The Enigma Cypher as a Cognitive–Cryptographic Primitive:

Gestalt Recognition, Holographic Projections, and Post-Quantum Combinatorics for Presence-as-Work

by Frank Dylan Rosario and Dr. Lin Wang

Abstract

We formalize the cybernetic foundations of the **Enigma cypher** as a cognitive–cryptographic primitive in which a human or agentic twin proves knowledge without disclosure. The core insight is that **gestalt perception**—our native ability to perceive *wholes*—can be harnessed as a cryptographic advantage when symbols from known, **enumerated alphabets** are embedded in a **stochastic manifold projection** with per-round entropy (a one-time, entropy-dependent projection, or OTP). The prover (Alice) owns a **private bijection ("private map")** from six semantic synonyms to six perceptual "leaves" on a projection canvas; each round she emits only a **masked witness**—a synonym, not a symbol—so the verifier (Bob) can check membership while learning nothing about the secret sequence. The design deliberately implements Shannon's **perfect secrecy** intuition at the transcript layer (no reusable residue, no mutual information), while cryptographically enshrining **diffusion** (many letters per holographic leaf/zone) and **confusion** (private map permutations across six hyperplanes, yielding $6! = 720$ bijective variants). The manifold itself admits **ring-rotation diffusion** of multiple alphabets across $30 \times 30 \times 12 = 10{,}800$ micro-states per ring, compounding the **combinatorial explosion** that submerges the witness from an adversary's view. We analyze false-accept probabilities, attacker hypothesis dynamics, and post-quantum hardness. We show why human/agent recognition collapses a high-dimensional search to a near-constant-time perceptual act, while adversaries—classical or quantum—cannot exploit structure and face search spaces that remain astronomically large even under quadratic speedups. Finally, we connect the primitive to **cybernetic presence-as-work**, deepfake-resilient identity, and agentic security, arguing that the Enigma cypher provides a **future-proof** authentication substrate that eliminates reliance on biometrics and external MFA tokens.

Cognitive/agent search for hyperplane membership (Enigma cypher)

In the **Enigma cypher**, "hyperplanes" are the six disjoint perceptual leaves that partition the capsule's projection. The search problem is: given the next secret character s_i and the capsule's one-time layout f_t, decide **which hyperplane currently contains** s_i and emit the private synonym for that region. Formally, let $H = \{H_1, \ldots, H_6\}$ denote the leaves.

$$M_i(H_j) = \mathbf{1}\{\, s_i \in f_t^{-1}(H_j) \,\}$$

TTS: "M sub i of H j equals one if s sub i is mapped by f sub t into hyperplane H j, otherwise zero."

Stage 1 — Category priming (coarse filter).
Humans (and agent twins) first narrow attention by **category**—e.g., lowercase vs uppercase vs digit/punct—so only relevant glyphs compete for attention. Write C for the active category and g_C for the filter that hides non-members. This reduces visual clutter and prevents timing leakage while preserving balance across leaves.

Stage 2 — Gist and salience (preattentive pass).
The visual system extracts the **alphabetic gist**—ordered arcs like $A \to Z$—and computes a fast salience score per leaf,

$$S(H_j \mid C) = \text{pop-out strength of category-}C \text{ tokens in } H_j,$$

then homes in on 1–2 candidate leaves with highest S.

TTS: "S of H j given C is the pop-out strength for category C in leaf H j."

Stage 3 — Template match (fine pass).
Within the top candidate leaf(s), the prover performs a quick template check: "Is the target s_i on this arc?" This is a **meaning-first** scan (order awareness), not letter-by-letter enumeration, so cognitive cost stays low. In enclave form, the agent uses the capsule seed to compute the same membership predicate deterministically.

Stage 4 — Decision and synonyming.
Select the winning leaf

$$\ell^\star = \arg\max_{H_j \in H} \big\{\, S(H_j \mid C) \cdot M_i(H_j) \,\big\},$$

then apply the **private bijection** $\pi : H \to \mathcal{S}$ to produce the masked witness $\sigma_i = \pi(\ell^\star)$.

TTS: "ell star equals the arg max over H j of salience times the membership bit; sigma sub i equals pi of ell star."

Stage 5 — Emission under capsule timing.
The prover outputs only σ_i within a quantized window. The verifier checks membership for the **current** capsule; because layouts and masks change per round, synonyms have no stable public meaning.

Why this is robust.
For the prover, the pipeline—**filter** → **gist** → **confirm** → **synonym**—is $O(1)$ perceptually: one or two glances. For an attacker lacking π and the live f_t, the same task explodes into a combinatorial search across hyperplanes, ring rotations, and private-map permutations. The human/agent exploits **gestalt** to collapse a high-dimensional scene to a binary decision, while the adversary faces unstructured hypotheses that cannot be pruned without oracle access.

1. Introduction: From Possession to Perception

For half a century, mainstream authentication has treated identity as a **thing you have** or a **thing you are**. The first bucket—**possessed artifacts**—includes passwords, one-time passwords (OTPs), recovery codes, smart cards, and private keys. The second bucket—**inherent traits**—centers on biometrics such as face, voice, fingerprint, and gait. Both buckets share a structural weakness: they **export** either a secret or a template, leaving residue a determined adversary can capture, replay, correlate, or coerce. Passwords traverse keyboards and networks; OTPs traverse messaging channels and screens; private keys reside in devices that must be provisioned, backed up, and occasionally repaired by someone else; biometric templates are stored, matched, and sometimes leaked. Even modern "phishing-resistant" passkeys ultimately certify **who controlled the I/O** of a particular device at a particular instant, not whether the human or agent behind that device actually **knows** what they claim to know.

The **Enigma cypher** proposes a different organizing principle. It reframes identity as **present-tense capability**: not a static artifact or an immutable template, but the ability to **recognize** and **respond** to a **time-keyed projection** using either private human cognition or a sealed, attested morphism running in an enclave. In a typical ceremony, the prover perceives a projection of enumerated alphabets arranged on rotating rings and partitioned into six disjoint regions ("leaves"). The prover alone holds a **private bijection**—their **private map**—from those six leaves to six private synonyms. In each round, they emit just a **masked witness**: a synonym, never the symbol it points to. The verifier checks **membership** for the current capsule (the session's time-bound, entropy-seeded instance) and moves on. No secret is exported; no reusable artifact survives; every session is freshly keyed and visually re-drawn.

This shift from possession to perception is more than aesthetic—it is **structural**. The crux is **cognitive asymmetry**. Humans excel at **gestalt**: the ability to perceive *wholes*—ordered runs of letters, spatial groupings, familiar symmetries—without scanning element by element. When the projection's geometry is **meaningless without the private map yet instantly legible with it**, the honest prover experiences a rapid "pop-out," while the observer perceives only clutter. By design, the Enigma cypher cultivates this asymmetry and ties it to a one-time, capsule-keyed schedule of placements and masks. A round can be guessed only with probability $p_{\text{guess}} = 1/6$; across h strict rounds, a blind attacker's success chance shrinks as $(1/6)^h$.

"p sub guess equals one over six; across h rounds the chance is one over six to the h."

