INL/MIS-26-91487

NUCLEAR POWER
FOR DATA CENTERS
PLAYBOOK

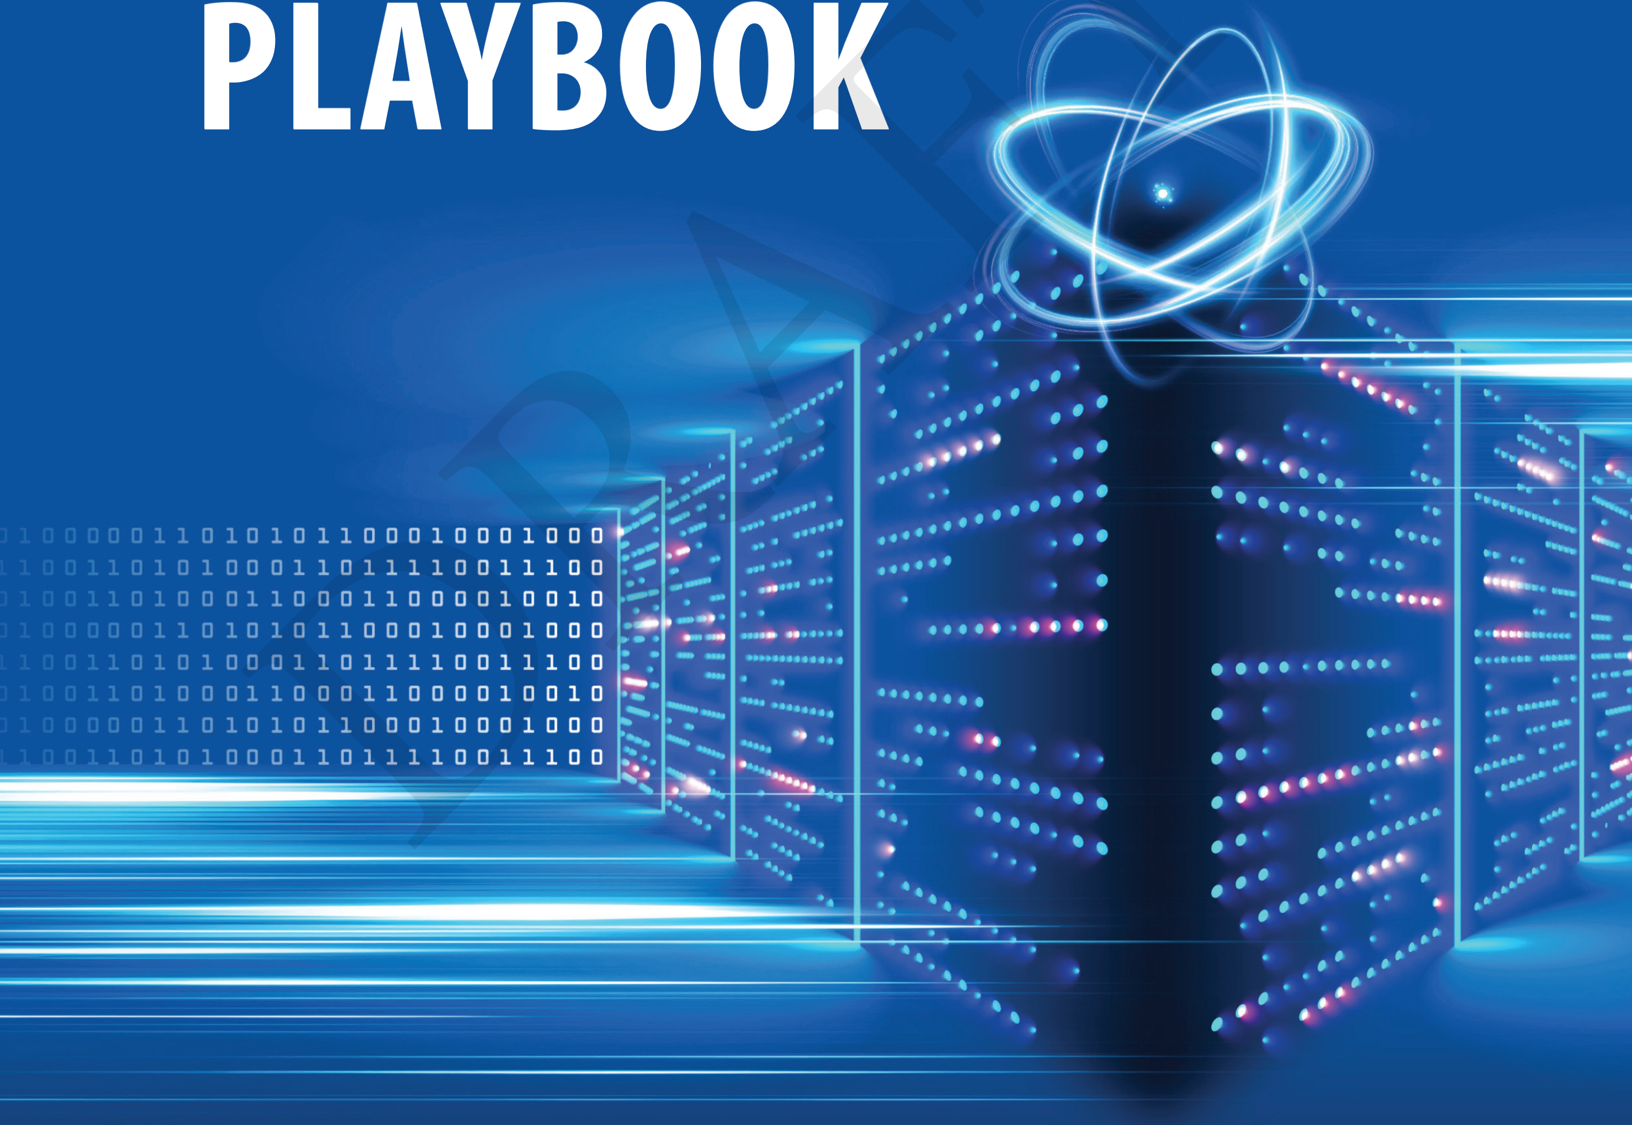

Battelle Energy Alliance manages INL for the
U.S. Department of Energy's Office of Nuclear Energy

INL Idaho National Laboratory

INL/MIS-26-91487

NUCLEAR POWER FOR DATA CENTERS PLAYBOOK

April 2026

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

AUTHORS AND CONTRIBUTORS

Authors

Aaron Epiney[1]
Bikash Poudel[1]
Brian Smith[1]
Chris Dietz[1]
Gabriel J. Soto[1]
Iza Lantgios[1]
Jonathan Tacke[1]
Kara Cafferty[1]
Lori Braase[1]
Megan Culler[1]
Nicolas Stauff[2]
Paul Talbot[1]
Rami M. Saeed[1]
Ryan Cobey[1]
So-Bin Cho[2]
Temitayo Olowu[1]
Tim McJunkin[1]
Will Jenson[1]

Contributors

Alexis David[1]
Barney Hadden[1]
Botros Hanna[1]
Brandon Miller[1]
Byung-Hee Choi[1]
Chris Lohse[1]
Eric Whiting[1]
Femi Omitaomu[3]
Jacob Bryan[1]
Jake P. Gentle[1]
Junyung Kim[1]
Neil Harward[1]
Nikki Peterson[1]
Ryan Spanger[1]
Sally Glass[1]
Tate Shorthill[1]

1 Idaho National Laboratory
2 Argonne National Laboratory
3 Oak Ridge National Laboratory

Idaho National Laboratory

Argonne NATIONAL LABORATORY

OAK RIDGE National Laboratory

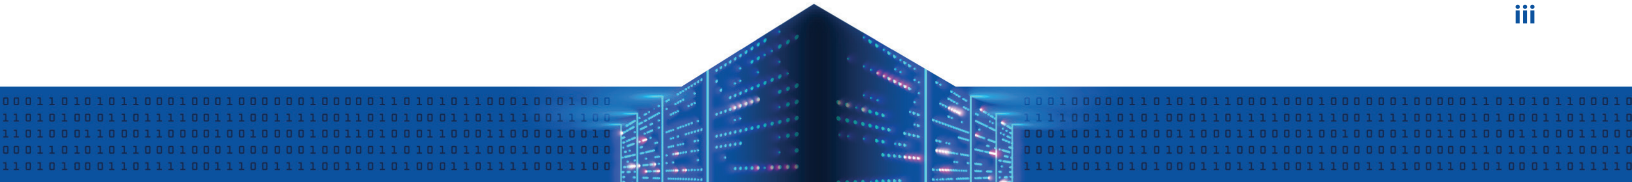

TABLE OF CONTENTS

FIGURES

TABLES

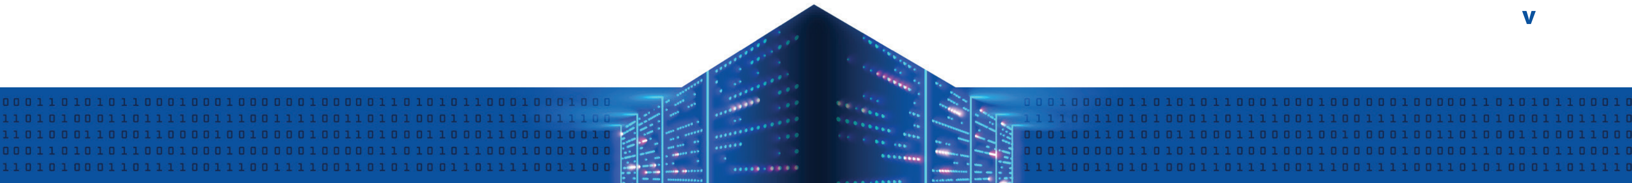

ACRONYMS

AC	Alternating current
ADVANCE	Accelerating Deployment of Versatile, Advanced Nuclear for Clean Energy
AI	Artificial intelligence
ANL	Argonne National Laboratory
ASCE	American Society of Civil Engineers
ATS	Automatic transfer switches
AVR	Automatic voltage regulator
BESS	Battery energy storage systems
BLS	Bureau of Labor Statistics
BOP	Balance of plant
BWR	Boiling-water reactor
CAISO	California Independent System Operator
CAPEX	Capital expenditure
CCP	Combined cooling and power
CFR	Code of Federal Regulations
CHIL	Control and validation through control hardware-in-loop
CHP	Combined heat and power
COP	Coefficient of performance
CRAC	Referred to as computer room air conditioners
CRAH	Computer room air handlers
DOE	Department of Energy
EDA	Economic Development Administration
EDF	Energy Dominance Financing
EIA	Energy Information Administration
EO	Executive orders
EPA	Environmental Protection Agency
EPC	Engineering, procurement, and construction
EPRI	Electric Power Research Institute
EPU	Extended Power Uprates
FEED	Front-end engineering design
FFR	Fast frequency response
FOAK	First-of-a-kind
GAIN	Gateway for Accelerated Innovation in Nuclear
HALEU	High-assay, low-enriched uranium
HERON	Holistic Energy Resource Optimization Network
HP	Heat pump
HPC	High-performance computing
HTGR	High-temperature gas-cooled reactors
HUD	Housing and Urban Development
INL	Idaho National Laboratory
IRA	Inflation Reduction Act
IRP	Integrated resource planning
ISO	Independent system operators
ITC	Investment Tax Credit
ITIC	Information Technology Industry Council
LCOE	Levelized cost of electricity
LEU	Low Enriched Uranium
LR	Learning rate
LTRA	Long-term reliability assessment
LWR	Light-water reactors
MACS	Microgrid Automated Control System
MAGNET	Microreactor Agile Non-Nuclear Experiment test
MIB	Microgrid in a Box
MSR	Molten salt reactors
MUR	Measurement Uncertainty Recapture
NE	Nuclear Energy
NERC	North American Electric Reliability Corporation
NLR	National Laboratory of the Rockies
NPP	Nuclear power plant
NPV	Net present value
NRC	Nuclear Regulatory Commission
O&M	Operation and maintenance
OCC	Overnight capital cost
PC	Planning Coordinators
PDU	Power-distribution units
PIR	Peak-IT-to-reactor
PJM	Pennsylvania-New Jersey-Maryland
PNNL	Pacific Northwest National Laboratory
POI	Point of interconnection
PPA	Power-purchase agreements
PQ	Power quality
PTC	Production Tax Credit
PUC	Public utility commission
PWR	Pressurized water reactor
RMS	Root-mean square
RTDS	Real-Time Digital Simulator
RTO	Regional transmission organizations
SA&I	System Analysis and Integration
SFR	Sodium-cooled fast reactors
SPP	Southwest Power Pool
SPU	Stretch Power Uprates
SR	Small reactors
T&D	Transmission and distribution
TEA	Technoeconomic analysis
TES	Thermal-energy storage
TP	Transmission Planners
TRISO	Tristructural- isotropic
UK	United Kingdom
USACE	United States Army Corps of Engineers
VAR	Volt-amp reactive
VC	Vapor-compression chiller
VHTR	Very high temperature reactors
WECC	Western Electricity Coordinating Council

INTRODUCTION

Purpose and Scope

The “Power for Data Centers Playbook” is a comprehensive guide designed to consolidate Department of Energy (DOE) research and provide actionable insights for data-center developers who aim to integrate nuclear energy into their power strategies. This playbook serves as a resource for the data center industry to navigate the complexities of nuclear power adoption, offering practical, research-backed information tailored to various use cases across scales and integration models. The intention is for this playbook to help guide digital infrastructure companies to develop their specific, individual roadmaps towards reliable power, including nuclear.

Recognizing the importance of industry collaboration, this playbook has been developed in partnership with data center companies to ensure it remains grounded in real world applications and avoids being created in a DOE/ government “vacuum.” Feedback from industry stakeholders has been a critical component, shaping the playbook’s recommendations and ensuring its practicality.

This playbook consolidates DOE research from programs spanning multiple national laboratories and incorporates industry insights, providing a balanced, actionable resource for data-center developers to explore nuclear-power solutions tailored to their specific needs.

How to Use This Playbook

The playbook is designed to be modular, allowing users to focus on specific chapters relevant to their needs. Each chapter can be read independently, making it easy for data-center developers to find targeted guidance without needing to navigate the entire document. For example:

- A developer interested in grid-connected solutions can focus on Chapter 1: “The Role of the Grid”
- A company exploring off-grid or hybrid solutions can refer to Chapter 2: “Power Beyond the Grid”
- Those seeking to understand the pathway to nuclear-energy adoption can turn to Chapter 4: “Road to Nuclear.”

To facilitate this modular approach, the playbook includes a Use-Case Matrix, which maps various data center scenarios—ranging from distributed small-scale facilities to hyperscale operations —and links them to relevant chapters. It also outlines different power-integration models, from grid-connected to fully islanded systems, ensuring that users can easily identify the sections most applicable to their specific circumstances.

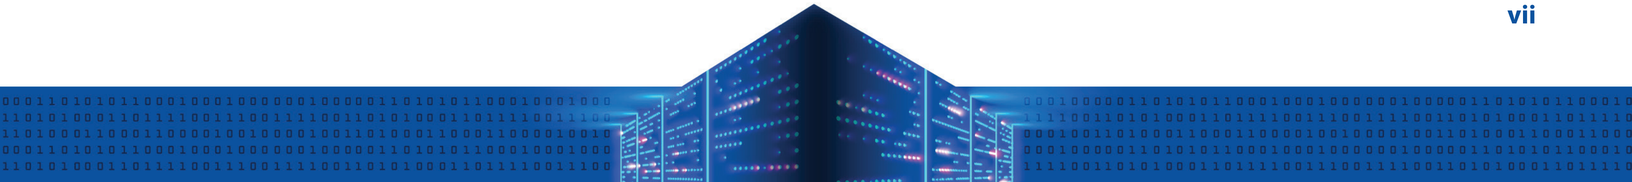

AUDIENCE

This playbook is intended for:

- Data center companies that seek to develop roadmaps for nuclear-power integration
- Energy planners and consultants who support data center operations
- Policy and decision makers who are interested in the intersection of data center growth and energy abundance.

The audience also includes stakeholders from the energy sector, including utilities, nuclear developers, and technology providers, who can use the playbook to align their offerings with data center needs.

Use-Case Matrix

The matrix in Table 1 outlines key data-center use cases and links them to relevant chapters of the playbook:

Table 1. Relevant chapters by use.

Use Case	Scale	Integration Model	Relevant Chapters
Distributed Small Data Centers	Small (multiple sites)	Grid-connected	Chapter 1
Hybrid Power Solutions	Medium	Grid + On-site Generation	Chapter 1, Chapter 2
Islanded Data Centers	Medium to Large	Completely Islanded	Chapter 2
Hyperscale Data Centers	Large	Grid-connected or Hybrid	Chapter 1, Chapter 2, Chapter 3
Nuclear-Only Power Solutions	Any	Fully Nuclear	Chapter 3

The matrix highlights the diversity of data-center needs and how the playbook's chapters address them. Each use case links directly to the relevant chapters, ensuring users can navigate to the most pertinent information.

Chapter 1

THE ROLE OF THE GRID

The electrical grid has long been understood as a background feature of modern life—reliable, invisible, and taken for granted. But as data centers have grown from modest server rooms into colossal campuses consuming power at a scale once reserved for industrial cities, that quiet background has been pushed to the foreground. A question that was once largely operational has become strategic: What role does the grid play in powering the data centers of tomorrow, and how is that role being redefined?

This chapter is organized to build understanding from the ground up. It opens by examining how large electrical loads are characterized and classified. How big is "big," and does a facility's consumption behavior matter as much as its size? The section on **Large-Load Classification and Characterization** introduces a taxonomy developed by INL to address these questions. From there, the internal standards that govern data center reliability are explored in the section on **Data Center Performance Requirements**, including how power architectures built around backup systems and cooling equipment interact with, and can sometimes disrupt, the broader grid.

The chapter then turns outward, to the processes and pressures that shape a data center's relationship with the grid over time. The section on **Interconnection Planning and Communications** walks through the formal process by which a facility is connected to the grid, while **Grid Stability and Reliability Requirements** examines how large, electronically intensive campuses can affect system frequency and voltage, and how those risks are being managed. The section on **Utility Planning Considerations** addresses a tension that nearly every data-center developer will encounter: the mismatch between a hyperscale facility's 18-month construction window and a utility's decade-long planning cycle.

Finally, the section on **Capacity Expansion and Market Studies** situates the data center buildout within the broader national picture of electricity supply and demand—where demand is growing fastest, where transmission constraints are binding, and how regional market structures shape the options available to developers seeking reliable, long-term power. Together, these sections are intended to provide a clear view of the grid as it currently exists: its capabilities, its constraints, and the pressures now being placed upon it.

1.1 Large Load Classification and Characterization

Introduction

The United States (U.S.) electrical grid is always contending with the introduction of ever larger and more diverse load types, as well as the need for enhanced stability that these loads introduce. As the scale of industrial operations and technological development has expanded, these high demand electricity customers have come to be termed "large loads." Grid operators and utilities consistently adapt and evolve planning and operational practices to serve these customers of increasing size, introducing advanced load forecasting, demand-side management programs, real time monitoring, and more. Today, new types of large loads are proliferating rapidly, across wide geographic areas, and reaching unprecedented levels of energy demand. This comes alongside the emergence of distributed energy systems. Together the changing paradigm has broad implications for grid reliability.

Why Does This Matter for Data Centers?

Despite the significant impact of large loads on the grid, no universal definitions or categorizations exist for these singular electricity customers; rather, classification varies regionally and by organization. Facilitating greater levels of integration creates an increasing need for a common approach to characterize large loads and understand their behaviors and attributes. Notably, the North American Electric Reliability Corporation (NERC) has created the Large Load Task Force, which has published a white paper that discusses the characteristics and risks of large loads. This document begins to detail differences amongst these types of loads, but emphasizes the need for greater examination moving forward, including "a framework for classifying large loads."[1]

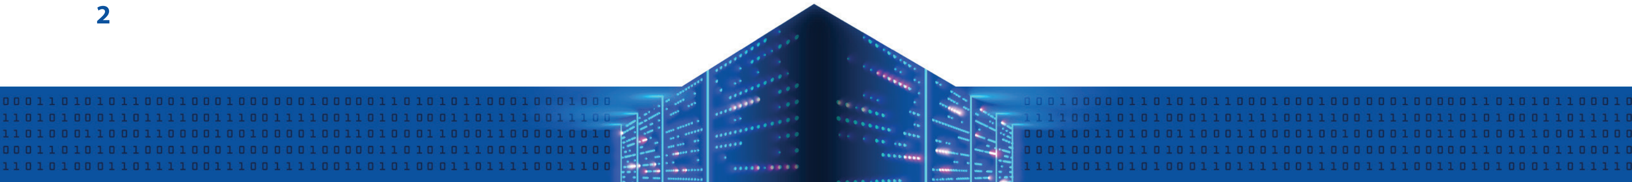

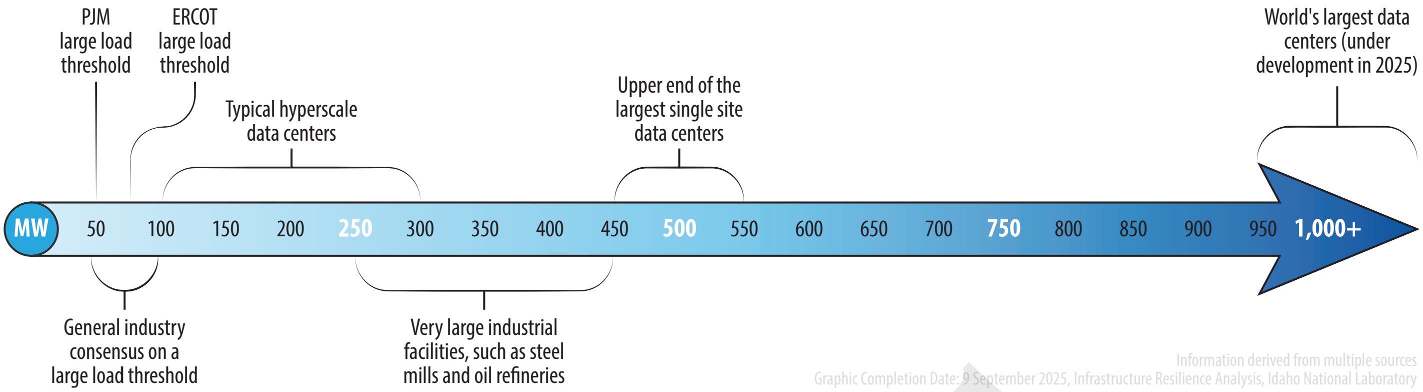

Information derived from multiple sources
Graphic Completion Date: 9 September 2025, Infrastructure Resilience Analysis, Idaho National Laboratory

Figure 1. Large load landscape.

Research and Findings

Data centers have emerged as technological and economic pillars with significant energy and reliability demands. The operational behavior of data centers can differ greatly from traditional industrial loads, presenting new challenges and potential opportunities for the stability of the grid. The fundamental role of any data center is to provide large quantities of reliable data storage, management, and processing, but these functions are applied in different ways and for different end results, depending on the data center.[2] For example, some, grouped as enterprise data centers, are dedicated to a single organization's use while some, referred to as colocation data centers, exist to serve numerous clients. Edge and modular data centers tend to be smaller and serve specialized use cases, while hyperscale data centers seek to maximize size and computing power.

To facilitate operational planning, the Idaho National Laboratory (INL) has developed a taxonomy to classify large loads. This taxonomy provides a common language to describe how diverse load types, including data-center campuses, will interact with the bulk energy system. While a full catalogue of all the different operational characteristics of different load types would be quite lengthy, from a grid integration perspective, loads can be described according to a small set of operational characteristics. The taxonomy proposed herein classifies large loads based on three characteristics: load size, load variability, and load flexibility.

Load size serves as a simple and fundamental quantitative differentiator amongst large loads. Intuitively, load size describes the net impact a facility may have on the grid, for better or worse. As depicted in Figure 1, large loads can extend from tens of megawatts (MWs) to hundreds of MWs, with gigawatt-sized loads on the horizon.[3] The larger the load, the greater its operation affects the system, and the greater potential instability effects are associated with it.[4] At the same time, however, these significant loads can potentially offer larger-scale stability services to the grid. Nameplate capacity is used to measure the size of a load, but there are nuances to using this single metric for size. Although most loads will rarely operate at their full nameplate capacity, and many large loads take several years to build up to this full capacity, this metric provides consideration for the maximum impact the load could have on the connected system, which allows for conservative stability analysis.

NERC notes that other factors, such as interconnection voltage, size of the load relative to the local system, and overall connection could also affect the impact the load has on the system.[1] This is the sort of nuance that plays a key factor when performing interconnection studies. For the purposes of this taxonomy, however, large loads are divided into three basic groups, based on size thresholds. Where loads interconnect and at what voltage will largely correspond with these groupings. Table 2 describes these ranges in more detail

Table 2. Load size.

Size 1: 50-99 MW	These loads typically include large commercial facilities, data centers, and mid-sized industrial operations. While substantial, they generally integrate into existing grid infrastructure with moderate upgrades. Their size allows for relatively flexible siting and interconnection, making them a common entry point for load growth in both urban and rural areas.
Size 2: 100-999 MW	This category encompasses major industrial facilities: aluminum smelters, steel mills, and large-scale hydrogen production plants. These loads require substantial grid infrastructure and may require dedicated transmission upgrades or new substations.
Size 3: 1,000+ MW	These are rare but transformative loads, including multi-gigawatt data center campuses, direct air-capture facilities, or large-scale electrification of industrial clusters. Their integration can reshape regional transmission planning, resource adequacy, and even market dynamics.

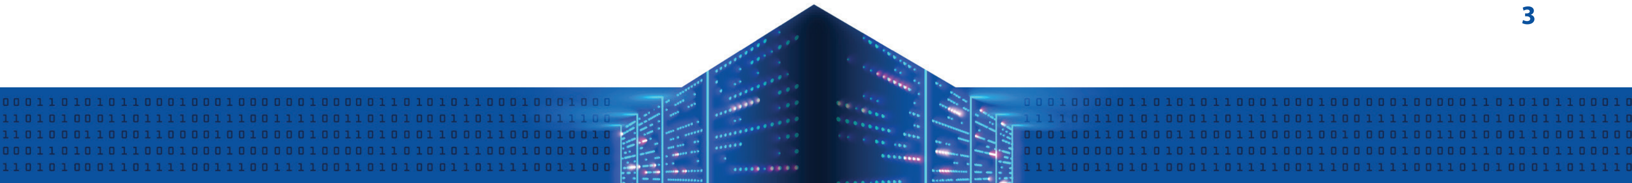

Load variability reflects an electricity customer's power use behavior, especially as it relates to the degree and consistency of power demand. Some loads have a relatively "flat," predictable power demand, even if that demand is high. This load variability is representative of steady, sustained operations with limited ups and downs in power demand. Variations in power use will occur, but they are well-known or predictable and do not introduce stability concerns. Alternatively, a load's power usage can fluctuate significantly and rapidly, and do so in ways that are unforeseen by utilities and balancing authorities. This can include sudden large spikes and significant drops in power usage.

Load profiles for large electricity customers can be important to stability because significant and rapid changes can cause unacceptable voltage or frequency deviations. However, some fast-acting loads have the potential to support the grid if they can be controlled. The taxonomy captures these important operational differences by identifying a load as having either a slow or fast varying load profile, as described in Table 3.

Table 3. Slow vs. fast loads

Slow Load Profile	A load with consistent and predictable power demand that rarely experiences significant increases and/or decreases in power usage over short periods. Such a load can be said to vary slowly. An example is a cloud storage data center that seeks to maintain a very level load profile for 24×7 operations. For purposes of this taxonomy, a "significant" power increase or decrease is equal to at least 50 MW,[i] and a "short period" is considered to be 60 minutes or less. "Rare" occurrences of these rapid power swings are considered to be less than 12 times per year or only during grid instability events. In practice, these thresholds will vary by interconnection[i], but in general these values are those that would spur a balancing authority to take remedial action.
Fast Load Profile	A load that, by virtue of its business operations, experiences significant increases and/or decreases in power usage over short periods. An example is a steel mill with arc furnaces that can widely vary their heat production function during the course of a production run.

Load flexibility also relates to increases and decreases in power usage, but from the standpoint of voluntary action to support grid stability. This can involve both ramping down and ramping up over defined periods in response to grid needs. Some key factors that support flexibility include the presence of adequate onsite power generation, single tenancy,[ii] and the ability to work around issues of criticality. The definition of flexibility in this taxonomy combines considerations from NERC to classify loads based on firm or flexible behavior and consideration of behind-the-meter (BTM) co-located generators.[1]

A load can be considered flexible if it can alter its power consumption from the grid in a controlled way. The key is that a flexible load can perform these power-demand changes (either up or down) without harming its essential operational and business functions, and in response to requests from the grid operator. This is analogous to dispatchability in generators. This characteristic is of interest in grid stability because it determines whether a given load represents a potential asset or liability for balancing authorities. Table 4 describes load flexibility as defined in the taxonomy.

Table 4. Flexible vs. inflexible loads.

Flexible Load	A load that can automatically raise or lower its net power consumption based on requests from the grid (i.e., ancillary services), while avoiding its own operational degradation.
Inflexible Load	A load that is unable to adjust power consumption quickly for purposes of ancillary services due to resultant operational impacts. This inability can be driven by factors such as lack of onsite storage or generation, lack of grid communications capability, mission criticality, and/or organizational rules and policies.

While many factors can influence how large loads are distinguished, the three fundamental characteristics presented above allow for a basic taxonomy. This taxonomy can be used to begin categorizing large loads into like groups from the standpoint of potential grid stability challenges, but also ancillary service market participation. The large load taxonomy is presented below in Figure 2. Based on the three fundamental characteristics of load size, load variability, and load flexibility, the taxonomy can result in 12 basic large load types. For example, a slow, inflexible load that has a peak demand of 200 MW would be a **Type B-2**, while a fast, flexible load with peak demand of 80 MW would be a **Type C-1**.

Figure 2. Large load taxonomy.

			Load Size		
	Load Variability	Load Flexibility	50-99 MW	100-999 MW	>1,000 MW
Type A	Slow	Flexible	1	2	3
Type B	Slow	Inflexible	1	2	3
Type C	Fast	Flexible	1	2	3
Type D	Fast	Inflexible	1	2	3

Provided below are two examples of hyperscale data centers and how they would be categorized within the taxonomy as large loads.

Example 1: A hyperscale cloud-services data center operates at a 400 MW capacity. The facility delivers continuous computing capacity for such global services as web browsing, streaming, and hosting enterprise software used for various business processes. Its baseline electricity demand is consistently high, resulting in a relatively predictable, static load profile that experiences

---

[i] According to FERC Order 842 (RM16-6), the maximum allowable frequency deadband for primary frequency response is ±0.036 Hz (0.06% of 60 Hz). If system frequency deviates beyond this range, generators must respond. Since frequency deviation is proportional to the imbalance between generation and load, a simultaneous 0.06% change in all loads across an interconnection would trigger a system-wide generator response within one second. This implies that a single load varying by 0.06% of the interconnection's total capacity can have the same effect. Therefore, any load that fluctuates by more than this threshold—especially faster than the spot market's typical 1-hour clearing time—requires additional reserves and is considered "variable." To simplify, we assume a 50 MW threshold already accounts for interconnection size. Thus, any load capable of changing by more than 50 MW per hour under normal conditions is classified as "fast-moving" in our taxonomy.

[ii] For example, data center owners and operators that house and serve multiple tenants are generally not in a position to make flexible power commitments on behalf of their tenants.

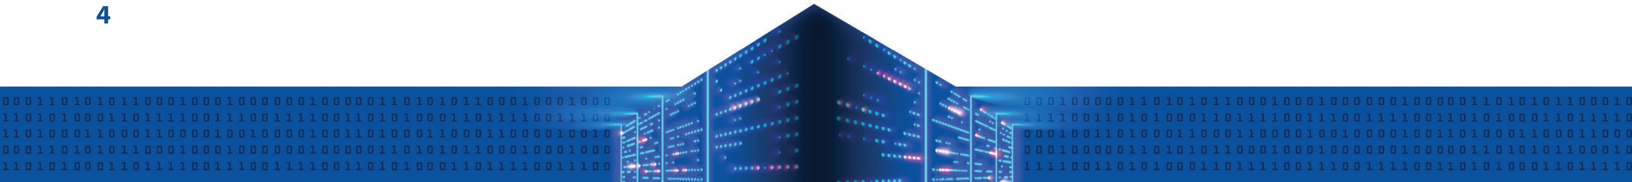

very little fluctuation. Any decrease in external power automatically activates the data center's uninterruptible power supply (UPS) battery systems, which can sustain the facility for several minutes, allowing for backup diesel-powered generators to cycle on and keep servers online for much longer periods of time without support from the grid. As a result, it provides grid operators with a stable base load that can still flex in response to contingencies. This large load is a Type A-2 within the taxonomy. Examples of actual facilities that may fit this profile include Google's large data center in Council Bluffs, Iowa, and Amazon's data center clusters in Virginia and Oregon.

Example 2: A large hyperscale data center operates at a 1.5 GW capacity and houses technically advanced information-technology infrastructure designed to handle computing-intensive tasks needed to train artificial intelligence (AI) models. Unlike hyperscale cloud data centers (which run steady, customer-facing workloads), this AI-training facility typically runs jobs in batches that can cause large fluctuations in electrical demand. Training runs may last for days or weeks, but they also feature idle periods or ramp-down windows between jobs. Though it experiences large, rapid power fluctuations, the facility's infrastructure allows for flexible scheduling to manage AI jobs and rapidly shift non-urgent tasks to reduce peak load. This large load is a Type C-3 within the taxonomy. Examples of actual facilities that may fit this profile include Meta's planned AI-training data center in Richland Parish, Louisiana, and the Vantage Frontier Campus project in Shackelford County, Texas.

Key Insights

1. Like all loads, the impact of a large load on power system operations depends on its size.
2. In addition to size, the effect of a large load on system stability is determined by its speed of variation and dispatchability.
3. Other considerations such as interconnection voltage are important when considering local stability impacts.

1.2 Data Center Performance Requirements

Introduction

The mission profile of data centers, which underpins cloud services, financial transactions, healthcare records, and AI computing, demands continuous service with extremely tight tolerances on power quality, thermal conditions, and latency. The Uptime Institute established a tiering system for data centers more than 30 years ago, and these definitions remain an industry standard.[5]

The tiers are tied to particular business functions and define criteria for maintenance, power, cooling, and fault capabilities. Both topology and operational sustainability are used to establish these criteria:

1. A Tier I data center has basic capability to support IT for an office setting and beyond. Even a Tier I facility requires a UPS for power usage, outages, and spikes; a dedicated cooling system; and backup generation for power outages. Tier I facilities require 99.671% uptime per year, which translates to a maximum of 28.8 hours of downtime per year.[6]
2. A Tier II data center requires additional redundant capabilities for power and cooling, including both engine generators and energy storage. Tier II facilities are common for small- to medium-sized businesses. A Tier II data center requires 99.741% uptime annually, which translates to a maximum of 22 hours of downtime.
3. A Tier III facility offers additional reliability in the form of N+1 redundancy, so performance is not impacted if any single component fails. Tier III facilities should not require shutdowns for equipment maintenance or replacement due to the built-in redundancy. Tier III data centers have 99.982% uptime, which corresponds to less than 1.6 hours of downtime per year.
4. A Tier IV data center is 2N+1 redundant, meaning it has twice the operational capacity required, and an additional backup in case a failure occurs while using the secondary system. Tier IV facilities have 99.995% annual uptime, corresponding to less than 26.3 minutes of downtime per year.

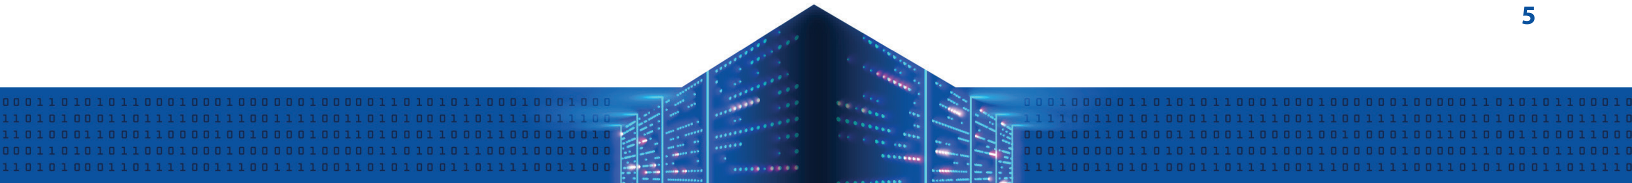

Why Does This Matter for Data Centers?

Given the high-performance requirements, data centers have complex power-management and distribution infrastructure, along with sensitive power-quality monitoring. The remainder of this section discusses data-center parameters that affect the need for high uptime and power-quality requirements and how these requirements impact data center behaviors as they interact with the grid. While much of this information will be familiar to data center stakeholders, this section serves as a useful primer for utilities seeking to interconnect these loads.

Research and Findings

Power Architecture Supporting Uptime

A typical data center's power distribution begins at the utility's point of interconnection (POI) and continues through high-voltage switchgear, step-down transformation, medium-voltage distribution, UPS systems, power-distribution units (PDUs), and finally the IT racks and critical mechanical systems: e.g., cooling, air handling, and pumps (see Figure 3).

Several components in this power-management system are critical to the behavior of the data center. For example, the computer-based equipment, which are the critical loads at the end of the power-distribution chain, are governed by the Information Technology Industry Council (ITIC) curve, which defines the alternating current (AC) voltage limits that most IT equipment can experience without unexpected shutdowns or malfunctions. However, event analysis has shown that behaviors in the real world do not align entirely to the ITIC curve, leading to multiple instances of partial loss of load.[6] As another example, the ride-through behavior of large loads is also affected by the ride-through characteristics of cooling equipment. A NERC incident review found that cooling load relies on variable-frequency drives and variable-speed drives, which include low-voltage protection for internal components.[7] Even brief voltage surges, resulting from faults normally cleared by the grid, can trip cooling-equipment load, which can quickly lead to IT equipment's overheating and forcing a complete shutdown.

UPS Systems

UPS systems are the first line of defense against sub-cycle voltage depressions and short outages, providing bridge power while generators start and protecting sensitive loads from flicker, sags, and inrush. Across facilities, three design families dominate, with distinct grid-interaction signatures documented in incident analyses[7]:

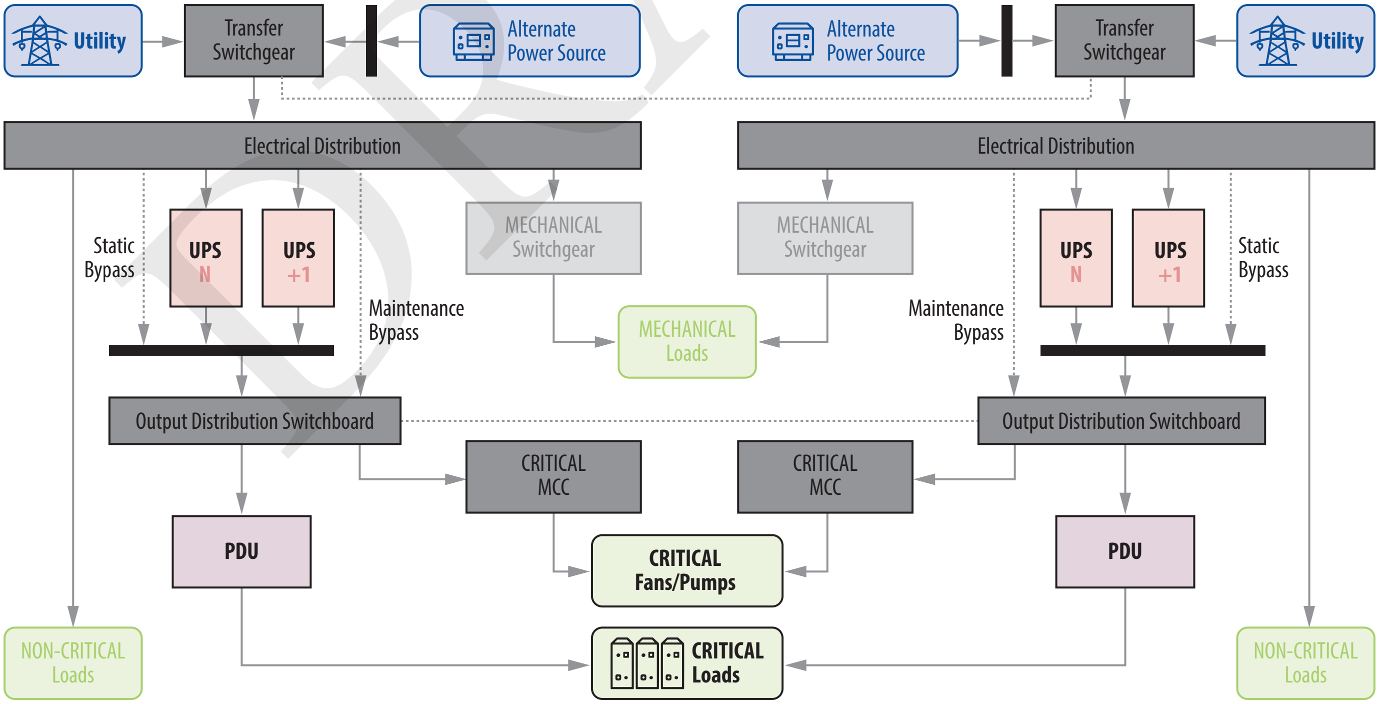

Figure 3: Centralized UPS power distribution. Source: NERC Incident Review: Considering Simultaneous Voltage-Sensitive Load Reductions.

1. Static centralized UPS (2–5 MW blocks). Battery-backed converters instantaneously take over during a transient. Because most utility faults clear in ~20–100 milliseconds (ms), centralized, static UPS sites usually return to grid quickly, so the system "sees" a short load loss that restores within seconds.
2. Static decentralized UPS (rack-level, 3–4 kilowatt (kW) each). Thousands of small UPS units act in aggregate; behavior is broadly similar to centralized static UPS, but can vary by vendor controls, firmware, and local settings.
3. Dynamic rotary UPS (DRUPS). A flywheel carries the load immediately; a coupled diesel engine then assumes prime-mover duty. Critically, DRUPS sites do not auto-reconnect quickly after the grid disturbance clears; reconnection is often deliberate and manual. This produces longer-duration load drops from the grid perspective.

Even with UPS systems in place to regulate power quality, many facilities implement disturbance-counting schemes, for example, "three voltage depressions within 1 minute." When that threshold is hit, sites transfer to backup generation and remain islanded until manual reconnection even if each sag lasts only tens of milliseconds.[7] This behavior contributed to more than 1,260 MW remaining offline during a 2024 Eastern Interconnection event, despite correct utility protection clearing.

Voltage Sensitivity, Load Disconnection Behavior, Ride-Through Expectations, and Grid Impacts

Data centers are voltage-sensitive, power-electronics-heavy loads. Their protective philosophy prioritizes mission continuity: avoid server brownouts and cooling trips that could damage equipment or corrupt data. UPS and transfer logic thus favor fast isolation from transient instability. However, the scale of modern campuses means that the aggregate effect of these transfers can present as bulk load loss to the grid, with frequency and voltage excursions and reactive imbalances that operators must counteract in real time.[2] This has led to multiple approaches for utilities to mitigate the risk of large load loss, including ride-through requirements, underfrequency load shedding, and pushes for higher power-usage effectiveness to reduce non-computational loads.

Large, sudden transfers to backup cause frequency rises, overvoltage conditions, and stressed reactive controls. Balancing the needs for grid stability with power quality and availability for data center creates a cross-sector opportunity that must be addressed through technical and policy solutions.[10]

Backup Power Systems for Extended Outages

UPS designs guarantee instantaneous continuity, but sustained uptime during grid unavailability depends on backup generation and storage sized to the mission. Onsite backup-power options often include some form of synchronous generation, like diesel or natural-gas generators, along with energy-storage systems. Table 5 provides an overview of characteristics from various backup power options, though it is not exhaustive.[11]

Table 5. Backup-power options.

