Mission-Energy: Homeostasis / Allostasis
The Self-Regulating Control Doctrine for Resilient Energy Ecosystems
Author: Michael S. Berger
Affiliation: Beech Creek Power & Energy, LLC
Date: May 2026
Version: 1.0

© 2026 Michael S. Berger. All Rights Reserved.

This work is protected under U.S. copyright law. No portion of this publication may be reproduced, distributed, transmitted, or used in any form or by any means, electronic or mechanical, without prior written permission from the author, except for brief quotations used in reviews, academic citation, or scholarly analysis.

Civilization does not survive because it believes harder.... It survives because it solves the physics in front of it.

— Michael S. Berger

MISSION-ENERGY
HOMEOSTASIS / ALLOSTASIS

THE SELF-REGULATING CONTROL DOCTRINE FOR RESILIENT ENERGY ECOSYSTEMS

AI/ML DIGITAL TWIN GOVERNANCE FOR TACTICAL AND INSTALLATION-SCALE MICROGRIDS

AUTHOR: Michael S. Berger, Beech Creek Power & Energy

MISSION-FOCUSED
Protect what matters most.

ADAPTIVE
Anticipate. Adjust. Stay ahead.

RESILIENT
Withstand stress. Recover faster.

INTELLIGENT
Physics-informed. AI/ML-driven.

TRUSTED
Explainable. Secure. Human-authorized.

GENERATION
Conventional & Hybrid

STORAGE
BESS & Energy Management

RENEWABLES
Solar & Other Distributed Inputs

POWER CONVERSION
Inverters / PCS & Controls

DISTRIBUTION
Protection, Switching & Priority Loads

COMMUNICATIONS
Secure Networks & Telemetry

OPERATORS
Human Authority & Oversight

GOVERNED POWER. MISSION ASSURED. RESILIENCE DELIVERED.

FROM RUCKSACK POWER TO INSTALLATION NODES AND BEYOND.

Mission-Energy Homeostasis / Allostasis

Beech Creek Power & Energy — Michael S. Berger, 2026

1. Abstract / Executive Summary

Mission-Energy Homeostasis / Allostasis is a proposed control doctrine for tactical and installation-scale energy ecosystems that must preserve mission function under stress. It applies the logic of biological regulation to military and critical-infrastructure power systems: homeostasis restores balance after drift or disturbance, while allostasis adapts the system before forecasted stress becomes failure.

In this framework, energy resilience is not defined only by generation capacity, battery duration, fuel supply, or installed equipment. It is defined by whether the system can preserve mission-valid equilibrium when loads change, assets degrade, communications weaken, fuel becomes constrained, faults emerge, cyber trust degrades, or operating priorities shift.

This paper advances the argument from prior work on fuel discipline and energy advantage. *Fossil-Smarter* established that fuel remains strategically necessary but should not be wasted through routine, inefficient combustion. *Solving for Energy Advantage* expanded that position into a broader architecture for coordinated generation, storage, power conversion, protected distribution, telemetry, cybersecurity, human authority, and grid interaction. Mission-Energy Homeostasis / Allostasis now defines the supervisory control behavior required to make that architecture adaptive, bounded, explainable, recoverable, and trusted.

The central claim is straightforward: a physics-informed, confidence-gated AI/ML digital twin can help preserve energy balance across tactical and installation-scale microgrids by sensing drift, modeling consequence, recommending or executing bounded corrective and anticipatory actions, informing the human operator, and returning the system to a mission-valid operating state. The AI/ML-DT is not proposed as unchecked autonomy. It does not replace protective relays, hardwired interlocks, inverter inner loops, battery management protections, local safe-state logic, or human command authority. It operates above deterministic control layers as a bounded supervisory intelligence.

The doctrine is designed around five recurring functions: **detect, understand, act, inform, and return.** The system detects drift, stress, abnormal behavior, degraded confidence, cyber-relevant anomalies, or emerging risk. It understands the condition through physics-informed modeling, telemetry fusion, mission context, operating limits, confidence scoring, and forecasted consequence. It acts only within bounded authority through pre-approved corrective or anticipatory control pathways. It informs the operator with clear explanation, confidence level, action status, authority boundary, and escalation requirements. It returns the energy ecosystem to mission-valid equilibrium by preserving priority loads, protecting reserve, stabilizing the system, restoring function in sequence, and capturing evidence for after-action review and revalidation.

This paper also argues that the AI/ML-DT itself must be treated as a system under test. Its predictions, classifications, confidence scores, recommendations, autonomous actions, refusal-to-act behavior, escalation decisions, fallback transitions, and explanations must be validated under realistic disturbance conditions — including load changes, equipment faults, communications degradation, stale or corrupted telemetry, model/plant mismatch, cyber-relevant events, operator override, and out-of-envelope scenarios.

The operational value is significant. At the MPTaPS or LOTaPS level, the doctrine can extend silent-watch endurance, reduce unnecessary generator runtime, preserve battery reserve, manage recharge windows, protect critical loads, reduce signature exposure, and reduce operator burden. At installation scale, it can coordinate distributed energy nodes, reduce single-point fragility, support recovery from localized faults, and preserve mission function without relying on one centralized power asset.

Mission-Energy Homeostasis / Allostasis is the research foundation — the R — of a deliberate RDT&E progression. It establishes the doctrine, defines the framework, bounds the authority model, and motivates the program. It does not produce validated confidence thresholds, tested operating envelopes, or certified system behavior. Those outputs cannot be written. They must be measured. The gap between doctrine and validated implementation closes through Development, Test, and Evaluation conducted against real hardware, real telemetry, real degraded conditions, and edge cases no architecture document can fully anticipate.

The future of mission energy resilience will not belong to whoever simply installs the most generation, storage, or automation. It will belong to systems that can govern power with discipline under stress. In high-consequence environments, governed power is the difference between energy that exists and energy that can be trusted when normal conditions disappear.

2. Problem Statement and Technical Gap

Modern tactical and installation-scale energy systems are becoming more capable, but not necessarily more resilient. More generation, more storage, more inverters, more solar, more telemetry, and more dashboards do not automatically produce mission assurance. A power system can contain advanced equipment and still fail to preserve the right loads, at the right time, under the right constraints, with the right authority, the right reserve posture, and the right recovery path.

The core problem is not simply energy supply. It is **governed energy behavior under stress.**

Current microgrid practice is strong in many areas — generation, storage, interconnection, islanding, power conversion, protection, monitoring, and supervisory control. However, a gap remains between

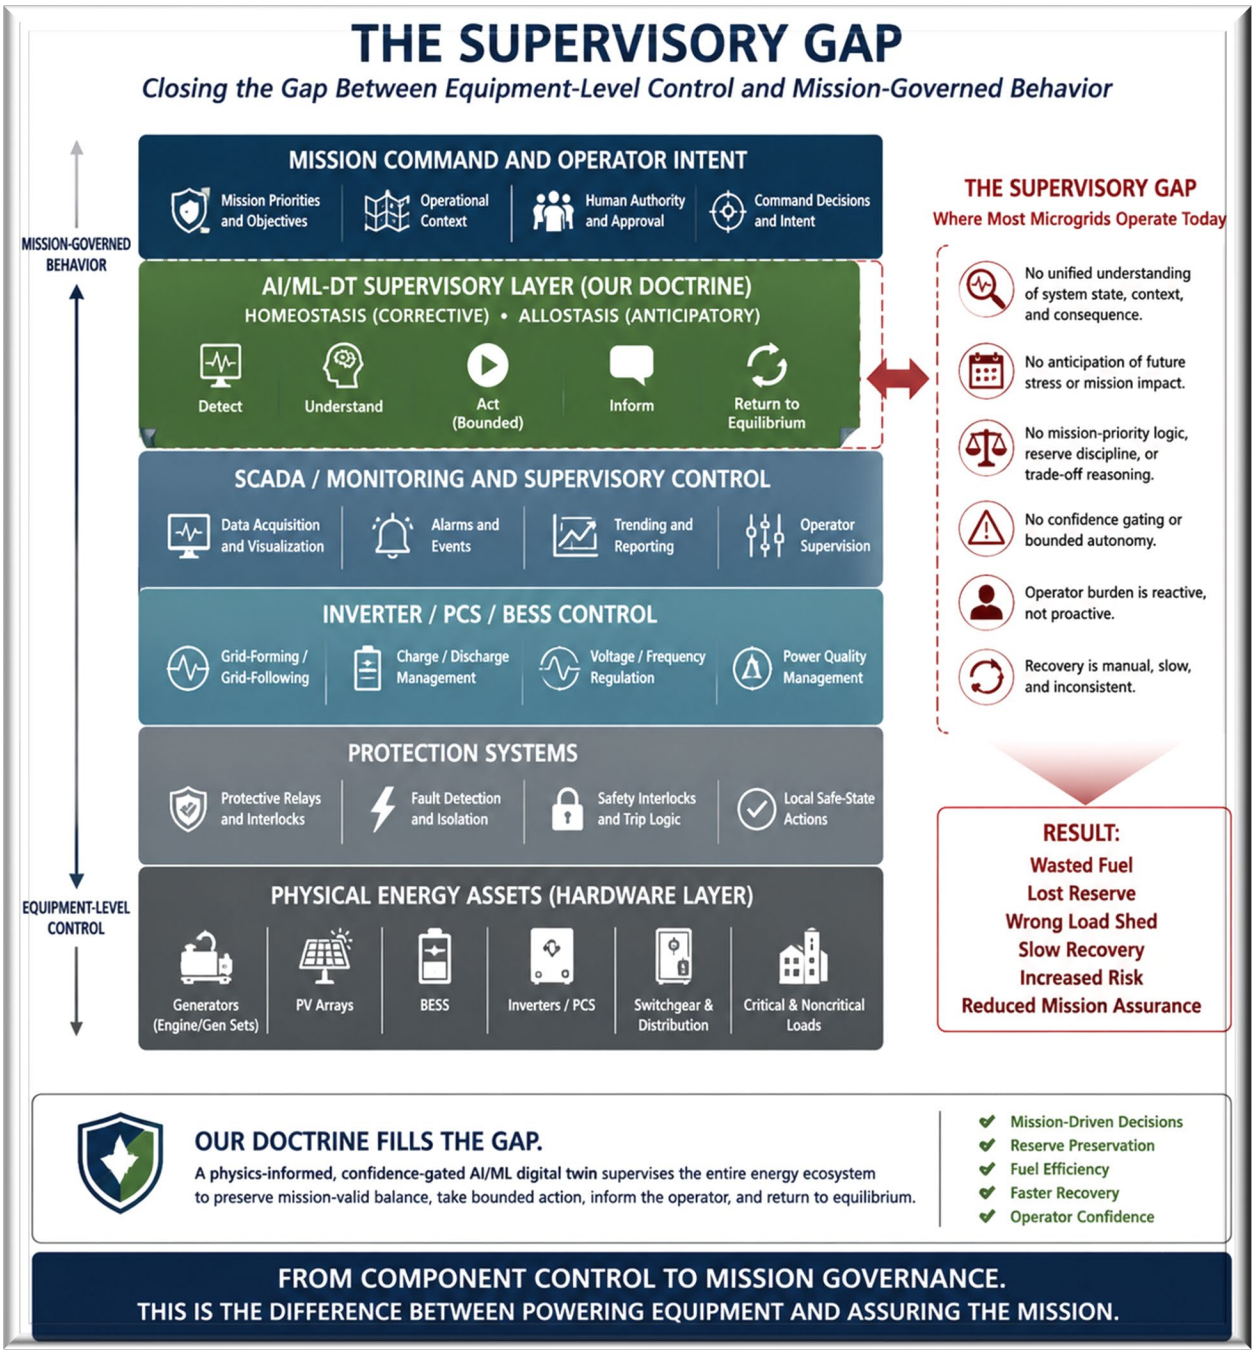

equipment-level control and mission-governed behavior. Existing systems may regulate voltage, frequency, power flow, and source dispatch, but they often do not explicitly answer the higher-order mission questions:

- Which loads must survive first?
- How much reserve must be protected for the next mission phase?
- When should fuel be burned, preserved, or delayed?

- When should storage carry the load instead of starting the generator?
- When does a degraded node require support from another node?
- When is telemetry too stale, corrupted, or unauthenticated to trust?
- When has the digital twin drifted too far from the physical system?
- When should the system stop optimizing and fall back?
- When must the human operator approve the next action?
- How does the system prove it returned to mission-valid equilibrium?

Five distinct technical gaps are addressed by this doctrine:

Gap 1 — Mission-governed supervisory behavior. These questions are not answered by hardware alone. They require a supervisory doctrine that connects physical state, mission priority, confidence, authority, cybersecurity, recovery logic, and evidence.

Gap 2 — Bounded AI/ML digital-twin control. AI and digital twins are often presented as tools for prediction, visualization, or optimization. In high-consequence energy systems, the AI/ML-DT must also understand what it is allowed to do, when confidence is sufficient, when deterministic protection takes precedence, when human approval is required, and when the safest action is no autonomous action at all.

Gap 3 — Evidence discipline. Resilience claims made through architecture diagrams or conceptual descriptions are not enough. A mission-energy system must prove its behavior through testable outcomes: critical-load continuity, reserve preservation, detection latency, confidence calibration, stale-state suppression, fallback success, and equilibrium return time.

Gap 4 — Cybersecurity and data trust. A self-regulating energy ecosystem cannot depend on blind trust in telemetry, commands, configurations, models, or communications pathways. Cyber uncertainty is operational uncertainty.

