DATA CENTRES AS “VIRTUAL POWER PLANTS”: EMERGING GRID CODE REQUIREMENTS

Prepared by: Kumar M

24 October 2025

INDEX

ABBREVIATION

Abbreviation	Full Form / Description
AURA	Automated Utility Response and Allocation
AURA EDGE DC	AI-Based Data Centre Grid Orchestrator
AI	Artificial Intelligence
API	Application Programming Interface
ATS	Automatic Transfer Switch
BESS	Battery Energy Storage System
CMS	Condition Monitoring System
CMMS	Computerised Maintenance Management System
DC	Data Centre
DG	Diesel Generator
DNP3	Distributed Network Protocol 3
EMS	Energy Management System
EPM	Energy Process Management
EDGE	Intelligent Bidding Engine for Decentralised Grid Economics
FRT	Fault Ride-Through
HPC	High Performance Computing
HVRT	High Voltage Ride-Through
ICT	Information and Communication Technology
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet Protocol
LVRT	Low Voltage Ride-Through
ML	Machine Learning
MV	Medium Voltage
NMS	Network Management System
OT	Operational Technology
PCS	Power Conversion System
PF	Power Factor

PLC	Programmable Logic Controller
PMU	Phasor Measurement Unit
PPC	Power Plant Controller
PR	Performance Ratio
PV	Photovoltaic
RMS	Root Mean Square
ROCOF	Rate of Change of Frequency
SCADA	Supervisory Control and Data Acquisition
SDN	Software Defined Networking
SGA	Smart Grid Analytics Pvt. Ltd.
SOC	State of Charge
SOH	State of Health
TCP/IP	Transmission Control Protocol / Internet Protocol
TSO	Transmission System Operator
UPS	Uninterruptible Power Supply
VPP	Virtual Power Plant
VPN	Virtual Private Network

1. ABSTRACT

In the push toward digitalization and decarbonization, Data Centres have become **grid-relevant electrical entities** capable of influencing system frequency, voltage, and stability at scale. Once treated as benign loads, hyperscale campuses now function as **controllable infrastructure** whose behavior during disturbances can either buffer or amplify system stress. Regulators in **ERCOT (Texas), AESO (Alberta), Fingrid (Finland), and CRU (Ireland)** are converging on a new operating reality: Data Centres should **remain connected** during common disturbances (LVRT/HVRT), **provide reactive/active support**, **participate in demand-response and frequency services**, and **coordinate on-site generation and storage** with the TSO/ISO [2]–[5]. In effect, the world's largest digital facilities are being asked to behave like **power plants**, with measurable performance limits, remote-dispatch interfaces, and audit-ready compliance.

This paper consolidates those developments and derives the **technical implications** for electrical architecture and control: **coordinated UPS behavior**, **staged reconnection and ramp-rate limits**, **curtailment interfaces**, and **microgrid orchestration** that blends BESS, fuel cells or engines, and flexible IT loads. We also examine how **AI/HPC workloads** introduce ultra-fast, nonstationary demand profiles that stress conventional control and stability assumptions. Finally, we frame a research agenda around **synthetic inertia from UPS fleets**, **multi-timescale hierarchical control**, and **standardized grid-compliance test suites**, positioning the future Data Centre as a **dispatchable, standards-certified, revenue-earning** participant in the power system. [1], [4], [6], [7]

Keywords— Data Centre; Grid Code; Fault Ride-Through (FRT); Virtual Power Plant (VPP); Demand Response; Energy Storage; Inertia Emulation; Load Flexibility.

2. INTRODUCTION

The digital backbone of the twenty-first century is built on a quiet but voracious consumer of electricity: Data Centres. Each click, query, and computation hides a stream of electrons flowing through facilities that rival small towns in power demand. Driven by exponential growth in cloud computing, artificial-intelligence inference, and edge-computing workloads, global Data Centre electricity consumption is projected to climb sharply. In the United States alone, analysts forecast an increase from roughly 3–4 % of total electricity use today to nearly 12 % by 2030 [1]. Emerging AI clusters in Europe and Asia add comparable stresses to their national grids. What began as a niche industrial load has now become an active participant in national energy planning.

Unlike conventional industries, modern Data Centres operate with tight electrical tolerances. Their reliance on uninterruptible power supplies, static transfer switches, and sophisticated cooling systems means that even minor grid disturbances can trigger rapid shifts in demand. A coordinated cluster of hyperscale facilities switching simultaneously to backup generation or battery mode can withdraw hundreds of megawatts from the grid in milliseconds. When those same facilities reconnect to recharge their UPS banks, the load surge can mirror the impact of a large power-plant trip, but in reverse. These oscillations complicate frequency and voltage stability and challenge system operators who once considered such loads benign.

