Compliance Is Not Safety

A Demonstrable Safety Argument for Energy Storage Systems and the Control Architecture That Sustains It

Eduardo Mayer Fagundes, M.Sc.
Electrical Engineer — Energy and Intelligent Systems
nMentors Engenharia
São Paulo, Brazil
eduardo.mayer@nMentors.com.br

Abstract — On 13 August 2026, a 150 MW battery storage project in England was refused planning permission on appeal despite full compliance with the applicable fire safety guidance and the use of equipment from a leading manufacturer. The inspector identified no breach of any standard; she found insufficient robust evidence that the proposed controls would be effective. This paper argues that this distinction — between compliance achieved and a demonstrable safety argument — defines the binding constraint on Brazilian storage projects contracted from the December 2026 capacity auctions onward, and that the capacity to sustain such an argument is a property of the control architecture rather than of the equipment. Three propositions organise the analysis. Physically, the allocation of a control function is determined by the time constant of the phenomenon it governs, not by the sophistication of the algorithm implementing it. Economically, market optimisation is already a mature, remotely delivered commercial service; what edge computation produces that cannot be sourced elsewhere is the operating envelope derived from the physical state of the asset. Normatively, actuation authority must reside in a component whose correctness is demonstrable independently of what the learning model has learned. The argument draws on four forensic investigations of documented incidents, on verified 2026 market signals, on the Brazilian regulatory framework approved in June 2026, and on published findings from an operating 30 MW/60 MWh transmission-connected asset in São Paulo, where the constraint on additional service stacking was identified in protection settings and control algorithms rather than in rated power or energy. The paper states explicitly what it does not demonstrate.

Index Terms — *battery energy storage systems, edge computing, functional safety, grid-forming inverters, safety argument, technical specification.*

I. Introduction

Brazil's Energy Research Company (EPE) registered 6,091 battery storage projects, totaling 296.8 GW, for the two Capacity Reserve Auctions scheduled for December 2 and 4, 2026 — a historic participation record for a tender that will contract only a fraction of that capacity. Neither capital nor projects are in short supply in the Brazilian storage market. What will separate the ventures that actually operate from those that never clear licensing, financing, or insurance is something else entirely.

On August 13, 2026, planning inspector Helen Heward dismissed, in a seventeen-page decision, Clearstone Energy's appeal for a 150 MW storage system in Axminster, England. The developer had followed National Fire Chiefs Council guidance to the letter and used equipment from a leading manufacturer. The inspector cited no breach of any standard. She found that the evidence submitted did not permit confidence that the proposed fire-safety and pollution-control measures would be effective — and since the pollution-control evidence was contingent on the fire-safety evidence, the insufficiency of the first dragged down the second.

That decision marks a boundary the Brazilian sector has yet to internalize. Regulatory compliance is a checked list. A safety argument is the demonstration that a system's behavior stays within known limits when something fails. These are different things, and only the second one sustains a license, financing, and an insurance policy.

This paper's thesis is that this capacity for demonstration is a property of the control architecture, not of the equipment — and that it rests on three propositions. The first, physical in nature: the allocation of a control function is determined by the time constant of the phenomenon it governs, not by the sophistication of the algorithm implementing it. The second, economic in nature: market optimization is already a mature, remote, and commercially available service; what edge computation produces that cannot be sourced elsewhere is the operating envelope derived from the asset's physical state. The third, normative in nature: actuation authority must reside in a component whose correctness is demonstrable independently of what the model has learned.

The paper is organized as follows. Section II addresses the Brazilian regulatory framework and its consequences. Section III examines the anatomy of failure across three layers of evidence. Section IV establishes the hierarchy of time constants. Sections V through VII describe the perception, decision, and shielding layers; Sections VIII and IX address the computational envelope and validation. Section X consolidates the reference architecture, and Section XI translates the argument into specification items. Section XII states what this analysis does not demonstrate.

II. The Brazilian Regulatory Framework and Its Consequences

On June 2, 2026, Brazil's National Electric Energy Agency (ANEEL) approved Normative Resolutions No. 1,161 and No. 1,162, published in the Official Gazette on June 24, and the Ministry of Mines and Energy issued Normative Ordinance No. 136, establishing the guidelines for the first Capacity Reserve Auction dedicated exclusively to battery storage systems, with supply beginning August 1, 2028 [1]–[3].

Resolution 1,161/2026 defines the requirements and procedures for granting authorization to standalone energy storage systems (ESS), creates the Authorization Request Registration Order, and sets a 54-month deadline for entry

into commercial operation. Resolution 1,162/2026 governs billing for plants with co-located ESS and for standalone ESS, establishing cumulative charges on both the consumption and injection portions, and brings storage systems under the agency's inspection and penalty regime.

That last point is often read as an administrative detail. It is not. By making the storage operator a regulated agent subject to penalties, the rule converts control error into regulatory liability. A controller that discharges outside its contracted window, that fails to deliver committed availability, or that violates connection limits no longer just produces a loss — it produces exposure to sanction, and consequently demands traceability of every decision.

On the technical side, the joint ONS/EPE technical note establishes grid-forming technology as the operational standard for connection to the basic grid, detailing requirements for voltage-source behavior with autonomy, dynamic active- and reactive-power response, passivity and damping of systemic oscillations, negative-sequence currents, and transient overcurrent capability [5] — requirements that parallel international grid-forming interconnection standards such as IEEE Std 2800-2022 [4]. Unlike a grid-following inverter, which synchronizes to an external reference, a grid-forming unit imposes voltage and frequency and must respond within fractions of a cycle during severe disturbances. No architecture that depends on a round trip to a remote center can meet that requirement. This is the first architectural constraint imposed by the standard, not by the designer's preference.

