MONITORING TRUTH & FLEET RECONCILIATION SERIES · EC-WP-1107

The Reconciliation Layer

RenewOps: what it does, what it refuses to do, and what it does not claim.

A capability reference for RenewOps, the fleet reconciliation platform Energy Compliance builds. What the product does, marked shipped, in remediation or deferred against a dated observation. The refusal catalog that defines it. And the boundary, item by item, of what it does not claim.

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 the absence of regulatory citations

The body of this series carries no requirement text, no requirement number, no standard version, and no effective date, and it cites no order, protocol section, or rule. The omission is deliberate. The subject of the series is how a monitoring platform represents measured reality. That question comes before, and stands apart from, which requirement is being satisfied. Regulatory standard families are named only where the connection to a specific finding is material, and they are named as families in the text. Where one is relevant, the Sources and Authorities page at the back identifies the version in effect on the review date so a reader can go to the source. Nothing here should be read as a statement of what any standard requires. A reader who needs operative language must obtain it from the current enforceable source for the relevant registered entity and audit period.

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. In this reference the platform measured is RenewOps, the platform Energy

Compliance builds, so the defects are ours. They are written up on the same terms as the rest of the series.

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 when measured 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. The product observations in this edition are those made between 24 and 28 August 2026 and were not re-measured for it. The external references were checked against their sources on October 17, 2026.

Non-affiliation

NERC, FERC, the United States Department of Energy, the Public Utility Commission of Texas, ERCOT, PJM, the International Society of Automation (ISA), the American National Standards Institute (ANSI), 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., ISA, ANSI, or any NERC 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. I have worked enforcement matters from inside the regulator's process. For the last several years I have advised registered entities directly through the firm I founded.

Here is the thing about a screen. 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, but 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 that was a vendor defect, 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. The gap is not a failure of any product; it is a category that does not yet exist in most fleets.

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.

The body of these references carries no regulatory citations, and that is deliberate; the reasoning is in the Disclaimer. They are also not marketing material dressed as content. I wrote them the way I talk 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.

This last reference is the product, and I have put it last on purpose. The seven that precede it establish that the problem is real and that it is structural. Read before that argument lands, a capability description is a brochure. Read in order, it answers a question the reader has already asked. It carries a refusal catalog and a chapter listing what RenewOps does not claim, and those two chapters are longer than most vendors' entire honesty budget.

Energy Compliance, Inc. exists because most of the consulting offered to registered entities is structured for billable hours rather than for outcomes. Every engagement is overseen by a senior practitioner. 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 and 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.

EC-WP-1100 · The Measurement Layer: Missing Is Not Zero
The six-class failure taxonomy and the five design rules that answer it. This is the foundational reference for the series; 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, and how a rules engine earns the right to keep a rule. This is the measurement that breaks the monotonic-growth ratchet in an alarm catalog.

EC-WP-1104 · The Procedure Layer: A Machine May Draft It
Switching order integrity when a machine prepares the draft. Five integrity requirements, each with the failure that motivates it, 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 is treated as a per-site position (ready, partial, blocked, or undefined) with the blocker named.

EC-WP-1106 · The Evidence Layer: Displayed Is Not Citable
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 (this reference)
RenewOps itself. What the platform does, what it deliberately refuses to produce, and what it does not claim. It is the capability reference for the layer the other seven show a need for.

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 both false alarm and false comfort.

Specific to this reference: everything here was observed on one live production instance between 24 and 28 August 2026 and is dated. A capability marked in remediation is not a capability. A platform under active development is a different platform six weeks later, in both directions. Nothing in this reference is a commitment, a specification, or a roadmap.

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.

About This Document

This is a capability reference for RenewOps, the fleet reconciliation platform Energy Compliance builds. It describes what the product does, what it deliberately refuses to do, and what it does not claim.

Every capability below was observed on a live production instance between 24 and 28 August 2026. Each is marked with what it is: **shipped** and verified on a stated date, **in remediation** where a defect is identified and the fix specified, or **deferred** where the question is open and the answer is not yet known. Nothing is described in the present tense because it is planned.

All observations are anonymized. No client, site, site code, counterparty, vendor, product other than RenewOps, asset tag or person is identified. Regulatory standard families are named only where the connection is material. No requirement text, requirement number, standard version or effective date appears in the body of this reference.

What this document is not. It is not a specification, a roadmap, or a commitment. A capability marked in remediation is not a capability. A figure carries the date it was measured, because a platform under active development is a different platform six weeks later, in both directions.

The seven diagnostic references in this series, EC-WP-1100 through EC-WP-1106, establish the failure classes this platform is built to answer. This reference assumes none of them, at the cost of restating each failure briefly where a capability answers it.

Executive Summary

The layer nobody built

The common arrangement on a renewable fleet is two systems that do not speak to each other. One is a monitoring platform, which displays what the plant is sending. The other is a filing cabinet of executed agreements and compliance obligations, which states what the plant is required to deliver and prove. Nothing sits between them.

RenewOps is that layer. It reconciles what the monitoring reports against what the contracts require and what the standards demand. The part that distinguishes it is this: it **refuses to produce a number it cannot defend**, visibly, with the reason on the face of the tile.

The three questions it answers

Question	What the layer returns
Is this number a measurement?	Provenance, age, denominator and plausibility on the face of every displayed value. Where any of the four cannot be established, a refusal with its reason.
Is this the number the contract asks for?	A per-site computability position (ready, partial, blocked or undefined) reconciled against the executed agreement, with the specific blocker named.
Would this survive being cited?	A provenance class per compliance-relevant signal, sorting issued from seeded from derived, and an obligation calendar that computes a next-due date only from a recorded last-performed date.

What it looks like in practice

On the fleet where it was measured in August 2026, the layer produced results that a monitoring platform structurally cannot produce, because a monitoring platform has never been asked to read a contract or check a value against a rating:

- **367 telemetry samples refused in twenty-four hours at one site**, with zero false positives elsewhere in the fleet. The readings ran around 13,700 MW and 9,280 kV and had been rendering as live on a plant where both are impossible.
- **Eleven of twenty-three sites returning blocked** on availability, most of them because coefficients the executed agreement left to be mutually agreed were never populated. Nobody can compute them, including the counterparty, and the screen says so.

