WATTNESS

Position Paper
Response to RFC #1
August 2026

Wattness Response to Grid 2.0 RFC #1

An Open Standard for Grid Service Access and Allocation

Physics-Grounded Agents, Verifiable Intelligence,
and the Path from Grid 1 to Grid 2

Wattness Response to Grid 2.0 RFC #1

Physics-Grounded Agents, Verifiable Intelligence,
and the Path from Grid 1 to Grid 2

AUTHORS

Wattness Team

CONTACT

Email: hello@wattness.ai
X: @wattness_ai | Website: wattness.ai

Abstract

We welcome the Grid 2.0 Task Force's RFC #1 as a serious attempt to operationalise connect-and-manage principles for large loads, batteries and flexible resources. This response affirms the core problem statement, clarifies factual and incentive-design issues, situates Grid 2.0 within a longer lineage of prior art, and advances Wattness's complementary thesis: a hybrid architecture of deterministic, physics-grounded agents at the edge, selective cryptographic verification, and ISO-compatible automated strategies.

We argue that Grid 2.0's strongest path forward is as an open, voluntary operational overlay that remains fully compatible with existing Grid 1 requirements, while a layer such as Wattness supplies the verifiable intelligence that makes such an overlay economically attractive to asset owners and operationally acceptable to grid operators. Our single most substantive technical comment concerns incentive compatibility: allocation proportional to the declared request is manipulable, and RFC #1's explicit commitment to coordination without prices removes the mechanism that comparable systems use to keep declarations honest. We propose a price-free substitute and offer to test it.

On Method and Citation

We read RFC #1 in full and worked through it section by section; Appendix A preserves that pass in the document's own order so each comment can be attached to the passage that prompted it. We cite RFC #1 by section and quote only the operative phrase rather than reproducing passages at length. Every external factual claim carries a source in Appendix C, and where we could not verify something to our own satisfaction we say so rather than rounding it into confidence.

Every factual claim was checked against a primary or named secondary source.

1. Introduction And Scope

Grid 2.0 proposes a voluntary, permission-based, real-time protocol standard that decouples system access from traditional peak-reliability planning while automatically protecting the grid. We agree with the diagnosis.

> *the electricity industry must decouple system access from traditional reliability planning while automatically protecting the grid*
>
> RFC #1, Summary and Benefits

Reliability planning remains the fundamental safeguard, and the phrase *decouple* deserves care because read quickly it sounds like a proposal to plan less. What has failed to keep pace is not planning itself but the coupling between planning horizons and access decisions. System-access protocols of the kind G2TF proposes work best as an **enhancing overlay**: the enhancement carries the upside, the underlying layer carries the risk, and adoption stays rational for a participant even under uncertainty about whether the enhancement performs. That is the same structure a trader uses when overlaying candidate setups on a chart to decide where, when and how large to size a position — the overlay improves decisions without ever removing the underlying risk controls. It is also why Phase 0 Shadow Mode is the strongest idea in the deployment plan.

The phrase *automatically protecting the grid* carries two complementary meanings, and the distinction matters for where an organization like ours can contribute.

- **Passive.** A fail-safe default that never violates a Grid 1 constraint. RFC #1 specifies this well: on loss of context or peers, a Member Element reverts to its Grid 1 baseline within aggregate ramp limits.
- **Proactive.** Autonomous execution by edge agents whose decisions can be audited after the fact. Protection here derives not from trusting the agent but from the agent's decisions being reconstructible.

Wattness occupies the second meaning. Our agents are deterministic quantitative models grounded in physics — marginal cost, degradation, thermal state, co-optimisation across energy and ancillary products — and they are designed to remain fully compliant with QSE operations and control-room practice. Selective cryptographic proofs provide the audit rail without forcing a full on-chain market, and Section 8 sets out that architecture in detail.

We are a small team, we are for-profit, and we say so plainly wherever our commercial interest touches a comment. Under the RFC's own Focused Stakeholder Outreach heading we would fall into the category of researchers and company labs actively modelling these dynamics, at the small end of it.

2. Clarifications On Policy Context

2.1 Federal Action on Large Loads

RFC #1 refers both to FERC show cause orders and to a section 403 directive. Both references are correct, and the sequence connecting them is worth stating precisely because it is favourable to G2TF.

On **23 October 2025** the Secretary of Energy, acting under section 403 of the Department of Energy Organization Act (42 U.S.C. § 7173(a)), directed FERC to consider an Advance Notice of Proposed Rulemaking on the timely and orderly interconnection of large loads to the interstate transmission system. FERC opened **Docket No. RM26-4-000**, defining large loads for that purpose as demand greater than 20 MW, and committed to act by June 2026.

On **18 June 2026** FERC issued tailored show cause orders under section 206 of the Federal Power Act to all six RTOs and ISOs — PJM, MISO, SPP, CAISO, ISO-NE and NYISO — in dockets **EL26-67 through EL26-72**. Each operator and its transmission owners were given 30 days to file an informational report on how adequate generation will be available to serve large loads, and 60 days either to justify their existing tariffs or to file changes. Five categories of reform were named. One of them is *providing new transmission services for*

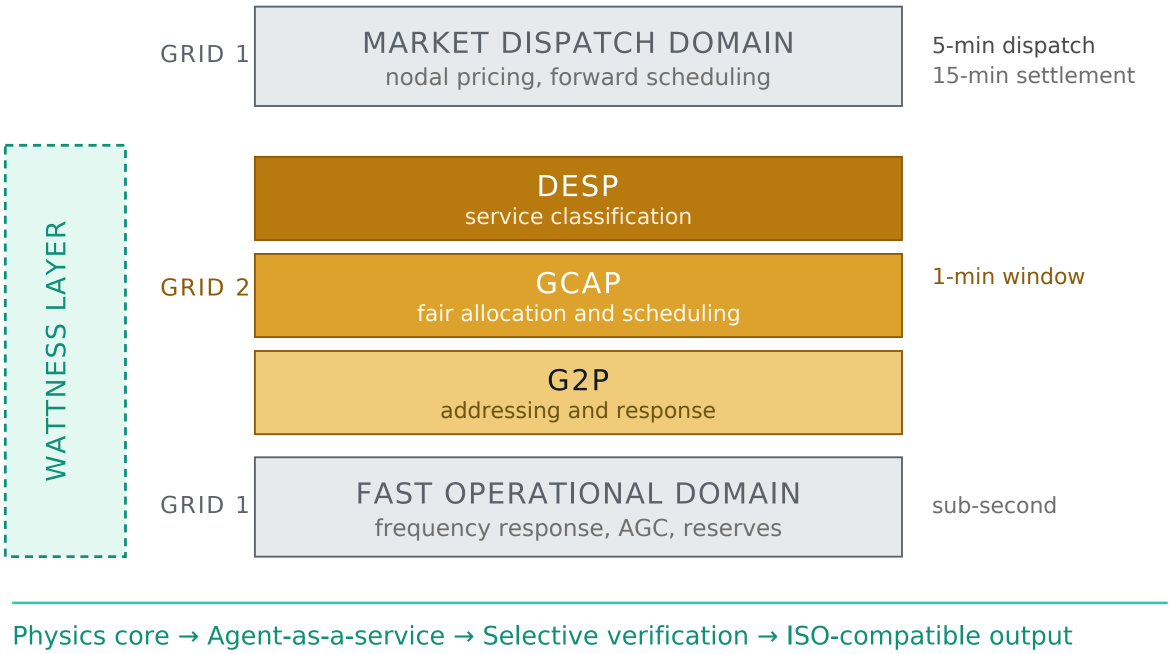

Figure 1. Where the two proposals meet. Grid 2 occupies the band between Grid 1's market dispatch domain and its fast operational domain. The Wattness layer spans that band, taking the access Grid 2 grants and turning it into a dispatch decision that an existing QSE can submit unchanged.

flexible large loads.

