


Cryptography that's *measurably ahead.*

A working post-quantum cryptographic stack — already **thousands of times faster** than the published prior best at strictly higher security — sitting on top of a new mathematical framework with no prior name in the literature. The cryptography is shipped. The framework is partially proved. The applications below follow from both.

| | |
|--------|----------------------------------|
| TIER A | Shipped cryptography |
| TIER B | Proved mathematics |
| TIER C | Enabled applications |
| TIER D | Certified adversarial robustness |
| TIER E | Frontier (explicitly disclaimed) |



What's already running, with measurable advantages.

Each item below is in working code, with benchmarks against the published state of the art. These are not roadmap items. They are the moat. The numbers are real and verifiable.

Practical post-quantum program obfuscation

[SHIPPED]

A working virtual machine that runs **obfuscated programs at sub-second timing** while producing a small cryptographic proof of correct execution. Programs can be shipped to customer infrastructure or untrusted hardware without revealing what they do, with mathematical guarantees. Obfuscation has been theoretically interesting for two decades and practically impossible — until now.

BENCHMARK

≈ 8,400× faster than the published prior best system, at strictly higher security.
≈ 0.6 second evaluation per call · ≈ 10KB cryptographic execution proof.

Sub-700-byte cryptographic proofs, faster than industry standard

[SHIPPED]

A proof-compression layer that produces zero-knowledge cryptographic proofs **under 700 bytes** that verify faster than the established benchmark. Smaller blockchain transactions, smaller credential checks, smaller everything-that-needs-a-zero-knowledge-proof. Direct relevance to verifiable compute, on-chain identity, and any system bottlenecked on proof size or verification cost.

BENCHMARK

637-byte median proof size · faster verification than the industry standard.

Exact, deterministic alternative to “approximate-and-round”

[SHIPPED]

A cryptographic primitive that **removes a class of failure modes** that has historically caused real-world cryptographic breaks. Where lattice-style cryptography has relied on bounded statistical noise, ours uses exact algebraic structure. No rounding, no heuristic, no probabilistic argument.

WHY IT MATTERS

Removes the most common attack surface in modern post-quantum cryptography.

Constant-memory streaming proof system

[SHIPPED]

Most cryptographic proof systems require memory proportional to the size of the computation being proved. Ours doesn't. **Constant memory** for proofs of arbitrary-size computations. Direct implication: cryptographic proof generation works on commodity hardware, mobile devices, embedded systems, and edge nodes — not just datacenter racks.

WHY IT MATTERS

Unlocks proof-generation on hardware that previously couldn't host it.

A complete post-quantum verifiable VM

[SHIPPED]

A complete software stack — separately-engineered Rust modules — that lets developers build applications using all of the above without needing to be cryptographers. Production tooling, not a research demo. Designed for integration into existing engineering workflows.

STATUS

31-module Rust workspace · post-quantum proving targets · operational tooling.

Bit-level cryptographic risk accounting

[SHIPPED]

Replaces the standard “asymptotic security argument” with **hard numbers**. For the first time, an engineer can budget cryptographic risk the way they budget memory or bandwidth — exact bits, with a calculation that survives expert review. This is what underwrites the security claim of every other item on this page.

WHY IT MATTERS

Audit-grade quantitative cryptographic accounting. Strictly stronger than asymptotic claims.

Quantitative side-channel leakage analyzer

[SHIPPED]

A side-channel analysis tool that quantifies leakage from physical observation channels (clocks, power traces, timing variations, error-retry counts) at the granularity of **bits per operation**. Tested on a NIST-standardized post-quantum signature scheme; produced specific bit-exposure numbers per accepted signature for each parameter set.

STATUS

Already produced concrete vulnerability disclosures for production cryptographic schemes.

A new mathematical framework with no prior literature.

The cryptography on the previous page is partially built on a new class of mathematical framework. Each item below is a mathematical theorem with no prior published name. They are durable foundational IP — the kind of asset that survives competitive replication because the underlying machinery doesn't yet exist outside this work.

A new class of mathematical analysis

[PROVED]

Classical analysis tracks signals. This framework tracks **signals and the residual obstruction** the operation couldn't remove — as independent first-class objects, each with its own algebra. Two invariants where there was one. Both are computable. The field has no prior published name because nobody had unified the components into one framework.