- **Eight of twenty-five sites earning a citable voltage-band figure.** The other seventeen publish nothing rather than a percentage measured against a band derived from the plant's own history.
- **Sixty-five recurring obligations reading "cannot compute: last performed not recorded"** rather than a manufactured next-due date.
- **Twenty historical availability rows withheld**, blank, with the reason on the row, because the denominator was wrong and the percentage cannot be corrected.

Every one of those is a number the fleet used to have and now does not. That is the product. It is a strange thing to sell, and it is the only honest position available.

How this reference is organized

Chapter 1 sets out the category and why the layer is not a monitoring platform. Chapter 2 gives the five design rules the whole product is built to. Chapters 3 through 8 take the layers in turn: measurement, alarms, operators, rules and procedures, contracts, evidence. **Chapter 9 is the refusal catalog**, and if you read only one chapter, read that one. Chapter 10 states what RenewOps does not claim, item by item, including the controls that are specified and not yet enforced. Chapter 11 covers deployment and what an assessment produces.

1 Not a monitoring platform

The most common first reaction to RenewOps is that it sounds like a monitoring platform with better validation. It is worth being precise about why it is not, because the distinction determines what it can be held responsible for.

1.1 What a monitoring platform is built to do

A monitoring platform acquires signals, stores them, and displays them. It is built for fidelity to the source: what the plant sent is what the screen shows. That is the correct design goal, and the good ones achieve it. A monitoring platform that silently altered readings to make them look sensible would be a much worse product.

The consequence is structural rather than a defect. A platform built for fidelity to the source has **no independent reference** to check the source against. It validates three things well: that data arrived, that the link is up, and that the value parses as the right type. It rarely validates the fourth, which is that the value is physically possible for *this* asset. That check requires a per-asset reference the platform was never given. A 200 MW reading is unremarkable on a 300 MW plant and impossible on a 20 MW one, and a generic range check cannot tell the difference.

1.2 What the reconciliation layer adds

Reference	What it makes checkable
Asset ratings	Nameplate, voltage class, expected operating band. Turns "the value arrived" into "the value is possible."
Executed agreements	The availability construct, its denominator, its exclusions and its clock. Turns "we computed 99.2%" into "we computed the number this agreement defines, or we could not, and here is why."
The document of record	The issued voltage schedule, the operating practice, the controlled procedure. Turns a displayed band into a citable one, or refuses.
The obligation calendar	What is due, when, and on what evidence. Turns a signal into evidence, or states that it is not evidence.
The system's own history	Age, sample counts, unanswered polls, refusal ledgers. Turns a screen into something that can describe its own health.

The practical division. The monitoring platform stays. RenewOps reads from it. Where the two disagree, RenewOps does not overwrite the monitoring platform and does not assert that it is wrong. It publishes both figures with the gap labeled and names which question each one answers. On one site on one day, the platform's own availability figure, the upstream contractual figure and the upstream technical figure read 66.0 percent, 84.76 percent and 51.45 percent. None of the three is wrong at its own source.

1.3 Where responsibility sits

This matters for what can fairly be asked of the layer. RenewOps does not fix a dead pyranometer, restore a link, or improve a plant's production. It cannot make an uncomputable contract computable. Where an agreement left a coefficient blank, no software resolves that. A person negotiates it.

What it changes is that all of the above become **visible, dated, attributable and countable** instead of being absorbed into a number that looks fine. The exposure existed before the layer was installed. The layer makes it surface on your schedule rather than on somebody else's.

2 Five rules, and what each one costs

The whole product reduces to five rules. Each is easy to state, and each carries a cost that explains why most platforms do not follow them.

2.1 Missing renders as unknown, never zero

The cost: a dash in a tile looks like a broken product and a zero looks like a working one. Every commercial pressure in software pushes toward the zero.

As implemented: three populations of false zeros were converted to unknown on 25 August 2026. They were a maximum-possible-power column carrying 3,758 zero rows, the measured-insolation column feeding performance ratio, and a telemetry voltage column. A curtailment determination that was never made is set to *unknown* rather than false, because a false all-clear during a real curtailment is worse than no reading. An infinite insulation-resistance value now renders blank. Infinity there is the absence of a fault path, and it had been rendering as an unusually healthy reading.

2.2 Every tile carries its own denominator

The cost: denominators are ugly, they take space, and they invite the question the tile was built to avoid.

As implemented: site availability publishes as an unweighted mean, *named as unweighted*, with the per-asset minimum and the count below threshold beside it. No rollup weighting was invented, because whether an agreement averages by asset count or weights by nameplate is a contract term. Every DC-loss row states how many paired hours it rests on; a day whose coverage swings more than twenty percent labels itself UNEVEN COVERAGE. Alarm parent rows carry the full child count and the priority mix.

> **In remediation**
>
> An authoritative fleet definition exists, carrying a separate flag for each thing a tile might really be counting. As of 26 August 2026 no tile cites it, and screens still state fleet counts that differ from one another. This is the largest open item against rule two.

2.3 Every number shows its provenance and its age

The cost: age is the field that makes a product look bad. A tile that says "as of eleven days ago" is telling the truth about a situation nobody wants to see.

As implemented: every state stamp carries the full date and an explicit age in days. All of them route through a single formatter, so the whole product was corrected at once rather than card by card. The stalest feed on the fleet reads 796 days. Values resolve to the latest non-null reading *per signal*, each carrying its own as-of. Two battery sites carry a ROLLUP LIVE, UNITS SILENT verdict: the site number refreshes every fifteen minutes while every underlying unit stopped answering days earlier.

2.4 When it cannot compute, it says why

The cost: "cannot compute" is an admission, and it is an admission on a screen somebody is paying for.

As implemented: curtailment tiles print a dash with the reason on the face. Twenty availability rows are withheld with the reason on the row and explicitly not estimated. Sixty-five recurring intervals read cannot-compute rather than showing a manufactured date. Four sites are excluded from a DC comparison, each with its own reason. An event fetcher that hits a response cap halves its window recursively and returns an error rather than writing a partial day that looks whole.

2.5 Refusals are recorded and visible

The cost: a refusal ledger is a public list of the things your product would not do. It is the single least flattering artifact a platform can publish.

As implemented: 367 refused samples at one site in twenty-four hours, each auditable, with the raw value preserved beside the blank. The record of a bad reading is the evidence the gate is needed. Every rejected tracker wind sample is named in an audit view. Two alarm-catalog promotions are held, with the reason held as data. A priority override the classifier cannot read is refused at the point of writing rather than stored where it would silently do nothing.