A separate section 403 letter dated 29 August 2025 concerns the draft Updated Certificate Policy Statement for interstate natural gas pipelines under the Natural Gas Act. It uses the same statutory instrument but is not about electricity large-load interconnection, and the two should not be conflated.

The conclusion we draw, and which we return to in Section 10, is that **the federal regulator has just instructed six system operators to create the service that Grid 2 is designed to operate**. That is not a competing approach. It is the demand signal, and we think it belongs near the front of RFC #2 rather than in a comparison table near the back.

2.2 Fragmented and Manual Mechanisms

> *but these options remain fragmented, somewhat manual*
>
> RFC #1, Summary and Benefits

We concur, and would sharpen the word *manual*. It does not mean that labour is involved; it means a human decision sits in the loop at a cadence slower than the phenomenon being managed. That is the real cost. PJM's Interim Resource Adequacy Service, SPP's Conditional High Impact Large Load Service and ERCOT's Provisional Controllable Load Resource each advance the frontier, and each still depends on processes that resolve in hours or days.

Our own approach is constrained by the same operational reality, deliberately. Deterministic models can be wired in ahead of the control room, but they must comply with QSE operating practice and leave the operator's authority intact. Any proposal requiring a control room to trust a model it cannot inspect will not be adopted, which is why the transition from manual to automated has to be shadowed, measured and proven before anything is submitted live.

3. What Grid 2.0 Inherits

RFC #1 positions itself against a set of approaches it describes as piecemeal and bespoke. We think the document would be stronger, not weaker, if it were more generous about its own lineage. Every one of Grid 2's core moves has a working precedent somewhere, and those precedents come with recorded failure modes a standards effort can learn from at no cost. Naming them also disarms the first objection a regulator will raise, which is that this has been tried.

Grid 2 Element	Closest Precedent	What the Precedent Teaches
Connect and manage	Great Britain adopted an enduring Connect and Manage transmission access regime in 2010–11, letting generators connect ahead of full network reinforcement with residual constraints managed operationally.	The term is borrowed and should be attributed. Ofgem's monitoring shows connection dates brought forward by years, and constraint costs rising materially as a direct consequence. RFC #2 should say who bears the analogous cost.
DESP classes A / B / C above firm	Kubernetes QoS classes (Guaranteed, Burstable, BestEffort), which set eviction order under node pressure. Loosely, DiffServ per-hop behaviours and the ATM service categories.	Class systems scale precisely because they are stateless, and they deliver relative priority rather than absolute guarantee. *Assured* risks promising more than a class mechanism can structurally deliver. Kubernetes also has one administrator and no strategic actors; a grid has thousands.
Per-member reservation	IntServ and RSVP, which reserved resources per flow and did not scale because every router held per-flow state. DiffServ displaced it.	The stateless choice Grid 2 has already made is the right one, and it is worth claiming explicitly since it is a real design decision with an evidence base.
Protocol stack over IP for grid devices	IEEE 2030.5, a RESTful application-layer protocol over HTTPS and TLS with an XML data model, mandated for smart-inverter DER interconnection under California Rule 21 through the CSIP profile.	The closest existing grid analogue to G2P, and already cited in the RFC. It is a device control and telemetry protocol, not a market bidding protocol — a distinction Section 10 returns to.
Flexible large-load access, live today	ERCOT's Provisional Controllable Load Resource under PGRR145 and NPRR1325, approved by the PUCT on 18 June 2026 and effective 10 July 2026. SPP's CHILLS non-firm service, accepted by FERC effective 1 July 2026.	PCLR splits a load into a firm floor and a dispatchable band and does much of what DESP describes, in production, in the market Wattness works in. Grid 2's marginal contribution should be stated relative to it.
Differential service during scarcity	PJM's May 2026 special report sets out a Differential Reliability path that rations scarce capacity by customer class, geography or contractual contribution.	An incumbent RTO independently proposing tiering in the same year is the strongest available evidence that DESP is directionally right.

Table 1. Prior art we believe RFC #2 should name. Sources in Appendix C.

None of this makes Grid 2.0 unoriginal. What we could not find anywhere is the specific combination it proposes: sub-minute allocation windows, explicit service classes above firm, and an IETF-style open standardisation process, applied to large loads on the bulk power system. Each ingredient exists separately. Saying that plainly is a stronger claim than novelty, because it is checkable.

4. The Central Technical Critique: Incentive Compatibility

> *Equitable dispatch is proposed as proportional to connection size and differential request size*
>
> RFC #1, The Grid 2.0 Proposal, Description

This is the sentence we would most like the Task Force to revisit, because everything downstream depends on it and because it is the one place where the internet analogy imports a property the grid does not have.

In a packet network, an endpoint that asks for more bandwidth than it can use is disciplined by the fact that unused capacity is simply not consumed; the bytes never arrive and the sender gains nothing. In an allocation mechanism where a member's share of scarce headroom rises with the size of its declared request, over-declaring is not wasteful. It is profitable. A member that declares four hundred megawatts of Class B demand it does not intend to draw is not punished by the protocol. It is rewarded with a larger proportional slice of every constrained interval, at the direct expense of members that declared honestly.

This is not a hypothetical failure mode. Proportional rationing is a textbook manipulable rule, and the theoretical result is unfriendly: for divisible resources, truthfulness is generally incompatible with pure proportionality, and even well-behaved mechanisms bound rather than eliminate the gain from misreporting. Figure 2 shows the arithmetic on a three-member domain.

The networks Grid 2 draws on all solve this, and none of them solve it with proportionality alone. DiffServ pairs its stateless classes with edge conditioners that police traffic against a profile and drop or re-mark whatever exceeds it. ATM's available bit rate service ran a closed feedback loop in which the network told sources what rate they were actually permitted. Telecom admission control simply refuses the call. Wholesale electricity markets use price. In every case there is a second mechanism whose job is to make the declaration mean something. RFC #1 currently has the classes without the conditioner.

A Price-Free Correction

The Task Force has committed to coordinating without prices, so we have tried to propose something that respects that constraint rather than arguing around it.

Let each member's allocation weight be its declared request multiplied by a **delivery score** derived from its own measured history: the ratio, over a trailing window, of what it actually delivered when curtailed to what it had declared it would deliver. Inflating a declaration then lowers the multiplier and the inflation cancels. No money changes hands, no price is published, and the Allocator still runs a proportional rule — it simply runs it over demonstrated capability rather than over stated ambition.

This is the DiffServ conditioner rendered as reputation, and it has three properties we think matter to the Task Force. It is stateless in the sense that matters, since the Allocator needs one scalar per member rather than a model of that member. It degrades gracefully, because a new member with no history starts at a neutral score.

