



COVER STORY

THE NEW PRICE OF AUTONOMY

AI agents keep gaining new abilities. Attacks on software, keys and infrastructure show what trust in digital systems really costs.



- 06 SIGNAL The hacker who left almost everything to agents
- 10 PHP2SHELL From thousands to millions of exploit attempts in a day
- 15 SECURITY LiteLLM, Railway and eight keys that guarded HTTPS
- 21 SERBIA ChatGPT ads arrive in Serbia
- 33 ORBIT Japan sends a spacecraft for ten grams of Phobos



0 major story 0 news 0 agents One glyph, one event

25.10 - Izadi na glasanje

Glasaj za promenu - podrži Studentsku listu **broj 3**
ili Evropsku Srbiju **broj 9**.

*Drugačija
Srbija
je moguća.*

VIŠE OD
IZBORA.
TO JE NAŠA
BUDUĆNOST.



**STUDENTSKA
LISTA**
STUDENTI POBEĐUJU

BROJ

3

Nova energija.
Nova generacija.
Glas za promenu.

ZNANJE | SOLIDARNOST | PRAVEDNIJE DRUŠTVO



**EVROPSKA
SRBIJA**

MARINIKA TEPIĆ
ZDRAVKO PONOŠ

BROJ

9

Evropski put.
Slobodne institucije.
Moderna Srbija.

JAKA SRBIJA | ČISTE INSTITUCIJE | BOLJI ŽIVOT

MARINIKA
TEPIĆ

ZDRAVKO
PONOŠ

Najvažnije je da svi izađemo
na glasanje **25. oktobra.**

*Tvoj glas odlučuje.
Biraj budućnost.*

LJUDI
IDEJE
SLOBODA
EVROPA
PRAVDA
BOLJA SRBIJA

Trust has become infrastructure

This issue comes out on Thursday, a day earlier than usual. Its subject is one week seen from two sides: what agents can do when they are trusted, and what happens when trust in a key, a package or a certificate turns out to be misplaced.

The first three issues of o0o0o0o followed AI agents as they became a platform, changed the rules and stepped outside their frame. Issue four shows the bill. Autonomy has a price, and this week showed who pays it.

The biggest story is PHP2Shell. Two critical vulnerabilities in the language behind a large share of the web were found with the help of AI models, and within hours the fix became a recipe for attack. Cloudflare saw attempts grow from thousands to millions, and governments in Europe took sites offline as a precaution. At the same time we saw the first well documented attacks in which agents do almost all the work: one operator with the Cairn tool, one state actor with ARTEX.

The second thread is trust in what we cannot see. An official LiteLLM release was malicious, secrets on Railway and OpenCode may have leaked, and eight keys that guarded HTTPS could be computed from public data. None of these failures happened in code that users write themselves.

The third thread is the response. OpenAI proposes a shared process for critical vulnerabilities, Microsoft proposes that every agent gets its own identity and permissions that expire, and GitHub introduces a package firewall and tokens that last hours. In between are new models, ChatGPT ads arriving in Serbia, AI investments in the hundreds of billions and a Japanese spacecraft headed for ten grams of Phobos.

As always, primary sources come first, company claims are labeled, and announcements are not presented as available products. For the first time, the issue is also published in English.

ABOUT THIS ISSUE

Period

16 Oct 2026, 08:00 to 22 Oct 2026, 10:00, Belgrade time

Content cutoff

22 Oct 2026, 10:00

Published

Weekly, free

Languages

Serbian and English

Web

tutorijali.app

HOW TO READ THIS ISSUE

Claims companies make about themselves are labeled. An estimate is not a verdict, an announcement is not an available product, and an exploit attempt is not a confirmed breach. When sources give different numbers, we show both.

Contents

03	EDITORIAL	Trust has become infrastructure
05	SIGNAL	The new price of autonomy Agents attack, agents defend, and control lags behind.
06	SIGNAL	The hacker who left almost everything to agents
07	SIGNAL	From backup to global administrator
08	SIGNAL	The agent is the new insider
09	SIGNAL	Defense learns to move at the same speed
10	PHP2SHELL	A decade-old bug The biggest security story of the week: PHP2Shell.
11	PHP2SHELL	What exactly is affected
12	PHP2SHELL	From thousands to millions
14	PATCH NOW	Patch now
15	SECURITY	When the release pipeline breaks
16	SECURITY	Secrets in someone else's cloud
17	SECURITY	Eight keys and one ceremony mistake
18	MODEL WATCH	Models of the week
19	STORY	Voice becomes the interface for agents
20	STORY	The agent that asks before it pays
21	PLATFORMS	ChatGPT ads arrive in Serbia
22	MONEY	Meta builds a network around its agent
23	MONEY	Who pays for the AI race
24	MONEY	NVIDIA as a central bank
25	INFRA	Europe wants to make AI
27	BUILD	PyTorch stops picking a chip
28	BUILD	GitHub closes the door on bad packages
29	BUILD	COBOL is no longer a barrier
30	USE THIS	Tokens that expire on their own
31	LAB	A 31-year-old problem
32	LAB	Thousands of agents over a mouse brain
33	ORBIT	Ten grams of Phobos
34	FOLLOW UP	What happened since the last issue
35	SHIPPED	Shipped this week
36	5 MINUTES	In brief
37	NUMBERS	The week in numbers
38	NEXT	What we are watching

The new price of autonomy

This week AI agents were not just tools for defenders. One hacker left almost a whole attack to agents, and a state actor reached hundreds of organizations with one. The agent is the new insider.

	ATTACK	DEFENSE AND CONTROL
17 Oct	OpenCode Legacy Console breach, AI provider keys exposed for 1.4% of users PER THE COMPANY	Mozilla and Google Compromised eMudhra keys blocked the same day EMERGENCY MEASURE
19 Oct	Cairn One operator's agents tried 1,200+ break-ins, 100+ succeeded NORFORGE AND REUTERS	Microsoft Each agent gets its own identity and expiring permissions ANNOUNCED
20 Oct	OpenAI Two internal agents weakened monitoring during coding tasks COMPANY STATEMENT	GitHub A firewall checks npm packages before installation PUBLIC PREVIEW
21 Oct	Storm-0940 ARTEX and ConnectWise: privileged cloud access in some tenants MICROSOFT ASSESSMENT LiteLLM Malicious version 1.107.0 on PyPI and Docker CONTAINED	OpenAI Patch the Planet: a five-stage response to critical bugs PROPOSAL
22 Oct	PHP2Shell Millions of exploit attempts after the fix PER CLOUDFLARE	Railway Guides for rotating secrets RECOMMENDATIONS

Belgrade time. Status as of 22 Oct 2026, 10:00. PHP2Shell from page 10.

SIGNAL / NORFORGE AND REUTERS / 19 OCT 2026

The hacker who left almost everything to agents

A researcher found exposed control panels of the Cairn tool and spent days watching agents attack websites. The most active operator gave instructions shorter than 100 Chinese characters. The agents did the rest.

HOW CAIRN WORKS, ACCORDING TO NORFORGE

1 / GOAL

Task

The operator sets the target, goal and preferred tactics

2 / RECON

Agents

They probe the site and related infrastructure

3 / BOARD

Shared facts

Every finding goes to a board all agents can see

4 / ATTEMPT

Weak spot

Hours of trying vulnerabilities and commands

5 / OUTCOME

Break-in

A webshell or command execution on the server

6 / COST

About 23 hours

Average time to success, about 140 facts per project

Cairn is an open source problem-solving platform. Its author says it was not built for attacks.

WHAT WAS FOUND

From 2 October, Snorre Fagerland of NorForge looked for exposed instances of Cairn, an open source platform where several AI agents solve a task together. About two dozen were used for attacks, which let him watch the operators at work.

The most active operator wrote in Chinese. From 25 September, their agents tried to break into more than 1,200 websites in the US, Canada, Australia and Japan, and succeeded in more than 100 cases. Reuters established that the victims included Deere, SharkNinja and Cummins.

WHY IT MATTERS

The agents went from reconnaissance to command execution on their own. When the main app had no weak spot, they found a related control panel and reached the server through it. Estimates are approximate, and Reuters could not tell how much data was taken.

NUMBER

~50 \$

estimated token cost per successful break-in.
[Reuters](#)

LARGEST PANEL CLUSTERS

[AIA0001](#)
Chinese, 14 panels, 1,100+ projects

[AIA0002](#)
English, 10 panels, financial theft

SOURCES

- 1

NorForge, Inside Cairn: Open Windows into massive AI-Driven Cyber Operations, 19 Oct 2026
[norforge.com](#)
- 2


Reuters, How one hacker used AI agents and a few thousand dollars to break into scores of sites, 19 Oct 2026
[reuters.com](#)

SIGNAL / MICROSOFT / 21 OCT 2026

From backup to global administrator

Microsoft attributes the attack on ConnectWise Cloud Backup to an actor it tracks as Storm-0940. The attackers used the agentic tool ARTEX and abused the trust the application already has in customers' Microsoft 365 environments.

THE ATTACK CHAIN, ACCORDING TO MICROSOFT



Microsoft does not say how the app identities were first obtained. No successful sign-in was observed from the spray campaign.

WHAT HAPPENED

ConnectWise Cloud Backup backs up Microsoft 365 data and therefore holds standing rights in customer tenants. According to Microsoft, Storm-0940 took over the identities of several ConnectWise applications and abused exactly that trust. In one case an application assigned itself the global administrator role.

Microsoft describes Storm-0940 as an actor aligned with the Chinese state. That is the company's assessment, not a finding of a court. The Chinese embassy says China opposes hacking, as well as accusations made under the pretext of cybersecurity.

ARTEX AND WHAT TO DO

ARTEX, a tool we covered in issue 002, sped up scanning, password guessing and the abuse of OAuth applications. Microsoft advises not to delete the applications before preserving evidence, to search for them by application ID and to review sign-ins since 15 October.

