THE BARATELLI INSTITUTE

MENTORING AT SCALE

FLAGSHIP WORKING PAPER · INSTITUTE POLICY DISCUSSION

Energy Matters

The American Energy Stack: Eleven Tools, One Grid, and the Constraints That Actually Matter.

Published July 2026

Baratelli Institute working paper. Not investment advice.

Every quantitative reference traces to a filed SEC document, publicly-reported transaction disclosure, or industry-standard practitioner benchmark.

Philip A. Baratelli, CPA, MBA · Founder, The Baratelli Institute

PONTE VEDRA BEACH, FLORIDA · BARATELLIINSTITUTE.COM

Contents

A Baratelli Institute Working Paper *July 2026*

> *"Do what you can, with what you have, where you are."* — *Theodore Roosevelt*

The philosophy behind every section in this paper is that the American energy discourse is dominated by ideological framings — solar-versus-gas, renewables-versus-nuclear, transition-versus-continuity — that no serious operator actually uses when they have to build something. Every serious operator uses the full stack. This paper walks the stack, one tool at a time, and names the specific constraint that stops each tool from being deployed faster than it is.

Disclosure

The author holds a position in Cleveland-Cliffs Inc. (NYSE: CLF) as of July 18, 2026, which is discussed in Section 1 (steel-in-panels), Section 8 (metallurgical coal), and referenced in cross-links.

The findings here admit opposing readings. Nothing in this paper is a recommendation to buy or sell any security, and nothing is a recommendation to file, litigate, appropriate, or legislate anything. The paper is analysis. The reader is entitled to weigh the disclosure accordingly. We welcome the reader distributing this paper to any and every individual or entity they think will help move the American energy industry forward. You are also welcome to link to the Baratelli Institute as you wish. Every ticker mentioned in the operator subsections is provided so the practitioner reader can locate the filed disclosures and form their own view.

About this paper

The Baratelli Institute publishes to share ideas that help move the American energy industry forward. We do not compete with the operators, developers, utilities, financiers, and manufacturers whose work is described in these pages — they are the audience. We are a publisher, and our contribution is analysis: a considered read of the constraints, tradeoffs, and mechanics that a working practitioner has to navigate, offered so that the people actually building things have one more careful reference to work from.

The point of this paper

Energy Matters is a plain-English overview of the American energy stack, written for a serious reader who is not already an energy specialist. The paper walks the eleven tools that make up the stack; describes the cost, the deployment mechanics, and the specific constraint that gates each one; names the operators worth watching; and closes each section with the Institute's read on what a practitioner should take away.

The purpose is understanding — not investment advice, not policy prescription, not a bid to impress the energy industry. A reader who finishes this paper should know how the American energy industry actually works, where its costs are, and what is stopping each part of it from being built faster than it is. The Institute's contribution is the framework and the synthesis — a considered read that helps a serious outside reader see how the pieces fit together.

The Eleven Tools

The American energy stack is not a single technology decision. It is eleven distinct tools, each with its own cost curve, its own operator set, its own deployment constraint, and its own political-economy problem. A serious grid uses all of them, in different proportions, on different timelines, for different jobs. This is not an editorial claim; it is what the actual grid mix looks like when you read the EIA data. The eleven tools this paper covers:

Table 1 — The Full Stack Summary

Table 1.

#	TOOL	PRIMARY JOB	DEPLOYMENT TIMELINE	THE CONSTRAINT
1	Solar (utility + distributed)	Bulk energy, cheapest incremental kWh	2-4 years	Interconnection queue
2	Wind (onshore + offshore)	Bulk energy, complementary generation profile	3-6 years onshore, 6-10 years offshore	Permitting, PTC uncertainty

#	TOOL	PRIMARY JOB	DEPLOYMENT TIMELINE	THE CONSTRAINT
3	Storage (BESS, pumped)	Firming, arbitrage, capacity	12-24 months BESS, 8-12 years pumped	Duration economics past 4 hours
4	Nuclear — utility-scale	Firm baseload, low-carbon	8-15 years new-build, 2-4 years restart	Capex, NRC licensing, PPA structure
5	Nuclear — SMR (small modular)	Distributed firm power, industrial + defense	5-8 years post-certification	Pre-commercial: no operating US SMR
6	Natural Gas — combined-cycle	Firm dispatchable, mid-carbon	3-5 years	Turbine backlog (GE Vernova)
7	Natural Gas + Diesel — mobile	Emergency, remote, bridge power	60-180 days deploy	Fuel logistics
8	Coal	Legacy baseload, metallurgical	(existing fleet retiring; no new-build)	Political / permitting for power; met coal is a different market
9	Hydro	Legacy baseload + pumped storage	(no material new large hydro; pumped storage 8-12 years)	Siting, tribal consultation, environmental review
10	Grid Transmission	The precondition for every other tool	8-12 years new HVDC line	State-level siting, cost allocation
11	Efficiency / Demand-Side	The cheapest kWh not consumed	12-36 months per project	Fragmented decision-makers

The point of the table above is not scorekeeping. It is that every serious American utility uses many of these tools in combination; the largest data-center operators are moving toward using most of them; and the reader who thinks in terms of a preferred subset — solar-and-wind-only, or nuclear-and-gas-only, or coal-forever — is reading the discourse, not the actual grid. Every tool has a real job. The paper's job is to walk each one plainly.

The rest of this paper walks each tool in turn.

Deployment Modes — Behind the Meter, In Front of the Meter, and Everything in Between

Before walking the eleven tools, one framing concept has to be established because it determines which tools are usable, on what timeline, by whom. That concept is the meter. Every energy asset in America is deployed in one of two positions relative to the utility meter that separates the customer's premises from the utility's grid: **in front of the meter** (IFM), meaning the generator connects to the transmission or distribution grid and sells electricity as a wholesale or retail resource, or **behind the meter** (BTM), meaning the generator sits on the customer's side of the meter and delivers power directly to the load without transiting the utility grid.

The distinction is not academic. A behind-the-meter generator never enters the utility interconnection queue. That single sentence explains why behind-the-meter deployment has become the fastest-growing category of American electricity infrastructure in the 2024–2026 window. When the queue takes eight years and a project needs power in eighteen months, the project builds behind the meter.

Behind-the-meter deployment splits into three scales, each with its own economics, its own regulatory environment, and its own operator set. The rest of this paper's tool-by-tool discussion references each scale where relevant, so a short walk of all three is warranted before proceeding.

Large-scale BTM — hyperscaler power islands

The largest single BTM category in America in 2026 is the hyperscaler data-center power island. The mechanic is that a generator (gas turbine, small nuclear reactor, or on-site solar-plus-storage) is co-located on the same physical site as the data center's IT load, and the generator's output feeds the data center directly without ever touching the utility grid. As of mid-2025, forty-six data-center projects totaling roughly fifty-six gigawatts of announced behind-the-meter generation had been publicly disclosed — ninety percent of them announced in 2025 alone (Institute compilation from Lawrence Berkeley National Laboratory data-center tracking and Wood Mackenzie hyperscaler reporting). Notable examples: Meta's roughly twenty-seven-gigawatt behind-the-meter target across multiple sites (as reported in Meta's public sustainability disclosures and press coverage); Amazon's Talen Energy nuclear PPA at the Susquehanna facility (a hybrid IFM/BTM structure); Microsoft's Constellation Three Mile Island Unit 1 restart PPA (also a hybrid structure, since TMI-1 also delivers to the PJM grid). The dedicated BTM subset of these deals is what avoids the eight-year queue entirely — the load and the generation share a substation that never needs a grid interconnection agreement for the primary energy flow.

Regulatory posture matters greatly here. Some states have moved to permit BTM at large industrial scale explicitly (Texas, Georgia, Virginia); some are still working through whether BTM is subject to standby charges from the local utility (Pennsylvania, Ohio); and some have imposed material friction on it (California, at least at scale). Practitioners underwriting a BTM data-center project should treat the specific state and utility service territory as the first analytical block.

Medium-scale BTM — industrial and campus

At medium scale — five to two hundred megawatts of on-site generation — behind-the-meter takes the form of factory microgrids, hospital cogeneration, university district energy systems, oil refinery on-site generation, and paper mill combined-heat-and-power. This category has existed for a hundred years and remains the largest installed base of BTM generation in the United States by megawatt count. Its 2020s renewal is driven by three factors: rising commercial electricity rates in restructured markets, resilience concerns after regional grid failures, and IRA §48 and §45 tax credits that extended to BTM projects for the first time. Operators in this space are typically privately held, but adjacent public equities — the industrial gas turbine and reciprocating engine manufacturers (Cummins, Rolls-Royce, Caterpillar, Wärtsilä) and the systems integrators (Bloom Energy, PowerSecure) — are the readable market signal.

Small-scale BTM — residential and small commercial

At the smallest scale, behind-the-meter is what most Americans intuitively understand as “solar on my roof and a battery in my garage.” Residential rooftop solar, home battery storage, and small commercial rooftop are all BTM by construction — they deliver power to the building’s load, and export to the grid only when generation exceeds consumption (and even that export is optional in some designs). Distributed rooftop solar economics are worse than utility-scale on a levelized basis, as Section 1 discusses, but the resilience value of a home battery in a wildfire-prone or hurricane-prone region has become a material driver of adoption independent of the economics. Section 11 (Efficiency and Demand-Side) discusses the parallel small-scale story of demand-side management, which does not generate electricity but reduces the electricity that needs to be delivered.

