ENERGY COMPLIANCE, INC. | *Field Reference for Monitoring and Fleet Data*

MONITORING TRUTH & FLEET RECONCILIATION SERIES · EC-WP-1102

What the Incoming Operator Cannot See

Turnover, attribution, and the record that has to answer when the control room is not in the room.

Seventeen standing Critical alarms, zero rendered, and a turnover screen reporting none. Acknowledgment time equal to archive time in 4,791 of 4,791 sampled rows, every one naming a person. Operator KPI coverage of one percent. Five failures of situational awareness, each measured, treating awareness as a testable property of the record rather than a quality of a dashboard.

The Operator 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 series 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 against the current standards, which are listed in Sources and Authorities at the back. 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. Several 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 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.

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. 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. 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.

These references carry no requirement-level citations in the body, and that is deliberate; the reasoning is on the previous page. 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.

This one is personal. I have taken turnover in a control room and I have audited the logs afterward, and the gap between the two is where a great deal of enforcement risk lives. The finding that stayed with me is the attribution one. 4,791 acknowledgment records out of 4,791 had an acknowledgment time equal to their archive time to the microsecond, and every one of them named a person. Nobody falsified anything. The system manufactured an attribution and then displayed it as a fact about a human being.

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 (this reference)	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	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: these five tests establish whether an incoming operator could reconstruct the plant's state from the record alone. They say nothing about whether the operators are good at their jobs. A crew with an excellent verbal turnover culture fails these tests as badly as a careless one, because the tests read the record, not the room. And the record is what answers when the room is not available.

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. This reference is a companion to EC-WP-1100 and EC-WP-1101 and assumes neither.

Executive Summary

The turnover screen said zero critical

At the moment it was measured, the plant carried seventeen standing Critical alarms. Among them was a high-voltage breaker open with a main power transformer isolated. The screen an incoming operator reads at shift change said: *0 critical. Oldest critical: none.*

Nothing was broken in the ordinary sense. The alarms existed, correctly, in the database. Every query returned them. The screen was reading a field that had been quietly emptied by an administrative action six weeks earlier, and no mechanism existed by which anyone could have discovered that.

Situational awareness is not a dashboard. It is a specific, testable property: **can an operator arriving cold reconstruct what the plant is doing, what has been done to it, and what is outstanding, from the record alone, without asking anybody?** This document takes that property apart on a live system and finds it failing in five distinct ways.

Failure	Observed
Invisible standing conditions	17 Critical alarms present, 0 rendered. Turnover screen reported none.
Manufactured attribution	Acknowledgment time equaled archive time **to the microsecond in 4,791 of 4,791** sampled rows. Every one named a person.
Unusable response metrics	Operator KPI coverage **1%**. A published mean time to acknowledge of 237 minutes over 923 events, computed from a population that was not acknowledgments.
Fractured identity	One person holding two operator codes; one code used by two people; one row named "E2E Test". Sites addressed across six incompatible namespaces.
Acknowledgment that ends the record	Acknowledging an alarm wrote an archive timestamp, collapsing "I have seen this" into "this is over."

Every one of these is a record-keeping defect that presents as an operational fact. That is what makes them dangerous. Nobody is looking for a wrong number, because nobody knows a number is being produced.

Chapters 2 through 6 take one failure each: mechanism, measurement, consequence, repair. Chapter 7 specifies what a defensible turnover pack actually contains. Chapter 8 is ten questions you can put to your own control room this week. Chapter 9 describes what RenewOps, our own platform, now does about each failure and what it still does not do.

