An AI Tool Faces Questions After America's Nationwide Ground Stop, AI Agents Get Their Own DNS, and Nvidia Reports Without Huang

The FAA is checking whether its newly expanded SMART air-traffic tool played any part in the November 13 ground stop, the Linux Foundation launches Agent Name Service to give AI agents verifiable identities, SC26's TOP500 list shows supercomputer turnover at a record low, Nvidia reports its first quarter under interim CEO Colette Kress, and Orkas is the local-first agent team developers are starring.

Themes: Governance & Safety · Infrastructure & Economics

Note: this issue relies on search summaries of the original pages. The cause of the FAA outage is still under investigation, and no evidence published so far links it to the AI tool. Nvidia's results were due after the US market close on November 17 (early on November 18 in Bangkok), so figures below are guidance and analyst estimates, not reported results. Treat secondary figures as unverified.

1) An AI Air-Traffic Tool Under the Microscope After a Nationwide Ground Stop

On November 13 the FAA stopped departures across the United States after a communications outage hit all 22 of its air route traffic control centers. It was the first nationwide ground stop since January 2023, and FlightAware counted more than 3,400 cancelled flights by the next morning. Three days earlier the FAA had expanded SMART (Strategic Management of Airspace, Routes and Trajectories), an AI scheduling tool built by Air Space Intelligence, to all 22 centers. Transportation Secretary Sean Duffy said the tool played no role and blamed an 'incredibly old' system that still depends on analog telecom lines. DHS Secretary Markwayne Mullin said officials were checking any possible link, and the developer says its technology was not responsible.

The concept, explained simply: SMART is a decision-support tool, not an autopilot. It predicts congestion and suggests reroutes and departure times so controllers can act before a jam forms rather than after: predictive instead of reactive management. The lesson today is about two things engineers confuse easily. Correlation is not causation: a failure that happens soon after a change is suspicious, but only logs and a timeline can prove the link. And in safety-critical systems, new software is normally rolled out in stages, with a 'change freeze' and a tested way to switch it off, so that when something breaks you can rule the new part in or out quickly.

Why it matters: Days 229 and 232 covered kill switches and fail-closed designs for AI. This is the real-world version of that question, in one of the most safety-critical systems there is. Whatever the investigation finds, the timing has already hurt public trust, which shows that deploying AI into old infrastructure needs evidence ready to show the public: clear logs, staged rollout and a documented rollback path. Expect calls for independent review before AI tools are extended further across public infrastructure.

2) Agent Name Service: The Linux Foundation Gives AI Agents a Verifiable Identity

At KubeCon + CloudNativeCon North America on November 10, the Linux Foundation formally launched Agent Name Service (ANS) as an open-source project, five months after announcing its intent in June. ANS is an open standard built on the Domain Name System (DNS). It ties a versioned agent name to a domain the agent's operator controls, publishes discovery information through DNS, and records signed lifecycle events (such as registration, key rotation and retirement) in a public transparency log. The specification has seven documents covering registration, DNS publication, log auditing and agent-to-agent authentication, and a technical steering committee will govern it. The foundation's AGNTCY Agent Directory Service can already verify ANS identities. Coverage so far mostly repeats the foundation's announcement, and no large production deployment has been described.

The concept, explained simply: Today, when an agent knocks on a website's door, the site cannot tell whose agent it is. ANS works like a passport plus a phone book. The name (for example, a booking agent at a company's domain) proves which organisation stands behind the agent, because only the domain owner can publish its DNS records. The transparency log is a public, append-only record, like the certificate logs behind the padlock in your browser, so anyone can check whether an agent's key has been revoked. Importantly, ANS answers 'who are you?', not 'what are you allowed to do?': permissions still have to be decided separately.

Why it matters: Day 230 covered airlines, retailers and banks blocking third-party agents, and on November 11 United rewrote its contract of carriage to cancel bookings made by unauthorised AI agents, even when the customer approved them. A big part of that pushback is that businesses cannot verify agents. A shared identity layer is a precondition for letting agents back in on clear terms, and it fits the agent-payments thread from Day 232 (ChatGPT Wallet). For builders, the practical step is to give each agent its own named, revocable identity now rather than letting it borrow a human's login.

3) SC26: The Supercomputer List Freezes While AI Reshapes Science Computing

The 68th TOP500 list of the world's fastest supercomputers was released on November 16 at the SC26 conference in Chicago (November 15-20). The top 10 is unchanged from June: China's LineShine at NSCS Shenzhen stays first, ahead of El Capitan at Lawrence Livermore and Frontier at Oak Ridge. Only 27 new systems joined the list, a record low. The biggest new entry is reported to be Lux at Oak Ridge, built with AMD EPYC CPUs and Instinct MI355X GPUs, though outlets disagree on its exact rank. AMD says its chips now power 202 systems, more than 40% of the list, and it used SC26 to announce an AI Factory for the Department of Energy's Genesis Mission on AI-driven science, plus an open-source agentic assistant for HPC users.

