

SECURITY DESIGN PRINCIPLES IN PRACTICE

A Systems Security Engineering Case Study of the PreVeil Architecture

Dr. Ron Ross¹

CEO, RONROSSECURE, LLC

Abstract

A trustworthy secure system cannot depend on every server, administrator, device, credential, and supporting service remaining trustworthy throughout the system life cycle. This architecture-focused case study examines the documented PreVeil end-to-end encrypted email and file-sharing architecture through the lens of NIST SP 800-160, Volume 1, Revision 1. It finds that client-side cryptographic operations, exclusion of plaintext private keys from servers, granular key wrapping, threshold Approval Groups, separation of signing and encryption, key versioning and rekeying, and protected audit trails instantiate a coherent set of design principles. Domain Separation and Least Sharing reduce concentrated trust; Distributed Privilege constrains consequential administrative action; Least Privilege and Mediated Access bound interaction; and Protective Recovery and Continuous Protection support restoration after lost devices or exposed keys. The result is not encryption added to a conventional service, but a redesign of where trust resides and what compromise can accomplish. The analysis also preserves the boundaries of the 2021 PreVeil white paper: a passive-server cryptographic threat model, visible operational metadata, consequential endpoint and Approval Group dependencies, and absence of any claim of certification or complete assurance.

1. Introduction

During the year since my retirement from NIST, I have had the opportunity to advise and work with a variety of technology-sector organizations, from early start-ups to established providers of security products and services. Throughout my career, I have advocated for a fundamental change in how we approach cybersecurity: breaking down the traditional cybersecurity silo and compliance-focused approaches and going “below the waterline” to address the protection of systems at the engineering and architectural levels.

In that spirit, I have encouraged technology providers to offer greater transparency into the security design decisions, safeguards, and development practices that contribute to trustworthy secure component products—products that customers can acquire and integrate into mission-critical systems with justified confidence. My advisory relationship with **PreVeil** provided an opportunity to put that proposition into practice by examining how the documented architecture of a commercial security product applies selected security design principles and trustworthy secure systems concepts from NIST SP 800-160, Volume 1. The more architectural and security design information available to customers, the better equipped they will be to make risk-informed decisions when acquiring and using components within complex, mission-critical systems.

¹ **Disclaimer.** The author provides cybersecurity advisory services to PreVeil. As part of those services, the author independently proposed and received approval to examine how the documented PreVeil security architecture aligns with selected systems security engineering concepts and design principles in NIST SP 800-160, Volume 1. The company gave the author full analytical and editorial independence. PreVeil was provided an opportunity to verify descriptions of its technology for factual accuracy but did not control the analytical framework, findings, conclusions, or publication decision. The analysis, interpretations, conclusions, and any errors are solely those of the author. This article is an architecture-focused case study; it is not a comprehensive evaluation, certification, or conformity assessment of the PreVeil product. PreVeil is examined as an illustrative case; the analytical approach is intended to be applicable to other technology providers and does not imply that PreVeil is the only company making meaningful systems security engineering investments.

Enterprise email and file sharing create a difficult *concentration-of-consequence* problem. Conventional designs can place protected content, decryption capability, administrative authority, and recovery power within the same service boundary, creating an architectural *path* by which compromise of a server, privileged operator, credential store, or key-management service can expose many users at once.

PreVeil begins from a different premise: centralized locations are attractive targets and may eventually be compromised by a sufficiently motivated attacker [1]. Its architecture performs encryption and decryption on user devices, keeps private decryption keys away from the server, and distributes sensitive recovery and administrative authority across multiple approvers. These choices reflect the *Engineering for Compromise (EfC)* objective: distinguish component compromise from system compromise and unacceptable loss, then use architecture to break the *causal chain* between them [2].

2. From Security Features to System Architecture

NIST SP 800-160 treats security as an *emergent property* of the system rather than a property supplied by a single product or mechanism [4]. A cryptographic algorithm may be strong while the surrounding system remains vulnerable because keys are centralized, authority is excessive, interfaces are bypassable, recovery reintroduces old secrets, or endpoints expose plaintext. Accordingly, this analysis does not ask only whether PreVeil encrypts data. It asks how the architecture allocates trust, authority, information, and failure consequences across clients, servers, administrators, approvers, storage services, and enterprise processes.