STATUS

Foundational theorem proved · finite specializations proved · framework name reserved.

Topological data analysis with an algebraic axis

[PROVED]

Standard topological data analysis sees holes. Our extension **distinguishes exploitable holes from microstructure noise** using an algebraic obstruction axis rather than just geometric scale. With a proved stability theorem. Direct relevance: shape comparison in biology, materials science, fluid topology, image analysis — anywhere the data has algebraic structure that pure geometry misses.

STATUS

Stability theorem proved · awaiting library packaging for industrial deployment.

A counting structure with zero free parameters

[PROVED]

When you measure a system in this framework, every measurement is determined by a **single observed quantity**. No fitting, no choice of regularizer, no convention. Most counting frameworks have at least one parameter to tune. This one doesn't. **Every structural rule the framework relies on — how measurements decompose, compose, and combine — has now been proved end to end.** Mathematically unusual; cryptographically valuable because it removes a class of "configuration mistake" failure modes from any system that uses it.

STATUS

Zero-parameter property proved · all foundational rules of the framework verified · numerical invariants to machine precision.

Finite-to-infinite mathematical bridge with explicit, verifiable conditions

[PROVED]

Solves a **known structural problem** in arithmetic geometry that mainstream tools handle ad hoc. We give explicit conditions for when finite local data can be assembled into infinite global data, with verifiable hypotheses. One of the building blocks for the cryptographic side, and a paper-grade contribution in its own right.

STATUS

Conditional theorem proved · finite chart-level closure proved.

Self-improving refinement with guaranteed contraction

[PROVED]

Most iterative refinement schemes hope for convergence. This one **either contracts at a stated rate or stops with a precise diagnostic**. No silent failure, no false confidence. Direct application to numerical solvers, machine-learning training loops, adaptive mesh refinement, and any iterative scheme where you need to know whether progress is real.

STATUS

Contraction theorem proved · fail-closed protocol formalized.

Coupled multi-norm budget framework

[PROVED]

A transformation is accepted only when **every measurement axis improves simultaneously**. Replaces single-number quality scores with a coupled ledger; mathematically stronger than any single-norm criterion. Includes an explicit guard against signal-hiding — the failure mode where a transformation looks good on one metric while quietly destroying another.

STATUS

Five-channel coupled criterion proved · entropy-hiding flat-abort guard formalized.

What this combination makes possible.

The combination of practical post-quantum cryptography plus the mathematical framework enables a long list of applications that don't currently exist as products. Each item below is buildable with engineering effort. Several are direct extensions of what's already shipped.

Sealed AI agents

[ENABLED]

Ship agent code — prompts, tool wiring, fine-tunes — into customer infrastructure with mathematical guarantees that the customer cannot extract the agent's logic. Solves the **"I want to use your AI but I can't share my data, and you can't share your model"** problem at the cryptographic level. Every AI lab and serious agent company currently has this problem and is solving it badly with hardware enclaves that keep getting broken.

MARKET

AI infrastructure · enterprise AI deployment · regulated-industry AI integration.

Provably watermarked AI models

[ENABLED]

Per-output proofs that a token came from a specific model, plus model fingerprints that survive distillation attempts. Solves "is this AI-generated?" definitively. Regulators (EU AI Act, U.S. executive orders) will mandate this within years; the technical capability to comply doesn't exist anywhere else today.

MARKET

AI provenance · regulatory compliance · anti-distillation IP protection.

Multimodal deepfake attribution

[ENABLED]

The same per-output structural fingerprinting **generalizes verbatim to image, audio, and video** generative models. The fingerprint sits in the structural residue of the generated content, not in surface statistics, so it survives compression, format conversion, common post-processing, and downstream distillation. Current deepfake detectors break under simple post-processing; ours doesn't. Solves "is this AI-generated, and from which model?" across the largest modality of the deepfake problem — a multi-billion-dollar regulatory and trust-and-safety market in 2026.

MARKET

Content moderation · news and election integrity · regulatory compliance · anti-fraud.