1 What turnover is actually for

Shift turnover is the only moment in continuous operations where the record has to stand on its own. Everything else can be patched by asking the person next to you. At 06:00 that person is leaving. The incoming operator needs four things, each a question about the record rather than the plant:

Question	Answered by	Fails as
What is wrong now?	Standing conditions, by priority, with age from the source clock.	Invisible criticals (Ch. 2)
What has been done?	Operator actions, attributed, timestamped, reconstructable.	Manufactured attribution (Ch. 3)
What is outstanding?	Acknowledged-but-unresolved conditions, held decisions, deferred items.	Acknowledgment collapsed into archive (Ch. 6)
What can I trust?	Feed health, monitor health, and which figures are computed versus withheld.	Everything in EC-WP-1100

1.1 The property, stated testably

> **The turnover test**
>
> Take an operator who has been away for two weeks. Give them the record and nothing else: no colleague, no phone call, no institutional memory. Can they state, without qualification, the standing conditions and their ages, the actions taken and by whom, the items awaiting decision, and which displayed figures are trustworthy?

If any of the four requires asking a person, that part of situational awareness lives in people, not the system, and it will not survive a resignation.

1.2 Why this is a compliance question

Event analysis, mitigation adequacy, and any reconstruction of what a control room knew all draw on the same record. The current transmission operations and operations planning standards list dated operator logs among the evidence a Transmission Operator may offer to show what it did and when, and the event reporting standard names an operator log as one way to evidence that a report was submitted. Every one of those measures assumes the log can be reconstructed and trusted. The failures in this document do not make operators worse at their jobs. They make the *record* unreliable, and that surfaces months later, with the operators not present to explain it.

2 The seventeen alarms nobody could see

Seventeen standing Critical alarms, zero rendered, and a turnover screen reporting none. The mechanism is administrative, not technical, which is why it survived.

2.1 What was standing

Anonymized, the invisible population included four feeder breakers at one site showing general relay trip with a collector breaker open; two breakers at a second site showing relay trip with a transformer breaker open; lockout relay trips at a third with the interconnecting utility's breaker open; two feeder breakers open at a fourth; and at a fifth, **a high-voltage breaker open with the main power transformer isolated.**

Those are not nuisance conditions. Several describe a plant that is disconnected.

2.2 The mechanism

Six bulk archive operations over a two-week period, each recorded with the reason baseline_reset, archived between 3,307 and 7,425 rows apiece, wholesale. The intent was housekeeping after a classification change.

The defect is what happened next. **Alarms that had never actually cleared retained their alarm key.** On the following ingest cycle the reconciler saw a key it already held, treated the condition as known, and did not re-latch it. The condition was still true in the plant. It was no longer true in the record.

Calculation 1 · Conditions that cannot re-latch

```
-- Archived rows whose condition is STILL asserting upstream.
-- Every row returned is a live condition invisible to the operator.
SELECT a.site_id, a.event_name, a.priority, a.archived_at, a.archive_reason, u.still_active_upstream,
ROUND(EXTRACT(epoch FROM (now() - a.first_seen_at))/86400, 1) AS days_asserting
FROM alarms a
JOIN upstream_active_keys u ON u.alarm_key = a.alarm_key
WHERE a.state IN ('archived','auto_resolved') AND u.still_active_upstream IS TRUE
ORDER BY a.priority, days_asserting DESC;
```

Run that query against your own system before reading further. It takes seconds, and it is the single highest-yield check in this document.

2.3 Why the turnover screen could not have caught it

The screen computed its Critical count from the alarm board. The board was correct about its own contents. The defect sat upstream of the board, in what the board had been permitted to contain, and a count is structurally incapable of reporting the absence of rows it was never given.

This generalizes, and the general form is worth holding onto: **a count cannot detect its own incompleteness.** Any figure derived by counting needs an independent statement of coverage alongside it, or it is asserting something it cannot know.

The design rule

A bulk archive must do one of two things: verify that each condition has genuinely cleared at source before archiving it, or invalidate the keys it archives so the condition re-latches on the next cycle.

And the turnover view must carry a **reconciliation line**. Not only "32 Critical standing," but "32 Critical standing · 8,412 upstream active keys reconciled · 0 archived rows still asserting upstream." The numbers are illustrative. The third clause is the one that would have caught this, and it is the one nobody builds.

