MONITORING TRUTH & FLEET RECONCILIATION SERIES · EC-WP-1105

Six Ways to Compute 99%

Availability monitoring: constructs, computability, and the denominator that doubled overnight.

More than a dozen sites, one guarantee, six structurally different definitions of it. Coefficients left to be mutually agreed and never populated. Excursions below guarantee with no physical event behind them. The contractual half and the computational half, with the arithmetic for both, and computability returned as a per-site position rather than a score.

The Contract 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 the absence of regulatory citations

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, never as citations. 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. The few general technical definitions this reference relies on (time-based availability, performance ratio, DC and AC capacity) were checked against the public technical sources listed in Sources and Authorities at the back.

On the findings

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

builds, and one of them appears in this reference. 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, and 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. 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 one vendor's failure, 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 regulatory citations, and that is deliberate; the reasoning is in the Disclaimer at the front. 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.

More than a dozen sites, one guarantee, six structurally different definitions of it. On several of those sites the availability equation depended on coefficients the agreement left to be mutually agreed upon later, and they were never populated. Neither party could have computed the number they had both signed. Nobody did anything wrong. The construct came out of the exhibit and went into a platform that was never asked to read the exhibit.

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

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

Rob Smith

Founder and Principal, Energy Compliance, Inc.

Bloomington, Minnesota

How to Use This Reference

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

Reference	Title	What it covers
EC-WP-1100 · The Measurement Layer	Missing Is Not Zero	The six-class failure taxonomy and the five design rules that answer it. The foundational reference; every paper that follows takes one class apart on a live system.
EC-WP-1101 · The Alarm Layer	The Board That Deleted Its Own Alarms	Alarm management treated as an information system: load, priority, latching, and the monitor that stops. Seven failure modes, each with the arithmetic that exposes it.
EC-WP-1102 · The Operator Layer	What the Incoming Operator Cannot See	Situational awareness as a testable property of the record rather than a quality of a dashboard. Turnover, attribution, and identity, measured five ways.
EC-WP-1103 · The Rules Layer	A Rules Engine That Measures Its Own Usefulness	The not-needed rate: how a rules engine earns the right to keep a rule. The measurement that breaks the ratchet of ever-growing alarm catalogs.
EC-WP-1104 · The Procedure Layer	A Machine May Draft It	Switching order integrity when a machine prepares the draft. Five integrity requirements, the failure behind each, and the one field a drafter must never write.
EC-WP-1105 · The Contract Layer	Six Ways to Compute 99% (this reference)	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: computability is a statement about whether a number can be produced from the terms as written and the data as held. A site marked ready is a site where the arithmetic runs. It is not a site where the counterparty has agreed the arithmetic, and it is certainly not a legal opinion about what the agreement means. Where a construct is ambiguous, this reference says so and stops. Resolving the ambiguity is counsel's work.

About this document

The findings here come from a live production platform and from executed agreements read in full. No contract term, rate or commercial figure from any executed agreement is reproduced. What travels is the structure of each failure and the arithmetic that exposes it.

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.

Executive Summary

Everyone guarantees 99%. Nobody means the same thing.

Across more than a dozen sites in one portfolio, availability was defined **six structurally different ways**. These were not six formulas expressing one idea. They were six constructs, with different denominators, different exclusions and different clocks. Every solar site guaranteed 99%.

On several of those sites the contractual equation depended on coefficients the agreement left *to be mutually agreed upon*. They were never populated. Neither party could have computed the guarantee they had both signed, and neither tried. Two further exhibits referenced attachments that were not present in the executed copy.

Set against that, one site's agreement allowed termination on any miss below its guarantee, with **no cure period**. That site's availability construct was among the ones that could not be computed.

If you guarantee a number, somebody should be able to compute it. That is a low bar, and it is not being cleared.

The two halves of the problem

Availability fails in two independent places, and most organizations are only looking at one of them.

Half	The question	Where it fails
Contractual	What does the executed agreement actually require, and can it be evaluated at all?	Coefficients never populated, missing exhibits, constructs that differ per site
Computational	Given a construct, is the arithmetic being performed correctly against the right population?	Denominators that move without a physical event; two asset populations conflated; guarantees compared against figures that do not exist

The second half produces failures that are almost impossible to see from the outside, because the output is always a plausible percentage.