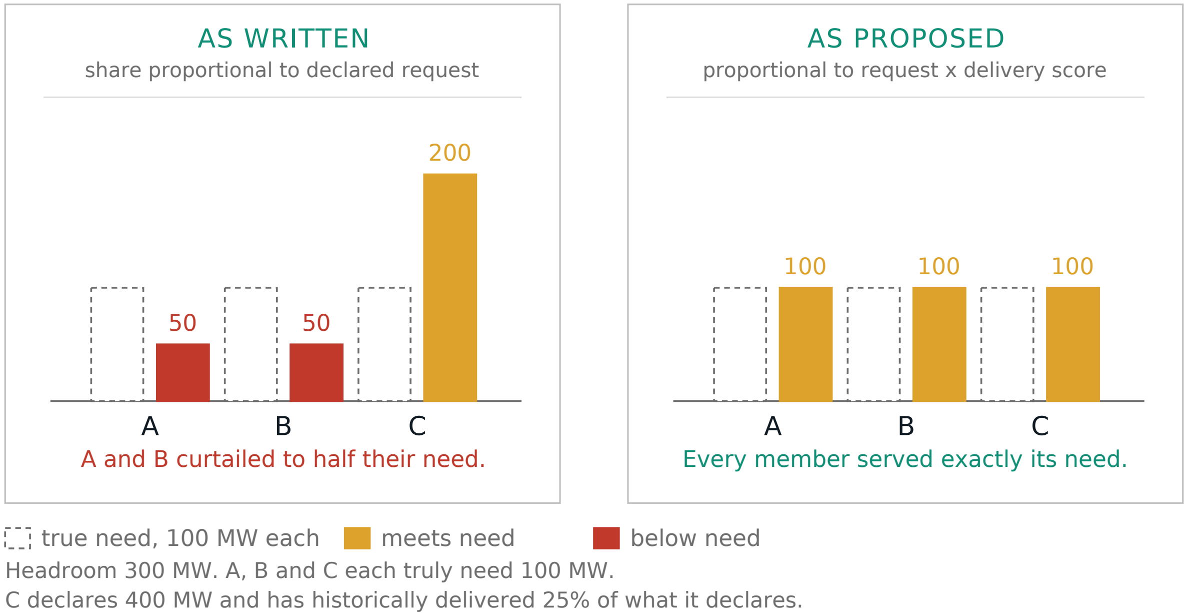

Figure 2. Three members, each genuinely needing 100 MW, competing for 300 MW of headroom. Member C declares 400 MW. Under the rule as written (left) the honest members are the ones curtailed and the inflating member is over-served. Weighting each declaration by the member's measured delivery history (right) cancels the inflation without introducing a price.

And it answers, with a number, an open question RFC #1 raises about itself — whether flexibility commitments will in fact materialise at scale — rather than deferring it to a later document.

It also has a cost we should name. It requires measurement that participants cannot unilaterally assert, which is one of the things we build. We are consequently proposing a mechanism that happens to need part of what we sell. The mechanism should be judged on whether it works rather than on who proposed it, and we would consider it a good outcome if someone else built the measurement.

Recommendation to G2TF. Specify the substitute mechanism — policing, penalties, admission control, or a delivery-weighted rule of the kind above — that GCAP will use to deter request inflation. Without one, the protocol's economic robustness remains incomplete in a way that a hostile reviewer will find before a friendly one does.

5. On Coordination Without Prices

> *Technical protocols handle real-time coordination; maintains current commercial settlements decoupled on slower clock*
>
> RFC #1, Key Design Principles and Characteristics

We understand the motivation. Price formation is where power market reform proposals go to die, and a protocol that avoids touching settlement can be adopted without a rate case. As a strategy for reaching Phase 1 this is astute.

As a design principle it needs more precision, because one phrase is covering two different claims. The weak claim is that Grid 2 does not itself clear a market or set a settlement price. That is true, uncontroversial and worth saying. The strong claim is that price information plays no part in the allocation decision. That is the claim we do not think survives contact with the reference implementation.

The ERCOT Counterexample

ERCOT's Provisional Controllable Load Resource is the mechanism RFC #1 points to as an initial deployment target. A PCLR is studied at its full requested peak, receives a firm floor called Low Power Consumption and a ceiling called Maximum Power Consumption, and must follow SCED dispatch instructions for everything in between. This is DESP's structure almost exactly: a firm tier, and a flexible band served on a best-efforts basis.

The relevant detail is how the flexible band is rationed. NPRR1325 introduces **Adjusted Bid Caps**, which cap a PCLR's energy bid curve based on constraint shadow prices and that load's locational impact on the binding constraint. The rationing signal is derived from the price of the constraint and from network topology. ERCOT also states plainly that it does not guarantee the requested maximum can ever be served in real time.

So the closest thing to Grid 2 that actually exists rations flexible access using a price-derived, locational signal, and it does so because that is the signal which already encodes where the scarcity is. We are not arguing that Grid 2 must adopt nodal pricing. We are arguing that a standard proposing to be the operational layer for PCLR cannot describe itself as price-free without explaining what it does with the shadow prices PCLR already uses — and that treating all members in a zone as fungible, when SCED does not, will produce allocations the underlying market then has to undo.

Two Specific Requests

1. State whether Grid 2 means it does not *set* prices, or that it does not *read* them. The first is a scoping decision and is fine. The second is a design constraint with consequences that should be argued for.
2. If the answer is the second, explain how GCAP distinguishes between two members in the same zone with different locational impact on the binding constraint. A proportional rule that ignores shift factors will allocate headroom to the member whose consumption most aggravates the constraint.

On Scale Without a Singular Operator

A smaller point in the same table. The stated objective is scale without a singular operator, and the rationale is that every asset runs the same algorithm against the same observable signals so that collective behaviour emerges. We would put it differently: what Grid 2 achieves is scale without a single point of *control*, which is a different and more defensible claim.

Operators do not disappear. Every Member Element is operated by someone with a commercial position, an asset warranty and a view about which hours matter. Emergent good behaviour from identical algorithms assumes identical objectives, and members will not have them. Naming the operators and their incentives makes the design more robust, not less elegant.

6. Scale, Scope And Physical State

> *It serves and sits inside the Grid 1 context at any scale or scope*

RFC #1, Summary and Benefits

This is the boldest sentence in the document and we would like to see it either defended or narrowed. The difficulty is not conceptual — an overlay can in principle nest at any level — but physical.

Proportional to Connection Size Is Not Physically Neutral

Consider two grid-scale batteries with identical nameplate connection capacity, both registered as Buffer members in the same Allocation Domain. One sits at 80% state of charge, 22 degrees ambient, early in its warranty. The other sits at 25% state of charge, has been cycling hard through an afternoon of price volatility, and is close to its thermal derate threshold. Under a rule proportional to connection size and requested amount these two assets receive the same commitment percentage. They cannot deliver the same energy, they cannot sustain it for the same duration, and the marginal cost to their owners of delivering it differs by an order of magnitude.

The same argument applies to loads. A data centre's deferrable fraction is not a constant; it is a function of what is currently running, and a training job at hour nine of a checkpoint interval is not as interruptible as the same job at hour one. RFC #1 handles this correctly in one place — it says the Use member decides at the edge how to map internal workloads onto service classes — and then discards it in another, by making the Allocator's arithmetic depend only on the size of what was asked for.

Portfolio composition compounds the problem. Ten batteries on one site and a thousand across a zone are not the same allocation object: irradiance, wind direction, distance to the point of interconnection and local latency all change what the aggregate can actually deliver in a given minute. A fleet of small reactors, or a population of electric vehicles behaving as distributed storage on wheels, changes it again.

Proposed: A Capability Descriptor in DESP

We suggest DESP carry, alongside the per-class request, a small and deliberately minimal capability descriptor: sustainable duration at the requested level, permitted ramp, and a single normalised headroom scalar that the member computes locally from whatever physics governs it — state of charge and thermal margin for storage, deferrable fraction for compute, process constraints for industry.

The Allocator does not need to know how the number was produced. It needs the number to exist, to be comparable across members, and to be checkable after the fact. This keeps intelligence at the edge, which is Grid 2's design commitment and ours. It also gives the delivery score of Section 4 something concrete to be scored against: a member that declares a two-hour sustainable duration and browns out after twenty minutes has made a checkable, falsifiable statement.

On Scope

Scope raises a separate question that RFC #1 leaves open. A Grid 2 Allocation Domain sits inside a jurisdiction, and jurisdictions differ in ways that bear on the protocol. ERCOT is contained within a single state and its own regulator, which is why Batch Zero moved from board approval to PUCT approval to effect within weeks. A domain spanning PJM answers to a multi-state governance process that is itself currently under federal scrutiny. A standard intending to span both should say which layer resolves a conflict between an Allocator's clearing decision and a state regulator's order, because that conflict will happen.

