Cell Reports Sustainability

CellPress
OPEN ACCESS

Perspective

The role of small modular reactors in powering data centers for a sustainable digital economy

Roshni Anna Jacob,[1,3] Honglin Li,[1,3] Ali Mahboub Rad,[1] Bikash Poudel,[2] Yaoyu Li,[1] Mucun Sun,[2] Timothy R. McJunkin,[2] and Jie Zhang[1,*]

[1]The University of Texas at Dallas, Richardson, TX 75080, USA
[2]Idaho National Laboratory, Idaho Falls, ID 83415, USA
[3]These authors contributed equally
*Correspondence: jiezhang@utdallas.edu
https://doi.org/10.1016/j.crsus.2026.100820

SUMMARY

This paper presents a critical perspective on the potential of small modular reactors (SMRs) to address the rapidly increasing energy demands of the digital economy, particularly those driven by advancements in artificial intelligence (AI). Drawing on expertise in techno-economic analysis and the modeling and control of advanced nuclear reactors and integrated energy systems, we explore the prospects for deploying SMR technology to power energy-intensive data centers. We examine practical implementation pathways and identify the key technical, regulatory, and operational challenges that must be addressed for effective integration. The SMR-data center discussion is grounded on SMR designs, deployment configurations, grid integration strategies, and interactions with data center cooling systems. This perspective aims to assess the potential of SMR-powered data centers as a scalable and forward-looking solution to escalating electricity demands, emphasizing their role in enhancing energy resilience and sustainability.

INTRODUCTION

The accelerating growth of artificial intelligence (AI), cloud services, and high-performance computing (HPC) is transforming data centers into mission-critical infrastructures with unprecedented energy demands. As shown in Figure 1, in 2023, United States (US) data centers, approximately 4,500 facilities consuming roughly 176 TWh or 4.4% of national electricity, are projected to rise to 6.7%–12% by 2028, nearly tripling their share within 5 years. Requests to utilities for new hyperscale campuses now range between 500 MW and 2 GW, rivaling the demand of large metropolitan areas.[1] Planned expansions are highly concentrated, with announced new capacity already surpassing 15 GW in Virginia, 8 GW in Pennsylvania, and 5 GW in Texas, while many other states report no public figures. These developments indicate that data centers will constitute not only one of the fastest-growing loads but also an essential pillar of the future energy system, requiring integrated planning alongside other large-scale demands such as transportation and industry. This rapid load growth, alongside the electrification of transportation, heating, and new industrial processes, intensifies pressure on power systems, already constrained by limited transmission capacity, renewable integration challenges, and affordability concerns. Meeting such simultaneous demands without undermining grid reliability or driving up customer energy costs has emerged as a defining challenge for utilities and policymakers.

From a sustainability perspective, rapidly growing data-center electricity demand is not only a capacity-planning challenge but also a lock-in problem for digital infrastructure. If new campuses rely on constrained grids, carbon-intensive backup generation, and inefficient thermal management, AI-driven expansion could deepen emissions and infrastructure burdens. Small modular reactor (SMR)-powered data centers therefore warrant evaluation as integrated energy systems rather than as isolated power-procurement decisions. Co-located firm low-carbon generation may reduce transmission losses and exposure to bottlenecks while improving resilience, especially where campus loads grow faster than grid infrastructure. Reactor heat and recoverable data-center waste heat could also support cooling and thermal-energy services, placing SMR-data center integration within a broader sustainability strategy.

In response, advocacy has grown for on-site generation strategies that can decouple the energy consumption of data centers from the congested bulk grid. Among these, SMRs have attracted interest due to their modular scalability, low-carbon generation profile, siting flexibility, and ability to provide not only electricity but also integrated heating and cooling.[2–8] Most operating data centers today, however, remain primarily grid-connected facilities, typically supported by uninterruptible power supply (UPS) systems and conventional backup generation rather than dedicated nuclear assets.[9] In this context, conventional large-scale nuclear plants and SMRs should be understood as distinct rather than interchangeable options for supplying firm low-carbon power to digital infrastructure. Large gigawatt-scale reactors benefit from commercial maturity, economics of scale, and the ability to support very large

This is an open access article under the CC BY-NC license (http://creativecommons.org/licenses/by-nc/4.0/).

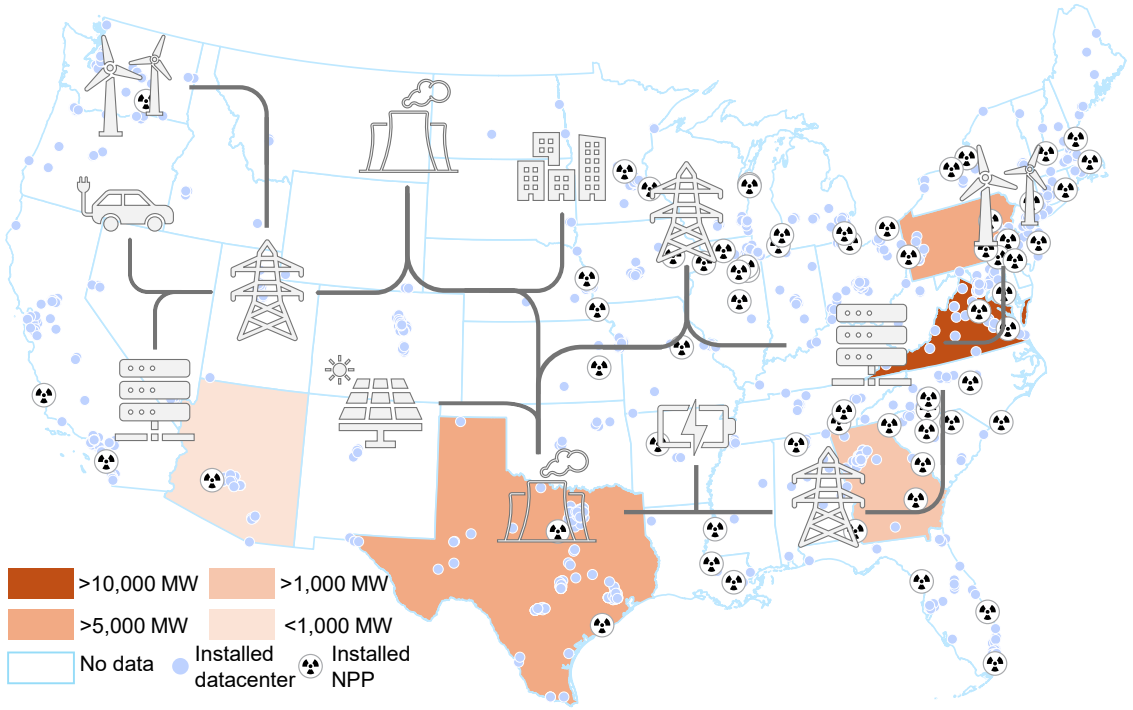


Figure 1. Overview of the current and emerging US landscape for SMR-powered data centers

Blue dots indicate currently operating data centers. State shading represents publicly announced projected data-center construction capacity over the next 5 years, grouped by capacity range, while unshaded states indicate that no public capacity estimate was identified. The map also shows the locations of currently operating large-scale nuclear power plants in the US, as indicated in the legend, to highlight their spatial relationship to existing data-center clusters and planned high-demand regions.

aggregated state- or metropolitan-scale loads through the bulk grid; however, their high absolute capital requirement, large site and cooling-water needs, long development timelines, and coarse unit sizing can make them less suitable for phased hyperscale campus expansion projects, especially when new facilities are emerging in the 500 MW to 2 GW range. SMRs, by contrast, are more naturally aligned with incremental campus growth, multi-module redundancy, closer siting to load, and integrated provision of electricity, heat, and cooling, although they also face important disadvantages, including lower commercial maturity, greater early-deployment cost uncertainty, and unresolved licensing and supply-chain risks. Critics often emphasize that SMRs are not yet commercially available at scale, with deployment pathways still under development. However, this uncertainty parallels the long and capital-intensive timelines associated with expanding transmission and distribution infrastructure. Moreover, SMR-data center hubs offer locational and functional advantages that grid build-out alone cannot provide, such as direct co-siting of generation with load,[10] integration of heat and cooling,[11] and resilience against systemic grid disturbances.[12] Accordingly, the central planning question is not whether SMRs should categorically replace large nuclear plants, but which nuclear form better matches local load concentration, phasing needs, infrastructure constraints, and the desired degree of co-located autonomy. Thus, rather than considering SMRs as an alternative to grid investment, the productive view is to explore them as a complementary solution, broadening the portfolio of options to meet the AI-driven data center demand.

Figure 2 highlights four advanced SMR designs with potential for data center integration: integral pressurized water reactor (iPWR), liquid-metal reactor, high-temperature gas-cooled reactor (HTGR), and molten salt reactor (MSR). All share common attributes critical for deployment near digital infrastructure, such as compact modular construction, and inherent or passive safety features. Their distinguishing features lie in coolant choice, operating pressure, and outlet temperature. iPWRs operate with light water at approximately 15 MPa and outlet temperatures near 300°C. Their integral design incorporates the steam generators and pressurizer within the reactor vessel, with the secondary water or steam loop routed externally to the balance of plant. Liquid-metal reactors operate at low pressure, typically below 1 MPa, with outlet temperatures of about 500°C–550°C for sodium and 480°C–560°C for lead or lead-bismuth coolants. A pool-type core configuration is commonly employed, incorporating an intermediate liquid-metal loop to transfer heat while isolating the primary circuit from water and thereby enhancing safety. HTGRs use helium coolant at roughly 7 MPa with outlet temperatures of 750°C–900°C. A representative design is the pebble-bed core, in which graphite spheres embedded with tristructural isotropic (TRISO)-coated fuel particles circulate slowly through the vessel, enabling continuous online refueling. This configuration provides strong fission-product retention and high thermal inertia. MSRs circulate fuel salt at near

