MARCH 2026

£25

THE PAPER FUEL

BUSINESS MAGAZINE

ISSN 2978-8587

INDIA'S AI IMPACT SUMMIT 2026

The World Meets at New Delhi, and AI Will Never Be the Same

THE PAPER FUEL
BUSINESS MAGAZINE

TABLE OF CONTENTS

THE PAPERFUEL LTD
96 RHODESIA ROAD
FAZAKERLEY
LIVERPOOL
UNITED KINGDOM - L 9 9BX

WWW.THEPAPERFUEL.COM

EDITOR'S NOTE

AI innovation meets global economic transformation

As we step further into 2026, two forces continue to redefine the global business landscape: artificial intelligence and economic recalibration. This month, we explore the rapid rise of AI-driven systems — from automated decision-making platforms to the growing influence of AI in content creation and strategic planning. The surge in AI submissions across industries signals not just innovation, but a structural shift in how businesses operate, compete, and scale.

At the same time, the global business economy remains in a delicate phase of adjustment. Shifting trade alliances, evolving monetary policies, and emerging markets are reshaping opportunity maps for investors and entrepreneurs alike. Resilience, adaptability, and informed leadership have never been more critical.

In this edition, we bring you insights from industry leaders, analysts, and innovators who are navigating this transformation firsthand. The future is not waiting — it is being built now.

Here's to informed decisions and bold thinking.

SIBIN CYRIAC
CHIEF EDITOR
THE PAPER FUEL

sibin cyriac

CREATIVE TEAM

MARCH 2026 Edition

EDITOR IN CHIEF
SIBIN CYRIAC

MANAGING EDITOR
PHILIP DEVASIA

SENIOR EDITOR
ADAM FRANK

ASSOCIATE EDITOR
NAYANTHARA P V

ART DIRECTOR
ASWANI M MENON

DESIGN
NEO DIGITAL

DIRECTOR & CEO
BENJITH BENNY

MULTIMEDIA PRODUCER
HAINA K

PHOTOGRAPHY
PETER GRAY

TECH DIRECTOR
KEVIN TOM

IT
TGI TECHNOLOGIES

magazine@thepaperfuel.com
www.thepaperfuel.com

The Paper Fuel
ISSN 2978-8587
Frequency : Monthly
Publisher : The Paper Fuel Ltd
Liverpool, United Kingdom

The Paper Fuel Ltd
96 Rhodesia Road
Fazakerley
Liverpool
United Kingdom – L 9 9BX

INDIA AI IMPACT SUMMIT 2026

The World Meets at New Delhi, and AI Will Never Be the Same

On the morning of February 19, hundreds of delegates stood in long, stalled queues outside Bharat Mandapam, New Delhi's vast convention complex, baking under a winter sun that felt entirely too warm. There was no food, no water, and no clear word on when the gates would open. Prime Minister Narendra Modi was due to inaugurate the main summit, and security had effectively frozen the venue. Bloomberg reported that delegates, some of whom had flown in from three continents, were simply left waiting.

It was, unintentionally, a perfect metaphor for the India AI Impact Summit 2026: enormous ambition, occasional chaos, and underneath it all, something genuinely significant trying to emerge.

By the time the week ended, that something had. Over six days at Bharat Mandapam — the summit ran a day longer than planned, overwhelmed by public interest — the world's largest AI gathering produced a multilateral declaration signed by 89 countries, investment pledges exceeding $250 billion, the launch of India's first homegrown large language models, and a clear message to Washington and Beijing: the future of artificial intelligence belongs to more than two countries.

"AI Is Not a Race Between Two Nations"

The India AI Impact Summit was the fourth in a series that began at Bletchley Park in 2023, continued through Seoul in 2024, and passed through Paris in early 2025. Each summit reflected the anxieties of its moment. Bletchley worried about existential risk. Seoul focused on safety frameworks. Paris pivoted, somewhat awkwardly, toward commerce.

New Delhi went somewhere different.

As the United States and China battle to dominate artificial intelligence, India used the summit to

highlight that there are other pathways to navigate the silicon surge. The theme "Welfare for all, Happiness of all" drew from the Sanskrit philosophical tradition, and it was plastered across banners throughout the capital. Cynics rolled their eyes. But the framing worked, at least rhetorically.

"Long term, it's good for the world that AI is not just viewed as a race between the U.S. and China, and I think that India is right now the player that most confidently says, 'We reject this dynamic,'" said Jakob Mökander, director of policy research at one international think tank.

That rejection carried real diplomatic weight. The summit drew over 20 heads of government, representatives from 118 countries, more than 100 global AI CEOs and CXOs, and over 500,000 participants, making it one of the largest AI gatherings in history. French President Emmanuel Macron and UN Secretary-General António Guterres addressed the opening ceremony. The who's-who of global tech arrived in force: Sundar Pichai, Sam Altman, Dario Amodei, Demis Hassabis, Mukesh Ambani.

It was also critically the first in this global summit series to be hosted by a country from the Global South. That matters more than it sounds.

The $250 Billion Question

The money talked loudly. Reliance Industries pledged $110 billion over seven years for sovereign computing infrastructure, and the Adani Group committed $100 billion by 2035 for renewable energy-powered AI data centres. Google announced an estimated $15 billion investment, including a full-stack AI hub in Visakhapatnam. OpenAI signed a partnership with Tata Consultancy Services, making TCS's data centre business its first such customer. Anthropic announced a partnership with Infosys and opened an office in Bangalore.

The government set an audacious target: attract $200 billion in AI investment over the next two years. Whether that number is achievable is a separate debate. What is clear is that global technology capital is taking India seriously in a way it simply did not five years ago.

Alphabet CEO Sundar Pichai and OpenAI boss Sam Altman both spoke enthusiastically about India's advantages, its enormous talent pool and its large consumer market. "The excitement here has been incredible to watch," Altman told CNBC on the sidelines.

The CEOs also said something more unsettling. Altman claimed the world might be "only a couple of years away from early forms of superintelligence." Demis Hassabis said artificial general intelligence could be achieved within five years, an apparent halving of his projected timeline from the previous year. Amodei, for his part, suggested that advanced AI could produce 25% annual GDP growth for India. He acknowledged that might "sound absurd." It did, but no one in the room asked him to leave.

Made in India: The Models Have Arrived

Beyond the deal announcements, the summit's most consequential development may have been quieter: the emergence of a credible Indian AI technology stack.

Sarvam AI, a Bengaluru-based laboratory, launched a new generation of large language models, including 30-billion and 105-billion parameter models using a mixture-of-experts architecture, as well as text-to-speech, speech-to-text, and vision models. The company also unveiled Kaze, a smart glasses product that Prime Minister Modi tried on at the expo — a rare moment where the theatre of a state summit and the reality of a startup intersected usefully.

The government-backed BharatGen launched Param2, a 17-billion parameter model supporting 22 Indian languages with multimodal capabilities —

designed for government services, education, and healthcare. The vision is a sovereign AI stack: trained in India, hosted in India, serving India's languages and needs.

To power all of this, the government announced plans to add more than 20,000 GPUs to its existing base of 38,000, with a target to cross 100,000 GPUs by the end of 2026. It also unveiled the MANAV framework — Moral, Accountable, National Sovereignty, Accessible, Valid as India's ethical scaffolding for AI development, and announced IndiaAI Mission 2.0, the next phase of its national AI strategy.

The message was deliberate: India does not want to be merely a consumer of AI built elsewhere. It wants to build.

The Delhi Declaration and Its Discontents

The diplomatic centrepiece of the summit was the New Delhi Declaration on AI Impact, endorsed by 89 countries and international organisations, including the United States and China, making it the broadest multilateral consensus on AI to date. It called for democratic diffusion of AI, equitable access to foundational resources, trusted frameworks, and collaborative research.

It also joined India to the Pax Silica initiative, the U.S.-led coalition to secure semiconductor supply chains — a significant geopolitical signal about which side of the great technology divide India is choosing to stand on, even as it presents itself as a neutral bridge between blocs.

But critics were pointed. Mark Brakel, director of policy at the Future of Life Institute, noted that "so many risks, from child safety to national security risks to loss of control, were discussed in the corridors with greater urgency than ever but didn't make it to the official outcome." Others questioned whether the Declaration's voluntary, non-binding language would ever translate into action. Some policymakers described the event as a natural progression from the Paris summit — one that left them feeling more like a trade fair than a diplomatic gathering.

That tension — between India's genuine desire to shape AI governance and its equally genuine desire to attract investment and project economic power — ran through the entire week. An analysis by TechPolicy.Press argued that the summit's structure granted "multinational corporations parity with sovereign governments" through the CEO Roundtable and Leaders' Plenary, while providing no equivalent high-level platform for civil society, labour leaders, or human rights defenders.

The summit also had its embarrassments. A university exhibited a Chinese-manufactured robot dog as an indigenous creation, sparking viral ridicule on social media. IT Secretary S. Krishnan directed the university to vacate its stall, and an apology followed. A Bengaluru entrepreneur alleged that prototypes were stolen from his booth inside the high-security venue — the devices were eventually recovered by Delhi Police. On the last day, opposition youth activists staged a protest inside the hall. India's IT minister had already apologised for "problems" on day one.

None of this is entirely surprising for an event of this scale and political symbolism. But it pointed to a gap that India will need to close: between the ambition it projects and the execution it can currently deliver.

What Comes Next: A World Looking Over the Horizon

The torch now passes to Geneva, and the world prepared for weighty tests and a fuuture where geopolitics and AI coverage

The baton now passes to Geneva, where Switzerland will host the fifth global AI summit in 2027. By then, the world will have had two years to test whether the Delhi Declaration's voluntarism produces anything real, whether India's $250 billion in pledged investment materialises into actual infrastructure, and whether the homegrown models can compete with OpenAI, Google, and the wave of Chinese laboratories advancing rapidly.

The stakes are larger than any single summit. For all the summit's focus on inclusion and development, some delegates wondered whether the Indian government was only posturing — with little mention made of the possibility that India's vast technology workforce might be uniquely vulnerable to dislocation by AI tools that the very CEOs onstage were promoting.

That is the honest paradox at the centre of this moment. AI, as Altman and Amodei described it in New Delhi, is moving faster than any governance framework can follow. Superintelligence in two years. AGI in five. GDP growth curves that"-sound absurd." And yet the summit chose as its tagline a phrase from a Sanskrit scripture about the welfare and happiness of all.

There is something worth holding onto in that framing. The Bletchley-Seoul-Paris-Delhi sequence has shifted the conversation from abstract existential risk to practical, measurable impact — healthcare tools in rural India, multilingual models for 22 Indian languages, AI-driven agriculture applications for smallholder farmers. That shift is real progress. The question is whether the world's governments can move fast enough, and honestly enough, to make that progress mean something before the technology outruns them.

New Delhi made one thing unmistakably clear: the Global South is no longer waiting to be invited to this table. It has built its own.

What it does at that table — over the next twelve months, and the twelve after that — will define not just India's technological future, but the shape of the world that AI builds.

> *"The AI question is no longer just about technology, but the shape of the world that AI builds."*

OPEN AI O1

Talking AI to Thinking AI

SUMMARY

In late 2024, OpenAI quietly shifted direction. Instead of chasing faster, more fluent language models, it focused on building AI that slows down, reasons step by step, and checks its own logic. That shift produced o1, a reasoning-first model designed for complex tasks like coding, mathematics, planning, and decision-making.

Unlike earlier models, o1 prioritises correctness over speed. It uses an internal planning phase to break problems into steps, identify errors, and revise outputs before responding. This makes it far more reliable in real work scenarios, especially in agentic AI systems where models must choose tools, adapt mid-task, and avoid cascading failures. o1 is enabling a new class of AI agents that act like junior digital workers rather than chatbots. The future of AI is not louder or faster, but more deliberate, auditable, and trustworthy.

