TechLife NEWS

Nov 29, 2025

IPHONE PRO CAMERA SET

THE MOST ADVANCED OPTICAL AND COMPUTATIONAL POWER

GOVERNMENT AGENCIES

AMAZON COMMITS UP TO $50B TO BUILD AI COMPUTING CAPACITY

Console PC

STEAM BECOMES A MAJOR REVENUE CHANNEL FOR PLAYSTATION

Your Business Is Invisible Where It Matters Most

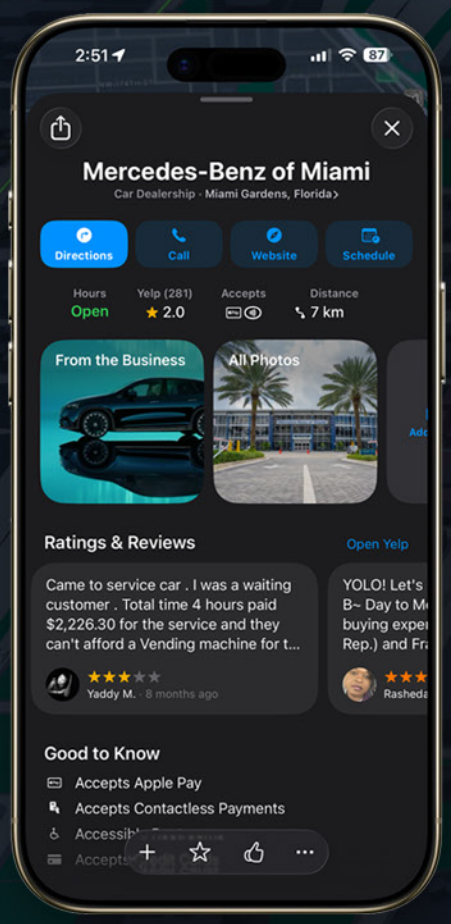

iPhone Is the Primary Gateway to Local Businesses

Every day, millions of people search for products, services, and stores using their iPhone. With 56% adoption in the U.S., Apple users represent the most valuable consumer segment—and yet most businesses have never claimed their place card on Apple Business Connect, leading to:

- ⚠ Outdated or incomplete information on Apple Maps
- ⚠ Missed visibility in Siri search results
- ⚠ No access to Showcases for promotions or deals
- ⚠ No branded email identity tied to each location
- ⚠ No control over customer reviews or business details
- ⚠ Lost opportunities to reach 186 million iPhone users who rely on Apple every day

Engage the most loyal and high-value audience in the U.S.
Claim your business. Connect your location.

One Platform. Full Control. Maximum Exposure.

From Apple Maps to Siri, from branded emails to AppleMagazine promotions—your entire presence becomes unified, updated, and strategically managed.

- Set up your business card and brand details
- Claim and verify all locations
- Publish deals, promotions, offers, and events
- Improve visibility on Apple Maps and Siri
- Use branded emails to strengthen credibility and trust
- Manage customer reviews and provide timely first responses
- Enable Apple Pay for faster, secure customer transactions
- Access AppleMagazine self-service brand promotion & advertising

Free business and brand setup included.
Marketing management starting as low as $89.99/month.

Get Started

Up to 3 locations included.

PublishChannel | Business Connect Third-Party Agency

Promote Smarter. Rank Higher. Lead the Market.

Luminar 4

Perfect Photo, Sucky Skies?

AI-powered Sky Replacement - Perfection in Seconds

© Sherwin Magsino

Say Goodbye to Bland and Boring Skies

Get Luminar 4, powered by new and innovative technologies like Accent AI and AI Sky Replacement. Instantly replace skies to change the mood of your photo. No complicated masking or brushes required!

Wherever your creative vision takes you, Luminar has the tools, smarts and style to help you achieve your absolute best.
Your Photography. Elevated.

Grab a copy at l.skylum.com/applemag
Use promo code **APPLEMAG** for $10 off your purchase

SKYLUM Apple Magazine
OFFICIAL PARTNERS

SUMMARY


TESLA

TESLA REVISES ROBOTAXI RAMP PLAN FOR AUSTIN AS INTERNAL MODELING ADJUSTS EARLY DEPLOYMENT EXPECTATIONS

Tesla has updated its internal ramp targets for the first wave of robotaxi operations planned for Austin, refining earlier projections as engineering teams continue to evaluate deployment timelines, service-area density and supporting infrastructure. The update follows new comments from Elon Musk, who noted that Tesla's earliest robotaxi fleet in Austin will scale more gradually than initial estimates suggested. The revision reflects ongoing analysis inside the company as it aligns its fully autonomous vehicle program with manufacturing readiness, software maturation and regulatory conditions in Texas.

According to the updated figures Musk posted, Tesla's previous planning cycles projected higher early fleet counts for Austin once the company introduced its dedicated robotaxi platform. Internal modeling now points to a more sequential approach, with smaller initial clusters designed to validate routing behavior, operational reliability and service predictability. While Musk reaffirmed Austin as a core launch market, the revised numbers indicate Tesla is focusing on structured early-stage operations rather than broad coverage from the outset.

The adjustment comes as Tesla continues to test advanced iterations of its autonomy software, refine the design of its purpose-built robotaxi vehicle and evaluate infrastructure options around Gigafactory Texas. The company has emphasized that early robotaxi deployments must integrate vehicle availability, charging logistics and service-area mapping with high precision, making recalibrations expected as new data emerges from simulation tools and on-road testing.

UPDATED FLEET RAMP REFLECTS A MORE MEASURED STARTING POINT

The revised timeline outlined by Musk resets expectations for Austin's initial robotaxi density, with Tesla moving toward an approach that prioritizes the reliability of early service nodes before broad expansion. Tesla's planning teams have been iterating models that incorporate traffic patterns, projected rider demand, vehicle turnover rates and charging cycles. These factors influence how many vehicles a market can realistically support during the earliest operational windows.

While earlier estimates envisioned faster scaling, the updated outlook favors deploying smaller fleets that can be monitored closely as Tesla validates behavior across constrained service zones. These adjustments follow expanded internal testing in Texas, where updated versions of Tesla's Full Self-Driving system are being evaluated on a near-continuous basis. Engineers are analyzing lane-selection behavior, interaction with local signage and the influence of Austin's varied driving environments on autonomous-system consistency.

Musk's update does not alter the company's broader plan to introduce robotaxi services in multiple U.S. cities, but it signals Tesla's intention to anchor the Austin rollout on datasets that emerge from incremental scaling rather than aggressive initial saturation. The revised expectations also align with Tesla's continued work on its next-generation platform, which will support the robotaxi program once manufacturing transitions at Gigafactory Texas enter higher-volume production phases.

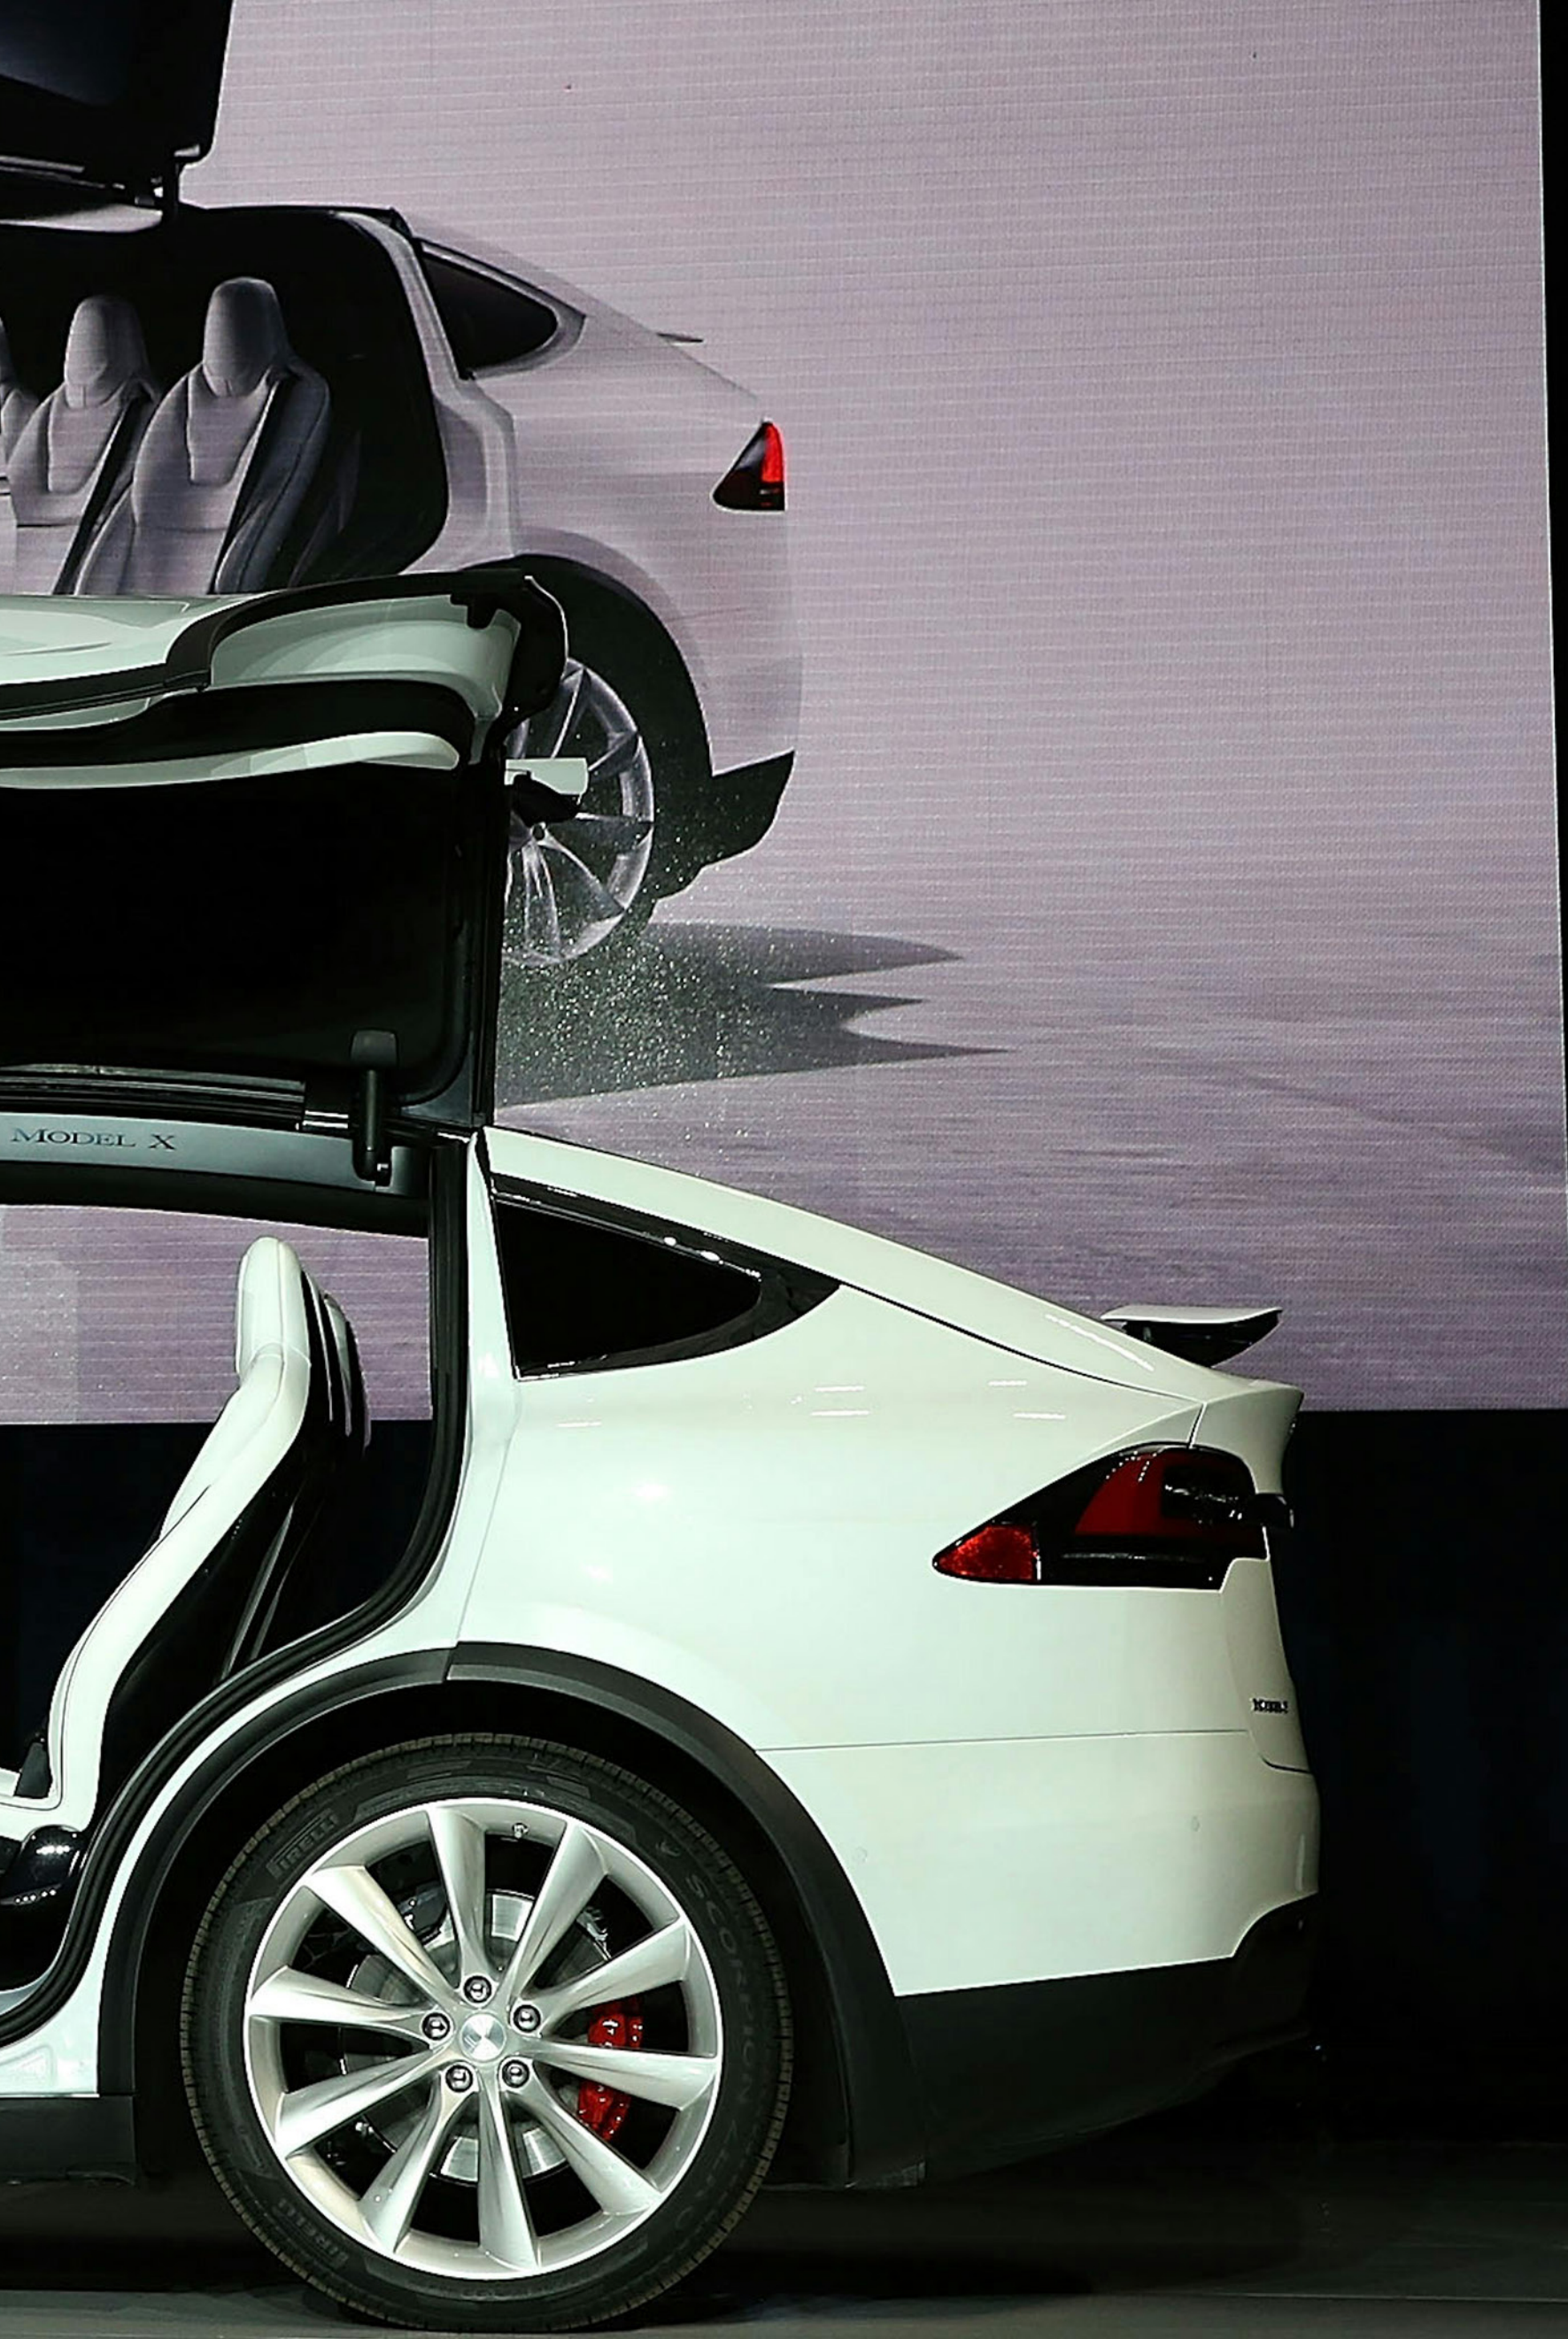
MODEL X

Tesla's internal communications to teams describe the change not as a reduction in long-term targets, but as an adjustment in how the initial deployment sequences will unfold. The company is analyzing when its next-generation vehicles will reach sufficient maturity for robotaxi service and how to structure launch conditions so early vehicle clusters can operate with consistent performance before expanding into larger networks.

SOFTWARE MATURITY, MANUFACTURING SCHEDULING AND INFRASTRUCTURE PREPARATION SHAPE THE TIMELINE

The updated fleet ramp is intertwined with Tesla's software-development cadence. As the company transitions into newer autonomy architectures that combine planning, perception and control into unified models, real-world evaluations feed directly into internal models that project how robotaxi deployments will function. These evaluations influence decisions around vehicle counts, service-area size and the level of redundancy required for early-stage fleets.

Tesla's manufacturing schedule at Gigafactory Texas also plays a role. The robotaxi program will rely on a next-generation vehicle platform designed for high-volume production, cost efficiency and autonomous-optimized packaging. As these vehicles progress through prototype stages and approach final validation, Tesla's operations teams determine how many units can enter early service programs without disrupting production ramps meant for other models.

Charging infrastructure is another variable. Robotaxi fleets require high-turnover charging patterns, consistent station availability and service centers capable of handling the maintenance cadence of vehicles operating for long continuous periods. As Tesla expands its Supercharger network and begins deploying next-generation charging hardware, the infrastructure supporting robotaxi vehicles becomes increasingly complex. Austin's layout requires distributed charging nodes integrated with Tesla's fleet-management software so vehicles can be routed efficiently for power cycles.

Regulatory frameworks in Austin and across Texas guide how quickly Tesla can activate wide-scale service. While the state remains open to autonomous-vehicle pilots, each program must align with safety reporting standards, local road-usage frameworks and ongoing coordination with municipal authorities. Those requirements influence the scale and configuration of early deployments, which Tesla's planning teams incorporate into ramp models that Musk referenced in his post.

AUSTIN REMAINS CENTRAL TO TESLA'S AUTONOMOUS-VEHICLE ROLLOUT

Despite the reduced early fleet projections, Austin remains a focal point for Tesla's long-term robotaxi ambitions. The city's role stems from multiple factors: proximity to Gigafactory Texas, ongoing autonomy testing in the region, and the city's participation in programs that evaluate emerging mobility systems. Tesla's presence in Austin also gives the company access to dense driving data that reflects a range of real-world scenarios, from


TESLA

urban corridors to multilane interchanges and suburban routes.

Tesla continues to test autonomy software extensively on Texas roads, capturing data on interactions with cyclists, pedestrians, construction zones and lane-sharing conditions common in the region. These datasets refine behavior models for the upcoming robotaxi platform and shape decisions about where early services can operate most effectively. Internal teams analyze which areas of the city provide the most predictable routing environments and where additional validation is required before robotaxis can begin serving passengers.

Musk's updated targets also reflect broader operational considerations. Early robotaxi services will depend on data-center coordination and fleet-management systems capable of routing and monitoring large vehicle clusters with real-time adjustments. Tesla's Austin deployment will serve as a foundation for testing these systems, which coordinate navigation, charging, dispatch timing and fallback modes under varying road conditions and customer demand patterns.

In addition to its manufacturing and testing significance, Austin continues to attract technology-sector development and mobility-focused initiatives, giving Tesla a supportive ecosystem as it advances its autonomous project. The city's evolving infrastructure, combined with Tesla's operational footprint, makes Austin an early proving ground for the company's long-term robotaxi network.

PLAY YOUR FAVORITES

Stream more than 15,000 episodes of your favorite daytime, primetime and late night CBS shows, catch up on classic hits or enjoy exclusive CBS All Access Original Series.

WATCH IT LIVE

Watch live sports, shows and news at home or on-the-go with your local CBS station, CBSN, CBS Sports HQ and ET Live 24/7.

CBS CBSN CBS SPORTS HQ ET LIVE

OBCESS OVER ORIGINALS

Get hooked on critically-acclaimed CBS All Access Original Series.

STREAM ACROSS SCREENS

It's easier than ever to watch on your phone, tablet, desktop or streaming media player.

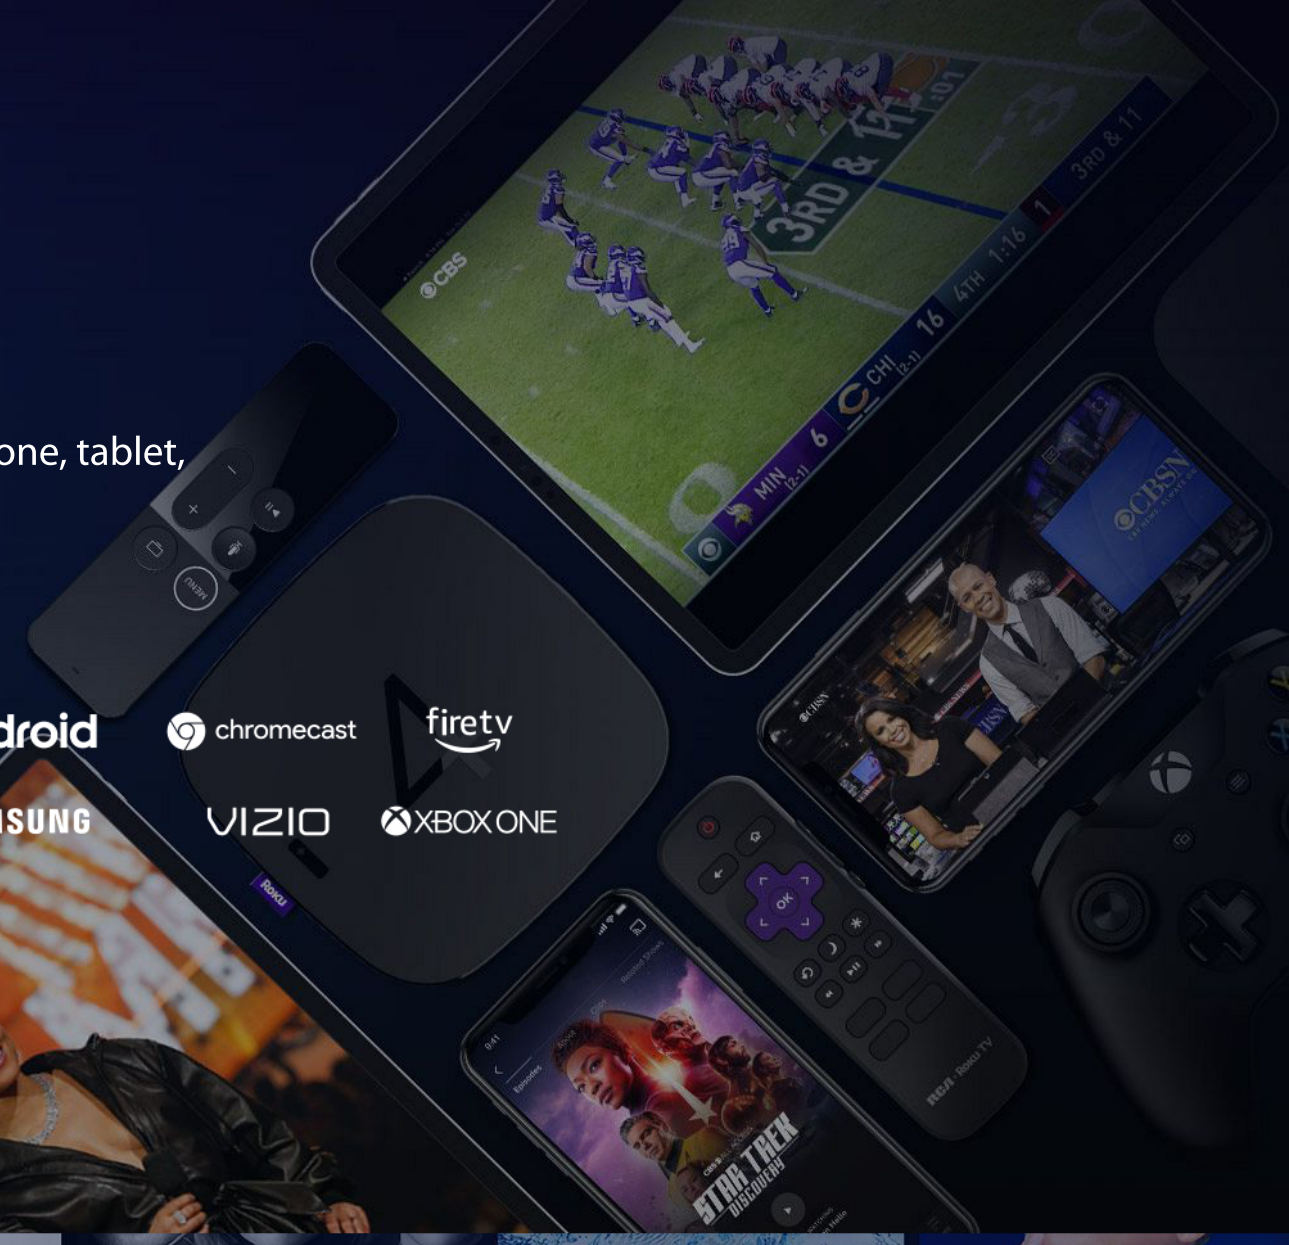

iPad iPhone tv android chromecast firetv

LG PS4 Roku SAMSUNG VIZIO XBOX ONE

xfinity

THOUSANDS OF EPISODES. LIVE TV. ORIGINAL SERIES.

Paramount+

Plans start at $4.99/month. Cancel anytime.

TRY IT FREE

ANDROID ADDS NEW DUAL-BAND WI-FI HOTSPOT OPTION AS CONNECTIVITY FEATURES EXPAND

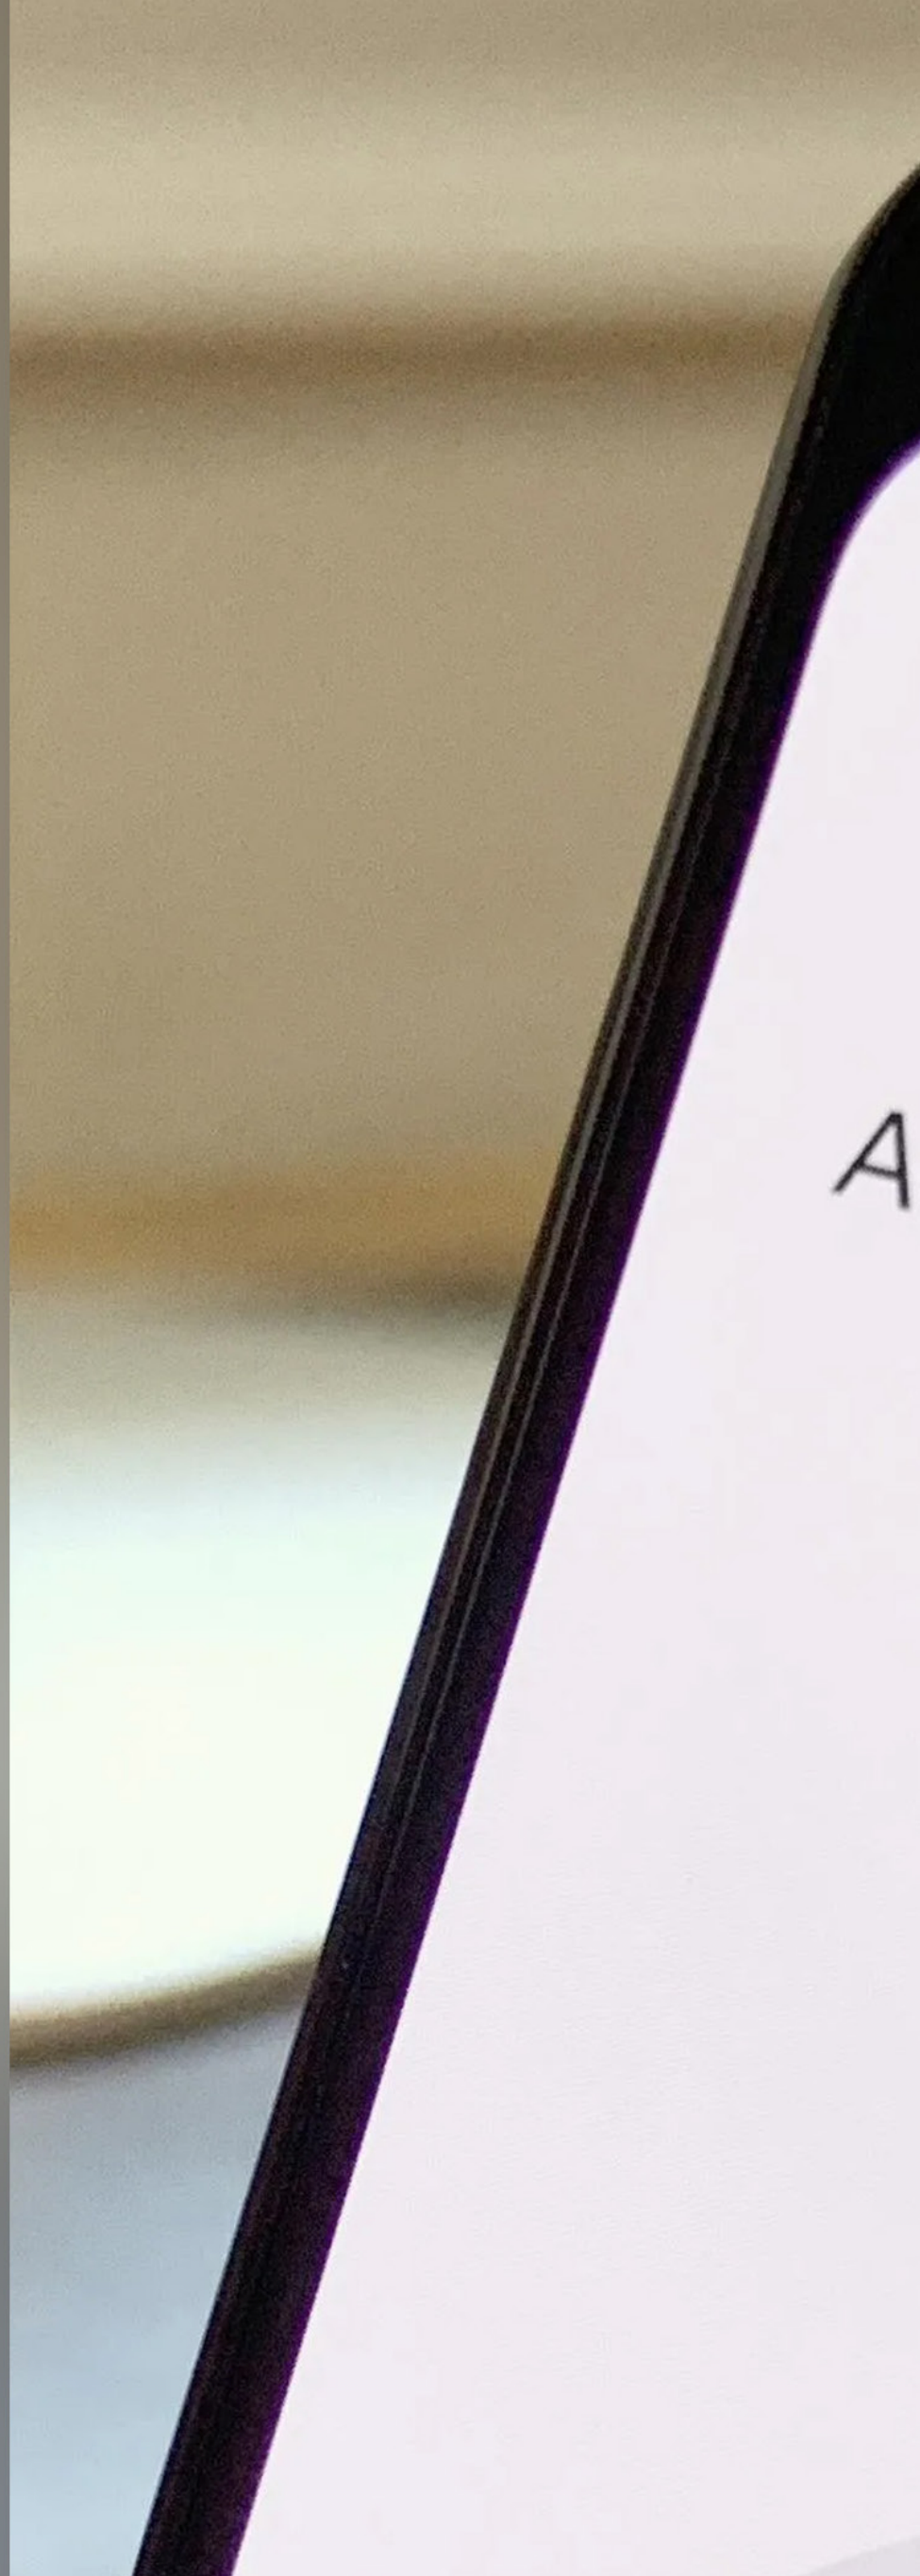

A new dual-band hotspot option has begun appearing in Android, giving users the ability to broadcast 2.4GHz and 5GHz connections simultaneously from a single device. The feature, first identified inside recent Android builds, introduces a more flexible hotspot configuration that aligns with how many modern routers manage multi-band support. Its arrival reflects Google's ongoing adjustments to Android's networking stack as the platform incorporates more advanced tethering and device-to-device connectivity tools.

10:00

LTE 100

dding Wifi point to your Wi-Fi network

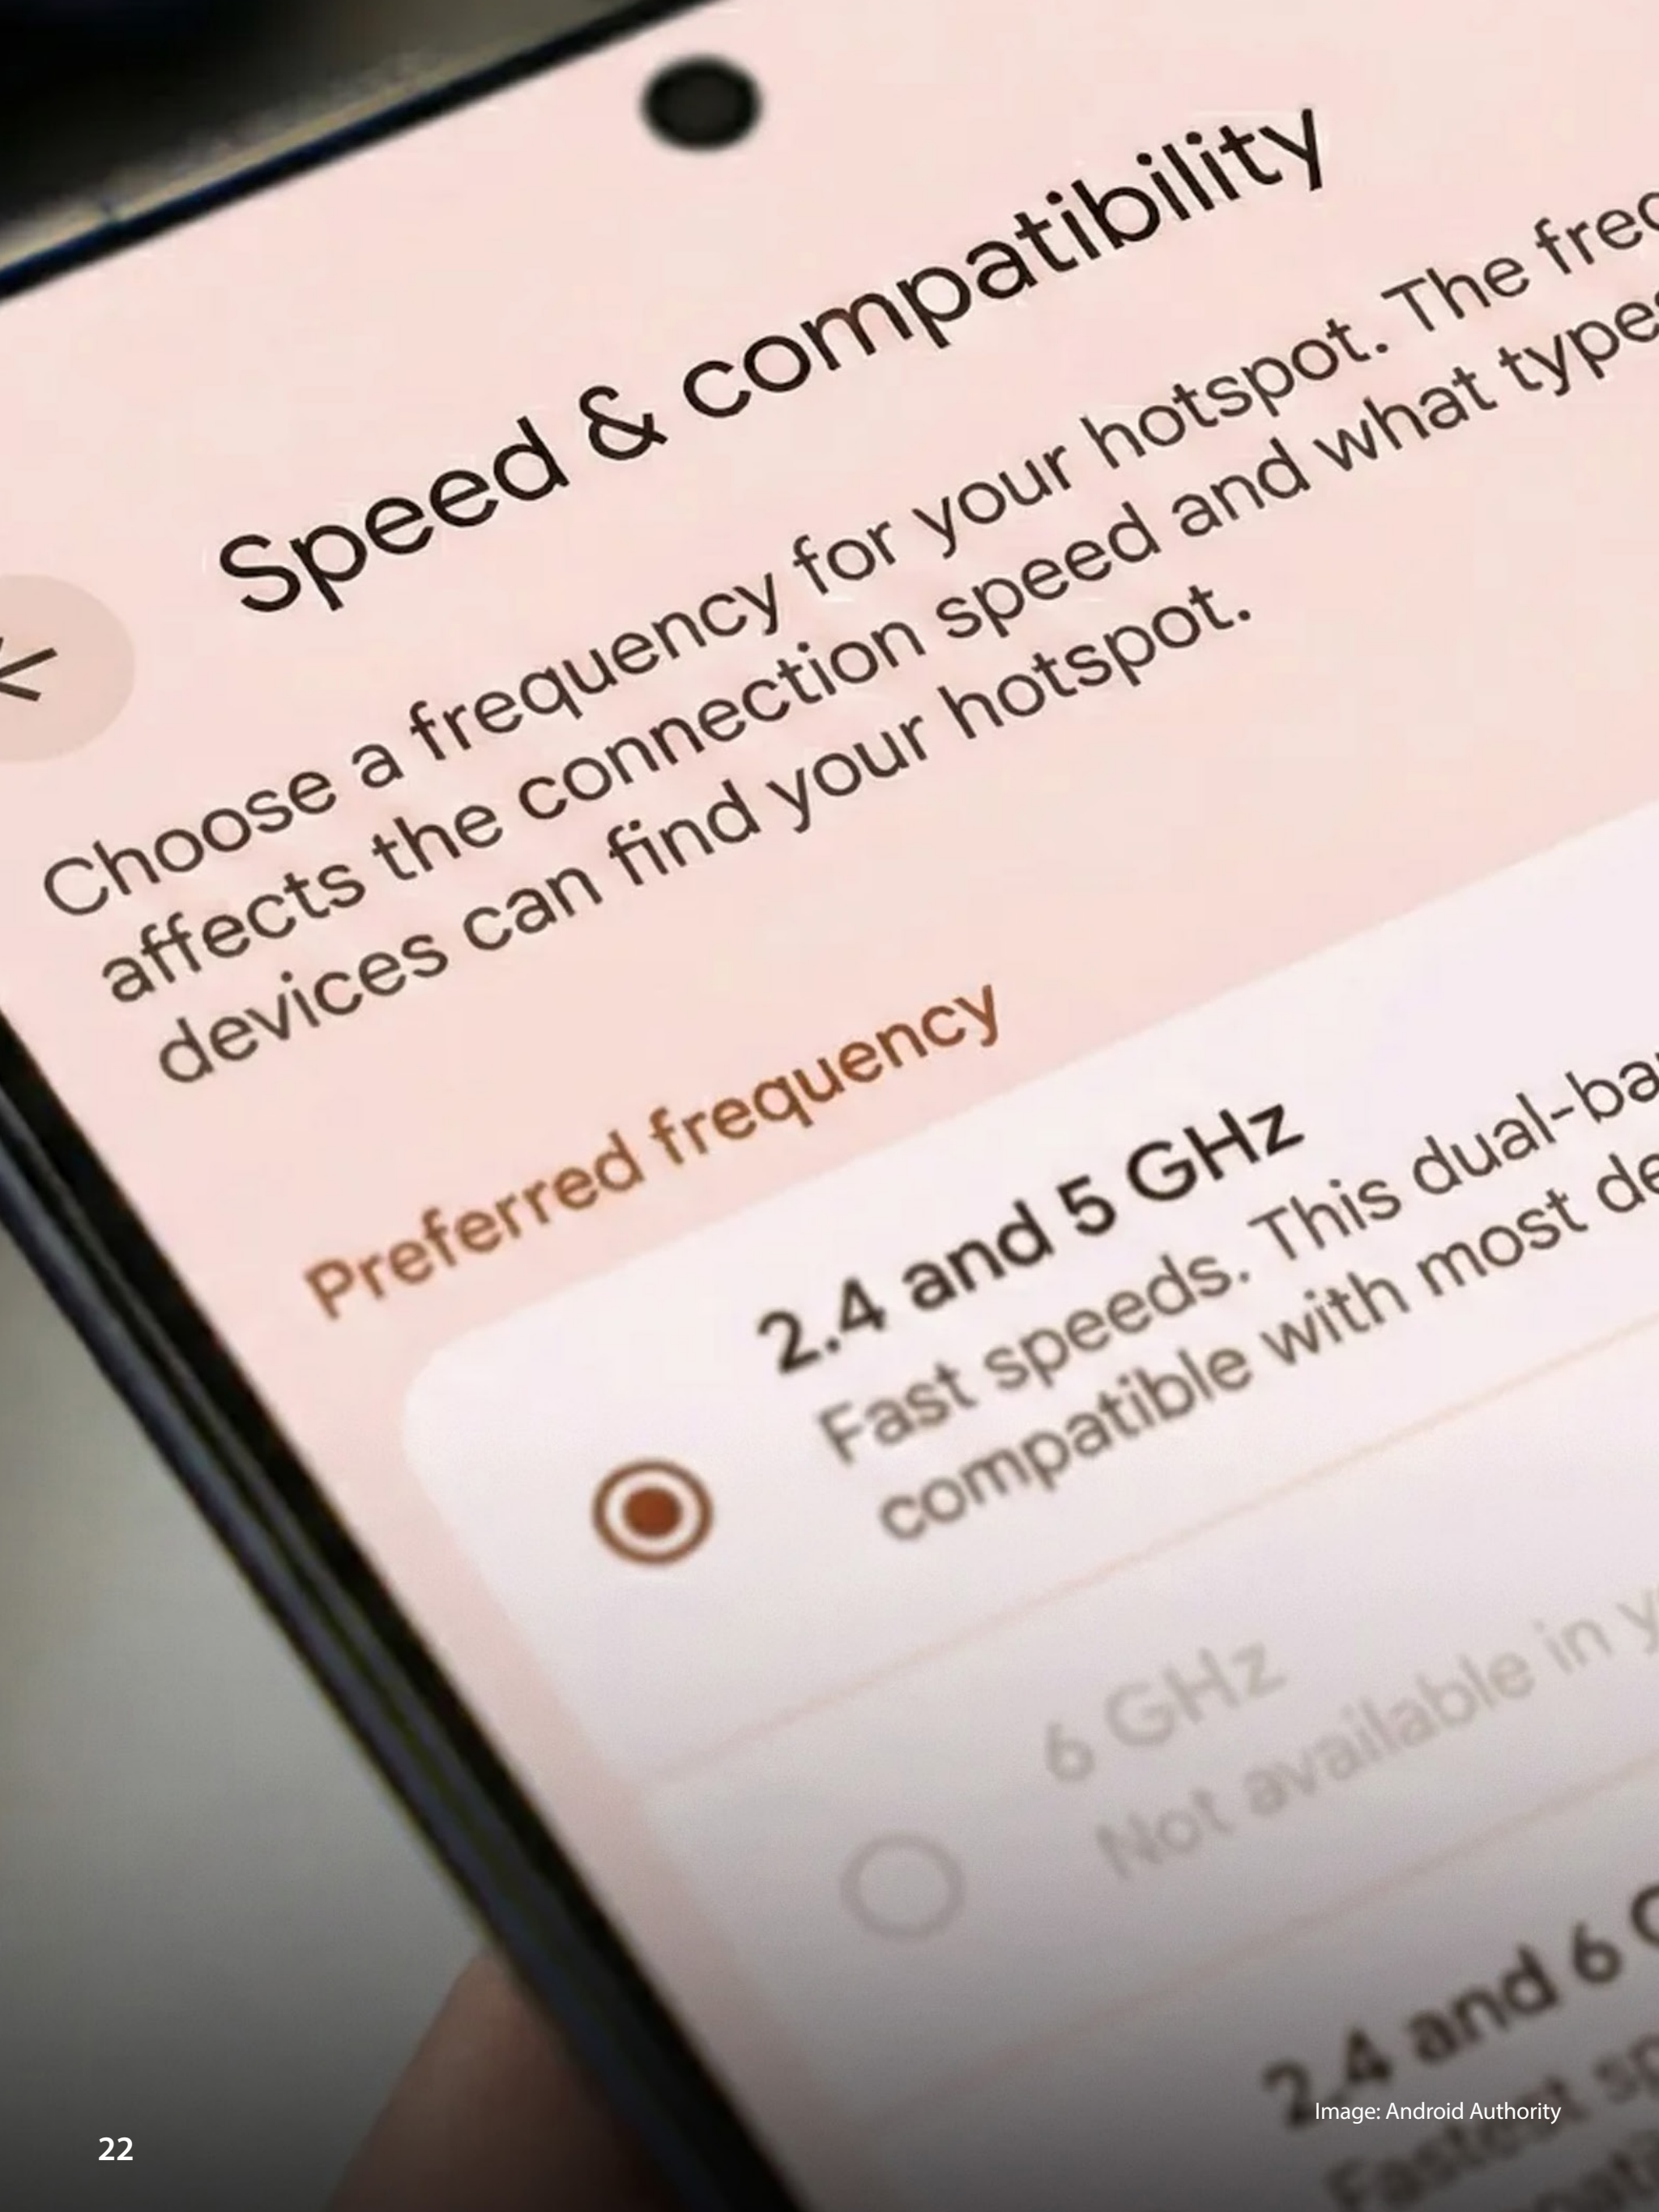

Image: Android Authority

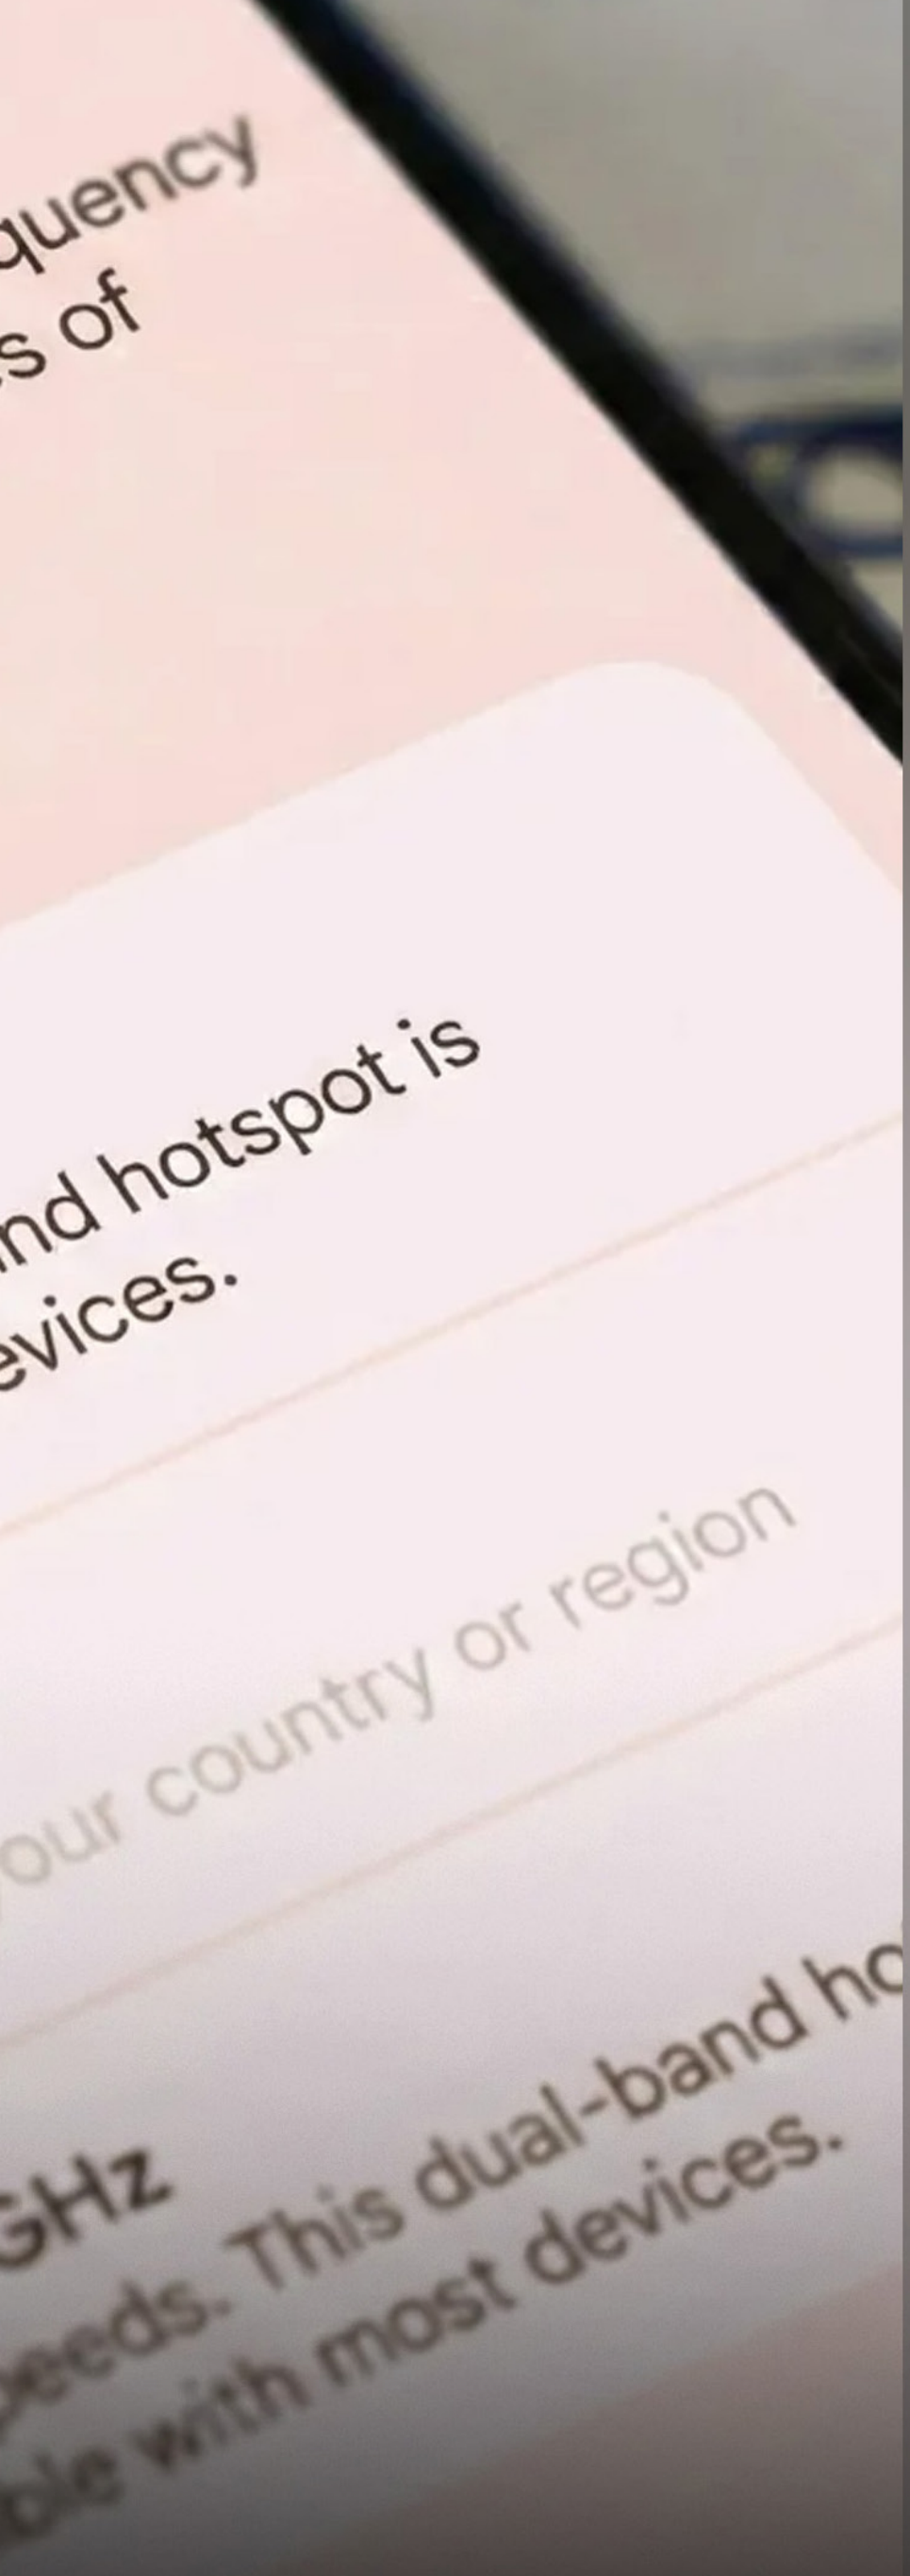

Android has long supported switching between 2.4GHz and 5GHz bands for mobile hotspots, but users previously had to choose one band at a time. The updated implementation provides a concurrent mode that allows older devices and low-power accessories to connect through the 2.4GHz channel while laptops, tablets and newer phones can use the faster 5GHz option. The change benefits households, travelers and remote workers who rely on mixed hardware and need a single tether that accommodates varying wireless capabilities.

