

# Privacy-Preserving Zero Trust Network Access Control System Using Chameleon Hash Functions and Fully Homomorphic Encryption\*

Ankita Jagadish, Hideaki Miyaji, and Noriaki Kamiyama<sup>†</sup>

Ritsumeikan University, Ibaraki, Osaka, Japan  
ankita.jagadish@gmail.com, {h-miyaji, kamiaki}@fc.ritsumei.ac.jp

## Abstract

Zero Trust Network Access (ZTNA) represents a significant advancement in network security by adopting a "never trust, always verify" approach. In ZTNA, users request access to the system, while verifiers determine whether access should be granted. Determining whether a user is authorized to access the system while preserving their privacy remains a challenging task. There are methods to protect the privacy of communication networks using Fully Homomorphic Encryption and Attribute-Based Encryption. However, their approach does not include a method for verifiers to detect tampering, making it insufficient to achieve ZTNA. To address this issue, we propose a practical privacy-preserving ZTNA system by simultaneously applying Chameleon Hash Functions and Homomorphic Encryption. Our system operates under a security model with certain assumptions and adversary capabilities expanded further in the paper. Our approach compares between two types of homomorphic encryption: Partial Homomorphic Encryption (PHE), which allows only addition or multiplication, and Fully Homomorphic Encryption (FHE), which imposes no restrictions on homomorphic addition and multiplication. We implement our system using Fully Homomorphic Encryption (FHE) with the Krawczyk-Rabin chameleon hash scheme, wherein the system is modeled using the Role-Operation-Target (ROT) framework and by employing a chameleon hash function for access control data, we realize a method to prevent tampering by verifiers, thereby achieving ZTNA. To validate our approach, we provide formal security analysis including theoretical security proofs and comprehensive threat resistance evaluation against various attack scenarios. We also conduct numerical evaluation of the proposed ZTNA system with twenty policy combinations to measure its real-time performance and implementation accuracy. The results are used to discuss the feasibility of the proposed approach and its potential for practical implementation in real-world scenarios.

**Keywords:** Zero Trust Network Access (ZTNA), Fully Homomorphic Encryption (FHE), Partial Homomorphic Encryption (PHE), Chameleon Hash Function (CH)

## 1 Introduction

Traditional network security models have relied heavily on perimeter-based defenses, operating under the assumption that threats primarily originate from outside organizational boundaries [1]. This paradigm has become increasingly obsolete in today's distributed computing landscape, where resources are accessed from various locations, devices, and networks. The concept of "Zero Trust" has emerged as a response to these changing dynamics, advocating for a security model that eliminates implicit trust and continuously verifies every access request regardless of its origin [2].

---

\*Proceedings of the 9th International Conference on Mobile Internet Security (MobiSec'25), Article No. 35, December 16-18, 2025, Sapporo, Japan. © The copyright of this paper remains with the author(s).

<sup>†</sup>Corresponding author

## 1.1 Zero Trust Network Access

Zero Trust Network Access (ZTNA) is the system for authentication and access control. These mechanisms reliably verify user identities and enforce precise permission specifications for operations on protected resources [3]. Importantly, ZTNA is not a single network topology or a standalone technology. It constitutes a strategic assembly of multiple principles, designed to safeguard enterprise assets [4]. Consequently, achieving access control that simultaneously identifies multiple users and devices while preserving their privacy is a challenging task. It requires the adoption of ZTNA designs and implementations that inherently incorporate privacy-preserving measures [5].

While ZTNA remains primarily a conceptual framework rather than a specific technology, its implementation requires innovative approaches to access control that go beyond conventional methods. Current attempts to implement Zero Trust principles often struggle with a fundamental contradiction: they require access to plaintext credentials, policies, and resource identifiers to perform verification, thereby creating security vulnerabilities at the very point of access control.

## 1.2 Homomorphic Encryption and Chameleon Hash functions

Homomorphic Encryption is a cryptographic technology that enables computation on ciphertexts without requiring decryption, allowing calculations to be performed on data while it remains encrypted [6, 7]. Thus, incorporating homomorphic encryption into ZTNA architectures enables the evaluation of access control policies over encrypted user attributes and contextual data [8]. However, since homomorphic encryption supports only limited operations such as addition and multiplication over numerical data, its integration into ZTNA architectures remains technically challenging.

Chameleon Hash functions, introduced by Krawczyk and Rabin at NDSS 2000 [9, 10], are powerful cryptographic primitives, particularly useful in digital signature schemes and non-interactive commitment protocols, and build upon earlier concepts of trapdoor commitments [11]. They are collision-resistant hash functions with a trapdoor property that enables only the holder of a secret key to efficiently find collisions, while preserving a uniform distribution over the output space. Incorporating the collision-resistant property of Chameleon Hash functions into ZTNA makes it possible to perform access control using a common ciphertext derived from different keys held by administrators and multiple users. This capability facilitates the design of ZTNA architectures with enhanced usability. However, implementing ZTNA systems that integrate Chameleon Hash functions remains a significant challenge.

## 1.3 Our contribution

Our research introduces a novel privacy-preserving access system that embodies Zero Trust principles through the strategic integration of Fully Homomorphic Encryption (FHE) and Chameleon Hash functions. By leveraging FHE, we achieve blind policy evaluation where access control decisions are computed without the system ever seeing user attributes, policy conditions, or resource identifiers in plaintext, fundamentally solving the plaintext exposure vulnerability inherent in traditional ZTNA systems.

