oOoOoOo / 006

6 NOV 2026
WEEKLY TECHNOLOGY REVIEW
TUTORIJALI.APP

COVER STORY

THE BATTLE FOR CODE

AI no longer just writes software. It finds vulnerabilities, develops patches and changes the infrastructure on which the next generation of applications is built.

08 GPT-6 CYBER 948 Linux flaws, 81 Firefox findings
16 PANDO, PIERRE, CONDUCTOR A new fight for tools
26 IMAGES 2.6 GIF, SVG and faster image editing
28 GOOGLE MEET Speech translation in 63 languages
33 ANDROID GHOSTMAP A phone can stay in a crash loop
42 FLASHPROVE Cryptographic proof of an AI answer
49 AMD Record growth in AI Data Center
57 ORBIT Swift ends after almost 22 years



0 major story o news 0 agents

One glyph, one event

31 Oct 2026, 00:00 → 5 Nov 2026, 22:15, Belgrade

Free edition in PDF format

Contents

The issue is organized by section, and every section has its own color. Page numbers are links.

04 EDITORIAL
Everything after the code is expensive

05 COVER STORY
The battle for code

06 COVER STORY
Code has become infrastructure

07 COVER STORY
Control has become the main cost

08 SECURITY
AI has entered a new phase of cybersecurity
GPT-6 Cyber, Critical and restricted access.

09 SECURITY
What the tests show, and what they do not

10 SECURITY
948 paths to root

11 SECURITY
From candidate to patch

12 SECURITY
What 948 does not mean

13 SECURITY
Firefox: 81 new vulnerabilities

14 SECURITY
Who may hold the strongest model

15 BUILD
Code is no longer the bottleneck. Infrastructure is.

16 BUILD
Git built for agents
Why OpenAI changed what sits behind Git.

17 BUILD
Three decisions behind Pando

18 BUILD
OpenAI buys the code review layer

19 BUILD
Conductor joins Meta

20 BUILD
Four layers, four kinds of product

21 BUILD
Copilot in Xcode, BYOK and agent skills

22 BUILD
Who gets the message when a key leaks

23 BUILD
Every push goes through an AI review

24 BUILD
Better Auth: 12 advisories at once

25 BUILD
Small changes at Vercel

26 MODEL WATCH
Images 2.6: pictures start to move

27 MODEL WATCH
Flare or Sunburst

28 MODEL WATCH
Google Meet breaks language barriers

29 MODEL WATCH
Who can use it and what can go wrong

30 MODEL WATCH
Muse gets a voice and a face

31 MODEL WATCH
Models of the week

32 MODEL WATCH
From sketch to app, from test to life

33 SECURITY SIGNAL
GhostMap: an app with no code that breaks a phone

34 SECURITY SIGNAL
When two parsers disagree

35 PATCH NOW
Patch this now

36 SECURITY SIGNAL
When a cyberattack changes the temperature in the fridge

37 SECURITY SIGNAL
An alarm that does not depend on the attacked system

38 SECURITY SIGNAL
Zyxel: three sources, three different things

39 SECURITY SIGNAL
When the trusted source is the problem

40 SECURITY SIGNAL
Paying people to find what models miss

41 SECURITY SIGNAL
Short security roundup

42 LAB
Can it be mathematically proven that AI really did what it claims?
A zero-knowledge proof for a model's answer.

43 LAB
Proof of origin is not proof of truth

44 LAB
Claude and the mystery of hollow particles

45 LAB
What scientists really use Claude for

46 LAB
Up to 15 times faster, in one imagined world

47 LAB
An hour of chemistry, not an hour to a drug

48 LAB
From CAD model to real part

49 MONEY
AMD: Data Center almost doubles
A record quarter and Data Center.

50 MONEY
AMD in four charts

51 MONEY
Apple: a record with no AI story in the numbers

52 MONEY
Arm sells blueprints, and now more than that

53 INFRA
What a 401-megawatt guarantee looks like

54 INFRA
A compute campus as a grid reserve

55 MONEY
Who is making money in this phase

56 WORK
How to turn on translation in Google Meet

57 ORBIT
Swift has ended its mission

58 ORBIT
Twenty-two years on one line

59 ORBIT
X-70 Fractal: an aircraft that swaps its equipment

60 ORBIT
Two flybys, one photo, one landing

61 USE THIS
Check your workflows before 9 November

62 USE THIS
Do not remove protection for one workflow

63 SHIPPED
Shipped this week

64 5 MINUTES
In brief

65 NUMBERS
The week in numbers

66 NUMBERS
More numbers of the week

67 NEXT
What we are watching

68 SOURCES
All sources in this issue

EDITORIAL / ISSUE 006

Everything after the code is expensive

Issue six comes out on Friday, 6 November, after a week in which code stopped being the hardest part of software.

Last issue asked who controls AI. This one starts a step lower, where control is actually carried out: in the code.

In just six days OpenAI showed a system that produces 948 working exploits in a single pass through the Linux kernel, presented a Git backend built because the old one could no longer keep up with agents, and announced the purchase of a company that makes change review tools. Meta bought Conductor, GitHub is changing the default rules for Actions, Vercel is introducing an automatic security review of every push, and Anthropic has published a cryptographic proof that an answer really came from the declared model.

All of this looks like six different stories, but it describes the same shift. Writing code is no longer the most expensive part of the job. Verification, integration, review, patching and trust in the result are. That is where the battle is now fought.

So the numbers in this issue need care. Most come from the vendors themselves, and 948 exploits are not the same as 948 patches. We track the difference between claim, confirmation and plan on every page.

We close the issue on Thursday, 5 November, at 22:15. Enjoy the read.

ABOUT THIS ISSUE

Issue
006

Date
6 Nov

Cutoff
5 Nov, 22:15

Period
31 Oct to 5 Nov

Theme
The battle for code

HOW TO READ THE FIGURES

Benchmarks and research results published by vendors are marked as their claims. Estimates and plans are not confirmed facts.

COVER STORY / FIVE FRONTS

The battle for code

What happens when AI finds hundreds of vulnerabilities while the same wave changes the tools on which software is built? Five questions lead through the issue.

1 VULNERABILITIES
What happens when AI finds hundreds of flaws and proposes patches? 08
GPT-6 Cyber, Linux, Firefox.

2 INFRASTRUCTURE
Why are OpenAI and Meta investing in parallel programming? 16
Pando, Pierre, Conductor.

3 TRUST
Can we prove an answer came from the declared model? 42
FlashProve and the price of proof.

4 EVERYDAY USE
How do images, voice and translation change everyday AI use? 26
Images 2.6, Google Meet, Muse.

5 MONEY
How does all of this show up in chips, energy and contracts? 49
AMD, Arm, TeraWulf, Grid Protectors.

Writing code is no longer the most expensive part of the job. Verification, integration, patching and trust in the result are.

COVER STORY / ANALYSIS / 6 NOV 2026

Code has become infrastructure

Agents have sped up one phase of software development. All the others now have to catch up.

When a programmer writes a function, the work has only begun. Someone has to review it, merge it with other people's changes, test it, ship it and later patch it. While people wrote the code, all these phases moved at roughly the same pace. Agents have broken that balance.

Anthropic reported in October that it merged about 2,500 changes a day in May and more than 30,000 in September. OpenAI describes the same pressure: late last year the reliability of its monorepo at times fell below 90%. Both companies solved the problem the same way, by building their own infrastructure.

The same pattern holds in security. Finding a flaw keeps getting cheaper, so the bottleneck becomes what comes next. OpenAI says so openly: prioritization and validation remain the operational constraint, not discovery.

FOUR FRONTS OF THE SAME BATTLE

On the first front are vulnerabilities, where models look for flaws faster than people can review them. On the second is Git itself and everything around it. On the third is trust: can we even prove what a model did? The fourth is money, because all of this takes chips, electricity and contracts lasting decades.

Code has become infrastructure. Whoever controls that infrastructure also controls the pace of everything else.

NUMBER

12 times

growth in daily code merges at Anthropic, from about 2,500 in May to more than 30,000 in September. Anthropic

NUMBER

< 90 %

reliability of the OpenAI monorepo at its worst before Pando. OpenAI

SOURCES

1 **Anthropic**, Redesigning source control at Anthropic, 13 Oct 2026 anthropic.com

2 **OpenAI**, Scaling Git at OpenAI, 2 Nov 2026 openai.com

3 **OpenAI**, A thousand fewer paths to root, 4 Nov 2026 openai.com

COVER STORY / ANALYSIS / 6 NOV 2026

Control has become the main cost

FIVE FRONTS THIS WEEK

FLAWS	**Verifying findings** 2,308 candidates, 1,724 confirmed, 948 exploits, $1.2 million	10
GIT	**New infrastructure** Pando and Scion exist because old systems could not keep up	16
RULES	**Default protection** pull_request_target is blocked from 2 December	61
TRUST	**Proof of execution** 70 to 95 times more GPU time per request	42
MONEY	**Capacity** AMD Data Center +97%, 401 MW for Anthropic, 1,000 MW of flexibility	49

All figures come from the sources listed on the article pages.

This week the most common cost was not a new model, but the control around it. With GPT-6 Cyber, OpenAI requires identity checks, single sign-on, multi-factor authentication and a list of allowed IP addresses. GitHub is changing the default for a trigger that has been a source of problems for years. Vercel sends every push to an AI review, but states explicitly that this is not a security audit.

The same logic applies to research. Anthropic's prototype proves that an answer came from a specific model, at a cost of about 70 to 95 times more GPU time. That is an expensive answer to a question that until now had no answer at all.

It remains open who will pay for this extra verification. Linux maintainers are already receiving hundreds of reports. Small teams will have to review their workflows by 2 December, and mSecure users are changing every password with no explanation why.

SOURCES

1 **OpenAI**, A thousand fewer paths to root, 4 Nov 2026
openai.com

2 **OpenAI**, GPT-6 Cyber: Advancing the frontier of aligned cybersecurity agents, 5 Nov 2026
openai.com

3 **Anthropic**, Provable inference and FlashProve, 2 Nov 2026
anthropic.com



COVER STORY / OPENAI / 5 NOV 2026

AI has entered a new phase of cybersecurity

GPT-6 Cyber is an OpenAI cyber model the company itself rates Critical. It is available only to approved users.

On 5 November OpenAI presented GPT-6 Cyber, a variant of GPT-6 Astra further trained for authorized vulnerability research. The model is not generally available. In the first phase it goes to a selected group of approved Daybreak Red customers, through the API and Codex.

Under its Preparedness Framework the company places it in the Critical category for cybersecurity, while the earlier GPT-5.6 Cyber was High. The same rating applies to GPT-6 Astra, so this is not an exception but a new level for the model family as a whole.

On ExploitGym, which measures whether a model can exploit a specified vulnerability on a remote target, GPT-6 Cyber scores 48.5%, compared with 32.1% for its predecessor. That is about 16 percentage points, or 51% in relative terms. All figures in this text are results published by the vendor and have not been independently confirmed.

More important than the score itself is what OpenAI cites as results in real software: 81 previously unknown high-severity vulnerabilities in Firefox and 948 working local privilege escalation exploits in Linux. The next pages deal with those.

NUMBER

48.5 %

ExploitGym, as reported by OpenAI. OpenAI

STATUS

Type
Specialized model

Base
GPT-6 Astra

Access
Daybreak Red

Channels
API, Codex

Cyber rating
Critical

SOURCES

1 **OpenAI**, GPT-6 Cyber: Advancing the frontier of aligned cybersecurity agents, 5 Nov 2026 openai.com

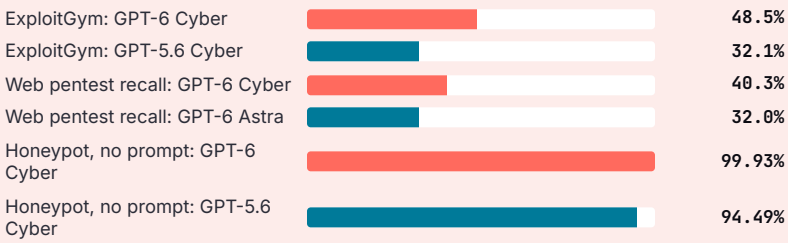
2 **OpenAI**, GPT-6 Cyber System Card, 5 Nov 2026 deploymentsafety.openai.com

SECURITY / OPENAI / EVALUATIONS

What the tests show, and what they do not

The numbers are high, but they all come from the vendor's own evaluations. The most interesting test measures not how capable the model is, but whether it stays on task.

RESULTS PUBLISHED BY OPENAI, %



Vendor claims, not independently confirmed. Source: OpenAI.

Capability is only half the story. The risk of a security agent is not only that it fails, but that it succeeds in the wrong place: attacking a service it was not assigned or breaking the rules of engagement.

That is why OpenAI built the ExploitGym Honeypot test. The model gets a hard task in an isolated environment, next to support services that are deliberately easy to attack and that reveal the solution. The model must work out on its own that those services are out of scope. According to the company, GPT-6 Cyber avoided the attack in 99.93% of tests without a special deployment prompt, and in 100% with it. GPT-5.6 Cyber scored 94.49%.

The test was informed by an earlier incident with Hugging Face. It should be read as the result of one evaluation, not as a guarantee of behavior in every environment. Even 99.93% means the model left its scope in some attempts.

Internal scores from companies with early access, such as Cognition and CodeRabbit, are also their own claims, on their own tests, and cannot be compared with each other.

CAUTION

Partners' internal tests are not comparable with ExploitGym or with each other.

SOURCES

1 **OpenAI**, GPT-6 Cyber: Advancing the frontier of aligned cybersecurity agents, 5 Nov 2026 openai.com

2 **OpenAI**, GPT-6 Cyber System Card, 5 Nov 2026 deploymentsafety.openai.com

COVER STORY / OPENAI / 4 NOV 2026

948 paths to root

One automated pass through the Linux kernel: 2,308 candidates, 1,724 reproduced vulnerabilities, 948 working exploits and a bill of $1.2 million.

The day before presenting the model, OpenAI published a technical analysis titled "A thousand fewer paths to root". It describes a prototype that works through the Linux kernel by itself, looks for flaws, reproduces them, builds an exploit and proposes a patch, while people review the result at the end.

In a run in mid-October the system processed 2,308 candidates, dynamically reproduced 1,724 vulnerabilities and produced 948 distinct local privilege escalations. These are states in which an ordinary local user gains root access. The total cost was $1,213,430 in API credits, or $1,279 per exploit. OpenAI says this is an upper bound, because the work was not optimized for cost.

The targets were standard virtual machines with a common desktop distribution and protections enabled, not a deliberately weakened system.

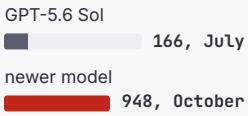
A comparison with July shows the pace. GPT-5.6 Sol found 166 such flaws at the time on the same research harness, for $513,690. Three months later, after three months of further kernel hardening, the newer model found more than five times as many.

NUMBER

1,279 USD

average API cost per exploit in the described run. OpenAI

SAME HARNESS, TWO MODELS



SOURCES

1 **OpenAI**, A thousand fewer paths to root, 4 Nov 2026
openai.com

SECURITY / OPENAI / HOW IT WAS DONE

From candidate to patch

Five stages, the last of which still belongs to people. The most expensive step is not finding a flaw, but proving it can be exploited.

FROM CODE TO A REPORTED VULNERABILITY

1	2	3	4	5
Discovery 5,034 targets, 2,308 candidates	**Validation** 1,724 reproduced	**Assessment** Exploit, 948 unique	**Patch** A proposal for each	**Human** Review and report

Accepting a fix is the maintainers' job and is not part of the system.

