STEBR: A Timed-Erasure, Threshold-Gated Backup Ratchet

for Bounding Recovery-Key Compromise in Signal-Style Encrypted Backup Systems

Shaurya Pratap Singh
Department of Computer Science and Automation,
Indian Institute of Science,
Bengaluru, India
shaurya.pratap@fsid-iisc.in

Abstract

The Signal Protocol's Double Ratchet and X3DH/PQXDH handshake give in-transit messages forward secrecy and post-compromise security: compromising a session key does not expose past traffic, and the protocol self-heals after a fresh Diffie–Hellman step. Encrypted *backups*, by contrast, are commonly protected by a single static secret — a "Backup Recovery Key" — generated once and held constant until manually rotated. We show, with an explicit attack, that this baseline design provably fails even a minimal forward-secrecy-style security notion: disclosure of the key at any time exposes the entire backup history, with no self-healing. This is not a hypothetical concern: a June 26, 2026 joint FBI/CISA advisory attributes exactly this exploitation pattern to two Russian intelligence-linked clusters, tracked as UNC5792 and UNC4221, who obtained victims' Backup Recovery Keys through impersonation-based social engineering rather than cryptanalysis. We propose STEBR (Secure Timed-Erasure Backup Ratchet), a backup-key architecture built from three composable layers: (1) a self-erasing hash-chain key ratchet, so that compromise of the current backup key exposes only a bounded, recent window of history rather than the full archive; (2) (t, n) threshold secret sharing of the current epoch key across independently held devices/guardians, so that no single credential extracted in one social-engineering interaction is sufficient; and (3) an interactive, rate-limited, out-of-band confirmation gate on any restore request, so that possession of valid recovery material is necessary but not sufficient to complete a restore. We give formal security definitions for each property and prove them via standard reductions (PRF security of the key-derivation function, IND-CPA security of the backup AEAD scheme, the information-theoretic secrecy of Shamir sharing, and the authenticity of the existing ratchet-protected control channel). All three layers compose on top of existing Signal Protocol primitives; none requires modifying the Double Ratchet, X3DH/PQXDH, or the wire format of message envelopes. This is a proposal for hardening the backup-key management layer specifically; we make no claim that the Signal Protocol's transport-layer cryptography is broken or requires replacement — the cited advisory itself states plainly that it is not.

Keywords: secure messaging protocols, Signal Protocol, Double Ratchet, encrypted backups, key management, forward secrecy, post-compromise security, key-insulated encryption, threshold secret sharing, Shamir secret sharing, social engineering, phishing, applied cryptography

Contents

1 Introduction

1.1 Motivation

Modern secure messengers built on the Signal Protocol achieve two properties for in-transit communication that are, by now, close to a gold standard in applied cryptography: *forward secrecy* (compromise of a current key does not reveal past plaintexts) and *post-compromise security* (the protocol recovers security after compromise, once a fresh Diffie–Hellman exchange occurs) [2, 1, 3]. These properties hold for the message stream. They do not, in typical deployed designs, extend to *encrypted backups* of that message stream, which are usually protected by deriving a single, long-lived symmetric key from a secret the end user is expected to generate once, record, and safeguard indefinitely.

This split matters because it creates exactly one artifact whose disclosure defeats the archive's confidentiality in full, and because that artifact is, by design, meant to be manageable by a human being — written down, stored in a password manager, or read aloud. Any component with that shape is a natural target for social engineering, independent of how strong the surrounding cryptography is.

This is not a theoretical worry. A joint advisory from the U.S. Federal Bureau of Investigation (FBI) and Cybersecurity and Infrastructure Security Agency (CISA), published June 26, 2026, describes a campaign by two Russian intelligence-linked clusters, publicly tracked as `UNC5792` (associated with Russian FSB-affiliated officers) and `UNC4221` (linked to Russian military intelligence), that escalated from stealing verification codes and abusing linked-device enrollment to specifically harvesting victims' Backup Recovery Keys via impersonated "support" messages inside the app. The advisory is explicit that the messaging application's encryption itself was not broken; the compromise is entirely at the level of a single, statically-issued, human-held secret. The U.S. State Department's Rewards for Justice program has since offered a reward of up to $10 million for information on `UNC5792`. Targets described in the advisory include current and former government officials, military personnel, political figures, journalists, and officials in Ukraine.

1.2 Contributions

This paper makes the following contributions.