7. Temporal Design And Market Portability

> *below the typical five-minute power market clearing domains*
>
> RFC #1, Proposed Scope

The one-minute window is, in our view, technically achievable and well chosen. Our comment is about the clock it is measured against.

Five-minute market clearing conflates dispatch cadence with settlement cadence, and the distinction matters for a protocol whose whole value proposition is temporal nesting. In ERCOT, security-constrained economic dispatch produces locational prices at least every five minutes, but real-time energy settles on **fifteen-minute intervals**, with the settlement price formed as a time-weighted average of the SCED prices within the interval. A Grid 2 window is therefore nested inside the dispatch clock but is one-fifteenth of the interval on which money actually changes hands. Any argument that Grid 2 leaves commercial settlement undisturbed has to be made against the settlement interval, not the dispatch interval.

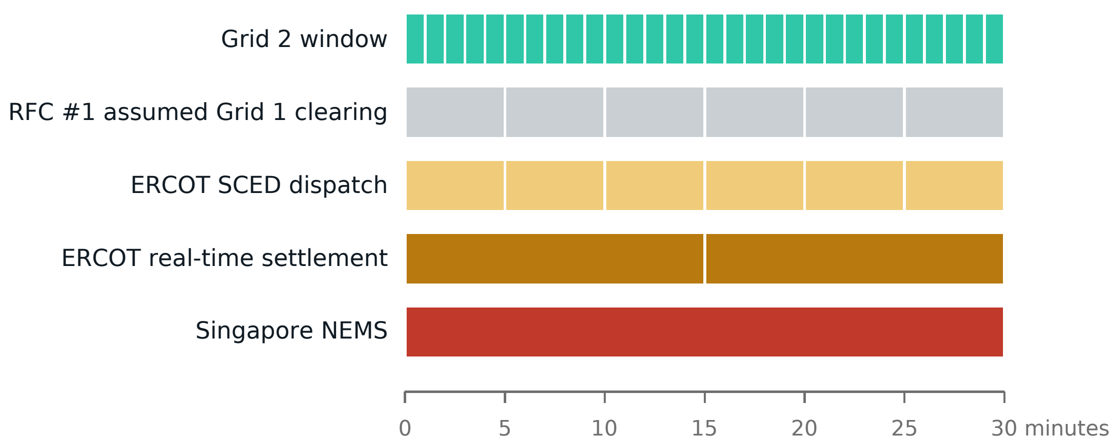

A one-minute window nests inside a five-minute dispatch clock, but not inside every settlement clock

Figure 3. Grid 2's proposed window against the clocks it must nest inside. The second row is the RFC's stated assumption. The third and fourth rows are the same market measured two different ways: ERCOT produces prices every five minutes but settles real-time energy on fifteen-minute intervals. The fifth row is a market where neither cadence appears.

On the Claim That Grid 2.0 Is Not Market-Specific

> *Grid 2.0 is not intended to be specific to any market or nation*
>
> RFC #1, Focused Stakeholder Outreach

We would move this sentence to the front of the document and then either evidence it or scope it. As it stands it appears on the penultimate page, after a paper whose analytical frame is US wholesale power throughout: FERC dockets, NERC BAL standards, RTO tariffs, five-minute clearing, and studies of the US interconnections.

Singapore, among others, is a concrete case. The National Electricity Market of Singapore dispatches and settles on half-hourly periods, forty-eight per trading day. The nesting argument that makes Grid 2 elegant in ERCOT — one minute inside five inside fifteen — has no counterpart there. That does not make Grid 2 inapplicable to Singapore. It means the temporal argument has to be re-derived for each market rather than transplanted, and a standard intending to travel should supply the template for doing so.

Concretely, we suggest RFC #2 define a short **market profile** — dispatch interval, settlement interval, gate closure, nodal or zonal pricing, and who holds curtailment authority — that any adopting region completes before claiming conformance. Five fields would do it, and it would convert a broad claim into a testable one.

8. The Wattness Thesis: Physics-Grounded Agents And Verifiable Intelligence

Everything above is criticism, which is the cheap half of a comment letter. This section is the part we would like to be held to.

8.1 The Gap We Think We Can Fill

RFC #1 states that in Member-to-Allocator interactions, members advertise their requests, Allocators clear them, and members are responsible for verified self-dispatch. In the Allocator row of the same table it states, admirably plainly, that the Allocator has no internal intelligence.

Read together, these two statements describe a protocol that allocates access and then stops. It tells a data centre or a battery how many megawatts it may draw or deliver in the next minute. It does not tell that asset **what to do with the allocation**, and it explicitly declines to price the decision. Yet the economics of participating in Grid 2 — whether a battery should discharge now or hold state of charge for a later scarcity interval, whether accepting a Class C allocation is worth the degradation it causes, whether declaring more Class A demand is worth the flexibility obligation attached — are exactly what determine whether an asset owner opts in at all.

That decision layer is what Wattness builds. It is not a competing allocation rule, and it does not require Grid 2 to change its architecture. It is the thing that makes a permission-based overlay economically attractive to the party being asked to volunteer.

8.2 Hybrid Architecture

Our architecture is deliberately hybrid, because the parts of this problem that need to be deterministic and the parts that benefit from cryptography are not the same parts.

Layer 01 — Deterministic physics core

Marginal-cost, degradation-aware, co-optimised offer curves and flexibility schedules. For storage this means state of charge, state of health and thermal headroom enter the optimisation directly rather than as after-the-fact constraints, so that the marginal cost of a discharge decision includes the capacity fade it causes. The models are deterministic and inspectable. This matters commercially as well as technically: a control room will not run what it cannot interrogate, and an asset owner will not accept degradation they cannot attribute.

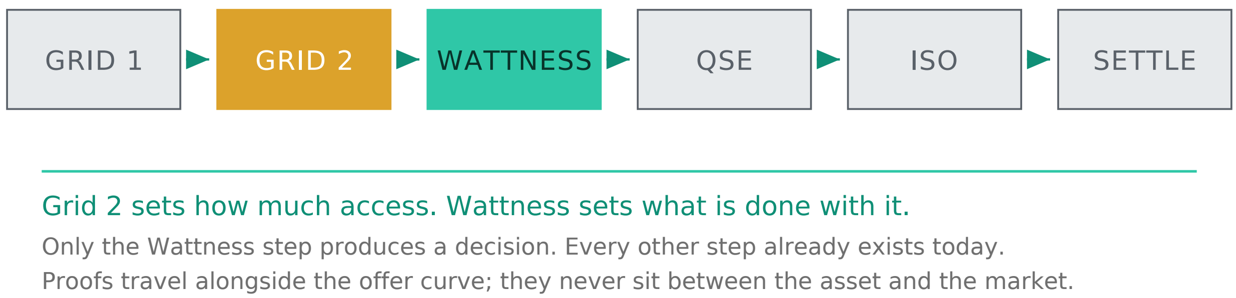

Figure 4. Conceptual flow. Grid 1 sets the firm baseline at interconnection; Grid 2 grants tiered permission classes on top of it; the Wattness agent turns that permission into an offer curve and a flexibility schedule; the QSE submits it through existing workflows; the ISO clears; settlement follows, with an optional secondary verification layer for bilateral flexibility or RECs.