One note on the counts before the table. The figures in this reference were taken from four populations that overlap but are not the same list: the sites whose executed agreements were

read in full and the sites in the asset crosswalk table (each somewhat more than a dozen), and the site trees in the upstream source and the sites in the platform's own fleet list (each roughly two dozen). Each figure below and in the chapters names the population it was counted against, and Chapter 9 states plainly that the lists have not yet been reconciled. On the assessed fleet:

Observed	Measure
Site availability excursions with no physical event (100% → 50% → 100% at the site in 4.1)	More than one site
Cause, at the site in 4.1: denominator moving from the inverter rows (34) to the module rows (68) while the numerator never moved	34/34 → 34/68 → 34/34
One site's denominator briefly becoming 300 (source not yet confirmed: AC nameplate in megawatts, or a module-row count)	300
Fleet inverter count, as reported across views	1,993 vs 2,293
Crosswalk-table sites where a single asset table holds two different populations	Most, about ten
Sites with a guarantee comparison published against an uncomputable figure (nearly all blocked, one partial)	About a dozen
Historical rows withheld rather than corrected after the fix	20

Chapter 1 sets out the four families of availability construct. Chapters 2 and 3 take the contractual half. Chapters 4 through 6 take the computational half. Chapter 7 states what an availability figure has to carry to be defensible, Chapter 8 is ten questions for your own portfolio, and Chapter 9 describes what RenewOps, our own platform, does about it and what it does not yet do.

1 Four ways to compute availability, and why the choice is contractual

Availability is not one metric. It is a family, and which member of the family applies to a given site is a question about that site's executed agreement, not about engineering preference.

1.1 The four constructs

Construct	Measures	Sensitive to
Time-based	Fraction of period the asset was capable of operating.	Definition of "capable"; exclusion handling
Production-based	Actual energy against energy that would have been produced but for unavailability.	The counterfactual: how expected output is derived
Capability / capacity	Available capacity against installed capacity at each interval.	The denominator (Chapter 4)
Performance-ratio-linked	Delivered performance against a reference model.	Irradiance and nameplate accuracy

The time-based form is the one used in the public PV study cited at the back. The Department of Energy's Federal Energy Management Program, in a 2021 assessment of 75 federal PV systems, defines availability as actual measured hours of production divided by modeled hours of production, and performance ratio as measured production divided by modeled production during timesteps when the system is available. Those are one study's working definitions. They are not what any particular agreement says, which is the point of this chapter.

Calculation 1 · The four, stated

```
-- Time-based. The exclusion set is where the argument lives.
A_time = (period_hours - unavailable_hours + excluded_hours) / period_hours
-- Production-based. The counterfactual is where the argument lives.
A_prod = actual_MWh / (actual_MWh + lost_MWh_due_to_unavailability)
-- Capability. The DENOMINATOR is where the argument lives.
-- available_capacity summed over INSTALLED units, never over healthy ones.
A_cap = SUM(available_capacity_MW) / SUM(installed_capacity_MW)
-- Performance-ratio linked. Two references, both fallible.
A_pr = PR_measured / PR_reference
```

1.2 Why the choice cannot be made by engineering

Each construct produces a different number from identical plant behavior, and the differences are not small. A site with half its inverters offline but full grid connection scores near 100% on a naive time-based construct and near 50% on a capability construct. Both are correct implementations of their own definition.

Only the executed agreement can say which one is owed. Where the monitoring platform computes one and the contract requires another, the reported number and the contracted number are different quantities that happen to share a name and a unit. On the assessed fleet the difference had never been quantified, because nobody had been asked to.

> **The question to ask internally this week**
>
> Does the availability figure you report to each counterparty use that counterparty's construct, or your internal one? My view: the two can differ, and until someone computes the delta, nobody in the organization knows there are two numbers.

2 Computability: the state before the number

Before asking what availability *is*, ask whether it can be computed at all. On the assessed portfolio, for a substantial fraction of sites, the honest answer was no.

2.1 Four states, per site

State	Means	Typical blocker
Ready	Construct readable, inputs present, figure computed.	None
Partial	Construct readable; one or more inputs incomplete.	Instrument coverage below what the method assumes.
Blocked	Construct readable but cannot be evaluated.	Model coefficients left "to be mutually agreed" and never populated; or the equation is absent from the text layer of the executed copy.
Undefined	No construct can be established from the executed document.	Exhibit prints "[see attached]" with nothing attached; or cites a section that does not exist in the document.

