured significantly over time which makes it a relatively strong platform for building decentralized applications. It (Ethereum) supports widely used token standards such as ERC-20 and ERC-721 [20], has numerous tools available for developers, and has strong community engagement. All these factors combined make Ethereum an excellent base for developing decentralized applications as many developers rely on this platform often.
Despite having many advantages, Ethereum has a few challenges it must overcome to remain as one of the leaders in the industry. Due to the overwhelming demand for transaction processing on its blockchain, Ethereum tends to have high transaction fees. Furthermore, as Ethereum continues to grow and expand in popularity many of the newer blockchain architectures offer far greater total throughput than Ethereum's current base layer. In turn, multiple projects are building additional Layer-2 scaling options on top of Ethereum to help alleviate some of these issues.

4.2 Polygon

Polygon is an innovative blockchain that enhances the usability and scale of Ethereum-compatible applications through the use of a sidechain based solution [17] to ensure interoperability with the Ethereum network while providing very affordable transaction costs and faster confirmation times than the Ethereum mainnet.
Polygon utilizes a Proof-of-Stake (PoS) consensus mechanism and allows developers to build and run all of their Ethereum-based applications within the same smart contract environment provided by Ethereum. When building decentralized applications(s) with Polygon, developers have the ability to deploy their applications directly onto the Polygon network with very little changes to their existing code, thus providing them with many performance benefits as well.
As far as scalability is concerned, Polygon is highly advantageous to developers; however, because Polygon is built on top of a sidechain architecture, it has different security assumptions compared to Ethereum since validation on the Polygon network is performed by an independent group of validators (validator set).

4.3 Arbitrum

Arbitrum is a Layer 2 scaling solution to improve the scalability of Ethereum via optimistic rollup technology [11], by processing transactions as optimistically assumed valid through execution off of the Layer 1 blockchain (Ethereum), with the upon completion being submitted for periodic batch inclusion as an alternative means to get cheaper transactions than completing via Layer 1 (the cost of completing transactions via Layer 1) while having security from Layer 1, as well as provide a higher throughput than what is available through Layer 1.
This optimistic nature of arbitrage means that, unless challenged, an optimistically assumed valid transaction will not be challenged unless it's during a dispute period, which would allow for higher overall throughput than would be achieved through using Layer 1 alone or as well as providing lower overall transaction fees than would be available through Layer 1.
This Layer 2 architecture has created a new level of complexity in terms of bridging mechanisms between Layer 2 and Layer 1 networks, including time delays due to potential dispute resolution time periods.

4.4 Base

Base is an Ethereum Layer-2 solution and a scalable and decentralized application network. Base will facilitate off-chain transaction processing by utilizing rollup technology and committing a portion of results back to Layer 1 (the Ethereum mainnet).
Base has been designed with developer experience and integration with current Ethereum tools in mind so that the desired application can be developed using the tools the developer is already familiar with, thus helping the developer to create scalable and decentralized applications more easily.
Like most Layer-2 architectures, Base has some dependencies on the base Layer-1 chain, and therefore the movement of assets between both networks will require a bridge.

4.5 Hedera

Hedera is a distributed ledger solution that implements a consensus mechanism based upon hashgraph algorithms rather than using conventional blockchain structures [1]. The implementation utilizes a Direct Acyclic Graph (DAG) and gossip protocols to establish consensus to order transactions [4].
This structure has the ability to provide a high-volume throughput along with rapid finalization of transactions. Hedera's goal is to support businesses with predictable transaction costs and consistent operational characteristics at an enterprise level.
In addition to providing standard transaction services, the Hedera platform includes built-in capabilities to support tokenization of assets and to create logs of events occurring on the network. Each of these capabilities allows for the ability to develop business applications that generate digital assets and log events occurring in an ordered fashion through the network. The design of the Hedera platform is focused upon efficiency of operational capacity/performance along with the administration of a network of distributed governing bodies that have authority over the operating integrity of the network.

5. COMPARATIVE ANALYSIS

This section includes a comparative assessment of the blockchain and distributed ledger systems, as previously discussed in the preceding analyses, with a focus on the most important features that will affect whether decentralized applications can be built at scale. The major characteristics that affect scalability are the architecture of the system, the transaction throughput of the system, the cost of the transactions to be executed on the system (i.e., the gas cost), the finality of the transactions processed by the system, and the degree of maturity of the developer ecosystem associated with the platform.
This comparison is not an attempt to determine which is the best platform overall, but should illustrate how the different technological solutions have advantages and disadvantages compared with one another, and what types of applications may be more suited to different types of platforms based on how the platforms operate and the specific requirements of each type of application.

5.1 Architectural Comparison

The architectural model of a blockchain platform will define how the platform operates and how it will perform, and will influence the scalability of the platform, along with its ability to perform transactions and achieve decentralization (Table 1).
Typically, layer-one blockchain platforms will prioritize decentralization and security when designing their platforms, while layer-two platforms based on sidechains will improve scalability by offloading some or all of the computation from the base layer blockchain. Many sidechain platforms operate independently of the base layer blockchain and may have a different validator set than the base layer blockchain. Other types of distributed ledger technologies, such as the hashgraph algorithm, typically focus on performance and have predictable operational costs.

5.2 Transaction Throughput and Scalability

Transaction throughput is also one of the most significant factors utilized to determine the suitability of a blockchain platform for high volume applications (Table 2). Practically, transaction throughput is defined as the maximum number of transactions (as defined by the platform's transaction execution model) that can be executed in a given time frame (e.g., per minute).
Layer-2 solutions significantly increase throughput compared to base-layer networks by batching transactions before committing them to the main chain. Alternative distributed ledger models can achieve higher throughput by using different consensus mechanisms and network communication models.

5.3 Transaction Cost Comparison

Transaction cost is a critical factor for systems that generate frequent events or transactions. High operational cost can significantly limit the feasibility of decentralized applications (Table 3).
Transactions generated on a layer, or L1, may have variable transaction fees that fluctuate based on the demand of the network. Conversely, layer two networks may charge far less in transaction fees since they aggregate transactions into the layer below prior to committing them to the main network, while certain types of distributed ledger-based architectures establish predictable pricing for transactions.

5.4 Transaction Finality

The transaction finality of a specific transaction from the time of submission until that transaction is irreversible is another essential aspect of many applications that require an immediate confirmation of an event.
The consensus model that a specific application uses will determine the type of finality that exists once the consensus has been met (Table 4). If deterministic finality is utilized to determine whether consensus occurs, an event will receive immediate assurance once the consensus is met. If a probabilistic model is used for consensus, an event may only receive increasing confirmations regarding the event as new blocks are added to the network.

5.5 Developer Ecosystem Comparison

The maturity of the developer ecosystem supporting a given platform impacts the amount of access developers have to develop software and contribute to the overall growth of an application's developer ecosystem (Table 5). A platform that provides significant volumes of written documentation, tools, and community support to help developers complete development cycles and onboard to the platform will have greater overall adoption.
Ethereum currently maintains the largest developer ecosystem and the most mature set of development frameworks. Layer-2 networks benefit from compatibility with Ethereum tooling, while alternative distributed ledger platforms often provide specialized SDK-based development environments.

5.6 Summary of Platform Trade-offs

Based on a comparative review, the various platforms selected their design objectives differently. Layer-1 blockchains center on decentralising the network and establishing a mature ecosystem; whereas, Layer-2 solutions focus on increasing scalability and reducing costs. Depending on the independence of validator networks, sidechains provide high throughput; however, sidechains do not share the same characteristics of multi-chain solutions.
Therefore, whether a specific platform is the best suited for a particular Target Application depends heavily upon that application's requirements. Target Applications needing high throughput and predictable transaction costs may be best suited for an alternative distributed ledger architecture. In contrast, target applications that value ecosystem maturity and interoperability may be best suited to Ethereum-based solutions.

6. MULTI-CRITERIA DECISION ANALYSIS

The preceding comparison demonstrates that blockchain platforms differ substantially across transaction cost, throughput, finality, security assumptions, and ecosystem maturity. However, examining these criteria independently does not provide a complete basis for platform selection. A platform that performs strongly according to one criterion may perform poorly according to another, and the relative importance of each criterion depends on the operational requirements of the target application.
To provide a more detailed analysis, this section applies a multi-criteria decision analysis to the five evaluated platforms. The objective is not to identify a universally superior platform, but to examine how platform suitability changes when different application priorities are considered.

6.1 Analytical Model

Five decision criteria were selected based on the evaluation framework presented in Section 3:

(1) transaction cost efficiency;
(2) transaction throughput;
(3) transaction finality;
(4) security and decentralization;
(5) developer ecosystem maturity.

Each platform was assigned a normalized ordinal score from 1 to 5 for every criterion, where 1 represents relatively weak suitability and 5 represents relatively strong suitability. The scores were derived from the architectural and operational characteristics presented in the preceding comparative analysis.
The overall suitability score for each platform was calculated using Equation 1.

$$S_p = \sum_{i=1}^{n} w_i r_{p,i} \tag{1}$$

Table 1. Architectural Comparison

Platform	Architecture Type	Consensus Mechanism	Smart Contract Support
Ethereum	Layer-1 Blockchain	Proof-of-Stake	Yes
Polygon	Sidechain	Proof-of-Stake	Yes
Arbitrum	Layer-2 Rollup	Optimistic Rollup	Yes
Base	Layer-2 Rollup	Optimistic Rollup	Yes
Hedera	Distributed Ledger (DAG-based)	Hashgraph	Limited / Service-based

Table 2. Transaction Throughput and Scalability

Platform	Approximate TPS	Scalability Model
Ethereum	15–30	Base Layer Scalings
Polygon	1000+	Sidechain Scaling
Arbitrum	2000+	Rollup Aggregation
Base	2000+	Rollup Aggregation
Hedera	10,000+	DAG-based Parallel Processing

Table 3. Transaction Cost Comparison

Platform	Average Transaction Cost	Cost Stability
Ethereum	$1–$20+ depending on congestion	Variable
Polygon	<$0.01	Moderate
Arbitrum	∼$0.10–$0.50	Moderate
Base	∼$0.05–$0.30	Moderate
Hedera	∼$0.0001	Predictable

Table 4. Transaction Finality

Platform	Finality Type	Approximate Finality Time
Ethereum	Probabilistic	∼12 seconds to several minutes
Polygon	Probabilistic	Few seconds
Arbitrum	Probabilistic	Minutes depending on dispute window
Base	Rollup-based	Minutes depending on dispute window
Hedera	Deterministic	3–5 seconds

where S_p represents the total suitability score of platform p, w_i represents the weight assigned to criterion i, and $r_{p,i}$ represents the normalized score of platform p for that criterion. The sum of all criterion weights is equal to 1, as shown in Equation 2.

$$\sum_{i=1}^{n} w_i = 1 \tag{2}$$

This weighted model reflects the fact that platform selection depends on application requirements rather than on a single performance characteristic.
Table 6 presents the normalized platform scores used throughout the decision analysis.
The normalized scores indicate that no platform dominates all evaluation criteria. Ethereum receives the highest scores for security, decentralization, and ecosystem maturity, but performs poorly in transaction cost and base-layer throughput. Hedera receives the highest performance-related scores but has a less mature general-purpose smart contract ecosystem. Polygon provides a comparatively balanced profile, while Arbitrum and Base offer Ethereum compatibility with improved transaction cost and throughput characteristics.

6.2 High-Volume Credential Scenario

The first decision scenario represents a high-volume credential management system. Such a system may issue, update, revoke, and verify large numbers of credentials. Consequently, transaction cost, throughput, and finality have greater operational importance than general ecosystem maturity.
The following criterion weights were applied:

—transaction cost efficiency: 30%;
—transaction throughput: 20%;
—transaction finality: 20%;
—security and decentralization: 20%;
—developer ecosystem maturity: 10%.

