Closing the Generation Gap: Renewables and the Data Center Boom

As utilities struggle to meet electricity demand from power-hungry data centers, energy transition projects have the most to gain.

Edward Levin
Co-head, Direct Infrastructure

Thomas Emmons
Co-head, Direct Infrastructure

VOYA® INVESTMENT MANAGEMENT

Executive summary

The rapid growth of large cloud service providers, also known as hyperscalers, and the tech industry's shift to power-hungry Graphic Processing Units (GPUs) are spurring a seismic shift in the way America consumes electricity. Thanks in large part to data centers, U.S. power demand is accelerating for the first time since 2007—and the grid is struggling to keep up.

This sudden and critical need for new generation capacity is causing skyrocketing demand for wind and solar projects, which are fast to build and have high margins (the fuel is free). Alongside them, a whole new energy infrastructure sector has emerged: utility-scale battery storage.

The capital needed over the next 5+ years to deliver this capacity will be massive—senior and subordinated debt, tax equity, sponsor equity—providing ample opportunities for nimble lenders who are experienced in the relevant structures.

The demand side

- Over the past 10 years, U.S. electricity peak demand has been flat. The next five years are forecast to see growth of 4.7%—almost double previous forecasts of 2.6%.[1]
- The U.S. has 16.8 GW of currently operational data centers and another 24.8 GW in development. Put in perspective, the entire state of Georgia's peak electricity demand in 2023 was 17 GW.
- The strongest growth is centered around data center "hotspot" areas such as Northern Virginia, Dallas, Silicon Valley and Phoenix, but secondary markets are also growing as developers seek access to power.
- Tech has never been able to resist increases in processing power, and GPUs use 2–10x the power of CPUs—making new data centers much more power hungry than old ones.
- Utility forecasts for data center demand remain conservative, suggesting significant upside to demand forecasts if even 25% of the current pipeline comes to fruition.
- While data center demand is one of the biggest growth drivers, it is not the only one: EVs and the electrification of office, manufacturing and industrial facilities also contribute.

The supply side

- Wind and solar are the most popular new generation projects (by a wide scale) because of their proven technology, rapid construction and strong economics.
- Utility-scale battery storage is a new but increasingly important component of solar projects, which raises both their price tag and their value to the grid.
- Regulatory incentives favor energy transition projects over new fossil fuel generation. Attractive project economics, strong state-level support, and corporate preference for clean energy all help mitigate political risk.
- Significant fossil fuel generation capacity is scheduled for decommissioning, which could further tighten supply.
- Due to their intermittent nature, 4–8 MW of solar or wind capacity must be built (plus storage) for every 1 MW of new peak demand—and, for every 1 MW of fossil fuel capacity retired, 2–4 MW of new renewable capacity is needed.
- The demand rise is sharp enough that any renaissance in baseload fossil fuel generation would not harm the energy transition outlook.

[1] FERC form 714 grid planner forecast data from 2022 versus 2023, as quoted in GridStrategies, The Era of Flat Power Demand Is Over, 12/01/23.

After decades of flat growth and underinvestment, the U.S. power market is struggling to meet data centers' surging electricity needs.

Electricity demand has been stagnant for over a decade… until now.

The five big hyperscalers are planning $350 billion in capex over the next two years.

Data centers: Bigger, faster, hungrier

For the past 14 years, U.S. electricity demand has been largely stagnant, with small fluctuations driven mainly by periods of extreme weather (Exhibit 1). Yet suddenly, everyone is hiking their power demand estimates—from Mid-Atlantic regional transmission organization PJM to usually staid bodies such as the North American Electric Reliability Corporation (NERC), the Energy Information Administration (EIA) and the Federal Energy Regulatory Commission (FERC).

The catalyst, they say, is data centers. Why do these server farms suddenly matter so much?

A brief history of hyperscaling

While data centers have existed almost as long as computers themselves, their current era of prominence began in 2006 with the launch of Amazon Web Services (AWS). As cloud computing accelerated from 2010 through 2018, the amount of computing done in data centers jumped 550%. Fortunately, data centers were getting way more efficient just as usage was exploding, leading to a mere 6% rise in related energy consumption during that time period (Exhibit 2).[2]

Demand was generally predictable, efficiency was on track—and then 2020 happened.

The growth in data center power demand post 2020 and its forecast continuation are broadly attributable to a few factors:

- The pandemic accelerated the rollout and adoption of cloud services generally—a simple example is the proliferation of streaming video platforms. Netflix runs entirely off AWS, where it is said to use more than 100,000 server instances.[3]

Exhibit 1: U.S. power demand—up, up and away

BkWh

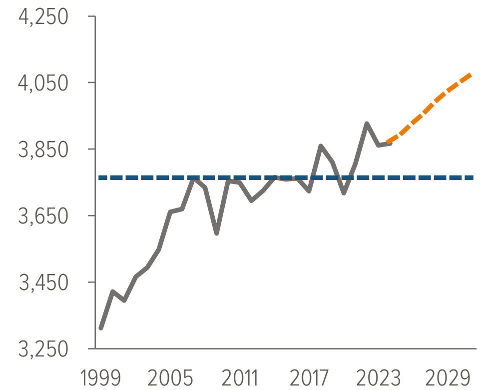

As of 07/20/24. Source: Energy Information Administration.

Exhibit 2: Data center power demand is outrunning efficiency improvements

U.S. data center power consumption (LHS, GW) vs. power usage effectiveness (RHS) (lower is better)

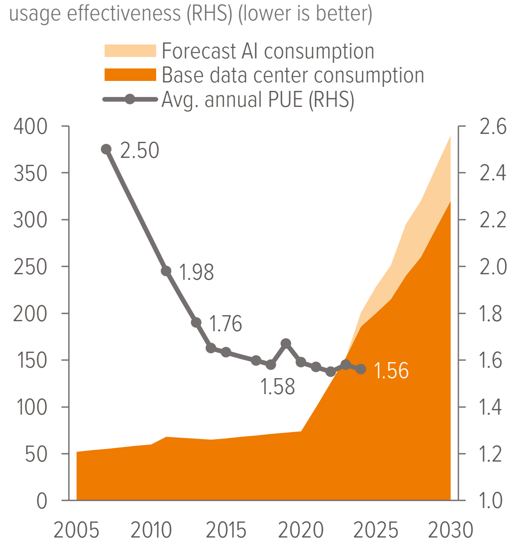

As of 07/31/24. Source: International Energy Agency, Uptime Institute, BCG, Electric Power Research Institute, Lawrence Berkeley National Laboratory, Voya IM estimates. Power usage effectiveness (PUE) is the ratio of a data center's total energy consumption to the energy used to run (and cool) IT equipment. The best possible PUE is 1, which signifies that all electricity consumed by the data center is used for computational needs.

