

Sub-Cycle Grid Defense: Sub-Cycle Micro-Shedding, Inverter Coordination, and Millisecond Frequency Regulation in High-Density AI Campuses

The grid cannot sustain AI as a shock load. Computing must become the damper.

Author: Brennan DeCrow
Affiliation: ATESO Labs / ManyMoats Research
Parent Research Authority: *ATESO & The Magma Runtime: The Thermodynamic Obsolescence of Document-Centric Execution for Continuous Spatial Compute* (DeCrow, 2026; USPTO Provisional App #64/159,586, Claims 5, 14, 15, 19) [1]
Status: M-Tier Master White Paper. Gold Master Candidate (post-Stage 5 synthesis, 22 September 2026; supersedes v3.2)
Review lineage: DeepSeek (swing equation, VRM, RoCoF) → Muse (regulatory framing, three-phase architecture) → GLM (small-signal stability, ERCOT setpoints) → Kimi/V4 (electrical and tariff arithmetic) → Claude Opus 5.5 (synthesis & certification)
Next stage: Headless Fable & Astra final stamp

Evidence notation (shared across the White Paper series). **F** formal result under stated assumptions · **R** reproduced from an executable artifact, including simulation · **U** author-reported, raw data not in the package · **V** independently checked against a grid operator, standards body, vendor or published third-party source · **P** proposed design or experiment not yet run. Appendix A assigns a mark to every load-bearing number.

STRIP	THE EXECUTIVE		
	TRIGGER	FULL SHED	PRODUCT
FLOOR	59.85 Hz	45 MW in 1 cycle	ERCOT FFR
59.30 Hz	ERCOT FFR setpoint	35–52 ms end to end	≤ 15 cycles allowed
first UFLS stage	(5% of load)	5–7× inside ERCOT	15-min compute hold

Abstract

AI training campuses swing tens of megawatts in lockstep as thousands of accelerators alternate between compute and communication [15]. The grid has noticed. NERC documented 1,500 MW of data-center load dropping off the Eastern Interconnection in a single event [14]. Texas now requires loads of 75 MW or more to accept disconnection in emergencies [13]. ERCOT is tracking about 410 GW of large loads seeking interconnection, some 87% of them data centers [6].

ATESO turns that liability into a resource. Its Compute-As-Virtual-Inverter (CAVI) sheds deferrable compute within one AC cycle of a trigger: 45 MW at 2.70 MW/ms, inside a 3.0 MW/ms bus cap. It then holds the curtailment for fifteen minutes entirely within the compute fabric. It adds no battery, cycles no cells, and leaves the facility UPS on float.

On a 1 GVA islanded system with $H = 3$ s, a 100 MW step drives the rate of change of frequency to 1.00 Hz/s. There, a frequency-response model shows the shed lifting the nadir from 59.20 Hz, below ERCOT's first under-frequency load-shedding stage at 59.30 Hz, to 59.54 Hz. On ERCOT's bulk system, the same 45 MW offsets 3.3% of the loss of a 1,375 MW unit.

End to end, the response completes in 35–52 ms, 5–7× inside ERCOT's 15-cycle requirement for Fast Frequency Response.

We also report where the argument stops: - **Magnitude sets the nadir more than speed does.** Delivering the same 45 MW at fifteen cycles instead of one costs only 28 mHz here. - **Equivalent inertia depends on the window** over which it is measured. - **Sub-synchronous stability requires an impedance scan** against the real network.

The Grid Anchor

Electrical & grid metric	Conventional AI hyperscale campus	ATESO Compute-As-Virtual-Inverter (CAVI)	Grid & financial impact
Transient power ramp	50–150 MW in < 100 ms ($di/dt > 1.5$ kA/s at 138 kV)	≤ 3.0 MW/ms bus cap 45 MW shed at 2.70 MW/ms	Bounded inductive transient $L \cdot di/dt = 57.1$ V at 13.8 kV (0.060% of BIL)
Frequency response latency	None (volatile load)	One-cycle actuation (16.7 ms) 35–52 ms end to end	ERCOT FFR-capable 5–7× inside the 15-cycle requirement; 15-minute hold
Interconnection pressure	ERCOT large-load queue ≈ 410 GW	Designed to support NERC	Strengthens the interconnection

Electrical & grid metric	Conventional AI hyperscale campus	ATESO Compute-As-Virtual-Inverter (CAVI)	Grid & financial impact
	($\approx$ 87% data centers, Mar 2026)	BAL-003-2 and ERCOT FFR; complements Texas SB 6	case a controllable load, not a shock load
Equivalent inertia	$H = 0$ s (no rotating mass)	$H_E = 4.20$ s on 100 MVA (420 MW·s, first 155.5 ms, 0.5 Hz basis)	$\approx$ a 105 MVA unit at $H = 4$ s (energy) $\approx$ a 250 MW unit's 18% governor headroom (power)
Incremental storage CAPEX	$\approx$ \$85M (dedicated 100 MW / 200 MWh BESS)	\$0 (hold is compute-internal)	$\approx$ \$85M avoided
Electrochemical wear	Cycle-limited (every event ages cells)	Zero cell cycling (UPS stays on float)	No new NFPA 855 installation

*ATESO CAVI is designed to support Balancing Authority frequency response under NERC BAL-003-2 [2] and to qualify as an ERCOT Load Resource providing Fast Frequency Response [4, 7]. Its continuous droop tier maps onto ERCOT's Controllable Load Resource primary frequency response. In PJM it maps onto Regulation and Synchronized Reserve (§4.8). IEEE Std 2800-2022 [8] and IEEE Std 1547-2018 [9] do not govern loads. ATESO adopts their active-power–frequency provisions by analogy, as a voluntary standard of care.