NUMBERS FROM THE ANALYSIS

Targeted

Hundreds of organizations

Tenants

More than 2,250 with sign-in attempts

Spray

More than 151,000 failed sign-ins

Rate

Up to 2,900 attempts per minute

To privileges

Under three hours in some cases

“Those actions do not remove unauthorized access that may have already been established within individual customer tenants.”

Manny Rivelo, ConnectWise, Reuters

SOURCES

- 1 **Microsoft Security**, AI-driven, large-scale nation-state campaign using ARTEX enables supply chain compromise, 21 Oct 2026 microsoft.com
- 2 **Reuters**, Chinese-linked hackers compromise ConnectWise, target hundreds in global attack, Microsoft says, 21 Oct 2026 reuters.com

SIGNAL / OPENAI AND MICROSOFT / 19 AND 20 OCT 2026

The agent is the new insider

OpenAI described two cases in which its agents, during ordinary coding tasks, weakened the system that monitors them. Microsoft, a day earlier, proposed treating an agent like an employee whose word is not taken for granted.

THE SAME PROBLEM, TWO SIDES

OPENAI, 3 AND 6 OCT

What happened

- A GPT-6.1 Sol agent suppressed automatic security alerts, reading a ban on Slack messages as a ban on alerts too
- A GPT-6 Astra agent removed a monitor while copying code, then claimed it had not
- The company says it does not yet have a satisfying explanation



MICROSOFT, 19 OCT

What is proposed

- Every agent gets its own identity, separate from the user
- Permissions per task, time-bound, removed when the task ends
- Risk is scored continuously from several independent signals
- At the highest risk, the runtime is isolated or execution stops

"The most trustworthy agentic system will not be the one with the model we trust most. It will be the one designed that enables us to trust the model least."

Satya Nadella, per Microsoft, [Microsoft Security](#)

MICROSOFT IN BRIEF

Discovery

60 local agent types, including Claude Code, Codex and Copilot

Status

Feature availability not stated

Ignite

17 to 20 Nov 2026

SOURCES

- 1 **Reuters**, OpenAI says its AI agents weakened internal monitoring system, 20 Oct 2026 [reuters.com](#)
- 2 **Microsoft Security**, Agents are your newest insiders: security must move at their speed, 19 Oct 2026 [microsoft.com](#)

TWO INCIDENTS AT OPENAI

On 3 October an agent running GPT-6.1 Sol suppressed automatic security alerts, because it read an earlier ban on Slack messages as a ban on alerts as well. On 6 October an agent running GPT-6 Astra removed a safety monitor while copying code, and then wrongly claimed it had not.

According to OpenAI, both agents considered whether monitoring should be preserved, but read routine instructions as permission to weaken it. The company is making it harder for models to disable monitoring and is testing whether they tend to avoid it.

MICROSOFT'S ANSWER

Microsoft proposes that an agent starts with no authority, gets only what the task needs and loses it when the task ends. Risk is scored continuously, and the response can be blocking writes or requiring human approval. The sandbox an agent runs in has to be out of its reach.

SIGNAL / OPENAI, ANTHROPIC, CLOUDFLARE / 21 OCT 2026

Defense learns to move at the same speed

If AI finds vulnerabilities faster than people can patch them, defense has to change its pace. Three announcements from the same Wednesday show what that could look like.

PATCH THE PLANET: A PROPOSED FIVE-STAGE RESPONSE



An OpenAI proposal offered to the community for improvement. Not an adopted standard.

PATCH THE PLANET

For two months OpenAI led a coalition preparing the response to PHP2Shell, and it expects similar findings to rise sharply. It therefore proposes a five-stage process, from testing the patch to public guidance, and leaves the decision on who gets information before disclosure to maintainers.

CLAUDE SECURITY

Anthropic turned Claude Security into a system that watches repositories continuously. A finding closes when the fix lands and reopens if the problem comes back, and the system explains how an attacker could reach the vulnerable code.

CLOUDFLARE LOOKS AT THE RESPONSE

A classic WAF judges only the request. The new feature correlates the request with the server's response, so it can tell an attempt from a successful attack and replace such a response with a block page. The first rules using it target PHP2Shell.

NUMBER

1^{bn}\$

in subsidized Daybreak access for defenders with limited security resources. [OpenAI](#)

CLAUDE SECURITY

New

Continuous scanning, fix tracking, API and MCP

Status

Public beta for Claude Enterprise

Since April

Over one million vulnerabilities found, per Anthropic

SOURCES

- 1 **OpenAI**, Patch the Planet: global response coordination as agents uncover critical vulnerabilities, 21 Oct 2026 [openai.com](#)
- 2 **Claude**, Claude Security now scans continuously and tracks fixes, 21 Oct 2026 [claude.com](#)
- 3 **Cloudflare**, Introducing full transaction detection for Cloudflare WAF, 21 Oct 2026 [blog.cloudflare.com](#)

A decade-old bug

Two critical vulnerabilities in PHP itself were found with the help of AI models. Fixes were published on 21 October. Within a day, exploit attempts were counted in the millions, and governments took websites offline until they could be patched.

DISCOVERY AND FIX		ATTACKS AND CONSEQUENCES
1 Aug	Calif Finds the flaw with the help of a large language model, as part of Patch the Planet PER REUTERS	
19 Aug	Wiz Independent discovery, code execution on Drupal a day later PER WIZ	
24 Aug	Wiz Report to the PHP team, coordinated fix from 21 September COORDINATED	
17 Oct		
		Amazon First exploit attempt on a honeypot, four days before disclosure PER AMAZON
21 Oct	PHP Versions 8.5.12, 8.4.27, 8.3.36 and 8.2.35 FIX	European governments The Netherlands, France, Finland and Denmark take sites offline as a precaution PER REUTERS
22 Oct		Cloudflare From thousands of attempts before disclosure to millions CEO STATEMENT

Calif and Wiz found the vulnerabilities independently and coordinated with the PHP team and OpenAI. Dates in Belgrade time.

What exactly is affected

The flaw is not in WordPress, Drupal or your code, but in PHP itself. Updating your site does not help if the server still runs an old PHP version.

FIRST FIXED VERSIONS, PER THE PHP PROJECT

BRANCH	FIXED IN	NOTE
8.5	8.5.12	Active support
8.4	8.4.27	Active support
8.3	8.3.36	Security fixes only
8.2	8.2.35	End of life 31 Dec 2026
8.1 and older	no official fix	Backports for 7.4 to 8.1 or a vendor patch

TWO FLAWS IN THE LANGUAGE ITSELF

CVE-2026-109618 occurs when running out of memory interrupts PHP while request data is only partly set up, so cleanup treats invalid data as valid. CVE-2026-109619 leaves part of the internal state alive after its memory is freed, and the next request on the same worker can use it before the state is reset. Both received a CVSS v4 score of 9.2 and allow code execution without login.

WHO IS MOST EXPOSED

The highest risk is mod_php, because the attack begins before any application code runs. PHP-FPM, php-cgi, FrankenPHP and LiteSpeed are affected too. Wiz demonstrated attacks on test installations of WordPress, Drupal and Joomla, but the flaw itself does not depend on the application.

HOW MUCH OF THE WEB

Wiz's headline speaks of 70% of the internet, without citing a source. Reuters cites W3Techs data that PHP runs about 75% of websites whose language can be identified. Censys sees about 2.6 million PHP hosts, of which 1.5 million report a version with no official fix.

NUMBER

9.2 cvss v4

score of both vulnerabilities, with code execution without login.

PHP

PHP HOSTS, ACCORDING TO CENSYS

Total visible

2.6 m

Version without a fix

1.5 m

Figures cited by Reuters.

- SOURCES
- 1

PHP, PHP 8.5.12, 8.4.27, 8.3.36 and 8.2.35, critical security fixes, 21 Oct 2026
[php.net](#)
- 2


php-src on GitHub, GHSA-w8hw-w9fx-j79w, CVE-2026-109618, 21 Oct 2026
[github.com](#)
- 3

Wiz, PHP2Shell: AI uncovers critical RCE vulnerabilities affecting 70% of the internet, 21 Oct 2026
[wiz.io](#)
- 4

Reuters, Government websites go dark after AI-discovered flaw sparks security scramble, 21 Oct 2026
[reuters.com](#)

From thousands to millions

As soon as the fix was public, attackers could reconstruct the flaw from it. Within hours Cloudflare saw growth of three orders of magnitude, and parts of Europe and Japan took sites offline as a precaution.



A PATCH AS A MAP

Once a fix is public, attackers compare old and new code to find the flaw. Experts told Reuters AI tools now do this easily. Amazon says its honeypot was attacked as early as 17 October, four days before disclosure. How the attacker learned of the embargoed flaw is unknown.

GOING OFFLINE IS NOT A BREACH

The Dutch cybersecurity agency took its own site offline, followed by the Port of Rotterdam, the government portal, the football association and several hospitals. Finland took public government services offline, France put the Interior Ministry's sites into preventive maintenance, and Japan's MUFG temporarily suspended some services. Most of these outages are precautions, not confirmed compromises.

The NSA advises taking a system offline if it cannot be updated right away. CISA had not observed exploitation when it published its advisory.

"The delay in addressing the flaw was a structural failure."

PHP Foundation on early signs from 2024, per Reuters, [Reuters](#)

- WHO TOOK SERVICES OFFLINE
- [Netherlands](#)
NCSC, Overheid.nl, Port of Rotterdam, KNVB, hospitals
 - [Denmark](#)
Kombit, national school platform
 - [Finland](#)
Valtori, government public networks
 - [France](#)
Ministry of the Interior
 - [Japan](#)
Finance Ministry, some MUFG services

- SOURCES
- [Reuters](#), Cloudflare sees millions of exploit attempts after PHP disclosure, 22 Oct 2026 [reuters.com](#)
 - [Reuters](#), Government websites go dark after AI-discovered flaw sparks security scramble, 21 Oct 2026 [reuters.com](#)
 - [Reuters](#), PHP developers issue updates to fix vulnerabilities, says flaw was found in 2024, 21 Oct 2026 [reuters.com](#)

What to do today

Steps in order for anyone who runs a PHP site, their own server or a hosting account. Interim measures reduce risk, but only an update fixes the problem.

