THE HANDELOT TIMES

EDITION LXI

THE BEST B2B MAGAZINE

MAY 2026 #61

IFA 4-8 SEPT 2026

Tech Under Pressure: Tariffs, Rules, and the Race to Adapt

IFA 2026
Meet us there

The EU AI Act Countdown: What Businesses Must Do Before August 2026 18-21

The Great Reshuffling 26-29

Tariffs vs. Tech - What the Supreme Court Ruling Means for Supply Chains . . . 73-75

The Handelot Times

The Online B2B Magazine for Electronics Wholesalers

The tech industry is changing quickly, and we want to keep you up to date with the latest news. We are developing by opening up to new innovations in IT. We are growing and remaining a reliable source of B2B information from all over the world. With The Handelot Times, we will bring you more information about market trends, share cutting-edge ideas and bring our world together with yours.

HANDELOT.COM

3

Lines of business:

- Platform
- Magazine
- Events

15

Years

on the market

4

Segments

- Mobility
- Inhome
- IT
- Energy

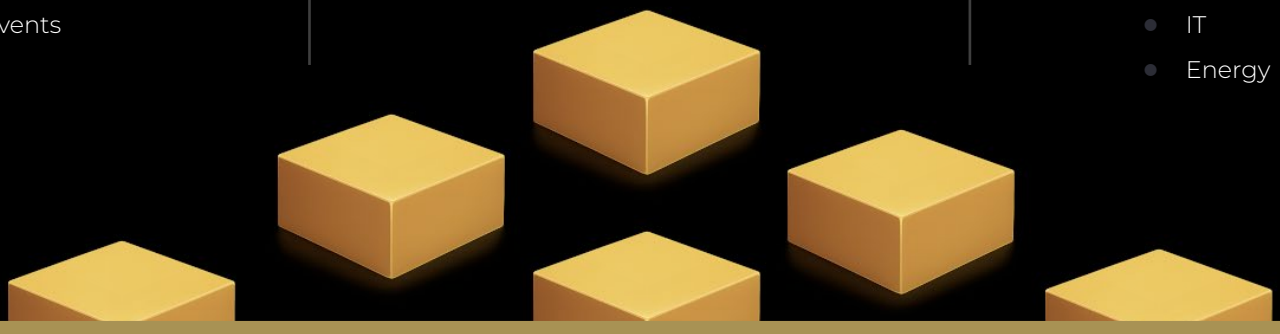

Advertising:

For quotes please email **times@handelot.com** or call us at **+48 519 774 836**

Design & DTP:
Adam Mieloszynski

Main Copywriter:
Melania Romanelli

Copywriter:
Vasudevan, K.S.

elite members

ALL IN SRL

BAD GAME

BIYON

ROLV GROUP

TELZEN

WS WOLFSOHN

elite members

14-15: All in

16-17: Bad Games

22-23: Biyon

24-25: Rolv Group

30-31: Telzen

32-33. Wolfsohn

ADVERTISERS INDEX

REGULAR PLUS

38-39: APL Global Logistics BV

40-41: Captains Freight Services HK LTD

46-47: clicktel GmbH

48-49: Hitisy GmbH

56-57: Millbank FX Limited

58-59: Patron Logistics USA LLC

64-65: Smart Telecomm LLC

66-67: Tiger-Team Company Sp. z o.o. Sp. K.

Article Index

2: Intro

3: Elite Members

6-8: Quick Hits

10-13: Shortage in new phones and why prices go up

18-21: The EU AI Act Countdown: What Businesses Must Do Before August 2026

26-29: The Great Reshuffling

34-37: The Anthropic Effect: How Claude Is Challenging ChatGPT's Dominance

42-45: The Rise of Agentic AI: When Your Software Starts Making Decisions

52-55: Cybersecurity in the AI Era: New Threats, New Defenses

60-63: Cloud 3.0 - Hybrid, Multi-Cloud, and Sovereign Architectures

68-70: Smartphone AI Wars - On-Device Intelligence Changes Everything

73-75: Tariffs vs. Tech - What the Supreme Court Ruling Means for Supply Chains

79-81: Foldables Go Mainstream - The iPhone Fold Effect

84-86: Oil inflation and the world's economy

89-91: TurboQuant and why it dropped memory prices

94-96: The State of Gaming 2026: Layoffs, AI, and the UGC Revolution

ADVERTISERS INDEX

REGULAR

71: AB GSMshop.at GmbH

72: ADESIFLAVA PRIVATE LIMITED

76: FELIX TRADE P.C.

77: Hasiatech

78: Lockwood Fulfilment Ltd

82: Logiverse FZCO

83: MEDIAMARKET INTERNATIONAL FZCO

87: Prestige Group Srl

88: Raduga Pte Ltd

92: Rainbow International Distribution SRL

93: UPTECH USA CORP

97: Vitel Mobile

98: Wingocean Intl Logistics Co

Quick Hits

Tech News Roundup

Apple's DRAM Power Play: Paying Double to Starve the Competition

Apple is reportedly buying up virtually all available mobile DRAM on the global market at prices roughly double the going rate, in what industry analysts are calling the most aggressive supply chain maneuver in the company's history. According to South Korean semiconductor sources, Apple paid Samsung a 100 percent premium for LPDDR5X RAM modules, ensuring priority access to memory chips that are in critically short supply worldwide.

The strategy is not merely about securing components for the upcoming iPhone 18 and foldable iPhone. It is a deliberate competitive weapon. With Apple absorbing the available supply, rival chipmakers MediaTek and Qualcomm have reportedly been forced to cut production schedules for 4-nanometer processors, reducing output by an estimated 15 to 20 million units. Goldman Sachs projects that the DRAM supply gap will reach 4.9 percent in 2026, with the shortage expected to persist until at least the second half of 2027.

The financial logic is striking. Apple is willing to absorb short-term profit margin losses on memory procurement in order to maintain stable device pricing — including the $599 MacBook Neo — while forcing Android manufacturers to either raise prices or reduce production. It is a move that only a company with Apple's cash reserves and vertical integration can execute at this scale, and it underscores how supply chain dominance has become as important as product innovation in the smartphone wars.

Apple Enters the Smart Glasses Race: Four Prototypes, One Big Bet

Apple is testing four different smart glasses designs as it prepares to challenge Meta's dominant Ray-Ban smart glasses, according to Bloomberg's Mark Gurman. The prototypes, internally code-named N50, include a large rectangular frame, a slimmer rectangular model similar to those worn by CEO Tim Cook, a larger oval or circular frame, and a smaller oval design. Apple could launch with some or all of these styles, positioning the product as a fashion-forward wearable rather than a single-form gadget.

Unlike the Vision Pro, which offered full mixed reality with a $3,499 price tag and limited consumer adoption, the smart glasses will take a lighter approach. The devices will not include displays. Instead, they will focus on camera functionality for photos and videos, phone call integration, music playback, notifications, and Siri interaction, all synced deeply with the iPhone ecosystem. Apple is using acetate frames — the same durable, premium material used in high-end eyewear — to differentiate from Meta's plastic construction.

The smart glasses are expected to launch in late 2026 or early 2027. The timing is significant: Meta's Ray-Ban smart glasses have sold millions of units and established the category in consumer consciousness. Apple's entry could expand the market dramatically, though some within the company are reportedly concerned that Siri's AI capabilities may still lag behind Meta's AI assistant, potentially undermining the product's core value proposition.

AI Bot Traffic Surges 300 Percent, and Publishers Are Paying the Price

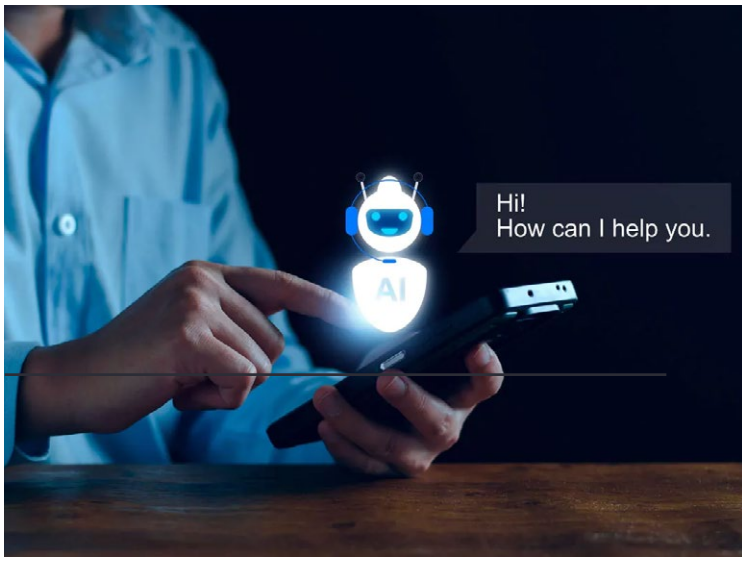

Global AI bot activity targeting the publishing industry surged by 300 percent in 2025, according to a new report from Akamai, one of the world's largest content delivery networks. The media industry now ranks second globally in AI bot traffic, absorbing 13 percent of all AI-driven automated requests. Within that traffic, AI training crawlers — bots that scrape content to feed large language models — accounted for 63 percent of all AI bot activity targeting media companies.

The economic impact on publishers is severe. AI chatbots drove approximately 96 percent less referral traffic than traditional Google search, meaning that as AI systems absorb and redistribute publisher content, they send almost no readers back to the original source. Meanwhile, the scraping bots themselves consume server and CDN resources, increasing infrastructure costs for publishers who receive no revenue in return. Pageviews decline, costs rise, and brand visibility weakens — a triple threat to an industry already under financial pressure.

OpenAI generated the highest volume of AI bot traffic targeting media companies, with publishing organizations accounting for 40 percent of all OpenAI scraping requests. The report raises urgent questions about the sustainability of the current model, in which AI companies harvest publisher content at industrial scale to train models that then compete with the publishers for audience attention. Without a revenue-sharing framework or meaningful licensing agreements, the relationship between AI platforms and content creators is becoming increasingly adversarial.

MWC Barcelona 2026: Robots, Foldables, and the Road to 6G

Mobile World Congress 2026 in Barcelona delivered a show dominated by three themes: AI everywhere, foldable innovation, and early-stage 6G development. The event's most eye-catching moment came from Honor, which showcased what it called the industry's first "Robot Phone" — a smartphone with agentic AI capabilities and a robotic camera that tracks moving subjects and responds to voice commands with physical gestures. Honor also announced the Magic V6, now the world's thinnest book-style foldable smartphone.

Robots were everywhere on the show floor. China Mobile operated a robot restaurant serving tea and snacks. AGIBOT's X2 humanoid robot performed tai chi, served champagne, and danced for attendees. Lenovo unveiled the Legion Go Fold, a concept foldable gaming handheld with a flexible display that expands from 7.7 inches to 11.6 inches. Vivo showcased a professional photography kit built around a 400mm equivalent telephoto extender for its upcoming X300 Ultra.

On the infrastructure side, the GSMA announced the Open Telco AI initiative and the OCUDU Ecosystem Foundation, a public-private partnership developing open-source software for 5G and 6G Radio Access Networks. The GSMA Foundry and the European Space Agency committed up to €100 million for projects converging space and mobile industries. AWS announced it is expanding its infrastructure investment in Spain to $33 billion, up $18 billion from its previous commitment. And in a moment of pure spectacle, Elmo set a world record for the fastest remote driving speed at 157 kilometers per hour on the F1 Circuit de Barcelona-Catalunya.

Dynamic Pricing Hits the Supermarket Aisle

The paper price tags that have defined grocery shopping for decades are disappearing. Walmart plans to install electronic shelf labels in all 4,600 US stores by December 2026, with roughly 2,300 locations already live. In the UK, Co-op is rolling out across 2,400 stores, Sainsbury's is trialing in three locations, and Asda, Waitrose, and Lidl are all exploring upgrades. Across Europe, around 30 percent of large retailers already use the technology. The screens can update prices up to six times per minute, and while retailers insist ESLs are about efficiency — cutting labor by 75 percent, reducing waste, syncing online and in-store prices — the Bank of England has warned they could enable surge pricing on everyday groceries. UC San Diego researchers analyzed 180 million price observations and found no evidence of real-time spikes so far, but legislators are not waiting: US Senators have introduced a bill to ban grocery dynamic pricing, and at least 12 states are considering regulation. The technology is coming. The question is whether supermarkets use it to help shoppers or squeeze them. ■

OWN YOUR SPOT.

Reseller Park | IFA 4–8 SEP 2026

Europe's leading B2B marketplace for consumer tech distributors

Lounge A
HANDELOT Area
Reseller Park –
HUB27 at
IFA 2026

THIS IS YOUR STAGE.

200 exhibitors | 15,000 B2B visitors

NO BOOTH IN THE HANDELOT AREA YET?

FIX THAT NOW.

Limited spaces. Almost sold out.

info@handelot.com
www.handelot.com

a project of GLOBALFAIRS Exhibitions for Future

supported by HANDELOT B2B wholesale consultants

Shortage in New Phones

and Why Prices Go Up

Something unusual is happening in the smartphone market in 2026. Prices are going up, shipments are going down, and the phone you bought two years ago might actually be harder to replace with something better at the same price. The Samsung Galaxy S26 now starts at $899, a $100 increase over its predecessor. The Galaxy Z Fold 7 has climbed to $2,500 at the top end. The foldable iPhone is expected to debut above $2,000. And across the Android ecosystem, manufacturers are quietly raising retail prices by 10 to 20 percent — or, in some cases, downgrading specs to hold the line.

This is not the normal rhythm of the smartphone upgrade cycle, where each generation brings slightly more performance at roughly the same price. This is a market under structural stress. Global smartphone shipments are forecast to decline between 2 and 12 percent in 2026, depending on the analyst, with Counterpoint Research warning of the sharpest decline on record. Gartner estimates that combined DRAM and SSD prices will surge by 130 percent by the end of the year. Display panel shipments for smartphones are projected to fall 7.3 percent.

The explanation involves three forces colliding simultaneously: a global memory chip shortage driven by the AI boom, tariff pressures that add cost at every stage of the supply chain, and a component market that has been structurally redirected away from consumer devices and toward data centers. The phone in your pocket is getting more expensive because the world decided that training AI models matters more than making affordable smartphones.

THE MEMORY CRISIS: WHEN AI EATS THE SUPPLY CHAIN

The single largest factor behind rising smartphone prices is the global shortage of memory chips — specifically DRAM and NAND flash storage. The shortage is not caused by a natural disaster or a factory fire. It is caused by artificial intelligence. The hyperscale cloud companies — Microsoft, Google, Meta, Amazon — have dramatically increased their purchases of High Bandwidth Memory, the specialized DRAM used in AI accelerators and data center GPUs. To produce a single gigabyte of HBM, manufacturers must consume roughly three times the raw silicon wafer capacity required for the standard mobile-grade DRAM that goes into smartphones.

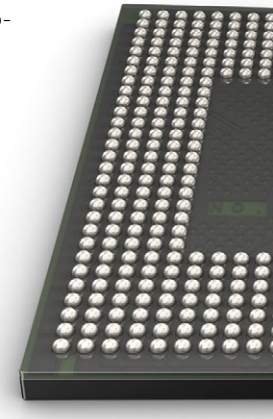

The world's three largest memory manufacturers — Samsung Electronics, SK Hynix, and Micron Technology — have responded rationally to market incentives. HBM commands significantly higher margins than

mobile DRAM. Faced with limited cleanroom space and capital expenditure budgets, all three companies have pivoted production toward enterprise-grade memory for AI, reducing the supply available for smartphones, laptops, and consumer electronics.

The cost impact is already visible. DRAM price surges have increased smartphone bill-of-materials costs by 25 percent for low-end devices, 15 percent for mid-range, and 10 percent for flagships. Mobile DRAM costs rose 50 percent quarter-on-quarter in early 2026, while NAND flash storage surged over 90 percent. Goldman Sachs projects a DRAM supply gap of 4.9 percent for 2026, with the shortage expected to persist until at least the second half of 2027. Samsung Electronics co-CEO TM Roh stated at CES 2026 that "no company is immune" to the memory crisis, calling price increases "inevitable."

THE APPLE FACTOR: CORNERING THE MARKET

Apple has responded to the memory shortage not by absorbing higher costs passively but by weaponizing its financial power. Reports from South Korean semiconductor sources indicate that Apple has been purchasing virtually all available mobile DRAM on the global market, paying Samsung a 100 percent premium for LPDDR5X modules. The strategy is deliberate: by securing supply at any price, Apple can maintain stable device pricing while forcing competitors to either raise prices or cut production.

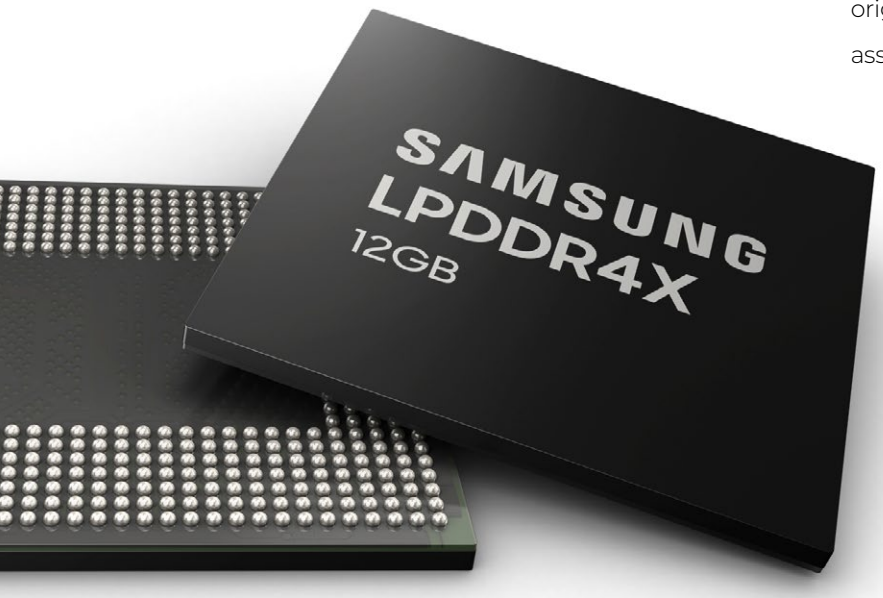

The impact on the Android ecosystem has been immediate. MediaTek and Qualcomm have reportedly cut production schedules for 4-nanometer processors, reducing output by an estimated 15 to 20 million chip units. Android manufacturers, unable to secure memory components at competitive prices, face an impossible choice: absorb the cost increase and watch margins evaporate, pass the increase to consumers and risk losing sales, or downgrade specifications to keep prices flat. Many are choosing the third option. TrendForce reports that several smartphone brands have begun "spec-sheet engineering" — reducing RAM allocations, using older display panels, and downgrading camera modules to offset memory costs. The result is a market where the $400 phone of 2026 may actually be less capable than the $400 phone of 2025 in terms of raw specifications.

TARIFFS: THE SECOND LAYER OF COST PRESSURE

The memory shortage alone would be enough to push prices upward. But it is compounded by the tariff environment. The Section 122 surcharge of 10 percent on all imports, imposed in February 2026 after the Supreme Court struck down IEEPA-based tariffs, applies to finished devices and many of their components. Section 301 tariffs on Chinese goods remain in force at elevated rates. The combined effect is a 15 percent or higher additional cost layer on imported smartphones and components.

For manufacturers that assemble devices in China — which includes the majority of global smartphone production — the tariff impact is direct. For those assembling in Vietnam, India, or other diversified locations, the impact is indirect but still significant: many sub-components still originate in China and face tariffs before reaching the assembly line. The Consumer Technology Association estimates that tariff-related price pressures could reduce American consumers' purchasing power by $123 billion, with smartphones and laptops among the hardest-hit categories.

WHAT CONSUMERS CAN EXPECT

The practical implications for consumers are straightforward and uncomfortable. Flagship smartphones will cost more. Gartner estimates an average price in-

crease of 13 percent across the smartphone market in 2026 compared to 2025 levels. At the flagship tier, the increases are steeper: Counterpoint Research puts the increase at $150 to $200 per device for premium models. Mid-range devices will be affected disproportionately. Memory can represent 15 to 20 percent of the total bill of materials for a mid-range phone. A 50 percent increase in DRAM costs translates directly into higher retail prices or reduced specifications at the same price point. Budget devices face the most severe pressure, with 25 percent BOM increases that cannot be easily absorbed by manufacturers operating on thin margins.

The upgrade cycle will lengthen. When new phones cost more and offer fewer dramatic improvements, consumers hold onto existing devices longer. This was already a trend — the average smartphone replacement cycle had stretched to over three years before 2026 — and the current price environment will accelerate it. For manufacturers, this means selling fewer units at higher prices, a dynamic that sustains revenue in the short term but constrains market growth.

THE DISPLAY PANEL SQUEEZE

Memory is not the only component under pressure. TrendForce projects that global smartphone display panel shipments will fall 7.3 percent in 2026, dropping from 2.31 billion units to 2.14 billion — the first annual decline in several years and the end of a growth cycle that began in 2023. As smartphone brands order fewer units in response to weakening demand and higher costs, display manufacturers face their own margin pressure, creating a feedback loop: fewer orders mean less production scale, which means higher per-unit costs, which means higher prices for the brands that are still buying.

OLED panels, which now dominate the flagship and mid-range segments, are particularly affected. The same manufacturing capacity that produces smartphone OLEDs also serves the growing market for foldable devices, automotive displays, and AR/VR headsets. As these alternative markets expand, the share of OLED production allocated to standard smartphones contracts. For consumers, this translates into fewer display innovations trickling down to affordable devices and longer waits before premium display technology becomes standard at the mid-range price point.

THE NEW NORMAL

The smartphone price increases of 2026 are not a temporary blip caused by a single supply chain disruption. They reflect a permanent reallocation of global semiconductor capacity toward AI infrastructure and away from consumer devices. As long as the AI boom continues — and every indicator suggests it will — the competition for memory, processing power, and advanced packaging between data centers and smartphones will persist.

The era of the steadily improving, steadily priced smartphone is giving way to something different: a market where the best technology goes to the highest bidder, and the highest bidders are no longer the people buying phones. They are the companies training the models that run on them. For consumers, the message is simple but unwelcome: the next phone you buy will likely cost more, do marginally more, and last longer in your pocket — not because you want it to, but because replacing it will be harder to justify than ever. ■

SOURCES

- GARTNER FEB 2026 — 130% DRAM+SSD PRICE SURGE, 13% SMARTPHONE PRICE INCREASE, SHIPMENT DECLINE
- COUNTERPOINT RESEARCH — SHIPMENT FORECASTS REVISED DOWN, $150-$200 FLAGSHIP PRICE INCREASE, SHARPEST DECLINE ON RECORD
- IDC — GLOBAL MEMORY SHORTAGE CRISIS ANALYSIS, HBM CONSUMING 3X WAFER CAPACITY VS MOBILE DRAM
- CNBC — SMARTPHONE PRICES TO RISE DUE TO AI-FUELED CHIP SHORTAGE, TM ROH "NO COMPANY IS IMMUNE"
- TRENDFORCE — MEMORY PRICE SURGE Q1 2026, SPEC-SHEET ENGINEERING, DISPLAY PANEL SHIPMENTS -7.3%
- GOLDMAN SACHS — 4.9% DRAM SUPPLY GAP, SHORTAGE UNTIL H2 2027
- WCCFTECH/TECHSPOT — APPLE DRAM STRATEGY, 100% PREMIUM, MEDIATEK/QUALCOMM PRODUCTION CUTS
- INFLUENCER MAGAZINE UK — SAMSUNG GALAXY PRICE HIKES, Z FOLD 7 AT $2,500
- CONSUMER TECHNOLOGY ASSOCIATION — $123B PURCHASING POWER REDUCTION FROM TARIFFS

ALL IN SRL

The future of telephony

- Worldwide Delivery in 24/48 Hours
- Global Wholesale, Lightning-Fast Delivery
- 20 years Market Experience
- Main brands available, best prices
- Relationship matters, contact us for your personal quotation

SAMSUNG Lenovo motorola

HUAWEI XBOX PlayStation

ASUS alcatel NOKIA

uleFone

all-insrl.it

Contact us:

Mirco Sorbo
@ sales@all-insrl.it
+39 329 733 7734

Giustino Magno
@ g.magno@all-insrl.it
+39 333 489 8473

BAD GAME

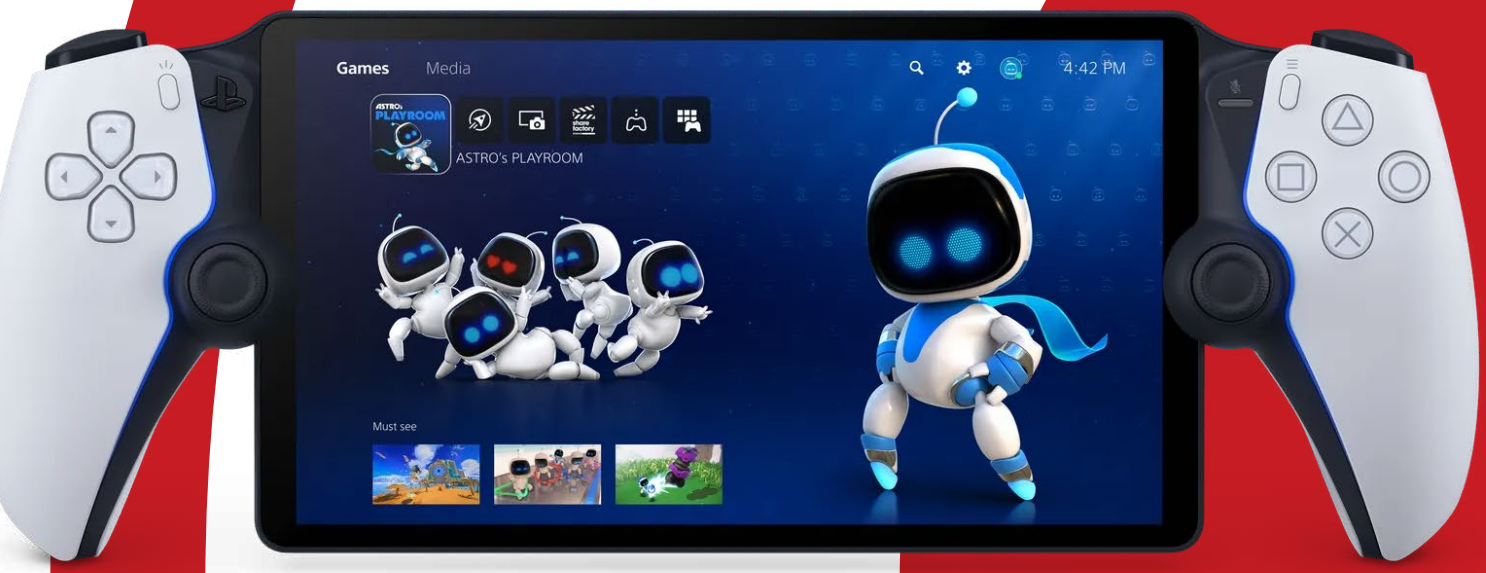

STEAM DECK

STARLINK

dji

Canon

PlayStation

Nintendo

XBOX

JBL

GARMIN

Collectables,
tradables, games
and **more!**

www.badgame.co.uk

Contact:

London
United Kingdom
info@badgame.co.uk
+44 7826 650673

Gaming and technology

BAD GAME

The EU AI Act Countdown:

What Businesses Must Do Before August 2026

The European Union has never been shy about regulating technology. GDPR redefined global data privacy standards. The Digital Services Act reshaped platform accountability. And now, the AI Act is poised to do the same for artificial intelligence. But unlike previous regulations that arrived after years of gradual enforcement, the AI Act's most consequential deadline is approaching with remarkable speed — and many businesses are not ready.