3 Attribution that was manufactured

Acknowledgment time equaled archive time to the microsecond in 4,791 of 4,791 sampled rows. Every one of them named a person.

3.1 The finding

An alarm table carried an acked_by value on exactly 71,791 rows. That number lined up suspiciously well with the archived population, which made it the first thing worth checking. It was not a coincidence.

Measurement	Result
Sampled rows where acked_at = archived_at exactly	4,791 of 4,791
Rows carrying acked_by	71,791
Rows in a 2,000-row sample archived by bulk reset *and* naming a person	493
Rows in the genuine acknowledgment log	703
Genuine log last written	five days before review
Rows in the real operator action trail	1,822, current

Two timestamps agreeing to the microsecond across nearly five thousand rows is not operator behavior. It is a single UPDATE writing both columns from one now(). The archive routine had been populating the acknowledgment fields as a side effect, and it had been naming whichever operator triggered the routine. That included rows archived by an automated baseline reset, where no acknowledgment of any kind occurred.

Nearly five hundred rows in a two-thousand-row sample recorded a named person acknowledging an alarm they had never seen. The person was real. The alarm was real. The act was not.

3.2 Why this is worse than missing data

An empty attribution field is honest. It says we do not know who did this. A populated one asserts a fact about a named individual's conduct, and it will be read as one: in a performance review, in an event analysis, in a regulatory proceeding.

The failure mode here is the same one that runs through EC-WP-1100: **an absence rendered as a value.** This time the value is a person's name, and that raises the stakes considerably.

Calculation 2 · Is your attribution real?

```
-- Test 1: identical timestamps = one UPDATE, not two human acts.
SELECT COUNT(*) AS sampled, COUNT(*) FILTER (WHERE acked_at = archived_at) AS identical,
COUNT(*) FILTER (WHERE acked_at IS NOT NULL AND archive_reason LIKE '%reset%') AS
acked_by_a_bulk_job, ROUND(100.0 * COUNT(*) FILTER (WHERE acked_at = archived_at) /
NULLIF(COUNT(*),0), 1) AS pct_identical
FROM alarms WHERE acked_at IS NOT NULL;
-- Test 2: does the attribution reconcile to an independent action trail? A coverage_pct near zero
means the attribution has no corroboration.
SELECT COUNT(DISTINCT a.alarm_id) AS alarms_claiming_ack, COUNT(DISTINCT l.alarm_id) AS
alarms_in_action_log, ROUND(100.0 * COUNT(DISTINCT l.alarm_id) / NULLIF(COUNT(DISTINCT
a.alarm_id),0), 2) AS coverage_pct
FROM alarms a LEFT JOIN operator_action_log l ON l.alarm_id = a.alarm_id AND l.action =
'acknowledge'
WHERE a.acked_by IS NOT NULL;
```

On the assessed system, test 2 returned a coverage of about **1%**. Seventy-one thousand rows claimed an acknowledgment. Seven hundred could be corroborated by an independent trail.

The design rule

An operator action is an event, not a column.

Write it to an append-only action log with actor, action, target and timestamp, plus a reason wherever the action is a judgment. Derive any column on the alarm row *from* that log, never in parallel with it. A denormalized convenience column that can be written by anything other than the logging path will eventually be written by something else.

And **no automated process may write an actor field.** If a bulk job archives rows, the actor is the job, named as the job. A job that signs a person's name is forging a record, whatever the intent.

4 Metrics computed from a population that was not the population

A published mean time to acknowledge of 237 minutes, computed over 923 events, from a population in which acknowledgment time was a copy of archive time. The number was arithmetically correct and meant nothing.

4.1 The shape of the error

Three endpoints served operator performance figures. Two of them correctly returned null, because the data they required did not exist. The third returned numbers. That asymmetry is the finding. When two of three implementations refuse and one answers, the one that answers is not the good one. It is the one that failed to check, and it is the one that gets quoted.