III. Anatomy of Failure

The predictable objection to any proposed architecture is the absence of empirical validation. The answer is that, in safety engineering, the primary source of evidence is not the prototype: it is the forensic investigation of incidents. Aviation and process industries built their safety regimes exactly this way. Demonstrating that a failure mode exists does not require a system in operation; it requires the reports. This section organizes the evidence into three layers.

A. International Forensics: Four Events

The first case sits closest to the thesis. On July 30, 2021, during commissioning of the Victorian Big Battery in Australia, a coolant leak in a Megapack caused a short circuit that escalated into thermal runaway, destroying two units. The regulator Energy Safe Victoria's finding is almost verbatim: the absence of monitoring and protection systems that would have been available had the unit not been switched to off-line service mode allowed the initial fault to go undetected. The unit operated for thirteen hours before that switch, which disabled monitoring and prevented alarms from triggering [6].

The detection capability existed. The authority to act had been removed by mode selection. A safety function that can be switched off is not a safety function — it is a feature. The supplementary technical report added a second relevant finding: the UL 9540A burn test envelope covers winds up to roughly 12 mph, while the region recorded gusts up to 35 mph on the day of the event — a gap between the test envelope and the real operating envelope.

The second case supplies the timescale. On April 19, 2019, at Arizona Public Service's McMicken facility, an internal fault in a single cell — identified at cell pair 7, module 2, rack 15 — triggered cascading thermal runaway. Roughly two hours elapsed between the first report of suspected fire, at 5:48 p.m., and the explosion, around 8:04 p.m. Smoke detectors did operate; gas venting from the cascading cells created a flammable atmosphere inside the container, and deflagration occurred when the door was opened, injuring firefighters [7].

Those two hours are the architecture's time budget. It is the window in which multimodal perception and deterministic authority would have changed the outcome. It is worth noting that the same utility had already suffered, in 2012, an incident whose root cause was traced to defective control logic.

The third case addresses the symmetrical error. In September 2021, at Moss Landing, California, the operator's own investigation concluded that the water-suppression system was armed in response to low smoke levels and, due to failed hose and piping couplings, sprayed water onto the racks; the water damaged the batteries and caused overheating, generating more smoke, which released more water and damaged additional batteries. A positive feedback loop between detector and actuator: the safety system produced the very incident it existed to prevent.

The fourth case is a positive one, and it anchors the latency argument. On December 14, 2017, before the Loy Yang A unit had even completed its shutdown, the Hornsdale Power Reserve had already injected 7.3 MW to arrest a frequency drop below 49.80 Hz. Data from the Australian operator show a response four seconds ahead of the generator then contracted to provide frequency-control services; the actual response was on the order of milliseconds — too fast for the record to capture.

Table I — Mapping between documented incidents and architectural layers.

Event	Failure mode	Layer addressed
Victorian Big Battery, 2021	Protection removed by mode switch; fault undetected for 13 h	Shielding — non-switchable authority
McMicken, 2019	Single-cell fault; ~2 h to explosion; gas accumulated without action	Perception — multimodal precursor
Moss Landing, 2021	Suppressor false positive caused the damage; positive feedback	Perception — operating point
Hornsdale, 2017	Local millisecond response, ahead of contracted resource	Decision — allocation by time constant

B. Market Signal: 2026

This case record is not ancient history. On July 22, 2026, a 1.5 MW storage system housed in four containers caught

fire in eastern Germany, requiring evacuation and prolonged cooling. On the underwriting side, a survey by insurer QBE of 49 UK fire brigades found 212 fires linked to photovoltaic systems in 2025 — up 133% from 91 in 2022 — while the number of installations grew only 52%; fire spread beyond the originating installation rose from 33 cases in 2022 to 88 in 2025. A separate study by the same insurer found a lithium-battery fire every five hours in the country in 2025, with containment requiring up to ten times more water than a conventional fire [8].

This underwriting data matters for a specific reason: claims are growing faster than the installed base. That means the per-unit failure rate is rising, not falling, as the fleet ages and installation becomes more widespread. No actuarial model based on product compliance captures this dynamic.

On the load side, the nature of the disturbance has changed. AI data centers require systems capable of riding through interruptions without any momentary degradation in power quality, and they deploy ever-larger, high-density battery banks, often inside or adjacent to buildings that house irreplaceable computing infrastructure. Training clusters do not draw power gradually: when hundreds of megawatts of processing units come online simultaneously, they create pulse loads that swing site demand within milliseconds, and generators cannot ramp fast enough to respond. Conventional lifetime testing does not capture these cycling patterns.

Finally, a signal that draws the boundary of the argument. AI-driven economic optimization of storage is already a mature commercial product: independent operators manage more than one gigawatt of batteries across multiple European markets, providing market access, balancing responsibility, and coordinated generation-and-storage trading behind a single connection point, active in primary, secondary, and tertiary reserve, voltage control, intraday, and spot markets. None of that touches the fast loop. The economic decision layer is solved, and it is remote; the operational-constraint layer is not.

C. The Brazilian Case: Operation Without a Case Record

Brazil has a track record of grid-storage operation that predates the 2026 regulatory framework. Commercial operation of the first utility-scale system on the Brazilian transmission grid began on November 28, 2022, at the Registro substation in São Paulo, with the first peak-shaving event executed on December 31 of that year, at 7:21 p.m. That is nearly four summers of effective dispatch, with public reports of performance during critical high-demand periods.

