MONITORING TRUTH & FLEET RECONCILIATION SERIES · EC-WP-1106

Displayed Is Not Citable

Compliance evidence capture: provenance, the tautology trap, and the obligation nobody could mark done.

Twenty-five sites; on August 25, 2026, eight had a citable voltage schedule. Three sites scoring exactly 100% in band against a threshold derived from themselves. Sixty-five recurring intervals, none able to compute a next-due date. An analysis of the class of failure that makes correct numbers uncitable, and the provenance discipline that answers it.

The Evidence Layer

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, engineering advice, or formal interpretation of any NERC Reliability Standard, FERC order, Department of Energy directive, RTO or ISO protocol, or state commission rule.

This publication may not be used or cited as authority in any audit, enforcement, regulatory, or legal proceeding.

On regulatory citations

The body of this reference carries no requirement text, no requirement number and no effective date, and it quotes no order, protocol section, or rule. That is deliberate. The subject of the series is how a monitoring platform represents measured reality, and that question comes before, and stands apart from, which requirement is being satisfied. Standard families are named only where the connection to a specific finding is material, and they are named as families. The handful of family-level statements the text does make were checked on October 17, 2026 against the current enforceable standards and the NERC Rules of Procedure, which are listed in Sources and Authorities at the back. Nothing here should be read as a statement of what any standard requires. Requirement-level citation belongs in a deliverable prepared against the current enforceable source for a specific registered entity and audit period, with the version and effective date verified at the time of writing. A reader who needs operative language must obtain it there.

On the findings

Every defect described here was measured against a live production monitoring platform. The observations are anonymized: no client, site, site code, counterparty, vendor, product, asset tag, or person is identified, and figures are reported at a level of aggregation that does not permit re-

identification. Some of the defects described are in RenewOps, the platform Energy Compliance builds. Those are marked as ours and written up on the same terms as the rest.

On currency

A measurement is a statement about one system at one moment. The platforms described are under active remediation. A figure that was true on the date given may not describe the same system today, and it does not describe any other system at all. Where a matter was open, unresolved, or unverified at the time of writing, it is identified as such. Energy Compliance, Inc. does not represent that any condition described here persists, or that any condition absent here does not exist elsewhere.

Non-affiliation

NERC, FERC, the United States Department of Energy, the Public Utility Commission of Texas, ERCOT, PJM, 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, the Department of Energy, the Public Utility Commission of Texas, ERCOT, PJM Interconnection, L.L.C., or any NERC Regional Entity. The views expressed in this publication do not represent the views of NERC, FERC, or any Regional Entity. RenewOps is a product of Energy Compliance, Inc.

Energy Compliance, Inc. assumes no responsibility or liability for errors, omissions, or outcomes resulting from the use of this publication. Compliance obligations depend on facts and circumstances unique to each registered entity, and the reader should consult qualified counsel and subject-matter experts when making compliance decisions.

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

In a control room the screen is not a report on the system. For the eight hours you sit in front of it, the screen *is* the system. You act on it. You hand over on it. You write it into the log, the log becomes the record, and the record becomes the evidence. Everything downstream inherits whatever the screen got wrong, and nothing downstream is equipped to notice.

That is the subject of this series. Not whether a fleet is performing. Whether the number on the screen is a measurement of the fleet at all.

The series came out of an assessment. In 2026 we pointed a defect battery at a live production monitoring stack running a multi-gigawatt utility-scale PV and storage fleet and wrote up everything it found: twenty-plus verified defects, each with a source and a timestamp. What I did not expect was the shape of the result. The defects were not a list of unrelated bugs. They were six recurring structures, and every one of them was a

system reporting an absence as a value. A missing irradiance reading rendered as zero. A denominator that shrank as equipment failed, so a half-dead site scored a perfect hundred. An ingest timestamp wearing an event timestamp's label. A reading four hundred days old displayed as current. A revenue meter an order of magnitude above nameplate rendered live, because nothing in the system had a nameplate to check it against. A voltage band seeded from a plant's own history, indistinguishable on screen from a schedule issued by the entity of record.

None of the third-party findings was a product defect in the usual sense, and I want to be careful here. The easy version of this series is a hit piece on monitoring platforms, and that version would be wrong. A monitoring platform displays what it is given. It was built to display. It was never asked to reconcile what it displays against what a contract requires, what a standard demands, or what physics permits. Nobody built that layer. What is missing is a category that does not yet exist in most fleets, not a failure of any product.

Two of the defects in the original assessment were in our own platform. We wrote those up the same way and shipped the report with them in it, and they appear in these references marked as ours. I mention it because it is the only credential in this series that is hard to fake. Any vendor can show you a findings list from someone else's product. A vendor who has never found a fault in their own has not looked.

These references carry no requirement-level citations in the body, and that is deliberate; the reasoning is in the Disclaimer. They are also not marketing material dressed as content. They are written in the same voice I use in a control room or in front of a Regional Entity audit team: direct, grounded in evidence, and honest about what a measurement can and cannot tell you.