Gap 5 — Recovery. Backup is not the same as resilience. True resilience requires the system to stabilize, isolate, rebalance, restore, learn, and return to a mission-valid operating state.

The technical gap is not the lack of microgrid hardware. The gap is the lack of a bounded, testable, mission-governed supervisory doctrine that allows tactical and installation-scale energy systems to detect stress, understand consequence, act within validated authority, inform the human operator, fall back safely, preserve evidence, and return to equilibrium.

3. Definitions and Key Terms

This paper uses the following terms to establish a common vocabulary for Mission-Energy Homeostasis / Allostasis. These definitions are intended to support engineering discussion, test planning, authority partitioning, and future RDT&E work.

Term	Definition
Mission-Energy Homeostasis / Allostasis	The control doctrine by which a physics-informed, confidence-gated AI/ML digital twin supervises a tactical or installation-scale energy ecosystem to preserve mission-valid balance, take bounded corrective or anticipatory action, preserve human authority, and return the system to equilibrium.
Homeostasis	The corrective function. It detects deviation from a validated operating state and initiates bounded action to restore stability.
Allostasis	The anticipatory function. It adapts energy posture before forecasted stress becomes failure.
AI/ML-DT	The combined artificial intelligence, machine learning, and digital-twin supervisory layer. It operates above deterministic control layers as a bounded supervisory intelligence.
Physics-Informed Digital Twin	A digitally synchronized model of the real energy system, constrained by actual equipment topology, operating limits, control modes, load hierarchy, reserve state, thermal behavior, communications status, and measured system condition.
Confidence-Gated Control	A control approach where no recommendation or autonomous action is permitted unless telemetry quality, model synchronization, scenario validity, authority boundaries, safety conditions, cybersecurity posture, and mission rules satisfy predefined criteria.
Validated Operating Envelope	The tested range of configurations, loads, equipment states, communications conditions, cybersecurity assumptions, and control actions where evidence demonstrates acceptable system behavior.
Left-Right Boundaries	Operational limits defining what the AI/ML-DT may observe, recommend, or execute. These include electrical, thermal, cyber, communications, fuel, reserve, safety, protection, mission-priority, and human-authority limits.
Mission-Valid Equilibrium	A stable operating state where required loads, reserve margins, power-quality limits, thermal limits, cybersecurity conditions, operator-authority rules, and recovery posture are satisfied for the active mission context.
Bounded Autonomy	The condition where the system may execute only pre-approved actions inside a validated operating envelope and must escalate, constrain itself, or fall back when confidence, authority, cybersecurity, or safety conditions are insufficient.
Deterministic Fallback	Hard-coded, validated, non-AI protective and safe-state logic that preserves safety and critical function when the AI/ML-DT is uncertain, impaired, compromised, or outside its validated operating envelope.

Term	Definition
Human-in-the-Loop (HITL)	An operating condition where specified actions require human approval before execution.
Human-on-the-Loop (HOTL)	An operating condition where the system may execute bounded actions while a human maintains supervisory awareness, override authority, and accountability.
Priority Load	A load whose continuity directly supports mission function, safety, life support, communications, command and control, security, medical operation, or other designated mission need.
Deferrable Load	A load that may be delayed, reduced, shifted, or shed without immediate mission failure.
Reserve Floor	The minimum required stored energy, fuel, or support capacity that must be preserved for a defined mission condition, recovery requirement, or contingency window.
Equilibrium Return Time	The measured time required to restore the system from a disturbed or degraded state to mission-valid equilibrium.
Distributed Energy Node	A local generation-storage-conversion-load support unit that can operate independently or as part of a coordinated energy ecosystem.
Energy Ecosystem	The full mission-relevant set of generators, storage, conversion equipment, distribution assets, loads, sensors, communications, controls, operators, cybersecurity controls, governing policies, and recovery processes.
Signature-Aware Energy Posture	An operating mode that accounts for acoustic, thermal, electromagnetic, fueling, maintenance, runtime-pattern, and observability consequences, not only electrical sufficiency.
Model/Plant Mismatch	A condition where the digital twin no longer accurately represents the physical system due to equipment degradation, configuration changes, telemetry errors, sensor drift, or undocumented field modifications.
Revalidation Trigger	Any change, event, fault, behavior, or configuration update requiring renewed testing or review before the AI/ML-DT is allowed to retain or expand control authority.
Trust	Demonstrated, evidence-backed confidence in system behavior inside a validated operating envelope. Trust is earned through testing, explanation, fallback behavior, operator authority, cybersecurity, auditability, and repeatable performance.

4. Introduction: Establishing Mission-Energy Balance

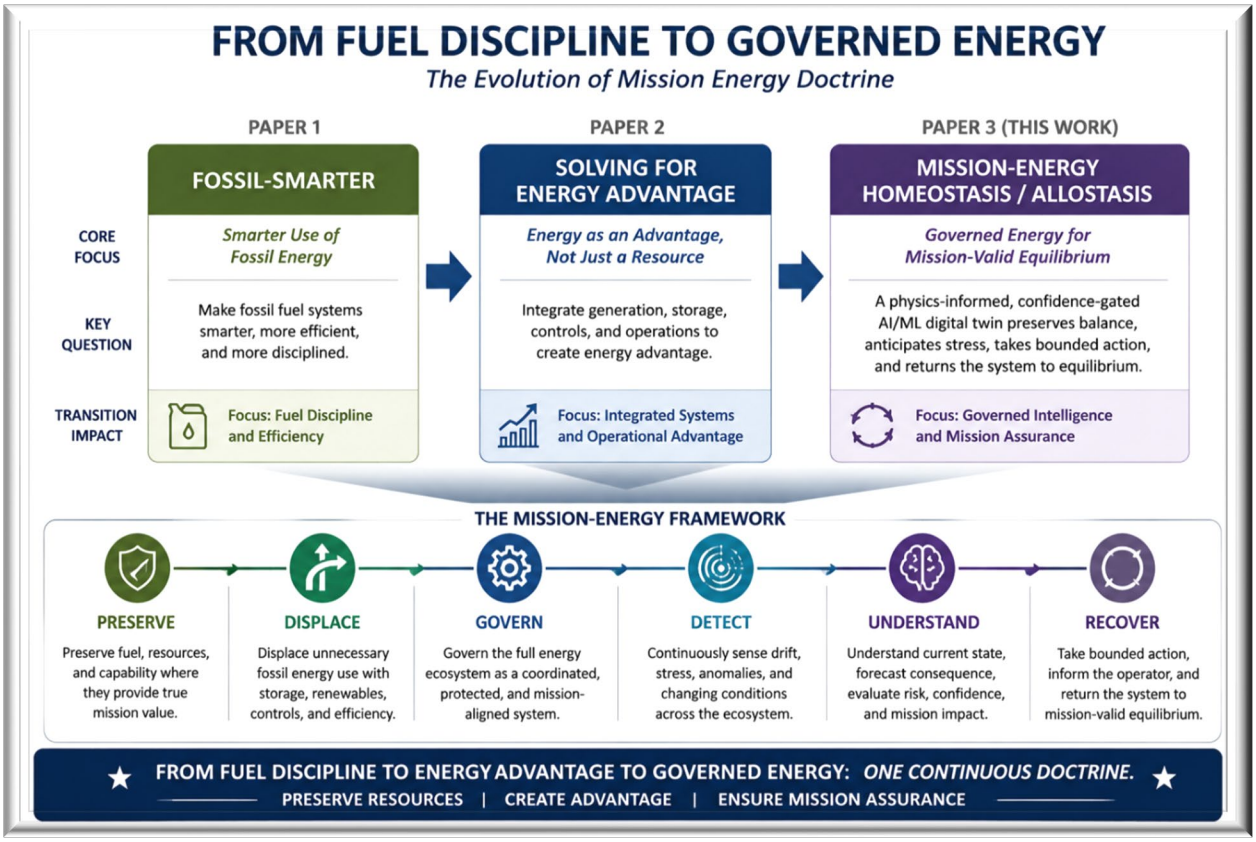

4.1 From Fuel Discipline to Governed Energy

Fossil-Smarter began with a practical truth: the future of energy is not built by pretending fuel disappears from military, industrial, or critical-infrastructure systems overnight. Fuel remains strategically necessary, especially where logistics are contested, grid access is unreliable, mission loads are mobile, and operational timelines cannot wait for ideal infrastructure.

The issue is not whether fuel still matters. It does.

The issue is whether fuel is burned with discipline.

That argument can be summarized in three words:

Preserve. Displace. Govern.

Preserve fuel where it provides true mission value. Displace unnecessary runtime where storage, solar, controls, load shaping, or better dispatch can reduce waste. Govern the full energy ecosystem so generation, storage, distribution, telemetry, cybersecurity, human authority, and mission priorities operate as one coordinated system.

Solving for Energy Advantage extended that argument beyond fuel discipline. It positioned energy not as isolated equipment, but as a governed ecosystem where generation, storage, power conversion, protected

distribution, telemetry, cybersecurity, human authority, recovery behavior, and grid interaction must be coordinated under stress.

This paper takes the next step. It defines the control doctrine required to keep that ecosystem stable, adaptive, trusted, recoverable, and mission-valid as conditions change.

4.2 Biological Logic Applied to Mission Energy

In biological systems, homeostasis preserves internal stability when conditions drift. Allostasis prepares the organism for expected stress before that stress causes damage. Mission energy systems face a similar challenge.

A tactical microgrid, installation microgrid, or distributed energy architecture must do more than produce power. It must sense when operating conditions are changing, determine whether those changes threaten mission function, take bounded corrective or anticipatory action, explain its behavior to the operator, and return the system to a stable mission-valid state.

In this doctrine, homeostasis is the corrective function. It detects when the energy system is drifting outside a validated operating state and initiates bounded action to restore stability.

Allostasis is the anticipatory function. It adapts the energy posture before forecasted stress becomes failure. This may include preserving reserve before a mission window, pre-charging storage before expected load growth, delaying generator start until an efficient recharge window, shaping discretionary loads, or coordinating distributed energy nodes before an installation-level stressor reaches critical load.

Together, Mission-Energy Homeostasis / Allostasis turns the microgrid from a collection of assets into a governed energy ecosystem.

4.3 Why Balance Is Not Enough

Traditional energy resilience often focuses on availability: how much generation exists, how much storage is installed, how much fuel is on hand, and how long the system can operate under assumed conditions. Those measures matter, but they are incomplete.

A system can have power and still fail to protect the right load.

A system can have storage and still consume reserve too early.

A system can have a generator and still waste fuel through inefficient runtime.

A system can have telemetry and still leave the operator without decision-grade understanding.

A system can have automation and still become unsafe if authority is not bounded.

Mission-energy balance is different. It requires the system to preserve the correct operating state for the active mission context. That means protecting priority loads, preserving reserve floors, maintaining power quality, respecting thermal and protection limits, accounting for cybersecurity confidence, informing the operator, and retaining a recovery path.

Balance is not simply electrical stability. It is mission-valid equilibrium.

4.4 The Central Thesis

A physics-informed, confidence-gated AI/ML digital twin can preserve energy balance across tactical and installation-scale microgrids by sensing drift, modeling consequence, taking bounded corrective or anticipatory action, and returning the system to mission-valid equilibrium while preserving deterministic protection and human authority.

The better question is no longer only:

How much power is available?

The better question is:

Can the energy ecosystem preserve mission-critical function under changing conditions, with enough reserve, enough confidence, enough control authority, and a clear recovery path?

4.5 Detect, Understand, Act, Inform, Return

The AI/ML-DT governs the energy ecosystem through five recurring functions.

It must detect drift, stress, abnormal behavior, degraded confidence, cyber-relevant anomalies, or emerging risk.

It must understand the condition through telemetry, the digital twin, mission context, operating limits, confidence scoring, and forecasted consequence.

It must act only inside validated authority, or recommend action when human approval is required.

It must inform the operator with clear explanation, confidence level, consequence, action status, authority boundary, and recovery path.

It must return the system to mission-valid equilibrium by preserving priority loads, protecting reserve, stabilizing the system, restoring function in sequence, and capturing evidence for after-action review and revalidation.

This is the core operating logic of Mission-Energy Homeostasis / Allostasis.

The future of tactical and installation energy resilience will not be defined only by how much power is installed. It will be defined by how well that power is governed.

5. Prior Foundations: From Fossil-Smarter to Energy Advantage

5.1 Integrated Progression of the Three Papers

The three papers form a deliberate progression:

Paper	**Core Argument**
Fossil-Smarter	Fuel is too valuable to waste. Preserve. Displace. Govern.
Solving for Energy Advantage	Energy assets should operate as a coordinated ecosystem, not isolated equipment.

Paper	Core Argument
Mission-Energy Homeostasis / Allostasis	That ecosystem should self-regulate, anticipate, recover, and return to equilibrium under mission stress.

Mission-Energy Homeostasis / Allostasis is the behavioral framework that allows the AI/ML-DT to govern the architecture described in the first two papers. The earlier papers also established that resilience must be proven through measured behavior — generator runtime avoided, fuel preserved, starts prevented, reserve protected, critical loads maintained, recovery time reduced, and mission function preserved.

That principle carries directly into this paper. **Trust is earned, not given.**

6. Core Requirements

6.1 Engineering-Control Doctrine Requirement

Mission-Energy Homeostasis / Allostasis must be treated as an engineering control doctrine, not a loose autonomy concept. Requirements must be explicit, testable, bounded, measurable, and tied to mission consequence.

6.2 Mission-Load Continuity Requirement

The architecture must identify, protect, and prioritize the loads that preserve mission function. Critical loads, mission-supporting loads, deferrable loads, and shed-ready loads must be defined before the system is placed in operation.

