



ZEUS X-Trust X1



Beyond authentication: an autonomous, self-healing and self-defending post-quantum security class. The valid secret exists only while the system runs — nothing to extract, nothing to attack offline.

Trauth Research · Distribution via GeoFold AI (planned, not yet incorporated)

Stefan Trauth — Founder & Owner

www.zeusxtrust.com

ZEUS X-TRUST — PRODUCT

live.zeusxtrust.com

LIVE TEST ACCESS — X1 PROTOTYPE

www.trauth-research.com

TRAUTH RESEARCH — COMPANY

Trauth Research · info@trauth-research.com · ORCID 0009-0003-9852-9788 · GeoFold AI (planned) ·

ZEUS X-Trust X1

The secret is never stored. It exists only while the system runs — nothing to extract, nothing to attack offline. X1 detects attacks from its own live-state response and heals itself bit-exactly.



A 6-character password carries what 20+ characters carry today

~725,000


ATTACK ATTEMPTS

0

BREACHES

30

ATTACK CLASSES



Validation from X-Trust v3.11 to X1 — ~140 M measured iterations, including the X1 corpus (~25 M; 13,500 attacks, 20 classes), the Expanded Validation (45 M forwards; 450,000 attack attempts) and Pre-X2 Testing (60 M forwards; 249,000 attacks). Zero accepted attacks; legitimate access 216/216, also under lockout.

2^{256} effective work factor · no stored static key · no offline verification object · bit-exact self-healing after attack

Nothing to brute-force, nothing to steal



No stored key

Access = a password-controlled positional map × the live amplitude profile × the enrolled neural instance. There is no hash and no static key to lift.




Live-state binding

The valid state exists only inside the running instance. Reproducing it means reproducing the whole substrate — not guessing a string.



Fail-closed

Any restart or tamper destroys the seal. The volatility of the live state is itself the tamper-evidence.



Authentication decision logic — **ACCEPT** only when the positional map and the live amplitude profile both match the enrolled instance.

~725,000

ATTACKS

0

BREACH

30

CLASSES

2²⁵⁶

TARGET*

*6-char password, effective security modelled at 2²⁵⁶ class — now backed by ~140 million measured iterations from X-Trust v3.11 to X1.

ROADMAP

X1 live (autonomous ACCEPT/DENY, source-bound delay, staged key regeneration) → **Product release X1.0** → **Next: external black-box validation of a fixed, hashed X1 instance**

The invariant is measured, not assumed

Attacks and legitimate inputs form two separate invariant relations. No attack class ever reproduces the legitimate signature.

Invariant response law

Across 20 attack classes and 5 blocks the attack invariant reproduced at $V \approx 4.112115$ with sub-ppm cross-block spread — while raw amplitudes moved by more than an order of magnitude between substrates.

Zero accepted attacks, at scale

V6EXT: 0 of 150,000 forwarded attacks accepted across 27 classes; the seal of the true password reproduced exactly (distance 0.0) in all 15 blocks, no drift over 5 M forwards, bit-identical on re-run. V7: 0 unauthorized releases on all three verification paths (0/103,402 · 0/567 · 0/191,198).

Legitimate access, under every condition

V7: 216/216 legitimate releases — also during lockout, after restart and after resume; lockout mechanics as specified in all 30 configurations (10 failures → lock, 11th blocked without escalation, auto-reset); binding tests B1–B6 fail-closed in all 30 configurations; 900-s production cycle passed.

0 / ~725,000

ATTACKS ACCEPTED

216 / 216

LEGITIMATE RELEASES

30 / 30

CONFIGURATIONS FAIL-CLOSED

2²⁵⁶

EFFECTIVE WORK FACTOR

The substrate is the security boundary. Structurally different attacks collapse onto one invariant response law at 99.9% explained variance, while legitimate input restores the valid state bit-exactly — on top of an effective work factor modelled in the 2²⁵⁶ class, the order of the 256-bit security level of SHA-512.

Beyond Authentication — DOI 10.5281/zenodo.21967777 · Expanded Validation — DOI 10.5281/zenodo.22736928 · Pre-X2 — DOI 10.5281/zenodo.22926782 · Raw data ~100 GB, DOI 10.5281/zenodo.22737465 (restricted access) · X1 corpus ~25 M iterations, 62 runs · V6EXT 15 M · V7 30 M

Volatile by design — recoverable by choice

Live-state loss or protected recovery is a deployment parameter — the security class stays the same in both modes.

Volatile: loss as a feature

Launch-code-class deployments run pure live-state: any restart or tamper destroys the seal, and loss of the instance is the intended fail-closed outcome. Nothing persists — so nothing can be stolen.

Recovery without an oracle

Enterprise profiles keep a protected baseline and a password-bound amplitude reference, stored separately or offline. A reloaded instance must reproduce the live-state relation before it is accepted.

Rebirth after attack

Any incident that isolates the system or exfiltrates weights is treated like a disclosed hash — even though no extraction path from ZEUS X-Trust is known. After recovery, a fresh initialisation is mandatory.

12

COUPLED RELATIONS

×100

AMPLITUDE SHIFT

0

KNOWN EXTRACTION PATHS

2²⁵⁶

EFFECTIVE WORK FACTOR

A stolen instance dies with the re-roll. Re-initialisation shifts every amplitude by up to two orders of magnitude, devaluing anything an attacker observed or exfiltrated. A leaked hash stays attackable forever — a re-rolled instance leaves no known attack surface behind, on top of an effective work factor modelled in the 2²⁵⁶ class, the order of the 256-bit security level of SHA-512.

Deployment note — a 6-character password addresses 12 coupled position–amplitude relations; to date, no inversion from recovery references to the password and no hash-like offline search path, classical or quantum, is known or has been demonstrated by third parties. Volatile vs. recoverable is configuration, not architecture.