CYBERSECURITY · CIP SERIES · EC-WP-203

High Impact CIP: The Definitive Guide

High Impact CIP for Control Center BES Cyber Systems. Scale, architecture, and audit defensibility.

High impact CIP is not medium impact with extra rules. It covers the Control Centers that run the grid, it carries requirement parts no other tier carries, and it fails in architectural ways that paperwork cannot fix. This reference covers the requirements in force in October 2026, where high impact programs break under audit, and the changes already approved for 2028.

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 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, and the reader should consult qualified counsel and subject matter experts when making compliance decisions.

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.

Standard versions, effective dates, and regulatory citations in this edition were checked against the primary sources listed in Sources and Authorities as of October 2026. Standards change. Confirm the current version before relying on any requirement described here.

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.

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 most 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

High impact CIP is not medium impact with extra rules. It is a different reliability conversation, with different consequence assumptions and different audit expectations. Under Attachment 1 of CIP-002-5.1a, high impact reaches only the BES Cyber Systems used by and located at certain Control Centers and backup Control Centers of Reliability Coordinators, Balancing Authorities, Transmission Operators, and Generator Operators. The bright-line criteria get you classified. What gets you defensible is everything that happens after.

High impact systems carry requirement parts that no other tier carries: two or more different physical access controls into the Physical Security Perimeter, where technically feasible; log review at least every 15 days; baseline monitoring at least every 35 days; testing changes, in a test environment or carefully in production, before implementing them, where technically feasible; an active vulnerability assessment at least every 36 months, where technically feasible; and a recovery exercise at least every 36 months. Those parts are where the scale of a high impact environment turns into evidence volume.

The points below are my read of where high impact programs succeed and fail. Each one is developed, with its sources, in the chapters that follow.

- If your program treats high impact like upgraded medium, you've miscalibrated. The standard isn't more of the same.
- ESP sprawl is the architecture problem most high impact programs eventually confront. The first version was clean. Five years of exceptions later, the boundary has eroded.
- Access control is where operations pushes back hardest. NERC's 2025 annual report lists CIP-010, CIP-007, and CIP-004 as the most frequently reported CIP noncompliance, and all three carry access or configuration discipline at their core.
- Compliance in a high impact environment is documentation discipline. The control might be perfect; without evidence, the audit produces a finding.
- The mature programs aren't audit-driven. They're risk-driven. The audit is a checkpoint, and the program is the work.
- At high impact scale, manual reconciliation stops being practical. Tooling helps, but only if the tool's outputs are themselves defensible.

The ground is also moving. The virtualization revisions to the CIP standards, a revised CIP-002-8, and new definitions of the Electronic Security Perimeter and Electronic Access Point take effect on July 1, 2028. Internal network security monitoring under CIP-015-1 follows on October 1, 2028. High impact programs should be planning for both now.

1 What Actually Makes High Impact Different

The most dangerous assumption in high impact CIP is that it is simply a more demanding version of medium impact. It isn't. High impact differs in architecture, interdependency, and failure consequence.

1.1 The Bright-Line Threshold Is Just the Beginning

CIP-002-5.1a, Attachment 1, Section 1 defines high impact by function, not by asset type. Each criterion covers the BES Cyber Systems used by and located at a Control Center or backup Control Center performing a specific set of functional obligations.

Criterion	Control Center or backup Control Center used to perform the functional obligations of
1.1	The Reliability Coordinator
1.2	The Balancing Authority, for generation of an aggregate 3,000 MW or more in a single Interconnection, or for one or more assets that meet criterion 2.3, 2.6, or 2.9
1.3	The Transmission Operator, for one or more assets that meet criterion 2.2, 2.4, 2.5, 2.7, 2.8, 2.9, or 2.10
1.4	The Generator Operator, for one or more assets that meet criterion 2.1, 2.3, 2.6, or 2.9

Two things in that table get missed. Backup Control Centers are in, on the same terms as primary ones. And the Guidelines and Technical Basis in CIP-002-5.1a make the point that the criteria follow functional obligations, not the name on the building: where a Transmission Operator delegates some of its obligations to a Transmission Owner, the Transmission Owner's Control Center systems performing them are subject to categorization as high impact.

What entities often discover too late is that the categorization itself is only the first challenge. Once a system is high impact, the architecture of every supporting system, every communication path, every access mechanism, and every vendor relationship becomes relevant. Protected Cyber Assets inside the same Electronic Security Perimeter take the high impact rating by definition. The Electronic Access Control or Monitoring Systems (EACMS) and Physical Access Control Systems (PACS) associated with the high impact system carry their own requirement parts. The categorization sets off a chain of applicability determinations that reaches far deeper into the operational technology environment than most entities anticipate.

1.2 Different Failure Modes, Not Just More Requirements

Medium impact environments fail in predictable ways: incomplete asset inventories, missing patch evaluations, access review documentation gaps. These are governance failures, process problems with process solutions.