Historically, grid codes have focused almost exclusively on generation assets. Power plants were obligated to provide fault ride-through capability, frequency regulation, reactive-power control, and synchronization discipline, while large loads faced little scrutiny. The rise of concentrated Data Centre developments has forced regulators to rethink that balance. Transmission system operators in regions such as the Electric Reliability Council of Texas (ERCOT), Alberta Electric System Operator (AESO), Fingrid Oyj (Finland) and Commission for Regulation of Utilities (CRU, Ireland) now recognize that these digital infrastructures must be governed by “power-plant-like” operational rules [2]–[4]. Data Centres are therefore being reclassified as grid-interactive assets—sometimes described as virtual power plants (VPPs) or even negative generators—that must actively coordinate with the grid instead of simply drawing from it.

This paradigm shift carries deep engineering and policy implications. Data Centres can no longer be treated as passive consumers of energy. They must coordinate, communicate, and, in some cases, contribute back to the grid. Their design philosophy is moving from “always on” to “always adaptive,” capable of responding to external grid conditions without compromising internal reliability. This transition redefines not only power-system operations but also the future architecture of digital infrastructure itself.

For operators, integrators, and technology providers such as Smart Grid Analytics, whose Energy Process Management (EPM) platform integrates SCADA, EMS, and advanced analytics, this evolution represents both challenge and opportunity. The demand is not merely for more resilient Data Centre power systems but for intelligent systems capable of understanding grid signals, prioritizing internal loads, and offering services like demand response, frequency stabilization, and synthetic inertia. As regulatory frameworks tighten, such intelligence will become a prerequisite for interconnection and reliable operation.

The objectives of this paper are therefore threefold:

1. To survey global regulatory and grid-code developments that increasingly classify large Data Centres as quasi-generation facilities with explicit compliance obligations.
2. To analyze the technical and operational implications of these policies on electrical architecture, control hierarchies, and integration of on-site generation and storage.
3. To identify emerging research areas—from ultra-fast AI load dynamics to standardized grid-compliance testing, that remains underexplored yet are essential for shaping the next generation of grid-interactive Data Centres.

In the sections that follow, we examine the current regulatory landscape across major geographies, discuss the technical mechanisms through which Data Centres can satisfy these new obligations, and propose a forward-looking framework that positions them as stabilizing rather than destabilizing forces in modern power systems.

3. GLOBAL REGULATORY & GRID-CODE LANDSCAPE

A. North America

1) Electric Reliability Council of Texas (ERCOT), Texas, USA — Senate Bill 6 (2025) & ERCOT Protocol Updates

Scope: Large-load sites of 75 MW or more at a single location are subject to the new rules when seeking connection after December 31 2025. [5]

Key Obligations:

- The grid operator may order remote curtailment or full disconnection during system emergencies and may require the facility to switch to on-site backup generation.
- The customer must pay an interconnection study fee (minimum USD 100,000), accept cost responsibility for required upgrades, disclose any duplicate interconnection applications, and report on-site generation and co-location arrangements. [6]
- ERCOT has formalised a Large-Load Interconnection Process (NPRR-1234 / PGRR-115) that standardises modelling, system-impact studies and queue treatment for large loads. [7]

Implication for data centres: Facilities must be treated as dispatchable loads; they should support curtailment, coordinate UPS/generator transitions, and provide planning and telemetry data on par with a mid-sized power plant.

2) Alberta Electric System Operator (AESO), Alberta, Canada — Interim “Large Load Integration” (2025)

System Context: With nearly 29 Data Centre applications and requests totaling thousands of megawatts in the pipeline, AESO has allocated an interim cap of approximately 1,200 MW of new large-load capacity for connection between 2025 and 2028 to protect grid reliability. [8]

Scope: Single-site projects of 75 MW or larger may qualify in the first allocation stage only if they require no new major transmission reinforcements (Phase-1). [9]

Key Obligations:

- Applicants must submit detailed technical characteristics (including rapid ramp behavior and UPS/engine performance) for connection studies. AESO retains the right to stage or sequence commissioning. [9]

Implication: Developers should prepare for phased energization, provide full dynamic modelling, and recognize that connection access will be sequenced to maintain system stability.

B. Europe

3) Commission for Regulation of Utilities (CRU), Ireland—Proposed decision on new data-center connections (2025)

Core Rule: New Data Centres must provide on-site or proximate generation and/or storage capacity equal to their requested demand capacity, and that asset must be capable of market participation and reporting on renewable use and emissions annually. [10]

Implication: This effectively demands a “bring-your-own power-plant or storage” strategy, aligned with the data-centre’s growth, market-participation readiness and transparency.
4) Fingrid Oyj (Finland) — KJV2026 Draft: Grid-Code for Demand Connections (2025)
Finland is implementing generator-type requirements for large consumption facilities, including data centres. The draft code (KJV2026) details disturbance-ride-through and reconnection behaviour.