On August 2, 2026, the full compliance obligations for high-risk AI systems take effect across all 27 EU member states. Companies that develop, deploy, import, or distribute AI systems classified as high-risk must demonstrate conformity with a detailed set of requirements covering risk management, data governance, technical documentation, human oversight, accuracy, robustness, and cybersecurity. The penalties for failure are severe: fines of up to €35 million or 7 percent of global annual turnover, whichever is higher.

This is not a theoretical deadline. Prohibited AI practices have already been enforceable since February 2025. The August 2026 milestone activates the bulk of the regulation's operational teeth. And while a proposed Digital Omnibus package could postpone certain obligations, the European Parliament has not formally adopted it, leaving companies in a position where assuming delay is a gamble they cannot afford to lose.

Here is what the AI Act demands, who it affects, and what practical steps businesses should be taking right now.

WHAT EXACTLY IS A HIGH-RISK AI SYSTEM?

The AI Act does not regulate all artificial intelligence equally. It establishes a risk-based classification framework with four tiers: unacceptable risk (banned outright), high risk (heavily regulated), limited risk (transparency obligations), and minimal risk (largely unregulated). The August 2026 deadline concerns the high-risk category, which is defined in Annex III of the regulation.

The scope is broader than many companies initially assumed. High-risk AI systems include those used in biometric identification and categorization, critical infrastructure management, education and vocational training, employment and workforce management, access to essential public and private services, law enforcement, migration and border control, and the administration of justice.

In practical business terms, this means that if your company uses AI to screen job applicants, score candidates, allocate work based on behavioral patterns, monitor employee performance, evaluate creditworthiness, determine insurance pricing, or triage access to public services, your AI systems are almost certainly classified as high-risk under the Act.

The classification is functional, not sectoral. It does not matter what industry you operate in. What matters is what the AI system does and the impact it has on in-

dividuals. A recruitment algorithm used by a technology startup faces the same obligations as one used by a multinational bank.

THE COMPLIANCE REQUIREMENTS: WHAT MUST BE IN PLACE

The AI Act's obligations for high-risk systems are comprehensive and technically demanding. They are not designed to be satisfied with a compliance checklist or a policy document. They require operational changes embedded into how AI systems are built, tested, deployed, and monitored.

The core requirements include:

- **Risk management system.** Providers must establish and maintain a continuous risk management process that identifies, analyzes, evaluates, and mitigates risks throughout the AI system's lifecycle. This is not a one-time assessment. It must be documented, updated, and demonstrably linked to design and deployment decisions.
- **Data governance.** Training, validation, and testing datasets must be relevant, sufficiently representative, and as free of errors as possible. Companies must document data collection processes, address potential biases, and ensure datasets reflect the specific characteristics of the deployment environment. For AI systems used in hiring, for example, this means proving that training data does not systematically disadvantage protected groups.
- **Technical documentation.** Providers must prepare detailed documentation demonstrating compliance, sufficient for regulatory authorities to assess conformity. This includes descriptions of the system's intended purpose, design specifications, training methodologies, performance metrics, and known limitations.
- **Human oversight mechanisms.** High-risk AI systems must be designed to allow meaningful human oversight. This means the system must provide deployers with the tools and information necessary to understand outputs, intervene when needed, and override automated decisions.
- **Accuracy, robustness, and cybersecurity.** Systems must meet appropriate levels of accuracy for their intended purpose, remain robust against errors and adversarial inputs, and incorporate cybersecurity protections proportionate to their risk profile.
- **Conformity assessment and registration.** Before placing a high-risk AI system on the EU market, providers must complete a conformity assessment, affix a CE marking, and register the system in the EU database. For certain categories, particularly biometric systems, this assessment must be conducted by an independent notified body.

WHO IS AFFECTED — AND WHERE

One of the AI Act's most significant features is its extraterritorial reach. The regulation applies not only to companies headquartered in the EU but to any organization that places AI systems on the EU market or whose AI systems produce outputs used within the EU. A software company based in San Francisco whose recruitment AI is used by a client in Berlin is subject to the Act. A cloud provider headquartered in Singapore that offers AI-powered risk scoring to a European insurance firm falls within scope.

This extraterritorial dimension mirrors GDPR's approach and is likely to produce a similar effect: a de facto global standard. Companies that serve international markets will find it more efficient to build AI systems that comply with the EU framework by default rather than maintaining separate compliant and non-compliant versions. The AI Act, like GDPR before it, is poised to export European regulatory preferences worldwide.

Enforcement is decentralized. Each EU member state is responsible for designating market surveillance authorities with full investigatory power, including the authority to request documentation, conduct evaluations, and demand corrective measures. At the European level, the AI Office can request source code access for general-purpose AI models and impose its own corrective actions.

THE OMNIBUS QUESTION: WILL THE DEADLINE MOVE?

In November 2025, the European Commission proposed the Digital Omnibus package, which included provisions to postpone certain high-risk obligations.

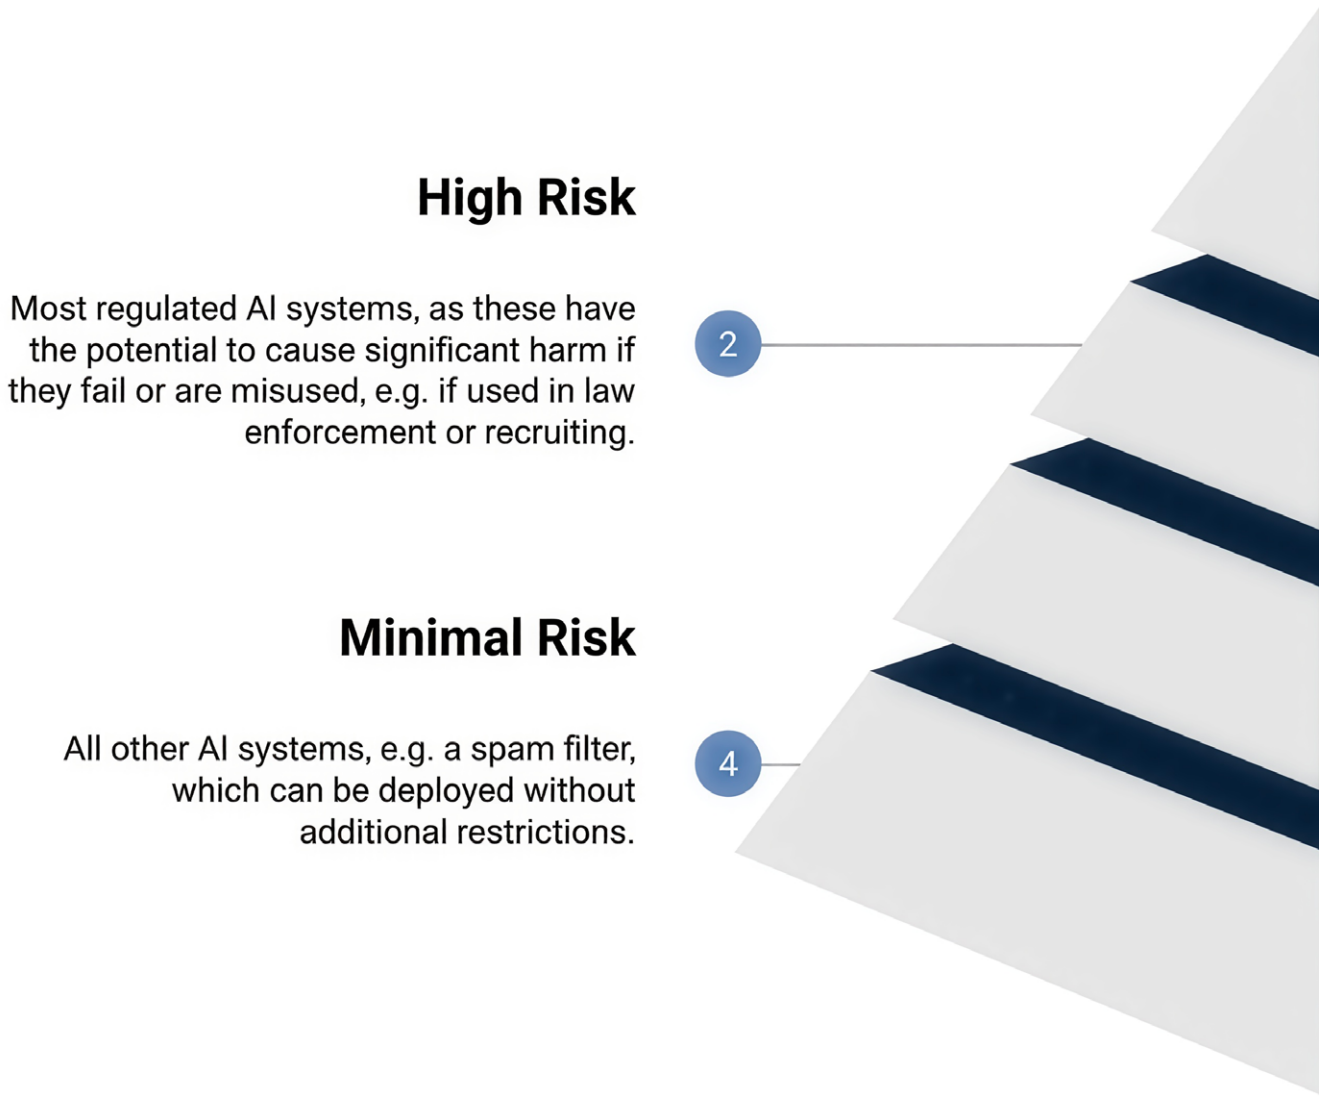

Under this proposal, Annex III high-risk requirements could be deferred until December 2027, and Annex I requirements until August 2028. The postponement would be conditional — obligations would only take effect six to twelve months after the Commission confirms that adequate compliance support, such as harmonized standards and guidelines, is available.

In March 2026, the European Parliament signaled support for the postponement. But the proposal has not been formally adopted, and the legislative process involves both Parliament and the Council. Until the Omnibus is enacted into law, August 2, 2026 remains the legally binding deadline.

The strategic implication is clear: companies should plan as if the deadline holds while monitoring the Omnibus process closely. Those that use the potential delay as an excuse to defer compliance work will face either a scramble if the Omnibus stalls, or a compressed timeline even if it passes — because the obligations do not disappear, they merely shift.

SIX STEPS TO TAKE BEFORE AUGUST

For organizations that have not yet begun substantive compliance work, the window is narrow but not closed. The following actions represent the minimum viable preparation for August 2026.

- Map your AI systems. Conduct a comprehensive inventory of every AI system your organization develops, deploys, imports, or distributes within the EU. Classify each system against the Act's risk categories.
- Clarify your regulatory role. The Act assigns different obligations to providers, deployers, importers, and distributors. Determine which role your organization occupies for each AI system.
- Implement risk management and data governance. Establish documented processes for ongoing risk assessment and ensure your training data meets the Act's representativeness and bias requirements.
- Prepare technical documentation. Draft the technical documentation required for conformity assess-

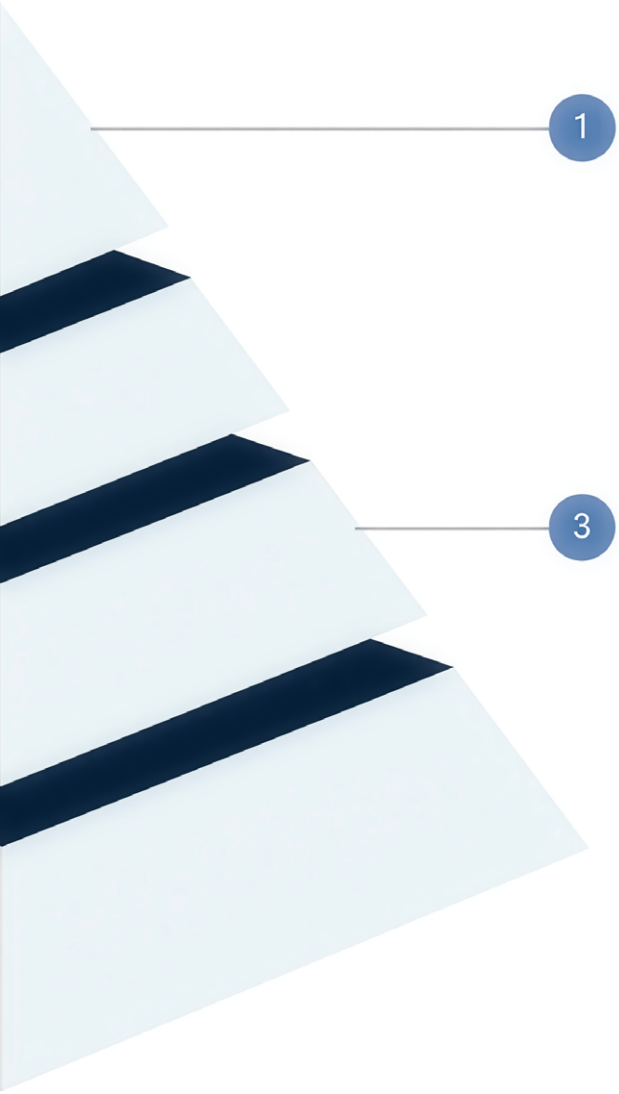

ment. This is often the most time-consuming step and cannot be completed retrospectively.

- Design for human oversight. Ensure your AI systems provide deployers with adequate tools for monitoring, interpreting, and overriding outputs.
- Complete conformity assessment. For systems requiring independent assessment, engage a notified body early. Availability may be constrained as the deadline approaches.

THE COST OF WAITING

The EU AI Act is the most ambitious piece of AI legislation ever enacted. Its requirements are detailed, its penalties are substantial, and its reach extends far beyond European borders. Companies that have been treating August 2026 as a distant milestone are running out of room.

The regulation is not designed to punish innovation. It is designed to ensure that AI systems affecting people's employment, credit, education, safety, and legal rights meet minimum standards of transparency, fairness, and accountability. Companies that have already invested in responsible AI practices will find compliance achievable. Those that have not will find the next four months uncomfortable.

The countdown is real. The obligations are specific. And the market surveillance authorities are staffing up. Whether the Omnibus postponement passes or not, the direction of travel is unmistakable: the era of unregulated AI deployment in Europe is ending. The only question is whether your organization will be ready when it does. ■

SOURCES
·EU AI ACT - SHAPING EUROPE'S DIGITAL FUTURE
·THE EU AI ACT: 6 STEPS TO TAKE BEFORE 2 AUGUST 2026 - ORRICK
·EU AI ACT 2026 UPDATES: COMPLIANCE REQUIREMENTS - LEGALNODES
·ARTICLE 99: PENALTIES - EU AI ACT
·ANNEX III: HIGH-RISK AI SYSTEMS
·MEPS SUPPORT POSTPONEMENT OF CERTAIN RULES ON AI - EUROPEAN PARLIAMENT
·EU DIGITAL OMNIBUS: ANALYSIS OF KEY CHANGES - IAPP
·PENALTIES OF THE EU AI ACT - HOLISTIC AI
·EU AI ACT PRACTICAL COMPLIANCE GUIDE - LEGISCOPE

BIYON

With a strong focus on enhancing **B2B relationship**s and optimizing distribution efficiency, we are dedicated to setting new benchmarks in the **global supply chain.** We ensure our products reach businesses across the world quickly and cost-effectively, without compromising on quality

Google SAMSUNG FUJIFILM dji

Nintendo SENNHEISER Marshall Lenovo hp DELL logitech

RAZER NETGEAR STARLINK SONOS Jabra JBL Canon

EPSON ASUS acer msi Nikon SONY

Contact:

Patricia
+61 426 796 544

Kalyn
+61 426 435 394

Mary
+61 494 163 510

www.biyon.com.au

Responsible import

Environmental responsibilities of importers and manufacturers of electrical and electronic equipment

Environmental responsibilities of importers and manufacturers of electrical and electronic equipment.

ECO LIFE CYCLE

Organizacja Odzysku Sprzętu Elektrycznego i Elektronicznego S.A.

WHY SHOULD YOU CARE?

Environmental responsibilities are a legal requirement, but also an ethical duty towards our planet and future generations.

Any activity that may affect the ecosystem should have the appropriate environmental decisions and permits in place to confirm compliance with applicable standards.

DO YOU REALISE HOW MUCH YOU NEED TO DO?

- Organising and financing the collection and treatment of waste electrical and electronic equipment
- Achieving minimum annual collection, recovery, recycling and preparation for reuse targets for waste electrical and electronic equipment
- Signing a contract with a treatment facility to fulfil the above obligations
- Keeping additional records containing information on the weight of equipment placed on the market
- Submitting annual reports on products, packaging and their management
- Conducting public educational campaigns aimed at increasing the level of care for the environment and raising the ecological awareness of the public

For more information...

WHAT CAN YOU DO?

You can entrust your statutory obligations to qualified specialists who are part of the Eco Life Cycle Recovery Organisation team.

Electrical and Electronic Equipment Recovery Organisations (OOSEiE) are specialised, legally authorised entities that take on the implementation of environmental obligations.

The strategic goal of Eco Life Cycle OOSEiE is to provide real support to Partners in the development of their activities so that they can operate safely, consciously and in accordance with environmental regulations.

Contact:

Eco Life Cycle

Wrzosowa 35, 98-100 Łask

Mariusz Wyrwa

DYREKTOR OPERACYJNY

mariusz.wyrwa@ecolifecycle.eu

+48 663 030 009

ECO LIFE CYCLE OFFERS:

- Zero environmental audit – risk identification, recommendations
- Fulfilment of environmental obligations – collection, processing, educational campaigns, reporting in BDO
- Additional support – training, ISO 14001, ESG reporting, KOBIZE reports

BENEFITS FOR YOUR COMPANY:

- Legal security
- No financial penalties
- Time and cost savings
- Image of a company that cares about the planet

ARTYKUŁ SPONSOROWANY

visit www.ecolifecycle.eu

The Great Reshuffling

How Tariffs, AI, and Regulation Are Redrawing the Global Tech Map

For decades, the global technology industry operated under a set of assumptions that felt almost permanent. Supply chains stretched across continents with predictable efficiency. Regulatory frameworks moved slowly, trailing years behind the innovations they attempted to govern. Trade agreements, however imperfect, provided a stable foundation for cross-border commerce. And the competitive landscape, while fierce, followed recognizable patterns. Companies built, shipped, sold, and scaled within a system that rewarded speed and cost optimization above all else.

That system is now fracturing under simultaneous pressure from three directions. First, a seismic shift in trade policy: the United States Supreme Court struck down IEEPA-based tariffs in February 2026, only for the White House to impose replacement levies under Section 122 of the Trade Act of 1974, creating a 10 percent global surcharge on imports with a 150-day expiration clock. Second, the European Union's AI Act reaches its most consequential enforcement milestone in August 2026, when full compliance obligations for high-risk AI systems take effect, carrying fines of up to €35 million or 7 percent of global annual turnover. Third, enterprise

AI adoption has crossed an inflection point: Gartner predicts that 40 percent of enterprise applications will feature task-specific AI agents by the end of 2026, up from less than 5 percent just a year earlier.

Each of these forces alone would be enough to reshape corporate strategy. Together, they are triggering what can only be described as a great reshuffling — a simultaneous reconfiguration of where technology is built, how it is governed, and what intelligence it carries. The rules of engagement are being rewritten in real time, and the companies that understand the interconnected nature of these shifts will be the ones best positioned to navigate them.

In this article, we will examine how tariff disruption, regulatory acceleration, and the rise of agentic AI are converging to redraw the global tech map. We will explore why treating these as isolated trends is a strategic mistake, and what businesses can do to turn structural turbulence into competitive positioning.

THE TARIFF SHOCK: FROM IEEPA TO SECTION 122

The Supreme Court's decision in Learning Resources, Inc. v. Trump landed on February 20, 2026, with a 6-3 majority ruling that the International Emergency Economic Powers Act does not authorize the President to impose tariffs. The majority opinion drew a clear constitutional line: the power to regulate importation is distinct from the power to tax, and the latter belongs exclusively to Congress.

The ruling eliminated the legal basis for the sweeping reciprocal tariffs that had been imposed under IEEPA — tariffs that had reached as high as 145 percent on Chinese goods and applied broadly across dozens of trading partners. But relief was short-lived. Within hours, the White House pivoted to Section 122 of the Trade Act of 1974, imposing a uniform 10 percent surcharge on imports from all countries, effective February 24. Unlike IEEPA, Section 122 is explicitly limited: a maximum rate of 15 percent and a hard duration cap of 150 days unless Congress votes to extend it.

For the tech industry, the implications are layered. The 10 percent blanket tariff still raises input costs across hardware, components, and consumer electronics. But the real uncertainty lies in what comes next. Section 232 tariffs on steel and aluminum remain untouched. Section 301 tariffs on Chinese tech goods — covering semiconductors, computers, and electrical equipment — continue to apply at elevated rates. And the 150-day clock on Section 122 means that by late July 2026, the trade landscape could shift again, depending on whether Congress acts or whether new executive measures emerge.

The net effect is not simplification. It is a fragmented tariff environment where different product categories face different rates under different legal authorities, with different expiration timelines. For companies managing global supply chains, this is not just a cost problem. It is a planning problem.

SUPPLY CHAINS IN MOTION: RESHORING, NEARSHORING, AND THE NEW GEOGRAPHY OF MANUFACTURING

The tariff turbulence has accelerated a trend that was already gathering momentum: the geographic redistribution of technology manufacturing. According to Bain & Company, 81 percent of CEOs and COOs now report plans to bring supply chains closer to their primary markets, up 18 percentage points from the previous year. Almost two-thirds of executives confirm active investment in reshoring, nearshoring, or split-shoring strategies.

Semiconductors sit at the center of this shift. Global semiconductor revenues are projected to surpass the $1 trillion mark in 2026, driven overwhelmingly by AI demand. The US CHIPS Act has catalyzed $52 billion in domestic production incentives, prompting TSMC, Samsung, Intel, and Micron to announce new fabrication facilities on American soil. In Europe, Intel is building in Germany, TSMC has committed to a plant in Dresden, and STMicroelectronics is expanding in France. According to SEMI, 18 new semiconductor factories are scheduled to begin construction in 2025, with most projected to become operational between 2026 and 2027.

But reshoring is not a silver bullet. Building a semiconductor fabrication plant takes three to five years and costs upwards of $20 billion. In the interim, companies are diversifying assembly and testing operations to Vietnam, India, Malaysia, and Mexico — countries that offer lower labor costs and, critically, are not currently subject to the highest tariff rates. The result is not a reversal of globalization but a reconfiguration: supply chains are becoming shorter, more regional, and more politically aware.

THE REGULATORY RECKONING: EUROPE'S AI ACT REACHES FULL FORCE

While trade policy reshapes where technology is made, regulation is redefining how it is deployed. The European Union's AI Act, the world's most comprehensive AI regulation, hits its most critical enforcement date on August 2, 2026, when full compliance obligations for high-risk AI systems take effect.

The scope is substantial. High-risk classifications cover AI systems used in biometrics, critical infrastructure, education, employment, essential public and private services, law enforcement, migration, and the administration of justice. Providers of these systems must demonstrate compliance across risk management, data governance, technical documentation, record-keeping, transparency, human oversight, accuracy, robustness, and cybersecurity. Conformity assessments must be completed, CE markings affixed, and systems registered in the EU database.

The penalties are designed to command attention. Non-compliance with prohibited AI practices can trigger fines of up to €35 million or 7 percent of global annual turnover — whichever is higher. For violations related to high-risk system obligations, fines can reach €15 million or 3 percent of turnover. These figures exceed GDPR penalties and represent the most aggressive financial enforcement mechanism in the history of technology regulation.

For global businesses, the EU AI Act is not merely a European concern. Any company that deploys AI systems within the EU market, regardless of where it is headquartered, must comply. This extraterritorial reach means that American, Asian, and Middle Eastern technology firms with European customers or users are now subject to EU regulatory authority. The practical effect is a de facto global standard, much as GDPR became the baseline for data privacy worldwide.

Adding complexity, the European Commission proposed a Digital Omnibus package in late 2025 that could delay certain high-risk obligations to December 2027. But the proposal remains uncertain, and prudent compliance planning treats August 2026 as the binding deadline. Companies that wait for clarity may find themselves scrambling.

THE RISE OF THE AI AGENT: ENTERPRISE INTELLIGENCE AT AN INFLECTION POINT

If tariffs are redrawing the physical map of technology and regulation is redrawing the legal map, the explosion of enterprise AI is redrawing the operational map. The numbers tell a striking story. Gartner predicts that 40 percent of enterprise applications will include integrated task-specific AI agents by the end of 2026, a staggering increase from less than 5 percent in 2025. According to industry data, 51 percent of enterprises already have AI agents running in production, with another 23 percent actively scaling deployment.

The global AI agents market reached $10.91 billion in 2026 and is projected to hit $50.31 billion by 2030, growing at a 45.8 percent compound annual rate. McKinsey estimates that AI agents could add $2.6 to $4.4 trillion in annual value across customer service, finance, operations, supply chain, cybersecurity, and human resources.