In late 2024, something subtle but important changed inside OpenAI. The company stopped chasing only bigger, more fluent language models and started asking a harder question. What if AI slowed down, reasoned step by step, and chose its moves like a junior analyst rather than a fast autocomplete engine?

That shift led to o1, OpenAI's reasoning-first model, now shaping how agentic AI systems are being built going into 2026.

This is not about prettier text or smarter chat. It is about AI that plans, checks itself, and acts with intent.

Where o1 Came From

The idea behind o1 was formed from frustration, not ambition. Engineers noticed that even powerful models failed at multi-step tasks like debugging production code, planning a supply chain change, or solving structured math problems. The models sounded confident, but collapsed under complexity.
.
Internal prototypes experimented with a simple but risky move. Let the model think longer before answering. Give it a private reasoning phase. Penalise speed. Reward correctness.

Early results were ugly. Latency spiked. Costs went up. Some answers became worse before they got better. But then came the inflexion point. In internal testing, o1 started outperforming earlier models on tasks that looked less like chat and more like work. Writing test cases. Catching logical gaps. Revising its own output when it spotted errors. One OpenAI researcher described it as the first time the model "felt cautious".

That caution mattered. Instead of guessing, o1 paused, decomposed the problem, and checked assumptions. For teams building AI to operate tools, write code, or make decisions, this was the missing piece.

The human insight was simple but uncomfortable. Intelligence was not about speed. It was about restraint.

What We Can Take from o1

How It Works, Minus the Jargon

o1 is trained to reason in steps before producing an answer. This internal planning phase is invisible to users but critical to outcomes. Think of it as an AI that drafts, reviews, and then responds, instead of blurting out the first thought.

This makes o1 especially strong in agentic systems, where AI must decide what tool to use, what action to take, and when to stop.

Why This Unlocks Agentic AI

Before o1, most AI agents were brittle. One bad step and the whole workflow broke. With reasoning models, agents can recover mid-task, re-plan, and adapt.
That is why companies are now building AI agents for finance ops, customer support escalation, software testing, legal review, and research synthesis. These are not chatbots. They are junior digital workers.

The Business Models It Enables

Reasoning-first AI shifts pricing from usage to outcomes. Founders are moving away from per-seat SaaS toward task-based pricing, per-decision pricing, or revenue-linked automation.

This matters because enterprises are no longer impressed by demos. They want AI that replaces cost centres.

Competitive Edge Over Other Models

Most competitors still optimise for fluency and scale. o1 competes on correctness under pressure. In regulated, high-risk environments like finance, healthcare ops, or infrastructure, this matters more than charm.

That is the real moat. Trust earned through fewer mistakes.

Adoption Friction Is Real

Latency and compute cost remain challenges. Founders must design workflows where reasoning time is acceptable. Not everything needs o1. Smart teams mix fast models for simple tasks and reasoning models for critical steps.

This hybrid architecture is becoming the default playbook.

Regulatory Implications

Reasoning models are easier to audit. Their step-wise structure supports explainability, a big deal for regulators in the EU and US. In 2026, this could make reasoning-first AI the safer choice for compliance-heavy industries.

What Entrepreneurs Should Copy

The biggest lesson is product discipline. OpenAI did not chase benchmarks. It chased failure modes. o1 exists because the team obsessed over where AI breaks in real workflows.

Founders building AI products should do the same. Stop asking what the model can say. Ask what the system must survive.

o1: The Next AI Wave

Most competitors still optimise for fluency and scale. o1 competes on correctness under pressure. In regulated, high-risk environments like finance, healthcare ops, or infrastructure, this matters more than charm.

That is the real moat. Trust earned through fewer mistakes.

What is OpenAI o1?

OpenAI o1 is a new series of advanced AI models introduced by OpenAI, designed to excel in complex reasoning tasks such as mathematics, coding, and scientific problem-solving. Unlike previous models like GPT-4, o1 emphasises critical thinking and multi-step reasoning, enabling it to tackle intricate challenges with greater accuracy and depth.

Key Features

OpenAI o1 models are trained to spend more time reasoning before generating responses, mimicking human-like problem-solving strategies. This approach allows the models to refine their thinking, recognise mistakes, and provide contextually accurate answers. They also incorporate self-fact-checking through a chain-of-thought reasoning process, reducing errors and misinformation.

The o1 models are particularly useful for:

- Developers: Debugging and generating complex code efficiently.
- Researchers: Solving advanced problems in physics, biology, and chemistry.
- Educators and Students: Assisting with complex problem-solving and personalised learning.

