AI DIGEST

Five teams with AI find critical holes in NGINX

Anthropic sold $81.3bn of shares, an attacker hijacked a Microsoft PyPI package, Gemini 4 Neon is out.

Thursday, 19 November 2026

IN THIS ISSUE

1. NGINX: two critical holes let an attacker take over a server without a password
2. Nadella: AI finds holes so fast that outages will continue
3. Anthropic sold $81.3 billion of shares, trading starts today
4. Coalition of scientists: "Anthropic is not a safe AI company"
5. Someone hijacked the Microsoft agent-framework package on PyPI and stole credentials
6. Google released Gemini 4 Neon: third place in the AA index at $3/$15
7. OpenAI: models improved AMD GPUs by 14% in five months [single source]
8. OpenAI: agents complete almost half of the research tasks of 2-4 hours
9. Nvidia: interim CEO calms people, board under pressure
10. KPMG: more than 100 AI agents ran a trial audit and asked better questions

---

TOPIC 1

NGINX: two critical holes let an attacker take over a server without a password

SOURCES Wiz, 18.11 Wiz · Xint, 18.11 Nginx2root · @gdb, 19.11 @gdb · @levelsio, 18.11 @levelsio · HN Hacker News · HN Hacker News

NGINX stands in front of websites and APIs like a front door. According to Wiz, it serves more than 225 million sites. On 18.11 F5, the company behind NGINX, released patches for two critical vulnerabilities: CVE-2026-123633 and CVE-2026-123007. Both have a CVSS score of 9.2. Both let an attacker run code on the server without authentication, if HTTP/2 is on.

Wiz reported the first hole (use-after-free) to F5 on 21.10. F5 replied that another team had already sent the same report. Wiz found the second hole (stream window counter overflow) on 16.11. This code has been in NGINX since version 1.9.5, that is, since 2015. Wiz showed an attack on a server that only serves static files, with no backend and no third-party modules.

Five teams found the holes independently: Wiz, Ikotas Labs, Metnir, Rapid7 and Xint. Wiz searched with its own AI harness, Atlas. Ikotas says it worked with OpenAI from discovery to disclosure. Xint published a full set called NGINX2Root: five NGINX vulnerabilities with

working exploits and a separate hole in the Linux kernel. Together they give root on the host. A Shodan snapshot from 10.11 shows that 11.8 million of 11.94 million hosts with a visible NGINX version ran a version in the range 1.9.5-1.31.8.

The fix is version 1.31.9 or a patch from the distribution vendor. After the update, restart NGINX. The command `nginx -s reload` does not load the new binary. If an update is not possible now, Wiz and Xint advise a temporary HTTP/2 shutdown. Levels says his servers updated on their own because he has automatic updates on.

Why it matters: this is the third such case in a month: PHP2Shell, Java4Shell, now NGINX. The pattern is the same. A hole sits for years in code that the best people reviewed. Then several teams with AI find it almost at the same time. Wiz states that code complexity no longer protects against discovery. The practical lesson for a server owner is that the window between disclosure and attack gets shorter. Thus automatic updates and a fast restart matter more than a perfect process.

TOPIC 2

Nadella: AI finds holes so fast that outages will continue

SOURCES The New York Times, 18.11 NYT · The Guardian, 18.11 Guardian · The Guardian, 18.11 Guardian · BBC, 18.11 BBC · Marginal Revolution, 18.11 Marginal Revolution

Yesterday the digest wrote that rushed patches caused the outages at banks, airlines and the NSA. Now officials repeat this explanation. Director of National Intelligence Jay Clayton told a House committee: "These are separate service outages associated with maintenance and technical changes". He said there is no coordinated attack. Secretary of State Marco Rubio said the US is "not under systemic attack".

Satya Nadella went further at the DealBook Summit. He said AI finds vulnerabilities at a furious pace. Companies must apply patches "instantaneously", because attackers already exploit the holes. There is less time to test patches, so outages may become more frequent. Nadella compared the situation to the year 2000 problem: "Y2K was one bug, and one day". He does not know when this will end. Microsoft explains last week's Outlook outage with a latent bug in an update that was tested on old hardware. At the worst moment, more than 80% of servers did not respond.