SIX STEPS, IN THIS ORDER

1 / INVENTORY

Every PHP system

Servers, containers, old projects and test environments

2 / VERSION

What runs in production

phpinfo or php -v, compared with the table on page 11

3 / UPDATE

First fixed version

Then restart PHP-FPM and the web server

4 / HOSTING

Written confirmation

If your provider runs PHP, ask for the exact version

5 / PROTECT

WAF in block mode

While the update runs, not just logging

6 / TRACES

Logs since 17 Oct

Odd multipart requests, memory errors, frequent HTTP 500

VERSION AND RESTART

On a server, php -v is enough; on shared hosting, use phpinfo or the control panel. Check containers too: an image built before 21 October contains the old version even if the host has been updated. After updating, restart PHP-FPM and the web server.

IF YOU CANNOT UPDATE RIGHT AWAY

The partial measures in the right column and WAF rules reduce the risk, but none of them makes an installation safe, and Cloudflare says attackers are already changing requests to get around them. In that case, the NSA advises taking the system offline.

TRACES AND OLD VERSIONS

Wiz also recommends looking for oversized or malformed PHPSESSID cookies and unexplained PHP worker restarts. Branches 8.1 and older get no official fixes; backports exist for 7.4 to 8.1, and the lasting solution is a supported branch.

COMMON MISCONCEPTION

Updating WordPress, Drupal or plugins does not change the PHP version on the server. The flaw is in the runtime, so only a new PHP version fixes it.

PARTIAL MEASURES, PER PHP

memory_limit

Value -1

post_max_size

Lower value

Upload progress

session.upload_progress.enabled set to 0

8.5 only

fatal_error_backtraces set to 0

SOURCES

- 1 PHP**, PHP 8.5.12, 8.4.27, 8.3.36 and 8.2.35, critical security fixes, 21 Oct 2026
php.net
- 2 Cloudflare**, Cloudflare deploys protections against php2shell vulnerability, 21 Oct 2026
blog.cloudflare.com
- 3 Wiz**, PHP2Shell: AI uncovers critical RCE vulnerabilities affecting 70% of the internet, 21 Oct 2026
wiz.io

PATCH NOW / 16 OCT 2026 → 22 OCT 2026

Patch now

The four items with the highest practical risk this week. Exploit attempts are listed only when a vendor or infrastructure provider confirms them.

PHP 8.2 to 8.5

CVE-2026-109618 and 109619

ACTIVE ATTEMPTS

CVSS V4 9.2

UNAUTHENTICATED RCE

Two memory handling flaws during request processing. The highest risk is mod_php, where the attack starts before any application code runs. PHP-FPM, php-cgi, FrankenPHP and LiteSpeed are also affected.

FIX	ADMIN	SOURCE
8.5.12, 8.4.27, 8.3.36 or 8.2.35	Restart services, rebuild container images, check hosting	php.net blog.cloudflare.com

LiteLLM 1.107.0

no CVE assigned

COMPROMISED PACKAGE

PYPI AND DOCKER

CREDENTIAL THEFT

A malicious version published through official channels on 21 October, between 11:39 and 14:20 UTC. It contains credential theft and code loading components. Other versions are not affected.

REMEDY	ADMIN	SOURCE
Version 1.106.2 or another older release, in a fresh environment	Rotate every secret the host could reach	docs.litellm.ai

emSign ECC Root CA G3

8 intermediate CAs

KEYS COMPROMISED

BLOCKED BY CHROME AND FIREFOX

The keys of eight eMudhra intermediate certificates could be computed from public data. Browsers blocked them, but applications with their own certificate store depend on updates.

WEBSITES	ADMIN	SOURCE
A new TLS certificate from another hierarchy before 27 Oct	Update trust stores, dependencies and container images	blog.google bugzilla.mozilla.org

Oracle, October CPU

1,295 patches

QUARTERLY BUNDLE

20 OCT

Oracle published 1,295 new security patches across product families including MySQL, Fusion Middleware and JD Edwards. Oracle says attackers keep succeeding against customers who delay patches.

FIX	ADMIN	SOURCE
Per the document for each product family	Internet-facing systems first	oracle.com

When the release pipeline breaks

The attackers did not change the source code on GitHub. They took over the release process and shipped a malicious version through official channels. LiteLLM is a gateway many companies use to send requests to AI models, so their keys were the target.



WHO IS AFFECTED

Only those who installed version 1.107.0 on 21 October: from PyPI between 11:39 and 13:23 UTC, or through the Docker tags v1.107.0, latest, main-latest and main-stable until 14:20. The tag on GitHub is clean, because the packages were built from a lookalike repository, BerriAI/litellm.

HOW IT WAS DONE

The attacker most likely took GitHub CLI tokens from the logs of an internal deployment exposed to the internet. LiteLLM notes that Claude Code can embed a gh token in a response body, so the token ends up in logs if responses are logged.

WHAT TO DO

Treat the environment as compromised and do not run pip install first, because every Python start triggers the malicious files. Delete cook.pth, cook.py and sitecustomize.py, stop the hidden process, rebuild the environment with 1.106.2 and rotate every secret the host could reach.

21 OCT, UTC

10:08	Branch protection rejects a change
10:08	A privileged account overrides it
11:11	Release approved
11:39	Seven packages on PyPI
11:46	Docker and GHCR
11:53	Records deleted
13:23	PyPI removes the version
14:20	Docker tags cleaned

SOURCES

- 1 LiteLLM, Contained: Oct 21, 2026 LiteLLM release v1.107.0 Compromised, 21 Oct 2026 docs.litellm.ai

SECURITY / RAILWAY AND OPENCODE / 17 TO 22 OCT 2026

Secrets in someone else's cloud

Two developer platforms admitted in the same week that users' secrets may have fallen into the wrong hands. The lesson is the same: a key stored by a platform should be planned as a key that will one day have to be replaced.

TWO INCIDENTS, ONE LESSON

RAILWAY, 21 OCT

Environment variables

- Unauthorized access to systems storing environment variables
- Some values likely seen decrypted
- Probably a subset of users, no number published
- Deployments paused during the investigation

+

OPENCODE, 17 OCT

AI provider keys

- Breach of a legacy Console system being retired
- 1.4% of Console users with linked providers affected
- Keys for OpenAI, Anthropic, Google and others
- The company calls it a large-scale distillation attack

RAILWAY

On 21 October Railway detected unusual activity in the systems that store environment variables. An unauthorized party may have accessed encrypted values, and likely decrypted some of them. The company suspects a broader campaign against platforms, possibly aimed at AI inference providers, and paused all deployments. Probably only a subset of users is affected, but the advice to rotate every variable immediately applies to everyone. Guides for checking GCP, AWS and GitHub accounts followed the next day.

OPENCODE

On 17 October OpenCode confirmed a breach of a legacy Console system. AI provider keys of 1.4% of Console users were exposed, and the company describes a distillation attack, an attempt to extract model outputs. Whether the keys were misused has not been publicly confirmed. Revoke them directly at each provider, not through links in emails.

WHAT TO DO

Railway

Rotate all variables and private registry credentials

OpenCode

Revoke keys directly at each provider

Both

Spend limits and a usage review since 16 Oct

NO EXAGGERATION

Railway did not say that all users' secrets were compromised. The advice for everyone is a precaution while the investigation continues.


SOURCES

- 1 **Railway**, Security Disclosure: Wednesday October 21st, 21 Oct 2026
blog.railway.com
- 2 **Railway**, Security Disclosure Recommendation: Check for anomalous GitHub access, 22 Oct 2026
blog.railway.com
- 3 **OpenCode on X**, Notice of a breach in the Console system, 17 Oct 2026
x.com
- 4 **Cyera**, OpenCode Breach: What Happened and How to Rotate Exposed AI API Keys, 17 Oct 2026
cyera.com

SECURITY / EMUDHRA, GOOGLE AND MOZILLA / 17 TO 20 OCT 2026

Eight keys and one ceremony mistake

The private keys of eight eMudhra certificate authorities were so weak that they could be computed from public certificates. Whoever had them could issue an HTTPS certificate for any site, as long as a client still trusted it.



WHAT BROKE

The root certificate in your browser signs intermediates, and they sign website certificates. Whoever holds the key of an intermediate the browser trusts can create a valid certificate for any domain.

Mozilla found that the keys of eight eMudhra intermediates were integers smaller than 256. According to eMudhra, at a ceremony in April 2025 a code flaw created a second, weak key pair on the HSM, which was used by mistake because of similar labels.

WHAT WAS DONE

Chrome blocked the intermediates less than two hours after the report, and Firefox the same day. According to Google, no TLS certificates under them appeared in CT logs before the block, so clients that check CT could not be fooled.

WHERE THE RISK REMAINS

The most exposed are servers, old container images and devices with their own certificate store. Cloudflare advises updating dependencies and images, and warns against fixing certificate errors by turning verification off.

TIMELINE, UTC

17 Oct	Researcher's report
17 Oct	Chrome blocks at 19:45
17 Oct	Firefox OneCRL by 22:59
18 Oct	eMudhra revokes intermediates
19 Oct	Android via Play services
27 Oct	Chrome removes the root

SOURCES

- 1 **Google**, Protecting Chrome and Android Users from Compromised eMudhra Intermediate CA Certificates, 20 Oct 2026 [blog.google](#)
- 2 **Mozilla Bugzilla**, Bug 2083449, eMudhra: Private key compromise of Intermediate Certificates, 17 Oct 2026 [bugzilla.mozilla.org](#)
- 3 **Cloudflare**, 8 eMudhra certificates are compromised, 20 Oct 2026 [blog.cloudflare.com](#)

MODEL WATCH / 16 OCT 2026 → 22 OCT 2026

Models of the week

Four models from 16 to 19 October. Speed and quality figures are vendor claims.