Table 7 presents the weighted suitability scores for the high-volume credential scenario.
The results show that Hedera receives the highest suitability score for a high-volume credential workload. Its strong result is primarily attributable to predictable transaction pricing, high throughput, and rapid deterministic finality. These characteristics are particularly important when every credential issuance, update, or revocation creates a ledger transaction.
Polygon ranks second because it combines low transaction costs, comparatively high throughput, and compatibility with Ethereum development tools. Although Polygon does not provide the same security assumptions as the Ethereum base layer, its balanced characteristics make it suitable for applications that require both scalability and compatibility with established smart contract environments.
Ethereum receives the lowest score in this scenario. This result does not indicate that Ethereum is technically incapable of supporting credential management. Instead, it demonstrates that direct execution on the Ethereum base layer may become economically inefficient when the application generates a large number of low-value transactions.
Arbitrum and Base occupy an intermediate position. Both platforms significantly improve transaction cost and throughput compared with Ethereum while retaining Ethereum Virtual Machine compatibility. However, their rollup architectures introduce additional operational considerations, including sequencer dependencies, bridge management, and differences between Layer-2 confirmation and Layer-1 settlement.

6.3 Security-Oriented Application Scenario

The second scenario represents an application in which decentralization, security, and ecosystem maturity are more important than minimizing transaction cost. Examples may include high-value ownership records, regulated digital assets, or systems in which broad validator participation is a primary requirement.
The following criterion weights were applied:

—transaction cost efficiency: 10%;
—transaction throughput: 10%;
—transaction finality: 20%;
—security and decentralization: 40%;

Table 5. Developer Ecosystem Comparison

Platform	Ecosystem Maturity	Developer Tooling	Community Size
Ethereum	Very High	Extensive	Very Large
Polygon	High	Extensive	Large
Arbitrum	High	Extensive	Large
Base	Growing	Good	Growing
Hedera	Moderate	SDK-based	Moderate

Table 6. Normalized platform scores used in the multi-criteria decision analysis

Platform	Cost Efficiency	Throughput	Finality	Security and Decentralization	Developer Ecosystem
Ethereum	1	1	3	5	5
Polygon	5	4	4	3	4
Arbitrum	3	4	2	4	4
Base	4	4	2	3	3
Hedera	5	5	5	3	3

Table 7. Weighted results for a high-volume credential system

Platform	Weighted Suitability Score	Rank
Hedera	4.40	1
Polygon	4.10	2
Arbitrum	3.30	3
Base	3.30	3
Ethereum	2.60	5

Table 8. Weighted results for a security-oriented decentralized application

Platform	Weighted Suitability Score	Rank
Ethereum	3.80	1
Hedera	3.80	1
Polygon	3.70	3
Arbitrum	3.50	4
Base	3.00	5

—developer ecosystem maturity: 20%.

Table 8 presents the weighted results for the security-oriented scenario.
The ranking changes substantially when security and decentralization receive the greatest weight. Ethereum moves from the lowest position in the high-volume scenario to the highest-ranked group. Its broad validator participation, mature protocol, extensive operational history, and established developer ecosystem compensate for its cost and throughput limitations.
Hedera also receives a high score because of its deterministic finality and strong consensus characteristics. However, the interpretation of security differs between Ethereum and Hedera. Ethereum emphasizes permissionless validator participation and economic security, whereas Hedera combines hashgraph consensus with a council-based governance structure. Therefore, their equal weighted scores do not imply equivalent governance or decentralization models.
Polygon remains competitive because of its performance and Ethereum compatibility, but its independent validator set introduces a different trust assumption from that of the Ethereum base layer. Arbitrum benefits from Ethereum settlement but depends on additional Layer-2 components. Base receives a lower result because its ecosystem and decentralization characteristics are less mature relative to the more established platforms considered in this study.

Table 9. Weighted results for a developer-oriented decentralized application

Platform	Weighted Suitability Score	Rank
Polygon	3.95	1
Ethereum	3.80	2
Hedera	3.70	3
Arbitrum	3.65	4
Base	3.15	5

6.4 Developer-Accessibility Scenario

The third scenario represents a development team that prioritizes rapid implementation, availability of libraries, smart contract frameworks, documentation, and community support.
The following criterion weights were applied:

—transaction cost efficiency: 15%;
—transaction throughput: 10%;
—transaction finality: 10%;
—security and decentralization: 20%;
—developer ecosystem maturity: 45%.

Table 9 presents the weighted suitability scores for the developer-oriented scenario.
Polygon obtains the highest score in this scenario because it provides relatively low operational costs while remaining compatible with Ethereum development tools. Developers can use established programming languages, libraries, wallets, testing frameworks, and smart contract deployment workflows.
Ethereum ranks second because it provides the most mature developer ecosystem, but its lower transaction efficiency reduces its overall weighted result. It may remain the preferred platform for prototypes, complex decentralized finance applications, and systems requiring access to established standards and protocol integrations.
Arbitrum also benefits from Ethereum Virtual Machine compatibility. However, developers must account for Layer-2-specific operational behavior, including bridging, withdrawal periods, sequencer operation, and cross-layer messaging.
Hedera provides software development kits and native network services that can simplify specific enterprise use cases, but its development model differs from the dominant Ethereum-based environment. This difference may create an additional learning requirement for teams already experienced with Solidity and Ethereum tooling.

Table 10. Platform rankings across application scenarios

Platform	High-Volume Credential System	Security-Oriented System	Developer-Oriented System
Ethereum	5	1	2
Polygon	2	3	1
Arbitrum	3	4	4
Base	3	5	5
Hedera	1	1	3

Base provides familiar Ethereum-compatible tools and may become increasingly attractive as its ecosystem matures. However, according to the relative maturity values used in this study, it currently receives a lower developer-accessibility score than more established networks.

6.5 Cross-Scenario Analysis

Figure 1 compares the weighted suitability scores of the five platforms across the three application scenarios and illustrates how changes in application priorities affect their relative positions.
The visualization shows that Ethereum exhibits the largest increase between the high-volume and security-oriented scenarios, whereas Hedera performs most strongly when transaction cost, throughput, and finality receive greater weight. Polygon remains comparatively stable across all three scenarios, supporting its characterization as the most balanced platform in the cross-scenario analysis.
Table 10 summarizes the rankings obtained under the three application profiles.
The cross-scenario results demonstrate that platform suitability is highly sensitive to application priorities. Ethereum changes from fifth place in the high-volume scenario to first place in the security-oriented scenario. This represents the clearest example of why an isolated comparison based only on transaction throughput or transaction cost is insufficient.
Polygon produces the most consistently balanced result. It ranks within the top three in every evaluated scenario. This consistency results from its combination of low fees, high throughput, Ethereum compatibility, and an established development ecosystem. However, its independent consensus model must still be evaluated against the security requirements of the intended application.
Hedera performs particularly strongly when cost, throughput, and finality are prioritized. Its lower ecosystem score does not substantially affect its result in high-volume applications, but it becomes more important in developer-oriented scenarios.
Arbitrum and Base occupy intermediate positions. Their principal advantage is the ability to retain Ethereum compatibility while reducing transaction costs and increasing throughput. Their suitability is therefore strongest for applications that require Ethereum tooling but cannot economically operate directly on the Ethereum base layer.

6.6 Trade-Off Analysis

The results reveal four major trade-offs that should be considered when selecting a blockchain platform.

6.6.1 Cost versus Decentralization. Networks with the lowest transaction costs do not necessarily provide the same validator participation or security assumptions as the Ethereum base layer. Lower fees may be achieved through independent validator sets, rollup execution, alternative governance structures, or different consensus mechanisms.
Therefore, transaction cost should not be evaluated independently. An application must determine whether its trust model permits reliance on a sidechain, a Layer-2 sequencer, or a council-governed distributed ledger.

6.6.2 Throughput versus Architectural Complexity. Layer-2 platforms increase transaction capacity by moving execution away from the base layer. However, this benefit introduces architectural components that do not exist in a conventional Layer-1 deployment. These components may include:

—sequencers;
—bridge contracts;
—fraud-proof or dispute mechanisms;
—cross-layer message delays;
—separate execution and settlement states.

Consequently, higher throughput may reduce transaction costs while increasing system integration and operational complexity.

6.6.3 Finality versus Settlement. The finality values presented in platform documentation may refer to different stages of transaction processing. A transaction may receive rapid execution confirmation on a Layer-2 network but require additional time before it is fully settled on the underlying Layer-1 blockchain.
Similarly, probabilistic and deterministic finality should not be treated as directly equivalent. Deterministic finality provides a defined point after which a transaction cannot be reversed under the normal assumptions of the protocol. Probabilistic finality becomes stronger as additional blocks or confirmations are added.
Applications must therefore define which finality event is operationally significant:

—local execution confirmation;
—Layer-2 transaction inclusion;
—Layer-1 settlement;
—irreversible consensus ordering;
—application-level acceptance.

6.6.4 Ecosystem Maturity versus Specialized Performance. Ethereum-compatible platforms benefit from extensive tooling, standards, wallet support, smart contract libraries, and developer experience. Alternative distributed ledger platforms may provide superior performance for specialized event-processing workloads but may require platform-specific integration and development knowledge.
Organizations must therefore consider not only network fees but also engineering cost. A platform with a slightly higher transaction cost may still produce a lower total implementation cost if the development team can use existing libraries, frameworks, and operational tools.

6.7 Sensitivity Analysis

A sensitivity analysis was conducted to determine whether the platform rankings remained stable when the criterion weights changed. The results show that Hedera remained the preferred platform when the combined weight of transaction cost, throughput, and finality exceeded approximately 60% of the total decision weight. Polygon remained competitive across a wider range of balanced weighting configurations because it did not receive the lowest score in any major technical criterion.
Ethereum became the preferred platform when security, decentralization, and developer ecosystem maturity jointly accounted for the

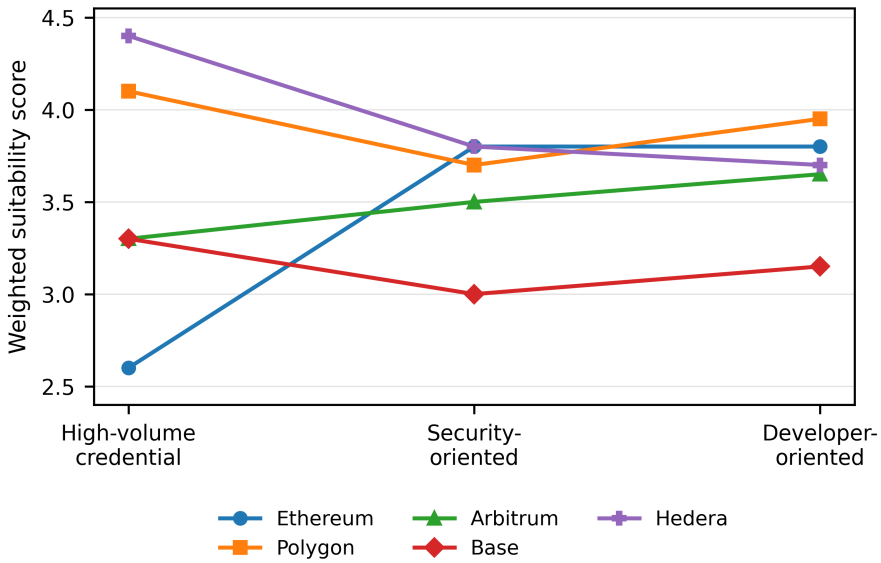

Fig. 1. Weighted platform suitability scores across the high-volume credential, security-oriented, and developer-oriented application scenarios.

majority of the total decision weight. This indicates that Ethereum's suitability depends less on transactional efficiency and more on its security model, operational history, and ecosystem network effects. Arbitrum and Base became more competitive when Ethereum compatibility was treated as a mandatory requirement and direct Layer-1 execution was excluded because of cost. In this restricted decision space, the comparison changes from selecting among fundamentally different architectures to selecting among Ethereum-compatible scaling environments.