1. We give a formal security definition, *Bounded-Exposure security* (Definition 4), that captures the property a backup-key scheme should have if a single leaked recovery credential is to expose only a limited window of archived history rather than the full archive.

2. We prove, via an explicit distinguishing attack (Claim 1), that the baseline single-static-key backup design used in current deployments does *not* satisfy Bounded-Exposure security — with advantage 1, using only the credential itself.

3. We construct STEBR, a three-layer backup-key architecture (Section 5) that achieves Bounded-Exposure security (Theorem 1), a threshold-secrecy property against sub-threshold credential compromise (Lemma 1, Corollary 1), and a Live-Phishing Containment property (Theorem 2) that bounds the success probability of a live social-engineering interaction regardless of its persuasiveness.

4. We give a composed theorem (Theorem 3) bounding the overall advantage of an adversary restricted to a single social-engineering interaction and no control over the victim's other enrolled devices, and we discuss the usability trade-offs the construction introduces (Section 6).

All three layers are built from primitives already present in Signal-style systems (a KDF, an AEAD scheme, and the existing authenticated control channel); no new cryptographic primitive or change to the transport protocol is proposed or required.

1.3 Related Work

Signal Protocol analysis. The Double Ratchet [2] and X3DH [1] specifications, and their formal treatment by Cohn-Gordon et al. [3], establish forward secrecy and post-compromise security for the *messaging* layer under standard DH and PRF/KDF assumptions. Our work does not modify this layer; it observes that the backup layer, as commonly deployed alongside it, does not inherit these properties and proposes bringing analogous guarantees to that layer specifically.

Forward-secure and key-insulated cryptography. The general technique of evolving a secret key over discrete time periods so that compromise at time t does not expose periods before t traces to forward-secure signatures and encryption [4, 5], and to key-insulated cryptography [6], which additionally splits trust between a long-lived "helper" and a short-lived device key. Layer 1 of STEBR (Section 5.1) is a direct, backup-specific instantiation of the forward-secure key-evolution paradigm; unlike prior key-insulated schemes, it does not require an always-online helper, only periodic re-derivation on the user's own device(s).

Threshold and proactive secret sharing. Shamir's (t, n) secret sharing [7] is the classical primitive for eliminating single points of credential failure; proactive secret sharing [8] additionally refreshes shares periodically so that an adversary who compromises shares at different times cannot accumulate a threshold. Layer 2 (Section 5.2) applies static Shamir sharing to the current epoch key; adopting a proactive refresh schedule on top of Layer 1's epoch ratchet is noted as a natural extension in Section 6.

Out-of-band confirmation and delay-based defenses. Delay/notify patterns are used elsewhere in security engineering (e.g. mandatory holds before account-recovery or high-risk-transaction completion) as a way of trading a small amount of availability for a large reduction in the success rate of live social-engineering attacks, which typically rely on manufactured urgency and cannot be sustained across a multi-hour or multi-day window. Layer 3 (Section 5.3) is an application of this pattern to backup restoration specifically, riding on the already- authenticated Double Ratchet control channel for delivery.

Positioning. To our knowledge, existing public analyses of the UNC5792/UNC4221 campaign (security-vendor advisories and the FBI/CISA PSA itself) are descriptive: they document the social-engineering tactic and recommend user-facing hygiene (rotate the key, verify linked devices, distrust in-app "support" messages). We are not aware of a prior proposal that reformulates the underlying design gap as a precise security game and gives a constructive scheme with a matching proof. That is the intended contribution here.

2 Preliminaries

Definition 1 (PRF security)**.** *A function family $F : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$ is a secure pseudorandom function (PRF) if for all PPT distinguishers $\mathcal{D}$,*

$$\mathsf{Adv}_F^{\mathsf{PRF}}(\mathcal{D}) := \left| \Pr_{k \leftarrow \mathcal{K}}[\mathcal{D}^{F(k,\cdot)} = 1] - \Pr_{f \leftarrow \mathrm{Funcs}(\mathcal{X},\mathcal{Y})}[\mathcal{D}^{f(\cdot)} = 1] \right|$$

is negligible in the security parameter. We assume HKDF *(instantiated, e.g., as HKDF over HMAC-SHA256 or similar) is a secure PRF in this sense when keyed by its first argument.*

