DIPLO SPACE · FIELD NOTE 02

The Grid Has a Heartbeat

Why the electricity grid keeps a clock — and what happens to that clock when the satellite signal it trusts goes quiet

Applied-research field note · Edition 2.0 · Diplo Space, Inc. (Virginia 501(c)(3), EIN 99-1402470) · Every claim traced to a public primary source · The lab takes no policy position

By Brandon N. Eickhoff · CEO, Diplo Space · howdy@diplospace.org

The short version. Most people know GPS as the thing that draws the blue dot on a map. Its quieter job — and, for the electricity grid, its more important one — is keeping time. To keep the lights on, the grid relies on machines all over the country agreeing on the exact time, to within about a millionth of a second (a microsecond). Most of them get that time from GPS. If the GPS signal is lost, each machine keeps time on its own internal clock — but only for a while. A grid built around a good atomic clock can coast for about a month; one built around a cheap clock, for about 23 seconds. That gap — a month versus 23 seconds — is the whole story, and it was decided when the equipment was bought, not when the signal was lost.

ABSTRACT Modern power grids are monitored by Phasor Measurement Units (PMUs) — devices at thousands of locations that sample the grid's electrical state and timestamp each sample to microsecond precision using GPS. This shared, precise time is what lets operators compare measurements across a wide area and see the grid as a single system. This field note explains, for a non-specialist governance audience, why that timing dependency matters and what a loss of GPS timing would mean. It derives — for readers who want the math — a closed-form estimate of "holdover": how long a PMU can keep usable time after losing GPS. The answer ranges from roughly a month for a rubidium clock to about 23 seconds for a low-cost oscillator, a span of six orders of magnitude fixed by a procurement decision. It then situates this timing dependency in a grid under unprecedented stress — surging demand, a shifting generation mix, and capacity prices at administrative caps — and argues that the timing question, like the GPS-interference question of Field Note 01, is ultimately a space-governance question.

How to read this note. Section 1 explains, in plain terms, what the grid's clock does and why one microsecond matters. Section 2 derives the holdover number and is walled off for readers who want the derivation — it can be skipped without losing the argument. Sections 3–6 place the timing dependency in the context of a stressed grid: a changing generation mix, rising demand, the limits of wide-area monitoring, and the market's own scarcity signal. Section 7 explains why this is a space-governance question, Section 8 poses questions for policymakers, and Appendix A defines a simple grid-stress index. A glossary precedes the references.

1 µs timing precision the grid needs	month → 23 s holdover: best vs cheapest clock	+69% NERC peak-risk capacity, one year	~800% PJM capacity-price jump, one year

1 What depends on the timestamp

The electricity grid is often described as the largest machine ever built. Keeping it stable requires knowing, moment to moment, the electrical state at points spread across a continent — and comparing those points as if they were measured at the same instant. That comparison is what Phasor Measurement Units make possible. A PMU samples voltage and current dozens of times a second and stamps each sample with a precise time, so that a measurement in one state and a measurement a thousand miles away can be placed on a single, common timeline. Deployed at thousands of locations and standardized under IEEE/IEC 60255-118-1, PMUs are the instruments that let operators watch the whole grid breathe.[1]

The “synchro” in synchrophasor is the crucial part, and it comes, in almost every case, from GPS. The free civilian GPS signal carries not just position but time, disciplined by the atomic clocks aboard the satellites.[2] That is why the grid keeps a clock, and why that clock is, at root, a space asset.

Why one microsecond? It is worth pausing on the tolerance, because it is easy to dismiss as excessive. The grid in North America runs at 60 cycles per second. To compare the *phase* of the electrical wave at two distant points — which is what reveals whether power is flowing smoothly or building toward instability — the two measurements must be aligned to a small fraction of one cycle. A microsecond is about two-hundredths of one degree of a 60-Hz cycle. That is the level of agreement that lets a continent’s worth of measurements be read as one coherent picture; lose it, and the picture blurs into noise. The precision is not gold-plating — it is the threshold at which wide-area situational awareness works at all.

> **The one-sentence takeaway.** The grid does not merely use GPS for convenience; wide-area situational awareness — the operator’s ability to see the whole system as one — is built on a shared microsecond-accurate clock that comes from space.

2 How the holdover number is derived (for the technically inclined)

This section derives the holdover estimate quoted above. A reader who accepts “a month versus 23 seconds” can skip to Section 3 without loss.