Asset Type	Duration	Speed of Response	Notes
UPS	Minutes to hours	<1s	UPS systems are only intended to sustain operations until longer duration assets can be deployed
Diesel	Hours to days	5–10 s	Long duration diesel readiness, ramp coordination, and segmented load pickup are essential to avoid overshoot/undervoltage during transitions back to grid service
Natural Gas	Months to years	minutes	Industrial gas engines/turbines offer longer runtimes and preferential characteristics, but introduce pipeline dependency and gas system outage risk
Battery Energy Storage System	1–4 hours	<1s	Controls must coordinate with UPS, generator governors, and facility segmentation to prevent forced oscillations or converter driven instabilities during islanded operation
On-site Nuclear	Months to years	5–15 minutes	Nuclear power provides continuous power that reduces reliance on frequent transfers and can sustain operations through extended grid stress, but must be coordinated with other resources for tight load-following capabilities due to limited ramp rates

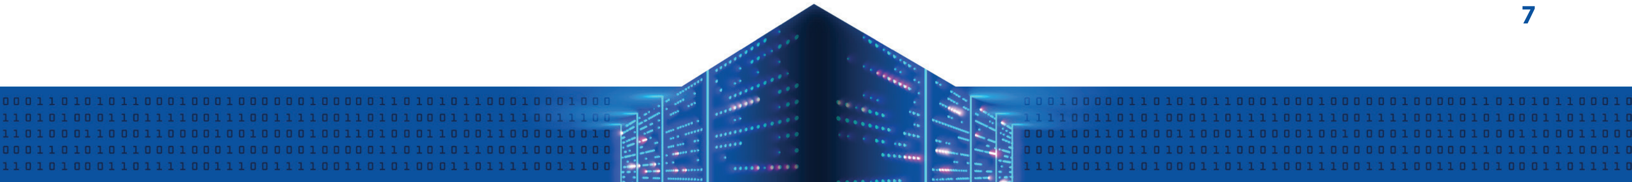

Key Insights

1. Voltage disturbances can trigger large, sudden data-center load losses, primarily due to protective behavior of UPSs, cooling systems, and transfer controls.
2. UPS architecture drives grid impacts because, for example, static UPS systems return quickly while DRUPS and disturbance-counting controls can lead to sustained islanding.
3. Reconnection must be managed carefully because uncontrolled or simultaneous restoration can create frequency and voltage challenges equal to, or greater than, the initial disturbance.
4. Current modeling and visibility tools lag behind real behavior, making improved telemetry, shared forecasting, and validated dynamic models essential for planning and operations.

Further Exploration

Future work should focus on improving data center and grid-operator coordination around disturbances, reconnection, and long-duration backup power in order to protect both data center uptime requirements and power-system reliability. Opportunities exist to develop standardized guidance for ride-through, undervoltage settings, and cooling-system protection. Continued collaboration among utilities, equipment manufacturers, national laboratories, NERC working groups, and data-center operators can help refine models, validate UPS and cooling-load behavior, and create reconnection practices that ensure stable voltage and frequency recovery.

Exploration of such advanced backup solutions as long-duration storage, hybrid microgrids, and emerging nuclear systems to enhance resilience during extended outages would provide high value. Parallel efforts to expand high-resolution telemetry, improve dynamic load models, and establish secure data-sharing arrangements between operators and facilities will help close gaps in visibility and forecasting as data centers grow in scale and complexity.

1.3 Interconnection Planning and Communications

Introduction

Large electrical loads like many new hyperscale data centers must be formally integrated into the transmission or distribution system in a way that ensures reliability, safety, and operational compatibility with the electric network. When a data center connects with the grid, several factors need to be accounted for to ensure overall grid reliability. These factors are focused primarily around the following questions:

1. How much maximum power will the data center need?
2. Is the data center's power demand expected to fluctuate or change?
3. Will changes in power demand happen frequently and/or randomly?

In essence, interconnection involves a structured engineering process that ensures very large, continuous, and often fast-growing customers like data centers can be added to the grid without compromising system performance or reliability. The process for connecting data centers to the power grid involves a coordinated technical and contractual pathway in which the grid operator and the local utility study how the new load will affect power flows, equipment ratings, stability, and power quality. Interconnection agreements form the backbone of this process. Through an interconnection agreement, the data center and the power provider define responsibilities, costs, technical requirements, and operating rules.

Why Does This Matter for Data Centers?

The following information is designed to provide data-center operators with a general understanding of the steps that power providers and certain balancing authorities require to properly interconnect with the grid. To set the stage, some key components to data center connections are provided first, followed by a high-level description of how data centers, utilities, and independent system operators (ISOs) and regional transmission organizations (RTOs) negotiate and carry out data-center development and interconnection processes. Certain nuances to this process are then discussed, which include specific characteristics of hyperscale data centers and the state of current standards and regulatory requirements.

Research and Findings

While data centers can vary substantially in their power needs based on their operational characteristics, most data centers contain specific componentry and exhibit similar BTM setups. Figure 4 provides a visual diagram of typical data center components and shows how they are generally connected to external power sources. These components can be categorized into the following groups:

1. **BTM supply** consists of switchgear or automatic transfer switches (ATS) that receive power from an external transformer and distribute power across data center components. This main power distribution unit generally feeds into an UPS which ensures the data center stays temporarily operational through battery storage. The UPS is connected to one or several PDUs which regulate voltage and power networking equipment and servers.[8]
2. **Energy storage systems and backup generation** are generally connected to the data center's main power distribution unit/ ATS, and at a minimum support the vital IT loads of the data center. Many data centers use diesel or natural gas generators for backup purposes. Battery energy storage systems (BESS) are also often used as a method for meeting instantaneous power needs in the event of a loss of external power.[8]
3. **Primary server components** are composed of the servers, storage, and networking equipment that form the operational backbone of data centers.
4. **Temperature control systems** regulate the ambient temperature of servers and IT equipment, and are often referred to as computer room air conditioners with built-in condensers or computer room air handlers that require an external cooling system, such as a water chiller plant. These systems receive power from the data center's main distribution unit.[8]

The way interconnection agreements are applied to a data center's specific BTM power resources can vary substantially based on how these resources are configured and their potential to cause or exacerbate a grid contingency.[iii] Impact studies, carried out in step two of Table 6, are a fundamental part of the interconnection process and play a pivotal role in determining the overall time it takes for large loads like data centers to become operational. These studies are subject to certain baseline requirements defined in NERC Standard FAC-002-4, which establishes criteria for assessing the reliability of generation, transmission, and end-user facilities that plan to interconnect with the grid.[12]

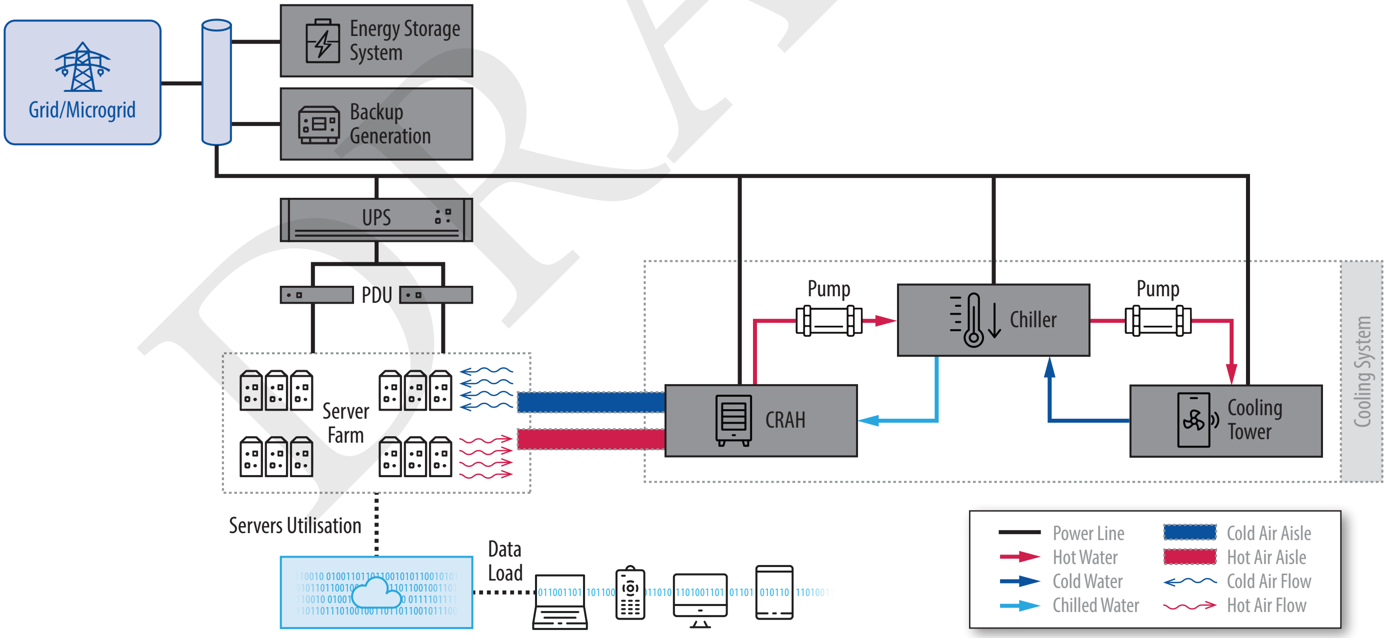

Figure 4. BTM power flow to different components in a data center.[8]

[iii] A contingency refers to an unanticipated failure of a specific system component, such as a generator, transmission line, circuit breaker, or switch.

Table 6. General process of an interconnection agreement.

Establish Contact	Data center contacts power provider to request grid connection, provides expected load size (MW), project timeline, and other relevant power characteristics (backup generation & storage, cooling, HVAC, etc.). Power provider assesses its capacity to support the data center with current infrastructure. Data center and power provider discuss the likely complexity, costs, and timeline for interconnection.
Conduct Impact Studies[*][13,14]	Data center provides technical details on equipment (UPS systems, cooling, inverters) to the power provider. Power provider conducts engineering studies on local lines, substations, and protection systems. This entails checking thermal limits, voltage, fault levels, and equipment ratings. ISOs and RTOs may also assess the need for new lines or transformers, and conduct grid stability tests.[†]
Sign Interconnection Agreement	Power provider issues a standardized interconnection agreement[‡] that includes: Specific points of connection, equipment needed for the connection, and who is responsible for equipment on either side of the connection. Required upgrades (lines, transformers, substations) and which party is responsible for financing each upgrade. Operating rules, which include day-to-day communication, data sharing, safety protections, managing variable power – keeping harmonics and voltage within required limits, and specific clauses for grid emergencies (temporary curtailment, on site backups, ancillary services). Data center agrees to these terms and agreement is signed.
Build Facilities	Power provider builds the 'grid-side' infrastructure. This can include high-voltage substation equipment, transmission lines, feeders, metering, and telecommunications links. Data center builds the 'customer-side' infrastructure. This can include on-site transformers, UPS/inverters, backup generators, and power distribution for chillers or cooling equipment.
Go Live	Power provider tests protection systems, verifies equipment ratings, and confirms telecommunication links for supervisory control and data acquisition (SCADA) and metering. Data center demonstrates compliance with power quality standards (voltage, harmonics) and confirms that backup systems work properly (i.e., will not shock the grid with sudden load changes). ISO/RTO issues final approval to energize the facility.

While FAC-002-4 does not list explicit requirements for data centers and other large loads, it requires ISOs and transmission organizations to study the impacts of a large load under steady state, short-circuit, and dynamic conditions.[12] This helps ensure that the bulk power system can perform sufficiently under the added demand of a large load, while also considering the ability of both the bulk system and large load to respond to and recover from certain disruptions without further destabilizing the grid. These studies typically require the data center to provide specific information on planned BTM assets and their configuration. For example, a data center implementing a BESS can be required to provide dynamic models for their inverter-based resources so that grid operators can properly assess how these backup resources will behave during a fault scenario.[15] Likewise, data centers utilizing backup diesel generators may be required to provide the specifications of the generator's automatic voltage regulator for grid operators to assess how voltage and reactive power is controlled under different contingencies, though these requirements vary depending on individual or combined generation capacity.[16] The most explicit requirements generally deal with power quality and ride-through capabilities.

Power quality refers to how consistent, and steadily power flows between the grid and a data center's BTM assets. Power quality is typically measured through harmonics, and limits are placed on the level of distortion produced by interconnected facilities.[17] Regional authorities like ISO-NE, Midcontinent ISO (MISO), PJM and vertically-integrated utilities like Dominion Energy follow IEEE 519, which requires that a large load's BTM assets remain below acceptable distortion levels.[18] Many of these power authorities have also incorporated specific characteristics of data centers into their interconnection requirements, emphasizing that impact studies include power quality assessments for UPS assets and non-linear connections such as chillers and temperature control systems.[19,20,21] Furthermore, power authorities like California Independent System Operator (CAISO) are implementing increasingly stringent power quality standards to help ensure grid stability in areas where traditional turbine-based generators providing natural inertia for the grid are being phased out or highly supplemented by inverter-based wind and solar generation sources.[22]

[*] NERC Industry Recommendations on large load interconnections specifically calls out transmission planners and planning coordinators, stating that these entities "should establish a comprehensive interconnection and system-wide study process using steady-state, dynamic, and short-circuit models to assess reliability impacts of Large Loads." (NERC 2025)

[†] The Energy Reliability Council of Texas (ERCOT) has pointed out that 'new large loads' (mainly data centers) typically have much shorter timelines to interconnect compared to traditional large loads, which has led to power providers seeking alternatives to modeling load impact to the grid. This means conducting studies that have historically taken a year or longer in just a matter of months (ERCOT Large Load Interconnection Process: Current Status, Lessons Learned, and Future Challenges)

[‡] Many ISOs/RTOs are in the process of reevaluating traditional interconnection procedures. The CAISO, for example, has traditionally used its Generator Interconnection and Deliverability Allocation Procedures (GIDAP) to manage large load interconnect compared that include on-site generation. However, in light of increasingly common data center projects, CAISO is working to enhance this process to align more closely with data centers (CAISO, 2023).

Ride-through refers to the ability of a data center's BTM assets to remain connected to the grid for a period of time during and following any grid disturbances leading to voltage or frequency imbalances.[23] While no national standard exists, ISOs and TOs across the country are increasingly implementing region-specific requirements that exceed ITIC thresholds by which large loads have traditionally based ride-through. For example, ERCOT requires that large loads remain connected to the grid at voltages above 90%, at which point operations can be curtailed if voltage drops to the 80–90% range or completely disconnected if voltage drops further than 80%.[23]

In many cases, data centers can help speed up interconnection processes by knowing ahead of time the performance requirements of specific ISOs and TOs in their region and preparing necessary data points that will be required for mandatory grid impact studies.

Key Insights

Large load interconnections remain a multistep, study-driven process, but data centers have shifted the planning paradigm. Regions now face saturation, prompting policy innovation, stronger reliability guidance, and a potential federal framework to standardize procedures for ≥20 MW loads. In the absence of specific Federal Energy Regulatory Commission (FERC) rules, data centers must navigate region specific manuals and utility standards while designing for stringent power quality, protection, and reliability requirements. Thus, when (or even prior to) entering into the interconnection process, data-center developers should be prepared to answer the following questions:

1. Does expected power consumption and/or on-site power generation capacity fall within the limits of what the regional power authority considers a large load?
2. What data points will be expected to provide for the power authority to carry out mandatory impact studies?
3. Does the power authority offer options to facilitate a quicker or more agile interconnection process?
4. How can key reliability requirements for power quality and ride-through during periods of grid instability be anticipated and addressed in advance?
5. Are BTM assets going to be configured in a way that might allow supporting the grid during periods of instability?

Interconnection process timelines often do not align with the speed at which new data centers expect to be energized and fully operational. Nevertheless, opportunities exist for accelerating these timelines, and data centers can facilitate this through a comprehensive understanding of the interconnection process, their expectations, and the potential incentives they can provide.

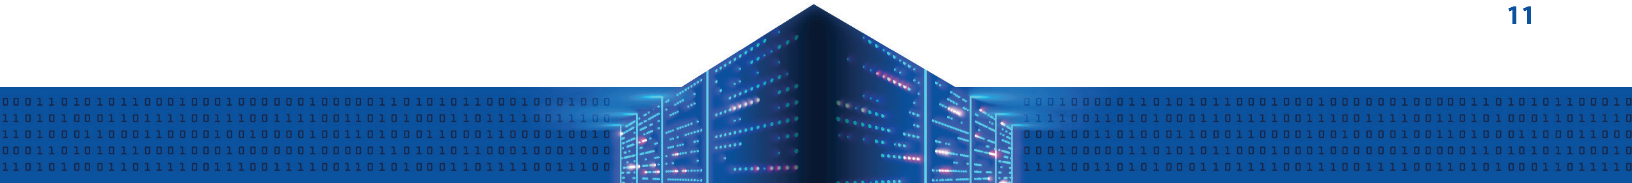

1.4 Grid Stability and Reliability Requirements

Introduction

As large data centers and AI campuses grow, they increasingly shape the real-time conditions under which the electric grid will operate. Stability and reliability, once dominated by predictable load shapes and synchronous generator behavior, are now influenced by highly dynamic, power electronics-rich facilities that can change consumption patterns in seconds.

Why Does This Matter for Data Centers?

Grid operators must maintain frequency, voltage, and system damping across wide geographic regions, even as these loads perform fast ramps, rapid transfers to backup power, and temporary islanding. For data-center developers and nuclear-generation partners, understanding these operational realities and regulatory requirements is crucial to design facilities that not only withstand grid disturbances, but also avoid exacerbating them. This section highlights the core stability challenges and emerging solutions identified across recent INL research, with the goal of equipping both operators and developers with actionable practices.

Research and Findings

Frequency Control Under Rapid Load Movement

Grid-frequency stability depends on a delicate balance between generation and load at every moment. Traditionally, this balance was supported by synchronous inertia and relatively slow-moving industrial and commercial loads. IT equipment, cooling systems, and UPS interactions in large, modern data centers can cause sudden changes in demand due to rapid downturns during protective transfers or steep increases during staged restoration or AI-workload shifts. Events involving load losses of hundreds of megawatts within seconds have demonstrated that frequency can deviate quickly, outpacing conventional primary frequency control.[9]

INL's analyses emphasize that frequency disturbances are no longer driven solely by generator trips; fast load changes are equally capable of producing significant deviations. As inertia declines with the retirement of synchronous generation, the grid becomes more sensitive to these changes, making fast frequency response (FFR) a critical operational requirement.[2] Industrial-scale batteries, utility-scale energy storage, and responsive industrial loads can all provide this service, but their activation must be coordinated. For data-center operators, preparing to participate in FFR or other ancillary services can help mitigate their own impact on system frequency.

Moreover, nuclear developers exploring data center co-location should understand that while nuclear provides steady power, its governor response is slower than inverter-based technologies. Pairing nuclear generation with fast-acting storage or automated control schemes can ensure the combined system enhances, rather than burdens, overall frequency resilience.[10]

Voltage Performance and Reactive Power Management

Voltage-stability challenges increasingly arise in regions with large clusters of loads driven by power electronics. Data centers produce unique voltage signatures: rapid transitions between grid-connected and islanded modes, sudden shifts in reactive-power consumption during cooling-system behavior, and protective responses that can disconnect large blocks of load.[2] These behaviors can sharply influence local voltage profiles and stress reactive-compensation systems, particularly in corridors already operating near their reactive limits.

Typical voltage-control devices—capacitor banks, transformer tap changers, and legacy Volt-amp reactive (VAR) schedules—are too slow to counteract the sub-second swings associated with large-campus behavior. Operators increasingly require dynamic VAR support, such as static synchronous compensators, capable of responding nearly instantaneously to electronic-load characteristics.[24] At the same time, data centers can reduce voltage risk by providing visibility into their cooling-system sensitivities and coordinating to maintain adequate voltage ride-through in plant controls.

Nuclear developers, especially those pursuing microreactor-based campuses, must also consider how local voltage regulation interacts with protective relays and inverter-based auxiliary systems. Validation environments and pilots, along with technical databases and model libraries for modern data center loads are essential technical needs to perform stability studies and support targeted mitigations.[25,26]

Operational Visibility and Coordination

Grid reliability hinges not only on control hardware but also on situational awareness. Operators frequently lack the detailed, near-real-time load data needed to make accurate decisions. Data centers may vary consumption significantly depending on workload type, cooling demand, and protective system behavior, yet operators often see only aggregated, slow-updating telemetry.

Improved visibility, including sub-second data at the POI, short-term load forecasts, advanced metering, and secure data-sharing protocols, allows operators to anticipate ramps, prepare reserves, and adjust voltage and frequency setpoints proactively. Data centers that provide this transparency can reduce the risk of over-procurement of reserves, unanticipated disturbances, and costly misoperations. For developers, this transparency may unlock preferential interconnection timelines or improved tariff structures.

Recent work also stresses the need for structured coordination frameworks that allow utilities, data centers, and nuclear developers to jointly establish operating envelopes, disturbance-response expectations, and restoration procedures.[10] Without this shared understanding, even well-designed systems may underperform in real time simply because assumptions differ between parties.

Key Insights

1. Seconds-scale load ramps demand purpose-built FFR stacks and governor/automatic generation control retuning in areas with clustered campuses.
2. Converter-driven and forced oscillations are an operational risk; POI-level monitoring plus screening of campus controls/filters is essential.
3. Sub-second telemetry and short-term load forecasts are prerequisite for stable operations, enabling smarter reserve and VAR provisioning.
4. Cross-stakeholder coordination templates accelerate adoption, turning research into repeatable operating practices across utilities and campuses.

Further Exploration

Near-term priorities include joint operator and developer pilots to validate FFR trigger logic, VAR dispatch driven by phasor-measurement units, and oscillation alarms in corridors with multiple campuses. These should publish datasets and parameter ranges that other balancing areas can reuse.

Parallel work should formalize secure telemetry and forecast exchanges and codify stability of operating envelopes (frequency and voltage) tailored to electronics-dense regions, so dispatch systems, VAR controls, and campus-workload schedulers operate to common limits. National laboratories can host templates and parameter guidance derived from these pilots for reuse across utilities and developers.

A final of concern highlighted in recent research is the increased risk of forced or poorly damped oscillations.[1] These oscillations may be subtle under normal conditions but grow under certain system configurations, particularly in areas with weak grids, long transmission lines, or series-compensated corridors. Utilities seek better understanding of oscillatory behaviors so that they can detect and mitigate these risks to stability.

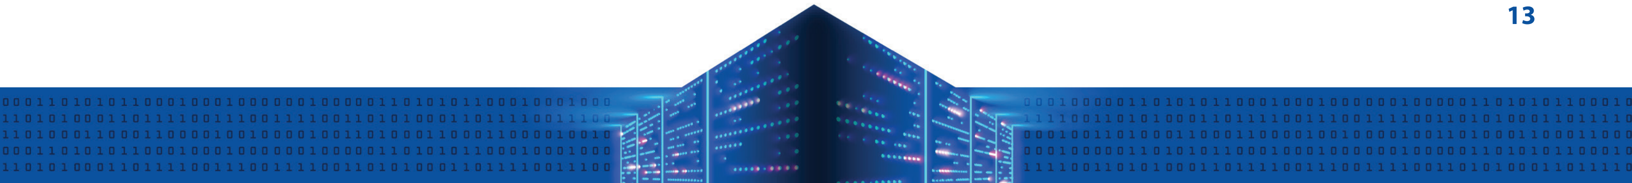

1.5 Utility Planning Considerations

Introduction

Data centers are now among the fastest-growing categories of load in North America, with dozens of regions experiencing simultaneous requests for campuses in the range of hundreds-of-MW to GW. This scale of demand reshapes utility planning because the grid was not historically designed for rapid, multi-GW increases in localized load. Growth in the U.S. has been relatively flat for the last 20 years, but is now seeing increases in forecasts at an ever-increasing pace.[27] Even well-resourced utilities face multi-year timelines to build new generation, reinforce transmission, or expand substations. At the same time, hyperscale developers often request service within 18–36 months, far quicker than legacy grid expansion cycles.[10]

Why Does This Matter for Data Centers?

These competing timeframes create friction across the sector, making it crucial for developers, operators, and policymakers to understand how utilities plan, approve, and recover costs for major infrastructure upgrades.

Research and Findings

The Strain of Large Load Growth and Forecasting Uncertainty

Traditional forecasting tools, based on population, economic trends, and incremental industrial growth, struggle to capture the surge in large electronic demand. Developers often evaluate multiple states or utility territories simultaneously, which leads to concerns over phantom load, where requests are submitted but do not materialize. Phantom load requests still consume resources to start evaluating proposals and run studies, causing some states like Texas and Ohio to enact rules requiring proof of site control, minimum contract lengths, and study fees to deter speculative queue entries.[28] Utilities must plan for the possibility that these loads will materialize because under-building carries major reliability risks, but over-building exposes ratepayers to stranded costs. This dynamic creates a structural tension; hyperscale developers want optionality and speed while utilities need certainty and long-term commitments before launching billion-dollar upgrades.

Limitations on Fast Grid Expansion

Even when utilities agree that new assets are necessary, the path from need to energization is inherently long because multiple sequential processes must align across planning, regulation, permitting, and delivery. Most investor-owned and public utilities operate within established integrated-resource-planning cycles that refresh major generation and transmission decisions every 2–3 years.[25] This cadence means large requests that arrive midcycle often wait until the next planning update or proceed only with interim measures that do not fully resolve capacity or corridor constraints.

Cost recovery adds a second layer of timing. Significant grid projects typically require rate-case approval from public-utility commissions to ensure customers pay only for prudent, necessary investments and to assign those costs in a manner consistent with "cost-causation" principles. Those proceedings involve testimony, discovery, and public input, routinely stretching 12–24 months, even under favorable conditions. States responding to rapid large-load growth have begun to formalize how big customers share costs; Ohio's tariff settlement for >25 MW loads and Georgia's public service commission rules on full cost responsibility are examples. Nevertheless, those decisions still move on regulatory timescales.[28]

On the technical side, lack of accurate dynamic models and detailed power-consumption data for large loads has been noted as an area of concern. Accurate simulations to determine the potential impact of a large load and apply appropriate mitigations in the interconnection-approval process are difficult when the models do not reflect the real-world behavior of those loads under non-steady-state conditions. Building these models to assess both the impact of large loads on the grid and the potential ways that large load dynamic capabilities can be used to support grid stability is an active area of INL research.

Permitting and siting are equally determinative. New transmission lines, high-voltage substations, or large generation assets must clear federal and state impact reviews, secure rights-of-way, and coordinate across jurisdictions—steps that can add years to a project in complex corridors. Recent federal executive actions have attempted to compress these durations by calling for modernization of permitting and prioritizing firm, dispatchable resources to meet surging AI and data center driven demand, but the practical acceleration is only beginning to manifest.[28]

Finally, generation and transmission expansion faces supply-chain and construction realities. Lead times for large power transformers frequently span 18–36 months, specialized labor is scarce in several regions, and outages must be coordinated to avoid compromising day-to-day reliability. Even with alignment across all steps, end-to-end timelines for substantial grid reinforcement or new generation commonly run 5–10 years. Until new procedures are adopted to speed approval timelines and secure resources necessary for infrastructure and labor costs, data-center developers continue to look for alternatives, such as bring-your-own power models and direct partnerships with generators in the form of power-purchase agreements (PPAs) or onsite construction.

Emerging Policy Mechanisms and Regulatory Reforms

Recognizing the mismatch between grid timelines and large-load expectations, regulators and system operators are deploying targeted mechanisms to add speed to power while preserving reliability and equity. At the federal level, FERC's approvals and directives have begun to delineate fast-track options for large loads and co-located generation.[29] In the Southwest Power Pool, the high impact large loads framework introduces a 90-day study path and conditional service models that can energize sooner while reserving the right to curtail temporarily to protect the system; this approach creates near-term service with transparent tradeoffs, rather than indefinite queue delays. Pennsylvania-New Jersey-Maryland (PJM) is also working toward options to bring large loads online faster, including provisions around onsite generation and transmission-service options.

MISO is exploring faster studies, interim service constructs, and demand-side participation through their Expediated Resource Addition Study. States are moving in parallel. Several states are adopting laws that attempt to balance the goals of protecting ratepayers and keeping energy costs down while also expediting permitting for energy infrastructure and maintaining regulatory oversight. The guiding principle across these reforms is not to bypass reliability, but to stage capacity additions, using conditional service, curtailed non-firm arrangements, and coordinated onsite resources while long-lead assets proceed through their normal cycles.

Cultural Mismatch: Utility Timelines vs. Data Center Expectations

Utilities plan and steward infrastructure on multi-decadal horizons, under mandates to maintain reliability and affordability for all customers. Their processes—integrated resource planning (IRP), rate cases, siting reviews, and cost allocation—exist to ensure prudence and fairness, and they move at institutional speed. Hyperscale developers, by contrast, operate on competitive timeframes measured in quarters: siting, negotiating, and commissioning campuses in 18–36 months to meet cloud, AI, or enterprise demand. Nuclear developers bring a third perspective to the conversation, as both an attractive option to serve large amounts of steady power for hyperscale data centers and background as a safety and regulation focused industry that can take years to reach approvals. This difference in cadence is not merely procedural; it reflects distinct risk cultures.[10]

INL research and engagement with industry stakeholders has identified that misaligned expectations produce reactive planning and strained operations, particularly when multiple large campuses cluster on constrained corridors. The most-promising remedies are hybrid approaches that meet in the middle: conditional or staged service paired with clear curtailment frameworks; co-located generation and storage that shoulder part of the adequacy burden; and pre-permitted, ready-to-serve substation zones in utilities' long-range plans that can accept large load blocks with reduced lead time. These approaches are gaining traction as RTOs and state reforms clarify cost responsibility, study requirements, and interim service options.

Key Insights

1. Large loads stress traditional planning paradigms by overwhelming forecasting tools and outpacing available generation and transmission capacity.
2. Infrastructure cannot scale quickly due to long regulatory, permitting, cost-recovery, and construction timelines—often stretching 5–10 years.
3. New policy mechanisms are emerging—including conditional service, co-located generation rules, and state-level large-load tariffs—to accelerate speed-to-power.
4. Cost causation and reliability protection are central themes, with regulators increasingly requiring long-term contracts, minimum billing, and developer-funded upgrades.

The cultural mismatch between utilities and data-center developers remains a core barrier, making collaboration, transparency, and early engagement essential.

Further Exploration

Future work should focus on coordinated planning frameworks that bring utilities, developers, and regulators into earlier, more-transparent dialogue. Pilot programs should test new large-load tariff structures, conditional service arrangements, hybrid supply architectures (e.g., nuclear + storage + grid), and accelerated permitting pathways. Another area of exploration is expanding load-flexibility programs that align campus operations with system needs without compromising uptime. National laboratories, RTOs, and state commissions will be key partners in evaluating how these mechanisms affect reliability, cost, and customer lifetime value.

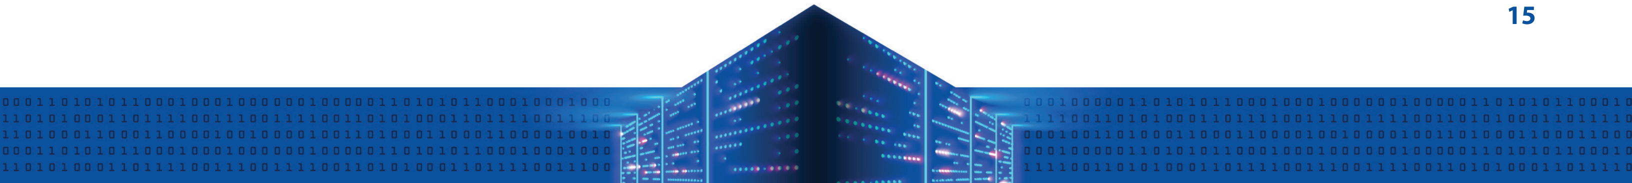

1.6 Capacity Expansion and Market Studies

Introduction

The surge in hyperscale data-center demand reflects a structural shift toward large, modular campuses optimized for AI and high-performance computing workloads. These firms accelerate deployment through repeatable 50–60 MW building blocks, long-term PPAs, repeatable procurement strategies, and vertically integrated infrastructure models while investment capital has flowed aggressively into the sector.[30]

Why Does This Matter for Data Centers?

Hyperscale data centers are fundamentally reshaping U.S. capacity expansion and market planning by driving one of the most concentrated and structurally significant load increases in power-system history. Planning challenges now center, not just on how much capacity is needed, but on where, how quickly, and whether transmission infrastructure can physically deliver under both normal and stressful conditions. With medium- to high-likelihood projects pushing projected national data-center demand beyond 300 GW by the early 2030s, heavily clustered in Texas, Virginia, and select high-growth states, the buildout is both geographically concentrated and front-loaded, far exceeding historical commercial load growth and outpacing the capacity additions and transmission upgrades considered by U.S. grid entities.[31]

Research and Findings

Navigating U.S. Market Structure for Large-Load Deployment

Electricity procurement and pricing vary by region, shaped by whether markets are regulated (cost-of-service) or competitive (RTO/ISO-based), which directly affects data-center contracting options, price risk, and speed to power.[32] About one-third of the U.S. load is served by vertically integrated utilities with public utility commission (PUC)-approved, cost-based rates that provide stable cost recovery while roughly two-thirds operates in competitive wholesale markets, where prices are set through supply and demand dynamics and locational marginal pricing, often alongside capacity and ancillary-service markets. Retail deregulation does not always align with wholesale-market participation, creating a spectrum of hybrid structures across the U.S.

Regional Adequacy Challenges

The 2025 long-term reliability assessment (LTRA) shows a dramatic acceleration in projected peak demand, far beyond historical forecasts. Over the 10-year assessment period (2026–2035), aggregated summer peak demand is forecast to rise by more than 224 GW while winter peak is projected to increase by over 245 GW. Large commercial and industrial loads, especially new data centers supporting AI and the digital economy, account for most of the projected increase in North American electricity demand over the next decade.[33] Load forecasts collected for the 2025 LTRA reveal a substantial buildout of data centers, particularly in Texas, PJM, and the Western Electricity Coordinating Council (WECC) assessment areas. In some WECC balancing authorities, planned data centers account for as much as 40% of the demand forecast.

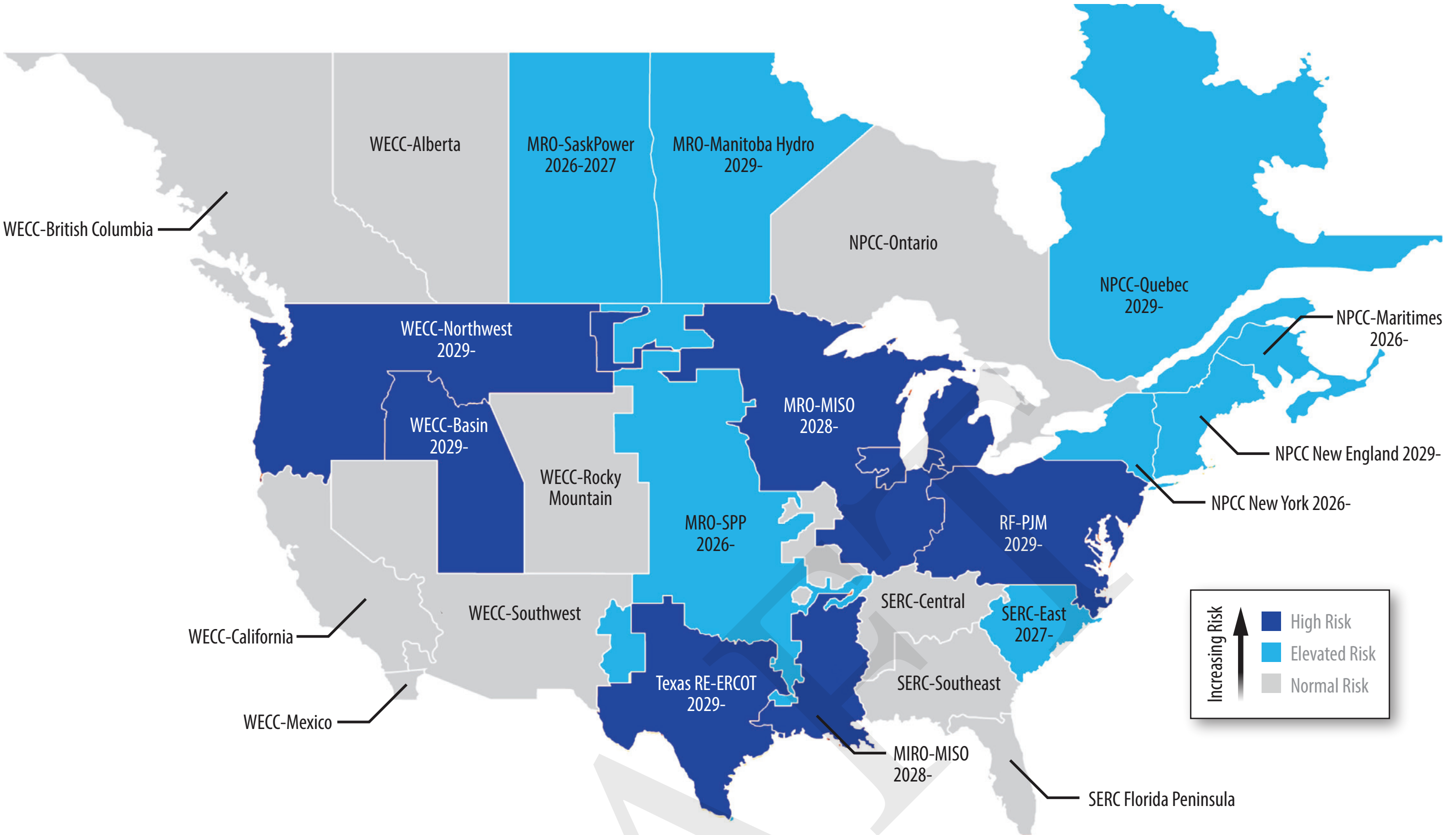

Figure 5. NERC 2025 LTRA regional risk outlook (2026–2030). High risk (dark blue), elevated (bright blue), and normal (gray) resource adequacy risk by region, reflecting tightening margins driven by load growth, retirements, and fuel constraints.[33] High-risk regions are projected to exceed resource adequacy limits with expected shortfalls, Elevated-risk regions meet targets under normal conditions, but are vulnerable under stress, and normal-risk regions remain within thresholds.

High-growth regions—ERCOT, PJM, MISO, and parts of WECC—will experience the strong demand increases.[33] In several regions, planning-reserve margins trend downward through the late 2020s and early 2030s absent major incremental additions and transmission expansion (see Figure 5). Other regions (e.g., the Southwest Power Pool [SPP] and parts of the Southeastern Electric Reliability Corporation [SERC]) may maintain comparatively higher margins in the near term although seasonal and fuel risk-driven shortfall concerns still appear under stressed conditions.

Strategic Load Clustering Across Electricity Markets and Utilities

NERC divides the U.S. into multiple reliability regions, and current data-center capacity is highly concentrated in a few key ISO/RTO markets, rather than evenly distributed nationwide. As shown in Figure 6, PJM leads by a wide margin, with roughly 16–17 GW of connected data-center load, representing the highest share of total regional demand (approaching ~17%). ERCOT and MISO follow with approximately 4–4.5 GW each, accounting for a growing share of their system load, while CAISO and SPP show moderate, but meaningful penetration.[31] In contrast, Independent System Operator-New England (ISO-NE) and New York Independent System Operator have comparatively low connected capacity and small percentages of total demand.

This distribution indicates that data centers are clustering in regions offering favorable market structures, transmission access, land availability, and scalable generation, particularly PJM and ERCOT, while higher-cost or more space-constrained regions in the Northeast show limited growth. Strategically, hyperscalers reinforce this concentration by securing long-term PPAs, investing in utility-infrastructure upgrades, and prioritizing interconnection speed and grid reliability, further reshaping regional load profiles and utility planning dynamics.

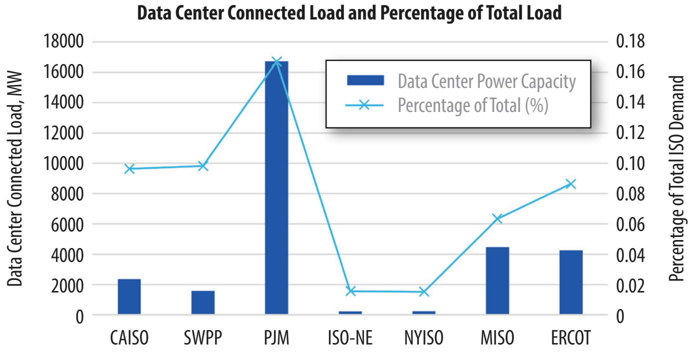

Figure 6. Data-center capacity in MW and percent-share across different ISOs.

Deploying Nuclear Uprates as a Near-Term Bridge

Meeting near-term supply–demand gaps require solutions that align with the rapid pace of large-load growth, particularly AI-driven data centers that can be built in 18–24 months, far faster than new generation and transmission infrastructure. Nuclear uprates offer a practical bridge because they leverage existing, licensed plants and grid interconnection points, enabling incremental additions of firm capacity on shorter timelines than new projects. This makes uprates one of the few scalable near-term options capable of strengthening reserve margins while longer-lead infrastructure is developed. However, their ultimate value depends on deliverability: even where data center clusters are geographically close to nuclear facilities that can be uprated, such as in the Mid-Atlantic, transmission congestion, interconnection backlogs, and contractual structuring determine whether incremental output can reliably reach high-growth load centers.

The geographic distribution of data-center development across the U.S. is highly uneven and concentrated within specific RTO footprints. Active facilities, under-construction projects, and announced developments are predominantly clustered in PJM (particularly Northern Virginia), ERCOT (Texas), and portions of the Southeast Virginia and the Carolinas (VACAR) and Southeastern regions). Secondary concentrations appear in the Midwest (MISO/Midwest Regional Operator [MRO]) and parts of the West (Northwest Power Pool and California-Mexico [CA-MX]), though at comparatively lower density.[34]

The most-pronounced clustering (Figure 7) occurs in Northern Virginia and the broader Mid-Atlantic region, where the highest density of active and upcoming data centers coincides with multiple nuclear units in Virginia, Maryland, and Pennsylvania. This region also exhibits some of the largest projected data-center demand growth (exceeding 20–25 GWe in the PJM footprint), while retaining measurable uprate potential across existing reactors. The spatial overlap suggests theoretical opportunities for structured long-term offtake arrangements. However, the same region is characterized by heavy transmission usage, queue backlogs, and deliverability constraints. Therefore, physical proximity alone does not guarantee feasible interconnection; transmission topology, congestion patterns, and contractual structure will be decisive.

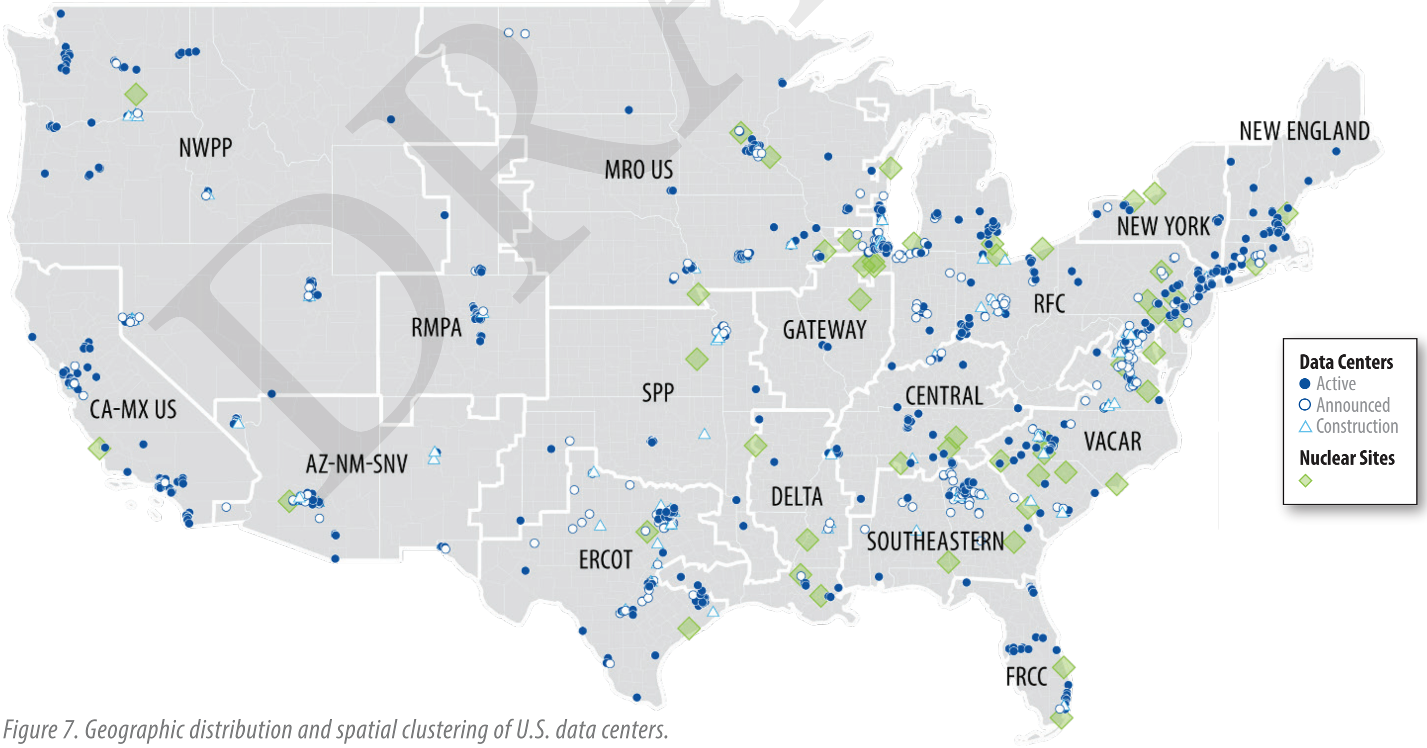

Figure 7. Geographic distribution and spatial clustering of U.S. data centers.

