Superseded edition

Report 001

PUBLISHED AUGUST 9, 2026 / REPLACED AUGUST 11, 2026

This PDF preserves the exact August 9 edition for the public corrections record. It is not the current Report 001. Claims, product status, version references, and service positioning reflect the edition as published on August 9, 2026.

CURRENT REPORT

https://ssx360.com/research/001

CORRECTION RECORD

https://ssx360.com/corrections

ARCHIVED FROM SOURCE COMMIT A7F62CBD7F42643533C631378205C824D5964689

REPORT 001 · AUGUST 9, 2026 · PERMANENT ARCHIVE

The Cyber Risk Institute publishes two separate frameworks with two separate counts, and the EU AI Act transparency duty that took effect on August 2, 2026 was not deferred alongside the high-risk rules.

Every count below carries the page it came from. Open the sources and check them against what your own readiness pack claims, because that comparison is the whole reason this sits on a permanent URL.

Two products from one institute

The Cyber Risk Institute maintains the CRI Profile and the FS AI RMF. They cover different ground on different timelines, and they carry different counts. A readiness claim that cites one number for "the CRI framework" has already lost the distinction, and an assessor reading that claim has to ask which product it meant.

CRI Profile v2.2

Cybersecurity risk management for financial services, aligned to NIST CSF v2. It harmonises 3,500+ regulatory expectations into **318 diagnostic statements**, and carries about 40 mappings to regulatory references and industry standards. Source: CRI Profile Overview.

FS AI RMF

A separate AI risk framework from CRI and the Financial Services Sector Coordinating Council (February 2026), structurally aligned to the NIST AI RMF and built with more than 100 financial institutions. It defines **230 control objectives** across four deliverables, organised by AI adoption stage. Source: CRI FS AI RMF, announced by the US Treasury on February 19, 2026.

Where the 318 goes wrong

CRI words it precisely once you read the source. The Profile Overview says the framework harmonises regulatory expectations into "318 control objectives, called diagnostic statements". One set of 318 items carries two names.

That phrasing produces a predictable double count downstream. A summary reads the sentence as two populations and reports control objectives alongside a separate tally of diagnostic statements, which counts one framework twice and inflates the apparent scope of an assessment. Take any vendor guide or internal readiness deck you rely on, find its CRI number, and hold it against the Overview wording. If the deck lists both populations with different totals, the deck is describing a framework that does not exist.

A second reading error sits next to it, and an earlier draft of this report made it. The 318 is the whole set. CRI's Impact Questionnaire determines which of those statements apply to a given institution by impact tier, so quoting the figure as a Tier 1 total states the wrong thing about scoping. The number to put in front of a board is the count that survived tiering for your institution, alongside the 318 it came from.

Article 50 applies now

The EU AI Act transparency obligations in Article 50 took effect on **August 2, 2026** and bind providers and deployers of interactive and generative AI systems. Systems already on the market before that date have until December 2, 2026 on the Article 50(2) machine-readable marking duty, and that transition covers marking alone.

The Digital Omnibus, in force July 27, 2026, moved the Annex III high-risk obligations to December 2, 2027 and Annex I product-embedded AI to August 2, 2028. Article 12 record-keeping rides those high-risk dates, so it stays a readiness exercise for most teams. Article 50 kept its original date, and the deferral headline travels further than that detail: a team reads that the AI Act slipped and files transparency alongside the articles that actually moved. Source: Commission FAQ on Article 50 transparency obligations.

What Article 50 asks of a provider is narrow. Disclose that a person is dealing with an AI system where that would not be obvious, and mark synthetic audio, image, video or text so a machine can detect it, as far as that is technically feasible. It obliges no development team to sign its commits, and anyone selling it as a commit-provenance deadline is describing a duty the Act does not contain. The Code of Practice on marking AI-generated content, published in June 2026, is where signed metadata enters the picture: it expects marking in several layers and names digitally signed metadata as one of them, beside imperceptible watermarking.

Why a count decides whether evidence is complete

An assessment maps collected evidence to control language, and the map is only as good as the control set going into it. Miscount that set and the resulting gaps are invisible, because a row nobody expected cannot show up as missing. Detecting gaps in an evidence trail starts with agreeing the denominator, which is why we open an engagement by reconciling the framework counts a client is already reporting rather than accepting them.

What ships today

The open protocol is Matrix Scroll, free permanently. The shipping pin is `matrixscroll==0.6.2` on PyPI and GitHub. Ed25519 signing and offline verification ship in that release line. SE050 M1 signing and the Pico 2 W display bring-up are bench-validated prototypes for evaluation, with no general availability and no roadmap commitment, and live SE050 signing on the display bring-up firmware is fail-closed.

Scope of our assessments

No single accreditation scheme covers logistics, defense, finance and climate intelligence together. Ours is an independent assessment, not a certification against a published standard.

We do not sell a compliance automation platform, a monitoring SaaS, or a hardware product, and we do not intend to.

We hold no trading positions in the markets we assess, and we do not operate a root of trust for anyone. We audit and sign the evidence. We do not produce the underlying keys for a client's systems.

Post-quantum signing in Matrix Scroll implements the ML-DSA and SLH-DSA algorithms specified in FIPS 204 and FIPS 205, through liboqs. That is an algorithm implementation, not a CMVP-validated cryptographic module, and we will not describe it as FIPS validated. liboqs itself states that it should not be relied on in production or to protect sensitive data, which is a limit we repeat rather than bury.

Working with defense and government suppliers is not the same as holding an accreditation. We hold no facility clearance, no CMMC certification and no ITAR registration, and we say so before anyone asks.

Matrix Scroll is free and stays free. Gating the protocol would remove the reason anyone trusts the audit.

Independence

We do not sell software to the organisations we assess. Any tool that would create that conflict lives outside this firm or does not get built.

We do not take an engagement that compromises our position as an independent third party, at any price.

Where our research scores or ranks an organisation, any commercial relationship with that organisation is disclosed alongside the score.

We publish corrections. A finding we got wrong is amended in public, where everyone who read the original can see the change.

Fees are fixed once scoped. A mid-engagement renegotiation is a negotiable finding.

Corrections

Disagree with a count here? Write to mission@ssx360.com and we amend the page in place and say what changed, on this URL, where everyone who read the original can see it.

Reconciling framework counts for a vendor review or a diligence timeline now? Request a call or read the engagements.

Where we name a scheme such as PCI DSS, SOC 2 or the EU AI Act, we are describing evidence mapping, not a certification claim. Certification under any scheme named here comes from that scheme's own accredited assessors.

PERMANENT LINK · HTTPS://SSX360.COM/RESEARCH/001 · MATRIXSCROLL 0.6.2

© 2026 SSX360. Security Services X-Ray. Risk prevention and security for you and your company's peace of mind.