ENERGY COMPLIANCE, INC. | *Field Reference for Compliance Programs*

SPECIAL TOPICS SERIES · EC-WP-602

AI and Automation in Reliability Compliance

Where AI and automation help compliance, where they do not, and how to deploy responsibly.

A practitioner's guide to AI and automation in a NERC compliance program: what tools do well, which judgment calls must stay with an accountable person, and how to deploy and revalidate automation so its output holds up at audit.

Energy Compliance, Inc. Professional Reference

energycomplianceinc.com **OCTOBER 2026**

Disclaimer & Legal Notice

This publication is the work product of Energy Compliance, Inc. and is provided for general information and educational purposes only. It does not constitute legal advice, regulatory guidance, or a formal interpretation of any NERC Reliability Standard, FERC order, or Regional Entity directive.

This publication may not be used or cited as authority in any audit, enforcement, regulatory, or legal proceeding. Compliance with NERC Reliability Standards depends on facts and circumstances unique to each registered entity. Standard versions, requirement text and effective dates cited here are as of October 25, 2026, and they change. Confirm the current text before relying on it, and consult qualified counsel and subject matter experts when making compliance decisions.

Nothing in this publication endorses or evaluates any third-party software product. References to AI and automation tools describe categories of capability, not specific vendors.

NERC, FERC, and the names of Regional Entities are referenced here for descriptive purposes only. Energy Compliance, Inc. is not affiliated with, sponsored by, or endorsed by the North American Electric Reliability Corporation, the Federal Energy Regulatory Commission, or any NERC Regional Entity.

Energy Compliance, Inc. assumes no responsibility or liability for errors, omissions, or outcomes resulting from the use of this publication.

Foreword

This professional reference is one of a series Energy Compliance, Inc. publishes for registered entities and the people who run their compliance programs.

I've spent more than thirty years on every side of the bulk electric system. I've operated control centers as a Reliability Coordinator, Transmission Operator, and Power System Operator. I've audited grid facilities and signed off on findings as a senior compliance auditor. I've worked enforcement matters from inside the regulator's process. For the last several years I've advised registered entities directly through the firm I founded.

The entities that do reliability well share a common habit. They take the standards seriously without confusing them with reliability itself. They know that a NERC Reliability Standard is a floor, not a ceiling. They know that compliance is something an auditor evaluates, but reliability is something a system either delivers or doesn't. They prepare for audits by building programs that survive real questions, not binders that look thick.

That's the perspective these references try to share. Each one focuses on a single topic: a standard family, an operational function, a regulatory framework, or an emerging industry challenge. Each one walks through how the topic actually works.

This one takes on AI and automation. Every compliance team is being pitched a tool right now, and some of those tools are genuinely useful. The question I care about is narrower than the pitch: when an auditor asks who made a decision and why, does the program still have an answer?

These references are written for the compliance manager who wants to understand the system, not just memorize requirements. For the legal counsel who has to brief a board honestly. For the senior operator who's been told that compliance and reliability are the same thing and suspects they aren't. And for the new compliance hire who got handed a binder and told good luck.

These references aren't marketing material disguised as content. They're the result of three decades of doing this work and watching it succeed and fail. I've written them in the same voice I use in a control room or in front of a Regional Entity audit team: direct, evidence-grounded, honest about what the standards do and do not require.

Energy Compliance exists because, in my view, too much of the consulting offered to registered entities today is structured for billable hours rather than for outcomes. Every engagement is overseen by a senior practitioner. We don't bring five people to a meeting that needs one. We automate the work that should be automated. We apply senior judgment to the work that requires it. If that approach matches what you're looking for in a compliance partner, the back of this reference has our contact information.

Rob Smith

Founder and Principal, Energy Compliance, Inc.

Executive Summary

There is a wide gap between what AI vendors say their tools can do and what those tools deliver in a production compliance program. Closing it takes a clear view of three things: what current tools are actually good at, where automation gives a compliance team real gains, and where it creates risk that wasn't there before. The wrong deployment doesn't simply fail. It makes the program look more efficient while the quality of the underlying compliance work erodes.

- **Explainability is the test.** If a tool can't tell you why it produced an output, the output isn't audit-defensible. The auditor will ask. The tool can't answer. The entity has to.
- **Automate the clocks and the cross-checks.** High-volume, rule-based, deadline-driven work is where automation pays. The CIP standards are full of it: a 35-day patch evaluation cycle and a separate 35-day action clock in CIP-007-6, quarterly and 15-month access checks in CIP-004-7, a 30-day baseline update and a 15-month vulnerability assessment in CIP-010-4. Those three standards topped NERC's list of the most frequently reported CIP noncompliance in 2025 (Chapter 2).
- **Some work must stay with a person.** Classification decisions, scope determinations, self-report strategy and audit-defense narrative need accountability a tool can't provide. CIP-002-5.1a already requires the CIP Senior Manager or a delegate to approve the BES Cyber System identifications (Chapter 2).
- **Unreviewed output is not evidence.** Every automated output needs a documented human review: who, when, what was decided, and on what basis. Automation that produces faster than people can verify produces noise.
- **Validate, then revalidate.** Test a tool against known cases before you rely on it, and again whenever the standards change. The virtualization revisions to the CIP standards take effect July 1, 2028, and CIP-015-1 internal network security monitoring follows on October 1, 2028. Screening logic built on today's versions will need rework (Chapter 5).
- **The auditor will ask who made the decision.** "The tool" is not an acceptable answer.