The *Trustworthy Secure Design* (TSD) framework provides the companion design lens. TSD starts with intended behavior, identifies states and conditions that can lead to asset loss, and applies the *Security Design Order of Precedence*: eliminate or reduce loss potential through design before relying on engineered features, monitoring, procedures, or training [3][4]. For sensitive collaboration, an illustrative unacceptable loss is disclosure or unauthorized alteration of protected enterprise content. Hazardous conditions include a service holding both ciphertext and usable decryption keys; one administrator being able to export all data; a removed user retaining access to future content; or a recovery mechanism recreating a single master secret.

A Loss-Driven Claim Boundary

A defensible claim boundary is narrower than “PreVeil is secure.” Based on the PreVeil white paper, the architectural claim examined here is that PreVeil reduces the ability of compromise of specified server-side or administrative elements to produce unauthorized disclosure of protected content, while preserving controlled collaboration, recovery, and enterprise administration [1]. That claim must be qualified by the architecture's stated threat assumptions, the integrity of clients and key distribution, the selection and behavior of approvers, the handling of exposed metadata, and the correctness of implementation.

Method and Standard of Assessment

The method begins with documented architectural facts and the stated threat model, identifies trust dependencies and plausible compromise conditions, maps mechanisms to selected SP 800-160 design principles, and then tests whether those mappings support bounded system-level claims. Company statements are treated as inputs to be examined, not as conclusions. A principle is credited only where a concrete design decision appears to realize its engineering intent; stronger assurance claims are reserved for cases supported by corresponding implementation, test, or operational evidence. The analysis reports strengths, qualifications, and evidence gaps using the same standard.

The decisive question is not whether a server can be compromised. It is whether compromise of that server yields the authority and keys needed to convert stored ciphertext into enterprise-wide loss.

The Architectural Trust Shift

PreVeil moves cryptographic operations and private-key control to user devices. The server retains encrypted content, public keys, encrypted key material, routing information, and operational metadata needed to deliver the service [1]. This does not eliminate trust. Rather, it changes the trust dependency. Confidentiality no longer rests on the server's ability to protect plaintext private keys it does not possess; it rests more heavily on endpoint integrity, authentic public-key distribution, correct cryptographic protocols, user sharing decisions, and the inability of insufficient approvers to reconstruct protected secrets. That distinction is central to EfC and SP 800-160. Trust is not a marketing label attached to a system element. It is a dependency created when failure or malicious behavior of that element can invalidate a system-level claim [2]. A sound architecture deliberately minimizes, distributes, and substantiates those dependencies.

3. The PreVeil Security Architecture

The white paper describes a closed client-server collaboration network in which PreVeil software performs cryptographic operations on enrolled user devices. Documents, file names, email content, attachments, subjects and snippets are encrypted before server storage. Unique content keys and a hierarchy of wrapped keys support granular sharing without giving the server plaintext decryption capability. User and device keys authenticate requests, while distinct signing and encryption key pairs separate identity assertions from confidentiality functions [1].

Table 1. PreVeil architectural mechanisms viewed as loss-control decisions [1].

Architectural Element	Security Function	Compromise Consequence Being Reduced
Client-side encryption and decryption	Keeps plaintext operations and private keys on user devices	A server-side data theft yields encrypted content rather than immediately readable enterprise data
Granular key wrapping	Encrypts content keys through user-, directory-, collection-, and block-oriented relationships	Access can be allocated and changed without one server-held master decryption key
Approval Groups	Requires an x-of-y set of authorized participants for recovery or sensitive enterprise actions	Compromise or misuse of one administrator is insufficient for the protected operation
Key versioning and rekeying	Rejects obsolete authorization for new operations and creates new key material after selected events	After rekeying, possession of a superseded key does not automatically confer access to information protected under the new key version
Signed, hash-chained logs	Protects the integrity and accountability of activity records	Unauthorized modification of audit history becomes more detectable

Approval Groups as Architectural Control

Approval Groups are the clearest example of a security design principle realized directly in the architecture. The white paper describes an x-of-y authorization model, using “Shamir Secret Sharing” for user-key recovery, that can also be composed into more complex approval structures. Encrypted shards are assigned to different approvers; the required subset must act before the secret can be reconstructed. After recovery, rekeying is required so the same recovered material cannot be reused without renewed authorization [1].