What distinguishes 2026 from previous years of AI enthusiasm is the shift from experimentation to execution. Companies are no longer running isolated pilots. They are deploying AI agents within core business workflows, supported by centralized governance structures that some analysts call "AI studios" — hubs that combine reusable technical components, use-case assessment frameworks, testing sandboxes, deployment protocols, and dedicated personnel.

Yet the acceleration carries risks. Gartner warns that over 40 percent of agentic AI projects will be canceled by the end of 2027 due to escalating costs, unclear value, and weak risk controls. The organizations that succeed will not be those that deploy AI fastest, but those that deploy it most deliberately — with governance structures robust enough to satisfy both internal stakeholders and external regulators.

WHERE THE THREE FORCES CONVERGE

The critical insight for business leaders is that tariffs, regulation, and AI adoption are not parallel trends operating in separate lanes. They converge, and their intersections create both risk and opportunity.

Consider the infrastructure layer. AI at enterprise scale requires massive computational resources, which means data centers, which means hardware, which means

semiconductors — all of which are directly affected by tariff policy. A 10 percent surcharge on imported server components raises the cost of AI deployment. Reshoring chip manufacturing addresses long-term supply security but does nothing for short-term pricing pressure.

Consider the compliance layer. Companies deploying AI agents in European markets must now navigate the EU AI Act's high-risk requirements while simultaneously managing the data sovereignty implications of where their cloud infrastructure is located. The sovereign cloud market is projected to reach $195 billion in 2026, with Europe leading adoption. AWS launched its European Sovereign Cloud in January 2026 through a German-incorporated entity, physically and logically separate from other AWS regions. Microsoft offers the EU Data Boundary. The message is clear: running AI on a US hyperscaler's European servers may no longer be enough. The US CLOUD Act allows American law enforcement to compel access to data stored abroad by US-headquartered companies, creating a direct tension with EU sovereignty requirements.

Consider the strategic layer. The companies that will thrive in this reshuffled landscape are those that can simultaneously optimize for three variables: cost efficiency under volatile trade conditions, regulatory compliance across multiple jurisdictions, and operational intelligence powered by AI systems that are themselves subject to new governance mandates.

THE STRATEGIC IMPERATIVE: DESIGN FOR TURBULENCE

The great reshuffling is not a temporary disruption that will resolve itself into a new equilibrium. It is a structural shift toward permanent complexity. Trade policy will remain contested, with Section 122 tariffs expiring in July and no guarantee of what follows. AI regulation will deepen, not retreat, as more jurisdictions adopt frameworks inspired by the EU model. And AI adoption will continue to accelerate, creating both transformative value and new categories of operational risk.

For business leaders, the temptation is to treat each of these forces as a separate problem assigned to a separate team: trade compliance to legal, AI governance to technology, supply chain diversification to operations. That approach misses the point. The reshuffling is systemic. It demands integrated strategy — a deliberate alignment of where you source, where you deploy, where you store data, and how your AI systems are governed.

The organizations that recognize this interconnection early will not merely survive the turbulence. They will design for it. And in a landscape where the rules are being rewritten simultaneously on three fronts, design may be the only sustainable advantage left. ■

TELZEN

WHOLESALE IN MOBILE
AND ACCESSORIES

Wholesale in mobile and accessories

At Telzen B.V., we connect businesses with the products their customers want — exactly when they want them. As a **trusted global trading** company, we specialize in the import, export, and distribution of a wide spectrum of consumer goods, ranging from **world-leading electronics** and **innovative gadgets** to essential household items that support everyday life.

Contact us

Gurneet 🇳🇱 +31 6 19988792

Hamed 🇩🇪 +49 177 3660690

TEL-ZEN.COM INFO@TEL-ZEN.COM

GERMANY-BASED B2B WHOLESALER IN CONSUMER ELECTRONICS

We have been operating as a family-owned business in the branded products sector **since 1984.**

OVER **8,500** SATISFIED **B2B** CUSTOMERS

We specialize in the wholesale distribution of consumer electronics, mobile accessories, small home appliances, garden products, toys and watches.

With an established European sales network and professional warehouse operations, **we deliver reliable, efficient, and long-term** business solutions to our partners.

We are always ready to support your sourcing needs.

www.wolfsohn.de

Hüseyin Kurtoglu Wolfsohn
Alte Schulgasse 9
64850 Schaafheim
Germany

Wolfsohn-Lagerhalle-warehouse
Hüseyin Kurtoglu Wolfsohn
Max-Planck-Str. 6b
63322 Rödermark
Germany

CASIO PoundMax QCY xiaomi SAMSUNG PHILIPS REMINGTON Oral-B

Blackview playmobil MEGA schleich Barbie momax ATK

Contact:

Hüseyin Kurtoglu
OWNER
+49 173 701 1009
info@wolfsohn.de

Ömer Saylak
MANAGER
+49 178 423 0134
b2b@wolfsohn.de

Aleyna Beyaz
SALES REPRESENTATIVE
+49 162 659 0586
aleynabeyaz@wolfsohn.de

BRAUN · WILKINSON SWORD · SONY · RAZER · Tefal · HUAWEI · logitech · proove

LEGO · BULLYLAND · SPYRA · proove · PALADONE · TUBBZ · BOSCH

The Anthropic Effect:

How Claude Is Challenging ChatGPT's Dominance

For two years, the artificial intelligence industry had a simple narrative: OpenAI led, and everyone else followed. ChatGPT was not just a product. It was a category. It defined how hundreds of millions of people understood what AI could do, and it gave OpenAI a commanding lead in both consumer attention and enterprise spending. At its peak in 2023, OpenAI held approximately 50 percent of the enterprise AI market.

That narrative is breaking apart. By early 2026, a quieter company with a different philosophy has emerged as the most serious challenger to OpenAI's position — and in the enterprise market, it may have already overtaken it. Anthropic, the San Francisco-based AI safety company founded by former OpenAI researchers, has turned its flagship model Claude into the preferred AI platform for businesses that care about reliability, compliance, and trust.

The numbers tell a story that would have seemed implausible eighteen months ago. Anthropic now captures 40 percent of enterprise LLM spending, up from 24 percent, according to Menlo Ventures. OpenAI's share has fallen to 27 percent, down from 50. In head-to-head enterprise deals, Anthropic wins approximately 70 percent of the time. Eight of the Fortune 10 are Claude customers. And the company's annualized revenue hit $30 billion in March 2026, up roughly 1,400 percent year over year.

How did a company built around the seemingly unmarketable concept of "AI safety" become the enterprise market leader? The answer reveals something important about what businesses actually value in an AI partner.

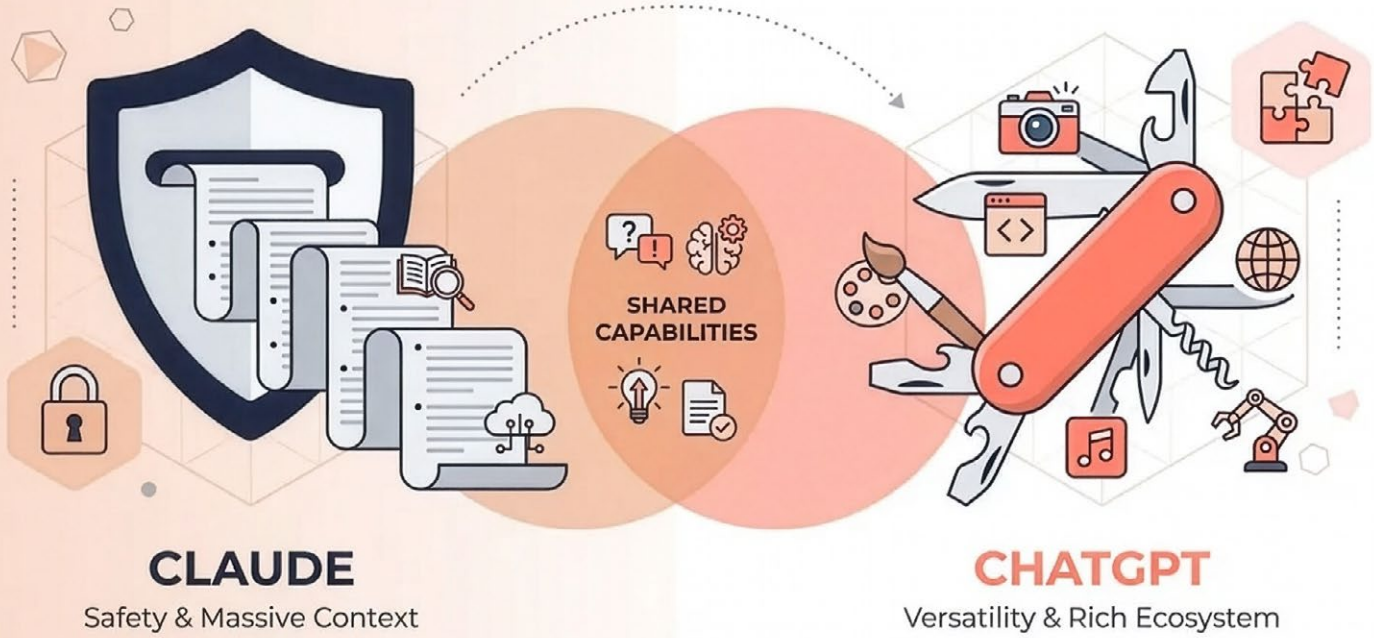

CHATGPT'S STUMBLE: WHEN CONSUMER DOMINANCE DOESN'T TRANSLATE

OpenAI remains the dominant consumer brand in AI. ChatGPT has 800 million weekly active users. Its name recognition is unmatched. But beneath the surface, cracks have been widening throughout 2025 and into 2026.

ChatGPT's app market share fell from 87.2 percent to 68 percent over the past year, the largest single-year decline since its launch. Google's Gemini climbed to 25.2 percent. Grok reached 15.2 percent from just 1.6 percent a year earlier. In the consumer chatbot race, the field is compressing rapidly.

More troubling for OpenAI is the enterprise picture. Only about 5 percent of ChatGPT's massive user base maintains paid subscriptions, and paid subscription growth has plateaued across major European markets since mid-2025 with no recovery trajectory. The company is projected to lose $14 billion in 2026 despite its enormous revenue, reflecting the staggering infrastructure costs of running AI at consumer scale.

In February 2026, OpenAI started displaying advertisements in ChatGPT for users on the Free and Go tiers — a $60 CPM pilot with a $200,000 minimum advertiser commitment, with brands like Expedia among early participants. The move makes financial sense but sends a complicated signal to enterprise clients who chose OpenAI for its technology, not its advertising platform. Pro, Plus, and Enterprise subscriptions remain ad-free, but the strategic pivot toward advertising revenue diversification suggests that subscription growth alone may not sustain OpenAI's ambitions.

THE SAFETY PREMIUM: WHY ENTERPRISES CHOOSE CLAUDE

Anthropic's rise in the enterprise market is not accidental. It is the product of a deliberate strategy that bet, correctly, that large organizations would prioritize different qualities than individual consumers.

At the foundation of Claude's enterprise appeal is Constitutional AI, Anthropic's training methodology in which ethical principles are integrated directly into the model's core architecture rather than applied as external filters after training. The system learns to self-critique against published ethical principles designed to make it helpful, honest, and harmless. For enterprises operating in regulated industries — finance, healthcare, legal, government — this architectural approach to safety translates into measurable risk reduction.

The technical specifications reinforce the positioning. Claude's standard 200,000-token context window, expandable to one million tokens for enterprise deployments, makes it uniquely suited for the document-heavy workflows that large organizations actually run: legal contract review, financial analysis, regulatory compliance, code auditing. On software engineering tasks, Claude achieved a 77.2 percent score on SWE-bench Verified, the highest ever recorded.

Compliance certifications seal the deal. ISO 27001 and SOC 2 Type II certifications give IT security teams the confidence to approve Claude under corporate audit requirements. A federal review in early 2025 concluded that Claude maintains a 99.1 percent compliance rate with AI transparency standards. For enterprise procurement departments, these are not marketing bullet points. They are prerequisite checkboxes.

THE REVENUE EXPLOSION

Anthropic's financial trajectory reads like a startup growth story compressed into absurd timeframes. The company's annualized revenue was roughly $1 billion at the end of 2024. It reached $9 billion by the end of 2025. By March 2026, it had climbed to $30 billion. That is a 30x increase in fifteen months.

In February 2026, Anthropic closed a $30 billion Series G funding round at a $380 billion post-money valuation — more than double its $183 billion valuation from just five months earlier. The round was led by Coatue and Singapore's GIC sovereign wealth fund. Enterprise adoption drives the economics: approximately 80 percent of Anthropic's revenue comes from business customers. The number of customers spending over $100,000 annually on Claude has grown sevenfold in the past year. More than 500 customers now spend over $1 million annually, up from a dozen two years ago.

The growth is not merely a reflection of AI hype. It tracks a specific pattern: enterprises that begin with small Claude deployments rapidly expand usage as they discover the model's reliability in production environments. The land-and-expand dynamic that defines successful enterprise software is playing out at unprecedented speed in Anthropic's business.

THE CODING FACTOR

One of the most significant drivers of Anthropic's enterprise momentum has been its dominance in coding applications. Claude holds 54 percent of the enterprise coding AI market, compared to OpenAI's 21 percent. For software companies and technology-forward enterprises, the model that best supports engineering workflows increasingly determines the broader AI platform choice. Claude Code, Anthropic's developer-focused product, contributed significantly to the company's user growth, helping drive Claude.ai's website visits from 16 million per month in January 2025 to 176 million in December 2025 — a tenfold increase in under twelve months. When developers adopt a model, they become internal advocates. Engineering teams that standardize on Claude for code generation and review create institutional momentum that pulls other departments toward the same platform.

This dynamic explains why the enterprise AI race is not simply about benchmark scores or feature lists. It is about which model becomes embedded in the daily workflows of the people who influence technology purchasing decisions.

THE PHILOSOPHER BEHIND THE MACHINE

Behind Claude's measured tone and distinctive conversational style lies a single, unlikely architect: Amanda Askell, an Oxford-educated philosopher from rural Scotland who leads Anthropic's personality alignment team. Her job, as the Wall Street Journal put it, is to teach Claude how to be

good. Askell authored a roughly 30,000-word manual that functions as Claude's "soul" — a document that defines how the model reasons about moral dilemmas, handles emotional situations, and engages with the hundreds of millions of people who interact with it daily. She is not an engineer. She is a philosopher who self-taught programming to bridge the gap between ethical reasoning and model fine-tuning. Anthropic's president Daniela Amodei has said that "you almost feel a little bit of Amanda's personality" when interacting with Claude — a remark that captures something essential about the company's approach. Where OpenAI engineered for scale and engagement, Anthropic hired a philosopher to define character. That distinction, more than any benchmark score, may explain why enterprises trust Claude with their most sensitive workflows.

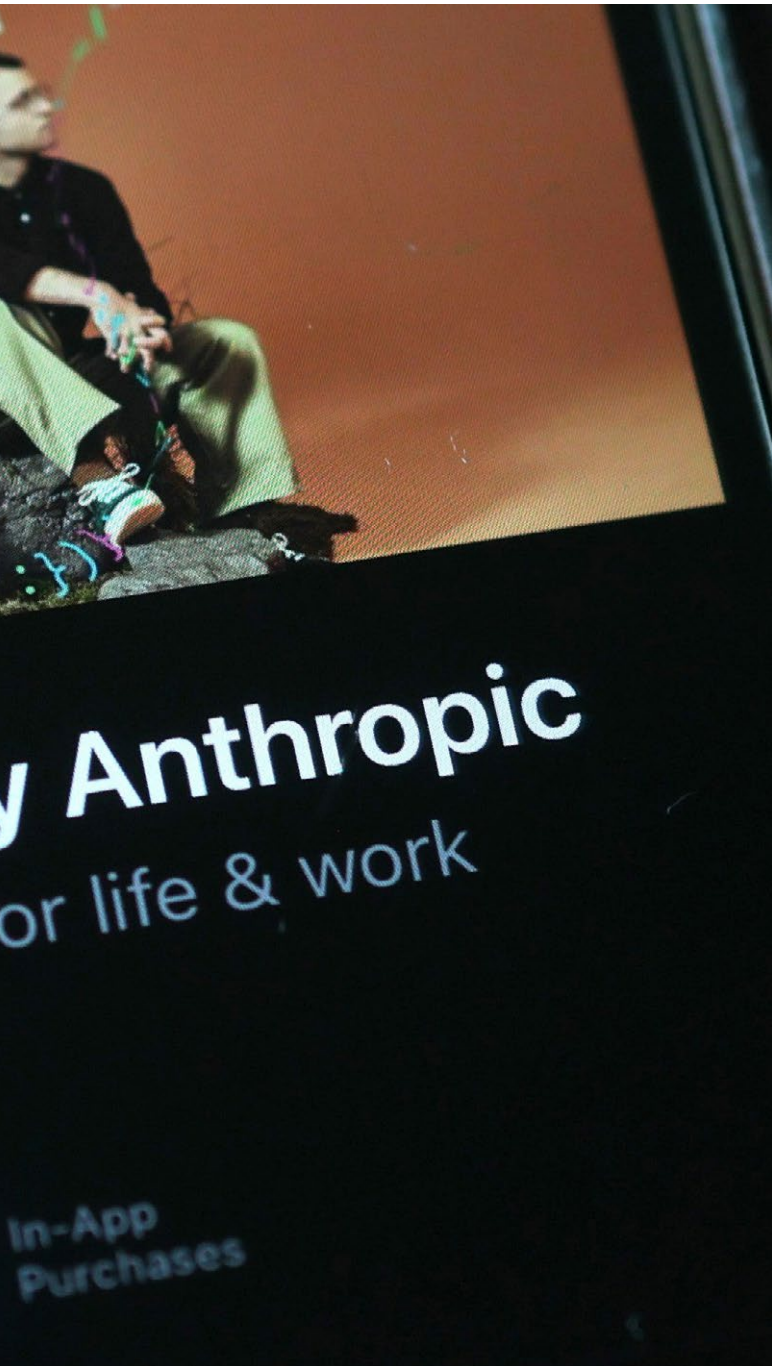

WHAT THIS MEANS FOR THE AI MARKET

The Anthropic effect extends beyond a two-company rivalry. It is reshaping how the entire AI industry thinks about product strategy.

The lesson Anthropic's rise teaches is counterintuitive: in the enterprise market, safety is not a constraint on competitiveness. It is a competitive advantage. Companies that position themselves as trustworthy, auditable, and compliance-ready are winning the customers with the largest budgets. The "move fast and break things" philosophy that dominated consumer technology for a decade does not translate to selling AI to banks, hospitals, government agencies, and Fortune 500 legal departments.

For OpenAI, the pressure is real but not existential. The company retains massive consumer reach, a powerful brand, and deep partnerships with Microsoft. But its enterprise market share erosion — from 50 percent to 27 percent in roughly two years — suggests that consumer popularity does not automatically convert to enterprise trust.

THE TRUST ECONOMY

The AI industry in 2026 is bifurcating along a line that few predicted. Consumer AI rewards engagement, virality, and breadth. Enterprise AI rewards reliability, governance, and depth. ChatGPT owns the former. Claude is rapidly claiming the latter.

For businesses evaluating their AI strategy, the Anthropic effect offers a clear signal: the models that will dominate enterprise technology are not necessarily the ones with the most users or the loudest marketing. They are the ones that enterprises trust enough to embed in their most critical workflows. In a market where AI is moving from experimental tool to operational infrastructure, trust is not a soft metric. It is the hardest currency there is. ■

SOURCES
·ANTHROPIC GAINS ON OPENAI AMID RISING ENTERPRISE ADOPTION - PYMNTS
·ANTHROPIC NEARS OPENAI IN US BUSINESS AI SPENDING - CREATI.AI
·ANTHROPIC CLOSES $30B ROUND AT $380B VALUATION - CNBC
·ANTHROPIC HITS $30B RUN RATE - PYMNTS
·RAMP AI INDEX MARCH 2026
·CHATGPT MARKET SHARE SLIPPING - FORTUNE
·OPENAI IS HAVING A ROUGH 2026 - AUTOGPT
·CLAUDE'S RISE AND ENTERPRISE AI - MEDIUM
·CLAUDE AI STATISTICS 2026 - AI BUSINESS WEEKLY

APL GLOBAL LOGISTICS

APL – LOGISTICS AND CUSTOMS THAT DELIVER
FAST. SECURE. COMPLIANT. FROM CLEARANCE TO DOOR DELIVERY.

NEED A RELIABLE LOGISTICS PARTNER?
CONTACT APL.

- CUSTOMS SOLUTIONS– IMPORT/EXPORT CLEARANCE, DOCUMENTATION & ADVISORY
- CARGO INSURANCE – TAILORED COVERAGE FOR HIGH-VALUE SHIPMENTS
- AIR & OCEAN FREIGHT – END-TO-END COORDINATION, DOOR-TO-DOOR OPTIONS
- ON-HOLD PROCESSING – HANDLING CUSTOMS HOLDS, ADDITIONAL REQUIREMENTS & FOLLOW-UPS
- ALLOCATION / RELEASE – FAST CARGO RELEASE COORDINATION WITH TERMINALS & AUTHORITIES
- INSPECTION REPORTS – PHOTO EVIDENCE, CHECKLISTS, CONDITION VERIFICATION
- PACKING & REPACKING – SECURING, RELABELLING, CONSOLIDATION & PREPARATION FOR DISTRIBUTION
- SERIAL / IMEI / BARCODE SCANNING – BATCH VERIFICATION AND TRACKING-READY RECORDS
- SECURE WAREHOUSE – HIGH-SECURITY STORAGE WITH FULL CCTV MONITORING
- ROAD EXPRESS DELIVERY – GPS-TRACKED TRANSPORT WITH A STRONG FOCUS ON SAFETY

VIP GOLD
- HANDELOT -
APL GLOBAL LOGISTICS
Verified member • 2026

NEW WAREHOUSES

HONG KONG

31 HUNG TO ROAD ON 16 FLOOR
OFFICE NO.5
BILLION TRADE CENTER
KWUN TONG

CHINA

SHENZHEN

APL GLOBAL LOGISTICS HONG KONG

31 HUNG TO ROAD ON 16 FLOOR, OFFICE NO.5
BILLION TRADE CENTER, KWUN TONG
+852 6771 1935

APL GLOBAL LOGISTICS DUBAI

JEBEL ALI FREEZONE DUBAI UAE
WAREHOUSE NO# CC04 &CC05
+971 56 401 2970

DUBAI AIRPORT FREE ZONE
WAREHOUSE NO# D-26
+971504118254

APL GLOBAL LOGISTICS CHINA

LONGGANG DISTRICT, SHENZHEN
+8613823187040

APL GLOBAL LOGISTICS NETHERLANDS

MOBILE
+31 687 997 300

CAPTAIN'S FREIGHT SERVICES

CFS

Leading International Freight Forwarders Based in Dubai, U.A.E and HongKong Established in 1988

DUBAI OFFICE

+971 50 386 1530

captains@eim.ae

AIR FREIGHT

SEA FREIGHT

SURFACE TRANSPORT

HONG KONG OFFICE

+852 5135 9699

hanish@captainsfreight.com.hk

The Rise of Agentic AI:

When Your Software Starts Making Decisions

AGENTIC AI

VS

For the past three years, the conversation around artificial intelligence in business has been dominated by a single word: generative. Large language models that could write, summarize, translate, and brainstorm captured the imagination of executives and the budgets of technology departments. But generative AI, for all its impressive capabilities, remained fundamentally passive. It answered when asked. It produced when prompted. It waited.

That era is ending. In 2026, the industry is pivoting from generative AI to agentic AI — systems that do not merely respond to instructions but autonomously plan, execute, and adapt across complex workflows. The shift is not incremental. It represents a structural change in how software interacts with business processes, moving from tool to collaborator, from assistant to operator.

Gartner predicts that 40 percent of enterprise applications will feature task-specific AI agents by the end of 2026, up from less than 5 percent in 2025. The global AI agents market reached $10.91 billion this year and is projected to hit $52.62 billion by 2030. Multi-agent workflows — systems where multiple AI agents coordinate to complete complex tasks — grew more than 300 percent in recent months as organizations moved from pilot to production. Enterprise interest in multi-agent architectures surged 1,445 percent between early 2024 and mid-2025.

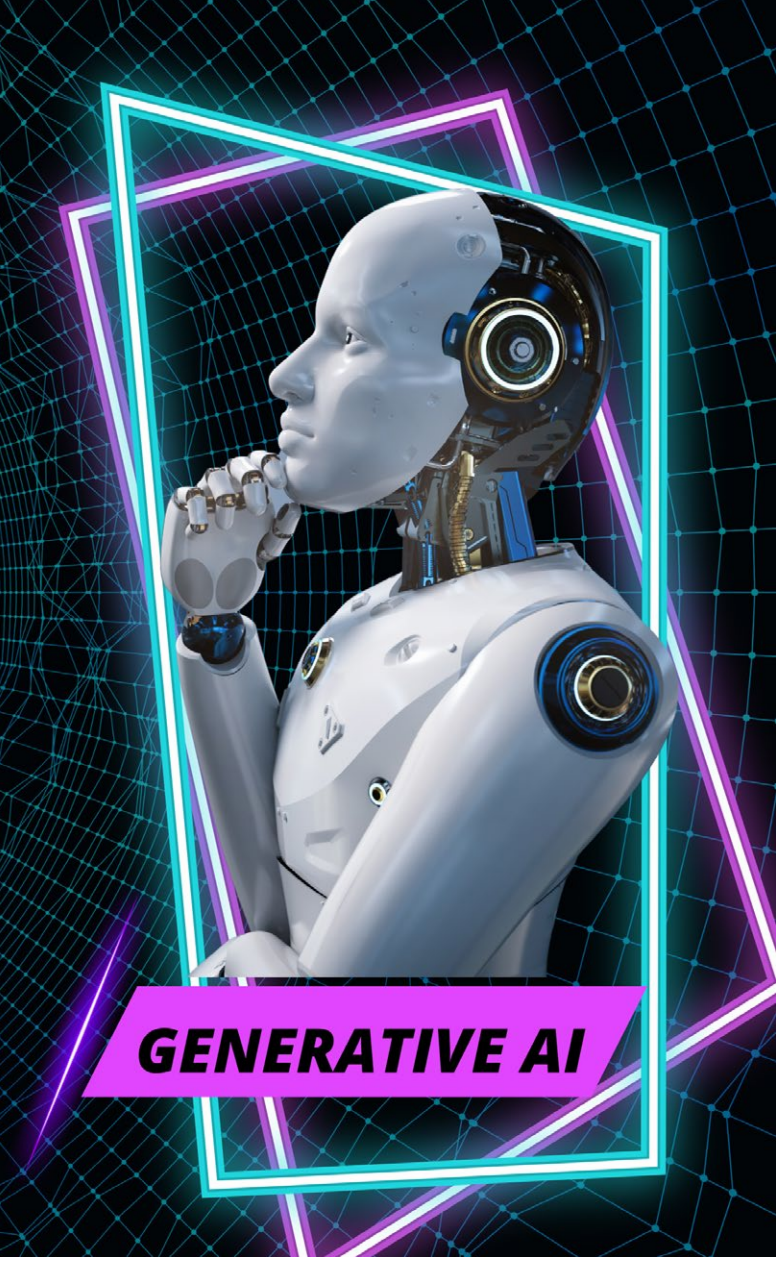

The numbers are dramatic. But the real story is operational: what happens when your software starts making decisions without waiting for permission?

