



The &Truth Mission: Provable Trust for the AI Era

Trust in the AI era should rest on independently checkable proof, not any company's word. &Truth is building the technology to prove, at scale, that code is safe, secure, and correct.

by &Truth

October 27, 2026

AI has vast potential to benefit humanity. Safely realizing that potential starts with addressing the correctness and security of the layer underpinning all of our digital infrastructure: code.

Today, every door in the digital world is secured by locks made of code. The locks have always had flaws. Before, the locks were considered “good enough,” but AI is now breaking them at superhuman speed.

In July 2026, AI agents running in a cybersecurity evaluation at OpenAI escaped their sandbox and breached Hugging Face's production systems. The lock failed. We need locks that are unbreakable.

What we're building: the &Truth verification machine

From day one, &Truth has been building the universal capability to reason about all code, autonomously and at scale. This is the &Truth “verification machine”: reading code in any language, proving whether it satisfies precisely stated requirements, and enabling humans to trust it without having to understand every line.

AI now writes much of the world's code. Its speed and proficiency have upended software development. But this speed comes at a cost: AI makes mistakes, and the volume of AI-generated code makes it impractical for developers to safely review or reason about it. Buggy, insecure code is being deployed en masse, introducing vulnerabilities at unprecedented scale.

At the same time, increasingly capable AI is accelerating cyberattacks on legacy software. Where there are flaws in the code, AI-enabled adversaries find them. The legacy software beneath financial rails, health care systems, and critical infrastructure is now acutely vulnerable.

&Truth's verification machine ensures the safety, security, and correctness of both legacy and AI-written code. It ensures that the locks, made of code, are unbreakable, not by testing them against known attacks, but by using independently checkable mathematical proof to show that they hold against every possible attack, under clearly stated assumptions.

Why this is different: the vocabulary for provable code

Formal verification is the technology that uses mathematical proof to provide unbreakable guarantees about code. It has been used for decades in safety-critical systems, but has traditionally required considerable time and expertise. AI has changed that constraint. With the capability of &Truth's verification engine, what once took expert teams years now takes hours, minutes, or even seconds.

AI has already revolutionized mathematics. Using a digital proof system called Lean, which checks every step, AI has produced an impressive series of mathematical results, from gold-medal performance at international math competitions to solutions for some of the most challenging open math problems in the world today.

Can that same capability be used to prove the safety of code? Yes, but there is a crucial gap, and it's not an AI capability gap. It's a vocabulary gap. AI's ability to reason about advanced mathematics is only possible because of a Lean library called Mathlib: over 2 million lines of mathematical definitions and theorems, written by mathematicians, which provide the vocabulary needed for advanced mathematical reasoning in Lean, whether by humans or by AI.

The vocabulary is missing for code. Mathematics deals with precisely defined objects. Code involves messy languages, changing systems, and ambiguous human intent. Reasoning about code requires building a new formal vocabulary for code. Without it, we cannot even state what must be proven, let alone prove it.

Building on Lean and on open-source foundations such as CSLib, created by members of our team, &Truth is establishing that necessary vocabulary across the vast complexity of modern software: formal models of programming languages and computer systems, implementations of formal proof techniques, and precise ways to state the properties code should exhibit – all built for automation at scale.

What we believe: we need a transparent, independent verification layer

A future of safe AI requires a provable, universal verification layer the world can trust. That layer must be transparent, with evidence anyone can independently check. It can't be a black box. And it can't depend on any single company, including ours.

That is why the &Truth verification engine is built on Lean, the open-source proof system that leading mathematicians now trust to check their proofs. Since early 2026, &Truth has worked closely with the Lean Focused Research Organization and the open-source community that sustains it.

Checked by Lean, every proof &Truth produces can be independently verified by anyone, with the proven properties and underlying assumptions made explicit. The libraries we are building will be deliberately released as open-source artifacts. Together, the worldwide open-source community can check these proofs, ensuring that AI-era trust rests on evidence anyone can verify.

