oOoOoOo / 007

13 NOV 2026
WEEKLY TECHNOLOGY REVIEW
TUTORIJALI.APP

COVER STORY

PROOF INSTEAD OF PROMISES

AI agents proved Signal's code. Companies sign numbers that can be checked.

06 ALPHAPROOF CODE Signal proven down to machine code
11 FABLE 5.2, GEMINI 4 Stronger models, lower prices
16 CLAUDE, CHATGPT AI moves into sheets and meetings
22 CONTROL ROOM A stop within 35 seconds
25 CISA KEV Barracuda, OpenSSH, Outlook
28 ANTHROPIC IPO 900.9 billion in commitments
33 LAB A robot in a plant, DNA by cell
37 ORBIT NASA starts building a Moon base



0 major story o news 0 agents

One glyph, one event

6 Nov 2026, 00:00 → 13 Nov 2026, 05:37, Belgrade

Free edition in PDF format

Contents

The issue is split into sections, each with its own colour. Page numbers are links.

04 EDITORIAL
What can be checked

05 COVER STORY
Proof instead of promises
Five questions through the issue.

06 COVER STORY
Signal, proven down to machine code
AlphaProof Code and libsignal.

07 COVER STORY
Specification, agents, Lean's kernel

08 COVER STORY
What is proven, and what is not

09 COVER STORY
Two old bugs as an exam

10 COVER STORY
Years of work become weeks

11 MODEL WATCH
Fable 5.2 and Mythos 5.2: one model, two doors
A cheaper, faster top Claude.

12 MODEL WATCH
What the numbers show, and whose they are

13 MODEL WATCH
Gemini 4: Carbon for everyone, Hydrogen for speed
Two models in three days.

14 MODEL WATCH
An expensive planner, a cheap worker

15 MODEL WATCH
Models of the week

16 WORK
AI no longer waits in a chat window.

17 WORK
Claude gets sheets, boards and data
Sheets, Whiteboards and Data.

18 WORK
Claude in the meeting: who may listen

19 WORK
Agents inside your perimeter

20 WORK
Businesses enter ChatGPT, ChatGPT enters AWS

21 WORK
Agents built for an industry

22 SECURITY
A kill switch that works when everything else fails
OpenAI Control Room.

23 SECURITY
Why the switch sits in the inference layer

24 SECURITY
AI finds bugs. So do attackers.
OpenSSH, Outlook and Hacktron.

25 PATCH NOW
Patch immediately

26 SECURITY
Eleven confirmed attacks in five days

27 SECURITY
Two chains: root on a box and a hash from a message

28 MONEY
Anthropic heads to market as ANT
Roadshow, price range and S-1/A.

29 MONEY
Revenue grows, but not like this forever

30 MONEY
900.9 billion dollars of commitments

31 MONEY
CoreWeave: record revenue, record debt

32 MONEY
Nebius: customers pay for the build up front

33 LAB
A humanoid on the assembly line
Gemini Robotics 2.5 and Apollo 2.

34 LAB
Dexterity, speed and safety around people

35 LAB
AlphaGenome 2 reads DNA cell by cell
From tissue averages to single cells.

36 LAB
Twice as many causal variants

37 ORBIT
Moon Base becomes a procurement programme

38 ORBIT
A rover that survives the lunar night

39 ORBIT
Contracts arrive, missions slip

40 ORBIT
ESCAPADE is on its way to Mars

41 BUILD
Secrets in logs and alerts that switch off

42 USE THIS
An agent a company can trust
A four-step guide with TypeScript examples.

43 USE THIS
A kill switch that fails closed

44 USE THIS
The agent never sees the key or picks the address

45 USE THIS
Irreversible waits for a person, everything is logged

46 SHIPPED
Arrived this week

47 5 MINUTES
In brief

48 NUMBERS
The week in numbers

49 NUMBERS
More numbers of the week

50 NEXT
What we are watching

51 SOURCES
All sources

EDITORIAL / 0000000 #007 / 13 NOV 2026

What can be checked

This week the biggest stories were not about what AI will be able to do. They were about what can be checked.

Last issue followed the battle for code. This one goes a step further, to the question that comes after any code: how do we know it does what we claim?

On 10 November Google DeepMind said AI agents had mathematically proved key security properties of the real libsignal code, at the machine-code level. The same day Anthropic put numbers into an amended IPO filing that it must stand behind in front of a regulator. OpenAI described a kill switch that stops models within 35 seconds and fails closed. In five days CISA added eleven entries to its catalogue of actively exploited vulnerabilities.

What they share is that instead of a promise they offer something that can be checked: a proof in Lean, a prospectus with legal liability, a measured stop time, confirmed exploitation.

So on every page of this issue we separate what is proven from what a vendor claims and from what is only planned. Model test results are marked as vendor claims, and dates are publication dates.

We close the issue on Friday, 13 November, at 5:37 CET. Enjoy the read.

NUMBER

6.5 M

lines of Lean proof for libsignal. Google DeepMind

NUMBER

35 s

maximum time to stop in OpenAI Control Room. OpenAI

NUMBER

11

new CISA KEV entries from 6 to 10 November. CISA

SOURCES

1 **Google DeepMind**, Introducing AlphaProof Code: a breakthrough in verifying complex software at depth, 10 Nov 2026 deepmind.google

2 **SEC EDGAR**, Anthropic, PBC, Form S-1/A (Amendment No. 2), 10 Nov 2026 sec.gov

3 **OpenAI**, Control Room: A central system to track and stop model workloads, 9 Nov 2026 openai.com

4 **CISA**, Known Exploited Vulnerabilities, JSON feed, 10 Nov 2026 cisa.gov

COVER STORY / FIVE PROOFS

Proof instead of promises



What changes when AI does not have to persuade but can prove? Five questions guide the issue, from a mathematical proof for Signal to a stock market prospectus.

1 CODE
Can real code be proven secure? 06
AlphaProof Code and Signal.

2 MODELS
What are vendor-published numbers worth? 11
Fable 5.2, Mythos 5.2, Gemini 4.

3 CONTROL
What does it mean for a kill switch to work when everything else fails? 22
Control Room, KEV, patches.

4 MONEY
What must Anthropic write under legal liability? 28
S-1/A, CoreWeave, Nebius.

5 WORK
What does an agent a company can trust look like? 42
Claude, OpenAI, Use this guide.

Trust is no longer asked for up front. It is built from things that can be checked: proofs, measurements and documents with liability attached.

COVER STORY / GOOGLE DEEPMIND / 10 NOV 2026

Signal, proven down to machine code

AI agents used Lean to prove key security properties of the real Signal protocol implementation, on the code that actually runs on devices.

On 10 November Google DeepMind introduced AlphaProof Code, a framework in which people and AI agents formally verify software together. The first target was libsignal, the reference implementation of the Signal protocol, which protects conversations in the Signal app and, in variants, in WhatsApp, Messenger and Google Messages.

According to the technical paper, the agents wrote about 6.5 million lines of Lean proof, checked by Lean's kernel. They rest on about 26,000 lines of specifications and AArch64 processor semantics, and those lines are what people must review and trust.

The proofs cover the post-quantum PQXDH handshake and the Triple Ratchet: forward secrecy, authentication, post-compromise security and handshake secrecy. DeepMind says this is the first time the implementation of such a widely used protocol has been verified by any method, with or without AI.

NUMBER

6.5 M

lines of proof checked by Lean. Google DeepMind

AT A GLANCE

Target
libsignal v0.105.4

Level
Machine code, AArch64

Duration
4 weeks

Agents
Gemini Carbon

Human review
about 1% of code

SOURCES

1 **Google DeepMind**, Introducing AlphaProof Code: a breakthrough in verifying complex software at depth, 10 Nov 2026 deepmind.google

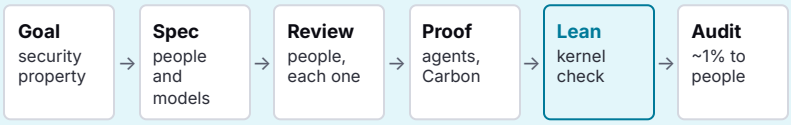
2 **Google DeepMind**, AlphaProof Code: Verifying Signal's Security at the Machine-Code Level, 10 Nov 2026 google-deepmind.github.io

COVER STORY / HOW IT WORKS / LEAN

Specification, agents, Lean's kernel

People decide what to prove. Agents prove it. Lean accepts only what passes the check.

THE ALPHAPROOF CODE WORKFLOW



Based on Google DeepMind's description and technical paper.

The framework splits the work like a research team. Researchers working with several models turn Signal's security goals into formal mathematical specifications, and a person reviews every specification. Only then does a team of autonomous agents, powered by Gemini Carbon according to DeepMind, write the proofs.

The key decision is that a proof is accepted only if Lean's kernel checks it against the approved specification. An agent cannot declare success, and a broken proof does not pass.

At the end the agents review the proofs themselves and flag the parts a person should look at. DeepMind puts that at about one percent of the code, and the whole job took four weeks.

Why machine code? Most earlier verification stopped at the protocol design or at versions of the code written specially for verification. Bugs often hide in the gap between design and implementation, so a proof over the code that actually runs closes exactly that gap. That is why the Rust compiler is not in the trusted base: if it gets something wrong, the proof fails.

“Having the core Signal protocol rigorously verified gives us even greater confidence that our code delivers on its privacy guarantees.”

Ehren Kret, CTO of Signal, quote shortened, Google DeepMind

SOURCES

1 **Google DeepMind**, Introducing AlphaProof Code: a breakthrough in verifying complex software at depth, 10 Nov 2026 deepmind.google

2 **Google DeepMind**, AlphaProof Code: Verifying Signal's Security at the Machine-Code Level, 10 Nov 2026 google-deepmind.github.io

3 **Google DeepMind**, Formal Verification of Signal, technical paper (PDF), 10 Nov 2026 deepmind.google

COVER STORY / LIMITS / TECHNICAL PAPER

What is proven, and what is not

A proof holds exactly as far as its assumptions do. The authors list them themselves, and that is the most valuable part of the paper.

PROOF SCOPE ACCORDING TO THE TECHNICAL PAPER

AREA	STATUS
PQXDH handshake	**Proven, symbolic and computational model**
Triple Ratchet	**Proven: forward secrecy, post-compromise security**
SPQR post-quantum healing	*Not proven, the computational theorem is classical*
Many sessions, variable length	*Simplified: one session, fixed lengths*
Groups, sealed sender, server	*Out of scope*
Cryptographic primitives	*Idealised contracts, not verified*

Highlighted cells are proven, grey ones are assumptions or out of scope. Source: Google DeepMind.

In the discussion section the authors themselves say where the proof stops. The computational theorem is classical, and the ratchet theorems let the attacker choose SPQR keys. So they do not prove the post-quantum healing SPQR is designed to provide. The models are simplified, for example to one session per handshake, and group messaging, sealed sender and server-side protocols are not covered.

The trusted base consists of the AArch64 semantics, Lean's model of a Linux process, contracts for a few external routines, properties of the cryptographic primitives, the build tooling and Lean itself. The authors write that the base can be reduced and the scope extended in future work.

For readers the message is simple: the sentence "Signal is proven secure" is not accurate. What is accurate is that specific properties of a specific version of the code are proven under stated assumptions.

NUMBER

26,000

lines people must review: specifications and semantics. Google DeepMind

SOURCES

1 **Google DeepMind**, AlphaProof Code: Verifying Signal's Security at the Machine-Code Level, 10 Nov 2026
google-deepmind.github.io

COVER STORY / CHECK / OLD BUGS

Two old bugs as an exam

A proof that code is correct is only worth something if the same method would catch code that is not.