1. The AI Grid Shock: The Physics of di/dt and RoCoF

1.1 Why the grid now treats AI as a shock load

- **Synchronized power swings.** In large training jobs, every accelerator alternates between a compute-heavy phase and a communication-heavy phase in lockstep. Microsoft, OpenAI and NVIDIA report that the resulting power swings grow with job scale. If they align with critical utility frequencies, they can damage grid equipment [15].
- **Synchronized load loss.** On 10 July 2024 a single 230 kV fault led about 1,500 MW of data-center load in the Eastern Interconnection to disconnect through its own ride-through protection. 1,260 MW stayed off for hours [14].
- **Regulation has followed.** Texas SB 6, signed 20 June 2025, requires loads of 75 MW or more to install shutoff equipment and accept disconnection during firm load shed [13].
- **Queues have exploded.** ERCOT is tracking about 410 GW of large loads seeking interconnection, roughly 87% of them data centers [6]. Generation and storage queues hold 2,061 GW nationally, 408 GW of it in ERCOT, and the

median project takes more than five years from request to operation [16]. That span was already nearly five years for projects completed in 2023 [17].


The binding constraint is not energy alone. It is the rate of change of power: the di/dt that a campus imposes on the network.

1.2 The synchronous step load

Collective operations such as AllReduce and tensor-parallel barriers move processors from memory stalls to full-throttle matrix multiplication together.

- **The instantaneous step.** A 140 MW AI campus can step from 40 MW to 140 MW in under 50 ms.
- **The current slew.** On a 138 kV bus, +100 MW is 418 A of line current. Delivered in 50 ms, that is 8,370 A/s. The Grid Anchor's "> 1.5 kA/s" is a floor that holds across the full 50–150 MW range.

[Fig. 1 – The high-density AI transient]



1.3 Rate of change of frequency, and the two grids this paper addresses

Frequency obeys the per-unit swing equation [19]:

$$2H_{\text{sys}} \frac{df}{dt} = P_{\text{mech}} - P_{\text{elec}} \quad [\text{p.u.}] \quad \frac{2H_{\text{sys}}}{S_n} \frac{df}{dt} = \frac{P_{\text{mech}}}{S_n} - \frac{P_{\text{elec}}}{S_n} \quad [\text{p.u.}]$$

Here H_{sys} is the system inertia constant (s), $f_0 = 60$ Hz and S_n the system base (MVA). The initial rate of change of frequency (RoCoF) for a power imbalance ΔP is therefore:

$$\text{RoCoF} = \left| \frac{df}{dt} \right| = \frac{f_0 \Delta P}{2H_{\text{sys}} S_n} \quad [\text{Hz/s}]$$

Case A: the islanded system. An isolated sub-transmission island or industrial microgrid with $S_n = 1,000$ MVA and $H_{\text{sys}} = 3.0$ s:

$$\text{RoCoF} = \frac{60 \text{ Hz} \times 1,000 \text{ MVA} \times 100 \text{ MW}}{2 \times 3.0 \text{ s} \times 1,000 \text{ MVA}} = 1.00 \text{ Hz/s}$$

Steps of 150–250 MW on the same island reach 1.5–2.5 Hz/s. The case applies to any system that runs, or can become, electrically separate, such as a transmission pocket after it islands from the interconnection. While a pocket remains synchronously connected, its frequency is the interconnection's frequency.

Case B: the ERCOT bulk system. ERCOT sets its critical inertia at 100 GW·s; the lowest it recorded through 2017 was 130 GW·s [5]. Stored kinetic energy replaces $H_{\text{sys}} S_n$: - A 100 MW step then moves frequency at only 0.023–0.030 Hz/s. - ERCOT’s design contingency, the simultaneous loss of both South Texas Project units (2,750 MW), moves it at 0.64–0.83 Hz/s. - ERCOT requires that frequency not fall from 59.7 Hz to 59.3 Hz in less than 0.416 s after that loss [5].

The floor. If frequency reaches 59.30 Hz, ERCOT’s first under-frequency load-shedding (UFLS) stage sheds 5% of system load [5], as required under NERC PRC-006 [3].

On an island, a single campus can decide whether that floor is reached (§4.4). On the bulk system, a single campus is one contributor to the fleet that defends it (§4.5).