Displayed is not citable. If one sentence from this series survives, I would like it to be that one. A number can be correct, current, well formatted, and produced by a system everybody trusts, and still be unusable the moment an auditor asks where it came from. On August 25, 2026, eight of twenty-five sites classified in RenewOps had a citable voltage schedule. The rest were not shown to be wrong. They could not be traced to a complete schedule from the entity of record for the point being measured, and that is a different kind of failure. It is the kind you cannot repair after the fact.

Energy Compliance, Inc. exists because most of the consulting offered to registered entities is structured for billable hours rather than for outcomes. A senior practitioner oversees every engagement, and we do not bring five people to a meeting that needs one. We automate the work that should be automated and apply senior judgment to the work that requires it. If that matches what you are looking for, the back of this reference tells you how to reach us.

If not, the reference still belongs to you. Take what is useful. Apply it well. And keep asking the only question that matters about a number on a screen: what would have to be true for this to be wrong, and would I be able to tell?

Rob Smith

Founder and Principal, Energy Compliance, Inc.

Bloomington, Minnesota

How to Use This Reference

This reference is one of eight in the Energy Compliance Monitoring Truth and Fleet Reconciliation Series. The series is organized layer by layer. Each reference stands on its own and assumes none of the others, at the cost of a small amount of repeated framing where two layers meet.

Reference	Title	What it covers
EC-WP-1100 · The Measurement Layer	Missing Is Not Zero	The six-class failure taxonomy and the five design rules that answer it. The foundational reference; every paper that follows takes one class apart on a live system.
EC-WP-1101 · The Alarm Layer	The Board That Deleted Its Own Alarms	Alarm management treated as an information system: load, priority, latching, and the monitor that stops. Seven failure modes, each with the arithmetic that exposes it.
EC-WP-1102 · The Operator Layer	What the Incoming Operator Cannot See	Situational awareness as a testable property of the record rather than a quality of a dashboard. Turnover, attribution, and identity, measured five ways.
EC-WP-1103 · The Rules Layer	A Rules Engine That Measures Its Own Usefulness	The not-needed rate: how a rules engine earns the right to keep a rule. The measurement that breaks the ratchet of ever-growing alarm catalogs.
EC-WP-1104 · The Procedure Layer	A Machine May Draft It	Switching order integrity when a machine prepares the draft. Five integrity requirements, the failure behind each, and the one field a drafter must never write.
EC-WP-1105 · The Contract Layer	Six Ways to Compute 99%	Availability constructs reconciled against executed agreements. Computability as a per-site position (ready, partial, blocked, or undefined) with the blocker named.
EC-WP-1106 · The Evidence Layer	Displayed Is Not Citable (this reference)	Provenance as the property that separates a number you can show from a number you can cite. The tautology trap, the seeded schedule, and the obligation nobody could mark done.
EC-WP-1107 · The Platform	The Reconciliation Layer	RenewOps itself: what the platform does, what it deliberately refuses to produce, and what it does not claim.

How to read a figure in this series

Read every number with its source, and with its measurement date where the text gives one. This series does not make claims about systems it has not measured. Where a value could not be verified, it is marked as deferred with the specific question that would resolve it, rather than softened into an estimate.

What these tests cannot tell you

Every diagnostic in this series is a test of the *record*, not of the plant. A clean result means the platform's representation of the fleet is internally consistent, sourced, and current. It does not mean the equipment is healthy, the revenue is correct, or the program is compliant. The reverse also holds: a failed test is evidence that a number cannot be defended, not evidence that the underlying asset is in trouble. Several of the worst findings in the original assessment sat on top of plants that were running perfectly well. Conflating the two is the most common misreading of this work, and it produces false alarm and false comfort in equal measure.

Specific to this reference: a provenance review sorts your signals into citable and not. It does not tell you whether a citable signal shows compliance. A fully traceable record of an exceedance is exactly as citable as a record of conformance, and rather more likely to be cited. Provenance is the precondition for a defensible position, not the position itself.

The findings in this reference are drawn from an assessment of one fleet. They are offered as a taxonomy to test your own fleet against, not as a statement about it. The restriction on requirement-level citation, explained in the Disclaimer, is not caution for its own sake. It is the same discipline this paper argues for.

Companion to EC-WP-1100 Missing Is Not Zero, whose provenance-erasure class this paper develops in full.

Executive Summary

Displayed is not citable

Your monitoring platform is not an evidence system. At audit, that is the whole conversation.

A setpoint band inferred from your own plant history is not a schedule your operating entity issued, yet on most screens the two look identical. If the screen does not say which is which, you cannot safely cite either.

This failure separates a data problem from an evidence problem. Other failures in this series make numbers wrong, and a wrong number can be recomputed. This one makes **correct numbers uncitable**, and recomputation cannot fix it, because what is missing was never captured. **The underlying operation may have been perfect. That is what makes this class expensive: you did the right thing and cannot demonstrate it.**

What the measurements found

Observed	Measure
Sites with a citable voltage schedule, August 25, 2026	8 of 25
Sites scoring exactly 100.0% in band, tautologically	3
Sites published in a monthly report despite being explicitly disabled	2
Regulation service-status field populated, fleet-wide	0 rows
Compliance obligations loaded	209
Applicability gates unanswered	33, gating 70 obligations
Recurring intervals with a last-performed date	0 of 65
Controlled documents carrying a template document number	99 of 121

The three tautological sites were scored against a band derived from the median of their own series. A measurement compared against a threshold computed from itself always passes, so the result is an identity rather than a compliance result.