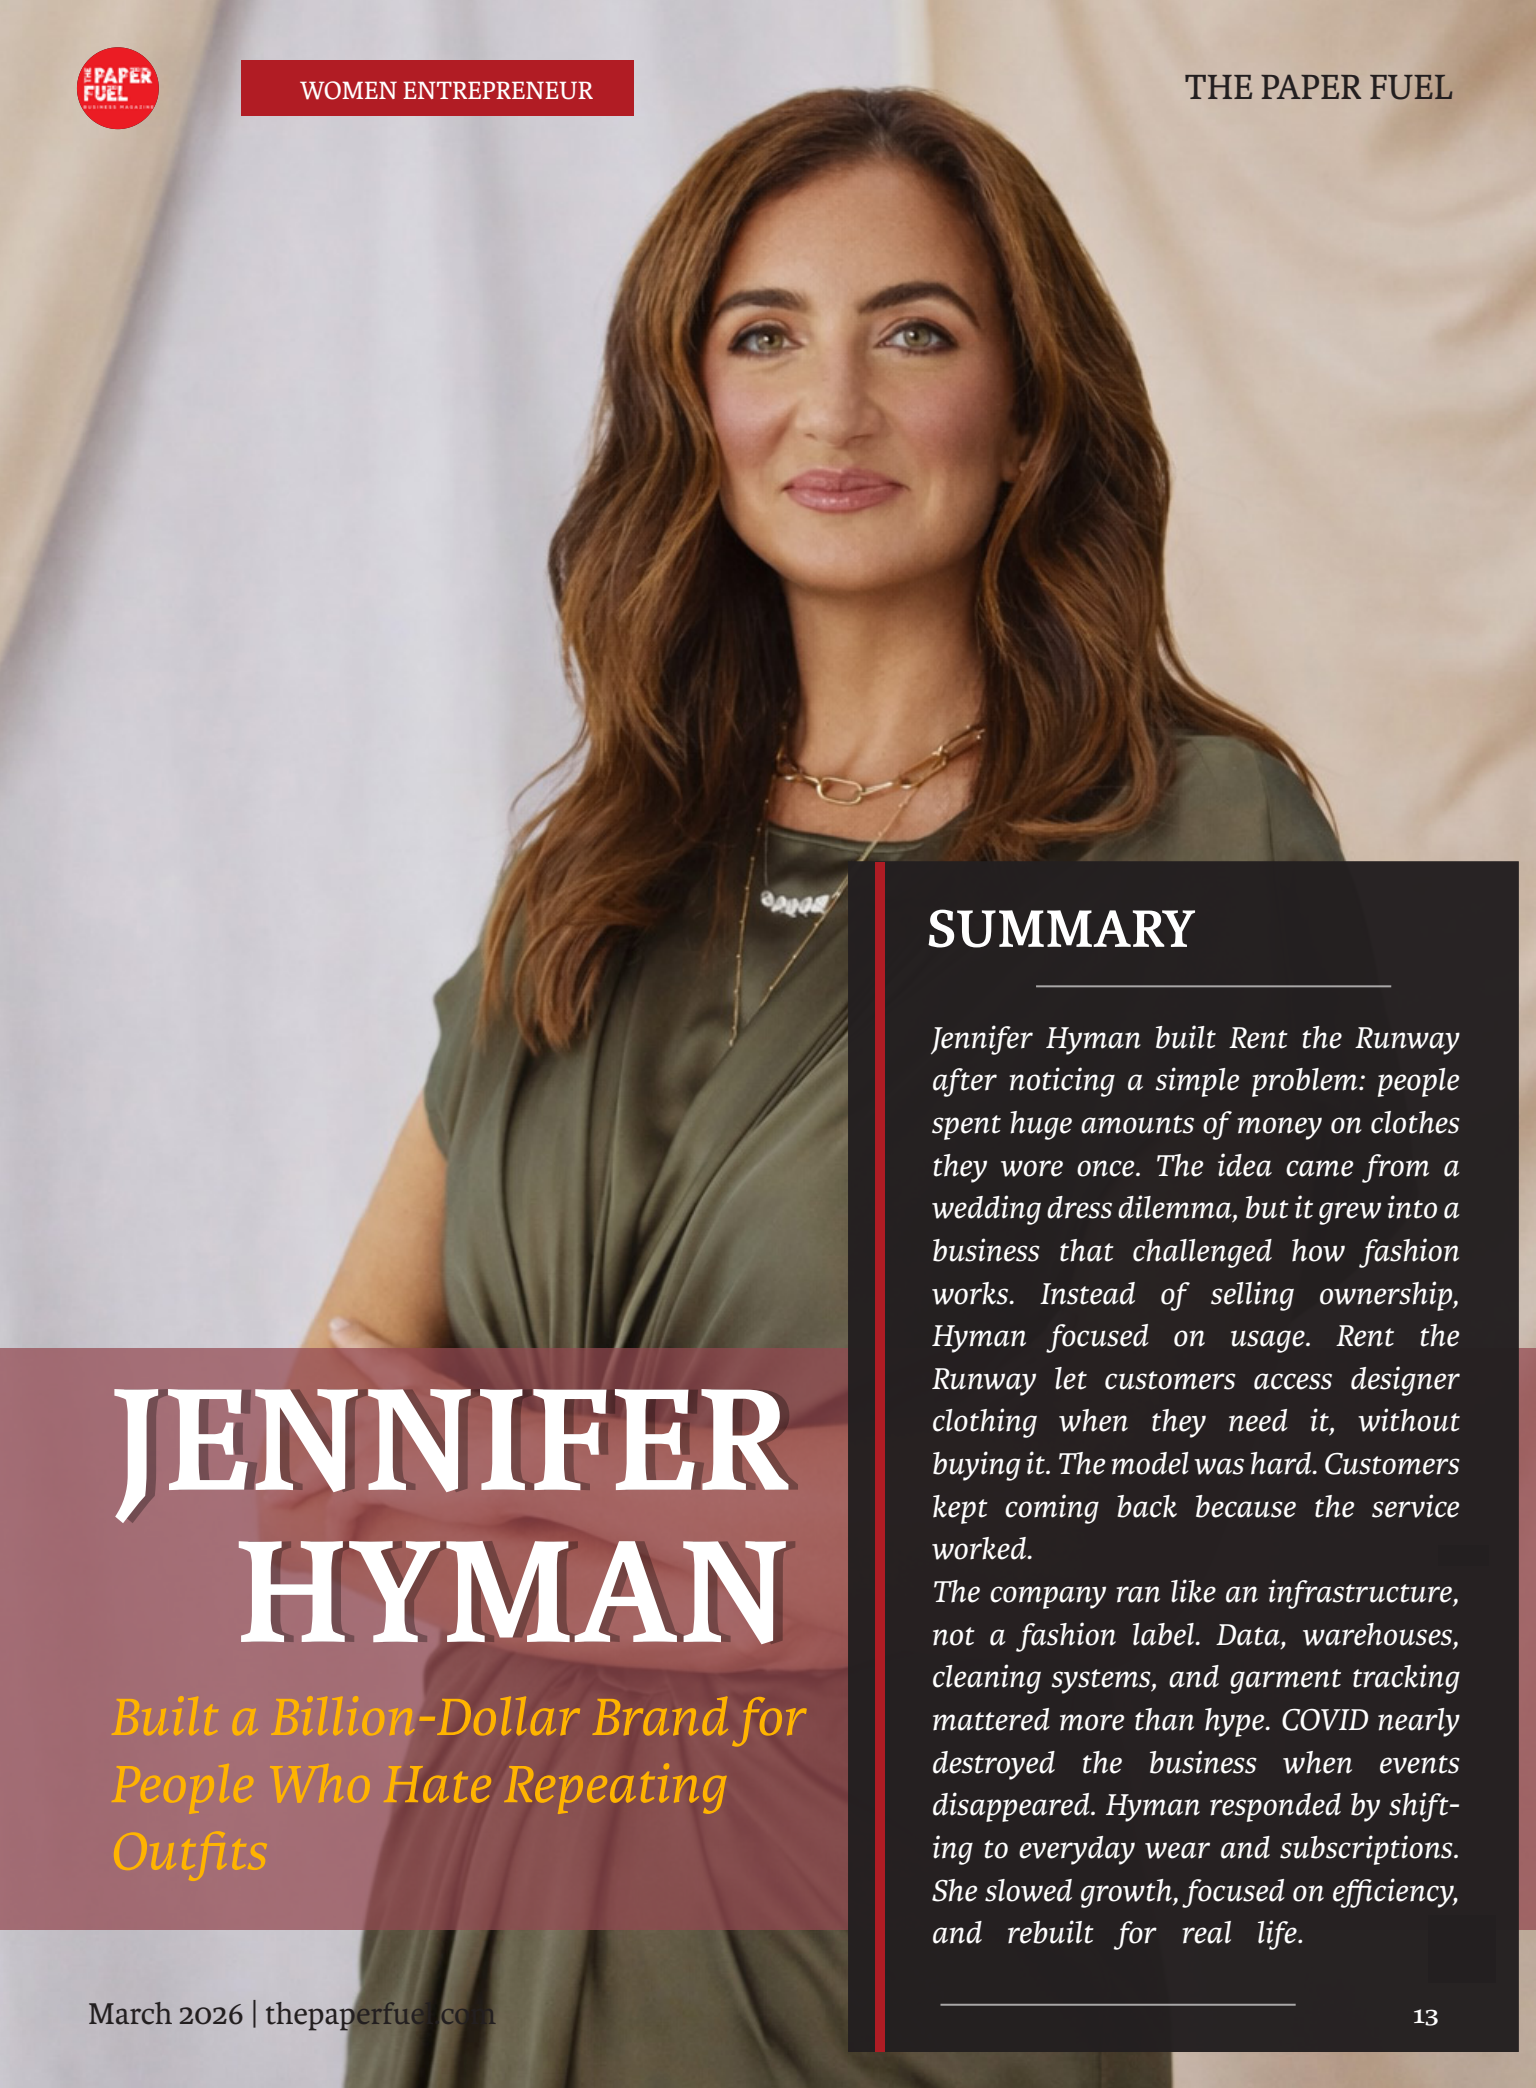

JENNIFER HYMAN

Built a Billion-Dollar Brand for People Who Hate Repeating Outfits

SUMMARY

Jennifer Hyman built Rent the Runway after noticing a simple problem: people spent huge amounts of money on clothes they wore once. The idea came from a wedding dress dilemma, but it grew into a business that challenged how fashion works. Instead of selling ownership, Hyman focused on usage. Rent the Runway let customers access designer clothing when they need it, without buying it. The model was hard. Customers kept coming back because the service worked.

The company ran like an infrastructure, not a fashion label. Data, warehouses, cleaning systems, and garment tracking mattered more than hype. COVID nearly destroyed the business when events disappeared. Hyman responded by shifting to everyday wear and subscriptions. She slowed growth, focused on efficiency, and rebuilt for real life.

From Bedroom to Boardroom

The story does not begin in a boardroom. It begins in a bedroom, with a wedding invite, a price tag that made no sense, and a question that lingered longer than expected. Why did ambition in fashion demand ownership, even when the moment passed in a single night?

For Jennifer Hyman, that question never really left. She grew up around retail, understood consumer desire early, and learned how emotional spending could collide with financial reality. At Harvard Business School, surrounded by spreadsheets and certainty, she fixated on an irrational system that everyone had accepted as normal.

The idea of renting designer fashion sounded clever. In reality, it was inconvenient, capital-intensive, and deeply unfashionable to investors. But Hyman was not chasing elegance. She was chasing usage.

A Closet, a Wedding, and a Crack in the System

Hyman's origin story is well told, but its emotional core still matters. As a Harvard MBA student, she watched her younger sister panic over what to wear to a wedding. The dress she wanted was wildly expensive for something she would wear once. The moment exposed a quiet inefficiency in fashion, aspiration was priced like ownership, even when use was fleeting.

Hyman did not come from fashion royalty. She grew up in Connecticut, worked early retail jobs, and learned firsthand how consumer desire collides with price ceilings. At Harvard Business School, surrounded by future financiers and operators, she sharpened the idea with her co-founder Jennifer Fleiss. Why couldn't women access designer fashion the way companies access cloud software, on demand, without owning the asset?

When Fashion People Said No

Early meetings were rough. Designers worried about dilution. Investors worried about logistics. Consumers worried about hygiene. Almost everyone worried about scale.

What kept Hyman moving was repetition. The first customers came back. Then they came back again. Not for novelty, but for reliability. A dress that arrived clean, on time, and fit as promised built trust faster than any ad campaign.

That was the real beginning. Not the pitch deck. Not the press. The moment when behaviour proved belief.

A Closet Built Like a Factory

From the outside, Rent the Runway looked like a lifestyle brand. Inside, it ran more like industrial infrastructure. Warehouses, cleaning technology, RFID tracking, garment lifecycle data, all of it had to be invented or customised.

While fast fashion chased volume, Hyman chased utilisation. How many lives could one dress live? That single question shaped inventory buys, pricing, and supplier relationships. It also made the business brutally hard.

Building a Company, Not a Campaign

Her fundraising strategy reflected that long view. Rent the Runway raised significant capital, but always to build systems, not just growth. When the company went public, the timing was brutal. Public markets were turning. Loss-making consumer tech was out of favour. The stock struggled.

Instead of chasing hype, Hyman did something unfashionable. She slowed expansion, tightened inventory buys, and refocused on profitability. The era of growth without operational clarity is over. Hyman adjusted early.

The Pandemic That Killed the Party

Then came the moment that should have ended everything. COVID nearly broke the business. Events vanished overnight. Events disappeared. Weddings stopped. Offices closed. Rentals collapsed overnight.

Instead of freezing, Hyman rewired the business. Occasion wear stepped aside. Everyday workwear moved in. Subscriptions were rebuilt around real life, not fantasy life. The company stopped selling dreams and started supporting routines.

Layoffs were painful. Decisions were public. But the reset worked. By the time social life returned, the company had a broader base and a sharper spine.

Culture That Knows Its Limits

This is not a company that worships speed. Engineers argue with merchandisers. Data teams influence fashion taste. Meetings are slower than Silicon Valley would like, but mistakes are fewer. Hyman built a culture that respects constraint. For operators, this matters. Culture is not vibes. It is how a company behaves when margins get thin and optimism runs out.

Why 2026 Belongs to This Model

As we move into 2026, Hyman's relevance sharpens. Consumers are cautious. Sustainability is no longer optional. Owning less feels smart, not sacrificial.

The next phase is about efficiency, not evangelism. Smarter pricing. Better demand forecasting. Designer partnerships that treat rental as discovery, not discounting. Global expansion that respects logistics instead of chasing flags on a map.

Hyman is not trying to win fashion. She is trying to rewire consumption.

Staying Is the Strategy

Jennifer Hyman did not build the loudest company in fashion. She built one that stayed. Through scepticism, public markets, and a pandemic that erased her core use case.

For founders reading this in 2026, the lesson is quietly radical. The future will not reward those who grow the fastest. It will reward those who understand what their business can endure.

Sometimes, the most disruptive thing you can do is refuse to leave the room.

THE BATTERY THAT LETS YOUR GADGETS LIVE LONG

Solid-state batteries are the quiet breakthrough shaping gadgets in 2026

By 2026, the most important upgrade in your phone, laptop, earbuds or wearable will not be the camera, the chip or the screen. It will be invisible. It will sit quietly under the casing and change how long your device lasts, how fast it charges, and how safe it feels in your hand.

Solid-state batteries are moving out of research labs and into real products. Not as hype. As hardware.

SUMMARY

By 2026, the most meaningful upgrade in everyday gadgets will be invisible. Solid-state batteries are moving from labs into real products, quietly reshaping how devices last, charge, and age. Unlike traditional lithium-ion batteries, which rely on flammable liquid electrolytes, solid-state batteries use solid materials that are safer, more stable, and more durable. This shift stems from years of breakthroughs in materials science and advancements in the manufacturing discipline. For consumers, this means thinner devices with longer battery life, faster charging, and less heat anxiety. For founders, batteries become a design and business unlock, enabling refurbished hardware, subscription models, and longer device lifecycles. Solid-state batteries are not flashy, but they quietly change everything built on top of them.

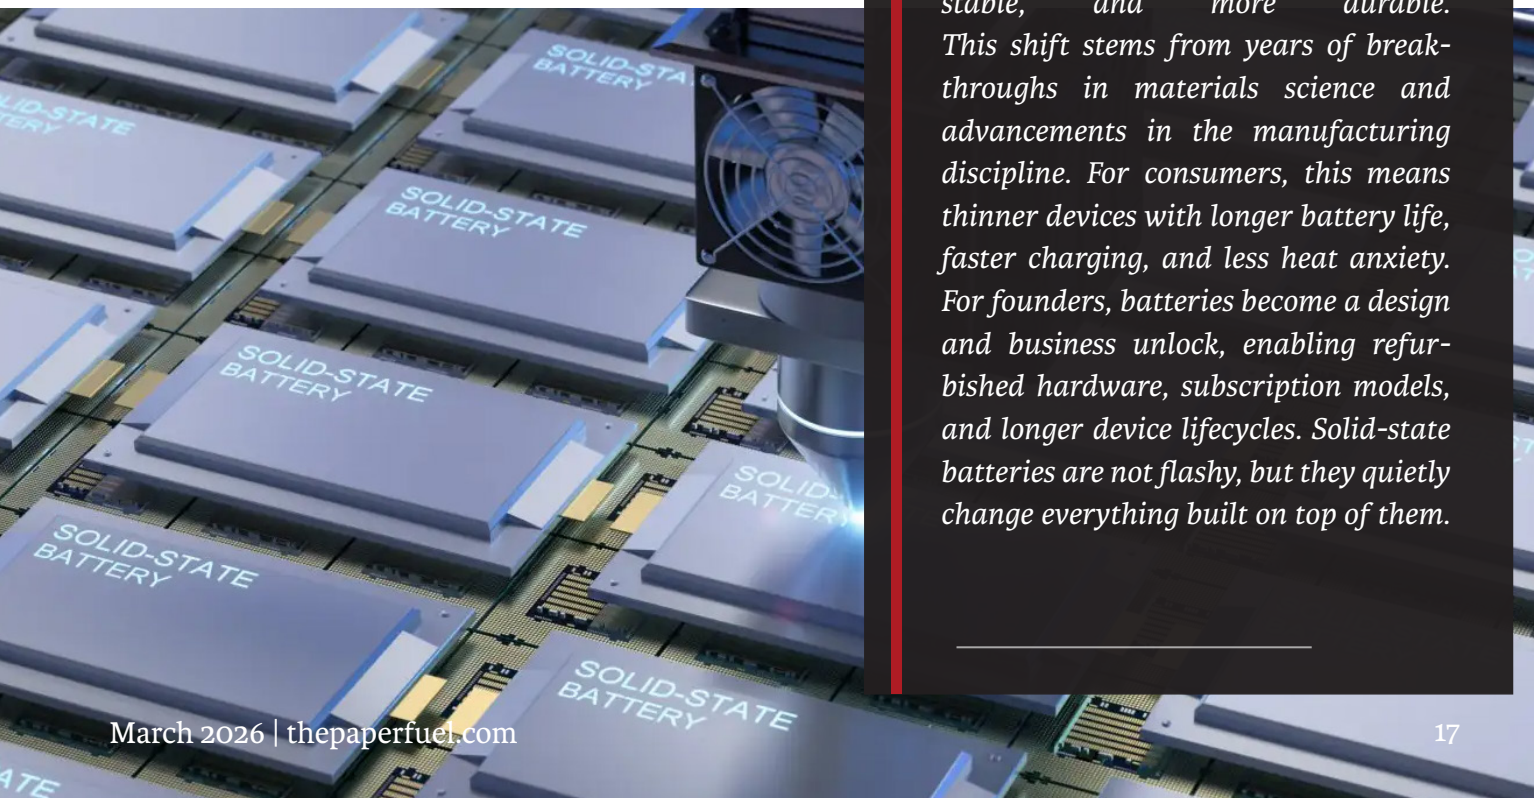

How solid-state batteries broke out of the lab

For decades, lithium-ion batteries ruled everything. Phones, laptops, drones, power banks. They were good enough, cheap enough, and improving just fast enough to avoid replacement. But they came with limits everyone learned to live with, slow degradation, fire risks, bulky safety layers, and charging anxiety baked into daily life. The push for solid-state batteries did not start as a consumer dream. It started as a safety problem.

