


European Power: Generation, Grid and Price

Electricity Pricing in the Age of AI, on a Continent That Cannot Connect Anything

Bernd Fastenrath

Edition 2026 · 15 chapters · 28,638 words

.



Chapter 01, the preface, feeds all three.

The document as a single-line diagram. Three busbars, one feeder per chapter.

Contents

Opening

- 01 Preface: Why the American Primer Does Not Port

All About European Power Plants

- 02 Fundamentals: The Same Equation, Different Binding Constraints
- 03 The Power Plant, With Carbon In The Stack
- 04 To Build A Power Plant (Or A Data Centre) In Europe
- 05 Case Study: Sines
- 06 Where The Demand Actually Went

How To Trade European Power

- 07 Power Markets in Europe
- 08 Case Study: Germany
- 09 How The Price Is Actually Computed
- 10 Practical Approximations
- 11 Types of Power Trades

The Grid That Has To Deliver It

- 12 The Layer Nobody Prices
- 13 What Automation Actually Buys
- 14 The 2026 to 2030 Window
- 15 Where the Large Loads Actually Land

Where to start

This is a long document and most readers need a third of it. The chapters below are the ones each kind of reader has found worth their time; the rest is there when you want it.


If you invest in or finance European energy assets. Chapter 04 on why merchant thermal does not finance and what the tenders replaced it with, Chapter 05 on acquiring a site with an existing connection and the four conditions that decide whether the capacity comes with it, Chapter 08 on what capacity mechanisms actually clear at and why that price is a policy output, and the connection contribution in Chapter 15, which is a large line item routinely missing from German models.

If you run or plan a distribution network. Part 3, and within it Chapters 14 and 15. Chapter 14 is the argument about the replacement cycle and the regulatory treatment of recurring cost; Chapter 15 covers connection charging, the large loads arriving on your network, and the German customer-installation ruling. Parts 1 and 2 are written for a different reader and you can skip them.

If you develop, operate or buy data centre capacity. Chapter 04 on queues and permits, Chapter 05 on brownfield sites, Chapter 06 on where European demand actually went and what the Irish and German rules now require, the ancillary services section of Chapter 11 on why onsite flexibility earns less than the pitch decks claim, and Chapter 15 on connection cost and the flexible connection.

If you want the argument rather than the reference. Chapter 01, then Chapter 12, then What's Next.

A note on what you are reading. The author advises companies in the medium-voltage equipment sector, which is the subject of Part 3's recommendation. That disclosure is set out in Chapter 01 and is worth reading before Part 3 rather than after it.



Preface: Why the American Primer Does Not Port

Neel Somani's *Power 2026* is the best short introduction to US power markets I have read for a non-specialist audience. Almost none of its operating conclusions survive contact with Europe.

That is not a criticism of the piece. It is a statement about how differently the two systems are built. The American primer is, at bottom, a pricing document. It teaches you that the power price is the marginal cost of the last unit dispatched, that the last unit is usually a gas turbine, and that if you understand the turbine you understand the price. Everything downstream, siting, financing, hedging, follows from that.

Europe's binding constraint is not the marginal unit. It is permission.


In the FLAP-D cluster (Frankfurt, London, Amsterdam, Paris, Dublin) a new large load waits an average of seven to ten years for a grid connection, stretching to thirteen years in the most congested primary markets. Construction takes eighteen to twenty-four months. Ireland now requires a large energy user to bring its own dispatchable capacity to the connection, the Netherlands restricted hyperscale development nationally with a preparatory decision in February 2022 and permanent rules in 2023, and Frankfurt is severely constrained into the next decade. The European Union Agency for the Cooperation of Energy Regulators put direct grid congestion management costs at €4.3 billion in 2024, and that figure excludes the economic cost of everything the queue delayed.

So the American question is: *what will power cost, and can I hedge it?* The European question is: *will I be allowed to plug in, and what will I pay on top of the commodity when I do?*

Both are answerable. They require different tools.

This document is self-contained. Where it refers to American practice it explains what that practice is, so nothing here requires you to have read the US primer first. If you want it, it is the first entry in the sources at the end, and it is a better introduction to the American system than anything in these pages.

Equipment and standards



Market and measurement

Two clocks run on the same decade and nobody has synchronised them. The specifications that decide whether a station can measure anything are rewritten between 2025 and 2027, once, on the F-gas timetable. The German coverage requirement bites at the end of 2028, after the window has closed. Equipment chosen inside the band is still in service in the 2060s.

Where I am writing from

I run a strategy consultancy working with European industrial operators, mostly in DACH and Benelux. Before that I was VP Strategy at ABB E-Mobility, which put me on the equipment side of the problem this document describes: high-power charging is one of the loads now arriving at the medium and low-voltage layer that Part 3 is about, and I spent that period watching connection capacity rather than hardware become the constraint on deployment.

Two disclosures, because Part 3 ends in a recommendation and a reader should be able to discount it properly. First, I advise companies in the medium-voltage equipment and grid technology segment, including on go-to-market strategy, so I have a commercial interest adjacent to the argument in Chapter 14, which is that Europe should put measurement into the equipment it is already replacing. Second, the structural framing in Chapter 13, the levels of automation capability and the grouping of what the equipment does, is broadly how that industry organises its own thinking. I have consolidated and genericised it rather than invented it. Neither disclosure makes the Chapter 14 argument wrong. Both are reasons to check the numbers it rests on rather than take them on trust.

What the American picture gets right about Europe

- Marginal pricing.** Every generator that wants to run offers into an auction, the offers are stacked cheapest first, and the price everyone receives is set by the most expensive unit needed to meet demand. That stack is the *merit order*, and the unit that sets the price is the *marginal unit*. Europe runs exactly this, as a single auction across most of the continent at once.
- The balance equation.** Supply equals demand plus net exports plus change in storage plus losses, at every location, all the time. Electricity is not stored in the wires; the books balance every second or the frequency moves.

3. **Startup costs and the duck curve.** A thermal plant cannot be switched on instantly or for free, so what matters is not only the price but how fast it moves. Plot a sunny day's demand net of solar output and it sags in the middle and rises steeply at dusk: the shape is called the duck curve, and the steepness of its neck is what punishes slow plant. Europe has the same problem, worse in Iberia, worse still in the Nordics on a seasonal basis.
4. **The optimisation.** The auction is solved as a large constrained optimisation. The price that falls out for each location is the amount total system cost would rise if you demanded one more megawatt-hour there. The solver differs between continents; the mathematics does not.
5. **Missing money.** A market that pays only for electricity delivered, and nothing for standing ready, is an *energy-only* market. It does not pay enough to finance plant that runs a few hundred hours a year but must exist for the days it is needed, which is *firm capacity*. Europe discovered this too, and responded differently.

Where it breaks

1. **Europe prices by zone, not by node.** A nodal market publishes a separate price at every point on the network, so a congested line shows up immediately as a price gap between the two ends, and traders can take positions on that gap. Europe publishes one price for a whole bidding zone, often an entire country. Congestion inside the zone does not appear in the price at all. It appears as a cost: the operator pays one plant to turn down and another to turn up, and bills everyone. That is *redispatch*, and Chapter 08 puts a number on it. There is no European equivalent of the nodal price, the locational marginal price, to trade.
2. **Europe's marginal fuel is imported.** America's gas benchmark, Henry Hub, prices a domestic supply curve. Europe's benchmark, TTF, prices gas arriving at a Dutch trading hub in competition with every other buyer of liquefied natural gas on earth, at a premium for delivery. European power prices are, structurally, a derivative of Asian gas demand and Middle East shipping risk.
3. **Carbon is a third input in the merit order.** A European generator must surrender one emission allowance, an EU Allowance or EUA, per tonne of CO₂ it emits, bought at the market price. At an EUA around €85/tonne that adds roughly €31/MWh to a modern combined cycle gas turbine, a CCGT, the workhorse plant of the European system, and roughly €72/MWh to a hard coal unit. This is not a footnote. It is the mechanism that reordered the European stack.
4. **The data centre is a planning object, not an economic actor.** In the US a hyperscaler is a counterparty that can underwrite a plant. In much of Europe it is an applicant in a queue. It is subject to an efficiency directive, a municipal vote, and increasingly a connection policy that requires it to bring its own generation.
5. **Europe's binding layer is distribution, and it is the layer with the least measurement on it.** Around 70% of new European renewable capacity connects below the transmission grid, as does essentially all new EV and heat pump load. Chapter 12 sets out the voltage levels, which are worth having in front of you, because almost every argument in this document turns on which one a thing connects at. The American primer can treat the grid as a transmission constraint because its marginal-unit logic lives at that level. In Europe the delivery problem sits one or two voltage levels

down, in a network whose low-voltage observability runs between 40 and 65% even among advanced operators. Part 3 of this document is about that layer, because the market instruments in Part 2 that reach into it quietly assume it can be seen. The ones that stay above it, where a connecting party is obliged to fund its own telemetry, do not have that problem and are held up by something else.

One more asymmetry worth naming at the start, because it is the uncomfortable one. Europe's wholesale power is more expensive than America's, and Europe's contracted offtake prices are *lower*. Nordic wind PPAs clear around €35 to €50/MWh and the average European solar PPA signed in Q1 2026 was roughly €55/MWh. American colocation leases for AI tenants imply a willingness to pay for powered capacity that has no European equivalent, though the figures are not comparable with a generation offtake price and should not be set against one: they are a different product on the other side of the meter. What is comparable is the direction. European developers face worse revenue per MWh than their US peers while facing higher construction and connection costs. That gap is not a European bargain. It is a European financing problem, and Chapter 04 puts numbers on both sides of it.

Part 1: All About European Power Plants



Chapter 02

Fundamentals: The Same Equation, Different Binding Constraints

In this chapter

- The commodity balance equation holds identically in Europe. What differs is which term binds.
- Europe's gas price is set at the margin by seaborne LNG, not domestic production, which imports geopolitical volatility directly into the power price.
- Carbon is a priced input to the merit order, and the carbon price is a policy variable subject to legislative revision.

Start with the same identity:

$$\text{Generation} = \text{Consumption} + \text{Net Exports} + \text{Change in Storage} + \text{Losses}$$

In the US, the interesting term is usually generation. In Europe, the interesting terms are net exports and losses, because Europe is not one system. It is roughly two dozen bidding zones stitched together by interconnectors with finite capacity, operated by separate transmission system operators, under a coupled market algorithm that runs once a day.

Gas: the imported margin

Europe's benchmark gas price is TTF, the Dutch virtual trading point. Unlike Henry Hub, which reflects a domestic supply curve with export capacity as a constraint on the upside, TTF reflects what Europe must pay to attract LNG cargoes away from Asia, plus what Norway sends through pipelines, plus whatever storage policy requires.

The practical consequence showed up in July 2026. Front month TTF rose roughly a third over the month on renewed Persian Gulf hostilities and concern about Europe's ability to refill storage before winter, trading near €62/MWh at the end of July before easing to around €57 to €58/MWh in early August. EU storage sat around half full in early July, roughly fourteen percent below the five-year average for the date.

That is a supply shock in a commodity Europe does not produce, transmitted immediately into the power price. There is no domestic production response. There is only demand destruction and coal.

Carbon: the third leg

Every tonne of CO₂ emitted by a European power plant requires a European Union Allowance. EUAs traded above €85/tonne in July 2026, the highest since January, supported by higher energy prices and by the Commission's proposed overhaul of the Emissions Trading System. The proposal would slow the annual cap reduction to 3.7% over 2031 to 2035 and 1.7% from 2036, and extend free allocation for industry to 2037 conditional on companies investing in decarbonisation projects inside Europe.

Note what that means for anyone modelling forward power curves: a meaningful part of your fuel cost is a legislative outcome. There is no equivalent in the US stack outside California and RGGI, and neither of those sets the price for a third of the continent.

Marginal pricing, and Europe's ambivalence about it

Europe clears power the same way. All accepted offers receive the clearing price set by the marginal unit. In 2022 that mechanism transmitted a gas crisis into every household bill on the continent, and Europe has spent the years since building a contract layer on top of the spot market rather than changing the spot market itself: two-way contracts for difference for new generation, encouragement of long-term PPAs, and in December 2025 formal Commission guidance on how to design both.

This matters for how you read European prices. The spot price is still a marginal cost signal. It is increasingly not the price anyone actually receives or pays. Between the spot and the counterparty sit contracts for difference (CfDs), under which a state tops the generator up to an agreed strike price and claws back anything above it, power purchase agreements (PPAs), capacity payments, network tariffs, levies, and in Germany's case a direct federal subsidy to transmission tariffs. Modelling the spot correctly and the contract layer incorrectly will lose you more money than the reverse.



The Power Plant, With Carbon In The Stack

In this chapter

- Europe quotes generator efficiency as a percentage, not a heat rate. The conversion is trivial and constantly gets fumbled.
- The clean spark spread, what a gas plant earns per MWh after paying for its fuel and its carbon, and the clean dark spread, the same figure for a coal plant, are the two numbers that determine what runs.
- In the Nordics and the Alps the marginal unit is often not a fuel cost at all. It is an opportunity cost.

Efficiency, not heat rate

An American trader quotes a 7 heat rate unit. A European quotes a 55% efficiency unit. They are the same statement. Heat rate in MMBtu/MWh equals 3.412 divided by efficiency, so a 7 heat rate is roughly 48.7% efficient and a modern CCGT at 60% is a 5.7 heat rate.


Europe's convention is more convenient, because European gas is already priced in €/MWh of thermal content. Your fuel cost is simply the gas price divided by efficiency. No conversion factor.

The short-run marginal cost of a European CCGT

Take August 2026 numbers. TTF at €57/MWh thermal, EUA at €85/tonne, a modern CCGT at 55% efficiency.

- **Fuel:** $57 / 0.55 = \text{€}103.6/\text{MWh}$ electrical
- **Carbon:** natural gas emits roughly 0.2016 tCO₂ per MWh thermal, so $0.2016 / 0.55 = 0.367$ tCO₂/MWh electrical. At €85: **€31.2/MWh**
- **Variable O&M:** call it **€3 to €5/MWh**

Short-run marginal cost: roughly €138/MWh.



Carbon is not a footnote in the European stack. It is 31 euros per megawatt hour on a modern gas unit and 72 on a hard coal unit, which is the whole reason the merit order reordered. At an EUA of 85 euros a tonne the two meet when TTF falls to roughly 38 to 40 euros. The coal figure is a ballpark rather than a verified print: run it against a live API2 quote before trading on it.

That is why most major European markets printed weekly average day-ahead prices above €110/MWh in the third week of July 2026, and why Italian spot approached €500/MWh during the heatwave. It also gives you the single most useful sensitivity in European power: **in an hour that gas sets, a €1 move in TTF moves the German power price by roughly €1.80 at 55% efficiency, and by €2 at 50%.** Divide by efficiency. Attach the condition, though, because it is doing real work: gas is marginal at the top of the residual load curve, not across the year, and lignite and hard coal sit below it for a large share of hours. Quoted against a peak hour the rule is sound. Applied to a baseload position it overstates the sensitivity, and a desk would always ask what share of hours gas actually set before using it.

Coal, and the switching level

Hard coal emits roughly 0.34 tCO₂/MWh thermal. A typical European hard coal unit runs around 40% efficient, so 0.85 tCO₂ per MWh electrical. At €85/tonne that is **€72/MWh of carbon before a single tonne of coal is bought.**

Assume coal at the API2 benchmark, the reference price for thermal coal delivered into north-west Europe, in the region of \$100 to \$110/tonne, roughly €11 to €13/MWh thermal, so €28 to €33/MWh electrical. Total short-run marginal cost around €100 to €105/MWh.

Setting the two equal gives the fuel-switching level. At an €85 EUA, an efficient CCGT beats an efficient hard coal unit when TTF is below roughly €38 to €40/MWh. Above that, coal is back in merit.

Two caveats, and they are large. First, my coal assumption above is a ballpark rather than a verified print, so run it with a live API2 quote before you trade on it. Second, and more importantly, much of Europe's coal fleet is legally retired or under a binding phase-out schedule, so the switching signal now expresses itself mostly in Germany and Poland. It is a real economic signal with a shrinking physical fleet behind it. That is precisely the reverse of MISO, where recent EIA analysis found coal units outperforming gas on profitability.


The units that ignore the price

Three categories of European generation do not behave like a fuel-cost merit order at all, and they matter:

Reservoir hydro (Norway, Sweden, the Alps). A hydro reservoir has no fuel cost. Its bid is the *water value*: the expected profit from holding the water for a later hour or a later month, discounted. That makes the Nordic price a function of reservoir levels, inflow forecasts, and the forward curve, not a fuel cost. If you model NO2 or SE4 with a thermal stack you will be wrong all year.

Heat-led combined heat and power, or CHP (Germany, Denmark, Poland, the Netherlands). These plants serving district heating networks run when it is cold, not when the power price is high. In a January cold snap a large slice of the thermal fleet is dispatching against a heat obligation and treating power as a byproduct. This flattens the winter supply curve in ways a pure merit order does not capture.

French nuclear. Roughly 60 GW of capacity that is neither fully must-run, meaning obliged to generate whatever the price does, nor price-responsive. It is availability-constrained: outage schedules, corrosion inspections, and in summer, river temperature limits. The 2026 heatwave cut EDF output at exactly the moment demand peaked. French nuclear availability is, in my view, the single largest forecast error source in continental European power.



Chapter 04

To Build A Power Plant (Or A Data Centre) In Europe

In this chapter

- Merchant thermal new build is not financeable in Europe on spreads alone. This is why capacity tenders exist.
- Site selection in Europe is a queue-position problem before it is a land problem.
- The scarcest asset in the European power system is a site where the network has already been built to carry serious power. The connection right itself does not simply transfer.

Scoping

The American developer asks what kind of plant, then where. The European developer asks what the member state will pay for, then what kind of plant.

This is not a bureaucratic complaint. It is a financing reality. At current spreads and current EUA prices, a new European CCGT running a few hundred hours a year against a duck curve cannot cover its fixed costs from energy revenue. That is the missing money problem, and Europe has answered it almost everywhere with administrative procurement rather than scarcity pricing: capacity markets in Great Britain, Poland, Italy, Belgium, France and Ireland, and now a tender programme in Germany.

If you are building generation in Europe, your first question is which mechanism you are bidding into.

Site selection

Site selection anywhere begins with deconfliction: overlaying the candidate sites on everything that rules them out, which is mapping work before it is anything else. Europe adds two layers to the American version. Natura 2000, the EU's network of protected habitats, covers roughly 18% of the Union's land area, so the protected-species constraint is not a local obstacle but a continental filter. And under the Electricity Market Design reform, transmission and distribution operators must publish available connection capacity, which has produced a growing set of *hosting capacity maps*. Those maps are now the first document you open, before land.

