The Co-Located Load Solution

Michael Kormos

JULY 2024

Table of Contents

The Co-Located Load Solution

An influx of large new loads is projected to seek to connect to the transmission grid over the next several years—including most prominently hyperscale data centers. The standard utility models are struggling to accommodate large data center projects on a workable timeline, and can take five years or longer.[1] These delays are caused in part by the increasing size of data centers, which in the past did not typically exceed 100 megawatts (MWs) but now can be up to ten times larger to accommodate artificial intelligence and other sophisticated applications.

Connecting a new gigawatt-sized load to the power grid would almost certainly require construction of new transmission lines, which is one of the most difficult challenges faced by the power sector. Based on my experience as chief operations officer of PJM Interconnection and related roles in the power industry, it can take up to a decade to plan, design, permit, and construct new transmission lines if they are contested (as are most large projects). These long delays risk impeding economic growth as well as technological advances.

In order to bring large data center projects online efficiently and equitably, the electric industry should be focused on finding solutions that best manage reliability, affordability (for all), efficiency, and speed. In the restructured markets where sellers compete to serve new demand, one such solution is for the new load to serve its own power needs "off the grid," and the most promising configuration is for this load to co-locate behind-the-meter with an existing power plant.

By *not taking service from the transmission grid*, the new load expedites the timeline but must pay for its behind-the-meter delivery facilities and assume the costs and risks of not being served by the grid. Because the new load has no ability to "lean" on the grid,[2] there is no need for expensive new network transmission projects to connect the load or the associated regional cost allocations to other customers. In addition, by partnering with an existing plant, the load avoids the long lead time for grid interconnection and the generator secures a steady customer, which can be critical for plants needing predictable, long-term

---

1 *See*, *e.g.*, "AI, data center electricity demand could drive advanced nuclear investment: NERC head Jim Robb," Utility Dive (June 6, 2024) ("it takes me about four years to build a substation," according to David Schleicher, Northern Virginia Electric Cooperative CEO), https://www.utilitydive.com/news/ai-data-center-electricity-demand-advanced-nuclear-investment-Robb/718181/; "Elk Grove Mayor Meets With ComEd on Substations for Data Centers," Journal & Topics (Mar. 14, 2024) (Mayor of Elk Grove, Illinois, "expressed frustration with ComEd, saying the electricity utility was 'dragging their feet' in building substations to support the growing data center industry"), https://www.journal-topics.com/articles/elk-grove-mayor-meets-with-comed-on-substations-for-data-centers/.

2 To be clear, the co-located behind-the-meter load configuration I suggest here is one where the load is unable to take energy or other services from the grid and in fact pays for and installs equipment to automatically disconnect in the event its co-located power supply trips. The load instead must rely on its co-located generator(s), batteries or other back-up resources to meet its needs. While other co-located load configurations may be considered, I have not done so here.

revenues to ensure continued operation and justify renewal of operating licenses and potential uprates.