What the country lacks is a failure case record. There is no Brazilian incident-investigation report for a grid-storage system comparable to those underpinning the previous subsection. The difference is not a technicality: a record of successful operation and a forensic record produce different kinds of knowledge, and only the second one informs the design of the safety layer. A fleet that has never failed has not demonstrated that it will not fail — it has demonstrated that it has not yet encountered the combination of conditions that makes it fail.

There is a second discontinuity, one of operating regime, and it matters more than the first. Registro is a single, purpose-built asset, dispatched seasonally, owned by the transmission segment, and operated neutrally with respect to arbitrage. The fleet contracted in the December auctions will be plural, multi-service, market-exposed, subject to grid-forming requirements, and continuously cycled. Experience accumulated under one regime does not automatically transfer to the other. It is worth recalling that the Victorian Big Battery's failure mode occurred during commissioning — precisely at the transition between operating regimes, when the protections were configured for a state that no longer matched the actual use.

That same asset, however, has already produced this paper's single most important finding. It is a 30 MW/60 MWh LFP system with a 0.5C discharge rate, exceptionally authorized under ANEEL Authorization Resolution No. 10,892/2021 and installed as an infrastructure reinforcement in place of a new transmission line that would have crossed an environmental preservation area. Because it operates on only a handful of days each year, a study evaluated four additional applications for its idle capacity: operational flexibility under contingencies, islanded operation, system-restoration support, and expanded hosting capacity for renewables [9].

The conclusion that matters for this paper is explicit in the study: several of the proposed applications would require updating the installation's protection settings and control algorithms, along with auxiliary software and hardware upgrades. The physical asset is already built and sits idle most of the year; what limits service stacking is neither power nor energy — it is the control layer. This paper's central thesis finds its verification here, in a Brazilian asset already in commercial operation.

The same study yields a second finding, this one about the value of the envelope. In dynamic restoration simulations, primary frequency response limited to a 2.6 MW/s ramp raised the frequency nadir by 3.1 mHz; with steeper ramps, oscillatory behavior emerged that could compromise system stability [9]. Faster is not better. There is an upper bound on response aggressiveness imposed by stability, and determining where that bound lies, for each connection point, is exactly the function of the layer this paper calls the envelope. A controller that optimizes performance without that constraint finds instability before it finds the optimum.

It is also worth noting that, for the islanded-operation application, the study assumes the system operates in grid-forming mode, with V/f control and black-start capability — confirming, in a concrete case, the requirement discussed in Section II. The systemic context reinforces the urgency: the national grid operator projects a 33% loss of inertia in the National Interconnected System by 2030.

The national engineering foundation for these requirements already exists and is citable. There is a dynamic storage-system model implemented in ANATEM following the operator's modeling requirements, with frequency, voltage, active-, and reactive-power control [10]. There is a supervisory control system developed and implemented in a 100 kVA/200 kWh pilot plant, covering

power control, state of charge, ramp control, and communication [11]. There is domestic battery development with cell-level management, state-of-charge and state-of-health estimation, balancing, and thermal dissipation [12]. And there is a published technical-specification methodology for procuring storage systems, covering minimum supply requirements, type tests, factory tests, commissioning, operation, maintenance, and warranties [13] — a direct reference for Section XI.

What the Brazilian body of work lacks is not a foundation. It is a failure case record and characterization of the embedded layer under local thermal conditions. And the calendar leaves no time to accumulate either through first-hand experience.

The conclusion that follows from the calendar is straightforward. Brazil is moving from tens of megawatts in a single purpose-built asset to gigawatts under a multi-service regime, within a two-year window, with no incident base of its own. The choice is not between learning from accidents or not learning: it is between learning from other people's accidents beforehand, or from its own afterward.

IV. The Hierarchy of Time Constants

The debate over edge intelligence is often framed as an ideological choice between cloud and local. In fact, it is a consequence of physics. The criterion for allocating a control function is the time constant of the phenomenon it governs.

Table II — Allocation of control functions by the time constant of the governed phenomenon.

Phenomenon	Order of magnitude	Allocation
Cell protection and PCS current loop	µs to hundreds of µs	Dedicated hardware / firmware
Grid-forming response to disturbance	Sub-cycle (<16.7 ms at 60 Hz)	Deterministic local controller
Computing-cluster pulse load	Milliseconds	Deterministic local controller
Per-cell state acquisition and estimation	1 to 10 s	Embedded controller
Thermal control and module balancing	Seconds to minutes	Embedded controller
Dispatch optimization and arbitrage	Minutes to hours	Remote market service
Model requalification and fleet analysis	Days to months	Cloud

The table dispels two opposite oversimplifications. The first is that everything should go to the cloud: the first four rows are incompatible with long-distance network latency and its variability. The second is that everything should stay at the edge: market optimization and cross-asset fleet comparison benefit from centralized aggregation, and that is exactly where commercial offerings already work well.

There is also a matter of volume. Large-scale operators report needing per-cell state-of-charge and state-of-health readings at intervals of 1 to 10 seconds, or more frequently during cycling, across installations with thousands of containers. Carrying that raw volume off-site is costly, slow, and unnecessary: what needs to travel is inference, not the raw sample.

V. Perception Layer

A. State Estimation Under Embedded Constraints

State-of-charge estimation in lithium iron phosphate cells — the dominant chemistry in stationary storage — is complicated by a flat open-circuit-voltage plateau, strong thermal dependence, and parameter drift from aging. Purely empirical estimators degrade under non-stationary operating profiles.