Chapters 1 to 4 cover provenance; 5 to 7, the obligation calendar, reporting readiness and controlled documents; 8, ten questions; 9, our own platform, RenewOps, including the bug we shipped.

1 Data, evidence, and the gap between them

Operational data becomes evidence when it can answer four questions about itself. Most operational data can answer none of them.

Question	Requires	Usually stored?
What is this value?	The number.	Yes
Where did it come from?	Source, authority, date obtained.	No
What kind of value is it?	Measured, modeled, inferred or issued.	No
Was it in force at the time?	Effective and supersession dates.	Rarely

Only the first is normally stored, which is entirely reasonable for a system built to operate a plant. It becomes a problem the moment that system is relied on as the substrate for a compliance demonstration. That happens by default, and silently, because the data is there and nothing else is.

The fourth question is not academic. Under the NERC compliance program, an entity's data must show compliance for the entire period a Compliance Audit covers. A value with no effective date cannot be placed in that period.

1.1 The four statuses

Status	Means	Evidentiary weight
Issued	Provided by the external entity of record, with source and date captured.	Citable without qualification
Measured	An instrument read it, and the instrument's health is known.	Citable with instrument provenance
Modeled	Computed from other quantities by a stated method.	Citable only with the method and its inputs
Inferred	Derived from the asset's own history because nothing better was available.	**Not evidence.** Operationally useful.

The whole of this chapter reduces to one observation: **a database column has one type, and these four have four different meanings.** Storing them together without a companion status

field destroys the distinction for good, because a seeded band and an issued schedule are the same datatype holding a similar number.

The rule. Status is stored alongside value, mandatory on write, and rendered alongside value, always. A write path that accepts a compliance-relevant value without a source is the defect. The missing source on the row that results is only the symptom.

1.2 Why seeding is not the villain

When you stand up a platform across a fleet, some sites will not have an issued schedule available on day one. For a generator, the voltage schedule of record comes from its Transmission Operator under the VAR family of standards, and that document does not always arrive before go-live. Seeding a plausible band from historical median is the correct engineering call. It makes the screen operable, and as long as the value is labeled as seeded, it harms nothing.

The defect is never the seeding. It is that the seeded value was not marked as seeded. Eighteen months later, when someone needs to demonstrate operation within an issued schedule, there is no way to separate the sites where that can be demonstrated from the sites where it cannot, short of a document-by-document reconstruction.

2 The tautology: scoring against a band derived from yourself

Three sites read exactly 100.0% in band. Not approximately. Exactly. The band had been computed from the median of the series being scored against it.

2.1 The mechanism

Where no issued schedule existed, the platform derived a band from the site's own historical voltage, which is a reasonable operational default. That band then became the threshold against which the same site's voltage was scored for in-band percentage.

A series compared against a band constructed from its own central tendency will fall inside that band almost always, and where the band is wide enough, always. The resulting figure is not a measurement of compliance. It measures nothing, and it renders in the same column, in the same format, as figures from sites that were being scored against a genuine issued schedule.

Calculation 1 · Detect self-referential scoring

```
-- A band derived from the scored series is a tautology.
-- 100.0% exactly, with a seeded source, is the signature.
SELECT s.site_id,
s.schedule_source, -- issued | seeded | absent
s.setpoint_kv, s.band_low_kv, s.band_high_kv,
v.pct_in_band,
v.samples,
(s.schedule_source <> 'issued') AS band_self_derived,
(v.pct_in_band = 100.0 AND s.schedule_source <> 'issued') AS tautological,
CASE WHEN s.schedule_source <> 'issued'
THEN 'INDICATIVE ONLY - not evidence of schedule compliance'
END AS caveat
FROM voltage_schedule s
JOIN voltage_in_band_monthly v USING (site_id, period)
ORDER BY tautological DESC, band_self_derived DESC;
```

2.2 The mirror-image failure

The same fleet produced the opposite error at another site: **37.6% out of schedule**. That is a dramatic-looking non-compliance figure, and it existed because the site was being scored against a collector-bus band while its actual schedule applied to a different bus at a materially different voltage class.

Both errors have the same root. The scored quantity and the band came from different places, and nothing recorded that they were supposed to correspond. One direction produces a reassuring number; the other produces an alarming one. Neither is a measurement.

Ours, and the correct handling. One site's schedule setpoint had itself been seeded from a corrupt meter. The value sat three orders of magnitude above the site's actual operating voltage and was still marked enabled. The plausibility layer **refuses to use** that setpoint and surfaces the row in a coverage view. It was deliberately not deleted. Schedules pass only through a function that demands a directive reference, so retiring it is a decision with a named owner rather than a cleanup.