CellPress
OPEN ACCESS

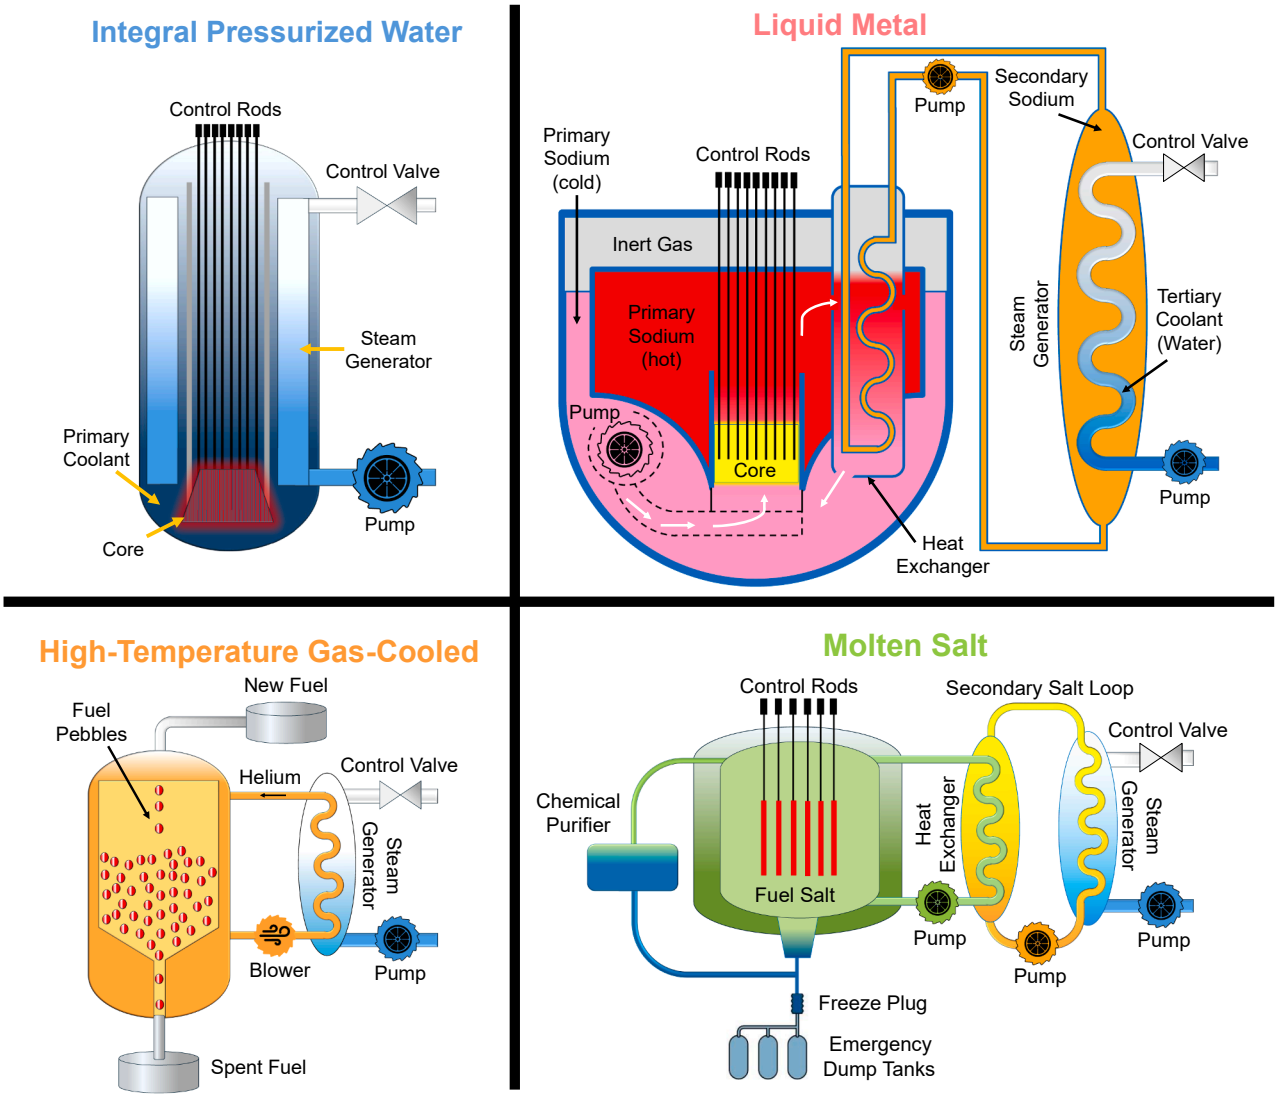


Figure 2. Four types of SMRs: Integral pressurized water, liquid metal, high-temperature gas-cooled, and molten salt

atmospheric pressure with 600°C–700°C outlet temperatures. A typical configuration includes a primary fuel-salt loop, a secondary clean-salt loop feeding a steam generator, and a dedicated drain tank that provides passive safety functions and can interface directly with thermal energy storage (TES). These characteristics create pathways to couple each design with data center power supply and heat recovery systems.

To organize these interdependent questions, we frame SMR-data center feasibility as a layered decision problem spanning four coupled dimensions: economic viability (reliability premiums, service-aligned metrics, first-of-a-kind (FOAK)-to-nth-of-a-kind (NOAK) cost divergence, and co-product revenues), architectural viability (grid coupling, storage, and cooling configuration), operational viability (dispatch, workload flexibility, and control), and institutional viability (licensing, market rules, and governance). This perspective does not attempt to quantify these dimensions through a single new case-study model; rather, it synthesizes them as a common framework for evaluating future SMR-data center pathways, including grid-connected, hybrid, and islanded modes. Regulatory requirements act as boundary conditions within this framework, shaping not only institutional feasibility but also system architecture, operational flexibility, and economic viability. In particular, licensing, siting, and grid-interconnection constraints determine which deployment configurations are feasible and directly influence project costs, development timelines, and overall system feasibility. The merit of any proposed configuration therefore depends on the consistency of these layers rather than on strong performance in any single dimension.

This paper aims to assess the potential of SMR-powered data centers as a scalable and forward-looking solution to escalating electricity demands, discussing challenges and opportunities across techno-economic analysis (TEA), advanced reactor modeling and control, integrated energy-system grid integration,

and regulation and governance. The remainder of this paper is organized as follows. Future of data centers powered by SMRs: Techno-economic feasibility and investment value section discusses the techno-economic feasibility of SMR-powered data centers, emphasizing the need to account for reliability premiums, broadened evaluation metrics beyond levelized cost of electricity (LCOE), FOAK-NOAK cost divergences, and multi-product value streams. Design reliable and energy efficient SMR-integrated data centers: Architecture and operation section turns to architectural and operational considerations, outlining how different grid-integration modes, storage technologies, and advanced cooling strategies can shape reliable and efficient SMR-data center designs. SMR-data center hubs: Dynamic interaction with the grid and AI coordination section expands the lens to system-level interactions, analyzing how SMR-powered hubs interact dynamically with the grid, how workload scheduling can enable demand response, and how AI may support the control of SMR-data center facilities. Regulation and governance section highlights the regulatory and governance environment, detailing licensing challenges, institutional actors, and evolving policy frameworks that shape deployment. Finally, conclusions section synthesizes the perspective, positioning SMR-data centers as multi-energy, cyber-physical infrastructures and pointing to future research and implementation pathways.

FUTURE OF DATA CENTERS POWERED BY SMRs: TECHNO-ECONOMIC FEASIBILITY AND INVESTMENT VALUE

Conventional performance metrics for data centers fail to capture the full spectrum of resilience and sustainability challenges. The most common, power usage effectiveness (PUE), the ratio of total facility energy demand to information technology (IT) energy use, offers only a narrow measure of cooling and power delivery efficiency, while ignoring IT infrastructure performance. Variants such as data center infrastructure efficiency add little new insight, while newer metrics like effective power utilization and grid usage effectiveness extend the scope to include workload efficiency and grid dependence.[13] In parallel, tier classifications (tier I–IV) provide a standardized view of redundancy and fault tolerance, but they remain limited to availability, offering little connection to efficiency or sustainability. Beyond these operational metrics, current TEA practices also fall short. They seldom treat reliability as a priced design constraint, e.g., the redundancy ($N + 1/N + 2$), staggered maintenance, and outage bridging strategies required for tier-level availability, leading to an underestimation of the true "reliability premium" of SMR-powered architectures. Similarly, generation-centric measures such as the LCOE overlook the multi-product nature of data center energy services, where cogenerated heat can offset electric cooling and substantially reshape both cost and emissions. Addressing these gaps requires service-aligned metrics (e.g., a levelized cost of computing) that integrate power, cooling, and availability into a single decision frame. Yet, evaluating data center configurations remains challenging, as no single indicator captures the interplay of efficiency, reliability, sustainability, and adaptability. This gap underscores the need for a holistic perspective, particularly for emerging SMR-data center ecosystems.

Reliability premium as a critical economic dimension

Hyperscale data centers require exceptional power reliability. Tier IV certification, for example, demands 99.995% annual availability,[14,15] corresponding to fewer than 26 min of allowable downtime per year. This reliability standard is not merely a technical benchmark, and it reflects the staggering economic and reputational losses incurred by service interruptions in sectors such as finance, healthcare, and cloud computing. Even with the historically high capacity factors of nuclear, often exceeding 90%, single-unit configurations remain insufficient for meeting tier IV thresholds due to periodic planned outages (e.g., refueling and inspections) and unplanned contingencies (e.g., equipment failure and cyber-physical disruptive events).[16,17] Unlike conventional fossil plants, SMRs face unique constraints: reactor refueling requires complete or partial shutdowns, and the limited size of modular units reduces the feasibility of internal load redistribution without dedicated redundancy.

To meet these uptime requirements, SMR-powered systems must adopt redundancy architectures, typically $N + 1$ or $N + 2$ configurations, wherein additional modules are installed solely to ensure uninterrupted power delivery in the event of primary unit unavailability. For instance, powering a 300 MW hyperscale data center under an $N + 1$ architecture may require 400 MW of installed SMR capacity, with the extra 100 MW unit idling under normal operations or sending power to the power grid. This engineering necessity introduces a substantial "reliability premium," inflating capital expenditures (CAPEX) by 25%–30% without yielding additional computational throughput or revenue. Standard LCOE assessments typically omit this premium, thus underestimating life cycle costs and mischaracterizing the competitiveness of nuclear-powered architectures compared with alternative supply models (as seen in Table 1).[18–23]

Battery energy storage systems (BESSs) can bridge short-duration interruptions—on the order of seconds to minutes—thereby reducing dependency on spinning or hot standby modules. However, BESS cannot substitute for full-module redundancy during extended shutdowns such as scheduled refueling, which may last multiple days or weeks. Similarly, hybrid integration with the grid or diesel-based backup generators introduces external dependencies and emissions risks, undermining the rationale for SMR co-location in the first place. Moreover, reliance on external grids may not guarantee tier IV performance due to grid congestion, cyber vulnerabilities, or extreme weather events—a growing concern under climate uncertainty.

From a TEA perspective, accurately quantifying the reliability premium is essential to understanding true system costs. This requires explicit modeling of capacity oversizing, redundant module deployment schedules, and load-balancing strategies under various failure and maintenance scenarios. In addition, TEA models should evaluate the impact of staggered refueling cycles (e.g., offsetting refueling windows among multiple modules) and advanced modular dispatch strategies (e.g., rotating baseload duty among reactors) to reduce idle capacity. These design choices may optimize the trade-off among CAPEX,

CellPress
OPEN ACCESS

Table 1. Reliability requirements and potential premium across data center types

Type/scale	Typical tier	Reliability importance	Potential premium
Hyperscale (cloud)	III–IV	mission-critical, global services; downtime = multi-million losses	very high; justify 2N redundancy and long backup
Financial	IV/dual-active	zero-tolerance; seconds of outage disrupt markets	extremely high; premium >$M/hour avoided outage
Large enterprise	III	core business systems; limited downtime acceptable	medium-high; pay for N+1 redundancy
Govt. (critical)	III–IV	public safety, national security	high; subsidized premium for dual-grid, priority power
Govt. (non-critical)	II–III	admin/web services; tolerable short outages	low-medium; budget constrained
Small and medium-sized enterprises (SMEs)	I–II	limited IT; higher tolerance, but outages risky	low absolute; rely on cloud for reliability
Edge	I–II	single node less critical; network redundancy key	medium; premium for distributed resilience

operating expenditure (OPEX), and availability, especially when paired with TES or flexible load-shifting algorithms.

Expanding evaluation metrics for integrated systems

Traditional TEA for nuclear systems often revolves around the LCOE, a metric that estimates the per-megawatt-hour cost of generating electricity over a facility's lifetime. While useful for comparing centralized generation assets in bulk power markets, LCOE provides an incomplete picture for integrated applications such as SMR-powered data centers. Electricity represents only part of the energy service demanded by digital infrastructure. In hyperscale facilities, cooling typically accounts for 30%–40% of total energy consumption, sometimes even more in hot and humid climates, making thermal management an operational and economic priority on par with electrical supply.[24]

SMRs inherently produce significant quantities of low-grade thermal energy as a byproduct of electricity generation. In most conventional deployments, this heat is rejected through cooling towers or condensers, contributing neither to revenue nor to efficiency. However, when co-located with data centers, this thermal output can be harnessed through absorption chillers, heat-driven refrigeration cycles, or hybrid cooling loops to satisfy part or all of the facility's cooling demand. This integrated approach not only enhances overall energy utilization but can also reduce reliance on electric chillers, which are both capital- and energy-intensive.

In cogeneration or trigeneration configurations, total system efficiency can reach 70%–85%, compared with roughly 35%–40% in electricity-only systems.[24,25] Such efficiency gains directly translate into reduced primary energy input per unit of useful output, with implications for both operating cost and emissions footprint. In the context of carbon-constrained digital infrastructure, these thermal synergies are economically and environmentally consequential.

To capture this broader service value, we propose the levelized cost of computing (LCOC) as a more appropriate metric than LCOE. LCOC measures the cost per unit of computational service, such as compute-hour or teraflop-hour, by integrating the cost of both electricity and cooling services into a single, service-oriented denominator.[26] This reframing enables direct comparison between different power architectures in terms of their actual contribution to data center operations. For example, if SMR-powered absorption chillers displace 30% of electric cooling loads, the LCOC could be significantly lower than a standalone LCOE-based estimate, even if the LCOE itself is marginally higher than grid benchmarks.