When a PMU loses its GPS reference, it keeps time using its own internal oscillator, whose frequency drifts slowly. Approximating that drift by its one-day Allan deviation σ_y (a standard stability measure), the accumulated phase error after a time *t* is:

$$\Delta\theta(t) = 2\pi f \, \sigma_y \, t$$

where *f* is the nominal system frequency (60 Hz). Wide-area monitoring is generally considered unreliable once the phase error crosses about half a degree — a deliberately conservative threshold:

$$\theta_{\text{thr}} = 0.5^\circ = 8.73 \times 10^{-3} \ \text{rad}$$

Setting the accumulated error equal to the threshold and solving for the time to mis-operation gives a compact closed form:

$$t_{\text{misop}} = \frac{\theta_{\text{thr}}}{2\pi f \, \sigma_y} = \frac{2.316 \times 10^{-5}}{\sigma_y} \ \text{s}$$

Two caveats keep this honest. First, the half-degree threshold is conservative; some applications tolerate more, which would lengthen holdover. Second, representing an oscillator's drift by a single Allan-deviation number is an approximation — real clocks age and respond to temperature — so these figures are order-of-magnitude estimates, not guarantees. With that understood, substituting the stability of the three oscillator classes found in deployed PMUs yields:

$$\sigma_y \sim 10^{-11}\ \text{(Rb)} \ \Rightarrow\ t_{\text{misop}} \approx 27\ \text{days}$$

$$\sigma_y \sim 10^{-9}\ \text{(OCXO)} \ \Rightarrow\ t_{\text{misop}} \approx 6.4\ \text{h}$$

$$\sigma_y \sim 10^{-6}\ \text{(TCXO)} \ \Rightarrow\ t_{\text{misop}} \approx 23\ \text{s}$$

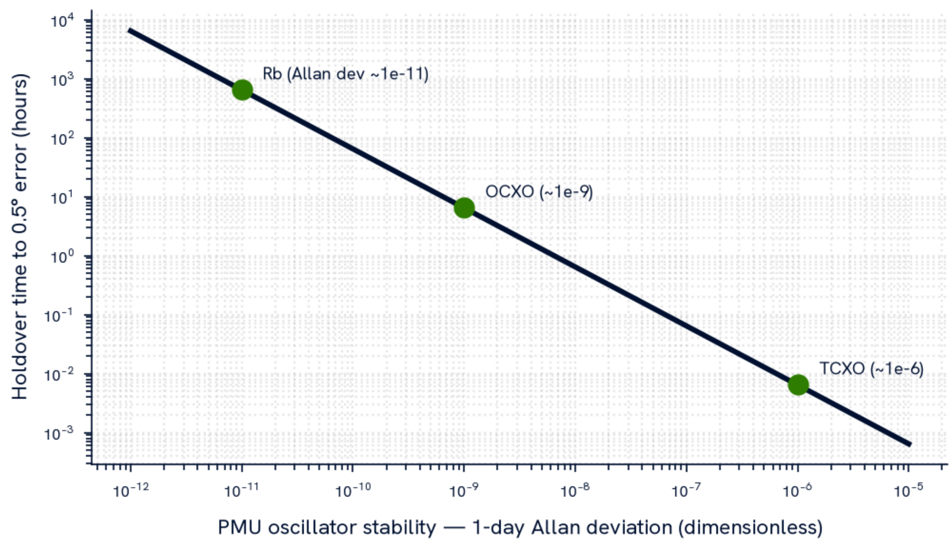

Figure 1. Holdover time to a half-degree phase error as a function of PMU oscillator stability (the closed-form result above). The vertical span is six orders of magnitude — from about a month for a rubidium clock to about 23 seconds for a low-cost oscillator — set entirely by which clock was installed.

The practical meaning is stark. A grid whose PMUs carry rubidium clocks can absorb a multi-week GPS outage with time to spare; a grid whose PMUs carry the cheapest oscillators loses wide-area visibility within seconds. The difference is not operational — it is a procurement decision made years earlier, usually without the resilience implications in view.

3 Why the timing matters more each year

The grid's timing dependency would matter even in a static system. But the system is not static — it is changing in ways that make precise, wide-area visibility more important, not less.