The sensitivity analysis confirms that the rankings are not absolute. Small changes in criterion weights may alter the order of platforms with similar scores, particularly Polygon, Arbitrum, and Base. In contrast, larger architectural differences produce more stable conclusions. Ethereum consistently performs poorly when transaction cost and throughput dominate, while Hedera consistently performs strongly under high-volume transactional requirements.

6.8 Platform-Selection Guidelines

Based on the multi-criteria analysis, the following platform-selection guidelines can be derived:

—**Ethereum** is most appropriate when decentralization, security history, smart contract composability, and ecosystem maturity are more important than transaction cost.

—**Polygon** is suitable when an application requires low fees, relatively high throughput, and compatibility with Ethereum development tools.

—**Arbitrum** is suitable when Ethereum settlement and tooling compatibility are required together with lower execution costs.

—**Base** is appropriate for applications seeking Ethereum-compatible Layer-2 execution and integration with a growing ecosystem, subject to an evaluation of its operational dependencies and maturity.

—**Hedera** is suitable for high-volume event recording, credential issuance, and applications requiring predictable fees and rapid deterministic finality.

These guidelines should not be interpreted as universal recommendations. Platform selection must also account for governance, legal requirements, availability of infrastructure providers, data privacy, interoperability, and the long-term sustainability of the network.

6.9 Analytical Discussion

The detailed analysis demonstrates that the comparison of blockchain platforms cannot be reduced to a ranking based exclusively on transactions per second or average transaction fees. Platform suitability is multidimensional and depends on how architectural properties influence the complete operational model of an application.

For credential management, Hedera and Polygon provide strong operational characteristics because credential issuance, update, and revocation events may generate high transaction volumes. However, Ethereum may remain preferable when credentials interact with an established smart contract ecosystem or when permissionless validation is treated as a primary requirement.

Layer-2 platforms provide a compromise between Ethereum compatibility and transaction efficiency. Nevertheless, they do not eliminate architectural trade-offs. Instead, they move part of the complexity from transaction execution to sequencing, bridging, settlement, and cross-layer coordination.

The analysis also demonstrates that platform rankings vary significantly under different weighting configurations. Hedera produces the strongest result for high-volume applications, Polygon provides the most balanced cross-scenario performance, and Ethereum becomes preferable when security, decentralization, and ecosystem maturity receive greater importance.

Therefore, blockchain selection should be performed as an application-specific decision process. A platform should be selected only after defining the required trust model, expected transaction volume, acceptable finality delay, development constraints, interoperability requirements, and long-term operating budget.

7. SIMULATION-BASED EXPERIMENTAL EVALUATION

The multi-criteria decision analysis presented in the previous section evaluates platform suitability according to architectural and operational characteristics. To complement that analysis with quantitative results, this section presents a controlled simulation-based evaluation of Ethereum, Polygon, Arbitrum, Base, and Hedera under a representative decentralized credential management workload.

The evaluation was designed to examine how platform characteristics affect transaction cost, processing capacity, workload completion time, finality delay, and cost stability as transaction volume and network congestion increase. The simulation does not attempt to reproduce every internal network mechanism. Instead, it applies a consistent workload model to all five platforms so that their relative operational behavior can be compared under equivalent conditions.

7.1 Experimental Objectives

The simulation-based evaluation was designed to answer the following research questions:

—**RQ1:** How does transaction volume affect the operational cost of each platform?

—**RQ2:** How effectively can each platform process low-, medium-, and large-scale credential workloads?

—**RQ3:** How does network congestion influence transaction cost and effective throughput?

—**RQ4:** Which platforms provide the lowest workload completion time for high-volume decentralized applications?

—**RQ5:** How stable are platform costs across repeated simulations with changing network conditions?

7.2 Workload Model

The evaluated application represents a decentralized credential management system in which institutions issue, update, revoke, and verify digital credentials. Each credential operation generates one ledger transaction.

Three workload sizes were evaluated:

—**Small workload:** 10,000 credential transactions;

—**Medium workload:** 100,000 credential transactions;

—**Large workload:** 1,000,000 credential transactions.

The workload consisted of four transaction categories:

—credential issuance: 25%;

—credential verification: 60%;

—credential update: 10%;

—credential revocation: 5%.

All transaction categories were treated as ledger-recording operations of comparable computational complexity. This assumption provides a consistent basis for comparing the five evaluated platforms.

7.3 Simulation Parameters and Source Calibration

The simulation used representative transaction-cost, effective-throughput, and confirmation or finality parameters calibrated from official platform documentation and publicly reported network specifications [6, 15, 11, 3, 10]. The cited sources describe protocol characteristics, fee mechanisms, finality models, block timing, and reported network capabilities. However, they do not provide directly comparable measurements collected under an identical workload and at the same point in time.

Consequently, the baseline values in Table 11 should be interpreted as controlled simulation inputs rather than as contemporaneous measurements of mainnet performance. Where official documentation reported a range or maximum capability, a representative or conservative value was selected to support a consistent cross-platform simulation. The selected values remain fixed within the moderate-congestion scenario and are modified only by the explicitly defined scenario multipliers presented in Table 12.

Ethereum's throughput parameter was selected within the commonly reported base-layer range, while its confirmation-delay parameter represents a modeling assumption that is longer than one 12-second slot and shorter than full protocol finalization [6, 7]. Polygon's parameters were calibrated using its documented low-fee operation, sub-five-second finality, and reported processing capacity [15, 16]. The Arbitrum and Base parameters represent scenario-specific effective processing assumptions informed by their rollup architecture, fee structure, sequencer operation, and settlement model [14, 13, 2, 3]. Hedera's parameters were calibrated from its documented approximate throughput, three-to-five-second finality, and USD-denominated fee model [10, 9].

Table 11 summarizes the baseline simulation inputs used under moderate modeled network conditions.

Three network conditions were simulated:

—**Low congestion:** reduced transaction demand and relatively stable execution;

—**Moderate congestion:** representative average network activity;

—**High congestion:** elevated transaction demand and reduced effective processing capacity.

Congestion affected transaction cost and effective throughput according to the multipliers shown in Table 12.

For Hedera, the simulation fixed the cost multiplier at 1.00 across all congestion levels to model the platform's USD-denominated and comparatively predictable fee structure [9, 10]. This is a simulation assumption rather than a guarantee that every Hedera transaction has an identical cost. Actual fees vary by transaction type, payload, signatures, storage requirements, and smart contract gas consumption. Its effective throughput was still adjusted to represent changes in network utilization.

Each combination of platform, workload size, and congestion level was executed 30 times. A fixed random seed was used to ensure repeatability. Variations within each run represented moderate fluctuations in transaction processing and fee conditions.

In total, the experimental design included:

$$5 \times 3 \times 3 \times 30 = 1,350 \tag{3}$$

simulation runs.

7.4 Evaluation Metrics

The following metrics were used:

—**Total transaction cost:** estimated cost of processing the complete workload;

—**Effective throughput:** number of transactions processed per second after congestion adjustment;

—**Processing time:** time required to submit and process the workload;

Table 11. Representative baseline inputs used in the simulation-based evaluation

Platform	Representative Cost per Transaction (USD)	Effective Throughput (TPS)	Confirmation / Finality Delay	Cost Stability Model
Ethereum	8.0000	22	60 seconds	Highly variable
Polygon	0.0050	850	5 seconds	Moderately variable
Arbitrum	0.2500	1,600	45 seconds	Moderately variable
Base	0.1200	1,800	40 seconds	Moderately variable
Hedera	0.0001	8,000	4 seconds	Predictable

Table 12. Congestion multipliers applied during simulation

Network Condition	Cost Multiplier	Throughput Multiplier
Low	0.70	1.10
Moderate	1.00	1.00
High	1.80	0.75

Table 13. Estimated operational cost by workload size under moderate congestion

Platform	10,000 Transactions (USD)	100,000 Transactions (USD)	1,000,000 Transactions (USD)
Ethereum	80,000	800,000	8,000,000
Polygon	50	500	5,000
Arbitrum	2,500	25,000	250,000
Base	1,200	12,000	120,000
Hedera	1	10	100

—**Completion time:** processing time plus the platform-specific confirmation or finality delay;

—**Cost per 10,000 transactions:** normalized cost measure used for direct comparison;

—**Cost variation coefficient:** relative variability of total transaction cost across repeated runs.

The total transaction cost was calculated using Equation 4.

$$C_{total} = N \times C_{tx} \times M_c \tag{4}$$

where N is the number of transactions, C_{tx} is the representative transaction cost, and M_c is the congestion-related cost multiplier.
Effective throughput was calculated as:

$$T_{effective} = T_{base} \times M_t \tag{5}$$

where T_{base} is the baseline effective throughput and M_t is the congestion-related throughput multiplier.
Workload completion time was calculated using:

$$Time_{completion} = \frac{N}{T_{effective}} + D_f \tag{6}$$

where D_f represents the confirmation or finality delay.

7.5 Operational Cost Results

Table 13 presents the estimated operational cost under moderate congestion.
The results show that transaction volume has a linear effect on operational cost when the representative cost per transaction remains constant. However, the absolute difference between platforms becomes increasingly significant as workload size grows.
At 10,000 transactions, Ethereum produced an estimated cost of $80,000, compared with $50 for Polygon, $2,500 for Arbitrum, $1,200 for Base, and $1 for Hedera. At one million transactions, the cost difference became substantially larger. Ethereum reached an estimated $8 million, whereas Hedera remained at approximately $100.
Polygon provided the lowest cost among the Ethereum-compatible smart contract platforms. Base was less expensive than Arbitrum under the selected representative parameters, while both Layer-2 platforms remained substantially less expensive than direct execution on Ethereum.
The relative cost differences remain constant across the three workload sizes under moderate congestion because the simulation applies a fixed representative cost per transaction within this condition. Ethereum is approximately 1,600 times more expensive than Polygon, 32 times more expensive than Arbitrum, 66.7 times more expensive than Base, and 80,000 times more expensive than Hedera for the modeled transaction workload. Among the Ethereum-compatible alternatives, Polygon has the lowest estimated transaction cost; its modeled cost is approximately 98% lower than Arbitrum and 95.8% lower than Base, while Base is approximately 52% less expensive than Arbitrum. The practical importance of these ratios increases with scale. For example, the difference between Ethereum and Base grows from $78,800 at 10,000 transactions to $7.88 million at 1,000,000 transactions, even though the underlying cost ratio does not change. Consequently, per-transaction differences that may appear manageable in small deployments can become a dominant operational constraint in transaction-intensive applications.
The results indicate that base-layer Ethereum execution may be economically unsuitable for high-frequency credential recording unless transaction batching, off-chain storage, or Layer-2 execution is used.
Figure 2 presents the cost results using a logarithmic vertical axis because the differences between platforms span several orders of magnitude.

