Grid Integration of AI Data Centers and Small Modular Reactors (SMR): The Role of Energy Storage

Dr. Wencong Su (Chair and Professor)
Dept. of Electrical & Computer Engineering
University of Michigan-Dearborn
www.SuWencong.com
Email: wencong@umich.edu



2026 Nuclear Energy and Applications Workshop
University of Illinois Urbana-Champaign
October 19–20, 2026

Power Engineering Research at UM-Dearborn ECE (1)



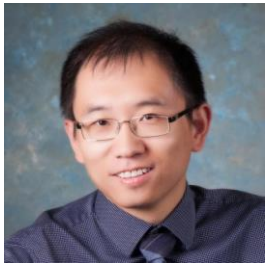
Wencong Su,
Professor and Chair
(optimization, control)



John O'Donnell,
Professor of Engineering Practice



Maggie Wang,
Professor
(power electronics)



Xuan (Joe) Zhou,
Associate Professor
(energy storage)



Junho Hong,
Associate Professor
(cybersecurity)



Taehyung Kim,
Associate Professor
(motor drive)



Guilherme Hollweg,
Assistant Professor
(control theory)

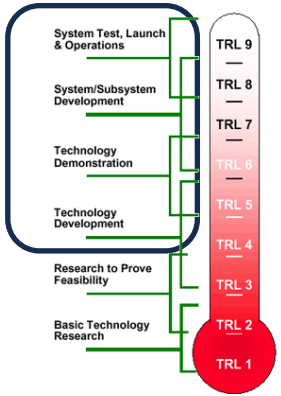


Van-Hai Bui,
Assistant Professor
(machine learning)



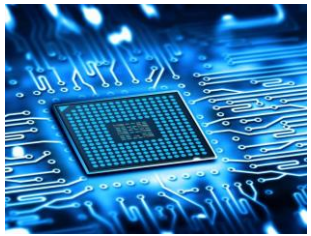
Apurv Shukla
Assistant Professor
(optimization & control)





UM-Dearborn sits at the center of the automotive industry, surrounded by leading OEMs & suppliers.

Power Engineering Research at UM-Dearborn ECE (2)



Power Electronics and Motor Drive



Artificial Intelligence in Energy Systems



Cybersecurity in Energy Systems

* A partial list of research sponsors

- Conducting cutting-edge federal-sponsored research projects (e.g., DoE, DoD, NSF).
- Supporting small business in Michigan (e.g., NSF SBIR/STTR).
- Strengthening academia-industry collaboration in Southeast Michigan.
- Providing hands-on research experiences to students.
- Translating lab-based research into commercially viable product.
- ……

Everybody Wants Them. Nobody Wants Them Next Door

Growing interest

~945 TWh

Global data center electricity by 2030, more than double 2024; U.S. data centers on course for nearly half of U.S. demand growth this decade.

Source: IEA, Energy and AI (2025)

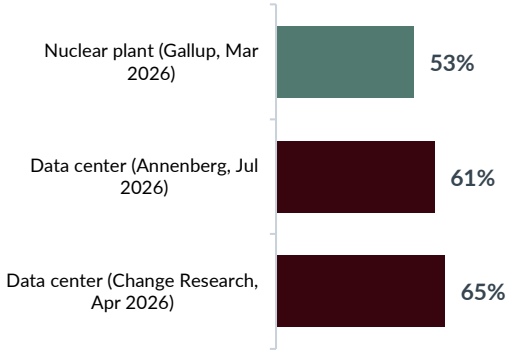
7 + 51

SMR designs operating or under construction, plus designs in licensing across 15 countries; over 10 GW of hyperscale new-nuclear commitments.

Source: Yale Clean Energy Forum (2026)

But not in my backyard

"Would you oppose one in your local area?"



Nuclear's paradox: in-principle support is at a record high (46% want more emphasis), yet a majority still opposes a plant nearby.

