

Ercole Rovida's Newsletter

October 2026



MESSAGE FROM THE AUTHOR

The Physical Layer Strikes Back

AI's next constraint is not code, capital or chips — it is glass, copper, concrete and the grid.

For three years, the story of artificial intelligence has been told in the language of software: model releases, benchmark scores, agent frameworks, token prices. September 2026 told it in a different vocabulary — megawatts, route miles, satellites, substations and permits. The message running through the month's news was consistent and, for anyone who sells infrastructure, deeply familiar: the digital economy still runs on physical things, and somebody has to build them.

Consider the evidence. Oracle reported that it delivered more than 300,000 GPUs and 850 MW of new data centre capacity in a single quarter — and that demand for AI cloud services is still growing faster than supply. Ciena closed its fiscal third quarter with an \$8.5 billion backlog, and supply constraints and lead times, not demand, now limit how quickly orders become revenue. Zayo has secured fibre cable from Corning years ahead of need because it expects the glass itself to become scarce. In Europe, the €15.6 billion IRIS² constellation moved from political ambition to industrial contracts. Each story says the same thing: the bottleneck has migrated from silicon to everything around it.

The friction is growing too. In the United States, local opposition blocked or delayed data centre projects worth nearly \$200 billion in the first half of 2026. In Spain, the Government's draft decree on data centres would tie grid access to additional renewable generation and digital-sovereignty commitments — a proposal the industry association SpainDC says puts €9 billion of investment at immediate risk. Whatever one thinks of the policy, the signal is unmistakable: power grids have become the arena where AI ambitions are negotiated. Energy regulators, planning authorities and local communities now sit at the same table as technology buyers — and their timelines rarely match a product roadmap.

And as infrastructure becomes more valuable, it becomes more attractive to attackers. The operational technology that runs pumps, substations, cooling plants and building management systems — much of it designed decades before anyone imagined it would be connected — is now squarely in the crosshairs. This month's featured story looks at why utilities, and the data centres that depend on them, face a security problem that cannot simply be patched.

For ICT vendors, integrators and their customers, this shift changes the conversation. The questions that matter are no longer only which

model or which cloud, but where the power comes from, how the fibre gets there, who secures the controllers and how long the permits take. Value is moving towards those who can connect the digital and physical layers — networks, energy, OT security and systems integration — and deliver them as one project rather than five.

The AI era was supposed to make the physical world invisible. Instead, it has made it the decisive competitive factor. The companies that will lead the next phase are those that understand both sides of the socket: the workloads that consume power, and the infrastructure that delivers it — securely and on time.

GLOBAL MARKET SNAPSHOT

Optical networking has become AI's second act. Ciena reported record fiscal Q3 2026 revenue of \$1.67 billion, up 37% year on year, with adjusted EPS up 215%. Direct cloud provider revenue jumped 82% and now represents 53% of sales, while pluggable optics revenue more than doubled. Backlog rose by \$800 million in a single quarter to \$8.5 billion, and management expects at least 30% growth again in fiscal 2027. Two customers alone accounted for more than 40% of quarterly revenue — a reminder of how concentrated hyperscaler demand remains. Every GPU cluster needs a matching investment in the transport that connects it — inside the building, across the campus and between regions.

The physical supply chain is tightening accordingly. Zayo, which plans to add 15,000 new route miles in North America by 2030 — including more than 8,000 miles of long-haul fibre developed with NVIDIA — has secured a major portion of the cable it expects to need through the end of the decade under an expanded agreement with Corning. When operators contract for raw materials years ahead, it tells you where the next shortage is expected.

Yet building is getting harder. According to Data Center Watch figures reported in September, at least 45 US data centre projects worth nearly \$68 billion were blocked or delayed between April and June, on top of 75 projects worth around \$130 billion in the first quarter — close to \$200 billion disrupted in six months. Water, energy use and community acceptance are now as material to project timelines as chip allocations.

In the radio access network, AI is finally producing measurable

returns. At the Intelligent RAN Forum on 22 September, Optus described live deployments with Ericsson covering link adaptation, coverage prediction and automated coverage compensation, targeting spectral efficiency gains of 10–15% and reaching 25% in selected field cases — on existing hardware.

The human dimension remains the hardest. EY's latest Telco of Tomorrow study, based on interviews with more than 90 telecom leaders, found that 81% rank operating model overhaul as a leading priority, and 69% expect three-quarters of today's workforce to be significantly upskilled or replaced within five years.

Capital markets are rewarding the builders. AMD crossed \$1 trillion in market value on 21 September, while Oracle's contracted backlog reached \$664 billion — both covered later in this edition. Infrastructure suppliers, not only model developers, are capturing the value of this cycle.