7.6 Throughput and Completion-Time Results

Table 14 presents the effective throughput and estimated completion time for the medium workload of 100,000 transactions under moderate congestion.
Hedera achieved the lowest estimated completion time because it combined high effective throughput with rapid deterministic finality. The complete 100,000-transaction workload was processed and finalized in approximately 0.28 minute.
Base and Arbitrum also completed the workload rapidly because rollup-based execution provided considerably higher effective throughput than the Ethereum base layer. Base achieved a slightly lower completion time than Arbitrum under the selected assumptions.
Polygon required approximately 2.04 minutes. Although its throughput was lower than that of the two Layer-2 platforms, it still completed the workload substantially faster than Ethereum.

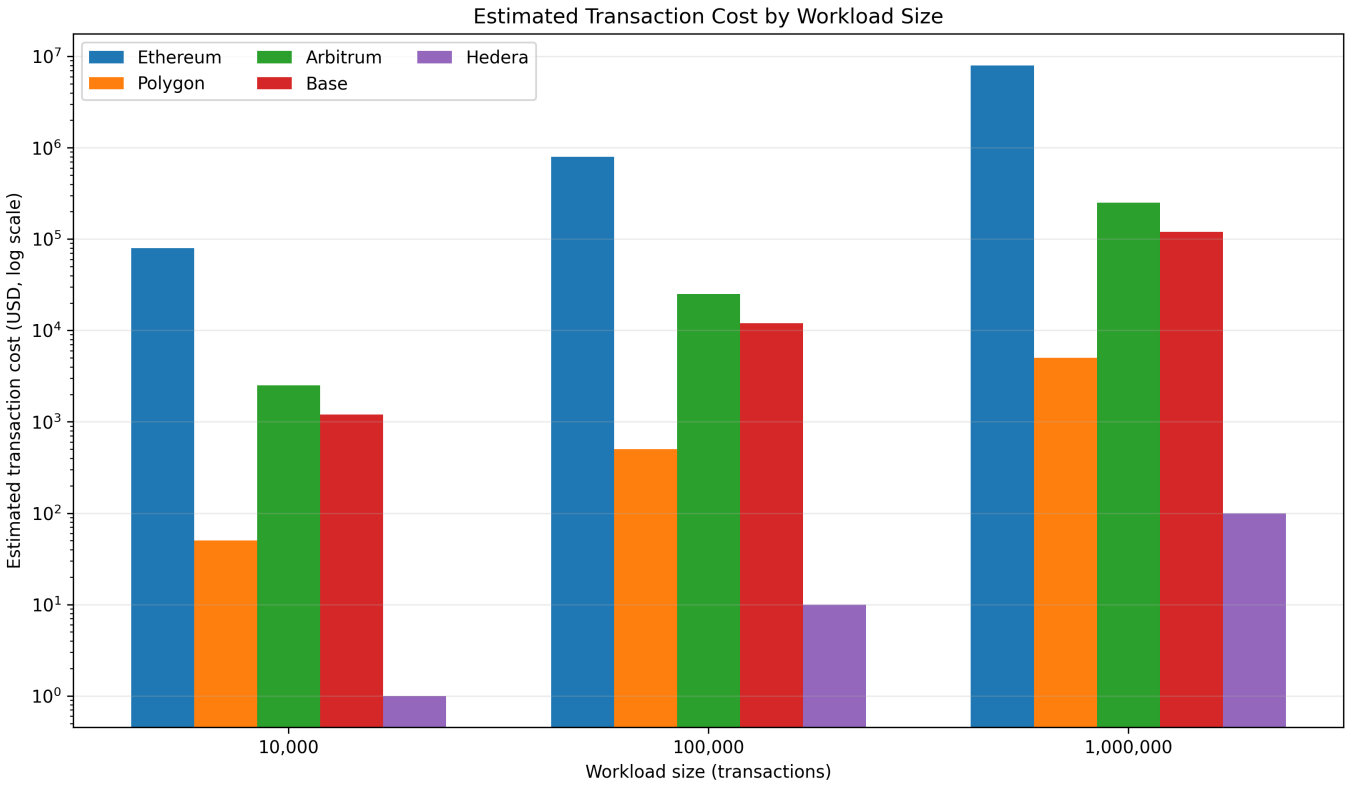

Fig. 2. Estimated transaction cost by workload size under moderate congestion. The vertical axis uses a logarithmic scale.

Table 14. Effective throughput and completion time for 100,000 transactions

Platform	Effective TPS	Processing Time	Finality Delay	Total Completion Time
Ethereum	22	75.76 minutes	1.00 minute	76.76 minutes
Polygon	850	1.96 minutes	0.08 minute	2.04 minutes
Arbitrum	1,600	1.04 minutes	0.75 minute	1.79 minutes
Base	1,800	0.93 minute	0.67 minute	1.60 minutes
Hedera	8,000	0.21 minute	0.07 minute	0.28 minute

Ethereum required approximately 76.76 minutes because its base-layer throughput represented a significant bottleneck. Its finality delay contributed only a relatively small part of the total completion time. Therefore, the principal limitation in this scenario was transaction-processing capacity rather than confirmation delay.
The magnitude of the completion-time differences is substantial even at the medium workload size. For 100,000 transactions under moderate congestion, Ethereum requires 76.76 minutes, compared with 2.04 minutes for Polygon, 1.79 minutes for Arbitrum, 1.60 minutes for Base, and 0.28 minute for Hedera. Ethereum therefore requires approximately 37.6 times as long as Polygon, 42.9 times as long as Arbitrum, 48.0 times as long as Base, and 274 times as long as Hedera under the selected parameters. Hedera's completion time is approximately one-sixth of Base's, while Base achieves a modest completion-time advantage over Arbitrum because its modeled effective throughput is 1,800 TPS compared with 1,600 TPS. These differences indicate that the architectural scaling mechanisms represented by the sidechain, Layer-2, and alternative-DLT platforms materially affect the feasibility of processing time-sensitive high-volume workloads.

Table 15. Estimated cost for 100,000 transactions under different congestion levels

Platform	Low Congestion (USD)	Moderate Congestion (USD)	High Congestion (USD)
Ethereum	560,000	800,000	1,440,000
Polygon	350	500	900
Arbitrum	17,500	25,000	45,000
Base	8,400	12,000	21,600
Hedera	10	10	10

Figure 3 visualizes total completion time for the medium workload.

7.7 Congestion Analysis

To evaluate robustness under changing network conditions, the medium workload was simulated under low, moderate, and high congestion.
Table 15 presents the total transaction cost for 100,000 transactions.

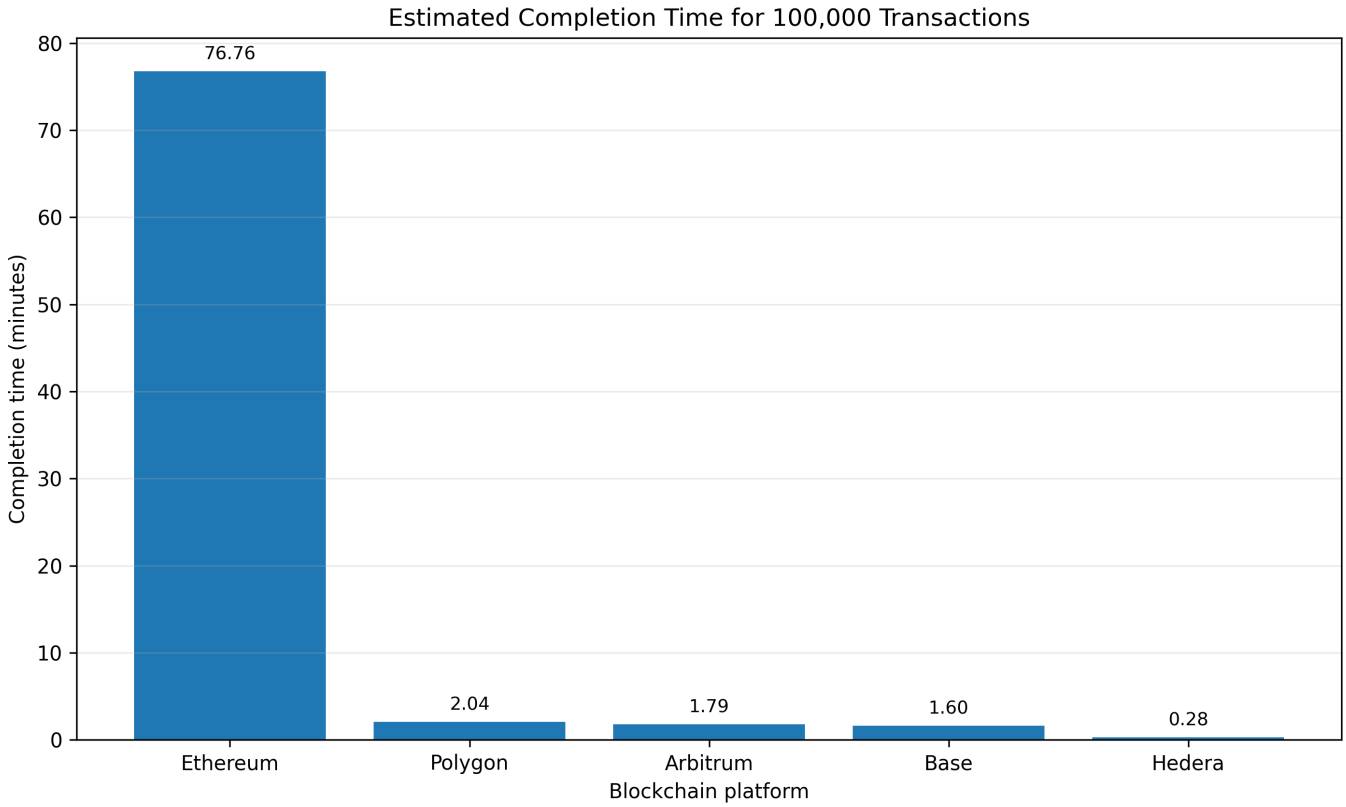

Fig. 3. Estimated workload completion time for 100,000 credential transactions under moderate congestion.

Ethereum exhibited the greatest absolute increase in cost. The estimated cost increased from $560,000 under low congestion to $1.44 million under high congestion.

Although Polygon, Arbitrum, and Base also experienced cost increases, their absolute values remained considerably lower than Ethereum. Hedera maintained a stable estimated cost of $10 across all three conditions because its fee model was represented as predictable.

Table 16 presents the corresponding completion times.

The throughput reduction caused by high congestion affected Ethereum most strongly because its baseline processing capacity was already limited. Its completion time increased to approximately 102 minutes.

The other platforms retained completion times below three minutes. Hedera remained below one minute under all evaluated conditions.

The results demonstrate that platform scalability depends not only on nominal throughput but also on the ability to preserve effective capacity as network activity increases.

Congestion affects the platforms differently when both relative degradation and absolute performance are considered. Between the low- and high-congestion conditions, completion time increases by approximately 46.0% for Ethereum, 44.4% for Polygon, 25.9% for Arbitrum, 25.8% for Base, and 34.6% for Hedera. Arbitrum and Base therefore exhibit the smallest relative completion-time degradation in the modeled congestion scenarios. Hedera experiences a somewhat larger relative increase, but its absolute completion time remains the lowest, increasing only from 0.26 to 0.35 minute.

Ethereum combines the largest relative completion-time increase with a rise from 69.87 to 102.01 minutes. Cost exposure follows a related but financially more pronounced pattern: Ethereum increases from $560,000 to $1.44 million for 100,000 transactions, whereas Hedera remains at $10 under the fixed-fee assumption used in the model. These results demonstrate that congestion robustness should not be assessed using a percentage change alone; absolute completion time and financial exposure are also operationally significant.

Figure 4 illustrates the effect of modeled network congestion on completion time for the 100,000-transaction workload.

