

New Dimensions in Zero-Information Cryptography

Structured Ambiguity Commitments, Entropic Hyperplane Proofs, and the ENI6MA Embodiment

Frank Dylan Rosario and Lin (Grant) Wang

June 2026

Index Terms — structured ambiguity, zero-information proof of knowledge, Shannon secrecy, Vernam equivalence, sealed verifier, ENI6MA, post-quantum cryptography

NEW DIMENSIONS IN ZERO-INFORMATION CRYPTOGRAPHY

Abstract

This paper presents a unified account of the Rosario / Wang New Dimensions construction: a commitment and proof architecture in which the secret is not compressed into the public record and then protected by computational difficulty, but is absent from the observable channel under stated balance, freshness, and sealing assumptions. The construction joins three formal objects. A **Structured Ambiguity Container** (SAC) publishes a scaffold $\mathcal{F}$ that displays k^n plausible selector paths while withholding the private path σ . An **Entropic Hyperplane Protocol** (EHP) derives session-specific balanced manifolds and lets the prover answer in a synonym vocabulary through a private morphism $\varphi : \mathcal{H} \rightarrow \Sigma$. A **Secure Verification Environment** (SVE) performs membership checks and accumulation internally and emits only a nonce-scoped event bit Y . Under the Four Pillars P1-P4, namely path-complete projection, uniform private synonyms, zero mutual information, and session-bound freshness, the observable channel satisfies a Shannon-style zero-information condition:

$$I(J; O \mid C) = I(\varphi; O \mid C) = I(\sigma; O \mid C) = I(M; O \mid C) = 0.$$

Equation Narrative: This Shannon secrecy equation establishes that the conditional mutual information between any of the private coordinates (the active secret index J , private morphism φ , selector path σ , or underlying message content M) and the observable transcript O , given public context C , is exactly zero, ensuring complete privacy of all hidden parameters.

The claim is not that deployed software automatically inherits its ideal secrecy, nor that hash-based binding or side-channel resistance become information-theoretic. Rather, the ideal SAC/EHP/SVE model isolates a channel-shaping route to the same class of zero mutual information associated with the Vernam one-time pad, while the ENI6MA operational protocol supplies an engineering embodiment with explicit operational caveats: nonce reservation, τ binding, burn-before-validate semantics, proof payload tiers, and constant-time verification obligations. We give definitions, constructions, theorem statements, proof sketches, comparison tables, parameter guidance,

figure prompts for later visual generation, and appendices collecting notation, theorem dependencies, deferred proofs, and implementation constraints.

1. Introduction

The public-key revolution changed cryptography by changing where the secret could safely reside. Before public-key methods, the secret was often entangled with the whole mechanism; after Diffie-Hellman and RSA, the mechanism could be public and the secret could survive as a trapdoor hidden behind a function that was easy to compute and hard to invert. That move made the modern network possible. It also left a structural residue: in a computational public artifact, the secret remains determinately related to what is public. The adversary may not be able to compute the inverse today, but the inverse exists as a mathematical object. Security is therefore a statement about time, algorithms, and hardware.

The construction developed here pursues a different target. It asks whether a prover can publish or interact with a public object while leaving no recoverable signal about the hidden state in the public record. The aim is not to make inversion expensive. It is to arrange the channel so that the inverse problem is not well formed from the adversary's data. This is the relocation thesis of the New Dimensions line of work: long-horizon cryptographic confidence should be moved from the adversary's clock to the information content of the transcript, as illustrated in Figure 2.

The relevant precedent is Shannon's treatment of secrecy. In the Vernam one-time pad, perfect secrecy follows because every plaintext is equally consistent with a ciphertext when the key is uniform, secret, message-length, and single-use. In mutual-information form, $I(M; C) = 0$. That result does not depend on any bound on the adversary's computation. The cost is operational: a key at least as long as the message, distributed out of band, used exactly once. The Rosario / SAC-EHP-SVE construction is not an encryption pad and should not be described as the same mechanism. It is an orthogonal route to a channel-level equality. Instead of uniformity over a message-length key, it uses uniformity over a balanced public projection and a private synonym morphism. Instead of making the key single-use, it makes the session geometry and nonce single-use.

The construction has three layers. The first layer is a public

structured ambiguity scaffold. A table, board, matrix, wheel, holographic carrier, or digital manifold displays every candidate route. The private fact is not which cell exists, because all cells exist; it is which path the prover privately selects. The second layer is a proof interaction over an entropy-derived hyperplane manifold. The prover locates secret symbols in the current manifold but emits only private synonyms rather than public leaf labels. The third layer is a sealed verifier. The verifier inverts the synonym internally, checks membership internally, accumulates internally, and emits only a final event bit. The observer may record the challenge, the responses, and the verdict; in the ideal model, this record has surface entropy but zero effective mutual information about the hidden state.

This paper synthesizes the narrative statement of the construction, the formal SAC/EHP/SVE proof apparatus, and the ENI6MA implementation vocabulary. It follows the rhythm of contemporary cryptographic papers: introduction, technique overview, preliminaries, definitions, construction, theorem statements with proof sketches, comparison, engineering embodiment, historical perspective, and appendices. The emphasis is precision rather than maximal claim strength. Hiding of the selector path in the public scaffold is information-theoretic under the path-completeness and balance model. Binding of a published tag is computational when implemented by a hash. Replay resistance is strict in the ideal session-type model and ledger-enforced in deployment. The reference verification pipeline specifications are evidence of an implementation strategy, not a substitute for the ideal security proof.

1.1 The Cryptographic Inflection Point: From Perishable Computation to Structural Certainty

The four pillars are not a minor refinement of conventional cryptographic engineering. They describe a structural break with the assumption that has organized the entire field since the introduction of public-key methods in the 1970s. Conventional cryptography secures a message by compressing a secret into a ciphertext artifact and relying on the computational difficulty of inverting a mathematical calculation. Under this paradigm, the security of the system is not a property of the channel, it is a wager on the asymmetry between forward and inverse computation. This wager is, by construction, perishable. Every increase in computational throughput, every algorithmic improvement, every advance in quantum hardware reduces the time required to traverse the supposedly hard direction. The protection is not absolute; it is a depreciating asset whose half-life is set by Moore’s law in the classical regime and by Shor’s algorithm in the quantum one. A ciphertext encrypted today under an RSA, ECC, or even lattice-based assumption is, in effect, a time-bomb whose detonation date is whenever the relevant inverse calculation becomes tractable.

The alternative pursued by the four-pillar architecture moves the locus of security off the calculation and onto the channel itself. The objective is not to make the inverse calculation difficult but to ensure that the calculation simply has no input, that the public transcript contains zero correlation with

the hidden state. This is the standard set by Claude Shannon in 1949 when he proved that the one-time pad achieves *perfect secrecy*: an attacker who intercepts the ciphertext is left with every plaintext equally likely, and the recording therefore carries no extractable signal regardless of how much computation is brought to bear against it. Shannon’s result is the high-water mark of cryptographic guarantee, and it is unattackable by any amount of computation, classical or quantum, present or future, because there is nothing to compute against. The cost of this guarantee in OTP form, however, is the requirement of a key as long as the message, used exactly once, and distributed out of band. This cost is the reason perfect secrecy has remained a mathematical landmark rather than an operational technology, and it is the reason the entire modern stack has retreated to computational hardness as a substitute.

The four pillars described in this document recover Shannon’s perfect-secrecy guarantee without paying the OTP’s logistical price. They achieve this by relocating the source of the guarantee from the *unguessability of a key* to the *distribution geometry of the public channel*. The hidden state is never compressed into the wire; instead, the public scaffold simultaneously displays every possible secret, and the prover’s interaction with it leaves a transcript whose marginals are mathematically uniform under every possible interpretation. The trust assumption shifts from “the attacker cannot perform a calculation” to “the calculation has no purchase on the observable,” from a depreciating computational wager to a structural property of the channel itself. This relocation is not an optimization. It is a category change in what cryptographic security *is*, as illustrated in Figure 1.

Contributions

1. We define the SAC/EHP/SVE stack as a single primitive family with explicit public and hidden variables.
2. We state the Four Pillars P1-P4 as axioms of the zero-information channel.
3. We give construction algorithms for setup, commit, challenge, response, accumulation, and verification.
4. We state and sketch the main theorems: decomposed hiding, binding/hiding orthogonality, synonym uniformity, zero-information/Vernam equivalence, replay non-reusability, forgery bounds, flat-amplitude Grover defense, multi-secret anti-correlation, and reverse-accumulator soundness.
5. We separate ideal claims from deployed ENI6MA operational protocol obligations, including nonce_uuid, τ , proof payload tiers, reserve/burn lifecycle, and constant-time behavior.

2. Technique Overview

The central observation is that an authentication transcript can be useful to a verifier without being useful to an observer if the verifier possesses private structure that is never emitted. Public-key cryptography uses a trapdoor to make an otherwise

Two Roads to Shannon Perfect Secrecy ($I = 0$)




Figure 1. The Two Roads from Shannon. Conventional cryptography (top branch) compresses the secret into a public artifact and bets that the inversion remains expensive; security degrades monotonically as the adversary’s clock speeds up. The structured-ambiguity / sealed-verifier construction (bottom branch) shapes the channel so that the observable transcript is statistically independent of the hidden state; security is invariant under any future advance in computational power because the recorded data does not contain the secret in the first place.




Figure 2. The Relocation Thesis: Show the contrast between computational secrecy (inversion) and zero-information channel shaping (ambiguity).




Figure 3. Surface Bits Versus Effective Bits: Distinguish recordable transcript entropy from secret-relevant mutual information.