The literature converges on hybrid architectures in which the neural network does not replace the filter but feeds it. One representative formulation couples a lightweight Transformer to an aging-aware second-order Kalman filter: the network processes current, voltage, temperature, and a health indicator to generate time-varying equivalent-circuit parameters, noise statistics, and bias terms, while the filter remains the sole state estimator. This choice preserves interpretability and uncertainty propagation, and puts learning to work exactly where it adds value — parameter adaptation.

For state-of-health estimation, physics-informed networks embed electrochemical degradation constraints into the loss function, typically anchored on the growth kinetics of the solid-electrolyte interphase layer. The payoff is physical consistency and generalization outside the training distribution [14], [15]. For an asset with a contracted lifetime measured in decades, physically plausible extrapolation is worth more than a low mean error on a test set.

Embedded feasibility has already been demonstrated: 8-bit integer-quantized models running with latencies on the order of 2 ms and footprints of roughly 1.2 kB of RAM and 11 kB of flash on Cortex-M0+ microcontrollers, with projection-based compression techniques allowing the approach to scale from cell level to module level. The embedded constraint is not an obstacle — it is a design requirement.

B. Imbalance as an Operational Variable

A practitioner from the industrial-automation side frames the problem in operational terms: imbalance is one of the biggest current challenges to maximizing revenue, and the design must identify, isolate, and correct imbalance problems both within containers and racks and across containers. This is state that exists only locally and does not survive aggregation — a rack average will not reveal a divergent cell.

C. Early Detection of Thermal Runaway

Decomposition of commercial liquid electrolyte begins around 150°C, and thermal stability varies across cathode chemistries, with LFP at the more favorable end. The critical design point is that temperature is a lagging indicator: gas emission, pressure change, deformation, and acoustic signature frequently precede any detectable temperature rise, which argues for multimodal fusion rather than an isolated thermal threshold. Reconstruction-error methods using bidirectional recurrent networks, and approaches that integrate a physical model with a deep network, point in the same direction: detect the precursor, not the event [16].

D. The Operating Point Is an Engineering Decision

Early detection is a classification problem with asymmetric costs. A false negative is a fire in a high-value asset, with risk to life and loss of license. A false positive is unnecessary shutdown, unavailability, and penalty — and, as the Moss Landing case demonstrates, it can be the damage itself. Choosing the operating point on the characteristic curve is a safety-engineering and risk-analysis decision, subject to the discipline of layers of protection, not a hyperparameter to be tuned by accuracy optimization.

The practical consequence is immediate for specification purposes. A vendor who presents aggregate detection accuracy without stating the false-negative rate in the operating region, detection latency, and behavior under sensor failure has not delivered the information that matters.

VI. Decision Layer

A. Degradation-Aware Predictive Control — and Its Objection

Model-predictive control is the natural formulation for storage dispatch because it explicitly handles constraints and horizon. The degradation-aware optimization literature imposes penalties on C-rate magnitude, on the state-of-charge operating window, and on imbalance between units. The obstacle is computational: high-fidelity degradation models are too expensive for a real-time receding horizon, and the practical workaround is a surrogate model that preserves degradation behavior at compatible cost.

This line of work faces an objection worth recording. Practitioners on the operations side argue that LFP has been meeting or beating cyclic-degradation schedules, and that the aggregate long-term capacity-loss picture matters more to operators and financiers than the marginal degradation from charge-discharge behaviors that do not repeat consistently; in ten- to twenty-year projects, the decisive determination would be whether the DC block is expanded or replaced once, twice, or three times.

The objection is well taken, and it bounds the argument. Penalizing degradation in the embedded controller is not justified by capturing fractions of a cent in lifetime per cycle. It is justified as an operational and thermal constraint — preventing an excessive C-rate for the current state of health, preventing operation outside the safe window, preventing inter-unit imbalance from worsening. It is an envelope function, not fine-grained optimization. This distinction runs through the rest of the paper.

B. Where Economic Optimization Actually Lives

The pressure to capture marginal revenue is real and growing: peak pricing in the Texas market fell from US$4,853/MWh in 2024 to US$345/MWh in 2026, making every slice of margin more relevant as markets mature. But that is a trading-platform competency, it operates on an intraday scale, and it is already offered commercially at gigawatt scale.

Under Resolution 1,162/2026, the cost function must internalize the tariff modality, contracted demand, operating windows, and the relationship with CCEE, ONS, or the local utility, on top of the degradation cost per MWh cycled. An optimizer that maximizes spread while ignoring Brazilian tariff structure will produce confidently wrong decisions — but that is a platform-localization problem, not an embedded-architecture one.

What the edge produces that cannot be sourced elsewhere is something else: the envelope within which any optimizer, local or remote, can operate. Live constraints derived from physical state — imbalance across racks and containers, thermal gradient, admissible C-rate for the current state of health, failure precursors, connection envelope. Without that envelope, the remote optimizer operates on an abstraction that grows stale between one update and the next.

VII. Shielding Layer

A. Why the Probabilistic Controller Cannot Sustain the Argument

Reinforcement-learning-based controllers do not, on their own, offer safety guarantees, and that limitation is the principal obstacle to their adoption in critical infrastructure. The objection is not conservatism: a learned policy offers, at best, a probabilistic guarantee over the training distribution, while operation demands certifiable compliance with hard constraints, including under rare disturbances. A system's safety argument cannot depend on the quality of its policy.