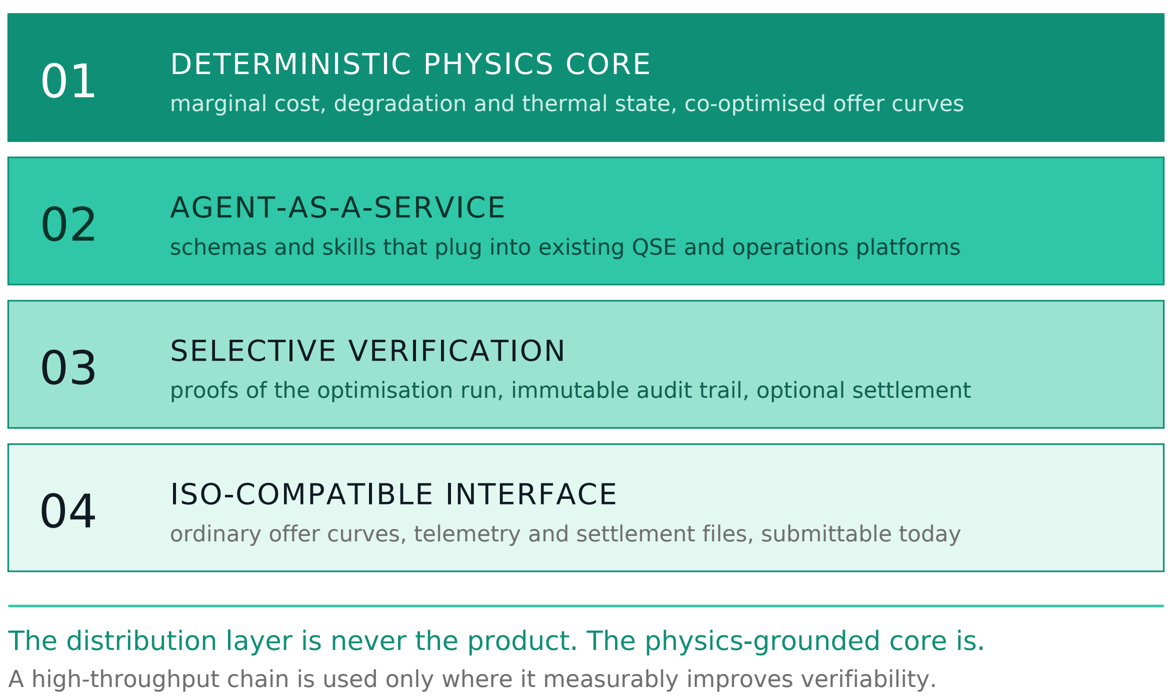

Figure 5. The Wattness stack. Layer 01 is the product; layers 02 to 04 exist so that layer 01 can reach an asset without asking anyone to change a tariff, a control-room procedure or a market protocol.

Layer 02 — Agent-as-a-service

Proprietary schemas and skills that plug into existing QSE and operations platforms. The customer keeps whatever front-end agent framework they already use; we are agnostic across the current field, and expect that field to keep changing. **The distribution layer is never the product. The physics-grounded core is.** We state this plainly because the opposite mistake — selling the wrapper and hoping the model follows — is common in this category right now.

Layer 03 — Selective verification

Zero-knowledge or trusted-execution proofs of the optimisation run, an immutable audit trail, and optional secondary settlement of bilateral flexibility or renewable energy certificates. The operative word is *selective*. A high-throughput chain is used only where speed and cost measurably improve verifiability, and never as a bidding engine. We are not proposing that anyone clear a power market on-chain, and we would regard a design that required it as a worse design.

The purpose of this layer is narrow and, we think, genuinely useful to the Task Force: it makes an edge decision reconstructible by someone who was not present when it was made. That is what turns a self-asserted delivery claim into a delivery score of the kind Section 4 proposes.

Layer 04 — ISO-compatible interface

Outputs are ordinary offer curves, telemetry and settlement files that any QSE can submit today under existing protocols. Nothing in the stack requires a tariff change, a protocol revision, or a regulator's permission to try. This is a deliberate constraint on ourselves, and it is the same constraint that makes Grid 2's voluntary framing sound.

8.3 Where This Meets Grid 2.0

Three concrete intersections, in ascending order of how much they involve us.

1. **The capability descriptor.** Section 6 proposes that DESP carry a minimal, locally computed statement of what a member can actually sustain. That number is an output of Layer 01 for any asset we run. We would be glad to help specify the field so it is comparable across asset classes rather than vendor-specific.
2. **The delivery score.** Section 4's price-free correction needs measurement that a participant cannot unilaterally assert. Layer 03 produces exactly that, and the scoring rule matters more than it looks. Electricity outcomes are heavy-tailed: a small number of scarcity intervals dominate everything, so a score averaged uniformly across the year ranks members on the eight thousand hours that do not matter. We use threshold-weighted continuous ranked probability score, weighted towards the scarcity region, which remains strictly proper — weighting the score towards the hours the grid cares about does not reintroduce a way to game it. We would write that up properly for the Task Force if it is wanted.
3. **Adversarial analysis of GCAP.** How much can a member gain by misdeclaring under the rule as written, under the delivery-weighted variant, and under a small number of plausible collusion patterns. We would run this properly and publish whatever it finds. We are aware this is the offer most likely to be unwelcome and most likely to be useful, and we would rather run it and share it than have it run later by someone with a reason to want Grid 2 to fail.

8.4 Sequence, and What We Have Not Built

Our path is deliberately incremental, and it starts where the risk is lowest.

1. Shadow analysis on public and partner data, producing verifiable uplift reports against outcomes that have already settled.
2. Pilots with one to three independent or developer-owned storage sites, through existing QSE partners rather than around them.
3. Productisation as agent-as-a-service.
4. Extension to fleets and aggregators under FERC Order 2222 pathways.
5. Direct QSE registration, optionally, once capital and operational maturity allow.

We think a paper about verifiability should be precise about its own status. The physics core and the scoring pipeline run today against historical and live ERCOT data. Nothing we have built sits inside anyone's dispatch path. No Grid 2 Allocator exists for us to integrate with, and we have not run any of this against a Grid 2 protocol stack because there is not yet one to run against. What we are offering the Task Force is a design contribution and an implementation we would build alongside a shadow deployment, not a finished product waiting for a customer.

9. Responses To The Task Force's Open Questions

RFC #1 lists its own open questions and invites the requisite technical studies. Below are those where we think we can say something useful, with our answers stated as positions rather than conclusions.

Open Question (RFC #1)	Wattness Response
How can we be assured that flexibility commitments will in fact materialise at scale? How should incentives be structured to deliver compliance?	This is the question our whole submission is aimed at. Measure it: score declared against delivered on every window, publish the score, and make the next window's allocation weight a function of it. Compliance then needs no penalty schedule, because enforcement is proportional to the miss and automatic.
How do we ensure coincidental load shedding does not cause system stability challenges?	Part of the risk is correlated physical state, not just correlated commands. Storage assets across a domain often share weather, ambient temperature and price signal, so their thermal derate thresholds cluster. With the capability descriptor of Section 6, an Allocator can see that clustering forming in the declarations before it becomes a simultaneous failure to deliver.
Is there sufficient headroom to absorb coincident swings?	We do not know, and we would not trust a general answer. This is precisely what Phase 0 with an evidentiary record can settle empirically for a given domain, which is a better outcome than settling it by argument.
Which existing signals from transmission owners are best to use for inputs, and how can these be standardised?	Start from what already binds. Constraint shadow prices and shift factors are produced continuously, are already used by ERCOT to cap PCLR bids, and encode locational impact that a size-proportional rule discards. Dynamic line ratings, which the RFC already names in the Allocator row, are the natural second input: conductor rating is weather-dependent, and a domain that observes it rather than assuming a static seasonal rating has materially more headroom to allocate in exactly the hours the constraint binds.