Understanding the Deliverability Bottleneck

Transmission and interconnection bottlenecks emerge as the binding constraint on large-load integration. In regions like PJM, where data center concentration is highest, interconnection timelines have stretched from less than 2 years to more than 8, creating a structural mismatch between the 18–24 month build cycle of data centers and the decade-long process required to permit and construct generation and transmission infrastructure. Even where incremental nuclear uprates are technically feasible, their economic value depends on deliverability and market access, which are increasingly constrained by queue backlogs and limited transfer capability. Although NERC projects a substantial increase in planned transmission mileage over the next decade,[33] much of it remains early-stage and faces siting and cost-allocation hurdles. Broader assessments suggest that tens of gigawatts of additional interregional transfer capability could materially improve reliability under stress, reinforcing that today's adequacy challenge is as much about transmission as it is about local capacity shortfalls.

At a high level, transmission capability determines whether new supply or large loads can be reliably integrated. The 500 kilovolt (kV) + network serves as the bulk-power backbone for long-distance, high-volume transfers while 345 kV lines function as the regional workhorse system, connecting generation hubs to metropolitan demand centers. Medium- to lower-voltage sub transmission layers often create the decisive "last-mile" constraints that govern actual deliverability to concentrated load pockets, such as hyperscale data-center campuses. As a result, reliability and adequacy challenges are frequently driven by transfer capability and localized network limits rather than purely by aggregate megawatt shortages, underscoring the importance of coordinated transmission planning alongside generation additions.

Grid Issues with Large-Load Integration

The rapid and highly concentrated growth of hyperscale data centers is creating localized grid stress that regional averages often mask. In dense corridors, such as Northern Virginia and parts of ERCOT, multiple hundred-MW campuses can form load pockets where transmission outages, substation constraints, or interface congestion have outsized reliability implications. As a result, the integration challenge is frequently geographic and operational rather than purely system-wide adequacy.

Reliability risks are bidirectional. Rapid load additions tighten reserve margins and strain interconnection pathways, but sudden loss of very large load blocks can also create operational disturbances. A recent Eastern Interconnection event involving the loss of roughly 1,500 MW of load during a transmission fault illustrates how large-load behavior, particularly customer-side protection and ride-through settings, can contribute to frequency and voltage excursions.[36] Such events underscore that both load growth and load volatility must be incorporated into planning studies.

These dynamics elevate the importance of explicit voltage ride-through standards, dynamic load modeling, fault sensitivity analysis, monitoring of coincident load loss, and controlled reconnection-ramp protocols. Integration is further constrained not only by transmission capacity, but also by infrastructures such as transformer ratings, reactive-power capability, and substation topology. In high-growth regions where interconnection timelines extend far beyond data center construction schedules, congestion and transfer capability, rather than aggregate MWs, are increasingly the decisive factors in reliable integration.

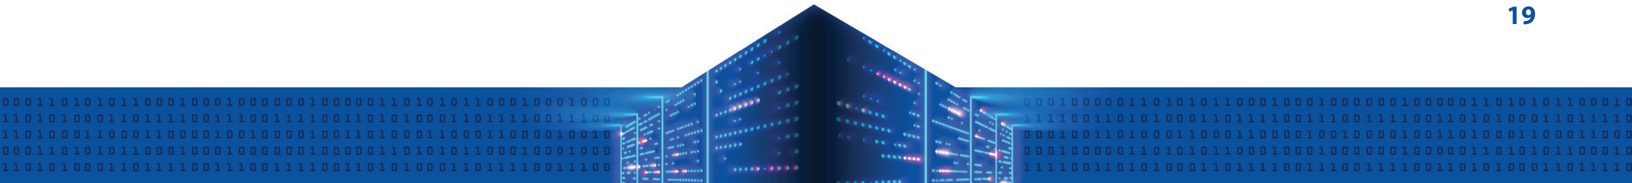

Key Insights

The analysis of data center deployment, market structure, regional adequacy, and transmission constraints points to several structural shifts in how capacity expansion must be evaluated. These insights frame how nuclear uprates and other firm resources should be positioned within rapidly evolving load centers.

1. Data centers have become a system-level driver of capacity expansion, with hyperscale and AI loads accelerating demand in concentrated regions (notably PJM and ERCOT)
2. Transmission and interconnection constraints—not just generation—are now the binding reliability limits, with queue delays and congestion determining deliverability
3. There is a structural mismatch between fast load build cycles (18–24 months) and slow infrastructure timelines (up to a decade)
4. Nuclear uprates offer a rare near-term, firm capacity option using existing assets, but their value depends on transmission and market alignment
5. Physical proximity between data centers and generation does not guarantee deliverability; sub-regional and backbone transfer limits are decisive
6. Market structure (regulated vs. competitive) materially shapes contracting pathways, risk exposure, and speed-to-power.

Strategic Priorities for Investors, Utilities and Policymakers

Given these structural dynamics, capacity expansion and offtake strategies should be evaluated through a focused set of screening questions. These are intended to guide siting decisions, market studies, and nuclear uprate alignment assessments:

- Where is the binding constraint: generation adequacy, regional congestion, interregional transfer limits, or last-mile infrastructure?
- Can incremental nuclear output be delivered to the target load zone without major transmission reinforcements?
- Do interconnection timelines realistically align with data center commissioning schedules?
- What transmission upgrades are required, and are they permitted, funded, and sequenced appropriately?
- Does the regional market structure support long-term firm offtake arrangements?
- How exposed is the project to queue backlogs, cost-allocation disputes, or siting delays?
- Are large-load operational risks explicitly modeled and contractually addressed?
- Is the opportunity addressing a local reserve margin issue, a deliverability constraint, or an interregional transfer limitation?

Chapter 2

POWER BEYOND THE GRID, INTEGRATED SYSTEMS, MICROGRIDS

For data-center developers navigating the realities of grid interconnection, a broader question is increasingly being asked: what options exist beyond, or even alongside, the grid? This chapter explores the growing range of options for generating and managing power outside of (or in combination with) the traditional utility relationship. Nuclear energy figures prominently in that picture, but the path to nuclear integration is shaped by a set of technical, economic, and operational choices that deserve careful examination on their own terms, and will be explored in more depth in Chapter 4.

This chapter opens with the question of how a nuclear power plant and a data center can actually be connected. Five distinct electrical coupling configurations are examined in the section on **Data Center Electric-Coupling Scenarios with Nuclear**, each carrying different implications for cost, reliability, grid dependence, and reactor sizing. From there, the discussion moves into thermal integration—specifically, the opportunity to use reactor heat, not just reactor electricity, to drive data center cooling. The section on **NPP/DC Thermal Coupling Through Absorption Chillers** explores how this approach can meaningfully expand the computing capacity a given reactor can support, and what tradeoffs are involved in different chiller configurations.

Water is then taken up as a resource that connects nuclear and data center operations in ways that are often underappreciated. The section on **Energy-Water Nexus** examines how both industries draw on shared water supplies, how that demand is regulated across different jurisdictions, and why water availability is increasingly a binding constraint on siting decisions. The section on **Controls** addresses the operational challenge of managing a complex microgrid (one that may include a reactor, battery storage, and variable loads) safely and reliably, and describes the testing frameworks being developed to validate these systems before deployment. Finally, the section on **Direct Current Power for Data Center Racks** examines an emerging shift in how power is distributed within data centers themselves, and why DC architectures may offer meaningful advantages as rack densities continue to rise.

2.1 Data Center Electric-Coupling Scenarios with Nuclear

Introduction

A data center requires energy in the form of electricity for its IT equipment: both computational and cooling, predominantly supplied by electric chillers that produce and circulate chilled water. Electricity can be provided directly by the nuclear power plant (NPP), either through a direct connection or through the grid. The different coupling options are driven by technical and economic considerations influenced by region-specific regulatory and grid constraints. This chapter summarizes those options, detailing some of their implications for nuclear projects and data centers.

Why Does This Matter for Data Centers?

The type of electrical connection considered will directly influence how much a data center relies on the electrical grid, how much nuclear-reactor capacity is needed, and the number of units or backup required to meet reliability requirements. Reducing data center reliance on the grid can potentially benefit long-term economics through reduced transmission costs and faster deployment due to current interconnection-queue backlogs and high rejection rates in some regions of the U.S.

Research and Findings

Five different options for electrical coupling of NPPs to data centers were defined, based on the type of grid connections to the coupled nuclear/data center system, as illustrated in Figure 8. Benefits and challenges of each option are summarized in Table 5-3 of [37], with highlights summarized here.

In **Option 1**, the NPP and data center are electrically interconnected through bulk-transmission lines only. Through a non-physical delivery PPA, the data-center owner can agree to purchase power from the NPP at set prices while physical settlement with the local utility or market operator proceeds separately. This option provides flexibility in terms of nuclear siting—because no proximity with the data center is required—and in terms of power matching. Although data-center capacity can be built in 1–3 years, which is faster than any firm source of energy generation or transmission expansion, NPP deployment need not match the data-center timeline provided the grid has sufficient reserve capacity. The two facilities can operate on independent schedules, with the NPP eventually contributing to peak data-center demand. The NPP and data center can maintain independent outage and operation schedules. This option (grid connection) benefits from the high capacity factors and firm baseload-power characteristic of nuclear generation while grid inertia contributes to high power quality for the data center. However, it requires sufficient existing or new transmission and distribution infrastructure, and because the data center would draw from the broader grid mix, it is not possible to track from which kind of source the electricity consumed originates.

Options 2, 3, 4, and 5 are different types of co-location or BTM configurations, all requiring some proximity between the NPP and data center to accommodate direct electrical connection and share a substation and its interconnections to the high- or extra-high-voltage transmission system. In Option 4, the NPP and data center are isolated from the grid and operating in an island mode. In Options 2 and 3, the NPP maintains a connection to the grid to sell excess electricity. In Options 2 and 5, the data center maintains a direct connection to the grid to serve as backup or complement to an NPP. The benefits and challenges of these options can be summarized as follows:

- **Option 2** connects both the NPP and data center to the grid, allowing the NPP to sell excess electricity and improving project economics while the data center retains grid backup for reliability. Transmission costs are minimized although proximity between facilities is needed, and load transients must be carefully managed to avoid impacts on the NPP or grid.
- **Option 3** similarly allows the NPP to sell excess power to the grid while minimizing transmission costs, but the data center has no direct grid connection. This places greater reliance on the NPP (potentially together with other energy infrastructure) to meet data-center power demand, necessitating additional onsite redundancy that adds to overall project costs.
- **Option 4** isolates both the NPP and data center from the grid entirely, eliminating transmission costs and enabling deployment in remote or off-grid locations. However, the absence of any grid support requires significant backup generation capacity and potential additional reactor units to meet data center reliability requirements, increasing overall costs.
- **Option 5** maintains a direct grid connection for the data center as backup while the NPP operates solely to serve the data center load without selling excess electricity to the grid. This improves reliability and can reduce electricity costs, though it foregoes potential NPP revenue streams and remains subject to regional regulatory constraints.

Within each coupling-path option, the power distribution path within the data center may be designed for varying degrees of availability by incorporating features for redundancy, maintainability without disruption, and fault tolerance. To meet reliability requirements, this distribution system may have separate redundant circuits and UPSs serving IT equipment, along with onsite generation capacity to meet or exceed the peak data-center power demand. The internal redundancy of the data center distribution system is designed independent of the reliability of any external power supplies, including the utility electricity supply and any generation that is not under direct control of the data-center operator.

The type of coupling option will affect significantly what type and number of nuclear reactors that may be needed to power a data center. This was illustrated in Stauff by applying the reactor sizing optimization framework developed from Hanna, with results summarized in Table 7.[37,38]

The analysis uses generic cost estimates for nuclear reactors of varying sizes, based on next commercial-offering cost overnight capital costs, operation

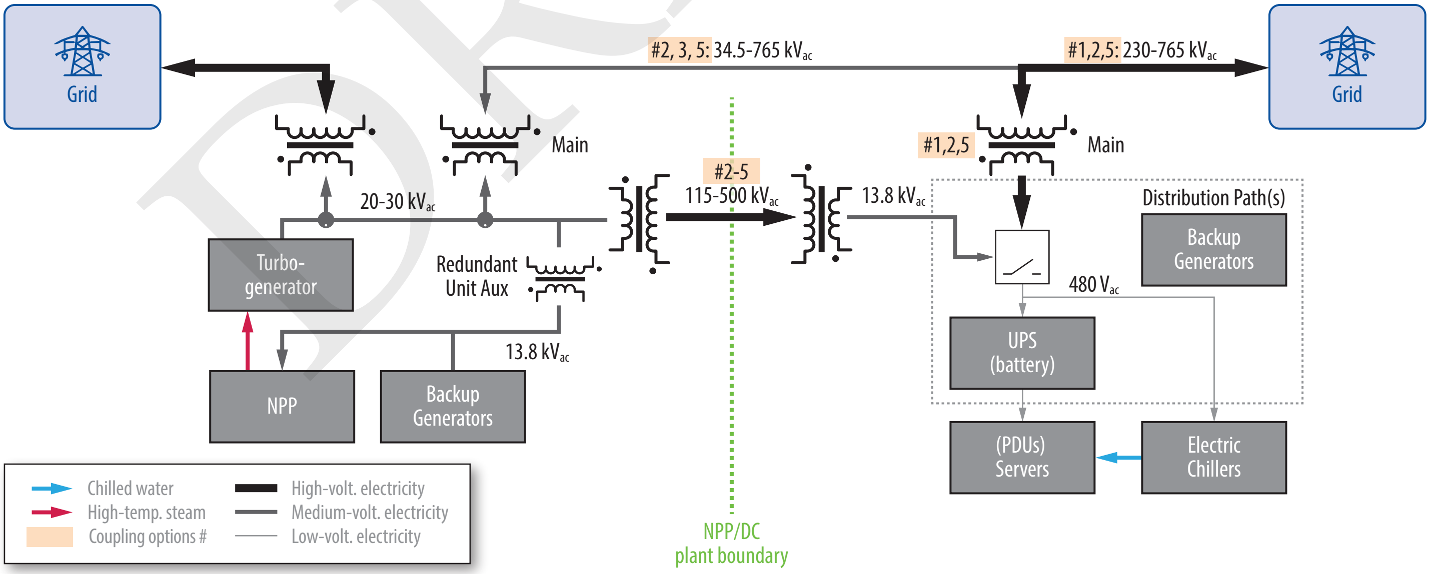

Figure 8. Illustration of different options for electrical connection of an NPP to a data center.

and maintenance (O&M) costs, and construction durations derived from curve-fitting of literature data. It is important to note that these cost estimates carry significant uncertainty, and the results should be interpreted in terms of relative trends rather than absolute cost predictions. Additionally, the analysis considers only the cost of power supplied by the nuclear plant itself and does not account for transmission and distribution costs or battery-storage costs that may be part of a real deployment. A fixed data-center demand is assumed, with load fluctuations expected to be managed on the data-center side.

The details of this methodology can be found in [38] together with the data used (costs, construction duration, learning rate [LR], refueling interval). Learning across multiple sites is not considered in this model, which focuses on a single site project. Depending on the configuration of the NPP and data center, total demand, economies of scale, and electricity price, the ideal reactor size varies for each application.

Key Insights

The main trends observed in the results show that larger reactor sizes will be preferred to meet higher data-center power demand while benefiting from economies of scale. In the absence of grid backup (as in Options 3 and 4), additional reactor units or other onsite power production will be needed to meet reliability requirements, leading to potentially smaller nuclear-reactor sizes. When the nuclear reactor is connected to the grid to sell excess electricity (as in Options 1, 2, and 5), higher electricity prices favor over-capacity builds. Additionally, higher interest rates have a more significant impact on larger reactors, owing to their longer construction times.

Further Exploration

Further details on the different connection options are provided in the System Analysis and Integration (SA&I) report on nuclear-powered data centers sponsored by DOE's Office of Nuclear Energy (NE).37 Table 5-4 of [37] report especially summarizes the main benefits and challenges associated with these different coupling options. Other relevant resources are the Electric Power Research Institute (EPRI) reports on nuclear-powered data centers39 and behind-the-grid operation.40 More details on the INL optimization framework for nuclear sizing to match data-center power demand are available in [37]. Ongoing effort by the SA&I campaign involves case-study analysis of different deployment scenarios for NPPs with data centers on specific sites. Other ongoing work from the Integrated Energy Systems program has explored the reliability and cost of grid-connected and islanded microgrid configurations that include nuclear power.

Table 7. Summary of reactor sizing optimization analysis for different NPP/DC coupling options. This table is intended to illustrate (1) how the optimum reactor size varies with data-center demand and coupling configuration and (2) the relative differences in levelized cost of electricity (LCOE) across options, considering only NPP costs.

DC Demand (MWe)	Option 1, 2, and 5			Option 3			Option 4		
	Through the grid (or direct connection with grid backup)			Direct connection to DC with option to sell excess to the grid			Direct connection to DC without connection to the grid		
	Optimum Reactor Size (MWe)	# Units	LCOE* ($/MWh)	Optimum Reactor Size (MWe)	# Units	LCOE ($/MWh)	Optimum Reactor Size (MWe)	# Units	LCOE ($/MWh)
50	50	1	147	10	6	166	5	11	174
100	50	2	129	20	6	145	10	11	153
500	500	1	97	100	6	108	50	11	116
1000	1000	1	87	500	3	95	50	21	103
2000	1000	2	**79**	1000	3	83	100	21	92
6000	1000	6	72	1000	7	73	400	16	77

* The LCOE values presented here do not represent the full cost of powering a data center with nuclear energy because they exclude transmission and distribution costs and other project-specific factors.

2.2 Economic Outlook of NPP/Data Center Thermal Coupling with Different Cooling Configurations

Introduction

Coupled nuclear data-center systems involve tightly integrated electrical and thermal subsystems that have their performance dependent on both infrastructure design and operational behavior. Power generation, cooling technology, thermal-energy storage (TES), and electrical storage interact in ways that influence system sizing, utilization, and cost. Cooling configuration is a primary driver of these interactions. Air-cooled systems rely primarily on electricity and minimize water use, but can increase electrical demand. Liquid-cooled systems improve heat removal efficiency, often reducing electrical demand while increasing water and infrastructure requirements. Absorption-cooled systems use thermal energy to drive cooling, partially substituting electrical demand with thermal input and introducing direct coupling between reactor-heat output and facility-cooling needs. Importantly, these cooling-driven electrical demands are distinct from the IT (compute) load, representing an additional and configuration-dependent component of total facility power demands. These differences in electricity, thermal power, and water requirements shape how generation and storage assets are deployed and utilized.

Data-center load profiles add another key dimension. While total demand is often predictable, temporal variability can differ significantly depending on workload type. High-performance computing and AI workloads, for example, may exhibit periods of high volatility, while others operate closer to steady baseload. This variability directly affects dispatch, storage utilization, and overall system performance.

Tools such as the Holistic Energy Resource Optimization Network (HERON) represent time-dependent, dispatch-based technoeconomic analysis (TEA) frameworks that capture these effects by co-optimizing system design and operation under time-varying conditions.[41] These methods co-optimize component sizing and operational dispatch over time, allowing system performance to be assessed under realistic, time-varying load conditions.[42,43,44] By resolving operation at an hourly (or similar) timescale, dispatch-based TEA captures key dynamics such as storage utilization, load-following behavior, and interactions between thermal and electrical subsystems.

Why Does This Matter for Data Centers?

Economic performance in coupled systems depends on how effectively infrastructure responds to time-varying demand and integrated thermal and electrical requirements. Static TEA approaches, based on averaged behavior, can overlook these dynamics—particularly when multiple storage systems and cooling strategies are involved. Time-dependent, dispatch-based TEA captures how variability in compute and facility loads drives the use of generation, battery storage, and TES. This is especially important when comparing cooling configurations because each imposes different temporal demand patterns and resource dependencies. Electrically driven cooling may increase sensitivity to electrical load fluctuations while absorption cooling can introduce additional flexibility—and constraints—through thermal coupling.

This framework enables clearer identification of system-level tradeoffs, including storage value, generation sizing, and cooling-strategy selection, all of which can vary with load variability and system design. It also occupies a practical middle ground: more representative than steady-state TEA, but far more tractable than detailed transient modeling. For data-center stakeholders, this supports early-stage decision making by enabling consistent comparisons across configurations, highlighting sensitivities to load variability, and identifying promising system designs for further analysis.

Research and Findings

INL's technoeconomic models investigated absorption-chiller performance with a pressurized water reactor (PWR)-type small modular reactor (SMR), comparing it with air-cooled and conventional water-cooled systems.[45] The study used the HERON tool to optimize nuclear-powered data-center configurations with a focus on reliability and economic performance. HERON conducts two optimizations simultaneously: an inner-loop optimization on hourly (or other timestep) resource utilization subject to technical constraints and load-time signals informs an outer-loop selection of component capacities to maximize economic performance. Loading scenarios were synthesized from 1 year of hourly data-center computing and facility load made available from the INL high-performance computing (HPC) for 2023. As shown in Figure 9, these synthesized time series represented a high- and low-volatility loading scenario for the data center to capture the range of possible operating data-center schemes. Cooling loads were handled separately, as a variable to be solved in the problem.

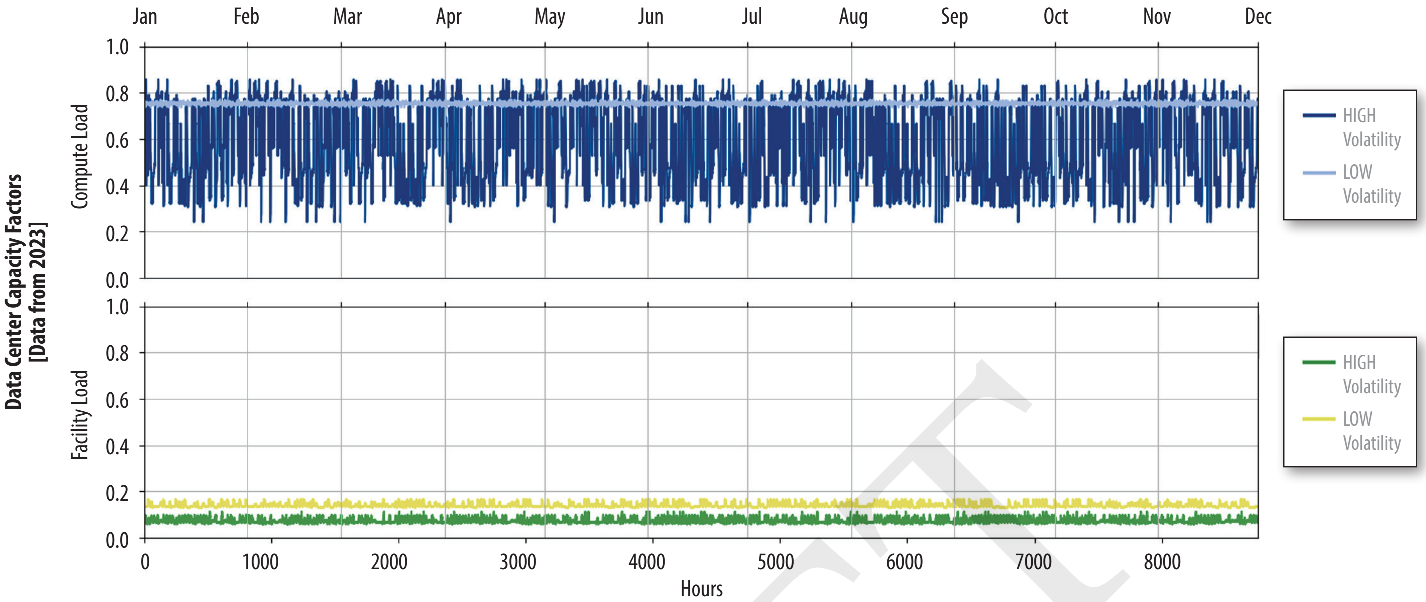

Figure 9. Normalized compute (top) and facility (bottom) electricity loads for simulated data center in hourly timesteps. A high and low variability signal is shown for each load type.

Microgrid conditions were simulated assuming a data center powered by a collection of 100 MWe SMR units with battery storage and TES. Different coupled configurations were modeled using air-cooled, conventional water-cooled, and absorption-cooled systems. Net present value (NPV) results were calculated for each configuration coupled to a 1 GW data center, as shown in Figure 10. Economic outcomes, based on generalized cost assumptions, should be interpreted as trends rather than finalized results. Each column represents a different cooling technology, and subplots display NPV (top row), required SMR units (middle row), and thermal capacity for TES and cooling (bottom row) as functions of battery capacity under high- and low-variability scenarios. Increasing battery capacity generally reduces costs, especially in high-volatility load scenarios, but diminishing returns are observed up to a certain battery size. For the low-volatility, liquid-chilled configuration, larger batteries reduce the need for one SMR unit, leading to lower costs. Across the configurations evaluated, the liquid-cooled system achieved the lowest costs under the given assumptions. It should be noted that water-related costs were not included because this analysis was not tied to a specific geographic location. Additionally, the absorption-cooled configuration was modeled using simplified, linear performance representations, which may not fully capture system behavior under all operating conditions. The following section explores absorption-chillers in more detail.

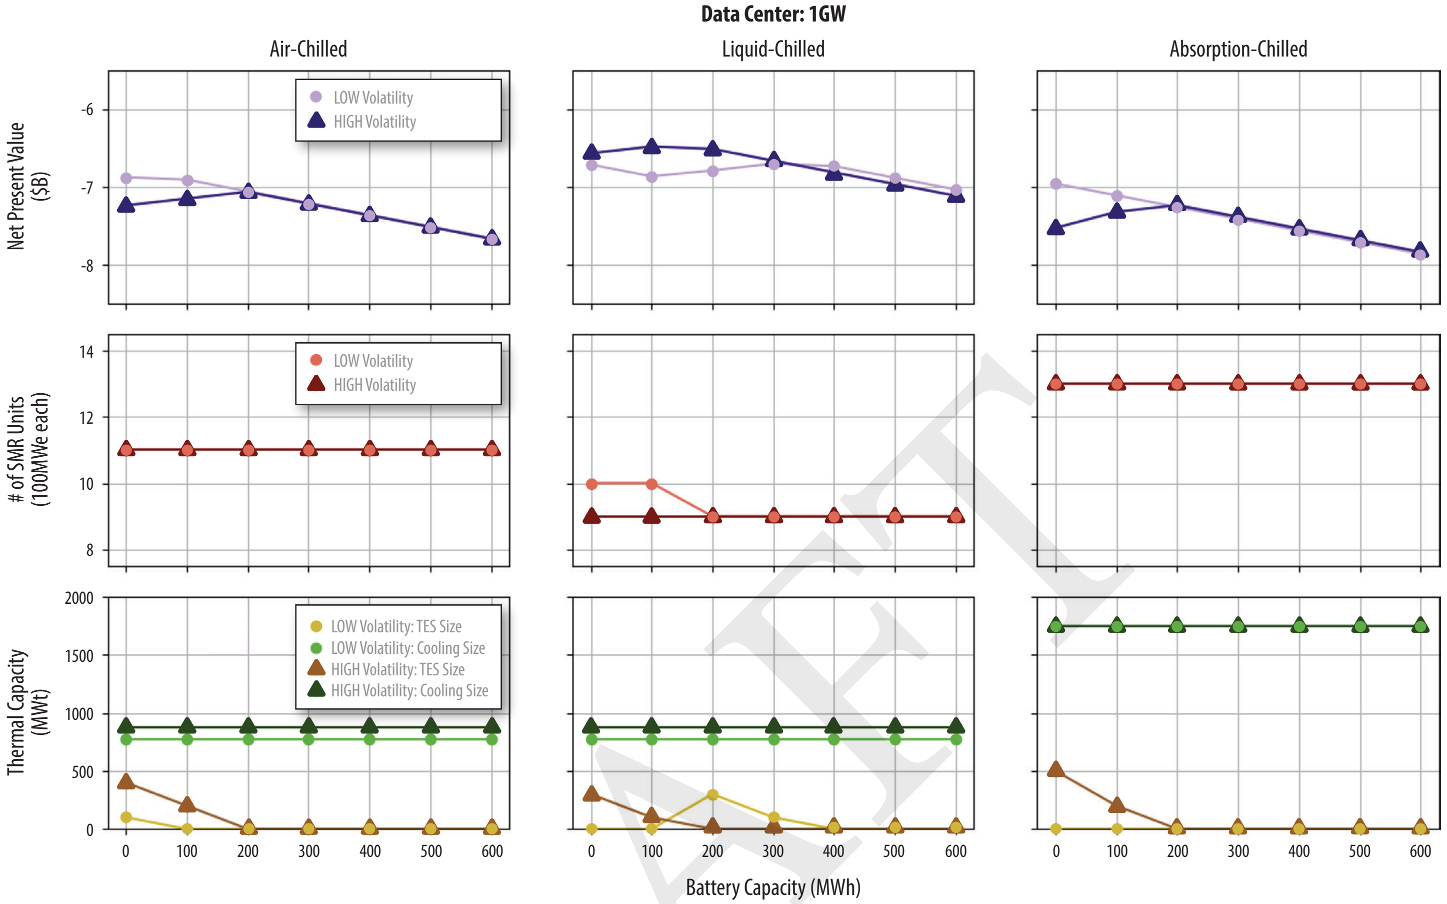

Figure 10. NPV shown as battery storage is increased for the 1-gw data center and the three cooling configurations. Corresponding number of SMRs and thermal capacity requirements also shown in the rows.

Key Insights

1. **Cooling configuration is a primary driver of system economics.** Across scenarios, liquid-cooled configurations achieved the lowest NPVs; however, water-related costs were not investigated, reflecting the balance between reduced electrical demand and manageable infrastructure requirements relative to air- and absorption-cooled systems.
2. **Battery storage provides the greatest value under high load variability.** Increasing battery capacity reduces system costs by smoothing temporal mismatches between generation and demand, with the strongest benefits observed in high-volatility load scenarios for the 1 GW data center. However, diminishing returns emerge beyond a certain storage capacity.
3. **Storage and cooling choices directly influence generation capacity needs.** Under certain scenarios, increased battery capacity can offset the need for additional SMR units, demonstrating the importance of co-optimizing storage and generation sizing rather than treating them independently.
4. **Thermal and electrical coupling introduces additional design tradeoffs.** Absorption-cooled systems create opportunities to use thermal energy but require careful integration and more-detailed modeling to fully capture potential efficiency gains, which were not fully represented in this analysis.
5. **Results are sensitive to assumptions and highlight the need for site-specific analysis.** Generalized cost inputs and simplified cooling models mean results should be interpreted as directional trends. Factors such as water costs, improved thermodynamic performance, and multi-year load data could materially shift outcomes.

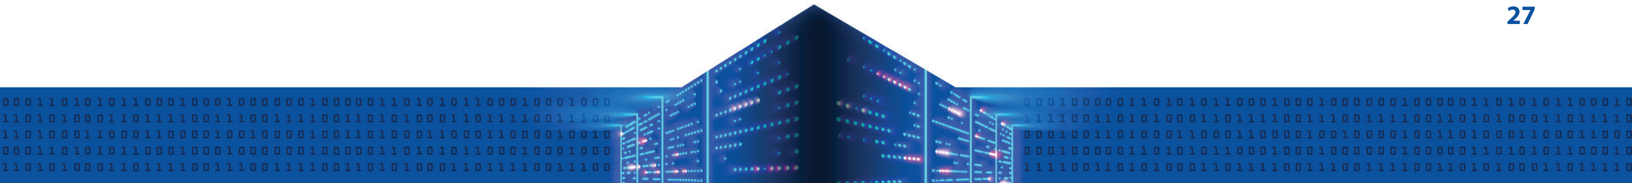

Further Exploration

Several extensions would improve fidelity and applicability to real-world deployments. Incorporating anonymized, site-specific data from existing data centers—particularly multi-year time series of IT load, facility demand, and regional weather—would better capture operational variability. Time series for thermal loads and water usage would also improve validation of results. Additional detail on cooling configurations, efficiency metrics, e.g., power usage effectiveness, and O&M costs would further refine the analysis.

Expanding the temporal scope to multiyear operation would enable evaluation of long-term economic performance under evolving conditions. Relaxing constraints such as fully islanded operation would also allow exploration of hybrid grid-connected strategies aligned with data-center requirements. Increasing temporal resolution below the hourly scale would enable modeling of faster-acting components (e.g., flywheels, backup generators) and additional operational constraints such as frequency and voltage response.

Finally, scenario-based analysis in HERON can be used to capture uncertainty in system performance. In the HERON framework, multiple realistic scenarios can be generated from historical data while preserving key temporal characteristics and evaluated in parallel.[41] Each scenario produces a corresponding economic outcome, resulting in a distribution of NPVs rather than a single value. This allows results to be summarized using statistics such as the mean or range and supports decision-making based on expected performance across a range of operating conditions. With sufficiently rich multi-year datasets, these models can be trained to reflect realistic system behavior, enabling more complex analyses while still producing simple, decision-relevant performance metrics.

2.3 NPP/DC Thermal Coupling Through Absorption Chillers

Introduction

Within a typical data center, cooling high-density server racks alone accounts for 30 to 50% of total site electricity consumption,[46] a load met almost exclusively by electrically driven vapor-compression chillers. Nuclear technology—large, SMRs or microreactors—offers a compelling case for powering data centers with continuous, reliable electricity with high (>90%) capacity factors.[47] Crucially, nuclear reactors also present an opportunity for thermal cogeneration.[48] By shifting the cooling burden from electricity to reactor heat—for example, by using absorption chillers rather than electrically driven ones—data-center operators can meaningfully improve energy utilization and reducing electrical-grid dependence.[49]

Nevertheless, several important gaps remain unaddressed in the existing literature. Studies focused on data-center waste-heat recovery have been constrained to low-temperature sources (typically below 80°C) which limits applicability to low-efficiency, single-effect chillers. Commercial absorption chillers require driving heat across a specific temperature spectrum—from 70 to 230°C depending on absorption-chiller type. Because nuclear reactors can readily supply steam that comprehensively covers these required temperatures (from 45 up to 560°C, spanning from PWRs to high-temperature gas-cooled reactors [HTGRs]), they are uniquely well-suited to drive absorption chillers. Recent research has mapped this design space by examining the electricity-heat tradeoff that arises when steam is extracted from a nuclear plant's balance of plant (BOP).[49] Extracting higher-temperature steam enables efficient multi-effect chillers with strong coefficient of performance (COP) values, but accelerates turbine-efficiency losses, while lower-temperature extraction reduces the electrical penalty at the cost of lower chiller efficiency and a much-larger thermal draw to meet the same cooling load.

Why Does This Matter for Data Centers?

A tightly integrated nuclear-powered data center uses both electricity and heat as primary energy carriers, fundamentally transforming its energy profile and infrastructure economics. In this architecture, reactor heat drives high-efficiency absorption chillers in place of conventional electrically driven units. Shifting this load from electricity to reactor heat directly frees up electrical capacity for revenue-generating IT equipment, including AI and HPC workloads.

A further infrastructure benefit follows from this shift. Because conventional data centers must treat cooling as a critical load, they are required to oversize their UPSs and backup diesel generators to cover both IT and cooling demands during outages.[50] Replacing electric chillers with heat-driven cooling could reduce the required capacity and cost of expansive battery-type backup systems, as well as strengthen the site's resilience by diversifying its backup infrastructure: e.g., through the inclusion of TES.

Research and Findings

To investigate this trade-off systematically, ongoing national-laboratory research is evaluating the benefits of tight integration of nuclear with data center. The research, led by Argonne National Laboratory (ANL), examines direct thermal coupling between several advanced-reactor technologies—PWRs, sodium-cooled fast reactors (SFRs), HTGRs, and very high temperature reactors (VHTRs)—with single-, double-, and triple-effect absorption chillers. A rapid-assessment framework seeks to identify the optimal steam-extraction rate and temperature, effectively probing the theoretical limits of BOP optimization where the electrical savings from displacing conventional chillers are maximized against the generation penalty.

To contextualize these improvements, the baseline electric-chiller configuration assumes a highly efficient vapor-compression system with a COP of 5.2. Using a metric, the peak-IT-to-reactor (PIR) ratio, the study quantifies that tightly integrated NPPs with absorption cooling systems can support 6–16 % more peak IT load per unit of reactor thermal capacity compared to baseline electric cooling configurations, as shown in Table 8. To demonstrate this potential, rapid thermal modeling of proposed coupling architectures is underway, offering a technically and economically grounded pathway toward maximizing plant utilization. Due to data availability, the HTGR was examined for double- and triple-effect configurations only; results below 100°C would likely align with those of the SFR. For each reactor type, the preferred absorption chiller configuration depends on the reactor outlet temperature, house-load thermal requirements, and condenser/heat sink temperature.[51]

The VHTR delivers the greatest data-center support capacity, outperforming other configurations by up to 16%. While VHTRs represent a longer-term commercialization horizon relative to near-term deployments of other advanced reactors, they are included to establish the upper bound of tight integration performance, defining the limits of design and operational space. Because its direct Brayton cycle rejects high-quality waste heat at temperatures up to 125°C, it ideally suits the driving requirements of standard to high-efficiency absorption chillers with no turbine extraction penalty; chiller-type selection will therefore hinge primarily on capital cost.

Table 8. PIR ratio improvement for several nuclear-reactor types with different chiller technologies.

	Electric Chiller (Base) [IT Load MWe/ NPP Capacity MWt]	**Single-Effect** [Improvement %]	**Double-Effect** [improvement %]	**Triple-Effect** [Improvement %]
PWR	0.231	6%	7%	6%
SFR	0.274	7%	8%	6%
HTGR	0.267	-	6%	6%
VHTR	0.293	16%	16%	16%

To determine total site-wide cooling duty, the framework assumes all IT power and UPS losses convert fully to heat, equating 1 MWe of electrical load to 1 MWt of thermal load. Absorption chillers are modeled across their optimal driving temperature ranges: ~80–100°C for single-effect, 100–130°C for double-effect, and 140–170°C for triple-effect systems. Using a bottom-up approach[49] that aggregates IT, auxiliary, and cooling demands, the framework yields a typical power-usage effectiveness of 1.4.

Key Insights

The derived efficiency improvements are broadly technology-agnostic: reactors with comparable thermal outlet conditions and BOP configurations tend to yield similar gains in IT support capacity, suggesting that these results are transferable across reactor designs rather than tied to any single technology. The VHTR represents the upper bound of integration performance given its high outlet temperature, but near-term reactor designs achieve meaningful improvements as well. Across configurations, the framework consistently shows that displacing conventional electric cooling with heat-driven absorption systems meaningfully expands the IT capacity a given reactor can support. This points to a generalizable pathway for maximizing nuclear plant utilization in data-center applications.

Further Exploration

The use of absorption refrigeration for cooling has been studied across several contexts although it has not yet been deployed in tight integration with nuclear systems. While successfully deployed globally, nuclear cogeneration lacks large-scale commercial demonstration in the U.S. Pioneering this tight integration will require navigating first-of-a-kind (FOAK) licensing challenges, specifically regarding the safety of directly coupling a data-center chiller system to a reactor's BOP. Absorption chillers operate on the same thermodynamic principles as vapor-compression systems but use thermal energy rather than mechanical work to drive the cycle.[52] Existing studies have explored harvesting low-grade waste heat from data centers to power absorption chillers, highlighting measurable reductions in electrical load.[35,53] Other work has examined absorption-chiller integration with natural-gas turbines, using exhaust waste heat to pre-cool compressor-inlet air and thereby improve turbine output.[54] A separate body of literature addresses solar-assisted lithium-bromide water-absorption chillers, confirming that multi-effect configurations achieve substantially higher COPs when supplied with higher-temperature heat.[55]

EPRI explicitly examines steam-driven absorption cooling as a strategic technical pathway for nuclear integration.[56] A thermal modeling study by INL examines how nuclear power plants can simultaneously supply electricity and cooling for hyperscale data centers (Table 9). Using Aspen Plus, it compares vapor-compression refrigeration with water/lithium bromide (LiBr) absorption chillers across multiple BOP coupling options. The study introduces COP of electricity (COPe) and COP of heat (COPx) to translate steam-extraction penalties and waste-heat recovery into comparable effective cooling performance. Chiller efficiency and plant impacts are shown to be governed primarily by evaporator or condenser temperatures and steam-diversion point within the turbine cycle. Table 9 shows that coupling an extraction-condensing turbine with a double-effect absorption chiller yields a COPe in the range of approximately 4.78–6.89, with an optimum at modest steam-extraction levels (≈0.1–0.2). However, at this operating point, the configuration limits the IT load that can be supported for a given reactor size. In contrast, a backpressure-turbine configuration can increase COPe to as high as 7.43 at higher cooling fractions, but this improvement comes at the cost of greater electricity-generation loss due to increased steam extraction.

Using the information of Table 9 and other reports, as described in [45], cost estimates for absorption chillers vapor-compression chiller (VC) chillers of 2000 RT (~7033 kW) with operation time of 99.5% were performed and tabulated in Table 10. A steam price of $3.5/MMBTU is derived from natural-gas price of $3/MMBTU using a conversion table (Table 9), and electricity cost is assumed to be 0.0886 $/kWh, the average industrial-electricity cost in the U.S. during June of 2025.[57] The installed cost of a VC chiller is less than an absorption chiller. However, energy cost for operation of absorption chillers is less than for VC chillers. The difference is even greater if waste heat or other readily available heat can be used in the absorption chiller. In the U.S., because the natural-gas price for heat is inexpensive relative to other countries, the advantages of the absorption chiller are enhanced.

Industry Commentary

Forthcoming

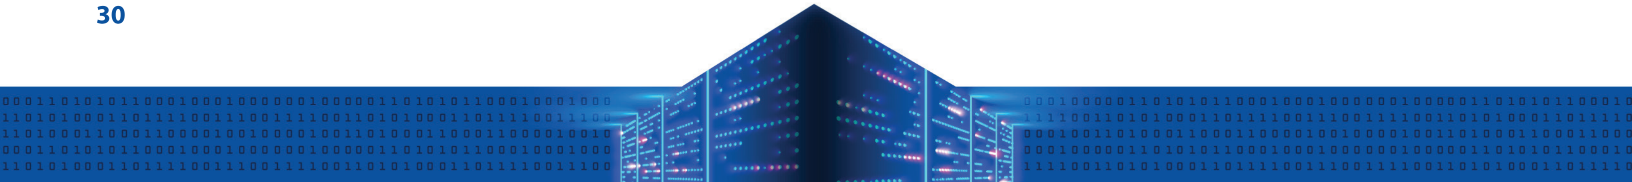

Table 9. Summary of combined cooling, heating and power configurations for data center with reactor.