The same construct supports enterprise operations such as data export and e-discovery, changes to approval groups, and selected administrative actions. This is significant because enterprise usability requirements often defeat end-to-end security by demanding an omnipotent recovery or compliance account. PreVeil instead attempts to satisfy those requirements while retaining *distributed authority*.

Dynamic Rather Than Permanent Authority

Key versioning and event-driven rekeying make authorization dependent on the system's current cryptographic state rather than on the passage of time. The server recognizes current key versions for new actions, while older keys may be wrapped under newer keys to preserve authorized access to historical information. Security-relevant events, including account recovery, membership changes, device revocation, or changes to Approval Groups can trigger the generation and distribution of new key material [1]. From an EfC perspective, this creates a key-state containment boundary: after rekeying is completed and the new key version is enforced, possession of a superseded key should not authorize new operations or provide access to information protected solely under the new key. Cryptography does not independently determine when rekeying should occur; protection therefore depends on reliably detecting the triggering event, completing the rekeying process, and rejecting obsolete key versions.

4. Domain Separation and Least Sharing

Domain Separation requires the system to establish distinct domains for information, authority, control, and trust so that interaction occurs only through defined relationships [4]. PreVeil's most consequential separation is between the server storage and processing domain and the client cryptographic domain. Encrypted content may reside on infrastructure controlled or observed by a server-side adversary, while the private keys needed to decrypt that content remain on authorized user devices. The separation is cryptographic and functional, not merely a network boundary [1].

The design also separates user identity, device identity, signing, decryption, collections, directories, blocks, and logs through different keys and metadata structures. Separating signing from encryption is particularly important. The white paper explains that compromise of a decryption key does not automatically provide the signing capability needed to masquerade as the user, and enterprise export can use decryption-related material without copying a user's signing identity [1]. This is a practical application of the design principle of *Structural Decomposition and Composition*: functions with different consequences are allocated to different cryptographic elements.

Least Sharing as Compromise Containment

The security design principle of *Least Sharing* minimizes unnecessary common mechanisms and shared resources because sharing creates coupling and common-mode compromise paths [4]. PreVeil reduces sharing of plaintext key material by generating and using private keys on client devices and by wrapping secrets for specifically authorized recipients. The server need not share a universal decryption secret with administrators or application services. Collections apply a common sharing policy to their contents, while granular permissions—read, write, view, share, and log—constrain what authorized users can do [1].

The architecture does not eliminate sharing; collaboration is its purpose. It makes sharing explicit and cryptographically represented. This is the right interpretation of *Least Sharing*: not isolation for its own sake, but elimination of unnecessary coupling and precise mediation of necessary coupling. Each additional recipient, device, approver, or enterprise export path expands the trust and exposure surface and should therefore be treated as an *engineering* decision.

Compromise Containment Boundaries

EfC defines a *compromise containment boundary* as the architectural boundary intended to limit the information, authority, control, functionality, and influence available through a compromised system element [2]. PreVeil creates several candidate boundaries: (1) a server boundary that excludes plaintext keys; (2) per-user authorization boundaries; (3) per-device authentication boundaries; (4) collection-based sharing boundaries; and (5) multi-person approval boundaries. These boundaries are meaningful only to the extent that alternate paths cannot bypass them.

Encryption is the mechanism. Separation of ciphertext from usable authority is the architectural property.

The server boundary is therefore strongest for the passive data-exposure scenario described in the white paper. If an attacker can read server memory and storage but cannot alter key distribution or protocol execution, the attacker obtains ciphertext and selected metadata, not user plaintext [1]. A fully active server capable of substituting public keys, manipulating software delivery, or changing protocol behavior is a different adversary condition and requires additional controls and evidence.

5. Distributed Privilege, Least Privilege, and Mediated Access

PreVeil constrains high-consequence operations by distributing authority among multiple parties, limiting the privileges assigned to each entity, and mediating necessary access through authenticated cryptographic and policy-enforcement mechanisms.

Distributed Privilege

The security design principle of *Distributed Privilege* requires multiple authorized entities to act in a coordinated manner before an operation is allowed to occur [4]. PreVeil's *Approval Groups* are a direct technology realization of this principle. An x-of-y threshold separates the authority needed for account recovery or other sensitive actions among multiple participants. Defeating the protection requires compromise, coercion, or collusion across the threshold rather than compromise of a single privileged account.