1 The State of AI in Compliance: Reality vs. Hype

Start with what the tools do well, because that list is real. Then look hard at where they fall short, because that list is where the enforcement exposure lives.

1.1 What Current AI Actually Does Well

The tools delivering real value in compliance programs have one thing in common. They are pointed at problems with large volumes of structured or semi-structured data, repetitive pattern recognition, and well-defined outputs.

An evidence management platform that ingests log files, configuration records and access review data, then flags items outside defined parameters, is doing something useful. So is a tool that cross-references the asset inventory against firewall rules and finds the discrepancies far faster than a manual reconciliation. A system that tracks deadlines across a portfolio of requirements and escalates items as their windows close prevents the patch and access review failures that come from plain administrative oversight.

These are not trivial applications. A high impact CIP environment can carry hundreds of applicable Cyber Assets and dozens of requirements running on independent clocks. Keeping that evidence audit-ready is a heavy coordination job. Tools that take over the tracking and flagging give compliance staff back the time they need for the analytical work that decides whether the program holds up under scrutiny.

1.2 Where AI Falls Short in Compliance Contexts

The failures cluster in three areas: regulatory interpretation, contextual judgment, and audit defensibility. Those happen to be the areas where compliance expertise matters most.

Regulatory interpretation. Deciding what a requirement demands in a specific operational context, how an ambiguous standard applies to an edge case, or whether a control design is defensible on the evidence available takes more than pattern recognition. It takes knowledge of how the standard was developed, how Regional Entities have applied it, how enforcement actions have characterized similar situations, and how a particular audit team is likely to approach the question. In my view no current AI system holds that context in a form that produces reliable regulatory guidance. A system that generates interpretations without it produces answers that look authoritative and may be wrong in ways that create enforcement exposure.

Contextual judgment. Why does a control gap exist? Is it a genuine reliability risk or a documentation artifact? How should it be characterized in a self-report that is accurate and also sound strategy? Those are human calls. A tool that spots a gap and recommends a response

by pattern matching can recommend something technically correct and strategically damaging. The professional who understands the regulatory dynamics, the entity's relationship with its Regional Entity, and the full context of the finding will make the better decision.

1.3 The Audit Defensibility Problem

This is the biggest limitation. When an auditor asks who made a determination and how, "the AI system flagged it" is not a defensible answer. Classification decisions, patch exception justifications, access review outcomes and incident characterizations all have to trace back to an identifiable person with articulable reasoning. An automated process that produces compliance outputs without human review and sign-off creates an evidence problem: the output exists, but the accountable determination does not.

That follows directly from how the framework is built. Under 18 C.F.R. § 39.2(b), the entities subject to FERC's reliability jurisdiction must comply with the applicable Reliability Standards. The obligation sits with the registered entity, not with its software. When compliance decisions come out of an automated system with no documented human review, the chain of accountability is broken. I treat that as a program deficiency in its own right, separate from whatever the system decided.

1.4 End-of-Chapter Summary

The tools generating real value today work on high-volume, pattern-dependent administrative tasks: evidence tracking, deadline management, anomaly flagging and cross-reference validation. They fall short on regulatory interpretation, contextual judgment and audit-defensible decision-making, which is exactly where experienced compliance professionals earn their keep. Knowing the difference is the prerequisite for using AI without creating new risk.

> **From the Field**
>
> The compliance functions where AI works are narrow and specific. The functions where AI fails quietly are broader and more consequential. Knowing which is which is the program's job, not the vendor's.

2 High-Value Automation Opportunities in CIP Compliance

The activities best suited to automation share a profile. They are high in volume, rule-based and deadline-sensitive, and they produce outputs an experienced professional can review quickly. Finding those activities inside a CIP program is the starting point for an automation strategy that delivers.

2.1 Evidence Management and Cross-Reference Validation

Cross-referencing evidence across control areas is one of the most time-consuming parts of CIP audit preparation. The asset inventory behind your CIP-002 categorization has to agree with the systems in your CIP-005 Electronic Security Perimeter documentation, the systems covered by your CIP-007 patch records, and the systems in your CIP-010 baseline configurations. When those four evidence streams describe different populations of assets, and they can, you have both an audit risk and a manual investigation on your hands.