However, FHE alone presents critical limitations in dynamic environments where policies require frequent updates, as traditional FHE implementations would necessitate complete re-encryption of all policies whenever changes occur, creating substantial computational overhead and potential service disruptions during policy modification periods.

The Chameleon Hash integration directly addresses these deficiencies by enabling dynamic policy integrity wherein policies can be legitimately updated by authorized administrators through controlled hash collisions while maintaining cryptographic proof of integrity and detecting any unauthorized tampering through hash verification processes.

This combination ensures that policy updates can occur efficiently without full re-encryption while providing continuous integrity assurance that pure FHE cannot deliver. Together, these technologies

realize cryptographically enforced Zero Trust eliminating trust assumptions and ensuring maximum privacy. This paper makes the following key contributions:

1. **Propose a privacy-preserving ZTNA system:** Integrates Fully Homomorphic Encryption (FHE) and Chameleon Hash to create a privacy-preserving access system that enforces Zero Trust principles without exposing any sensitive data. Thus, our system can allow verification operations to be performed on encrypted credentials and policies.
2. **Implement and evaluate our proposed system:** We implement and analyze the performance of our system using Fully Homomorphic Encryption (FHE) and Partial Homomorphic Encryption (PHE), providing optimization insights and benchmarks.
3. **Optimization strategies and Future research Opportunities:** We proposed a plausible solutions for future work on optimizing encryption-based access control mechanisms, research directions for improving efficiency and scalability in Zero Trust implementations using FHE and Chameleon Hash functions.

It is known that FHE is more computationally challenging than PHE, due to which practical deployment of FHE-based systems is still underway. In this study, we demonstrate that by appropriately tuning the weight parameter used in PHE, it is possible to construct an FHE-based system that not only outperforms its PHE-based counterpart in terms of efficiency but also operates within a practical time frame.

## 1.4 Paper organization

The remainder of this paper is organized as follows: Section 2 surveys related work in access control systems, ZTNA, Homomorphic encryption, and Chameleon Hash applications. Section 3 details our proposed architecture and implementation approach. Section 4 presents experimental results and performance evaluations. Section 5 talks about the discussions and outlines for future research. Finally, Section 6 concludes the paper outlining a concise summary of the project.

## 2 Existing research and definition

This section examines key research contributions relevant to our privacy-preserving access system that integrates Homomorphic encryption and Chameleon Hash functions to implement Zero Trust principles.

### 1. Zero Trust

The concept of Zero Trust, introduced by Forrester Research [12], advocates a "never trust, always verify" approach to security, eliminating implicit trust regardless of connection source. Google's BeyondCorp and NIST's Special Publication [13] have provided practical frameworks for implementing these principles. However, these implementations typically perform verification on plaintext data, creating privacy vulnerabilities that contradict core Zero Trust principles.

### 2. Fully Homomorphic Encryption

Fully Homomorphic Encryption (FHE) is a cryptographic technique that allows computations to be carried out on encrypted data without first decrypting it. This means sensitive information can remain confidential throughout the entire processing workflow, making FHE highly suitable for privacy-preserving applications such as cloud computing. In FHE, operations like addition and multiplication can be performed directly on ciphertexts. These operations are mathematically equivalent to those on the original plaintexts, which ensures that the final decrypted result matches what would have been obtained by computing on the unencrypted data. These equations express the idea that computations can be carried out directly on encrypted data, and the result once decrypted is the same as if the operations were performed on the original, unencrypted data [14].

$$\text{Encrypt}(m) \oplus \text{Encrypt}(n) = \text{Encrypt}(m + n) \quad \dots(1)$$

$$\text{Encrypt}(m) \otimes \text{Encrypt}(n) = \text{Encrypt}(m \times n) \quad \dots(2)$$

- (a) Equation (1): Homomorphic Addition This means that if two plaintext messages  $m$  and  $n$  are individually encrypted, their encrypted forms can be added together, and the result will decrypt to  $m+n$ . The operator  $\oplus$  represents a homomorphic addition in ciphertexts.
- (b) Equation (2): Homomorphic Multiplication Similarly, multiplying two ciphertexts will produce a new ciphertext that decrypts to  $mn$ . The operator  $\otimes$  represents a homomorphic multiplication.

### 3. Chameleon Hash

The Chameleon Hash (CH) function, introduced by Krawczyk and Rabin in NDSS 2000[10], is a special type of hash function that allows controlled collisions using a secret key. It operates with a public key and produces both a hash value  $h$  and a randomness value  $r$  for a given message  $m$ . If the secret key holder wants to modify the hash to correspond to a different message  $m'$ , they can compute a new randomness value  $r'$  such that both message-randomness pairs— $(h, r)$  and  $(h, r')$ —remain valid under the same hash function.

In the traditional CH scheme, the hash value is computed as  $h = g^m pk_{ch}^r$ , where  $g$  is a generator and the public key is defined as  $pk_{ch} = g^{sk_{ch}}$ . To adapt the hash for a new message  $m'$ , the secret key owner calculates  $r'$  such that the equation  $m + sk_{ch} \cdot r = m' + sk_{ch} \cdot r'$  holds [10].