Endpoint	Returned	Correct?
Aggregate operator statistics	923 acks · 237 min MTTA	No. Computed over manufactured attribution.
Second KPI endpoint	null	Yes
Third KPI endpoint	null	Yes

It gets worse. The same underlying figures were labeled **mean time to acknowledge** on one screen and **mean time to resolve** on another. Those measure different things, so at least one of the two labels was wrong on every render.

4.2 Defining the two metrics properly

Calculation 3 · TTA and TTR, defined so they cannot be confused

```
-- Time to acknowledge: onset to the moment a HUMAN saw it. The source must be the action log, never a column on the alarm row.
tta_minutes = (ack_event.occurred_at - alarm.first_seen_at) / 60
-- Time to resolve: onset to the moment the CONDITION ended. Resolution is a plant fact, not a screen fact.
ttr_minutes = (alarm.condition_cleared_at - alarm.first_seen_at) / 60
-- Both are undefined unless the terminating event is independently evidenced. Neither may fall back to archived_at.
```

The rule embedded there is short: **archive time terminates neither metric.** Archiving says nothing about when a person looked or when a condition ended, and any metric that falls back to it is measuring administration.

4.3 Coverage as a first-class output

The correct response to 1% attribution coverage is neither to suppress the metric nor to publish it unqualified. It is to **publish the coverage alongside the metric and attach a machine-readable usability flag**, so a consumer cannot accidentally rank people on it.

> **The design rule**
>
> An operator response-time view emits four things per operator, not one. The **count** of corroborated acknowledgment events, meaning those with a matching entry in the operator action log, not merely a name written into the alarm row. The **central tendency**, mean and median both, over that corroborated population only. The **coverage percentage**: corroborated events as a fraction of all acknowledgments claimed under that operator's code. And a **usability boolean**, computed from the first three, stating whether this operator's number may be used for ranking at all.
>
> Set the boolean with two conditions, and both must hold: a minimum event count, so a person with four data points is never ranked, and a minimum coverage floor, so a person whose acknowledgments are mostly uncorroborated is never ranked either. Pick the thresholds to suit the fleet and write them down where the consumer of the number can see them.

The usability boolean is the important output, and it is the one most systems omit. A median response time computed from 1% of the population is not a weak signal. It is a different measurement entirely, drawn from a self-selected sample, and the boolean is what tells every downstream consumer so.

The management consequence, stated plainly. Operator performance metrics get used in conversations that affect careers. A metric built on manufactured attribution will, at some point, be used to tell a specific operator that their response times are poor, on the strength of rows a bulk archive job generated. Coverage is the control that prevents that conversation.

5 Identity: one person, two codes; one code, two people

Attribution cannot be better than identity. On the assessed system one person held two operator codes, one code was used by two people, and one row was attributed to "E2E Test".

5.1 Operator identity

An ingest path correctly rejected an operator code that was not on the roster. Good design, working as intended. The complication was that the same individual's earlier work had been recorded under a *different* code, never on the roster either, and two dozen controlled documents carried that code as their uploader.

Neither record is wrong about the person. Both are unusable for reconstruction, because neither resolves to one identity that joins to anything.

Calculation 4 · Identity integrity

```
-- Codes appearing in action data that are not on the roster.
SELECT DISTINCT actor_code, COUNT(*) AS actions, MIN(occurred_at), MAX(occurred_at) FROM operator_action_log l WHERE NOT EXISTS (SELECT 1 FROM operator_roster r WHERE r.operator_code = l.actor_code) GROUP BY actor_code ORDER BY actions DESC;
-- People holding more than one code, and codes held by more than one person.
SELECT full_name, COUNT(DISTINCT operator_code) AS codes FROM operator_roster GROUP BY full_name HAVING COUNT(DISTINCT operator_code) > 1;
SELECT operator_code, COUNT(DISTINCT full_name) AS people FROM operator_roster GROUP BY operator_code HAVING COUNT(DISTINCT full_name) > 1;
```

5.2 Site identity: the same failure, one level up

Sites were addressed across **six incompatible namespaces**: an internal site key, a forecasting slug, a document-management code, a switching-order identifier, a free-text site name, and the upstream platform's object identifier. Nothing joined them. The observable consequences on the assessed system:

- An administrative board listing **37 rows for a 25-site fleet**.
- A **phantom site carrying 612 live alarms**, produced by a namespace collision.
- Three real sites missing from views entirely, their key in one namespace having no counterpart in another.
- Forecast tables that could not be joined to operational tables at all.