Then the queue. Some numbers to calibrate:

- FLAP-D: seven to ten years average, up to thirteen in the worst primary markets
- Italy: by end 2024, Terna's connection queue held roughly 30 GW of data centre projects, about half of Italian peak demand, with about 80% of those requests filed in the preceding twelve months. By 31 August 2026 the same queue held 101.3 GW, close to twice the all-time Italian peak of 60.5 GW, and the clearest illustration in Europe of how fast a queue figure goes stale
- Ireland: roughly €5.8 billion of stranded data centre investment, meaning fully permitted projects on purchased land that cannot energise
- Great Britain: the contracted demand queue tripled from 41 GW to 125 GW in seven months before emergency reform

Europe's policy response arrived on 10 December 2025 as the **European Grids Package**: a revision of TEN-E, the regulation that designates cross-border energy infrastructure as projects of common interest and gives them faster permitting, plus a permitting acceleration directive, guidance on efficient and timely grid connections, and guidance on two-way CfD design. The connection guidance is the operationally important piece. It pushes member states away from first-come-first-served toward allocation based on project maturity and progress, and toward *flexible connection agreements*, where the user accepts curtailment when the network is constrained in exchange for connecting sooner. A Citizens Energy Package followed in March 2026. The Grids Package was flagged as a priority file for 2026, so expect movement.

Read that carefully if you are a large load. Flexible connection is Europe telling you that you can have your connection now if you accept interruptibility.

Which workload can accept that is worth getting right, because the intuitive answer is the wrong way round. A training run is a batch job. It checkpoints, it can be paused and resumed, and in a multi-site fleet it can be moved. Interrupting it costs utilisation on expensive hardware, which is a question of price rather than feasibility. Inference serves live users under latency commitments, and cannot be interrupted at all without failing over to somewhere else, which most European deployments cannot do. Google, which has more operating evidence here than anyone, has been signing demand response agreements aimed at exactly this, in the United States and with Elia in Belgium, and is explicit that what it cannot flex is the live-serving side.

So the curtailable workload is the one people assume is not. Chapter 10 adds the complication that makes this interesting rather than simple: the same training cluster that can most easily be switched off is also the one that moves in lockstep at second scale, which is what makes its behaviour hard for an operator to forecast and therefore hard to write a bound around. The workload most able to accept curtailment is the one hardest to promise anything about. Chapter 15 sets out where that leaves the contract.

Financing

Before the instruments, the numbers. The primer has so far discussed what a plant earns per MWh without saying what one costs to build, which makes every financing claim unverifiable.

The reference most useful for Europe is Fraunhofer ISE's levelised cost study, because it is denominated in euros, built on European assumptions, and unusually covers conventional plant alongside renewables. LCOE is the levelised cost of energy: total lifetime cost divided by total lifetime output, which is the number that makes technologies with different cost shapes comparable. Figures below are 2024, in real euros. Solar is quoted per kWp, the peak output under standard test conditions, and its yield in kWh per kWp a year; everything else is quoted per kW and in full load hours, the hours a year the plant would run at nameplate to produce what it actually produces. The two are not directly comparable, which is itself worth knowing when someone puts them in one table. Costs are given with the real weighted average cost of capital, or WACC, the study assumes for each technology. The WACC is the blended return the project has to earn to satisfy both its lenders and its equity, so it is the price of the money rather than the price of the kit.

Technology	Capital cost	LCOE	Annual yield	Real WACC
Utility-scale PV, above 1 MWp	€700 to €900 per kWp	€41 to €69/MWh	935 to 1,280 kWh per kWp	3.5%
Onshore wind	€1,300 to €1,900 per kW	€43 to €92/MWh	1,800 to 3,200	3.9%
Offshore wind	€2,200 to €3,400 per kW	€55 to €103/MWh	3,200 to 4,500	6.0%
CCGT	€900 to €1,300 per kW	€109 to €181/MWh	3,000 to 6,300	6.4%
Gas turbine, peaking	€450 to €700 per kW	€154 to €326/MWh	500 to 3,000	6.4%
Hard coal	€1,700 to €2,300 per kW	€173 to €293/MWh	3,000 to 5,200	6.8%
Battery, within a PV system	€400 to €1,000 per kWh	Not reported standalone	n/a	2.2 to 2.5%

Three things in that table are worth more than the rest of it.

The WACC column is the European story in one line. Offshore wind is financed at 6.0% and utility solar at 3.5%, while a gas turbine carries 6.4% and coal 6.8%. The cost of capital, not the cost of equipment, is what separates these technologies in Europe, and it is set by how firmly the revenue is contracted. That is why the contract-for-difference discussion later in this chapter is a financing discussion rather than a subsidy discussion.

The battery row looks like a counter-example and is not. It is costed here as storage attached to a PV system that already has a contracted revenue line, so it inherits that contract's risk profile rather than carrying merchant exposure of its own. A standalone merchant battery finances nothing like this.

A peaker is cheap to build and ruinous to run. At €450 to €700 per kW it is the least capital-intensive plant on the list, and at 500 to 3,000 full load hours its LCOE runs to €326/MWh. This is the missing money problem in capital terms: the asset the system most needs for scarcity hours is the one whose economics are worst without a capacity payment.

A CCGT at €900 to €1,300 per kW puts a 1 GW plant at roughly €1 billion of capital. Set that against the short-run marginal cost of about €138/MWh from Chapter 03, and against the spark spread that has to cover it.

One caveat on the source. The study is internally inconsistent in places, giving different PV LCOE and battery cost ranges in different sections. Use the headline table, not the body text, and do not treat any single figure as precise to the euro.

The European cost premium, quantified

The preface claimed European developers face worse revenue and higher costs than American peers. The International Renewable Energy Agency's 2025 cost survey, IRENA from here on, puts figures on the second half of that.

Globally in 2025, onshore wind reached a total installed cost of USD 976/kW, roughly €900/kW at the 2025 average rate, at an LCOE of USD 33/MWh, utility solar USD 667/kW at USD 44/MWh, and offshore wind USD 2,931/kW at USD 78/MWh. Utility-scale four-hour battery installed cost fell close to 30% in a single year, to around USD 140/kWh.

Against those global averages, **German solar came in at USD 65/MWh and German onshore wind at USD 52/MWh**. That is roughly a 50% premium on solar and a 58% premium on wind, against a global average that includes the cheapest markets on earth. IRENA notes that the United States and Europe retain a persistent though narrowing cost premium, without putting a figure of their own on it.

So the asymmetry in the preface is real and it compounds: European projects cost more to build, are financed at a higher cost of capital than the equipment cost alone would imply, and contract at offtake prices below what American projects achieve. None of those three facts is individually fatal. Together they are why so little gets built merchant.

Financing instruments

Same underwriting logic as the US, different instruments at the top of the stack.

- **Two-way CfD.** The state, or a state entity, pays the difference between a strike price and a reference market price, and claws back above it. This is now the EU's preferred support instrument for new low-carbon generation, with Commission design guidance published alongside the Grids Package.
- **Corporate PPA.** Pay-as-produced, baseload, pay-as-nominated, or hybrid. Volumes and structures are covered in Chapter 11.
- **Capacity payment.** Where a mechanism exists, this is the fixed revenue that makes the debt work.

There is a fourth structure that is less a financing instrument than the place this logic eventually arrives, and Belgium has just demonstrated it. On 30 April 2026 the Belgian State and ENGIE signed a letter of intent opening exclusive negotiations for the State to acquire the entire Belgian nuclear business of ENGIE and Electrabel: all seven reactors, the personnel, every nuclear subsidiary, and all associated assets and liabilities including decommissioning and dismantling obligations. Heads of terms were targeted for 1 October 2026. No price has been disclosed, and the parties agreed the transaction should not unduly affect ENGIE's overall financial position in either direction, which tells you the negotiation concerns the transfer of liability rather than the valuation of an asset.

Set that against the rest of this chapter. Merchant thermal does not finance in Europe at current spreads, so the state fills the gap through tenders and contracts for difference. Nuclear is the same argument carried to its conclusion: the private operator exits, and the state takes the asset and the liability together. Belgium is simultaneously targeting 4 GW of nuclear in its mix and has announced ambitions for a further 4 GW beyond that. France and the United Kingdom reached structurally similar positions by different routes.

Europe has effectively concluded that nuclear is a state balance-sheet item. The transaction is at letter-of-intent stage and may not complete on these terms, so treat the specific deal as indicative rather than settled. The direction of travel is not seriously in doubt.

The tender mechanism is less reliable than it looks

One qualification to the claim that Europe builds generation through tenders, because the mechanism has been failing in public.

Germany's offshore wind auction for the centrally pre-examined sites N-10.1 and N-10.2, 2.5 GW between them, **received no bids at all** by its deadline of 1 August 2025. Denmark's 3 GW tender had already drawn none in December 2024. Both used zero-subsidy negative bidding, where developers compete for the right to build with no revenue support and, in the German design, pay for the site itself.

Both countries are now moving to two-way contracts for difference. Denmark retendered with a capability-based CfD carrying a state payment cap of DKK 55.2 billion across at least 2.8 GW on three sites. Nine member states were already on CfDs, and the Netherlands has legislated the shift with first tenders expected around 2027.

The comparison that settles the argument is the United Kingdom's Allocation Round 7, announced in January 2026, which cleared **8,245 MW of offshore wind at £65.45/MWh in 2012 prices**, roughly £91.20 in 2024 money. Same technology, same North Sea, broadly the same supply chain, and a revenue stabilisation mechanism in place of merchant exposure.

The lesson is not that tenders fail. It is that auction design decides whether they work, and that merchant risk is unbearable even for the most mature and best-understood renewable technology in Europe. The Belgian nuclear transaction above reaches the same conclusion from the other extreme: where offshore wind could not be financed without a state-backed price, nuclear could not be financed at all, and the state took the asset itself.


European data centre PPA volumes actually *fell* from 4.2 GW in 2024 to 2.6 GW in 2025, even as capacity buildout accelerated, driven by offshore wind delays and by growing difficulty agreeing price points as capture rates, the average price a plant actually realises against the average market price, fall and negative price hours multiply. Spain leads on volume with more than 2 GW across twelve deals, anchored by Amazon's 596 MW hybrid wind and solar agreement, plus a 213 MW solar-plus-storage deal from Merlin Properties with Solaria on a reported forty-year term. The Netherlands is second, helped by Equinix's 250 MW nuclear PPA with ULC Energy.

Note the pattern in those two headline structures: a forty-year term and a nuclear offtake. Both are buyers reaching for something baseload-shaped because intermittent renewables alone do not match an always-on load. That is the direction of travel.

Construction

Same global turbine backlog: GE Vernova, Siemens Energy, Mitsubishi Power. Europe hosts two of the three OEMs and still cannot get slots.

The European-specific bottleneck is worse and less discussed: high-voltage transformers, HVDC converter stations, and subsea cable. Lead times on large power transformers have stretched to multiple years, and the European grid buildout, the offshore wind programme, and the data centre queue are all bidding for the same factories. If you are modelling a European connection date, the transformer is more likely to be your critical path than the turbine.



Chapter 05

Case Study: Sines

In this chapter

- Sines is a retired coal site redeveloped for AI compute. It reused the plant's seawater cooling; its grid capacity was awarded by the transmission operator in two tranches. The same move is being made in the US, at Homer City in Pennsylvania among others.
- The trade is explicit: you exchange queue risk for zonal price and interconnection risk.

Portugal closed the Sines coal plant in 2021. It was the country's largest, sitting on the Atlantic coast beside a deepwater port and a set of transatlantic subsea cable landings.

Start Campus is redeveloping the site, with Microsoft and Nscale expanding capacity there. The deployment runs from roughly 12,600 NVIDIA Blackwell Ultra GPUs starting in early 2026 toward more than 66,000 Rubin-generation GPUs from late 2027, on a campus permitted for 1.2 GW.

Why it works

The retired coal plant came with seawater cooling infrastructure, an industrial site already zoned and served, and a local authority that had just lost its largest employer. That combination collapses the seven-to-ten-year queue problem into a permitting and upgrade problem.

One correction to the way this case is usually told, because the usual telling is the reason people over-read it. Start Campus's own account of what it repurposed from the decommissioned plant is the seawater cooling system. It does not claim to have inherited the plant's grid connection, substation or electrical capacity. The campus's power was awarded by REN, the Portuguese transmission operator, in stages: an initial secured capacity of 495 MW and then an increase to 1.2 GW announced in August 2024, alongside a decree recognising the area as a zone of high demand. The brownfield site was the reason the award was possible on that timescale. It was not a substitute for the award. Add Iberian solar, some of the cheapest in Europe, and seawater cooling in a mild coastal climate.

This is the single most transferable insight in the European market. The scarcest asset in the European power system is not land, not gas supply, not turbines. It is an existing high-voltage connection at a site where the network has already been built to carry serious power. Europe has hundreds of them attached to retired coal plants, closing aluminium smelters, shrinking refineries, and paper mills that will not survive the decade. Whoever systematically buys those sites is buying a head start on the only thing the queue cannot manufacture.

The conditions that come with it

Having said that, the way this idea usually gets repeated is wrong, and the wrong version will lose you a year. The capacity right is the start of the negotiation, not the end of it. Four conditions decide whether a given site is actually convertible, and they should be screened before anyone signs anything.

A generation connection is not a load connection. The retired plant exported power. A data centre imports it. That is a different service, a different agreement and, in general, a new study, because the network was planned around power flowing out of that node rather than into it. Even the most optimistic published treatment of coal-to-data-centre conversion puts it no higher than that the existing interconnection point may make the study simpler and the upgrades smaller. Nobody credible describes it as automatic.

The equipment faces the wrong way. A generating station's transformers step voltage up toward the network. A load needs the opposite. Protection settings assume fault current arriving from a particular direction. The reactive power support that the plant's own machines used to provide to the local network leaves with the plant, and the new load does not replace it. None of these is insurmountable and all of them cost time and money that the headline version of this idea does not budget for.

The right may already have gone. Unused capacity is increasingly reclaimable rather than bankable. The Netherlands has codified exactly this: a decision of April 2024 by the Dutch regulator, on unused transport rights and known as GOTORK, allows operators to cut back contracted capacity that a connected party cannot show it needs, on medium, high and extra-high voltage connections in congestion areas. A site that has been dark since 2021 should be assumed to have lost something until the operator confirms otherwise, and that confirmation is the first phone call, not the last.

Somebody else may own the terminal. The high-voltage switchgear at the site boundary, meaning the circuit breakers and isolators through which the site connects, is often the network operator's asset rather than the plant owner's, and the land purchase does not come with it.

A Dutch judgment in April 2026 shows what these conditions look like when they bite. A developer had sought a 70 MW connection for a data centre near Schiphol, applying in 2021, reaching a preliminary agreement and making advance payments in 2023, and being told in late 2025 that constraints would block it after all. It went to court to force the work to resume. The claim was dismissed. The court found the regional network genuinely full, that no definitive connection agreement had ever been concluded, and that a professional party aware of the current energy landscape should not have assumed the connection would proceed. Five years, a preliminary agreement and money paid bought no enforceable capacity.

Sines worked, and it is the case most often cited as proof that the right transfers, which the case study above sets out why it does not show. What the site actually supplied was cooling water, zoned industrial land and a municipality that wanted the project. The electrical capacity was awarded, in two tranches, by the transmission operator, against a national government that wanted the project. The brownfield was necessary and it was not sufficient, and that is the correct reading of every example in this category. Screen for the asset, then price the four conditions above, and treat any seller who tells you the right simply transfers as someone who has not asked.

Why it is risky

Iberia is also the zone with Europe's deepest capture rate problem. In the thirty days to 13 July 2026, EU bidding zones recorded roughly 399 hours of negative day-ahead prices, with Spain alone accounting for 85 of them. That is excellent for a buyer taking spot exposure and terrible for the PPA counterparty trying to finance the generation.

And Iberia has thin interconnection to the rest of the continent, roughly 2.8 GW across the Pyrenees against an EU target several times higher. On 28 April 2025 the peninsula went dark. The expert panel convened by ENTSO-E, the association of Europe's transmission system operators (the TSOs, who own and run the high-voltage grid and keep the system in balance), published its final report attributing the event to a combination of interacting factors: oscillations, gaps in voltage and reactive power control, differing voltage regulation practices, rapid output reductions and generator disconnections in Spain, and uneven stabilisation capabilities across the system. Renewable penetration is notably not the headline finding. Voltage control gets the bulk of the recommendations.

So the Sines trade is legible: you solve the queue, and you accept concentration in a weakly interconnected zone with a demonstrated stability incident and the continent's worst negative price problem. Whether that is a good trade depends entirely on your hedge structure, which is Chapter 11.



Where The Demand Actually Went

In this chapter

- European electricity demand is still below its 2021 level. The growth narrative and the realised data point in opposite directions.
- What Europe has is composition change at roughly flat volume, which is a locational problem rather than a volume one.
- European data centre demand is small in aggregate and destabilising locally.
- The EU is simultaneously restricting connections and tendering seven gigafactories. Nobody has reconciled those.

Start with what did not happen

Everything that follows in this chapter describes demand arriving. Before any of it, the realised figure.

European electricity demand in 2025 was still around **7% below its 2021 level**. EU27 demand grew 1% in 2024 and 2% in the first half of 2025 against the same period a year earlier, with 2025 coming in broadly flat on 2024. More than half the shortfall against pre-crisis levels is industrial. German industrial power consumption fell 13% between 2021 and 2023, and German industrial production fell a further 4% year on year through 2024.

So the honest position is uncomfortable for the narrative. Europe has spent four years discussing a demand surge while consuming less electricity than it did before the crisis.

What that actually means

Two readings are available and only one survives contact with the rest of the evidence.


The first is that demand growth is a story the industry tells itself. That reading is wrong, because the forward commitments are real. The connection queues later in this chapter are not hypothetical, Belgium's transmission operator has stated plainly that data centre and battery growth exceeded its own expectations, and the electrification of transport and heating is a legal obligation rather than a forecast.

The second reading is that Europe is not experiencing demand growth at all. It is experiencing **composition change at roughly flat volume**. Industrial load has been destroyed across the chemicals, aluminium and fertiliser belt, a good deal of it permanently, while new load arrives as data centres,

charging infrastructure and heat pumps. The aggregate barely moves. Everything underneath it moves enormously.

The second reading is the one the rest of this document rests on, and it is worth stating without hedging.

Europe does not have an electricity shortage. It has electricity in the wrong places.



Europe does not have an electricity shortage. It has electricity in the wrong places. Aggregate volume barely moves while everything underneath it moves enormously, and the load that left was connected at high voltage in a place the load that is arriving does not want to be. That is a locational problem, not a volume one.

A smelter that closed in 2022 was connected at high voltage, drew a flat industrial profile, and sat where it sat for forty years because that is where the connection was. The data centre that replaces it on the national balance sheet is somewhere else entirely, frequently at a different voltage level, and in a queue. The heat pumps and chargers replacing the remainder are spread across a low-voltage network that was never built for them and, as Chapter 12 shows, largely cannot see them.