Why the pushback

Data centers

- Electricity bills (88% concerned); grid reliability (83%)
- Water and land use (80%); noise, diesel backup
- Few permanent jobs for the tax breaks granted

Nuclear / SMRs

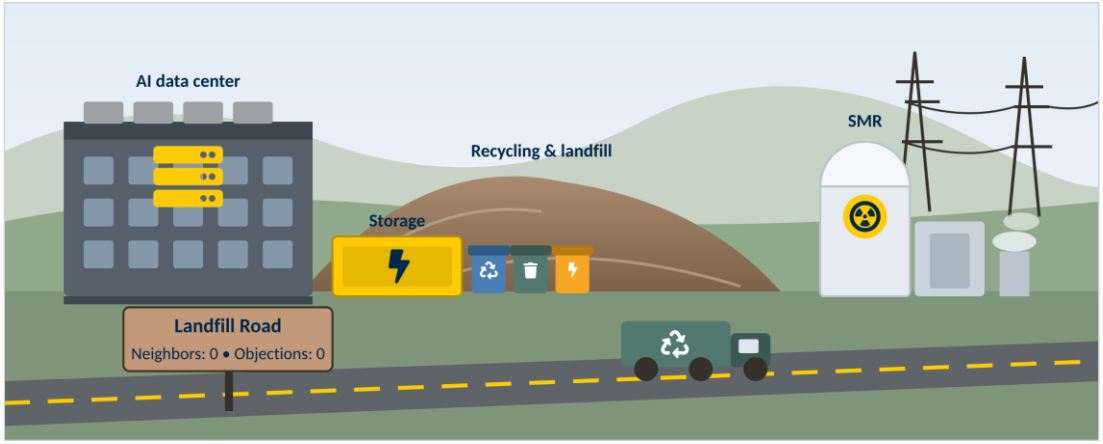
- Safety and waste legacy: Three Mile Island, Fukushima
- Cost overruns and schedule slips on recent builds
- Trust gap: 67% of women, 59% of over-55s oppose a local plant

A fun question for the room:

If a majority will fight a data center and a majority will fight a reactor, where in America could you site a 300 MW SMR and a 300 MW AI data center side by side without a single objection? *Shout out your answer; mine is on the next slide.*

Site Selection, Solved

Q: Where could you site an SMR and an AI data center side by side without a single objection?



A: The landfill.

- Industrial zoning is already in place
- Transmission, road access, and often water rights are already permitted
- The garbage trucks got there first, so nobody moved in for the quiet
- Reactor, servers, and storage on one bus is exactly the configuration the rest of this talk is about

The serious version

Brownfields, closed landfills, and retired coal plants are the real siting opportunity. DOE's coal-to-nuclear study found that a large majority of retired and operating coal plant sites could host advanced reactors, reusing the switchyard, cooling water, rail, and local workforce. SMR developers are already targeting former coal sites, and data center developers are following the same transmission.

Two Emerging Players, One Shared Challenge

AI data center: a "wild" load

- 100 MW to GW-scale campuses, concentrated at single points of interconnection
- Synchronized GPU training steps => sub-second power swings of tens of percent of rated load
- Near-zero tolerance for interruption => today backed by diesel + UPS
- Utilities and ISOs now scrutinizing ramp rates, ride-through, and flexibility of large loads

SMR: a "stiff" generator

- Factory-built reactors, roughly 50 to 300 MWe per module, sited near load
- High capital cost, low fuel cost => economics favor flat, full-power operation
- Xenon dynamics, thermal cycling, and licensing limit load following
- Grid increasingly needs flexibility, not just more baseload

Thesis

Energy storage is the shock absorber that lets each connect to the grid on the grid's terms. For the data center it removes volatility from a load; for the SMR it adds flexibility to a source. The same hardware, two opposite jobs.

AI data centers	SMRs	Co-location	Research
Load behavior; storage roles by timescale	Flexibility limits; storage as a decoupler	SMR + data center + storage as a microgrid	Open questions for the next 3–5 years

What Makes AI Data Center Load Different (1)

10×
More energy per AI query vs Google search

3.4mo
AI training workload doubling period

>100kW
Rack-level power in AI DCs

1,500 MW
Sudden load drop (Virginia grid event)

Key Grid Challenges Introduced by AI Data Centers

Rapid Power Fluctuations

Sub-second power spikes from training jobs differ fundamentally from traditional predictable loads, stressing upstream grid infrastructure.

Frequency Stability

Large DCs transitioning to backup can cause sudden GW-scale load drops, triggering frequency deviations and operator intervention.