High impact environments fail differently. The failure modes at this tier are architectural. An Electronic Security Perimeter scoped too narrowly leaves systems inside the operational environment without the protections they need. One scoped too broadly creates an evidence burden that cannot be sustained across an audit cycle. Jump hosts added for operational convenience but never identified as EACMS create access pathways nobody is controlling. Shared accounts that predate CIP and were never inventoried become indefensible during an account management review. Legacy systems that cannot accept patches, and were never covered by a dated mitigation plan, accumulate enforcement exposure silently.

These are not documentation problems. They are engineering problems that show up as compliance problems. The distinction matters because the solution is different. Adding more spreadsheets does not fix a perimeter scoping error. Only architecture does.

1.3 The Scale Problem Is Real

A medium impact entity may manage dozens of BES Cyber Assets across a handful of ESPs. A high impact Control Center may manage hundreds of Cyber Assets across a multi-layered architecture that includes External Routable Connectivity, Interactive Remote Access for several vendors, and interdependencies with systems that touch both BES and non-BES environments.

At that scale, the administrative burden of maintaining consistent, traceable evidence across every requirement is significant. This is where high impact programs most commonly build up invisible technical debt: evidence that exists on paper but cannot be substantiated when an auditor asks for cross-referenced documentation across CIP-005, CIP-007, and CIP-010 at the same time.

1.4 What High Impact Adds

Most CIP requirement parts apply to both high and medium impact systems. The parts below apply to high impact systems and not to medium. Each one is a recurring obligation, which means each one generates evidence on a clock.

Standard and part	High impact obligation
CIP-004-7 R5.3	Revoke a terminated individual's non-shared user accounts within 30 calendar days of the termination
CIP-004-7 R5.4	Change passwords for shared accounts known to a terminated or transferred user within 30 calendar days
CIP-006-6 R1.3	Where technically feasible, use two or more different physical access controls for unescorted access into the PSP
CIP-007-6 R4.4	Review a summarization or sampling of logged events at intervals of no more than 15 calendar days

Standard and part	High impact obligation
CIP-009-6 R2.3	Test each recovery plan through an operational exercise at least once every 36 calendar months
CIP-010-4 R1.5	Where technically feasible, test changes in a test environment or a production environment that minimizes adverse effects before implementing them
CIP-010-4 R2.1	Monitor for changes to the baseline configuration at least once every 35 calendar days
CIP-010-4 R3.2	Where technically feasible, perform an active vulnerability assessment at least once every 36 calendar months

High impact systems also sit inside the supply chain plan required by CIP-013-2, and every high impact BES Cyber System sits at a Control Center, which brings the data links between Control Centers under CIP-012-2 (effective July 1, 2026).

1.5 Chapter Summary

High impact is a different class of compliance challenge: architectural, interdependent, and less forgiving of the governance gaps that medium impact entities routinely manage through process correction. Understanding that distinction is the starting point for a program that survives audit.

> **From the Field**
>
> Classify the backup Control Center, and any Transmission Owner Control Center performing delegated Transmission Operator functions, with the same rigor as the primary. Attachment 1 does.

2 Architectural Truths That No One Writes Down

The CIP standards describe what must be protected. They do not describe how to architect a defensible protection scheme in a real operational environment. That gap, between what the standard requires and what the engineering actually supports, is where most high impact failures originate.

2.1 ESP Sprawl Is the Primary Architecture Problem

The Electronic Security Perimeter is the foundational construct of CIP cybersecurity. Under CIP-005-7 R1, every applicable Cyber Asset connected to a network by a routable protocol must reside within a defined ESP. All External Routable Connectivity must pass through an identified Electronic Access Point. Each EAP must require inbound and outbound access permissions, with the reason for granting access, and deny all other access by default. For high impact systems, each EAP must also have one or more methods for detecting known or suspected malicious communications, inbound and outbound. The architecture of the ESP therefore determines the architecture of the entire compliance program.

In practice, ESPs in high impact environments grow organically rather than by design. A system is added here, a communication path is established there, a vendor connection is opened for a maintenance window and never formally closed. Over time, the documented ESP boundary stops reflecting the actual network topology. That is ESP sprawl, and in my view it is the most common architectural problem in high impact environments.

When auditors ask for current network diagrams under CIP-005 and compare them against firewall rule sets, access control lists, and system inventories, ESP sprawl becomes immediately visible. The question that follows is simple: why does this system appear in your firewall rules but not in your ESP documentation? There is no good answer when the discrepancy has been accumulating for years.

2.2 Paper Segmentation Is Not Segmentation

A related problem is what practitioners sometimes call paper segmentation: architectures that look segmented in compliance documentation but are not functionally segmented in the operational environment. It shows up most often in two ways.

The first is firewall rules that technically restrict access but, in practice, permit connections broad enough to make the restriction meaningless. A rule permitting access from an entire VLAN to an ESP-protected system on a specific port does not provide meaningful segmentation if that VLAN contains dozens of systems with different levels of access control. A rule like that is also hard to square with the CIP-005-7 R1.3 requirement to document the reason for each

permission. Auditors who pull firewall rule sets and cross-reference them against system inventories will find this pattern.