Figure 4. Four Pillars Dependency Graph: Visualize P1-P4 as a causal dependency graph ending in a zero-information plus no-replay event.

infeasible inversion feasible for the holder. The present construction uses a private morphism to make an impossible recovery unnecessary for the holder. The holder does not recover the hidden map from the transcript; the holder already has it. The observer does not merely lack enough compute to recover it; under the ideal balance assumptions, the observer lacks any statistical signal that distinguishes one map from another. The SAC layer begins with a public function

$$\mathcal{F} : I \times L \rightarrow V^d,$$

Equation Narrative: This function defines the public scaffold mapping $\mathcal{F}$, which maps each coordinate index from the set I and alphabet label from the set L to a d -dimensional payload vector in the vector space V^d . where $I = \{1, \dots, n\}$ indexes positions, $L = \{\ell_1, \dots, \ell_k\}$ is a public label set, and V^d is a payload space. A selector path $\sigma : I \rightarrow L$ chooses one label per position. Its induced section is

$$s_{\sigma}(i) = \mathcal{F}(i, \sigma(i)).$$

Equation Narrative: This equation defines the induced section $s_{\sigma}(i)$ at position i , mapping position index i to a public scaffold value using the private selector path $\sigma(i)$. The public scaffold displays every (i, ℓ) payload. If labels are balanced and no visible marker singles out the chosen label, then the scaffold alone leaves k^n paths equiconsistent. The conditional min-entropy of the path is $H_{\infty}(\sigma \mid \mathcal{F}) = n \log_2 k$ in the idealized public-scaffold model. A binding tag $D = H(\text{enc}(\sigma) \parallel \text{enc}(s_{\sigma}))$ may bind a committer to the path, but that binding layer is logically distinct from the hiding layer. The EHP layer turns this commitment into a proof of knowledge. Each session derives a manifold

$$\mathcal{M}_t = (H, \{L_i\}_{i=1}^H, \{W_i^t\}_{i=1}^H)$$

Equation Narrative: This defines the EHP session manifold $\mathcal{M}_t$ for round t , consisting of the foliation cardinality H , par-

tition labels L_i , and balanced witness sets W_i^t . from hidden entropy E , session time or nonce τ , salt ρ , round index t , and structural parameters Θ . The witness sets W_i^t are balanced:

$$|W_i^t| = \frac{|\mathcal{L}|}{H}$$

Equation Narrative: This defines the balanced witness set condition, specifying that each witness set W_i^t has a size exactly equal to the alphabet cardinality $|\mathcal{L}|$ divided by the foliation cardinality H . for every leaf i and round t , in the exact case. The prover locates a secret symbol in one of the leaves but does not name the leaf. Instead the prover emits a synonym

$$v_r^t = \varphi(h_{z_r^t})$$

Equation Narrative: This equation defines the generated synonym symbol v_r^t for round t at coordinate position r , mapped via private bijective morphism φ from the public leaf label $h_{z_r^t}$ for selected leaf index z_r^t . under a private bijective morphism $\varphi : \mathcal{H} \rightarrow \Sigma$, where the synonym vocabulary Σ is disjoint from both the public leaf labels and the eigenvalue rings. The SVE layer closes the feedback channel. It receives the submitted witness array $\mathbf{W}$, inverts each synonym internally, checks membership internally, and emits

$$Y = \mathbf{1}\{\Lambda = 1\} \in \{0, 1\}.$$

Equation Narrative: This defines the SVE output event bit Y as the indicator function of whether the internal verification accumulator Λ is equal to 1, yielding a single binary output. No mismatch index, partial score, round-by-round residue, or error localization leaves the sealed boundary. In the ideal model, this eliminates the per-round oracle that ordinary interactive protocols expose. In practical systems, this requirement becomes an engineering obligation: constant-time comparisons, fixed-shape error responses, burn-before-validate ledger behavior, and proof payloads that do not carry hidden witness material, as illustrated in Figure 8.

3. Preliminaries and Notation

Let λ denote a security parameter. Let $I = \{1, \dots, n\}$ be an index set, L a label set of size k , and $\Pi = L^I$ the selector-path space. Let $\mathcal{H} = \{h_1, \dots, h_H\}$ be the set of public hyperplane labels and $\Sigma = \{v_1, \dots, v_H\}$ a disjoint synonym set. The private morphism φ is a bijection from $\mathcal{H}$ to Σ . Unless otherwise stated, φ is sampled uniformly from the $H!$ bijections. We distinguish four kinds of data. The **public context** C includes the scaffold, public challenge geometry, non-secret parameters, circuit identifier, and, in implementation, the ledger-issued nonce reference. The **hidden substrate** includes σ , φ , E , ρ , Θ , and any selected secret family $\mathcal{K}$. The **observable transcript** O or $\mathbf{W}$ includes challenge material visible to the adversary, submitted synonyms or bearings, public proof payload fields, and final verdict visibility. The **sealed computation** consists of state and intermediate predicates visible only inside the SVE. The adversary is allowed full knowl-




Figure 5. Theorem Dependency Graph: Show how the main lemmas and theorems depend on one another.




Figure 6. Security Parameter Surface: Communicate how H , L , and R affect classical and Grover-bounded forgery probability.

edge of the algorithms and public parameters, complete passive recordings of all public messages, ledger-visible metadata, and unbounded offline computation. Quantum computation is included in this offline model. The adversary may also attempt online forgery and replay, but online attempts are mediated by verifier rate limits, nonce lifecycle, and the SVE acceptance predicate. Physical extraction of sealed state, compromise of entropy, malicious verifier design, UI capture that changes the human's choices, and side channels that violate sealing are outside the ideal proof and must be treated as implementation risks. We use Shannon mutual information in the standard way. A channel is zero-information for hidden variable X against observation O given public context C if

$$I(X; O | C) = 0.$$

Equation Narrative: This represents the general Shannon channel hiding condition, stating that the conditional mutual information between any hidden variable X and the observable transcript O is zero given public context C . Equivalently, the posterior distribution of X after observing O and C equals the prior distribution of X conditioned on C . For practical systems with small imbalance or side-channel residuals, we write $I(X; O | C) \leq \epsilon_X$, where ϵ_X is an explicit leakage budget, not an aesthetic margin. We also distinguish **surface bits** from **effective bits**. If a user emits one of H bearings, the observer records $\log_2 H$ surface bits per symbol. Those bits are not necessarily secret-relevant. The effective leakage is the mutual information between the hidden state and the observation. The

The Hyperplane Collapse: ~ 12.72 bits of board $\rightarrow \sim 2.585$ bits of decision.




Figure 7. The Hyperplane Collapse: how ~ 12.72 bits of board geometry are surjectively destroyed into a ~ 2.585 -bit decision. The three rotation rings ($\delta_1, \delta_2, \delta_3$) form a configuration space of $N_{\text{rot}} = 26 \cdot 26 \cdot 10 = 6760$ board states with $H(G) \leq 12.72$ bits of latent entropy. The non-invertible map $\kappa(a, z) = \chi(z)$ projects this state onto a 3×6 leaf array and finally collapses it into one of six color decisions, yielding $H(c^*) \leq 2.585$ bits. The amplification factor $\eta = 12.72 / 2.585 \approx 4.92$ board-bits discarded per decision-bit is exactly the structural absence of inversion exploited by the SAC/EHP/SVE construction: the observer’s transcript carries the colored decision residue while the latent rotation state remains information-theoretically irrecoverable.

observer’s notebook can fill with surface bits while the effective leakage remains zero in the ideal model:

$$H_{\text{surface}}(\text{bearing}) = \log_2 H, \quad I_{\text{effective}}(\text{secret}; O \mid C) = 0.$$

Equation Narrative: This relation separates surface entropy from effective leakage, establishing that while the surface entropy of observed bearings is positive at $\log_2 H$ bits, the effective leakage mutual information between the secret and the transcript O remains exactly zero.

4. Threat Model, Observer-Bit Budget, and Non-Claims

The construction is easiest to misunderstand when the adversary model is stated too weakly. A system whose claim is “the attacker cannot compute the answer quickly enough” can often describe the adversary in terms of polynomial time, circuit size, or classical versus quantum resources. A system whose claim is “the answer is absent from the transcript” must do something different: it must give the adversary every reasonable computational advantage and then account for what the adversary can actually observe. The proof is therefore not a race against an algorithm. It is an accounting exercise over random




Figure 8. SAC-EHP-SVE Layering: Depict the three-layer construction as a pipeline from public ambiguity to private synonym proof to sealed event bit.

Objects and Notation




Figure 9. Objects and notation overview. The public substrate (scaffold $\mathcal{F}$, hyperplane labels $\mathcal{H}$, challenge carrier C_τ , public payload fields, final event bit Y) is shown above the sealed boundary; the private substrate (selector path σ , private morphism φ , hidden entropy E , salt ρ , structural parameters Θ , secret family $\mathcal{K}$ with active index J) is shown below the boundary. Membership predicates, score functions, and accumulator state (Definition 9) live inside the SVE. The boundary is operative: anything that crosses the public side without preserving the uniform-marginal constraint is a leak, not a feature.

variables. Let $\text{View}_{\mathcal{A}}$ be the adversary’s view of one session. In the ideal model it may include the public scaffold $\mathcal{F}$, public challenge C_τ , public labels $\mathcal{H}$, submitted synonyms $\mathbf{W}$, public payload fields, final verdict visibility, and any ledger-visible status that is intentionally published. It does not include σ , φ , E , ρ , Θ , unredacted private configuration, SVE intermediate predicates, or the selected index J in a multi-secret family. The zero-information claim is a statement about the relation between hidden variables and $\text{View}_{\mathcal{A}}$:

$$I((J, \varphi, \sigma, M); \text{View}_{\mathcal{A}} \mid C) = 0$$

Equation Narrative: This equation establishes that the conditional mutual information between the adversary’s total view $\text{View}_{\mathcal{A}}$ and all hidden variables (secret index J , private morphism φ , path σ , and message M) is exactly zero given public context C . under exact balance and sealing. The condition is deliberately conditional on C , because public context is not hidden. The challenge geometry may be large, complex, and visually rich. Ledger records may prove that a nonce existed and was consumed. A screen recording may capture every bearing. None of this contradicts zero information if those variables are independent of the hidden state after conditioning on the public

challenge. This distinction is the **observer-bit model**. The observer may record positive surface entropy. With H bearings, every submitted symbol has at most $\log_2 H$ recordable bits. A length- L proof ceremony therefore exposes roughly $L \log_2 H$ surface bits plus public metadata and one verdict bit. These are bits in the notebook, not necessarily bits about the secret. Effective leakage is measured by mutual information. The design objective is:

$$H(\text{View}_{\mathcal{A}} \mid C) > 0 \quad \text{while} \quad I(\text{Hidden}; \text{View}_{\mathcal{A}} \mid C) = 0.$$