2.3 The best compliance writing in the platform

Worth naming what worked. The schedule table carried a suspect_note field: plain-language text explaining, per row, why a given schedule was or was not trustworthy. Not a flag, not an enum. A sentence a compliance lead could read and act on.

That field did more for the defensibility of the fleet's voltage evidence than any calculation in the platform, and it cost nothing but the discipline of writing down what was already known.

3 Qualifiers that do not survive the export

The platform knew which schedules were seeded, which meters were suspect, and which sites were explicitly disabled. None of that reached the monthly report.

3.1 The finding

Three compliance-facing endpoints (a monthly evidence extract, a statistics view, and the monthly report itself) returned **no source column, no suspect flag, and no explanatory note.**

Worse, the monthly report **ignored the enabled flag entirely**, publishing high-nineties in-band percentages for two sites that had been disabled because they were measured on the wrong bus. Neither figure meant anything, and neither said so.

Where the qualifier existed	Where it was lost
Schedule source (issued vs seeded)	Dropped from all three compliance endpoints
Meter suspect flag	Dropped; the underlying trust view was never joined
Schedule enabled flag	Ignored by the report's own filter
Plain-language suspect note	Present in the table, absent from every export

A qualifier that does not travel with the number is not a qualifier. It is a private opinion held by a database.

3.2 Why this happens, structurally

Compliance exports are usually built after the operational views, by someone picking the columns that look relevant. Status columns are not the number, so they are not picked. The export is then reviewed for whether the numbers are right. They are.

The design rule: an evidence export carries the qualifier or it does not carry the value. Implement it as a constraint, not a convention: the export selects value and status *as a pair*, and a value with no status renders as withheld, with the reason. Then **generate the export and read it as an auditor would.** Every question the export cannot answer about a figure is a column that should have been in it.

One view in the same platform got this right: schedule_is_compliance_grade on every excursion, plus a plain-English statement of the obligation, so a reader could tell at once whether it was citable evidence or an operational observation. The difference was not sophistication. Somebody had asked what the output was *for*.

4 The evidence source that does not exist

At every site carrying automatic voltage regulation, the service-status field was null across every row. The display read a dash on every card in the fleet.

4.1 Why this is the most complete form of the failure

The other findings in this paper concern values whose status is unclear. This one concerns a value that does not exist. The VAR family puts a notification clock on a change in a generator's voltage-regulator status, and that clock had nothing behind it anywhere in the platform.

The correct action is not to find a substitute. It is to state the position plainly in the compliance narrative, identify the source that would supply the value (a plant control system point, a manual log, an operating record), and then go and get it. Discovering this during an audit is the expensive path to the same conclusion.

4.2 Three empties, three meanings

Rendering	Should mean	Usually reads as
Blank cell	Nothing. Ambiguous by construction.	*Fine*, which is the danger
Dash	We do not have this value right now.	Not yet loaded
No evidence source	No mechanism exists to obtain this value at all.	Correctly, as a gap

The three are operationally different and must render differently. A field that is temporarily unavailable is a monitoring issue. A field for which *no source exists* is a program gap that needs a decision, a budget and an owner, and it will never get them while it renders as a dash among other dashes.

Calculation 2 · Find the fields with no source

```
-- A 100% null rate on a compliance-relevant field is not sparse data.
-- It is an evidence source that does not exist.
SELECT 'regulation_state.service_status' AS field,
COUNT(*) AS rows_total,
COUNT(service_status) AS populated,
COUNT(*) - COUNT(service_status) AS null_rows,
ROUND(100.0*(COUNT(*)-COUNT(service_status))/NULLIF(COUNT(*),0),1) AS pct_null,
CASE WHEN COUNT(service_status) = 0
THEN 'NO EVIDENCE SOURCE - program gap, needs an owner'
WHEN COUNT(service_status) < COUNT(*) * 0.5
THEN 'sparse - monitoring gap'
ELSE 'populated' END AS verdict
FROM regulation_state;
```

Run that shape against every compliance-relevant field you hold. The fields returning *no evidence source* are your program gaps, and the list is almost always shorter than people fear and more specific than they expect.

4.3 Acknowledgment deadlines, and evidence of timeliness

Five regulation state changes on the same fleet were acknowledged after their deadline. That finding is only possible because three things were captured: when the state changed, when it was acknowledged, and what the applicable window was.

Most platforms capture the first two and not the third. Lateness is then computable only by someone who already knows the requirement, which means it gets computed during an audit rather than before one. **The window belongs in the data**, expressed in plain language alongside the obligation, so that timeliness is a query rather than an expert judgment.

5 The obligation calendar, and the field nothing could write

209 obligations loaded. 65 recurring intervals. **Zero with a last-performed date, so not one next-due date was computable.** The field existed. Nothing in the application could write to it.

5.1 What was loaded

Component	Count
Obligations	209
Of which, hard dates	86
Of which, recurring intervals	65
Of which, event-triggered clocks	58
Applicability gates, all unanswered	33
Hard dates gated behind those unanswered gates	70 of 86
Intervals with a last-performed date	0 of 65
Open deferrals recorded	37
Primary sources cataloged	64

5.2 Three structural observations

Applicability is upstream of everything. Thirty-three unanswered gates suppressed seventy of eighty-six dated obligations. Until an entity establishes which requirements apply to which registered functions and assets, a calendar cannot tell it what is due. Applicability analysis is therefore the first compliance task, and every later task depends on it.

