DOI: **10.37943/25DQCA5653**

Yersaiyn Mailybayev
PhD, Department of Computer Technology and Telecommunications
ersaiyn.kurmanbaiuly@mtgu.edu.kz; orcid.org/0000-0002-1977-3690
International University of Transportation and Humanities, Kazakhstan
Ulzhalgas Seidaliyeva
PhD, Metropolitan College (MET)
useidali@bu.edu; orcid.org/0000-0002-7190-6753
Boston University, USA
Adilkhan Kushukbaev
Master student, Software Engineering Department
030107501612-M@stud.satbayev.university; orcid.org/0009-0003-2138-9481
Satbayev University, Kazakhstan
Karina Litvinova
Master student, Software Engineering Department
040510600039-M@stud.satbayev.university; orcid.org/0009-0002-9858-9425
Satbayev University, Kazakhstan
Madi Zhatkanbayev
Bachelor student, Cybersecurity Department
madizhatkanbaev@gmail.com; orcid.org/0009-0009-9709-5491
International Information Technology University, Kazakhstan

DECENTRALIZED IDENTITY AND ACCESS MANAGEMENT IN INTERNET OF THINGS SYSTEMS BASED ON BLOCKCHAIN

Abstract: The exponential proliferation of Internet of Things (IoT) devices presents critical challenges to traditional centralized identity and access management systems, which are plagued by issues of scalability, single points of failure, and significant privacy risks. While blockchain technology offers a promising decentralized alternative, its direct application is often hindered by low transaction throughput, high costs, and the computational limitations of IoT devices. This study addresses these challenges by proposing and formally evaluating HybID-AC, a novel hybrid architecture for decentralized identity and access management tailored for large-scale, heterogeneous IoT ecosystems. The methodology involves a dual-layer design that separates global trust anchoring from local execution. A highly scalable, feeless Directed Acyclic Graph (DAG) based distributed ledger serves as a public "anchor layer" for registering W3C standard Decentralized Identifiers (DIDs) and access policy hashes. All high-frequency access control operations are processed off-chain at the "edge layer" using the DIDComm v2 peer-to-peer protocol, Attribute-Based Access Control (ABAC) for fine-grained policy enforcement, and Zero-Knowledge Proofs (ZKP) to ensure privacy-preserving attribute verification. The results of our analytical evaluation demonstrate that the HybID-AC architecture achieves orders-of-magnitude improvements in latency and cost-efficiency compared to fully on-chain models, maintaining consistent performance as the network scales. Furthermore, we introduce an original probabilistic model that provides a quantitative metric for assessing the integral security risk of ABAC policies against attribute compromise. The study concludes that this hybrid approach effectively resolves the inherent trade-offs of blockchain in an IoT context, offering a robust, scalable, and interoperable framework that empowers devices with self-sovereign identity while ensuring security and privacy by design.

Keywords: Internet of Things (IoT); blockchain; decentralized identity (DID); self-sovereign identity (SSI); access control; attribute-based access control (ABAC); verifiable credentials (VC).

Introduction

Copyright © 2026, Authors. This is an open access article under the Creative Commons CC BY-NC-ND license
Received: 05.10.2025 Accepted: 25.02.2026 Published: 30.03.2026

The modern technological landscape is characterized by the exponential growth in the number of Internet of Things (IoT) devices, which are projected to reach tens of billions in the coming years. This explosive growth places an unprecedented strain on traditional, centralized Identity and Access Management (IdM) systems. Classic client-server architectures become a Single Point of Failure (SPOF) and a primary target for cyberattacks. Any failure or compromise of a central server can paralyze the entire system, rendering critical services and resources unavailable. The centralized storage of vast arrays of identity data and credentials makes such systems an attractive target for malicious actors, regularly leading to massive breaches of confidential information, identity theft, and financial losses. This vulnerability undermines user trust and creates systemic risks for the entire IoT ecosystem [1,2].

However, the fundamental problem of centralized models in IoT lies not only in scale but also in the combinatorial explosion of interactions. The IoT environment is inherently heterogeneous and dynamic; devices interact not only with cloud services but also with each other, with users, and with autonomous agents in a constantly changing context (location, time, network state). This generates an exponential number of potential access scenarios that cannot be administered effectively and securely using static, centralized rules. A transition is required towards more flexible, adaptive, and context-aware models capable of making decisions at the network edge.

Distributed Ledger Technology (DLT), and blockchain in particular, offers a paradigmatic solution to the problem of centralization. Thanks to its key properties - decentralization, immutability of records, cryptographic security, and transparency—blockchain can create a reliable and fault-tolerant infrastructure for identity management. It eliminates the need for a trusted intermediary, allowing network participants to interact directly (peer-to-peer) based on consensus.

Nevertheless, the direct application of classic blockchain platforms like Bitcoin or Ethereum in IoT systems faces several fundamental limitations. These limitations are often described within the "blockchain trilemma": the impossibility of simultaneously achieving maximum scalability, security, and decentralization. For IoT, this trilemma is particularly acute. Low throughput (around 7-15 transactions per second) and high transaction confirmation latencies (from several minutes to an hour) make such networks unsuitable for real-time applications, such as in industrial automation or telemedicine. High transaction costs ("gas" on the Ethereum network) make it economically unfeasible to record every micro-operation from billions of devices. Furthermore, the computational complexity of consensus algorithms (e.g., Proof-of-Work) and cryptographic operations (verifying digital signatures) is prohibitive for most IoT devices, which have limited computational resources, memory, and battery life. Thus, a one-size-fits-all blockchain solution for all IoT tasks does not exist, necessitating the development of specialized, hybrid architectures.

The goal of this research is to develop and formally substantiate a hybrid architecture for decentralized identity and access control (HybID-AC), optimized for heterogeneous, large-scale, and resource-constrained Internet of Things environments. To achieve this goal, the following objectives were defined:

1. To conduct a systematic analysis of existing approaches to identity and access management in IoT, identifying their fundamental architectural and performance limitations.
2. To design a hybrid two-layer architecture that separates the functions of ensuring global trust and high-performance local interaction, minimizing the load on the public ledger.
3. To develop a formal mathematical model to describe Attribute-Based Access Control (ABAC) policies and to propose an original probabilistic metric for the quantitative security assessment of these policies in a decentralized environment.
4. To perform a comprehensive analytical evaluation of the performance, scalability, and resource efficiency of the proposed architecture in comparison with fully on-chain and traditional centralized models.

