THE CURRENT

10 NOVEMBER 2026 / SCIENCE, SOFTWARE AND THE EVIDENCE BEHIND THE CLAIMS

Proof reaches the running code

AlphaProof Code links security claims to a compiled Signal implementation. Its most important boundary is what the proof actually covers.

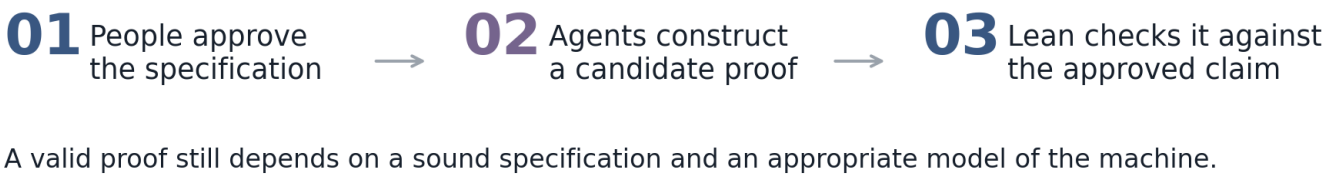

Original explanatory diagram. Human judgment remains in the specification and trusted assumptions.

Google DeepMind announced AlphaProof Code on November 10: an agent framework that proves selected security properties of Signal’s cryptographic implementation. The ambition is to bridge a familiar gap. A protocol may be sound on paper while the program implementing it behaves differently. [1]

Researchers say the process took four weeks. AI agents generated millions of lines of Lean proof code, while humans concentrated on the statements to be proved and the assumptions beneath them. Lean, a proof assistant, checks the result rather than accepting a model’s assurance that it is correct. Two previously fixed bugs were reintroduced into an isolated version; the system proved that those versions violated the specified security properties. [1]

What the technical paper adds

The study targets libsignal v0.105.4 and SPQR v1.6.0, compiled to AArch64 machine code through a command-line harness. It verifies session initiation, sending and receiving; key generation is outside that verification. Group messaging is left for future work. [2]

[1] DeepMind announcement · 10 November

[2] Technical paper and verification artifacts

The trust boundary moves

The technical account reports about 6.5 million lines of proofs resting on roughly 26,000 lines of trusted specifications and machine semantics. Cryptographic primitives and some called routines are represented by contracts. The proof therefore still depends on those assumptions, the machine model and Lean’s kernel. [2]

Its symbolic and computational results also differ. The computational analysis does not cover quantum attackers or multi-device state, and its bound becomes uninformative for multiple long sessions with many epochs. The symbolic treatment has its own exclusions, including the post-quantum healing property of SPQR. This is not a proof of every security property of every Signal application. [2]

Why it matters

The practical lesson is that a convincing explanation and a checkable proof are different products. An automated proof writer can be fallible if the acceptance mechanism reliably rejects invalid work. But a flawless proof of a weak claim can still leave the real concern untouched.

For future verification announcements, ask three questions: which executable was checked, what exact property was established, and which components remain trusted? Those answers tell a reader much more than the size of the proof or the confidence of the agent that wrote it.

Editorial analysis. The proof build was not independently rerun for this edition.

02 / AI IN PRACTICE

A smaller bill and a wider boundary

A production classification case and an incident notice show why model choice and data handling belong in the same conversation.



Reported cost per million tickets, USD. Amazon’s workload and list-price assumptions; not an independent benchmark. [3]

The cheaper job is a narrower job

Amazon’s November 9 Strands case study describes replacing Sonnet 4.6 with Strands Decider E4B for a ticket-labeling workload. Instead of writing a JSON answer, the decision model selects from a fixed catalog. The engineers report costs falling from about $22,200 to $500 per million tickets, with average call time dropping from 4,215 to 897 milliseconds. [3]

The change was more than a model swap. Twenty labels were rewritten as short criteria, a “none” option was added, and the oldest comments were dropped from long tickets to fit the smaller context window. Requests used sorted JSON keys after ordering alone shifted output probabilities by as much as 0.19. [3]

The team first ran both systems on the same traffic while retaining Sonnet’s decisions. The estimated savings depend on keeping the hosted hardware busy; this is not a price for arbitrary reasoning or writing. The published account does not establish a universal quality advantage. [3]

A useful deployment test

Original analysis: count the expensive mistakes, not just disagreements. A classifier can look efficient while routing the rare urgent ticket incorrectly. Keep the label set and input policy explicit, inspect uncertain cases, and measure the cost of review alongside inference. A narrower task can justify a smaller model without implying that a general assistant should become one.