Harmonic Distortion

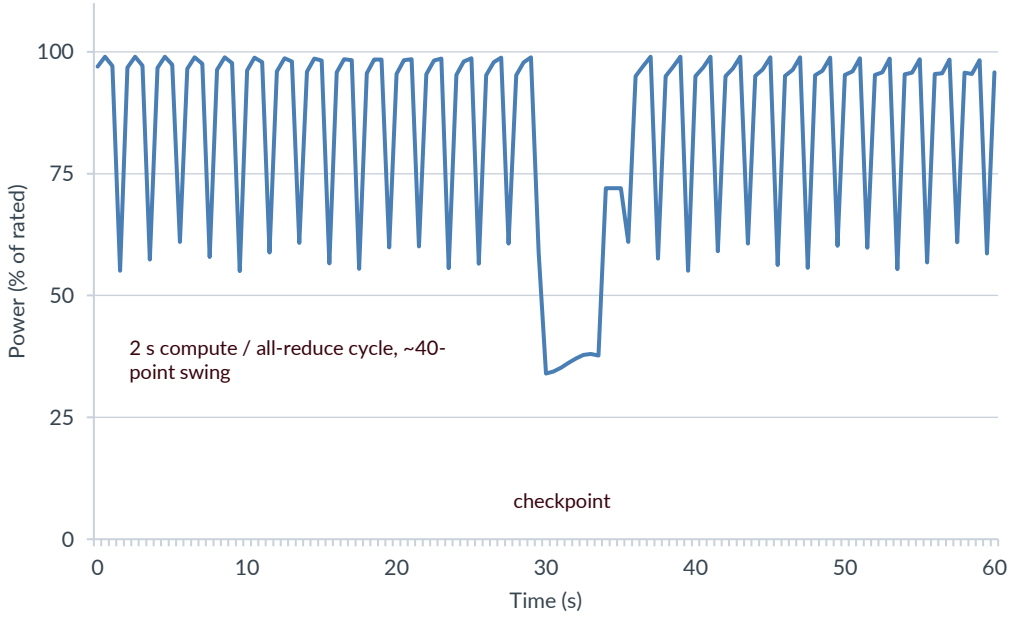
High-density power converters introduce power quality issues at the utility level, requiring active compensation strategies.

Reserve Margin Pressure

Highly uncertain AI DC loads force grid operators to carry additional spinning reserve, increasing operational costs.

What Makes AI Data Center Load Different (2)

Illustrative GPU-cluster power during distributed training



Synthetic profile for illustration.

Synchronized training steps

Thousands of GPUs compute, then wait on communication, together. Swings of tens of percent every few seconds, sub-second edges.

Checkpoints and job transitions

Step changes of 30–60% of site load lasting seconds to minutes.

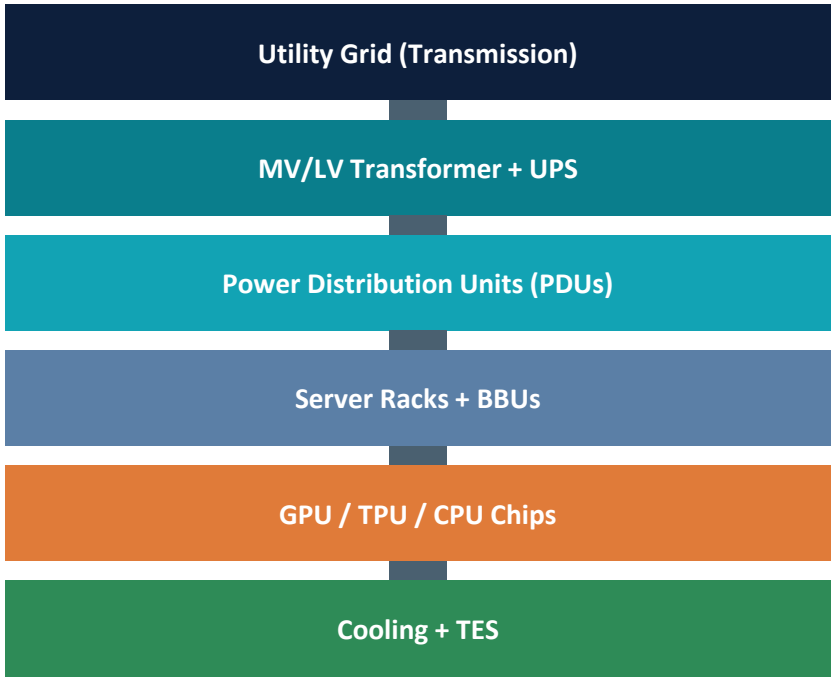
Grid-side consequences

Voltage flicker, frequency deviation on weak grids, interaction with protection and inverter controls.

Emerging requirements

Ramp-rate limits, ride-through and curtailment expectations in large-load interconnection rules.