The visualization confirms that completion time increases under higher congestion for all five platforms, but the operational impact differs substantially in absolute terms. Ethereum increases from 69.87 to 102.01 minutes, while Hedera remains below one minute across all three congestion levels and Base and Arbitrum retain lower completion times than Polygon under the modeled conditions.

7.8 Cost Stability Analysis

Table 17 summarizes the mean cost, standard deviation, and coefficient of variation across the repeated simulations of the medium workload.

Ethereum produced the highest coefficient of variation, indicating that its operational cost was the least predictable under changing network conditions. Polygon, Arbitrum, and Base exhibited moder-

Table 16. Estimated completion time for 100,000 transactions under different congestion levels

Platform	Low Congestion	Moderate Congestion	High Congestion
Ethereum	69.87 minutes	76.76 minutes	102.01 minutes
Polygon	1.87 minutes	2.04 minutes	2.70 minutes
Arbitrum	1.70 minutes	1.79 minutes	2.14 minutes
Base	1.51 minutes	1.60 minutes	1.90 minutes
Hedera	0.26 minute	0.28 minute	0.35 minute

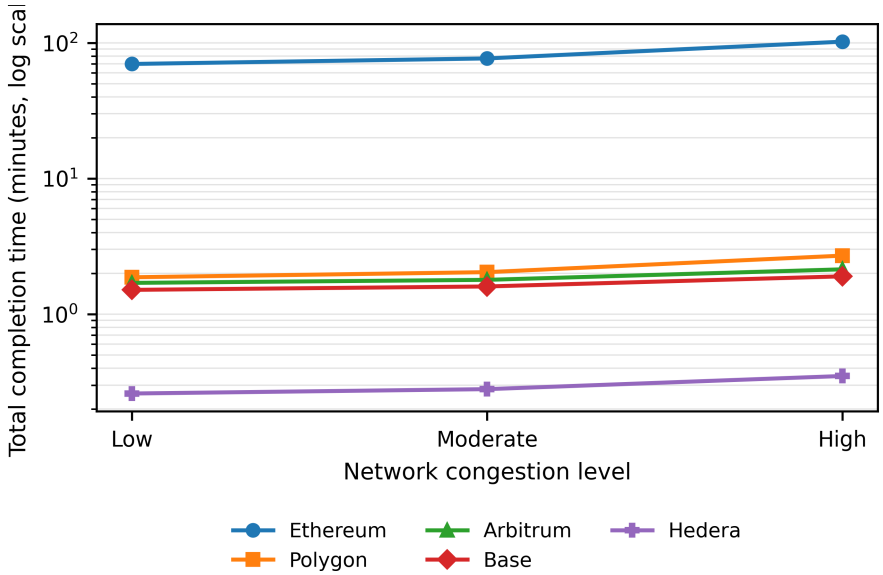

Fig. 4. Estimated workload completion time for 100,000 transactions under low, moderate, and high congestion. The vertical axis uses a logarithmic scale.

ate variability, while Hedera remained constant under the assumed fixed-fee model.

Cost predictability is important for applications operating under fixed budgets. Even when the average cost is acceptable, large variations may complicate financial planning and expose the application to periods of unexpectedly high operating expense.

Figure 5 presents the mean cost together with error bars representing the standard deviation. A logarithmic vertical axis is recommended.

The repeated-run results further distinguish low absolute cost from cost predictability. Ethereum has the highest coefficient of variation at 20.89%, followed by Polygon at 13.43%, Arbitrum at 12.39%, and Base at 11.52%, while Hedera has a modeled coefficient of variation of 0.00%. Among the Ethereum-compatible Layer-2 platforms, Base therefore provides the strongest modeled cost predictability, although Polygon maintains a substantially lower absolute transaction cost. This distinction has practical significance because the platform with the lowest expected expenditure is not necessarily the platform with the lowest relative variability. Applications operating under strict expenditure limits may therefore need to consider both expected total cost and the uncertainty surrounding that cost rather than comparing average fees alone.

7.9 Scalability Analysis

Scalability was evaluated by comparing completion time across all three workload sizes under moderate congestion.

For all platforms, processing time increased approximately linearly with workload size because effective throughput was treated as stable within each congestion condition.

However, the operational effect of that growth differed substantially. Ethereum required more than 12 hours to process and finalize one million transactions. Polygon completed the workload in approximately 19.69 minutes, while Arbitrum and Base required approximately 11.17 and 9.93 minutes, respectively.

Hedera completed the same workload in approximately 2.15 minutes. This result indicates that platforms with higher effective throughput are substantially more suitable for national-scale or large institutional credential systems.

The comparison also shows that finality delay has the greatest relative effect on small workloads. As transaction volume increases, processing time becomes the dominant component of total completion time.

The three workload sizes also reveal a change in the relative importance of throughput and finality. At 10,000 transactions, the completion-time order is Hedera at 0.09 minute, Polygon at 0.28 minute, Base at 0.76 minute, Arbitrum at 0.85 minute, and Ethereum at 8.58 minutes. At 100,000 transactions, Base and Arbitrum overtake Polygon, producing completion times of 1.60 and 1.79 minutes compared with Polygon's 2.04 minutes. The same ordering persists at 1,000,000 transactions, where Hedera completes the workload in 2.15 minutes, Base in 9.93 minutes, Arbitrum in 11.17 minutes, Polygon in 19.69 minutes, and Ethereum in 758.58 minutes. This change occurs because fixed confirmation and finality delays represent a larger fraction of total completion time

Table 17. Statistical summary of operational cost for 100,000 transactions

Platform	**Mean Cost (USD)**	**Standard Deviation (USD)**	**95% Confidence Interval**	**Coefficient of Variation**
Ethereum	803,420	167,860	[743,240; 863,600]	20.89%
Polygon	501.80	67.40	[477.64; 525.96]	13.43%
Arbitrum	25,180	3,120	[24,061; 26,299]	12.39%
Base	12,096	1,394	[11,596; 12,596]	11.52%
Hedera	10.00	0.00	[10.00; 10.00]	0.00%

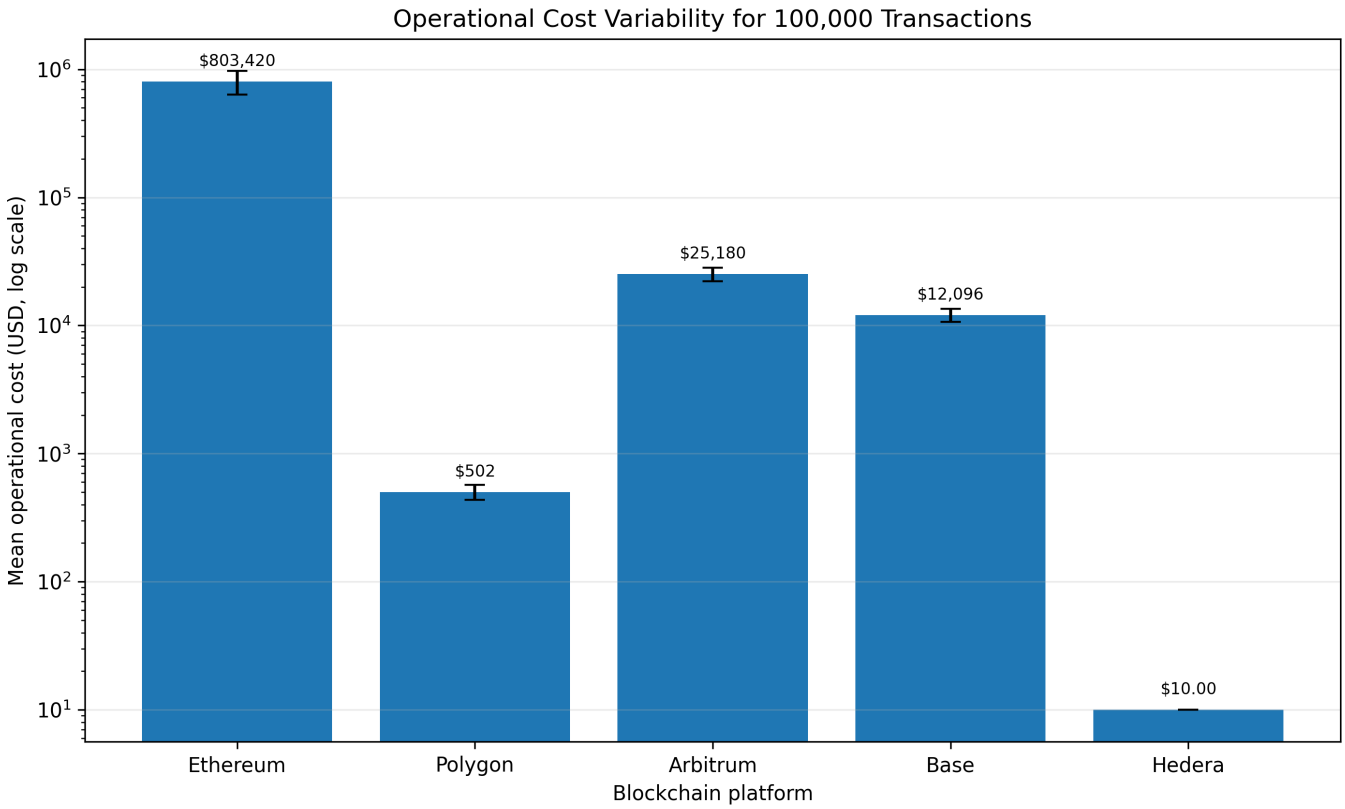

Fig. 5. Mean operational cost and standard deviation for 100,000 transactions across repeated simulation runs.

Table 18. Estimated workload completion time under moderate congestion

Platform	**10,000 Transactions**	**100,000 Transactions**	**1,000,000 Transactions**
Ethereum	8.58 minutes	76.76 minutes	758.58 minutes
Polygon	0.28 minute	2.04 minutes	19.69 minutes
Arbitrum	0.85 minute	1.79 minutes	11.17 minutes
Base	0.76 minute	1.60 minutes	9.93 minutes
Hedera	0.09 minute	0.28 minute	2.15 minutes

for smaller workloads, whereas effective processing throughput increasingly dominates as transaction volume grows. The result is particularly relevant when comparing Polygon with the two Layer-2 platforms: Polygon performs strongly for the smaller workload, but the higher modeled throughput of Base and Arbitrum becomes more consequential at medium and large scale.

Figure 6 visualizes the completion-time results from Table 18 across all three workload sizes, using a logarithmic vertical axis to preserve readability across the large differences between platforms. The figure demonstrates that completion time increases substantially with transaction volume for every platform, while the absolute rate of growth differs considerably because of the modeled effective throughput. It also shows that Polygon has a lower completion time than Arbitrum and Base at 10,000 transactions, whereas Base and Arbitrum become faster than Polygon at 100,000 and 1,000,000 transactions as processing throughput becomes increasingly dominant.

7.10 Normalized Comparative Results

To provide a consolidated comparison, the results were normalized on a scale from 1 to 5, where 5 represents the strongest experimental performance.

