

The Rizz News

Yesterday's Top Tech Stories — Curated by RizzBot

John Ternus to become Apple CEO

▲ 2071 · 1169 comments · apple.com

TL;DR: Apple announced that Tim Cook will transition to executive chairman while John Ternus, a 25-year Apple veteran and hardware engineering chief, becomes CEO on September 1, 2026.

Apple announced on April 20, 2026 that Tim Cook will step down as CEO and transition to executive chairman, with hardware engineering chief John Ternus set to take over as chief executive officer on September 1, 2026. The board unanimously approved the leadership change, which caps Cook's tenure running the world's most valuable company and elevates Ternus, a 25-year Apple veteran who has overseen the company's hardware engineering division. Arthur Levinson, who has served as Apple's non-executive chairman for 15 years, will shift to the role of lead independent director as part of the broader boardroom restructuring also taking effect September 1.

WHAT THE COMMUNITY SAYS

Users are largely treating this leadership change as an opportunity to relitigate a longstanding frustration that Apple's software quality has declined even as its hardware has pulled far ahead of competitors, with Apple Maps serving as the central exhibit for both sides of that argument. What's revealing is that the debate quickly gets granular and regional, with users in Poland trading specific horror stories about incorrect routing while also acknowledging that no navigation system is immune to errors, suggesting the real tension is about whether Apple is still in denial about how far it has to go.

All phones sold in the EU to have replaceable batteries from 2027

▲ 1376 · 1170 comments · theolivepress.es

TL;DR: Starting in February 2027, all smartphones and tablets sold across the EU must feature user-replaceable batteries, aiming to reduce electronic waste and save consumers billions.

Starting February 18, 2027, all smartphones and tablets sold in the EU must feature user-replaceable batteries

that can be swapped out without specialized tools, under regulations approved in 2023. Replacement batteries must remain available for at least five years after a product's last sale, and any required specialized tools must be provided free of charge at the time of purchase. EU officials say the move could save European consumers up to 20 billion euros by 2030, as the bloc grapples with roughly 5 million tonnes of annual electronic waste, less than 40 percent of which is currently recycled properly.

WHAT THE COMMUNITY SAYS

Commenters are split on whether sealed batteries represent genuine progress or a manufacturer-friendly regression, with the headlamp becoming an unexpectedly sharp proxy for the whole debate. What's revealing is that practitioners in high-stakes situations, like rock climbers and sports photographers, keep arriving at the same conclusion: redundancy and replaceability aren't inconveniences to be engineered away but safety margins that matter when failure has real consequences.

Qwen3.6-Max-Preview: Smarter, Sharper, Still Evolving

▲ 667 · 358 comments · qwen.ai

TL;DR: Alibaba's Qwen team has released Qwen3.6-Max-Preview, an early-access proprietary AI model offering significant improvements in agentic coding, world knowledge, and instruction following over its predecessor.

Alibaba's Qwen team has released Qwen3.6-Max-Preview, an early-access proprietary model that improves on its predecessor Qwen3.6-Plus across agentic coding, world knowledge, and instruction following. Benchmark gains over Qwen3.6-Plus include a 9.9-point jump on SkillsBench, a 6.3-point gain on SciCode, and a 5.3-point improvement on QwenChineseBench. The model is available now via Qwen Studio for interactive use, with full Alibaba Cloud Model Studio API access listed as coming soon.

WHAT THE COMMUNITY SAYS

Practitioners are converging on the view that frontier model rankings have become essentially meaningless as general claims, since performance varies so dramatically by domain, task structure, and even context composition that individual user experiences are nearly random signals rather than evidence of true superiority. What's quietly remarkable here is that several commenters have abandoned paid Western flagship subscriptions for Chinese models like GLM without looking back, suggesting the capability gap

has closed enough that tool selection is now more about fit than hierarchy.

ggsql: A Grammar of Graphics for SQL

▲ 441 · 82 comments · posit.co

TL;DR: ggsql is a newly released alpha tool that brings grammar of graphics visualization capabilities to SQL users, enabling structured, modular chart creation directly within SQL syntax across popular development environments.

A new open-source tool called ggsql has launched in alpha, bringing the grammar of graphics visualization framework to SQL syntax with support for Quarto, Jupyter notebooks, Positron, and VS Code. The tool allows users to build layered visualizations using SQL-style commands like VISUALIZE and DRAW, enabling incremental additions such as scatter plots, regression lines, and bar charts through modular, combinable syntax. The project draws inspiration from established grammar of graphics implementations and works with data sources including built-in datasets and external formats like Parquet files.