AI Data Center Architecture Overview



AI DCs *v.s.* Traditional DCs

Attribute	Traditional DC	AI DC
Power Level	< 30 MW	> 100 MW – GW
Rack Power	Few kW	> 100 kW+
Grid Connection	Distribution	Transmission
Load Profile	Predictable	Highly dynamic
Cooling	Air-cooled	Liquid / hybrid
Location	Near population	Remote / semi-rural

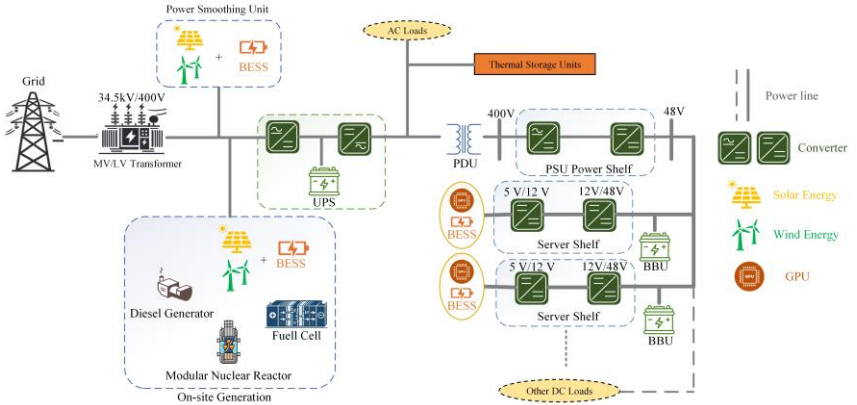


Fig. I. Overall AI DC architecture with multi-layer ESS

For more information, please see our paper: "Grid Integration of AI Data Centers: A Critical Review of Energy Storage Solutions," available at https://arxiv.org/abs/2603.00415

Grid-Interactive UPS (GiUPS) Systems

Traditional UPS → Grid-Interactive UPS

Standard Operation

100% grid supply. Battery on standby.

Partial Disconnect

Grid + battery share the load (e.g., 75/25).

Full Disconnect

100% battery supply for islanded operation.

Recharge Mode

Battery recharging from grid or RES.

Energy Export

Bidirectional: inject energy back to grid.

Key GiUPS Capabilities

Fast Frequency Response

Responds within 0.5–10 s of a grid event; dynamic or static injection modes.

Reactive Power Support

±250–300 kVAr per 1 MVA unit; improves PCC power factor to >0.99.

Voltage Ride-Through

LVRT/HVRT compliance; delivers 2–3 pu reactive current within 2–5 ms during sags.

Grid-Forming Mode

Can regulate terminal voltage/frequency directly and transition to islanded operation.

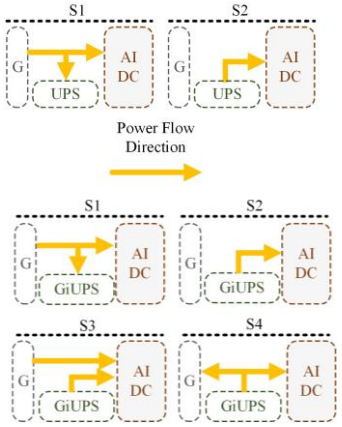
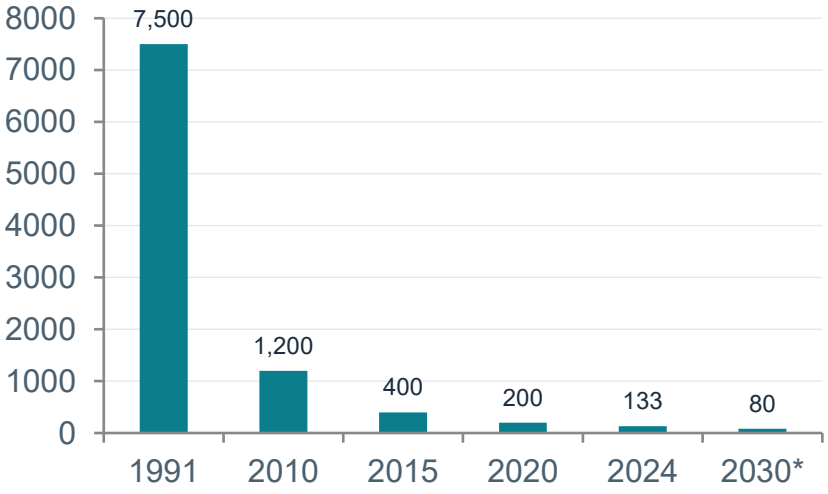


Fig. 2. UPS and GiUPS operation modes

Battery Energy Storage Systems (BESS) in AI DCs

BESS Cost Trajectory ($/kWh)



*projected

BESS Grid-Support Roles

Power Smoothing

Attenuates rapid AI DC demand variations; presents smooth profile to grid. GFM BESS effectively handles training/inference fluctuations.