Scope thresholds: Demand facilities > 30 MW (Power Class F/G) and data centres & electric boilers > 10 MW (Power Class E) must comply. [11]

Selected Hard Limits & Behaviours:

- Rate of change of frequency (RoCoF) at the point-of-common-coupling (PCC) remains 2 Hz/s. [11]
- A voltage phase jump of ±30° must not lead to disconnection for eligible equipment (inverters, UPS, VFDs). [11]
- On-over-voltage ride-through (OVRT): equipment must remain connected through specified excursion curves (e.g., 1.00 pu = 118 kV @110 kV system, and 1.00 pu = 400 kV @400 kV system). [11]
- During sag conditions: active current must be limited when supply falls below 90 % V and blocked when below 50 % V (with exceptions for some power-factor–adaptive responses). [11]
- Must ride through ten separate 100 ms bolted faults within 90 s (accounting for reclosing and multiple events). [11]
- Post-fault active-power recovery must follow explicit ramp-rate limits to avoid system stress; operator-signal coordination is required. [11]
- TSOs may issue emergency control signals to demand facilities; dynamic voltage/reactive control expectations apply. [11]

Implication: Data-centre UPS, rectifier and reconnection logic must be engineered to remain grid-connected through disturbances and to recover smoothly, analogous to plant LVRT/HVRT and ramp-rate rules.

5) European Union Baseline — Demand Connection Code (DCC)

The Network Code on Demand Connection provides the baseline requirements for large consumption facilities: power-quality limits, capability to withstand system events, telemetry and coordination. Member-state TSOs add local numeric curves and frameworks. [12]

C. Summary Table - Changes

Region / Code	Who's In Scope	Must-Have Capabilities	Key Numeric Limits / Triggers
ERCOT (TX, USA) – SB 6; NPRR-1234 / PGRR-115	New "large loads" ≥ **75 MW** (post **Dec 31 2025** connections)	Remote **curtail/disconnect**; ability to **island on backup**; full interconnection modelling; cost-sharing	**USD 100k** study fee; disclosure of duplicate requests & on-site gen; ERCOT emergency authority to **shed** or **switch** to backup.
Alberta (AESO) – Large Load Integration (Interim)	Large-load DC projects ≥ **75 MW**	Phased access; detailed dynamic data; **no new Tx reinforcements** for Phase-1	**Cap: 1,200 MW (2025–2028)** total; sequencing by AESO for reliability.
Ireland (CRU) – Proposed DC Connection Policy (2025)	All **new** DC connections	**Gen/Storage = Load (1:1)** on-site or **local**; **market participation**; annual **renewables/emissions** reporting; location constraints apply	Match **requested demand** with dispatchable gen/ storage; operators in constrained zones face tighter scrutiny.
Finland (Fingrid) – KJV2026 Draft	Demand sites **> 30 MW**; DCs & e-boilers **> 10 MW**	**LVRT/HVRT**, **phase-jump** immunity, **multi-fault ride-through**, controlled **post-fault ramp**, dynamic voltage/ reactive behavior, TSO signaling	**RoCoF 2 Hz/s**; **± 30°** phase-jump: no trip; **10× 100 ms LVRT within 90 s**; **limit active current < 0.9 pu**, **block < 0.5 pu** V; OVRT per kV base; staged recovery.
EU (DCC baseline)	Large demand facilities	Power-quality, coordination, telemetry; withstand events	National TSOs add numeric curves; DCC underpins Fingrid's KJV.

D. Summary Table - Changes w.r.t Telemetry, Ancillary services

Region	Fault-ride-through (V/F)	Reconnection / Ramp Behaviour	Remote Disconnect / Load-shed	On-site Gen / Storage Obligations	Modelling, Telemetry, Planning Data	Ancillary / Market Participation
ERCOT (Texas, USA)	Emerging ride-through expectations for large loads; studied via new interconnection process	Controlled restoration expected; can be instructed to island on backup before reconnect	**Yes** — remote curtail/disconnect authority for new ≥75 MW sites (post-Dec 31, 2025)	Disclosure of on-site gen; co-location scrutinised	**Yes** — full interconnection studies, queue transparency, duplicate-request disclosure	**Encouraged/possible** via demand response programs for large loads
Alberta (AESO, Canada)	Studied case-by-case in connection assessment	Phased/ staged energization; ramp impacts reviewed	**Possible** — operational constraints and staged access	**Encouraged** (province promotes self-supply; AUC approvals if paralleled)	**Yes** — detailed technical characteristics and dynamic behaviour required	**Possible** (depends on market registration/ arrangements)
Ireland (CRU)	Must remain grid-friendly; specifics aligned with SO requirements	Staged reconnection consistent with market participation	**Requested** via TSO flexibility calls in tight conditions	Required — Gen/Storage sized ≈ requested demand (1:1), on-site or proximate, market-participating	**Yes** — annual reporting on renewables/ emissions; full connection study inputs	**Yes** — on-site assets must be market-capable
Finland (Fingrid)	**Explicit** LVRT/ HVRT & disturbance limits for large demand (e.g., phase-jump ±30°, multiple sag events, RoCoF tolerance)	**Explicit** controlled post-fault active-power recovery / ramp limits	**TSO signaling** for emergency control; site must respond	Not mandatory but often paired with UPS/BESS tuned to code	**Yes** — equipment data, dynamic models, operating schedules	**Expected/ feasible** (e.g., reactive/ voltage support; frequency services via UPS/BESS)
EU Baseline (DCC)	Baseline "withstand events" & power-quality obligations	Coordination on restoration	Member-state specific; curtailment frameworks exist	Not mandated at EU level	**Yes** — coordination, telemetry, quality limits	**Possible** in many markets via load participation frameworks