Equation Narrative: This joint condition illustrates the observer-bit model, showing that while the adversary’s view possesses strictly positive entropy given context C , its mutual information with any hidden parameters is exactly zero. The inequality on the left is unavoidable in a usable protocol. If the verifier and prover interact over a public channel, something is visible. The equality on the right is the cryptographic objective. A camera can record that a user pressed “up” or “left”; the question is whether that direction identifies the chamber, the symbol, the path, the morphism, or the active secret. Under a private uniformly sampled morphism and balanced foliation, the bearing is a uniformly distributed synonym for ev-

every possible chamber. The observer sees a symbol but lacks the private dictionary that would turn it into a chamber, and the public geometry gives no statistical preference to any hidden path. This also explains why the sealed verifier is not an implementation detail but a formal boundary. If the verifier emitted per-round mismatch indices, it would create an oracle. If it revealed partial scores, it would let the observer update a posterior distribution. If it returned distinguishable errors for “wrong chamber,” “wrong position,” or “wrong hash,” then the public channel would contain more than the final event bit. The SVE abstraction requires that all such distinctions remain internal. In deployment, the same requirement appears as constant-time comparison, fixed-shape responses, uniform failure messages, and careful logging policy, as illustrated in Figure 3. The model grants the adversary unbounded offline computation, but it does not grant impossible physical powers. If the adversary extracts the private morphism from memory, compromises the entropy source, reads the user’s secret directly, or subverts the verifier, then the hidden variables are no longer hidden. That is not a failure of Shannon’s equality; it is a failure of the implementation boundary that the equality assumes. Similarly, if a carrier renders the morphism visually, or if UI timing correlates with chamber search difficulty, then P2 or SVE sealing has been violated. The correct response is not to weaken the theorem, but to mark the deployment as outside the theorem’s hypotheses. The construction also does not claim that binding is information-theoretic when D is implemented by a hash. Binding and hiding are intentionally decomposed. The scaffold hiding argument says that $\mathcal{F}$ does not identify σ . The tag D makes later equivocation difficult under computational assumptions. If the hash function were broken, a committer might attack binding while the public scaffold would still fail to reveal the original path. Conversely, learning σ through endpoint compromise would not produce a hash collision. This orthogonality is a strength only when it is stated precisely. Nor does the construction claim that all zero-knowledge systems are obsolete. Conventional zero-knowledge protocols solve relation-proof problems with mature composition theory and efficient verification in many settings. The SAC/EHP/SVE architecture has a different objective: it minimizes secret-relevant information in an authentication or proof-of-knowledge transcript under a particular geometric and sealed-verifier model. The comparison is therefore about security basis and transcript semantics, not about replacing every proof system. Finally, the phrase “Vernam equivalence” must be read at the channel level. The construction is not a one-time pad, does not encrypt arbitrary messages by XORing them with a message-length key, and does not inherit every operational property of OTP. The equivalence is the mutual-information equality under the stated public context. OTP obtains $I(M; C) = 0$ through uniform single-use keys. SAC/EHP/SVE obtains $I(\text{Hidden}; O \mid C) = 0$ through path-complete projection, uniform private synonyms, and sealed residue-only verification. The mechanisms differ; the Shannon shape of the observable channel is the point of comparison.

Carrier independence is part of the same threat discipline. A carrier projection $\pi : \mathcal{M} \rightarrow \mathcal{C}$ is acceptable only if it preserves three invariants: public labels are accessible to the legitimate prover, witness elements are detectable, and the morphism remains concealed. A printed card, a mechanical wheel, a holographic plate, an acoustic cue, or a digital UI may all instantiate the same abstract manifold if they preserve those invariants. If a carrier makes labels hard to distinguish, completeness fails for honest users. If it makes witness elements ambiguous, soundness and completeness become noisy. If it embeds the synonym map in visible layout, the zero-information proof fails. Carrier substitution is therefore not a claim that any medium works; it is a theorem schema saying that security depends on invariants rather than material. The ENI6MA embodiment should be evaluated in this language. Its challenge matrix, bearings, proof hash, nonce UUID, and ledger receipts are public context. Its private morphism, sealed configuration, entropy, and membership checks are hidden substrate. Its validate-proof outcome is a public event bit. A good deployment keeps these categories separate in code, payload schemas, logs, UI state, and operational dashboards. A bad deployment accidentally turns hidden substrate into public context and then wonders why the theorem no longer applies.

5. Formal Objects

Definition 1 (Structured Ambiguity Container). A Structured Ambiguity Container is a tuple $(\mathcal{F}, D, I, L, V^d)$ where $I = \{1, \dots, n\}$, $L = \{\ell_1, \dots, \ell_k\}$, V^d is a payload vector space, $\mathcal{F} : I \times L \rightarrow V^d$ is public, and D is a binding tag computed from a hidden selector path. The public commitment is $C_{\text{SAC}} = (\mathcal{F}, D)$. The structure is illustrated in Figure 11.

Definition 2 (Selector Path and Induced Section). A selector path is a function $\sigma : I \rightarrow L$. Its induced section is $s_\sigma(i) = \mathcal{F}(i, \sigma(i))$. The path space has size $|\Pi| = |L|^I = k^n$.

Definition 3 (Ambiguity Class). The ambiguity class of a scaffold is

$$\mathcal{S}(\mathcal{F}) = \{s_\sigma : \sigma \in L^I\}.$$

Equation Narrative: This equation defines the ambiguity class $\mathcal{S}(\mathcal{F})$ of the scaffold $\mathcal{F}$ as the set of all possible induced sections s_σ generated by all possible selector paths σ in L^I . The class is path-complete when every $\sigma \in L^I$ is represented by a valid induced section and no public marker privileges one section over another. **Definition 4 (Binding Tag).** A binding tag is a public value

$$D = H(\text{enc}(\sigma) \parallel \text{enc}(s_\sigma)).$$

Equation Narrative: This equation defines the binding tag D as the cryptographic hash H of the concatenated serializations of the selector path σ and its induced section s_σ , representing computational binding. In the standard instantiation, D provides computational binding under collision resistance of H . It is not the source of information-theoretic hiding. The **Structured Ambiguity Commitment** layer should therefore be read separately from proof and verification. The public commit-




Figure 10. Flow of information in a structured-ambiguity authentication system. The prover holds a private selector path σ and a private morphism φ , neither of which is ever transmitted. The public scaffold $\mathcal{F}$, the per-session challenge C_τ , the submitted synonyms $\mathbf{W}$, and the final event bit Y flow over the open channel. The morphism φ is shared with the verifier or the SVE through an out-of-band channel (dashed). The recorded transcript carries surface entropy but, by Theorem 4, zero mutual information about σ or φ .

ment contains the whole scaffold of possible paths, plus any public tag that binds an opening, while the selected path σ is not observable. EHP and SVE later use a session foliation and private morphism φ to prove knowledge against that commitment; they do not make σ or φ public. Changing the foliation cardinality H changes the EHP challenge geometry and proof surface, but the SAC path ambiguity remains governed by n, k , and the private selector σ . **Definition 5 (Private Morphism).** A private morphism is a bijection $\varphi : \mathcal{H} \rightarrow \Sigma$ such that $\Sigma \cap \mathcal{H} = \emptyset$ and $\Sigma \cap A_j = \emptyset$ for every eigenvalue ring A_j . It is shared by prover and verifier or held by prover and SVE, but is never emitted, hashed into a public recoverable value, or rendered on the public carrier. **Definition 6 (Balanced Foliation).** For witness sets W_i^t on H leaves in round t , the foliation is balanced when

$$|W_i^t| = \frac{|\mathcal{L}|}{H}$$

Equation Narrative: This defines the balanced witness set condition, specifying that each witness set W_i^t has a size exactly equal to the alphabet cardinality $|\mathcal{L}|$ divided by the foliation cardinality H . for all i and t , and when placement is uniform conditioned on the public challenge. With imbalance Δ around target count κ , practical leakage is bounded as a function of Δ/κ . **Definition 7 (Entropic Hyperplane Protocol).** EHP is the interactive protocol in which session geometry is derived from $(E, \tau, \rho, t, \Theta)$, the prover finds leaf labels contain-

ing secret symbols, translates those labels under φ , and submits synonyms to an SVE. **Definition 8 (Secure Verification Environment).** An SVE is a sealed stateful verifier with hidden inputs $(K, \varphi, E, \tau, \rho, \Theta)$, submitted witness array $\mathbf{W}$, internal geometry reconstruction, internal morphism inversion, internal membership checks, and public output $Y \in \{0, 1\}$. The SVE is residue-only if no intermediate predicate, score, mismatch index, timing variation, or exception class is externally distinguishable, as illustrated in Figure 12.

Definition 9 (Reverse Witness Accumulator). For a secret family $\mathcal{K} = (K_1, \dots, K_m)$, a reverse witness accumulator is a deterministic sealed procedure

$$\text{RWA} : (\mathcal{P}, \mathcal{K}, E, \rho, \Theta) \rightarrow \{0, 1\}$$

Equation Narrative: This maps the Reverse Witness Accumulator function RWA from proof payload $\mathcal{P}$, secret family $\mathcal{K}$, hidden entropy E , salt ρ , and parameters Θ to a binary verification decision in $\{0, 1\}$. that parses a proof payload, derives the session manifold from τ , rebuilds witness sets, inverts synonyms, computes score functions over candidate secrets, and




Figure 11. The structured-ambiguity scaffold $\mathcal{F}$. All $n \times k$ values are public and openly displayed. The selector path σ (highlighted) is private. From $\mathcal{F}$ alone, all k^n paths are equiconsistent; the secret is information-theoretically hidden. A binding tag $D = H(\text{enc}(\sigma) \parallel \text{enc}(s_\sigma))$ binds the committer to the chosen path under collision resistance of H but does not identify the path itself.

emits a single verdict. Its score function may be written

$$\text{Score}(K_j) = \sum_{t=1}^R \sum_{r=1}^L \mathbf{1}\{K_j[r] \in W_{\hat{z}_r^t}^t\}.$$

Equation Narrative: This equation defines the candidate secret score function $\text{Score}(K_j)$ as the sum over rounds t and secret coordinates r of whether each candidate character $K_j[r]$ is contained in the reconstructed witness set $W_{\hat{z}_r^t}^t$ for inverted synonym index $\hat{z}_r^t$. The index $J^* = \arg \max_j \text{Score}(K_j)$ remains sealed.