The overarching signal: demand is not the problem. Execution capacity — fibre, optics, permits and people — is.

MORE INFORMATION

[Ciena Q3 Revenue Jumps 37% as AI and Cloud Drive Optical Demand](#)

[Zayo Locks In Fiber Supply as AI Infrastructure Demand Intensifies](#)

[Nearly \\$200 Billion of Data Center Projects Blocked or Delayed](#)

[AI-RAN Gets Real — Optus Claims Major Gains with Ericsson](#)

[Telecoms Have an AI Vision. Now Comes the Hard Part \(EY\)](#)

REGIONAL FOCUS – SPAIN AND EUROPE

Spain Rewrites the Rules for Data Centres — and the Industry Pushes Back

Spain's data centre boom has collided with the limits of its electricity grid. Data centre projects have obtained more than 12 GW of grid access and connection rights, while estimated demand for 2030 is between 3.5 and 4 GW. To close that gap, the Ministry for the Ecological Transition has put out for consultation a draft Royal Decree that changes how capacity is allocated: access would no longer be first come, first served, but conditional on energy efficiency, new renewable generation and digital sovereignty.

The requirements are demanding. Facilities of 1 MW or more would have to cover at least 80% of their hourly consumption with additional renewable generation commissioned no earlier than 18 months before start-up. Non-compliance could bring surcharges of 100% to 500% on network charges, and loss of grid access if coverage falls below 60% for five consecutive years. Operators would also have to be established in the EU, keep operational data within the Union and resist foreign authority requests contrary to European law.

On 11 September, SpainDC called the draft a “de facto veto”, warning that €9 billion of investment is at immediate risk against a sector pipeline of around €67 billion for 2026–2030. The Ministry, which received more than 600 submissions, defends an “orderly and efficient” development, and Merlin Properties says its projects already comply.

In telecoms, Zegona, owner of Vodafone Spain, announced on 24 September a €400 million shareholder return — a €200 million dividend and

a €200 million buyback — after cutting leverage from 3.1x to 2.3x EBITDAaL since the 2024 acquisition, crediting its PremiumFiber and FiberPass fibre joint ventures. Zegona now expects Vodafone Spain to self-fund its growth strategy from annual cash flow.

European Pulse

Europe's sovereign satellite constellation became industrial reality. At the International Space Summit in Paris on 10 September, Airbus signed the initial contract to build the first layer of IRIS² satellites for Eutelsat — at least 66 LEO satellites for delivery in 2029 — with Thales Alenia Space supplying secure government payloads. The programme now totals 348 satellites and €15.6 billion, with Spain's Hispasat leading the ground segment. Meanwhile, Spain has still not transposed NIS2: after the Commission referred it to the EU Court of Justice in July, the national cybersecurity law remained at draft stage in late September.

MORE INFORMATION

[SpainDC denuncia un veto al sector y duda del encaje jurídico del nuevo Real Decreto](#)

[España pone freno a la fiebre de los centros de datos](#)

[Zegona \(Vodafone España\) repartirá 400 millones a sus accionistas](#)

[Airbus Starts the Development of IRIS² Satellites](#)

[NIS2 en España: qué exigirá la ley a los proveedores de la Administración](#)

FEATURED STORY

When the Grid Becomes the Target: OT Security for Utilities in 2026

For most of the last decade, cyberattacks on utilities followed a familiar script: criminals breached corporate IT, encrypted servers, demanded a ransom and — in most cases — never touched the systems that actually move water or electricity. That script is changing. Attackers are no longer content to hold data hostage. Increasingly, they are reaching the controllers that run physical processes, and learning how to manipulate them.

A Month of Warning Signs

September showed both scenarios. On 1 September, Stadtwerke Landsberg, a German municipal utility, suffered a criminal attack that encrypted its central IT systems. Its crisis team cut internet links and, crucially, electricity, water, district heating and fibre services kept running — segmentation working as designed.

In Colorado, the picture was darker. State officials disclosed in mid-September that foreign actors had penetrated the operational technology of two small water utilities in late August, modifying pumping cycles, altering device parameters, disabling alarms and cutting remote access. The incidents follow a summer campaign in which more than 30 Minnesota water systems were hit, and utilities in at least seven US states reported intrusions into internet-exposed programmable logic controllers. The techniques resemble those of Iran-linked groups, though attribution remains unconfirmed.

The vulnerability landscape is equally active. Schneider Electric's September advisories include a critical authentication flaw (CVE-2026-3869, CVSS 9.2) in certain

Modicon M580 controllers, while CERT Polska confirmed that two MikroTik RouterOS flaws were being chained in live attacks — the kind of edge equipment that often sits at the boundary of industrial networks.