Early builds showing the feature indicate that Google is restructuring the settings interface to surface the dual-band option in a straightforward configuration panel. The interface reveals separate toggles for each band, along with background handling that switches channels based on device compatibility. While the precise version rollout may differ across manufacturers, the functionality is tied to Google's core tethering framework rather than a vendor-specific enhancement, making it likely to appear across a range of upcoming Android releases.

THE NEW HOTSPOT MODE TARGETS MIXED-DEVICE CONNECTIVITY NEEDS

The addition of a dual-band hotspot addresses a long-standing issue for users whose devices span several generations of Wi-Fi hardware. Older accessories, including budget IoT devices and legacy laptops, often struggle with 5GHz connectivity, while newer devices gain the full benefit of higher throughput and lower interference on the 5GHz band. Android's updated hotspot mode allows both categories to operate under one broadcast,

reducing the need to manually switch bands when connecting different devices.

Developers who have examined the early code suggest that Android's networking service automatically manages how connections route through each band, depending on the capabilities the connecting device advertises. This approach ensures that older devices retain compatibility while newer ones leverage the performance improvements associated with higher-frequency channels. It also reduces friction when users share connections with groups, such as in hotel rooms, vehicles or temporary workspaces where multiple devices compete for a single mobile data link.

Users who often tether laptops or tablets may see the dual-band mode as a more stable alternative to switching manually between hotspot modes. The feature works within Android's existing tethering architecture and reflects a broader trend in which mobile operating systems adopt behaviors that resemble dedicated routers. As more devices rely on mobile hotspots in environments where traditional Wi-Fi access is limited or unavailable, the ability to support multiple bands concurrently becomes increasingly important.

MANUFACTURER ADOPTION WILL SHAPE HOW QUICKLY THE FEATURE REACHES USERS

Like many platform-level changes introduced in new Android releases, availability will depend partly on how device manufacturers implement the feature within their custom interfaces. Some vendors may integrate the dual-band toggle directly into their Wi-Fi settings panels,

while others may tie it to performance or battery-optimization menus. Manufacturers with strong emphasis on connectivity features, such as Samsung and Google's Pixel line, are expected to adopt the option early, though the specific rollout timing may differ across device families.

Carriers may also influence how the feature appears. Tethering behavior can vary depending on carrier policies, and certain hotspot configurations may be enabled or restricted based on data-plan requirements. The dual-band feature is implemented at the OS level, but final visibility in consumer menus may still reflect carrier guidelines. Developers who have tested pre-release builds note that the underlying functionality is present even when UI switches are hidden, indicating it may be broadly available in the system regardless of how it is surfaced.

As Android moves toward more unified networking components in upcoming releases, the dual-band hotspot mode is expected to become a standard part of the tethering feature set. This aligns with Google's effort to modernize core interfaces that handle local connectivity, device discovery and peer-to-peer interactions. The company has recently expanded features related to device-to-device transfers, local authorization prompts and cross-platform switching, suggesting that a more router-like approach to hotspot management is part of a larger roadmap.

POSITIONING WITHIN ANDROID'S EVOLVING CONNECTIVITY STRATEGY

The introduction of dual-band hotspot functionality comes as Android continues refining how devices handle wireless links that

bridge mobile data and local networks. Over recent releases, Android has integrated more granular controls over Wi-Fi power settings, roaming preferences, mesh participation and private network access. The dual-band update fits into this progression by mirroring capabilities that users expect from home networking equipment but applied to a handset-based environment.

For developers working with tethering and local-network applications, the feature also signals additional flexibility. Apps that rely on stable local connections—such as device-setup tools, local-transfer utilities or media-casting services—may benefit from environments where low-bandwidth accessories maintain compatibility while higher-throughput devices access faster connections. The ability to support both bands concurrently may reduce interference patterns and improve overall stability in dense wireless environments.

Future Android builds are likely to expand on this foundation, particularly as more devices ship with Wi-Fi 6 and Wi-Fi 7 support. These standards introduce additional spectrum and multi-link capabilities, which may influence how Android manages simultaneous channels in future hotspot modes. While the current implementation focuses on traditional 2.4GHz and 5GHz pairing, the groundwork may support more advanced configurations over time as new devices enter the market.

11:35 AM | 0.4KB/s

Wi-Fi

Dual band Wi-Fi speed boost

Primary network Wifi_Supremecy

AUXILIARY WI-FI NETWORK

Wifi_Supremacy_5G 5G
Tap to share password

Available auxiliary networks

Add network

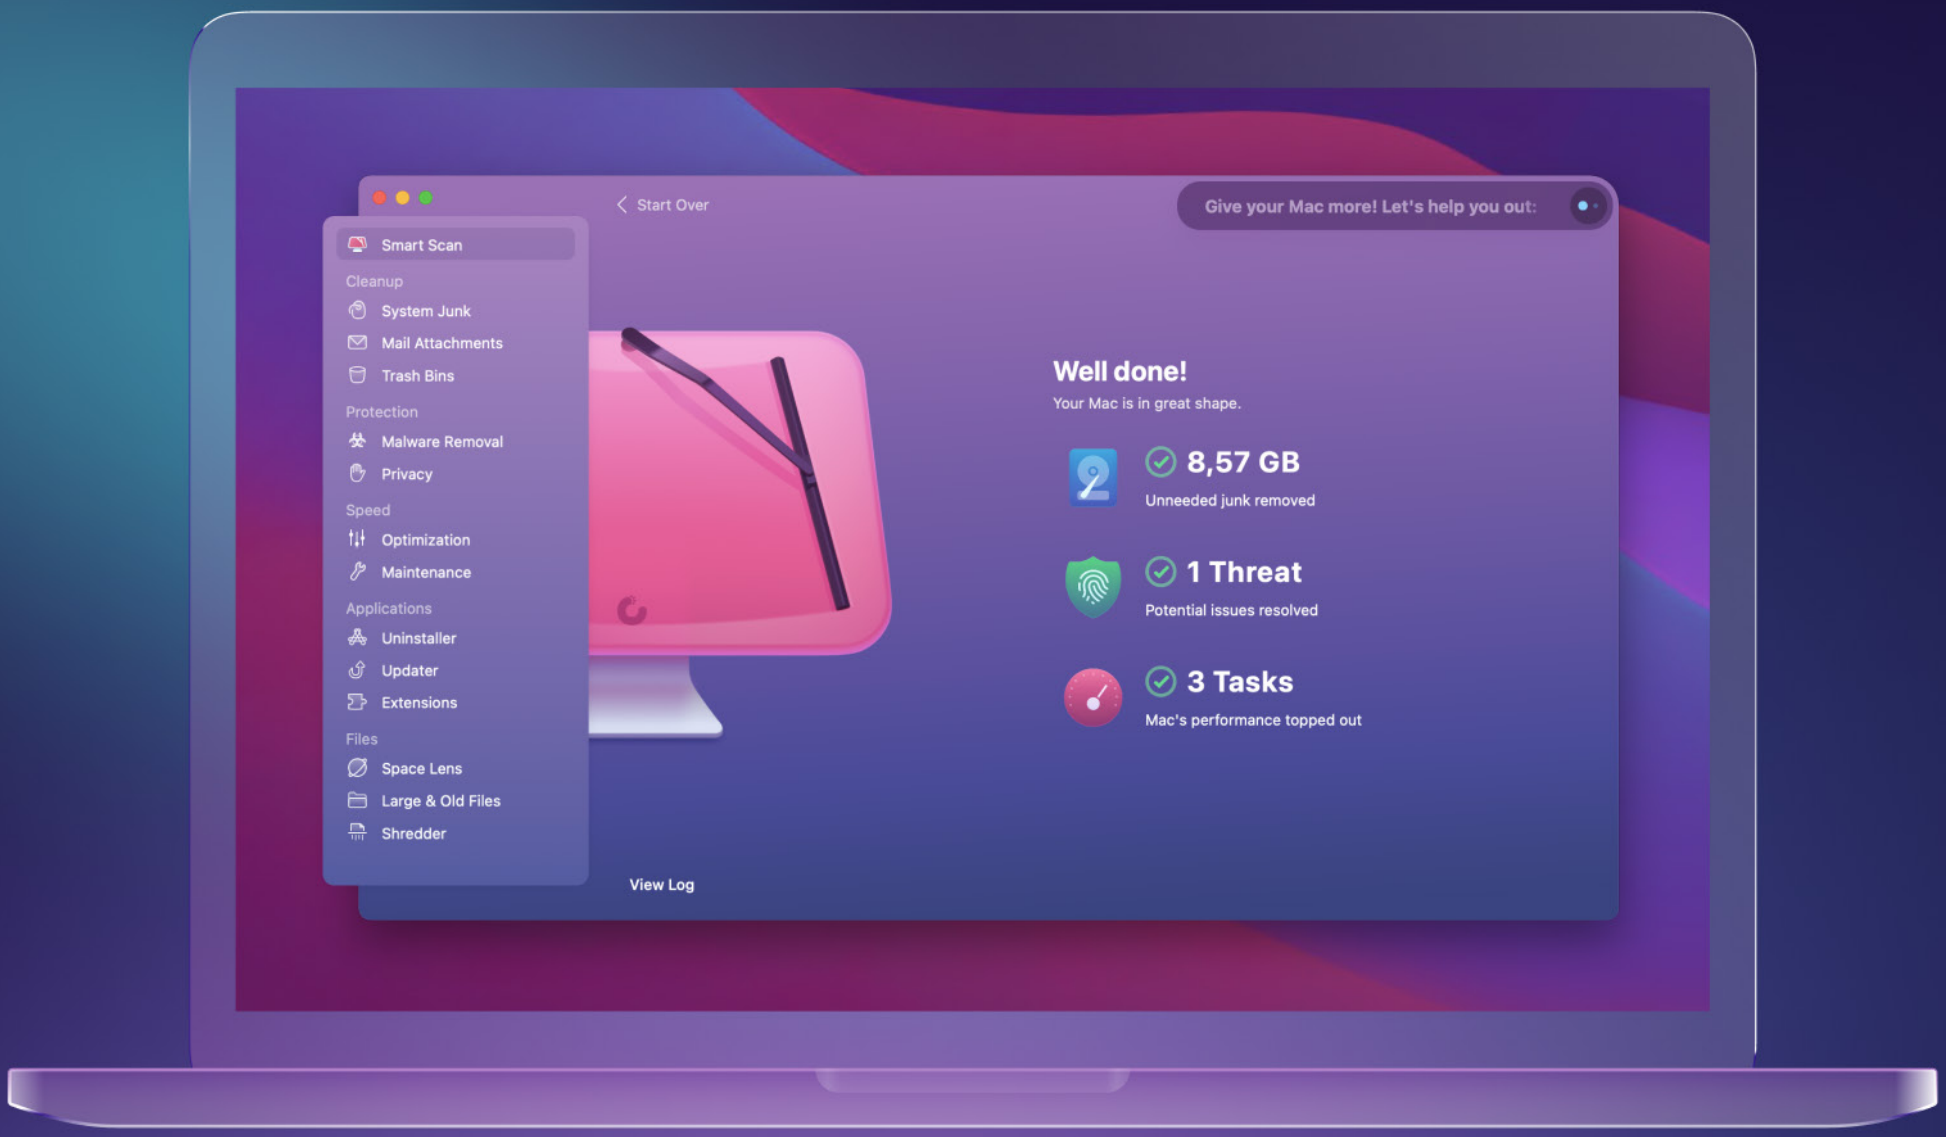

CleanMyMac X

Your Mac.
As good as new.

Learn More

Coca-Cola store

FREE SHIPPING
ON ALL COKE STORE ORDERS OVER $60!

www.cokestore.com


12:16
What's happening?!
everywhere
0 Following

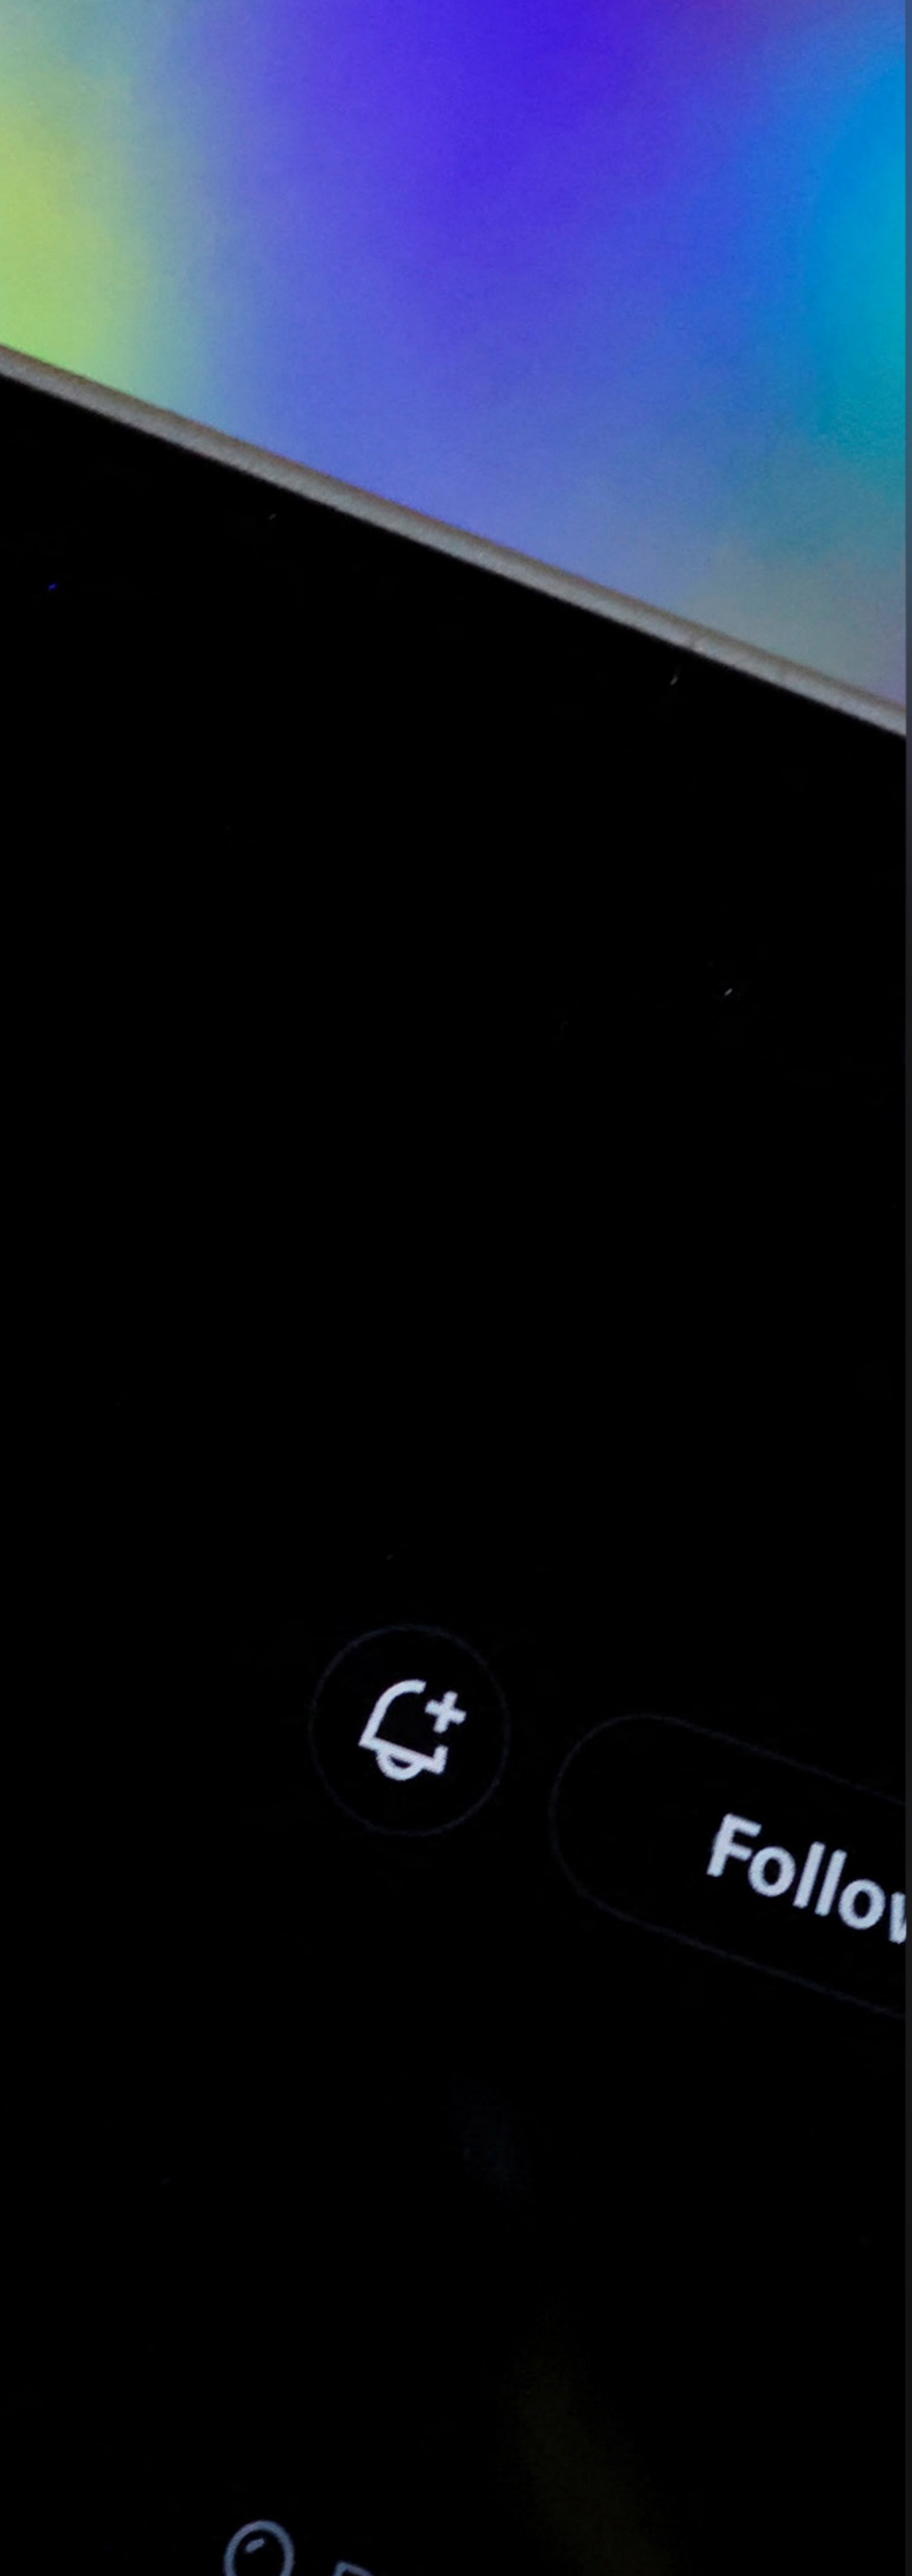

X INTRODUCES NEW “FOLLOWING FEED” RANKING SYSTEM POWERED BY GROK AI

X has launched a redesigned “Following Feed” that integrates ranking intelligence from its Grok model, marking a shift in how posts from accounts users follow are organized inside the platform’s primary timeline. The update, described in new reporting, adjusts the feed from a purely chronological view to one guided by relevance signals determined by X’s in-house AI system. The change reflects the company’s ongoing effort to weave Grok into more of the

platform's core functions, including search, recommendations and content surfacing.

The new feed keeps its familiar structure but no longer presents posts strictly in time order. Instead, Grok evaluates posts from followed accounts and adjusts their placement based on indicators tied to engagement patterns, topic alignment and user interaction history. X states that users can still access unranked chronological views, but the default experience now incorporates ranking designed to highlight posts the system identifies as most likely to be relevant to each individual user.

The rollout continues a broader strategy by X to expand Grok beyond its chatbot format and into foundational platform mechanics. Earlier updates incorporated Grok into trending-topic analysis and certain discovery surfaces. With the Following Feed now adopting similar underlying models, the platform is repositioning its internal AI as a structural component of how information flows through the network instead of a standalone assistant.

THE NEW FEED CHANGES HOW POSTS FROM FOLLOWED ACCOUNTS SURFACE

The introduction of Grok-based ranking alters a familiar part of the platform's user experience. The Following Feed has long been the place where users could view posts from accounts they follow in sequential order, distinct from the For You feed that surfaced algorithmic recommendations. With the new system, Grok uses signals derived from user interactions such as post viewing behavior, reply frequency, content preferences and historical topic interest to reorder posts from followed accounts.

Users who prefer chronological sequencing can still switch to it, though X now positions the Grok-ranked view as the updated default. People who have tested the feature noted that the feed now groups posts by perceived relevance, placing certain followed accounts more prominently while shifting others lower in the timeline. These changes reflect the model's ongoing attempt to evaluate which posts users are most likely to engage with or find useful based on previous behavior.

The company communicated that the ranking does not introduce additional sources beyond followed accounts. Instead, it adjusts the relative ordering inside that set. For creators and public figures who rely on follower timelines for distribution, this change may influence how posts are scheduled and how engagement patterns shift under the new ranking conditions. Observers expect additional adjustments as X collects more performance data and updates Grok's feed-ranking parameters.

The change also arrives during a period when major social platforms continue reworking how chronological views coexist with AI-driven recommendation layers. While X maintains a separate For You feed built heavily around Grok's analysis, the introduction of ranking into the Following Feed expands the AI's role even in streams historically reserved for direct user selection.

GROK'S EXPANDING ROLE IN X'S PLATFORM INFRASTRUCTURE

X has continued broadening Grok's integration into the platform's underlying systems.


23:07
Search
X
Blaze your glory!
823K RATINGS
4.8
AGE
17+
Years Old
CHART
No. 1
News
What's New
Version History
Version 10.0
We made improvements and squashed bugs so X is even better for you.
Preview
Blue
Blue looks good on you. Enjoy your verified account and blue checkmark.
See what's included in Blue and manage your settings
Creator Subscript
Subscribe to
Today

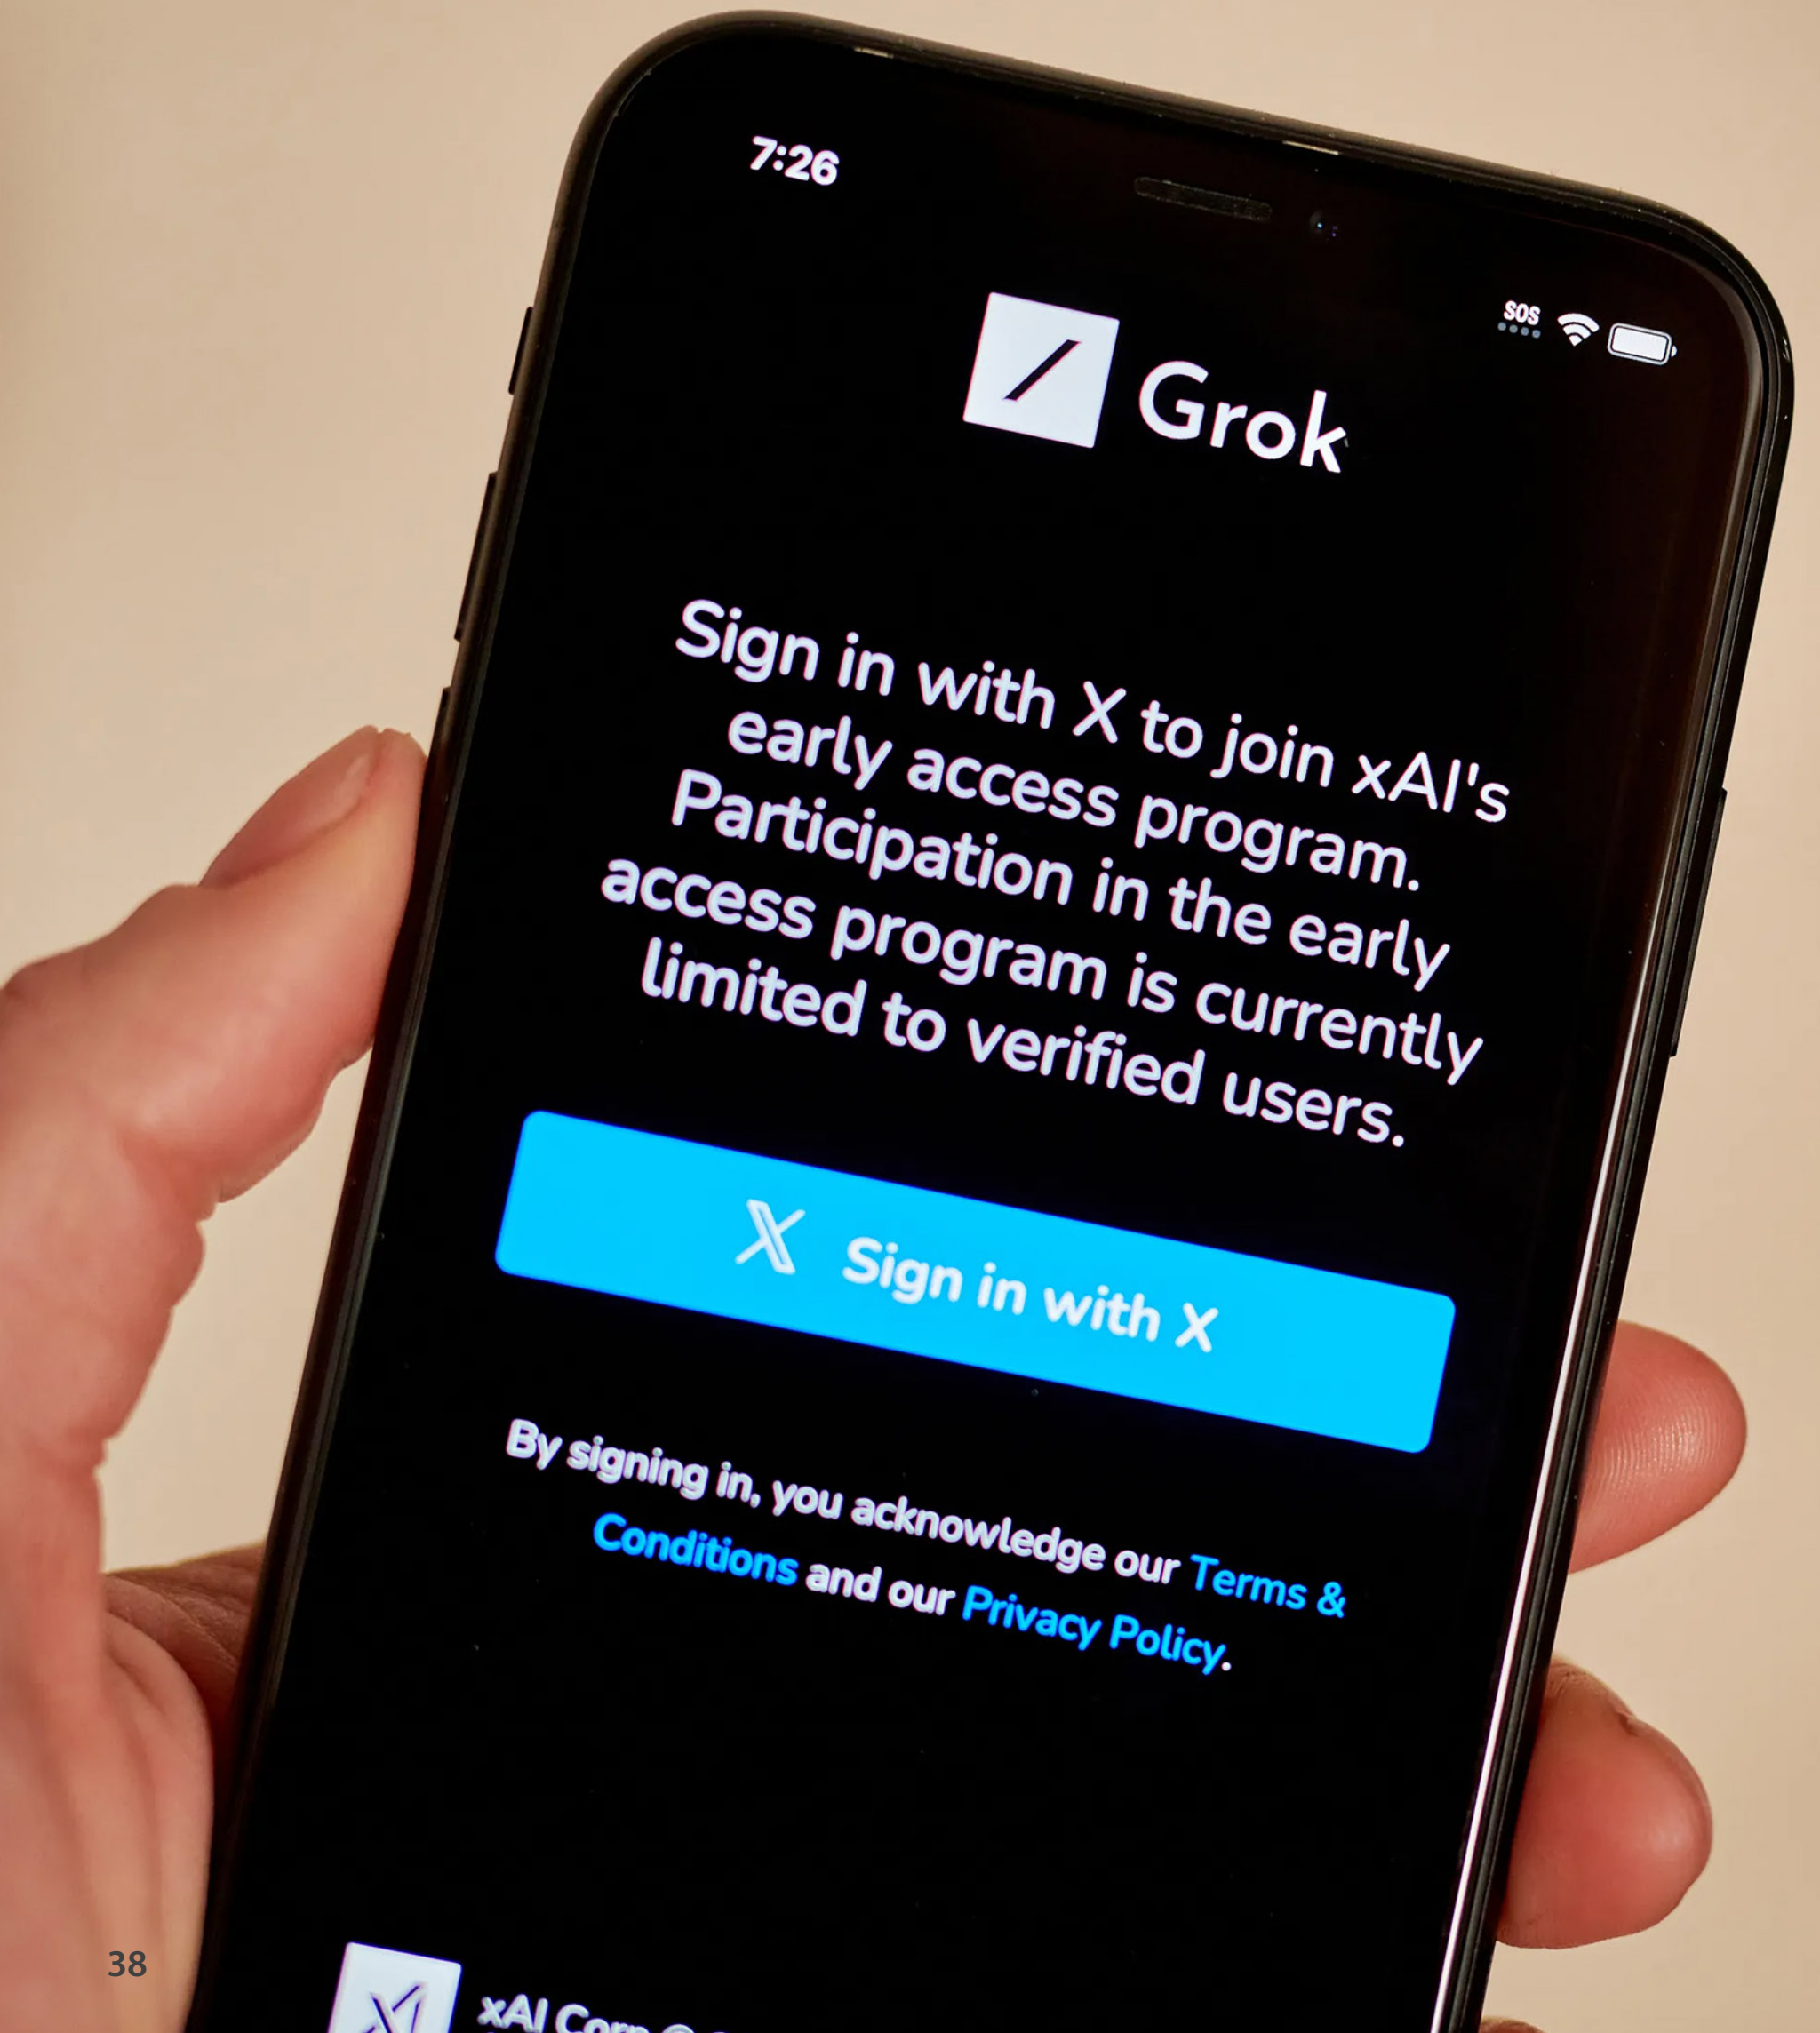
7:26
SOS
Grok
Sign in with X to join xAI's early access program. Participation in the early access program is currently limited to verified users.
Sign in with X
By signing in, you acknowledge our Terms & Conditions and our Privacy Policy.

What began as an AI assistant accessible through side panels and conversation threads has now evolved into a tool that influences search results, trending assessments and recommendation flows. The feed-ranking update places Grok at the center of a core user experience pathway that historically relied on simpler sequencing logic.

The growing integration reflects X's internal shift toward using Grok as a unifying engine for organizing platform content. The system's deployment across multiple surfaces gives the company a way to maintain consistent evaluation logic for posts while gathering performance data that feeds back into model refinement. X's engineering teams have gradually expanded the model's operational footprint to manage both conversational features and high-volume feed operations.

Developers familiar with X's infrastructure note that embedding Grok into the feed introduces new challenges related to latency, scalability and real-time evaluation. Ranking millions of posts across global time zones requires models that can process large content volumes quickly, adjust to emerging events and maintain responsiveness. As the updated feed becomes more widely used, X will continue tuning the ranking model to account for posting behavior, regional patterns and shifts in user activity.

The expansion of Grok's role inside X also intersects with the company's broader efforts to redefine how AI systems support social platforms. Grok's ranking layer does not replace manual feed controls but adds a new tier of interpretation that influences how users navigate their followed accounts. This tier shapes how often users encounter specific

content, how quickly posts surface in the timeline and how the platform identifies signals of user interest.

USER ADOPTION AND FUTURE ADJUSTMENTS EXPECTED AS ROLLOUT CONTINUES

As the updated Following Feed reaches more users, adoption patterns will shape how X refines Grok's ranking behavior. Early reactions vary depending on user preference for chronological feeds versus algorithmic ordering. People accustomed to sequential timelines may switch back to chronological mode, while others may find the ranked version surfaces posts they would otherwise miss. X is monitoring usage to evaluate how the system affects interaction rates and overall time spent inside the app.

Creators and public figures are also paying attention to how the update affects visibility. Changes in ranking behavior can influence how their posts circulate to followers, though the impact will depend on how Grok weighs engagement, topic relevance and posting frequency. As X continues adjusting ranking logic, the platform's communication to creators about best practices may evolve to align with the new system.

Future updates to the feed will depend on how Grok performs across different regions and content categories, and how the system interprets variations in user behavior over time. X has indicated that continued refinement is expected as the company gathers more data, and as additional features tied to Grok's conversational and multimodal capabilities enter the platform.

Gemini 2

The intelligent duplicate file finder

Learn more >

SHARP | Simply Better Living

SUPERSTEAM+™ BUILT-IN WALL OVEN

SSC2489DS

The Sharp® SuperSteam+ Built-In Wall Oven is the start of a cooking revolution. With Wi-Fi enabled IoT features, the innovations within this steam oven are a perfect match for modern cooking needs.

While regular steam only reaches 212°F, the SuperSteam+ oven can create superheated steam up to 485°F. Steam this hot can roast meats and caramelize sugars so your food can be brown and crispy on the outside, tender and juicy on the inside. With the Sharp SuperSteam+ Oven, you can grill without smoke, roast without drying, and get the roasty-toasty, tasty results you desire.

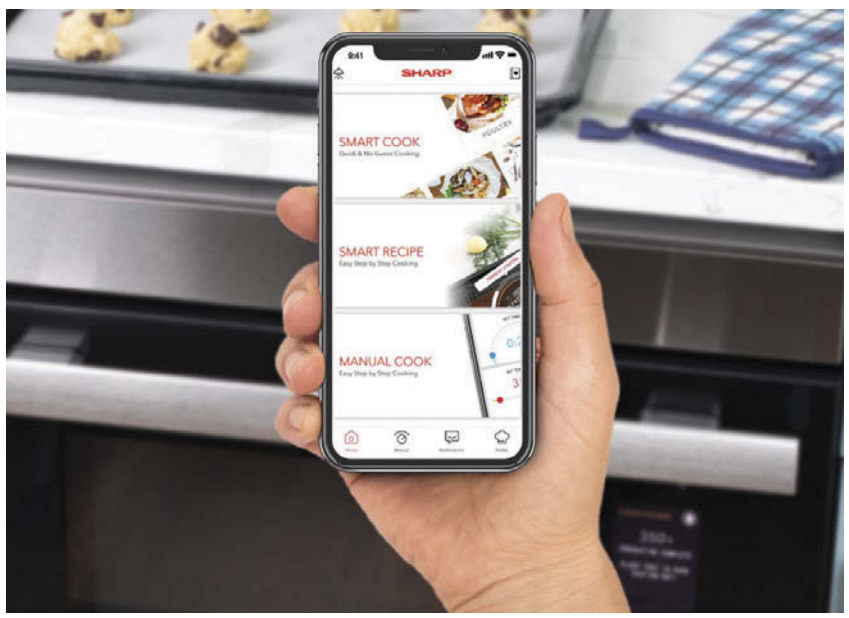

SEE FOR YOURSELF

Get started right away with built-in recipes and download the Sharp SuperSteam+ Oven app* to enable the smart features and access custom recipes powered by SideChef.

The new Sharp SuperSteam+ Built-In Wall Oven features Steam Bake for superior breads, and Water Bath for cheesecakes, custards and puddings.

www.sharpusa.com | simplybetterliving.sharpusa.com

*Mobile Application and Home Assistant Skill available upon commercial release. Product specifications and design are subject to change without notice. Internal capacity calculated by measuring maximum width, depth and height. Actual capacity for holding food is less.

GOOGLE EXPANDS CIRCLE TO SEARCH WITH NEW GEMINI-POWERED “CIRCLE TO GEMINI” ANALYSIS FEATURE

A new update is bringing Gemini-driven screen analysis to Android devices through an expanded version of Circle to Search, allowing users to draw on the screen and ask questions that trigger Gemini’s multimodal interpretation capabilities. The feature, appearing on some devices ahead of a wider rollout, represents the next phase of Google’s effort to merge real-time on-screen interactions with AI tools that can understand images, text and context without requiring app switching or manual copy-and-paste steps.

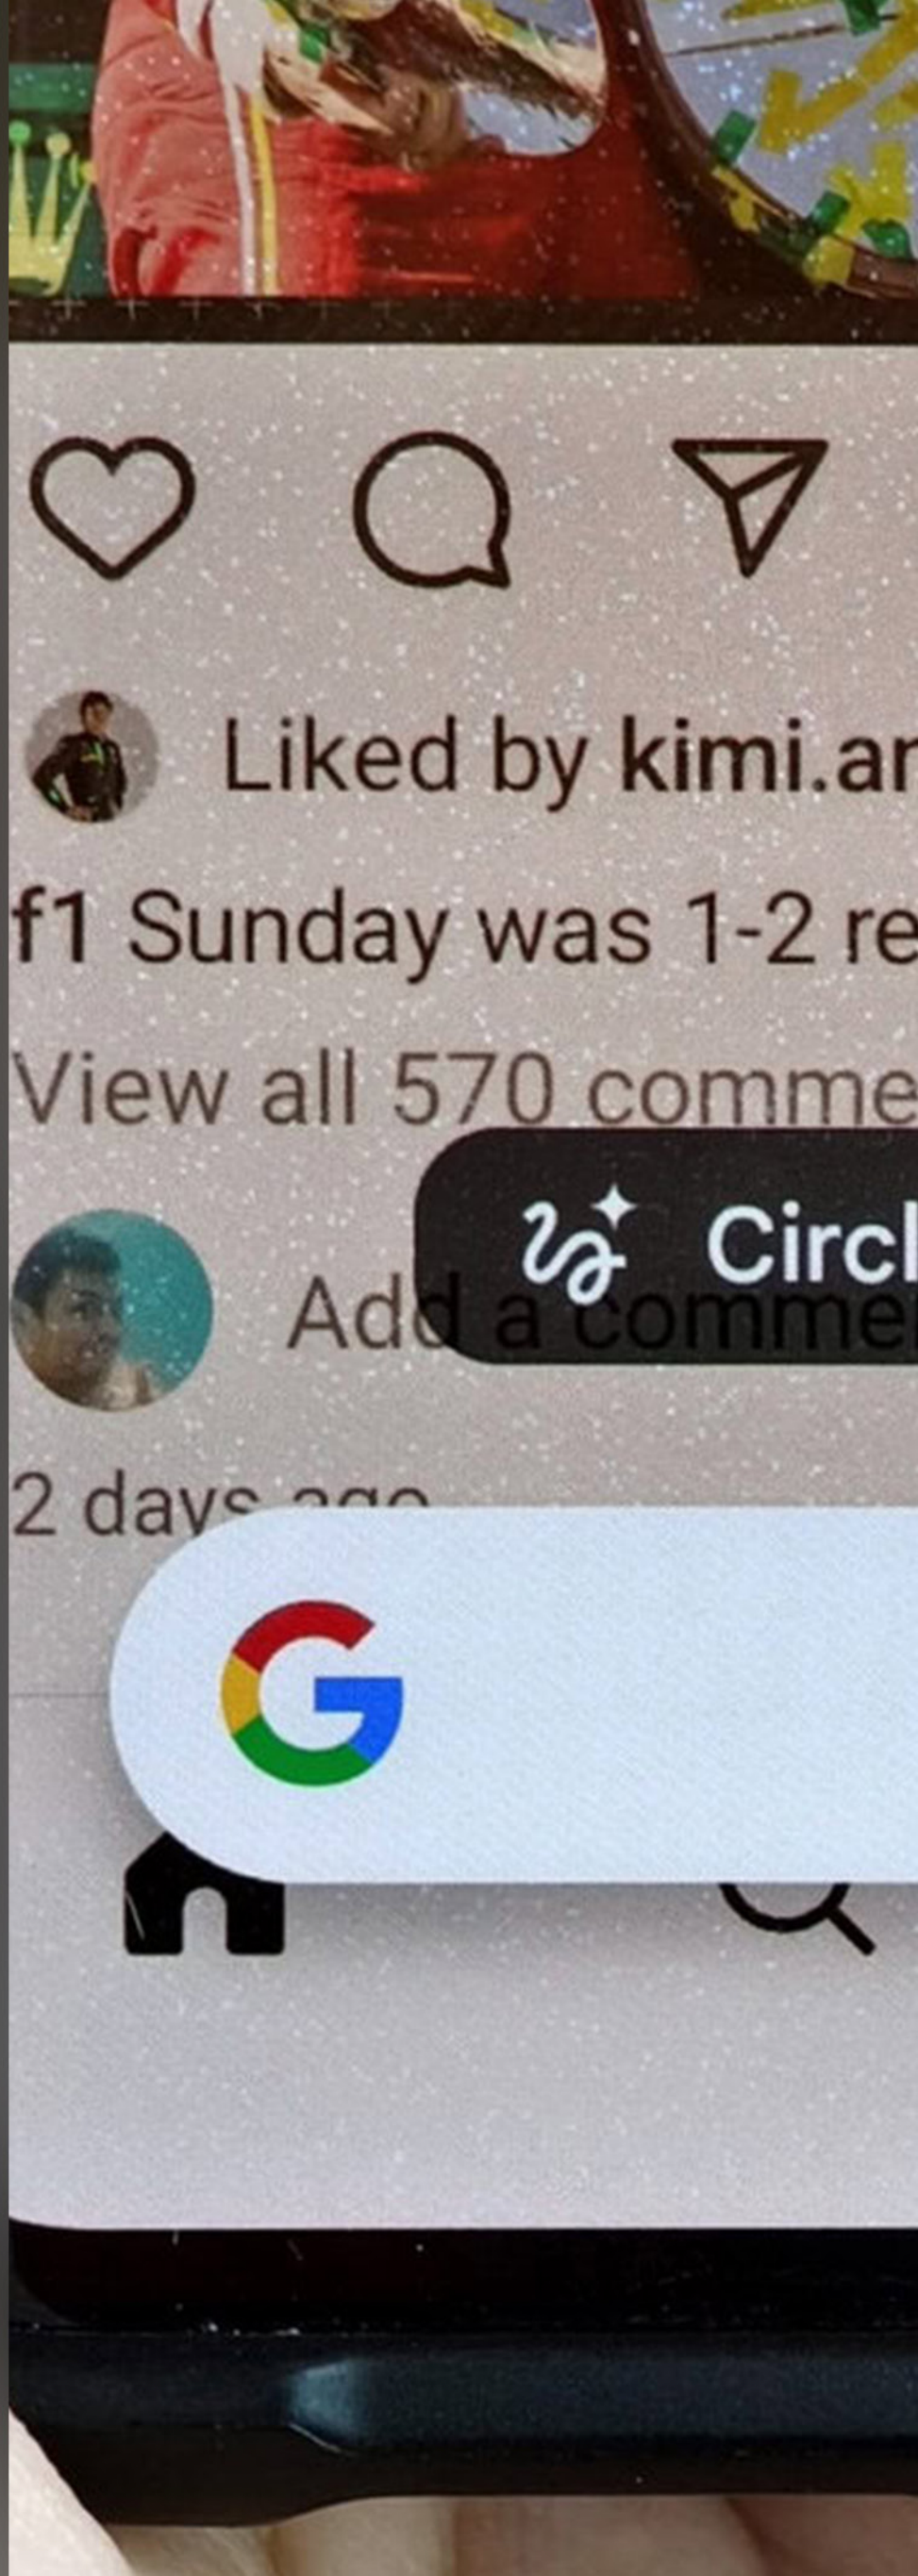

ntonelli and 587 405 others

emember for Ferrari 😉... more

nts

le or tap anywhere to search

nt...

Early users report that the new experience appears as "Circle to Gemini" inside the interface, replacing the familiar Circle to Search animation on supported devices. Instead of redirecting a query to Google Search, the gesture now opens a Gemini panel capable of offering broader interpretations, follow-up questions, and multimodal responses tied directly to whatever is displayed on the screen. The update is being delivered through server-side changes in combination with new beta app versions, with availability varying by region and device model.

GEMINI BECOMES THE NEW LAYER FOR ON-SCREEN ANALYSIS

The shift from Search to Gemini reflects Google's broader integration strategy across Android, where Gemini is gradually becoming the primary system assistant for devices running newer versions of the operating system. With Circle to Gemini, users can highlight content by circling, scribbling or tapping, and Gemini will analyze the selected area to provide contextual answers. These interactions include identifying objects, interpreting diagrams, evaluating text, parsing screenshots, or offering explanations tied to on-screen material.

Users testing the feature have noted that Gemini can handle more ambiguous or complex queries than Circle to Search, especially when the question requires additional reasoning beyond simple lookup. The panel includes space for follow-up prompts and offers contextual awareness of the selected portion of the screen, treating the image as part of a broader conversation. These capabilities extend Gemini's multimodal

12:45 100%

travelwithrachel.com

Google

TRAVEL WITH RACHEL

Add to your search

l backpack
n the desert

features that have been rolling out across Android since early releases of Google's system-wide assistant transition.

The system also incorporates direct links to continue conversations in the full Gemini app, allowing users to expand an on-screen question into a longer dialogue or request that Gemini perform actions related to the content. Because the new model supports image and text fusion, Circle to Gemini can interpret items such as charts, menus, maps, or visual instructions without requiring users to capture additional screenshots.

ROLLOUT REMAINS LIMITED AS GOOGLE TESTS DEVICE COMPATIBILITY

The early rollout is appearing on Pixel devices and select Android phones where Circle to Search has already been fully deployed. Several users have reported that the feature activates only after updating both the Google app and the Gemini app to recent beta versions, though activation still depends on server-side flags. Google has not yet published a timeline for a broad release, but testing patterns mirror previous staged expansions of Gemini features across Android.

Some devices display alternating behaviors, where Circle to Search remains active for certain queries while Gemini handles others. This suggests that Google is evaluating how different categories of questions should be routed, particularly in cases where standard search results remain more effective than conversational outputs. The hybrid state may continue until Google finalizes a unified flow for image-based queries across the platform.