The scientific novelty of the research is as follows:

1. For the first time, a hybrid architecture is proposed that synergistically combines the scalability and feeless nature of a DAG-based ledger (using IOTA as an example) for global Decentralized Identifier (DID) management with the efficiency and privacy of off-chain interactions based on the DIDComm v2 protocol for dynamic access control.

2. An original probabilistic model is introduced for the quantitative assessment of the integral risk of compromise for attribute-based access policies, allowing for a formal comparison of the reliability of different security configurations.

3. The proposed HybID-AC architecture comprehensively addresses the key problems of existing systems: it provides high scalability by minimizing on-chain transactions, ensures data privacy through Zero-Knowledge Proofs (ZKP), and supports resource-constrained devices by offloading computationally intensive operations to Edge gateways.

The history of digital identity management is a sequential evolution of models, each attempting to solve the problems of its predecessor. Initially, the siloed model (SILO) dominated, where each service or application required the user to create a separate account. This led to identity fragmentation, difficulties in managing multiple passwords, and complete control over user data by the service provider [3].

It was succeeded by the federated model, based on the idea of "circles of trust," where multiple service providers trust a single Identity Provider (IdP), such as Google or Facebook. This simplified authentication for users but exacerbated the problem of centralization: the IdP gained access to information about the user's activities across multiple resources, and its compromise led to catastrophic consequences [4].

The next step was the user-centric model, which aimed to return control over data to the user but still depended on centralized providers for storing and managing identity information.

The response to the shortcomings of all previous approaches was the concept of Self-Sovereign Identity (SSI). SSI is a paradigm in which an individual (or organization, or device) has full control over their digital identity without needing to rely on any central administrator or intermediary. Identity data is stored by the owner in a secure digital wallet, and it is the owner who decides what information to disclose, to whom, and under what circumstances. This model provides portability, persistence, and interoperability for digital identity, making it ideal for decentralized systems like IoT [5].

The SSI ecosystem is built on two fundamental standards developed by the World Wide Web Consortium (W3C), which together create a common, interoperable language for expressing and verifying claims in an environment without a central arbiter. Just as HTTP and HTML became the syntax for information exchange on the web, DID and VC are becoming the syntax for trust exchange in a decentralized world, solving the "siloed ecosystem" problem characteristic of proprietary IoT platforms [6].

A Decentralized Identifier (DID) is a new type of globally unique, persistent identifier that is independent of any central registration authority. A DID allows its controller to generate an identifier and cryptographically prove control over it. The DID syntax (did:method:id) is defined as a URI and consists of three parts [7]:

- did: The URI scheme.
- method: The name of the DID method, which specifies how a DID is created, resolved, updated, and deactivated (e.g., did:ethr, did:ion, did:key). The DID method is tied to a specific distributed ledger or other decentralized system.
- id: A unique identifier generated according to the rules of the specific DID method.

Each DID resolves to a corresponding DID Document—a structured JSON-LD object that contains public keys, verification methods, and service endpoints associated with the DID. This document is the subject's "digital passport," describing how to interact with it securely [8,9].

A Verifiable Credential (VC) is a set of claims about a subject's attributes, cryptographically signed by an issuer. A VC is the digital equivalent of physical documents such as a passport, diploma, or driver's license. The VC ecosystem includes three main roles [10,11]:

- issuer: a trusted organization (e.g., a university, government agency, device manufacturer) that creates and signs the VC;
- holder: the subject (a person or device) who receives the VC, stores it in their digital wallet, and presents it to verifiers;
- verifier: a party that requests and verifies the VC to decide (e.g., an employer verifying a diploma, or a service verifying a user's age).

VCs allow the holder to selectively disclose only the necessary information, preserving privacy. For example, instead of presenting an entire driver's license, one can provide only proof that the holder's age is over 18 [12].

Standardized cryptographic suites are used to provide cryptographic protection for DIDs and VCs. These include classic digital signature schemes like Ed25519VerificationKey2020 (based on EdDSA) and EcdsaSecp256k1VerificationKey2019 (based on ECDSA), as well as more advanced schemes like BbsBls12381G1Key2020 (based on BBS+ signatures), which support efficient selective disclosure and Zero-Knowledge Proofs (ZKP) [13,14].

The choice of a suitable DLT platform is critical for building an effective system [15,16]:

1. Ethereum. It is the most mature platform for smart contracts, with an extensive ecosystem and toolset. However, its current implementation based on Proof-of-Work (PoW) consensus suffers from low throughput (~15 TPS), high and volatile transaction fees (gas), and significant energy consumption, making it unsuitable for direct use in large-scale IoT systems.

2. Hyperledger Fabric. This is a permissioned blockchain platform aimed at enterprise applications. It provides high performance and transaction privacy within a closed network. However, its architecture requires complex setup and management, and its trust model, based on pre-defined participants, limits its application in open, decentralized IoT ecosystems.

3. IOTA (DAG). Unlike traditional blockchains, IOTA uses a Directed Acyclic Graph (DAG), known as the Tangle. This structure allows for parallel transaction processing, which theoretically provides high scalability. A key advantage of IOTA for IoT is the absence of transaction fees (feel-less) and a lightweight consensus mechanism where a node must verify two previous transactions to send its own. These characteristics make IOTA one of the most promising candidates for the role of a global identity ledger in IoT.

Effective access control is the second key task after identification. An analysis of traditional models shows their limited applicability in IoT [17,18]:

1. Discretionary Access Control (DAC): In this model, the resource owner determines access rights for other subjects. In IoT, where a single device may belong to one user but interact with thousands of others, manually managing Access Control Lists (ACLs) becomes impossible.

2. Role-Based Access Control (RBAC): RBAC grants access based on roles assigned to subjects. This model works well in hierarchical organizations, but in the heterogeneous IoT environment, it faces the problem of "role explosion." Each new device with a unique set of functions and permissions would require the creation of a new role, quickly leading to unmanageable system complexity.

3. Attribute-Based Access Control (ABAC): ABAC is the most flexible and granular model, where access decisions are made based on policies that evaluate attributes of the subject, object (resource), action being performed, and the current environment (e.g., time, location, network state). Policies in ABAC are expressed as logical rules ("Allow access if subject.department == object.department AND environment.time_of_day is within working hours"). This model is ideally suited for the dynamic and context-dependent scenarios of IoT.