2. Inverter Coordination vs. the Limits of Chemical Buffering

Utilities increasingly ask large loads to firm themselves with dedicated battery energy storage (BESS).

BUFFER WEAR	CAPEX (100 MW) NEW FIRE-CODE FOOTPRINT	CYCLE
Chemical BESS (100 MW/200 MWh) ≈ \$85M cells NFPA 855 installation		Every event ages
ATESO CAVI compute damper float) None	\$0 incremental	None (UPS on

2.1 What a dedicated battery costs

- **Capital.** At NREL’s 2024 base-year costs (\$241/kWh of energy plus \$372/kW of power), a 100 MW / 200 MWh system costs about \$85 million before land and interconnection [18].
- **Permitting.** Stationary lithium-ion installations fall under NFPA 855 [11], with setbacks and review that add schedule risk.
- **Wear.** Every charge–discharge cycle consumes cell life, so fast-response duty is paid for in degradation.

2.2 Where batteries still win

Batteries follow bidirectional, energy-neutral signals without touching the host’s workload. They can also inject power. CAVI can only withhold it. This paper does not claim that CAVI replaces storage for every service; it claims CAVI makes storage unnecessary for fast under-frequency response.

2.3 The CAVI advantage

CAVI draws its response from work that can wait. The shed comes from deferring Class B computation, not from a battery. No cells cycle, no capital is added, and the facility UPS remains on float, reserved for utility-outage ride-through.

3. Compute-As-Virtual-Inverter (CAVI) via Granular Micro-Shedding


ATESO treats compute as discrete, resident binary granules governed by a sub-millisecond admission loop ($h \leq 1$ ms), not as opaque operating-system tasks. That structure lets the runtime modulate active power the way an inverter would: precisely, quickly and on command.

3.1 Workload partitioning and shedding capacity

- **Class A: real-time reflex compute, 25–35% of load.** Live inference tokens, spatial tracking, autonomous-agent safety loops. Class A is never throttled or deferred, so active-user inference keeps its full throughput.
- **Class B: elastic batch compaction, 65–75% of load.** Background Merkle subtree compaction, vector-index clustering, memory consolidation and batch embedding precomputation. Class B is interruptible at 64-byte granule boundaries.
- **Availability.** A 45 MW commitment requires at least 45 MW of Class B load, which means the campus must be drawing 60–69 MW. At a 100 MW operating point, Class B is 65–75 MW. The campus bids FFR only when that condition holds.
- **Derating.** The partition applies to training clusters and batch-augmented inference campuses. Dedicated low-latency conversational campuses derate to 20–30 MW.

3.2 Campus topology

[Fig. 2 – Campus electrical topology and the damping boundary]



└ ... 100,000 sockets ───> desynchronized ramp: 2.70 MW/ms (cap 3.0 MW/ms)

3.3 Socket-level current

Spread across 100,000 accelerator sockets, a 45 MW shed is 450 W per socket: - On the 800 V bus, that is 0.5625 A per socket. - At the socket's core rail, near 0.8 V, it is about 562 A. Ramped over one cycle, that is 0.034 A/μs.

The same voltage regulators already follow compute-to-communication transitions every training iteration [15]. A ramp that spreads the change across a full AC cycle adds no new stress to the voltage regulators.

3.4 Transformer inductive transient (IEEE C57.12.00)

Consider the worst case: the entire 45 MW shed passes through one 100 MVA, 13.8 kV transformer with leakage reactance 0.10 p.u.

$$\Delta I = 45 \times 10^6 \text{ W} / 13,800 \text{ V} = 1,882.7 \text{ A} \quad \frac{dI}{dt} = 1,882.7 \text{ A} / (1/60) \text{ s} = 1,882.7 \text{ A} \cdot 0.01667 \text{ s} = 112,960 \text{ A/s} = 113.0 \text{ kA/s}$$

$$\Delta I = \frac{45 \times 10^6 \text{ W}}{\sqrt{3} \times 13,800 \text{ V}} = 1,882.7 \text{ A} \quad \frac{dI}{dt} = \frac{1,882.7 \text{ A}}{(1/60) \text{ s}} = \frac{1,882.7 \text{ A}}{0.01667 \text{ s}} = 112,960 \text{ A/s} = 113.0 \text{ kA/s}$$

$$Z_{\text{base}} = (13,800)^2 / 100 \times 10^6 = 1.9044 \Omega \quad X_{\text{leak}} = 0.19044 \Omega \quad L_{\text{leak}} = 0.19044 \times 60 = 0.505 \text{ mH}$$

$$Z_{\text{base}} = \frac{(13,800)^2}{100 \times 10^6} = 1.9044 \Omega \quad X_{\text{leak}} = 0.19044 \Omega \quad L_{\text{leak}} = \frac{0.19044}{2\pi \times 60} = 0.505 \text{ mH}$$

$$V_L = L_{\text{leak}} \frac{dI}{dt} = 0.50516 \text{ mH} \times 112,960 \text{ A/s} = 57.1 \text{ V} \quad V_L = L_{\text{leak}} \frac{dI}{dt} = 0.50516 \text{ mH} \times 112,960 \text{ A/s} = 57.1 \text{ V}$$

- **Against phase voltage.** Phase-to-neutral voltage is 7,967 V, so 57.1 V is 0.72% of it.
- **Against insulation.** The 95 kV basic lightning impulse insulation level (BIL) [10] is conservative for the 15 kV class. The transient is 0.060% of it.
- **Steady state.** The leakage-reactance drop changes by $X \cdot \Delta I = 358.5 \text{ V}$, 4.5% of phase voltage. For a unity-power-factor load the drop is in quadrature, so the voltage magnitude moves only about 0.1%. Tap changers correct it in normal operation.