Aggregate demand below 2021, and specific nodes that cannot connect anybody. That is the definition of a locational problem rather than a volume problem, and it is why Part 3 of this document is about the distribution network rather than about generation capacity.

One risk this creates

The gap between realised and projected demand is itself something to price, and almost nobody is pricing it.

Network investment, capacity mechanisms and connection queue reform are all being sized against projections that have not yet shown up in the aggregate figures. If industrial load continues to leave faster than new load connects, Europe will have built for a demand level it does not reach, and the stranded cost lands in exactly the network tariffs Chapter 12 discusses.

That is not a prediction. It is the risk sitting on the other side of a forecast that most of the sector is treating as settled.


The numbers

European data centres consumed roughly 96 TWh in 2024, about 3% of total demand. Continentally, manageable. Locally, not: the regulator’s own analysis of census-year metering data put data centres at roughly half of all metered electricity consumption in the Dublin and Meath region in 2024, against a national Irish share of 22%, and Amsterdam, London and Frankfurt run somewhere between roughly a third and over 40% of local demand. Ireland is the extreme case because 97% of its data centre capacity sits in one city region.

Forward projections: European data centre capacity growing from 16 GW in 2024 to around 36 GW by 2030, with about a third of that added in the final two years. The European Data Centre Association puts required investment at €176 billion cumulatively from 2026 to 2031, roughly €25 to €26 billion a year, on a 17% compound growth rate in European IT power demand through 2031.

What it costs to run

Take a 100 MW campus drawing roughly 876,000 MWh a year, which is 100 MW held flat for 8,760 hours and counts total facility draw including cooling, not IT load. All-in annual energy cost runs from roughly €38 to €55 million in the Nordics to roughly €118 to €149 million in Dublin. Taking the midpoint of each range, that is roughly €870 million over ten years for the same hardware, and the gap runs from about €630 million to €1.1 billion depending on which ends you pair.



The gap is not the wholesale price. It is grid use-of-system charges of roughly 4 to 28 euros a megawatt hour depending on market, demand charges at 15 to 25 per cent of the electricity cost, renewable levies and the cooling multiplier. The most common European planning error is underwriting a site on the wholesale curve and discovering the network tariff afterwards.

The gap is not wholesale price alone. It is grid use-of-system charges of roughly €4 to €28/MWh depending on market, demand charges, which bill the highest power the site ever draws rather than the energy it consumes, typically 15 to 25% of the electricity cost, renewable levies, and the PUE multiplier on cooling. The most common planning error in Europe is underwriting a site on the wholesale curve and discovering the network tariff afterwards.

Government responses

Broadly restrictive, with a large and unresolved contradiction sitting on top.

The restrictive layer starts with Ireland. Under the large energy user connection policy decided in December 2025, a data centre at or above 10 MVA, megavolt-amperes being the apparent power a connection is rated for, must deliver matching dispatchable generation or storage before it can ramp to full capacity, and any data centre at or above 1 MVA must cover 80% of annual demand with additional Irish renewables on a six-year glide path. Between 1 and 10 MVA the requirement is an onsite generating unit covering the site's full contracted capacity, discounted for the share of it the system can actually rely on. Two details decide whether this applies to you: the regime governs applications made after the decision was published, and anything already in process continues under the 2021 direction. Alongside it sit the Dutch national decree limiting hyperscale development outside two northern municipalities, Amsterdam's threshold on projects of 70 MW and above, and severe constraint in Frankfurt. On the regulatory side, the Energy Efficiency Directive requires data centres above 500 kW to report power usage effectiveness (PUE), which is total facility energy divided by the energy reaching the computing equipment, along with renewable energy factor and energy reuse factor, annually, and Germany's Energieeffizienzgesetz goes further with mandatory PUE targets and waste heat assessments above 1 MW.

The expansionary layer arrived on 30 July 2026, when the EuroHPC Joint Undertaking opened the **AI Gigafactories call**: up to seven facilities, backed by as much as €10 billion in combined EU and national funding, intended to mobilise at least €20 billion more in private capital. Each site must hold at least 100,000 advanced AI processors, roughly four times the most powerful of the EU's existing AI factories, the smaller national supercomputing sites the gigafactories are meant to supersede. AMD, NVIDIA and Qualcomm have signed letters of intent to supply. The tender closes 12 November 2026, bidders are announced in early 2027, and selected facilities must be operational within eighteen months of contract signature.

Do the arithmetic. Contracts signed in 2027, operational by 2028 or 2029, at 100,000-accelerator scale, in a system quoting seven-to-ten-year connection queues. Either the gigafactory programme is an implicit queue-jumping programme, or the timeline is fiction. My read is that it is the former and nobody wants to say so, because saying so means admitting that connection capacity is allocated politically rather than procedurally. Which it now is.

Off-grid and behind the meter

The American play is: build your own generation on site, skip the interconnection queue, accept the operational complexity. In Europe this is materially harder.

Islanded or partially islanded operation runs into national grid codes and connection conditions that were not written with 500 MW private thermal generation in mind. More importantly, self-supply exemptions from network charges and levies have been narrowed in several member states precisely because they erode the socialised cost base that funds the grid. If you build behind the meter to escape network tariffs, you are in a policy fight, not an engineering project.

The genuinely European versions of the workaround are three:

1. **Go where the queue is short.** Finland, Poland, Greece and increasingly the Baltics. Check this list before you use it, because it decays fast: Italy was on it in 2024 and its queue has since passed 100 GW, Iberian demand-side access has tightened, and Norway has turned politically hostile to the load even where the network is clear. Capital has already started reallocating on exactly this basis, with power-secured sites commanding a premium and sites without credible power access treated as speculative regardless of location.
2. **Take a flexible connection agreement** and design the workload around curtailment. More viable for batch training than for live inference, which is the opposite of the common assumption and is set out in Chapter 04.
3. **Buy the site, then price the connection.** Acquire a retired thermal or shrinking industrial site where the network was built to carry serious power. Chapter 05 sets out the four conditions that decide whether the capacity is actually recoverable, and Sines shows why the site is necessary and not sufficient.

Propensity to buy

The demand signal is thinner in Europe than the headlines suggest. PPA volumes fell year on year. Some hyperscalers have moderated clean energy commitments while racing to secure power at all. And European offtakers have shown no willingness to pay US-style prices.

That is the crux. A European developer looking at a data centre anchor tenant is having a €50 to €70/MWh conversation. At €138/MWh short-run marginal cost for a CCGT, no anchor tenant is underwriting new European thermal generation on energy revenue. The state is.

One clarification before leaving these prices, because getting it wrong leads to a whole category of mistaken assumptions about who is buying what. In a hyperscale colocation structure the tenant usually procures its own power. The hyperscaler signs the PPA, often directly with a generator and often in a different country from the campus, and the landlord passes supply through and bills network charges and levies on top. The developer is therefore not selling power to the tenant and is not the counterparty to the offtake. The developer is selling a connection with a defined firmness, a delivery date and a cost pass-through structure.

That matters for everything downstream. The prices in this chapter tell you what the tenant can get in the wholesale and PPA market, which is what it will compare your campus against. They do not tell you what a landlord earns, because the landlord is not in that trade. AI labs and the neocloud operators are the exception: several of them take space and power together, which is why their deals look more like the American ones. For the large established hyperscalers, the power conversation and the real estate conversation happen with different people on different contracts.

One further German development belongs here, because it changes what a multi-tenant campus is in law.

Germany long had a category called the *Kundenanlage*, a customer installation that distributes power inside a site without its operator becoming a regulated network. It carried real advantages. In November 2024 the Court of Justice held that the German definition was too wide, and in May 2025 the Federal Court of Justice read it down accordingly: an installation that distributes electricity and sells it to customers is a distribution network, and the additional German criteria cannot rescue it. The practical dividing line is the sale. Supplying yourself inside one undertaking is still safe. On-selling power to tenants is not.

A colocation operator that bills tenants for power is squarely in the second category. The legislature has bought time rather than solved it: installations already connected on or before 23 December 2025 are exempt from network regulation until the start of 2029. Anything connected after that date has no transitional cover and no settled framework, because the permanent regime has not been written. For a campus now in design this is a live structuring question, and the alternatives, such as operating as a closed distribution system, carry obligations of their own.

Which brings us to Germany.

An unsolicited recommendation

Five, offered as my view rather than as analysis.

Make connection capacity tradeable. A large share of European queue volume is speculative option-holding. Britain is the case with public numbers behind it, where a generation queue reached 738 GW against a system that will never need a fraction of it, and both the system operator and the regulator have moved to commitment fees and maturity gates in response. Elsewhere the scale is inferred rather than measured, which is itself part of the problem. Europe is trying to clear speculation with paperwork, milestone tests and maturity criteria. The instrument that clears speculation is a price. Allow connection rights to be transferred and priced, with hard use-it-or-lose-it milestones, and the queue will shorten faster than any permitting directive will shorten it.

Germany has already moved part of the way, and it is worth knowing what it chose. From April 2026 the four transmission operators allocate new transmission-level connections for large consumers, storage and data centres through a maturity procedure rather than by order of arrival. Applications are scored against verifiable evidence of land, permits, technical design and the applicant's own capacity to build, a non-refundable application fee of €50,000 is payable per request, and a realisation security of €1,500 per megawatt is credited against the connection contribution if the project proceeds. There is no capacity threshold, which distinguishes it from the separate regime for large generating plant. Two details matter operationally: applications are batched rather than rolling, so missing a cycle means waiting for the next one, and part of the score rewards sharing a connection point. Co-locating storage or a second load on the same node is the one lever an applicant can pull to improve its own ranking. Pure generation stays outside it, under the renewables and plant-connection rules.

That is a hybrid of the two instruments this section has been contrasting. It clears speculation with evidence, as a maturity test does, but it also attaches a price to holding a place, which is the part that actually bites. It stops short of letting anyone sell a position. Whether a deposit of that size deters an option-holder with a serious balance sheet is not really an open question. Britain, facing the same problem, has consulted on commitment fees in the region of £237,500 to £712,500 per megawatt, which is two to three orders of magnitude above the German figure. A deposit at the German level is a filing discipline rather than a price on holding a place, and the maturity scoring is doing most of the work.

The tradeable version has an objection to beat, and it is a good one, so it is worth stating properly rather than leaving it for the reader to raise. Making a queue position tradeable hands a windfall to whoever currently holds one, and the whole premise above is that a large share of those holders are speculators. Europe would be paying the people who caused the congestion to go away. That is not a detail regulators have overlooked. It is most of the reason they have not done it.

Two things answer it. The windfall is a one-off transfer, while the queue distortion is a recurring cost paid by every project behind it and by every consumer waiting for the capacity, so the comparison is a single payment against an annual loss. And the windfall is capturable if anyone wants to capture it: auction released capacity rather than grandfathering it, tax the transfer, or make the right saleable only back to the operator at a regulated price. A design that is politically difficult is a different problem from one that does not work.

There is a second objection that is harder, and it is specific to the distribution network, the lower-voltage layer beneath the transmission grid that Part 3 is about. Above the primary substation a transfer is close to a contractual matter. Below it, connection capacity is entangled with connection charging and with the statutory duty of non-discriminatory access, so making it tradeable is a legislative programme rather than a pricing tweak. The recommendation is therefore much more realistic at transmission level than at distribution level, and anyone proposing it should say which one they mean.

Screen brownfield sites, then price the four conditions. For investors, the acquisition screen is retired thermal sites and heavy industrial sites where the network was built for a large load. What you are buying is a head start, not a right: Chapter 05 sets out why a generation connection is not a load connection, why the equipment faces the wrong way, how unused capacity is reclaimed, and who owns the terminal. Price those four before you price the site. The category is real and it is being bought; whether it is cheap is not a question this document can answer, because no European transaction has published a price per megawatt of connection.

Procure voltage support, do not mandate it. The Iberian report's centre of gravity is voltage and reactive power control, not renewable share. Grid-forming inverters, synchronous condensers and reactive capability are services with a cost. They should be bought in a market, not imposed as an unremunerated connection condition, or they will be delivered to specification and no further.

Two developments since make that concrete rather than theoretical. Spain wrote **€615 million for e-STATCOMs and controlled reactances** into the third modification of its statutory network plan, agreed in July 2026, which is a member state putting a figure on exactly this and funding it through the planning instrument rather than a connection condition. And the four Nordic transmission operators

published a joint position on grid-forming converter requirements in 2025, which is the clearest forward statement of system operation philosophy anywhere in Europe and the constructive answer to what the Iberian panel found.

If Germany will not split its bidding zone, price location some other way. The economics are covered in Chapter 08 and they point one way, though by a narrower margin than the advocates imply: the annual congestion saving is on the order of 1% of German wholesale market value, and the published estimates carry neither the transition cost nor the higher hedging cost of a fragmented reference curve. If the political answer stays no, then locational network tariffs and locational connection charges are the second best, and they should be pushed hard rather than treated as a consolation prize.

Say the quiet part out loud. Somani's recommendation for the US is to temporarily relax resistance to dirtier, faster generation while building startup-speed development capability. Europe will not do this, and I do not think it should be asked to abandon its Emissions Trading System to chase inference workloads. But a Europe that will not build fast, dispatchable capacity is choosing to import the inference instead. That may well be the right choice. It should be made explicitly, with the tradeoff stated, rather than arrived at through seven years of queue attrition while ministers give speeches about digital sovereignty.

Part 2: How To Trade European Power



Chapter 07

Power Markets in Europe

In this chapter

- Europe has no single institution running a regional market end to end, the way an American independent system operator does. Market operation, network operation and regulation sit in separate bodies.
- The traded unit is the bidding zone, and zones are mostly national borders rather than congestion boundaries.
- Cross-border capacity in the Core region is allocated flow-based, not by simple transfer limits.

Institutional structure

A US regional market is run by a single institution. An independent system operator, or ISO, plans the network, runs the auction, dispatches the plants and settles the money, all under one roof; PJM, covering thirteen states and the District of Columbia, is the largest. Europe has no equivalent. Those same functions are split across:

- **Transmission system operators** (TenneT, Amprion, 50Hertz, TransnetBW, RTE, Elia, Terna, REE, Statnett, Svenska kraftnät and others), who own and operate the grid and manage congestion. Great Britain splits the two roles: NESO operates the system and the transmission owners hold the assets
- **Nominated electricity market operators**, in practice the power exchanges (EPEX SPOT, Nord Pool, OMIE, GME, and others), who collect orders
- **Regional coordination centres**, who compute the cross-border capacity domains
- **National regulators and ACER** at EU level, who set the rules
- **ENTSO-E**, the TSO association, who runs the studies

The day-ahead clearing itself is done once for the entire coupled region by a shared algorithm, EUPHEMIA, under **Single Day-Ahead Coupling (SDAC)**. Gate closure, the moment after which no further bids are accepted, is 12:00 CET for delivery the next day. Intraday runs continuously under SIDC with intraday auctions layered in. Balancing is being harmonised through the PICASSO platform for automatic frequency restoration reserve and MARI for manual.

The structural consequence: **the entity that sets the price and the entity that keeps the lights on are different entities with different incentives.** Congestion inside a zone is therefore not a price signal to market participants. It is an operational cost that the TSO incurs and recovers through network tariffs. Hold that thought for Chapter 08.

The zones

A partial tour of what matters:

DE-LU. The largest and most liquid zone in Europe, and the default answer if you had to guess where a European power trader operates. Enormous renewable penetration, a north-south transmission constraint, and no capacity market. Chapter 08.

France. Nuclear-dominated, the continent's swing exporter when availability is good and a large importer when it is not. Availability is the whole trade.

The contracting framework changed at the end of 2025 and anyone pricing French supply needs to know it, because the thing that replaced ARENH is not a replacement in any sense a buyer would recognise.

The ARENH, the regulated-access scheme that gave suppliers a claim on part of EDF's existing nuclear output, obliged EDF to sell up to 100 TWh a year at €42/MWh. It was heavily oversubscribed to the end: the final window drew requests of about 135 TWh against the cap, so buyers received roughly 74% of what they asked for. It expired on 31 December 2025.

What took its place, the Versement Nucléaire Universel, runs in the opposite direction. It is not an entitlement to volume at a price. It is a tax on EDF's nuclear revenue, 50% of revenue above €78/MWh and 90% above €110/MWh, redistributed to all final consumers as a reduction on the bill. The design has an obvious property that is worth stating plainly: it does nothing at all unless nuclear revenue exceeds the lower threshold. For 2026 the regulator estimated that revenue at €66.08/MWh, which is €12 below the trigger, and set the unit tariff at zero. The scheme's first year delivers nothing.

So from 2026 a large French buyer negotiates. The instrument that has emerged for industrial-scale offtake is a bilateral allocation contract over ten to fifteen years, under which the buyer takes a share of the existing nuclear fleet's output and, with it, a share of the fleet's costs and risks. Reported terms sit in the region of €65 to €70/MWh, against the €42 that ARENH guaranteed. There is already a data centre precedent: EDF signed a twelve-year contract with Data4 in September 2025 for 40 MW on that basis.

Read the change properly. The headline is a move from roughly €42 to roughly €65 to €70. The structural change is larger than the price: a regulated entitlement has been replaced by a negotiated contract in which the buyer carries part of the operational and investment risk of a nuclear fleet. That is a different thing to put in front of a credit committee, and it is why French industrial consumers objected to it rather than to the number.

The Netherlands and Belgium. Gas-heavy, tightly coupled to Germany, both with significant offshore wind. Belgium runs a capacity remuneration mechanism that has become the reference design others cite.

Italy. Internally zonal (NORD, CNOR, CSUD, SUD, CALA, SICI, SARD), which makes it the closest thing continental Europe has to a locational market. Structurally the most expensive large market, gas-set, with a large battery procurement programme underway and a data centre connection queue that passed 100 GW in 2026.

Iberia (MIBEL). Spain and Portugal clear together. Cheapest solar in Europe, deepest negative price problem, thinnest interconnection, one blackout.

The Nordics. Norway (NO1 to NO5), Sweden (SE1 to SE4), Denmark (DK1 and DK2), Finland. Hydro-dominated in Norway and northern Sweden, with a north-to-south internal constraint that produces enormous intra-country price spreads. Traded against a **system price** with area differentials, which is a structure that does not exist anywhere else and which I will return to in Chapter 11 because it is where people lose money.

Poland. Coal-dominated, capacity market, rapidly building renewables, and increasingly a destination for data centre developers who cannot get connected further west.

Great Britain. Left SDAC after Brexit and now couples to the continent through less efficient arrangements on the interconnectors. National pricing retained after the Review of Electricity Market Arrangements considered and rejected zonal pricing, with locational signals to be pursued through network charging reform instead. GB runs a capacity market and a contracts-for-difference regime that predates and informed the EU's.