Four of these five rules make the product look worse in a demonstration than a competitor that follows none of them. That is not a flaw in the rules. It is the reason the category does not exist yet.

3 The measurement floor

Everything else in the platform rests on one question: is this value a measurement? The measurement floor is the set of checks that answer it before any screen is allowed to display anything.

3.1 The plausibility gate

Three independent tests run on every point-of-interconnection sample:

1. A **per-site power ceiling** derived from the plant's own rating rather than a generic range.
2. An **absolute bulk-system voltage ceiling**, plus a per-site band derived from the voltage schedule of record.
3. A **shape test** that rejects a twenty-four-hour span wider than one and a half times its own center. It catches an intermittently corrupt channel that individually plausible samples would pass.

On the day it landed, the gate refused 367 samples at one site in twenty-four hours, with worst raw values around 13,700 MW and 9,280 kV, and **zero false positives anywhere else in the fleet**. That last figure is the one that makes the gate deployable. A plausibility check that fires on healthy plants gets switched off within a week.

The check that checks itself. Where a site's own voltage schedule is suspect, the platform **refuses to use that setpoint as a banding input** and surfaces the site in a coverage view. Six sites carried no assertable band when measured. They are stated as open per-site questions rather than defaulted to a plausible nominal. Banding against a number you do not trust is not a check. It is a check-shaped object.

The raw value is never modified. A refused reading renders blank *with the raw value retained beside it*, because a maintenance engineer needs to see 9,280 kV to understand that a tag is mapped wrong. Deleting it would remove the only evidence of the defect.

3.2 Methods that validate themselves

A derived figure is only as good as the assumption underneath it, so the platform builds the check into the method rather than adding it afterward:

Method	Its own validity test
DC output per MW installed	Eight of nine sites put DC actual within 0.94 to 1.11 of the revenue meter for the same window. That ratio is what makes the rest usable. A site reading 0.82 is physically impossible, because it is a partial plant measured against a whole-plant meter, and it is excluded with that reason stated.
Comparison buckets	Both sides must carry comparable asset coverage before the comparison runs. A day whose coverage swings more than twenty percent labels itself uneven.

Method	Its own validity test
Voltage schedule	Reconciled against the plant controller's own setpoint, band and output, side by side. Agreement at all eight sites where both were available.
Availability ingest	Verified by two independent pulls reproducing the same six site figures exactly. A figure that cannot be reproduced on demand is not a measurement.

3.3 The first reading after a start is not an observation

This rule was adopted after it cost a wrong figure: the first poll following a deployment is discarded. A sweep forty-six minutes old caught twenty-seven of forty-one inverters at zero output. It drove a loss figure that described the poller's cold start rather than the plant.

A related rule runs opposite to what most systems do: **inverters reading zero are counted, not filtered out.** An inverter at zero against a non-zero expectation is precisely the loss being measured. Three earlier versions of the calculation filtered them, and each produced a reassuring number.

3.4 Dead fields inside live rows

A row reads live when *any one* of its fields is fresh. Underneath, per-field liveness with age and unanswered-poll counts found fields dead for between 41 and 293 days, sitting behind between 289 and 413 unanswered polls. The schedule-versus-controller comparison now returns **CANNOT COMPARE** on a dead point. Before, it returned agreement, a pass produced on evidence that does not exist.

> **In remediation as of 27 August 2026**
>
> The per-field liveness migration is staged and had not been run at the time of writing. It is described here as staged, not shipped.

3.5 Units are provenance too

Power factor was stored in two conventions under one field. A fleet ranking consequently named one site the worst on the system when its value was identical to the best. The convention is now settled per site by derivation from measured real and reactive power. Where a site cannot be settled, the fleet readout **refuses and names the reason**. A unit convention is not a formatting detail. It is a statement about what the number means.

4 The alarm layer

An alarm system is an information system whose output is human attention. The layer treats it as one, and measures it accordingly.

4.1 One definition of the board

There are two states and no others: unacknowledged-active and acknowledged. Acknowledging does not archive. An acknowledged alarm that is still asserting carries an **ACK'D · STILL ASSERTING** badge, because "I have seen this" and "this is over" are different statements, and they had been recorded as one. The audible feed is scoped to unacknowledged-active only, so an acknowledged alarm goes quiet and stays visible.

4.2 Status is not an alarm, and cannot latch again

The board went from 7,839 rows to 1,175, with 48 of 6,965 status and informational rows held deliberately. One daylight status event had appeared on an operator's board 2,092 times. Criticals went from seventeen present and none rendering to thirty-two rendering.

The durable half of that fix is a rule at ingest so those rows **never latch again**. A new arrival of that class lands archived, with an explicit archive reason, rather than landing on the board to wait for the next cleanup.

4.3 Priority resolves in the order authority runs

Approved override comes first. Then the reviewed cause catalog, matched by event name and then by code. Then a classifier fallback. A priority posted by a browser can no longer contradict a cataloged code. A vocabulary check reports any priority word outside the known set, so a new upstream value cannot sort quietly to the bottom of every board.

Grouped alarms take the **worst priority among their children plus the full child count**, and the parent row carries a priority-mix qualifier on its face. An operator who reads "High (High, Medium)" knows the group is not uniform. Nothing is discarded. A proposed deduplication would have kept three assets at High *or* eleven at Medium and thrown away the other. It was withdrawn before it ran. Under-reporting count and severity is worse than staleness, because staleness at least announces itself in a timestamp.

4.4 Suppression measured against its own window

Suppression rate publishes with the window it was measured over, which runs from the later of twenty-four hours ago and the moment suppression tagging began. That correction moved a reported 14.2 percent to a measured **80.9 percent**. The behavior had not changed. The denominator had. Even at that rate, 1,131 alarms per hour were still reaching the board.

Deferred

Alarm load is computed against a threshold of sixty alarms per operator-hour, the value configured in this platform. It is not a published standard value. The current ANSI-approved alarm-management standard for the process industries is ANSI/ISA-18.2-2016, *Management of Alarm Systems for the Process Industries*. A 2017 trade article by the two co-chairs of the ISA18 committee gives far lower average rates per operator or operator console: fewer than six an hour as very likely acceptable, about twelve as the maximum manageable, and more than thirty as very likely too demanding. The article presents those as the authors' guidance, not as text of the standard, and we have not read them in the standard itself. Neither they nor the configured sixty are stated as the standard anywhere in this series.