Regional factors also affect availability. Gemini-based features roll out on a country-by-country

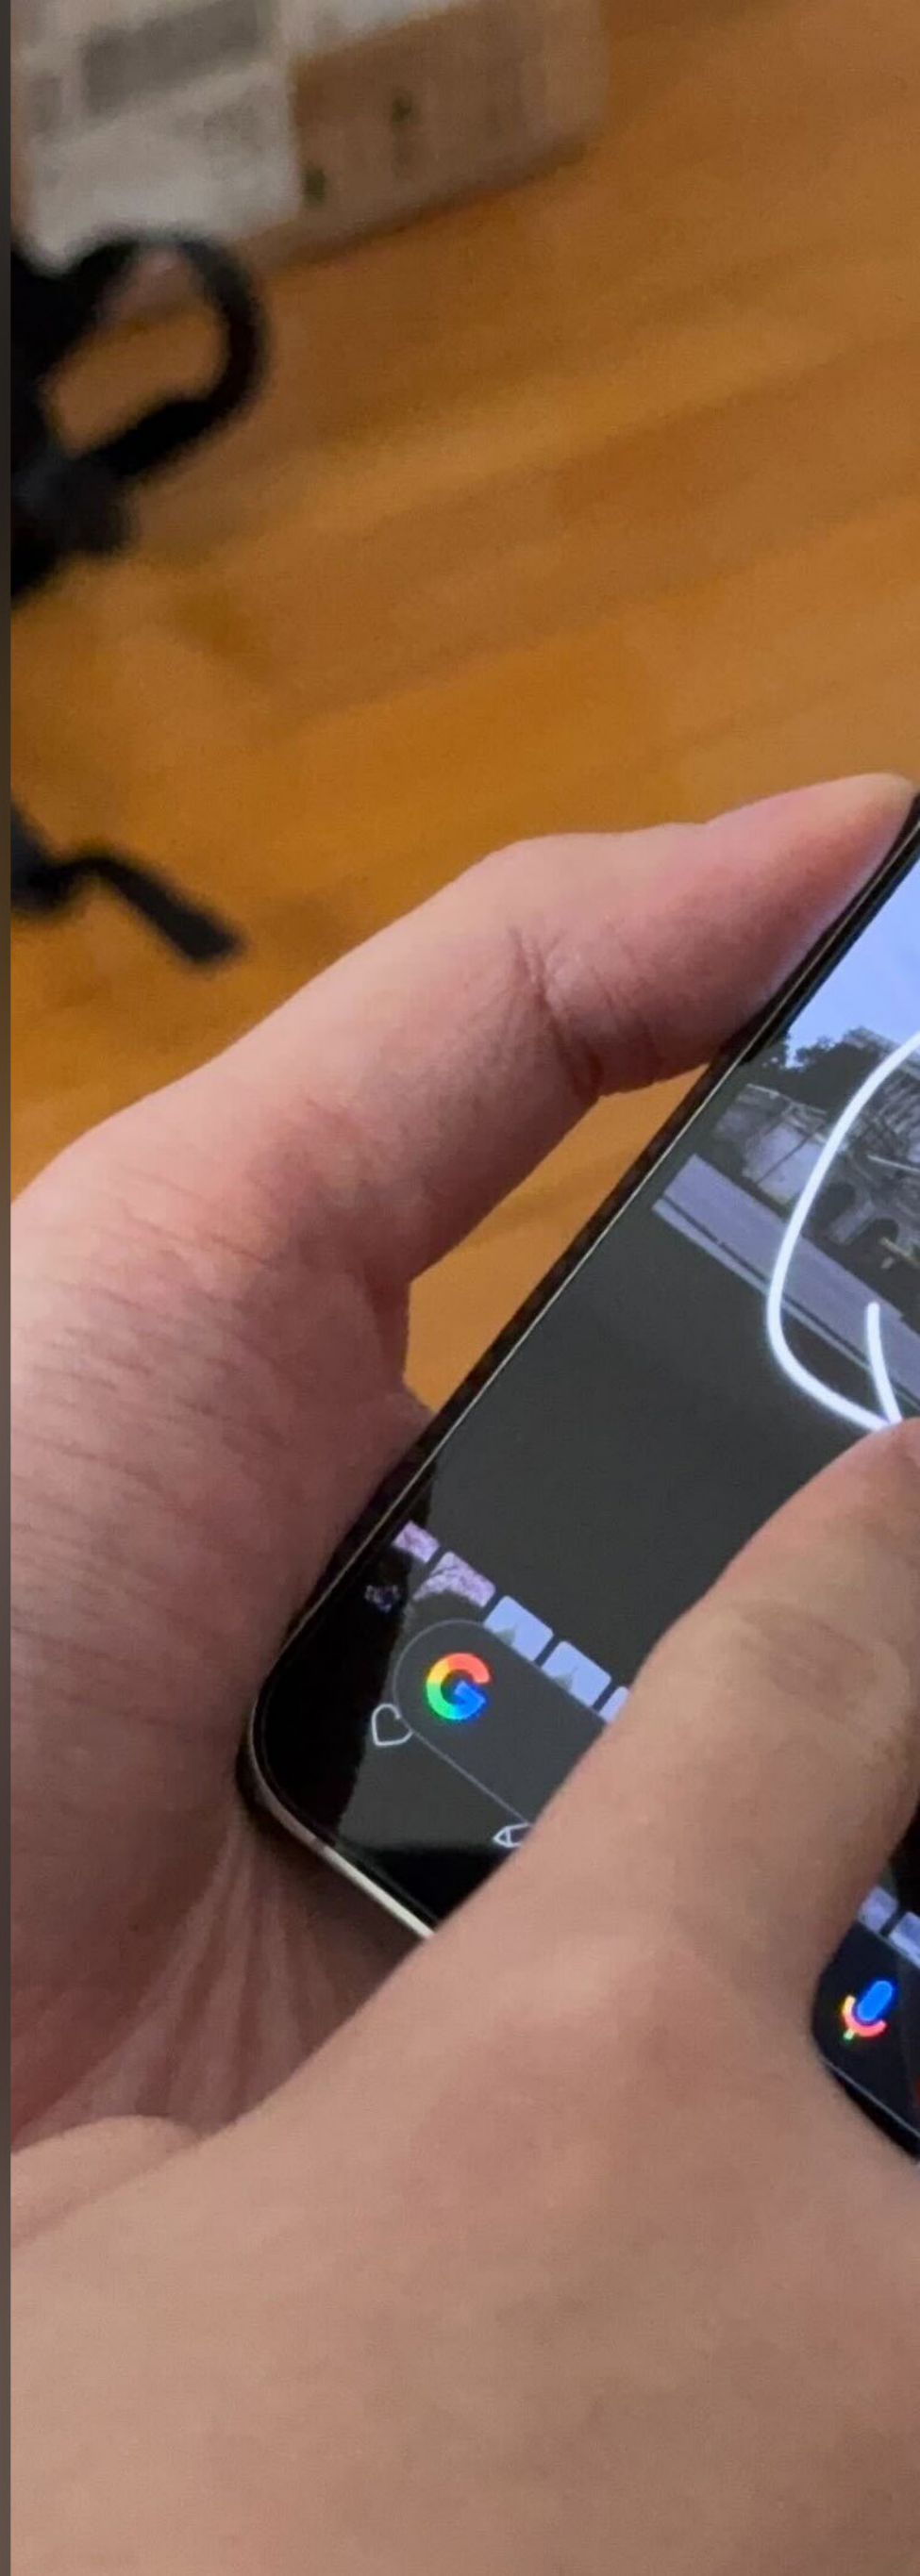


11:33
Google
98%


Ask Gemini

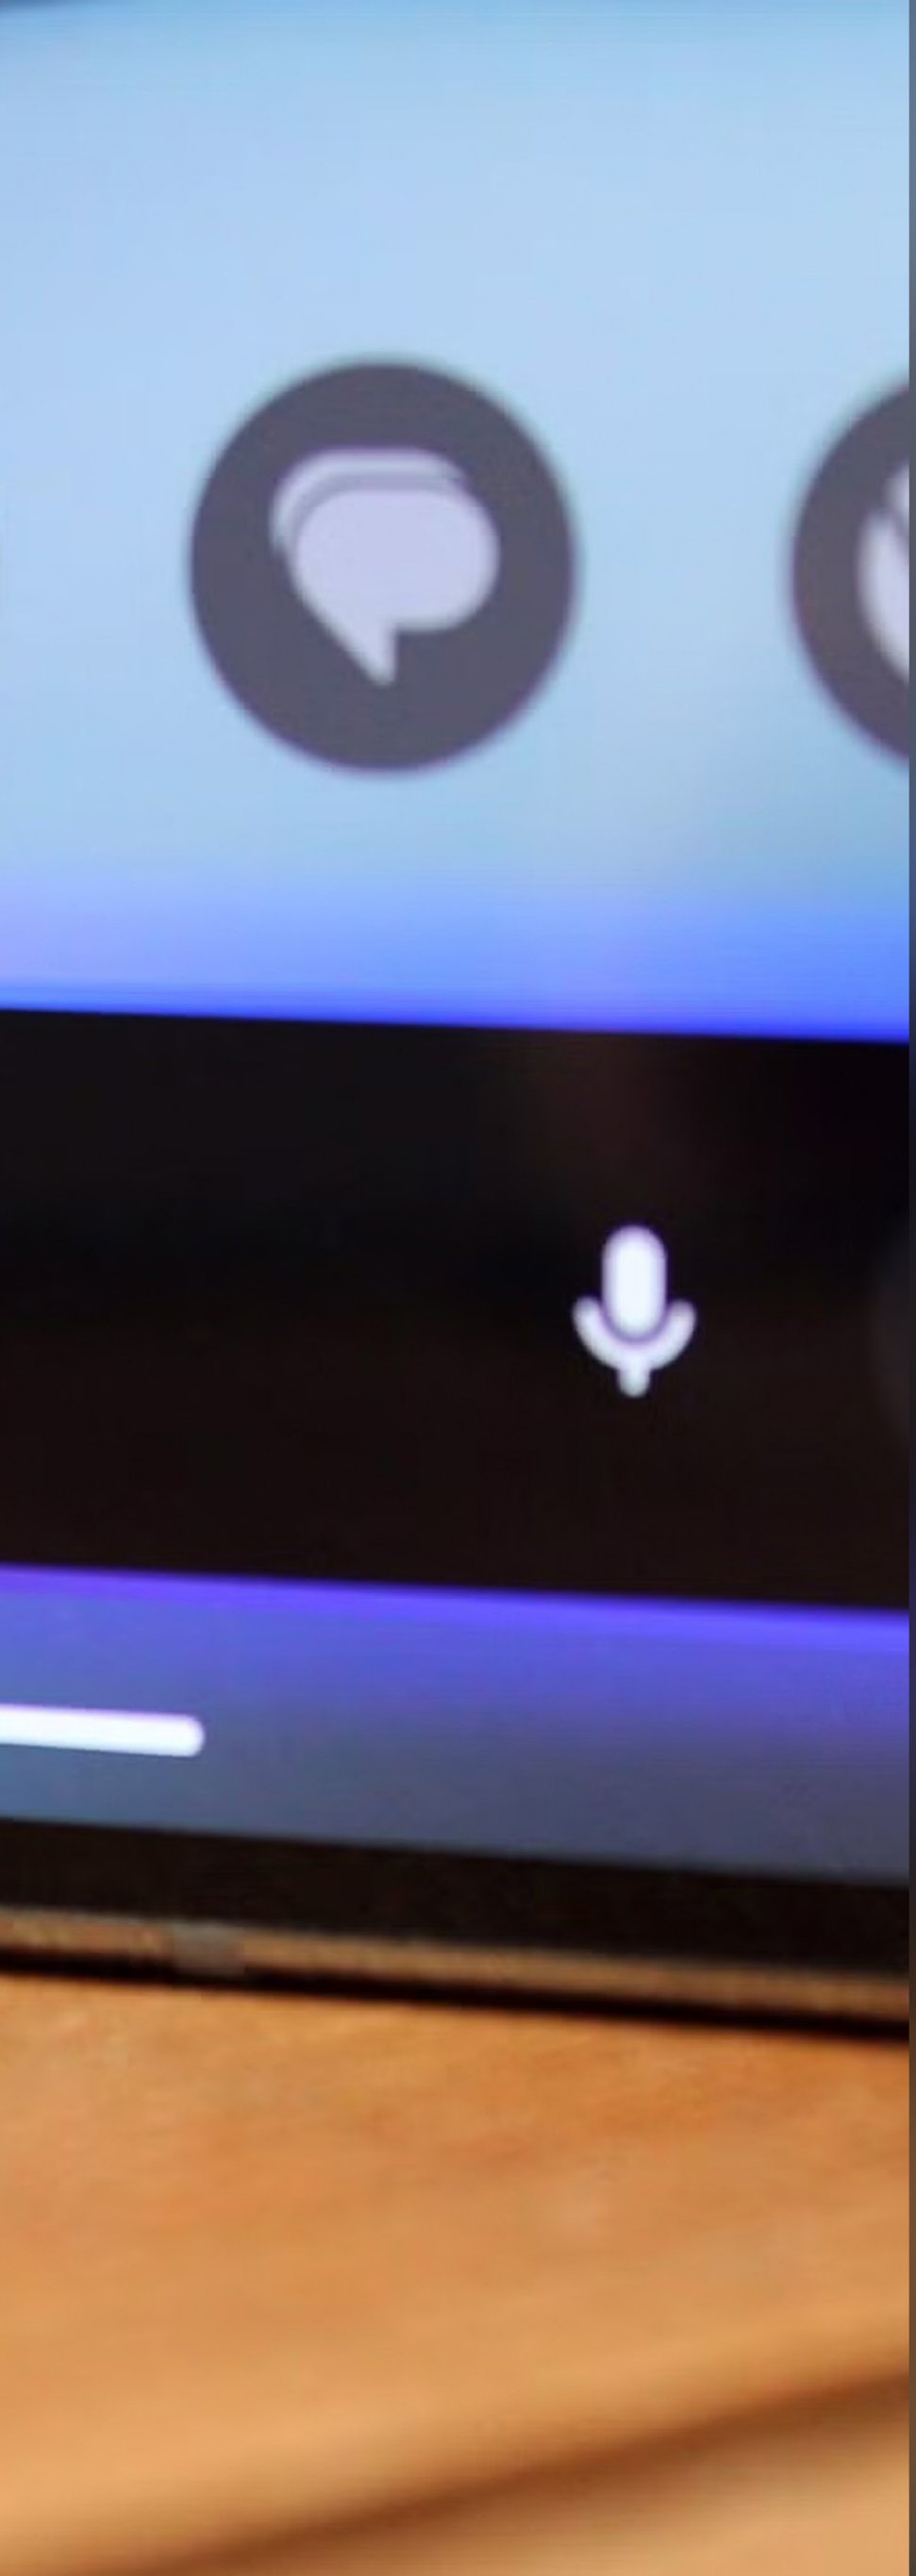

basis depending on language support, regulatory considerations and Google Play Services updates. The current rollout has been spotted primarily in English-language regions, with broader distribution expected once Google completes additional localization work.

POSITIONING WITHIN GOOGLE'S TRANSITION TO A GEMINI-NATIVE ANDROID

The appearance of Circle to Gemini fits into Google's ongoing transition away from Assistant toward a Gemini-native Android experience. Over recent months, Gemini has been integrated into system-level features such as voice interactions, document interpretation, device setup guidance and multimodal query handling. Circle to Gemini expands that footprint by positioning Gemini as the default handler of visual analysis and contextual reasoning triggered directly from the screen.

This deeper tie to on-screen content aligns with Google's broader goal of making Gemini the central intelligence layer of Android, capable of interpreting visual, textual and environmental signals without requiring users to move between multiple apps. Circle to Gemini serves as a key part of that shift: it allows users to select any element on the screen and ask questions in place, turning casual interactions into multimodal prompts.

For developers and manufacturers, the update reflects a growing emphasis on system-wide AI hooks across the platform. OEMs that previously integrated search-based analysis tools may need to adapt interfaces to accommodate Gemini's larger response panels and conversational workflows. These integrations also signal that future Android builds will

treat multimodal AI not as an optional add-on but as a baseline expectation for device interaction.

WHAT THIS MEANS FOR USERS AS THE FEATURE CONTINUES TO EXPAND

As Circle to Gemini becomes available on more devices, users will gain access to a tool capable of interpreting complex visual content with minimal friction. Tasks that previously required switching between apps—such as identifying an object, comparing products, extracting text, or interpreting diagrams—can now be handled within a single gesture. The system's conversational design allows follow-up queries that build directly on the initial selection, making it more adaptable than traditional image recognition tools.

The expansion also signals how Google is positioning its mobile AI efforts in relation to competing platforms. By embedding Gemini directly into gesture-driven system features, Google is extending its multimodal capabilities beyond chat interfaces and into everyday interactions across Android. The staged rollout suggests that additional features may attach to Circle to Gemini over time, particularly as Gemini's models grow more capable of handling layered, context-rich tasks anchored to what users see on their devices.

As testing continues across beta builds and server-side updates, more Android users should begin seeing Circle to Gemini appear as part of their standard interface. The company is expected to refine the feature as feedback arrives from early adopters, adjusting how queries are routed and how the model interprets different classes of visual content.


Gemini

ClearVPN

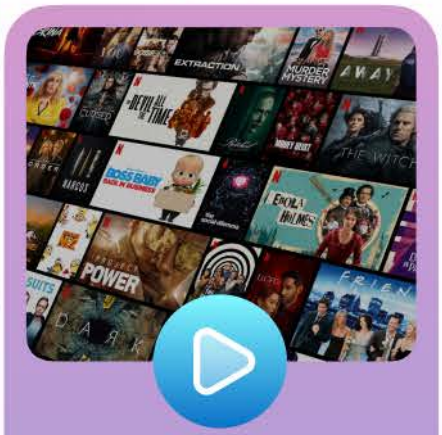

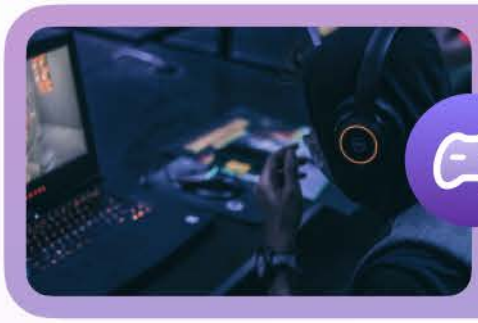

The only VPN you'll ever need

Get ClearVPN

Lightinthebox.com

One World, One Market

Offering high quality products in best price for global customers

STEAM BECOMES A MAJOR REVENUE CHANNEL FOR PLAYSTATION

Sony's presence on Steam has grown from a secondary experiment into a substantial piece of the company's software business, supported by data that outlines more than a billion dollars in gross revenue from PlayStation titles released on Valve's PC storefront. New figures compiled by Alinea Analytics estimate that Sony's PlayStation Studios catalog on Steam has now generated more than $1.5 billion in total sales,

teampowered.com

STEAM

with Sony's net share sitting just under $1.2 billion after storefront fees. The findings point to Steam becoming a consistent revenue channel for Sony, despite the company's historical dependence on exclusive console releases.

The revenue scale reflects a broader shift in Sony's strategy as it positions a portion of its first-party catalog for audiences beyond the PlayStation hardware ecosystem. For years, PC releases arrived well after the original console launch, framed as a way to reach niche players or extend a franchise's lifecycle. The new analysis suggests that Sony's participation in the PC market has expanded far beyond that initial rationale, signaling a structural change in software distribution rather than a short-term experiment. As more titles arrive on Steam—and as some franchises gain traction with PC-native audiences—the platform is now influencing Sony's longer-term decisions around scheduling, pricing, and multi-platform support.

HELLDIVERS 2 AND THE RISE OF REPEATABLE PC PERFORMANCE

The roster of PlayStation Studios games available on Steam spans action, adventure, open-world, and cooperative genres, but one title stands out as a major driver of Sony's PC revenue: Helldivers 2. According to Alinea Analytics, the game is estimated to have sold more than 12 million copies on Steam, generating roughly $400 million in gross sales on that platform alone. The game's cooperative structure, ongoing updates, and strong performance on PC contributed to sustained player engagement, marking it as Sony's most commercially successful PC release to date.


PlayStation

Beyond Helldivers 2, several earlier entries in Sony's PC catalog established a baseline that proved the audience potential. Horizon Zero Dawn, one of Sony's earliest major PC ports, reached more than four million Steam sales and generated well over $100 million in revenue. God of War (2018) achieved similar traction, with estimates placing its Steam performance above four million units and strong revenue totals that positioned it among the top-performing Sony titles available on PC. These early successes created a foundation that helped justify subsequent releases and likely influenced internal resource allocation for PC optimization, porting, and update support.

The data also illustrates how Steam's tiered fee structure benefits high-volume publishers. Valve's revenue share begins at 30 percent but falls to 25 percent after a title reaches $10 million in sales and then to 20 percent after surpassing $50 million. Because several Sony titles exceed those thresholds, their net revenue share increases over time. The analysis suggests that Sony's retention rates improve substantially once a title crosses major revenue milestones, reinforcing the financial appeal of Steam for large-scale publishers with established franchises.

However, Alinea's research indicates shifts in demand patterns over time. Follow-up titles in existing franchises have not always matched the velocity of prior releases. For instance, first-year sales for the PC version of God of War's sequel appear lower than the earlier installment over a similar release window. This suggests that while PC audiences remain strong, performance is increasingly sensitive to factors such as pricing, release timing relative to the console version,

and how effectively a title aligns with PC player expectations around technical options, optimization, and ongoing support.

WHAT STEAM'S GROWING REVENUE MEANS FOR SONY AND THE PC MARKET

The scale of Sony's PC revenue introduces new considerations for its broader software strategy. PlayStation's traditional identity has centered on hardware exclusivity and platform-defining titles, but the emerging revenue from Steam provides an additional axis for long-term planning. While Sony continues to position PlayStation as its core ecosystem, the success of its PC releases shows that many players engage with Sony franchises without owning a console. This dual-market presence increases visibility, strengthens franchise recognition, and diversifies revenue sources.

Steam's role as a distribution platform reinforces this shift. With a large active user base and consistent engagement across genres, Steam offers Sony a dependable market for titles that suit PC audiences. The storefront's reach allows Sony to target players who prefer PC hardware, broader customization, and community-driven environments. As cross-platform gaming expands and game engines evolve to support multi-platform output more efficiently, Sony's participation in PC gaming becomes part of a long-term approach that complements rather than disrupts its console strategy.


LIBERATOR

The relationship also influences how Sony approaches live-service development. Cooperative, online-oriented games have historically performed well on PC due to strong modding cultures, replayability, and consistent update engagement. With multiple live-service projects under development across PlayStation Studios, the performance of Helldivers 2 on Steam provides a measurable indicator of the potential reach for similar titles. While console exclusivity continues to define some of Sony's biggest releases, analysts note that multiplayer-driven titles may increasingly appear on PC earlier—possibly alongside the PlayStation launch—to maximize their network effects.

For Valve, Sony's expanded presence underscores Steam's importance as a platform capable of supporting major AAA publishers from outside the traditional PC-first ecosystem. The partnership highlights how Steam continues to serve as a critical distribution channel even as alternative storefronts, subscriptions, and cloud-gaming models evolve. Sony's participation contributes to Steam's broader catalog diversity and creates additional incentive for players to remain within the platform's ecosystem.

LONG-TERM OUTLOOK FOR PLAYSTATION'S PC STRATEGY

As Sony continues evaluating the performance of its PC releases, several trends appear to shape the long-term trajectory. The first is timing. While earlier PC ports often arrived years after the console launch, newer releases have seen narrower gaps between platforms. Shorter delays may continue as Sony tests how


GOD
DIGITAL DEL


OF WAR
™
UXE EDITION

simultaneous or near-simultaneous launches affect overall revenue and player engagement. This approach aligns with market behavior in which players increasingly expect parity across platforms, especially for titles that rely on online participation.

Another trend involves production planning. Porting to PC requires technical adjustments, optimization work, and ongoing patch support. As PC participation becomes predictable rather than experimental, Sony's internal teams may allocate more resources to multi-platform production pipelines. This would reduce the manual overhead required for PC conversions and support more efficient updates post-launch.

Market conditions also play a role. While the PC market remains large, competition for attention has grown, and players have become more demanding about the performance of AAA titles on varied hardware. Sony's future success on Steam may depend on maintaining high technical standards, optimizing across hardware tiers, and supporting titles through consistent updates. The expectations around PC versions differ materially from console releases, where uniform hardware makes optimization more straightforward.


PlayStation®


ST

Sony's top-earning PC titles demonstrate that strong narrative franchises and action-driven experiences have consistent appeal among PC players. However, live-service titles are likely to drive future growth. The long-term viability of PC revenue will depend on how effectively Sony can connect its upcoming slate of multiplayer and ongoing-update projects to the audiences that have responded so strongly to Helldivers 2.

The steady income from Steam gives Sony additional flexibility in how it positions its major releases and builds franchise longevity. While the company's strategic emphasis remains rooted in PlayStation hardware, the PC market has become a supplemental avenue that influences software planning, production decisions, and audience reach.

Digiarty MacXDVD

56% OFF

MACX DVD RIPPER PRO

Convert any DVD Movies to Videos in 5 Minutes

Learn More >>

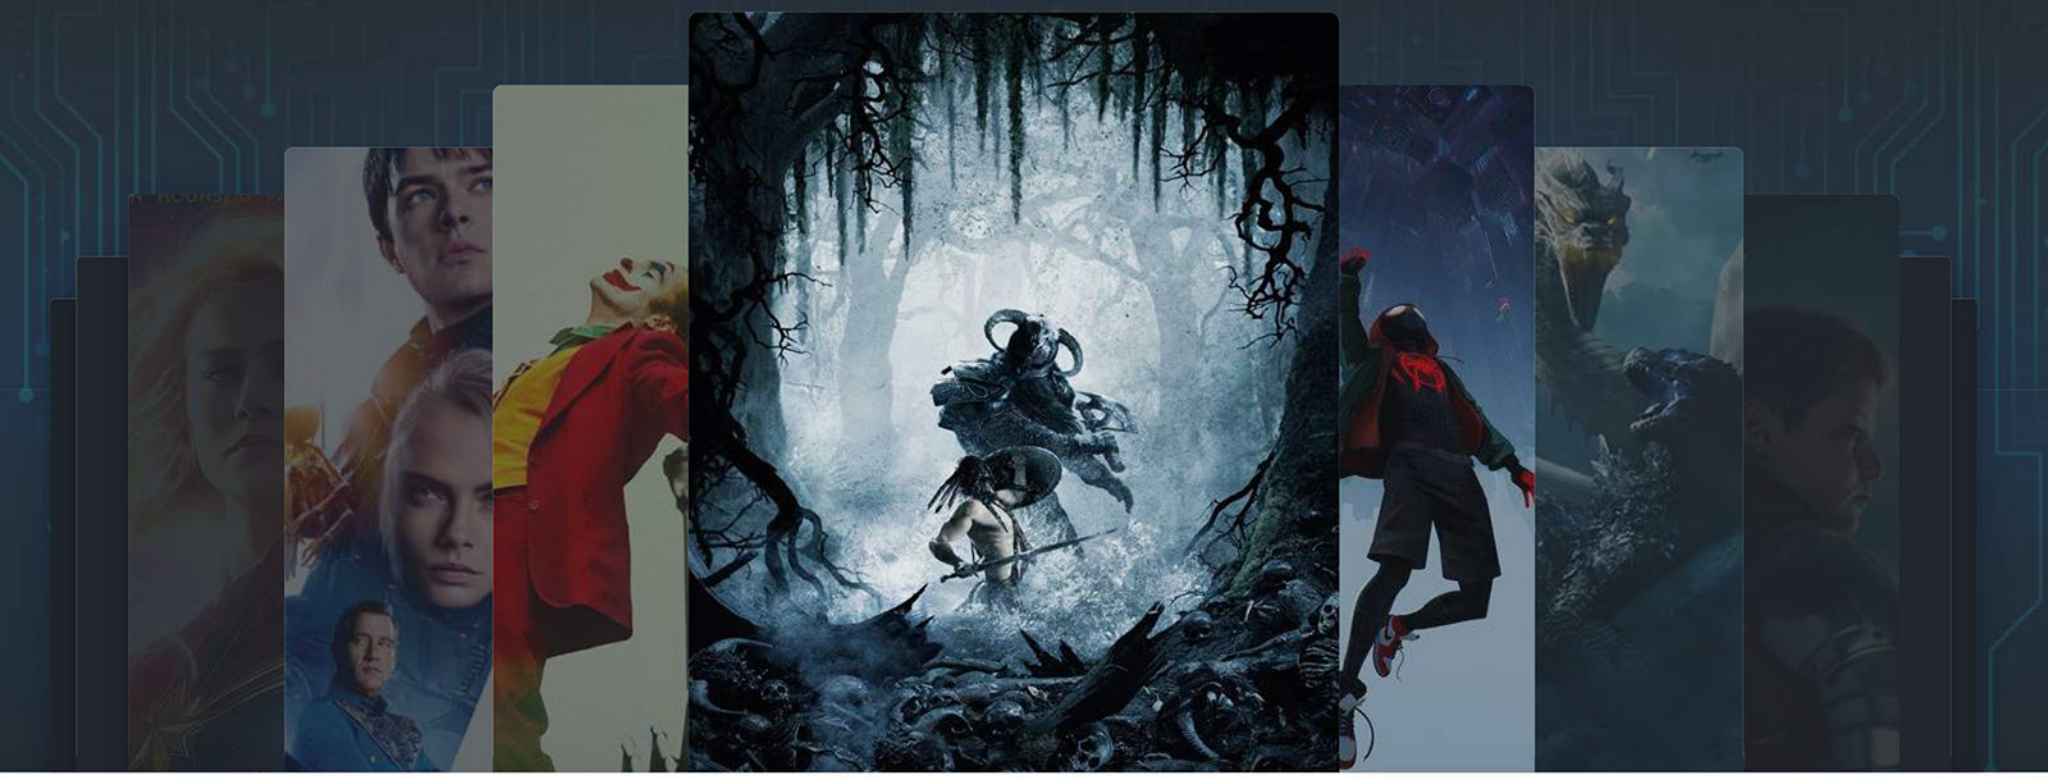

- Convert all old and new DVD movies to MP4, MOV.
- Convert DVD to all iPhone, iPad, Apple TV, Android, etc.
- Bypass DVD copy protections of all types of DVDs.
- Edit, cut, crop, merge video, add subtitle to video.
- No.1 fast speed, no quality loss thanks to GPU acceleration.

Buy Full Version
~~$67.95~~ **for only $29.95**

30 Day Money Back | Free Lifetime Upgrade | 7X24 Tech Support

Digiarty MacXDVD

50% OFF

MACX VIDEO CONVERTER PRO

The Fastest Tool to Convert, Resize, Cut, Download 4K, HD Videos.

Learn More >>

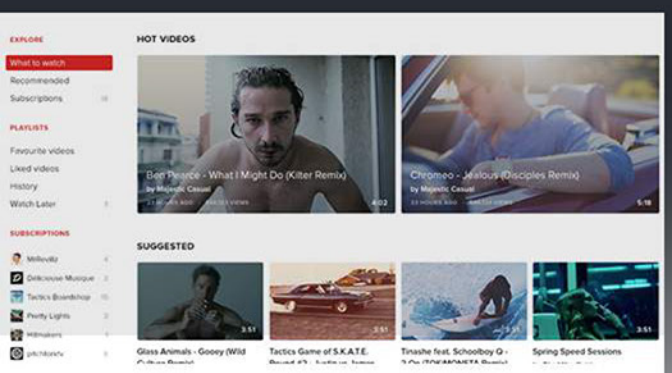

Download

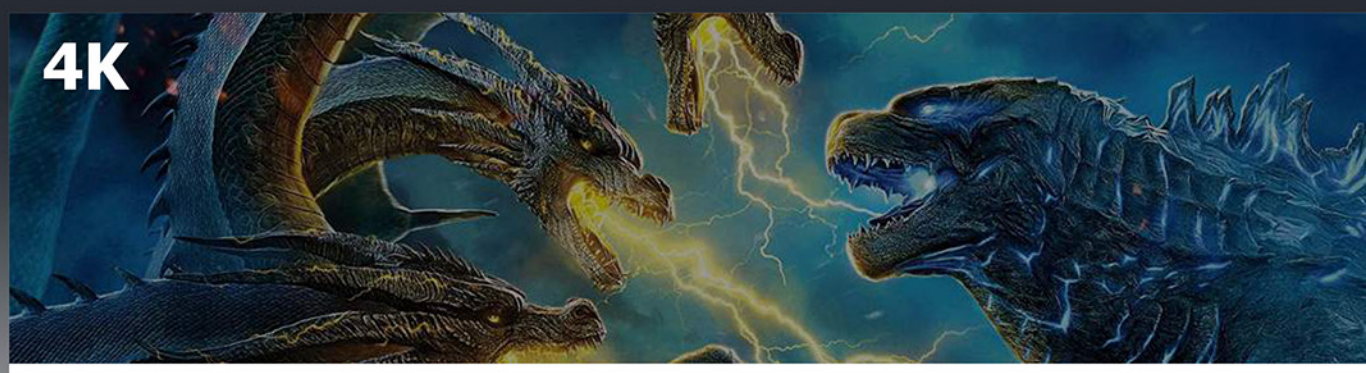

4K Support

Cut Video

Compress

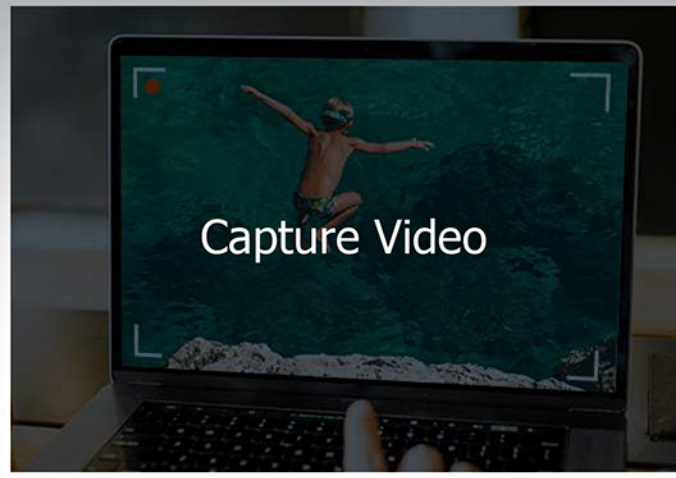

Capture Video

- Convert 4K/HD videos to MP4, MOV for Apple, Android
- Download videos from YouTube and 1000+ sites
- Compress video file size without quality loss
- Edit video clips with a few clicks
- GPU acceleration tech, no. 1 fast speed.

Buy Full Version
~~$59.95~~ **for only $29.95**

30 DAY 30 Day Money Back

Free Lifetime Upgrade

7X24 Tech Support


Chat
AI

USER SURVEY FINDS SEVERAL CHATBOTS OUTRANK CHATGPT IN VALUE METRICS

A recent user-centric study of generative-AI chatbots suggests that several alternatives to OpenAI's ChatGPT are perceived as stronger in certain conversational and reasoning dimensions, even though ChatGPT remains the market leader by user volume. The research, conducted by Prolific and publicly shared via Hugging Face, evaluated 28 distinct models using a head-to-head framework called "Humaine" and found that Google's Gemini 2.5 Pro emerged at the top of the ranking, followed by models from DeepSeek, Mistral AI and xAI. ChatGPT-4.1 was ranked eighth. While the survey reflects a specific methodology and user base, its results highlight shifting perceptions about what users value in AI chatbots beyond raw popularity.


AI


ChatGPT

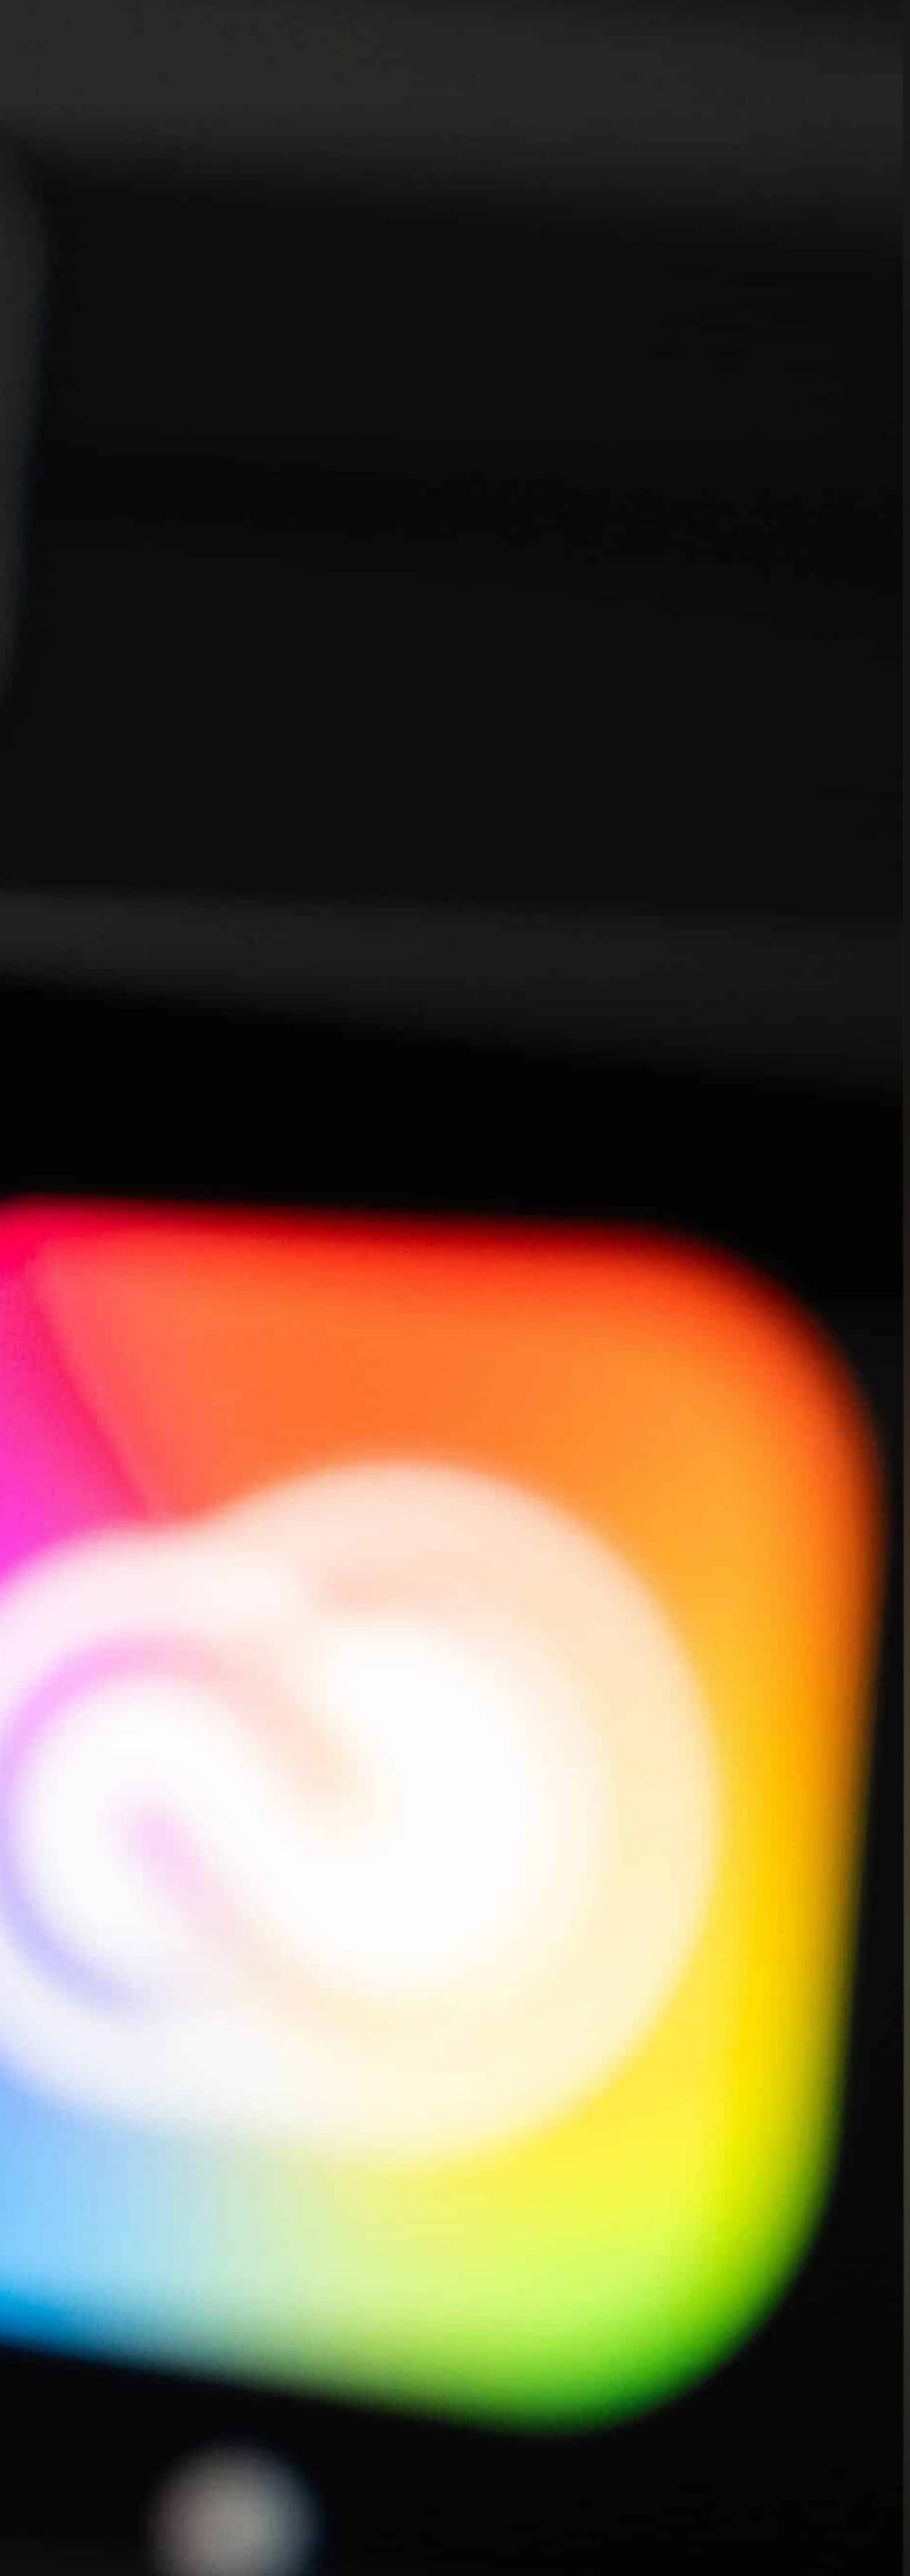

The study emphasised metrics such as how well a chatbot understands real-world dialogue, adapts to changes in context, presents information clearly, and operates in a manner users trust. These human-centric dimensions differ from technical benchmarks that focus on speed or accuracy alone. Because of this approach, a model that excels in reasoning or large-scale question-answering may still score lower if it struggles with conversational nuance or clarity. The findings invite closer scrutiny of how AI systems are evaluated, deployed and branded in a landscape where conversational fluency and trust increasingly matter.

WHAT USERS VALUE AND HOW MODELS STACK UP

The Humaine survey sought to move beyond traditional evaluation frameworks by anonymising pairs of chatbots and asking participants to judge which they preferred across a series of real-life conversational scenarios. These included multi-turn chats, clarification requests, shifting topics and conversational detours. The scoring metrics were split into four main categories: Core Task Performance & Reasoning, Interaction Fluidity & Adaptiveness, Communication Style & Presentation, and Trust, Ethics & Safety.

In the latest results, Google's Gemini 2.5 Pro achieved a top score, averaging about 18.75 out of a possible 27 in the head-to-head tests. It was followed by DeepSeek v3, which ranked second, and Mistral's Magistral Medium in third place—both demonstrating strengths in conversational style and adaptiveness. Notably,

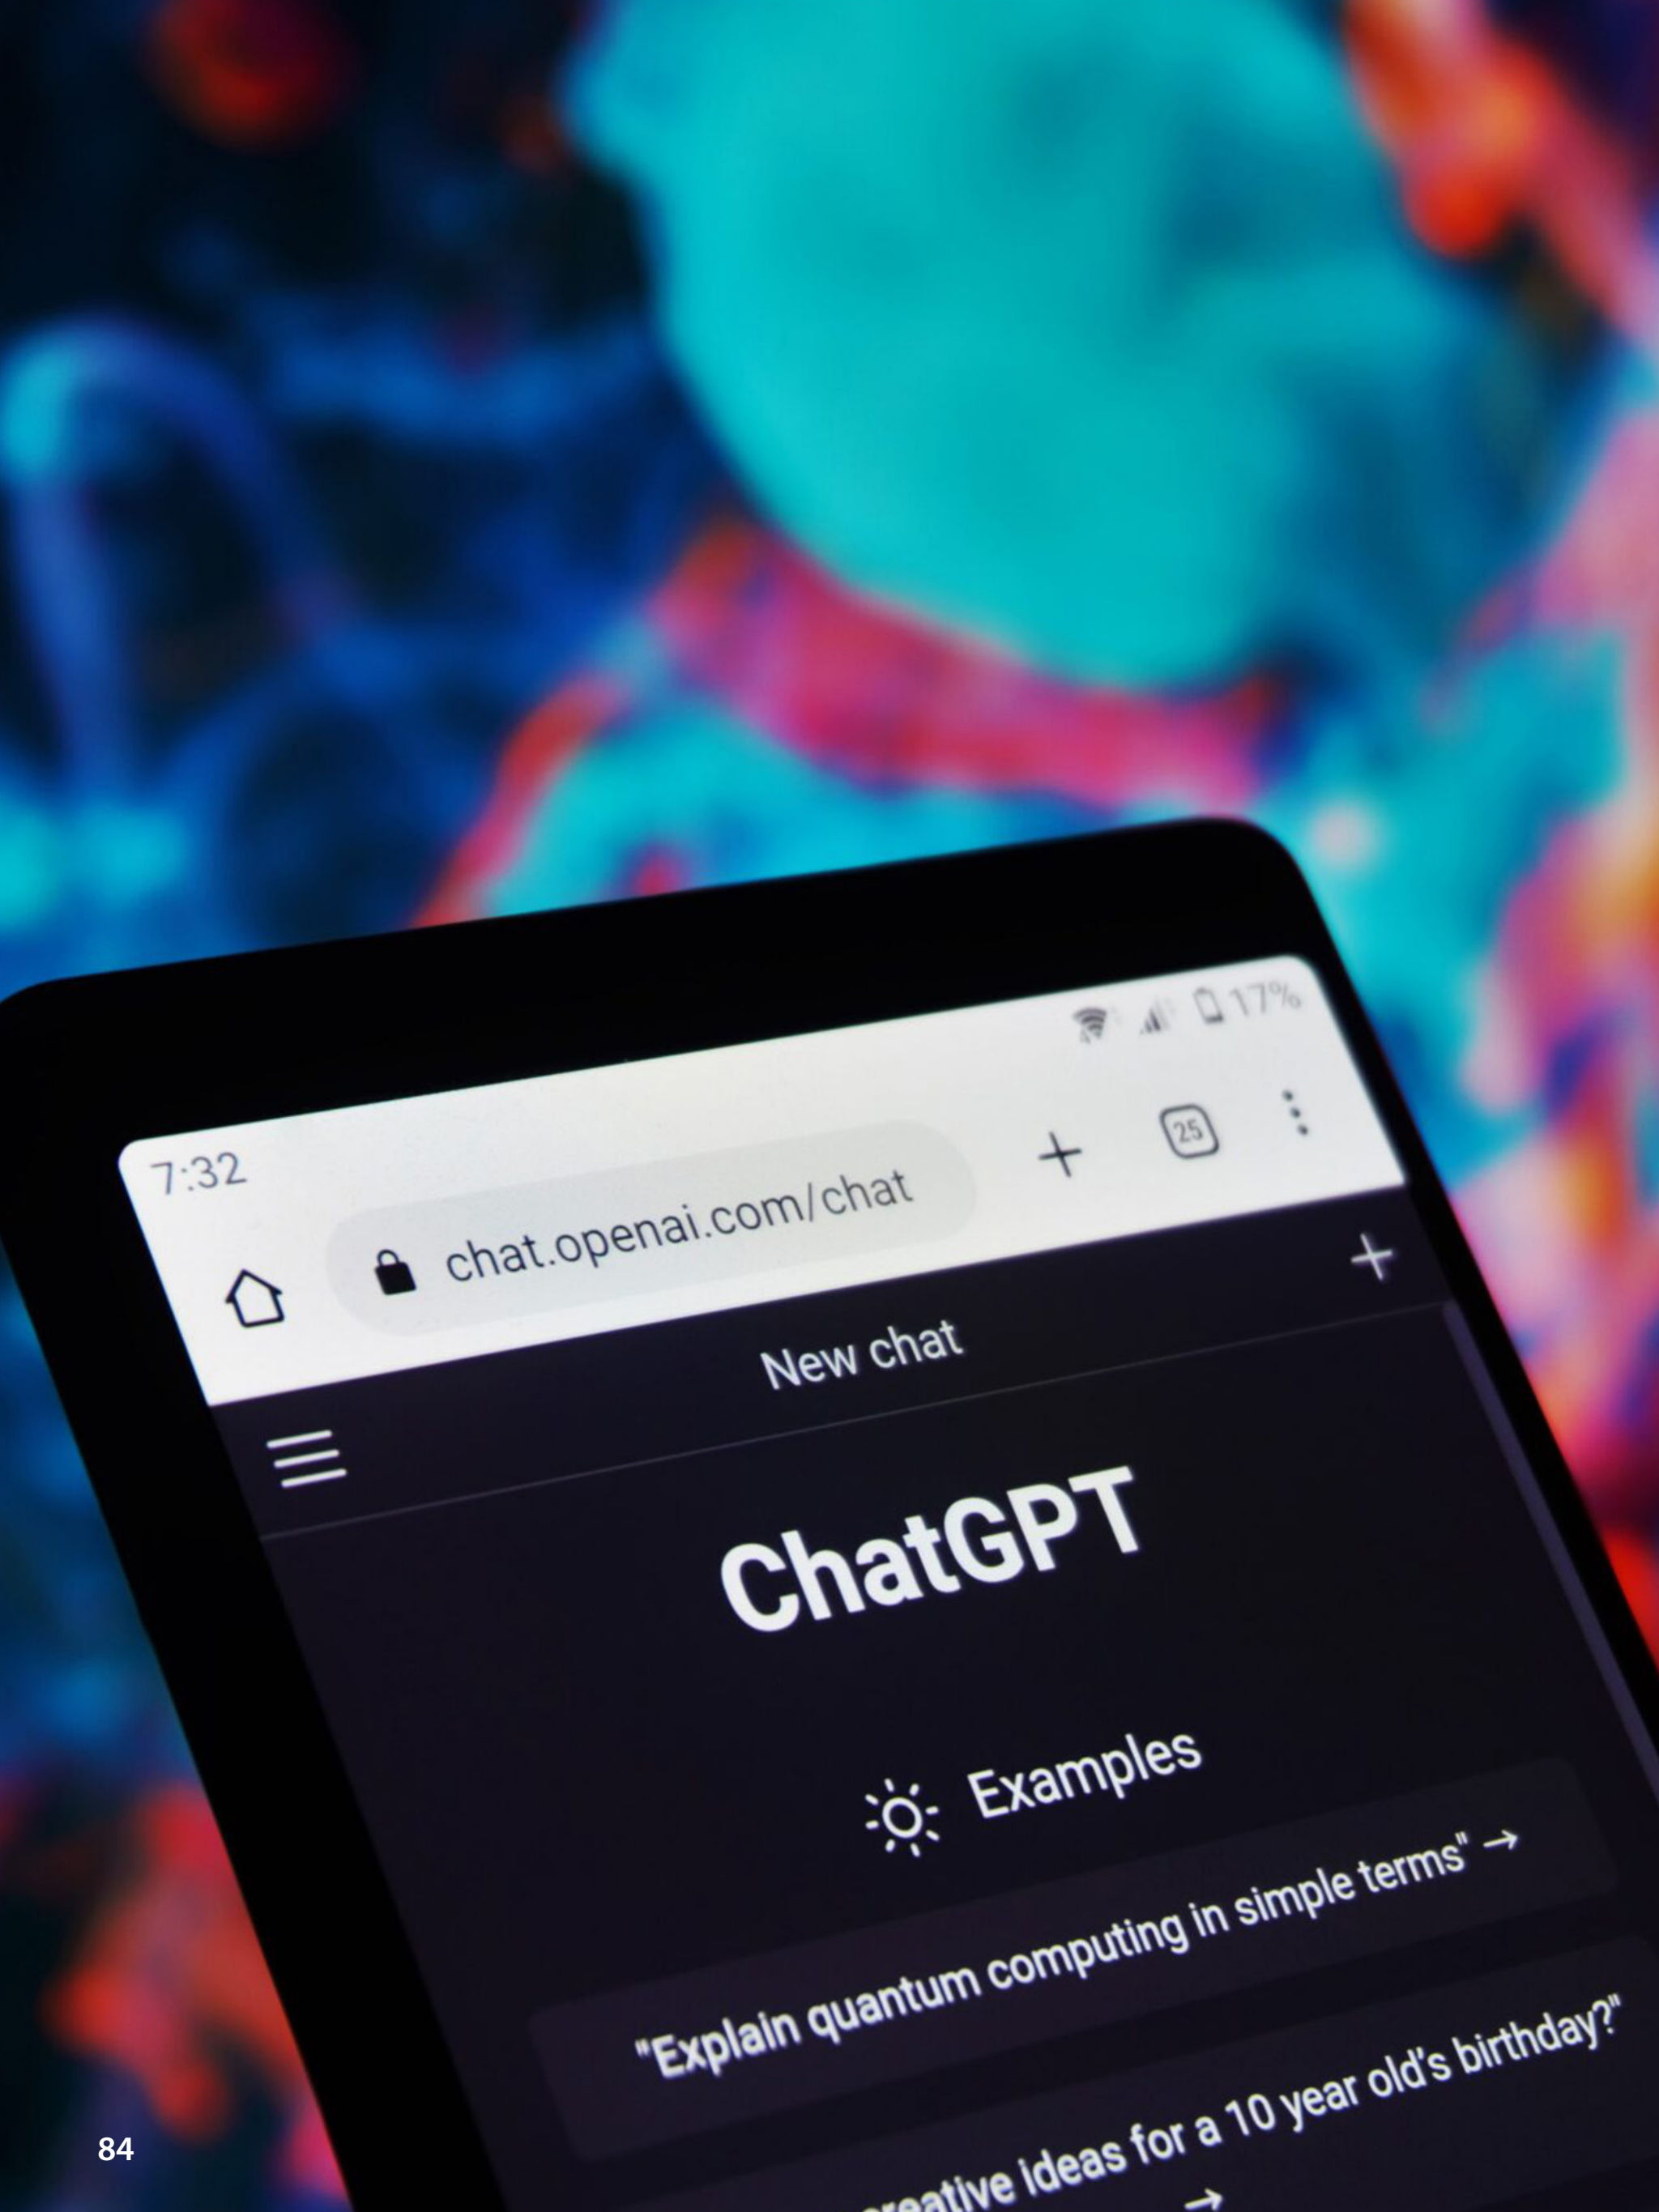
7:32
chat.openai.com/chat
New chat
ChatGPT
Examples
"Explain quantum computing in simple terms" →
creative ideas for a 10 year old's birthday?"

OpenAI's ChatGPT-4.1, despite dominating market share, placed eighth in this ranking. While the survey does not argue that ChatGPT is functionally inferior, it indicates that users judged other models more favorably in the tested categories.

The results also illustrate variation across demographic segments. Users from different age groups rated model performance differently, with younger participants placing slightly higher value on conversational fluency and older users emphasizing clarity and trust-worthiness. While overall ranking order remained consistent across segments, these subtle differences hint at how user expectations shift with age and experience.

The study further observes that the large scale of ChatGPT's user base may contribute to brand inertia, but that does not automatically translate into higher user-preference scores in structured tests.

IMPLICATIONS FOR THE AI CHATBOT MARKET AND OUTLOOK

The user-preference data adds an interesting layer to the competitive dynamics in the generative-AI space. While ChatGPT continues to dominate in terms of usage and ecosystem reach—accounting for nearly half of all chatbot interaction globally according to referenced figures—the gap in user-judged performance suggests that perception and experience may allow other platforms to gain traction.

For developers and enterprise users, this means choice of model may increasingly depend on qualitative dimensions rather than sheer brand visibility.