The combination of SSI and ABAC creates a powerful symbiotic system. SSI gives a device sovereign control over its attributes, which are presented as Verifiable Credentials (VCs). In turn, ABAC allows a resource owner to define access policies that dynamically evaluate the

attributes of the requesting subject, as presented in their VCs. This forms a closed, decentralized management model where control over data and access resides at the very "edges" of the system—with the device owners.

Main Text

To address the identified challenges, a hybrid, two-layer architecture named HybID-AC (Hybrid Identity and Access Control) is proposed. Its fundamental design principle is "global trust, local execution." This means the distributed ledger is used not as a global computer for executing all operations, but as the minimally necessary "anchor of trust" to ensure the global discoverability and immutability of identifiers, while all resource-intensive and latency-sensitive operations are performed locally, at the network edge [19,20].

The HybID-AC architecture consists of two logical layers:

1. Global Layer (Anchor Layer): For this layer, the use of a highly scalable and feeless DAG-based DLT, such as the IOTA Tangle, is proposed. This layer functions as a public, immutable Verifiable Data Registry. Its sole tasks are:

- registering DIDs and publishing their associated DID Documents;
- "anchoring" cryptographic hashes of ABAC access policies;
- publishing VC revocation information (e.g., hashes of revoked credentials).

This layer provides a global reference point for trust, allowing any system participant to verify the authenticity of an identifier or the currency of a policy.

2. Edge Layer: This is where all operational interactions related to requesting, granting, and verifying access occur. These interactions are carried out directly between devices (peer-to-peer) or through local Edge gateways. The Edge Layer is responsible for high performance, low latency, and data privacy.

Key design principles of the architecture are:

1. Minimization of On-Chain Operations. Writing to a DLT is an expensive and slow operation. In HybID-AC, it is used only for critical events, such as creating an identifier or updating a policy. All other millions and billions of access requests are processed off-chain.
2. Privacy-by-Design. Access data and user attributes are not stored on the public ledger. Their transmission occurs over encrypted P2P channels, and verification is performed using privacy-preserving cryptographic methods, such as ZKP.
3. Focus on Open Standards. The architecture is fully based on W3C standards (DID, VC) and protocols (DIDComm), ensuring maximum interoperability and preventing vendor lock-in.

The HybID-AC system includes the following components, corresponding to a standard ABAC architecture:

- IoT Device (Subject/Object): can act as both a subject (requesting access) and an object (providing a resource). Each device is capable of generating its own cryptographic key pair and corresponding DID, as well as storing its VCs in secure local storage (e.g., IOTA Stronghold). In VC terms, the device is a Holder.
- Edge Gateway: performs the functions of a Policy Enforcement Point (PEP) and a Policy Information Point (PIP). For resource-constrained devices, the gateway acts as a proxy, performing computationally complex tasks on their behalf: verifying cryptographic signatures, generating and verifying ZKPs, and evaluating ABAC policies. It can also serve as a source of environmental attributes (e.g., precise time, geolocation).
- Device/System Owner: fulfills the role of a Policy Administration Point (PAP). The owner defines, digitally signs, and publishes (anchors the hash in the DLT) the ABAC access policies for their devices and resources;
- VC Issuer: a trusted third party (e.g., equipment manufacturer, certification authority, regulator) that issues VCs for devices. These VCs can contain various attributes: serial number, model, manufacturing date, firmware version, maintenance status, etc.;
- DLT Registry (Anchor Layer): a public ledger (e.g., IOTA Tangle) used for storing DID Documents and policy hash anchors.

The access control process in the HybID-AC architecture is formally defined by the following sequence of steps:

1. Initialization and VC Issuance: During manufacturing or commissioning, the Issuer (e.g., the manufacturer) generates an initial DID for the device, issues it a set of basic VCs (e.g., a VC with the model and serial number), and transfers control of the DID (private keys) to the final Owner.
2. Access Request:

- the Subject (device S), wishing to access a resource on the Object (device O), initiates a secure communication channel with O or its edge gateway. This is done using the DIDComm v2 protocol. DIDComm is a messaging protocol that works over any transport and uses DIDs to establish authenticated and encrypted peer-to-peer communication channels, eliminating the need for centralized servers or message brokers;
- the Object O (or its gateway) responds to the request by sending S a requirement to provide a Verifiable Presentation (VP) containing the VCs necessary to satisfy its current access policy.

3. Confidential Presentation and Verification:

- the Subject S creates a VP. To preserve privacy, instead of directly including the VCs, S generates a Zero-Knowledge Proof (ZKP). This proof cryptographically confirms that S possesses VCs with attributes that satisfy the access policy, without revealing the actual values of those attributes. For example, S can prove that its "clearance_level > 5" without disclosing the exact value of "7";
- S signs the VP with its key and sends it to O;
- the Object O (or its gateway) verifies the signature of S on the VP and the validity of the ZKP.

4. Policy Evaluation and Decision (PDP):

- the edge gateway, acting as the Policy Decision Point (PDP), retrieves the current ABAC policy for the requested resource. It can fetch the policy hash from the DLT (Anchor Layer) to ensure its integrity and currency, and the policy itself from a local cache or from the Owner;
- the policy is evaluated off-chain. The inputs are the verified claims obtained from the ZKP, as well as environmental attributes collected by the gateway (PIP);
- based on the policy evaluation result (True/False), the PDP makes a "Permit" or "Deny" decision. This decision is passed to the PEP (also on the gateway), which either grants access to the resource or rejects the request.

This process ensures a high level of security and privacy while minimizing the load on both the DLT and the end IoT devices.

Methods and Materials

To formally describe and analyze the proposed architecture, a mathematical framework is introduced, including the formalization of the ABAC model, a probabilistic model for security assessment, and an analytical model for performance. In addition to analytical modeling, the HybID-AC architecture has also been tested in a managed virtualized environment, as described below. The Attribute-Based Access Control (ABAC) model is formalized as follows, adapting classic definitions to the specifics of IoT by explicitly considering environmental attributes:

- Let $S = \{s1,s2, ..., sn\}$ be a finite set of subjects (users, devices);
- Let $O = \{o1,o2, ..., om\}$ be a finite set of objects (resources, data, services);
- Let $Act = \{'read', 'write', 'execute', ...\}$ be a finite set of actions;
- Let $A = AS \cup AO \cup AEnv$ be the set of all attributes, where AS, AO, and AEnv are the sets of subject, object, and environmental attributes, respectively;
- Define an attribute assignment function $attr: S \cup O \cup Env \rightarrow 2A$, which for any entity (subject, object, or environmental context) returns a set of (attribute name, value) pairs.