The distinction between **blocked** and **undefined** matters commercially. Blocked is a bilateral gap. Both parties signed something evaluable in principle, and neither supplied the inputs. In our view, agreement on the inputs can close a blocked gap. Undefined means the agreement does not, on its face, define the thing it guarantees. That likely needs an amendment, and that is counsel's call.

2.2 The rule that follows

> **The design rule**
>
> **A site that is not *ready* publishes no availability figure, and no comparison against guarantee.**
>
> It publishes its state, the specific blocker, and the input that would resolve it. That is more useful than a number, and it is the only defensible output.

2.3 The failure this rule exists to prevent

On the assessed platform, a comparison against guarantee was populated for **about a dozen sites whose contractual figure could not be computed**, nearly all of them blocked and one

partial. Chapter 9 reports a blocked count of about a dozen against the roughly two dozen site trees the upstream source holds. Six of them read exactly +1.0.

A variance of exactly one point, on a site whose construct cannot be evaluated, is not a measurement. It is an artifact of a default. And it rendered in the same column, in the same format, as genuine variances from sites that were computing correctly.

Calculation 2 · Gate the guarantee comparison on computability

```
SELECT c.site_id, c.computability_state, c.blocking_input,
-- the figure itself, only where the construct is evaluable
CASE WHEN c.computability_state = 'ready' THEN a.availability_pct END AS
contractual_availability,
-- and the comparison, gated the same way
CASE WHEN c.computability_state = 'ready' THEN ROUND(a.availability_pct - g.guarantee_pct, 2)
ELSE NULL END AS vs_guarantee,
CASE WHEN c.computability_state <> 'ready' THEN 'not computable: ' || c.blocking_input END AS
withheld_reason
FROM availability_computability c
LEFT JOIN availability_monthly a USING (site_id, period)
LEFT JOIN site_guarantee g USING (site_id)
ORDER BY c.computability_state, c.site_id;
```

Note that withheld_reason is a column, not a footnote. It travels with the row into every export and board pack, and a footnote does not.

3 Reading the construct out of the executed document

Five questions, per site, answered from the executed agreement rather than from what everyone assumes it says.

3.1 The five questions

1. **What exactly is the denominator, in the contract's own words?** Installed capacity, contracted capacity, capacity available for dispatch, or an asset count. These give different answers, and the agreement usually specifies one.
2. **Which exclusions apply, and can each be identified in your data?** Force majeure, grid outage, curtailment by the offtaker, scheduled maintenance within an agreed window. An exclusion you cannot identify in the data is an exclusion you cannot take.
3. **Is every coefficient the equation needs actually populated somewhere?** This is the question that returned *no* most often.
4. **Does the exhibit reference an attachment, and is that attachment in the executed copy you hold?**
5. **Does the number you report to this counterparty use this construct, or your internal one?**

3.2 Question three, in detail

Several sites specified a regression-based expected-output model whose coefficients the agreement left to be mutually agreed after execution. The coefficients were never agreed. The practical consequence:

- Neither party can compute the contractual figure.
- Any figure being reported is therefore computed some other way, which by definition is not the contracted way.
- Neither party can produce the contractual figure, so neither can rely on it without first resolving the coefficients. Whether the guarantee can be enforced is a question for counsel.
- Whoever raises it first controls how it is characterized.

> **The asymmetry worth internalizing**
>
> Discovered internally, an uncomputable guarantee is a remediation project on your own timeline, resolved by a short bilateral agreement on coefficients. Discovered by a counterparty during a dispute, it is a credibility problem attached to a number you have been reporting for years. The gap is identical. The position is not.

3.3 Exclusions you cannot identify

An exclusion is worth nothing unless the data can evidence it. On the assessed fleet the curtailment record could not support a curtailment exclusion at all. The flag indicating curtailment was *false on every row* across ten thousand rows. The directed-curtailment quantity was null on every row. Coverage extended to three site keys, and one of them was not a fleet site.

During a live curtailment event at one site, with output walking from 300 MW to 37.5 MW inside fifty minutes, the record went on reporting *not curtailed*.

Ours, and the correct handling

The remediation set that flag to NULL wherever no determination had been made. *False* asserted "not curtailed," a claim the data never supported, and a false all-clear during a real curtailment is worse than no reading. A companion view now names, per site, exactly which input is missing.