On the privacy axis, the protocol embraces Shannon's ethos: a transcript must **not leak** the secret. Let S denote the secret array and T the public transcript (commitments, salted index digests, accept/reject bits). The goal is $I(S;T) = 0$,
"I of S semicolon T equals zero."
so that an eavesdropper, even after recording billions of ceremonies, cannot extract **mutual information** about S. The transcript carries no symbols, only minimal, capsule-scoped facts that the verifier needs to decide. Because the capsule re-keys placements and masks per round and reseeds per session, **no stable residue** accumulates across time. The same synonym a user speaks today may target a different leaf tomorrow; the same leaf today contains a different mix of characters tomorrow. Frequency and correlation have nothing stationary to grip.

The engineering surface follows from these principles. Each ceremony lives inside a **capsule**, parameterized by a time tag T_t, fresh entropy seed_t, and auxiliary policy flags. The projection map f_t is capsule-keyed, yielding a balanced layout: each leaf receives roughly equal symbol mass so that guessing stubbornly remains $1/6$. The user's private bijection $\pi : \Omega! \to !\mathcal{S}$ (from leaves Ω to private synonyms $\mathcal{S}$) never leaves the mind or sealed agent. A one-time per-round mask ρ_i (derived from seed_t) ensures the **visible synonym** is a masked witness that **cancels** under the verifier's schedule—intuitively, a geometric XOR that blocks oracle construction. The only exported data are fixed-shape, salted digests of indices and the boolean outcomes. By construction there is no deterministic function f the adversary can query repeatedly outside the live capsule epoch, denying both classical black-box probing and quantum superposition queries.

Threat-modeling within this frame looks different from legacy protocols. **Eavesdropping or recording** yields uncorrelatable, one-time views. **Relays and man-in-the-middle** puppets cannot "forward" authority because there is no transferable token—only a live act requiring the conjunction of π **and** the current capsule state. **Deepfakes and biometric spoofs** fail because appearance is not an input to acceptance; only **correct membership** matters. **Key-loggers and screen scrapers** starve for structure: the sole artifact is a synonym with user-specific semantics, and an **attested, minimalist viewer** renders the canvas while exporting only salted indices/digests (no frames, no cursor telemetry). **Oracle probing** collapses because the projection redraws and masks change; there is no stable input–output map to interrogate.

The Enigma cypher is deliberately **carrier-agnostic**. The perceptual task can be presented visually, auditorily, or haptically, with the same verification math. This improves accessibility and blunts modality-specific side channels. Usability remains front-of-mind: category priming (e.g., the prover mentally filters for "lowercase" before searching) reduces perceptual load; margins and large targets reduce motor slips; timing is quantized into bins that are generous for humans yet coarsened for leakage. A threshold acceptance rule (e.g., accept if $\sum_i M_i \geq \theta$ for $\theta < h$) tolerates occasional misses while keeping the false-

accept tail negligible via standard binomial bounds.
"Accept if the sum of M sub i is at least theta."

Crucially, Enigma integrates **governance** into the cryptographic fabric. Acceptance is bound to a **policy digest**, producing audit statements like: "Operation X executed because $\Lambda = 1$ under policy hash $H(\text{policy})$ at time T_t." This decouples UI narratives from evidence—no pop-ups or screenshots are required to prove consent—and supports **privacy-by-design**: the ledger is thin, self-authenticating, and free of PII. Revocation is first-class: rotate π, reseed the capsule, and the ceremony's meaning changes instantly without resetting identities or replacing hardware. Because secrets never leave the mind or sealed agent, **data at rest** become a minimal, tamper-evident record rather than a honeypot.

From a systems perspective, Enigma **complements** rather than rejects hardware roots. An attested viewer can live in a TEE, SEV, or measured container. Organizations may still maintain passkeys for transport security while moving the **authorization primitive** "up" to **presence-as-work**: no session keys are derived until the human/agent succeeds **now**. After acceptance, both sides may derive an ephemeral key $K_t = \mathrm{KDF}(EC_t, T)$ that expires with policy; if acceptance fails, **no key exists** to steal or relay.
"K sub t equals K D F of the capsule and the transcript."

Finally, the paper situates the primitive within a broader **post-quantum** worldview. We do not seek a new algebraic assumption; instead, we remove the structured oracles that quantum algorithms rely upon. Even if one modeled the attacker's hypothesis space as $N = n \cdot 6! \cdot 72^{20}$ (with n possible secrets), Grover's quadratic speedup would still leave $\sqrt{N}$ beyond practical reach. More importantly, Enigma denies the coherent oracle **in principle**: outside the capsule there is no repeatable function to query; inside the capsule the adversary lacks control of the inputs and visibility of the unmasked state. Security is thus anchored in **one-time unpredictability**, **minimal outputs**, and **cognitive asymmetry**, not in promises of multi-decade algebraic hardness.

Roadmap. The remainder of the work proceeds in six movements. **(i) Cognitive foundations** formalize the gestalt mechanisms—preattentive pop-out, perceptual grouping, and category priming—that make the honest prover's task rapid and low-load. **(ii) The formal model** specifies capsules, projections, leaves, private maps, masks, and acceptance, including balance conditions that keep guessing at 1/6. **(iii) Security analysis** develops completeness, soundness, and **perfect-secrecy** style leakage claims, with explicit tails for strict and threshold acceptance. **(iv) Combinatorics** quantifies the hypothesis space under multiple alphabets, rings, and hyperplane maps, and extends to the case of n encoded secrets where the active secret is unknown. **(v) Post-quantum considerations** explain why removing oracles—and not merely enlarging exponents—yields durable resistance to both classical and quantum adversaries. **(vi) Deployment** addresses cybernetic presence, attested viewers, policy binding, duress mappings, entropy health, and human-factors tuning. The result is a practical, privacy-forward, agent-compatible primitive that authenticates **without exporting identity** and authorizes **without accumulating residue**—a

shift from possession to perception that meets the realities of deepfakes, surveillance, and an oncoming quantum era.

2. Cognitive Foundations: Gestalt as a Cryptographic Lever

2.1 Perceptual Grouping and Preattentive Pop-Out

The visual system performs **parallel, preattentive grouping**: edges, clusters, and familiar orders (e.g., Latin $A! \to !Z$) "pop out" even in clutter. The Enigma cypher uses **ordered, enumerated alphabets** and arranges their tokens along **rotating rings** and **six leaves** (e.g., cardinal/diagonal bearings). For a user who *knows* the alphabet and holds a **private synonym map** (e.g., "up, down, left, right, forward, back"), the correct leaf becomes immediately salient: the mind searches *by meaning*, not by pixels.

2.2 Category Priming and Alphabet Filtering

Before each micro-search, the prover primes on **category**: upper vs lower case, digit vs letter, punctuation vs glyph. The projection intentionally **mixes categories** across leaves, but category priming shrinks the prover's perceptual search to a relevant sub-cloud, **reducing cognitive load** without emitting any side-channel.

2.3 Low-Load, High-Assurance Interaction

The prover's loop is **see** $\to$ **apply private map** $\to$ **emit synonym**. There is no arithmetic, no string entry, no motor signature to replay. The **witness** is purely semantic (a synonym), and the **map** that gives it meaning remains private. The verifier's loop is **check membership** for that round; the transcript stores only a **yes/no** result and a salted digest of indices, yielding no stable residue.