Chiller Type	Features	ECE [%]	COP Cooling (COPvc or COPac)	COPe Effective Cooling)	COPx (CCP)	COPhp (HP)	Final Heat Released Temp. [°C]	Note
Integration of Chiller and BOP for Data Center								
Vapor compression	Te of 9.22 °C	30.96	2.24-7.33	—	—	—	32-65	COP highly depends on evaporator and (low pressure) condenser temperature
	Te of 15 °C		2.24-11.8					
	Te of 25 °C		3.72-32					
Absorption + Extraction condenser turbine	Double effect Te of 15 °C	24.1-30.3	1.36-1.45	4.78-6.89	5.03-8.76	—	32-65	Optimal cooling capacity exists with extraction ratio of 0.1–0.2 regarding COPe, in which the amount of cooling capacity is less than the amount of electricity generated. So, auxiliary cooling equipment is required for data centers. Thus, it is not good for large capacity as reducing turbine reliability and COPe.
Absorption + Backpressure turbine	Double effect Te of 15 °C	14.7	1.36-1.45	0.63-7.43	5.41-9.70	—	32-120	COPe is dramatically improved when large cooing capacity. High temperature waste heat can be obtained without a heat pump but resulting in sacrificing cooling efficiency a lot.
Absorption + Direct extraction before turbine	Double effect	10.6-29.6	1.429	4.06-4.50	4.06-4.50	—	32-65	COPe does not change much but is lower than other configurations.
	Triple effect		1.8	5.12-5.67	5.12-5.67	—	32-65	It is a rough estimation Need multiple effect absorption chiller for efficient use of heat
	Quadruple effect		2.2	6.26-6.93	6.26-6.93	—	32-65	
With Heat Augmentation for Reuse of Waste Heat								
Vapor Compression	Cascade heat boosting	30.96	—	—	3.24-4.48	2.12-2.74	105	COP is a strong function of evaporator and condenser temperature
Absorption	Cascade heat boosting	23.15	1.382	6.50	1.75	4.131	105	Using an absorption condenser temperature of 60°C. Cascade heat boosting is used to avoid high temperatures and pressure ratios
Absorption	Direct steam compression	23.15	1.382	6.50	1.78	5.081	105	Using an absorption condenser temperature of 60°C. More efficient direct steam compression is used but has a high pressure ratio and high max temperature.

Energy (electricity) conversion efficiency (ECE), absorption chiller (AC), combined cooling and power (CCP), temperature of evaporator (Te)

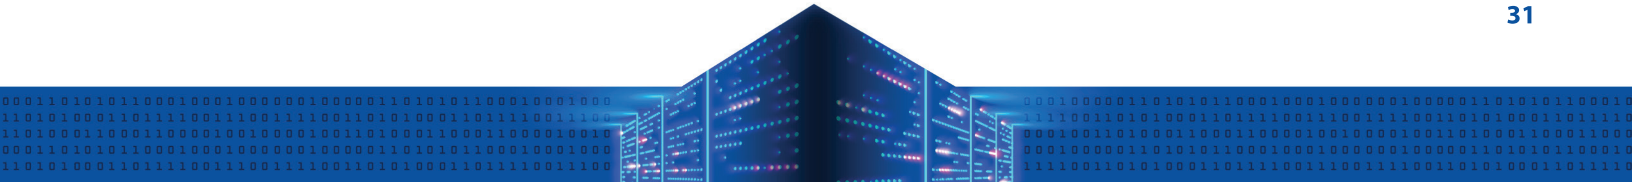

Table 10.Capital- and operational-cost estimates of absorption chiller (from DOE EERE 2024 and 2017) (Department of Energy 2024).

Design	SIngle effect†							Double effect†						
Heat source	Hot water						Steam (low p)	Steam			Steam (high p)		Exhaust fired	
Nominal Cooling Capacity (tons)	50	250	500	1000	50	440	1320	500	1000	2000	330	1320	330	1000
Installed cost ($/ton)	4508	2318	1863	1590	6000	2300	1800	2162	1671	1314	3000	2200	3300	2000
└ Equipment ($/ton)	2373	1129	860	697	2010	930	820	1086	787	649	1190	1000	1330	930
└ Installation ($/ton)	2135	1189	1003	893	3990	1370	980	1076	884	665	1810	1200	1970	1070
O&M cost (¢/ton-hr)*	3.75	1	0.75	0.5	0.6†	0.2	0.1	0.9	0.65	0.4	0.3	0.1	0.3	0.1
Reference	a	a	a	a	b	b	b	a	a	a	b	b	b	b

* *O&M costs do not include energy costs required for operation but do include all maintenance requirements associated with an absorption chiller, including periodic purging of non-condensable gases and monitoring of the cooling tower and of chilled water quality.*

† *There is huge difference between 2024 and 2017 data, cost of 0.5-1¢/ton-hr from 2024 data except small size (50 ton) is a reasonable estimation; thus, interpolation of these data is used for calculation of cost estimation done in Table 11.*

Table 11.Cost estimate for absorption chiller and compression chiller of 2000 RT (~7033 kW) with operation time of 99.5%

	Absorption chiller[*]		Compression Chiller[†]	
	Single effect, Steam	Double effect, Steam	Full load	60% load
COP	0.8	1.4	6.5	9.6
Installed cost	$2.56M[‡] ($1,280/ton, extrapolated)	$2.63M[‡] ($1,314/ton)	$2M ($1,000/ton)	$3.3M ($1,000/ton)
Energy cost	$916k/year	$523k/year	$836k/year	$566k/year
O&M cost	$53k/year (0.3 ¢/ton-hr, assumed)	$70k/year (0.4 ¢/ton-hr)	—[§]	—[§]

[*] *Steam price of $3.5/MMBTU, derived from natural-gas price of $3/MMBTU.*

[†] *Electricity cost, 8.86 ¢/kWh.*

[‡] *Not considered for installation of steam boiler.*

[§] *No data available*

2.4 Energy-Water Nexus

Introduction

The intricate linkage between water and energy is referred to as both the "water-energy" nexus and the "energy-water" nexus. This nexus acknowledges that water is used in all phases of energy generation and production while energy is required to treat and transport water. Conveying and treating water is energetically expensive while producing energy consumes large quantities of water. This water-energy interaction was initially explored by DOE in the report "Water-Energy Nexus: Challenges and Opportunities Overview and Summary." The report served as a call to action to improve water-efficient energy technologies and promote energy-efficient water management. The report highlighted the need for improved data on water usage within the energy sector. As a "preliminary assessment," the report identified the following water-for-energy technologies for more in-depth analysis: cooling, waste-heat recovery, process-water efficiency and quality, alternatives to freshwater in energy production, and hydropower.[58]

In the absence of a unified approach to the DOE recommendations, attempts have been made on a sector-by-sector basis to address these complex relationships, with limited success. For example, Cafferty et al. provided an update to cooling water demands in the commercial nuclear fleet in their report, "Analysis of Water-Energy Issues for Nuclear Power with Industry Perspective," as quoted in [58]. Other efforts on the process of water and waste heat have seen some success.

Why Does This Matter for Data Centers?

Historically, interactions between energy and water systems have been developed, managed, and regulated independently, for the most part. While the interdependence between water and energy systems has long been recognized, there is growing importance as water and energy demand increase for data centers. Shehabi et al. reported direct water consumption of data centers tripling, from 21.2 billion liters in 2014 to 66-billion-liters in 2023. Similarly, they estimated the indirect water consumption (i.e., energy-water nexus) of data centers in 2023 at 800 billion liters, more than 100 times greater than the direct use. [59] Furthermore, WestWater Research forecasts that data center related water consumption in the U.S. is likely to increase by 170% by 2030.[60]

Research and Findings

It is reported that currently, 20% of data centers rely on watersheds under moderate to high stress due to drought. Although data centers typically use less water than agriculture or municipal sectors, they can still represent a significant demand and impact on localized areas from direct and indirect demands. For example, multiple largescale centers are clustered in Arizona's Tech Corridor in Mesa and Silicon Valley's Santa Clara. Clustering large demands individually may have benign impact, but if considered as a group, the impact can be influential. Finally, while work is employing site-cooling configurations that rely less on water, these configurations can require substantially more power and involve less-established technologies. New power—particularly thermal plants—that come online to support data centers have their own significant water demands for cooling.

In the U.S., water has historically been managed as a public resource, meaning it is held in trust by the government for common use rather than private ownership. Water resources encompass lakes, streams, groundwater, coastal water, wetlands, and even newly emerging alternative sources like treated (e.g., reused, recycled water) wastewater. This contrasts with other resources, like minerals, that can be extracted and exploited for profit. Similarly, energy in the U.S. is not a public resource because over 80% of energy infrastructure is privately owned, operated by investor-owned utilities, and driven by private investment. While public power and cooperatives exist, the bulk system is primarily market driven. As a result, many U.S. energy-sector water decisions are made by private entities. Conversely, public or state entities have the majority of the authority over water use and allocation policies and decisions.[61]

Although water is part of a connected system, or "water cycle," it is typically regulated based on its source—i.e., surface water, groundwater, brackish water have different rules—but even within the category of "surface water," water regulations can vary depending on whether the surface water is perennial, ephemeral, or anthropogenic. In general, water law—the field of governing ownership, control, allocation, and use—can be divided into two main areas: rights to use water and restrictions on quality degradation. More specifically, water law balances public rights and private rights to use water with the relative rights of individual water users and water quality and the regulation of discharges to water. In the U.S., water allocation is dominated by thermoelectric power generation, with over 40% withdrawals (total diverted) and roughly 28% consumption.[62]

The American Society of Civil Engineers reports on the status of American infrastructure. In 2025, the overall grade for water infrastructure was a C-. The report notes that despite recent increases in federal funding, the U.S. water infrastructure remains aged and underfunded. The Environmental Protection Agency (EPA) reports that the deficit of funds to be almost $625 billion over 20 years, which is $150 billion more than the 2018 assessment. However, as reported in common reporting standard (CRS) R46471, "Federally Supported Projects and Programs for Wastewater, Drinking Water, and Water Supply Infrastructure," for decades, Congress has authorized, modified, and funded federal programs and projects to help communities address water supply and water-infrastructure needs administered by the Bureau of Reclamation, the United States Army Corps of Engineers (USACE), the United States Department of Agriculture (USDA), EPA, the Department of Housing and Urban Development, and the Economic Development Administration. Despite federal support of traditional financial tools—including project grants, formula grants, capitalization grants, direct and guaranteed loans—many alternative financing approaches to encourage private-sector investments and public-private partnerships (P3) are being considered and could be an opportunity for data-center developers.

Regional and local energy- and water-resource opportunities vary based on several factors, including (1) the type of water and the local and regional significance of water use, (2) regional conditions and policies for management of source-water and/or waste-water, and (3) energy type being produced and/or used.

Key Insights

1. Communities are concerned about their ability to maintain water services to both existing and new demands, including data centers. Maintaining services means balancing demand growth with necessary quality and quantity requirements as regulated under prescriptive regulatory laws, but also within structural financial limitations.
2. Zero-water evaporation and water-free cooling alternatives being considered as a means to reduce water consumption need to consider energy-water nexus consumption. The savings may or may not be net positive if advanced-manufacturing needs for ultrapure water and more-critical material mining are required to supplement a greater number of resources needed in next-generation central processing units.
3. Alternative water sources, like reclaimed and recycled water, are additional opportunities that should be explored as source water for data-center developers considering construction in arid and semi-arid environments. However, the energy use from the treatment requirements must be considered, as should zero-water evaporation and water-free cooling indirect water use.
4. New legislation policies are pushing for more reporting requirements on water transparency; more consistent reporting conditions across state boundaries would put regions and states on similar playing fields for comparison.
5. As regions trend to multi-scalar water-governance frameworks, like in the Colorado River basin, net water evaluation could be important to show holistic impact. A standard or certification program may bridge the gap needed for energy-water nexus considerations. Exploring existing programs, for example Leadership in Energy and Environmental Design that uses its own criteria for water efficiency but also includes ISO 14044, "Life Cycle Analysis," and ISO 14046, "Water Footprint," could benefit certain data centers.

Further Exploration

Continued opportunities allow water entities to embrace innovation and improvement. Many water utilities actively improve infrastructure through innovations like asset monitoring, which improves the ability to identify issues before they become failures. For example, the U.S. drinking-water infrastructure is composed of more than 2 million miles of underground transmission and distribution lines. Each year an estimated 33 trillion gallons is lost due to leaks, breaks, and failures. Recovering a small fraction of these system losses would cover the projected data-center need. Ironically, asset monitoring and prediction innovations use data and will need data centers to implement.

Similarly, there is a disconnect in understanding the scale of water use in terms of localized impact. The energy-water nexus must show individualized influence at the community level if public resources will be accepted for data-center needs. Attempts at the county level have been made by Pacific Northwest National Laboratory, but more work is needed. Finally, data centers must do a better job of showing public benefit. As the public understands the scale and scope of data centers, acceptance will follow. If data centers are used to advance medical discovery or improve accountancy, those successes, properly communicated to the public, will increase the public's support of the use of its resources.

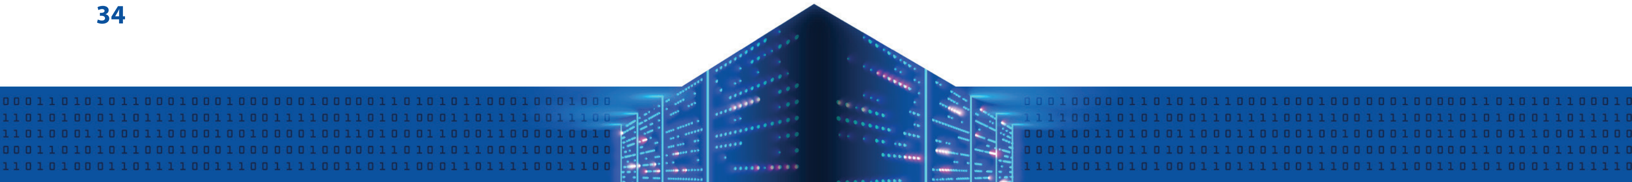

Industry Commentary

Forthcoming

2.5 Controls for Data Center Microgrids

Introduction

Microgrids are emerging as an effective solution for meeting diverse energy needs for data centers. Integrating small reactors (SRs) with traditional thermal generation or novel high efficiency fuel cells operating on natural gas would offer a reliable option for decentralized power generation. Dedicated operation of power sources alongside locally opportunistic power generation and storage systems could produce power for operation of data centers capable of meeting the extreme uptime requirements of some data centers with the flexibility to assist the distribution and bulk electric systems when not islanded.

Reducing development risk will increasingly emphasize physical testing and validation of data-center microgrids, bridging analytical studies with fielded demonstration. These demonstration platforms will provide critical validation of dynamic behavior and system safety and reliability across multiple scenarios, reducing the technical and institutional risks associated with integrating power sources and storage into operational data center microgrids.

As this playbook has shown the path for establishing the data-center application with adequate controls needs to follow a framework. The illustration in Figure 11 shows such a framework with a nuclear-powered application as an example.

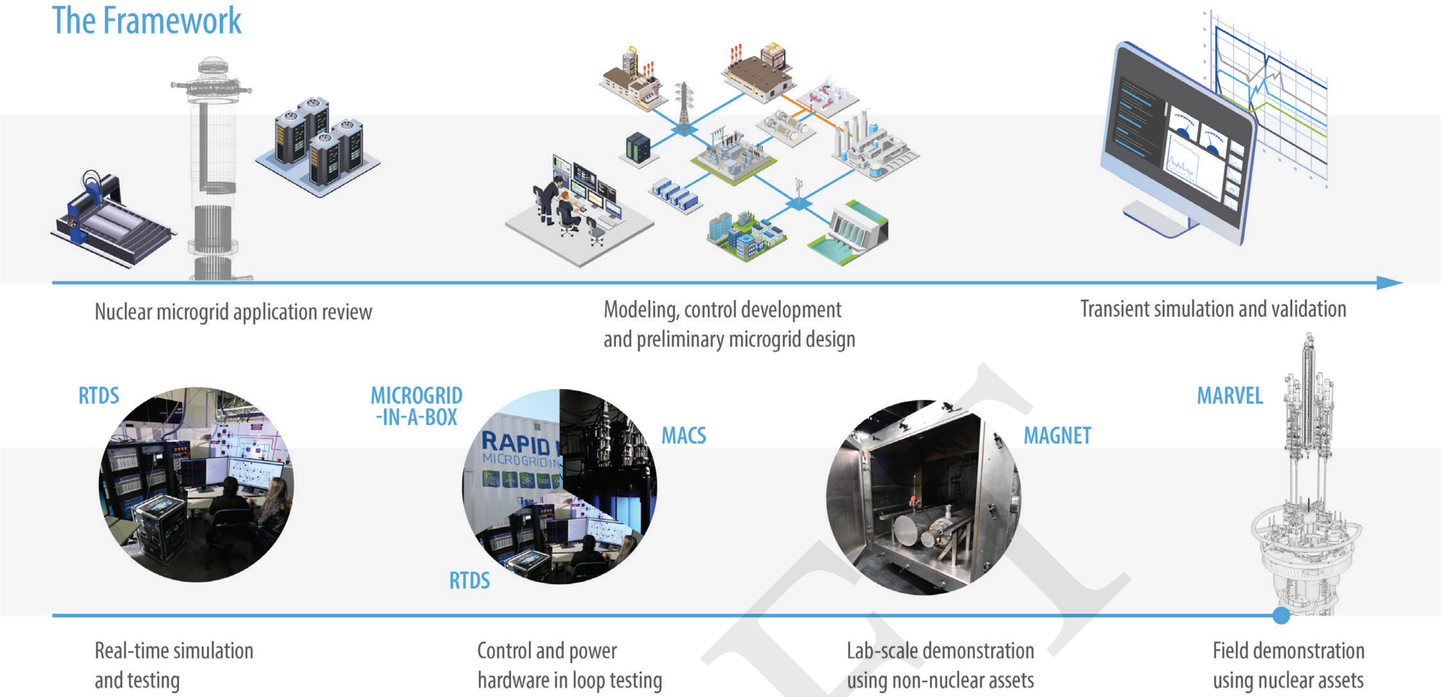

Figure 11. Data-center microgrid feasibility and testing framework.

The framework supports the full range of pre-deployment activities from early-stage modeling and TEA to real-time simulation, hardware-in-the-loop testing, and physical implementation. Ultimately, this initiative aims to demonstrate the deployability of microgrids to serve emerging high-demand applications such as data centers. To achieve this objective, this framework focuses on the following key research and development areas:

- Modeling and Simulation Resources: Develop a comprehensive library of simulation models, and root-mean square and cost-estimation tools to support analysis of data center microgrid applications.
- Design Configurations for Microgrids: Create a library of energy system configurations for power resources in microgrids including nuclear that support both electricity and heating needs integrating thermal and electric storage integration in grid-connected and islanded configurations.
- Technology and Market Database: Establish a centralized database of power resources with detailed technical and economic characteristics to support early-stage screening and application matching.
- Energy-Management System: Design an operational framework to optimize dispatch of generation and storage, and manage loads based on dynamic capabilities, ensuring safe and efficient performance in both normal and emergency scenarios.
- Integrated Control and Protection-Coordination Framework: Develop a generalized, standards-compliant control architecture for coordinating nuclear, thermal, and electrical subsystems, with AI/machine learning (ML)-enhanced autonomous control and validation through control hardware-in-loop (CHIL) and physical-component testing. Develop strategy to coordinate protection of power resources, thermal, and electrical subsystems under variety of design configurations that ensure safe and resilient operation during transients, aligned with standards.
- Testing and Demonstration Platforms: Build and use platforms to CHIL simulations, lab-scale data-center microgrid applications.

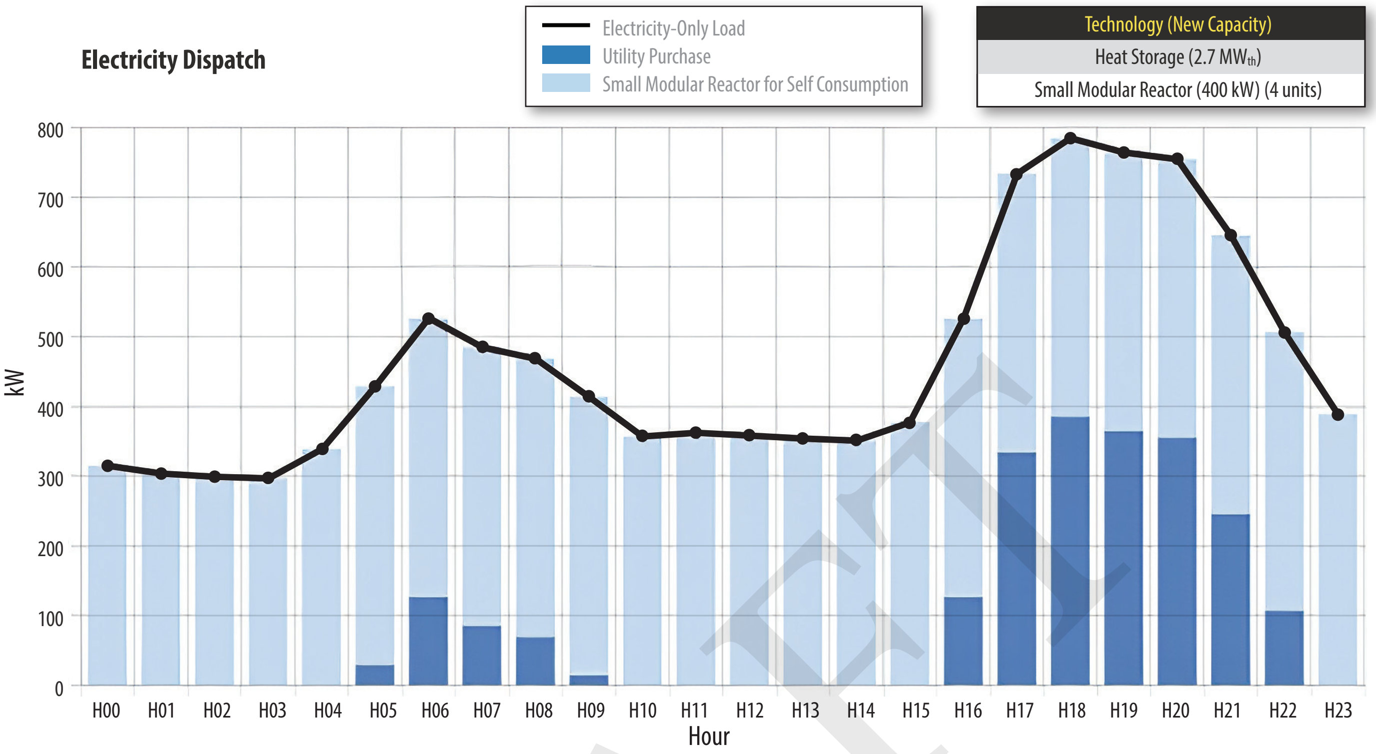

Electricity Dispatch

Technology (New Capacity)
Heat Storage (2.7 MW_{th})
Small Modular Reactor (400 kW) (4 units)

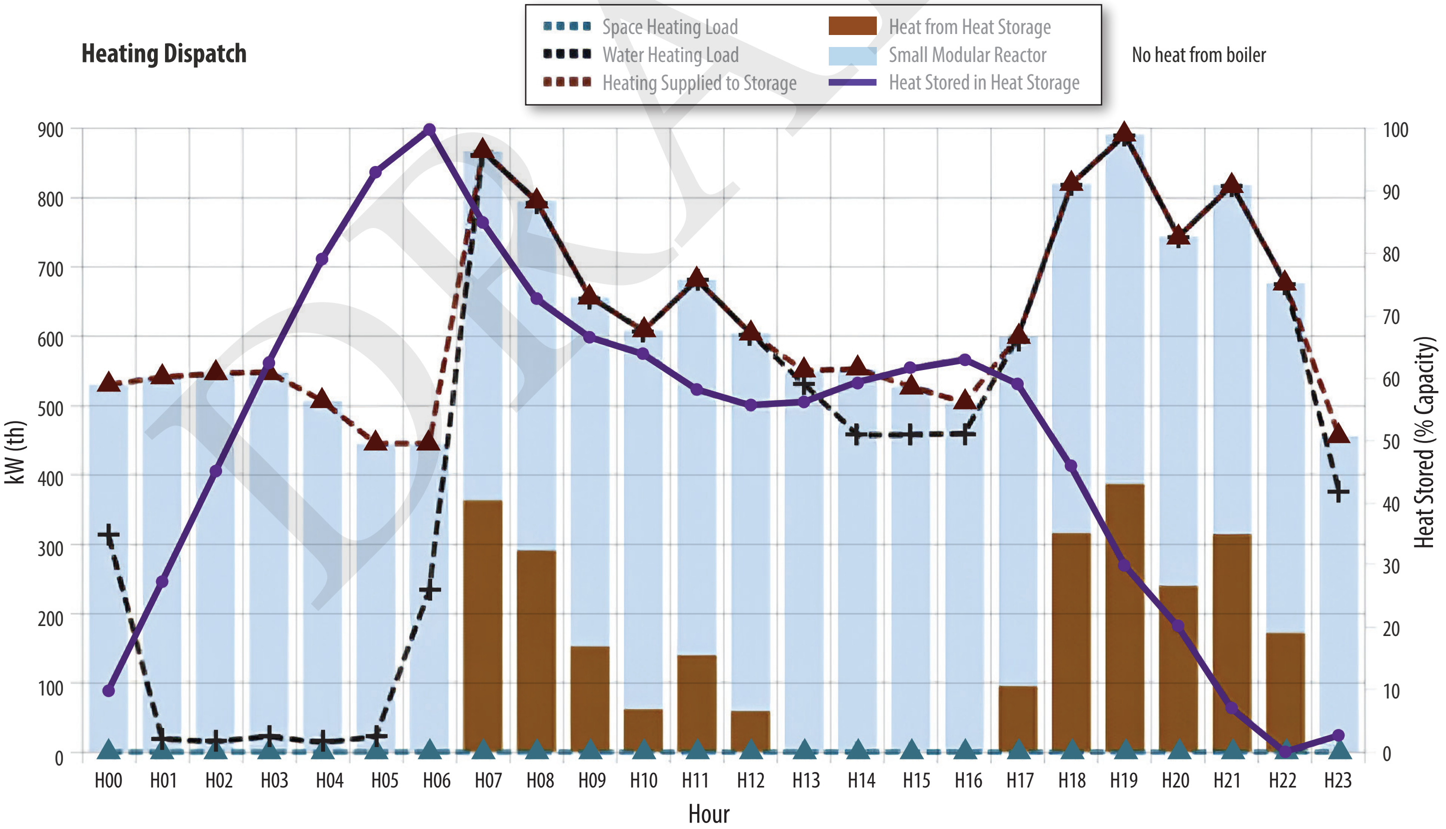

Heating Dispatch

No heat from boiler

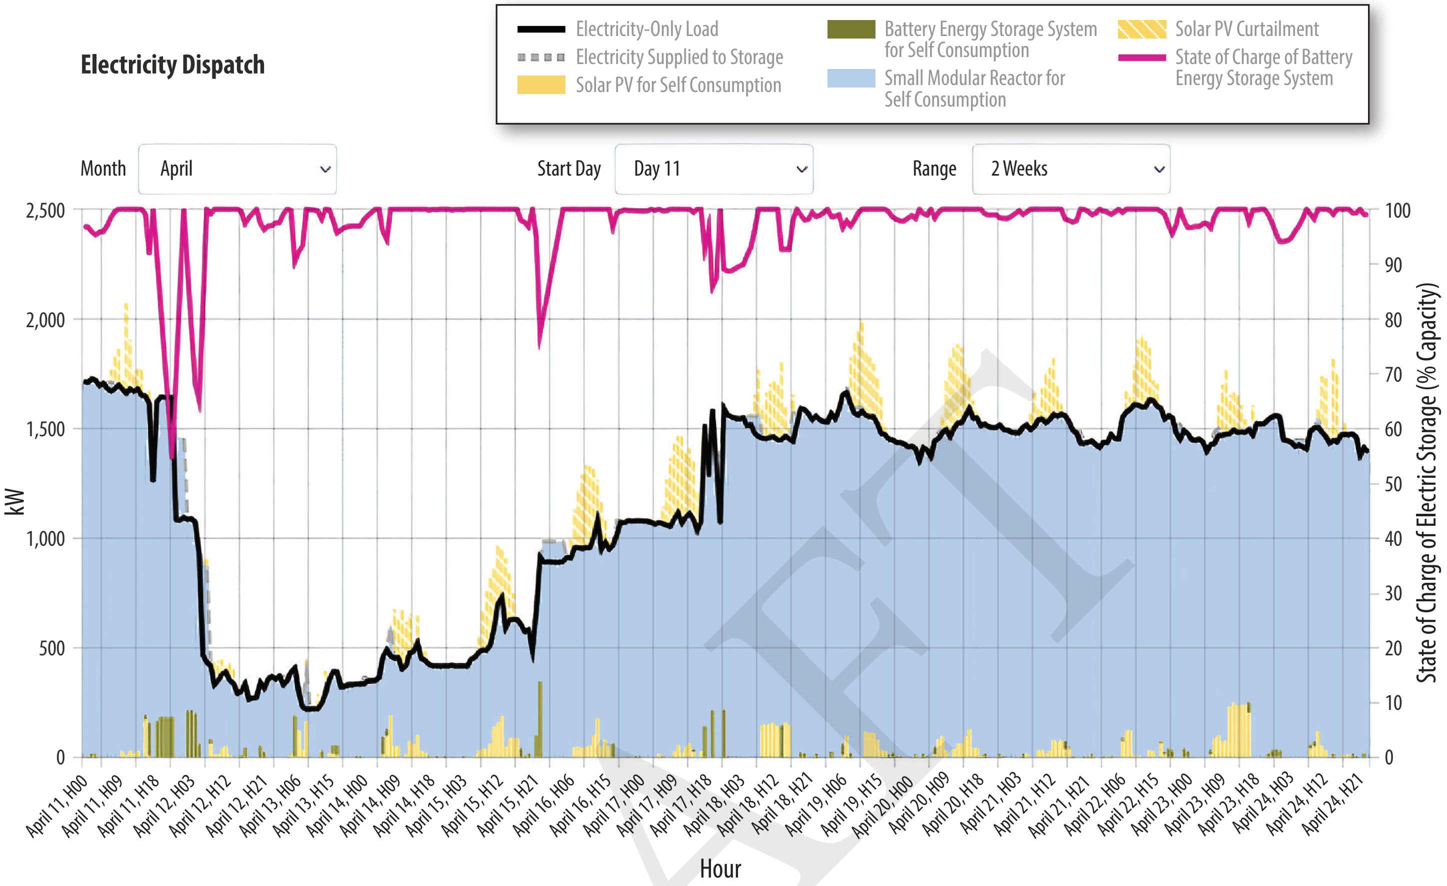

Month: April
Start Day: Day 11
Range: 2 Weeks

Figure 12. Two use-case examples. Previous page: a case explores the integration of a nuclear reactor or traditional fossil fuel-based heat and power generation, BESS and photovoltaic (PV) for off-grid data center facility. Above: a use case exploring the integration of an SR and TES to support heat and electricity demand of a grid-connected microgrid.

Technoeconomic optimization typically leads to hour dispatch solutions for designing and potentially operating the system at the hourly time scale (see Figure 12). However, this is a design that works in theory but needs to be tested in real time implementation of the control system that must manage the large slower-moving energy resources alongside the agile systems required to maintain frequency and voltage stability in islanded mode.

Successful deployment of data-center application microgrids requires tight coordination of traditional electricity and combined heat and power (CHP) control systems with advanced control strategies capable of coordinating large energy producers and locally available energy sources. Figure 13 illustrates the main blocks of the combined SR-based microgrid system, the low-level control systems, and the overarching application management system that operates applications of the microgrid. The integration of SRs into microgrids introduces a new type of highly reliable, continuous electricity and heat generation. Advanced control strategies integrate across the system and must ensure operational flexibility, safety, and regulatory compliance while maximizing the microgrid's resilience and performance needs. Data centers demand highly reliable, safety-critical controls, especially when operating within microgrids that include variable generation and storage. In such settings, control systems must support autonomous operation, islanding and resynchronization, and dynamic load balancing across both dispatchable and intermittent resources.

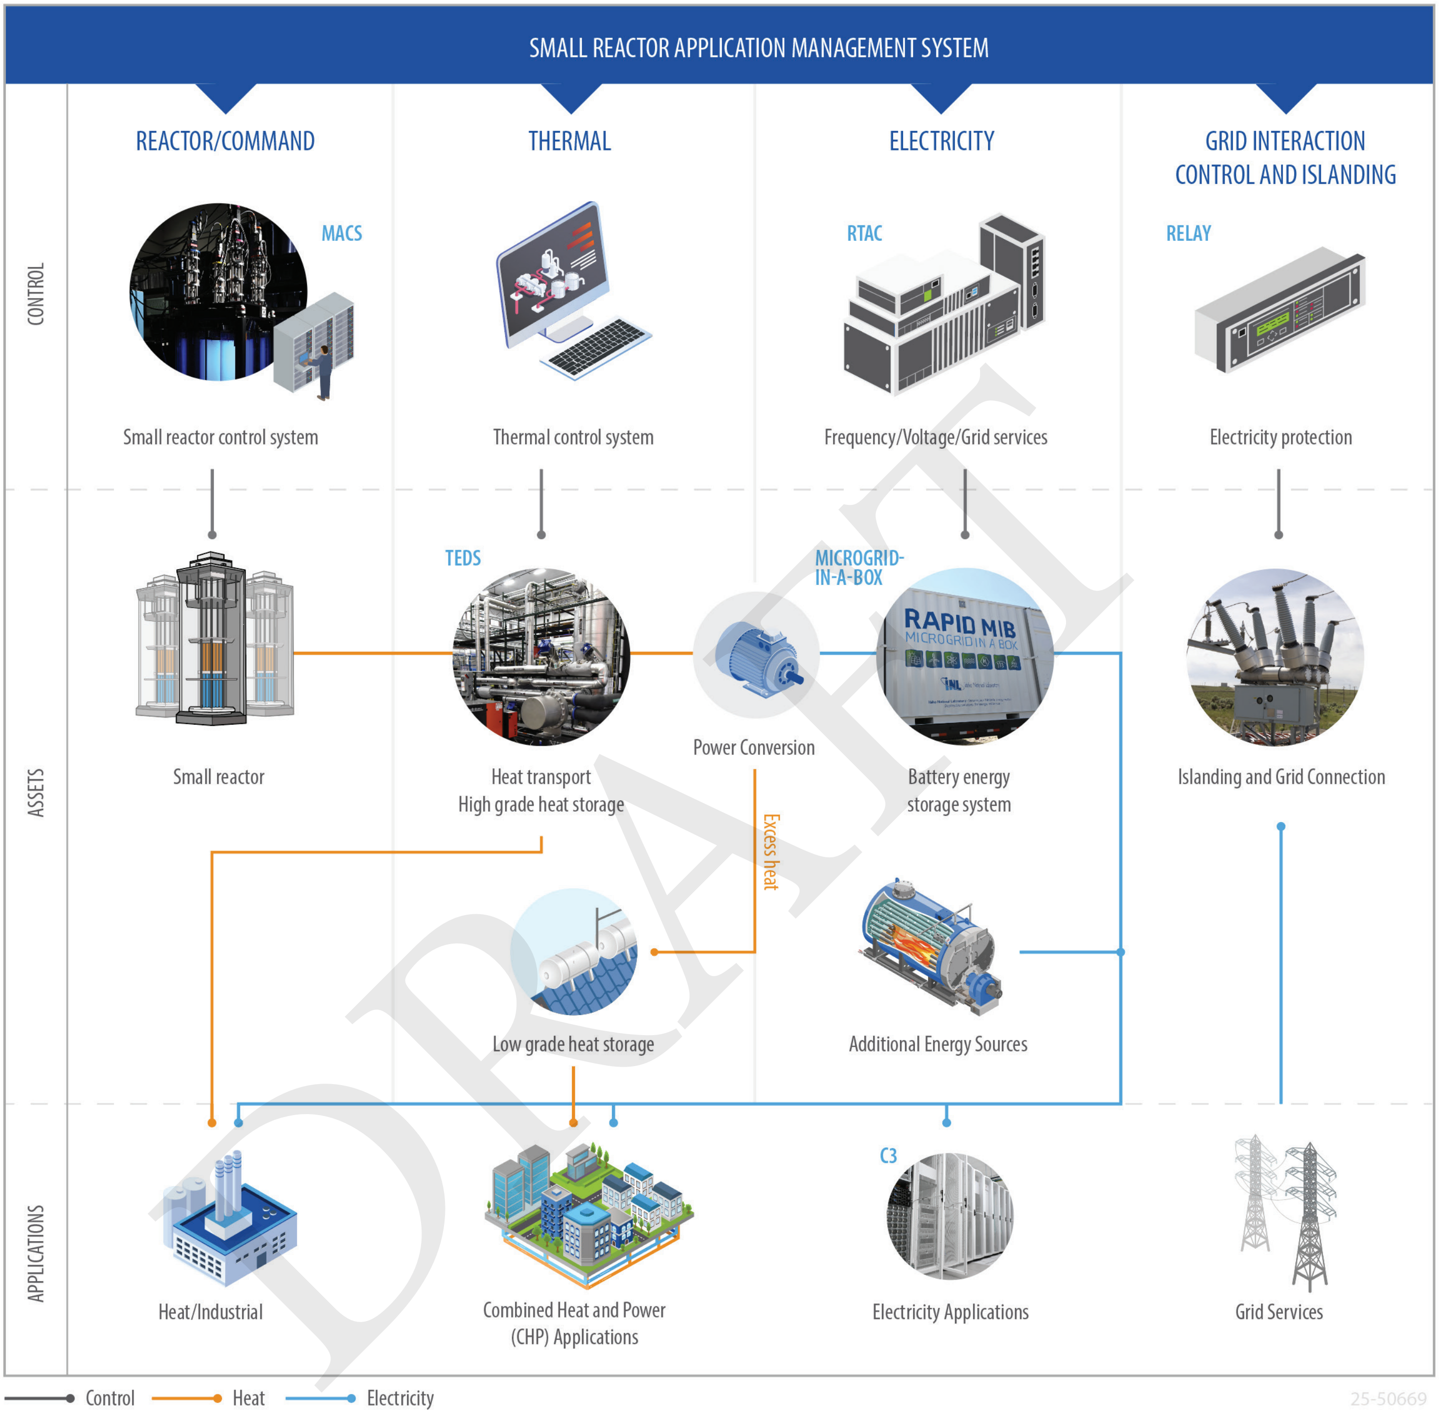

Figure 13. Application management system. The system must both leverage existing electric and CHP microgrid systems that can support facility and campus as well as power for computation in the data center.

Why Does This Matter for Data Centers?

- Many data centers require very high reliability of the grid.
- Data centers have been projected to have volatile power usage in some applications.
- Data centers are in a position of being required to "bring your own power" and can in that case optimize over scenarios where the grid is connected.
- Control systems of complicated energy systems require the ability to be tested with the intended application being represented.

Research and Findings

The definition of a process that is exemplified in the development of a small nuclear-energy-based data-center microgrid is shown below.[63]

SRs, including microreactors and SMRs, can power microgrids for years without refueling. They can supply reliable electricity to remote areas, data centers, and mission-critical facilities. INL is accelerating the safe and affordable deployment of microgrids powered by small nuclear reactors. Using an approach like black starts of run-of-the-river hydroelectric systems, INL is using its real-time simulation capabilities to help nuclear reactor vendors and microgrid developers mitigate technical and financial risks (Figure 14). INL researchers develop representative models for small nuclear reactors and other energy assets such as battery storage, diesel generators, and flexible load with advanced control systems in a real-time simulation environment.

Researchers conduct hardware-in-the-loop tests with models, gradually replacing simulated components with physical hardware, avoiding the risks and costs associated with direct experimentation on one-of-a-kind equipment.

Because any part of the system can be simulated in real-time, it can easily be reconfigured to incorporate both simulated and physical distributed-energy assets and adapt to different microgrid topologies. This flexibility allows developers to test a wide range of configurations and control strategies before physical deployment, reducing risk and accelerating design iteration.

Real-time grid emulation

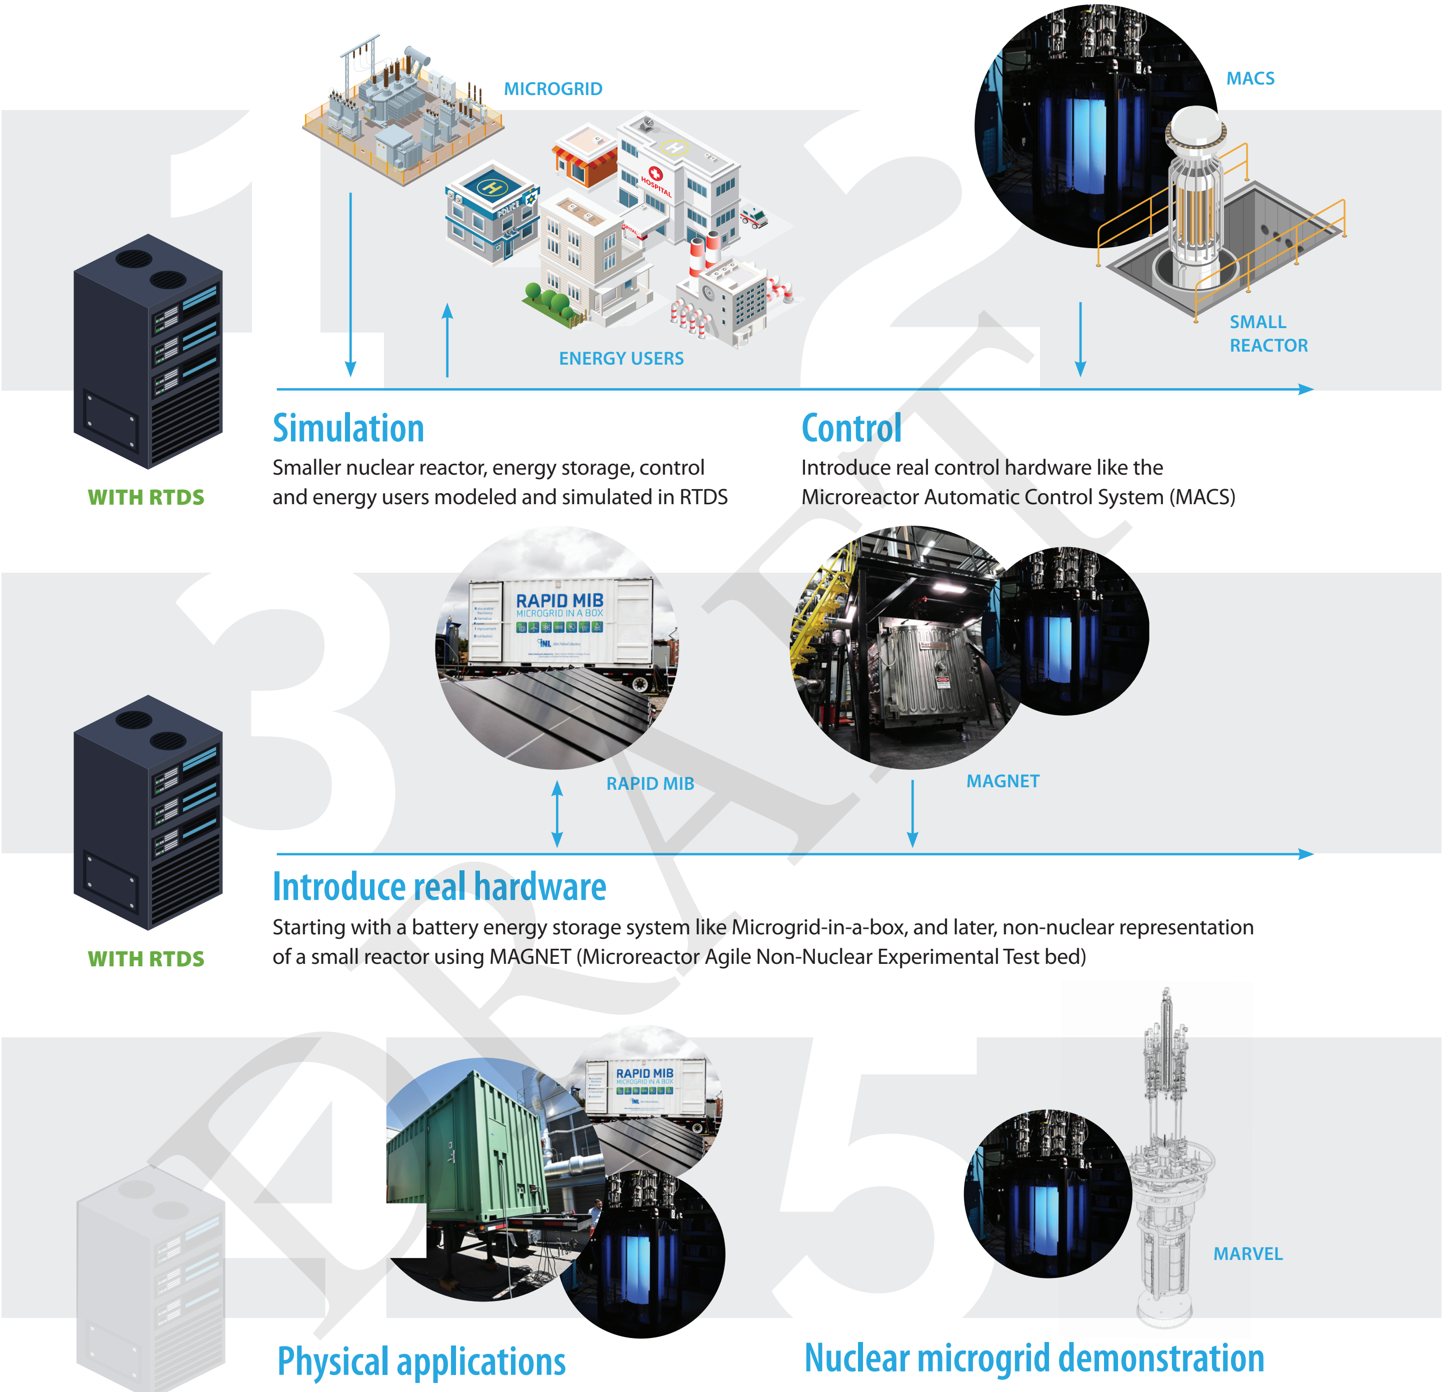

Simulation

Smaller nuclear reactor, energy storage, control and energy users modeled and simulated in RTDS

Control

Introduce real control hardware like the Microreactor Automatic Control System (MACS)

Introduce real hardware

Starting with a battery energy storage system like Microgrid-in-a-box, and later, non-nuclear representation of a small reactor using MAGNET (Microreactor Agile Non-Nuclear Experimental Test bed)

Physical applications

Introduce the real world devices the microgrid is designed to power, like INL's Data Center-in-a-Box

Nuclear microgrid demonstration

Integration with MARVEL, controlled by MACS

The MARVEL (Microreactor Applications Research Validation and Evaluation) project, led by INL under the Department of Energy Microreactor Program, is a nuclear microreactor demonstration test bed.