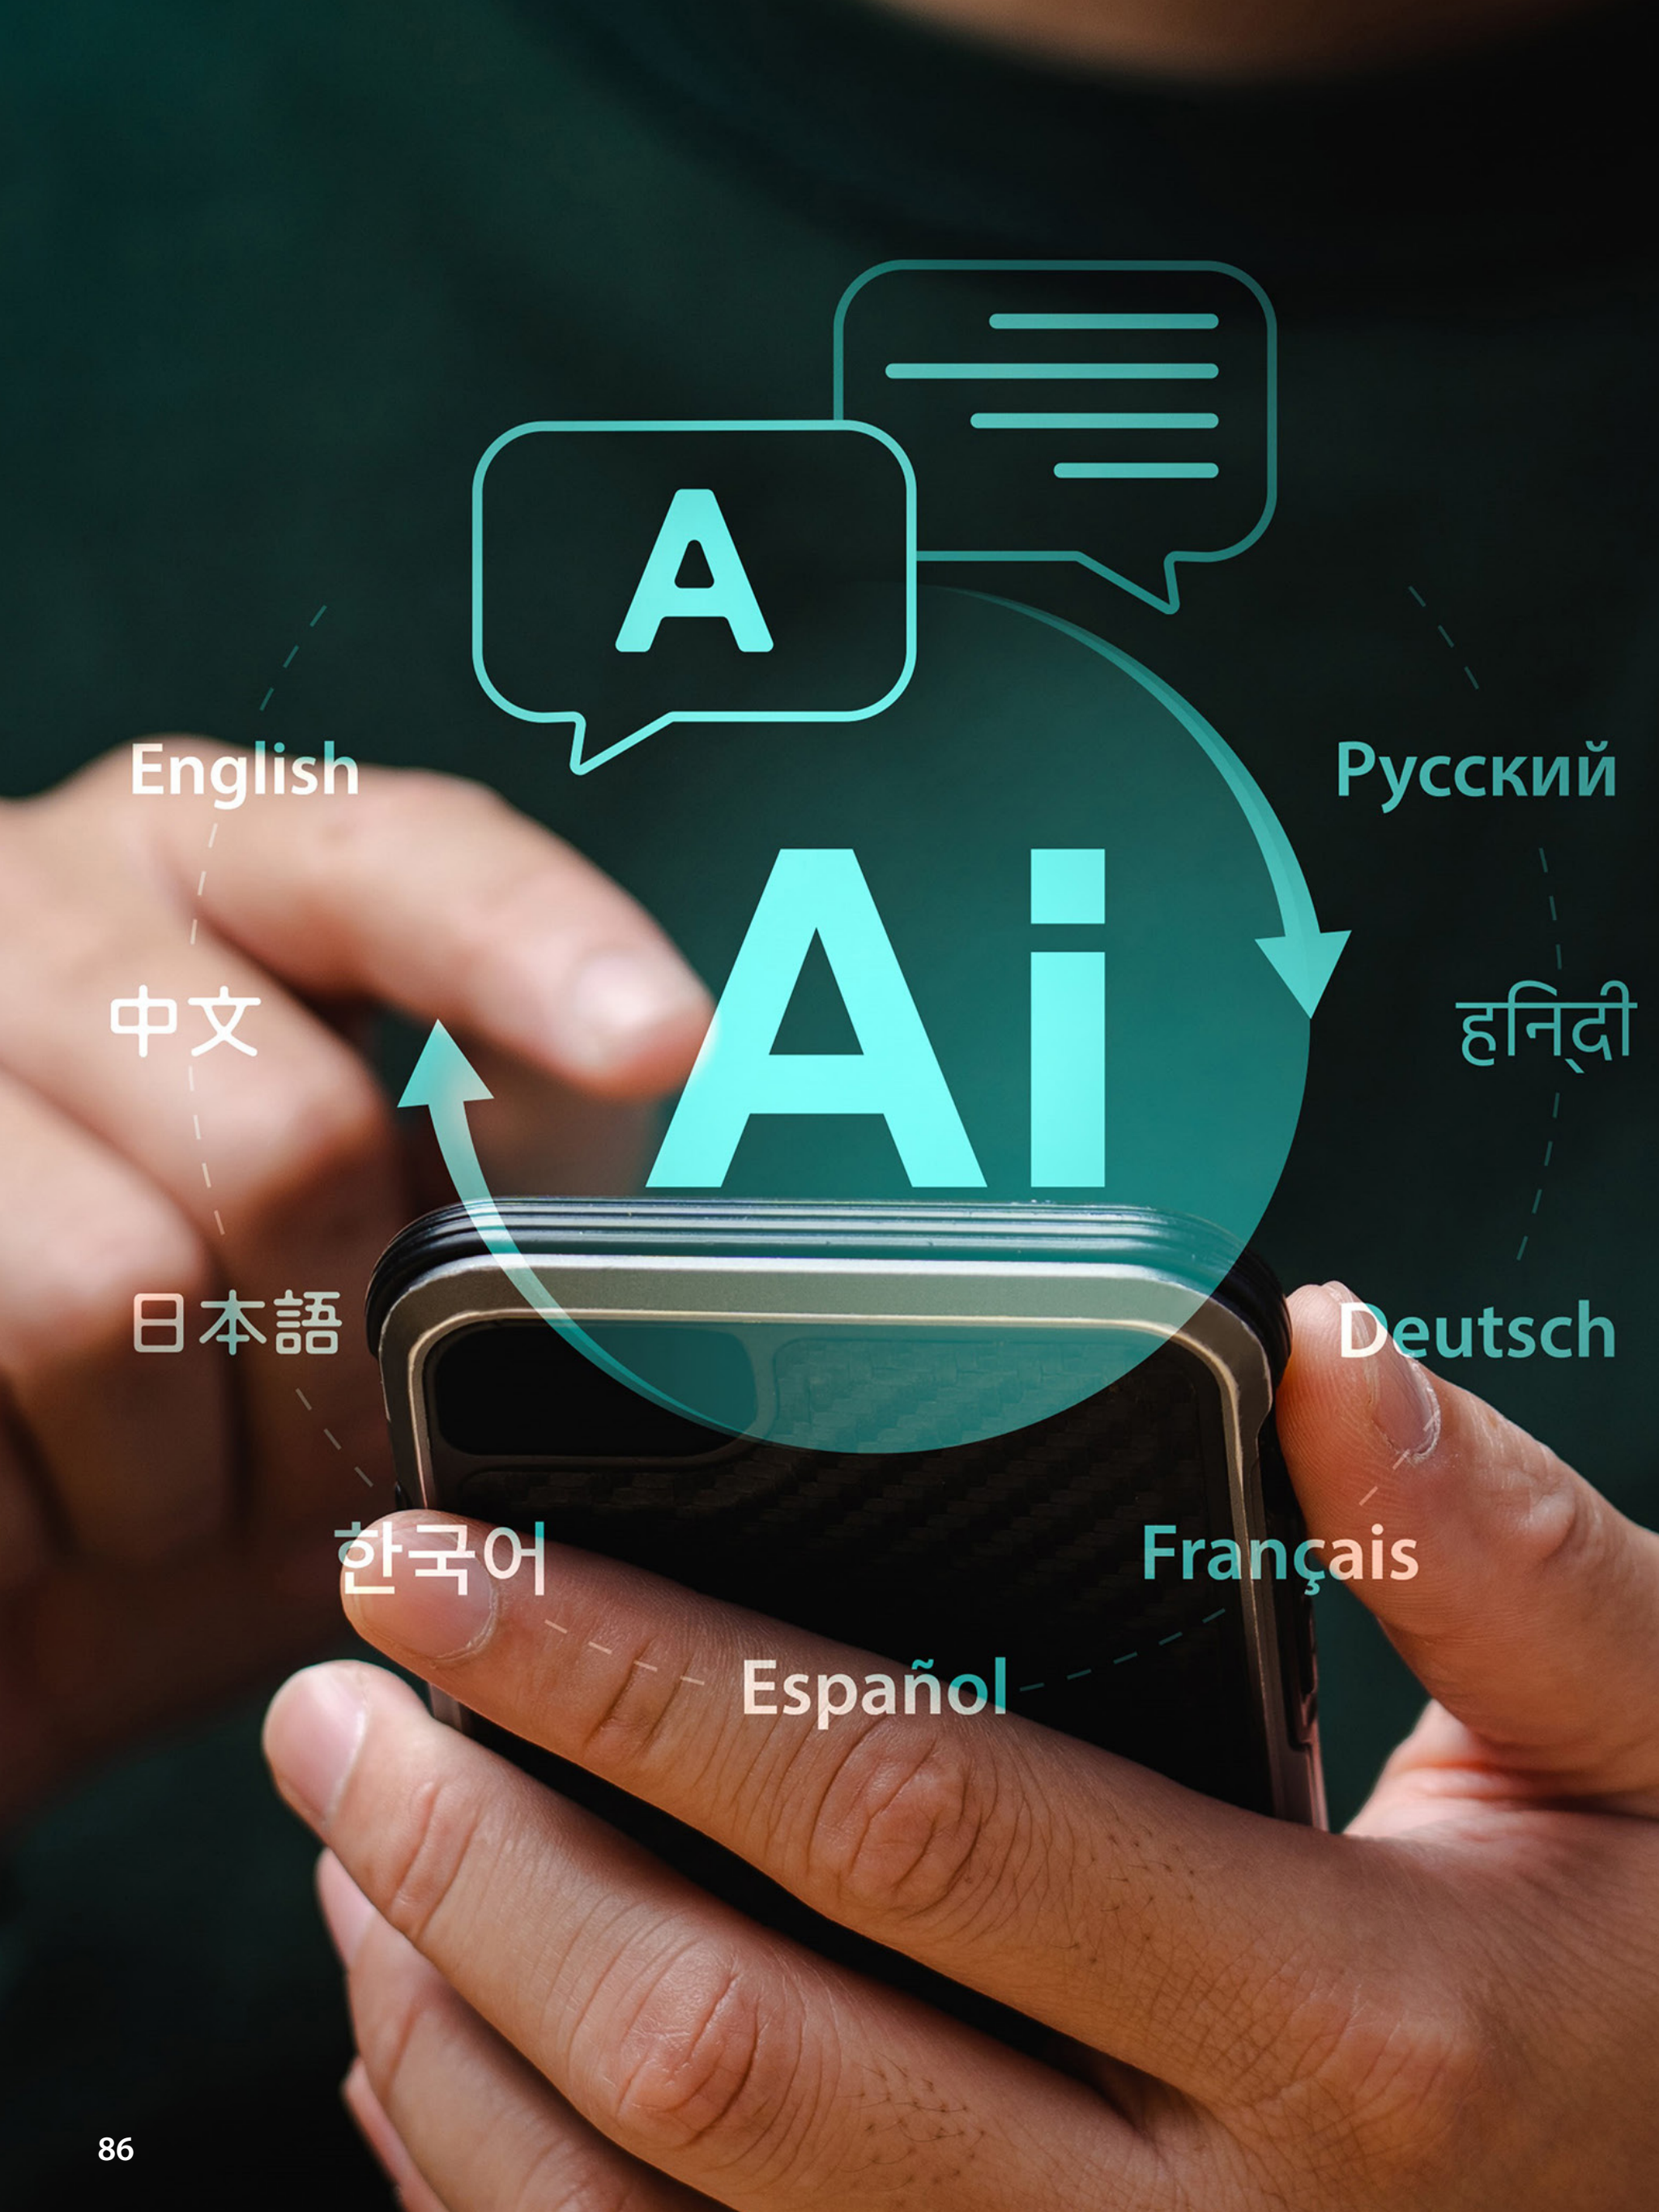
A
Ai
English
Русский
中文
हिन्दी
日本語
Deutsch
한국어
Français
Español

The results may prompt firms to refine how they present chatbot capabilities, moving away from headline statistics toward emphasizing conversation quality, adaptability and trust. For instance, enterprises deploying chatbots in customer-service or internal-knowledge workflows may prioritize models that rate higher on clarity and adaptiveness rather than just speed or scale. For providers, the findings suggest that message, branding and user-experience investment could become as important as parameter size or raw reasoning benchmarks.

Still, several caveats apply. The Humaine survey targets a specific user population and covers a subset of possible conversational scenarios; it does not replace technical performance evaluations nor does it guarantee one model will perform better than another in all use cases. Additionally, deployment context, fine-tuning, access to compute resources and customisation all influence real-world outcomes. In short, a model rated highest by users in a study may not deliver better results for every enterprise or domain.

Nonetheless, the study indicates that user experience remains a key differentiator in chatbot development. As AI systems proliferate across consumer and enterprise settings, expectations around fluency, adaptability and trust will continue to shape how models are selected, trained and marketed. The presence of strong challengers to ChatGPT also implies that the era of a single dominant chatbot provider may give way to a more diversified landscape.

Run Your Private Cloud Office

All-in-one suite for your business

Make use of ONLYOFFICE online editors integrated with collaboration platform:
mail, CRM, projects, calendar and much more.

180-Day Free Trial

GET ONLYOFFICE NOW

ONLYOFFICE by Ascensio System SIA offers a powerful office suite that comprises online editors for text documents, spreadsheets and presentations highly compatible with Microsoft Office and OpenDocument file formats. ONLYOFFICE provides users with multiple editing tools and collaborative features ensuring greater team workflow, and seamless work with complex formatting and objects within your web solution.

100% compatible with MS Office formats

www.onlyoffice.com

ONLYOFFICE

iPhone ›

iPad ›

Mac ›

Apple Watch ›

AirPods ›

Beats ›

Apple TV ›

Accessories ›

HomePod ›

Pro Cameras

THE MOST ADVANCED OPTICAL
AND COMPUTATIONAL POWER

The iPhone has long been a trailblazer in smartphone technology, transforming pocket-sized devices into veritable powerhouses of visual storytelling. With the iPhone 17 Pro, Apple has elevated this legacy to unprecedented heights. Boasting a trio of 48MP Fusion cameras, groundbreaking ProRes RAW video capabilities, and seamless integration into professional workflows, the iPhone 17 Pro is a studio in your hand.

Image: Samuel Angor

PRO-CALIBER MASTERY

The new device bridges the gap between casual snapshots and cinematic masterpieces, powered by the most sophisticated blend of optical engineering and computational wizardry yet seen in a smartphone. As we dive into its capabilities, we'll trace the device's evolutionary journey, dissect its technical prowess, and explore how it empowers creators from hobbyists to Hollywood pros. Whether you're chasing the perfect sunset or syncing footage on a film set, the iPhone 17 Pro delivers raw, unfiltered potential.

The iPhone's camera system didn't conquer the world overnight. Its ascent is a testament to Apple's relentless iteration, blending hardware innovation with software smarts to redefine mobile imaging. Let's rewind to 2007, when the **original iPhone debuted with a solitary 2MP rear camera,** no flash, no autofocus, just a basic sensor capable of VGA video at 30fps. It was revolutionary for its time, turning phones into point-and-shoots, but by today's standards, it was rudimentary. Fast-forward to the iPhone 3G and 3GS (2008–2009), which nudged resolution to 2MP and added video recording, albeit grainy and fixed-focus.

The **iPhone 4 in 2010 marked a pivotal shift: a 5MP rear sensor with LED flash and 720p HD** video, plus a front-facing VGA camera for the era's nascent FaceTime. But it was the iPhone 4S that first whispered computational photography with improved low-light performance. By iPhone 5 (2012), we saw an 8MP sensor, 1080p video, and panorama mode, hallmarks of a maturing system. The **dual-camera era dawned with the iPhone**

Image: MacRumors

7 Plus (2016), packing a 12MP wide-angle and telephoto lens for 2x optical zoom and Portrait mode, Apple's inaugural foray into depth-sensing bokeh effects via software. This hybrid approach, optical hardware augmented by computational algorithms, became the blueprint. The iPhone X (2017) refined this with dual 12MP rear and a 7MP front camera, introducing TrueDepth for Animoji and Face ID, while **Night mode arrived in the iPhone 11 (2019) to conquer dark scenes** through multi-frame stacking.

Enter the Pro era: **iPhone 11 Pro (2019) tripled up with ultra-wide, ushering in 4K 60fps video** and Deep Fusion for pixel-level detail enhancement. LiDAR scanner in iPhone 12 Pro (2020) supercharged low-light autofocus, and Cinematic mode in iPhone 13 Pro (2021) brought rack-focus video magic. The tetraprism periscope lens in iPhone 15 Pro Max (2023) delivered 5x optical zoom, a leap from digital crops. ProRAW followed in 2021, enabling manual-like control over exposure and color. The **iPhone 16 Pro (2024) pushed 48MP across main and ultra-wide lenses,** with 4K 120fps ProRes. Now, the iPhone 17 Pro (2025) completes the trifecta: all three cameras at 48MP, an 8x optical zoom via an upgraded tetraprism with a 56% larger sensor, and ProRes RAW for uncompressed video. This evolution isn't haphazard; it's a consistent march toward versatility, higher resolutions, broader focal ranges, and deeper software integration, making pros out of everyday users. From 2MP to 48MP quad-pixel sensors, Apple's cameras have grown not just in megapixels but in intelligence, capturing life's nuances with surgical precision.

THE 48MP SENSOR REVOLUTION

At the core of the iPhone 17 Pro's prowess lies its 48MP Pro Fusion camera array: a main, ultra-wide, and telephoto, each a 48MP marvel. The primary camera sports a 24/48mm focal length (1x/2x), f/1.78 aperture, and 2.44µm quad-pixel binning that fuses four pixels into one for superior low-light sensitivity (effective 1.22µm pixels). The ultra-wide (13mm, 0.5x/ macro) mirrors this at f/2.2 with 1.4µm quad-pixels, enabling razor-sharp macro shots down to 2cm. The telephoto steals the show:

100/200mm (4x/8x), f/2.8, with the same quad-pixel tech for lossless cropping. Default output? Stunning 24MP photos across all lenses, balancing file size with detail. This sensor upgrade (56% larger on the tele) delivers the longest iPhone telephoto ever, with up to 16x total optical zoom. Stabilization is sensor-shift OIS on all three, quelling shakes for buttery-smooth shots and video.

Pure optics have always been Apple's anchor, but the iPhone 17 Pro's lenses push boundaries. **The tetraprism design folds light for a**

48MP

4x Fusion Telephoto camera

100 mm focal length

ƒ/2.8 aperture

1.4 µm quad pixel

0.7 µm pixel

Tetraprism design

Hybrid Focus Pixels

3D sensor-shift OIS

12MP

8x Telephoto

200 mm focal length

ƒ/2.8 aperture

0.7 µm pixel

Tetraprism design

Hybrid Focus Pixels

3D sensor-shift OIS

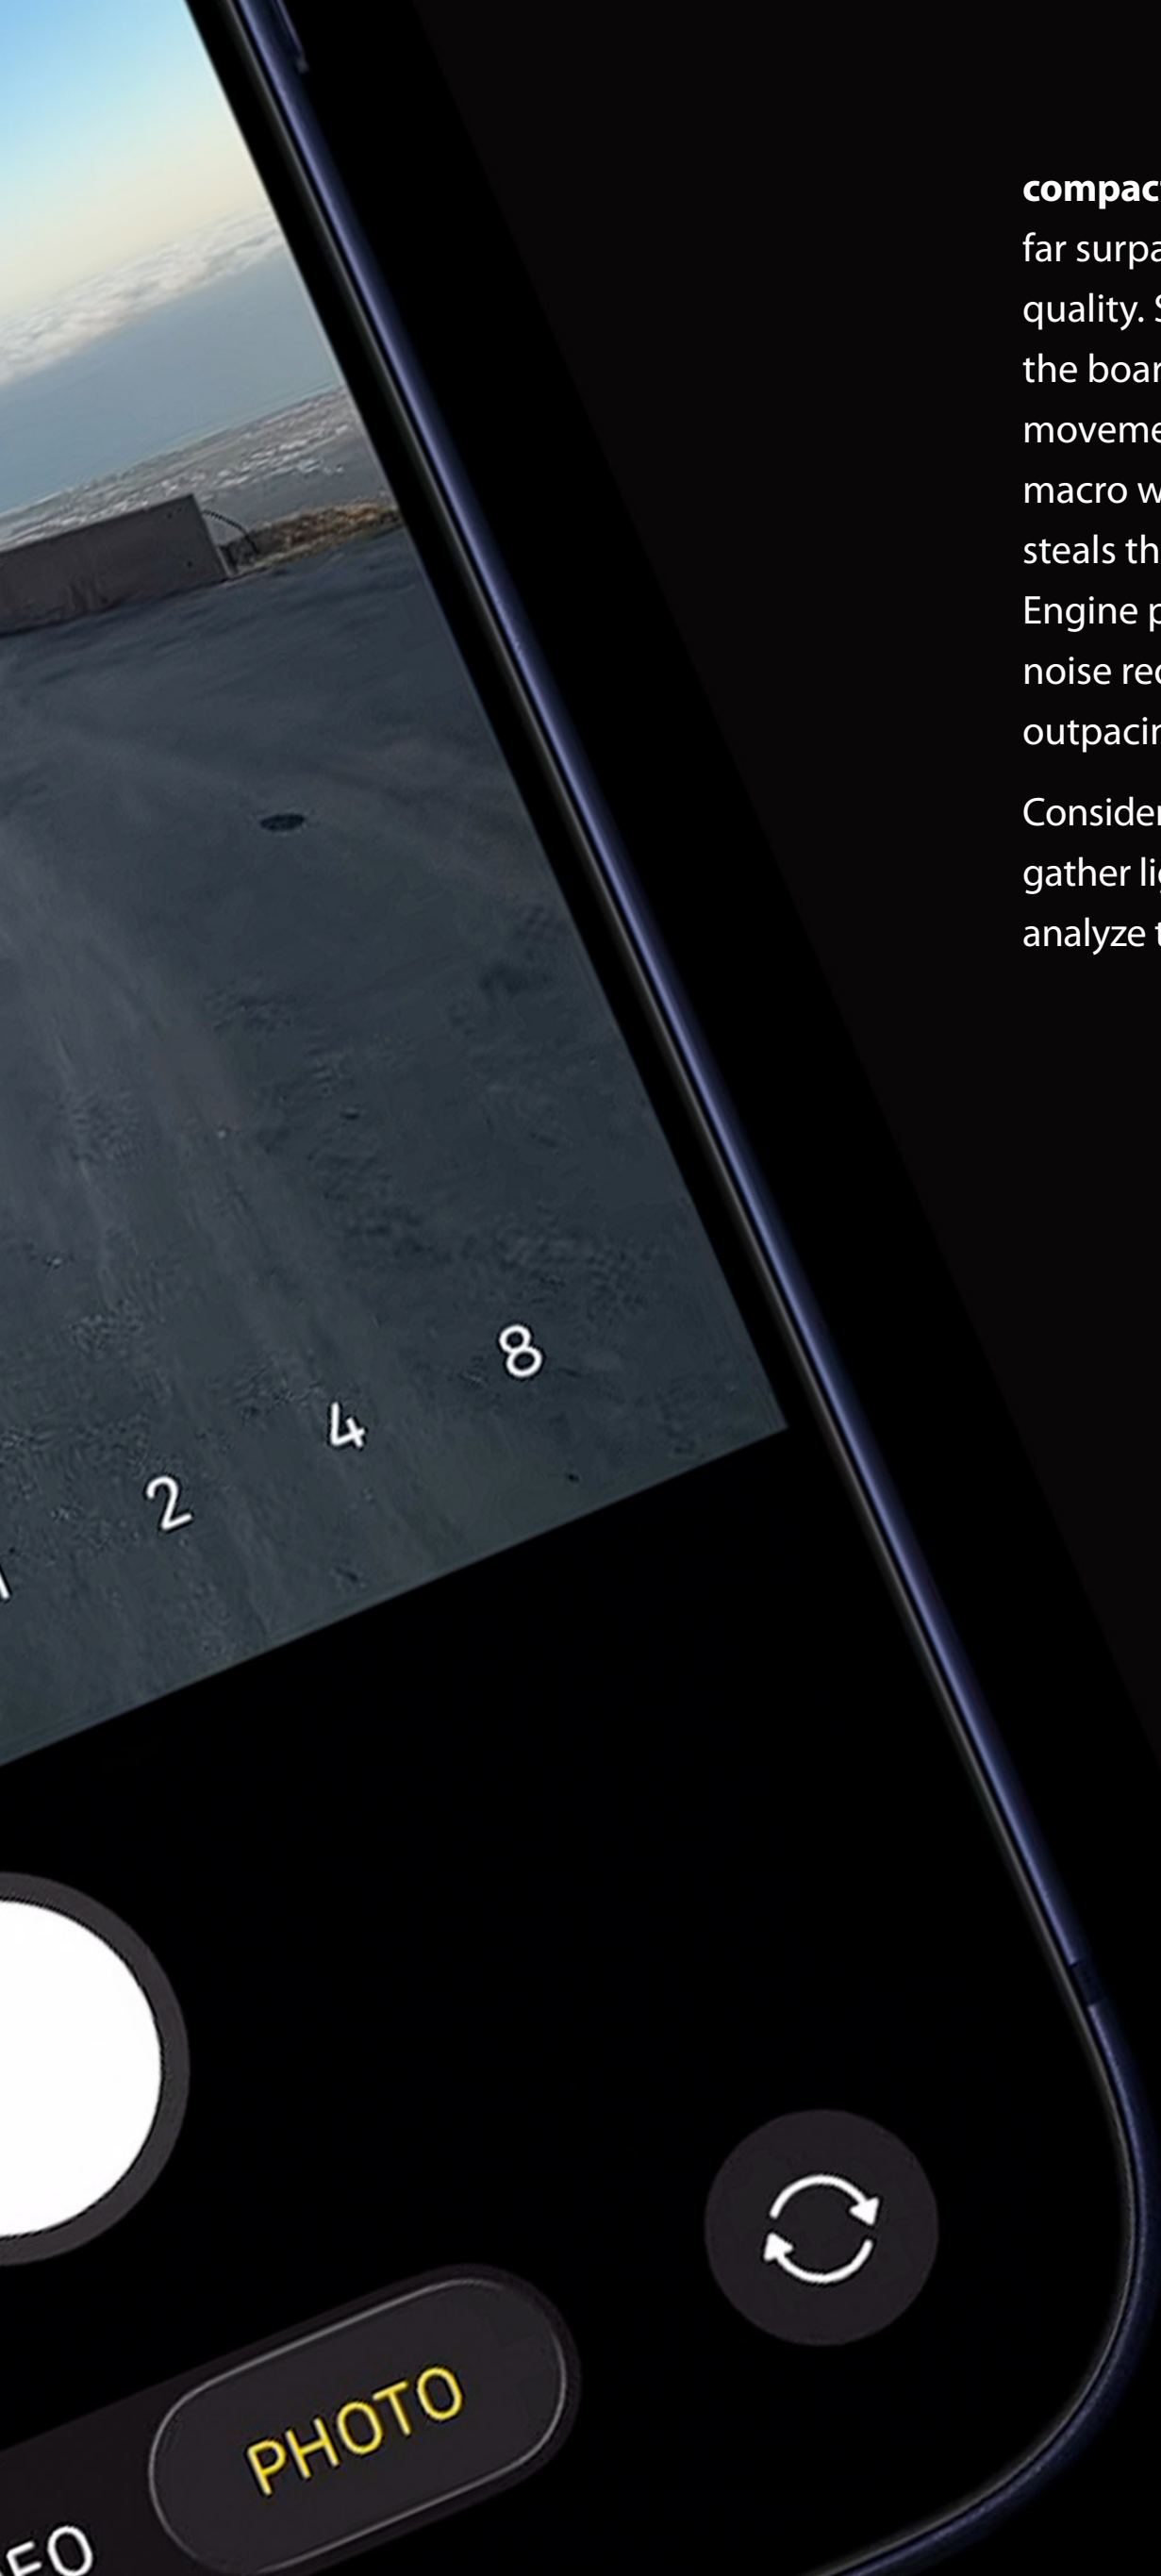

compact 8x optical zoom without distortion, far surpassing digital crops that degrade quality. Sensor-shift stabilization, now across the board, compensates for up to five axes of movement, ideal for handheld telephoto or macro work. Yet, computational photography steals the spotlight here. Features like Photonic Engine process multi-frame bursts for HDR, noise reduction, and edge detection, often outpacing raw optics in dynamic scenes.

Consider a low-light portrait: Optical lenses gather light, but algorithms like Deep Fusion analyze textures pixel by pixel, blending

MACRO

0.5×
13 MM

1×
24 MM

1.2×
28 MM

1.5×
35 MM

2×
48 MM

4×
100 MM

8×
200 MM

exposures for a realism that rivals DSLRs. Vs. Pure optics? Hardware excels in controlled zooms (e.g., 8x tetraprism yields sharper 200mm shots than cropped 1x), but computation shines in versatility. Night mode fuses 9+ frames, turning no-light into pro-grade. The synergy? Unbeatable: 8x optical for reach, AI for polish. As one reviewer notes, "The iPhone 17 Pro's '3=8' lens mathusing the 48MP sensor for cropped 8x—blurs lines, delivering optical-quality without extra glass." It's hardware as the foundation, software as the artist.

BLURRING THE LINES

Where does "good enough" end and professional begin? The iPhone 17 Pro erodes that divide. Casual users revel in point-and-shoot 24MP HEIFs with auto HDR, but pros unlock ProRAW: 48MP DNG files with editable metadata, non-destructive adjustments in Lightroom. The real game-changer? ProRes RAW video, he first on a smartphone. This Apple-developed codec captures uncompressed sensor data at up to 4K 120fps, preserving dynamic range for post-production grading without baked-in limits.

For eventual (casual) photography, it's effortless: Auto modes handle exposure, while computational tools like Semantic Rendering isolate subjects for ethereal bokeh. But Pro use demands more: ProRes RAW external recording to SSDs via USB-C, Apple Log 2 for 16+ stops of latitude, and Open Gate for full-sensor framing. Limits? Battery drain during RAW shoots (up to 4x faster depletion) and storage (a minute of 4K ProRes RAW eats 6GB). Yet, for filmmakers, it's liberating: Sync with DaVinci Resolve for color science that rivals ARRI cameras, all in a $1,199 device. **As Austin Mann's Dolomites review raves,** "The 48MP telephoto's low-light prowess turns alpine twilight into canvas." The threshold? Intent: Snap for Instagram? Casual bliss. Edit for Sundance? Pro paradise.

ZOOM

CENTER STAGE FOR CREATORS AND SELFIES

Don't sleep on the front: An 18MP Center Stage camera with ultra-wide field-of-view tracking. Tap to expand framing, rotate seamlessly from portrait to landscape, or pull in friends for epic "friendsies." For content producers, it's gold, auto-framing vloggers during walks, stabilizing group calls. Selfies? Smarter AI enhances skin tones, adds Portrait effects, and supports Night mode for dimly lit hangs. Paired with Dual Capture (simultaneous front- and rear-facing video), it's a creator's dream: film a reaction while capturing the scene, all in 4K.

The iPhone 17 Pro video is cinematic sorcery.

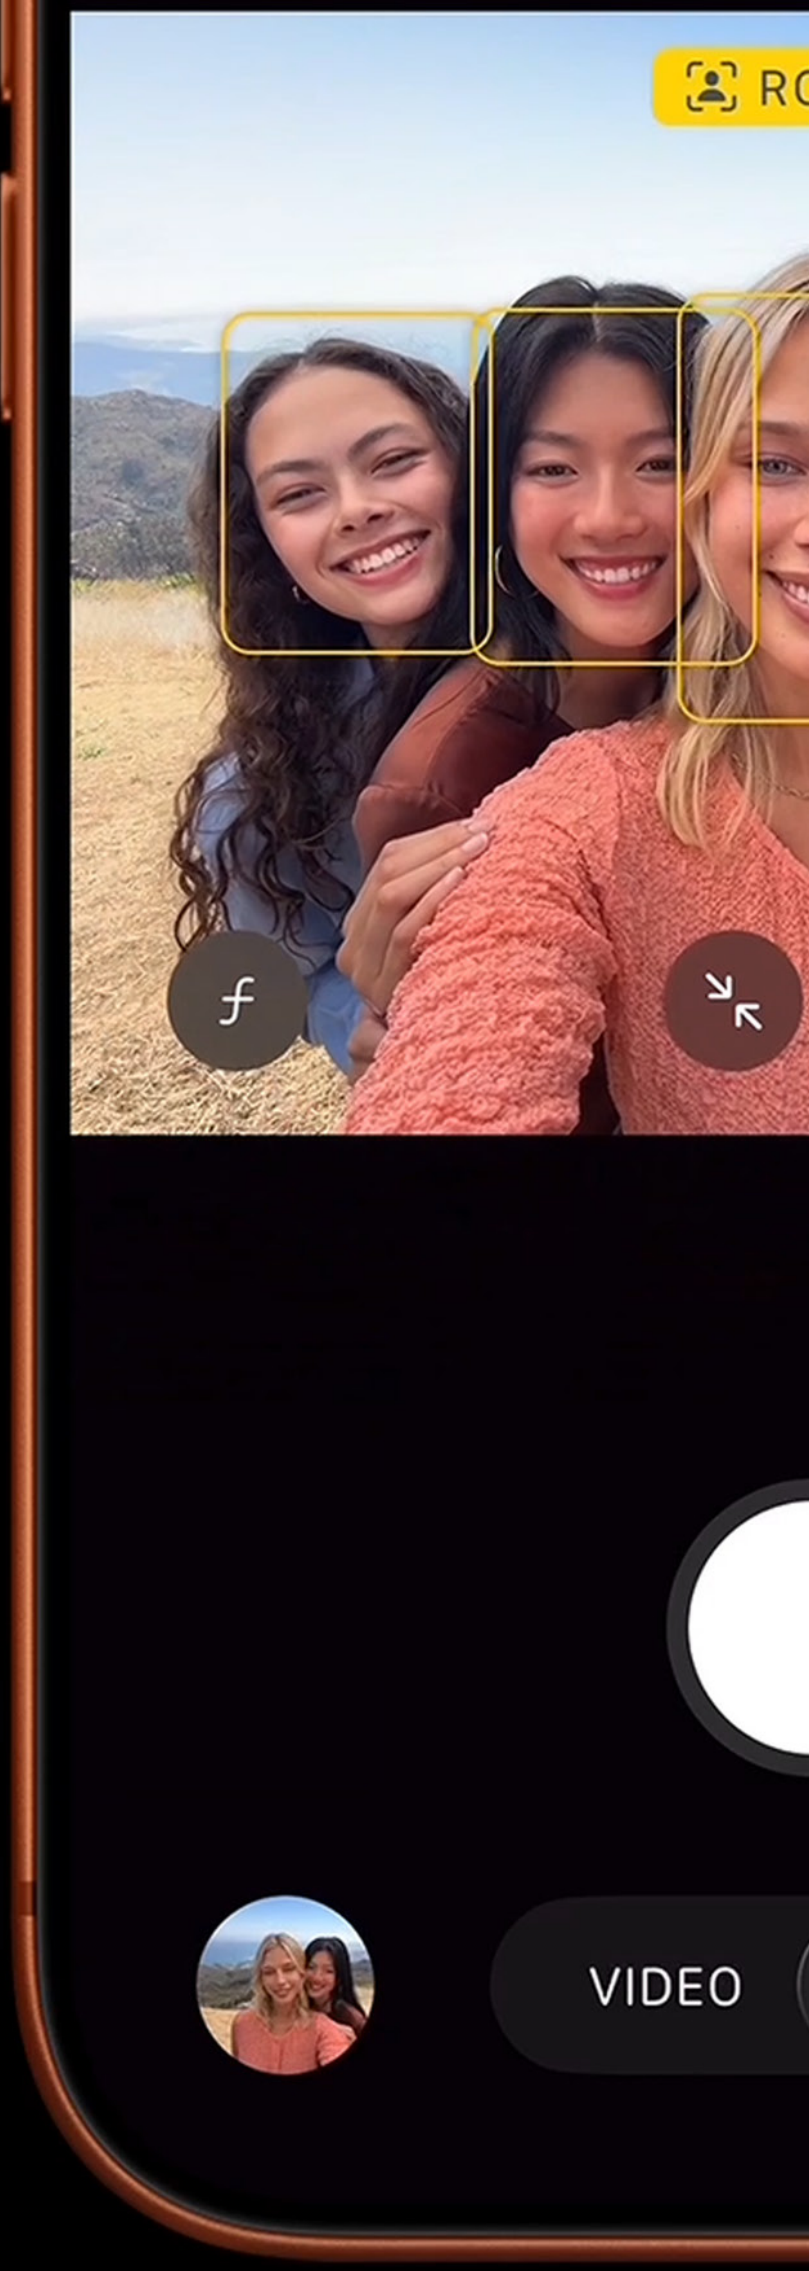

Dual Capture lets you record front and back simultaneously, perfect for BTS footage or tutorials. Ultra-stabilized video, leveraging Action mode and 4K 120fps, smooths extreme motion, think skiing or biking without gimbals. Apple's ProRes codec, now with RAW, is the crown jewel: Internal 4K 120fps ProRes, external unlimited via Thunderbolt, with Academy Color Encoding for broadcast compliance. Apple Log 2 extends shadows and highlights, enabling pro grading. Stabilization? Next-gen sensor-shift plus electronic image stabilization (EIS) delivers gimbal-free glide, even at 8x zoom.

The iPhone 17 Pro integrates into pro pipelines like never before. Genlock syncs video to

external sources via BNC (with Blackmagic's ProDock), while timecode embeds SMPTE metadata for multi-cam edits in Final Cut Pro or Premiere. Final Cut Camera 2.0 app adds broadcast rates, Open Gate, and telephoto ProRes support. Wire to an Atomos recorder for monitored RAW output, or dock into rigs for virtual production. Is it a replacement for REDs, or just a B-cam or specialty angle? Invaluable.

SUSTAINING THE POWER

Heavy lifting like ProRes RAW demands endurance. The iPhone 17 Pro's larger battery, up to 4 more hours than the 15 Pro Max, fuels extended shoots, thanks to A19 Pro chip efficiencies. The unibody titanium-aluminum frame (forged for strength) pairs with a laser-welded vapor chamber that dissipates heat 30% faster than predecessors. **iFixit teardowns reveal it keeps temps under 40°C during 4K loops,** averting throttling. More room inside means denser cells without bulk, ensuring sustained 120fps bursts.


A19
PRO

The iPhone 17 Pro's camera system is aspirational. From its evolutionary roots to ProRes RAW's unyielding fidelity, it democratizes pro tools while honoring optical purity.

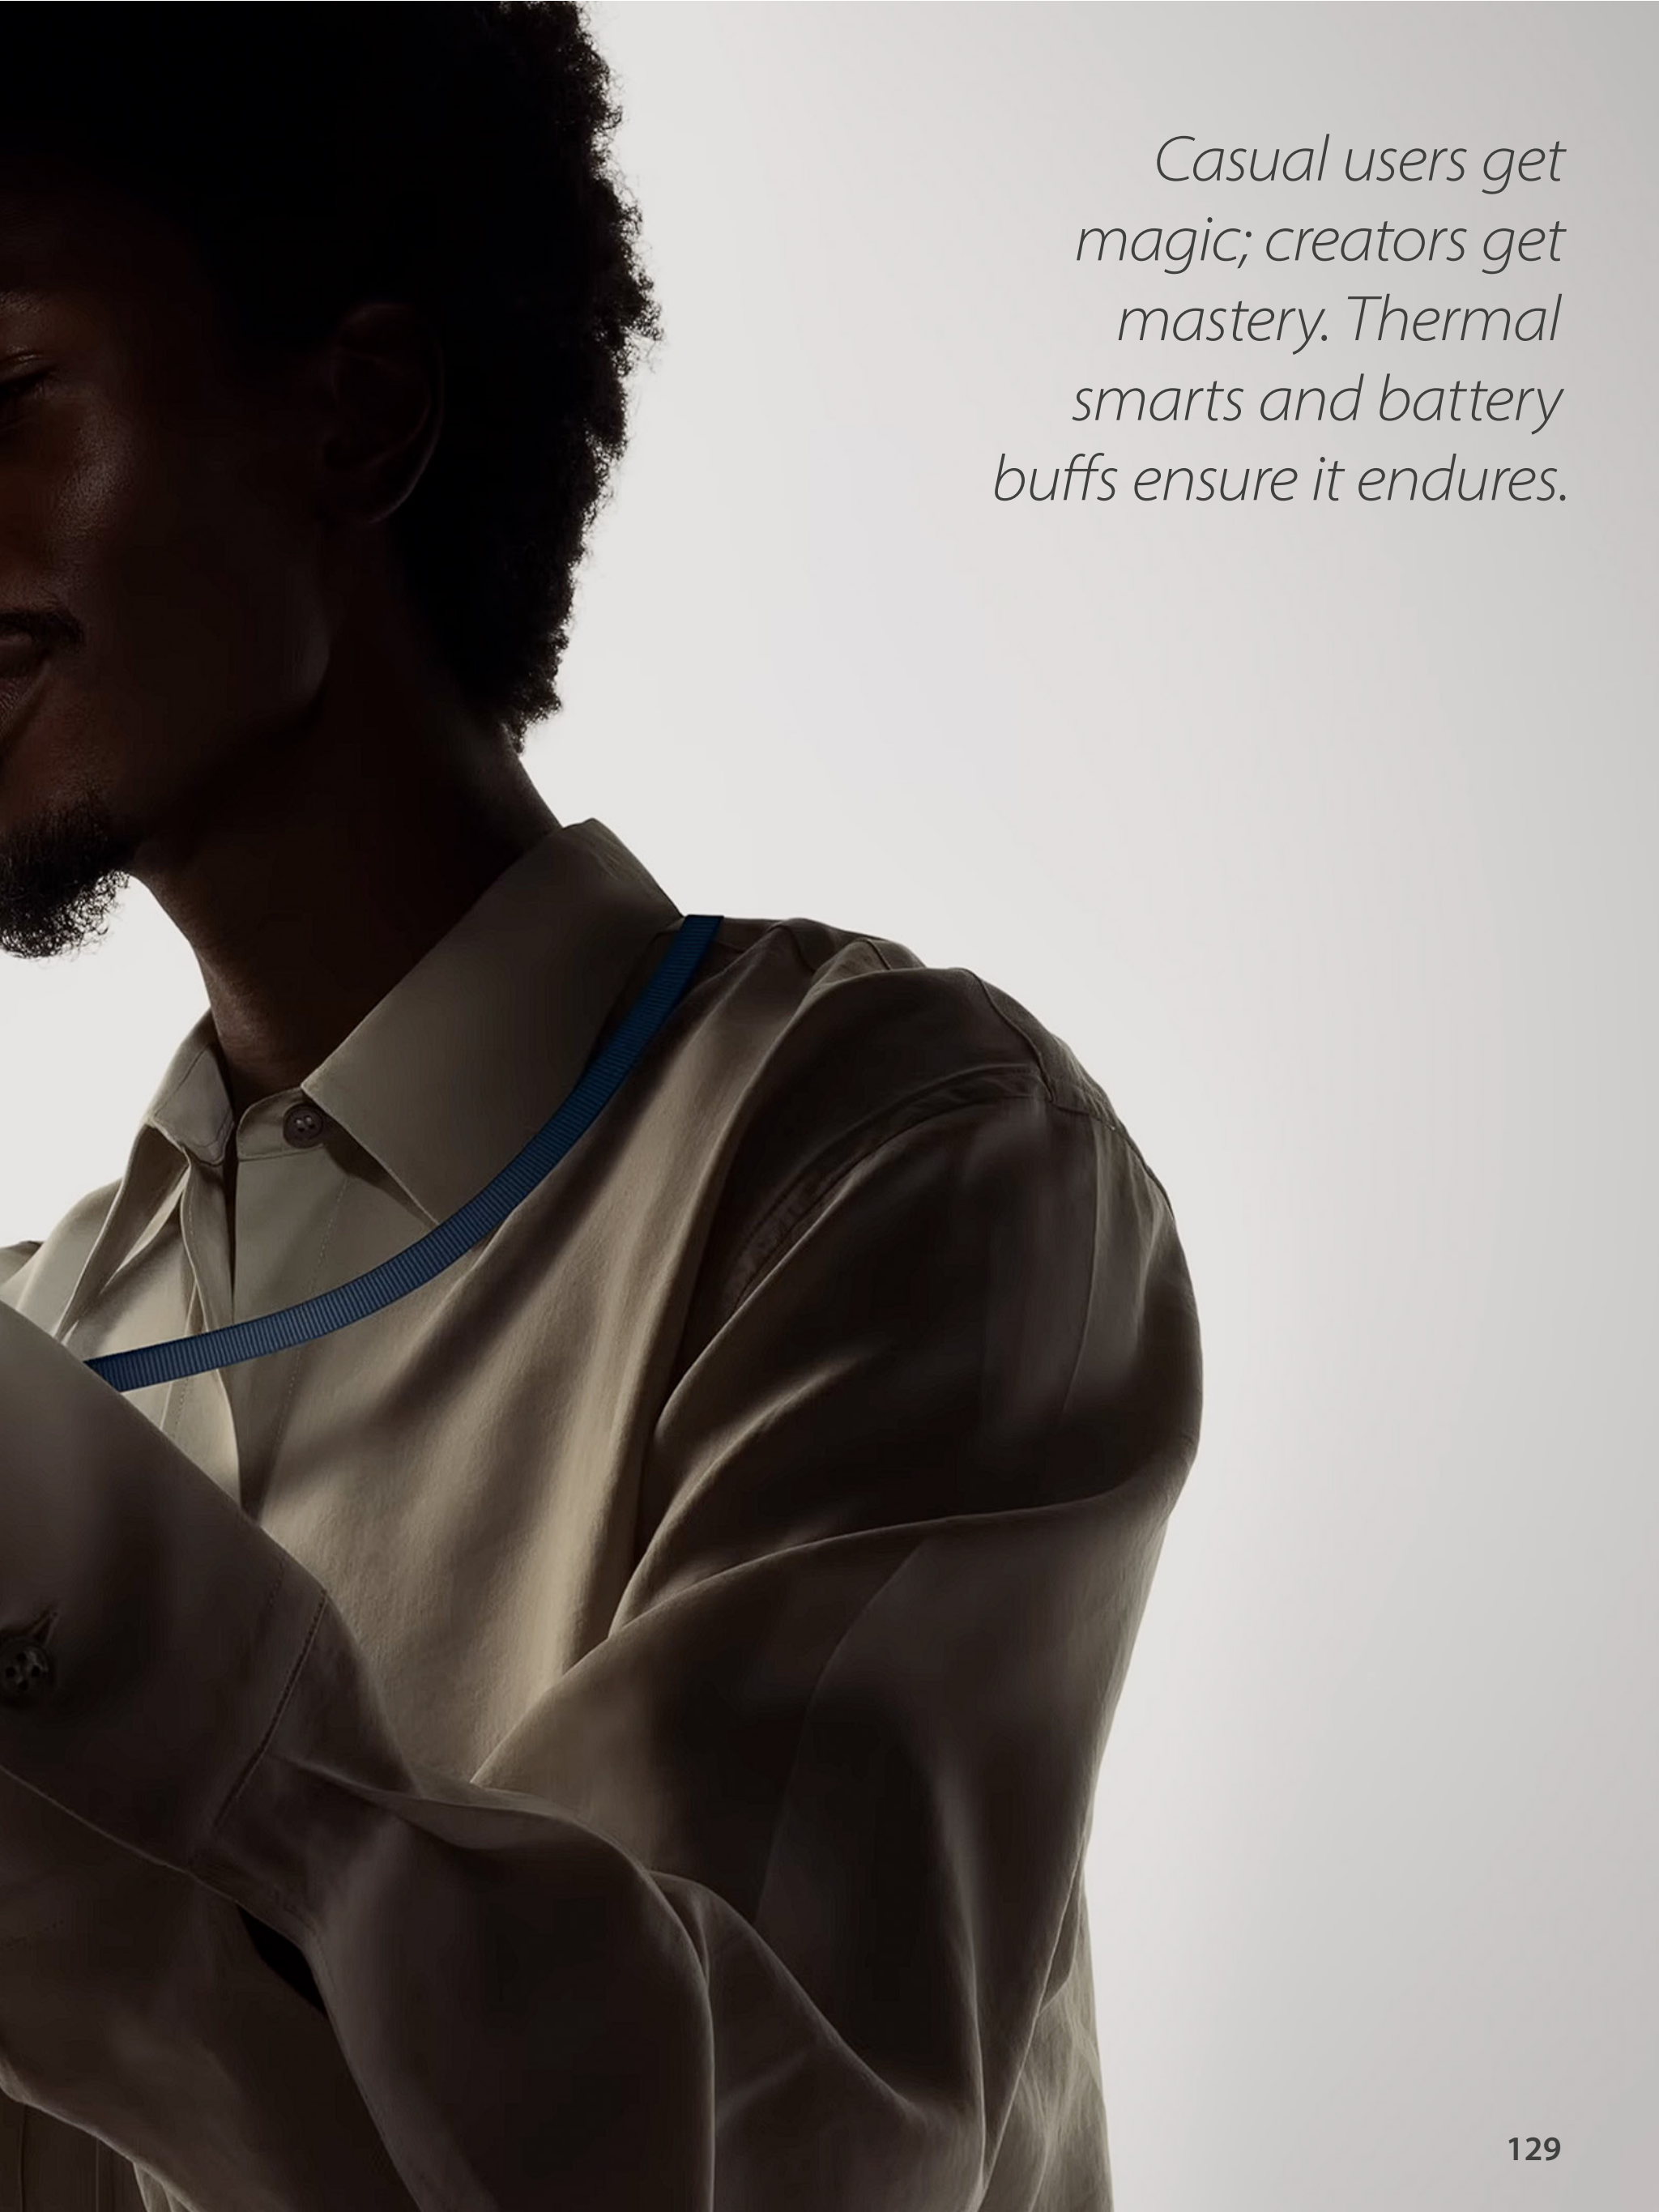

Casual users get magic; creators get mastery. Thermal smarts and battery buffs ensure it endures.

What do I think? It's Apple's boldest camera bet yet, a pocket Arca-Swiss that could upend indie filmmaking. In a world of fleeting trends, this is timeless power.