6. Four Pillars and the Zero-Information Channel

The construction is easiest to audit through four pillars. They are not four unrelated defenses. They form a dependency chain in which each condition supports the next. **P1, Path-Complete Geometric Projection.** The public challenge C enumerates every possible path that any member of the relevant secret family could take. For path space $\Pi(C)$,

$$\forall j \in \{1, \dots, m\}, \quad \sigma_j \in \Pi(C)$$

Equation Narrative: This represents the path-completeness projection property, asserting that the private selector path σ_j for every secret j in the family belongs to the path space $\Pi(C)$ defined by public challenge C . In every valid session. If the scaffold omitted a path, the absence would carry information. Completeness is the precondition for neutrality. **P2, Uniform**

Private Synonym Marginals. Conditioned on the public challenge, every private synonym has uniform marginal probability:

$$\Pr[v_r^t = v_k \mid h_{z_r^t} = h_i, C] = \frac{1}{H}$$

Equation Narrative: This represents the uniform private synonym marginals property, showing that the conditional probability of emitting any synonym symbol v_k given selected leaf label h_i and challenge C is exactly $1/H$. for all i, k, r, t . Balanced foliation and uniform morphism sampling are the structural sources of this equality. Frequency analysis fails because the transcript has no preferred direction. **P3, Zero Mutual Information.** Given public context C , the observable transcript carries zero mutual information about the active secret index J , morphism φ , selector path σ , and message or content variable M :

$$I(J; O \mid C) = I(\varphi; O \mid C) = I(\sigma; O \mid C) = I(M; O \mid C) = 0.$$

Equation Narrative: This asserts the zero mutual information axiom, indicating that the conditional mutual information between any private parameters (index J , morphism φ , path σ , message M) and public transcript O is exactly zero given context C . This is the structural-absence claim. The data needed for recovery is not hidden in the transcript; it is not present. **P4, Session-Bound Freshness and Non-Replay.** Every proof event is bound to a freshness coordinate τ and session geome-




Figure 12. The sealed verifier as an information-theoretic one-way channel. Per-round membership results $M_{r,t} = \mathbf{1}\{K[r] \in W_{z_t}^t\}$ are accumulated internally as $a_i = H(a_{i-1} \parallel M_i)$; the only observable output is the final event-validity bit Y . No per-round oracle is exposed, so the recorded exchange is statistically independent of the hidden morphism ϕ . The constant-time comparison at the end ensures that a mismatch at the first position and a mismatch at the last position produce indistinguishable outputs (in time and shape).

try. Ideally,

$$\Pr[\text{Verify}(O_{\tau_a}, \tau_b) = 1] = 0 \quad \text{for} \quad \tau_a \neq \tau_b.$$

Equation Narrative: This defines the session-bound freshness and non-replay property, establishing that the probability of the verification function accepting a transcript O_{τ_a} under a different session coordinate τ_b is exactly zero. In deployed systems this property is approximated and enforced by nonce reserve, TTL, τ binding, burn-before-validate semantics, and proof reconstruction against the live session context. Together P1 and P2 imply the uniform transcript distribution required by P3. P4 prevents the only useful thing an observer might try with a captured transcript: replay. In OTP language, secrecy and freshness are coupled through key single-use. Here secrecy is carried by projection and synonym uniformity, while freshness is carried by the per-session scaffold and nonce life-cycle, as illustrated in Figure 4.

7. Construction

We now specify the construction in algorithmic form. The notation is intentionally close to modern cryptographic syntax while keeping the hidden/public split visible. The single-character data path through the construction is illustrated in Figure 13, and the synonym-zero-capacity property that justifies the proof-of-knowledge guarantee is illustrated in Figure 14.

Algorithm 1, $\text{SAC.Setup}(1^\lambda, n, k, H)$. Choose public parameters (I, L, V^d, H) and balance constraints. Select or gen-

erate a public scaffold $\mathcal{F} : I \times L \rightarrow V^d$ satisfying path completeness and observable balance. Publish the scaffold format and verification parameters. Keep no secret in setup unless a particular embodiment uses sealed generation state. **Algorithm 2,** $\text{SAC.Commit}(\mathcal{F}, \sigma)$. Given a selector path $\sigma : I \rightarrow L$, compute $s_\sigma(i) = \mathcal{F}(i, \sigma(i))$ for all i . Compute $D = H(\text{enc}(\sigma) \parallel \text{enc}(s_\sigma))$. Output public commitment $(\mathcal{F}, D)$, where $\mathcal{F}$ still displays every path in L^I , and retain σ as private witness. This is a commitment step, not a verification transcript; proof soundness is supplied later by EHP/SVE membership checks. **Algorithm 3,** $\text{EHP.Challenge}(E, \tau, \rho, \Theta, R)$. For each round $t \in \{1, \dots, R\}$, derive ring offsets

$$o_{j,t} = \text{PRF}(E_{j,t}, \tau, \rho, t, j) \bmod |A_j|.$$

Equation Narrative: This equation computes the ring rotation offset $o_{j,t}$ for ring j in round t using a pseudorandom function PRF on private round entropy $E_{j,t}$, session freshness timestamp τ , salt ρ , round t , and ring index j , modulo the ring alphabet size $|A_j|$. for each eigenvalue ring A_j . Rotate the rings by those offsets,

$$A_j^{(t)} = \text{Rotate}(A_j, o_{j,t}).$$

Equation Narrative: This equation defines the rotated ring alphabet $A_j^{(t)}$ as the base ring alphabet A_j rotated by the computed offset $o_{j,t}$. Then foliate each rotated ring into H balanced

Anatomy of One Character Entry




Figure 13. Single-character pipeline. A secret character $K[r]$ is located in one of H balanced witness sets $W_{z_r^t}^t$ at round t ; the prover never names the leaf z_r^t but emits $v_r^t = \varphi(h_{z_r^t})$ from a synonym vocabulary Σ disjoint from the public labels and the eigenvalue rings (equation 4). The SVE recomputes the geometry from (E, τ, ρ, Θ) , inverts the synonym internally as $\hat{z}_r^t = \varphi^{-1}(v_r^t)$, and appends $M_{r,t} = \mathbf{1}\{K[r] \in W_{\hat{z}_r^t}^t\}$ to a fixed-shape accumulator. Only the final indicator $Y \in \{0, 1\}$ leaves the sealed boundary.

zones:

$$A_j^{(t)} = Z_{j,1}^{(t)} \sqcup \cdots \sqcup Z_{j,H}^{(t)}.$$

Equation Narrative: This represents the zone foliation of the rotated alphabet ring $A_j^{(t)}$ as a disjoint union $\sqcup$ of H balanced zone slices $Z_{j,w}^{(t)}$. The stochastic random pool, session freshness coordinate τ , salt ρ , round index, and ETP-style freshness derivation supply the entropy for the modulo offsets. In the ENI6MA codebook generator this appears operationally as: rotate each alphabet ring by a per-session offset, then split the rotated alphabet into hyperplane_count zones. Publish the challenge carrier C_τ and public labels. Do not publish E , ρ , Θ , or φ . **Algorithm 4**, EHP.Respond(K, φ, C_τ). For each secret position r and round t , find a leaf index z_r^t such that $K[r] \in W_{z_r^t}^t$. Emit $v_r^t = \varphi(h_{z_r^t})$. Output the witness array $\mathbf{W} = (v_r^t)_{r,t}$ and any public integrity hash required by the payload contract. **Algorithm 5**, SVE.Accumulate($\mathbf{W}, K, \varphi, E, \tau, \rho, \Theta$). Reconstruct the same witness sets internally. For each submitted synonym, compute $\hat{z}_r^t = \varphi^{-1}(v_r^t)$ and

$$M_{r,t} = \mathbf{1}\{K[r] \in W_{\hat{z}_r^t}^t\}.$$

Equation Narrative: This defines the coordinate match bit $M_{r,t}$ at round t position r as the indicator function $\mathbf{1}$ showing whether the secret candidate character $K[r]$ belongs to the reconstructed witness set $W_{\hat{z}_r^t}^t$ for the inverted synonym index $\hat{z}_r^t$. Accumulate with a fixed-shape operation such as

$$a_0 = H(C_\tau), \quad a_i = H(a_{i-1} \parallel M_i), \quad \Lambda = \bigwedge_{r,t} M_{r,t}.$$

Equation Narrative: This set of equations outlines the SVE internal accumulation: the initial hash state a_0 is the hash H of the challenge carrier C_τ ; intermediate state a_i is the hash of the previous state a_{i-1} concatenated with the round match M_i ; and the final verification accumulator Λ is the logical conjunction $\bigwedge$ of all coordinate matches. No intermediate $M_{r,t}$ or a_i leaves the SVE. **Algorithm 6**, Verify. Emit $Y = \mathbf{1}\{\Lambda = 1\}$, or in threshold deployments emit

$$Y = \mathbf{1} \left\{ \max_j \text{Score}(K_j) \geq T_{\text{accept}} \right\}.$$

Equation Narrative: This defines the threshold verification output event bit Y as the indicator function of whether the maximum candidate secret score is at least equal to the accep-

The Witness Synonym: a Zero-Capacity Channel




Figure 14. Synonym channel zero-capacity. Conditioned on the public challenge C , the marginal distribution of every emitted synonym is uniform on Σ (Lemma 3). The diagram shows the joint distribution of the hidden leaf $h_{z_r^t}$ and the emitted synonym v_r^t as a doubly-stochastic matrix in which every row and every column sums to one: the channel from leaf to synonym carries zero mutual information about the leaf, and therefore zero mutual information about the secret. This is the structural reason Theorem 4 holds without any computational assumption.

tance threshold T_{accept} . The output is nonce-scoped. A proof accepted under τ is not a reusable credential.

8. Theorems and Proof Sketches

Theorem 1 (Decomposed-Layer Hiding). For a path-complete SAC with n coordinates and k labels per coordinate,

$$H_{\infty}(\sigma \mid \mathcal{F}, D) \geq n \log_2 k - \varepsilon_{\text{bind}},$$

Equation Narrative: This represents the decomposed-layer hiding min-entropy theorem, establishing that the conditional min-entropy of selector path σ given public scaffold $\mathcal{F}$ and binding tag D is bounded below by the ideal entropy $n \log_2 k$ minus the computational binding leakage budget $\varepsilon_{\text{bind}}$, where $\varepsilon_{\text{bind}}$ accounts for leakage through a non-ideal binding tag.