4 The denominator that doubled overnight

More than one site showed availability excursions across four daily readings with no physical event behind them. At the site below, availability moved 100% → 50% → 100%. The numerator never moved.

4.1 The observation

Day	Online	Denominator	Availability	vs guarantee
Day 1	34	34	100.0%	+1.0
Day 2	34	68	50.0%	−49.0
Day 3	34	68	50.0%	−49.0
Day 4	34	34	100.0%	+1.0

A forty-nine point excursion below guarantee, on a site where nothing happened. Other sites showed the same kind of excursion; the count is pending confirmation. On one site the denominator briefly became **300**. Whether that is the site's AC nameplate in megawatts substituted for a count of inverters, or a module-row count (Site B below holds 300 module rows), has not yet been confirmed against the measurement record.

4.2 The root cause

The asset crosswalk table carried **two different populations in one table**: inverters, and the modules that make them up. It already distinguished them with a boolean flag. **The ingest was not applying the filter.**

Whether a given snapshot picked up the inverter population or the module population depended on which rows the load happened to touch. Hence the exact multiples, and hence the reversion when the next load behaved differently.

Site	Combined	Inverters	Modules
Site A	224	32	192
Site B	375	75	300
Site C	450	375	75
Site D	205	41	164

Site	Combined	Inverters	Modules
Site E	153	51	102
Site F	102	34	68

Most of the crosswalk-table sites, about ten, were mixed in this way. Counting on the flag gives a correct fleet-wide inverter population of **2,293**. A separate view reported 1,993. The three-hundred-unit discrepancy resolved entirely to one site, where a view filtered on an asset type that means "inverter" at every other site and something else entirely at that one.

Naming is not a reliable guide and must not be used as one. At one site the rows named as inverters are the inverters. At another, the rows named as modules are the ones carrying the inverter flag. Any logic inferring population from a naming convention will be right most of the time and silently wrong at the exceptions.

4.3 The fix, and the twenty rows that were not fixed

A single function became the sole source of the site inverter count, and a guard trigger stamps the population flag on every row at write time.

Then the hard part. **Twenty historical rows were withheld rather than recomputed.** The reasoning is worth stating in full, because the temptation to backfill is strong and wrong.

A wrong denominator means the online count was taken over the *same* wrong population. Correcting the denominator alone does not repair the ratio. It produces a new ratio between two quantities that were never measured over the same set.

The correct output is NULL with the reason on the row: *denominator population uncertain for this date; figure withheld.* Twenty days of withheld availability is a defensible record. Twenty days of recomputed availability is a fabricated one, and nobody can detect the fabrication afterward.

Calculation 3 · Denominator stability: run this on your own history

```
-- A denominator should only change when assets are commissioned or decommissioned. Any other movement is a data defect.
WITH d AS ( SELECT site_id, as_of_date, denominator_count, LAG(denominator_count) OVER (PARTITION BY site_id ORDER BY as_of_date) AS prev FROM availability_daily )
SELECT site_id, as_of_date, prev, denominator_count, denominator_count - prev AS delta,
ROUND(denominator_count::numeric / NULLIF(prev,0), 3) AS ratio,
NOT EXISTS (SELECT 1 FROM asset_commissioning_events e WHERE e.site_id = d.site_id AND e.event_date = d.as_of_date) AS unexplained
FROM d WHERE prev IS NOT NULL AND denominator_count <> prev
ORDER BY unexplained DESC, ABS(denominator_count - prev) DESC;
```

Read the ratio column first. Exact ratios such as 2.000, 0.500 and 1.500 are the signature of a population swap rather than a physical change. Commissioning moves a denominator by the units added; it does not double it overnight and halve it back.

5 Nameplate, and the reference that poisons everything downstream

Nameplate is the reference used by capability availability, by performance ratio, by plausibility bounds, and by every per-MW commercial calculation. A wrong nameplate is wrong in four places at once, and the four disagree in ways nobody traces back.

5.1 Observed discrepancies

Comparing contracted capacity against the capacity held in the operational platform, two sites stood out: one showed a **31% gap**, another an 8.5% gap in the other direction.

The 31% site was a performance-ratio-linked availability site. That is the worst possible place for the discrepancy to land. The nameplate sits in the denominator of the performance calculation, so if the construct reads the platform's figure, a 31% gap flows proportionally into a contractual number.