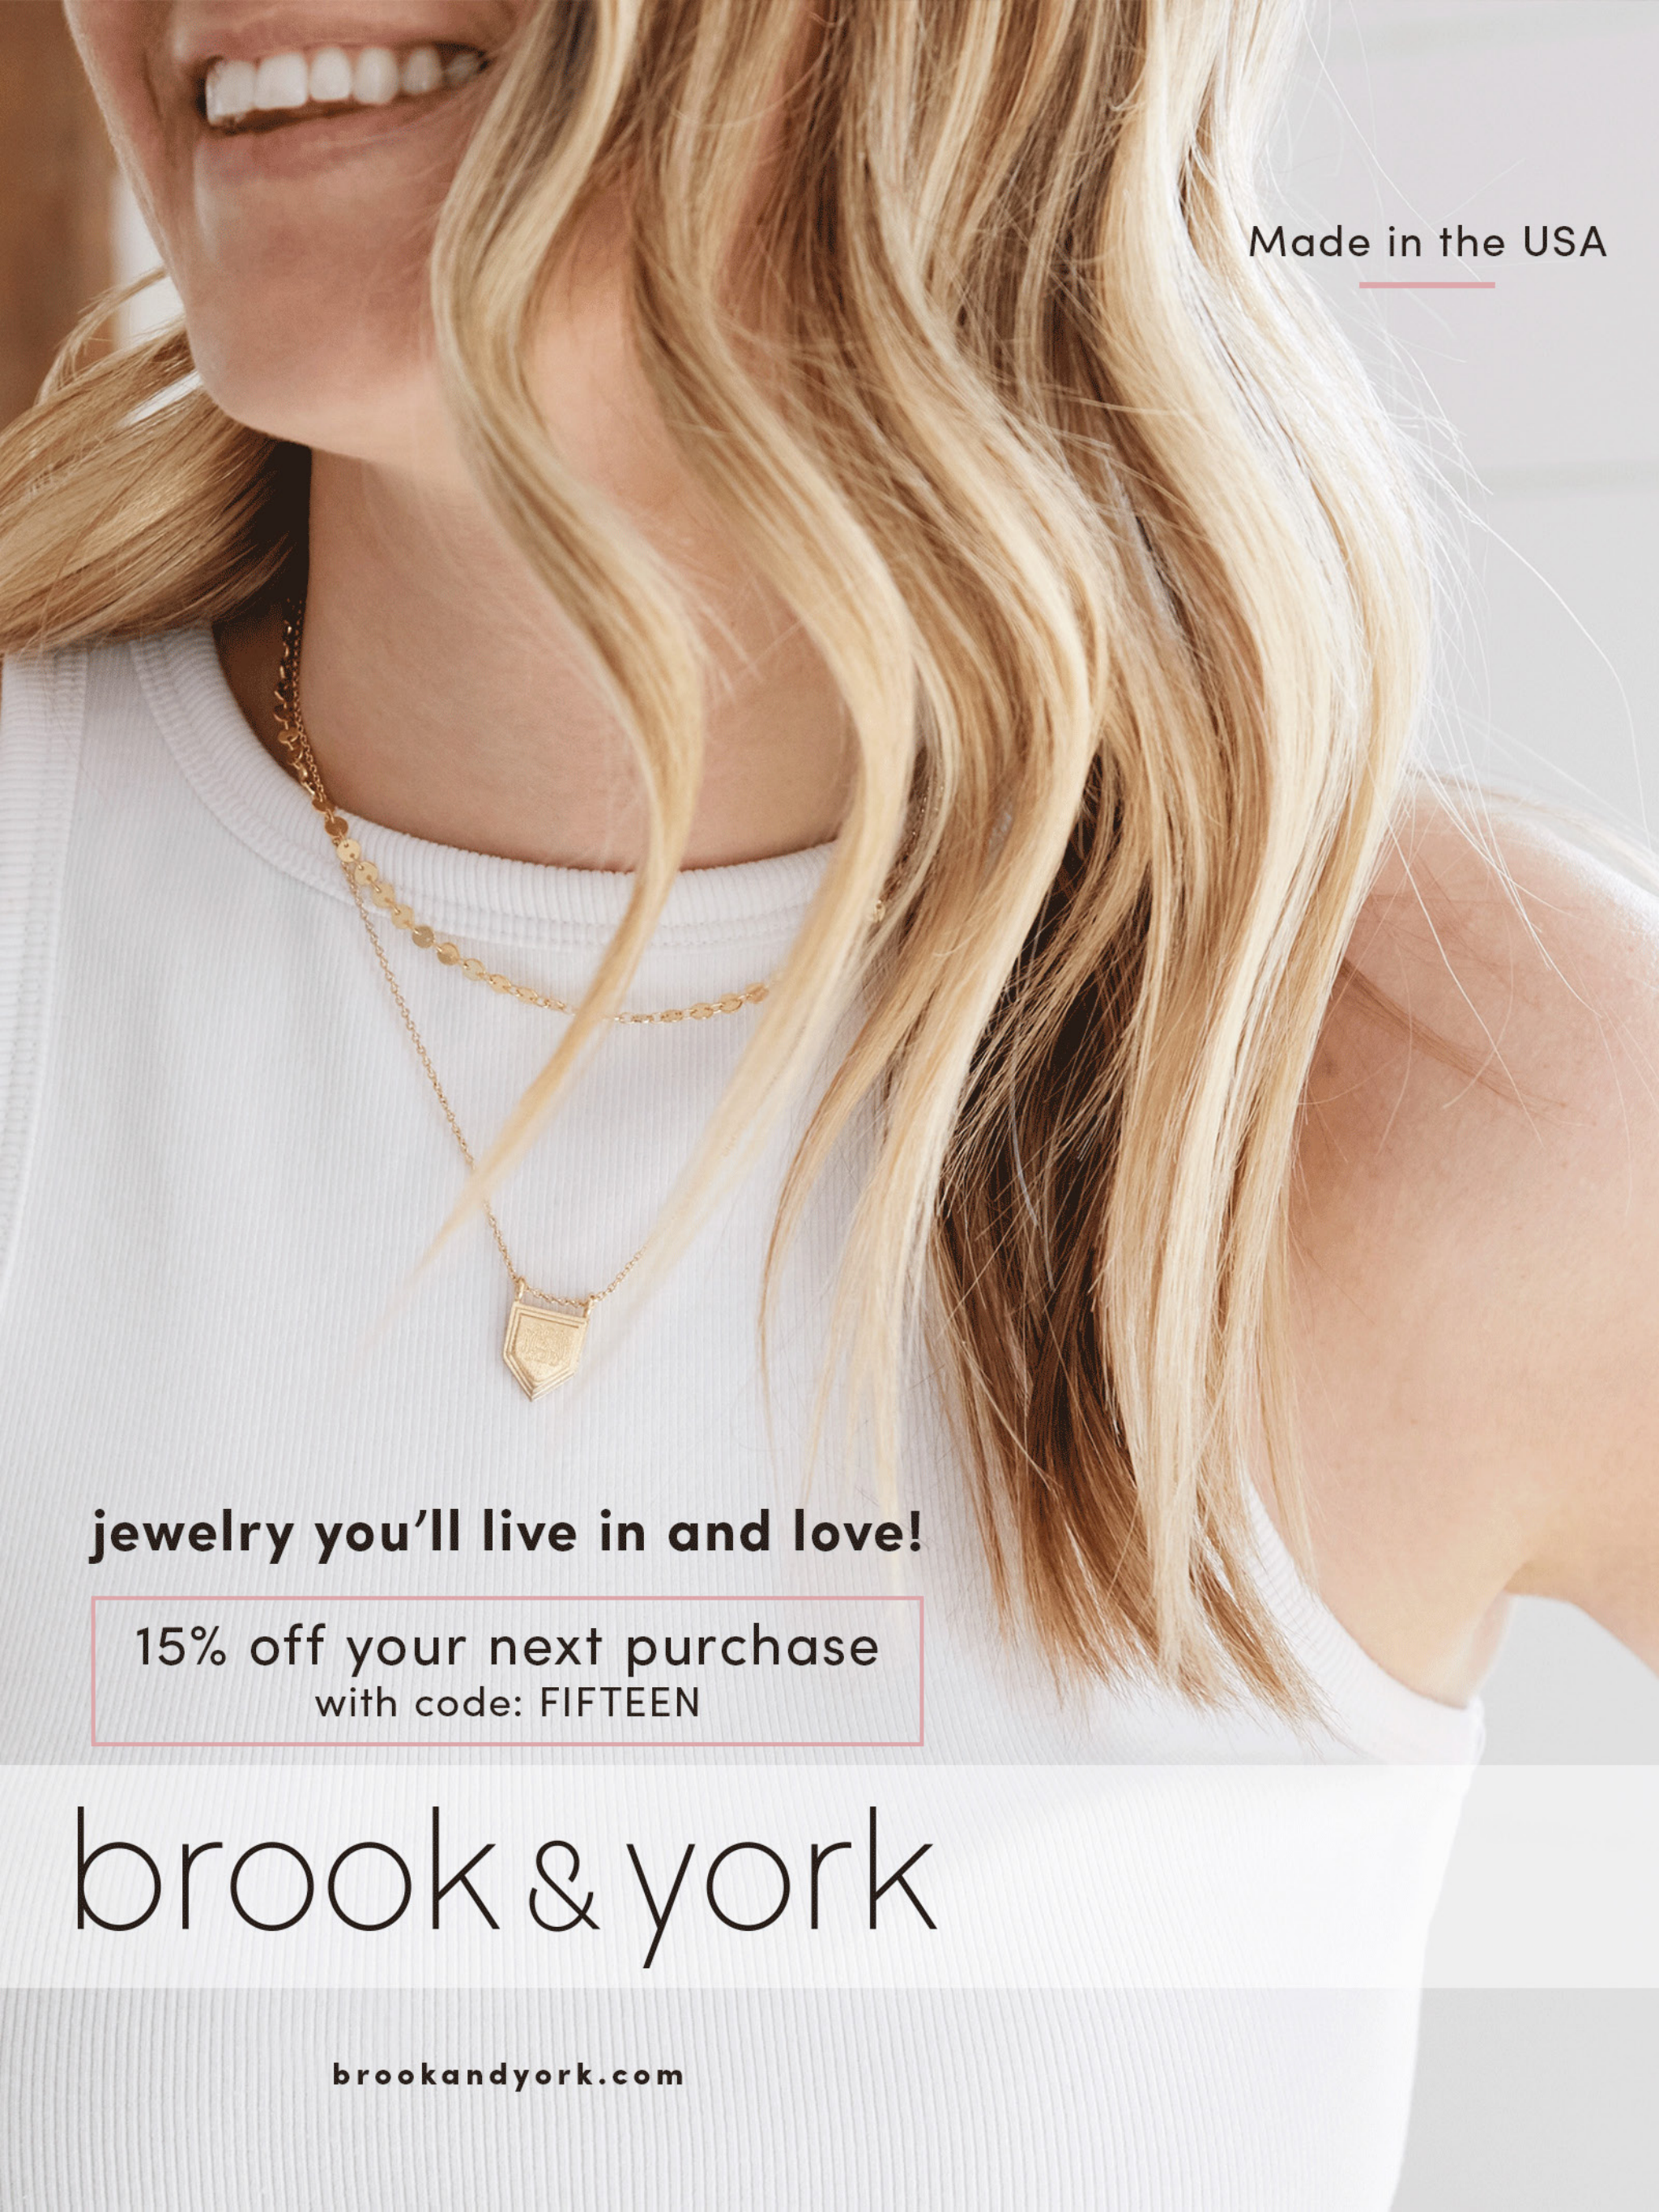
Made in the USA
jewelry you'll live in and love!
15% off your next purchase
with code: FIFTEEN
brook&york
brookandyork.com

fromyouflowers.com

20% OFF SITEWIDE »

All gifts eligible for no contact delivery.

APPLE MAKES RARE JOB CUTS AS SALES TEAMS FACE RESTRUCTURING

Apple has initiated a small round of job cuts affecting portions of its sales organization, marking an uncommon workforce adjustment at a company known for maintaining employment stability through fluctuating economic cycles. The reductions appear limited in scope and are concentrated within teams responsible for channel operations and certain regional sales functions. While the total number of affected roles remains modest, the move is notable given Apple's longstanding practice of avoiding layoffs even during periods of broad industry contraction.

Employees impacted by the restructuring were informed that their roles were no longer required as part of an organizational shift aimed at streamlining internal processes. Apple has not issued public commentary on the changes, but internal communications described the adjustments as part of operational alignment rather than a broad workforce reduction. The company has typically favored reassignment, hiring pauses and attrition during previous downturns, making any instance of direct job cuts unusually visible to industry observers.

The timing of the move coincides with ongoing adjustments in Apple's regional sales and channel strategies, including efforts to refine how the company works with distribution partners, enterprise clients and emerging markets. While the scale of the cuts remains far smaller than those seen across the technology sector in 2024 and 2025, the decision has drawn attention because of its rarity and the fact that it touches a side of the business not usually associated with staff reductions.


Genius Bar

SALES ORGANIZATION SHIFTS AND INTERNAL ALIGNMENT EFFORTS

The affected employees span a subset of Apple's global sales operations, including positions tied to channel coordination and partner management. People familiar with the matter said the changes were linked to the consolidation of overlapping responsibilities within regional divisions. Apple appears to be centralizing certain sales workflows and reallocating strategic planning tasks to fewer teams. While the overall headcount impact is small, the restructuring signals a refinement of operational roles rather than a broad financial retrenchment.

Apple's sales organization supports a global network of carriers, retailers, distributors and enterprise partners. As the company expands manufacturing and distribution footprints in markets such as India and Southeast Asia, parts of the sales structure have evolved to reflect regional differences and new partner relationships. The consolidation seen in this restructuring reflects the company's effort to create consistency across these regions, according to individuals familiar with the process.

Some employees were offered opportunities to apply for other internal roles, which aligns with past practices. Apple has historically prioritized reassignment rather than layoffs when possible. However, in this instance, the company acknowledged that not all affected employees would be placed in new roles, and some were informed that their last day would arrive within weeks.

The restructuring also reflects broader operational adjustments Apple has made in recent years. With product portfolios stabilizing and services revenue continuing to grow, the company has gradually shifted how it manages enterprise accounts, reseller programs and partner ecosystems. Sales teams across multiple regions have increasingly focused on software-and-services-related engagements rather than solely on hardware volumes. This transition may have contributed to the reassessment of certain roles tied to older sales models.

CONTEXT WITHIN THE INDUSTRY AND APPLE'S HISTORICAL APPROACH

The decision to make even small reductions carries significance because Apple has largely avoided layoffs during moments when other major technology firms cut thousands of jobs. In 2022 and 2023, companies across the tech sector undertook large workforce reductions as markets adjusted from pandemic-era demand. While Alphabet, Meta, Amazon, Microsoft and others announced layoffs, Apple held steady, reducing new hiring and slowing expansion without conducting mass job cuts.

Apple's workforce discipline has long been a distinguishing factor. The company historically maintained leaner team structures compared with many peers, making widespread reductions both operationally unnecessary and strategically undesirable. When Apple has made targeted reductions—such as in corporate retail functions or certain engineering support teams—they have generally been absorbed quietly and involved reassignment rather than full separation.

The latest cuts therefore stand out not due to scale, but due to Apple's cultural and historical aversion to layoffs. The affected sales roles represent a narrow slice of the company's global workforce, which exceeds 160,000 employees. Apple continues to hire in other divisions, including machine learning, operations, product design and key engineering disciplines.

Industry analysts examining the move view it as part of broader adjustments in how major technology firms are aligning sales structures to reflect shifts in enterprise demand, product portfolios and regional markets. Apple's business mix has changed considerably in recent years, with services accounting for a larger share of revenue and hardware upgrades following extended consumer replacement cycles.

As Apple expands its enterprise presence with offerings such as Apple Business Essentials and increases its global reach through new retail and reseller networks, the company may also be eliminating redundancies created by overlapping sales channels. This type of organizational refinement is not unusual for a company of Apple's scale, though layoffs remain an atypical method of achieving it.

MARKET ENVIRONMENT, REGIONAL SALES DYNAMICS AND FUTURE OUTLOOK

The broader backdrop for the restructuring includes modest but steady demand for Apple's flagship product lines, a growing services division and a shifting competitive environment in global smartphone markets. While Apple continues to see strength in North America,

Japan and Western Europe, it has also expanded in India and Southeast Asia through new distribution models, manufacturing transitions and localized financing programs. These regional differences have required changes in sales-management structures, particularly in how Apple coordinates with distributors and channel partners.

The restructuring appears aligned with efforts to refine sales strategies in these areas. As Apple strengthens its foothold in India through expanded iPhone production and retail presence, regional sales responsibilities may be shifting from older distribution-led models to hybrid approaches that involve direct oversight from centralized teams. This can reduce the need for certain redundant roles while increasing coordination across partners.

The global smartphone landscape has also introduced pressure on mid-range competitors, contributing to different rhythms in sales performance across regions. Apple's premium positioning has insulated it from some of these fluctuations, but the company's ability to maintain market share has created new expectations for operational efficiency in regional sales divisions. Streamlining parts of the sales organization may reflect an effort to match staffing structures with the company's updated geographic strategies.

At the same time, Apple continues to invest in areas such as marketing, retail expansion, supply-chain operations and machine learning. The workforce adjustments within its sales organization do not appear connected to broader financial strain, and the company's

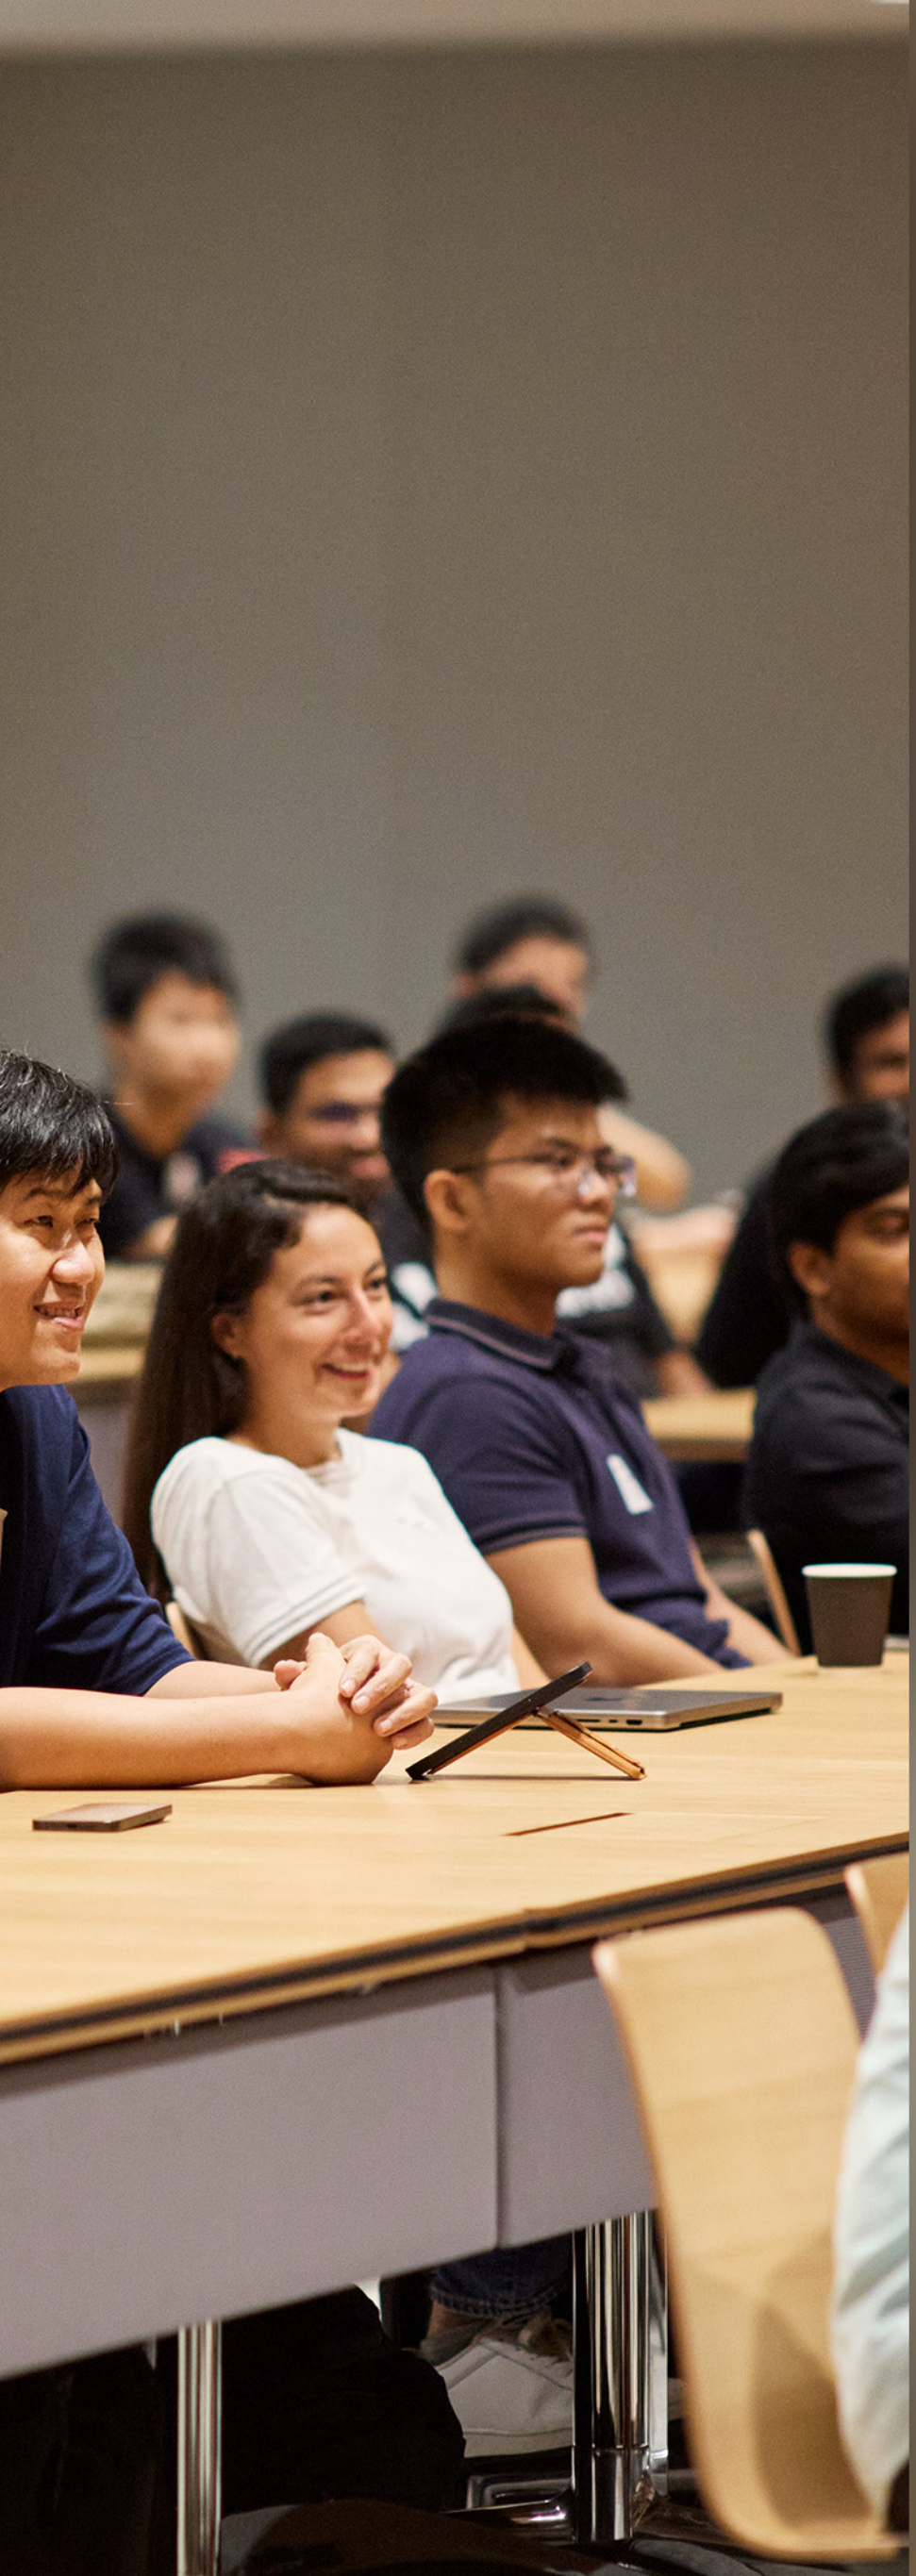

most recent earnings results indicate stable performance across major product categories. Apple's hiring plans in engineering and software categories remain active, supporting long-term initiatives across services, silicon development, AI and device ecosystems.

For suppliers, carriers and retail partners, the changes in Apple's internal teams may lead to updated communication pathways, reorganized account management structures and shifts in how Apple coordinates campaigns or product rollouts. These operational adjustments are part of ongoing evolution in how the company interacts with its global network and are unlikely to affect end-user experiences.

While the workforce shift is small, it signals an operational recalibration within parts of Apple's sales infrastructure. The company's long-term approach remains centered on maintaining lean staffing models while adjusting structures to align with product, regional and ecosystem shifts that continue to evolve across global markets.

PLUSH BEDS

CALIFORNIA

SHOP NOW

AMAZON COMMITS UP TO $50 BILLION TO BUILD AI AND SUPERCOMPUTING CAPACITY FOR U.S. GOVERNMENT AGENCIES

Amazon Web Services is preparing one of the largest infrastructure expansions in its history, pledging to invest up to $50 billion to support artificial-intelligence and high-performance-computing capabilities for U.S. federal agencies beginning in 2026. The announcement reflects Amazon's intent to strengthen its long-standing position in the public-sector cloud market and address the rising demand for specialized compute across intelligence,


DCO DH 1-1
CART 04

Image: Noah Berger

defense, and civilian agencies managing data-intensive workloads. With plans to expand capacity across AWS's Top Secret, Secret, and GovCloud (US) regions, the investment marks a significant escalation in the scale of AI-focused infrastructure dedicated to government use.

AWS said the buildout will add almost 1.3 gigawatts of new compute capacity to these environments, signaling a shift in how much power federal clients may require as artificial-intelligence workloads grow more complex. The company framed the initiative as a multiphase infrastructure program that will incorporate AI-optimized hardware, expanded networking systems, secure data-center environments and integrated access to machine-learning tools already used in the commercial AWS ecosystem. For Amazon, the effort extends a long-running strategy to embed its cloud services into the daily operations of federal agencies, but at a scale that reflects how AI models and high-performance computing are reshaping government-technology priorities.

A SCALE INCREASE FOR CLASSIFIED AND HIGH-SECURITY CLOUD REGIONS

AWS said that the investment will be distributed across the company's controlled-access regions: AWS GovCloud (US), AWS Secret and AWS Top Secret. These regions are designed to host workloads requiring elevated protection, including intelligence analysis, national-security operations, scientific modeling, and classified research. Amazon noted that the expanded infrastructure will support advanced tools such as Amazon SageMaker for model training, Amazon Bedrock for the deployment

of foundation models, and Amazon's own Nova models, alongside hardware that includes AWS Trainium AI chips and clusters of NVIDIA GPUs.

The scale of the proposed 1.3-gigawatt buildout indicates a much broader set of use cases than traditional cloud storage or elasticity-based workloads. Federal agencies have increased their reliance on large-scale simulations, geospatial modeling, cybersecurity analytics and other workloads that demand dense compute environments. According to AWS, agencies will be able to consolidate these systems rather than operate smaller, siloed hardware, which often requires significant maintenance and long procurement cycles. Centralizing these resources in high-security cloud regions may reduce the operational friction associated with managing distributed hardware fleets.

AWS executives described the investment as a response to persistent capacity constraints and growing demand for environments capable of supporting larger models. While earlier cloud regions were built around virtual-machine deployment and secure storage, the next phase is being designed with optimization for AI clusters, low-latency networking, and specialized accelerators. Amazon said that as model sizes increase, agencies require consistent access to hardware that can support long training cycles and high-bandwidth inferencing without competing for commercial-region capacity. The dedicated nature of these regions allows AWS to isolate workloads, apply tailored security controls, and deliver predictable compute availability.

Image: Stephane Mahe


230 230
Precist 0.001 us
Stabil 0.225 ppm

The buildout will likely require significant expansion of physical facilities, including new power arrangements, cooling systems and interconnect infrastructure capable of handling AI-optimized clusters. AWS did not specify exact site locations but noted that construction, installation and deployment would unfold across multiple states. Each region will remain separated from the public AWS cloud and will continue to follow the compliance standards set out by U.S. agencies overseeing classified-data environments.

IMPLICATIONS FOR FEDERAL AI ADOPTION AND COMPETITION IN CLOUD SERVICES

The scale and timing of the investment reflect how federal agencies are reorienting around AI as a core operational tool. Government organizations have increasingly explored uses that range from large-scale document analysis to satellite-imagery processing and advanced logistics modeling. Many of these workloads require compute that cannot be supported efficiently by older systems or by on-premises hardware that lacks the scalability of cloud-based infrastructure. AWS said the expansion could reduce the time needed for certain simulations from weeks to hours, depending on how agencies adopt and configure the new resources.

For the federal government, the announcement signals an opportunity to standardize AI development environments across agencies that have historically relied on a mix of legacy systems and vendor-specific solutions.

Agencies with separate data-center footprints often face inconsistent hardware availability and high maintenance costs. The unified framework offered by AWS's secure regions may reduce fragmentation and create a common foundation for model training, testing, and deployment. It also creates a pathway by which large agencies can run multiple generations of AI models concurrently, eliminating bottlenecks associated with resource allocation.

The infrastructure expansion occurs at a moment when AWS faces rising competitive pressure in the public-sector cloud market. Google Cloud and Oracle Cloud have recently increased investment in AI-focused infrastructure for government clients, emphasizing secure-region offerings and specialized hardware. Amazon's $50-billion commitment reinforces its intent to remain the dominant provider of classified and high-security cloud capabilities, a position it has held for more than a decade. Analysts following the industry consider this investment part of a broader trend in which cloud providers seek to define long-term relationships with agencies that increasingly rely on AI-enhanced workflows.

AWS's long-standing base of more than 11,000 government customers provides a significant advantage, but the landscape has shifted as AI-driven requirements have grown more demanding. Compute-intensive workloads associated with model training and inferencing require far more power than traditional databases or line-of-business applications. The investment suggests that AWS views these workloads as durable drivers of demand rather than temporary trends.


CATCHER SIDE
PDC-1.1A4
WARNING
FAN NO.5
FAN NO.10

Image: Jonathan Raa

It also signals to agencies that Amazon intends to support multi-year modernization programs that rely heavily on advanced compute.

The partnership implications extend beyond infrastructure. Agencies may increasingly look to integrate proprietary and open-source models from commercial developers through AWS services. The expansion of Amazon Bedrock into secure regions, along with support for third-party models from partners such as Anthropic, could provide agencies with a more diverse set of tools for building applications. This approach may further shift how government software is developed and maintained, integrating generative-AI capabilities into processes such as risk assessment, logistics analysis, and internal knowledge retrieval.

HOW THE INVESTMENT INTERSECTS WITH NATIONAL STRATEGY AND INDUSTRY ECOSYSTEMS

The commitment aligns closely with national policy objectives that identify AI as a strategic capability requiring secure, domestically controlled infrastructure. As global competition around AI intensifies, the U.S. government has emphasized the need for scalable, resilient compute environments designed to manage sensitive and classified datasets. Cloud regions with robust isolation measures provide a mechanism for agencies to handle these tasks while maintaining compliance with federal security standards.

The scale of the planned infrastructure introduces ripple effects for companies connected to the AI-compute ecosystem.

Suppliers of AI chips, networking hardware, data-center components, fiber-optic technology and facility-construction services may see sustained demand as AWS ramps up deployment. The volume of hardware necessary to power 1.3 gigawatts of new capacity indicates long-term procurement pipelines for accelerator chips, storage subsystems and system integration, many of which operate under multi-year contracts. As federal agencies adopt more complex AI models, vendors may increasingly be required to certify hardware through government-specific review processes, adding further depth to the supply chain.

At the policy level, the effort may influence how agencies budget for modernization. Multi-year infrastructure investments give agencies a consistent environment for long-term planning, reducing the need to build internal data centers or compete for commercial cloud resources. Agencies that previously relied on hybrid arrangements combining on-premises systems with limited cloud deployments may expand into fully cloud-based approaches as AWS's capacity grows.

For commercial customers outside the federal space, AWS's government-focused investment may shape expectations about future AI offerings. While the infrastructure is not shared with commercial regions, Amazon's technical advancements in high-performance clusters, model-deployment optimization and power-efficient data-center design often flow downstream into the public cloud. Enterprise buyers may watch this investment as an indicator of how Amazon intends to evolve its AI services more broadly. However, pricing, availability and hardware composition will differ because of the classified nature of government workloads.

The broader context is one in which AI infrastructure is increasingly treated as a strategic asset comparable to transportation networks or energy systems. The scale of Amazon's planned expansion reflects how federal agencies now rely on AI to manage functions that extend from national security to public welfare administration.

Taken together, the investment suggests that AI and high-performance computing have become structural components of government operations, prompting cloud providers to build capacity that can match long-term federal requirements.

OPENAI AND JONY IVE RAMP UP HARDWARE PUSH AS APPLE LOSES ENGINEERING TALENT

OpenAI and designer Jony Ive are advancing their collaborative hardware initiative amid growing momentum and a reported exodus of engineering talent from Apple Inc.. The project, which involves an “elegantly simple” device described by the team in public remarks, is now confirmed to have functioning prototypes in development. At the same time, public filings and multiple media reports suggest that OpenAI has recruited over forty senior hardware engineers from Apple in approximately a month, underscoring the seriousness and urgency of OpenAI’s entry into consumer-hardware territory.

In a recent interview moderated by Laurene Powell Jobs at the Emerson Collective Demo Day, Ive and OpenAI CEO Sam Altman confirmed the existence of working prototypes and indicated a target of under two years for a market introduction—a timeline somewhat accelerated from earlier expectations. Ive described the device as stripped of extraneous complexity, and Altman suggested it offers deep personalization, referencing the concept of a system “knowing everything you’ve ever thought about, read, and said.” At the same time, the influx of Apple hardware engineers to OpenAI has raised internal concern at Apple, where the shift is viewed as a potential competitive threat as the company pursues its own AI hardware ambitions.

PROTOTYPE STATUS AND DESIGN AMBITION

According to comments made during the public interview, the hardware effort has moved beyond concept and into advanced prototype stages. Ive stated the objective was to create a device "you want to use almost without thought," invoking a design ethos focused on simplicity and emotional appeal. Altman supported this framing by indicating a belief that previous devices lacked the intuitive quality the team is now targeting. The unit's form factor remains undisclosed, but speculation suggests something smaller than today's typical smartphone and possibly without a conventional screen interface. The emphasis is on ambient intelligence and minimal distraction.

Behind the scenes, the hardware team has reportedly expanded rapidly. Media coverage notes that OpenAI's devices group has hired more than forty people in a month, with recruiters targeting senior directors, camera-engineering specialists, silicon-design experts, device-testing and manufacturing engineers—all drawn from Apple's hardware ecosystem. Apple insiders cited late-2025 attrition among key hardware teams focused on iPhone, Mac and wearable devices. The talent movement signals that OpenAI is assembling a high-volume, high-excellence hardware pipeline rather than a side project, addressing consumer ambitions rather than enterprise only.

The acquisition path has reinforced that view. Earlier in 2025, OpenAI acquired "io Products, Inc.," the startup co-founded by Ive and several former Apple executives, in a deal valued

around $6.4-$6.5 billion. This deal brought in design talent, formerly independent hardware teams and a symbolic transfer of momentum into OpenAI's hardware ambitions. These developments coincide with Apple's own reported delays in internal AI hardware efforts, increasing the perceptual gap between the two companies.

IMPLICATIONS FOR APPLE, OPENAI AND THE HARDWARE ECOSYSTEM

For OpenAI, the device strategy represents a shift from pure-software dominance into physical product territory. With hardware built around AI agents rather than apps, the company aims to create a new category distinct from smartphones and traditional computers. The timing—under two years from public prototype—reflects an appetite to accelerate the transition to consumer-facing AI hardware while competitors are still defining their approaches.

Apple, meanwhile, faces an exacerbated challenge. The reported outflow of hardware engineers comes while the company is progressing several hardware-and-software initiatives tied to its AI strategy, including smart home devices, mixed-reality accessories and robotics. Losing talent with CMOS, camera, and manufacturing expertise may disrupt those roadmaps or raise the execution risk. Apple's leadership may need to move quickly to retain critical skillsets and maintain its hardware innovation lead.

EMO
AY

Emerson Collective

nAI

The hardware supply chain stands to feel broader shifts as well. If OpenAI's device achieves scale, manufacturing partners formerly tied heavily to Apple could diversify. Equipment vendors, component suppliers and contract manufacturers may begin aligning with the new entrant's needs, potentially reshuffling long-standing relationships in consumer electronics. Established OEMs accustomed to working with Apple's cadence could face competitive pressure to meet OpenAI's timeline and volume demands.

For consumers, the hardware announcement raises questions about what "AI-first" devices will look like. The design framing emphasizes minimalism, ambient intelligence and integration of large-language models into everyday interaction. If successful, the product may redefine expectations for how hardware, AI software and interaction combine—but open issues remain around privacy, power consumption, pricing and the ecosystem environment. Separation from screen-based models implies a shift not only in design but in how users engage with intelligent systems.

These moves add urgency to the hardware race among major technology firms. With OpenAI and Ive mobilizing rapidly, companies like Apple, Google, Meta and others may accelerate their own efforts to build AI-optimized devices. This dynamic could shorten product-development cycles for hardware that integrates advanced AI, increasing competition for talent, manufacturing capacity and consumer attention.

IVE AND ALTMAN REVEAL FIRST PROTOTYPE OF MYSTERIOUS AI HARDWARE

During a public interview at Emerson Collective's Demo Day, designer Jony Ive and Sam Altman, CEO of OpenAI, confirmed they have developed the first prototypes of their new AI hardware project and aim to bring the product to market within two years. Moderated by investor and philanthropist Laurene Powell Jobs, the discussion offered rare insight into one of the most secretive collaborations in tech today—though the specific features of the device remain largely unknown.

The announcement follows earlier reports of OpenAI's acquisition of Ive's design startup io Products for approximately $6.5 billion. During the on-stage conversation, both Ive and Altman emphasized minimalist design architecture and a user experience that differs from conventional smartphone or computer interfaces. While skeptical observers awaited detail on functionality, the pair emphasized the product's role in redefining how users engage with AI rather than simply extending existing platforms.

In describing their progress, Altman said they believed the design had reached a breakthrough when the prototype evoked an instinctive desire to pick it up—or "take a bite out of it," in Ive's phrase. He added that previous versions lacked that intuitive appeal. Ive elaborated that his design goal was crafting a device "you feel no intimidation about, you want to use almost without thought, that it's just a tool." The session closed with the timetable: they expect a market-ready product "in less than two years."

PROTOTYPE STAGE AND DESIGN PHILOSOPHY

The hardware development is still in early production stage but has advanced beyond concept. The duo said they have multiple working prototypes internally and are now refining design, materials and interaction paradigms. While they declined to share dimensions, clues suggest a device roughly the size of a smartphone but potentially lacking a traditional screen, aligning with speculation that the product might use ambient display or projection technologies rather than rely on conventional form factors.

Image: Kyle Grillot

Ive, formerly chief design officer at Apple Inc., drew upon his past hardware work but said this project required a distinct approach. The pair described an emphasis on direct interaction with AI—less distraction, more ambient assistance. Altman described current smartphones as "walking through Times Square, all the distractions it produces," suggesting their device aims for a calmer interface. Ive noted that ceramic and other advanced materials were under consideration, pointing to design research focused on touch, weight, and tactility.

This hardware effort reflects OpenAI's broader ambition to move from software-only models into tangible devices. The challenge of aligning hardware, software, and design at such a scale has attracted major attention across the industry. Ive and Altman referenced the need to engineer a system where hardware doesn't get in the way of the AI experience, but supports it. "If you make something new, there will be consequences unforeseen," Ive said, adding that he feels a responsibility given his role in earlier device ecosystems.

STRATEGIC AND MARKET IMPLICATIONS

The announcement signifies that OpenAI and its design partner are positioning themselves to compete beyond cloud-based AI models and toward consumer-facing hardware. While the device remains unnamed and its ultimate features unconfirmed, the timetable and internal prototypes suggest the project is more serious than typical rumor. For the hardware ecosystem, the move indicates a potential shift in how AI hardware might integrate into everyday life—moving from apps on a phone to dedicated devices optimized for AI interaction.

For hardware suppliers and the tech supply chain, the project could influence component sourcing, manufacturing strategies and competitive dynamics. If OpenAI delivers a new category device that taps into AI-native interaction rather than legacy PC or smartphone paradigms, it may accelerate investment in new materials, sensor arrays, interaction modules, and design tooling. The choice of partner designers and the level of secrecy suggest the project is a strategic priority rather than an experiment.

Brand-wise, the involvement of Jony Ive adds symbolic weight to the effort. Ive's association with iconic products such as the iPhone and MacBook raises expectations around design quality and product experience—but also increases scrutiny. The fact that Laurene Powell Jobs moderated the session and has previously backed design and innovation firms lends additional gravitas and signals investor confidence in the initiative's trajectory.

Sam Altman and Jony Ive with Laurene Powell Jobs

At the same time, the competitive landscape in consumer AI hardware is crowded and uncertain. Previous efforts from companies seeking dedicated AI hardware have faced challenges in achieving mass adoption, aligning ecosystem partners, and defining clear value propositions. By moving the timeline to under two years, the project faces a tight window to mature, partner, supply, and launch ahead of both hardware and software competition.

The significance for OpenAI extends beyond one product. If the hardware succeeds, the company may have a physical platform for its models—a device that could tie directly to its AI ecosystem, differentiate from competitors, and create a new category of hardware-plus-AI service. For users, a shift away from screens and toward ambient or dedicated AI interaction could redefine how computing is accessed and consumed.

META'S WHATSAPP SEES DECLINE IN THIRD-PARTY AI CHATBOT ACTIVITY AS DEVELOPERS SHIFT TO IN-APP SYSTEMS


Meta

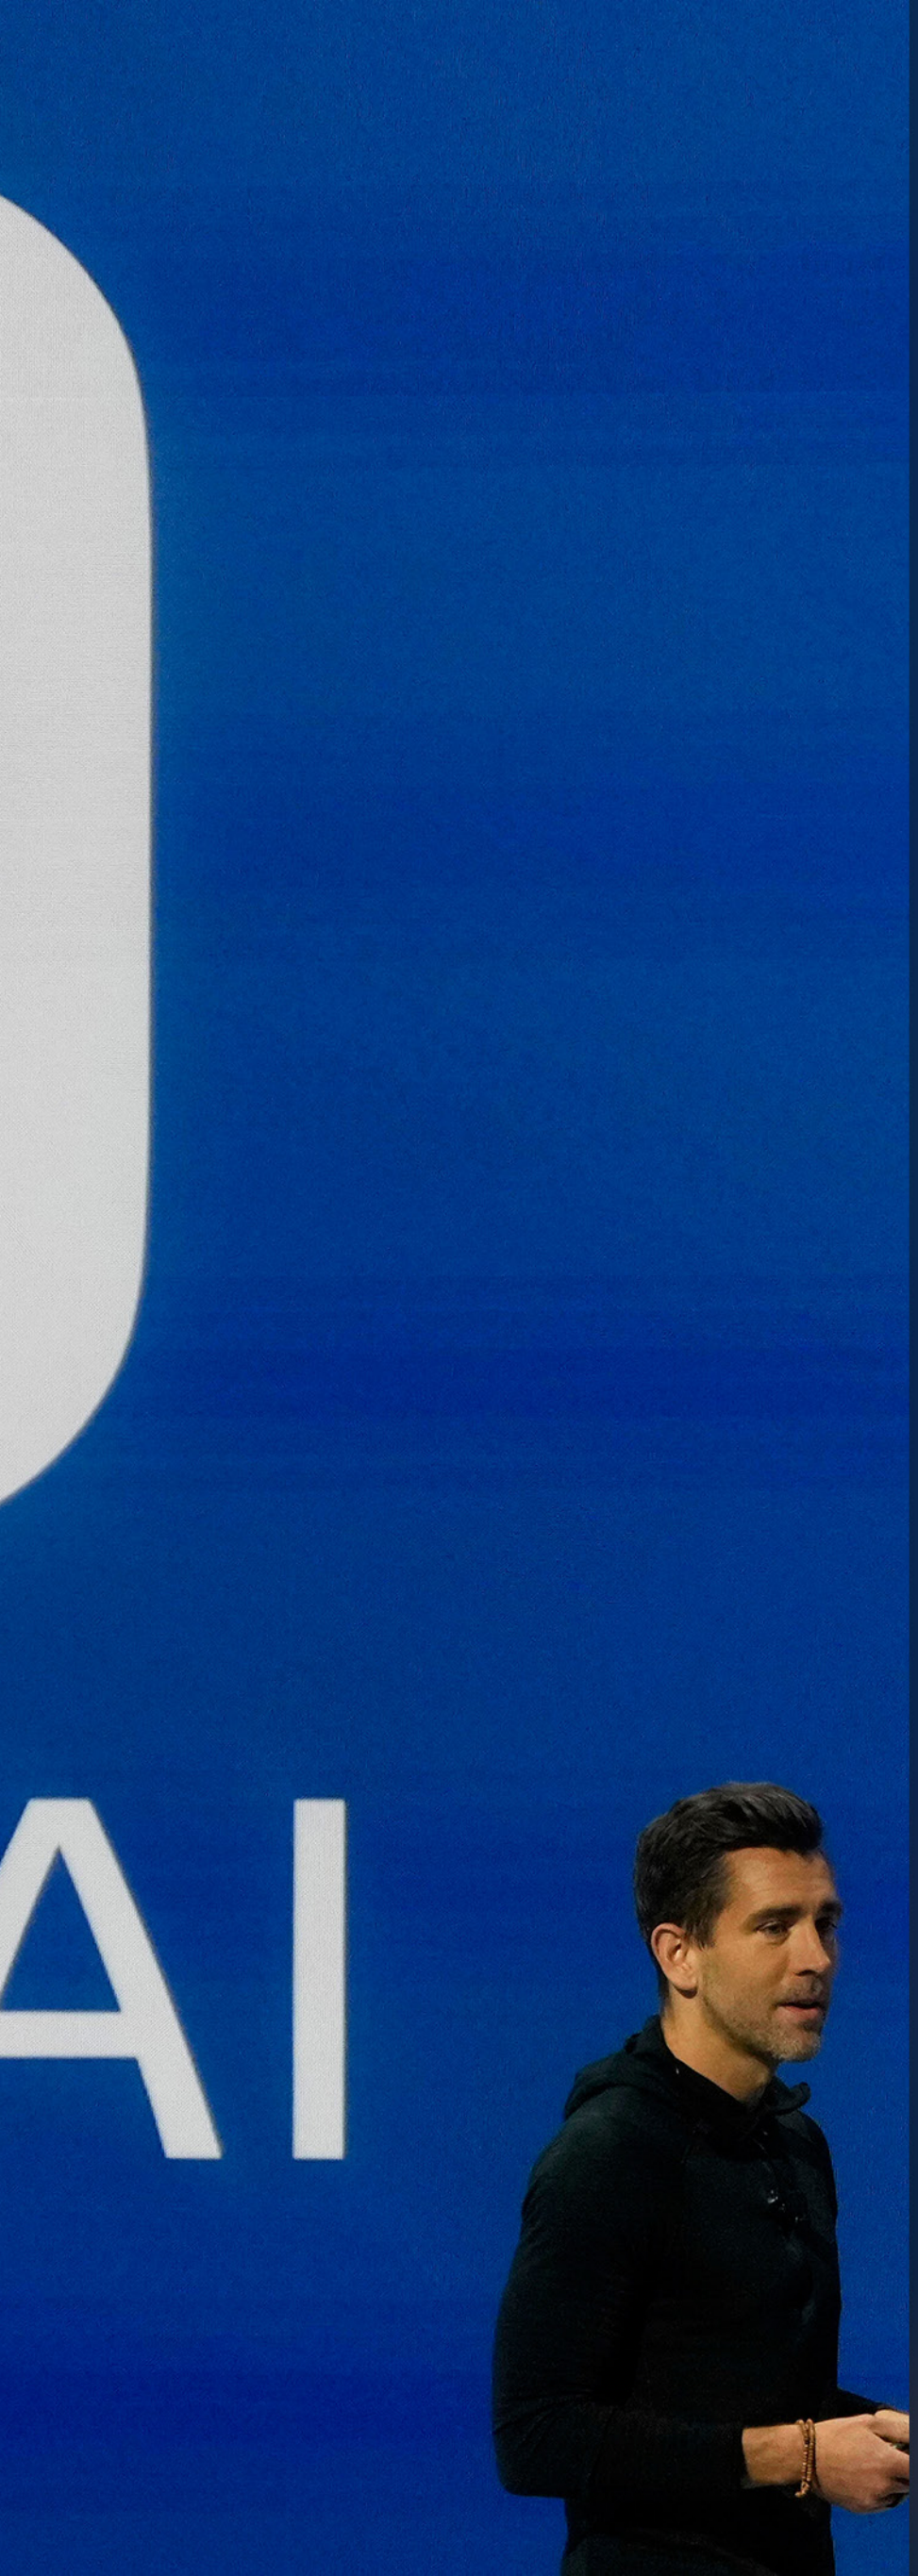

WhatsApp is experiencing a drop in activity from third-party AI chatbots built on external large language models. The decline reflects how builders of conversational systems are increasingly redirecting their efforts toward Meta's native assistant framework rather than relying on external integrations accessed through WhatsApp's existing bot interfaces. The trend is emerging at a time when messaging platforms across the industry are reassessing how third-party AI services operate inside closed communication ecosystems.

Developers who previously used WhatsApp as a channel for distributing conversational bots now report lower engagement as Meta expands access to its own assistant across the WhatsApp interface. These builders have noted that usage patterns began shifting as Meta consolidated its AI features within the app, giving users a direct way to access built-in tools for summarization, Q&A, drafting and media generation. Several creators told The Verge that the new structure reduces the need for external chatbots routed through unofficial or semi-supported entry points, changing how WhatsApp fits into the broader AI-assistant ecosystem.

The shift has created a measurable redistribution of activity among developers who were early adopters of WhatsApp as an AI interface. While some continue to maintain their bots for existing user bases, others report that they are focusing on integrations in environments where third-party models retain clearer access paths, such as standalone apps, web clients or enterprise-facing workflows. The change also reflects Meta's broader effort to integrate AI features directly into its products at the interface level, diminishing reliance on external connectors.


WhatsA

DEVELOPERS REASSESS WHATSAPP AS META EXPANDS ITS BUILT-IN AI ASSISTANT

Meta's integration of its own AI assistant into WhatsApp has created a new interaction model that intersects with how third-party chatbots previously operated. Built-in assistant features run within WhatsApp's main input field and can be invoked spontaneously across chats, giving users immediate access to tasks that external bots once provided. This shift has made third-party bots less central to common use cases, according to developers who track engagement metrics.

Several builders described how the introduction of Meta's assistant reoriented user behavior. Tasks such as translation, content drafting or question answering, which previously required navigating to a separate bot conversation, can now occur directly within the app's interface. For developers relying on WhatsApp as a distribution layer, this reduces the interactions that once sustained user retention. As Meta continues rolling out upgrades to the assistant, including multimodal inputs and more contextual interactions, developers expect this trend to deepen.

The move has implications for how WhatsApp functions as a channel for external AI tools. WhatsApp was once used widely as an informal interface for third-party models because of its global reach and low barrier to entry for bot deployment. Developers could connect bots to large user bases, especially in regions where WhatsApp is a primary communication tool. But as Meta tightens integration with its own systems, the platform is shifting toward

a model where bot developers must rely on Meta's official business tools rather than consumer-facing conversational entry points.

These changes align with Meta's approach across its messaging platforms. Messenger and Instagram have also received deeper AI features, moving toward a coordinated strategy that positions Meta's assistant as a central entry point for AI interactions. As developers encounter these changes, many are reconsidering how much engineering investment to allocate to WhatsApp, focusing instead on environments where external models maintain fuller control over user interaction pathways.

MODEL PROVIDERS SEE ENGAGEMENT DIVERGENCE ACROSS PLATFORMS

The decline in WhatsApp-based third-party bot usage has contributed to a broader pattern observed by providers of large language models. Several model developers have reported uneven engagement levels depending on how messaging platforms structure access. In environments where native assistants dominate the interface, third-party bots experience less organic discovery and lower sustained usage. Conversely, platforms with open bot frameworks or direct API access continue to generate more stable activity.

Developers working with multiple platforms say WhatsApp's shift has contributed to a redistribution of user attention rather than an overall drop in model usage. Traffic that previously flowed through WhatsApp bots is now redirected to web apps, standalone interfaces or alternative messaging platforms where integration points

remain open. This redistribution aligns with how AI developers often adapt to platform constraints by shifting their tools to environments where they can control functionality, user onboarding and data handling policies more fully.

The divergence is also affected by regional patterns. In countries where WhatsApp is the dominant messaging service, third-party developers had historically treated it as a primary deployment target. As Meta's assistant gains adoption, these developers report that customer onboarding and retention now depend on strategies outside WhatsApp, such as companion web apps or mobile installations. In markets where multiple messaging services compete, developers have shifted more quickly between platforms depending on access conditions.

The dynamic reflects a broader tension across messaging apps as they expand in-app AI features. Companies deploying their own assistants increasingly position these capabilities as core product elements, reducing the need for users to rely on external models. Developers who built earlier integrations are now responding to ecosystems that limit visibility for third-party tools unless deployed through business-focused frameworks rather than consumer chat environments.

