

HR-CD-Auth: Consent-Bound Cross-Domain Pseudonymous Authentication for Vocational Counseling and Employment Matching

Haewon Byeon¹, Changmin Keum^{2*}

Department of Future Technology, Korea University of Technology and Education (KOREATECH),
Cheonan 31253, South Korea¹

Department of Employment Service Policy, Korea University of Technology and Education (KOREATECH),
Cheonan 31253, South Korea²

Abstract—Vocational counseling platforms move sensitive records across counseling centers, enterprise human-resource systems, public employment services, and automated matching engines. Direct reuse of cross-domain authentication leaves identifiers linkable, separates authentication from purpose and consent, and gives revocation no protocol-level effect on cached credentials. HR-CD-Auth combines domain-scoped pseudonyms, compartmented data keys, consent-bound capabilities, and an ephemeral X25519 handover. Automated analysis of the initial capability-only fast path with ProVerif 2.05 found four failures: protected-record secrecy, accepted-session-key secrecy, and both injective-authentication correspondences. The repaired protocol delivers a fresh one-time handover key in separately protected client and gateway tickets, authenticates the first handover message, and mixes that key with the X25519 secret in HKDF. All four ProVerif queries then returned true. An OpenSSL-backed Python proof of concept executed three independent runs of 10,000 in-memory handovers on an AMD Ryzen 7 8845HS. Run-level medians were 0.2773-0.2812 ms, p95 latency was 0.4669-0.4820 ms, sequential throughput was 3041.5-3119.9 handovers/s, and the three online messages totaled 444 bytes. These figures isolate the cryptographic path and exclude network, TLS, database, policy-engine, and ledger delay. Identity resolution is separated from the online trust service through a two-authority, 2-of-2 threshold profile. The result is a measured and symbolically checked design whose remaining deployment and governance assumptions are explicit.

Keywords—Anonymous authentication; consent management; cross-domain authentication; employment matching; forward secrecy; pseudonymization; vocational counseling

I. INTRODUCTION

Vocational counseling has become a distributed information service. A single case may touch a mobile counseling application, a psychological-assessment provider, an enterprise human-resource gateway, a public employment database, and an automated matching engine. The resulting transcript is unusually sensitive: job-change intention, mental-health indicators, disability or health constraints, salary preferences, employment disputes, and interview availability can all influence a person's prospects. A breach is harmful; silent correlation across otherwise legitimate domains can be just as damaging.

The protocol problem is therefore narrower, and harder, than secure login. A destination gateway must recognize an authorized client without receiving a durable cross-domain identifier. It must also decide whether the present role, purpose, data class, consent version, and retention rule permit the requested operation. Authentication that ends with 'this is a valid user' is incomplete in this setting. The security decision is 'this pseudonymous session may perform this operation, for this purpose, under this consent state, until this expiry time.'

CD-AKA-IoV offers a useful starting point because it separates first-domain authentication, cross-domain notification, and fast subsequent handover [1]. Its vehicle, roadside-unit, and cloud-server assumptions do not transfer unchanged. The published construction assigns a pseudonym during registration and uses the central server to prepare state for a new destination. That is reasonable for its stated vehicular model. In an employment service, the same choices create a linkable handle, a concentrated metadata view, and no cryptographic response to consent withdrawal. These are deployment gaps, not a claim that the source protocol is broken in its original domain.

HR-CD-Auth addresses the gap through four linked mechanisms. First, a keyed, domain-scoped pseudonym is rotated with the session and consent version. Second, the trust service issues a short-lived handover capability that binds destination, purpose, data class, policy version, and expiry. Third, each domain handover performs an ephemeral X25519 exchange; long workflows then use a one-way key ratchet with erasure of prior state. Fourth, access events are committed through Merkle roots while the protected record stays in an encrypted off-chain vault. The ledger is outside the authentication critical path.

The article makes four contributions. First, it defines a threat model that joins a Dolev-Yao network adversary with legitimate-but-overreaching staff and removes unilateral pseudonym resolution through a two-authority threshold profile. Second, it specifies the protocol fields, one-time handover key, key schedule, authorization rule, revocation transition, and audit boundary. Third, an executable ProVerif model identifies an attack in the capability-only draft and verifies four secrecy and injective-authentication queries after the repair. Fourth, a reproducible OpenSSL-backed proof of concept reports

operation timings, three-message payload size, handover latency, and sequential throughput without treating local cryptographic measurements as production HR-system performance.

II. RELATED WORK AND DESIGN GAP

A. Cross-Domain Authentication and Fast Handover

Lightweight vehicular protocols typically combine hashes, nonces, pseudonyms, and symmetric protection to limit work on constrained endpoints. Representative studies address privacy-preserving VANET authentication [2], honey-list authentication [3], three-party fog key agreement [4], blockchain-based multidomain authentication [5], and confidential-computing key transfer [6]. Recent Industry 5.0 work also treats availability as a first-class authentication property rather than a secondary operational metric [7]. Their shared insight is sound: the first interaction with a foreign domain can justify a trust-service exchange, whereas repeated handovers should reuse bounded authorization state.

The differences lie in what that state authorizes and what it reveals. A mobility protocol often needs to prove that a vehicle is registered and may enter a service region. Vocational counseling must carry a much narrower statement. A client authorized for an assessment summary is not automatically authorized to disclose clinical notes, disability details, or salary history. A fast handover token that omits purpose and consent is efficient but semantically overbroad.

B. Private Matching, Access Control, and Audit

Privacy-preserving matching research shows that matching can be performed without exposing every attribute. Matchmaking encryption reveals little beyond whether two policies match [8]. Encrypted task matching has also been combined with revocation [9] and fine-grained worker attributes [10]. These schemes pursue stronger content privacy than HR-CD-Auth, but they do not by themselves provide cross-domain session continuity, gateway authentication, or an operational

consent lifecycle. The framework developed here is complementary: it protects the control plane and limits which pseudonymous features reach a matching service.

Sensitive-data systems provide a closer authorization analogy. Blockchain-assisted e-health protocols have coupled authentication with access control [11], examined data minimization on ledgers [12], and hidden access policies through attribute-based techniques [13]. Patient-oriented authentication work likewise treats identity privacy as distinct from channel secrecy [14]. The transferable lesson is not that employment data should be placed on a healthcare blockchain. It is that authentication, least privilege, revocation, and audit need a single protocol context, while raw records remain under the data controller's storage policy.

A protected counselor terminal can reduce exposure of long-term keys, but a trusted execution environment is an assumption rather than a proof shortcut. SGX-based authentication schemes illustrate how hardware isolation can protect credential operations [15], and the broader SGX literature documents the attendant side-channel, rollback, and lifecycle concerns [16]. HR-CD-Auth therefore treats the terminal as a bounded trust component and keeps compromise of ordinary gateway storage inside the adversary model.