6.3 Deterministic Protection Precedence Requirement

Protective relaying, hardwired interlocks, inverter and PCS inner-loop protections, BMS safety limits, anti-islanding safeguards, emergency stops, and local safe-state logic remain authoritative. The AI/ML-DT operates above these layers and must not override protective intent.

6.4 Bounded Supervisory Authority Requirement

The AI/ML-DT must have a clearly defined authority partition: deterministic only, advisory, bounded autonomous (after validation), and human-approval required. Authority must follow evidence and must never expand through software update alone.

6.5 Validated Operating Envelope Requirement

Every autonomous or semi-autonomous action must remain inside tested left-right boundaries. When the system reaches the edge of its validated envelope, it must constrain itself, fall back, or escalate.

6.6 Confidence-Gated Control Requirement

The AI/ML-DT must evaluate telemetry quality, model synchronization, scenario validity, communications status, control authority, reversibility, consequence, and mission-rule compliance before recommending or executing action. Confidence must be tied to measurable data quality, calibration, and tested behavior — not treated as a vague assertion.

6.7 Human Authority Requirement

The operator must understand what the system sees, what it believes is happening, what action it recommends or has taken, what confidence exists, what boundaries apply, and when approval is required. The system reduces operator burden without hiding authority, consequence, or accountability.

6.8 Deterministic Fallback Requirement

The system must remain safe when prediction is wrong, telemetry is incomplete, communications degrade, cyber confidence weakens, or the operating condition exceeds the validated envelope. Fallback behavior may include local PLC/RTU logic, fixed protection settings, conservative load shedding, reserve-floor preservation, generator safe mode, BESS protection, and manual control.

6.9 Cybersecurity and Data Trust Requirement

The AI/ML-DT must rely only on authenticated telemetry, authorized commands, traceable configuration baselines, time-synchronized event records, and controlled interfaces. Role-based access control, cryptographic integrity protection, tamper-evident logging, and signed firmware and model artifacts are design requirements.

6.10 Model Governance and Configuration-Control Requirement

The system must track model versions, training data, firmware, controller settings, relay settings, BMS parameters, and mission-rule updates. Any change affecting behavior must trigger review and, where required, revalidation.

6.11 AI/ML-DT Assurance Requirement

The AI/ML-DT itself must be treated as a system under test. Predictions, classifications, confidence scores, recommendations, autonomous actions, refusal-to-act behavior, escalation decisions, fallback transitions, and explanations must be validated under representative disturbances.

6.12 Recovery and Equilibrium Return Requirement

The system must support a controlled recovery cycle: stabilize, isolate, rebalance, restore, learn, and return to mission-valid equilibrium. Recovery performance is measured through critical-load continuity, reserve preservation, isolation time, recovery-sequence accuracy, fallback success, and equilibrium return time.

6.13 Evidence Discipline Requirement

Every meaningful event must produce an auditable record capturing telemetry state, model state, confidence level, recommendation logic, action taken, human approval or override, fallback behavior, load impact, recovery sequence, and final equilibrium status.

6.14 Scalable Deployment Maturity Requirement

The architecture matures through a crawl-walk-run pathway, beginning with observation and advisory support and progressing to bounded autonomous actions only after sufficient test evidence exists.

6.15 Standards-Aware Traceability Requirement

Each deployment should build a candidate standards and certification path without casually claiming compliance before evidence exists.

6.16 Limitations Honesty Requirement

The doctrine must state what it does not claim: no universal autonomy, no AI replacement of protection engineering, no uniform authority across all installations, no trust outside a validated envelope.

7. Reference Architecture: AI/ML-DT as the Supervisory Nervous System

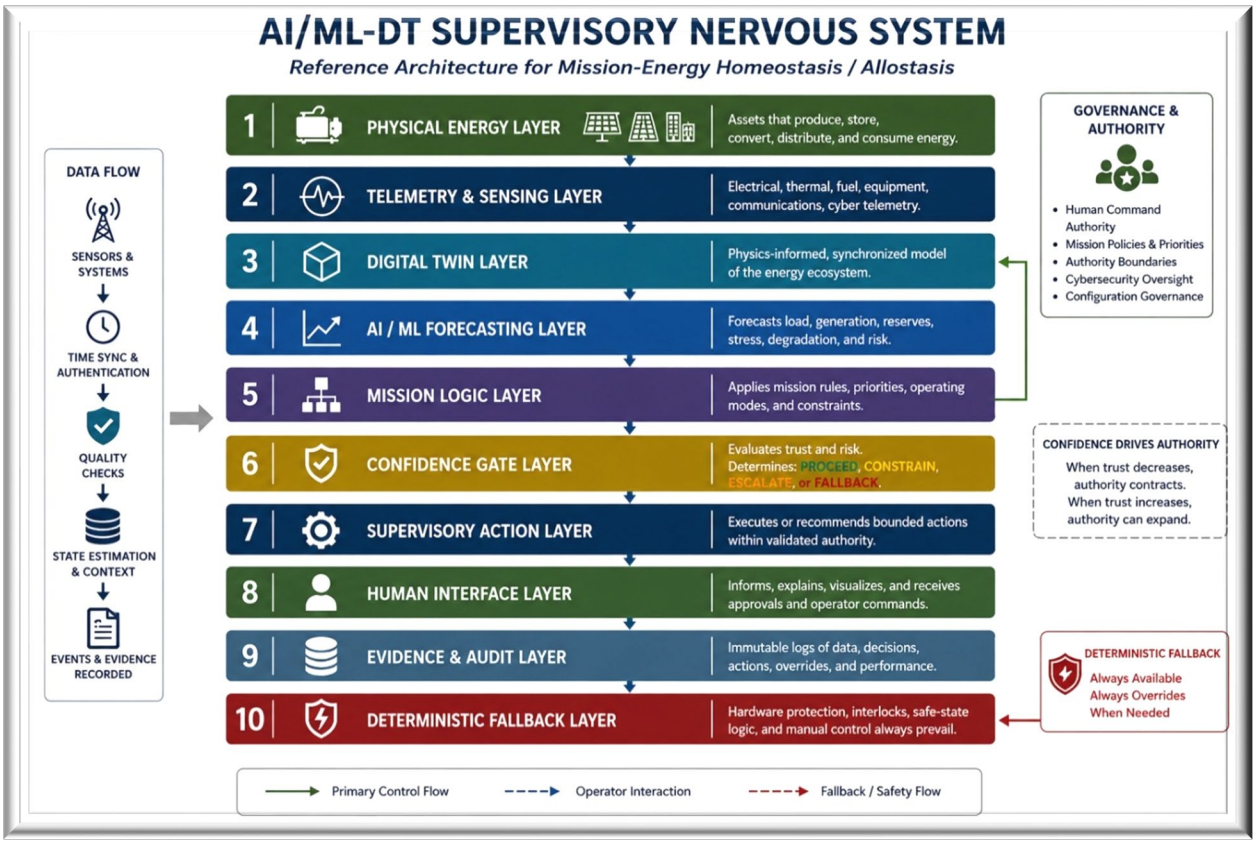

7.1 Architecture Purpose

If the physical microgrid is the body of the energy ecosystem, the AI/ML-DT is its supervisory nervous system.

Generators, batteries, inverters, PV inputs, switchgear, relays, sensors, communications pathways, protected distribution, and mission loads form the physical architecture. The AI/ML-DT provides the supervisory intelligence that allows the energy ecosystem to understand its own condition, anticipate stress, recommend or execute bounded action, inform the operator, and return to mission-valid equilibrium.

The AI/ML-DT continuously answers four operational questions:

What is the system doing now?

What is the system likely to do next?

What action best preserves mission-energy balance?

Is the system confident and authorized enough to act, or must it notify the human operator?

The architecture is not designed to replace protection engineering, inverter controls, BMS protections, relay logic, or human authority. It is designed to operate above those deterministic layers as a bounded supervisory control and decision-support layer.

7.2 Architecture Logic

The reference architecture follows a simple governing logic:

Telemetry becomes state.
State becomes consequence.
Consequence becomes bounded action.
Bounded action becomes evidence.
Evidence becomes trust.

This sequence is central to the doctrine. Raw data alone is not enough. The system must convert measured electrical, thermal, fuel, reserve, cyber, communications, and load conditions into decision-grade understanding. That understanding must then be filtered through mission rules, confidence gates, deterministic protection boundaries, and human authority before any recommendation or action is permitted.

7.3 Architecture Layers

Physical Energy Layer — Generators, BESS modules, PV input, inverter/PCS equipment, protected output panels, load interfaces, thermal controls, fuel systems, and local safety devices. The AI/ML-DT must be built around this layer, not imagined above it. Physical constraints define the control reality.

Sensing and Telemetry Layer — Collects voltage, frequency, current, phase balance, power quality, SOC, SOH, generator output, fuel state, inverter loading, PV production, thermal condition, breaker status, relay events, load demand, and communications quality. Telemetry must be authenticated, time-aligned, quality-checked, and tagged with source identity. This layer answers: **What is happening now?**

Digital Twin Layer — Transforms telemetry into operational context. It represents the actual power architecture, including equipment limits, topology, control modes, load hierarchy, reserve floors, thermal limits, protection boundaries, communications status, cybersecurity assumptions, and mission rules. It also detects model/plant mismatch. This layer answers: **What does the current state mean?**

AI/ML Forecasting Layer — Estimates what is likely to happen next, including load growth, reserve depletion, generator underperformance, battery degradation, thermal stress, communications degradation, inverter overload, or future mission demand. Forecasts must include confidence bounds, calibration history, data-quality status, and scenario-validity constraints. This layer answers: **What is likely to happen if nothing changes?**

Policy and Mission Logic Layer — Applies load-priority rules, reserve-floor requirements, operating modes, signature constraints, fuel rules, cybersecurity posture, safety boundaries, operator approval rules,

recovery logic, and site-specific command authority. This layer answers: **What should the system do, and why?**

Confidence Gate — The trust checkpoint before recommendation or action. The system evaluates telemetry reliability, digital twin synchronization, scenario validity, communications adequacy, action authorization, deterministic protection status, cybersecurity posture, reversibility, mission consequence, and operator role. When confidence degrades, autonomous authority must shrink. This layer answers: **Is the system confident and authorized enough to act?**

Supervisory Action Layer — Implements approved actions through bounded control pathways. Actions may include generator start/stop recommendations, recharge scheduling, reserve preservation, demand shaping, noncritical load shedding, inverter dispatch adjustments, thermal-management responses, support-node coordination, and recovery sequencing. Higher-consequence actions require human approval unless specifically validated and pre-authorized. This layer answers: **What bounded action is permitted now?**

Human-Machine Interface Layer — Preserves operator authority by explaining condition, consequence, confidence, action taken, recommended action, approval requirements, fallback status, reserve impact, and recovery path. A useful HMI does not say only "Generator fault." It says:

"Generator 2 output is degrading relative to expected response. Local BESS can support priority loads for 42 minutes at current demand. Neighboring Node 3 has reserve available. Recommended action: isolate Generator 2 from dispatch, shed noncritical Load Group C, request Node 3 support, and notify maintenance. Human approval required for cross-node support."

Evidence and Audit Layer — Records every meaningful system event, including telemetry state, model state, confidence level, recommendation logic, action taken, operator approval or override, fallback behavior, load impact, recovery sequence, cybersecurity state, and equilibrium-return status. This layer provides the evidence required for after-action review, troubleshooting, model improvement, revalidation, and future authority expansion.

Deterministic Fallback Layer — Protective relaying, fixed safety limits, BMS protection, inverter protection, local PLC/RTU logic, hardwired interlocks, conservative load shedding, manual control, and predefined recovery states. Fallback is not separate from trust. It is one of the reasons trust is possible.

7.4 Data and Control Flow

The AI/ML-DT does not control the energy ecosystem by jumping directly from prediction to action. It follows a governed flow.

Telemetry enters the system from authenticated sensors, controllers, relays, meters, inverters, BMS devices, generators, fuel systems, thermal sensors, communications nodes, and operator inputs. The digital twin converts this telemetry into a synchronized model of the current system state. The AI/ML layer forecasts likely near-term behavior and identifies risk. The policy and mission logic layer determines which outcomes matter most for the active mission context. The confidence gate determines whether the system has enough trust, authority, and scenario validity to recommend or execute action.

If confidence is sufficient and the action is pre-approved, the supervisory action layer may execute bounded control. If confidence is partial, the system may provide an advisory recommendation. If confidence is insufficient, the system constrains itself, falls back, or escalates to the operator.

Every action, recommendation, refusal-to-act, fallback transition, and operator override is recorded as evidence.

7.5 Architecture Control Cycle

The architecture operates through a recurring control cycle:

Detect → Understand → Decide → Act → Inform → Fallback if needed → Return

The system detects changes in physical state, mission demand, confidence, communications, cyber posture, or equipment behavior.

It understands the consequence by comparing current state against the digital twin, mission rules, reserve requirements, safety limits, and forecasted demand.

It decides whether action is permitted, advisory, blocked, or requires human approval.

It acts only within validated authority.

It informs the operator in plain operational terms.

It falls back when confidence, communications, cybersecurity, or operating conditions are insufficient.

It returns the energy ecosystem to mission-valid equilibrium and records the evidence needed to maintain or adjust trust.

7.6 Tactical and Installation-Scale Applicability

At the MPTaPS level, the reference architecture supports silent-watch preservation, runtime compression, reserve-floor protection, efficient recharge windows, load prioritization, generator underperformance detection, and reduced operator burden.