[2] Masanet et al., "Recalibrating global data center energy-use estimates," *Science*, 02/28/20.
[3] AWS case studies: Netflix. Accessed 07/19/24.

- As a result, many more data centers are being built. AWS, Microsoft Azure, Google Cloud, Meta and Oracle are collectively planning infrastructure capital expenditure of $350 billion over the next two years to cope with expected growth.[4] There is also growth in on-premises and co-location data centers.

Data centers consume as much power as the state of Georgia.

- At the same time, each of those data centers is getting hungrier due to the step change in power requirements when servers move to graphics processing units (GPUs)—which consume ten times the power of central processing unit (CPU) clusters and require considerably more cooling (Exhibit 3).[5]

AI is certainly driving some of this data center growth. But whether you're an AI bull or bear, there's a load of confirming evidence from industry.

Because data centers are almost universally described by their power draw rather than their square footage or server count, their future need for electricity is well documented. That need is staggering—16.8 GW of capacity already exists, nearly equal to the peak summer demand of the entire state of Georgia (Exhibit 4). If the full 24.8 GW of data center capacity currently in the pipeline materializes, data centers' electricity appetite will start closing in on California's.

EVs and the electrification of manufacturing and industry are also driving demand growth.

The power companies are well aware of it, too: When a data center project is started, its developer calls up the local utility and says, "Hey, I'm going to need 100 MW of power pretty much 24/7, starting 18 months from now," and the utility says, "Okay." Once the utility gets enough of those calls, it starts raising its own demand figures—and discussing data centers' impact on analyst calls (Exhibit 5).

Exhibit 3: Watts up, doc—GPUs can require 10x the power of CPUs

Thermal design power of currently available processors (watts)

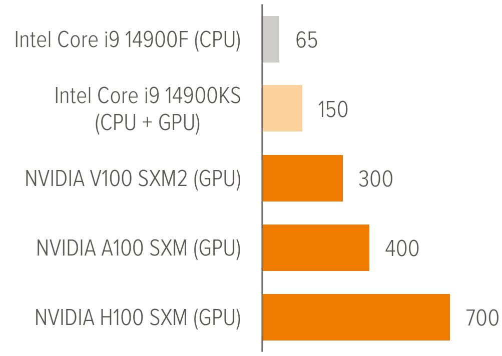

As of 07/20/24. Source: Intel, Nvidia. Fun fact: OpenAI used 25,000 NVIDIA A100 GPUs to train GPT-4, over 100 days.

Exhibit 4: Data centers already consume more power than many U.S. states

U.S. data center capacity vs. selected states' peak electricity demand (GW)

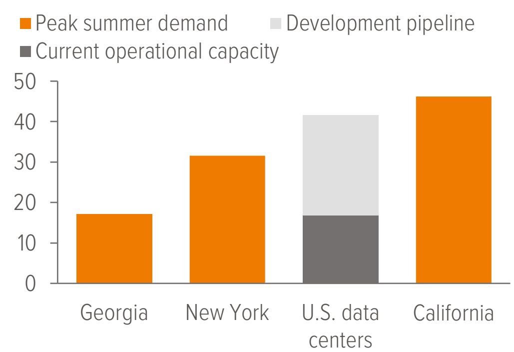

As of 07/20/24. Source: Cushman & Wakefield, Georgia Power, NYISO, CAISO.

Exhibit 5: Everybody's talking about it

Mentions of "data center" in utilities' earnings calls

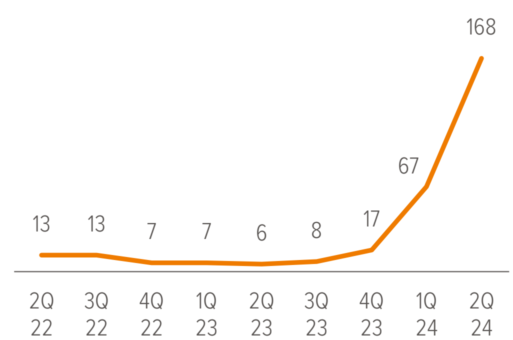

As of 07/31/24. Source: LPL Research, Bloomberg.

[4] This is slightly more than the annual GDP of Egypt. Capex forecast from Dgtl Infra, Cloud and Hyperscale Capex Expenditures in 2024, 05/30/24. Google, Microsoft, and AWS reported quarterly cloud revenue growth of 28%, 21%, and 19% YoY recently, with Amazon CEO Andre Jassy noting that with an estimated 90% of global IT spend still on-prem, there remains significant growth potential in cloud services going forward.

[5] Newmark, 2023 U.S. Data Center Market Overview & Market Clusters, 01/15/24.

A grid in growth mode

Data centers are by far the most discussed reason for power demand to move firmly into growth territory, but they are not the only factor. The San Francisco Bay Area's PG&E cites higher-than-forecast electric vehicle ownership as its largest load growth driver, followed by its 3.5 GW of data center pipeline.[6] Texas's ERCOT and the Southwest's SPP are dealing with the electrification of the Permian Basin, which by itself is forecast to increase grid load by 13 GW from 2022 to 2032.[7]

Much of the growth in commercial electricity demand over the past five years, especially around data centers, has been concentrated in about 10 states.[8] Yet that growth has been strong enough and continual enough that the EIA has raised its electricity demand forecasts nearly monthly in 2024 (Exhibit 6).

The speed of data center development has caught utilities by surprise.

At present, America's largest data center market is Northern Virginia—this is due to its proximity to national government and to a confluence of data cables from across the U.S. and overseas. The concentration of data centers across Loudoun, Fairfax, Prince William and Fauquier counties presents a useful case study for just how much the past year's data center demand upswing has caught even the experts off balance—and how it affects the broader regional grid.

The local utility covering those counties is Dominion Energy—a smaller utility that covers eastern Virginia and a couple counties in North Carolina. 24% of Dominion's demand is from the circa 300 data centers in its territory, which together had a peak capacity of more than 4 GW.[9] In short, Dominion should know how to cope with data centers.

Exhibit 6: U.S. commercial electricity demand forecasts keep being revised upwards

BkWh

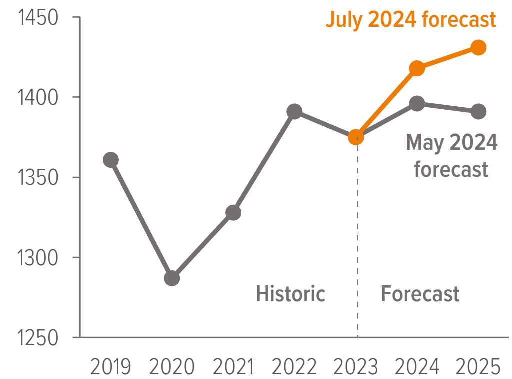

As of 07/20/24. Source: Energy Information Administration.