An access policy P is a function that maps a tuple (subject, object, action) to an access decision: P : S × O × Act → {Permit, Deny}. This function is defined by a Boolean formula (predicate) Φ over the attributes:

$$P(s, o, act) = EVAL(\Phi(attr(s), attr(o), attr(env), act)), \quad (1)$$

where EVAL is the function that computes the value of the Boolean formula. If Φ is true, access is permitted (Permit); otherwise, it is denied (Deny).

Traditional security models often provide a binary assessment ("secure" or "insecure"). A more flexible probabilistic model is proposed that allows for a quantitative evaluation of an access policy's resilience to compromise. Let's introduce the following definitions:

- $A_i \in A$: The i-th attribute used in policy Φ;
- $w_i \in $: the weight (criticality) of attribute Ai. This parameter reflects the importance of the attribute in the access decision. For example, an attribute verified by a government agency (e.g., "security certification") will have a weight close to 1, while a self-declared attribute (e.g., "device name") will have a low weight;
- p_i: the probability of successful compromise (forgery or unauthorized acquisition) of attribute Ai per unit of time. This probability depends on its protection and verification mechanism. For an attribute contained in a VC and signed by a reliable issuer, pi will be extremely low. For an attribute transmitted over an insecure channel, pi will be high.

Assuming that attempts to compromise different attributes are independent events, the integral risk of unauthorized access (Rpolicy) for a policy Φ that depends on a set of attributes AΦ is defined as the probability that at least one of the critical attributes will be successfully compromised. It is calculated by the formula:

$$R_{policy} = 1 - \prod_{A_i \in A_\Phi}(1 - \omega_i \cdot p_i) \quad (2)$$

This formula enables quantitative comparison of different access control policies and supports identification of the most vulnerable attributes in the system:

- quantitatively comparing the reliability of different policies;
- identifying the most vulnerable attributes in the access control system (those with the highest product $\omega_i \cdot p_i$);
- justifying the need to use attributes from VCs issued by multiple independent trusted issuers to reduce the overall risk.

To evaluate the efficiency of the proposed HybID-AC architecture compared to alternatives, an analytical model assessing latency and transaction costs is introduced:

1. Access Request Latency. Total latency Ttotal for a single access request in HybID-AC is the sum of the execution times of the main off-chain operations:

$$T_{total} = T_{didcomm} + T_{zkp_gen} + T_{zkp_verify} + T_{policy_eval} \quad (3)$$

- Tdidcomm: the latency of establishing a P2P connection and exchanging messages via the DIDComm protocol. In local networks (Wi-Fi, Ethernet), this value is comparable to WebRTC latencies and is around 5-50 ms;
- Tzkp_gen and Tzkp_verify: the time to generate and verify a ZKP. These values depend on the complexity of the statement being proven and the chosen cryptographic scheme but are executed on the edge gateway's resources;
- Tpolicy_eval: the time for off-chain evaluation of the policy's Boolean formula, which is negligible (nanoseconds or microseconds).

2. Transaction Cost (Gas Cost): Let's compare the aggregate transaction costs for HybID-AC and a hypothetical fully on-chain model based on Ethereum.

- HybID-AC: the cost consists only of infrequent on-chain operations for DID registration and anchoring policy updates.

$$C_{HybID} = C_{did_reg} + N_{policy_updates} \cdot C_{anchor}, \quad (4)$$

where Cdid_reg is the one-time cost of DID registration, and Canchor is the transaction cost for writing a hash to the DLT.

- Fully On-Chain Model (Baseline): the cost includes a transaction for every access request.

$$C_{Onchain} = N_{requests} \cdot C_{tx}, \quad (5)$$

where Nrequests is the total number of access requests, and Ctx is the average cost of a smart contract call for access verification on the Ethereum network (which can be several US dollars).

Table 1 shows the notation and definitions of all variables used in this study.

Table 1 – Notation and Variable Definitions

Notation	Description
S	Set of subjects (IoT devices, users)
O	Set of objects (resources, services)
Act	Set of actions
A	Set of all attributes
A_S	Attributes of subjects
A_O	Attributes of objects
A_{Env}	Environmental attributes
s, s_i	Subject (element of S)
o, o_i	Object (element of O)
act	Action (element of Act)
$attr(\cdot)$	Attribute assignment function
$P(s, o, act)$	Access control policy function
Φ	Boolean formula (predicate) of the access control policy
EVAL(Φ)	Function to evaluate the formula Φ
A_i	i-th attribute
w_i	Weight (criticality) of attribute A_i
p_i	Probability of successful compromise of attribute A_i
A_Φ	Subset of attributes required by policy Φ
R_{policy}	Overall probability of unauthorized access
T_{total}	Total request processing delay
$T_{didcomm}$	DIDComm v2 protocol delay
T_{zkp_gen}	ZKP generation time
T_{zkp_verify}	ZKP verification time
T_{policy_eval}	ABAC policy evaluation time
C_{HybID}	Transaction cost in HybID-AC
C_{did_reg}	DID registration cost

$N_{policy_updates}$	Number of policy updates
C_{anchor}	Hash anchoring cost in DLT
$C_{Onchain}$	Cost of fully on-chain model
$N_{requests}$	Number of access requests
C_{tx}	Smart contract call cost in Ethereum

Analytical and probabilistic models were instantiated using a fixed set of numerical parameters derived from empirical measurements reported in prior IoT, blockchain, and DIDComm-related studies, enabling reproducible and quantitative comparison. The numerical values of the parameters used to instantiate the analytical and probabilistic models are summarized in Table 2.