### 4. Research Gap

While Xie et al. [15] established the theoretical foundation for combining homomorphic properties with Chameleon Hash functions, there remains a significant gap in applying these techniques to build practical privacy-preserving access control systems. Our work extends their theoretical model into a practical implementation that maintains data privacy throughout the entire verification process, thereby creating an access system that truly embodies Zero Trust principles without compromising on privacy or performance.

## 3 OUR PROPOSED SYSTEM

This section presents our novel privacy-preserving access system that integrates Homomorphic Encryption and Chameleon Hash functions to implement Zero Trust principles. We begin with a high-level overview of the system architecture, followed by detailed descriptions of each component and their interactions.

### 3.1 Overview

This section presents our novel privacy-preserving access system that integrates Homomorphic Encryption and Chameleon Hash functions to implement Zero Trust principles. We begin with a high-level overview of the system architecture, followed by detailed descriptions of each component and their interactions.

### Threat Model Definition

Our system operates under a security model with the following assumptions and adversary capabilities:

**Adversary Model:** We consider a computationally bounded adversary  $\mathcal{A}$  with the following capabilities:

- **Honest-but-Curious Server:** The access control server follows the protocol correctly but may attempt to learn information from encrypted data and intermediate computations.

- **Passive Network Observation:** The adversary can observe all network communications between clients and the access control server.
- **Limited Query Manipulation:** The adversary cannot arbitrarily modify encrypted queries in transit due to Chameleon Hash integrity verification.
- **No Cryptographic Key Access:** The adversary does not have access to FHE secret keys or Chameleon Hash private keys.

**Trust Assumptions:**

- The client-side key generation and encryption processes are secure and not compromised.
- The Chameleon Hash verification component operates in a trusted environment.
- Secure channels exist for initial key distribution and policy setup.

### 3.1.1 Security Limitations and Known Vulnerabilities

**Active Attack Vulnerability:** Our system inherits the well-known limitation that FHE schemes are vulnerable to active attacks where malicious adversaries can manipulate ciphertexts to potentially recover secret key information. In the context of our ZTNA implementation, a malicious client could theoretically craft specially designed encrypted queries that, when processed by the FHE evaluation engine, might leak information about the secret key through timing attacks, error patterns, or chosen-ciphertext analysis.

**Semi-Honest Security Limitation:** The current implementation provides security guarantees only against semi-honest (honest-but-curious) adversaries. Fully malicious adversaries who can actively manipulate the protocol execution, inject false queries, or corrupt the evaluation process are outside our current security model.

**Query Replay Attacks:** Without additional timestamp or nonce mechanisms, the system is susceptible to replay attacks where previously valid encrypted queries could be resubmitted to gain unauthorized access.

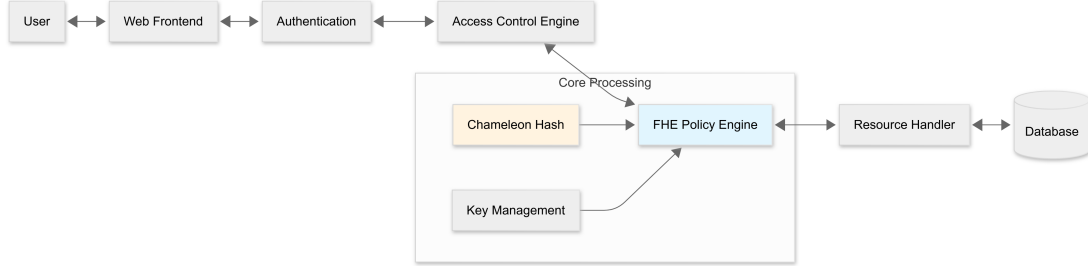


Figure 1: Flowchart of the proposed privacy preserving ZTNA system using Chameleon Hash and Homomorphic Encryption

## 3.2 System Architecture Overview

The proposed Zero Trust Network Access (ZTNA) system employs a hybrid cryptographic approach combining Fully Homomorphic Encryption (FHE) for privacy-preserving policy evaluation and Chameleon Hash functions for integrity verification. The architecture consists of six primary components operating in a layered fashion to ensure both security and operational efficiency.

The system architecture follows a modular design where each component serves a specific security function while maintaining interoperability with other modules. The frontend layer provides user interaction capabilities through a Flask-based web interface, while the backend processing occurs through encrypted computation modules that never expose sensitive policy information in plaintext.

### 3.3 Mathematical Notation and Definitions

Before detailing the FHE-based policy evaluation framework, we establish the mathematical notation used throughout this methodology:

Table 1: Mathematical Notation for FHE-based Policy Evaluation Framework

Symbol	Description
<b>Encryption and Cryptographic Operations</b>	
$Enc_{FHE}(x)$	Fully Homomorphic Encryption of data $x$
$Enc_{AES}(x)$	Advanced Encryption Standard encryption of data $x$
$pk_{FHE}, sk_{FHE}$	FHE public and secret key pair
$CH(P, r)$	Krawczyk-Rabin Chameleon Hash: $CH(m, r) = H(m) \cdot r^e \bmod n$
$tk$ (or $d$ )	RSA private exponent serving as trapdoor key
$n$	RSA modulus for chameleon hash operations
$e$	RSA public exponent for chameleon hash operations
<b>Policy and Attribute Notation</b>	
$P_i$	Access control policy $i$
$condition_i$	Boolean function defining access conditions for policy $i$
$action_i$	Set of permitted operations specified by policy $i$
$attr_i$	User attribute $i$
$A$	Complete set of user attributes
$r_i$	Random value used in Chameleon Hash computation
$h_i$	Computed hash value for policy $i$
<b>Homomorphic Operations</b>	
$HomEQ(x, y)$	Homomorphic equality comparison
$HomPrefixMatch(x, y)$	Homomorphic prefix matching operation
$\otimes$	Homomorphic multiplication (logical AND for binary values)
$\oplus$	Homomorphic addition (logical OR for binary values)
$HomEval(A, P)$	Homomorphic evaluation of policy $P$ against attributes $A$
<b>System Components</b>	
$StoredPolicy_i$	Complete stored policy structure including encrypted policy and hash
$match_i$	Boolean result of policy $i$ evaluation
$decision$	Final access control decision (grant/deny)