Definition 2 (IND-CPA security)**.** *A symmetric encryption scheme* $(\mathsf{Enc}, \mathsf{Dec})$ *is IND-CPA secure if for all PPT adversaries* $\mathcal{A}$*, the advantage in the standard left-or-right indistinguishability game is negligible. We assume the backup AEAD scheme (e.g. AES-256-GCM or ChaCha20-Poly1305, as already used for message encryption) satisfies this.*

Definition 3 (Shamir (t, n) secret sharing [7])**.** *A dealer holding secret* $K \in \mathbb{Z}_p$ *picks a uniformly random polynomial* $f \in \mathbb{Z}_p[x]$ *of degree* $t-1$ *with* $f(0) = K$*, and issues share* $s_\ell = f(\ell)$ *to party* $\ell \in \{1, \ldots, n\}$*. Any* t *shares determine* f *(and hence* K*) via Lagrange interpolation; any* $t-1$ *or fewer shares are distributed independently of* K*.*

3 The Baseline Design and Its Failure

We model the deployed baseline as follows. A single master secret $K_{\text{recovery}} \in \{0,1\}^\lambda$ (in practice, a 64-character string) is generated once at backup setup. All backed-up content $M_1, \ldots, M_n$, produced at arbitrary times $t_1 \le \cdots \le t_n$, is encrypted under keys derived solely from K_{recovery}:

$$C_i = \mathsf{Enc}_{K_i}(M_i), \qquad K_i = \mathsf{HKDF}(K_{\text{recovery}}, i).$$

K_{recovery} itself is never rotated except by explicit manual user action, and no other value is required to decrypt any C_i given K_{recovery}.

Definition 4 (Bounded-Exposure security)**.** *Let* Π *be a backup-key scheme with epoch-indexed message keys* $M_0, M_1, \ldots$ *and a designated "current credential"* κ_j *available to the user (and hence exfiltratable by an adversary) at epoch* j*. For retention window* W*, adversary* $\mathcal{A}$*, and security parameter* λ*, define game* $\mathsf{BE\text{-}Exp}_{\Pi,W}$*:*

1. *Challenger runs* Π *to a random current epoch* $j \le q(\lambda)$ *and gives* $\mathcal{A}$ *the current credential* κ_j*.*
2. $\mathcal{A}$ *selects a target epoch* $i^* \le j - W$ *and a pair of equal-length messages* (m_0, m_1)*.*
3. *Challenger samples* $b \leftarrow \{0,1\}$ *and returns* $c^* = \mathsf{Enc}_{M_{i^*}}(m_b)$*.*
4. $\mathcal{A}$ *outputs* b' *and wins if* $b' = b$*.*

Π *has Bounded-Exposure security if* $\left|\Pr[b' = b] - \frac{1}{2}\right| \le \mathsf{negl}(\lambda)$ *for every PPT* $\mathcal{A}$ *and every* $W \ge 1$*.*

Claim 1. *The baseline scheme of Section 3 does not satisfy Bounded-Exposure security for any* $W < j$*: there is an adversary* $\mathcal{A}$ *with advantage exactly* $\frac{1}{2}$ *(i.e. it wins* $\mathsf{BE\text{-}Exp}$ *with probability* 1*), for every choice of target epoch* i^**.*

Proof. Given $\kappa_j = K_{\text{recovery}}$, $\mathcal{A}$ computes $M_{i^*} = \mathsf{HKDF}(K_{\text{recovery}}, i^*)$ directly — this is possible for *every* $i^* \le j$, since the baseline scheme places no restriction on which epoch keys are derivable from the current credential. $\mathcal{A}$ then computes $\mathsf{Dec}_{M_{i^*}}(c^*)$ and compares against m_0, m_1 to recover b with certainty. Since this holds for arbitrary i^*, including $i^* \ll j - W$ for any retention window W, the scheme fails Definition 4 maximally: there is no window outside of which exposure is bounded. □

This formalizes, in the language of a security game, the property underlying the June 2026 advisory: obtaining the single current recovery credential is unconditionally sufficient to decrypt the entire archive, independent of how long ago any given message was sent or how the credential was obtained (social engineering, device theft, or otherwise). The remainder of the paper constructs a scheme that avoids Claim 1 by construction.