Co-Located Behind-the-Meter Configuration

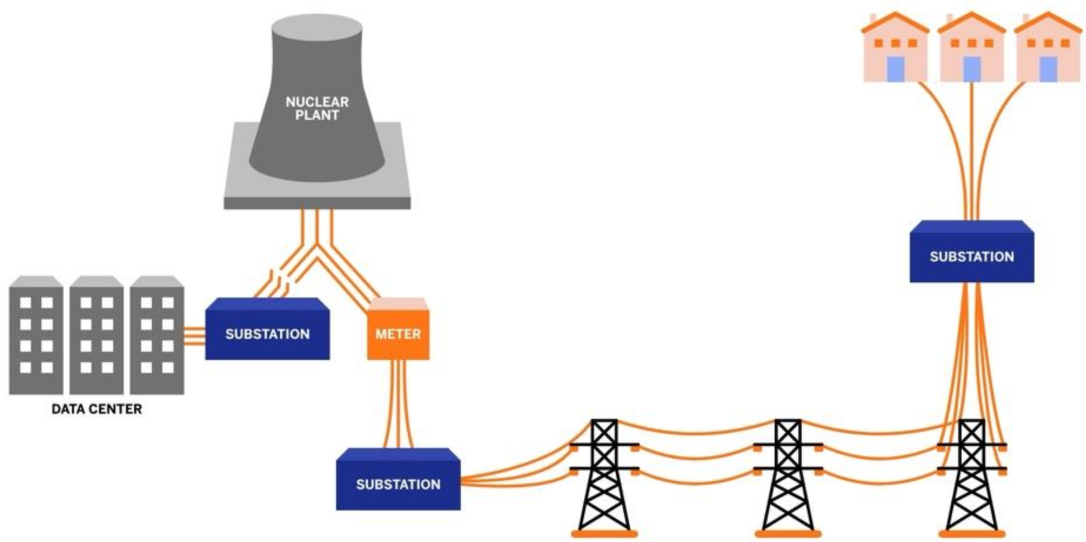

Grid-Supplied Front-of-Meter Configuration

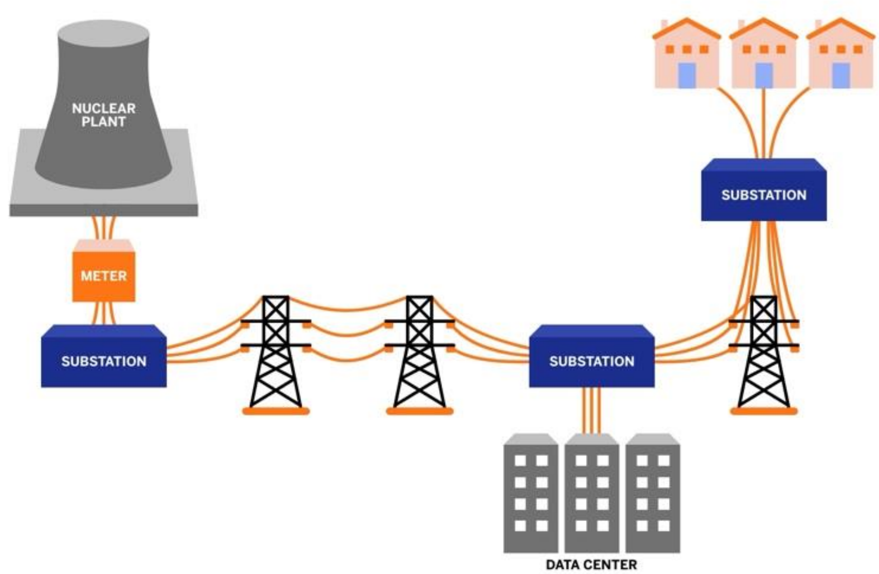

While co-location could involve the pairing of any type of power plant with any large load, existing nuclear plants provide some of the best opportunities for data centers. Nuclear plants are large, often with multiple units, carbon-free and sustainable, and capable of and preferring to run at maximum power for up to 18 to 24 months, which matches perfectly with the data center load profile.[3] Nuclear units have the highest reliability and availability of any of the existing resources.

Despite their benefits, nuclear units have, in the recent past, faced economic challenges that led many to announce retirements. In the last decade, over 10 GW of capacity retired, mostly due to economic factors, and 20 nuclear units representing 20.3 GW of capacity avoided retirement by seeking and obtaining state-based economic support. Although the federal government has stepped in to prevent further retirements through enactment of the nuclear production tax credit (PTC), that program expires in 2032 at which time federal agencies project a new wave of nuclear retirements far exceeding the generating capacity lost over the last decade.[4]

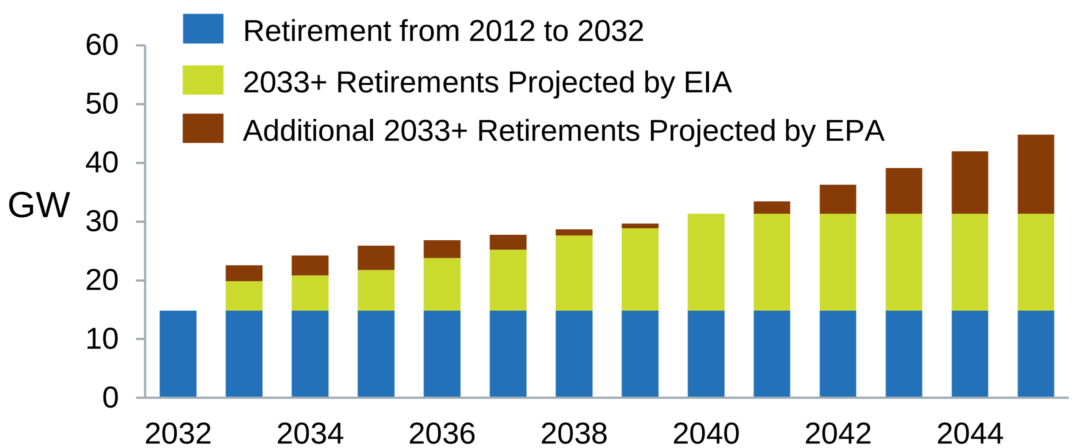

Under a co-location configuration, the data center gets the carbon-free electricity it wants without lengthy delays (but must pay for any on-site delivery facilities), and the nuclear plant gets a steady customer (forestalling premature retirement and enabling NRC license extension and potential uprates). And with the nuclear unit now supplying the data center load and not some distant network load, deliverability on the transmission grid is freed up for other existing and newly-interconnecting resources, typically wind and solar projects.

---

3 Nuclear plants have the added benefits of usually being remote, secure, and well-buffered, minimizing the potential for noise, visual or other impacts that can be associated with data center development.

4 U.S. Energy Info. Admin., *Annual Energy Outlook 2023, Table 9 (Electric Generating Capacity, Reference Case)*, https://www.eia.gov/outlooks/aeo/data/brower/#/?id=9-AEO2023&cases=ref2023&sourcekey=0; EPA, *Power Sector Modeling, Post-IRA 2022 Reference Case* (last updated Mar. 1, 2024), https://www.epa.gov/power-sector-modeling/post-ira-2022-reference-case; EPA, *Power Sector Modeling, Pre-IRA 2022 Reference Case* (last updated Dec. 13, 2023), https://www.epa.gov/power-sector-modeling/pre-ira-2022-reference-case.