Automation can shrink that investigation dramatically. A tool that ingests the inventory, the network documentation, the patch records and the baselines, and flags any system that shows up in one stream but not another, takes the hunting out of that job. The important part: the tool is not making a compliance determination. It is surfacing information for a compliance professional to evaluate. That distinction is what keeps the output usable in an audit.

2.2 Deadline and Timeline Management

CIP compliance runs on many independent clocks. Tracking those clocks, escalating items as windows close, and keeping a live view of deadline status is not sophisticated AI. It is basic workflow management applied to a structured problem. It also goes straight at the most-reported problems. NERC's ORCP and CMEP Annual Report (February 2026) names CIP-010, CIP-007 and CIP-004 as the most frequently reported CIP noncompliance in 2025 and notes that all three involve high volume and high frequency conduct. WECC's May 2026 Oversight Trends Update reports that CIP self-reports in the West rose 50 percent over four years, with patch management under CIP-007-6 R2 among the notable upward trends.

The failure that tracking prevents can be mundane: a patch evaluation that was done but documented after the window closed, or an access verification that happened but left no dated record inside the quarter. Those aren't complex control failures. They are administrative failures, and automation is well suited to preventing them.

A few of the clocks that matter most for automation:

Requirement	What has to happen	Clock
CIP-007-6 R2, Part 2.2	Evaluate security patches released since the last evaluation for applicability	At least once every 35 calendar days
CIP-007-6 R2, Part 2.3	For applicable patches, apply them, create a dated mitigation plan, or revise an existing plan	Within 35 calendar days of completing the evaluation
CIP-004-7 R4, Part 4.2	Verify that individuals with active electronic or unescorted physical access have authorization records	At least once each calendar quarter
CIP-004-7 R4, Part 4.3	Verify that user accounts, groups or role categories and their privileges are correct and necessary	At least once every 15 calendar months
CIP-010-4 R1, Part 1.3	Update the baseline configuration for a change that deviates from it	Within 30 calendar days of completing the change
CIP-010-4 R3, Part 3.1	Conduct a paper or active vulnerability assessment	At least once every 15 calendar months

Each of these applies to defined high and medium impact systems; check the applicable systems column of the standard before you build a calendar around it. In a high impact environment, the number of clocks running at any moment is substantial.

2.3 RSAW Population and Evidence Mapping

A Reliability Standard Audit Worksheet is a guide from the ERO that describes the types of evidence a registered entity may use to demonstrate compliance with a standard, along with how the ERO may assess that evidence. NERC is explicit that RSAWs don't require specific evidence and can't change the scope or intent of a standard. Still, they are the best public map of what an audit team will look for, and preparing against them can be partly automated.

Tools that pull evidence from defined sources, organize it against the RSAW structure, and flag where expected evidence is missing or out of bounds cut the manual work and make the output more consistent. The limitation is the same as everywhere else in this paper. The tool can assemble and format evidence. It can't decide whether that evidence satisfies the requirement given the entity's operations and control design. That call takes human review. The value of automation is the time it saves on assembly, which a compliance professional can then spend on the review that makes the package defensible.

2.4 Misclassification Screening

Asset inventories in large CIP environments never sit still. Systems are added, modified and retired continuously. Automation that screens the active inventory against the CIP-002-5.1a Attachment 1 criteria on a rolling basis, and flags systems whose characteristics suggest a possible misclassification, gives you continuous coverage that a periodic manual review can't match.

The design rule for this application is simple: the tool produces a flag, not a classification. A tool that marks a system as possibly meeting medium impact criteria is raising a question for a

compliance professional. It is not making the decision. CIP-002-5.1a R2 requires the entity to review its identifications at least once every 15 calendar months and to have its CIP Senior Manager or delegate approve them on the same cycle. The standard itself puts a named human at the end of the process. Automation can make that process faster. It can't make it automatic.

2.5 End-of-Chapter Summary

The highest-value automation targets in CIP compliance are evidence cross-reference validation, deadline and timeline management, RSAW-structured evidence assembly, and misclassification screening. Each one strips out administrative work and frees professionals for the judgment the tool can't supply. In every case the automation produces inputs for human review, not outputs that replace it.

> **From the Field**
>
> Patch tracking, evidence collection, log review aggregation and inventory reconciliation are good automation candidates. Judgment-heavy activities are not. Automation should expand what your compliance team can do, not replace what it does.

3 What Must Not Be Automated

Some compliance functions can't be responsibly automated because human judgment, regulatory expertise and institutional context aren't a supplement to the process. They are the process. Automating them doesn't improve outcomes. It introduces audit exposure that didn't exist before.

3.1 Regulatory Interpretation and Standard Application

Every CIP program hits questions the standard's language doesn't settle. Is this system inside or outside the Electronic Security Perimeter? Does this operational practice constitute Interactive Remote Access? Does this account change call for a new access review, or does it fall within the existing one? Answering takes regulatory judgment: knowledge of the standard's intent, familiarity with enforcement history, and an understanding of how Regional Entities have treated similar questions.