5.2 Why nameplate discrepancies are structural rather than careless

The same site has several legitimately different capacity figures, and they are not interchangeable:

Figure	Is	Used by
DC nameplate	Sum of module ratings.	Performance ratio, per-MW pricing
AC nameplate	Sum of inverter ratings.	Capability availability
Point-of-interconnection limit	Contractual export cap.	Curtailment, plausibility
Contracted capacity	Whatever the agreement names.	The guarantee

The first two are standard industry definitions. NREL's Annual Technology Baseline rates a PV system either in MWdc, aggregating all modules' rated capacities, or in MWac, aggregating all inverters' rated capacities, and calls the ratio between the two the inverter loading ratio. A plant built with a DC-to-AC ratio above one will always show two honest nameplates that disagree.

Hybrid sites make it worse. Where an agreement prices a generation and storage site as a single line, the operational platform frequently holds only the generation half. The two then disagree by exactly the storage nameplate, and the difference looks like an error rather than a scope mismatch.

The design rule

Four named capacity fields, never one. Each carries its own provenance and its own source document reference. Any calculation naming "nameplate" without qualifying which one is refused at review.

And **a site with no capacity reference is a gap on a list, not an exemption.** On the assessed fleet, a site with no nameplate row had a corrupt meter contributing 39% of the fleet rollup with nothing to check it against. That is the plausibility void of EC-WP-1100, arriving through the availability reference table.

6 Expected output: why peer referencing beats a curve

Production-based availability needs a counterfactual: what would have been produced but for unavailability. How that counterfactual is derived determines how the metric behaves when instruments fail.

6.1 The two methods

Method	Derives expected output from	Fails by
Curve-based	A modeled response to measured irradiance and temperature.	Consuming a bad irradiance reading and producing confident, precise, wrong numbers.
Peer-referenced	Comparison of like assets against each other over the same interval.	Degrading in precision as the peer set shrinks, visibly if coverage is published.

Curve-based is more elegant and, with healthy instruments, more precise. On paper it wins.

Then come eleven days in which the measured-irradiance feed reports zero at every site while rows continue to arrive on cadence. A curve method fed a zero does not fail. It *succeeds*. It publishes eleven days of confident expected-power figures and every derate classification underneath them.

On the assessed fleet the same defect surfaced in completed-day form: site-days where measured insolation was recorded as zero against real generation, the clearest being **1,059.6 MWh generated against zero reported sunlight.**

A method that gets less precise when the world gets uncertain beats one that stays precise and stops being true.

Calculation 4 · Peer-referenced expected output, with coverage

```
-- Expected output for a unit = median normalized output of its peers. Coverage is published, not assumed.
WITH peers AS ( SELECT p.interval_ts, p.site_id, p.unit_id, p.output_kw / NULLIF(p.rated_kw,0) AS normalized FROM unit_output p WHERE p.state = 'generating' ),
ref AS ( SELECT interval_ts, site_id, PERCENTILE_CONT(0.5) WITHIN GROUP (ORDER BY normalized) AS peer_median, COUNT(*) AS peer_n FROM peers GROUP BY interval_ts, site_id )
SELECT u.interval_ts, u.site_id, u.unit_id, u.output_kw,
ROUND(r.peer_median * u.rated_kw, 1) AS expected_kw, r.peer_n AS peers_in_reference,
-- refuse the comparison on a thin reference set
CASE WHEN r.peer_n < 5 THEN NULL ELSE ROUND(100.0 * u.output_kw / NULLIF(r.peer_median * u.rated_kw,0), 1) END AS pct_of_expected,
CASE WHEN r.peer_n < 5 THEN 'withheld: peer reference below 5 units' END AS withheld_reason
FROM unit_output u JOIN ref r USING (interval_ts, site_id);
```

Two properties make that query defensible. **The peer count is published on every row**, so a reader can see when the reference thinned. And **the comparison is withheld below a floor** rather than computed against two surviving units. That is the denominator-drift failure of EC-WP-1100, caught at the point where it would otherwise enter.

6.2 Where curve methods remain correct

Peer referencing needs peers. A single-inverter site, or a fleet where all units share one failure mode, has none. The honest position is not that one method is superior. It is that **the method must be stated with the number**, and where a curve method is used, its irradiance input needs an independent plausibility check ahead of it: the completed-day reconciliation of generation against reported insolation, run daily.