Separating Co-Location Facts from General Concerns About New Load

Some have raised questions about whether pairing data centers with nuclear behind-the-meter increases costs for other customers and ratepayers via higher prices, system upgrades, or cost shifts, or otherwise harms reliability. Based on my years of experience operating power systems, I think these concerns have very little to do with co-location itself. In most cases the questions raised are a consequence of adding load *anywhere* on the grid and not *how* it is served, whether connected to the grid or co-located behind-the-meter.

1. Serving any new load will affect market prices

One frequently mentioned issue relating to co-located load is the impact on energy and capacity markets by "removing" an existing generator that is currently serving network load. All things being equal, co-locating data centers with nuclear units will not raise network prices any more than serving the same load in front of the meter in the same general location.[5] Any new demand will affect price, so unless we assume the new load *would disappear or not otherwise be served*, how the load is served does not materially change the effect on market prices. While every situation is unique, dedicating a portion of existing generation to a particular customer behind-the-meter through a direct connection will have the same effect on the supply/demand dynamic as serving the same amount of new load through deliveries over the transmission system from participating in the market or from a remote generator under a power purchase agreement. This new load, regardless of configuration, can be met using existing market and system planning tools.

2. Serving any new load may affect infrastructure costs

Another common misconception is that if an existing generator is used to serve a behind-the-meter load, new infrastructure (most likely transmission) will be needed and existing customers will have to pay for it. In my experience, it is far more likely that connecting a data center to the grid in front of the meter will require more transmission upgrades than co-locating it behind a generator, which does not rely upon the grid for service.

Larger loads (those approaching 1000 MWs) must connect to the grid on an extra high voltage line (230 kV or greater). This requires an extra high voltage substation and associated facilities to accommodate the front-of-the-meter data center. The construction costs for this type of transmission project will depend on the circumstances but can range from $150 million up to $250 million.[6] In a co-location configuration, these costs will be paid

---

5 Given both nuclear plants and data centers operate generally as baseload facilities, the impact of adding new load or "removing" such supply will have virtually the same effect on market prices.

6 *See* "Transmission Expansion Advisory Committee – PPL Supplemental Projects," April 2, 2004 at https://www.pjm.com/-/media/committees-groups/committees/teac/2024/20240402/20240402-item-08---ppl-supplemental-projects.ashx (identifying $244 million PPL supplemental project); *see also* PJM, Transmission Expansion Advisory Committee: AEP Supplemental Projects, at 7 (June 4, 2024),

by the co-located load. In a front-of-the-meter configuration, by contrast, some of these costs will be allocated to grid customers.

Moreover, in a front-of-the-meter configuration, there will be additional costs associated with other transmission upgrades required to get power to the new substation built for that customer. The costs can be significant. In Northern Virginia, for example—where no data centers have co-located—data center growth combined with generator retirements has required over $5 billion of transmission investment to reliably serve the new load. The costs of those new transmission facilities, like all transmission facilities, will be *shared by all customers* in accordance with PJM's and the transmission owner's tariffs.

As discussed later, serving that same load behind-the-meter at a power plant can reduce the need for new transmission lines compared to a grid-connected configuration. The nuclear unit is also giving up use of the transmission system to others. Some transmission investment might be required in a co-location configuration to replace the generation with excess or new units. But if the load were in front of the meter, it is far more likely that a greater amount of transmission investment would be needed. The overall costs to grid customers should be significantly less in a co-location configuration.

So again, unless we are simply assuming these large new loads will not be served, any addition of that load will result in incremental infrastructure costs and the co-located configuration ensures that the costs incurred to serve the new load are paid for by the data center, not by customers of the surrounding utilities.

3. Serving any new load may affect reliability

The next common refrain is reliability, and that we should not let nuclear resources "take megawatts off the grid" to serve new loads because we cannot adequately replace the critical generation in a timely manner. Again, the issue is not whether we are serving the load behind-the-meter versus using the transmission grid to serve the load in front of the meter but instead is about serving new load, period. And again, unless one assumes the new load *would not otherwise be built or delayed for years,* the effect on grid reliability is the same. Tools already are in place in every restructured market to bring on any new resources needed to serve network load.

As an aside, the challenges of integrating new generation resources to replace retiring fossil generation and to meet all new load are a rightful focus of policymakers. In my view, markets should provide better incentives to attract new resources and the generation interconnection process should be significantly improved to get needed resources timely connected. What should not happen, however, is discrimination against one type of large new load (data centers) or one type of generation (nuclear). It cannot be that, after many years of financial struggles by nuclear units, we are now troubled by arrangements with counterparties that are willing to contract to ensure continued operation of those units. We

---

https://www.pjm.com/-/media/committees-groups/committees/teac/2024/20240604/20240604-item-05---aep-supplemental-projects.ashx (identifying $155.69 million AEP supplemental project).

must find mechanisms to allow all new types of load to be connected in their preferred configuration and on their timeline and not claim some new load can connect but others must wait or not be served at all.