The generation mix is shifting from large spinning machines toward inverter-based resources (IBRs) — solar, wind, and battery storage that connect through power electronics. Conventional generators carry physical inertia: the momentum of their spinning mass resists sudden changes in frequency, buying the system precious moments during a disturbance. IBRs have little or no inherent inertia; they behave according to their control software. That makes the grid faster-acting and, in some respects, less forgiving — and it puts a premium on the fast, wide-area measurement that PMUs provide.

The consequences are visible in the disturbance record. In June 2022, a fault near Odessa, Texas, caused a large, near-simultaneous reduction in solar-PV output — the kind of correlated IBR response that conventional generation does not exhibit. NERC and ERCOT documented the event as part of a series of IBR-performance disturbances across 2021–2024, several in the Western Interconnection.[4,3] The point for a governance audience is not any single event but the pattern: as IBRs grow, the grid's failure modes are changing, and the tools for seeing those failures depend on the timing signal this note is about.

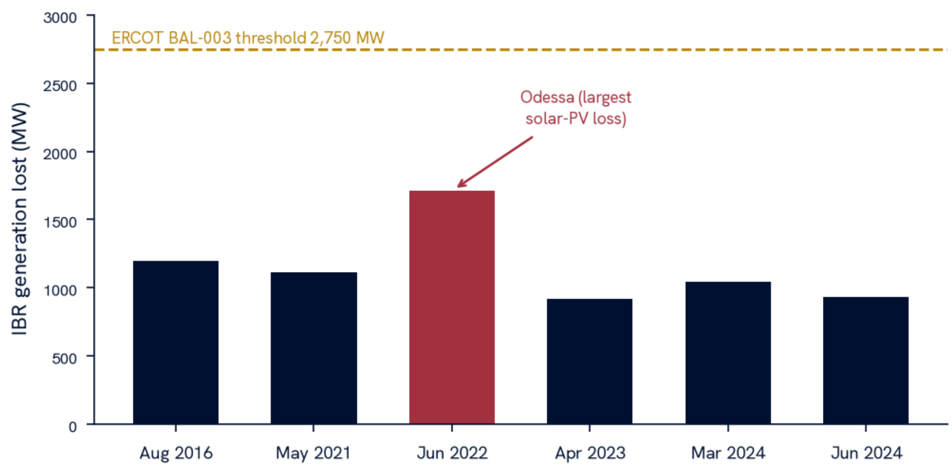

Figure 2. Inverter-based-resource generation lost in selected disturbances, 2016–2024, against ERCOT's BAL-003 frequency-response threshold. The June 2022 Odessa event was among the largest solar-PV losses on record. These events motivated new ride-through standards (IEEE 2800-2022).

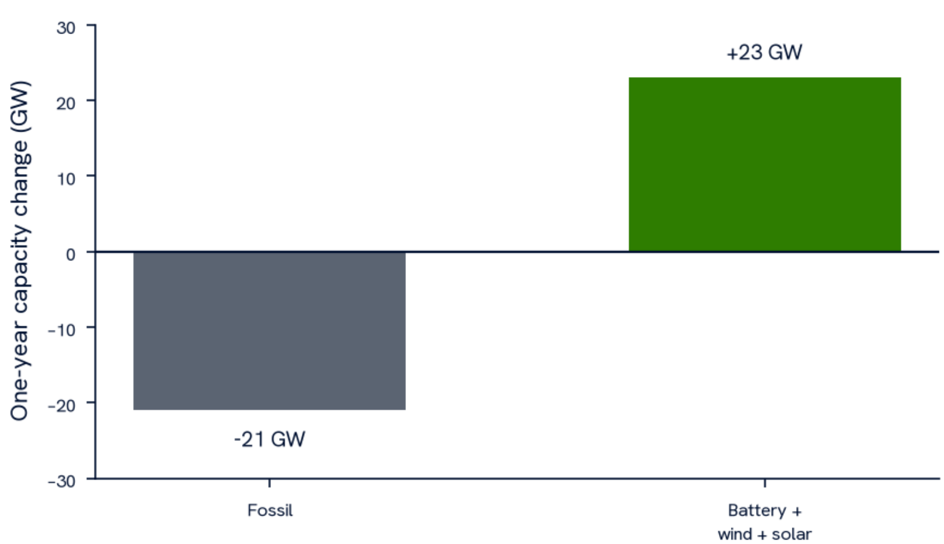

Figure 3. One-year change in the North American resource mix from NERC's 2025 Long-Term Reliability Assessment: conventional (fossil) capacity retiring while battery, wind, and solar capacity is added. The mix is shifting toward exactly the inverter-based resources whose disturbances PMUs are relied upon to observe.