This architecture changes the unit of *administrative risk*. In a conventional system, the compromise margin for a high-consequence operation may be a single administrator credential. With a properly configured Approval Group, the margin becomes the number and independence of additional approvers that must be subverted before the threshold is reached. The benefit depends on group composition. Five approvals offer little independence if all approvers use the same managed endpoint, identity provider, administrator, or communication channel. Distributed privilege must be accompanied by distributed failure dependencies.

Least Privilege

The design principle of *Least Privilege* limits each entity to the minimum authority needed for its assigned function [4]. PreVeil applies this principle in several ways. The server can store and route encrypted data without possessing general plaintext authority. A read-only recipient receives less capability than a user allowed to edit or share. A device has its own key and can be disabled independently. Data export can be structured to use decryption capabilities without transferring user signing keys. Approval thresholds constrain administrators who otherwise might act alone [1].

These mechanisms reduce the amount of adversary capability inherited after compromise. A compromised server, administrator, device, recipient, or export process does not necessarily acquire every capability held elsewhere in the system. The architecture should still be evaluated under hostile use: if every legitimate privilege of a particular device or authorized user were exercised maliciously, what protected information and future operations become reachable?

Mediated Access

Mediated Access requires all accesses to system resources to be governed by a policy-enforcing mechanism that cannot be bypassed [4]. PreVeil mediates collaboration through signed client requests, active user and device keys, wrapped content keys, collection permissions, and server checks on key status and association [1]. Cryptographic possession and server-side authorization work together: a user must obtain the appropriate wrapped secret and present acceptable authenticated requests.

Table 2. Principle-to-evidence questions for authorization architecture.

Design Principle	PreVeil Realization	Engineering Test
Distributed Privilege	Threshold <i>Approval Groups</i> for recovery and sensitive operations	Can one person, account, or device still complete the consequential action through an alternate path?
Least Privilege	Separate permissions, device keys, signing and decryption keys	What is the maximum loss reachable through each legitimate privilege if its holder is hostile?
Mediated Access	Signed requests, key status checks, wrapped keys, collection rules	Are all normal, administrative, recovery, and degraded-mode access paths covered and non-bypassable?

6. Protective Failure, Recovery, and Continuous Protection

Trustworthy secure design assumes that protective functions can fail for intentional or unintentional reasons. The engineering objective is not simply rapid restoration of service; it is restoration of a known, authorized, and sufficiently trustworthy state [3][4]. PreVeil's device revocation, account recovery, key versioning, and rekeying mechanisms provide an architectural basis for responding to lost devices, compromised keys, membership changes, and recovery events [1].

Protective Failure

The security design principle of *Protective Failure* requires failures to occur in a state that does not produce unacceptable loss [4]. In PreVeil, unavailability of an individual device or approver can deny or delay access, but the design avoids resolving that failure by placing a universal plaintext recovery key on the server. Threshold groups can preserve recoverability despite loss of some participants while still denying recovery below the required threshold. This represents an explicit *security-availability* trade: authorization should fail closed when required approvals or current keys are absent, even though collaboration may be delayed. The white paper also describes replicated collections and event logs that permit server responsibilities to move after infrastructure failure [1]. That availability design is complementary, but distinct. Replication protects service continuity; it does not substitute for the confidentiality and authorization properties of client-side keys and threshold approval.

Protective Recovery

The design principle of *Protective Recovery* restores the system to a state that is consistent with security objectives [4]. PreVeil's recovery process reconstructs a user secret only after sufficient approvers provide their shards through protected exchanges. The recovered account is then rekeyed, limiting later reuse of

the same approval event and preventing an older key from authorizing new information [1]. This is stronger than simply restoring a backup: recovery changes the *security state* so the compromise condition is not blindly carried forward.

Continuous Protection and Least Persistence

The security design principle of *Continuous Protection* requires protection to remain effective across system states and modes, including initialization, normal operation, maintenance, recovery, and transition [4]. Key versioning and event-driven rekeying help maintain this continuity as devices, users, collections, and Approval Groups change. The related principle of *Least Persistence* is reflected in limiting the post-rekey usefulness of superseded key material. Older keys may remain available, wrapped under the current key, for authorized access to historical information, but the server should reject obsolete key versions for new actions [1]. The effectiveness of this protection depends on the timely initiation and successful completion of rekeying following a security-relevant event.

Recovery is trustworthy only when it removes or bounds the condition that made recovery necessary.