WHAT YOU MUST TRUST, AND WHAT LEAN CHECKS

CHECKED BY LEAN	**6.5M lines of proof**	models, refinement, games, function contracts
REVIEWED BY PEOPLE	**about 26,000 lines**	specifications and AArch64 semantics
ASSUMPTION	**Crypto primitives, tools, Lean**	idealised contracts and the Lean kernel

Simplified view of the trusted computing base section of the paper.

To check that the specifications were not too weak, the researchers put back into an isolated copy two bugs that Cheval and colleagues had found earlier and that libsignal had already fixed. The first is a so-called weak state: a failed message could store the signed prekey secret. The second is a small subgroup check, where the PQXDH responder does not validate the initiator's key, so a replayed handshake creates a second session with the same secret.

For both buggy versions there are complete Lean proofs that the security theorem does not hold. In other words, the same tool that says the correct code is secure also shows that the buggy code is not.

This matters because the most common objection to formal verification is that it proves the wrong thing. The test does not rule that out, but it shows the specification catches at least two known real-world bugs.

REINJECTED BUGS

Weak state
fix cd36118

Small subgroup
fix 5293caa

Result
proof it is insecure

SOURCES

1 **Google DeepMind**, AlphaProof Code: Verifying Signal's Security at the Machine-Code Level, 10 Nov 2026 google-deepmind.github.io

2 **GitHub**, google-deepmind/code-verification, 10 Nov 2026 github.com

COVER STORY / CONTEXT / THE COST OF PROOF

Years of work become weeks

Formal proof was long reserved for the most critical systems because it cost too much. This autumn the price is falling in code and in mathematics alike.

SCALE AND DURATION OF LARGE FORMALISATIONS

PROJECT	WHAT	DURATION
seL4	microkernel, people	*20 person-years*
CompCert	C compiler, people	*6 person-years*
Fermat's theorem	Claude, 13M lines of Lean	**11 days, September**
AlphaProof Code	libsignal, 6.5M lines	**4 weeks**

Human effort and agent runtime are not the same measure, so the rows are not directly comparable. Source: Google DeepMind, Anthropic.

The technical paper recalls that verifying the seL4 microkernel took about twenty person-years and the CompCert C compiler about six. Even then the code was written with verification in mind, while libsignal is existing production code.

The same trend shows up in mathematics. On 4 September Anthropic said Claude had, largely on its own, written the first complete computer-checked proof of Fermat's Last Theorem in 11 days: 13 million lines of Lean. In the 9 November announcement of Fable 5.2 the company says the model formalised the same proof in 30 hours during testing, which is a vendor claim.

The AlphaProof Code authors also warn about cost: machine-code-level proofs are long, and much of them deals with registers and memory. Working at a higher language level would be cheaper, but would require more trust in the compiler.

NUMBER

11 days

for Claude's proof of Fermat's theorem, published 4 September. Anthropic

WHY THIS IS HERE

The Fermat proof is from September. We cite it as context, not as news from this week.

SOURCES

1 **Google DeepMind**, AlphaProof Code: Verifying Signal's Security at the Machine-Code Level, 10 Nov 2026 google-deepmind.github.io

2 **Anthropic**, Formalizing Fermat's Last Theorem, 4 Sep 2026 anthropic.com

3 **Anthropic**, Introducing Claude Fable 5.2 and Claude Mythos 5.2, 9 Nov 2026 anthropic.com

MODEL WATCH / ANTHROPIC / 9 NOV 2026

Fable 5.2 and Mythos 5.2: one model, two doors

Anthropic refreshed its models for the hardest technical work. The price is lower, output is faster, and access depends on what you need the model for.

According to Anthropic, Claude Fable 5.2 and Claude Mythos 5.2 are the same underlying model. Fable has strong cybersecurity and biology safeguards and is available to all users. Organisations doing defensive security or biomedical research can apply for Mythos through the Life Sciences and Cyber Verification programmes.

The price is 8 dollars per million input tokens and 40 dollars per million output tokens, with output up to 30% faster than version 5.1. An optional fast mode speeds up generation by up to 2.5 times and costs 16 and 80 dollars.

Anthropic still recommends starting with Opus 5.5 for most coding and office work and moving up to Fable for the most open-ended and demanding tasks. Fable 5.2 can also act as an advisor to smaller models or as an orchestrator that hands them work.

NUMBER

8 / 40 $

per million input and output tokens. Anthropic

AVAILABILITY

`Fable 5.2`
all users

`Mythos 5.2`
by approval

`Fast mode`
$16 / $80, up to 2.5x

`GitHub Copilot`
from 9 Nov

SOURCES

1 **Anthropic**, Introducing Claude Fable 5.2 and Claude Mythos 5.2, 9 Nov 2026
anthropic.com

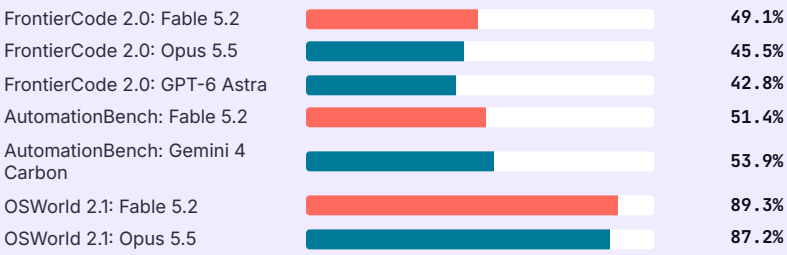
2 **GitHub Changelog**, Claude Fable 5.2 in GitHub Copilot, 9 Nov 2026
github.blog

MODEL WATCH / BENCHMARKS / VENDOR

What the numbers show, and whose they are

The table is impressive, but it was compiled and measured by the company selling the model. One row still shows Fable does not win everywhere.

SELECTED RESULTS FROM ANTHROPIC'S ANNOUNCEMENT, %



Vendor claims, not independently verified. Results at the highest available effort. Source: Anthropic.

On FrontierCode 2.0, which measures whether open-source maintainers would merge a model's changes, Fable 5.2 scores 49.1% at maximum effort. Anthropic also says that at high effort it scores 45.3%, roughly what Opus 5.5 reaches at its maximum, for about half the cost per task.

In the same table Gemini 4 Carbon does better on AutomationBench, 53.9% against 51.4%. That row deserves attention precisely because Anthropic published it even though a competitor wins it.

The biggest jump is on the Remote Labor Index, from 17.9% for Fable 5.1 to 41.3%. Anthropic also gives an example from its own work: a team of agents moved a production service from Python to Rust, and it now starts about 26 times faster. That is an anecdote, not a measurement.

BENCHMARK STATUS

VENDOR. Every number on this page was published by Anthropic.

NUMBER

41.3 %

Remote Labor Index, Fable 5.2, vendor claim. Anthropic

SOURCES

1 **Anthropic**, Introducing Claude Fable 5.2 and Claude Mythos 5.2, 9 Nov 2026 anthropic.com

MODEL WATCH / GOOGLE / 10 AND 12 NOV 2026

Gemini 4: Carbon for everyone, Hydrogen for speed

In three days Google opened its strongest model to all paying developers and introduced a cheap model meant to work alongside it.

Since 10 November Gemini 4 Carbon has been available through the API in Google AI Studio and in Gemini for Business for all paying developers and companies. The introductory price is 2 dollars per million input and 10 per million output tokens, the same as Argon, and according to a footnote in the announcement it runs until 30 November. Two weeks earlier a group of cyber defenders got Carbon through the Fairwind programme.

Two days later came Gemini 4 Hydrogen, a model for fast everyday work and agentic workflows. It costs 1.50 dollars per million input and 7.50 per million output tokens, flat across the whole one-million-token context. It is available on every AI Studio tier, in the Gemini API, the Gemini app and Search.

Google says Hydrogen scores 51 on the Artificial Analysis Intelligence Index at its default thinking level, against 50 for Argon.

PRICES PER MILLION TOKENS

Carbon
$2 / $10, intro

Hydrogen
1,50 / 7,50 $

Context
1M tokens

NUMBER

57 %

Hydrogen on Terminal-Bench 4.1, against 18.2% for 3.8 Flash, per Google. Google

SOURCES

1 **Google**, Gemini 4 Carbon is now available, 10 Nov 2026
blog.google

2 **Google**, Introducing Gemini 4 Hydrogen, 12 Nov 2026
blog.google

MODEL WATCH / ANALYSIS / NOVEMBER 2026

An expensive planner, a cheap worker

This week Google and Anthropic are selling the same pattern: the strongest model thinks and checks, while a faster, cheaper one does the routine work.

ONE PATTERN AT TWO COMPANIES

GOOGLE

Carbon + Hydrogen

- Orchestrator in Antigravity switches models mid-session
- /swarm, preview: an agent team from one prompt
- Carbon coordinates, Hydrogen works and verifies

↔

ANTHROPIC

Fable + Opus, Sonnet, Haiku

- Fable as an advisor to smaller models
- Fable as an orchestrator that hands out work
- Advice: start with Opus 5.5

Based on company announcements of 9, 10 and 12 November.

Google says Orchestrator, which moves work between Carbon and Hydrogen within one session, came close to Carbon's own success rate on Terminal-Bench 4.0. It used on average 24% less quota and was 19% faster. Those are vendor figures.

For Fable 5.2 Anthropic describes two roles: an advisor that smaller models consult when they get stuck, and an orchestrator that hands simpler tasks to smaller models.

For you, this means price per token is no longer enough for a comparison. What matters more is how many tokens and how many calls to the expensive model a whole task uses. Before choosing, it is worth measuring the real cost per completed task on your own examples, because vendor claims refer to their tests.

NUMBER

24 %

less quota for Orchestrator on Terminal-Bench 4.0, per Google. Google

SOURCES

1 **Google**, Introducing Gemini 4 Hydrogen, 12 Nov 2026
blog.google

2 **Google**, Gemini 4 Carbon is now available, 10 Nov 2026
blog.google

3 **Anthropic**, Introducing Claude Fable 5.2 and Claude Mythos 5.2, 9 Nov 2026
anthropic.com

MODEL WATCH / 6 NOV 2026 → 13 NOV 2026

Models of the week

Models that appeared or expanded from 6 to 13 November. Quality and speed figures are vendor claims.

Claude Fable 5.2

ANTHROPIC · 9 NOV 2026

Refreshed model for the hardest technical work, cheaper and up to 30% faster than 5.1. [Anthropic]

AVAILABILITY	PRICE PER 1M	BENCHMARK	BENCHMARK STATUS
All users, API, Copilot	$8 in, $40 out	FrontierCode 2.0: 49.1%	VENDOR

Claude Mythos 5.2

ANTHROPIC · 9 NOV 2026

Same model as Fable, with fewer limits for cyber and bio work. [Anthropic]

AVAILABILITY	BENCHMARK	BENCHMARK STATUS
By approval, verification programmes	CVE Rediscovery: 89.6%	VENDOR

Gemini 4 Carbon

GOOGLE · 10 NOV 2026

The strongest Gemini, now open to all paying developers. [Google] [GitHub Changelog]

AVAILABILITY	PRICE PER 1M	BENCHMARK	BENCHMARK STATUS
API, Ultra, Copilot	$2 / $10 intro	#1 CursorBench 4.0, per Google	VENDOR

Gemini 4 Hydrogen

GOOGLE · 12 NOV 2026

Fast model for agentic work, on every tier. [Google] [Vercel]

AVAILABILITY	PRICE PER 1M	BENCHMARK	BENCHMARK STATUS
API, app, Search, AI Gateway	1,50 $ / 7,50 $	AA index: 51	VENDOR

Gemini Robotics 2.5

GOOGLE DEEPMIND · 10 NOV 2026

Vision-language-action model for robots. Details on the Lab page. [Google DeepMind]

AVAILABILITY	PRICE	BENCHMARK	BENCHMARK STATUS
Rolling out to partners	*not published*	+30 to 100% success	VENDOR

VENDOR means the result was published by the vendor and has not been independently verified.