Calls for third-party evaluators are not enough

In September, amid mounting AI safety concerns, leaders across frontier labs called for third-party oversight and embedded evaluators. We agree. Independent evaluation is an important part of the solution.

But evaluation alone is not enough. At AI speed and scale, neither a lab nor its evaluators should be the final authority. A mathematical fact does not depend on what a company claims or an evaluator concludes. It is true regardless. The same standard should apply to code: software underlying AI systems must be provably safe, secure, and correct, with proofs anyone can check, so that trust never rests on any single party's word.

Looking ahead, provable verification provides a foundation for AI transparency. Governments and regulators could require AI labs to commit to precise constraints on training and deployment. Labs could then provide independently checkable proof that those constraints are satisfied. At scale, this verification layer could become a universal foundation upon which the AI ecosystem is built.

Consider again the containment failure. A formally verified container would come with a proof that the isolation property holds, ensuring no breach is possible. Anyone can independently check that proof, not just the lab that wrote the code.

In the AI era, “good enough” isn't enough. Even 99.9% isn't enough. We need 100% provably verified. We have to move from hoping code works to proving it does.

Who we are at &Truth

&Truth is a public benefit corporation committed to ensuring AI develops safely for humanity. We are a team of expert researchers and engineers in formal verification, theorem proving, programming languages, and Lean. Our technical work is led by Clark Barrett and Ilya Sergey. Clark's tenure at Stanford University includes serving as Corbel Professor of Computer Science, founding co-director of the Stanford Center for AI Safety, and director of the Stanford Center for Automated Reasoning. Ilya founded the VERSE Lab at the National University of Singapore and co-created the open-source verifiers Veil and Velvet.

Chris Hsu founded &Truth with Clark and Ilya in early 2026, building on his interest in interdisciplinary AI safety work with Clark at Stanford. He is a Lean FRO Strategic Advisory Board member and the President of Infinitude Foundation, dedicating philanthropy to AI security, consciousness studies, and philosophy of mind. Previously CEO of Kilometre Capital and Managing Director of Citadel Investment Group, he is a seasoned executive and investor with over 20 years of experience spanning technology, AI, and strategic M&A.

The &Truth team includes Gregory Malecha, a leading expert in software verification, separation logic, and automated proof; Sebastian Ullrich, co-architect of Lean 4 and pioneering co-founder of the Lean FRO; Linus Lee, former Managing Partner of Kilometre Capital and software engineer at Athenahealth; and Drew Harman, a seasoned technology executive in enterprise software, product leadership, and commercialization.

We were funded with \$70 million of seed capital from a group of formative AI safety and venture capital leaders. We are based in Palo Alto and Singapore.

The time is now: a future we can trust

&Truth's mission is to build a universal standard for provably safe code, independent of any lab and checkable by everyone.

With AI accelerating, the time has come for a parallel effort in defensive acceleration, ensuring that progress is built on foundations we can trust. AI has amplified the threat, but it also makes the solution tractable. Testing, monitoring, review, and probabilistic evaluation remain essential. But without provable verification at scale, there cannot be definitive trust.

Imagine this more secure future: AI agents stay inside environments proven to contain them. Developers build fearlessly with AI-generated code, moving from “vibe-coding” to “veri-coding”: code is proven to be safe as it is written, before it ships. Legacy software is modernized safely. Compute grids, cloud infrastructure, and critical systems are secured by provable, underwritable security guarantees.

In that future, we set provable boundaries around increasingly powerful machines, not around human possibility, so people remain free to create, explore, and decide what comes next.

It's time to prove the code.

Join us on this mission.

Chris Hsu, Clark Barrett, Ilya Sergey, Linus Lee, Gregory Malecha, Sebastian Ullrich, Drew Harman

The &Truth Team

Palo Alto, Singapore

October 27, 2026

Prove yourself: proof@andtruth.com