Engineers working across automotive and electronics companies kept running into the same wall. Liquid electrolytes inside lithium-ion cells were unstable under heat, pressure and fast charging. Every thin phone, every faster charger increases the risk. Fires in phones, laptops and even aeroplanes were warning signs the industry could not ignore. The early solid-state idea was simple on paper. Replace the flammable liquid electrolyte with a solid material. In reality, it was brutal to execute.

Researchers began combining ceramic and polymer electrolytes, creating hybrid solids that were both stable and flexible. At the same time, advances in precision manufacturing made it possible to produce ultra-thin solid layers without microscopic defects. Companies like Samsung and Toyota kept funding long-term programs even when timelines slipped. Inside labs, teams stopped chasing perfect performance and focused on manufacturable performance. Slightly lower energy density was acceptable if the battery could survive real-world abuse.

Prototype devices stopped being fragile demos. They survived drops. They charged fast. They aged predictably. That was the moment solid-state batteries stopped being a science project and started becoming a product strategy.

What should learn from this shift

The tech, explained simply

A solid-state battery replaces the liquid electrolyte with a solid material that conducts ions. This removes the main fire risk, allows tighter internal design, and enables higher energy density over time. Fewer safety layers mean more usable space inside the device.

For product builders, this is not just a battery change. It is a design unlock.

What changes in real gadgets

Expect thinner phones that actually last longer. Wearables that run for days, not hours. Laptops that degrade more slowly over five years instead of two. Faster charging without heat anxiety.

For hardware startups, battery reliability becomes a competitive advantage, not a compromise.

New business models it unlocks

Longer battery life enables subscription hardware models. Devices that are leased, upgraded, and resold. It also strengthens the refurbished electronics market because batteries no longer die first.

For enterprise devices, think industrial handhelds, medical tools, and field equipment; longer life reduces service contracts and downtime.

Capital and scaling realities

Solid-state batteries are capital-intensive. Manufacturing requires precision equipment, clean environments, and long validation cycles. This favours partnerships over solo scale-ups. Founders should think ecosystem-first, not factory-first.

The winning startups are building materials, processes, or integration layers, not entire battery plants.

Adoption hurdles to watch

Cost remains higher than lithium-ion. Early adoption will appear in premium devices first. Consumer trust will also take time. Batteries fail silently, so brands must educate without over-promising.

Competitive dynamics

Big incumbents will dominate volume early. Startups win by specialising, flexible form factors, faster charging chemistries, extreme safety certifications.

Regulatory tailwinds

Governments are tightening safety standards for batteries in transport and consumer electronics. Solid-state tech aligns naturally with these rules, creating quiet regulatory momentum.

Why this matters for 2026 and beyond

Solid-state batteries are not flashy. They do not demo well on stage. But they reshape everything built on top of them.

In 2026, the world is looking forward to devices that last longer, waste less, and feel safer. This innovation answers all three without asking users to change behaviour.

For founders, the lesson is clear. The biggest shifts often come from boring problems solved deeply. Not new features, but new foundations. The battery is finally growing up. And the future of gadgets is growing with it.

MUD JEANS

Concerned about your Jeans after they wear out

SUMMARY

MUD Jeans is a rare fashion brand built to sell less, not more. Founded in 2012 by Bert van Son, a veteran denim executive disillusioned with fashion's waste, the company challenged a core assumption of retail: ownership. Instead of pushing endless new products, MUD introduced "Lease A Jeans", letting customers wear denim for a monthly fee and return it when worn out. Those old jeans become raw material for new ones, turning circular fashion from a marketing claim into a working system. The brand prioritised product quality over ideology, sticking to timeless designs to extend lifespan and protect economics. As regulation tightens and cotton supply becomes unstable, MUD's model looks less radical and more inevitable. Its bet is simple but disruptive: the future of fashion belongs to brands that stay responsible long after checkout.

Most fashion startups try to sell you more clothes. One Dutch brand decided to do the opposite.

In a market obsessed with speed, trends and volume, MUD Jeans built its business around a radical idea: what if customers never truly owned their jeans, and brands stayed responsible for what happens after they are worn out?

As the fashion industry enters a reckoning year in 2026, MUD Jeans stands out as one of the clearest experiments in what the future of consumer goods could look like.

A Founder Tired of Fashion's Waste

MUD Jeans was founded in 2012 by Bert van Son, a former denim executive who had spent decades inside the traditional fashion system. He had seen the waste up close: overproduction, dead stock, landfills filled with barely worn jeans, and supply chains that rewarded speed over responsibility. The personal trigger came when van Son realised that even "sustainable" denim still followed the same broken logic. Brands sold products, washed their hands of responsibility, and moved on.

The early constraint was credibility. Circular fashion sounded good on panels, but almost nobody had made it work commercially. Recycled denim was expensive. Supply was inconsistent. And convincing customers to rethink ownership was risky .

The breakthrough came with the "Lease A Jeans" model. Customers could lease a pair of jeans for a monthly fee. When worn out, they returned them. MUD reused the cotton fibres to make new jeans. Suddenly, circularity was not a slogan. It was a system.

The idea became a real company when customers actually sent their jeans back. .

Turning Circular Theory Into Real Economics

Product first, ideology second

MUD Jeans made a smart early decision. The jeans had to look and feel good before they could be ethical. If the fit failed, the model would collapse. Design stayed classic. Trends were avoided. That reduced returns and extended product life, which directly improved unit economics.

A business model built for reuse

The leasing model gave MUD predictable revenue and direct access to raw materials, used jeans. Returned denim became input, not waste.

This closed-loop system reduced dependence on volatile cotton markets, a major advantage as supply chains faced climate shocks and price spikes.

Margins were thinner early on, but improved as reuse rates increased and recycling technology matured.

Growth without blitzscaling

Unlike fast-fashion brands, MUD did not chase rapid global scale. Expansion was deliberate. Partnerships with conscious retailers and pilot programs in key European markets replaced aggressive store rollouts.

Fundraising followed the same logic. Capital was raised to improve recycling processes and supply chain resilience, not to flood the market with ads

Competing in an Industry Built on Speed

MUD operates in one of the hardest consumer sectors in the world. Denim is crowded, brand loyalty is weak, and price competition is brutal.

Big brands now talk about circularity, but most still rely on partial recycling and voluntary take-back schemes. MUD's advantage is structural. The business is designed for circularity from day one.

A near-fail moment came when recycled fibre quality dipped during early scale-up. Instead of shipping lower-quality products, MUD slowed production, reworked blends, and absorbed short-term losses. That decision protected brand trust and avoided long-term damage.

Hiring stayed lean. Teams were cross-functional, with designers working closely with supply chain and recycling partners. Silos were avoided because circular systems break easily when departments don't talk.

Competing in an Industry Built on Speed

Fashion is under pressure from every direction. Regulators in Europe are pushing extended producer responsibility laws. Consumers are questioning overconsumption. Cotton supply is increasingly unstable.

By 2026, brands will not be judged only on aesthetics or price, but on what happens after purchase. .

MUD Jeans is positioning itself for that world. Its next phase focuses on scaling recycling technology, improving fibre recovery rates, and partnering with other brands to license its circular systems. The long-term ambition is not to dominate denim, but to prove that ownership-free fashion can work at scale.

For founders, the lesson is uncomfortable but powerful. The future does not belong to companies that sell the most units. It belongs to those that take responsibility for the full lifecycle of what they create.

MUD Jeans is not trying to win fashion. It is trying to change the rules. And by 2026, the rest of the industry may be forced to follow.

WHO GIVES A CRAP

The Toilet Paper Company Now a Global Brand

SUMMARY

Who Gives A Crap began in 2012 with a viral stunt, locking a founder in a toilet until enough eco-friendly toilet paper was pre-sold to fund the business. What started as a joke quickly proved a serious insight: even boring products can win if they speak like humans. The founders' breakthrough was simplicity, a direct-to-consumer model, bold packaging, subscriptions without lock-ins, and a promise to donate 50 per cent of profits to global sanitation projects.

The company prioritised cash-flow discipline over hype-driven growth. It slowed expansion when logistics broke, protecting customer loyalty. As 2026 approaches, Who Gives A Crap is evolving into a broader ethical household brand, proving that impact-led businesses can scale without losing credibility or soul.

In 2012, three friends in Australia made a strange bet. They locked one of them in a toilet until they had pre-sold enough eco-friendly toilet paper to fund a startup. It sounded like a joke. It wasn't. That stunt became the origin story of Who Gives A Crap, now one of the most quietly successful consumer brands of the last decade. And as 2026 approaches, it is turning into something bigger than a toilet paper company.

From Bathroom Joke to Real Business

The founders, Simon Griffiths, Danny Alexander and Jehan Ratnatunga, were not FMCG insiders. What they shared was frustration. Toilet paper was wasteful, boring, and deeply unsustainable. Forests were being cut for a product used for seconds and forgotten.

The early constraint was obvious. Toilet paper is a low-interest category with brutal logistics. It is bulky, low-margin, and dominated by giants. No one wakes up excited to buy it.

The first breakthrough came before the product even shipped. The viral pre-sale stunt raised enough money to fund the first production run and proved something important: people were willing to care, if you spoke to them like humans. The habit that shaped the company early was simplicity. Fewer SKUs. Direct-to-consumer only. Bold packaging that people did not hide in cupboards. And a clear promise: 50 per cent of profits would fund sanitation projects around the world.

The company became real when repeat orders started coming in. Customers were not just buying once for the joke. They were subscribing.

Why This Business Works When Others Struggle

Turning a boring category into a brand

Who Gives A Crap understood early that sustainability alone does not sell. Tone does. The brand used humour, colour and honesty to stand out in a category obsessed with discounts and bulk pricing.

That emotional connection allowed premium pricing in a commodity market. Margins were protected not by cost leadership, but by loyalty.

Subscription done the right way

The product's natural replenishment cycle made it perfect for subscriptions. But unlike many DTC brands, Who Gives A Crap did not force lock-ins. Customers could pause, change frequency, or reorder manually.

That flexibility built trust and lowered churn, even if it slowed short-term growth.