WORK / AI AS A BUSINESS PARTNER

AI no longer waits in a chat window.

In four days Anthropic, OpenAI, IBM and Microsoft introduced tools that put AI into spreadsheets, meetings, procurement and company infrastructure. Most are still in beta.

	WORK TOOLS	CONTROL AND PROCUREMENT
9 Nov	**OpenAI** ChatGPT Enterprise in AWS Marketplace AVAILABLE	**OpenAI** Control Room for research jobs ~~INTERNAL SYSTEM~~
10 Nov	**Anthropic** Claude Sheets, Whiteboards, Data BETA, DATA GA **OpenAI** Businesses in ChatGPT BETA	**IBM, Microsoft** Imagine Enterprise VIA IBM CONSULTING
12 Nov	**Anthropic** Claude Meetings BETA, GOOGLE MEET **Anthropic** Claude for Manufacturing PLUGIN	**Anthropic** Managed Agents in your environment PRIVATE BETA

Status as of the cut-off date, 13 November.

WORK / ANTHROPIC / 10 NOV 2026

Claude gets sheets, boards and data

Work done with Claude no longer has to be copied into another program. The team can edit it together, inside Claude.

On 10 November Anthropic released Claude Whiteboards and Claude Sheets in beta, and Claude Dashboards became generally available under a new name, Claude Data. All three run on paid plans, on web, desktop and mobile.

Whiteboards is a shared board where a team brainstorms with Claude, and Claude groups notes by theme and drafts action items. Sheets is a spreadsheet with tabs, formulas and drop-downs, for now with basic functions. For full Excel there is export to .xlsx and CSV.

Claude Data adds certified queries, so the whole team uses the same metric definitions, and live sharing with refreshed numbers. Those two features are still in beta. A viewer sees only the results of queries the author approved and cannot run others with the author's access.

On Enterprise plans Claude Data turns on by default on 19 November, and an admin can change that.

STATUS

Whiteboards
beta

Sheets
beta

Claude Data
generally available

Certified queries
beta

Plans
all paid

CONNECTORS

From the same day interactive connectors open as full apps, for example Adobe, Atlassian, Box and Notion.

SOURCES

1 **Claude**, Introducing Claude Whiteboards, Claude Sheets, and Claude Data, 10 Nov 2026
claude.com

2 **Claude**, Your apps can now open in Claude, and Claude can work in them, 10 Nov 2026
claude.com

WORK / ANTHROPIC / 12 NOV 2026

Claude in the meeting: who may listen

Claude Meetings joins a video call as a visible participant. The default rules are more interesting than the feature itself.

Claude Meetings has been in beta since 12 November for Team and Enterprise plans, for now only on Google Meet. Claude follows the conversation, answers questions, in writing by default, and during the meeting builds documents, boards and notes with action items. With permission it can also use tools you have connected.

WHAT IS ON BY DEFAULT

When it joins, Claude shows up as a participant and posts a notice in the call chat. It keeps a transcript, a summary and screen descriptions, but not audio. If someone from outside the organisation is invited, Claude stays out by default, and it leaves if such a guest joins later. An admin can change this.

WHAT YOU HAVE TO DO

Anthropic explicitly asks you to check that everyone on the call is fine with being transcribed before you add Claude or turn on auto-join. During the beta the feature is not available to organisations with zero data retention, a HIPAA configuration or customer-managed encryption keys.

AT A GLANCE

Status
beta

Plans
Team, Enterprise

Platform
Google Meet

Audio
not stored

Outside guests
out by default

SOURCES

1 **Claude**, Introducing Claude Meetings, 12 Nov 2026
claude.com

WORK / ANTHROPIC / 12 NOV 2026

Agents inside your perimeter

A company can now keep the whole agent service in-house, from sandbox to session history. The four design principles matter more than the location.



Based on Anthropic's description of 12 November.

Self-hosted Claude Managed Agents has been in private beta since 12 November. You run every part of the service, in your own cloud or on premises: the harness, the sandboxes, the credentials and the session history. Model requests go wherever you choose, directly to Anthropic, through Amazon Bedrock, Google's Agent Platform, Microsoft Foundry or your own gateway.

Anthropic explains that an agent running unattended in your code cannot rely on a person checking every step, so the limits have to be built into the system. That is why the same principles apply no matter who hosts the service.

For an agent that fixes broken builds, the allowlist might be just your git host and your package registry. We use these four principles as the basis for the Use this guide at the end of the issue.

STATUS

Stage
private beta

Session history
your database

Models
Anthropic, Bedrock, Google, Foundry

SOURCES

1 **Claude**, Operate Claude Managed Agents inside your security perimeter, 12 Nov 2026
claude.com

WORK / OPENAI / 9 AND 10 NOV 2026

Businesses enter ChatGPT, ChatGPT enters AWS

In two days OpenAI opened ChatGPT to businesses from both sides: as a place where customers find them and as a tool bought through an existing AWS contract.

Businesses in ChatGPT has been in beta since 10 November. Through Business Hub a company can claim its profile, update its details and add a product feed with prices and availability. In return it gets aggregated insights: how often it was mentioned, how many clicks its website got and which topics it shows up around. OpenAI says it does not share user conversations with businesses and that insights appear only above certain thresholds.

US businesses in e-commerce and online services can also build a Business Agent, which a customer can talk to separately from their own ChatGPT conversation. The agent can also be a paid destination for ChatGPT ads, clearly labelled.

A day earlier, on 9 November, ChatGPT Enterprise arrived in AWS Marketplace. Companies pay for it through their AWS account, and the spend can count toward existing AWS commitments.

NUMBER

1.2 bn

people use ChatGPT each week, per OpenAI. OpenAI

STATUS

Business Hub
beta

Business Agents
US, selected categories

AWS Marketplace
available

SOURCES

1 **OpenAI**, Introducing Businesses in ChatGPT, 10 Nov 2026
openai.com

2 **OpenAI**, ChatGPT Enterprise now available in AWS Marketplace, 9 Nov 2026
openai.com

WORK / IBM, MICROSOFT, ANTHROPIC / 10 AND 12 NOV 2026

Agents built for an industry

A general assistant is not enough for a bank or a factory. Two announcements this week offer agents prepared for specific jobs.

On 10 November IBM and Microsoft introduced Imagine Enterprise, a joint platform for building, running and governing agents for industry-specific processes. IBM Consulting brings more than 6,000 agents and other assets it has already built and used, and Microsoft brings its AI, cloud and security tools, including Foundry, GitHub Copilot, Copilot Studio and Fabric.

The first industry frameworks cover financial services, insurance and retail. IBM Consulting can deploy and run the platform, and self-management is also possible. The announcement gives no measured production results.

On 12 November Anthropic announced Claude for Manufacturing, a plugin with 17 skills and four connectors for manufacturers. It is meant for drafting quotes, quality reports and supplier requests from the systems a factory already runs.

What they share is that a package is being sold: a model, ready-made workflows and someone to roll them out.

NUMBER

6,000+

IBM Consulting agents and assets in Imagine Enterprise. IBM

NUMBER

17

skills in the Claude for Manufacturing plugin. Claude

SOURCES

1 **IBM**, Imagine Enterprise, 10 Nov 2026
ibm.com

2 **CIO Dive**, Microsoft, IBM team up on AI agent platform, 10 Nov 2026
ciodive.com

3 **Claude**, Introducing Claude for Manufacturing, 12 Nov 2026
claude.com

SECURITY / OPENAI / 9 NOV 2026

A kill switch that works when everything else fails

OpenAI described a system that stops models across its research infrastructure within 35 seconds at most. If the system itself fails, the models stop.

Control Room is a central system that tracks every research job in which models generate output: reinforcement learning, evaluations, sampling and data processing. OpenAI says it is already deployed across its research infrastructure, where, it says, the need is greatest.

The system works with misalignment monitoring. Every large reinforcement learning job has a monitor, another model that watches what the trained model does and raises an alert when it sees a problem. If the monitor decides a job should stop, Control Room enforces it.

Two numbers and one decision matter most. From the order to the moment further output is no longer possible takes at most 35 seconds, because of an authorisation cache, and the goal is zero. Every research job with agentic generation must have an approved safety case. And the system fails closed: if Control Room goes down, every request is refused.

NUMBER

35 s

maximum from order to stopped output. OpenAI

AT A GLANCE

Scope
research

Failure
everything refused

External serving
separate system

SOURCES

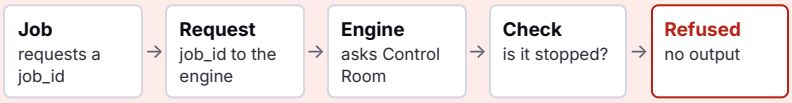
1 **OpenAI**, Control Room: A central system to track and stop model workloads, 9 Nov 2026 openai.com

SECURITY / HOW IT WORKS / FAIL CLOSED

Why the switch sits in the inference layer

Research job code changes all the time and is a poor place for a safety control. The places where models actually run are few and uniform.

ONE REQUEST'S PATH THROUGH CONTROL ROOM



Based on OpenAI's description. If the service does not answer, the result is the same as a stop.

At the start a job gets a job_id from Control Room. Every time it asks a model for tokens, it sends that ID to the engine, the program that runs the model on GPUs. Before every request the engine asks Control Room whether the job has been stopped. If it has, the engine refuses the request.

OpenAI explains the choice of location: job code is varied, decentralised and constantly changing, while the paths inference goes through are very few. If the control lives there, it applies to all code, including buggy code. The company also runs agentic and programmatic sweeps that look for engines with the check switched off, by accident or on purpose.

The lesson is the same for much smaller systems: do not put the control in code that an agent or a team changes often, but at the bottleneck every call has to pass through. That is also the first step of the Use this guide.

WHAT IT DOES NOT COVER

Control Room is for research. For public models OpenAI says it already has a centralised way to take deployments down quickly.

SOURCES

1 **OpenAI**, Control Room: A central system to track and stop model workloads, 9 Nov 2026 openai.com

SECURITY / OPENSSH, MICROSOFT, OPENAI / 5 TO 10 NOV 2026

AI finds bugs. So do attackers.

Three different sources in the same week describe the same thing: bugs found by AI are soon found by someone else too. Patch speed becomes part of security.

In the release notes for OpenSSH 10.6p2 on 5 November the team writes that it has recently received a large number of reports, many from AI models or made with their help. They also noticed that a bug found by AI is often found independently by another researcher. The conclusion: attackers who do not report bugs can probably find them too, so OpenSSH will ship fixes more often instead of waiting for a planned release.

On 10 November Microsoft fixed an Outlook vulnerability that had already been exploited. The acknowledgements for that vulnerability say it was reported by Benoît Sevens, using Claude.

On 6 November OpenAI announced a partnership with Hacktron, whose researchers build AI agents for code review and penetration testing. Together they will test how OpenAI systems connect to third-party services.

“Adversaries who do not report bugs to OSS projects are likely to be able to discover these bugs too.”

