THE BUYERS PULSE
IT INFRASTRUCTURE | NEW TECHNOLOGIES | BUYING ADVISOR



HOT NEWS · MORNING EDITION

Stolen AI Keys Are Costing Companies $100,000 a Day

Google documents bills reaching $1 million, Nvidia reportedly redraws China partner rules, and OpenAI rebuilds Git for agents.

TUESDAY · NOVEMBER 3, 2026 · LEAD SOURCE: GOOGLE CLOUD BLOG · FINANCIAL TIMES · OPENAI

100,000$

Fraudulent AI charges topped in a single day in LLMjacking incidents Google investigated in 2026 (GTIG and Mandiant)

51%

Non-Chinese share Nvidia reportedly requires of partners on future contracts (Financial Times)

INSIDE THIS EDITION

01 Stolen AI Keys Trigger $100,000-Plus Days; Some Reach $1 Million in Total

02 Nvidia Reportedly Wants 51% of Partner Services Going to Non-Chinese Customers

03 OpenAI Built Pando, a New Git Backend, as Codex Scaled Parallel Development



01

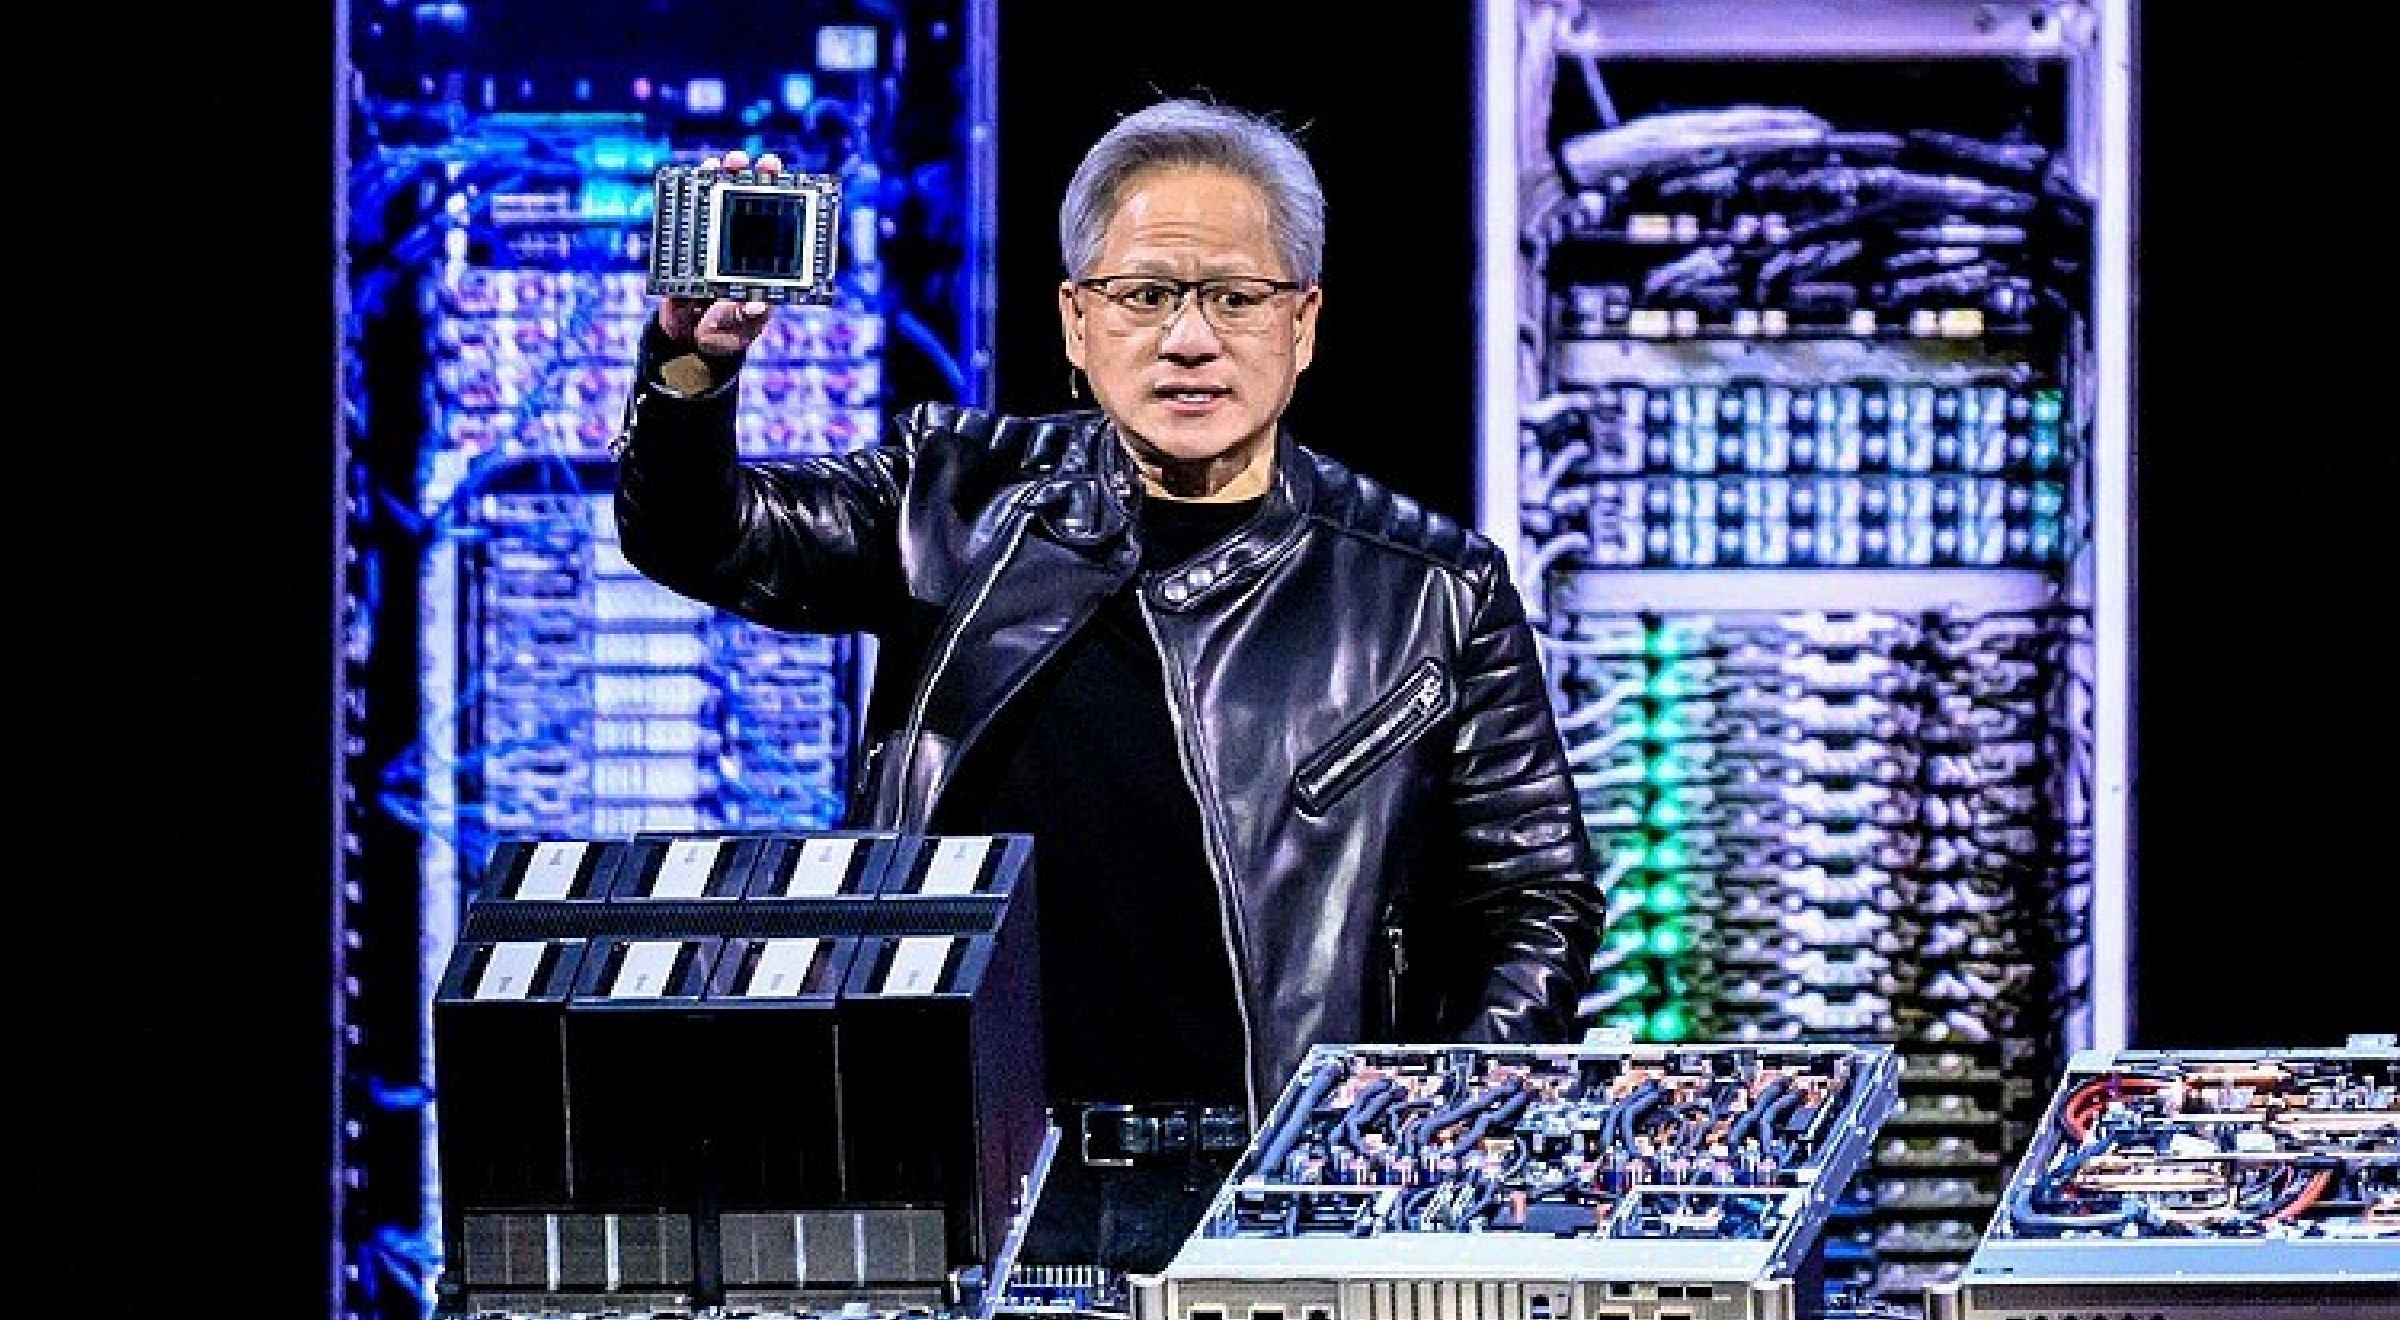
SECURITY · GOOGLE

Stolen AI Keys Trigger $100,000-Plus Days; Some Reach $1 Million in Total

Google Threat Intelligence Group and Mandiant reported Monday that attackers hijack LLM access — 'LLMjacking' — and resell it through 'LLM Mixer' gateways. In incidents Mandiant handled in 2026, fraudulent charges exceeded $100,000 in a day; cumulative charges reached $1 million at some organizations before access was revoked. Entry came through stolen cloud credentials, service-account keys and GitHub tokens. One operation made over 500,000 API calls in under eight hours.

