

ODA3 INSTITUTE

EXECUTIVE BRIEF

GAISSF™ v1.0 Energy Sector Executive Brief

Leadership decisions for secure, resilient and accountable energy-sector AI

GAISSF-SEC-044 | Informative sector guidance | Draft for Publication

Version 1.0 | Public | External Review Open | Updated 6 October 2026

© 2026 ODA3 Pvt Ltd

Methodology note

Evidence sourcing follows ODA3’s controlling T1-T4 evidence-sourcing model; this brief does not define a separate T1-T4 scale. Implementation evidence should be proportionate to system authority, consequence and assurance objective and does not alter GAISSF Foundational / Operational / Optimized conformance tiers. The Applicable-Control Baseline is all 52 D1-D8 controls plus D9 where Physical AI or cyber-physical actuation is in scope.

Where D9 is in scope, use GAISSF D9 alongside PAI-SF™ v1.0 as applicable; CSX-PAISF-004 is mapping only and does not establish equivalence, evidence sufficiency or automatic conformance.

Executive assessment

Energy organisations should govern AI according to the operational authority it can exercise and the consequence of error, manipulation or dependency failure. The highest-priority boundary is between an AI output and a physical, market or customer-impacting action.

Five decisions leaders must make

- Which AI systems can influence operational, safety, market or essential-customer outcomes?
- Which actions require positive human approval, dual authorisation or independent machine verification?
- What deterministic fallback and safe state remain available when the model, data, cloud or communications fail?
- What evidence is required before release and during operation to demonstrate control effectiveness?
- Which residual risks and open specialist reviews prevent deployment or require executive acceptance?

Highest-priority actions

- Complete the AI inventory and classify authority, criticality, safety relevance and reversibility.
- Block unauthorised actuation and separate recommendations from approved commands.
- Preserve independent protection, manual override, fallback and trusted rollback.
- Validate representative and degraded scenarios; monitor drift, rejected actions, overrides and near misses.
- Require supplier provenance, secure updates, incident notification, continuity and exit evidence.

Illustrative 90-day sequence (non-prescriptive)

The timing below is illustrative only. Adapt sequence and duration to system authority, consequence, programme phase, operating constraints, applicable safety or reliability processes, and jurisdiction-specific obligations.

Window	Leadership outcome
Days 1-30	Scope, inventory, accountable owners, authority classification, N/A rules and critical containment.
Days 31-60	Priority controls, supplier review, release gates, evidence repository and fallback tests.
Days 61-90	Exercises, degraded-mode tests, monitoring thresholds, residual-risk decisions and independent challenge.

Notably Absent

- No universal evidence that AI improves reliability or safety.
- No threshold, maturity level, profile or sequencing label establishes a universal energy-sector acceptance criterion.
- No claim of legal compliance, regulatory approval, nuclear approval or certification.
- No assumption that human presence alone constitutes effective oversight.
- Use-case mappings and scenario treatments remain subject to specialist validation.