Anomaly Detection and Accountability

PreVeil records operations in encrypted logs; collection log entries are described as signed and hash-chained to make modification evident, and organization-wide events support enterprise visibility [1]. This aligns with the security design principle of *Anomaly Detection* and accountability objectives in SP 800-160. Logging alone does not provide timely detection. Effective application requires defined expected behaviors, analytic rules, protected time and identity context, alerting, response authority, and evidence that gaps or truncation are detectable.

7. Compositional and Substantiated Trustworthiness

A secure cryptographic primitive does not automatically create a trustworthy secure system. The design principle of *Compositional Trustworthiness* asks whether the trustworthy properties of system elements are preserved when those elements interact [4]. For PreVeil, composition spans endpoint software, operating systems, cryptographic libraries, device storage, key-generation processes, server authorization logic, public-key lookup, synchronization, cloud storage, approval workflows, administrative consoles, recovery processes, software updates, and user actions.

The architecture has a disciplined functional decomposition: block keys protect file blocks; directory and collection keys organize access; user and device keys identify authorized entities; signing and decryption are separated; log keys protect audit information; and *Approval Groups* distribute selected authority [1]. That decomposition supports reasoning, but every interface introduces assumptions. A public key must belong to the intended recipient. A client must faithfully encrypt before upload. A recovered key must be delivered to the right device. A revoked device must be unable to regain authority through stale state. A software update must not move plaintext or secrets across the server boundary.

Evidence Required for the System Claim

The security design principle of *Substantiated Trustworthiness* requires objective evidence commensurate with the consequences of failure [4]. The PreVeil white paper is valuable architectural evidence, but it is not by itself sufficient to establish implementation correctness or current operational assurance. A complete evidence package would include:

- Traceability from unacceptable losses and security constraints to architecture, protocol requirements, implementation, and verification results.
- Protocol analysis covering confidentiality, authenticity, replay, rollback, key substitution, stale-key acceptance, threshold authorization, and recovery transitions.
- Verification that plaintext private keys and protected content do not cross prohibited server interfaces, including telemetry, logs, crash reports, backups, and support workflows.
- Independent review of cryptographic design and implementation, secure key generation and storage, randomness, library configuration, and cryptographic-agility decisions.
- Adversarial testing of clients, servers, administrative consoles, enrollment, device transfer, revocation, rekeying, export, and Approval Group workflows.
- Supply-chain and update assurance, including build provenance, code signing, dependency controls, rollback resistance, and response to compromised distribution infrastructure.
- Operational evidence for alert completeness, recovery testing, approver independence, configuration control, vulnerability remediation, and continued validity of threat assumptions.

Assurance Cases

An *assurance case* can organize this evidence as claims, arguments, assumptions, and supporting artifacts. A top-level confidentiality claim might be decomposed into client-side encryption, absence of server-held plaintext keys, authentic key distribution, correct access mediation, endpoint protection, and cryptographic strength. Each subclaim would identify the evidence and residual uncertainty. This prevents an important architectural design consideration—servers do not hold private decryption keys—from being expanded into an unsupported universal claim about every attack scenario.

The assurance case should be a living engineering artifact. Changes to cryptographic algorithms, identity and enrollment, browser or mobile clients, storage providers, metadata processing, recovery policy, export, or enterprise administration can invalidate earlier assumptions even if the high-level architecture appears unchanged.

8. Engineering for Compromise Applied to PreVeil

EfC asks engineers to assume a selected system element is hostile and determine the maximum loss reachable through its legitimate and illegitimate influence [2]. Applied to PreVeil, this approach reveals both meaningful containment and residual trust dependencies.

Table 3. Illustrative compromise-propagation analysis based on the stated architecture [1][2].

Assumed Compromised System Element	Architectural Containment	Residual Concern / Next Barrier
Server storage or passive server observer	Content and many descriptive fields remain encrypted; private user keys are not stored in plaintext on the server	Traffic, relationship, timing, count, and size metadata remain; availability and replay/rollback behavior require analysis
One enterprise administrator	Threshold Approval Groups can prevent unilateral protected operations	Common identity, endpoint, or organizational dependencies may reduce real independence