The concept, explained simply: TOP500 ranks machines by one test, LINPACK, which measures 64-bit (FP64) maths, the high-precision arithmetic that weather and physics simulations need. AI training mostly uses much lower precision (16-, 8- or even 4-bit numbers), which is far cheaper per operation. So the money flowing into AI data centres does not necessarily show up on TOP500: many of the world's largest AI clusters are never submitted at all. A record-low number of newcomers says less about computing slowing down and more about where the money now goes: into AI clusters measured in tokens per second, not FP64 flops.

Why it matters: Day 233's Pax Silica pact and earlier issues on the compute race described compute as a strategic resource. SC26 shows the scientific side of that shift: national labs are rebuilding around 'AI for science', in which AI models run alongside traditional simulations and agents help scientists run experiments. Nvidia has said early data suggests agentic AI can raise simulation demand by up to ten times. For readers, the lesson is to look past headline rankings and ask what precision and workload a performance claim is actually measuring.

4) Nvidia's First Earnings Without Jensen Huang

Nvidia was due to report results for its fiscal third quarter (ending October 25) after the US market close on November 17, its first report since Jensen Huang resigned as CEO (Day 230). Colette Kress, who remains CFO and is now interim CEO, will present numbers she guided as CFO in August: revenue of $108 billion, plus or minus 2%, with a gross margin of about 74%, excluding China data-centre sales. Analysts' consensus sits around $109 billion in revenue and $2.47 adjusted earnings per share, with data-centre revenue near $102 billion; Kress has said Vera Rubin would be about 20% of data-centre revenue. Live coverage also reports that Huang may be planning a comeback, which the company has not confirmed. Anthropic's IPO is expected to price around November 18.

The concept, explained simply: A share price already reflects what investors expect, so a company can grow fast and still see its stock fall if it misses the 'whisper number' or gives a cautious forecast. That forecast is guidance: the company's own estimate for next quarter. For Nvidia, guidance matters more than the past quarter because it is the best public signal of how much the cloud giants plan to spend on AI chips. Gross margin, the share of revenue left after the cost of making the chips, shows whether Nvidia still has pricing power as AMD and custom chips such as Google's TPUs compete.

Why it matters: Nvidia's report is the market's best single reading on the AI build-out. A strong forecast supports the capex plans behind every model launch this series covers; a weak one would hit the AI trade in the same week Anthropic lists on Nasdaq. Watch three things: next-quarter guidance (one bullish analyst expects about $126 billion), gross margin, and what Kress says about leadership and China. For businesses, chip supply and pricing feed directly into what you will pay for AI tokens next year.

VIRAL / OPEN-SOURCE SPOTLIGHT

Orkas: a local-first desktop app where one AI agent leads a team

Orkas (github.com/Orkas-AI/Orkas) is an MIT-licensed desktop app for running a team of AI agents from a single chat. A 'Commander' model turns your goal into a plan and hands steps to specialist sub-agents, which can work in parallel or one after another. Existing command-line agents such as Claude Code, Codex, OpenCode, OpenClaw and Hermes-Agent can join the team alongside its built-in specialists. Conversations, files, knowledge bases and

custom agents stay on your own disk by default, you bring your own API keys, and a local OpenAI-compatible model server also works. It runs from source with Node 20+ and Python 3, or from packaged installers. Trending write-ups in November reported about 7,000 stars with more than 1,000 on its first day, though star counts differ by source and date, and some listings mention a paid Pro edition.

Why it is taking off: Developers now use several coding agents at once and want one place to coordinate them without sending everything to another cloud service. Orkas applies the 'orchestrator-workers' pattern this series covered on multi-agent days to a desktop app anyone can install, and its local-first design suits teams worried about the supply-chain and credential incidents of the past week (Days 228 and 233). Treat it as an experiment: give each plugged-in agent its own narrow keys, check what tools it can run, and start on a sandbox project.

MARKET SIGNAL

This is the AI trade's biggest week of the quarter. Nvidia's results and forecast (November 17, US time) and Anthropic's IPO pricing (expected around November 18, with Nasdaq trading as ANT around November 19, according to press reports) will show whether investors still pay up for AI growth after a shaky month. Under the market story runs a trust story: an outage that may or may not involve AI grounded US flights, airlines are cancelling agent-made bookings, and the industry's answer is identity and audit standards such as ANS. The signal: money keeps flowing to compute, but the next round of adoption depends on proving AI systems are safe to put in charge of real things, and on showing who is responsible when they fail.