The standards system is responding. IEEE 2800-2022 established interconnection and ride-through requirements for IBRs, and FERC Order No. 901 directed the development of reliability standards for them.[5] But standards apply going forward and depend on the deployed fleet actually meeting them — which, again, is something wide-area monitoring is needed to verify.

4 Rising demand raises the stakes

For two decades, U.S. electricity demand was roughly flat, and grid planning could assume a stable load. That assumption has collapsed. In its 2025 Long-Term Reliability Assessment, released in January 2026, NERC raised its ten-year summer peak-demand growth estimate to about 224 GW — up from roughly 132 GW in the prior year's assessment, an increase of about 69% in a single year, and the steepest such revision in decades.[6] Winter peak-growth estimates rose comparably. NERC identified a majority of assessed areas as facing elevated or high reliability risk over the horizon.

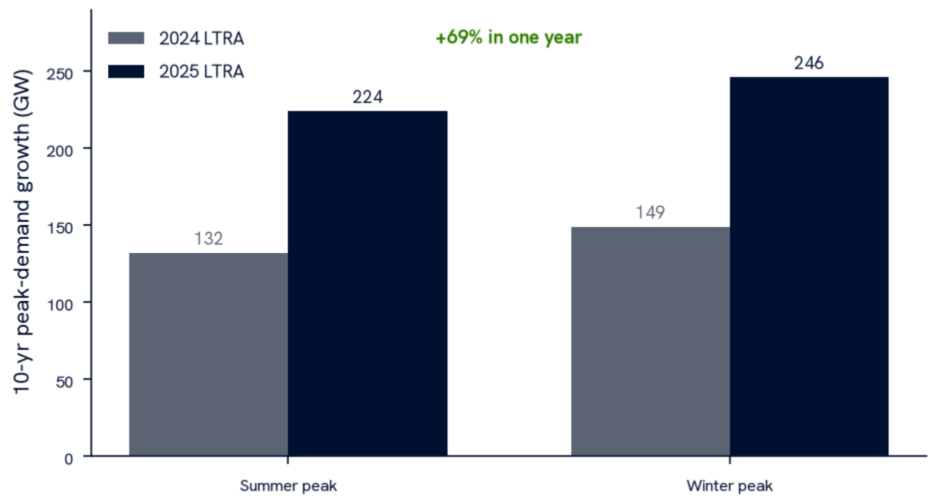

Figure 4. Ten-year peak-demand growth as projected by NERC's 2024 and 2025 Long-Term Reliability Assessments. The one-year jump — roughly 69% for the summer peak — reflects data-center load, electrification, and manufacturing, and is the steepest revision in decades.

It is worth stating the disagreement fairly. Some analysts — including Grid Strategies, in a March 2026 review prepared for several environmental organizations — argue that NERC's outlook is too pessimistic, cautioning that it may overstate data-center load and count interconnection requests that will not all materialize.[7] That is a disagreement about magnitude, not direction: even the more conservative view accepts that demand is rising after two flat decades. For the purpose of this note, the relevant fact is that a grid absorbing its fastest demand growth in a generation, while shifting toward inverter-based resources, is a grid that can less afford to lose its wide-area eyes.

5 Why wide-area monitoring cannot stop a fast cascade

A natural response to the foregoing is: if PMUs can see the whole grid, can they not also protect it? The answer — and it is the crux of the timing argument — is no, and the reason is a mismatch of time scales.

Local protection acts in a few cycles — on the order of tens of milliseconds — because it must clear a fault before it damages equipment or spreads. Wide-area monitoring operates a tier slower: PMU data is reported and concentrated over tens of milliseconds, and the wide-area applications that use it (state estimation, oscillation detection) run on time scales of seconds. Those are the right speeds for situational awareness, but they are too slow to interrupt a fast cascade, which can propagate across a region in well under a second.