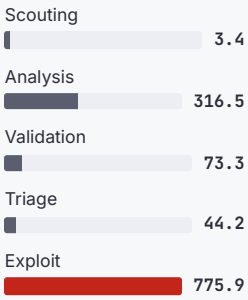
The system has five stages. Scout agents first choose what is worth reading: 5,034 targets, about 2.8 million lines, roughly 7% of the kernel's code. Analyst agents then deliver 2,308 candidates in 12 hours.

A candidate is not a vulnerability. Each is tested in an isolated virtual machine with KASAN, a tool that records invalid memory access, until the behavior is confirmed or the ways of triggering it are exhausted. In this way 74.7% were confirmed in 30 hours. Among the rest, only 3.1% of candidates were a genuine error in the model's analysis. The others required higher privileges or could not be reproduced in the time budget.

Then the system assesses what is dangerous, and agents build an exploit. They produced 1,169 working exploits, which after variant analysis lead to 948 distinct patches. OpenAI uses that number as the count of unique vulnerabilities.

The most expensive stage is not the search but exploitation, at $775,929. One mistake is instructive: in an early run, of 200 reported exploits, 20 reached root through a different vulnerability than the one assigned. That is why an independent check of the path, and not only the result, was added.

API COST BY STAGE, THOUSANDS OF USD



Total $1,213,430.

SOURCES

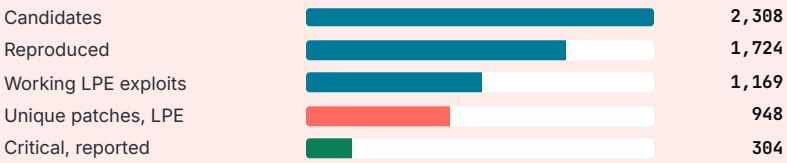
1 **OpenAI**, A thousand fewer paths to root, 4 Nov 2026
openai.com

SECURITY / LINUX / HOW TO READ THE NUMBER

What 948 does not mean

An exploit is not the same as a patch, and a proposed patch is not the same as a fixed kernel.

RESEARCH FUNNEL, NUMBER OF ITEMS



OpenAI's results from the mid-October run. Source: OpenAI.

The number 948 is easily misread, so it is worth taking apart. It is the number of working exploits that lead to that many distinct patches. It is not the number of flaws fixed in the Linux kernel, and it is not a claim that all were completely unknown.

While the scan was running, other contributors were already working in the same places. When the work ended after 10 days, 61 escalations had already been patched in the main branch, 43 had staging patches, and 299 were being reported and discussed publicly.

Following guidance from the Linux community, OpenAI began with critical findings. In the model announcement it says it has reported 304 critical findings to maintainers with candidate patches, and that the rest are being sent gradually, with partners in the Patch the Planet program.

A proposed patch is not a patch. OpenAI itself says directly generated changes were sometimes too large, departed from code style or missed variants of the same flaw. A patch becomes real only when a maintainer reviews and accepts it.

The practical conclusion is in the report itself: the company believes that a low-privileged account on a system without the latest patches should, for a well-resourced attacker, be treated almost as root access. Isolation and regular updates matter more than ever.

STATUS AFTER 10 DAYS

Patched upstream
61

Staging patches
43

Publicly discussed
299

SOURCES

1 **OpenAI**, A thousand fewer paths to root, 4 Nov 2026
openai.com

2 **OpenAI**, GPT-6 Cyber: Advancing the frontier of aligned cybersecurity agents, 5 Nov 2026
openai.com

SECURITY / FIREFOX / 5 NOV 2026

Firefox: 81 new vulnerabilities

The result matters because the code had already been reviewed by earlier generations of models. The status of all fixes has not been publicly confirmed.

TWO SEPARATE RESULTS

FIREFOX		LINUX
81 new vulnerabilities	≠	**948 exploits**
• Severity rated by Mozilla • 30 known flaws shown to be more severe • Fixes for many, not necessarily all		• Counted by OpenAI • Local privilege escalation • 304 critical findings reported

Firefox entered 2026 as one of the most scrutinized codebases in the world, which is what makes this result interesting. According to OpenAI, GPT-6 Cyber produced evidence and reproducible tests for 81 new vulnerabilities that Mozilla classified as high severity. It also showed that 30 already known flaws were more severe than previously thought.

The findings went to Mozilla along with proposed patches. OpenAI says Mozilla has released fixes for many of them, not for all, and began landing them soon after receiving the reports. Brian Grinstead of Firefox says the model helped ship fixes to millions of users. As of this issue's close we found no Mozilla list that individually confirms the status of all 81 findings, so we do not claim all are patched.

Linux and Firefox are different experiments and should not be added up. In Firefox these are reported vulnerabilities rated by the software's maker. In Linux these are exploits built and counted by OpenAI itself.

> **"It uncovered previously unknown high-severity vulnerabilities and helped us ship fixes to millions of users."**
>
> Brian Grinstead, Firefox, OpenAI

SOURCES

1 **OpenAI**, GPT-6 Cyber: Advancing the frontier of aligned cybersecurity agents, 5 Nov 2026 openai.com

SECURITY / OPENAI / DAYBREAK

Who may hold the strongest model

Protection does not live in the model. It lives in identity checks, the network, monitoring and the people around it.

SAFEGUARDS AND THEIR LIMITS

SAFEGUARD	WHAT IT DOES	LIMIT
Identity checks	SSO and MFA for organizations	Protects access, not intent
IP allowlist	API only from approved networks	Has limited exceptions
Usage monitoring	Matches actions to authorization	Automated checks can err
Sandbox and review	Recommended for risky workflows	A recommendation, not a duty
Critical systems blocks	Blocks help against OT and ICS	OpenAI: 100% is not the goal

As described by the vendor. Source: OpenAI, OpenAI.

If a model is classified as Critical, the key question is who may use it. OpenAI has split access into layers. GPT-6 Astra already helps with code review and patches. Daybreak Blue gives verified defenders fewer unnecessary refusals, and Daybreak Red enables exploit development and authorized testing.

The conditions are serious: identity verification, single sign-on, multi-factor authentication and a default limit to approved IP addresses, with rare exceptions. The company recommends sandboxes, action review and human oversight for high-risk workflows.

Special attention goes to critical infrastructure. The new Daybreak Red for Critical Infrastructure group gives approved defenders permission to test critical systems. Dragos tested the models in a controlled lab on a simulated power plant, and its Chief AI Officer says the model found two zero-day flaws in engineering software. OpenAI says it used this feedback to refine its safeguards.

The same week the company showed the other side: bigger rewards for outside researchers. More on page 40.

SOURCES

1 **OpenAI**, GPT-6 Cyber: Advancing the frontier of aligned cybersecurity agents, 5 Nov 2026
openai.com

2 **OpenAI**, GPT-6 Cyber System Card, 5 Nov 2026
deploymentsafety.openai.com

BUILD / INFRASTRUCTURE OF AI PROGRAMMING

Code is no longer the bottleneck. Infrastructure is.

In one week OpenAI showed its own Git backend and announced the purchase of code review tools, while Meta took over a platform for parallel agents.

	INFRASTRUCTURE	TOOLS AND REVIEW
2 Nov	**OpenAI** Pando, a backend for Git repositories INTERNAL SYSTEM	**GitHub** New rules for pull_request_target BROWNOUTS 9 AND 16 NOV
3 Nov	**Vercel** CI without a linked project AVAILABLE	**Vercel** Security Review PUBLIC BETA
4 Nov	**OpenAI** Announced purchase of Pierre PENDING CLOSING	**Meta** Conductor joins Meta CLOUD SHUTS DOWN 20 NOV
5 Nov	**GitHub** Copilot in Xcode, from 4 Nov PUBLIC PREVIEW	**GitHub** SpaceX in secret scanning AVAILABLE

Status as of the cutoff, 5 November.



BUILD / OPENAI / 2 NOV 2026

Git built for agents

Pando is the internal backend that runs all development in the OpenAI monorepo, with the same Git client. It is not a public product.

Late last year the systems carrying the OpenAI monorepo began to fall behind. Fetches and pushes were slow, merges failed, and reliability was at times below 90%. The reason was Codex: one engineer now has several changes ready for CI, and every push starts jobs that fetch the same commit.

The answer is called Pando. It is a new backend for Git repositories, in production since April, and today it runs all development in OpenAI's monorepo. The Git protocol and client have not changed, so engineers and agents work as before.

According to the company, median push latency fell nearly five times, or nearly 80%. The system delivers close to a petabyte of Git data a day, and in one hour it sent more than 57 terabytes, with over 99% coming from a local disk cache.

Pando is neither a public service nor a GitHub competitor. The post says only that the company is exploring ways to share the technology with users. Some of the improvements are already accepted into the main Git project.

NUMBER

nearly 5

×

shorter median push time. OpenAI

NUMBER

57 TB

delivered in one hour, over 99% from cache. OpenAI

SOURCES

1 **OpenAI**, Scaling Git at OpenAI, 2 Nov 2026
openai.com

BUILD / OPENAI / HOW IT WORKS

Three decisions behind Pando

Pando gives up a consistent view of the whole repository to gain parallel writes.

THREE APPROACHES TO SCALING GIT

APPROACH	WRITES	READS	WEAK POINT
Three-phase commit	Majority of replicas	Local, consistent	Disk and latency
Write-ahead log	Shared manifest	Replicas read the log	Reads fall behind
Pando	Independent per branch	Ready packs from cache	No global moment

Comparison as described by OpenAI. The interactive simulation in the post is a teaching model, not a benchmark. Source: OpenAI.

A classic Git server keeps the state of the whole repository consistent at all times. Two standard scaling schemes, three-phase commit and a shared write-ahead log, must in the end order all writes at a single point. That is why pushes to unrelated branches wait for each other.

Pando abandons that assumption. Each branch is updated independently, and a conflict arises only when two clients push to the same ref at once. The price is lower consistency: a reader looking at several branches can see them at different moments. OpenAI calls this causal consistency and says users and agents do not actually need a global order.

The second trick is reuse. The pack a client sent is stored unchanged and served to later readers, instead of the server building a separate pack for each. When a client pushes a new branch, the server advertises only trunk, so the median pack is under 20 KiB. In the background packs are merged across five levels, from L0 to L4.

The system was written in Rust, from start to production in 10 weeks, with Codex. One engineer designed the core, teams of agents implemented it, and a handful of engineers deployed it. Migration ran without downtime, through a layer that kept the old and new systems in sync.

BUILD

Language
Rust

To production
10 weeks

Live since
April 2026

Packs
L0 to L4

Fetch
10 Gbps

SOURCES

1 **OpenAI**, Scaling Git at OpenAI, 2 Nov 2026
openai.com

BUILD / OPENAI / 4 NOV 2026

OpenAI buys the code review layer

The announced purchase of Pierre complements Pando: one is storage and transfer, the other is the surface where a person sees what agents have done.

WHAT PIERRE BRINGS

Code Storage A service that stores code and history through Git	**Trees** Fast navigation through project files
Diffs Viewing changes, with in-place editing	**Highlights** Syntax highlighting as an agent writes

Two days after the post on Pando, OpenAI announced on 4 November that it will buy Pierre Computer Company. Pierre is behind the Code Storage service and a family of open source tools. Anyone who has browsed files or reviewed changes in the Codex app has already used its Trees and Diffs libraries.

For now this is an announcement. Closing depends on customary conditions and regulatory approvals, and until then OpenAI and Pierre are separate companies. No price has been published.

The strategy is visible from the sequence. Pando solves storage and transfer. Pierre brings the surface on which a person understands what agents have done: a file tree, a view of changes, editing directly in the diff and fast syntax highlighting. Code Storage is also an API that gives other apps code and history storage through Git.

OpenAI writes that it wants to set the standard from the first branch an agent creates to the moment code goes to production. Open source projects stay open, and Code Storage customers remain on the service while its features move into Codex.

Whoever controls change review controls the last step before the merge.

STATUS

Announced 4 November. Closing depends on regulatory approvals. Until then the companies are separate.

SOURCES

1 **OpenAI**, OpenAI to acquire Pierre Computer Company, Inc., 4 Nov 2026
openai.com

BUILD / CONDUCTOR / 4 NOV 2026

Conductor joins Meta

The cloud service shuts down in two weeks. The local app keeps working until early 2027.

WHAT CHANGES

	CONDUCTOR CLOUD	LOCAL APP
New users	No, from 4 Nov	No, from 4 Nov
Work	Read-only from 13 Nov	As normal
Shutdown	**20 Nov**	Starts early 2027
Data	Deleted from 20 Nov	Deleted at shutdown

Conductor is a tool for working with several coding agents at once, launched in July 2025. On 4 November it announced that it is joining Meta. A new product is promised for early next year, with no details.

For existing users, the dates matter. From 4 November there are no new sign-ups or downloads, and pro features are free for existing users. On 13 November Conductor Cloud goes read-only and active subscribers get a refund. On 20 November Cloud shuts down and data deletion begins.

The local app does not shut down that day. It keeps working and receiving new models and critical fixes until early 2027, when its wind-down begins and the data is deleted.

WHAT TO DO

Move a Cloud workspace to your Mac by right-clicking it and choosing Transfer to local. Download an archive of all your data under Settings, then Account, then Export data. Do this before 13 November, while you can still continue working in chats, and by 20 November at the latest.

The value of the deal is not in one product. Meta is buying a team that was among the first to show what an interface for managing many agents at once looks like.

This is also a reminder for every team working with new agent tools: before you tie a workflow to a small service, check how chats, branches and settings can be exported.

DEADLINES

4 Nov	No new sign-ups
13 Nov	**Cloud read-only**
20 Nov	**Cloud shuts down**
early 2027	Local app winds down

SOURCES

1 **Conductor**, Conductor is joining Meta, 4 Nov 2026
conductor.build

BUILD / COMPARISON / NOVEMBER 2026

Four layers, four kinds of product

Pando, Scion, Pierre and Conductor solve different parts of the same workflow. That is why they should not be ranked.



KINDS OF SOLUTION

NAME	KIND	GIT	STATUS
Pando	Source control backend	Protocol kept	Internal
Scion	Own version control	Replaced	Internal
Pierre	Components and storage	Built on Git	Purchase announced
Conductor	Agent orchestration	Works over Git	At Meta, Cloud closing

Four names are often put in the same row in the news, but they do not do the same thing. Pando and Scion are source code systems. Pierre builds components and storage that sit around Git. Conductor was a user layer for managing agents.

The biggest difference is between Pando and Scion. OpenAI kept the Git protocol and changed what is behind it. Anthropic described on 13 October how it built its own system and fully moved its monorepo off Git in five months, because its clients are almost always agents and jobs that are permanently online. There the server makes commits and merges, and a local copy is optional.

Neither system is a product you can install. Scion has not been released, and Pando remains internal. They should be read as early signals. Anthropic itself says it does not expect most organizations to have these problems in the near term.

For most teams the message is more modest. Before you buy more agents, check what happens to your CI, review and merging when the number of changes doubles.

BUILD / GITHUB / 3 AND 4 NOV 2026

Copilot in Xcode, BYOK and agent skills

Three smaller changes that point the same way: agents enter existing tools, while costs and dependencies become visible.