The second is jump hosts that were never treated as what the NERC Glossary says they are. A jump host that controls Interactive Remote Access is an Intermediate System, and the Glossary includes Intermediate Systems in the definition of EACMS. It must not be located inside the ESP. As an EACMS associated with a high impact system, it carries its own CIP-007 obligations: patch evaluation, account controls, and security event logging. The failure pattern looks like this: a jump host that sits on the wrong side of the boundary, runs software that has not been evaluated for patches on the required cycle, holds local accounts that predate the CIP program, and logs nothing the security event monitoring requirements call for. The perimeter exists on paper. The controls do not exist in practice.

2.3 The Systems Nobody Identified

The current CIP standards have no requirement that says, in one place, "identify your EACMS, PACS, and Protected Cyber Assets." Many requirements depend on doing it anyway. NERC's 2025 ORCP and CMEP Annual Report explains how the ERO Enterprise handles the gap: when an entity fails to identify a PACS, PCA, or EACMS and violates other CIP requirements as a result, the failure is processed as a single violation under CIP-010 R1, with every other missed requirement documented in that violation and covered in its mitigation. NERC describes that practice as temporary, pending a permanent fix through standards development. On October 15, 2025, the Standards Committee authorized drafting under Project 2021-03 (CIP-002).

For a high impact program, the practical point is that an unidentified EACMS is not one gap. It is every requirement that should have applied to it. The jump host from the previous section is the classic example.

2.4 The Vendor Architecture Problem

High impact Control Centers almost universally depend on vendors with some form of remote access. Energy management systems, SCADA platforms, and protection engineering tools need periodic vendor support that cannot practically be delivered on site every time.

CIP-005-7 R2 sets the rules for that access. All Interactive Remote Access must go through an Intermediate System, so the initiating device never reaches an applicable Cyber Asset directly. The session must be encrypted to the Intermediate System. Multi-factor authentication is required for every Interactive Remote Access session. The entity must also have methods to determine when vendor remote access sessions are active, including system-to-system sessions, and methods to disable them. In many high impact environments, the vendor access architecture predates these provisions. Connections that were established early and never re-evaluated against the current standard are both a control gap and an evidence gap.

The problem compounds when a vendor's remote access tool cannot support multi-factor authentication, or when operational constraints make it impractical to route a vendor through a fully monitored intermediate system. These are real constraints. They are not exceptions to the requirement. CIP-005-7 R2 applies "where technically feasible," and Appendix 4D of the NERC Rules of Procedure sets out the formal process for a Technical Feasibility Exception. Its footnoted list, current as of July 1, 2016, names CIP-005-5 R2 Parts 2.1 through 2.3, and NERC

keeps the current list on its website. An entity that treats a vendor limitation as an implicit exception, without a TFE and without documented compensating measures, will be in a difficult position at audit.

2.5 The Architecture Changes Already Approved

Two sets of changes will reshape high impact architecture work in the next three years.

The virtualization revisions developed under NERC Project 2016-02 take effect on July 1, 2028. They include CIP-005-8, retitled "BES Cyber System Logical Isolation," and revised Glossary definitions. The new ESP definition adds "a logical boundary defined by one or more Electronic Access Points," and the new EAP definition describes "an electronic policy enforcement point or a Cyber Asset interface on an Electronic Access Control or Monitoring Systems that controls routable communication to and from one or more BES Cyber Systems or their associated Protected Cyber Assets." CIP-002-8, approved by FERC in Docket No. RD25-8-000, also takes effect on July 1, 2028. CIP-002-7, the virtualization version, is retired immediately before CIP-002-8 takes effect, so CIP-002-8 is the categorization standard that applies from that date.

CIP-015-1, internal network security monitoring, follows on October 1, 2028. It requires documented processes for monitoring networks inside the ESPs of high impact BES Cyber Systems and medium impact systems with External Routable Connectivity, with network data feeds chosen on a risk-based rationale and methods to detect and evaluate anomalous activity. CIP-015-2, which extends the scope in response to FERC's directive in Order No. 907, is listed for October 1, 2029. An entity that cannot describe its own ESP accurately today will struggle to decide where those data feeds belong.

2.6 Chapter Summary

The architectural failures that expose high impact entities in audits are rarely the result of ignorance. They come from operational environments that evolve faster than compliance programs adapt. Defensible architecture needs a periodic, honest check of whether documented boundaries still match operational reality, done continuously and not only during audit preparation.

> **From the Field**
>
> Walk the firewall rules before the auditor does. Every rule that reaches into the ESP should have a documented reason and a system on the inventory at both ends.

3 Control Design at Scale: CIP-007, CIP-010, and CIP-011 Interplay

The system security management, configuration management, and information protection requirements in CIP-007, CIP-010, and CIP-011 form an interdependent control layer. It works only when the asset inventory, the baseline configurations, and the information protection procedures agree with each other. In high impact environments, keeping that agreement across hundreds of assets is the central challenge.