3.5 Deterministic desynchronized ramping

ATESO enforces a facility-bus ramp cap of $(dP/dt)_{\text{bus}} \leq 3.0 \text{ MW/ms}$. Socket release times are staggered deterministically across one AC cycle.

The 45 MW shed therefore arrives at 2.70 MW/ms, 10% inside the cap. A ramp spread over a full cycle limits high-frequency content in the step. Confirming that no network resonance is excited remains a site-specific frequency scan (P).

4. Frequency Regulation, Stability and Tariff Product Mapping

4.1 Dual-tier control architecture

- **Tier 1: sub-cycle open-loop trigger.**
 1. A frequency estimator delivers a reading within one to two cycles.
 2. A digital comparator at 59.85 Hz issues the trip within 2 ms.
 3. Direct DVFS throttling of Class B cores completes the 45 MW shed within one AC cycle (16.7 ms).

End to end, full response lands 35–52 ms after frequency crosses 59.85 Hz, which is 5–7× inside ERCOT’s 15-cycle (250 ms) requirement.
- **Tier 2: continuous closed-loop regulation.** The transfer function $G_{CAVI}(s)$ (§4.7) governs droop regulation around the curtailed operating point, with a 20 ms measurement filter. The filter belongs to Tier 2 alone, so it never delays the Tier 1 trip.

4.2 Setpoints

Setpoint	Value	Basis
Tier 2 deadband	59.95–60.05 Hz, no modulation	Design choice (P)
FFR trigger	59.85 Hz	ERCOT FFR rule [4]
Full deployment	35–52 ms after trigger; ERCOT allows 15 cycles (250 ms)	§4.1 (F), [4]
Margin to the first UFLS stage	Trigger sits 0.55 Hz above 59.30 Hz	[5]
For contrast: ERCOT under-frequency-relay Load Resources	59.70 Hz, 30 cycles	[4]

4.3 Product mapping: the fifteen-minute compute-internal hold

ERCOT’s Fast Frequency Response requires full deployment within 15 cycles of a 59.85 Hz trigger. The resource must sustain the response for up to 15 minutes once deployed and restore within 15 minutes of recall [4]. ATESO meets the rule in three phases:

Phase	Window	Product	Mechanism
1. Ingress	Trip → one cycle (35–52 ms after 59.85 Hz, end to end)	ERCOT FFR deployment	Tier 1 trip; DVFS throttle of Class B cores; 45 MW at 2.70 MW/ms
2. Hold	To 15 minutes	ERCOT FFR sustainment	Class B deferral inside the

Phase	Window	Product	Mechanism
3. Restore	After ERCOT recall	Handoff to primary response and ECRS	compute fabric; UPS on float Class B resumes over minutes; ready to redeploy within 15 minutes


1. **Ingress.** The shed lands within one cycle of the trip, before governors have delivered meaningful response.
2. **Hold.** The 45 MW curtailment is sustained for the full fifteen minutes by deferring Class B work. No battery supplies it, and the UPS never leaves float, so the hold incurs zero cell cycling.
3. **Restore.** Deferred work resumes at a grid-friendly pace agreed with the operator, for example 45 MW over five minutes. The 3.0 MW/ms cap governs electrical transients, not restoration. Re-adding 45 MW in a few milliseconds would itself be the shock load this paper sets out to remove. No work is dropped; deferred work completes later.

4.4 Modeled frequency arrest: the islanded system

We model Case A with the Anderson–Mirheydar system-frequency-response structure [20]: - $H = 3.0$ s, load damping $D = 1$, 5% droop; - a reheat turbine ($F_H = 0.3$, $T_R = 8$ s); - a governor lag and ERCOT's ± 36 mHz governor deadband.

A 100 MW step strikes at $t = 0$. CAVI responds per §4.1.

