VEREYE

FACT SHEET · 2026

An invisible seal for what's real.

Vereye writes a unique, invisible key into an image. Any copy, even cropped, screenshotted, compressed or edited, can later be checked from the copy alone: Vereye reads the key, names the registered owner, and shows the areas where the copy no longer matches what was registered. No machine learning is involved, and a false exact match and a wrong owner both have computed bounds.

$1 \text{ in } 2^{256}$

False exact match: every pixel checked with SHA-512

$< 10^{-12}$

Wrong owner per read, for 100 billion registered images

98.5%

Edits found from the copy alone (our tamper set), no untouched copy flagged

0

False owners on 46,467 attacked copies of unmarked images

WHAT IT UNLOCKS

- News that can prove a photo is an exact match to what was sealed
- Photographers and artists proving a work is theirs
- Real and AI-generated content told apart, once both are sealed
- Leak tracing with a unique key per copy
- Evidence for courts, insurers and platforms

PUBLIC BENCHMARKS

Benchmark	Vereye	For comparison
SynthID-Image (Google DeepMind), 512 px, 300 new photos	**99.7%**	99.96% SynthID, Google's own images
Same, most severe	**95.8%**	98.84% SynthID, Google's own images
Watermark Anything (Meta), exact key, COCO / 2K	**98.7% / 99.05%**	70.3% / 71.3% WAM (Meta)
False alarms, same protocol	**0**	0.9–2.7% WAM and TrustMark

Google's closed SynthID is measured by Google on its own images, on detecting a watermark; Vereye is scored on the stricter task of naming the exact owner.

Patent pending

Adam@vereye.ai · vereye.ai