Tamper-evident sensor capture

[ENABLED]

Recording devices — police body cameras, dashcams, broadcast cameras, drones, scientific instruments — that capture frames with cryptographic proofs the data conforms to **declared physical state-transition laws**: temporal continuity, optical consistency, sensor-specific dynamics. Post-capture alterations that violate those laws fail verification. Generative inpainting introduces residues incompatible with the declared physics and is detectable. Combines the streaming proof primitive with physical-law-conformant defect analysis. Direct application to courtroom evidence, broadcast journalism, body-cam manufacturer integrity, insurance claim verification, and autonomous-vehicle accident reconstruction.

MARKET

Law enforcement · broadcast media · insurance · courtroom evidence · body-cam vendors · autonomous vehicles.

Encrypted algorithm marketplace

[ENABLED]

Trading strategies, machine learning models, drug-discovery models can be sold where buyers can run them but cannot extract them. **Mathematical, not hardware-trusted.** Currently impossible without trusted execution enclaves, which keep getting broken. Quantitative finance, pharmaceutical R&D, and competitive ML labs are direct buyers.

MARKET

Quant finance · pharma · ML model licensing.

Cryptographic DRM that actually works

[ENABLED]

Software and media licensing where the decryption logic *is* the cryptography rather than embedded code that can be reverse-engineered. The DRM industry has wanted this for 25 years; current schemes (Widevine, FairPlay) get cracked because the decryption logic is in the binary. **Ours puts it in the math.**

MARKET

Game studios · streaming platforms · enterprise software licensing.

Defense-grade autonomous-system protection

[ENABLED]

Captured drones currently leak everything when reverse-engineered. Our framework lets logic be embedded in hardware in a way that **cannot be extracted even with full physical access.** Direct strategic relevance to defense primes building autonomous systems where adversarial reverse-engineering is a known failure mode.

MARKET

Defense primes · classified compute · autonomous-system manufacturers.

Confidential smart contracts that hide logic, not just inputs

[ENABLED]

Existing privacy-preserving blockchains hide who-did-what; ours can additionally **hide what-logic-was-executed.** New product category in DeFi, RWA, private market making, sealed-bid auctions, MEV-immune trading.

MARKET

DeFi infrastructure · privacy-preserving L2s · institutional on-chain finance.

Compliance as cryptographic proof

[ENABLED]

A program proves it respects HIPAA / SOC 2 / PCI-DSS / EU AI Act policies **cryptographically,** without anyone needing to read the program. Multi-trillion-dollar regulatory market. Long sales cycles, but every Fortune 500 has the budget for provable compliance.

MARKET

Healthcare · finance · enterprise regtech · cross-border data flow.

Quantitative side-channel certificates

[ENABLED]

A masking scheme can be certified to **leak no more than X bits per operation**, with mathematical backing. Strictly stronger than current pass/fail t -tests and mutual-information attacks. Already operational; needs industrial packaging.

MARKET

Hardware security modules · secure enclave vendors · cryptographic engineering audit.

Quantum error-correcting codes with guaranteed contraction

[ENABLED]

A code synthesis tool that **emits new quantum error-correcting codes alongside the proof** of how fast errors converge. Tightens magic-state distillation schedules in fault-tolerant quantum computing. Direct relevance to quantum hardware companies racing toward fault tolerance.

MARKET

Quantum hardware vendors · quantum software stacks · fault-tolerance research.

Numerical solvers with guaranteed conservation

[ENABLED]

Climate, computational fluid dynamics, molecular dynamics, plasma simulation — all currently estimate conservation of mass, energy, momentum, and charge. Ours can **certify it**. Solver-grade adaptive mesh refinement, with a mathematical backstop on whether the solver is actually conserving what it says it is.

MARKET

Computational science · weather/climate modeling · industrial simulation.