Renewable Integration

Stores excess RES generation; enables solar/wind co-generation; supports PPAs and carbon goals.

Frequency Regulation

Participates in primary freq. regulation, FFR, and synthetic inertia via GFL/GFM modes.

Power-Smoothing Architectures — Trade-offs

GFM BESS

GFM+GFL BESS

E-STATCOM+BESS

Supercapacitors

Extends UPS for 1–4 hours; can replace diesel generators for sustained outages.

UPS-Integrated BESS: Operation Modes

S1 Normal Operation

AI DC load 100% from utility grid. BESS and UPS on standby.

S2 Full Disconnection (BESS OK)

Site isolated. BESS supplies 100% of load for 1–4 hours.

S3 Full Disconnection (BESS Low)

Long outage, BESS insufficient — UPS assumes critical load support.

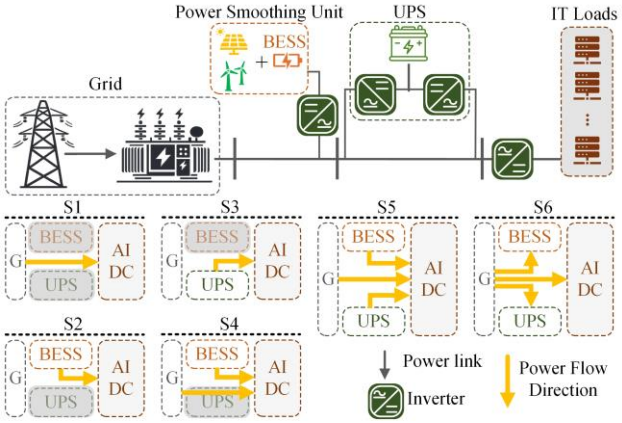


Fig. 3. UPS-Integrated BESS operation modes

S4 Partial Disconnection

BESS supplements grid to support frequency regulation & smoothing.

S5 Coordinated Smoothing

Grid + BESS + UPS jointly supply AI DC. UPS handles fast transients.

S6 Recharge Mode

After over-frequency: UPS recharges first, then BESS from grid or RES.

Thermal Energy Storage (TES) in AI DCs

Cooling = 30–40% of AI DC energy consumption. ~90% of GPU power becomes waste heat.

Sensible TES (Chilled Water)

✓ Advantages

· Lowest cost per kWh stored

· Mature, widely deployed

· Simple integration with chilled-water systems

✗ Challenges

· Low volumetric energy density

· Large footprint required

Latent TES (Phase-Change / Ice)

✓ Advantages

· Higher volumetric density than sensible

· Constant-temperature storage

· Proven in large-scale district cooling

✗ Challenges

· More complex heat exchanger

· Higher capital cost

Thermochemical TES

✓ Advantages

· Highest energy density of all three

· Lowest standby losses

· Potential for long-duration storage

✗ Challenges

· Complex and expensive

· Not yet mature for DC deployments

RES Integration: During surplus generation or low electricity prices, chillers charge TES systems. This shifts cooling-related electrical demand in time, enhancing AI DC flexibility and grid decarbonization.

Rack- & Server-Level Energy Storage

Rack-Level BBUs

- Distributed DC ride-through: OCP Open Rack V3 (ORV3) BBU shelves with 5+1 redundancy
- Timescale: seconds to minutes — bridges gap to BESS/gensets
- Coordinated charging reduces recharge power spikes by ~80%
- Selective deployment (high-priority racks) reduces CapEx
- Li-ion safety considerations: fire/thermal runaway near IT equipment requires conservative SoC windows
- NOT ideal for sub-second smoothing (accelerates aging)

Server/GPU-Level Capacitors

- Supercapacitors handle microsecond-to-millisecond GPU power swings
- Tolerate extremely high charge/discharge cycle counts with minimal degradation
- Sized for transient energy of rapid GPU swings — not ride-through
- NVIDIA GB300 NVL72: firmware-level 'power burner' mode for ramp-rate control
- NVML API exposes enforced GPU power limits and real-time telemetry
- Cross-stack co-design: hardware buffers + firmware shaping + software control

Aspect	Rack BBUs	Server Supercaps
Timescale	Seconds – minutes	Microseconds – ms
Purpose	Ride-through / backup	Sub-second smoothing
Cycle Life	Limited (aging concern)	Very high (minimal aging)
Energy Density	High (Li-ion)	Low

Second-Life Battery Energy Storage Systems (SLBESS)

EV batteries are retired at ~80% capacity. Repurposed for stationary AI DC storage — lower cost, circular economy.