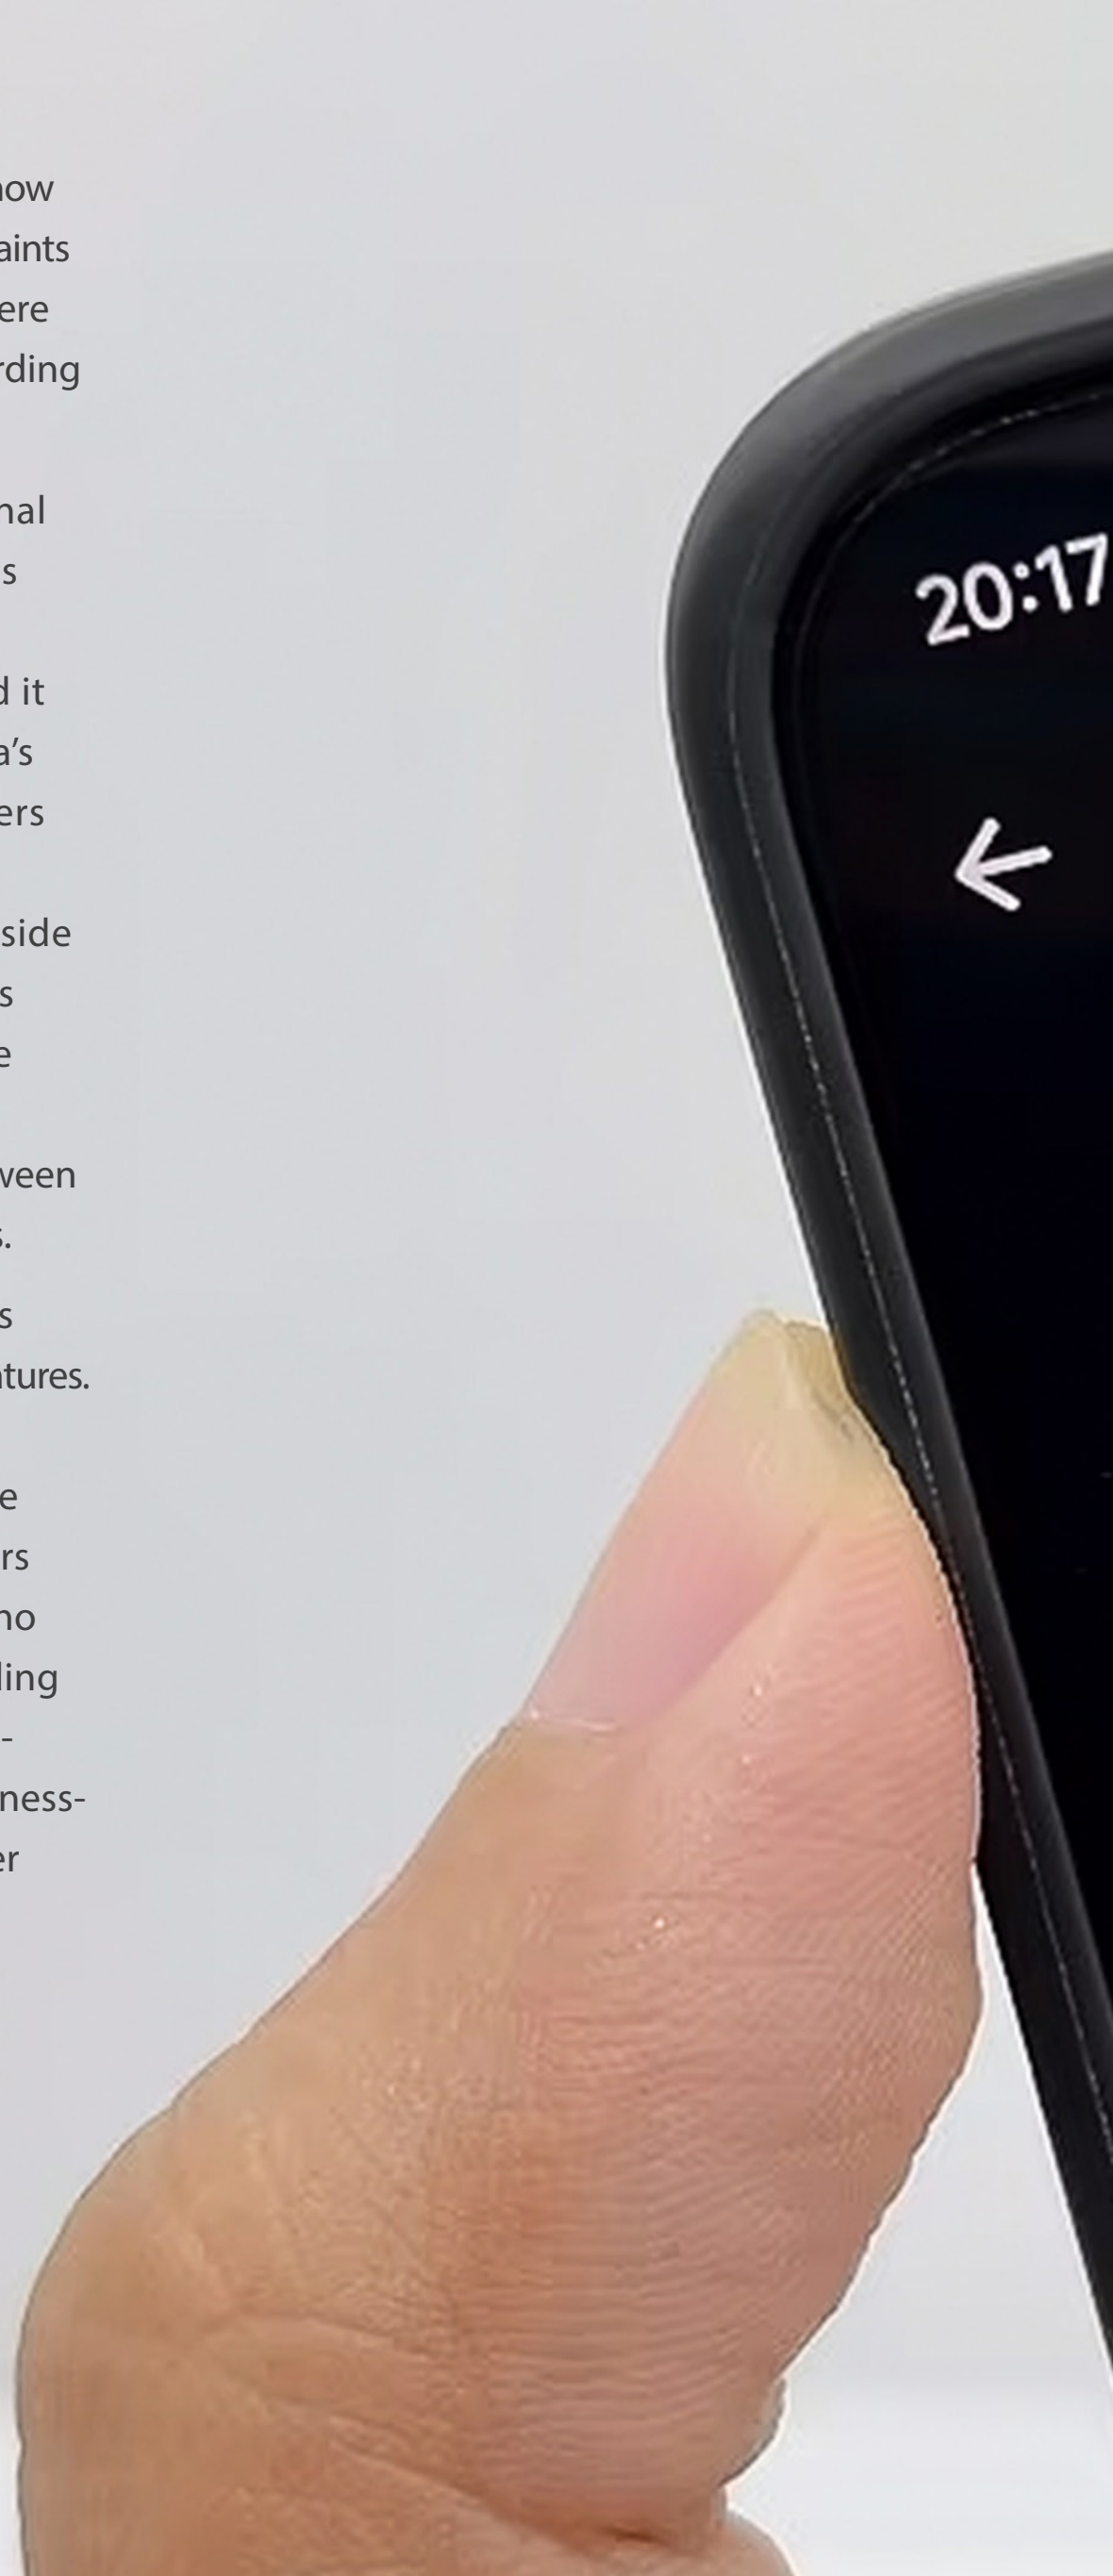

0.00
KB/s

Meta AI
with Llama 3.2

Share

Message

Messages are generated by AI. Some
inaccurate or inappropriate.

WHATSAPP'S AI LANDSCAPE EVOLVES AS META'S ASSISTANT EXPANDS

WhatsApp's changing environment illustrates how the integration of platform-native AI assistants shifts the ecosystem for developers. As Meta refines its assistant and expands capabilities across languages and regions, usage patterns continue to evolve. The reduced activity in third-party chatbot engagements reflects how these built-in tools reshape user expectations and alter the pathways developers use to reach WhatsApp's large global audience.

Developers following the shift expect additional changes as Meta adjusts business APIs, messaging surfaces and integration models. The long-term trajectory will depend on how Meta balances its own assistant's capabilities with opportunities for third-party developers building tools for WhatsApp's business ecosystem, and how those tools coexist with broader messaging-platform AI strategies.

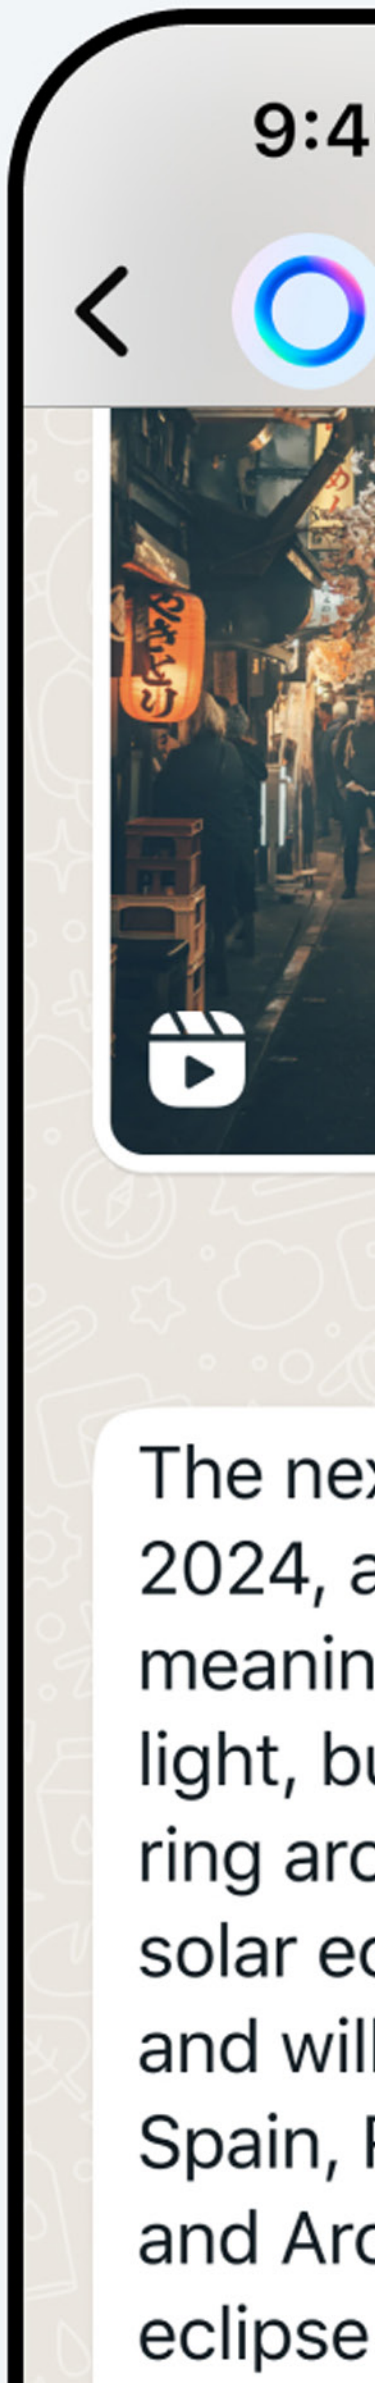

1

Meta AI

with Llama 3

Next eclipse 9:41

xt solar eclipse is on October 2,
and it will be an annular eclipse,
g the moon will block the sun's
ut the sun will still be visible as a
ound the moon [1] [2]. The next total
clipse will be on August 12, 2026,
l be visible in Greenland, Iceland,
Russia, and parts of the Atlantic
ctic Oceans [1] [2]. The next total solar
visible from North America will be

APPLE POISED TO BECOME THE WORLD'S TOP SMARTPHONE MAKER AS MARKET MOMENTUM SHIFTS

Apple is on track to surpass Samsung as the world's largest smartphone maker by annual unit shipments, according to new industry projections that indicate a notable shift in global market dynamics. The development follows a series of quarterly gains by Apple and continuing pressure on Samsung from intensifying mid-range competition in key regions. Analysts tracking shipment flows across 2024 and 2025 indicate that if present trends hold through the current quarter, Apple could close the year with its highest global market

share on record, marking a turning point in a competition that has defined the smartphone landscape for more than a decade.

The projection reflects structural changes in the premium device segment, shifts in emerging-market consumer patterns, and supply-chain repositioning across Asia. Although Samsung continues to operate the industry's broadest portfolio of smartphones, the data signals tightening margins across its mid-tier and entry-tier lines. Apple's more concentrated product range and ecosystem positioning appear to be translating into steadier global demand. Several analysts cited by regional media point to Apple's ability to maintain volume despite price tiers that historically limited its share in markets outside North America and Western Europe.

If confirmed by full-year shipment tallies, the shift would highlight changing consumer behavior that increasingly rewards premium ecosystems with consistent upgrade paths, long software support windows and brand continuity. It also underscores how the geographic balance of smartphone demand is moving toward regions where Apple has recently raised production and expanded distribution.

APPLE'S PATH TO THE TOP THROUGH PRODUCT STRATEGY, GEOGRAPHIC GROWTH AND SUPPLY-CHAIN REALIGNMENT

Apple's anticipated rise in global ranking is grounded in patterns that have unfolded across multiple quarters. One central factor is the company's approach to managing both premium flagship cycles and extended


Listening...
9:41

device life spans. In mature markets such as the United States, Canada, Japan and Western Europe, upgrade cycles have lengthened, but replacement decisions continue to favor Apple because of its integrated services and hardware ecosystem. This leads to consistent demand even when global volumes fluctuate.

Another driver is Apple's expansion in emerging markets. India, in particular, has shown rapid growth in iPhone adoption as the company increased retail channels, upgraded in-country assembly capacity and introduced financing options tailored for price-sensitive markets. The rise of trade-in programs, combined with increased production of entry-tier iPhone models in India, has brought Apple into demographic segments that were previously difficult to access due to pricing. Southeast Asia and parts of Latin America have also shown measurable growth, contributing to Apple's broader geographic base.

Jointly, these factors have strengthened Apple's global volume more than earlier projections anticipated. Analysts note that even where local economic pressures persist, Apple's brand and resale stability enable it to maintain shipment momentum in regions where many Android competitors face steep volatility.

On the supply-chain side, Apple's multi-year effort to diversify assembly across India and Vietnam has reshaped its production resilience. Initial transitions began as risk-management exercises, but by 2024 and 2025 the manufacturing expansion had reached meaningful scale. These shifts have reduced reliance on any single region and given Apple

flexibility to manage logistics, component sourcing and labor constraints. Suppliers such as Foxconn, Pegatron and Tata Electronics have increased output capacity for iPhones and accessories, allowing Apple to better match global demand cycles.

The company's ability to condense production timelines for flagship launches has also supported its volume growth. Faster ramp-ups enable early-quarter shipments to accumulate more rapidly than in past cycles, contributing to stronger year-end totals. Analysts point to these operational changes as one reason Apple has maintained steadier shipments even during supply-chain disruptions that affected competitors across components, display technologies and semiconductor availability.

SAMSUNG'S POSITION AS COMPETITION RISES IN MID-RANGE AND ENTRY SEGMENTS

Samsung remains one of the world's most diversified smartphone manufacturers, offering devices across a wide range of price tiers and form-factor niches. However, the pressures it faces differ from Apple's. While Samsung's flagship Galaxy S and Galaxy Z lines continue to perform globally, their growth has been moderated by increased saturation in core markets. At the same time, Chinese brands have accelerated their challenge in the mid-range and entry tiers, which historically served as Samsung's volume drivers.

Companies such as Xiaomi, Vivo, Oppo and Honor have expanded aggressively in India, Southeast Asia, Africa and parts of Europe with

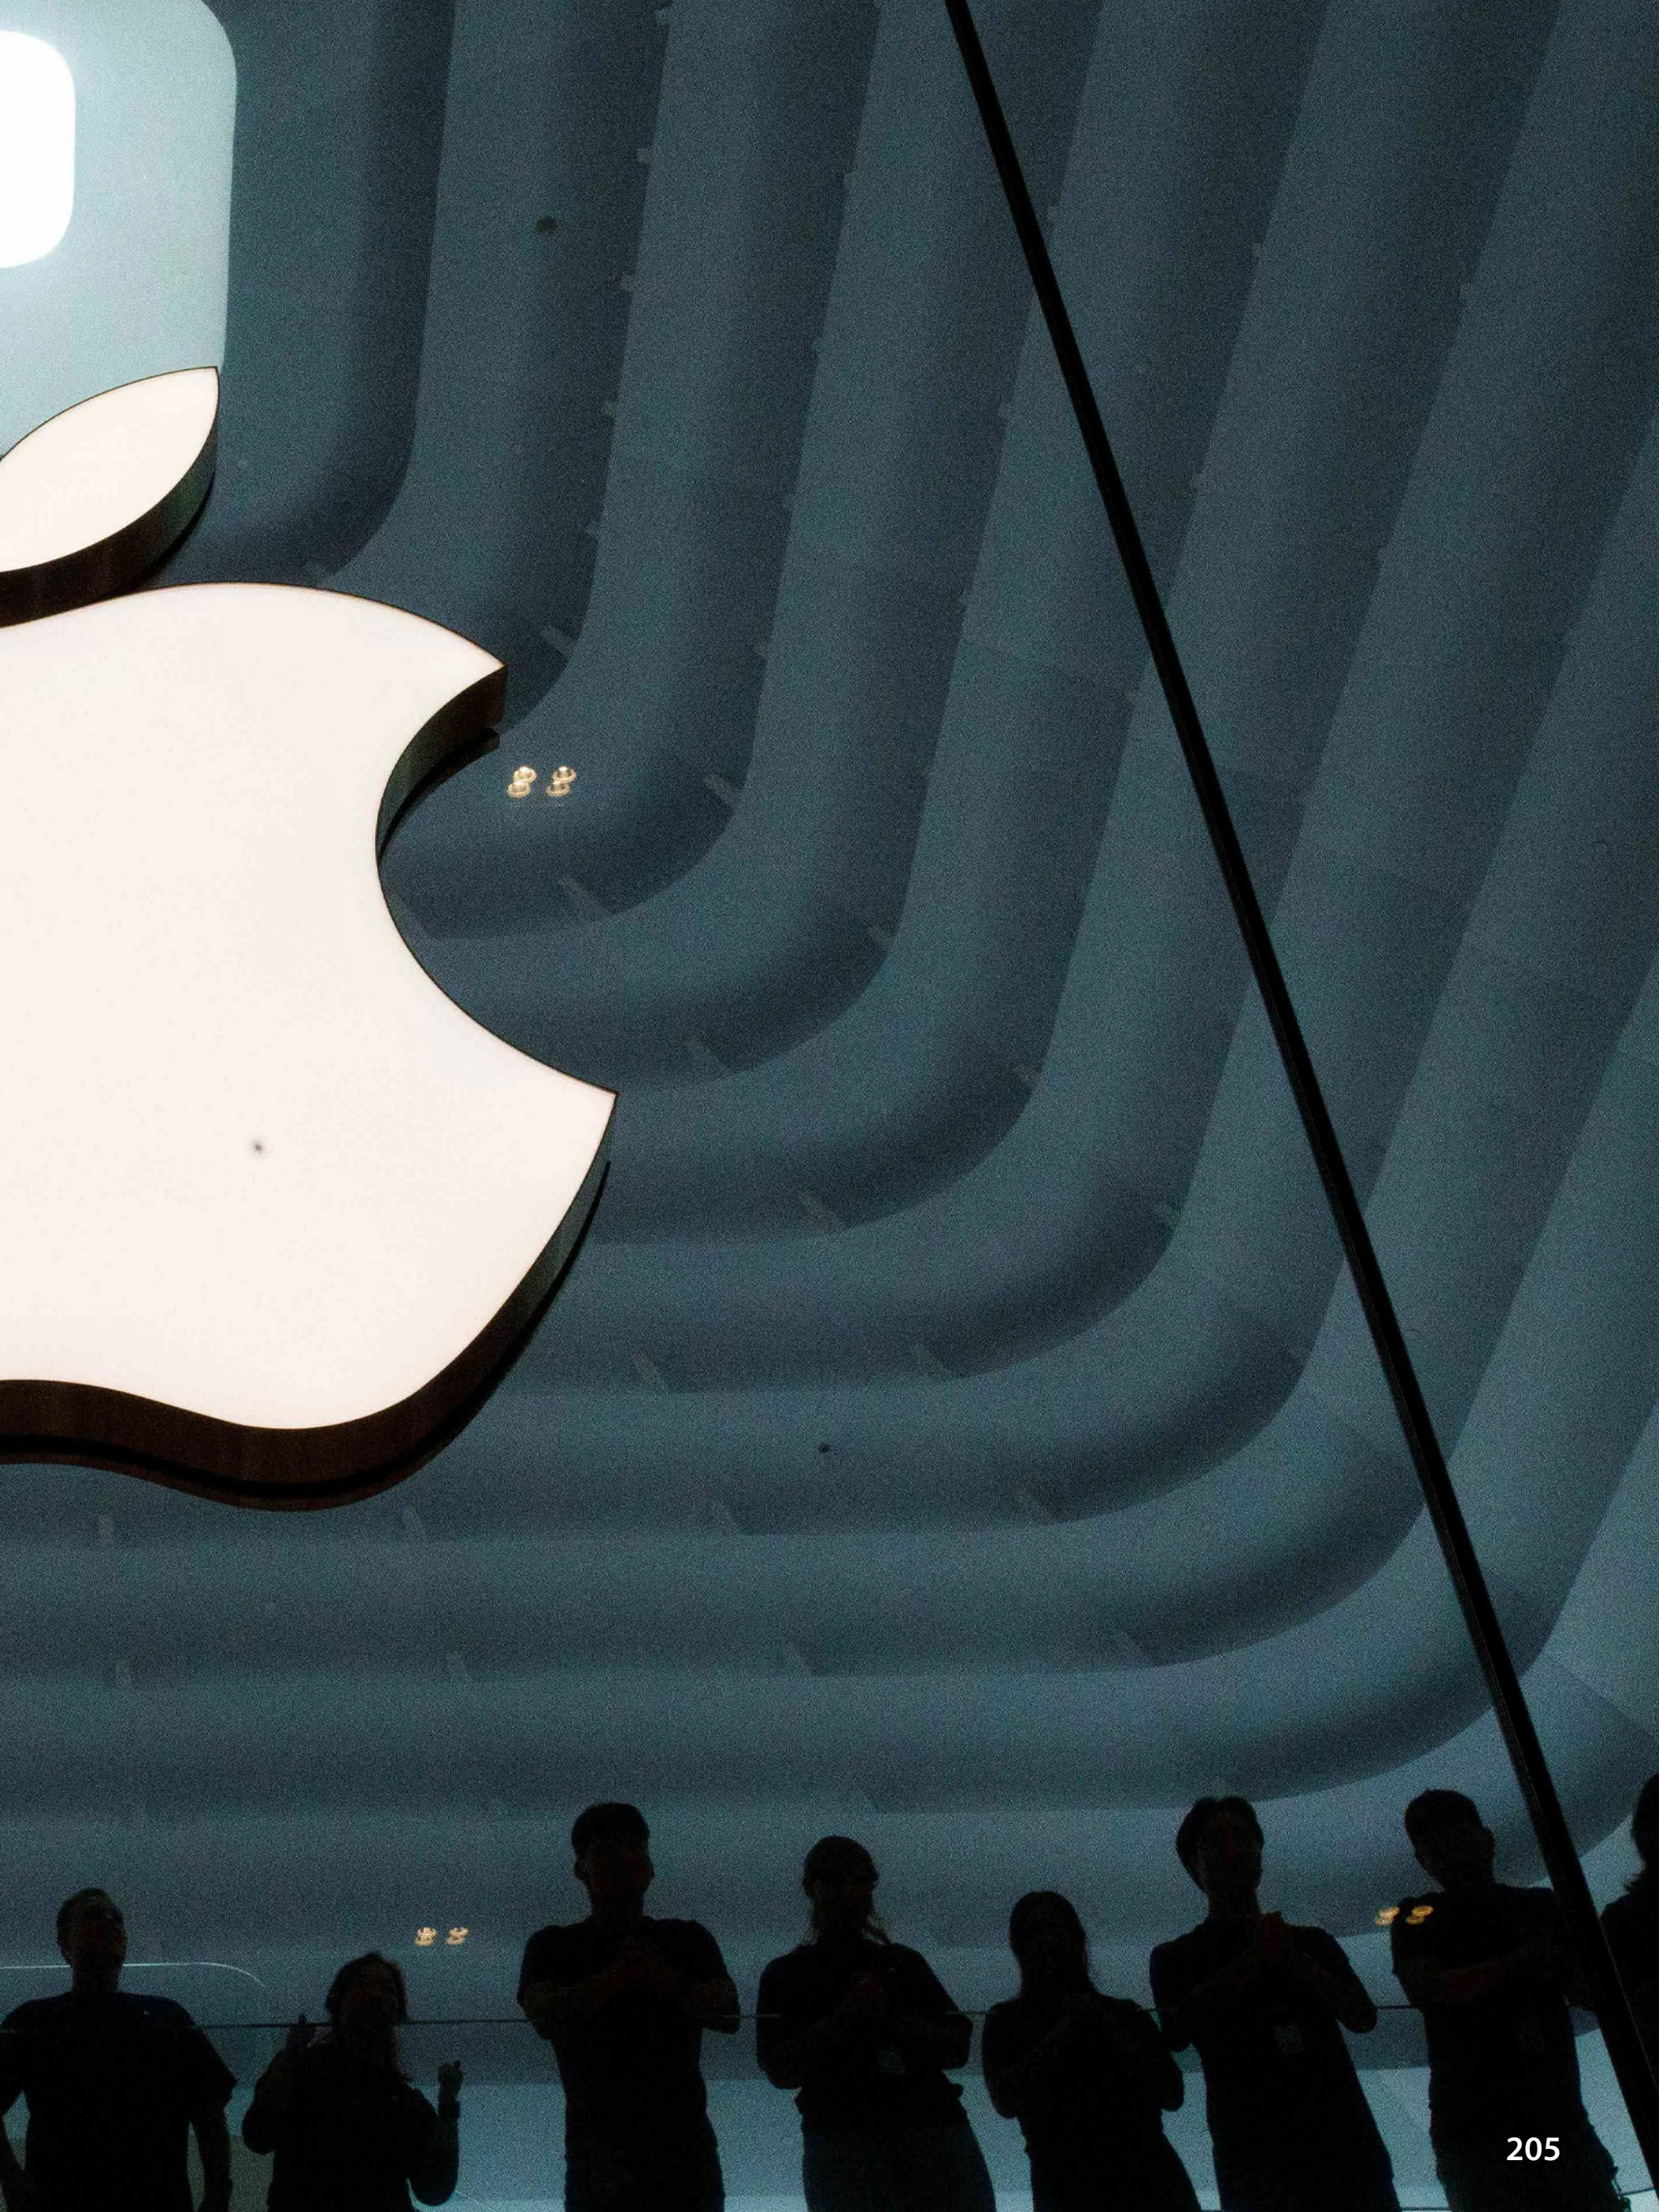

devices that offer competitive specifications at lower prices. These brands have localized features, distribution and marketing strategies, creating pressure on Samsung's A-series and M-series lines. As consumers in key growth regions weigh cost against premium functionality, Samsung's broader portfolio has encountered price-sensitive competition that narrows margins and affects global shipment totals.

Analysts also point to Samsung's complex product lineup as both strength and friction. Broad diversification allows reach across demographics, but it also introduces higher manufacturing, marketing and inventory burdens. Frequent model refreshes across multiple tiers create additional cost challenges, particularly in volatile currency environments. As Samsung spreads resources across dozens of product variants, its ability to absorb competitive pricing moves from Chinese rivals becomes limited, affecting shipment competition with Apple at the top rank.

Samsung's foldable-device category, once positioned as a breakthrough differentiator, has maintained steady sales but has not grown quickly enough to offset softness elsewhere. While the Galaxy Z Fold and Z Flip models attract attention, their pricing and durability considerations have slowed wider adoption, leaving much of Samsung's volume tied to conventional slab smartphones where price competition is most intense.

In contrast, Apple's narrower product family allows more concentrated engineering, software and supply-chain focus, facilitating

cost efficiencies at scale. As Apple's volumes rise in markets once dominated by Samsung, the competitive landscape becomes more balanced, reducing Samsung's historic lead in global unit share.

INDUSTRY WIDE EFFECTS AS GLOBAL SMARTPHONE DEMAND EVOLVES

The potential leadership shift illustrates several broader industry patterns. First, premiumization continues to influence global purchasing behavior. Even in regions where average income remains below global benchmarks, the value placed on long-term support, integrated ecosystems and device longevity is increasingly affecting upgrade decisions. Apple has capitalized on this through extended iOS update timelines and close hardware-software integration.

Second, global supply chains are undergoing reconfiguration. The location and scale of manufacturing now play a larger role in determining shipment performance, particularly during periods of geopolitical tension or logistics disruptions. Apple's expansion into India and Vietnam reduces risk concentration and aligns with government incentives in those regions, contributing indirectly to shipment stability.

Third, marketing influence is increasingly tied to ecosystem services rather than hardware alone. Apple's services portfolio, including iCloud, Apple Music, Apple TV+ and AppleCare, strengthens retention among existing users and helps convert first-time buyers. Samsung maintains its own ecosystem strategy, but the

sheer breadth of its product range can dilute the perception of unified software continuity, affecting long-term retention differently across markets.

For carriers, component suppliers and software developers, the implications of an Apple-led shipment market are material. Carrier promotion cycles often shift to align with the strongest-selling flagship category, affecting subsidy allocation and upgrade-device visibility. Component suppliers may adjust output priorities to align with Apple's demand patterns, which can in turn influence availability and pricing for smaller manufacturers. Developers may continue to prioritise iOS early in feature rollouts if the platform's installed base and upgrade rate remain proportionally strong compared with Android's fragmented update cycles.

Finally, the smartphone market overall remains stable in annual units, but the competitive composition within it is changing. With Chinese manufacturers growing in volume and Apple strengthening the premium segment, global competition increasingly hinges on ecosystem control, brand identity and cost efficiencies across supply chains.

Should the year-end data confirm Apple's new position at the top, the shift would reflect a longer-term rebalancing of the industry rather than a single-cycle anomaly. It would also set the stage for new dynamics in flagship competition, services integration and geographic expansion across 2026 and beyond.


TIX
ROBOTAXI

TESLA REVISES ROBOTAXI FLEET TARGETS IN AUSTIN AS DEVELOPMENT PRIORITIES SHIFT

Tesla has adjusted its robotaxi deployment expectations in Austin, reducing internal fleet projections as the company reassesses timelines, operational planning and early-stage infrastructure needs for its autonomous-vehicle program. According to new information reported this week, the revised outlook reflects updated modeling inside Tesla that narrows the number of vehicles expected to operate in the city during the initial rollout phase. The change illustrates how Tesla continues to recalibrate its approach as it works toward a commercial network built around fully autonomous vehicles.


TESL

People familiar with the updated planning described the adjustment as part of a broader internal review of deployment timelines, regulatory pathways and resource allocation across multiple markets. Tesla's early robotaxi plans included more expansive fleet estimates for Austin, where the company has strong manufacturing and testing ties, but the newest internal projections reflect a more measured starting point. The shift appears connected to updated assumptions around operational density, permitting coordination and early routing frameworks rather than a retreat from the program itself.

Tesla has maintained its public ambition to introduce a dedicated robotaxi vehicle and scale autonomous operations in select US markets. Austin remains a core location for testing and early commercial activity, supported by proximity to the company's Gigafactory Texas complex and the city's participation in multiple mobility initiatives. The revised fleet estimates indicate that Tesla is evaluating how a concentrated deployment could shape the earliest stages of service before expanding the network more broadly.

UPDATED FLEET MODELS REFLECT OPERATIONAL AND PLANNING CONSIDERATIONS

The reduced Austin estimates align with Tesla's internal recalibration of expected vehicle volumes for early robotaxi deployments. While earlier projections included larger fleets operating during initial phases, the company has since refined assumptions about route planning, vehicle turnaround times, service-area boundaries and staged expansion. Adjustments were also tied to

how autonomous systems perform under real-world conditions, with the company conducting extensive internal scenario modeling to determine viable operating density.

The decision to narrow the initial fleet size coincides with ongoing development of Tesla's next-generation autonomous vehicle, designed specifically for robotaxi use. The company continues to work through prototype iterations, sensor-suite evaluations and software-integration cycles as it prepares the platform for scaled production. These engineering processes influence the timing and quantity of vehicles available for deployment.

Austin's role as an early robotaxi market remains unchanged, but Tesla's updated targeting reflects the company's approach to infrastructure readiness. The early fleet must integrate with service depots, charging hubs and maintenance cycles that still require staging. People familiar with the planning noted that Tesla's internal models now emphasize the value of reliability, route predictability and service continuity during the initial phase, which favors a concentrated fleet rather than broad saturation.

Tesla's internal communications to staff also framed the revised fleet expectations as part of a data-driven refinement process. As autonomous-system testing continues in Austin and other regions, operational insights shape the models used to determine optimal fleet size. These models account for traffic patterns, regulatory checkpoints, population density, vehicle availability and local mapping data.


04·26
ROBOTAXI

ALIGNMENT WITH REGULATORY, MANUFACTURING AND SOFTWARE TIMELINES

The updated fleet estimates in Austin intersect with several parallel factors shaping Tesla's robotaxi roadmap. Regulatory considerations remain central, with municipal and state frameworks determining how autonomous services can operate. Austin continues to engage with companies testing autonomous vehicles, but approvals, reporting standards and safety protocols influence how quickly high-density deployments can begin.

Tesla's manufacturing pipeline for next-generation vehicles also plays a role. The robotaxi platform is expected to rely on a variant of Tesla's latest structural components and battery systems, requiring synchronization between Gigafactory Texas production timelines and software-readiness cycles. As development advances, the alignment between manufacturing output and software maturity affects how many vehicles Tesla can commit to early markets.

The company's ongoing work on its autonomous-driving software informs these decisions as well. Tesla continues to test incremental updates to its Full Self-Driving system, incorporating new perception, planning and behavior models to support autonomous capability. Integration between the robotaxi hardware platform and the evolving software stack will influence deployment timing and initial fleet size.

Tesla's internal planning documents indicate that early robotaxi operations are expected to rely heavily on tightly defined service areas. This structure allows the company to refine routing


Robotaxi

models, assess real-world system performance and calibrate vehicle behavior within controlled geographic boundaries. Smaller fleet sizes support that scaled-testing approach and reflect how early deployments often prioritize operational clarity over broad footprint expansion.

IMPLICATIONS FOR TESLA'S BROADER AUTONOMOUS STRATEGY

The adjustment to Tesla's Austin robotaxi fleet projections does not alter the company's long-term aim of deploying autonomous-vehicle networks, but it signals how Tesla is sequencing its efforts as development advances. The company continues to refine launch markets, infrastructure plans and the software-hardware alignment that supports autonomous operations. These refinements shape how initial fleets are sized, how operating areas are selected and how expansion stages are defined.

Austin's significance within Tesla's autonomous plans remains tied to proximity to the company's manufacturing operations, the availability of test environments and the city's willingness to work with emerging mobility technologies. As testing and infrastructure preparation progresses, the company expects to introduce additional stages of deployment that expand vehicle count and service area over time. Early operational data from the refined fleet will inform planning for other markets under evaluation.

MUSK LINKS REAL-WORLD DATA FROM X AND TESLA TO FUTURE AGI DEVELOPMENT


TESLA
FDC
FIRE SPRINKLER RISER ROOM
FIRE ALARM CONTROL PANEL
SERVICE ENTRANCE
SERVICE ENTRANCE

Elon Musk has outlined a connection between his companies' real-world data operations and long-term artificial-general-intelligence research, describing how information gathered across Tesla's autonomous-driving systems and the social interactions hosted on X could contribute to the development of advanced AI models. The remarks, referenced in new reporting, reflect Musk's view that large-scale datasets representing human behavior, situational complexity and real-time decision patterns will shape how next-generation systems are trained.

Musk's comments situate Tesla's vehicle-sensor network and X's conversational activity as complementary elements in a broader AI strategy. Tesla collects extensive driving data from vehicles operating across numerous road conditions, while X hosts continuous public discourse across millions of posts. Musk described these streams as separate but mutually relevant for long-horizon model development. The framing highlights how data variety—from physical-world perception to textual interactions—remains central to building increasingly adaptive systems.

Industry observers tracking Musk's companies note that Tesla and X operate under distinct constraints, engineering requirements and commercial models. The former builds automotive hardware and software, while the latter runs a social-media platform with its own technical and operational realities. Musk's comments do not indicate any merging of datasets but suggest an intellectual link between the complexity each platform captures and the types of training pathways Musk envisions for future AI systems.

TESLA'S SENSOR DATA REMAINS CENTRAL TO ITS AUTONOMOUS-VEHICLE SYSTEMS

Tesla's fleet generates continuous real-world data through camera arrays, ultrasonic sensors and onboard computing systems installed in millions of vehicles. These data streams support the company's autonomous-driving development, allowing engineers to analyze traffic patterns, environmental conditions, edge cases and driver reactions. The resulting datasets help refine perception systems, route-planning modules and decision-making behavior as the company evaluates new iterations of its software.

Musk has long described Tesla's fleet as a distinctive asset due to the scale at which the vehicles operate across varied environments. The combination of recorded driving sequences and simulated scenarios contributes to training cycles for Tesla's next-generation autonomy models. The scale of this dataset positions Tesla differently from developers relying solely on controlled test environments or limited pilot fleets.

The company continues to work on the evolution of its autonomous-driving stack, including updates to planning algorithms and neural networks used to interpret road environments. Each software revision undergoes internal validation processes, with data from the fleet shaping how the system responds to unpredictable or low-frequency events. As Musk discusses AGI in broader terms, Tesla's real-world driving data remains firmly connected to its automotive roadmap rather than any cross-platform integration.


ESLA
F K
PUTIN
California
8SAP741

X PROVIDES A SEPARATE STREAM OF HUMAN-GENERATED INTERACTIONS

In parallel, X operates as a high-volume platform for written and visual interactions. The site hosts conversations, debates, news commentary and social observations that generate large quantities of public posts. These interactions form a dataset reflective of human communication, opinion formation and social context. Musk referenced this activity when discussing potential inputs for long-range AI research.

X's data profile differs substantially from Tesla's. While Tesla's datasets capture physical-world complexity, X captures linguistic behavior, user responses and broad social engagement. Any connection Musk draws between the two appears conceptual rather than operational; each company maintains its own data policies and legal frameworks, and there is no indication of any blending of datasets.

Developers and analysts familiar with AI-model training note that platforms with varied content streams often contribute to different types of model capabilities. Text-based environments such as X can support language understanding and conversational reasoning work, whereas sensor-rich environments like Tesla's fleet are used for perception and navigation. Musk's remarks connect these domains only in the context of examining how diverse forms of complexity might shape wider AGI research efforts.

MUSK'S LONG-TERM FRAMING OF AGI WORK ACROSS SEPARATE COMPANIES

Musk's commentary reflects a pattern in which he discusses AI development not as a single-

project undertaking but as an interconnected field informed by multiple real-world systems. His companies approach AI with distinct engineering aims: Tesla focuses on automotive autonomy, while X maintains conversational features and content-moderation tools. The conceptual relationship Musk draws between them underscores his broader view that AGI will require exposure to both structured physical environments and diverse human-generated content.

The comments arrive during a period when major technology firms continue refining how they collect, manage and apply data for model training. Discussions surrounding data privacy, regulatory frameworks and platform governance shape how datasets can be used. Musk's framing of Tesla and X in relation to AGI reflects his long-standing interest in advanced AI systems but does not point toward any operational crossover.

Future developments in Musk's AI efforts will depend on how each company progresses through its respective data, software and product roadmaps. Tesla continues to expand testing of its autonomy software while X adjusts its platform features and engagement strategies. As these systems evolve, Musk's public comments will continue to be examined for how they contextualize the relationship between real-world data and long-horizon AI research.

NVIDIA–GOOGLE AI CHIP COMPETITION INTENSIFIES AS META EXPLORES NEW OPTIONS

Industry dynamics in artificial-intelligence hardware are shifting again as new reports suggest Meta is evaluating alternatives to NVIDIA's data-center accelerators, prompting renewed attention on Google's position in the market. The development adds another dimension to the long-running competition between NVIDIA and Google in high-performance AI computing, an area central to training and running large-scale models. While NVIDIA remains the dominant supplier for advanced compute clusters, Meta's exploratory


NVIDIA

discussions highlight how major technology firms are reassessing supplier diversity and long-term infrastructure strategies.

A new report describes ongoing conversations in which Meta leaders have been weighing new hardware pathways that could complement or eventually offset reliance on NVIDIA's GPU roadmap. Google's cloud division, with its internally designed tensor processing units, has emerged as one avenue under consideration. These discussions, according to individuals familiar with the matter, reflect Meta's continued effort to evaluate multiple hardware options for its expanding AI workloads. The talks have not resulted in any finalized deals, but they underscore the evolving competitive landscape as demand for compute accelerators increases across the industry.

NVIDIA's standing in the market is still anchored by its top-to-bottom ecosystem of GPUs, interconnects and software stacks that developers rely on. The possibility of Meta exploring external alternatives illustrates broader pressures within high-performance AI computing, where companies are balancing model growth, cost structures and supply-chain constraints. For Google, offering TPU-based cloud services to high-profile partners would extend a strategy aimed at expanding adoption of its vertically integrated compute platform.

NVIDIA'S POSITION REMAINS CENTRAL EVEN AS COMPETITION EXPANDS

NVIDIA continues to lead the AI-accelerator market by a wide margin, supported by consistent performance gains across GPU generations and the industry's reliance on its


oogle

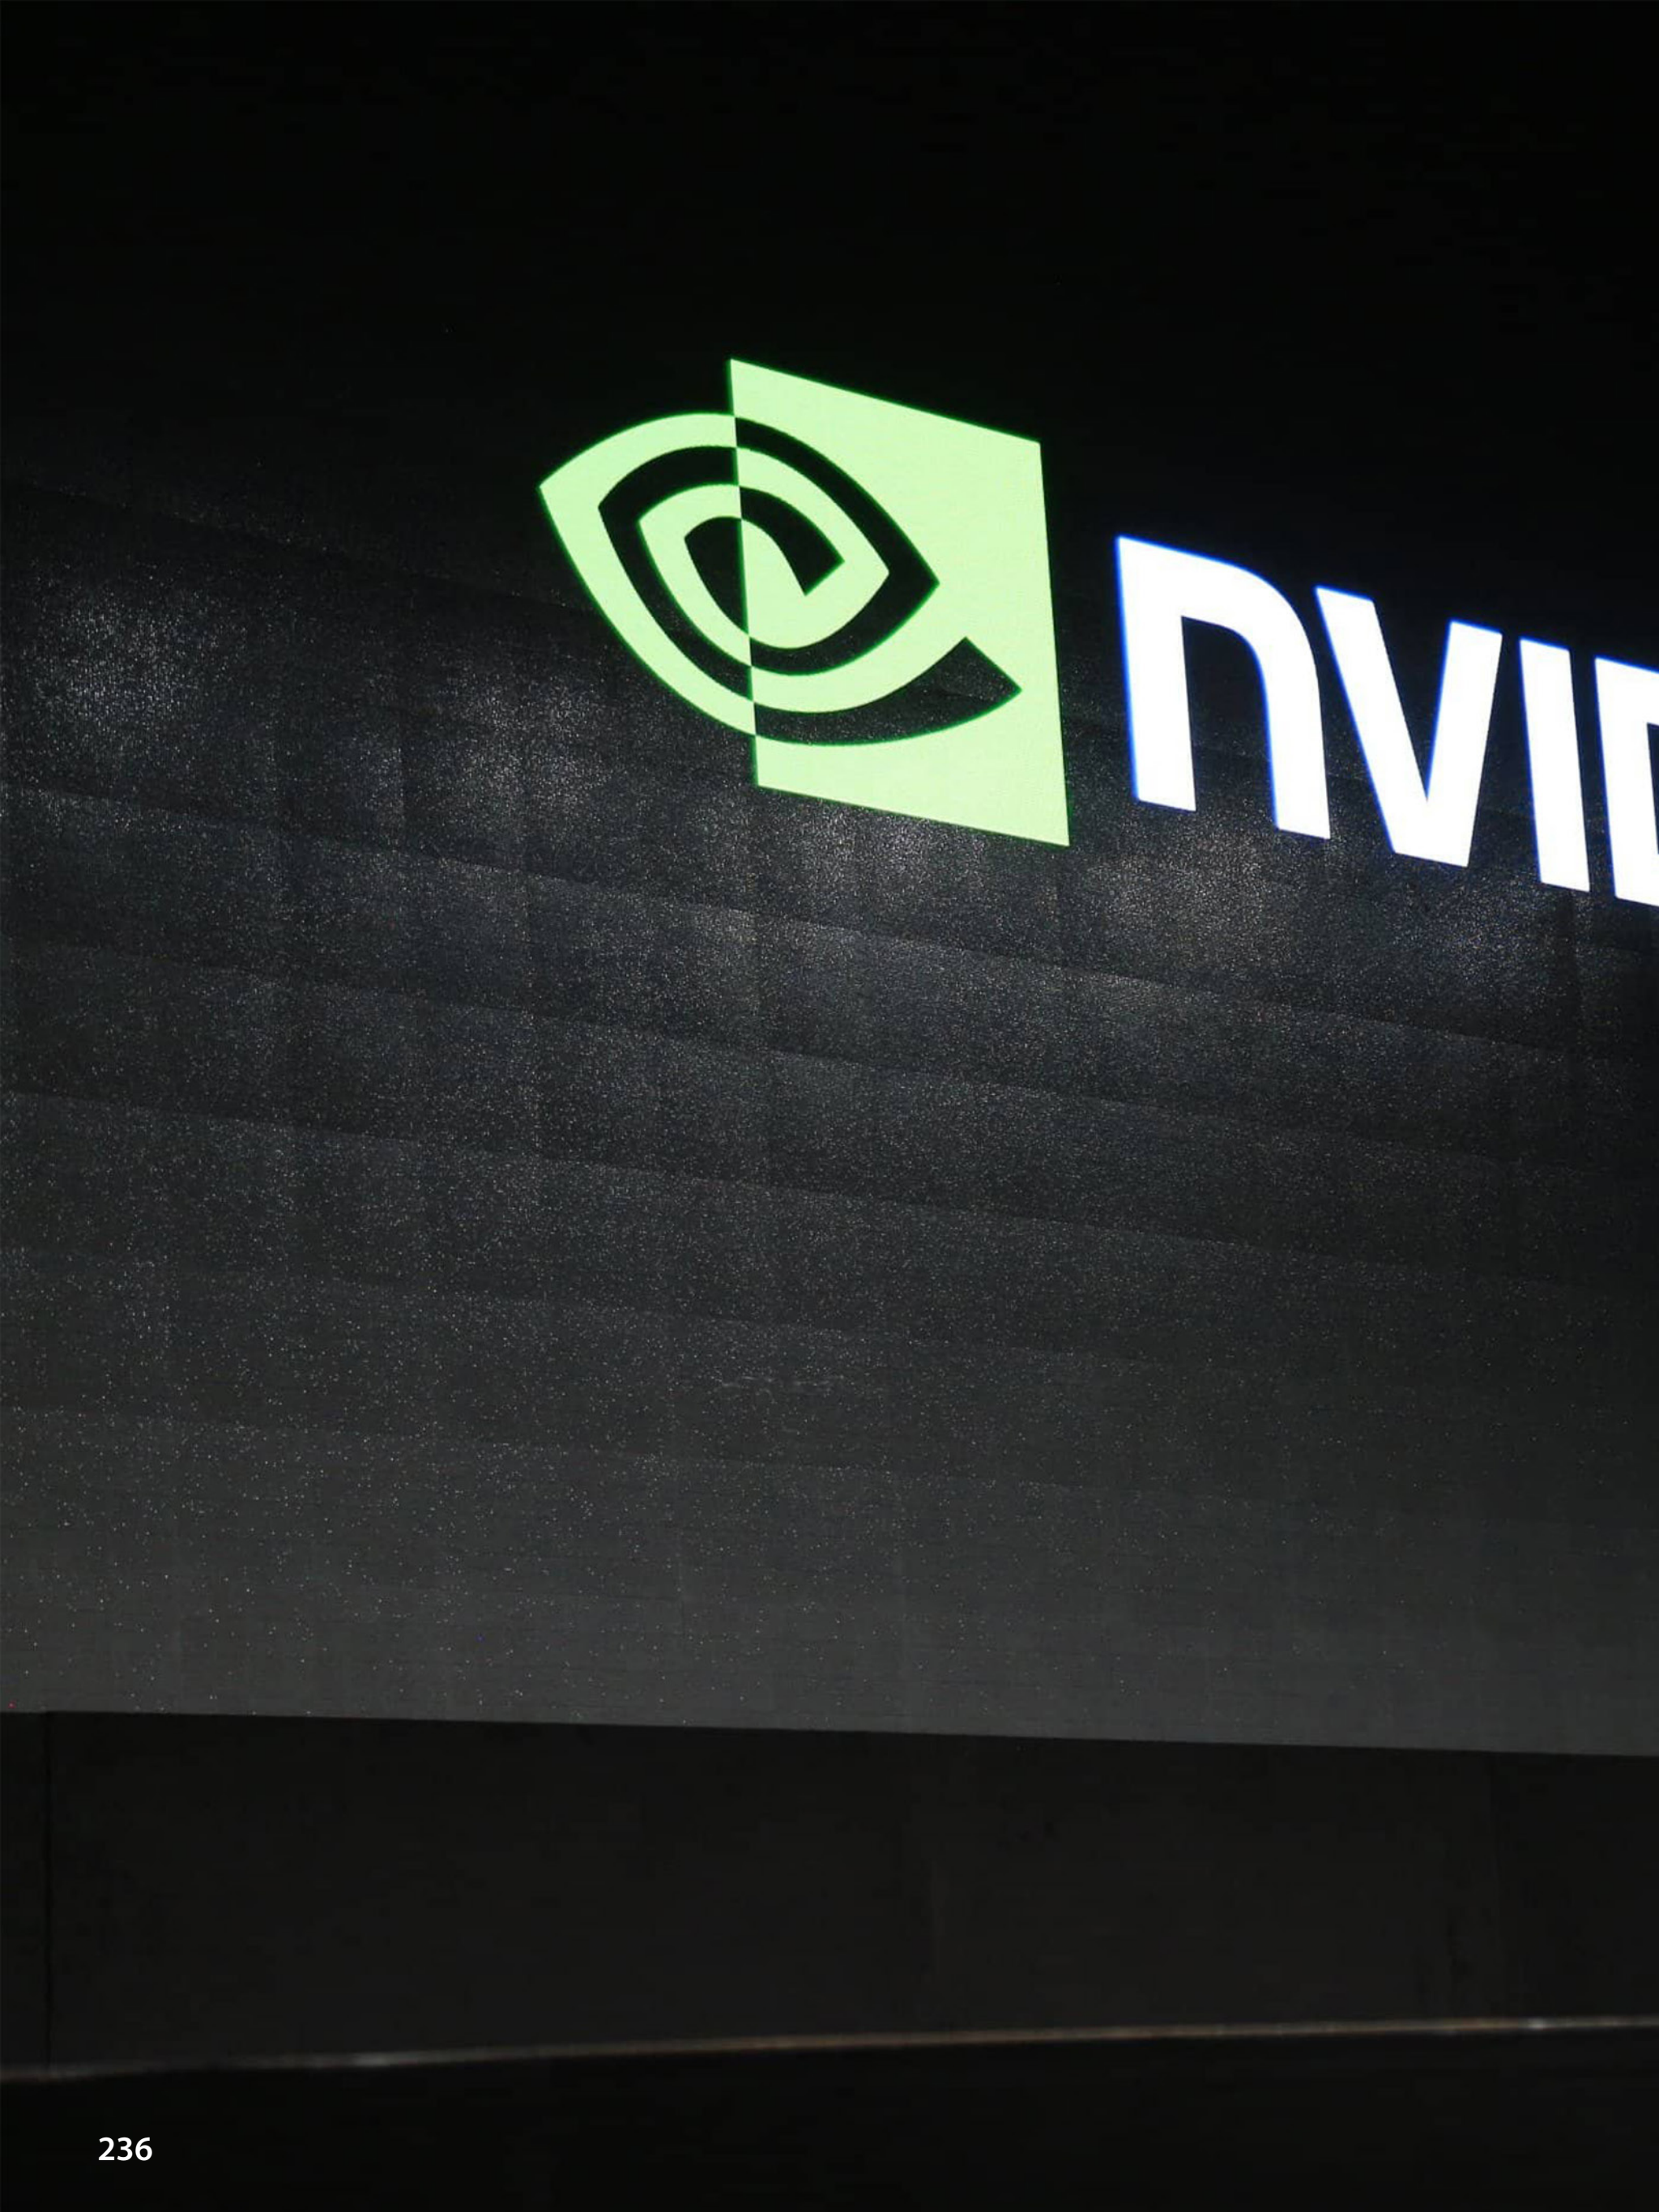
NVI

CUDA software environment. Training and deploying multimodal and large-language models at scale still depend heavily on NVIDIA's high-end systems, which remain core to clusters operated by hyperscalers, research institutions and enterprise clients. NVIDIA's most recent hardware releases have maintained this trajectory, offering increased compute density and interconnect bandwidth.

The company's lead is also reinforced by the maturity of its developer ecosystem. CUDA, libraries for distributed training, inference optimization tools and long-established integration across machine-learning frameworks make NVIDIA's platform a default choice for organizations building large-scale pipelines. This ecosystem effect continues to shape procurement decisions, even among firms evaluating alternates for specific workloads or long-term diversification.

However, surging demand for high-performance accelerators has sharpened interest in supplementary hardware paths. Companies with expansive AI roadmaps face fluctuating supply availability, cost structures shaped by multi-year hardware cycles and strategic questions about the balance between in-house development and dependence on external suppliers. These pressures have led multiple firms to explore custom silicon or cloud-based alternatives, prompting new opportunities for competitors emphasizing different compute architectures.

GOOGLE'S TPU ROADMAP AND META'S EXPLORATION OF ALTERNATIVE OPTIONS

Google has spent years refining its TPU architecture, using it internally for search, ads, large-scale training and experimental research. The company's cloud business now provides TPU access to external clients, offering a hardware-software stack tuned for high-volume machine-learning operations. For organizations exploring hardware diversity, Google's TPUs present an option structured around deep integration with Google Cloud and a dedicated compiler toolchain rather than the CUDA ecosystem.

People familiar with Meta's internal evaluations indicated that the company's leadership has periodically reviewed TPU performance characteristics, cost considerations and deployment models. These reviews occur alongside Meta's ongoing internal chip development programs, which aim to create accelerators aligned with its long-term AI and infrastructure needs. Meta has also continued using NVIDIA's latest systems for training and inference across several model families, including generative and recommendation workloads.