Geometrically-optimal AI training

[ENABLED]

Defect-aware geometric analysis of the gradient-descent landscape gives a **precise diagnostic** for whether each training step is real progress or noise. Replaces the empirical “did the loss go down?” check with a structural “did the contraction rate match the theoretical bound for this noise structure?” Applies to every gradient-based optimizer in production AI training (SGD, Adam, AdamW, momentum and learning-rate schedules) and to every iterative numerical solver beyond AI (Newton’s method, conjugate gradient, fixed-point iteration). The deeper result is the **first exact statement of how good any iterative solver can possibly be against a given noise structure** — a Shannon-style fundamental limit for iterative refinement, analogous to channel capacity for communication. The tradeoff between convergence rate and stability has been intuited for 350 years (Newton has “difficult cases”; trust regions are heuristic patches; Adam’s hyperparameters are empirical) but never stated in exact form. The framework gives the exact form. Compute savings of even $1.5\times$ at frontier-model scale: billions of dollars across the industry.

MARKET

Frontier AI labs · cloud GPU providers · ML training infrastructure · scientific computing · numerical optimization libraries.

Codecs for scientific data with conservation certificates

[ENABLED]

A compressed climate dataset ships with a **cryptographic proof** that mass / energy / momentum are preserved to a stated tolerance. Direct relevance to climate-science consortia, NASA, NOAA, weather services, and any organization moving large scientific datasets where conservation guarantees are scientifically required.

MARKET

Climate science · earth observation · scientific data infrastructure.

Pollution-free eigenvalue computation

[ENABLED]

Spurious eigenvalues are a known failure mode in computational chemistry, materials science, and continuum mechanics. Our framework eliminates them with a **provable count of true vs spurious negatives**. Eigensolver vendors (COMSOL, ANSYS) have wanted this for decades.

MARKET

Computational chemistry · materials science · finite-element analysis software.

Federated learning with consistency certificates

[ENABLED]

Cross-client agreement is a **legal requirement**, not just a metric, in regulated medical and financial federations. Our framework provides cryptographic certificates that cross-client model behavior is consistent at a stated tolerance, backed by mathematics rather than testing.

MARKET

Healthcare ML federations · cross-bank fraud models · regulated AI consortia.

Verifiable scientific simulations

[ENABLED]

Every simulation ships a **cryptographic ledger** of which conservation laws were preserved during the run, and to what tolerance. Directly addresses the reproducibility crisis in computational science. Per-timestep audit trail; verifiable by a third party.

MARKET

Government research · pharma · academic computational science.

Certified compression with multi-invariant guarantees

[ENABLED]

A new class of lossy compression that **preserves declared mathematical invariants** of the data, not just bits. A single compressed file can carry guarantees for several conservation laws simultaneously — uniformly certified, mathematically backed.

MARKET

Climate · computational fluid dynamics · molecular dynamics data infrastructure.

Provable scientific machine learning

[ENABLED]

Operator-inference and Koopman-style models that come with **theorem-grade generalization bounds** from the underlying contraction guarantees. Applies to neural operators, dynamical-system identification, reduced-order modeling — any setting where the model's behavior on test data needs a mathematical certificate.

MARKET

Scientific ML · neural-operator software · industrial digital twins.

The other side of the same math: red-team certification.

The mathematical framework is direction-symmetric. The same theorems that certify a transformation preserves an invariant also **characterize what would destroy it**. We frame these as defensive certification and adversarial-robustness audit tools — same mathematics, defensive packaging. Defenders learn the certified attack budget; attackers gain nothing they couldn't already compute.

Worst-case adversarial perturbation budgets

[RED-TEAM]

For any deployed machine learning model, computes the **smallest input perturbation that flips its decision**, with a mathematical lower bound rather than heuristic search. Generalizes existing adversarial-example tools. The defender's certified robustness budget *is* the attacker's certified attack-cost lower bound.

USE CASE

Model-card-grade adversarial-robustness certificates · ML deployment audits.

Anti-detection waveform topology characterization

[RED-TEAM]

Given a publicly-known radar / lidar / sonar receiver specification, characterizes the family of physical shapes whose signature falls **outside the receiver's certified detection envelope**. Defensive: tells radar designers which shapes their receiver cannot reliably classify — actionable for system improvement, not for offensive design.

USE CASE

Radar capability assessment · sensor-system robustness analysis · defense vendor R&D.

Provable jamming-energy bounds

[RED-TEAM]

For a communication channel with a declared receiver, computes the **minimum energy required to reduce its capacity** by a stated amount. Defensive use: tells radio designers how much jamming-resistance their hardware actually has. Strictly stronger than current empirical jamming-tolerance estimates.