For an operator this is not an abstraction. It means the alarm count on one screen and the alarm count on another are computed over different fleets, and nothing anywhere says so.

The design rule

One identity per entity, and every alias is an explicit mapping row.

Pick the canonical namespace and migrate to it. That is a real project, and it deserves an honest schedule. Until it completes, every cross-namespace join goes through a published alias table with a *confidence* column, and any figure computed across an unresolved alias is withheld rather than estimated.

And **an unrostered actor code is refused at write time**, not cleaned up later. An action attributed to an identity that does not exist is not attributable at all, and it cannot be repaired after the fact without asking a person to remember.

6 Acknowledgment is not resolution

Acknowledging an alarm wrote an archive timestamp. Two distinct operator intentions, *I have seen this* and *this is over*, collapsed into one act.

6.1 The three states an alarm actually has

State	Means	Operator intent
Active, unacknowledged	Condition true; nobody has looked.	Horn sounds. Demands attention.
Acknowledged, still asserting	Condition *still true*; a person has seen it and taken ownership.	Silent, and still outstanding.
Resolved	Condition ended, evidenced at source.	Off the board.

The middle state is the one systems drop, and turnover depends on it entirely. **"Acknowledged but still asserting" is the definition of an outstanding item.** It is the population an incoming operator most needs, and it is the population that vanishes when acknowledgment writes an archive time.

6.2 The corrected behavior

Four changes, each small, together restore the middle state:

- Acknowledgment writes state = 'acked' and **does not write an archive timestamp**.
- The board's definition includes both active-unacknowledged and acknowledged rows, so an acknowledged condition stays visible while it is still true.
- The audible feed is scoped to **unacknowledged only**. An acknowledged alarm is silent but present, which is the entire purpose of acknowledgment.
- An explicit **ACK'D · STILL ASSERTING** badge makes the state legible at a glance instead of leaving the operator to infer it from the absence of a sound.

Calculation 5 · The outstanding-items population

```
-- The single most important query in a turnover pack.
SELECT a.site_id, a.event_name, a.priority, l.actor_code AS acknowledged_by, l.occurred_at AS
acknowledged_at, ROUND(EXTRACT(epoch FROM (now()-a.first_seen_at))/3600,1) AS
hours_asserting, ROUND(EXTRACT(epoch FROM (now()-l.occurred_at))/3600,1) AS
hours_since_ack, a.last_confirmed_at
FROM alarms a
JOIN LATERAL (SELECT actor_code, occurred_at FROM operator_action_log WHERE alarm_id =
a.alarm_id AND action = 'acknowledge' ORDER BY occurred_at DESC LIMIT 1) l ON true
WHERE a.state = 'acked'
ORDER BY a.priority, hours_since_ack DESC;
```

An alarm acknowledged eleven hours ago that is still asserting is the most informative row on a turnover screen. It says a colleague saw this, took it on, and it is not finished. That is precisely the handover the outgoing operator would give verbally, and it is precisely what a collapsed acknowledgment destroys.

6.3 A scoping detail that matters more than it looks

An "acknowledge all" control on the assessed system was scoped to *loaded* rows rather than to the bucket the operator believed they were acknowledging. With pagination in play, an operator who pressed it acknowledged the visible page and believed they had acknowledged the category.

The resulting record is worse than incomplete. It is *wrong in a specific direction*. It shows an operator acknowledging some members of a category and, by omission, appearing to have deliberately left the others. After the fact, there is no way to tell that apart from a considered decision.