### 3.4 FHE-Based Policy Evaluation Framework

#### 3.4.1 Theoretical Foundation

The core innovation of this system lies in its ability to evaluate access control policies without revealing the policy contents, user attributes, or resource characteristics to the processing engine. This is achieved through a structured homomorphic evaluation framework based on the Role-Operation-Target (ROT) model.

#### 3.4.2 Policy Decomposition Strategy

Each access control policy is decomposed into three fundamental components:

- **Role Component (R):** Defines the user identity, organizational role, or attribute requirements necessary for access. This component encapsulates both static attributes (job title, department) and dynamic attributes (current project assignments, temporary permissions).

- **Operation Component (O):** Specifies the permitted actions on the target resource, including read, write, modify, delete, or execute permissions.
- **Target Component (T):** Identifies the specific resources, resource categories, or resource patterns that the policy governs.

### 3.5 Mathematical Framework and Implementation

We describe the Mathematical Framework.

**Attribute Encryption and Storage.**  $Enc_{FHE}(attr_i)$  is computed as

$$Enc_{FHE}(attr_i) = FHE.Encrypt(pk_{FHE}, attr_i).$$

Also,  $Enc_{FHE}(A)$  is computed as

$$Enc_{FHE}(A) = \{Enc_{FHE}(a_1), Enc_{FHE}(a_2), \dots, Enc_{FHE}(a_n)\}.$$

**Policy Representation and Storage.**

$$\begin{aligned} P_i &= (condition_i, action_i) \\ h_i &= CH(P_i, r_i) \\ StoredPolicy_i &= (Enc_{AES}(P_i), h_i, r_i) \end{aligned}$$

Where  $condition_i$  represents a boolean function on user attributes,  $action_i$  specifies permitted operations, and  $r_i$  serves as the random value for chameleon hashing.

**Homomorphic Evaluation Process.**

**Phase 1: Component Encryption and Verification** For each relevant policy  $P_i$ , the system first verifies policy integrity:

$$CH(P_i, r_i) \stackrel{?}{=} h_i$$

User request components are encrypted using the same FHE scheme to ensure homomorphic compatibility.

**Phase 2: Homomorphic Comparison Operations** The system performs encrypted comparisons for each ROT component:

$$\begin{aligned} Enc_{FHE}(R_{result}) &= HomEQ(Enc_{FHE}(User_{Role}), Enc_{FHE}(Policy_{Role})) \\ Enc_{FHE}(O_{result}) &= HomEQ(Enc_{FHE}(Requested_{Op}), Enc_{FHE}(Policy_{Op})) \\ Enc_{FHE}(T_{result}) &= HomPrefixMatch(Enc_{FHE}(Requested_{Target}), Enc_{FHE}(Policy_{Target})) \end{aligned}$$

Then,  $Enc_{FHE}(match)$  can be written as Equation (1).

$$\begin{aligned} Enc_{FHE}(match) &= HomEQ(Enc_{FHE}(attr), const) \\ &\otimes HomPrefixMatch(Enc_{FHE}(resource), prefix) \end{aligned} \tag{1}$$

**Phase 3: Result Aggregation** The final access decision combines all component results through homomorphic AND operations:

$$Enc_{FHE}(Final_{Decision}) = Enc_{FHE}(R_{result}) \otimes Enc_{FHE}(O_{result}) \otimes Enc_{FHE}(T_{result})$$

For multiple applicable policies, results are combined using homomorphic OR operations:

$$Enc_{FHE}(decision) = CombineResults(\{Enc_{FHE}(match_1), Enc_{FHE}(match_2), \dots\})$$

The final decision is decrypted in a secure environment:

$$decision = FHE.Decrypt(sk_{FHE}, Enc_{FHE}(decision))$$

### 3.5.1 Multi-Policy Evaluation

Each policy undergoes independent ROT evaluation, and encrypted decisions are aggregated:

$$Enc_{FHE}(decision) = Enc_{FHE}(match_1) \oplus Enc_{FHE}(match_2) \oplus \dots$$

### Illustrative Example

To illustrate the system operation, consider an HR Manager requesting read access to employee payroll data. The system processes this request as follows:

- **Input:** User="HR\_Manager", Resource="Employee\_Payroll\_Q1\_2025", Operation="Read"