B. Runtime Shielding

The consolidated answer is runtime shielding: a high-level policy proposes actions; a deterministic shield filters out unsafe actions through fast forward simulation, imposing safety as a runtime invariant, independent of policy quality and of the training distribution [17]. Alternative formulations verify actions through a safety layer built on set-based backward reachability analysis, with a temporal constraint that explicitly encodes the islanding contingency [18], or employ control barrier functions to guarantee invariance of the safe set.

The common denominator is an asymmetry of authority: the learned component proposes, the verified component disposes. In a storage system, the shield materializes as already-known, verifiable limits — per-cell voltage and temperature windows, maximum C-rate for the current state of health, admissible thermal gradient between modules, connection power envelope, islanding-contingency constraints.

C. Where the Constraint Set Comes From

A legitimate question is how the set of limits the shield imposes is derived. Enumerating them by engineering intuition covers the obvious cases and omits precisely the ones that cause accidents — control actions that are individually valid yet dangerous in combination, or correct in one system state and unsafe in another.

System-Theoretic Process Analysis (STPA) was proposed specifically for grid storage with this purpose in mind, treating battery-management-system failures, flammable gases, electrical arcing, fire, explosion, and component interactions as a problem of inadequate control rather than of isolated component failure [19]. The method produces an inventory of unsafe control actions — commanding when one should not, failing to command when one should, commanding too early or too late, commanding for an inadequate duration. That inventory is precisely the specification for the shielding layer, and the Brazilian finding of oscillatory behavior under excessive ramp rate, discussed in Section III, is an example of a control action that is unsafe only by duration and magnitude, not by nature.

D. Normative Anchoring

This architectural pattern is not an invention of the power sector. ISO/IEC TR 5469:2024 addresses the properties, risk factors, and methods related to using artificial intelligence within safety-related functions, to using non-AI-based functions to safeguard AI-controlled equipment, and to using AI systems in the development of safety functions. Among the architectural patterns it describes is that of a safe, if suboptimal, backup function designed with non-AI techniques [20]. The shielding described above is the instantiation of that pattern within the storage domain.

The complete normative chain combines, at minimum: IEC 61508 for the functional-safety lifecycle and the determination of the safety integrity level [21]; IEC 62933-5-2:2025 for the safety requirements of grid-integrated electrochemical systems [22]; the IEC 62443 series for automation-and-control cybersecurity [23]; and IEC TR 63069 for coordinating functional safety with cybersecurity [24]. The ongoing revision of IEC 61508 incorporates cybersecurity threats into the safety analysis and explicitly references ISO/IEC TR 5469 — a sign that convergence among the three vocabularies is no longer merely an academic topic.

The relevance of this anchoring to the business model is concrete, and Axminster demonstrates it by negative example. Insurers, financiers, and licensing authorities assess risk based on demonstration, not declaration. An architecture that maps each AI component to a TR 5469 pattern and each safety function to an IEC 61508 integrity level produces auditable evidence; one that offers only product certificates and model performance metrics does not.

E. Design Rule

In summary: no learning-based component holds direct write authority over actuators. The safety function remains in verifiable deterministic logic, executed on hardware segregated from the inference hardware, with the constraint set encoded independently of the model, and with no operating mode capable of disabling it. AI operates within the envelope; it does not define the envelope.

The rule has a cost: performance is bounded by the shield's conservatism. It is a cost worth paying, because it converts a non-demonstrable system into a demonstrable one, and because that conservatism can be relaxed gradually as operational evidence accumulates. The reverse is not true.

VIII. The Computational Envelope

Running models inside a storage container is subject to size, weight, and power constraints that simply do not exist in the cloud. The enclosure is sealed, the environment is hot, dusty, and subject to vibration, and heat dissipation competes with the battery bank's own cooling. Under this regime, processor thermal throttling reduces clock frequency, degrading inference throughput precisely under the most critical operating conditions.

Two classic limits reassert themselves. Dark silicon requires a share of transistors to remain idle to avoid exceeding the thermal design power [25]. And the memory-bandwidth bottleneck — the memory wall described three decades ago — is the actual limiting factor for inference of attention-based models on embedded platforms, more so than arithmetic throughput [26].

There is, however, a less-discussed and more decisive constraint for control: what matters is not average inference latency, but worst-case execution time. An estimator whose inference occasionally takes three times its typical duration is not a slow estimator — it is one that intermittently breaks the control loop in ways that are hard to diagnose. Embedded design for control requires a demonstrable upper bound, which pushes the choices toward fixed-topology, quantized models with no dynamic memory allocation and no data-dependent execution paths. The acceptance criterion is conformance to the timing envelope, not accuracy in isolation.

IX. Validation

Trial-and-error training on a live bank is unacceptable: it exposes a high-value asset to induced degradation and safety risk, and the events that matter most — cell failure, sensor failure, loss of cooling, transition to islanding — are rare by construction.

The practicable alternative unfolds in three stages. First, an electrothermal digital twin coupled to a degradation model, calibrated with bench data, in which policies are trained and exposed to extreme scenarios. Second, hardware-in-the-loop testing, in which the real embedded controller — with its actual firmware, scheduler, and thermal envelope — runs against a simulated plant in real time; this is the stage at which worst-case execution-time problems surface, not before. Third, commissioning in supervisory mode, with the controller issuing logged recommendations but no actuation authority, until accumulated evidence justifies closing the loop.