4.5 The monitor that stops

The platform publishes its own heartbeat, monitor health and scheduled-job health, and it labels an unconfirmed alarm rather than clearing it. Two clocks are named separately, minutes since the newest onset and minutes since the newest confirmation, so a stopped feed cannot hide behind a calm board.

What that surfaced. A close-out sweep had been destroying every alarm the platform raised itself: **548 voltage-band and 246 schedule-deviation alarms raised and deleted**, each stamped as though the plant had returned to normal, none ever seen by an operator. The sweep now touches only alarms that originated upstream, identified by origin. On 25 August 2026, for the first time, sixteen voltage alarms stood across nine sites. Twelve were acknowledged within the hour.

4.6 A publish that refuses itself

The board refuses to commit unless four assertions pass. Children summed across all parents must equal the count of unacknowledged-active alarms. Duplicate keys must read zero. Every dependent view must rebuild with its permissions re-asserted. Both layers must complete a concurrent refresh. The duplicate-key counter is checked on every refresh, and a non-zero reading is visible *before* the board freezes rather than after.

That guard exists because of what it found. One pre-computed layer had failed fifty-seven consecutive refreshes, and a second had never once refreshed since it was created. Both were serving last-good contents behind a page that looked live.

5 The operator layer

Situational awareness is a property of the record, not of the dashboard. The test is whether an operator arriving cold can reconstruct the plant from the record alone, without asking anybody.

5.1 What the record answers

Question	Answered by	Status
What is wrong now?	Standing conditions by priority, aged on the source clock with any fallback labeled. 2,096 of 2,126 genuinely still asserting when measured; 257 asserting for over a week; the oldest a critical breaker status at 458 days.	Shipped 25 Aug
What has been done?	A separate operator action log holding what a human actually did, apart from the alarm row a bulk operation can rewrite. Bulk acknowledgment writes its own record in the same transaction.	Shipped 25 Aug
What is outstanding?	Acknowledged-but-still-asserting badged on the row. Events whose end time is *assumed* counted separately from events observed to end.	Shipped 26 Aug
What can I trust?	Per-feed staleness in one place: assets, answered in three hours, answered in twenty-four, age of newest sample. Twelve of thirty feeds stale on enable, the worst at 796 days.	Shipped 25 Aug
Who did it?	Recorded on every controlled write, from a caller-supplied identity.	**Not verified**

5.2 Acknowledge what you say you are acknowledging

A bucket acknowledgment acknowledges the whole bucket server-side and returns the count actually changed. It refuses an unattributed call, refuses an unrecognized priority, and never writes an archive stamp. The confirmation dialog shows the true bucket size. Before the fix, "Ack all (1,707)" acknowledged the hundred rows the browser had loaded and reported success.

5.3 Metrics that decline to rank people

Operator response coverage publishes at **one percent**, and time-to-acknowledge is explicitly marked **not usable for ranking**. The platform works to a simple rule. A response-time view emits four things: the corroborated event count, the central tendency over that corroborated population only, the coverage percentage, and a usability flag computed from the first three. A median drawn from one percent of a population is a different measurement taken from a self-selected sample. A footnote does not travel with a number into whatever chart somebody builds next. A flag does.

The honest version of an operator metric is often the one that says it cannot be used to rank people. It is harder to publish than the median, and it is the only version that survives being cited.

5.4 Events that never closed

About 6,600 events stand open and were never closed. Of those, roughly 200 are over a year old, about 4,400 are between ninety days and a year, and about 2,000 are recent. A separate, unambiguous flag catches night events stuck open: 1,866 fleet-wide, including 779 low-irradiance-at-night events across seven sites, the oldest at 335 days.

The analysis stops there. It would be tidier to conclude that stuck events are what drives the availability figures. The sites do not line up cleanly enough to support that, so the claim is not made.

6 NORA: rules, prompts and procedures

NORA is the operator response layer. It raises prompts from a guarded vocabulary of signal types, holds the controlled procedure library, and drafts switching orders. It also keeps score of how often its own output was ignored.

6.1 The not-needed rate

Every rule is scored by the share of its prompts an operator answered "not needed." On 25 August 2026 six enabled rules all measured between zero and 1.5 percent, with answer counts of 66, 21, 13, 12, 7 and 6. Each was reported as earning its place, and a number rather than an opinion is what earned it.

A rule has three legible states, not two. It is **enabled**; or **disabled, with the reason written onto the rule record**; or **NOT WIRED**, meaning it exists but has no candidate branch and could never fire under any input. That third state is the one that matters most. On 24 August 2026 four of seven rules structurally could not fire while the screen printed *quiet: no independent check available*. An operations manager would reasonably read that as a clean fleet on four dimensions nothing was checking. Enabled rules went from six to nine the following day, and wiring state now renders separately from firing state.

6.2 Two gates on every candidate branch

A rule comparing a partial day's measurement against a full day's model made the whole fleet look faulted every morning and recovered every afternoon. Any candidate branch now requires a **completed-day gate** and a **persistence test**. Re-enabled on completed days, that rule caught four genuine site-days in seven: sites generating between 1,000 and 1,800 MWh against zero reported sunlight.

6.3 The rules it declines to enable

Rule	Why it is off
Availability below guarantee	Deliberately disabled until the capability score is persisted server-side. A rule firing off a number computed in a browser session produces a prompt nobody can audit afterward.
Alarm flood	The only available count does not mean what the rule needs it to mean. Pointing the rule at it would put a fabricated alarms-per-hour figure in front of an operator, so the signal is left unconsumed.
Plausibility	The permitted signal vocabulary contains no signal for *arrived on time, from a healthy link, and is not physically possible*. Widening it governs what NORA may say at all, so it is an owner decision rather than an engineering repair.

The inverted rule is worth recording because the failure is so easy to reproduce. A communications rule counted comms-loss alarms, so it fired on healthy sites and stayed silent on dark ones. **A site that has gone dark cannot send an alarm saying it has gone dark.** Its replacement asks how long it has been since the feed last said anything. That covers dark sites by construction, and it flagged twelve of thirty feeds stale on enable.

6.4 The controlled procedure library

121 controlled documents (operating practices, alarm response guides, switching practices) are live and retrievable, filed by document type, with post-load reconciliation reporting nothing missing. Three behaviors matter more than the count:

- **A company procedure is never returned as the standard it implements.** The document type says what the document is; the standard linkage is a separate field. All fields are supplied explicitly on ingest, because auto-detection previously misfiled eight documents as standards.
- **Supersession runs on ingest.** A replacement retires its predecessor, and a reload retires the entire prior load. There is no state in which two versions are both live.
- **Ingest re-verifies the operator against a roster on every call** and writes an audit row per ingest. An unrostered code was rejected rather than reconciled later.