The Fully Homomorphic Encryption (FHE) evaluation proceeds:

$$\begin{aligned} Enc(R_{result}) &= HomEval.Compare(Enc("HR\_Manager"), Enc("HR\_Manager")) = Enc(1) \\ Enc(F_{result}) &= HomEval.Match(Enc("Employee\_Payroll\_*"), \\ &Enc("Employee\_Payroll\_Q1\_2025")) = Enc(1) \\ Enc(O_{result}) &= HomEval.Check(Enc("Read|Write"), Enc("Read")) = Enc(1) \\ Enc(match) &= Enc(1) \otimes Enc(1) \otimes Enc(1) = Enc(1) \end{aligned}$$

- **Result:** Access Granted

### 3.6 Chameleon Hash Integration

The system implements the Krawczyk-Rabin chameleon hash scheme based on the RSA cryptosystem. For a message  $m$  and randomness  $r$ , the chameleon hash is computed as:

$$CH(m, r) = H(m) \cdot r^e \bmod n$$

where  $n = p \cdot q$  is an RSA modulus,  $e$  is the public exponent,  $d$  is the private exponent satisfying  $ed \equiv 1 \pmod{\phi(n)}$ , and  $H(\cdot)$  is a cryptographic hash function (SHA-256).

The trapdoor property allows computing collisions: given  $(m_1, r_1)$  and  $m_2$ , one can find  $r_2$  such that  $CH(m_1, r_1) = CH(m_2, r_2)$  using the private key  $d$ .

#### 3.6.1 Dynamic Policy Update Mechanism

**Policy Update Protocol:**

1. **Collision Parameter Computation:**

$$r'_i = r_i \cdot \left( \frac{H(P_i)}{H(P'_i)} \right)^d \bmod n$$

such that:

$$CH(P_i, r_i) = CH(P'_i, r'_i) = H(P_i) \cdot r_i^e = H(P'_i) \cdot (r'_i)^e \pmod{n}$$

2. **Policy Replacement:**

$$(Enc_{AES}(P_i), r_i) \rightarrow (Enc_{AES}(P'_i), r'_i)$$

3. **Hash Preservation:** The hash value  $h_i$  remains unchanged.

### 3.7 Authentication and Session Management

#### 3.7.1 User Authentication Pipeline

The authentication module operates independently, supporting multiple authentication methods and enforcing policy evaluation only for authenticated users.



### 3.7.2 Encrypted Session Context

Each session includes encrypted context information (role, project, time-based permissions), which influences encrypted access decisions.

## 3.8 Performance Optimization Strategies

### 3.8.1 Preprocessing Techniques

- Policy templates are pre-encrypted and cached.
- Frequently used roles and resource patterns are preprocessed.

### 3.8.2 Evaluation Caching

Encrypted results of previous identical evaluations are cached securely to reduce computation time.

## 4 Evaluation and Performance Analysis

This section presents a comprehensive performance evaluation of our FHE and Chameleon Hash-based access control system. Experiments were conducted on an Intel i5-10300H CPU @ 2.5GHz using the Microsoft SEAL library with 128-bit security parameters. Our evaluation focuses on two key aspects: system performance characteristics and comparison with Partial Homomorphic Encryption (PHE).

### 4.1 System Performance Analysis

We evaluated our FHE+CH hybrid approach by measuring the processing time for policy evaluations ranging from 1 to 20 concurrent access requests. Each policy evaluation computes our core formula:

$$P(\text{role}, \text{operation}, \text{target}) = R(\text{role}) \times O(\text{operation}) \times T(\text{target})$$

#### 4.1.1 Evaluation Dataset Configuration:

Our benchmark employs a realistic enterprise access control scenario with the following parameters:

- **Roles:** 3 distinct roles (HR, Finance, Admin)
- **Resources:** 3 resource categories (payroll, budget, company\_records)
- **Operations:** 3 operation types (read, write, delete)
- **Policy Set:** 20 diverse access control policies covering different role-resource-operation combinations

#### 4.1.2 Benchmark Policy Examples:

The 20 concurrent evaluations represent realistic access control scenarios such as:

- HR personnel requesting read access to payroll data
- Finance staff attempting write operations on budget documents
- Admin users performing delete operations on company records
- Cross-departmental access attempts (e.g., HR accessing budget data)
- Unauthorized operation attempts (e.g., non-Admin users requesting delete permissions)

Each evaluation in the "20 concurrent" benchmark represents a simultaneous access control decision, where multiple users across different departments request various operations on different resources. This models realistic enterprise scenarios where multiple access requests occur simultaneously during peak operational periods.




Figure 2: Hybrid Performance graph of Fully Homomorphic Encryption and Chameleon Hash for Dynamic access control system

Figure 2 illustrates the performance characteristics of our FHE+CH system across multiple benchmark runs. The x-axis represents the number of concurrent policy evaluations (1 to 20), while the y-axis shows processing time in milliseconds. The results demonstrate the scaling behavior and system-level effects encountered in real-world homomorphic encryption deployments.

#### 4.1.3 Performance Range and Scaling:

Processing times range from approximately 26–33 ms for single policy evaluation to 498–683 ms for 20 concurrent evaluations. The three benchmark runs show consistent scaling patterns with performance variations attributable to system-level effects and cryptographic operation complexity.

#### 4.1.4 Performance Variation Analysis:

The non-perfectly-linear scaling observed across runs can be attributed to several factors:

- **Thermal effects:** Sustained FHE operations generate significant CPU load, leading to thermal throttling in extended benchmarks. This explains the progressive performance degradation observed in consecutive test runs.