7 What a defensible turnover pack contains

Eight sections. Each answers a question the incoming operator would otherwise have to ask a person, and each can be produced by query.

#	Section	Must carry
1	**Standing conditions**	By priority, with age from the *source* clock, and a label wherever a local fallback clock was used.
2	**Reconciliation line**	Upstream active keys reconciled; archived rows still asserting upstream. The check that catches Chapter 2.
3	**Outstanding items**	Acknowledged-but-still-asserting, with who took it and how long ago.
4	**Actions this shift**	From the append-only action log. Actor, action, target, time, and a reason where the action was a judgment.
5	**Feed and monitor health**	Per channel, not per site. Newest onset and newest confirmation as a pair.
6	**Withheld figures**	What could not be computed, and the specific input that would resolve each.
7	**Held decisions**	Anything deliberately not actioned, with the reason recorded as data.
8	**Coverage statement**	Sites reporting of sites expected. A count that cannot state its own coverage is asserting something it does not know.

Section 6 is the one that gets argued about. Publishing a list of things the system could not compute feels like advertising weakness. It is the opposite. An operator who knows that three sites have no trustworthy voltage reading can work around it. An operator who is shown a number for those sites cannot. The withheld list is the difference between a degraded system and a lying one.

7.1 The reconstruction test

Once the pack exists, test it the only way that means anything. Take a real incident from three months ago. Using **only** the record, reconstruct what was standing, what was done, by whom, in what order, and what was outstanding at each shift boundary.

Every point at which you have to ask a person is a gap. It will still be there when the question comes from outside the organization, and by then the person may not be.

8 Ten questions for your own control room

#	Question	What a bad answer means
1	Are any archived alarms still asserting upstream?	Every one is a live condition invisible to the operator. Run this first.
2	Does your turnover view publish a reconciliation line, or only a count?	A count cannot detect its own incompleteness.
3	Do acknowledgment and archive timestamps ever match exactly?	Microsecond agreement is one UPDATE, not two human acts.
4	What fraction of claimed acknowledgments appear in an independent action log?	In my judgment, below about 80% no response metric derived from them is usable.
5	Can an automated job write an operator's name?	If yes, some rows in your record are forgeries, however unintentional.
6	Is the same figure labeled MTTA on one screen and MTTR on another?	At least one label is wrong on every render.
7	Does any actor code appear in your data that is not on the roster?	Those actions are unattributable and cannot be repaired later.
8	How many namespaces address a site, and what joins them?	More than one without a published alias table means your screens count different fleets.
9	Does acknowledging an alarm remove it from the board?	Then you have no outstanding-items population, and that population is the core of turnover.
10	Is "acknowledge all" scoped to the bucket or to the loaded page?	Page scoping produces a record showing deliberate omissions that never happened.

Situational awareness is not a screen. It is whether the record can answer for the control room when the control room is not in the room.

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.

Five failures of situational awareness, and the RenewOps behavior built to answer each. The test throughout is the one from Chapter 1: can an operator arriving cold reconstruct the plant from the record alone?

9.1 The standing conditions are visible, and stay visible

Acknowledging an alarm no longer archives it. The board carries unacknowledged-active and acknowledged alarms, and an alarm acknowledged while still asserting carries an explicit **ACK'D · STILL ASSERTING** badge. The incoming operator sees what the outgoing operator saw, and also sees what they did about it. Those are two different columns.

Standing alarms are banded by asserting age **on the source clock**, and any row where that clock is a fallback is labeled as such. On the day it was measured, 2,096 of 2,126 were genuinely still asserting, 257 had been asserting for over a week, and the oldest was a critical breaker status at **458 days**.

The Criticals that were present and rendering as zero now render. On the day of measurement there were thirty-two of them, and all thirty-two were on the board.