Chapter 08

Case Study: Germany

In this chapter

- DE-LU is an energy-only market with a structural internal congestion it is not allowed to price.
- The cost of that decision shows up as congestion management, redispatch plus the compensation paid for curtailed renewables, which has run between €2.8 and €4.2 billion a year since 2022.
- Germany's answer to missing money is not a scarcity adder. It is a state tender.
- The Belgian mechanism Germany proposes to copy is specified below, along with why storage does not substitute for firm capacity in it.

The closest North American comparison is Alberta, which runs an energy-only market, meaning generators are paid for electricity delivered and nothing for merely being available, alongside a large and growing renewable fleet. Germany is the same experiment at roughly ten times the scale, with one additional twist: Germany refuses to let its internal congestion appear in the price at all.

The setup

Wind is in the north. Industrial load is in the south. Nuclear is gone. Coal is going. Wind and solar supplied 30% of EU electricity in 2025, overtaking fossil generation at 29% for the first time, and Germany is well above that average.

The transmission capacity to move northern wind to southern load does not exist in sufficient quantity, and building it has taken longer than anyone planned.

In a nodal market, that constraint would bind and prices would separate: cheap in the north, expensive in the south, with the difference telling everybody exactly where to build generation, storage and load. That is what a transmission constraint does in a nodal market. When the line between two points fills up, the two points stop sharing a price, and the gap between them is the signal.

Germany is a single bidding zone. The constraint cannot bind on price.

So it binds on cost instead

The mechanism is **redispatch**. The TSO clears the market as if the grid were unconstrained, then looks at the resulting physical flows, discovers they are infeasible, and instructs northern wind down and southern gas up, compensating both. German congestion management, which covers redispatch and the curtailment of renewable output when the network cannot carry it, cost about €4.2 billion in 2022, €3.3

billion in 2023 and €2.8 billion in 2024. The 2022 peak was driven by fuel prices rather than by more congestion. Treat the published volumes with care: the regulator reports redispatch-adjusted feed-in and renewable curtailment as overlapping categories rather than one additive total.

Every euro of that is socialised through network charges. None of it is a price signal to anyone deciding where to site a 300 MW load.

The bidding zone fight

Europe has a formal process for this. Under the Electricity Regulation, bidding zone borders are supposed to follow structural congestion, and ACER mandated a **Bidding Zone Review** to test alternatives.

The TSOs published their study on 28 April 2025 across fourteen configurations. Every German-Luxembourg split configuration came out more economically efficient than the status quo, ranging from about €251 million to €339 million for the 2025 target year, with a five-zone split performing best. ACER then reviewed the study and concluded the TSOs had understated the gains, putting them at roughly €450 to €540 million a year, and criticising an unrealistic assumption about how well remedial actions would be coordinated.

Berlin's answer is no. German politicians and grid operators oppose the split on grounds of liquidity fragmentation, industrial cohesion, new hedging basis risk, and transition cost. Those are real arguments, not stupid ones. The liquidity concern in particular is genuine: DE-LU is the reference contract for continental European power, and splitting it fragments the most-traded curve on the continent.

They are also losing on arithmetic. And EU law makes it very difficult to impose a split on a member state that does not want one, so the arithmetic will probably keep losing.

Missing money, German edition

Alberta solves missing money with scarcity rent under a C\$1,000/MWh cap. Texas solves it with an explicit price adder tied to the probability of lost load, up to \$5,000/MWh.

Germany's political tolerance for a €500/MWh print is approximately zero. So Germany cannot solve missing money with a price. It solves it with procurement.

The **Strom-Versorgungssicherheits- und Kapazitätengesetz (StromVKG)** passed the Bundestag on 9 July 2026 and cleared the Bundesrat the following day, implementing the long-delayed power plant strategy. This is no longer a proposal: the regulator opened the first tender on 21 July and bids closed on 8 September, with awards due by early November. The design:

- Two tender rounds in 2026, the first closed on **8 September** and the second set for **29 December**, for a combined **9 GW** of new “long-term capacity,” 4.5 GW each, with a ceiling price of €244,000 per MW of derated capacity per year in the first round, where long-term, confusingly, describes the

asset's endurance rather than its contract: it means capable of running at full output continuously for at least ten hours

- A further **2 GW** without the long-term criterion on **18 May 2027**, which opens the door to batteries
- Further tenders in 2027 and 2029, including hydrogen conversion tenders, together intended to cover Germany's full firm capacity requirement for 2031
- A broader technology-neutral **capacity market** with decisions to be taken in 2027 and operation from **2032**, with a centralised Belgian-style mechanism under discussion

Plants are to be hydrogen-ready with a path to decarbonisation by 2045.

The honest reading: Germany has concluded that an energy-only market with 60%-plus variable renewables does not finance the firm capacity it needs, and has stopped pretending otherwise. Alberta reached the same conclusion from the other side of the argument: it started as the most committed energy-only market in North America and has been adding capacity-style instruments rather than removing them.

What Germany would actually be copying

Since the Belgian mechanism is the stated template, it is worth knowing what is in it rather than taking the reference on trust.

Belgium's capacity remuneration mechanism rests on a reliability standard: a loss of load expectation of no more than three hours a year. Loss of load expectation is the number of hours in an average year during which available supply is expected to fall short of demand somewhere in the system, averaged across simulated weather and outage years rather than observed in any one of them. That figure is what "adequate" means in Belgian law, and every auction volume is calibrated back to it. Three auctions are held for each delivery period, at four years, two years and one year ahead, which lets the system correct a volume error twice before delivery rather than discovering it once. The participation threshold is 1 MW of derated capacity with aggregation permitted, so demand response and small storage bid alongside a CCGT. And projects requiring significant investment can apply for multi-year contracts of three, eight or fifteen years, graded by capital intensity.

That last provision is the one that matters for Germany, because it is the part that actually finances construction. A one-year capacity payment reduces volatility. A fifteen-year contract underwrites a loan. Nine European countries had market-wide capacity mechanisms in place as of 2023, and the design differences between them are mostly differences in contract tenor. Spain has since joined them, through Orden TED/966/2026 of 15 September 2026.

The scale of what the Belgian mechanism is buying is instructive as well. Without support, roughly 1,600 MW of existing Belgian thermal capacity, mostly older units, is at risk of leaving the market over the coming decade, rising to as much as 2,400 MW if a further nuclear lifetime extension goes ahead and depresses running hours. The mechanism is not principally procuring new plant. It is paying to stop the existing fleet closing, which is the missing money problem in its purest observable form.

Derating, and why storage is not firm capacity

One finding from the Belgian adequacy work deserves to be read by anyone who assumes batteries dissolve this problem.

A storage asset's contribution to adequacy is counted after derating, which is the share of nameplate that can be relied on during a scarcity event. In 2028, roughly 400 MW of new Belgian battery capacity is required to close a 200 MW adequacy gap, a derating factor of about 50%. By 2034, an additional 8,300 MW of batteries is assessed as contributing only around 2,200 MW of effective capacity, a derating factor of roughly 25%.

The mechanism is saturation. A four-hour battery covers a four-hour scarcity event, so once storage is deployed at scale it flattens the short peaks and what remains are the longer events it cannot ride through. Each additional unit therefore contributes less to the metric that actually binds, which is the reverse of how storage is usually discussed in policy documents.

None of that makes batteries a poor investment, and Chapter 10 covers the arbitrage case on its own terms. It does mean storage and firm capacity are not substitutes on the adequacy ledger, and a capacity market that treats them as interchangeable will systematically under-procure.

The pattern, not the pair

Belgium and Germany on their own could be two countries sharing a problem. Spain makes it a pattern.

Red Eléctrica's national coverage analysis of October 2025 found a loss of load expectation of **4.08 hours a year in 2028** and 2.41 hours in 2030, against a Spanish reliability standard of 1.5 hours. It identified roughly **3,524 MW of gas turbines that were economically unviable** on energy revenue alone, and concluded that the energy market would not deliver adequate security of supply by itself. Eleven months later the capacity market existed.

That is the same sequence three times over: an energy-only market, an adequacy study showing the standard will be breached, a finding that existing thermal cannot cover its costs, and a capacity mechanism. The Netherlands is currently at step two, with TenneT's 2026 supply security monitor reporting that the country breaches its own 4 hour standard from 2030 and reaches **46 hours a year by 2035** under high demand.

Energy-only markets with deep renewable penetration do not finance firm capacity. Europe has now run that experiment in enough jurisdictions that treating it as an open question is a choice rather than a judgement.

What capacity actually clears at

A capacity mechanism without a price attached is a brochure, so here is the price, and the direction of travel in it matters more than the level.

Great Britain runs the most liquid of these markets and publishes everything. Its four-year-ahead auction for the 2029/30 delivery year cleared at **£27.10/kW/year** for 40.1 GW of derated capacity. The two auctions before it cleared in the £60 to £65 range, and the auction for 2026/27 delivery cleared at £63/kW/year, which at the time was an all-time high.

Set that against Ireland, which ran its four-year-ahead auction for the same 2029/30 delivery year at **€135,500/MW/year**, roughly four times the British price. Belgium's four-year-ahead result for 2029/30 came in near €27,300/MW/year and Italy's most recent auction cleared at €47,000/MW/year. Resist the urge to read that spread as a pure scarcity ranking. These mechanisms do not price the same product: some clear uniformly and some are pay-as-bid, the delivery bases differ, and each market derates capacity by its own method. The spread is real, but a meaningful part of it is design rather than adequacy.

One detail inside the Irish result is worth more than the headline, because it is the locational signal this document keeps saying Europe does not have. The clearing price is not what most Irish capacity is actually paid. Additional capacity was procured to satisfy constraints in Northern Ireland, in Ireland and specifically in Greater Dublin, and those awards are paid at their own offer prices, which pushed the average awarded price to around €151,000/MW/year with individual units well above that. Ireland is paying a located premium for firm capacity in the places where its network is tight. It is doing so inside a capacity mechanism rather than through the energy price, which is exactly the workaround a zonal market reaches for when the energy price cannot express location.

That British halving, from £63 to £27.10 in a single year, is not a sign that adequacy improved. The demand target was cut from 43.7 GW to 39.4 GW, and a rule change lowered the capital threshold for three-year refurbishment agreements from £170/kW to £65/kW, which let around 14 GW of existing gas plant into the auction on refurbishment terms. At £27 almost nobody wanted a multi-year commitment: roughly 11 GW took one-year agreements against about 1.5 GW on three-year terms, and no new-build gas cleared at all.

Read that as a warning about the instrument rather than about Britain. A capacity price is set by an administratively chosen demand curve and by eligibility rules, which means it is a policy output at least as much as a scarcity signal. If you are underwriting an asset on capacity revenue anywhere in Europe, the sensitivity that matters is not the clearing price you model. It is the possibility that the regulator moves the curve or widens the eligible pool, which it can do between auctions, and which halved the British price in a single year.

What the system is actually insuring against

Capacity mechanisms exist because of a weather event, and it is worth knowing its real shape rather than invoking it as a word.

The German term that has entered English is *Dunkelflaute*, a period of simultaneously low wind and low solar output. The clean operational statistics come from a 2021 study of the North and Baltic Sea region. Events lasting more than 24 hours occur almost exclusively in November, December and

January, accounting for 50 to 100 hours in each of those months, with individual countries seeing roughly **2 to 10 such events a year** and some running three to five days. Interconnection helps substantially: for a single country the share of hours in such an event runs between roughly 3 and 9% depending which country, while across eleven interconnected European countries the figure converges to around 3.5%. Interconnection does not help every country equally: it pulls the worst-exposed down toward the European figure and leaves the best-placed roughly where they were.

The tail is considerably worse than those figures suggest. A 2026 study in Communications Earth and Environment found the most severe pan-European renewable drought on record, in the winter of 1996/97, ran **55 days even assuming perfect interconnection**. Those two descriptions are measuring different things and the gap between them is large enough to be worth stating plainly. The two-to-ten-events figure counts acute events, days when wind and solar together fall below a low output threshold. The duration figures below count something else: the longest run of consecutive days in which output stayed under a seasonal average, which catches long mediocre spells as well as still ones. On that second measure Germany's longest single-country winter drought ran 109 days and Spain's 131. Combining solar with onshore and offshore wind cuts maximum drought duration by 64%, 52% and 47% respectively, and perfect European interconnection cuts the longest portfolio drought by 65% against isolated national systems.

One caveat, which the authors state themselves and which anyone citing this should repeat. The results depend heavily on the threshold chosen to define a drought, so any headline figure of events per year is a convention rather than a physical constant. Treat the direction as robust and the specific numbers as threshold-dependent.

The fiscal layer

One more German specific that anyone underwriting a long-dated German load must understand.

Transmission network charges in Germany were set to rise sharply. The federal government intervened with a €6.5 billion annual subsidy to the transmission operators, and the four TSOs published preliminary 2026 tariffs dropping from an average of roughly 6.65 to 2.86 ct/kWh, a 57% reduction, conditional on the legislation passing. The electricity tax for manufacturing and agriculture was set to the EU minimum of 0.05 ct/kWh. Separately, the Bundesnetzagentur must replace the entire grid charging framework from 2029 following a European Court of Justice ruling, and is consulting on locational and time-variable structures.

So a large part of the German industrial power price is now a discretionary fiscal transfer with a four-year horizon, sitting on top of a tariff architecture that is being rebuilt for 2029.

If you are signing a fifteen-year German offtake, you are not underwriting a tariff. You are underwriting a subsidy renewal and a regulatory redesign. Price it accordingly, or get an indexation clause.



How The Price Is Actually Computed

In this chapter

- EUPHEMIA is the same class of problem as a production cost model, with zonal balance constraints instead of nodal ones.
- Cross-border capacity in Core is allocated flow-based, using power transfer distribution factors on critical network elements.
- Intrazonal congestion is deliberately excluded, then fixed afterwards by redispatch.

The optimisation

SDAC maximises social welfare across the coupled region subject to:

- **Zonal balance constraints**, one per bidding zone per market time unit
- **Cross-zonal capacity constraints**, either simple net transfer capacities or, in the Core region, a flow-based domain
- **Order constraints**, including block orders, linked blocks, exclusive groups and complex orders

Every zone carries a constraint saying supply must equal demand inside it. The shadow price on that constraint, the amount the total system cost would rise if you demanded one more megawatt-hour there, is that zone's clearing price. This is the same mathematics American nodal markets use. The difference is granularity: Europe gets one price per zone, not one price per node.

Block orders are Europe's unit commitment. A block order is all-or-nothing across a set of hours, which is how a thermal plant expresses a minimum run time and startup cost. Blocks are what turn EUPHEMIA from a linear program into a mixed integer problem, and they are the reason the algorithm occasionally struggles and the coupling occasionally fails to a fallback procedure.

The market time unit moved to fifteen minutes across the coupled day-ahead market in 2025. Check the current status for your specific zone before you build anything on it, but the direction is settled and it materially changes intraday shape economics.

Flow-based allocation

Simple net transfer capacity treats each border as a pipe with a maximum. That is a poor model of an AC network, where a trade between two zones loads lines in a third.

Flow-based market coupling instead defines a set of **critical network elements** and computes **power transfer distribution factors** describing how a net position change in each zone loads each element. The feasible set of net positions becomes a polytope, and EUPHEMIA optimises within it. Shift factors are the same device American nodal markets use. The European difference is that flow-based coupling makes them the primary constraint formulation rather than an approximation applied afterwards.

A minimum share of cross-zonal capacity, the 70% minRAM requirement (minimum remaining available margin), must be made available to cross-border trade, which is a legally mandated floor rather than a physical one and is itself a source of redispatch.

What is not in the price

Intrazonal congestion. Deliberately.

The zonal design is a political and liquidity choice: fewer, larger zones mean deeper markets, simpler hedging, and no politically explosive regional price differences within a country. The cost is that the clearing result is physically infeasible and has to be repaired afterwards, at socialised expense.

Congestion rent from cross-zonal constraints accrues to the TSOs and is partly recycled into long-term transmission rights, the European analogue to the American financial transmission right. See Chapter 11.

What a European trader actually needs

You do not need a nodal solver. You need four things:

1. A **residual load model** per zone
2. A **merit order with carbon**
3. A **cross-border model**, flow-based or NTC-approximated
4. A **hydro water value model** for the Nordics and the Alps

Everything else is refinement.



Practical Approximations

In this chapter

- Residual load plus merit order gets you most of the European price most of the time.
- Over the hours that set prices, data centre load flattens the residual load curve rather than adding to peak, which inverts the standard political assumption. Over seconds it does nothing of the kind.
- Your largest forecast errors will be French nuclear availability, Nordic hydrology and LNG, in that order.

Residual load

Residual load = demand – wind – solar – must-run

Must-run means French nuclear at its scheduled availability, heat-led CHP against the temperature, and run-of-river hydro against flow. Once you have residual load, sort the remaining dispatchable fleet by short-run marginal cost including carbon, and read off the marginal unit. This is the European version of merit order dispatch and it is remarkably serviceable.

Demand

Temperature and calendar regression, plus a slow electrification trend, plus data centre load.

That last term deserves care, and it deserves a distinction the debate usually skips, because “a data centre is a flat load” is true and false at the same time depending on the timescale you are asking about.

On the timescale that sets prices, which is the hourly day-ahead auction, a continuously utilised campus is close to flat. Expensive accelerators are not left idle, so utilisation stays high around the clock and the hourly energy profile looks like a plateau rather than a duck.

On the timescale that matters to a system operator, it is not flat at all. A large training cluster runs its accelerators in lockstep, so job starts, checkpoint entry and exit, fault recovery and collective-communication barriers move the whole campus together. Published work puts the resulting ramps at tens of megawatts per second, rising toward hundreds at the gigawatt-scale American campuses where it has been measured, which is in any case faster than automatic generation control can respond. In July 2024 roughly 1.5 GW of voltage-sensitive data centre load in Northern Virginia tripped offline within seconds, an event NERC compared to the abrupt loss of a large nuclear plant, and one American regional entity now describes AI campuses in the language it uses for steel mills. Inference is the gentler case: smaller sites, distribution voltages, and a profile that follows human activity on a daily and weekly cycle.

Hold both. For the price argument that follows, the hourly plateau is the relevant fact. For the connection and flexibility arguments in Part 3, the second-by-second behaviour is, and the two should not be confused with each other.

Taking the hourly view, then: adding flat load to a system with a duck curve **raises the trough more than the peak**. In practice that means fewer negative price hours, higher midday prices, better capture rates for solar, and a better case for running CCGTs through the day rather than cycling them.

The same effect has been observed in American markets, and it cuts directly against the European political framing, which assumes data centres straightforwardly raise consumer bills. Sometimes they do, through network costs and local congestion. On the wholesale energy component, in a high-renewables zone, a flat load is closer to a stabiliser. Spain, with 85 negative price hours in a single month, is the clearest case.

I would not overstate it. The network cost effect is real and the local congestion effect is real. But anyone modelling European price impact should model both signs.

Capture rates, in both directions

Chapter 11 names falling capture rates as the exposure European PPA buyers most often under-price. The paragraph above says flat load improves them. Both are true, and the tension is worth resolving here rather than leaving the reader to notice it.

