

NERC CIP-to-AI Governance Mapping Framework

Mapping NERC Critical Infrastructure Protection standards
to AI agent governance in energy infrastructure.

AUTHOR

Gabriel Ohaike

*Founder & Chief Architect, Agentic AI Runtime Governance
PhronEdge*

PUBLISHED

June 2026

VERSION

1.0

DISTRIBUTION

Public

Executive Summary

NERC Critical Infrastructure Protection standards define the cybersecurity requirements for the North American bulk electric system. These standards were written for a world where humans operate the grid, traditional software manages assets, and vendors supply physical and digital components.

That world is changing. AI agents are entering energy infrastructure: optimizing dispatch, forecasting generation, trading in wholesale markets, monitoring equipment, and analyzing compliance obligations. These agents access the same BES Cyber Systems that CIP standards were designed to protect.

The standards do not mention AI. This creates a governance gap that every registered entity deploying AI systems must address.

This framework provides the first comprehensive mapping of every applicable CIP standard to its AI agent equivalent. For each standard, it defines the original requirement, the equivalent obligation when an AI agent is involved, a concrete operational scenario, and the technical controls required for enforcement.

Why This Matters Now

REGULATORY CONVERGENCE

Three forces are converging. **FERC** is examining algorithmic participation in wholesale markets, with AI-driven bidding strategies increasingly common across ISOs. **NERC** has not issued AI-specific guidance, but CIP standards are technology-neutral by design. An AI agent accessing BES Cyber Systems is a cyber asset that must be identified, classified, and protected under existing rules. **The EU AI Act** classifies energy grid management as high-risk AI, creating compliance obligations for US companies with European operations.

DEPLOYMENT REALITY

AI adoption in energy is not theoretical. ERCOT added 11 GW of new generation and 6 GW of battery storage in 2025. PJM projects 70% load growth over two decades driven by data centers. Battery dispatch optimization, renewable forecasting, demand response management, and compliance analysis are already running on AI systems at utilities, ISOs, and independent power producers across North America.

The question is not whether these systems need governance. The question is whether governance frameworks will exist before the first enforcement action forces them into existence.

The Mapping

Nine CIP standards have direct applicability to AI agent governance. Each entry includes the original intent, the AI agent equivalent, a concrete operational scenario, and the technical enforcement mechanism.

CIP-002 BES Cyber System Categorization

ORIGINAL REQUIREMENT

Identify and categorize all BES Cyber Systems by impact level.

AI AGENT EQUIVALENT

Register every AI agent that accesses BES infrastructure in the asset inventory. Classify by impact level using the same criteria applied to traditional cyber assets. An agent dispatching generation is High Impact. An agent generating internal reports is Low Impact. The classification determines which governance controls apply.

SCENARIO

A utility deploys an AI forecasting agent for wind generation scheduling. The agent reads real-time SCADA data from High Impact BES Cyber Systems. Under this mapping, the agent is classified as High Impact and inherits all corresponding CIP controls.

ENFORCEMENT

Agent registry with impact classification. Every agent declared with its BES access scope before deployment. Unregistered agents blocked from accessing BES data.

CIP-003 Security Management Controls

ORIGINAL REQUIREMENT

Document and implement cybersecurity policies for all BES Cyber Systems.

AI AGENT EQUIVALENT

Define behavioral policies for each AI agent as machine-readable, version-controlled rules enforced at runtime. Policies specify what the agent can access, what actions it can take, what outputs it can produce, and what constraints it must respect. Documentation alone is insufficient. Policies must be executable.

SCENARIO

An AI trading agent operating in the ERCOT day-ahead market is governed by a policy that caps energy offers at \$2,000/MWh per FERC Order 831, restricts API calls to authorized ERCOT endpoints, and prohibits including market position data in any response.

ENFORCEMENT

Constitutional policy engine. Policies defined as code, version-signed, enforced at the point of every AI decision. Policy violations blocked before the request reaches the model.

CIP-004 Personnel and Training

ORIGINAL REQUIREMENT

Background check, train, and authorize every person with BES Cyber System access.

AI AGENT EQUIVALENT

Every AI agent receives a verified identity, scoped permissions, and periodic access review. The agent's prompt, model version, and behavioral constraints serve as its training record. Permissions reviewed on the same cycle as human personnel.

SCENARIO

A predictive maintenance agent has read access to transformer sensor data across 14 substations but cannot issue control commands. When the model is retrained, the updated version goes through the same approval process as onboarding a new employee with BES access.