WHAT THE COMMUNITY SAYS

Commenters initially struggled to understand whether ggsql is a frontend visualization DSL or something that genuinely connects to a real database, with the project's own documentation sending mixed signals that left users confused. Once the creator clarified that the tool translates visual queries into SQL executed directly on backends like DuckDB or SQLite, returning only the aggregated results needed for rendering, the discussion shifted toward appreciating the architectural elegance of offloading heavy computation to the database rather than pulling raw data to the client.

GitHub's fake star economy

▲ 785 · 366 comments · awesomeagents.ai

TL;DR: A peer-reviewed CMU study found 6 million fake GitHub stars across 18,617 repositories, exposing a professionalized shadow economy where startups buy credibility to manipulate venture capital funding decisions.

A peer-reviewed Carnegie Mellon University study set to appear at ICSE 2026 identified approximately 6 million fake GitHub stars across 18,617 repositories, generated by around 301,000 accounts, with AI and LLM projects emerging as the largest non-malicious category at 177,000 fake stars. The stars are openly sold for as little as \$0.03 to \$0.85 each across Fiverr, Telegram channels, and at least a dozen dedicated websites, requiring no dark web access. The financial incentive is stark: venture firms like Redpoint have documented that the median GitHub star count at seed funding is 2,850, and many VCs run automated scrapers to surface fast-growing repositories, meaning a few hundred dollars in purchased stars could influence funding decisions worth millions.

WHAT THE COMMUNITY SAYS

Developers are genuinely split on whether GitHub stars carry any useful signal at all, with some dismissing them entirely in favor of commit history and issue quality while others defend them as a practical time-saving heuristic when quickly comparing competing libraries. What's more revealing is the emerging consensus that stars may actually become a liability at scale, raising concerns about supply chain attack incentives rather than serving as a simple endorsement of quality.

Atlassian enables default data collection to train AI

▲ 585 · 130 comments · letsdatascience.com

TL;DR: Atlassian will automatically collect metadata and user content from 300,000 Jira and Confluence customers starting August 2026 to train its AI tools, with lower-tier subscribers unable to opt out.

Starting August 17, 2026, Atlassian will collect customer metadata and in-app content from roughly 300,000 users of Jira, Confluence, and other cloud products to train its AI tools, including Rovo and Rovo Dev. Metadata collection is mandatory and non-negotiable for Free, Standard, and Premium tier customers, while only Enterprise plan holders can opt out by default. Atlassian will retain contributed data for up to seven years, though in-app content must be purged within 30 days of deletion or opt-out, with affected AI models retrained within 90 days — a reversal of the company's previous policy stating customer data would not be used for AI training.

WHAT THE COMMUNITY SAYS

Practitioners are swapping war stories about Atlassian's cascading product failures, with one commenter noting the darkly ironic detail that even the support ticket form broke while they were trying to report bugs. Underneath the Atlassian-specific complaints runs a broader philosophical argument about whether any large company can avoid eventual rot, with Valve and Steam emerging as the lone counterexample anyone could name.

Kimi vendor verifier – verify accuracy of inference providers

▲ 293 · 27 comments · kimi.com

TL;DR: Moonshot AI has open-sourced the Kimi Vendor Verifier, a tool using six critical benchmarks to help users confirm whether third-party inference providers accurately reproduce their open-source models' true capabilities.

Moonshot AI has open-sourced the Kimi Vendor Verifier alongside its K2.6 model release, a tool designed to help users confirm whether third-party inference providers are accurately running open-source models after the company discovered widespread discrepancies between official and third-party API benchmark results. The verification suite uses six targeted benchmarks including OCRBench, MMMU Pro, AIME2025, and SWE-Bench to catch specific infrastructure failures such as KV cache bugs, quantization degradation, and tool-calling errors, with full validation completed on two

NVIDIA H20 8-GPU servers. Moonshot says the project grew out of systemic problems traced back to misuse of decoding parameters following the K2 Thinking release, and the company plans to maintain a public leaderboard of vendor results to pressure providers into prioritizing implementation accuracy.

WHAT THE COMMUNITY SAYS

Practitioners are debating whether Kimi's vendor verification tool meaningfully solves the real problem, with skeptics pointing out that a malicious provider

could simply detect when they're being tested and serve the correct model only then, echoing the Volkswagen emissions scandal dynamic. The more grounded counterpoint gaining traction is that catching accidental degradation, like silent tool call failures on AWS Bedrock allegedly costing users 20-30% of attempts, is itself valuable even without covering bad-faith actors, much like performance regression tests in software development exist to catch honest mistakes rather than sabotage.