The practitioner posture — build behind the meter first

The clearest example of BTM-first operating posture in the American energy market is Elon Musk. Across Tesla, SpaceX, and xAI, the pattern is consistent: when a Musk enterprise needs power for a new factory, launch facility, or compute cluster, the default starting question is not “how long is the utility interconnection queue” but “what generation can we put on the site directly." The Tesla Gigafactory in Nevada was designed with rooftop solar and on-site Powerpack storage. Starbase in Boca Chica operates on a substantially self-generated microgrid. The xAI Colossus supercomputer in Memphis reportedly deployed with on-site gas turbines and batteries because the utility interconnection timeline was incompatible with the training-cluster schedule. Musk's personal acquisition of APR Energy in 2026 — confirmed by a Federal Trade Commission filing in May 2026 and reported at an estimated $1 billion (Section 7) — brought a mobile behind-the-meter fleet of gas and diesel turbines with more than 1.1 gigawatts of on-site power under his direct control. APR advertises 30-to-60-day deployment timelines for infrastructure that would otherwise take years to interconnect through a utility. The transaction sits inside a broader pattern of Musk-adjacent BTM decisions: xAI's Colossus supercomputer in Memphis already runs on on-site methane turbines (a deployment that has drawn Clean Air Act litigation from advocacy groups). We read the APR acquisition as, at least in part, a bet that mobile BTM generation is the fastest way to put megawatts on the ground for a customer who cannot wait for a grid connection.

The point is not that Musk's specific projects are the model for every operator. The point is that the operator who has actually built things at scale in the 2020s treats behind-the-meter as the default, not the exception. That practitioner posture — build the generation with the load, treat the grid connection as an optional add-on rather than the precondition — is the deployment mode that most of the American energy discourse still treats as unusual. The eleven tools below can each be deployed either in front of the meter or behind it. The operator who understands which deployment mode is available to which tool at what scale has a material informational advantage.

One more framework — the regulated-utility layer (IRP)

Alongside behind-the-meter and in-front-of-the-meter, a third framework governs a large share of American energy deployment: **Integrated Resource Planning (IRP)**. Regulated utilities in most states file multi-year IRPs with their state public utility commissions, describing what generation they intend to build, retire, or purchase over a rolling planning horizon. Every regulated-utility investment discussed in this paper runs through an IRP proceeding, a public utility commission approval, and a rate case or PPA before it is built. The reader who wants to understand where a regulated utility's next generation will come from should read that utility's most recent IRP filing before doing anything else.

A related distinction worth naming: within the behind-the-meter world, the difference between a **power island** (a fully behind-the-meter site with no grid tie for the primary energy flow) and an **on-site-with-grid-tie** deployment (a hybrid site where some generation flows to the customer's load and some flows to the grid) is significant. Most hyperscaler deals reported as "behind the meter" are in practice hybrid. The pure power-island cases exist but are less common than press coverage suggests.

The reader's map

When the sections below discuss operator economics, deployment timelines, and constraints, the BTM/IFM framing determines which numbers apply. A Section 6 (natural gas combined-cycle) BTM data-center project bypasses the interconnection queue but still needs a gas pipeline. A Section 5 (SMR) BTM project at a data-center site bypasses the queue but still needs NRC certification. A Section 1 (solar) BTM residential project bypasses the queue entirely and interacts with the utility only through a net-metering tariff. Where the meter sits changes the whole story.

1. Solar

The cost anchor

Utility-scale solar photovoltaic is the cheapest incremental kilowatt-hour available to the American grid in 2026. Lazard's Levelized Cost of Energy Analysis 17.0 (2024) puts unsubsidized utility-scale solar PV at $29–$92 per MWh, with a mid-point around $60/MWh — below new combined-cycle gas across most of the interconnection queue, and below new nuclear by a factor of two to four. Applying the Investment Tax Credit under the Inflation Reduction Act pulls the subsidized range down to $0–$40/MWh at the low end. Distributed rooftop solar is materially more expensive: $122–$284/MWh unsubsidized. The premium is not primarily hardware. A commercial rooftop installation has a physical floor near $0.41 per watt and a market price near $1.71 per watt — and roughly seventy-six percent of that price is organization: permitting, customer acquisition, sales margin, installer overhead. When practitioners talk about "cheap solar," they mean utility-scale, not rooftop.

How a project actually gets built

A utility-scale solar project moves through five overlapping stages: site control and land assembly, interconnection application and study, offtake contracting (a Power Purchase Agreement with a utility or a corporate offtaker), permitting at the county or state level, financing (a mix of tax-equity investors, debt, and sponsor equity), and construction. Construction itself, once financing closes, is twelve to twenty-four months for a utility-scale project. But the full path from site control to commercial operation runs closer to thirty-six to forty-eight months under current conditions, and the excess time above construction is spent almost entirely on the interconnection queue.

The most active developer in the United States is NextEra Energy Resources, the unregulated subsidiary of NextEra (NEE).

The deployment constraint

The deployment constraint is not price and it is not panel supply. The constraint is the interconnection queue. In PJM, the timeline from interconnection application to commercial operation rose from under two years in 2008 to over eight years by 2025. Average interconnection costs in the PJM queue rose eightfold, from $29/kW to $240/kW. FERC's target for executing an interconnection agreement is eight to eleven months; the national average is twenty-five months, and PJM averages forty. Projects sited in data-center load-growth zones wait three to four years just for the study process to complete. For contrast: ERCOT's average time to an executed interconnection agreement is roughly twenty months — half of PJM's forty. The queue is not physics; it is procedure.

The Federal Energy Regulatory Commission's Order 2023, issued in July 2023, is the first overhaul of interconnection procedures since 2003. It moved the process to a cluster-study model with firm study deadlines and penalties for missed timelines. The reforms are still being operationalized at the RTO level in 2026, and their throughput impact is not yet fully visible in queue metrics. Until Order 2023 flows through and transmission is built to absorb the projects that clear the queue, cheap solar does not translate into deployed solar. This is exactly the argument the Institute develops at length in *The Deployment Plan* — Move 1 (connect and manage under FERC Order 845's provisional service and ERIS), Move 2 (surplus interconnection at retiring thermal plants), and Move 3 (reconductoring the existing grid).

Operators and tickers

The public equities that give practitioners direct or indirect exposure to utility-scale solar in the United States:

Table 1.1 — Solar operators (representative, not exhaustive)

Table 2.

TICKER	COMPANY	EXPOSURE
NEE	NextEra Energy	Largest US utility-scale solar developer via NextEra Energy Resources; regulated Florida utility (FPL)

TICKER	COMPANY	EXPOSURE
FSLR	First Solar	US-based thin-film module manufacturer; §45X production credits
RUN	Sunrun	Largest US residential solar installer
ENPH	Enphase Energy	Microinverters + residential storage
CSLR	Complete Solaria	Acquirer of SunPower's residential assets (2024 bankruptcy)
AGR	Avangrid	Iberdrola-controlled US utility with significant renewables
DUK, SO, XEL, AEP	Regulated utilities	Direct solar deployment inside regulated rate-base
CLF	Cleveland-Cliffs	Sole domestic producer of grain-oriented electrical steel; see Institute view below and disclosure at front matter

2024-26 activity

The 2024–2026 window has been the largest solar deployment period in American history. The Energy Information Administration reported that solar accounted for roughly 60% of all new generation capacity added to the US grid in 2024, a share that held into 2025. First Solar's US thin-film manufacturing has continued to expand under the IRA's §45X advanced-manufacturing production credits, which materially reshaped domestic module economics; First Solar, Enphase, and battery cell producers have all cited §45X directly as changing their capital plans. NextEra brought material new solar capacity online each of the last two years — the specific figure appears in NEE's annual filings and investor presentations, which practitioners should read directly rather than relying on secondary summaries.

The Institute view

Solar is the cheapest generating asset on the American grid on a levelized basis, and it is being throttled by the interconnection queue and by the transmission that isn't there. The correct read for practitioners is not "is solar cheap enough to compete." That question is already answered. The correct read is "how fast can it deploy against the specific constraints that gate it." Those constraints are, in order: interconnection queue reform, transmission siting and permitting, transformer manufacturing capacity, and workforce.

The Institute's separate disclosure of a position in Cleveland-Cliffs is noted here for completeness. It is not, however, a claim that CLF is a bottleneck-resolving investment for solar deployment. *The Deployment Plan*, Move 4, examines the transformer bottleneck at length and concludes — from Cleveland-Cliffs' own disclosures — that grain-oriented electrical steel is not the constraint. Cliffs' fiscal 2025 electrical-steel shipments fell three percent to 552,000 net tons, and Cliffs itself told the Secretary of Commerce in September 2025 that its electrical-steel market is soft and its largest segment stagnant with excess inventory. The transformer constraint is transformer-manufacturing capacity and intellectual property, held by GE Vernova (GEV), Eaton (ETN), Siemens Energy, and Hitachi Energy. Readers drawn to the CLF story through grid-transformer coverage should read *The Deployment Plan* Move 4 before forming a view.

Cross-links

- *The Deployment Plan* — the full companion working paper covering Move 1 (connect and manage), Move 2 (surplus interconnection), Move 3 (reconductoring), Move 4 (transformer manufacturing), and Move 5 (clamped-export residential)
- Cleveland-Cliffs case study — for the operating and financial profile of the electrical-steel and met-coal producer
- Berkshire Hathaway Energy — the regulated-utility comparator for solar rate-based deployment inside an integrated utility

2. Wind

The cost anchor

Onshore wind is competitive with utility-scale solar on a levelized basis. Lazard 17.0 puts unsubsidized onshore wind at \$27–\$73/MWh with a mid-point near \$50/MWh — comparable to solar. Subsidized onshore wind under the Production Tax Credit falls to \$0–\$45/MWh. Offshore wind is a different economic reality: unsubsidized offshore wind lands at \$74–\$139/MWh per Lazard 17.0, roughly two to three times onshore. Post-IRA supply chain inflation (turbines, monopile foundations, installation vessels) has pushed real US offshore economics further above the Lazard estimates, which is why multiple US offshore projects were cancelled or restructured in 2023–2025.