Beyond improved cost granularity, LCOC internalizes key design trade-offs. For instance, higher CAPEX for SMR-cooling integration may be offset by lower operational energy costs, fewer emissions-related penalties, and smaller chiller footprints, all of which have implications for site planning and life cycle profitability. LCOC also enables apples-to-apples comparison with alternative architectures. While renewable-plus-storage systems have achieved impressive reductions in LCOE, they often rely on electrically powered chillers, requiring additional battery capacity or diesel backup to meet tier IV reliability and thermal loads. These auxiliary systems introduce complexity, cost, and space constraints that are difficult to capture in LCOE alone.

Adopting LCOC (or similar service-aligned metrics) also enhances the rigor of TEA modeling. Instead of modeling electricity production in isolation, LCOC frameworks evaluate end-use alignment, thermal load displacement, and system-level performance indicators. This approach better reflects the interdependent nature of digital infrastructure energy demands, particularly in designs where heat, electricity, and reliability are co-optimized. From this perspective, combined storage-and-cooling configurations are economically justified not by technical integration alone, but when the additional capital cost, control complexity, and maintenance requirements of TES, absorption chillers, chilled-water storage, and associated balance-of-plant equipment are offset by measurable reductions in electric cooling demand, avoided chiller oversizing, lower peak-power exposure, and improved resilience value. Their attractiveness is therefore expected to be greatest in facilities with large and sustained cooling loads, reactor heat available at usable temperature levels, and operating environments in which electricity prices, carbon constraints, or reliability premiums are sufficiently high for these avoided costs to matter materially. Sensitivity analyses can be structured around key LCOC drivers, e.g., thermal recovery ratio, coefficient of performance (COP) of absorption

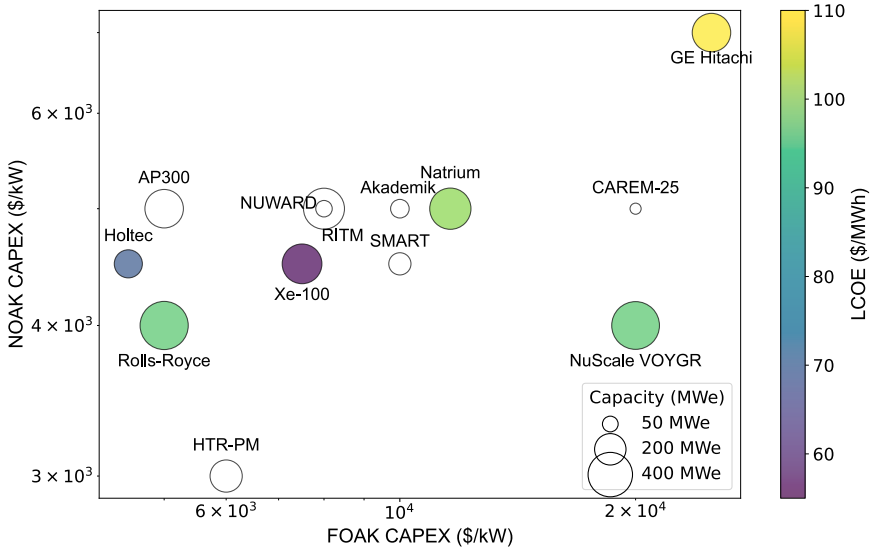


Figure 3. Comparison of FOAK and NOAK CAPEX for selected SMR projects

Bubble size is proportional to reactor capacity (MWe), while color shading represents the LCOE. Unfilled bubbles indicate projects for which LCOE data are unavailable.[36–40]

chillers, and marginal cost of chiller capacity to inform investment decisions and modular system sizing.

Cost estimation across system life cycles

A defining challenge in SMR TEA is the stark divergence between FOAK and NOAK costs.[27] While proponents of SMR deployment often highlight NOAK LCOE estimates in the range of 44–54 $/MWh, these figures are premised on long-term cost improvements enabled by factory fabrication, modular assembly, serial deployment, and standardized regulatory pathways.[28,29] In contrast, early FOAK deployments have exhibited cost profiles that are not only substantially higher but also more volatile, shaped by unproven supply chains, one-off engineering workarounds, and protracted licensing timelines. For example, the NuScale-UAMPS (Utah Associated Municipal Power Systems) Carbon Free Power Project, once hailed as a flagship modular deployment, was ultimately canceled after projected LCOE estimates escalated to nearly 89 $/MWh, and capital costs surpassed $20,000/kW due to inflationary pressures, supply-chain constraints, and shifting demand projections.[30–32] Similarly, the BWRX-300 FOAK deployment at Darlington in Canada has reported all-in capital costs approaching $17,000/kW, more than triple the typical capital intensity of mature nuclear installations.[33,34]

For hyperscale operators evaluating SMR integration, these cost gaps are not abstract, and they represent direct economic exposure. Acting as anchor offtakers or early co-investors places data center developers in the role of *de facto* venture financiers, absorbing risks typically borne by governments or utilities in traditional nuclear programs. This includes exposure to budget overruns, multi-year schedule delays, evolving safety requirements, and first-time commissioning uncertainties, all of which complicate financial modeling and increase the cost of capital.[35] Moreover, the lack of historical performance data for FOAK SMR units reduces lender confidence, making debt financing more expensive or unavailable. Without explicit treatment of FOAK-NOAK cost divergence, TEA results risk conveying an overly optimistic view of project viability that may fail under real-world implementation.

A credible TEA framework must therefore model FOAK and NOAK scenarios as distinct cases with separate cost curves, risk profiles, and financing structures. This includes incorporating FOAK-specific capital contingencies, stress-tested discount rates, insurance premiums, and flexible commissioning timelines. Cash flow models should account for milestone-dependent disbursements, stepwise CAPEX release tied to construction progress, and delayed revenue realization under phased deployment. Scenario analysis can be used to evaluate sensitivities to CAPEX escalation, licensing delays, and module failure rates—factors that disproportionately affect early-stage projects (Figure 3).

Mitigating these challenges requires the active involvement of institutional mechanisms designed to de-risk early deployment. Cost-sharing programs such as the US Department of Energy's Advanced Reactor Demonstration Program (ARDP) provide direct capital support and accelerated licensing pathways, lowering FOAK project barriers.[36] Federal loan guarantees and credit support facilities reduce the cost of borrowing and improve bankability. A long-term power purchase agreement with indexed pricing and take-or-pay structures can ensure predictable cash flows for project sponsors and investors. In parallel, mechanisms such as milestone-based cost reimbursement or production tax credits may be used to attract private anchor clients while avoiding full risk internalization.

In this context, hyperscale data centers can serve as catalytic "anchor loads," helping to launch commercial SMR programs by aggregating long-term, high-value demand. However, such participation is only viable if TEA frameworks explicitly reflect the heightened financial and operational exposure faced by first movers. Rather than viewing SMR-powered data centers as plug-and-play applications of generic nuclear economics, analysts and investors must recognize the transitional nature of FOAK deployments and design financial structures that allocate risk commensurately while preserving the upside potential of subsequent cost decline. A rigorous and differentiated treatment of FOAK versus NOAK conditions is thus indispensable to any serious techno-economic evaluation of SMR integration into digital infrastructure.

Recent literature suggests that the economics of SMRs are best interpreted through prospective cost and competitiveness benchmarks rather than realized operating profits. A recent European review reports an average SMR capital cost of about 7,031 €/kW and an average LCOE of about 85 €/MWh, with capital costs on average 41% higher than those of large reactors.[41]

CellPress
OPEN ACCESS

More detailed bottom-up assessments of advanced SMR designs report overnight capital costs in the range of 3,985–4,844 $/kW and LCOE values of 80.6–89.6 $/MWh across light-water, gas-cooled, and molten-salt concepts.[42,43] In renewable-intensive European system modeling, SMRs become economically competitive when construction costs remain below approximately 7,400 €/kW and operating costs below about 230 €/kW-year.[44] At the project level, additional value may arise under favorable market and policy conditions: for example, a recent coal-to-nuclear analysis for Texas found that brownfield asset reuse and tax credits could reduce modeled SMR LCOE to 43–52 $/MWh.[45] At the same time, recent workforce-based operating-cost analysis indicates that even the labor component of SMR operation remains non-negligible, with the net present cost of operating-worker salaries estimated at roughly $431 million for a representative 300 MWe iPWR case.[46] These estimates should nevertheless be interpreted with caution. Although a very small number of SMR-type projects have recently entered operation, the global operating base remains too limited and too recent to provide representative long-term data on average market revenues, operating margins, or gross profit. As a result, the literature still relies primarily on modeled indicators such as LCOE, operation and maintenance (O&M) assumptions, net present value (NPV), and system-level competitiveness thresholds rather than mature empirical operating benchmarks.[47,48]

Capturing multiproduct value streams

Positioning SMRs solely as electricity providers undervalues their systemic role in digital ecosystems. At the same time, it is important to distinguish between capabilities that are intrinsic to SMR technology and synergistic benefits that arise specifically from co-location with data centers. When co-located with data centers, SMRs function as multi-output energy hubs, simultaneously delivering firm electricity, low-grade thermal energy, resilience services, and in certain configurations valuable co-products such as hydrogen, desalinated water, or district heating.[49,50] These co-products are not unique to data-center applications per se; rather, their value in SMR-data center systems depends on whether co-location creates additional benefits through steady anchor demand, shared infrastructure, and tighter coupling among electricity, heat, cooling, and resilience services. This multi-vector capacity not only expands the functional utility of the reactor but also creates synergistic opportunities for cost reduction, emissions mitigation, and spatial efficiency within tightly integrated campus-style infrastructure, particularly when local demand conditions and system design support such integration.

The steady baseload profile of SMRs provides a unique advantage in anchoring energy-intensive digital loads, enabling a high degree of grid independence and reducing the need for upstream infrastructure reinforcement. Unlike renewable sources that suffer from intermittency, or gas-fired generators that incur emissions and fuel risk, SMRs offer predictable, round-the-clock output that aligns well with the continuous power needs of tier III and tier IV certified data centers. This stability facilitates smoother internal load scheduling, more efficient cooling-system operation, and smaller battery storage footprints for power quality management.

Thermal integration further amplifies the economic case. By using reactor heat to drive absorption chillers or heat-assisted cooling loops, data centers can displace large-scale electric chillers, significantly reducing both capital investment and operating costs associated with high-efficiency air conditioning systems.[51] In cold climates or shoulder seasons, this thermal energy may also support district heating networks or building heating, ventilation, and air conditioning (HVAC) systems within adjacent facilities, creating cascading benefits across the campus or industrial zone.

Accordingly, the rationale for each additional value stream must be assessed case by case. District heating is most compelling where there is proximate and sustained thermal demand and an appropriate local network or campus context. Desalination is more relevant in water-stressed or coastal regions where treated water has clear local value. Hydrogen is most attractive when off-peak reactor output can be redirected to electrolysis and where either long-duration resilience needs or external hydrogen demand justify the added conversion, storage, and reconversion infrastructure. In each case, the relevant techno-economic question is not only whether the SMR can technically provide the service, but whether temporal matching, spatial adjacency, and the cost of additional components support its inclusion in an integrated data-center system.

Importantly, the firm excess capacity available during off-peak computing hours, particularly at night, can be redirected to produce green hydrogen via high-temperature electrolysis or support participation in wholesale electricity markets and ancillary services.[49,52] Hydrogen production co-located with SMRs benefits from thermal integration (increasing electrolysis efficiency), capacity factor maximization, and infrastructure sharing. Within SMR-powered data center systems, hydrogen is more appropriately viewed as a strategic co-product and long-duration energy vector than as a default balancing resource. Its value arises most clearly when high-capacity-factor SMR operation produces surplus energy that can be redirected to electrolysis, particularly where thermal integration improves process efficiency and where hydrogen has downstream value in resilience provision or external end uses. In contrast, the operational needs of short-timescale balancing, ride-through support, and day-to-day load smoothing are generally better served by batteries and thermal storage. From this perspective, the significance of hydrogen lies not primarily in round-trip efficiency, but in its potential to broaden the economic and functional scope of SMR-data center integration within a multi-energy infrastructure context. In regions with market mechanisms for spinning reserve, frequency regulation, or ramping support, SMRs can be operated flexibly, especially when paired with on-site battery storage or thermal inertia buffers, to generate ancillary revenue streams that further improve asset utilization.