4 Threat Model

Assumption 1 (Adversary capabilities)**.** *The adversary $\mathcal{A}$ is a remote, PPT party who may:*

1. *Conduct one social-engineering interaction with the victim (e.g. an impersonated support conversation) of bounded real-time duration d, extracting some secret material the victim is induced to disclose during that interaction — modeled as at most one recovery credential or share.*

2. *Observe and/or actively control the network path between the victim and the backup storage service.*

$\mathcal{A}$ is assumed not *to: break the PRF/KDF or AEAD primitives in use; solve the discrete-log/DH problem in the deployed group where relevant; compromise t or more of the n independently held key shares (Section 5.2) within one interaction; or control any of the victim's already-enrolled devices. This models the demonstrated real-world attack (credential phishing) rather than a cryptanalytic break, consistent with the advisory's own characterization of the campaign.*

5 The STEBR Construction

5.1 Layer 1: Self-erasing epoch ratchet

Time is partitioned into epochs $i = 0, 1, 2, \ldots$ (e.g. one per backup snapshot, or per fixed wall-clock interval Δ). Given a master secret sampled once at setup,

$$\begin{aligned} K_0 &= \mathsf{HKDF}(\text{master secret}, \texttt{"stebr-epoch-0"}), \\ K_{i+1} &= \mathsf{HKDF}(K_i, \texttt{"stebr-epoch-next"}), \\ M_i &= \mathsf{HKDF}(K_i, \texttt{"stebr-message-key"}). \end{aligned}$$

Content produced in epoch i is encrypted under M_i. As soon as K_{i+1} has been derived, K_i is deleted everywhere it was computed (client, and any server-side re-encryption worker, if used); only the current epoch key is ever retained or exposed to the user as the "current credential." A retention policy periodically re-wraps content older than W epochs forward under the then-current key (discarding the old ciphertext), or, where an organizational legal-hold requirement demands full history, re-wraps it under a separate escrow key that is architecturally outside the end-user recovery path entirely and is not reachable via κ_j.

Theorem 1. *If* HKDF *is a secure PRF and* Enc *is IND-CPA secure, then the Layer 1 construction satisfies Bounded-Exposure security (Definition 4).*

Proof. We proceed by a hybrid argument over the epochs between the target i^* and the exposed current epoch j, where $j - i^* > W$ by the definition of the game.

For $k = 0, 1, \ldots, j - i^*$, define hybrid H_k: sample $K_{i^*}, \ldots, K_{i^*+k-1}$ as independent uniform random strings (instead of via the HKDF chain), and compute $K_{i^*+k}, \ldots, K_j$ honestly from K_{i^*+k} onward exactly as in the real scheme, so that the distribution of K_j handed to $\mathcal{A}$ is preserved bit-for-bit across all H_k. H_0 is the real experiment.

Adjacent hybrids are PRF-indistinguishable. H_k and H_{k+1} differ only in how K_{i^*+k+1} relates to K_{i^*+k}: in H_k, K_{i^*+k} is random and $K_{i^*+k+1} = \mathsf{HKDF}(K_{i^*+k}, \cdot)$ is computed honestly from it; in H_{k+1}, K_{i^*+k+1} is itself replaced by an independent random string (with the honest chain resuming one step later). A distinguisher between H_k and H_{k+1} yields a PRF distinguisher $\mathcal{D}$ against $\mathsf{HKDF}(\cdot, \texttt{"stebr-epoch-next"})$: $\mathcal{D}$ samples K_{i^*+k} itself, queries its PRF oracle (real or

random) on the fixed label to obtain a candidate K_{i^*+k+1}, continues the rest of the chain honestly from there, and runs the H_k/H_{k+1} distinguisher on the result. Hence

$$|\Pr[H_k = 1] - \Pr[H_{k+1} = 1]| \leq \mathsf{Adv}^{\mathsf{PRF}}_{\mathsf{HKDF}}(\lambda).$$

Summing over the (at most $q(\lambda)$, polynomially bounded) hops from H_0 to H_{j-i^*},

$$|\Pr[H_0 = 1] - \Pr[H_{j-i^*} = 1]| \leq q(\lambda) \cdot \mathsf{Adv}^{\mathsf{PRF}}_{\mathsf{HKDF}}(\lambda) = \mathsf{negl}(\lambda).$$

Final hybrid reduces to IND-CPA. In H_{j-i^*}, K_{i^*} (and hence $M_{i^*} = \mathsf{HKDF}(K_{i^*}, \cdot)$, again by PRF security of HKDF applied once more, absorbed into the same negligible term) is distributed independently of K_j, which is all $\mathcal{A}$ receives. An adversary distinguishing b from $c^* = \mathsf{Enc}_{M_{i^*}}(m_b)$ in H_{j-i^*} is then exactly an IND-CPA adversary against Enc under a key unknown to and independent of its view, giving

$$\left|\Pr[b' = b \mid H_{j-i^*}] - \tfrac{1}{2}\right| \leq \mathsf{Adv}^{\text{IND-CPA}}_{\mathsf{Enc}}(\lambda).$$