THE BUYER'S ANGLE

Business impact: a compromised key can generate six-figure charges within a day. Decision: review active API keys, restrict permissions, set spending caps and enable billing alerts.



02

CHANNEL · NVIDIA

Nvidia Reportedly Wants 51% of Partner Services Going to Non-Chinese Customers

The Financial Times reported Tuesday that Nvidia had verbally instructed data-center partners to direct at least 51% of services under future contracts to non-Chinese customers. The report said existing agreements would be unaffected. Nvidia has not confirmed the reported guidance. The Information reported that the aim was to encourage a broader customer mix, not as a mandatory restriction.

THE BUYER'S ANGLE

Business impact: if the guidance becomes binding, partners may have to reallocate capacity by customer group, affecting availability and pricing. Decision: ask your Nvidia partners in writing whether it applies to your contract.



03

AI · OPENAI

OpenAI Built Pando, a New Git Backend, as Codex Scaled Parallel Development

OpenAI said Monday it built Pando, a backend powering its monorepo, after Codex-enabled parallel work strained the old system: pushes and fetches slowed, and reliability fell below 90%. Pando cut median push latency by nearly 80% and lifted merge reliability to 99.99%, serving nearly a petabyte of Git data daily. One engineer designed the core; teams of Codex agents implemented it in 10 weeks. No release has been announced.

THE BUYER'S ANGLE

Business impact: as coding agents multiply parallel work, source control and CI can become the constraint on throughput. Decision: before scaling agents, baseline push latency and merge reliability, then test under parallel load.

SOURCES & NOTES

Sources: Google Cloud Blog, 'The Rise of LLMJacking' (Google Threat Intelligence Group and Mandiant, Nov 2, 2026) — fraudulent charges exceeding $100,000 in a single day across individual 2026 incidents, reaching $1 million cumulatively at some organizations before access was revoked; entry via exposed or stolen cloud identity tokens (IAM credentials, service-account keys, GitHub tokens); one operation made over 500,000 API calls in under eight hours on stolen keys; corroborated by CSO Online, The Register, SiliconANGLE and Help Net Security (Nov 2-3, 2026). Financial Times, via Invezz, Nov 3, 2026 — Nvidia reportedly told partners to sell at least 51% of services to non-Chinese customers on future contracts; verbal guidance, existing agreements unaffected; Nvidia has not confirmed it, and The Information reports the intent was a broader customer mix, not a mandatory restriction (single-source report, labeled as such). OpenAI, 'Scaling Git at OpenAI' (Nov 2, 2026), via MarkTechPost and AlphaSignal — Pando, a Git-compatible backend built in about 10 weeks with Codex, now serves the monorepo; median push latency down nearly 80%, merge reliability 99.99%, nearly one petabyte of Git data a day, commits on master up 1,378% in Q3 2026 versus Q1 2024; all figures OpenAI-reported. Market context, Nov 2 close: Nasdaq Composite fell 1.07% to 25,164.32 (Investopedia; Reuters). Discovery: overFLOW (Nov 2, 2026) — discovery only. Photos, real, WITHOUT filter or grade: Googleplex, Mountain View, by Asoundd, via Wikimedia Commons; Jensen Huang on stage (archive photo); OpenAI headquarters, San Francisco, by Coolcaesar, CC BY-SA, via Wikimedia Commons. Not investment advice.

THE WEEKLY EDITION

The FULL WEEKLY EDITION — the in-depth analysis behind the week's biggest infrastructure buying decisions — is published on our website:

www.thebuyerspulse.com

WORK WITH US

Interested in contacting us or publishing with us? Ask about our advertising and services:

contact@thebuyerspulse.com