At the LOTaPS or multi-node level, it supports distributed-node coordination, shared reserve awareness, cross-node support logic, staged recovery, signature-aware operating posture, and local fallback when communications degrade.

At installation scale, it supports mission-load mapping, distributed energy node coordination, feeder-aware recovery logic, transformer and protection-boundary awareness, cybersecurity-informed supervisory control, and human approval for high-consequence switching or reconfiguration.

The architecture scales because the doctrine remains the same: understand the real system, preserve mission-valid equilibrium, bound authority, protect deterministic control, inform the human operator, and earn trust through evidence.

8. Authority Partition and Validated Operating Envelope

8.1 Central Principle

Autonomy must never exceed validated trust.

The AI/ML-DT is a bounded supervisory layer operating above deterministic protection and below human command authority. Its authority exists only inside a tested operating envelope supported by evidence, configuration control, cybersecurity discipline, and approved mission rules.

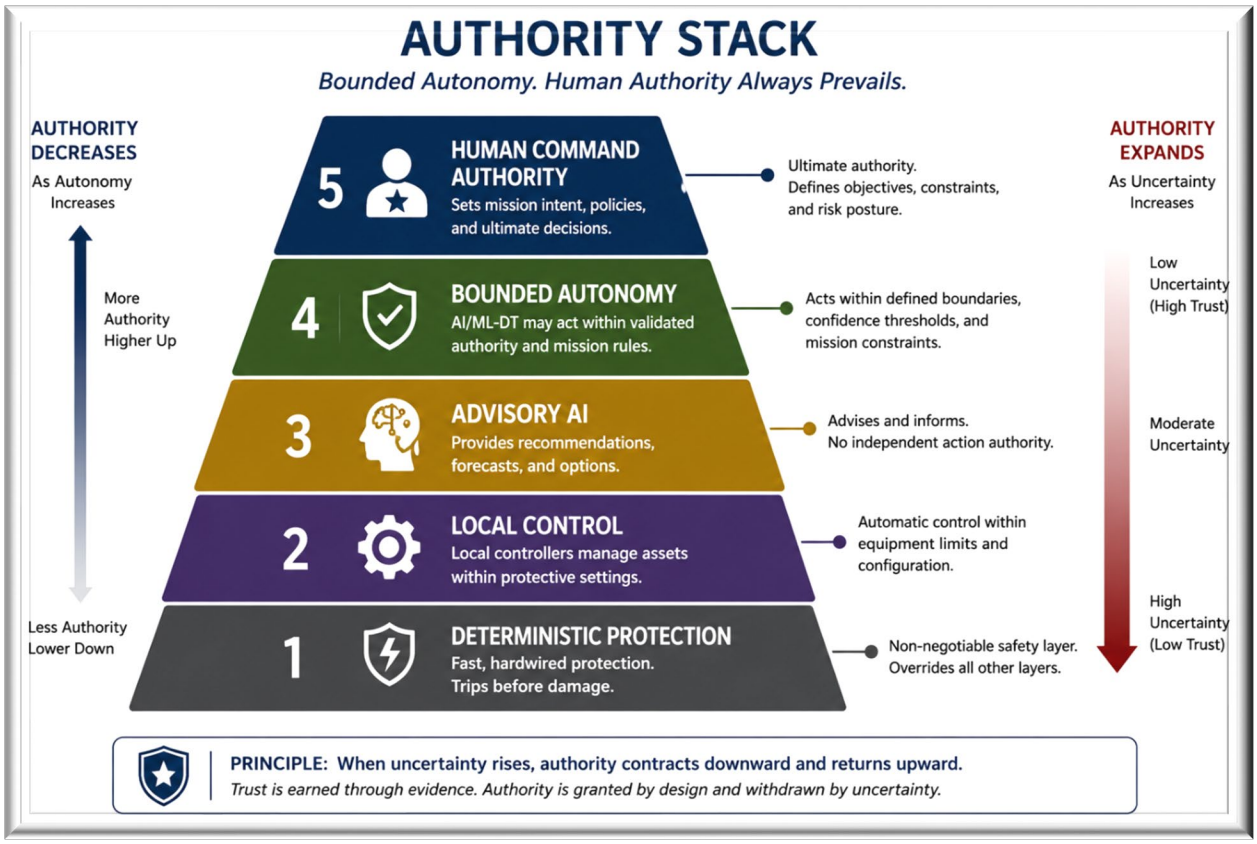

8.2 The Five Authority Layers

Layer 1 — Deterministic Protection Authority *(always authoritative)* Protective relays, inverter protection logic, BMS safety protections, hardwired interlocks, breaker trip logic, emergency-stop systems, anti-islanding protections, thermal shutdown logic, arc-flash protections, overcurrent protection, and undervoltage protection. The AI/ML-DT must not bypass or override this layer.

Layer 2 — Local Deterministic Control Authority Inverter inner-loop controls, generator governors, voltage regulators, PCS response logic, local PLC control, thermal-management systems, local battery balancing, and predefined safe-state transitions. The AI/ML-DT may request bounded state changes through approved interfaces but does not directly replace these loops.

Layer 3 — Supervisory Advisory Authority *(preferred initial deployment mode)* The AI/ML-DT may observe, forecast, classify, prioritize, recommend, and explain — but not autonomously execute high-consequence actions. The operator retains approval authority. Examples include reserve-preservation

recommendations, predicted load-stress alerts, recharge-window recommendations, and proposed load-shedding sequences.

Layer 4 — Bounded Supervisory Autonomous Authority Permits limited autonomous action only inside a validated operating envelope. Actions must be pre-approved, reversible where possible, consequence-bounded, test-validated, cybersecurity-reviewed, and tied to explicit mission rules. Examples include starting a recharge cycle, transitioning between predefined operating modes, preserving reserve floors, shedding designated noncritical loads, and managing silent-watch preservation.

Authority must shrink under uncertainty, not expand.

Layer 5 — Human Command Authority *(always final authority)* The operator, commander, or facility authority retains override control, approval authority, emergency intervention authority, and mission-priority authority. Actions that always require human approval include: cross-feeder reconfiguration, mission-priority changes, reserve-floor override, black-start initiation, major load-shedding actions, parallel-node isolation, actions affecting life-safety systems, and operation outside validated boundaries.

8.3 Envelope Expansion Philosophy

The architecture earns greater authority through evidence, following a controlled maturity pathway:

Observation → Advisory Support → Limited Bounded Autonomy → Expanded Bounded Autonomy → Coordinated Multi-Node Supervisory Behavior

This mirrors aviation, industrial controls, nuclear operations, and other high-consequence engineering domains where autonomy matures gradually through validation, operator trust, and demonstrated reliability.

9. Cybersecurity, Data Trust, and Model Governance

9.1 Foundational Principle

Cybersecurity and data trust are not supporting features of the architecture. They are foundational operating requirements.

A self-regulating energy ecosystem cannot govern itself safely if it cannot trust its own inputs, commands, configurations, models, time references, or communications pathways. The AI/ML-DT must assume degraded, manipulated, stale, incomplete, conflicting, or malicious data can exist at any time.

Cyber uncertainty is operational uncertainty.

When trust weakens, the system must become more conservative, not more autonomous.

9.2 Data Trust as a Control Input

Telemetry is not trustworthy simply because data exists. Every data stream used by the AI/ML-DT must be treated as a control-relevant input with an assigned trust state.

Voltage, frequency, current, SOC, SOH, generator output, fuel state, relay status, inverter loading, breaker position, thermal condition, communications quality, and operator commands must be authenticated, time-aligned, source-tagged, quality-checked, and tied to a known configuration baseline.

The system should evaluate:

Data source identity.

Time synchronization quality.

Message integrity.

Sensor health.

Configuration consistency.

Expected-versus-observed behavior.

Communications latency.

Model/plant synchronization.

Cybersecurity posture.

When one or more of these factors degrades, the AI/ML-DT must reduce confidence and limit its authority.

9.3 Cyber-Aware Confidence Gating

The confidence gate must include cybersecurity and data-trust conditions, not only electrical or performance conditions.

The AI/ML-DT should not recommend or execute autonomous action unless it can establish sufficient confidence in the data, model, authority pathway, control interface, and operating context. Unexpected command behavior, stale telemetry, mismatched configuration files, unauthorized parameter changes, abnormal traffic patterns, timing irregularities, unexplained sensor disagreement, or model/plant divergence must reduce confidence.

A cyber-aware confidence gate should produce one of four outcomes:

Proceed — The system has sufficient confidence, authority, and scenario validity to recommend or execute bounded action.

Constrain — The system may continue operating, but with reduced authority, narrower control options, or conservative reserve posture.

Escalate — The system requires human review or approval before action.

Fallback — The system exits supervisory autonomy and transitions to deterministic safe-state, advisory-only mode, local control, or manual operation.

9.4 Degraded-Trust Operating Modes

The architecture must define what happens when trust degrades. “Cybersecurity alert” is not an operating mode. The system must have predefined degraded-trust behaviors.

Possible degraded-trust modes include:

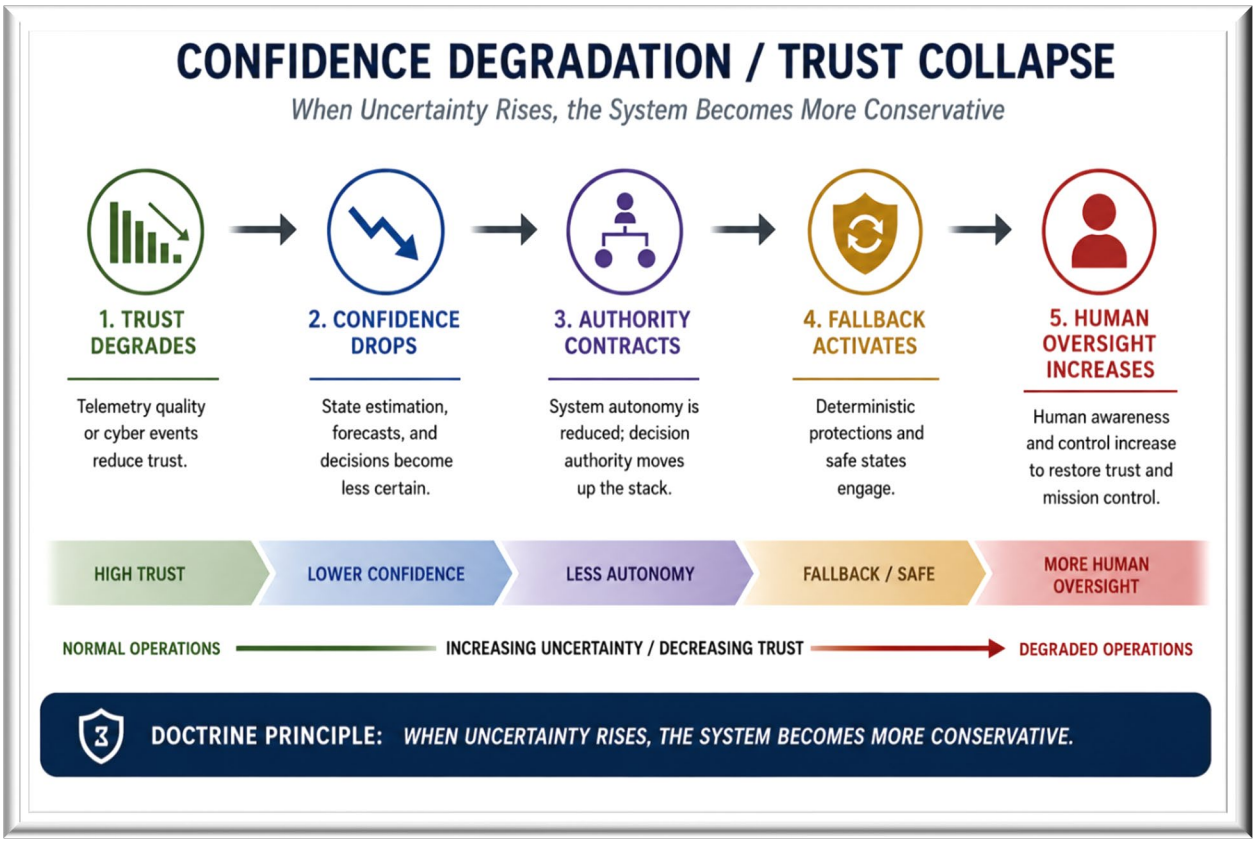

Advisory-Only Mode — The AI/ML-DT continues observing, forecasting, and explaining, but does not execute autonomous actions.

Local Safe Mode — The affected node relies on local deterministic controls, protection settings, reserve floors, and manual operator actions.

Quarantine Mode — A sensor stream, interface, model instance, controller, communications path, or distributed node is isolated or excluded from supervisory decision logic until trust is restored.

Rollback Mode — The system returns to a previously approved firmware, model, configuration, mission-rule set, or controller setting.

Manual Control Mode — Human operators assume direct control where supervisory authority is no longer trusted.

Conservative Reserve Mode — The system raises reserve floors, reduces discretionary loads, delays nonessential actions, and protects mission-critical loads until confidence is restored.

The point is not to prevent every cyber event. The point is to preserve mission-valid energy behavior despite uncertainty.

9.5 Segmentation and Authority Isolation

The AI/ML-DT must not operate as a flat, unrestricted control entity with broad access across the energy ecosystem. Distributed architecture increases survivability only if segmentation prevents one compromised node, interface, or model instance from propagating effects across the system.

Segmentation should separate monitoring, advisory functions, control execution, configuration management, model updates, operator access, remote maintenance, and external communications. Command pathways should be explicitly authorized, logged, and constrained by role, mission state, control authority, and validated operating envelope.