> **Audit finding raised against the library itself, ahead of its effective date**
>
> 99 of the 121 controlled documents still carry a placeholder document number, and the approver and wet-signature register could not be loaded.

6.5 Machine-drafted switching orders

NORA drafts switching orders. The order number issues from a controlled sequence at submission; it is not typed, and it is not minted at drafting. As measured on 24 August 2026, two related controls were enforced: controlled numbering, and the document provenance discipline above.

The integrity requirements themselves are **specified and not yet enforced**. Four of them bear on this register: expected state required on every step; step provenance and disclosure of any inferred step; separation of drafter, approver and executor; and a refusal recorded as a refusal rather than persisted as the order body. The register as measured on 24 August 2026 held twenty records. Thirteen were cancelled and seven were drafts. **None had ever been submitted, approved or executed.** Expected state was null on every step of every order, and one cancelled record's body was a model's refusal to write it.

Stated plainly because it will be checked. An integrity control that has never been exercised is a design, not a control. The first competent question at a demonstration is "show me an order approved by somebody other than the person who drafted it," and on that register there is no approved order at all. Chapter 10 carries this in the does-not-claim list.

7 The contract layer

Availability is not a number. It is a construct defined in an executed agreement, and the layer's output is whether that construct can be computed at all.

7.1 Computability as the output

State	Meaning
Ready	Construct fully specified, every input held. The number can be produced.
Partial	Construct specified; one or more inputs missing or unreliable, and named.
Blocked	Construct specified and not computable by anybody, including the counterparty. A term was left unpopulated, or the equation is absent from the executed text.
Undefined	The agreement defines no construct: an exhibit referencing an attachment that is not attached, or a section that does not exist in the executed copy.

On the fleet measured, **eleven of twenty-three sites read blocked**. The dominant cause is coefficients left "mutually agreed upon" and never populated. The panel says so. It does not substitute a market convention, and it does not average around the gap.

7.2 Three figures, side by side, gap labeled

The platform's own inverter-count availability publishes beside the upstream contractual figure and the upstream technical figure. The platform's own figure carries a standing disclaimer that it is not the contractual figure. One site on one day: 66.0, 84.76 and 51.45 percent.

The technical figure is ingested, stored and exposed, and it is **consumed by no calculation anywhere**, because its behavior contradicts its name. It read 0.00 percent at night and 11.58 percent at midday against a contractual 98.47 percent on the same asset and day. It is held until somebody with authority defines it. Storing a figure and refusing to compute on it is a legitimate state, and a more honest one than finding a use for it.

7.3 The denominator, and twenty rows that cannot be repaired

The inverter denominator was corrected fleet-wide by counting only rows carrying an inverter flag. The upstream source holds inverters and their sub-modules in one population, and ten of seventeen sites are mixed. **Naming is explicitly rejected as a guide.** At one site the rows named as inverters are the inverters. At another it is the rows named as modules that carry the flag. Only the flag is authoritative.

Twenty historical rows were computed against the wrong denominator, which means the online count was taken over the same wrong population. The percentage cannot be corrected. Those rows are withheld, return blank, carry the reason on the row, and are **not estimated**.

7.4 A figure pulled early is not a settled figure

Contractual availability is not final on the day. Where the contractual and technical figures agree to three decimals, no exclusions have yet been applied. A "yesterday" figure pulled too early is a pre-exclusion number wearing a post-exclusion label. This is flagged and under observation. Ingest is idempotent on its natural keys and accepts a date argument for backfill.

7.5 The calendar behind the guarantee

There are two calendars, one click from the compliance header. The regulatory calendar holds 209 obligations: 86 hard dates, 65 recurring intervals and 58 event clocks. The contractual calendar holds 490: 287 hard dates and 203 intervals, carrying four times the regulatory deadline load inside sixty days.

- **Provisional dates are badged, not hidden.** 84 of 287 dated rows rest on an unconfirmed anchor, badged *date provisional · anchor not confirmed* with the anchor text printed beneath. Nine sites have no usable anchor at all, including two agreements executed with the effective date left blank. That is a legal finding, not a data problem.
- **A cadence is not a date.** 163 obligations that recur without a fixed day load as recurring *windows*, carrying the decision still owed and a confidence rating. They do not become 163 deadlines nobody agreed to.
- **Applicability gates are carried as open questions**, not guessed: 33 unanswered gates control 70 of the 86 regulatory hard dates.

8 The evidence layer

Provenance is what separates a number you can show from a number you can cite. The evidence layer sorts one from the other and refuses to blur them.

8.1 Provenance classes

All twenty-five sites are classified against the voltage schedule of record into six states: a complete document from the operating entity; target only, with no tolerance band; wrong bus; not in the source of record; source unusable; and no row at all. Two live sites had been hiding in that last state, because a per-row grade cannot grade a row that does not exist.

Eight of twenty-five receive a citable in-band percentage, all from complete documented schedules. One site that had been publishing 62.9 percent against a placeholder band on the wrong bus now publishes nothing and reads WRONG BUS.

What the classification fixed. Compliance exports previously carried no provenance flags at all. They published in-band percentages for wrong-bus and explicitly disabled sites, including **tautological readings of exactly 100.0 percent**, where the band had been derived from the median of the very series being scored against it. A perfect score against yourself is the cleanest illustration available of a number that is arithmetically correct and evidentially worthless.

A band auto-detected from a site's own median is labeled **INDICATIVE ONLY** wherever it appears. Each schedule row quotes its source inline: file, path, server-modified date, document key. The source document's own defects are recorded rather than repaired, because silently cleaning a source destroys the evidence that it needed cleaning.

8.2 One write door

The schedule of record may only be set through a single controlled path. That path demands a directive reference, refuses an implausible setpoint, writes history and stamps the source. Two honest notes. First, as of 25 August 2026 **it had no caller**, because schedules arrived through document load; the path is correct and unexercised. Second, a corrupt seeded setpoint at an impossible value was **left in place rather than edited**, and the plausibility layer refuses to use it. Changing a voltage schedule is a person's decision made against a directive. The platform is built so it cannot become a data-quality cleanup.

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 are fabricated evidence, one reassuring and one alarming.