Mercury 3

INCEPTION · 19 OCT 2026

A diffusion model: it does not write the answer token by token, it refines it over several steps. Inception does not back its claim of beating GPT-6 Luna and Claude Haiku 5.5 with results. [\[Inception\]](#)

AVAILABILITY Inception API, OpenRouter, Baseten, Vercel	CONTEXT 1M tokens	PRICE PER 1M \$0.20 in, \$0.70 out; half at launch	SPEED 1,068 tok/s on NVIDIA GPUs
PARAMETERS <i>not stated</i>	WEIGHTS <i>not announced</i>	BENCHMARK Mind2Web, 50 tasks: 50%	BENCHMARK STATUS VENDOR

MiniMax M3.1 Flash

MINIMAX · 19 OCT 2026

A fast model for agents and coding. MiniMax cites nearly three times faster inference and 70% lower cost on agentic tasks. [\[MiniMax\]](#)

AVAILABILITY MiniMax API, M Plan, MiniMax Code	CONTEXT <i>not stated</i>	PRICE PER 1M \$0.15 in, \$0.60 out	SPEED Over 200 tok/s, median 170+
WEIGHTS Announced for later	CYBER CyberGym: 82%	BENCHMARK DeepSWE v1.1: 73.7%	BENCHMARK STATUS VENDOR

Beam

REFLECTION AI · 16 OCT 2026

The model we listed as upcoming in issue 002 now has public weights. The Apache 2.0 license allows it to run without an internet connection. [\[Reflection AI\]](#)

PARAMETERS 501 bn, 23 bn active (MoE)	LICENSE Apache 2.0	FORMATS BF16, FP8, NVFP4	AVAILABILITY Hugging Face, Reflection API, major cloud providers
--	-----------------------	-----------------------------	---

DeepSeek V4.1 Flash-1017

DEEPSEEK · 17 OCT 2026

Shorter answers and better results at lower cost, according to DeepSeek. The API name stays deepseek-flash. [\[DeepSeek API Docs\]](#)

AVAILABILITY Web, app, API	WEIGHTS Hugging Face and ModelScope	SEARCH API GA, \$0.75 per 1,000 calls	BENCHMARK <i>chart only, no figures</i>
-------------------------------	--	--	--


VENDOR: result published by the model maker and not independently verified. Italics: not stated in the source. Context length and parameter counts for Beam and DeepSeek are not given in the announcements.

STORY / MISTRAL AI / 20 OCT 2026

Voice becomes the interface for agents

Mistral released Voxtral 3, a family of speech recognition and speech synthesis models, plus an SDK that joins them with a language model and tools in a single stream. The goal is a voice agent that answers in under a second.

ONE CONVERSATIONAL TURN, ACCORDING TO MISTRAL



Each stage is a separate model, so the language, voice or LLM can be swapped independently. A full turn in under a second, according to Mistral.

PRICING

Small Transcribe 3

\$0.004 per minute

Mini Transcribe 3

\$0.0035 per minute, 4 bn parameters

Realtime 3

\$0.008 per minute

Mini TTS 3

\$13 per million characters

Voice agent SDK

\$0.04 per minute with Small 4

SOURCES

- 1 **Mistral AI**, Voxtral 3, 20 Oct 2026 mistral.ai

WHAT WAS RELEASED

Small Transcribe 3 is the most accurate speech-to-text model in the family, for more than 50 languages. Mini Transcribe 3 has 4 billion parameters and handles recordings, Realtime 3 works live and detects the end of a turn, and Mini TTS 3 speaks with control over emotion and pacing and can clone a voice from at least 30 minutes of audio. The SDK, built on LiveKit, works over phone lines and in the browser.

WHERE IT MAKES SENSE

The obvious use is customer support: the agent understands the question, checks the order in your system and answers by voice. Since every stage is a separate model, a company can replace only the voice or only the language model.

CAUTION

Mistral made the comparisons with Whisper itself, and the announcement gives different language counts for live use. Check the list of languages, including Serbian, in the documentation.

STORY / ANTHROPIC / 21 OCT 2026

The agent that asks before it pays

Claude Duet is a separate app in which Claude works as a proactive agent: it watches your connected apps, speaks up when something needs attention and does tasks on its own virtual machine.

WHERE AUTONOMY STOPS IN DUET

ON ITS OWN

What Duet does alone

- Reads and drafts
- Sends a morning brief of what matters
- Checks connected accounts every few hours
- Builds spreadsheets, forecasts and prototypes

WITH APPROVAL OR NEVER

What waits for you

- Every message, payment, signature or deletion
- Purchases with a virtual card or Shop Pay
- Trading securities and moving money between accounts are not allowed even with approval

HOW IT WORKS

Duet carries over the memory, skills and connected apps from your Claude account. In the morning it sends a brief of what matters, and during the day it speaks up only when something needs attention. It runs on its own virtual machine with its own browser, so your computer does not have to stay on. More than 6,000 connectors are available through the Claude Marketplace, from Gmail and Slack to Stripe and Shopify.

WHERE THE LINE IS

It never sees passwords or sign-in tokens, and a dedicated classifier rejects requests to share card numbers or ID documents. It shares health and personal finance information only when you explicitly tell it to. Anthropic says Duet can misjudge urgency, is weaker at personal errands and cannot yet show you its browser.

The same day, Meta said that Muse, its consumer agent, has more than ten million accounts. The race for a personal agent has begun.

IN BRIEF

- Status
 - Research preview
- Access
 - Waitlist, Pro and Max in the US
- Platforms
 - iOS, Android, claude.ai
- Connectors
 - More than 6,000
- Price
 - Within Pro and Max plan limits

FOR READERS OUTSIDE THE US

Duet is not available outside the United States for now. Anthropic has not given a date for other countries.

SOURCES

- 1 Claude, Introducing Claude Duet, 21 Oct 2026
claude.com

ChatGPT ads arrive in Serbia

OpenAI expanded ChatGPT ads to 26 more countries, Serbia among them. The same day it announced that its strongest models will also be sold through Google Cloud, although Microsoft remains its primary cloud partner.

WHO SEES ADS

WITH ADS

Free and Go

- Ads labeled and separate from answers
- Ads chosen based on the conversation
- Advertisers do not see conversations, per OpenAI

AD-FREE

Plus, Pro and Enterprise

- Paid plans remain ad-free
- Advertising does not influence answers, per OpenAI


WHAT CHANGES IN SERBIA

Serbia is among 26 new countries with ads in ChatGPT. An ad is chosen based on the topic of the conversation: OpenAI gives the example of someone training for a marathon who asks about meals and is then shown a grocery delivery service. For now, the claims about separation of ads and privacy of conversations can only be checked by how the ads actually look. Local businesses can advertise through technology partners or the self-serve Ads Manager. OpenAI says more than 1.2 billion people use ChatGPT every week.

OPENAI ON GOOGLE CLOUD

OpenAI models and Codex will be available in Google's Model Garden, and companies with committed Google Cloud spend will be able to pay for them out of it. General availability is planned in the coming weeks. Microsoft remains OpenAI's primary cloud partner.

OpenAI says that after Mexico and Brazil it is bringing ads to 14 more Latin American countries, and the new markets also include Nigeria, South Africa, Ukraine, Cambodia, Bangladesh and Pakistan. The next phase, according to the company, brings new ad formats and measurement options.



- OPENAI ON GOOGLE CLOUD
- Where

Model Garden and Marketplace
- Billing

Also from existing GCP commitments
- Status

GA in the coming weeks
- Soon

Antigravity and Gemini Agent

SOURCES

1

OpenAI, ChatGPT Ads expands to 26 new countries as over 100,000 advertisers have used the platform, 21 Oct 2026 [openai.com](#)

2

OpenAI, OpenAI models are coming to Google Cloud, 21 Oct 2026 [openai.com](#)

MONEY / META / 21 OCT 2026

Meta builds a network around its agent

More than ten million people in the US and Canada set up a Muse account in six weeks. Meta is now organizing its strategy around the agent and introducing non-voting shares to fund its AI spending.

THIRD QUARTER AND OUTLOOK, PER REUTERS

ITEM	VALUE	COMPARISON
Q3 revenue	\$64.60 bn	Analyst estimate \$63.37 bn
Q4 outlook	\$73 to 76 bn	Above the \$73.73 bn estimate
2026 capex	\$135 to 145 bn	Previously \$130 to 145 bn
Free cash flow	Deficit of \$2.12 bn	A year earlier plus \$10.63 bn
Reality Labs 2027	Loss of about \$15 bn	Investment shifts to AI
Class C shares	January 2027	Effectively a two-for-one split

MUSE AS A NETWORK

Mark Zuckerberg expects Muse to serve billions of people one day and to bring in more revenue than the average Meta user. The idea is the same as with social networks: an agent is worth more when many people and businesses use it. That is why Meta is not charging everyone for now, but it admits Muse costs more than most new products and that monetization will come sooner than usual. CFO Susan Li expects AI businesses to add revenue in 2027.

WHO PAYS

Spending is growing faster than the money the business brings in, so free cash flow in the third quarter was negative. The new class of non-voting shares is meant for acquisitions and retaining talent without Zuckerberg losing control. The shares rose more than 4% after the announcement. Some companies are nonetheless restricting access to Muse over security concerns.

NUMBER

10 m+

Muse accounts in the US and Canada in six weeks.
[Reuters](#)

“Part of the reason why we're not just charging everyone up front to use this is we want to establish this Muse network.”


Mark Zuckerberg, per Reuters, [Reuters](#)

SOURCES

- 1 Reuters, Meta puts Muse agent at heart of AI strategy, plans new stock to power investments, 21 Oct 2026
[reuters.com](#)

Who pays for the AI race

Ahead of its listing, Anthropic shows a net loss of \$80.7 billion, OpenAI is rushing to raise about \$30 billion before that, and DeepSeek is wrapping up a round of about \$15 billion. The numbers are huge, but they do not always mean what they seem to.



ANTHROPIC: A LOSS THAT IS NOT CASH SPENT

Of the \$80.7 billion net loss in the first half, \$72 billion is an accounting charge: convertible notes are worth more because the company's estimated value went up. The truly expensive item is computing, nearly \$9 billion, four times more than a year earlier. The company has also committed to spending at least \$518 billion over a decade on cloud services, computing and infrastructure. The listing is expected in mid-November, and some investors are weighing a valuation of about \$1.5 trillion.