Exhibit 7: PJM keeps having to raise its forecasts because of data centers

Summer peak demand forecast (GW)

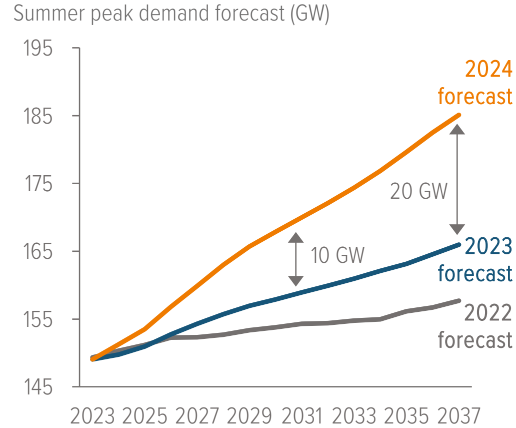

As of 01/08/24. Source: PJM long-term load forecasts, 2022-2024.

Yet the 1.24 GW of additional data center capacity under construction seems to have caught Dominion on the hop.[10] Between 2023 and 2024, it has more than doubled its summer load growth forecasts for the next 10 years—from 2.2% to 5.5%. Dominion's forecast 2028 demand figure alone was raised by 6 GW. That's more than twice the growth of any other part of the area's regional transmission organization (RTO).[11]

[6] PG&E, 2024 Investor Update (q4cdn.com), 06/12/24. PG&E's Jason Glickman admits later in the earnings call that they have only factored a couple hundred megawatts out of that 3.5 GW data center potential into their forecasts.
[7] S&P Global, prepared for the ERCOT planning committee, Electrifying the Permian Basin, March 22, 2023.
[8] The top data center markets in the U.S. at year-end 2023 were, according to Cushman & Wakefield: Northern Virginia (5,491 MW operational / 1540 MW under construction); Phoenix (1,941/235); Atlanta (1,664/405); Oregon (1440/542); Dallas (982/115); Silicon Valley (903/315); Chicago (844/432); Columbus (504/278). Vacancy rates are below 5% across the board.
[9] Dominion Energy, 2023 Annual Report (03/21/24) and 1Q earnings presentation (05/02/24).
[10] Newmark, op. cit. Plans have been filed for a further 3 GW of data centers in the area.
[11] PJM, 2024 Load Forecast Report, 02/01/24.
[12] Ibid.

In their own words: The grid's growing pains

Recent statements from public utilities, transmission organizations and government show the strain on today's grid.

North American Electric Reliability Corporation: "Electricity peak demand and energy growth forecasts over the 10-year assessment period are higher than at any point in the past decade. The aggregated assessment area summer peak demand forecast is expected to rise by over 79 GW, and aggregated winter peak demand forecasts are increasing by nearly 91 GW. Furthermore, the growth rates of forecasted peak demand and energy have risen sharply since the *2022 LTRA*, reversing a decades-long trend of falling or flat growth rates."[13]

Federal Energy Regulatory Commission: "Our country is facing an unprecedented surge in demand for affordable electricity while confronting extreme weather threats to the reliability of our grid and trying to stay one step ahead of the massive technological changes we are seeing in our society. Our nation needs a new foundation to get badly needed new transmission planned, paid for and built."[14]

Energy Information Administration: "After reviewing information and projections from utilities and grid operators in the areas of the country with rapid data center development, we revised our forecasts upward for commercial electricity demand through 2025... Nationally, we expect U.S. sales of electricity to the commercial sector will grow by 3% in 2024 and by 1% in 2025. Data center developments are evolving rapidly, and we plan to reevaluate our upcoming forecasts as we receive more information."[15]

Dominion Energy: "Historically, a single data center typically had a demand of 30 megawatts or greater. However, we're now receiving individual requests for demand of 60 megawatts to 90 megawatts or greater, and it hasn't stopped there. We get regular requests to support larger data center campuses that include multiple buildings and require total capacity ranging from 300 megawatts to as many as several gigawatts."[16]

DigitalBridge: "Power is really the constraining factor... We started talking about this over two years ago at the Berlin infrastructure conference, when I told the investor world, "We're running out of power in five years." Well, I was wrong about that. We're running out of power in the next 18 to 24 months."[17]

PJM: "The demand ... reflects growth from end-use electrification, electric vehicles and data centers. Recent history of this anticipated growth has proven unprecedented and dynamic. Average growth estimates for PJM's summer peak, for example, have increased by 375% between the 2022 and 2024 load forecasts—from 0.4% per year to 1.6% per year."[18]

PG&E: "We have demand... for 3.5 GW of new data centers. There are 19 customers that are currently in our pipeline in the Bay Area to bring these data centers online over the course of this next five-year time horizon ... What's currently in our plan is just a few hundred megawatts. If even a portion of this ends up coming to fruition, it is significant upside opportunity from a load growth perspective."[19]

ERCOT: "The plan that we are going to be using coming out of this study work has a load growth projection of 152 GW by 2030... That is about 40 GW higher than the same forecast a year ago, which is a tremendous growth ... We see 85% of the expressed interest for generation supply in Texas continuing to be made up of solar and battery resources."[20]

Georgia Power: "Because of the rapid pace of economic development in Georgia since January 2022, the company's current projections reflect load growth of 6,600 MW through the winter of 2030/2031, which is approximately 17 times greater than that previously forecasted ... The size of many of these projects far exceeds historical annual norms, with some individual projects surpassing 1,000 MW. In addition to the size of the large load presented by these new projects, many of the projects reflect a higher load factor and around-the-clock operations, which requires a substantial amount of generation and consistent energy delivery throughout the day and night, as opposed to only during specific times."[21]

In turn, this has caused Dominion's RTO, PJM, to say it has "observed unprecedented data center load growth" in multiple areas, with the potential for "all remnant capacity on the transmission system" to be utilized, resulting in regional reliability issues.[12]

Due to data centers, electric vehicles and the electrification of industrial sectors, PJM has doubled its 10-year annual growth projection for summer peak demand each year since 2022, from 0.4% to 0.8% to its current 1.6% (Exhibit 7).