Combining both bounds by the triangle inequality proves the theorem. □

Remark 1. *Theorem 1 is deliberately pessimistic in one respect: the challenger in Definition 4 is allowed to answer the challenge ciphertext using a key it has, in the real system, already erased. This models an adversary who could somehow inspect deleted memory, and is included so the proof does not implicitly rely on erasure being physically perfect; the argument goes through on cryptographic indistinguishability of the hybrid alone.*

5.2 Layer 2: (t, n) threshold recovery

Rather than exposing K_j directly to the user as a single string, STEBR Shamir-shares it: the current holder of K_j samples a random degree-$(t-1)$ polynomial f over $\mathbb{Z}_p$ with $f(0) = K_j$ and issues $s_\ell = f(\ell)$ to n independent parties (e.g. the user's other enrolled devices, and optionally a designated recovery contact). Reconstruction of K_j requires any t of the n shares.

Lemma 1 ([7])**.** *For any set of $t' < t$ shares, the joint distribution of those shares is independent of K_j: for every candidate value K', there is exactly one degree-$(t-1)$ polynomial consistent with the given t' points and $f(0) = K'$.*

Corollary 1. *Under Assumption 1 (a single interaction yields at most one share, so $t' \leq 1$), for any $t \geq 2$ the adversary's view after a successful social-engineering interaction is, by Lemma 1, statistically independent of K_j. Consequently such an adversary gains no advantage in $\mathsf{BE\text{-}Exp}$ (Definition 4) beyond guessing: obtaining $\kappa_j = K_j$ in that game requires t shares, which Assumption 1 precludes in one interaction.*

We recommend $t = 2, n = 3$ in practice (e.g. new device, one existing enrolled device, one optional guardian/server-held share gated by Layer 3), keeping the common single-companion-device case a two-party approval while making a single phished value information-theoretically useless.

5.3 Layer 3: Interactive, rate-limited, out-of-band gate

A restore request R for account A, even carrying a valid reconstructed K_j, triggers the following before being honored.

Algorithm 1 Restore-gate protocol (executed by the backup-storage service)

1: Receive restore request R with reconstructed-key claim and requesting-device identifier D_{new}.
2: Send an authenticated notification over each enrolled device's existing ratchet-protected control channel: "A restore was requested from a new device. If this was not you, tap Deny."
3: Start delay timer τ (e.g. 24–72h).
4: **if** Deny received from any enrolled device **then** reject R permanently, invalidate K_j (forcing an out-of-cycle Layer 1 ratchet step), and alert the user via a separate registered channel.
5: **else if** Approve received from any enrolled device **then** release the archive to D_{new} (optionally after a short confirmation window).
6: **else if** τ elapses with no response **and** this is A's first restore request in ≥ 30 days **then** allow (bounded fail-open, to preserve availability for users who have lost every enrolled device).

Definition 5 (Live-Phishing Containment)**.** *A recovery scheme achieves Live-Phishing Containment with delay τ if any adversary who obtains valid recovery credentials via one social-engineering interaction of duration $d \ll \tau$ completes a restore only with probability bounded by the probability that no enrolled device of the victim acts on the corresponding notification within τ.*

Theorem 2. *The Restore-gate protocol achieves Live-Phishing Containment (Definition 5) with delay τ, assuming (a) the control channel delivers authenticated notifications with probability at least $1 - \epsilon$ within τ, and (b) at least one enrolled device is checked by the legitimate user with probability at least $1 - \delta$ within τ.*

Proof. By Assumption 1, the adversary does not control any enrolled device and has not broken the ratchet primitives protecting the control channel; hence it can neither forge an Approve nor suppress delivery of the notification undetectably. Completing a restore therefore requires one of: (i) an explicit Approve from a device the adversary does not control, which by assumption does not occur within a single interaction; or (ii) τ elapsing with the fail-open condition satisfied, whose probability is bounded, by a union bound over non-delivery (ϵ) and non-response by an attentive user (δ), by $\epsilon + \delta$. Hence the adversary's success probability is at most $\epsilon + \delta$, independent of d or of the persuasiveness of the interaction that produced the credential. □

5.4 Composed guarantee

Theorem 3. *Under Assumption 1 and the hypotheses of Theorems 1 and 2, an adversary restricted to one social-engineering interaction succeeds in recovering any backup content with probability at most*

$$\mathsf{negl}(\lambda) \;+\; (\epsilon + \delta),$$

and any content it does recover lies within the most recent W-epoch retention window.