Proof sketch. The scaffold exposes all k^n paths symmetrically. In the ideal tag model, D does not identify which path was selected; it only binds a later opening. The adversary's posterior over σ therefore remains uniform over the ambiguity class. Any tag leakage is isolated in $\varepsilon_{\text{bind}}$. ■ **Theorem 2 (Binding/Hiding Orthogonality).** Breaking the computational binding tag does not reduce the information-theoretic hiding of the scaffold, and learning the selector path from a compromised endpoint does not imply a hash collision. **Proof**

sketch. Hiding is a property of the public distribution of $\mathcal{F}$ over paths. Binding is a property of collision resistance for D . The two games depend on disjoint mechanisms. A collision can let a malicious committer equivocate, but it does not mark the original path in $\mathcal{F}$. Conversely, endpoint disclosure of σ gives no method for finding a second preimage for D . ■

Lemma 3 (Uniform Synonym Marginals). Under balanced foliation and uniformly sampled φ ,

$$\Pr[v_r^t = v_k \mid h_{z_r^t} = h_i, C] = \Pr[v_r^t = v_k \mid C] = \frac{1}{H}.$$

Equation Narrative: This expresses the uniform synonym marginals lemma, establishing that the probability of emitting any synonym v_k is exactly $1/H$, completely independent of the selected leaf label h_i under public challenge C . **Proof sketch.** Condition on any leaf identity. Since φ is uniform over all $H!$ bijections, exactly $(H-1)!$ morphisms map that leaf to each synonym. The probability of every synonym is therefore $(H-1)!/H! = 1/H$. Balance ensures that the leaf identity itself is not biased by observable geometry. ■ **Theorem 4 (Zero-Information / Vernam Equivalence).** Under P1-P3, the observable channel is Shannon-equivalent to a Vernam channel

at the level of mutual information:

$$I(J; O | C) = I(\varphi; O | C) = I(\sigma; O | C) = I(M; O | C) = 0.$$

Equation Narrative: This represents the zero mutual information theorem, showing that the conditional mutual information between any private variables (index J , morphism φ , path σ , message M) and public transcript O is exactly zero given context C . **Proof sketch.** The transcript distribution is a product of uniform synonym draws conditioned on C . For the morphism coordinate,

$$I(\varphi; \mathbf{W} | C) = H(\mathbf{W} | C) - H(\mathbf{W} | \varphi, C) = LR \log_2 H - LR \log_2 H = 0.$$

Equation Narrative: This shows that the conditional mutual information between private morphism φ and submitted synonym array $\mathbf{W}$ given C is exactly zero, as the entropy $H(\mathbf{W} | C)$ is equal to the conditional entropy $H(\mathbf{W} | \varphi, C)$, which both equal $LR \log_2 H$ for secret length L and rounds R . The same symmetry applies to selector paths and active-secret indices because P1 makes all paths present and P2 makes observable synonyms uniform over the path choices. This matches the Shannon form of the OTP proof while using a different source of uniformity. ■ **Theorem 5 (Nested Factorial Reconstruction Barrier).** An observer attempting to recover hidden state from $\mathbf{W}$ faces

$$W_{\text{recon}} = \Omega((H!)^2)$$

Equation Narrative: This represents the reconstruction work bound W_{recon} as asymptotically bounded below by $\Omega((H!)^2)$ for foliation cardinality H , representing the factorial barrier to recovering the private morphism. candidate morphism and leaf-assignment pairs, all equiconsistent with the transcript, as illustrated in Figure 15 and the synonym-masking diagram in Figure 16.

Proof sketch. There are $H!$ candidate morphisms. For each morphism, balanced geometry leaves a factorial-scale assignment ambiguity over leaf interpretations. By Theorem 4 the transcript does not rank these candidates. The point is not merely the size of the space; it is the absence of a distinguishing statistic over it. ■ **Theorem 6 (Flat-Amplitude Grover Defense).** Grover search gives no advantage for morphism recovery under the ideal zero-information model. **Proof sketch.** Grover requires an oracle that marks targets differently from non-targets. Define an oracle that marks candidate morphisms consistent with $\mathbf{W}$. Under Theorem 4, every morphism is equally consistent, so the oracle is flat. If all states are marked, amplitude amplification has no direction. For online forgery, Grover can still halve a brute-force exponent; parameter selection must account for that separately. ■ **Theorem 7 (Replay Non-Reusability).** If session geometry is bound to τ and verifier state consumes τ after first use, then a transcript accepted for τ_a is rejected for $\tau_b \neq \tau_a$ except with bounded implementation failure probability:

$$\Pr[\text{Verify}(O_{\tau_a}, \tau_b) = 1] \leq 2^{-\lambda}.$$

Equation Narrative: This bounds the probability of a replay attack succeeding across different sessions by $2^{-\lambda}$, where λ is the security parameter. **Proof sketch.** The witness sets are functions of (E, τ, ρ, Θ) . Changing τ changes the membership relation the transcript must satisfy. In an ideal type model the old transcript does not type-check against the new geometry. In deployed ENI6MA, the ledger additionally burns the nonce before validation, returning replay or expiry errors for reused or stale nonces. ■ **Theorem 8 (Forgery Bound).** For an adversary lacking K and φ ,

$$\Pr[\text{forge}] \leq H^{-LR}$$

Equation Narrative: This defines the forgery probability bound, showing that the probability of an online forge attempt succeeding without knowledge of the private path is bounded above by H^{-LR} for hyperplane foliation H , proof length L , and rounds R . classically and $\Pr[\text{forge}] \leq H^{-LR/2}$ under a Grover-bounded online search model. **Proof sketch.** Each of the LR checks has success probability at most $1/H$ without the leaf-to-synonym correspondence and secret membership knowledge. Independence gives H^{-LR} . Quantum search can reduce the exponent for a yes/no online target but cannot create transcript information where mutual information is zero. ■ **Theorem 9 (Multi-Secret Anti-Correlation).** For an equal-length secret family $\mathcal{K} = (K_1, \dots, K_m)$ with private session index J , the observable transcript distribution is invariant in J :

$$I(J; O | C) = 0.$$

Equation Narrative: This states that the mutual information I between the active secret index J and the observable transcript O is exactly zero given context C . Moreover, the single-generator predicate needed by frequency and correlation attacks is indeterminate. **Proof sketch.** Equal-length secrets generate the same number and shape of submitted synonyms. Balanced foliation and private morphism uniformity make the marginal transcript distribution identical for every K_j . Since the observer cannot determine whether two sessions used the same secret or different secrets, aggregation has no stable generator to correlate against. ■ In the ENI6MA example configuration discussed below, the equal-length family for this theorem is 000000, 123456, and SECRET. The additional example secret Ab1Cd2Ef3Gh4 has length 12 and is outside the equal-length theorem unless it is normalized, padded, or placed in a separate equal-length family. **Theorem 10 (Reverse-Accumulator Soundness).** In strict mode, a forger lacking every secret in $\mathcal{K}$ is accepted by the reverse witness accumulator with probability at most

$$m \cdot H^{-LR},$$

Equation Narrative: This represents the union-bounded forgery probability for an active family of m secrets, bounded by $m \cdot H^{-LR}$ where H is the foliation cardinality, L is the secret length, and R is the number of rounds. tightened under anti-




Figure 15. The coupled factorial reconstruction barrier. The morphism space ($H!$ bijections) and the per-round leaf assignment space ($H!$ permutations for each round, balanced under P2) are coupled through the synonym channel into a $\Omega((H!)^2)$ candidate space. Crucially, by Theorem 4 (zero-information / Vernam equivalence), every one of these candidates is exactly equiconsistent with the recorded transcript $\mathbf{W}$. The barrier is therefore not merely combinatorial but flat-amplitude: there is no statistic over the candidate space that the observer could compute to distinguish one element from another, which is precisely the condition that defeats Grover amplification (Theorem 6).

correlation assumptions to the stated family-amplified bound.

Proof sketch. The reverse accumulator tests each candidate secret internally and emits only whether some score reaches threshold. A union bound over m secrets gives mH^{-LR} in strict mode. The anti-correlation condition reduces the adversary's usable predicate space because transcript families are not attributable to a single generator. The matched-secret index remains sealed, so acceptance does not reveal which candidate, if any, matched. ■

The logical dependencies between our core lemmas and theorems are illustrated in Figure 5, and an alternative proof-lemma graph emphasizing the information-theoretic core is given in Figure 17.

8.1 Unbounded Composition and the Anti-Unicity Property

A perfect-secrecy primitive is only as valuable as its composition behavior. The OTP achieves perfect secrecy on a single message but degenerates under composition: reuse of the key collapses the guarantee, and the operational solution, fresh keys per message, is precisely what makes the OTP impractical for high-volume authentication traffic. The present architecture handles composition through a different mechanism entirely, one that does not require fresh secrets across sessions and one that exhibits a property the OTP cannot: the *inversion*

of the unicity-distance dynamic.

The composition primitive is the **fresh nonce triple**: a per-session injection of micro-timestamps, per-round entropy $\mathcal{E}_{j,t}$, and a regenerated geometric scaffold C_τ . Each of these three components contributes independent freshness. The micro-timestamps bind every proof event to a unique moment in the verifier's clock; the per-round entropy reshuffles the manifold projection on every round of every session; the scaffold regeneration ensures that the very geometry against which the prover interacts is renewed each session. The freshness is structural at three layers, not one, and crucially, the morphism φ itself need not be refreshed for any of this to hold. This is the operational substance of Axiom P4 and the source of Theorem 7's replay non-reusability property: the freshness has been moved off the key and onto the session, which is exactly the relocation that frees the construction from the OTP's single-use constraint.

The composition behavior has a striking inverse property, which we will call the **anti-unicity property**. In conventional cryptography, the *unicity distance*, Shannon's own measure of how much ciphertext an attacker must accumulate before the key is uniquely determined, is finite and breachable by accumulation. Every additional intercepted ciphertext narrows the candidate key space until, at the unicity distance, only one key remains consistent with the observed traffic. The attacker's job is to wait for enough material. The present architecture in-

Two Independent Hiding: Synonym (cognitive) then One-Time Pad (transport).




Figure 16. Synonym masking from witness to bit. For each round t , the prover's true leaf $h_{z_r^t}$ is masked by the private morphism φ before it reaches the wire; the observer sees only v_r^t , an element of a synonym vocabulary disjoint from both the public labels and the eigenvalue rings (equation 4). Because φ is uniform over the $H!$ bijections, the channel from leaf to synonym is a uniform double-stochastic mask, and the recorded synonyms are statistically independent of the underlying leaf. The bit that finally leaves the SVE (Y) is a function of all sealed match indicators, not of any individual synonym.

verts this dynamic entirely. Because each session's transcript is conditionally independent of the secret given the scaffold ($I(J; O | C) = 0$), no amount of accumulation reduces the candidate set. An attacker who records ten transcripts and an attacker who records ten million transcripts are in the same epistemic position: every secret in the family is equally consistent with every transcript, and the joint distribution over multiple sessions remains a product of uniform marginals.

The accumulation of data does not narrow the search space; it merely confirms, with increasing statistical confidence, that the channel contains no signal to extract. The architecture moves trust from the hope that a calculation is difficult to invert to the **structural certainty that the communication channel contains no extractable information**, and that certainty is invariant under any future advance in computational power. The per-character ledger view of this phenomenon is given in Figure 18, the ceremony-growth view in Figure 19, the pass/fail accumulator structure in Figure 20, and the two-ledger separation in Figure 21.