Why OT Is Different

The difficulty is not awareness but the physics of operations. Controllers stay in service for 15 or 20 years. Patching requires maintenance windows, testing and sometimes site visits — in Minnesota, restoring locked-out controllers meant opening cabinets, pulling backup batteries and performing hard resets.

Convergence makes it harder. Every smart meter rollout, remote substation monitoring project and cloud-connected SCADA upgrade improves efficiency, but also adds pathways between the corporate network and the process network. Third-party remote access — vendors and integrators maintaining equipment across hundreds of sites — remains one of the least mature control areas in the sector.

The problem extends to the data centres AI depends on. Claroty's Team82 analysed more than 750,000 cyber-physical assets and found that 88% of building management systems use insecure protocols, while 41% of power distribution units and 32% of cooling systems are exposed or one network hop from risky connections.

The Regulatory Clock

NIS2 places energy, water and digital infrastructure among essential entities, with board-level accountability, incident reporting and supply chain obligations. In Spain, public tenders already cite NIS2 requirements ahead of the national

law. In the United States, the Cybersecurity Information Sharing Act of 2015, which protects companies that share threat data with the government, has been extended only until 30 September 2026 — leaving a question mark over the information-sharing model on which critical infrastructure defence relies.

What Good Looks Like

The playbook is well understood, if not yet widely implemented: remove controllers from the public internet, enforce multi-factor authentication on remote access, segment IT from OT using IEC 62443 zones and conduits, maintain continuous asset visibility and rehearse manual operation.

For leaders in utilities, telecoms and data centre operations, the conclusion is direct. The attack surface now includes pumps, breakers and chillers. Securing it is no longer an IT project — it is a condition of keeping the lights on, and procurement teams will increasingly demand proof that suppliers can deliver it.

[Daily OT Security News: September 09, 2026](#)

[Colorado Water Utilities OT Attacks: Investigative Report](#)

[Daily OT Security News: September 08, 2026](#)

[State of CPS Security: Data Center Exposures \(Claroty\)](#)

LISTEN ON THE GO

Lisa Su: The Engineer Who Rebuilt AMD

On 21 September 2026, Advanced Micro Devices crossed a threshold that would have seemed absurd a decade earlier: a market capitalisation of \$1 trillion. Three days later, its chair and chief executive, Lisa Su, was seated near President Trump at the White House state dinner for Chinese President Xi Jinping. For an executive who spent years avoiding the spotlight, it was a September to remember — and the culmination of one of the great turnarounds in technology.


Lisa Su was born on 7 November 1969 in Tainan, Taiwan, and moved to the United States as a young child. She studied electrical engineering at the Massachusetts Institute of Technology, earning her bachelor's, master's and doctoral degrees there. Her doctoral work on silicon-on-insulator technology addressed the problem that would define her industry for decades: making transistors faster and more efficient.

Her career combined research depth with operational responsibility. At IBM she led semiconductor research and development, and she later served as chief technology officer of Freescale Semiconductor before joining AMD in 2012 as a senior vice president. Two years later she became CEO.

The company she inherited was in serious difficulty: losing money, heavily indebted and trailing Intel in processors and NVIDIA in graphics, with shares trading around \$2. Su's diagnosis was disciplined and unglamorous — focus on high-performance computing, build a small number of excellent products and deliver them on time, every time. The Zen architecture, launched in 2017, and the Ryzen and EPYC

processors that followed restored AMD's credibility in PCs and then, decisively, in the data centre. Her bet on chiplets — building processors from smaller modular dies — proved one of the most consequential design decisions of the decade.

She then turned to AI. The Instinct accelerators, the acquisition of Xilinx and an open software ecosystem made AMD the principal alternative to NVIDIA for hyperscalers seeking a second source. In October 2025, AMD agreed to supply OpenAI with six gigawatts of GPU capacity, reframing the company as a core AI infrastructure supplier. At AMD's Advancing AI conference in July 2026, she presented new accelerators and rack-scale platforms aimed squarely at the largest AI clusters.



The numbers now reflect that shift. In the second quarter of 2026, AMD reported record revenue of \$11.5 billion, up 50% year on year, with data centre revenue more than doubling to \$6.7 billion. It expects server CPU revenue to grow 80% in the second half of 2026 — a sign that agentic AI is making general-purpose compute important again.

Her leadership style is quiet, precise and relentlessly engineering-driven. She masters technical detail, sets clear roadmaps and delivers against them — a reputation for reliability that has become a competitive asset with

customers planning multi-year infrastructure investments.

"Run toward the hardest problems."

Measured by market value, AMD has grown from a few billion dollars when she took charge to more than \$1 trillion — one of the largest value creations by any chief executive of her generation. Su has also become one of the industry's most consistent advocates of open software ecosystems, arguing that customers should never depend on a single company for the computing platforms on which AI is built.