Open Question (RFC #1)	Wattness Response
Are there cybersecurity or NERC CIP concerns?	Yes, and one specifically. Grid 2 creates a new remote influence path over aggregate load, so a compromised Allocator is a systemic actor. Signed commitments help here for a reason worth stating precisely: they do not prevent a compromise, but they make one detectable and attributable after the fact, which is the property incident response actually needs. Separately, long-lived signing keys in field assets should be planned with post-quantum migration in mind given the service life of grid hardware.
What might commercial terms look like, particularly regarding enforcement of flexible capabilities?	We would rather not opine on tariff design. We will say that a published delivery score is a better foundation for a commercial term than an attestation, because it is a number both parties can audit and neither party produces alone.
If a new battery is needed to support new Grid 2 participating loads, how is the investment decision made?	On the same basis as any other storage investment, but with one input the current framing omits: the degradation cost of the duty cycle Grid 2 would impose. A buffer sized against energy and power alone will be mis-sized against throughput. This is squarely in Layer 01 of our stack and we would be glad to contribute the method.
Should there be guidance or limits on Grid 2 differential service classes?	Three classes is right, and the pressure will come from members whose internal workload structure does not map onto three. The useful guidance is not more classes but a documented mapping procedure, plus permission to decline: a member should be able to participate in G2P and GCAP while declaring that it cannot honestly classify a portion of its load.

Table 2. Responses to selected open questions from RFC #1.

10. Comments On Current Industry Approaches

This table is where we most often disagreed with the document, and on reflection our disagreement is with its framing rather than its content. The table evaluates each alternative on whether it delivers what Grid 2 delivers, which makes the comparison circular in places. Most of these entries are complements, and the paper would be more persuasive if it said so.

On the FERC directives row

We initially read this row as arguing against the regulator, which would be both unnecessary and unwise. Rereading it, the intended point is narrower and correct: a directive is not an operational protocol. But the row as written says a directive will not necessarily result in a unified interconnection approach, which reads as a criticism of the instrument rather than a description of the gap Grid 2 fills.

Section 2 sets out the sequence. The federal regulator has just told six system operators to create new transmission services for flexible large loads. That is the demand signal for this standard, and we would rewrite the row to say so.

On the energy markets row

The row rejects reliance on energy markets on the grounds that data centre operators can tolerate energy costs many multiples of a gas heat-rate benchmark, so price does not discipline them. The observation is right and the conclusion does not follow. That a hyperscaler is price-insensitive at the level of the energy bill does not mean it is insensitive to the locational price of the constraint it is aggravating, which is a different and much smaller number attached to a specific hour and a specific line. Discarding the whole price signal because one participant class is inelastic to one component of it gives up information the protocol could use. Section 5 is the concrete version of this objection, and we would ask why a win-win is not available here.

On the custom rate structures row

The row calls bespoke hyperscaler demand-response agreements piecemeal. They are, but they are also the strongest existing proof that the thesis works: Google announced demand-response agreements for machine-learning workloads with Indiana Michigan Power and the Tennessee Valley Authority in August 2025 and has since extended the approach to further utilities.

That these were achieved through bilateral contracts inside existing tariffs is the sharpest question facing Grid 2 — not whether flexible large load works, but whether a standard is needed to deliver what contracts already deliver. Our answer is that contracts do not compose and do not aggregate: ten bilateral agreements in a zone cannot be cleared against a single constraint, and the eleventh has to be negotiated from scratch. The case for a protocol should be made on *composability* rather than on capability. We think that argument is available to the Task Force and is not currently made.

On the off-grid loads row

The stated reason for deprioritising off-grid large loads is that many intend to interconnect eventually. We do not think that follows. A load that will interconnect later is a future member, and the years it spends off-grid are years in which it is running exactly the kind of self-contained generation-and-load balancing Grid 2 wants to standardise. The stronger version of the Task Force's own position is that these operators are an ideal Phase 0 population: they can run the full protocol stack in shadow against their own island, at no risk to anyone's grid, and arrive at interconnection with a scored operating history. We offer that reframing as a gift rather than a criticism.

On IEEE 2030.5 as a foundation (RFC #1, p. 6)

A clarification rather than a disagreement, because we had to look this up ourselves. IEEE 2030.5, also called Smart Energy Profile 2.0, is an application-layer protocol carried over HTTPS with TLS, using a RESTful resource model and an XML data model rather than JSON. It is the default protocol for smart-inverter communication under California Rule 21 through the CSIP profile.

Critically for the RFC's purposes, it is a device configuration, telemetry and control protocol for physical assets. It is not a market bidding protocol and has no notion of a day-ahead offer. That makes it a plausible foundation for G2P and an implausible one for anything touching settlement, which is worth stating explicitly since the RFC gestures at it as a foundation without specifying for which layer.

11. Recommendations To The Grid 2.0 Task Force

Summarising what we are asking for, in descending order of how much we think it matters and ascending order of how much it involves us.

1. **Specify the truthfulness mechanism.** Either add a term that makes over-declaration unprofitable — polic-

ing, penalties, admission control, or the delivery-weighted rule of Section 4 — or state explicitly that Grid 2 relies on membership agreements to police declarations, and accept that the protocol is then only as honest as its contracts.

2. **Disambiguate coordination without prices.** Say whether Grid 2 declines to set prices or declines to read them, and reconcile the answer with ERCOT's use of constraint shadow prices to cap PCLR bids.
3. **Name the prior art.** Great Britain's Connect and Manage, IEEE 2030.5, ERCOT PCLR, SPP CHILLS, PJM's differential reliability work, and the DiffServ and Kubernetes lineages. Grid 2's contribution is the combination, and that claim is stronger for being narrower.
4. **Clarify the relationship to live constructs.** State precisely how Grid 2 service classes map onto PCLR's Low and Maximum Power Consumption bands, onto CHILLS, and onto PJM's Path B. A mapping table would answer half the questions a stakeholder will ask.
5. **Add a market profile.** Five fields defining the temporal and jurisdictional context an adopting region operates in, so the not-specific-to-any-market claim becomes testable.
6. **Correct the temporal framing.** Distinguish dispatch cadence from settlement cadence, and state which one the decoupling argument is made against.
7. **Keep Grid 2 an overlay that never relaxes a Grid 1 obligation.** This is already the design intent and it is the single most important thing to keep saying, in every document, without qualification.
8. **Invite complementary intelligence layers** rather than attempting to embed all decision logic inside the protocol. The protocol should decide how much access an asset has. What the asset does with that access is a competitive layer, and Grid 2 will be adopted faster if that layer is open to competition.

12. Conclusion

Grid 2.0 correctly identifies the need to decouple access from peak planning while protecting the grid. It is unusual for a launch document to publish its own open questions as candidly as RFC #1 does, and it is that candour, rather than any disagreement, that made a detailed response worth writing.

Its success will depend on three things: rigorous incentive design, particularly around the allocation rule; honest acknowledgment of the prior art it builds on; and openness to complementary layers that asset owners and QSEs can adopt without friction. The first is where we have concentrated our comments, because it is the one a reviewer will find first and the one that is cheapest to fix now.

Wattness is building the complementary layer: physics-grounded deterministic agents, selective cryptographic witness, and ISO-native interfaces that require no tariff change to try. We would like to demonstrate through shadow and live pilots that verifiable edge intelligence is a practical accelerant for what Grid 2.0 is trying to do, and we would rather be told where we have misread the RFC than be agreed with politely.

We invite G2TF and the wider community to engage with this response through the G2TF process and at X: @wattness_ai.

Appendix A. Point-By-Point Commentary On RFC #1

This appendix preserves our reading pass in the order of the document, so each comment can be attached to the passage that prompted it. Where a point is developed in the main body we note the section and keep the entry short. Quotations are shortened to the operative phrase.

A1 "decouple system access from traditional reliability planning"

Summary and Benefits

We read this as the paper's central claim and we agree with it, with one qualification about emphasis. Reliability planning is not the obstacle; it is the guardrail that has kept the system running. What has failed to keep pace is the coupling between planning horizons and access decisions.