BTM Backup Generation

SLBESS qualifies as behind-the-meter source. Hybrid configs combine used & new packs. BMS balances heterogeneous capacities via active/passive balancing.

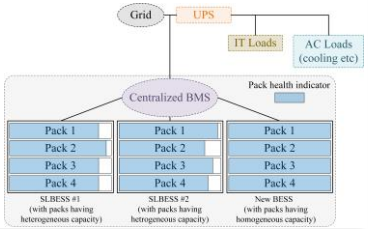


Fig. 4. SLBESS heterogeneous pack configuration for AI DC UPS backup

Hybrid GiUPS Support

SLBESS for short-duration backups; new BESS handles FFR. UPS controller uses SLBESS to recharge new BESS — reduces grid energy import.

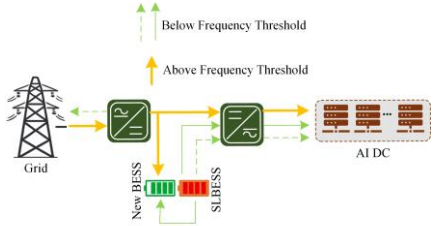


Fig. 5. FFR by utilizing new BESS and SLBESS in GiUPS architecture

Grid-Scale FTM SLBESS

Shaves AI DC load peaks, acts as standby generation, participates in ancillary markets. Enables hyperscalers (Meta, Microsoft, Apple) to manage PPA surplus.

Sizing & Operational Challenges

Narrow SoC Window

SLBs operated 30–70% SoC (vs 10–90% for new) to limit degradation — requires oversizing in both kWh and kW.

SoH Heterogeneity

Variable state-of-health across packs causes capacity dispersion; must be accounted for in system-level sizing.

Degradation Monitoring

More rigorous lifecycle tracking required vs new BESS; Battery Passport (EU) and UL 1974 compliance mandatory.

Cost Dynamics

Stationary storage cost fell $169→$108/kWh (2020–2025); declining new-BESS prices may reduce SLBESS cost advantage.

Multi-Layer ESS Architecture: Putting It Together

Layer 1: Chip/GPU Level

Supercapacitors | Firmware shaping

µs – ms

Layer 2: Server / Rack Level

BBUs (OCP ORV3) | Local DC bus

Seconds – minutes

Layer 3: Facility UPS / GiUPS

Bidirectional UPS | FFR | VRT | Peak shaving

Minutes

Layer 4: On-Site BESS / SLBESS

GFM / GFL BESS | Renewable integration | Grid services

Minutes – hours

Layer 5: Grid-Scale / FTM

Transmission-connected BESS | Reserve | Market participation

Hours – days

What Makes SMRs Hard to Flex

Baseload economics

High capital cost and very low fuel cost mean every MWh not produced is lost revenue. The plant wants a straight line.

Reactor physics

Xenon-135 transients after power changes constrain how fast and how often the core can maneuver, especially in cycle.