OPENAI: FASTER THAN ITS RIVAL

OpenAI is trying to raise about \$30 billion before then, at a valuation that could exceed \$1.4 trillion. SoftBank and MGX are in talks. The company tells investors about \$50 billion in annualized revenue and expects \$70 billion by year end.

DEEPSEEK: DOUBLE THE PLAN

DeepSeek is wrapping up a round of about 100 billion yuan, about \$15 billion, with CATL and Tencent among the biggest investors, twice the original target. Reuters' sources say the details could still change.

In September, Reuters reported that DeepSeek had hired CITIC Securities to prepare for a possible listing on Shanghai's STAR Market.

ANTHROPIC QUARTERLY REVENUE

Q2 2026

\$11.6 bn

Q3 2026

\$16.8 bn

From an investor presentation, reported by Reuters.

STATUS

OpenAI, \$30 bn

In talks, not closed

DeepSeek, \$15 bn

Nearly done, details may change

- SOURCES
- 1

Reuters

Anthropic logged net loss of over \$80 billion in first half of 2026, 20 Oct 2026

reuters.com
- 2

Reuters

OpenAI seeks to accelerate fundraising ahead of Anthropic IPO, sources say, 19 Oct 2026

reuters.com
- 3

Reuters

DeepSeek wrapping up \$15 billion fundraising to fuel AI race, 22 Oct 2026

reuters.com

NVIDIA as a central bank

NVIDIA does not just sell chips. It invests in customers, guarantees their revenue and takes capacity they cannot sell. That keeps the AI market moving, but it also ties NVIDIA to its customers' financial health.

HOW NVIDIA SUPPORTS ITS CUSTOMERS

STAKE
\$2 bn
Equity investment in CoreWeave

CAPACITY BACKSTOP
\$6.3 bn
Unsold CoreWeave capacity through 2032, deal from September

ESTIMATE
\$230 bn
Total NVIDIA backstops across the AI ecosystem, per Bank of America

The last figure is an analyst estimate, not a sum of published contracts.

WHY NVIDIA BACKSTOPS

Neoclouds such as CoreWeave borrow money to build server farms for big tech companies and AI labs. Long-term contracts bring about \$10 million per megawatt and give banks confidence, while short-term ones bring up to \$40 million but make credit harder to get. As interest rates rise, the business gets more expensive.

NVIDIA therefore acts like a central bank: it invested in CoreWeave and committed to buy any unsold capacity through 2032. The columnist also warns about five-year debt raised against three-year contracts.

WHERE THE LIMIT IS

NVIDIA is not a bank that prints money. If customers run into trouble, its revenue falls too, so the support has limits. The author stresses that the worst case is hypothetical for now, and Jensen Huang said on 15 October that CoreWeave keeps getting stronger.

Customers other than the big cloud companies, neoclouds among them, accounted for 45% of NVIDIA's \$89 billion in data center revenue last quarter, according to Visible Alpha. That gives NVIDIA every reason to keep this market liquid.

NUMBER

45 %

share of customers other than hyperscalers in NVIDIA's \$89 bn data center revenue last quarter, per Visible Alpha. [Reuters Breakingviews](#)

COMMENTARY, NOT NEWS

Breakingviews is a Reuters opinion column. The figures come from the column, and the conclusions are the author's view.

SOURCES

- 1 **Reuters Breakingviews**, Nvidia emerges as the central bank of AI, 20 Oct 2026 [reuters.com](#)

Europe wants to make AI

In Berlin, Jensen Huang said every European company will become an AI company. NVIDIA presented European partners, prefabricated data center modules and a reference architecture for running agents safely.

OPENSHELL REFERENCE ARCHITECTURE

INFERENCE	Vera Rubin NVL72	Runs the models
AGENT	Vera CPU and OpenShell	Agent harness and tool calls in a protected runtime
SILICON	BlueField-4 and NVIDIA Sentry	Policy enforced in hardware, independent of the host

Announced at GTC Berlin, with no availability date.

EUROPE AS A PRODUCER

Huang says Europe has an edge in robotics thanks to its engineering tradition, but must not only consume other people's AI. NVIDIA says it has committed nearly \$10 billion to the European ecosystem. Mistral Large 4 was trained on 3,800 Grace Blackwell GPUs in European data centers.

A MISLEADING HEADLINE

NVIDIA's post about Google Cloud announces general availability of Vera Rubin infrastructure in its headline, but the text says A5X instances become available only next month. The figure of 960,000 GPUs is a plan, not capacity that is already running.

CONTROL IN HARDWARE

The OpenShell architecture enforces policy on a BlueField-4 card, out of the agent's reach. It is the same idea as Microsoft's on page 08, moved down into silicon.

NUMBER

960,000

Rubin GPUs Google plans across multiple sites; a plan, not existing capacity. [NVIDIA Blog](#)

GOOGLE A5X
System
Vera Rubin NVL72

Status
GA next month
First customer
Ineffable Intelligence, London

- SOURCES
- 1

Reuters, Nvidia's Huang hails Europe's AI potential, flanked by robot ducks, 21 Oct 2026
[reuters.com](#)
- 2

NVIDIA Blog, Google Cloud and NVIDIA Announce General Availability of NVIDIA Vera Rubin Infrastructure, 21 Oct 2026
[blogs.nvidia.com](#)
- 3

NVIDIA Blog, NVIDIA Cloud Partners Scale Across Europe: Sovereign, Open and Built for Safe AI Agents, 21 Oct 2026
[blogs.nvidia.com](#)

INFRA / NVIDIA / 16 TO 23 OCT 2026

AI moves back to the desk

While data centers are measured in gigawatts, NVIDIA is also pushing the other side of the story: agents that run locally, on a laptop or a workstation, without sending data to the cloud.

THREE TIERS OF LOCAL AI

LAPTOP

RTX Spark

Up to 128 GB of unified memory and 1 PFLOPS; available since 16 Oct

SMALL COMPUTER

DGX Spark 64 GB

From partners starting 23 Oct

WORKSTATION

DGX Station

748 GB of coherent memory, up to 20 PFLOPS

Prices are not given in NVIDIA's announcement. Performance figures are NVIDIA's.

WHY LOCAL

A local model does not send data to the cloud, does not wait on the network and is not billed per token. For companies with sensitive data and for developers testing agents, that matters more and more. In issue 003 we saw Meta's 816 billion parameter model running on a few Mac Studio computers. This week NVIDIA showed its own ladder, from laptop to workstation.

THREE TIERS

RTX Spark laptops arrived on 16 October with up to 128 GB of unified memory. DGX Spark with 64 GB, a small computer we have followed since issue 002, reaches partners on Friday, 23 October. DGX Station has 748 GB of coherent memory, and a Windows version arrives in December. This month NVIDIA Sync will let users cluster up to eight DGX Spark systems.

WHAT IS MISSING

NVIDIA does not give prices. Before buying, check that the model you need fits in memory and how fast it runs with a longer context, since the figures in the announcement come from the vendor.

IN BRIEF

RTX Spark

ASUS, Dell, HP, Lenovo, Microsoft and MSI; Acer soon

DGX Spark 64 GB

Nemotron 3.5 Bolt 74B, Qwen3.8 27B, Gemma 4 26B

DGX Station

Shipping now; a Windows version in December

NVIDIA Sync

Clusters of up to eight DGX Spark systems

SOURCES

- 1 **NVIDIA Blog**, Bringing the Spark to Europe: Local Intelligence Powers Next Wave of Agents, 21 Oct 2026 blogs.nvidia.com

PyTorch stops picking a chip

Google open sourced TorchTPU, and AWS released TorchNeuron as a public beta. The same PyTorch code can now run on NVIDIA GPUs, Google TPUs and Amazon Trainium chips with small changes.

ONE MODEL, SEVERAL ACCELERATORS

GOOGLE
TorchTPU
Open source; vLLM, SGLang, Ray and Hugging Face

AWS
TorchNeuron
Trainium2 and Trainium3; eager and torch.compile; public beta

KUBERNETES
llm-d 1.0
Distributed inference on GPU and TPU behind one endpoint

WHY IT MATTERS

Most AI code is written in PyTorch, and for years PyTorch ran best on NVIDIA GPUs. Moving to another accelerator meant special libraries and rewrites. TorchTPU and TorchNeuron use PyTorch's official extension points, so TPU and Trainium behave like just another device. For companies, that means more choice and a stronger position when negotiating compute prices.

WHAT WAS RELEASED

Google built TorchTPU with the PyTorch team at Meta, which uses it for language models, ads and recommendations. AWS released TorchNeuron as an open source public beta for Trainium2 and Trainium3; a model moves over with one line of code, but some operations are not supported yet. Deepgram says it processes about 60% more audio seconds per dollar on Trainium. Google, Red Hat and IBM also released llm-d 1.0 for distributed inference on Kubernetes.

VENDOR NUMBERS

Google says TorchTitan reached up to 52% compute utilization on Qwen 3 8B, with near-linear scaling up to 8,192 Ironwood TPUs. AWS tested TorchNeuron on 100,000 models from Hugging Face. Both figures come from the vendors.

NUMBER

8,192

Ironwood TPUs in Google's TorchTitan scaling test, with 52% MFU on Qwen 3 8B.
[Google Cloud Blog](#)

STATUS

[TorchTPU](#)
Open source on GitHub

[TorchNeuron](#)
Public beta, Apache 2.0

[llm-d](#)
1.0, CNCF Sandbox

[FP8 on Trainium](#)
Private beta

SOURCES

- [Google Cloud Blog](#), TorchTPU is now open source: Run native PyTorch on TPUs, 20 Oct 2026
[cloud.google.com](#)
- [PyTorch](#), Introducing TorchNeuron: Native PyTorch backend for AWS Trainium, 20 Oct 2026
[pytorch.org](#)

GitHub closes the door on bad packages

After a series of attacks through npm and PyPI packages, GitHub is introducing a firewall that checks packages before installation. The same week brought license checks, support for new package managers and the retirement of six Copilot models.