3.1 The Patch Management Trap

CIP-007-6 R2 is technically straightforward. Part 2.1 requires a patch management process that identifies the sources the entity tracks for each applicable Cyber Asset. Part 2.2 requires the entity to evaluate, at least once every 35 calendar days, the security patches released since the last evaluation. Part 2.3 then gives 35 calendar days from the completed evaluation to apply the applicable patches, create a dated mitigation plan, or revise an existing one. Under Part 2.4, the mitigation plan must be carried out within its own timeframe unless the CIP Senior Manager or a delegate approves a revision or an extension.

NERC's 2025 annual report ranks CIP-007 second among the most frequently reported CIP noncompliance, and WECC's May 2026 Oversight Trends Update flags an upward trend in CIP-007-6 R2.

In my view, the failure pattern is rarely that entities ignore patches. The process breaks down between the security team that tracks patch releases and the engineering team that evaluates operational impact. Evaluations start. Then they stall, because the engineering review needs a maintenance window that has not been scheduled, or the vendor has not tested the patch against the production version, or the change management process is waiting on approvals.

Thirty-five calendar days pass quickly in an operational environment. When the window in Part 2.3 closes with no patch applied and no dated mitigation plan, the entity is out of compliance. The mitigation plan has to exist, with its date, when the decision not to patch is made. A plan reconstructed during audit preparation is not the same evidence, and an experienced audit team will see the difference.

3.2 Baseline Configuration: The Foundation That Crumbles

CIP-010-4 baseline configuration requirements create a dependency that runs through the entire control architecture. Under Part 1.1, the baseline records each applicable Cyber Asset's operating system or firmware, installed software, logical network accessible ports, and security patches applied. Every change management action under CIP-010, every patch decision under CIP-007, and every vulnerability assessment traces back to the accuracy of that baseline.

In high impact environments with hundreds of Cyber Assets, baseline accuracy is a continuous maintenance job. Systems are updated. Software versions change. Settings drift. Part 1.3 requires the baseline to be updated within 30 calendar days of completing a change that deviates from it, and an installed patch is such a change. For high impact systems, Part 2.1 adds a 35-day monitoring cycle that catches the changes nobody recorded.

When auditors compare baselines against actual system configurations, any discrepancy raises an immediate question: was this change authorized and documented, or did it happen outside the change management process? The answer decides whether the finding is a baseline documentation gap or an unauthorized change. Neither is a good audit outcome. In my view, the second carries far more enforcement exposure.

3.3 CIP-011 and the Information That Falls Through the Cracks

CIP-011-3 R1 requires an information protection program with methods to identify BES Cyber System Information and methods to protect and securely handle it. R2 covers reuse and disposal of Cyber Assets that hold it. In a high impact environment that category is extensive: network diagrams, ESP boundary documentation, access control lists, firewall rule sets, account management records, and patch evaluation records all belong in the conversation. Access to BCSI is authorized and verified separately under CIP-004-7 R6.

Most entities have CIP-011 procedures. They just don't apply them to the full scope of BCSI. Configuration documents sit on file shares with broader access than the procedure contemplates. Network diagrams circulate in email threads outside controlled distribution. Vendor communications with system configuration details live in mailboxes with none of the controls the procedure describes.

When auditors ask how BCSI is protected in storage, in transit, in use, and at reuse or disposal, an entity that has applied its CIP-011 procedures unevenly will struggle to produce evidence covering the full scope. The gap between the written procedure and the operational practice is where CIP-011 findings are born.

3.4 Chapter Summary

The controls in CIP-007, CIP-010, and CIP-011 are interdependent. A weakness in one creates evidence inconsistencies that spread to the others. High impact entities that run these requirements as separate workstreams, rather than as one control architecture, will keep finding that audit preparation surfaces contradictions they cannot resolve without operational disruption.

> **From the Field**
>
> An asset inventory that doesn't reconcile with a baseline configuration that doesn't reconcile with information protection isn't three separate problems. It's one architecture problem with three audit symptoms.

4 Identity and Access: Where Audits Are Won or Lost

Access control is the area where, in my view, high impact programs most often lose an audit. Not because the requirements are ambiguous, since they are among the most explicit in the CIP family, but because the operational environments where they have to be implemented resist the discipline they demand. CIP-004 is third on NERC's 2025 list.

4.1 The Shared Account Problem Has Not Gone Away

The access requirements are split across standards, and that alone causes confusion: CIP-004-7 for authorization, verification, and revocation, CIP-007-6 R5 for system authentication, and CIP-005-7 R2 for remote access.

CIP-007-6 R5 does not ban shared accounts. It requires the entity to identify and inventory all known enabled default or other generic account types (Part 5.2) and to identify the individuals who have authorized access to each shared account (Part 5.3). For high impact systems, CIP-004-7 R5.4 adds a hard clock: when a user who knows a shared account password is terminated or no longer needs the access after a transfer, the password has to change within 30 calendar days.