Senators Maria Cantwell and Mike Rounds wrote a letter to Clayton. They write that AI finds weak points faster than organizations can respond. Testing patches in connected systems takes weeks or months. Economist Tyler Cowen advises against looking for a pattern in every outage and recalls Taleb's book "Fooled by Randomness".

Why it matters: the official version is now one: there is no attack, there is an overloaded patch cycle. Item 1 shows that the flow of new critical holes does not weaken. For any operations team this means a choice between two risks: apply a patch without a full test, or leave the hole open. The advantage goes to those who can roll back a change fast.

TOPIC 3

Anthropic sold $81.3 billion of shares, trading starts today

SOURCES Anthropic, 18.11 Anthropic · The New York Times, 18.11 NYT · Semafor, 18.11 Semafor · BBC, 18.11 BBC · @AnthropicAI, 19.11 @AnthropicAI · HN Hacker News

Anthropic set the IPO price at $83 per share. The earlier range was $70-78. The company sells 980 million Class A shares. Of these, 956 million are its own and 23.6 million come from individual shareholders. Trading on Nasdaq under the ticker ANT starts on 19.11. The underwriters have a 30-day option for 147 million more shares.

NYT counts this as the biggest IPO in history: $81.3 billion against SpaceX's $75 billion in June. With a $26 billion private placement, Anthropic receives $107 billion. At this price the company is valued at more than $1.5 trillion.

Semafor writes that retail investors can buy about $250 million of shares on the first day. This is twice the SpaceX record. BBC calculated that the stake of each of the seven founders will be worth at least $19.3 billion at the IPO price.

Yesterday the digest wrote that demand was seven times the supply. The company raised the price above the range although tech stocks fell after the midterm elections and the Nvidia scandal.

Why it matters: the first frontier lab becomes a public company. From this day it has quarterly reports, shareholders and a share price that reacts to every release and every incident. Market pressure on the release pace will grow. Item 4 shows that critics of the company already use this moment.

TOPIC 4

Coalition of scientists: "Anthropic is not a safe AI company"

SOURCES Business Wire, 18.11 Businesswire · @TheRundownAI, 18.11 @TheRundownAI · HN Hacker News

One day before the IPO, the Center for AI Safety (CAIS) published the MisAnthropic Statement. The text is short.

Anthropic warns that AI that improves itself can concentrate power and slip out of human control. At the same time, the company races to build such an AI first. The authors write that Anthropic weakened its safety promises and lobbied for regulation that suits it.

The list of signers is very varied. The scientists include Steven Pinker, Nobel laureate Oliver Hart and Turing Award winner Judea Pearl. The lawyers Laurence Tribe and Lawrence Lessig signed. From the industry, there are Thomas Wolf of Hugging Face, Igor Babuschkin, Max Tegmark and George Hotz. Guillaume Verdon, founder of the e/acc movement, also signed. CAIS Executive Director Dan Hendrycks says: "People in this coalition disagree on almost everything".

CAIS names three pieces of evidence. In February Anthropic rewrote its Responsible Scaling Policy and removed the unilateral commitment to stop development. The company lobbied against SB 1047, the RAISE Act and the AEGIS Act. It presents AI research automation to investors as a competitive advantage. The sources contain no answer from Anthropic.

Why it matters: criticism of safety usually comes from one group. Here, supporters of strict regulation and accelerationists who do not want regulation stand under one text. They share one thing: distrust of a company that calls itself safe and at the same time speeds up the automation of AI research. After the IPO, such statements also become a risk for shareholders.

TOPIC 5

Someone hijacked the Microsoft agent-framework package on PyPI and stole credentials

SOURCES Socket, 18.11 Socket · The Register, 18.11 Theregister · HN Hacker News

On 17.11 versions 1.26.1 and 1.26.2 of the agent-framework package appeared on PyPI. This is the official Microsoft framework for building AI agents, with about 200 thousand downloads per month. Both versions install a `.pth` file that Python runs at every interpreter start. A user does not need to import the package. It is enough that it is installed.