Dominion isn't alone. PG&E, the utility that provides power to another data center hotspot—the San Francisco Bay Area—announced in June that its current data center pipeline will add 3.5 GW of load to the grid and contribute 0.5-1.5% of its forecast 2-4% CAGR load growth from 2023 to 2040. (The other major contributors are electric vehicles and building electrification.)[22]

We expect demand to reach more states as hyperscalers and developers seek to expand into secondary markets with more readily available real estate and power. Duke Energy (North and South Carolina) and Georgia Power, for example, have hiked forecasts significantly over the past year in response to data center "mega-projects."

Duke Energy: "The economic development pipeline over the period that we've shared through 2028, data centers represent about 25% of that pipeline. As we get out to 2030 and beyond, that 25% grows."[23]

In terms of national estimates, the results of FERC form 714 submissions show that grid planners expect overall U.S. peak power demand to rise by 38 GW between now and 2028. Between 2022 and 2023 submissions, cumulative growth forecasts for the next five years almost doubled, from 2.6% per year to 4.7% per year.[24]

Energy transition projects are forecast to see the largest growth, due to economics, speed of construction, customer preference and government policy.

The supply struggle

If electricity demand jumps faster than expected—which clearly seems to be happening in various hotspot regions—power companies can't simply pull more supply out of thin air. It takes around 2–3 years to build a new solar photovoltaic plant and 4–5 years for a combined cycle natural gas plant. Before shovels meet dirt, there can be a further 2–5 years of waiting in interconnection queues. Getting that power from where it's generated to where it's needed is also a time-consuming process involving raising capex and securing approvals to build transmission lines.

There are only three short-term responses to power demand rising faster than projected supply:[25]

1. **Pull more energy from peaker plants,** which are high-cost emergency generating plants that run on fuel oil or natural gas. The result is usually a mix of lower utility margins and angry customers with sky-high power bills.
2. **Delay plans to decommission coal and natural gas plants.** This assumes both that the specific utility has a lot of plants scheduled for decommission

[13] NERC, 2023 Long-Term Reliability Assessment, 12/15/23, page 33.
[14] FERC Summer Assessment 2024 presentation at the Commission's open meeting by chairman Willie Phillips, 05/23/24.
[15] Energy Information Administration, Today in Energy, 06/28/24.
[16] Bob Blue, CEO of Dominion Energy, during the utility's 1Q24 earnings call, 05/02/24.
[17] DigitalBridge CEO Marc Ganzi, during the company's 1Q24 earnings call, 04/30/24. DigitalBridge has 5 GW of data centers in its pipeline, with over 2 GW in construction.
[18] PJM, Energy Transition in PJM, 06/24/24, p.4. PJM is a regional transmission organization (RTO). Its initials stand for Pennsylvania-New Jersey-Maryland Interconnection, and its area includes those states as well as Ohio, Virginia, West Virginia, Delaware, large parts of Kentucky, and bits of North Carolina, Illinois (Chicagoland), Indiana, and Michigan. Crucially, PJM (via Dominion Energy, which falls within its area of oversight) delivers electricity to the biggest concentration of data centers in the U.S., which are spread across four counties in Northern Virginia. Dominion's current peak load is 22.8 GW (15% of PJM's total peak demand), estimated to rise to 42.8 GW (23% of total) by 2037.
[19] PG&E EVP Jason Glickman, during the utility's 06/12/24 investor update webcast. PG&E covers the San Francisco Bay Area. Its RTO is CAISO (California Independent System Operator).
[20] ERCOT CEO Pablo Vegas, during the RTO's board of directors meeting, 04/23/24. ERCOT is the Electric Reliability Council of Texas, the RTO for most of the state of Texas.
[21] Georgia Power, 2023 Integrated Resource Plan Update, October 2023.
[22] PG&E, 2024 Investor Update (q4cdn.com), 06/12/24.
[23] CEO Lynn Good, Duke Energy quarterly earnings call, 08/06/24. The famously fossil fuel-centric Duke has recently signed MOUs with Amazon, Google, Microsoft, and Nucor to develop structures that will allow them to invest in clean energy technologies for customers with large-scale energy needs.
[24] FERC form 714 data, as quoted in GridStrategies, The Era of Flat Power Demand Is Over, 12/01/23.
[25] Our colleagues on the PCIG Energy & Infrastructure desk covered this beautifully in a recent quarterly: Power Generation's Wake-Up Call.

Data centers draw nearly their full power capacity 24/7/365.

Solar and wind are set to grow from 15% of U.S. generation to 43% by 2033.

and plant maintenance hasn't been abandoned in the run-up to planned decommission.

3. **Let the grid fall over.** Nobody likes this. We don't recommend it.

The supply struggle is further complicated by data centers' unique power demand profile. Normal commercial and residential power demand rises and falls throughout the day, with the highest-use periods typically in the morning and early evening (with some seasonal and weather fluctuations) before tapering off at night as people go to sleep.

Data centers never sleep. They're generally expected to have "seven nines" uptime—99.99999% working. In high-utilization areas such as the Northern Virginia market, this translates to data center loads almost never dipping below 90% of peak capacity.[26]

Even in data center markets where utilization is lower, a data center will have a high baseline power need, likely between 65% and 80%, because servers idle at about half their peak power draw and cooling systems still need to run.

Data centers also experience power demand swings, at unpredictable times and without warning, up towards that 100% capacity mark.[27] Utilities can't anticipate whether a data center will suddenly require all of its capacity at 3 a.m. on a Tuesday morning or at 6 p.m. on Friday night in the middle of a heatwave. Thus, the utility has to plan its capacity around the data center drawing its full capacity 24/7/365, even if that's not the data center's actual load profile.

In sum: America needs more power, and fast. And how are grids scrambling to meet that growth? Overwhelmingly, the answer is with renewables (Exhibit 8).

Why renewables?