An interval with no baseline cannot compute. A next-due date is last-performed plus interval. With last-performed empty, the correct output is *cannot compute: last performed not recorded*, which is what the platform correctly displayed. A default would have been worse. A next-due date computed from a blank is arbitrary, and it will be believed.

Deferrals are first-class. Thirty-seven open deferrals, recorded as data rather than living in email, is a far stronger position than thirty-seven items quietly omitted. A deferral with a reason and an owner is a managed position. The same item undocumented is a gap.

5.3 Completion as a record, not a checkbox

Calculation 3 · What a completion must capture, and what it must refuse

```
-- hard date -> closes THAT occurrence only. Next month is its own row.
-- interval -> sets LAST PERFORMED, which makes next-due computable.
CREATE FUNCTION compliance_complete_obligation(
p_obligation_id text,
p_performed_on date,
p_note text
) RETURNS TABLE (out_obligation_id text, out_next_due date)
LANGUAGE plpgsql AS $$
DECLARE v_actor text := current_operator(); -- session, never a parameter
BEGIN
IF v_actor IS NULL THEN
RAISE EXCEPTION 'unattributed completion refused';
END IF;
IF p_performed_on > current_date THEN
RAISE EXCEPTION 'completion date in the future refused';
END IF;
-- ... writes completed_at, completed_by, performed_on, note
-- ... and, for intervals, last_performed_at
END $$;
```

Four properties make that a record rather than a checkbox. It captures **who** from the session rather than from a parameter. It captures **when performed** separately from when recorded. It **refuses an unattributed call**. And it **refuses a future date**. Reopening requires a written reason that stays on the record permanently. Nothing is quietly un-completed.

That function is ours, and its first version shipped broken. The bug, and what it taught us about testing controls, is in Chapter 9.

6 Reporting readiness, field by field

Periodic operational reporting fails in a specific way. Not all at once, but three fields at a time, discovered at submission.

6.1 The five states

The remedy is a standing, field-level readiness position rather than an annual scramble. Each field in the reporting set sits in one of five states, and anything not live carries the specific route to resolution.

State	Means	Resolved by
Live	Populated from a trusted source, current.	Nothing further
Partial	Populated for some assets or periods.	Extend coverage; state the gap meanwhile.
Blocked	Source identified, not yet flowing.	An integration task with an owner.
Defect	Populated but known wrong.	A fix. Until then the field is withheld, not submitted.
Needs input	No system holds it.	A registry lookup, an engineering document, or a business decision.

The fifth state carries the most value, because it routes work correctly. A field that needs a registry lookup, a field that needs an engineering document, and a field that needs somebody to make a decision are three different tasks for three different people. Lump them together as "missing data" and all three go to the wrong person.

6.2 Withhold, do not caveat

The discipline that makes this work. A field in *defect* is **withheld from submission**, not submitted with a caveat. A known-wrong value in a regulatory submission is a different act from an omission with a stated reason, and the difference is not one of degree.

7 Controlled documents as evidence

A procedure library is an evidence source. That imposes four requirements most document stores do not meet.

Requirement	Because
Audited write path	Who ingested what, and when, checked against a roster on every call. A library with an unaudited write path is not an evidence source.
Explicit supersession	A replacing document names what it retires, and the retired version stays queryable. An audit covers a whole period, so you must be able to show what was in force on a past date.
Effective dating	A future-dated library exists and is not yet in force. Answering an operator today from a procedure effective next month is wrong in a way that looks right.
Real document numbers	An unnumbered controlled document cannot be cited unambiguously, which defeats the reason it is controlled.

7.1 Two findings from one load

Ninety-nine of 121 documents still carried a template document number rather than a real one. Two had genuine numbers. That was raised ahead of the library's effective date rather than discovered afterward, which is the entire value of looking.

The library was future-dated. Every document carried an effective date ahead of the load. The library existed and was not yet in force, and the retrieval layer must carry that distinction into every answer it gives.

7.2 The classification decision

None of the ingested documents was a reliability standard. They were the entity's own procedures, which *implement* standards. An earlier load let automatic detection read internal document numbers as standard numbers and filed eight procedures as standards.

Had that stood, a question about what a standard requires would have been answered with the entity's own procedure. That is provenance erasure at its most consequential: **an entity citing itself as the source of a requirement.** A procedure describes what an entity chose to do. A standard describes what is required. Confusing the two in front of an audit team is fatal to the position.

The rule. Document type is declared at ingest, never inferred from the document's contents. Any linkage to a standard is a separate field. Automatic detection may *propose* a classification. It may never *assign* one.