OpenSSH team, 10.6p2 release notes, OpenSSH

SOURCES

1 **OpenSSH**, OpenSSH 10.6p2 release notes, 5 Nov 2026 openssh.org

2 **Microsoft MSRC**, CVE-2026-124194, Microsoft Outlook Spoofing Vulnerability, 10 Nov 2026 msrc.microsoft.com

3 **OpenAI**, Partnering with Hacktron to find vulnerabilities in OpenAI's systems, 6 Nov 2026 openai.com

PATCH NOW / 6 NOV 2026 → 13 NOV 2026

Patch immediately

CISA added every item on this page to its catalogue of actively exploited vulnerabilities between 6 and 10 November. Patches exist.

Barracuda ESG, WAF and others

CVE-2026-116606, CVE-2026-111560

CISA KEV 10 NOV · DEADLINE WAS 13 NOV

Pre-auth authorisation bypass and command injection. Chained, they give root on the device. On 9 November Barracuda confirmed attacks on a small number of organisations.

PATCH
Latest firmware and patches, e.g. BNSF-40425

ADMIN
Take the Admin UI off the internet

SOURCE
trust.barracuda.com
cisa.gov

OpenSSH on Linux

CVE-2026-123158

CISA KEV 9 NOV · AUTH BYPASS

According to OpenSSH, it affects some Linux systems with musl libc and PAM disabled. The root cause is outside sshd, but it is exploitable through it.

PATCH
OpenSSH 10.6p2

ADMIN
Check Alpine and other musl distros

SOURCE
openssh.org
cisa.gov

Microsoft Outlook

CVE-2026-124194

CISA KEV 10 NOV · PATCH TUESDAY

Opening a crafted message can send the NTLMv2 hash to an attacker, without the usual warning. The preview pane is not an attack vector.

PATCH
November Office updates

CISA DEADLINE
24 Nov

SOURCE
msrc.microsoft.com
cisa.gov

SonicWall SonicOS SSLVPN

CVE-2026-19868

CISA KEV 6 NOV · MFA

A user with valid credentials for another sign-in method can bypass SAML, including the MFA that SAML enforces.

PATCH
SNWLID-2026-0014

CISA DEADLINE
20 Nov

SOURCE
psirt.global.sonicwall.com
cisa.gov

SECURITY / CISA KEV / 6 TO 10 NOV 2026

Eleven confirmed attacks in five days

A KEV entry means CISA has evidence of active exploitation. It is the shortest list of what to check first.

NEW CISA KEV ENTRIES, CATALOGUE VERSION 2026.11.10

ADDED	PRODUCT	TYPE	DUE
6 Nov	SonicWall SonicOS SSLVPN	MFA bypass	20 Nov
9 Nov	GPUStack	unauth admin	**12 Nov**
9 Nov	OpenSSH (Linux)	auth bypass	**12 Nov**
9 Nov	SignalWire FreeSWITCH	code injection	**12 Nov**
9 Nov	IBM Aspera Faspex	privilege escalation	**12 Nov**
10 Nov	Microsoft Outlook	spoofing	24 Nov
10 Nov	Openwave Webtop Transcoder	resource transfer	24 Nov
10 Nov	Cisco BroadWorks	XXE	24 Nov
10 Nov	Barracuda (2 CVE)	bypass and commands	**13 Nov**
10 Nov	Google Chromium	use-after-free	24 Nov

The deadline applies to US federal agencies but is a good guide for everyone. Highlighted deadlines have passed or expire on publication day. Source: CISA.

From 6 to 10 November CISA added eleven vulnerabilities, and Barracuda has two entries in one table row. Six of the eleven have a deadline of just three days, which CISA uses for the most urgent cases.

It is striking how many entries touch AI and communications infrastructure: GPUStack, a tool for managing GPU clusters and models, the FreeSWITCH VoIP server and OpenSSH. For the Chromium flaw CISA notes that it also affects other browsers built on it, such as Edge and Opera.

The entries refer to directive BOD 26-04, which sets patch priorities by risk.

NUMBER

6 of 11

entries with a three-day deadline. CISA

SOURCES

1 **CISA**, Known Exploited Vulnerabilities, JSON feed, 10 Nov 2026 cisa.gov

2 **CISA**, Known Exploited Vulnerabilities Catalog, version 2026.11.10, 10 Nov 2026 cisa.gov

3 **Google Chrome Releases**, Stable Channel Update for Desktop, 10 Nov 2026 chromereleases.googleblog.com

SECURITY / BARRACUDA, MICROSOFT / 10 NOV 2026

Two chains: root on a box and a hash from a message

The most serious items of the week are not single bugs but combinations: one opens the door, the other hands over control.

BARRACUDA: BYPASS PLUS COMMAND

Barracuda published its advisory on 2 November. CVE-2026-116606 in the Email Security Gateway, CVSS 9.3, lets an attacker who reaches the admin interface call the configuration import without signing in and restore a crafted backup. On 5 November Barracuda corrected its own rating: no authentication is required. CVE-2026-111560, CVSS 8.6, is command injection through a time zone field, in six products. Chained, they give root access.

Alongside the patch, the best defence is to keep the admin interface off the internet. If firmware is several versions behind, more than one update in a row may be needed.

OUTLOOK: A HASH WITHOUT A WARNING

CVE-2026-124194 is the only flaw in Microsoft's November updates marked as already exploited. The attacker sends a crafted message. When the user opens it, Outlook connects to the attacker's location without the external-content warning and can disclose the user's NTLMv2 hash. Microsoft rates the severity as moderate, but exploitation has already been detected.

NUMBER

9.3 CVSS

Barracuda CVE-2026-116606, per Barracuda. Barracuda

PATCH TUESDAY

Action1 counts about 1,094 fixes in the November release, with one zero-day. Counts differ by source.

SOURCES

1 **Barracuda**, Pre-Authentication Authorization Bypass and Command Injection Vulnerabilities, 9 Nov 2026 trust.barracuda.com

2 **Microsoft MSRC**, CVE-2026-124194, Microsoft Outlook Spoofing Vulnerability, 10 Nov 2026 msrc.microsoft.com

3 **Qualys**, Microsoft Patch Tuesday, November 2026 Security Update Review, 10 Nov 2026 blog.qualys.com

4 **Action1**, Patch Tuesday November 2026, 10 Nov 2026 action1.com

MONEY / ANTHROPIC / 10 NOV 2026

Anthropic heads to market as ANT



Anthropic launched its roadshow and, in an amended SEC filing, set a price range of 70 to 78 dollars per share. These are numbers the company signs with legal liability.

On 10 November Anthropic launched the roadshow for its initial public offering of Class A shares. The offering covers 825 million shares, of which 801.35 million are sold by the company and 23.65 million by existing shareholders. Underwriters also have an option on another 123.75 million shares. The range is 70 to 78 dollars, and the company has applied to list on the Nasdaq Global Select Market under the symbol ANT.

At the same time Google, NVIDIA, AMD, Broadcom, Micron, Microsoft and TSMC will buy a total of 26 billion dollars of Class N shares at the offering price. Those shares carry no vote. Including them, estimated net proceeds at the midpoint are about 84.6 billion dollars.

According to TechCrunch, CNN and Fortune, pricing is expected on 18 November and trading on 19 November. Media put the implied value at up to about 1.4 trillion dollars, but those figures are not from the SEC filing.

NUMBER

70–78 $

price range per Class A share. Anthropic

SOURCES

1 **Anthropic**, Anthropic launches roadshow for proposed initial public offering, 10 Nov 2026 anthropic.com

2 **SEC EDGAR**, Anthropic, PBC, Form S-1/A (Amendment No. 2), 10 Nov 2026 sec.gov

3 **TechCrunch**, Anthropic sets IPO price range of $70 to $78 per share, 10 Nov 2026 techcrunch.com

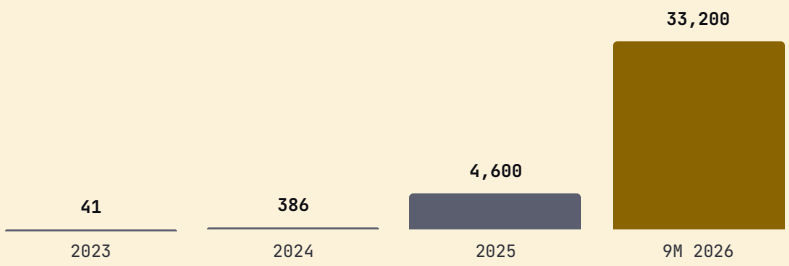
4 **CNN**, We've never seen anything like the Anthropic IPO, 10 Nov 2026 cnn.com

5 **Fortune**, Anthropic IPO price range, 10 Nov 2026 fortune.com

MONEY / S-1/A / PROSPECTUS NUMBERS

Revenue grows, but not like this forever

The prospectus shows growth from 41 million to 33.2 billion dollars in under three years. The most important sentence next to those numbers is the company's own warning.



2026 covers the first nine months, to 30 September. Source: SEC EDGAR.

According to the prospectus, revenue was 41 million dollars in 2023, 386 million in 2024 and 4.6 billion in 2025. In the first nine months of 2026 it reached 33.2 billion, up 1,206% on the same period a year earlier. The company immediately adds that it does not expect revenue to keep growing at a similar pace and cannot guarantee that it will grow at all.

In the third quarter Anthropic had an operating loss of 567 million dollars, after operating income of 974 million in the second quarter. The prospectus names compute costs as the main reason.

No single customer brought in more than 10% of revenue, but the top five together account for about 25%. Many of them have no long-term contracts and can cut spending at any time.

NUMBER

1,206 %

nine-month revenue growth, year on year. SEC EDGAR

NUMBER

567 $M

operating loss in the third quarter. SEC EDGAR

SOURCES

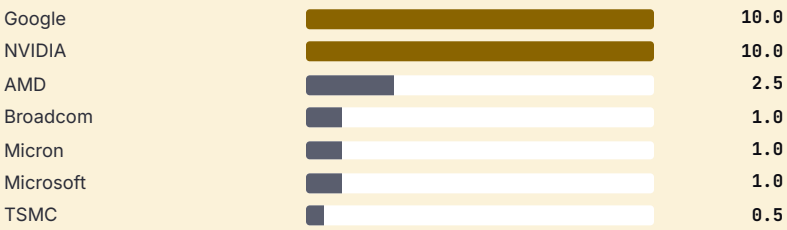
1 **SEC EDGAR**, Anthropic, PBC, Form S-1/A (Amendment No. 2), 10 Nov 2026
sec.gov

MONEY / S-1/A / COMMITMENTS

900.9 billion dollars of commitments

The biggest number in the prospectus is not revenue or valuation but the total of compute contracts the company has already signed.

WHO IS BUYING CLASS N SHARES, $BN



Amounts at 74 dollars, the midpoint of the range. Total 26 billion. Source: SEC EDGAR.

According to the prospectus, as of and after 30 September Anthropic had contracted purchases and leases worth about 900.9 billion dollars in total, almost all for compute and infrastructure. About 451.6 billion is expected to be spent by the end of 2030, and the rest mostly from 2031 to 2045. About 70.7 billion falls due in the first twelve months.

For comparison, on 30 September the company had 146.6 billion dollars of liquidity, before proceeds from the offering.

The list of Class N buyers tells the same story: almost all of them supply chips, memory or cloud. Part of the offering money flows back to the very companies Anthropic buys capacity from.

WHEN COMMITMENTS FALL DUE

12 months	about $70.7bn
to 2030	about $451.6bn
2031. do 2045.	mostly the rest

Amounts are cumulative, from the prospectus.

SOURCES

1 **SEC EDGAR**, Anthropic, PBC, Form S-1/A (Amendment No. 2), 10 Nov 2026
sec.gov

MONEY / COREWEAVE / 10 NOV 2026

CoreWeave: record revenue, record debt

CoreWeave raised its forecast again. The same report shows what it costs to build that growth.

