Nvidia Loses Its Founder-CEO, Anthropic's IPO Book Fills on Day One, and Big Websites Lock Out AI Agents

Jensen Huang resigns as Nvidia's CEO with CFO Colette Kress stepping in as interim leader, Anthropic's $70-$78 IPO is reported covered on its first roadshow day with a 20% retail slice, and airlines, retailers and banks push back on third-party AI agents acting for customers.

Themes: Enterprise & Strategy · Infrastructure & Economics

Note: direct page fetches were blocked today, so this issue relies on search summaries. The Nvidia resignation is confirmed by the company's 8-K filing with the SEC; IPO book details come from reports citing people familiar with the matter, and the agent-blocking roundup dates from October. Treat secondary figures as unverified.

1) Nvidia's Jensen Huang Resigns; CFO Colette Kress Is Interim CEO

Nvidia's board accepted Jensen Huang's resignation as president, chief executive and director on November 11, effective immediately, according to an 8-K filed with the SEC; the company announced it on November 12. CFO Colette Kress becomes interim CEO while keeping the finance role, and Harvey Jones becomes Nvidia's first-ever independent board chair. The board says it will search internal and external candidates for a permanent CEO. The departure followed days of reporting about Huang's personal conduct that coverage describes as raising national-security questions; those details are still developing and are not set out in the filing. Shares fell more than 13% in early trading on Thursday before recovering part of the drop.

The concept, explained simply: Key-person risk is what happens when one person is so tied to a company's strategy, relationships and reputation that losing them suddenly moves the share price. Huang co-founded Nvidia in 1993 and ran it for over three decades. Splitting the jobs of chair (who runs the board that supervises management) and CEO (who runs the company) is the textbook protection: the board can act independently when the boss is the problem.

Why it matters: Nvidia supplies most of the world's AI training chips, so its leadership is now an AI-industry risk, not just a company story. Kress is well known to investors, which helps continuity, but chip roadmaps (Rubin, Vera), China export talks and giant customer deals were often negotiated founder-to-founder. Expect customers and analysts to watch whether delivery schedules, pricing and partnerships stay on track. For any organisation, the lesson is to name successors and separate oversight from execution before a crisis.

2) Anthropic's IPO Order Book Is Reported Covered on Day One

Anthropic launched its IPO roadshow on November 10, offering 825 million Class A shares at $70-$78 each for a Nasdaq listing under the ticker ANT (covered here yesterday). CNBC, citing people familiar with the matter, reported the book was already covered at that range on the first day, and Bloomberg reported orders covering the shares on offer; one market note describes demand as several times oversubscribed. Anthropic is reportedly reserving about 20% of the deal for retail investors, far above the usual single-digit share. At the range the raise would be about $58-64 billion; pricing is expected around November 18-19 with trading to follow.

The concept, explained simply: Book-building is how a big IPO finds its price. Banks spend the roadshow collecting orders from investors: 'I'll take X shares at up to $Y.' If orders exceed the shares on offer, the book is 'covered'; if they exceed it several times, it is 'oversubscribed', and the company can raise the price or the range. A retail allocation is the slice reserved for ordinary investors through brokerage apps instead of large funds.

Why it matters: Strong demand on day one, even on the day Nvidia's CEO quit, says public investors still want direct exposure to a frontier lab despite near-trillion-dollar compute commitments. A large retail slice also means millions of individuals will own AI-lab shares, which raises the stakes for disclosure about safety, costs and revenue. Watch whether the final price lands at, above or below the range; that is the cleanest signal of how markets value frontier AI.

3) Airlines, Retailers and Banks Push Back on Third-Party AI Agents

A late-October roundup and follow-up reports say users of consumer agents such as Meta's Muse, xAI's Grok and OpenAI's Dots and Instinct have hit blocks on major sites. Delta says it has no agreements allowing outside agents to book flights and wants to protect customers from unauthorised automated activity; Amazon and Yelp are named as blocking; JPMorgan Chase and Charles Schwab say agents cannot touch customer accounts. United has pointed to terms of use that ban bots and unauthorised automated processes, though a specific block was not confirmed. Travel analysts note automated search and price-monitoring traffic generates heavy load without bookings.

The concept, explained simply: When an AI agent visits a website for you, the site sees a robot, not a person. Websites already use bot-management tools to stop scrapers and scalpers, and those same tools now catch helpful agents too. The missing piece is agent identity and authorisation: a way for the agent to prove 'I act for this verified customer, with these limits', like a signed letter of authority, so the site can say yes to good agents and no to bad bots.

Why it matters: Agentic commerce only works if the shops let the agents in. This is the next stage of the robots.txt opt-out story covered last week: first sites decided whether AI could read them, now they are deciding whether AI can act on them. Expect a split between 'walled garden' companies that only accept their own or partnered agents, and open protocols for signed agent identity. If you are building agents, plan for official APIs and partnerships rather than screen-scraping your way through someone else's checkout.