C. Design Gap

The literature leaves a specific gap. Cross-domain AKA designs optimize mobility but seldom bind the handover credential to consent version, role, purpose, and data class. Private matching protects attributes but does not establish a pseudonymous session across administrative domains. Ledger-based access control can improve accountability, yet many designs put ledger confirmation or excessive metadata on the online path. HR-CD-Auth combines these concerns without claiming that any one component is novel in isolation. Its contribution is the protocol binding among them (Table I).

TABLE I. FROM CD-AKA-IOV TO THE HR-CD-AUTH SECURITY CONTRACT

Source property	Employment-domain gap	HR-CD-Auth control
Registered pseudonym supports anonymous vehicular access	A pseudonym reused across counseling, assessment, HR, and public services becomes a profile key	Keyed pseudonym scoped to domain, purpose, consent version, counter, and fresh nonce
Central server prepares first contact with a destination	The trust service can observe a concentrated handover graph	Short-lived destination capability; no online T call; identity resolution requires two independent authorities
Authentication establishes a session key	Authentication alone does not limit role, purpose, data class, or retention	Session-bound access token and compartment-specific wrapped data keys
Fast path relies on hashes and symmetric operations	Hash-only key derivation cannot substantiate forward secrecy after long-term key compromise	Ephemeral X25519 at each handover, followed by an erased one-way symmetric ratchet
Security analysis targets network attacks	A valid staff credential can still be used outside assigned purpose	Insider model, explicit authorization predicate, and tamper-evident audit commitment

III. SYSTEM, TRUST, AND THREAT MODEL

A. Actors and Trust Boundaries

The client C is the data subject and holder of a protected registration credential. A source gateway G_s operates a counseling service; a destination gateway G_d represents an enterprise HR system, public employment service, assessment provider, or another approved domain. The National Labor Data Trust T registers clients and gateways, distributes policy material, and issues bounded capabilities, but it does not retain a complete identity-resolution key. Resolution authorities R_1 and R_2 hold independent shares of a 2-of-2 key. Secure counselor terminals protect staff-side credentials; encrypted vaults hold records; the permissioned audit service accepts commitments only (Fig. 1).

T is trusted to perform registration and policy issuance correctly, but it cannot read raw counseling content or resolve a stored pseudonym mapping by itself. A single compromised resolution authority reveals no usable decryption key; R_1 and R_2 must approve the same case-bound request. Collusion between both authorities remains a privacy failure. Gateways are protocol-correct but may be curious, misconfigured, or later compromised. Staff members can hold valid credentials and still exceed purpose or scope. The audit service may delay publication, but it does not decide whether client and gateway authentication completes (Table II).

TABLE II. ACTOR RESPONSIBILITIES AND PROTECTED STATE

Actor	Trust assumption	Protected state or output
Client C	Credential store may be protected; host OS can be exposed	Registration key, counter, current ratchet state
Gateway G_s / G_d	Protocol-correct but curious; ordinary storage can be compromised	Gateway trust key, ticket cache, authorization decision
Trust service T	Correct registration and policy issuance; no unilateral identity resolution	Threshold-encrypted mapping, gateway registry, policy keys
Resolution authorities R_1 / R_2	At least one authority remains honest; approvals are case-bound	Independent key shares and signed resolution approvals
Secure counselor terminal	Hardware boundary assumed intact within the model	Staff key, role certificate, local decision evidence
Encrypted vault	Ciphertext and metadata may be observed	Compartmented counseling, assessment, and matching records
Audit service	Append-only consensus; may be delayed or unavailable	Policy hashes, consent-token hashes, Merkle roots, signatures




Fig. 1. HR-CD-Auth actors and trust boundaries. T issues policy and tickets but cannot resolve identity alone; raw records remain in encrypted vaults.

B. Adversary Capabilities

The network adversary follows the Dolev-Yao model: it can intercept, replay, reorder, delay, modify, and inject messages on public channels [17]. Session compromise follows the spirit of the Canetti-Krawczyk model [18]. The adversary may reveal an established session key, compromise ordinary client or gateway storage, obtain an expired token, or corrupt one gateway after a completed session. It cannot break X25519, HKDF, HMAC, authenticated encryption, or the collision resistance of the selected hash. It also cannot recover an erased ephemeral scalar.

The insider adversary is an authenticated counselor, HR officer, gateway administrator, public-service employee, or operator of T or one resolution authority. The insider may request a field outside the active purpose, reuse a cached token after withdrawal, correlate timing metadata, or alter a local audit file. The model allows compromise of either R_1 or R_2 , but not both in the same resolution event. Ordinary login therefore does not end the security analysis; authorization, key release, and resolution approval remain separate decisions.

Two exclusions are explicit. The framework does not prevent a recipient from retaining plaintext that it was legitimately shown before consent was withdrawn. It also does not hide network-layer identifiers, timing, message size, or external information that may correlate sessions. Those risks require transport anonymity, traffic shaping, endpoint controls, and organizational sanctions beyond the protocol.

C. Security and Privacy Properties

The required properties are mutual entity authentication; agreement on destination, purpose, data class, consent version, and policy version; session-key secrecy; forward secrecy for completed handovers; replay resistance; cross-domain pseudonym unlinkability against ordinary gateways; rejection of withdrawn or expired capabilities; least-privilege data-key release; detection of audit-log alteration; and resistance to

unilateral identity resolution after compromise of T or one resolution authority. Availability is narrower: a pre-authorized handover must not wait for T, the resolution authorities, or ledger finality.

IV. HR-CD-AUTH PROTOCOL

A. Context Binding and Notation

Each protocol decision is tied to one context rather than a free-standing identity. Let d denote the destination domain, $purpose_s$ the declared use, $dataClass$ the requested compartment, $consentVersion$ the client-approved policy instance, and $epoch_s$ a bounded validity period. C and G generate fresh nonces N_C and N_G and ephemeral X25519 key pairs (x_C, X_C) and (x_G, X_G) . T shares K_{CT} with the protected client credential and K_{TG} with each registered gateway. For a pre-authorized handover, T also creates a fresh one-time key K_{HO} and places it only in separately protected client and gateway tickets. HMAC, HKDF, and AEAD denote a message authentication code, a key-derivation function, and authenticated encryption with associated data.