9.2 Two clocks, both named

The buckets stamp **onset**. The event log stamps the **latest real transition** (raised, acknowledged, or returned to normal) and names the kind of transition on the row. A live time-base readout publishes minutes since the newest onset alongside minutes since the newest confirmation.

That second number answers the question this reference opens with. A board can be calm because nothing is wrong, or calm because nothing is arriving. Minutes since the newest confirmation separates the two, and it sits on the screen rather than in a health dashboard somebody would have to think to open.

9.3 Attribution that does not manufacture a person

The finding in Chapter 3 was that acknowledgment time equaled archive time to the microsecond in 4,791 of 4,791 sampled rows, every one of them naming a person. The answer is not a better acknowledgment timestamp. It is a **separate operator action log** that records what a human actually did, held apart from the alarm row that a bulk operation can rewrite.

Operator response coverage is published at **one percent**, and time to acknowledge is **explicitly marked not usable for ranking**. RenewOps works to the rule from Chapter 4: a response-time

view emits the corroborated event count, the central tendency over that corroborated population, the coverage percentage, and a usability flag computed from all three. A median drawn from one percent of a population is not a weak signal. It is a different measurement, taken from a self-selected sample, and a footnote does not travel with the 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. Publishing that is harder than publishing the median, and it is the only version that survives being cited.

9.4 Identity that is checked against a roster

The controlled document ingest re-verifies the operator against a roster on every call and writes an audit row per ingest. In testing, an unrostered operator code was rejected rather than accepted and reconciled later.

> **Deferred, and it is the significant one**
>
> Content is guarded well across the platform's controlled writes. **Identity is not.** On most paths the actor is a string the caller supplies, so who set, approved, changed, acknowledged or responded is recorded but not independently verified. Until that closes, every attribution in the system is a claim rather than a fact, and this reference will not describe it as anything else.

9.5 What an incoming operator can now reconstruct

Question	Answered by	Status
What is wrong now?	Standing conditions by priority, aged on the source clock, with the fallback labeled; parent rows carrying child counts and priority mix.	Shipped 25 Aug
What has been done?	A separate operator action log; acknowledgment no longer collapsed into archive; bulk acknowledgment writes its own record.	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 the newest sample. Twelve of thirty feeds were stale on the day the rule was enabled, the worst at 796 days.	Shipped 25 Aug
Who did it?	Recorded on every controlled write, from a caller-supplied identity.	**Not verified**

9.6 Events that never closed

About 6,600 events stand open and were never closed. Banded by age: roughly 200 over a year, about 4,400 between ninety days and a year, and about 2,000 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. Eleven months open is not a plant condition.

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

Figures in this chapter measured 24–26 August 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
Acknowledgment	An operator's recorded statement that they have seen an alarm and taken ownership of it. In this paper it is distinct from resolution and from archiving.
Alarm key	The identifier a monitoring system uses to recognize a recurring alarm condition across ingest cycles.
Alias table	A published mapping that ties one entity's identifiers in different namespaces to a single canonical identity, with a confidence value per mapping.
Append-only action log	A record of operator actions to which entries are added but never edited or overwritten.
Archive	A record-keeping act that removes an alarm row from the active board. It says nothing about whether the condition ended or anyone looked.
Baseline reset	A bulk archive operation used for housekeeping, such as after an alarm classification change.
Coverage	The fraction of a population a figure is actually computed over, such as corroborated acknowledgments as a share of all claimed.
Latch / re-latch	An alarm latches when the system registers it as standing. Re-latching is registering it again on a later cycle while it is still true.
MTTA (mean time to acknowledge)	Average time from alarm onset to a human acknowledgment, evidenced in the action log.
MTTR (mean time to resolve)	Average time from alarm onset to the end of the underlying condition, evidenced at source.
Namespace	One scheme for naming or keying an entity, such as a site key, a document-management code or a free-text name.
Outstanding item	An alarm that has been acknowledged and is still asserting.
Reconciliation line	A statement published with a count showing what the count was checked against: upstream active keys reconciled, archived rows still asserting upstream.
Roster	The authoritative list of valid operator codes and the people who hold them.
Source clock	The timestamp of an event as recorded by the originating device or system, as opposed to the time the monitoring platform received it.
Usability flag	A machine-readable boolean carried with a metric that states whether the metric may be used for a given purpose, such as ranking operators.