[3] Strands production case study · 9 November

A transcript can identify itself

Notion’s current Fireworks incident notice says fewer than 1% of its workspaces may have been affected. It describes possible exposure of Notion AI chat transcripts used for its Insights feature. The records lacked user or workspace labels and tool-call results, but words inside a conversation could still contain sensitive information. [4]

The notice says Fireworks supplied further confirmation on November 8. Notion has contacted affected workspace owners and provides an incident-review area. It advises treating the affected data as exposed and rotating credentials if they appeared in those transcripts. The notice says no action is needed now if neither the user nor the workspace owner has heard from Notion. [4]

The vendor’s account

Fireworks traces the incident through a third-party cloud provider and development machines on October 23–24. Its update lists affected categories including resource metadata, billing information and subsets of some customers’ model weights and datasets. It says payment-card numbers were not stored there and were not affected. [5]

The company reports no detected impact to its default-zero-retention Inference, Training SDK and Nexus services, while stressing that the investigation continues. That statement and Notion’s transcript warning describe different data paths; they should not be flattened into a blanket claim that all inference traffic leaked, or that no customer text did. [5]

[4] Notion incident notice · checked 10 November

[5] Fireworks investigation update

03 / SPACE AND ASTRONOMY

A slingshot and a hidden companion

Two November 9 reports turn timing into evidence: one puts spacecraft on course; the other explains a star that keeps erupting.

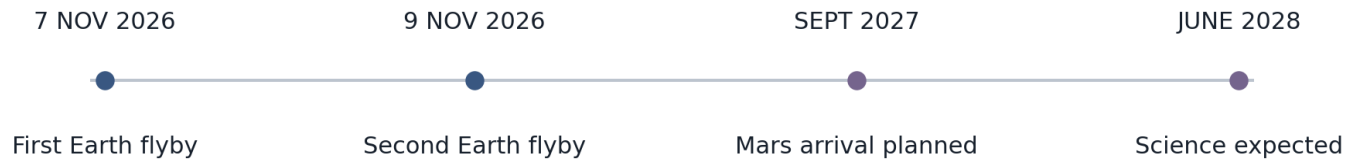

ESCAPADE milestones reported by NASA. Future dates are plans, not completed events; spacing is schematic. [6]

Mars by way of Earth

NASA’s twin ESCAPADE spacecraft have completed powered Earth flybys and are heading toward Mars. The first maneuver occurred on November 7 and the second on November 9. Each combined Earth’s gravity with an approximately nine-minute engine burn. [6]

After launching in November 2025, the pair loitered on a trajectory around the Earth–Sun L2 region rather than traveling directly to Mars. Returning to Earth at the right alignment let them obtain the boost for the next leg. NASA describes the approach as a way to loosen the usual constraints on launch timing. [6]

Mars orbit insertion is planned for September 7 and 9, 2027, with science operations expected in June 2028 after further orbital adjustments. The instruments will study the solar wind and Martian atmospheric loss. Coordinated observations from two locations should help separate changes over time from differences between places. [6]

What this milestone is

Original analysis: the news is successful navigation, not a completed atmospheric discovery. The scientific payoff still requires the cruise, arrival burns, instrument operations and useful paired measurements. Separating those stages protects an important engineering achievement from a claim it has not yet earned.