USE CASE

Communications hardware certification · spectrum-warfare resilience audit.

Cryptanalysis attack-budget bounds

[RED-TEAM]

For any cryptographic primitive whose security argument fits the framework, computes **provable lower bounds on the attack budget** required to break it. Strictly stronger than current asymptotic security claims because it produces actual numbers. Useful as both a design tool and a competitive analysis tool against existing schemes.

USE CASE

Cryptographic scheme audit · post-quantum migration risk analysis · design-time safety.

Differential-privacy attack lower bounds

[RED-TEAM]

For any declared differential-privacy mechanism, computes the smallest **auxiliary-information injection** that defeats it with stated probability. Defenders compute this to set their epsilon budget; attackers gain a known information-theoretic floor.

USE CASE

DP-deployed ML systems · privacy-engineering audit · regulatory privacy guarantees.

Steganographic capacity bounds

[RED-TEAM]

For any declared content-detection system, computes how many bits of payload can be **provably hidden** in cover data such that the detector cannot find it within its declared budget. Useful for both privacy-preserving communication and content-moderation system design.

USE CASE

Content moderation system design · privacy-preserving communication audit.

Hash-collision and commitment-collapse design analysis

[RED-TEAM]

For algebraic hash schemes, characterizes the **structural conditions** under which collision attacks of a stated budget exist. Tool for designers of new hash schemes; identifies dangerous design patterns before deployment.

USE CASE

New-cryptosystem design review · structural cryptanalysis.

Worst-case input families for certified iterative solvers

[RED-TEAM]

Numerical solvers that advertise contraction rates (most commercial CFD, ML training, optimization software) get tested against **input families designed to violate the contraction claim**. Automated certification testing for “this solver actually does what it advertises.”

USE CASE

Numerical-solver certification · scientific-software audit · AI training stability.

Adversarial topology against ML feature extractors

[RED-TEAM]

Identifies the family of inputs that any deployed feature extractor **cannot embed faithfully**. Defensively: a model-card-grade certificate of where the model fails. Without gradient access — works on black-box models.

USE CASE

Black-box ML audit · model-failure characterization · regulated ML deployment.

Side-channel optimal trace-collection schedule

[RED-TEAM]

Given a publicly-declared masking scheme, computes the **minimum number of physical traces an attacker needs** to recover a secret with certified probability. Used by both sides of the masking-vs-attack arms race. Quantitative replacement for current heuristic estimates.

USE CASE

Side-channel attack budgeting · masking-scheme effectiveness audit.

Compression-bomb and training-data-poisoning attack design

[RED-TEAM]

Identifies the **minimum dataset perturbation** that maximally degrades a target ML system's declared invariant. Defensive use: certify a model's poisoning-resistance; offensive analysis: identify avoidable attack patterns to harden against.

USE CASE

ML training pipeline security audit · adversarial robustness certification.

*Defenders learn the certified attack budget by computing it. Attackers gain nothing from secrecy that the calculus does **not already give defenders**. These tools are framed as red-team and certification — never as operational guidance for offensive use against deployed systems.*

What we're aiming at, knowing we won't reach it.

The cryptography and the mathematical framework on the prior pages were extracted from a much larger research program. We are explicit about what we are and are not claiming as a deliverable: the structural targets below are **not products and not promises**. They are research directions whose *pursuit* produces the byproducts you have already seen — and the same pursuit will produce more byproducts over the coming months, regardless of whether the headline targets ever close.

An adelic non-commutative geometry framework

[RESEARCH]

A unified framework that ties the cryptographic stack, the mathematical analysis, and the certification tooling under a single structural identification rooted in the Connes / Bost-Connes adelic NCG tradition. **Partially closed at chart level** — the finite chart-algebra construction is proved with formal verification witnesses; the infinite-limit identification remains an active research target.

STATUS

Chart-level Hamiltonian construction proved · 60+ theorem candidates · infinite-limit lift open.