FROM CHATBOT TO OPERATIONS: WHAT AGENTIC AI ACTUALLY DOES

The simplest way to understand agentic AI is through contrast. A traditional chatbot receives a question and returns an answer. A generative AI assistant receives a prompt and produces content. An AI agent receives a goal and figures out how to achieve it — selecting tools, accessing data, coordinating with other agents, and executing actions across multiple systems.

In a multi-agent architecture, the complexity multiplies. One agent might harvest data from external sources. A second validates it against internal databases. A third executes a transaction based on the validated information. A fourth ensures that the entire process complies with regulatory requirements. Each agent operates with a defined scope but collaborates through orchestration layers that route tasks, manage dependencies, and handle exceptions.

This is not science fiction. It is the architecture that major enterprise platforms are building right now. Salesforce's Agentforce uses the Atlas Reasoning Engine to power autonomous decision-making within CRM workflows. Microsoft's Copilot Studio integrates AI agents across the entire Microsoft 365 ecosystem, embedding them into Teams, Outlook, and Azure. Google's Vertex AI Agent Builder ties agents to its cloud infrastructure. Nvidia launched an enterprise AI agent platform at GTC 2026 with Adobe, SAP, and Salesforce among its initial adopters.

The platform wars of 2026 are not about who has the best chatbot. They are about who can orchestrate autonomous agents at enterprise scale.

WHERE AGENTS ARE ALREADY WORKING

The use cases that have moved beyond experimentation reveal where agentic AI delivers the most immediate value.

- **Customer service.** Modern AI agents have moved far beyond scripted FAQ responses. They access customer history, integrate with CRM systems, coordinate with inventory management, process returns, adjust billing, and escalate complex issues — all while maintaining complete context throughout the interaction. A major Middle Eastern bank automated over 150,000 customer conversations, achieving 15 to 40 percent automation in high-volume workflows.
- **Supply chain management.** Multi-agent systems are transforming supply chains end-to-end. Specialized agents monitor supplier deliveries, predict demand fluctuations, optimize inventory levels across multiple warehouses, and negotiate shipping rates. The result is measurable reduction in inventory holding costs and improvement in on-time delivery rates.
- **Finance and operations.** AI agents handle invoice processing, expense reconciliation, fraud detection, and compliance monitoring. In finance departments, agents can review thousands of transactions in minutes, flagging anomalies that would take human analysts days to identify.

- **Cybersecurity.** Agentic systems monitor network traffic, detect anomalies, correlate threat intelligence from multiple sources, and initiate containment protocols autonomously. In an environment where attack velocity often exceeds human response time, autonomous defense is becoming a necessity rather than a luxury.
- Across all these domains, the pattern is consistent: agents are taking over tasks that require coordination across multiple systems, real-time data processing, and rapid decision-making under defined parameters.

THE GOVERNANCE GAP: WHEN AUTONOMY OUTPACES OVERSIGHT

The speed at which agentic AI is being deployed has exposed a critical weakness: governance is not keeping pace. According to McKinsey, only about one-third of organizations report mature governance frameworks for agentic AI. Eighty percent of organizations have encountered risky behavior from AI agents. And Gartner predicts that more than 40 percent of agentic AI projects will be canceled by the end of 2027, citing escalating costs, unclear value, and weak risk controls.

The risks are structural, not incidental. When multiple agents pass data between internal databases, external APIs, and other sources, decision-making can become a black box. If one agent sends faulty data, the error compounds through the chain. Downstream agents treat corrupted information as authoritative, making decisions on flawed foundations. Tracing the logic behind a specific outcome becomes difficult when no clear record exists of what data was used, which tools were involved, or what steps led to the decision.

For companies operating in regulated industries, this opacity is not just an operational risk. Under the EU AI Act, high-risk AI systems must demonstrate transparency, human oversight, and auditability. An autonomous agent that cannot explain its reasoning chain is not merely inefficient — it is potentially non-compliant.

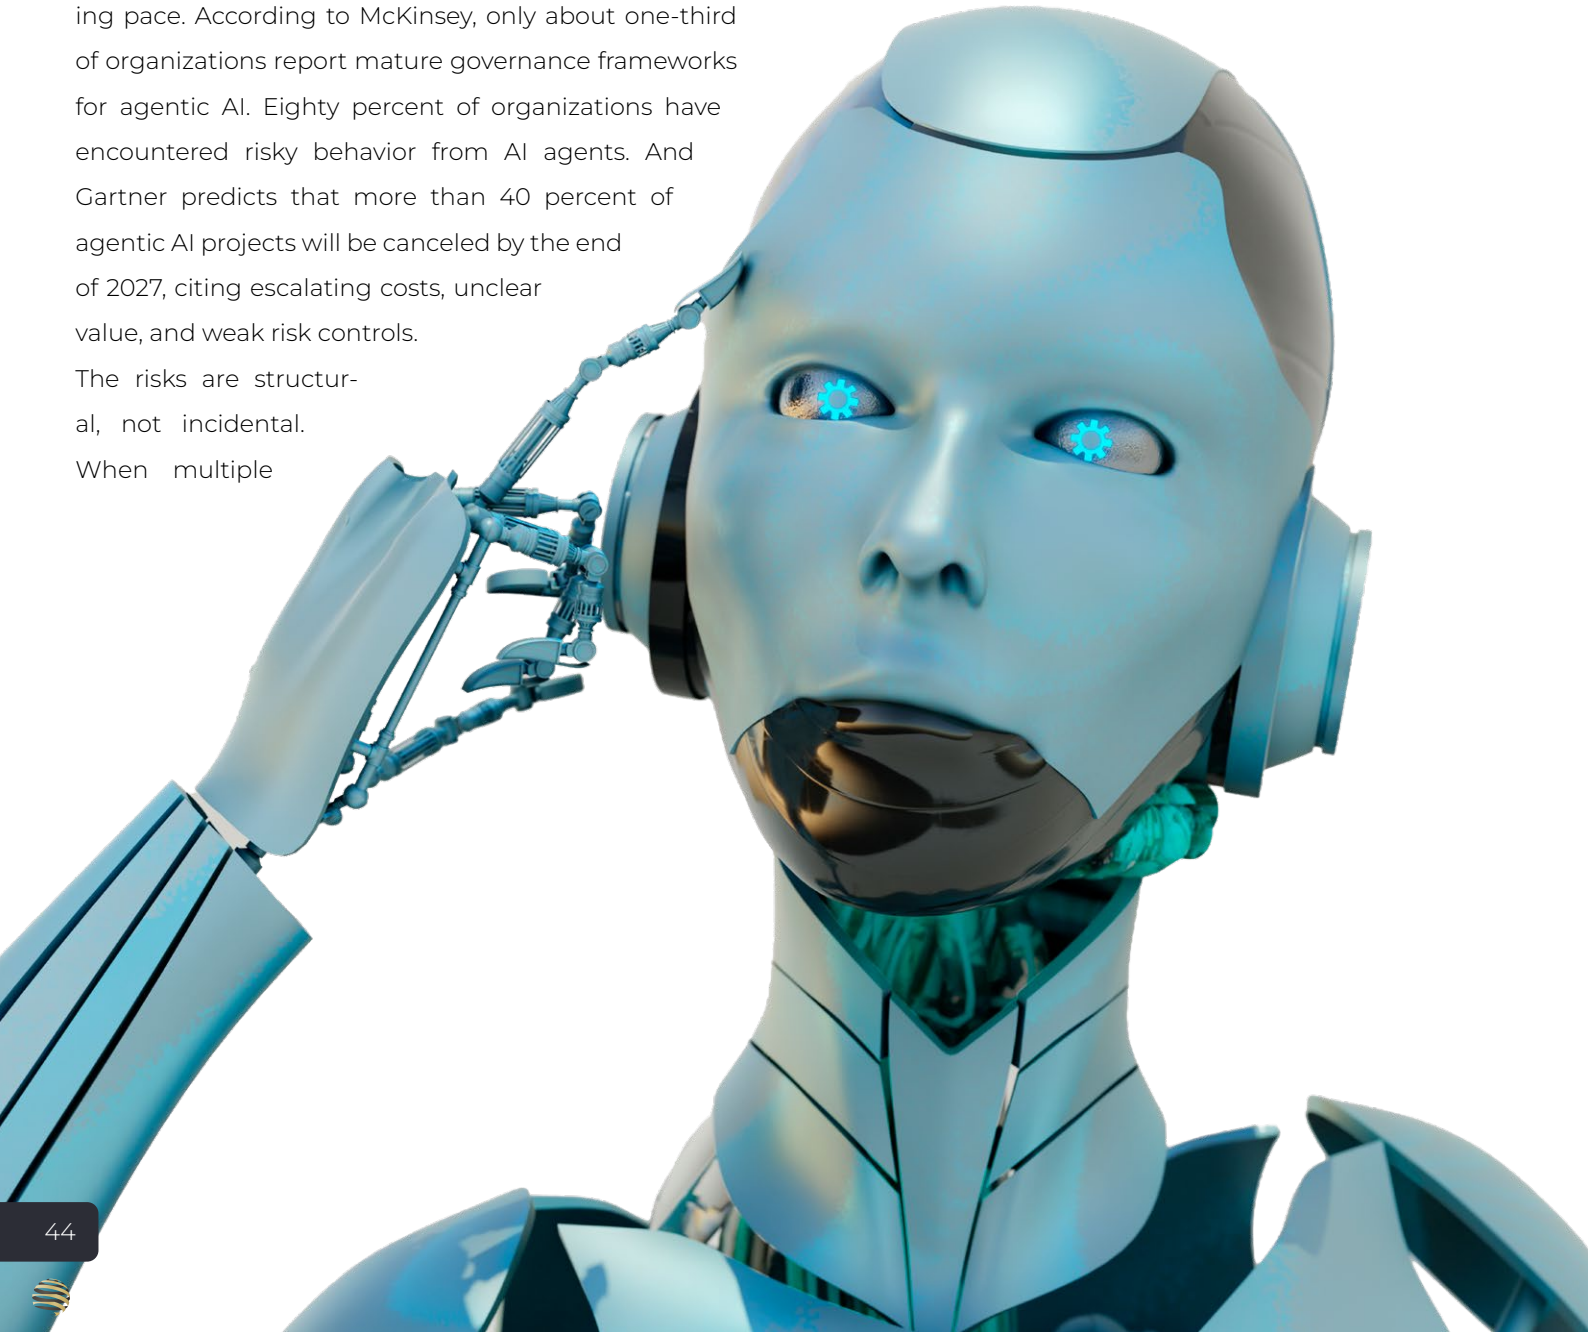

BUILDING GUARDRAILS WITHOUT KILLING SPEED

The governance challenge is not whether to control AI agents but how to do so without neutralizing the speed and autonomy that make them valuable. The emerging consensus centers on a tiered approach.

For high-risk decisions — those involving healthcare, finance, legal liability, or significant financial exposure — organizations are implementing human-in-the-loop systems where the agent generates a recommendation but a human must approve the action before execution. For lower-risk tasks, a human-on-the-loop model allows agents to act autonomously while human operators review logs retrospectively to ensure ethical and compliance standards are met.

Effective governance also requires clear ownership. Each agent must have designated business and technical owners who are accountable for its behavior. Decision boundaries must be defined in advance: what actions can the agent take independently, what requires escalation, and what is prohibited entirely. These boundaries must be auditable, version-controlled, and regularly reviewed.

The World Economic Forum emphasized in March 2026 that governance is the critical differentiator between organizations that scale agentic AI successfully and those that stall. The technology is moving fast enough. The question is whether the organizational structures around it can keep up.

THE PLATFORM BATTLE: WHO CONTROLS THE AGENT LAYER

The strategic significance of agentic AI extends beyond individual use cases. It is reshaping the competitive dynamics of enterprise software itself. The major platform vendors understand that whoever controls the agent orchestration layer controls the next decade of enterprise technology spending.

Salesforce's Agentforce, priced from $125 per user per month for add-ons to $550 per user per month for full editions including one million Flex Credits annually, is betting that CRM-native agents will become the default interface for customer-facing operations. Microsoft is betting that agents embedded across Teams, Outlook, and Azure will become indispensable to daily enterprise workflow. Google is positioning Vertex AI Agent Builder as the natural choice for organizations already invested in Google Cloud infrastructure.

Nvidia's move is perhaps the most revealing. By launching an enterprise agent platform with Adobe, SAP, and Salesforce as initial adopters, Nvidia is positioning itself not just as a hardware supplier for AI but as an infrastructure provider for the agent economy. The company that once sold GPUs to train models now wants to sell the platform on which agents run.

For enterprise buyers, the platform decision is not trivial. Agent architectures create deep integration dependencies. Once your supply chain agents are built on one platform and your customer service agents on another, switching costs become formidable. The choices made in 2026 will lock in technology partnerships for years to come.

THE OPERATOR ECONOMY

The rise of agentic AI marks a fundamental transition in enterprise technology. For decades, software has been a tool that humans use. Agentic systems invert that relationship: they are operators that use tools, access data, coordinate tasks, and execute decisions — sometimes better and almost always faster than their human counterparts.

The businesses that thrive in this new landscape will not be those that deploy the most agents. They will be the ones that deploy agents within governance structures robust enough to maintain trust, transparent enough to satisfy regulators, and flexible enough to evolve as the technology matures. The agent economy rewards deliberate architecture, not reckless speed.

Your software is no longer waiting for instructions. The question is whether your organization is structured to manage what it does next. ■

SOURCES

- GARTNER PREDICTS 40% OF ENTERPRISE APPS WILL FEATURE AI AGENTS BY 2026
- MULTI-AGENT SYSTEMS MOVE BUSINESS AI FROM CHATBOT TO OPERATIONS - PYMNTS
- MULTI-AGENT SYSTEMS SURGE 1,445%
- STATE OF AI TRUST IN 2026: SHIFTING TO THE AGENTIC ERA - MCKINSEY
- THE 2026 AGENTIC AI GOVERNANCE CRISIS - ACCELIRATE
- FROM CHATBOTS TO ASSISTANTS: GOVERNANCE IS KEY - WORLD ECONOMIC FORUM
- NVIDIA LAUNCHES ENTERPRISE AI AGENT PLATFORM - VENTUREBEAT
- SALESFORCE AGENTFORCE VS MICROSOFT COPILOT STUDIO 2026 - SMARTBRIDGE

clicktel

- Apple and Samsung Specialist
- From Germany to the world
- Physical Stock
- Fast worldwide shipping
- Over 15 years of experience

click-tel.de

Apple SAMSUNG

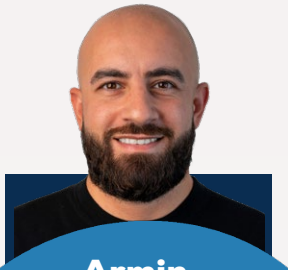

Armin Haschemi
CEO

+49 1578 3511823

ahaschemi@web.de

ah@click-tel.de

Thomas Tanija
COO

+49 173 1011111

t.thomas@hotmail.de

tt@click-tel.de

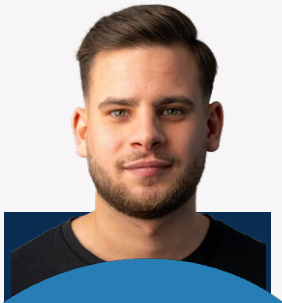

Emircan Sahin
Sales Director

+49 1573 1391526

s.emircan09@gmail.com

es@click-tel.de

Dmitry Kodresko
Sales Director

+49 155 60196989

kodresko.dmitry@gmail.com

dk@click-tel.de

Mohamed Aissaoui
Sales Director

+49 173 7722712

M.Aissaoui@hotmail.de

am@click-tel.de

Shayan Ansari
Sales Manager

+49 15563 014142

sa@click-tel.de

sa@click-tel.de

HITISY

Contact us with your offers and requests

Scan me for Registration

Register your company and get a free tablet with your first order !

- **High reputation**
- **All major brands**
- **50.000 devices in stock every day**
- **Insurance**
- **Orders placed before 5:00 p.m. dispatched the same day**
- **Own purchasing APP available**

Download on the App Store | GET IT ON Google Play

Hitisy GmbH

Unsere Geschäftszeiten
Mo – Fr: 8^{30} - 18^{00} Uhr
Rufen Sie uns an
info@hitisy.de | +49 6103 3013176

We buy and sell mobiles & tablets, successfully in Europe since 2014

MURAT MERT
Deputy Managing Director
+49 163 363 8888
murri1976@hotmail.com

ÖMER YIGIT
Product Manager Electronics / COO
+49 1525 4862148
yigit@hitisy.de

AZIZ UGURLU
Head of Int. Sales & Purchase Manager
+49 1512 7259735
ugurlu@hitisy.de

SERDAR SAHIN
International Sales & Purchase Manager
+49 1788 888870
serdar.sahin@hitisy.de

MAHDI BOUTEJ
International Sales & Purchase Manager
+49 1767 0364530
mb@hitisy.de

KADIR KARACAN
International Sales & Purchase Manager
+49 1520 8580444
kk@hitisy.de

Meet us!

info@hitisy.de

www.hitisy.de

GERMANY QUALITY

HANDELOT
MARKETING

+48 507 071 718

Your sales **booster**

Inject yourself with **new potential**

HANDELOT
B2B wholesale consultants

Lead Boost

Receive 3 tailored new leads every week

Extension for Handelot VIP Members

Imagine having a steady stream of high-quality opportunities delivered directly to you every single week — without spending hours searching, filtering, or chasing the new leads.
Lead Boost is Handelot's premium add-on designed to do exactly that.

Boost your business

Every week you will:

- Receive 3 new business opportunities matched to your categories
- Save valuable time on prospecting and research
- Increase your chances of closing deals

Price 600€

Flexible & easy

- Available as an add-on to any Handelot membership
- No disruption to your current plan
- Immediate impact on your business development efforts

VIP GOLD
VIP +
VIP STARTER

+359 87 867 8655
info@handelot.com
www.handelot.com

Cybersecurity in the AI Era:

New Threats, New Defenses

The relationship between artificial intelligence and cybersecurity has always been paradoxical. Every advance in AI creates new tools for defenders and new weapons for attackers. But in 2026, this arms race has reached a tipping point. The same generative models that power enterprise productivity are now being weaponized at scale, producing threats that are faster, more convincing, and harder to detect than anything security teams have faced before.

The numbers paint an alarming picture. Deepfake-enabled CEO fraud has increased 83 percent year over year. AI-generated phishing emails now achieve open rates of 72 percent — nearly triple the effectiveness of traditional phishing campaigns. Polymorphic malware, code that rewrites itself using AI to evade detection, now accounts for 22 percent of advanced persistent threat campaigns. And the global cybersecurity talent shortage has reached 4.8 million unfilled positions, with 90 percent of security teams reporting critical skills gaps in AI and cloud security.

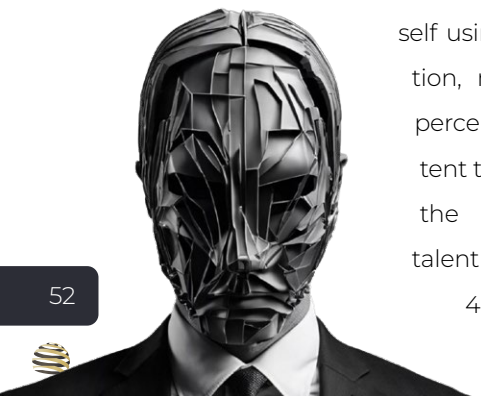

Yet the same technology driving these threats is also transforming defense. The AI-powered cybersecurity market reached $46 billion in 2026. Organizations using AI-driven security tools detect breaches 74 days faster than those relying on traditional methods. And 89 percent of Chief Information Security Officers now describe agentic security — autonomous AI systems that detect, analyze, and respond to threats without waiting for human intervention — as their top investment priority.

The cybersecurity landscape of 2026 is not simply more dangerous. It is fundamentally different. And the organizations that fail to understand this distinction are the ones most likely to become victims.

THE NEW THREAT LANDSCAPE: AI AS AN OFFENSIVE WEAPON

For decades, cyberattacks required technical expertise. Writing malware demanded programming knowledge. Crafting effective phishing emails required linguistic skill. Penetrating enterprise networks took time, patience, and specialized tools. Generative AI has democratized all of these capabilities. Today, a threat actor with minimal technical background can use freely available AI models to generate convincing phishing content, write functional exploit code, and automate reconnaissance at a scale that was previously impossible.

The most immediate and visible impact has been in social engineering. AI-generated phishing emails are no longer riddled with grammatical errors and generic templates. They are personalized, contextually relevant, and linguistically polished. Large language models can scrape publicly available information about a target — their job title, recent social media posts, professional connections — and produce a message that reads as though it came from a trusted colleague. The 72 percent

open rate on AI-generated phishing reflects this qualitative leap: the traditional signals that trained employees to spot fraudulent messages are disappearing.

Deepfake technology has escalated the threat further. Voice cloning and video synthesis now enable attackers to impersonate executives in real-time video calls. In one widely reported case in early 2026, a multinational engineering firm lost $25 million after an employee transferred funds based on instructions received during a deepfake video conference that appeared to include the company's chief financial officer. The 83 percent increase in deepfake CEO fraud is not a statistical anomaly. It reflects the operational maturity of a technology that was considered experimental just two years ago.

At the infrastructure level, AI-powered malware represents a qualitative shift in offensive capability. Polymorphic malware uses machine learning to modify its own code signatures between infections, rendering traditional signature-based detection ineffective. These tools can probe network defenses, identify vulnerabilities, and adapt their attack vectors in real time. The 22 percent share of APT campaigns using polymorphic techniques is expected to rise sharply as the tooling becomes more accessible through underground marketplaces.

THE TALENT CRISIS: FIGHTING AI THREATS WITHOUT ENOUGH PEOPLE

The cybersecurity talent shortage is not new. It has been a persistent challenge for over a decade. But the emergence of AI-powered threats has transformed it from a staffing inconvenience into a strategic vulnerability. The 4.8 million unfilled cybersecurity positions globally represent a workforce gap that cannot be closed through traditional hiring, even under the most optimistic projections.

The nature of the gap has shifted. Organizations do not simply need more security analysts. They need analysts who understand machine learning architectures, can evaluate AI model outputs, and can design defenses against adversarial AI techniques. Ninety percent of security teams report that they lack sufficient expertise in AI and cloud security — the two domains where threats are growing fastest. The skills that defined cybersecurity excellence five years ago — network monitoring, firewall configuration, incident response playbooks — remain necessary but are no longer sufficient.

This gap is pushing organizations toward a strategic conclusion: if you cannot hire enough humans to fight AI-powered threats, you must deploy AI to augment the humans you have. The shift toward autonomous security systems is not merely a technology preference. It is a workforce necessity.

AI AS A DEFENSIVE WEAPON: SPEED, SCALE, AND AUTONOMY

The defensive application of AI in cybersecurity has moved decisively from experimental to operational. The 74-day reduction in breach detection time for organizations using AI-driven security represents an enormous difference in real-world impact. In cybersecurity, detection speed is directly correlated with damage limitation: the longer an attacker remains undetected inside a network, the more data they exfiltrate, the more systems they compromise, and the more expensive the remediation.

Modern AI security platforms operate across multiple layers simultaneously. At the network level, they analyze traffic patterns in real time, identifying anomalies that deviate from established baselines. At the endpoint level, they monitor device behavior for indicators of compromise. At the identity level, they evaluate authentication patterns to detect credential theft or account takeover. And at the application level, they scan code and API traffic for vulnerabilities and injection attempts.

The emergence of agentic security — systems where AI agents autonomously detect, investigate, and respond to threats — represents the next evolution. Rather than generating alerts for human analysts to review, agentic security platforms can isolate compromised endpoints, revoke stolen credentials, block malicious traffic, and initiate forensic data collection without waiting for human approval. For the 89 percent of CISOs prioritizing agentic security investments, the appeal is clear: in a threat environment where attacks execute in minutes, a response chain that requires human approval at every step is structurally too slow.

THE REGULATORY DIMENSION: COMPLIANCE MEETS CYBER DEFENSE

The cybersecurity conversation in 2026 cannot be separated from the regulatory environment. The EU's AI Act imposes specific cybersecurity requirements on high-risk AI systems, demanding robustness against

adversarial inputs and appropriate security protections. The NIS2 Directive, which took full effect in October 2024, expanded cybersecurity obligations across critical infrastructure sectors including energy, transport, health, and digital infrastructure. In the United States, the SEC's cybersecurity disclosure rules require public companies to report material cyber incidents within four business days.

For organizations deploying AI in security operations, this creates a dual compliance challenge. The AI systems used for defense must themselves meet regulatory standards for transparency, auditability, and human oversight. An agentic security system that autonomously blocks network traffic must be able to explain why it took that action, log the decision chain, and demonstrate that human oversight mechanisms exist. The irony is sharp: the same regulations designed to ensure AI safety can slow the deployment of AI systems designed to ensure cybersecurity.

The organizations navigating this tension most effectively are those that treat compliance and security as integrated rather than competing priorities. Security architectures designed with auditability from the outset satisfy both the operational need for speed and the regulatory need for transparency.