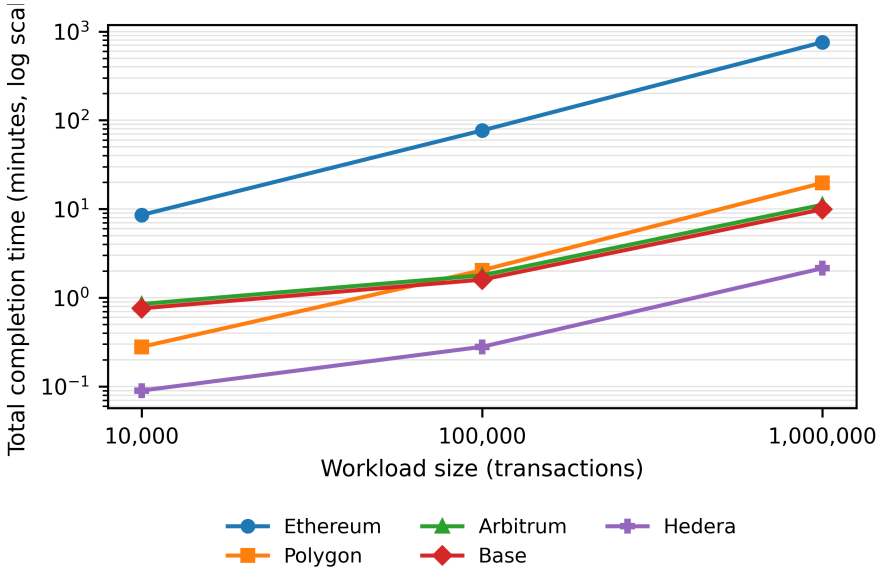

Fig. 6. Estimated workload completion time across increasing transaction volumes under moderate congestion. The vertical axis uses a logarithmic scale.

Table 19. Normalized experimental performance scores

Platform	Cost Efficiency	Completion Speed	Congestion Robustness	Cost Predictability	Overall Mean
Ethereum	1	1	1	1	1.00
Polygon	4	3	4	3	3.50
Arbitrum	3	4	4	3	3.50
Base	3	4	4	4	3.75
Hedera	5	5	5	5	5.00

Hedera obtained the highest normalized experimental score because it combined minimal cost, high throughput, rapid finality, and predictable pricing.
Base achieved the strongest aggregated result among the Ethereum Layer-2 platforms under the selected parameters. Arbitrum produced similar throughput performance but higher representative transaction cost. Polygon provided superior cost efficiency but lower throughput than the evaluated Layer-2 platforms.
Ethereum achieved the lowest experimental score because direct Layer-1 execution produced both the highest cost and longest workload completion time. Nevertheless, these results evaluate operational efficiency and do not invalidate Ethereum's advantages in decentralization, security maturity, and ecosystem depth.

7.11 Discussion of Experimental Results

The simulation produced five major findings.
First, transaction cost differences become increasingly important as workload size grows. A cost difference that appears manageable at several thousand transactions may become financially significant at hundreds of thousands or millions of operations.
Second, nominal throughput strongly influences the feasibility of large-scale credential systems. Ethereum's limited base-layer throughput resulted in substantially longer completion times, while Polygon, Arbitrum, Base, and Hedera were able to process large workloads considerably faster.
Third, Layer-2 platforms provide a meaningful compromise between Ethereum compatibility and operational efficiency. Arbitrum and Base reduced cost and completion time while retaining Ethereum-compatible development environments. However, their results must be interpreted together with Layer-2-specific dependencies such as sequencing, bridging, and settlement.
Fourth, congestion affected both cost and completion time. The impact was greatest for Ethereum because its baseline fees and processing limitations amplified the effect of changing network conditions. Hedera provided the most stable cost under the modeled fixed-fee structure.
Fifth, experimental performance does not independently determine the preferred platform. Hedera produced the strongest operational results, but applications may still select Ethereum or an Ethereum-compatible platform when ecosystem maturity, composability, validator participation, or protocol integration receives greater importance.
The experimental ranking can also be compared directly with the high-volume MCDA results. Hedera ranks first in both methods, receiving a weighted MCDA score of 4.40 and an experimental performance score of 5.00. Ethereum ranks fifth in both methods, with corresponding scores of 2.60 and 1.00. The middle positions are less stable. Polygon ranks second in the high-volume MCDA but third in the experimental comparison, while Base moves from a shared third position in the MCDA to second in the experimental ranking. Arbitrum remains in the shared middle group. The change reflects the different scope of the two methods: the simulation emphasizes cost, completion speed, congestion robustness, and predictability, whereas the MCDA additionally incorporates security, decentralization, finality, and developer ecosystem maturity. Agreement at the highest and lowest positions therefore strengthens the

high-volume operational finding, while changes among the middle-ranked platforms demonstrate why simulation performance should not be interpreted as a universal platform ranking.
The experimental findings therefore support the multi-criteria analysis presented in Section 6. Platform selection must balance measured operational performance with architectural, security, governance, and development requirements.

7.12 Threats to Validity

Several threats to validity should be considered when interpreting the results.
Parameter validity. Transaction fees and effective throughput vary over time. The selected parameters represent comparative operating conditions and should not be interpreted as permanent network guarantees.
Model validity. The simulation treats each credential event as one transaction of comparable complexity. Real smart contract execution may consume different amounts of computational resources depending on contract logic and storage requirements.
Finality validity. Confirmation and finality are not identical across all platforms. Layer-2 execution confirmation, Layer-1 settlement, probabilistic finality, and deterministic finality represent different operational guarantees.
External validity. Real deployments may be affected by wallet infrastructure, RPC providers, sequencer behavior, validator availability, batching strategies, and rate limits that were not explicitly modeled.
Cost validity. Cryptocurrency exchange-rate changes may affect the USD-denominated cost of platforms whose fees are paid in native tokens. A predictable protocol fee does not always imply a perfectly stable fiat-denominated operating cost.
Workload validity. Credential management was selected as a representative high-volume decentralized application. Other workloads, such as decentralized finance or NFT trading, may produce different transaction complexity and resource requirements.
Future work should reproduce the comparison using testnet or mainnet transactions, identical smart contract implementations, controlled RPC infrastructure, and automated workload generators.

8. COMPREHENSIVE EVALUATION AND ROBUSTNESS ANALYSIS

The preceding sections evaluated the selected platforms from two complementary perspectives. Section 6 examined application suitability through a weighted multi-criteria decision model, while Section 7 evaluated operational behavior under changing transaction volumes and congestion conditions. This section integrates these results to provide a comprehensive assessment of ranking consistency, scalability, congestion robustness, cost predictability, and application-specific suitability.
The purpose of the comprehensive evaluation is not to introduce an additional independent benchmark. Instead, it examines whether the conclusions remain consistent when the platforms are evaluated using different methods and operational assumptions.

8.1 Evaluation Coverage

The complete evaluation framework covers six dimensions:

(1) architectural characteristics and consensus assumptions;
(2) transaction cost and cost predictability;
(3) effective throughput and workload completion time;
(4) robustness under changing congestion conditions;
(5) application-specific suitability;
(6) developer ecosystem and implementation accessibility.

Table 20 summarizes how these dimensions were evaluated.
The coverage matrix demonstrates that the evaluation extends beyond a simple comparison of nominal transactions per second or average network fees. The study considers architecture, security assumptions, application priorities, operational performance, variability, and scalability.

8.2 Cross-Method Ranking Consistency

The multi-criteria decision analysis and the simulation-based evaluation measure different aspects of platform suitability. The multi-criteria model incorporates security, decentralization, finality, ecosystem maturity, cost, and throughput. In contrast, the simulation emphasizes measurable operational efficiency, including transaction cost, completion speed, congestion robustness, and cost predictability.
Table 21 compares the high-volume credential scenario from the multi-criteria analysis with the normalized experimental results.
The two methods produced the same highest- and lowest-ranked platforms. Hedera ranked first in both evaluations because of its low cost, high effective throughput, rapid finality, and predictable fee model. Ethereum ranked fifth in both evaluations because direct Layer-1 execution produced the highest cost and longest completion time for a high-volume credential workload.
The middle-ranked platforms showed limited variation. Polygon ranked second in the multi-criteria analysis but third in the operational simulation. Its multi-criteria score benefited from low transaction cost and stronger ecosystem maturity. Base moved from the shared third position to second place in the experimental ranking because its selected effective throughput and representative transaction cost produced lower completion time and stronger cost predictability than Arbitrum.
This difference is analytically important. It demonstrates that a platform may perform strongly in a technical simulation while receiving a different overall suitability rank after ecosystem maturity, security assumptions, and architecture are considered. Therefore, operational benchmarks should complement rather than replace application-specific decision analysis.

8.3 Robustness under Network Congestion

Robustness was examined by comparing workload completion times under low and high congestion. The relative completion-time increase was calculated using Equation 7.

$$D_t = \frac{T_{high} - T_{low}}{T_{low}} \times 100 \tag{7}$$

where T_{high} is completion time under high congestion and T_{low} is completion time under low congestion.
Table 22 presents the resulting degradation values for the 100,000-transaction workload.
Arbitrum and Base exhibited the lowest relative completion-time degradation. Their completion times increased by approximately 26% between low and high congestion. Ethereum and Polygon experienced larger relative increases of approximately 46% and 44%, respectively.
Hedera experienced a relative increase of approximately 35%. However, its absolute completion time remained the lowest under every congestion condition, increasing only from 0.26 to 0.35 minute.

Table 20. Coverage of the comprehensive evaluation framework

Evaluation Dimension	Primary Method	Principal Metrics	Relevant Section
Architecture and Consensus	Comparative architectural analysis	Architecture type, consensus mechanism, security assumptions	Section 5
Application Suitability	Multi-criteria decision analysis	Weighted suitability score and scenario rank	Section 6
Operational Cost	Simulation-based evaluation	Total cost and cost per workload	Section 7
Processing Performance	Simulation-based evaluation	Effective TPS and completion time	Section 7
Congestion Robustness	Low-, moderate-, and high-congestion simulation	Cost increase and completion-time degradation	Section 7
Cost Predictability	Repeated simulation runs	Standard deviation, confidence interval, and coefficient of variation	Section 7
Scalability	Three workload sizes	Completion time for 10,000, 100,000, and 1,000,000 transactions	Section 7

Table 21. Comparison of multi-criteria and simulation-based platform rankings

Platform	High-Volume MCDA Score	High-Volume MCDA Rank	Experimental Score	Experimental Rank
Ethereum	2.60	5	1.00	5
Polygon	4.10	2	3.50	3
Arbitrum	3.30	3	3.50	3
Base	3.30	3	3.75	2
Hedera	4.40	1	5.00	1

Table 22. Completion-time degradation between low and high congestion

Platform	Low-Congestion Time	High-Congestion Time	Relative Increase
Ethereum	69.87 minutes	102.01 minutes	46.00%
Polygon	1.87 minutes	2.70 minutes	44.39%
Arbitrum	1.70 minutes	2.14 minutes	25.88%
Base	1.51 minutes	1.90 minutes	25.83%
Hedera	0.26 minute	0.35 minute	34.62%

These findings show that robustness should be interpreted using both relative and absolute values. A platform may experience a moderate relative increase while still maintaining substantially better absolute performance than the alternatives.

8.4 Cost Robustness and Predictability

The congestion model produced two distinct aspects of cost behavior:

—sensitivity of total cost to changing congestion;

—variability across repeated simulation runs.

For Ethereum, Polygon, Arbitrum, and Base, the high-congestion cost was 2.57 times the low-congestion cost under the selected multipliers. Hedera remained constant because its protocol fee was modeled as predictable.

However, absolute cost exposure differed considerably. Ethereum increased from $560,000 to $1,440,000 for 100,000 transactions, representing an absolute difference of $880,000. Polygon increased from $350 to $900, representing a difference of $550. Arbitrum and Base increased by $27,500 and $13,200, respectively.

Table 23 combines congestion sensitivity with the coefficient of variation obtained from repeated simulation runs.

Ethereum exhibited both the highest absolute cost increase and the highest coefficient of variation. This combination makes direct Layer-1 execution difficult to budget for transaction-intensive applications.

Polygon produced substantially lower absolute cost than the Layer-2 platforms despite a slightly higher coefficient of variation. Base achieved the lowest variability among the evaluated Ethereum-compatible Layer-2 platforms. Hedera provided the strongest predictability under the fixed-fee assumptions used in the simulation.