Two forces act on solar capture at the same time and in opposite directions. Every additional gigawatt of solar depresses the price in precisely the hours when solar generates, which is why capture rates have been falling and why PPA negotiations stalled. Every additional gigawatt of flat load lifts the midday floor, which pushes the other way.

At the European aggregate, the contest is not close. Chapter 12 puts the renewable build-out at roughly 510 GW to 2030. Chapter 06 puts European data centre capacity at 16 GW in 2024, growing to around 36 GW by 2030, so perhaps 20 GW of new flat load against 510 GW of new generation. Cannibalisation wins by more than an order of magnitude. A solar developer underwriting on the stabilising effect will be disappointed.

Locally it can invert, and that is the whole point of this document. The stabilising effect scales with the load's share of its own bidding zone, measured over the hours that set the price, not with the continental total. That share is extremely uneven. Data centres took roughly half of metered consumption in the Dublin and Meath region in 2024 while accounting for 22% of Irish demand nationally; Amsterdam, London and Frankfurt run between roughly a third and over 40% of local demand. In a zone like that, flat load is a material support to midday prices. In a zone with heavy solar and little coincident load it is not, and the capture rate keeps falling regardless.

So the rule is to model both signs, at the zone the asset actually sits in. The continental average tells you nothing, which is the same conclusion Chapter 12 reaches from the engineering side and Chapter 06 reaches from the demand side. The two things have to land in the same bidding zone for the effect to

exist at all, and Europe's connection queues are currently deciding where that happens.

Cross-border flows

Treat them exogenously where you can. France exports broadly when nuclear availability is good. The Nordics export south when hydrology is wet. Iberia is largely bottled up behind the Pyrenees regardless.

Water value


The Nordic approximation that works without building a stochastic dynamic program: regress area price against reservoir deviation from the seasonal norm, snowpack, and the forward curve. It will not capture the tail, but the tail in Nordic hydro is a dry-year event you should be hedging rather than forecasting.

Batteries

Charge midday, discharge evening. Growing quickly in Great Britain, Germany, Italy and increasingly Iberia. As in Alberta, batteries compress the daytime-to-evening spread, which is the spread most European solar-plus-storage projects are underwriting. If you are long that spread through a storage asset, you are also short its own build-out. Chapter 08 gives the adequacy-side version of the same effect, where a battery's derated contribution falls as more storage is added to the system.

Where it goes wrong

- **French nuclear availability.** The largest single error source in continental Europe. Outage schedules slip, inspections extend, and summer river temperatures curtail.
- **Nordic hydrology.** A dry year moves Nordic prices by a factor, not a percentage.
- **LNG.** Not a European variable. A global one, transmitted through TTF within hours.
- **Policy.** Unbounded. In 2026 alone: an ETS overhaul proposal, a German capacity tender law, a German network tariff subsidy, a grids package in trilogue, and a gigafactory tender. None of these are in your weather model.



Types of Power Trades

In this chapter

- The clean spark and clean dark spreads are the European working instruments.
- Nordic area basis and PPA capture rate risk are the two exposures that most often go unhedged.
- A large flexible load is technically a legitimate ancillary services provider, and in Germany three constraints mean almost none of that revenue is reachable.

Where it trades

Futures on EEX for German, French, Italian, Spanish, Dutch, Belgian, Austrian and other continental zones. Nasdaq Commodities for the Nordics. Products run day, weekend, week, month, quarter, season and calendar year, in **baseload** and **peakload**, where continental peak is 08:00 to 20:00 Monday to Friday. The Nordics trade base almost exclusively.

Spot is day-ahead and intraday, plus imbalance settlement against the TSO. The same day-ahead against real-time structure American markets run, with the balancing market playing the real-time role.

Clean spark spread

The European working instrument. Sell power, buy gas, buy carbon:

Clean spark spread = Power price – (Gas price ÷ efficiency) – (emission factor ÷ efficiency × EUA)

Quoted by efficiency, so a “50% clean spark” or a “55% clean spark.” This is the exact hedge a CCGT operator sells to lock in their plant economics, and the exact hedge a large gas-exposed load buys.

Clean dark spread

Same structure with coal, priced off API2, at roughly 0.34 tCO₂/MWh thermal and typically 36% to 40% efficiency.

The difference between the two, sometimes called the climate spread or switching spread, is the fuel-switching trade, and it is a direct bet on the EUA price and on TTF. It is a smaller trade than it used to be because the coal fleet is shrinking on a legal schedule, but it still moves the German and Polish stacks.

Zone spreads and transmission rights

The basis between two zones is the European congestion trade. DE-FR, DE-NL, FR-ES, IT-CH, and the Nordic internal spreads are the liquid ones.

The financial instrument is the **long-term transmission right**, auctioned by the Joint Allocation Office on behalf of the TSOs, in either physical form, a physical transmission right that entitles you to move power across the border, or financial form, which pays out the price difference without you moving anything. Congestion rent from the day-ahead coupling funds the payout.

Nordic EPADs, and the mistake

An EPAD, an Electricity Price Area Differential, is the one instrument an American trader will not have seen, and it is where Nordic data centre hedges go wrong.

The Nordic market clears a **system price**, which is the price that would result if there were no internal transmission constraints, alongside **area prices** for each bidding zone. Futures trade against the system price. To hedge an actual physical position in a specific area, you need to add an **EPAD** (Electricity Price Area Differential), a contract on the difference between the area price and the system price.

A Nordic PPA or hedge built only against the system price leaves you fully exposed to area basis. Given that the Nordic system's defining feature is a large north-to-south constraint, with cheap hydro in SE1 and SE2 and expensive load in SE4 and southern Norway, that basis is exactly the risk you were trying to hedge.

If you take one operational point from this document: **hedge the area, not the system.**

Corporate instruments

PPAs come in several shapes, and the shape determines who holds which risk:

- **Pay-as-produced.** You take whatever the asset generates. You hold profile risk, volume risk and balancing risk.
- **Baseload.** The seller delivers a flat shape and holds the shaping risk. More expensive, obviously.
- **Pay-as-nominated.** You nominate, the seller delivers or buys the difference.
- **Hybrid and firmed.** Wind plus solar plus storage, structured to a flatter delivery profile. This is where the market is going, and it is why Amazon's Spanish deal is a hybrid and Merlin's is solar-plus-storage.

The exposure everyone under-prices is **capture rate**. A solar PPA at €55/MWh looks cheap against a €110/MWh day-ahead average until you notice that the solar asset produces its energy in the hours when the price is €0 or below. The seller's revenue is capture-weighted, not average-weighted, and falling

capture rates are the primary reason European PPA price negotiations have stalled and volumes have fallen. Chapter 10 gives the force that pushes the other way, and the conditions under which it is strong enough to matter.

Two-way CfDs are the state-side equivalent, now with formal Commission design guidance. If you are a generator with a CfD, you have handed your upside to the state in exchange for a floor, and your exposure to the price level is largely removed. Your exposure to how the contract settles is not: the reference period, the treatment of negative-price hours and your dispatch against that reference are where a CfD-backed asset is still won or lost, and they differ enough between schemes to be worth reading before signing.

Capacity market obligations are a cost line for a large load in Great Britain, Poland, Italy, Belgium, France, Ireland and now Spain, and will become one in Germany from 2032. Budget it.

Ancillary services and flexibility

Frequency containment reserve, automatic and manual frequency restoration reserve, and increasingly local flexibility markets run by distribution system operators, the regional companies that own the local network beneath the transmission grid.

It is worth knowing what the system is actually buying, because “flexibility” appears across European policy documents as though it were a single product. Belgium’s transmission operator splits it by response time and puts volumes against each. Slow flexibility, responding within about five hours, absorbs forecast updates between the day-ahead market and intraday, with a requirement exceeding 4 GW by 2036. Fast flexibility, activated within 15 minutes, is expected to exceed 3 GW by 2036, roughly double today’s level. Within five minutes the requirement is 0.5 GW. In aggregate the Belgian system will need 6 to 7 GW able to react in the last hours before real time, against a national peak that is a fraction of Germany’s.

Those quantities give the word a scale, and the value sits alongside them. Unlocking additional end-user flexibility is assessed as worth 40 to 120 million euros a year in avoided capacity mechanism costs by 2030, rising to 90 to 240 million by 2036, before counting any saving on grid investment. A flexible load is not being thanked for good behaviour. It is displacing procurement the system would otherwise have to pay for, and that is the basis on which it should be negotiating.

A data centre with uninterruptible power supply (UPS) systems, batteries and backup generation is a technically legitimate provider of several of these products, and prequalification routes exist in Germany, Great Britain and the Nordics. The Grid Connection Guidance explicitly encourages flexible connection agreements and time-of-use and locational network tariff signals, which means the regulatory direction is toward paying flexible loads more.

Very few European data centre operators are monetising this, and it is worth being clear that this is mostly not because nobody has thought of it. Three constraints do the work, and only one of them is about willingness.

The first is physical, and it is the one that quietly disqualifies most sites. European frequency containment requires a provider to sustain full activation continuously for at least fifteen minutes. That is the minimum delivery period in the system operation regulation, and when German transmission operators proposed raising it to thirty minutes the regulator refused, so fifteen minutes is the floor rather than a conservative reading of it. A data centre UPS is not sized for that. Autonomy has been falling for years as operators lean on fast generator start instead, from around thirty minutes a decade ago to roughly three to five minutes in current designs. A stock UPS misses the requirement by a factor of three to five. It is not a paperwork problem and no amount of commercial enthusiasm closes it. Closing it means buying additional battery energy, which is a real capital line rather than a free option on an asset you already own.

This is why the working European examples are where they are. Ireland's system services products are specified in seconds and minutes rather than quarter-hours, and Microsoft's Dublin campus has had UPS batteries bid into fast frequency response through an aggregator for several years, responding inside 150 milliseconds. A Nordic site in Finland does something similar with an hour of lithium-ion behind it. None of this is because the Irish and the Nordics are more imaginative; it is that their products were specified around equipment that already exists, and the German one was not.

The second constraint is the permit, and it is sharper than most models assume. German data centre approvals cap standby generation by purpose and by hours, for the plant as a whole rather than per machine. One large Berlin campus is permitted for a total of 200 operating hours a year in emergency mode across all seventeen of its generators, with testing separately limited and confined to a one-hour window on weekday mornings. Another, much larger, is capped at 300 hours a year across forty-eight units with simultaneous operation expressly prohibited. Both restrict the plant to supplying the site when the grid supply fails. Neither forbids grid services in so many words. Both make them practically unreachable, and running outside the stated conditions requires the authority's prior agreement.


The third constraint is contractual, and here I want to be careful, because it is the one most often asserted and least often evidenced. In a colocation structure the party that would do the bidding is usually not the party that owns the availability promise, and the tenant whose equipment would be bid has no particular reason to agree. That split incentive is well described in the literature. What is not documented, at least not publicly, is a standard lease clause prohibiting it. Reviews of colocation contracting practice simply do not discuss grid services at all, which is suggestive of an industry that has not yet written the clause either way rather than proof of a prohibition. Treat it as an open drafting question on a new build and as a negotiation on an existing one.

What this does not support is a story about collapsed revenue. German frequency containment prices fell hard through 2025, bottoming near €7/MW/h in December, and then recovered through 2026 to sit above €20/MW/h for several consecutive months. The market is not dead. It is simply supplied, mostly by batteries built for the purpose, with prequalified battery capacity now well above the volume Germany procures. A campus arriving late into a market that storage developers reached first, carrying equipment that misses the technical floor and a permit written for outages, is not being held back by a lack of ideas.

The honest version of the opportunity is therefore narrower and more specific. It sits in capacity and congestion products rather than in fast frequency response, it is far easier to design into a new build and its leases than to retrofit, it is worth more in jurisdictions whose products match the equipment, and it should not be underwritten on the idea that it improves your standing in a connection queue. Queue position is scored against published criteria, as the German maturity procedure in Chapter 06 shows. Goodwill is not one of them.

That closes the trading side, and it leaves a question sitting underneath every instrument in it. The flexible connection, the locational tariff, the long-term transmission right, the flexibility product just priced: each one is a contract about the state of the network at a particular place and at a particular time. Part 3 is about who can actually see that state. The answer splits by voltage level, and the split decides which of these instruments work today and which are waiting on something nobody has funded.

Part 3: The Grid That Has To Deliver It



Chapter 12

The Layer Nobody Prices


In this chapter

- Roughly 70% of Europe's new renewable capacity connects at distribution level, not transmission.
- The distribution network is where the load actually lands, and it is the layer with the least measurement on it.
- Every market reform in Part 2 that reaches below the primary substation presupposes observability that does not yet exist.

Part 1 priced generation. Part 2 priced zones. Both of them rest on something neither of them looked at: that the network underneath can deliver what the price says it should, and that somebody can tell whether it is doing so. This part is about that layer. The short version is that at the top of it, where the large loads and the large plants connect, the instrumentation is mandated and paid for by whoever connects. Below that, where most of the new capacity is actually arriving, it is nobody's obligation and largely absent.

Start with the ladder, because the rest of this document is unreadable without it.

Where a thing connects decides almost everything about it



Distribution is everything below the transmission grid. Secondary distribution is the last stage of it: the medium-voltage network and the small transformer stations that step down to the socket. Part 3 of this document is about that layer. Boundaries vary between countries and the bands overlap, so treat these as working ranges rather than legal definitions. What does not vary is the consequence: a transmission connection is slow, visible and politically contested, and a distribution connection is none of those things, which is why almost all of the new load went there and why nobody was watching when it did.

Europe's public argument about the grid is conducted almost entirely in transmission terms: interconnectors, bidding zones, 400 kV corridors, the queue for a high-voltage connection. That is where the political drama is, because a transmission line is visible, contested and slow. But if you follow the actual megawatts, most of the new ones never touch a transmission tower. Around 70% of new renewable capacity to 2030 connects at distribution level, against a build-out of roughly 510 GW. The EV fleet, projected at 50 to 70 million vehicles by 2030, connects at low voltage. So do 40 to 50 million heat pumps.

The size of it

Eurelectric's current assessment puts European distribution grid investment needs at around €67 billion a year to 2050, against a historical run rate near €33 billion, with perhaps €55 billion achievable if the sector takes the efficiency measures the study proposes. Read that as a doubling of the annual requirement. An earlier Eurelectric figure of €375 to €425 billion to 2030 is still widely quoted, including in places that should know better; it works out near the historical run rate rather than the required one, so treat any plan built on it as underfunded by about half. The associated employment is 440,000 to 620,000 jobs a year across the EU27 and UK, which is the same labour pool the transmission build-out and the offshore programme are already bidding for, and which Chapter 15 returns to.

Then there is the age profile, which is the part that rarely makes a headline. About one third of EU grids are already more than 40 years old, and that share is projected to pass 50% by 2030. A network built for one-directional flow into predictable domestic load is being asked to host bidirectional flow, high-power charging and industrial-scale new connections, and it is doing so while entering the replacement phase of its own life cycle.


What a kilometre actually costs

Aggregates are hard to sanity-check, so here is the unit data. ACER publishes investment cost indicators drawn from actual European projects and inflation-adjusted to current price levels. These are the most useful reference figures in this document for anyone evaluating a network claim.

Asset	Average	Range
400 kV overhead line, two circuits	€1.46m per km	€0.32m to €4.87m
220 kV underground cable, two circuits	€3.67m per km	€1.95m to €5.24m
AC substation, 6 to 9 bays	€14.77m per asset	€5.39m to €41.27m
HVDC converter station	€0.21m per MW	€0.13m to €0.27m

Two things follow. The ranges are enormous, which makes cross-country comparisons of grid spend per kilometre close to meaningless without knowing terrain, voltage and whether the line is strung or buried. And underground cable at 220 kV costs roughly two and a half times an overhead line at 400 kV per kilometre, while carrying less power. Those are different voltages, so this is not a clean like-for-like test of burying against stringing, and ACER does not publish one. It is still the comparison the undergrounding debate actually runs on, because 400 kV overhead and 220 kV buried are the two things usually on the table for the same corridor. That ratio is how a public acceptance problem converts directly into a cost problem.

ACER unit costs, inflation-adjusted, from actual European projects



Underground cable at 220 kV costs roughly two and a half times an overhead line at 400 kV per kilometre, while carrying less power. That ratio is the whole economics of the undergrounding debate, and it is how a public acceptance problem converts directly into a cost problem. Note also the ranges: they make cross-country comparisons of grid spend per kilometre close to meaningless without knowing terrain, voltage and whether the line is strung or buried.

Who actually pays

The European grid debate is conducted almost entirely in the abstract, and the thing it abstracts away is the thing that decides whether any of it happens.

Network costs do not come from a budget. They come from tariffs, and tariffs land on bills. German grid costs run to roughly **€37 billion a year and account for about 30% of a household electricity bill**, with household network charges at 10.72 ct/kWh in 2025. Across the EU, ACER reports that two thirds of member states have recently revised or plan to revise their network tariff methodologies. That is not coincidence. It is what happens when a cost base rises faster than anyone can explain it.

Every recommendation in this document costs money that somebody pays through a tariff. The distribution investment above, whichever figure you take, is not free, the observability programme in Chapter 14 is not free, and capacity mechanisms appear on bills as a levy. A reader who takes away only the engineering will misread why European reform moves slowly. It is not principally that regulators are timid, or that operators prefer building to measuring, though in several regimes the remuneration does push that way. Chapter 14 sets out why the German version of that claim is wrong. It is that the person who ultimately approves the investment experiences it as a line on an invoice rather than as infrastructure.

That is also why Germany's AgNeS reform matters well beyond Germany. AgNeS is the regulator's project to rebuild the entire network charging framework from scratch, covering roughly €37 billion a year of grid cost and due in force on 1 January 2029. It is the first serious European attempt to allocate these costs according to where they originate rather than spreading them flat, which is the only route by which the bill becomes defensible.

A distinction the queue debate keeps missing

One variable shapes connection queue behaviour more than any allocation rule, and European policy documents rarely name it.

When a developer connects, somebody pays for the network reinforcement that the connection triggers. Under **shallow** charging the developer pays only for the physical connection and the reinforcement is socialised. Under **deep** charging the developer pays for the reinforcement too. Super-shallow variants, common offshore, socialise even the connection asset.

The incentives follow directly, and Germany's answer is the connection contribution set out in Chapter 15, which is a deep-charging instrument in all but name. Under shallow charging a speculative application costs very little, which is a substantial part of why the British generation connection queue reached 738 GW while the separate demand queue tripled from 41 GW to 125 GW in seven months. Under deep charging applications are expensive and queues are shorter, but the first mover at a constrained node pays for reinforcement that later entrants then use for nothing, which deters precisely the connection you most want.


Neither is obviously correct and Europe has not converged. There is also no single published source setting out which member state uses which, which is a striking gap given how much of the current reform agenda quietly assumes one. ACER's 2025 tariff practices report defines the terms and shows which states apply deep charges to generation, and that is as close to a European map as presently exists.