The overlay-with-fallback pattern is a good one: the enhancement carries the upside, the underlying layer carries the risk, and adoption stays rational under uncertainty. It is why Phase 0 is well designed. On *automatically protecting the grid*, see Section 1 for the passive and proactive readings; Wattness occupies the second.

A2 "FERC's recent show cause orders call for aggressive action"

Summary and Benefits

Accurate, and worth stating with more precision than the RFC currently does, because the precision is favourable. Section 2 gives the sequence: the section 403 directive of 23 October 2025, FERC docket RM26-4-000, and the six section 206 show cause orders of 18 June 2026 in dockets EL26-67 through EL26-72.

The word *aggressive* is FERC's own, from the Commission's news release. Reading it as an invitation for untested approaches would be a stretch; reading it as a signal that the Commission expects proposals rather than objections seems right, and it is a genuine opening for work of this kind.

A3 "but these options remain fragmented, somewhat manual"

Summary and Benefits

Consolidated in Section 2.2. In brief, *manual* means a human decision sits in the loop at a cadence slower than the phenomenon being managed, which is the real cost rather than the labour.

A4 "voluntarily adopt orchestrated management . . . in a real-time, permission-based paradigm"

Summary and Benefits

Voluntary adoption is the right strategy and shadow testing is the right proof mechanism. It also gives a clean sequence for anyone building alongside the standard: shadow, measure, then participate live.

On *permission-based*, we would resist reading this as opposed to permissionless systems. The interesting axis is not who is allowed to participate but whether participants must be trusted to be honest about what they did. Grid 2 can be entirely permissioned in membership and still verifiable in operation, and we think it should be. Membership is a policy question for the Task Force; verifiability is an engineering property, and only the second is our concern.

A5 "an unprecedented level of electric grid resiliency and efficiency"

Summary and Benefits

Aspirational language that we would soften, not because it is wrong but because it costs credibility the paper does not need to spend.

One substantive point on resiliency. The mechanism optimises each Allocation Domain against its own constraints, and domains that are individually well-optimised can still compose badly, because the constraint binding a region is often not local to any one domain. The federation design in Figure 6 of the RFC anticipates this; the allocation rule does not yet. This is the network-scale problem we work on and the one we would most like to study jointly.

A6 "fully recognized in economic dispatch and unit commitment paradigms"

Summary and Benefits

We agree that catalysing rather than replacing is the right posture.

One terminological note offered gently. Unit commitment is the scheduling of generating units over hours and is distinct from the real-time dispatch of distributed and behind-the-meter assets. The sentence reads as though Grid 2 would be recognised in both, which raises the question of what a one-minute protocol contributes to an hours-ahead commitment decision. If the intended claim is that Grid 2 participation should be visible to unit commitment as a modelled resource, that is a substantial and interesting claim deserving its own paragraph.

A7 "codifying connect and manage methods into a single dependable protocol standard"

Summary and Benefits

Codification is exactly right as a strategy: it converts a bilateral negotiation into an object outsiders can reason about and contribute to. The term itself has a specific origin worth attributing — Great Britain adopted an enduring Connect and Manage regime in 2010–11 on the generation side. See Section 3.

A8 "proven modern flow control . . . like TCP/IP, and . . . such as Kubernetes"

Summary and Benefits

The analogies are apt but load-bearing, and each carries a lesson the RFC does not yet claim. IntServ failed on per-flow state and DiffServ won by being stateless; Grid 2 has already made the right choice and should say so. Kubernetes QoS classes map onto DESP almost exactly, including eviction ordering, and the disanalogy is that a cluster has one administrator and no strategic actors. See Table 1 and Section 4.

The general caution we share: an analogy to a successful protocol is a hypothesis about which properties transfer, not evidence that they do.

A9 "only recently possible . . . due to the scaling of battery storage"

Summary and Benefits

We agree entirely; this is our thesis too. The capability arrived before the coordination layer, which is why the opportunity exists at all, and it is self-reinforcing: as storage and dispatchable large load scale, the untapped headroom that software can address grows faster than the hardware being added.

The blocker is not technical. Where a utility's regulated return accrues to capital deployed, a protocol that defers capital by using existing assets harder works against the incumbent's own incentive, and a five-to-ten-year build remains the most tried-and-tested path for a planner. We do not raise this cynically — it is why the affordability framing in the RFC is strategically correct — but it means adoption will be driven by the parties bearing the cost of delay, which is the large loads themselves. Designing for that population first, and letting each site maximise its own advantage inside a protocol that keeps the system whole, is in our view the realistic path to system-level benefit.

A10 "without the need for multiple stakeholder consensus processes"

Summary and Benefits

The PJM disputes of 2025–2026 concern capacity prices and stakeholder governance: Pennsylvania's complaint at FERC in docket EL25-46-000, the subsequent price collar settlement, several states publicly raising the prospect of leaving PJM, and FERC's own scrutiny of PJM governance during 2026. These are federal and multi-state matters. The Public Utility Commission of Texas governs ERCOT and is a separate track, where the relevant action is Senate Bill 6 and the Batch Zero rulemaking.

The substantive point survives. A voluntary protocol that routes around a stalled consensus process is genuinely valuable, and PJM is the live demonstration of why. It is also a reason to be candid that Grid 2 routes around governance rather than resolving it, which is a legitimate design choice but not a neutral one.

A11 "does not need to be universally applied"

Summary and Benefits

Agreed, and under-sold. Optionality is what makes the proposal adoptable by parties who cannot afford to be wrong — system operators, asset owners with warranty exposure, offtakers with delivery obligations. It deserves more prominence than a single clause.

A12 "at any scale or scope"

Summary and Benefits

Developed in Section 6. In brief: scale is not neutral because identical connection capacities do not imply identical deliverable capability, and portfolio composition changes the combinatorics of what a domain can actually promise. Scope is not neutral because jurisdictions differ in who holds curtailment authority.

A13 "Flexibility is implemented via curtailment . . . in one-off contracts"

Grid 1 vs Grid 2 comparison table

Broadly fair as a characterisation of the status quo, but it now understates the field. ERCOT's PCLR is neither one-off nor manual: it is a registered resource category following SCED base points in real time. The comparison table would be stronger if the Grid 1 column described 2024 rather than 2020. Alternative supply strategies, including small modular nuclear developers targeting co-located load, are also being attempted as a route to flexibility.

A small precision: curtailment is not solely a flexibility instrument. It is also how frequency is held, how thermal limits are respected, and how contingencies are survived. Framing curtailment as the primitive form of flexibility slightly obscures that some of it is not optional at all.

A14 "systematically validated by transmission owners and regional grid operators"

Grid 1 vs Grid 2 comparison table

This is the sentence that most directly motivates our submission. Validation by the transmission owner and grid operator is necessary and, we would argue, not sufficient in a design where the Allocator is explicitly stated to have no internal intelligence.

Our proposal is not that these institutions be distrusted; it is that verification be cheap enough to extend downward. A distributed asset at the edge of a domain should be able to demonstrate what it delivered without relying solely on an intermediary's attestation, and an intermediary should be able to demonstrate that it allocated as specified. The design principle we borrow from distributed systems is that the coordination layer should be the part hardest to corrupt: you can build centralised applications on a verifiable substrate, but not the reverse. Decentralisation in that sense is about political, geographic and technological resilience rather than about counting participants. See Section 8.

A15 "automatically and proportionally orchestrating participating loads"

Grid 1 vs Grid 2 comparison table

Proportionally to what, measured how, is the question Section 4 develops at length. Our position is that asset longevity — thermal headroom and state of health for storage in particular — belongs in the allocation input, not only in the owner's private optimisation.