PRACTICAL TAKEAWAYS

1) Roll out AI tools in stages with an off switch.

Before extending an AI tool to every site, run it in a few locations first and test that you can turn it off in minutes.

2) Separate correlation from cause in incidents.

When something breaks right after a change, collect logs and a timeline before blaming or clearing the new component.

3) Give every agent its own identity.

Register agents under your own domain with separate, revocable keys instead of letting them reuse a person's login.

4) Read the terms before letting agents transact.

Check a site's terms for AI-agent clauses, since United now cancels agent-made bookings even when the customer approved them.

5) Ask what a benchmark measures.

Check the precision and workload behind any compute claim, because FP64 rankings and AI token throughput are different things.

6) Watch guidance, not just results.

In AI earnings, focus on next-quarter guidance and gross margin, the clearest signals of future chip supply and pricing.

7) Try an agent team on a sandbox project.

Install Orkas locally with narrow test keys to see how a lead agent splits work among sub-agents before using it on real code.

SELF-TEST

Q1. What is the FAA's SMART tool, and has it been shown to cause the November 13 ground stop?

A1. An AI scheduling tool that predicts congestion and suggests reroutes; officials are investigating, but no published evidence links it to the outage, which the Transportation Secretary blamed on old telecom equipment.

Q2. What does Agent Name Service do?

A2. It gives AI agents verifiable names tied to domains their operators control, publishes discovery data through DNS and logs signed lifecycle events in a public transparency log.

Q3. Does an ANS identity tell a website what an agent may do?

A3. No. ANS proves who an agent is; permissions and authorisation must still be decided separately.

Q4. Why might AI's growth not show up on the TOP500 list?

A4. TOP500 measures high-precision FP64 maths, while AI mostly uses low-precision arithmetic, and many large AI clusters are never submitted.

Q5. Why does Nvidia's guidance matter more than its past-quarter revenue?

A5. Guidance is the company's forecast for next quarter and the best public signal of how much cloud giants plan to spend on AI chips.

TOMORROW

What We're Watching: the FAA's findings on the cause of the ground stop and whether it pauses or reviews SMART.

What We're Watching: Nvidia's actual results, next-quarter guidance and any statement on Jensen Huang's future role.

What We're Watching: Anthropic's final IPO price and its first day of trading as ANT.

What We're Watching: the first named companies registering agents under Agent Name Service and whether airlines or banks accept them.

What We're Watching: further SC26 announcements on AI-for-science systems and the Genesis Mission.

TOPICS COVERED SO FAR (RECAP)

Recent days covered software supply-chain attacks through coding-agent extensions (Cursor), friend-shoring of AI supply chains (Pax Silica), enterprise agent catalogues (Imagine Enterprise), the memory wall and SRAM inference chips (Fractile), open-source AI security operations (AiSOC), agent payments (ChatGPT Wallet), gateway kill switches, fast workhorse model tiers (Gemini 4 Hydrogen), prompt-injection defences, websites blocking AI agents, IPO book-building, fail-closed research kill switches, specialised TPUs, the Fable/Mythos/Haiku 5.x refresh and attacks on AI developer platforms. Today added AI in safety-critical public infrastructure and the difference between correlation and cause in incidents (FAA SMART), DNS-based identity for AI agents (Agent Name Service), FP64 supercomputing versus low-precision AI compute (TOP500 at SC26), reading earnings guidance as an AI-spending signal (Nvidia), and local-first multi-agent orchestration (Orkas).

Sources: washingtontimes.com, click2houston.com, mensjournal.com, futurism.com and yahoo.com (FAA ground stop and SMART); linuxfoundation.org, cloudnativenow.com, securitybrief.co.uk and unite.ai (Agent Name Service); skift.com, cbsnews.com, phocuswire.com and businesstravelnews.com (United agent restrictions); hpcwire.com, nextplatform.com, tomshardware.com, newsroom.amd.com, casrai.org and nvidianews.nvidia.com (SC26 and TOP500); cnbc.com, finance.yahoo.com, fxstreet.com, techtimes.com and invezz.com (Nvidia earnings preview); renaissancecapital.com and thenextweb.com (Anthropic IPO terms); github.com/Orkas-AI/Orkas, osrepos.com and codepick.dev (Orkas). Figures are secondary reports and unverified unless from the company's own page or filing.