Authority isolation ensures a degraded or compromised function cannot automatically become a system-wide control problem.

9.6 Model Governance and Configuration Control

The AI/ML-DT is only as trustworthy as the model, data, software, configuration, and assumptions behind it.

Every model must have version traceability, training-data provenance, calibration history, configuration traceability, test history, validation boundaries, operational approval status, rollback capability, and authorized deployment records.

Configuration control must cover:

AI/ML model versions.

Digital twin assumptions.

Training and validation data sets.

Firmware.

Controller logic.

Relay settings.

BMS parameters.

Inverter and PCS settings.

Mission rules.

Load hierarchy.

Reserve floors.

Cybersecurity policy.

Communications pathways.

Operator-role permissions.

Any change affecting system behavior must trigger review and, where required, revalidation. Authority must not expand through software update alone.

9.7 Signed and Controlled Software Artifacts

Firmware, control logic, configuration files, relay settings, AI models, digital twin updates, cybersecurity policies, and mission-rule files must be signed, version-controlled, reviewable, and recoverable.

The system should be able to answer:

What version is running?

Who approved it?

When was it deployed?

What changed?

What operating envelope does it support?

What authority does it allow?

What rollback option exists?

What evidence supports its use?

Without those answers, the system may still function, but it should not be trusted with expanded supervisory authority.

9.8 Recovery of Trust

The system must define how trust is restored after degradation.

Full supervisory autonomy should not automatically resume just because communications return or an alert clears. Trust restoration should require reconciliation of telemetry state, time synchronization, configuration baseline, model state, event logs, command history, operator actions, and cybersecurity posture.

Depending on severity, the system may require operator acknowledgement, engineering review, configuration rollback, model recalibration, functional test, or formal revalidation before returning to full authority.

Trust must degrade quickly when evidence weakens and recover deliberately when evidence supports recovery.

9.9 Standards-Aware Cybersecurity Alignment

The architecture should be developed with awareness of operational technology and information-system cybersecurity frameworks, including NIST SP 800-82, NIST SP 800-53, IEC 62443, DoDI 8510.01, and applicable RMF processes.

These references do not create automatic compliance. They provide a traceability path for cybersecurity engineering, control selection, system categorization, assessment, authorization, monitoring, incident response, and lifecycle governance.

The objective is not paperwork compliance alone. The objective is a cyber-aware control architecture that preserves mission-energy behavior under degraded trust conditions.

9.10 Cybersecurity as Mission-Energy Survivability

Cybersecurity in this doctrine is not limited to confidentiality, access control, or network defense. It is part of mission-energy survivability.

A compromised or uncertain data stream can lead to the wrong generator start, the wrong load shed, the wrong reserve decision, the wrong recovery sequence, or the wrong confidence level. In a high-consequence energy system, bad data can become bad control.

For that reason, the AI/ML-DT must treat cybersecurity, data trust, model governance, fallback behavior, and human authority as inseparable elements of safe supervisory control.

The system earns trust by proving it knows when to act, when to ask, when to constrain itself, when to fall back, and when not to trust itself.

10. Planning and Integration

10.1 Foundational Planning Principle

The digital twin must first understand the real twin.

Mission-Energy Homeostasis / Allostasis cannot be added as an afterthought. It must be engineered into the architecture from the beginning.

10.2 Eight-Step Planning Process

Step 1 — Mission-Function Analysis Identify which loads are critical, mission-supporting, deferrable, and shed-ready. The first question is not how many kilowatts to install — it is which mission functions must survive, for how long, under what conditions, and with what recovery path.

Step 2 — Physical Mapping and One-Line Analysis Obtain or develop the electrical one-line. Transformers, feeders, switchgear, breakers, relays, meters, ATS equipment, panels, grounding systems, interconnection points, and final-mile distribution paths must be understood before the AI/ML-DT is trusted to model or recommend action.

Step 3 — Equipment Assessment Evaluate existing infrastructure for age, rating, condition, interoperability, protection behavior, data availability, and cyber suitability. The AI/ML-DT must not be asked to control equipment whose behavior or protection logic is not understood.

Step 4 — Interconnection and Distributed-Node Planning Position distributed energy nodes to support mission-relevant load pockets, reduce single-point vulnerability, and use existing internal infrastructure where technically appropriate.

Step 5 — Model Construction Build the AI/ML-DT around the actual architecture: sources, storage, conversion equipment, protected distribution, load hierarchy, interconnection points, reserve floors, thermal limits, fuel constraints, communications pathways, cybersecurity controls, operator rules, and deterministic fallback logic.

Step 6 — Authority Definition Define what the AI/ML-DT may observe, recommend, execute, or escalate. No fielded system should enter operation with vague control rights.

Step 7 — Communications and Cyber Integration Authenticate telemetry, authorize and log command pathways, control configuration baselines, sign and govern model updates, and retain safe local operation when communications degrade.

Step 8 — Phased Implementation

Phase	Behavior
Crawl	AI/ML-DT observes, models, compares, alerts, and supports operator insight without executing control actions.
Walk	AI/ML-DT recommends actions while operators approve or supervise execution.
Run	AI/ML-DT executes only bounded, pre-approved actions inside a validated operating envelope with full operator awareness and override authority.

11. Testing, Validation, and Revalidation

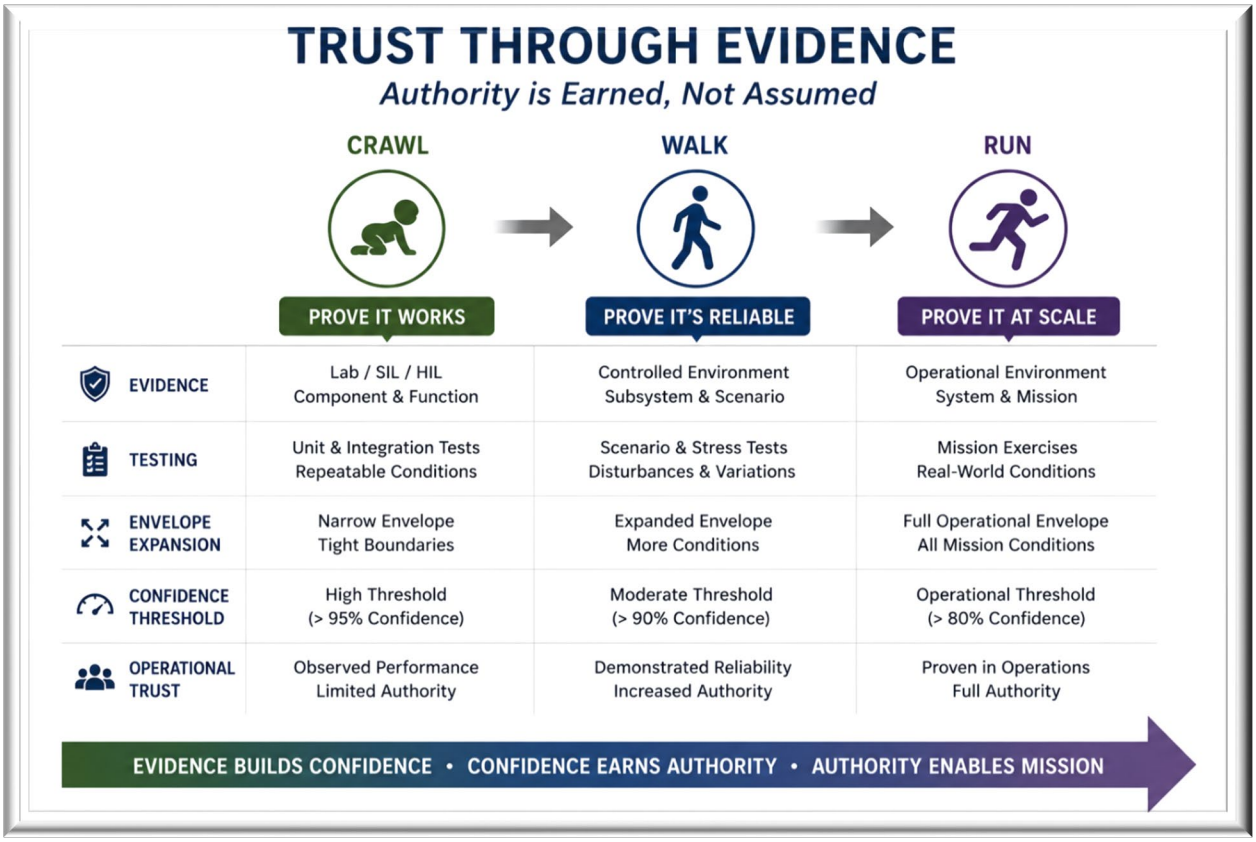

11.1 Foundational Validation Principle

The supervisory intelligence is part of the operational system and must therefore be validated like the operational system.

Mission-Energy Homeostasis / Allostasis cannot be accepted on the strength of concept diagrams, control logic descriptions, or assumed model accuracy. The doctrine must prove its behavior against real hardware, real telemetry, degraded operating states, representative mission loads, communications loss, cybersecurity uncertainty, operator intervention, and out-of-envelope conditions.

The purpose of validation is not simply to show that the AI/ML-DT can optimize energy use during normal operation. The purpose is to prove that it behaves correctly when conditions are uncertain, stressed, degraded, conflicting, or unsafe.

The architecture must demonstrate that it can detect drift, classify system condition, estimate consequence, preserve priority loads, protect reserve, recommend or execute bounded action, inform the operator, reduce authority under uncertainty, transition safely to fallback modes, recover from disturbance, and return to mission-valid equilibrium.

11.2 What Must Be Validated

Validation must apply to both the physical energy system and the AI/ML-DT supervisory layer. The test program should verify at least the following behaviors:

The system detects abnormal conditions within acceptable latency.

The digital twin remains synchronized with the physical system under normal and degraded conditions.

The confidence gate properly suppresses action when telemetry, model state, cybersecurity posture, or scenario validity is insufficient.

The AI/ML-DT does not override deterministic protection, inverter protections, BMS limits, relay logic, hardwired interlocks, or emergency-stop behavior.

The system preserves priority loads according to the approved mission-load hierarchy.

The system protects reserve floors during normal, stressed, and degraded operation.

The system distinguishes between actions it may execute, actions it may recommend, and actions requiring human approval.

The system explains condition, confidence, consequence, action status, authority boundary, and recovery path to the operator.

The system transitions to deterministic fallback when confidence, communications, data trust, or operating conditions fall outside validated limits.

The system captures auditable evidence sufficient for after-action review, troubleshooting, model improvement, and revalidation.

11.3 Progressive Validation Philosophy

The architecture must never begin with maximum autonomy. Authority must be earned through evidence.

During the crawl phase, the AI/ML-DT observes without autonomous action. It models equipment behavior, compares predicted response against actual response, measures telemetry quality, identifies

model/plant mismatch, and generates recommendations for operator review. The purpose is to build confidence in the digital twin and supervisory logic without allowing the system to affect operation.

During the walk phase, the AI/ML-DT begins participating in bounded supervisory functions under human oversight. It may recommend reserve preservation, recharge sequencing, generator runtime compression, recovery sequencing, noncritical load reduction, or degraded-mode operation, but operators remain the approval authority. The purpose is to validate decision logic, operator interaction, explanation quality, fallback behavior, and mission impact.

During the run phase, the AI/ML-DT executes only pre-approved bounded actions inside the validated operating envelope. Deterministic protection, local control logic, human override, cybersecurity controls, and fallback systems remain authoritative throughout. The purpose is not to remove the operator from the system. The purpose is to reduce operator burden while preserving accountability, authority, and safe recovery.

11.4 Test Conditions and Stress Cases

Validation must include normal operation and deliberate disturbance testing. Test conditions should include load-bank testing, generator start/stop sequencing, battery discharge and recharge cycles, inverter response, reserve-floor enforcement, priority-load preservation, noncritical load shedding, thermal stress, communications degradation, stale telemetry, corrupted or missing data, model/plant mismatch, unauthorized command attempts, operator override, and out-of-envelope scenarios.

For MPTaPS and LOTaPS applications, validation should also test runtime compression, silent-watch preservation, generator underperformance, distributed-node coordination, cross-node support logic, and signature-aware recharge timing.

For installation-scale applications, validation should include feeder constraints, transformer limits, protection coordination, distributed node failure, localized outage recovery, islanding behavior, restoration sequencing, and human approval of high-consequence switching actions.

11.5 Revalidation Triggers

Trust is not permanent. Any change that can alter system behavior must trigger review and, when required, revalidation.

Revalidation triggers include changes to hardware, firmware, inverter settings, relay settings, BMS parameters, load hierarchy, mission rules, reserve floors, communications pathways, cybersecurity configuration, AI model version, digital twin assumptions, topology, control authority, operating mode, or site infrastructure.

The system must also be revalidated after significant faults, unexplained behavior, cyber-relevant events, repeated model/plant mismatch, operator override trends, degraded performance, or expansion of autonomous authority.

11.6 Evidence Output

The final output of testing is not simply a pass/fail result. It is an evidence package.

That evidence package should document the tested operating envelope, confidence thresholds, authority boundaries, fallback behavior, recovery sequence, operator approval points, event logs, telemetry records, model performance, cybersecurity assumptions, unresolved limitations, and conditions requiring revalidation.

This evidence package becomes the bridge between doctrine and trust. It defines what the system has actually proven, what remains advisory only, what may be executed autonomously, and where human authority remains mandatory.

In this doctrine, trust is not declared. Trust is measured, bounded, documented, and maintained.