Shared accounts become a problem when nobody can say who uses them. A credential shared by several operators, or used both by an automated process and by people, with no list of authorized individuals, fails Part 5.3 on its face. That problem persists because it predates CIP. Control systems commissioned before individual account management were built around shared credentials, and changing that takes coordination among operations, engineering, IT, and the vendor. The operational risk of changing production systems is real. The compliance risk of not changing them is just as real. An entity that has deferred this work for years will find the deferred risk harder and harder to explain as an "ongoing remediation effort."

4.2 Service Accounts: The Invisible Access Problem

Service accounts, the accounts used by applications, automated processes, and system-to-system communications, are often overlooked in access management programs. They do not appear in HR systems. They are not tied to individuals on access review rosters. Their permissions were often set at commissioning and never reviewed against least privilege.

In a high impact environment with a complex operational technology architecture, service accounts can outnumber individual user accounts. CIP-004-7 R4.3 requires verification, at least once every 15 calendar months, that all user accounts, account groups, or role categories and their privileges are correct and are those the entity determines are necessary. When auditors compare an account inventory with the access review records, service accounts missing from

the review are an immediate problem. When they then ask what permissions those accounts hold and who is responsible for reviewing them, an entity without a documented answer has nothing to show the auditor.

4.3 MFA: The Requirement That Creates Operational Friction

Multi-factor authentication for Interactive Remote Access is not new. It was in CIP-005-5 R2 when the version 5 framework took effect on July 1, 2016, and it remains in CIP-005-7 R2.3. Entities whose remote access architecture still does not support it have had a long time to close the gap. Yet in high impact environments MFA gaps persist, mostly because the operational technology environment resists change and because vendors whose tools lack native MFA create friction that compliance teams cannot resolve on their own.

Don't wait for the vendor. Put an Intermediate System at the boundary, a jump host or secure access gateway that enforces MFA whether or not the end system supports it. That is the architecture CIP-005-7 R2.1 already requires. Where a real technical limit remains, Section 2.4 covers the TFE path and the compensating measures that go with it.

4.4 Access Review Evidence: The Documentation That Falls Apart

For high impact systems, CIP-004-7 R4.2 requires verification at least once each calendar quarter that everyone with active electronic access or unescorted physical access has an authorization record. R4.3 adds the 15-month privilege verification described above. R5.1 requires removal of unescorted physical access and Interactive Remote Access within 24 hours of a termination action. These reviews generate the evidence auditors scrutinize most closely in this domain. The documentation has to show what was reviewed, who reviewed it, what was decided, and what was done as a result.

In practice, access review documentation often shows that a review happened without showing that anyone made a decision. A spreadsheet of account names with a date stamp does not show that the reviewer evaluated whether each account's access was still appropriate. Auditors will ask to see an account that was found to have excess permissions and was adjusted. If the answer is "we reviewed all accounts and no changes were needed," that answer needs corroboration: evidence that the review was rigorous enough to find a problem if one existed.

4.5 Chapter Summary

Access control is the most auditor-visible domain in high impact CIP. It is also the domain most often undermined by operational constraints that were allowed to persist as implicit exceptions instead of documented, managed risks. Entities that treat access management as a standing operational discipline, not a periodic compliance exercise, are the ones whose evidence auditors find credible.

5 Evidence Strategy for High Impact Environments

Compliance in a high impact environment is ultimately a documentation discipline. Not documentation for its own sake, but documentation that tells a coherent, traceable story about what controls exist, how they are implemented, how they are verified, and what happens when they fail. When that story has gaps or contradictions, findings follow.

5.1 The Sampling Problem

In high impact environments with large numbers of applicable Cyber Assets, auditors do not review evidence for every asset. They sample. NERC's CIP Evidence Request Tool User Guide describes the mechanics: Level 1 requests collect populations on detail tabs, and the audit team selects a sample size and a set of samples from them for Level 2 requests asking for detailed implementation evidence. The guide says this sampling is conducted according to the Compliance Monitoring and Enforcement Manual.

In my view, the choice of what to sample is rarely random in practice. It follows risk, complexity, and patterns found in earlier evidence. Systems that appear often in firewall rules but rarely in patch records draw attention. So do accounts that appear in access logs but not in access review records, and baselines updated recently after a long period with no updates.

That means the weakest evidence, the documentation produced most hastily for the system that got the least attention, is the evidence most likely to be examined closely. In my view, entities that prepare by bringing evidence quality up evenly across all applicable assets, rather than polishing the most visible systems, fare better at audit.

5.2 Traceability: The Test That Catches Everything

The most reliable indicator of a mature high impact program is traceability: the ability to follow a thread from a requirement to the implementing control, to the evidence that the control was implemented, to the review or test that verified it works. An auditor who can trace that thread without hitting gaps or contradictions has far less to write up.