3. Formal Elements of the Enigma Cypher

We model the ceremony at round index $i \in 1, \ldots, h$ within a **capsule** that carries time and entropy.

3.1 Capsule, Projection, Leaves

- **Capsule:** $EC_t = (T_t, \text{seed}_t, \text{aux}_t)$ binds the ceremony to time T_t, fresh entropy seed_t, and policy/UX flags aux_t.
- **Projection:** A capsule-keyed map f_t projects high-dimensional symbol embeddings onto a 2D canvas with **six disjoint leaves** $\Omega = \uparrow, \rightarrow, \downarrow, \leftarrow, \nearrow, \searrow$ (notation illustrative).

- **Balance:** Placements are balanced so that **random guessing** succeeds with probability 1/6 per round.

Equation (per-round random-guess pass):

$$p_{\text{guess}} = \frac{1}{6}.$$

"The probability of passing a single round by guessing is one divided by six."

3.2 Private Map and Mask

- **Private map (bijection):** $\pi : \Omega! \to !\mathcal{S}$ where $\mathcal{S} =$ `up,down,left,right,forward,back` or any six private synonyms. π is **secret** and session-local across rounds.
- **Masking (XOR intuition):** The verifier samples a one-time mask ρ_i from the capsule; Alice's visible synonym acts as a **masked** statement that **cancels** under the verifier's schedule. Geometrically, this "XOR" ensures **non-members** cannot correlate observed outputs to unmasked witness bits.

3.3 Acceptance Rule

Let $M_i \in 0, 1$ be the result of the i-th membership check. Under **strict** acceptance:

$$\Lambda = \bigwedge_{i=1}^{h} M_i.$$

"Lambda equals the logical AND of the round results from one to h."

Under **threshold** acceptance (human error tolerance), accept if $\sum_{i=1}^{h} M_i \geq \theta$ for a fixed θ.

4. Security: Soundness, Completeness, and Perfect-Secrecy Style Leakage

4.1 Completeness

An honest prover who knows the alphabet, possesses the private map π, and recognizes the next secret character in each round will pass with high probability. Practical designs choose h and margins so honest error rates remain low (e.g., < 1

4.2 Soundness and False-Accept Bounds

Under strict acceptance, a random guesser's **false-accept** probability is:

$$\Pr[\mathrm{FA}_{\mathrm{strict}}(h)] = \left(\frac{1}{6}\right)^h.$$

"The probability of a false accept with strict acceptance after h rounds is one over six to the h."

Concrete tails:

- $h = 8: \quad (1/6)^8 \approx 5.9537 \times 10^{-7}$,
  "About five point nine five times ten to the minus seven."
- $h = 10: \quad (1/6)^{10} \approx 1.6538 \times 10^{-8}$,
  "About one point six five times ten to the minus eight."
- $h = 12: \quad (1/6)^{12} \approx 4.5939 \times 10^{-10}$.
  "About four point five nine times ten to the minus ten."

Under threshold acceptance with parameter θ, the false-accept rate for a non-member is a **binomial tail**:

$$\Pr[\mathrm{FA}_{\mathrm{thr}}(h, \theta)] = \sum_{k=\theta}^{h} \binom{h}{k} \left(\tfrac{1}{6}\right)^k \left(1 - \tfrac{1}{6}\right)^{h-k},$$

"The threshold false-accept probability equals the sum from k equals theta to h of 'h choose k' times one-sixth to the k times five-sixths to the h minus k." and can be upper-bounded via Chernoff bounds for design tuning.

4.3 Perfect-Secrecy Style Leakage

Each round exports only a **private synonym**, not the symbol. Capsules **re-randomize** placements and masks per round and per session. Let S be the secret sequence and T the public transcript (commitment, masked indices digest, accept/reject). The design objective is **information-theoretic non-leakage** at the transcript interface:

$$I(S; T) = 0.$$

"The mutual information between the secret and the transcript equals zero."

Intuitively, because each session's projection is one-time (OTP-like) and because the private map is never revealed, an eavesdropper's view is distributionally **independent of the secret**. This is the **Shannon** ethos: even unlimited passive observation yields **no reusable signal**.

4.4 Diffusion and Confusion

- **Diffusion:** many letters per holographic leaf/zone; a single leaf never "means" a single symbol. This **destroys** frequency fingerprints.
- **Confusion:** the **private map** π permutes meanings across **six hyperplanes**, yielding $6! = 720$ equally likely bijections. No public feature binds synonyms to leaves.

$$|\{\text{bijections}\}| = 6! = 720.$$

"The number of bijections from six leaves to six synonyms equals six factorial, which is seven hundred twenty."

Together with per-round masks, diffusion and confusion ensure transcripts look like structureless outputs to outsiders.

5. Combinatorics of the Search Space (20-Character Secret, n Encoded Secrets)

5.1 Baseline: Unknown Secret and Unknown Map (active secret among n)

Assume the attacker knows neither Alice's private map nor which **one** of her n encoded secrets is active in the session. Each secret is a length-20 string over a 72-symbol alphabet. The count of length-20 strings is

$$|\mathcal{P}| \;=\; 72^{20} \;=\; 14{,}016{,}833{,}953{,}562{,}607{,}293{,}918{,}185{,}758{,}734{,}155{,}776 \;\approx\; 1.40168{\times}10^{37}.$$

"Seventy-two to the twentieth is about one point four zero one six eight times ten to the thirty-seven."

Including the 6! private-map permutations and the ambiguity over **which** of the n secrets is in use, the initial (active-secret, bijection) hypothesis space is

$$|\mathcal{H}_0| \;=\; n{\cdot}6!{\cdot}72^{20} \;=\; n{\cdot}10{,}092{,}120{,}446{,}565{,}077{,}251{,}621{,}093{,}746{,}288{,}592{,}158{,}720 \;\approx\; n{\cdot}1.00921{\times}10^{40}.$$

"H naught equals n times six factorial times seventy-two to the twentieth, approximately n times one point zero zero nine two one times ten to the fortieth."

Under the crude i.i.d. pruning model where each round eliminates about 5/6 of the wrong hypotheses, the rounds h required satisfy $6^h \approx |\mathcal{H} * 0|$. Since $6^{52} \approx 2.90981 \times 10^{40}$ covers the single-secret case, the multi-secret case adds roughly $\log *6n$ more rounds:

$$h \;\approx\; 52 \;+\; \lceil \log_6 n \rceil.$$

"h is approximately fifty-two plus the ceiling of log base six of n."

Operational note. Real systems prevent such perfect accumulation by limiting visible rounds per capsule, enforcing short acceptance windows, and **re-seeding between sessions**, which erases carry-over constraints before the attacker can realize $\log_6 n$ additional pruning.