Connections to the Riemann Hypothesis

[RESEARCH]

The framework is structurally adjacent to several existing approaches to RH (Connes spectral, Bombieri explicit, Tate adelic). We use RH as a **tractor**: aiming at it pulls out a long list of intermediate tools and structural results that are valuable regardless of whether RH ever closes. We do **not** claim a path to RH. We do claim that the byproducts of aiming at it are what is on the prior pages — the cryptographic stack, the mathematical framework, and the certification tooling.

HONEST FRAMING

RH-tier closure: explicitly disclaimed. Connes-tier byproducts: shipping.

Multiple structural-identification candidates

[RESEARCH]

Several candidate identifications are under active testing — Bost-Connes adelic thermodynamics, Patterson-Sullivan exponents, minimal-model character coincidences, hyperbolic-flow connections. Each is tested with **predeclared falsifying criteria** and either closes as a real structural identification or is explicitly refuted with a CAS witness. The latter outcome — a clean negative — is also a contribution we ship.

DISCIPLINE

Predeclared falsifying criteria · CAS-witnessed refutations · honest negatives recorded as deliverables.

The continuous research engine itself

[RESEARCH]

The framework runs as a continuous AI-assisted research loop with cross-LLM peer review, manifest-bound proof promotion, and predeclared negative-result criteria. **The discipline architecture itself is a contribution:** most LLM-assisted mathematics drifts silently; ours catches its own regressions, produces honest negatives at the same rate as positives, and survives independent audit. That architecture is itself a research output we intend to publish.

CADENCE

Hundreds of cycles · ≈ 1 theorem candidate per cycle · canonical constants stable to 18 digits across 19+ cycles.

*We aim at the long-standing structural problems of analytic number theory **knowing we will not solve them**. The value is not in arriving — it is in what falls out of the tree as we shake it. The cryptographic stack and the proved mathematical framework on the prior pages are what fell out of the first six months. The coming months will produce more.*

Calibrated confidence.

No overclaim.

The work is real and the benchmarks are verifiable, but we are explicit about which categories of claim are at which confidence level. The list below is what gets said out loud, and what doesn't.

What we do claim

- ✓ The cryptographic benchmarks are real, measured, and reproducible. **Independent verification under NDA is available.**
- ✓ The mathematical theorems we say are proved are proved, with formal CAS witnesses where applicable.
- ✓ The "enabled" applications can be built with engineering effort, with development cost varying by item.
- ✓ The defensive / red-team capabilities can be packaged as certification tools today.
- ✓ The framework is genuinely novel as a unification — the individual mathematical tools mostly exist; the combination doesn't.
- ✓ The cryptographic stack is a **commercializable asset today**, independent of how the mathematics evolves.

What we don't claim

- ✘ Any solution to the Riemann Hypothesis or other century-old mathematical open problems.
- ✘ An exact, fully-proved adelic theory. A precise finite-section calculus is in hand; the infinite lift remains an active research target.
- ✘ Anything labeled "AGI." Anything labeled a quantum computer. Anything labeled artificial general anything.
- ✘ Replacement of standard cryptographic review processes — **independent expert audit is required** before any deployment claim.
- ✘ Operational guidance for offensive use against deployed real-world systems. Capabilities are framed as defensive, certification, and red-team.
- ✘ Probabilistic estimates dressed up as deterministic guarantees.

The five tiers, with honest labels

Each tier above corresponds to a different evidence level. We don't conflate them. A Silicon Valley investor evaluating this work should understand that Tier A is shipped product, Tier B is mathematical research-grade output, Tier C is buildable application roadmap, Tier D is buildable defensive tooling, and Tier E is frontier research that we explicitly disclaim as a deliverable.

TIER A Shipped

Working code, measured benchmarks, reproducible by third parties under NDA.

TIER B Proved

Mathematical theorems with formal verification witnesses where applicable.

TIER C Enabled

Buildable applications with engineering work; not yet packaged as products.

TIER D Certified

Defensive / red-team capabilities, framed as certification tools.

TIER E Frontier

Research targets explicitly disclaimed as deliverables. Optionality, not promise.

*Cryptography that's already running.
Mathematics that's already proved.*

The combination is the moat.

ASTREA FOUNDATION

research@astrea.systems

INDEPENDENT VERIFICATION UNDER NDA IS AVAILABLE.