Traceability breaks most often at the handoff between technical and compliance teams. The engineer who runs the patch process knows which patches were applied and when. The compliance team knows what the documentation requirement says. If the engineer's patch records and the compliance team's evaluation records use different asset identifiers, different version conventions, or different date formats, the chain breaks. An auditor who cannot reconcile the two sets of records will treat the gap as a finding.

5.3 Audit Readiness Is Not Audit Preparation

There is a real difference between an entity that keeps audit-ready evidence continuously and one that runs an intensive preparation effort in the weeks after an audit notification. Both produce evidence. Only one produces evidence that survives close scrutiny.

Evidence assembled under time pressure tends to have tells: inconsistent naming across documents assembled by different people, dates that cluster close to the audit notification, documentation that addresses the requirement but not the actual practice. Experienced auditors recognize these patterns. They are not violations in themselves, but they prompt a closer look at the underlying controls.

The entity that keeps consistent, contemporaneous evidence presents a different posture altogether. Patch evaluations are completed inside the 35-day cycle. Access reviews happen on schedule and are documented when they happen. Baselines are updated promptly after authorized changes. That entity is not reconstructing its compliance history under pressure.

5.4 Chapter Summary

Evidence strategy in a high impact environment means keeping a continuous, traceable record of how controls are implemented and verified, consistent and corroborated across the control areas that interact. Entities that do this give auditors far less to find. Entities that don't get audits that show what their program actually looks like from the outside.

> **From the Field**
>
> A program that's clean in the sampled portion and dirty everywhere else gets caught eventually. Documentation that exists only in someone's head isn't documentation. It's a single point of evidence failure.

6 How Mature High Impact Programs Actually Operate

Across three decades of involvement in bulk electric system operations, compliance oversight, and regulatory proceedings, the pattern has been consistent: entities with mature high impact CIP programs share structural characteristics that set them apart from entities whose programs are mostly reactive. Those characteristics are worth describing precisely, because they are achievable targets, not abstract ideals.

6.1 Governance That Connects Operations and Compliance

In mature high impact programs, the people responsible for CIP compliance have direct, working relationships with the engineers and operators who run the systems the program covers. This is not proximity on an org chart. It is operational integration. Compliance staff understand the systems well enough to see when a proposed operational change has CIP implications before it is made. Operations staff understand CIP well enough to escalate a potential compliance issue instead of solving it operationally in a way that leaves a documentation gap.

That integration does not happen by itself. It comes from deliberate organizational design: regular joint reviews, shared documentation tools, and a norm that treats CIP compliance as part of how operations are run rather than an outside reporting obligation. In my view, programs built this way tend to catch their own problems before an auditor does.

6.2 Continuous Compliance vs. Event-Driven Panic

The defining feature of a mature high impact program is that it does not change much when an audit notification arrives. The documentation is current because it is maintained continuously. The evidence is consistent because it is produced at the time. The controls are verified because verification is an ongoing operational activity, not a pre-audit sprint.

Getting there takes investment in process infrastructure: systems that track the 35-day patch evaluation cycle and escalate approaching deadlines, access review workflows that produce real decisions instead of spreadsheet entries, change management that updates baseline documentation as authorized changes are completed. None of these are CIP requirements. The standards do not prescribe how an entity builds its program. They are, in my view, the structural features that separate programs that sustain high impact compliance over time from programs that achieve it once and then degrade.

6.3 The Role of the Compliance Consultant at High Impact

Outside expertise has a legitimate role in high impact CIP compliance, especially where deep standards knowledge, audit experience, and architectural assessment are hard to keep in

house. Program gap assessments, mock audits, evidence review, and enforcement response support help entities no matter how mature their internal programs are.

What outside expertise cannot replace is internal ownership. When a consultant writes the documentation and nobody inside owns it, the entity's people often cannot explain, defend, or maintain it. When an auditor asks a subject matter expert to walk through the logic of a control and the SME defers to a document someone else wrote, the audit posture is compromised on the spot. The best use of outside support is to build internal capability, not to substitute for it.

6.4 Chapter Summary

What sets mature high impact CIP programs apart is how deeply compliance is built into operational culture. They treat CIP as an engineering discipline with governance implications, not a reporting obligation with engineering implications. That one distinction drives every structural difference between programs that hold up under audit scrutiny and programs that show their gaps the moment scrutiny arrives.

Glossary

Definitions marked with a NERC Glossary reference are quoted or closely paraphrased from the definitions currently subject to enforcement. Several of them change on July 1, 2028.