The protocol uses accountable pseudonymity rather than absolute anonymity. Ordinary gateways cannot compute the mapping from their transcripts. T stores only a threshold-encrypted mapping and cannot resolve it alone; R_1 and R_2 must authorize and complete the 2-of-2 resolution procedure.

$$PID_s = \text{HMAC}_{KC}(ID_C \parallel d \parallel purpose_s \parallel consentVersion \parallel ctr_C \parallel N_C) \quad (1)$$

where, ctr_C is a monotonic client counter protected with the registration credential. Including the fresh nonce avoids deterministic output if a device restores an earlier counter. The counter advances after a successful handover, a consent transition, or an abnormal termination. The context digest is

$$context_s = H(PID_s \parallel ID_G \parallel d \parallel purpose_s \parallel dataClass \parallel consentVersion \parallel policyVersion \parallel epoch_s) \quad (2)$$

Both endpoints place $context_s$ in the transcript and reject a mismatch before data-key release. This single digest is the binding point among authentication, authorization, and audit.

B. Registration and Policy Provisioning

C enrolls with T over an authenticated protected channel. The direct identity mapping is encrypted under the joint resolution public key, and the complete resolution secret is not retained by T. T stores the threshold ciphertext, credential status, pseudonym counter state, and current consent-policy reference; it stores no counseling record. Each gateway registers ID_G , domain type, supported data classes, policy version, and K_{TG} . Counselor terminals receive staff credentials carrying role and organizational scope.

Policy provisioning produces two independent authorization objects. A handover capability authorizes one destination to

authenticate a pseudonymous client under a bounded context; a session access token authorizes a data operation after authentication. T also generates a fresh K_{HO} for each capability. The client receives K_{HO} inside its K_{CT} -protected ticket, while G_d obtains the same value only by decrypting HOT_s under K_{TG} . K_{HO} never appears on the public channel. Compromise of a capability alone therefore neither authenticates a client nor unwraps a record key.

C. First-Domain Authentication

On first contact or after a policy refresh, C sends PID_s , N_C , X_C , $purpose_s$, $consentVersion$, and an authenticator derived from K_{CT} . G adds N_G , X_G , its identity, requested data class, and a gateway authenticator, then forwards the protected request to T. T verifies registration status, gateway status, consent version, purpose, and policy version. A successful response contains a gateway ticket and a client ticket over the same context. Neither ticket contains raw counseling data (Table III).

C and G compute the ephemeral shared secret and derive the session key from the full transcript:

$$Z_s = X25519(x_C, X_G) = X25519(x_G, X_C) \quad (3)$$

$$SK_s = \text{HKDF}(Z_s \parallel K_{HO}, H(\text{transcript}_s), context_s) \quad (4)$$

The endpoints exchange role-separated confirmation tags under SK_s and erase x_C , x_G , and Z_s after confirmation. A later compromise of K_{CT} , K_{TG} , or a completed session state does not reveal the erased ephemeral secret. Forward secrecy therefore rests on the hardness of X25519, correct randomness, HKDF, and actual erasure; it is not inferred from nonce freshness alone.

D. Pre-Authorized Cross-Domain Handover

When a counseling workflow is allowed to move to G_d , T creates a short-lived capability and a fresh one-time K_{HO} . The capability may be delivered during first contact or cached through an authenticated gateway channel. Its protected payload is

$$HOT_s = \text{AEAD}_{KTG}(PID_s \parallel ID_{Gd} \parallel purpose_s \parallel dataClass \parallel consentVersion \parallel policyVersion \parallel exp_s \parallel K_{HO}) \quad (5)$$

G_s transfers HOT_s and the encrypted state needed for the expected client. C rotates its pseudonym and sends N_C , X_C , $context_s$, HOT_s , and $Auth_C^{HO} = \text{HMAC}_{KHO}("C\text{-auth"} \parallel PID_s \parallel N_C \parallel X_C \parallel context_s \parallel H(HOT_s))$. G_d decrypts HOT_s , checks expiry, policy, and consent versions, and verifies $Auth_C^{HO}$ before performing X25519. Both endpoints derive SK_s from Z_s , K_{HO} , the transcript hash, and $context_s$. The gateway response and final client message carry role-separated confirmation tags under SK_s . This remains a three-message client-destination exchange; T, the resolution authorities, and the ledger are absent from the latency-critical path (Fig. 2).

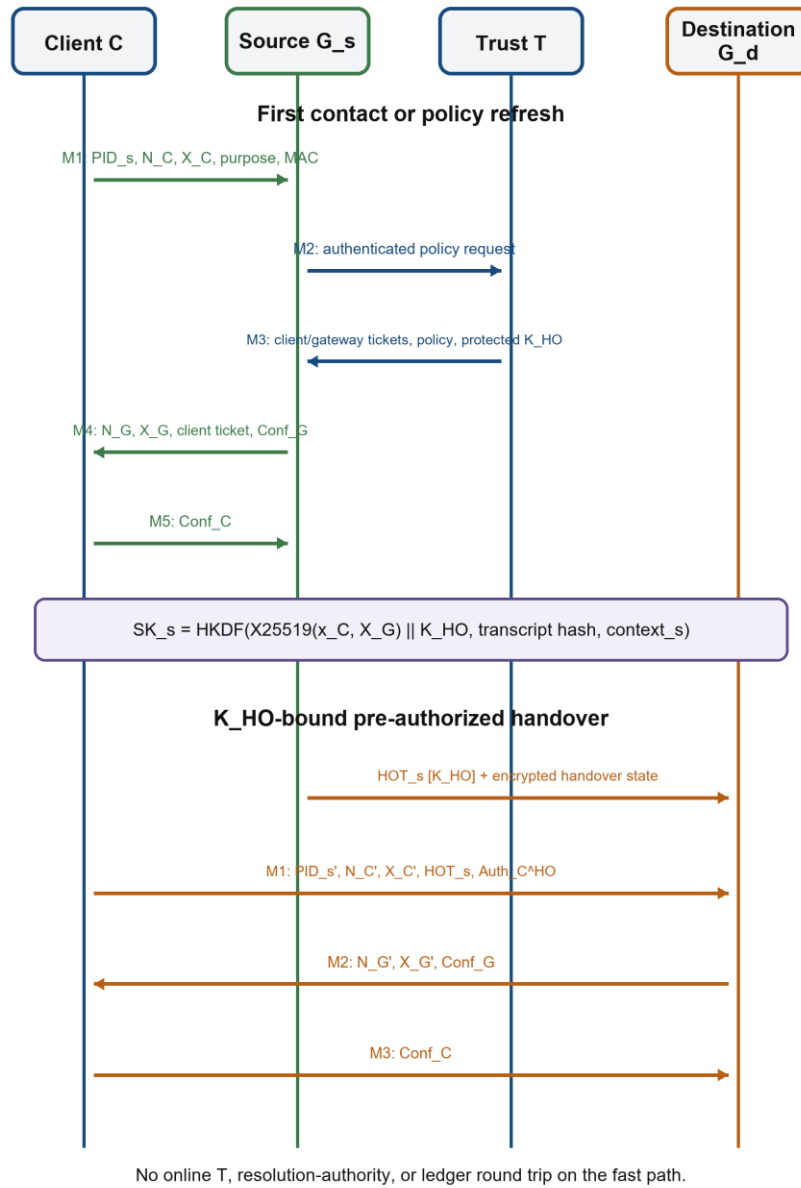


Fig. 2. First-contact and K_{HO}-bound pre-authorized handover sequence. The three-message fast path authenticates both ephemeral shares without an online T, resolution-authority, or ledger round trip.