An AI tool trained on published standards, guidance documents and enforcement actions can produce answers that sound right and are wrong. In a regulatory interpretation, a wrong answer isn't just an error. It becomes a documented compliance decision that will be examined at audit. An entity whose classification determinations, access review methods or ESP scoping decisions were guided by AI-generated interpretations has a problem bigger than any one determination. It has a process failure in how it makes compliance decisions.

3.2 Audit Interactions and SME Preparation

The human interactions that decide audit outcomes can't be delegated to software: subject matter expert interviews, preliminary findings discussions, formal responses to requests for information. Auditors are trained to judge whether the people they interview understand and own the controls they describe. An SME who was briefed by an AI tool on what to say, rather than prepared by a compliance professional who understands both the control and the regulatory context, won't hold up under follow-up questions.

There is a more basic issue too. The relationship between a registered entity and its Regional Entity runs across many audit cycles, enforcement matters and standards projects. How the entity manages it, through the candor of its disclosures, the quality of its corrective actions and the credibility of its representations, shapes how the Regional Entity approaches the program over time. That takes human judgment, institutional knowledge and professional credibility. You can't outsource it to a tool.

3.3 Enforcement Response and Self-Report Strategy

When a gap turns up, whether through internal review, an audit or an event investigation, the decisions that follow are among the most consequential a program makes. Whether to self-report, and when. How to characterize the scope and duration of the violation. How to frame the

corrective actions. How to engage the Regional Entity through the enforcement process. Each of those carries financial, reputational and regulatory consequences, and each calls for experienced judgment about the dynamics involved.

A tool that recommends an enforcement strategy by pattern matching against past enforcement actions is offering background. It is not offering a recommendation anyone should act on without independent analysis. The situations that lead to enforcement are specific, fact-intensive and shaped by context no training set captures. Responding to them takes an experienced person.

3.4 The Accountability Imperative

One principle runs under all of these limits: accountability. The obligation to comply rests with the registered entity, and in my view that means decisions have to be made by identifiable people with reasons they can state, not by algorithms producing pattern-matched outputs. The analysis behind a classification, the reasoning for an exception to a patch timeline, the characterization of a self-reported violation: each has to reflect human judgment that can be explained and defended.

Automation that produces those outputs without the human layer does more than create audit risk. It builds a culture that mistakes the appearance of process for the substance of accountability. The programs that hold up across multiple audit cycles are the ones where people own the decisions. They know why a control is designed the way it is, why an exception was granted, and why a self-report was characterized as it was.

3.5 End-of-Chapter Summary

Regulatory interpretation, audit interactions, enforcement response strategy and any decision that must be attributed to an identifiable person should not be automated. In each area, the efficiency automation appears to offer hides a loss in the quality that sustains a program under scrutiny. Recognizing that boundary, and automating only where it improves compliance rather than substituting for it, is the mark of a mature program.

> **From the Field**
>
> Classification decisions, scope determinations and audit-defense narrative need judgment, accountability and the ability to explain reasoning under questioning.

4 Implementation: How to Deploy Automation Without Creating New Risk

Knowing which functions benefit from automation is one thing. Deploying it so the benefit shows up is another. The common failures are tools that don't connect to the existing evidence process, automation with no review step built around it, and systems that generate compliance-looking documents with no compliance underneath. They are common, and they are avoidable.

4.1 Start With the Process, Not the Tool

The most common failure is buying a tool before understanding the process it is supposed to support. Deploy a deadline tool in an organization that hasn't documented the timelines for each requirement, and it will track the wrong clocks or none at all. Deploy a cross-reference tool where the evidence streams don't share a consistent asset identifier, and it will produce results nobody can reconcile. The tool amplifies the process. If the process is broken, the tool amplifies the break.

Good implementation starts with process documentation. Map every compliance activity, define its inputs and outputs, and identify where automation would cut the burden without degrading quality. That mapping shows where the process is inconsistent, where the data is weak, and where the human review steps that preserve accountability need to be formalized before any tool goes in. The tool comes after the process is understood and written down.

4.2 Human Review Is Not Optional

Every automated output, whether a flagged inconsistency, a populated evidence section, a deadline escalation or a possible misclassification, has to pass through a defined human review step before anyone treats it as a compliance determination. That review has to be documented: who reviewed it, when, what was decided, and on what basis. Without that record, the output isn't a compliance record. It's a data point with no one accountable for it.

Design the review step into the workflow before deployment. Don't assume reviewers will add it on their own. That choice is the difference between automation that strengthens the program and automation that creates the illusion of it. The review step isn't bureaucratic overhead. It is the mechanism that makes the output defensible.