Figure 14. Approach used to incrementally build confidence in control system design and hardware integration towards demonstration and deployment of a data center microgrid.

Key Insights

Using models early in the development process allows microgrid developers to reduce deployment risks. Testing components with simulated models and physical hardware allows developers to evaluate system behavior, validate control strategies and identify integration challenges well before a physical reactor is built and tested. Studying the reactor's performance in real-time, alongside simulated and physical microgrid equipment, allows researchers to assess whether the system architecture, asset mix and control logic are appropriately configured for FOAK deployment.

Real-Time Digital Simulator (RTDS) and grid emulation enable the modeling of complex or in-development components, such as SRs described by neutronics, gas-Brayton heat-transfer mechanisms, and power-conversion systems. This allows testing with high-fidelity representation of the system components without putting one-of-a-kind equipment at risk.

1. Control systems integration including controller-system hardware such as the Microgrid Automated Control System (MACS), the Microgrid in a Box (MIB) controller and the overarching microgrid-management system or distribution energy-management system. The functionalities provided by these controllers include temperature-control and load-following capabilities.
2. Testing of real-world equipment through integrating physical hardware into the RTDS test bed, as it becomes available, to gain operational experience. In this case, the Microreactor Agile Non-Nuclear Experiment test bed (MAGNET) will be paired with MIB via RTDS for real-time physical demonstration of grid scenarios, to be as prepared as possible for full demonstration using the Microreactor Applications Research Validation and Evaluation (MARVEL). Ongoing real-time simulation and hardware-in-the-loop testing with assets such as MAGNET, MARVEL and MIB, along with advanced controllers like MACS, will support the safe, efficient and scalable deployment of microreactor-based energy systems under realistic grid condition.
3. Looking ahead, this capability will support a broader range of emerging nuclear technologies, including those being demonstrated at INL and deployed by industry partners across the United States.

Further Exploration

- Idaho National Laboratory. 2025. "Accelerating safe deployment of small nuclear reactors for grid use." https://inl.gov/content/uploads/2025/12/25-315_Accelerating-safe-deployment-of-SMRs-for-grid-use.pdf.
- Allied Market Research. 2022. "Microgrid Market Size, Share, Competitive Landscape and Trend Analysis Report, by Connectivity, Type and End User: Global Opportunity Analysis and Industry Forecast, 2020–2030." Allied Market Research.
- Hatch, Cory. 2023. "Microgrid in a Box opening new possibilities in defense, utilities, disaster relief." Idaho National Laboratory. https://inl.gov/integrated-energy/microgrid-in-a-box-opening-new-possibilities-in-defense-utilities-disaster-relief/.
- Idaho National Laboratory. 2023. "Power and Energy Real-Time Laboratory: Modernizing power systems with at-scale testing and emulations." https://inl.gov/content/uploads/2023/07/Power-and-Energy-Real-Time-Laboratory.pdf.
- Idaho National Laboratory. 2025. "Dynamic Energy Transport and Integration Laboratory." https://inl.gov/document/dynamic-energy-transport-and-integration-laboratory/.
- IEEE Standards Association. 2018. "IEEE 1547-2018: IEEE Standard for Interconnection and Interoperability of Distributed Energy Resources with Associated Electric Power Systems Interfaces," IEEE. https://standards.ieee.org/ieee/1547/5915/.
- Ojo, Yemi et al. 2023. "Hydropower Black Start: A Guidebook for Retrofitting Grid Dependent Hydropower." Idaho National Laboratory. https://doi.org/10.2172/2500383.

- O’Neil, Connor. 2023. “City of Cordova and U.S. Department of Energy Push Envelope of Microgrid Design.” National Laboratory of the Rockies. https://www.nrel.gov/news/detail/program/2023/city-of-cordova-and-us-department-of-energy-push-envelope-of-microgrid-design.
- Poudel, Bikash et al. 2023. “Small Reactors in Microgrids: Technology Modeling and Selection.” INL/RPT-23-73046, Idaho National Laboratory, Idaho Falls, ID. https://www.osti.gov/biblio/1988420.
- Poudel, Bikash et al. 2023. “Small Reactors in Microgrids: Technical Studies Guidance.” INL/EXT-21-64616, Idaho National Laboratory. https://www.osti.gov/biblio/1829672.
- Spangler, Donna. 2025. “MARVEL milestone: Microreactor project completes coolant system testing.” Idaho National Laboratory. https://inl.gov/feature-story/marvel-milestone-microreactor-project-completes-coolant-system-testing/.
- • Walton, Rob. 2024. “Five Questions for the 20s: The Most Compelling Microgrid Challenges of the Decade.” *Microgrid Knowledge*.

2.6 Direct Current Power for Data Center Racks

Introduction

Data-center power architectures commonly employ high-voltage AC, typically delivered at 600 or 480 V. This supply is subsequently transformed to 208 or 120 V for distribution to equipment racks supporting servers and other information-technology loads. A UPS with associated energy-storage elements, such as batteries, provides ride through capability and protection from power disturbances. In conventional designs, the incoming AC is rectified to direct current (DC) for storage and then inverted back to AC for distribution through the facility’s electrical system and routed to PDUs supplying individual racks.[64]

However, DC power distribution to computer racks is increasingly being adopted in modern data centers, particularly in high-density computing environments such as hyperscale facilities, AI clusters, and edge computing systems. DC-power architecture centralizes conversion and distributes DC directly to server racks, reducing conversion losses and simplifying rack-level power infrastructure. With the rapid growth of power-intensive workloads (e.g., graphical processing units, AI accelerators), DC distribution is emerging as a viable solution to improve energy efficiency, scalability, and operational resilience.

A typical medium-voltage DC (MVDC) and low-voltage DC (LVDC) architecture for datacenter loads and flexible DC loads and sources is shown as Figure 15. Because servers, batteries, fuel cells, and many wind and solar sources are inherently DC-based, DC distribution reduces the number of power-conversion stages and removes harmonic propagation associated with AC rectification. DC buses can provide fast capacitive buffering that smooths sub-second load ramps while droop-controlled DC converters improve stability when feeding constant power loads. Furthermore, MVDC systems allow for faster fault isolation using solid-state breakers and enable better integration with BESS, improving dynamic response and overall efficiency. As data centers scale toward hundreds of megawatts, especially for AI-driven applications, MVDC and LVDC can provide improved stability, reduced losses, simplified control, and enhanced compatibility with wind and solar generation and storage systems. The overall system will serve to mitigate any power-quality issues or unpredictable-load profile disruption to the bulk electric system when tied to the grid.

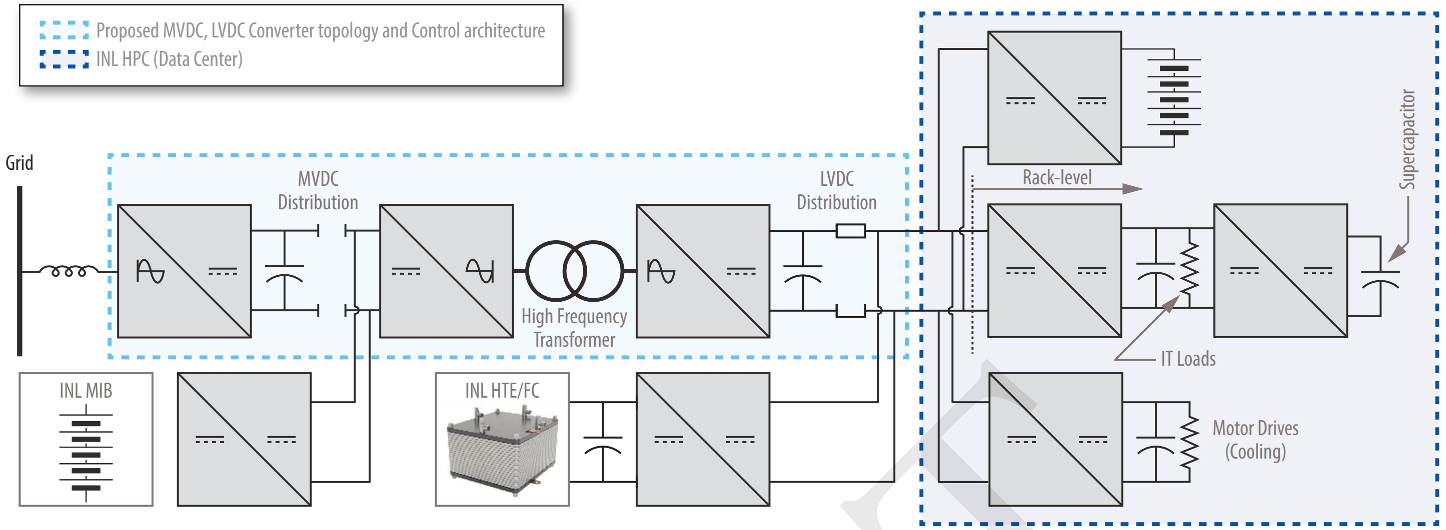

Figure 15. Single-line diagram for MVDC/LVDC distribution design. [65]

Why Does This Matter for Data Centers?

Developing DC-power architecture for data centers posits the following advantages:

- **Increased Efficiency**: DC power distribution eliminates multiple AC-DC and DC-AC conversion stages typically present in conventional architectures. This reduces cumulative conversion losses, improving overall power-usage effectiveness and lowering operational energy costs.
- **Higher Rack Power Density**: By removing individual server power supplies (which can exceed 100 units per rack in traditional systems), DC architecture frees up physical space and reduces thermal constraints. This enables higher computing density per rack, which is critical for AI/ML and HPC workloads.
- **Enhanced Reliability**: Fewer components at the rack level translate to fewer points of failure. Centralized power conversion systems can be designed with redundancy and better monitoring, improving system uptime and maintainability.
- **Reduced Infrastructure Cost**: DC distribution can reduce the size and complexity of busways, cabling, and PDUs. This leads to potential savings in both capital expenditure (CAPEX) and installation complexity.
- **Improved Thermal Management**: Eliminating distributed power supplies reduces localized heat generation within racks, simplifying cooling design and improving thermal efficiency.

Research and Findings

Recent studies and industry deployments indicate that DC-powered data centers can achieve efficiency improvements of 5–15% compared to traditional AC systems, depending on architecture and load profile.[66] Key findings include:

- Conversion Efficiency Gains: DC-powered conversion stages can achieve efficiencies above 97–98%, compared to cumulative efficiencies of 85–92% in multistage AC-DC conversion systems. The center of expertise for data-center energy reports that the overall efficiency of AC-DC conversion systems used for servers and other IT products can be as low as 50%.[67]
- Reduced Component Count: Eliminating rack-level power supplies significantly reduces hardware complexity and maintenance requirements.
- Scalability for High-Power Loads: DC systems are better suited for emerging high-power racks (>50–100 kW), where traditional AC distribution becomes increasingly inefficient and bulky.
- Integration with Energy Storage: DC architecture enables more seamless integration with BESS, reducing conversion steps during backup or peak-shaving operations.
- Industry Trends: Key industry players leading the 800V DC charge include Vertiv, Eaton, Delta, and SolarEdge while hyperscalers like Nvidia (with its Vera Rubin Ultra Kyber platform) are setting the hardware standards that are pulling the entire ecosystem toward DC-native power distribution.[68]

Key Insights

1. **Voltage level matters**: Higher DC distribution voltages (e.g., 380 or 800 VDC) reduce current levels, minimizing I^{2}R losses and conductor-size requirements.
2. **Centralized vs distributed architectures**: Centralized power architectures improve system efficiency and reduce hardware redundancy but require robust redundancy schemes and advanced protection strategies to ensure reliability whereas distributed architectures offer better fault isolation at the expense of lower efficiency and higher component count.
3. **Compatibility with IT equipment**: Adoption depends on server and equipment manufacturers supporting DC input standards, which is increasingly becoming available.
4. **Safety and standards**: While DC systems present challenges such as arc suppression and protection coordination, modern standards and protection technologies address these concerns.
5. **Enabling the use of intermittent energy resources to improve efficiencies**: DC architecture supports direct integration of intermittent energy sources, such as wind and solar, and DC-power sources such as fuel cells. Natural-gas-fueled solid-oxide fuel cells can achieve electrical efficiencies approaching 60–70%, relative to approximately 30–40% for conventional gas-turbine generation. When directly coupled to DC-powered data-center racks, this high-efficiency, DC-native generation pathway minimizes conversion losses and translates into reduced fuel consumption and lower operational expenses.

Further Exploration

- **Standardization efforts**: Continued development of industry standards—e.g., EMerge Alliance, ETSI, Open Compute Project—is critical to accelerate adoption
- **Hybrid AC/DC data centers**: Transitional architecture combining AC and DC distribution may provide a practical pathway for retrofitting existing facilities
- **Advanced protection systems**: Research into DC fault detection, isolation, and arc mitigation remains an important area for safe deployment
- **AI and high-density workloads**: As AI workloads continue to scale, DC distribution may become essential for supporting ultra-high-density racks efficiently
- **Economic analysis**: Further cost-benefit studies comparing life-cycle costs (CAPEX + operational expense) of DC vs AC systems are needed across different deployment scales.

Chapter 3

IDENTIFYING STAKEHOLDERS

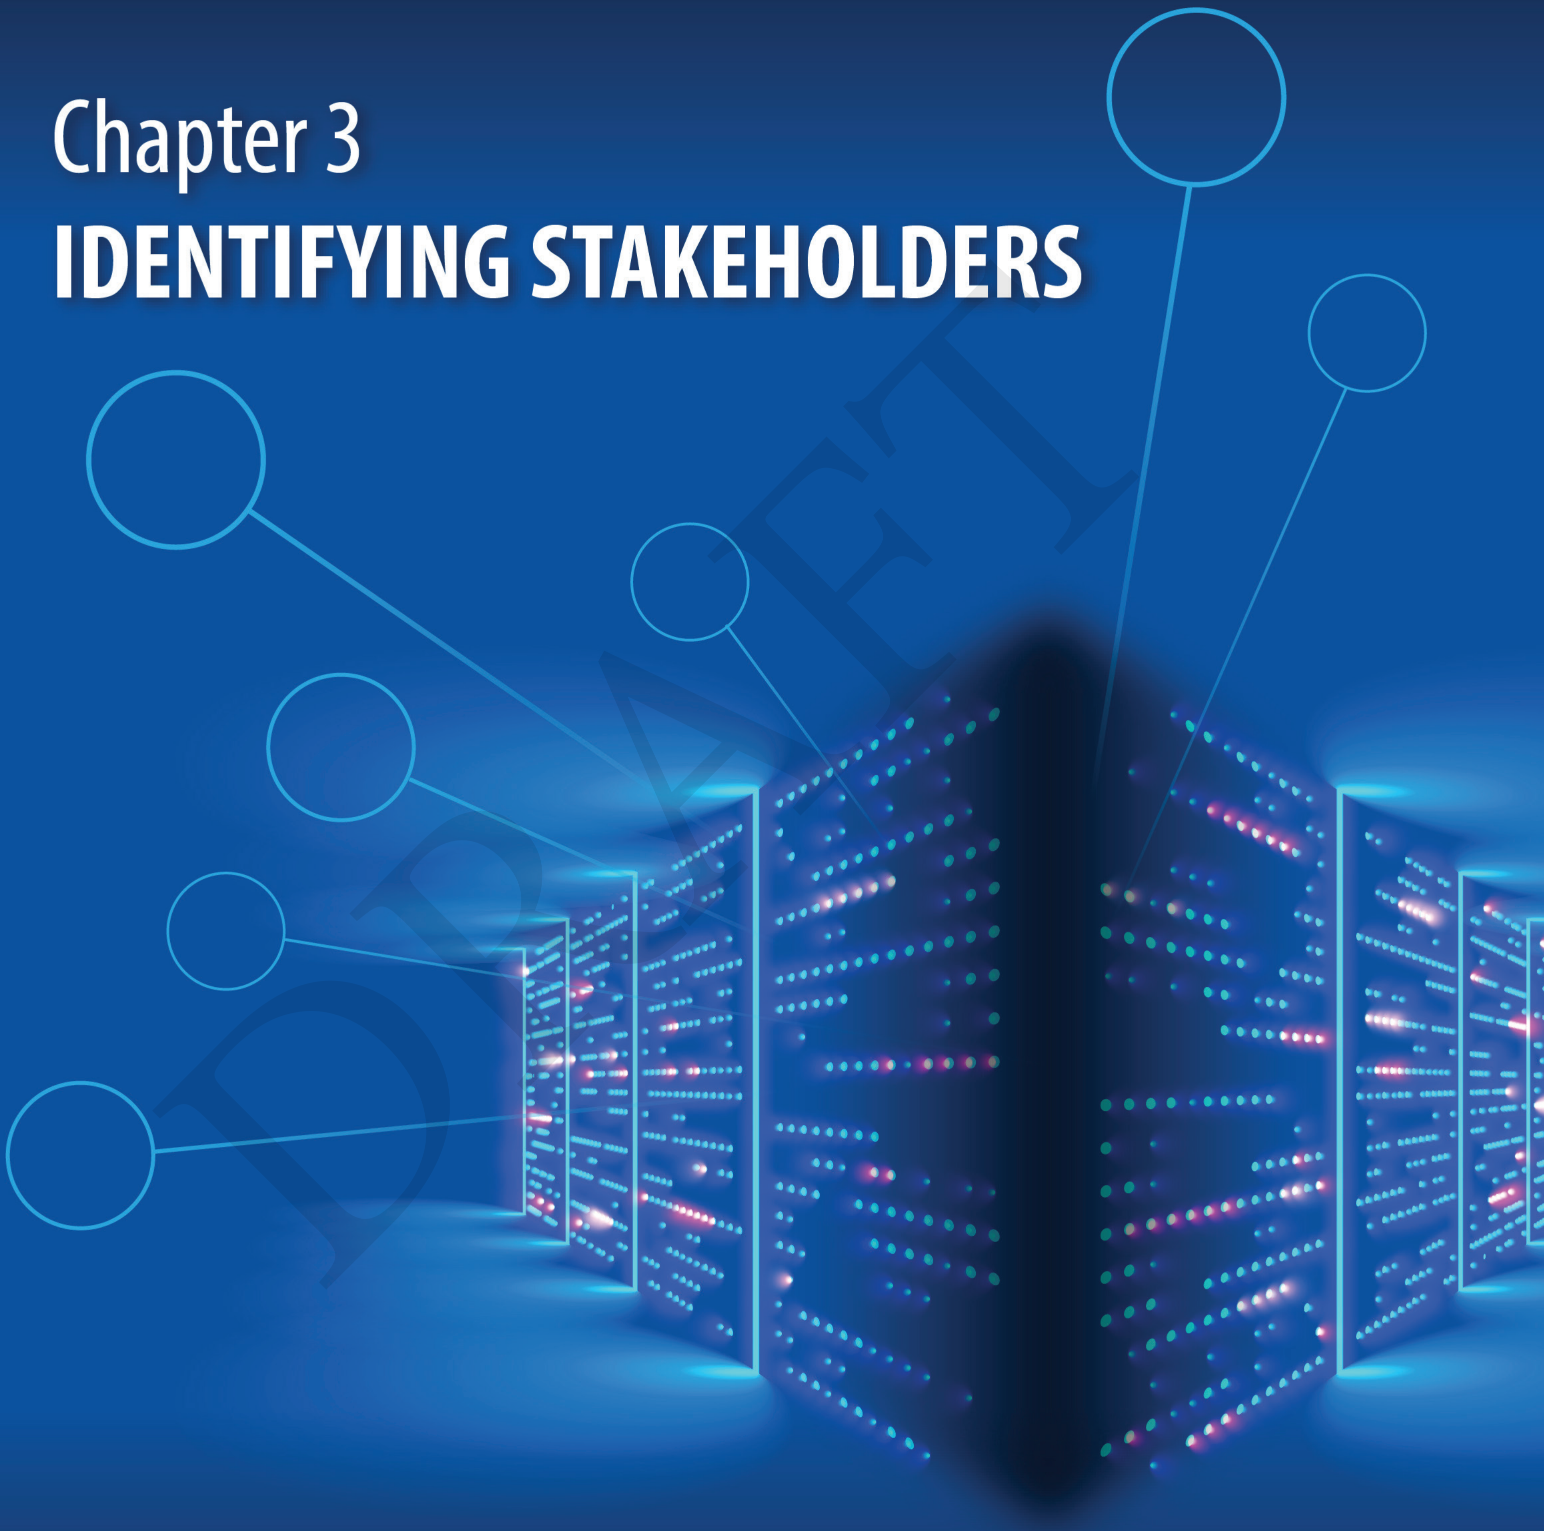

Introduction

Powering a data center with nuclear energy is not a bilateral transaction between a developer and a reactor vendor. It is a multistakeholder undertaking that requires sustained engagement across utilities, engineering and construction partners, nuclear developers, government entities at every level, and a range of financial actors. Identifying the right stakeholders early and understanding their distinct roles, incentives, and constraints is one of the most consequential decisions a data-center developer will make. Misjudging who needs to be at the table, or engaging them too late, is a leading cause of project delays.

This chapter provides an overview of the primary stakeholder categories relevant to nuclear-powered data-center projects, with practical guidance on why each matters and when to engage them.

Why Does This Matter for Data Centers?

Nuclear integration is unlike other power-procurement strategies. A long-term PPA with a gas peaker or an alternative energy developer involves a relatively contained set of counterparties. By contrast, a nuclear-powered data-center project may implicate a regulated utility's integrated resource plan, a state PUC, multiple federal agencies with overlapping jurisdiction, a FOAK engineering challenge, and a financing structure that requires government credit support to be bankable. Each of these relationships carries its own timeline, approval process, and organizational culture. Data-center developers who approach nuclear with the same stakeholder cadence they use for conventional power procurement will find themselves behind schedule and surprised by the complexity.

The good news is that the stakeholder landscape, while broad, is mappable. The sections below identify the key actors, the questions each will need answered, and the points of leverage available to data-center developers.

3.1 Utilities

Electric utilities are often the first institutional stakeholder a data-center developer encounters, and they remain relevant throughout the project life cycle in ways that depend heavily on the market structure in which the project is located.

Regulated (Vertically Integrated) Markets

In states with traditionally regulated, vertically integrated utilities—common across much of the Southeast, Mountain West, and parts of the Midwest—a single utility owns and operates generation, transmission, and distribution assets and recovers costs through rates approved by the state PUC. For a data-center developer in these markets, the utility is simultaneously the interconnection gatekeeper, the potential power supplier, and a regulated actor that must justify any significant new investment to its regulator and ratepayer base.

Engaging a regulated utility around nuclear power requires an understanding of its IRP process. Most regulated utilities file IRPs on 2–4 year cycles, and a large load, like a hyperscale data center, can materially affect those plans. Developers who engage utilities proactively—sharing load projections early, participating in IRP comment processes, and offering to collaborate on demand forecasting—are better positioned to influence how the utility accounts for new nuclear capacity in its planning. Regulated utilities also have established cost-recovery mechanisms, and a data-center developer willing to sign a long-term contract can provide the utility with the demand certainty it needs to justify a capital investment to its regulator.

Deregulated (Restructured) Markets

In deregulated markets, including ERCOT, PJM, and ISO-NE, and parts of the Southwest, generation is separated from transmission and distribution. Power is procured through competitive wholesale markets. In these markets, the utility's role shifts: a transmission and distribution utility manages the wires and handles interconnection while the data-center developer procures power through bilateral contracts (PPAs), direct market participation, or ownership structures.

For nuclear integration in deregulated markets, the critical utility relationship is with the transmission operator managing interconnection. Interconnection queues in major RTOs and ISOs have grown significantly in recent years, with median wait times for large loads now exceeding 4 years in some regions. Data-center developers in these markets should treat interconnection as a long-lead item and engage the relevant transmission operator and any applicable RTO/ISO early and continuously. The interconnection process, timeline, and cost can vary dramatically by region and by the specific configuration of the nuclear facility (BTM, co-located, or grid-connected).

Whether in regulated or deregulated markets, utilities also play a role in resilience planning. Data centers have uptime requirements typically targeting Tier III or IV availability, meaning 99.982–99.995% uptime, that create specific requirements around backup power, redundancy, and grid stability. Understanding how the utility's grid characteristics interact with those requirements is an important early diligence step.

3.2 Engineering, Procurement, and Construction Firms

Engineering, procurement, and construction (EPC) contractors are the organizations that translate a nuclear-powered data-center concept into a physical facility. Their role spans feasibility and front-end engineering design (FEED), equipment procurement, construction management, and systems integration, and their selection and engagement timeline can determine whether a project stays on schedule and on budget.

The Challenging Interface Between Nuclear and Data Centers

Few EPC firms have deep experience at the intersection of nuclear plant BOP engineering and hyperscale data-center construction. These are historically separate disciplines with different safety cultures, contracting norms, and regulatory environments. Data-center construction has historically moved fast—a hyperscale shell can be delivered in 18–24 months—while nuclear construction operates on longer timescales driven by regulatory requirements, quality-assurance programs (including Nuclear Quality Assurance-1 compliance), and the complexity of safety-related systems.

Data-center developers pursuing nuclear integration should expect to engage multiple EPC partners: one with nuclear BOP and civil expertise, and one with data-center infrastructure experience. Understanding where the interfaces between those scopes lie, particularly around electrical interconnection, thermal systems, and control architectures, is a critical early engineering task.

Front-End Engineering and FEED Studies

Before committing to a project, most developers benefit from a structured FEED study that characterizes site conditions, defines the coupling configuration between the NPP and the data center (see Chapter 2), identifies long-lead equipment, and produces a preliminary cost estimate with meaningful confidence bounds. FEED studies for nuclear-adjacent projects are more expensive and time-consuming than for conventional power infrastructure, but they surface project-specific risks, particularly around site hydrology, geotechnical conditions, and interconnection, before they become costly surprises during construction.

Procurement and Long-Lead Equipment

Nuclear projects involve long-lead equipment—reactor pressure vessels, steam generators, large transformers, and specialized electrical switchgear—that must be ordered years before it is needed on site. EPC firms with established nuclear-procurement relationships can provide access to qualified suppliers and help manage the supply-chain risks discussed in Chapter 4. Data-center developers should specifically ask prospective EPC partners about their experience with nuclear Tier 1 and 2 supplier qualification, as well as their relationships with the domestic and international nuclear supply chain.

3.3 Advanced Nuclear Developers

The companies designing and commercializing SMRs and microreactors are stakeholders distinct from traditional large-reactor utilities or EPC firms, and their engagement model with data-center developers is evolving rapidly.

Several advanced-reactor companies have already announced partnerships with hyperscale data-center operators, and more are in discussion. These relationships vary in structure: some involve direct power-offtake agreements, some involve equity co-investment, and others are structured as design partnerships in which the data-center operator provides input into plant configuration (particularly around thermal coupling, discussed in Chapter 2). Data-center developers evaluating advanced-reactor partnerships should assess each developer against several dimensions: regulatory maturity (i.e., has the design received a Nuclear Regulatory Commission [NRC] design certification or combined license application?), fuel supply security, manufacturing readiness, and the credibility of the developer's cost and schedule commitments.

It is also worth noting that the advanced-nuclear developer community is itself a stakeholder that needs the data-center industry's engagement to succeed. Firm, long-term power-offtake commitments from creditworthy data-center operators are among the most-powerful tools available to de-risk advanced-reactor projects and unlock project financing. The relationship is genuinely symbiotic, and developers who approach it as such, rather than as a purely transactional procurement, tend to achieve better commercial outcomes.

3.4 Government

Government engagement for a nuclear-powered data-center project is not a single conversation. It spans multiple agencies at the federal, state, and local levels, each with distinct authorities, timelines, and priorities. Treating government engagement as a compliance exercise rather than a relationship building opportunity is a common and costly mistake.

Federal Government

DOE is the federal agency most directly engaged with the data-center–nuclear nexus. DOE-NE funds much of the research summarized throughout this playbook and administers programs that can support early-stage project development, including cost-shared demonstrations and technical assistance. DOE's Gateway for Accelerated Innovation in Nuclear (GAIN) program provides a structured mechanism for industry to access national-laboratory expertise and facilities. Data-center developers pursuing nuclear integration are encouraged to engage GAIN early, both for technical support and for connections to the broader DOE research ecosystem.

NRC has exclusive federal authority over the licensing and regulation of civilian nuclear facilities in the U.S. For data-center developers, the most important NRC engagement questions are the licensing pathway that applies to the specific reactor technology under consideration, and a the realistic timeline. The NRC offers pre-application engagement, such as meetings, white-paper reviews, and regulatory-audit programs that can significantly reduce uncertainty before a formal application is submitted. Developers should expect licensing timelines of 5–10 years for novel reactor designs under current frameworks although the NRC actively works to streamline reviews for advanced-reactor technologies under The Accelerating Deployment of Versatile, Advanced Nuclear for Clean Energy Act of 2024.

FERC has jurisdiction over wholesale electricity markets, interstate transmission, and the interconnection of generation facilities in most of the U.S. (Texas and certain non-interstate systems excepted). Data-center projects with grid-connected nuclear configurations will encounter FERC's interconnection rules, and developers operating in RTO/ISO markets should monitor FERC's ongoing interconnection reform proceedings, which are reshaping how large loads and generators access the grid.

EPA and **USACE** are together involved in environmental permitting, particularly for water intake and discharge associated with nuclear-plant cooling. Section 404 of the Clean Water Act (USACE jurisdiction) and Section 316 of the Clean Water Act (EPA jurisdiction over thermal discharge) are the most-commonly encountered authorities. Early coordination with these agencies, informed by the water-energy nexus considerations discussed in Chapter 2, can avoid late-stage permitting delays.

State Government

State governments play a pivotal role in nuclear-powered data-center siting, permitting, and economic development. Their influence operates through several channels:

PUCs, in regulated states, must approve any significant changes to a utility's resource plan, including the addition of new nuclear generation contracted to serve large loads. Data-center developers should understand the PUC's procedural requirements and engage early with utility partners on how a nuclear project will be presented to the commission.

State energy offices and economic development agencies have established programs in many states to attract data-center investment, including property-tax abatements, sales-tax exemptions on equipment, and expedited permitting processes. Some states, particularly those with existing nuclear fleets or active advanced nuclear programs, have also established state-level incentives specifically for nuclear energy development. Mapping the incentive landscape in target states early in site selection can meaningfully affect project economics.

Environmental and water permitting authorities in the states administer many of the environmental permits relevant to both nuclear facilities and data centers, including water-withdrawal permits, air-quality permits, and state environmental policy act reviews. As noted in Chapter 2, water availability is often the binding constraint on siting, and state water law, which varies significantly across regions, governs how water rights are allocated and transferred.

Local Government

The influence of **county commissions, municipal authorities, and regional planning bodies** over large energy and data-center projects is frequently underestimated. Zoning approvals, local permitting, road-use agreements for construction traffic, and tax-increment financing arrangements are all within local government authority. More importantly, local elected officials and community leaders shape the public perception environment in which a project operates, a topic addressed in Chapter 4.

Data-center developers with experience in conventional siting know that community opposition can derail even well-capitalized projects. Nuclear projects carry additional public-perception complexity, and early, transparent engagement with local officials, especially before formal applications are filed, is strongly recommended. Communities adjacent to proposed nuclear sites will have questions about safety, jobs, tax revenue, and water use. Developers who arrive with answers and who engage in genuine dialogue, rather than one-way communication, build the durable local relationships that sustain projects through multi-year development timelines.

3.5 Financing

The capital intensity of nuclear-powered data-center projects that combine two of the most capital-intensive industries in the modern economy requires a financing strategy that goes beyond conventional corporate debt or equity. A well-structured financing approach draws on public-sector credit support, tax-incentive programs, and private capital structures designed to share risk appropriately among parties.

Department of Energy Office of Energy Dominance Financing

The DOE Office of Energy Dominance Financing (EDF) is among the most-consequential federal tools available to advanced nuclear and energy-infrastructure developers. Under its Title XVII Innovative Energy Loan Guarantee Program, DOE-EDF can provide loan guarantees for projects that employ innovative advanced energy technologies—a category that explicitly includes advanced nuclear reactors. Loan guarantees from EDF can unlock financing at rates and tenors unavailable in the private market alone, and they carry an implicit federal endorsement that can improve the project's profile with commercial lenders and equity investors.

EDF's engagement process is rigorous, but the agency has made significant investments in building its nuclear expertise and has signaled strong interest in supporting nuclear-powered data-center projects. Early pre-application engagement with EDF staff is strongly encouraged because the agency can provide guidance on eligibility, structuring, and the information it will require to evaluate a project.

Federal Tax Incentives

The Inflation Reduction Act of 2022 introduced and expanded several tax incentives relevant to nuclear-powered data-center projects:

- The **Production Tax Credit (PTC) for nuclear** (Section 45U of the Internal Revenue Code) provides a per-kilowatt-hour tax credit for electricity generated by existing nuclear facilities. While this credit applies to existing plants rather than new construction, it is relevant to data-center developers pursuing PPAs with operating nuclear plants or capacity-expansion projects at existing sites.
- The **Investment Tax Credit (ITC)** and **Advanced Energy Project Credit** (Section 48C) may apply to certain advanced nuclear components and qualifying manufacturing investments associated with the project. The applicability of these credits is fact-specific and depends on the technology, ownership structure, and project configuration.
- **Transferability and direct pay provisions** in the Inflation Reduction Act (IRA) allow tax credits to be monetized by parties that cannot use them directly, which presents a significant structural tool for project finance transactions involving tax-exempt entities or investors with limited tax capacity.

Data-center developers should engage tax counsel with energy-project-finance experience early in project development to map the applicable credit landscape and structure the project to capture available incentives.

Power Purchase Agreements and Offtake Structures

Long-term PPAs are among the most important tools available to data-center developers pursuing nuclear power. A creditworthy, long-term offtake commitment, typically 15–25 years for nuclear projects, provides the revenue certainty that lenders and equity investors require to underwrite project finance. For data-center companies with investment-grade credit ratings, a direct PPA with a nuclear developer or utility can serve as the anchor around which the rest of the project's financing is assembled.

BTM configurations (Options 2–5 in Chapter 2) create additional structural options: the data-center operator can participate as an equity co-investor in the nuclear facility, capturing both the economic upside and the security of onsite power while the nuclear developer benefits from a committed anchor customer.

Public-Private Partnerships

Several states and municipalities have expressed interest in P3 structures for energy infrastructure that serves both public and private needs. In the context of nuclear-powered data centers, P3 structures could take several forms: a state economic-development authority co-investing in site infrastructure, a municipal utility participating as a minority power-offtake agent alongside the data-center operator, or a regional economic-development compact supporting workforce training and community infrastructure around a new nuclear site.

P3 structures require patient capital and complex multi-party negotiation, but they can unlock public resources—including state revolving funds, federal economic-development grants, and tax-exempt bond financing—that are not available to purely private projects. Developers operating in regions where state and local governments are actively competing for data center and energy investment are well-positioned to explore these structures.

Additional Information Sources

The stakeholder landscape for nuclear-powered data centers is evolving rapidly, and the following resources provide additional context:

- DOE GAIN: https://gain.inl.gov
- DOE-EDF: https://www.energy.gov/EDF
- NRC Pre-Application Engagement: https://www.nrc.gov/reactors/new-reactors/pre-application-activities.html
- FERC Interconnection Reform resources: https://www.ferc.gov.

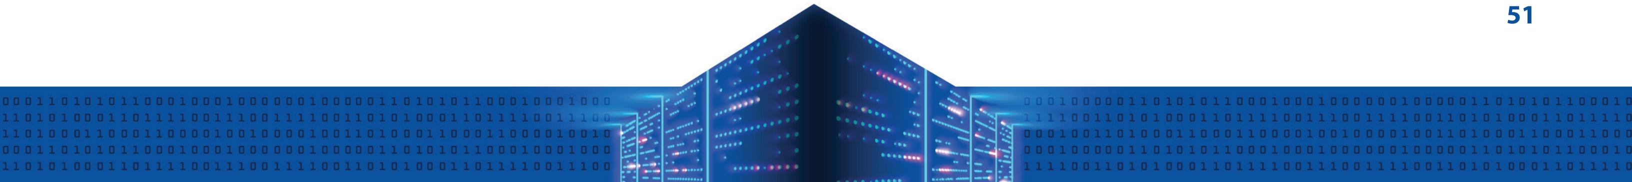

Chapter 4

ROAD TO NUCLEAR

Nuclear energy is increasingly recognized as one of the few power sources capable of meeting the scale, reliability, and longevity requirements of the most demanding data-center applications. But what does it actually take to get there, and where does a developer begin? Recognizing the potential and realizing it are very different things. The road to nuclear may be technically complex, and shaped by forces that extend well beyond the reactor itself, from supply chains and site conditions to public perception and project finance. This chapter is intended to serve as a practical guide to that road, covering the full range of considerations a data-center developer is likely to encounter.

The chapter opens with the near-term opportunity. The section on **Capacity Expansion of Existing Power Plants** examines how power uprates at currently-operating nuclear facilities can add firm capacity in 3–7 years—faster than any new-build option—and what site-specific factors determine whether a given plant is a viable candidate. For developers with a longer horizon, the section on the **Advanced Reactor Landscape** surveys the range of designs currently under development, from large light-water reactors (LWRs) to SMRs and beyond, with attention to how differences in fuel type, modularity, and design maturity translate into practical differences in deployment timeline and operational reliability. The section on **New-Build Plant Costs** then addresses the economics of nuclear construction directly, including the significant uncertainty that surrounds cost projections for FOAK designs and the factors—learning rates, economies of scale, financing structure, and incentives—that most influence the all-in cost of power.

The remaining sections broaden the lens. **Nuclear Integration: Dynamics of Policy and Public Perception** examines the federal, state, and local regulatory environments that shape where and how nuclear projects can be developed, as well as the community engagement strategies that are increasingly recognized as essential to project success. **Nuclear Supply Chain Considerations** maps the equipment, fuel, and workforce dependencies that nuclear integration introduces, including areas of overlap with data center supply chains that can create both efficiencies and competition. The section on **Nuclear Plant Reliability and Availability** translates the nuclear industry's well-established reliability metrics into terms relevant to data center uptime requirements. The section on **Site Planning and Selection** draws on a national siting analysis to identify where co-located nuclear and data-center development is most feasible, with particular attention to the role of water availability as a binding constraint. Finally, the section on **Workforce and Socioeconomic Impact** examines the employment, labor income, and broader economic contributions that a co-located campus can generate for its host community, an asset that is increasingly relevant to permitting strategy and stakeholder relations alike.

4.1 Capacity Expansion of Existing Power Plants

Introduction

A nuclear power uprate refers to the process of increasing the licensed thermal power of an existing power plant, which allows for an increase in electricity production. The uprate process is mature; since 1977, the NRC has approved 171 uprates, adding roughly 8 GW of capacity—comparable to eight new large reactors.[69] In addition to uprates, power increases can be achieved through thermal performance upgrades, which improve the efficiency with which a plant converts thermal energy into electricity. This section discusses the different types of power uprate, the potential for future power uprates and thermal performance upgrades, and the anticipated cost and timeline.

Why Does This Matter for Data Centers?

For data-center developers seeking nuclear power, expanding output at existing nuclear plants is one of the fastest pathways to additional firm, always-on nuclear capacity. Power uprates can add meaningful capacity in 3–7 years depending on the scope of the uprate and regulatory timeline, in contrast to the deployment of a new advanced reactor, which may require up to 10–15 years. Uprates are also generally less expensive on a per-kWe basis than new builds, since major civil and structural infrastructure is already in place. Further, uprate projects typically carry lower regulatory and construction risk, since the plant is already licensed and operating. As a result, power uprates are an important option for data-center developers to consider as part of a broader nuclear power strategy, particularly for those seeking to add firm capacity on a shorter timeline than new builds allow.

Research and Findings

A nuclear plant's licensed thermal power, the highest power level at which it is allowed to operate, is regulated by the NRC. Operators can submit an application to the NRC requesting that the licensed thermal power of an existing plant be increased, so long as plant safety is maintained. The NRC defines three different types of power uprate, which vary in impact and complexity:

- **Measurement Uncertainty Recapture (MUR) Power Uprates:** MURs recover conservatism embedded in reactor-power measurements by using higher-accuracy instrumentation and analysis methods. Projects generally require minimal hardware changes, have the shortest implementation timelines, and result in a licensed thermal power increase of 1.5–2%.
- **Stretch Power Uprates (SPU)**: SPUs achieve moderate power increases (2–7%) through targeted modifications, typically on the secondary side of the plant: e.g., turbines, feedwater systems. SPUs involve modest capital investment and limited BOP adjustments.
- **Extended Power Uprates (EPU):** EPUs enable the largest output increases (typically up to 20%), but require significant upgrades to major plant systems, such as steam generators or turbines. EPUs often deliver 200–400 MW per unit and involve multiyear engineering, procurement, and installation efforts.

There is also potential for power increase through thermal performance upgrades, which can be achieved through refurbishment or replacement of major equipment in the BOP. This differs from power uprates in that power increases are achieved through Rankine cycle efficiency gains rather than increases in licensed thermal power. As such, NRC approval is not required, which may allow for shorter timelines compared to uprates.

Uprate Potential and Cost

While much of the U.S. fleet has already undergone power uprates, there is remaining potential for future uprates. The amount that a plant can be uprated is site specific and will be dependent on many factors such as prior uprates, design constraints, regulatory approvals, supply-chain considerations, and sustained multiyear project execution. Determining the amount which a specific plant can be uprated will require detailed engineering and financial analysis. Further, the decision to uprate a plant and the extent of the uprate will depend on factors such as the cost, remaining plant lifetime, market forecast, and perceived risk.

Figure 16 and Figure 17 show the amount of power already gained through past uprates and an estimate of the amount of untapped power available through further uprates. For the boiling-water reactor (BWR) fleet, it is estimated that the licensed thermal power can be increased to 130% of the original value, which would result in approximately 5 GW of additional power produced. For the PWR fleet, it is estimated that licensed thermal power can be increased to 125% of the original licensed thermal power, which would result in approximately 12 GW of power.