How a project actually gets built

Onshore wind: land assembly (often through lease optionality), Federal Aviation Administration and Fish & Wildlife review, county-level zoning, PPA, tax-equity financing, construction. Typical timeline 3–6 years. Offshore wind: Bureau of Ocean Energy Management (BOEM) lease auction, environmental review under NEPA (multi-year), Construction and Operations Plan approval, port infrastructure buildout, vessel contracts, and a transmission interconnection that in most cases requires a purpose-built substation onshore. Typical timeline 6–10 years and rising.

The deployment constraint

Onshore wind's binding constraint is now permitting and local opposition, not economics. Counties in the wind belt have adopted more restrictive setbacks and moratoriums, and Fish & Wildlife Service consultations on eagle and bat takes have lengthened project timelines. Offshore wind's binding constraint is supply chain: not enough Jones Act-compliant installation vessels, no domestic monopile fabrication at the scale needed, and turbine suppliers (Siemens Gamesa, GE Vernova, Vestas) reporting negative margins on offshore contracts through 2024. Federal PTC extensions under the IRA improved economics but did not cure the supply chain gap.

Operators and tickers

Table 2.1 — Wind operators and turbine suppliers

Table 3.

TICKER	COMPANY	EXPOSURE
NEE	NextEra Energy	Largest US wind operator with roughly 24 GW of wind capacity
GEV	GE Vernova	Largest US-based wind turbine OEM (post-spin from GE)
AGR	Avangrid	Iberdrola-controlled; significant onshore and offshore US portfolio
TPIC	TPI Composites	Turbine blade manufacturer
BKR	Baker Hughes	Turbine services and adjacent industrial gas exposure
DNNGY	Ørsted (ADR)	Danish offshore wind developer; large US offshore portfolio

TICKER	COMPANY	EXPOSURE
BEP, BEPC	Brookfield Renewable	Diversified renewables including wind

2024-26 activity

Onshore US wind additions moderated in 2024–2025 to roughly 6–8 GW per year — well below the 2020–2021 peak of 14 GW. Offshore is materially behind schedule versus the Biden administration's 30-GW-by-2030 target: as of mid-2026, roughly 250 MW of offshore wind is operating in US waters (Vineyard Wind 1 partial, South Fork), against about 3 GW under construction, and multiple projects (Ørsted's Ocean Wind 1 and 2, Avangrid's Commonwealth Wind) either cancelled or restructured. Empire Wind proceeded under a renegotiated PPA. The pattern in the offshore data is that projects with locked-in PPAs signed before the 2022 inflation shock are proceeding; projects that need new PPAs are struggling to clear economically.

The Institute view

Onshore wind is a mature, deployable, cost-competitive resource whose deployment is now gated by county-level land-use politics, not by economics. Offshore wind is currently uneconomic in the US at post-IRA supply chain prices in most locations. Practitioners underwriting US offshore should treat 2026-vintage cost estimates as the anchor, not 2020-vintage estimates. The offshore story rebuilds around 2028 when domestic vessels and monopile fabrication scale up — not before.

Cross-links

- Berkshire Hathaway Energy — the utility-scale wind operator in the Institute case library
- GE Vernova WACC page — the cost of capital analysis for the largest US wind turbine OEM
- *The Deployment Plan* — the permitting and interconnection arguments apply to wind as well as solar

3. Storage (Battery Energy Storage Systems)

The cost anchor

Utility-scale lithium-ion battery storage moved from a boutique technology in 2020 to a mainstream grid asset by 2026. Lazard's Levelized Cost of Storage 9.0 (2024) puts unsubsidized four-hour lithium-ion battery storage at $130–$296 per MWh delivered. The IRA's §48E standalone storage Investment Tax Credit, applied in 2023, meaningfully changed the economics by extending the ITC to storage without requiring co-location with solar. The §45X advanced-manufacturing production credit further improved cell economics for domestic producers. On a four-hour duration basis, storage has crossed the threshold at which it is deployed as a peaker replacement in most ERCOT, CAISO, and PJM markets. Longer-duration storage — eight-hour and above — has not crossed that threshold; it awaits a different chemistry.

How a project actually gets built

BESS deployment is fast by grid-asset standards. Once financing closes and interconnection is secured, a utility-scale BESS project reaches commercial operation in twelve to twenty-four months — materially faster than any thermal or renewable generation asset. Site control is simpler because BESS has a small footprint relative to solar or wind (an acre or two for a 100 MW project rather than hundreds). Permitting is generally faster because the local land-use footprint is smaller. Interconnection remains the largest single time block, and that block is shared with solar and wind: the queue is the queue.

Increasingly, storage is being deployed as a co-located asset either with a solar project (which shares the interconnection point and enables both ITC eligibility and hybrid PPA structuring) or as a standalone project. The Deployment Plan's Move 2 — surplus interconnection at retiring thermal plants — is a natural fit for standalone storage, which can be sized to the retired plant's grid interconnection and dispatched around whatever remains of the thermal asset.

The deployment constraint

The binding constraint on battery storage is duration economics past four hours. Lithium-ion cell chemistry does not scale linearly with duration; an eight-hour lithium-ion project is not simply twice a four-hour project because the cell energy density does not favor extended dispatch cycles. The industry's answer is a different chemistry: flow batteries (vanadium redox, iron flow) and long-duration technologies (iron-air, compressed air, thermal, hydrogen) targeting one hundred hours or more. None of these long-duration technologies is at unsubsidized cost parity with lithium-ion four-hour today, and the twenty-billion-dollar-plus infrastructure investments implied by scaling them are still being made. Practitioner readers should treat any project claim on hundred-hour storage economics as pre-commercial until the specific project reports first-year operating data.

Operators and tickers

Table 3.1 — Storage operators, integrators, and manufacturers (representative)

Table 4.

TICKER	COMPANY	EXPOSURE
TSLA	Tesla	Tesla Energy — Megapack utility-scale storage; residential Powerwall
FLNC	Fluence Energy	Utility-scale storage integrator; joint venture Siemens/AES origin
ENPH	Enphase Energy	Microinverters and residential storage
EOSE	Eos Energy	Zinc-hybrid stationary battery (long-duration)
RUN	Sunrun	Residential storage installer
STEM	Stem Inc.	Behind-the-meter AI-optimized storage
GEV	GE Vernova	Grid integration and storage systems
ETN	Eaton	Storage-adjacent power management and grid equipment
—	Form Energy (private)	Iron-air 100-hour long-duration; DOE-backed
NEE	NextEra Energy	Largest US storage deployer via NEER

2024-26 activity

The 2024–2026 window has been the largest storage deployment period in American history. Standalone storage additions roughly tripled from 2024 to 2025 as the §48E ITC took full effect (EIA Form 860). Tesla’s Megapack sold through 2027 production in mid-2025. Fluence and Wärtsilä both reported record project backlogs. Regulatory frameworks under FERC Order 841 (issued 2018 but implemented over the following years) now require every wholesale market to allow storage to bid as a resource on the same terms as generation — a structural change that took roughly five years to work through market-by-market. In parallel, several state public utility commissions (California, Arizona, Massachusetts) approved procurement targets that materially raised expected demand for storage.

The Institute view

Battery storage is the enabling technology that makes utility-scale solar and wind fully deployable across the American grid. Four-hour lithium-ion is economic today and being deployed. Longer duration — eight-hour and above — is the open technology question, and the answer is not yet known. Practitioners looking at storage as an asset class should focus on: (i) which markets have implemented Order 841 in a way that allows storage to earn capacity as well as energy revenue, (ii) which manufacturers have secured domestic cell supply under §45X, and (iii) which developers have interconnection queue positions at retiring thermal sites where surplus interconnection makes deployment fast. The technology is real. The bottleneck is the same interconnection queue that gates solar and wind.

Cross-links

- *The Deployment Plan*, Move 2 — surplus interconnection at retiring thermal plants (co-located storage is one of the most natural applications)
- Section 1 (Solar) above — solar-plus-storage hybrid economics
- Berkshire Hathaway Energy — utility-scale storage deployment inside a regulated utility

---

4. Nuclear — Utility-Scale

The cost anchor

New utility-scale nuclear is expensive on a levelized basis. Vogtle Units 3 and 4 — the two most recent US new-build reactors, completed 2023 (Unit 3) and 2024 (Unit 4) — came in at approximately $36 billion for 2,234 MW of capacity, roughly $16,000 per kW of installed capacity. That is materially higher than the original estimates and is the reference point every subsequent US new-build discussion runs against. On a levelized-cost basis, new-build nuclear from the Vogtle experience is above $150 per MWh unsubsidized, and higher on a first-of-a-kind basis than a subsequent similar reactor would be. The Inflation

Reduction Act's §45U production tax credit for nuclear provides operating support for the existing fleet at up to $15 per MWh (subject to gross-receipts phase-out), which extends existing plant economics but does not by itself justify new build.

Restart economics are fundamentally different from new-build economics. A restart of a recently retired reactor — Palisades, Three Mile Island Unit 1 — reactivates an existing licensed site with existing grid interconnection, existing containment, and a substantial fraction of the original capital cost already sunk. Restart timelines run two to four years and total costs run in the low single-digit billions rather than the tens of billions of a greenfield project.

How a project actually gets built

A new-build utility-scale reactor moves through: site selection and Early Site Permit, Combined Construction and Operating License application to the Nuclear Regulatory Commission (multi-year review), engineering and procurement (with long-lead reactor pressure vessel and turbine orders), construction, fuel load, and commissioning. Total timeline eight to fifteen years is the historical experience. A restart moves through: NRC license reinstatement, fuel procurement (with enriched uranium supply becoming increasingly complex post-Russia sanctions), condition assessment and refurbishment, and grid re-interconnection. Restart total timeline two to four years.