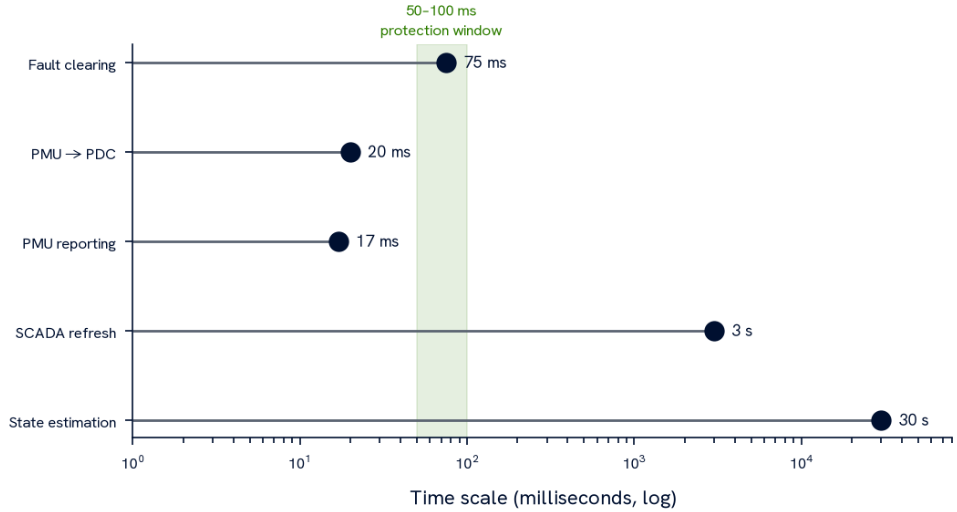

Figure 5. Time scales of grid protection and monitoring. Local fault clearing happens within the shaded protection window (tens of milliseconds); wide-area monitoring functions operate seconds later. The gap is why PMUs provide awareness, not real-time protection — and why losing timing degrades the ability to understand a cascade, even though it would not by itself cause one.

This distinction matters for how one reasons about a timing outage. Losing GPS timing does not, by itself, trip the grid; local protection keeps working. What it degrades is the operator's ability to *see* and *understand* the system — to detect the slow build toward instability, to run the wide-area applications that anticipate trouble, and to reconstruct an event afterward. In a grid that is faster-acting and more heavily loaded than before, losing that visibility at the wrong moment is a serious handicap, even though it is not a direct cause of failure.

6 The capacity market as an independent stress signal

If the engineering case for grid stress were speculative, the market would not be pricing it — but it is. PJM Interconnection, the largest U.S. grid operator, runs a forward auction (the Base Residual Auction) that procures capacity for a future delivery year; its clearing price is a direct market signal of expected scarcity. That signal has moved violently. The price rose from $28.92/MW-day for 2024/25 to $269.92 for 2025/26 — an increase of roughly 800% in a single year — then to $329.17 for 2026/27 and $333.44 for 2027/28, the last of which cleared at the administrative price cap.[8]

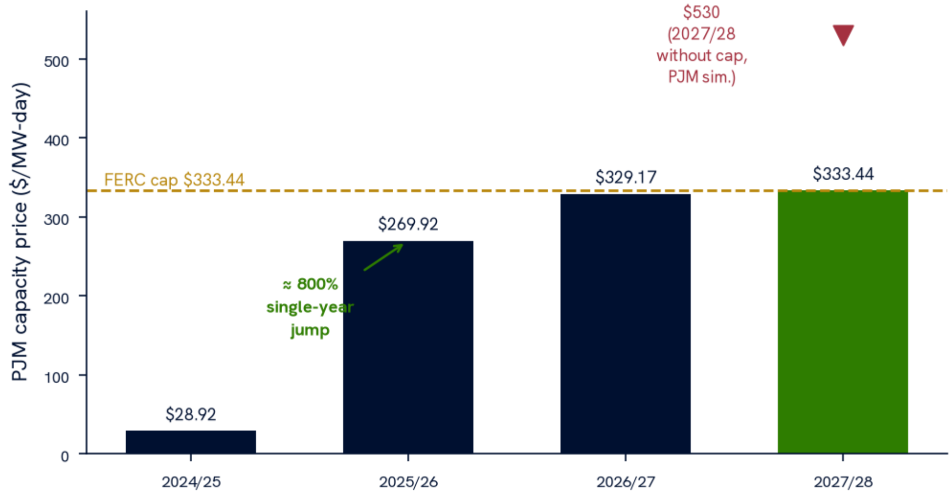

Figure 6. PJM Base Residual Auction clearing prices, 2024/25–2027/28. The roughly 800% single-year jump and the subsequent clearing at the FERC-approved cap are market confirmation of the tightening the reliability assessments describe. The dashed marker shows the price a PJM simulation estimated the 2027/28 auction would have reached without the cap.