8.3 Obligations that cannot be marked done dishonestly

Recurring intervals compute a next-due date **only from a recorded last-performed date**. With no baseline, all sixty-five regulatory intervals read *cannot compute: last performed not recorded*. Not a default, not the start of the calendar year, not silently overdue. The prompt to record the baseline is itself raised by the rules engine.

Completion behaves by type. A hard date closes that occurrence only. A recurring obligation records LAST PERFORMED, the field that unblocks next-due. Both refuse an unattributed call and a completion date in the future. Reopening requires a written reason that stays on the record. A completed view shows who recorded what, when, how many days early or late, and every reopen with its reason.

After the completion function was verified, **the test rows were deleted**, leaving the performed list deliberately empty. A fabricated performance record on a compliance calendar is worse than no feature. The fabrication is indistinguishable from evidence, and somebody acting in good faith will produce it as evidence.

Notifications compose into an outbox and hold as **held-without-recipient** until an address is configured. They are composed and readable, and never reported as sent.

8.4 Can the evidence survive the platform?

On 28 August 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. There were zero schema errors and zero function errors, and all 220 relations executed.

Two findings came with it. First, the export was **not self-sufficient and did not say so**. It granted 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 nothing recorded. Second, the materialized views must refresh in dependency order. Refreshed alphabetically, four of eighteen fail and break ten ordinary views, and nothing on screen says the numbers are stale. The screens simply report nothing.

> **Deferred, stated precisely**
>
> What was proven is a schema and objects restore from a schema-only export. A data restore has not been demonstrated, and this reference does not describe the result as a proven backup of the evidence.

9 The refusal catalog

If you read one chapter of this reference, read this one. A platform is defined more precisely by what it declines to produce than by what it displays, because anything can be displayed.

Every entry below is a place where RenewOps has a value available, or could readily compute one, and does not. Each was observed on a live production instance between 24 and 28 August 2026.

9.1 Refusals of measurement

What is refused	What appears instead
Implausible telemetry	Blank, with the raw value retained beside it as evidence the gate is needed. 367 samples in twenty-four hours at one site, zero false positives elsewhere.
A setpoint the platform does not trust	Not used as a banding input; the site surfaces in a coverage view. Six sites carry no assertable band, stated as open questions rather than defaulted.
Tracker wind outside a physical band	Rejected, with every rejected sample named in an audit view. 162 controllers at one site reported wind speeds near Mach 2.2, and wind stow is decided off that signal.
An infinite insulation reading	Blank. Infinity is the absence of a fault path, not a perfect reading.
A comparison across uneven coverage	The day labels itself UNEVEN COVERAGE and the bucket does not compare.
A site whose derived total is physically impossible against its meter	Excluded from the comparison with the reason stated: a partial plant measured against a whole-plant meter.
The first poll after a deployment	Discarded. A cold start is not an observation of a plant.
A comparison against a dead point	CANNOT COMPARE, rather than agreement. It had been producing a pass on evidence that does not exist.
A power factor whose convention cannot be settled	Refused, with the reason named, rather than ranked against values in a different convention.

9.2 Refusals of computation

What is refused	What appears instead
A curtailment figure with no directed history	A dash, with the reason on the face of the tile.

What is refused	What appears instead
A determination that was never made	Unknown, never false. A false all-clear is worse than no reading.
Twenty availability rows on a wrong denominator	Blank with the reason on the row. Not estimated, not interpolated, not dropped from the series to keep a chart continuous.
A site availability rollup weighting	An unweighted mean, named as unweighted, with the minimum and the count below threshold. The weighting is a contract term, not a display choice.
A next-due date with no recorded baseline	Cannot compute: last performed not recorded. Sixty-five of sixty-five.
A cadence with no fixed day	A recurring window carrying the decision still owed, not a manufactured deadline. 163 of them.
A partial day written as a whole one	The fetcher halves its window recursively and returns an error.
A figure whose meaning contradicts its name	Stored, exposed, and used in nothing until an owner defines it.

9.3 Refusals of evidence

What is refused	What appears instead
An in-band percentage against a seeded band	Nothing. Seventeen of twenty-five sites publish no figure. A band the operating entity never issued is not evidence of conformance to one it did.
A schedule row for a site with no telemetry	Not written. A band against a dead channel yields silence or a 100 percent excursion the moment a tag appears, and both are fabricated.
Editing a corrupt schedule to make it sensible	Left in place; the plausibility layer refuses to use it. Changing a schedule is a person's decision against a directive.
Repairing a garbled source document	The defect is recorded on the row. Cleaning a source destroys the evidence that it needed cleaning.
Test rows on a compliance calendar	Deleted after verification, leaving the performed list deliberately empty.
A notification with no configured recipient	Held-without-recipient. Composed, readable, never reported as sent.
A completion dated in the future, or unattributed	Refused at the point of writing.

9.4 Refusals of configuration

What is refused	What appears instead
A priority override the classifier cannot read	Rejected on write. An override that cannot take effect is worse than one never made, because somebody will later cite it.
A rule evaluation interval that would make it useless	Refused outright.
A rule pointed at a signal that does not mean what it needs	Left unconsumed rather than approximated. It would have put a fabricated alarms-per-hour figure in front of an operator.
Two catalog promotions	Held, with the reason held as data rather than in a ticket. One would have moved an isolated main transformer off Critical. The other would have buried thirty-two Criticals under a night-time flood.
A board publish that does not reconcile	Refused unless four assertions pass: child counts reconcile, duplicate keys read zero, dependent views rebuild with permissions re-asserted, and both layers refresh.
Widening the rules engine's signal vocabulary	Not done by engineering. It governs what the engine may say, so it is an owner decision.
Write operations on the upstream integration	Denied by default. Verified externally on 27 August 2026: insert, delete and point-create all refused.

9.5 Refusals of conclusion

The last category cannot be built into software, and it does more than any other to decide whether the rest is worth anything. Three examples from the same six weeks:

- A deduplication that would have kept three assets at High *or* eleven at Medium and discarded the other was **withdrawn before it ran**.
- A ratio returning exactly 101.8 percent on every unit at two sites was recorded as a **constant, not a capacity metric**. That is a negative result, written down so nobody repeats the mistake.
- The link between stuck events and the availability figures was **not claimed**. The sites do not line up cleanly enough, and saying it would have been tidier than the evidence.

Refusal costs money on every demonstration and earns it back exactly once, in a room where somebody asks where a number came from.