- **Cryptographic operation complexity:** Different policy combinations require varying computational paths through HomEQ and HomPrefixMatch operations, resulting in non-uniform processing times even for the same number of concurrent evaluations.
- **System resource management:** Memory allocation patterns, garbage collection cycles, and OS scheduling effects introduce performance variability typical of resource-intensive cryptographic applications.
- **FHE noise management:** While noise accumulation occurs within individual evaluations rather than across runs, the complexity of noise management operations varies based on ciphertext characteristics and operation sequences.

**Real-World Deployment Implications:** The performance characteristics demonstrate several important considerations for FHE deployment in enterprise environments:

- **Thermal management:** Production systems require adequate cooling and thermal throttling considerations for sustained FHE workloads.
- **Performance predictability:** The 2–3x variation in processing times necessitates conservative capacity planning and SLA definitions.
- **Concurrent evaluation efficiency:** The sub-linear scaling for multiple concurrent evaluations (20 evaluations taking less than 20x single evaluation time) demonstrates reasonable efficiency for batch processing scenarios.

Despite the computational overhead and performance variability, these processing times remain acceptable for high-security enterprise environments where privacy preservation through homomorphic encryption is prioritized over raw performance. The consistent sub-second response times for up to 20 concurrent policy evaluations support practical deployment in access control systems with moderate to high security requirements.

**Benchmark Methodology:** All measurements were conducted on a Windows system using Python 3.x with TenSEAL library. Multiple runs were performed to capture thermal effects and system variability. The presented data represents realistic operational conditions including system warm-up effects and thermal considerations typical of production FHE deployments.

## 4.2 Comparison with Partial Homomorphic Encryption (PHE)

To enable comparison with PHE, we implemented an alternative approach using weighted addition:

$$P = w_1 \cdot R + w_2 \cdot O + w_3 \cdot T$$

This adaptation was necessary because PHE schemes (such as Paillier cryptosystem) cannot perform ciphertext-to-ciphertext multiplication required by our original formula  $P = R \cdot O \cdot T$ , unlike FHE which supports both homomorphic addition and multiplication operations needed for our HomEQ and HomAND computations. The weighted linear combination provides a computationally feasible alternative that leverages PHE’s additive homomorphic properties while maintaining meaningful policy evaluation semantics.

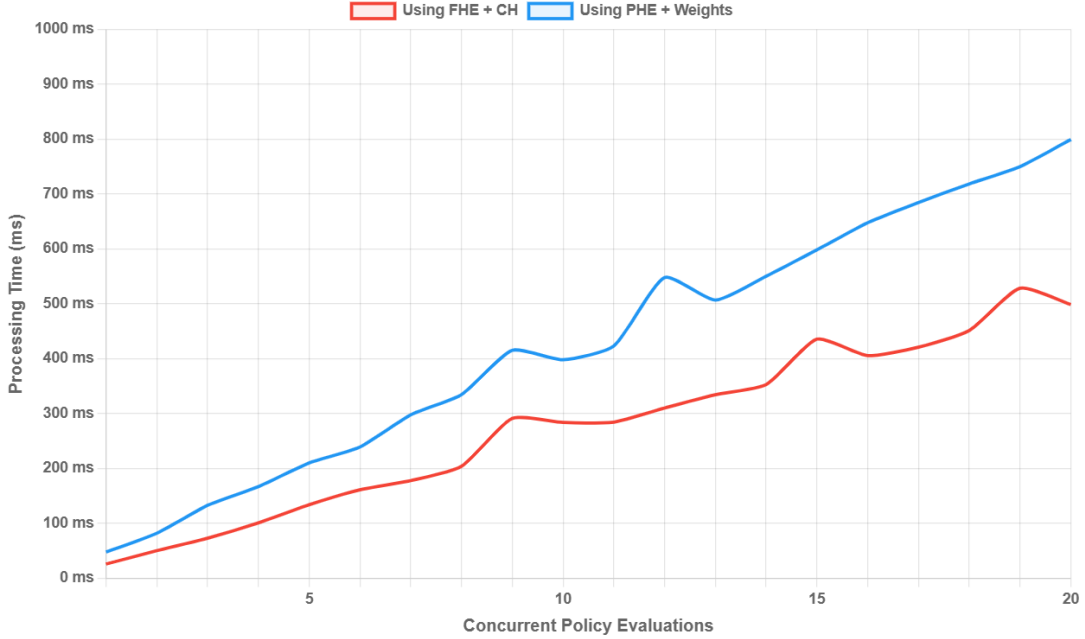


Figure 3: Performance comparison of Partial Homomorphic Encryption and Fully Homomorphic Encryption

Figure 3 presents a comparison between our FHE approach and a PHE-based implementation using weights. Contrary to expectations, our FHE + Chameleon Hash approach demonstrates superior performance, achieving approximately 1.5 to 2 times faster execution than the PHE + Weights implementation.

The superior performance of FHE over PHE in our implementation occurs because PHE requires multiple complex computational operations to approximate the Boolean logic that FHE performs natively. Specifically, our Zero Trust access control system demands precise Boolean conjunctions (AND operations) for policy evaluation, which FHE handles directly through multiplicative homomorphic operations.

In contrast, the PHE implementation must simulate these Boolean operations through weighted arithmetic approximations, requiring additional homomorphic addition operations, weight calibration computations, and threshold determination processes for each policy evaluation.