Capital discipline over hype

While many consumer brands chased venture capital aggressively, Who Gives A Crap grew steadily, funding expansion largely through cash flow. When it did raise external capital later, it did so to support international expansion, not to fix broken economics.

That discipline mattered during capital tightening. The company never had to panic-cut marketing or compromise on product quality.

Navigating Today's Pressures

The brand sits at the centre of multiple global shifts

First, consumer behaviour. Buyers increasingly want ethical products, but they want them to feel normal, not worthy. Who Gives A Crap delivers impact without moral lectures.

Second, supply chain pressure. Recycled paper sourcing is volatile. The company invested early in long-term supplier relationships and diversified sourcing to reduce shocks during freight and pulp price spikes.

Third, rising competition. Big FMCG players are launching recycled lines. Who Gives A Crap's defence is brand depth, community, and operational learning built over years.

A near-fail moment came during rapid international expansion when shipping costs ballooned and delivery times slipped. Instead of pushing growth harder, the company slowed expansion, fixed logistics, and reset expectations. That pause protected customer trust.

2026: Beyond Toilet Paper

As 2026 approaches, Who Gives A Crap is quietly changing shape. It is expanding into broader household essentials, tissues, paper towels, and bathroom-adjacent products. But the real shift is strategic.

The company is positioning itself as a default ethical household brand. One product in the bathroom becomes a relationship across the home.

For founders, the lesson is sharp. You do not need a revolutionary product to build a meaningful business. You need clarity, patience, and the courage to make something boring feel human. By 2026, Who Gives A Crap may no longer be known just for toilet paper. It may be known for proving that impact-driven businesses can scale without losing their soul.

WILD

Convinced People to Refill Deodorant

In 2020, when most consumer startups were freezing hiring and cutting spend, a small UK brand was betting on something unfashionable at the time: patience. While the world was locked indoors, Wild launched a refillable deodorant with no retail presence, no celebrity backing, and no guarantee that customers would change a deeply ingrained daily habit.

Five years later, Wild is one of the most closely watched consumer brands in Europe. And heading into 2026, it offers a sharp case study in how the next generation of consumer companies will be built.

SUMMARY

Launched in 2020, Wild built a serious consumer business by betting on patience over hype. Instead of chasing fast growth, the UK startup focused on changing one small but sticky habit: disposable deodorant. Its insight was simple. Sustainability would only work if the product felt premium. So Wild treated packaging as the product, built a refillable aluminium case people wanted to keep, and made refills easy, flexible, and commitment-light.

That approach unlocked strong repeat behaviour. Customers came back for refills, not discounts. Wild also moved carefully from DTC into selective retail, protecting margins and customer data. As refillable personal care products become increasingly crowded, Wild's edge lies in brand trust and operational depth. Its growth came from changing habits, not shouting louder.

A Small Habit, A Big Problem

Wild was founded by Charlie Bowes-Lyon and Freddy Ward, two friends who kept circling the same frustration. Bathroom shelves were full of plastic. Deodorant, used daily and discarded often, was a quiet environmental disaster hiding in plain sight.

The idea was simple but uncomfortable: could you convince people to refill deodorant?

Early constraints were real. Refillable systems were expensive to design, harder to explain, and slower to scale than single-use plastic. Early suppliers pushed back. Unit costs looked ugly. And there was no existing playbook for subscription-led personal care in the UK at that scale.

The breakthrough came from focusing obsessively on one thing: make the product feel premium first, sustainable second. The aluminium case felt solid. The scents were modern. The refill experience was frictionless. Sustainability was present, but not preachy.

The company became real when customers came back. Not for discounts. For refills.

Why Wild Worked When Many Didn't

The product was the marketing

Wild's most important decision was treating packaging as a product, not a cost. The refillable case became the brand. Customers displayed it. Talked about it. Gifted it.

That choice allowed Wild to rely less on paid ads and more on word-of-mouth at a time when CACs were exploding across DTC.

Subscription without fatigue

Wild avoided the common subscription trap. Customers were not locked in. Refills were flexible. Skips were easy. That trust-first approach improved retention and reduced churn, even if it slowed short-term revenue recognition.

For founders, this is the lesson: control beats coercion.

Capital discipline in a tight market

While many consumer brands rose aggressively in 2021, Wild stayed cautious. It raised enough to build supply chain resilience and expand SKUs, but avoided hiring for growth at all costs.

That discipline mattered in 2023–24 when freight costs spiked, and consumer demand softened. Wild could absorb shocks without slashing quality or brand spend.

Navigating the New Consumer Reality

Wild sits at the intersection of several global shifts that need to be understood.

First, consumer behaviour has changed. Shoppers want better products, but not lectures. Wild sells aspiration, not guilt.

Second, supply chains are no longer invisible. Wild invested early in supplier relationships and inventory planning, which helped avoid stockouts during logistics disruptions.

Third, retail is back. After proving the model DTC, Wild moved selectively into physical retail. This hybrid GTM gave the brand reach without losing customer data or margin control.

Competition Is Rising, So Are the Stakes

Refillable personal care is no longer niche. Big FMCG players are experimenting. Copycats exist. Wild's defence is brand depth and operational learning.

Hiring remains lean. Product launches are spaced out. The team prefers fewer SKUs that move fast over wide catalogues that confuse customers.

A near-fail moment came when an early overseas expansion strained fulfilment timelines. Instead of pushing harder, Wild pulled back, fixed logistics, and relaunched later with stronger partners. That pause likely saved the brand from reputational damage.

What 2026 Will Look Like for Wild

By 2026, Wild is expected to move beyond deodorant. Body care, hair care, and household refills are all in play. But the bigger shift is strategic.

Wild is positioning itself not as a product company, but as a system company. One refillable habit leads to another. One trusted product opens the door to many.
.
For founders, the lesson is clear. The next wave of consumer brands will not win by being louder. They will win by being repeatable, resilient, and quietly essential.

Wild did not change the world overnight. It changed a habit. And in 2026, that may be the most powerful kind of growth there is.

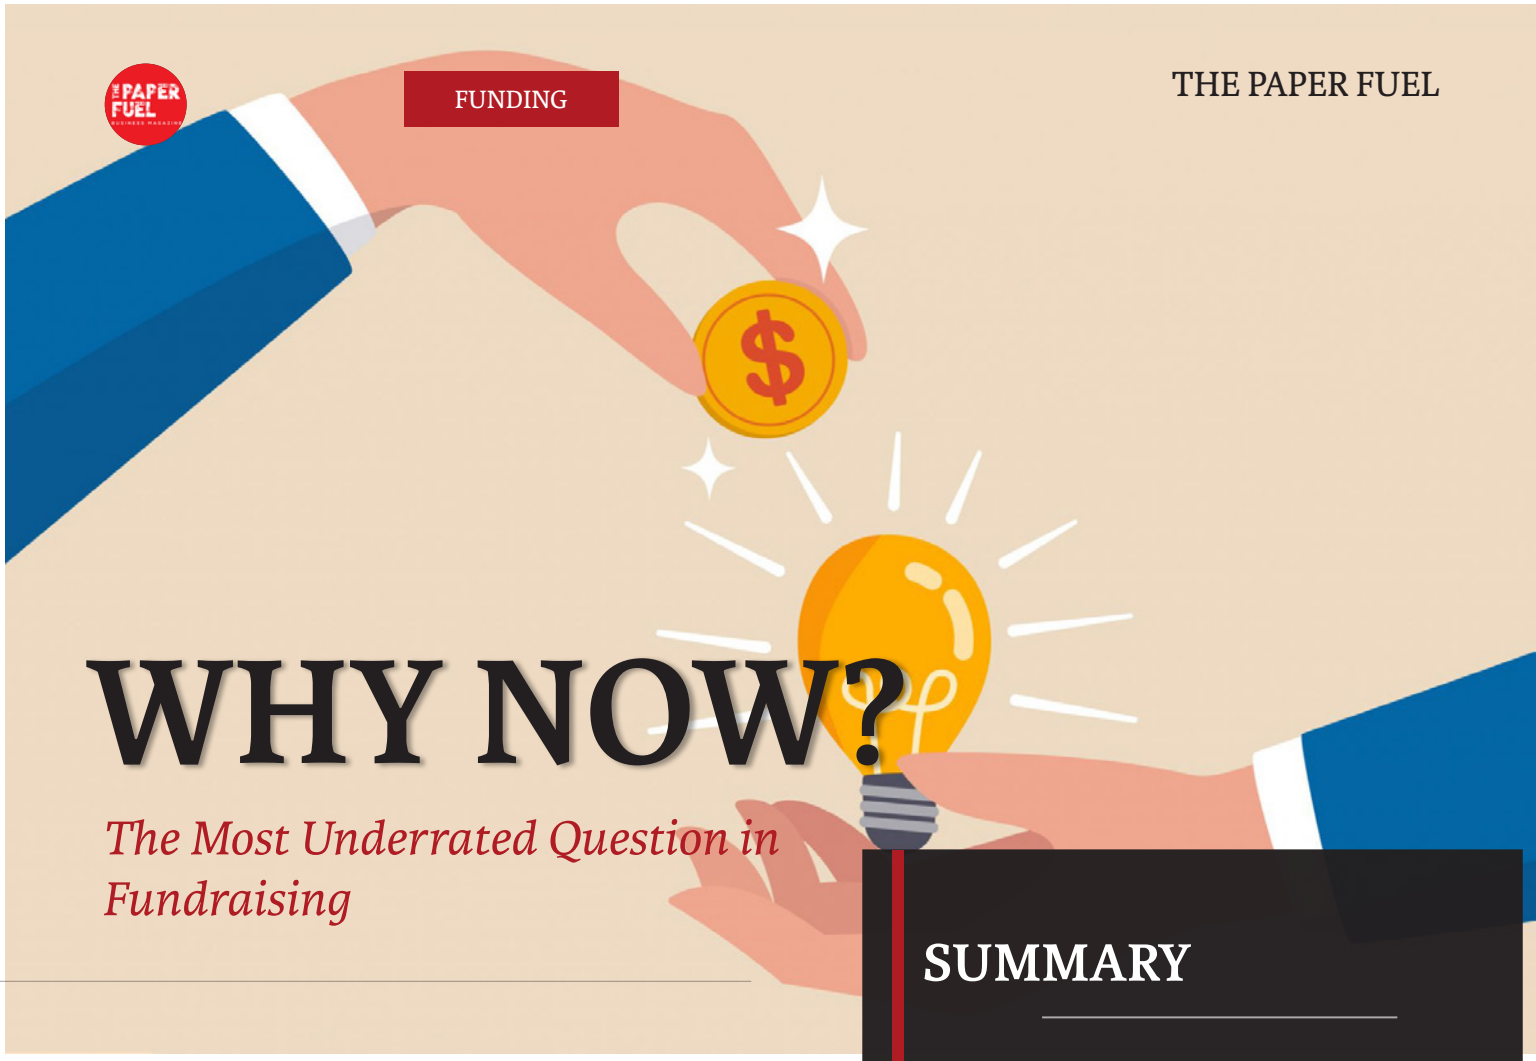

WHY NOW?

The Most Underrated Question in Fundraising

SUMMARY

Timing is one of the most overlooked factors in fundraising. Investors are not just backing ideas or founders, they are backing moments. The key question is not why the idea exists, but why it must exist now. A strong "why now" explains what has changed in the market, whether it is technology costs, regulation, customer behaviour, or broken legacy systems. Trends alone are not enough. Timing sits in a narrow window where conditions align and urgency is real. Customer behaviour is the strongest proof that the market is ready. A credible timing story also shows founder insight and advantage. When founders clearly explain why their startup matters right now, fundraising becomes easier and more focused.

Founders often walk into investor meetings armed with vision, ambition, and beautifully built decks. What many still miss is the single question that quietly decides the outcome of the room. Why now?

Not why this idea. Not why you. Why this company, at this exact moment in the market? In a capital environment where investors have more options and less patience, timing is no longer a footnote. It is the thesis.

The Market Doesn't Fund Ideas. It Funds Moments

Every great company looks obvious in hindsight. In real time, what makes it investable is context.

Airbnb was not just about spare rooms. It arrived when travel was becoming social, payments were trusted online, and a recession pushed people to monetise assets. Stripe did not invent payments. It arrived when

developers were finally building internet-first businesses and needed infrastructure, not banks.

Your startup exists inside a moment. Investors want to understand that moment clearly. If your pitch works in any year, it will excite no one.

What Changed? Start There

A strong "why now" always begins with change.

Something in the market shifted, and that shift created urgency. It could be obvious or subtle, but it must be real. Common triggers include:

- A technology breakthrough that reduced cost or complexity
- A regulatory change that opened or forced a market
- New customer behaviour driven by culture, demography, or habit
- Distribution channels that suddenly scale faster or cheaper
- Legacy players are becoming slow, expensive, or misaligned

The key is specificity. Saying "AI adoption is growing" is weak. Explaining how falling inference costs suddenly make your product viable at scale is strong.

Timing Is a Knife Edge, Not a Trend Slide

Founders often confuse trends with timing. Trends describe direction. Timing explains inevitability.

If you are too early, customers are not ready, and the infrastructure is missing. Too late, and the market is crowded with better-funded players. The "why now" sits in the narrow window where conditions align and speed matters. Investors listen carefully for this. They want to know whether you are riding a wave or trying to create one from scratch.

Customers Are the Loudest Proof of "Now"

Nothing validates timing like customer behaviour.

Are buyers actively searching for solutions? Are budgets shifting? Are purchase cycles shortening? Are customers hacking together workarounds because nothing fits their needs? These signals matter more than market size charts. When customers are already moving,

investors know the market is awake. If your early users say, “We’ve been waiting for this,” you are on solid ground.

Why You, Why Now, Why This Order

A credible “why now” also explains founder advantage.

Why are you seeing this shift before others? What exposure, pain point, or unfair insight puts you ahead? Timing without insight feels lucky. Insight without timing feels academic.

The strongest stories combine both. You noticed the change early, understood its second-order effects, and built for where the market is heading, not where it was.

Don’t Confuse Vision With Readiness

Many founders oversell the future and undersell the present.

Investors are not funding your end state. They are funding the next 18 months. The “why now” must connect directly to what you are building today, not just what it could become someday. Show how current market conditions make your next milestone achievable now, not eventually.

Stress-Test Your Timing

Before pitching, pressure-test your own story. Ask yourself:

- What would have made this fail two years ago?
- What makes it harder to ignore today?
- What will make it crowded in two years?

If you cannot answer these cleanly, neither can your investors.

HEMANT TANEJA

MANAGING PARTNER AND CEO, GENERAL CATALYST

- **Background and role:** Long-time operator turned investor, now CEO and Managing Partner at General Catalyst, one of the most influential global growth investors today.
- **Thesis and focus:** AI-first businesses, healthcare transformation, fintech infrastructure, enterprise software, climate-adjacent systems, and category-defining platforms.

- **Net worth and money trail:** Estimated $500M+ largely from carried interest at General Catalyst, early-stage equity, and long-term fund economics. Notable exposure includes Stripe and HubSpot.
- **Cheque size and stage:** $1M to $50M+, spanning Series A to late-stage growth, with conviction-led follow-ons.
- **Signature portfolio:** Stripe, HubSpot, Canva, Gusto, Airbnb, Snap, Livongo.

What he looks for in founders

Taneja backs founders who think in systems, not features. He prefers leaders who can explain how an industry actually works under the hood and how they plan to rewire it. Clarity of thought matters more than charisma. He expects founders to understand incentives, regulation, distribution, and second-order effects early. He leans toward founders who want to build enduring companies, not fast exits.

Decision-making style and risk appetite

High-conviction, high-context. He takes big bets when the market is structurally broken and technology creates a real opening. Comfortable with long timelines and complexity, especially in regulated or slow-moving sectors like healthcare and financial services. He avoids incremental SaaS clones and shallow arbitrage plays.

How he adds value beyond capital

General Catalyst under Taneja operates like a platform, not just a fund. He brings operating partners, policy expertise, GTM support, healthcare and fintech networks, and board-level guidance on scaling responsibly. He is deeply involved in shaping strategy, governance, and leadership transitions as companies grow.

Red flags and why he passes

He walks away from founders who chase hype cycles without understanding real-world constraints. Weak ethics, fuzzy data integrity, and "AI-washing" are hard no's. He also avoids founders who want capital without accountability or who underestimate execution complexity in regulated markets.

Reputation in the ecosystem

Seen as one of the most thoughtful long-term investors globally. Founders respect him for intellectual honesty and staying power. He is known for backing companies through difficult phases rather than forcing early exits. Policy makers and large enterprises also take his calls, which matters for scale-heavy businesses.

What should know before pitching

Do not pitch surface-level metrics. Come with a clear map of the ecosystem you are entering. Be ready to discuss second-order risks, regulation, and how AI or technology actually changes outcomes, not just costs. He values preparation, intellectual depth, and founder maturity.

What to expect next

Taneja is pushing General Catalyst deeper into "responsible scale." Expect more capital flowing into AI-native companies that touch real-world systems like healthcare delivery, financial infrastructure, public services, and climate resilience. He is likely to back fewer but bigger bets, with tighter integration between capital, policy, and operations. By 2026, his influence will be strongest in companies that blur the line between technology, institutions, and society.

If you are building for the long haul in a hard, consequential market, this is one of the most aligned investors you can have on your cap table.

SUBSCRIBE NOW

INSIDE THE STRATEGIES, RISKS AND DECISIONS
DRIVING INDUSTRY LEADERS

www.thepaperfuel.com

Monthly £25 | Instant access | Latest editions in your whatsapp

AMAN RESORTS

From a Coconut Plantation to Luxury Holiday Home

SUMMARY

Aman Resorts began in 1985 as a quiet gamble. Adrian Zecha, walking along Phuket's Pansea Beach, chose to build a tiny, ultra-private resort instead of a holiday home. Banks refused to fund it, so he self-financed the idea. When Amanpuri opened in 1988, at prices far above the market.

Under Vladislav Doronin, Aman expanded carefully into cities, real estate, private clubs, and now a luxury yacht, without diluting its core promise. Design stayed site-specific, service remained invisible, and scarcity became strategy. Today, Aman earns heavily from branded residences, loyal repeat guests, and gated access. With the launch of Janu and deeper bets in the Middle East, Aman is testing a rare question for luxury in 2026: can a brand built on smallness grow without losing its soul?

In 1985, Aman Resorts began almost by accident. Adrian Zecha was walking along Pansea Beach in Phuket, looking for land to build a holiday home. He found a coconut plantation with a perfect view. Instead of building one house, he imagined something radical for its time: a tiny resort built for privacy, calm, and silence.

Banks refused to fund it. The math did not work. So Zecha and three friends paid for it themselves. That decision reshaped global luxury.

One Beach, One Risk, One New Idea

When Amanpuri opened in 1988, it charged five times more than nearby hotels. Guests still came. There was no front desk drama, no loud lobby, no scripted service. Rooms felt like private homes. Staff knew your name. Silence was the main feature. People came not for what Aman offered, but for what it removed.

Within four years, Aman opened properties in Indonesia, Bora Bora, and the French Alps. A new kind of traveller appeared. Guests who planned entire trips only around Aman resorts. They were called Amanjunkies.

Growth, Drama, and Ownership Chaos

The early success did not mean stability. In 1992, Zecha sold control of the company, keeping only a minority stake. A few years later, a major investor takeover led to lawsuits and Zecha's exit. He returned briefly, then left again. Ownership changed hands multiple times. By 2007, India's DLF bought Aman. In 2014, Russian developer Vladislav Doronin acquired it. Zecha stepped away for good. Through all this, something strange happened. The brand did not break. The idea stayed intact

A Superfan Takes Control

Doronin had stayed at Aman since the early 1990s. He understood the obsession.

As CEO, his promise was simple: grow without ruining what makes Aman special. He moved the headquarters to Switzerland. He opened Aman Tokyo in 2014, proving Aman could work in cities, not just beaches. In 2022, Aman New York opened in the Crown Building. Critics mocked the $200,000 club membership. Rooms are still sold out months ahead.

Design Is the Product

Aman never copies itself.

Architect Ed Tuttle created the original Aman design language: minimal, local, deeply tied to place. Others followed, but the rule stayed the same. Every property must feel like it could exist only there.

Amanjiwo in Java mirrors ancient temples. Amangiri in Utah looks carved from desert rock. No templates. No repetition.

That restraint is expensive. It is also the brand.

The Numbers Behind the Calm

By 2025, Aman will operate 36 properties across 20 countries.

Average room rates hover around $3,500 a night. Some properties cross $5,000. Occupancy stays high. Nearly 70 per cent of guests return. Staff often outnumber guests six to one.

There is still no formal check-in desk. That detail matters.

Luxury Is Changing, So Aman Is Too

Between 2019 and 2023, luxury prices doubled. Millions of customers dropped out. Growth shifted away from the US and Europe and toward places like India and the Middle East.

Aman's answer was not discounts. It was diversification.

Not cheaper. Broader.

A Floating Aman: The Yacht Play

In 2026, bookings open for Amangati, Aman's first yacht.

This is not a cruise ship. It is a floating Aman with just 47 suites, massive terraces, a spa, four restaurants, and a one-to-one crew ratio. It is designed to dock directly in Venice, where most luxury yachts cannot.

Backed by Saudi investors, Amangati will target private charters during events like Cannes and Monaco. Prices will be extreme. That is the point. Aman believes loyalty built on land will follow at sea.

Real Estate Is the Real Engine

Here is the most important number: nearly half of Aman's revenue comes from branded residences.

Apartments in Miami and Beverly Hills sell at huge premiums and often sell out before completion. Owners then live inside the Aman ecosystem, paying for services, dining, and experiences.

Aman earns upfront sales and long-term fees.

Owners become walking advertisements. This is luxury as an asset, not just an experience.

Janu: A Younger, Louder Sibling

Aman knows it cannot stretch forever.

So it created Janu, a separate brand aimed at younger luxury travellers. More rooms. Lower prices. Social spaces. Music. Energy.

Janu Tokyo opened in 2024. More are coming. The idea is simple. Start guests at Janu in their 30s. Graduate them to Aman later. Keep the core brand untouched.

It is risky. But deliberate.

Following the Money East

Saudi Arabia, the UAE, and nearby markets are building luxury tourism from scratch. They need global brands. Aman fits perfectly.

Billions in funding are flowing in. New properties are planned across the region. The yacht project is part of this strategy.

For Aman, this is not a side bet. It is the future.

Paying Just for Access

Aman's private club costs $200,000 to join.

Members get priority bookings, private events, and guaranteed access during peak seasons. Renewal rates are above 90 percent. Members now drive most room nights.

This is not a loyalty program. It is a gate. And people are paying happily.

Sustainability, Without Guilt

Aman talks about solar power, local sourcing, and cultural restoration. It does the work quietly. Guests still arrive by private jet. Pools are heated. Luxury remains uncompromised.

The deal is unspoken. Aman will do better behind the scenes. Guests will not be asked to sacrifice comfort.

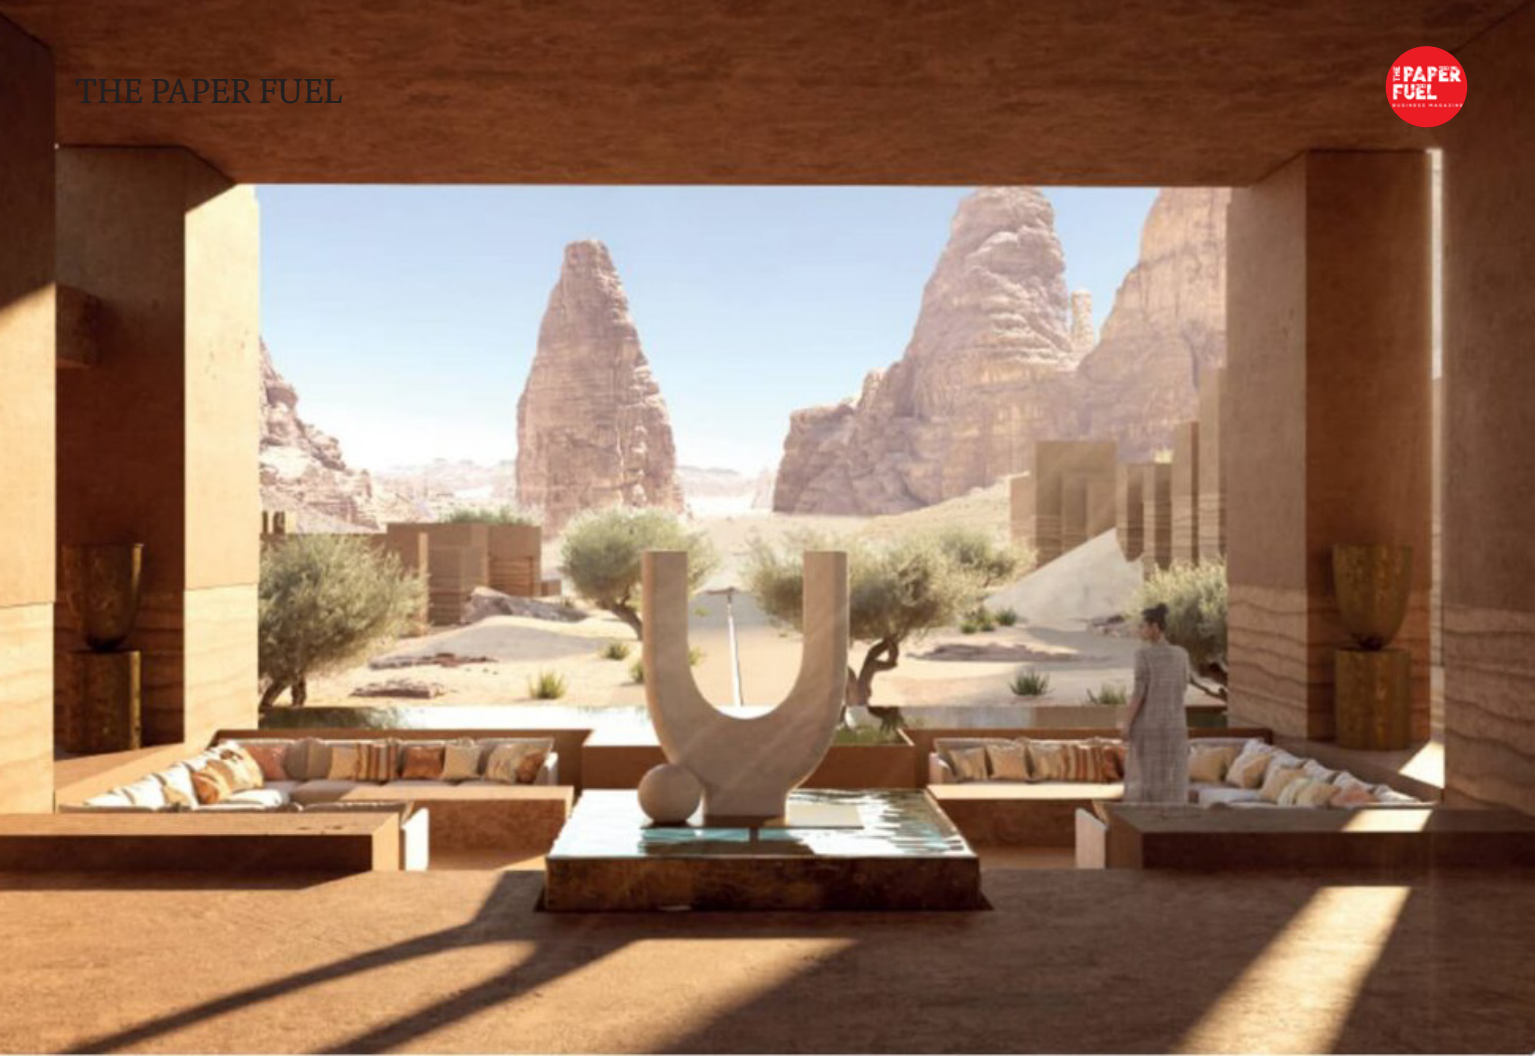

Technology That Stays Invisible

Most bookings happen directly. The website is clean. There is no noisy app.

Behind the scenes, data tracks guest preferences across the world. Pillows, food, spa choices follow you. AI exists, but it never announces itself.

Technology serves the illusion of effortlessness.

What Can Learn

Aman proves three things.

First, strong brands survive ownership chaos if the promise stays clear.

Second, expansion works when the brand travels naturally into new spaces, not when it stretches thin.

Third, scarcity is a strategy. Fewer rooms create more value. Scale kills mystique.

The Big Question of 2026

By the end of 2026, Aman will have new resorts, a yacht nearing launch, major real estate sales, and a growing younger brand.

The entire luxury industry is watching. Can a brand built on smallness grow without losing its soul?

Once, it was just a coconut plantation with a view. Now it is a global experiment in how far quiet luxury can go without breaking.

The answer will shape luxury's next decade.

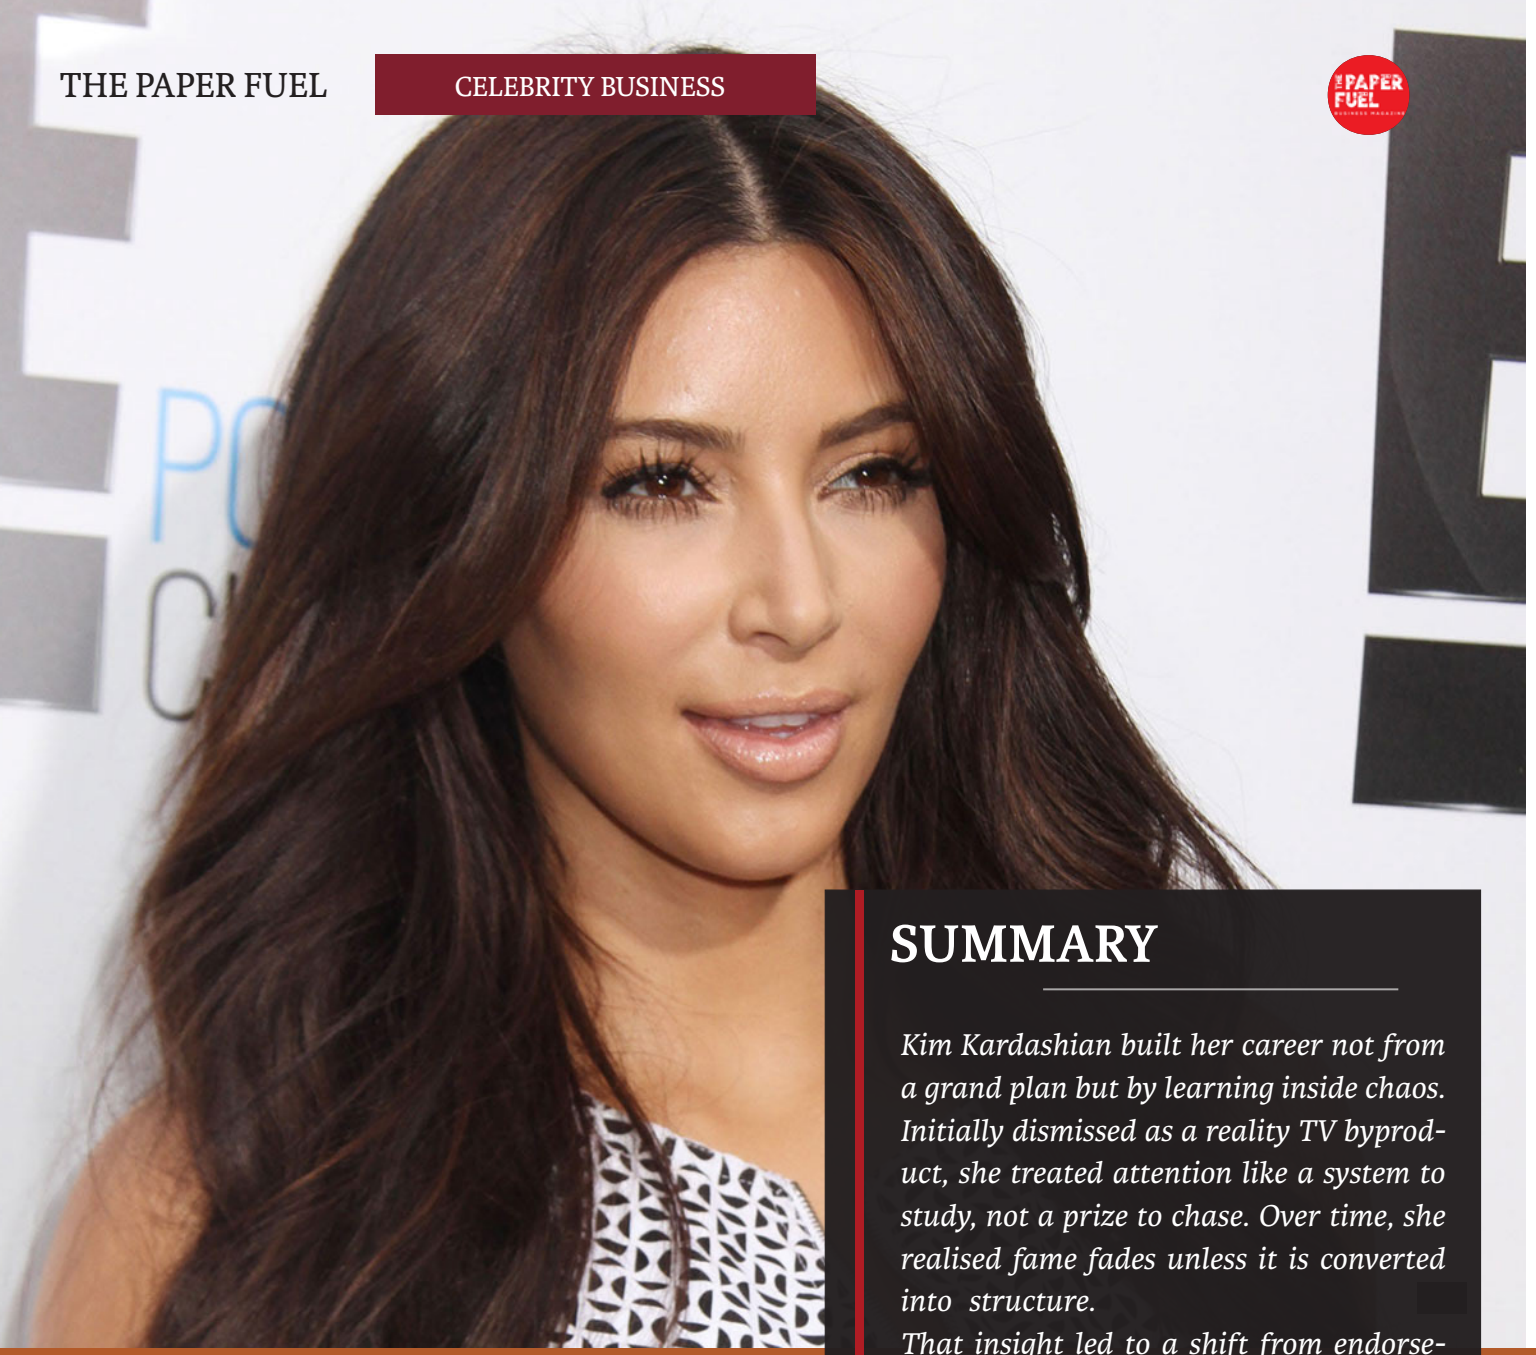

KIM KARDASHIAN

Made Solutions For Every Body

SUMMARY

Kim Kardashian built her career not from a grand plan but by learning inside chaos. Initially dismissed as a reality TV byproduct, she treated attention like a system to study, not a prize to chase. Over time, she realised fame fades unless it is converted into structure.