Consequently, the core TEA question evolves: not whether SMR-sourced electricity is cost-competitive with grid supply on a marginal basis, but whether the integrated system architecture reduces the total cost of computing, cooling, resilience, and carbon mitigation in aggregate. This reframing is essential for accurately assessing project feasibility and strategic fit. Marginal LCOE comparisons ignore co-product value, opportunity cost offsets, and cross-vector synergies—factors that increasingly dominate in high-performance digital energy hubs.

This system-level framing is also consistent with broader decarbonization strategies, which increasingly favor multi-vector energy platforms capable of delivering heat, power, and molecules from a single site.[53,54] In this paradigm, value is not additive but compound, reflecting the benefits of co-location, demand co-optimization, and shared infrastructure. For instance, a co-located SMR-H_2-data center complex may yield higher overall economic return and lower life cycle emissions than any of its components evaluated in isolation.

DESIGNING RELIABLE AND ENERGY-EFFICIENT SMR-INTEGRATED DATA CENTERS: ARCHITECTURE AND OPERATION

Grid integration versus autonomy: Defining the design of SMR-data centers

Most large-scale data centers today operate as grid-connected facilities, relying on the bulk power system for primary supply and/or redundancy.[55] The integration of SMRs, however, raises a fundamental architectural question: should these facilities evolve into self-sufficient micro-nuclear islands, operating independently of the grid, or should they remain grid-coupled nodes, leveraging both the SMR and the power network? Evaluating this choice through the lenses of system resilience, operational economics, and sustainability metrics underscores the trade-offs that will shape the future role of SMR-powered data centers.

A grid-connected SMR-data center configuration positions the facility as both a significant energy consumer and a dispatchable supplier of electrical or thermal energy within the broader energy ecosystem. Among potential architectures, the grid-coupled arrangement is expected to remain the predominant near-term pathway for SMR-data center integration.[13,56,57] In this configuration, the SMR units primarily meet the facility's energy demand while maintaining interconnection with the power grid. Grid integration provides important reliability benefits, including redundancy during SMR maintenance, scheduled refueling outages, or unscheduled reactor trips, as well as the ability to meet unexpected peak demand during normal operations. The grid interconnection also allows the SMR-data centers to participate in ancillary service markets,[57] generating additional revenue streams while enhancing grid stability, a particularly valuable role in renewable-heavy regions. However, the sustainability performance of a grid-connected facility is closely tied to the regional energy mix. If the surrounding grid is dominated by "brown" generation, the effective carbon intensity of a nuclear-powered data center may be diluted by imported electricity. Finally, grid-connected nuclear remains the established norm, but the regulatory and jurisdictional frameworks specific to SMRs are still evolving.[56]

In contrast, off-grid or islanded SMR-data center operation remains a less common design choice. From the standpoint of autonomy and resilience, such configurations may be attractive: they reduce exposure to cyber intrusions, the propagation of grid-wide disturbances, and grid price volatility, while simplifying on-site control and coordination.[58] Yet autonomy comes with significant technical demands. Because the SMR must serve as the sole energy source, the system often requires oversizing key components such as reactor modules, battery banks, hydrogen fuel cell systems, or thermal storage, to ensure continuity of supply.[59,60] These buffers are essential to ride through spikes in demand patterns with AI training workloads and to maintain operations during scheduled maintenance or unplanned outages. However, achieving the levels of redundancy expected in higher-tier facilities, such as $N + 1$ or $2N$ resilience standards,[13] is challenging under an island architecture.

A more promising pathway forward is the hybrid configuration, in which the SMR-data center operates in grid-connected mode under normal conditions but can seamlessly transition to islanded operation during grid disturbances.[58] This architecture blends the flexibility of grid participation with the resilience of autonomous operation, positioning SMR-powered data center facilities as adaptable nodes within the electrical grid. The rationale is especially compelling for AI-intensive and mission-critical workloads in defense, finance, and national security, where the propagation of grid disturbances can have severe consequences. In such contexts, SMR-powered data centers function as nuclear enclaves, capable of sustaining operations independent from the grid. Figure 4 provides an overview of the electrical configuration for an SMR-integrated data center, illustrating grid and on-site interconnections together with storage technologies. It is important to note that this figure does not represent a prescriptive system architecture, but rather illustrates a range of possible electrical configurations and optional components that may be incorporated depending on design objectives. In particular, hydrogen production and storage are included as one potential pathway to enable long-duration energy buffering and enhance system autonomy. Such subsystems can complement SMR operation by providing multi-timescale energy balancing, especially in islanded or resilience-focused configurations. However, their adoption depends on techno-economic considerations, including system cost, efficiency, and the degree of reliance on grid support.

Storage as a critical enabler in SMR-data center architectures

Storage plays a pivotal role in SMR-data center architectures by bridging the temporal and operational gaps between reactor output, highly variable IT load dynamics, and fluctuating grid conditions. The central question, however, is whether storage should remain a passive backup resource or evolve into active infrastructure that provides reliability, improves efficiency, and contributes to operational economics. Rather than functioning solely as insurance against rare contingencies, storage can be reconceptualized as a dynamic asset, supporting local resilience needs, participating in grid services, and arbitraging between SMR generation, grid conditions, and compute loads.[1,60] The guiding question for future design is not whether storage can play these roles, but how to balance and prioritize them to define the next generation of SMR-data center integration.

Batteries are central to the reliability and resilience of the SMR-data center infrastructure, providing fast-response capabilities for power quality support, peak shaving, and black-start functionality.[59,60] These services ensure continuity of critical operations even under volatile workloads or unexpected outages. In grid-connected or hybrid configurations, batteries can also be deployed to enhance grid stability through ancillary services,[60]

CellPress
OPEN ACCESS

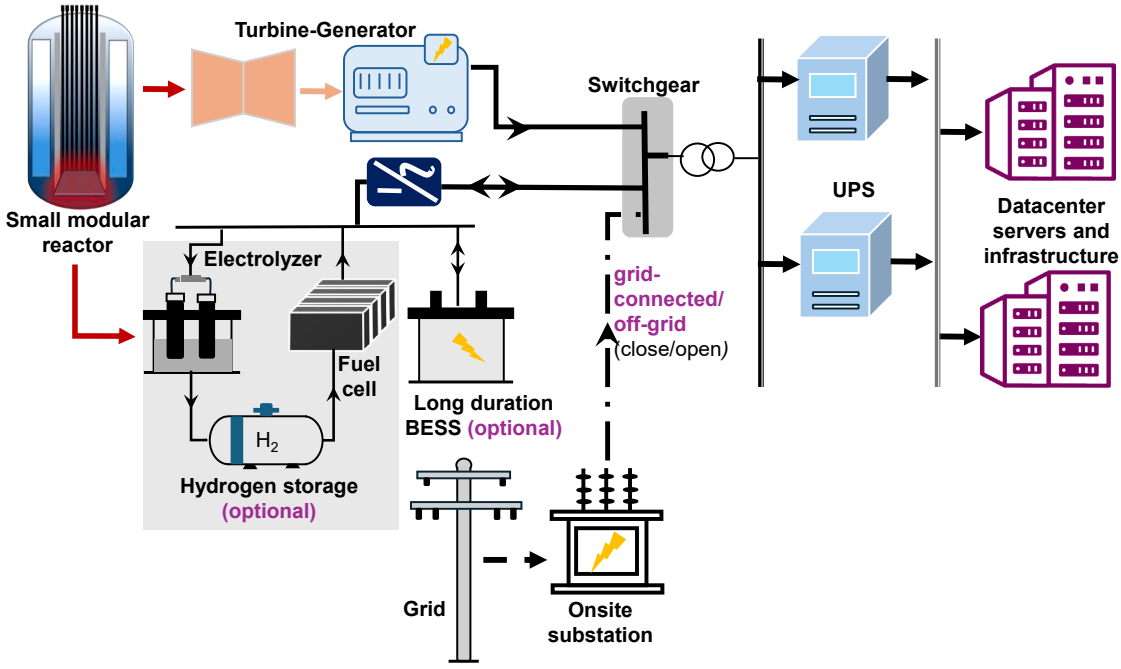


Figure 4. SMR-data center configuration with optional long-duration BESS, hydrogen storage system, and grid interconnection

while simultaneously generating new revenue streams for operators. In addition, batteries enable energy arbitrage, storing electricity from the grid during periods of low cost or low carbon intensity and discharging it during peak demand or carbon-intensive intervals.[61] However, their limited discharge duration confines their role to short-term response, underscoring the need to integrate longer-duration storage solutions such as hydrogen or thermal energy systems.

TES, by contrast, aligns more directly with the cooling demands of data centers, adding a unique dimension to SMR integration. TES enables shifting of thermal loads across time, reducing cycling stresses on the SMR and aligning cooling availability with demand variability. TES can be used to capture excess or waste heat during low-demand periods and then release it to drive absorption chillers or supply district energy networks.[62,63] At the cooling-system scale, lower-cost chilled-water storage can provide complementary load shifting,[64] buffering the chilled-water supply, and further reducing peak electrical demand. TES utilization is found to improve overall energy efficiency (PUE), reduce reliance on mechanical chillers, and strengthen resilience by ensuring continuity of cooling during peak IT workloads, grid disruptions, or reactor maintenance cycles.[65] TES can also promote flexibility of the data center by enabling arbitrage of thermal energy. The perspective challenge is whether TES should be designed primarily to stabilize on-site data center cooling, to valorize waste heat for broader energy services, or to provide flexibility, and how these roles can be effectively balanced in practice.

Hydrogen storage extends the resilience horizon of SMR-data center systems by providing long-duration energy buffering that batteries and TES alone cannot achieve.[66] Surplus reactor output can be directed to electrolysis, producing hydrogen that can be stored and later reconverted to electricity through fuel cells or turbines.[67] This capability enables data centers to ride through extended outages, refueling cycles, or sustained grid disruptions, thereby ensuring continuity of operations for critical workloads. Its deployment, however, faces significant challenges, including low round-trip efficiency, high infrastructure costs, and safety considerations,[67] which constrain its near-term feasibility for data center use. The value proposition of hydrogen storage goes beyond the arbitrage and resilience benefits offered by other technologies: it provides a versatile energy vector and a strategic bridge to cross-sector integration (transportation and chemical industries), positioning SMR-data centers as active players in the emerging hydrogen economy.

The power distribution architecture in a data center defines how electricity is routed from generation or grid entry points through UPS systems, switchgear, and distribution units to the IT load. In conventional data centers, these architectures are designed primarily for reliability, with redundant paths and UPS topologies (centralized or distributed)[13] ensuring tier III/IV uptime against grid failures or local faults. With the integration of SMRs and diverse storage systems, however, this paradigm requires reframing. The role of distribution expands from providing reliability and resilience functions to coordinating multiple supply pathways, SMR output, grid imports, and storage, while dynamically following the varying and unpredictable profiles of IT workloads. This shift moves distribution from a static, redundancy-driven hierarchy toward an adaptive infrastructure layer that balances reliability, efficiency, and flexibility. In practice, this implies the evolution of a multi-path intelligent backbone, enabled by advanced power-electronic controls capable of arbitrating in real time between nuclear, grid, and storage resources.

Cooling configurations in SMR-data centers: From overhead to synergistic systems

Thermal management in data centers has evolved from conventional air-based systems to advanced liquid and hybrid configurations. In legacy facilities, computer room air conditioners and computer room air handlers circulate chilled air, typically supported by chilled-water loops and cooling towers.[62,63] These systems remain dominant in existing data centers because of their proven reliability and technological maturity. Their efficiency, however, is fundamentally constrained by the low heat capacity of air and the narrow supply temperature tolerances required by IT hardware. The resulting high parasitic cooling load inflates the PUE and, in the context of SMR integration, offers no opportunity to leverage nuclear heat. In this configuration, cooling remains an overhead burden on the reactor's electricity output rather than a synergistic subsystem.

Unlike air, liquid offers much higher heat capacity and thermal conductivity, enabling significantly more efficient heat removal. Rear-door heat exchangers, which mount chilled-water panels at the back of server racks, and cold plates with embedded liquid channels mounted directly on CPUs and GPUs provide direct thermal contact with high-power chips.[63,65] Immersion cooling extends this principle further by submerging entire servers in dielectric fluids.[65,68] In single-phase designs, the liquid absorbs heat and is circulated through heat exchangers, while in two-phase systems, a low-boiling-point fluid vaporizes at the chip surface and condenses on a cooled coil and achieves highly efficient phase-change cooling. These approaches achieve higher heat flux removal, which is critical for AI/machine learning (ML) workloads and high-power GPUs, while elevating waste heat to more useful temperature levels (40°C–75°C). As a result, they reduce auxiliary fan energy, enable higher rack densities (more compute per unit space), and create opportunities for waste heat reuse.