Summary

The boundary between load and generation is rapidly fading. Large data centres are increasingly treated as active grid participants: they now have responsibilities to ride through faults, reconnect in a controlled manner, accept remote curtailment, provide detailed operational modelling and telemetry, and in some jurisdictions deploy on-site dispatchable capacity and participate in market mechanisms [8]–[11].

4. TECHNICAL AND OPERATIONAL IMPLICATIONS FOR DATA CENTRES

The evolving grid-code landscape compels Data Centres to shift from electrically isolated islands into **grid-synchronous ecosystems**. Each new regulatory obligation translates into engineering and control requirements that span electrical architecture, operations, and workload management. The following subsections parse these implications in detail.

A. Fault Ride-Through (FRT) and Disturbance Behavior

For decades, only generation assets were required to ride through voltage sags, frequency deviations, or transient faults to protect system stability. Large Data Centres are now entering the same responsibility domain.

A typical hyperscale Data Centre may draw 50–300 MW through multiple medium-voltage feeders and parallel UPS strings. During a short-duration sag, an instantaneous transfer of all UPS strings to battery mode can cause a load drop of hundreds of megawatts. The grid perceives this as a negative generation event: frequency accelerates, protection relays may trip and nearby generators may destabilise.

Modern grid codes now expect the facility to remain grid-connected through such events:

- **Voltage tolerance**: Internal buses and UPS systems must remain connected to as low as ~0.2 pu voltage for at least 100–150 ms, in line with generator LVRT curves.
- **Frequency tolerance**: The facility should continue operations through ±2 Hz deviation from nominal frequency without tripping.
- **UPS coordination**: UPS strings should follow staggered transfer logic instead of an “all-or-nothing” switch, preserving partial grid interface during the disturbance.
- **On-site generation coordination**: Diesel, gas, or battery backup must synchronise transitions with no significant overlap or surge at reconnection.

To verify compliance, facility electrical models (UPS inverters, rectifiers, PFC banks, feeders) should be incorporated into dynamic stability simulations run by the TSO. This ensures the data-centre will not behave as a “negative generator” during faults [12]–[13].

B. Load Restoration and Controlled Reconnection

After a disturbance, reconnection can pose greater risk than disconnection. A 100 MW site that immediately reloads to full demand can trigger grid stress comparable to a generation loss. Grid operators now define maximum ramp-rates (for example, 10–20 MW/min) and expect staged restoration logic.

A modern Energy Process Management (EPM) system must choreograph this process:

1. **Prioritised sequencing**: Mission-critical IT racks and cooling systems are restored first; non-critical batch compute nodes reconnect later in defined intervals.
2. **Ramp control**: On-site BESS offset part of the draw while grid-supply ramps up smoothly.
3. **TSO coordination**: The EPM must accept a “reconnect permit” from the grid operator before restoring full load.

In practice this means implementing micro-islanding logic: finite-state controllers that transition through *fault* → *hold* → *partial restore* → *full restore*, each state with explicit timing and ramp-limits. The outcome is a grid-friendly facility whose reconnection does not exceed the system’s absorption capability [14].

C. Remote Disconnect, Demand Response and Dispatchable Loads

One of the most significant changes in current grid-codes is the expectation that large loads must become controllable by the system operator. TSOs now issue remote-curtailment or disconnect commands through secure SCADA or IEC 61850 links.

To comply, a data-centre must be equipped with:

- Automated feeder relays or breakers capable of switching to backup supply within seconds to a few minutes.
- Telecommand integration (IEC 60870-5-104, DNP3) enabling TSOs to signal load reduction or “switch-to-on-site” directives.
- Hierarchical load-shedding logic within the EPM: mission-critical racks remain active while less-critical clusters or chillers are shed in descending priority.
- Sufficient on-site generation and storage redundancy to maintain uptime during a grid-requested disconnection.

With these in place, data-centres effectively become **dispatchable loads** – capable of responding to grid frequency events within seconds. Some operators already monetise this through ancillary-service programs, offering primary frequency response via battery discharge or fast reserve through load shedding under 30 seconds [15].

D. On-Site Generation, Energy Storage and Microgrid Architecture

The next logical step for grid-compliance is self-generation and storage deployment. Many hyperscale campuses now include on-site gas turbines, fuel cells or sizable BESS capable of powering the campus for hours. This converts the data-centre into a **microgrid**, able to operate in island mode and resynchronise seamlessly with the main grid.