The malicious code collects environment variables, files with credentials and the environments of other processes. It also actively requests fresh Azure tokens and, in version 1.26.2, OIDC tokens from GitHub Actions. It compresses all this and sends it to a webhook.site address. The Register, citing OX Security, writes that 1.26.1 was rolled back, but then 1.26.2 came out through the official Microsoft publishing account. Thus the attacker still had access to the account.

The agent-framework repository on GitHub is now unavailable. On the same day the repository of ONNX Runtime, a popular engine for running models, also disappeared. On HN people write that this breaks package builds in distributions. The sources contain no official explanation from Microsoft.

Why it matters: agent packages are often installed in CI and in cloud environments, where the most valuable tokens live. Here the malicious code asks for new tokens, so even short-lived credentials reach the attacker in time.

If one of these versions is in an environment, remove it and replace all credentials that the process could access.

TOPIC 6

Google released Gemini 4 Neon: third place in the AA index at $3/$15

SOURCES Google, 18.11 Blog · @GoogleDeepMind, 18.11 @GoogleDeepMind · @GoogleDeepMind, 18.11 @GoogleDeepMind · @lennysan, 19.11 @lennysan · HN Hacker News

Neon sits in the Gemini 4 line between Carbon and Hydrogen. Google writes that at the highest reasoning level the model scores 61 in the Artificial Analysis Intelligence Index. This is third place, behind larger and slower models.

The price is $3 per million input tokens and $15 per million output tokens. The context is 1 million tokens.

The other numbers also come from the Google announcement. Neon is first in the Artificial Analysis Coding Agent Index (58.6), scores 65.2% on Terminal-Bench 4.1 and is second in the Vals Index (70.9%). In the same announcement, the company Glean writes that Neon is as good as Carbon in quality, but three times faster and costs 60% of the Carbon price. The model is available in the Gemini API, Antigravity and the Gemini app for Pro and Ultra subscribers.

First reviews are mixed. Lenny Rachitsky writes that for now Carbon seems better to him than Neon. On the same day, Google opened Genie 4 in Project Genie (details in misc).

Why it matters: in two months three labs released several models each. The price per task now matters as much as the top lines of the rankings. Neon is presented as a mid-price model with flagship quality. There are no independent measurements yet, except the AA index, which Google itself cites.

TOPIC 7

OpenAI: models improved AMD GPUs by 14% in five months [single source]

SOURCES OpenAI, 18.11 OpenAI · @OpenAI, 18.11 @OpenAI · @gdb, 18.11 @gdb · HN Hacker News

OpenAI described how its models design chips. In a partnership with AMD, a model rewrote the circuit description in SystemVerilog for a flagship GPU. The improvement in PPA (power, performance, area) reached 14% against baseline versions that AMD engineers optimized for years. OpenAI estimates this as about one year of progress in less than 5 months. AMD already included part of the changes in a design that it sent to production.

On Jalapeño, its own inference chip, the changes from the models cut the standard cell area by about 13% in the vector unit and by 10% in the matrix unit. OpenAI writes that these changes are already checked in silicon. The model checked each change with Synopsys tools. Now OpenAI trains a separate model, GPT-Synopsys, together with Synopsys.

Why it matters: chip design has formal verification: a simulation either passes or fails. Thus AI gives a result here that is easy to measure. If AMD confirms the numbers, this will speed

up the hardware cycle for the whole industry. For now, all numbers come from the OpenAI blog.

TOPIC 8

OpenAI: agents complete almost half of the research tasks of 2-4 hours

SOURCES OpenAI, 18.11 OpenAI · @OpenAI, 19.11 @OpenAI

OpenAI published a second report on how AI speeds up its own research. From March to October, the share of tasks of 2-4 hours of human work that an agent completes alone grew from 23% to 46%. For tasks of 4-8 hours, it grew from 13% to 35%. The picture is uneven: agents close 79% of training monitoring tasks, 35% of training launch tasks and only 25% of inference code tasks.

In October the median researcher used agents worth $2,400 per day at API prices. This is 62% more than in September.

The survey gave one more number. To keep agents for three months, the median researcher would give up 30% of the team compute or 40% of the people on the team. Work weeks on safety and alignment grew by 77% from January to September.