10 What RenewOps does not claim

A capability reference that lists only capabilities is a brochure. This chapter is the boundary. It is stated at the same level of specificity as the rest, and dated the same way.

10.1 Controls specified but not enforced

Claim not made	Position as of 28 August 2026
Enforced separation of drafter, approver and executor on switching orders	No order on the register has reached submitted, approved or executed. The separation has never been exercised, and an empty register does not demonstrate it.
Expected state enforced on every step	Null on every step of every order on record. Specified, not enforced at write.
Step provenance and inferred-step disclosure	No order carries either field populated.
Refusal recorded as a refusal	One cancelled record currently holds a model's refusal as its order body.
A capability grade that respects its own denominator gate	The gate is deployed and confirmed present in the running bundle. **It does not fire.** One site shows a top grade with 50 of 75 inverters reporting.
An auditable capability score	Computed per browser session, per operator, not persisted. Two operators can hold two numbers with no way to reconcile them.

10.2 Attribution

Content is guarded well across every controlled write in the platform. **Identity is not.** On most paths the actor is a string the caller supplies. Who set a schedule, approved an override, marked an obligation complete or answered a prompt is recorded but not independently verified. Every attribution in the system is therefore a claim rather than a fact. An evidence layer that has solved provenance of *values* while leaving provenance of *actors* open has solved half the problem, and this reference will not describe it as more than that.

10.3 Open items carried, not closed

- **Fleet qualifiers are not on the tiles.** The authoritative fleet definition exists and no screen cites it. Counts still differ between screens.
- **Eleven days of historical false zeros remain written as zeros.** New false zeros are stopped. The historical remediation decision is open, and those rows feed everything downstream.
- **131 of 249 cataloged alarm codes disagree with the priority in use**, including a breaker trip-coil monitor sitting below Critical. Documented, not reconciled.

- **The governed priority override store is empty.** The path exists and refuses bad writes, but nothing reads it yet. Applying an override today writes an audit trail and changes nothing.
- **No rule has been retired on the not-needed rate.** The lifecycle evidenced so far runs enabled → disabled-with-reason → not-wired. A ratchet that has never had to release is not yet proven to release.
- **The per-condition quiet period is owed.** One user answering eleven prompts in ninety-nine seconds muted twenty-eight standing protection conditions at once.
- **99 of 121 controlled documents carry placeholder numbers**, and the approver register could not be loaded.
- **No contract terms are held as structured data.** Computability positions are produced by a person reading an executed exhibit. The platform holds the result, not the reasoning.
- **The upstream source holds twenty-three site trees; the fleet holds twenty-five.** The overlap has not been reconciled. This is stated before fleet-wide ingest rather than discovered afterward.
- **A data restore has not been demonstrated.** What was proven on 28 August 2026 is a schema and objects restore from a schema-only export.

10.4 What the layer cannot do at all

It does not fix a dead sensor, restore a link, or improve production. It cannot make an uncomputable agreement computable; where a coefficient was left blank, a person negotiates that, not software. 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. And a clean result from any of its diagnostics says the *record* is consistent, sourced and current. It does not say the plant is healthy. Several of the worst findings on the fleet measured sat on top of plants running perfectly well.

Conflating the record with the plant is the most common misreading of this work. It produces false alarm in one direction and false comfort in the other.

11 How it is deployed, and where to start

The layer reads from what you already run. It does not replace a monitoring platform, and a fleet that installs it on the assumption that it will is buying the wrong thing.

11.1 The shape of a deployment

1. **Read from the existing monitoring stack.** The layer acquires signals read-only. The monitoring platform stays, and where the two disagree the layer publishes both with the gap labeled rather than overwriting anything.
2. **Establish the references.** Asset ratings, the executed agreements, the documents of record, the obligation set. This step makes checking possible, and it is the step that takes the time. Most of it is a document exercise, not an integration.
3. **Resolve signals to roles.** Decide what each signal is for, per asset, so a value can be checked against something rather than merely stored.
4. **Turn the floor on before the screens.** Plausibility, staleness, denominator and provenance come first. Screens built on an unchecked floor produce confident numbers, which is the problem being solved.
5. **Enable rules on measured branches only.** A rule goes live when its candidate branch has been tested against completed data and shown to fire on real conditions and stay quiet on healthy ones.

11.2 What an assessment produces

Most fleets should not begin with a deployment. They should begin by finding out whether they need one, and the honest way to find out is to look.

Engagement	What comes back
Fleet Exposure Read	Your availability constructs reconciled against your executed agreements, returning a per-site computability position (ready, partial, blocked or undefined) with the specific blocker named for every site that cannot compute.
Monitoring Truth Assessment	The full defect battery against your live stack. Every finding comes back with a source, a timestamp, and the question that resolves it. Where something cannot be verified, it comes back marked deferred with the specific question, not softened into a maybe.
Evidence Provenance Review	Your compliance-relevant signals traced to source and sorted into citable and not, with the seeded, the derived and the self-referential separated from the issued.

You keep the findings either way. If you never buy anything else, you still walk away knowing which of your numbers you can defend. That is not a concession. It is the point. Nobody buys a reconciliation layer before they know they need one.

11.3 Five questions to put to your own fleet first

None of these requires us. Somebody who already has access can answer all five in a week.

1. **Pick a tile showing a percentage. What is its denominator, and does the tile say?** If you have to ask an engineer, an auditor will have to ask one too.
2. **Find a value that is currently zero. Can you tell whether it is a measured zero or an absent one?** If the screen renders both identically, you have the first failure class somewhere in your fleet.
3. **Take the oldest reading on any screen. Does the screen say how old it is?** A time of day with no date is the cheapest version of this defect, and in my view the most common.
4. **Take one availability guarantee out of one executed agreement and try to compute it from your monitoring data.** Not approximately. Exactly as written, with that denominator, those exclusions and that clock.
5. **Pick one compliance-relevant number you would cite at audit and trace it to source.** Is the source a document somebody issued to you, or a value your own system derived from itself?

If all five come back clean, you do not need this layer, and we will say so. On the fleet behind this series, none did.