TABLE III. ONLINE MESSAGES AND THEIR SECURITY PURPOSE

Phase	Flow	Essential fields	Security purpose
First contact	C → G _s	PID _s , N _C , X _C , purpose, consent version, authenticator	Client freshness and context proposal
First contact	G _s → T	Client block, N _G , X _G , gateway ID, data class	Registration and policy decision
First contact	T → G _s → C	Gateway ticket, client ticket, policy version, expiry	Bind both endpoints to one authorized context
Handover setup	G _s → G _d	HOT _s with K _{HO} ; K _{HO} separately protected for C and G _d	Pre-position destination authorization and one-time authentication key
Online handover	C ↔ G _d	Rotated PID, nonces, key shares, Auth _C ^{HO} , confirmation tags	K _{HO} -bound forward-secret mutual authentication

E. Consent-Bound Access and Data-Key Release

Authentication does not release a record. The policy service issues a session access token only after evaluating the staff credential, client consent, purpose, requested data class, and retention rule. The token is bound to the hash of SK_s , which prevents replay in a different authenticated session:


$$AT_s = \text{HMAC}_{K_{\text{policy}}}(\text{PID}_s \parallel \text{role} \parallel \text{purpose}_s \parallel \text{dataClass} \parallel \text{consentVersion} \parallel \text{exp}_s \parallel H(SK_s)) \quad (6)$$

A gateway unwraps a compartment key only when the complete predicate holds:

$$\begin{aligned} \text{allow} = & \text{valid}(AT_s) \wedge \text{role} \geq \text{requiredRole} \wedge \text{purpose}_s \in \\ & \text{allowedPurpose} \wedge \text{consentVersion} = \text{activeVersion} \wedge \\ & \text{now} < \text{exp}_s \end{aligned} \quad (7)$$

Counseling notes, assessment summaries, matching features, public-service transfer fields, and audit evidence use separate data-encryption keys. An HR gateway may receive a job-tag match and interview token without receiving raw psychological notes. A counselor may receive an assessment summary without learning enterprise feedback. The policy decision therefore changes the available key material, not only the user interface.

Withdrawal or expiry invalidates HOT_s and AT_s , advances ctr_C , and triggers rewrapping or destruction of retained compartment keys according to policy. The protocol prevents new decryptions under managed keys. It cannot recall plaintext already exported to a recipient before withdrawal; that residual obligation belongs to endpoint governance and enforcement (Fig. 3).



The ledger stores commitments, not the consent document or protected record.

Fig. 3. Consent-state transition and its cryptographic effects. The ledger stores a commitment to the transition, not the consent document or protected record.

F. Forward-Secure Ratchet and Audit Commitments

Within a long counseling workflow, an endpoint may ratchet the confirmed session key at an epoch or data-class boundary and erase the previous value:

$$SK_{\text{next}} = \text{HKDF}(SK_{\text{current}}, \text{epoch}_{\text{next}} \parallel \text{consentVersion}) \quad (8)$$

This one-way update protects earlier epoch keys from later state exposure when erasure is effective. It does not replace the ephemeral X25519 exchange at a domain handover; the two mechanisms address different compromise windows.

Each authorized operation creates an audit leaf that omits the direct identity and record content:

$$L_e = H(\text{PID}_s \parallel \text{role} \parallel \text{action} \parallel \text{dataClass} \parallel \text{consentVersion} \parallel t_e \parallel \text{recordID} \parallel \text{gatewayID}) \quad (9)$$

$$R_m = \text{MerkleRoot}(L_1, L_n) \quad (10)$$

Merkle trees provide compact evidence that an event belongs to a committed batch [19]. The permissioned ledger receives R_m , $H(\text{policyVersion})$, $H(\text{consentToken})$, the gateway signature, and the publication epoch. Authentication may finish before ledger finality; a bounded publication deadline and a locally signed queue cover temporary ledger unavailability. Raw psychological, health, disability, salary, employment, and interview data never enter the ledger payload.

V. SECURITY ANALYSIS

A. Property-by-Property Argument

Replay attempts fail at two layers. The authentication transcript contains fresh nonces, ephemeral public keys, destination identity, context_s, and an expiry-bounded ticket. Reusing an earlier client message either hits the destination replay cache or produces a transcript hash that does not match the new gateway contribution. Reusing AT_s in another session fails because the token contains $H(SK_s)$. Reusing HOT_s after a consent change fails because the active consent version and revocation state no longer match.

Client impersonation in the pre-authorized path requires K_{HO} to forge $Auth_C^{HO}$; gateway impersonation requires K_{HO} as an input to the session KDF. Replacing X_C or X_G changes the transcript, and a chosen public share does not yield SK_s without K_{HO} . A compromised gateway can still impersonate that gateway until its trust key and outstanding capabilities are revoked. Short expiry and gateway-specific keys contain, but do not erase, that exposure.

Completed sessions have forward secrecy against later compromise of registration and gateway trust keys because SK_s depends on an erased X25519 shared secret. This statement excludes compromise of x_C or x_G before erasure, weak randomness, endpoint snapshots that preserve the scalar, and side-channel recovery. The within-session ratchet protects past ratchet epochs from later current-state disclosure, again only when old keys are erased.

Ordinary gateways cannot compute whether two PID_s values correspond to the same client without K_C or the threshold-encrypted mapping. Domain, purpose, consent

version, counter, and nonce change the HMAC input. The claim is protocol-level unlinkability, not traffic-analysis anonymity. A shared IP address, timing pattern, rare job attribute, or external database can still correlate sessions. Re-identification requires both resolution authorities; collusion of R_1 and R_2 defeats that protection.

Insider resistance follows from key release, not from staff authentication alone. A valid counselor credential that does not satisfy the predicate in Eq. (7) receives no wrapped compartment key. Altering a local decision record changes the corresponding leaf and breaks verification against R_m . This mechanism makes deletion or modification detectable; it does not prove that a malicious gateway faithfully logged an event it chose never to record. Independent gateway signing, monitored publication deadlines, and cross-checks with vault access logs address that residual risk.