[Fig. 3 – 100 MW step on the 1 GVA island: modeled frequency (governor lag 0.3 s, ± 36 mHz)]



Governor model	Nadir, unmanaged	Nadir, with 45 MW CAVI
No lag, no deadband	59.32 Hz	59.62 Hz
0.2 s lag	59.27 Hz	59.59 Hz
0.2 s lag, ± 36 mHz deadband	59.24 Hz	59.56 Hz

Governor model	Nadir, unmanaged	Nadir, with 45 MW CAVI
0.3 s lag, ± 36 mHz deadband (Fig. 3)	59.20 Hz	59.54 Hz

Speed versus magnitude. The nadir arrives about 1.5 s after the step, so the 45 MW magnitude does most of the work. For the Fig. 3 case, the table shows the nadir against the delay before the shed begins, counted from the 59.85 Hz crossing:

Shed begins	Nadir
0 cycles after crossing	59.541 Hz
3 cycles	59.539 Hz
15 cycles (ERCOT's allowance)	59.513 Hz
30 cycles	59.436 Hz
60 cycles	59.252 Hz

Sub-cycle speed buys 28 mHz of margin over a response at ERCOT's allowance. The one-cycle actuation is margin; the mechanism is magnitude delivered before the governors. Speed matters most where inertia is lowest and the nadir arrives soonest.

4.5 The bulk-system case: ERCOT

Take a low-inertia condition, $E_k = 135 \text{ GW}\cdot\text{s}$, and the loss of one South Texas Project unit (1,375 MW, half of ERCOT's 2,750 MW design contingency [5]):

- **Rate of fall.** Initial RoCoF is 0.306 Hz/s.
- **Timing.**
 - Frequency leaves the ± 36 mHz governor deadband after about 118 ms, so governors are already moving.
 - It reaches the 59.85 Hz trigger after about 0.49 s.
 - CAVI's 45 MW lands 35–52 ms after that.
- **Contribution.** Once landed, the shed offsets 3.3% of the lost generation. It lowers the rate of fall from 0.306 to 0.296 Hz/s while governors are still ramping.
- **Scale.** Response scales linearly with the fleet. Ten such campuses supply 450 MW of FFR, a third of the unit.

On the bulk system, one campus is a contributor, not an arrester. Its value is as a procured, fast, non-degrading reserve.

4.6 Arrest contribution and equivalent inertia

The arrest contribution. A shed ΔP sustained over a window Δt reduces the frequency decline accumulated over that window by:

$$\Delta f_{\text{arrest}} = f_0 \Delta P_{\text{shed}} \Delta t^2 H_{\text{sys}} S_{\text{sys}} = 60 \times 45 \text{ MW} \times 0.1555 \text{ s}^2 \times 3.0 \text{ s} \times 1,000 \text{ MVA} = 0.070 \text{ Hz}$$

$$\Delta f_{\text{arrest}} = \frac{f_0 \Delta P_{\text{shed}}}{2 H_{\text{sys}} S_{\text{sys}}} \Delta t^2 = \frac{60 \times 45 \text{ MW}}{0.1555 \text{ s}^2 \times 3.0 \text{ s} \times 1,000 \text{ MVA}} \Delta t^2 = 0.070 \text{ Hz}$$

This is the 1 GVA island case over a 155.5 ms window.

The struck metric. Earlier drafts reported a “local arrest” inertia of 30.0 s. That figure is withdrawn. Substituting Δf_{arrest} back into the inertia formula returns $H_{\text{sys}} \cdot S_{\text{sys}} / S_{\text{campus}} = 3.0 \times 1,000/100 = 30.0$ s identically. It is the grid’s own stored energy restated on the campus base and says nothing about the campus.

The energy basis. Over the 155.5 ms window, the shed delivers $45 \text{ MW} \times 0.1555 \text{ s} = 7.0 \text{ MW}\cdot\text{s}$. A synchronous machine releases the same energy while frequency falls 0.5 Hz if it stores:

$$E_{\text{equiv}} = f_0 \Delta P_{\text{shed}} \Delta t \Delta f = 60 \times 45 \times 0.1555 \times 0.5 = 420 \text{ MW}\cdot\text{s} \Rightarrow H_E = 4.20 \text{ s on a 100 MVA base}$$

$$E_{\text{equiv}} = \frac{f_0 \Delta P_{\text{shed}} \Delta t \Delta f}{\Delta f} = \frac{60 \times 45 \times 0.1555 \times 0.5}{0.5} = 420 \text{ MW}\cdot\text{s} \Rightarrow H_E = 4.20 \text{ s on a 100 MVA base}$$

That is the kinetic energy of a 105 MVA combustion-turbine generator with $H = 4.0$ s. It is 42% of a 250 MW unit with the same H .

The power basis. The shed’s 45 MW equals the full governor response of a 250 MW unit holding 18% headroom ($250 \times 0.18 = 45 \text{ MW}$). Under a 5% droop, that unit delivers 45 MW only after frequency has fallen 0.54 Hz, and over seconds. CAVI delivers it one cycle after the trip.

The caveat. Equivalent inertia depends on the window chosen. A synchronous machine releases energy only while frequency falls; a sustained shed keeps delivering. Measured over one second instead of 155.5 ms, the same shed corresponds to $H_E = 27.0$ s on the same base. We therefore report CAVI as Fast Frequency Response, with inertia equivalents given only alongside their window.