Recognition has followed. Time magazine named her CEO of the Year in 2024, and she has received the IEEE Robert N. Noyce Medal, one of the semiconductor industry's highest honours. In a curious twist, she and NVIDIA's Jensen Huang — last month's profile — are distant relatives, a connection both have acknowledged.

For enterprise and telecom buyers, AMD's rise matters for a simpler reason too: a credible second supplier improves pricing, availability and negotiating leverage across the entire AI infrastructure stack.

The trillion-dollar milestone settles the question that hung over AMD in 2014: whether it would survive. It did — and it now helps define what AI infrastructure looks like.

[AMD CEO Lisa Su Is Having a September to Remember](#)

[AMD's Path to \\$1T Market Cap: A Timeline of Lisa Su's Big Wins](#)

[Lisa Su: The Definitive 2026 Profile](#)

COMPANY OF THE MONTH

Oracle: The \$664 Billion Backlog

Few technology companies have reinvented themselves as dramatically, or as expensively, as Oracle. For four decades it was synonymous with the database — the software layer that quietly ran the world's banks, airlines and governments. In 2026 it has become something else: one of the largest builders of AI computing capacity on the planet.

The scale was confirmed on 10 September, when Oracle reported first-quarter fiscal 2027 results; the shares, down more than 5% during the session, rebounded after hours. Total revenue grew 30% to \$19.3 billion. Cloud revenue rose 62% to \$11.6 billion, driven by Oracle Cloud Infrastructure, up 121% to \$7.4 billion — the ninth consecutive quarter of accelerating infrastructure growth. Remaining performance obligations reached \$664 billion, up \$209 billion year on year, after Oracle booked more than \$30 billion of new AI cloud contracts in the quarter.

The operational numbers are just as striking. Oracle delivered more than 300,000 GPUs to AI customers and added 850 MW of data centre capacity, almost triple the capacity delivered in the previous quarter, with GPU utilisation near 98%. Management raised full-year guidance to at least \$90 billion in revenue and says demand for AI training and inference is still growing faster than supply.


The strategy carries real risk, and Oracle does not hide it. Free cash flow was negative \$23.7 billion for fiscal 2026, and in the latest quarter the company sold \$20 billion of common stock to fund its buildout. To limit the capital burden, many large contracts are structured so that

customers prepay for GPUs or supply them directly — arrangements totalling \$75 billion by the end of fiscal 2026. Concentration is another concern: much of the backlog depends on a handful of very large AI customers.

The transformation has also reshaped Oracle's leadership. In September 2025 the company named Clay Magouyrk, who built its cloud infrastructure business, and Mike Sicilia as co-chief executives, with Safra Catz becoming executive vice chair — a clear signal that infrastructure execution sits at the centre of the strategy. Oracle is also a core partner in OpenAI's Stargate programme, among the largest AI construction projects ever undertaken.

What makes Oracle interesting for enterprise buyers is not only the infrastructure. The company is betting that AI will ultimately be applied to private enterprise data — and that much of that data already lives in Oracle databases. Its multicloud database services, running inside Amazon, Google and Microsoft data centres, have become its fastest-growing business. Larry Ellison argues that companies must first describe their data and processes precisely, in an “enterprise ontology”, before AI agents can automate them.

Oracle is running two bets at once: that AI demand will keep outrunning supply long enough to justify its capital programme, and that its database franchise will make it the natural home for enterprise AI. The market rewarded the first quarter. The next three will show whether power and capacity remain the only constraints that matter.




[Oracle Announces Q1 FY27 Results](#)

[Oracle Q1 FY2027 Slides: Cloud Infrastructure Surges 121%](#)

[Oracle Record Q4 and FY 2026 Results](#)

EVENTS

- **20ENISE** (21–23 October 2026, León, Spain): INCIBE's International Information Security Meeting celebrates its 20th edition, bringing together Spain's cybersecurity ecosystem — companies, public administrations and international delegations — around critical infrastructure protection, AI, regulation and cyber resilience, with a dedicated space for startups and business meetings.
- **Smart City Expo World Congress 2026** (3–5 November 2026, Fira Barcelona Gran Via, Spain): The world's leading urban innovation event, with more than 1,000 exhibitors and 600 speakers focused on housing, mobility, energy efficiency and the digital transformation of cities — a key meeting point for public sector ICT, connectivity and Smart Building.
- **Enlit Europe 2026** (10–12 November 2026, Vienna, Austria): Europe's flagship energy transition event, gathering 15,000 professionals and 750+ exhibitors around grids, digitalisation, metering and utility cybersecurity.



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!