7 What an availability figure must carry

Seven fields. A figure missing any of them is exposed in a dispute, whatever its arithmetic.

7.1 The seven fields

#	Field	Because
1	**Construct**	Which of the four, named. The same plant produces different correct answers under different constructs.
2	**Source of construct**	Contractual (with the exhibit reference) or internal. Never ambiguous.
3	**Denominator**	The population and its size, on the face of the figure. Not in a footnote.
4	**Exclusions applied**	Which, for how long, and the evidence for each.
5	**Coverage**	Intervals with data against intervals in the period.
6	**Computability state**	Ready, partial, blocked or undefined, with the blocker named where not ready.
7	**Method for expected output**	Where the construct is production-based. Peer-referenced or curve, with peer count or irradiance source.

> **The reproduction test**
>
> Hand your monthly availability figure to someone with access to your raw data and the executed agreement, and ask them to reproduce it. Every question they have to ask you is a field missing from the list above. In a dispute, the person asking will not be someone you can answer informally.

7.2 Reporting the internal and contractual figures side by side

Where an internal method is carried alongside the contractual one, which is reasonable and often operationally necessary, publish both, labeled, with the delta computed and explained.

What is never acceptable is one column labeled "availability" whose method varies by site depending on what could be computed. That column will be averaged, and the resulting fleet figure is an average of quantities measured different ways.

8 Ten questions for your own portfolio

#	Question	What a bad answer means
1	Can anyone compute the availability guarantee for every site, from the executed agreement, today?	If the answer is "our platform reports availability," that is a different number.
2	How many distinct constructs exist across your portfolio?	More than one, and a fleet average is an average of different quantities.
3	Is every coefficient your contractual equations need actually populated?	"To be mutually agreed" that was never agreed leaves neither party able to compute the guarantee. Whether it can be enforced is a question for counsel.
4	Does any exhibit reference an attachment absent from your executed copy?	That site's construct is undefined, not merely blocked.
5	Has any site's denominator changed without a commissioning event?	Exact ratios mean a population swap. Run the query in Chapter 4.
6	Does one asset table hold more than one population?	If a flag distinguishes them, verify every consumer applies it.
7	How many capacity fields do you hold per site, and are they named?	One field called "nameplate" is being used for four incompatible purposes.
8	Is a guarantee comparison ever published for a site that cannot compute?	A variance against a figure that does not exist is an artifact of a default.
9	Can you evidence every exclusion you claim?	An exclusion the data cannot support fails when the other party inspects it.
10	When an availability figure was wrong historically, was it corrected or withheld?	Recomputing a ratio whose numerator was measured over a different population fabricates a record.

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.

What follows is what the RenewOps contract layer does with an availability figure it cannot compute. On the fleet measured, that was about a dozen of the roughly two dozen site trees the upstream source holds.

9.1 Computability is the output, not availability

The contract layer returns, per site, one of the four states set out in 2.1, with the specific blocker named.

On the fleet measured, **about half of the site trees held by the upstream source read blocked**. The panel names the blocker and routes the site to counsel. It does not fill a missing coefficient with a market convention, and it does not average around it.

> **Why a position beats a score**
>
> A score compresses "I computed this and it is 97.2" and "I could not compute this so I assumed the coefficient" into the same cell. A position cannot. Blocked is not a bad number. It is the absence of a number, attributable to a specific line in a specific agreement, and it routes to counsel rather than to operations.

9.2 No rollup was invented

The upstream source returns availability **per asset**. Whether a given agreement averages by asset count or weights by nameplate is a term of that agreement, so RenewOps does not choose. The site figure publishes as an **unweighted mean, named as unweighted**, with the per-asset minimum and the count below threshold beside it. A reader who needs the contractual weighting has everything required to apply it and is told plainly that it has not been applied.

9.3 Three figures, side by side, with the gap labeled

The platform's own inverter-count availability is published **beside** the upstream contractual figure and the upstream technical figure. One site on one day reads 66.0 percent, 84.76 percent and 51.45 percent.

None of the three is wrong at its own source. They answer three different questions, and the screen's job is to say which is which, not to pick a winner and hide the others. The platform's own figure carries a standing disclaimer stating it is not the contractual figure.