Solar and wind—currently the dominant types of renewable generation—compose 15% of net electricity generation in the U.S. They are forecast to grow to 43% of net generation by 2033 (Exhibit 9). And it's already happening: 58% of planned U.S. grid additions in 2024 are solar, followed by battery storage (23%) and wind (13%).[28]

Exhibit 8: 2020 vs. 2023 EIA estimates for renewables capacity

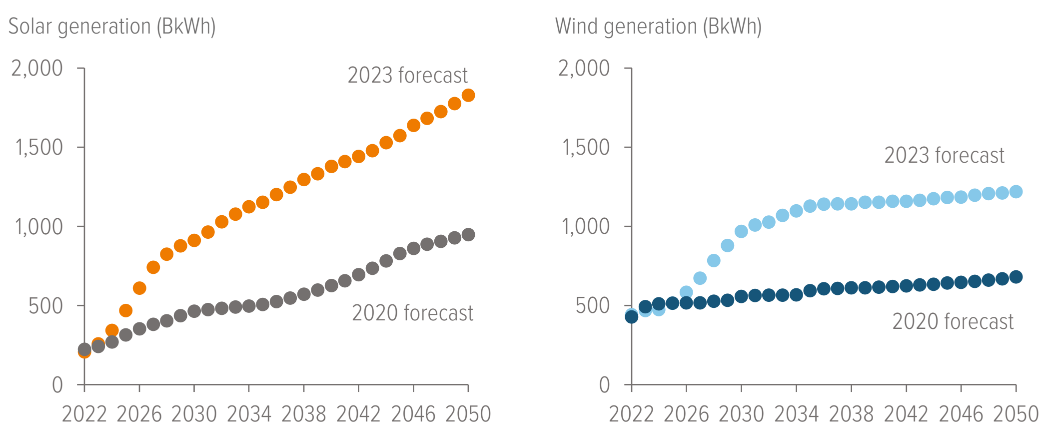

As of 07/20/24. Source: Energy Information Administration.

[26] Google provides an illustrative chart of a data center's baseline load here, which shows its rather extreme intraday load swings.
[27] Energy Information Administration, "Solar and battery storage to make up 81% of new U.S. electric-generating capacity in 2024," 02/15/24.
[28] PJM, 2024 Load Forecast Supplement, p. 27, 01/15/24.

Reasons for the growth of renewable generation aren't complicated: Wind and sunshine are free, new plants are fast to build, consumers prefer them, and current regulation (both federal and state) is supportive.

The heart of renewables' appeal is basic economics.

Exhibit 9: Solar and wind are poised to triple their share of U.S. generation over the next decade

Net electricity generation by fuel (BkWh)

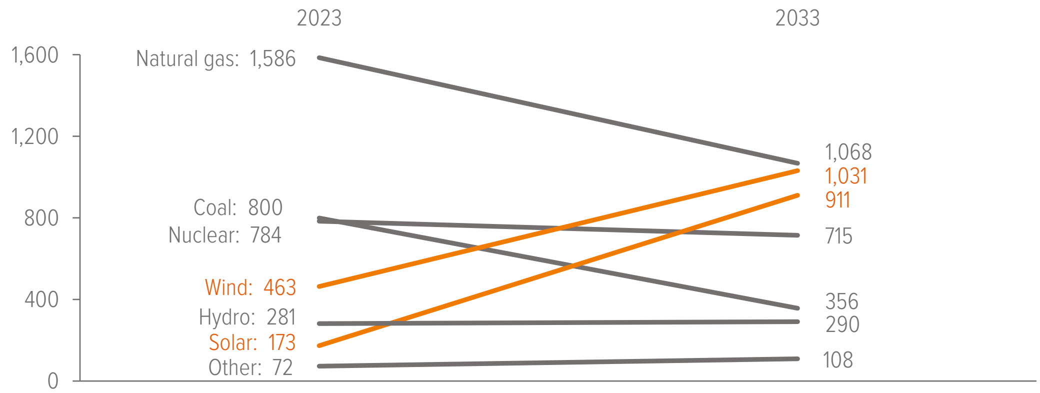

As of 07/20/24. Source: Energy Information Administration.

Solar and wind PPA prices have significantly outpaced inflation from 1Q20-1Q24:

U.S. inflation: +17%

Average PPA prices: +107%

As of 07/20/24. Source: Federal Reserve, LevelTen Energy, Voya IM estimates.

Renewable fuel is intangible and has no incremental cost

U.S. electricity is still overwhelmingly supplied by combined cycle natural gas plants (Exhibit 9). Generating plants buy natural gas and then sell electricity at a small spread above that natural gas price.

Wind and solar generating facilities get fuel for free and then sell electricity via power purchase agreements (PPAs), with fixed initial prices per megawatt-hour plus an annual fixed escalation. Over the past four years, those PPA prices have increased at about 6x the rate of inflation (see sidebar).

The heart of renewables' appeal is basic economics: Sponsors like the attractive margins of renewable generation and the nice, long PPAs with investment grade offtakers, which create highly predictable cash flows.

There aren't a lot of risks for this particular driver. Even if oil and gas prices were to fall, it wouldn't change the preference for renewables, since sponsors know fuel prices go up as well as down in the time it takes to build a plant.

Renewable generation is fast to build

Solar projects are based on proven, reliable technology and take only 25 months (on average) from interconnection agreements to commercial operation.[29] That's roughly two years faster than the average combined cycle natural gas plant.[30]

If you're a utility who needs power today, you may be more inclined to go for the generation asset that will be ready sooner—even if you're based in a state famous for fossil fuel production, such as Texas. In fact, Texas has the largest total wind and solar generation levels in the U.S., at 140 TWh in 2023—nearly double California's 77 TWh.[31]

The proof is in the interconnection queues: At the end of 2023, solar, storage and wind made up 95% of active capacity in U.S. interconnection queues, with natural gas coming in at only 3% (79 GW).[32] Picking on ERCOT again as our avatar of fossil fuel positivity, its 2024-2030 interconnection queue is: 15 GW of natural gas generation, 35 GW of wind, 155 GW of solar and 141 GW of battery storage.[33]

[29] Lawrence Berkeley National Laboratory, "Queued Up: 2024 Edition. Characteristics of power plants seeking transmission interconnection as of the end of 2023," April 2024.
[30] Bolson et al., "Capacity factors for electrical power generation from renewable and nonrenewable sources," PNAS, 12/20/22.
[31] ERCOT, Interval Generation by Fuel Report, 03/07/24; CAISO website, accessed 08/06/24.
[32] Lawrence Berkeley National Laboratory, op. cit. Adorably, Texas' solar boom has also created significant new opportunities for sheep ranchers.
[33] ERCOT board of directors meeting, 04/23/24.

It's not just Texas. In the past year or so, solar has overwhelmingly become the generation project of choice (Exhibit 10).

95% of the U.S. inter-connection queue is wind, solar and storage.