9. Comparison and Parameters

The following comparison separates security basis, public-record exposure, quantum posture, and replay behavior. Representative parameters show the separation between hiding in the scaffold and soundness in the proof interaction.

Choosing the Foliation Cardinality H

The foliation cardinality H is selected at commitment or configuration time, not fixed by the theory. A deployment chooses H according to the desired security level, prover usability, alphabet balance, and carrier constraints. Larger H increases per-check online forgery resistance and enlarges the private morphism space, but it can make a human carrier harder to scan and can make alphabet balancing harder when ring sizes are not multiples of H . Smaller H may improve ergonomics or display clarity, but requires compensating with larger L , larger R , stricter rate limits, or a narrower operating context. For example, $H = 6, L = 7, R = 10$ gives a classical online forgery floor $H^{-LR} = 6^{-70} \approx 2^{-181}$ and a Grover-bounded floor $H^{-LR/2} = 6^{-35} \approx 2^{-90}$. Increasing H, L , or R changes soundness without changing the zero-information proof structure. Separately, $n = 128, k = 4$ gives $4^{128} = 2^{256}$ possible selector paths in the public scaffold model, as illustrated in Figure 6.

Worked Concrete Example

To make the construction definite, take a small scaffold with $n = 4$ positions and $k = 4$ labels $\{A, B, C, D\}$ per position; the public table $\mathcal{F}$ therefore displays $n \times k = 16$ values in the open.




Figure 17. Proof and lemma dependency graph (information-theoretic view). Lemma 3 (uniform synonym marginals) is the structural fixed point: from it follow Theorems 4 (zero-information / Vernam equivalence), 5 (nested factorial reconstruction barrier), and 6 (flat-amplitude Grover defense). Theorem 7 (replay non-reusability) is independent of the hiding chain and rests on session geometry plus nonce burn. Theorems 8–10 (forgery, multi-secret, reverse accumulator) build on Lemma 3 and on Theorem 4 to give bounded online forgery and union-bounded multi-secret soundness.

Let the committer’s private path be $\sigma = (B, D, A, C)$, selecting one label at each position. From $\mathcal{F}$ alone an observer sees all sixteen values but cannot tell which label is “true” at any position; the number of equiconsistent paths is $k^n = 4^4 = 256$, and the unconditional min-entropy is $H_\infty(\sigma | \mathcal{F}) = n \log_2 k = 4 \cdot 2 = 8$ bits. At deployment scale ($n = 128, k = 4$) the same expression yields 256 bits, and the equiconsistent count $4^{128} \approx 10^{77}$ is of the order of the number of atoms in the observable universe.

For the proof of knowledge, fix $H = 6$ hyperplane leaves and draw a private morphism φ uniformly from the $H! = 720$ bijections onto a disjoint synonym vocabulary $\Sigma = \{v_1, \dots, v_6\}$ satisfying $\Sigma \cap \mathcal{R} = \emptyset$ and $\Sigma \cap A_j = \emptyset$ for every eigenvalue ring A_j . With $L = 7$ trials per round over $R = 10$ rounds, the forgery bound $\Pr[\text{forge}] \leq H^{-LR}$ gives $6^{-70} \approx 2^{-181}$, while uniform synonym marginals make every one of the 720 candidate morphisms produce an identical transcript distribution. An eavesdropper who records the entire session is left with the coupled barrier $W_{\text{recon}} = \Omega((6!)^2) \approx 5.2 \times 10^5$ candidate pairs, *every one of which is equally consistent with what was recorded*. The toy parameters already deliver 181 bits of classical and 90 bits of post-quantum security from a layout small enough to be realized with physical cards on a

table.

Information-Theoretic Security and the P versus NP Question

It is worth recalling how the appeal to computational complexity arose, because the present construction represents a return to an older and firmer standard. During the era of public-key cryptography the security of a system was argued by showing that an attack would require solving a problem in a class believed to be intractable. A function was said to belong to the class **P** if it could be computed in time bounded by a polynomial in the length of its input, and to the class **NP** if a proposed solution could be verified in polynomial time. The great open question, whether **P** = **NP**, remains unsettled, and the security of the deployed systems rests, ultimately, on the *belief* that certain problems lie outside **P**. The general cryptanalytic problem was observed to be no harder than **NP**, since a key can be guessed nondeterministically and verified in polynomial time, and certain cryptographic problems were shown to be **NP**-complete. This was a powerful framework, but it purchased its guarantees on credit: every such guarantee is contingent on an unproven separation of complexity classes and on

One Character: Two Ledgers.




Figure 18. Per-character ledger view of zero-information accumulation. For each emitted synonym, the observer’s private posterior over the underlying leaf remains the uniform distribution on $\mathcal{H}$. Successive characters do not narrow this posterior because each is independently uniform conditioned on the public challenge. The ledger of LR characters therefore records $LR \log_2 H$ surface bits while contributing zero effective bits about the secret, in exact analogy with Shannon’s per-symbol mutual-information accounting for an OTP transcript.

the continued absence of a fast algorithm.

The present construction does not draw on the theory of computational complexity at all. It draws instead on the theory of communication, and specifically on Shannon’s result that the mutual information $I(X; Y)$ between a source X and an observation Y is an upper bound on the rate at which any receiver can learn about X from Y . This bound is **non-relativizing**: it holds for classical receivers, for quantum receivers, for receivers with oracle access to functions we believe to be uncomputable, and for receivers we have not yet imagined. It does not assume a polynomial bound on the adversary, or any resource bound whatever. When a channel is shaped so that $I(\varphi; \mathbf{W}) = 0$, the bound on extractable information is zero, and the impossibility of recovery is not a conjecture about complexity classes but a theorem about channels.

This reframing has a direct bearing on the quantum question. Two quantum algorithms threaten the deployed systems. Shor’s algorithm efficiently solves the integer-factorization and discrete-logarithm problems and so breaks RSA, finite-field DH, and elliptic-curve systems by attacking the algebraic group inversions on which they rest. The present construction performs no algebraic group inversion; there is no modulus to factor, no logarithm to compute, no curve point to invert. Shor’s algorithm is therefore not resisted but rendered *cate-*

gorically inapplicable: it has no structure in the construction to attach to. Grover’s algorithm provides a quadratic speed-up for search problems equipped with an oracle that distinguishes a target from a non-target. For the task of forging an accepting transcript, Grover halves the security exponent, from H^{-LR} to $H^{-LR/2}$, and this is compensated simply by doubling the round count. For the task of recovering the morphism, Grover obtains *no advantage at all*, because the oracle it would amplify is flat: by Lemma 3 every candidate morphism produces the same observable statistics, so the oracle returns the same value for all $H!$ candidates, the amplitude rotation is zero, and the quantum search collapses to classical behaviour over a space whose elements are in any case equiconsistent. Grover requires a signal to amplify; the construction denies it one.

10. Engineering Embodiment and Caveats

The reference verification pipeline specifications give an engineering embodiment of the proof lifecycle. They should be read as implementation grounding, not as automatic proof of the ideal theorem. The ideal model says that under P1-P4, balanced foliation, private morphism secrecy, sealed residue-only verification, and fresh session geometry, the observable transcript is zero-information. A deployed service must earn the approximation by enforcing contracts at every boundary.




Figure 19. Ceremony growth: classical and Grover-bounded forgery floors as functions of LR for representative H . The classical exponent grows as $LR \log_2 H$ and the Grover-bounded exponent halves it, confirming Theorem 8. The ceremony length is engineering-tunable independent of the zero-information property: increasing L or R improves online forgery resistance without changing the structural absence of secret-relevant information from the transcript.




Figure 20. Pass/fail accumulator schematic. Per-round match indicators $M_{r,t}$ are folded into a fixed-shape state $a_i = H(a_{i-1} \parallel M_i)$ inside the SVE. The final indicator $Y = \mathbf{1}\{\Lambda = 1\}$ is the only quantity that crosses the sealed boundary; constant-time comparison ensures that a mismatch at the first position produces an output indistinguishable from a mismatch at the last position.




Figure 21. Two ledgers, two roles. The public nonce ledger is read-write at the verifier and read-only to the world; it stores nonce_uuid, τ , reserve receipts, and burn status. The internal SVE ledger is sealed; it stores per-round match bits and accumulator state. Public visibility cannot reserve, consume, or inspect the internal ledger, and internal state never leaks to the public ledger. The separation is a structural prerequisite for Theorem 7 (replay non-reusability) and for the residue-only property of Definition 8.

Primitive	Security basis	Secret in public data?	Shor relevance	Grover relevance	Replay model
RSA / finite-field DH	factoring or discrete log hardness	Yes	Breaks core assumption	Mild search speedup	External protocol layer
Elliptic-curve systems	EC discrete log hardness	Yes	Breaks core assumption	Mild search speedup	External protocol layer
Lattice / hash PQC	assumed post-quantum hardness	Yes, computationally hidden	No known direct break	Parameter impact	External protocol layer
Conventional ZK/SNARK authentication	algebraic proof soundness and setup assumptions	Transcript carries proof structure	Depends on curve/setup	Parameter impact	Usually application-level nonce
Vernam OTP	information-theoretic key uniformity	No	Inapplicable	Inapplicable	Single-use key axiom
SAC/EHP/SVE ideal model	information-theoretic channel shape	No under P1-P3	No group inversion target	No morphism-recovery signal; online forgery exponent halves	Session geometry plus nonce freshness

Parameter family	Symbol	Example	Security role	Ideal condition
SAC coordinates	n	128	Path ambiguity	$n \log_2 k \geq \lambda$
Labels per coordinate	k	4	Hiding entropy	Uniform public display
Hyperplane leaves	H	Chosen configuration, e.g. 6 or 8	Per-check soundness and morphism space	Balanced leaves
Secret length	L	6 to 64	Number of position checks	Equal-length family for multi-secret
Rounds	R	4 to 10	Repeated checking	Fresh geometry per round
Balance deviation	Δ/κ	0 ideal	Leakage budget	$\Delta = 0$ exact
Nonce freshness	τ	microsecond timestamp plus UUID	Replay non-reusability	Burn before validation

Worked ENI6MA Codebook Example

The compiled-circuit reference architecture describes a compiled Rust circuit whose private configuration, entropy pool, zone mappings, hyperplane count, and bearing labels are embedded at build time, with the cloud acting as a factory rather than a vault. It also summarizes six operational variants of the same proof core: Standard, Hybrid, Sovereign, Tunnel P2P, Tunnel Proxy, and A2A Oscillation. The worked codebook example here uses the Standard circuit vocabulary only to instantiate the paper’s abstract carrier model; the variant list is an implementation taxonomy, not a theoretical restriction on H . The example private configuration contains three public alphabet rings: an uppercase/symbol ring, a lowercase/symbol ring, and a digit/symbol ring. In that file, `hyperplane_count` is set to 6, so this example uses $H = 6$ as its chosen configuration. This does not make six leaves a constant of the SAC/EHP/SVE theory. A different commitment or circuit configuration may choose a different H if its alphabet balance, usability, and carrier geometry support it. For each session, the codebook generator derives a challenge from the timestamp τ and entropy-bearing configuration, computes modulo rotation offsets for each ring, and renders the rotated alphabets as concentric rings. Each rotated ring is foliated into H zones, matching the formal

equations $A_j^{(t)} = \text{Rotate}(A_j, o_{j,t})$ and $A_j^{(t)} = Z_{j,1}^{(t)} \sqcup \dots \sqcup Z_{j,H}^{(t)}$. The generated sheet is a public carrier for the challenge geometry. It is not the private morphism ϕ , does not reveal σ , and does not prove secrecy by itself; secrecy depends on the carrier preserving the public/hidden split and on SVE verification remaining residue-only, as illustrated in Figure 22 and the corresponding hyperplane collapse in Figure 7.