5.2 Shortlist Variant (attacker knows the exact n candidate secrets)

If the attacker improbably knows the **exact content** of the n candidate secrets (all length-20), the initial space collapses to

$$|\mathcal{H}_0^{(\text{short})}| \;=\; n \cdot 6! \;=\; 720\, n,$$

"H naught short equals seven hundred twenty times n."
and the rounds required in one uninterrupted capsule are

$$h \;\approx\; \lceil \log_6(720\, n) \rceil \;=\; \lceil \log_6 720 \rceil \;+\; \lceil \log_6 n \rceil \;\approx\; 4 \;+\; \lceil \log_6 n \rceil\,,$$

"About four plus the ceiling of log base six of n."
still blocked in practice by capsule rotation, attested viewing, and transcript-only logging.

5.3 Ring-Rotation Diffusion: $30 \times 30 \times 12$

Ring-rotation diffusion is unchanged but compounds the ambiguity multiplicatively. One ring yields

$$30 \times 30 \times 12 \;=\; 10{,}800$$

"Thirty times thirty times twelve equals ten thousand eight hundred."
micro-configurations; with R rings,

$$N_R \;=\; (10{,}800)^R$$

"N sub R equals ten thousand eight hundred to the R."
must be contemplated **per round**, on top of the $n \cdot 6!$ and symbol choices.

5.4 Design Summary

- **Legitimate prover:** cognitive effort $\sim O(h)$, unchanged; the user simply knows which secret they are proving.
- **Adversary:** hypothesis space scales **linearly in** n ($n{\cdot}6!{\cdot}72^{20}$), so required co-epoch observations increase by $\approx \log_6 n$ rounds; **capsule rotation** resets progress before those rounds accrue.

6. Why Traditional Cryptanalysis Fails (with n Encoded Secrets)

- **Frequency analysis:** demands a stationary map and stable symbol–leaf skew. Here, **diffusion** and balancing keep per-leaf mass flat; adding n candidate secrets multiplies ambiguity about *which* sequence is active, further diluting any frequency signal.
- **Correlation attacks:** need cross-round or cross-session alignment. **Per-round shuffling** and **per-session re-seeding** already break alignment; uncertainty over the **active secret among** n adds another layer—observations that might fit one secret under one map equally explain others under different maps.
- **Dictionary/known-plaintext oracles:** require symbol-level feedback. The Enigma cypher emits **only private synonyms**, never symbols; even if an attacker carried multiple hypothesized secrets, there is no oracle to select which of the n is being proven within a fresh capsule.
- **Side-channels (timing, kinematics):** are constrained by **low cognitive load**, **quantized response windows**, and **attested viewers** that export only discrete, salted digests—insufficient to disambiguate among n plausible secrets.

Bottom line: introducing n encoded secrets multiplies the attacker's initial hypotheses by n and increases the co-epoch round requirement by $\approx \log_6 n$, while the observable surface remains minimalist, balanced, and continuously refreshed—leaving frequency and correlation with **nothing stationary to grip**.

7. Post-Quantum Considerations (20-Character Secret, n Encoded Secrets)

Quantum search (e.g., **Grover's algorithm**) offers at most a **quadratic** speedup over unstructured search. If we model the adversary's work as scanning a hypothesis space of size N, Grover reduces $O(N)$ to $O(\sqrt{N})$. For the baseline in this setting,

$$N \;=\; |\mathcal{H}_0| \;=\; n \cdot 6! \cdot 72^{20} \;\approx\; n \cdot 1.00921 \times 10^{40}.$$

"N equals n times six factorial times seventy-two to the twentieth, approximately n times one point zero zero nine two one times ten to the fortieth."

The corresponding Grover lower bound is

$$\sqrt{N} \;=\; \sqrt{n}\,\sqrt{6!}\,72^{10} \;\approx\; \sqrt{n}{\times}26.833{\times}3.743906242624487424{\times}10^{18} \;\approx\; \sqrt{n}{\times}1.005{\times}10^{20}.$$

"Square root of N is approximately square-root of n times twenty-six point eight three three times three point seven four three nine zero six two four two six two four four eight seven four two four times ten to the eighteen, which is about square-root of n times one point zero zero five times ten to the twenty."

Even **with** quadratic speedup, $\widetilde{O}(10^{20}\sqrt{n})$ coherent queries are infeasible for any realizable quantum architecture. Moreover, the Enigma cypher withholds the very **oracle** Grover would need: the capsule's **leaf schedule and masks are one-time**, and **membership** cannot be evaluated **without** the private map π and the live per-round schedule. There is no black-box function f that an adversary can query in superposition; only a human/agent's **masked** synonym is observable, per round and per capsule. The structure required to realize amplitude amplification is **absent by design**.

Further, **combinatorial compounding** from ring rotations and 6! permutations pushes realistic N far beyond $n \cdot 10^{40}$ for inversion attempts that must account for layout degrees of freedom. With, for example, $R = 6$ independently parameterized rings:

$$N \gtrsim n{\cdot}6!{\cdot}72^{20}{\cdot}(10{,}800)^6 \approx n{\cdot}(1.00921{\times}10^{40}){\cdot}(1.586874322944{\times}10^{24}) \approx n{\cdot}1.60{\times}10^{64}.$$

"N is at least n times one point zero zero nine two one times ten to the forty multiplied by one point five eight six eight seven four three two two nine four four times ten to the twenty-four, which is about n times one point six zero times ten to the sixty-four."

Grover on this compounded space still leaves $\sqrt{N} \approx \sqrt{n}{\times}1.26{\times}10^{32}$—astronomical. And again, this presumes a coherent oracle that does not exist in Enigma's one-time capsule model. Therefore, the Enigma cypher remains **post-quantum robust**: it deprives attackers of structured oracles and anchors security in **one-time unpredictability** and **cognitive asymmetry**.

8. Human and Agentic Workflows (Cybernetics of Presence)

8.1 Human Protocol (Alice)

1. **Orient:** Recognize the six-leaf layout (stable geometry; randomized contents).
2. **Prime:** Determine the category of the next secret character (e.g., lowercase).
3. **Search:** Use gestalt to spot the alphabet order along rings/leaves.
4. **Map:** Apply private bijection π to convert the correct leaf to a **private synonym**.
5. **Emit:** Output only the synonym (e.g., "down"), not the symbol.

6. **Repeat:** For h rounds; verifier checks M_i and forms Λ.

8.2 Agentic Twin (AI)

An enclave-sealed morphism plays the role of π. The agent, given the capsule commitment, derives per-round masks and placements internally and emits only the **masked** witness outcome. The enclave never exports π or unmasked bits; only **indices** and Λ exit. Attestation binds the capability to a code hash, mirroring the human's "mind boundary" with a **sealed boundary**.

8.3 Presence-as-Work

Authority is minted **iff** $\Lambda = 1$ **now**. There is no key in transit, no token to relay. After acceptance, both sides may derive an **ephemeral session key** K_t from a KDF over capsule and transcript digests; absent $\Lambda = 1$, **no key materializes**. Presence becomes an **irreproducible act**, not a transferable artifact.

9. Minimal Leakage and the Two-Layer Shield

The Enigma cypher reduces the witness to **one private synonym per round**, erecting a **two-layer shield**:

1. **Bijection ambiguity:** 6! possible π; the same synonym word need not mean the same leaf across users (or even across epochs).
2. **Secret ambiguity:** Multiple plausible secrets may be in play; the same sequence of synonyms could correspond to different secrets under different maps.