The offtake question is central. A conventional utility PPA structure — a utility takes the output at a regulated rate — is what Vogtle 3 and 4 use. The 2024 innovation is the hyperscaler PPA: Microsoft's twenty-year offtake at the restarted Three Mile Island Unit 1 (announced September 2024, targeting a 2028 restart) commits Microsoft to take the entire output of the plant for its data-center loads. This structure — a corporate offtaker with a compute-driven demand curve underwriting a nuclear restart — is new, and it is one of the few structures under which nuclear economics arithmetic actually closes at 2024–2026 cost levels.

The deployment constraint

The binding constraint on new-build nuclear is capital cost combined with construction-timeline risk. Vogtle 3 and 4 finished materially over budget and years late; that experience is what a bank underwriting a new-build project reads first. The NRC licensing process is a separate constraint but not the binding one — it has been shortened under recent reforms and is not the reason new projects stall. Enriched uranium supply is emerging as a real constraint given the 2022 disruption of Russian supply; HALEU (high-assay low-enriched uranium) production capacity is still being built out and is the specific supply-chain question for the SMR generation of reactors (see Section 5).

For restart projects, the binding constraint is different: it is the specific commercial arrangement between the plant owner, the offtaker, and the state utility regulator.

Operators and tickers

Table 4.1 — Nuclear operators and adjacent equities (representative)

Table 5.

TICKER	COMPANY	EXPOSURE
CEG	Constellation Energy	Largest US nuclear operator (~22 GW); Microsoft TMI PPA counterparty
VST	Vistra	Comanche Peak nuclear; also Texas gas fleet
SO	Southern Company	Georgia Power parent — Vogtle 3 & 4 owner
DUK	Duke Energy	Nuclear fleet across Carolinas
XEL	Xcel Energy	Nuclear fleet in Minnesota
ETR	Entergy	Nuclear fleet in the South
PEG	Public Service Enterprise Group	Salem and Hope Creek nuclear
PCG	PG&E	Diablo Canyon (life extension in progress)
CCJ	Cameco	Uranium mining and fuel processing
BWXT	BWX Technologies	Naval reactor and defense nuclear

2024-26 activity

The 2024–2026 window is arguably the most active period for US nuclear since the 1970s. Vogtle Unit 4 achieved commercial operation in April 2024. Constellation announced the Microsoft Three Mile Island Unit 1 restart in September 2024, with a targeted 2028 return to service and a twenty-year PPA. Palisades' restart, backed by Holtec International, continues to work through regulatory and commercial questions. PG&E secured the Diablo Canyon life extension. Congressional support for advanced reactor demonstrations under the ADVANCE Act (2024) provided regulatory clarification for SMRs (Section 5). Hyperscaler interest in nuclear PPAs has broadened beyond Microsoft to include announcements from Amazon and Google exploring similar structures.

The Institute view

Nuclear is back — but the return is bifurcated. Restarts and existing-fleet life extensions have entered a period of genuine growth, underwritten by a combination of §45U tax credit support, state-level clean-energy policy, and hyperscaler PPAs providing corporate offtake with pricing above what a state utility rate would allow. New-build economics remain challenging outside of specific structural cases: a strategic hyperscaler offtaker willing to underwrite a twenty-year fixed-price PPA, or a state legislature committed to a rate-based cost recovery. Practitioners looking at nuclear as an investment thesis should treat the restart universe as a different asset class from the new-build universe.

Cross-links

- Section 5 (Nuclear — SMR) below — the pre-commercial small modular story
- Constellation Energy — the operating case study on the largest US nuclear operator and Microsoft's PPA counterparty
- Berkshire Hathaway Energy — regulated utility comparator on how nuclear fits into a rate-based portfolio

5. Nuclear — Small Modular Reactors (SMR)

The cost anchor

Small modular reactors are pre-commercial in the United States. No SMR is operating on the American grid as of mid-2026, and no US SMR has completed the full NRC construction and operating license process. Cost anchor claims for SMR are therefore projections, not observations. Vendor projections put mature-deployment SMR levelized cost of energy in the $60–$120 per MWh range, with the low end dependent on serial-manufacturing learning curves that have not yet been demonstrated at any US facility. The single US SMR project that reached advanced design and site selection — NuScale's Utah Associated Municipal Power Systems (UAMPS) project — was cancelled in November 2023 when the projected cost per MWh rose past its offtaker's acceptance threshold. That cancellation is the reference data point for the current commercial reality: SMR cost projections are highly sensitive to first-of-a-kind execution and to the specific offtake structure.

How a project actually gets built

The SMR development path is: design certification with the NRC (which several vendors have achieved or are pursuing), site selection with an Early Site Permit, Combined License Application (COLA), engineering and procurement, module manufacturing at a centralized facility, module transport to site, on-site assembly, fuel load, and commissioning. Timelines for a first-of-a-kind US SMR realistically run five to eight years from a signed offtake and secured financing — which is materially faster than a conventional new-build nuclear reactor but not fast in the context of near-term data-center power needs. HALEU fuel supply is a specific constraint for the advanced-reactor designs (Oklo, X-energy, others) that require enrichment above conventional 5% U-235.

The deployment constraint

The binding constraint on US SMR deployment is not technology and it is not NRC licensing — it is offtake commitment at a price that clears vendor economics. The NuScale-UAMPS cancellation happened not because the reactor could not be built but because the projected LCOE (which had risen from an earlier $58/MWh estimate to $89/MWh by cancellation) exceeded what the municipal utility offtaker was willing to pay. The 2024–2026 window has seen offtake commitments emerge from a new counterparty class: hyperscaler data-center operators willing to sign twenty-year fixed-price PPAs at $80–$120/MWh in exchange for guaranteed low-carbon baseload. Amazon Web Services signed with X-energy in October 2024. Google signed with Kairos Power in October 2024. Both are behind-the-meter or hybrid deployments at data-center sites. This is why the SMR deployment discussion is now inseparable from the hyperscaler capital-expenditure discussion.

Operators and tickers

Table 5.1 — SMR developers and adjacent equities (representative)

Table 6.

TICKER	COMPANY	EXPOSURE
SMR	NuScale Power	Light-water SMR; NRC-certified 50 MWe design; pursuing commercial-scale VOYGR
OKLO	Oklo Inc.	Sodium-cooled fast reactor (Aurora); HALEU-fueled; DOE site permit
BWXT	BWX Technologies	Naval reactor supplier; Advanced Reactor Demonstration Program work

TICKER	COMPANY	EXPOSURE
GEV	GE Vernova	BWRX-300 boiling-water SMR; Canadian first deployment path
—	X-energy (private)	HTGR (Xe-100); Amazon investment; Dow industrial deployment
—	Kairos Power (private)	Fluoride-salt-cooled; Google offtake; DOE-supported
—	Terrestrial Energy (private)	Integral molten salt reactor
CCJ	Cameco	Uranium fuel supply

2024-26 activity

The 2024–2026 window shifted SMR economics from municipal-utility-led (which failed at NuScale-UAMPS) to hyperscaler-led. Amazon's $500 million investment in X-energy and 5 GW of anticipated deployment (October 2024) and Google's Kairos offtake (also October 2024) both target behind-the-meter or hybrid deployment at data-center sites. The ADVANCE Act of 2024 provided the NRC with new authority and appropriations to streamline advanced reactor licensing. DOE loan guarantees have flowed toward the leading SMR vendors. The Tennessee Valley Authority's Clinch River site continues to work through licensing for BWRX-300. In parallel, Canadian SMR deployments (BWRX-300 at Darlington) are proceeding faster than any US site and will produce operational data by 2028.

The Institute view

SMR is not a near-term grid-impact story. The earliest realistic US commercial operation date for a new SMR is 2029–2030, and material fleet deployment is a 2030s question. What is new in the 2024–2026 window is the emergence of hyperscaler behind-the-meter offtake, which changes the answer to the question "who pays for the first SMR." That is a genuine unlock and it is why the hyperscaler capital-expenditure story matters for SMR economics. Practitioners underwriting the SMR investment thesis should treat pre-2030 revenue as effectively zero and value the equities on demonstrated design certification progress, secured hyperscaler offtake, and manufacturing capacity buildout.

Cross-links

- Section 4 (Nuclear — Utility-Scale) above — the near-term nuclear story that is real today
- Deployment Modes section — hyperscaler BTM offtake is the specific unlock for SMR economics
- *The Deployment Plan* Move 4 — the transformer bottleneck applies to SMRs as much as to solar

6. Natural Gas — Combined-Cycle

The cost anchor

Combined-cycle gas turbines are the workhorse of American firm dispatchable generation and the most economically dispatched thermal resource in most of the grid. Lazard's LCOE 17.0 puts unsubsidized combined-cycle gas at \$45–\$108 per MWh, with a mid-point around \$75/MWh. That is above utility-scale solar and wind on a levelized basis but firm — a solar project can only produce when the sun shines, whereas a combined-cycle plant runs on demand at a specified capacity factor. The distinction is why combined-cycle gas remains the marginal price-setter in most competitive power markets even as the levelized cost of solar has fallen below it. Behind-the-meter deployments at hyperscaler data-center sites are a specific and rapidly growing subset of combined-cycle deployment, discussed below.

How a project actually gets built

A utility-scale combined-cycle project moves through: site control, air-quality permit (a specific and often binding step in EPA nonattainment areas), gas supply pipeline arrangement, water rights (combined-cycle plants use water for steam cycle and cooling), grid interconnection or behind-the-meter siting, turbine and heat-recovery-steam-generator procurement (with multi-year lead times as of 2026), construction, and commissioning. Total timeline is typically three to five years from site control to commercial operation. Behind-the-meter data-center deployments compress this timeline further by eliminating the grid interconnection block, but the turbine procurement block remains the binding constraint. A GE Vernova heavy-duty gas turbine ordered in 2026 has a delivery date measured in years, not months.