CoreWeave reported third-quarter revenue of 3.661 billion dollars, about 168% more than a year earlier, and an adjusted EBITDA margin of 62%. The net loss was 709 million. In the quarter alone the company spent 9.44 billion on property and equipment, and 23.56 billion over nine months.

The 2026 outlook was raised to revenue of 13.0 to 13.3 billion and capex of 37 to 39 billion. Contracted power exceeds 4.6 gigawatts, with 1.7 active.

Contracted but not yet delivered revenue stands at about 132.5 billion dollars. That is what gives confidence in the growth, and debt of about 46.9 billion is what makes it risky.

NUMBER

62 %

adjusted EBITDA margin. CoreWeave

SOURCES

1 **CoreWeave**, CoreWeave Reports Strong Third Quarter 2026 Results, 10 Nov 2026 investors.coreweave.com

MONEY / NEBIUS / 12 NOV 2026

Nebius: customers pay for the build up front

Nebius trimmed this year's outlook because of site delays and announced capex of more than 40 billion dollars for next year.

In its third-quarter shareholder letter Nebius reports revenue of 981 million dollars, 571% more than a year earlier. Since the start of the quarter it has signed contracts worth 12 billion in total.

The financing is the most interesting part. Customer prepayments for contracts signed in the third quarter covered more than 100% of the related investment, against 50 to 60% in the second quarter. The company expects prepayments and financing against contracts with investment-grade customers to cover 75 to 85% of capex in 2027.

There is bad news too. Because delivery dates at several leased sites moved into the first half of 2027, Nebius now expects 2026 revenue at the low end of its 3.0 to 3.4 billion range. The 7 billion annual run rate it expected in December will arrive around February.

For 2027 it guides to revenue of 12 to 13 billion, capex of 40 to 45 billion and pre-tax breakeven by the end of that year.

NUMBER

981 $M

third-quarter revenue, +571%. Nebius Group / SEC

NUMBER

40 to 45 $bn

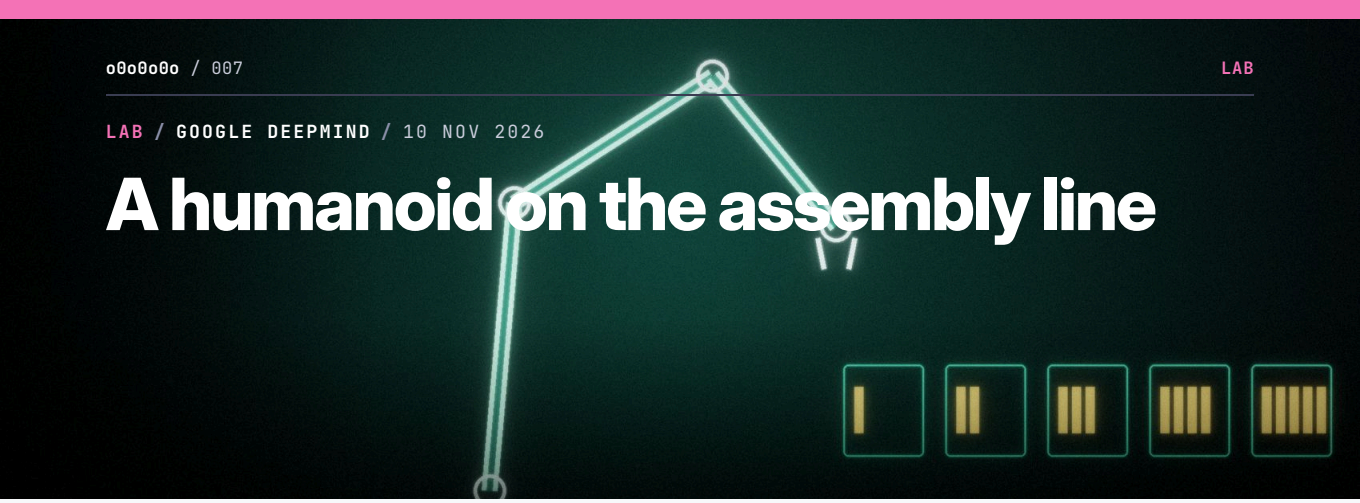
planned capex for 2027. Nebius Group / SEC

SOURCES

1 **Nebius Group / SEC**, Letter to shareholders, Q3 2026 (Exhibit 99.2), 12 Nov 2026
sec.gov

LAB / GOOGLE DEEPMIND / 10 NOV 2026

A humanoid on the assembly line



Gemini Robotics 2.5 was tested outside the lab, on production lines. The 96% result is a proof of concept, not series production.

On 10 November Google DeepMind introduced Gemini Robotics 2.5, a model that combines vision, language and action. The company tested it on more than 200 different tasks, from serving food with tongs to tying a bow and pipetting liquid.

The most concrete test was in Mercedes-Benz plants in Kecskemét, Hungary, and Sindelfingen, Germany. There Apptronik's Apollo 2 humanoid robot did kitting on its own: it took metal brackets from parts bins and placed them in trays in the exact order needed for each vehicle. The success rate was 96%.

The model works together with the Gemini Robotics ER reasoning model, which tracks the workflow, while Gemini Robotics 2.5 controls the hands. For now it is rolling out to robotics partners.

NUMBER

96 %

kitting success in the proof of concept. Google DeepMind

AT A GLANCE

Tasks
200+

Robot
Apptronik Apollo 2

Hand
5 fingers, 22 DoF

Status
partners

SOURCES

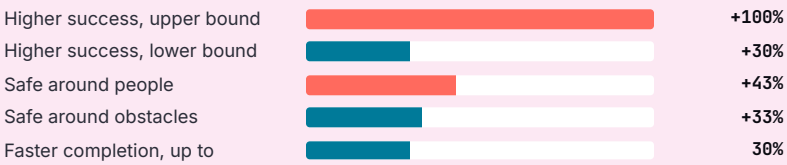
1 **Google DeepMind**, Gemini Robotics 2.5 accelerates real-world general dexterity, 10 Nov 2026
deepmind.google

LAB / NUMBERS / VENDOR

Dexterity, speed and safety around people

Every number comes from the vendor's own evaluations, but what was chosen for measurement is interesting: not just success, but behaviour around people.

IMPROVEMENTS REPORTED BY DEEPMIND, %



Vendor claims, against previous baselines. Not independently verified. Source: Google DeepMind.

DeepMind says Gemini Robotics 2.5 has a 30 to 100% higher success rate than previous models, depending on the type of task, and completes tasks 20 to 30% faster.

For safety the model uses guidance at run time that steers it around physical hazards. In evaluations this raised safe task completion around obstacles by 33% and around nearby people by 43%.

DeepMind is also releasing ASIMOV-Jailbreak, an open set of several hundred prompts that hide unsafe requests inside harmless-looking scenes. According to the company, Gemini 4 Argon refuses 99.5% of such attempts. The set is open, so others can check this result.

BENCHMARK STATUS

VENDOR. ASIMOV-Jailbreak is open, so others can reproduce it.

SOURCES

1 **Google DeepMind**, Gemini Robotics 2.5 accelerates real-world general dexterity, 10 Nov 2026
deepmind.google

LAB / GOOGLE DEEPMIND / 11 NOV 2026

AlphaGenome 2 reads DNA cell by cell

The new model predicts how genetic variants act in hundreds of human cell types, not just in tissue averages.

Every cell in the body carries the same DNA, but each cell type switches different genes on and off. The first AlphaGenome predicted variant effects in tissue averages. AlphaGenome 2, introduced on 11 November, works at the level of individual cell types and states.

The key change is the architecture. The model is now conditional: alongside a million-letter DNA sequence it receives a target biological state, from a catalogue or from a researcher's own experiment. It was trained on more than 107 million single-cell profiles from humans and mice, and besides gene expression it also predicts DNA methylation.

The model can adapt to new cell data without retraining. It is available through the AlphaGenome API for non-commercial research and in private preview for commercial research.

NUMBER

107 M

single-cell profiles in training. Google DeepMind

“What makes AlphaGenome 2 exciting is how it adapts directly to new single-cell data without retraining.”

Sarah Teichmann, Human Cell Atlas, quote shortened, Google DeepMind

SOURCES

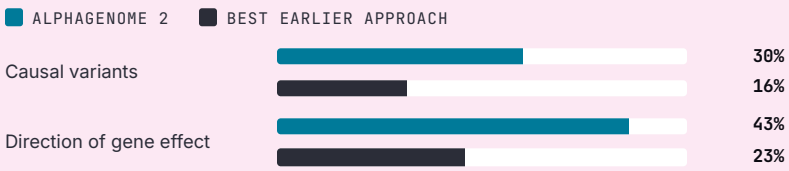
1 **Google DeepMind**, AlphaGenome 2: Deciphering how cells read DNA to better understand disease, 11 Nov 2026 deepmind.google

LAB / RESULTS / TENK10K

Twice as many causal variants

On a map of 28 immune cell types the new model finds almost twice as many variants that really change genes.

TENK10K, 28 IMMUNE CELL TYPES, %



Results reported by Google DeepMind, with a paper. Compared with the first AlphaGenome fine-tuned on single-cell data and other tools. Source: Google DeepMind.

On the TenK10K dataset, a large map of 28 primary immune cell types, AlphaGenome 2 reliably identified 30% of causal variants, against 16% for the best earlier approach. It got the direction of the gene effect right in 43% of cases, against 23%.

DeepMind also gives a concrete example. A recent study showed that a variant in a distant regulatory switch drives inflammatory bowel disease by overactivating the ETS2 gene in gut macrophages. AlphaGenome 2 reproduced that finding computationally: it singled out the causal variant and showed that it boosts ETS2 specifically in macrophages, while other cells and genes stay largely untouched.

That is a reproduction of a known result, not a new discovery, but exactly such tests show whether a model sees biology or just statistics.

NUMBER

68 tissues

in the training data. Google DeepMind

SOURCES

1 **Google DeepMind**, AlphaGenome 2: Deciphering how cells read DNA to better understand disease, 11 Nov 2026 deepmind.google

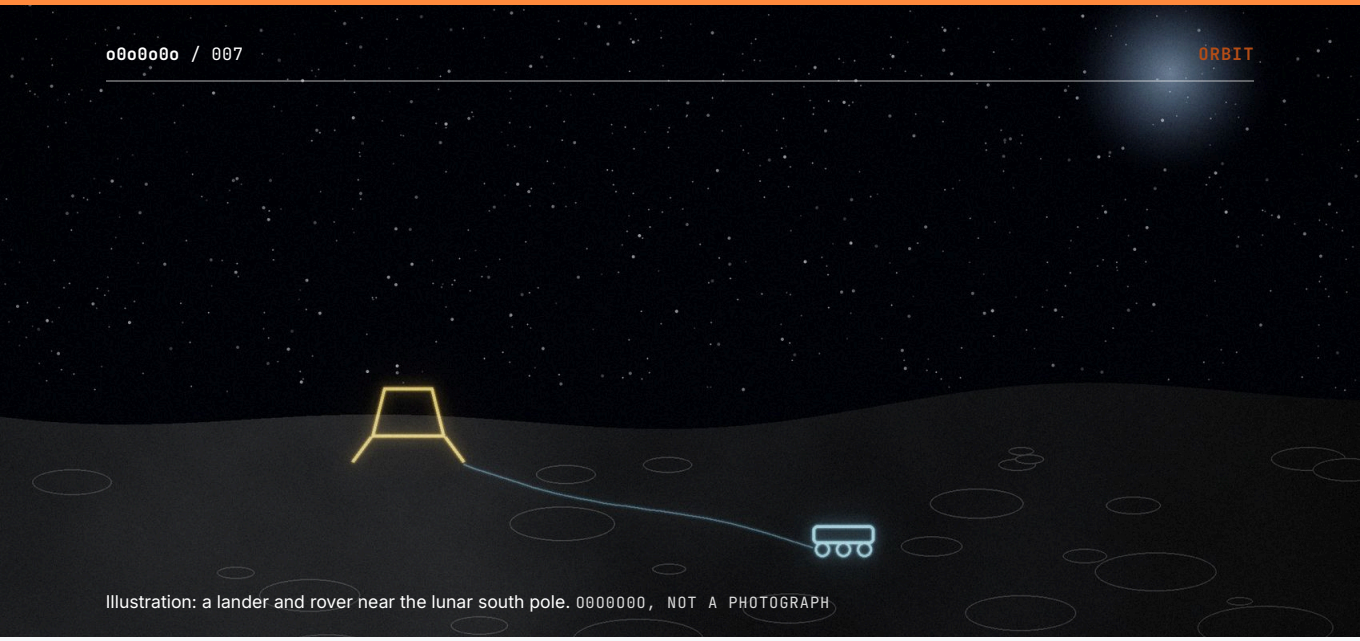


Illustration: a lander and rover near the lunar south pole. 0000000, NOT A PHOTOGRAPH

ORBIT / NASA / 7 NOV 2026

Moon Base becomes a procurement programme

At the MAX POWER event at Kennedy Space Center NASA awarded contracts for lunar transport, power and science. At the same time the mission schedule is slipping.