Key design principles include:

- **Synchronous control**: On-site generator or inverter systems must meet IEEE 1547-2020 and local interconnection standards, aligning voltage phase and frequency before reconnection.
- **Fast transfer transitions**: Automatic transfer switches (ATS) or static switches must manage sub-cycle transfers between grid-connected and island modes.
- **Storage dispatch coordination**: Batteries evolve from backup roles to **fast frequency responders**, injecting or absorbing power within 100–200 ms.
- **Hierarchical EPM orchestration**: The EPM coordinates generation, storage and load such that during grid disturbance the facility behaves as a coherent **virtual power plant (VPP)**.

Over time, aggregated microgrids of multiple data-centres may collectively provide frequency containment, spinning reserve and even black-start support – turning a cluster of data-centres into a **stabilising force** rather than a liability [16]–[17].

E. Flexible Compute Workloads as Grid Assets

Unlike traditional industrial loads, data-centres have a unique lever: **compute flexibility**. They can shape their electrical demand through workload scheduling. By aligning IT demand with grid conditions they can act as **virtual energy storage**:

- **Load-shifting**: Non-time-critical batch jobs (model training, rendering) can execute when renewable generation is abundant or electricity prices are low.
- **Rapid curtailment/migration**: Workloads may be paused or shifted across geographic regions in response to grid-curtailment signals.
- **Carbon-aware scheduling**: Compute intensity can align with renewable-output windows, reducing grid strain and carbon emissions concurrently.

To enable this, the EPM must integrate with workload orchestration systems (e.g., Kubernetes, Slurm) and link compute job allocation with available power. Electrical, cooling and IT subsystems must tolerate dynamic power modulation without compromising service-level agreements. This approach of **"demand-as-control"** is a largely untapped asset for balancing renewables-heavy grids [18]–[19].

F. Power Quality, Telemetry and Market Participation

Finally, the foundation of compliance is **visibility and power-quality control**. Generators have long met standards of harmonic distortion (<5% THD), power-factor limits (0.98 lag to 0.98 lead), and reactive-power control. Large data-centres are now required to meet equivalent performance. That means:

- **High-fidelity metering and telemetry**, streaming real-time data (1–4 second intervals) on active/reactive power, voltage, frequency to the TSO.

- **Dynamic power-factor control** via UPS and inverter firmware.
- **Harmonic filtering** to maintain THD within IEEE 519 or EN 61000 limits.
- **Cyber-secure SCADA and communications**, segregated from business data networks.

When these data streams are integrated into market systems, data-centres can **bid flexibility**—either reducing load or exporting stored power—to participate in frequency-regulation and demand-response markets. Thus, compliance becomes a revenue opportunity rather than a cost [20].

Synthesis

Technically, all of the above leads to a single convergence: a data-centre is no longer a passive consumer but a **controllable electro-digital organism**—part computer, part power-plant. Its electrical, mechanical and IT subsystems must be engineered not only for uptime, but for **grid citizenship**: the ability to stay online, recover gracefully and actively contribute to system balance.

5. EMERGING AND UNDER-EXPLORED RESEARCH TOPICS

The transformation of large-scale computing infrastructure into grid-interactive assets introduces a frontier of questions that neither the traditional power-systems community nor the data-centre industry has yet answered. The following themes outline fertile ground for future IEEE papers, doctoral projects, and industrial R&D.

A. Ultra-Fast Load Dynamics and Stability in AI Data Centres

AI and high-performance-computing (HPC) facilities are fundamentally different from conventional enterprise loads. They operate with rapid, burst-mode GPU clusters that can swing from idle to full utilisation in milliseconds. Recent modelling suggests that a multi-gigawatt AI campus could reach ramp rates above 1,000 MW s^{-1}, a figure that dwarfs most generator or industrial-load transients [21].

At these speeds, traditional control cycles—typically one to two seconds—are far too slow. The next generation of control architectures must therefore operate on sub-millisecond time steps, possibly embedded within inverter firmware itself.

The resulting research questions are non-trivial:

• What are the maximum permissible ramp rates that maintain frequency and voltage stability under varying system inertia conditions?
• How can hierarchical controllers—ranging from micro-second power-electronic loops to second-scale supervisory controls—be coordinated without instability?
• Can AI-data-centre load profiles be formally embedded in dynamic-stability models such as ROCOF and frequency-nadir analysis, and what thresholds prevent cascading trips?

A comprehensive theoretical framework would bridge power-electronics transient theory and IT-workload scheduling, an intersection rarely explored in current literature [21], [27], [38].

B. Hierarchical Control Architecture for Data-Centre Microgrids Interacting with TSOs

Data centres are rapidly evolving into nested microgrids—each combining uninterruptible power supplies (UPS), battery systems, diesel or gas generation, and bidirectional inverters [22]. Their participation in grid-support programs requires a hierarchical control structure spanning four temporal layers:

1. Device layer (< 1 ms): inverter current loops, protection relays, and UPS control.
2. Site layer (1 ms–1 s): local energy-management and fault-ride-through logic.
3. Aggregation layer (1 s–5 min): multi-site coordination, reactive dispatch, and voltage regulation.
4. Market layer (5 min–24 h): energy trading, ancillary-service bidding, and forecasting.