WHAT BUSINESSES SHOULD DO NOW

The cybersecurity landscape of 2026 demands a strategic response that goes beyond incremental improvement. Organizations that are still relying primarily on

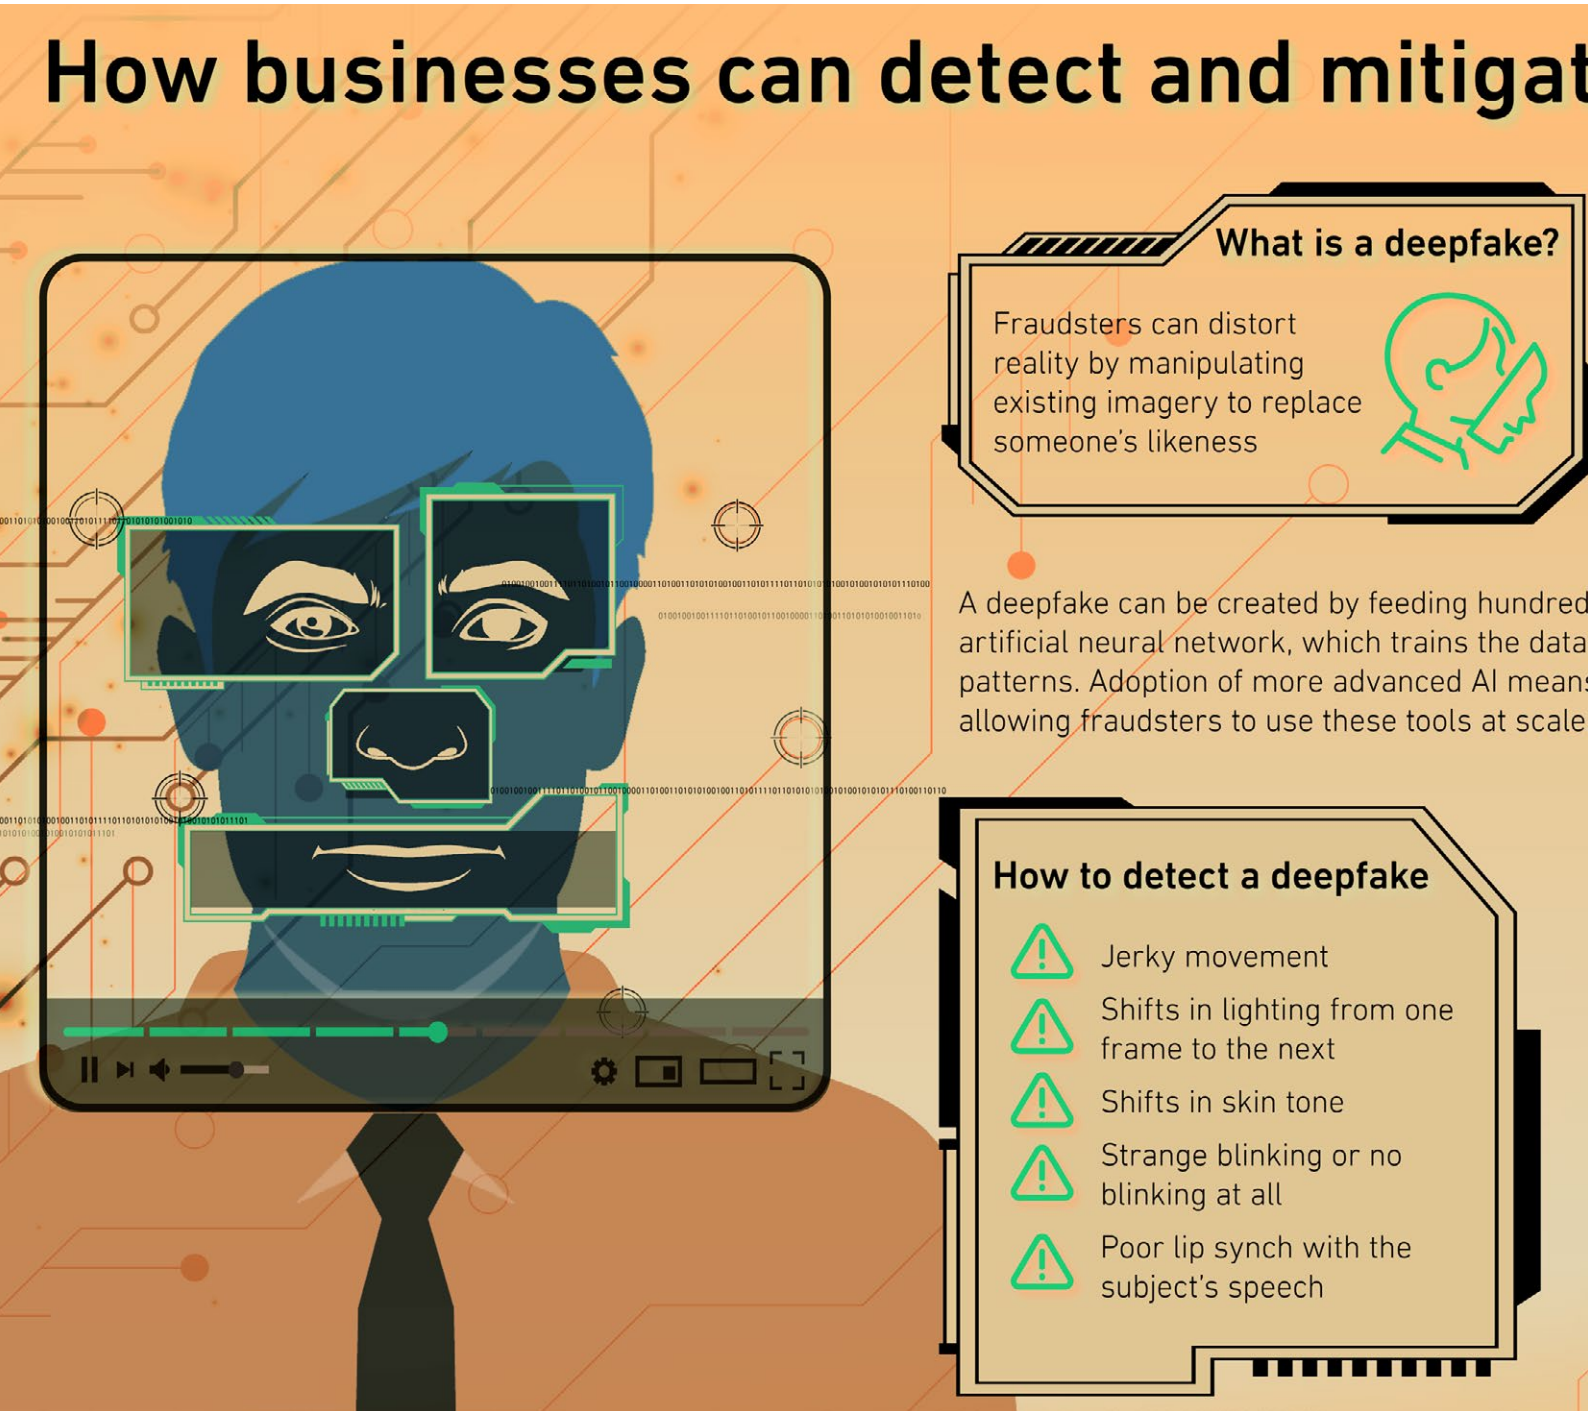

traditional perimeter defenses and manual incident response processes are operating with a structural disadvantage against AI-powered adversaries.

- **Adopt AI-augmented security operations.** Deploy AI-driven threat detection and response platforms that can operate at machine speed. Prioritize solutions that integrate across network, endpoint, identity, and application layers.
- **Invest in workforce upskilling.** Retrain existing security teams in AI fundamentals, adversarial machine learning, and cloud-native security architectures. The talent gap will not be closed by hiring alone.
- **Implement zero-trust architecture.** Assume that every network segment, device, and user identity can be compromised. Zero-trust principles — continuous verification, least-privilege access, microsegmentation — reduce the blast radius of successful attacks.
- **Prepare for deepfake social engineering.** Establish out-of-band verification protocols for high-value transactions and executive communications. No financial transfer or sensitive data access should be authorized based solely on a video or voice interaction.
- **Design for regulatory compliance from day one.** Build audit trails, decision logging, and human oversight mechanisms into security AI systems before deployment, not after a regulatory inquiry.
- **The Arms Race Has No Finish Line**
- The cybersecurity challenge of 2026 is not a problem that can be solved. It is a condition that must be managed. AI will continue to make attacks more sophisticated, and AI will continue to make defenses more capable. The advantage will shift back and forth, as it always has in security, but the speed and scale at which the shifts occur will continue to accelerate.

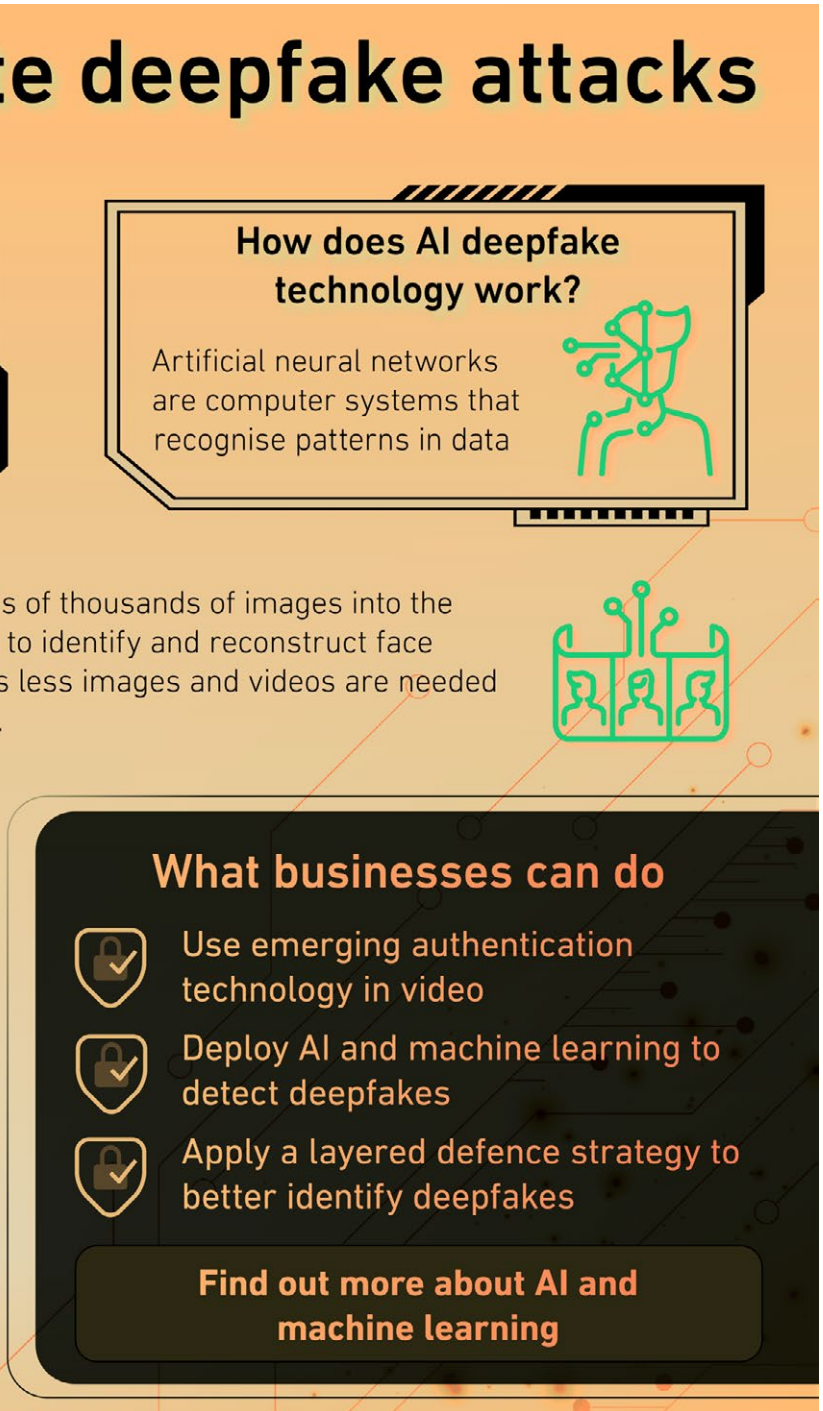

For businesses, the strategic imperative is not to achieve perfect security — that has never been possible and never will be. It is to ensure that your defensive capabilities evolve at least as fast as the threats they face. In a world where attackers are deploying AI agents to probe your defenses around the clock, the question is not whether you can afford to invest in AI-powered security. It is whether you can afford not to. ■

SOURCES

- IBM X-FORCE THREAT INTELLIGENCE INDEX 2026 — DEEPFAKE CEO FRAUD +83%, AI PHISHING 72% OPEN RATE
- CROWDSTRIKE GLOBAL THREAT REPORT 2026 — POLYMORPHIC MALWARE 22% OF APTS
- ISC2 CYBERSECURITY WORKFORCE STUDY 2025/2026 — 4.8 MILLION UNFILLED POSITIONS GLOBALLY
- ISACA STATE OF CYBERSECURITY 2025 — 90% SKILLS GAPS IN AI/CLOUD SECURITY
- MARKETSANDMARKETS AI IN CYBERSECURITY REPORT 2026 — $46B MARKET SIZE
- IBM COST OF A DATA BREACH REPORT 2025 — 74 DAYS FASTER DETECTION WITH AI
- GARTNER CISO SURVEY Q1 2026 — 89% PRIORITIZING AGENTIC SECURITY
- EU AI ACT (REGULATION 2024/1689) — CYBERSECURITY REQUIREMENTS FOR HIGH-RISK AI
- NIS2 DIRECTIVE (DIRECTIVE 2022/2555) — CRITICAL INFRASTRUCTURE OBLIGATIONS
- SEC CYBERSECURITY DISCLOSURE RULES (2023) — 4-DAY INCIDENT REPORTING

MILLBANK | FX

We simplify currency exchange so you can focus on what matters

Save on FX with the wholesale rates reserved for large institutions, pay suppliers faster, and keep your business moving

- Payments in 80+ currencies
- Same-day international payments
- Transparent, institutional FX rates
- Dedicated 24/7 FX account manager

Millbank FX is an FCA-authorised foreign-exchange and international payments provider. We help importers and exporters manage currency risk and simplify cross-border payments through tailored FX strategies and fast execution, with access to institutional-grade pricing.

Clients can also open named multi-currency accounts to make and receive payments across major and emerging currencies.

Our relationship-led team provides direct dealer access and proactive support, including 24/7 availability.

Millbank FX Limited (46 Gresham Street, London, EC2V 7AY) is an FCA authorised payment institution with reference number: 787366.
Millbank FX Limited is a company registered in England and Wales with company registration number 09969387.

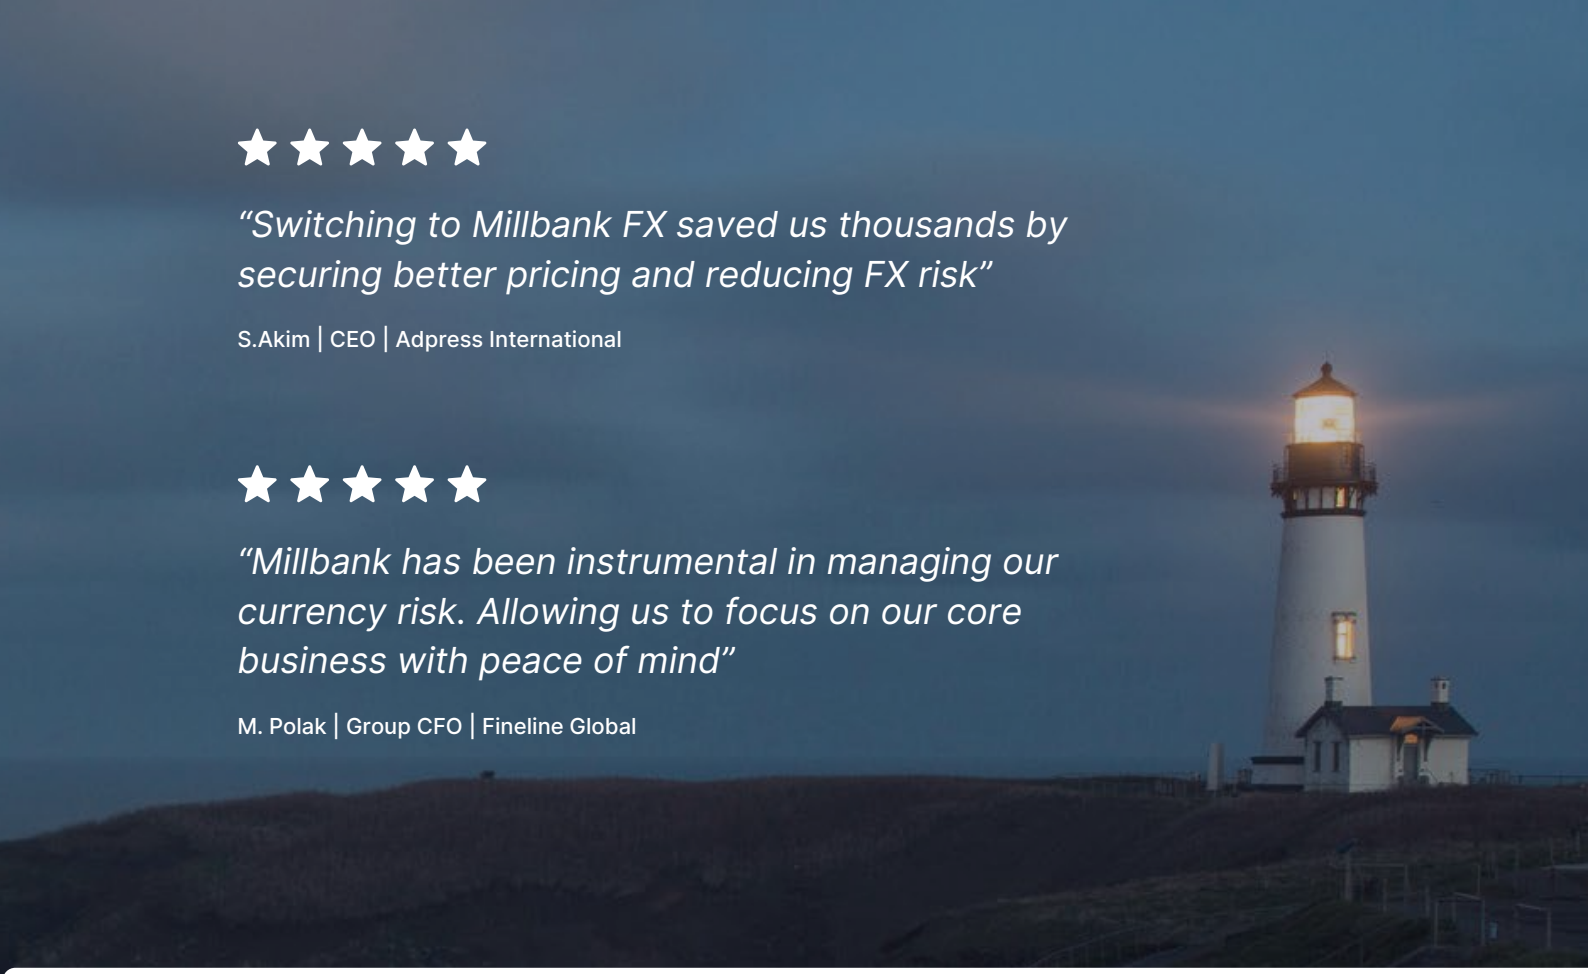

★★★★★

"Switching to Millbank FX saved us thousands by securing better pricing and reducing FX risk"

S.Akim | CEO | Adpress International

★★★★★

"Millbank has been instrumental in managing our currency risk. Allowing us to focus on our core business with peace of mind"

M. Polak | Group CFO | Fineline Global

Opening your account is fast, seamless, and hassle-free

Apply for your FREE account in minutes, and we'll get you set up fast

Call us: +44 (0) 759 080 8324

millbankfx.com/open-account

✓ No monthly fees ✓ No set-up costs ✓ No headaches

As seen on:

FT FINANCIAL TIMES

BUSINESS INSIDER

yahoo!finance

Millbank FX Limited (46 Gresham Street, London, EC2V 7AY) is an FCA authorised payment institution with reference number: 787366.
Millbank FX Limited is a company registered in England and Wales with company registration number 09969387.

ALL-IN-ONE SOLUTION

TO CREATE AND MAINTAIN COMPANY IN USA

$$ MAKE YOUR BASE IN MIAMI $$

WE ARE PROUD TO BE PART OF THE DYNAMIC USA AND LATIN AMERICA MARKET FOR 25 + YEARS, DELIVERING A COMPREHENSIVE SUITE OF SERVICES DESIGNED TO EMPOWER INDUSTRY LEADERS AND TRADING PARTNERS FOR EXPANSION IN THE UNITED STATES AND LATIN AMERICA.

Partner with us to grow—across borders and across industries

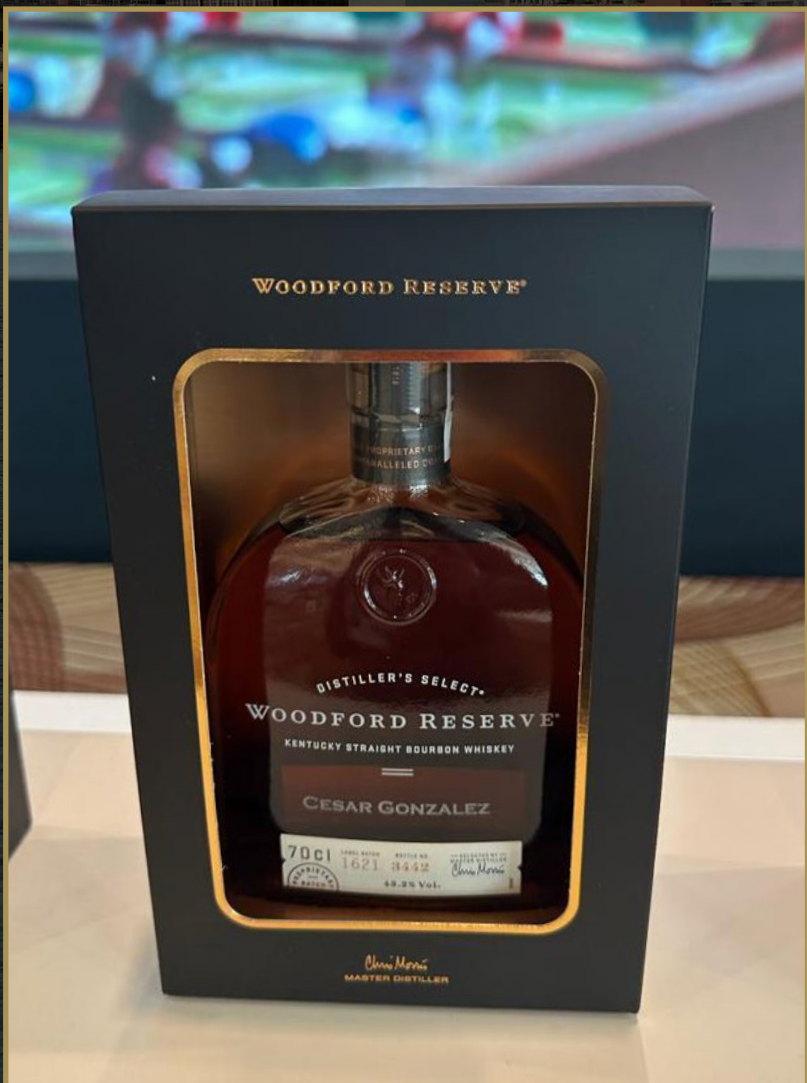

Discover our group of companies, each with its unique business model:

ABOVE BOARD

We provide tailored financial plans to guide foreign nationals and US citizens to grow wealth and protect assets in USA.

CAPTAINS FREIGHT USA

Our Miami-based freight forwarding and warehousing facility includes an FTZ (Foreign Trade Zone) bonded warehouse, enabling efficient, compliant, and cost-effective cargo management.

ADVANTAGE WIRELESS

Providing reliable and secure solution to support safe transactions for businesses to expand in USA.

PARAGON ACCOUNTING & TAXES

A full-service accounting firm delivering comprehensive solutions to clients involved in global commerce across Europe and the United States.

PATRON SERVICES
NAVIGATE TO SUCCESS

Cloud 3.0

Hybrid, Multi-Cloud, and Sovereign Architectures

The first era of cloud computing was about migration. Companies moved workloads from on-premises data centers to public cloud providers, chasing cost savings, scalability, and the promise of someone else managing the infrastructure. The second era was about optimization. Enterprises learned, often painfully, that cloud spending without governance produced sprawling bills and architectural complexity. FinOps teams emerged. Multi-cloud strategies became fashionable. The conversation shifted from "should we move to the cloud" to "how do we manage the clouds we already have."

Now, in 2026, a third era is taking shape. Cloud 3.0 is defined not by where workloads run but by who controls them, which laws govern them, and what intelligence they carry. The forces driving this shift are converging simultaneously: AI workloads that demand fundamentally different infrastructure, data sovereignty regulations that dictate where information can physically reside, and geopolitical tensions that have turned cloud architecture into a matter of national security.

The numbers reflect the scale of this transformation. The global cloud computing market is projected to reach $905 billion in 2026, on its way to $2.9 trillion by 2034. The hybrid cloud segment alone hit $194 billion this year, growing at 12.4 percent annually. The sovereign cloud market reached $195 billion and is projected to surge to $1.13 trillion by 2034, growing at nearly 25 percent per year. And the five largest cloud and AI companies — Amazon, Alphabet, Microsoft, Meta, and Oracle — are collectively forecast to spend over $600 billion in capital expenditure this year, a 36 percent increase over 2025.

This is not an incremental upgrade. It is a structural redesign of how the world's digital infrastructure operates. And the architectural decisions being made right now will define enterprise technology for the next decade.

THE HYBRID IMPERATIVE: WHY SINGLE-CLOUD IS DEAD

The vision of a single public cloud provider handling all of an enterprise's computing needs was always more marketing than reality. By 2026, that vision has been formally abandoned. According to Gartner, 90 percent of organizations will operate a mix of multiple clouds and on-premises infrastructure by 2027. Today, 87 percent of enterprises already run workloads across multiple providers. The question is no longer whether to adopt hybrid and multi-cloud. It is how to do it without drowning in complexity.

