AAIGF-E

Adaptive AI Governance Framework for the Electric Sector

Version 1.0 | Executive Brief

> **The only control-mapped, sector-specific AI governance framework built for the electric sector, designed to extend your existing compliance program into AI-specific risk territory.**

Operationalizing governance through integrity, assurance, and audit-ready controls.

The Problem

AI is operating inside the Bulk Electric System: load forecasting, grid optimization, contingency assessment, DER coordination, and OT anomaly detection. **The governance frameworks that protect critical infrastructure were not built for this.**

The gap is structural. No existing mandatory or sector-specific standard governs how AI systems are assessed, controlled, monitored, and responded to within electric sector operations.

Standard	The blind spot for AI in electric sector operations
NERC CIP	Mandatory but predates AI. No requirements for model governance, adversarial threats, drift detection, or AI output influence on operators.
NIST AI RMF	Voluntary and sector-agnostic. No mandatory standing in electric sector compliance. No OT operational specificity.
ISA/IEC 62443	Governs OT cybersecurity. Does not govern AI model behavior, lifecycle management, or AI-specific failure modes.
ISO/IEC 42001	Generic AI management system. No electric sector grounding. No enforcement pathway in the US regulatory structure.

The Solution

AAIGF-E provides 111 governance controls across 11 risk domains, structured around seven concurrent lifecycle anchors, and explicitly mapped to NERC CIP, NIST AI RMF, ISA/IEC 62443, and ISO/IEC 42001.

It is an overlay, not a replacement. For organizations subject to NERC CIP, CIP-aligned processes remain the system of record. AAIGF-E extends that posture into AI-specific risk territory that CIP does not address, without creating parallel obligations or rebuilding what already works.

What Makes AAIGF-E Different

Differentiator	Why it matters
Electric sector only	Every control is grounded in BES reliability, OT constraints, and operator decision influence. Not generic AI ethics.

Differentiator	Why it matters
Adversarial AI coverage	MITRE ATLAS adversarial threat taxonomy is integrated. Model poisoning, input crafting, and inference manipulation are addressed explicitly, not by analogy.
Concurrent lifecycle model	Seven anchors run simultaneously. Governance and monitoring never pause. Response and recovery are triggered by events and run alongside the others, not instead of them.
Audit-ready by design	Every control has a testable objective, a defined Evidence Required field, and an Audit Verification statement. Written for assessors, not architects.
CIP overlay positioning	NERC CIP stays as the compliance backbone. AAIGF-E closes the AI-specific gap above it. No duplication. No conflict.
Organization-defined scope	Controls use organization-defined thresholds and intervals, consistent with NIST SP 800-53 convention, making them enforceable and audit-defensible.

The Lifecycle Anchor Model

This is not a phase model. Every anchor runs continuously. Governance does not pause during incidents. Monitoring does not stop during changes. Respond and Recover activate on events and operate alongside the other anchors, because that is how grid operations actually work: concurrent, continuous, and event-driven.