Interoperability across these layers demands adherence to open standards such as IEC 61850 for real-time communication [23], IEEE 2030.5 for distributed-energy coordination [24], and IEC 62559 for use-case modelling [25]. Stability proofs must demonstrate that aggregated microgrids do not create oscillatory or voltage-droop conflicts with TSO controls [22], [37]. This represents an emerging niche for both control-theory and cybersecurity research [37].

C. Standardisation of Grid-Compliance Testing and Certification for Data Centres

Traditional generation plants undergo rigorous grid-compliance testing—covering LVRT, HVRT, frequency-ride-through, reactive-power control, and black-start verification [26]. No equivalent exists for large data centres, despite their increasing grid significance.

Future work could define a Data-Centre Grid-Compliance Test Suite, including:

• Simulated voltage-sag and frequency-dip ride-through tests.
• Controlled load-drop and staged-reconnection verification.
• Telemetry latency and remote-disconnect responsiveness benchmarks.

The result would be a pass/fail certification framework aligned with IEEE 1547 [26], IEC 61000, and national TSO grid-codes [36]. Regulators could issue a “Grid-Interactive Data Centre (GIDC) Gold/Platinum” certification, akin to LEED or ISO 50001, but focused on electrical stability and interoperability [36].

D. Business and Market Models for Data Centres as Grid Assets

Economic models have not kept pace with the technical evolution. Data centres now have the capability to provide demand-response, frequency regulation, and reserve services, yet pricing and settlement mechanisms remain underdeveloped [28]. Future research must quantify:

• The marginal market value (USD per MW-shift) of deferrable compute or storage workloads.
• Contract structures that equitably distribute value among TSOs, cloud providers, and aggregators.
• The effect of *space-time load-shifting*—executing jobs where and when renewable power is abundant—on market clearing and carbon accounting [28], [29], [30].

Preliminary field trials by hyperscale operators have already shown measurable frequency-response capability using idle UPS and BESS assets [29], [35]. Formalising these business pathways could transform a regulatory obligation into a revenue-positive grid-service model [30].

E. Waste-Heat Recovery and Grid Load Mitigation

Most discussions on data-centre energy impact focus on megawatts consumed, not megawatts avoided. In cold climates, waste-heat recovery can displace up to 30–50 MW of district-heating demand per campus [31], [32]. This effectively offsets electrical load elsewhere and improves net grid efficiency [33], [34].

Future research directions include:

• Quantifying avoided electrical load from recovered thermal energy.
• Developing dual-network models linking electricity and thermal flows to assess combined grid stability [33], [40].
• Proposing policy frameworks that reward data centres for thermal-sector decarbonisation as part of grid-integrated performance [32], [34].

Integrating these cross-vector interactions could make grid models holistic—treating data centres as **bi-directional energy hubs** rather than single-vector consumers.

Synthesis

Collectively, these topics represent a paradigm shift from “data-centre efficiency” to **“data-centre integrability.”** The research frontier now lies at the intersection of control theory, market economics, and thermal-electrical coupling [39], [40]. Each area supports the central question: **How can intelligent computing infrastructure stabilise rather than stress the energy systems that power it?**

6. RECOMMENDATIONS FOR DATA CENTRE OPERATORS & ECOSYSTEM STAKEHOLDERS

Based on the foregoing analysis, the following are recommended steps for data-centre operators, equipment vendors, system integrators (such as your team at Smart Grid Analytics) and TSOs/utilities to anticipate and act upon.

1. **Early Engagement with TSOs/Utilities** – Data-centre owners should engage grid operators **at the earliest stages** of site selection, load-growth forecasting and interconnection planning. Doing so ensures that potential obligations (remote disconnect, telemetry, ride-through, curtailment) are identified and baked into design rather than retrofitted [41].
2. **Design for Grid-Friendly Behaviour from Day-One** – From project inception, electrical designs should incorporate ride-through capability, staged reconnection logic, remote-disconnect controls and real-time telemetry (UPS, BESS, on-site generation, SCADA/EPM). The goal is for the facility to behave as an **active grid participant**, not merely a large passive load [42].
3. **Integrate EPM (Energy Process Management) Platform with IT Workload Flexibility** – Workload-control and energy-control must converge. The EPM platform should interface with IT schedulers so that non-critical compute is shifted or shed in grid events, enabling the data centre to offer grid services rather than simply consume power [43].
4. **Run Simulated Grid-Fault & Load-Shedding Tests** – Prior to commissioning, simulate key scenarios: voltage/frequency sag ride-through, full-scale remote disconnect, step-load recovery profiles, remote curtailment events. Capture performance data, validate behaviours and document results for regulatory or TSO review [44].
5. **Monitor and Report Load Characteristics Continuously** – Operate with full transparency: collect telemetry on load ramp-rates, power-quality metrics (harmonics, power-factor), time-to-shed when signalled, time-to-reconnect, ramp-back profiles. Such data supports compliance, settlement and builds TSO confidence [45].
6. **Explore Ancillary Service Participation & Business Models** – Consider monetising flexibility: offer demand-response or load-flex services, shift compute loads or run on-site generation during high-price periods. Treat grid-support as a revenue stream rather than purely obligation [46].
7. **Stay Ahead of Standardisation Trends** – Actively engage with working groups at Institute of Electrical and Electronics Engineers (IEEE), International Electrotechnical Commission (IEC) or industry consortia that develop data-centre grid-compliance standards. Early alignment offers competitive advantage and ensures future regulatory alignment [6], [47].