8 Ten questions for your own evidence position

8.1 The questions

#	Question	What a bad answer means
1	Does a provenance column exist on your compliance-relevant tables?	If not, every value is of unknown status until reconstructed from documents.
2	What proportion of your setpoints and schedules were issued by the entity of record?	The rest are operationally useful and are not evidence.
3	Is any compliance percentage scored against a band derived from the same series?	Exactly 100.0% with a seeded source is a tautology, not a result.
4	Do your compliance exports carry source, suspect and enabled flags?	A qualifier that does not travel with the number is a private opinion.
5	Does any report ignore a disabled flag?	You are publishing figures somebody deliberately marked untrustworthy.
6	Which compliance-relevant fields are 100% null?	Those are program gaps needing an owner, not monitoring gaps.
7	How many applicability gates are unanswered, and what do they gate?	Until they are answered your calendar cannot tell you what is due.
8	How many recurring intervals have a last-performed date?	Without one, next-due is not computable and any date shown is arbitrary.
9	Can completion be recorded, attributed, and reopened only with a written reason?	Otherwise the calendar is a list, and completions are unattributable.
10	For each control in your compliance stack, has the success case been tested, or only the refusals?	A function that refuses everything passes every negative test ever written.

9 What RenewOps does about this

The reconciliation layer. It reads your monitoring stack and checks it against ratings, agreements and the documents of record. It declines to publish what it cannot defend.

Provenance is the difference between a number you can show and a number you can cite. This chapter is the provenance discipline as RenewOps implements it, including which of the twenty-five sites earned a citable figure as of August 25, 2026, and why the rest did not.

9.1 Every voltage schedule carries a provenance class

On August 25, 2026, all twenty-five sites were classified into one of six states:

Class	What it means for a citation
Complete document from the operating entity	Target and tolerance band both from the entity of record. **Citable.**
Target only, no band	The entity gave a target and no tolerance. An in-band percentage requires a band, so none is published.
Wrong bus	A schedule exists and describes a different bus than the one being measured.
Not in the source of record	No schedule for this site appears in the document set at all.
Source unusable	A document exists and its extract is garbled. The defect is recorded, not repaired.
No row	The site has no schedule row of any kind. Two live sites had been hiding in this state, because a per-row grade cannot grade a row that does not exist.

As of August 25, 2026, eight of twenty-five sites received a citable in-band percentage. All eight come from a complete documented schedule. One site that had been publishing 62.9 percent against a placeholder band on the wrong bus now publishes nothing at all and reads WRONG BUS instead.

9.2 Seeded is labeled, permanently

A band auto-detected from a site's own historical median is not a schedule, and RenewOps labels it **INDICATIVE ONLY** wherever it appears. The August 25, 2026 classification left eight of twenty-five sites with a complete documented schedule. Monthly evidence output now requires complete documented provenance before it will produce a figure. A band the operating entity never issued is not evidence of conformance to a band the operating entity issued.

What that fixed. Before the classification, compliance exports carried no provenance flags at all. They published in-band percentages for wrong-bus sites and for sites explicitly disabled, including **tautological readings of exactly 100.0 percent**.

9.3 Citations are recorded as citations

Each schedule row quotes its source inline: the file, the path, the server-modified date, and the document key. The source document's own defects are recorded on the row rather than repaired in it. A garbled extract stays garbled and says so, because silently cleaning a source document destroys the only evidence that it needed cleaning.

9.4 The write door

The schedule of record may only be set through **one controlled path**. That path demands a directive reference, refuses an implausible setpoint, writes history, and stamps the source. The schedule is never written directly.

Two honest notes on that control. First, as of August 25, 2026 **it had no caller**. Schedules arrived through document load, so the path is correct and unexercised. Second, the corrupt seeded setpoint described in Section 2.2 stays in place and out of the banding, for the reasons given there.

The evidence it declined to create. Two live sites have complete documented schedules and deliver no voltage telemetry. No schedule row was written for either. A band against a dead channel yields either silence or a 100 percent excursion the moment a tag appears. Both of those are fabricated evidence, one reassuring and one alarming.

9.5 A dead field inside a live row

Per-field liveness is recorded with the age of each field and its count of unanswered polls, because a row reads live when *any* one of its fields is fresh. Behind that, fields dead between 41 and 293 days sat concealed behind between 289 and 413 unanswered polls. The schedule-versus-controller comparison now returns **CANNOT COMPARE** on a dead point rather than agreement. It had been producing a pass on evidence that does not exist, which is the worst of the three possible outcomes.

9.6 The obligation nobody could mark done

Recurring intervals compute a next-due date **only from a recorded last-performed date**. With no baseline recorded, all sixty-five regulatory intervals read *cannot compute: last performed not recorded*. The platform does not substitute a default, a calendar-year start, or an overdue flag. Recording the date once makes the interval computable, and the rules engine itself raises the prompt to record it.

Marking an obligation complete behaves differently by type. A hard date closes that occurrence only and leaves the next untouched. A recurring obligation records LAST PERFORMED, the field that unblocks the next-due computation. Both carry the refusals and the reopen rule set out with Calculation 3 in Section 5.3. A completed view shows who recorded what, when, how many days early or late, and every reopen with its reason.