On 7 November NASA picked eight companies for the CLPS 2 lunar cargo contract: Blue Origin, Firefly Aerospace, Intuitive Machines, ispace-U.S., Lockheed Martin Space, Northrop Grumman Space Systems, SpaceX and Voyager Technologies. The contract runs ten years with a combined maximum value of 10 billion dollars. The first CLPS grew from 2.6 to 4.2 billion.

Under the VADR contract NASA awarded six task orders worth about 1.3 billion dollars to Arrow Science and Technology, Blue Origin and Impulse Space. SpaceX will launch communications satellites toward cislunar space on Starship in 2028. Honeybee Robotics, Redwire and Yank Technologies were chosen for a communications and solar power tower.

The same weekend NASA opened applications for a new astronaut class.

CONTRACTS

CLPS 2
up to $10bn, 10 yrs

CLPS 1
$2.6bn to $4.2bn

VADR
6 orders, ~$1.3bn

PROMISE
$94M + option

SOURCES

1 **NASA**, NASA Awards Moon Base Transportation, Technology, Science Projects, 7 Nov 2026 nasa.gov

2 **SpacePolicyOnline**, NASA Shares Moon Base Progress, Contract Awards, 9 Nov 2026 spacepolicyonline.com

3 **NASA**, NASA Seeks Next Generation of Explorers to Join Astronaut Corps, 7 Nov 2026 nasa.gov

ORBIT / NASA, JPL / PROMISE

A rover that survives the lunar night

A lunar night lasts about 14 days, and solar rovers usually do not survive it. PROMISE will stay warm with a radioisotope source.

PROMISE is a six-wheeled rover meant to land near the lunar south pole. NASA is targeting a late-2028 launch and an 18-month mission. JPL will run the mission and build the rover largely from spare flight hardware and engineering models from the Curiosity and Perseverance Mars rovers.

Its drill should reach about one metre below the surface and sample regolith, looking for ice and other resources. According to SpacePolicyOnline, the rover uses a radioisotope thermoelectric generator to stay warm through the roughly 14-day night that solar landers and rovers usually do not survive.

Blue Origin won a 94 million dollar base contract for delivery, with an option worth another 140 million, and will use its Mark 1 lander. According to SpacePolicyOnline, there is a "pretty firm" commitment of 400 million dollars for the JPL modifications.

PROMISE

Launch
late 2028

Duration
18 months

Drill
~1 m

Power
radioisotope

Lander
Blue Origin Mark 1

SOURCES

1 **NASA**, NASA Awards Moon Base Transportation, Technology, Science Projects, 7 Nov 2026 nasa.gov

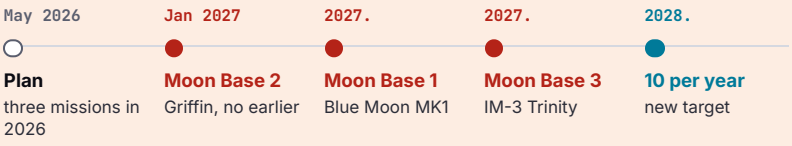
2 **SpacePolicyOnline**, NASA Shares Moon Base Progress, Contract Awards, 9 Nov 2026 spacepolicyonline.com

ORBIT / NASA / SCHEDULE

Contracts arrive, missions slip

At the same event where contracts were awarded, NASA acknowledged that the first three Moon Base missions will not fly this year.

THE FIRST MOON BASE MISSIONS



According to SpacePolicyOnline, from the 7 November event.

According to SpacePolicyOnline, Moon Base 1, Blue Origin's Blue Moon MK1 Endurance lander, has moved to 2027. The same goes for Moon Base 3, Intuitive Machines' IM-3 Trinity lander. Moon Base 2, Astrobotic's Griffin, is still in engine testing and will not fly before January.

For 2027 NASA now expects four CLPS missions, these three plus one from Firefly. The goal of ten missions a year, according to Nicky Fox, starts only in 2028.

The record of past missions explains the caution. According to the same source, Astrobotic did not reach the Moon, both Intuitive Machines landers tipped over, and only Firefly's Blue Ghost Mission 1 was a complete success.

NUMBER

4

CLPS missions expected in 2027.
SpacePolicyOnline

SOURCES

1 **SpacePolicyOnline**, NASA Shares Moon Base Progress, Contract Awards, 9 Nov 2026
spacepolicyonline.com

ORBIT / NASA / 7 AND 9 NOV 2026

ESCAPADE is on its way to Mars

In the last issue the manoeuvres were a preview. Now they are done: both spacecraft used an Earth flyby and an engine burn to head for Mars.

THE ESCAPADE ROUTE

13 Nov 2025	7 Nov	9 Nov	7 and 9 Sep 2027	June 2028
Launch	**First craft**	**Second craft**	**Orbit insertion**	**Science**
New Glenn	~1,000 miles, 9 min	~650 miles, 9 min	~11 min	two orbiters

According to NASA, 9 November release.

The first spacecraft passed about 1,000 miles above Earth around 5:00 EST on Saturday, 7 November, and fired its engine for about nine minutes. The second did the same on Monday, 9 November, around 6:00, at about 650 miles. NASA says this is the first time the agency has used a powered Earth flyby to send a spacecraft toward Mars.

Before the flybys the spacecraft spent almost a year waiting in an orbit around L2, about a million miles from Earth. About ten months of flight follow, and the Mars orbit insertion burns are planned for 7 and 9 September 2027. Science begins in June 2028, when the two spacecraft will together measure how the solar wind affects the Martian atmosphere.

> **“This is the first time NASA has used a powered Earth flyby to propel a spacecraft toward Mars.”**
>
> Asal Naseri, NASA, NASA Science

SOURCES

1 **NASA Science**, NASA's ESCAPADE Completes Powered Earth Flybys, Now En Route to Mars, 9 Nov 2026 science.nasa.gov

BUILD / GITHUB, VERCEL / 9 TO 12 NOV 2026

Secrets in logs and alerts that switch off

Several small changes need action before 8 December. Two concern secrets, and two concern repositories nobody touches.

WHAT CHANGES AND WHAT TO DO

CHANGE	STATUS	WHAT TO DO
Secrets in Actions logs	public preview	Turn on the enterprise policy
Cloudflare service tokens	available	Check new alerts
Dependabot pause	**from 8 Dec**	A push or manual re-enable restores alerts
Actions ARM32	**ends by 8 Dec**	Move jobs to 64-bit runners
Vercel Sandbox	available	Snapshots up to 10x faster

According to the GitHub and Vercel changelogs. Source: GitHub Changelog, GitHub Changelog, GitHub Changelog, GitHub Changelog, Vercel.

Since 9 November secret scanning can check the logs of completed GitHub Actions workflows, in public preview for GitHub Enterprise Cloud with Secret Protection or Advanced Security. The policy is off by default, old logs are not scanned, and alerts send no notifications, so you have to watch the list yourself. Since 12 November Cloudflare service tokens are detected too.

From 8 December Dependabot alerts switch off on most repositories with no activity in the past year. Old alerts are archived and come back with the first push or in settings. By the same date GitHub is also retiring ARM32 on hosted Actions runners for Linux and Windows, because Node 24 does not support ARM32.

On 9 November Vercel made Sandbox snapshot creation up to ten times faster, and Gemini 4 Carbon and Hydrogen were added to AI Gateway.

NUMBER

08.12.

date for the Dependabot pause and end of ARM32 runners. GitHub Changelog

SOURCES

1 **GitHub Changelog**, Secret scanning detects secrets in GitHub Actions workflow logs, 9 Nov 2026 `github.blog`

2 **GitHub Changelog**, Secret scanning detects Cloudflare service tokens, 12 Nov 2026 `github.blog`

3 **GitHub Changelog**, Dependabot alerts will pause on inactive repositories, 10 Nov 2026 `github.blog`

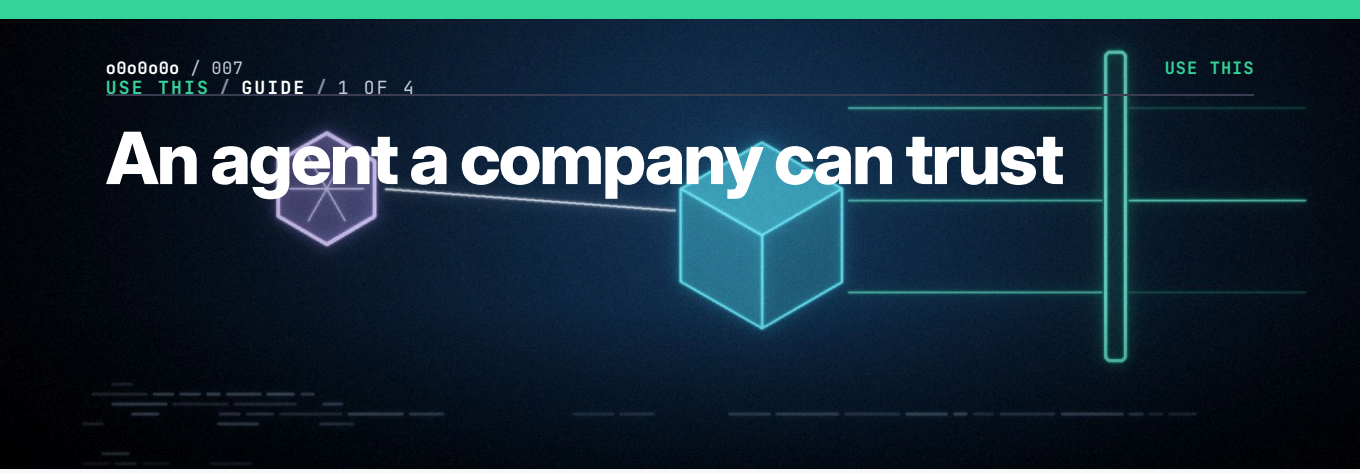
4 **GitHub Changelog**, Actions for ARM32 Linux & Windows retirement, 12 Nov 2026 `github.blog`

5 **Vercel**, Vercel Sandbox snapshot creation is now up to 10x faster, 9 Nov 2026 `vercel.com`

6 **Vercel**, Gemini 4 Carbon now available on AI Gateway, 10 Nov 2026 `vercel.com`

7 **Vercel**, Gemini 4 Hydrogen now available on AI Gateway, 12 Nov 2026 `vercel.com`

An agent a company can trust



This guide turns what OpenAI and Anthropic described this week into four rules for an agent that runs unattended. The examples are in TypeScript and deliberately short.