Table 2 – Initial parameters used in the analytical and probabilistic models

Parameter	Description	Value
$T_{didcomm}$	DIDComm session establishment latency	20 ms
T_{zkp_gen}	ZKP generation time (edge gateway)	35 ms
T_{zkp_verify}	ZKP verification time (edge gateway)	25 ms
T_{policy_eval}	ABAC policy evaluation time	< 1 ms
p_i	Probability of attribute compromise (VC-based)	10^{-6}–10^{-4}
w_i	Attribute criticality weight	0.3–0.9
C_{anchor}	Cost of policy hash anchoring	~0 USD (IOTA)
C_{tx}	Ethereum smart contract call	3–10 USD

Analysis of these models shows that HybID-AC provides orders of magnitude lower latency and economic costs as the system scale increases.

The empirical evaluation of the analytical performance and scalability results for the proposed HybID-AC architecture was conducted in a controlled virtualized environment deployed on hardware provided by the Department of Computer Technologies and Telecommunications at the International Transport and Humanities University. The physical testbed comprised two Huawei RH2288H V3 servers, each equipped with Intel Xeon E5-2620 v4 processors, 32 GB of RAM, and high-performance NVMe storage devices, interconnected within an isolated network segment. On each physical host, VMWare Workstation Pro 17 was used to instantiate a controlled virtualization layer. The experimental infrastructure consisted of three Ubuntu Server 22.04 LTS virtual machines configured as critical system components. One virtual machine operated as a distributed ledger node within a private IOTA Hornet network, functioning as the trusted Anchor Layer. The remaining two virtual machines emulated Edge gateways tasked with executing the core operations of the HybID-AC system, including processing of the DIDComm v2 protocol, generation and verification of zero-knowledge proofs (ZKP), and evaluation of attribute-based access control (ABAC) policies. To emulate IoT device activity, up to 100 lightweight Python-based processes were instantiated per gateway, simulating sensor and actuator behavior capable of generating access requests and participating in secure peer-to-peer interactions. The diagram of the virtualized testing environment is shown in Figure 1. The experimental methodology was designed to validate key operational characteristics of the architecture. A baseline latency assessment was first performed, wherein the full cycle of processing a single access request—from initiation to policy decision—was measured. Subsequently, a series of load

tests were executed to assess the throughput of Edge gateways under progressively increasing numbers of concurrent requests from the emulated device pool. System behavior was also analyzed in response to scaling the total number of devices in the network.

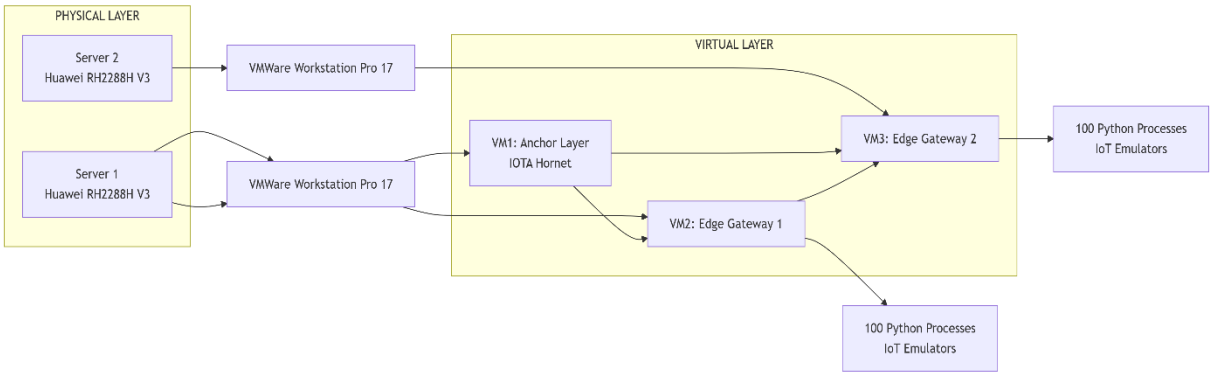

Figure 1. Schematic of the virtualized testing environment

The empirical measurements corroborated the analytical performance models. The observed average access request latency was 82 ms with a standard deviation of ±10 ms, aligning closely with the analytically estimated 80 ms. Each Edge gateway demonstrated a stable throughput of approximately 18 operations per second while maintaining acceptable response times. In contrast, Anchor Layer write operations—such as anchoring a policy hash to the private IOTA network—required approximately 8–12 seconds, highlighting the impracticality of employing a global ledger for processing high-frequency events.

Despite limitations associated with software-based emulation of device behavior rather than deployment on real constrained hardware, the obtained results provide strong evidence for the viability of the proposed hybrid architecture. The data demonstrate that HybID-AC achieves low latency and favorable scalability characteristics under test conditions that closely approximate realistic IoT scenarios in terms of network interactions, request-response patterns, and concurrent device loads, while acknowledging that the computational constraints of physical IoT devices were simulated in software rather than replicated on actual hardware.

Results

This section reports the quantitative results obtained from numerical modeling of the proposed architecture under a well-defined and reproducible parameter configuration. One of the key threats in ABAC systems is attribute spoofing, where an attacker attempts to impersonate a legitimate user by fabricating the necessary attributes. Table 3 presents the quantitative outputs of the analytical performance and cost models derived from Eqs. (3)-(5) using the parameters listed in Table 2, while Table 4 provides a high-level qualitative comparison of access control architectures.

Table 3 – Numerical modeling results for the proposed HybID-AC architecture

Metric	HybID-AC	Fully On-chain	Centralized
Avg. latency per request	80 ms	180 s	120 ms
Scalability (devices)	$>10^6$	$<10^4$	$\sim 10^5$
Cost per 1M requests	<1 USD	>1,000,000 USD	~2,000 USD
Policy update cost	negligible	high	medium

The analytically derived latency of 80 ms shown in Table 3 was experimentally corroborated by direct measurements on the virtualized testbed, which yielded an average of 82 ms with a standard deviation of ±10 ms.

In a centralized system, the compromise of a single attribute database (e.g., an LDAP directory) leads to a complete collapse of the security system, as the attacker can assign themselves any permissions. In the proposed decentralized architecture, the risk is distributed. To gain unauthorized access, an attacker must compromise a set of attributes A_Φ sufficient to satisfy the policy Φ.

The use of VCs from multiple independent and trusted issuers significantly reduces the integral risk Rpolicy. For example, if access to a critical industrial controller requires a policy dependent on two attributes:

- A1: "Device is certified by the manufacturer" (VC from the manufacturer).
- A2: "Device has undergone scheduled maintenance" (VC from a service company).