Result: the observer's posterior over (secret, map) pairs remains **broad**; balanced and re-randomized placements prevent posterior sharpening across sessions. No reusable residue accumulates.

10. Threat Model and Defenses

- **Eavesdropping/recording:** Useless; each session's capsule re-keys schedules and masks.
- **Relay/MitM:** No token to forward; the puppet cannot compute masked answers without π and live capsule state.
- **Deepfakes/biometric spoofs:** Irrelevant; the face/voice is just a carrier for a **projection task**. Only **correct membership** matters.
- **Key-loggers and screen scrapers:** The exported artifact is a synonym; without π, it has no universal meaning. Attested viewer eliminates rich telemetry and exports only digests.

- **Oracle probing:** Denied by tokenless design; no deterministic function from inputs to outputs is callable outside the capsule epoch.

- **Eavesdropping/recording:**
  Passive capture buys an attacker nothing because every ceremony is sealed inside a fresh **capsule** with new entropy. The projection layout, per-round placements, and masking schedule are all re-keyed from the capsule seed by a PRF, so two recordings—even of the same user—are statistically independent views. The only exported data are (i) a minimal commitment, (ii) a salted digest of indices, and (iii) accept/reject bits. No symbol values or map hints ever leave the boundary. Formally, the transcript satisfies an information-theoretic goal: $I(S;T) = 0$, i.e., the mutual information between the secret S and public transcript T is zero. Because the **private bijection** π never appears and the one-time masks ρ_i change each round, there is no stable residue to correlate across sessions. Even high-fidelity video of the screen and audio of the verbal synonym cannot be aligned between capsules, since the same synonym may reference different leaves and different symbol populations on each new draw.

- **Relay/MitM:**
  A man-in-the-middle cannot "forward" authority because there is no transferable token—only **presence-as-work**. In each round i, Alice must combine the live capsule schedule (fresh placements and mask ρ_i) with her secret bijection π to produce the correct **synonym**. A relay puppet sees the projection but lacks π; a second puppet might hold a guess of π but lacks the live per-round schedule tied to the current capsule. Without both, the masked witness cannot be computed. Latency bounds and short response windows further constrain "wormhole" attacks: by the time a relayed prompt reaches a remote accomplice and returns, the capsule may have advanced or closed. Because Bob's check is a local membership test on the **current** canvas, any replayed synonym from a prior frame fails. The result is non-forwardability by construction: capability exists only at the location where the mind (or sealed agent) and the live capsule coincide.

- **Deepfakes/biometric spoofs:**
  Visual or audio likeness is irrelevant because the ceremony does not authenticate *appearance*; it authenticates **recognition under hidden geometry**. The user's face/voice is merely a carrier for emitting a private synonym; Bob verifies only that the synonym corresponds to the leaf that currently contains the next character, not that the speaker looks or sounds like Alice. A deepfake can mimic timbre or facial features, but it cannot infer the private bijection π nor the per-round mask ρ_i and placements derived from the capsule. With $6! = 720$ possible π permutations and balanced, re-randomized layouts, the probability that a faker's guessed synonym matches across h rounds is $(1/6)^h$, which collapses rapidly for modest h. Thus, biometric spoofing provides no leverage: the **membership proof** depends on internal knowledge applied to the live manifold,

not on external traits. Degrade or replace the carrier, and the check still holds; remove π, and the deepfake fails immediately.

- **Key-loggers and screen scrapers:**
  Traditional spyware hunts for stable artifacts—keystrokes, clipboard contents, UI coordinates—that map to secrets. Enigma exports **no secret text** and **no coordinates** with reusable meaning. The only outward artifact is a short **synonym** selected from a private codebook; without the user's π, that synonym has no universal semantics (e.g., "down" today may map to a different leaf tomorrow). An **attested viewer** renders the canvas and emits only salted, fixed-shape digests (e.g., index tuples) to the verifier, stripping rich telemetry like cursor paths, hover timings, or pixel buffers that screen scrapers need for reconstruction. Timing is quantized to reduce variance leakage, and layouts are balanced so frequency patterns remain flat. Even if malware records every frame and every uttered synonym, the data cannot be stitched across capsules because the one-time schedule and diffusion ensure non-alignment. Net effect: recorded inputs lack the latent structure necessary to replay or infer the secret.

- **Oracle probing:**
  Effective probing requires a deterministic, queryable function f that maps inputs to outputs in a way the attacker can call repeatedly (classically or in quantum superposition). Enigma withholds such an oracle. The only "output" per round is a human/agent's **masked synonym**, computed using the private bijection π and the capsule's one-time schedule ρ_i and placements. Outside the live capsule epoch, the same query cannot be posed because the projection is re-drawn and the masks are refreshed; inside the epoch, the attacker cannot drive the input space without controlling the capsule. There is no symbol-level response, no gradient, and no stable feedback to support black-box learning or amplitude amplification. Consequently, classical adaptive queries and quantum algorithms like Grover cannot be mounted: there is no f to coherently evaluate across a superposition. Security arises from **tokenless design**, **one-time unpredictability**, and **minimal outputs** that resist being turned into an oracle.

11. Design Parameters and Tuning

11.1 Rounds h and Threshold θ

Choose h and θ against desired **false-accept** (FAR) and **false-reject** (FRR). For example, $h = 10$ strict gives

$$\text{FAR} \approx (1/6)^{10} \approx 1.65 \times 10^{-8}.$$

"FAR is about one point six five times ten to the minus eight."

Thresholding (e.g., accept if ≥ 9 of 10 rounds pass) reduces FRR for humans while keeping FAR negligible (binomial tail).

11.2 Alphabet Size and Category Mix

Larger alphabets improve **entropy per round** but may increase **human error** if visual density rises. Category priming (e.g., switching to digits for a digit round) keeps load low. Carrier-agnostic stimuli (visual/audio/haptic) preserve the core math while serving accessibility.

11.3 Leaf Balance and Margins

Leaves should be **balanced** (near-uniform symbol mass) and separated with **margins** that absorb motor imprecision. Think of a geometric **Hamming ball**: slight selection jitter doesn't flip membership.

11.4 Capsule Windows

Windows must be short enough to defeat **wormhole relays** yet generous enough for human timing. Quantized response bins reduce timing side-channels without harming UX.

12. Formal Leakage Claims (Sketch)

Let V denote all public outputs visible to an adversary: $(\text{commit}, \text{aux}, \text{salted index digests}, \Lambda)$. Because (i) placements and masks are PRF-derived from seed_t and (ii) π never appears in V, the distribution $V|S = s$ is **identical** across s in the secret family (up to parameter-announced aux). Formally, for any two secrets s, s':

$$V|S = s \stackrel{d}{=} V|S = s'.$$

"The distribution of the public view given secret s is equal in distribution to that given secret s-prime."

This implies $I(S;V) = 0$ and aligns with **perfect secrecy** at the transcript layer. (Implementation caveats: enforce fresh seeds, balanced placements, constant-shape outputs, and suppression of rich telemetry.)

13. Comparative Analysis

- **Passwords/OTPs:** Export secrets; replay/relay viable; side channels abound. Enigma exports only **masked synonyms**; no reusable residue, no token to relay.
- **Biometrics:** Spoofable; non-revocable; privacy risk. Enigma is **revocable** (rotate private map), **carrier-agnostic**, and **non-replayable**.