The part that should worry a policymaker

Now put the investment figure next to the measurement figure.

Smart meter penetration across the EU averages around 56%, with Austria, Italy and Spain close to full rollout and a long tail well below the average. At the network level the picture is worse and less discussed. In Eurelectric’s maturity assessment, medium-voltage observability runs somewhere in the 60 to 85% range depending on how advanced the operator is, while low-voltage observability reaches only 40 to 65% even in advanced systems and is close to negligible in basic ones.

Be careful with that number, because it is cited throughout this document and it is softer than it looks. It is the share of the relevant network an operator can see the electrical state of, feeder by feeder, in something close to real time. It is not a count of installed devices, and it is not audited: the published figures come from a maturity self-assessment across operators at different stages, using bands rather than a common asset register. No European body publishes a verified fleet-level measurement of this, which is itself one of the findings, and is why the third recommendation in Chapter 14 is simply to make operators report it. The share of digitalised low-voltage substations ranges from under 40% in emerging implementations to above 65% in leading ones. Only 20 to 40% of distribution system operators, the DSOs that own and run everything below the transmission grid, have reached emerging-to-advanced status on flexibility management at all.



Every instrument in the Grids Package is a contract about the state of the network at a particular place and time: flexible connection agreements, locational and time-of-use tariffs, congestion management, non-firm settlement. Each assumes the line on the right wherever the capacity it contracts sits below the primary substation. Large loads at high voltage are the exception: their telemetry is mandated and self-funded. For everything smaller, Europe is designing a market in locational scarcity while being unable to see most of the locations, and that is a specification and procurement gap rather than a technology one.

Read that against what Europe is currently legislating. The Grids Package pushes flexible connection agreements, where a user accepts curtailment when the network is constrained in return for connecting sooner. Locational and time-of-use network tariffs are recommended. The Netherlands has gone furthest, with ACM making binding arrangements with its seven system operators in 2026 to expand flexible contract offerings, increase the use of congestion management and improve grid utilisation transparency, tracked through a quarterly dashboard.

Every one of those instruments is a contract about the state of the network at a particular place and time. A contract of that kind is only as good as the measurement behind it. You cannot sell non-firm capacity you cannot observe, you cannot enforce a curtailment obligation you cannot verify, and you cannot settle a dispute about either without a time-stamped record of what actually happened.

So the honest statement of Europe's position is this: it is designing a market in locational scarcity while being unable to see most of the locations. That is not a technology gap in the sense of something waiting to be invented. The equipment exists, is mature, and is comparatively cheap. It is a specification and procurement gap, and the following chapters are about how it gets closed.



Chapter 13

What Automation Actually Buys

In this chapter

- Distribution automation is a portfolio decision across levels of capability, not a migration from one state to another.
- In a network with excellent outage figures the case is not reliability, it is how many new connections a feeder can absorb.
- Of everything automation delivers, measured grid state per feeder is the one the rest of this document depends on.

The equipment industry describes distribution automation as a ladder, from a station that indicates a fault locally to a network that reconfigures itself. The ladder is a useful sales structure and a poor planning tool, because no operator climbs it. A rural station feeding forty customers stays at the bottom for its whole life and should. A node under a dense urban cable ring belongs near the top. The planning question is which stations sit on which step, and almost no European operator can answer that for its whole fleet with evidence rather than assumption.

What matters for this document is not the ladder but what each step buys, and that depends on where you are standing.

The reliability argument does not travel

The usual case for automation is restoration time: a fault that took hours to find and switch now takes seconds. In a network with poor supply quality that argument carries a budget on its own.

It does not carry one in Germany. German supply interruption averages twelve to thirteen minutes per customer per year, among the best figures in the world, and in a compensated medium-voltage cable network with fault indicators and manual switching the healthy sections are typically restored inside an hour and a half. There is very little outage time left to buy, the quality element of the regulatory formula pays correspondingly little for buying it, and a German network planner who is pitched on restoration time will say so.

The argument that does travel is throughput. How many heat pumps, how many charging points, how much rooftop generation can be connected to a given feeder before someone has to dig up a street. That question is answered by knowing what the feeder is actually carrying, and it is the question every German distribution operator is now being asked faster than it can answer.

What the equipment actually does

Strip the vendor taxonomy down and the work falls into four groups.

Finding and clearing faults. Locating a fault before the crew leaves the depot, handling earth faults in a compensated network, isolating the faulted section and restoring the rest automatically, and keeping a record good enough to establish afterwards what happened.

Holding supply up. Transferring a critical node to its second feed without waiting for an instruction, and switching remotely rather than by driving to the station.

Seeing the network. Measuring load and voltage per feeder, watching the condition of the asset and its enclosure, and recording power quality.

Making the rest possible. The unglamorous group: safe access and safe switching, measurement built into the cable connector rather than bolted on afterwards, provision for sensing in equipment being installed for other reasons, and protection that works without an auxiliary supply.

That fourth group is where the cost argument of the next chapter lives, because those are the things that are nearly free at the moment a station is open and expensive on a visit of their own.

One honest limit

None of this creates capacity. If a corridor is physically full at peak, the best-instrumented network in Europe will tell you so with great precision and restore you very quickly, and it will not move one additional megawatt through the conductor. What measurement and automation do is let you use the headroom that already exists for most of the year, contract for non-firm access against it, and stop sizing everything for a worst case nobody has measured. That is a large prize. It is not the same prize as building, and anyone selling it as a substitute is overselling.

The keystone

Of everything above, one item carries the rest of this document.

Measured grid state per feeder is what makes a flexible connection agreement enforceable, a locational tariff defensible, a congestion management scheme auditable and a non-firm contract settleable, wherever the capacity being contracted sits below the primary substation. Chapter 12 showed that low-voltage observability sits somewhere between 40 and 65% even among advanced operators. Large loads connected at high voltage are the exception, and Chapter 14 sets out why: their telemetry is mandated and self-funded. For everything smaller, that 40 to 65% is the ceiling on every market reform in the Grids Package, and it is barely mentioned in the package itself.



Chapter 14

The 2026 to 2030 Window

In this chapter

- The F-gas Regulation does not replace Europe's switchgear fleet, the switches and circuit breakers that connect and disconnect the network. It resets the catalogue, which forces every utility technical standard for secondary switchgear to be rewritten now.
- Germany's §14a regime sets a deadline rather than a measurement mandate: from the end of 2028 an operator that cannot show the coverage thresholds loses a power it relies on, and the station is one route to those thresholds rather than the only one.
- The specification window closes before the measurement deadline arrives, and equipment specified without sensing provision stays blind for a 40-year asset life.

This chapter contains the largest sum of avoidable money in the document, and almost nobody outside the equipment industry is discussing it.

The forcing function

Regulation (EU) 2024/573 phases SF6 out of switchgear on a staged timetable. Sulphur hexafluoride is the gas the industry has used for decades to insulate switchgear and quench the arc when a switch opens, which is what lets the equipment be small. It is also the most potent greenhouse gas the IPCC lists, which is why it is being removed. Medium-voltage switchgear for primary and secondary distribution rated at or below 24 kV is prohibited from 1 January 2026. Above 24 kV and up to 52 kV

follows on 1 January 2030. On the high-voltage side, 52 to 145 kV with rated current at or below 50 kA is prohibited from 1 January 2028, and above 145 kV or above 50 kA from 1 January 2032. Member States were required to have penalties in place by 1 January 2026.

Be precise about what this does and does not do, because the overstated version of this argument is easy to make and a DSO reader will dismantle it immediately. The regulation does not require anyone to rip out working SF6 switchgear. It prohibits putting new F-gas switchgear into operation, it applies to equipment ordered after 11 March 2024, and it carries derogations: equipment ordered before that date can still be commissioned with evidence of the order, parts may be installed to repair or service existing switchgear without triggering the prohibition, and relief exists where compliant equipment cannot be secured.

What it does instead is reset the catalogue. From 1 January 2026, every new and replacement secondary station at or below 24 kV has to be built from a new generation of equipment. Manufacturers redesign their ranges. More importantly, every utility technical standard that specifies secondary switchgear, the German Werknorm, a utility's own binding in-house specification for what it will buy and install, being the clearest case, has to be rewritten to name the replacement. Those rewrites are happening between roughly 2025 and 2027, once, and they govern what gets installed for the following decade.

The volume flowing through those specifications is large for reasons that have nothing to do with F-gas. Around one third of the EU fleet is already past forty years. Roughly 70% of new renewable capacity connects at distribution. EV charging, heat pumps and large-load connections all arrive at this layer. So the regulation does not create the replacement wave; it determines what the replacement wave is allowed to install, at the exact moment the technical standards are open.

Why the timing is the whole argument

Work out what a retrofit actually costs.

The expensive parts of instrumenting a secondary station are not the sensors. They are the planned outage, the switching programme, the crew visit, the civil and cable work, the commissioning and the paperwork. Those are labour and access costs, in exactly the trades Chapter 15 identifies as the scarcest in the European grid workforce, and in a market where crew time is the binding constraint rather than capital.

Every one of those costs is already being paid, by someone, on every F-gas-driven replacement between now and 2030. The station is open. The crew is there. The outage is booked. The cable terminations are being made up. Adding current and voltage measurement, a controller, a communications path and the provision for a motor drive at that moment is a marginal cost on a visit already funded. Doing it later, as a dedicated retrofit, means paying the whole access cost a second time.

This is also where the previous chapter's unglamorous fourth group stops being a list of components and becomes a strategy. If the measurement lives in the cable connector that is being installed anyway, the incremental scope is close to zero.

The second forcing function

Germany has already created the demand side of this, and it is instructive because it is specific where the European guidance is general.

Since 1 January 2024, under §14a EnWG and the Bundesnetzagentur determination BK6-22-300 of 27 November 2023, a network operator curtailing controllable consumption devices has to base that action on a determination of grid state. The devices in scope are non-public EV charging points, heat pumps, cooling systems and battery storage in its consumption phase.

Note what the determination requirement actually says, because the detail is the whole point. Measurement everywhere is not mandated. Operators may use network models and calculations built on current measurements, which is the sensible engineering answer. But the determination carries coverage thresholds: grid-oriented control counts as technically feasible where the operator has minute-level grid state data from at least 15% of network connection points in the affected area, or from at least 7% plus measurement at the distribution transformer. Below that, the operator falls back on preventive control, and the fallback is explicitly temporary. It runs to the end of 2028 at the latest, is limited to 24 months per grid area from first use, and is capped at two hours per device per day.

Read those two thresholds carefully, because the relationship between them is the part this chapter has so far undersold. Both are counted at network connection points, not at the station. The station measurement does not satisfy the requirement on its own; what it does is cut the share of connection points you need from 15% to 7%. And connection-point measurement at minute granularity is, in practice, the metering rollout: the intelligent metering systems being installed at customer connections and administered through the gateway operator. That programme is running on its own schedule, with its own budget, and is not waiting for anybody's secondary substation.

This matters for the argument in two directions. It means an operator has a route to compliance that does not pass through the station at all, which is a fair answer to anyone reading this chapter as a case for buying sensors. It also means the station measurement is a lever on the cost of compliance rather than the only road to it: halving the required connection-point coverage is worth real money in a network area where the metering rollout is slow, and worth little where it is fast. A planner weighing the two is making a portfolio decision area by area, not a yes or no on instrumentation, and the chapter's recommendation should be read that way.

So Germany has written a measurement deadline rather than a measurement mandate, and the deadline is the end of 2028. An operator that reaches it without the coverage thresholds in a given network area loses the ability to curtail controllable load there in a defensible way, which in a grid absorbing heat pumps and home charging is not a theoretical inconvenience.

Put the two clocks side by side. The technical standards that decide whether a station can measure anything are being rewritten in 2025 to 2027. The measurement coverage requirement bites at the end of 2028. The specification window closes first, and whether those two facts get connected inside procurement departments is being decided unevenly, in a way worth separating carefully, because the honest picture is not that nobody is doing anything.


The large group operators have named programmes and are running them at scale. Westnetz passed its thousandth digital station of the year in November 2024 and had around 2,400 in service at that point, against roughly 60,000 local substations in its area, of which it intends about 20,000 to transmit data and 12,000 to be remotely controllable. It was installing at roughly twenty-five a week at about €100,000 per station. Bayernwerk commissioned its five-thousandth in August 2025, added around 3,400 in 2025, and is targeting some 13,000 more by 2030, split about evenly between new digital stations and retrofits. These are not organisations that need to be told the argument in this chapter, and their group technical standards already carry telecontrol and sensing provisions. Writing as though they do not would be wrong, and a reader inside one of them would stop here.

That €100,000 is worth pausing on, and so is what it does not tell you. As reported it is a programme cost per station for a full digital conversion: the equipment, the visit, the outage, the connectivity, the integration and the commissioning, carried across a rollout of thousands. It is not the cost of adding a sensor to a station you are already opening.

Three numbers matter here and only one of them is public. The first is that programme cost. The second is the cost of retrofitting measurement into a station as a job of its own, which is lower than the first but still carries the visit and the outage, the two expensive parts. The third is the premium for specifying a unit sensing-ready at the moment it is being replaced for other reasons, which carries neither, and which is the number this chapter's recommendation actually rests on.

I do not have the third number from a source I can cite, and I am not going to infer it from the first. The argument for sensing-ready provision is that it moves the cost from the second category to the third, and the size of that move is precisely what a procurement department would need before changing a specification. A credible table of those three figures, built from framework tender data rather than from vendor material, is the single most useful thing anyone could add to this argument, including me.

Two things survive the correction. The first is that even these programmes are early against their own targets, though less early than the snapshot above suggests: Westnetz at 2,400 stations was already at roughly 12% of the 20,000 it intends to have transmitting data, and the figure is from late 2024. The E.ON group's German operators together passed ten thousand digital stations during 2025. The direction of travel is faster than the vendor literature implies. What has not changed is the ordering: the specification window still closes before most of the installation happens. The second is the long tail. Germany has over eight hundred distribution operators, and survey work across German, Austrian and Swiss operators found fewer than one in ten describing their networks as largely automated, with full implementation of the controllable-load rules in the low single digits. A separate survey of German operators put the share that have substantially integrated smart grid technology into their network architecture at around 3%, and reported that the smallest municipal operators struggle most to find the specialist staff. Neither survey breaks the result down by operator size in a way that would let me assert the contrast as a measured fact, so take it as the direction the evidence points rather than as a proven split. But the group programmes above are self-reported numbers from the largest operators in the country, and they are not the fleet average.



The specification window closes before the measurement deadline arrives. A Werknorm is opened once on the F-gas clock, between roughly 2025 and 2027, and what it says about sensing provision governs what gets installed for the following decade. Equipment specified in that window without sensing provision stays blind for a 40-year asset life.

The window closes

A secondary station has an asset life measured in decades, and roughly one third of the EU fleet is already past forty years. Equipment installed between 2026 and 2030 will still be in service in the 2060s.

So the decision being made by default, in thousands of individual procurement specifications over the next four years, is whether Europe's distribution network becomes observable during the replacement cycle it is already paying for, or whether it stays blind until the cycle after next.

What I would do about it

Three things, offered as a view.

Put sensing-ready into the technical standards now being rewritten. Not full automation, which would be disproportionate for a rural single-feed station and would rightly be resisted on cost. Sensing-ready: the measurement provision, the space, the auxiliary supply and the communications path, so the later step is a configuration change rather than a second outage. The intervention point is the specification rather than the station, which is what makes this cheap. A Werknorm is a document, it is being opened anyway on the F-gas clock, and it governs thousands of installations.

Recognise what it costs to run a sensor, not what it costs to buy one. The OECD and ACER both frame this as CAPEX bias, and in several European regimes that framing is right. In Germany, which is where the replacement cycle is largest, it is wrong, and getting it wrong is the fastest way to lose a network planner's attention.

Since the 2016 reform, §10a ARegV has given distribution operators a Kapitalkostenaufschlag: capital spent after the base year enters the revenue cap without waiting for the next regulatory period, applied for annually and partly on planned figures. A sensor, an RTU and a motor drive are activatable assets. They ride the same channel as the circuit. Nobody in Germany is choosing copper over instrumentation because only copper earns.

The bias runs the other way. What a monitored station costs is not mainly the hardware. It is the connectivity subscription, the data platform, the cybersecurity obligations that attach to a remotely reachable endpoint, and the commissioning engineer who sets it up and keeps it configured. Those are operating costs, and operating costs are fixed to the base-year level for the whole regulatory period, so an operator who instruments a thousand stations carries the recurring cost until the next cost review resets the baseline. NEST addresses this with an operating-cost adjustment rather than with TOTEX treatment, and that is the mechanism worth arguing about.

So the recommendation is narrower than it first looks, and more actionable for it: make the recurring cost of an observable network recognisable in the year it is incurred. The capital side is already handled.

Two qualifications, because the regulatory ground is moving and an operator will know it better than I do. The reform now under way does introduce an operating-cost adjustment for distribution operators, but it is triggered by growth in the supply task, the connection points and load an operator has to serve, rather than by digitalisation as such. Whether a connectivity subscription or a data platform qualifies is the live question, not a settled one, and it is the question worth putting to the regulator rather than the general case for recognising operating cost.

The second qualification is the harder one and it is the reason a network planner hesitates even when the money is recognised. Operating cost entering the revenue cap also enters the efficiency comparison against peers. An operator that instruments a thousand stations while its comparators do not carries that spend as measured inefficiency, and the benchmark trims it over the period. Until digitalisation cost is either carried by everyone or shielded from the comparison, the first mover is penalised for moving first. That is a sharper statement of the obstacle than anything about capital bias, and it is the one I would take into a regulatory conversation.

Get the observability metric defined properly, because it is already being drafted. Operators report reliability indices because they are required to. Almost nobody reports what share of their secondary stations can measure load per feeder. The German regulator is now consulting on a digitalisation index as the final step of its incentive-regulation reform, which means the question is no longer whether such a metric exists but what it counts. That is a better fight and a shorter one. An index that rewards counting connected devices will be met by counting connected devices; one that asks what share of feeders can be observed at minute granularity, per operator, published annually, measures the thing that actually constrains flexibility. It is cheap to collect, it is the ceiling on every flexibility instrument in the Grids Package, and what gets measured gets budgeted.

Germany illustrates the problem by moving the other way. Its periodic report on the state and expansion of the distribution networks ended after the 2022 edition, and the findings now appear as articles on the regulator's data portal instead. That is a change of format rather than an abandonment of the monitoring,

but the effect on this particular question is the same: the one recurring document in which a distribution observability figure could have lived, and been compared year on year, no longer exists. That is a small administrative decision with a long shadow, and it happened in the member state with the most ambitious network tariff reform in Europe.