12. Concept of Operations (CONOPS)

12.1 Operational Foundation

The doctrine is built around one recurring sequence:

Detect. Understand. Act. Inform. Return.

12.2 CONOPS 1 — Standalone MPTaPS Silent-Watch Operation

The AI/ML-DT monitors battery state of charge, load demand, thermal state, inverter behavior, reserve floor, mission duration, and expected future demand. When reserve approaches the silent-watch floor, the system first evaluates whether noncritical loads can be deferred, discretionary charging reduced, or operating mode adjusted before recommending a generator start. If confidence is high and the action is pre-approved, the system may shed noncritical loads, preserve the reserve floor, and delay generator start until a validated recharge window — notifying the operator with projected endurance, reserve status, and recommended follow-on decisions.

12.3 CONOPS 2 — Two-Node MPTaPS Operation with Generator Underperformance

The AI/ML-DT detects abnormal generator behavior by comparing measured output, fuel behavior, voltage and frequency response, and recharge efficiency against expected performance. The digital twin determines whether Node A's BESS can carry the priority load, how long the ride-through window will last, and whether Node B can support load transfer without violating reserve limits. If communications between nodes degrade, Node A reverts to local reserve-preservation logic and notifies the operator rather than assuming external support.

12.4 CONOPS 3 — Installation-Scale Distributed Node Failure

One node experiences a PCS failure, generator fault, or feeder-isolating protection event. The AI/ML-DT detects the event, confirms equipment state, checks local BESS ride-through capacity, evaluates neighboring-node reserve, and determines whether validated pathways exist to support affected loads — checking feeder limits, transformer constraints, protection boundaries, cyber confidence, and operator-approval rules before acting. Events requiring feeder reconfiguration, major switching, or reserve-floor override escalate to human approval.

12.5 CONOPS 4 — Communications Degradation and Local Safe Mode

As communications quality degrades, the AI/ML-DT reduces confidence in cross-node coordination and disables cross-node autonomous actions. Local load-priority logic and reserve floors are maintained. Affected nodes shift into local-safe or advisory-only mode. Coordinated autonomy is not re-enabled until authenticated communications recover, time synchronization is verified, node states are reconciled, and pending actions are reviewed.

12.6 CONOPS 5 — High-Load Mission Window with Allostatic Reserve Prepositioning

Before a known mission event, the AI/ML-DT forecasts a high-load window that will stress the current posture if no anticipatory action occurs. Rather than waiting for failure, the system may recommend or execute pre-approved allostatic actions: pre-charging storage, preserving generator runtime, reducing discretionary loads, staging a support node, or locking a higher reserve floor before the event begins.

12.7 CONOPS 6 — Cyber or Data-Trust Degradation

The AI/ML-DT detects suspicious telemetry, unauthorized command attempts, unexpected configuration mismatch, or model/plant divergence. Autonomous authority is reduced. Affected interfaces, nodes, or sensor streams may be quarantined or downgraded. The system shifts to trusted local telemetry, deterministic fallback, advisory mode, or manual operator control depending on severity — treating cyber uncertainty as operational uncertainty. It returns to full coordinated autonomy only after data integrity, configuration state, time synchronization, and authorized control pathways are restored.

13. Recovery and Resilience

13.1 Defining Resilience

Resilience is the controlled ability to absorb disturbance, preserve mission-critical function, adapt under stress, recover in sequence, and return to mission-valid equilibrium.

Backup power alone does not accomplish this. A resilient system must understand what changed, preserve priority function, stabilize operating state, constrain consequence, reduce cascading effects, recover in sequence, and restore trust in system state.

13.2 Recovery Priority Sequence

When stress emerges, the architecture preserves in order:

1. Life and safety
2. Mission-critical loads
3. Command-and-control continuity
4. Protected reserve
5. System stability
6. Recovery capability
7. Broader restoration

13.3 Key Recovery Principles

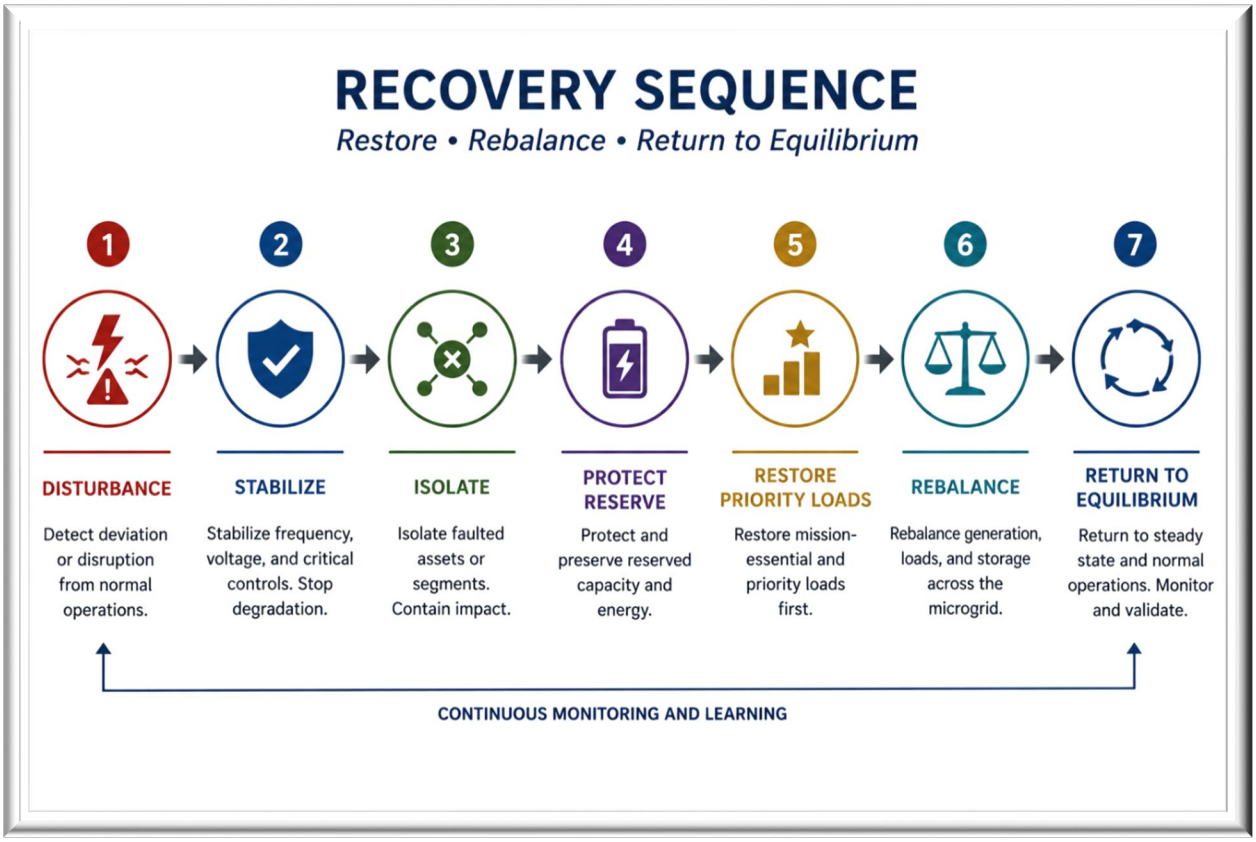

Stabilization — After recognizing disturbance, the architecture stabilizes the energy ecosystem before attempting full restoration by preserving reserve floors, isolating degraded equipment, maintaining frequency and voltage stability, and limiting distributed-node interaction.

Graceful Degradation — In conditions exceeding normal operating capability, lower-priority loads reduce first, reserve is preserved intentionally, nonessential functions defer, supervisory authority shrinks, and operators remain informed throughout the event.

When uncertainty rises, the system becomes more conservative. This is one of the defining principles of the doctrine.

14. Operational Impact

Mission-Energy Homeostasis / Allostasis reframes energy as a **governed survivability function**, delivering the following operational effects:

Tactical Endurance and Silent-Watch Extension — Longer mission endurance without proportional increases in fuel consumption or generator runtime by compressing generator runtime into efficient recharge windows, deferring noncritical demand, and pre-positioning energy before expected stress. The operational effect is increased survivability and mission duration.

Reduced Operator Burden — The AI/ML-DT converts raw telemetry into supervised operational understanding — evaluating consequence rather than simply displaying voltage, frequency, SOC, and

runtime — allowing operators to spend less time interpreting data and more time making mission decisions.

Improved Reserve Discipline — Reserve is treated as a strategic survivability margin rather than consumed reactively, especially critical for denied logistics, Arctic operations, expeditionary operations, and degraded fuel-access conditions.

Distributed Survivability — The AI/ML-DT provides the supervisory coordination that allows distributed nodes to behave as a governed ecosystem rather than isolated equipment, especially important for military installations, radar sites, and resilient communications infrastructure.

Recovery Acceleration — Earlier disturbance detection, intentional recovery sequencing, reduced secondary instability, and preserved operator awareness produce more stable restoration — not just faster restoration.

Cyber-Aware Operational Behavior — Cyber uncertainty becomes operational uncertainty. When telemetry trust weakens, the AI/ML-DT reduces authority and transitions toward conservative behavior, creating a survivability-oriented operational posture.

Signature-Aware Energy Management — The AI/ML-DT may preserve silent-watch reserve before surveillance windows, compress recharge cycles into lower-risk periods, and reduce discretionary load during sensitive operations — reducing predictability and improving operational discretion.

Mission-Energy as Command Freedom — The most important operational impact is conceptual: energy becomes a mission-enabling control function. Commanders and operators make decisions with improved awareness of endurance, reserve, survivability, recovery posture, confidence, and operational consequence. The architecture does not eliminate risk. It reduces uncertainty-driven fragility.

15. Standards and Certification Path

15.1 Certification Philosophy

Candidate alignment first. Compliance only after evidence.

Mission-Energy Homeostasis / Allostasis should be standards-aware from the beginning, but it must not casually claim compliance before testing, inspection, documentation, cybersecurity assessment, and approval-authority acceptance support that claim.

This doctrine does not replace established certification pathways for inverters, battery systems, generators, switchgear, protective devices, communications systems, cybersecurity controls, or installation electrical work. Instead, it provides a supervisory-control framework that must respect certified equipment limits, approved protection settings, safety requirements, interconnection rules, cybersecurity obligations, and authority boundaries.

The correct standard is not one universal checklist. Applicable standards depend on system scale, voltage class, storage chemistry, inverter type, interconnection method, military use case, installation authority, utility interface, data architecture, and cybersecurity categorization.

15.2 Standards as Evidence Pathways

Standards should be treated as evidence pathways, not decorative references. Each applicable reference should help answer a practical engineering question:

What safety requirement applies?

What control behavior must be tested?

What interface must be verified?

What cybersecurity control must be implemented?

What inspection or test record proves the claim?

What authority accepts the evidence?

This is especially important for the AI/ML-DT. The supervisory layer cannot be considered trusted because it is intelligent, predictive, or well-modeled. It becomes trusted only when its behavior is mapped to requirements, tested against representative conditions, bounded by authority, documented through evidence, and accepted by the responsible approval authority.

15.3 Candidate Standards Mapping

Domain	Candidate References	Evidence Contribution
Electrical safety and installation	NFPA 70 / National Electrical Code	Supports safe installation, conductor sizing, grounding, overcurrent protection, equipment clearances, wiring methods, and inspection basis.
Distributed energy interconnection	IEEE 1547-2018; IEEE 1547.1-2020	Supports DER interconnection requirements, interoperability, abnormal-condition response, and conformance test procedures.
Microgrid controller requirements	IEEE 2030.7-2017	Supports functional requirements for microgrid controllers, operating modes, transitions, dispatch, and coordination.
Microgrid controller testing	IEEE 2030.8-2018	Supports test procedures for controller behavior, transitions, resilience functions, and operating-mode verification.
Power quality and harmonics	IEEE 519-2022	Supports assessment of harmonic limits, power-quality impacts, and electrical compatibility.
Inverter / PCS equipment	UL 1741	Supports certification path for inverters, converters, controllers, and DER interconnection equipment.

Domain	Candidate References	Evidence Contribution
Energy storage system safety	UL 9540; UL 9540A; NFPA 855	Supports ESS safety, thermal-runaway evaluation, installation requirements, spacing, fire protection, and hazard mitigation.
Operational technology cybersecurity	NIST SP 800-82 Rev. 3; IEC 62443 Series	Supports OT cybersecurity architecture, segmentation, access control, monitoring, incident response, and control-system security.
Information-system cybersecurity	NIST SP 800-53 Rev. 5	Supports security and privacy control selection, implementation, assessment, and continuous monitoring.
AI risk management	NIST AI RMF 1.0	Supports AI risk framing, governance, measurement, management, transparency, explainability, and trustworthiness considerations.
DoD RMF / authorization	DoDI 8510.01	Supports cybersecurity categorization, risk management, assessment, authorization, and continuous monitoring for DoD systems.
Installation energy management	DoDI 4170.11	Supports alignment with DoD installation energy resilience, energy management, and mission assurance objectives.
Test and evaluation	DoDI 5000.89	Supports disciplined test planning, evaluation, operational relevance, data collection, and evidence-based decision making.
Power utility automation communications	IEC 61850	Supports structured communications for power utility automation, interoperability, and substation/grid automation concepts.
Time synchronization	IEC/IEEE 61850-9-3	Supports precision time protocol requirements for event ordering, synchronized telemetry, and power-system automation timing.

15.4 AI/ML-DT Assurance Path

The AI/ML-DT requires its own assurance path because it influences supervisory decisions. This does not mean the AI replaces certified hardware, protection engineering, or human authority. It means the AI/ML-DT must be evaluated as a control-relevant system element.