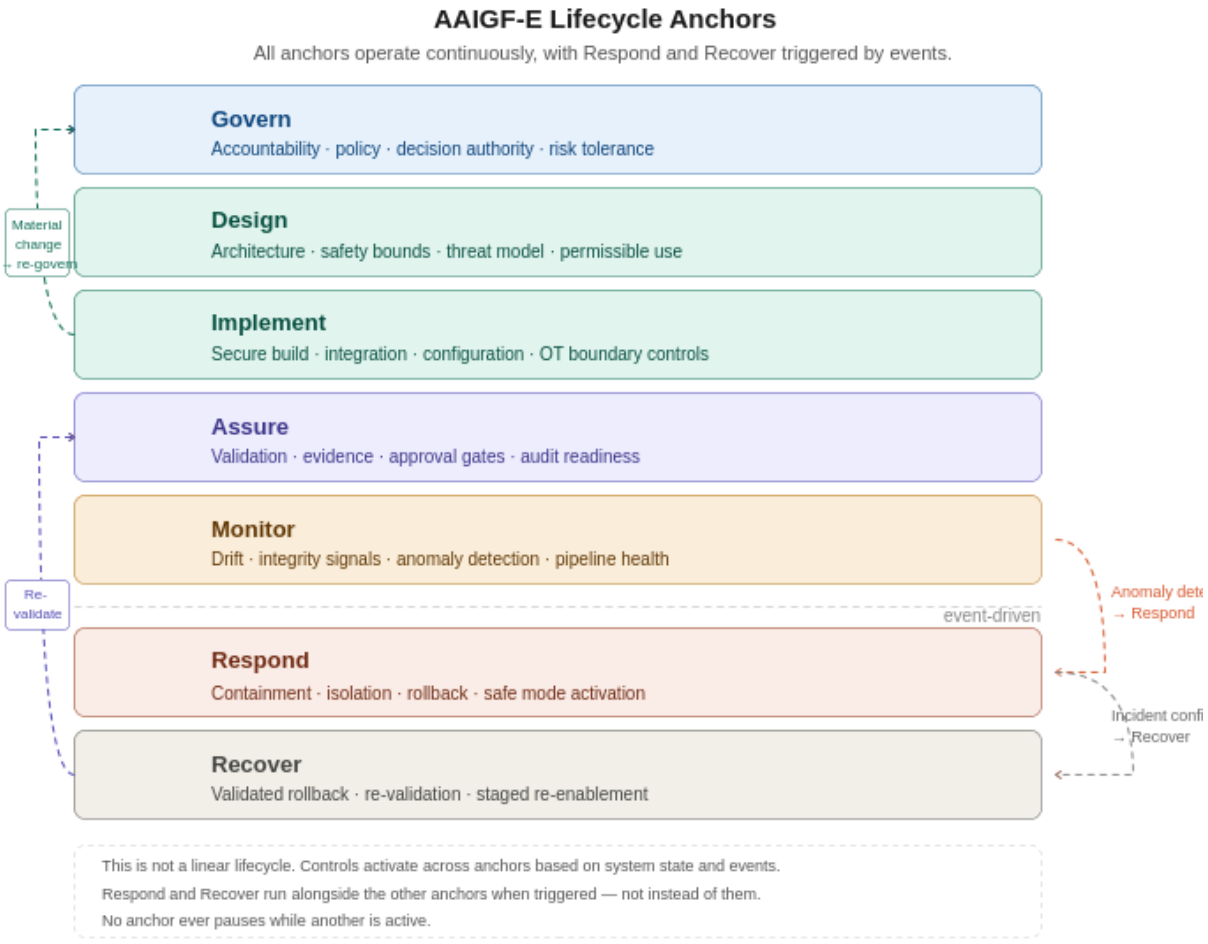

Figure 1: AAIGF-E governance reflects how electric sector operations actually work: continuous, concurrent, and event-responsive.

Adoption Model

Four-tier criticality classification

Tier	Level	AI system scope
3	**High**	AI outputs directly influence Bulk Electric System (BES) reliability, operator actions, or automated workflows.
2	**Medium**	AI shapes operational plans, schedules, or maintenance. No direct real-time control.
1	**Low**	AI supports analytics or reporting with limited operational consequence.
0	**Sandbox**	Experimental or prototype systems not used for production decisions.

Minimum viable adoption: three required artifacts

1. AI system risk register with tier assignment and acceptance criteria.
2. Vendor security package for all third-party AI components.
3. AI incident response playbooks with evidence of tabletop exercises.

Governance requirement

AAIGF-E requires authorization at CRO, GRC leadership, or AI governance function level before technical implementation. It cannot be adopted bottom-up.

National Security Dimension

AI governance failure in the electric sector is not a compliance problem. It is a national security risk. A model that drifts, produces manipulated outputs, or is compromised through its training pipeline can influence operator decisions before detection is possible, with physical consequences on infrastructure that underpins national defense, healthcare, water, and financial systems.

AAIGF-E integrates MITRE ATLAS adversarial AI threat taxonomy to address model poisoning, adversarial input crafting, and inference manipulation: attack surfaces that existing cybersecurity frameworks were not built to govern.

Who Should Act on This Now

Role	Why AAIGF-E is relevant to your function
CRO / Enterprise Risk Leadership	You own the exposure. No sector-specific AI governance standard exists. AAIGF-E closes that gap before a regulator or incident does.
Head of GRC / Compliance Director	You need something audit-defensible. AAIGF-E is built to evidence requirements, not just state them.
CISO / OT Security Leader	Your CIP program does not govern model integrity, drift, or adversarial AI inputs. AAIGF-E adds that layer without disrupting what you already have.
Grid Operator / ISO / RTO	AI is influencing decisions on infrastructure you are responsible for. AAIGF-E gives you the governance structure to control that influence.
Regulator / Standards Body	No mandatory AI governance standard exists for this sector. AAIGF-E provides the reference architecture for what one should look like.

> **AI is already operating in your grid. The governance question is not whether to address it. It is whether you address it before an incident, a regulatory finding, or a compromise forces you to.**

AAIGF-E Version 1.0 is available in full. For assessment support, pilot program inquiries, or framework review, contact the authors.

Authors: Suhail Ahmad Rana, Clint Bodungen