There is a second issue underneath it. If a storage owner elects to accept accelerated degradation to capture a real-time price spike, that is a rational commercial decision with a calculable cost, and it is not the protocol's business — unless it is, because the domain was counting on that asset's declared buffering capability. Where that boundary sits should be stated rather than discovered in a dispute.

Relatedly, we would ask whether a member may participate while declining to accept the standard classification of part of its load. We think the answer should be yes, and saying so would broaden adoption.

A16 "a protocol stack to unify emerging flexibility proposals"

Grid 1 vs Grid 2 comparison table

We are aligned on the goal and bullish on the storage and renewables framing. The open question is what happens to flexibility as the standard succeeds: if the adopted parts harden into requirements, the overlay becomes part of Grid 1 and the optionality that made it adoptable disappears. That is not necessarily bad, but it is a trajectory to choose deliberately and state, because participants are being asked to adopt on the strength of it being voluntary. It also raises the question of whether competing task forces or protocols are expected, and whether Grid 2 intends to be one standard among several or the standard.

A17 "near real-time dispatch windows of one minute"

Proposed Scope

Plausible and achievable. Our comments concern the clocks it must nest inside rather than the window itself;

see Section 7.

A18 "version 26.0 of the protocol is focused on units of energy"

Proposed Scope

The substantive question is what happens when the focus broadens, as the RFC says it will with voltage support. A protocol whose scope expands will eventually overlap with protocols designed for the domains it expands into, and stating the intended boundary early is cheaper than negotiating it later. We would also welcome a sentence on why energy was prioritised; ease of implementation is a perfectly good reason and we would simply rather see it stated than inferred.

A19 "Stakeholders in ERCOT and PJM are already identifying paths"

Proposed Scope

We would like to know who, to the extent that can be said publicly, because it determines whether this is a live coordination effort or a statement of intent. We would like to participate in the ERCOT thread specifically, and can bring a network-optimisation and physical-state perspective to it.

A20 "the increase in non-firm revenues also helps drive down electricity rates"

Proposed Scope

The adoption path through hyperscalers is sound: they have the balance sheet to run pilots and the strongest exposure to interconnection delay.

The rate-reduction claim will be examined most closely, because it is the promise made to ratepayers. Brattle's March 2026 analysis (RFC #1, p. 12), which the RFC cites, models a representative mid-sized utility and finds that a ten percent improvement in system utilisation could reduce rates by roughly three and a half percent, all else equal, with modelled national savings above one hundred billion dollars over a decade. That is a real and useful number that depends on assumptions about pass-through. We suggest quoting it with its conditions attached rather than as a general expectation.

A21 "internet-like statistical multiplexing"

Proposed Scope

This is where the paper's target user finally comes into focus, and we would move the framing much earlier. Statistical multiplexing is the clearest one-line statement of what Grid 2 actually does.

A22 "the way laptop and smart phone hardware and software support TCP/IP"

Proposed Scope

The device-scale vision is attractive and the induction is doing a lot of work. TCP/IP scales across device classes because the abstraction hides the physical layer completely. Grid 2 cannot hide its physical layer: a thermostat and a two-hundred-megawatt data centre differ in ramp rate, metering granularity, latency tolerance and legal responsibility for the load. The claim that one standard operates at any scale needs the argument written out rather than assumed from the analogy.

A23 “initial work is underway . . . through the IEEE 2030.5 standard”

Proposed Scope

Consolidated in Section 10. In brief: the standard is IEEE 2030.5, its data model is XML over a RESTful HTTPS interface rather than JSON, and it is a device control and telemetry protocol rather than a market bidding protocol. That makes it a plausible foundation for G2P and an implausible one for anything touching settlement, and the RFC should say which layer it means. We intend to look at whether 2030.5 can serve as common ground for the capability descriptor proposed in Section 6.

A24 “ensuring electricity affordability and rapid interconnections”

Proposed Scope

We are aligned on both objectives and it is worth being clear that our motivations are commercial. We work on affordability through better price and dispatch intelligence, which serves asset owners who do not have a trading desk of their own — large trading firms already have this capability in-house and generally keep it there. We work on interconnection through physical-state-aware use of headroom, which increases what an existing connection can safely deliver. We are for-profit and the nodes we optimise for are the ones that pay us. We do not think that conflicts with the RFC’s public objectives, and we would rather state it than have it inferred.

A25 “permission-based use of (mostly surplus) grid transmission and power capacity”

The Grid 2.0 Proposal

This confirms our reading that Grid 2’s object is surplus. Our interest runs to the deficit side as well: what happens in the intervals when there is no surplus to allocate, which are the intervals that determine both reliability outcomes and the bulk of economic value. A protocol optimised for surplus allocation should say what it does when headroom is negative, because that is when its allocation decisions matter most and when its failsafe is most likely to be exercised simultaneously across members. Our own models are deliberately variable-agnostic in this respect — scarcity and surplus are the same optimisation with different signs.

A26 “mirroring the flow control fundamentals of other constrained networks”

The Grid 2.0 Proposal

We researched this specifically. The constrained networks the phrase points to are, in order of closeness: DiffServ, which pairs stateless classes with edge policing; the ATM available bit rate service, which ran a closed rate-feedback loop; weighted fair queuing; telecom admission control; and Kubernetes resource requests with QoS classes. Table 1 summarises what each teaches.

The single most important finding is negative and is developed in Section 4: none of these systems allocates proportionally to the declared request without a second mechanism that disciplines the declaration. Proportional allocation is the visible half of every one of these designs; policing is the half that makes it work.

A27 “from local transmission congestion zones . . . into regional RTOs/ISOs”

The Grid 2.0 Proposal

The federation is the hard part and the RFC treats it briefly. Local congestion is not independent of balancing-authority conditions; the same physical flows appear in both, so a local Allocator clearing against a local constraint may be allocating headroom a regional Allocator has already committed. The announce-and-listen federation sets up the communication but not the precedence.

Our concrete question: when a local and a regional Allocator produce inconsistent commitments within the same minute, which one does the Member Element follow, and how does it know? A stated precedence rule would resolve this.

A28 "allocate percentage dispatch … proportional to their requested amounts"

The Grid 2.0 Proposal, Structure and Figure 3

We like the figure; the mapping from Allocators to Grid 1 constraints is the right thing to put at the centre.

On proportionality, beyond the incentive problem of Section 4, there is a functional-form question. Linear proportionality is an assumption about the right allocation function, not a derivation. Whether a linear rule dominates a learned or physically-conditioned rule is an empirical question, and Phase 0 with an evidentiary record is precisely the apparatus that could settle it for a given domain. We would rather see linear proportionality framed as a testable default than as a design conclusion.

A29 "the Grid 2 protocol stack and its position between Grid 1 domains"

The Grid 2.0 Proposal, Structure and Figure 4

The layering is clear and we have mirrored it in our Figure 1 so the two documents can be read side by side.

We intend to review the reference implementation as it appears in the Task Force's public repository and to comment at the level of the code rather than the prose, including how the one-minute cadence is structured and how the stack sits against the applicable NERC BAL reliability standards. We would rather submit that as a second, narrower comment than speculate here.

A30 "Protocol Interactions"

The Grid 2.0 Proposal, Protocol Interactions table

Edge intelligence for Use, Buffer and Source members is our largest common ground with the Task Force, and the honesty of the Allocator row — stating plainly that it has no internal intelligence — is what let us locate our contribution precisely. See Section 8.

The complementary claim we would make is that intelligence at the edge and verification at the edge have to arrive together. An edge that decides for itself and reports for itself is an edge that cannot be checked, and the failure will not be visible until a scarcity interval, when it is expensive. This is also where we would ask whether the DESP classification is negotiable: a member may agree to participate and still disagree with how its load maps onto three classes.

A31 "Key Design Principles and Characteristics"

The Grid 2.0 Proposal, