

Ercole Rovida's Newsletter

September 2026



MESSAGE FROM THE AUTHOR

The Compliance Inflection Point

AI regulation moves from policy to enforcement — and reshapes how enterprises must operate.

For the past two years, the dominant conversation in enterprise AI has been about adoption: which models to deploy, which workflows to automate, which use cases deliver ROI. That conversation has not ended — but in August 2026, a second and equally demanding conversation has moved to the foreground. How do you govern what you have already built?

The trigger is concrete. August 2, 2026 was the binding enforcement date for high-risk AI system obligations under the EU AI Act, covering provider requirements under Articles 9 to 17 and deployer requirements under Article 26. While the European Parliament's June 2026 approval of the Digital Omnibus on AI deferred compliance for standalone high-risk systems to December 2, 2027, a critical trap remains: core user-transparency rules under Article 50 — disclosing AI chatbots and synthetic media to users — still apply on schedule. Non-compliance carries penalties of up to 7% of global annual turnover. The practical challenge for most organisations is not legal — it is operational. Providers must complete conformity assessments,

register systems in the EU AI database, implement quality management systems, and activate post-market monitoring. Deployers must implement human oversight mechanisms, retain automated logs for at least six months, and conduct Fundamental Rights Impact Assessments where required. Enterprise compliance programs lag far behind the pace of AI deployment.

The numbers tell an uncomfortable story. First-year compliance costs for large enterprises are estimated between €8 million and €15 million. Per-system costs for high-risk applications run to approximately \$1 million annually. And the Act's extraterritorial scope means any company deploying AI within EU markets — regardless of where it is headquartered — must comply. This creates a Brussels Effect that forces global technology firms to align products with EU standards even outside Europe, making AI governance a global operational requirement, not a regional one. The simultaneous arrival of NIS2 enforcement in Spain and across the EU adds a second compliance front. Essential entities — operators of critical infrastructure including energy, transport, banking, health, and digital infrastructure —

face mandatory incident reporting, supply chain security requirements, and board-level accountability for cybersecurity failures. Cloud computing services, managed service providers, and managed security service providers are explicitly captured, regardless of size in many cases. This is the defining tension of the September edition. The age of deployment is not over — but it now runs in parallel with the age of accountability. For IT leaders, telecom operators, public sector organisations, and industrial enterprises alike, this convergence of regulatory pressure and operational complexity represents both a cost and an opportunity. Those who build compliance as infrastructure — rather than treating it as a legal exercise — will find it becomes a trust signal that opens procurement doors, particularly in public sector and regulated enterprise markets. Governance is no longer a constraint on innovation — it is a prerequisite for sustained market access. The enterprises that will define this next phase are not those moving fastest, but those moving most deliberately: embedding accountability into their architectures before regulators, customers, or auditors force it upon them.

GLOBAL MARKET SNAPSHOT

In their Q2 2026 earnings calls, Microsoft, Alphabet, and Meta shifted focus from aggregate capex to time-to-energy, large-scale networking, power procurement, and the speed of converting infrastructure into revenue-generating compute. The hyperscaler race is no longer about how much money to spend — it is about how fast the grid can absorb it. Amazon indicated it will increase AI infrastructure spending this year, yet still expects capacity to trail customer demand, and is planning Project Eagle — a four-building data centre campus in Wharton County, Texas. Meta advanced two large-scale initiatives, including the expansion of its Hyperion campus in northeast Louisiana into what the company describes as a 5 GW AI supercluster — an initiative with implications that extend well beyond conventional data centre planning. The capital expenditure of the 14 largest publicly owned data centre operators globally is approaching \$750 billion in 2026, against just under \$450 billion last year. Over 23 gigawatts of data centre capacity was under construction globally at the end of Q3 2025, with approximately three quarters concentrated in the United States. Analyst expectations for spending by these same operators in 2027 have been revised upward by 56% since August 2025 — a signal that the market sees no near-term plateau. The constraint is no longer capital or demand. AWS data centre withdrawal in several US locations signals growing political and community challenges — with local opposition, grid access queues, and permitting delays emerging as the

binding limits on expansion in mature markets. The energy bottleneck is forcing a rethink of where and how AI infrastructure is built.

In Europe, Eutelsat is positioning 5G non-terrestrial networks at the core of the EU's expanded sovereign satellite project, explicitly targeting Starlink and Amazon's LEO ambitions — a signal that European policymakers now view satellite-terrestrial integration as strategic infrastructure, not an experiment.

The talent constraint is proving equally binding. An industry-wide survey published by TM Forum in August 2026 found that 70% of global telecom executives identify the scarcity of AI-fluent network engineers as the primary impediment to their automation strategies. The traditional silos between network engineering and data science are proving a major organisational hurdle that capital investment alone cannot solve.

The overarching signal for ICT vendors and enterprise buyers: infrastructure spending has never been higher, but the bottleneck has moved from capital to execution — power grids, permitting, and people with the skills to operate AI at scale.

MORE INFORMATION

[New Data Center Developments: August 2026](#)

[AI Data Center Build Advances at Full Speed: Five Things to Know](#)

[AI in Telecom: From 5G to AI-Native Networks — August 2026 Update — AI Conference London](#)

[Industry Outlook: Telecoms & Connectivity — Week of August 10, 2026](#)

REGIONAL FOCUS – SPAIN AND EUROPE

Spain's Digital Economy Reaches 26% of GDP — and Faces Its Compliance Moment.

The headline number out of Spain this August is striking. The digital economy, heavily driven by AI, now accounts for 26% of national GDP — equivalent to €414 billion — reflecting 17% year-on-year growth. Spain has climbed to 4th place globally in AI-related projects, and has leapt from 10th to 2nd worldwide for receiving greenfield ICT and internet infrastructure projects, with data centres capturing over 20% of global greenfield investment. Madrid and Catalonia rank as the 2nd and 3rd most attractive European regions for international ICT investment.

But August has brought Spain face-to-face with the compliance implications of this growth. Three regulatory currents are now converging simultaneously: the EU AI Act reaches its broadest application date on 2 August 2026, the NIS2 Directive requires national transposition and enforcement in 2026, and the España Digital 2026 agenda is channelling public funding and procurement standards toward AI governance and cybersecurity. For ICT vendors serving the public sector and regulated industries, compliance is now a sales prerequisite, not a post-sale concern. According to Spain's 2026 Digital Decade Country Report, 78% of Spaniards consider digital policy a high EU priority, with protecting online privacy and security identified as the top priority for the next decade by 92% of respondents.

Spain's national digital roadmap comprises 67 measures with a total budget of €33.8 billion — 87% of which expire by end of 2026, meaning the next six months represent the execution deadline for the bulk of Spain's digital transformation agenda.

On the telecom front, Spain's market consolidation is entering a new chapter. Telefónica and MasOrange have held informal talks regarding a potential deal for Vodafone Spain, with one scenario involving separating Vodafone Spain's fixed-line and mobile assets to address potential antitrust concerns. If it materialises, the move would reduce Spain's major operators from four to three and concentrate infrastructure ownership further — with significant implications for wholesale access, enterprise connectivity pricing, and network investment cycles.

European Pulse

The EU AI Act's Digital Omnibus deferred standalone high-risk AI systems — recruitment tools, credit scoring, education, law enforcement, critical infrastructure — to full compliance by 2 December 2027. Transparency obligations for AI-generated content however apply from August 2026 without delay, and penalties reach 7% of global annual turnover for the most serious breaches.

MORE INFORMATION

[Digital Economy - Invest in Spain](#)

[Spain's 2026 Digital Decade Country Report](#)

[The EU AI Act in Spain: What Every Tech Startup Must Know in 2026](#)

[EU AI Act: What Actually Applies on 2 August 2026](#)

FEATURED STORY

When the Agent Becomes the Target: Cybersecurity in the Agentic AI Era

For years, enterprise cybersecurity was built on a relatively stable set of assumptions: protect the perimeter, control human identities, monitor network traffic, respond to known threat signatures. Those assumptions are now structurally obsolete. The arrival of agentic AI — autonomous systems that plan, execute, and interact with enterprise infrastructure with minimal human oversight — has created an attack surface that legacy security models were never designed to handle.

The numbers confirm the shift. A 2026 Dark Reading readership poll found that 48% of cybersecurity professionals now rank agentic AI as the top attack vector heading into this year, outranking deepfakes, passwordless adoption, and board-level cyber recognition. Gartner has named agentic AI the single most important cybersecurity trend of 2026. And Forrester has gone further, predicting that agentic AI will cause a public breach in 2026 significant enough to result in executive dismissals — the moment when theoretical risk becomes corporate liability.

The New Attack Surface

What makes agentic AI different from previous technology risks is not its novelty but its permission level. AI agents are designed to act — executing code, querying databases, sending communications, moving files, and interacting with external APIs — across multiple systems simultaneously. Every agent introduced into an organisation creates what security researchers call a Non-Human Identity: a service account, API key, or session token that grants the agent the access it needs to operate.

These non-human identities have become the fastest-growing attack vector in enterprise security, according to the Huntress 2026 data breach report.

Security researchers have identified five primary attack vectors targeting agentic systems. Prompt injection embeds malicious instructions inside documents, emails, or web pages that an AI agent processes — causing it to execute attacker-controlled commands while appearing to operate normally. Memory poisoning corrupts the persistent context that agents carry between sessions, introducing false beliefs that systematically distort future decisions. Tool misuse and privilege escalation exploit the elevated permissions agents hold to access systems beyond their intended scope. Cascading failures occur when compromised agents trigger actions in downstream agents, multiplying the blast radius of a single breach across entire automated workflows. Identity impersonation uses stolen session tokens or API keys to masquerade as a legitimate agent — invisible to monitoring systems that see valid credentials and assume authorised activity.

The Shadow AI Compounding Factor

The security challenge is amplified by a parallel phenomenon: shadow AI. Employees are importing unsanctioned AI tools into work environments without security oversight, bypassing procurement, legal review, and identity governance processes. More than a third of data breaches now involve unmanaged shadow data connected to these unsanctioned deployments — creating attack entry points that security teams cannot monitor because they do not know the assets exist.

The Governance Imperative

The response requires a fundamental rethinking of enterprise security architecture. Treating agent identity as a first-class security concern

means extending the same governance applied to human users — least privilege, continuous authentication, behavioural monitoring — to every non-human identity in the environment. This is not a configuration change but an architectural shift, and it must be built before deployment, not retrofitted after a breach.

For IT leaders and CISOs in telecoms, utilities, financial services, and public sector organisations — the sectors most aggressively deploying agentic AI and most exposed to NIS2 and EU AI Act compliance obligations — the calculus is direct. Speed of AI adoption and strength of security governance are not in tension. In 2026, one is a prerequisite for the other.

[Agentic AI: Biggest Enterprise Security Threat for 2026](#)

[2026: The Year Agentic AI Becomes the Attack-Surface Poster Child](#)

[Top Cyberthreats in 2026: Agentic AI Will Trigger a Breach](#)

[CTop Agentic AI Security Threats in Late 2026](#)

[Agentic AI Security: Protecting Enterprises from the #1 Threat of 2026](#)

LISTEN ON THE GO

Jensen Huang: The Engineer Who Made AI Infrastructure

There is a question Jensen Huang has been asked for most of his career: what is NVIDIA, really? For its first decade, the answer seemed simple — a company that made graphics chips for video games. Then it became a high-performance computing company. Then the backbone of deep learning research. Then the engine of the generative AI boom. In 2026, Huang has given the clearest answer yet. Speaking at Davos in January, he put it without ambiguity: "AI is infrastructure. There's not one country in the world I can't imagine that you need to have AI as part of your infrastructure. Because every country has its electricity, you have your roads — you should have AI as part of your infrastructure."