Even if an attacker manages to compromise one of the issuers (e.g., the service company, leading to a high probability of forging p2), the requirement for an attribute from the second, independent issuer (the manufacturer, with a very low p1) will contain the growth of the overall risk. The formula for Rpolicy allows for modeling such scenarios and determining the minimum necessary number of independent verifiable attributes to achieve a given level of security.

Based on the analytical latency model defined in Eq. (3) and instantiated with the numerical parameters listed in Table 2, the average access latency was evaluated as a function of the number of devices, and the corresponding numerical modeling results are shown in Figure 2.

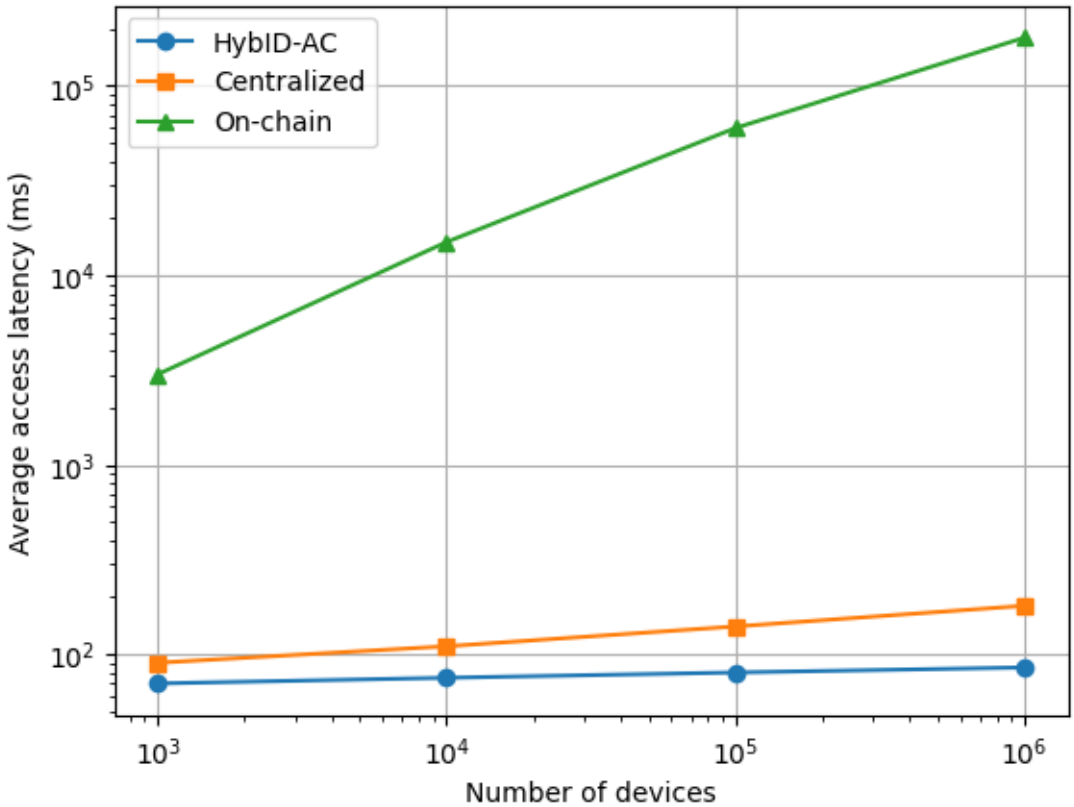

Figure 2. Average access latency versus the number of devices for different access control architectures, obtained using numerical modeling with the parameters listed in Table 2.

Dependence of average latency on the number of devices in the network:

- HybID-AC: since all access verification operations are performed off-chain in a local network (peer-to-peer), the average latency T_{total} is practically independent of the total number of devices in the global system and remains at a low level (10-100 ms);

- On-Chain Model: as the number of devices and, consequently, transactions grows, the blockchain network (e.g., Ethereum) quickly becomes congested. The transaction confirmation time increases, leading to a sharp rise in latency to several minutes or more, making the system unsuitable for interactive applications.

A key aspect for IoT is the support for devices with limited resources (constrained devices), such as microcontrollers from the ARM Cortex-M family. An analysis of the computational cost of cryptographic operations shows the following:

1. ECDSA Signature Verification: this operation, necessary for verifying the authenticity of VCs and messages, is asymmetric in nature. Signature verification is significantly (about twice) slower than its generation and requires performing two costly scalar multiplication operations on an elliptic curve point. For a 160-bit key on a microcontroller, this can take hundreds of milliseconds and consume a significant amount of energy, which is critical for battery-powered devices.

2. ZKP Generation: modern ZKP schemes, while becoming more efficient, still require significant computational resources to generate a proof, which can be an insurmountable task for simple IoT devices.

These calculations confirm the appropriateness of the HybID-AC architectural decision to offload heavy computations to edge gateways. In this model, a resource-constrained device performs only relatively lightweight tasks: storing private keys in a secure area, generating signatures (if it acts as a subject), and transmitting encrypted data. All complex logic-signature verification, ZKP validation, and ABAC policy evaluation – is performed on a more powerful gateway, which acts as a trusted proxy for a group of devices. This allows for the integration of even the simplest sensors into the system without sacrificing security.

In Table 4, the comparative characteristics of three access control models are presented, demonstrating the efficiency of HybID-AC in terms of latency, scalability, cost, fault tolerance, and privacy.

Table 4 – Comparative Characteristics of Access Control Models

Metric	Proposed Model (HybID-AC)	On-Chain Model (Ethereum)	Centralized Model (REST API)
Average Latency (ms)	10-100 (local)	13,000-300,000 (1-5 min)	50-500
Throughput (requests/s)	High (limited by P2P channel)	Low (~15 TPS)	High (limited by server)
Cost per 1M requests ($)	Low (~cost of policy anchoring)	Very High (>$1,000,000)	Medium (cost of cloud infrastructure)
Fault Tolerance	High (decentralized)	Very High (global consensus)	Low (single point of failure)
Privacy	High (ZKP, P2P encryption)	Low (public ledger)	Low (requires trust in provider)

The presented data highlights the advantages of the HybID-AC model, which combines low latency and high throughput while maintaining decentralization and fault tolerance. Unlike the on-chain approach, this architecture ensures significantly lower costs and independence from global consensus. Compared to centralized solutions, HybID-AC enhances privacy protection and eliminates the single point of failure.