The independence of this signal is what makes it valuable. Reliability assessments are engineering judgments; capacity prices are the aggregated bets of market participants with money at stake. When the two point the same way — and they do — the case that the grid is tightening rests on two legs, not one. That the most recent auction cleared at its cap, and that a PJM simulation estimated a materially higher uncapped price, indicates that the administrative ceiling, not abundant supply, is holding the number down.

7 Why this too is a space-governance question

It might seem that grid reliability is a purely domestic, terrestrial matter — a question for utility regulators and reliability organizations. But the thread running through this note leads back to space. The timing that makes wide-area monitoring possible comes from GPS; the vulnerability of that signal to interference is the subject of Field Note 01; and the interference in question crosses borders and implicates the international spectrum regime. A grid's timing resilience is therefore partly a function of a space asset and the governance around it.

The connection is likely to deepen. As complementary and alternative PNT sources mature — including timing from low-Earth-orbit constellations — the grid's clock may in the future draw on a wider set of space-based sources, each with its own governance. And the frontier extends further still: as activity on and around the Moon grows, the question of how to keep time across cislunar space is being taken up by the same community that governs terrestrial space activity. The humble grid timestamp thus sits at one end of a chain that runs all the way to the governance of timekeeping beyond Earth. That is why a note that begins with a power-grid clock ends, like its companion, as a space-governance question — and why a neutral layer able to translate across the domains is useful.

Four lenses on the grid's clock. The timing dependency is not only an engineering matter; it reads across every instrument of national power. **Diplomacy:** the clock comes from a space asset governed internationally, and as timing begins to draw on low-Earth-orbit constellations — and, further out, on cislunar timekeeping — the governance widens with it. **Information:** GPS time degrades from more than one cause, including natural space weather that perturbs the signal independent of any adversary. **Military:** power-system infrastructure — including the control and timing systems that keep it synchronized — becomes a target in conflict, as the sustained strikes on Ukraine's grid since 2022 have shown. **Economics:** the market is already pricing the stress, with capacity prices up roughly 800% in a single year and the most recent auction clearing at the administrative cap. A governance audience needs all four at once — the integrated view no single-instrument institution is built to hold.

8 Questions for policymakers

As in the companion note, this section closes with questions rather than recommendations — offered to structure deliberation, not to steer it.

Q1. Does your grid's wide-area monitoring depend on a single space-based timing source, and what is the holdover budget — hours or days — if that source is denied region-wide?

Q2. Which of your protection and monitoring systems assume precise time is always available, and would degrade silently rather than fail visibly if it were not?

Q3. Wide-area monitoring observes a fast cascade but cannot stop one. Where in your system does the gap between observation (tens of milliseconds to seconds) and protection (a few cycles) leave you exposed?

Q4. As inverter-based resources grow as a share of supply, is your disturbance-ride-through standard (e.g., IEEE 2800) keeping pace, and who tracks whether the deployed fleet actually meets it?

Q5. Capacity markets are now pricing scarcity at or near administrative caps. Is that price signal being read as an independent confirmation of the engineering stress your reliability planners describe — or discounted as a market artifact?

Q6. Demand projections have risen sharply (NERC's ten-year peak-growth estimate jumped roughly 69% in a single year). Which assumption — data-center load, electrification, or executed interconnections — is your planning most sensitive to, and how would you know if it is wrong?

Q7. Grid timing increasingly comes from, and could in future be complemented by, space-based sources beyond GPS. Is your resilience planning connected to the international PNT-governance conversation, or handled purely as a domestic reliability matter?

9 Conclusion

The grid keeps a clock, and that clock comes from space. How long the grid can keep good time without it is not a matter of chance but of a procurement choice — a month or 23 seconds — made when the monitoring equipment was bought. That timing dependency sits inside a grid under its greatest stress in a generation: a mix shifting toward inverter-based resources, demand rising at its fastest rate in decades, and a market pricing scarcity at its cap. Wide-area monitoring lets operators see this system, but only if the clock holds. Seeing the whole picture — the timing, the stress, and the space asset underneath both — is exactly the cross-domain translation this series exists to provide, and, as ever, we stop at the translation and leave the decisions to those who must make them.