It is worth dispelling a recurring objection here — that relying on simulation and hardware-in-the-loop testing is a concession made in the absence of field data. Under the

current Brazilian regime, the relationship runs the other way: C-HIL validation is the procedure the national grid operator itself requires before connection. The joint ONS/EPE technical note, in its May 2026 revision, covers ANATEM and PSCAD models, C-HIL validation, requirements for grid-forming and grid-following operating modes, voltage and frequency control, black start, and commissioning tests [5]. Closed-loop testing with the real controller does not substitute for field evidence for lack of anything better — it is the normatively required evidence, and for good reason: the events that determine the asset's safety are too rare to be expected to show up in operation.

The same design is consolidating internationally. The Australian grid operator established, for 2026, a pre-commissioning test plan for grid-forming storage systems, including the creation and testing of a standardized hardware-in-the-loop test procedure to validate functionality before grid connection [27].

Transfer from simulation to reality depends on domain randomization: systematically varying, during training, the spread of capacity and internal resistance across cells, sensor noise and bias, ambient temperature, and actuation characteristics [28]. The goal is not the model that performs most precisely in the nominal simulator, but the policy most robust to the gap between simulator and plant.

X. Reference Architecture

The proposal consolidates into four vertical layers, crossed by a horizontal model-governance cycle.

1. Instrumentation. Per-cell and per-module acquisition of voltage, current, temperature, pressure, and strain; instrumentation of the power-conversion system and the cooling circuit; heterogeneous industrial interfaces normalized at a single point.
2. Perception. Embedded state-of-charge and state-of-health estimation through a hybrid network-filter estimator; multimodal detection of thermal-failure precursors; characterization of intra- and inter-container imbalance; local forecasting of net load and co-located generation.
3. Decision. Composing the operating envelope from the current physical state, and local arbitration among competing requirements within that envelope, interfacing with the remote market-optimization platform.
4. Shielding. Deterministic verification of every proposed action — whether of local or remote origin — against the set of hard constraints, executed on segregated hardware, with no disabling mode, and with logging of active constraints for traceability and regulatory oversight.

The governance cycle collects selective telemetry — not raw samples — monitors estimator performance drift against observed behavior, and feeds periodic requalification off-site, with redeployment subject to the same three-stage validation procedure.

Two properties deserve emphasis. The first is that the flow leaving the facility carries inference and events, not continuous samples, which simultaneously reduces communication cost, cyber-exposure surface, and dependence on connectivity. The second is that the shielding layer is the only one with actuation authority, which allows the upper layers to evolve — including by swapping models and optimization vendors — without reopening the system's safety argument.

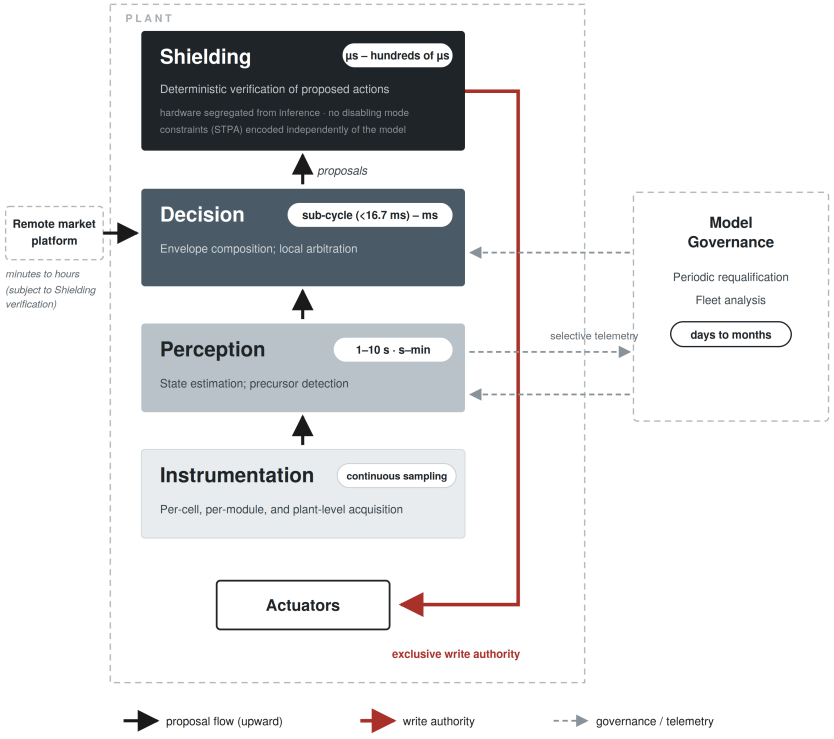

Fig. 1. Reference architecture for BESS control, integrating instrumentation, perception, decision, deterministic shielding, and model governance across distinct operational timescales.

XI. From Architecture to Specification

Axminster sets the standard by which this section should be read. The developer complied with applicable guidance, used equipment from a recognized manufacturer, and lost the project because the licensing authority lacked confidence that the controls would be effective. The question a buyer needs to put to a vendor is not whether it is compliant. It is what evidence it presents that behavior stays within known limits when something fails.

Eight items translate the argument into specification requirements. Each states what to demand and what the absence of an answer means.