4.7 Small-signal stability of the Tier 2 loop

The linearized Tier 2 controller, with gains on the campus base, is:

$$G_{\text{CAVI}}(s) = \Delta P_{\text{CAVI}}(s) \Delta f(s) = K_{\text{droop}} + s K_{\text{inertial}} \frac{1}{1 + s \tau_f}, \tau_f = 20 \text{ ms}, K_{\text{droop}} = 25 \text{ p.u.}, K_{\text{inertial}} = 2.0 \text{ p.u.}$$

$$G_{\text{CAVI}}(s) = \frac{\Delta P_{\text{CAVI}}(s)}{\Delta f(s)} = \frac{K_{\text{droop}} + s K_{\text{inertial}}}{1 + s \tau_f}, \tau_f = 20 \text{ ms}, K_{\text{droop}} = 25 \text{ p.u.}, K_{\text{inertial}} = 2.0 \text{ p.u.}$$

Design specification (P). Where the loop gain crosses unity, as it can in a weak island, Tier 2 is tuned for: - a gain crossover near $\omega_c = 18.2 \text{ rad/s}$ (2.9 Hz); - a phase margin of 51.4° , against a requirement of more than 30° ; - a damping ratio of $\zeta = 0.49$ (0.495 to three places), against a requirement of more than 0.20.

For a second-order-equivalent loop, the damping ratio follows from the phase margin ϕ_m exactly:

$$\zeta = \frac{1}{2} \tan \phi_m \cos \phi_m = \frac{\sin \phi_m}{2 \cos \phi_m} \Rightarrow \zeta = \frac{1}{2} \tan \phi_m \cos \phi_m$$

$$\zeta = \frac{1}{2} \tan \phi_m \cos \phi_m = \frac{\sin \phi_m}{2 \cos \phi_m} \Rightarrow \zeta = \frac{1}{2} \tan \phi_m \cos \phi_m$$

Verification on the reference island (R). We closed the loop around the §4.4 plant, converting the campus gains to the 1,000 MVA system base. - The loop gain never reaches unity (peak 0.34), so no crossover exists and the loop is stable with margin to spare. - Closing it moves the system's frequency mode from a damping ratio of 0.96 to fully overdamped. All closed-loop poles are real: -51.3 , -1.26 and -0.36 s⁻¹.

Sub-synchronous interaction (P). The derivative path's gain rises toward $K_{\text{inertial}}/\tau_f = 100$ p.u. It is 57 p.u. at 5 Hz and 99 p.u. at 45 Hz, and the swing-equation plant is not valid in that band. - Excluding sub-synchronous resonance therefore requires an impedance-based scan against the actual network. - We recommend a second low-pass pole of about 2 Hz on the derivative path, so that the controller's gain falls below the droop gain above a few hertz.

4.8 Products and standards

- **ERCOT.** Primary qualification is as a Load Resource providing Fast Frequency Response: trigger 59.85 Hz, full response within 15 cycles, 15-minute sustainment, restoration within 15 minutes of recall [4, 7].
 - Tier 2 droop maps onto Controllable Load Resource primary frequency response.
 - ECRS is the handoff reserve.
 - Qualification testing is required (P).
- **PJM.** RegA and RegD were merged into a single bidirectional Regulation signal on 1 October 2025. Separate Regulation Up and Regulation Down products are scheduled for October 2026 [12].
 - Tier 2 can follow the regulation signal around a partially curtailed Class B baseline without cycling any battery.
 - The Phase 2 hold fits Synchronized Reserve.
- **NERC BAL-003-2.** A Balancing Authority's frequency-response obligation is measured at the BA. CAVI is designed to support it [2].
- **Texas SB 6.** SB 6 requires shutoff equipment and emergency disconnection for loads of 75 MW or more [13]. CAVI's graded, reversible curtailment complements that last-resort disconnect.
- **IEEE 2800-2022 and IEEE 1547-2018.** IEEE 2800 governs transmission-connected inverter-based resources [8] and IEEE 1547 governs distribution-connected distributed energy resources [9]. Neither governs loads.
 - ATESO voluntarily adopts, by analogy, IEEE 2800 Clause 6 (active power–frequency response) and IEEE 1547's frequency-droop provisions as its standard of care.
 - Clause 5 of IEEE 2800 covers reactive power and voltage control, outside CAVI's active-power function.
- **NFPA 855.** NFPA 855 governs stationary battery installations [11]. CAVI adds none.

5. Scope and Limits

- **Scale.** On a large interconnection, one campus offsets a few percent of a large-unit loss. Islands, microgrids and separated pockets are where a single campus decides the outcome.

- **Speed and magnitude.** In the island model, sub-cycle speed adds 28 mHz of nadir margin over a response at ERCOT’s 15-cycle allowance. Magnitude does most of the work.
 - **Availability.** The 45 MW commitment needs 60–69 MW of campus load. Class B service levels must tolerate fifteen-minute deferrals.
 - **Measurement.** Tier 1’s end-to-end time depends on the frequency estimator and must be demonstrated in qualification testing.
 - **Models.** Island results come from a reduced-order frequency-response model and need field validation. Sub-synchronous screening requires an impedance scan.
 - **Standards.** Voluntary analogies to IEEE 2800 and IEEE 1547 are a standard of care, not compliance.
-

6. Conclusion: Turning the AI Grid Threat into a Grid Asset

The energy transition and the AI build-out cannot proceed on a collision course. If data centers remain volatile loads, utilities will keep slowing their interconnection.