Term	Meaning in this paper
BES Cyber Asset	A Cyber Asset that, if rendered unavailable, degraded, or misused, would within 15 minutes of its required operation, misoperation, or non-operation adversely impact one or more Facilities, systems, or equipment which, if destroyed, degraded, or otherwise rendered unavailable when needed, would affect the reliable operation of the Bulk Electric System. Redundancy is not considered. (NERC Glossary)
BES Cyber System	One or more BES Cyber Assets logically grouped by a responsible entity to perform one or more reliability tasks for a functional entity. (NERC Glossary)
BES Cyber System Information (BCSI)	Information about the BES Cyber System that could be used to gain unauthorized access or pose a security threat to it, such as network topology or collections of network addresses. Individual pieces of information that pose no threat by themselves are excluded. (NERC Glossary)
Bright-line criteria	The impact rating criteria in Attachment 1 of CIP-002 that sort BES Cyber Systems into high, medium, and low impact. They are applied as written, not by entity judgment.
Control Center	One or more facilities hosting operating personnel that monitor and control the BES in real time to perform the reliability tasks, including their associated data centers, of a Reliability Coordinator, a Balancing Authority, a Transmission Operator for transmission Facilities at two or more locations, or a Generator Operator for generation Facilities at two or more locations. A revised definition takes effect July 1, 2028. (NERC Glossary)
Electronic Access Control or Monitoring Systems (EACMS)	Cyber Assets that perform electronic access control or electronic access monitoring of the ESP or BES Cyber Systems, including Intermediate Systems. (NERC Glossary)
Electronic Access Point (EAP)	A Cyber Asset interface on an ESP that allows routable communication between Cyber Assets outside the ESP and Cyber Assets inside it. (NERC Glossary)
Electronic Security Perimeter (ESP)	The logical border surrounding a network to which BES Cyber Systems are connected using a routable protocol. (NERC Glossary)
External Routable Connectivity (ERC)	The ability to access a BES Cyber System from a Cyber Asset outside its associated ESP through a bi-directional routable protocol connection. (NERC Glossary)

Term	Meaning in this paper
High impact BES Cyber System	A BES Cyber System used by and located at a Control Center or backup Control Center that meets criterion 1.1, 1.2, 1.3, or 1.4 of CIP-002 Attachment 1.
Interactive Remote Access (IRA)	User-initiated access by a person using a remote access client or other remote access technology over a routable protocol, originating from a Cyber Asset that is not an Intermediate System and is not inside any of the entity's ESPs or at a defined EAP. It does not include system-to-system process communications. (NERC Glossary)
Intermediate System	A Cyber Asset or collection of Cyber Assets performing access control to restrict Interactive Remote Access to only authorized users. It must not be located inside the ESP. (NERC Glossary)
Jump host	Common industry term for the system that brokers remote access into an ESP. When it controls Interactive Remote Access, it is an Intermediate System and therefore an EACMS.
Notice of Penalty (NOP)	The filing NERC submits to FERC and other applicable governmental authorities after an entity accepts a notice of alleged violation, a violation is affirmed, or a settlement is reached. Appendix 4C of the NERC Rules of Procedure, Section 5.9, governs it.
Physical Security Perimeter (PSP)	The physical border surrounding locations in which BES Cyber Assets, BES Cyber Systems, or EACMS reside, and for which access is controlled. (NERC Glossary)
Protected Cyber Assets (PCA)	Cyber Assets connected by a routable protocol within or on an ESP that are not part of the highest impact BES Cyber System in that ESP. Their impact rating equals the highest rated BES Cyber System in the same ESP. (NERC Glossary)
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 do not require specific evidence.
Subject Matter Expert (SME)	The person designated to speak with authority on a specific technical or operational topic during an audit interview.
Technical Feasibility Exception (TFE)	The formal exception process in Appendix 4D of the NERC Rules of Procedure, available only for CIP requirements that apply "where technically feasible" or that FERC has directed be subject to it.

How Energy Compliance Can Help

High impact programs live or die on architecture that matches the documentation and evidence that matches the architecture. Energy Compliance services related to Critical Infrastructure Protection include, but are not limited to:

- CIP applicability and scope analysis (CIP-002 through CIP-014)
- Cyber asset and BES Cyber System identification and impact classification
- Electronic Security Perimeter and Physical Security Perimeter boundary analysis
- CIP governance and program assessments
- Integration of cybersecurity oversight with broader reliability programs
- Audit and enforcement support for CIP findings (non-advocacy)
- CIP framework reviews, gap analyses, and improvement plans
- Training focused on CIP framework, requirements, and audit expectations
- Executive and board-level CIP awareness briefings

For high impact programs specifically, EC provides program assessment and gap analysis, ESP architecture review and boundary documentation evaluation, mock audit preparation and RSAW evidence review, SME preparation and audit coordination, and independent compliance monitoring program readiness evaluations. Around that core: audit defense with Notice of Penalty response and settlement posture, enforcement response support and mitigation planning, data request strategy, NERC compliance program support, interpretation, and audit preparation, CIP advisory on high impact architecture, evidence strategy, and enforcement response, senior advisory on complex reliability questions, and industry engagement in standards development and working-group participation.

WattOps Compliance is EC's NERC compliance evidence platform (wattopscompliance.com). It keeps the evidence behind your non-CIP obligations organized while your team carries the high impact load. EC will run one standard at one facility through WattOps at no cost (CIP standards excluded).

RenewOps is EC's renewable operations monitoring and analytics platform, with a Data Center module. Any operations-data path that leaves a high impact Control Center belongs on the ESP diagram and the EAP rule set before it goes live.