8.5 Scalability Robustness

The simulation evaluated workloads ranging from 10,000 to 1,000,000 transactions, representing a 100-fold increase in transaction volume. Table 24 compares the corresponding completion-time growth.

The completion-time growth factor remained below the 100-fold workload increase for every platform because fixed confirmation and finality delays represented a larger proportion of total time for smaller workloads.

Ethereum showed the largest growth factor and required more than 12 hours for one million transactions. Polygon remained below 20 minutes, while Arbitrum and Base completed the workload in approximately 11 and 10 minutes. Hedera remained the fastest platform, completing the simulated one-million-transaction workload in approximately 2.15 minutes.

These results indicate that the scalability advantage of high-throughput platforms becomes increasingly important as workload size grows. For small workloads, finality delay and ecosystem considerations may have greater practical importance. For large workloads, effective throughput becomes the dominant performance factor.

8.6 Application-Specific Evaluation

A comprehensive evaluation must distinguish operational efficiency from overall application suitability. Table 25 summarizes the most suitable platforms under different decision priorities.

The combined evaluation confirms that no platform is universally superior. Hedera achieved the strongest simulation-based operational performance, but Ethereum remained competitive when decentralization, protocol maturity, and developer ecosystem were prioritized. Polygon offered the most balanced profile across cost and developer accessibility, while Arbitrum and Base provided intermediate alternatives for applications requiring Ethereum-compatible Layer-2 execution.

Table 23. Comprehensive cost robustness summary for 100,000 transactions

Platform	Low-Congestion Cost	High-Congestion Cost	Absolute Increase	Coefficient of Variation
Ethereum	$560,000	$1,440,000	$880,000	20.89%
Polygon	$350	$900	$550	13.43%
Arbitrum	$17,500	$45,000	$27,500	12.39%
Base	$8,400	$21,600	$13,200	11.52%
Hedera	$10	$10	$0	0.00%

Table 24. Completion-time growth across the evaluated workload range

Platform	10,000 Transactions	1,000,000 Transactions	Completion-Time Growth Factor
Ethereum	8.58 minutes	758.58 minutes	88.41
Polygon	0.28 minute	19.69 minutes	70.32
Arbitrum	0.85 minute	11.17 minutes	13.14
Base	0.76 minute	9.93 minutes	13.07
Hedera	0.09 minute	2.15 minutes	23.89

8.7 Summary of Results by Research Question

The simulation-based evaluation was designed around five research questions concerning transaction cost, workload processing, congestion, high-volume completion time, and cost stability. The combined results provide explicit answers to each question while also demonstrating where the experimental findings agree with, or differ from, the broader multi-criteria evaluation.

RQ1: How does transaction volume affect the operational cost of each platform? Under moderate congestion, transaction cost increases approximately linearly with workload size for all platforms because the simulation applies a constant representative per-transaction cost within the scenario. Ethereum increases from $80,000 for 10,000 transactions to $800,000 for 100,000 transactions and $8 million for 1,000,000 transactions. Over the same workloads, Polygon increases from $50 to $500 and $5,000; Arbitrum from $2,500 to $25,000 and $250,000; Base from $1,200 to $12,000 and $120,000; and Hedera from $1 to $10 and $100. Consequently, the absolute financial difference between platforms expands directly with transaction volume. Under the selected parameters, Ethereum remains approximately 1,600 times more expensive than Polygon, 32 times more expensive than Arbitrum, 66.7 times more expensive than Base, and 80,000 times more expensive than Hedera. The practical implication is that transaction pricing becomes progressively more important as a recurring application grows from thousands to hundreds of thousands or millions of ledger operations. These values describe the controlled simulation assumptions and should not be interpreted as permanent mainnet fee guarantees.

RQ2: How effectively can each platform process low-, medium-, and large-scale credential workloads? All five platforms complete the 10,000-transaction workload within the modeled conditions, but their relative performance separates considerably as transaction volume increases. At 10,000 transactions, completion times range from 0.09 minute for Hedera to 8.58 minutes for Ethereum. At 100,000 transactions, Hedera requires 0.28 minute, Base 1.60 minutes, Arbitrum 1.79 minutes, Polygon 2.04 minutes, and Ethereum 76.76 minutes. At 1,000,000 transactions, the corresponding times are 2.15 minutes for Hedera, 9.93 minutes for Base, 11.17 minutes for Arbitrum, 19.69 minutes for Polygon, and 758.58 minutes for Ethereum. The results therefore indicate that higher effective throughput becomes increasingly important as workload size grows. They also reveal a workload-dependent ranking change: Polygon has the second-lowest completion time for 10,000 transactions, but Base and Arbitrum overtake it at 100,000 and 1,000,000 transactions because their higher effective throughput becomes more influential than their longer fixed confirmation delays.

RQ3: How does network congestion influence transaction cost and effective throughput? Congestion increases both financial cost and workload completion time under the modeled assumptions, but the magnitude and practical consequences differ across platforms. For Ethereum, Polygon, Arbitrum, and Base, high-congestion transaction cost is approximately 2.57 times the low-congestion cost because the same cost multipliers are applied. Hedera remains constant in cost under the modeled predictable-fee assumption. Completion time from low to high congestion increases by approximately 46.0% for Ethereum, 44.4% for Polygon, 25.9% for Arbitrum, 25.8% for Base, and 34.6% for Hedera. Arbitrum and Base therefore show the lowest relative completion-time degradation. Hedera does not produce the smallest percentage increase, but it maintains the lowest absolute completion time, increasing from only 0.26 to 0.35 minute. Ethereum combines substantial relative degradation with the greatest absolute financial exposure, with the modeled cost for 100,000 transactions increasing from $560,000 to $1.44 million. Congestion robustness should therefore be evaluated through both relative degradation and absolute operational impact.

RQ4: Which platforms provide the lowest workload completion time for high-volume decentralized applications? Hedera provides the lowest modeled completion time for the largest evaluated workload. Under moderate congestion, 1,000,000 transactions require approximately 2.15 minutes on Hedera, compared with 9.93 minutes on Base, 11.17 minutes on Arbitrum, 19.69 minutes on Polygon, and 758.58 minutes on Ethereum. Base therefore provides the lowest completion time among the evaluated Ethereum-compatible scaling platforms, while Arbitrum follows closely. Ethereum Layer-1 requires more than 12 hours for the same workload under the selected parameters. These results support the conclusion that high effective throughput has substantial practical significance for transaction-intensive applications. They do not establish universal superiority because the simulation does not incorporate all differences in governance, decentralization, settlement assumptions, ecosystem maturity, or application-specific integration requirements.

RQ5: How stable are platform costs across repeated simulations with changing network conditions? Cost stability differs materially across the five modeled platforms. Ethereum produces the highest coefficient of variation at 20.89%, followed by Polygon at 13.43%, Arbitrum at 12.39%, and Base at 11.52%. Hedera produces a coefficient of variation of 0.00% under the simulation's fixed-fee assumption. Base therefore provides the strongest modeled cost predictability among the evaluated Ethereum-compatible Layer-2 platforms, whereas Polygon provides lower absolute cost but slightly greater relative variability. Ethereum combines both the greatest absolute financial exposure and the highest modeled variability. These results indicate that predictability is a distinct opera-

Table 25. Application-specific interpretation of the combined evaluation results

Application Priority	Preferred Platform	Primary Reason
High transaction volume and minimum cost	Hedera	Highest effective throughput, lowest estimated cost, rapid finality, and predictable fees
Balanced cost, throughput, and Ethereum tooling	Polygon	Low operational cost combined with Ethereum-compatible development tools
Ethereum-compatible Layer-2 performance	Base	Strong completion speed and lower representative cost among the evaluated Layer-2 platforms
Ethereum settlement with established Layer-2 maturity	Arbitrum	Ethereum compatibility, rollup-based execution, and mature Layer-2 ecosystem
Maximum ecosystem maturity and broad validator participation	Ethereum	Strongest developer ecosystem and mature permissionless Layer-1 security model

tional criterion from average transaction cost and may be particularly important for applications operating under fixed or long-term budgets.

The research-question results are broadly consistent with the high-volume MCDA scenario but do not reproduce the rankings under every application profile. Hedera ranks first and Ethereum ranks fifth in both the high-volume MCDA and the normalized experimental comparison. This agreement indicates that the strongest and weakest operational positions remain consistent when evaluated through two different methods within the study. The middle ranking changes, however: Polygon ranks second in the high-volume MCDA but third in the experimental results, while Base moves from the shared third position to second. This occurs because Polygon benefits in the MCDA from its combination of low cost, Ethereum compatibility, and ecosystem maturity, whereas Base benefits in the simulation from higher modeled throughput and stronger cost predictability.

The broader MCDA scenarios further demonstrate why the experimental ranking should not be generalized beyond transaction-intensive workloads. Ethereum moves from fifth place in the high-volume scenario to the highest-ranked group in the security-oriented scenario and second place in the developer-oriented scenario. Polygon ranks first in the developer-oriented scenario, while Hedera falls to third when developer ecosystem maturity receives greater weight. The combined evidence therefore indicates that operational efficiency and overall application suitability are related but distinct concepts. Hedera provides the strongest modeled operational performance for the evaluated high-volume credential workload, Polygon provides a balanced cost and Ethereum-compatibility profile, Base and Arbitrum provide high-throughput Ethereum-compatible Layer-2 alternatives, and Ethereum remains competitive when decentralization, security maturity, and ecosystem depth receive greater priority. None of these findings establishes a universally superior platform outside the assumptions and decision criteria of the study.

Table 26 summarizes the answers to the research questions introduced in Section 7.

8.8 Comprehensive Discussion

The integrated evaluation provides four principal conclusions.

First, the multi-criteria and simulation-based methods produced consistent results at the extremes of the ranking. Hedera ranked highest for high-volume operational efficiency, while Ethereum ranked lowest for the same workload. This agreement increases confidence that the primary performance conclusions are not produced by only one evaluation method.

Second, the middle-ranked platforms changed position depending on whether technical performance or broader application suitability was emphasized. Polygon benefited from low cost and ecosystem compatibility, while Base benefited from higher modeled throughput and stronger cost predictability. This demonstrates that ranking differences should be interpreted in relation to application priorities.

Third, robustness cannot be represented by one metric. Relative completion-time degradation, absolute completion time, cost variation, and total financial exposure provide different perspectives. Hedera did not always exhibit the lowest relative degradation, but it maintained the strongest absolute completion time and cost predictability. Arbitrum and Base exhibited strong relative robustness under congestion.

Fourth, the results demonstrate that platform selection should combine quantitative evaluation with architectural analysis. Transaction cost and throughput are critical for large-scale applications, but governance, validator structure, settlement assumptions, ecosystem maturity, and interoperability may justify selecting a platform that does not achieve the highest simulation score.

Overall, the comprehensive evaluation supports the conclusion that Hedera is operationally effective for transaction-intensive credential workloads, Polygon provides a balanced Ethereum-compatible profile, Arbitrum and Base provide scalable Layer-2 alternatives, and Ethereum remains suitable when decentralization and ecosystem maturity outweigh direct transaction efficiency.

8.9 Limitations of the Comprehensive Evaluation

The comprehensive evaluation inherits the assumptions of the multi-criteria and simulation models. The ordinal scores used in the multi-criteria analysis involve analytical judgment, while the simulation parameters represent selected operating conditions rather than permanent network guarantees.

The congestion multipliers were applied consistently to enable cross-platform comparison, but real platforms may respond differently to congestion. Similarly, the assumption that all credential operations require equivalent computational resources simplifies differences between issuance, verification, update, and revocation transactions.