ENFORCEMENT

Agent identity management. Cryptographic identity per agent. Permission scoping enforced at the gateway. Access audit logs maintained for regulatory review.

CIP-005 Electronic Security Perimeters

ORIGINAL REQUIREMENT

Control all traffic crossing the Electronic Security Perimeter around BES Cyber Systems.

AI AGENT EQUIVALENT

Every API call an AI agent makes crosses a boundary. Outbound calls to external data providers and inbound responses from model inference endpoints must be logged, validated against authorized endpoints, and blocked if unauthorized.

SCENARIO

A renewable forecasting agent pulls weather data from a third-party API. Every call is logged, the endpoint is validated against an allowlist, and any attempt to reach an unauthorized service is blocked. If the provider changes its domain, the agent is blocked until the new endpoint is approved.

ENFORCEMENT

API-level enforcement at the network perimeter. Request and response logging. Endpoint allowlisting. Automatic blocking of unauthorized destinations with alert to the security team.

CIP-007 System Security Management

ORIGINAL REQUIREMENT

Patch management, malware prevention, and security event monitoring for all cyber assets.

AI AGENT EQUIVALENT

Model version tracking replaces traditional patch management. Every model update is treated as a configuration change with a cryptographic hash. Integrity verification before deployment replaces malware scanning. Output anomaly detection replaces security event monitoring.

SCENARIO

A demand forecasting model is retrained quarterly. Before deployment, the new model's hash is compared against the approved baseline. If it does not match, deployment is blocked. Post-deployment, the system monitors output distributions and alerts if predictions deviate beyond established thresholds.

ENFORCEMENT

Model integrity verification via cryptographic signing. Hash validation before deployment. Output anomaly detection at runtime. Automated alerting for behavioral drift.

CIP-008 Incident Reporting and Response

ORIGINAL REQUIREMENT

Detect, respond to, and report cybersecurity incidents.

AI AGENT EQUIVALENT

When an AI agent makes a consequential decision, the full decision chain must be auditable: what data was received, what model version processed it, what policy constraints were evaluated, and

what output was generated. The response plan must include rollback capability.

SCENARIO

A substation monitoring agent recommends emergency shutdown of a transformer based on dissolved gas readings. The audit trail records: sensor data from unit X at timestamp Y, model version Z, policy constraints checked, recommendation generated. A NERC auditor can replay the entire decision chain. Without this, the utility cannot explain why it took the unit offline.

ENFORCEMENT

Immutable audit log with cryptographic signing. Full decision chain captured: input, model version, policy evaluation, output. Tamper-evident records. Rollback capability via policy revocation.

CIP-010 Configuration Change Management

ORIGINAL REQUIREMENT

Document baseline configurations and approve all changes before implementation.

AI AGENT EQUIVALENT

Model retraining is a configuration change. Prompt modifications are configuration changes. New data source integrations are configuration changes. Each requires documented rationale, testing in an isolated environment, approval by the designated owner, and a signed baseline before production deployment.

SCENARIO

A utility fine-tunes its outage prediction model with two years of new storm data. The change follows CIP-010 change management: documented rationale, test results, approval by the model owner, signed baseline of the new model hash, and production deployment only after sign-off. Unauthorized model versions are rejected at deployment.

ENFORCEMENT

Configuration baseline management. Model hash signing on every approved version. Change approval workflow integration. Deployment gating: unsigned or unapproved versions rejected automatically.

CIP-011 Information Protection

ORIGINAL REQUIREMENT

Protect BES Cyber System Information and Critical Energy Infrastructure Information from unauthorized disclosure.

AI AGENT EQUIVALENT

AI agents trained on or accessing CEII and BCSI data must not disclose that data in outputs. The agent cannot include raw generation figures, grid topology details, or security vulnerability information in responses to unauthorized users. Data classification must be enforced at the output layer.

SCENARIO

A generation forecasting model was trained on historical output data for 3 gas turbine plants totaling 2,400 MW. That training data is CEII. The system enforces: the model cannot reproduce raw training data in responses, cannot be queried for specific unit output figures, and cannot be accessed by users without CEII clearance.

ENFORCEMENT

Output filtering for CEII and BCSI patterns. User authorization validation on every request. Data classification enforcement at the response layer. Automated redaction of sensitive content before delivery.

CIP-013 Supply Chain Risk Management

ORIGINAL REQUIREMENT

Assess and mitigate cybersecurity risks from vendors with access to BES Cyber Systems.

AI AGENT EQUIVALENT

AI model providers are supply chain vendors. Any organization whose model processes BES data is a vendor requiring risk assessment proportional to the data sensitivity and decision authority of the AI system. Contract clauses must address model update notification, data handling, breach notification, right to audit, and geographic restrictions.