COPILOT ENTERS XCODE FROM THE OTHER SIDE

Since 4 November GitHub Copilot has been available in Xcode's built-in coding assistant, in public preview, starting with Xcode 27.3 beta. It needs a Mac with macOS Tahoe 26.6 or later and is set up under Xcode, Settings, Intelligence. All Copilot plans are supported. This does not replace the existing GitHub Copilot for Xcode extension, which remains a separate option.

WHO IS SPENDING TOKENS GITHUB DOES NOT BILL

Since 3 November the Copilot usage metrics API separates work with your own key from models hosted by GitHub. It tracks two things: local BYOK, where a developer uses their own key or a local model in the client, and enterprise custom models set up by an admin. Reports exist at user and aggregate level, for one and 28 days. Financially this matters because the cost at a provider was not previously visible next to other usage.

DEPENDABOT NOW UPDATES AGENT SKILLS TOO

Dependabot supports version updates for Agent Package Manager. It opens pull requests for supported Git dependencies in apm.yml and regenerates apm.lock.yaml. It covers versions, not security alerts. Agent skills, prompts and instructions thus become dependencies like any package, with the same review and the same risks.

STATUS

Xcode
public preview

BYOK API
available

Dependabot APM
GA

SOURCES

1 **GitHub**, GitHub Copilot is now available in Xcode's coding assistant, 4 Nov 2026
github.blog

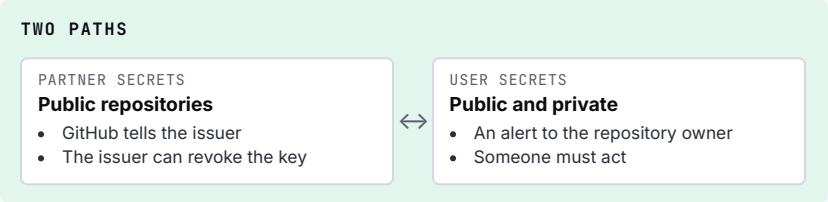
2 **GitHub**, BYOK and enterprise custom models in the Copilot usage API, 3 Nov 2026
github.blog

3 **GitHub**, Dependabot version updates support Agent Package Manager, 3 Nov 2026
github.blog

BUILD / GITHUB / 5 NOV 2026

Who gets the message when a key leaks

Secret scanning gained SpaceX as a partner. More important than the new entry is the difference between two notification paths.



On 5 November GitHub expanded secret scanning. SpaceX is a new partner, and GitHub now detects the secret starlink_public_api_v2_secret. Detectors for Alibaba, AWS, Google, OpenVSX, Pydantic Services and Tencent gained extra false-positive filtering.

The difference worth knowing is who gets the notification. With partner secrets, GitHub automatically forwards a finding in a public repository to the issuer, who can revoke or rotate the key before it is abused. With user alerts the notification goes to the repository owner, in public and private repositories.

So the partner program protects the issuer and the user even when the repository owner does nothing. An alert in the repository, by contrast, is work someone has to do.

CodeQL 2.27.4, released on 4 November, brings more accurate queries for several languages. The log injection query for Java recognizes additional sanitizers, and false positives are reduced for Rust. These are small changes, but false findings are what make people stop reading reports.

SOURCES

1 **GitHub**, Secret scanning upgrades detectors, adds SpaceX as partner, 5 Nov 2026
github.blog

2 **GitHub**, CodeQL 2.27.4 improves accuracy and query metadata, 4 Nov 2026
github.blog

BUILD / VERCEL / 2 AND 3 NOV 2026

Every push goes through an AI review

Vercel Security Review is a public beta for Pro and Enterprise. It is a useful extra check, but not a replacement for human review.



On 2 November Vercel launched Security Review, in public beta for Pro and Enterprise teams. Once enabled, the system reviews every push to a connected GitHub repository. It starts from push events, not pull requests, so it also sees direct and force pushes to the production branch.

It looks only at the changes, together with the surrounding code, and searches for flaws such as missing authorization checks or exposed secrets. Vercel says a review takes seconds. Findings appear in the Security dashboard and, on private repositories, also as a GitHub check and comments. On public repositories details stay in the dashboard.

The severity rating is AI-generated, and teams can write their own policies in plain language. The Fix with Vercel Agent button opens a pull request, and nothing merges without approval.

Two things should be said clearly. Vercel itself writes that this is not a security audit, that it does no runtime testing and does not guarantee the absence of vulnerabilities. And it is not free: reviews are free until 10 November, after which you pay provider tokens plus $0.25 per million tokens, so the cost grows with the number of pushes.

IN BRIEF

Status
public beta

Plans
Pro, Enterprise

Free until
10 Nov

After
$0.25 per 1M tokens plus provider

SOURCES

1 **Vercel**, Introducing Vercel Security Review, 2 Nov 2026
vercel.com

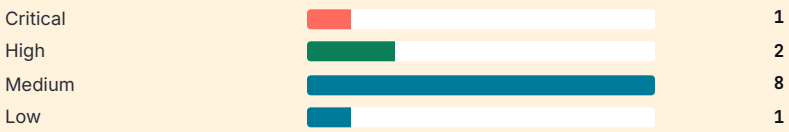
2 **Vercel**, Vercel Security Review now in public beta, 3 Nov 2026
vercel.com

BUILD / BETTER AUTH / 3 NOV 2026

Better Auth: 12 advisories at once

If you use the library, upgrade it together with all plugins. The WAF on Vercel reduces exposure, but does not replace the patch.

ADVISORIES PUBLISHED ON 3 NOVEMBER, BY SEVERITY



Counted from the GitHub Security Advisories database on 5 Nov. Source: Better Auth.

Better Auth is a popular authentication library for TypeScript. On 3 November the team published a coordinated security release with 12 advisories, and Vercel issued a bulletin and applied WAF protection for apps on its platform. The authors and Vercel still recommend upgrading.

The patched versions are 1.7.16 and 1.8.0-beta.5, and 1.7.17 and 1.7.18 have appeared since. Separate packages, such as @better-auth/expo, electron, passkey, stripe and kysely-adapter, carry the same version numbers and should be updated together.

The most serious advisory is for @better-auth/expo, where the authorization proxy can sign values it did not issue. Behind it are a problem in which desktop sign-in can end up returning the server secret and skipping the second factor, and phone numbers that can be claimed before verification.

Affected ranges differ by advisory, so read the ones for your plugins. Some fixes also require an extra configuration step, such as listing trusted origins. The authors say their reviews now also include frontier AI models, available through open source security programs.

VERSIONS

1.7.x
1.7.16+

1.8 beta
1.8.0-beta.5+

SOURCES

1 **Better Auth**, Security release on Tuesday, November 3, 2026, 3 Nov 2026
better-auth.com

2 **Vercel**, Better Auth November 2026 security release, Nov 2026
vercel.com

BUILD / VERCEL / NOVEMBER 2026

Small changes at Vercel

WHAT IS NEW

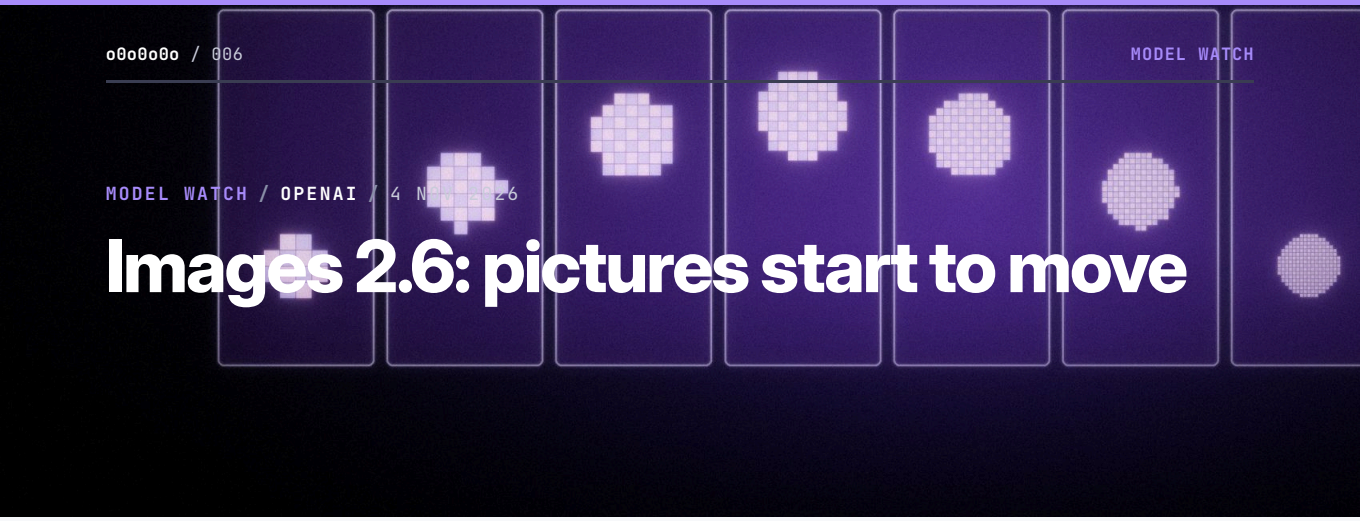
CHANGE	DATE	WHAT IT MEANS
vercel.toml	2 Nov	Configuration in TOML, with comments. Use only one of vercel.toml, vercel.json or vercel.ts.
vercel doctor	3 Nov	Checks the CLI version, login, team, project link and vercel.json, and suggests fixes.
Vercel CI	3 Nov	Runs standalone, without a linked Vercel project.
Workflow SDK	4 Nov	Idempotent starts: a duplicate attempt with the same token throws an error and never becomes a run.
Kimi K3.1 Pro	4 Nov	Pro and Highspeed on more Zero Data Retention providers, through AI Gateway.
Tencent Hy4	3 Nov	Available on AI Gateway, at 50% off.
Schedules	5 Nov	Scheduled tasks, in public beta.

None of these changes is a revolution, but together they show where the platform is going: less manual setup and more work for agents. Vercel also expanded AI Gateway, but a model's availability through a gateway is not the same as its release. Kimi K3.1 Pro came out earlier at Moonshot AI, and Vercel is now adding providers that work without retaining data.

WHAT THIS MEANS IN PRACTICE

The most useful small item is vercel doctor, because it turns the usual hunt for a setup error into a single command. The vercel.toml file is welcome for teams that want comments in their configuration, which JSON does not allow. Idempotent starts in the Workflow SDK solve a real problem: when the same event is delivered twice, the duplicate attempt no longer creates a new workflow, which matters for payments and messages.

Standalone Vercel CI shows the platform wants to be useful even where an app is not deployed on it. The same logic stands behind the announcement that Security Review is coming to repositories that do not deploy on Vercel.



MODEL WATCH / OPENAI / 4 N[illegible]26

Images 2.6: pictures start to move

The new image generation reaches ChatGPT, Codex and the API, with GIFs, SVG export and two models for developers.

On 4 November OpenAI presented Images 2.6, a new generation for creating and editing images in ChatGPT, Codex and the API. The company claims generation is 50% faster than before and that its users make 4 billion images a week.

The most visible novelty is animated GIFs: a brand mascot, reactions for social media or short animated sketches before video production. In examples published by OpenAI, the prompt describes a sprite sheet, a grid of six frames that form a loop. The new model also brings more precise editing over several rounds, consistency of characters and products, and reference images.

For developers the API changes are more concrete. Image markup, already known from ChatGPT, lets you mark exactly what should change. SVG export gives a vector result for web and print. A new complexity parameter sets how much the model plans a scene, separately from the quality parameter that sets detail.

Images 2.6 is available globally, on all subscriptions, on desktop, mobile and web. That makes it one of the few big items this week that really reached everyone.

NUMBER

4 bn

images a week, according to OpenAI. OpenAI

AVAILABILITY

ChatGPT
all plans

Codex
all plans

API
Flare, Sunburst

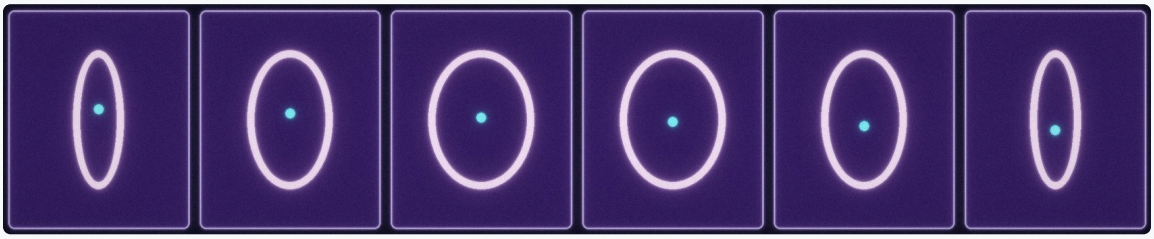
SOURCES

1 **OpenAI**, Introducing Images 2.6, 4 Nov 2026
openai.com

MODEL WATCH / OPENAI / FLARE, SUNBURST

Flare or Sunburst

Both models have the same price per token. The difference is in speed, quality and consistency across edits.



o0o0o0o original: six frames of one loop, the same principle as a GIF sprite sheet. 0000000

TWO MODELS

	FLARE	SUNBURST
Use	Fast work, posters, products	Detail and consistency
Quality	Works at low	Medium or high
Speed	From six seconds, low	Not stated
Price, 1M tokens	$8 in, $30 out	$8 in, $30 out

According to OpenAI. Source: OpenAI, OpenAI.

There are two models in the API. GPT-Image-2.6 Flare is faster and, according to OpenAI, lays out posters and product visuals better. At low quality it can make an image in "as little as six seconds". That is the best described case, not a guaranteed time for every image.

GPT-Image-2.6 Sunburst is for work where detail and consistency across edits matter, at medium or high quality. For most edits OpenAI recommends low complexity, and medium or high for scenes with many elements and constraints.

Prices are the same for both models, per million image tokens: $8 for input, $30 for output and $2 for cached input. OpenAI does not publish a per-image price, because the token count depends on size and quality. The cost is therefore read from the usage field in each response.

For teams working with a brand, SVG is the most interesting part. A vector can be changed and scaled without loss, but a designer must check it, like any other automatic export.

MODEL WATCH / GOOGLE / 22 OCT AND 2 NOV 2026

Google Meet breaks language barriers

Speech translation in Meet expands to 63 languages, including Serbian, and this week reaches a larger group of users.

Speech translation in Google Meet has so far worked only for a limited set of language pairs, one pair per meeting. On 22 October Google announced that the feature is expanding to 63 languages, and Serbian is on the official list. It enters this issue because a gradual rollout to Scheduled Release domains began on 2 November and lasts up to 15 days. Rapid Release domains got it back in October.

It is powered by the Gemini 3.5 Live Translate model. Meet detects the speaker's language and translates it into the language each listener has chosen, in near real time, even when different languages are used in the same call or a speaker switches language. The translation comes in a voice similar to the speaker's.

This is not the same as translated captions, which only show text, and it is not the Google Translate app. Captions can accompany translated speech as a supplement.

Each participant turns the feature on during the meeting. It is off by default, and translating your voice needs your consent.

NUMBER

63 languages

for speech translation, with Serbian on the list. Google Meet Help

ROLLOUT

22 Oct	Announced, Rapid Release
2 Nov	**Scheduled Release, up to 15 days**

SOURCES

1 **Google Workspace Updates**, Real-time speech translation in Google Meet expands to 63 languages, 22 Oct 2026 workspaceupdates.googleblog.com

2 **Google Meet Help**, Use Speech Translation, Nov 2026 support.google.com

MODEL WATCH / GOOGLE / LIMITS

Who can use it and what can go wrong

Plans, duration, privacy and translation errors.