Exhibit 10: Both solar capacity and project size are rising rapidly

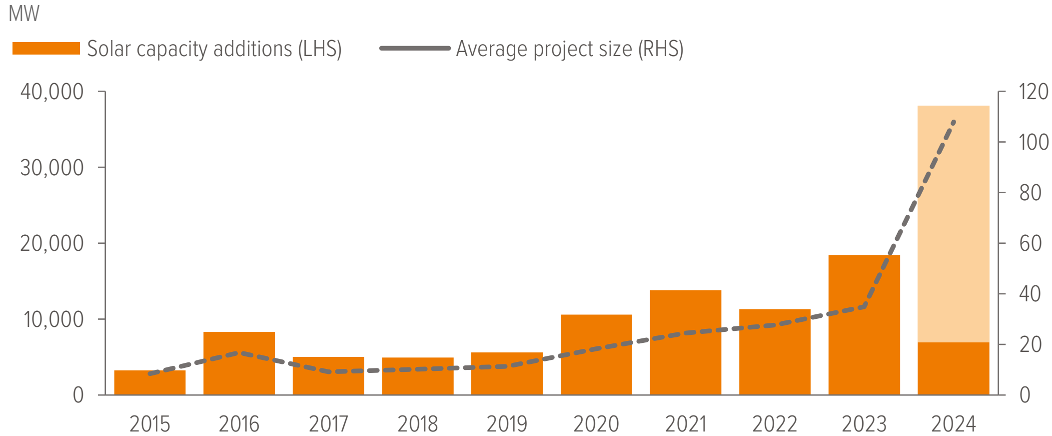

As of 07/20/24. Source: Energy Information Administration.

The risks around this are several—but unlikely. If FERC walks back its recent reforms meant to speed interconnection permits, that will hurt all new generating projects equally, not just renewables.

The U.S. has a healthy domestic solar panel and wind turbine manufacturing industry, mitigating potential effects from a hike to tariffs on imported Chinese panels and turbine parts. Tariffs may cause shortages in supply as U.S. facilities ramp up production, but that would likely be a temporary frustration rather than a permanent setback.

Most data center operators and customers have aggressive 100% renewable energy commitments.

The consumers driving demand growth prefer renewables

Tech companies are, by and large, committed to green power—especially the ones that make extensive use of data centers. Microsoft Azure, Google, AWS, Oracle, Apple and Meta (plus other data center operators, such as Iron Mountain, Digital Realty Trust, QTS Realty Trust, and Switch) all have commitments to use 100% renewable energy within the next 10 years.[34]

DigitalBridge: "A big piece of the power puzzle centers around renewables. This is an area of intense interest from our portfolio of company customers. Again, it's a customer-driven opportunity and solution—who all have aggressive net-zero targets for their compute and connectivity footprints, and from our institutional LPs as well. They want to see green electrons increasingly power their data center investments."[35]

In many cases, these companies are investing directly in the construction of renewable generation assets.

The risk here is that hyperscalers decide to stop worrying and learn to love fossil fuels—or go all in on nuclear. Their promises of a near-term green future are already being tested: The sharp rise in computing power needed to innovate in AI at scale has already caused Microsoft and Google to miss their renewables and emissions targets.[36]

[34] S&P Global, "Data center companies continue renewable buying spree, surpassing 40 GW in US," 03/28/23. Google is legitimately trying to go carbon free by powering all its operations directly by renewable power, and the company has written about it quite extensively. Its 2018 paper, "Moving Toward 24x7 Carbon-Free Energy at Google Data Centers: Progress and Insights," helps illuminate many of the practices of and issues faced by large data center operators in going fully renewable. Iron Mountain and Microsoft have also made 24/7 clean power pledges.
[35] DigitalBridge CEO Marc Ganzi, speaking during the company's 1Q24 earnings call, 04/30/24.
[36] Computerweekly.com, "Microsoft and Google's GHG emissions gains call viability of net-zero targets into question," 07/09/24.

Both companies admitted in their 2023 annual reports that their greenhouse gas emissions had actually risen year on year. However, the companies also reiterated their goals for net-zero emissions in 2030 as they announced the missed 2023 targets, so an abandonment of renewable goals seems unlikely.

The breadth and momentum of pro-renewable government policy makes it unlikely to be affected by changes in the Oval Office.

Government policies favor renewable generation

Recent policy actions have strongly favored renewables. The Inflation Reduction Act of 2022 (IRA) increased tax credits for investment in renewables, the Bipartisan Infrastructure Law of 2021 provided over $10 billion in funding for transmission and storage, and a suite of new EPA emissions restrictions essentially rendered new baseload natural gas and coal power plants uneconomic.

Because the IRA is law, a new president couldn't simply cancel it with an executive order. Congress would have to vote on it, which is a time-consuming and contentious process. And there is broad-based industry support for the IRA, which is likely to keep Congress from turning against it.

EPA emissions restrictions could be at risk, as they're likely to be heavily litigated no matter who is in the Oval Office in 2025. However, given the state of the grid and the usefulness of balancing the intermittency of renewable resources with combined cycle generation, any delay or walk-back of these regulations is not likely to affect the energy transition growth curve.

It's worth noting, too, that much of the regulation forcing retirement of fossil fuel generating plants is happening at the state level, and state laws aren't affected by changes in federal administration. The Illinois Climate and Equitable Jobs Act and the New Jersey CO_2 rule, for example, are both causing decommissioning of major coal and natural gas generating assets in these states (and are making it difficult to build new ones).

Likewise, significant state regulation exists to promote clean energy production. All 50 states have some sort of renewable energy policy and incentives—and 24 states, plus Puerto Rico and the District of Columbia, have goals for 100% clean electricity by 2050 or earlier.[37]

Exhibit 11: Capacity factor varies widely among fuels

U.S. electric generation capacity factors, 2023

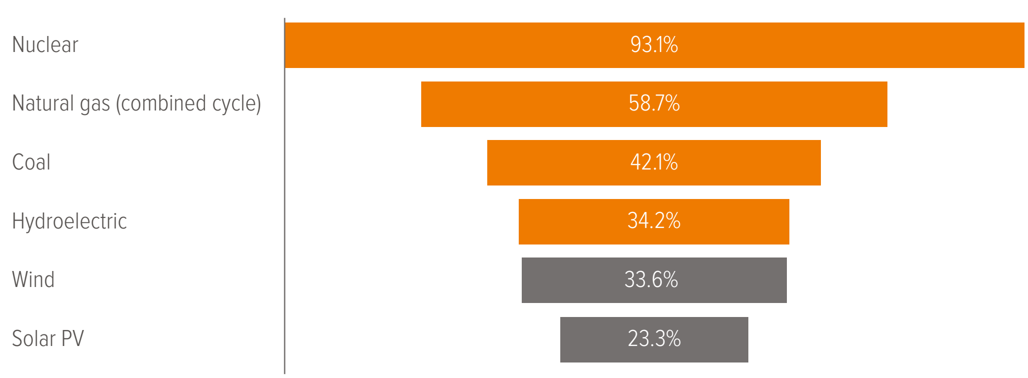

As of 07/20/24. Source: Energy Information Administration.