Appendix A A simple grid-stress index

To make the several stress signals discussed above legible at a glance, one can combine them into a single illustrative index. The following weighted combination is offered only as a communication device, not as an operational metric:

$$\mathrm{GSI} = 0.25\,\frac{\Delta L}{L} + 0.50\,\mathrm{IBR}_{\mathrm{frac}} + 0.80\,(1 - C_{2800}) + 0.50\,\frac{N_{\mathrm{anom}}}{N_{\mathrm{base}}}$$

where $\Delta L/L$ is the relative load-growth revision, IBR_{frac} the inverter-based-resource share, C_{2800} the fraction of the fleet compliant with IEEE 2800, and N_{anom}/N_{base} the rate of disturbance anomalies relative to baseline. The weights are illustrative; the index is meant to show how independent signals can be read together, not to produce an authoritative score.

Glossary of key terms

PMU. Phasor Measurement Unit — a device that measures the voltage and current at a point on the grid many times a second and stamps each measurement with a precise time, so measurements from across a wide area can be compared on a common clock.

Synchrophasor. A time-stamped measurement of the grid's electrical state. The "synchro" refers to the shared, precise time — usually from GPS — that lets measurements taken hundreds of miles apart be aligned to the same instant.

Holdover. How long a clock can keep accurate time on its own after losing its GPS reference, before its error grows past a usable limit. Better internal oscillators hold over longer.

Allan deviation. A standard measure of how much a clock's frequency drifts over a given interval. A smaller value means a more stable clock and a longer holdover.

IBR. Inverter-Based Resource — a generator, such as solar or battery storage, that connects to the grid through power electronics rather than a spinning machine. IBRs behave differently from conventional generators during disturbances.

Ride-through. The ability of a generator to stay connected and keep operating through a brief grid disturbance rather than tripping offline. Standardized ride-through behavior (IEEE 2800) reduces the risk that one event cascades into many.

BRA. Base Residual Auction — PJM's forward auction that procures generating capacity for a future delivery year. Its clearing price is a market signal of how tight supply is expected to be.

LTRA. Long-Term Reliability Assessment — NERC's annual ten-year outlook on whether North America will have enough electricity supply to meet projected demand.

References & sources

1. IEEE/IEC 60255-118-1:2018, *Measuring relays and protection equipment — Synchrophasor for power systems* (defines PMU measurement and reporting requirements).
2. North American SynchroPhasor Initiative (NASPI) and U.S. Department of Energy technical resources on PMU timing, GPS dependence, and holdover performance.
3. ERCOT, Nodal Operating Guides and BAL-003 frequency-response obligations; and ERCOT event analyses of inverter-based-resource (IBR) disturbances.
4. NERC, *2022 Odessa Disturbance Report* and related IBR-performance disturbance reports (2021–2024 Western Interconnection events).
5. IEEE 1547 and IEEE 2800-2022, *Standard for Interconnection and Interoperability of Inverter-Based Resources (IBRs)*; and FERC Order No. 901.
6. NERC, *2025 Long-Term Reliability Assessment (LTRA)*, released January 2026 (ten-year demand-growth and resource-adequacy projections).

7. Grid Strategies LLC, review of NERC's 2025 Long-Term Reliability Assessment, prepared for Earthjustice, NRDC, Sierra Club, and EDF, March 2026.
8. PJM Interconnection, 2027/2028 Base Residual Auction results and reports, December 2025; and prior-delivery-year BRA results (2024/25–2026/27).
9. U.S. Department of Energy, emergency orders issued under Section 202(c) of the Federal Power Act, 2025–2026; and FERC statements on resource adequacy, 2026.
10. *Primary-source verification (July 2026): U.S. electricity-generation figures cross-checked against the U.S. Energy Information Administration (EIA) API. NERC 2025 LTRA demand and PJM 2027/28 Base Residual Auction figures are cited to those bodies' published reports (refs 6 and 8); they were not independently re-confirmed against primary records in this pass.*

About Diplo Space. Diplo Space is a neutral, nonprofit translation layer for space governance: we turn contested, technical, space-dependent systems into material a non-specialist official can navigate — every claim traced to a public primary source, and no policy position taken. This note applies that method to a dependency most grid discussions leave implicit: the bulk power system's reliance on the one-microsecond GPS timestamp, why that timing matters for reliability, and what the engineering and the markets are signaling. Its companion, *When GPS Lies*, sets out the broader positioning, navigation, and timing (PNT) picture.