Poseidon Cryptanalysis Initiative 2024–2026

State of the Art

Dmitry Khovratovich

Ethereum Foundation

SPRING, 10 May 2026

Confidence in Symmetric Crypto Is Earned, Not Assumed

In symmetric cryptography, a primitive gains trust only through **sustained public scrutiny**: independent researchers trying — and failing — to break it.

The SHA-3 competition (2007–2012) is the canonical example:

- **64 submissions** received by NIST
- **51** advanced to Round 1
- **14** advanced to Round 2
- **5** finalists; **1 winner** (Keccak)
- Dozens of candidates eliminated by published cryptanalysis — collisions, preimages, 2nd-preimages found within months

The SHA-3 Zoo (`ehash.isec.tugraz.at`) catalogues every attack on every candidate. The survivors earned their place.

What public scrutiny requires

1. A **clear, narrow problem statement** that researchers can attack
2. **Open access** to specifications, parameters, and test vectors
3. **Incentives** for the community to invest effort
4. **Time** — years, not months

For ZK hash functions

The field is new (∼2016), the designs are algebraic, and the attack surface (CICO, density, collision, mode) is different from classical hash functions. The Poseidon Initiative is building the scrutiny infrastructure from scratch.

Program Overview

Mission: systematic, community-driven cryptanalysis of Poseidon hash functions deployed in high-value Ethereum applications.

Phase 2 (2024–2026):

- Expanded scope: Poseidon2 → Poseidon1 (KoalaBear)
- Higher bounty budget, new problem types (CICO, density, zero-test, collision)
- Formal grant programme for foundational research
- Regular workshops attached to major venues

EFPG team: Kadianakis, Khovratovich, Sanso

Advisory board: Aumasson, Ben-Sasson, Hopwood, Lubarov, Rothblum

Note: Poseidon authors recused from grant & bounty decisions.

Budget overview

Programme	Budget
Bounty 2025	\$130K
Poseidon1 Collision Prize	\$992K
Bounty 2026	\$150K
Attack Reward 2026	\$90K
Short-Term Grants 2026	$\leq$ \$160K
Total	$>$ **\$1.5M**

Bounty Program 2025 — What Was Claimed

Target: partial preimage of **0** (first capacity element). Three instances: Poseidon-256 (BLS12-381, d=5), Poseidon-64 (Goldilocks, d=7), Poseidon-31. $130K total; closed **1 Dec 2025**.

Instance	Parameters	Level	Rounds (R_P)	Bounty	Claimed
Poseidon-256	BLS12-381, d=5, t=3	24-bit	8	$4K	**9 Dec 2024**
Poseidon-31	KoalaBear, d=3, t=16	24-bit	–	$4K	**30 Nov 2024**
Poseidon-31	KoalaBear, d=3, t=16	28-bit	–	$6K	**29 Nov 2024**
Poseidon-31	KoalaBear, d=3, t=16	32-bit	–	$10K	**5 Dec 2024**
Poseidon-31	Mersenne-31, d=7, t=16	24-bit	–	$4K	**29 Nov 2025**
Poseidon-31	Mersenne-31, d=7, t=16	40-bit	–	$15K	**5 Nov 2025**
Poseidon-64	Goldilocks, d=7, t=8	24-bit	7	$4K	**23 Apr 2025**
Poseidon-64	Goldilocks, d=7, t=8	28-bit	8	$6K	**27 Apr 2025**
Poseidon-64	Goldilocks, d=7, t=8	32-bit	10	$10K	**24 May 2025**
Poseidon-256	BLS12-381, d=5, t=3	28-bit	9	$6K	**31 Dec 2024**

Best attack: **interpolation**. Gröbner basis improvement bonus — *not collected*.

Verification repository

`https://github.com/khovratovich/poseidon-tools` — the authoritative source for verifying all bounty submissions. Authors who disclose new solutions are expected to submit there; accepted claims appear in the repo.

Bounty Program 2025 — Unclaimed Records

Program closed 1 Dec 2025. The following levels were never broken; they survive as **open attack records**.

Unclaimed at close

Instance	Rounds (R_P)	Value
Poseidon-256	R_P=11, 32-bit	\$10K
Poseidon-256	R_P=16, 40-bit	\$15K
Poseidon-64	R_P=13, 40-bit	\$15K

\$40K in records remain unbeaten.

What it means

- Interpolation reaches R_P=10 on Goldilocks, R_P=9 on BLS
- Above that the polynomial degree exceeds feasibility
- Gröbner basis could push further — open research question
- See Grants 2026: TU Graz task on GB complexity

Why We Moved from Poseidon2: Round Skipping Attack

"Skipping Class" — Merz & Rodríguez García (ETH Zurich), ePrint 2026/306, Feb 2026.

- Exploits the non-MDS internal matrix of Poseidon2 (chosen for circuit efficiency, branch number $b < t+1$)
- By spending **4 degrees of freedom per round**, the attack can **skip** $\approx t/4$ **full rounds**, effectively reducing the security margin
- New algebraic preimage attack: easier than the corresponding CICO problem
- First algebraic *collision* attack outperforming the preimage counterpart
- Does *not* break recommended parameters outright, but substantially erodes the security margin

Response: Poseidon Initiative pivots to **Poseidon1** (KoalaBear, MDS matrix) for Bounty 2026.