VI. CONCLUSION

The growth of Data Centres, especially those supporting AI and HPC workloads—presents both a significant challenge and a major opportunity for power systems. Rather than passive consumption, data centres are increasingly being called upon to become **active grid-participants**: behaving like virtual power plants, providing stability services and being controllable by system operators. The line between load and generation is becoming blurred.

To succeed in this transformed landscape, data-centre operators must shift mentality: build electrical and control systems with **grid-support capabilities**, integrate workload flexibility and qualify as “good grid citizens.” At the same time, this evolution unlocks value: data centres can monetise flexibility, provide ancillary services and improve resilience.

From a research standpoint, the most fertile ground lies in ultra-fast load dynamics of AI data centres, standardising grid-compliance tests for data centres and developing business/market models for compute-load flexibility as a grid asset. For companies operating at the interface of digital infrastructure and power systems (such as Smart Grid Analytics), this is a golden opportunity to lead—architect future-ready systems, define standards and enable data centres to meaningfully contribute to grid reliability while sustaining their primary mission.

REFERENCES

[1] McKinsey & Company, "Data Centres and AI: How the energy sector can meet power demand," 2024.

[2] Electric Reliability Council of Texas (ERCOT), "Large load interconnection process," 2025.

[3] Alberta Electric System Operator (AESO), "Interim approach to large-load connections," 2025.

[4] Fingrid Oyj, "KJV2026 draft grid-code for demand connections," 2025.

[5] Electric Reliability Council of Texas, "Senate Bill 6 — large load reforms and interconnection requirements," 2025.

[6] Legal Energy Alert – Weil Gotshal & Manges LLP, "Texas SB 6 reforms interconnection and colocation rules for data centres and large loads," July 2025.

[7] ERCOT Board of Directors, "Approval of large-load interconnection process (NPRR-1234 / PGRR-115)," 2025.

[8] Alberta Electric System Operator, "Interim approach to large-load connections: data centre applications," June 2025.

[9] Bennett Jones LLP, "Large load integration on Alberta's electricity grid: the new framework for data centres," 2025.

[10] Commission for Regulation of Utilities (Ireland), "New electricity connection policy for data-centre developers," February 2025.

[11] Fingrid Oyj, "KJV2026 draft: grid-code specifications for demand connections," June 2025.

[12] European Network of Transmission System Operators for Electricity (ENTSO-E), "Implementation progress report for the Demand Connection Code," 2024.

[8] Uptime Institute. (2025). Grid reliability and data-centre ride-through failures. New York, NY.

[9] Lawrence Berkeley National Laboratory. (2023). Modeling data-centre load flexibility for grid services. Berkeley, CA.

[10] European Commission. (2024). Data centres energy efficiency and grid integration policy paper. Brussels, Belgium.

[11] International Renewable Energy Agency (IRENA). (2025). *The role of digital infrastructure in grid stability.* Abu Dhabi, UAE.

[12] IEEE Standard 519-2014. (2014). IEEE recommended practice and requirements for harmonic control in electrical power systems. IEEE, New York, NY.

[13] IEEE Standard 1547-2020. (2020). Standard for interconnection and interoperability of distributed energy resources with associated electric power systems interfaces. IEEE, New York, NY.

[14] Milanović, J. V., & Zamora, R. (2023). Hierarchical control and stability assessment for inverter-dominated grids. *IEEE Transactions on Power Systems*, 38(4), 2921-2935. https://doi.org/10.1109/TPWRS.2023.3249821

[15] Google. (2025, August 4). How we're making data centres more flexible to benefit power grids. Google Blog.

[16] Microsoft Research. (2024). Virtual Battery: Data centres as grid-interactive flexible loads.

[17] IEEE Power & Energy Society. (2024). Emerging trends in grid-code harmonization for large loads.