B. BAN-Logic Belief Analysis

BAN logic is used here for a limited purpose: checking whether the protected tickets and confirmation tags support mutual belief in an authorized context and possession of the derived key [20]. It does not prove X25519 hardness, HMAC unforgeability, traffic-analysis resistance, or consent-law compliance. The notation $P \models X$ means that P believes X , $P \sim X$ means that P once said X , $\#(X)$ denotes freshness, $P \Rightarrow X$ denotes jurisdiction, and $P \leftrightarrow_K Q$ denotes a key shared by P and Q .

The authentication-critical messages are idealized as follows:

$$M_2: G \rightarrow T: \text{Enc}_{\text{KCT}}(\text{PID}_s, N_C, X_C, \text{context}_s), \text{Enc}_{\text{KGT}}(N_G, X_G, G) \quad (11)$$

$$M_3: T \rightarrow G: \text{Enc}_{\text{KGT}}(N_C, N_G, X_C, X_G, \text{context}_s, \text{exp}_s), \text{Ticket}_C \quad (12)$$

$$\text{Ticket}_C = \text{Enc}_{\text{KCT}}(N_C, N_G, X_C, X_G, G, \text{context}_s, \text{exp}_s) \quad (13)$$

$$M_4: G \rightarrow C: \text{Ticket}_C, \text{HMAC}_{\text{SKS}}(\Gamma\text{-}\chi\text{ov}\phi\text{r}\mu \parallel H(\text{transcript}_s)) \quad (14)$$

$$M_5: C \rightarrow G: \text{HMAC}_{\text{SKS}}(X\text{-}\chi\text{ov}\phi\text{r}\mu \parallel H(\text{transcript}_s)) \quad (15)$$

The initial beliefs are that C shares K_{CT} with T , G shares K_{GT} with T , each endpoint trusts T 's jurisdiction over the identity-to-context binding, and N_C and N_G are fresh:

$$C \models C \leftrightarrow_{\text{KCT}} T; G \models G \leftrightarrow_{\text{KGT}} T \quad (16)$$

$$C \models \#(N_C); G \models \#(N_G) \quad (17)$$

$$C \models T \Rightarrow \text{Bind}(C, G, \text{context}_s); G \models T \Rightarrow \text{Bind}(C, G, \text{context}_s) \quad (18)$$

Applying the message-meaning and nonce-verification rules to M_3 and Ticket_C gives each endpoint the belief that T recently asserted the same key shares and context. Jurisdiction then yields

$$C \models \text{Bind}(C, G, \text{context}_s); G \models \text{Bind}(C, G, \text{context}_s) \quad (19)$$

Both endpoints compute SK_s from the authenticated key shares and transcript. The role-separated tags in M_4 and M_5 establish the remaining possession beliefs:

$$C \models G \models \text{Possess}(\text{SK}_s); G \models C \models \text{Possess}(\text{SK}_s) \quad (20)$$

The analysis supports mutual agreement on the authorization context and key possession under the stated belief assumptions. Computational secrecy and forward secrecy still depend on the cryptographic and erasure assumptions discussed above.

C. Automated Symbolic Verification and Design Repair

An applied-pi-calculus model was executed with ProVerif 2.05 [21] over unbounded sessions and a public Dolev-Yao channel. Protected tickets, ideal HMAC and hashing, the Diffie-Hellman commutation equation, transcript-bound HKDF, and injective correspondence events are represented explicitly. Four queries test protected-record secrecy, secrecy of an accepted session key, client acceptance only after a matching gateway run, and gateway acceptance only after a matching client run.

The capability-only draft failed all four queries. ProVerif reconstructed an active trace in which the adversary replays HOT_s , substitutes a chosen ephemeral share, derives the corresponding session key, and completes key confirmation without a matching peer run. This was a protocol defect, not a limitation of BAN logic.

The repaired model delivers a fresh K_{HO} to C inside the K_{CT} -protected client ticket and to G_d inside HOT_s under K_{TG} . $\text{Auth}_{C^A}\text{HO}$ authenticates the first online message, and HKDF takes both Z_s and K_{HO} as input. After this change, ProVerif returned true for all four queries. Fig. 4 preserves the actual result lines for the rejected and repaired models.

The result is scoped to the symbolic abstraction (Table IV). It does not establish constant-time implementation, X25519 input validation, secure erasure, wall-clock expiry, replay-cache consistency, or the threshold-resolution workflow. Those properties require code-level and deployment tests.

```

ProVerif 2.05 - HR-CD-Auth automated verification
$ proverif hr_cd_auth_capability_only.pv
RESULT not attacker(authorizedRecord[]) is false.
RESULT not (event(gateway_accept(pid_3,nc_2,nc_2,ng_2,ctx_3,sk_2)) && attacker(sk_2))
is false.
RESULT inj-event(client_accept(pid_3,nc_2,nc_2,ng_2,ctx_3,sk_2)) ==>
inj-event(gateway_begin(pid_3,nc_2,nc_2,ng_2,ctx_3,sk_2)) is false.
RESULT inj-event(gateway_accept(pid_3,nc_2,nc_2,ng_2,ctx_3,sk_2)) ==>
inj-event(client_begin(pid_3,nc_2,nc_2,ng_2,ctx_3)) is false.
model sha256: 0cee7862c57a09d26070db...

$ proverif hr_cd_auth_repaired.pv
RESULT not attacker(authorizedRecord[]) is true.
RESULT not (event(gateway_accept(pid_3,nc_2,nc_2,ng_2,ctx_3,sk_2)) && attacker(sk_2))
is true.
RESULT inj-event(client_accept(pid_3,nc_2,nc_2,ng_2,ctx_3,sk_2)) ==>
inj-event(gateway_begin(pid_3,nc_2,nc_2,ng_2,ctx_3,sk_2)) is true.
RESULT inj-event(gateway_accept(pid_3,nc_2,nc_2,ng_2,ctx_3,sk_2)) ==>
inj-event(client_begin(pid_3,nc_2,nc_2,ng_2,ctx_3)) is true.
model sha256: 432d52db5017e3fc288273a4...

```

Fig. 4. Actual ProVerif 2.05 summaries. The capability-only draft fails 4/4 queries; the K_{HO} -bound repair passes 4/4.