FOUR LAYERS AROUND AN AGENT

1	**A switch at the bottleneck**	every model call is checked, failure means stop
2	**Credentials and network outside the agent**	a proxy adds the key, an allowlist of hosts
3	**Irreversible needs a person**	delete, pay and publish wait for approval
4	**A log the agent cannot change**	every step, input and result

Editorial recommendations based on the Control Room and Managed Agents descriptions.

The rules are not tied to one vendor. They apply whether the agent runs on Claude, Gemini or GPT, and whether you host it yourself.

The examples assume Node.js 22 and a single service that every model and tool call passes through. If you do not have such a service yet, that is the first task, because without a bottleneck none of the four rules can be enforced reliably.

SOURCES

1 **OpenAI**, Control Room: A central system to track and stop model workloads, 9 Nov 2026 openai.com

2 **Claude**, Operate Claude Managed Agents inside your security perimeter, 12 Nov 2026 claude.com

USE THIS / STEP 1 / KILL SWITCH

A kill switch that fails closed

As in Control Room: a job has an ID, and before every model call you check whether it was stopped. If the check service does not answer, the call is refused.

```
type Status = "active" | "stopped";
const TTL_MS = 5_000; // longest stop delay
const cache = new Map<string, { s: Status; at: number }>();

async function jobStatus(jobId: string): Promise<Status> {
  const hit = cache.get(jobId);
  if (hit && hit.at + TTL_MS > Date.now()) return hit.s;
  try {
    const res = await fetch(`${CONTROL_URL}/jobs/${jobId}`,
      { signal: AbortSignal.timeout(2_000) });
    if (!res.ok) return "stopped"; // an error means stop
    const { status } = (await res.json()) as { status: Status };
    cache.set(jobId, { s: status, at: Date.now() });
    return status;
  } catch {
    return "stopped"; // service not answering: stop
  }
}

export async function callModel(jobId: string, body: unknown) {
  if ((await jobStatus(jobId)) !== "active") {
    throw new Error(`job ${jobId} is stopped`);
  }
  return modelClient.create(body); // the only path to the model
}
```

WHAT TO WATCH FOR

The cache is why OpenAI has a 35-second delay. In your case the longest delay is TTL_MS plus one call, so choose it deliberately. Only the service with callModel should hold the model key, so nothing can call the model around it.

TEST

Shut down the check service and start the agent. If the agent keeps working, the switch does not work.

SOURCES

1 **OpenAI**, Control Room: A central system to track and stop model workloads, 9 Nov 2026 openai.com

USE THIS / STEP 2 / CREDENTIALS AND NETWORK

The agent never sees the key or picks the address

As in Managed Agents: a proxy outside the sandbox adds the key, and a request to a host that is not on the list is blocked and logged.

```
const ALLOWED = new Set(["github.com", "registry.npmjs.org"]);

export async function proxyFetch(url: string, init: RequestInit = {}) {
  const { hostname, protocol } = new URL(url);
  if (protocol !== "https:" || !ALLOWED.has(hostname)) {
    await audit.append({ type: "blocked", hostname });
    throw new Error(`host ${hostname} is not allowed`);
  }
  const headers = new Headers(init.headers);
  headers.delete("authorization"); // the agent cannot send its own keys
  headers.set("authorization", `Bearer ${await vault.tokenFor(hostname)}`);
  // no redirects: the key must not travel to another host
  return fetch(url, { ...init, headers, redirect: "error" });
}
```

WHY EXACTLY THIS WAY

The agent gets only the proxyFetch function, not the key. If someone tricks it into "sending the token to my address", there is nothing to send. The redirect: "error" option matters: without it, a redirect from an allowed host could carry the key header further.

For an agent that fixes builds, Anthropic gives an example list of just two hosts: the git server and the package registry. If the agent runs in a sandbox, set the same limit in the firewall too, so the proxy is not the only protection.

DO NOT FORGET DNS

A list in code does not help if the agent can open its own network connection. Hence the firewall.

SOURCES

1 **Claude**, Operate Claude Managed Agents inside your security perimeter, 12 Nov 2026
claude.com

USE THIS / STEPS 3 AND 4 / APPROVAL AND LOG

Irreversible waits for a person, everything is logged

The last two rules fit in one function: a tool that cannot be undone needs approval, and every tool call goes into the log.

```
type Tool = { name: string; irreversible: boolean;
  run: (input: unknown) ⇒ Promise<unknown> };

export async function runTool(jobId: string, tool: Tool, input: unknown) {
  if (tool.irreversible && !(await approvals.has(jobId, tool.name))) {
    return { ok: false, out: "human approval required" };
  }
  const result = await tool.run(input).then(
    (out) ⇒ ({ ok: true, out }),
    (err) ⇒ ({ ok: false, out: String(err) }),
  );
  await log.append({ jobId, tool: tool.name, input, ok: result.ok,
    at: new Date().toISOString() }); // append only, no edits
  return result;
}
```

HOW TO CHECK IT WORKS

Keep the log where the agent cannot write, for example in a separate database or storage with append-only rights. Managed Agents exports the same record over OpenTelemetry, so you can send it to the monitoring tools you already have.

Before going live, run four tests: shut down the check service, try a request to a host outside the list, call an irreversible tool without approval and delete a log row as the agent. Every attempt must fail.

FOUR TESTS

1. check service down
2. host off the list
3. no approval
4. log deletion

SOURCES

1 **Claude**, Operate Claude Managed Agents inside your security perimeter, 12 Nov 2026 claude.com

2 **OpenAI**, Control Room: A central system to track and stop model workloads, 9 Nov 2026 openai.com

SHIPPED / 6 NOV 2026 → 13 NOV 2026

Arrived this week

What is actually available, with platform and limits. Private betas and announcements are not on this page.

ANTHROPIC · 9 NOV

Claude Fable 5.2

All users, API and Copilot. Mythos by approval only.

AVAILABLE Anthropic · GitHub Changelog

GOOGLE · 10 NOV

Gemini 4 Carbon

Paying developers and Ultra. Intro price until 30 November.

AVAILABLE Google

GOOGLE · 12 NOV

Gemini 4 Hydrogen

Every AI Studio tier, the Gemini app and Search.

AVAILABLE Google

ANTHROPIC · 10 NOV

Claude Data

Paid plans. Certified queries and sharing are still beta.

GA Claude

ANTHROPIC · 10 NOV

Claude Sheets, Whiteboards

Paid plans. Off by default on Enterprise.

BETA Claude

ANTHROPIC · 12 NOV

Claude Meetings

Team and Enterprise, Google Meet only.

BETA Claude

OPENAI · 10 NOV

Businesses in ChatGPT

Business Agents only for selected US categories.

BETA OpenAI

OPENAI · 9 NOV

ChatGPT Enterprise on AWS

Bought through AWS Marketplace, subject to eligibility.

AVAILABLE OpenAI

GITHUB · 9 NOV

Secrets in Actions logs

Enterprise Cloud, policy off by default.

PUBLIC PREVIEW GitHub Changelog

DEEPMIND · 10 NOV

Gemini Robotics 2.5

Robotics partners only.

ROLLOUT Google DeepMind

5 MINUTES / 6 NOV 2026 → 13 NOV 2026

In brief

Short news. If a topic has its own article, the card points to it.

ANTHROPIC · 12 NOV

Managed Agents on-prem

Private beta, the whole service in your environment. Page 19.

Claude

ANTHROPIC · 12 NOV

Claude for Manufacturing

17 skills and four connectors. Page 21.

Claude

OPENAI · 6 NOV

Hacktron

A vulnerability-hunting partnership, with access via Daybreak Red. Page 24.

OpenAI

OPENSSH · 5 NOV

More frequent releases

Because of AI reports, fixes ship before planned releases. Page 24.

OpenSSH

GITHUB · 6 NOV

Copilot in Slack and Teams

Generally available.

GitHub Changelog

GITHUB · 11 NOV

Issue types for personal accounts

Types and fields are no longer just for organisations.

GitHub Changelog

NASA · 7 NOV

A new astronaut class

Applications open, the same weekend as the lunar contracts. Page 37.

NASA

CISA · 10 NOV

Chromium

Use-after-free in the XML component, also affects Edge and Opera. Page 26.

CISA · Google Chrome Releases

THIS WEEK IN NUMBERS / 6 NOV 2026 → 13 NOV 2026

The week in numbers

6.5M

lines of Lean proof for libsignal

GOOGLE DEEPMIND · 10 NOV

~1%

of code flagged for human review

GOOGLE DEEPMIND · 10 NOV

4 weeks

the Signal verification took

GOOGLE DEEPMIND · 10 NOV

20 pers.-yrs

the seL4 verification took

GOOGLE DEEPMIND · PAPER

8 / 40 $

Fable 5.2 per million tokens, in and out

ANTHROPIC · 9 NOV

1.50 $

Hydrogen per million input tokens

GOOGLE · 12 NOV

35 s

maximum stop time in Control Room

OPENAI · 9 NOV

11

new CISA KEV entries

CISA · 6 TO 10 NOV

THIS WEEK IN NUMBERS / 6 NOV 2026 → 13 NOV 2026

More numbers of the week

33.2 $bn

Anthropic revenue for nine months

SEC EDGAR · S-1/A

900.9 $bn

Anthropic contracted commitments, mostly compute

SEC EDGAR · S-1/A

26 $bn

Class N shares for seven companies

SEC EDGAR · S-1/A

132.5 $bn

CoreWeave backlog

COREWEAVE · Q3

96 %

robot success in the plant, proof of concept

GOOGLE DEEPMIND · 10 NOV

107 M

cell profiles for AlphaGenome 2

GOOGLE DEEPMIND · 11 NOV

10 $bn

maximum value of the CLPS 2 contract

NASA · 7 NOV

~650 miles

altitude of the second ESCAPADE flyby

NASA SCIENCE · 9 NOV

NEXT / 6 NOV 2026 → 13 NOV 2026

What we are watching

Future events after the 13 November cut-off. None of them had happened when the issue closed.

→ 18 NOV

Anthropic share price

According to media, not the SEC filing.

TechCrunch · CNN

→ 19 NOV

ANT on Nasdaq

Expected start of trading, also per media.

TechCrunch · Fortune

→ 19 NOV

Claude Data

Turns on by default on Enterprise.

Claude

→ 20 NOV

KEV deadline: SonicWall

For US federal agencies.

CISA

→ 24 NOV

KEV deadline: Outlook, Chromium, Cisco

And Openwave Webtop Transcoder.

CISA

→ 30 NOV

Gemini 4 Carbon intro price

According to Google, valid until this date.

Google

→ NOVEMBER

Euclid DR1-Foundation

ESA plans data for about 1,900 square degrees of sky, with no exact date.

ESA

→ 8 DEC

Dependabot and ARM32

Pause on inactive repositories and end of ARM32 runners.

GitHub Changelog · GitHub Changelog

→ JAN 2027

Moon Base 2

Griffin, no earlier than January.

SpacePolicyOnline

→ 7 SEP 2027

ESCAPADE at Mars

First orbit insertion burn, the second on 9 September.