4.3 Validate Before You Trust

Test every tool against known cases before you rely on it in production. Run the cross-reference tool against an evidence package with planted inconsistencies and confirm it finds them. Run

the classification screen against a facility whose categorization you already know and confirm it flags, or declines to flag, the right systems. Run the deadline tool against a past compliance calendar and confirm it would have escalated each item inside its window.

Validation isn't a one-time event. Repeat it when requirements change, when the operating environment changes, and when the tool itself is updated. A tool that was accurate at deployment can go wrong the day a new standard version changes the applicable systems, terms or clocks. The coming CIP revisions described in Chapter 5 are exactly that kind of change. The person responsible for the tool is responsible for knowing when revalidation is due and making sure it happens.

4.4 The Expert Remains Indispensable

The most important principle is this: automation works best when it is built by people who understand the compliance domain it serves. A patch tracking tool has to be configured by someone who knows the two CIP-007-6 R2 clocks in the Section 2.2 table. A classification screen has to be designed by someone who knows the CIP-002 criteria well enough to write the screening logic correctly. An RSAW evidence tool has to be built by someone who knows what evidence each requirement calls for.

Without that expertise, the tool will systematize the team's misunderstandings instead of its knowledge. Systematized misunderstandings produce consistent, voluminous, audit-ready documentation of compliance failures.

The compliance expert who understands the framework, has seen audits from both sides of the table, and can work through the enforcement process isn't made obsolete by AI. That person becomes more valuable. Their expertise is what gets the automation to the right outputs, and their judgment, applied in the review step, is what makes those outputs defensible. When the audit interviews start, no tool stands in for the relationships and credibility that person brings.

4.5 End-of-Chapter Summary

Effective implementation starts with process documentation, builds in human review, validates against known cases before production, and depends on compliance expertise in both design and oversight. Automation deployed without those elements is more dangerous than not automating at all, because its failures are harder to detect and harder to correct.

> **From the Field**
>
> Automation brings two new risks into a compliance program: the tool can fail, and the organization can come to depend on it. The mature deployment integrates automation with existing evidence management. The immature one runs it in parallel. Parallel systems diverge, and the divergence becomes a finding.

5 The Next 12 to 24 Months: What Is Changing and What Is Not

The direction of AI in compliance is clear. The pace is not. Capabilities emerging now will be more mature and more widely deployed within two years. The programs best positioned for that are building automation on a foundation of process discipline and human expertise, not replacing the foundation with automation. What follows is my read on the trend, plus the regulatory changes that are already on the calendar.

5.1 What Will Change

Evidence processing will improve. Tools that ingest unstructured evidence, such as audit logs in proprietary formats, configuration files from operational technology, and communication records, and pull structured, compliance-relevant information out of them are getting more capable. That matters, because much of the manual labor in CIP audit preparation is translating operational records into a form an auditor can evaluate. Automating the translation frees professionals for the analysis.

Anomaly detection will improve. Systems that find patterns in large evidence sets that don't fit compliant behavior will become more reliable and more available: access patterns that suggest shared account use, configuration drift that suggests unauthorized change, patch records that cluster in ways that suggest back-dated documentation. These capabilities don't replace the judgment about whether an anomaly is a finding. They surface the anomalies that deserve that judgment.

Integration will deepen. Evidence that is assembled by hand today, such as access logs from identity management, baselines from change management platforms, and patch records from vulnerability management tools, will increasingly be assembled through system integrations. The team that spends its time on evidence assembly now will spend it on evidence evaluation instead.

5.2 What Is Already on the Calendar

Four CIP changes are approved and scheduled, and all of them bear directly on compliance tooling. NERC's Reliability Standards listing shows them as subject to future enforcement.

The currently enforceable versions cited in this paper (CIP-002-5.1a, CIP-003-9, CIP-004-7, CIP-005-7, CIP-007-6 and CIP-010-4) remain in effect until July 1, 2028. CIP-003-11 (FERC Order No. 918) supersedes CIP-003-10 on July 1, 2029.

Change	Standards	Effective date
Virtualization revisions (NERC Project 2016-02; FERC Docket RM24-8-000)	CIP-003-10, CIP-004-8, CIP-005-8, CIP-007-7.1 and CIP-010-5, with companion revisions across the CIP family	July 1, 2028
Further CIP-002 revision (FERC Docket RD25-8-000)	CIP-002-8 (supersedes CIP-002-7, which is retired immediately before CIP-002-8 takes effect)	July 1, 2028
Internal network security monitoring (NERC Project 2023-03; FERC Docket RM24-7-000)	CIP-015-1	October 1, 2028
INSM extension	CIP-015-2, which extends monitoring to the Applicable Systems in its Table R1, including EACMS, PACS and Shared Cyber Infrastructure associated with high impact and medium impact (with External Routable Connectivity) BES Cyber Systems	October 1, 2029

For a compliance program running automation, the practical point is revalidation. Any screening logic, evidence map or deadline calendar built on today's versions has to be checked against the new text before the effective date, not after. CIP-015-1 matters for another reason. It requires documented internal network security monitoring inside the Electronic Security Perimeters of high impact BES Cyber Systems and medium impact BES Cyber Systems with External Routable Connectivity, with the stated purpose of improving the probability of detecting anomalous or unauthorized network activity. That is anomaly detection becoming a requirement: the tool can surface the activity, and a person has to evaluate it and own the response.