4. Behind-the-meter service imposes no new costs on grid customers

The last general claim is that behind-the-meter load causes "cost shifts" that negatively impact existing customers. For example, in the recent Exelon/AEP protest to the Susquehanna Interconnection Service Agreement amendment, a consultant from Concentric claimed up to $140 million in "cost shifts" from the data center to grid customers. This calculation is nothing more than the revenue that the transmission owner would have been paid if the data center had connected in front of the meter where it would take grid service and benefit from being connected from the grid. But in the co-location scenario, the extension cord to the grid is cut – the data center does not cause grid costs to be incurred, cannot take any service from the grid, and is not a customer of the transmission owner. There are no costs to shift in the co-location scenario.

It is true that existing grid customers cannot share the costs of their service with the data center supplying its own service, but no one can reasonably expect to share costs with someone *who is not taking service*.

Not only is there no cost shift, but co-location can help grid customers save money. The existing grid customers do not pay for new costs assumed by the data center or the costs for any upgrades identified in the host generator's updated interconnection studies. As elaborated below, many risks (*e.g.*, outage risk) are directly the responsibility of the loads in this co-located configuration, alleviating the need for other customers to share in costs to cover such risks.

* * * * *

We must not compare serving new data center load behind-the-meter with not serving it at all, as some suggest. Adding new load in any location may add new costs, so the question is how the configurations at issue—grid-connected or behind-the-meter—result in efficiency, reliability, and affordability for *all customers*, including the new data center load.

The Benefits of Co-Location

The beneficiaries of co-location behind-the-meter extend well beyond the data center and the host generator. In many cases, these configurations are better for all than if the data center were served in front of the meter. Co-location behind-the-meter helps to:

1. Serve the Load

Remember that co-location allows us to serve large new loads that otherwise would have to wait years for service, and thus to meet the technology needs and fuel economic growth on a much more expedited basis. As I noted above, connecting a new gigawatt-sized load to the power grid will almost always require construction of new transmission infrastructure, which can take years to plan, design, permit, and construct. Our goal should always be to serve customers when and how they want to be served, and co-location offers an excellent alternative.

2. Improve Grid Efficiency

The best and most cost-efficient way to supply a new large load is by generating as close as possible to that customer. The further the generation is from the load, the more expensive it becomes to move the power and the more at risk the system is to overloading existing transmission lines. By placing a data center where it can be directly served by a generator versus locating the load somewhere remote from the ultimate generation needed to serve it, the grid requires fewer upgrades to serve that new demand.

Similarly, the geographic proximity inherent in co-location also is likely to reduce energy losses ("line losses") resulting from transmitting long-distances from the power plant to data center, which at the transmission level are in the range of 1-3 percent. For a gigawatt-sized data center, that would avoid the loss of 90,000-260,000 megawatt hours of electricity. Given that large data center customers generally seek sources of carbon-free power, this prevents the unnecessary loss of an increasingly important commodity: clean electricity.

3. Transfer Risk from the Grid to the Co-Located Load

Greater risks equal greater costs and the co-locating data center load bears its own risks, while imposing no incremental risks on the network.

The risks the data center itself takes on are significant. If the data center were supplied from the grid, it could expect supply certainty in 99+% of hours given the diversity of resources in a region like PJM. This is very different for load connected to a single resource, even a very-well run nuclear resource that can be expected to run ~93% of the time.[7] To the extent a co-located load seeks to improve the reliability of the individual resource(s) it is co-located with,

---

[7] *See* DOE, Office of Nuclear Energy, *5 Fast Facts About Nuclear Energy* (June 11, 2024), https://www.energy.gov/ne/articles/5-fast-facts-about-nuclear-energy#:~:text=Nuclear%20energy%20is%20one%20of%20the%20most%20reliable%20energy%20sources%20in%20America (noting that nuclear power plants operated at full capacity more than 93% of the time in 2023).

it must pay for additional back-up (*e.g.*, on site back-up generators, batteries, etc.). Some customers might also choose to size to one unit at a dual unit site and use the second unit during refueling outages, *i.e.*, have the second unit be part of its additional supply.

Another way to view this is from the perspective of the data center itself: a grid-supplied data center avails itself of the redundancy and high availability inherent in a wholesale market overseen by an RTO as well as all the grid-supplied ancillary services. If the data center were in front of the meter and thus on the network, if one grid supply resource fails, another would be started to make up for the lost output from the failed resource, allowing grid connected load to continue seamless service under most scenarios.

In contrast, a co-located center must accept the outage risk associated with a discrete resource which will always be higher than grid power even if that resource is individually highly reliable. It also must carry—and pay for—its own reserves. This is a critical distinction between the two configurations: the services co-located data centers receive from a generator are not the same as the services the data center would receive if supplied by the power grid.

4. Charge Data Centers Instead of Grid Customers for More of the Transmission Facilities

Not only are risks transferred to the data center, so too are many of the costs. The data center pays for the private behind-the-meter delivery facilities as well as the electricity. And as I discussed above, in my experience it is very likely that a data center connecting to the grid will impose significantly more grid upgrade costs that will be socialized than a data center supplied behind-the-meter. The costs for this can be substantial (easily in the billions) and can be spread beyond the data center customer to all other customers.