The discussions involving Google reflect an environment where hyperscalers and large technology firms test multiple hardware options to assess how different compute architectures align with organizational priorities. TPU clusters offer tightly controlled environments, high levels of integration with Google's cloud infrastructure and performance optimization for specific

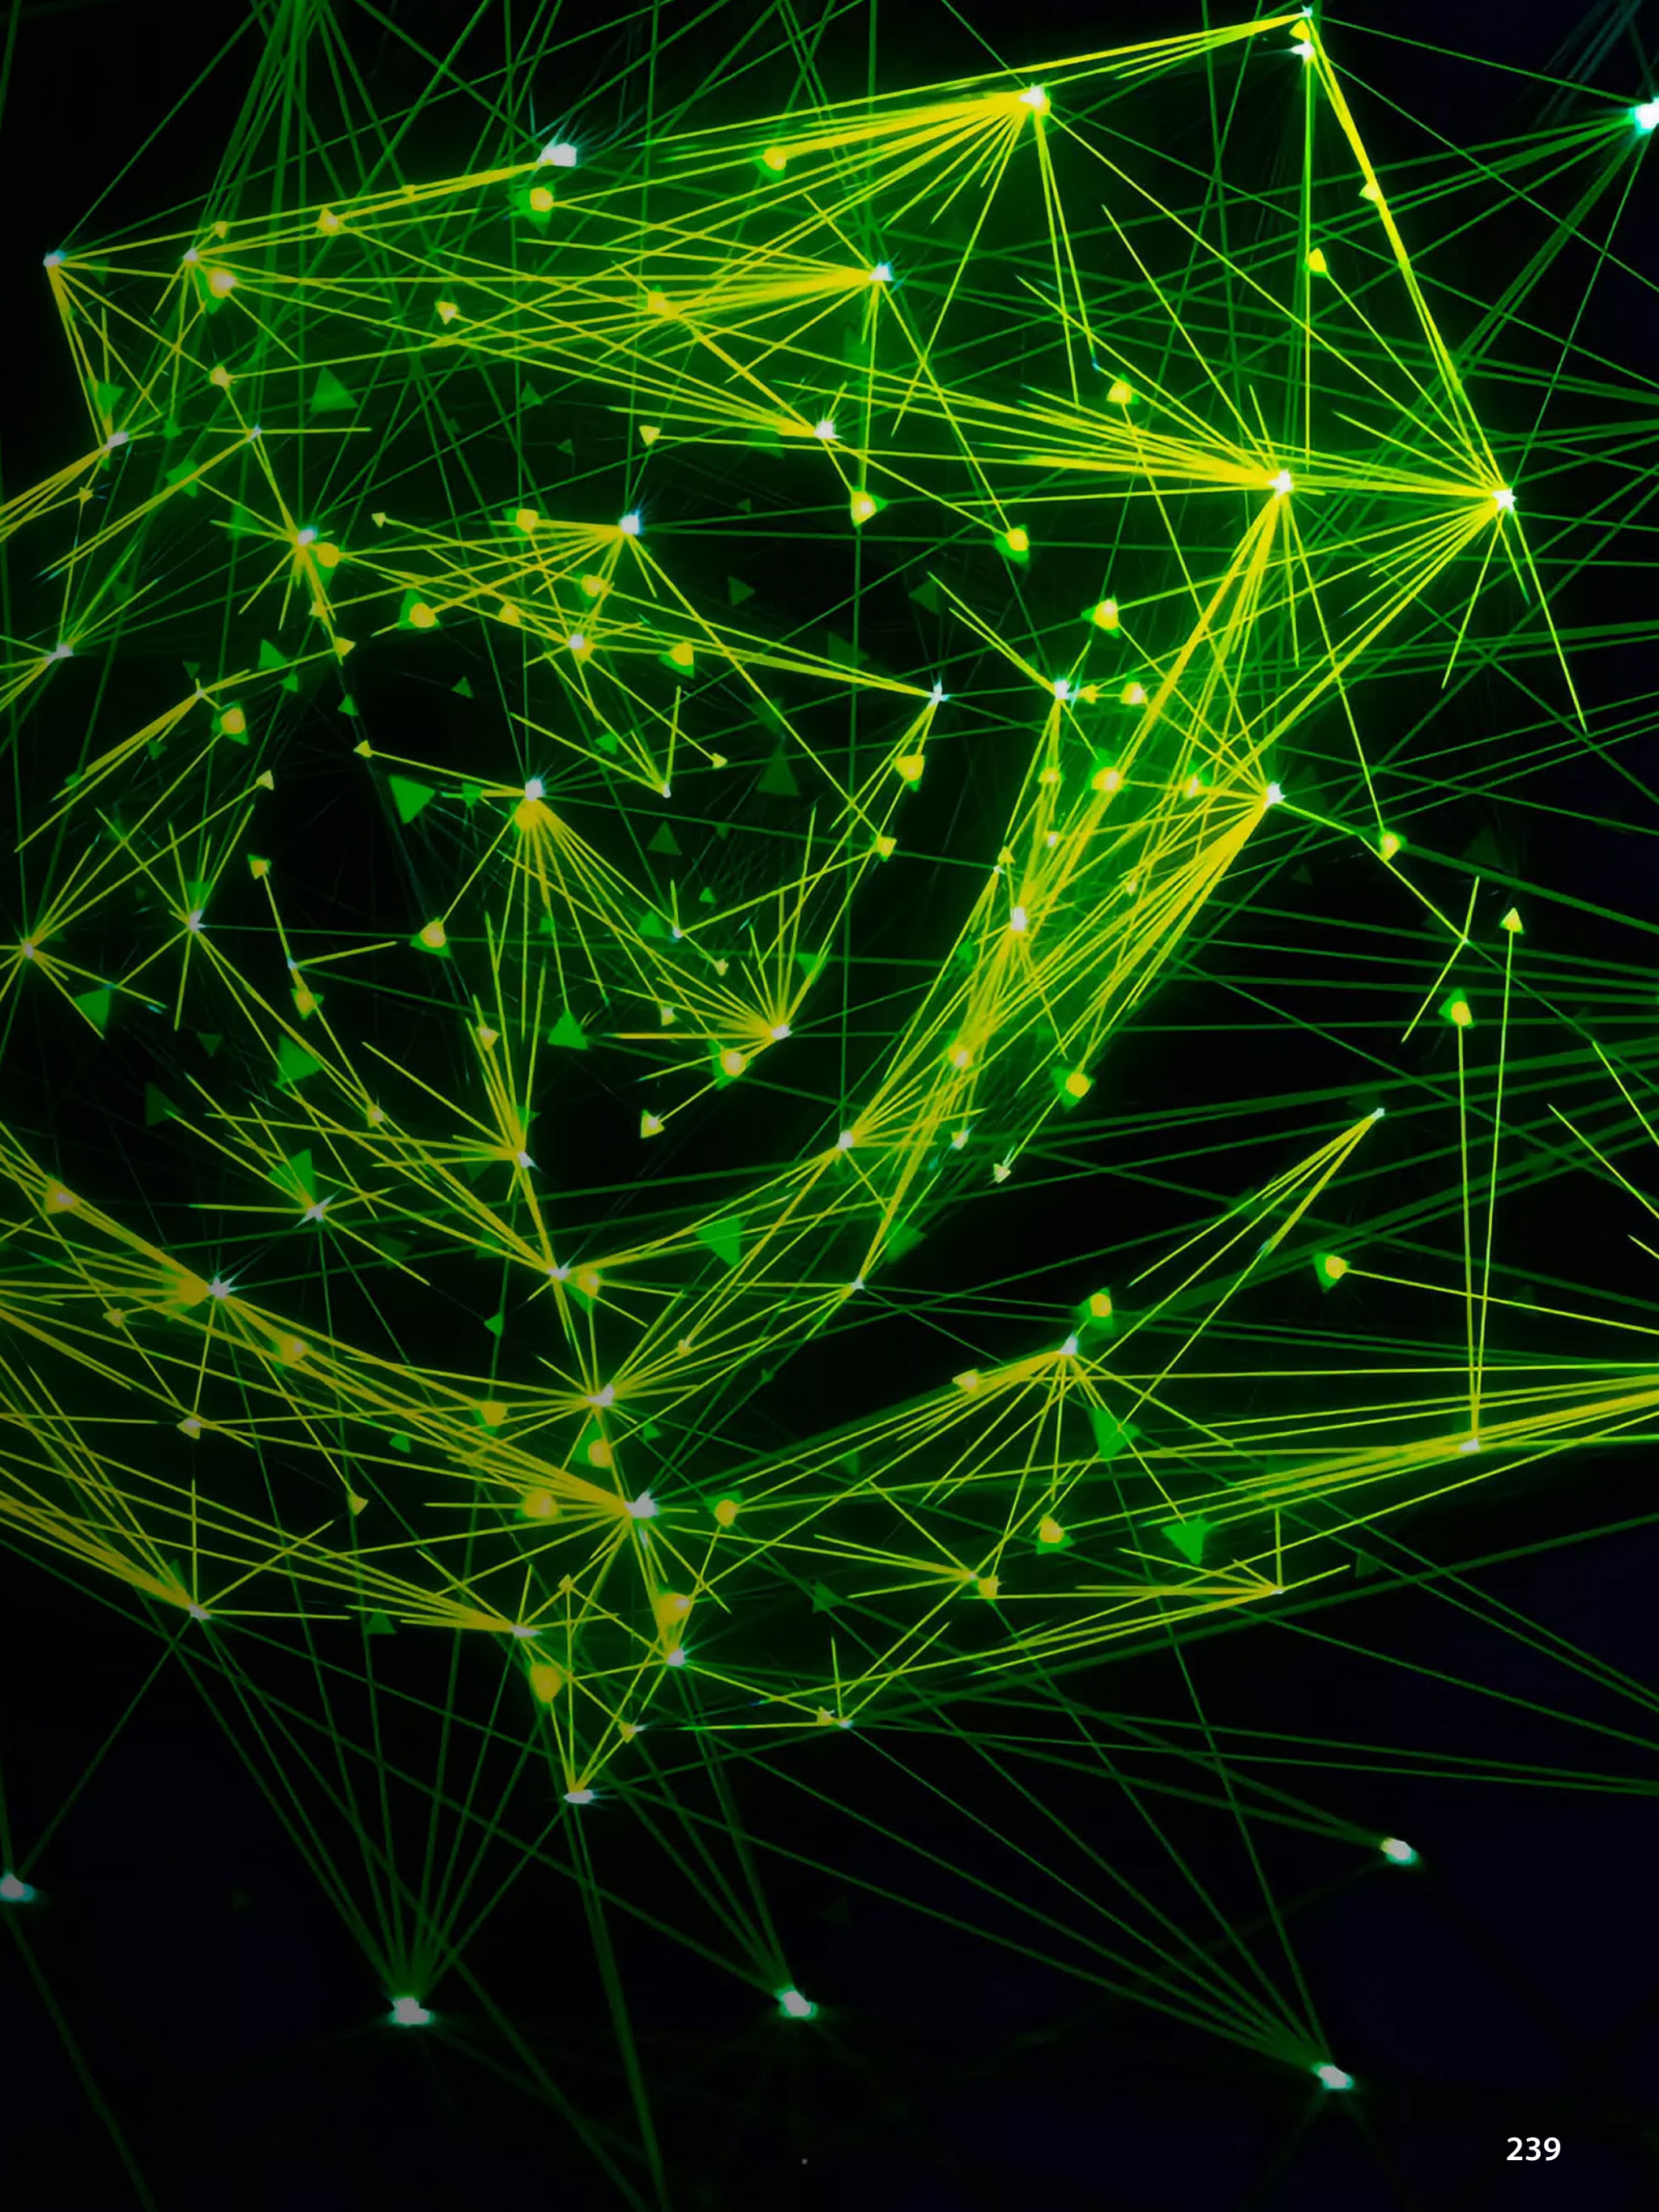

model architectures. For Meta, exploring such options allows benchmarking across hardware families while assessing whether alternate paths could support parts of its AI pipeline.

Google's efforts to expand TPU adoption outside its organization align with trends in cloud-based AI infrastructure. Cloud providers are emphasizing differentiated silicon strategies to support clients building high-volume compute workloads. Google's TPU evolution, combined with its large-scale deployment experience, positions it as one of the few entities able to compete with NVIDIA in select areas of high-performance model training.

A SHIFTING INFRASTRUCTURE LANDSCAPE AS MODEL DEMANDS INCREASE

The competitive dynamic between NVIDIA and Google reflects the broader scale of AI development. Model architectures continue to grow in parameter count, context window length and multimodal complexity, increasing compute requirements for training and inference. Organizations building products around generative models face rising demand for high-density data-center capacity as well as concerns about power availability, operational efficiency and long-term cost predictability.

As a result, many firms have expanded conversations with multiple hardware suppliers, internal silicon teams and cloud platforms. The exploration of alternatives does not necessarily indicate reductions in NVIDIA use, especially given the scale and maturity of its ecosystem. Instead, these evaluations illustrate

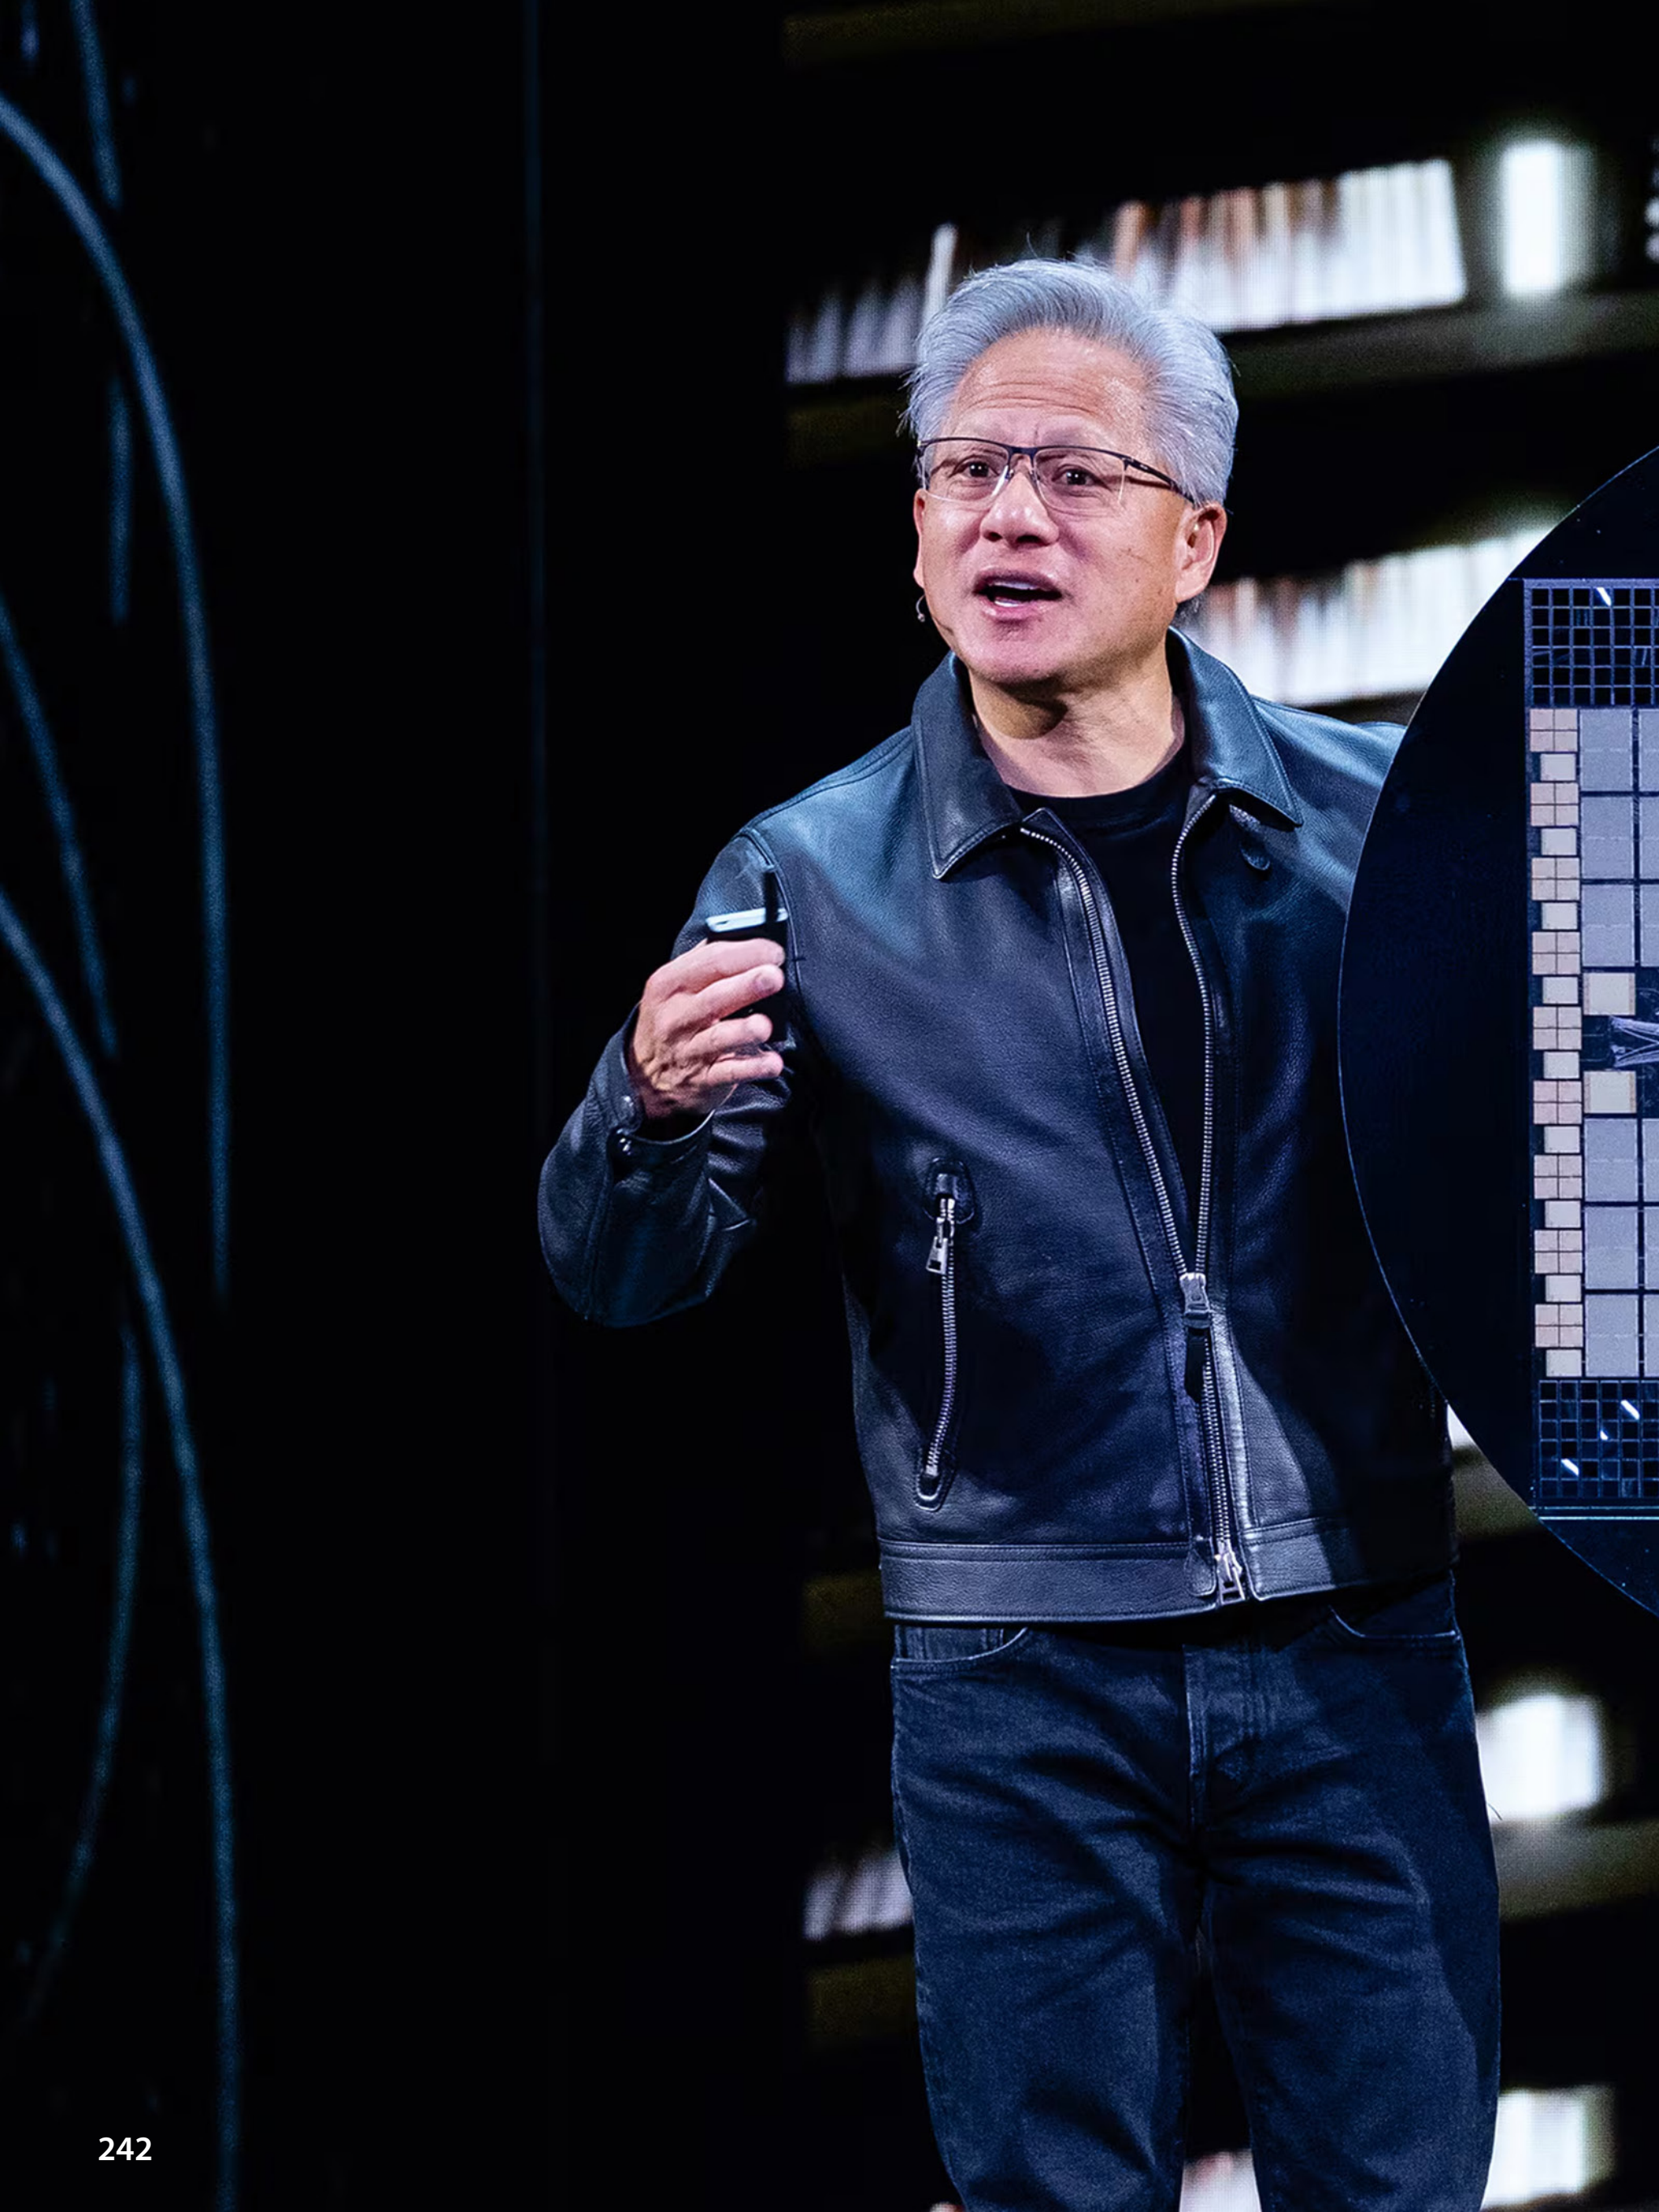

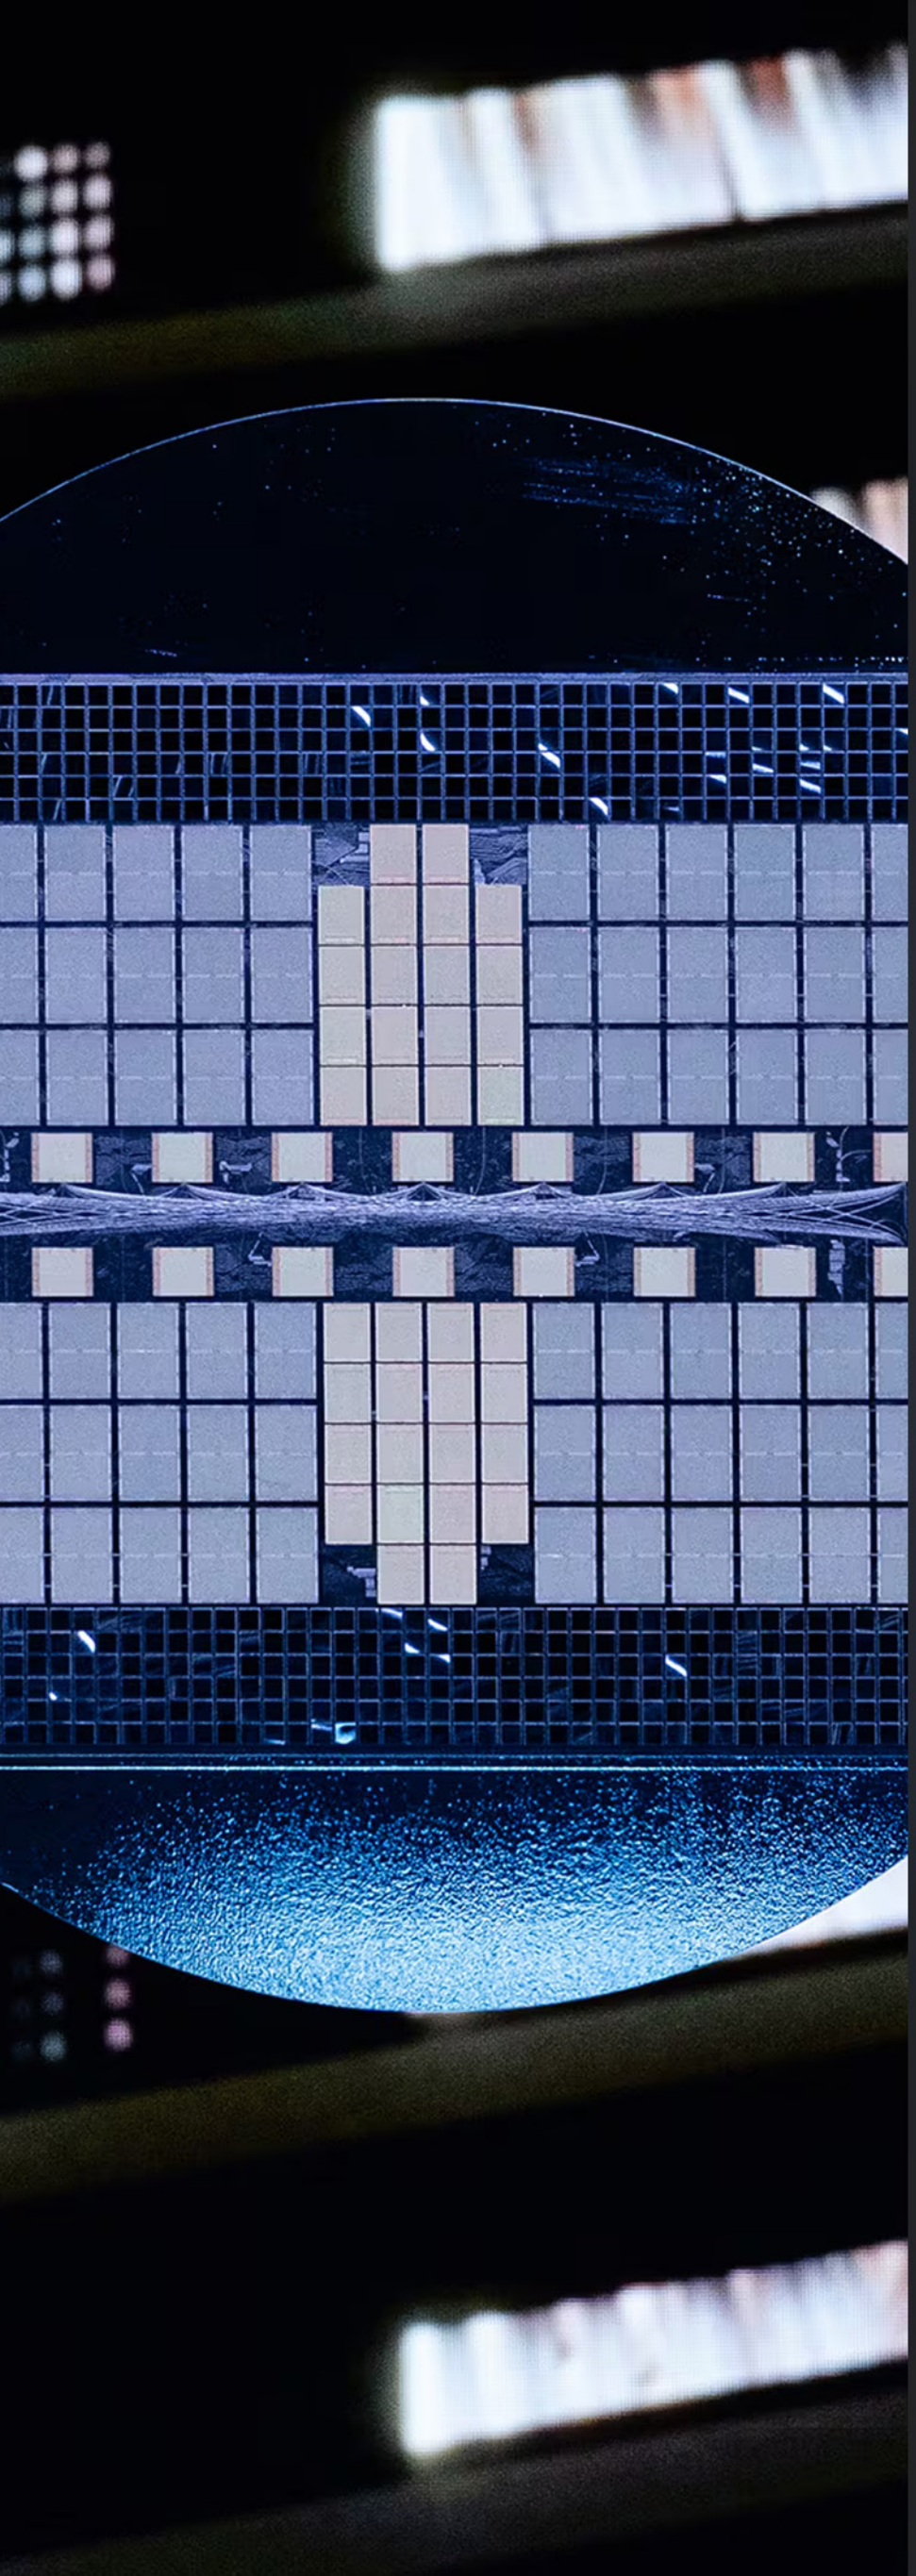

how companies are seeking diversified and adaptable infrastructure plans as compute needs escalate.

For competitors such as Google, these conditions create openings. Cloud-based AI services increasingly rely on specialized hardware offerings that integrate silicon, networking and software tooling. As Google advances its TPU roadmap and increases external availability, it gains opportunities to position its platform for organizations requiring additional compute pathways.

The ongoing developments also intersect with supply-chain factors. Securing long-term access to high-performance accelerators requires multi-year planning, coordinated manufacturing allocation and tight relationships with chip suppliers. As organizations expand AI workloads across global data-center footprints, the ability to secure sustained hardware access has become a critical part of infrastructure strategy.

These conditions shape the discussions between companies like Meta and potential hardware partners. As compute needs rise and infrastructure planning becomes increasingly complex, organizations continue evaluating multiple hardware choices across both cloud and internally hosted environments. Additional reviews, partnerships and testing cycles are expected as firms refine their long-term pipelines.

GOOGLE REASSERTS ITS POSITION IN AI AS PRODUCT MOMENTUM BUILDS


Google

Google is entering a new phase of activity in the artificial-intelligence sector after a period in which external perceptions often portrayed the company as lagging behind competitors. Recent developments across its product stack, internal infrastructure and research initiatives indicate a coordinated effort to accelerate the pace of deployment and restore a level of momentum that had defined the company's earlier leadership in machine learning. According to industry analysis referenced by the Financial Post, Google's core technical capabilities remain strong, and its renewed operational posture reflects the scale of its internal resources and long-term research investments.

The shift follows a year of debate across the technology industry about whether Google's approach to AI development had become too cautious relative to entrants pursuing more aggressive release cadences. While the discussion surrounding model safety, reliability and public trust has influenced many of Google's decisions, the company has simultaneously been preparing a set of platform upgrades that underscore its infrastructure depth and the breadth of its research pipeline. As these components move into active rollout cycles, Google's position within the competitive environment is evolving, supported by continued investment in engineering teams, data-center capacity and specialized silicon.


oogle AI
https://ai.google
e AI
AI
Responsibility
Disc
cing AI for eve
UR LATEST NEWS


Goog

INTERNAL CAPABILITIES AND INFRASTRUCTURE GIVE GOOGLE A STRONG TECHNICAL FOUNDATION

Google's extensive infrastructure forms the backbone of its ability to scale AI deployments across products and services. The company has long operated one of the world's most advanced distributed computing networks, built around custom-designed tensor processing units and data-center clusters optimized for large-scale training workloads. These systems support model development across multiple research groups, from language and vision to multimodal reasoning.

The Financial Post analysis highlights that the notion of Google "falling behind" does not fully account for these infrastructure advantages. Google has consistently produced breakthroughs in model architecture, multimodal representation and large-scale optimization. The company's research labs continue to publish foundational work cited across the field, and these developments inform upcoming product lines scheduled to be integrated into consumer and enterprise services.

In practice, Google's infrastructure advantages allow it to operate multiple parallel model families with high levels of redundancy and resource allocation flexibility. The training schedules for next-generation models can run continuously across distributed regions, enabling faster iteration cycles than would be possible with smaller compute networks. This setup also allows Google to test multi-track developmental approaches, where different variants of a model can evolve in parallel until

one demonstrates performance or efficiency characteristics suitable for deployment.

Google's model-safety teams, including groups dedicated to alignment, red-team testing and interpretability research, factor heavily into its developmental posture. While these processes introduce additional review cycles, they also reflect the company's approach to evaluating model behavior and assessing how different architectures perform under real-world usage conditions. As newer systems move through testing, evaluations increasingly include assessments of robustness, long-context reasoning, interaction stability and invariance across multilingual inputs.

Across its infrastructure stack, Google continues to integrate new forms of hardware acceleration. Recent internal utilization reports show increased dependence on next-generation TPU clusters as well as experimental compute modules tested in limited deployments. These systems form part of Google's plan to scale inference across search, cloud services, consumer applications and enterprise AI flows, creating a unified environment where model updates propagate across multiple product layers.

PRODUCT MOMENTUM RETURNS AS GOOGLE EXPANDS AI FEATURES ACROSS ITS ECOSYSTEM

Momentum across Google's consumer and enterprise products has been rising as new AI features enter public release cycles. The company has introduced expanded multimodal capabilities into search, communication tools

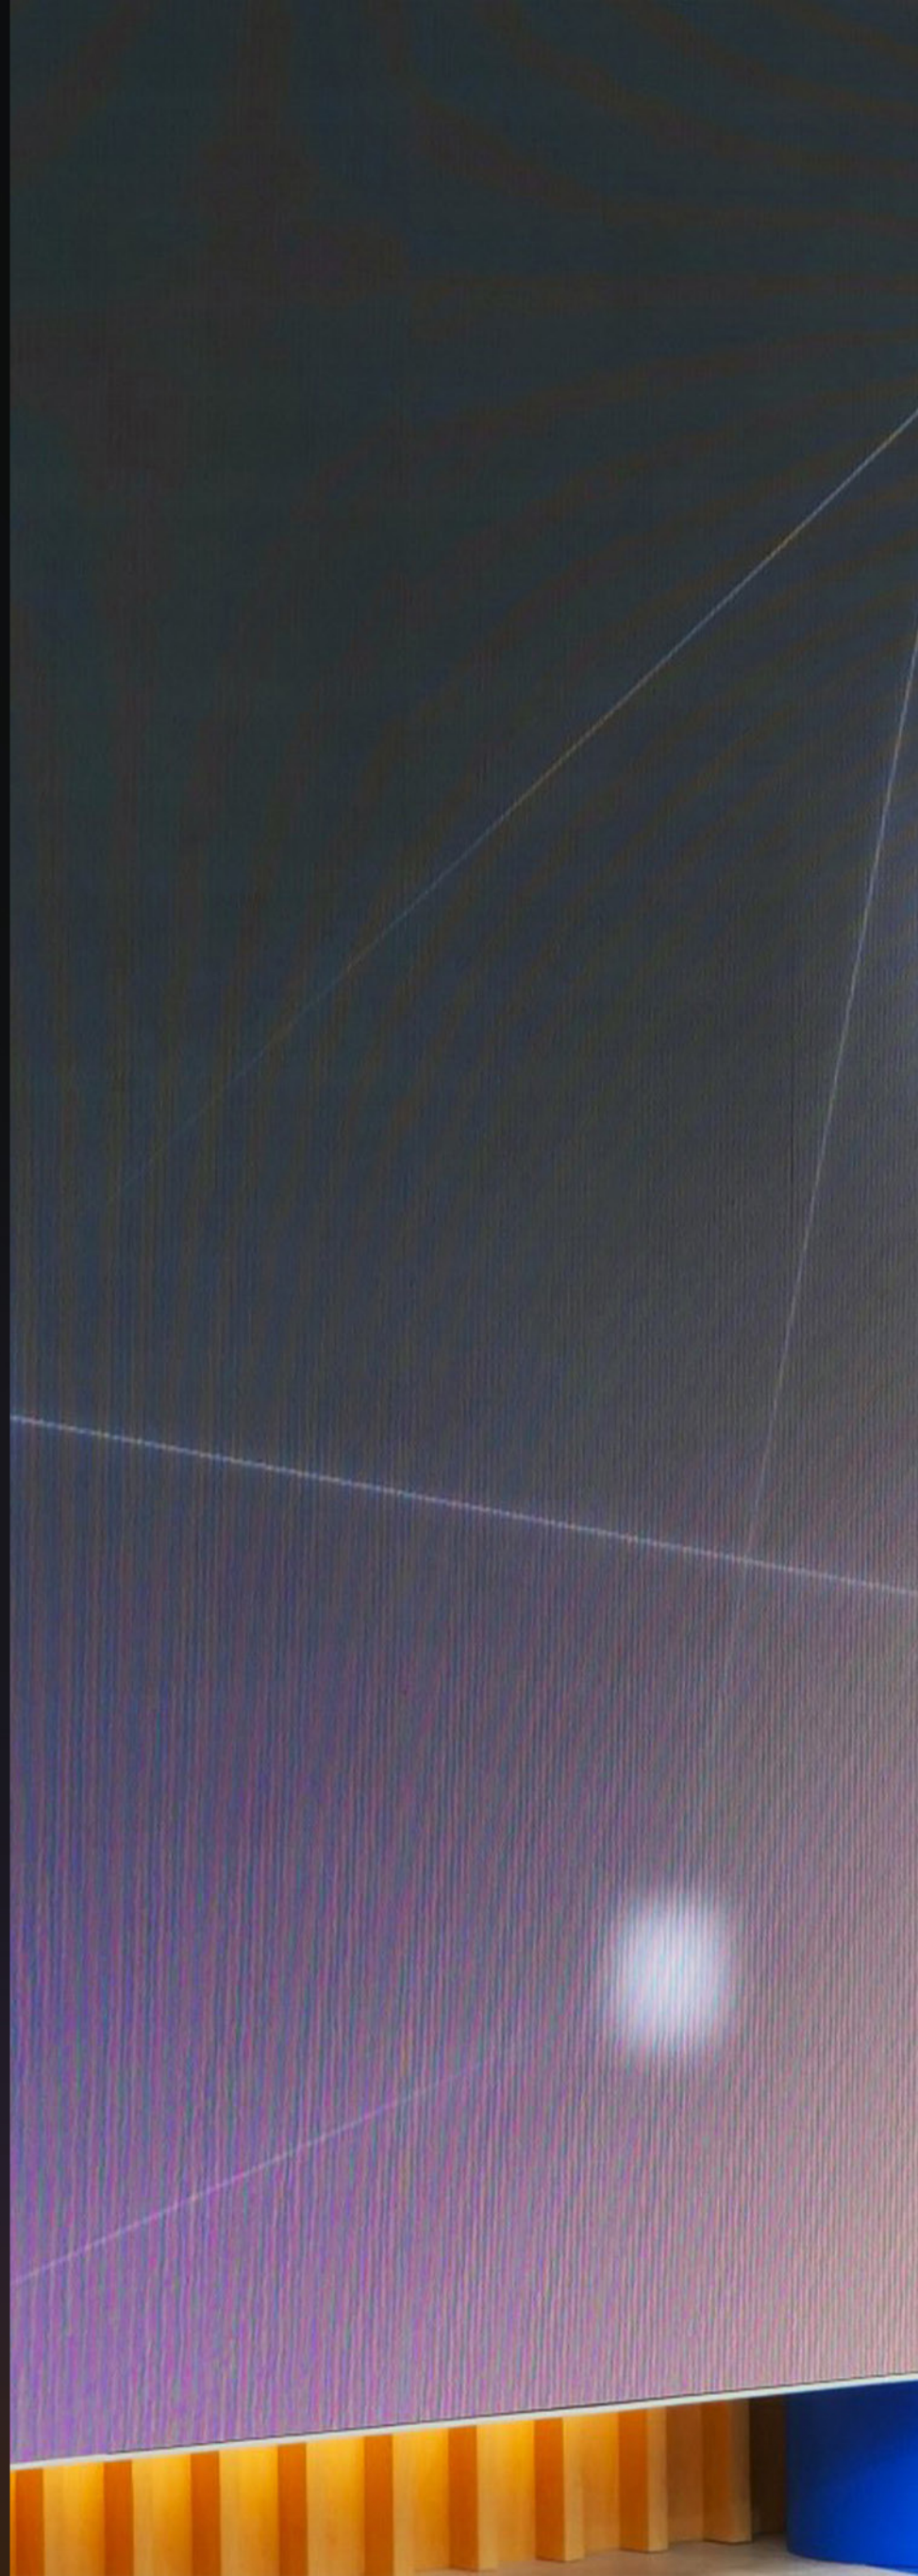


Gemini


Google
Celine Beunier

and cloud applications, reflecting a shift from small embedded features toward integrated system-level upgrades. These changes follow a year in which competitive pressure accelerated industry timelines for conversational AI, prompting rapid iteration among major developers.

Within Google's product organization, recent updates indicate closer coordination between research groups and application teams. This structure allows trained models to enter testing earlier in the product cycle and shortens the interval between research milestones and public release. While product readiness still depends on safety assessments and infrastructure scaling, the tighter integration has contributed to a steady stream of updates.

One significant development has been the integration of advanced AI capabilities into Google Cloud, where enterprise clients have access to tools for application development, code assistance, data analysis and workflow automation. The company's cloud division continues to expand APIs and platform extensions that connect customers to models with multimodal, reasoning and summarization abilities. These services are designed to operate across industries with varying regulatory requirements, making deployment flexibility a critical part of the design process.

Search also remains a central focus for Google's AI rollout. Upgrades in test environments include systems capable of synthesizing multimodal inputs, contextualizing queries across longer input windows and providing structured results that balance model-generated

information with verifiable reference content. These features reflect ongoing refinements to how large models interact with the information ecosystems that power Google's core business.

Hardware integrations have further contributed to the company's AI resurgence. Google's mobile and computing products now incorporate on-device model components for tasks such as image analysis, real-time translation and voice interactions, with cloud-linked systems expanding their capabilities. These integrations highlight Google's long-term strategy of distributing AI workloads between device and cloud in ways that reduce latency and preserve user privacy when possible.

The Financial Post report emphasizes that Google's reaccelerated product activity is not the result of abrupt strategic pivots but rather the visible output of internal efforts that have been underway for several years. As the company's development cycles align across hardware, software and cloud services, the resulting product cadence gives Google the capacity to reassert a stronger presence in the broader AI market.

COMPETITIVE POSITIONING SHIFTS AS GOOGLE INCREASES DEVELOPMENT PACE

The broader competitive environment around AI has been shaped by multiple companies deploying large-scale models across consumer and enterprise markets. While Google's pace previously appeared slower relative to firms opting for faster release schedules, the Financial Post notes that Google's reemergence reflects


Google


19:09
Google
Google LLC
4.3★
667K reviews
10M+
Downloads
Install
Try Gemini
an experimental

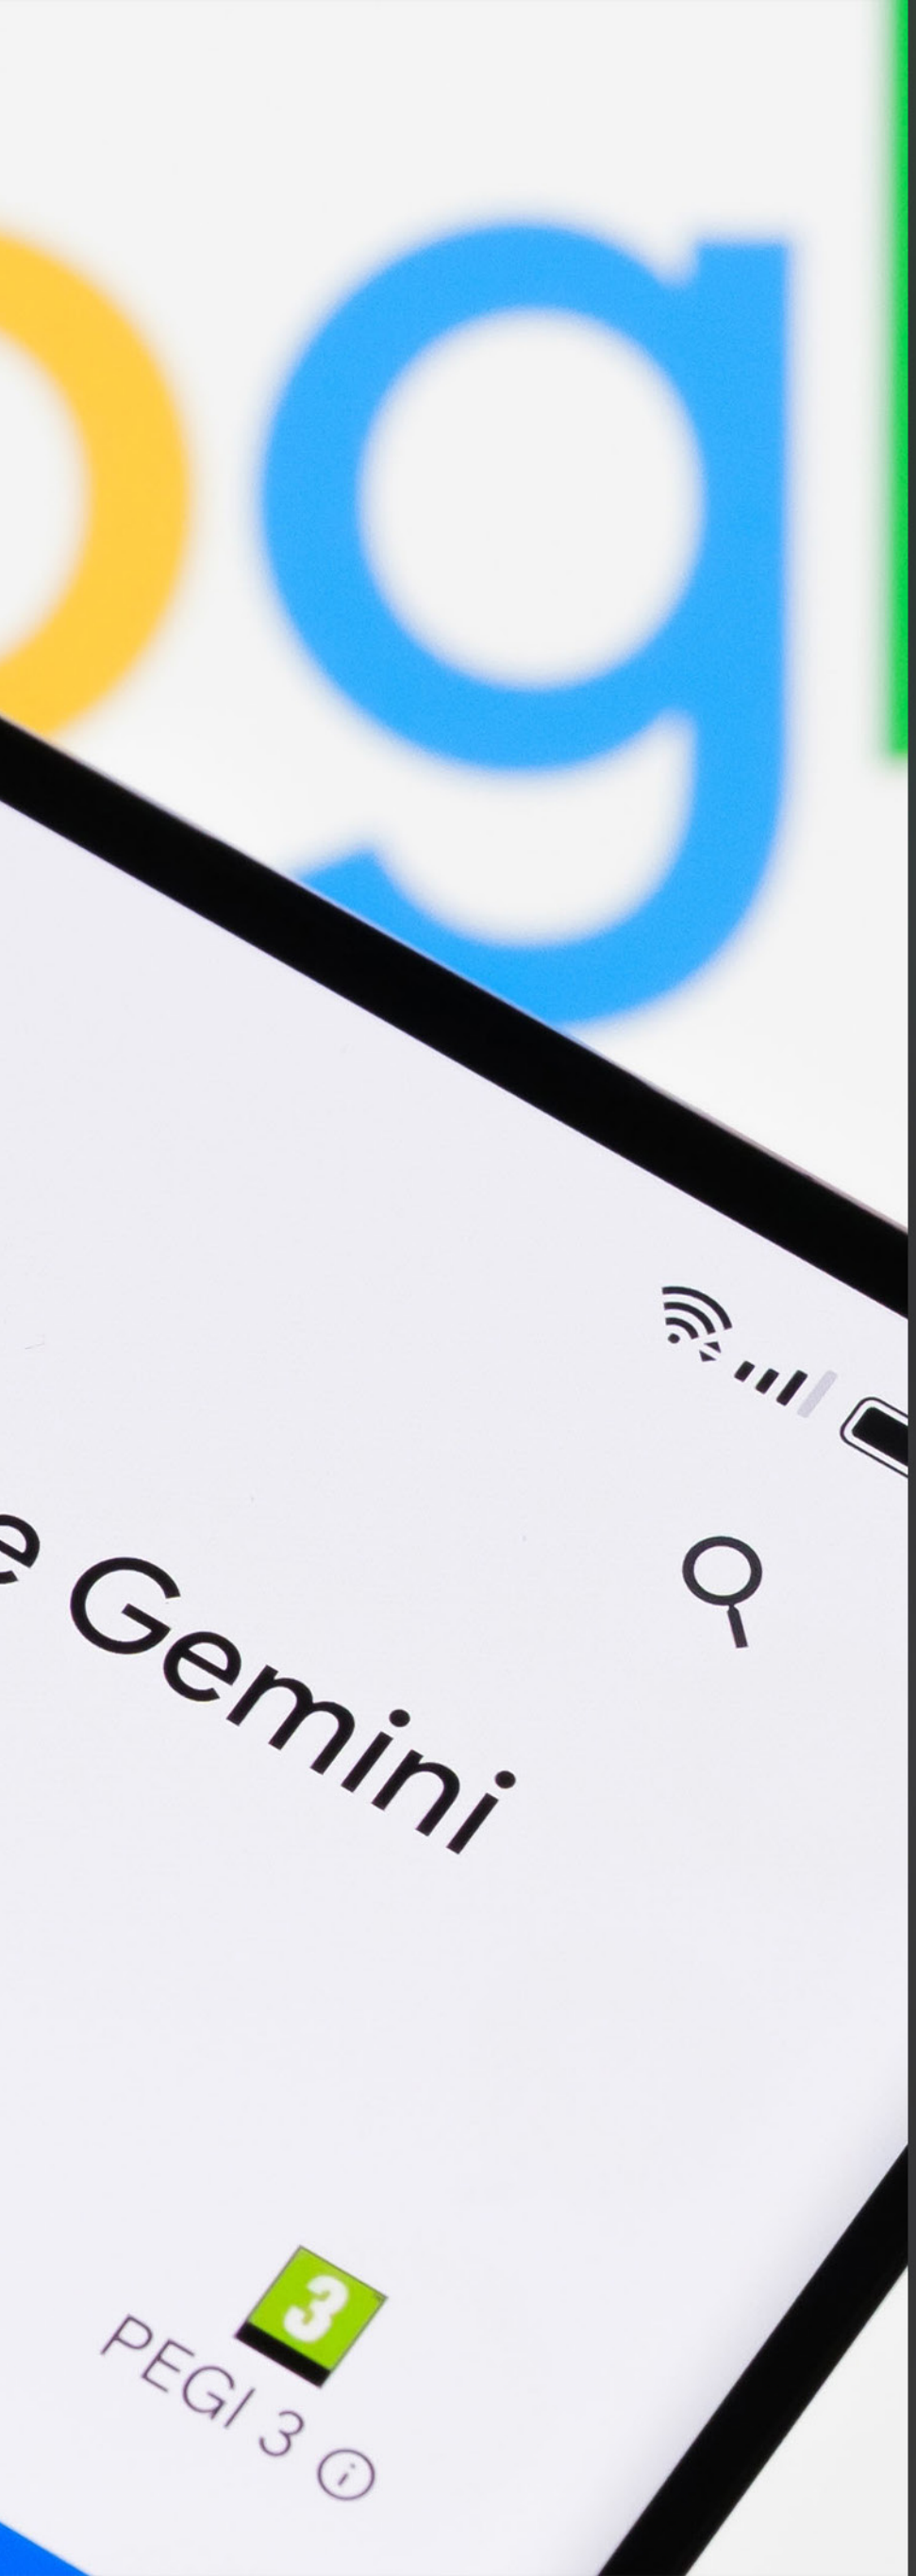

a convergence between its historical research strengths and its recent operational acceleration.

Google's ecosystem reach gives it a distinct positioning advantage: search, maps, mobile software, productivity applications and cloud infrastructure provide integrated environments where AI can be deployed at scale. As new models reach production readiness, they can be incorporated into multiple products in parallel, allowing Google's internal teams to refine feedback loops and adjust model behavior based on real-world usage patterns.

Enterprise analysts observing the current cycle have described Google's competitive posture as a function of its infrastructure depth, research continuity and product breadth. While competing systems continue to evolve rapidly, Google's multipronged development strategy provides a stable platform for multi-year AI deployment plans. This framework allows the company to distribute workloads across data centers, assess model performance across diverse usage patterns and transition upgrades into both cloud services and consumer-facing products.

The dynamics across the sector continue to shift as new models, hardware strategies and regulatory requirements emerge. For Google, the work now underway will influence how its AI systems integrate into search, cloud services, mobile platforms, enterprise workflows and developer ecosystems as it refines the next stages of its strategy.

APPLE DEPLOYS OVER 100 CAMERAS TO CAPTURE REAL MADRID IN ULTRA-IMMERSIVE FORMATS FOR VISION PRO

Apple has expanded its immersive-media push with a large-scale production effort centered on Real Madrid, using more than 100 dedicated cameras across Santiago Bernabéu Stadium to generate multi-angle footage tailored for Apple Vision Pro. The project integrates high-resolution volumetric capture, spatial video formats and synchronized camera arrays intended to create a viewing experience far beyond conventional broadcast coverage. According to production personnel involved, the system is designed to capture matches, player motion and crowd environments with enough spatial depth to support next-generation sports experiences on Vision Pro.