5.3 What Will Not Change

The framework that defines compliance won't change because AI is available to help meet it. The NERC Reliability Standards set mandatory requirements. The Compliance Monitoring and Enforcement Program monitors and enforces them. Audits evaluate whether they've been met, whatever tools the entity uses.

The premium on experienced judgment won't shrink. In my view it will grow. As AI makes professional-looking documents easier for less experienced teams to produce, telling real control apart from automated assembly matters more, and the experience to make that call will be worth more.

The human relationships that carry programs through audits and enforcement won't change either. The auditor judging whether SMEs own their controls and the enforcement staff judging whether a self-report is candid are making human calls, and the people who keep those relationships provide value no automation replicates.

5.4 End-of-Chapter Summary

Evidence processing, anomaly detection and integration will improve, and the 2028 CIP revisions will force revalidation of every tool built on today's standards. The framework won't change. The programs that come out ahead will use AI to extend experienced professionals, not replace them.

Glossary

Term	Meaning in this paper
Artificial Intelligence (AI)	In the compliance context, software systems that use machine learning, natural language processing, or pattern recognition to assist with compliance functions. Current AI tools are best suited to high-volume, pattern-dependent tasks and are not a substitute for regulatory judgment or accountable decision-making.
Audit Defensibility	The quality of a compliance determination or process that lets it withstand auditor scrutiny. A defensible determination is attributable to an identifiable person, supported by documented reasoning, and consistent with the applicable Reliability Standard.
Automation	The use of software to perform compliance functions that would otherwise take manual effort, such as deadline tracking, evidence cross-referencing, and anomaly detection. Most valuable on high-volume, rule-based tasks; most risky on judgment-intensive decisions.
CIP Senior Manager	The individual a Responsible Entity must identify by name under CIP-003-9 R3, with any change documented within 30 calendar days. CIP-002-5.1a R2 requires the CIP Senior Manager or delegate to approve the entity's BES Cyber System identifications at least once every 15 calendar months.
Compliance Monitoring and Enforcement Program (CMEP)	Under the NERC Rules of Procedure, Appendix 2, either the NERC CMEP in Appendix 4C (or a Regional Entity's Commission-approved program), or the NERC or Regional Entity organization responsible for monitoring and enforcing registered entities' compliance with Reliability Standards.
Evidence	Documentation showing that a required control exists and was implemented during the relevant period, consistently. AI tools can help assemble and organize evidence, but its validity depends on human review and attestation.
Human Review	The documented evaluation of an automated output by an identifiable compliance professional who takes responsibility for the determination the output reflects. It is the accountability mechanism that makes automated outputs defensible.
Internal Network Security Monitoring (INSM)	Monitoring of network activity inside the Electronic Security Perimeter, with the stated purpose in CIP-015-1 of improving the probability of detecting anomalous or unauthorized network activity. CIP-015-1 takes effect October 1, 2028.
Machine Learning	A form of AI in which systems learn patterns from large datasets and apply them to new inputs. Compliance uses include anomaly detection, evidence classification, and timeline prediction. Outputs need human validation before they are treated as compliance determinations.

Term	Meaning in this paper
Reliability Standard Audit Worksheet (RSAW)	A guide provided by the ERO that describes types of evidence registered entities may use to demonstrate compliance with a Reliability Standard and how the ERO may assess that evidence. RSAWs exist across the standard families, not only CIP. They do not require specific evidence and may not change the scope or intent of a standard.
Validation	Verifying that an automated tool produces correct outputs in the compliance context where it is deployed. Done before initial deployment and repeated when requirements, the operating environment, or the tool change.
Workflow Automation	Software that manages the sequence and timing of compliance tasks, including deadline tracking, escalation, evidence collection triggers, and review routing. Among the highest-value and lowest-risk uses of automation in a compliance program.

How Energy Compliance Can Help

Everything in this paper comes down to two jobs: building automation on a sound process, and keeping a named person accountable for every determination. That is the work we do.

Automation implementation advisory. Process design before tool selection. We map your compliance activities, document the clocks and inputs each requirement runs on, and design the human review steps into the workflow before any tool is switched on. When the 2028 CIP revisions arrive, we revalidate the logic against the new text.

Evidence management architecture and quality review. We reconcile the asset populations across your CIP-002, CIP-005, CIP-007 and CIP-010 evidence, check that every automated output carries a documented review, and find the gaps before an auditor does.