TABLE IV. SECURITY CLAIMS, MECHANISMS, AND RESIDUAL ASSUMPTIONS

Claim	Primary mechanism	Required assumption	Residual exposure
Replay resistance	Fresh nonces, transcript hash, expiry, replay cache	Fresh randomness and bounded clock/epoch checks	Large distributed replay caches require synchronization
Mutual authentication	K _{HO} client authenticator, K _{HO} -bound HKDF, and role-separated confirmation	Unforgeable HMAC/AEAD and authenticated gateway registry	A live stolen gateway key impersonates that gateway until revocation
Forward secrecy	Fresh ephemeral X25519 plus scalar erasure	Hardness of X25519 and effective erasure	Pre-erasure endpoint compromise reveals the session
Gateway unlinkability	Keyed domain/purpose/session pseudonym	K _C and the two threshold key shares remain protected	Traffic correlation or collusion of both resolution authorities
No unilateral resolution	Threshold-encrypted mapping and two signed approvals	At least one of R ₁ and R ₂ remains honest	Collusion, enrollment capture, or reconstruction-boundary compromise
Consent enforcement	Consent-version HOT _s and AT _s ; key rewrapping	Gateways consult current revocation state	Previously exported plaintext cannot be recalled
Audit integrity	Signed leaves and Merkle-root commitment	At least one independent verifier retains a commitment	A malicious gateway may omit an event before logging

VI. ANALYTICAL AND EMPIRICAL EVALUATION

A. Online Cost Model

The evaluation begins with an analytical cost model and then measures the repaired fast path. Let T_X denote one X25519 scalar multiplication, T_H one hash, T_M one HMAC, T_K one HKDF invocation, and T_A one AEAD operation. Database lookup, replay-cache access, and policy evaluation remain separate because their cost depends on deployment scale. Table V reports the operation counts; Table VI reports local measurements from the executable proof of concept.

At first contact, C performs one pseudonym HMAC, one X25519 operation, one HKDF, and confirmation or transcript MACs. A pre-authorized handover removes T from the online path but retains one fresh X25519 operation at each endpoint, verifies HOT_s and Auth_C^{HO}, and mixes K_{HO} into HKDF. The public-key work is intentional: a hash-only handover cannot

support the forward-secrecy claim, while a capability-only exchange does not authenticate the chosen key shares.

$$B_{\text{handover}} = B_{\text{PID}} + 2B_N + 2B_X + B_{\text{HOT}} + 3B_{\text{MAC}} + B_{\text{context}} \quad (21)$$

Eq. (21) gives the analytical payload form. In the proof-of-concept encoding, the three online messages contain 332, 80, and 32 bytes, respectively, for a total of 444 bytes. The pre-positioned gateway transfer and asynchronous ledger publication are not additional client-destination rounds.

CD-AKA-IoV reports a two-round, 1024-bit fast cross-domain phase for its narrower vehicular contract [1]. HR-CD-Auth carries consent and purpose state, authenticates both ephemeral shares with K_{HO}, and provides forward secrecy. The local proof-of-concept therefore demonstrates feasibility, not a claim that HR-CD-Auth is faster than CD-AKA-IoV or a deployed HR service.

TABLE V. ANALYTICAL ONLINE COST PROFILE

Path	Online messages	Client operations	Gateway operations	Online external dependency
First contact/policy refresh	5	$T_X + T_K + \text{HMAC/H}$ for pseudonym and confirmations	$T_X + T_K + \text{ticket and MAC processing}$	Trust service T
Pre-authorized handover	3 client-destination; 1 prior gateway transfer	$T_X + T_K + K_{\text{HO-authenticated HMAC/H}}$	$T_X + T_K + \text{HOT/Auth}_C^{\text{HO}} \text{ verification} + \text{HMAC/H}$	None
Within-session ratchet	0 or application piggyback	$T_K + \text{erasure}$	$T_K + \text{erasure}$	None
Audit publication	Asynchronous batch	None	Leaf hash + signature + queue	Permissioned ledger, off critical path

B. Proof-of-Concept Benchmark

The executable prototype uses Python 3.12.13, cryptography 49.0.0, and OpenSSL 4.0.1 on Windows 11 with an AMD Ryzen 7 8845HS. Each of the three independent runs performed 500 warm-up handovers followed by 10,000 measured handovers. `time.perf_counter_ns` recorded the complete in-memory three-message path: binary construction and parsing, HOT's AES-GCM verification, `Auth_C^HO`, two ephemeral X25519 key generations and exchanges, two HKDF-SHA256 derivations, and role-separated HMAC confirmations.

Run-level median latency ranged from 0.2773 to 0.2812 ms; p95 ranged from 0.4669 to 0.4820 ms. Sequential throughput ranged from 3041.5 to 3119.9 completed handovers/s. The fixed binary encoding produced 332-byte `M_1`, 80-byte `M_2`, and 32-byte `M_3` messages, or 444 online bytes. Table VI reports the operation-level measurements, and Fig. 5 preserves the console summary.

These values measure cryptographic feasibility, not enterprise HR-service latency or capacity. Network transport, TLS, database and policy-engine calls, distributed replay-cache synchronization, concurrency, disk I/O, trust-service issuance, and ledger publication were deliberately excluded.

TABLE VI. MEASURED PROOF-OF-CONCEPT PERFORMANCE

Measured item	Median range (ms)	p95 / rate
X25519 key generation	0.0515-0.0516	0.0773-0.0790 ms
X25519 shared-secret exchange	0.0489	0.0709-0.0720 ms
HKDF-SHA256 (32-byte output)	0.0064	0.0115-0.0122 ms
HMAC-SHA256 (256-byte input)	0.0033	0.0057-0.0063 ms
AES-GCM HOT's decryption	0.0015-0.0016	0.0027-0.0031 ms
Complete three-message handover	0.2773-0.2812	0.4669-0.4820 ms
Sequential throughput	-	3041.5-3119.9 /s
Online payload	-	444 B (332/80/32)

```
HR-CD-Auth PoC - measured OpenSSL-backed benchmark
$ python benchmark_hr_cd_auth.py [three independent runs]
environment: AMD Ryzen 7 8845HS; Windows 11; Python 3.12.13; cryptography 49.0.0; OpenSSL 4.0.1
4.0.1 9 Jun 2026
design: 3 independent runs x 10,000 handovers; 500 warm-up iterations/run
end-to-end median range: 0.2773-0.2812 ms; p95 range: 0.4669-0.4820 ms
sequential throughput range: 3041.5-3119.9 handovers/s
online bytes: M1=332, M2=80, M3=32, total=444
scope: single-process, in-memory cryptographic path; network/DB/TLS excluded
```

Fig. 5. Actual OpenSSL-backed PoC console summary across three independent 10,000-handover runs.