Chapter 15

Where the Large Loads Actually Land

In this chapter

- The AI load is reported as a transmission story and experienced on site as a distribution story.
- A large campus is a private distribution network with the same work to do as a utility's.
- At high voltage the flexible connection is held up by contract terms rather than by missing measurement, and the statutory frame for those terms already exists.

Chapter 06 left the data centre question at the connection queue. This chapter picks it up one voltage level down, because that is where the engineering actually sits and where the policy conversation is weakest.

Two different problems wearing one name

A hyperscale campus takes a high-voltage connection, frequently at 110 kV, and increasingly owns the substation that terminates it. Behind that substation it runs its own medium-voltage network across the site: a campus backbone, distribution to halls, transformation to the equipment, and in a growing number of cases onsite generation, storage and grid-forming inverters behind the meter.

That private network has precisely the same work to do as a utility's. It needs safe access and safe switching, fault location, automatic transfer, event records, power quality measurement. The difference is the loss function. For a utility, an outage is a regulated reliability statistic. For a campus, uptime is the product, so continuity dominates, and the operator will pay for automatic transfer at nodes where a utility would settle for remote monitoring.

Meanwhile the connection application lands in a grid that was planned for a fraction of the requested load. Data centres and EV hubs are requesting connections into networks sized for the domestic and light-commercial demand of the area, which is the structural situation behind the queue figures in Chapter 06.

Why this matters for the flexible connection debate

The Grids Package answer to the queue is the flexible connection: connect sooner, accept curtailment when the network is constrained. The Dutch have gone further, offering alternative and time-dependent transmission rights in exchange for lower network costs and faster access.

Now look at what that contract actually requires from each side, and, more importantly, at which voltage level it is being signed, because the answer differs far more than the debate allows. Getting it wrong is the most common mistake in this argument.

A 200 MW campus connects at 110 kV or above. At that level the measurement is not missing and never was. German high-voltage connection conditions require the connecting party to install telecontrol equipment at its substation and deliver, to the network operator's control centre, the phase voltages and currents at the connection point, active and reactive power and plant status, transmitted cyclically and on change with timestamps fine enough to reconstruct a sequence of events. The connecting party pays for it as a condition of being connected at that voltage.

One distinction matters here and the operator-published specifications draw it clearly. What a high-voltage connection standardly delivers is measurement. The graduated setpoint channel, where the operator sends an active power instruction and receives confirmation of what was applied, is specified for generating plant. A pure load connection is not built with that control path as a matter of course, and curtailment of a large consumer today runs through operational instructions under the grid-security provisions rather than through a telecontrol setpoint. So the actuation half is a gap, and it is a smaller and more tractable gap than the measurement one, because the standard already exists next door.

So for a large load the sensor is not the obstacle. Everything a flexible connection needs in order to be measured and settled is already specified in the connection conditions, and has been for years. What is missing is the other half of the arrangement: the control path to instruct the load, which is standard for generation and not for consumption, and then the terms on which the load is curtailed.

Germany has in fact legislated those terms. Section 17 of the energy act now lets a network operator offer a flexible connection agreement that limits withdrawal or feed-in, and it specifies what such an agreement must contain: the level of the limit, the periods in which it applies, the duration of the agreement, the technical requirements, and the connecting party's liability if the limit is exceeded. Those are, almost exactly, the four clauses a financeable non-firm product needs. The statutory frame is not the gap. The forecast obligation is the genuinely hard clause, and Chapter 10 explains why. A training campus moves in lockstep at tens to hundreds of megawatts per second, faster than automatic generation control responds, so what the operator needs committed is not a smooth curve but a bounded envelope and a notice period. Nobody has standardised what that clause should say. Large loads discount non-firm access heavily, not because curtailment is unacceptable in principle and not because nobody can see the network, but because an interruption obligation with no defined bound cannot be put in front of a credit committee. An inference workload can genuinely tolerate interruption, as Chapter 06 argued. It cannot tolerate an unbounded one.

For large loads at high voltage, the flexible connection is a contract-design problem, not a measurement problem. That distinction matters because the two have different owners and different timelines. And the gap is narrower and more awkward than a drafting problem, because the drafting has been done. The German statute has the fields. The Dutch alternative and time-bound transport rights are working instruments of the same kind. What is missing is numbers in the fields.

That is a commercial standoff rather than a legal vacuum. The network operator will not commit to a bound on curtailment hours it cannot forecast, because the forecast depends on connection decisions it has not yet made and on a load whose draw moves in lockstep at second scale. The connecting party will not accept an obligation without a bound, because an uncapped interruption cannot be financed. Both positions are rational and neither side can move first without the other. Naming it as a standoff rather than an oversight matters, because the two have different remedies: an oversight is fixed by someone drafting, and a standoff is fixed by a regulator setting a default or by one party pricing the risk the other will not carry.

The measurement problem is real. It simply lives one layer down, and the reason it lives there is structural rather than technical. At 110 kV there is a connecting party with a balance sheet, a connection agreement and a technical obligation to fund its own telemetry. Further down it thins out and then stops. A larger medium-voltage customer still has a transfer station and can be required to provide telecontrol under the connection rules that govern that level. Below that there is nobody to ask. A heat pump, a charging point, a rooftop inverter and a small battery do not install a telecontrol gateway, and the operator cannot charge any one of them for the cost of seeing the feeder they sit on. That cost comes out of the tariff or it does not get spent, and mostly it has not been spent. It is also why the replacement cycle in Chapter 14 matters so much: it is the only moment at which the expensive part of the job, the access to the station, is already being paid for by something else.

What the connection itself costs

Before the campus argument, one line item that is missing from most European data centre models and is large enough to change a site decision.

Germany charges a **Baukostenzuschuss**, a connection contribution, on the capacity a load books. It is not a cost-plus charge for the physical connection assets, which are billed separately. It is a per-kilowatt payment for the network capacity being reserved, calculated as a published rate multiplied by the booked capacity, and it falls on withdrawal points.

That last word does more work than it looks. Pure generation pays nothing on its feed-in capacity. A grid-scale battery does pay, on the capacity it books to withdraw, and in July 2025 the Federal Court of Justice confirmed that charging it is lawful and not discriminatory: for the purpose of sizing a connection, a battery is a consumer like any other. Anyone modelling a European storage project in Germany should carry this line. On a 100 MW battery it is roughly €10 million, against perhaps €25 to €30 million of capital cost for a two-hour system, and it is routinely left out.

The rates are public and the arithmetic is immediate. The regulator's own worked example puts a 600 MW connection at roughly €60 million. Published operator schedules sit where you would expect against that: TransnetBW lists €96.91/kW at extra-high voltage for 2026, and Bayernwerk's distribution schedule lists €97.26/kW at high voltage and €125.31/kW at medium voltage. A 100 MW campus at high voltage is therefore looking at something near €10 million and a 300 MW campus at around €29 million, payable on the capacity booked rather than the capacity used. The rate is a five-year average of the underlying capacity price rather than the current year's, which matters because transmission capacity prices fell sharply into 2026 and the averaging deliberately stops a connecting party from timing its contract into that trough.

Four consequences follow, and the last two are the ones usually missed.

The first is that this is a real number in a model that often omits it entirely, and it lands early, before any revenue.

The second is that booking headroom is expensive, and that developers do it anyway on purpose. A campus that books 300 MW to leave room for a later phase pays for the headroom now. Read naively that looks like a forecasting error, and sometimes it is. More often it is an option purchase: in a regime where you cannot come back for more capacity at the same node without rejoining a queue, the contribution is the premium on an option to grow, and against a queue measured in years it is a cheap one. That reframes the flexibility argument earlier in this chapter. A developer who can credibly offer flexibility is not just buying goodwill; it is trying to buy the same option for less.

The third is that the per-kilowatt rate does not fall monotonically as you go down the voltage levels. Medium voltage is the most expensive band in the Bayernwerk schedule, more expensive per kilowatt than high voltage. That is not an anomaly, it is the charge reflecting how much network sits between the connection point and the rest of the system. It also means the usual intuition, that connecting lower down is cheaper, is wrong on this line item, which matters because the rest of this document has been arguing that the layer below the primary substation is where the pressure is.

The fourth is the interesting one, and it is the closest thing in German practice to the locational price signal this document keeps saying Europe does not have. Transmission operators are allowed to vary the contribution by where you connect, across a small number of bands set by how much grid expansion and congestion management your connection point causes. The regulator's published example runs from the full rate down to a fifth of it. On the regulator's own example that is the difference between roughly €60 million and roughly €12 million for the same 600 MW. A locational signal worth tens of millions already exists; it simply arrives as a one-off connection charge rather than as a price per megawatt-hour, and it is visible only to parties large enough to be connecting at that level.

Distribution operators do not do this, because the regulator's guidance does not provide for it on the reasoning that locational differences within a distribution area are generally not material enough. That guidance is a position paper rather than a binding determination, which means it is arguable, and on the evidence in Part 3 it is worth arguing with. That is a defensible position about the historical network and

an increasingly awkward one about the network described in Part 3, where the differences between one feeder and the next are exactly what everybody is trying to price. The instrument exists, it works, and it is switched off precisely at the level where this document has argued the constraint now binds.

The campus as a grid asset

One further point, because it is where the political framing is furthest behind the engineering.

A campus with automatic transfer, onsite storage, grid-forming inverters and the ability to shed or shift a defined block of load on a verifiable signal, is a materially different object from a campus that simply draws its contracted maximum. The first can relieve a constraint, provide a service and support local voltage. The second can only consume.

European policy currently treats them identically. Both are a connection application of X megawatts in a queue, assessed against a worst-case coincident peak. There is no standard way for a large load to be credited for the flexibility it can actually provide, no common measurement basis on which such a credit could be settled, and in several member states no route for a behind-the-meter asset to be remunerated for a grid service at all.

System operators, meanwhile, have already begun pricing the difference that policy refuses to recognise. Belgium's transmission operator, in its adequacy and flexibility study covering 2026 to 2036, assumes data centres are 20% flexible as its central case, with a low case of zero and a high case of 30%, and applies that central figure across every scenario it runs. That is a European TSO putting a number on precisely the question this chapter is about. Twenty per cent of a 300 MW campus is 60 MW of dispatchable load that currently earns nothing and appears nowhere in the connection assessment.

The same study is unusually candid about why the question became urgent, noting that the growth of battery storage and data centres has far exceeded expectations. Belgian data centre consumption stood at roughly 3.2 TWh in 2024 and has grown at about 10% a year since 2018, with projections reaching somewhere between 10 and 30 TWh by 2050, equivalent to a 12 to 35% increase in the country's entire present electricity consumption. Across Europe, data centres account for around 3% of total energy consumption today.

That is a design choice, not a law of physics, and it is the cheapest capacity Europe is currently refusing.

Where this touches the workforce problem

A brief note, because the connection is real and the omission would be odd.

Automation substitutes for field labour. Each capability added removes truck rolls per fault: the crew stops driving to diagnose and stops driving to reconfigure, and is dispatched only to repair. In a system where crews are the binding constraint and retirements outnumber entrants in grid trades by roughly 1.4 to 1, that substitution is not a cost saving so much as a way of making a shrinking workforce cover a growing network.

But it shifts the demand rather than removing it. Anything beyond a local indicator needs commissioning engineers, protection and control specialists, and people who can configure and test an automation scheme. Those are the scarcest categories in the European grid workforce, harder to train than a field crew and in direct competition with the transmission build-out and the offshore programme for the same people.

Automation is therefore both a partial answer to the labour constraint and an additional claim on its scarcest part. Any programme built on the argument in Chapter 14 has to fund the commissioning capability alongside the hardware, or it will install equipment that never gets commissioned, because nobody is available to bring it into service. I have seen that outcome more than once and it is an expensive way to buy a panel indicator.

What's Next

The American story is a generation story: demand is arriving, prices will move, build fast and trade the response.

The European story is an allocation story. The generation will get built, slowly, mostly through state-backed tenders and CfDs rather than merchant investment. The grid will get built more slowly still, and will remain the binding constraint through the end of the decade. The prices will be set by imported gas and a legislated carbon price, both of which are more volatile and less forecastable than a domestic supply curve.

The opportunity is not in predicting the marginal unit. Europe's marginal unit is a CCGT and everyone knows it. The opportunity is in the four places where Europe's institutional design creates a gap between value and price:

1. **Connection capacity**, which is scarce, valuable and not tradeable. It is already priced, locationally, at transmission level in Germany through the connection contribution in Chapter 15. What is missing is a visible market rather than any market at all. Positions do change hands, usually through the sale of the company holding the connection agreement, and the people doing it know what a megawatt is worth. There is no published price, no obligation to offer, and no mechanism for an operator to recover capacity held for its option value rather than built on.
2. **Locational value**, which is real, estimated at somewhere between €251 and €540 million a year in Germany alone depending on whose study you take, and deliberately excluded from the market price.
3. **Flexibility**, which is increasingly remunerated and structurally under-supplied.
4. **Observability below the primary substation**, which is the precondition for the other three wherever they act on distribution-connected capacity, and which is being decided by default, in procurement specifications, during a replacement cycle that closes in 2030. At high voltage the equivalent constraint is not measurement but contract design, which is cheaper to fix and nobody is fixing.

Whether a European data centre build raises or lowers prices, whether the gigafactories get connected, whether Germany ever splits its zone: unsatisfyingly, it depends.

But you are now equipped to argue about it properly.

Sources

1. Neel Somani, *Power 2026: Electricity Pricing in the Age of AI* · <https://power2026.ai/>
2. IEA, “Overcoming energy constraints is key to delivering on Europe’s data centre goals” · <https://www.iea.org/commentaries/overcoming-energy-constraints-is-key-to-delivering-on-europe-s-data-centre-goals>
3. IEA, *Electricity 2026: Grids* · <https://www.iea.org/reports/electricity-2026/grids>
4. Ember, *Grids for data centres in Europe* · <https://ember-energy.org/app/uploads/2025/06/Report-Grids-for-data-centres-in-Europe-PDF.pdf>
5. Euronews, “Europe is hungry for AI data centres but its energy grid cannot feed them,” 5 May 2026 · <https://www.euronews.com/next/2026/05/05/europe-is-hungry-for-ai-data-centres-but-its-energy-grid-cannot-feed-them>
6. Global Data Center Hub, “Grid Interconnection Delays Repricing Europe’s Data Center Market” · <https://www.globaldatacenterhub.com/p/grid-interconnection-delays-repricing>
7. Avanza Energy, “The \$176 Billion Detour” · <https://avanzaenergy.substack.com/p/the-176-billion-detour>
8. CNBC, “High energy prices could derail Europe’s AI race,” 18 May 2026 · <https://www.cnbc.com/2026/05/18/europe-ai-energy-electricity-costs-data-centers-china-us.html>
9. European Commission, European Grids Package · https://energy.ec.europa.eu/topics/infrastructure/european-grids_en
10. Eurelectric, “Under the EU policy tree: the European Grids Package unwrapped” · <https://www.eurelectric.org/blog/european-grids-package-energy-highways/>
11. EP Think Tank, “The European grids package: Guidance on CfDs and grid connections” · <https://epthinktank.eu/2026/02/07/the-european-grids-package-guidance-on-contracts-for-difference-and-grid-connections/>
12. FSR, “Flexibility and grids: tackling grid capacity and connection queues” · <https://fsr.eui.eu/flexibility-and-grids-understanding-available-grid-capacity-connection-queues-and-what-can-be-done/>
13. EuroHPC JU, “The EuroHPC Joint Undertaking launches the AI Gigafactories Call,” 30 July 2026 · https://www.eurohpc-ju.europa.eu/eurohpc-joint-undertaking-launches-ai-gigafactories-call-2026-07-30_en
14. European Commission, AI Gigafactories · https://commission.europa.eu/topics/competitiveness/competitiveness-coordination-tool-projects/ai-gigafactories_en
15. Technology.org, “EU Commits €10 Billion to Seven AI Gigafactories,” 1 August 2026 · <https://www.technology.org/2026/08/01/eu-seven-ai-gigafactories-10-billion-tender/>
16. STL Partners, “The EU’s AI Gigafactory Initiative” · <https://stlpartners.com/articles/data-centres/eu-ai-gigafactory-initiative/>

17. ENTSO-E, Bidding Zone Review · https://www.entsoe.eu/network_codes/bzr/
18. ACER, “ACER will evaluate electricity system operators’ bidding zone review” · <https://www.acer.europa.eu/news/acer-will-evaluate-electricity-system-operators-bidding-zone-review-alignment-eu-regulatory-framework>
19. Montel, “TSOs underestimated gains of German bidding zone split, Acer” · <https://montelnews.com/news/a8869b1a-b9cc-4015-95ad-99e5ae935cab/>
20. Clean Energy Wire, “German government approves new gas power plant auctions,” 13 May 2026 · <https://www.cleanenergywire.org/news/german-government-approves-new-gas-power-plant-auctions-back-wind-and-solar>
21. White & Case, *Infrastructure & Energy Europe Regulatory Compass Q2 2026* · <https://www.whitecase.com/insight-alert/infrastructure-energy-europe-regulatory-compass-q2-2026>
22. PwC Legal, “German Government reaches general agreement on power plant strategy with the European Commission” · <https://legal.pwc.de/en/news/articles/german-government-reaches-general-agreement-on-power-plant-strategy-with-the-european-commission>
23. Trio, “Grid-fee subsidy and power-tax cut in Germany 2026” · <https://www.trioadvisory.com/resources/grid-fee-subsidy-and-power-tax-cut-germany-2026-what-energy-intensive-sites-will-really-save>
24. Oxera, “Reform of the German electricity grid tariff system” · <https://www.oxera.com/insights/agenda/articles/reform-of-the-german-electricity-grid-tariff-system-should-producers-pay-grid-charges/>
25. Cube Concepts, “Rising redispatch costs” · <https://cubeconcepts.de/en/rising-redispatch-costs/>
26. Rystad Energy, “European data center PPAs fall even as capacity buildout accelerates” · <https://www.rystadenergy.com/insights/european-data-center-ppas-fall-even-as-capacity-buildout-accelerates>
27. Voltstack, “European Gas & Power in Numbers” · <https://voltstack.energy/insights/european-gas-power-statistics-2026>
28. Trading Economics, EU Natural Gas and EU Carbon Permits · <https://tradingeconomics.com/commodity/eu-natural-gas> and <https://tradingeconomics.com/commodity/carbon>
29. Captured Carbon, *Weekly Carbon Report*, week 31, 27 July 2026 · <https://www.capturedcarbon.com/wp-content/uploads/2026/07/Week-31-Carbon-Report.pdf>
30. Enlit World, “ENTSO-E releases final report into Iberian blackout” · <https://www.enlit.world/library/final-entso-e-reports-sheds-light-on-multiple-factors-that-caused-iberian-blackout>
31. Rabobank, “Facts and lessons learned from the Iberian blackout” · <https://www.rabobank.com/knowledge/d011479255-facts-and-lessons-learned-from-the-iberian-blackout>
32. Tech Plus Trends, “AI Data Center Energy Cost Europe: 2026 Market & Cost Guide” · <https://techplus-trends.com/ai-data-center-energy-cost-europe-market-guide/>