The assurance path should address:

Model purpose and intended use.

Training and validation data provenance.

Scenario limits.

Model/plant synchronization.

Confidence calibration.

Prediction accuracy.

False-positive and false-negative behavior.

Refusal-to-act behavior.

Out-of-envelope detection.

Cyber-aware confidence gating.

Human explanation quality.

Operator approval points.

Fallback transitions.

Event logging and auditability.

Revalidation triggers.

The purpose is not to certify "AI" in the abstract. The purpose is to prove the AI/ML-DT behaves safely, predictably, explainably, and conservatively within a defined mission-energy operating envelope.

15.5 Compliance Boundary

The paper should make a clear boundary between standards alignment and compliance.

Alignment means the architecture is designed with awareness of applicable standards, uses those references to shape requirements, and prepares evidence for later evaluation.

Compliance means the system has been tested, inspected, documented, assessed, and accepted by the responsible authority against applicable requirements.

Mission-Energy Homeostasis / Allostasis can claim candidate alignment at the doctrine stage. It cannot claim final compliance until the system has produced evidence through development, test, evaluation, cybersecurity assessment, inspection, and approval.

That distinction matters. It protects the doctrine from overclaiming and strengthens the credibility of the RDT&E path.

15.6 Certification Does Not Equal Mission Trust

Certification is necessary, but not sufficient.

An inverter can be certified and still be misapplied.

A battery system can meet safety requirements and still be operated without mission reserve discipline.

A microgrid controller can satisfy functional requirements and still lack mission-governed authority partitioning.

A cybersecurity control can be documented and still fail to preserve decision confidence under degraded data conditions.

Mission trust requires more than component compliance. It requires integrated system behavior under stress.

For that reason, the certification path must be paired with operational validation. The architecture must demonstrate critical-load continuity, reserve-floor protection, fallback behavior, cyber-aware authority reduction, operator approval logic, evidence capture, recovery sequencing, and equilibrium return.

15.7 Practical Certification Roadmap

The practical path should proceed in sequence:

First, identify applicable standards, approval authorities, inspection requirements, cybersecurity categorization, and operational constraints.

Second, map each standard or requirement to a specific design feature, control behavior, test method, inspection record, or evidence artifact.

Third, test component-level behavior, including inverter response, ESS safety behavior, generator response, protection logic, communications reliability, and deterministic fallback.

Fourth, test integrated system behavior under representative mission loads, degraded conditions, cyber-relevant uncertainty, operator override, and recovery scenarios.

Fifth, document the validated operating envelope, confidence thresholds, authority boundaries, fallback modes, known limitations, and revalidation triggers.

Sixth, submit evidence to the relevant engineering, cybersecurity, safety, installation, utility, or operational approval authority.

This path keeps the doctrine honest. It does not claim trust before evidence exists. It builds trust by turning standards into testable requirements, test results into evidence, and evidence into bounded operational authority.

16. Limitations and Boundaries

Mission-Energy Homeostasis / Allostasis is a bounded control doctrine, not a claim of universal autonomy. The following limitations are stated explicitly, as they strengthen the doctrine by demonstrating engineering discipline:

- **Deterministic protection remains authoritative.** The AI/ML-DT may supervise, recommend, forecast, and execute bounded actions where validated — but must not override protective intent.

- **Authority is mission-dependent.** Not every system should receive the same level of autonomous authority. Authority must be matched to mission consequence, system maturity, validation evidence, and site-specific risk.
- **Infrastructure limitations apply.** Legacy switchgear, aging transformers, undocumented modifications, or incomplete one-lines may restrict the AI/ML-DT to observation, monitoring, and advisory support until upgrades are completed.
- **Prediction is not certainty.** When confidence degrades, authority must shrink.
- **Resilience does not mean invulnerability.** Some disturbances may exceed available reserve, physical capacity, or repair capability. In those cases, the objective is controlled degradation, critical-load preservation where possible, and safe recovery.
- **Cybersecurity perfection is not claimed.** The architecture is designed around survivability under degraded conditions, not the prevention of all attacks.
- **Standards compliance requires evidence.** Candidate standards provide a traceability path; they do not create compliance without test results, inspection evidence, documentation, and approval authority acceptance.
- **Trust is envelope-bounded.** The AI/ML-DT is trusted only within tested configurations, known assumptions, governed models, verified data paths, and controlled authority boundaries. Hardware, firmware, mission rule, topology, or model version changes may require revalidation.

These limitations are not weaknesses. They are the reason the architecture can be trusted at all.

17. Deployment Roadmap

17.1 Deployment Philosophy

Test. Refine. Expand. Earn trust incrementally.

Mission-Energy Homeostasis / Allostasis must mature through disciplined deployment, not assumption. The system should begin by observing, modeling, comparing, and informing. It should then progress to advisory recommendations, bounded supervised execution, and eventually limited autonomous action only where evidence supports that authority.

The roadmap is not a software-release schedule. It is a trust-building pathway. Each phase should produce evidence that defines what the system may observe, recommend, execute, refuse, escalate, or fall back from.

17.2 Persistent Governance Bands

Three governance bands remain active across the full lifecycle.

Persistent Band 1 — Deterministic Protection and Safe-State Authority
Protective relays, hardwired interlocks, BMS protections, inverter and PCS inner loops, local safe-state logic, emergency stops, and manual override authority remain active in every phase. The AI/ML-DT may supervise or recommend, but it must not override deterministic protection.

Persistent Band 2 — Human Authority and Configuration Governance
HITL/HOTL boundaries, operator approval thresholds, escalation logic, signed releases, model governance, cybersecurity posture, configuration control, and authority partitioning remain visible across every phase. Human authority remains final where mission consequence, safety, cybersecurity, or out-of-envelope behavior requires judgment.

Persistent Band 3 — Evidence and Trust Accumulation
Testing, validation, revalidation, cybersecurity evidence, AI assurance evidence, operational replay, event logging, auditability, operator feedback, and lifecycle learning accumulate across the roadmap. Trust is an earned engineering condition, not a software assumption.

17.3 Six-Phase Maturity Roadmap

Phase	Name	Purpose	Primary Evidence Output
1	**Baseline**	Establish the known physical and operational system. Confirm electrical baseline, sensors, control inventory, one-line architecture, deterministic boundaries, mission-load hierarchy, reserve floors, communications pathways, cybersecurity posture, and operator authority.	Baseline system description, one-line reference, load hierarchy, authority boundary, configuration baseline, initial risk register.
2	**Crawl**	Build observation-only digital twin trust. The AI/ML-DT observes, models, compares, alerts, and supports operator insight without executing control actions.	State-estimation results, measured-versus-expected comparison, telemetry-quality assessment, model/plant mismatch log, confidence-calibration data.
3	**Walk**	Validate advisory control and operator interaction. The AI/ML-DT recommends actions while operators approve or supervise execution.	Advisory recommendation accuracy, operator response records, explanation-quality review, fallback validation, load-bank results, degraded-data test results.
4	**Run**	Enable bounded autonomous actions inside the validated operating envelope. Actions remain pre-approved, consequence-bounded, logged, reversible where possible, and subject to operator override.	Validated operating envelope, confidence thresholds, autonomous-action logs, unsafe-action suppression results, fallback transition evidence.

5	**Scale**	Expand to LOTaPS, multi-node MPTaPS, or installation-representative pilots. Validate distributed-node coordination, cross-node support logic, feeder-aware recovery, cybersecurity segmentation, and high-consequence human approval points.	Multi-node test data, distributed recovery sequence, cyber-degradation behavior, standards traceability, RMF/certification evidence package.
6	**Sustain**	Govern updates, revalidation, lifecycle learning, model drift, configuration changes, and authority maintenance. Trust is maintained only through continued evidence discipline.	Revalidation records, signed model/configuration releases, drift-monitoring reports, equilibrium performance trends, after-action review archive.

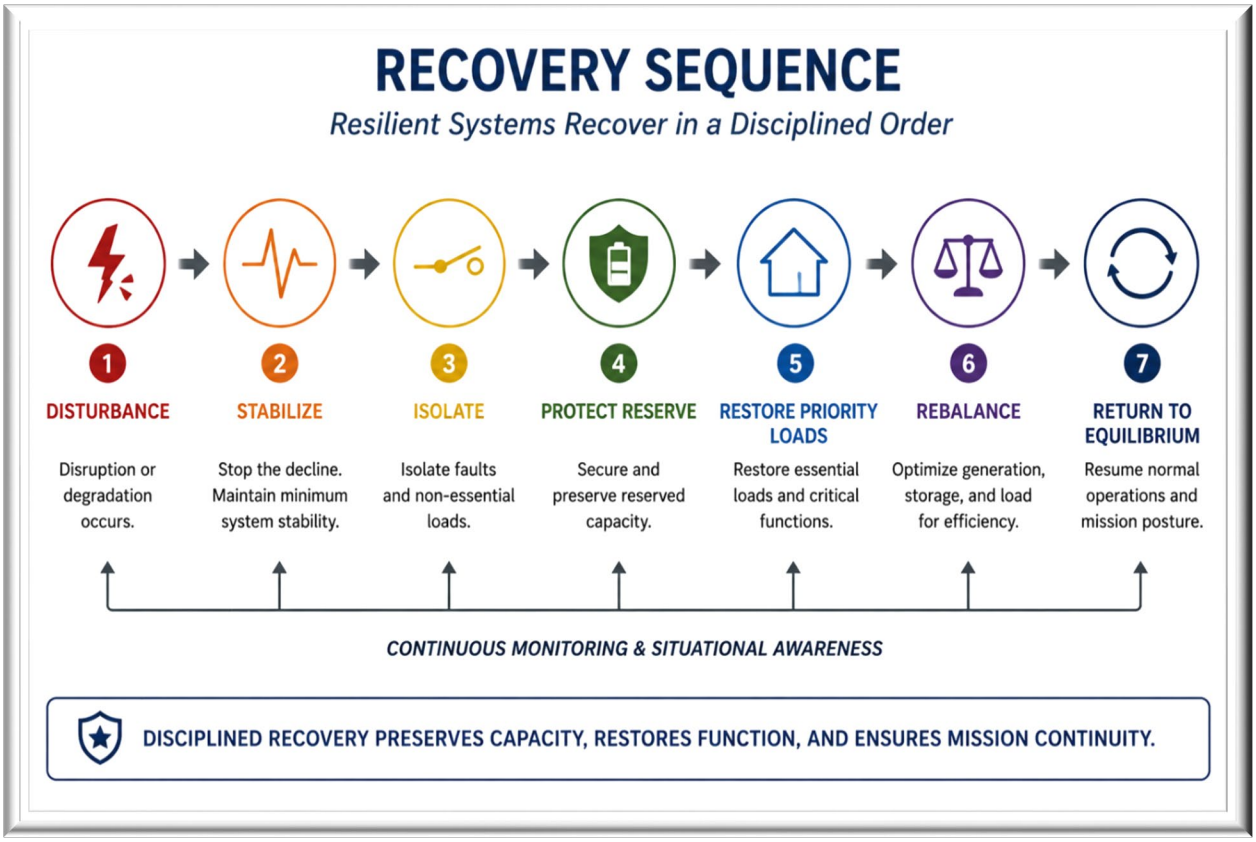

17.4 Deployment Entry and Exit Logic

Each phase should have clear entry and exit criteria. A system should not move from observation to advisory, from advisory to bounded action, or from single-node operation to coordinated multi-node behavior merely because the software is available.

Progression should require evidence that the previous phase performed correctly under representative conditions.

At minimum, phase advancement should consider:

Whether telemetry is reliable, authenticated, time-aligned, and sufficient.

Whether the digital twin accurately represents the physical system.

Whether the AI/ML-DT detects drift, mismatch, degraded confidence, and out-of-envelope conditions.

Whether operator explanations are clear enough to support decision-making.

Whether deterministic fallback works as designed.

Whether cyber-relevant uncertainty causes authority to shrink.

Whether reserve floors and priority-load rules are preserved.

Whether event logging supports audit, troubleshooting, and revalidation.

Whether operators understand when the system is advising, acting, escalating, or falling back.

17.5 Authority Expansion Rules

Authority expansion must be deliberate, documented, and reversible.

The AI/ML-DT should gain authority only when testing demonstrates that it can act safely, predictably, explainably, and conservatively within a defined operating envelope. Authority must not expand through software update alone. It must expand through approved evidence.

Authority should shrink automatically when telemetry trust, model synchronization, communications quality, cybersecurity posture, scenario validity, or operating conditions degrade. Expansion should require approval; contraction should be automatic.

This principle protects the architecture from the most dangerous failure mode in high-consequence autonomy: acting with confidence it has not earned.

17.6 Deployment Sequence by Use Case

For **MPTaPS**, deployment should begin with single-node observation, telemetry validation, load-bank testing, generator/BESS behavior characterization, runtime-compression modeling, reserve-floor validation, and advisory recharge-window recommendations. Bounded autonomous behavior should be limited to pre-approved actions such as noncritical load reduction, reserve-floor preservation, and recharge-window execution after test evidence supports the behavior.

For **LOTaPS**, deployment should extend the MPTaPS evidence base into larger-capacity operation, multi-node coordination, support-node logic, signature-aware energy posture, cross-node communications degradation, and coordinated fallback. LOTaPS should not assume greater authority simply because it has greater capacity.

For **installation-scale pilots**, deployment should begin with mission-load mapping, electrical one-line validation, feeder and transformer constraint analysis, protective-device coordination review, communications and cybersecurity assessment, distributed-node placement, and advisory-only digital twin operation. High-consequence actions such as feeder reconfiguration, black-start sequencing, reserve-

floor override, and major load-shed events should remain human-approved unless specifically validated and authorized.