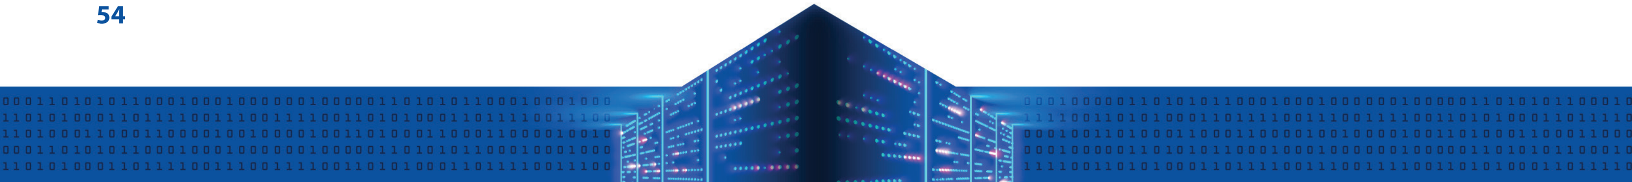

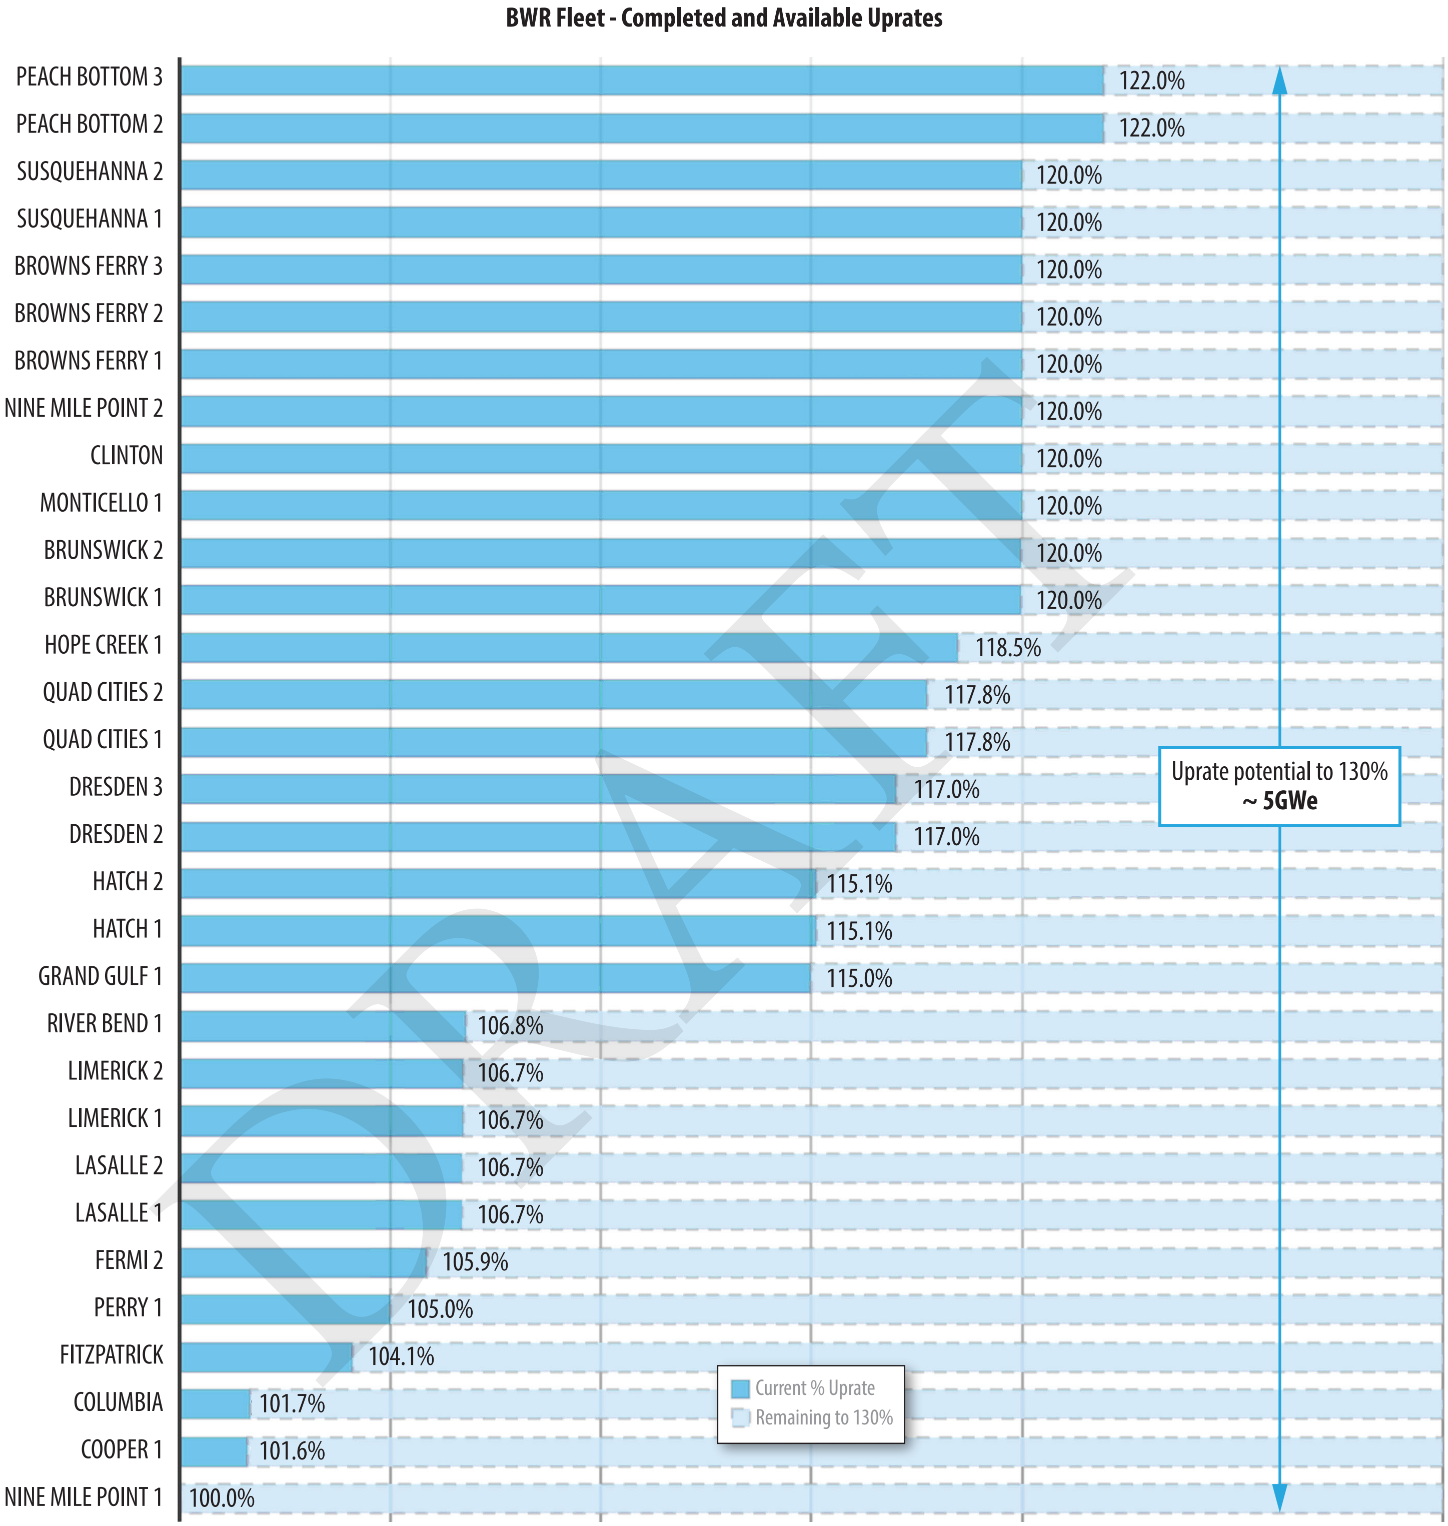

Figure 16. Potential uprate capacity for BWRs.[70]

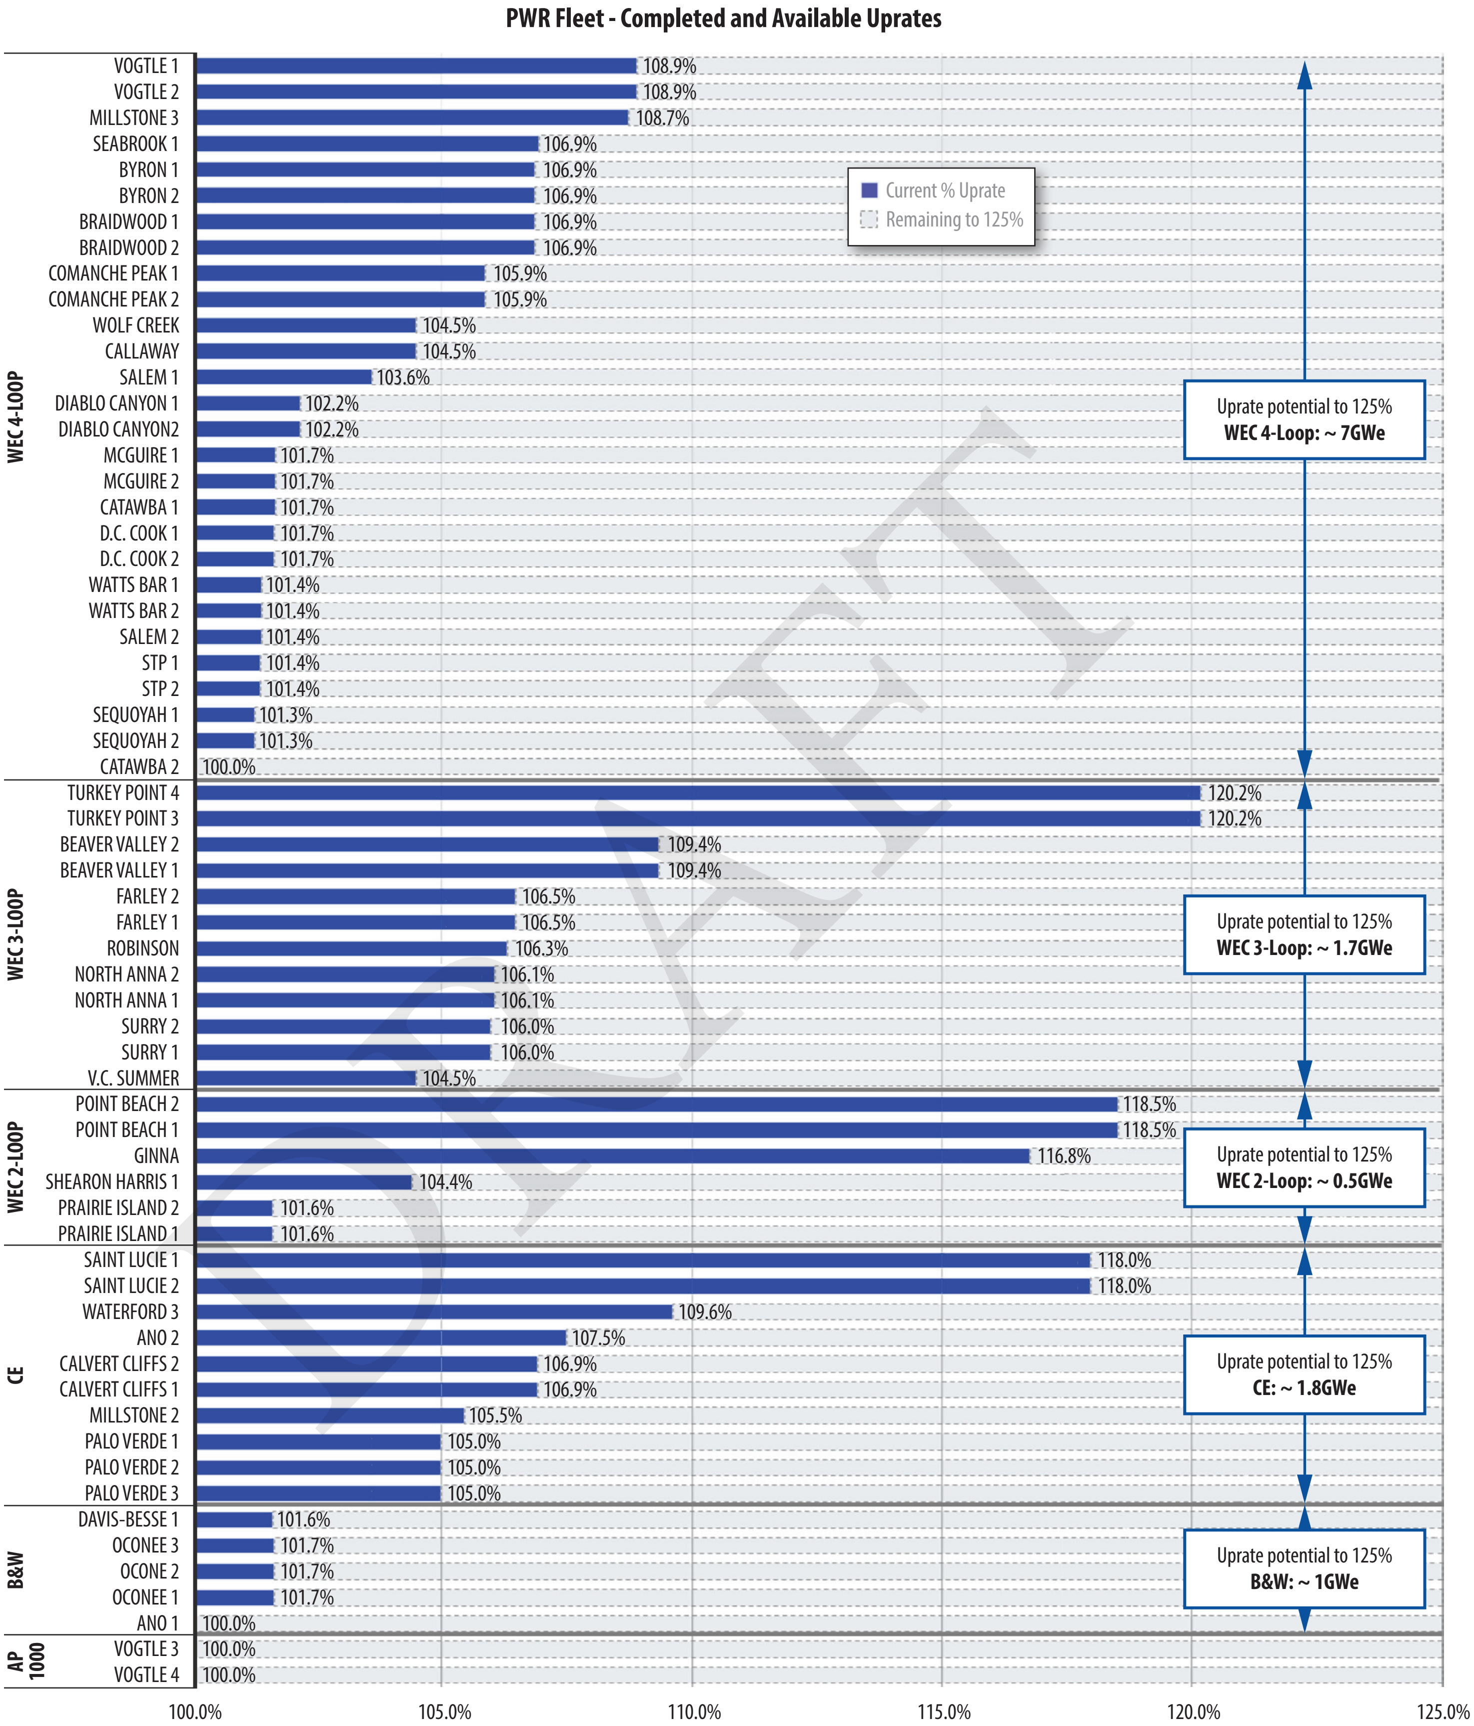

Figure 17. Potential uprate capacity for PWRs.[70]

The cost of a power uprate will depend on the type of uprate and the modifications required. For any uprate being considered, a feasibility study should first be conducted, which will include a detailed analysis of the plant to determine the modifications required to achieve a given power increase and the cost associated with those modifications.[71]

The difference between a straightforward uprate and a high-risk one is rarely the reactor island; the risk can come from the surrounding plant and site. Three constraints dominate feasibility and cost: cooling and heat rejection (condenser performance and water/ambient limits), switchyard and interconnection capacity, and BOP condition (how much equipment must be replaced, rather than modified). The same site-specific factors that shape scope also drive a wide cost distribution. The cost of most uprate projects have not been shared publicly, though the cost of some EPU projects has been published, which are shown in Table 12. Costs for future uprates will be plant-specific and should be evaluated on a case-by-case basis.

Table 12. Representative EPU Costs (2023 USD).[72]

Plant	Plant Type	Year	Capacity Added	Cost ($/kW)
Browns Ferry (3 units)	BWR	2019	465 MW	$1,200
Monticello	BWR	2013	71 MW	$10,350
Grand Gulf	BWR	2012	178 MW	$6,440
St. Lucie Unit 1	PWR	2013	145 MW	$7,570
St. Lucie Unit 2	PWR		132 MW	
Turkey Point Unit 3	PWR		~120 MW	
Turkey Point Unit 4	PWR		~120 MW	
Brunswick Unit 1	BWR	2002	116 MW	$1,250
Brunswick Unit 2	BWR		117 MW	

Capacity Outlook

Uprates can move faster than new plants can be built, but near-term additions are still constrained by project timelines. Pathways to delivering 1.5–2.3 GW by 2030 from licensed uprates and thermal-performance improvements have already been identified, rising to 7–10 GW total by 2035 with sustained investment, and these projections are evolving rapidly. These timelines may be accelerated with targeted actions to align policy goals with action-oriented financial incentives. Ongoing efforts at DOE national laboratories are investigating the extent to which timelines can be accelerated, and at what cost. Developers can improve outcomes by engaging early and by structuring commitments around gated milestones. Prioritizing the following diligence items can mitigate project risks and improve outcomes:

- Contract around milestones (e.g., license-amendment request submitted and approved; outage complete) and define uprate scope and life-cycle management work
- Validate cooling and switchyard/interconnection constraints early
- Stress-test the outage plan and understand the utility's regulatory and cost-recovery pathway.

The timeline for capacity expansion will depend on the chosen pathway. Thermal performance upgrades are the fastest option, typically achievable in 3–4 years, because they do not require NRC license amendments. However, capacity gains are limited, generally 50–150 MW per unit. MUR and SPU projects can also deliver additional power quickly, but they provide a smaller impact in generation capacity. Large EPU projects will likely take longer than 6 years to complete, but sustained investment and timely project initiation can deliver large generation gains by 2035.

Further Exploration

- EPRI. July 2023. "Facilitating Power Uprates at Nuclear Power Plants: Feasibility Study Guideline."
- Larsen, L. M. et al. 2023. "Assessing the impact of the Inflation Reduction Act on Nuclear Plant Power Uprate and Hydrogen Cogeneration." Idaho National Laboratory, Idaho Falls, ID.
- International Atomic Energy Agency, "Power Uprate in Nuclear Power Plants: Guidelines and Experience" Vienna, 2011.

4.2 Advanced-Reactor Landscape: What Data Center Developers Need to Know

Introduction

The 94 operating reactors in the U.S. are all LWRs: either PWRs or BWRs. The operational characteristics of the existing fleet are well known and based on over 50 years of operating experience. Additionally, the supply chains needed to service, maintain, and fuel these reactors are well established. The new generation of advanced reactors is expected to include a range of advanced-reactor technologies, including large LWRs, LWR SMRs, HTGRs, SFRs, and molten salt reactors (MSRs). Unlike the existing fleet, these designs vary in terms of coolants, fuel types, enrichment levels, and power outputs. While these design differences may not directly matter to a data-center developer whose primary need is reliable electricity, they have significant downstream consequences for the factors that do matter: deployment timeline, cost, and reliability and availability. This section discusses the advanced-reactor landscape, focusing on the characteristics that vary across reactor types and are most likely to affect deployment timelines, cost, and reliability for data-center applications.

Why Does This Matter for Data Centers?

Data-center developers considering nuclear should understand the differences between advanced-reactor designs and how these differences can impact deployment timelines, cost, and reliability and availability for data-center applications. This is particularly important for data-center developers considering a dedicated nuclear plant for BTM power.

Research and Findings

Many different advanced-reactor designs are being actively developed; all of these can serve as a reliable source of abundant energy for data-center applications. However, these reactors vary in terms of design characteristics, design maturity, supply-chain readiness, and other factors, which can lead to differences in deployment timelines, cost, and operational reliability.

To highlight the variation in advanced-reactor designs, Table 13 below lists advanced-reactor designs that have been deployed or are currently under development, including designs that have reached the pre-application licensing stage or and equivalent stage in the licensing process in the United Kingdom (UK), excluding microreactors.

Key Differences Between Advanced Reactor Designs

Table 13 illustrates the breadth of advanced-reactor designs under development and highlights the key characteristics that vary across reactor types. The following discusses how these characteristics translate into practical considerations for data-center developers.

Power Output and Modularity: For any reactor that can be deployed with multiple units, the amount of power produced per reactor will have an impact on availability and the optimal microgrid design. The backup power requirements will vary depending on the power output and number of modules because units will undergo outages for refueling, maintenance, or other potential issues. More units may provide more resilience by enabling units to be taken offline in stages, rather than needing to take the entire plant offline at the same time.

Power output and modularity will also influence capital costs of the plant. While it is difficult to project the cost of specific reactor types, general trends related to the differences between large and small reactors are known. For example, the per-kilowatt cost for large reactors is expected to be less than for SMRs due to economies of scale while SMRs are expected to benefit from faster learning and cost reductions due to economies of multiples.[38] Additionally, SMRs are expected to have shorter construction times than do large reactors due to the potential for factory fabrication and concurrent manufacturing, which can decrease the amount of interest accrued during construction and reduce overall capital expense.

Fuel Enrichment and Fuel Type: The existing fleet benefits from a well-established fuel supply chain; however, new advanced-reactor designs, particularly of non-LWRs, may require different supply chains that do not yet exist. The existing fleet uses low-enriched (less than 5% enrichment) uranium (LEU) fuel while many advanced-reactor designs will require high-assay, low-enriched uranium (HALEU), which is enriched to levels close to 20%. Currently neither this enrichment infrastructure nor a HALEU supply chain exists in the U.S. The DOE has awarded funds to companies to stand up this capability,[73] but when this will be commercially available at scale to support a significant rollout of advanced nuclear reactors requiring HALEU is unclear.

Table 13. Advanced reactor designs and characteristics.

Large Light Water Reactor

Reactor Type	Coolant	Fuel Type	Fuel Enrichment	Design	Developer	Power Output (MWe)	Licensing Stage
PWR	Water	Uranium Oxide	LEU	AP1000	Westinghouse	1,120 MWe	Several units deployed

Small Modular Reactor

Reactor Type	Coolant	Fuel Type	Fuel Enrichment	Design	Developer	Power Output (MWe)	Licensing Stage
PWR	Water	Uranium Oxide	LEU	SMR-300	Holtec	320 MWe	Pre-Application and Limited Work Authorization for Pioneer 1 & 2 (Michigan) under review
				NuScale Power Module	NuScale	77 MWe	Standard Design Approval Application Approved
				Rolls Royce SMR	Rolls-Royce	470 MWe	UK Generic Design Assessment Step 3
				AP-300	Westinghouse	300 MWe	Pre-Application
BWR	Water	Uranium Oxide	LEU	BWRX-300	General Electric-Hitachi	300 MWe	Construction Permit Application for TVA Clinch River Site under review
HTGR	Helium	TRISO	HALEU	Xe-100	X-energy	320 MWe	Construction Permit Application under review
	Helium	Uranium Carbide	HALEU	Energy Multiplier Module	General Atomics	265 MWe	Pre-Application
SFR	Molten Sodium	Metallic Uranium Zirconium	HALEU	ARC-100	ARC Clean Energy	100 MWe	Pre-Application
				Aurora	Oklo	75 MWe	Pre-Application
				Natrium	TerraPower	345 MWe	Construction Permit Application Granted
MSR (Salt-Cooled)	Molten Salt	TRISO	HALEU	KP-FHR	Kairos	140 MWe	Pre-Application. Have two construction permits approved for demonstration facilities
MSR (Salt-Fueled)	Molten Salt	Molten Salt	HALEU	Integral Molten Salt Reactor	Terrestrial Energy	195 MWe	Pre-Application
			LEU	MSR-100	Natura Resources	100 MWe	Pre-Application

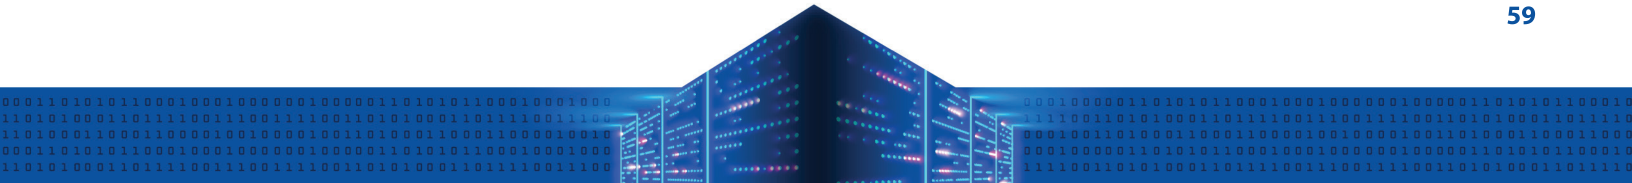

In addition to differences in enrichment, some of these new reactors may require different fuel forms than do LWRs. These newer fuel forms may not be fully scaled for commercial production. While the U.S. has deployed metallic fuels for SFRs in the past, no current commercial demand exists for metallic SFR fuels. TRISO fuel does have some commercial capability, but fuel-fabrication facilities and infrastructure are insufficient to support large-scale commercial operations. The enrichment and fuel-fabrication infrastructures could play key roles in commercially scaling these technologies on a timeframe that is needed for data center buildouts. Additionally, the fuel supply chain must be vetted to understand fuel availability and future pricing of fuel around the life-cycle cost of the plant. If a fuel unique in enrichment or fabrication is employed, the long-term commercial viability of that fuel supply chain must be established to ensure future operations of the reactor.

Refueling Method and Maintenance Strategy: Traditional reactor types and most advanced-reactor types refuel offline, meaning that they must periodically power down for refueling. At a periodicity of 18–24 months, the existing fleet of LWRs typically refuel their reactor cores, which requires a reactor to be shut down to take care of refueling and other normal inspection and maintenance activities. Refueling outages typically last a month or less, but during this time, power is unavailable. Some newer advanced-reactor types, such as pebble-bed reactors, can be refueled online, which is advantageous for data-center applications because they have no need to shut down for refueling. However, the need to perform inspection and maintenance activities will periodically require the reactor to be shut down, so the maintenance strategy and plan will influence the availability of the plant.

Design and Operational Maturity: To date, the non-LWR advanced-reactor designs have not been demonstrated. The U.S. DOE has awarded funds through the Advanced Reactor Demonstration Program[74] to support commercial demonstrations of some designs, like TerraPower's Natrium or X-energy's Xe-100. Others are receiving support from the DOE through the reactor pilot program to provide a DOE authorization pathway to demonstrate and de-risk these systems prior to their commercial scale up.[75] That these designs have not yet been demonstrated in these configurations means that there is some risk to commercial scale up. This risk could manifest in a true design risk that leads to a reactor's not being commercially viable, or it could be on the operations-risk side, where unknown conditions lead to a reactor that runs, but may not be as reliable in operations. It would thus be more costly to operate than traditional reactors. Understanding the technology maturation plan will help stakeholders comprehend potential timelines of availability and risk of certain technologies.

While the new advanced reactors may boast certain unique features, the maturity of the design and operational characteristics may mean that operational learning is happening in a commercial setting. LWR designs are well tested, with significant operational knowledge to drive high-capacity factors with limited unplanned maintenance. The newer advanced designs may not yet offer the reliability of LWR performance until they operate and develop the operational processes that have made the LWR fleet the backbone of electricity generation with greater than 90% capacity factors.

Commercial Market Penetration: A further consideration is the commercial maturity of reactor technology and its supporting ecosystem; the current fleet of LWRs benefits from a high level of commercial market penetration, with a supply chain that services the needs of that specific reactor technology. Newer designs may not yet have achieved the market penetration needed for additional infrastructure to develop around them. Data-center developers should consider the overall risk of a limited vendor that will service, fuel, and maintain advanced nuclear designs that have not yet achieved significant commercial deployment because these chokepoints can affect both the timeline to bring a plant online and its long-term operability.

Implications for Data Center Developers

The advanced-reactor landscape offers data-center developers a range of technology options, each with distinct characteristics that affect deployment timeline, cost, reliability, and availability. Selecting among reactor technologies involves more than comparing power outputs and price quotes; the design characteristics discussed in this section have real consequences for when power can be delivered, its cost over the plant's lifetime, and the reliability and availability generation.

The deployment timeline—the speed at which a nuclear plant can begin supplying power to a data center—will be influenced by factors such as design completion, supply-chain maturity, fuel availability, partnership agreements, licensing, and expected construction durations. The deployment timeline should also consider how risks identified in earlier demonstrations are incorporated into the final commercial design such that risks can be managed before building the commercial design.

Cost will also vary across reactor designs. Capital cost is driven primarily by reactor size, with large reactors benefiting from economies of scale and SMRs offering potential cost reductions through economies of multiples and shorter construction timelines. Beyond capital costs, the long-term cost of operating a nuclear plant is shaped by the maturity of its fuel supply chain, the depth of the commercial ecosystem available to service it, and the operational track record of the design—all of which are less certain for new designs than for established LWR technologies.

Reliability and availability depend on power-unit modularity, the refueling and maintenance strategy, and design and operational maturity. The frequency and duration of scheduled outages will depend on refueling and maintenance strategies while power-unit modularity will determine the fraction of NPP that is unavailable during these outages. Unscheduled outages are more difficult to predict for non-LWR designs, which are less mature and do not have the benefit of decades of operating experience. Data-center developers who are early customers of less-mature designs should consider the impact of design maturity on plant reliability and plan backup capacity accordingly.

4.3 New-Build Plant Costs

Introduction

As data-center developers evaluate nuclear energy for their long-term power needs, understanding how to estimate and contextualize reactor costs is essential. However, due to the early stage of advanced nuclear reactors, the economic considerations have a high degree of uncertainty with many factors to consider.

Nuclear is uniquely capital intensive. Unlike natural gas or other alternative energies, for which fuel or equipment costs are spread over time and tied to market prices, the majority of the lifetime cost of nuclear plants is determined upfront by the cost to build the plant. This means that getting the cost estimate right and understanding the uncertainty around it matters enormously. In addition to using line-item estimates (such as bottom-up cost estimates), estimating reactor costs for data-center applications should include a transparent set of modifiers that reflect scale, delivery method, learning, financing, and site integration. This method moves beyond single-point values and encompasses many factors on a micro- and macroeconomic scale. The work completed by Abou-Jaoude et al.[76] presents ranges of advanced-reactor cost estimates in this fashion, taking bottom-up cost estimates and adding several economic assumptions to produce an informed set of cost ranges. This work also documents how multi-unit plants, construction timelines, lifetime, and ramp rates affect cost and performance, all of which are relevant to hyperscale computing campuses.

This section discusses the main cost-differentiating contributors that should be considered when performing nuclear energy planning for data centers. This discussion will include learning (cost reductions over time), economies of scale, construction risk and overruns, standardization, and financing and incentives.

Why Does This Matter for Data Centers?

For data-center developers evaluating nuclear as a long-term power source, plant economics are unavoidably central to the decision. Current advanced-reactor designs have not yet been demonstrated or built at commercial scale, so cost projections rely on engineering estimates and modeling rather than observed data. As a result, estimates vary widely and carry substantial uncertainty. A developer who cannot critically evaluate a vendor's cost projections, or who treats a single-point estimate as reliable, risks making a multi-decade, multi-billion-dollar capital commitment based on assumptions that may not stand up over time.

Research and Findings

Cost Estimation

The overnight capital cost (OCC) is the upfront cost to build a nuclear plant, excluding financing cost. This can include pre-construction costs, direct costs, indirect services costs, personnel costs from before the commercial operation date, and supplementary costs. A meta-analysis of advanced-reactor cost estimates analyzed a number of historic estimates to project the cost of new nuclear builds.[76] This report analyzed historic bottom-up reactor cost estimates, both OCC and O&M costs, to develop projections for future advanced-reactor costs. The report separated large reactors (~1000 MWe) and SMRs (~300 MWe) and provided cost estimates with varying levels (advanced, moderate, and conservative) of conservatism, as shown in Table 14. The cost estimates are not specific to any reactor type or technology and are based on a variety of estimates, including HTGRs, LWRs, and SFRs. The estimates are normalized per kilowatt of electricity produced; they therefore scale with reactor size.

Table 14. OCC for large NPPs and SMRs.[76]

Account	Title	Large Reactor (~1,000 MWe)			Small Modular Reactor (~300 MWe)		
		Advanced ($/kWe)	Moderate ($/kWe)	Conservative ($/kWe)	Advanced ($/kWe)	Moderate ($/kWe)	Conservative ($/kWe)
10	**Capitalized Pre-Construction Costs**	$369	$404	$545	$280	$394	$509
11	Land and Land Rights	$3	$3	$4	$11	$16	$20
12	Site Permits	$19	$21	$28	$7	$9	$12
13	Plant Licensing	$182	$199	$269	$134	$189	$245
14	Plant Permits	$20	$22	$29	$16	$22	$29
15	Plant Studies	$28	$31	$41	$22	$31	$40
16	Plant Reports	$14	$15	$20	$11	$16	$21
17	Community Outreach and Education	$0	$0	$0	$0	$0	$0
18	Other Pre-Construction Costs	$61	$67	$91	$45	$63	$81
19	Contingency on Pre-Construction Costs	$42	$46	$62	$34	$48	$62
20	**Capitalized Direct Costs**	$3,441	$3,769	$5,080	$3,553	$5,007	$6,461
21	Structures and Improvements	$849	$930	$1,253	$1,099	$1,549	$1,999
22	Reactor System	$695	$761	$1,026	$796	$1,121	$1,447
23	Energy Conversion System	$206	$226	$304	$212	$299	$386

Account	Title	Large Reactor (~1,000 MWe)			Small Modular Reactor (~300 MWe)		
		Advanced ($/kWe)	Moderate ($/kWe)	Conservative ($/kWe)	Advanced ($/kWe)	Moderate ($/kWe)	Conservative ($/kWe)
24	Electrical Equipment	$332	$364	$491	$520	$733	$946
25	Initial Fuel Inventory	$118	$130	$175	$165	$233	$300
26	Miscellaneous Equipment	$118	$130	$175	$165	$233	$300
27	Material Requiring Special Consideration	$325	$356	$480	$0	$0	$0
28	Simulator	$0	$0	$0	$0	$0	$0
29	Contingency on Direct Costs	$797	$873	$1,177	$596	$839	$1,083
30	**Capitalized Indirect Services Cost**	$1,128	$1,235	$1,665	$1,431	$2,016	$2,601
31	Factory & Field Indirect Costs	$387	$424	$572	$435	$613	$791
32	Factory & Construction Supervision	$276	$303	$408	$166	$234	$302
33	Startup Costs	$23	$25	$33	$43	$60	$78
34	Shipping and Transportation Costs	$0	$0	$0	$41	$58	$74
35	Engineering Services	$393	$430	$580	$221	$311	$401
36	PM/CM Services	$49	$53	$72	$162	$229	$295
39	Contingency on Indirect Services Cost	$0	$0	$0	$362	$511	$659
40	Capitalized Pre-COD Personnel Costs	Not included					
50	Capitalized Supplementary Costs	$312	$341	$460	$236	$333	$429
51	Taxes	$0	$0	$0	$0	$0	$0
52	Insurance	$82	$90	$121	$0	$0	$0
53	Spent Fuel Storage	$0	$0	$0	$0	$0	$0
54	Decommissioning	$9	$10	$13	$12	$17	$22
55	Other Owners' Costs	$0	$0	$0	$0	$0	$0

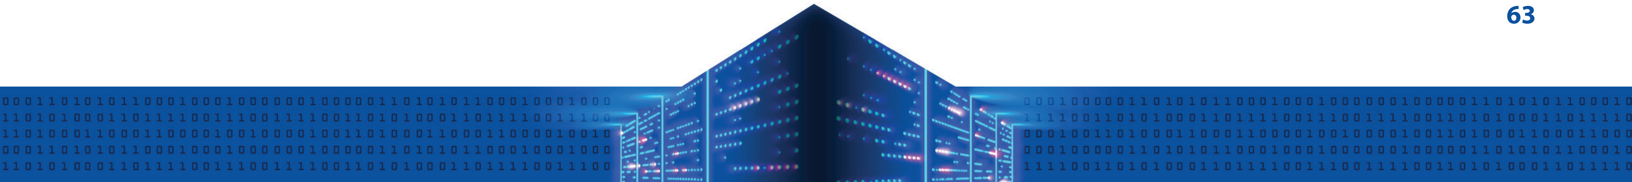

Account	Title	Large Reactor (~1,000 MWe)			Small Modular Reactor (~300 MWe)		
		Advanced ($/kWe)	Moderate ($/kWe)	Conservative ($/kWe)	Advanced ($/kWe)	Moderate ($/kWe)	Conservative ($/kWe)
56	Fees	$0	$0	$0	$0	$0	$0
57	Management Reserve	$0	$0	$0	$0	$0	$0
59	Supplementary Contingencies	$221	$242	$326	$182	$257	$331
	Total	**$5,250**	**$5,750**	**$7,750**	**$5,500**	**$7,750**	**$10,000**

O&M costs for nuclear plants are commonly categorized as fixed non-fuel O&M costs, variable non-fuel O&M costs, and fuel costs. Fixed O&M costs are largely driven by ongoing labor costs: operator, engineering, and security costs. Variable O&M costs include costs such as consumables and chemicals. Table 15 below provides assumptions for O&M costs based on [76].

Table 15. Estimated operations and maintenance cost and fuel costs for large reactors and SMRs.[76]

	Large Reactor (~1,000 MWe)			Small Modular Reactor (~300 MWe)		
	Advanced	Moderate	Conservative	Advanced	Moderate	Conservative
Fixed Non-Fuel O&M ($/kW-yr)	126	175	204	118	136	216
Fixed O&M ($/MWh @ 93% CF)	15.5	21.5	25.1	14.5	16.6	26.5
Variable Non-Fuel O&M ($/MWh)	1.9	2.8	3.4	2.2	2.6	2.8
Nuclear Fuel Cost ($/MWh)	9.1	10.3	11.3	10.0	11.0	12.1
Nuclear Fuel Cost ($/MBTU)	0.88	0.99	1.09	0.97	1.06	1.17
Total O&M ($/MWh)	**26**	**35**	**40**	**27**	**30**	**41**

Learning from First-of-a-Kind to *N*th-of-a-Kind

One of the most-important cost modifiers to understand is where a given reactor offering sits on the learning curve—specifically, whether it is a FOAK, a between-of-a-kind, or an nth-of-a-kind unit. These descriptions represent the maturity of a specific reactor cost offering based on the number of reactors (n) deployed. These titles are less of a specific design choice or parameter, but rather they represent a culmination of factors that are learned over time through deployment experience that contribute to lower costs.

FOAK costs are expected to be substantially higher than nth-of-a-kind costs, with the delta driven by non-recurring engineering costs, supply-chain inefficiencies, construction-learning penalties, regulatory uncertainty, and the absence of an established workforce trained on the specific design. The meta-analysis performed by Abou-Jaoude et al.[76] quantifies LRs from literature and applies them to project how costs may decline over successive deployments. Historical nuclear-LRs have varied dramatically. Experience from Korea and France, where standardized designs were deployed in rapid succession, showed meaningful cost declines while the fragmented U.S. program of the 1970s and 1980s saw cost escalation due to design changes mid-construction, regulatory interventions, and lack of standardization. The work by Abou-Jaoude et al. reports LRs ranging from about 5–15%[77] while follow on work in Abou-Jaoude et al.[76] estimates 8% for large reactors and 9.5% for SMRs. The cost of the nth reactor can be expressed as:

$$\text{Cost}_N = \text{FOAK} \times (1-\text{LR})^{\log_2 N}$$

where the cost of reactor n is a function of the FOAK cost, the LR, and the reactor number (N).

For a data center, the practical implication remains that the earlier in the deployment sequence a reactor is built, the higher the uncertainty (and potentially higher cost) one should expect. To combat this, organizing

large consortium-based purchasing agreements can help move the developer down the cost curve and lower the order-book average cost for all purchasers. Data-center developers considering nuclear should understand where in the production sequence the quoted units fall and what assumptions drive their cost projections regarding LRs and standardization.

Economies of Scale

Traditional nuclear power benefited from economies of scale: larger reactors produced more power with proportionally lower capital cost per kilowatt. The move to SMRs and microreactors deliberately trades some scale advantage for other benefits: modularity, faster construction, better financing profiles, and operational flexibility.

Analysis by Abou-Jaoude et al. distinguishes broadly between large (~1,000 MWe) reactors and small reactors (~300 MWe) reactors, with OCC ranges consistently higher on a dollar-per-kilowatt, electric basis for smaller units.[76] The key question for data center planners is whether the premium for smaller, modular units is offset by advantages in deployment speed, financing cost, reduced construction risk, and operational flexibility.

Economies of scale can be partially recovered through multi-unit deployments on a single site. Rather than building one large reactor, a campus could deploy multiple SMR modules, capturing volume discounts on shared infrastructure (site preparation, security perimeter, control systems, switchyard, cooling systems) while retaining modular-construction benefits. Estimates from Abou-Jaoude et al. estimate reductions in OCC and O&M for multi-unit plants up to 30 and 33%, respectively (Table 16).[76] Note that when applying the multi-unit reduction factor, the standard deployment package should be treated as a single unit. For example, if the standard plant design is a four-pack of reactors, then the four-pack counts as a single plant and the multi-unit reduction factor should only be applied if multiple four-packs are deployed at a single site.

Table 16. OCC and O&M cost reductions for multi-unit nuclear plants.

Number of Units	OCC Cost Reduction	O&M Cost Reduction
1	1.0	1.0
2	0.9	0.67
4	0.8	0.67
8	0.7	0.67
10	0.7	0.67

(Illustrative multipliers; actual values will vary by technology, site, and contract.)

Plant Performance and Capacity Factors

While not directly contributing to the OCC of a reactor, the capacity factor can directly impact life-cycle cost of a reactor. Given the large capacity factor of electricity from nuclear energy, the high-capacity factor of the current fleet of LWRs are the one reason why nuclear is still considered a necessity for energy security. The analysis of Abou-Jaoude et al.[76] did not distinguish between capacity factors between large and small reactors and used a value of 93% for all options. In reality, the 93% value for the capacity factor is directly applicable to LWR technology. When considering other technologies like SFRs, HTGRs, or MSRs, there will be a learning curve to reach the same level of performance and decades of performance of the current LWR fleet.

For an understanding of some of the capacity factors for certain reactors, a review of research reactors that operated in the past can be done. The U.S. previously operated research facility SFRs, notably the Experimental Breeder Reactor II (EBR-II), which operated from 1964 to 1994 and produced power as well as running experiments. Initially, the capacity factor was quite low, typically under 50% for the first decade of operation, even reaching as low as 20% in one year.[78] After this initial period, the capacity factor increased and exceeded 70% although it could have been 5–10% higher if its main purpose was power generation, rather than experimental operations.[78] A similar trend was seen with the operation of the BN-600 SFR in Russia, which has been operating for decades. Initially, the capacity factor was around 56%, but this steadily increased over time and exceeds 75% now.[79] These historical examples show that there are FOAK risks to a new technology. While there is historical experience, there will still most likely be startup issues with these new reactors.

Construction Risk, Overruns, and Standardization

Historically, nuclear-construction cost overruns have been the primary driver of project failures and uncompetitive economics. The U.S. projects at Vogtle Units 3 and 4 (the only recently completed large nuclear construction in the U.S.) experienced material cost and schedule overruns, with total costs significantly above original budgets and commercial operation dates years behind schedule. Construction risk is a function of many factors, including design maturity, regulatory stability, labor availability, supply chain depth, site conditions, and project management capability. These factors, specific to a nuclear project, should be strongly considered before decisions are made.

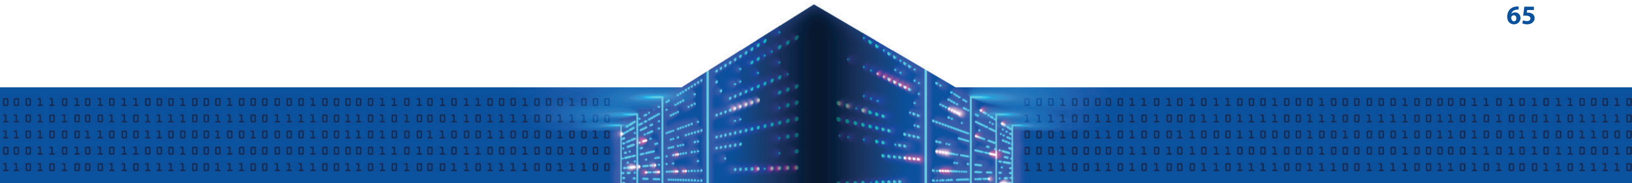

While most observed construction-risk data come from large LWR projects, one of the primary value propositions of advanced SMR designs is the shift from field construction to factory fabrication, where manufacturing occurs in controlled factory environments and is then shipped and assembled on site. This approach promises tighter quality control, shorter onsite construction schedules, reduced skilled labor requirements in the field, and a more predictable cost profile.