C. Implementation Boundaries

The client counter and ratchet state need rollback protection. A secure element, platform monotonic counter, or server-checked counter window can provide it; simply storing `ctr_C` in an ordinary application database is insufficient. Gateways need a bounded replay cache keyed by ticket identifier and epoch. Policy and revocation updates should be signed, versioned, and distributable during partial trust-service outages. The protocol must also define canonical encoding so that every participant hashes the same byte sequence.

A deployment may substitute another authenticated ephemeral key exchange or a post-quantum KEM if the transcript and confirmation rules remain intact. The substitution changes operation and payload counts and must be re-evaluated. Similarly, the audit service can be a consortium ledger or another independently verifiable append-only log. The security requirement is durable, cross-organizational commitment; it is not allegiance to a particular blockchain platform.

VII. GOVERNANCE, PRIVACY, AND OPERATIONAL FIT

A. Regulatory and Standards Alignment

HR-CD-Auth is an engineering control, not a finding of legal compliance. Under the General Data Protection Regulation, health-related fields, profiling, purpose limitation, data minimization, and meaningful information about consequential automated decisions require particular attention [22]. The Korean Personal Information Protection Act likewise makes lawful processing, safeguards, access control, and access-record management deployment questions rather than properties conferred by a cryptographic protocol [23]. The design supports those duties by reducing stable identifiers, binding purpose and consent to key release, and separating raw records from durable audit commitments.

ISO/IEC 27001:2022 provides the organizational information-security context for key management, access control, supplier interfaces, logging, and incident response [24]. ISO/IEC 29100 supplies a privacy vocabulary for purpose specification, use limitation, data minimization, openness, and accountability [25]. These standards help define operating controls around HR-CD-Auth; they do not certify the protocol or the deploying organization.

When a matching engine uses machine learning, the NIST AI Risk Management Framework is relevant to model governance, explanation, monitoring, and contestability [26]. Authentication should pass only the pseudonymous features needed for the approved matching purpose, and the audit log should identify the model and policy version. Zero-trust guidance reinforces the decision to evaluate every data request against current identity, device, policy, and resource context rather than rely on network location [27]. None of these controls detects biased training data or an unfair ranking rule by itself.

B. Integration and Emergency Access

An implementation can expose registration, ticket, handover, authorization, revocation, audit, and dual-resolution approval endpoints through an OpenAPI-described interface [28]. The specification must distinguish control-plane objects from protected payloads, define canonical encodings, and forbid

raw sensitive fields in logs. Gateway adapters should map local roles and data labels to the shared policy vocabulary rather than broaden them silently.

Emergency access needs its own purpose code, short expiry, senior authorization, user notification when lawful, and a separately reviewable audit leaf. It must not become an unlogged administrator bypass. Public employment services that must resolve civil identity submit a case-bound request containing purpose, legal basis, requesting officer, expiry, and the hash of the target pseudonym. R_1 and R_2 each verify the request and produce an independently signed approval. Only matching approvals permit 2-of-2 threshold reconstruction inside a protected execution boundary; the reconstructed key is not exported. This use of split key custody follows the threshold principle of Shamir secret sharing [29]. A 2-of-2 rule prevents unilateral resolution but reduces availability; a 2-of-3 deployment is a possible resilience variant.

VIII. LIMITATIONS AND VALIDATION AGENDA

The proof of concept validates the repaired cryptographic handover, not a commercial HRIS, applicant-tracking system, counseling application, or public employment service. Its latency and throughput exclude network, TLS, database, policy-engine, replay-cache synchronization, disk I/O, concurrent clients, and ledger publication. The results are therefore local cryptographic-path measurements rather than production TPS or user-perceived latency.

BAN logic checks beliefs about tickets and key confirmation. ProVerif adds an automated symbolic analysis of four secrecy and injective-authentication queries, all of which pass for the repaired model. Neither result is a computational proof of the implementation. The executable model idealizes HMAC, hashing, ticket encryption, and Diffie-Hellman, and it does not model wall-clock expiry, distributed replay-cache state, counter rollback, threshold resolution, side channels, or source-code memory erasure.

The privacy boundary still depends on governance. T or either resolution authority alone can decrypt the identity mapping, but collusion between R_1 and R_2 , capture during enrollment, or compromise of the protected reconstruction boundary can re-identify a client. The 2-of-2 profile also sacrifices availability when either authority is offline. A recipient may retain plaintext received before withdrawal, network metadata can correlate sessions, and a matching model can still discriminate.

Future work should deploy the client, gateway, and trust-service roles over TLS with real policy and database calls; repeat the benchmark on ARM and Raspberry Pi hardware under concurrent load; test expiry, replay-cache synchronization, rollback, and key erasure by fault injection; and add a game-based proof of the final wire encoding. The two-authority resolution profile also needs an HSM-backed implementation, collusion exercises, availability measurement, and comparison with a 2-of-3 variant.

IX. CONCLUSION

HR-CD-Auth treats a valid identity as the start, not the end, of a security decision. The repaired protocol rotates

pseudonyms, binds the first handover message with a one-time K_{HO} , combines K_{HO} with an ephemeral X25519 secret, and releases compartment keys only after a role-and-purpose authorization check. Identity resolution requires two independent authorities, while the audit service records commitments rather than raw records.

The automated analysis first rejected the capability-only draft and then proved all four secrecy and injective-authentication queries for the K_{HO} -bound repair. Across three 10,000-handover runs, the local cryptographic path recorded 0.2773-0.2812 ms median latency, 0.4669-0.4820 ms p95 latency, 3041.5-3119.9 sequential handovers/s, and 444 online bytes. These are measured feasibility results, not production-service claims. Their value lies in exposing the design repair, empirical cost, trust boundary, and remaining deployment risks in terms a reviewer can reproduce and challenge.

DECLARATION ON GENERATIVE AI

Generative artificial-intelligence tools were used to assist with language editing under the author's supervision. The author is responsible for the technical claims, equations, references, and final manuscript.

ACKNOWLEDGMENT

This study was supported by the Education and Research Promotion Program of KOREATECH in 2026.