5. Protect Reliability

Reliability always is paramount, and the data center/nuclear pairing will be studied appropriately to ensure reliability is maintained on the electric grid. If anything, co-location helps reliability by not trying to move more power a further distance.

The independent grid operator responsible for reliability studies the impacts of the new configuration at the nuclear facility to ensure reliability is maintained. In the PJM process, for example, PJM conducts a "necessary study" to determine whether a generator's modification to include behind-the-meter load has any reliability impact on the generator's interconnection with the grid.[8] If so, the generator pays for changes at its facility and any necessary network upgrades to cure the potential issue. PJM also reduces the available capacity for sale through the reduction of the capacity interconnection rights of the generator to reflect any behind-the-meter sales. Changes to the Interconnection Service

[8] PJM has provided detailed guidance for the current process to connect co-located load and all the steps required to ensure reliability. *See* PJM Guidance on Co-Located Load (March 22, 2024) (Updated April 17, 2024), https://pjm.com/-/media/markets-ops/rpm/rpm-auction-info/pjm-guidance-on-co-located-load.ashx. Numerous proposals to adjust these rules have been raised at PJM and in other fora.

Agreement on file at FERC may also be made to clarify reliability procedures given the behind-the-meter configuration.

Because a co-located data center served in a behind-the-meter scenario has equipment to prevent it from taking service from the grid, the grid will not be planned for or otherwise accommodate the load. Data centers almost always have their own backup power supply, further reducing demands that would exist if their backup supply was the network, or if the backup supply itself relied on the network for transmission services.

In the case of a nuclear unit, it will also remain connected to the network, as required by law. Any power at the nuclear station not committed to the behind-the-meter load still will be available for sale to the network or to others. The electricity is not disappearing or retiring. It is being used to serve load, just as it would if it were injecting into the grid and delivering to the data center over the transmission system.

6. Facilitate Steady Customers for Nuclear Units to Remain in Service

Speaking of nuclear, the behind-the-meter configuration supports long-term investment in nuclear power plants and the grid and environmental benefits they provide. It creates the financial security needed to support a subsequent license renewal (the application itself cost tens of millions of dollars), making it more likely the nuclear plant and its emissions-free output remain available well into the future and able to consider uprate projects to increase output.

Over the last decade, nuclear resources faced significant financial uncertainty and one-third of the fleet either retired or obtained state support to preserve the assets and prevent higher prices and increased emissions.[9] The federal government followed with the PTC, which provides near-term financial stability for the nuclear fleet. However, it expires in 2032 just as a large portion of the fleet will be going through the NRC regulatory process to extend operating licenses, which can take five years or longer to complete. Thirty percent of the merchant nuclear units will need to renew operating licenses in the next 10 years to prevent shutdown, increasing to 50 percent over the following decade. Hosting a data center with a long-term power sales agreement would certainly play the key role in a unit owner's decision

[9] *See*, *e.g.*, D. Murphy & M. Berkman, The Brattle Group, The Impacts of Illinois Nuclear Power Plants on the Economy and the Environment (2019) (prepared for Ill. IBEW State Council and Ill. AFL-CIO), https://www.brattle.com/wpcontent/uploads/2021/05/17147_the_impacts_of_illinois_nuclear_power_plants_on_the_economy_and_the_environment.pdf; D. Murphy & M. Berkman, The Brattle Group, Pennsylvania Nuclear Power Plants' Contribution to the State Economy (2016) (prepared for Penn. Building and Construction Trades Council, et al.), https://www.brattle.com/wpcontent/uploads/2017/10/5732_pennsylvania_nuclear_power_plants_contribution_to_the_state_economy.pdf; D. Murphy & M. Berkman, The Brattle Group, Salem and Hope Creek Nuclear Power Plants' Contribution to the New Jersey and Local Economies (2020) (prepared for PSEG), https://www.brattle.com/wpcontent/uploads/2021/05/20628_salem_and_hope_creek_nuclear_power_plants_contribution_to_the_new_jersey_an d_local_economies.pdf.

of whether to undertake the multi-year regulatory process and related investments needed to relicense.