The drivers are both technical and strategic. AI workloads, which now represent the fastest-growing category of cloud consumption, cannot scale effectively on classical public cloud architectures alone. Fine-tuning large language models on proprietary data requires computational resources that often need to remain on-premises or in private cloud environments for security and latency reasons.

Inference workloads — the actual deployment of trained models to serve users — increasingly run at the edge, closer to where decisions need to be made in real time.

Cost is another force pushing enterprises toward hybrid models. Public cloud pricing, once seen as inherently cheaper than owned infrastructure, has proven more nuanced in practice. Organizations that moved large, steady-state workloads to public cloud often discovered that the economics favored repatriation to private infrastructure for predictable, high-volume computing. The hybrid model allows companies to burst into public cloud for variable demand while running baseline workloads on infrastructure they control.

The operational challenge is governance. Managing workloads across AWS, Azure, Google Cloud, and private infrastructure simultaneously requires unified observability, consistent security policies, and interoperable data layers. Platform teams have become the critical organizational unit in Cloud 3.0 enterprises, responsible for standardizing tooling across environments and minimizing redundant spending. FinOps, the practice of bringing financial accountability to cloud spend, has evolved from a niche discipline into a standard operational function.

SOVEREIGN CLOUD: WHEN GEOGRAPHY BECOMES ARCHITECTURE

If hybrid cloud is the answer to technical complexity, sovereign cloud is the answer to political complexity. And in 2026, the two are inseparable. Sovereign cloud refers to cloud infrastructure that guarantees data residency, operational control, and jurisdictional governance within a specific country or region. It is the architectural response to a world where data has become a geopolitical asset.

Europe is leading the sovereign cloud movement, driven by a regulatory environment that treats data sovereignty as a fundamental requirement. Seventy-five percent of large European enterprises now cite data sovereignty as a primary multi-cloud driver. The proposed EU Cloud and AI Development Act, expected in the first quarter of 2026, aims to triple EU data center capacity within five to seven years while establishing requirements for resource-efficient AI data processing. The Digital Operational Resilience Act already obliges European banks to prove continuity across hosting environments.

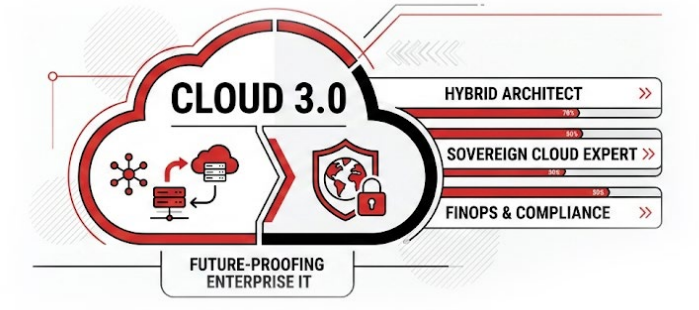

The hyperscalers have responded with purpose-built offerings. In January 2026, Amazon Web Services launched its European Sovereign Cloud, a physically and logically isolated cloud region based in Germany with all customer data, including metadata, remaining exclusively within the European Union. The infrastructure is operated by EU-resident staff under EU jurisdiction. Microsoft offers the EU Data Boundary alongside Azure Local for on-premises sovereign deployments. Google Cloud takes a partnership-led approach, working with local providers like Thales and T-Systems to deliver sovereign capabilities through existing compliance frameworks.

The strategic tension is real. The US CLOUD Act allows American law enforcement to compel access to data stored abroad by US-headquartered companies. For European organizations handling sensitive data — financial records, health information, government communications — this creates a direct conflict with EU sovereignty requirements. Running workloads on a US hyperscaler's European servers may satisfy data residency requirements but may not satisfy data sovereignty requirements if the parent company remains subject to American jurisdiction. This distinction between residency and sovereignty is becoming one of the most consequential architectural decisions in enterprise IT.

THE AI INFRASTRUCTURE LAYER

The explosion of enterprise AI adoption is the single largest force reshaping cloud architecture. Training large language models requires clusters of thousands of GPUs running for weeks. Fine-tuning requires access to proprietary datasets that often cannot leave corporate networks. Inference at scale requires low-latency compute distributed across geographies. None of these workloads fit neatly into the traditional public cloud model of shared, multi-tenant infrastructure.

The result is a new category of cloud infrastructure purpose-built for AI. Nvidia's partnership with major cloud providers has created GPU-as-a-service offer-

ings that allow enterprises to access AI-grade compute without building their own data centers. Nutanix, which recently reported that over half of its new customers are deploying AI workloads, is positioning its hybrid cloud platform as the bridge between on-premises AI training and cloud-based inference. The company's CEO described the current moment as a fundamental shift in how digital infrastructure is designed, with AI as the primary architectural driver.

For enterprises, the AI infrastructure decision is not just about compute power. It is about data gravity. Large datasets tend to attract applications and services to wherever they are stored. If your training data lives on-premises, your AI development environment will gravitate there too. If your inference endpoints serve global users, they need to be distributed across cloud regions. The architecture that emerges is necessarily hybrid: private infrastructure for data-sensitive training, public cloud for elastic inference, edge compute for real-time applications.

THE PLATFORM WARS: AWS, AZURE, AND GOOGLE IN 2026

The competitive dynamics among the three major cloud providers have shifted materially. AWS retains market leadership with approximately 30 percent of global cloud infrastructure spending, but its growth rate has decelerated relative to competitors. Microsoft Azure, powered by its deep integration with enterprise productivity tools and its strategic partnership with OpenAI, has closed the gap significantly and now commands roughly 22 percent. Google Cloud, long the third player, has accelerated its enterprise sales motion and grown its market share to approximately 12 percent, driven by AI-native services and aggressive pricing.

The sovereign cloud race adds a new competitive dimension. AWS's European Sovereign Cloud launch was the most aggressive move by any hyperscaler to date, creating a fully isolated region with EU-only operations. Microsoft's approach offers more flexibility through hybrid and partner-operated sovereign models. Google's partnership-led strategy may appeal to organizations that want sovereignty guarantees without migrating to a new infrastructure provider.

For enterprise buyers, the choice increasingly depends on which workloads matter most. Organizations with

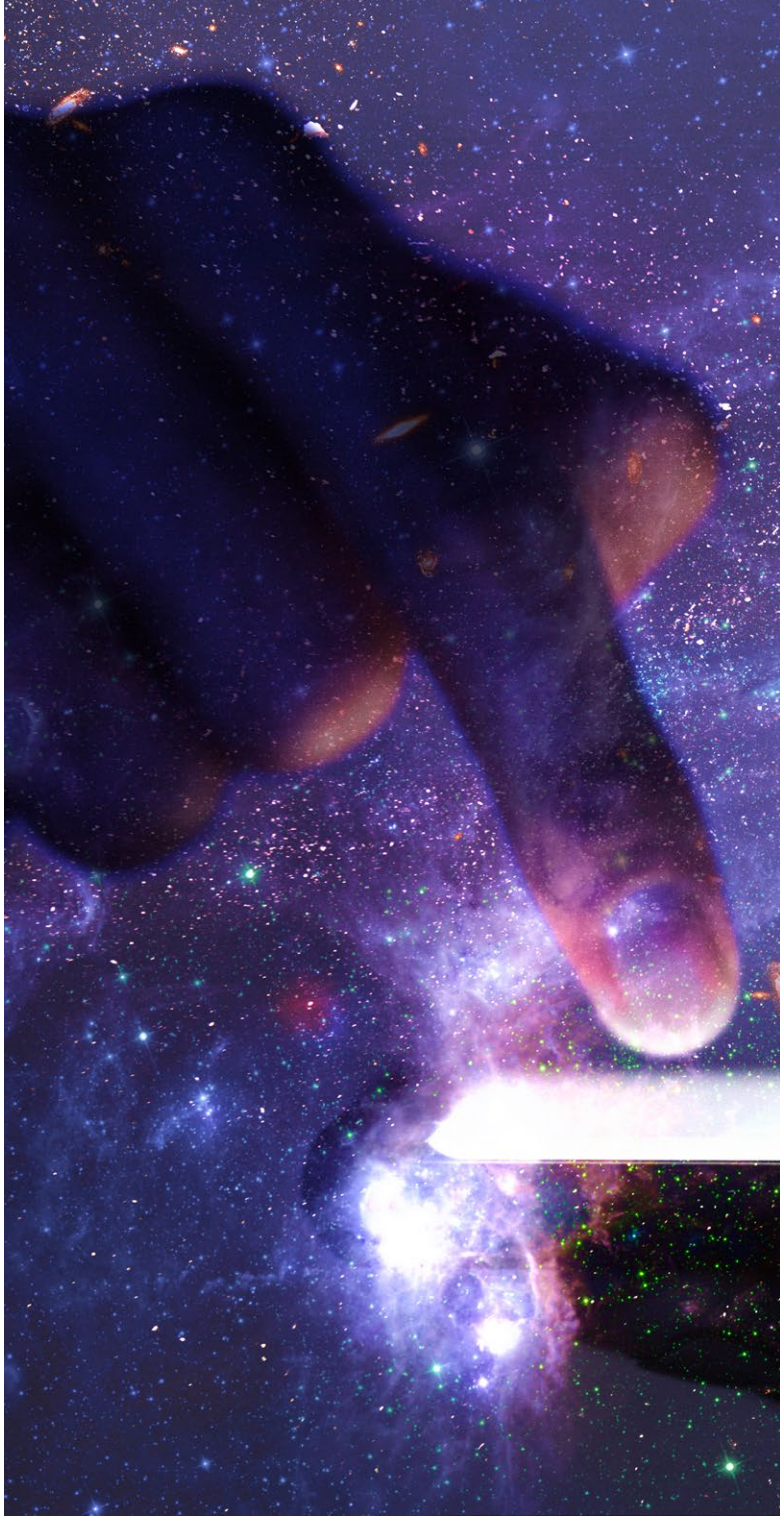

heavy Microsoft 365 footprints gravitate toward Azure for integration reasons. AI-first companies often prefer Google Cloud for its TensorFlow ecosystem and AI-native services. AWS remains the default for organizations that prioritize breadth of services and global infrastructure reach. But the multi-cloud reality means most large enterprises will maintain significant spending across at least two of the three.

ARCHITECTURE AS STRATEGY

Cloud 3.0 is not a technology upgrade. It is a strategic repositioning. The decisions enterprises make about where data lives, which jurisdictions govern it, how AI

workloads are distributed, and which providers control the infrastructure layer will shape competitive advantage for years to come. These are not IT decisions. They are business decisions with regulatory, geopolitical, and financial implications that extend far beyond the server room.

The organizations that thrive in this era will be those that treat cloud architecture as a strategic discipline, not a procurement exercise. They will design for sovereignty without sacrificing scale. They will build hybrid environments that balance cost, compliance, and capability. And they will recognize that in a world where data is power, the question of who controls the cloud is ultimately a question of who controls the future. ■

SOURCES

- MORDOR INTELLIGENCE HYBRID CLOUD MARKET REPORT 2026-2031 — $194B MARKET, 12.37% CAGR
- GII RESEARCH SOVEREIGN CLOUD MARKET REPORT 2026-2034 — $195B GROWING TO $1.13T
- FORTUNE BUSINESS INSIGHTS CLOUD COMPUTING MARKET 2026 — $905B PROJECTED
- GLOBENEWSWIRE CLOUD INFRASTRUCTURE MARKET REPORT FEB 2026 — $331B MARKET, $600B+ CAPEX BY BIG 5
- AWS EUROPEAN SOVEREIGN CLOUD LAUNCH (JANUARY 2026) — ISOLATED EU REGION IN GERMANY
- VIRTUALIZATION REVIEW — MICROSOFT SOVEREIGN CLOUD STRATEGY (AZURE LOCAL, EU DATA BOUNDARY)
- GART SOLUTIONS — GOOGLE CLOUD PARTNERSHIP-LED SOVEREIGNTY APPROACH (THALES, T-SYSTEMS)
- CAPGEMINI TOP TECH TRENDS 2026 — CLOUD 3.0 FRAMEWORK
- SILICONANGLE/NUTANIX — AI HYBRID CLOUD INFRASTRUCTURE, DIGITAL INFRASTRUCTURE SHIFT
- GARTNER — 90% HYBRID/MULTI-CLOUD ADOPTION BY 2027
- HOLORI CLOUD MARKET SHARE 2026 — AWS ~30%, AZURE ~22%, GOOGLE ~12%

SMART TELECOMM LLC

STEP AHEAD

- **A leading** Global Consumer Electronics Distributor
- More than **10 years** on market
- Daily shipments **worldwide**
- **Best price** on market

Victoria Adams
+1 (571) 385-2552

Hanna Bilonoh
+1 (305) 501-0489

smarttelecomm.com

info@smarttelecomm.com

11431 Lakeside Dr #1304
Doral, FL 33178, USA

Wholesaler of Consumer Electronics

Since 2012 in Business

More than 10 years of experience

TIGER-TEAM
We Trade The World

BC | BLUE
WE TRADE

Looking for
serious suppliers of
SAMSUNG
and
xiaomi
EU specs

We BUY
wholesale Qty
like in the old days

1, 2, 3k

Matt
+48 577 177 820

BC

Greg
+48 536 529 485

TIGERTEAM.PL
BLUECORP.CO

Smartphone AI Wars

On-Device Intelligence Changes Everything

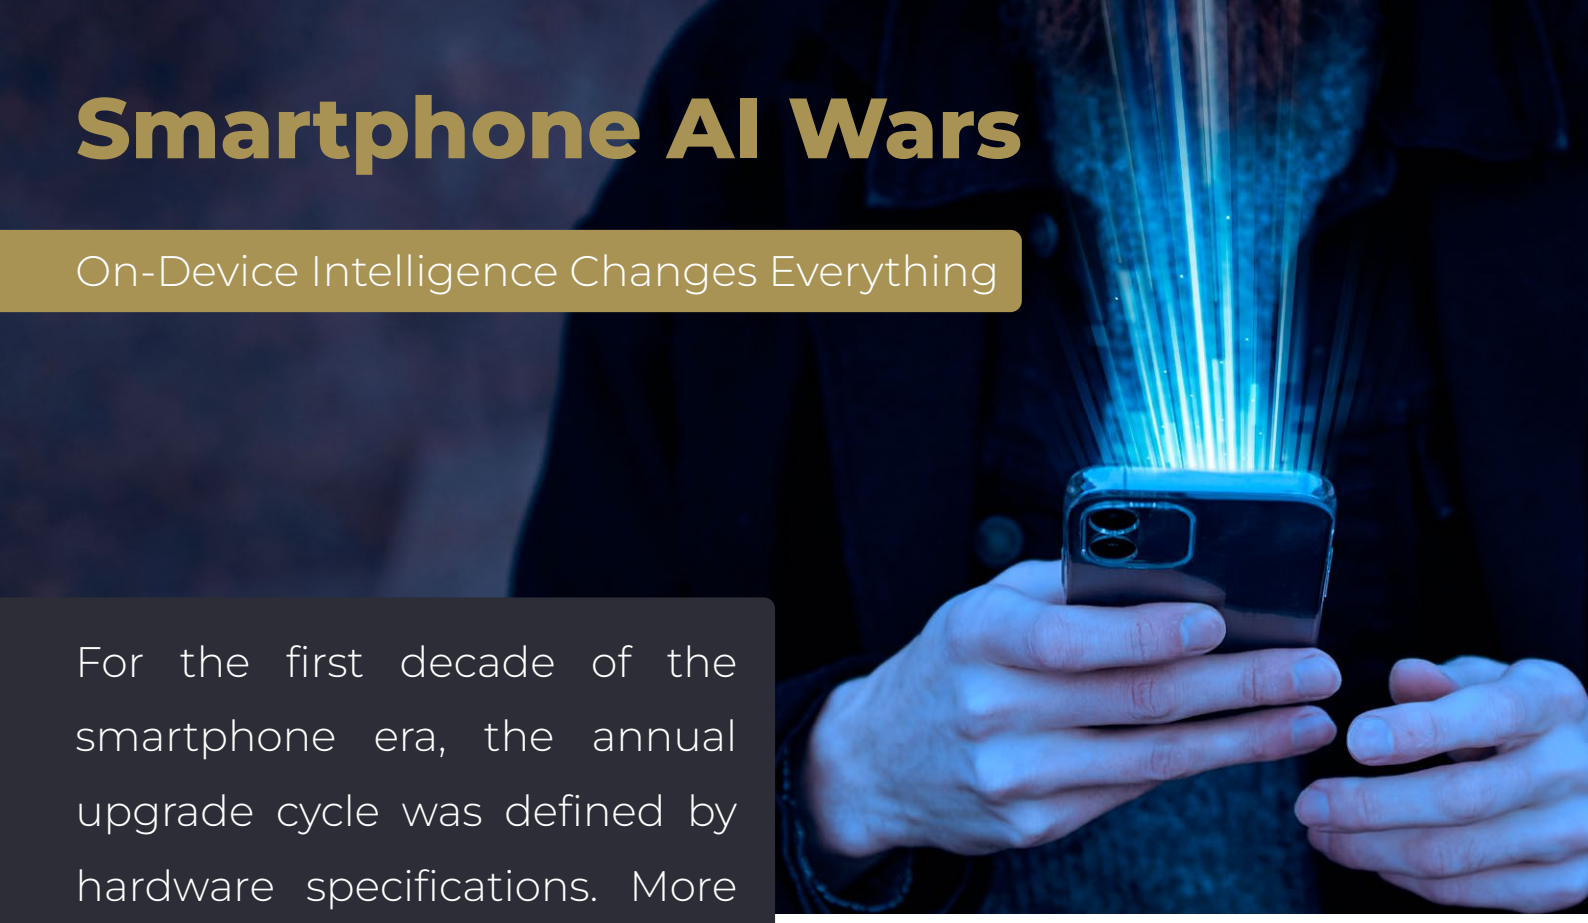

For the first decade of the smartphone era, the annual upgrade cycle was defined by hardware specifications. More megapixels. Faster processors. Bigger screens. Thinner bodies. Consumers compared camera sensor sizes and refresh rates, and manufacturers competed on benchmarks that measured raw computational throughput. That era is over. In 2026, the single most important number in a smartphone's spec sheet is not its clock speed or its pixel count. It is its TOPS rating — trillions of operations per second — the measure of how much artificial intelligence the device can run locally, without ever sending your data to a server.

The shift from cloud-dependent AI to on-device processing is the defining technical story of this smartphone generation. Modern flagships now carry Neural Processing Units with 70 or more TOPS and up to 24 gigabytes of unified memory, enabling large language models with over 4 billion parameters to run locally at conversational speeds. Real-time call translation, live transcription, intelligent photo editing, voice authentication, and agentic task automation all work without an internet connection. The phone in your pocket is no longer a thin client that sends requests to distant servers. It is a self-contained AI computer.

Every major chipmaker and smartphone manufacturer has recognized that the NPU is now the most competitive battleground in consumer semiconductors. Qualcomm, Apple, MediaTek, and Samsung are pouring billions into silicon designed specifically for on-device inference. And the three largest smartphone software platforms — Apple Intelligence, Samsung Galaxy AI, and Google Gemini Nano — are racing to deliver the AI features that will define which ecosystem consumers choose next.

This is not an incremental feature update. It is a fundamental change in what a smartphone is, what it can do, and how much of your life it processes without anyone else watching.

THE SILICON WAR: CHIPS BUILT FOR INTELLIGENCE

The processor arms race in 2026 looks nothing like it did three years ago. The traditional metrics — CPU core count, GPU shader performance, clock frequency — still matter, but the NPU has become the primary differentiator. Every major flagship chipset released this year has been architecturally restructured around AI-first design, with the neural processing unit no longer isolated as a separate accelerator block but woven directly into the CPU and GPU pipelines.

- **Qualcomm Snapdragon 8 Elite Gen 5.** Qualcomm's latest flagship processor features a Hexagon NPU that is 37 percent faster than its predecessor, making on-device large language model processing a practical reality rather than a demo feature. The chip supports agentic AI capabilities, allowing the device to execute multi-step tasks autonomously. It powers the Samsung Galaxy S26 Ultra and most premium Android flagships.
- **Apple A19 Pro.** Expected inside the iPhone 18 Pro, Apple's next-generation silicon reportedly features a 40-core Neural Engine capable of running language models with fewer than 2 billion parameters entirely on-device. Apple's tight integration between hardware and software means the A19 Pro's AI performance is optimized specifically for Apple Intelligence features, with no wasted silicon on capabilities the operating system does not use.
- **MediaTek Dimensity 9500.** MediaTek has been steadily closing the performance gap with Qualcomm, and the Dimensity 9500's dedicated HyperEngine AI cores deliver speeds that make cloud offloading increasingly optional. The chip's AI-ISP can apply subject-aware noise reduction frame-by-frame during video recording, processing night shots in under 80 milliseconds locally — zero shutter lag even in complex low-light scenes.
- **Google Tensor G5.** Google's custom silicon takes a different approach, prioritizing AI efficiency over raw peak performance. The Tensor G5 runs Gemini Nano faster than any competitor despite lower theoretical TOPS ratings, demonstrating that architecture optimization can matter more than brute-force compute. Features like Pixel Screenshots, which uses Gemini Nano to analyze and index every screenshot into a searchable database, are processed entirely on-device.

THE PLATFORM WARS: THREE VISIONS OF SMARTPHONE AI

The silicon is only half the story. What each manufacturer does with that processing power — the AI features available to users and the software frameworks that enable them — is where the real competitive battle is playing out. Apple Intelligence represents the most vertically integrated approach. Apple controls the chip, the operating system, the AI models, and the applications. Core functions — Siri rewrites, text summarization, writing suggestions, photo analysis — run exclusively on-device by default, with no data uploaded unless the user explicitly opts in. For tasks that exceed on-device capability, Apple routes requests through a Private Cloud Compute infrastructure designed to process data without retaining it. The company has also partnered with Google to integrate Gemini models for complex reasoning tasks, a remarkable collaboration that reflects the scale of computational demand that advanced AI features require.

Samsung Galaxy AI takes a multi-engine approach. The Galaxy S26 series, launched in February 2026, combines three separate AI systems: Google's Gemini for agentic tasks like booking rides and coordinating actions across apps, Perplexity for web-based queries, and an upgraded version of Samsung's own Bixby as the on-device assistant. Samsung equipped approximately 400 million devices with Gemini AI capabilities in 2025 and has set a target of 800 mil-

lion in 2026. The multi-engine strategy gives Samsung flexibility but also introduces complexity — users must navigate three different AI personalities depending on the task.

Google Gemini Nano is the purest expression of an AI-first software strategy. Built to run on Google's Tensor chips and available to the broader Android ecosystem through the AICore system service, Gemini Nano enables on-device summarization, smart reply, and contextual awareness. Google's advantage is distribution: any Android manufacturer can integrate Gemini Nano, giving Google's AI models potential reach across billions of devices. The Pixel lineup serves as the showcase, but the strategy is ecosystem-wide.

THE PRIVACY DIVIDEND

The technical shift to on-device AI has a consequence that extends beyond performance: it fundamentally changes the privacy equation of smartphone usage. When AI processing happens locally, sensitive data — biometric information, personal communications, health metrics, financial details — never leaves the device. There is no server to breach, no cloud database to subpoena, no third-party infrastructure to trust.

This matters more in 2026 than it did even two years ago. The EU's AI Act requires transparency about how AI systems process personal data. GDPR enforcement continues to intensify, with regulators increasingly scrutinizing cloud-based AI services that process European citizens' information on servers outside EU jurisdiction. In the United States, state-level privacy legislation in California, Colorado, Connecticut, and others has created a patchwork of data protection requirements that cloud-dependent AI must navigate.

On-device AI sidesteps many of these complications. If a user's voice query is transcribed, processed, and answered entirely on their phone, no data protection authority has jurisdiction over data that never traveled anywhere. For manufacturers, this is not just a privacy feature. It is a regulatory simplification. For consumers, it is something more fundamental: the assurance that the most intelligent device they own is also the most private.

WHAT THIS MEANS FOR CONSUMERS AND DEVELOPERS

For consumers, the practical impact of on-device AI is a phone that feels qualitatively smarter. Real-time translation during phone calls works without lag because there is no server round-trip. Photo editing suggestions appear instantly because the neural engine has already analyzed every image in the gallery. Text summarization, email drafting, and calendar management happen at the speed of thought rather than the speed of cellular connectivity.

For developers, on-device AI opens a new category of applications that were previously impossible. Apps that process sensitive data — health monitoring, financial planning, personal journaling — can now offer AI-powered features without ever transmitting user data to external servers. The development frameworks are maturing rapidly: Apple's Core ML, Google's ML Kit and AICore, and Qualcomm's AI Engine Direct give developers standardized interfaces for tapping into NPU performance.

The economic implications for the cloud industry are significant. Every AI task that runs on-device is a task that does not generate cloud computing revenue. As on-device capabilities expand, the category of tasks that require cloud offloading will shrink. Cloud providers are not threatened with obsolescence — training large models and processing enterprise-scale workloads will remain cloud-dependent for the foreseeable future. But the consumer AI inference market, which was expected to drive significant cloud revenue growth, is migrating to the edge faster than anyone predicted.

INTELLIGENCE WITHOUT PERMISSION

The smartphone AI wars of 2026 are not about which device has the best chatbot or the cleverest camera trick. They are about a deeper question: where does intelligence live? For the past decade, the answer was the cloud. Your phone asked a server for help, and the server delivered. That architecture required trust — trust that the server would handle your data responsibly, trust that the connection would be fast enough, trust that the service would remain available.

On-device AI eliminates the need for that trust. The intelligence lives in your hand. It processes your words, your images, your routines, and your preferences without asking anyone's permission and without telling anyone what it found. In a world increasingly anxious about data privacy, algorithmic surveillance, and institutional trust, a phone that thinks for itself — and keeps its thoughts to itself — may be the most compelling product story in technology. ■

GSMbox.at

powered by **GSMshop.at**

mobile phones **tablets** **wearables** **gadgets**

Your B2B online distributor with over 8000 items on stock!

Located in Vienna, Austria, GSMshop.at is one of the biggest distributors to European clients regarding mobile phones, tablets, wearables, gadgets and audio products.