This computational overhead, combined with the inherent complexity of managing weight parameters and decision boundaries in the encrypted domain, results in the counterintuitive performance outcome where the theoretically more complex FHE system outperforms the simpler PHE approach in our specific access control application.

#### 4.2.1 PHE with Weighted Addition: Limitations

While this weighted approach allows PHE to approximate policy evaluation, it introduces significant limitations:

- **Loss of semantic meaning:** Additive weights cannot fully capture the logical conjunctions required in access policies, leading to imprecise policy enforcement.
- **Threshold ambiguity:** Determining secure and meaningful weights and decision thresholds requires complex calibration and remains vulnerable to edge cases.

- **Computational overhead:** The weighting mechanism introduces additional computational complexity, as evidenced by our performance results showing PHE + Weights being 1.5-2 times slower than FHE + Chameleon Hash.
- **Scalability challenges:** As the number of attributes increases, managing appropriate weights becomes increasingly complex and error-prone.
- **Attack vulnerability:** The additive model with fixed weights is more susceptible to inference attacks through pattern analysis of encrypted values.
- **Policy expressiveness:** Complex nested conditions and hierarchical policies cannot be accurately represented through simple weighted sums.

### 4.3 Security Performance vs Tradeoff Analysis



Figure 4: Security vs Performance Trade-off Analysis positioning our FHE+CH system relative to existing access control approaches.

To provide a quantitative assessment of our FHE+CH system’s security performance, we developed a composite security index based on threat resistance across multiple attack categories. The methodology consists of three systematic steps:

#### 4.3.1 Step 1: Threat Category Identification

We identified five fundamental threat categories for access control systems based on literature review and our established threat model. Each category is assigned a weight reflecting its relative importance in our ZTNA security framework:

1. **Policy Exposure** (Weight: 25%) - Protection against unauthorized access to policy content
2. **Attribute Leakage** (Weight: 25%) - Prevention of user attribute disclosure during processing
3. **Query Manipulation** (Weight: 20%) - Detection and prevention of unauthorized query modifications
4. **Timing Attacks** (Weight: 15%) - Resistance to information inference through processing time analysis
5. **Cryptanalytic Attacks** (Weight: 15%) - Protection against cryptographic analysis attempts

## 5 Discussion and Future Work

### 5.1 Performance Analysis and Implications

The experimental results reveal several critical insights about FHE and Chameleon Hash functions-based access control systems. The non-linear scaling of processing times, ranging from 26 ms for single policy evaluations to upto 700 ms for twenty concurrent evaluations, reflects the inherent computational complexity of homomorphic operations. This performance characteristic is consistent with theoretical expectations for FHE systems, where noise accumulation and ciphertext expansion create super-linear growth patterns.

The comparison with PHE reveals an unexpected outcome that challenges conventional assumptions about homomorphic encryption performance trade-offs. Our FHE + Chameleon Hash approach demonstrates superior performance, achieving approximately 1.5–2 times faster execution than the PHE + Weights implementation across all policy set sizes. This finding indicates that the computational overhead introduced by the weighting mechanism required for policy evaluation in PHE settings can outweigh PHE's theoretical advantages in additive-only operations.

More importantly, this comparison validates our architectural decision to prioritize functional completeness and semantic accuracy over theoretical performance expectations. While PHE requires complex workarounds through weighted addition that compromise policy expressiveness and introduce additional computational overhead, FHE naturally supports the multiplicative operations essential for robust access control policies.

### 5.2 Security and Privacy Implications

Our system achieves end-to-end privacy preservation where the server processes access control decisions without learning user roles, requested resources, or intended operations. This represents a significant advancement in privacy-preserving access control, particularly relevant for cloud computing environments and zero-trust architectures. The computational cost of FHE is justified by these unprecedented privacy guarantees, especially in high-security domains such as healthcare, finance, and government applications.

The superior performance of our FHE approach compared to the weighted PHE alternative further strengthens the case for FHE adoption, as it eliminates the traditional performance penalty argument against FHE while maintaining full semantic accuracy in policy evaluation.

### 5.3 Practical Deployment Considerations

The processing times observed in our evaluation, with sub-second response times for most practical policy set sizes, are well within acceptable bounds for enterprise access control applications. For context, enterprise access control decisions typically occur during user login or resource access initiation, where even multi-second delays are tolerable given the security benefits provided.

The system's current capacity of handling twenty concurrent policy evaluations within 700 ms suggests strong suitability for medium to large-scale deployments. The relatively modest computational requirements make the system deployable in standard enterprise environments without specialized hardware.

### 5.4 Performance Optimization Strategies

Several optimization avenues warrant investigation:

- **Algorithmic Improvements:** Implementing SEAL's batching capabilities could enable parallel evaluation of multiple policies, potentially reducing per-evaluation overhead. Parameter optimization, including fine-tuning polynomial modulus and coefficient sizes, could yield significant performance gains for specific use case requirements.

- **Caching and Precomputation:** Common role-resource combinations could be precomputed and cached, reducing real-time computational overhead. This approach would be particularly effective for organizations with stable role hierarchies and resource structures.
- **Hardware Acceleration:** Exploring specialized FHE hardware accelerators or GPU-based implementations could provide substantial performance improvements. Recent advances in FHE-optimized hardware suggest potential for order-of-magnitude speedups.