The deployment constraint

The binding constraint on new combined-cycle deployment is now turbine manufacturing capacity, not gas supply, permitting, or economics. GE Vernova's gas turbine order backlog extended to 2028–2029 as of mid-2026, and lead times for large-frame gas turbines have doubled since 2022. Siemens Energy reports similar backlog dynamics. Mitsubishi Heavy Industries and Wärtsilä provide additional capacity but face similar constraints. The turbine supply chain includes specific components — advanced cast blades, HRSG heat exchangers, generators — where the bottleneck is specific fabricator capacity, not just aggregate manufacturing.

Gas supply is a secondary constraint. US natural gas production remains abundant (the shale complex added roughly 15 Bcf/d of production 2020–2025 despite regulatory headwinds), but pipeline access to specific demand pockets — particularly in the Northeast, where interstate pipeline capacity has been politically contested for a decade — is a real constraint for some sites.

Operators and tickers

Table 6.1 — Combined-cycle operators, turbine suppliers, and gas supply chain

Table 7.

TICKER	COMPANY	EXPOSURE
GEV	GE Vernova	Largest US-based gas turbine OEM (post-spin from GE); heavy-duty and aeroderivative
ENR.DE	Siemens Energy	Global gas turbine OEM; substantial US installed base
MHI	Mitsubishi Heavy Industries	Gas turbine OEM; strong presence in Southeast US
VST	Vistra	Largest US independent power producer with combined-cycle fleet
NRG	NRG Energy	Combined-cycle fleet and retail electricity
CEG	Constellation Energy	Combined-cycle plus the nuclear fleet
EQT	EQT Corporation	Appalachian natural gas producer
LNG	Cheniere Energy	LNG export capacity — affects domestic gas price
KMI	Kinder Morgan	Interstate gas pipeline systems

TICKER	COMPANY	EXPOSURE
WMB	Williams Companies	Interstate gas pipeline systems
BKR	Baker Hughes	Turbine services and industrial gas

2024-26 activity

The 2024–2026 window has been a record period for US gas turbine orders. GE Vernova reported that gas power orders in 2025 exceeded the sum of 2022 and 2023 combined. Vistra secured multi-year gas turbine slots to support its Texas and PJM fleet expansion. Hyperscaler behind-the-meter deployments at data-center sites have driven a specific spike in aeroderivative turbine orders — smaller units, faster to deploy, well-suited to sites without heavy heat-recovery steam cycles. The 2024 rescission of certain EPA power sector regulations reduced regulatory friction for new gas builds in several states. Interstate pipeline capacity additions remain slow, with Mountain Valley Pipeline as the most notable multi-year delay finally reaching partial commercial service in 2024.

The Institute view

Combined-cycle gas is the workhorse of both the front-of-meter grid and behind-the-meter data-center power islands. Where solar and wind have won on levelized cost, combined-cycle gas has won on firm dispatchability and on turbine supply chain lead times short enough to meet 2026-vintage load growth timelines. The binding constraint is not fuel and not economics — it is the specific serial-manufacturing capacity for large-frame gas turbines, which is held by three global suppliers. Practitioners underwriting a combined-cycle-heavy thesis should treat the GE Vernova, Siemens Energy, and Mitsubishi Heavy backlog data as the primary leading indicator. The Institute view is that gas turbines will remain the marginal firm-capacity addition through the end of the decade, and that the specific turbine supply chain is what a serious reader should be watching on the industrial side of the story.

Cross-links

- GE Vernova WACC page — the cost of capital analysis for the largest US gas turbine OEM
- Deployment Modes section — hyperscaler BTM data-center power islands are combined-cycle's fastest-growing subsegment
- Section 7 (Natural Gas Mobile) below — the emergency-and-remote adjacent story

7. Natural Gas + Diesel — Mobile Turbines

The cost anchor

Mobile gas and diesel turbines do not compete on levelized cost of energy. Their price per megawatt-hour delivered is materially higher than any grid-connected utility asset — typically $200 to $400 per MWh, depending on fuel supply and operating hours. Practitioners who evaluate mobile turbines against a Lazard LCOE chart will always conclude that mobile deployment is the most expensive way to buy a kilowatt-hour. That conclusion is correct on a levelized basis and irrelevant to the actual deployment decision. Mobile turbines are not purchased for their per-MWh cost. They are purchased for their speed of deployment and their independence from the grid.

The rest of the section walks the specific economic role mobile turbines play — as bridge power for hyperscaler behind-the-meter data-center buildouts, as emergency response for disaster and outage situations, and as remote-site generation where no grid interconnection exists at all.

How a project actually gets deployed

Mobile turbine deployment is measured in weeks, not years. A single-unit deployment — an aeroderivative gas turbine on a skid, roughly 30 to 60 MW, with an on-site fuel system — moves from contract to commercial operation in thirty to sixty days. The mechanic is that the turbine is trucked or barged to the site, connected to a fuel supply (natural gas pipeline, LNG regasification, or diesel storage), and interconnected to the customer's load. Because the deployment is behind the meter by construction, no utility interconnection agreement is required. Larger deployments — multi-unit installations at 100+ MW scale — take longer but still compress the timeline dramatically compared to any grid-connected alternative.

Fuel logistics is where the operator does most of the work. A gas turbine running at full load consumes roughly 250 MMBtu per hour per 30 MW; the fuel supply arrangement (pipeline connection, LNG delivery contract, diesel storage refill) is often the binding constraint on how long the turbine can run continuously. This is why the mobile turbine industry sells "power plus fuel" as a combined service rather than turbines alone.

The deployment constraint

The binding constraint on mobile turbine deployment is not equipment supply or fuel; it is regulatory. Mobile turbines burn fossil fuels on-site, which places them under EPA Clean Air Act stationary-source rules once they operate for more than temporary durations. The specific test — whether a mobile turbine deployed for months at a data-center site qualifies as a "temporary" installation or as a "stationary source" requiring full air permitting — is being litigated in real time. Earthjustice and other environmental advocacy groups have filed Clean Air Act challenges against xAI's Colossus supercomputer deployment in Memphis, alleging that on-site methane turbines have exceeded the temporary threshold. The outcome of that litigation and its progeny will shape how much of the hyperscaler BTM buildout can be legally powered by mobile turbines in 2027 and beyond.

The secondary constraint is fuel. Mobile turbines that run on diesel are constrained by fuel logistics and by long-run cost. Mobile turbines that run on natural gas need pipeline access — which brings back a version of the same permitting problem that affects Section 6 combined-cycle plants, just at smaller scale.

Operators and tickers

Table 7.1 — Mobile turbine operators and adjacent equities (representative)

Table 8.

TICKER	COMPANY	EXPOSURE
—	APR Energy (private, Musk-owned as of May 2026)	1.1+ GW mobile fleet; largest independent mobile power operator in the US; explicit AI data-center BTM focus
—	Aggreko (private)	Global mobile-power operator; historic leader in emergency and industrial deployment
—	VoltaGrid (private)	On-site natural gas generation for oil & gas, data centers
CAT	Caterpillar	Diesel generator manufacturer and rental operator
CMI	Cummins	Diesel and gas gensets; hydrogen fuel cells
WRTBF	Wärtsilä	Reciprocating engine manufacturer for peaker and mobile applications
RYCEY	Rolls-Royce	Aeroderivative turbine manufacturer; Trent MT30 and industrial turbines

TICKER	COMPANY	EXPOSURE
BE	Bloom Energy	Solid oxide fuel cells for BTM (adjacent technology, non-turbine)

2024-26 activity

The most consequential 2024–2026 transaction in this segment is Elon Musk's personal acquisition of APR Energy, confirmed by a Federal Trade Commission filing in May 2026 and reported at an estimated $1 billion. That transaction consolidated the largest independent US mobile-turbine fleet — more than 1.1 gigawatts of nameplate capacity — under an owner whose stated data-center power needs (xAI's Colossus expansion, xAI's next-generation training cluster, Tesla's manufacturing footprint) point directly at behind-the-meter deployment. In parallel, VoltaGrid expanded its natural gas mobile fleet to serve both oilfield and hyperscaler customers. Aggreko continued to serve global emergency and industrial deployment. Regulatory friction increased over the same period: Earthjustice's Clean Air Act lawsuit against xAI's Colossus is the highest-profile case, and analogous challenges are moving forward at other data-center sites where mobile turbines have been deployed for extended durations.

The Institute view

Mobile turbines are a specialty tool with a specific and important role in the American energy stack: they are the fastest way to put megawatts on the ground when the customer cannot wait for a grid connection. The role is not going away — if anything, it is expanding as hyperscaler load growth outstrips utility interconnection queue throughput. The regulatory question — how much mobile deployment will Clean Air Act enforcement allow before the "temporary" characterization no longer applies — is the specific analytical block that separates a serious practitioner reading from a superficial one. The APR acquisition tells the reader that at least one very well-informed operator has concluded the tool is worth $1 billion at current fleet scale. That is a signal, not a conclusion, but it is a signal.

Cross-links

- Deployment Modes section — mobile turbines are BTM by construction
- Section 6 (Natural Gas Combined-Cycle) above — the utility-scale peer with different economics
- Section 11 (Efficiency and Demand-Side) below — the load-side alternative that reduces the need for emergency generation

8. Coal

The cost anchor

Coal-fired power in the United States is a legacy fleet in managed decline. No new coal-fired power plant has been permitted for construction in the United States since 2015, and no new-build LCOE is meaningfully relevant to any current investment decision. Lazard's LCOE 17.0 puts existing coal generation at \$69–\$168 per MWh unsubsidized — the wide range reflecting the fact that a fully depreciated plant with paid-off capital runs on marginal fuel cost, while a newer plant carrying capital costs runs at a much higher levelized figure. The economically dispatched portion of the coal fleet is shrinking, and coal's share of US electricity generation has fallen from roughly 50% in 2000 to about 16% in 2025.