REFERENCES

- [1] T.-Y. Wu, H. Wu, M. Tang, S. Kumari, and C.-M. Chen, "CD-AKA-IoV: A provably secure cross-domain authentication and key agreement protocol for Internet of Vehicle," *Computers, Materials & Continua*, vol. 85, no. 1, pp. 1715-1732, 2025, doi: 10.32604/cmc.2025.065560.
- [2] X. Li, T. Liu, M. S. Obaidat, F. Wu, P. Vijayakumar, and N. Kumar, "A lightweight privacy-preserving authentication protocol for VANETs," *IEEE Systems Journal*, vol. 14, no. 3, pp. 3547-3557, 2020, doi: 10.1109/JSYST.2020.2991168.
- [3] J. Lee, G. Kim, A. K. Das, and Y. Park, "Secure and efficient honey list-based authentication protocol for vehicular ad hoc networks," *IEEE Transactions on Network Science and Engineering*, vol. 8, no. 3, pp. 2412-2425, 2021, doi: 10.1109/TNSE.2021.3093435.
- [4] S. A. Eftekhari, M. Nikooghadam, and M. Rafighi, "Security-enhanced three-party pairwise secret key agreement protocol for fog-based vehicular ad-hoc communications," *Vehicular Communications*, vol. 28, Art. no. 100306, 2021, doi: 10.1016/j.vehcom.2020.100306.
- [5] Y. Yang, L. Wei, J. Wu, C. Long, and B. Li, "A blockchain-based multidomain authentication scheme for conditional privacy preserving in vehicular ad-hoc network," *IEEE Internet of Things Journal*, vol. 9, no. 11, pp. 8078-8090, 2022, doi: 10.1109/JIOT.2021.3107443.
- [6] C.-M. Chen, Z. Li, S. Kumari, G. Srivastava, K. Lakshmana, and T. R. Gadekallu, "A provably secure key transfer protocol for the fog-enabled Social Internet of Vehicles based on a confidential computing environment," *Vehicular Communications*, vol. 39, Art. no. 100567, 2023, doi: 10.1016/j.vehcom.2022.100567.
- [7] Y. Gao, T. Zhou, W. Zheng, H. Yang, and T. Zhang, "High-availability authentication and key agreement for Internet of Things-based devices in Industry 5.0," *IEEE Transactions on Industrial Informatics*, vol. 20, no. 12, pp. 13571-13579, 2024, doi: 10.1109/TII.2024.3414443.
- [8] G. Ateniese, D. Francati, D. Nuñez, and D. Venturi, "Match me if you can: Matchmaking encryption and its applications," *Journal of Cryptology*, vol. 34, no. 3, Art. no. 16, 2021, doi: 10.1007/s00145-021-09381-4.
- [9] J. Shu, K. Yang, X. Jia, X. Liu, C. Wang, and R. H. Deng, "Proxy-free privacy-preserving task matching with efficient revocation in crowdsourcing," *IEEE Transactions on Dependable and Secure*

- Computing, vol. 18, no. 1, pp. 117-130, 2021, doi: 10.1109/TDSC.2018.2875682.
- [10] J. Xu, Z. Lin, and J. Wu, "Privacy-preserving task-matching and multiple-submissions detection in crowdsourcing," *Sensors*, vol. 21, no. 9, Art. no. 3036, 2021, doi: 10.3390/s21093036.
- [11] X. Xiang, J. Cao, and W. Fan, "Decentralized authentication and access control protocol for blockchain-based e-health systems," *Journal of Network and Computer Applications*, vol. 207, Art. no. 103512, 2022, doi: 10.1016/j.jnca.2022.103512.
- [12] R. Mukta, H.-Y. Paik, Q. Lu, and S. S. Kanhere, "A survey of data minimisation techniques in blockchain-based healthcare," *Computer Networks*, vol. 205, Art. no. 108766, 2022, doi: 10.1016/j.comnet.2022.108766.
- [13] Y. Zhang, X. Wei, J. Cao, J. Ning, Z. Ying, and D. Zheng, "Blockchain-enabled decentralized attribute-based access control with policy hiding for smart healthcare," *Journal of King Saud University - Computer and Information Sciences*, vol. 34, no. 10, pp. 8350-8361, 2022, doi: 10.1016/j.jksuci.2022.08.015.
- [14] M. Rasslan, M. M. Nasreldin, and H. K. Aslan, "Ibn Sina: A patient privacy-preserving authentication protocol in medical Internet of Things," *Computers & Security*, vol. 119, Art. no. 102753, 2022, doi: 10.1016/j.cose.2022.102753.
- [15] X. Liu, Z. Guo, J. Ma, and Y. Song, "A secure authentication scheme for wireless sensor networks based on DAC and Intel SGX," *IEEE Internet of Things Journal*, vol. 9, no. 5, pp. 3533-3547, 2022, doi: 10.1109/JIOT.2021.3097996.
- [16] N. C. Will and C. A. Maziero, "Intel Software Guard Extensions applications: A survey," *ACM Computing Surveys*, vol. 55, no. 14s, pp. 1-38, 2023, doi: 10.1145/3593021.
- [17] D. Dolev and A. Yao, "On the security of public key protocols," *IEEE Transactions on Information Theory*, vol. 29, no. 2, pp. 198-208, 1983, doi: 10.1109/TIT.1983.1056650.
- [18] R. Canetti and H. Krawczyk, "Analysis of key-exchange protocols and their use for building secure channels," in *Advances in Cryptology - EUROCRYPT 2001*, LNCS 2045, pp. 453-474, 2001, doi: 10.1007/3-540-44987-6_28.
- [19] R. C. Merkle, "A digital signature based on a conventional encryption function," in *Advances in Cryptology - CRYPTO 1987*, LNCS 293, pp. 369-378, 1988, doi: 10.1007/3-540-48184-2_32.
- [20] M. Burrows, M. Abadi, and R. Needham, "A logic of authentication," *ACM Transactions on Computer Systems*, vol. 8, no. 1, pp. 18-36, 1990, doi: 10.1145/77648.77649.
- [21] B. Blanchet, "Modeling and verifying security protocols with the applied pi calculus and ProVerif," *Foundations and Trends in Privacy and Security*, vol. 1, nos. 1-2, pp. 1-135, 2016, doi: 10.1561/3300000004.
- [22] European Parliament and Council of the European Union, "Regulation (EU) 2016/679 (General Data Protection Regulation)," *Official Journal of the European Union*, L119, pp. 1-88, 2016.
- [23] Republic of Korea, "Personal Information Protection Act," *Korean Law Information Center*, Act No. 10465, as amended.
- [24] ISO/IEC 27001:2022, *Information Security, Cybersecurity and Privacy Protection - Information Security Management Systems - Requirements*, ISO, 2022.
- [25] ISO/IEC 29100:2011, *Information Technology - Security Techniques - Privacy Framework*, ISO, 2011.
- [26] National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, 2023, doi: 10.6028/NIST.AI.100-1.
- [27] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, *Zero Trust Architecture*, NIST Special Publication 800-207, 2020, doi: 10.6028/NIST.SP.800-207.
- [28] OpenAPI Initiative, *OpenAPI Specification, Version 3.0.3*, Linux Foundation, 2020.
- [29] A. Shamir, "How to share a secret," *Communications of the ACM*, vol. 22, no. 11, pp. 612-613, 1979, doi: 10.1145/359168.359176.