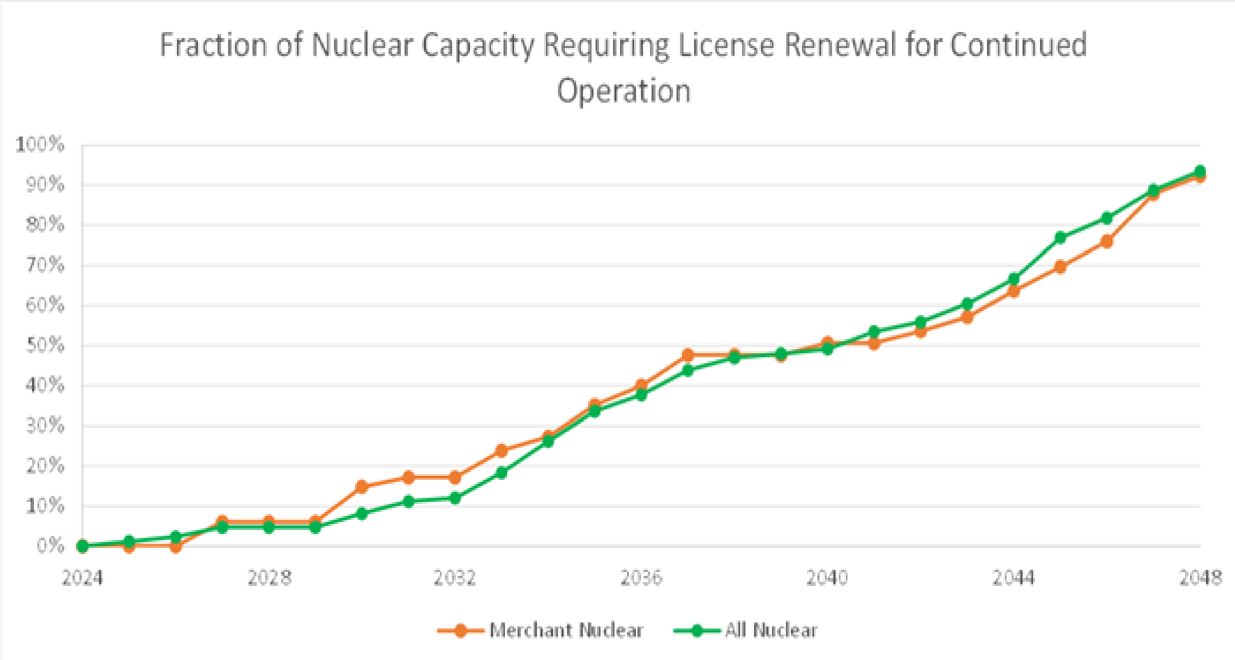

7. Increase Interconnection Opportunities for Renewables

Finally, other beneficiaries of co-location would be renewable projects being curtailed because of inadequate capacity on the transmission system to deliver from remote locations to load centers. Co-location of data center and nuclear generation frees up transmission headroom for new generator resources by making available the transmission capacity currently used by the host power plant. This benefit can be significant, such as in Illinois where transmission congestion is projected to increase in the next decade as the state adds more wind and solar generation to meet decarbonization goals. Co-locating data centers at Illinois nuclear plants could reduce curtailment of wind and solar output by over 80%.[10]

10 Analysis provided by Constellation Energy reflecting anticipated wind and solar expansion needed to comply with the Climate and Equitable Jobs Act.

Renewable Curtailments – Annual, GWh (Bars)
Renewable Curtailments – Cumulative, GWh (Lines)

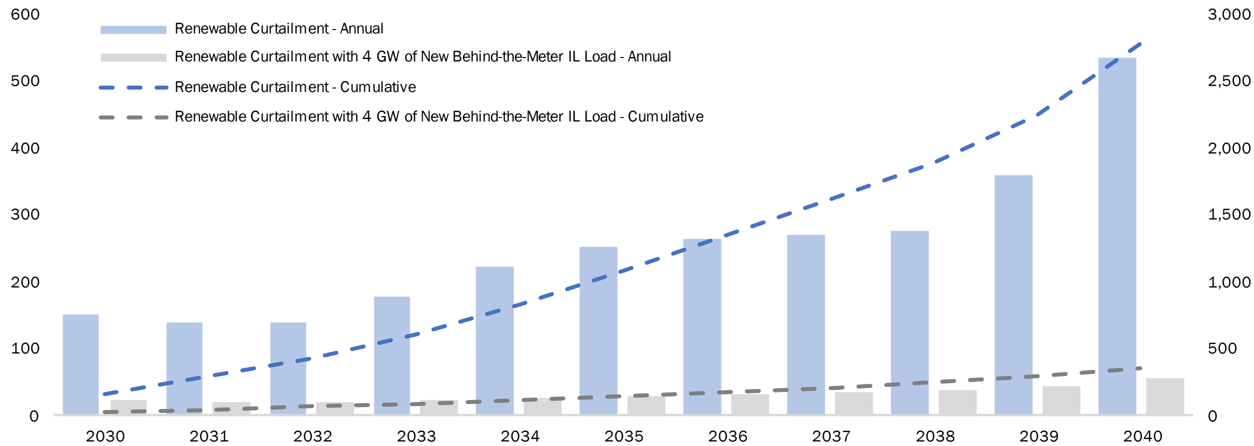

Co-locating data centers at nuclear plants also could benefit renewable projects idling in the interconnection queue. As mentioned above, PJM's rules, for example, require a co-locating generating resource to forgo its capacity interconnection rights which can then free up capacity rights for another resource, which based on PJM's current queue is likely to be wind, solar or batteries. This can allow those resources to more quickly connect to the transmission grid with no (or fewer) transmission upgrades needed.