How Energy Compliance Can Help

Everything in this reference comes down to one question: when the control room is not in the room, can the record answer for it? Here is where we fit if the answer is no, or not yet.

Look first

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 each site that cannot compute.

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

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 out from the issued.

The platforms

RenewOps. EC's renewable operations monitoring and analytics platform, with a Data Center module. Chapter 9 covers what it does about each failure and what it still does not do. Every tile carries its own denominator, its own provenance and its own age, and it refuses to compute, visibly, when it cannot.

WattOps Compliance. EC's NERC compliance evidence platform, at wattopscompliance.com. We will run one standard at one facility through WattOps at no cost, CIP standards excluded, so you can see how your own record holds up before you commit to anything.

The services

Audit defense. When a Regional Entity asks what the control room knew and when, the answer has to come from the record. We prepare the position, the evidence package and the narrative, and we stay through the engagement.

Compliance program management. Monthly evidence production, RSAW narrative development and mitigation strategy, structured so the record can be reconstructed and trusted the day somebody asks.

Where to start. Run the query in Chapter 2 against your own system. If it returns rows, call me. If it returns nothing, run the ten questions in Chapter 8 and call me about whichever one you could not answer from the record alone.

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 a portfolio of compliance technology: WattOps Compliance, the firm's NERC compliance evidence platform, and RenewOps, the reconciliation layer described throughout this series. The firm is based in Bloomington, Minnesota.

Sources and Authorities

The findings in this reference are Energy Compliance's own dated measurements of one live production system and of RenewOps, as stated in the text. They are first-party observations and have no external primary source. The regulatory sources below support the family-level statements in Chapter 1 and the Disclaimer. All were accessed October 17, 2026; status, effective date and docket for each standard are as shown on its page under nerc.com/standards/reliability-standards. No standard is cited for any threshold or figure in this reference; the only threshold stated, the 80% coverage floor in Chapter 8, is the author's judgment.

NERC Reliability Standards

- NERC, Reliability Standard TOP-001-6, Transmission Operations. Status on the NERC standard page: Mandatory Subject to Enforcement, effective April 1, 2024 (FERC Docket No. RD22-2-000). Measure M1 and several later measures list dated operator logs among the evidence a Transmission Operator may provide. https://www.nerc.com/globalassets/standards/reliability-standards/top/top-001-6.pdf
- NERC, Reliability Standard TOP-002-5, Operations Planning. Status on the NERC standard page: Mandatory Subject to Enforcement, effective October 1, 2025 (FERC Docket No. RD24-1-000). Measures M3 and following list dated operator logs or email records as evidence. https://www.nerc.com/globalassets/standards/reliability-standards/top/top-002-5.pdf
- NERC, Reliability Standard EOP-004-4, Event Reporting. Status on the NERC standard page: Mandatory Subject to Enforcement, effective April 1, 2019, inactive date September 30, 2027 (FERC Docket No. RM17-12-000). Applies to Reliability Coordinators, Balancing Authorities, Transmission Owners, Transmission Operators, Generator Owners, Generator Operators, and Distribution Providers. Measure M2 names an operator log among the evidence of event report submittal. https://www.nerc.com/globalassets/standards/reliability-standards/eop/eop-004-4.pdf
- NERC, Reliability Standard EOP-004-5, Event Reporting. Status on the NERC standard page: Subject to Future Enforcement, effective October 1, 2027 (FERC Docket No. RD26-8-000). EOP-004-4 governs until then. https://www.nerc.com/standards/reliability-standards/eop/eop-004-5

© Energy Compliance, Inc. · EC-WP-1102 · Status as of October 17, 2026