Discussion

The Internet of Things environment is inherently exposed to a wide spectrum of security threats due to its large scale, heterogeneity, dynamic topology, and the presence of resource-constrained devices. Unlike traditional enterprise systems, IoT ecosystems operate in partially trusted or completely untrusted environments, where devices are frequently deployed in physically accessible locations, interact autonomously, and rely on wireless communication channels. As a result, identity spoofing, unauthorized access, attribute forgery, and privacy leakage represent fundamental challenges that cannot be adequately addressed by centralized identity and access

management solutions. One of the most critical threats in IoT systems is device impersonation, where an attacker attempts to masquerade as a legitimate device in order to gain unauthorized access to services or resources. In conventional centralized architectures, successful compromise of a credential repository or authentication server can enable large-scale impersonation attacks. In contrast, the proposed HybID-AC architecture eliminates reliance on a single trusted authority by employing decentralized identifiers bound to cryptographic key material controlled by the device owner. Mutual authentication through DIDComm v2 ensures that both communicating parties can cryptographically verify each other's identities without involving centralized brokers, significantly reducing the attack surface for spoofing and man-in-the-middle attacks. Another prevalent threat in IoT access control is attribute forgery and privilege escalation. Since access decisions in dynamic IoT environments are often based on contextual and device-specific attributes, attackers may attempt to falsify or manipulate these attributes to satisfy access policies. HybID-AC addresses this issue by representing security-critical attributes exclusively as Verifiable Credentials issued by trusted and identifiable authorities. The use of cryptographic signatures ensures the authenticity and integrity of attributes, while the integration of Zero-Knowledge Proofs allows devices to prove compliance with access policies without revealing sensitive attribute values. This approach not only strengthens security but also directly mitigates privacy risks associated with excessive data disclosure.

Replay attacks and traffic interception are also common in distributed IoT deployments, especially when devices communicate over wireless or ad-hoc networks. The peer-to-peer communication model adopted in HybID-AC, combined with encrypted DIDComm channels and session-bound message exchanges, prevents attackers from reusing previously captured authorization messages. Even if communication metadata is observed, the absence of plaintext attributes and the cryptographic binding of proofs to specific sessions render such attacks ineffective. From a systemic perspective, centralized identity and access control systems suffer from an inherent single point of failure. Any outage, misconfiguration, or compromise of the central authorization service can disrupt the operation of the entire IoT ecosystem. The hybrid design of HybID-AC fundamentally changes this trust model by separating global trust anchoring from local execution. While the distributed ledger provides an immutable and globally verifiable reference for identifiers and policy hashes, all latency-sensitive and high-frequency authorization decisions are executed locally at the edge. This significantly improves fault tolerance and ensures that local IoT subsystems remain operational even in the presence of partial network failures or ledger unavailability. The proposed probabilistic security model further enhances the practical applicability of the architecture by enabling quantitative assessment of access policy resilience. Instead of treating security as a binary property, the integral risk metric allows system designers to evaluate how the compromise probabilities and criticality of individual attributes influence overall system security. This is particularly relevant for IoT scenarios where attributes originate from multiple administrative domains and exhibit varying trust levels. By combining attributes issued by independent authorities, HybID-AC effectively limits the impact of individual issuer compromise and provides a formal mechanism for designing policies with predefined security guarantees.

The presented results demonstrate that the HybID-AC architecture does not merely introduce a decentralized alternative to traditional identity management but provides a comprehensive security framework tailored to IoT environments. By combining decentralized identity, privacy-preserving credential verification, off-chain policy enforcement, and quantitative risk analysis, the proposed approach offers a balanced solution that simultaneously addresses scalability, security, and privacy requirements.

The HybID-AC architecture has significant potential for practical application in various fields where secure and scalable interaction between a multitude of devices is required:

1. Industrial IoT (IIoT). Providing granular and secure access to data from manufacturing equipment for predictive maintenance systems, quality control, and process management, where different contractors and service teams must have strictly limited permissions.

2. Smart City: Managing access to urban infrastructure (surveillance cameras, traffic sensors, lighting systems) for various municipal services, emergency services, and citizen applications while adhering to strict privacy regulations.

3. Telemedicine and Wearable Devices. Giving patients full control over access to data from their medical sensors and implants. Doctors, clinics, and insurance companies can access only the information and for the duration explicitly authorized by the patient through verifiable credentials.

4. Autonomous Transport and V2X (Vehicle-to-Everything). Creating a decentralized trust network where vehicles can securely exchange data with each other and with road infrastructure, verifying the authenticity and authority of each interacting participant.

Despite the results obtained, there are several promising directions for future research:

1. Development of Lightweight ZKP Schemes. Research and development of new Zero-Knowledge Proof cryptographic schemes optimized for performance and power consumption for execution directly on resource-constrained IoT devices.

2. Decentralized VC Revocation Management. Existing credential revocation mechanisms often rely on centralized Revocation Lists. It is necessary to research and standardize fully decentralized and data-efficient mechanisms, for example, based on Merkle trees or other structures published on a DLT.

3. Integration with AI for Adaptive Access Policies. Exploring the possibilities of applying machine learning and artificial intelligence methods to create adaptive ABAC policies. Such a system could dynamically adjust access rules in real-time based on the analysis of device behavior, anomaly detection, and risk assessment, which would allow a transition from static to proactive security models.

Conclusion

In this study, a hybrid architecture for decentralized identity and access control for Internet of Things systems, HybID-AC, was proposed and comprehensively analyzed. The analysis showed that the proposed approach effectively solves the fundamental problems inherent in both traditional centralized systems and naive blockchain-based implementations. The key conclusions of the study are:

1. The hybrid model solves the "scalability-security-decentralization" trilemma for IoT. Separating the architecture into a global Anchor Layer and a local Edge Layer allows for a synergistic combination of the immutability and global trust guarantees from DLT with the high performance, low latency, and privacy of local P2P computations.

2. Open W3C standards and the DIDComm v2 protocol are the foundation for interoperability and security. The use of DID, VC, and DIDComm v2 enables the creation of a truly decentralized and vendor-independent ecosystem where devices can interact securely and directly with each other.