The bug that shipped, and the lesson in it

> **Ours**
>
> The first version of the completion function in Calculation 3 shipped broken. Declaring RETURNS TABLE (obligation_id text, ...) creates an output variable of that name, which collides with the table column of the same name. PostgreSQL raised error 42702, as observed in the platform's captured output: column reference "obligation_id" is ambiguous. It could refer to either a PL/pgSQL variable or a table column.
>
> **Both refusal paths passed in production.** Unattributed operator: refused correctly. Future date: refused correctly. Everything tested was green.
>
> **The one thing not tested was the success case.** It was caught only by completing a real obligation.

The correction renamed the output columns and schema-qualified every column reference, and the self-test now exercises the success path against a real row and reverts it. That reversion matters. The test rows were cleared afterward, leaving the performed list deliberately empty. **A fabricated performance record on a compliance calendar is worse than no feature at all**, because the fabrication is indistinguishable from evidence and will be produced as evidence by somebody acting in good faith.

The lesson reaches well past this bug. **Testing that a control refuses correctly is only half a test.** A control that refused every call would pass all of those tests and still record nothing.

9.7 Notifications that never claim to have been sent

Notifications compose into an outbox and hold as **held-without-recipient** until an address is configured. They are composed, readable, and never reported as sent. A compliance system that reports a notification as delivered when no delivery occurred has manufactured the one fact the obligation turns on.

9.8 Can the evidence survive the platform?

On August 28, 2026 a restore was proven onto a clean, empty database from the exported files alone: 136 tables, 202 views, 95 functions, 18 materialized views, 255 indexes, 163 policies. Zero schema errors, zero function errors, and all 220 relations execute.

Two things must be said about that result. First, the export was **not self-sufficient and did not say so**. It granted privileges to roles it never created, so it failed on its first statement against an empty database. Anyone attempting a restore before that date would have concluded the export was broken. It was not broken. It was incomplete in a way that nothing recorded, which is the harder failure. Second, the materialized views must refresh in dependency order. Refreshed alphabetically, four of eighteen fail, which breaks ten ordinary views, and nothing on screen says the numbers are stale. The screens simply report nothing.

> **Deferred, and stated precisely**
>
> What was proven on August 28, 2026 is a schema and objects restore from a schema-only export. A data restore has not been demonstrated. This reference does not describe the result as a proven backup of the evidence, and neither should anybody citing it.

Provenance separates a record that survives being questioned from one that only survives not being questioned.

9.9 The gap that remains

Content is guarded well across every controlled write in the platform. **Identity is not.** Who set a schedule, approved an override, marked an obligation complete or answered a prompt is recorded from a string the caller supplies. Every attribution in the system is therefore a claim rather than a verified fact. An evidence layer that has solved provenance of *values* while leaving provenance of *actors* open has solved half the problem.

Figures measured August 24–28, 2026 on a live production instance. They describe that instance on those dates and are not a claim about any other fleet.

Glossary of Terms

Term	Meaning in this paper
Applicability gate	A yes-or-no question about an entity's registered functions or assets that must be answered before the platform knows whether a given obligation applies.
Automatic voltage regulator (AVR)	The generator control that holds terminal voltage or reactive output to a setpoint. Its in-service status is the field Chapter 4 found empty.
Band (tolerance band)	The upper and lower limits around a voltage setpoint within which operation is treated as on schedule.
Citable	Of a value: carrying enough provenance (source, status, and effective date) to be put in front of an auditor as evidence.
Compliance Audit	Under the NERC Rules of Procedure, a systematic, objective review and examination of records and activities to determine whether a Registered Entity meets the Requirements of applicable Reliability Standards.
Compliance grade	RenewOps' term for a voltage schedule whose target and band both come from a complete document issued by the entity of record.
Controlled document	A procedure or record managed under document control, with a unique number, an effective date and a supersession history.
Deferral	An obligation or field whose resolution is consciously postponed, recorded as data with a reason and an owner.
Entity of record	The external organization with authority to issue a value, for example the Transmission Operator that provides a generator's voltage schedule.
Generator Operator (GOP)	The NERC registered function that operates a generating facility. Under the VAR family it is the function held to the voltage or Reactive Power schedule its Transmission Operator provides.
Hard date	An obligation due on a fixed calendar date. Completing it closes that occurrence only.
Inferred value	A value derived from the asset's own history because nothing better was available. Operationally useful, not evidence.
Issued value	A value provided by the external entity of record, with its source and date captured.
Last-performed date	The date a recurring obligation was most recently carried out; the baseline from which its next-due date is computed.

Term	Meaning in this paper
Materialized view	A database view whose results are stored and must be refreshed, in dependency order, to stay current.
Measured value	A value read by an instrument whose health is known.
Modeled value	A value computed from other quantities by a stated method.
Provenance	The recorded origin of a value: who supplied it, of what kind it is, when it was obtained, and when it was in force.
Recurring interval	An obligation that repeats on a fixed cycle; its next-due date is the last-performed date plus the interval.
Seeded value	A starting value written into a platform from historical data so the screen can operate before an issued value is available.
Supersession	The act of one controlled document formally replacing another, with the retired version kept queryable.
Tautology (in scoring)	A compliance percentage computed against a threshold derived from the same series being scored. It always passes and measures nothing.
Transmission Operator (TOP)	The NERC registered function that operates the transmission system. Under the VAR family it is the function that provides a generator's voltage or Reactive Power schedule, and the one a Generator Operator notifies of a voltage-regulator status change.
Voltage schedule	The target voltage or Reactive Power, with any tolerance, that a generator is directed to maintain.