33. Voltarione, “How much electricity does an AI data center really consume? A 2026 European benchmark” · <https://volarione.com/blog/ai-data-center-electricity-consumption-europe-2026>
34. Solar Data Atlas, Solar PPA Prices Europe · <https://www.solardataatlas.com/en/insights/solar-ppa-prices-europe>

Part 3, the distribution layer

1. Eurelectric, “Grids for Speed” and “Connecting the dots” (distribution investment, €375–425bn to 2030, 70% of new renewables at distribution, jobs) · <https://www.eurelectric.org/publications/grids-for-speed/> and <https://www.eurelectric.org/news/pr-connectingthedots/>
2. Eurelectric, *Wired for Tomorrow* (grid observability and digitalisation maturity, smart meter penetration) · <https://powersummit2024.eurelectric.org/wp-content/uploads/2024/05/Wired-for-tomorrow-full-report.pdf>
3. Eurelectric, *Distribution Grids Handbook* · <https://www.eurelectric.org/wp-content/uploads/2024/11/Distribution-Grids-Handbook.pdf>
4. JRC, *Distribution System Operator Observatory 2022* · https://publications.jrc.ec.europa.eu/repository/bitstream/JRC132379/JRC132379_01.pdf
5. Regulation (EU) 2024/573 on fluorinated greenhouse gases, SF6 phase-out timetable for MV and HV switchgear · <https://www.nuventura.com/sf6-eu-regulations> and <https://www.esbnetworks.ie/services/get-connected/renewable-connection/f-gas-regulation>
6. PTR Inc., “Europe’s MV GIS Transformation: Charting a Sustainable Course with F-Gas Phase-Out” · <https://ptr.inc/europes-mv-gis-transformation-charting-a-sustainable-course-with-f-gas-phase-out/>
7. §14a EnWG and BNetzA determination BK6-22-300: measurement-based grid state determination for grid-oriented control · <https://www.a-eberle.de/wissen/14a-enwg-ortsnetzstation-abgangsueberwachung/> and <https://envelio.com/de/paragraph-14a-enwg-ueberblick>
8. ACM (NL), binding arrangements with system operators on flexible contracts, congestion management and grid utilisation transparency, 2026 · <https://www.acm.nl/en/publications/acm-and-system-operators-make-binding-arrangements-offering-flexible-contracts-use-congestion-management-and-increased-insight-grid-utilization>
9. TenneT, Time-dependent transport rights (TDTR) · <https://www.tennet.eu/nl-en/time-dependent-transport-rights-tdtr>
10. OECD, *Diagnostic Toolkit for Reducing Regulatory Barriers* (CAPEX bias, grid-enhancing technologies, TOTEX remuneration) · https://www.oecd.org/en/publications/diagnostic-toolkit-for-reducing-regulatory-barriers-to-solar-wind-and-pumped-hydro-storage-in-the-european-union_15f4aed4-en/full-report/grids_e2c2091b.html
11. ACER, *DSO revenue-setting report 2026* (CAPEX bias in distribution remuneration) · <https://www.acer.europa.eu/sites/default/files/documents/Publications/ACER-2026-DSO-revenue-setting-report.pdf>

12. IEA, *Electricity 2026: Grids* (grid-enhancing technologies: DLR, DTR, APFC, topology optimisation, storage as transmission asset) · <https://www.iea.org/reports/electricity-2026/grids>

Belgian source batch (capacity mechanisms, adequacy, flexibility)

1. Elia, *Adequacy & Flexibility Study for Belgium 2026-2036* (reliability standard, CRM design, adequacy gaps, derating factors, flexibility volumes, data centre flexibility assumptions)
2. Elia, *Balancing philosophy applicable to the Belgian electricity system*, January 2026 (flexibility taxonomy, FRCE restoration, aFRR and mFRR activation times)
3. Elia, *Federal Development Plan of the Belgian Transmission System 2024-2034*, Executive Summary
4. Belgian Federal Government, “Belgian State and ENGIE Group enter exclusive negotiations on the acquisition of ENGIE’s nuclear activities in Belgium”, 30 April 2026 · <https://premier.be/en/belgian-state-and-engie-group-enter-exclusive-negotiations-acquisition-belgium-engies-nuclear>

Demand, costs, auctions and weather (October 2026 batch)

1. Eurelectric, *Power Barometer 2025: In Shape for the Future* (EU demand still ~7% below 2021; 1% growth 2024, 2% H1 2025) · <https://www.eurelectric.org/publications/power-barometer-2025-in-shape-for-the-future/>
2. Eurelectric, “Electricity in 2025: solar surges, gas influence fades, but volatility and weak demand persist”, 2 January 2026 · <https://www.eurelectric.org/news/electricity-2025-solar-growth-volatility/>
3. Eurelectric, “Electricity in 2024” (industrial slowdown >50% of the demand shortfall; German industrial consumption –13% 2021 to 2023) · <https://www.eurelectric.org/news/electricity-in-2024-emissions-hit-historic-low-average-prices-declined-but-demand-did-not-pick-up/>
4. Ember, *European Electricity Review 2026*, 22 January 2026 · <https://ember-energy.org/latest-insights/european-electricity-review-2026/>
5. ACER, *Energy Infrastructure Unit Investment Cost Indicators*, 24 April 2026 (transmission and substation unit costs) · <https://www.acer.europa.eu/sites/default/files/documents/Publications/ACER-2026-Unit-Investment-Cost-Indicators.pdf>
6. ACER, *Getting the Signals Right: Electricity Network Tariff Methodologies in Europe*, 26 March 2025 (shallow vs deep connection charging at para 36 and Figure 7) · <https://www.acer.europa.eu/sites/default/files/documents/Reports/2025-ACER-Electricity-Network-Tariff-Practices.pdf>
7. Bundesnetzagentur, offshore wind auctions 2025, centrally pre-examined sites N-10.1 and N-10.2 (no bids by 1 August 2025) · https://www.bundesnetzagentur.de/DE/Beschlusskammern/BK06/BK6_72_Offshore/Ausschr_vorunters_Flaechen/2025/uebersicht_ausschreibungen_2025.html
8. Danish Energy Agency, “North Sea offshore wind tendering procedure: no bids received”, 5 December 2024, and the subsequent CfD retender · <https://ens.dk/en/press/north-sea-offshore-wind-tendering-procedure-danish-energy-agency-has-not-received-any-bids>

9. WindEurope, “Is negative bidding finally dead?”, 7 October 2025 · <https://windeurope.org/news/is-negative-bidding-finally-dead/>
10. DESNZ, *Contracts for Difference Allocation Round 7 Results*, 14 January 2026 · <https://www.gov.uk/government/publications/contracts-for-difference-cfd-allocation-round-7-results/contracts-for-difference-allocation-round-7-results-accessible-webpage>
11. Red Eléctrica, *Análisis Nacional de Cobertura del Sistema Eléctrico Peninsular Español*, October 2025 · https://www.ree.es/sites/default/files/14_OPERACION/Documentos/red-electrica-analisis-cobertura-2025.pdf
12. Orden TED/966/2026, 15 September 2026, creating the Spanish peninsular capacity market · https://www.boe.es/diario_boe/txt.php?id=BOE-A-2026-19329
13. TenneT, *Monitor Voorzieningszekerheid 2026*, 10 June 2026 (Dutch LOLE breaches from 2030, 46 h/yr by 2035) · <https://www.tennet.eu/nl/over-tennet/publicaties/rapport-monitor-voorzieningszekerheid>
14. Nordic TSOs (Energinet, Fingrid, Statnett, Svenska kraftnät), *The Nordic Position on Grid-Forming*, 2025 · https://www.svk.se/4aa98d/siteassets/om-oss/rapporter/2025/the-nordic-position-on-grid-forming_publication.pdf
15. M. Kittel and W.-P. Schill, “Multi-threshold time series analysis enables characterization of variable renewable energy droughts in Europe”, *Communications Earth & Environment*, 10 February 2026 · <https://www.nature.com/articles/s43247-026-03251-2>
16. B. Li, S. Basu, S. J. Watson and H. W. J. Russchenberg, “A Brief Climatology of Dunkelflaute Events over and Surrounding the North and Baltic Sea Areas”, *Energies* 14(20):6508, 2021 · <https://www.mdpi.com/1996-1073/14/20/6508>
17. Bundesnetzagentur and Bundeskartellamt, *Monitoringbericht 2025* (redispatch costs, grid charges, market concentration) · <https://www.bundesnetzagentur.de/DE/Fachthemen/ElektrizitaetundGas/Monitoringberichte/start.html>
18. Bundesnetzagentur, AgNeS draft determination on the general electricity network charge framework, 6 August 2026 · <https://www.bundesnetzagentur.de/1059162>
19. Fraunhofer ISE, *Levelized Cost of Electricity: Renewable Energy Technologies*, July 2024 (capital costs, LCOE and real WACC by technology, including conventional plant) · <https://www.ise.fraunhofer.de/en/publications/studies/cost-of-electricity.html>
20. IRENA, *Renewable Power Generation Costs in 2025*, July 2026 (global installed costs and LCOE; German solar and onshore wind premium) · <https://www.irena.org/Publications>
21. Commission for Regulation of Utilities (Ireland), *Large Energy User Connection Policy Decision CRU2025236*, 12 December 2025 (dispatchable capacity requirement at 10 MVA and above, 80% additional Irish renewables on a six-year glide path, ~50% of Dublin/Meath metered consumption attributable to data centres in 2024, 97% of Irish data centre capacity in Dublin) · <https://www.cru.ie/>

22. Central Statistics Office (Ireland), *Data Centres Metered Electricity Consumption* (22% of national metered electricity in 2024, 23% in 2025) · <https://www.cso.ie/>
23. Bundesnetzagentur, *Netzengpassmanagement* quarterly and annual reporting (congestion management cost €4.2bn 2022, €3.3bn 2023, €2.8bn 2024; redispatch and curtailment reported as overlapping categories) · <https://www.bundesnetzagentur.de/>
24. Bundesnetzagentur, *Netzkostenermittlung / Anreizregulierung* (base-year Ausgangsniveau, Kapitalkostenaufschlag under §10a ARegV applying to distribution operators without time lag) · https://www.bundesnetzagentur.de/DE/Fachthemen/ElektrizitaetundGas/Netzentgelte/Anreizregulierung/WesentlicheElemente/Netzkosten/Netzkostenermittlung_node.html
25. Terna, *Econnexion* connection queue portal, via QualEnergia (101.3 GW of data centre connection requests as at 31 August 2026) · <https://www.qualenergia.it/>
26. Westnetz GmbH, *Spezifikation fernwirktechnische Anbindung, Hochspannung and Technische Anschlussbedingungen Hochspannung* (mandatory IEC 60870-5-101 telecontrol at 110 kV connections: phase voltages and currents, P and Q at the connection point, availability, setpoint feedback in 10% steps, 10 ms timestamp accuracy, cyclic and on-change transmission) · <https://www.westnetz.de/content/dam/revu-global/westnetz/documents/bauen/ihr-weg-zum-netzanschluss/spezifikation-fwt-anbindung-hochspannung.pdf>
27. *From Barrier to Bridge: The Case for AI Data Center / Power Grid Co-Design*, arXiv 2605.03090 (synchronised training ramps of tens to hundreds of MW per second, faster than automatic generation control; the July 2024 Northern Virginia event in which roughly 1.5 GW of voltage-sensitive data centre load tripped within seconds; inference load profiles by contrast) · <https://arxiv.org/html/2605.03090v1>
28. BNetzA, *Positionspapier zur Erhebung von Baukostenzuschüssen* and press release of 20 November 2024 (Leistungspreismodell on contracted connection capacity, five-year average, locational differentiation 100% to 20% permitted for transmission operators only, worked example of ~€60m for a 600 MW connection) · https://www.bundesnetzagentur.de/SharedDocs/Pressemitteilungen/DE/2024/20241120_BKZ.html
29. TransnetBW, *Preisblatt Netzentgelte und Baukostenzuschuss 2026*, and Bayernwerk Netz, *Baukostenzuschuss 2025* (€96.91/kW extra-high voltage; €97.26/kW high voltage, €125.31/kW medium voltage; charged on booked capacity, withdrawal points only) · <https://www.transnetbw.de/>
30. Court of Justice of the European Union, Case C-293/23, judgment of 28 November 2024, and Bundesgerichtshof, EnVR 83/20, 13 May 2025 (the German Kundenanlage definition read down; sale of electricity to customers makes an installation a distribution network); EnWG §118(7) transitional exemption to 1 January 2029 · <https://eur-lex.europa.eu/legal-content/DE/TXT/?uri=CELEX%3A62023CJ0293>
31. Regulation (EU) 2017/1485 Art. 156(9) and BNetzA determination BK6-17-234 of 2 May 2019 (minimum FCR delivery period of 15 minutes; the transmission operators' proposal to extend it to 30 minutes refused) · https://www.clearingstelle-eeg-kwkg.de/sites/default/files/BK6-17-234_beschluss_190502_PRL_Speicher.pdf

32. Senatsverwaltung Berlin, immission-control permits for data centre standby generation: Nonnendammallee (26 November 2021, 200 operating hours a year across 17 units) and Im Marienpark (2 December 2024, 300 hours a year across 48 units, parallel operation prohibited) · <https://www.berlin.de/sen/uvk/umwelt/industrie-und-gewerbeanlagen/>
33. ISEA, RWTH Aachen and Enspired, *Battery Revenue Index* via pv magazine, 23 January 2026 (German FCR prequalified battery capacity ~0.81 GW early 2024 to ~1.35 GW early 2026 against demand of ~0.53 to 0.56 GW per product) · <https://www.pv-magazine.de/2026/01/23/erloespotenziale-fuer-stationaere-batteriespeicher-in-deutschland-sind-2025-zurueckgegangen/>
34. Eurelectric and Eaton on Microsoft's Grange Castle campus, Dublin (UPS batteries in EirGrid fast frequency response, sub-150 ms) and Aurora DC, Finland (FCR-N with one hour of lithium-ion) · <https://www.eurelectric.org/>
35. Westnetz and Bayernwerk Netz digital secondary substation programmes (Westnetz ~2,400 stations by November 2024 against ~60,000 in its area, ~25 a week at ~€100,000 per station; Bayernwerk 5,000th station August 2025, ~13,000 more targeted by 2030) · <https://www.energie-und-management.de/>
36. envelio and energate, *Digital Grid Insights 2025*, and AXXCON *Energiestudie 2025* (fewer than one in ten operators largely automated; ~3% with smart grid technology substantially integrated; smallest municipal operators report the greatest difficulty recruiting specialist staff)
37. NESO and EMR Delivery Body capacity market results, SEMO PCAR2930T-4, Elia CRM auction results of 30 October 2025, and Italian capacity market results for delivery year 2027 (GB T-4 2029/30 £27.10/kW/yr for 40.1 GW; Ireland T-4 2029/30 €135,500/MW/yr with average awarded ~€151,000/MW/yr on locational constraints; Belgium Y-4 €27,300/MW/yr; Italy €47,000/MW/yr) · <https://www.sem-o.com/>
38. CRE, délibérations 2025-265, 2025-268 and 2026-52 (ARENH end 31 December 2025; Versement Nucléaire Universel thresholds of €78/MWh at 50% and €110/MWh at 90%; estimated 2026 nuclear revenue €66.08/MWh and a 2026 unit tariff of zero) · https://www.cre.fr/fileadmin/Documents/Deliberations/2026/260226_2026-52_Tarif_unitaire_minoration_du_vnu_2026.pdf
39. ACM, *Codebesluit niet gebruikte transportrechten* (GOTORK), decision of 16 April 2024 (Dutch operators may reduce contracted transport capacity a connected party cannot show it needs, at medium, high and extra-high voltage in congestion areas) · <https://www.acm.nl/nl/publicaties/codebesluit-niet-gebruikte-transportrechten-gotork>
40. Pacific Northwest National Laboratory, *Redeveloping Coal Power Plants: Data Centers*, PNNL-SA-201505, September 2024 (reuse of an existing point of interconnection may require an interconnection study and equipment upgrades) · <https://www.pnnl.gov/sites/default/files/media/file/PNNL-SA-201505-CoaltoDataCenter.pdf>
41. Start Campus, project description (the Sines campus repurposes seawater cooling infrastructure from the decommissioned plant; secured capacity of 495 MW raised to 1.2 GW by award from REN announced August 2024) · <https://www.startcampus.pt/>
42. Rechtbank Gelderland, summary proceedings, April 2026, reported in NL Times and w.media (70 MW data centre connection near Schiphol refused; no definitive connection agreement concluded; applicant held to

the standard of a professional party aware of the energy landscape) · <https://nltimes.nl/2026/04/29/judge-rules-electricity-grid-haarlemmermeer-full-data-centre-connection-delayed>

43. 50Hertz, Amprion, TenneT and TransnetBW, *Reifegradverfahren* (from 1 April 2026, maturity-scored allocation of transmission-level connections for large consumers, storage and distribution operators; realisation deposit €1,500/MW credited against the connection contribution; generation remains under EEG and KraftNAV) · <https://www.transnetbw.de/de/transparenz/netzzugang-und-entgelt/reifegradverfahren>
44. BNetzA determination BK6-22-300 of 27 November 2023 and accompanying guidance on §14a EnWG network-state monitoring (minute-level data from at least 15% of network connection points, or 7% plus feeder measurement at the local network station; preventive control permitted to end-2028 at the latest) · https://www.bundesnetzagentur.de/DE/Beschlusskammern/1_GZ/BK6-GZ/2022/BK6-22-300/Beschluss/BK6-22-300_Beschluss_20231127.pdf
45. Terna and Il Sole 24 Ore on Italian peak demand (all-time peak 60.5 GW, set 2015) · <https://en.ilsole24ore.com/art/peak-electricity-consumption-hits-a-record-55-gw-in-2026-blackouts-continue-AIV3i6qD>
46. Herbert Smith Freehills Kramer, note on the Reifegradverfahren (scope covers large consumers and storage at transmission level with no capacity threshold; €50,000 application fee; €1,500/MW realisation security credited against the connection contribution) · <https://www.hsfkramer.com/>

Part 3 added 6 October 2026. Belgian source batch integrated 7 October 2026. Demand, cost, auction and weather batch integrated 7 October 2026. Figures current as of that date. Gas, carbon and power prices move daily; verify before using any number in a model. The coal cost calculation in Chapter 03 uses an assumed API2 range rather than a verified print.

Bernd Fastenrath runs Altered Principles, a strategy consultancy working with European industrial operators in DACH and Benelux. He was previously VP Strategy at ABB E-Mobility. He advises companies in the medium-voltage equipment and grid technology segment on go-to-market strategy, and readers should weigh that interest accordingly.

Edition 2026. Figures move; verify before using any number in a model. Full sources are listed at the end of the document.