Mock audit preparation with AI-assisted evidence package assessment. We assemble the evidence package with automation, then review it the way an audit team would, including SME preparation, so the people answering questions own the controls they describe.

WattOps Compliance. WattOps is EC's NERC compliance evidence platform (wattopscompliance.com). It gives the evidence, review records and deadlines in Chapters 2 and 4 one home. To see how your program fits, EC will run one standard at one facility through WattOps at no cost. CIP standards are excluded from that offer.

RenewOps. For renewable and storage operators, much of the operational record behind compliance starts in SCADA, alarms and availability data. RenewOps is EC's renewable operations monitoring and analytics platform, with a Data Center module, and it keeps that operations data in one place where your compliance team can find it.

Pick one automated output your program relies on today and ask who reviewed it, when, and on what basis. If the answer takes more than a minute, call me and we'll work through it together.

Website: energycomplianceinc.com · Email: rsmith@complianceforenergy.com

Phone: 763.438.4427

About the Author

Rob Smith is the Founder and Principal of Energy Compliance, Inc., with more than thirty years in the industry on every side of the system. He began in control-room operations as a Reliability Coordinator, Transmission Operator, and Power System Operator across RTO/ISO and utility control centers. On the regulator side, he was a senior compliance auditor and subject matter expert for NERC Reliability Standards. He audited grid facilities, evaluated mitigation adequacy, participated in risk-based oversight of utility mitigation activities, and supported the development of violation notifications and settlements in FERC-directed enforcement actions. He conducted a regulatory audit in the Sultanate of Oman against the Sultanate's Sector Law and Grid Code.

He founded Energy Compliance, Inc. to deliver senior regulator-side compliance authority without the layered staffing and billable-hour overhead of larger firms.

Education:

- MBA, Energy Management, University of Mary
- MSL, Compliance Law, Fordham Law University
- BAS, Energy Management, Bismarck State College
- BAS, Mechanical Engineering, Metallurgy minor, University of Florida
- BAS, Criminology, St. Petersburg College (State Law Enforcement Certification)
- AAS, Power Plant Technology, Bismarck State College
- AAS, Electrical Transmission System Technology, Bismarck State College

Certifications: NERC Lead Auditor training; previously NERC-certified at the Reliability Coordinator level; 1st Class A Boiler Engineer License, Minnesota.

About Energy Compliance, Inc.

Energy Compliance, Inc. is an independent consulting and advisory firm based in Bloomington, Minnesota, focused on electric reliability, cybersecurity reliability, and regulatory compliance for organizations connected to the North American Bulk Electric System. We support registered entities, including Generator Owners and Operators, Transmission Owners and Operators, Reliability Coordinators, Balancing Authorities, and Distribution Providers, across NERC Reliability Standards, FERC orders, RTO/ISO market participation rules, Regional Entity oversight, and state regulatory frameworks.

We do this work differently than larger firms. Engagements are overseen by a senior practitioner with regulator-side experience. We don't staff for billable hours; we staff for outcomes. Where automation can replace manual work, we build the automation. Where senior judgment is required, the senior is in the room. Our approach to compliance technology rests on one principle: automation designed by people who understand the regulatory framework produces compliance value, and automation designed without that understanding produces compliance risk. EC builds its own compliance technology: WattOps Compliance, its NERC compliance evidence platform (wattopscompliance.com), and RenewOps, its renewable operations monitoring and analytics platform with a Data Center module.

Compliance technology and program services

- Compliance program design and gap assessment
- Automation implementation advisory, process design before tool selection
- Evidence management architecture and quality review
- Mock audit preparation with AI-assisted evidence package assessment
- Ongoing compliance monitoring program support
- Senior advisory on complex regulatory questions that require expert judgment

Reliability and compliance services

- NERC reliability and compliance advisory support
- Reliability governance and program assessments
- Registration and applicability analysis
- Operational and engineering reliability alignment
- Compliance program design and improvement
- Audit and enforcement support (non-advocacy)
- Mitigation planning and Self-Report development
- Training and executive briefings on reliability frameworks
- Regulator-perspective program reviews

Practice areas

- **NERC Compliance.** Program support, interpretation, and audit preparation.
- **Compliance Technology Advisory.** Automation implementation guidance grounded in regulatory expertise.
- **Senior Advisory.** Direct engagement on complex reliability and enforcement questions.
- **Industry Engagement.** Standards development and working-group participation.
- **Audit Defense.** Notice of Penalty response and settlement posture.

Each engagement is scoped to the entity's role, function, and bulk system impact.

Energy Compliance is not affiliated with, sponsored by, or endorsed by the North American Electric Reliability Corporation, the Federal Energy Regulatory Commission, or any Regional Entity.