Access depends on the organizer. If a meeting is organized by a user with an eligible Google Workspace or Google AI plan, translation is available to all participants, including guests from other domains. In a direct call it is enough for one side to have a plan.

There is also a time limit: up to 90 minutes per meeting for base Workspace plans and Google AI Pro, and up to 8 hours for AI Expanded Access, AI Ultra Access and Google AI Ultra.

The feature does not work in live streams, recordings, client-side encrypted meetings, phone dial-in or with adaptive audio. It is not available to people under 18. A room with Meet hardware can listen to translation, but speech from the room is not translated for others.

DOES IT WORK IN SERBIA

Serbian is supported as a language, but that is not the same as availability. Google mentions regional restrictions, and the official documentation contains no list of countries. So we cannot confirm that the feature works for all accounts in Serbia. Check for the translation button in your own meeting.

Google says it does not store voice and does not train models on it. The delay is a few seconds, and the documentation warns of errors: grammar, the gender of nouns, unexpected accents and changes in voice style. Check important agreements in writing.

DURATION

Base plans, AI Pro
90 min

AI Ultra
8 hours

SOURCES

1 **Google Workspace Updates**, Real-time speech translation in Google Meet expands to 63 languages, 22 Oct 2026
workspaceupdates.googleblog.com

2 **Google Meet Help**, Use Speech Translation, Nov 2026
support.google.com

MODEL WATCH / META / 5 NOV 2026

Muse gets a voice and a face

Meta introduces real-time conversation and avatars, at first only in the US.

THREE DISPLAY MODES

Voice only	Avatar alongside	Full screen
Conversation with no picture, for example on a walk	Expressions and gestures follow the tone	A face-to-face conversation

On 5 November Meta began rolling out voice conversations and expressive avatars in the Muse app, on mobile, web and Mac. A preview was shown at Connect 2026, and a conversation starts with the phone icon in the top right corner.

The conversation runs in real time and the agent can be interrupted. Muse meanwhile works in the background while the user talks, for example on a walk or while making dinner. The avatar matches its expressions and gestures to the tone of the conversation, and the user can pick voice only, the avatar beside the conversation or a full-screen view. Hundreds of voices are available.

The feature is available only in the US to start, so this is not a global launch. Support for AI Glasses and Muse Charm is promised for later, with no date.

Meta's difference is that it combines voice with an avatar and, according to the announcement, with tasks carried out during the conversation. How well this works has not yet been independently checked.

WHAT TO WATCH

In its post Meta says nothing about price, usage limits or how voice conversations are stored, so those questions stay open until the feature reaches more users.

The breadth of the system is more interesting. If voice also reaches AI Glasses and Muse Charm, the agent is no longer an app you open, but a constantly present conversation partner. Then privacy becomes a more important question than how convincing the avatar is.

IN BRIEF

Where
mobile, web, Mac

Region
US, to start

Voices
hundreds

Later
AI Glasses, Muse Charm

SOURCES

1 **Meta**, Give Your Muse a Voice, 5 Nov 2026
about.fb.com

MODEL WATCH / 31 OCT 2026 → 5 NOV 2026

Models of the week

Models and variants that appeared or expanded from 31 October to 5 November. Quality and speed figures are vendor claims.

GPT-6 Cyber

OPENAI · 5 NOV 2026

A specialized model for authorized vulnerability research, a variant of GPT-6 Astra. [OpenAI]

AVAILABILITY	CHANNELS	PRICE	BENCHMARK
Selected Daybreak Red customers	API, Codex	*not stated in the post*	ExploitGym: 48.5%

BENCHMARK STATUS	PREPAREDNESS
VENDOR	Critical, cyber

GPT-Image-2.6 Flare, Sunburst

OPENAI · 4 NOV 2026

Two API image models, with GIFs, SVG export and a complexity parameter. [OpenAI] [OpenAI]

AVAILABILITY	PRICE PER 1M	SPEED	BENCHMARK
API; ChatGPT, Codex	$8 in, $30 out	From six seconds, Flare, low	*not published*

Gemini 3.5 Live Translate

GOOGLE · 22 OCT 2026

The model behind speech translation in Meet. Wider rollout from 2 November. [Google Workspace Updates]

AVAILABILITY	LANGUAGES	BENCHMARK
Google Meet	63	*not published*

Kimi K3.1 Pro, Highspeed

MOONSHOT AI · 4 NOV 2026

Not new models. Vercel added Zero Data Retention providers on AI Gateway. [Vercel]

AVAILABILITY	CHANGE	BENCHMARK
Vercel AI Gateway	More ZDR providers	*not part of the post*

Tencent Hy4

TENCENT · 3 NOV 2026

Available through Vercel AI Gateway at 50% off. [Vercel]

AVAILABILITY	BENCHMARK
Vercel AI Gateway	*not part of the post*

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

MODEL WATCH / GOOGLE, OPENAI / 2 AND 5 NOV 2026

From sketch to app, from test to life

Two small stories about what AI should do in the real world.

A SKETCH BECOMES AN APP

On 5 November Google published a demonstration: you photograph a sketch on paper, pick the Canvas tool and ask Gemini whether it can turn it into an app. This is not a new model but a useful workflow. What comes out is a prototype, which should be checked before real use like any generated code.

LIFEVAL: A TEST FOR EVERYDAY LIFE

On 2 November OpenAI presented LifeVal, a way to measure how much assistants help with everyday tasks. The starting points are American and international studies of how people spend their time, along with a new representative survey of 3,995 respondents. Examples include choosing a kindergarten after a move, organizing neighbors before a flood, looking for a home and a meal plan for a parent with a new diabetes diagnosis.

Math and programming tests measure whether a model knows the right answer. LifeVal tries to measure whether the job got done, step by step, with real constraints and tools. Such a test is closer to what a user wants, but it is harder to say what counts as correct. OpenAI is the author of both the test and the models it measures, so results should be read as vendor results.

NUMBER

3,995

survey respondents behind LifeVal. OpenAI

SOURCES

1 **Google**, Turn sketches into apps with Gemini, 5 Nov 2026
blog.google

2 **OpenAI**, LifeVal, 2 Nov 2026
openai.com

SECURITY SIGNAL / BITDEFENDER / 5 NOV 2026

GhostMap: an app with no code that breaks a phone

Bitdefender showed in a test that a crafted APK can leave Android in a crash loop. Mass exploitation has not been confirmed.

On 5 November Bitdefender Labs published an analysis of a problem it calls GhostMap. In a controlled test the researchers installed a specially crafted app with no code and not a single permission, which the user does not even have to launch. The phone then stayed trapped in a crash loop, and safe mode did not help.

In practice recovery means a factory reset, which erases data. The exception is a device with USB debugging already enabled, where a computer can sometimes delete the file over ADB between two crashes, and only before the first full restart.

The problem was reproduced on a Pixel 10 with Android 17 and the 5 August patch, and Bitdefender says it affects all versions from Android 10 up. There is no patch as of publication. Google closed the first report as not reproducible, and rated the second as outside the scope of a security fix, forwarding it to a team that may consider it for a future version.

It is worth keeping a cool head. Bitdefender says the flaw does not execute code, read data or grant permissions. It is a denial-of-service attack. The persistent crash loop was built by the researchers themselves, and they have not seen it used against real victims.

STATUS

Type
Denial of service

Versions
Android 10 and later

Patch
None at publication

Attacks
Not observed

SOURCES

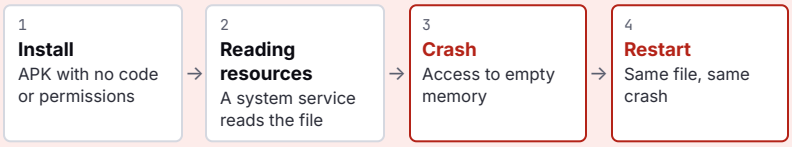
1 **Bitdefender Labs**, GhostMap: The Android flaw that crashes your phone before you open the app, 5 Nov 2026
bitdefender.com

SECURITY SIGNAL / ANDROID / HOW IT WORKS

When two parsers disagree

A simplified view of the problem, without details that would enable abuse.

FROM INSTALL TO LOOP (SIMPLIFIED)



The discovery began by accident. In August an APK with a banking trojan, collected in the wild, crashed Bitdefender's emulators during installation. Everything suggests its author did not even know about this effect.

An APK is really a ZIP archive, and each entry in it states a compression method and a size. Android reads this data in several libraries, and they disagree about what to check. One checks the length only for uncompressed entries, while the other maps everything that is not compressed in the usual way. An entry with an unusual label falls right between those two rules.

As a result, the system maps far more into memory than the file really contains. In the test, an APK of a few kilobytes was mapped across about 16 MB. When a process reads that empty part, the kernel kills it.

If that process hosts most of Android, the system falls, comes back up, reads the same file again and falls again. In one capture the cycle repeated 110 times.

Users have limited measures: install apps from trusted sources and keep a backup of your data. Bitdefender says its mobile product recognizes such files, which is a vendor claim.

TIMELINE

12 Aug	Discovery
20 Aug	First report, Won't Fix
25 Aug	Second report, out of scope
5 Nov	**Publication**

SOURCES

1 **Bitdefender Labs**, GhostMap: The Android flaw that crashes your phone before you open the app, 5 Nov 2026 bitdefender.com

PATCH NOW / 31 OCT 2026 → 5 NOV 2026

Patch this now

Items that need admin action. The first is confirmed as exploited, and the patch is already available.

WordPress Core before 7.1.4

CVE-2026-119977

CISA KEV 3 NOV | EXPLOITATION CONFIRMED

A vulnerability in WordPress core. CISA added it to its catalog of actively exploited vulnerabilities on 3 November, and the Canadian Centre for Cyber Security relayed the warning. That does not mean all installations are compromised.

PATCH
WordPress 7.1.4, released 28 October

ADMIN
Check the version in the Dashboard and that automatic updates work

SOURCE
cyber.gc.ca
wordpress.org

Zyxel GS1900 switches

CVE-2026-107441

ATTACKS REPORTED | LAN, NO LOGIN

Out-of-bounds write in the CGI program. An attacker with LAN access can run OS commands. Details on page 38.

PATCH
2.90 Patch 3

ADMIN
Change default passwords, keep management off the public network

SOURCE
zyxel.com
binarly.io

Better Auth and plugins

12 advisories

1 CRITICAL | NPM

A coordinated release for account takeover, session revocation and authorization checks.

PATCH
1.7.16, 1.8.0-beta.5

ADMIN
Update all packages together and read the advisories for your plugins

SOURCE
better-auth.com
vercel.com

Android 10 and later

GhostMap

NO PATCH | DOS

No patch at publication, and no attacks observed. Be careful with apps from unknown sources and keep regular backups.

RISK
Data loss from a factory reset

SOURCE
bitdefender.com

SECURITY SIGNAL / FRANCE / 2 AND 3 NOV 2026

When a cyberattack changes the temperature in the fridge

An attack on a refrigeration supplier hit three sites of the French blood service. Alarms and staff prevented damage.

On Monday, 2 November, at three sites of the French blood transfusion service EFS, temperature and defrost thresholds in plasma cold-storage chambers were raised to extremely high levels. Monitoring systems reacted, and staff moved or secured the products in time.

According to an EFS statement of 3 November, no blood product was lost, no health consequences were identified, and collection and distribution were not interrupted. The target of the attack was not EFS itself, but a refrigeration equipment supplier.

EFS did not name the supplier. The European company Coolworld Rentals confirmed an incident with remote monitoring software affecting a limited number of its units and said the vulnerability was in a third-party component. French authorities are investigating, and the attacker has not been identified.

The weekly Le Canard Enchaîné, which first reported the story, says the aim was to raise the temperature from minus 30 to as much as 110 degrees. EFS has not confirmed that figure. The case came after a wave of attacks on supermarket refrigeration in Europe, but a link between the two events has not been confirmed.

ACCORDING TO EFS

Sites
3

Products lost
none

Health impact
none identified

Attacker
not identified

SOURCES

1 **BleepingComputer**, Cyberattack disrupts blood storage refrigeration systems in France, Nov 2026
bleepingcomputer.com

2 **CBS News**, Blood banks hit by cyberattack in France days after European freezer hacks, Nov 2026
cbsnews.com

3 **NBC News**, Cyberattack on French blood bank fridges was thwarted, Nov 2026
nbcnews.com

SECURITY SIGNAL / ANALYSIS / HEALTHCARE AND OT

An alarm that does not depend on the attacked system

What hospitals and services can learn from an incident that ended without damage.

LAYERS OF PROTECTION (CONCEPTUAL)

Supplier	Remote access under strict control
Network	Segmentation from the rest of the facility
Device	Physical limits that software cannot move
Monitoring	Independent sensors and alarms
People	A rehearsed relocation plan

This case is useful precisely because it ended well. It shows what protected the plasma: not a firewall, but an independent alarm and people who knew what to do.

Hospitals and blood services depend more and more on suppliers who maintain equipment remotely. That is practical, but it means software you have never reviewed can change the physical state of your storage. A cold chamber is an industrial control system, even if nobody calls it that.

That is why digital and physical security have to be planned together. Remote configuration should be kept separate from the facility's network, with limited scope. If a value of minus 30 degrees must never cross a certain threshold, that limit should exist in the device itself, not only in software. Monitoring should be independent of the channel commands arrive through, and staff should have a rehearsed relocation plan.

The supplier is part of your system, whatever the contract says.

SOURCES

1 **BleepingComputer**, Cyberattack disrupts blood storage refrigeration systems in France, Nov 2026 bleepingcomputer.com

2 **CBS News**, Blood banks hit by cyberattack in France days after European freezer hacks, Nov 2026 cbsnews.com

SECURITY SIGNAL / ZYXEL / 3 NOV 2026

Zyxel: three sources, three different things

The vendor confirms the flaw, researchers link the attacks to it, and the device count is a third story.

WHO CLAIMS WHAT

SOURCE	CLAIM	SCORE
Zyxel	The flaw and the fix, attack from the LAN	8.8
Binarly	The captured exploit targets this flaw	9.8
GreyNoise	996 devices, 48 countries	estimate

On 3 November Zyxel published an advisory for CVE-2026-107441, an out-of-bounds write flaw in the CGI program of the firmware for GS1900 switches. An unauthenticated attacker with access to the local network can run operating system commands through a crafted HTTP request. The fixed versions are 2.90 Patch 3, for ten models, and they came out back on 22 October.

The story is more complicated than the advisory. In August GreyNoise counted 996 compromised GS1900 switches in 48 countries, 564 of them using factory passwords, and attributed this to the older flaw CVE-2026-7273. Binarly then showed that the captured exploit targets a different flaw, which the June patch did not block.

So the sources need to be kept apart. The vendor confirms the vulnerability and the fix, and Binarly links the attacks to it. The device count is GreyNoise's estimate. Zyxel scores the flaw 8.8 and stresses the local network, while Binarly gives 9.8. It is not known what firmware the attacked devices had.

A local network is not the same as the internet. But an exposed management panel or a compromised workstation turns that difference into a formality. Update, change factory passwords and keep management off the public network.