[18] Zhang, W., & Zavala, V. M. (2021). Remunerating space-time load-shifting flexibility from data centres in electricity markets. *arXiv*. https://doi.org/10.48550/arXiv.2105.11416

[19] ScienceDirect. (2025). Data centres as a source of flexibility for power systems. *Energy Reports*, 12, 1452-1467. https://doi.org/10.1016/j.egyr.2025.01.002
[20] Enel X. (2023). Supporting electricity grids with data centres. Rome, Italy.
[21] A. Peivandizadeh, "A theoretical framework for virtual power plant integration with gigawatt-scale AI Data Centres: Multi-timescale control and stability analysis," *arXiv preprint*, Jun. 2025. [Online]. Available: https://arxiv.org/abs/2506.17284
[22] J. V. Milanović and R. Zamora, "Hierarchical control and stability assessment for inverter-dominated grids," *IEEE Trans. Power Syst.*, vol. 38, no. 4, pp. 2921–2935, 2023.
[23] International Electrotechnical Commission, IEC 61850 – Communication networks and systems for power utility automation, Geneva, Switzerland, 2019.
[24] IEEE Standard 2030.5-2021, IEEE Standard for Smart Energy Profile Application Protocol, IEEE, New York, NY, 2021.
[25] IEC 62559-2:2015, Use-Case Methodology – Part 2: Definition of the Use-Case Template and Guidelines for Use-Case Development, IEC, Geneva, Switzerland, 2015.
[26] IEEE Standard 1547-2020, Standard for Interconnection and Interoperability of Distributed Energy Resources with Associated Electric Power Systems Interfaces, IEEE, New York, NY, 2020.
[27] Z. Wang, T. Hong, and N. Lu, "Aggregated modeling and control of large load clusters in power systems: A review," *IEEE Trans. Smart Grid*, vol. 14, no. 2, pp. 1659–1676, 2023.
[28] W. Zhang and V. M. Zavala, "Remunerating space-time load-shifting flexibility from Data Centres in electricity markets," *arXiv preprint*, May 2021. [Online]. Available: https://doi.org/10.48550/arXiv.2105.11416
[29] Google LLC, "How we're making Data Centres more flexible to benefit power grids," *Google Blog*, Aug. 2025. [Online]. Available: https://blog.google/inside-google/infrastructure/how-were-making-data-centers-more-flexible-to-benefit-power-grids/
[30] McKinsey & Company, "Data Centres and AI: How the energy sector can meet power demand," *McKinsey Insights*, Sept. 2024.
[31] "Data centres as a source of flexibility for power systems," *Energy Reports*, vol. 12, pp. 1452–1467, 2025. [Online]. Available: https://doi.org/10.1016/j.egyr.2025.01.002
[32] ScienceDirect Editorial Board, "Waste-heat recovery and thermal grid interactions in hyperscale data centres," *Energy Conversion and Management*, vol. 308, 2025.
[33] M. García and P. Rodríguez, "Thermal-electrical co-simulation of smart district networks with data centre integration," *IEEE Access*, vol. 12, pp. 110 223–110 239, 2024.
[34] International Energy Agency (IEA), Data Centres and Energy Efficiency – Tracking Clean Energy Progress, Paris, France, 2025.
[35] Enel X Global Services, "Supporting electricity grids with data centres," Rome, Italy, Dec. 2023.
[36] European Network of Transmission System Operators (ENTSO-E), "Towards harmonised compliance testing for demand facilities," *Technical Paper*, Brussels, 2025.
[37] R. Irvin and L. Vokkarane, "Cyber-secure communication and data integrity in hierarchical microgrid control," *IEEE Trans. Ind. Informat.*, vol. 19, no. 3, pp. 2904–2916, 2023.
[38] Y. Xu and S. H. Low, "Dynamic models for frequency regulation with controllable data-center loads," *IEEE Trans. Smart Grid*, vol. 15, no. 1, pp. 501–514, 2024.

[39] J. Wang, A. Srivastava, and N. Lu, “Modeling and control of flexible computing loads in integrated energy systems,” *IEEE PES GM Conference Proc.*, 2024.
[40] T. Huang, “Integrated thermal-electrical modeling for waste-heat utilisation in data centres,” *Renewable and Sustainable Energy Reviews*, vol. 194, 2025.
[41] U.S. Department of Energy, *Recommendations on Powering Artificial Intelligence and Data Centre Infrastructure*, July 2024. [Online]. Available: https://www.energy.gov/sites/default/files/2024-08/Powering%20AI%20and%20Data%20Center%20Infrastructure%20Recommendations%20July%202024.pdf
[42] K. Watson, “Data Centres – A good grid citizen,” Eaton White Paper, Jul. 2025.
[43] American Council for an Energy-Efficient Economy (ACEEE), Turning Data Centres into Grid and Regional Assets: Considerations and Recommendations, Oct. 2024.
[44] IEEE Industry Connections, “Data Centres: Standards Needs Analysis and Recommendations,” IEEE Std. 2870-2024, 2024.
[45] W. Wong, “Data Centre Regulation Trends to Watch in 2025,” *Data Centre Knowledge*, Nov. 2024.
[46] L. Lin and A. A. Chien, “Adapting Datacenter Capacity for Greener Datacenters and Grid,” *arXiv preprint*, Jan. 2023.
[47] “Data Centre Power Grid Demands Explained: Interconnection Bottlenecks, Large-Load Rules and Grid Stability,” *CAELED Blog*, Apr. 2025.