17.7 Roadmap End State

The end state is not unrestricted autonomy. The end state is governed, evidence-backed supervisory control.

A mature Mission-Energy Homeostasis / Allostasis deployment should know the physical system, understand mission consequence, preserve priority loads, protect reserve, reduce unnecessary generator runtime, shrink authority under uncertainty, explain itself to operators, fall back safely, recover in sequence, and maintain evidence for continued trust.

The deployment roadmap therefore converts the doctrine into a disciplined RDT&E pathway:

- **Research establishes the doctrine.**
- **Development builds the architecture.**
- **Test and Evaluation produce the evidence.**
- **Sustainment preserves trust.**

Thought for a couple of seconds

Replace Section 18 with this:

18. Conclusion

18.1 The Return to Mission-Energy Balance

This paper began with a simple argument: tactical and installation-scale energy systems do not become resilient merely because more equipment is added. More generation, more storage, more fuel, more telemetry, more dashboards, and more automation may increase capability, but capability without governed behavior is not resilience.

Resilience is proven when the system preserves mission-valid equilibrium under stress.

That is the central meaning of Mission-Energy Homeostasis / Allostasis. Homeostasis restores balance after drift. Allostasis prepares the system before forecasted stress becomes failure. Together, they define a doctrine for energy systems that must do more than produce power. They must sense, understand, act, inform, recover, and return.

A generator can exist and still waste fuel.

A battery can exist and still consume reserve at the wrong time.

A solar array can exist and still fail to protect the right load.

A microgrid can island and still lack recovery discipline.

A digital twin can visualize the system and still fail to support command judgment.

AI can forecast demand and still become unsafe without boundaries, confidence gating, cybersecurity discipline, deterministic fallback, and human authority.

The future of mission energy will not belong to systems that merely have more power. It will belong to systems that govern power with discipline.

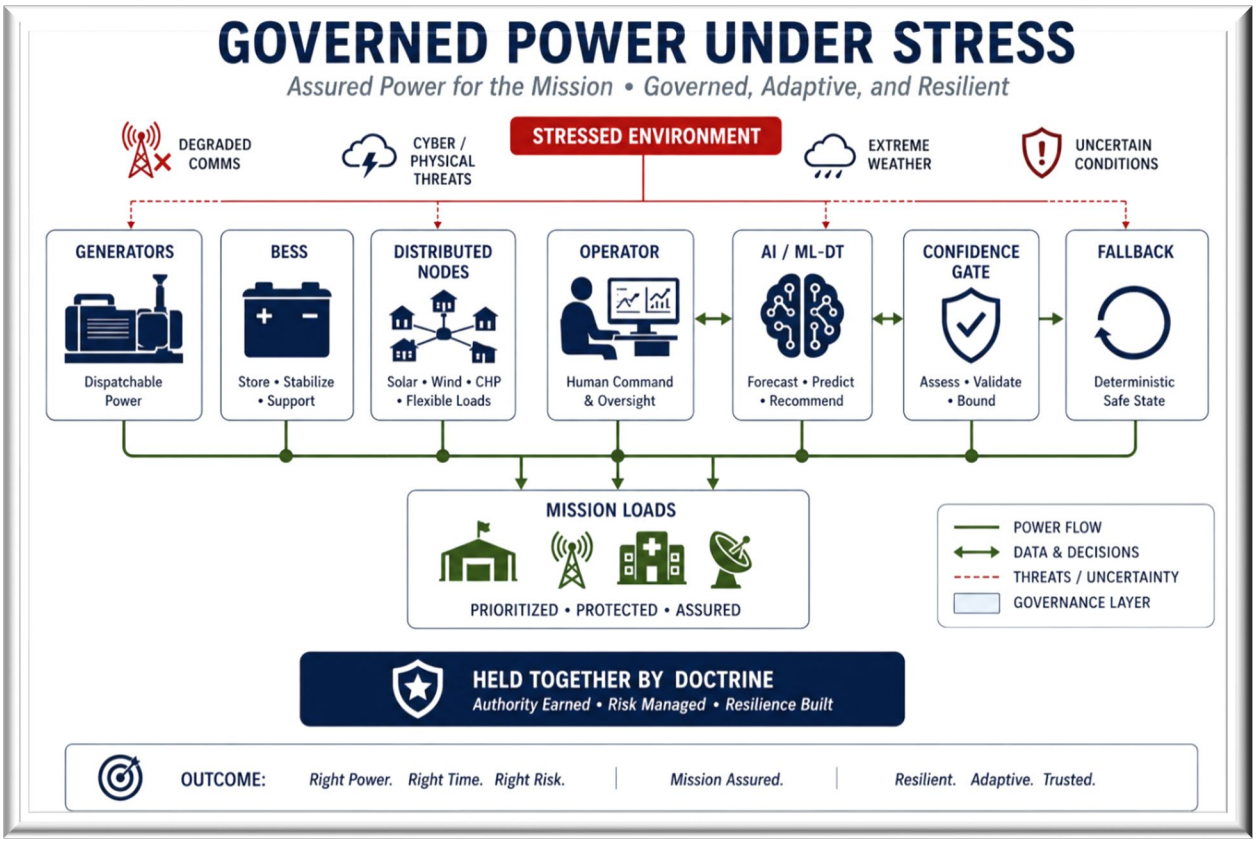

18.2 From Energy Availability to Command Freedom

The deeper issue is not energy availability alone. The deeper issue is command freedom.

Commanders and operators do not need abstract kilowatts. They need time, options, reserve, confidence, discretion, recoverability, and trust. They need to know which loads can survive, how long they can survive, what must be protected, what can be deferred, when fuel should be burned, when silence should be preserved, when automation should act, when it should ask, and when it should stop trusting itself.

That is the difference between energy that exists and energy that can be trusted.

Mission-Energy Homeostasis / Allostasis reframes energy as a governed survivability function. It moves the question beyond “how much power do we have?” and toward the harder operational question introduced at the beginning of this paper:

Can the energy ecosystem preserve mission-critical function under changing conditions, with enough reserve, enough confidence, enough control authority, and a clear recovery path?

That is the question this doctrine is built to answer.

18.3 The Epistemological Boundary of Doctrine

Doctrine has an honest boundary. It can define the framework, establish the vocabulary, describe the control logic, bound the authority model, and motivate the program. It can explain what the system should do and why the behavior matters.

What doctrine cannot do is prove trust.

Trust cannot be asserted by architecture. It cannot be assumed because the model is elegant, because the interface is clear, because the hardware is capable, or because the algorithm performs well in nominal conditions. Trust is not a narrative condition. It is an evidence condition.

The real confidence thresholds, the real authority partition boundaries, the real validated operating envelope, the real fallback behavior, and the real equilibrium return time are not written into existence. They are discovered through development, test, evaluation, degraded operation, fault injection, operator use, cybersecurity assessment, and the uncomfortable edge cases that only hardware, telemetry, and field conditions reveal.

That is why the crawl phase matters. It is not administrative caution. It is where the doctrine meets physical reality. It is where assumptions become measurements. It is where ambition is forced to answer to evidence.

18.4 The R in RDT&E

This paper represents the **R** — the research foundation — of a deliberate RDT&E progression.

The research establishes the doctrine.

Development must build the architecture.

Test and Evaluation must produce the evidence.

Sustainment must preserve trust.

That progression matters because Mission-Energy Homeostasis / Allostasis is not a claim of universal autonomy. It is not an argument for replacing protection engineering, operator judgment, certified equipment behavior, or command authority. It is an argument for disciplined supervisory intelligence: physics-informed, confidence-gated, cyber-aware, bounded, explainable, recoverable, and validated.

The doctrine’s strength is not that it promises the system will always know what to do. Its strength is that it demands the system know when it does not know enough.

In high-consequence energy systems, that distinction is everything.

18.5 Governed Power as the Future of Resilience

The next failure will not always announce itself as a blackout. It may begin as a small drift in load, a weak generator response, a stale telemetry stream, a battery reserve consumed too early, an inverter operating near its limit, a communications path losing integrity, a model no longer matching the plant, or an operator receiving data without decision-grade meaning.

Failure begins when imbalance is not recognized, consequence is not understood, authority is not bounded, recovery is not sequenced, and trust is assumed instead of earned.

Mission-Energy Homeostasis / Allostasis is a doctrine for preventing that failure path. It gives tactical and installation-scale energy systems a governing logic: detect drift, understand consequence, act within validated authority, inform the human operator, fall back when trust weakens, and return to mission-valid equilibrium.

That is the standard future systems must meet.

Not more power alone.

Not automation for its own sake.

Not resilience by brochure.

Governed power.

Measured trust.

Command freedom under stress.

References

Berger, Michael S. *Fossil-Smarter: Why the Future of Energy Is Not Less Fuel, But Less Waste.* Beech Creek Power & Energy, 2026.

Berger, Michael S. *Solving for Energy Advantage.* Beech Creek Power & Energy, 2026.

National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0).* NIST AI 100-1, January 2023.

Stouffer, Keith, et al. *Guide to Operational Technology (OT) Security.* NIST SP 800-82, Revision 3, September 2023.

Joint Task Force. *Security and Privacy Controls for Information Systems and Organizations.* NIST SP 800-53, Revision 5, September 2020.

Department of Defense. *DoD Instruction 8510.01: Risk Management Framework for DoD Systems.* July 2022.

Department of Defense. *DoD Instruction 4170.11: Installation Energy Management.* December 2009.

Department of Defense. *DoD Instruction 5000.89: Test and Evaluation.*

IEEE. *IEEE Standard 1547-2018: Interconnection and Interoperability of Distributed Energy Resources.* 2018.

IEEE. *IEEE Standard 1547.1-2020: Conformance Test Procedures for Equipment Interconnecting DERs.* 2020.

IEEE. *IEEE Standard 2030.7-2017: Specification of Microgrid Controllers.* 2017.

IEEE. *IEEE Standard 2030.8-2018: Testing of Microgrid Controllers.* 2018.

IEEE. *IEEE Standard 519-2022: Harmonic Control in Electric Power Systems.* 2022.

UL Standards & Engagement. *UL 1741: Inverters, Converters, Controllers and Interconnection System Equipment for Use With Distributed Energy Resources.*

UL Standards & Engagement. *UL 9540: Energy Storage Systems and Equipment.*

UL Standards & Engagement. *UL 9540A: Test Method for Evaluating Thermal Runaway Fire Propagation in Battery Energy Storage Systems.*

NFPA. *NFPA 70: National Electrical Code.*

NFPA. *NFPA 855: Standard for the Installation of Stationary Energy Storage Systems.*

IEC. *IEC 61850: Communication Networks and Systems for Power Utility Automation.*

IEC/IEEE. *IEC/IEEE 61850-9-3: Precision Time Protocol Profile for Power Utility Automation.*

IEC. *IEC 62443 Series: Industrial Communication Networks — Network and System Security.*

He, Xing, et al. "Preliminary Exploration on Digital Twin for Power Systems." arXiv:1909.06977, 2019.

Valibeygi, Amir, et al. "Microgrid Control Using Remote Controller Hardware-in-the-Loop Over the Internet." arXiv:1804.05998, 2018.

Morstyn, Thomas, et al. "Model Predictive Control for Distributed Microgrid Battery Energy Storage Systems." *IEEE Transactions on Control Systems Technology,* 2017.

Xu, Hanchen, et al. "Data-Driven Coordination of Distributed Energy Resources for Active Power Provision." *IEEE Transactions on Power Systems,* vol. 34, no. 4, 2019.

Watson, Jeremy, et al. "A Scalable Control Design for Grid-Forming Inverters in Microgrids." arXiv:2012.11556, 2021.

Li, Yitong, et al. "Revisiting Grid-Forming and Grid-Following Inverters: A Duality Theory." arXiv:2105.13094, 2021.

Ansari, O. A., et al. "Reliability Assessment of Microgrid with Renewable Generation and Prioritized Loads." arXiv:1709.07970, 2017.

Bidram, Ali, and Ali Davoudi. "Hierarchical Structure of Microgrids Control System." *IEEE Transactions on Smart Grid,* vol. 3, no. 4, 2012.

Guerrero, Josep M., et al. "Hierarchical Control of Droop-Controlled AC and DC Microgrids." *IEEE Transactions on Industrial Electronics,* vol. 58, no. 1, 2011.

Olivares, Daniel E., et al. "Trends in Microgrid Control." *IEEE Transactions on Smart Grid,* vol. 5, no. 4, 2014.

Lasseter, Robert H. "MicroGrids." *IEEE Power Engineering Society Winter Meeting,* 2002.

Li, Xiaojun, et al. "Data-Driven Thermal Anomaly Detection for Batteries Using Unsupervised Shape Clustering." arXiv:2103.08796, 2021.

Cai, Ting, et al. "Li-Ion Battery Fault Detection in Large Packs Using Force and Gas Sensors." arXiv:2010.13519, 2020.

Firoozi, Roya, et al. "Cylindrical Battery Fault Detection Under Extreme Fast Charging: A Physics-Based Learning Approach." arXiv:2105.02169, 2021.

National Renewable Energy Laboratory. *Microgrids.* NREL technical resources on microgrid research, controls, resilience, and distributed energy integration.

U.S. Department of Energy, Office of Electricity. *Microgrid Program Strategy.*

U.S. Department of Energy. *Energy Storage Safety Strategic Plan.*

Sandia National Laboratories. *Energy Storage Safety and Reliability Research.*