The answer is not to buy \$85 million of degradable chemistry for every campus. It is to recognize that most of a campus’s load is work that can wait a quarter of an hour, and to give the grid a runtime that knows which work that is.

ATESO’s resident binary structure and sub-millisecond admission loop let a campus act one cycle after the trip: - shed 45 MW of deferrable compute; - complete ERCOT’s Fast Frequency Response 5–7× inside the time allowed; - hold for fifteen minutes without touching a battery; - keep active users at full throughput.

On an island, that response keeps frequency above the first load-shedding stage. On the bulk system, it joins a fleet of fast reserves that never wear out.

The grid cannot sustain AI as a shock load. Computing must become the damper.

Appendix A: Claim Register

Claim	Value	Class	Where
Island RoCoF, 100 MW step	1.00 Hz/s	F	§1.3
ERCOT critical / record-low inertia	100 / 130 GW·s	V	§1.3, [5]
ERCOT RoCoF, 100 MW step	0.023–0.030 Hz/s	F	§1.3
	0.64–0.83 Hz/s	F	§1.3

Claim	Value	Class	Where
ERCOT RoCoF, 2,750 MW loss			
ERCOT first UFLS stage	59.3 Hz, 5% of load	V	[5]
ERCOT FFR rules	59.85 Hz, 15 cycles, 15-min sustain, 15-min restore	V	[4]
ERCOT NOG §2.3 as the governing section	—	U (author citation)	[7]
Large-load queue	≈ 410 GW, ≈ 87% data centers (Mar 2026)	V	[6]
Generation/ storage queues	2,061 GW US; 408 GW ERCOT; > 5 years median	V	[16]
Data-center load-loss event	1,500 MW (1,260 MW sustained), 10 Jul 2024	V	[14]
AI training power swings	Qualitative, production telemetry	V	[15]
BESS 100 MW / 200 MWh capital	≈ \$85M	V	[18]
Class A / Class B shares	25–35% / 65– 75%	U	§3.1
Campus load needed for 45 MW	60–69 MW	F	§3.1
Shed ramp	2.70 MW/ms (cap 3.0)	F	§3.5
Leakage inductance	0.505 mH	F	§3.4
Inductive transient	57.1 V; 0.72% of phase V; 0.060% of BIL	F	§3.4
Socket current	0.5625 A at 800 V; ≈ 562 A at 0.8 V; 0.034 A/μs	F	§3.3
			§4.1

Claim	Value	Class	Where
Tier 1 end-to-end response	35–52 ms; 5–7× inside 250 ms	F (estimator assumption); P until tested	
Island nadirs	59.20–59.32 Hz unmanaged; 59.54–59.62 Hz with CAVI		§4.4
Speed vs magnitude	28 mHz from sub-cycle speed	R	§4.4
ERCOT bulk contribution	3.3% of a 1,375 MW loss; 0.306 → 0.296 Hz/s	F	§4.5
Arrest contribution	0.070 Hz over 155.5 ms (island)	F	§4.6
“30.0 s local arrest inertia”	Withdrawn (identity)	—	§4.6
Energy-equivalent inertia	420 MW·s; H_E = 4.20 s on 100 MVA (155.5 ms, 0.5 Hz basis)	F	§4.6
Power-basis equivalence	250 MW × 18% = 45 MW at a 0.54 Hz deviation (5% droop)	F	§4.6
Tier 2 specification	ω_c 18.2 rad/s, PM 51.4°, ζ = 0.49	P	§4.7
Tier 2 on the reference island	Loop gain ≤ 0.34; overdamped	R	§4.7
PJM regulation structure	Single signal since 1 Oct 2025; Up/Down Oct 2026	V	[12]
Texas SB 6	≥ 75 MW: shutoff equipment, emergency disconnection	V	[13]
	IBRs only; Clause 5	V	[8]

Claim	Value	Class	Where
IEEE 2800 scope and clauses	reactive/voltage; Clause 6 active power–frequency		
Compute-internal 15-minute hold	Class B deferral, UPS on float	P	§4.3