[37] Clean Energy States Alliance, Table of 100% Clean Energy States, accessed 07/26/24.

For every megawatt of new peak demand, multiple megawatts of wind, solar and storage capacity are needed to meet it.

Solar's intermittency will require increasing mitigation to maintain a stable grid.

Building a stable clean energy grid is an asymmetric process

The relationship between electricity demand and supply is not linear. To cover 100 MW of new demand, you need more than 100 MW of new supply, regardless of the fuel type (yes, even nuclear).[38] There are three reasons for this: capacity factor, intermittency and resource adequacy.

Capacity factor measures a generating asset's actual production versus its nameplate capacity, and it varies widely among fuel and plant types (Exhibit 11). For renewables, it goes hand in hand with intermittency.

For every 1 GW of fossil fuel capacity retired, 2–4 GW of renewables are needed to replace it.

Intermittency measures the natural peaks and valleys in supply for a generating asset (Exhibit 12). Solar has generally predictable intermittency (from zero at night to near-peak production at high noon). Wind can blow day and night, and its intermittency is much less predictable than the sun. Nuclear has almost no intermittency; its generation profile is pretty much a straight, horizontal line.

Before about 2021, the way grids handled renewables' intermittency was to use combined cycle natural gas and/or coal as "old faithful"—providing a baseline level of power, which could be ramped up in low-wind periods or ramped down during the sunniest part of the day, when solar was churning out its nameplate capacity.

As more solar was built in the past four years, however, its generation profile started to become an issue for grids—especially California's. So much solar was coming in around midday that natural gas plants were being turned down to the point that their operational economics plummeted.

There were also times when transmission lines couldn't handle renewables' peak output. (Imagine a day that was both really sunny and really windy.) Grids such as California's started having to deliberately curtail their renewable generation to prevent oversupply and congestion and keep natural gas plants functioning at economic capacity levels.

Exhibit 12: Battery storage is critical to smoothing out solar intermittency

July's higher solar generation is balanced by more battery storage, smoothing out January's sharp midday drop in natural gas generation

CAISO generation profile: January 15, 2024

CAISO generation profile: July 15, 2024

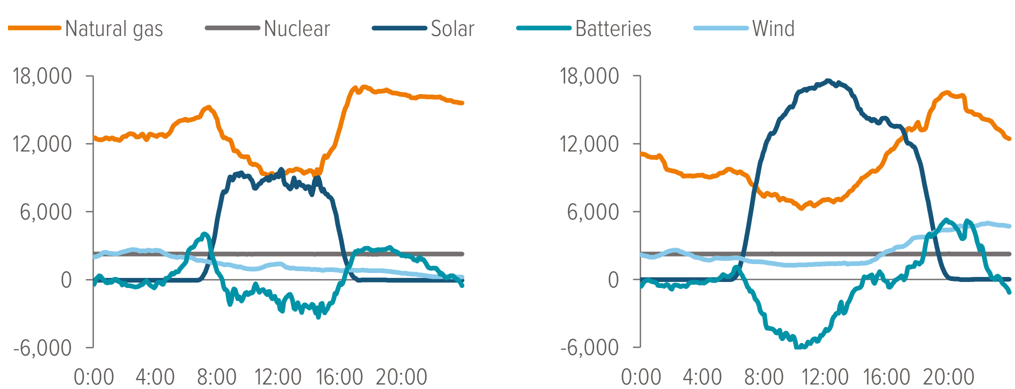

As of 07/20/24. Source: CAISO.

[38] A high capacity generating asset like nuclear requires significant backup generation capacity, to provide resiliency in case the nuclear plant goes offline unexpectedly. Otherwise, customers may face sudden restrictions in power supply.

The struggles of California's transmission authority, the California Independent System Operator (CAISO), have not yet been fully resolved, but the addition of significant battery storage—charged during the day with excess solar and discharged at peak morning and evening hours—has helped smooth out the precipitous dips in midday natural gas generation (Exhibit 12).[39]

Resource adequacy measures how well the grid can cope with sudden spikes in demand or unexpected shortfalls in supply. Grid stability is important to utilities and their RTOs, and most of them manage their supply/demand balance according to a risk metric known as loss-of-load expectation (LOLE). LOLE is a measure of how often, on average, the grid's capacity will fall short of demand (i.e., require rationing).

Due to renewables' lower capacity factors, LOLE has major implications for future requirements. In short, the more fossil fuel generation capacity you take away, the greater the multiple of renewable capacity you need to replace it. Mid-Atlantic RTO giant PJM provides an instructive example of this in a recent report.[40] PJM operates under an LOLE of one day in 10 years—in other words, approximately 2½ hours per year is the most the grid's demand should exceed supply.

PJM's current assumption is that 41 GW of its natural gas and coal capacity will be retired between 2022 and 2030. While this is broadly in line with its 2012-2022 average of 4.3 GW in annual retirements, it still represents 21% of PJM's entire installed capacity—and 25 GW are retirements driven by environmental policy (a mix of IL, NJ and federal/EPA regulations).[41]

In order to maintain resource adequacy, according to its LOLE projections, PJM expects to need 84 GW of new wind, solar, storage and hybrid (solar plus storage) capacity to replace the 41 GW of "dirty" power it plans to retire—about a 2:1 ratio (Exhibit 13).

Exhibit 13: Renewable generation capacity cannot be substituted 1:1 for fossil fuel capacity

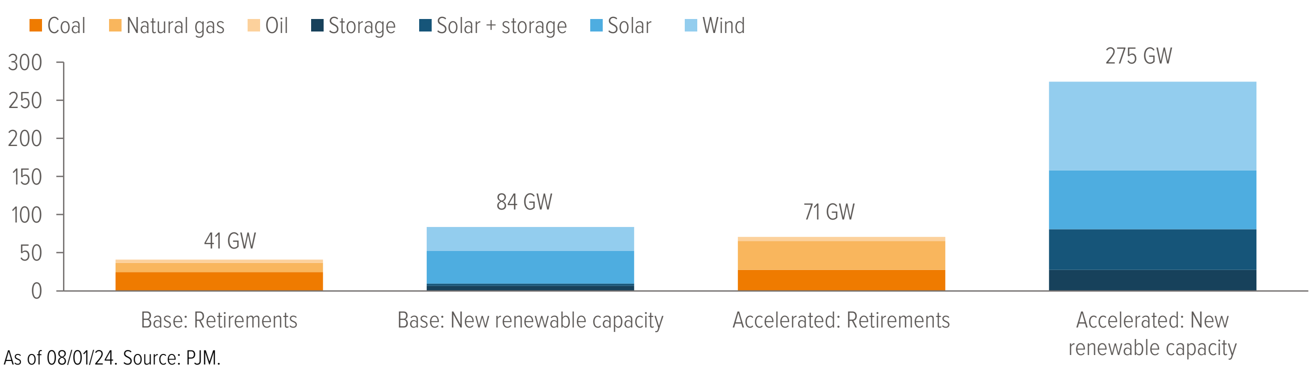

As of 08/01/24. Source: PJM.