For data-center developers, a key question is how construction and delivery risk is allocated between the developer and the customer. Different contract structures distribute this risk in different ways, and those choices directly influence pricing and project outcomes. Some structures place more of the schedule and cost exposure on the project developer, while others shift more variability to the buyer. Many projects also rely on consortium or partnership models, where multiple parties share responsibilities and coordinate risk allocation across engineering, construction, and technology providers. Irrespective of the structure, understanding how risk is shared is essential for buyers, and financial models should incorporate stress-testing of budgets and schedules.

Financing, Incentives, and Cost Savings

Due to the high upfront capital requirements and multiyear construction timelines, interest costs and equity returns can add substantially to the all-in project cost well beyond OCC alone. However, data centers looking to deploy nuclear for BTM power present a different financing dynamic than standard utility-scale electricity providers that sell into a wholesale market. In this situation, financing can be underwritten against a large corporate offtaker (such as a hyperscaler), potentially achieving lower financing costs due to large investment-grade balance sheets and credit ratings.

Incentives like ITCs or PTCs can also play a part in cost reductions. Over the last several years, many state and federal bills have incentivized NPPs and data centers. Depending on the locale, these incentives could significantly reduce capital, operational, or tax burdens.

Key Insights

When building a reactor cost model for data center planning purposes, the following factors should be treated not as fixed inputs but as scenario variables that can be modified through strategic decisions.

1. **First of a kind versus nth of a kind:** the deployment maturity of the technology may be the largest single driver of cost uncertainty and should be bounded with high and low scenarios.
2. **Multi-unit sites:** the number of units deployed on site or across a buyer's portfolio affects both shared infrastructure economics and volume discounts.
3. **Standardization and modularization:** the degree of standardization and factory modularization affects how much of the cost-reduction potential the developer can actually achieve.
4. **Construction risk:** the construction risk allocation in the contract structure determines who bears the financial exposure to schedule and cost overruns.
5. **Financing:** the financing structure, timeline, and incentives can drastically shift all-in economics relative to OCC alone.

Taken together, these factors mean that two projects deploying nominally identical reactor technology could have significantly different all-in power costs that differ by a factor of two or more. For data-center operators, getting the deployment strategy, contracting structure, and financing approach right is critical.

Further Exploration

- National Laboratory of the Rockies. 2024. “2024 Annual Technology Baseline.” Golden, Colorado. https://atb.nrel.gov/.
- Department of Energy. March 2023. “Pathways to Commercial Liftoff: Advanced Nuclear.” https://yardsale.energy/wp-content/uploads/2025/05/DOE-Liftoff-Advanced-Nuclear-2024.pdf.
- Shirvan, K. 2024. “Total cost projection of next AP1000.” Massachusetts Institute of Technology. https://web.mit.edu/kshirvan/www/research/ANP201%20TR%20CANES.pdf.
- Robb Stewart, W. and Koroush Shirvan. 2023. “Construction Schedule and Cost Risk for Large and Small Light Water Reactors.” Nuclear Engineering and Design 407: 112305. https://doi.org/10.1016/j.nucengdes.2023.112305.
- Spangler, Ryan et al. 2025. “Potential Cost Reduction in New Nuclear Deployments Based on Recent AP1000 Experience.” https://doi.org/10.2172/2571083.
- Bolisetti, Chandrakanth et al. 2024. “Quantifying Capital Cost Reduction Pathways for Advanced Nuclear Reactors.” https://doi.org/10.2172/2361138.

4.4 Nuclear Integration: Dynamics of Policy and Public Perception

Introduction

Nuclear projects must navigate myriad federal, state, and local policies that shape the design and outcomes of project proposals. Understanding the policy landscape at each of these jurisdictional levels can greatly inform planning, application, and construction timelines as the U.S. nuclear resurgence gains momentum. The following sections give a brief, summary-level overview of the regulations most pertinent to nuclease projects. While not exhaustive, these policy considerations provide potential starting points for data-center stakeholders to prioritize early actions and inform decision-making when considering location, siting, configuration, and design choice of nuclear-generation assets.

Nuclear projects have historically provided safe, reliable, and relatively inexpensive energy within the U.S. However, a handful of rare nuclear accidents around the world in decades past have helped perpetuate negative stigma surrounding nuclear technology. How the public perceives nuclear technology will be vital to integrating nuclear projects into data centers. This is based not only on the trends observed with the development of data centers themselves, but also perspectives from some local officials who believe advanced nuclear project will likely “face significant community resistance.”[80]

Why Does This Matter for Data Centers?

These concepts affect the planning, construction and operational phases. Public policy and perception represent a double-edged sword for data centers: clarity surrounding regulatory nuances and robust public support can enable such projects to achieve timely completion and streamline operations and future expansion down the road. However, without a keen understanding of these concepts, nuclear integration can face political hurdles and developmental setbacks due to an uncertain regulatory environment at each jurisdictional level.

Research and Findings

Federal Policy

The nuclear space is shaped by several federal policies, some of which span decades. Recent national priorities have culminated in a suite of executive orders (EOs) at the federal level that are geared toward revitalizing the U.S. nuclear-energy sector while pushing for innovation through adoption of new technologies such as SMRs and microreactors. This includes EO 14299, "Deploying Advanced Nuclear Reactor Technologies for National Security," EO 14300, "Ordering the Reform of the Nuclear Regulatory Commission," EO 14301, "Reforming Nuclear Reactor Testing at the Department of Energy," and EO 14302, "Reinvigorating the Nuclear Industrial Base." Together, these EOs outline priorities maximizing the capacity of existing nuclear generation facilities while fostering the addition of new generation sources through technological innovation and regulatory reform and, ultimately, quadrupling U.S. nuclear energy capacity from 100 to 400 GW by 2050. Understanding these directives and their purpose provides a valuable context in understanding the necessity of growing U.S. energy supply. Aligning these efforts to the needs of data center growth and other large loads may help nuclear generation projects achieve desired outcomes in feasible timeframes.

However, the brunt of federal regulatory agency resides with the NRC. This body is responsible for accepting, reviewing, approving and renewing applications for nuclear projects. The Atomic Energy Act of 1954 and Title II of the Energy Reorganization Act of 1974 codified the functions and authorities of the NRC. Additionally, the NRC is responsible to produce environmental impact statements for licensing all new and renewed nuclear plants, pursuant to the National Environmental Policy Act of 1969.[81] Specifically, there are certain sections of U.S. Code that pertain to the development and renewal of nuclear projects and licenses. Title 10 of the United States Code of Federal Regulations (CFR), Chapter 1, enumerates the roles and responsibilities of the NRC, making this entity a central fixture in the regulatory-review process through which any new nuclear project must go. The general review process of the NRC for new and renewed nuclear plant licensing is outlined in Figure 18.

Figure 18. New-reactor licensing process (source: NRC).

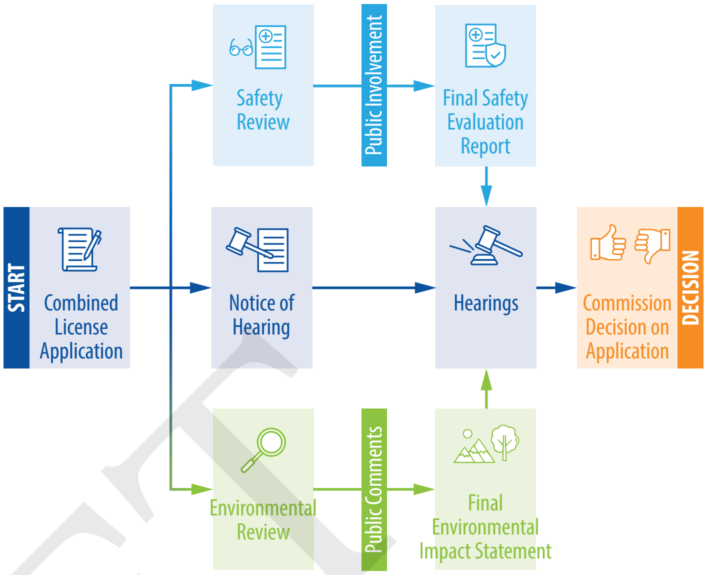

Specifically, 10 CFR 1 articulates the authorities of the NRC pertaining to the licensing-application process as well as requirements for various license applications. Of note, Parts 50–55 specify the requirements of production and utilization facilities (reactors). Understanding the provisions of these sections can help stakeholders identify key requirements, thereby streamlining the path to regulatory approval. Other related sections elsewhere in 10 CFR 1 may also provide guidance on regulations for other nuclear infrastructure, such as fuel production (Parts 40–49), waste disposal (Parts 60–69), and transport facilities (Parts70–79). For the immediate focus of nuclear reactors, Table 17 identifies Parts 50–55 and their corresponding requirements:

Table 17. Code provisions: 10 CFR 1, Parts 50–55.

10 CFR Chapter 1 Provision	Regulatory Topic
Part 50	Power reactor licensing
Part 51	Environmental protection
Part 52	New reactor licenses
Part 53	Risk informed, technology-inclusive regulatory framework for advanced reactors (draft)
Part 54	License renewal
Part 55	Operator licenses

Of note, some licensing options under Part 52 allow for application of early site permits, which may obtain approval for a reactor site without specifying the design of the reactor to be built, as well as certified plant designs that can be approved expeditiously approved.[82] This standardization may accelerate the regulatory process because these designs need only go through full review once, providing predictable metrics and outcomes with each subsequent application. Some designs have already gained this expedited approval from the NRC, and this trend is likely to continue as advanced-reactor designs evolve over the next few years.

While most of these sections are in effect, Part 53 represents a new section that seeks to expedite the review process for advanced SMRs and microreactors. The NRC recognizes the uniqueness of advanced reactors and is in the process of developing a final rule that would allow future commercial nuclear plants, including both LWRs and non-LWRs, to pursue licensing under a separate review process. As the title implies, this process would leverage a probabilistic risk assessment that applies risk-informed and performance-based review methods that are flexible and practicable for application across the multitude of different advanced-reactor designs.[82] This draft-rulemaking package for Part 53 was proposed in 2023, with the final version expected to be implemented by 2027.

General awareness of the various federal policies and procedures can influence planning efforts early in the data-center development process. This can lead to better-informed decision-making when selecting the type of nuclear technology most appropriate for a data-center project and potentially expedited regulatory approval processes as well.

State and Local Policy

While the federal regulatory environment is well-defined and supported by national priorities, a great amount of variation exists at the state and local levels. In fact, the variation becomes more evident at lower levels of jurisdiction. Support for nuclear technology at the state level has been relatively positive, with several states passing legislation that favors nuclear technology as a reliable-energy source in recent years. A 2023 study led by INL identified several nuclear-related policies and developments across several states.[83] Select findings from this study are outlined in Table 18.

Table 18. Nuclear-related policy and developments. Source: INL.

Policy and Developments	States
Nuclear adoption and extension support	AK, CT, GA, ID, IL, IN, KY, LA, MD, MI, MS, MT, NE, NC, UT, VA, WV, WY
Expanding policy definitions with nuclear as a preferable energy source	CA, CT, ID, IL, IN, NJ, VA
Repealed prohibitions on nuclear development or updated siting authority	AK, CT, KY, MT, WV
Proposing or advancing studies aimed at potential for siting or permitting advanced reactors	MD, MI, NE, NH, PA, SC, VA, WI

The advancement in state-level policy related to nuclear technology is likely to continue, with more states expected to develop their own regulatory frameworks and siting studies to accommodate new nuclear projects. These frameworks will likely have an impact, not only on nuclear siting processes, but on data-center siting as well, depending on the configuration and nuclear-integration arrangement. This is especially important for data centers seeking BTM arrangements—including co-location designs that must navigate both nuclear and data center regulatory frameworks. Tracking policy developments and economic drivers between states may lead to environments conducive to the development of data centers with nuclear integration. Still, regulatory challenges may reside within each state at the local level.

Despite being smaller jurisdictions than the state and federal levels of government, local government represents a key threshold that all projects, whether nuclear- or data center-related, must cross before construction begins. In many ways, local regulation is the gatekeeper for almost all of these projects. Hence, understanding the frameworks, processes, and requirements for development at this level should be a priority immediately after submitting a combined (construction and operation) license application with the NRC.

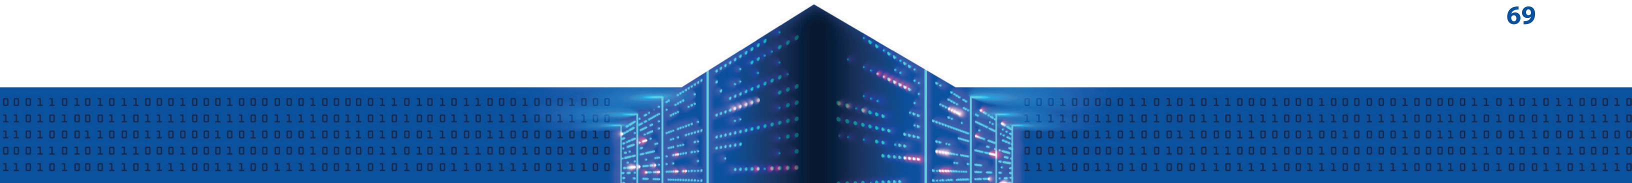

Regional economic, environmental, and social factors are far more nuanced at the local level than at the federal or even state levels. Within a single state, factors such as water availability and water rights, population density, and economic incentives vary greatly. All of these factors influence public policy. Localities with existing nuclear plants may be more conducive to additional nuclear projects. Others that may have significant existing data-center infrastructure may not have the regulatory framework in place for nuclear. This has the potential to lengthen the approval, and even construction processes, once approved. Local zoning and land-use codes have specific site-planning and performance requirements that vary between jurisdictions.

However, this complexity can be used to a data-center project's advantage. Variation in regulatory policy may translate to flexibility on the development side, allowing projects to choose from a variety of options and customizations to calibrate nuclear-integration projects to niche regulations within a local context. The wide variety of nuclear-generation options—such as the many advanced reactor designs—can be used as levers to optimize cost, data-center power demand, and local regulatory provisions. In this regard, a one-size-fits-all approach may actually hinder project advancement while adjusting and fine-tuning nuclear projects to meet the needs of local jurisdictions may ensure approval and expedite construction.

Another advantage for nuclear projects, specifically advanced reactors, may exist—somewhat ironically—in their timeline to commercial viability: data centers and nuclear projects are both heavily regulated industrial projects with similar scales in terms of spatial requirements, water usage, environmental impacts, and safety concerns. Significant overlap exists between data centers and nuclear in these areas. While widespread commercial availability of advanced reactors is expected to occur 5–10 years from now, many of the existing and planned data centers today will have already undergone these regulatory processes. This means advanced-reactor projects will have the advantage of identifying regions where data-center developments were successful and understanding the specific challenges involved in accommodating those developments. This, in turn, removes much uncertainty as to where nuclear projects would have a favorable chance of success. Leveraging lessons learned by current and planned data-center projects at the local level can provide valuable insights and best practices for integrating nuclear projects as well.

Several concepts—setbacks, transportation access requirements (especially for nuclear fuel delivery and storage) and construction access and aesthetic design—warrant consideration early in the process. Each jurisdiction has unique development requirements, and this environment is evolving due to the rise in data-center development. Incorporating these considerations into the data-center planning and development processes can expedite co-located nuclear-development approval and preclude heavy revisions to data-center campus-site design once advanced reactors become available. However, public policy at the local level is largely driven by public perception, a component in the regulatory context that should not be overlooked.

Public Perception

Understanding public perception is especially important at the local-government level for a few reasons. First and foremost, public perception influences election of public officials, whose offices drive policy, development standards, and regional priorities that can affect the regulatory environment. Second, the public has much greater access to elected officials, who are more intrinsically involved in specific site-planning processes than their state or federal counterparts. Third, federal, state, and local nuclear-regulatory processes require many public hearings during which individual residents and interest groups can speak either against or in favor of projects. Getting ahead of any potential concerns can reduce uncertainties that might extend the review process and thereby minimize public resistance. Finally, gauging public perception early in the planning process establishes a baseline of the peoples' understanding and clarifies areas where outreach and educational campaigns may quell many of the concerns harbored by the public.

Many of the same questions posed by the public surrounding data centers regard environmental impacts, resource depletion, and public-safety concerns. These map onto nuclear projects as well. One important distinction between data centers and nuclear plants is job creation; while data center employment surges during construction, the actual workforce needed long-term to operate data centers is relatively small. By contrast, nuclear projects must retain a sizeable workforce of operators, engineers, physicists, and other support staff to operate and maintain facilities long after construction is complete. Other concerns regarding water usage and public safety are among the top priorities for the public. Addressing these concerns may help clear objections and alleviate pressure placed on public officials by their constituents, thus aiding nuclear projects to clear regulatory review.

Here too, timeline can be used to the advantage. Whereas data-center developments must alleviate concerns in real time as new and proposed projects undergo regulatory and site-planning review, nuclear projects—particularly advanced-reactor projects—may be years away, allowing a valuable window through which to observe trends in public perception, impacts on data-center development, and the effectiveness of outreach and educational campaigns to inform the public on the safety and benefits of nuclear integration.

Data-center stakeholders can take immediate action to support nuclear-integration projects by actively engaging with public officials and local interest groups to open dialogue on priorities and concerns regarding future projects. This helps achieve amicable solutions that benefit all parties involved while demonstrating transparency and proactive stakeholder engagement. Educational and public-outreach initiatives can greatly inform the public while demystifying some of the stigma surrounding nuclear energy. Underscoring the technological advancement of nuclear energy and highlighting the safety, economic benefit of potential BTM configuration, and minimal environmental impacts of nuclear energy should serve as a solid starting point around which key messaging strategies can be built.

Key Insights

1. The efforts to expedite regulatory processes, standardize reactor designs, and build flexibility into new frameworks (like 10 CFR Part 53) will become increasingly important in public policy at the federal level.
2. The decentralized and varied nature of state and local policies introduces challenges that data-center developers must carefully navigate. This complexity presents both obstacles and opportunities for tailored solutions. Understanding the nuances of development processes and distinct requirements in each locality can greatly influence the implementation of nuclear-integration projects for data centers.
3. The timeline for advanced-reactor adoption provides an opportunity to analyze successful data-center projects, identify favorable regulatory environments, and leverage lessons learned to address public concerns and policy hurdles for future nuclear integration.
4. Public perception of nuclear energy represents a key issue that warrants early consideration and action, especially at the local level, requiring data-center developers to proactively engage communities early through educational initiatives that address safety, environmental impacts, and economic benefits to build trust and reduce opposition.

4.5 Nuclear Supply-Chain Considerations

Introduction

Supply chains are a key component of nuclear integration that should be considered early in the planning process for data-center development. Adequate planning and accommodation have the potential not only to shape the power supply of data centers, but their design. For this analysis, an important factor to note is that a marked difference exists between traditional nuclear power-generation assets, otherwise known as large NPPs, and advanced nuclear-generation assets, such as SMRs. Though the fission process is largely the same among all nuclear reactors, great variation is found among these nuclear-generation assets, which greatly influences their underlying supply chains, fuel life cycles, and the facility requirements that enable them. Some key supply-chain considerations include fuel life cycles, logistics requirements, and mutual supply chains that support both SMRs and data centers. Some supply chains that support nuclear-generation facilities meant to sustain data centers may also support the data centers themselves, potentially leading to inadvertent competition for resources and supplies to achieve the same goals and objectives. Such supply-chain considerations should be accounted for early in the planning process of data-center development. This will help identify potential pitfalls and inform strategic decision-making while building supply-chain resilience in the long term. Additionally, other factors, such as location, configuration, and workforce development, can also affect nuclear supply-chain resilience as well. The following sections offer key insights and considerations that account for variability in the numerous supply chains that will support nuclear integration with data centers and other large loads.

Why Does This Matter for Data Centers?

Supply chains for nuclear technology can be complex and diverse, depending on a number of factors including design, location, configuration, and fuel type. Understanding this diversity early on enables detailed analysis of nuclear integration in tandem with data-center development. Mapping the supply chains for fuels and other resource requirements can greatly inform selection of nuclear technology to be applied to these data centers. Incorporating these factors in the planning phase of data-center development can ensure that spatial and connectivity considerations, as well as configuration for nuclear technology are accounted for in the overall data-center design, potentially saving expense and regulatory difficulty of piecemealing the development of nuclear projects and data centers separately.

Research and Findings

Fuel Cycles of Large NPPs vs. SMRs

Nuclear fuel cycles should be one of the primary factors to consider when integrating nuclear generation with data centers. Tradeoffs certainly exist between large NPPs and SMRs, most of which are affected by the fuel used in electricity production. Fuel cycles between large NPPs and SMRs are expected to be relatively similar, with reactors operating between 1.5 and 2 years between refueling.84 Modularity of SMRs allows for timely refueling. When units are brought offline, they can be removed and immediately replaced during the refueling stage, allowing additional SMR units to continue operating through the refueling cycle. However, nuances exist between the type, quantity, and even enrichment levels of fuels. These variables warrant further investigation into which type of nuclear technology best suits the data-center project in question.

Large NPPs often have well-established fuel supply chains that, in many cases, have been in operation for decades. This fact—coupled with the largely ubiquitous fuel requirements and standardized transport, storage, and disposal infrastructure—leads to a robust, predictable fuel supply chain that accommodates the entirety of the fuel life cycle. By comparison, SMR fuel supply chains and life cycles are not yet well established, leading to challenges affecting CAPEX and operational expenditure estimations. This is because tremendous variation exists between SMR designs; no two designs are the same. This translates to niche supply chains for each, meaning that containment vessels, transport requirements, and storage and disposal protocols may not be as standardized as those of large NPPs.[85] Still, these advantages are relegated to regions where large NPPs already exist and have the capacity to support data centers.

Just as SMRs that are the FOAK are typically more expensive to develop, operate, and maintain, their fuel supply chains will likely also experience higher costs due to the high LR associated with FOAK development. As the fuel supply chains materialize and begin to solidify, associated costs are expected to reduce due to the growing economies of scale.

When SMRs are considered, understanding the duration and optimal phasing, or staggering, of unit refueling cycles can help inform which design and configuration best fits the needs and uptime requirements of data-center projects. Research conducted at INL identifies an optimal number of eight 300-MW SMRs needed to power a hyperscale datacenter with a load requirement of 1 GWe. The remaining two units may be utilized to allow for scheduled maintenance on a rotating basis without affecting overall capacity. However, the exact number reactors needed will vary among datacenters, based on factors such as load and availability requirements, type of SMR deployed, and interconnection agreements with utilities and ISOs.

Understanding the widespread application and availability of each fuel type, manufacturer, and source of upstream materials that factor into fuel production can provide visibility into how to build a resilient fuel supply chain. As SMR technology advances and becomes commercially viable, coordination between nuclear and data center stakeholders will become increasingly important in working toward greater standardization of fuel types, life cycles, and equipment, thereby collectively making the fuel supply chain more robust across the industry.

Shared and Indirect Supply Chains

Research reveals several commonalities in supply chains between data centers and nuclear generation plants. Both require intensive electricity transmission and distribution (T&D) equipment, often relying on the same electrical equipment supply chains to support them. Research has shown that transmission bottlenecks can delay nuclear and data-center deployment and transmission congestion directly drives supply-chain urgency.

This is especially true in cases where nuclear facilities and data centers are geographically concentrated in the same locality. Co-location of nuclear generation with data centers can reduce interconnection costs and allow shared use of existing infrastructure, such as substations, transmission lines, and water-system facilities. While this commonality may appear synergistic, it creates competition for common resources that, in many cases, are becoming increasingly scarce. This in turn can cause delays in data-center development. In fact, at least 20% of projected data-center additions expected through 2030 could face delays due to supply chain shortages in such equipment as turbines and transformers.[86]

The energy-sector industrial base is especially sensitive to heighted demand. Water represents another commodity for which data centers and nuclear generation plants may compete. Though water is often viewed as a siting, regulatory, or environmental issue—or any combination thereof—it also warrants consideration as a supply-chain concern. Both data centers and nuclear technology rely on light water that requires treatment and purification prior to use. Some SMR designs utilize light water for neutron mediation as well as cooling. These LWRs and some data centers employ light water as a coolant due to its thermal properties that can cool systems up to 30 times faster than air-cooling technology. Each type of SMR may have different water requirements, with SMR internal configurations varying just as much as those of data centers. Understanding the specific water requirements of the specific SMR model, and location of the data center in question, are both key considerations that affect water supply-chain dynamics.

Because of the requirement to treat water prior to use by either nuclear reactors or data centers, treatment facilities with adequate capacity must be in place to accommodate such heavy uses. In smaller localities and remote regions, this often entails the construction of new water treatment facilities to accommodate the significant increases in volume, along with the disinfectant and treatment supply chains, such as chlorine, to enable a constant, robust light-water supply. Heavy usage by both data centers and power plants within the same geographical region can not only exacerbate local and regional water shortages, but create significant supply-chain dependencies where fluctuations in water availability may strain data-center operations, depending on the center's design factors, scale, and employed cooling technology. This may occur either directly through data-center water usage or indirectly through electricity generation. Notably, the latter concern is not relegated to nuclear generation alone; it also impacts other generation sources that rely on steam turbine technology. Nevertheless, understanding the sourcing of water and potential supply-chain risks—such as disruptions to water-treatment plant operations or sudden drops in water supply—are important factors to consider when selecting which nuclear technology will support data-center development.

Even secondary and tertiary supply chains that do not directly feed into nuclear power generation have the potential to impact the overall supply chain security. One such example can be found in the context of international supply chains for data centers themselves. As of 2025, the U.S. has no domestic commercial large-scale supplier of HALEU, and instead comes as a byproduct of U.S. DOE's downblending and processing of high-enriched uranium stockpile.[37] Meanwhile, processors, rare-earth elements, fiber-optic cables, and other elements that enable data centers to operate may be strained by geopolitical dynamics, sanctions, or economic factors that may cause supplies to fluctuate. While such considerations clearly have a direct impact on data centers, demand for nuclear power is intrinsically tied to the demand sought by data centers and other large loads. If the projected growth of data-center development fluctuates or is truncated by the availability of materials and resources, so too is the demand for new technology such as SMRs, affecting economies of scale for building and deploying SMRs to meet this demand.

Other supply chains, such as those across the transportation sector, also impact the economic viability of nuclear integration. The price of fuel, bandwidth of manufacturing of specialty shipping and hauling vessels that accommodate emerging nuclear technology, and even construction equipment can have indirect impacts on the supply chain for nuclear assets—particularly SMRs. Conversely, these existing supply chains for data centers are relatively robust compared to those of SMRs; continued elevated demand for data centers may translate to nuclear supply chains' not being able to keep pace to satiate demand over the next 5–10 years. Though seemingly disconnected, these supply chains can have reciprocal influence on one another, warranting the need to monitor supply chains across different industries in order to adapt nuclear integration strategies accordingly. Table 19 demonstrates potential areas of mutual reliance on common supply chains between nuclear generation plants, data centers, and electricity T&D with the strongest commonalities bolded for additional consideration.

Other Supply Chain Considerations

Table 19. Mutual supply-chain dependencies. Nuclear generation plants, data centers, and electricity T&D.

Component / System	Nuclear	Data Centers	Electricity T&D
Prime mover	Steam turbine	None	Conduit
Electric generator	✓	No	No
Heat source	Nuclear fission	IT equipment	No
Steam generation	✓ (steam generators)	No	No
Condenser	✓	No	No
Cooling system	✓	✓	Sometimes
Cooling towers/heat rejection	✓	✓	No
Primary working fluid	Water/steam	Air/liquid	Oil
Feedwater system	✓	No	No
Pumps	✓	✓	No
Large rotating machinery	✓	Limited	No
Transformers	✓	Step-down/UPS	✓
Switchyard/substation	✓	✓	✓
Electrical distribution	✓	✓	✓
Emergency power	✓	✓	Sometimes
Backup generators	✓	✓	Sometimes
Energy storage	Limited	Batteries (UPS)	Sometimes

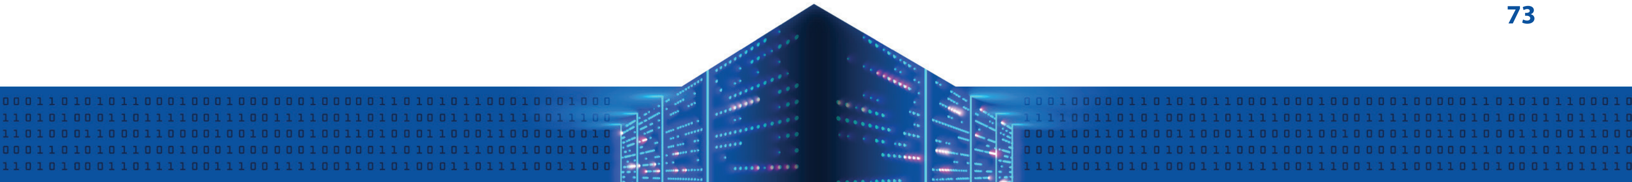

Component / System	Nuclear	Data Centers	Electricity T&D
Control room / NOC	✓	✓	✓
Instrumentation & control	✓	✓	✓
Protection systems	✓	✓	✓
Fire protection	✓	✓	Sometimes
Water treatment	✓	Sometimes	No
Fuel handling	Nuclear fuel	None	None
Waste handling	Rad waste	E-waste	Specialty Chemicals
Cranes/heavy lifts	✓	✓	✓
Civil structures	✓	✓	✓

Site planning is often influenced by supply-chain considerations. The size and demand of data centers can influence the quantity, type, and configuration of such onsite nuclear generation as SMRs, which inherently determine the type of generation, capacity, fueling cycles, water-supply requirements, and spatial design of co-located projects. Transportation connectivity is another key aspect that influences supply-chain viability for nuclear-powered data centers. Ensuring that surface-transportation corridors have the correct clearance, turning radii, and loading facilities to, from, and within the data-center development will be largely driven by SMR design. Site selection that incorporates robust and redundant transportation, such as maritime and intermodal facilities, can influence supply-chain security as well.

Workforce is yet another factor related to supply chains that will warrant additional consideration. Data centers, in combination with nuclear facilities, can support a significant number of local jobs. A GW-scale data center is expected to directly employ around 1,000 employees. A single GW-scale nuclear facility would employ at least 650 employees per reactor. Employment for SMRs is currently expected to require fewer workers per gigawatt of electricity. Nuclear generation facilities require a robust workforce that may not be large enough to accommodate the rapid increase in nuclear integration. In this regard, talent pipeline may be regarded as yet another supply chain that is often overlooked. For more information on workforce needs and development, please refer to the **Workforce and Socioeconomic Impact of Nuclear-Powered Data Centers** section of this Playbook.

Federal, local, and state regulatory approval processes should also be taken into consideration. Public hearings are typically required for any nuclear project, including new builds and expansions. Timing these processes—in alignment with lead times for construction materials, labor, and other capital—should be considered in supply-chain assessments. Under one scenario, labor contracts and early orders for fuels and construction materials could stall if the approval processes are prolonged. Likewise, expeditious approval prior to material procurement could impact project timelines and lead to cost overruns. To this end, close coordination with regulatory bodies at federal, state, and local levels should not be overlooked, but incorporated throughout the development and procurement processes. Even the construction of these projects can have adverse effects on localities, especially smaller communities that may lack robust infrastructure and housing stock to support construction activities. Here, too, supply chains are a key consideration because a rapid influx of construction personnel and families can constrict housing stocks or even leave a high number of vacant homes that these communities would have to address. While growth is typically welcomed in most communities, rapid fluctuations in housing availability and overbuilding represent a long-term risk for localities experiencing nuclear integration with data-center projects.

Key Insights

- While both traditional large NPPs and SMRs rely on similar processes, their supply-chain composition will be inherently different. Moreover, no two SMR designs are the same, meaning that designs that leverage common fuels, technology, and storage equipment may help mitigate supply-chain shocks.
- Mapping nuclear supply chains should include identifying those components which may overlap with supporting data centers, such as electrical equipment and water distribution.
- Secondary supply chains may still impact the speed of SMR adoption, such as rare-earth elements, processors, and transportation equipment availability. Understanding the effect of these supply chains on nuclear integration with data-center projects while seeking domestic production sources may mitigate some of the potential risk associated with international supply chains.
- Transportation and workforce requirements should be taken into consideration to reflect the size and scope of SMR integration with data centers and may influence ultimate project design.

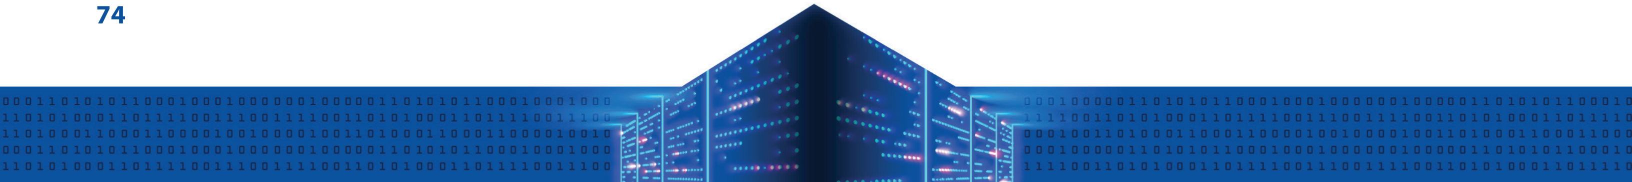

4.6 Nuclear Plant Reliability and Availability

Introduction

The nuclear industry has long relied on reliability or unreliability metrics (e.g., failure rates or frequencies) for the purposes of safety, design, regulation, and operation. These metrics help organizations identify weaknesses, prioritize maintenance, optimize designs, and reduce operational risk. Through its emphasis on reliability metrics, the nuclear industry has developed a reputation for safety and performance. Emphasis on quantifying and managing equipment has contributed to operational reliability and extended periods of uninterrupted operation. In fact, the average capacity factor of the commercial fleet regularly exceeds 90%.[87] Given this fact, nuclear power aligns well with data centers or sectors requiring uninterrupted, reliable power.

Why Does This Matter for Data Centers?

The demand from data centers is outpacing the supply of available power from the standard electric grid. This demand has created significant challenges for grid operators, including regional constraints on capacity and long lead times for establishing new connections to the electric grid. As a result, data-center developers are forced to consider alternate sources of electricity, such as onsite localized microgrids powered by nuclear generators. The design of such microgrid solutions is a balance between reliability and cost. Reliability can simply be described as the ability of a component to perform its required functions during a specified period.[88] Generally high reliability (e.g., the absence of failures) is expensive or unfeasible. Given this reality, failures are expected and managed through maintenance activities and system designs (e.g., redundant features). Thus, multiple, cost-effective components with decent reliability can meet the overall performance of a cost-prohibitive, highly reliable component. Understandably, these design decisions, as well as maintenance activities (i.e., planned downtimes), have an impact on the ability of a microgrid to support the steady power requirements of a data center. This next section discusses reliability metrics pertaining to a nuclear power source.

Research and Findings

Reliability is described as the ability (expressed as a probability) of a component's performing its required functions during a specified period.[88] Strictly speaking, reliability does not consider the concept of repair. By considering repair, the discussion of reliability becomes a discussion of availability. After an item fails, there is a period during which it is incapable of performing its required function. This downtime (i.e., unavailable period) encompasses the activities necessary for repairing and redeploying the item to its functional or ready state. Availability captures the downtime due to unscheduled maintenance (i.e., failures and repair) and scheduled maintenance. Thus, availability generally refers to the quality or state of being ready for use. This perspective of availability is often used interchangeably with "uptime" or the fractional time that a system is up and running during a given period. A similar concept, known as capacity factor, is related to uptime but is slightly different. Capacity factor represents the ratio of actual energy produced for a set period to the maximum energy that could have been produced during the same without any downtime.

In summary, a highly reliable power generator is one that operates for a long time without failure. A highly available power generator is one that is nearly always capable of providing power, but there is no limit or requirement for how much power it must provide. A power generator with high-capacity factor means that the generator nearly always operates at maximum power. In all cases, reliability metrics (used loosely here to describe metrics that influence reliability, availability, capacity factor, etc.), are important to monitor and track. Generally, metrics that are important are those that capture downtime, including failure rates and maintenance activities.

Unanticipated Failures

NPPs are complex machines with numerous components. To evaluate the reliability and availability of a nuclear power generator, it is necessary to consider the multiple ways by which a reactor may be offline or unavailable (for both scheduled and unscheduled maintenance). Presently, advanced and next-generation reactor data are limited; however, data on existing and historical large NPPs are less scarce. NRC data acquired from current and historical nuclear reactor operations can be used to determine an overall unscheduled trip rate.[89,90] Table 20 provides a list of events that would be considered the cause of unscheduled maintenance activities. Note that mean duration data was extracted from [90] and [91].

Table 20. NPP transient data (PWR).

Initiating Event	Mean Event Frequency (per year)	Mean Duration (hours)
Loss of vital AC bus (LOAC)	5.29E-03	192.00
Loss of vital DC bus (LODC)	6.23E-04	120.00
Partial loss of component cooling water (LCFW)	1.56E-03	348.00
Loss of feedwater (LFW)	6.98E-02	188.16
Partial loss of service water (LSW)	1.25E-03	408.00
PWR loss of instrument air (LOIA)	8.44E-03	59.04
PWR loss of heat sink (LOHS)	5.81E-02	340.80
PWR General Transient (GT)*	9.43E-01	137.52
Loss of coolant accident (LOCA)	1.56E-03	1200.00
Loss of offsite power (LOOP)	2.36E-02	18.11
Total events/year	1.1132	
Total Event down time per year (hours)	167.56	

Scheduled Outages

Traditional LWRs typically need to shut down and refuel every 18–24 months of operation for PWRs and BWRs although the possibility of extending this cycle is being explored. The U.S. Energy Information Administration (EIA) publishes information on the average length of refueling outages for operating plants in the U.S. based on information from the NRC.[92] The average length of refueling outage for plants in a given year ranged from approximately 30 to 45 days between 2000 and 2023.[92] However, refueling outage length can vary by plant. For example, Constellation has reported that the average refueling outage length for its fleet averaged 21.5 days in 2025, compared to the industry-wide average of approximately 37.5 days.[93]

The outage frequency and duration for advanced-reactor designs are expected to differ from those of the existing fleet. For example, some reactors are designed for online refueling, meaning that the reactor would not need to be shut down for refueling at all. For reactors that refuel online, the refueling cycle (i.e., refueling outage frequency) and outage duration can vary for different reactor designs. Table 21 lists the anticipated refueling cycle characteristics for some advanced reactor designs.

Table 21. Designed and upgraded fuel cycle length with the duration of the refueling and maintenance for different NPP types.

Reactor Type [Size]	Designed Fuel Cycle Length	Potential Upgrades for Fuel Cycle Length	Refueling and Maintenance Duration
PWR [1000 MW]/ AP-1000	18 months[94]	24 months	30-45 days[92]
BWR [1000 MW]	24 months	36 months	30-45 days[92]
Natrium [345 MW]	18 months[95]	N/A	N/A
Arc-100 [100 MW]	20 years[96]	N/A	N/A
NuScale [77 MW/ module]	21-24 months[97,98]	48 months[96]	10 days per module[97]
AP-300 [300 MW]	48 months[99]	N/A	N/A

Even reactors designed for online refueling will still require shutdown periodically for maintenance activities. It is expected that the maintenance schedule for the turbomachinery will be the driving factor for the frequency and duration of scheduled outages for reactors that refuel online. Turbomachinery maintenance schedules are typically dictated by the equipment supplier; therefore, they will vary. As an example of possible outage requirements, Table 22 provides the maintenance schedule provided by one of the major turbine manufacturers.

Table 22. Turbomachinery outage assumptions, based on expert solicitation.

Outage Type	Shutdown Time (days)	Maintenance Time (days)	Startup Time (days)	Outage Frequency (years)
Endoscope Inspection	1	1-3	1	2
Minor Maintenance	1	12	1	4
Major Maintenance	1	27	1	12

Key Insights

Ultimately, for the purposes of reliability and availability assessments, it is necessary to track events and data that influence the downtime of a given power source. Most of the random failure data for next-generation SMRs, especially candidates for microgrids, remain uncertain. Ideally this data should come from the nuclear vendors themselves.

Further Exploration

Further details for reliability and availability can be found in *Reliability Engineering and Risk Analysis: A Practical Guide.*[100]

4.7 Site Planning and Selection

Introduction

This section explains where and how NPPs and data centers can be built in tandem across the U.S., based on a comprehensive national siting analysis using the OR-SAGE tool. This analysis determines local suitability using criteria based on environmental, geotechnical, hydrological, and infrastructure factors.

Environmental and geotechnical constraints exclude regions where siting would be either impractical, prohibitively expensive, or detrimental to the safe operation of the infrastructure in question. These constraints include land within 100-year floodplains, wetlands and protected lands, areas with excessively steep terrain, and regions with a high risk of landslide or seismic hazard. Additionally, for nuclear facilities, the risk of radiation dose to nearby populations in the event of failure requires appropriate setback distances from dense population centers.

Site determinations are also based on access and proximity to various resources. For instance, cooling is essential for both reactors and data centers. Sites near rivers, lakes, or robust municipal water systems offer greater flexibility. In arid regions, dry cooling or alternative systems may be required. Other resources required for successful deployment of both nuclear and data center facilities include access to transmission infrastructure and roads, the presence of a suitable and sizeable workforce, and fiber-optic connectivity for data centers.

Why Does This Matter for Data Centers?

While nuclear reactors and data centers serve different purposes, they share similar requirements. Data centers require massive amounts of reliable, constant electricity, significant cooling resources, and large land areas. Advanced nuclear reactors, especially SMRs, can provide a reliable source of base-load electricity. Compared to large reactors, their compact footprint and relaxed criteria constraints offer a widespread region of suitability. When co-located, these systems can reduce strain on regional grids, improve energy security, enable AI growth, and revitalize existing energy communities. The configuration options based on this study is summarized in Table 23.

Table 23. Configuration options for co-location of SMR and data-center technology.

Configuration	Strengths	Best Use Case
SMR + hyperscale data center	Flexible siting, lower water demand	Midwest & Southeast
SMR + data center at brownfield site	Fastest deployment	Energy communities
Large LWR + GW data center	Maximum scale	Water-rich regions
Dry-cooled SMR + modular data center	Water-constrained areas	Select Western states

Research and Findings

Brownfield Revitalization: A Transformational Opportunity

One of the findings of the national study is the potential to repurpose existing infrastructure. Many operational nuclear sites have available land within a 5-mile radius that could support hyperscale data centers. Leveraging these locations can reduce development timelines, use existing transmission capacity, and build on established regulatory familiarity. Retiring power plants present exceptional opportunities for transformation because they come with preexisting connections to transmission infrastructure and water supply. These sites are already zoned for industrial use, and the surrounding community has become accustomed to energy operations. This also means that a workforce to operate these sites has already been established.

Repowering these sites with SMRs and co-locating data centers can preserve skilled jobs, stabilize local tax bases, and avoid greenfield land disturbance, all while supporting energy abundance. This pathway offers one of the fastest and most state-aligned strategies for deployment.

The additional findings of this analysis can be summarized as:

- Abundant land is suitable for data centers
- Advanced nuclear reactors significantly expand viable siting options
- Revitalizing brownfield is both possible and practical
- Water, not land, is often the primary limiting factor.

Basically, for brownfield development, two major options exist: existing nuclear sites and coal-fired power plants. There are total of 64 existing nuclear sites in total. Of these, 54 sites are active NPPs while 11 have recently retired. Eighty-seven percent of these sites can host a 200-acre hyperscale data center within 5 miles. Only 24% can host a 2,000-acre gigawatt data center.

With respect to coal-powered plant sites, the screening process yielded a total of 146 viable sites. Ninety-three percent of these can host co-located SMRs and hyperscale data centers; 48% can host co-located LWR and gigawatt data centers. Overall, the strongest states for brownfield development are Wyoming, Indiana, and Kentucky.