Proof. By Corollary 1, a single interaction yields at most $t - 1$ shares, which are statistically independent of K_j; the adversary therefore cannot reach the "adversary has κ_j" precondition of Theorem 1 directly, and any attempt to instead complete a full restore must go through the Restore-gate protocol, whose success probability is bounded by Theorem 2. Combining the two bounds via a union bound, and noting that Theorem 1 already restricts any content that *is* obtained to epochs within W of the current one, gives the stated bound. □

6 Practical Discussion

- **Usability cost.** The bounded fail-open in Layer 3 line 6 is a disclosed, tunable policy choice trading a small residual availability risk for the "lost every device" case; it should be configurable per risk profile, and disabled for accounts self-identified as high-risk (the categories named in the advisory: government officials, military personnel, journalists).
- **No single human-portable secret.** Because recovery requires an interactive, multi-party ceremony rather than typing a string into a prompt, there is no longer a value a fake "Support" conversation can simply request and receive in full — removing the specific mechanical step ('paste your key here') the advisory describes, without altering the underlying AEAD/KDF primitives.
- **Proactive refresh (extension).** Combining Layer 2 with proactive secret sharing [8] — periodically re-sharing K_j with fresh randomness without changing the secret — would further prevent an adversary from accumulating a threshold of shares phished across multiple, separated interactions; we leave a full treatment of this combination to future work.
- **Compatibility.** All three layers sit on top of existing AEAD/KDF primitives and the existing authenticated control channel used for messages; no change to X3DH/PQXDH, the Double Ratchet, or the wire format of message envelopes is required.
- **Scope.** Consistent with the advisory's own characterization, this proposal addresses an operational/key-management gap, not a cryptanalytic weakness in Signal's transport-layer protocol.

7 Conclusion

The June 2026 UNC5792/UNC4221 campaign succeeded against the one component of Signal-style secure messaging systems designed to be memorized, written down, or read aloud by a human: the Backup Recovery Key. We formalized why the corresponding baseline design provably fails even a minimal bounded-exposure notion (Claim 1), and constructed STEBR, which restores the same forward-secrecy and interactivity principles already used for *messages* to the *backup* layer: epoch-based key erasure bounds exposure (Theorem 1), threshold sharing removes any single phishable secret (Corollary 1), and an out-of-band gate makes even valid credentials insufficient on their own (Theorem 2). Every guarantee reduces to standard, well-studied assumptions — PRF/KDF security, IND-CPA security of the backup AEAD scheme, information-theoretic Shamir secrecy, and the authenticity of the already-trusted ratchet control channel — rather than to any new cryptographic primitive.

References

[1] M. Marlinspike, T. Perrin. *The X3DH Key Agreement Protocol.* Signal Technical Documentation, 2016.

[2] T. Perrin, M. Marlinspike. *The Double Ratchet Algorithm.* Signal Technical Documentation, 2016.

[3] K. Cohn-Gordon, C. Cremers, B. Dowling, L. Garratt, D. Stebila. *A Formal Security Analysis of the Signal Messaging Protocol.* IEEE EuroS&P, 2017.

[4] M. Bellare, B. Yee. *Forward-Security in Private-Key Cryptography*. CT-RSA, 2003.

[5] R. Canetti, S. Halevi, J. Katz. *A Forward-Secure Public-Key Encryption Scheme*. EUROCRYPT, 2003.

[6] Y. Dodis, J. Katz, S. Xu, M. Yung. *Key-Insulated Public Key Cryptosystems*. EUROCRYPT, 2002.

[7] A. Shamir. *How to Share a Secret*. Communications of the ACM, 22(11), 1979.

[8] A. Herzberg, S. Jarecki, H. Krawczyk, M. Yung. *Proactive Secret Sharing Or: How to Cope With Perpetual Leakage*. CRYPTO, 1995.

[9] FBI/CISA Public Service Announcement I-062626-PSA. *Russian Intelligence-Linked Actors Targeting Signal Backup Recovery Keys*. June 26, 2026.