The deployment marks one of Apple's most extensive sports-film capture operations in Europe, incorporating fixed and mobile camera units across the stadium's upper decks, sidelines, player tunnels, and elevated gantries. Engineers and production teams have coordinated placement to ensure overlapping fields of view that allow Apple's post-production systems to construct spatial layers required for immersive playback. Multiple camera rigs were tested in recent months as part of Apple's wider effort to expand Vision Pro sports content, which includes collaborations across soccer, baseball, basketball and motorsport.

The Real Madrid project builds on earlier work Apple completed for the MLS Season Pass program and for select live events captured with its spatial video technologies. However, the scale of the Bernabéu installation stands out because it spans a full 360-degree capture grid, allowing Apple to create high-fidelity match perspectives

optimized for its headset's dual-4K micro-OLED displays and spatial audio features.

INSIDE APPLE'S STADIUM-WIDE CAMERA NETWORK

The camera grid involves more than 100 units distributed strategically throughout the stadium. Each unit aligns with a set of production requirements that dictate both coverage and redundancy. Upper-deck placements provide panoramic reference angles for large-scale spatial mapping, while field-level cameras feed detailed player-motion sequences into Apple's processing pipeline. The system incorporates both fixed and mobile devices, enabling synchronized capture as players move between formations and ball trajectories shift during play.

Apple's production team coordinated with Real Madrid's stadium operations staff and European broadcast partners to place cameras without disrupting existing infrastructure. Engineers positioned rigs to maintain consistency across lighting conditions, camera height and lens calibration. The company has also paired select rigs with depth-sensing modules designed to analyze player positioning and spatial relationships essential for reconstructed viewpoints. These components extend Apple's earlier experiments with spatial sports content and introduce new methods for assembling multi-layer video formats suitable for Vision Pro.

Part of the workflow includes capturing ambient crowd and stadium elements that give the Bernabéu its distinct character. Microphone arrays complement video units, collecting

positional audio that later integrates with the headset's spatial environment. Apple's engineers previously relied heavily on controlled environments for spatial audio development, but the Bernabéu project shifts that work into a live-stadium setting at full match capacity.

The organization of more than 100 cameras requires continuous coordination, with real-time monitoring to ensure each feed remains synchronized throughout the match. Apple uses dedicated production trucks and off-site data links to move the footage into processing labs where volumetric reconstruction occurs. This approach parallels workflows Apple has tested in the United States but adapts them to European stadium architecture and lighting environments.

REAL MADRID'S ROLE IN APPLE'S EXPANDING IMMERSIVE SPORTS STRATEGY

Real Madrid's involvement reflects the club's ongoing engagement with emerging media formats. The Bernabéu has undergone years of renovation designed in part to support modern broadcast and digital-production requirements, with expanded cabling access, reinforced gantries, retractable roof structures and multi-use infrastructure. These features align with Apple's need for stable camera platforms, high-capacity data routes and tightly controlled environmental conditions.

The partnership also positions the club at the center of Apple's expanding spatial-content strategy. While the company has developed immersive features for various sports, soccer


ORID C.F.
Realmadrid
51

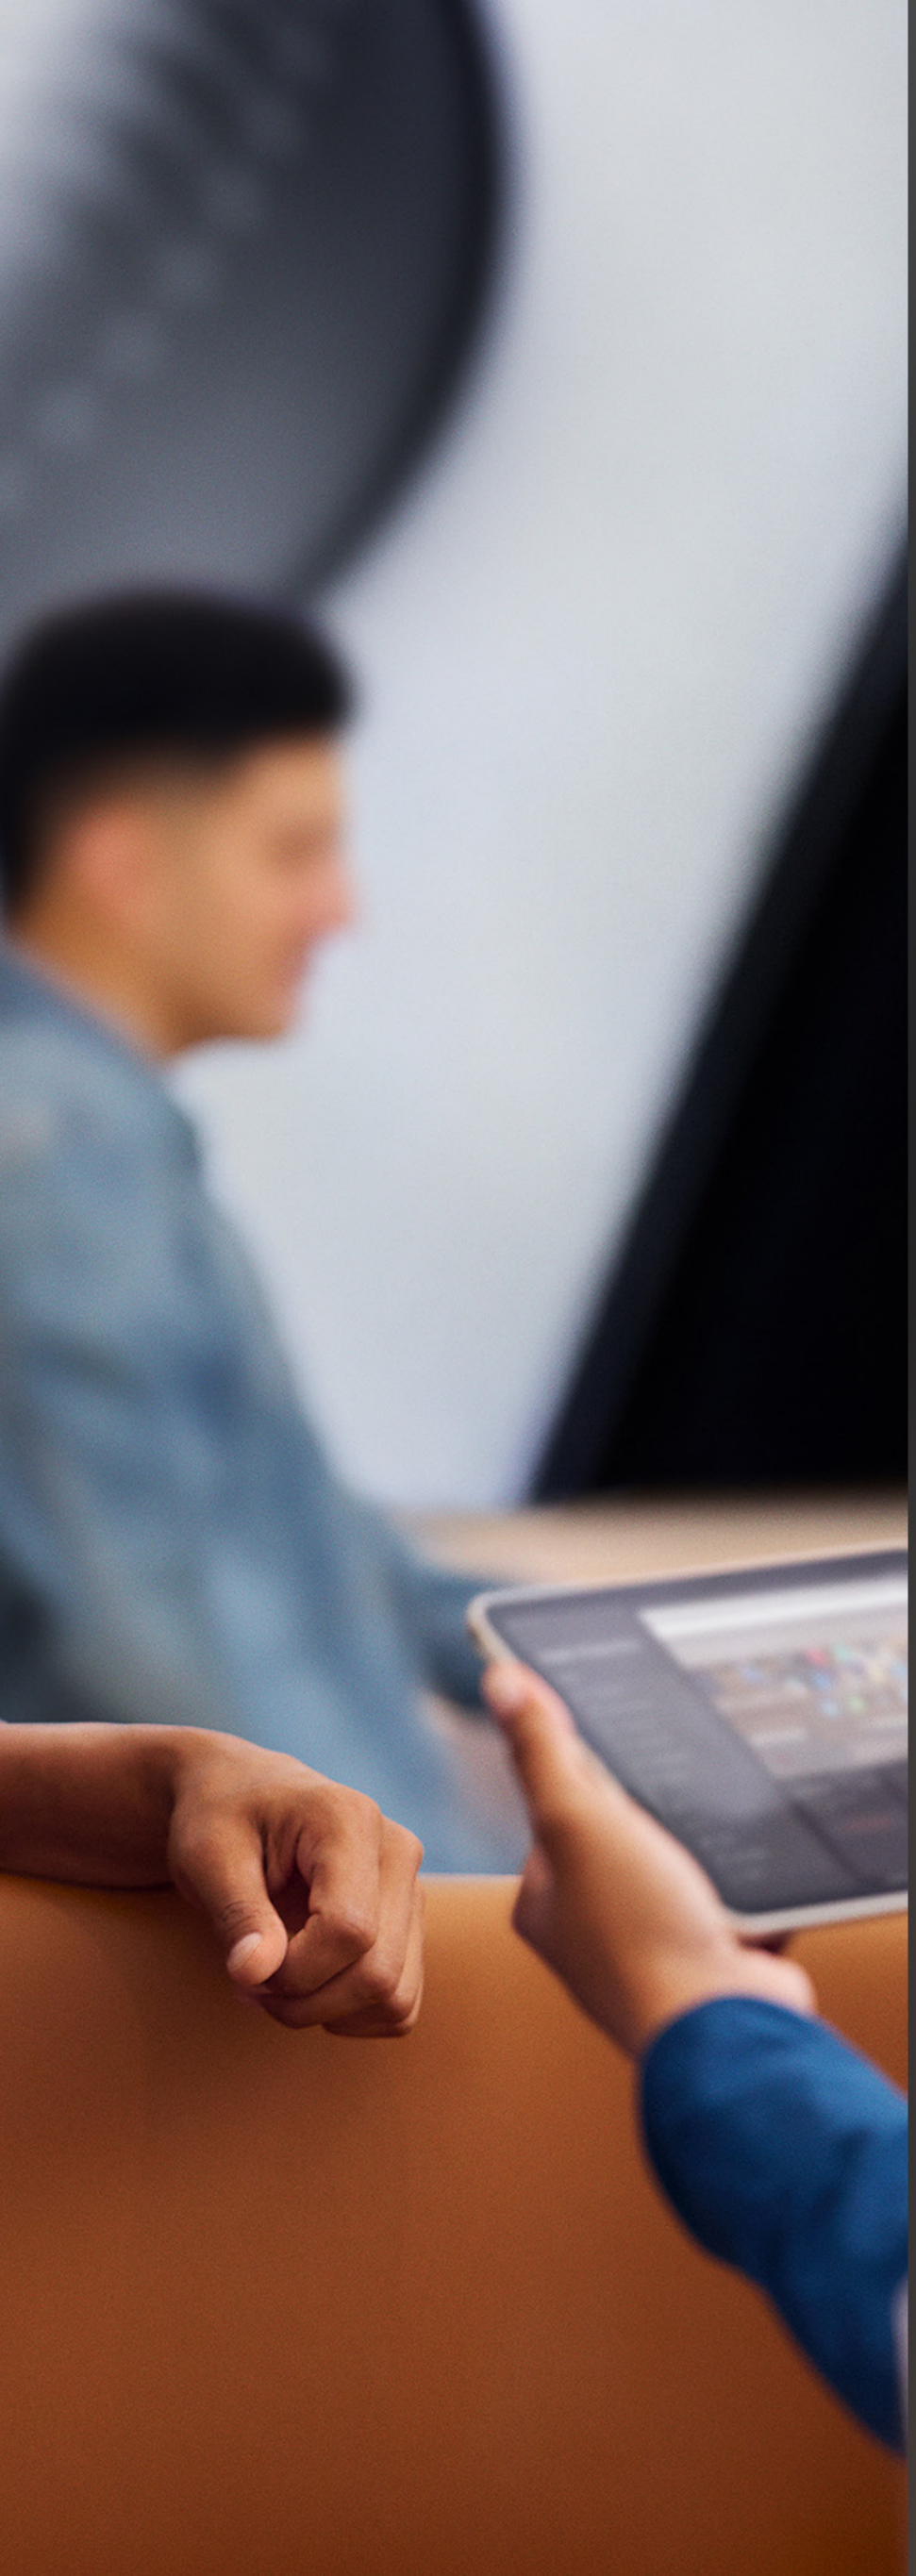

presents unique challenges due to field size, continuous play patterns and the high-speed distribution of players across large areas. Capturing the sport in formats suitable for immersive devices requires precise tracking across wide surfaces, consistent depth reference and reliable vantage points that can be stitched together without visual gaps.

Real Madrid's global following gives Apple a high-profile test environment for these techniques. The production team benefits from the stadium's lighting design, which has been engineered to meet stringent broadcast requirements and minimize color-shift variance that could complicate spatial reconstruction. Apple's systems use this controlled lighting environment to reduce calibration drift across camera arrays, helping to ensure uniform rendering quality during post-production.

Apple's move into large-scale stadium capture comes at a moment when technology and sports organizations across Europe are experimenting with advanced forms of fan engagement. Real Madrid's infrastructure enables Apple to test how immersive content can integrate with matchday experiences, digital-ticketing ecosystems and international distribution platforms. Although the project focuses on Vision Pro content, the data gathered across multi-camera rigs may inform Apple's long-term strategies for spatial frameworks in entertainment, training analysis and archival preservation.

WHAT THE PRODUCTION MEANS FOR APPLE'S SPATIAL VIDEO ECOSYSTEM

The addition of full-stadium spatial video from Real Madrid complements a growing catalog of sports content Apple has been preparing since Vision Pro's launch. The company has already tested spatial recording formats for basketball, baseball and motorsports, and it introduced early spatial sports sequences within Apple TV's immersive content line. The Bernabéu project contributes to a broader dataset designed to refine Apple's spatial capture algorithms, improve rendering fidelity and support new forms of narrative presentation.

Part of Apple's long-term work involves adapting immersive capture to differing venue geometries. Stadiums vary widely in architecture, lighting, seating elevation and field clearance, all of which influence how well multi-camera rigs can collect consistent spatial data. Santiago Bernabéu's recent renovation has given Apple a rare level of predictability in these conditions, allowing the company to expand testing at scale and refine camera-placement models that could later apply to other venues across Europe and North America.

Apple's spatial video development is also informed by the need to handle large volumes of data. Capturing a full match with more than 100 cameras across both fixed and mobile systems generates substantial storage, requiring high-throughput processing workflows. Apple uses internal compression and metadata-tagging systems to segment footage into layers, enabling more efficient reconstruction for Vision Pro's playback environment. Apple's engineers


VINI JR

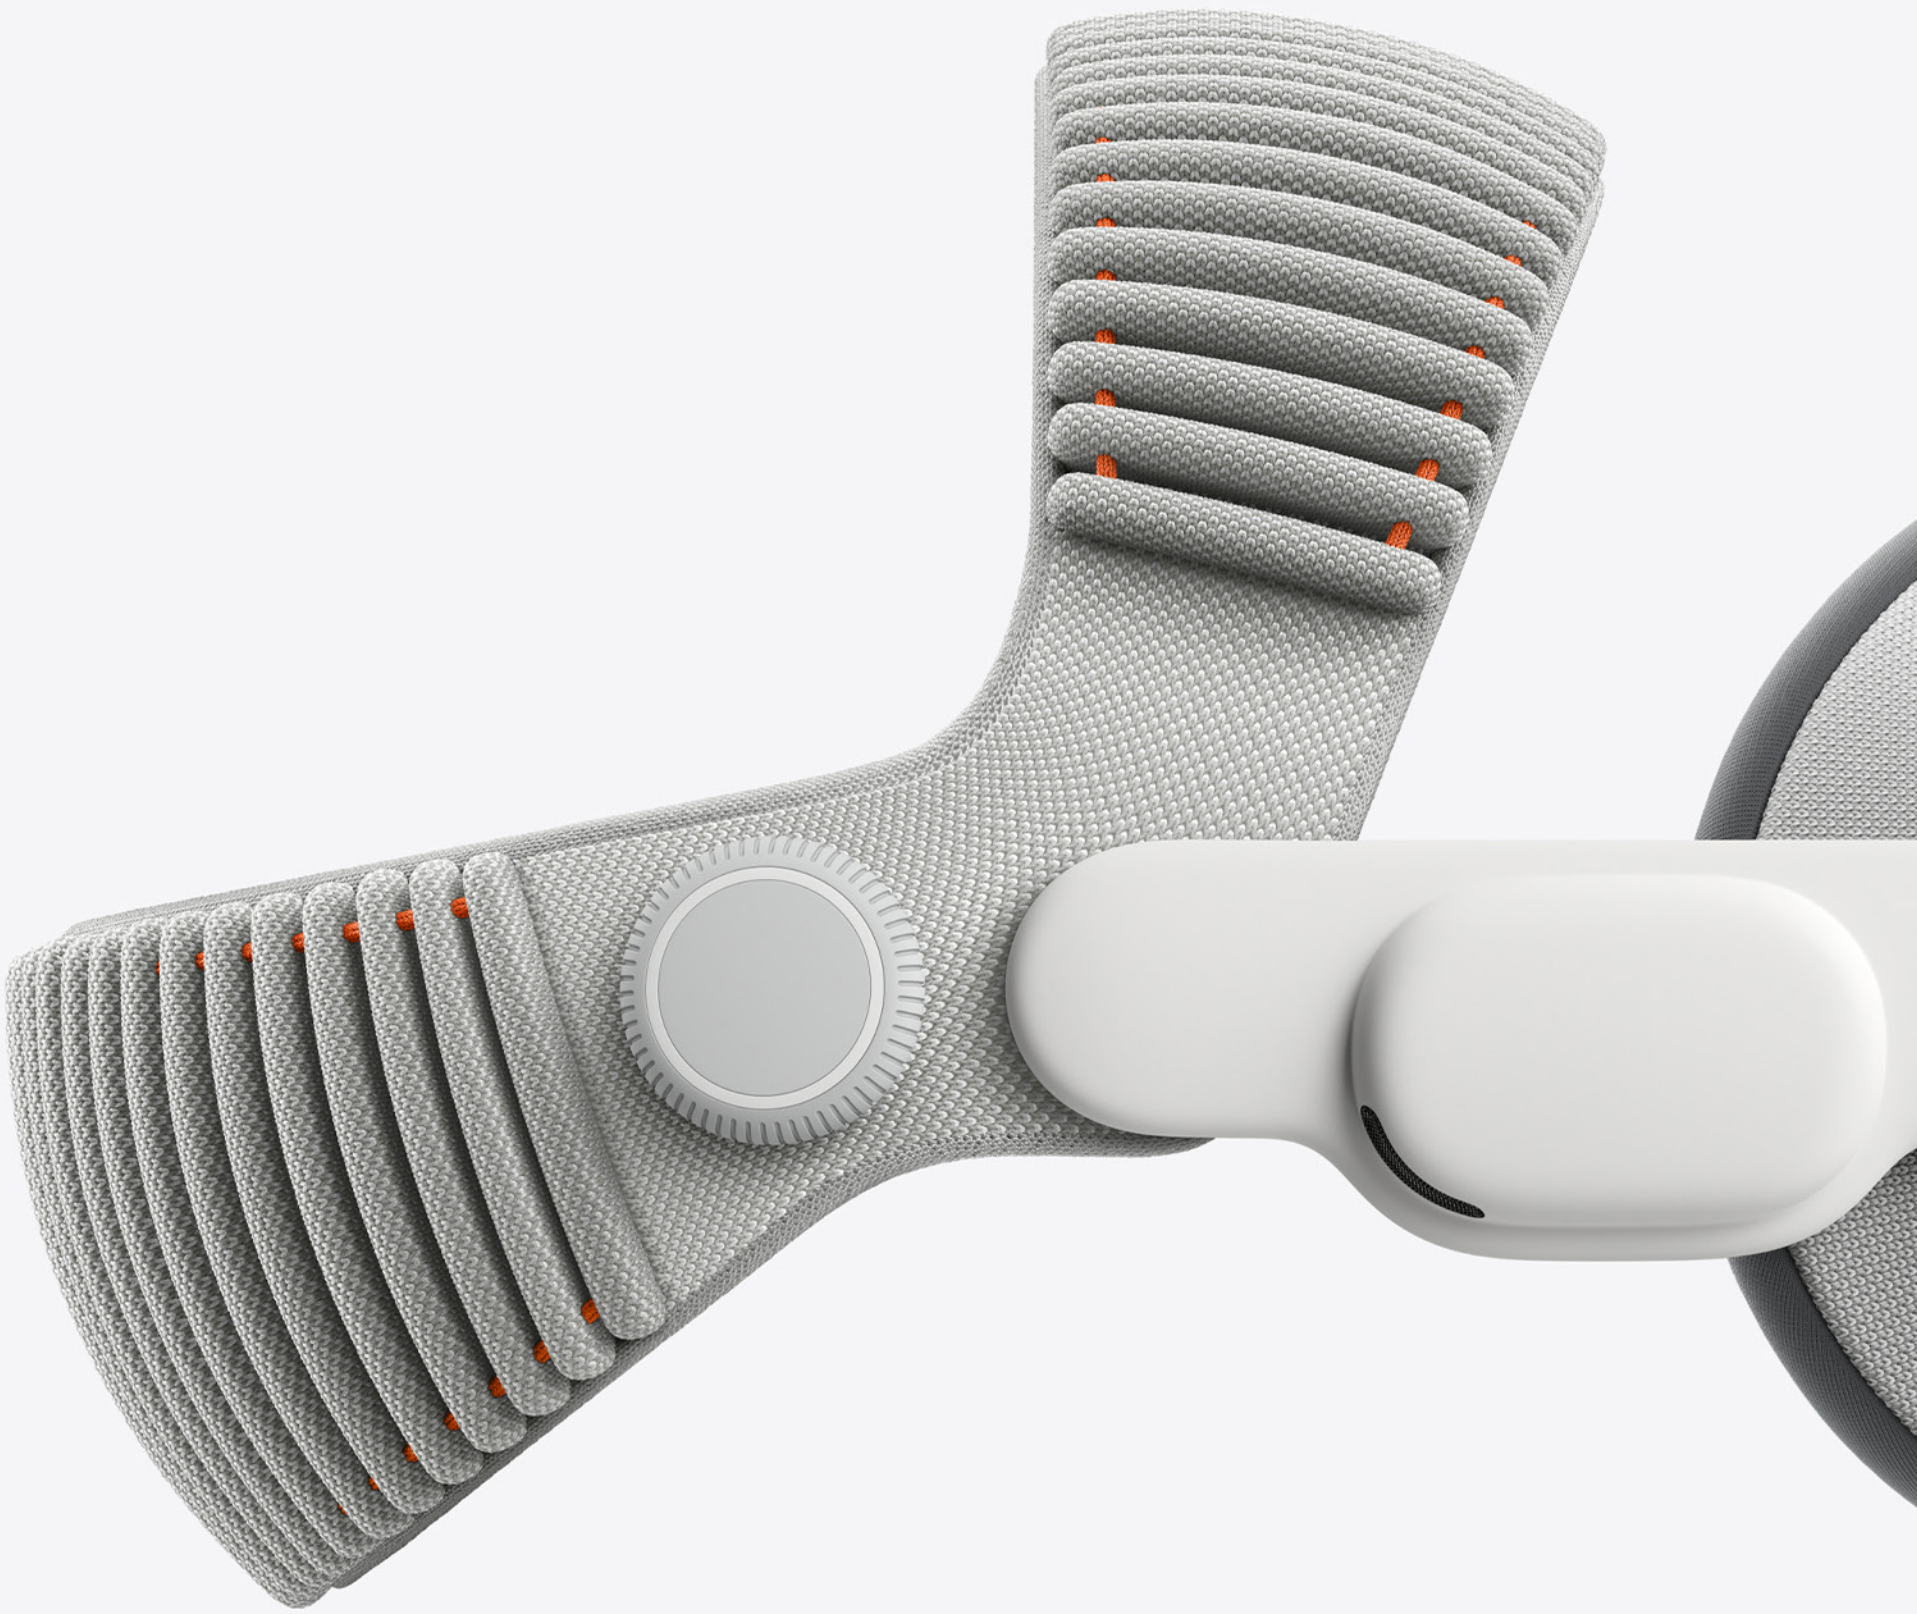

are testing how these layers interact with the headset's rendering pipeline, particularly its ability to generate smooth parallax effects and maintain clarity across motion-heavy sequences.

The presence of synchronized rigs across the stadium also supports experiments in viewpoint flexibility. Although the content is not expected to be fully volumetric in the sense of free movement across any point in the field, Apple's multi-camera approach allows for controlled perspective transitions that surpass what traditional broadcast feeds provide. These transitions can place the viewer closer to the field, behind the goal line, or in elevated sideline perspectives without the distortion that occurs when trying to artificially shift angles from a single camera feed.

Apple's team continues to adjust calibration models to address challenges such as player occlusion, rapid directional shifts and variations in crowd density. The data from each match contributes to Apple's efforts to build a more stable foundation for future immersive sports experiences on Vision Pro and similar devices.


NVIDIA

CHINESE REGULATORS RESTRICT BYTEDANCE FROM USING NVIDIA CHIPS, ALTERING AI SUPPLY DYNAMICS


Byte

ByteDance has been blocked by Chinese regulators from using certain NVIDIA chips in its data-center operations. The restriction marks a notable development in China's oversight of advanced semiconductor deployment, adding a layer of regulatory complexity for one of the country's largest technology firms and raising questions about how local companies will adjust their AI-compute strategies under tightened controls.

The decision, as described in the reports reviewed by Reuters, reflects regulators' concerns over how advanced accelerators are utilized inside large domestic platforms. While China continues to encourage AI development across multiple industries, the treatment of high-end chips used for large-scale model training and inference is subject to approval processes that vary by company and workload. ByteDance's restrictions appear specific to select NVIDIA models rather than the full range of available accelerators, creating a narrower but still consequential constraint on the company's compute roadmap.

ByteDance, the parent of TikTok and Douyin, operates extensive AI pipelines that power its content-ranking engines, recommendation systems and expanding generative-modeling research. The inability to use specific NVIDIA parts affects how the company can structure future training clusters and may shift more of its workloads toward either alternative domestic hardware or internal chip-development pathways. The regulatory decision places ByteDance in a distinct position relative to other Chinese firms that have received approval for certain NVIDIA products under localized guidelines.

REGULATORY SCRUTINY TIGHTENS AROUND HIGH-END ACCELERATORS IN CHINA

The restriction comes as Chinese authorities continue evaluating how advanced semiconductor hardware is deployed across large technology companies. Regulators examine how compute clusters are scaled, how data governance frameworks operate, and how hardware availability influences domestic AI development trajectories. In this context, ByteDance's use of high-end NVIDIA accelerators underwent a separate review that resulted in limits not imposed on all companies across the sector.

Officials have not publicly detailed the specific criteria behind the decision, but the action aligns with a broader regulatory environment in which chip utilization, data-processing practices and alignment with national guidelines are reviewed in parallel. These evaluations influence whether companies can access particular accelerator models and under what conditions cluster deployments are permitted to proceed.

For ByteDance, the restriction affects the company's near-term compute planning. Its data centers rely on large clusters to support both recommendation-engine tuning and experimentation with next-generation model architectures. The inability to source or deploy certain NVIDIA chips may require adjustments to its internal workload allocation, distribution of training cycles and hardware-specific optimization strategies. While Chinese firms have access to a growing number of domestic accelerators, performance and availability vary across manufacturers, creating operational

VIDIA®

considerations for companies accustomed to NVIDIA's mature software stack and ecosystem.

Other large Chinese technology companies have continued to explore mixed hardware strategies that combine domestic accelerators and NVIDIA products where allowed. The regulatory distinction applied to ByteDance shows how permissioning can diverge across organizations depending on their specific platform structures, data-scale realities and regulatory assessments.

IMPACT ON BYTEDANCE'S AI SYSTEMS AND COMPUTE STRATEGY

ByteDance runs some of the most demanding AI workloads in China, spanning content-ranking systems, generative-media tools, search features and moderation systems for multiple platforms. These workloads depend on the ability to train models continuously, update embeddings and integrate feedback loops that power personalized recommendations. Restrictions on certain NVIDIA chips could impact how these pipelines evolve, especially where training efficiency and inference density rely on accelerator classes known for high throughput.

The company's long-term efforts to develop or adapt its own compute solutions may become more prominent as a result. ByteDance has explored internal chip programs and partnerships with domestic semiconductor firms, though the maturity of these efforts is not uniform across model types. The regulatory constraint may accelerate the company's assessment of alternative hardware paths, requiring more reliance on domestic accelerators even where software-ecosystem support differs from NVIDIA's.


ByteDance

In the near term, ByteDance is expected to redistribute workloads toward available hardware, adjust batch-training cycles and evaluate how model architecture choices interact with the performance characteristics of domestic chips. These operational changes are common among Chinese technology companies navigating hardware constraints, but ByteDance's scale increases the impact of such adjustments on internal engineering planning.

The company's data-center architecture will also shape how quickly it can adapt. Large installations built around specific NVIDIA-optimized layouts may require reconfiguration to operate efficiently with mixed hardware sources. Decisions about new cluster construction or expansion may be influenced by the expectation of continued regulatory review, guiding ByteDance toward longer-term diversification of compute suppliers.

The development adds to the broader complexity of China's AI-hardware environment following multiple layers of international export controls and domestic oversight. Chinese firms continue to build out local accelerator ecosystems, with companies such as Huawei, Alibaba's T-Head and Biren developing chips aimed at competing with NVIDIA's offerings. The performance gaps between domestic and international accelerators vary by model class, influencing how companies allocate workloads.

Regulatory decisions that apply uniquely to specific firms create differentiated pressures across the ecosystem. Companies that receive approval for certain NVIDIA accelerators can structure training pipelines more like global peers, while firms facing restrictions may adopt hybrid

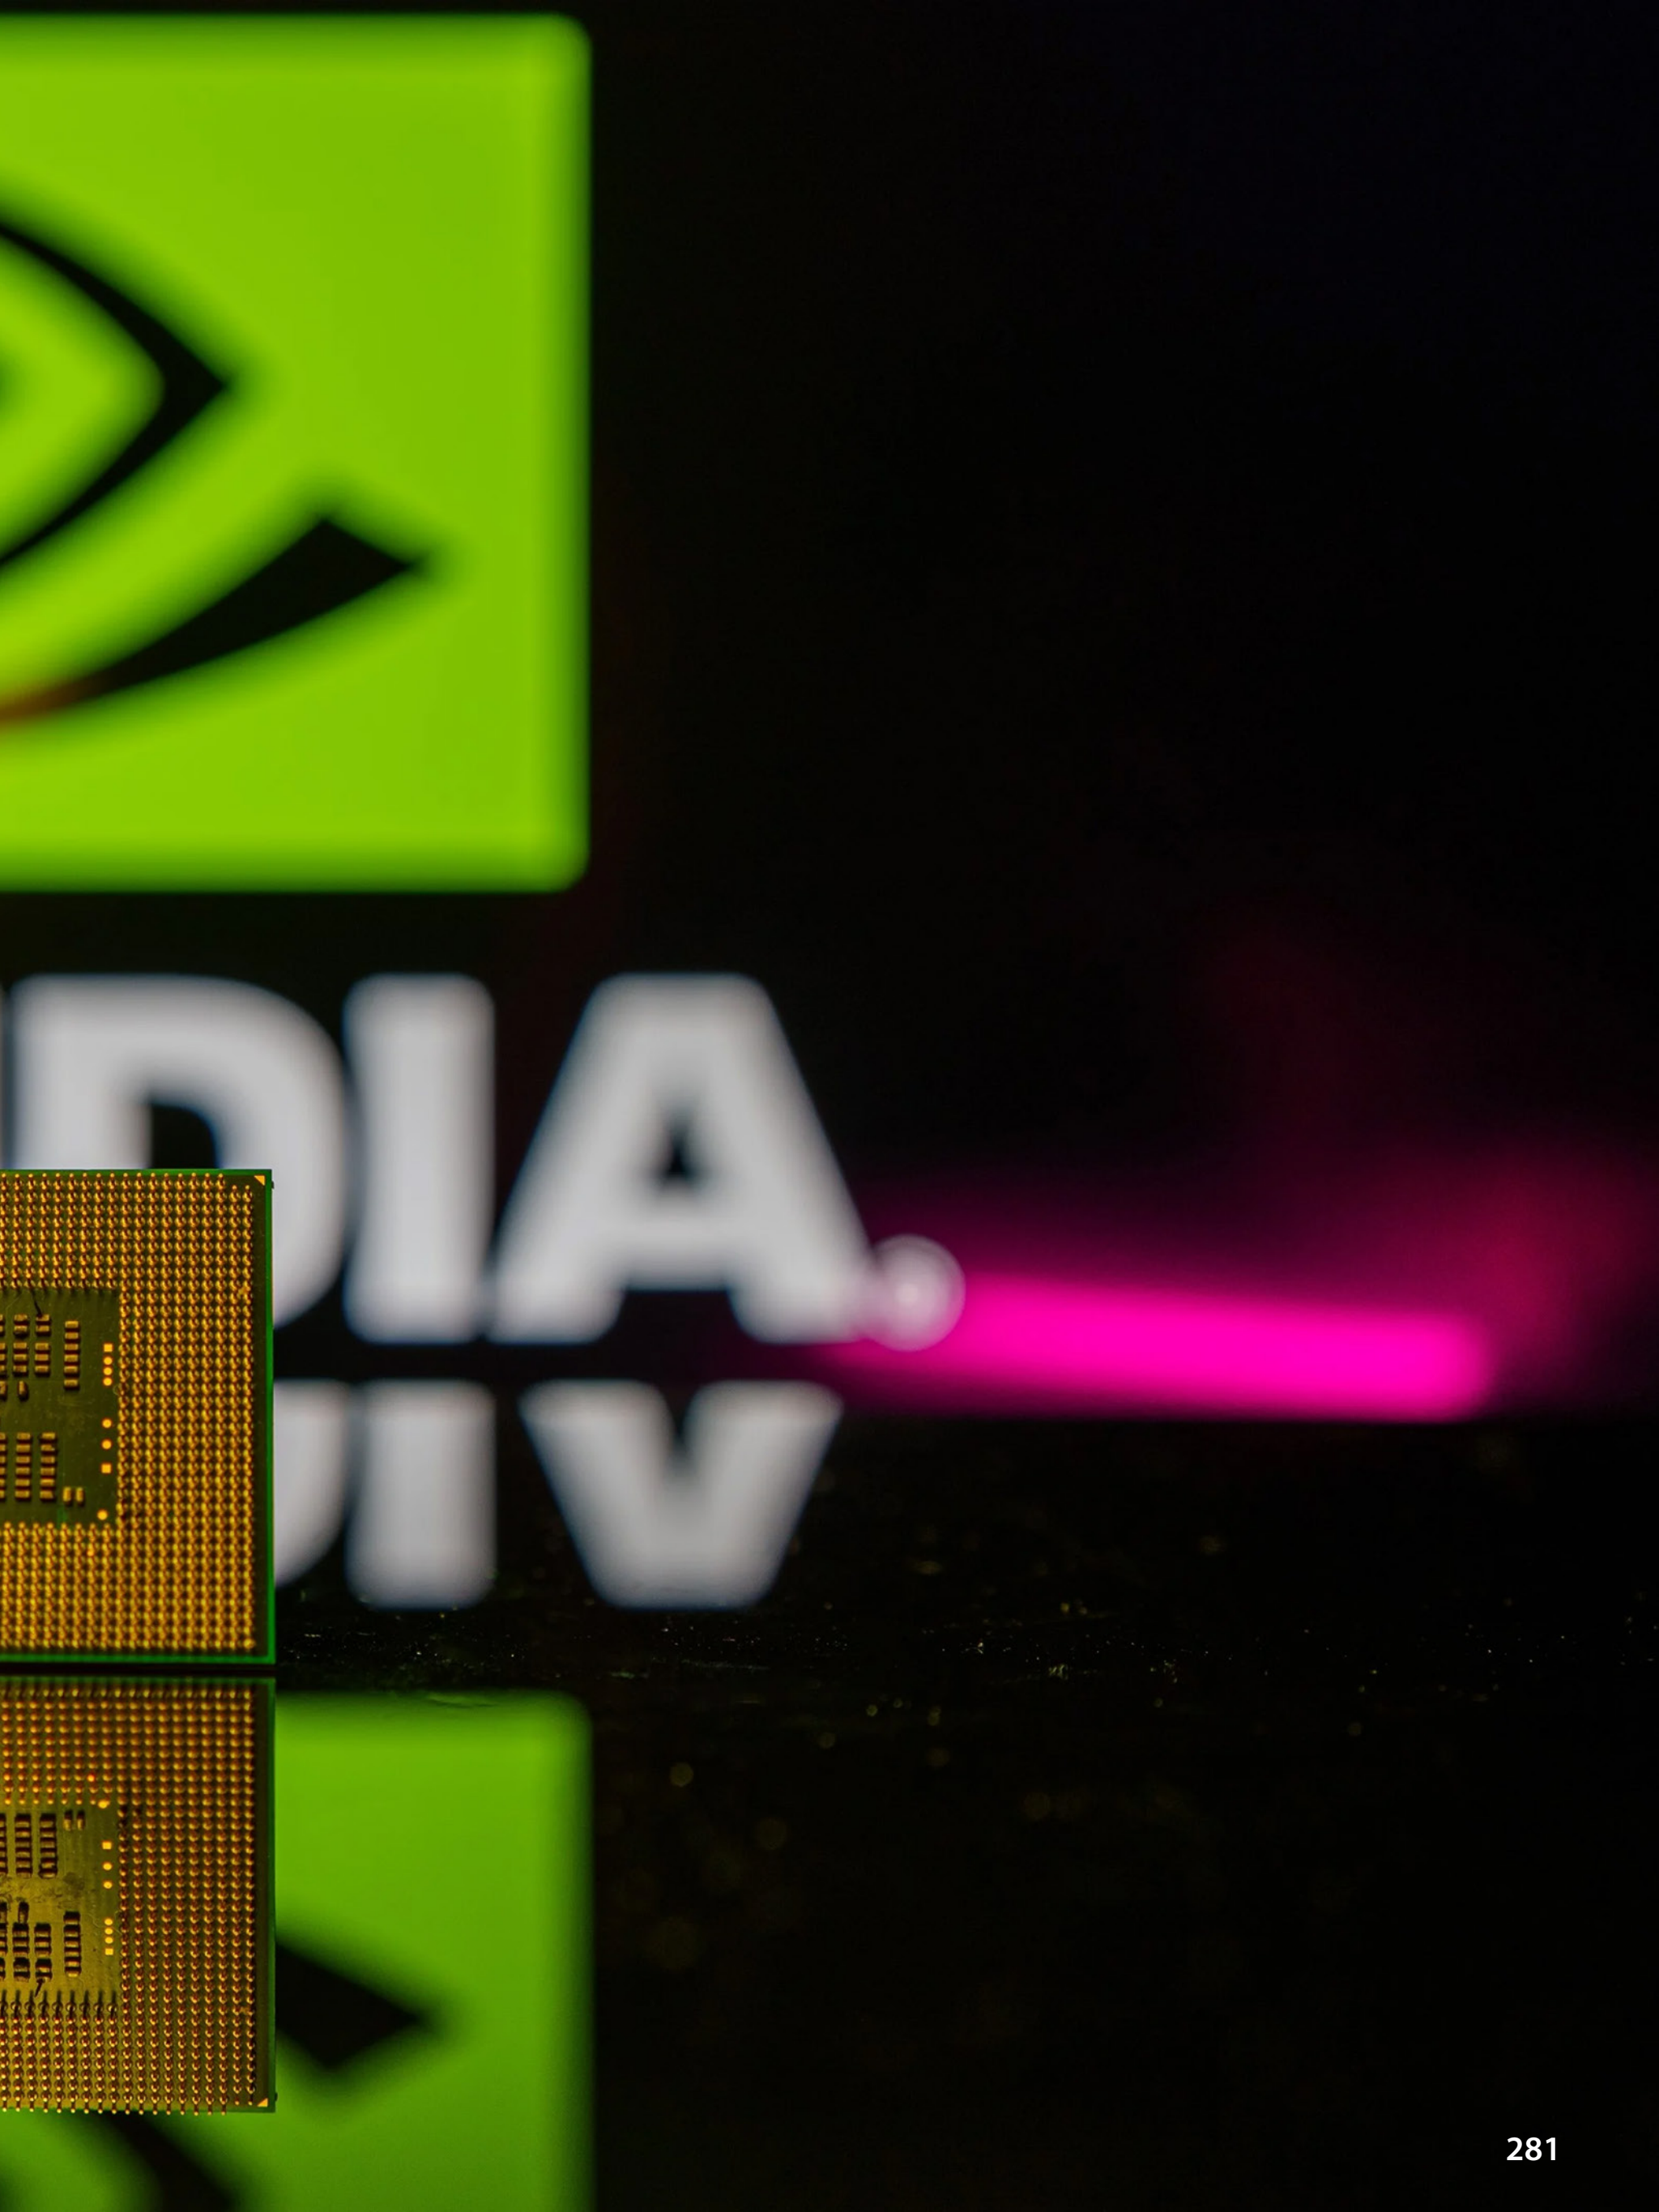

architectures sooner. These dynamics affect domestic chip demand, influencing investment flows across China's semiconductor sector.

ByteDance's response to the restriction will be watched closely by industry analysts, particularly as companies balance the need to scale large-model training with the realities of chip availability and regulatory review. The direction ByteDance takes—whether accelerating its own chip programs, expanding partnerships or increasing reliance on domestic accelerators—will shape how its platforms continue to evolve under these constraints.

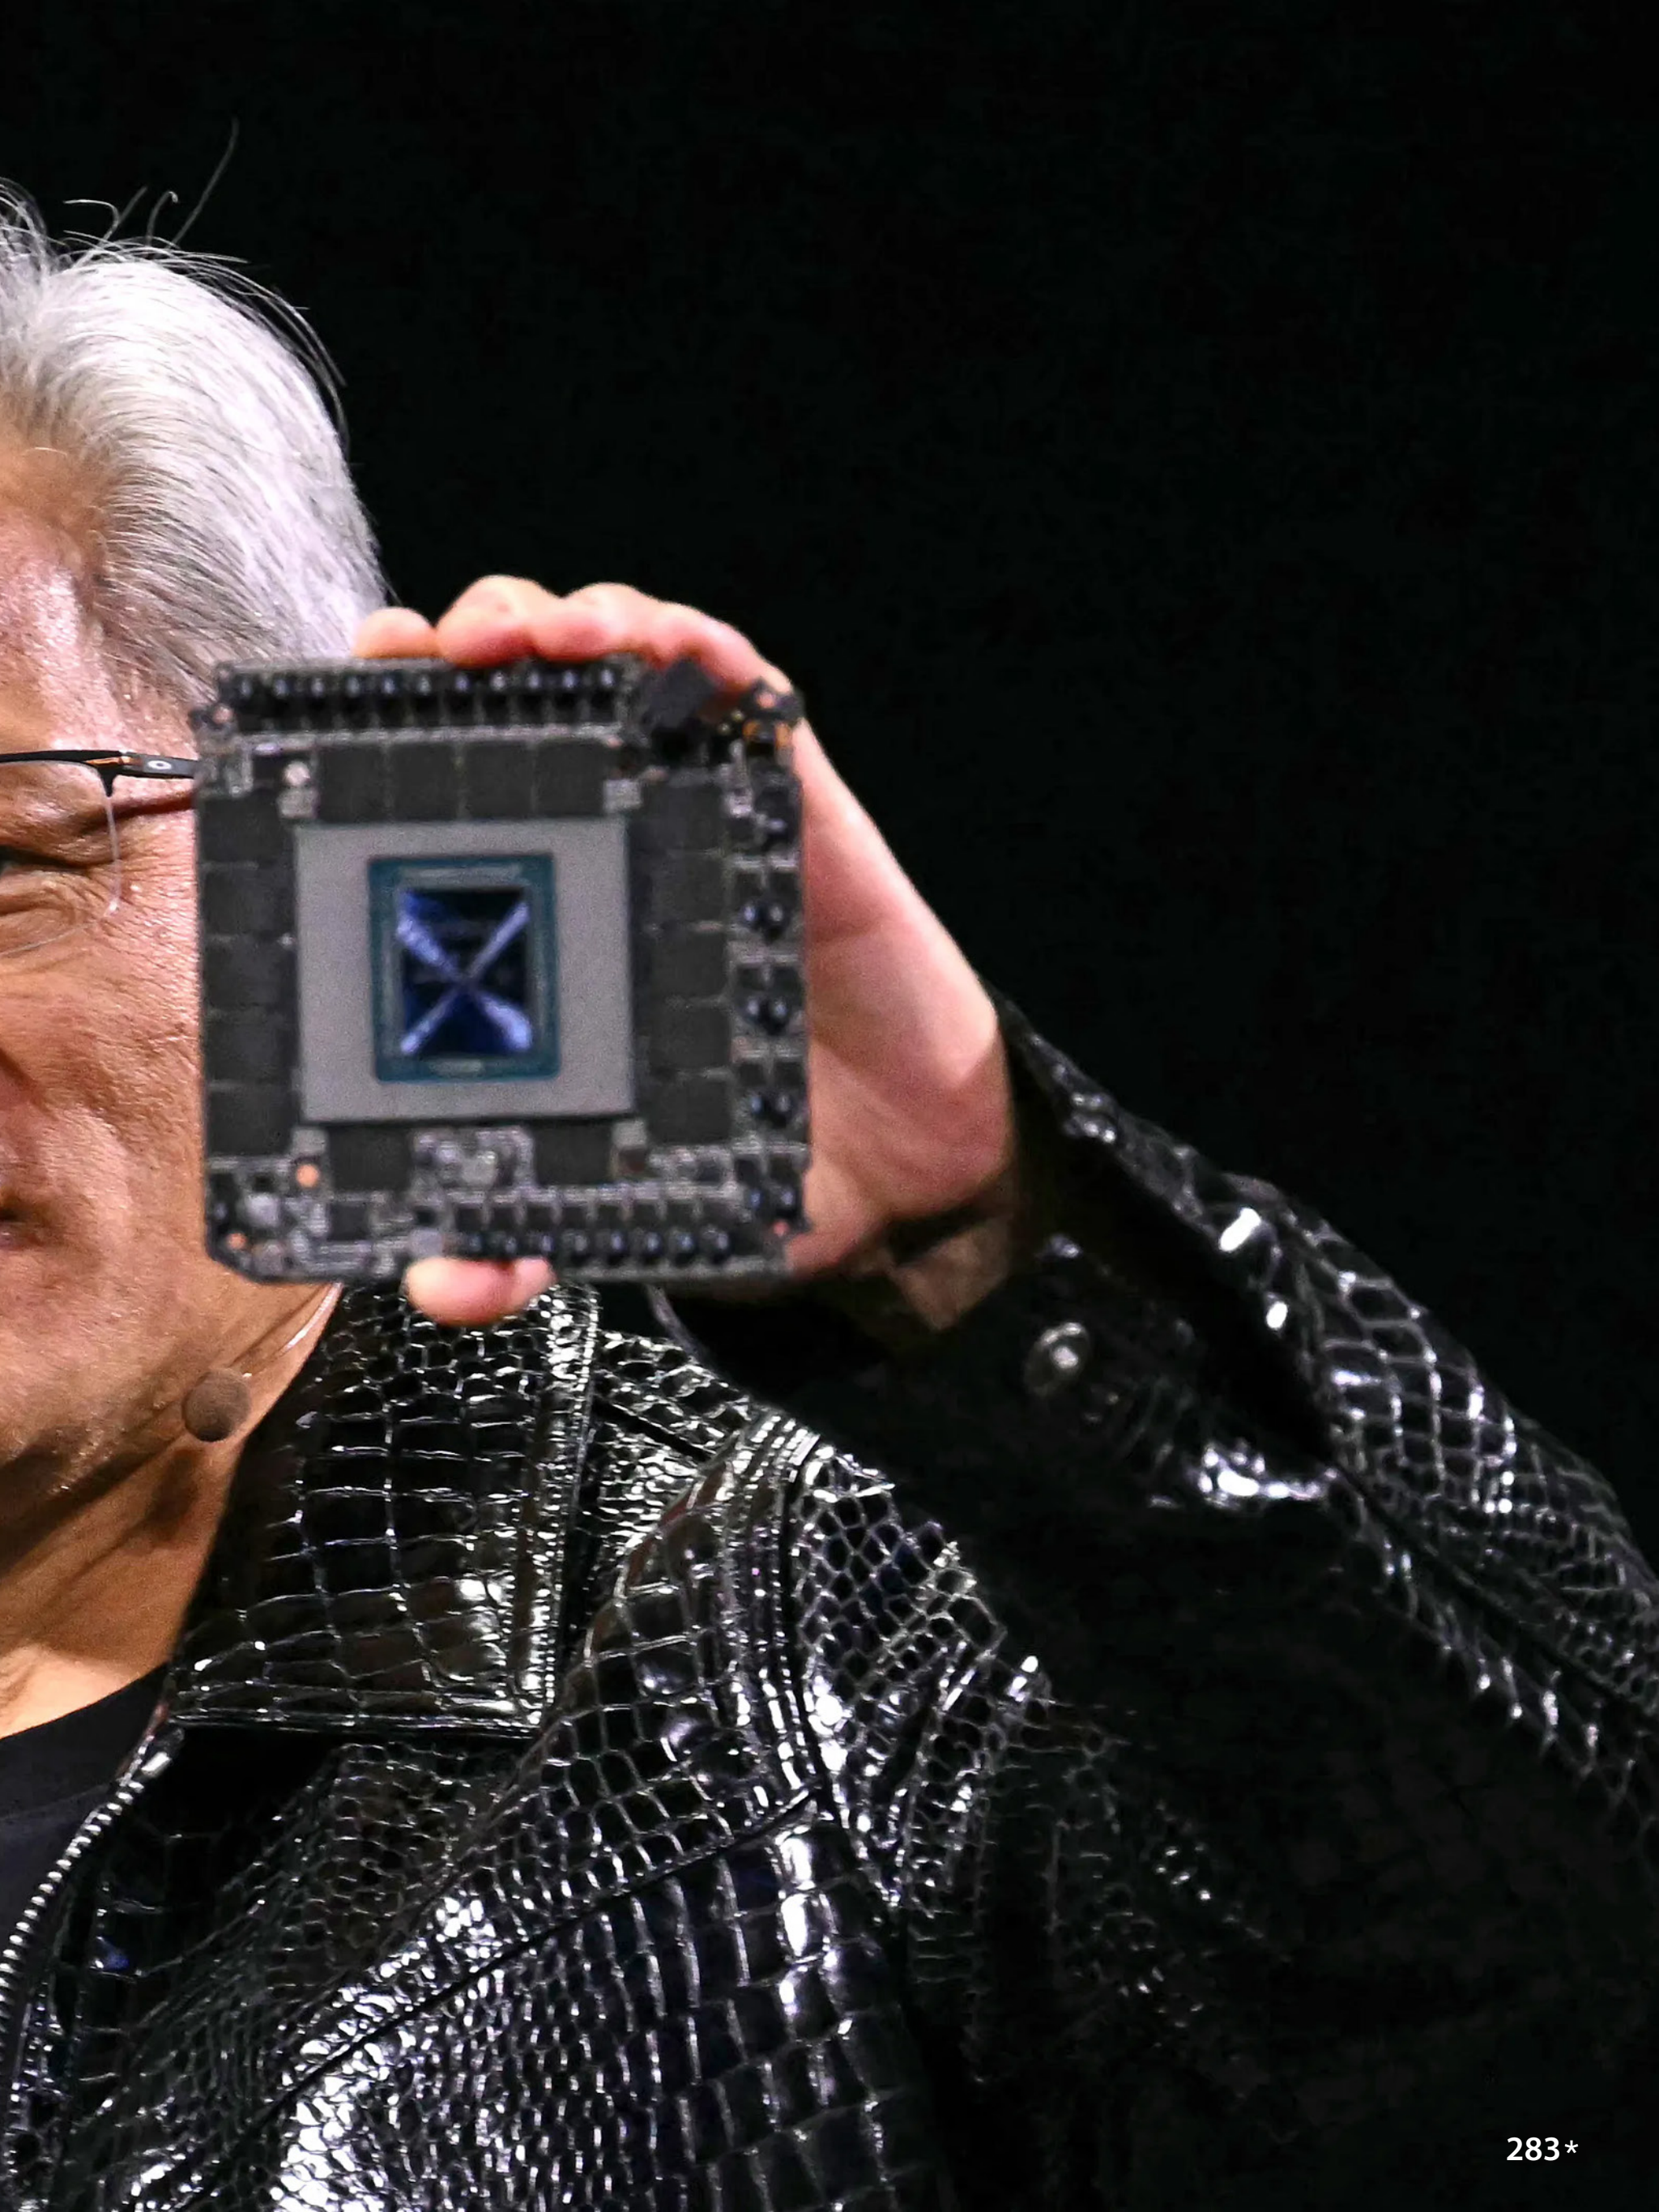

YOUR BRAND HAS NEVER BEEN IN SUCH GOOD COMPANY

Everything Apple, Every Week.

MASTHEAD

TECHLIFE NEWS INTERNATIONAL
TechLifeNewsMagazine.com

SOCIAL NETWORKS

X.com
@Techlife_News
x.com/#!/Techlife_News

Facebook
facebook.com/TechlifeNews

SUBSCRIPTIONS

TechLife News Website
techlifenewsmagazine.com

iTunes App Store™ / Newsstand
itunes.apple.com/app/techlife-news/id529054466?mt=8

CONTACTS

Executive Director - ceo@techlifenewsmagazine.com
Relationship Management - crm@techlifenewsmagazine.com
Editor - editor@techlifenewsmagazine.com
Advertise - ads@techlifenewsmagazine.com
Social Network - network@techlifenewsmagazine.com
Promotions - promo@techlifenewsmagazine.com

EXECUTIVE DIRECTOR

Ivan Castilho
Chief Executive Officer / Head of Design Conception

MINDFIELD DIGITAL ART & GRAPHIC DESIGN

Milene Mencari
Graphic Production Coordinator

Saulo Cardoso
Graphic Design Producer

Natielly Dias
Graphic Design Producer

Claudio Magalhaes
Graphic Design Producer

NEWS WEBSITE

Jack Crute
Newsroom Production

MAGAZINE EDITORIAL CONTENT

Benjamin Kerry
Chief Editor

WRITERS

Precise English Inc.
Gavin Lenaghan
Craig Lenaghan
Shanon O'Dovovan

COLUMNS

DIGITAL ENTERTAINMENT REVIEW
Gavin Lenaghan

FINANCE NEWS
Associated Press / Bloomberg

REVISION

Gavin Lenaghan
Jack Crute

PRODUCTION SUPPORT / COLLABORATORS

Rui da Costa
Segolene Vincent
Fabian Groenhof
Lise Berda
James Jarvis
Richard Sawyer
Lauren Brown
Matthew Coburn
Jeffrey Milks
Susan Kiesling
Raquel Serrano
Robert Fluellen
Lisa Swiniarski
Roger Gayalkar
Icy Ventura
Pilar Bescos
Candice Liu
Sharnae Hughes
Saikumar B
Ryan Szloch
Ryan Szloch
Ciera Chang
Trevor Carroll
Oscar Garcia
Jenny Marry
Sierra Simon
Sharon Ferraro
Devesh Kamal

PRIVACY POLICY

techlifenewsmagazine.com/privacy-policy/

TERMS OF USE

techlifenewsmagazine.com/terms-of-use/

TechLife NEWS

TechlifeNews.com