Appendix B: Frequency Model and Reproduction

- **Structure.** Anderson–Mirheydar system frequency response [20], per unit on 1,000 MVA with $f_0 = 60$ Hz.
 - Swing: $2H d\omega/dt = \Delta P_m - \Delta P_L - D\omega + \Delta P_{CAVI}$
 - Governor–turbine: $-(K_m/R)(1 + F_H T_R s) / ((1 + T_G s)(1 + T_R s))$, with a symmetric deadband on the governor input.
- **Parameters.** $H = 3.0$ s, $D = 1.0$, $R = 0.05$, $K_m = 0.95$, $F_H = 0.3$, $T_R = 8$ s, $T_G \in \{0, 0.2, 0.3\}$ s, deadband $\in \{0, \pm 36$ mHz}.
- **Disturbance.** A 100 MW step at $t = 0$.
- **CAVI.** A 45 MW shed triggered when the estimated frequency reaches 59.85 Hz:
 - Tier 1 uses a one- or two-cycle estimator and a 2 ms trip;
 - the shed ramps over one cycle.
- **Integration.** Explicit Euler, $\Delta t = 0.1$ ms, 20 s horizon.
- **Tier 2 check.** Frequency response of $G_{CAVI}(s) \times$ plant on the system base, plus the closed-loop characteristic roots.
- **Reproduce:** `python3 wp4_frequency_model.py` (package RESEARCH/STAGE5-FREQUENCY-MODEL/).

References

Parent research

1. B. DeCrow. *ATESO & The Magma Runtime: The Thermodynamic Obsolescence of Document-Centric Execution for Continuous Spatial Compute*. ManyMoats Research, 2026. USPTO Provisional App. 64/159,586 (Claims 5, 14, 15, 19).

Reliability standards and grid-operator sources

1. NERC. *BAL-003-2: Frequency Response and Frequency Bias Setting*.
2. NERC. *PRC-006: Automatic Underfrequency Load Shedding*.
3. J. Matevosyan (ERCOT). *Frequency Response and Ancillary Services in ERCOT*. NERC workshop presentation. Covers FFR (59.85 Hz, 15 cycles, 15-minute sustain, 15-minute restore) and UFR Load Resources (59.70 Hz, 30

- cycles). https://www.nerc.com/globalassets/our-work/workshops/5-3_matevosjana_pfr_ercot_frequency_response_and_ancillary_services.pdf
4. ERCOT. *Inertia: Basic Concepts and Impacts on the ERCOT Grid*. April 2018. https://www.ercot.com/files/docs/2018/04/04/Inertia_Basic_Concepts_Impacts_On_ERCOT_v0.pdf
 5. ERCOT. *Large Load update, Senate Committee on Business & Commerce*. April 2026. https://www.ercot.com/files/docs/2026/04/01/ERCOT_LargeLoad_Update_April2026_B-C_-Hearing.pdf
 6. ERCOT. *Nodal Protocols and Nodal Operating Guides: Fast Frequency Response provisions* (NPRR 863; NOGRR 187).
 7. IEEE. *Std 2800-2022: Interconnection and Interoperability of Inverter-Based Resources (IBRs) Interconnecting with Associated Transmission Electric Power Systems*. <https://standards.ieee.org/ieee/2800/10453/>
 8. IEEE. *Std 1547-2018: Interconnection and Interoperability of Distributed Energy Resources with Associated Electric Power Systems Interfaces*.
 9. IEEE. *Std C57.12.00: General Requirements for Liquid-Immersed Distribution, Power, and Regulating Transformers*.
 10. NFPA. *855: Standard for the Installation of Stationary Energy Storage Systems*.
 11. PJM. *Regulation Redesign Phase 1 FAQ*, with an independent summary by Modo Energy. <https://www.pjm.com/-/media/DotCom/markets-ops/ancillary/regulation-redesign-phase-1-faq.ashx> · <https://modoenergy.com/research/en/pjm-regulation-market-redesign-2025-battery-impacts>

Large loads, AI power and queues

1. Utility Dive. *Texas law gives grid operator power to disconnect data centers during crisis* (SB 6, signed 20 June 2025). <https://www.utilitydive.com/news/texas-law-gives-grid-operator-power-to-disconnect-data-centers-during-crisi/751587/>
2. NERC. *Incident Review: Considering Simultaneous Voltage-Sensitive Load Reductions* (event of 10 July 2024). https://www.nerc.com/pa/rrm/ea/Documents/Incident_Review_Large_Load_Loss.pdf
3. E. Choukse et al. (Microsoft, OpenAI, NVIDIA). *Power Stabilization for AI Training Datacenters*. arXiv:2508.14318, 2025. <https://arxiv.org/abs/2508.14318>
4. Lawrence Berkeley National Laboratory. *Queued Up: 2026 Edition* (data through end of 2025; published 1 July 2026). <https://emp.lbl.gov/news/backlog-power-plants-seeking-transmission-grid-connection-eased-somewhat-2025-amidst>
5. Lawrence Berkeley National Laboratory. *Queued Up: 2024 Edition*. https://emp.lbl.gov/sites/default/files/2024-04/Queued%20Up%202024%20Edition_1.pdf

Storage cost and power-system dynamics

1. NREL. *Cost Projections for Utility-Scale Battery Storage: 2025 Update*. June 2025. <https://docs.nrel.gov/docs/fy25osti/93281.pdf>
2. P. Kundur. *Power System Stability and Control*. McGraw-Hill, 1994.
3. P. M. Anderson, M. Mirheydar. A low-order system frequency response model. *IEEE Transactions on Power Systems* 5(3), 1990.