SCENARIO

A utility uses a third-party language model for compliance analysis of CIP audit data. The model provider is classified as a Tier-1 vendor. The assessment covers: where inference runs, whether query data is retained, model update notification timelines, breach notification SLAs, and right to audit. The same vendor questionnaire is extended with AI-specific questions on model provenance and training data governance.

ENFORCEMENT

Vendor dependency tracking for all model providers in the AI stack. Automated monitoring for model version changes. Contract compliance verification. Supply chain risk scoring for AI-specific dependencies.

Reference Table

Standard	Original Intent	AI Equivalent
CIP-002	BES Cyber System Categorization	Register every AI agent that accesses BES infrastructure in the asset inventory. Classify by impact level using the same criteria applied to traditional cyber assets.
CIP-003	Security Management Controls	Define behavioral policies for each AI agent as machine-readable, version-controlled rules enforced at runtime. Policies specify what the agent can access, what actions it can take, what outputs it can produce, and what constraints it must respect.
CIP-004	Personnel and Training	Every AI agent receives a verified identity, scoped permissions, and periodic access review. The agent's prompt, model version, and behavioral constraints serve as its training record.
CIP-005	Electronic Security Perimeters	Every API call an AI agent makes crosses a boundary. Outbound calls to external data providers and inbound responses from model inference endpoints must be logged, validated against authorized endpoints, and blocked if unauthorized..
CIP-007	System Security Management	Model version tracking replaces traditional patch management. Every model update is treated as a configuration change with a cryptographic hash.
CIP-008	Incident Reporting and Response	When an AI agent makes a consequential decision, the full decision chain must be auditable: what data was received, what model version processed it, what policy constraints were evaluated, and what output was generated. The response plan must include rollback capability..
CIP-010	Configuration Change Management	Model retraining is a configuration change. Prompt modifications are configuration changes.
CIP-011	Information Protection	AI agents trained on or accessing CEII and BCSi data must not disclose that data in outputs. The agent cannot include raw generation figures, grid topology details, or security vulnerability information in responses to unauthorized users.
CIP-013	Supply Chain Risk Management	AI model providers are supply chain vendors. Any organization whose model processes BES data is a vendor requiring risk assessment proportional to the data sensitivity and decision authority of the AI system.

Implementation Path

Phase 1: Inventory and Classification

Identify every AI system that accesses, processes, or influences BES Cyber Systems. Classify each by impact level using CIP-002 criteria. Document the agent's purpose, data access scope, decision authority, and model provider. No governance program can begin without a complete inventory.

Phase 2: Policy Definition

For each classified agent, define behavioral policies aligned with the CIP mapping. Policies must be specific and enforceable: what endpoints the agent can access, what data it can include in outputs, what actions it can take autonomously versus those requiring human approval. Policies that exist only in documentation provide no protection.

Phase 3: Runtime Enforcement

Deploy technical controls that enforce policies at the point of every AI decision. This includes API-level request validation, output filtering for CEII and BCSI content, model integrity verification before deployment, and immutable audit logging for every agent action. Enforcement must be continuous, automated, and tamper-evident.

Phase 4: Audit Readiness

Maintain continuous compliance evidence for every AI system in the inventory. Every agent decision should produce an auditable record: input data, model version, policy constraints evaluated, output generated. When regulators ask how your AI systems comply with CIP standards, the answer is a signed, immutable decision chain.

About the Author

Gabriel Ohaike has spent his career at the intersection of enterprise platform architecture and regulatory compliance. He has led the design and implementation of AI Centers of Excellence across energy, financial services, and critical infrastructure, helping organizations define the governance frameworks, operating models, and technical architectures required to deploy autonomous systems in regulated environments.

He founded PhronEdge after seeing the same gap across every enterprise engagement: organizations adopt AI agents faster than they can govern them. The distance between what regulators require and what teams can enforce at runtime continues to widen. PhronEdge was built to close that gap with constitutional controls that operate at the point of every AI decision.

About PhronEdge

PhronEdge provides regulatory intelligence and runtime governance for AI systems operating in highly regulated industries. The platform monitors regulatory changes globally and enforces compliance at the point of every AI decision through constitutional signing, policy enforcement, immutable audit logging, and decision chain auditability.

<https://www.phronedge.com/energy>

This framework is published for public distribution. PhronEdge encourages regulatory bodies, registered entities, and industry working groups to reference and build upon this mapping in the development of AI governance standards for critical infrastructure.