1. Declared separation of authority. Require the vendor to declare which functions are implemented by learning-based components and which remain in deterministic logic, with explicit mapping to the corresponding ISO/IEC TR 5469:2024 architectural pattern [20]. The absence of an answer indicates the boundary was inherited, not designed.
2. Demonstrated timing envelope. Require worst-case execution time, not average latency, for fast-loop functions, with hardware-in-the-loop test evidence under design thermal conditions. An answer expressed in average latency indicates the worst case was never measured.
3. Detection operating point. Require a stated false-negative rate in the operating region, detection latency, and behavior under sensor failure. Aggregate accuracy alone is not an answer.
4. Containment evidence, and at what level. Require the test report, not the certificate, and the level at which it was conducted — cell, module, unit, or installation — along with the propagation result. The test level determines what it demonstrates.
5. Non-disablement of the safety function. Require a declaration that no operational, service, or commissioning mode can remove the protection, and the applicable procedure when maintenance requires intervention. This is the item the Victorian Big Battery made mandatory.
6. Decision traceability. Require logging of active constraints and blocked actions, with defined retention. Under the penalty regime of Resolution 1,162/2026, this is what allows the system to demonstrate why it acted as it did.
7. Autonomy without connectivity. Require the specified behavior and minimum duration of autonomous operation in the absence of communication with the control center, with a corresponding acceptance test.
8. Interoperability and data ownership. Specify open interfaces appropriate to the integration point and ownership of operational data, preventing the controller's evolution from becoming captive to a single vendor.

None of these items is exotic, and all of them are frequently refused by vendors whose business model depends on keeping the intelligence remote and the asset peripheral. That is precisely the reason to specify them before contracting, while the buyer still holds the leverage.

There is an asymmetry worth naming. A certificate transfers to the buyer the conclusion of a test, not the competence to assess its scope. Knowing that a system passed a given test does not reveal at what level it was conducted, what scenarios it covered, or what remains outside the tested envelope. That asymmetry is not resolved by demanding more certificates; it is resolved by demanding evidence and knowing how to read it.

XII. What This Analysis Does Not Demonstrate

Honesty about the argument's limits is part of the argument. This analysis maintains that the failure mode is real, recurring, and documented; that fulfilled regulatory compliance is not sufficient to license a project; that the timing requirement follows from standards and from physics, not from design preference; that the separation between the learning component and the safety function is an established normative pattern; that closed-loop test validation is the procedure required before connection, not a substitute for field evidence; and that, in an operating Brazilian asset, the limit on service stacking was identified in the control layer, not in the sizing.

It does not demonstrate that the proposed integrated architecture works in the field. The individual components have published domestic precedents — dynamic modeling per the operator's requirements, supervisory control in a pilot plant, cell-level battery management — but the four-layer composition with segregated authority has not been built and tested as a system.

What is specifically still missing is the characterization of worst-case execution time for embedded inference under thermal throttling in a container operating in a tropical climate. That characterization does not exist publicly, and it cannot be inferred from results obtained in temperate climates, because thermal throttling is a function of the margin between junction temperature and ambient — a margin that the Brazilian climate compresses precisely during the hours of heaviest asset demand. Closing this gap requires an instrumented test bench with a representative thermal profile and measurement of the worst-case execution-time bound under throttling, not simulation. In this author's assessment, it is the most consequential experimental gap for the fleet entering operation from 2028 onward.

Nor does this analysis present any quantitative gain in revenue or lifetime, and it should not be read as an evaluation of any specific vendor or technology.

It is also worth recording the limits of the very technology this paper advocates for. The Australian operator noted, in its July 2026 annual risk review, that grid-forming inverter storage systems have demonstrated the ability to sustain voltage-waveform stability, but have not yet been confirmed capable of delivering protection-grade fault current at the standard required for minimum system-strength requirements — a question a dedicated test seeks to resolve [27]. Anyone advocating grid-forming technology as a substitute for synchronous machines needs to say so.

XIII. Conclusion

The June 2026 regulatory framework closed the question of whether storage is viable in Brazil, and opened a different, more technically demanding one. With grid-forming as the connection standard, revenue stacking as a condition for returns, and a penalty regime bearing on the storage operator, the control architecture has stopped being a subsystem and has started determining the asset's economics and risk.

The argument defended here is that the relevant boundary does not separate those who comply with standards from those who do not — nearly everyone will comply. It separates those who can demonstrate that behavior stays within known limits under failure from those who merely declare compliance. That capacity is architectural: it comes from allocating each function by the time constant of the phenomenon it governs, from locally producing the operating envelope that remote optimization cannot derive, and from keeping actuation authority in a verifiable component with no disabling mode.

The technical path is known and does not depend on speculative algorithmic breakthroughs: hybrid estimators, degradation-model surrogates, runtime shielding, and normative standards already exist and are documented. What is missing is integration — systems-engineering work, not fundamental research.

Two years remain until the first commercial operation. 6,091 projects have been registered, totaling 296.8 GW. The country is moving from a single demonstration asset to a gigawatt-scale fleet under a distinct operating regime, with no failure case record of its own. The choice is not between learning from accidents or not learning: it is between learning from other people's accidents now, or from its own later — and the window to correctly specify these assets closes before they come into existence.

References

[1] Brazilian National Electric Energy Agency (ANEEL), Normative Resolution No. 1,161, of June 2, 2026. Establishes requirements and procedures for granting authorization to energy storage systems. Brasília, Brazil: ANEEL, 2026.

[2] Brazilian National Electric Energy Agency (ANEEL), Normative Resolution No. 1,162, of June 2, 2026. Amends Normative Resolution No. 1,000 of December 7, 2021, to regulate billing and the regulatory classification of electric energy storage systems. Brasília, Brazil: ANEEL, 2026.

[3] Brazil, Ministry of Mines and Energy, Normative Ordinance No. 136, of June 1, 2026. Establishes the guidelines and procedures for the Capacity Reserve Auction through new battery energy storage systems — LRCAP 2026. Brasília, Brazil: Official Gazette of the Union, 2026.