OpenAI repeats that labs must be required to track and publish metrics of recursive self-improvement (RSI). The company states separately that the report does not mean an intent to go to full RSI.

Why it matters: this is direct context for item 4. AI self-improvement now has metrics in a lab report. The numbers show that a human is still needed, but the share of agent work doubles in half a year.

TOPIC 9

Nvidia: interim CEO calms people, board under pressure

SOURCES The New York Times, 18.11 NYT · Semafor, 18.11 Semafor · Semafor, 18.11 Semafor · @BillAckman, 18.11 @BillAckman · @dnapway, 18.11 @dnapway

Yesterday the Nvidia board said for the second time that Huang will not return. On 18.11 interim CEO Colette Kress held the first all-hands meeting of the company. NYT watched the broadcast. Kress said: "Our ambition will not change"

and "we've got this". Semafor writes that she promised pay raises. In one month the market value of Nvidia fell from almost $6 trillion to $3.66 trillion.

Pressure on the board grows. Investor Bill Ackman writes that the board broke its duty of care: it fired Huang in a hurry, without facts. Travis Kalanick promises to come to the headquarters with employees who are ready to quit in support. Semafor writes that the

candidates for CEO include executives of ARM and Marvell. Jeff Bezos advises to look for a replacement inside the company.

Separately, Nvidia stopped the sale of H200 chips to China. The White House allowed their export in January. Semafor links this to pressure from Congress before Huang's testimony next month.

Why it matters: the most important supplier of AI hardware lives without a permanent leader for the second week.

The report showed that demand does not fall. But the question of who decides on sales to China and on relations with Washington is open.

TOPIC 10

KPMG: more than 100 AI agents ran a trial audit and asked better questions

SOURCES Financial Times, 18.11 FT · European Business Magazine, 18.11 Europeanbusinessmagazine · HN Hacker News

KPMG UK ran an experiment called Project Exponential. More than 100 AI agents checked the accounts of an unnamed FTSE 250 company in parallel with a normal team of auditors. The agents collected data, analyzed it, wrote draft conclusions and proposed next steps. KPMG UK head of audit Cath Burnet says that in several cases the agents dug deeper than humans and asked questions that humans did not ask.

There are also limits. An auditor had to step in when an agent matched regulatory requirements to company documents incorrectly. The company did not use the results of the experiment for official approval of the accounts. KPMG plans to extend the approach to 20 large audits, but says the technology is not yet ready for thousands of checks.

Why it matters: audit rests on large teams of junior staff who do routine checks. Agents take over this part first. KPMG does not say whether prices for clients will fall.

---

misc

- **Opus 5.5 ran a "company":** researchers from EPFL and Tubingen let Claude Opus 5.5 run the Horizon-Labs account on Hugging Face for seven weeks. 38 models got more than 50 thousand downloads, and more than 100 projects on GitHub use them. @AWuhrmann · @TheRundownAI
- **Genie 4:** Google opened Genie 4 in Project Genie for AI Ultra subscribers. A user can change the world with a prompt during a walk, and the model adds sound in real time. Blog · HN Hacker News
- **Palo Alto Networks buys Applied Compute:** an all-stock deal worth $4.5 billion, WSJ reports. The startup of former OpenAI researchers helps companies train their own open-weight models.

WSJ

- **A bridge in Crimea without Starlink:** Ars Technica and FT write that Ukraine destroyed a span of the Kerch Bridge, and operators guided the missiles through a network of relay drones. The operators were more than 250 km away. Units deliberately did not use Starshield.

  Ars Technica · FT
- **A Berlin site lost 75% of its traffic:** the author of All About Berlin writes that Google AI answers took three quarters of visitors over seven years of work. Yesterday's RevScope experiment showed a 40% drop.

  Nicolasbouliane · HN Hacker News
- **Schools against AI in search:** New York, Los Angeles and Chicago consider another search engine. Google refused to give schools search without AI features.

  NYT
- **Word of the year:** Cambridge Dictionary chose "human" as the word of 2026.

  Ac · HN Hacker News