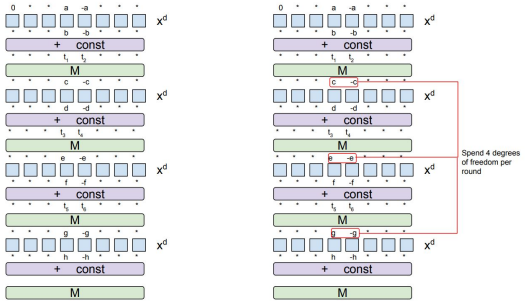

Truncated differential exploited by the round-skipping attack.

MDS Property Upper-Bounds Round Skipping

Theorem (Informal, Khovratovich 2026)

A trail skipping m rounds of an SPN with t state elements and linear branch number b requires:

$$b \leq \frac{4t}{m} \quad (m \text{ even}), \qquad b \leq \frac{4t-2}{m-1} \quad (m \text{ odd}).$$

Corollary: MDS kills 4-round skipping

An MDS matrix has branch number $b = t{+}1$. For a 4-round skip ($m{=}4$):

$$b \leq \frac{4t}{4} = t < t{+}1 = b_{\mathrm{MDS}}.$$

Contradiction $\Rightarrow$ no 4-round skipping trail exists for any MDS-based SPN.

3-round skipping — still open

For $m{=}3$: bound gives $b \leq 2t-1$, which MDS satisfies ($t{+}1 \leq 2t{-}1$ for $t \geq 2$). No 3-round trail found yet;

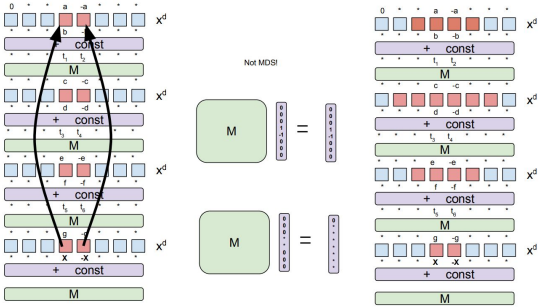

Poseidon2's non-MDS layer enables the high-probability truncated differential used in the skip.

Poseidon1 Collision Prize

New in Phase 2. Find a **partial collision** on Poseidon1 ($\mathbb{F}_p^{16} \to \mathbb{F}_p^1$, KoalaBear, d=3):

$$H(X)_{\{0,\ldots,q-1\}} = H(Y)_{\{0,\ldots,q-1\}}, \quad X \neq Y.$$

q outputs	Prize	Status
3	\$32K	**Claimed 6 Apr 2026**
4	\$64K	**Open**
5	\$128K	**Open**
6	\$256K	**Open**
7	\$512K	**Open**

Runs until 1 Jan 2029.

Why this matters

- A 7-output collision is a full collision on a widely used primitive
- \$512K prize signals real-world deployment stakes
- q=3 claim (Apr 2026) shows the problem is tractable at low q
- Gap between q=3 (claimed) and q=4 (open) is the current cryptanalytic frontier

Note

MDS structure of Poseidon1 means collision resistance relies on a different (stronger) algebraic argument than Poseidon2.

Bounty Program 2026 — CICO / Density / Zero-test

Target: Poseidon1 KoalaBear (d=3, t=16, R_F=6). \$150K total, closes **1 Jan 2027**.

CICO Problem

Find (X, Y) with constrained input and output.

Parameters	Prize	Status
R_P=6	\$6K	**Claimed 10 Apr 2026**
R_P=8	\$10K	**Submitted 6 May 2026**
R_P=10	\$15K	**Open**

Density Challenge

Maximize fraction of zero entries in a preimage. \$40K total; two-phase (Aug / Dec 2026). Current record: R_F=6, R_P=6.

Zero-Test Challenge

Find a non-trivial multivariate polynomial vanishing on the hash output. \$40K total; ranked by rounds. Current record: R_F=6, R_P=6.

Submissions & verification

Authors disclosing solutions submit to `https://github.com/khovratovich/poseidon-tools`; accepted claims and verification scripts are published there.

Why Poseidon1 / KoalaBear?