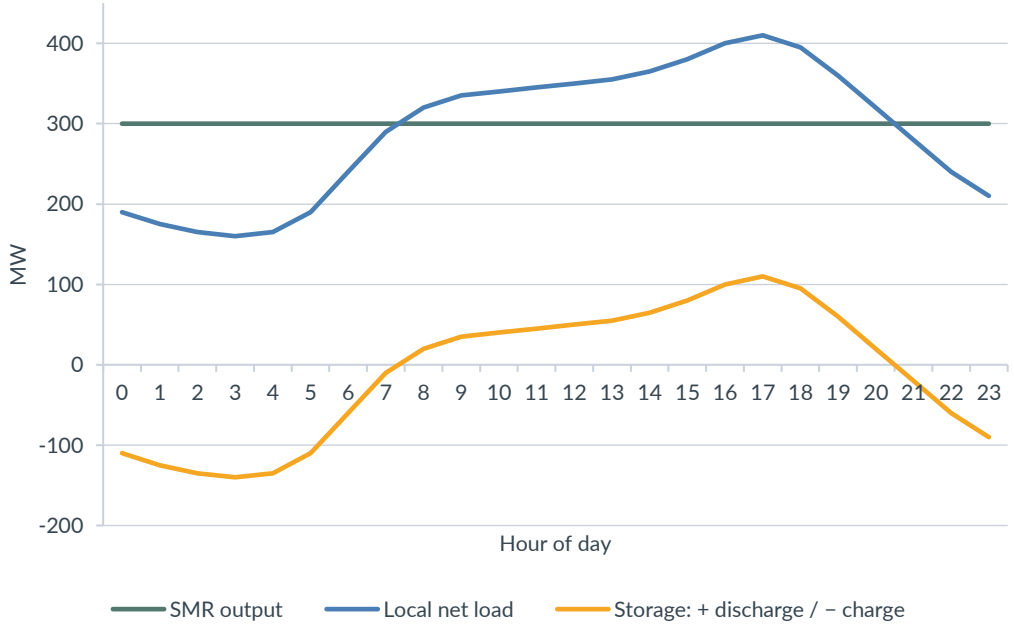
Materials and licensing

Thermal cycling drives fatigue in the primary system; the licensed operating envelope may not permit daily load following.

The mismatch

The grid increasingly rewards flexibility, and co-located loads are anything but flat.

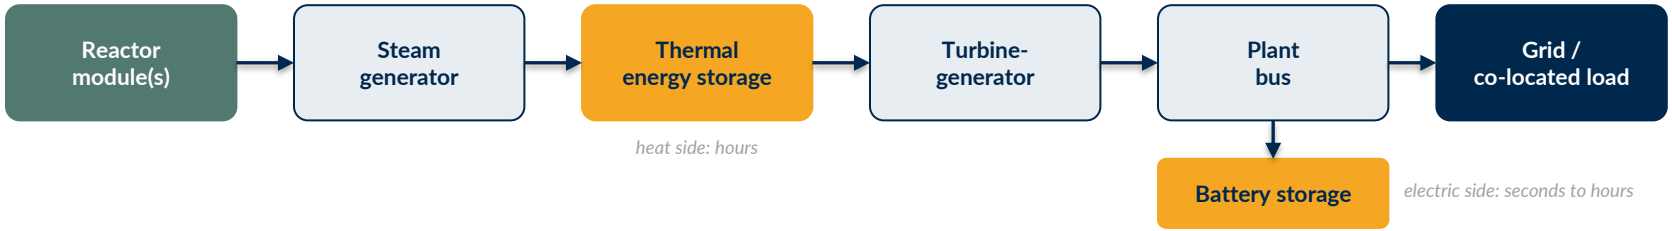
Flat reactor output versus a variable local net load (illustrative, MW)



Storage absorbs the difference so the reactor never moves; the area between the curves sizes the energy requirement.

Storage as a Decoupler Between Reactor and Grid

Where storage sits in an SMR plant



Thermal energy storage

- Molten salt, firebrick, or steam accumulators between steam generator and turbine
- Hours-scale flexibility without moving reactor power
- Oversized turbine lets the plant surge above reactor rating

Battery storage at the bus

- Fast transients, frequency response, voltage support
- Ride-through for reactor trips and turbine runbacks
- Smooths co-located loads before they reach the reactor

New grid interactions

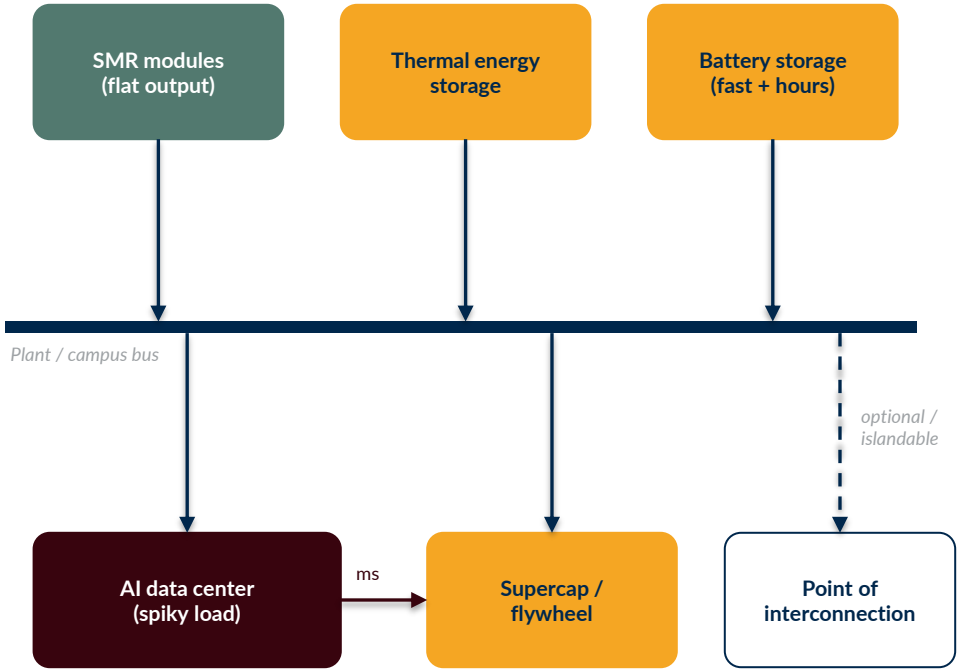
- Islanded operation and black-start capability
- Capacity, reserve, and ancillary-service revenue on top of energy
- Reactor stays at 100%; storage faces the market