NASA Science

SOURCES / 6 NOV 2026 → 13 NOV 2026

All sources

Sources are grouped by page, in order of first appearance. The primary source comes first. Every link is live.

04 EDITORIAL: WHAT CAN BE CHECKED

1. **Google DeepMind**. AlphaProof Code: Verifying Signal's Security at the Machine-Code Level. 10 Nov 2026
   google-deepmind.github.io/code-verification/
2. **OpenAI**. Control Room: A central system to track and stop model workloads. 9 Nov 2026
   openai.com/index/control-room/
3. **CISA**. Known Exploited Vulnerabilities, JSON feed. 10 Nov 2026
   cisa.gov/sites/default/files/feeds/known_exploited_vulnerabilit…
4. **Google DeepMind**. Introducing AlphaProof Code: a breakthrough in verifying complex software at depth. 10 Nov 2026
   deepmind.google/blog/introducing-alphaproof-code/
5. **SEC EDGAR**. Anthropic, PBC, Form S-1/A (Amendment No. 2). 10 Nov 2026
   sec.gov/Archives/edgar/data/2130357/000162828026072484/anthropi…

07 COVER STORY: SPECIFICATION, AGENTS, LEAN'S KERNEL

6. **Google DeepMind**. Formal Verification of Signal, technical paper (PDF). 10 Nov 2026
   deepmind.google/documents/36/Formal_Verification_of_Signal_4.pdf

09 COVER STORY: TWO OLD BUGS AS AN EXAM

7. **GitHub**. google-deepmind/code-verification. 10 Nov 2026
   github.com/google-deepmind/code-verification

10 COVER STORY: YEARS OF WORK BECOME WEEKS

8. **Anthropic**. Formalizing Fermat's Last Theorem. 4 Sep 2026
   anthropic.com/research/formalizing-fermats-last-theorem
9. **Anthropic**. Introducing Claude Fable 5.2 and Claude Mythos 5.2. 9 Nov 2026
   anthropic.com/claude-fable-and-mythos-5-2

11 MODEL WATCH: FABLE 5.2 AND MYTHOS 5.2: ONE MODEL, TWO DOORS

10. **GitHub Changelog**. Claude Fable 5.2 in GitHub Copilot. 9 Nov 2026
    github.blog/changelog/2026-11-09-claude-fable-5-2-in-github-cop…

13 MODEL WATCH: GEMINI 4: CARBON FOR EVERYONE, HYDROGEN FOR SPEED

11. **Google**. Introducing Gemini 4 Hydrogen. 12 Nov 2026
    blog.google/innovation-and-ai/models-and-research/gemini-models…
12. **Google**. Gemini 4 Carbon is now available. 10 Nov 2026
    blog.google/innovation-and-ai/models-and-research/gemini-models…

15 MODEL WATCH: MODELS OF THE WEEK

13. **GitHub Changelog**. Gemini 4 Carbon in GitHub Copilot. 10 Nov 2026
    github.blog/changelog/2026-11-10-gemini-4-carbon-in-github-copi…
14. **Vercel**. Gemini 4 Hydrogen now available on AI Gateway. 12 Nov 2026
    vercel.com/changelog/gemini-4-hydrogen-now-available-on-ai-gate…
15. **Google DeepMind**. Gemini Robotics 2.5 accelerates real-world general dexterity. 10 Nov 2026
    deepmind.google/blog/gemini-robotics-2-5-accelerates-real-world…

SOURCES / CONTINUED

17 WORK: CLAUDE GETS SHEETS, BOARDS AND DATA

16. **Claude**. Introducing Claude Whiteboards, Claude Sheets, and Claude Data. 10 Nov 2026
claude.com/resources/articles/sheets-whiteboards-data
17. **Claude**. Your apps can now open in Claude, and Claude can work in them. 10 Nov 2026
claude.com/resources/articles/interactive-connectors-full-tool

18 WORK: CLAUDE IN THE MEETING: WHO MAY LISTEN

18. **Claude**. Introducing Claude Meetings. 12 Nov 2026
claude.com/resources/articles/claude-meetings

19 WORK: AGENTS INSIDE YOUR PERIMETER

19. **Claude**. Operate Claude Managed Agents inside your security perimeter. 12 Nov 2026
claude.com/resources/articles/self-hosted-claude-managed-agents

20 WORK: BUSINESSES ENTER CHATGPT, CHATGPT ENTERS AWS

20. **OpenAI**. Introducing Businesses in ChatGPT. 10 Nov 2026
openai.com/index/introducing-businesses-chatgpt/
21. **OpenAI**. ChatGPT Enterprise now available in AWS Marketplace. 9 Nov 2026
openai.com/index/chatgpt-enterprise-aws-marketplace/

21 WORK: AGENTS BUILT FOR AN INDUSTRY

22. **IBM**. Imagine Enterprise. 10 Nov 2026
ibm.com/consulting/imagine-enterprise
23. **Claude**. Introducing Claude for Manufacturing. 12 Nov 2026
claude.com/resources/articles/claude-for-manufacturing
24. **CIO Dive**. Microsoft, IBM team up on AI agent platform. 10 Nov 2026
ciodive.com/news/microsoft-ibm-team-up-ai-agent-platform/874152/

24 SECURITY: AI FINDS BUGS. SO DO ATTACKERS.

25. **OpenSSH**. OpenSSH 10.6p2 release notes. 5 Nov 2026
openssh.org/releasenotes.html
26. **Microsoft MSRC**. CVE-2026-124194, Microsoft Outlook Spoofing Vulnerability. 10 Nov 2026
msrc.microsoft.com/update-guide/en-US/vulnerability/CVE-2026-12…
27. **OpenAI**. Partnering with Hacktron to find vulnerabilities in OpenAI's systems. 6 Nov 2026
openai.com/index/partnering-with-hacktron/

25 PATCH NOW: PATCH IMMEDIATELY

28. **CISA**. Known Exploited Vulnerabilities Catalog, version 2026.11.10. 10 Nov 2026
cisa.gov/known-exploited-vulnerabilities-catalog
29. **Barracuda**. Pre-Authentication Authorization Bypass and Command Injection Vulnerabilities. 9 Nov 2026
trust.barracuda.com/security/pre-authentication-authorization-b…
30. **SonicWall PSIRT**. SNWLID-2026-0014. 6 Nov 2026
psirt.global.sonicwall.com/vuln-detail/SNWLID-2026-0014

26 SECURITY: ELEVEN CONFIRMED ATTACKS IN FIVE DAYS

31. **Google Chrome Releases**. Stable Channel Update for Desktop. 10 Nov 2026
chromereleases.googleblog.com/2026/11/stable-channel-update-for…

27 SECURITY: TWO CHAINS: ROOT ON A BOX AND A HASH FROM A MESSAGE

32. **Qualys**. Microsoft Patch Tuesday, November 2026 Security Update Review. 10 Nov 2026
blog.qualys.com/vulnerabilities-threat-research/2026/11/10/micr…
33. **Action1**. Patch Tuesday November 2026. 10 Nov 2026
action1.com/patch-tuesday/patch-tuesday-november-2026/

SOURCES / CONTINUED

28 MONEY: ANTHROPIC HEADS TO MARKET AS ANT

34. **Anthropic**. Anthropic launches roadshow for proposed initial public offering. 10 Nov 2026
anthropic.com/news/roadshow-launch
35. **TechCrunch**. Anthropic sets IPO price range of $70 to $78 per share. 10 Nov 2026
techcrunch.com/2026/11/10/anthropic-sets-ipo-price-range-of-70-…
36. **CNN**. We've never seen anything like the Anthropic IPO. 10 Nov 2026
cnn.com/2026/11/10/investing/weve-never-seen-anything-like-anth…
37. **Fortune**. Anthropic IPO price range. 10 Nov 2026
fortune.com/2026/11/10/anthropic-ipo-price-range-ai-stock-sell-…

31 MONEY: COREWEAVE: RECORD REVENUE, RECORD DEBT

38. **CoreWeave**. CoreWeave Reports Strong Third Quarter 2026 Results. 10 Nov 2026
investors.coreweave.com/news/news-details/2026/CoreWeave-Report…

32 MONEY: NEBIUS: CUSTOMERS PAY FOR THE BUILD UP FRONT

39. **Nebius Group / SEC**. Letter to shareholders, Q3 2026 (Exhibit 99.2). 12 Nov 2026
sec.gov/Archives/edgar/data/1513845/000110465926126008/tm262957…

35 LAB: ALPHAGENOME 2 READS DNA CELL BY CELL

40. **Google DeepMind**. AlphaGenome 2: Deciphering how cells read DNA to better understand disease. 11 Nov 2026
deepmind.google/blog/alphagenome-2-deciphering-how-cells-read-d…

37 ORBIT: FROM VISITING TO BUILDING

41. **NASA**. NASA Awards Moon Base Transportation, Technology, Science Projects. 7 Nov 2026
nasa.gov/news-release/nasa-awards-moon-base-transportation-tech…
42. **SpacePolicyOnline**. NASA Shares Moon Base Progress, Contract Awards. 9 Nov 2026
spacepolicyonline.com/news/nasa-shares-moon-base-progress-contr…
43. **NASA**. NASA Seeks Next Generation of Explorers to Join Astronaut Corps. 7 Nov 2026
nasa.gov/news-release/nasa-seeks-next-generation-of-explorers-t…

40 ORBIT: ESCAPADE IS ON ITS WAY TO MARS

44. **NASA Science**. NASA's ESCAPADE Completes Powered Earth Flybys, Now En Route to Mars. 9 Nov 2026
science.nasa.gov/science-research/heliophysics/nasas-escapade-c…

41 BUILD: SECRETS IN LOGS AND ALERTS THAT SWITCH OFF

45. **GitHub Changelog**. Secret scanning detects secrets in GitHub Actions workflow logs. 9 Nov 2026
github.blog/changelog/2026-11-09-secret-scanning-detects-secret…
46. **GitHub Changelog**. Secret scanning detects Cloudflare service tokens. 12 Nov 2026
github.blog/changelog/2026-11-12-secret-scanning-adds-a-detecto…
47. **GitHub Changelog**. Dependabot alerts will pause on inactive repositories. 10 Nov 2026
github.blog/changelog/2026-11-10-dependabot-alerts-will-pause-o…
48. **GitHub Changelog**. Actions for ARM32 Linux & Windows retirement. 12 Nov 2026
github.blog/changelog/2026-11-12-actions-for-arm32-linux-window…
49. **Vercel**. Vercel Sandbox snapshot creation is now up to 10x faster. 9 Nov 2026
vercel.com/changelog/vercel-sandbox-snapshot-creation-is-now-up…
50. **Vercel**. Gemini 4 Carbon now available on AI Gateway. 10 Nov 2026
vercel.com/changelog/gemini-4-carbon-now-available-on-ai-gateway

47 5 MINUTES: IN BRIEF

51. **GitHub Changelog**. GitHub Copilot in Slack and Teams is generally available. 6 Nov 2026
github.blog/changelog/2026-11-06-github-copilot-in-slack-and-te…
52. **GitHub Changelog**. Issue types and fields are available for personal accounts. 11 Nov 2026
github.blog/changelog/2026-11-11-issue-types-and-fields-are-ava…

50 NEXT: WHAT WE ARE WATCHING

53. **ESA**. Euclid DR1 Timeline. 15 Jun 2026
cosmos.esa.int/web/euclid/dr1-timeline

o0o0o0o

/ 007

Issue seven. A weekly technology review, free, on tutorijali.app.

Issue 007
Date 13 Nov 2026
Period 6 Nov 2026, 00:00 → 13 Nov 2026, 05:37
Web tutorijali.app

tutorijali.app