That insight led to a shift from endorsements to ownership. SKIMS, launched in 2019, solved a real product gap around fit and inclusivity. She also reframed influence by entering criminal justice reform, signalling a more responsible use of power. Her core strength lies in turning audience attention into infrastructure, hiring experts, iterating quietly, and letting products lead. As she looks ahead, the focus is on depth over hype, building platforms that can outlast fame. Systems scale. Spectacle does not.

Kim Kardashian's career did not begin with a master plan. It began inside a media storm she did not control. Growing up in Los Angeles, surrounded by lawyers, stylists, and celebrity proximity, she learned early how image, access, and attention worked. Fame arrived fast and loud. For years, it felt accidental, even unserious. Critics dismissed her as a byproduct of reality TV. She absorbed the noise, stayed visible, and quietly learned.

The early turning point was not a show or a scandal. It was repetition. Season after season, she watched what stuck, what faded, and what audiences actually responded to. She noticed that attention behaves like a commodity. It spikes, crashes, and moves on unless you give it structure. Most celebrities burn out at that stage. Kardashian slowed down and studied the machinery.

The first Step

Her first serious experiments in business were messy. Beauty apps, fragrance launches, and endorsements that worked but taught limits. The real inflexion came when she stopped licensing her name and started building ownership. SKIMS, launched in 2019, was born from a personal frustration with shapewear that did not fit real bodies. That origin story mattered. It was specific, relatable, and rooted in lived experience. More importantly, she stayed deeply involved. Fabric choices, sizing logic, marketing language, customer feedback loops. This was not a vanity project.

Public perception began to shift. She was still a celebrity, but now one who read customer reviews at scale and redesigned products accordingly. At the same time, she took an unexpected pivot into criminal justice reform, studying law and advocating for sentence commutations. That move confused many, but it revealed something important. Kardashian was building credibility in multiple, very different arenas. She was learning how to carry influence responsibly

By the mid-2020s, she had turned from being underestimated to being strategically misread. People still saw the spectacle. Few noticed the operator underneath.

Kim converts attention into infrastructure.

Celebrity is not the moat. Systems are

SKIMS did not win because of her following alone. The brand won because it fixed real product problems, executed supply chains well, expanded sizing inclusively, and reinvested profits into product depth. Her celebrity accelerated distribution, but retention came from operational discipline. Founders should note the order. Trust first, amplification later.

She uses audience as a research engine

Kardashian treats her audience less like fans and more like a continuous feedback loop. Drops are tests. Campaigns are data points. Social media becomes a live focus group. This is creator-led commerce done with intent, not impulse. In 2026, more brands will copy this, but few will do it with her consistency.

Brand voice beats brand volume.

SKIMS marketing avoids loud celebrity theatrics. The tone is neutral, confident, and product-forward. That restraint is deliberate. Kardashian learned that credibility grows when the founder steps back and lets the product speak. Founders often over-index on personal visibility. Her lesson is counterintuitive. Pull back to scale trust.

Crisis response through silence and correction

She has faced backlash, missteps, and cultural criticism. Her approach is rarely defensive. Pause, assess, adjust, move on. In business, this shows up as rapid iteration rather than public argument. For operators in volatile markets, this is a playbook worth studying.

Team-building without ego protection

Kardashian hires specialists and lets them run. Designers, supply chain experts, operators. She does not need to be the smartest person in the room. That humility is rare among celebrity founders. It is also why her ventures feel institutional rather than personality-dependent.

Looking ahead to 2026, her next phase is about depth, not expansion for its own sake. Expect SKIMS to move further into global retail, category adjacencies like activewear and menswear at scale, and possibly infrastructure plays around manufacturing and logistics. The shift is from brand to platform. Less noise, more permanence.

There is also growing interest in how she blends social impact with enterprise. Legal reform work is not a side quest. It is shaping how she thinks about power, accountability, and long-term influence. For founders, this signals a broader trend. Celebrity-led businesses of the next decade will be judged not just on revenue, but on responsibility.

Kim Kardashian's real achievement is not reinvention. It is an accumulation. Of knowledge, systems, trust, and optionality. In 2026, the world will not be watching to see if she can launch another product. It will be interesting to see if she can build something that outlasts her own fame. That is the hardest transition of all, and the most instructive one for anyone building in public today.

AMAZON'S DEEPFLEET

The robot handling your Amazon orders

On a busy holiday weekend in late 2024, something quietly changed inside Amazon's fulfilment centres. There were no flashy humanoids walking the floor. No press demo with applause. Instead, packages started moving faster, with fewer traffic jams, fewer idle robots, and shorter human walking distances.

The shift came from Amazon Robotics' DeepFleet, a large-scale AI system designed to coordinate more than one million warehouse robots as a single intelligent fleet. It did not replace robots. It taught them how to work together.

By March 2026, DeepFleet will represent a turning point for global robotics. The breakthrough is not hardware. It is orchestration at an industrial scale, and that is where the next decade of robotics business is heading.

SUMMARY

Aman Resorts began in 1985 as a quiet gamble. Adrian Zecha, walking along Phuket's Pansea Beach, chose to build a tiny, ultra-private resort instead of a holiday home. Banks refused to fund it, so he self-financed the idea. When Amanpuri opened in 1988, at prices far above the market.

Under Vladislav Doronin, Aman expanded carefully into cities, real estate, private clubs, and now a luxury yacht, without diluting its core promise. Design stayed site-specific, service remained invisible, and scarcity became strategy. Today, Aman earns heavily from branded residences, loyal repeat guests, and gated access. With the launch of Janu and deeper bets in the Middle East, Aman is testing a rare question for luxury in 2026: can a brand built on smallness grow without losing its soul?

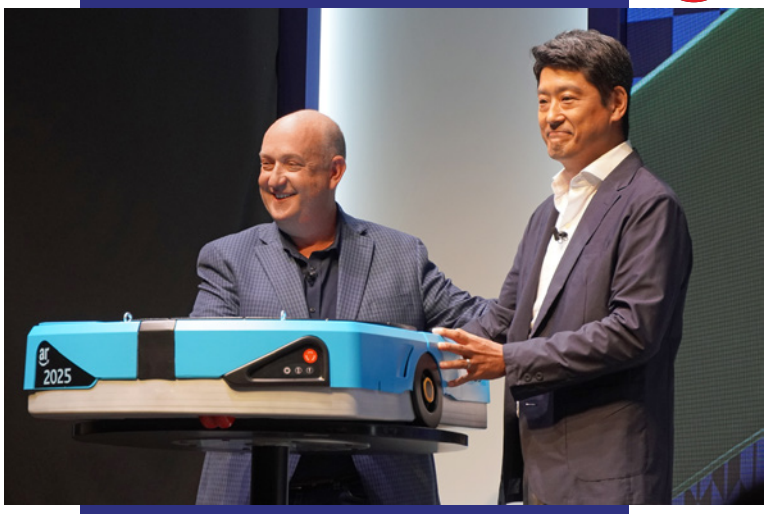

How the idea came to life

Amazon did not start out trying to build the world's largest robot brain. The story begins in 2012, when it acquired Kiva Systems and rolled out orange robots that carried shelves to human pickers. The robots worked, but coordination was basic. Each robot was smart on its own, dumb as a group.

As Amazon added new robot types, Sparrow for item picking, Proteus for autonomous movement around humans, and robotic arms for sorting, the system began to strain. Engineers noticed a strange problem. Individual robots were performing well in tests, but overall warehouse efficiency plateaued. Robots waited for each other. Aisles clogged. Humans walked more, not less.

One internal simulation showed a painful truth. Adding more robots without smarter coordination actually reduced throughput.

That was the inflexion point. A small team inside Amazon Robotics began treating the warehouse like a living system, not a collection of machines. They borrowed ideas from traffic engineering, swarm intelligence, and large-scale cloud scheduling. The early prototypes failed badly. In one test, robots optimised for speed created deadlocks that froze entire zones. In another, energy-efficient routing slowed urgent orders.

The breakthrough came when the team flipped the goal. Instead of optimising each robot, they optimised the fleet outcome. DeepFleet was trained on years of fulfilment data, weather spikes, demand surges, and human movement patterns. It learned when robots should rush, when they should wait, and when humans should be routed instead.

Engineers describe the moment it clicked as unsettling. The system began making counter-intuitive decisions, sending robots on longer paths that somehow improved overall speed. Trust came slowly. Only after months of shadow testing did DeepFleet go live.

What we can take

What DeepFleet actually does

In simple terms, DeepFleet is a large-scale decision engine. It sits above individual robots and assigns tasks, routes, priorities, and timing in real time. Think of it as air traffic control for warehouses, but powered by continuous learning. It does not control motors. It controls flow.

Why this matters in 2026

Most robotics companies still sell machines. Amazon is building systems. As robotics spreads into warehouses, hospitals, hotels, and cities, coordination becomes the bottleneck, not hardware.

By 2026, the competitive edge will belong to companies that own orchestration layers, not just robots.

Commercial and business model impact

DeepFleet strengthens Amazon's robotics-as-a-service economics. Faster fulfilment means fewer buildings, lower energy use, and higher asset utilisation. The same fleet moves more goods.

For startups, the lesson is clear. Value is shifting from selling robots to selling outcomes. Subscription software, usage-based pricing, and performance-linked contracts will dominate robotics revenue models.

Industries this unlocks

Warehouse logistics is just the start. Similar orchestration models are emerging in airport ground operations, hospital logistics, last-mile delivery hubs, and even agriculture. Anywhere multiple robots and humans share space, coordination beats brute force automation.

Competitive landscape

Players like Ocado, AutoStore, and emerging Chinese warehouse robotics firms are racing toward similar fleet intelligence. The moat is not algorithms alone. It is data density, deployment scale, and years of messy real-world learning.

Amazon's advantage is brutal. No lab can simulate the chaos of a million robots and humans working 24x7.

Adoption and regulatory hurdles

For most companies, the barrier is not trust in robots, but trust in autonomy at scale. Safety certification, labour integration, and system explainability remain challenges. Regulators will increasingly ask who is accountable when AI makes fleet-level decisions.

DeepFleet shows one answer. Keep humans in the loop, but let machines run the traffic.

Investment signals for 2026

Investors should look beyond humanoid demos. The real value is in software layers that sit above hardware. Startups building fleet intelligence, task allocation engines, and human-robot coordination systems will see strong demand, even without owning a single robot.

Why does this shape the future

DeepFleet is a quiet revolution. It proves that the future of robotics is not about making robots more human. It is about making systems more intelligent.

By 2026, the winners in robotics will not be those with the flashiest machines, but those who can choreograph thousands of imperfect machines into something that feels seamless. Amazon has shown the path. The rest of the industry now has to catch up.

THE PAPER FUEL

BUSINESS MAGAZINE

March 2026