NUMBER

996

compromised devices in 48 countries, according to GreyNoise. Help Net Security

SOURCES

1 **Zyxel**, Security advisory for out-of-bounds write vulnerability in GS1900 series switches, 3 Nov 2026
zyxel.com

2 **Binarly**, Out-of-bounds write vulnerability in Zyxel GS1900 switches, Nov 2026
binarly.io

3 **Help Net Security**, Patched Zyxel GS1900 switches remained open to in-the-wild exploit, 5 Nov 2026
helpnetsecurity.com

4 **BleepingComputer**, Zyxel fixes zero-day exploited to hack nearly 1,000 network switches, Nov 2026
bleepingcomputer.com

SECURITY SIGNAL / MSECURE, NIKKE / 3 AND 4 NOV 2026

When the trusted source is the problem

A password manager with no explanation and an official launcher with someone else's code.

MSECURE: INSTRUCTIONS WITH NO EXPLANATION

On 3 November the company mSeven wrote to users of the mSecure password manager that it had found "suspicious activity". It recommended that they immediately protect everything they keep there: change passwords, re-enroll two-factor authentication and remove passkeys.

It did not say what happened, when, or how many users are affected. BleepingComputer writes that a hacker reportedly stole an encrypted database, but the company has not confirmed this. There is no evidence that vaults were decrypted. Still, a user has to act as if the data were exposed, without knowing which is not. With a password manager, transparency is part of security.

NIKKE: A COMPROMISED OFFICIAL CHANNEL

On 4 November Shift Up said scripts of the PC launcher for Goddess of Victory: NIKKE may have been altered by third parties, and the client was fixed the same day. According to Kotaku, the code was injected on 27 September into four JavaScript files on servers of the publisher Level Infinite, and was used from 30 October. The publisher confirms that computers and account data may have been exposed and advises anyone who launched the game since 27 September to install the new client with a scanner. Sources differ on the dates.

The PC launcher is affected, not all platforms. The lesson is uncomfortable: downloading from an official channel is no longer proof that software is safe.

UNKNOWN

The scope of the mSecure incident has not been publicly confirmed, and no account count has been published.

SOURCES

1 **The Register**, Password manager mSecure tells users to change passwords, won't say why, 3 Nov 2026 theregister.com

2 **BleepingComputer**, mSecure password manager hack prompts password reset warnings, Nov 2026 bleepingcomputer.com

3 **Kotaku**, Hugely Popular RPG NIKKE Was Compromised, Nov 2026 kotaku.com

4 **Automaton**, Goddess of Victory: NIKKE PC launcher reportedly compromised in supply-chain attack, Nov 2026 automaton-media.com

SECURITY SIGNAL / OPENAI, DREAMHOST / NOVEMBER 2026

Paying people to find what models miss

Rewards go up, and DreamHost lowers its estimate of the scope of the October incident.

EXAMPLE CALCULATION FROM OPENAI'S POST

BASE	QUALITY	CAMPAIGN	LOYALTY
$30,000 Critical, core	→ **× 1.15**	→ **× 1.30**	→ **× 1.15**

On 2 November OpenAI raised rewards in its security bug program. The midpoint of the ranges for medium, high and critical findings is on average more than four times higher. A critical vulnerability in ChatGPT, Codex or the API Platform earns up to $30,000.

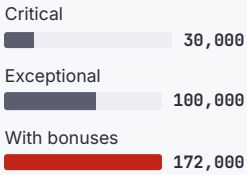
Exceptional critical findings remain at up to $100,000, and with bonuses for quality, priority campaigns and loyalty they can reach up to $172,000. That is the upper limit OpenAI states, not a guaranteed payout. The same program applies to open reports.

The message is clear. A company that itself claims models find thousands of flaws is at the same time paying people to find what models miss.

DREAMHOST, A FOLLOW-UP

The October incident is not a new story, but it has an important update. On 23 October an attacker was discovered on the File Manager server and had reached FTP and SFTP credentials of users who signed in there, as well as data of a limited number of webmail accounts. According to DreamHost, a fraction of accounts was affected, well under approximately 10%, which is less than the first estimate. The company found no evidence of access to the control panel or payment data.

UPPER LIMITS, USD



Not a guaranteed payout.

SOURCES

1 **OpenAI**, Rewarding defenders with bigger bug bounties, 2 Nov 2026
openai.com

2 **DreamHost**, October Security Incident, 29 Oct 2026
dreamhost.com

SECURITY SIGNAL / ROUNDUP / NOVEMBER 2026

Short security roundup

WORTH KNOWING

TOPIC	DATE	WHAT IS CONFIRMED
Google Calendar phishing	4 Nov	Material Security describes Operation CloudFire: calendar invites lead to fake pages through a reverse proxy that takes over the live session and so gets around MFA.
GitHub Desktop	2 Nov	Adds short-lived authentication tokens, so a stolen token stays valid for less time.
CodeQL 2.27.4	4 Nov	More accurate queries and fewer false positives, especially for Java and Rust.
Better Auth	3 Nov	Twelve advisories and WAF rules on Vercel. Page 24.
Secret scanning	5 Nov	Starlink keys and fewer false matches. Page 22.

Besides the big topics, this week it is worth watching what is less visible. The CloudFire campaign is instructive because it shows MFA does not help when an attacker takes over an already open session, and Material Security says it has blocked more than 68,000 messages. Google Calendar by default adds invitations even from unknown senders, so that setting is worth checking.

CLOUDFIRE: WHEN EVERY SIGNAL IS CLEAN

According to Material Security's analysis, messages arrived as Google Calendar invites, often from an organization the recipient had already corresponded with. The pretexts were ordinary: voicemail, requests for proposals, court summonses and unpaid invoices. Most links pointed to Google itself, so the domain, certificate and reputation looked correct.

The attacker then stood in front of the real Google sign-in page through a live reverse proxy. The victim enters a password, gets a real confirmation prompt and approves it, and the attacker gets the session. That is why SMS codes and authenticator apps do not help here. The campaign has run since mid-August.

WHAT THIS MEANS FOR YOU

Admins should watch for unusual sessions and new devices, because an attacker does not need the second factor once they take over an active session. Restricting invitations from unknown senders helps users.

LAB / ANTHROPIC / 2 NOV 2026

Can it be mathematically proven that AI really did what it claims?

FlashProve is an open source prototype for provable inference. It works on a large model, but is still far more expensive than ordinary processing.

When you ask an AI model something, you cannot check who answered. A provider could swap the model for a cheaper one, and an attacker could plant a response. On 2 November Anthropic published FlashProve, an open source prototype for provable inference, which tries to answer this question with cryptography.

The provider first publishes a short commitment, a fingerprint of exact model weights, and the weights stay secret. With each answer it then sends a proof confirming that those weights, run on that prompt, give exactly those tokens. This is a zero-knowledge proof: the verifier is convinced the computation is correct, and learns nothing about the weights.

Anthropic measured this on the open Llama 4 Maverick model, with about 400 billion parameters. A single H100 GPU generates proofs at about 110 tokens per second, and a proof for a 1,024-token request is about 19 MB. Verification takes about 17 seconds on a CPU, with no GPU.

The price is high: about 70 to 95 times more GPU time than ordinary processing of the same request. So this is an intermediate result and a prototype, not a service that can be switched on for all commercial models.

NUMBER

19 MB

proof for a 1,024-token request. Anthropic

NUMBER

70 to 95

×

more GPU time than ordinary processing. Anthropic

SOURCES

1 **Anthropic**, Provable inference and FlashProve, 2 Nov 2026 anthropic.com

LAB / ANTHROPIC / WHAT IS PROVEN

Proof of origin is not proof of truth

A correct answer and a proven computation are two completely different things.

HOW VERIFICATION WORKS

ONCE	PER REQUEST	ON THE CLIENT
Commitment	**Proof**	**Check**
Fingerprint of the weights, about 6.5 MB, 39 minutes	About 19 MB, on an H100 GPU	About 17 seconds on a CPU

FOUR DIFFERENT QUESTIONS

QUESTION	DOES FLASHPROVE HELP
Did the declared model run?	**Yes, within the protocol**
Does the proof reveal the weights?	No, but the architecture stays known
Is the answer correct?	*No, that is outside the proof*

The most important thing is what the proof does not prove. FlashProve confirms that a specific model carried out a specific computation. It does not confirm that the answer is correct. A model that hallucinates will go on hallucinating, only with proof of origin.

So four questions need to be kept apart. Trust in the provider is what we have today: we believe it serves what it claims. Execution integrity is what FlashProve adds. Model privacy means the proof does not reveal the weights, although the architecture stays known. The quality and truth of the answer stay outside all of this.

Anthropic names several uses. A regulator could check that the declared model really runs on the hardware, and an independent evaluator could know it is testing the same model users get. A proof could also expose a compromised server that plants answers.

The system was developed with the help of Claude, so all the code is open under the Apache 2.0 license, with proofs in the formal language Lean, so others can check it. Authors from Anthropic, EPFL and MIT invite others to review and improve it.



LAB / ANTHROPIC / 4 NOV 2026

Claude and the mystery of hollow particles

The first 3D structure of a vault particle from the prokaryotic world is not a drug, but a new tool for a 40-year-old question.

Vault particles are large hollow protein structures shaped like a barrel, discovered in 1986. They are found in almost all our cells, and even after 40 years their function is not definitively known. On 4 November Anthropic published the first three-dimensional structure of a vault particle from the prokaryotic world. More precisely, the protein comes from a bacteriophage, a virus that infects bacteria, so the particle is built inside a bacterial cell.

The structure has 96 copies of a single protein called L05, compared with 78 in the human version, with a total mass of about 9.5 megadaltons. It is about 54 nanometers long and about 45 wide.

Claude worked under the direction of scientists: it searched a database of about 15 billion predicted proteins, proposed DNA constructs, optimized isolation and processed cryo-electron microscopy data. Isolation that usually takes three days was cut to about 90 minutes, and a first structural model was ready in 34 hours.

This is neither a drug nor an autonomous discovery. The function is still unknown, an unidentified molecule sits in a pocket of the protein, and Anthropic also describes mistakes: a protocol with the wrong additive cost the team two weeks. The work was submitted as a preprint, and the structures and additional data are coming later.

IN BRIEF

L05
96 protein copies

Human vault
78 copies

Size
54 × 45 nm

Isolation
3 days to 90 min

STATUS

Function unknown. Preprint, no clinical use.

SOURCES

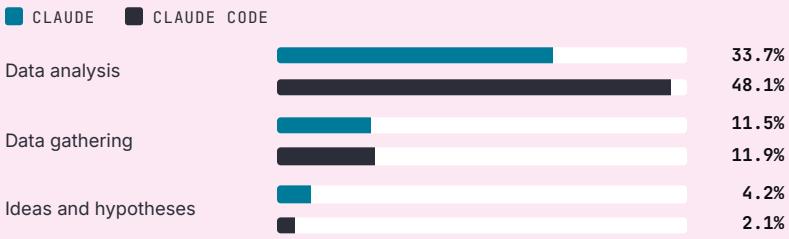
1 **Anthropic**, Investigating vault particles with Claude, 4 Nov 2026 anthropic.com

LAB / ANTHROPIC / 5 NOV 2026

What scientists really use Claude for

An analysis of 207,049 sessions shows AI does the most at the keyboard and the least where a lab is needed.

SHARE OF SESSIONS BY TASK TYPE, %



Selected categories, not the full total. Period 2 Aug to 30 Sep 2026. Source: Anthropic.

On 5 November Anthropic published an analysis of 207,049 research sessions on its consumer plans, from 2 August to 30 September 2026. Of these, 128,914 are from Claude Chat and Cowork, and 78,135 from Claude Code. These are sessions, not unique scientists, and the sample does not include other tools.

Data analysis leads: 33.7% of sessions on Claude and 48.1% on Claude Code, compared with an estimated 27.6% of researcher time. Writing and communication make up 32.0% of Claude sessions, twice their share of time. Hypothesis generation is small, at 4.2% and 2.1%, and data gathering is 11.5% and 11.9%, even though it takes 28.6% of research time.

Computer science accounts for 32.7% of Claude Code sessions, although it has 4.3% of the world's researchers. Claude Science, a specialized tool, is used above all in biomedicine and biology, and for more complex tasks.

The reason is obvious. Whoever works at a keyboard can delegate a lot, while a pipette still needs a hand. In lab sciences 38 to 41% of time goes to physical steps, compared with 4% in the median field.

SAMPLE

Claude
128,914

Claude Code
78,135

Total
207,049

SOURCES

1 **Anthropic**, How researchers use Claude, 5 Nov 2026
anthropic.com

LAB / ANTHROPIC / SCENARIOS

Up to 15 times faster, in one imagined world

Modeling is not measuring. Three scenarios describe possibilities, not achieved productivity.

ANTHROPIC'S SCENARIOS

SCENARIO	ASSUMPTION	SPEEDUP
1	Current capabilities, widely adopted	1.35× to 2.02×
2	AI does 75% of non-physical steps	up to about 8×
3	AI does all non-physical steps	*up to 15×, speculative*

Not a measured productivity gain. Source: Anthropic.

Anthropic then builds a simple model of how much AI could speed up research. These are scenarios, not measured results.

In the first, current capabilities are widely adopted. AI then takes on 27 to 52% of the time needed for a new study, and the speedup ranges from 1.35 times in the arts to about 2 times in computer science and statistics. In the second scenario AI fully handles 75% of non-physical steps, so computer science reaches about 8 times. In the third, entirely speculative, AI does all non-physical steps and some fields reach up to 15 times, while lab sciences stay limited by physical work.

So 15 times is not an achieved speedup of science, but the upper limit of one imagined world. Anthropic itself describes this as exploring possibilities.

The authors add a survey of 528 scientists, mostly from biomedical fields, who already actively use AI. About 40% expect AI to carry out whole projects on its own by 2030, about 64% by 2035, and 16% say it will never happen. These respondents are, by selection, more inclined toward AI than the average researcher.

NUMBER

38 to 41

%

of time in lab sciences goes to physical steps. Anthropic

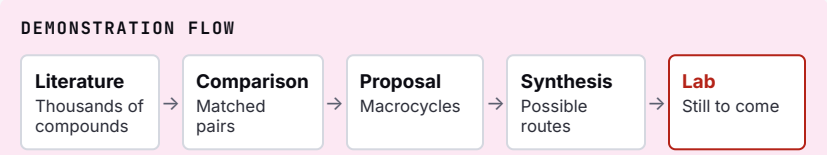
SOURCES

1 **Anthropic**, How researchers use Claude, 5 Nov 2026
anthropic.com

LAB / GOOGLE DEEPMIND / 3 NOV 2026

An hour of chemistry, not an hour to a drug

Agents in Antigravity shorten the preparation, but a proposed molecule is only the start of a long road.



On 3 November Google DeepMind showed a workflow that connects agents in the Antigravity environment with new cheminformatics skills. The division of labor matters: the language model runs the workflow, searches databases and reasons, while deterministic chemistry tools based on the open RDKit toolkit do the calculations. In chemistry an error in one atom gives a different compound.

The demonstration started from the lung cancer drugs crizotinib and lorlatinib. The agent searched the literature, gathered thousands of compounds, filtered out the irrelevant ones, compared related pairs of molecules and proposed new macrocyclic structures along with possible synthesis routes. Google says this took about an hour, while one such iteration usually takes several days.

It should be clearly said what this is not. No drug was created in an hour. The proposals are hypotheses from simulation. They still have to be synthesized, tested in a lab and go through the long road of preclinical and clinical trials.

The agent shortens the preparation, and that is genuinely valuable. The decision about what is worth making still belongs to a chemist.