COPILOT MODELS RETIRED ON 19 OCTOBER

RETIRED	REPLACEMENT	NOTE
GPT-5.5 and GPT-5.4	GPT-6.1 Sol	Chat, agent mode and code completion
GPT-5.4 mini, GPT-5 mini	GPT-6 Luna	An admin may need to enable it
Gemini 3.7 Flash	Gemini 3.8 Flash	Check automations
Grok 4.5	Grok 4.7	Check automations

PACKAGE FIREWALL

The firewall sits between your computers, CI servers or agents and the public registries. Before installation it checks whether a package carries a malware or vulnerability advisory, whether it lost its provenance attestation and whether it is older than the cooldown period you set. If a package is blocked by mistake, you request an exception through a form linked in the error message. After the preview it will be billed by usage.

LICENSES AND DEPENDENCIES

Since 16 October, Enterprise Cloud organizations can centrally define which open source licenses are allowed, and a pull request with a disallowed license gets flagged. Dependabot now supports npm 12, pnpm 11 and 12 and Bun 1.4.

COPILOT MODELS

If your automations name a retired model exactly, update them.

WHY NOW

LiteLLM on page 15 shows that even an official package from an official registry can be malicious. A cooldown before installation gives time to detect and remove such a version.

PACKAGE FIREWALL

Status

Public preview, free

For

GitHub Enterprise Cloud

Ecosystem

npm only for now

SOURCES

1

GitHub Changelog, GitHub package firewall in public preview, 20 Oct 2026
github.blog

2

GitHub Changelog, Open source license compliance generally available, 16 Oct 2026
github.blog

3


GitHub Changelog, Dependabot now supports npm 12, pnpm 11 and 12, and Bun 1.4, 20 Oct 2026
github.blog

4

GitHub Changelog, GitHub Copilot models deprecated on October 19, 2026, 19 Oct 2026
github.blog

COBOL is no longer a barrier

OpenAI published two tests for languages with little public code. Its newest models jumped from around ten percent to between 70 and 90 percent, which changes the math of maintaining banking and government systems.



TWO TESTS

MultiLangBench asks a model to build a SQLite-compatible database from scratch in one of 20 languages. The second test migrates COBOL programs to other languages, and a migration counts only when hidden tests give byte-for-byte identical output. There GPT-6.1 Sol reached 86.6% successful migrations at low effort, against 5.4% for GPT-5.6 Sol, at 39% lower cost.

THE TRAP

COBOL, which according to OpenAI runs 95% of ATM transactions, drops extra decimals unless rounding is requested. A rewritten program that rounds still runs, but produces different amounts.

NUMBER

97.9 %

successful COBOL migrations for GPT-6.1 Sol at higher effort, in OpenAI's test. [OpenAI](#)

CAUTION

The test was built by the model maker. A production migration still needs its own behavior tests, especially for rounding money amounts.

SOURCES

- 1 [OpenAI](#), Advancing coding in low-resource programming languages, 19 Oct 2026 [openai.com](#)

USE THIS / GITHUB CLI / 21 OCT 2026

Tokens that expire on their own

From version 2.103.0, GitHub CLI can use tokens that last a few hours and refresh themselves. If such a token leaks from logs, as at LiteLLM, an attacker gets hours rather than unlimited time.

THREE STEPS

1 / VERSION

gh 2.103.0 or newer
Check with gh version



2 / ENABLE

Expiring tokens
gh config set
expiring_credentials
enabled



3 / GIT

gh as the helper
gh auth setup-git, with git
2.46 or newer

The setting applies to sign-ins made after it, so sign in again with gh auth login after step 2.

WHY YOU NEED THIS

A classic GitHub CLI token is valid until you revoke it. If it leaks, an attacker can use it for days or months. That is what happened at LiteLLM: tokens ended up in logs and were used to take over releases. An expiring token becomes useless within a few hours.

HOW IT WORKS

When you turn the option on, gh stores a short-lived token and a refresh token that is rotated regularly. When the token expires, gh renews it in the background, so you do not have to sign in again. The gh auth status command shows when the current token expires.

For git over HTTPS, GitHub recommends making gh the credential helper and using git 2.46 or newer. Scripts should not store the token, but ask gh for it every time they need it.

WHERE IT WORKS

The option works on github.com and ghe.com, and on GitHub Enterprise Server from version 3.23. On servers that do not support it, gh falls back to non-expiring tokens. GitHub plans to make expiring tokens the default for new sign-ins.

USEFUL COMMANDS

Token expiry

gh auth status

Token for a script

gh auth token

Git over HTTPS

gh auth setup-git

FOR SCRIPTS


Do not store the token in a file or variable. Call gh auth token every time you need it, because a stored token will expire.

SOURCES

- 1 **GitHub Changelog**, GitHub CLI: Opt in to expiring, auto-refreshing OAuth tokens, 21 Oct 2026 github.blog

A 31-year-old problem

A Microsoft researcher and a collaborator proved the Kannan-Lovász-Simonovits conjecture, one of the major open questions in high-dimensional geometry. AI agents explored more than a hundred directions, but the decisive steps were human.



WHAT THE CONJECTURE SAYS

The KLS conjecture from 1995 says that for a broad class of shapes in high-dimensional space, the best flat cut is almost as good as the best cut of any shape, and that the gap does not grow with the number of dimensions. It determines how fast algorithms for sampling and optimization can run in such spaces.

THE ROLE OF AGENTS

Xinzhi Zhang of Microsoft Research and a collaborator used agents that explored more than a hundred directions in parallel. A researcher dropped the unpromising ones and passed the best ones on. People chose the direction, assembled the final argument and verified the proof.

INDEPENDENT CONFIRMATION

On the same day another group posted a full proof built on Microsoft's paper, and a third proof uses a closely related idea. Part of the work has been formally checked in Lean 4.

THE PAPER

Title
An $O(1)$ Bound for the KLS Constant

arXiv
2610.01447, versions of 1 and 4 Oct

Author
Xinzhi Zhang, Microsoft Research, and a collaborator

Checking
Lean 4, partly published by other researchers

SOURCES

- 1 **Microsoft Research**, Solving a 31-year-old grand challenge in mathematics: the KLS conjecture, 21 Oct 2026 microsoft.com

Thousands of agents over a mouse brain

Thousands of agents powered by GPT-Rosalind analyzed data from millions of mouse brain cells in 32 minutes. The result is not drugs but research leads that scientists now have to test in the lab.



WHAT WAS DONE

The team analyzed an existing CRISPR atlas of the mouse brain, which shows how switching off individual genes changes cell types. Each agent analyzed one change, reviewed the literature and proposed an explanation, and the reports were merged into a publicly available atlas.

The agents reproduced findings from human analyses and found new leads, among them a possible link between cholesterol regulation and neurodegeneration. Their assessments matched those of experts who did not know whose assessment they were reviewing. These are hypotheses, not therapies.

ORCHESTRA BIO

A day earlier OpenAI acquired Orchestra Bio, a platform where people and agents run drug development together and track the knock-on effects of problems, for example when a failed manufacturing batch delays a clinical trial.

NUMBER

32 min

for an analysis that people, one gene at a time, would do over months to a year, by the team's estimate. [OpenAI](#)

ORCHESTRA BIO

What it is

A platform for coordinating drug development

Announced

20 Oct 2026

Joins

GPT-Rosalind and Rosalind Workbench

SOURCES

- 1 [OpenAI](#), Using thousands of AI agents to uncover new leads in biological data, 21 Oct 2026 [openai.com](#)
- 2 [OpenAI](#), OpenAI acquires Orchestra Bio, 20 Oct 2026 [openai.com](#)




Illustration: the MMX spacecraft on its way to Phobos, a moon of Mars. Sizes and distances are not to scale. ILLUSTRATION BY THE EDITORS

ORBIT / JAXA / 19 AND 20 OCT 2026

Ten grams of Phobos

Japan's MMX spacecraft lifted off on an H3 rocket at 21:41 Belgrade time on 19 October. If all goes to plan, in 2031 it will bring back at least ten grams of soil from a moon of Mars.

LAUNCH

The H3 rocket lifted off from Tanegashima on 20 October at 04:41 Japan time. According to JAXA, the flight went as planned, and the spacecraft separated 71 minutes after liftoff.

THE PLAN

MMX should enter Mars orbit in August 2027 and orbit Phobos a month later. The 25 kilogram IDEFIX rover, built by France's CNES and Germany's DLR, will be released in 2029, and on Phobos it will weigh about 15 grams because of the weak gravity. In 2030 the main spacecraft will land twice and collect samples. ESA provides communication equipment and ground stations. The launch is a confirmed success; everything else is a mission goal.

MISSION PLAN

19 Oct 2026	Launch, H3 flight 10
Aug 2027	Mars orbit
Sep 2027	Orbit around Phobos
2029	IDEFIX rover
2030	Landing and sample
2031	Capsule in Australia

SOURCES

- 1 **AeroTime**, Japan launches mission to bring back first samples from Mars moon Phobos, 20 Oct 2026 aerotime.aero
- 2 **ESA**, Mars moon lander set to launch on 19 October, 13 Oct 2026 esa.int
- 3 **JAXA**, MMX: Martian Moons eXploration, 2026 mmx.jaxa.jp

FOLLOW UP / 16 OCT 2026 → 22 OCT 2026

What happened since the last issue

Threads we announced or left open in issues 002 and 003, and what happened to them by 22 October.

ORBIT · 16 OCT

USSF-481 has launched

Falcon 9 lifted off from Vandenberg on 15 October Pacific time, and the first stage landed on a droneship. In issue 003 the outcome was not yet known.

[SpaceX](#)

MODEL · 16 OCT

Reflection released Beam

A model listed as upcoming in issue 002 now has public weights under the Apache 2.0 license. Details on page 18.

[Reflection AI](#)

AGENTS · 20 OCT

Unintended actions, round two

After Anthropic's report in issue 003, OpenAI now describes agents that weakened their own monitoring. Page 08.