- MDS matrix ⇒ no 4-round skipping (see previous slides)
- KoalaBear ($p = 2^{31} - 2^{24} + 1$) widely adopted in STARKs
- d=3 gives lower circuit cost than d=5/7

Attack Reward Program 2026 & Gröbner Basis Exploratory

Attack Reward Program 2026

- **Budget:** \$90K total, \$5K minimum per award
- **Scope:** any improved algebraic attack on reduced-round Poseidon-256, Poseidon-64, or Poseidon-31
- **Condition:** attack published on ePrint by 31 Dec 2026
- Awards judged relative to records set in Bounty 2025
- Gröbner basis improvement on those records earns **bonus**

Gröbner Basis Exploratory Grant (2025)

- Deadline: April 1, 2025
- Goal: derive *precise* GB complexity formulas over all feasible (R_F, R_P) parameter combinations
- Output: closed-form bound enabling automated security margin analysis
- Partially delivered; ongoing work feeds into Short-Term Grants 2026

Workshops & Retreats 2025–2026

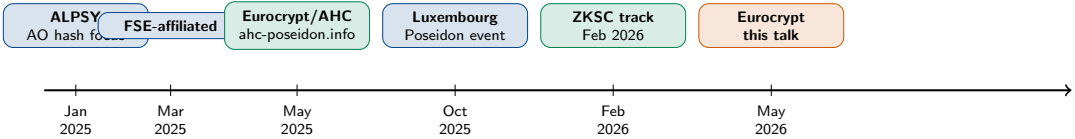

Completed (2025)

- **ALPSY Jan 2025** — algebraic & arithmetic hash focus; first presentation of Phase 2 bounty rules
- **FSE-affiliated Mar 2025** — progress report on GB analysis
- **Eurocrypt / AHC May 2025** — collision prize announced; community Q&A on Poseidon2 concerns
- **Luxembourg Oct 2025** — internal retreat; Poseidon2 decision

Active (2026)

- **ZKSC Feb 2026** — dedicated Poseidon track; round-skipping attack presented publicly for the first time
- **Eurocrypt May 2026** — this presentation; Bounty 2026 launch, Short-Term Grants announced

Short-Term Grants 2026

$20K–$40K each. Deadline: 1 June 2026.

Topic	Research question	Team
Fiat-Shamir attacks	Exploit Poseidon in a Fiat-Shamir-based protocol (e.g. Schnorr / Plonk transcript); find parameters where FS security is easier to break than standalone CICO.	**Aalto University**
Gröbner basis attacks	Extend the GB complexity analysis to all (R_F, R_P) combinations; produce automated tooling for security-margin certification.	**TU Graz**
Round-skipping bounds	Prove (or disprove) the existence of 3-round skipping trails for MDS-based SPNs; close the gap left by the informal bound.	**ETH Zurich**
Round-skipping (independent)	Independent research on round-skipping attacks across SPN designs (not restricted to Poseidon); exploring non-MDS and MDS variants.	**ongoing**
Quantum attacks	Analyse collision and preimage resistance of the Poseidon sponge and compression function against quantum adversaries (BHT, Grover, quantum walk).	*open call*

Outlook — Towards a Public Certificate of Confidence

Where we stand (May 2026):

- Interpolation ceiling reached on Poseidon-256/64; gap to full rounds remains
- CICO on Poseidon1-KoalaBear moving fast (R_P=8 just submitted)
- Collision prize: q=3 claimed, q=4 is the immediate frontier
- Round-skipping eliminated for MDS (4-round); 3-round case open
- Fiat-Shamir and quantum security largely unexplored

Phase 2 closes: December 2026

The central goal: a public certificate

A **public certificate of confidence** is a virtual, community-held conviction resting on three pillars:

1. **Breadth of attack coverage** — many independent cryptanalysis directions have been seriously tried: algebraic, differential, statistical, quantum, mode-of-operation
2. **Quality of scrutiny** — the best scholars in the field have looked at it and published their findings
3. **Well-scoped problem** — the security claim is narrow and precise enough that a failure would be unambiguous

No document encodes it. It emerges from the record.
We are building that record.

`https://poseidon-initiative.info`

How I Made These Slides with AI

All slides were generated by **GitHub Copilot** (Claude Sonnet) in VS Code from the following plain-English commands, issued in order:

1. *"I plan to give a talk about the state of the art of the Poseidon Initiative... Show me the plan first. It should cover: bounties claimed and alive, workshops, grants from 2025 and new ones, grants from 2026."*
2. *"For Slide 9 we have: Aalto University on Fiat-Shamir, TU Graz on Gröbner basis, ETH Zurich on round skipping. Also add 'independent research on round skipping – ongoing'. Slide 10: emphasize the need of a public certificate of confidence."*
3. *"Also add ePrint 2026/306 for a slide between 4 and 5: we decided to ditch Poseidon2 due to round-skipping attack that exploits the non-ideal diffusion. Use the Materials to show that MDS upper-bounds skipping rounds to 3."*
4. *"I have added the mds-skip repo to Materials/. Use it for the slide on the MDS property against round skipping."*
5. *"Mention the repo github.com/khovratovich/poseidon-tools as the final word on verifying bounty submissions. New submissions which authors disclose will appear there."*
6. *"Bounty 2025 program is closed – it is not alive!"*
7. *"The unbroken claims do not qualify for the 2026 program."*
8. *"Last slide: public certificate is a virtual thing that includes (1) admitting that many cryptanalysis directions have been tried, (2) best scholars have looked into it, (3) the problem was well defined and narrow enough. Also add a slide on public scrutiny in symmetric crypto with the SHA-3 example."*
9. *"Make a summary slide with my commands under the title 'How I made these slides with AI'."*

The AI fetched web pages, read local repos and PPTX files, wrote LaTeX, fixed overflows, and compiled the PDF – all autonomously.