Exhibit 14: The birth of a new infrastructure sector: solar storage

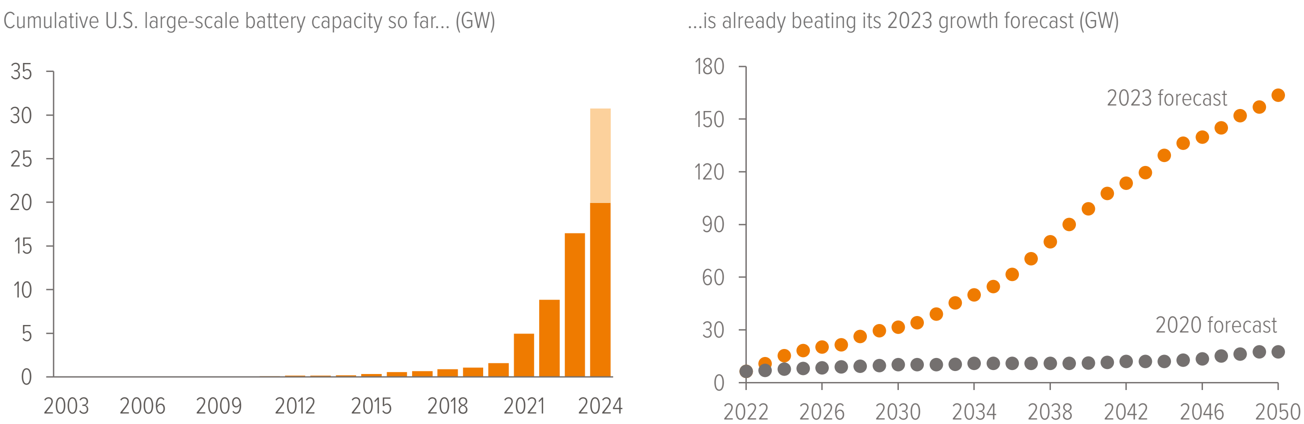

As of 08/01/24. Source: Energy Information Administration.

[39] California is rapidly adding solar capacity, such as the Edwards & Sanborn Solar + Energy Storage project in the Mojave Desert, which came online in January and has 875 MW of solar generating capacity plus 3,287 MWh of battery storage. (The batteries can discharge at the solar plant's full 875 MW capacity for about four hours straight, or at lower capacities for longer.)

[40] PJM, "Energy Transition in PJM: Flexibility for the Future," 06/24/24.

[41] PJM, "Energy Transition in PJM: Resource Retirements, Replacements & Risks," 02/24/23.

Hybrid "solar plus storage" projects are becoming increasingly common.

The RTO also models out an accelerated scenario in which 71 GW of fossil fuel generating capacity is retired. In that case, the ratio spikes to almost 4:1, requiring 275 GW of renewable and storage buildout. Even that 275 GW is dependent on a remaining baseload fossil fuel generation capacity of 79 GW to meet PJM's resource adequacy requirements.

So yes, renewable generation is set to boom, and transmission assets are greatly in demand. But there's one more piece to the energy transition puzzle: batteries.

Battery storage: Mitigating intermittency

As renewables' share of the grid has expanded over the past three years, standalone large-scale battery storage and hybrid projects (solar plus co-located storage) have exhibited explosive growth—which is forecast to continue (Exhibit 14).

Standalone storage is a smaller, but growing, segment of the market.

The battery technology currently used in both standalone and hybrid storage allows batteries to discharge for about four hours at their peak/nameplate capacity. For example, the Edwards & Sanborn hybrid project in the Mojave Desert, which became operational this year, pairs 875 MW of solar with 3,287 MWh of energy storage. Simple math shows that the batteries can dispense a little over 800 MW of power per hour for four hours before running out.

For hybrid projects, the batteries do much of what they're shown to do in Exhibit 12: They soak up extra solar power in the middle of the day, keeping it from overwhelming the grid, and then discharge it during the classic peak electricity demand "shoulders" of morning and evening.

Hybrid projects are becoming increasingly common, as they make larger solar generation projects possible by reducing their intermittency.

Utilities like these projects because they can buy both the electricity and the storage at the same time. Sponsors like them because when you have one customer, they want to sign a long-term PPA, and they're usually investment grade.

The price tag of hybrid projects is high—for example, the currently operating phases of Edwards & Sanborn are said to have cost $1.75 billion—but, at present, the financing market for energy transition projects remains strong.

Standalone utility-scale storage is also emerging as a financeable segment. However, these projects have different and more complex revenue streams than hybrid/co-located solar plus storage.

Standalone storage facilities have to buy power and then either 1) find someone to sell it to when it's discharging or 2) offer other services, such as system stability. These intangible services aren't all that intuitive to sponsors unless they are, say, large data centers that need ultra-reliable uptime.

In our team's experience, standalone storage doesn't make sense everywhere. It seems to work best in regions such as Texas and California, where high concentrations of wind and solar generation sources have led to a less stable grid. In these markets, having standalone storage at the ready is more valuable than in markets with low renewable penetration.

Conclusion

The next 5-10 years are likely to see a scramble for new generation and storage resources as U.S. electricity demand accelerates and fossil fuel plants are retired.

The energy transition sector is set for an era of tremendous growth. The sector's complex financing structures present attractive opportunities for creative, skilled capital providers.

These projects are especially compelling from a debt standpoint due to the predictability of PPA-derived cash flows, the high quality of counterparties and the presence of tax-driven government incentives.

Voya's energy transition team have been active lenders in this market for over 30 years, directly originating most deals ourselves due to our extensive contacts within the energy industry.

We would be happy to discuss the ways that renewable energy and infrastructure debt financing can benefit clients' private credit portfolios.

A note about risk

The principal risks are generally those attributable to bond investing. Holdings are subject to market, issuer, credit, prepayment, extension and other risks, and their values may fluctuate. Market risk is the risk that securities may decline in value due to factors affecting the securities markets or particular industries. Issuer risk is the risk that the value of a security may decline for reasons specific to the issuer, such as changes in its financial condition.

INVESTMENT
MANAGEMENT