Pull your firewall rules and your ESP diagram and lay them side by side. If they don't match, call us.

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 exclusively 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. Our work spans NERC Reliability Standards, FERC orders, RTO/ISO market participation rules, Regional Entity oversight, and state regulatory frameworks.

We do this work differently than larger consulting firms. A senior practitioner with regulator-side experience oversees every engagement. We staff for outcomes, not billable hours. Our deliverables are written to be operationally executable and audit-defensible, structured for the auditor's question rather than the consultant's binder. Where automation can replace manual work, we build the automation. Where senior judgment is required, the senior is in the room.

Energy Compliance operates independently of regulatory and oversight bodies. Every engagement is focused on clarity, defensibility, and practical execution, not theoretical compliance. 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.

Learn more at energycomplianceinc.com.

Sources and Authorities

- North American Electric Reliability Corporation (NERC), Reliability Standard CIP-002-5.1a, Cyber Security: BES Cyber System Categorization, including Attachment 1, Impact Rating Criteria, and Guidelines and Technical Basis (effective December 27, 2016). https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-002-5.1a.pdf
- NERC, Reliability Standard CIP-004-7, Cyber Security: Personnel & Training (effective January 1, 2024). 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) (effective October 1, 2022). https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-005-7.pdf
- NERC, Reliability Standard CIP-006-6, Cyber Security: Physical Security of BES Cyber Systems (effective July 1, 2016). https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-006-6.pdf
- NERC, Reliability Standard CIP-007-6, Cyber Security: System Security Management (effective July 1, 2016). https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-007-6.pdf
- NERC, Reliability Standard CIP-009-6, Cyber Security: Recovery Plans for BES Cyber Systems (effective July 1, 2016). https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-009-6.pdf
- NERC, Reliability Standard CIP-010-4, Cyber Security: Configuration Change Management and Vulnerability Assessments (effective October 1, 2022). https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-010-4.pdf
- NERC, Reliability Standard CIP-011-3, Cyber Security: Information Protection (effective January 1, 2024). https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-011-3.pdf
- NERC, Reliability Standard CIP-012-2, Cyber Security: Communications between Control Centers (effective July 1, 2026). https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-012-2.pdf
- NERC, Reliability Standard CIP-013-2, Cyber Security: Supply Chain Risk Management (effective October 1, 2022). https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-013-2.pdf
- NERC, Reliability Standards CIP-015-1 and CIP-015-2, Cyber Security: Internal Network Security Monitoring (subject to future enforcement October 1, 2028 and October 1, 2029). https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-015-1.pdf ; https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-015-2.pdf

- NERC, Reliability Standards: CIP standards listing with status and effective dates. https://www.nerc.com/standards/reliability-standards/cip
- NERC, Project 2016-02 Modifications to CIP Standards (virtualization), Implementation Plan (April 3, 2024). https://www.nerc.com/globalassets/standards/reliability-standards/cip/2016-02-virtualization-implementation-plan_clean_04032024.pdf
- NERC, Glossary of Terms Used in NERC Reliability Standards. https://www.nerc.com/glossary-of-terms
- NERC, Rules of Procedure, Appendix 4C, Compliance Monitoring and Enforcement Program, Section 5.9, Notice of Penalty (effective May 19, 2022). https://www.nerc.com/globalassets/who-we-are/rules-of-procedure/appendix-4c.pdf
- NERC, Rules of Procedure, Appendix 4D, Procedure for Requesting and Receiving Technical Feasibility Exceptions to NERC Critical Infrastructure Protection Standards (effective July 1, 2016). https://www.nerc.com/globalassets/who-we-are/rules-of-procedure/appendix-4d.pdf
- 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
- NERC, Critical Infrastructure Protection Evidence Request Tool User Guide (January 2026). https://www.nerc.com/globalassets/programs/compliance/cmep-resources/cip-evidence-request-tool-user-guide-v10.pdf
- NERC, Reliability Standard Audit Worksheet (RSAW) Review and Revision Process (effective March 1, 2018). https://www.nerc.com/globalassets/programs/compliance/rsaws/nerc-reliability-standard-audit-worksheet-rsaw-review-and-revision-process.pdf
- Western Electricity Coordinating Council (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
- Federal Energy Regulatory Commission (FERC), North American Electric Reliability Corporation; Order Approving Reliability Standard CIP-002-8, Docket No. RD25-8-000 (March 19, 2026). https://www.govinfo.gov/content/pkg/FR-2026-03-24/pdf/2026-05715.pdf
- FERC, Order No. 907, Critical Infrastructure Protection Reliability Standard CIP-015-1, Cyber Security: Internal Network Security Monitoring, Docket No. RM24-7-000 (issued June 26, 2025; 90 FR 28889, July 2, 2025). https://www.govinfo.gov/content/pkg/FR-2025-07-02/html/2025-12309.htm

© Energy Compliance, Inc. · EC-WP-203 · October 2026