Jensen Huang was born on February 17, 1963, in Tainan, Taiwan. His family moved to Thailand when he was a child, and he later immigrated to the United States for his education. He studied electrical engineering at Oregon State University and went on to complete a master's degree at Stanford. Before founding NVIDIA, he worked as a microprocessor designer at AMD and then as a director of LSI Logic — experience that gave him a rare combination of deep technical fluency and commercial instinct that would define everything he built afterward.

In 1993, Huang co-founded NVIDIA with Chris Malachowsky and Curtis Priem in a Denny's restaurant in San Jose — a detail he has recounted often, not for nostalgia but as a reminder of how unlikely the company's trajectory seemed at the start. The early years were genuinely precarious. Several product launches failed. The company came close to collapse before finding its footing with the RIVA 128 GPU in 1997, a chip that shipped in four months and saved the business.

What Huang understood before almost anyone else was that the GPU — designed to process graphics in parallel — was structurally suited to the kinds of matrix computations that underpin machine learning. In 2006, NVIDIA launched CUDA, a parallel computing platform that allowed researchers to programme GPUs for general scientific computation. It was a decision made years before the deep learning revolution it would eventually enable. When AlexNet demonstrated the power of GPU-accelerated neural networks in 2012, CUDA was already a mature platform with a global research community built on top of it. The advantage was not accidental — it was a decade in the making.

By 2026, NVIDIA had expanded far beyond chips. The company's Blackwell GB300 AI systems, NVLink networking, and Vera Rubin architecture represent the next generation of AI infrastructure, designed for increasingly sophisticated agentic and physical AI workloads. At GTC 2026 in San Jose in March, more than 30,000 attendees from 190 countries gathered for a conference that had become, in Huang's own words, "the epicenter of the AI industrial era." He described AI as a five-layer cake of infrastructure — energy, chips, infrastructure, models, and applications — and argued that every layer is advancing simultaneously, at a pace the world has never seen.

His leadership style is unusual for a technology CEO of his stature. He has kept NVIDIA private until 1999, maintained remarkable continuity of vision across thirty years, and built a culture that rewards deep technical expertise over managerial hierarchy. He has also been consistently willing to say things that seemed premature — and been proven right.

"Software is eating the world, but AI is going to eat software"

As of mid-2026, NVIDIA's data centre revenue for full fiscal year 2026 totalled approximately \$93.7 billion, up 75% year-on-year. Jensen Huang's net worth stands in the range of \$176 to \$200 billion. Both numbers will likely look modest in five years. He has spent thirty years making sure of it.

[Jensen Huang: The Leader Powering the Global AI Revolution](#)

[Jensen Huang: The Visionary Who Turned NVIDIA into the Engine of the AI Revolution \(2026 Edition\)](#)

[NVIDIA CEO Jensen Huang and Global Technology Leaders to Showcase Age of AI at GTC 2026](#)

[Jensen Huang Biography: NVIDIA CEO, AI Visionary, Career, Net Worth & Legacy](#)



COMPANY OF THE MOUNTH

SoftBank: The Biggest Bet in the History of Technology



Masayoshi Son has spent his career making bets that seemed reckless at the time. He invested \$20 million in Alibaba in 2000 when it was an obscure Chinese startup, a stake that eventually became worth over \$100 billion. He lost most of his personal fortune on dot-com stocks and rebuilt it. He poured billions into WeWork and watched it collapse. He turned Arm Holdings into the backbone of the global chip industry. Each cycle followed the same pattern: a vision dismissed as excessive, followed by reality catching up.