- **Passkeys/HSMs:** Strong transport but still artifact-centric. Enigma **gates capability** on **presence-as-work**, deriving keys only **after** acceptance.

- **Passwords/OTPs:**
  Conventional schemes **export secrets**: a password leaves the mind through a keyboard; a one-time code leaves a device over a channel. Once exported, they are phishable, replayable, relayable, and harvestable at scale (keylogging, shoulder surfing, SIM-swap, SS7 abuse, OCR from screenshots). Even when entropy is high, side-channels (keystroke dynamics, paste timing, autofill artifacts) leak signal. Datastore compromise yields credential stuffing; SMS/email OTPs invite real-time social engineering. By contrast, **Enigma never exports the secret**. Each round emits only a **masked synonym** derived via a private bijection π and a one-time capsule schedule; transcripts contain no reusable residue, and relays lack the live ρ_i masks. Formally, the design targets $I(S;T) = 0$: the mutual information between the secret S and transcript T is null. With no token to forward and no plaintext to steal, replay and relay evaporate, and the usual side-channel surfaces are starved of structure.

- **Biometrics:**
  Biometrics bind access to **inherent traits** (face, voice, fingerprint). They are **spoofable** (deepfakes, 3D masks, gummy fingers), **non-revocable** (you cannot change your face), and carry lasting **privacy risks** (cross-site linkage of templates, coercion, surveillance). Countermeasures—liveness tests, PAD sensors, model updates—raise cost but cannot deliver revocability or erase leakage once a template escapes. In Enigma, nothing about *who you look or sound like* is authoritative. The carrier (voice, click, gaze) merely transports a **private synonym**; only **correct membership** under the hidden geometry matters. If compromise is suspected, rotate the **private map** π and reseed the capsule—**revocation** by design. Because Enigma's proof is **non-replayable** (fresh placements and masks each session) and **carrier-agnostic** (visual, auditory, haptic), it avoids biometric lock-in, limits surveillance value, and restores user agency without sacrificing speed or usability.

- **Passkeys/HSMs:**
  Passkeys (FIDO2/WebAuthn) and HSM-backed keys provide **strong transport security** and phishing resistance, but they remain **artifact-centric**: authority is vested in possession of a private key sealed in hardware. That creates operational frictions (device loss, recovery, multi-device sync), residual relay vectors (e.g., real-time VNC "screen-sharing" scams), supply-chain trust assumptions, and at-rest key material that must be guarded for years. **Enigma gates capability on presence-as-work**: no long-lived private key is consulted until the human/agent demonstrates knowledge *now*. Only **after** acceptance ($\Lambda = 1$) is an ephemeral session key K_t derived from capsule and transcript digests; if proof fails, **no key ex-**

ists to steal or relay. Practically, Enigma complements hardware roots: use attested viewers for rendering, but shift the *authorization primitive* from "who holds a key?" to "who can solve this capsule's projection?". The result is artifact-light authentication with first-class revocability and minimal blast radius.

---

14. Practical Attacks and Residual Risks

- **Coercion:** Mitigate via **duress mappings**—a special perceptual pattern produces an apparently valid transcript that cryptographically restricts policy.
- **Entropy failure:** Health checks; multi-source mixing; abort on degradation.
- **Compromised viewer:** Deploy **attested** minimalist viewers that render only the projection canvas and output only indices/digests.
- **Human error:** Use threshold acceptance and training rounds; rotate witnesses periodically.
- **Coercion:**
  Coercion (rubber-hose, shoulder pressure) targets the human, not the math. Enigma mitigates with **duress mappings**: a reserved perceptual pattern or synonym sequence that yields a **valid-looking transcript** while silently constraining policy (reduced scopes, time-locked actions, write-only audit, silent alert). Because observers cannot distinguish normal from duress outputs—the same minimalist witness and acceptance record—Alice gains **plausible deniability** and a safe escape hatch. Additional controls include rate limits, step-up approvals from independent verifiers, and environment attestation (only accept from attested viewers). Organizational playbooks should encode automatic **break-glass** responses: freeze high-risk operations, rotate π, reseed frequently, and require multi-party continuation. None of this removes the human risk entirely, but it **shifts leverage**: attackers cannot cash out a captured transcript, and forcing a witness under duress **reduces** what can be authorized without detection.
- **Entropy failure:**
  One-time projections are only as strong as their randomness. Enigma treats entropy as a **first-class resource**: capsules derive masks and placements from a DRBG seeded by **multi-source mixing** (OS CSPRNG, on-chip TRNG, user-space jitter, hardware noise), with continuous **health checks** (FIPS 140/ISO fail tests) and backtracking resistance. If any source stalls or biases rise, the capsule **fails closed**: abort the ceremony,

rotate seeds, and log an auditable fault. Designs should integrate **entropy beacons** (where appropriate), per-boot reseeding, and periodic reseeding during long sessions. Deterministic PRFs expand high-quality seeds; they do not create entropy, so monitoring and **entropy accounting** are essential. Administratively, enforce minimum-quality thresholds, telemetry for entropy events (without leaking state), and **automated rollback** paths. By making randomness a monitored dependency rather than an assumption, Enigma constrains the failure mode to **no auth**, not **weak auth**.

- **Compromised viewer:**
  If the rendering environment is hostile, it might exfiltrate pixels, timing, or layout hints. Enigma narrows this surface with an **attested, minimalist viewer**: a small TEE/SEV/SGX-style enclave or measured container that (1) renders only the projection canvas, (2) accepts minimal input events, and (3) emits only **indices/digests**—no raw frames, no cursor telemetry. Remote attestation lets the verifier check code identity and configuration before accepting witnesses. The viewer should implement constant-shape outputs, timing quantization, UI hardening (large margins; no hover states), sealed configuration, and auto-update signed builds. Even if the host OS is noisy, the enclave boundary preserves the projection's **integrity** and the output's **minimalism**, depriving malware of the rich data it needs to reconstruct π or predict ρ_i. Operationally, treat viewer attestation as a policy gate: **no attestation, no ceremony**.

- **Human error:**
  People mis-click, hesitate, or momentarily lose focus. Enigma is engineered for **graceful human variance**. Use **threshold acceptance**—e.g., accept if $\sum_i M_i \geq \theta$ with $\theta < h$—to absorb occasional misses while keeping FAR negligible (binomial tail design). Provide **training rounds** that mirror production visuals but don't authorize actions, plus adaptive difficulty (adjust density, font, spacing) based on observed accuracy. Quantize timing to reduce pressure and leakage; add generous **margins** so small motor slips don't flip membership. Rotate witnesses and, periodically, the private map π to freshen cognitive context without relearning the entire ritual. Offer accessibility modes (contrast, modality switches, haptics) that preserve the same verification math. Finally, instrument **explainable failure**: if a session aborts, report non-sensitive causes (timeout, entropy fault, viewer attestation) so users learn the rhythm without being trained into insecure workarounds.