In conventional AI/ML data centers, liquid cooling loops are typically supplied by chilled-water systems, with compressors and pumps powered by grid electricity.[65] The IT load rejects heat into the liquid loop, which is ultimately discharged to the atmosphere through electric compression chillers and cooling towers. In this arrangement, cooling functions entirely as an electric load. With on-site SMRs, however, cooling can be more deeply integrated. Reactor electricity can drive pumps, control systems, and auxiliary chillers, while low-grade steam or hot water from the SMR can drive absorption chillers,[64] which provide chilled water for the liquid loop. Single-effect absorption chillers, operating with 70°C–105°C heat, typically achieve a COP of approximately 0.7, while double-effect designs, requiring higher-grade heat of 140°C–180°C, can reach COP values of 1.3–1.4. Since SMR steam temperatures span 275°C–850°C, the system can flexibly draw bypass heat from different steam stages, but must employ strategies to "step down" the temperature to absorption chiller-compatible levels. Therefore, the heat rejected by IT servers is absorbed into a refrigeration cycle powered by nuclear heat, reducing parasitic electrical demand, improving PUE, and decreasing reliance on grid electricity. Higher-grade nuclear steam can be coupled with TES or directed through organic Rankine cycles (ORCs)[69] to produce secondary electricity while simultaneously supporting cooling demands through absorption chillers. Recent studies on nuclear district-heating and SMR-ORC systems further support the feasibility of using ORC cogeneration to recover surplus nuclear heat and improve heat-source utilization, particularly when coordinated with storage or load-shifting strategies.[70,71] In such configurations, the cooling stream emerges as a co-equal pathway alongside electricity generation, opening new avenues for efficiency, flexibility, and resilience. On the other hand, the IT waste heat produced in immersion cooling can also be valorized directly by preheating absorption chillers, charging TES, supplying district-heating networks, or, under favorable temperature and market conditions, supporting ORC-based recovery.[72] In parallel, the steady low-grade steam or hot water from SMRs can complement IT waste heat, creating a dual-stream configuration in which nuclear and data center thermal outputs are valorized together. This synergy transforms the cooling plant from a passive overhead into a cogeneration hub, supporting reliable computing while simultaneously supplying useful thermal energy. Heat-pump coupling and TES can help manage temperature mismatches across these thermal streams.

Free cooling leverages favorable ambient conditions to dissipate heat, thereby reducing electricity consumption for cooling. Airside and waterside economization use outside air, cooling towers, or natural water bodies such as rivers, lakes, or seawater as heat sinks, either directly or indirectly through heat exchangers. Radiative cooling,[73] by contrast, exploits the sky's thermal radiation, particularly at night, to reject heat without relying on compression chillers. High-temperature-tolerant data centers extend the number of hours during which free cooling is viable,[74] enabling greater reductions in cooling demand. By lowering parasitic load, free cooling improves overall efficiency, reduces PUE, and enhances sustainability. In SMR-data centers, free cooling can address a portion of the cooling demand during favorable conditions such as cool nights, cold winters, or dry climates, while the SMR provides reliable base cooling through electricity or combined electricity and heat. Together, this hybrid approach can be used to balance resilience with efficiency, positioning free cooling as a complementary layer to SMR integration.

Figure 5 summarizes the cooling configurations for SMR-integrated data centers. Figure 5A depicts cooling treated solely as an electric load, with electric chillers powered by the SMR, the grid, or storage. Figure 5B illustrates different pathways for direct SMR integration with data center cooling, including TES, absorption chillers, and stratified chilled water tanks, with optional free cooling during favorable weather conditions. Figure 5C highlights opportunities for valorization of IT waste heat through rear-door heat exchangers and immersion cooling, which can be coupled with TES, preheaters for absorption chillers, or redirected to district and building heating networks. Cooling must therefore be redefined in SMR-data centers as an integrated energy service, aligned with electricity and heat flows, rather than a passive overhead. Appropriate thermal energy management configurations should be identified based on TEA and the targeted energy vectors.

SMR-DATA CENTER HUBS: DYNAMIC INTERACTION WITH THE GRID AND AI COORDINATION

Redefining demand response through SMR integration and workload scheduling

In the current era, where AI training and inference workloads have gained unprecedented prominence, hyperscale data

CellPress
OPEN ACCESS

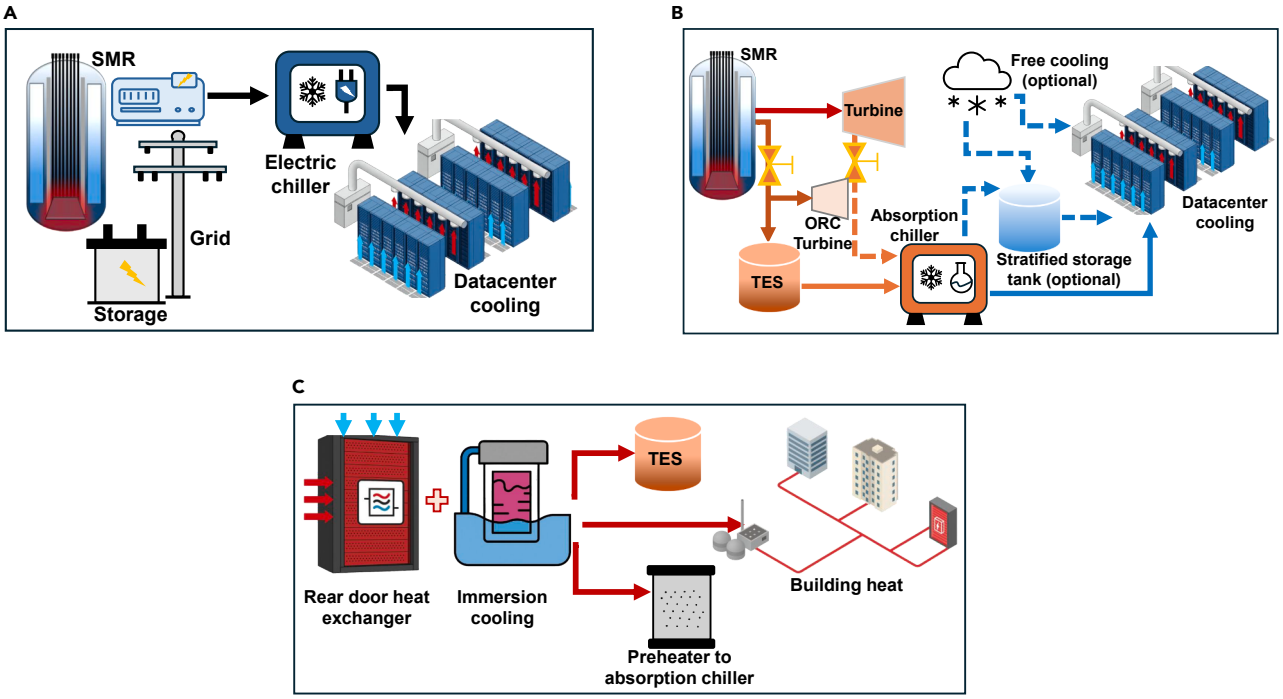


Figure 5. Cooling configurations in SMR-powered data centers

(A) Data center cooling powered as an electric load, with electric chillers supplied by SMR, grid, or storage.

(B) Data center cooling options integrated with SMR using absorption chillers, TES, and stratified chilled water storage.

(C) Valorization pathways for IT waste heat.

centers increasingly host heterogeneous computational tasks with widely varying requirements. Real-time inference and financial transactions, for example, demand immediate execution with strict latency constraints, whereas large-scale model training, data analytics, and backup processes are delay-tolerant and can be scheduled flexibly. This opens up opportunities for energy-aware task scheduling, where flexible workloads can be strategically time-shifted[61,75] or spatially allocated[76,77] to align with objectives such as minimizing operating costs, reducing carbon intensity, or enabling grid and storage arbitrage. From the perspective of the grid, this flexibility transforms the data center from a passive consumer into an active participant in demand response. In traditional grid-connected facilities, delay-tolerant tasks can be shifted to periods of low demand, high renewable penetration,[75] or low marginal prices, thereby reducing operating costs[60] and lowering the facility's carbon footprint. Conversely, workloads can be strategically increased during periods of electricity surplus, enabling arbitrage by absorbing excess generation and supporting overall grid stability.[78]

In SMR-coupled data centers, however, demand response acquires a fundamentally different meaning from that in grid-only facilities. Flexible workloads can be aligned with periods of excess SMR output,[79] such as at night, when cooling demand is reduced and more reactor capacity can be directed toward IT loads, or during intervals when storage reservoirs are fully charged and low cost or low-carbon grid imports are available. During periods of high immediate IT demand, workload scheduling must account for SMR ramping constraints, minimum/maximum power limits, and dynamics of coupled thermal subsystems, with storage acting as critical buffers to smooth mismatches. The overarching operational goal shifts from pure demand response to maximizing SMR utilization, preserving a high nuclear capacity factor, and reducing curtailment of the reactor output.

From this perspective, demand response in SMR-powered hubs is reframed: instead of responding passively to grid signals, the facility functions as a bidirectional energy node. Importantly, this capability should be understood as coordinated system-level flexibility rather than unrestricted high-frequency reactor load following. Reported SMR load-following rates include approximately 5% per minute for secondary frequency control and up to 20% per minute for scheduled power cycling or specific flow-rate-control strategies.[80] However, the allowable ramping envelope is highly plant-specific and depends on reactor type, balance-of-plant configuration, control strategy, fuel-cycle condition, and safety constraints.[81] Therefore, fast data-center load fluctuations would need to be buffered primarily by UPS systems, battery energy storage, thermal storage, workload scheduling, and grid exchange in grid-connected configurations.[80] The SMR would instead follow slower net-load variations within its allowable operating envelope, while the facility adjusts workloads internally, optimizes storage dispatch, and exports surplus electricity to the grid during periods of low IT demand or storage saturation. Quantifying these operational limits for low-, normal-, and high-demand data-center days would

require high-resolution data-center load traces coupled with plant-specific SMR dynamic models, and therefore remains an important direction for future work. Thus, while the traditional demand-response capability may be less pronounced in an SMR-integrated setting, the broader value proposition is enhanced: the system evolves from load flexibility alone into a co-optimization paradigm that simultaneously coordinates nuclear generation, multi-vector storage, and heterogeneous IT workloads.

Modeling and control of SMR-data center systems

Traditional grid-connected data centers draw electricity exclusively from the grid to power both IT loads and cooling systems. In contrast, integrating nuclear reactors, advanced storage and load-shifting technologies, and hyperscale data centers transforms the facility into a multi-energy system, where electricity, heat, and cooling function as interdependent vectors. Recent SMR-data center modeling studies support this framing by showing that reactor operation, data-center power demand, and hybrid cooling can be modeled as coupled subsystems rather than independent components.[82] Nuclear hybrid-energy studies further show that nuclear heat and electricity can be coordinated with data-center waste heat recovery, district heating, ORC cogeneration, hydrogen production, and thermal load shifting.[70,83] This coupling redefines the facility as a complex cyber-physical system, requiring real-time coordination across nuclear generation, IT workloads, storage or load-shifting assets, and thermal subsystems. Control can no longer be limited to conventional electrical load following; instead, it must simultaneously improve efficiency, enable flexibility through load shifting and storage arbitrage, and uphold the stringent operational standards of both nuclear systems and mission-critical computing.