The single-sheet generator command for the example carrier is:

```
[ ] python -m codebook sheet \ --private-json private.example.json \ --timestamp 1748208600123456 \ --hyperplanes 6 \ --layout landscape \ --mode gray \ --label-style greek \ --out ./output/new-dimensions-example-sheet.svg
```

This command generates visual carrier material for the paper’s worked example. It should not be read as a proof generator and should not be used to claim that the generated sheet alone establishes secrecy.

The challenge lifecycle is reserve, challenge, respond, validate. A registry asks a ledger to reserve a nonce. The ledger returns `nonce_uuid`, τ , and a reserve receipt. The circuit produces a challenge bound to τ . The user or client produces a proof payload, preferably minimal: τ , bearings, and a proof




Figure 22. ENI6MA Codebook Sheet Carrier (Public/Private Boundary). The left panel renders the public carrier sheet exactly as the observer sees it: three concentric rings (uppercase/symbols, lowercase/symbols, digits/symbols) each presented as a cyclic group of order H wedges, with per-ring offsets $\Delta_0, \Delta_1, \Delta_2$ defining the session rotation. The right panel shows the sealed boundary that the observer cannot cross: the private morphism $\phi : (\Sigma_0 \times \Sigma_1 \times \Sigma_2) \rightarrow M$ induces the rotations and the synonym mapping, but lives entirely outside the observable surface. The observer perceives only zone indices (bearings) or synonyms on the sheet – equivalence classes within zones, with no knowledge of ϕ . This is the operational embodiment of P1 (path-completeness) and P2 (uniform private synonyms): the carrier displays every legal path while concealing which one is being walked.

hash or circuit hash rather than witness internals. Validation begins by burning the nonce through the ledger. If the nonce is already spent, expired, or has a mismatched τ , validation stops with replay, expiry, or mismatch status. Only after a successful burn does the registry run `validate-proof`.

This burn-before-validate ordering matters. If validation ran first and burn happened second, a race could admit multiple attempts against the same geometry. If error messages exposed per-round failures, the SVE would become an oracle. If proof payloads carried zone numbers, witness sets, or hidden geometry, the payload would defeat the transcript-silence goal. If timing varied by first mismatch index, the sealed accumulator would leak structure. The implementation therefore has obligations that parallel the formal definitions, as illustrated in Figure 25:

The proof payload documentation distinguishes minimal and full formats. For this whitepaper’s theory, the minimal format is the cleanest expression of the observable channel: it carries public validation facts, τ , circuit hash or proof hash, and submitted bearings. A full proof format may be needed for

compatibility or audit, but every added field must be audited for whether it is public context, sealed state, or unintended witness leakage.

The observer-bit model sharpens the caveat. An observer can record a positive number of surface bits: screen geometry, bearings, timestamps, public verdicts, and ledger metadata. The ideal claim is not that the observer records nothing. It is that the recorded variables are independent of the hidden variables relevant to recovery. In the implementation, this becomes a measurable leakage budget. System tests should check uniform bearing distributions, rejection behavior under replay, fixed response shape, timing stability, and absence of hidden fields in payloads.

11. Historical Perspective

The construction can be read as a continuation of a long migration. Classical secret systems hid the method. Kerckhoffs moved security toward hidden keys with public methods. Public-key cryptography moved even more material into the public record, leaving only a trapdoor secret. The present

Formal requirement	ENI6MA-facing obligation	Failure mode if violated
SVE emits one bit	Return fixed-shape PASS/FAIL outcome	Per-round oracle leakage
Session freshness	Bind challenge and proof to nonce_uuid and τ	Replay or cross-session confusion
Burn before validate	Consume nonce before circuit execution	Race-enabled replay
Minimal transcript	Payload contains bearings and hashes, not witness internals	Hidden state serialized into public record
Constant-time checks	No mismatch-index timing or branch leakage	Side-channel correlation
Public/hidden split	Keep cypher_config, morphism, and entropy sealed	Morphism recovery by configuration exposure

work asks whether, for commitment and proof of knowledge, the public record can be made still less dangerous: not merely public but non-informative about the hidden state. This historical migration is summarized in Figure 26.

The Enigma lesson is useful because it was not only a story about key-space size. The machine’s operational weaknesses mattered, but the deeper cryptographic failure was structural betrayal. The output retained regularities that pointed back toward the hidden rotor state. Public-key artifacts do not betray through letter frequencies, but they retain a different structural relation: the public value determines the private value under an inverse relation, and security is the cost of computing that inverse. In long-horizon settings, that is still a kind of betrayal, delayed by complexity.

The SAC/EHP/SVE design removes the inverse relation from the recorded channel. The scaffold shows all paths. The synonym channel makes every emitted symbol equally likely under every hidden leaf. The SVE denies the observer an oracle. Session freshness makes the event non-reusable. The result belongs to the same family of conceptual moves as Shannon’s perfect secrecy and Diffie-Hellman’s public-key exchange: instead of optimizing an old attack surface, it changes the surface.

This does not make the construction magic. It moves the burden. Balance must be real. The morphism must stay private. The verifier must be sealed. The nonce lifecycle must be atomic. The carrier must preserve label accessibility and eigenvalue detectability while concealing the morphism. A theorem about an ideal channel is not a theorem about a compromised phone camera, a malicious verifier, or a leaked configuration file. The contribution is the clean separation: where the ideal proof ends, the engineering checklist begins.

The medium-independence of the construction has no parallel in the prior history. Cryptography has always been bound to its substrate: the rotor machine to its gears, the public-key system to its arithmetic. The present construction depends only on three projection invariants, namely that the public labels remain readable, that the value rings remain detectable, and that the private morphism remains concealed; any physical medium that can satisfy those three invariants is a valid carrier. The same protocol, with the same parameters and the same security argument, can be carried by a deck of cards on a glass table, by a wheel turning under a pointer, by a hologram,

by an acoustic signal, or by software on silicon. A commitment may thus be archived in physical form, a printed card sealed in a safe, and verified a hundred years hence by the same protocol on whatever substrate then exists, long after the computational cryptography of today has been broken. This carrier-agnostic property is illustrated in Figure 27 and the corresponding mechanical embodiment in Figure 28.

12. Conclusion

The unified New Dimensions / Rosario / ENI6MA thesis is simple: a cryptographic proof transcript need not contain a recoverable form of the secret. In the ideal SAC/EHP/SVE model, the public scaffold contains every path, the private morphism makes visible responses uniform, the SVE emits only one event bit, and session freshness prevents reuse. The observer receives surface bits but zero effective bits about the hidden state. This is the sense in which the construction is Vernam-equivalent at the channel level, while obtaining replay behavior through a different mechanism than the OTP’s single-use key axiom.

The practical value of the work depends on discipline. The formal model supplies definitions and bounds; the ENI6MA embodiment supplies a deployment vocabulary; neither excuses vague claims. A deployed system must demonstrate that the payload is minimal, the ledger burns before validation, τ is bound into reconstruction, verdicts are residue-only, timing is controlled, and hidden configuration remains hidden. If those obligations are met, the construction offers a distinct category of cryptographic confidence: not a promise that inversion will remain expensive, but a proof obligation that the channel carries no signal to invert. A summary of the unified picture is given in Figure 29.

Acknowledgments

The authors thank the cryptographic-engineering community whose scrutiny of transcript-level and side-channel attacks shaped the threat model adopted here, and acknowledge the foundational influence of C. E. Shannon, whose information-theoretic treatment of secrecy is the basis on which the present construction rests.




Figure 23. The ENI6MA paper print challenge interface. A physical/printed sheet plus transparent overlay card realizes the EHP/SVE protocol in a zero-connectivity environment: the alphabet rings are printed in the open, the per-session rotation offsets are inscribed on a transparent overlay that the prover lays over the sheet, and the prover’s response is the synonym written for each round. The same security parameters as in the digital embodiment apply, demonstrating that the protocol is medium-independent and survives in environments without electronic infrastructure.

Appendix A. Notation

Appendix B. Deferred Proof Details

B.1 Uniform synonym calculation. Fix a leaf h_i and synonym v_k . The number of bijections $\varphi : \mathcal{H} \rightarrow \Sigma$ satisfying $\varphi(h_i) = v_k$ is $(H - 1)!$. The total number of bijections is $H!$. Therefore

$$\Pr[\varphi(h_i) = v_k] = \frac{(H - 1)!}{H!} = \frac{1}{H}.$$

Equation Narrative: This equation calculates the uniform mapping probability of the morphism φ mapping a leaf label h_i to a synonym v_k , which simplifies to $(H - 1)!/H! = 1/H$ under uniform selection over all $H!$ bijections.

Because balanced foliation removes leaf-size bias, conditioning on the public challenge does not alter the marginal.

B.2 Zero mutual information. If $\mathbf{W}$ consists of LR independent uniform synonyms conditioned on C , then $H(\mathbf{W} | C) = LR \log_2 H$. If conditioning additionally on φ does not change the distribution seen by the observer because the true leaf remains hidden and balanced, then $H(\mathbf{W} | \varphi, C) = LR \log_2 H$. Their difference is zero. Practical imbalance replaces the equality with a leakage term.