Metallurgical coal is a different market entirely. Met coal is not burned for electricity; it is used as a reducing agent in blast-furnace steelmaking, where it is chemically necessary and not readily substitutable at scale. Global met coal demand has been stable or rising through the 2020s, and US met coal exports to Asia (Japan, South Korea, India) and Europe have been a bright spot for producers even as thermal coal power retires.

How the fleet actually retires

The US thermal coal fleet was 172 GW as of May 2025 and is planned to fall to approximately 145 GW by the end of 2028. Coal retirements happen through utility resource planning processes, regulatory proceedings, and — increasingly — through announcements that specific plants will retire or convert to gas by specific dates. Some plants have been converted to natural gas rather than retired (see Centralia in Washington State, discussed in *The Deployment Plan* Move 2), and the transmission interconnection at those sites is often more valuable than the plant itself. Met coal mines operate on multi-decade lives with capital budgets tied to specific reserves; retirement dynamics do not apply in the same way.

The deployment constraint

For thermal coal, the constraint is political and regulatory. New coal-fired construction is effectively impossible in the United States given current air permitting standards and financial market conditions. For met coal, the constraint is different: reserve depletion and permitting for mine expansion, particularly in Central Appalachia where the highest-grade metallurgical coal is produced.

Operators and tickers

Table 8.1 — Coal operators and adjacent equities

Table 9.

TICKER	COMPANY	EXPOSURE
BTU	Peabody Energy	Largest US coal producer; thermal and met coal
HCC	Warrior Met Coal	Pure-play US metallurgical coal exporter
AMR	Alpha Metallurgical Resources	Central Appalachian met coal producer
CEIX	Consol Energy	Northern Appalachian thermal coal producer
ARCH	Arch Resources	Diversified coal producer with met exposure
CLF	Cleveland-Cliffs	Integrated steelmaker that consumes US met coal; disclosed position — see paper front matter

The Cleveland-Cliffs reference is highlighted because CLF is not primarily a coal producer, but its steel operations are a major domestic buyer of US metallurgical coal, and the paper's author holds a position in CLF as disclosed at the front. Readers evaluating met coal producers as a category should consider that domestic met coal demand runs directly through the US steel industry, and that Cleveland-Cliffs' operating decisions materially affect that demand.

2024-26 activity

Thermal coal retirements continued at a moderate pace through 2024 and 2025, with some announced retirements pushed later due to concerns about resource adequacy in specific markets (PJM, MISO). PJM's 2024 and 2025 capacity auctions cleared at materially higher prices in part because coal retirements are running ahead of replacement firm-generation additions. On the met coal side, US exports remained a bright spot; Warrior Met Coal and Alpha Metallurgical continued to serve Asian and European steelmakers. Cleveland-Cliffs' domestic operations continued to consume met coal at scale.

The Institute view

Coal is not coming back for US electricity generation. Coal is not going to zero either. The right practitioner read is that thermal coal generation retires on a specific schedule, that the retirement pace has slowed modestly due to load growth and capacity concerns, and that the specific interconnection assets at retiring plants are valuable for surplus interconnection deployment of solar, storage, and gas (see *The Deployment Plan* Move 2). Metallurgical coal is a separate market that runs on global steel demand, and it will continue to have a domestic role as long as blast-furnace steelmaking exists in the United States. The two markets should be evaluated separately, and the equities on each side price different fundamentals.

Cross-links

- *The Deployment Plan* Move 2 — surplus interconnection at retiring coal plants
- Cleveland-Cliffs case study — the integrated steelmaker consuming domestic met coal
- Section 6 (Natural Gas Combined-Cycle) — the replacement generation for retiring coal

---

9. Hydro

The cost anchor

Existing hydroelectric generation is the cheapest source of electricity on the American grid on a levelized basis. Fully depreciated hydro facilities with paid-off capital, operating on fuel that arrives free from the sky, deliver electricity at $30–$60 per MWh in most cases — competitive with any new-build alternative. The catch is that this reference cost is entirely a legacy statistic. New large-scale hydroelectric construction is effectively impossible in the United States: the sites are permitted, the environmental review requirements are prohibitive, and the political economy has shifted toward dam removal rather than new dam construction. Pumped-storage hydro is a different story — new pumped-storage projects have been permitted and are proceeding, but on 8-to-12-year timelines and at costs materially higher than the legacy hydro fleet.

How a project actually gets built

New conventional hydro projects are not being built. Existing hydro projects are being relicensed through the Federal Energy Regulatory Commission on 30-to-50-year license cycles, with each relicensing subject to environmental review, fish passage requirements, and often material capital investment. Pumped storage projects — where excess grid electricity is used to pump water uphill for later release through generation turbines — are being permitted, financed, and constructed on a slower cadence. A pumped-storage project moves through: site identification (typically an existing upper reservoir with lower-elevation grid access), FERC preliminary permit, feasibility studies, environmental review, license application, financing, and construction. Total timeline 8 to 12 years is typical.

The deployment constraint

The constraint on hydro is siting. The best sites for large hydroelectric generation were developed by 1970, and remaining opportunities face insurmountable environmental and tribal consultation obstacles. Pumped storage has more opportunity because it uses existing water and reservoirs rather than blocking wild rivers, but even pumped storage projects face lengthy environmental review, fish passage requirements, and cost allocation questions in interstate markets.

Operators and tickers

Table 9.1 — Hydro operators (representative)

Table 10.

TICKER	COMPANY	EXPOSURE
BEP, BEPC	Brookfield Renewable	Largest US hydro operator; over 30 GW of installed capacity worldwide
DUK	Duke Energy	Regional utility with substantial hydroelectric portfolio
SO	Southern Company	Alabama Power hydro assets
PCG	PG&E	Extensive California hydro operations
IDA	IDACORP	Idaho Power — hydro-heavy generation portfolio
—	Bureau of Reclamation (federal)	Federal hydroelectric operator (WAPA power marketing)
—	Tennessee Valley Authority (federal)	Federal utility with extensive hydro

2024-26 activity

The Klamath River dam removal, completed in 2024, is the largest dam removal in US history and marks a specific inflection in the political economy of American hydro. Simultaneously, several pumped storage projects have moved through FERC preliminary permit into feasibility and permit stages, including projects in the Pacific Northwest and Appalachian regions. FERC has continued to work through licensing renewals for the existing hydro fleet, with environmental groups pushing for stricter fish-passage and minimum-flow requirements at several projects.

The Institute view

Hydro is not a new-build growth story. It is an under-appreciated existing asset that provides low-cost, firm, dispatchable, low-carbon generation to specific parts of the American grid, particularly in the Pacific Northwest and Southeast. Pumped storage is the specific segment where new capacity is possible and where the deployment mode fits behind-the-meter or grid-scale storage applications for solar and wind firming. Practitioners should evaluate hydro utilities primarily on their operating fleet quality and their relicensing risk, and evaluate pumped storage developers on specific project pipelines and financing arrangements.

A related category the paper does not fully cover: next-generation geothermal

A category adjacent to hydro and to the broader firm-baseload discussion is next-generation geothermal — enhanced geothermal systems that use hydraulic techniques adapted from shale oil-and-gas to access heat resources at commercial scale. Fervo Energy has signed a long-term geothermal offtake deal with Google, and the technology is progressing from demonstration to early commercial. The Institute does not treat next-generation geothermal as a full eleventh (or twelfth) tool in this paper because the technology is not yet at commercial deployment scale, but readers should be aware that the category exists and is likely to appear in future editions of this paper.

Cross-links

- Section 3 (Storage) — pumped storage as an alternative to lithium-ion for long-duration applications
- Berkshire Hathaway Energy — PacifiCorp hydro operations
- *The Deployment Plan* — no direct hydro reference but the transmission constraint applies

10. Grid Transmission

The cost anchor

Grid transmission is not a generating asset; it is the precondition for every generating asset in this paper. Transmission cost per mile varies from roughly $1 million per mile for lower-voltage AC lines to $10 million or more per mile for high-voltage direct current (HVDC) lines. The Department of Energy's National Transmission Needs Study (2023) estimated that meeting American electricity demand growth through 2035 will require approximately 47,000 miles of new high-voltage transmission — roughly double the current pace of construction. Without that transmission, the utility-scale solar, wind, storage, and nuclear discussed in Sections 1–5 cannot reach the loads they are meant to serve.

How a project actually gets built

A new transmission line moves through: route selection, environmental review under NEPA (multi-year), state siting and permitting proceedings (multiple states for interstate lines), tribal consultation where applicable, FERC ratemaking for cost recovery, financing, procurement, and construction. Total timeline for a new interstate HVDC line is 8 to 12 years, and can be longer when routing crosses multiple state jurisdictions with divergent permitting frameworks. Reconductoring existing lines — replacing conventional aluminum-conductor cables with advanced composite-core conductors that carry more current at higher temperatures — is materially faster and is a specific opportunity discussed at length in *The Deployment Plan* Move 3.

The deployment constraint

The constraints on transmission are siting and cost allocation. Every new transmission line crosses land owned by multiple parties in multiple political jurisdictions. State public utility commissions have historical authority over intrastate siting and often oppose interstate lines that provide primarily out-of-state benefits. Cost allocation — who pays for the line, and how the cost is recovered from ratepayers — is a specific FERC ratemaking question that has been in flux under FERC Order 2023 and successor orders. Regional Transmission Organizations conduct planning processes for interregional transmission, but historical throughput of RTO transmission planning has been materially lower than demand.

Operators and tickers

Table 10.1 — Transmission operators, developers, and construction contractors

Table 11.