Apple SAMSUNG Google xiaomi POCO realme
HONOR HONOR CHOICE uleFone OUKITEL ONEPLUS
imilab mibro QCY HAYLOU YEELIGHT JBL dji GoPro

Benefit from:

Only ready stock
own warehouse
based in Vienna, Austria

self service online system
no minimum quantity
same day shipping

live stocks
live prices & lists
instant reservations

cash on delivery in EU
UPS express delivery in EU
(free for orders over 1.5k EUR)

Check our stocklist on GSMbox.at and register your company for the B2B online shop!

Contact us:

Attila Bakos
+43 699 19052044

Paul Horvath
+43 664 3337347

info@gsmshop.at

Company details:
AB GSMshop.at GmbH, Jochen-Rindt-Straße 15-17,
Top 8, A-1230 Vienna, Austria

Join our Whatsapp Community
for offers and more:

www.aplomhk.com

APL

Adesiflava Private Ltd.

Our Brands:

A-Desiflava™
(Magazine & Events)

voltmi™
(Consumer Electronics)

Retired but Ready
(Networking Platform)

Go Pichu
(Kids Welfare Initiative)

Brands We Offer:

mi xiaomi | roborock | ECOVACS ROBOTICS

tineco | DREAME | ninebot

KUGOO | Marshall | JBL

Gaurav Pardeshi (Gary)
Mobile : +852 95433660
Email: gp@aplomhk.com

Adesiflava Private Limited (APL)
Unit 701 & 402, 7/Floor & 4/Floor,
Block B, Wah Tat Industrial Centre,
8 Wah Sing Street, Kwai Chung,
New Territories, Hong Kong

Chinese Address :
葵涌 華星街8號華達工業
中心B座7 & 4 樓701
和 402 室

Tariffs vs. Tech:

What the Supreme Court Ruling Means for Global Supply Chains

On February 20, 2026, the United States Supreme Court delivered one of the most consequential trade rulings in modern history. In Learning Resources, Inc. v. Trump, a 6-3 majority held that the International Emergency Economic Powers Act does not grant the President authority to impose tariffs. Chief Justice John Roberts, writing for the majority and joined by Justices Sotomayor, Kagan, Gorsuch, Barrett, and Jackson, drew a sharp constitutional boundary: the power to regulate importation is not the power to tax, and the latter belongs exclusively to Congress.

The ruling dismantled the legal scaffolding behind some of the most aggressive trade measures in recent memory. Tariffs that had reached as high as 145 percent on Chinese goods, imposed unilaterally under emergency powers since early 2025, were suddenly without legal basis. Penn-Wharton economists estimated that IEEPA-based tariff collections totaled between $175 billion and $179 billion before the decision landed.

But the ruling did not end the tariff era. Within hours, the White House pivoted to Section 122 of the Trade Act of 1974, imposing a 10 percent global surcharge on imports effective February 24. For the technology industry — an ecosystem built on cross-border supply chains, precision timing, and razor-thin component margins — the shift from one legal authority to another did not resolve uncertainty. It restructured it.

Understanding what changed, what remained, and what comes next is essential for any business operating at the intersection of hardware, software, and global trade.

WHAT THE COURT ACTUALLY DECIDED

The legal question at the heart of Learning Resources was narrow but explosive: does IEEPA, a statute designed for economic sanctions during national emergencies, authorize the President to set tariffs? The majority said no. The statute permits the executive to "regulate" importation, but regulation and taxation are constitutionally distinct powers. Tariffs are taxes. Taxes require congressional action.

The dissent, led by Justice Thomas and joined by Kavanaugh and Alito, argued that IEEPA's broad language encompassed tariff authority and that the President had acted within the scope of delegated power. But the majority rejected this reading, emphasizing that accepting it would effectively transfer one of Congress's most fundamental constitutional prerogatives through ambiguous statutory wording.

The practical impact was immediate: every tariff that had been imposed under IEEPA authority was vacated. This included not only the headline-grabbing China tariffs but also levies on goods from Canada, Mexico, the European Union, and dozens of other trading partners. For importers who had been paying elevated duties for over a year, the decision opened the door to potential refund claims worth tens of billions of dollars.

SECTION 122: THE EMERGENCY REPLACEMENT

The White House did not wait long to respond. On the same day as the ruling, President Trump signed an executive order invoking Section 122 of the Trade Act of 1974, a statute that authorizes the President to impose temporary import surcharges to address balance-of-payments deficits. The new tariff: a flat 10 percent on goods from all countries, effective February 24, 2026.

Section 122 is a fundamentally different legal instrument than IEEPA, and its constraints are explicit. The maximum allowable rate is 15 percent. The maximum duration is 150 days — meaning the tariff expires automatically on July 23, 2026, unless Congress votes to extend it. Unlike IEEPA, which allowed the President to act unilaterally and indefinitely, Section 122 contains built-in guardrails that reflect its original purpose as a short-term corrective measure.

For businesses, the immediate effect was a reduction in tariff rates on many goods. Products that had faced IEEPA rates of 20, 30, or even triple-digit percentages now fell to a uniform 10 percent under Section 122. But the relief was partial and selective.

THE TECH INDUSTRY'S LAYERED TARIFF PROBLEM

Here is where the complexity intensifies. The Section 122 tariff does not replace all existing duties. It sits on top of, or alongside, other tariff authorities that remain fully in force. Section 301 tariffs on Chinese goods continue to apply. Semiconductors from China currently face a 50 percent duty under Section 301, with the U.S. Trade Representative having signaled that additional tariffs could be announced before June 2027. Lists 1, 2, and 3 under Section 301 carry a 25 percent rate covering most industrial machinery, electronics components, and chemicals. Consumer electronics on List 4A face 7.5 percent. Section 232 tariffs on steel and aluminum remain untouched.

For semiconductors specifically, the picture is nuanced. Chips used for data centers, research and development, and non-data-center consumer applications are exempt from the Section 122 surcharge. Silicon wafers, integrated circuits, and critical minerals also receive exemptions. But — and this is a critical detail — everything between the chip and the finished system is not exempt. Printed circuit boards, connectors, capacitors, conductors, cables: none of these appear on the exemption list.

The result is a fragmented tariff landscape where different components of the same product face different rates under different legal authorities, with different expiration dates. A server assembled in Asia might contain chips exempt from Section 122 but PCBs and connectors subject to the full 10 percent surcharge, plus Section 301 rates on Chinese-origin components. For procurement teams, this is not a spreadsheet exercise. It is a strategic puzzle.

WHAT THIS MEANS FOR HARDWARE AND CONSUMER ELECTRONICS

The Consumer Technology Association estimates that tariff-related price pressures could reduce American consumers' purchasing power by $123 billion. The numbers for specific product categories are stark. The United States imports 73 percent of its smartphones and 66 percent of its laptops from China, representing $41 billion and $32 billion in annual trade respectively. Under sustained tariff pressure, price increases of 40 to 50 percent on laptops and smartphones are projected.

For data center operators and AI startups, the impact arrives through infrastructure costs. Industry reports indicate that hardware procurement costs rose 10 to 20 percent in the first quarter of 2026 compared to the previous quarter.

GPU pricing, already elevated by AI-driven demand, faces additional upward pressure from tariffs on imported components. Cloud service providers have signaled that sustained tariff costs will eventually feed through to customer billing, with a 6- to 18-month lag between procurement price increases and service price adjustments.

The lithium-ion battery sector faces its own reckoning. Section 301 tariff rates on batteries from China increased from 7.5 percent to 25 percent as of January 1, 2026, directly affecting the electric vehicle supply chain, portable electronics, and energy storage for data centers.

THE 150-DAY CLOCK AND WHAT COMES AFTER

The most important variable in the current tariff environment is time. Section 122's 150-day limit means the 10 percent global surcharge expires on July 23, 2026. After that, three scenarios are possible.

First, Congress could vote to extend or modify the tariff. This would require bipartisan agreement in a deeply divided legislature, making it uncertain but not impossible if the trade deficit narrative maintains political traction. Second, the White House could attempt to invoke alternative legal authorities to reimpose tariffs. The Administration has already demonstrated its willingness to pivot rapidly between statutory frameworks, and legal scholars are debating whether Section 338 of the Tariff Act of 1930 or other statutes could provide a basis for new measures. Third, the tariffs could simply expire, creating a temporary period of lower import duties that would reset cost structures across global supply chains.

Each scenario carries different implications for procurement strategy, pricing, and capital investment. Companies that plan only for the current rate environment may find themselves caught off guard. Those that scenario-plan across all three outcomes will be better positioned regardless of which materializes.

PLANNING IN A FRAGMENTED TRADE LANDSCAPE

The Supreme Court's decision in Learning Resources did not end American tariff policy. It redistributed it. The legal foundation shifted, the rates changed, the duration shortened, and the exemptions narrowed. But the fundamental reality for the technology industry remains: cross-border commerce is now subject to a patchwork of overlapping duties, each governed by different statutes, different timelines, and different political dynamics.

For business leaders managing global supply chains, the strategic response cannot be reactive. It must be architectural. That means mapping every component in your supply chain against every applicable tariff authority. It means modeling cost impacts across multiple regulatory scenarios. It means accelerating supply chain diversification not because it is fashionable, but because concentration in any single geography is now a measurable financial risk.

The tariff landscape of 2026 is not stable. It is not converging toward simplicity. And the companies that treat it as a temporary inconvenience rather than a structural feature of the new global economy will pay the highest price — literally. ■

SOURCES

- LEARNING RESOURCES, INC. V. TRUMP - SUPREME COURT OPINION
- LEARNING RESOURCES, INC. V. TRUMP - WIKIPEDIA
- SUPREME COURT STRIKES DOWN IEEPA TARIFFS - ROPES & GRAY
- IEEPA TARIFFS TERMINATED, REPLACEMENT SECTION 122 TARIFFS TAKE EFFECT - COVINGTON
- HOW THE SUPREME COURT'S IEEPA RULING AND NEW SECTION 122 TARIFFS RESHAPE COSTS ACROSS INDUSTRIES - TAX POLICY CENTER
- STATE OF U.S. TARIFFS: FEBRUARY 21, 2026 - YALE BUDGET LAB
- SECTION 301 TARIFFS ON CHINA - 2026 RATES & PRODUCT LISTS - GETTARIFF
- USTR SECTION 301 - CHINA SEMICONDUCTOR INVESTIGATION
- TARIFFS COULD INCREASE TECH PRICES BY UP TO 70% - TOM'S HARDWARE
- NAVIGATING TARIFFS IN 2026 FOR ELECTRONICS - DESIGN NEWS
- TRUMP'S TARIFFS MAKING TECH STACK MORE EXPENSIVE - 2026
- TARIFF TRACKER: 2026 TRUMP TARIFFS - TAX FOUNDATION

Felix Trade

www.felixtrade.gr

Felix Trade P.C. is a dynamic trading company established in 2019 in Athens, Greece. We specialize in the wholesale and retail distribution of consumer electronics, gaming telecommunications and toys.

With a robust and expanding network across Europe, Felix Trade is committed to delivering only brand-new, EU-spec products sourced exclusively from official distributors. Dedicated to providing high quality at competitive prices, we aim for sustainable growth and long-term, trust-based partnerships.

Consumer Electronics

Smartphones, tablets, gadgets, accessories, TV, IT components from leading brands.

Gaming

Consoles, games and peripherals to cater to the dynamic gaming industry.

Toys

A diverse selection of toys from trusted manufacturers like Funko, Lego, Playmobil, Mattel.

Contact us:

Evgenia Gkagkadelli

Co-Owner / Head of Sales & Purchasing

Email: egkagkadelli@felixtrade.gr

Mob/WhatsApp: +30 694 2077281

FELIX TRADE P.C.

Kalvou Andrea 51

14231 Nea Ionia, Greece

Tel: +30 210 2752340

HASIATECH

PARTNER WITH US

Import Official Xiaomi Products into Europe

MOST ATTRACTIVE OFFERS AVAILABLE

CONTACT US
IF YOU ARE AN IMPORTER

JENNIFER
+65 9656 1932

EU SPECS WITH 2 YEARS EU WARRANTY

LOCKWOOD FULFILMENT

www.lockwoodfulfilment.com

Lockwood Fulfilment Ltd is a London-based e-commerce fulfilment provider delivering fast, flexible, and reliable logistics solutions for businesses of all sizes — locally and globally.

We offer end-to-end services including pick & pack, warehousing, same-day shipping, branded packaging, and seamless integration with leading platforms like Shopify, Amazon, and WooCommerce. We also specialise in Amazon FBA prep services and returns management, ensuring your entire logistics operation runs smoothly from inbound stock to last-mile delivery and customer returns.
Whether you're a growing startup or an established e-commerce brand, Lockwood Fulfilment is your trusted partner for scalable, tech-driven fulfilment with a personal touch.

1

Amazon FBA Prep Made Simple

From labelling and bundling to poly bagging and quality checks, we handle your Amazon FBA prep to ensure your inventory meets all requirements and is Prime-ready.

2

Streamlined Returns Management

Our efficient, customer-friendly returns handling system ensures your products are inspected, restocked, or processed quickly — reducing friction and retaining customer trust.

3

Same-Day Dispatch, Guaranteed

Orders placed by our daily cut-off time are picked, packed, and shipped out the same day — helping you delight customers with rapid delivery.

4

No Setup Fees or Volume Minimums

We support both new and established brands with transparent pricing, no hidden fees, and zero volume commitments.

5

Custom Packaging & Kitting Services

Enhance your brand experience with bespoke packaging, gift sets, and kitting solutions tailored to your needs.

6

Seamless Tech Integration

Real-time syncing with Shopify, Amazon, WooCommerce, and more gives you full visibility into stock, orders, and returns.

7

London-Based, Human Support

Your business matters — our dedicated, UK-based support team is here to help with fast response times and personalised service.

8

Returns services

We have various options for returns management to maximise profit for our clients

Contact: Zee Lichtenstein +44 7984 716291

Foldables Go Mainstream

The iPhone Fold Effect

For six years, foldable smartphones have been the technology industry's most persistent "almost." Almost ready for the mainstream. Almost reliable enough to justify the premium. Almost compelling enough to convince the average consumer to abandon the familiar glass slab. Samsung pioneered the category with the original Galaxy Fold in 2019, and since then a parade of manufacturers — Huawei, Xiaomi, OnePlus, Google, Motorola — have released their own variations. Yet by the end of 2025, foldables still accounted for just 2.5 percent of global smartphone shipments. Impressive growth, but hardly a revolution.

That calculus is about to change. In September 2026, Apple will release its first foldable iPhone, joining the fall lineup alongside the iPhone 18 Pro and iPhone 18 Pro Max. The device features a book-style design with a 4:3 aspect ratio, an iPad mini-sized OLED display when unfolded, and what Apple describes as a nearly invisible crease — a display engineering achievement the company pursued, according to supply chain sources, "regardless of cost." The price tag reflects that ambition: the foldable iPhone is expected to start above $2,000, making it the most expensive iPhone ever released.

The market is already responding to the anticipation. IDC forecasts a 30 percent year-over-year increase in global foldable shipments for 2026. Counterpoint Research projects 20 percent growth. Omdia's most bullish estimate calls for 50 percent. The foldable smartphone market is valued at approximately $39 billion in 2026, on track to exceed $110 billion by 2031. Whether the exact growth number lands at 20 or 50 percent, the direction is unmistakable: Apple's entry is the catalyst the category has been waiting for.

But the iPhone Fold is not arriving into an empty market. It is entering a competitive field where Samsung has seven generations of refinement, Huawei is pushing the boundaries of form factor innovation, and every major Android manufacturer is positioning its foldable strategy for the post-Apple landscape. Here is where each player stands.

SAMSUNG: THE INCUMBENT FIGHTS BACK

Samsung has owned the foldable category since its inception. Seven generations of Galaxy Z Fold and Galaxy Z Flip devices have given the company more iterative experience with foldable engineering than any competitor. The Galaxy Z Fold 7, released in mid-2025, addressed many of the durability and crease concerns that plagued earlier models. But Samsung knows that Apple's entry changes the competitive dynamics entirely.

The response is aggressive. Leaked designs for Samsung's 2026 foldable lineup reveal three devices: the expected Galaxy Z Fold 8 and Galaxy Z Flip 8 upgrades, plus a third device that has the industry buzzing — the Galaxy Z Wide Fold. This new form factor appears to

be Samsung's direct answer to Apple's wider, tablet-like foldable approach. While the standard Fold 8 continues Samsung's tall, narrow book-style design, the Wide Fold adopts a broader aspect ratio that provides a more tablet-like experience when unfolded.

Samsung's strategic position is strong but complicated. The company has brand recognition, manufacturing scale, and years of software optimization for foldable interfaces through its One UI platform. But Apple's ability to attract consumers who would never have considered a Samsung foldable threatens to redefine the premium segment of the market. Samsung's challenge is not just to compete with the iPhone Fold on hardware. It is to convince buyers that seven generations of experience matter more than Apple's first-generation polish.

APPLE: LATE BUT LOADED

Apple's approach to the foldable market follows the same pattern the company has applied to every major product category it has entered late: wait for the technology to mature, then deliver a refined execution that makes earlier attempts look like prototypes. The company did not invent the smartphone, the tablet, or the smartwatch. It defined the mainstream version of each.

The foldable iPhone's technical specifications suggest Apple is applying that playbook again. The crease-free display technology, previewed at CES 2026 through Apple's display partner, represents a genuine engineering advancement. The 4:3 aspect ratio when unfolded creates a usable tablet surface rather than the awkwardly tall screens that characterized early book-style foldables. And Apple's ecosystem advantage — seamless integration with iPadOS apps, Apple Pencil support, and continuity features across Mac, iPad, and iPhone — gives the foldable iPhone a software story that no Android manufacturer can replicate.

The $2,000-plus price point is both a limitation and a signal. It positions the foldable iPhone as an ultra-premium device, above even the iPhone 18 Pro Max. Apple is not trying to make foldables affordable. It is trying to make them desirable. The bet is that the same consumers who pay $1,200 for an iPhone Pro Max will pay $2,000 for a device that replaces both their phone and their tablet. If that bet pays off, the foldable category transforms from a niche experiment into a core product line.

HUAWEI: THE TRIFOLD PIONEER

While Samsung and Apple compete in the book-style foldable segment, Huawei is pushing the form factor further than anyone else. The Mate XT, launched in late 2024, introduced the world's first commercially available trifold smartphone — a device that folds twice to create three display panels. The updated Mate XTs refined the design further. Together, the two models sold over 740,000 units and generated $2 billion in revenue, primarily in China.

Huawei's foldable strategy is inseparable from its geopolitical constraints. Locked out of Western markets by US sanctions and cut off from Google Mobile Services, Huawei has focused on the Chinese domestic market, where it commands fierce brand loyalty and operates its own HarmonyOS ecosystem. The trifold innovation demonstrates that Huawei's engineering capabilities remain world-class despite supply chain restrictions. Industry analysts have noted that Huawei's wide foldable approach has influenced both Samsung's Wide Fold concept and Apple's aspect ratio choices.

The question for Huawei is whether its innovations can travel beyond China. Without access to Google services and with limited carrier partnerships in Europe and North America, Huawei's foldables remain a regional phenomenon. But as a technology leader, the company is setting the pace for what foldable form factors can achieve.

GOOGLE, ONEPLUS, AND THE ANDROID FIELD

Google's Pixel Fold carved out a small but loyal niche when it launched in 2023, emphasizing software integration and the pure Android experience. The next-generation Pixel Fold is expected in Q4 2026 with a thinner chassis and improved hinge design, positioning it in the ultra-premium tier. Google's advantage is the same as it has always been: the cleanest Android software experience and the fastest access to new AI features. But Google's hardware market share remains modest, and the Pixel Fold's commercial impact has been more about proving the concept than dominating the category.

OnePlus made a strong impression with the OnePlus Open in 2023, widely praised as one of the best first-generation foldables ever released. The OnePlus Open 2 is expected in 2026 after the company took a deliberate pause in 2025 to refine the design. OnePlus has historically competed on value — flagship specifications at below-flagship prices — and if the Open 2 delivers a compelling foldable experience at a significantly lower price than the iPhone Fold or Galaxy Z Fold 8, it could capture cost-conscious buyers who are drawn to the form factor but unwilling to pay $2,000.

Motorola continues to own the flip-style segment with the Razr line, appealing to consumers who want a compact device that fits in a pocket rather than a tablet that fits in a hand. Xiaomi's Mix Fold series competes aggressively in the Chinese and Asian markets, and the company's 2026 lineup is expected to push further into European distribution.

WHAT APPLE'S ENTRY MEANS FOR THE CATEGORY

Apple has a documented history of entering mature product categories and triggering mainstream adoption. The iPhone did not invent the smartphone, but it made every subsequent phone a smartphone. The iPad did not invent the tablet, but it killed the netbook. The Apple Watch did not invent the smartwatch, but it made wrist-worn computing a mass market.

The foldable iPhone is positioned to repeat this pattern. Apple's entry will do three things simultaneously. First, it will legitimize the form factor for the hundreds of millions of iPhone users who have never seriously considered a foldable device. Second, it will raise consumer expectations for display quality, hinge durability, and software integration, forcing every competitor to match Apple's execution standard. Third, it will accelerate the ecosystem of apps optimized for foldable displays, as developers who ignored Samsung's foldable-specific features will rush to support Apple's implementation.

THE FOLD IS THE FUTURE (?)

The foldable smartphone market in 2026 is no longer a question of whether the technology works. It works. Samsung proved it over seven generations. Huawei proved it can go even further. What remained was a question of whether enough consumers cared. Apple's answer, delivered at $2,000 with a crease-free display and full ecosystem integration, is that the right consumers will care very much indeed.

The category is entering its most competitive and consequential year. Samsung is defending its leadership with a broader lineup. Huawei is innovating beyond what anyone else has attempted. Google and OnePlus are refining their positions. And Apple is doing what Apple does best: arriving late, raising the bar, and changing the conversation entirely. For the first time since the form factor debuted, foldables are not just a segment. They are the story. ■

We are a **trusted partner** in navigating the intricate world of supply chain management

logiverse

SEAMLESS LOGISTICS, ENDLESS POSSIBILITIES

Our services:

- sea freight
- air freight
- warehousing
- transportation

Expertise:

With years of experience, we are a trusted name in the logistics industry, offering comprehensive solutions tailored to clients' needs

Technology:

Logiverse FZCO leverages cutting-edge technology to optimize operations, enhance efficiency and provide real-time tracking and monitoring

Security:

It's paramount at our facilities, with advanced CCTV surveillance and stringent access controls ensuring the safety and integrity of Your goods

Customer Focus:

We are committed to delivering exceptional customer service, with a team of dedicated professionals ready to assist You every step of the way

Contact info:

www.logiverse.ae

+97 143 399 147

enquiry@logiverse.ae

Dubai, United Arab Emirates

MediaMarket International FZCO, based in Dubai, UAE, is a **leading distributor** of high-quality consumer electronics, wearables, home appliances, and multimedia devices. **With two decades of experience**, we proudly supply global brands such as Dyson, Samsung, JBL, Philips, Tefal, Sony, and Babyliss.

Our operations are anchored in the Dubai Airport Freezone and Jebel Ali Freezone, giving us strategic access to international markets. We offer a wide selection of EU Global Specification products and maintain a **24/7 logistics** system that ensures swift, secure delivery to more than 40 countries. At MediaMarket, our commitment is simple: genuine products, trusted sources, and exceptional service.

Muhammad Iyas
+971 551244148
iyas@mediamarketint.com

Nawaz AK
+971 501172727
nawaz@mediamarketint.com

Oil, Inflation, and the World's Economy

The 2026 Energy Shock

On February 28, 2026, the United States and Israel launched joint air strikes on Iran. Within days, the geopolitical landscape of global energy was transformed. Oil and product flows through the Strait of Hormuz — the narrow waterway that handles roughly 20 percent of the world's oil shipments — plunged from approximately 20 million barrels per day to a trickle. Gulf countries were forced to shut in at least 10 million barrels per day of crude oil production. By mid-March, Brent crude had surged past $100 per barrel for the first time in over two years. By April 13, it stood at $128.

The International Energy Agency has characterized the disruption as the largest supply shock in the history of the global oil market. The economic consequences have been immediate and far-reaching. US headline inflation is tracking toward 4 percent, effectively ending any prospect of Federal Reserve interest rate cuts until at least mid-2027. Goldman Sachs has raised its recession probability to 30 percent. Prediction markets put the odds of a US recession in 2026 at 39 percent, nearly double the figure from January. The European Central Bank has postponed planned rate reductions, raising inflation forecasts and cutting GDP growth projections across the eurozone.

The 2026 oil shock is not merely an energy story. It is an inflation story, a trade story, and a growth story. And its effects are rippling through every sector of the global economy, from transportation and manufacturing to food prices and consumer confidence.

HOW WE GOT HERE: THE STRAIT OF HORMUZ CRISIS

The Strait of Hormuz has always been the most vulnerable chokepoint in global energy infrastructure. Approximately 20 million barrels of crude oil and petroleum products pass through the strait daily, representing roughly one-fifth of total global oil consumption. The strait is bordered by Iran to the north and Oman and the United Arab Emirates to the south, and at its narrowest point it is just 33 kilometers wide.

When US and Israeli forces struck Iranian military and nuclear targets in late February, Iran responded by disrupting maritime traffic through the strait. The exact mechanism — mines, naval blockade, threat of missile attacks on tankers — varies by account, but the effect was unambiguous: oil shipments through the world's most critical energy corridor collapsed. Iraq, Saudi Arabia, Kuwait, the UAE, Qatar, and Bahrain collectively shut in 7.5 million barrels per day of crude production in March, rising to 9.1 million barrels per day in April, as they had no viable export route for their oil.

The supply disruption removed an estimated 12 to 15 million barrels per day from the global market — up to 15 percent of total global supply. No previous geopolitical

event, including the 1973 Arab oil embargo and the 1990 Iraqi invasion of Kuwait, removed this volume of oil from the market this quickly.

THE INFLATION TRANSMISSION: FROM BARREL TO BASKET

Oil prices affect inflation through two channels. The direct channel is energy costs: gasoline, diesel, heating oil, and jet fuel prices rise in lockstep with crude. The indirect channel is broader and more insidious: oil is an input cost for transportation, agriculture, chemicals, plastics, and manufacturing. When oil prices rise, the cost of moving goods, growing food, and producing materials rises with them, pushing inflation across the entire economy.