B.3 Replay non-reusability. The session relation verified

by SVE is indexed by τ . A transcript valid for relation R_τ is not generally a member of $R_{\tau'}$. In the ideal type model $R_\tau \cap R_{\tau'} = \emptyset$ for $\tau \neq \tau'$. In a deployed probabilistic model, the intersection probability is bounded by the entropy and collision properties of the geometry derivation plus the atomicity of nonce burn.

B.4 Reverse accumulator. In strict mode the acceptance predicate is $\max_j \text{Score}(K_j) = LR$. Without any K_j , the adversary guesses a correct leaf for each check with probability H^{-LR} for each candidate. A union bound gives mH^{-LR} . The sealed index J^* prevents the acceptance event from becoming a side channel over which candidate was closest.

Appendix C. Figure Program

The visual generation pass should create line-art or infographic figures from the prompt blocks embedded in the body:

1. The Relocation Thesis.
2. SAC-EHP-SVE Layering.
3. Four Pillars Dependency Graph.
4. Theorem Dependency Graph.
5. Security Parameter Surface.
6. ENI6MA Codebook Sheet Carrier.
7. ENI6MA Reserve-Challenge-Burn-Validate Lifecycle.

No generated image is included in this source. Each figure prompt is intentionally explicit about title, purpose, style,




Figure 24. Multi-verifier quorum and threshold accumulator. A k -of- n threshold verifier topology: each independent SVE emits a binary residue, and a binomial threshold gate accepts if and only if at least k residues are unity. The construction preserves the residue-only property at the verifier-quorum level: the public output is still a single bit, the per-verifier residues are sealed inside their respective SVEs, and the threshold parameter k is public context. Multi-verifier quorum is therefore an architectural lever that does not weaken the zero-information property of any individual verifier.

layout, labels, and instructions so that later generation can be audited against the manuscript’s claims.

Appendix D. Implementation Constraints Checklist

Appendix E. Parameter Examples and Audit Scenarios

The parameter tables in the body state the formulas; this appendix gives three audit scenarios that make those formulas operational. In each scenario, H is selected at commitment or configuration time and then audited against security level, usability, alphabet balance, and carrier constraints.

Scenario 1: Demonstration carrier. A small classroom or tabletop demonstration may use $H = 4$, $L = 8$, and $R = 2$. The classical online forgery exponent is $LR \log_2 H = 32$ bits, which is enough to show protocol mechanics but not enough for production authentication. The audit question in this setting is pedagogical: do observers see all paths, do all leaves have equal size, and does the response alphabet remain disjoint from the public labels? A demonstration may deliberately expose the morphism after the exercise to teach the construction, but that exposure should be labeled as didactic and outside the security claim.

Scenario 2: Human-scale strong authentication. A practical low-friction ceremony might choose $H = 6$, $L = 7$, and $R = 10$ at configuration time, yielding approximately 181 classical bits and 90 Grover-bounded bits against online forgery.

This particular $H = 6$ configuration can fit human directional or color vocabularies while preserving a nontrivial $6!$ morphism space, but another carrier could choose a different H . The audit focus shifts from arithmetic to ergonomics and leakage: can the user reliably detect witness membership, do failed attempts take the same shape and time as successful attempts, and does every replay return a ledger-level 409 or 410 before circuit validation begins?

Scenario 3: Machine or agent authorization. A machine-to-machine or high-assurance agent workflow can raise L and R because the prover is not constrained by human scanning speed. For example, $H = 8$, $L = 64$, and $R = 4$ gives $LR \log_2 H = 768$ classical bits before online-rate limits and policy controls. In such a setting the main risk is not the combinatorial bound but boundary discipline: whether the agent’s sealed configuration is protected, whether policy context is bound to the nonce, whether logs redact hidden substrate, and whether proof-of-record endpoints remain read-only.

Across all scenarios, the same sanity checks apply. First, hiding and soundness are distinct. Increasing H , L , or R improves online forgery resistance but does not repair a public scaffold that leaks σ . Second, balance is not optional. A beautiful UI with unequal chamber populations is not a zero-information channel; it is a biased channel with a leakage budget. Third, replay defense is not a final if statement. The




Figure 25. ENI6MA Reserve-Challenge-Burn-Validate Lifecycle: Explain the deployed nonce lifecycle and where replay is stopped, highlighting that burn occurs before validation.

nonce must be reserved before challenge generation, bound into the proof, consumed before validation, and made visible only through read-only status after the fact.

The safest way to review an implementation is therefore to ask four questions in order. Does every possible secret path remain present in the public projection? Are synonym marginals uniform under the public context? Does the SVE reveal only a nonce-scoped event bit? Does the nonce lifecycle make captured transcripts unusable after the event? If any answer is no, the system may still be useful as an authentication mechanism, but it should not be described as satisfying the ideal SAC/EHP/SVE zero-information theorem.

Author Biographies

Frank Dylan Rosario is the inventor of the structured-ambiguity commitment, the entropic hyperplane protocol, and the sealed-verifier proof of knowledge described in this paper, and is the author of the underlying United States patent specification. His work concerns information-theoretic foundations for cryptographic confidentiality and carrier-independent verification.

Dr. Lin (Grant) Wang contributed to the information-theoretic analysis and the comparative framing of the construction against deployed public-key, zero-knowledge, and post-quantum primitives.

Suggested Citation

F.D. Rosario and L. Wang, “New Dimensions in Zero-Information Cryptography: Structured Ambiguity Commitments, Entropic Hyperplane Proofs, and the ENI6MA Embodiment,” June 2026.

References

1. C. E. Shannon, “Communication Theory of Secrecy Systems,” *Bell System Technical Journal*, 1949.
2. C. E. Shannon, “A Mathematical Theory of Communication,” *Bell System Technical Journal*, 1948.
3. W. Diffie and M. E. Hellman, “New Directions in Cryptography,” *IEEE Transactions on Information Theory*, 1976.
4. R. L. Rivest, A. Shamir, and L. M. Adleman, “A Method for Obtaining Digital Signatures and Public-Key Cryptosystems,” *Communications of the ACM*, 1978.
5. S. Goldwasser, S. Micali, and C. Rackoff, “The Knowledge Complexity of Interactive Proof Systems,” *SIAM Journal on Computing*, 1989.
6. P. W. Shor, “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer,” *SIAM Journal on Computing*, 1997.
7. L. K. Grover, “A Fast Quantum Mechanical Algorithm for Database Search,” *Proceedings of STOC*, 1996.
8. F. D. Rosario, *Systems and Methods for Structured Ambiguity Commitments and Entropic Hyperplane Proof-of-Knowledge Verification*, patent specification materials, 2026.

Cryptography: Unconditional vs. Conditional Security




Figure 26. Cryptographic migration from hidden methods to hidden channels. Caesar/classical: hide the method. Kerckhoffs/rotor (Enigma): publish the method, hide the key. Diffie-Hellman / RSA / EC: publish the public value, hide the trapdoor. Lattice / hash PQC: substitute a post-quantum hardness assumption. Shannon (one-time pad): use information-theoretic key uniformity. SAC/EHP/SVE (this work): use information-theoretic channel shaping; the secret is absent from the recorded transcript by construction. Each step decreases the portion of the system that has to be protected from public view; the present construction reaches the limit at which the secret is not in the public data at all.

9. F. D. Rosario and L. Wang, *Zero-Information Proof-of-Knowledge: Structured Ambiguity Commitments and Entropic Hyperplane Protocols*, May 2026 manuscript.
10. *ENI6MA Reference Verification Architecture and Nonce Ledger Specification*, Technical Report ER-2026-04, ENI6MA Security Group, 2026.




Figure 27. Carrier-agnostic projection. The same mathematical manifold $\mathcal{M}$ projects onto physical carriers (printed cards, mechanical wheels, holographic grids, acoustic signals, or digital silicon) while preserving the three invariants: label accessibility, eigenvalue detectability, and morphism concealment. Security is therefore a property of the projection, not of the medium.




Figure 28. Mechanical wheel/rotor verifier. A purely mechanical embodiment of the EHP/SVE protocol: concentric rotating wheels carry the public alphabet rings, a pointer indexes the private leaf the prover selected, and a flywheel-driven entropy source supplies session randomness. The same security parameters that govern the digital embodiment apply unchanged, demonstrating the carrier-agnostic property in concrete physical form.




Figure 29. Unified summary of the New Dimensions construction. SAC supplies path ambiguity (k^n equiconsistent paths under the public scaffold), EHP supplies session foliation and synonym channelization (uniform marginals over $H!$ morphisms), and SVE supplies sealed residue-only verification (one event bit). Together they realize the four pillars P1–P4 and yield the channel-level Vernam equivalence $I(\text{Hidden}; O | C) = 0$. The engineering surface (nonce ledger, burn-before-validate, constant-time comparison, payload minimality) earns the approximation in deployment.

Symbol	Meaning
$\mathcal{F}$	Public structured ambiguity scaffold
I	Coordinate set $\{1, \dots, n\}$
L	Public label set of size k
σ	Private selector path $I \rightarrow L$
s_σ	Induced section selected by σ
D	Binding tag
$\mathcal{P}(\mathcal{F})$	Ambiguity class of scaffold sections
$\mathcal{H}$	Public hyperplane-label set
Σ	Disjoint synonym vocabulary
φ	Private bijective morphism $\mathcal{H} \rightarrow \Sigma$
H	Number of hyperplane leaves
L, R	Secret length / checks per round and number of rounds; context disambiguates L from label set
E, τ, ρ, Θ	Session entropy, freshness coordinate, salt, and structural parameters
W_i^t	Witness set in leaf i at round t
$\mathbf{W}$	Submitted witness synonym array
O	Observable transcript
C	Public context or challenge
Y	Final event-validity bit
$\mathcal{K}$	Secret family $(K_1, \dots, K_m)$
J	Active secret index
RWA	Reverse witness accumulator

Constraint	Required audit evidence
Path-complete scaffold	Enumerate legal paths or prove construction covers them
Balanced foliation	Show exact counts or bounded Δ/κ
Private morphism secrecy	Confirm no payload, log, UI artifact, or config leak exposes φ
Residue-only SVE	Verify no per-round diagnostics leave the sealed boundary
Constant-time verification	Measure timing across first-failure, last-failure, and pass cases
Nonce reserve/burn	Demonstrate 409/410 behavior on replay and expiry
τ binding	Show mismatch rejection before circuit validation
Payload minimality	Confirm public payload contains no witness internals
Ledger read/write separation	Confirm public visibility cannot reserve or consume nonces
Documentation caveat	Label ideal SAC/EHP/SVE claims separately from the reference verification pipeline deployment notes