There Is No Cost Shift Without a Fundamental Market Redesign

Notwithstanding these benefits and the fact that new load connecting behind-the-meter imposes no new costs on grid customers—and cannot use grid power—some opposing voices still claim that transmission and related services are being provided to the disconnected load through the co-located generator. They claim because the *generator* remains connected to the grid, any load connected to the generator is "free riding" on the network including because they are synchronized through the generator.[11] For proof, they point to ancillary and other grid services provided to generation from the grid and claim that the co-located load benefits. This argument is creative but reflects a fundamental misunderstanding of the open access rate design established by the Federal Energy Regulatory Commission (FERC) thirty years ago.

While it is true nuclear plants must remain connected to the grid, any surface appeal this "cost shift through the generator" argument has falls apart when we look at current practices and precedent. Generators always have been connected to the grid and have never been

11 Simply being synchronized to the grid through a generator which in turn is synchronized to the power grid does not mean the co-located load receives the same service as those relying on the grid for service. This misses the point elaborated below that the co-located load is not taking service from the grid, even if the co-located generator may be.

charged for most of the services in question. This is a core design choice made by FERC when it adopted its open access transmission requirements in Order No. 888.[12]

Like any grid-connected generator, the co-located generator will continue to deliver any generation not consumed by the data center onto the grid for the benefit of the network's customers. Reducing the output of these generators has no material bearing on the system or its costs. It would be discriminatory to charge only a few generators for these services and would take a fundamental market design shift to start charging all generators for things like fluctuations in their output before these alleged cost shifts could be credible. In other words, for there to be a credible cost shift here, two things would need to happen: *first*, FERC would need to overhaul the basic open access rates applicable to all generators to start charging them for these services; and, *second*, co-located generators would have to be the only generators that do not pay them.[13]

Such a market redesign also would require regulators to reevaluate how costs are allocated when generation is located behind a grid customer's meter, i.e., when it is *generation* that is behind-the-meter, not *load*. Load that owns generation, such as a municipal customer with city-owned generation or a retail customer with roof-top solar, use that generation to reduce the volume of grid-delivered electricity for which they pay. It would be highly discriminatory to apply grid charges to load co-located behind-the-meter of a generator without doing the same to the portion of a customer's load met by behind-the-meter generation.

To put all this in context, let's consider each type of grid service and how it applies to a co-located behind-the-meter load configuration where the generator is connected to the grid and the co-located load takes power only from the generator:

- ***Transmission service*** – Since the co-located load is not able to take power from the grid, no system power is transmitted from the grid to the load. Indeed, equipment is installed to automatically switch the load off in the event the nuclear resource and any backup generation supply become unavailable so that the co-located load is never served by the grid. Any power that is transmitted over the grid is an injection from the existing nuclear unit to and for the benefit of network load and—like all generating resources in every RTO/ISO market—the generator is not required to pay any transmission service, either network or point to point to serve internal network load. The generator has an

12 *Promoting Wholesale Competition Through Open Access Non-Discriminatory Transmission Services by Public Utilities; Recovery of Stranded Costs by Public Utilities and Transmitting Utilities*, Order No. 888, FERC Stats. & Regs. ¶ 31,036 (1996).

13 Prompted by the increase in variable-output renewable generation across the country, FERC undertook a multi-year investigation fifteen years ago to evaluate whether open access rules should be altered to assign generators the cost of certain ancillary services. FERC declined to do so, instead providing a framework for any public utility that might seek to impose such costs on generators that includes, among other things, requirements to justify with operational data any distinction between different types of generators. *Integration of Variable Energy Resources*, 139 FERC ¶ 61,246 at P 315-335 (2012).

interconnection service agreement, which will be amended as needed, that provides for the ability to inject power into grid. No cost shift.

- ***Distribution service*** – Since any power supplied by the co-located resources to the load would be provided over a privately-owned line and not a utility's distribution (or transmission) system, no distribution services are being provided by the utility. Sale of power from the generating facility to the co-located load remains subject to the state's authority, including how the state assesses charges used to fund infrastructure and social programs. Depending on the state, there could be state-level charges that co-located load would pay—not for utility distribution infrastructure, but for those state-jurisdictional services or programs. No cost shift.

- ***Capacity*** – Whether the new load is in front of the meter or behind-the-meter being served by existing resources, new resources or excess from existing resources will be needed to meet remaining customer demand. As described above, reducing the amount of output an existing generator sells into the market or adding new load to be served by that market has the same effect on capacity—one adds to the "demand" side of the equation, while the other subtracts from the "supply" side.[14] No cost shift.

- ***Energy*** – Since the co-located load would have no ability to draw power from the grid, there is no grid energy purchased. Any energy sold into the market by the generator in excess of what the co-located load is consuming is treated like every other generator connected to the grid, subject to both day-ahead and real-time pricing. As with any generator connected to the grid, generators have the right and the ability to sell energy to anyone on the grid including externally to other systems. Like capacity, any new load behind-the-meter or in front of the meter is going to increase energy requirements by either reducing the supply curve or increasing the demand curve but the effect on the overall market is the same. No cost shift.