How Energy Compliance Can Help

When an auditor points at a number, can you show where it came from? Most programs find out during the audit. We would rather you found out now.

Services that fit this problem

Audit defense. If an audit is on the calendar, we review your evidence the way the audit team will, sort what is citable from what only looks citable, and close the gaps before anything goes out the door.

RSAW production. An RSAW narrative is only as strong as the evidence it points to. We write narratives tied to specific, sourced records, and we tell you plainly where the record will not carry the claim.

Compliance program management. Applicability decisions, obligation calendars with recorded completions, deferrals held as data with an owner, and controlled documents that are numbered, effective-dated and superseded properly.

The platforms

WattOps Compliance is our NERC compliance evidence platform (wattopscompliance.com). We will run one standard at one facility through WattOps at no cost, so you can see your own evidence position before you commit to anything. CIP standards are excluded from that offer.

RenewOps is our renewable operations monitoring and analytics platform, with a Data Center module. Chapter 9 describes how it handles provenance, what it refuses to publish, and what it still does not do.

Where to start

Evidence Provenance Review. Your compliance-relevant signals traced to source and sorted into citable and not. The seeded, the derived, and the self-referential are separated out from the issued.

Monitoring Truth Assessment. The full defect battery against your live monitoring stack. Every finding comes back with a source, a timestamp, and the question that resolves it, **whether or not you buy anything afterward.** What we cannot verify comes back marked deferred, not softened into a maybe.

Nobody buys a reconciliation layer before they know they need one, and the only honest way to find out is to look. Pick the one signal you would least like to defend in front of an audit team and send it to us. We will tell you what it would take to cite it.

energycomplianceinc.com · rsmith@complianceforenergy.com · 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 regulatory compliance and advisory firm serving the energy sector, with a focus on NERC, FERC, and RTO/ISO compliance and a particular concentration in the ERCOT and Texas PUCT markets. The firm helps registered entities and prospective registrants through the full compliance lifecycle: registration, program design, evidence development, RSAW production, mitigation, and audit defense.

The approach is research-analyst-first. Every conclusion is tied to an authoritative source, every narrative is backed by evidence, and every deliverable is built to survive regulator scrutiny. A senior practitioner with regulator-side experience oversees every engagement. We do not 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.

That discipline extends to compliance technology: WattOps Compliance, our NERC compliance evidence platform, and RenewOps, the renewable operations monitoring and analytics platform that serves as the reconciliation layer described throughout this series.

Auditing teaches one habit that never leaves: before you believe a number, ask what it would look like if it were wrong.

Website: energycomplianceinc.com · Email: rsmith@complianceforenergy.com · Phone: 763.438.4427 · Bloomington, Minnesota

Sources and Authorities

The findings here are Energy Compliance's own first-party, anonymized measurements of one live production system and of RenewOps, not publicly available. The sources below, accessed October 17, 2026, support the family-level statements in Chapters 1, 4 and 7, the Glossary, and Chapter 9.

NERC Reliability Standards

- NERC, Reliability Standard VAR-002-4.1, *Generator Operation for Maintaining Network Voltage Schedules*, "Mandatory Subject to Enforcement," effective September 26, 2017. R2 (maintain the Transmission Operator's voltage or Reactive Power schedule) supports Chapter 1; R3 (notify the Transmission Operator of an AVR status change within 30 minutes) supports Chapter 4. https://www.nerc.com/pa/Stand/Reliability%20Standards/VAR-002-4.1.pdf

NERC Rules of Procedure

- NERC, Rules of Procedure Appendix 4C, *Compliance Monitoring and Enforcement Program*, effective May 19, 2022, Section 4.1.3.2 (Period Covered), for Chapters 1 and 7. https://www.nerc.com/globalassets/who-we-are/rules-of-procedure/appendix-4c.pdf
- NERC, Rules of Procedure Appendix 2, *Definitions Used in the Rules of Procedure*, effective June 27, 2024, definition of "Compliance Audit," for the Glossary entry. https://www.nerc.com/globalassets/who-we-are/rules-of-procedure/appendix-2.pdf

Technical reference

- *PostgreSQL Documentation*, Appendix A, Error Codes (SQLSTATE 42702, ambiguous_column). https://www.postgresql.org/docs/current/errcodes-appendix.html
- *PostgreSQL Documentation*, PL/pgSQL, "Variable Substitution" and "Declarations". https://www.postgresql.org/docs/current/plpgsql-implementation.html ; https://www.postgresql.org/docs/current/plpgsql-declarations.html

© Energy Compliance, Inc. · EC-WP-1106 · Status as of October 17, 2026