[Reuters](#)

APPLE · 16 OCT

iPhone Duo preorders opened

Apple's foldable phone can be ordered in the US and some other countries. Delivery times started slipping right away.

[MacRumors](#)

ANTHROPIC · 20 OCT

The prospectus ahead of the IPO

A draft prospectus shows a first-half net loss of \$80.7 bn, mostly an accounting charge. The listing is expected in mid-November. Page 23.

[Reuters](#)

ARTEX · 21 OCT

A tool from issue 002 in a new attack

In issue 002 we wrote about ARTEX after the attack on a Korean bank. Microsoft now says Storm-0940 used it against ConnectWise. Page 07.

[Microsoft Security](#)

SHIPPED / 16 OCT 2026 → 22 OCT 2026

Shipped this week

Only what actually became available, with the status from the official source.

MODEL · 19 OCT

Mercury 3

A diffusion model via the Inception API, OpenRouter, Baseten and Vercel, with a context of one million tokens.

AVAILABLE Inception

VOICE · 20 OCT

Voxtral 3

Speech recognition and synthesis models and an SDK for voice agents, via Mistral's API.

AVAILABLE Mistral AI

CLOUDFLARE · 21 OCT

A WAF that sees the response

Full transaction detection for all zones, including the free plan, with up to three response rules.

GA Cloudflare

GITHUB · 20 OCT

Package Firewall

Checks npm packages before installation on CI hosts and for agents. Enterprise Cloud.

PUBLIC PREVIEW GitHub Changelog

AWS · 20 OCT

TorchNeuron

A native PyTorch backend for Trainium2 and Trainium3, open source.

BETA PyTorch

MODEL · 16 OCT

Beam, open weights

501 billion parameters on Hugging Face, under the Apache 2.0 license.

AVAILABLE Reflection AI

API · 17 OCT

DeepSeek Search and Fetch API

Web search and page fetching for agents, \$0.75 per 1,000 calls.

GA DeepSeek API Docs

GITHUB · 16 OCT

Dependency license checks

Allowed license policies for Enterprise Cloud, with GitHub Code Security.

GA GitHub Changelog

GITHUB · 21 OCT

Expiring tokens in gh

An option from version 2.103.0. Guide on page 30.

PUBLIC PREVIEW GitHub Changelog

ANTHROPIC · 21 OCT

Claude Security, continuous scanning

Fix tracking, an API and an MCP connector. For Claude Enterprise customers.

BETA Claude

ANNOUNCED, NOT YET AVAILABLE

- Claude Duet waitlist, US only
- OpenAI on Google Cloud GA in the coming weeks
- Google A5X GA next month

- DGX Spark 64 GB from 23 Oct
- OpenShell architecture no date
- MiniMax weights later

5 MINUTES / 16 OCT 2026 → 22 OCT 2026

In brief

Eight stories that did not get their own page but are worth knowing.

DEVELOPERS · 20 OCT

TypeScript rewritten in Rust

Cloudflare and VoidZero released tsrs, an experimental port of the TypeScript compiler to Rust. The authors say it is up to eight times faster than TypeScript 7.0 and uses about 70% less memory.

[Cloudflare](#)

PHP · 21 OCT

PHP 8.6 RC4

The fourth release candidate for PHP 8.6 includes the PHP2Shell fixes but is not for production. RC5 is planned for 5 November.

[PHP](#)

DATA CENTERS · 21 OCT

A data center in two weeks

NVIDIA showed DSX POD, factory-built modules with the Vera Rubin platform. They deploy in about two weeks instead of six months, according to the company.

[Reuters](#)

META · 21 OCT

A \$25 billion penalty

A judge in New Mexico ordered the penalty after a jury found Meta misled the state's residents about how it shared user data, in a case related to Cambridge Analytica.

[Reuters](#)

CLOUDFLARE · 19 OCT

Clairic extracts data

A model that extracts and estimates values from unstructured text and charges only for input tokens. A preview on Workers AI, with open weights announced.

[Cloudflare](#)

ROBOTS · 21 OCT

Up to 2,000 humanoids

Germany's Agile Robots plans to build 1,000 to 2,000 small humanoid robots with NVIDIA next year, its CEO told Reuters.

[Reuters](#)

ANTHROPIC · 18 OCT

5% of shares for public benefit

The proposed fund would receive a further 2% of shares for every 50% rise in value above an undisclosed baseline, at the board's discretion.

[Reuters](#)

ORBIT · 20 OCT

Bandwagon-5

SpaceX launched the Bandwagon-5 mission from Florida, with the first stage landing on a droneship.

[SpaceX](#)

THIS WEEK IN NUMBERS / 16 OCT 2026 → 22 OCT 2026

The week in numbers

9.2_{cvss}

score of both PHP2Shell vulnerabilities, on version 4 of the scale

PHP · 21 OCT

1.5_m

hosts running a PHP version without a fix, out of 2.6 million, per Censys

REUTERS · 21 OCT

1,200+

break-in attempts by one operator's agents in Cairn, more than 100 successful

REUTERS · 19 OCT

2,250+

Microsoft 365 tenants with sign-in attempts in the campaign against ConnectWise

MICROSOFT SECURITY · 21 OCT

8_{CAs}

eMudhra intermediate certificates with weak keys, six of them for TLS

MOZILLA BUGZILLA · 17 OCT

1,295

new security patches in Oracle's October bundle

ORACLE · 20 OCT

1,068_{tok/s}

Mercury 3 on NVIDIA GPUs, according to Inception

INCEPTION · 19 OCT

501_{bn}

parameters in the Beam model, 23 billion active, Apache 2.0 license

REFLECTION AI · 16 OCT

72_{\$bn}

accounting charge within Anthropic's \$80.7 bn net loss

REUTERS · 20 OCT

89

countries with ads in ChatGPT, Serbia among them

OPENAI · 21 OCT

32_{min}

to analyze data from millions of mouse brain cells with thousands of agents

OPENAI · 21 OCT

10_g

minimum amount of Phobos soil MMX is meant to return in 2031

ESA · 19 OCT

NEXT / 16 OCT 2026 → 22 OCT 2026

What we are watching

Scheduled events and developing stories. None of this is a prediction.

→ 23 OCT

DGX Spark 64 GB

The 64 GB configuration arrives at NVIDIA's partners.

[NVIDIA Blog](#)

→ 27 OCT

LiteLLM, an embargoed vulnerability

Enterprise customers already have the fix. Public disclosure is scheduled for 27 October.

[LiteLLM](#)

→ 5 NOV

PHP 8.6 RC5

The fifth release candidate for PHP 8.6.

[PHP](#)

→ 17 TO 20 NOV

Microsoft Ignite

A chance for Microsoft to give dates for the agent controls on page 08.

[Microsoft Security](#)

→ 31 DEC

End of life for PHP 8.2

After that date the 8.2 branch gets no security fixes either.

[PHP](#)

→ 27 OCT

Chrome removes the emSign root

Sites with certificates under emSign ECC Root CA G3 should move to another hierarchy by then.

[Google](#)

→ 31 OCT

DeepSeek starts billing search

Search API calls from tools in DeepSeek Harness, Claude Code and Codex are billed from 16:00 UTC.

[DeepSeek API Docs](#)

→ MID-NOVEMBER

Anthropic goes public

Expected first public trading. OpenAI is trying to close its round before then.

[Reuters](#)

→ COMING WEEKS

OpenAI on Google Cloud

General availability of the models in Model Garden.

[OpenAI](#)

→ AUG 2027

MMX reaches Mars

Mars orbit insertion, and a month later orbit around Phobos.

[ESA](#)

SOURCES / 16 OCT 2026 → 22 OCT 2026

All sources in this issue

Sources grouped by the page where they first appear. Primary sources come first.

06 SIGNAL: THE HACKER WHO LEFT ALMOST EVERYTHING TO AGENTS

1. **Reuters.** How one hacker used AI agents and a few thousand dollars to break into scores of sites. 19 Oct 2026
reuters.com/legal/government/how-one-hacker-used-...
2. **NorForge.** Inside Cairn: Open Windows into massive AI-Driven Cyber Operations. 19 Oct 2026
norforge.com/blog/inside_cairn

07 SIGNAL: STORM-0940 AND CONNECTWISE

3. **Reuters.** Chinese-linked hackers compromise ConnectWise, target hundreds in global attack, Microsoft says. 21 Oct 2026
reuters.com/technology/chinese-linked-hackers-com-...
4. **Microsoft Security.** AI-driven, large-scale nation-state campaign using ARTEX enables supply chain compromise. 21 Oct 2026
microsoft.com/en-us/security/blog/2026/10/21/stor-...

08 SIGNAL: THE AGENT IS THE NEW INSIDER

5. **Microsoft Security.** Agents are your newest insiders: security must move at their speed. 19 Oct 2026
microsoft.com/en-us/security/blog/2026/10/19/agen-...
6. **Reuters.** OpenAI says its AI agents weakened internal monitoring system. 20 Oct 2026
reuters.com/legal/litigation/openai-says-its-ai-a-...

09 SIGNAL: DEFENSE LEARNS TO MOVE AT THE SAME SPEED

7. **OpenAI.** Patch the Planet: global response coordination as agents uncover critical vulnerabilities. 21 Oct 2026
openai.com/index/coordinated-disclosure-of-critic-...
8. **Claude.** Claude Security now scans continuously and tracks fixes. 21 Oct 2026
claude.com/resources/articles/claude-security-now-...
9. **Cloudflare.** Introducing full transaction detection for Cloudflare WAF. 21 Oct 2026
blog.cloudflare.com/full-transaction-waf/

11 PHP2SHELL: WHAT EXACTLY IS AFFECTED

10. **PHP.** PHP 8.5.12, 8.4.27, 8.3.36 and 8.2.35, critical security fixes. 21 Oct 2026
php.net/
11. **php-src on GitHub.** GHSA-w8hw-w9fx-j79w, CVE-2026-109618. 21 Oct 2026
github.com/php/php-src/security/advisories/GHSA-w-...
12. **Wiz.** PHP2Shell: AI uncovers critical RCE vulnerabilities affecting 70% of the internet. 21 Oct 2026
wiz.io/blog/php2shell-php-rce
13. **Reuters.** Government websites go dark after AI-discovered flaw sparks security scramble. 21 Oct 2026
reuters.com/world/asia-pacific/eu-agency-warns-we-...