Glossary

Term	Meaning in this paper
Applicability gate	An open question about whether an obligation applies to a site. Until it is answered, the dates it controls are carried as provisional rather than guessed.
ANSI/ISA-18.2	*Management of Alarm Systems for the Process Industries*, the alarm-management standard published by the International Society of Automation and approved by ANSI. The 2016 edition is current.
Availability construct	The definition of availability in an executed agreement: its equation, denominator, exclusions and clock.
Blocked	A computability state in which the agreement specifies a construct that nobody can compute, including the counterparty, because a term was left unpopulated or the equation is missing from the executed text.
Candidate branch	The logic path through which a rule can fire. A rule with no candidate branch is NOT WIRED.
Completed-day gate	A check that holds a rule until the day being evaluated is complete, so a partial day is never compared against a full day's model.
Computability position	The per-site result of reconciling an availability construct against the data held: ready, partial, blocked or undefined, with the blocker named.
Denominator	The population a percentage is computed over. In this series every tile is required to state it.
Deferred	A status for a question that is open and whose answer is not yet known, stated with the specific question that would resolve it.
In remediation	A status for a defect that has been identified and whose fix is specified but not yet in effect. It is not a capability.
Materialized view	A database view whose results are stored and refreshed, rather than computed on each read. A failed refresh leaves the last-good contents in place.
Monitoring platform	A system that acquires, stores and displays plant signals with fidelity to the source. It is not built to check those signals against contracts, documents of record or asset ratings.
NORA	The operator response layer in RenewOps. It raises prompts, holds the controlled procedure library and drafts switching orders.
NOT WIRED	A rule state: the rule exists but has no candidate branch and cannot fire under any input.

Term	Meaning in this paper
Not-needed rate	The share of a rule's prompts that operators answered "not needed." It is the measure used to decide whether a rule keeps its place.
Plausibility gate	The set of checks that decide whether a sample is physically possible for the asset that reported it: a per-site power ceiling, voltage ceilings and bands, and a shape test.
Provenance class	The classification of a compliance-relevant signal by its source, such as issued by the operating entity, seeded, or derived from the plant's own history.
Reconciliation layer	A layer between the monitoring platform and the documents of record that checks displayed values against asset ratings, executed agreements, issued documents and the obligation calendar, and refuses values it cannot defend.
Refusal ledger	The recorded, visible list of values the platform declined to display or compute, with the reason and, for telemetry, the raw value.
RenewOps	The reconciliation platform built by Energy Compliance, Inc.
Schedule of record	The voltage schedule issued by the entity responsible for it, as distinct from a band seeded or derived from the plant's own data.
Seeded band	A voltage band populated from a placeholder or from the plant's own history rather than from an issued schedule. It is labeled INDICATIVE ONLY and not treated as evidence.
Suppression rate	The share of alarm arrivals filed off the board with a recorded reason, over a stated measurement window.
Tautological reading	A result scored against a reference derived from the same data, such as a 100.0 percent in-band figure against a band built from the series' own median.
Undefined	A computability state in which the agreement defines no construct at all, for example where an exhibit references an attachment that is not attached.

How Energy Compliance Can Help

This reference describes a platform, but the questions behind it belong to whoever has to defend the numbers. Three pieces of our work fit them directly.

Assessment before deployment

Start where Chapter 11 starts. An assessment points our defect battery at your live monitoring stack and hands you the findings, each with a source, a timestamp, and the question that resolves it, **whether or not you buy anything afterward.** The Fleet Exposure Read reconciles your availability constructs against your executed agreements. The Monitoring Truth Assessment runs the full battery. The Evidence Provenance Review sorts your compliance-relevant signals into citable and not.

RenewOps

RenewOps is EC's renewable operations monitoring and analytics platform, and it includes a Data Center module. This reference is its capability statement, dated and bounded: what it does, what it refuses to do, and what it does not claim. Deployed against a fleet, it reads from the monitoring you already run and reconciles it against the references in Chapter 1.

Compliance evidence and audit defense

A refused number is only half the job. The other half is the evidence that stands behind the numbers you do cite. WattOps Compliance (wattopscompliance.com) is EC's NERC compliance evidence platform. We will run one standard at one facility through WattOps at no cost; CIP standards are excluded from that offer. Where the evidence is headed for an audit, our audit defense and RSAW production work picks it up from there.

Talk to us

Give us thirty minutes on your specific situation, and we will tell you inside the first ten whether we would find anything.

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.

The views expressed in this publication do not represent the views of NERC, FERC, or any Regional Entity.

About Energy Compliance, Inc.

Energy Compliance, Inc. is an independent regulatory compliance and advisory firm serving the energy sector. Its focus is NERC, FERC, and RTO/ISO compliance, with 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. Every engagement is overseen by a senior practitioner with regulator-side experience. 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 EC's compliance technology. WattOps Compliance is the firm's NERC compliance evidence platform. RenewOps, the reconciliation layer described throughout this series, is its renewable operations platform.

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

Rigorous Compliance. Defensible Programs.

Sources and Authorities

Alarm management

- American National Standards Institute webstore, ANSI/ISA-18.2-2016, *Management of Alarm Systems for the Process Industries*, listed as "Most recent" and revising ANSI/ISA-18.2-2009, with no amendments or replacement listed (accessed October 17, 2026). https://webstore.ansi.org/standards/isa/ansiisa182016
- Donald G. Dunn and Nicholas P. Sands (ISA18 committee co-chairs), "How to Solve Common Problems with Industrial Alarm Systems," Automation.com (February 13, 2017). Average alarms per hour per operator or operator console: fewer than six very likely acceptable, more than 12 about the maximum manageable, more than 30 very likely too demanding. Presented as the authors' guidance, not attributed to the standard. Secondary source. https://www.automation.com/article/solve-common-problems-industrial-automation-alarm-systems

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, per the NERC VAR standards listing (accessed October 17, 2026; https://www.nerc.com/standards/reliability-standards/var). The standard family behind the voltage schedules of record in Chapters 3, 8 and 9; R2 covers maintaining the schedule provided by the Transmission Operator. No successor version is listed; Project 2021-02 (Modifications to VAR-002) is low priority with no anticipated completion time in the 2026–2028 Reliability Standards Development Plan (December 4, 2025; https://www.nerc.com/globalassets/standards/rsdp/2026-2028-reliability-standards-development-plan_ferc.pdf). https://www.nerc.com/pa/Stand/Reliability%20Standards/VAR-002-4.1.pdf

Measured findings

- All system measurements in this reference (counts, rates, durations, timestamps, ratios and states) are Energy Compliance first-party observations of a live production RenewOps instance, made between 24 and 28 August 2026. They are anonymized under the Disclaimer & Legal Notice and are not available from a public source.

© Energy Compliance, Inc. · EC-WP-1107 · Status as of October 17, 2026