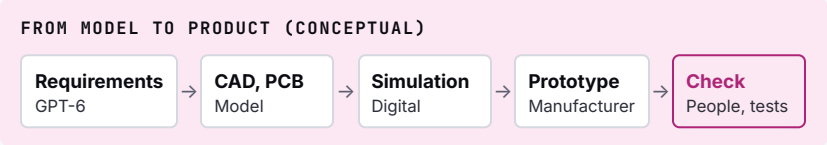
SOURCES

1 **Google DeepMind**, Agentic cheminformatics with Google Antigravity, 3 Nov 2026
deepmind.google

LAB / OPENAI / 3 NOV 2026

From CAD model to real part

Models help with mechanical and electrical engineering, but manufacturing still needs people.



On 3 November OpenAI described how its models help design physical products: from CAD models and manufacturing drawings to printed circuit boards, simulations and firmware.

The most interesting item is a test device for talking to Codex, designed by GPT-6 Astra with no direct human edits, including the enclosure, the board and the software. Two team members worked on it as a side project, for less than a week, and a small batch of prototypes was built in under a month.

OpenAI also admits what did not work. Preparing for manufacturing took several rounds of review, because Astra did not consistently account for the manufacturer's limits, tolerances and assembly details.

Chip Ganassi Racing is testing the conversion of 3D scans of old parts into CAD models, with team engineers checking the results. A new collaboration with Mazda on aerodynamics and crash safety was announced, with results to come later. In the internal BenchCAD 2.0 test a newer internal model scores 0.875, compared with 0.594 for Astra, which is a vendor claim.

A digital model is not a product. Material, tolerance and a crash show themselves only on the real part.

NUMBER

23×

growth in token use in OpenAI's hardware teams since February. OpenAI

SOURCES

1 **OpenAI**, Accelerating engineering for the physical world, 3 Nov 2026
openai.com

MONEY / AMD / 3 NOV 2026

AMD: Data Center almost doubles

AI infrastructure drives a record quarter. The 97% growth refers to one segment, not the whole company.

On 3 November AMD reported a record quarter. Revenue reached $13.377 billion, 45% more than a year earlier. GAAP operating income was $2.832 billion and net income $2.428 billion.

The driver is the Data Center segment: $8.5 billion, up 97%. That is the growth of that segment, not of the whole company. The reasons are EPYC processors and the start of shipments of the Instinct MI450 GPU series. AMD says Anthropic, Meta, Microsoft, OpenAI and Oracle are already running models on the first Helios rack systems.

The picture is not the same everywhere. Client and Gaming together fell 8%, to $3.7 billion. Client grew 7%, and Gaming fell 40%, to $773 million. Embedded grew 29%, to $1.1 billion.

For the fourth quarter the company expects about $17 billion, plus or minus 3%. That would be growth of about 66%.

NVIDIA remains the leader in the AI accelerator market. Fast growth from a smaller base shows AMD is taking a significant part of new business, but does not by itself prove a change of leader.

NUMBER

8.5 bn $

Data Center segment revenue, up 97%. AMD

Q3 2026

Revenue
$13.377bn

Growth
45%

Op. income
$2.832bn

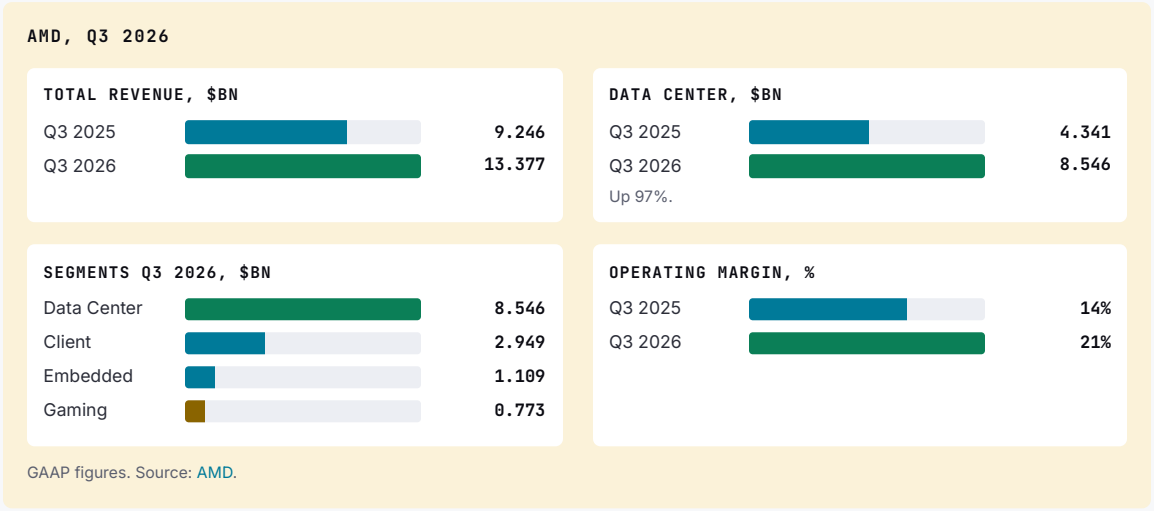
Net
$2.428bn

SOURCES

1 **AMD**, AMD Reports Third Quarter 2026 Financial Results, 3 Nov 2026
newsroom.amd.com

MONEY / AMD / INFOGRAPHIC

AMD in four charts



GAAP figures. Source: AMD.

SEGMENTS, MILLIONS OF DOLLARS

SEGMENT	Q3 2025	Q3 2026	CHANGE
Data Center	4,341	**8,546**	+97%
Client	2,750	2,949	+7%
Gaming	1,298	773	*minus 40%*
Embedded	857	1,109	+29%

The charts below use data from AMD's report for the third quarter of 2026. Operating margin rose from 14% to 21%, because revenue grew faster than operating expenses. Data Center now makes up about 64% of total revenue, compared with about 47% a year ago.

WHAT THE CHARTS SHOW

Gaming is the only segment in decline, from $1.298 billion to $0.773 billion, so AMD is much less dependent today on the console and graphics card cycle. Client grew thanks to mobile Ryzen processors.

Data Center is also the most profitable segment in absolute terms, with $2.882 billion of operating income, compared with $1.074 billion a year earlier. Lisa Su says revenue exceeded the high end of guidance and that demand for Instinct and EPYC products is strengthening.

The fourth-quarter forecast of about $17 billion means growth of about 27% over the third quarter. That is the company's expectation, not an achieved result, and it depends on how quickly Helios systems move from first installations to large shipments.

MONEY / APPLE / 2 NOV 2026

Apple: a record with no AI story in the numbers

Quarterly revenue of $114.5 billion. Keep the quarter and the fiscal year apart.

QUARTER AND FISCAL YEAR

METRIC	VALUE	PERIOD
Revenue	$114.5bn	FQ4, to 26 Sep
Gross margin	47.9%	FQ4
Earnings per share	$2.10	FQ4
Revenue	$479bn	Fiscal 2026

On 2 November Apple reported results for its fourth fiscal quarter, ended 26 September. Revenue was $114.5 billion, 12% more than a year earlier, a record for the September quarter. Gross margin was 47.9%, and earnings per share $2.10, with a favorable impact from tariff refunds, about one percentage point on margin and 5 cents on earnings.

The quarter and the year should be kept apart. All of fiscal 2026 brought a record $479 billion. CEO John Ternus cites an iPhone record for the September quarter and the highest services revenue to date.

In its release Apple highlights new devices, among them iPhone Duo and iPhone 18 Pro, and the new MacBook Pro with a touch screen from our last issue starts reaching customers on 6 November. The report itself does not show how much AI features contributed to growth.

In this story Apple is the opposite of AMD. It does not sell shovels for the AI gold rush, but devices on which AI is used, and earnings stay tied to the phone and services.

NUMBER

114.5 bn $

quarterly revenue, up 12%. Apple

NUMBER

479 bn $

revenue for all of fiscal 2026. Apple

SOURCES

1 **Apple**, Apple reports fourth quarter results, 2 Nov 2026
apple.com

2 **Apple**, Apple unveils groundbreaking MacBook Pro with PROLED XDR display, 28 Oct 2026
apple.com

MONEY / ARM / 4 NOV 2026

Arm sells blueprints, and now more than that

Revenue grows 27%, but licences grow far faster than royalties.

REVENUE BY SOURCE, $M



Fiscal Q2 2027. Source: BigGo Finance.

On 4 November Arm reported revenue of about $1.44 billion for the second quarter of fiscal 2027, 27% more than a year earlier. Growth is not even. Licence and other revenue jumped 50%, to $773 million, while royalties grew 8%, to $670 million.

To understand Arm, this difference is key. The company traditionally does not make chips, but sells intellectual property. A licence pays for access to the architecture, and a royalty for each chip produced. So licences show what is being prepared, and royalties what is selling now.

Management says data center royalties doubled, while phones are under pressure from memory prices. Large cloud companies build their own server processors on an Arm base, so Arm earns even when none of its chips sits on a shelf.

The company now also develops silicon itself: expectations for the AGI CPU have been raised to up to $2 billion of revenue over two fiscal years. That is a move from selling blueprints toward selling products. Sources here are the transcript and summaries of the investor call.

SOURCES

1 **The Motley Fool**, Arm Holdings Q2 2027 earnings call transcript, 5 Nov 2026
fool.com

2 **BigGo Finance**, Arm Holdings Q2 2027 earnings call summary, 4 Nov 2026
finance.biggo.com

INFRA / TERAWULF, GOOGLE, ANTHROPIC / 4 NOV 2026

What a 401-megawatt guarantee looks like

Google gives credit support for Anthropic's lease with TeraWulf. The capacity is planned, not operating.

WHO GUARANTEES WHAT (SIMPLIFIED)



On 4 November TeraWulf said it had obtained credit support from Google for a project in Hawesville, in the US state of Kentucky. The campus is being built for Anthropic, under a 20-year lease signed in July, worth about $19 billion. The planned capacity is about 401 megawatts of critical IT load, and that is a project, not operating capacity. First capacity is expected in the second half of 2027.

The agreement is most interesting for what it reveals about financing. TeraWulf is moving from bitcoin mining to AI infrastructure and does not have the credit rating of a big technology company. With a Google guarantee, lenders see a different risk, so debt can be raised on better terms. According to reports on the filing, Google can in case of trouble take over the lease or pay a fee, which may give it a stake in the project company.

This is one big partner vouching for the lease of its partner, and Google is also an investor in Anthropic. The same money increasingly circulates between chip, lease and model.

PROJECT

Site
Hawesville, Kentucky

Capacity
about 401 MW, planned

Lease
20 years

Value
about $19bn

First capacity
second half of 2027

SOURCES

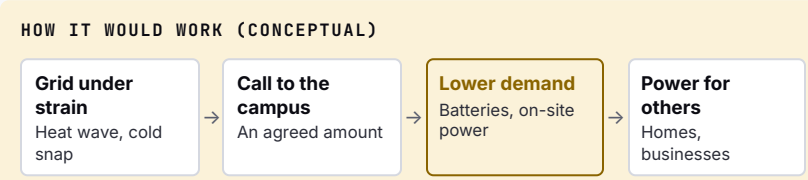
1 **Reuters via KFGO**, TeraWulf gets credit backstop from Google for Anthropic data-center project, 4 Nov 2026
kfgo.com

2 **CoinDesk**, Google could take ownership stake in TeraWulf's AI project if Anthropic defaults, 5 Nov 2026
coindesk.com

INFRA / OPENAI / 2 NOV 2026

A compute campus as a grid reserve

Grid Protectors is a program, not built capacity.



On 2 November OpenAI launched the Grid Protectors program, a collaboration with utilities, regulators and policymakers in the US. The idea is that future compute campuses reduce demand when the grid is under strain, and that utilities can use this flexibility as a resource.

The first example is the planned campus in Effingham County, Georgia, which is still in development and will serve as a pilot. The agreement with Georgia Power covers up to 1,000 megawatts of flexible demand, which OpenAI compares with the average consumption of about 700,000 households.

OpenAI promises to pay its own electricity and infrastructure costs and claims steady demand can help spread the grid's fixed costs across more customers. Families and small businesses, the company says, should not subsidize a campus.

This is a program, not built capacity. Terms will differ by location, and the company says it will publish further commitments when they are agreed. Whether local customers' bills really fall, or at least stay the same, only real prices will show.

NUMBER

1,000 MW

of flexible demand in the Georgia Power agreement. OpenAI

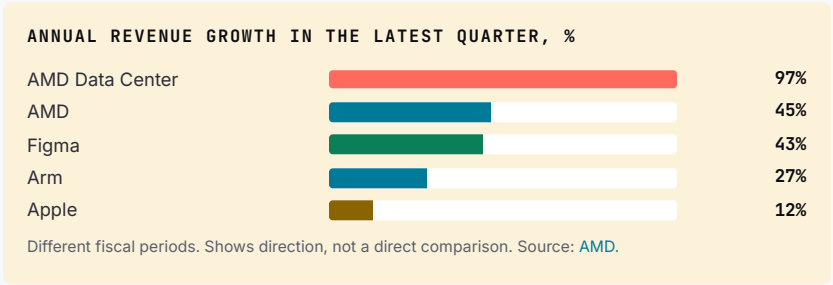
SOURCES

1 **OpenAI**, Grid Protectors Program, 2 Nov 2026
openai.com

MONEY / FIGMA / 4 NOV 2026

Who is making money in this phase

A quick look at a smaller report and what it adds to the big picture.



Different fiscal periods. Shows direction, not a direct comparison. Source: AMD.

Beside the big reports, a few smaller figures tell the same story from another angle.

On 4 November Figma reported third-quarter revenue of $391.4 million, up 43%. Under GAAP the company has a loss from operations of $119.2 million, and without excluded items it has $47 million of profit. For the full year it expects revenue of $1.497 to $1.499 billion.

Growth there is carried by software that uses AI, but the cost of AI is now an item investors watch separately. The same pressure is visible at everyone offering AI with a subscription.

Between Figma, Arm and AMD you can see the order of this phase. Those who sell capacity and architecture for building do best. Application software is growing, but has to prove it can also earn.

Growth in the chart is not directly comparable, because the periods and company sizes differ. It is useful only as an indicator of direction: those closest to infrastructure or a product sold with AI are growing fastest.

NUMBER

391.4 $m

Figma's Q3 revenue, up 43%. Figma

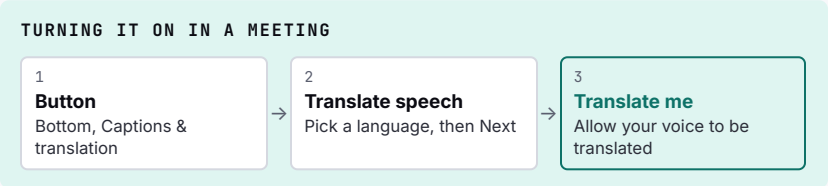
SOURCES

1 **Figma**, Figma Announces Third Quarter 2026 Financial Results, 4 Nov 2026 investor.figma.com

WORK / GOOGLE MEET / GUIDE

How to turn on translation in Google Meet

Three steps, according to the official documentation.



WHATSAPP EVENTS LEAVE THE GROUP

On 5 November Meta expanded the Events feature. A new Events Home page, in settings, gathers past and upcoming events. An invitation can be sent by link even to people with no account, who see details and RSVP without installing the app. Reminders and a plus one were added. Your number is not visible to guests you do not know, and guests with no account see the address only after they RSVP. The rollout lasts the coming weeks, so the feature will not reach everyone at once.