- ***Ancillary Services***

  - **Regulation and Frequency Response (aka Load Following)** – Variations in the co-located generator/load balance will change the amount of megawatts a generator is injecting into the grid. This is not the same as withdrawing power from the grid. While it is true the grid is absorbing the variations in the generator's injections, the grid does this for every generator regardless of whether on-site load is being served. Many resource types, particularly intermittent wind and solar, inject into the grid at levels that vary from moment to moment, with grid injections fluctuating up (or down) quite quickly and significantly in response to variable weather patterns or other factors. No generator—regardless of type—in any RTO/ISO region pays for regulation service to

[14] While current PJM rules do not give a generator serving behind-the-meter load the ability to retain its rights to sell that capacity into the PJM markets, those rules could evolve in a way that enables the generator to retain such rights (and the related obligations), making that alternative available to co-located data center customers.

ensure the grid is prepared to absorb their moment-to-moment fluctuations in power output. Nuclear units are no exception. It would require a fundamental policy change to begin charging all generators for this service, as discussed above. No cost shift without fundamental market redesign.

- **Operating Reserves** – Reserves are primarily to respond to a reduction in generator output or outages; they allow the grid to make up for a sudden change in the generation/load balance.[15] As with regulation service, generators do not pay for reserve services regardless of generator type and regardless of whether they have co-located, on-site load. As for the co-located load, it provides its own reserves through its supply arrangement with the host generator and has no ability to obtain reserves from the grid, nor does the presence of the co-located load increase or affect the existing grid reserves required in any way. Instead, if its co-located generator has an outage, the co-located load is on its own to either shed its load or provide its own reserves via interruptible power supplies and backup generation. No cost shift without fundamental market redesign.

- **Reactive Power** – Co-located generators are under the same NERC and RTO requirements to provide reactive power to the system as any other units. If, for example, PJM determines as part of its "necessary study" that any reactive deficiencies are caused by the co-location configuration, it will require the co-located generator to supply the reactive power—most likely through the installations of capacitors at the location that are paid for entirely by the generator—before the co-located load is connected to the generator. Generator already pays/contributes in kind.

- **Black Start** – Today no generator pays for black start, yet all generators (and all loads) benefit from the grid being restored. To the extent the nuclear plant's location is included and prioritized in any restoration plan, FERC could change the current paradigm and find that it may be appropriate for generators to pay a share regardless of whether it hosts co-located load. The cost would be minimal. Using PJM as an example, assigning a 1,000 MW nuclear unit a share of PJM's total $73 million annual black start costs would amount to a $590,000 per year charge for the nuclear unit that reduces costs assigned to others by 0.008 percent. Even if black start charges were assessed to generators, it does not follow that full network service should apply to co-located load. Generators could pay, but it would require fundamental market redesign.

- **Station Power Services and Emergency Services Supporting Nuclear Units** – To the extent that a nuclear resource is relying on grid-supplied power to meet its station power service needs and/or NRC license obligations, it should pay for those services. This is no different than how any other resource connected to the grid is treated and there is no reason to alter that treatment just because there is a load behind that

15 Rapid load changes or dispatch error can also cause the needs for reserves.

resource because that load is unable to avail itself of those services. Generator already pays.

The only way serving co-located behind-the-meter generation could create a cost shift to grid customers is if we implement market design changes to begin charging *all* generators for load following, reserves, and black start but then excused a generator supporting a behind-the-meter load from those charges. I am not aware of any current efforts to enact any such proposals as the general view is that the provision of these services to generators *benefits the grid much more than it costs*.

Next Steps

Policymakers, regulators, and industry have a basic obligation to serve load. There is no dispute that it will take years for large data centers and other loads to connect to the grid under the standard utility models. Our collective focus should be on finding solutions that best manage reliability, affordability (for all), efficiency, and speed, without standing in the way of data centers choosing to pay for and provide their own transmission facilities and power supply (and backup power supply). Under the co-location behind-the-meter model, data centers and nuclear plants can do so without risking reliability or shifting costs to network customers, including those in other states. This data center/nuclear pairing is not only symbiotic for the parties, but for anyone with an interest in the promises of the digital age and a sustainable future.

Biography

Michael Kormos is an electric utility consultant with extensive experience solving complex problems and leading large teams across the energy industry. Until 2021, Mr. Kormos was Sr. Vice President, Transmission & Compliance for Exelon Corporation, where he was responsible for the operations and long-term planning for transmission systems of Exelon utilities Atlantic City Electric, Baltimore Gas & Electric, ComEd, Delmarva Power, PECO, and PEPCO. Mr. Kormos shaped Exelon and its utilities' policy positions on issues related to the Federal Energy Regulatory Commission, Regional Transmission Organizations (RTOs), and national transmission matters.

Before joining Exelon, Mr. Kormos spent 27 years at PJM Interconnection, the RTO responsible for transmission systems in 13 Mid-Atlantic and Midwestern states, and Washington, D.C. As PJM's Chief Operations Officer, he supported the development and implementation of competitive wholesale markets in PJM and earlier in his career at PJM, he oversaw an operational expansion that tripled size of the system controlled by PJM. He also held management and engineering positions in transmission operations while at PJM.

Mr. Kormos has also served on the board of director for Reliability First Corporation; an officer of the Eastern Interconnection Planning Collaborative and the Eastern Interconnection Data Sharing Network; and as a member of the North America Energy Reliability Corporation operating committee.