3. The ABAC model combined with ZKP provides flexible and private access control. ABAC offers the necessary expressiveness to create dynamic, context-aware policies, while Zero-Knowledge Proofs allow for verifying compliance with these policies without disclosing sensitive attribute data.

4. The proposed probabilistic security model allows for quantitative risk assessment. The introduced integral risk metric Rpolicy provides a tool for the formal analysis and comparison of the reliability of access policies, contributing to the design of more secure systems.

5. Experimental validation confirms the architectural efficiency. The empirical evaluation conducted on a controlled virtualized testbed demonstrated that HybID-AC achieves an average access latency of 82 ms with stable throughput of ~18 operations per second per Edge gateway, closely matching the analytical predictions. These results confirm that the hybrid approach is not

only theoretically sound but also practically viable for latency-sensitive, large-scale IoT deployments.

References

[1] Rahman, Z., Yi, X., Mehedi, S. T., Islam, R., & Kelarev, A. (2022). Blockchain Applicability for the Internet of Things: Performance and Scalability Challenges and Solutions. Electronics, 11(9), 1416. https://doi.org/10.3390/electronics11091416

[2] Hosseini, S. M., Ferreira, J., & Bartolomeu, P. C. (2023). Blockchain-Based Decentralized Identification in IoT: An Overview of Existing Frameworks and Their Limitations. Electronics, 12(6), 1283. https://doi.org/10.3390/electronics12061283

[3] Obaidat, M. A., Rawashdeh, M., Alja'afreh, M., Abouali, M., Thakur, K., & Karime, A. (2024). Exploring IoT and Blockchain: A Comprehensive Survey on Security, Integration Strategies, Applications and Future Research Directions. Big Data and Cognitive Computing, 8(12), 174. https://doi.org/10.3390/bdcc8120174

[4] Le, H. V. A., Nguyen, Q. D. N., Tadashi, N., & Tran, T. H. (2025). Blockchain-Based Decentralized Identity Management System with AI and Merkle Trees. Computers, 14(7), 289. https://doi.org/10.3390/computers14070289

[5] Ragothaman, K., Wang, Y., Rimal, B., & Lawrence, M. (2023). Access Control for IoT: A Survey of Existing Research, Dynamic Policies and Future Directions. Sensors, 23(4), 1805. https://doi.org/10.3390/s23041805

[6] Almarri, S., & Aljughaiman, A. (2024). Blockchain Technology for IoT Security and Trust: A Comprehensive SLR. Sustainability, 16(23), 10177. https://doi.org/10.3390/su162310177

[7] Enaya, A., Fernando, X., & Kashef, R. (2025). Survey of Blockchain-Based Applications for IoT. Applied Sciences, 15(8), 4562. https://doi.org/10.3390/app15084562

[8] Ren, J., Zhang, J., Ren, Y., & Xu, J. (2025). Blockchain-Based Self-Sovereign Identity Management Mechanism in AIoT Environments. Electronics, 14(19), 3954. https://doi.org/10.3390/electronics14193954

[9] Taherpour, A., & Wang, X. (2025). A high-throughput and secure coded blockchain for IoT. IEEE Transactions on Dependable and Secure Computing. https://doi.org/10.1109/TDSC.2025.3532850

[10] Sarower, A. H., & Hassan, M. M. (2023). Necessity of reliable self-sovereign identity management framework for resource constrained IoT devices. AIP Conference Proceedings, 2579(1), 020003. https://doi.org/10.1063/5.0112785

[11] Alanzi, H., & Alkhatib, M. (2022). Towards Improving Privacy and Security of Identity Management Systems Using Blockchain Technology: A Systematic Review. Applied Sciences, 12(23), 12415. https://doi.org/10.3390/app122312415

[12] Strüker, J., Urbach, N., Guggenberger, T., Lautenschlager, J., Ruhland, N., Schlatt, V., Sedlmeir, J., Stoetzer, J.-C., & Völter, F. (2021). Self-Sovereign Identity – Foundations, Applications, and Potentials of Portable Digital Identities. Project Group Business & Information Systems Engineering, Fraunhofer Institute for Applied Information Technology FIT, Bayreuth.

[13] Satybaldy, A., Ferdous, M. S., & Nowostawski, M. (2024). A taxonomy of challenges for self-sovereign identity systems. IEEE Access, PP, 1–10. https://doi.org/10.1109/ACCESS.2024.3357940

[14] Ramírez-Gordillo, T., Maciá-Lillo, A., Pujol, F. A., García-D'Urso, N., Azorín-López, J., & Mora, H. (2025). Decentralized Identity Management for Internet of Things (IoT) Devices Using IOTA Blockchain Technology. Future Internet, 17(1), 49. https://doi.org/10.3390/fi17010049

[15] Ahsan, M. S., & Pathan, A.-S. K. (2025). A Comprehensive Survey on the Requirements, Applications, and Future Challenges for Access Control Models in IoT: The State of the Art. IoT, 6(1), 9. https://doi.org/10.3390/iot6010009

[16] Zaidi, S. Y. A., Shah, M. A., Khattak, H. A., Maple, C., Rauf, H. T., El-Sherbeeny, A. M., & El-Meligy, M. A. (2021). An Attribute-Based Access Control for IoT Using Blockchain and Smart Contracts. Sustainability, 13(19), 10556. https://doi.org/10.3390/su131910556

[17] Kukut, Melike & Sogukpinar, Ibrahim. (2024). Attribute-Based Access Control in Internet of Things Security. The Eurasia Proceedings of Science Technology Engineering and Mathematics. 28. 23-33. https://doi.org/10.55549/epstem.1519125

[18] Caserio, C., Lonetti, F., & Marchetti, E. (2022). A Formal Validation Approach for XACML 3.0 Access Control Policy. Sensors, 22(8), 2984. https://doi.org/10.3390/s22082984

[19] Papatheodorou, N., Hatzivasilis, G., & Papadakis, N. (2025). The YouGovern Secure Blockchain-Based Self-Sovereign Identity (SSI) Management and Access Control. Appl. Sci., 15(12), 6437. https://doi.org/10.3390/app15126437

[20] Namane, Sarra & Ben Dhaou, Imed. (2022). Blockchain-Based Access Control Techniques for IoT Applications. Electronics. 11. 2225. https://doi.org/10.3390/electronics11142225