[6] NASA ESCAPADE report · 9 November

The star that survives its explosions

AT 2016blu repeatedly brightens like a supernova without being destroyed. NASA’s November 9 account of a new Astrophysical Journal study points to a hidden companion: either a stellar-mass black hole or a neutron star drawing material from the massive star. [7]

The object lies in NGC 4559, about 30 million light-years away. Observers recorded at least 27 optical outbursts between 2012 and 2026, recurring roughly every 113 days. Chandra observations were timed to an outburst; its angular resolution let researchers separate the source from a nearby bright X-ray object. [7]

The proposed picture is an orbiting pair in which transferred material grows hot enough to radiate X-rays. The result narrows the explanation without identifying which kind of compact remnant is present. Calling it a confirmed black hole would outrun the reported evidence. [7]

The value of another kind of light

Original analysis: a repeating brightness pattern can suggest a clock without revealing the mechanism that winds it. A second wavelength adds a different physical clue. The interesting advance here is not simply a sharper photograph; it is an observation chosen to distinguish explanations at the moment they differ.

NASA’s release was read in full. The linked journal page was not accessible in this reporting pass, so the account above is attributed to NASA rather than presented as an independent audit of the paper.

[7] NASA Chandra report · 9 November

04 / CLIMATE AND THE RECORD

A dip is not yet a turning point

The new carbon budget projects lower fossil emissions in 2026, but a higher total once land-use change is included.

−0.4%	43.3 Gt	42.8 Gt
Fossil CO_2 change	Total CO_2 in 2026	Total CO_2 in 2025

Global Carbon Budget figures. 2026 values are projections; totals combine fossil and land-use CO_2, not all greenhouse gases. [8]

Two totals tell different stories

The Global Carbon Project's November 10 release projects fossil CO_2 emissions down 0.4% in 2026, to 38.2 billion tonnes. It attributes the decline to clean-energy expansion and the energy shock from the Iran conflict. Oil emissions are projected to fall 1.9%. [8]

That is not the whole carbon ledger. Adding land-use change brings projected total CO_2 to 43.3 billion tonnes, above 42.8 billion tonnes in 2025. The release connects the land-use increase with peat fires and forest degradation in Southeast Asia amid drier conditions. The 2026 budget is initially a preprint, not yet its eventual peer-reviewed version. [8]

A companion Nature Climate Change analysis was published November 10. Its calculations distinguish observed emissions from counterfactual emissions without solar, wind and electric vehicles, and test sensitivity to deployment assumptions. That distinction matters: avoided emissions are an estimate of what did not happen, not a directly measured flow. [9]

Progress can coexist with growth

The project's FAQ says 40 economies significantly reduced fossil CO_2 while growing over 2016–2025, compared with 22 in the preceding decade. It also notes the slow turnover of vehicles: only about 6% of the car fleet was electrified in 2025 when non-plug-in hybrids were excluded. Rapid sales growth takes time to change the installed stock. [10]

[8] Global Carbon Budget release · 10 November

[9] Le Quéré and colleagues · Nature Climate Change

[10] Global Carbon Budget 2026 FAQ

What would make the decline durable

Original analysis: a lower annual flow is encouraging only to the extent that its causes survive. Reduced demand during an energy disruption can reverse. Equipment that displaces fossil fuel use can keep doing so, but its effect competes with new demand and depends on how it is operated.

The accounting boundary matters as much as the direction of the headline. Fossil emissions, land-use emissions and atmospheric concentration answer different questions. A decrease in one component does not establish a decrease in the sum, and a smaller positive flow does not mean the atmosphere has stopped accumulating carbon.

Read this year's projection as a provisional balance sheet. The next useful evidence will be completed-year inventories and a sustained change in underlying sources, rather than a declaration of victory from a single percentage.

How this edition was reported

This issue follows the November 9 edition with newly checked November 9–10 reporting. Publication dates and event dates are separated where they differ.

Hacker News supplied leads. Jag's AI Heat Map was checked, but its article inventory was not available in retrieved text. Primary announcements, technical papers and institutional reports supplied the evidence.

Reported results were not independently reproduced. Climate reporting used primary-site search extraction where direct retrieval failed. Source links are embedded in both editions.

Discovery desk · Hacker News

Discovery desk · Jag's AI Heat Map