15. Ethics and Governance

The Enigma cypher reduces **data at rest** and **PII exposure**, supporting **privacy-by-design**. Governance binds policy digests into acceptance, yielding crisp audit statements of the form: "Operation X executed because $\Lambda = 1$ under policy hash $H(\text{policy})$ at time T_t." By decoupling UI narratives from

cryptographic evidence, we eliminate **consent theater** and favor verifiable **presence-as-work**.

The **Enigma cypher** operationalizes privacy-by-design by minimizing both **data at rest** and **PII exposure** across the entire lifecycle of authentication. Secrets never leave the mind or sealed agent; each ceremony is encapsulated in a fresh capsule whose outputs are strictly **boolean evidence** and salted indices, not identities or attributes. As a result, the residual dataset is a thin, audit-ready ledger rather than a honeypot of personal data. Formally, the transcript is engineered so that the mutual information between the secret and what is stored satisfies $I(S;T) = 0$, aligning with purpose limitation and storage minimization. *"I of S semicolon T equals zero: the transcript reveals no information about the secret."*

Governance is made cryptographically explicit: acceptance is bound to a policy digest so that every action is justified by a compact, tamper-evident clause—"Operation X executed because $\Lambda = 1$ under policy hash H(policy) at time T_t." This yields **crisp audit statements** that are verifiable independent of user interface narratives or screenshots.
"Operation X executed because lambda equals one under policy hash H of policy at time T sub t."

By **decoupling UI narratives from cryptographic evidence**, Enigma eliminates **consent theater**—the practice of relying on ambiguous pop-ups or click-throughs as proof of authorization. Instead, it privileges **presence-as-work**: a right to proceed exists only when a live, one-time proof succeeds. No enduring identifier is required; no biometric template is stored; and no reusable token is issued. Post-acceptance, any session key is derived ephemerally from the capsule and transcript (e.g., $K_t = \mathrm{KDF}(EC_t, T)$), then expires with policy. *"K sub t equals K D F of the capsule and the transcript."*

This architecture supports **revocability** (rotate the private map), **accountability** (policy-bound logs), and **least disclosure** (synonyms instead of secrets). It aligns naturally with regulatory principles—data minimization, storage limitation, integrity/confidentiality—while providing organizations a high-assurance audit trail that is small, comprehensible, and **cryptographically self-authenticating**. Presence becomes the measurable commodity; privacy becomes the default.

16. Synthesis: Why the Human (and the Twin) Win

The **Enigma cypher** operationalizes privacy-by-design by minimizing both **data at rest** and **PII exposure** across the entire lifecycle of authentication. Secrets never leave the mind or sealed agent; each ceremony is encapsulated in a fresh capsule whose outputs are strictly **boolean evidence** and salted indices, not identities or attributes. As a result, the residual dataset is a thin, audit-ready ledger rather than a honeypot of personal data. Formally, the transcript is

engineered so that the mutual information between the secret and what is stored satisfies $I(S;T) = 0$, aligning with purpose limitation and storage minimization. *"I of S semicolon T equals zero: the transcript reveals no information about the secret."*

Governance is made cryptographically explicit: acceptance is bound to a policy digest so that every action is justified by a compact, tamper-evident clause—"Operation X executed because $\Lambda = 1$ under policy hash H(policy) at time T_t." This yields **crisp audit statements** that are verifiable independent of user interface narratives or screenshots.
"Operation X executed because lambda equals one under policy hash H of policy at time T sub t."

By **decoupling UI narratives from cryptographic evidence**, Enigma eliminates **consent theater**—the practice of relying on ambiguous pop-ups or click-throughs as proof of authorization. Instead, it privileges **presence-as-work**: a right to proceed exists only when a live, one-time proof succeeds. No enduring identifier is required; no biometric template is stored; and no reusable token is issued. Post-acceptance, any session key is derived ephemerally from the capsule and transcript (e.g., $K_t = \mathrm{KDF}(EC_t, T)$), then expires with policy. *"K sub t equals K D F of the capsule and the transcript."*

This architecture supports **revocability** (rotate the private map), **accountability** (policy-bound logs), and **least disclosure** (synonyms instead of secrets). It aligns naturally with regulatory principles—data minimization, storage limitation, integrity/confidentiality—while providing organizations a high-assurance audit trail that is small, comprehensible, and **cryptographically self-authenticating**. Presence becomes the measurable commodity; privacy becomes the default.

17. Conclusion

Conclusion

The Enigma cypher advances a simple but decisive proposition: when identity is treated as **presence-as-work** rather than **possession of an artifact** or **resemblance to a template**, both security and privacy improve at once. Throughout this paper we showed how a human (or sealed agent) converts the ordinary act of **gestalt recognition**—seeing ordered alphabets amid clutter—into a cryptographic witness that never exports the secret itself. A private bijection π maps six perceptual leaves to six private synonyms; each round emits only a synonym, not a symbol. Because the projection is re-keyed per capsule with one-time masks and balanced placements, the transcript carries **no reusable residue**. In Shannon's terms, the design targets $I(S;T) = 0$: the stored record reveals nothing about the secret.

Two structural choices power the scheme. First, **diffusion**: multiple letters live within each holographic leaf or ring zone, erasing frequency fingerprints and ensuring that no leaf acquires stable semantic bias. Second, **confusion**:

the private map permutes meanings across six hyperplanes, yielding $6! = 720$ equiprobable bijections that deny an attacker any public anchor. These are embedded in a projection manifold with **ring-rotation diffusion** across $30 \times 30 \times 12 = 10{,}800$ micro-states per ring (and $(10{,}800)^R$ across R rings). The consequence is a **combinatorial wall**: the honest user's visual system collapses the search to a glance, while an adversary faces a branching inference tree whose size remains astronomical even under optimistic assumptions.

We quantified that wall for a 20-character secret and the realistic case where the user maintains n **encoded secrets**. Without knowing which secret is active, the attacker's initial hypothesis space scales linearly in n: $|\mathcal{H}_0| = n \cdot 6! \cdot 72^{20} \approx n \cdot 1.00921 \times 10^{40}$. Under the crude i.i.d. pruning model where roughly $5/6$ of wrong hypotheses die each round, isolating a unique explanation would demand about $h \approx 52 + \lceil \log_6 n \rceil$ **perfectly observed, co-epoch rounds**—a condition the protocol deliberately prevents via short capsule windows and mandatory reseeding. This asymmetry is the heart of Enigma: the legitimate prover's work stays light and local; the attacker's work either fails to accumulate or explodes.

The **post-quantum** stance arises not from a new algebraic hardness assumption but from **removing the oracle** that quantum algorithms require. Grover's search can quadraticly speed unstructured search *if* a coherent oracle exists. Enigma refuses to be that oracle. The only observable per round is a **masked synonym** grounded in π and the capsule's one-time schedule; outside the capsule epoch, the same query cannot even be posed. Thus, even though $\sqrt{n \cdot 6! \cdot 72^{20}}$ is already prohibitive (on the order of $10^{20}\sqrt{n}$), the more decisive fact is structural: there is **no function** for a quantum adversary to evaluate in superposition. Security rests on **one-time unpredictability** and **minimal outputs**, not on fragile promises of long-term algebraic hardness.