VIRAL / OPEN-SOURCE SPOTLIGHT

Worktrunk, git worktrees for running many coding agents in parallel

Worktrunk (max-sixty/worktrunk, by Maximilian Roos, known for PRQL and Xarray) is an open-source Rust command-line tool, invoked as *wt*, that manages git worktrees for parallel AI coding agents. A worktree is a second (or tenth) checkout of the same repository in its own folder, each on its own branch, so several agents can edit code at once without overwriting each other's files. One command, for example *wt switch -c -x claude feat*, creates the branch, the folder and launches the agent inside it; *wt list* shows every worktree with its status, and *wt remove* cleans up. It also adds hooks for local automation and per-branch CI status and summaries. Install with Homebrew or *cargo install worktrunk*.

Why it is taking off: GitHub's weekly trending page lists it at about 2,300 new stars this week, alongside claude-mem (cross-session agent memory, about 7,500 stars this week) and addyosmani/agent-skills. Developers have moved from one assistant in one editor to several agents working at once, and plain git worktrees are fiddly (the README notes you type the branch name three times). Worktrunk turns parallel agents from a juggling act into one command. Caution: star counts are from snapshots and the 'most popular worktree manager' claim is the project's own; parallel agents still need human review before branches are merged.

MARKET SIGNAL

The AI trade is learning to price people and plumbing, not just models. Nvidia's 13% early drop on a leadership change shows how concentrated AI infrastructure risk is in one supplier, while Anthropic's covered book on the same day shows investor appetite for frontier labs is intact. Underneath, financing is getting more creative: Naver's AI Factory with NVIDIA and Brookfield is funded through a Brookfield special-purpose vehicle that owns the GPUs and charges Naver by use (about $9 billion), with NVIDIA's $1 billion equity stake conditional on that financing. The signal: governance, succession and financing

structures now move AI valuations as much as benchmarks do.

PRACTICAL TAKEAWAYS

1) Map your key-person risks.
List the vendors and internal teams whose plans depend on one leader, and note what happens if that person leaves tomorrow.

2) Separate oversight from execution.
Make sure whoever approves an AI programme is not the same person who runs it, so problems can be escalated.

3) Diversify critical AI suppliers.
Where a single chip, cloud or model vendor is a hard dependency, keep a tested fallback route.

4) Read the book, not the hype.
Treat 'oversubscribed' as a demand signal, and judge an IPO by the final price, lock-ups and disclosed costs.

5) Build agents on permission, not scraping.
Use official APIs or partnerships for bookings and payments; assume websites will block unauthenticated agents.

6) Give each coding agent its own worktree.
When running agents in parallel, isolate them in separate branches and folders and review each before merging.

SELF-TEST

Q1. Who is Nvidia's interim CEO after Jensen Huang's resignation?

A1. CFO Colette Kress, who keeps her finance role while the board searches for a permanent CEO.

Q2. Why does separating the board chair and CEO roles matter?

A2. An independent chair lets the board supervise and, if needed, act against management, which is crucial when the leader is the problem.

Q3. What does it mean that an IPO book is 'covered' or 'oversubscribed'?

A3. Investor orders equal (covered) or exceed several times (oversubscribed) the shares on offer, which can push the price up.

Q4. Why are sites like Delta and banks blocking AI agents?

A4. They cannot verify who the agent acts for or limit what it does, and automated traffic loads their systems without bookings.

Q5. What problem does Worktrunk solve?

A5. It gives each parallel coding agent its own git worktree and branch with one command, so agents do not overwrite each other.

TOMORROW

What We're Watching: whether Nvidia names a permanent CEO or clarifies the circumstances of Huang's exit, and how customers react. What We're Watching: final pricing of Anthropic's IPO around November 18-19 and whether the range is raised. What We're Watching: whether airlines, retailers or banks announce official agent-access programmes or signed agent identity standards. What We're Watching: whether Naver's AI Factory signs more outside customers.

TOPICS COVERED SO FAR (RECAP)

Recent days covered fail-closed kill switches for AI research (OpenAI Control Room), specialised and confidential TPUs, Anthropic's S-1, IPO price range and terminable compute contracts, accessibility-tree desktop automation, the Fable/Mythos/Haiku 5.x model refresh, gated cyber releases, attacks on AI developer platforms, Australia's move toward an AI regulator, AI in cardiology, the AI layoff boomerang, the shared robots.txt opt-out, the Agent Client Protocol and always-on agents. Today added key-person risk and board independence at Nvidia, IPO book-building and retail allocations, websites blocking third-party AI agents and the need for agent identity, SPV-financed GPU factories, and git worktrees for parallel coding agents.

Sources: sec.gov (Nvidia 8-K, Nov 11), techcrunch.com, cnn.com, aljazeera.com and finance.yahoo.com (Nvidia CEO transition), anthropic.com/news/roadshow-launch, rttnews.com, fourweekmba.com (CNBC/Bloomberg book reports) and thenextweb.com (Anthropic IPO), pasqualepillitteri.it, aiweekly.co, androidheadlines.com and cheq.ai (sites blocking AI agents), nvidianews.nvidia.com, asiae.co.kr and theinvestor.co.kr (Naver AI Factory), github.com/trending, docs.rs/crate/worktrunk and jimmysong.io (Worktrunk), riorundown.substack.com and deepgains.substack.com (Nov 12 digests). Figures are secondary reports and unverified unless from the company's own page or filing.