The column that is used in nothing. The technical availability figure is ingested, stored, and exposed, and no calculation anywhere in the platform consumes it, because its behavior contradicts its name. It reads 0.00 percent at night and 11.58 percent at midday against a contractual 98.47 percent on the same asset and day. Until somebody with authority defines what it measures, it is held rather than used.

9.4 Twenty rows withheld, and not estimated

The inverter denominator was corrected fleet-wide by counting only rows carrying the inverter flag, for the reasons set out in 4.2. The twenty historical rows computed against the wrong denominator are withheld with the reason on the row, as 4.3 explains.

9.5 A figure pulled early is not a settled figure

Contractual availability is **not final on the day.** A manual check, not a platform calculation, gives the tell: where the contractual and technical figures agree to three decimal places, no exclusions have yet been applied. A "yesterday" figure pulled too early is therefore a pre-exclusion number wearing a post-exclusion label. The behavior is flagged and under observation, and the collection window is likely to move back a day.

The ingest itself is idempotent on its natural keys and accepts a date argument for backfill. Two independent pulls reproduced the same six site figures exactly.

9.6 The dates behind the guarantee

A contractual calendar carries 490 obligations (287 hard dates and 203 recurring intervals) alongside the regulatory calendar's 209. Two behaviors in it belong to this reference.

- **Provisional dates are badged, not hidden.** 84 of the 287 dated rows rest on an unconfirmed anchor. They load with their computed dates and a badge reading *date provisional · anchor not confirmed*, with the anchor text printed beneath the row. Nine sites have no usable anchor at all, including two agreements executed with the effective date left blank. That is a question for counsel, not a data problem.
- **A cadence is not a date.** Obligations that recur without a fixed day are loaded as recurring *windows* carrying the decision still owed and a confidence rating. They are not turned into deadlines nobody agreed to.

The temptation with an unpopulated coefficient is to use the market convention, because the market convention is usually right. It is also not what anybody signed, and the first time that distinction matters will be in a dispute.

9.7 What the contract layer does not yet do

No contract or guarantee data exists inside RenewOps as structured terms. The exposure figures in this reference came from **reading executed agreements by hand.** An obligations matrix that holds the terms as data is identified as the next step and is not built. Until it is, a person reading an exhibit produces each computability position, and the platform holds the result rather than the reasoning.

The upstream source and the platform's fleet each hold roughly two dozen sites, but the counts differ. The two lists are not the same, and nobody has reconciled the overlap yet. We are stating that before anything is ingested fleet-wide, rather than discovering it afterward in a number that does not add up.

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
AC nameplate	The sum of a PV plant's inverter ratings, expressed in MWac.
Availability construct	The specific definition of availability a site's executed agreement requires: its denominator, its exclusions and its clock. This reference groups constructs into four families: time-based, production-based, capability, and performance-ratio-linked.
Blocked	A computability state. The construct is readable but cannot be evaluated by anybody, because a required term was never populated or the equation is absent from the executed text.
Capability availability	Available capacity divided by installed capacity at each interval, summed over installed units rather than healthy ones.
Computability	Whether an availability figure can be produced from the terms as written and the data as held. Reported per site as ready, partial, blocked or undefined.
Contracted capacity	The capacity figure an agreement names for its guarantee. It may differ from both nameplates and from the interconnection limit.
Crosswalk table	A table mapping assets in an upstream source to the platform's asset records.
DC nameplate	The sum of a PV plant's module ratings, expressed in MWdc.
Denominator drift	A failure class from EC-WP-1100: a derived metric whose denominator moves for reasons other than a physical change in the installed population.
Exclusion	An event an agreement removes from the availability calculation, such as force majeure, grid outage, offtaker curtailment or scheduled maintenance in an agreed window.
Inverter loading ratio	The ratio of a PV system's DC capacity to its AC capacity.
Partial	A computability state. The construct is readable but one or more inputs are incomplete; the missing input is named.
Peer-referenced expected output	Expected output for a unit derived from the median normalized output of comparable units over the same interval, published with the peer count.
Performance ratio	Measured production divided by modeled production. Definitions differ in detail between studies and agreements, so the reference model must be named.
Plausibility void	A failure class from EC-WP-1100: no bound exists against which a value could be checked, so an impossible value passes through as if it were real.
Point-of-interconnection limit	The contractual cap on what a site may export at its interconnection point.