Website: energycomplianceinc.com · Email: rsmith@complianceforenergy.com

Phone: 763.438.4427

Sources and Authorities

The views on AI and automation in this whitepaper are the professional judgment of Energy Compliance, Inc. The requirements, definitions, dates and statistics it cites come from the primary sources below, checked October 25, 2026. Confirm the current text before relying on them.

Federal regulation

- 18 C.F.R. § 39.2, Jurisdiction and applicability, paragraph (b). https://www.ecfr.gov/current/title-18/chapter-I/subchapter-B/part-39/section-39.2

NERC Reliability Standards

- NERC, Reliability Standards, CIP family listing (status and future effective dates). https://www.nerc.com/standards/reliability-standards/cip
- NERC, Reliability Standard CIP-003-9, Cyber Security, Security Management Controls, Requirement R3. https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-003-9.pdf
- NERC, Reliability Standard CIP-002-5.1a, Cyber Security, BES Cyber System Categorization, Requirement R2 and Attachment 1. https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-002-5.1a.pdf
- NERC, Reliability Standard CIP-004-7, Cyber Security, Personnel & Training, Requirement R4, Parts 4.2 and 4.3. https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-004-7.pdf
- NERC, Reliability Standard CIP-005-7, Cyber Security, Electronic Security Perimeter(s). https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-005-7.pdf
- NERC, Reliability Standard CIP-007-6, Cyber Security, System Security Management, Requirement R2, Parts 2.2 and 2.3. https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-007-6.pdf
- NERC, Reliability Standard CIP-010-4, Cyber Security, Configuration Change Management and Vulnerability Assessments, Requirement R1, Part 1.3, and Requirement R3, Part 3.1. https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-010-4.pdf
- NERC, Reliability Standard CIP-015-1, Cyber Security, Internal Network Security Monitoring, Purpose and Requirement R1. https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-015-1.pdf

- NERC, Reliability Standard CIP-015-2, Cyber Security, Internal Network Security Monitoring, Requirement R1 and Table R1. https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-015-2.pdf
- FERC, North American Electric Reliability Corporation; Order Approving Reliability Standard CIP-002-8, Docket No. RD25-8-000, Federal Register Doc. 2026-05715 (March 24, 2026). https://www.federalregister.gov/documents/2026/03/24/2026-05715/north-american-electric-reliability-corporation-order-approving-reliability-standard-cip-002-8
- FERC, Order No. 919, Virtualization Reliability Standards, Docket No. RM24-8-000, Federal Register Doc. 2026-05716 (March 24, 2026). https://www.federalregister.gov/documents/2026/03/24/2026-05716
- FERC, Order No. 918, Critical Infrastructure Protection Reliability Standard CIP-003-11, Cyber Security, Security Management Controls, Docket No. RM25-8-000, Federal Register Doc. 2026-05711 (March 24, 2026). https://www.govinfo.gov/content/pkg/FR-2026-03-24/html/2026-05711.htm
- FERC, Order No. 907, Critical Infrastructure Protection Reliability Standard CIP-015-1, Cyber Security, Internal Network Security Monitoring, Docket No. RM24-7-000, Federal Register Doc. 2025-12309 (July 2, 2025). https://www.federalregister.gov/documents/2025/07/02/2025-12309
- NERC, Project 2016-02 Modifications to CIP Standards, Implementation Plan, April 3, 2024. https://www.nerc.com/globalassets/standards/reliability-standards/cip/2016-02-virtualization-implementation-plan_clean_04032024.pdf

NERC Rules of Procedure and audit tools

- NERC, Rules of Procedure Appendix 2, Definitions Used in the Rules of Procedure, effective June 27, 2024 (Compliance Monitoring and Enforcement Program). https://www.nerc.com/globalassets/who-we-are/rules-of-procedure/appendix-2.pdf
- NERC, Rules of Procedure Appendix 4C, Compliance Monitoring and Enforcement Program, effective May 19, 2022. https://www.nerc.com/globalassets/who-we-are/rules-of-procedure/appendix-4c.pdf
- NERC, Reliability Standard Audit Worksheets (RSAWs). https://www.nerc.com/programs/compliance/reliability-standard-audit-worksheets-rsaws
- NERC, Reliability Standard Audit Worksheet (RSAW) Review and Revision Process. https://www.nerc.com/globalassets/programs/compliance/rsaws/nerc-reliability-standard-audit-worksheet-rsaw-review-and-revision-process.pdf

Compliance trends

- NERC, ORCP and CMEP Annual Report (February 11, 2026). https://www.nerc.com/globalassets/programs/enforcement/cmep-and-vegetation-reports/2025-orcp-and-cmep-annual-report.pdf
- WECC, Oversight Trends Update, data as of May 1, 2026. https://www.wecc.org/sites/default/files/documents/program/2026/Oversight%20Trends%20Update%202026-05.pdf

© Energy Compliance, Inc. · EC-WP-602 · Status as of October 25, 2026