Assumed Compromised System Element	Architectural Containment	Residual Concern / Next Barrier
One user device	Access is limited to information and authority available to that user or device; another active device may disable it	Plaintext and user keys on an authorized endpoint are high-value targets; revocation latency and offline data matter
Old user key	After rekeying, current key versions govern new actions, and superseded keys cannot access material protected solely under the new keys	Previously decrypted or copied data cannot be cryptographically recalled
Authorized recipient	Collection permissions and key wrapping bound access to shared material	An authorized recipient can disclose content already granted; user intent and endpoint controls remain essential
Insufficient approvers	Below-threshold participants cannot reconstruct the protected secret	Collusion at threshold, coerced approval, or correlated compromise defeats the boundary

Compromise Margin

Compromise margin is the difference between the compromise already assumed and the additional independent failures needed to reach unacceptable loss [2]. *Approval Groups* make that concept measurable. If three of five independent approvers are required, compromise of one results in a nominal margin of two additional approvers. But numerical threshold alone is misleading unless independence is real. Shared credentials, managed devices, administrators, locations, recovery channels, and incentives can collapse multiple nominal barriers into one common-mode failure.

Preventing Propagation

Key wrapping, collection boundaries, differentiated permissions, device-specific authentication, and separated signing and decryption capabilities limit how authority propagates. Once completed and enforced, rekeying prevents an exposed or superseded key from authorizing new operations or accessing material protected solely under the new key version. These are EfC-compatible barriers because they constrain what an attacker inherits after compromise. The analysis should be performed recursively: compromise a client, server, approver, administrator, software update service, cryptographic dependency, and storage provider in turn; then examine combinations that could defeat the intended boundaries.

What Architecture Cannot Recover

No architecture can revoke knowledge already disclosed to an authorized endpoint or recipient. Rekeying can deny future access, not erase plaintext copied earlier. Likewise, end-to-end encryption protects content but cannot remove all metadata required for routing and synchronization. These are not reasons to reject the architecture. They define where complementary endpoint security, data minimization, insider-risk management, user education, retention policy, and operational monitoring remain necessary.

9. Applying the Concepts Across the System Life Cycle

SP 800-160 applies systems security engineering across the full life cycle. The PreVeil architecture should therefore be treated as a maintained set of security claims and constraints, not a one-time cryptographic design [4]. The architecture illustrates the concept of *Security Design Order of Precedence*. First, it reduces the loss potential through design by excluding plaintext private keys from servers and avoiding unilateral administrative authority. Second, it uses engineered features—including cryptography, wrapping, signing, permissions, revocation, and rekeying—to enforce the design. Third, protected logs and enterprise event

streams support detection and response. Finally, policies and trained approvers govern sensitive human decisions [3][4]. The ordering matters: procedures do not have to compensate for a server architecture that inherently holds universal decryption power.

The PreVeil architecture also supports change control as a key system security engineering concept. Security-relevant change includes more than cryptographic replacement. A new client platform, search function, telemetry pipeline, AI assistant, customer-support feature, export capability, administrator role, storage integration, or recovery option may create a new plaintext path, metadata disclosure, shared dependency, or bypass. Traceability should determine which claims, requirements, threat assumptions, tests, and authorization decisions must be revisited. The central architectural invariants should be explicit enough that engineers can recognize when a seemingly local product change violates them.

Table 4. Life-cycle application of trustworthy secure development concepts [1][4].

System Life Cycle Activity	Application to the PreVeil Solution
Mission analysis and stakeholder needs	Define content, metadata, availability, compliance, usability, recovery, export, and insider-risk needs; identify unacceptable losses and claim boundaries.
System requirements	Specify prohibited plaintext flows, key-location constraints, threshold rules, revocation and rekey behavior, metadata limits, logging integrity, and recovery outcomes.
Architecture and design	Allocate cryptography, keys, identity, policy enforcement, approval authority, storage, logs, and recovery to domains; analyze bypass and common-mode dependencies.
Implementation and integration	Preserve architectural constraints in client code, server APIs, crypto libraries, device platforms, administrative tools, cloud services, and update mechanisms.
Verification and validation	Test protocol properties, negative paths, threshold enforcement, stale-key rejection, metadata exposure, revocation, export, recovery, failover, and stakeholder mission outcomes.
Operation and maintenance	Monitor anomalies; manage approvers and devices; re-evaluate threats; patch dependencies; conduct recovery exercises; maintain the assurance case after every security-significant change.
Disposal	Revoke authorization, retire devices and keys, address retained ciphertext and backups, sanitize local plaintext, and preserve only required audit evidence.