Greenfield Development: An Expansion Opportunity

With regard to greenfield development, the most-flexible siting options for advanced reactors are SMRs cooled by city water. They show high suitability across the Midwest, Southeast, and parts of the Northwest, where their independence from natural cooling water is decisive. Conversely, fresh-water cooled reactors are heavily restricted west of the 100th Meridian. Large LWR technology, in particular, shows significantly reduced viability in the West due primarily to its high cooling-water demand. The dominance of water access as a siting constraint is a common throughline in this analysis.

Greenfield data centers are broadly suitable, with every state having 50% land suitability for city-water cooled data centers. Here again, freshwater-cooled data centers are more heavily constrained in the Southwest and West.

Based on the above insights, the best greenfield co-location regions include:

- The Midwest (Indiana, Illinois, Iowa, Ohio)
- The Southeast (Alabama, Georgia, Mississippi)
- The Northeast (Maine).

Key Insights

1. While both large NPPs and SMRs rely on similar processes, their supply chain composition will be inherently different. Moreover, no two SMR designs are the same, meaning that designs that leverage common fuels, technology, and storage equipment may help mitigate supply chain shocks.
2. Mapping nuclear supply chains should include identifying those components which may overlap with supporting data centers, such as electrical equipment and water distribution.
3. Secondary supply chains may still impact the speed of SMR adoption, such as rare-earth elements, processors, and transportation equipment availability. Mapping these supply chains—as they affect nuclear integration with data-center projects while seeking domestic production sources—may mitigate some of the potential risk associated with international supply chains.
4. Transportation and workforce requirements should be taken into consideration to reflect the size and scope of SMR integration with data centers and may influence ultimate project design.

Further Exploration

Further details on the national siting results are provided by Stauff et al.[37] Figures 7-1 and 7-2 of that study present a national outlook for data center siting across the 48 contiguous U.S. states based on the siting criteria developed and shown in Table 7-1. Figures 7-3 to 7-7 present a national outlook for advanced-reactor siting based on siting criteria shown in Table 7-2. The breakdown of the precent suitable land area is shown in Table 7-3. The co-location configuration analysis by percent land area per state is shown in Table 7-4. The summary of the brownfield analysis for co-location configuration is presented in Table 7-5.

More details about the OR-SAGE tool used for this national siting analysis are available.[101,102] Ongoing effort by the SA&I campaign involves case study analysis of different nuclear/data center deployment scenarios on specific sites.

4.8 Workforce and Socioeconomic Impact of Nuclear-Powered Data Centers

Based on [37].

Introduction

Co-locating a data center with a NPP does more than stabilize energy costs or strengthen grid resilience; it creates a shared industrial ecosystem with far-reaching community impacts. The economic benefits ripple outward—supporting thousands of jobs, expanding local supply chains, and generating meaningful tax revenue for regions aiming to diversify their energy portfolio. In both urban and rural settings, the combination of a continuously operating data center and a multi-decade nuclear facility anchors economic activity in a way few other industries can.

Drawing on recent analysis of recently published research on the colocation of data centers with NPPs, this section illustrates how these co-located campuses can reshape local economies, offering stability, opportunity, and long-term value for the communities that host them.

Why Does This Matter for Data Centers?

For data-center operators, the business case for nuclear co-location begins with cost structure. Power expenses account for between 58 and 64% of total data-center operating costs: the single largest expenditure by a wide margin. Labor, by contrast, represents only around 28%. Any strategy that secures long-term price-stable electricity directly addresses the industry's most significant financial exposure.

Nuclear PPAs are typically structured over 20–30-year terms at fixed prices, insulating operators from the volatility of natural-gas markets and fluctuating grid rates. High-profile agreements like the one between Microsoft and Constellation for the restart of Three Mile Island Unit 1 or Meta's agreement for the Clinton Clean Energy Center signal that this model is already being validated at scale by the industry's largest players.

Beyond cost stability, co-location offers operational advantages. Shared infrastructure for security, water systems, and administrative functions can reduce per-unit operating costs for both the nuclear facility and the data center. For hyperscale operators managing facilities that span millions of square feet and consume hundreds of MWs continuously, the ability to anchor power supply at the source, rather than depending on grid transmission, offers meaningful resilience benefits.

The socioeconomic dimension is equally important for site selection and regulatory strategy. Data-center developers increasingly face community and regulatory scrutiny around their local economic contributions. An NPP/data center campus generates substantial local employment and tax revenue, providing a tangible community benefit and associated narrative that can accelerate permitting, ease local opposition, and strengthen relationships with state and local governments. For companies with environmental, social, and governance commitments, the combination of advanced power generation and broad economic impact is a compelling story.

Research and Findings

Employment

Direct employment at a GW scale data center is estimated at approximately 1,000 workers while a co-located large reactor would employ around 650, and an SMR cluster, approximately 412. These figures, while meaningful, understate the true employment footprint. Input-output modeling reveals that every direct data center job supports an additional 5.3 jobs in the broader economy through supply-chain activity and employee spending—a total employment multiplier of 6.3. Large NPPs carry a multiplier of 4.0; SMR clusters, 5.6.

In urban settings, a combined NPP and data-center campus can support nearly 9,000 total jobs. In rural areas, this figure drops by approximately 45–47% due to thinner local supply chains, yielding roughly 4,900 total jobs—still a transformational presence in a rural county of 30,000 residents. Data-center supply-chain effects account for 60% of total employment impact, compared to 30% for nuclear facilities, reflecting the data center's higher demand for external goods and services during ongoing operations (Table 24).

Table 24. Nuclear power plant and data center employment comparison.

Facility Type	Direct Jobs	Total Jobs Supported (rural)	Total Jobs Supported (urban)
Gigawatt-scale data center	~1,000	~3,500	~6,400
Large NPP	~650	~1,360	~2,600
SMR cluster (~1,280 MWe)	~412	~1,050	~2,300
Combined data center + large NPP	~1,650	~4,900	~9,000

Workforce Comparison

For normal annual operations, data centers and NPPs both require a highly trained and dependable workforce. Based on U.S. Bureau of Labor Statistics (BLS) industry staffing pattern data, business and financial operations related jobs were the most similar across data centers and NPPs, making up 14% and 11% of the workforce, respectively.[103] Management related occupations make up 9% of the workforce at an NPP and 17% at a data center. The distribution of other occupation types is very different for the most part. Differences between the two workforces are especially evident in comparing the share of workers needed for engineering, computer, installation, production, protective service, and physical-science jobs.

Given the difference in occupational needs, it does not appear that there would be much competition for the same types of occupations between the two plants, with the exception of business and financial operations related work. There may be competition for jobs between NPPs and data centers that require lower amounts of experience and education or perhaps only require on-the-job training. This is one category of workforce competition that could benefit from further exploration. A detailed comparison of occupations needed at NPPs and data centers is available in Table 25.

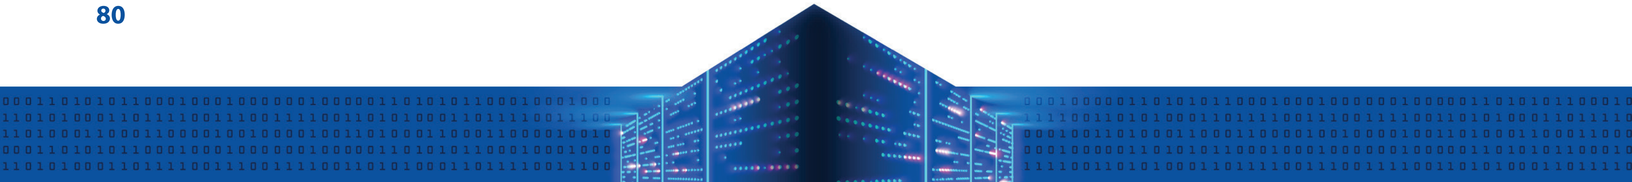

Table 25. Nuclear power and data center workforce comparison.

Occupation Category	Percentage of Workforce		Jobs Per 1,000 Workers	
	NPP	DC	NPP	DC
Architecture and engineering	16%	1%	161	9
Arts, design, entertainment, sports, and media	1%	2%	7	18
Building and grounds cleaning and maintenance	0%	0%	0	1
Business and financial operations	11%	14%	110	142
Community and social service	0%	0%	0	1
Computer and mathematical	2%	41%	21	405
Construction and extraction	3%	0%	27	0
Educational instruction and library	0%	0%	0	2
Healthcare practitioners and technical	0%	1%	0	8
Installation, maintenance, and repair	13%	1%	130	5
Legal	0%	1%	0	8
Life, physical, and social science	12%	0%	121	1
Management	9%	17%	86	172
Office and administrative support	3%	13%	34	128
Production	17%	1%	168	5
Protective service	12%	0%	121	2
Sales and related	0%	9%	0	90
Transportation and material moving	1%	0%	9	4

Labor Income

Nuclear facilities pay notably higher wages than most other electricity-generating technologies. A large NPP in an urban setting produces an estimated $208 million in direct annual labor income, and an SMR cluster generates approximately $174 million, despite each employing fewer workers than the co-located data center. The data center's direct labor income stands at $144 million, reflecting a broader, but generally lower-wage workforce.

When multiplier effects are applied, the picture shifts. The data center's higher economic multiplier drives total urban-labor income impact to $519 million annually—more than double the large reactor's $221 million contribution. In rural deployments, these figures decline substantially: the data center generates $221 million in total labor income, and the large NPP produces $146 million, underscoring the role of local economic density in determining ultimate impact.

Value Added and Total Output: Billion-Dollar Regional Anchors

Value-added impacts—representing net-new economic production stimulated within a regional economy—are substantial. In urban settings, the data center generates over $1 billion in total annual value added while the large reactor and SMR cluster each contribute approximately $664 million and $634 million, respectively. A combined campus exceeds $1.6 billion in annual urban value added. Total output impacts, including all supply chain and community spending effects, reach approximately $3.6 billion per year for a co-located urban campus: $2.4 billion from the data center and $1.2 billion from the nuclear facility.

Rural deployments see output reduced by roughly 30%, bringing combined totals to approximately $2.5 billion annually. While smaller in absolute terms, these figures can be economically transformative relative to the scale of rural-county economies.

Long-Term Stability as a Structural Advantage

Unlike boom-bust cycles associated with short-term industrial projects, co-located NPP/data center campuses are designed for multidecade operation. The 20–30-year nuclear power agreements that underpin these projects translate directly into predictable local tax revenues, stable employment, and durable economic anchoring for host communities. This stability is a structural advantage over other large-scale industrial investments and makes these facilities uniquely valuable as long-term community economic anchors.

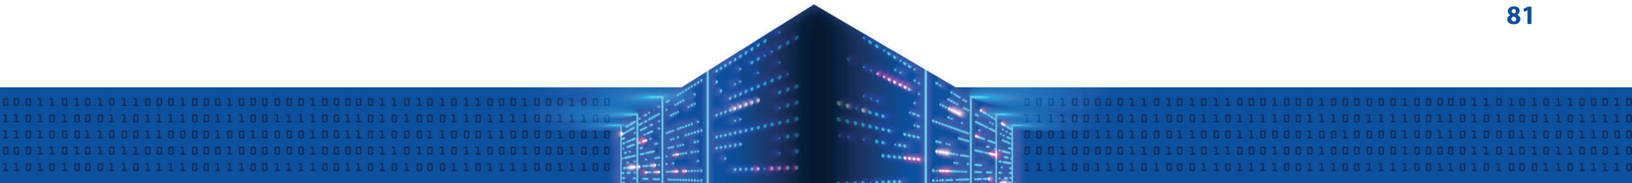

Key Insights

Beyond cost stability, co-locating with a nuclear plant enables shared infrastructure and grid independence. Crucially, the workforce profiles of the two facilities are largely complementary rather than competitive; nuclear plants skew toward engineering, physical sciences, and protective services while data centers lean heavily on computer, mathematical, and management roles.

The combined economic footprint of a co-located campus is striking: nearly 9,000 total jobs supported in urban settings, over $1.6 billion in annual value added and roughly $3.6 billion in total annual output. Rural deployments see figures about 30% lower, but remain transformational relative to local economies. Underpinning all of this is longevity; with nuclear agreements spanning 20–30 years, these campuses offer a rare combination of scale and stability that few other industrial investments can match.

CONCLUSION

The energy challenge facing the data center industry is real, urgent, and solvable, but solving it will require a fundamentally different approach to how the digital infrastructure industry thinks about power. This playbook is designed to serve as a starting point for that shift. Across its chapters, a consistent picture emerges: the scale and pace of data center growth have outrun the grid infrastructure designed to serve it, and the path forward demands new partnerships, technologies, and planning frameworks. Nuclear energy, both from the existing U.S. fleet and from the next generation of advanced reactor designs, is uniquely positioned to meet the reliability and scale requirements that define the most demanding data center applications. The conclusions below distill the key insights from this playbook and offer a framework for the decisions ahead.

The grid alone cannot keep pace. The U.S. electrical grid was not designed for the load growth now underway. Hyperscale data centers representing hundreds of megawatts, and increasingly gigawatt-scale campuses, are appearing faster than transmission lines, substations, and generation capacity can be built to serve them. Interconnection queues that once took months now stretch beyond four years in many regions. This is not a temporary backlog; it reflects a structural mismatch between the speed of digital infrastructure deployment and the multi-decade timelines of grid expansion. Data center developers who plan their power strategy around grid availability alone will find themselves competing for a constrained and increasingly expensive resource. The developers who move most successfully will be those who take a broader view of power—one that includes behind-the-meter generation, direct partnerships with power producers, and long-term contractual structures that provide the certainty both financiers and utilities require.

For more on grid constraints, interconnection timelines, and emerging policy mechanisms for large-load integration, see Chapter 1: The Role of the Grid, particularly the sections on Interconnection Planning and Communications, Utility Planning Considerations, and Capacity Expansion and Market Studies.

Nuclear power offers something the grid cannot guarantee: firm, always-on power at scale. Wind and solar are valuable contributors to the generation mix, but they are variable by nature. A data center requiring 99.982% uptime (the Tier III standard) cannot be reliably anchored to resources that depend on weather conditions. Nuclear plants in the United States run at capacity factors exceeding 90%, producing near-maximum output around the clock, year after year, with planned outages that are scheduled, predictable, and manageable. For AI training workloads, hyperscale cloud operations, and mission-critical enterprise applications, that consistency is not a preference, it is a requirement. The existing U.S. fleet of 94 reactors offers a near-term pathway through power uprates that can add meaningful capacity in 3–7 years without the licensing and construction timelines associated with new plants. For data center developers with near-term power needs, this is among the most actionable options currently available.

For more on nuclear plant reliability, capacity factors, and uprate pathways, see Chapter 4: Road to Nuclear, particularly the sections on Capacity Expansion of Existing Power Plants and Nuclear Plant Reliability and Availability. Data center uptime tiers and power architecture are discussed in Chapter 1 under Data Center Performance Requirements.

The economics of nuclear are complex but navigable, and the landscape is changing rapidly. Nuclear power is capital intensive, and cost projections for advanced reactor designs carry real uncertainty. First-of-a-kind projects cost more than nth-of-a-kind projects, and construction risk is a genuine consideration. However, the financing environment for nuclear-powered data center projects is shifting in meaningful ways. Federal loan guarantee programs, production and investment tax credits, and the demonstrated willingness of major hyperscalers to sign long-term power purchase agreements at investment-grade credit are combining to make nuclear project finance more tractable than it has been in decades. A data center operator with a strong balance sheet is, in many respects, an ideal anchor customer for a nuclear developer: creditworthy, long-tenured, and hungry for the kind of firm, abundant power that nuclear uniquely provides. When developers approach nuclear procurement as a partnership, rather than a commodity transaction, the economics tend to improve for both parties.

For more on nuclear construction costs, learning rates, financing structures, and tax incentives, see Chapter 4: Road to Nuclear, particularly New-Build Plant Costs. Financing mechanisms including DOE loan guarantees, the IRA's production and investment tax credits, and PPA structures are discussed in Chapter 3: Identifying Stakeholders, under the Financing section.

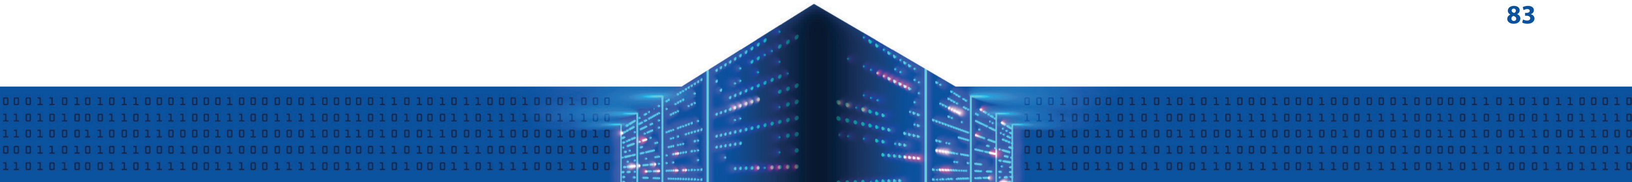

Siting, water, stakeholders, and supply chains deserve as much attention as the reactor technology itself. The technical choice of reactor design, while important, is often not the binding constraint on a nuclear-powered data center project. Water availability shapes where facilities can be built. Local zoning and community perception determine whether they can be permitted. Supply chain depth for specialized equipment—large transformers, reactor components, HALEU fuel—determines when they can be completed. The stakeholder landscape spans utilities, the Nuclear Regulatory Commission, state public-utility commissions, federal agencies, EPC contractors, and local governments. These require sustained, coordinated engagement that is categorically more complex than conventional power procurement. Projects that begin this engagement early, treat stakeholder relationships as long-term investments rather than compliance obligations, and build interdisciplinary teams capable of navigating nuclear, data center, and regulatory domains simultaneously are the ones most likely to close on time and on budget.

For more on siting analysis, water constraints, and co-location configurations, see Chapter 4: Road to Nuclear, under Site Planning and Selection, and Chapter 2 under the Energy-Water Nexus. Supply chain considerations are addressed in Chapter 4 under Nuclear Supply Chain Considerations. Stakeholder identification and engagement strategies are covered in full in Chapter 3: Identifying Stakeholders.

The workforce and community benefits of co-location are a strategic asset, not a side effect. A gigawatt-scale data center co-located with a nuclear facility can support nearly 9,000 total jobs in an urban setting and generate more than $1.6 billion in annual economic value, providing a community anchor that endures for decades. The workforce profiles of the two facilities are largely complementary rather than competitive: nuclear operations draw heavily on engineering, physical sciences, and protective services while data centers rely on computer science, mathematics, and management roles. For projects navigating permitting, community engagement, and state-level incentive negotiations, this economic footprint is a powerful asset. Communities that host these campuses gain not just construction activity but a durable, high-wage employment base—the kind of long-term economic contribution that sustains local and political support across multi-decade project lifetimes.

For more on employment multipliers, labor income, workforce profiles, and regional economic impact, see Chapter 4: Road to Nuclear, under Workforce and Socioeconomic Impact of Nuclear-Powered Data Centers.

This playbook is a starting point, not a finish line. The questions it raises do not have uniform answers. The right reactor technology, coupling configuration, financing structure, and stakeholder strategy will look different for a 50-MW edge deployment in the Mountain West than for a 2-GW AI campus in the Mid-Atlantic. What is consistent across every scenario is the need for rigorous, cross-disciplinary analysis grounded in the best available research, and for institutions capable of translating that research into project-ready guidance.

Idaho National Laboratory exists to accelerate exactly that kind of work. INL brings together expertise in nuclear engineering, grid integration, techno-economic modeling, site assessment, microgrid controls, and workforce development, and does so in direct partnership with industry. Through Department of Energy programs and direct research partnerships, INL offers data center developers access to technical depth, analytical tools, and a national laboratory network built to move projects from questions to decisions. INL invites data center developers, energy planners, and industry partners to engage: to test the findings in this playbook against specific project contexts, to identify the gaps that future research should address, to demonstrate novel energy campuses, and to build the partnerships that will shape how the next generation of digital infrastructure is powered.

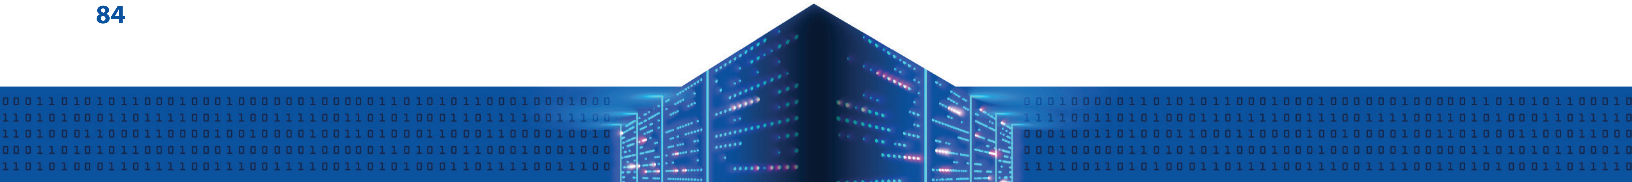

REFERENCES

1. NERC. 2025. "Characteristics and Risks of Emerging Large Loads: Large Loads Task Force White Paper." North American Electric Reliability Corporation, Atlanta, GA. https://www.nerc.com/comm/RSTCReviewItems/3_Doc_White%20Paper%20Characteristics%20and%20Risks%20of%20Emerging%20Large%20Loads.pdf.
2. L. F. Enriquez-Contreras et al. 2025. "MegaWatt Mayhem: Grid Operator Challenges with Large Loads," INL/RPT-25-87430. Idaho National Laboratory, Idaho Falls, ID.
3. R. Miller. 2024. "The Gigawatt Data Center Campus is Coming." Data Center Frontier, Accessed September 29, 2025. https://www.datacenterfrontier.com/hyperscale/article/55021675/the-gigawatt-data-center-campus-is-coming.
4. P. S. Kunder and O. P. Malik. 1994. *Power System Stability and Control*. McGraw-Hill.
5. "Tier Classification System." The Uptime Institute. Accessed March 4, 2026. https://uptimeinstitute.com/tiers.
6. CoreSite. n.d. "Breaking Down Data Center Tier Level Classification." Accessed March 4, 2026. https://www.coresite.com/blog/breaking-down-data center-tiers-classifications.
7. NERC. 2026. "Incident Review: Voltage-Sensitive Crypto Load Reductions," Accessed March 4, 2026. https://www.nerc.com/globalassets/our-work/reports/event-reports/incident_review_considering_voltage-sensitive_crypto_load_reductions.pdf.
8. R. Rahmani, I. Moser, I. and Mehdi Seyedmahmoudian. 2018. "A Complete Model for Modular Simulation of Data Centre Power Load." https://doi.org/10.48550/arXiv.1804.00703.
9. NERC. 2025. "Incident Review: Considering Simultaneous Voltage-Sensitive Load Reductions." Accessed March 4, 2026. https://www.nerc.com/globalassets/our-work/reports/event-reports/incident_review_large_load_loss.pdf.
10. P. Talbot et.al. 2025. "Navigating Integration: Key Challenges of Data Centers, Nuclear Stakeholders and Utility Operators." INL/RPT-25-87663, Idaho National Laboratory, Idaho Falls, ID. https://csdet.inl.gov/content/uploads/45/2025/10/Navigating-Integration-report.pdf.
11. C. Green et al. 2026. "Enabling Large Load Deployment through Grid Upgrades and Greenfield Solutions." INL/RPT-26-90183, Idaho National Laboratory, Idaho Falls, ID. https://inl.gov/content/uploads/2026/01/Library-Creation-and-LL-Deployment_white-paper_final-clean.pdf.
12. NERC FAC-002-3. "Facility Interconnection Studies." (approved November 17, 2022) https://nercipedia.com/standard-pending-enf/fac-002-3-facility-interconnection-studies/.
13. NERC FAC-002-3. "Facility Interconnection Studies." (approved November 17, 2022) https://nercipedia.com/standard-pending-enf/fac-002-3-facility-interconnection-studies/.
14. NERC. n.d. "Aggregated Report on NERC Level 2 Industry Recommendation: Large Load Interconnection, Study, Commissioning, and Operations." https://www.nerc.com/globalassets/programs/bpsa/alerts/2025/aggregated-report-level-2-large-load-interconnection-study-commissioning-and-operations.pdf.
15. NERC. 2023. "Reliability Guideline: Performance, Modeling, and Simulations of BPS-Connected Battery Energy Storage Systems." https://www.nerc.com/globalassets/who-we-are/standing-committees/rstc/irps/reliability_guideline_bess_hybrid_performance_modeling_studies.pdf.
16. NERC. n.d. "MOD-026-1: Verification of Models and Data for Generator Excitation Control System or Plant Volt/Var Control Functions." https://nercipedia.com/active-standards/mod-026-1-verification-of-models-and-data-for-generator-excitation-control-system-or-plant-voltvar-control-functions/.
17. A. Schirn. 2022. "IEEE 519-2022: Harmonic Control in Electrical Power Systems." American National Standards Institute. ANSI. https://blog.ansi.org/ansi/ieee-519-2022-harmonic-control-power-systems/.
18. IEEE. 2022. "IEEE Standard for Harmonic Control in Electric Power Systems." IEEE Power and Energy Society. https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=9848440&tag=1.
19. Dominion Energy. 2026. "Dominion Energy Virginia – Electric Transmission Facility Interconnection Requirements." https://www.dominionenergy.com/-/media/content/large-business-services/pdfs/virginia/facility-interconnection-requirements.pdf.
20. ATC. 2025. "Load Interconnection Guide." https://www.atcllc.com/wp-content/uploads/Load-Interconnection-Guide_Rev-14_021225_Executed.pdf.
21. PJM Decision Letter on Large Load Additions. 2026. https://www.pjm.com/-/media/DotCom/about-pjm/who-we-are/public-disclosures/2026/20260116-pjm-board-letter-re-results-of-the-cifp-process-large-load-additions.pdf.
22. California ISO. 2017. "FAQ Using Renewables to Operate a Low-Carbon Grid." https://www.caiso.com/documents/usingrenewablestooperatelowcarbongrid-faq.pdf.
23. ESIG. 2026. Large Loads Performance Requirements: Current Practices and Recommendations https://www.esig.energy/wp-content/uploads/2026/03/ESIG-Large-Loads-Performance-Requirements-report-2026c.pdf.
24. C. Mishra et al. 2025. "Understanding the Inception of 14.7 Hz Oscillations Emerging from a Data Center." *Sustainable Energy, Grids and Networks* p. 101735.
25. L. Enriquez-Contreras et al. 2025b. "Bridging the Gap for Powering Data Centers." INL/RPT-26-89901, Idaho National Laboratory, Idaho Falls, ID. https://inl.gov/content/uploads/2026/01/Bridging-the-Gap-for-Powering-Data centers.pdf.
26. M. Culler et. al. 2026. "Digital Assurance for Grid Reliability in the Era of Large Load Growth." INL/RPT-25-89589, Idaho National Laboratory, Idaho Falls, ID.
27. J. D. Wilson and Z. Zimmerman. 2023. "The Era of Flat Power Demand is Over." Grid Strategies. Accessed March 6, 2026. https://gridstrategiesllc.com/wp-content/uploads/2023/12/National-Load-Growth-Report-2023.pdf.
28. S. Morash and N. Oberski. 2025. "Power System Implications of Operating Large Loads: A Summary of Domestic and International Initiatives." INL/RPT-25-88021, Idaho National Laboratory, Idaho Falls, ID.
29. FERC. 2025. "FERC Directs Nation's Largest Grid Operator to Create New Rules to Embrace Innovation and Protect Consumers." [Online] Available: https://www.ferc.gov/news-events/news/ferc-directs-nations-largest-grid-operator-create-new-rules-embrace-innovation-and.
30. McKinsley & Company. 2023. "Investing in the rising data center economy." https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/investing-in-the-rising-data center-economy#/.
31. Aterio. U.S. Data Center Database, January 2026. https://www.aterio.io.

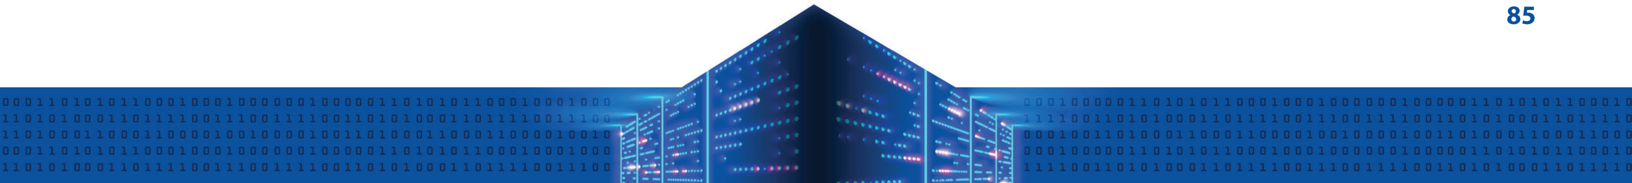

32. K. Cleary and K. Palmer. 2022. "U.S. Electricity Markets 101." Resources for the Future (RFF), Washington, DC, USA, Explainer, Mar. 3, 2020, updated Mar. 17, 2022. https://www.rff.org/publications/explainers/us-electricity-markets-101.
33. NERC. 2026. "2025 Long-Term Reliability Assessment (NERC)." https://www.nerc.com/globalassets/our-work/assessments/nerc_ltra_2025.pdf.
34. U. S. EPA. 2025. "eGRID2023 Technical Guide." https://www.epa.gov/egrid.
35. B. Patel, N. B. Desai, and S. S. Kachhwaha. 2017. "Optimization of waste heat based organic Rankine cycle powered cascaded vapor compression-absorption refrigeration system." *Energy Conversion and Management* 154: 576–590.
36. NERC. 2025d. "Incident Review: Large Load Loss Event." Princeton, NJ, Tech. Rep., https://www.nerc.com/globalassets/our-work/reports/event-reports/incident_review_large_load_loss.pdf.
37. N. E. Stauff et al. 2025. "Preliminary Analysis of Nuclear-Powered Data Center Scenarios." ANL/NSE-25/47 https://doi.org/10.2172/3001074.
38. B. N. Hanna. 2024. "Navigating Economies of Scale and Multiples for Nuclear-Powered Data Centers and Other Applications with High Service Availability Needs." Energies, 17: 5073. https://doi.org/10.3390/en17205073.
39. EPRI. 2025a. "Advanced Nuclear Technology: A Guide for Co-locating Data Centers with Nuclear Plants."
40. EPRI. 2025b. "Behind-the-Meter Large Electrical Loads: Challenges and Opportunities for Nuclear Power Plants."
41. P. W. Talbot et al. 2020. "HERON." Idaho Falls: USDOE Office of Nuclear Energy. https://doi.org/10.11578/dc.20200929.2.
42. A. Epiney et al. 2019. "Economic analysis of a nuclear hybrid energy system in a stochastic environment including wind turbines in an electricity grid." Idaho National Laboratory, Idaho Falls, ID. https://doi.org/10.1016/j.apenergy.2019.114227.
43. D. J. McDowell et al. 2021. "A technical and Econnomic Assessment of LWR Flexible Operation for Gneration and Demand Balancing to Optimize Plant Revenue." Idaho National Laboratory, Idaho Falls, ID. https://doi.org/10.2172/1844211.
44. Cho et al. 2025. "Expanding market opportunities: cogeneration strategies for integrated PWR and thermal energy storage systems." Idaho National Laboratory, Idaho Falls, ID. https://doi.org/10.1016/j.seta.2025.104631.
45. P. W. Talbot et al. 2025. "Firming Technologies for Nuclear Power Data Centers." Idaho Falls: Idaho National Laboratory.
46. Z. Cao et al. 2023. "Data center sustainability: Revisits and outlooks." *IEEE Transactions on Sustainable Computing* 9(3): 236–248. https://ieeexplore.ieee.org/document/10139829.
47. Energy Information Administration. n.d. "Nuclear explained: U.S. nuclear industry." Retrieved February 23, 2026, from https://www.eia.gov/energyexplained/nuclear/us-nuclear-industry.php.
48. V. Novotny et al. 2025. "Nuclear—Thermal Energy Storage Configurations for Industrial Combined Heat and Power Supply—Conceptual and Thermodynamic Study with High Temperature Gas-cooled Reactor." *Energy* 322(1): 135544.
49. M. Aziz et al. 2025. "Evaluating the Benefits of Nuclear–Data Center Tight Integration: Enhancing Plant Utilization via Heat-Driven Absorption Cooling." ANS Winter. Washington D.C.
50. N. Stauff et al. 2025. "Preliminary Analysis of Nuclear-Powered Data Center Scenarios." Lemont: Argonne National Laboratory.
51. N. Stauff et al. 2023. "Assessment of nuclear energy to support negative emission technologies." Lemont: Argonne National Laboratory.
52. R. Gomri. 2010. "Investigation of the potential of application of single effect and multiple effect absorption cooling systems." *Energy Conversion and Management* 51(8): 1629–1636.
53. A. Trafone et al. 2018. "Liquid air energy storage performance enhancement by means of organic Rankine cycle and absorption chiller." *Applied Energy* 228: 1810–1821.
54. C. B. Dorgan, C. E. Dorgan, and S.P. Leight. 1995. "ASHRAE`s new application guide for absorption cooling/refrigeration using recovered heat." *ASHRAE Journal* 37(7).
55. A. Shirazi et al. 2016. "A systematic parametric study and feasibility assessment of solar-assisted single-effect, double-effect, and triple-effect absorption chillers for heating and cooling applications." *Energy Conversion and Management* 114: 258–277.
56. EPRI. 2025c. "Nuclear Provided Steam for Data Center Cooling: A Feasibility Study on Integrating Low-Grade Nuclear Steam with Steam Absorption Chillers."
57. EIA. n.d. "Electric Power Monthly, Table 5.3. Average Price of Electricity to Ultimate Customers." Retrieved September 15, 2025, from https://www.eia.gov/electricity/monthly/epm_table_grapher.php?t=epmt_5_3.
58. DOE. 2014. "The Water-Energy Nexus: Challenges and Opportunities." https://www.energy.gov/sites/prod/files/2014/07/f17/Water%20Energy%20Nexus%20Full%20Report%20July%202014.pdf.
59. A. Shehabi et al. 2024. "2024 United States Data Center Energy Usage Report." Lawrence Berkeley National Laboratory. Report No: LBNL-2001637. http://dx.doi.org/10.71468/P1WC7Q. Retrieved from https://escholarship.org/uc/item/32d6m0d1.
60. R. Iftekhar and O. Browne. 2025. "Estimating Data Center Water Demand." Water Market Insider. Westwater Research. https://westwaterresearch.com/wmi-estimating-data-center-power-demand/.
61. K. G. Cafferty et al. 2025. "Community Water-Power Resilience: Ownership and Management Structures influence." Idaho DEQ Water Quality Workshop.
62. M. A. Harris and T. H. Diehl. 2019. "Withdrawal and Consumption of Water by Thermoelectric Power Plants in the U.S." Report No.: 2328-0328. US Geological Survey.
63. INL. n.d. "Accelerating safe deployment of small nuclear reactors for grid use." Idaho National Laboratory, Idaho Falls, ID. https://inl.gov/content/uploads/2025/12/25-315_Accelerating-safe-deployment-of-SMRs-for-grid-use.pdf.
64. M. Ton, B. Fortenbery, and W. Tschudi. 2008. "DC Power for Improved Data Centers Efficiency.(DC power demo project)." Lawrence Berkeley National Laboratory. https://datacenters.lbl.gov/sites/default/files/DC%20Power%20Demo_2008.pdf.
65. B. A. Ross and J. Follum. 2026. "Electromagnetic Transient Modeling of Large Data Centers for Grid-Level Studies." Pacific Northwest National Laboratory PNNL-38817. https://www.pnnl.gov/main/publications/external/technical_reports/PNNL-38817.pdf.
66. Data Center Knowledge. n.d. "Current debate: Will the data center of the future be AC or DC?" Accessed: March 27, 2026. https://www.datacenterknowledge.com/energy-power-supply/current-debate-will-the-data-center-of-the-future-be-ac-or-dc.
67. Lawrence Berkeley National Laboratory. n.d. "Calculating energy savings using high-voltage direct current (HVDC) power distribution in data centers." Accessed: March 27, 2026. https://datacenters.lbl.gov/resources/calculating-energy-savings-using-high-0.
68. D. Robb. 2026. "In Edison's Revenge, Data Centers Are Transitioning From AC to DC." IEEE Spectrum. https://spectrum.ieee.org/data-center-dc.
69. NRC. n.d. "Approved Applications for Power Uprates." https://www.nrc.gov/reactors/operating/licensing/power-uprates/status-power-apps/approved-applications.

70. F. C. Joseck et al. 2026. "Regional Market Assessment for Power Offtake Agreements." INL/RPT-26-91053." Idaho National Laboratory, Idaho Falls, ID (available on request).
71. Electric Power Research Institute (EPRI). 2023. "Facilitating Power Uprates at Nuclear Power Plants: Feasibility Study Guideline."
72. I. Lantgios. 2025. "A Technical and Economic Analysis of Integrated Nuclear-DAC Systems." Doctoral Dissertation, University of Pittsburgh.
73. DOE. 2026. "U.S. Department of Energy Awards $2.7 Billion to Restore American Uranium Enrichment." https://www.energy.gov/articles/us-department-energy-awards-27-billion-restore-american-uranium-enrichment.
74. DOE-NE. 2020. "Infographic: Advanced Reactor Development: Paving a Path to Commercialization." https://www.energy.gov/ne/articles/infographic-advanced-reactor-development.
75. U.S. DOE-NE. n.d. "U.S. Department of Energy Reactor Pilot Program." https://www.energy.gov/ne/us-department-energy-reactor-pilot-program.
76. A. Abou-Jaoude et al. 2024. "Meta-Analysis of Advanced Nuclear Reactor Cost Estimations." INL/RPT-24-77048. Idaho National Laboratory, Idaho Falls, ID: Gateway for Accelerated Innovation in Nuclear. https://doi.org/10.2172/2371533.
77. A. Abou-Jaoude et al. 2023. "Literature Review of Advanced Reactor Cost Estimates." https://doi.org/10.2172/1986466.
78. G. L. Lentz, H. W. Buschman, and R. N. Smith. 1985. "EBR-II: Twenty Years of Operating Experience." CONF-850722-1, DE85 013876. Idaho Falls, ID: Argonne National Laboratory.
79. International Atomic Energy Agency (IAEA). n.d. "Power Reactor Information System (PRIS)." Accessed February 16, 2026. https://pris.iaea.org/.
80. M. Turner. 2025. "Loudoun County, Virginia: Data Center Capital of the World." Loudoun County Virginia Board of Supervisors. https://www.loudoun.gov/ArchiveCenter/ViewFile/Item/13979.
81. National Archives and Records Administration. 2026. "Code of Federal Regulations: Title 10, Chapter I." Updated on March 2, 2026. https://www.ecfr.gov/current/title-10/chapter-I.
82. U.S. NRC. 2026. "Part 53—Risk Informed, Technology-Inclusive Regulatory Framework for Advanced Reactors." Accessed March 2, 2026. https://www.nrc.gov/reactors/new-reactors/advanced/modernizing/rulemaking/part-53.
83. S. E. Aumeier et al. n.d. "Microreactor Applications in U.S. Markets: Evaluation of State-Level Legal, Regulatory, Economic and Technology Implications." INL/RPT-23-71733, Idaho National Laboratory, Idaho Falls, ID. https://www.osti.gov/biblio/1964093
84. B. N. Hanna et al. 2025. "Navigating Economies of Scale and Multiples for Nuclear-Powered Data Centers and Other Applications with High Service Availability Needs." Energies 17(20): 5073. https://www.mdpi.com/1996-1073/17/20/5073.
85. J. S. Kang et al. 2025. "Comprehensive review of small modular reactor development focusing on challenges in the backend nuclear fuel cycle." Nuclear Engineering and Technology. October 24, 2025. DOI: 10.1016/j.net.2025.104000. https://www.sciencedirect.com/science/article/pii/S1738573325005686?via%3Dihub.
86. International Energy Agency (IEA). 2026. "Energy and AI." Accessed January 15, 2026. https://iea.blob.core.windows.net/assets/601eaec9-ba91-4623-819b-4ded331ec9e8/EnergyandAI.pdf.
87. S. Gallier. 2025. "U.S. nuclear capacity factors: Stability and energy dominance." Nuclear Newswire. American Nuclear Society. https://www.ans.org/news/2025-05-02/article-6967/us-nuclear-capacity-factors-stability-and-energy-dominance/.
88. Institute of Electrical and Electronics Engineers (IEEE). 2007. IEEE Std 493-2007, "Design of Reliable Industrial and Commercial Power Systems." Also called the "IEEE Gold Book."
89. N. Johnson and Z. Ma. 2024a. "Initiating Event Rates at U.S. Nuclear Power Plants: 2023." Update. INL/RPT-24-04723, Idaho National Laboratory, Idaho Falls, ID.
90. N. Johnson and Z. Ma. 2024b. "Analysis of Loss-of-Offsite-Power Events 2023." Update. INL/RPT-24-79969, Idaho National Laboratory, Idaho Falls, ID.
91. J. Doyle et al. 2020. "Highly Available Nuclear Power for Mission-Critical Applications." *Nuclear Technology* 206 (7), 1059–1074. DOI: 10.1080/00295450.2019.1699382.
92. U.S. IEA. 2023. "U.S. summer nuclear outages rose in 2023, returning to 2021 levels." https://www.eia.gov/todayinenergy/detail.php?id=60682.
93. Constellation. n.d. "Nuclear Reliability." Retrieved February 26, 2026, https://www.constellationenergy.com/work/generation/nuclear/reliability.html.
94. Westinghouse. n.d. "Improved Nuclear Power Plant Operations." Accessed February 27, 2026, https://westinghousenuclear.com/new-plants/ap1000-pwr/operations-and-maintenance/.
95. Y. Miao et al. 2019. "Report on Assessment of Advanced Reactor Fuel Performance." ANL/NSE-19/37. https://publications.anl.gov/anlpubs/2019/11/156360.pdf.
96. Gen IV International Forum. n.d. "Sodium Fast Reactor (SRF)." Accessed February 27, 2026, https://www.gen-4.org/generation-iv-criteria-and-technologies/sodium-fast-reactor-sfr.
97. NuScale. n.d. "NuScale Power Module." Accessed February 27, 2026, https://www.nuscalepower.com/products/nuscale-power-module#:~:text=Product%20Overview,Scalable%20Design.
98. T. Bergman et al. 2016. "The NuScale Design." https://www.nrc.gov/docs/ML1616/ML16161A723.pdf.
99. Westinghouse. n.d. "Only SMR based on Licensed, Operating & Advanced Reactor Technology." Accessed February 27, 2026, https://westinghousenuclear.com/new-plants/ap300-smr/#:~:text=The%20Westinghouse%20AP300%20SMR%20delivers,simplify%20operations%20and%20increase%20reliability.
100. M. Modarres, M. P. Kaminsky, and V. Krivstov. 2016. *Reliability Engineering and Risk Analysis: A Practical Guide*. 3rd edition. CRC Press.
101. O. A. Omitaomu et al. 2022. "Method and System for Siting Advanced Nuclear Reactors and Evaluating Energy Policy Concerns." *Progress in Nuclear Energy* 148: 10419.
102. R. J. Belles et al. 2012. Updated Application of Spatial Data Modeling and Geographical Information Systems (GIS) for Identification of Potential Siting Options for Small Modular Reactors." ORNL/TM-2012/403. Oak Ridge National Laboratory. https://info.ornl.gov/sites/publications/Files/Pub39008.pdf.
103. BLS. 2025. "Industry Occupation Matrix." Accessed March 4, 2026, https://www.bls.gov/emp/tables/industry-occupation-matrix-industry.htm.

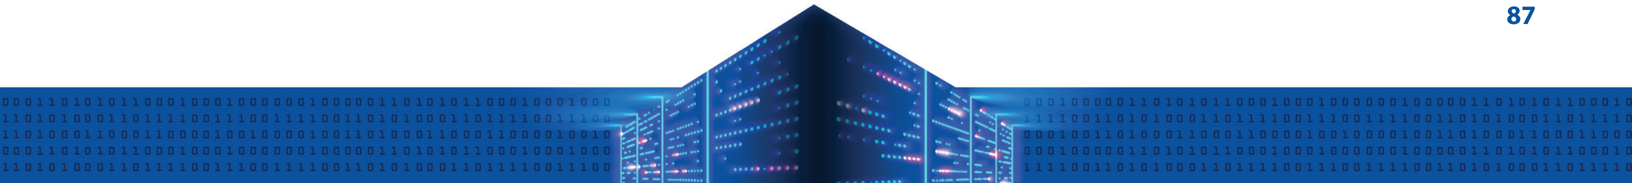

Idaho National Laboratory