This primitive reshapes the **threat surface**. Eavesdroppers and recorders gain nothing; each capsule redraws the geometry and masks. Relays and MitM puppets cannot "forward" capability because there is no transferable token, only a live act that requires π *and* the current capsule state. Deepfakes fail because likeness is not an input to acceptance; only **correct membership** under hidden geometry matters. Key-loggers and screen scrapers starve for structure: the artifact is a synonym with user-specific semantics, and the **attested, minimalist viewer** exports only salted indices/digests. Oracle probing collapses because no deterministic, repeatable input–output map exists outside the live epoch.

The **governance** and **privacy** benefits are equally central. By binding acceptance to a policy digest, we obtain crisp, self-authenticating audit statements: "Operation X executed because $\Lambda = 1$ under policy hash $H(\text{policy})$ at time T_t." This decouples UI narratives from cryptographic facts, eliminating "consent theater" and replacing it with verifiable authorization. Because secrets never leave the mind/agent and no biometric templates are stored, **data at rest** is reduced to a thin ledger of commitments, indices, and outcomes—valuable for accountability yet innocuous for privacy. Revocation is practical (rotate π; reseed capsules). Least disclosure is native (synonyms instead of symbols). Regulatory alignment follows from first principles: data minimization, storage limitation, integrity, and confidentiality are inherent rather than bolted on.

From an **engineering** perspective, Enigma is tunable. Designers pick rounds h and thresholds θ to meet false-accept/false-reject targets—for example, strict $h = 10$ yields FAR $(1/6)^{10} \approx 1.65\times10^{-8}$, while thresholding absorbs human slips without sacrificing security (binomial tails). Visual density, alphabet mix, and ring count adjust usability and per-round entropy; margins and timing quantization reduce side-channel leakage. An **attested viewer**—small, measured, and export-minimal—anchors the human–machine interface with a verifiable boundary, ensuring that even on a noisy host the projection remains trustworthy and the outputs stay lean.

We do not claim Enigma abolishes all risk. Humans can be coerced; devices can be compromised; entropy sources can degrade. But the primitive offers **clean levers**: duress mappings that produce valid-looking yet policy-restricted transcripts, "fail-closed" entropy health checks, and "no attestation, no ceremony" viewer gating. Crucially, when failure happens, it tends to fail by **withholding** authorization rather than by **leaking** it. That inversion of failure mode—no auth vs. weak auth—is a meaningful shift in real-world deployments.

Compared to incumbent approaches, the contrast is stark. Passwords and OTPs **export secrets** and invite replay and relay at industrial scale. Biometrics are **non-revocable** and privacy-heavy; liveness helps but cannot make a stolen template private again. Passkeys/HSMs are robust on transport but remain **artifact-centric**, leaving organizations to manage key custody, recovery, and long-lived risk. Enigma re-centers authority on a **momentary act**—recognition under hidden geometry—so that keys (when needed) are derived **after** acceptance and vanish with policy. Enigma can live **alongside** passkeys (attested viewer + presence-as-work), or **in front** of high-value actions as a second, human-cognizable factor that resists both phishing and deepfake mediation.

We also note the broader **cybernetic** implication: Enigma treats cognition as a programmable control surface. The human or agentic twin becomes a *capability switch* that closes only when the environment presents the correct, one-time manifold and the private map is applied correctly. That reframing is useful not just for login, but for approvals, step-up authorizations, automated agents acting under human oversight, and safe handoffs between humans and AI systems. Because the witness is minimal and the policy binding is explicit, compliance and audit gain a **clear, cryptographic spine**.

In sum, the Enigma cypher is not a tweak to passwords nor a cosmetic add-on to biometrics. It is a **different primitive**: a tokenless, appearance-agnostic, oracle-free protocol where **recognition becomes proof**. It turns human perceptual strength into a verifier-friendly ceremony, shields secrets behind diffusion and confusion, denies structure to classical and quantum adversaries, and leaves behind only a compact, policy-bound record of what was authorized and why. In an era saturated with synthetic media, scalable surveillance, and emerging quantum capabilities, the ability to authenticate without exporting identity—and to authorize without accumulating residue—is not merely elegant; it is necessary. If we adopt presence-as-work as the organizing principle for access, Enigma offers the rare combination of stronger guarantees, lighter data footprints, and better human ergonomics. That is a foundation worth building on.

Appendix A: Selected Equations (with TTS)

1. **Per-round guess pass:**
   $p_{\text{guess}} = 1/6$.
   "Per round, the chance to pass by guessing is one sixth."

2. **Strict false-accept:**
   $\Pr[\text{FA}_{\text{strict}}(h)] = (1/6)^h$.
   "False accept under strict acceptance is one over six to the h."

3. **Threshold false-accept:**
   $\Pr[\text{FA} * \text{thr}(h, \theta)] = \sum * k = \theta^h \binom{h}{k} (1/6)^k (5/6)^{h-k}$.
   "Sum from k equals theta to h of h choose k times one-sixth to the k times five-sixths to the h minus k."

4. **Initial hypothesis count (unknown secret, unknown map, n encoded secrets):**
   $|\mathcal{H}_0|; =; n \cdot 6! \cdot 72^{20}; \approx; n \cdot 1.00921 \times 10^{40}$.
   "H naught equals n times six factorial times seventy-two to the twentieth, about n times one point zero zero nine two one times ten to the fortieth."

5. **Rounds to isolate (order-of-magnitude, co-epoch):**
   $6^{52} \approx 2.90981 \times 10^{40}, \quad h \approx 52 + \lceil \log_6 n \rceil$.
   "Six to the fifty-two is about two point nine zero nine eight one times ten to the fortieth; h is roughly fifty-two plus the ceiling of log base six of n."

6. **Ring diffusion per ring:**
   $30 \times 30 \times 12 = 10{,}800$.
   "Ten thousand eight hundred micro-states per ring."

7. **Compound ring configurations:**
   $(10{,}800)^R$.
   "Ten thousand eight hundred to the power R."

8. **Perfect-secrecy style transcript condition:**
   $I(S; T) = 0$.
   "Mutual information between the secret and transcript equals zero."

9. **Grover lower bound (no oracle available in practice):**
   $\sqrt{N} = \sqrt{n}, \sqrt{6!}, 72^{10} \approx \sqrt{n} \times 1.005 \times 10^{20}$.
   "Square root of N is approximately square-root of n times one point zero zero five times ten to the twenty."

Practical Epilogue (Implementation Heuristics)

- Keep h in the 8–12 range; apply thresholding to absorb slips while achieving $< 10^{-8}$ FAR.
- Use **balanced leaves**, **large margins**, and **category priming** to preserve usability.

- Enforce **fresh** capsules per ceremony; never reuse seeds.
- Deploy a minimal, **attested viewer** that renders only the canvas and exports only salted **index digests**.
- Rotate **private maps** periodically; train users with gentle rounds that **do not** authorize actions.
- For duress resistance, include a **special mapping** that yields a valid-looking transcript with a **restricted policy bit**.

By respecting these principles, the Enigma cypher delivers exactly what the era demands: **knowing without showing**, secured by human gestalt and hardened by cryptography—robust against surveillance, deepfakes, and quantum futures alike.