TICKER	COMPANY	EXPOSURE
PWR	Quanta Services	Largest US electric infrastructure contractor; transmission construction leader
MTZ	MasTec	Infrastructure construction including transmission
MYRG	MYR Group	Transmission and distribution construction
PRIM	Primoris Services	Utility infrastructure construction
ETR	Entergy	Regulated utility with transmission assets
DUK	Duke Energy	Regulated utility with transmission
SO	Southern Company	Southern Transmission subsidiary
GEV	GE Vernova	HVDC converter station and grid technology
—	Grid United (private)	Interstate HVDC developer (Southern Spirit line)

2024-26 activity

FERC Order 2023, issued in July 2023, was the first overhaul of interconnection procedures since 2003 and is still being operationalized at the RTO level in 2026. The Department of Energy has released the National Transmission Needs Study identifying specific corridors requiring new investment. Multiple HVDC interstate lines have proceeded through permitting: Grid United's Southern Spirit line between Texas and the Southeast, TransWest Express in Wyoming, and multiple projects in MISO's tranche upgrades. Reconductoring pilots at existing transmission lines have demonstrated that advanced composite-core conductors can materially increase throughput on existing rights-of-way — a specific opportunity elaborated in *The Deployment Plan* Move 3.

The Institute view

Grid transmission is the specific binding constraint on the deployment of every other generating asset in this paper. Without more transmission, the interconnection queue does not clear, and every LCOE advantage that solar, wind, storage, and utility-scale nuclear enjoy remains hypothetical. The Institute view is that grid transmission is the highest-leverage single investment area in American energy in the second half of the 2020s and into the 2030s, and that the specific mix of new construction, reconductoring, and FERC ratemaking reform will determine the actual pace at which the rest of the stack deploys. *The Deployment Plan* Moves 1 and 3 elaborate the arguments for connect-and-manage as a near-term throughput multiplier and reconductoring as an opportunity to double the carrying capacity of existing lines within a decade.

Cross-links

- *The Deployment Plan* Move 1 — connect and manage under FERC Order 845
- *The Deployment Plan* Move 3 — reconductoring existing transmission
- Every generating-asset section above — transmission is the precondition

11. Efficiency and Demand-Side

The cost anchor

The cheapest kilowatt-hour is the one not consumed. Building efficiency, industrial process efficiency, and demand-side management — where loads shift or reduce in response to grid conditions or price — deliver "negawatts" at costs that are typically negative when properly accounted for. The American Council for an Energy-Efficient Economy estimates that utility efficiency programs deliver energy savings at $20–$40 per MWh, less than half the cost of new generation of any type. Demand-response programs, in which loads reduce during peak grid hours, deliver capacity at costs materially below the equivalent new-build peaker plant. Virtual power plants — aggregations of distributed batteries, EV chargers, smart thermostats, and other flexible loads — are beginning to bid into wholesale capacity markets on the same terms as conventional generation.

How a project actually gets deployed

Efficiency and demand-side deployment is fragmented by construction. A utility efficiency program moves through utility resource planning, public utility commission approval, program design, contractor procurement, customer marketing, installation, and measurement and verification. A commercial building energy retrofit moves through owner decision, engineering, financing, contractor selection, construction, and commissioning. An industrial process efficiency project moves through operations engineering, capital budgeting, procurement, installation, and commissioning. Timelines vary widely — 12 to 36 months per project is a reasonable central estimate.

Virtual power plant deployment is faster: enrollment in an aggregator program, contract signing, dispatch integration. Some VPP aggregators can enroll a residential customer within days.

The deployment constraint

The constraint on efficiency and demand-side is decision-maker fragmentation. Every efficiency project has a specific decision-maker (a building owner, a factory operator, a homeowner, a fleet manager) who benefits from the project on a specific timeline. The fragmentation makes efficiency deployment slow at aggregate scale even when individual project economics are excellent. Utility programs solve this at scale for residential and small commercial, but industrial and large commercial retrofit deployment remains stubbornly slower than the underlying economics would predict.

The secondary constraint is measurement. Because efficiency delivers negawatts rather than positive generation, its impact on the grid is measured against a counterfactual — what would consumption have been without the intervention. Measurement and verification methodologies are complex and are a specific source of professional friction between program administrators and program evaluators.

Operators and tickers

Table 11.1 — Efficiency and demand-side operators (representative)

Table 12.

TICKER	COMPANY	EXPOSURE
TT	Trane Technologies	HVAC systems and building efficiency
CARR	Carrier	HVAC systems and building automation
JCI	Johnson Controls	Building automation and efficiency
HON	Honeywell	Building controls and industrial automation
SU.PA	Schneider Electric	Electrical distribution and building energy management
STEM	Stem Inc.	AI-optimized behind-the-meter storage and demand management
ENPH	Enphase Energy	Residential storage and virtual power plant enrollment

TICKER	COMPANY	EXPOSURE
RUN	Sunrun	Residential VPP participant
—	Voltus (private)	Commercial and industrial demand-response aggregator
—	AutoGrid (private, part of Schneider)	VPP software platform

2024-26 activity

The 2024–2026 window has been notable for virtual power plant scale-up. Multiple state public utility commissions approved VPP programs at meaningful capacity — California, New England, New York, and several others. The IRA's §179D commercial building energy-efficiency deduction was expanded, and PACE (Property Assessed Clean Energy) financing scaled at both commercial and residential levels. Utility efficiency spending continued at modest growth rates. Demand-response participation in wholesale capacity markets grew as ISOs updated their tariffs to conform to FERC Order 2222, which requires wholesale markets to allow distributed energy resource aggregations to bid on the same terms as conventional generation.

The Institute view

Efficiency and demand-side management is under-invested relative to what its economics justify. Part of the reason is that there is no single-equity way to own the theme — the value is spread across HVAC OEMs, building automation companies, VPP aggregators, and utility programs, none of which trade purely on efficiency exposure. The theme is spread across HVAC OEMs, building automation companies, VPP aggregators, and utility programs, none of which trade purely on efficiency exposure. Practitioners looking for exposure should assemble a basket rather than picking a single name. The demand-response segment is the most concentrated near-term opportunity because Order 2222 implementation is creating a new revenue category (wholesale capacity payments to DER aggregations) that did not exist five years ago. The efficiency theme is the specific area where the American energy stack is materially under-invested relative to project-level economics — and it is under-invested because the ownership and decision-making structures do not match the technology.

Cross-links

- Section 3 (Storage) — residential storage as a VPP participant

- Deployment Modes section — small-scale BTM overlap with distributed solar and storage
- Every generating-asset section — a kilowatt-hour saved is a kilowatt-hour that does not need to be generated

The Practitioner's Read

If the reader takes only three things away from this paper, they should be these.

One: no single tool solves the American energy problem. Every serious utility, every hyperscaler, every state grid operator uses a combination of the tools described here. The paper's editorial spine is that the reader who thinks in terms of solar-only, or nuclear-only, or gas-only, is reading the ideological discourse, not the actual grid. The grid uses everything, in different mixes, for different jobs, on different timelines. Serious analysis follows the grid, not the discourse.

Two: the binding constraints are procedural, not technical. Solar is cheap enough. Nuclear can be restarted. Storage is deployable. Efficiency has excellent economics. What stops the American energy stack from deploying faster is interconnection queue reform, transmission construction, state siting authority, permitting delay, and workforce. These are all political-economy problems, not physics problems. That is where the Institute's *Deployment Plan* focuses.

Three: where the meter sits changes everything. The single most consequential change in American energy deployment in the 2020s has been the migration of large loads — data centers, industrial customers, remote operators — to behind-the-meter deployment. Behind-the-meter deployment bypasses the eight-year interconnection queue. The operators who understand this — Elon Musk with APR Energy, Meta with its behind-the-meter targets, Amazon and Google with hybrid nuclear PPAs — are moving faster than the ones who are still waiting for grid-connected utility service. Where the reader looks for opportunity in energy over the next five years, they should look at the deployment mode question first.

The Institute's contribution here is the framework, not a stock tip. What the reader does with the framework is their own decision. If the paper has helped a serious outside reader understand how the American energy stack actually works, it has done what it set out to do.

Where a reader could reasonably disagree with this paper: the "no single tool" claim understates how much marginal capacity gas and solar together account for in current deployment. The behind-the-meter growth narrative could plateau if Clean Air Act enforcement tightens materially. And SMR economics could clear earlier than the Institute expects if the hyperscaler offtake structure scales more aggressively than anticipated. The Institute welcomes reader engagement on all three points.

Ticker Appendix

Every publicly-traded equity mentioned in this paper, consolidated for practitioner reference. This is not a recommended portfolio; it is a reference list of the operators the paper walks through, so the reader who wants to research further has a starting point.