Term	Meaning in this paper
Production-based availability	Actual energy divided by actual energy plus the energy lost to unavailability. Its counterfactual decides the result.
Ready	A computability state. The construct is readable, every input is held, and the figure can be computed.
Time-based availability	The fraction of a period during which the asset was capable of operating, after the agreed exclusions.
Undefined	A computability state. No construct can be established from the executed document, for example because an exhibit refers to an attachment that is not attached.
Withheld	A figure deliberately returned blank, with the reason carried on the row, because it cannot be computed defensibly.

How Energy Compliance Can Help

This paper asks one question of every availability figure you report: can anybody reproduce it from the executed agreement and the data you hold? Here is where we come in on that question.

Fleet Exposure Read

Your availability constructs reconciled against your executed agreements. It returns a per-site computability position (ready, partial, blocked, or undefined) and names the specific blocker for each site that cannot compute. It is the Chapter 2 and Chapter 3 work, done for your portfolio by people who read the exhibits in full.

Monitoring Truth Assessment

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.** Where we cannot verify something, it comes back marked deferred with the specific question, not softened into a maybe. Denominator stability, population mixing and nameplate references from Chapters 4 and 5 are part of the battery. Nobody buys a reconciliation layer before they know they need one, and the only honest way to find out is to look.

RenewOps

RenewOps is our renewable operations monitoring and analytics platform, and it includes a Data Center module. Chapter 9 sets out what it does with availability today and, just as plainly, what it does not yet do.

Compliance program support and WattOps Compliance

When a figure from your monitoring stack ends up in a compliance record, a senior on our team picks it up: gap and risk assessments, audit preparation and defense, and Data Request management that maps every request to a deliverable before anything goes out.

The evidence behind those records lives in WattOps Compliance, our NERC compliance evidence platform (wattopscompliance.com). If you want to see how it handles your evidence before you commit to anything, we will run one standard at one facility through WattOps at no cost. CIP standards are excluded from that offer.

Start with the ten questions in Chapter 8. If you cannot answer the first one for every site today, call me.

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.

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

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. 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. RenewOps, the reconciliation layer described throughout this series, is one part of it. WattOps Compliance, our NERC compliance evidence platform, is another.

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

Sources and Authorities

The findings in this reference are Energy Compliance's own measurements and have no external primary source. The public sources below support the general definitions in Chapters 1 and 5 and the Glossary. All were accessed October 17, 2026.

Availability, performance and capacity definitions

- Andy Walker and Jal Desai, National Renewable Energy Laboratory, *Understanding Solar Photovoltaic System Performance: An Assessment of 75 Federal Photovoltaic Systems*, DOE/GO-102021-5627, U.S. Department of Energy, Federal Energy Management Program, December 2021. The Chapter 1 definitions. https://www.energy.gov/sites/default/files/2022-01/understanding-solar-photovoltaic-system-performance.pdf
- International Electrotechnical Commission, IEC TS 63019:2019, *Photovoltaic power systems (PVPS): Information model for availability*, edition 1.0, May 9, 2019. https://webstore.iec.ch/en/publication/27253
- International Electrotechnical Commission, IEC 61724-1:2021, *Photovoltaic system performance, Part 1: Monitoring*, edition 2.0, July 21, 2021. https://webstore.iec.ch/en/publication/65561
- National Renewable Energy Laboratory, "PV System 'Availability' as a Reliability Metric: Improving Standards, Contract Language and Performance Models," NREL/CP-7A40-67949, 2016 IEEE 43rd Photovoltaic Specialists Conference, June 5–10, 2016, Portland, Oregon. Records that the meaning of PV availability has long been unclear, with contract language one of three areas being clarified. https://www.osti.gov/biblio/1343680
- National Renewable Energy Laboratory, Annual Technology Baseline 2024, "Solar PV AC-DC Translation." https://atb.nrel.gov/electricity/2024/pv-ac-dc
- Kevin S. Anderson and others, National Renewable Energy Laboratory, *The Effect of Inverter Loading Ratio on Energy Estimate Bias*, NREL/CP-5K00-82812, July 2022. https://www.nrel.gov/docs/fy22osti/82812.pdf
- Energy Compliance, Inc., EC-WP-1100, *Missing Is Not Zero*, companion reference for the denominator drift and plausibility void failure classes and the eleven-day irradiance zero-fill and 39% meter cases cited in Chapters 5 and 6.

© Energy Compliance, Inc. · EC-WP-1105 · Status as of October 17, 2026