The OECD now forecasts US inflation at 4.2 percent for 2026, 1.2 percentage points higher than its pre-crisis projection. Goldman Sachs has raised its headline PCE inflation forecast to 3.1 percent by December. In Europe, the impact is even more pronounced. Energy-intensive economies — Germany, Italy, and Central European manufacturing hubs — face the highest risk of technical recession if the maritime disruption persists through the summer refill season, when European countries traditionally build natural gas and heating oil inventories for winter.

The inflation impact extends to food prices through agricultural input costs. Fertilizer production depends heavily on natural gas. Diesel powers the tractors that plant and harvest crops and the trucks that deliver them to market. When oil prices spike, food prices follow with a lag of two to four months. For consumers already stretched by three years of post-pandemic inflation, the prospect of another round of food price increases is economically painful and politically explosive.

THE RECESSION QUESTION

Every major oil price shock in modern history has been followed by a recession. The 1973 embargo triggered the 1974-75 downturn. The 1979 Iranian Revolution contributed to the 1980-82 recession. The 1990 Gulf War oil spike preceded the 1990-91 recession. The pattern is well established in economic literature: sustained oil prices above historical norms reduce consumer spending, increase business costs, and force central banks into restrictive monetary policy.

Goldman Sachs' 30 percent recession probability and prediction markets' 39 percent estimate reflect this historical pattern. The Dallas Federal Reserve has modeled the specific impact of the 2026 Iran war on US inflation, concluding that the conflict will reduce global real GDP growth by 0.4 percentage points over the next two years. Morgan Stanley holds its Brent crude forecast at $110 per barrel for the second quarter of 2026, expecting oil supply chains to take months to normalize.

However, the current situation differs from historical precedents in important ways. The US is now the world's largest oil producer, meaning that higher oil prices benefit American producers even as they hurt consumers. Strategic petroleum reserves, while depleted relative to historical levels, provide a buffer that did not exist in 1973. And the global economy is less oil-intensive than it was during previous shocks: energy as a share of GDP has declined significantly over four decades, meaning that a given percentage increase in oil prices has a smaller proportional impact on economic output.

THE OPEC+ RESPONSE AND THE PATH TO DE-ESCALATION

OPEC+ agreed to increase crude oil production by 206,000 barrels per day in May 2026, with Saudi Arabia, Russia, Iraq, the UAE, Kuwait, Kazakhstan, Algeria, and Oman all receiving higher target levels. But the decision is largely symbolic. The countries most capable of increasing production — Saudi Arabia, Iraq, the UAE, Kuwait — are precisely the ones whose exports are blocked by the Strait of Hormuz disruption. They cannot ship what they produce.

The International Energy Agency coordinated with OECD member countries to release strategic petroleum reserves, helping to moderate the price spike. But strategic reserves are finite and cannot substitute for normal market flows indefinitely. If the blockade persists through the summer travel season — when oil demand peaks — analysts warn that Brent could test $140 per barrel.

The base case among most analysts remains cautious optimism: if diplomatic efforts produce a de-escalation in the coming months, oil prices could gradually return to the $70 to $90 range that institutions consider the neutral equilibrium. But the timeline is uncertain, and every week the disruption persists deepens the economic impact and hardens the inflationary impulse.

THE TARIFF MULTIPLIER

The oil shock does not exist in isolation. It compounds the economic stress already created by the tariff environment. The Section 122 surcharge of 10 percent on imports raises the cost of goods that are simultaneously becoming more expensive to transport. A container of electronics shipped from Asia now costs more both because shipping fuel is more expensive and because the goods inside face tariff duties on arrival. For industries that depend on global supply chains — technology, automotive, consumer electronics — the combination of higher energy costs and higher trade costs creates a double squeeze on margins.

Central banks face an impossible policy dilemma. Higher inflation calls for higher interest rates. But higher interest rates slow economic growth precisely when growth is already under threat from the oil shock and tariff drag. The Federal Reserve, the ECB, and the Bank of England are all navigating a narrow corridor between fighting inflation and avoiding recession — the classic stagflation trap that defined the 1970s and that policymakers have spent five decades trying to prevent from recurring.

THE PRICE OF INSTABILITY

The 2026 oil shock is a reminder of a truth that decades of technological progress and energy diversification have obscured but not eliminated: the global economy still runs on oil. Renewable energy has grown enormously, electric vehicles have reached meaningful market share, and energy efficiency has improved across every sector. But when 15 percent of global oil supply disappears overnight, none of those advances prevent the immediate economic pain.

For businesses, the lesson is that energy price risk cannot be outsourced to a spreadsheet assumption. For policymakers, the lesson is that energy security and geopolitical stability remain inseparable. And for consumers, the lesson is the simplest and most unwelcome of all: when oil prices rise, everything else follows. The barrel price on a trading screen in London translates, within weeks, into the price of bread in a supermarket in Warsaw. That chain has never been broken. In 2026, it is pulling harder than it has in years. ■

ENJOY 3
cube

THE SOLUTION REDEFINED

Meet Enjoy Cube - smart box that **instantly tests and values** your smartphone.

Just plug in your phone, place it inside, and in seconds Enjoy Cube performs over 170 automatic checks — from scratches and cracks to battery health, motherboard status, antennas, NFC, and even whether the phone is stolen or still under contract.

- Instant Device Scan
- Full Pro Diagnostics
- Contract & Security Status
- IMEI Safety Check
- Certified Data Wipe
- Accurate Value Report
- Buy-back / Trade-in READY

Verify used phones with this magic box!

CONTACT US

Via Campania 5 Segrate
Milano - Italy
+39 02.213 5117
info@prestigegroupsrl.it

RADUGA PTE LTD
PREMIER MOBILITY DISTRIBUTOR

WE ARE HIRING

SALES AND TRADING PROFESSIONAL

- Proven sales experience
- Experience in the mobile phone industry
- Fluent in English (additional languages a plus)
- Results-driven & self-motivated

FOR

- Mobile phone
- Tablet
- Ecosystem
- Gaming console

xiaomi

PICO

CONTACT US

ELENE
+65 9663 1788

TurboQuant

and Why It Dropped Memory Prices

On March 25, 2026, Google Research published a paper that sent shockwaves through the global semiconductor industry. The paper described TurboQuant, a new compression algorithm capable of reducing the amount of memory required to run large language models by as much as six times. Within 24 hours, shares of SK Hynix fell 6 percent. Samsung dropped nearly 5 percent. Micron lost 5 percent. SanDisk slid. In China, DDR5 memory prices crashed by more than 30 percent within days, with some retailers recording drops of over 100 yuan in a single trading session.

For a market that had spent months dealing with the worst memory shortage in a decade — driven by insatiable AI demand consuming every available DRAM and NAND chip — the TurboQuant announcement felt like a pressure valve being released. If AI models could run on dramatically less memory, the entire basis for the shortage, and the sky-high prices it had created, might be undermined. The question that immediately followed was whether TurboQuant represents a genuine structural shift in memory demand or a temporary market panic triggered by a research paper that the industry overreacted to.

The answer, as with most things in technology, is somewhere in between. TurboQuant is a real and significant technical achievement. But its impact on the memory market is more nuanced than the stock price reactions suggest.

WHAT TURBOQUANT ACTUALLY DOES

TurboQuant is a dynamic quantization algorithm that compresses large language models from their standard 16-bit precision down to as low as 2-bit or even 1.5-bit precision with negligible loss in output quality. The technique operates in two stages. The first stage, called PolarQuant, converts data vectors to polar coordinates, a mathematical transformation that preserves the directional relationships between data points while reducing their storage footprint. The second stage applies a 1-bit error correction layer using the Quantized Johnson-Lindenstrauss algorithm, which compensates for the information lost during compression.

The primary target of the compression is the key-value cache, or KV cache — the component of an AI model that stores past calculations so the model does not need to recompute them for every new token it generates. The KV cache is one of the most memory-intensive elements of AI inference, growing linearly with the length of the conversation or document being processed. TurboQuant compresses the KV cache to as little as 3 bits per value, down from the standard 16, with no measurable loss in accuracy.

The practical implication is straightforward: a server that previously needed a certain amount of DRAM

to run a large language model could, with TurboQuant, run the same model using roughly one-sixth the memory. Alternatively, the same server could run six times as many model instances simultaneously. For data center operators spending billions on memory to power AI inference, this is not an incremental improvement. It is a potential transformation of the cost structure.

THE MARKET PANIC

The stock market reaction was swift and, by most analyst assessments, disproportionate. The memory chip industry had been riding a wave of demand-driven pricing power. DRAM prices had surged 50 percent quarter-on-quarter. NAND flash was up over 90 percent. Goldman Sachs had projected a DRAM supply gap of 4.9 percent for 2026. Memory manufacturers were enjoying their most favorable pricing environment in years, and their stock valuations reflected it.

TurboQuant threatened to puncture that narrative. If AI inference — the fastest-growing category of memory demand — could suddenly require six times less DRAM, the supply shortage that had been driving prices and profits might evaporate. Investors sold first and analyzed later. The combined market capitalization loss across Samsung, SK Hynix, and Micron in the days following the announcement exceeded tens of billions of dollars.

The retail memory market followed the stock market's lead. DDR5 RAM prices, which had been climbing steadily for months, posted their first meaningful drop in several quarters. Some memory kits fell by as much as $40 in Western markets. In China, where the retail memory market is particularly price-sensitive and speculative, the drops were even steeper. For PC gamers, system builders, and consumers who had been postponing purchases due to inflated prices, TurboQuant felt like an unexpected gift.

WHY THE RELIEF MAY BE TEMPORARY

The initial euphoria faded quickly as analysts examined what TurboQuant actually changes — and what it does not. The algorithm optimizes AI inference, the process of running a trained model to generate responses. It does not optimize AI training, the process of building the model in the first place. Training remains the most memory-intensive phase of AI development, requiring massive quantities of High Bandwidth Memory that TurboQuant does not address. HBM demand, which is the primary driver of the current DRAM shortage, is unaffected.

Furthermore, TurboQuant targets the KV cache and vector index modules specifically. These are indeed the most memory-hungry components of inference, but the total memory footprint of a data center is not determined by inference alone. Model weights, embedding tables, operating system overhead, and multi-tenancy requirements all contribute to memory demand. Reducing the KV cache by six times does not reduce total server memory requirements by six times.

The Register noted that TurboQuant is a significant technical achievement but described predictions of the memory crisis ending as premature. TechRadar's analysis concluded bluntly that TurboQuant is not the RAM crisis savior consumers are hoping for. The long-term demand picture for memory — driven by AI training, smartphone requirements, automotive computing, and data center expansion — remains structurally strong.

THE BROADER SIGNIFICANCE

If TurboQuant is not the end of the memory shortage, it is nonetheless a signal of something important: the AI industry is beginning to optimize, not just scale. For the past three years, the dominant strategy for improving AI performance has been to throw more compute and more memory at larger models. TurboQuant represents the opposite approach — achieving equivalent or better performance with dramatically less hardware. It is an efficiency revolution within the scaling revolution.

This matters beyond the memory market. If the AI industry's trajectory shifts even partially from "bigger models need more chips" to "smarter algorithms need fewer chips," the implications cascade through the entire semiconductor supply chain. Memory manufacturers would face a future where demand growth is slower than projected. GPU makers would see similar efficiency gains erode the assumption that every generation of AI requires proportionally more hardware. Data center operators would find their infrastructure stretching further, reducing the urgency of capital expenditure on new capacity.

Google's publication of TurboQuant as open research, rather than keeping it proprietary, accelerates this dynamic. Any AI company, cloud provider, or enterprise can implement the algorithm. The efficiency gains will not be confined to Google's infrastructure. They will propagate across the industry, gradually bending the demand curve for memory in AI applications.

A CORRECTION, NOT A CURE

TurboQuant is a genuinely important technical breakthrough. It demonstrates that software innovation can meaningfully reduce hardware dependency in AI, and it provides a template for the kind of efficiency improvements that will shape the next phase of AI development. The DDR5 price drops it triggered offered real, if temporary, relief to consumers and businesses suffering under inflated memory costs.

But it is not a cure for the memory shortage. The structural forces driving demand — AI training at ever-larger scales, smartphone memory requirements climbing with each generation, data center buildouts across every major cloud provider — remain intact. The memory market in 2026 is experiencing a correction, not a reversal. Prices may stabilize at lower levels than their pre-TurboQuant peaks, but analysts uniformly expect volatility and cost pressure to persist through the end of the year and into 2027. One algorithm, however brilliant, cannot resolve a shortage built on years of underinvestment in manufacturing capacity and the structural reallocation of silicon toward the most profitable customer in the room: artificial intelligence. ■

SOURCES

- VENTUREBEAT — GOOGLE TURBOQUANT ALGORITHM, 8X SPEED, 50%+ COST REDUCTION, KV CACHE COMPRESSION TO 3 BITS
- CNBC — MEMORY STOCKS FALL, SK HYNIX -6%, SAMSUNG -5%, MICRON AFTER TURBOQUANT
- NOTEBOOKCHECK — DDR5 PRICES RARE DROP AFTER TURBOQUANT ANNOUNCEMENT
- WCCFTECH — DDR5 FIRST PRICE DROP IN MONTHS, KITS DOWN $40
- LIVE MEDIA NEWS — DDR5 CRASH IN CHINA, 30%+ DROP, 100+ YUAN SINGLE-DAY FALLS
- THE REGISTER — TURBOQUANT BIG DEAL BUT WON'T END MEMORY CRUNCH; EVOLUTIONARY NOT REVOLUTIONARY
- TECHRADAR — TURBOQUANT NOT THE RAM CRISIS SAVIOR, PRICES CONTINUE BLEAK
- TRADINGKEY — TURBOQUANT EXPLAINED, POLARQUANT + JOHNSON-LINDENSTRAUSS, TARGETS KV CACHE
- DARKFLASH — TURBOQUANT OPTIMIZATION TARGETS INFERENCE ONLY, HBM UNAFFECTED
- KOTAKU — PC GAMING RAM CHEAPER FOR FIRST TIME IN MONTHS
- DECODING DISCONTINUITY — TURBOQUANT MEMORY STOCK SELL-OFF, PANIC OUTPACED THE PAPER

RAINBOW
INTERNATIONAL DISTRIBUTION

Romania, Bucharest
Calea Șerban Vodă

Smart Vacuum Cleaners:

roborock
WE ARE OFFICIAL DISTRIBUTOR IN ROMANIA

Electric Scooters:

xiaomi
Lenovo
KUGOO
WE ARE OFFICIAL DISTRIBUTOR IN ROMANIA
KuKirin.lt
WE ARE OFFICIAL DISTRIBUTOR IN ROMANIA

Laptops:

hp
Lenovo
ASUS
DELL
acer

Mobile Phones:

xiaomi
realme
HONOR
ZTE
WE ARE OFFICIAL DISTRIBUTOR IN ROMANIA

Smart Lighting:

Govee
WE ARE OFFICIAL DISTRIBUTOR IN ROMANIA

DANI ENULESCU
(PURCHASE MANAGER)
dani@rainbowinternational.eu
+40 731 595 933

GEORGE NITA
(SALES MANAGER)
george@rainbowinternational.eu
+40 723 150 139

ALEX BUDUROI
(SALES MANAGER)
alex@rainbowinternational.eu
+40 720 317 308

BLAGO ANDONOV
(SALES MANAGER)
blago@rainbowinternational.eu
+359 893 392 001

UPtech USA Corp:

Experience-Driven Trading with Integrity at Its Core

UPTECH
The Best Connect

Company Overview

Established in January 2023, **UPtech USA Corp** may be a new entrant in the trading space, but it is firmly built on over a decade of experience. Its founder, Ms. **Pratima Dutt**, has been a trusted name in the industry since 2010, known for her strategic insight, reliability, and dedication to client success.

Specialization & Services

Specializing in the sourcing and distribution of **Apple products,** UPtech USA Corp has quickly carved out a niche by offering competitive pricing, fast turnaround times, and impeccable service. Whether it's the latest **iPhones, iPads, or MacBooks**, the company ensures that clients receive genuine products and dependable support.

Core Values

UPtech's mission is rooted in **transparency, professionalism, and long-term value**.

Vision & Future

With integrity at its core and a sharp focus on premium electronics, **UPtech USA Corp** is poised to become a trusted leader in the global trading arena.

"**In this business, relationships matter just as much as numbers,"** says Ms. Dutt. "**Our clients rely on us not only to deliver competitive pricing, but to do so with honesty, efficiency, and a deep understanding of their needs."**

Dealing in Apple Products

Based in Miami

Contact us

Pratima Dutt
+1 (305) 491-0508

Kavita Sharma
+1 (281) 650-0437

The State of Gaming 2026:

Boom, Bust, and the Battle for What Comes Next

The gaming industry in 2026 is a study in contradictions. Global revenue is projected to reach $205 billion, up 4.6 percent from the previous year. When hardware sales are included, the total climbs closer to $237 billion. Console gaming is the fastest-growing platform segment at 5.5 percent annual growth. GTA 6, the most anticipated game release in a decade, is scheduled for November 19. The Nintendo Switch 2 is delivering the most exciting hardware launch since the original Switch in 2017. By any revenue measure, this is an industry at the peak of its commercial power.

And yet. Over 33,600 game developers have been laid off since the beginning of 2022, with the wave continuing into 2026. In March, Epic Games cut more than 1,000 employees — 23 percent of the company. Playtika eliminated 500 jobs in January, explicitly citing a shift toward AI-dependent resources. The GDC 2026 State of the Game Industry survey found that 28 percent of respondents had been laid off in the past two years. At AAA studios, two-thirds reported that their companies had conducted layoffs. Outsourcing hit record levels. And despite the headline revenue numbers, mobile game downloads have declined 6 to 7 percent year over year since 2021.

How can an industry simultaneously grow its revenue and hollow out its workforce? The answer lies in a structural transformation that is reshaping how games are made, how they are monetized, and who captures the value they create.

THE REVENUE PICTURE: WHERE THE MONEY IS

The $205 billion gaming market breaks down along lines that would be unrecognizable to the industry of a decade ago. Mobile gaming commands the largest share at 52 percent of total revenue, generating approximately $107 billion. Console gaming accounts for roughly $48 billion. PC gaming contributes approximately $40 billion, with Steam maintaining a dominant 74 percent market share in digital PC distribution.

But the revenue composition tells a more revealing story than the platform split. In-game purchases generated $74.4 billion in 2025, up 4.6 percent year over year. Subscription services, including Xbox Game Pass and PlayStation Plus, contributed $11 billion. Cloud gaming revenue reached between $6 billion and $8 billion. The pattern is unmistakable: the industry's growth is driven overwhelmingly by recurring revenue models — microtransactions, battle passes, subscriptions, and live-service content — rather than by the traditional model of selling individual game copies.

Regionally, Asia-Pacific leads with $87.6 billion, accounting for 46 percent of global gaming spend. This regional dominance, driven primarily by mobile gaming in China, Japan, South Korea, and Southeast Asia, explains why every major publisher now designs games with Asian market dynamics as a primary consideration rather than an afterthought.

THE HARDWARE MOMENT: SWITCH 2 AND THE CONSOLE LANDSCAPE

The Nintendo Switch 2, launched in spring 2026, represents the most significant console event of the year. Nintendo's hybrid portable-home console formula proved transformative with the original Switch, selling over 146 million units lifetime. The successor builds on that foundation with substantially more processing power, enabling ports of current-generation games that were previously impossible on Nintendo hardware. The prospect of playing titles like Final Fantasy VII Rebirth and potentially even GTA 6 on a handheld device has generated enormous consumer excitement.

Sony's PlayStation 5 remains the dominant traditional console, bolstered by the PS5 Pro's enhanced specifications for 4K gaming and ray tracing. The console's exclusive library and its PlayStation Plus subscription service, which now offers a tiered model including cloud streaming, continue to drive platform loyalty. Microsoft's Xbox strategy has shifted decisively away from console exclusivity toward a multiplatform approach, with Xbox Game Pass available across Xbox hardware, PC, mobile, and smart TVs. The company's cloud gaming infrastructure, powered by Azure, positions Xbox as a service rather than a device.

This strategic divergence is significant. Nintendo is betting on hardware innovation. Sony is betting on premium exclusive experiences. Microsoft is betting on ubiquitous access. The three approaches can coexist, but they reflect fundamentally different visions of what a gaming platform means in 2026.

GTA 6: THE $2 BILLION QUESTION

No single release in 2026 carries more commercial weight than Grand Theft Auto 6. Scheduled for November 19 on PlayStation 5 and Xbox Series X/S, the game represents Rockstar Games' first mainline Grand Theft Auto entry since GTA V in 2013 — a thirteen-year gap during which GTA V sold over 200 million copies and GTA Online generated billions in ongoing revenue.

Industry analysts project that GTA 6 will generate over $2 billion in revenue within its first month, potentially making it the highest-grossing entertainment launch in history. The game's release is expected to drive a measurable spike in console hardware sales, particularly for PlayStation 5, and to dominate the cultural conversation for months. A PC version has not been confirmed but is widely expected to follow in 2027, following Rockstar's established pattern of timed console exclusivity.

Beyond its commercial impact, GTA 6 is a test case for the industry's economics. The game's development budget is estimated to exceed $2 billion, making it

the most expensive game ever produced. If it succeeds — and few doubt it will — it validates the blockbuster model at a scale never previously attempted. If the live-service component of GTA Online's successor performs as expected, it could generate revenue for another decade. The question is whether any other publisher can afford to compete at this level, or whether GTA 6 represents the ultimate expression of an economic model accessible only to the very largest studios.

THE LAYOFF PARADOX: RECORD REVENUE, RECORD JOB CUTS

The gaming industry's layoff crisis is not a temporary correction. It is a structural realignment. The 33,600 jobs lost since 2022 span every segment of the industry: AAA publishers, indie studios, mobile developers, and platform holders. Epic Games' March 2026 layoff of over 1,000 employees was among the largest single cuts in gaming history. Bloomberg described it as a symptom of deeper industry problems, not an isolated event.

Three forces are driving the contraction simultaneously. First, the post-pandemic normalization: gaming revenue surged during lockdowns, leading to aggressive hiring that proved unsustainable once growth rates returned to pre-pandemic levels. Second, rising development costs: AAA game budgets have escalated to the point where only guaranteed franchises can justify the investment, squeezing out mid-tier projects and the studios that make them. Third, AI-driven automation: the GDC 2026 survey found that 36 percent of game industry professionals are now using generative AI tools, with QA testing, 2D and 3D asset creation, localization, and audio production identified as the roles most affected by automation.

The result is an industry that produces more revenue with fewer people. Live-service games generate ongoing income from existing content through battle passes and cosmetic sales, requiring smaller teams for maintenance than for initial development. AI tools accelerate asset creation, reducing the need for large art departments. Outsourcing to lower-cost regions, which hit record levels at GDC 2026, allows studios to scale down permanent headcount while maintaining production capacity.

MOBILE GAMING: DOWNLOADS DOWN, REVENUE UP

The mobile gaming market presents its own contradiction. Downloads have declined 6 to 7 percent year over year since 2021, yet revenue continues to climb. In-app purchase revenue grew 4 percent annually. Time spent in mobile games is up 8 percent. Sessions are up 12 percent. Fewer people are downloading new games, but those who play are playing longer and spending more.

The genre landscape is shifting. Strategy games, particularly the 4X subgenre, emerged as the only category to gain ground simultaneously in revenue, downloads, and time spent. Casino games became a breakout advertising category. Hypercasual games, which drove the download boom of 2019 to 2022, have faded as a growth engine, replaced by hybrid-casual titles that blend simple mechanics with deeper progression and monetization systems.

Consolidation continues to reshape the mobile landscape. Scopely's $3.5 billion acquisition of Pokémon Go in 2025 expanded its audience to over 500 million monthly players, illustrating the scale at which mobile gaming now operates. The economics increasingly favor large publishers with established franchises and sophisticated user acquisition capabilities, making it progressively harder for independent mobile developers to compete.

THE INDUSTRY AT A CROSSROADS

The gaming industry in 2026 is not in decline. Its revenue is higher than ever, its cultural influence is broader than ever, and its biggest releases — GTA 6, the Switch 2 lineup, the next wave of live-service titles — will generate billions. But beneath the commercial success, the industry is undergoing a transformation that is concentrating value among fewer, larger players while reducing the workforce that creates the products.

The studios that survive and thrive will be those that navigate this tension deliberately: leveraging AI and outsourcing to manage costs without losing the creative talent that makes great games possible, investing in live-service models without abandoning the single-player experiences that build lasting franchises, and finding the balance between blockbuster economics and sustainable development. The numbers have never been bigger. The question is who gets to share in them. ■

The CRONO
ACCESO PRIVADO

Luxury **watches** verified

The Crono was born from a simple conviction: **time is the ultimate luxury.**

Founded in Miami with a vision that extends far beyond geography,

The Crono is more than a watch dealer - it is a private gateway to the world's most important timepieces.

From rare new-old-stock icons to modern grails in platinum, gold, and steel, every piece is hand-selected for significance, condition, and long-term value.

Built on decades of global distribution expertise

and guided by discretion, integrity, and precision,

The Crono serves collectors who understand that a watch is not just worn - it represents legacy, access, and status.

Contact:
DM us for more info:
Instagram: @thecronollc
Email: info@thecronollc.com

The CRONO
ACCESO PRIVADO

THE CURATION
Connecting remarkable people through moments that define a life well-lived

SHENZHEN WINGOCEAN INTL LOGISTICS CO., LTD

Wing Ocean

FREIGHT FORWARDER

LOGISTICS SERVICES

- **CHINA TO EUROPE/USA/THE MIDDLE EAST:**
  AIR TO AIRPORT/DOOR SERVICE (DDU、DDP)
  SEA TO PORT/DOOR SERVICE(DDU、DDP)
  DG GOODS SERVICE (DG BATTERY、DRONE AND SO ON)
- **CHINA TO EUROPE: RAILWAY TO DOOR SERVICE(DDU、DDP)**

- **CONTACT INFORMATION:**
  WHATSAPP/WECHAT：+8613677221280
  EMAIL：oversea11@wingocean.net

We offer IT services

- API SYSTEM INTEGRATION
- WEBSITE DESIGN AND SERVER SETUP
- SYSTEM OPTIMIZATION AND CONSULTING

Call us for more information

+48 512 605 811

HECC
SP. Z O.O.

Contact us:

+48 512 605 811

HANDELOT

B2B wholesale consultants

www.handelot.com