[4] Institute of Electrical and Electronics Engineers (IEEE), IEEE Std 2800-2022: Standard for Interconnection and Interoperability of Inverter-Based Resources (IBRs) Interconnecting with Associated Transmission Electric Power Systems. New York: IEEE, 2022.

[5] National Electric System Operator (ONS) and Energy Research Company (EPE), Technical Note NT-ONS DPL 0111/2025 – EPE-DEE-NT-095-2025: Minimum Technical Requirements for the Connection of Battery Energy Storage Systems. Rio de Janeiro: ONS/EPE, 2025, rev. 4, May 26, 2026.

[6] Energy Safe Victoria, Statement of Technical Findings: Fire at the Victorian Big Battery. Melbourne: ESV, Sep. 28, 2021.

[7] D. Hill, McMicken Battery Energy Storage System Event Technical Analysis and Recommendations, report prepared for Arizona Public Service. Chalfont, PA: DNV GL Energy Insights, Jul. 2020.

[8] QBE Insurance, "Solar panel fires increasing at more than double the rate of installations," QBE European Operations, London, Aug. 2026.

[9] P. F. Torres, A. R. A. Manito, G. Figueiredo, M. P. Almeida, J. C. S. Almeida Neto, R. L. Cavalcante, C. C. V. F. A. Silva, and R. Zilles, "Energy storage as a transmission asset — assessing the multiple uses of a utility-scale battery energy storage system in Brazil," Energies, vol. 18, no. 4, art. 902, 2025, doi: 10.3390/en18040902.

[10] P. F. Torres, G. Figueiredo, M. P. Almeida, A. R. A. Manito, J. C. S. Almeida Neto, M. A. Cassares, and R. Zilles, "Dynamic model of battery energy storage systems for the provision of ancillary services," Eletrônica de Potência, vol. 28, no. 2, pp. 94–106, 2023, doi: 10.18618/REP.2023.2.0036.

[11] A. R. A. Manito, M. P. Almeida, M. A. Cassares, J. C. S. Almeida Neto, P. F. Torres, R. Zilles, and G. Figueiredo, "Development and implementation of a high-level control system — SCAN," Revista Brasileira de Energia Solar, vol. 13, no. 1, 2022, doi: 10.59627/rbens.2022v13i1.388.

[12] M. F. N. Rosolem et al., "Development of a lithium-ion battery for ancillary services," in Proc. 8th Brazilian Congress on Solar Energy (CBENS), 2020, doi: 10.59627/cbens.2020.910.

[13] V. B. Riboldi, P. A. B. Block, and T. R. Ricciardi, "Technical specification of energy storage systems connected to distribution grids," in Proc. Brazilian Congress on Automatic Control (CBA), 2020, doi: 10.48011/asba.v2i1.1559.

[14] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," J. Comput. Phys., vol. 378, pp. 686–707, 2019.

[15] F. Wang, Z. Zhai, Z. Zhao, Y. Di, and X. Chen, "Physics-informed neural network for lithium-ion battery degradation stable modeling and prognosis," Nat. Commun., vol. 15, no. 1, art. 4332, 2024.

[16] M. Naguib et al., "Thermal fault detection of lithium-ion battery packs through an integrated physics and deep neural network based model," Commun. Eng., 2025.

[17] G. Malik, "Hierarchical reinforcement learning with runtime safety shielding for power grid operation," arXiv preprint arXiv:2604.14032, 2026.

[18] M. Eichelbeck, H. Markgraf, and M. Althoff, "Contingency-constrained economic dispatch with safe reinforcement learning," arXiv preprint arXiv:2205.06212, 2022.

[19] Sandia National Laboratories, Analyzing System Safety in Lithium-Ion Grid Energy Storage. Albuquerque, NM: Sandia National Laboratories, 2016. [Online]. Available: https://www.osti.gov/servlets/purl/1257985

[20] International Organization for Standardization and International Electrotechnical Commission, ISO/IEC TR 5469:2024: Artificial Intelligence — Functional Safety and AI Systems. Geneva: ISO/IEC, 2024.

[21] International Electrotechnical Commission (IEC), IEC 61508: Functional Safety of Electrical/Electronic/Programmable Electronic Safety-Related Systems, parts 1–7. Geneva: IEC.

[22] International Electrotechnical Commission (IEC), IEC 62933-5-2:2025: Electrical Energy Storage (EES) Systems — Part 5-2: Safety Requirements for Grid-Integrated EES Systems — Electrochemical-Based Systems. Geneva: IEC, 2025.

[23] International Electrotechnical Commission (IEC), IEC 62443: Security for Industrial Automation and Control Systems, series. Geneva: IEC.

[24] International Electrotechnical Commission (IEC), IEC TR 63069: Industrial-Process Measurement, Control and Automation — Framework for Functional Safety and Security. Geneva: IEC.

[25] H. Esmaeilzadeh, E. Blem, R. St. Amant, K. Sankaralingam, and D. Burger, "Dark silicon and the end of multicore scaling," in Proc. 38th Int. Symp. Computer Architecture (ISCA), San Jose, CA, 2011.

[26] W. A. Wulf and S. A. McKee, “Hitting the memory wall: Implications of the obvious,” ACM SIGARCH Comput. Archit. News, vol. 23, no. 1, pp. 20–24, 1995.

[27] Australian Energy Market Operator (AEMO), General Power System Risk Review 2026. Melbourne: AEMO, Jul. 2026.

[28] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, “Domain randomization for transferring deep neural networks from simulation to the real world,” in Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS), Vancouver, Canada, 2017.