At the operational scale, beyond configuration design and component interconnection, the first step is typically to solve an optimal scheduling problem.[60,79] Here, single- or multi-objective optimization[84] is applied to steady-state equations that represent system behavior and constraints, with objectives such as minimizing operating costs, maximizing energy efficiency, reducing carbon intensity, and enabling trade-offs between nuclear output utilization, storage dispatch, and cooling performance. This framework can be extended to stochastic optimization[85] to account for uncertainties in IT workloads, grid price volatility, and ambient conditions. Model predictive control (MPC) adds a dynamic layer, anticipating workload and environmental fluctuations and adjusting storage or cooling setpoints proactively.[86] Reinforcement learning (RL) and multi-agent RL (MARL) can provide adaptability by learning control policies that coordinate SMR output, battery use, hydrogen electrolysis, and thermal loop management in real time.[61,87] A key limitation of these scheduling approaches, however, is that they often cannot fully capture transient dynamics. Short-term fluctuations or fast disturbances may exceed safety margins, particularly in nuclear-integrated systems. To ensure safety and compliance under such conditions, dynamic simulations are required as a complementary step, providing detailed insight into transient responses and validating that optimized schedules remain within safe operating envelopes.

Developing physics-based models based on differential-algebraic equations (DAEs) for SMR-data center facilities is essential for capturing the operational dynamics of the system, validating control and coordination strategies, and assessing the feasibility of proposed demand-response schemes. This approach is not new in the SMR domain, where integrated energy-system modeling has been performed using tools such as Modelica.[11,88] Looking ahead, SMR-data center configurations and benchmark models will provide the research community with standardized testbeds for advancing control and optimization frameworks, much like those developed for other emerging SMR systems.

With the emergence of digital twins, there is a growing consensus that data-driven methods, alongside physics-based models, are indispensable for real-time management and control of data centers.[89] Forecasting models trained on IT workload traces, cooling demand profiles, and electricity usage can predict AI training surges, electrical peaks, or thermal load variations.[90,91] RL models trained on physics-based models can be further used to provide real optimal control. In practice, digital twins serve as operational mirrors of the facility, enabling predictive simulation and real-time decision support, thereby enhancing reliability, efficiency, and flexibility.

From load driver to system orchestrator: AI in SMR-data centers

The rapid expansion of data centers is largely driven by the adoption of AI, which is therefore often regarded as one of the sector's primary energy consumers.[1,55] Yet a more forward-looking perspective is to regard AI not merely as a consumer, but as an orchestrator of cyber-physical systems, actively shaping the real-time interplay between energy and computation. For example, RL and multi-agent control frameworks enable AI to co-optimize the dispatch of SMR output, batteries, hydrogen fuel cells, and thermal subsystems as discussed previously. At the workload layer, AI can allocate tasks by assigning inference and mission-critical applications to dedicated resources, while flexibly scheduling training, analytics, etc., to align with reactor availability, storage states, or low-price grid windows.[61] At the grid interface, AI-enhanced digital twins can forecast grid conditions and estimate the type of ancillary services, such as frequency regulation, spinning reserves, or fast ramping support. Therefore, AI provides the much-needed intelligence and real-time control capabilities required in this complex symbiosis.

Viewing AI solely as a burden on the grid or on SMR capacity frames its rapid growth only as a challenge to sustainability and reliability.[92] Yet reframing AI as both a consumer and a contributor, an intelligence layer that actively supports control, coordination, and optimization, may open a complementary line of thinking. From this perspective, AI's role in enabling resilience and flexibility within the nuclear-digital ecosystem deserves further investigation as both a technical opportunity and a research frontier.

AI-based coordination and control in SMR-data center systems can be implemented with relatively modest computational overhead compared with overall facility demand. Such coordination tasks, including scheduling, forecasting, and real-time control, are typically executed on dedicated control infrastructure rather than on primary compute resources. Recent studies on RL-based data center control report computational requirements on the order of sub-kW GPU power and single-digit GB memory usage,

CellPress
OPEN ACCESS

which are small relative to MW-scale data center loads.[93] Despite this limited computational footprint, AI-enabled control can deliver significant system-level benefits by optimizing interactions across tightly coupled subsystems, including power supply, cooling, and storage. Prior work demonstrates that learning-based control strategies can improve thermal compliance and reduce cooling energy consumption through coordinated, real-time decision-making.[93] The effectiveness of such approaches depends on the operational architecture. In grid-connected configurations, AI coordination can improve economic dispatch and efficiency, whereas in islanded or autonomy-focused systems, it plays a more critical role in maintaining real-time balance and ensuring reliable operation under varying conditions.

REGULATION AND GOVERNANCE

Integrating SMRs into AI data centers presents substantial regulatory challenges, particularly in the US, where the federal permitting and construction approval process remains a significant obstacle. These regulatory requirements should be understood not as external considerations but as boundary conditions that shape system architecture, operational flexibility, and economic viability. Licensing, siting, emergency-planning, and grid-interconnection constraints determine whether a proposed SMR-data center configuration can be deployed as a grid-connected, hybrid, behind-the-meter, or islanded system, while also influencing project costs, development timelines, financing risk, and overall feasibility. This section therefore examines the regulatory barriers and governance pathways that may either constrain or enable SMR-data center deployment.

Three institutional actors are especially important for reducing regulatory hurdles for nuclear deployments serving data centers: the US Nuclear Regulatory Commission (NRC), the Federal Energy Regulatory Commission (FERC), and the North American Electric Reliability Corporation (NERC). The first two are federal regulatory agencies responsible for nuclear plant deployment and operation standards and electric-grid regulation, respectively, whereas NERC is a not-for-profit, non-governmental regulatory authority responsible for ensuring a reliable and secure bulk power system.

Current NRC regulatory frameworks under 10 Code of Federal Regulations (CFR) parts 50 and 52 govern the licensing of existing nuclear plants and new nuclear construction, respectively. NRC regulations primarily focus on traditional electricity generation and export to the electric grid. There is a notable absence of regulatory clarity for the use of nuclear reactors in non-grid applications such as data centers. These frameworks were developed for large, site-specific reactors and rely on siting assumptions under 10 CFR part 100 that favor low population density areas, limiting flexibility for modular, co-located, and distributed SMR deployment. To address this, the NRC has introduced 10 CFR part 53, which enables flexible licensing through probabilistic risk assessment.[94] It includes alternative siting criteria (10 CFR 53.530) permitting deployment in higher density or industrial areas, including near populations exceeding 25,000, along with provisions for remote operations, reduced staffing, and load following (10 CFR 53.740). While part 53 introduces greater flexibility, its application to co-located and behind-the-meter SMR-data center systems remains relatively untested and will require further implementation experience.

The licensing process is rigorous and can be time-consuming. Applicants must provide detailed safety analyses, including evaluations of potential accidents and operational transients. These requirements can significantly delay project timelines, which is a critical concern for data centers that need rapid deployment solutions. The NRC places significant emphasis on safety. Safety assessment in the context of co-located nuclear facilities and data centers will require extensive hazard analysis, including evaluation of risks related to radiation exposure, heat-transfer systems, and emergency preparedness protocols, as well as identification of new risks that could arise from co-location, all of which add layers of complexity to the licensing review.[95]

The NRC, under CFR part 50, also mandates that only licensed operators may manipulate the controls of any nuclear facility. Any apparatus or mechanisms affecting reactor reactivity or power levels must be operated with the knowledge and consent of a licensed operator or senior operator. This requirement introduces a new challenge for nuclear operators who must now also address the operational needs of data centers.

FERC and NERC are tasked with maintaining grid operating standards and depend on both existing and upcoming nuclear plants to fulfill adequacy and reliability requirements. Many regions overseen by NERC are currently experiencing high or elevated resource adequacy risks,[96] a situation that is likely to worsen with the anticipated increase in electricity demand driven by data centers and industrial electrification. Under these conditions, data centers may encounter significant resistance from FERC and NERC, who are inclined to prioritize grid services associated with new generation over the needs of data centers operating behind the meter. A notable example is FERC's rejection of a co-location amendment filed by Pennsylvania-New Jersey-Maryland Interconnection (PJM), which sought to expand the co-located load for data centers at the Susquehanna nuclear power plant. FERC cited the potential negative implications of such changes for system reliability, resource adequacy, and economic impacts on customers and markets.[97]

Although the International Atomic Energy Agency (IAEA) does not directly impose regulatory hurdles for the civilian use of nuclear energy for power production, some advanced reactor concepts that can use high-assay low-enriched uranium (HALEU) fuel may face additional complexities in meeting international non-proliferation and safeguards standards. HALEU fuel can contain up to 20% uranium-235, blurring the line between low-enriched uranium (LEU) used in conventional nuclear plant concepts and weapons-grade highly enriched uranium (HEU). The NRC is responsible for maintaining security protocols for nuclear fuels and regulates the use of large quantities of HALEU consistent with IAEA safeguards standards.

Regulatory hurdles may also arise under the National Environmental Policy Act (NEPA), under which nuclear reactor projects must provide detailed environmental impact statements for new nuclear facilities co-located with data centers or for changes in the intended use of existing nuclear facilities. NEPA also introduces public-involvement and environmental-justice considerations that could significantly delay approval and deployment of nuclear data center applications.

Regulatory stakeholders are making significant efforts to reduce barriers that impose substantial cost burdens and delays on nuclear developers. The Accelerating Deployment of Versatile, Advanced Nuclear for Clean Energy (ADVANCE) Act of 2024, which was passed with bipartisan support and signed into law in July 2024, mandates the NRC to streamline the flexible deployment of nuclear plants while maintaining the NRC's core safety and security missions.[98] The NRC is also developing new regulatory frameworks, deployment strategies, and guidance for advanced nuclear technologies, including microreactors and non-conventional applications such as data centers.

CONCLUSIONS

The accelerating rise of AI, cloud computing, and HPC is driving data centers to the forefront of energy sustainability challenges. This perspective paper presents the integration of SMRs with hyperscale data centers, reframing these facilities as multi-energy cyber-physical systems. The novelty of this perspective lies in reframing SMR-powered data centers not simply as an alternative electricity supply option, but as an integrated multi-energy, cyber-physical infrastructure problem that must be evaluated through reliability value, service-oriented metrics beyond LCOE, lifecycle cost evolution from FOAK to NOAK deployment, and multiproduct energy-service synergies. By bringing these dimensions into a single perspective, this work provides a more system-relevant foundation for future techno-economic assessment, design, control, and policy analysis of SMR-data center systems. The benefits of SMR-powered data centers are articulated through broader system-level perspectives: the ability to co-optimize electricity, heat, and cooling; to provide new pathways for waste heat valorization; to enhance resilience through on-site generation and storage integration; and to increase independence from the local grid while enabling active participation as a prosumer. Looking ahead, SMR-data center ecosystems have the potential to evolve into multi-energy hubs that couple nuclear generation with cooling and storage through advanced control, reshaping the trajectory of digital infrastructure while accelerating the sustainable deployment of SMR technology. Future work should translate this framework into site-specific quantitative studies that compare grid-connected, hybrid, and islanded configurations under different regional market, cooling, and licensing conditions.