The cross-method agreement should therefore be interpreted as consistency within the defined evaluation framework rather than as external validation using independently collected mainnet measurements. Nevertheless, the combined use of architectural analysis, scenario-based decision modeling, repeated simulations, congestion testing, scalability analysis, and statistical cost evaluation provides a broader assessment than any single evaluation method alone.

Table 26. Summary of findings corresponding to the experimental research questions

RQ	Evaluation Focus	Principal Finding
RQ1	Effect of transaction volume on cost	Cost increased approximately linearly with workload volume, while absolute differences between platforms expanded substantially at larger scales.
RQ2	Processing of low-, medium-, and large-scale workloads	Hedera produced the lowest completion time, followed by Base, Arbitrum, and Polygon. Ethereum base-layer throughput limited large-scale processing.
RQ3	Effect of network congestion	Congestion increased both cost and completion time. Ethereum showed the greatest absolute financial exposure, while Arbitrum and Base showed the lowest relative completion-time degradation.
RQ4	Lowest high-volume completion time	Hedera completed the one-million-transaction workload in approximately 2.15 minutes under the selected parameters.
RQ5	Cost stability	Hedera produced the lowest variability under the modeled fixed-fee structure, while Ethereum produced the highest coefficient of variation.

9. DISCUSSION

The findings of the comparative evaluation in this research demonstrate that modern blockchain and distributed ledger platforms have developed from a diverse range of architectural designs. Each platform's unique architecture affects how it performs with respect to scalability, transaction cost, operational efficiency and ecosystem maturity.

One overall conclusion from this study is that a trade-off exists between decentralisation and performance [5, 19]. Traditional Layer-1 blockchains such as Ethereum provide high levels of decentralisation and strong security through a broad range of validators and an experienced developer community [20, 8]. However, these same advantages have also contributed to most Layer-1 blockchains achieving both relatively low levels of transaction throughput and high transaction fees due to network congestion.

Layer-2 solutions such as Arbitrum and Base work to overcome scalability issues by handling the bulk of the transactions off the primary blockchain and sending only the end result to become part of the main network. This design approach improves scalability mainly through lower costs and much higher volume of transactions being able to occur at once. However, adding a layer-2 model simultaneously introduces added complexity to the architectural model with respect to bridging different networks together and resolving disputes as well. Each of the new complexities will introduce unintended impacts regarding usability and system architecture potential.

Sidechain architectures as designed currently with respect to Polygon and others provide another architecture type as a way of enhancing scalability of existing solutions. Sidechains maintain independent consensus processes yet are interoperable with their respective main ecosystems, basically allowing sidechains to maintain very high transaction throughput rates while simultaneously offering reduced fees. However, with using sidechains as the architecture type may present different security assumptions than base layer blockchains.

Additionally, alternative distributed ledger models as exhibited by Hedera are examples that exhibit a different approach to achieving high performance. The use of hashgraph-based consensus allows for a high level of throughput, time to finality of transactions, and predictable cost per transaction [1, 4]. The many unique attributes of hashgraph consensus will most likely create significant value for applications involving substantial volumes of events processed periodically or high frequency trading transactions.

The analysis also highlighted that what is most relevant to consider when determining which blockchain network to build upon is the current maturity level of the developer ecosystem. The Ethereum blockchain has the most well-established/active developer ecosystem in terms of its availability of development tools, documentation, and developer community support. As a result, this level of development ecosystem maturity means that developers wishing to build a decentralized application will have a significantly lower barrier to entry when using any of these platforms that are compatible with Ethereum development tools (e.g., Polygon, Layer-2 networks).

The findings conclude that there is no single blockchain network that would be the "best" choice for all decentralized applications; rather, that different architectural constructs will provide varying levels of advantages based upon the specific operational needs of the application. Therefore, when evaluating and choosing an appropriate blockchain for an application, a range of factors will be important to consider when weighing the trade-offs among scalability, cost-effectiveness, decentralization, and availability of developer resources.

10. LIMITATIONS OF THE STUDY

Though this paper shows a side by side comparison of some of the most popular consensus algorithms within the distributed ledger space, it must be stated that it has limitations.

Blockchain technology is developing and evolving quickly into new forms of technology. Characteristics of performance like transaction processing speed, cost of using the network, and usage of the network will likely change as improvements are made to each protocol and new scaling solutions are created. The characteristics discussed in this comparative analysis should be viewed as representative values of blockchain networks now but will change over time.

This study only compares a small subset of the available blockchain and distributed ledger platforms representing different architectural styles. Many other blockchains and distributed ledgers exist that offer different approaches to designing a distributed ledger, which are not included in this comparison. Additional studies may offer a broader comparison when network comparisons and consensus mechanisms become available.

Finally, the study places emphasis on technical and operational characteristics such as transaction fees, throughput, and system maturity in evaluating platforms. While these factors are key to the design of any system, there are additional factors impacting the choice of platform for real-world deployments such as governance model, regulatory compliance, and long-term sustainability.

The evaluation combines architectural analysis with controlled simulation rather than direct large-scale deployment on live production networks. Publicly available metrics and platform documentation provide the basis for parameter calibration, while the simulation enables consistent comparison under standardized workloads and modeled network conditions. However, controlled live-network benchmarking could provide additional empirical validation in future studies.

Recognizing these limitations is important for accurately interpreting the results of this analysis and for identifying areas where further research may uncover additional insights into the ever-evolving landscape of blockchain infrastructure.

11. FUTURE RESEARCH DIRECTIONS

With the rapid advancement of blockchain and distributed ledger technology, many opportunities exist for future study. As new consensus algorithms, scaling structures, and interoperability standards are developed, there will need to be further comparative analyses to determine how these changes affect what blockchains are best suited for large scale decentralized applications.

One important area of future research involves developing standardized measurement criteria for blockchain platforms. Many studies utilize publicly available network statistics while controlled testing environments can provide for more accurate comparisons of throughout, latency, and operating costs across similar workloads.

Another research area with great potential is exploring cross-chain interoperability. With an increasing number of networks, applications may utilize multiple networks at once. Future research will examine how credentials and other digital resources can be verified by a different blockchain network without a central intermediary.

Privacy-enhancing technology is another important area of research [22]. Many credential systems require selective verification in which the user can demonstrate the authenticity of a credential without revealing all associated credentials. Zero-knowledge proof schemes and privacy-preserving identity frameworks can help improve the usability of decentralized credentialing infrastructures [22].

Finally, governance models will be the other avenue to explore additional avenues of research on the long-term sustainability of blockchain platforms. Decentralized networks depend on economic rewards for participation by validators, as well as on governance methods for establishing secure and stable networks. Research into how these factors will affect the longevity of blockchain infrastructures will be an important area of future academic research.

12. CONCLUSION

This research investigates and compares a number of significant blockchain and distributed ledger infrastructure platforms, encompassing their design principles, system limitations, and levels of maturity within a respective space. Ethereum, Polygon, Arbitrum, Base, and Hedera were all assessed in terms of architectural attributes such as transaction fees, transaction throughput, transaction finality, accessibility to the developer ecosystem, and the models by which the consensus layer functions.

Findings reveal contemporary blockchain infrastructures exhibit varying degrees of architectural diversity. Within the blockchain ecosystem, L1's historically provide a greater degree of decentralisation and mature ecosystem, while L2 scaling solutions to improve overall transaction throughput and decrease total transaction costs, also incorporate more complexity into their design. In addition, many new technologies, referred to as sidechains, will have independent validation layers that can lead to improved performance characteristics. There are additional non-blockchain distributed ledgers that may also have high transaction throughput and quick finality than other non-blockchain technologies, only because these methods leverage alternative consensus algorithms.

This comparative analysis has discovered that there is no individualised best solution for all decentralised systems. Each network comprises a unique blend of decentralisation, scalability, performance, and maturity when determining an optimal platform. Decisions on the appropriate platform must therefore; be made on comprehensive evaluation of the operational requirements of the application for which an organisation will be building.

The multi-criteria analysis further demonstrates that platform suitability is sensitive to application priorities. Hedera achieved the strongest result for high-volume credential workloads, Polygon provided the most balanced cross-scenario profile, and Ethereum became preferable when security, decentralization, and ecosystem maturity received greater weight. These findings reinforce that platform selection cannot be based on a single performance metric and must instead reflect the complete operational and trust requirements of the target application.

As blockchain and distributed ledger technology continues to evolve, continued research into the comparative characteristics of blockchain infrastructure will be necessary to assist developers, organisations, and the research community in selecting appropriate platforms to develop and deploy decentralised applications. Continued investigation into interoperability, scalability, and privacy technologies will further shape the future landscape of distributed systems.

13. REFERENCES

[1] Leemon Baird. The hedera hashgraph consensus algorithm. *Hedera Whitepaper*, 2018.

[2] Base. Network fees, 2026. Base Documentation, accessed July 2026.

[3] Base. Transaction finality, 2026. Base Documentation, accessed July 2026.

[4] Karl Crary. Verifying the hashgraph consensus algorithm. *Carnegie Mellon University Technical Report*, 2021.

[5] Kyle Croman, Christian Decker, Ittay Eyal, et al. On scaling decentralized blockchains. In *International Conference on Financial Cryptography and Data Security*. Springer, 2016.

[6] Ethereum Foundation. Proof-of-stake, 2026. Ethereum Developer Documentation, accessed July 2026.

[7] Ethereum Foundation. Proof-of-stake frequently asked questions: Finality, 2026. Ethereum Developer Documentation, accessed July 2026.

[8] Arthur Gervais, Ghassan O. Karame, Karl Wüst, et al. On the security and performance of proof of work blockchains. *Proceedings of the ACM SIGSAC Conference*, 2016.

[9] Hedera. Hedera fee model, 2026. Hedera Documentation, accessed July 2026.

[10] Hedera. What is hedera?, 2026. Hedera Documentation, accessed July 2026.

[11] Harry Kalodner, Steven Goldfeder, Xiaoqi Chen, et al. Arbitrum: Scalable, private smart contracts. `https://offchainlabs.com/`, 2018.

[12] Arvind Narayanan, Joseph Bonneau, Edward Felten, Andrew Miller, and Steven Goldfeder. *Bitcoin and Cryptocurrency Technologies*. Princeton University Press, 2016.

[13] Offchain Labs. Inside arbitrum nitro: Fees, 2026. Arbitrum Documentation, accessed July 2026.

[14] Offchain Labs. Transaction lifecycle on arbitrum, 2026. Arbitrum Documentation, accessed July 2026.

[15] Polygon Labs. Polygon chain overview, 2026. Polygon Developer Documentation, accessed July 2026.

[16] Polygon Labs. Transaction finality, 2026. Polygon Developer Documentation, accessed July 2026.

[17] Polygon Technology. Polygon whitepaper. `https://polygon.technology/`, 2021.

[18] A. Tariq et al. Cerberus: A blockchain-based accreditation and degree verification system. In *International Conference on Blockchain Technology*. IEEE, 2019.

[19] Marko Vukolić. The quest for scalable blockchain fabric: Proof-of-work vs. bft replication. In *International Workshop on Open Problems in Network Security*. Springer, 2015.

[20] Gavin Wood. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum Yellow Paper*, 2014.

[21] Zibin Zheng, Shaoan Xie, Hong-Ning Dai, et al. An overview of blockchain technology: Architecture, consensus, and future trends. *IEEE International Congress on Big Data*, 2017.

[22] Andrej J. Zwitter and Joris Hazenberg. Decentralized network governance: Blockchain technology and the future of regulation. *Frontiers in Blockchain*, 2020.