Utilities and independent power producers - NEE — NextEra Energy — largest US utility-scale solar and wind operator (via NEER); regulated Florida utility - DUK — Duke Energy — regulated utility across the Carolinas; nuclear, gas, renewables - SO — Southern Company — Georgia Power parent; Vogtle 3 & 4 nuclear - XEL — Xcel Energy — Minnesota nuclear, Midwest wind - AEP — American Electric Power — Ohio Valley transmission and generation - ETR — Entergy — Southeast nuclear fleet - PEG — Public Service Enterprise Group — Salem, Hope Creek nuclear - PCG — PG&E — Diablo Canyon nuclear; California hydro - AGR — Avangrid — Iberdrola US utility - IDA — IDACORP — Idaho Power, hydro-heavy - VST — Vistra — largest US IPP with gas and nuclear - NRG — NRG Energy — gas fleet and retail electricity - CEG — Constellation Energy — largest US nuclear operator; Microsoft TMI PPA

Renewables, storage, and residential - FSLR — First Solar — US-based thin-film module manufacturer - RUN — Sunrun — residential solar installer - ENPH — Enphase Energy — microinverters and residential storage - CSLR — Complete Solaria — acquirer of SunPower's residential assets - BEP, BEPC — Brookfield Renewable — global hydro and renewables - FLNC — Fluence Energy — utility-scale storage integrator - EOSE — Eos Energy — zinc-hybrid stationary battery - STEM — Stem Inc. — behind-the-meter AI-optimized storage

Steel, coal, uranium - CLF — Cleveland-Cliffs — sole domestic producer of grain-oriented electrical steel; met-coal buyer; *Institute disclosure: paper's author holds a position* - BTU — Peabody Energy — largest US coal producer - HCC — Warrior Met Coal — US metallurgical coal exporter - AMR — Alpha Metallurgical Resources — Central Appalachian met coal - CEIX — Consol Energy — Northern Appalachian thermal coal - ARCH — Arch Resources — diversified coal with met exposure - CCJ — Cameco — uranium mining and fuel processing

Turbines, engines, and grid technology - GEV — GE Vernova — largest US-based gas and wind turbine OEM; grid technology - CAT — Caterpillar — diesel generator manufacturer and rental - CMI — Cummins — diesel and gas gensets; hydrogen fuel cells - BE — Bloom Energy — solid-oxide fuel cells - ETN — Eaton — power management and grid equipment - SMR — NuScale Power — light-water SMR - OKLO — Oklo Inc. — sodium-cooled fast reactor - BWXT — BWX Technologies — naval reactor and defense nuclear

Natural gas, LNG, and pipelines - EQT — EQT Corporation — Appalachian natural gas producer - LNG — Cheniere Energy — LNG export capacity - KMI — Kinder Morgan — interstate gas pipelines - WMB — Williams Companies — interstate gas pipelines - BKR — Baker Hughes — turbine services and industrial gas

Efficiency and demand-side - TT — Trane Technologies — HVAC and building efficiency - CARR — Carrier — HVAC and building automation - JCI — Johnson Controls — building automation and efficiency - HON — Honeywell — building controls and industrial automation - TSLA — Tesla — Tesla Energy Megapack storage; VPP program

Sports/services adjacent - PWR — Quanta Services — US electric infrastructure contractor - MTZ — MasTec — utility infrastructure construction - MYRG — MYR Group — transmission and distribution construction - PRIM — Primoris Services — utility infrastructure construction - TPIC — TPI Composites — wind turbine blades

Foreign-listed operators mentioned for reference - ENR.DE — Siemens Energy — global gas turbine OEM - MHI — Mitsubishi Heavy Industries — gas turbine OEM - WRT-BF — Wärtsilä — reciprocating engines - RYCEY — Rolls-Royce — aeroderivative turbines - DNNGY — Ørsted (ADR) — Danish offshore wind developer - SONY — Sony (not in this paper; parked here for future editions) - SU.PA — Schneider Electric — electrical distribution and building energy management

Berkshire Hathaway Energy exposure - BRK.A / BRK.B — Berkshire Hathaway (BHE is held privately through Berkshire; the equity path for reader exposure is Berkshire Hathaway itself)

Not a recommendation. Every ticker above is provided so a serious reader can locate the filed disclosures and form their own view.

Glossary

ADVANCE Act (2024) — federal legislation providing the Nuclear Regulatory Commission with new authority and appropriations to streamline advanced-reactor licensing.

AC (Alternating Current) transmission — the standard transmission mode across most of the American grid. Cost per mile is materially lower than DC but efficiency drops over long distances.

BESS (Battery Energy Storage System) — a grid-connected or behind-the-meter installation of batteries, typically lithium-ion, used to store and dispatch electricity.

Behind the meter (BTM) — generation sited on the customer's side of the utility meter, delivering power directly to load without transiting the utility grid.

Bcf/d (Billion cubic feet per day) — a unit for measuring natural gas production and pipeline flow.

Capacity market — a wholesale market mechanism in which generators are paid for the ability to be available to produce electricity, separate from the energy they actually produce.

Combined-cycle gas turbine — a natural-gas plant that pairs a gas turbine with a steam turbine driven by the gas turbine's exhaust heat, delivering higher efficiency than a simple-cycle plant.

COD (Commercial Operation Date) — the date at which a generating asset begins delivering power under contract.

DC (Direct Current) / HVDC transmission — a transmission mode used for long-distance and interregional lines. More expensive per mile but more efficient over distance.

DER (Distributed Energy Resource) — a small generating or storage asset connected to the distribution grid rather than the transmission grid.

EIA (Energy Information Administration) — the statistical agency of the US Department of Energy.

ERCOT — the Electric Reliability Council of Texas, the grid operator for most of Texas.

Ellison PIPE — a private-investment-in-public-equity commitment made in this paper's Institute case-study context by the Ellison Trust and syndicated to sovereign wealth partners.

FERC (Federal Energy Regulatory Commission) — the federal agency with jurisdiction over interstate transmission, wholesale power markets, and interstate gas pipelines.

FERC Order 845 (2018) — reform of the standard interconnection agreement requiring provisional interconnection service and other rights.

FERC Order 2222 (2020) — reform requiring wholesale markets to allow DER aggregations to bid as resources.

FERC Order 2023 (2023) — reform of interconnection procedures to a cluster-study model with firm deadlines.

GW (Gigawatt) — one billion watts; a typical utility-scale power plant is measured in gigawatts.

HALEU (High-Assay Low-Enriched Uranium) — uranium enriched above 5 percent U-235 but below 20 percent, required by many advanced-reactor designs.

HVAC (Heating, Ventilation, and Air Conditioning) — the systems that condition indoor air; a major share of building energy consumption.

IFM (In Front of the Meter) — generation connected to the utility grid rather than to a specific customer's load.

IRA (Inflation Reduction Act) (2022) — federal legislation providing extensive tax credits and other support for energy investment.

IRP (Integrated Resource Planning) — the state-regulated process by which utilities file multi-year plans for generation, transmission, and demand-side investments.

ISO (Independent System Operator) — an entity that operates a regional electricity market and grid; ERCOT is an example.

ITC (Investment Tax Credit) — a federal tax credit for a percentage of an investment's cost, applied to solar, storage, and other technologies.

kW (Kilowatt) — one thousand watts; a common unit for residential and commercial-scale systems.

Lazard LCOE / LCOS — the annual Levelized Cost of Energy (or Storage) analyses published by Lazard, an industry benchmark for cost comparisons.

LCOE (Levelized Cost of Energy) — the per-megawatt-hour cost of generation over a plant's lifetime, adjusted for capital cost, operating cost, and expected output.

LNG (Liquefied Natural Gas) — natural gas cooled to a liquid state for transport by tanker.

MW (Megawatt) — one million watts; a common unit for utility-scale generation.

MWh (Megawatt-hour) — the energy produced by a one-megawatt generator running for one hour.

NOL (Net Operating Loss) — a tax attribute allowing a company to offset future income with past losses.

NRC (Nuclear Regulatory Commission) — the federal agency regulating US commercial nuclear reactors.

Peaker plant — a generating asset (often gas or diesel) designed to run during peak-demand periods rather than continuously.

PILOT (Payment In Lieu Of Taxes) — a negotiated arrangement between a project and a host county to provide payments in place of standard property taxes.

PJM Interconnection — the regional transmission organization covering Mid-Atlantic and parts of the Midwest.

Power island — a behind-the-meter site with no grid tie for the primary energy flow.

PPA (Power Purchase Agreement) — a long-term contract under which a customer agrees to buy power from a generator at specified terms.

PTC (Production Tax Credit) — a federal tax credit paid per unit of energy produced, applied to wind, nuclear, and other technologies.

§45U — the IRA production tax credit for existing nuclear generation.

§45X — the IRA advanced-manufacturing production credit for domestic clean-energy component production.

§48E — the IRA standalone-storage investment tax credit.

§179D — the commercial-building energy-efficiency tax deduction.

RTO (Regional Transmission Organization) — a FERC-approved entity that operates a regional grid and market; PJM is the largest.

SMR (Small Modular Reactor) — a nuclear reactor designed to be manufactured in modules at a central facility and assembled on site, typically 50-300 MW per unit.

Ticking fee / Ticking consideration — a mechanism in a merger agreement providing additional payment to sellers if closing is delayed past a specified date.

Transformer, generator step-up (GSU) — the high-voltage transformer that steps generator output up to transmission voltage; a specific and constrained piece of grid hardware.

Utility-scale — a generating asset sized for wholesale sale rather than for a specific customer's load; typically 5 MW and above.

VPP (Virtual Power Plant) — an aggregation of distributed energy resources managed as a single grid-facing resource.

§197 intangible amortization — the section of the Internal Revenue Code governing the tax amortization of acquired intangible assets over 15 years.

§338(h)(10) election — a joint tax election in an acquisition treating a stock sale as an asset sale for tax purposes.

§382 (NOL limitation) — the section of the Internal Revenue Code limiting the use of net operating losses after an ownership change.

Index

Institute Cross-References

This paper sits at the intersection of multiple Baratelli Institute frameworks and case studies. Readers who want to go deeper on any single tool can follow the trails below.

Table 13.

SECTION	RELATED INSTITUTE WORK
Section 1 (Solar)	*The Deployment Plan* — full companion working paper on solar deployment mechanics
Section 1, 8 (Solar, Coal)	Cleveland-Cliffs case study — the electrical-steel and met-coal operator
Section 2 (Wind)	Berkshire Hathaway Energy — the operating case study on utility-scale wind
Section 3 (Storage)	*The Deployment Plan* Move 2 — surplus interconnection at retiring thermal plants
Section 4 (Nuclear)	Constellation Energy — largest US nuclear operator
Section 5 (SMR)	Deployment Modes section — hyperscaler BTM as the SMR unlock