AUTHOR CONTRIBUTIONS

R.A.J. and H.L. contributed equally to this work and both contributed to manuscript revision. H.L. drafted the techno-economic analysis section and related parts of the manuscript, including the introduction and conclusions. R.A.J. and A.M.R. contributed to the grid-integration content, with R.A.J. leading this component. A.M.R. and Y.L. contributed to the thermal-integration and thermal-management content. B.P., M.S., and T.R.M. contributed to the policy and regulatory analysis. J.Z. contributed to manuscript revision, supervised the work, and handled corresponding author responsibilities. All authors reviewed and approved the final manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

REFERENCES

1. Industry Technical Support Leadership Committee (ITSLC) Task Force on Data Center Growth and Grid Readiness (2025). Data Center Growth and Grid Readiness. Tech. Rep. PES-TR 131 (IEEE Power & Energy Society).
2. Vanatta, M., Patel, D., Allen, T., Cooper, D., and Craig, M.T. (2023). Technoeconomic analysis of small modular reactors decarbonizing industrial process heat. Joule *7*, 713–737. https://doi.org/10.1016/j.joule.2023.03.009.
3. Lu, J. (2024). Small modular reactors: A viable path to sustainable data centers in the age of artificial intelligence. Front. Soc. Sci. Technol. *6*, 58–62. https://doi.org/10.25236/FSST.2024.061110.
4. Shehabi, A., Hubbard, A., Newkirk, A., Lei, N., Siddik, M.A.B., Holecek, B., Koomey, J., Masanet, E., and Sartor, D. (2024). United States Data Center Energy Usage Report. Tech. Rep. LBNL-2001637 (Lawrence Berkeley National Laboratory). https://eta.lbl.gov/publications/2024-lbnl-data-center-energy-usage-report.
5. U.S. Department of Energy. (2024), Doe releases new report evaluating increase in electricity demand from data centers. https://www.energy.gov/articles/doe-releases-new-report-evaluating-increase-electricity-demand-data-centers.
6. CBRE Research. (2025). North america data center trends h1 2025, Tech. rep.. https://www.cbre.com/insights/reports/north-america-data-center-trends-h1-2025.
7. Miller, R. (2025). From cloud to concrete: How explosive data center demand is redefining commercial real estate. https://www.datacenterfrontier.com/machine-learning/article/55310824/from-cloud-to-concrete-how-explosive-data-center-demand-is-redefining-commercial-real-estate
8. Cushman & Wakefield, Data Center Advisory Group. (2025). 2025 global data center market comparison, Tech. rep. (Cushman & Wakefield). https://cushwake.cld.bz/globaldatacentermarketcomparison-05-2025-global-central-en-content.
9. U.S. Department of Energy (2024). Recommondation on Powering AI and data center infrastructure. https://www.energy.gov/sites/default/files/2024-08/Powering AI and Data Center Infrastructure Recommendations July 2024.pdf.
10. Buzzetti, R., Lo Frano, R.L., and Cancemi, A.S. (2024). Feasibility study of desalination plant powered by SMR. Nucl. Eng. Des. *418*, 112897. https://doi.org/10.1016/j.nucengdes.2023.112897.
11. Poudel, B., Gautam, M., Li, B., Huang, J., and Zhang, J. (2023). Design, modeling and simulation of nuclear-powered integrated energy systems with cascaded heating applications. J. Renew. Sustain. Energy *15*, 054103. https://doi.org/10.1063/5.0163557.
12. Bejarano, J., and Mandal, P. (2025). Dynamic modeling and fault response analysis of SMR-integrated power grid. In IEEE Green Technologies Conference (GreenTech), pp. 66–70. https://doi.org/10.1109/GreenTech62170.2025.10977618.
13. Huang, H., Lin, W., Lin, J., and Li, K. (2025). Power management optimization for data centers: A power supply perspective. IEEE Trans. Sustain. Comput. *10*, 784–803. https://doi.org/10.1109/TSUSC.2025.3542779.
14. Rostirolla, G., Grange, L., Minh-Thuyen, T., Stolf, P., Pierson, J.M., Da Costa, G., Baudic, G., Haddad, M., Kassab, A., Nicod, J.M., et al. (2022). A survey of challenges and solutions for the integration of renewable energy in datacenters. Renew. Sustain. Energy Rev. *155*, 111787. https://doi.org/10.1016/j.rser.2021.111787.
15. ISL. (2024). Nuclear-powered data centers, defines Tier IV and Tier V requirements including full system redundancy. https://www.islinc.com/wp-content/uploads/2024/10/ISL-White-Paper-Nuclear-Powered-Data-Centers.pdf.
16. Warden, J., and Koenig, R. (2025). Data centers: a new dawn for nuclear energy?, discusses capacity factor and redundancy via multiple SMR units. https://kernd.de/wp-contet/uploads/2025/03/atw-2025-2_Data-Centers.pdf.
17. Bhowmik, P.K. (2024). Nexus-dc: Nuclear energy expedition for us data centers. Tech. rep. (Idaho National Laboratory [INL]).

18. Ponemon, L. (2016), Cost of data center outages. Data Center Performance Benchmark Serie. https://planetaklimata.com.ua/instr/Liebert_Hiross/Cost_of_Data_Center_Outages_2016_Eng.pdf.

19. Goldman Sachs Global Investment Research. (2024). The push for the 'green' data center and investment implications. Equity research report, Goldman Sachs, gS SUSTAIN: AI/data centers' global power surge. https://www.goldmansachs.com/pdfs/insights/goldman-sachs-research/the-push-for-the-green-data-center/aidatacenters.pdf.

20. ITIC. (2024). Information Technology Intelligence Consulting (ITIC), 2024 hourly cost of downtime report. Survey report. https://itic-corp.com/itic-2024-hourly-cost-of-downtime-report/.

21. Li, Z., Liang, M., O'Brien, L., and Zhang, H. (2013). The cloud's cloudy moment: A systematic survey of public cloud service outages. Int. J. Cloud Comput. Serv. Sci. (IJ-CLOSER) *2*, 1–15. https://doi.org/10.11591/closer.v2i5.5125.

22. HostDime. (2025). What is a tier 4 data center? a technical deep dive. https://www.hostdime.com/blog/tier-4-data-center/.

23. Liu, R., Sun, W., and Hu, W. (2023). Placement of high availability geo-distributed data centers in emerging economies. IEEE Trans. Cloud Comput. *11*, 3274–3288. https://doi.org/10.1109/TCC.2023.3280983.

24. COGEN World Coalition. (2024). Coalition, Meeting the energy needs of data centres: What role for cogeneration?, Tech. rep. https://www.cogenworld.org/wp-content/uploads/2024/11/CWC_Report_on_Data_Centres_and_Cogeneration_2024.pdf.

25. U.S. Environmental Protection Agency. (2006). The role of distributed generation and combined heat and power (chp) systems in data centers. https://www.epa.gov/sites/default/files/2015-07/documents/the_role_of_distributed_generation_and_combined_heat_and_power_chp_systems_in_data_centers.pdf.

26. Xendee. (2025). How absorption chillers boost microgrid performance. https://xendee.com/insights/how-absorption-chillers-boost-microgrid-performance.

27. Vanatta, M., Stewart, W.R., and Craig, M.T. (2025). The role of policy and module manufacturing learning in industrial decarbonization by small modular reactors. Nat. Energy *10*, 77–89. https://doi.org/10.1038/s41560-024-01665-w.

28. Locatelli, G., Bingham, C., and Mancini, M. (2014). Small modular reactors: A comprehensive overview of their economics and strategic aspects. Prog. Nucl. Energy *73*, 75–85. https://doi.org/10.1016/j.pnucene.2014.01.010.

29. William, R.M., Madia, J., and Vine, G. (2015). Small modular reactors: A call to action for nuclear cost competitiveness. Tech. rep. (Hoover Institution Press) https://www.hoover.org/sites/default/files/research/docs/small-modular-reactors-a-call-for-action-by-william-j-madia-gary-vine-and-regis-matzie.pdf

30. NuScale Power. (2023). UAMPS, Carbon free power project (cfpp) updates. https://www.nuscalepower.com/.

31. World Nuclear News. (2023). Nuscale raises carbon free power project cost estimate to $9.3 billion. https://www.world-nuclear-news.org.

32. Reuters. (2023). Nuscale cancels first us small modular reactor project as costs rise. https://www.reuters.com/business/energy/nuscale-power-uamps-agree-terminate-nuclear-project-2023-11-08/.

33. O. P. Generation. (2023). Darlington new nuclear project – bwrx-300 small modular reactor. https://www.opg.com/projects/darlington-new-nuclear-project/.

34. Clean Prosperity. (2022). Prosperity, Smrs in canada: Economic opportunities and challenges, Tech. rep. https://cleanprosperity.ca.

35. Rosner, R., and Goldberg, S.M. (2011). Small modular reactors – key to future nuclear power generation in the u.s. Energy Policy Institute at Chicago (EPIC) Report. https://www.energy.gov/ne/articles/small-modular-reactors-key-future-nuclear-power-generation-us.

36. DOE Office of Nuclear Energy. (2020). U. D. of Energy, Advanced reactor demonstration program (ardp), Tech. rep. https://www.energy.gov/ne/advanced-reactor-demonstration-program.

37. OECD (2011). NEA, Current status, technical feasibility and economics of small nuclear reactors. Tech. rep. (OECD Nuclear Energy Agency). https://www.oecd-nea.org/ndd/reports/2011/current-status-small-reactors.pdf

38. OECD (2021). NEA, Unlocking reductions in the construction costs of nuclear: A practical guide for stakeholders. Tech. rep. (OECD Nuclear Energy Agency). https://www.oecd-nea.org/upload/docs/application/pdf/2020-07/7530-reducing-cost-nuclear-construction.pdf

39. IAEA (2023). Advanced reactors information system (aris). Tech. rep. (International Atomic Energy Agency). https://www.iaea.org/resources/databases/advanced-reactors-information-system-aris

40. IEA. (2023). Electricity market report – special focus on nuclear, Tech. rep. https://www.iea.org/reports/electricity-market-report-2023.

41. Van Hee, N., Peremans, H., and Nimmegeers, P. (2024). Economic potential and barriers of small modular reactors in europe. Renew. Sustain. Energy Rev. *203*, 114743. https://doi.org/10.1016/j.rser.2024.114743.

42. Asuega, A., Limb, B.J., and Quinn, J.C. (2023). Techno-economic analysis of advanced small modular nuclear reactors. Appl. Energy *334*, 120669. https://doi.org/10.1016/j.apenergy.2023.120669.

43. Black, G.A., Aydogan, F., and Koerner, C.L. (2019). Economic viability of light water small modular nuclear reactors: General methodology and vendor data. Renew. Sustain. Energy Rev. *103*, 248–258. https://doi.org/10.1016/j.rser.2018.12.041.

44. Qu, C., Bang, R.N., and Sandal, L.K. (2025). Construction and operational cost requirements for competitive small modular reactors in europe. Appl. Energy *401*, 126621. https://doi.org/10.1016/j.apenergy.2025.126621.

45. Wieser, E., Clarno, K., Haas, D., and Webber, M.E. (2025). The economics of small modular reactors at coal sites: A program-level analysis within the state of texas. Energy Policy *202*, 114572. https://doi.org/10.1016/j.enpol.2025.114572.

46. Egieya, J.M., Amidu, M.A., and Hachaichi, M. (2023). Small modular reactors: An assessment of workforce requirements and operating costs. Prog. Nucl. Energy *159*, 104632. https://doi.org/10.1016/j.pnucene.2023.104632.

47. Mignacca, B., and Locatelli, G. (2020). Economics and finance of small modular reactors: A systematic review and research agenda. Renew. Sustain. Energy Rev. *118*, 109519. https://doi.org/10.1016/j.rser.2019.109519.

48. Larsen, L.M., Biegel, K., Guaita, N., and Foss, A. (2024). Effects of the u.s. inflation reduction act on smr economics. Front. Nucl. Eng. *3*, 1379414. https://doi.org/10.3389/fnuen.2024.1379414.

49. Elkhalafy, M.S., Abd-Allah, M.A., Elgabalawy, M., and Said, A. (2024). Techno-economic assessment of small modular reactors and renewable energy for hydrogen production: An egyptian grid case study. In 25th International Middle East. Power System Conference (MEPCON) (IEEE), pp. 1–8. https://doi.org/10.1109/MEPCON63025.2024.10850191.

50. Rath, M., and Morgan, M.G. (2020). Assessment of a hybrid system that uses small modular reactors (smrs) to back up intermittent renewables and desalinate water. Prog. Nucl. Energy *122*, 103269. https://doi.org/10.1016/j.pnucene.2020.103269.

51. Modro, S., Fisher, J., Weaver, K., Reyes, J., Groome, J., Babka, P., and Carlson, T. (2003). Multi-application small light water reactor final report. Tech. rep. (Idaho National Lab. [INL]).

52. Li, J., Dong, Z., Li, B., and Huang, X. (2022). Dynamic modeling of nuclear hydrogen generation plant based on multi-modular high-temperature gas-cooled reactor. International Conference on Nuclear Engineering, *86502* (American Society of Mechanical Engineers), V015T16A006. https://doi.org/10.1115/ICONE29-88899.

53. Ghimire, L., and Waller, E. (2023). Small modular reactors: opportunities and challenges as emerging nuclear technologies for power production. J. Nucl. Eng. Radiat. Sci. *9*, 044501. https://doi.org/10.1115/1.4062644.

54. Boldon, L., Sabharwall, P., Rabiti, C., Bragg-Sitton, S.M., and Liu, L. (2016). Thermodynamic exergy analysis for small modular reactor in

nuclear hybrid energy system. EPJ Nuclear Sci. Technol. *2*, 23. https://doi.org/10.1051/epjn/2016011.