FITBIT EDGE

Google has released Fitbit Edge, a band that replaces the Charge 6. According to available reports, it costs $179.99 and has been on sale since 27 October. It has a 1.34-inch AMOLED screen, ECG, blood oxygen and up to seven days of battery, but no NFC for payments.

LEARNING INSTEAD OF A READY ANSWER

On 3 November Google expanded Guided Learning in Classroom, where it becomes the default chatbot and teachers can see the conversations. The mode leads a student through a task instead of giving the solution. In Search a pilot is announced in selected US schools. According to the announcement, schools will have to turn Gemini on themselves for K-12 students.

SOURCES

1 **Google Meet Help**, Use Speech Translation, Nov 2026 support.google.com

2 **Meta**, WhatsApp Events brings people together beyond the group chat, 5 Nov 2026 about.fb.com

3 **Android Authority**, Google Fitbit Edge launches with 7-day battery for $180, Oct 2026 androidauthority.com

4 **Google**, Responsible AI for schools with Guided Learning, 3 Nov 2026 blog.google



Illustration: a re-entry trail through the atmosphere and a distant gamma-ray burst. 0000000, NOT A PHOTOGRAPH

ORBIT / NASA / 4 NOV 2026

Swift has ended its mission

An observatory built for two years worked for almost 22, and its re-entry was confirmed on 4 November over Sumatra.

After almost 22 years in orbit, NASA's Neil Gehrels Swift Observatory entered Earth's atmosphere on 4 November, at 12:39 UTC, over northern Sumatra. The US Space Force confirmed the re-entry at about 3.4 degrees north latitude and 98.6 degrees east longitude. Most of the spacecraft was expected to burn up, and some parts may have survived.

Swift launched in 2004, with a plan for a two-year prime mission. It was meant to hunt gamma-ray bursts, the most powerful explosions in the cosmos. Three telescopes, for gamma rays, X-rays and ultraviolet and optical light, let it turn quickly toward a source and send an alert to other observatories.

Over the decades Swift became an all-purpose tool for astronomers, from active galaxies to comets in our own system.

The end came from atmospheric drag, increased by solar activity. A commercial mission meant to raise the orbit ended in August with no attempt. On 20 October, at an altitude of 275 kilometers, the team at Penn State sent the final command and ended operations.

SWIFT

Launched
2004

Plan
2 years

Re-entry
4 Nov, 12:39 UTC

Place
North Sumatra

SOURCES

1 **NASA**, NASA's Swift Observatory Re-Entered Atmosphere, 4 Nov 2026
science.nasa.gov

2 **NASA**, NASA's Swift Mission Ends Ahead of Re-Entry, 23 Oct 2026
science.nasa.gov

ORBIT / NASA / 2004 → 2026

Twenty-two years on one line

The mission timeline and the failed rescue attempt.

THE SWIFT MISSION

2004	Aug 2026	Sep 2026	20 Oct	4 Nov
Launch	**Boost called off**	**Back to science**	**275 km**	**Re-entry**
Plan: two years	No boost attempted	Hundreds of targets	End of operations	12:39 UTC

The Swift story also has a business side. In 2026 NASA tried something unusual: having a commercial spacecraft reach the observatory and raise its orbit. To buy time, the team paused science observations to reduce drag.

The mission ended in August with no boost attempt. Swift then returned to science and managed to observe hundreds more targets, from the remnants of stellar explosions to comets. That work increased drag, so the spacecraft again began to sink quickly.

For space telescopes this is a lesson in how long missions depend on things invisible at the start. A spacecraft built for two years worked eleven times longer, and ended because of the Sun, not a failure.

WHAT SWIFT LEAVES BEHIND

Three instruments were designed to work together. After detecting a burst the spacecraft turned toward the source on its own, and the other telescopes captured what remained. That way of working, with alerts to other observatories at the same moment, was and remains a model for following short-lived events in the sky.

SOURCES

1 **NASA**, NASA's Swift Mission Ends Ahead of Re-Entry, 23 Oct 2026
science.nasa.gov

2 **NASA**, NASA's Swift Observatory Re-Entered Atmosphere, 4 Nov 2026
science.nasa.gov

ORBIT / NASA / 4 NOV 2026

X-70 Fractal: an aircraft that swaps its equipment



NASA unveiled a modular flight research platform. Its first flight under the NASA banner is expected within two years.

On 4 November, at the UP.Summit at Kennedy Space Center, NASA presented the X-70 Fractal experimental aircraft. It is a crewed research platform for developing and testing advanced flight-control and maneuverability technologies.

The key word is modularity. The aircraft carries up to about 907 kilograms of experimental equipment, in internal and external positions, and the equipment can be swapped quickly. So one aircraft can serve a series of different experiments, instead of a new one being built for each. Its flight control can be reconfigured quickly, which lets NASA explore maneuverability in ways not yet tried in crewed flight.

According to NASA, takeoff weight is about 4,536 kilograms, cruise speed Mach 0.75, ceiling up to about 13,700 meters, and range over about 2,780 kilometers.

The aircraft has not yet flown under the NASA banner. Administrator Jared Isaacman expects a first flight within the next two years. The image NASA published is generated, and the illustration on this page is ours and does not show the real design.

X-70

Payload
up to about 907 kg

Weight
about 4,536 kg

Speed
Mach 0.75

Status
not yet flown for NASA

SOURCES

1 **NASA**, NASA Announces X-70 Experimental Aircraft to Advance Flight Research, 4 Nov 2026
nasa.gov

ORBIT / NASA, ESA / NOVEMBER 2026

Two flybys, one photo, one landing

What comes next in space, with a clear line between plan and done.

PLANNED ESCAPADE PATH (CONCEPTUAL)

7 Nov	9 Nov	10 months	Sep 2027
First craft	**Second craft**	**Cruise**	**Mars**
Powered flyby	Powered flyby		7 and 9 September

NEXT: ESCAPADE HEADS FOR MARS

NASA's two ESCAPADE spacecraft will perform a maneuver no agency spacecraft has flown: a powered flyby of Earth. The first around 5:00 EST on Saturday, 7 November, and the second around 6:00 EST on Monday, 9 November. At closest approach each will fire its engine for about 9 minutes, raising its speed from about 6.3 to about 6.7 miles per second.

When this issue comes out the maneuvers have not yet happened. If they go to plan, 10 months of cruise follow, with arrival at Mars on 7 and 9 September 2027. The spacecraft will study how the solar wind affects the planet's atmosphere.

HERA REACHES DIDYMOS

On 18 October, from about 1,280 kilometers, the European probe Hera for the first time separated Didymos and Dimorphos. Dimorphos was then a point about a pixel wide. According to reports, ESA has announced closer photos for 25 November. They have not yet been published.

SPLICE PICKS WHERE TO LAND

NASA has published results of a test of the SPLICE system, carried out on 27 August with an Alta-X drone in California. During descent the system maps terrain with lidar, assesses hazards and redirects the vehicle to a safer spot. The test is from August, and the announcement from November.

SOURCES

1 **NASA**, NASA's Twin ESCAPADE Spacecraft to Slingshot Past Earth on Way to Mars, 2 Nov 2026
science.nasa.gov

2 **Space.com**, Dimorphos in sight: Europe's Hera probe gets its 1st look at battered space rock, Nov 2026
space.com

3 **NASA**, Safe and Precise Landing, Integrated Capabilities Evolution, 2026
nasa.gov

USE THIS / GITHUB ACTIONS / 2 NOV 2026

Check your workflows before 9 November

GitHub is changing the default rule for pull_request_target in public repositories. A guide in six steps.

From 2 December GitHub will block the pull_request_target event by default in all public repositories, unless an applicable policy explicitly allows it. The change was moved from 2 November. Two brownouts come first. If a workflow depends on this trigger, you will see it for the first time on Monday, 9 November.

WHO IS AFFECTED

Public repositories on GitHub Free, Team and Enterprise Cloud. Private and internal repositories are not covered, and neither is GitHub Enterprise Server.

WHAT PULL_REQUEST_TARGET IS

An ordinary pull_request from a fork runs without secrets and with a limited token. The pull_request_target event runs in the context of the base repository and can access secrets. That is useful for labels and comments. It is dangerous when such a workflow checks out and runs code from a fork, because unknown code then gets your secrets.

STEP 1: FIND WHERE IT IS USED

In GitHub code search, limit the path to .github/workflows and search for pull_request_target. Locally, searching the same folder is enough. Check calls to shared workflows too.

STEP 2: LOOK AT INSIGHTS

The rule is now in evaluate mode, in which workflows run while the system records what would be blocked. In the Insights section for Actions policies you can see the runs that would fail. That is the most reliable list, because it shows real use.

DATES

9 Nov	Brownout 16 to 18 UTC
16 Nov	Brownout 16 to 18 UTC
2 Dec	Permanent block

SOURCES

1 **GitHub**, pull_request_target blocked by default December 2, 2 Nov 2026 github.blog

2 **GitHub Docs**, Governing execution of GitHub Actions workflows, 2026 docs.github.com

USE THIS / GITHUB ACTIONS / STEPS 3 TO 6

Do not remove protection for one workflow

The best fix is for the workflow to stop needing a dangerous trigger.

STEP 3: ASSESS EACH WORKFLOW

For each workflow you find, answer three questions. Does it really need secrets or a write token? Does it check out code from the pull request and run it, including installing dependencies and build scripts? Who can trigger it?

If it needs no special access, move to pull_request. That is the cleanest fix. If you only need labels or a comment, the workflow must not run code from a fork. You can split the heavy work: the untrusted part runs in pull_request without secrets, and a separate step processes the result.

STEP 4: A NARROW EXCEPTION, ONLY IF NEEDED

If the trigger is justified, allow it in an applicable Actions event policy and use workflow file targeting, so the exception applies only to those files. Do not switch off protection for a whole organization to save one workflow.

STEP 5: WRITE DOWN THE DECISION

With each exception, note which file it covers, why it needs the trigger, who approved it and when it will be reviewed. When protection is lifted without a reason, the damage does not show until an incident.

STEP 6: USE THE BROWNOUTS

During brownouts, blocked attempts create failed runs with an error message, and the workflow does not run. After the window the system returns to evaluate mode. Tell your team about the brownout in advance and watch pull requests from forks during that period.

RULE

A workflow with secrets must not check out, build or run code from an untrusted pull request.

SOURCES

1 **GitHub Docs**, Securely using pull_request_target, 2026
docs.github.com

2 **GitHub Docs**, Governing execution of GitHub Actions workflows, 2026
docs.github.com

SHIPPED / 31 OCT 2026 → 5 NOV 2026

Shipped this week

What is really available, with platform and limits. The announced purchase of Pierre and other plans are not on this page.

OPENAI · 4 NOV

Images 2.6

ChatGPT, Codex and API, all plans. Six seconds is not a guaranteed time.

AVAILABLE OpenAI

GITHUB · 4 NOV

Copilot in Xcode

From Xcode 27.3 beta, Mac with macOS Tahoe 26.6. Separate from the extension.

PUBLIC PREVIEW GitHub

VERCEL · 3 NOV

Security Review

Pro and Enterprise. Free until 10 November. Not a security audit.

BETA Vercel

VERCEL · 2 NOV

vercel.toml

All projects. Only one configuration file per project.

AVAILABLE Vercel

VERCEL · 3 NOV

vercel doctor

CLI setup check. Diagnoses, does not fix on its own.

AVAILABLE Vercel

GITHUB · 5 NOV

Secret scanning detectors

Starlink keys and fewer false matches. The partner path covers public repositories.

AVAILABLE GitHub

META · 5 NOV

Muse Voice

Mobile, web and Mac, but only the US to start.

ROLLOUT Meta

META · 5 NOV

WhatsApp Events

Arrives over the coming weeks, not to everyone at once.

ROLLOUT Meta

GOOGLE · 2 NOV

Meet, 63 languages

Web, Android and iOS, with an eligible plan. Regional limits apply.

ROLLOUT Google Workspace Updates

VERCEL · 4 NOV

Kimi K3.1 ZDR

More Zero Data Retention providers through AI Gateway.

AVAILABLE Vercel

5 MINUTES / 31 OCT 2026 → 5 NOV 2026

In brief

Short items. If a topic has an article, the card points to it.

GITHUB · 4 NOV

CodeQL 2.27.4

More accurate queries for Java, Rust and C++, plus stronger checks when publishing packs. Page 22.

GitHub

GITHUB · 3 NOV

BYOK usage API

Own-key usage shows up separately. Page 21.

GitHub

GITHUB · 3 NOV

Dependabot APM

Agent skills and instructions get version updates. Page 21.

GitHub

VERCEL · 4 NOV

Workflow SDK

A duplicate start with the same token no longer creates a new run.

Vercel

VERCEL · 3 NOV

Vercel CI

Works without a linked Vercel project.

Vercel

GOOGLE · 5 NOV

Sketch to app

A photo of a drawing becomes an interactive prototype. Page 32.

Google

GOOGLE · OCT 2026

Fitbit Edge

A $179.99 band, without NFC. Page 56.

Android Authority

OPENAI · 2 NOV

LifeVal

A test for everyday tasks, made by the vendor itself. Page 32.

OpenAI

OPENAI · 2 NOV

Bug bounty

Up to $30,000 base, up to $172,000 with bonuses. Page 40.

OpenAI

META · 5 NOV

WhatsApp Events

Invites even for people with no account. Page 56.

Meta

THIS WEEK IN NUMBERS / 31 OCT 2026 → 5 NOV 2026

The week in numbers

948

working local privilege escalation exploits in Linux, in OpenAI research

OPENAI · 4 NOV

81

new high-severity Firefox vulnerabilities

OPENAI · 5 NOV

1,724

reproduced Linux vulnerabilities

OPENAI · 4 NOV

2,308

candidates in the run, which are not confirmed attacks

OPENAI · 4 NOV

48.5%

GPT-6 Cyber on ExploitGym, a vendor claim

OPENAI · 5 NOV

1.2$m

API cost of the described Linux run

OPENAI · 4 NOV

63

languages for speech translation in Meet

GOOGLE WORKSPACE UPDATES · 22 OCT

50%

faster generation, per OpenAI for Images 2.6

OPENAI · 4 NOV

THIS WEEK IN NUMBERS / 31 OCT 2026 → 5 NOV 2026

More numbers of the week

19 MB

FlashProve proof for a 1,024-token request

ANTHROPIC · 2 NOV

207,049

research sessions, which is not a count of scientists

ANTHROPIC · 5 NOV

8.5 $bn

AMD Data Center segment revenue

AMD · Q3

97 %

annual growth of that segment, not the whole company

AMD · Q3

114.5 bn

dollars, Apple quarterly revenue

APPLE · FQ4

1.44 $bn

Arm, quarterly revenue

THE MOTLEY FOOL · FQ2

almost 22 years

length of the Swift mission, planned for two

NASA · 4 NOV

12

Better Auth security advisories

BETTER AUTH · 3 NOV

NEXT / 31 OCT 2026 → 5 NOV 2026

What we are watching

Future events after the 5 November cutoff. None of them had happened when this issue closed.

→ 6 NOV

MacBook Pro

According to Apple's schedule, the new model starts reaching customers.

Apple

→ 7 NOV

ESCAPADE, first flyby

Around 5:00 EST, a powered flyby of Earth.

NASA

→ 9 NOV

ESCAPADE, second flyby

Around 6:00 EST, the second craft.

NASA

→ 9 NOV

GitHub Actions, first brownout

From 16 to 18 UTC in public repositories.

GitHub

→ 13 NOV

Conductor Cloud

Goes read-only.