Same Technology, Opposite Jobs

	AI data center	SMR
Role of storage	Remove volatility from a load	Add flexibility to a source
Dominant timescale	Sub-second to minutes (power problem)	Hours (energy problem)
Primary technology	Supercapacitor / BESS	Thermal energy storage + Li-ion BESS
Control objective	Make the load look conventional at the point of interconnection	Keep the reactor at 100% while the plant follows the grid
Grid interaction	Ramp-rate compliance, ride-through, demand response	Islanding, black start, capacity and ancillary markets
Sizing driver	Step magnitude and ramp limit (MW); backup duration (MWh)	Area between flat output and net load (MWh); max gap (MW)

Put them on the same bus and the two jobs collapse into one: storage absorbs the difference between a straight line and a sawtooth. That is the co-location case on the next slide.

Co-Locating SMR, Data Center, and Storage



TES on the steam side feeds the turbine; shown at the bus for simplicity.

Storage absorbs the mismatch

A flat generator and a sawtooth load cannot share a bus without something in between. Fast storage handles the edges; thermal and battery energy handle the daily shape.

Islanding becomes credible

With storage covering transients and reactor trips, the grid connection becomes an option rather than a lifeline. That changes the interconnection conversation entirely.

The sizing question

How much fast storage, how much energy storage, and how much grid import is the cheapest mix for a target reliability? Nobody has a settled answer yet.

Research Opportunities and Open Challenges

1 Load modeling and forecasting for AI clusters

- Open, measured power datasets at sub-second resolution; workload-aware stochastic models usable in grid studies
- AI DC profiles are highly unpredictable (training and inference bursts); accurate forecasts are essential for ESS sizing and dispatch

2 Sizing, siting, and cost

- Joint power / energy sizing for hybrid portfolios, robust to unknown future workloads and tariffs
- Little public research on multi-layer ESS sizing in AI DCs: one hyperscale BESS or several grid-scale units?
- Long-duration storage (flow batteries, salt cavern) to firm on-site renewables is not yet mature

3 Hierarchical control and coordination

- Centralized EMS must orchestrate chip-, rack-, facility-, and grid-scale storage across millisecond-to-hour timescales
- Energy-aware GPU scheduling (DVFS) is immature; schedulers optimize throughput, not energy, and accelerate ESS cycling
- Stability of storage-heavy microgrids with a stiff nuclear source

4 Degradation under new duty cycles

- High-frequency cycling in AI environments accelerates aging; aging models must be validated for micro-cycling profiles
- Online SoH monitoring and RUL prediction, especially for second-life BESS
- Lifetime-aware dispatch that protects the energy battery

5 Standards, interconnection, and equipment

- Large-load ramp-rate and ride-through requirements; islanding, protection, and reactor-trip coordination rules
- Bidirectional flow from GiUPS / BESS causes thermal aging and voltage-regulation issues in unidirectional distribution transformers

6 Simulation, validation, and economics

- No tool today combines transient electrical with steady-state thermal simulation of AI DC EMSs; multi-solver integration is missing
- Value stacking across site, utility, and market
- Reliability indices for co-located SMR + data center systems

Takeaways

1

Opposite problems, one enabler

AI data centers need their volatility removed; SMRs need flexibility added.
Energy storage is the shock absorber in both cases.

2

Design by timescale, not by technology

Sub-second, minutes, and hours are separate problems with separate hardware.
Hybrid portfolios outperform one battery stretched across all three.

3

Co-location is where it gets interesting

A flat reactor, a spiky load, and storage on one bus make islanding credible and turn the grid connection into an option.