## 5.5 Scalability Enhancements

- **Distributed Architecture:** Implementing distributed policy evaluation across multiple nodes could enable horizontal scaling for large enterprise deployments. This approach would require careful consideration of load balancing and result aggregation strategies.
- **Hierarchical Policy Structures:** Developing hierarchical policy evaluation schemes could reduce the number of required homomorphic operations by eliminating obviously invalid combinations early in the evaluation process.
- **Advanced Chameleon Hash Optimizations:** Investigating more efficient Chameleon Hash implementations and batch verification techniques could further improve the performance advantage of our FHE + Chameleon Hash approach.

## 6 Conclusion

By conducting our research, it was found that the integration and successful implementation of FHE and Chameleon Hash is possible to achieve under ZTNA principles. The evaluation of our system provides comprehensive performance analysis and security validation of our FHE and Chameleon Hash-based access control system, demonstrating both its capabilities and practical viability. The security analysis reveals meaningful improvements across multiple threat categories. Our system achieves complete policy confidentiality protection where traditional systems provide none, along with improved query manipulation detection and timing attack resistance. These improvements were validated through both theoretical analysis and experimental testing with our implementation, demonstrating enhanced security compared to traditional access control approaches.

The experimental results reveal that our approach not only enables privacy-preserving access control capabilities but also achieves superior performance compared to PHE-based alternatives, challenging conventional assumptions about homomorphic encryption trade-offs. The comparison with PHE clearly establishes the advantages of FHE for our system architecture on multiple fronts: functional completeness through native support for multiplicative operations, semantic accuracy in policy evaluation, and better computational performance compared to weighted PHE approximations. This finding validates the architectural decisions made in system design.

Our formal security analysis provides theoretical foundations through security proofs demonstrating policy confidentiality and decision integrity properties. The implementation testing achieved high accuracy across all twenty defined access control policies, confirming that homomorphic operations maintain correct functionality while preserving privacy.

Our work contributes to the research in privacy-preserving access control by providing: (1) a practical integration of FHE with chameleon hash for access control, (2) formal security analysis with experimental validation, (3) comprehensive performance evaluation, and (4) concrete implementation demonstrating feasibility for deployment scenarios.

The identified future work directions provide clear pathways for further performance improvements and system capability extensions, positioning this research as a foundation for continued development in privacy-preserving access control systems.

## Acknowledgement

This work is supported by Japan Society for the Promotion of Science (JSPS) KAKENHI Grant Number JP24K20774, 23K21664, 23K21665, 23K28078, and Support Center for Advanced Telecommunications Technology Research (SCAT).

## References

- [1] G. UTU, M. ALKAN, . A. Doru, and M. Drterler, “Perimeter network security solutions: A survey,” in *2019 3rd International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT)*, pp. 1–6, 2019.
- [2] X. Yan and H. Wang, “Survey on zero-trust network security,” in *Artificial Intelligence and Security* (X. Sun, J. Wang, and E. Bertino, eds.), (Singapore), pp. 50–60, Springer Singapore, 2020.
- [3] M. Tsai, S. Lee, and S. W. Shieh, “Strategy for implementing of zero trust architecture,” *IEEE Transactions on Reliability*, vol. 73, no. 1, pp. 93–100, 2024.
- [4] N. F. Syed, S. W. Shah, A. Shaghaghi, A. Anwar, Z. Baig, and R. Doss, “Zero trust architecture (zta): A comprehensive survey,” *IEEE access*, vol. 10, pp. 57143–57179, 2022.
- [5] W. Huang, X. Xie, Z. Wang, J. Feng, G. Han, and W. Zhang, “Zt-access: A combining zero trust access control with attribute-based encryption scheme against compromised devices in power iot environments,” *Ad hoc networks*, vol. 145, p. 103161, 2023.
- [6] C. Gentry, “Fully homomorphic encryption using ideal lattices,” in *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pp. 169–178, 2009.
- [7] C. Gentry, “Computing arbitrary functions of encrypted data,” *Communications of the ACM*, vol. 53, no. 3, pp. 97–105, 2010.
- [8] K. Frikken, M. Atallah, and J. Li, “Attribute-based access control with hidden policies and hidden credentials,” *IEEE Transactions on Computers*, vol. 55, no. 10, pp. 1259–1270, 2006.
- [9] H. Krawczyk and T. Rabin, “Chameleon hashing and signatures.” Cryptology ePrint Archive, Paper 1998/010, 1998.
- [10] H. Krawczyk and T. Rabin, “Chameleon signatures,” in *Proceedings of the Network and Distributed System Security Symposium, NDSS 2000, San Diego, California, USA*, The Internet Society, 2000.
- [11] G. Brassard, D. Chaum, and C. Crpeau, “Minimum disclosure proofs of knowledge,” *Journal of Computer and System Sciences*, vol. 37, no. 2, pp. 156–189, 1988.
- [12] J. Kindervag *et al.*, “Build security into your networks dna: The zero trust network architecture,” *Forrester Research Inc*, vol. 27, pp. 1–16, 2010.
- [13] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, “Zero trust architecture,” 2020-08-10 04:08:00 2020.
- [14] A. Tsuji and M. Oguchi, “Comparison of the schemes and libraries for efficient cryptographic processing,” in *2024 International Conference on Computing, Networking and Communications (ICNC)*, pp. 584–590, 2024.
- [15] D. Xie, P. Haipeng, and L. Li, “Homomorphic signature from chameleon hash functions,” *Information technology and control*, vol. 46, 02 2017.