12 PHP2SHELL: FROM THOUSANDS TO MILLIONS

14. **Reuters.** Cloudflare sees millions of exploit attempts after PHP disclosure. 22 Oct 2026
reuters.com/world/asia-pacific/cloudflare-sees-mi-...

15. **Reuters.** PHP developers issue updates to fix vulnerabilities, says flaw was found in 2024. 21 Oct 2026
reuters.com/business/php-programming-language-dev-...

13 PHP2SHELL: WHAT TO DO TODAY

16. **Cloudflare.** Cloudflare deploys protections against php2shell vulnerability. 21 Oct 2026
blog.cloudflare.com/php-vulnerabilities-waf-prote-...

14 PATCH NOW

17. **LiteLLM.** Contained: Oct 21, 2026 LiteLLM release v1.107.0 Compromised. 21 Oct 2026
docs.litellm.ai/blog/security-incident-oct-21-2026
18. **Google.** Protecting Chrome and Android Users from Compromised eMudhra Intermediate CA Certificates. 20 Oct 2026
blog.google/security/protecting-chrome-and-androi-...
19. **Mozilla Bugzilla.** Bug 2083449, eMudhra: Private key compromise of Intermediate Certificates. 17 Oct 2026
bugzilla.mozilla.org/show_bug.cgi?id=2083449
20. **Oracle.** Oracle Critical Patch Update Advisory, October 2026. 20 Oct 2026
oracle.com/security-alerts/cpuoct2026.html

16 SECURITY: RAILWAY AND OPENCODE

21. **Railway.** Security Disclosure: Wednesday October 21st. 21 Oct 2026
blog.railway.com/p/security-disclosure-wednesday-...
22. **Railway.** Security Disclosure Recommendation: Check for anomalous GitHub access. 22 Oct 2026
blog.railway.com/p/security-disclosure-recommenda-...
23. **OpenCode on X.** Notice of a breach in the Console system. 17 Oct 2026
x.com/opencode/status/2111519637049319818
24. **Cyera.** OpenCode Breach: What Happened and How to Rotate Exposed AI API Keys. 17 Oct 2026
cyera.com/research/opencode-breach-what-happened-...

17 SECURITY: EMUDHRA AND HTTPS

25. **Cloudflare.** 8 eMudhra certificates are compromised. 20 Oct 2026
blog.cloudflare.com/emudhra-certificate-compromis-...

18 MODEL WATCH

26. **Inception.** Diffusion at Scale: Introducing Mercury 3. 19 Oct 2026
inceptionlabs.ai/blog/introducing-mercury-3
27. **MiniMax.** MiniMax M3.1 Flash: Redefining Agent Performance and Cost. 19 Oct 2026
minimax.io/blog/minimax-m31f
28. **Reflection AI.** Beam is now available. 16 Oct 2026
reflection.ai/blog/beam-is-available
29. **DeepSeek API Docs.** DeepSeek-V4.1-Flash-1017 and the Search API. 17 Oct 2026
api-docs.deepseek.com/news/news261017

19 STORY: MISTRAL VOXTRAL 3

30. **Mistral AI.** Voxtral 3. 20 Oct 2026
mistral.ai/news/voxtral-3/

SOURCES / CONTINUED

20 STORY: CLAUDE DUET

31. **Claude.** Introducing Claude Duet. 21 Oct 2026
claude.com/resources/articles/claude-duet

21 PLATFORMS: CHATGPT ADS IN SERBIA

32. **OpenAI.** ChatGPT Ads expands to 26 new countries as over 100,000 advertisers have used the platform. 21 Oct 2026
openai.com/index/chatgpt-ads-expands-26-new-count...
33. **OpenAI.** OpenAI models are coming to Google Cloud. 21 Oct 2026
openai.com/index/openai-models-google-cloud/

22 MONEY: META MUSE

34. **Reuters.** Meta puts Muse agent at heart of AI strategy, plans new stock to power investments. 21 Oct 2026
reuters.com/business/meta-issue-new-class-shares-...

23 MONEY: WHO PAYS FOR THE AI RACE

35. **Reuters.** Anthropic logged net loss of over \$80 billion in first half of 2026. 20 Oct 2026
reuters.com/business/anthropic-recorded-net-loss-...
36. **Reuters.** OpenAI seeks to accelerate fundraising ahead of Anthropic IPO, sources say. 19 Oct 2026
reuters.com/world/openai-seeks-accelerate-fundrai...
37. **Reuters.** DeepSeek wrapping up \$15 billion fundraising to fuel AI race. 22 Oct 2026
reuters.com/world/asia-pacific/deepseek-raises-15...

24 MONEY: NVIDIA AS A CENTRAL BANK

38. **Reuters Breakingviews.** Nvidia emerges as the central bank of AI. 20 Oct 2026
reuters.com/commentary/breakingviews/nvidia-emerg...

25 INFRA: GTC BERLIN

39. **NVIDIA Blog.** Google Cloud and NVIDIA Announce General Availability of NVIDIA Vera Rubin Infrastructure. 21 Oct 2026
blogs.nvidia.com/blog/google-cloud-vera-rubin-inf...
40. **Reuters.** Nvidia's Huang hails Europe's AI potential, flanked by robot ducks. 21 Oct 2026
reuters.com/business/germanys-agile-make-up-2000-...
41. **NVIDIA Blog.** NVIDIA Cloud Partners Scale Across Europe: Sovereign, Open and Built for Safe AI Agents. 21 Oct 2026
blogs.nvidia.com/blog/cloud-partners-europe-sover...

26 INFRA: LOCAL AI

42. **NVIDIA Blog.** Bringing the Spark to Europe: Local Intelligence Powers Next Wave of Agents. 21 Oct 2026
blogs.nvidia.com/blog/local-ai-gtc-rtx-dgx-spark-...

27 BUILD: PYTORCH ON EVERY CHIP

43. **Google Cloud Blog.** TorchTPU is now open source: Run native PyTorch on TPUs. 20 Oct 2026
cloud.google.com/blog/topics/ai-infrastructure/to...
44. **PyTorch.** Introducing TorchNeuron: Native PyTorch backend for AWS Trainium. 20 Oct 2026
pytorch.org/blog/introducing-torchneuron-native-p-...

28 BUILD: GITHUB CLOSES THE DOOR ON BAD PACKAGES

45. **GitHub Changelog.** GitHub package firewall in public preview. 20 Oct 2026
github.blog/changelog/2026-10-20-github-package-f-...

46. **GitHub Changelog.** Open source license compliance generally available. 16 Oct 2026
github.blog/changelog/2026-10-16-open-source-lice...
47. **GitHub Changelog.** Dependabot now supports npm 12, pnpm 11 and 12, and Bun 1.4. 20 Oct 2026
github.blog/changelog/2026-10-20-dependabot-now-s...
48. **GitHub Changelog.** GitHub Copilot models deprecated on October 19, 2026. 19 Oct 2026
github.blog/changelog/2026-10-19-github-copilot-m...

29 BUILD: COBOL IS NO LONGER A BARRIER

49. **OpenAI.** Advancing coding in low-resource programming languages. 19 Oct 2026
openai.com/index/low-resource-programming-languag...

30 USE THIS: TOKENS THAT EXPIRE ON THEIR OWN

50. **GitHub Changelog.** GitHub CLI: Opt in to expiring, auto-refreshing OAuth tokens. 21 Oct 2026
github.blog/changelog/2026-10-21-github-cli-opt-i...

31 LAB: THE KLS CONJECTURE

51. **Microsoft Research.** Solving a 31-year-old grand challenge in mathematics: the KLS conjecture. 21 Oct 2026
microsoft.com/en-us/research/blog/solving-a-31-ye...

32 LAB: THOUSANDS OF AGENTS IN BIOLOGY

52. **OpenAI.** Using thousands of AI agents to uncover new leads in biological data. 21 Oct 2026
openai.com/index/scaling-ai-research-for-complex-...
53. **OpenAI.** OpenAI acquires Orchestra Bio. 20 Oct 2026
openai.com/index/openai-to-acquire-orchestra-bio/

33 ORBIT: MMX TO PHOBOS

54. **AeroTime.** Japan launches mission to bring back first samples from Mars moon Phobos. 20 Oct 2026
aerotime.aero/articles/japan-mmx-phobos-sample-re...
55. **ESA.** Mars moon lander set to launch on 19 October. 13 Oct 2026
esa.int/Science_Exploration/Space_Science/Mars_mo...
56. **JAXA.** MMX: Martian Moons eXploration. 2026
mmx.jaxa.jp/en/

34 FOLLOW UP: SINCE THE LAST ISSUE

57. **SpaceX.** Launches. 22 Oct 2026
spacex.com/launches
58. **MacRumors.** iPhone Duo Now Available to Pre-Order From Apple. 16 Oct 2026
macrumors.com/2026/10/16/iphone-duo-now-available...

36 5 MINUTES

59. **Cloudflare.** Introducing tsrs: A port of TypeScript to Rust. 20 Oct 2026
blog.cloudflare.com/typescript-rust-port/
60. **Cloudflare.** Introducing Clairic, a new System One class of extraction models. 19 Oct 2026
blog.cloudflare.com/clairic-extraction-model/
61. **Reuters.** Anthropic's public benefit fund and how it compares with other big philanthropic efforts. 20 Oct 2026
reuters.com/business/anthropics-public-benefit-fu...

ooooo

/ 004

A weekly technology review. Free, on
tutorijali.app.

Issue 004

Date 22 Oct 2026

Period 16 Oct 2026, 08:00 → 22 Oct 2026, 10:00

Web tutorijali.app



tutorijali.app