In June 2026, Son made his largest bet yet. Standing alongside French President Emmanuel Macron in Paris on June 1, he announced that SoftBank would invest up to €75 billion — approximately \$87 billion — to build 5 gigawatts of AI data centre capacity in France, marking the company's largest AI infrastructure commitment in Europe. The first phase comprises €45 billion to deliver 3.1 GW of capacity in the Hauts-de-France region by 2031, in partnership with Schneider Electric. SoftBank's shares rose 14% in two days. Year-to-date gains reached approximately 80%, briefly pushing the company past Toyota as Japan's most valuable publicly traded company. Son did not understate the moment. "I think this is like more than 10x, probably 50x bigger than dot-com," he told CNBC in Paris. "This is the biggest revolution of technology and realisation that mankind ever

This is just like the beginning of the internet." On the question of an AI bubble, he was direct: he denied one exists, and described what is underway as a technological revolution that could last 50 to 100 years.

The France deal is not a standalone move. SoftBank announced in March a 10 GW data centre project in Ohio valued at approximately \$500 billion in total infrastructure. It has committed over \$60 billion to OpenAI. It holds a large stake in Arm Holdings, whose chip designs are now embedded in AI servers and data centres globally. It is co-building a 5 GW AI computing cluster in Abu Dhabi alongside G42, OpenAI, Oracle, NVIDIA, and Cisco. The strategic logic is consistent: SoftBank is no longer primarily a venture fund that writes cheques into technology startups. It is positioning itself as a builder and owner of the physical infrastructure required to run the next generation of AI systems at scale. The France commitment also carries a European significance that goes beyond the economics. At a moment when the EU is confronting its dependence on US and Chinese technology platforms through the Cloud and AI Development Act, SoftBank's investment offers an alternative model: non-American, non-Chinese capital building sovereign-grade AI infrastructure

EU, with hyperscalers as anchor customers and European industrial partners in the supply chain.

Whether this bet follows the Alibaba arc or the WeWork arc remains to be seen. Son himself has never claimed certainty — only conviction. "There's no choice," he said of Europe's need for AI infrastructure. "U.S. and China are racing." In 2026, SoftBank is not watching the race. It is funding the track.

<https://group.softbank/en>

[SoftBank becomes Japan's most valuable company, announces \\$53 billion French investment](#)

[SoftBank CEO doubles down on AI, stock market message for 2026](#)

[SoftBank plans up to €75 billion investment in French AI centers](#)

<https://cryptobriefing.com/softbank-masayoshi-son-ai-investment-spreel/>



Information Revolution

— Happiness for everyone

EVENTS

- **VMware Explore 2026** (| 31 Aug – 3 Sept 2026 Las Vegas, USA) One of the flagship enterprise infrastructure events, VMware Explore brings together cloud, networking, and security leaders to explore virtualisation, hybrid cloud, and AI-driven infrastructure at scale.
- **Nova Future Summit 2026** (28–30 September 2026, Napa, CA, USA): A closed-format senior executive summit with no press, no vendor pitches, and no public tickets — focused on business models, capital allocation, and governance for the next wave of digital infrastructure. A rare space for candid conversation among peers at the highest level of the industry
- **MSP Summit 2026** (28–30 September 2026, Orlando, FL, USA): The MSP Summit brings together managed service providers, channel partners, and solution providers to address the rapidly evolving landscape of AI-driven IT services, cybersecurity, and enterprise automation.





Ercole Rovida's Newsletter

DISCLAIMER:

The content shared in this newsletter may include articles, excerpts, and opinions from third-party sources. While I strive to provide accurate and reliable information, the views and opinions expressed in these external contributions are those of the authors and do not necessarily reflect the views of Ercole Rovida's Newsletter. All third-party content is used for informational purposes only, and we do not claim ownership of these materials unless otherwise stated. Please consult the original source for full context and details.

I VALUE YOUR INPUT:

I'm always looking to improve and make my newsletter more relevant to you! If you have any suggestions, article recommendations, or feedback, feel free to reach out to me via LinkedIn message.

I'd love to hear from you!