55. Davenport, C., Lee, B., Corbett, B., Ritchie, J., Jaiya, J., Venugopal, V., Cash, N., Halferty, O., Singer, B., Mackay, J., et al. (2024). Data centers and the coming US power demand surge, Tech. rep., Goldman Sachs Global Investment Research, Equity Research Report. https://www.goldmansachs.com/pdfs/insights/pages/generational-growth-ai-data-centers-and-the-coming-us-power-surge/report.pdf.
56. Hemphill, T.A. (2024). US AI data centers and deployment challenges for small modular reactors: proposed regulatory policy recommendations. Sci. Public Policy *51*, 999–1003. https://doi.org/10.1093/scipol/scae040.
57. Su, C., Wang, L., Sui, Q., and Wu, H. (2025). Optimal scheduling of a cascade hydro-thermal-wind power system integrating data centers and considering the spatiotemporal asynchronous transfer of energy resources. Appl. Energy *377*, 124360. https://doi.org/10.1016/j.apenergy.2024.124360.
58. Ishaq, S., Khan, I., Rahman, S., Hussain, T., Iqbal, A., and Elavarasan, R.M. (2022). A review on recent developments in control and optimization of micro grids. Energy Rep. *8*, 4085–4103. https://doi.org/10.1016/j.egyr.2022.01.080.
59. Thompson, C.C., Konstantinos Oikonomou, P.E., Etemadi, A.H., and Sorger, V.J. (2016). Optimization of data center battery storage investments for microgrid cost savings, emissions reduction, and reliability enhancement. IEEE Trans. Ind. Appl. *52*, 2053–2060. https://doi.org/10.1109/TIA.2016.2517149.
60. Zhang, Y., Tang, H., Li, H., and Wang, S. (2025). Unlocking the flexibilities of data centers for smart grid services: Optimal dispatch and design of energy storage systems under progressive loading. Energy *316*, 134511. https://doi.org/10.1016/j.energy.2025.134511.
61. Sarkar, S., Naug, A., Luna, R., Guillen, A., Gundecha, V., Ghorbanpour, S., Mousavi, S., Markovikj, D., and Ramesh Babu, A.R. (2024). Carbon footprint reduction for sustainable data centers in real-time. In Proceedings of the AAAI Conference on Artificial Intelligence, pp. 22322–22330. https://doi.org/10.1609/aaai.v38i20.30238.
62. Ebrahimi, K., Jones, G.F., and Fleischer, A.S. (2014). A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities. Renew. Sustain. Energy Rev. *31*, 622–638. https://doi.org/10.1016/j.rser.2013.12.007.
63. Hao, Y., Zhou, H., Tian, T., Zhang, W., Zhou, X., Shen, Q., Wu, T., and Li, J. (2025). Data centers waste heat recovery technologies: Review and evaluation. Appl. Energy *384*, 125489. https://doi.org/10.1016/j.apenergy.2025.125489.
64. Bragg-Sitton, S., Misenheimer, C., Frick, K., Doster, J.M., and Terry, S. (2018). Absorption chiller and stratified chilled-water storage tank configurations for coupling to a small modular reactor. Tech. rep. (Idaho National Lab. [INL]).
65. Alkrush, A.A., Salem, M.S., Abdelrehim, O., and Hegazi, A.A. (2024). Data centers cooling: A critical review of techniques, challenges, and energy saving solutions. Int. J. Refrig. *160*, 246–262. https://doi.org/10.1016/j.ijrefrig.2024.02.007.
66. Gan, C.L., Chung, M.-H., Huang, C.-Y., and Jao, L. (2024). Hydrogen as power storage technology, polymeric and interconnect material innovations for future ai datacenter applications: a review. J. Mater. Sci.: Mater. Electron. *35*, 1951.
67. Hassan, Q., Algburi, S., Sameen, A.Z., Jaszczur, M., and Salman, H.M. (2024). Hydrogen as an energy carrier: properties, storage methods, challenges, and future implications. Environ. Syst. Decis. *44*, 327–350. https://doi.org/10.1007/s10669-023-09932-z.
68. Azarifar, M., Arik, M., and Chang, J.-Y. (2024). Liquid cooling of data centers: A necessity facing challenges. Appl. Therm. Eng. *247*, 123112. https://doi.org/10.1016/j.applthermaleng.2024.123112.
69. Abushamah, H.A.S., Burian, O., Ray, D., and Škoda, R. (2024). Integration of district heating systems with small modular reactors and organic rankine cycle including energy storage: Design and energy management optimization. Energy Convers. Manag. *322*, 119138. https://doi.org/10.1016/j.enconman.2024.119138.
70. Ullmann, J., Mašata, D., Závorka, J., Jiřičková, J., and Škoda, R. (2026). Flexible nuclear district heating with the teplator concept: Orc cogeneration, load shifting, and data center integration. Sustain. Energy Technol. Assess. *87*, 104912. https://doi.org/10.1016/j.seta.2026.104912.
71. Deng, Y., Li, X., Guliyeva, S., and Kenjayeva, U. (2025). Optimization of district heating systems integrated with energy storage for enhanced economic performance. Nucl. Eng. Des. *444*, 114393. https://doi.org/10.1016/j.nucengdes.2025.114393.
72. Zhou, X., Xin, Z., Tang, W., Sheng, K., and Wu, Z. (2024). Comparative study for waste heat recovery in immersion cooling data centers with district heating and organic rankine cycle (orc). Appl. Therm. Eng. *242*, 122479. https://doi.org/10.1016/j.applthermaleng.2024.122479.
73. Aili, A., Long, W., Cao, Z., and Wen, Y. (2024). Radiative free cooling for energy and water saving in data centers. Appl. Energy *359*, 122672. https://doi.org/10.1016/j.apenergy.2024.122672.
74. Zhang, Y., Shan, K., Li, X., Li, H., and Wang, S. (2023). Research and technologies for next-generation high-temperature data centers – state-of-the-arts and future perspectives. Renew. Sustain. Energy Rev. *171*, 112991. https://doi.org/10.1016/j.rser.2022.112991.
75. Radovanović, A., Koningstein, R., Schneider, I., Chen, B., Duarte, A., Roy, B., Xiao, D., Haridasan, M., Hung, P., Care, N., et al. (2023). Carbon-aware computing for datacenters. IEEE Trans. Power Syst. *38*, 1270–1280. https://doi.org/10.1109/TPWRS.2022.3173250.
76. Jayanetti, A., Halgamuge, S., and Buyya, R. (2024). Multi-agent deep reinforcement learning framework for renewable energy-aware workflow scheduling on distributed cloud data centers. IEEE Trans. Parallel Distrib. Syst. *35*, 604–615. https://doi.org/10.1109/TPDS.2024.3360448.
77. Zhou, S., Zhou, M., Wu, Z., Wang, Y., and Li, G. (2024). Energy-aware coordinated operation strategy of geographically distributed data centers. Int. J. Electr. Power Energy Syst. *159*, 110032. https://doi.org/10.1016/j.ijepes.2024.110032.
78. Zare Ghaleh Seyyedi, A., Akbari, E., Mahmoudi Rashid, S., Nejati, S.A., and Gitizadeh, M. (2024). Application of robust optimized spatiotemporal load management of data centers for renewable curtailment mitigation. Renew. Sustain. Energy Rev. *204*, 114793. https://doi.org/10.1016/j.rser.2024.114793.
79. Yang, Y., Wang, D., Shi, J., Wu, C., and Han, Z. (2025). Uncertainty-aware day-ahead datacenter workload planning with load-following small modular reactors, ACM SIGENERGY Energy. Inform. Rev. *5*, 91–97.
80. Zhang, Z., and Jiang, J. (2024). On load-following operations of small modular reactors. Prog. Nucl. Energy *173*, 105274. https://doi.org/10.1016/j.pnucene.2024.105274.
81. Byun, H.-H., and Yim, M.-S. (2025). System modeling-based investigation of smr operations using variable t-avg control for cost-effective electricity generation in a grid. Nucl. Eng. Technol. *57*, 103304. https://doi.org/10.1016/j.net.2024.11.006.
82. You, Z., Wu, Y., Wang, J., Wei, Z., and Ling, J. (2025). Dynamic equilibrium mechanism: Integrating small modular reactors and optimized cooling for sustainable data centers. Ann. Nucl. Energy *224*, 111662. https://doi.org/10.1016/j.anucene.2025.111662.
83. Temiz, M., and Dincer, I. (2025). Thermal design and analysis of a floating small modular reactor and wind system for a sustainable community with floating data center. Therm. Sci. Eng. Prog. *61*, 103561.
84. Lian, Y., Li, Y., Zhao, Y., Yu, C., Zhao, T., and Wu, L. (2023). Robust multi-objective optimization for islanded data center microgrid operations. Appl. Energy *330*, 120344. https://doi.org/10.1016/j.apenergy.2022.120344.
85. Ding, Z., Cao, Y., Xie, L., Lu, Y., and Wang, P. (2019). Integrated stochastic energy management for data center microgrid considering waste heat recovery. IEEE Trans. Ind. Appl. *55*, 2198–2207. https://doi.org/10.1109/TIA.2018.2890789.

86. Wu, Y., Xue, X., Le, L., Ai, X., and Fang, J. (2020). Real-time energy management of large-scale data centers: A model predictive control approach. In 2020 IEEE Sustainable Power and Energy Conference (iSPEC), pp. 2695–2701. https://doi.org/10.1109/iSPEC50848.2020.9351010.

87. Kahil, H., Sharma, S., Välisuo, P., and Elmusrati, M. (2025). Reinforcement learning for data center energy efficiency optimization: A systematic literature review and research roadmap. Appl. Energy *389*, 125734. https://doi.org/10.1016/j.apenergy.2025.125734.

88. Jacob, R.A., and Zhang, J. (2023). Modeling and control of nuclear–renewable integrated energy systems: Dynamic system model for green electricity and hydrogen production. J. Renew. Sustain. Energy *15*, 046302. https://doi.org/10.1063/5.0139875.

89. Athavale, J., Bash, C., Brewer, W., Maiterth, M., Milojicic, D., Petty, H., and Sarkar, S. (2024). Digital twins for data centers. Computer *57*, 151–158. https://doi.org/10.1109/MC.2024.3436945.

90. Chen, J., Tan, R., Wang, Y., Xing, G., Wang, X., Wang, X., Punch, B., and Colbry, D. (2012). A high-fidelity temperature distribution forecasting system for data centers. In 2012 IEEE 33rd Real-Time Systems Symposium, pp. 215–224. https://doi.org/10.1109/RTSS.2012.73.

91. Yuan, H., Bi, J., Li, S., Zhang, J., and Zhou, M. (2024). An improved lstm-based prediction approach for resources and workload in large-scale data centers. IEEE Internet Things J. *11*, 22816–22829. https://doi.org/10.1109/JIOT.2024.3383512.

92. Bianchini, R., Belady, C., and Sivasubramaniam, A. (2024). Data center power and energy management: past, present, and future. IEEE Micro *44*, 30–36. https://doi.org/10.1109/MM.2024.3426478.

93. Naug, A., Guillen, A., Kumar, V., Greenwood, S., Brewer, W., Ghorbanpour, S., Babu, A.R., Gundecha, V., Gutierrez, R.L., and Sarkar, S. (2025). Lc-opt: Benchmarking reinforcement learning and agentic ai for end-to-end liquid cooling optimization in data centers. NeurIPS 2025.

94. U.S. Nuclear Regulatory Commission. (2026). 10 CFR Part 53 – Risk. Informed, Technology-Inclusive Regulatory Framework for Commercial Nuclear Plants. https://www.nrc.gov/reactors/new-reactors/advanced/modernizing/rulemaking/part-53.html.

95. U.S. Nuclear Regulatory Commission (NRC) (2025). Licensing considerations relating to use of nuclear energy for nonelectric applications. Tech. rep. (U.S. Nuclear Regulatory Commission) https://www.nrc.gov/docs/ML2505/ML25055A251.pdf

96. (2024). NERC. Tech. rep. https://www.nerc.com/globalassets/our-work/assessments/2024-ltra_corrected_july_2025.pdf

97. Federal Energy Regulatory Commission (FERC). (2024). Order rejecting amendments to interconnection service agreement (FERC eLibrary). https://elibrary.ferc.gov/eLibrary/filelist?accession_number=20241101-3061optimized=false.

98. U.S. Nuclear Regulatory Commission. (2025). About the advance act. https://www.nrc.gov/about-nrc/governing-laws/advance-act/about-advance-act.html.