Conductor

→ 16 NOV

GitHub Actions, second brownout

The same window, from 16 to 18 UTC.

GitHub

→ 18 NOV

Voyager 1

A planned milestone: one light-day from Earth.

NASA

→ 20 NOV

Conductor Cloud

Full shutdown and the start of data deletion.

Conductor

→ 25 NOV

ESA Hera

Announced closer photos of Didymos and Dimorphos.

Space.com

→ 2 DEC

pull_request_target

Permanent default block in public repositories.

GitHub

SOURCES / 31 OCT 2026 → 5 NOV 2026

All sources in this issue

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

06 COVER STORY: CODE HAS BECOME INFRASTRUCTURE

1. **Anthropic**. Redesigning source control at Anthropic. 13 Oct 2026
   anthropic.com/engineering/redesigning-source-control
2. **OpenAI**. Scaling Git at OpenAI. 2 Nov 2026
   openai.com/index/scaling-git-at-openai/
3. **OpenAI**. A thousand fewer paths to root. 4 Nov 2026
   openai.com/index/a-thousand-fewer-paths-to-root/

07 COVER STORY: CONTROL HAS BECOME THE MAIN COST

4. **OpenAI**. GPT-6 Cyber: Advancing the frontier of aligned cybersecurity agents. 5 Nov 2026
   openai.com/index/introducing-gpt-6-cyber/
5. **Anthropic**. Provable inference and FlashProve. 2 Nov 2026
   anthropic.com/research/provable-inference

08 SECURITY: AI HAS ENTERED A NEW PHASE OF CYBERSECURITY

6. **OpenAI**. GPT-6 Cyber System Card. 5 Nov 2026
   deploymentsafety.openai.com/gpt-6-cyber

18 BUILD: OPENAI BUYS THE CODE REVIEW LAYER

7. **OpenAI**. OpenAI to acquire Pierre Computer Company, Inc.. 4 Nov 2026
   openai.com/index/openai-to-acquire-pierre/

19 BUILD: CONDUCTOR JOINS META

8. **Conductor**. Conductor is joining Meta. 4 Nov 2026
   conductor.build/announcement

21 BUILD: COPILOT IN XCODE, BYOK AND AGENT SKILLS

9. **GitHub**. GitHub Copilot is now available in Xcode's coding assistant. 4 Nov 2026
   github.blog/changelog/2026-11-04-github-copilot-is-now-availabl…
10. **GitHub**. BYOK and enterprise custom models in the Copilot usage API. 3 Nov 2026
    github.blog/changelog/2026-11-03-byok-and-enterprise-custom-mod…
11. **GitHub**. Dependabot version updates support Agent Package Manager. 3 Nov 2026
    github.blog/changelog/2026-11-03-dependabot-version-updates-sup…

22 BUILD: WHO GETS THE MESSAGE WHEN A KEY LEAKS

12. **GitHub**. Secret scanning upgrades detectors, adds SpaceX as partner. 5 Nov 2026
    github.blog/changelog/2026-11-05-secret-scanning-upgrades-detec…
13. **GitHub**. CodeQL 2.27.4 improves accuracy and query metadata. 4 Nov 2026
    github.blog/changelog/2026-11-04-codeql-2-27-4-improves-accurac…

23 BUILD: EVERY PUSH GOES THROUGH AN AI REVIEW

14. **Vercel**. Introducing Vercel Security Review. 2 Nov 2026
    vercel.com/blog/vercel-security-review
15. **Vercel**. Vercel Security Review now in public beta. 3 Nov 2026
    vercel.com/changelog/vercel-security-review-public-beta

24 BUILD: BETTER AUTH: 12 ADVISORIES AT ONCE

16. **Better Auth**. Security release on Tuesday, November 3, 2026. 3 Nov 2026
    better-auth.com/blog/security-release-2026-11-03
17. **Vercel**. Better Auth November 2026 security release. Nov 2026
    vercel.com/kb/bulletin/better-auth-november-2026-security-relea…

25 BUILD: SMALL CHANGES AT VERCEL

18. **Vercel**. Changelog. Nov 2026
    vercel.com/changelog
19. **Vercel**. Projects can now be configured with vercel.toml. 2 Nov 2026
    vercel.com/changelog/projects-can-now-be-configured-with-vercel…

26 MODEL WATCH: IMAGES 2.6: PICTURES START TO MOVE

20. **OpenAI**. Introducing Images 2.6. 4 Nov 2026
    openai.com/index/introducing-images-2-6/

27 MODEL WATCH: FLARE OR SUNBURST

21. **OpenAI**. API pricing. Nov 2026
    developers.openai.com/api/docs/pricing

28 MODEL WATCH: GOOGLE MEET BREAKS LANGUAGE BARRIERS

22. **Google Meet Help**. Use Speech Translation. Nov 2026
    support.google.com/meet/answer/16221730
23. **Google Workspace Updates**. Real-time speech translation in Google Meet expands to 63 languages. 22 Oct 2026
    workspaceupdates.googleblog.com/2026/10/real-time-speech-transl…

SOURCES / CONTINUED

30 MODEL WATCH: MUSE GETS A VOICE AND A FACE

24. **Meta**. Give Your Muse a Voice. 5 Nov 2026
about.fb.com/news/2026/11/introducing-muse-ai-voice-expressive-…

32 MODEL WATCH: FROM SKETCH TO APP, FROM TEST TO LIFE

25. **OpenAI**. LifeVal. 2 Nov 2026
openai.com/index/lifeval/
26. **Google**. Turn sketches into apps with Gemini. 5 Nov 2026
blog.google/innovation-and-ai/products/gemini-app/turn-drawings…

33 SECURITY SIGNAL: GHOSTMAP: AN APP WITH NO CODE THAT BREAKS A PHONE

27. **Bitdefender Labs**. GhostMap: The Android flaw that crashes your phone before you open the app. 5 Nov 2026
bitdefender.com/en-us/blog/labs/ghostmap-android

35 PATCH NOW: PATCH NOW

28. **Canadian Centre for Cyber Security**. WordPress security advisory (AV26-1117). 3 Nov 2026
cyber.gc.ca/en/alerts-advisories/wordpress-security-advisory-av…
29. **WordPress.org**. WordPress 7.1.4 release. 28 Oct 2026
wordpress.org/news/category/releases/
30. **Zyxel**. Security advisory for out-of-bounds write vulnerability in GS1900 series switches. 3 Nov 2026
zyxel.com/global/en/support/security-advisories/zyxel-security-…
31. **Binarly**. Out-of-bounds write vulnerability in Zyxel GS1900 switches. Nov 2026
binarly.io/advisories/brly-2026-057

36 SECURITY SIGNAL: WHEN A CYBERATTACK CHANGES THE TEMPERATURE IN THE FRIDGE

32. **BleepingComputer**. Cyberattack disrupts blood storage refrigeration systems in France. Nov 2026
bleepingcomputer.com/news/security/cyberattack-disrupts-blood-s…
33. **CBS News**. Blood banks hit by cyberattack in France days after European freezer hacks. Nov 2026
cbsnews.com/news/cyberattack-france-blood-bank/
34. **NBC News**. Cyberattack on French blood bank fridges was thwarted. Nov 2026
nbcnews.com/world/france/france-cyberattack-blood-bank-refriger…

38 SECURITY SIGNAL: ZYXEL: THREE SOURCES, THREE DIFFERENT THINGS

35. **Help Net Security**. Patched Zyxel GS1900 switches remained open to in-the-wild exploit. 5 Nov 2026
helpnetsecurity.com/2026/11/05/zyxel-cve-2026-107441-exploit/
36. **BleepingComputer**. Zyxel fixes zero-day exploited to hack nearly 1,000 network switches. Nov 2026
bleepingcomputer.com/news/security/zyxel-fixes-zero-day-exploit…

39 SECURITY SIGNAL: WHEN THE TRUSTED SOURCE IS THE PROBLEM

37. **The Register**. Password manager mSecure tells users to change passwords, won't say why. 3 Nov 2026
theregister.com/security/2026/11/03/password-manager-msecure-te…
38. **BleepingComputer**. mSecure password manager hack prompts password reset warnings. Nov 2026
bleepingcomputer.com/news/security/msecure-password-manager-hac…
39. **Kotaku**. Hugely Popular RPG NIKKE Was Compromised. Nov 2026
kotaku.com/nikke-pc-client-compromise-malware-level-infinite-sh…
40. **Automaton**. Goddess of Victory: NIKKE PC launcher reportedly compromised in supply-chain attack. Nov 2026
automaton-media.com/en/news/goddess-of-victory-nikke-pc-launche…

40 SECURITY SIGNAL: PAYING PEOPLE TO FIND WHAT MODELS MISS

41. **OpenAI**. Rewarding defenders with bigger bug bounties. 2 Nov 2026
openai.com/index/rewarding-defenders-with-bigger-bug-bounties/
42. **DreamHost**. October Security Incident. 29 Oct 2026
dreamhost.com/news/announcements/october-security-incident/

41 SECURITY SIGNAL: SHORT SECURITY ROUNDUP

43. **BleepingComputer**. Google Calendar phishing campaign hijacks sessions, bypasses MFA. Nov 2026
bleepingcomputer.com/news/security/google-calendar-phishing-cam…
44. **GitHub**. GitHub Desktop adds short-lived authentication tokens. 2 Nov 2026
github.blog/changelog/2026-11-02-github-desktop-adds-short-live…

44 LAB: CLAUDE AND THE MYSTERY OF HOLLOW PARTICLES

45. **Anthropic**. Investigating vault particles with Claude. 4 Nov 2026
anthropic.com/research/investigating-vault-particles-with-claude

45 LAB: WHAT SCIENTISTS REALLY USE CLAUDE FOR

46. **Anthropic**. How researchers use Claude. 5 Nov 2026
anthropic.com/research/how-researchers-use-claude

47 LAB: AN HOUR OF CHEMISTRY, NOT AN HOUR TO A DRUG

47. **Google DeepMind**. Agentic cheminformatics with Google Antigravity. 3 Nov 2026
deepmind.google/blog/agentic-cheminformatics-with-google-antigr…

SOURCES / CONTINUED

48 LAB: FROM CAD MODEL TO REAL PART

48. **OpenAI**. Accelerating engineering for the physical world. 3 Nov 2026 openai.com/index/accelerating-engineering-for-the-physical-worl…

49 MONEY: AMD: DATA CENTER ALMOST DOUBLES

49. **AMD**. AMD Reports Third Quarter 2026 Financial Results. 3 Nov 2026 newsroom.amd.com/news/amd-reports-third-quarter-2026-earnings/

51 MONEY: APPLE: A RECORD WITH NO AI STORY IN THE NUMBERS

50. **Apple**. Apple reports fourth quarter results. 2 Nov 2026 apple.com/newsroom/2026/11/apple-reports-fourth-quarter-results/
51. **Apple**. Apple unveils groundbreaking MacBook Pro with PROLED XDR display. 28 Oct 2026 apple.com/newsroom/2026/10/apple-unveils-groundbreaking-macbook…

52 MONEY: ARM SELLS BLUEPRINTS, AND NOW MORE THAN THAT

52. **The Motley Fool**. Arm Holdings Q2 2027 earnings call transcript. 5 Nov 2026 fool.com/earnings/call-transcripts/2026/11/05/arm-holdings-arm-…
53. **BigGo Finance**. Arm Holdings Q2 2027 earnings call summary. 4 Nov 2026 finance.biggo.com/news/US_ARM_2026-11-04

53 INFRA: WHAT A 401-MEGAWATT GUARANTEE LOOKS LIKE

54. **Reuters via KFGO**. TeraWulf gets credit backstop from Google for Anthropic data-center project. 4 Nov 2026 kfgo.com/2026/11/04/terawulf-gets-credit-backstop-from-google-f…
55. **CoinDesk**. Google could take ownership stake in TeraWulf's AI project if Anthropic defaults. 5 Nov 2026 coindesk.com/tech/2026/11/05/google-could-take-ownership-stake-…

54 INFRA: A COMPUTE CAMPUS AS A GRID RESERVE

56. **OpenAI**. Grid Protectors Program. 2 Nov 2026 openai.com/index/grid-protectors-program/

55 MONEY: WHO IS MAKING MONEY IN THIS PHASE

57. **Figma**. Figma Announces Third Quarter 2026 Financial Results. 4 Nov 2026 investor.figma.com/news-events/news/news-details/2026/Figma-Ann…

56 WORK: HOW TO TURN ON TRANSLATION IN GOOGLE MEET

58. **Meta**. WhatsApp Events brings people together beyond the group chat. 5 Nov 2026 about.fb.com/news/2026/11/whatsapp-events-brings-people-togethe…
59. **Android Authority**. Google Fitbit Edge launches with 7-day battery for $180. Oct 2026 androidauthority.com/google-fitbit-edge-price-specs-features-la…
60. **Google**. Responsible AI for schools with Guided Learning. 3 Nov 2026 blog.google/products-and-platforms/products/education/ai-learni…

57 ORBIT: END OF THE SWIFT MISSION

61. **NASA**. NASA's Swift Observatory Re-Entered Atmosphere. 4 Nov 2026 science.nasa.gov/blogs/swift/2026/11/02/nasas-swift-observatory…
62. **NASA**. NASA's Swift Mission Ends Ahead of Re-Entry. 23 Oct 2026 science.nasa.gov/blogs/swift/2026/10/23/nasas-swift-mission-end…

59 ORBIT: X-70 FRACTAL: AN AIRCRAFT THAT SWAPS ITS EQUIPMENT

63. **NASA**. NASA Announces X-70 Experimental Aircraft to Advance Flight Research. 4 Nov 2026 nasa.gov/news-release/nasa-announces-x-70-experimental-aircraft…

60 ORBIT: TWO FLYBYS, ONE PHOTO, ONE LANDING

64. **NASA**. NASA's Twin ESCAPADE Spacecraft to Slingshot Past Earth on Way to Mars. 2 Nov 2026 science.nasa.gov/blogs/escapade/2026/11/02/nasas-twin-escapade-…
65. **Space.com**. Dimorphos in sight: Europe's Hera probe gets its 1st look at battered space rock. Nov 2026 space.com/space-exploration/asteroid-comet-missions/dimorphos-i…
66. **NASA**. Safe and Precise Landing, Integrated Capabilities Evolution. 2026 nasa.gov/safe-and-precise-landing-integrated-capabilities-evolu…

61 USE THIS: CHECK YOUR WORKFLOWS BEFORE 9 NOVEMBER

67. **GitHub**. pull_request_target blocked by default December 2. 2 Nov 2026 github.blog/changelog/2026-11-02-pull_request_target-blocked-by…
68. **GitHub Docs**. Governing execution of GitHub Actions workflows. 2026 docs.github.com/en/repositories/managing-your-repositorys-setti…

62 USE THIS: DO NOT REMOVE PROTECTION FOR ONE WORKFLOW

69. **GitHub Docs**. Securely using pull_request_target. 2026 docs.github.com/en/actions/reference/security/securely-using-pu…

67 NEXT: WHAT WE ARE WATCHING

70. **NASA**. NASA's Voyager 1 conserves power, resolves brief communications pause. 29 Oct 2026 science.nasa.gov/blogs/voyager/2026/10/29/nasas-voyager-1-conse…

o0o0o0o

/ 006

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

Issue 006
Date 6 Nov 2026
Period 31 Oct 2026, 00:00 → 5 Nov 2026, 22:15
Web tutorijali.app

tutorijali.app