10. Findings and Conclusion

Principal Findings

The supplied architectural white paper from PreVeil describes a technology solution that meaningfully applies several NIST SP 800-160 security design principles. *Domain Separation* is realized by dividing client cryptographic authority from server storage and processing. *Least Sharing* is reflected in recipient-specific wrapped keys rather than a commonly held plaintext master secret. *Distributed Privilege* is directly implemented through threshold *Approval Groups*. *Least Privilege* and *Mediated Access* appear in differentiated permissions, device-specific keys, signed requests, and server checks. *Protective Recovery*, *Continuous Protection*, and *Least Persistence* are supported by recovery, versioning, revocation, and rekeying. *Compositional and Substantiated Trustworthiness* define the engineering work needed to demonstrate that these properties survive real implementation and operation.

Viewed through the lens of *Engineering for Compromise*, these mechanisms create barriers between the compromise of a server, administrator, device, or old key and enterprise-wide loss. Viewed through the lens of *Trustworthy Secure Design*, they reduce loss potential by changing system structure before relying on monitoring and procedure. The architecture's deepest idea is therefore the deliberate placement of *trust*: the service should perform useful, centralized coordination without centralizing all of the secrets and authority that make catastrophic compromise possible.

PreVeil's architecture demonstrates a central systems security engineering lesson: reduce the consequences of compromise by refusing to concentrate the authority that makes compromise catastrophic.

Conclusion

PreVeil is best understood as an architectural security solution rather than an encryption feature. Its client-side cryptography, distributed approval, granular key relationships, separated capabilities, dynamic keys, and protected accountability mechanisms collectively reshape the system's trust structure. That structure aligns strongly with the loss-driven, life-cycle-oriented, and evidence-based approach of NIST SP 800-160 and with the practical objectives of *Engineering for Compromise* and *Trustworthy Secure Design*.

The remaining obligation is assurance: demonstrate that the implemented and operated system preserves those architectural invariants across every client, interface, state, mode, recovery path, administrative function, and change. Subject to the stated evidence and scope limitations, PreVeil provides a useful real-world case study of a company making security design decisions that materially reflect SSE principles. When the corresponding evidence is explicit and continuously maintained, such cases show how security design principles can move from publication language into operational technology solutions—and how rigorous analysis can recognize meaningful engineering without becoming product promotion.

An Industry Call for Architectural Transparency

This case study is also intended to encourage technology providers that have invested in trustworthy secure product design and development to conduct and publish comparable architectural analyses. Rather than relying principally on generalized security claims, providers should identify the specific security design principles they have applied, explain how those principles are realized in system architecture, disclose the threat assumptions and trust dependencies that bound their claims, and describe the evidence used to substantiate them.

Such analyses would substantially improve transparency into the interacting system elements needed to build mission-resilient systems—including hardware, software, cryptographic components, identities, administrative authorities, communications, supply-chain dependencies, enabling systems, recovery mechanisms, people, and procedures. They would help customers distinguish between products that merely incorporate security features and systems deliberately engineered to control the consequences of adversity and compromise.

Architectural transparency does not require disclosure of intellectual property, proprietary implementation details or information that would assist an adversary. It requires sufficient clarity for customers, integrators, assessors, and authorizing officials to understand where trust is placed, how privilege is constrained, what boundaries limit compromise, how recovery preserves security intent, and what evidence supports the resulting claims.

PreVeil's willingness to permit this architecture-focused analysis provides an opportunity to lead by example. The larger objective is not to confer a special designation on one company, but to encourage an industry practice in which technology providers explain and substantiate—in concrete engineering terms—how their products contribute to trustworthy secure and mission-resilient systems.

References

- [1] PreVeil, *PreVeil Security and Design: A Description of the PreVeil System Architecture*, Version 2.0, 2021.
- [2] Ross, R. “Engineering for Compromise: Designing Systems That Remain Trustworthy, Secure, and Resilient Under Adversity.” RONROSSECURE, LLC, LinkedIn, August 2026.
- [3] Ross, R. “Trustworthy Secure Design: Engineering Security from the Inside Out.” RONROSSECURE, LLC, LinkedIn, August 2026.
- [4] Ross, R., Winstead, M., and McEvelley, M. *Engineering Trustworthy Secure Systems*. NIST Special Publication 800-160, Volume 1, Revision 1. National Institute of Standards and Technology, November 2022. <https://doi.org/10.6028/NIST.SP.800-160v1r1>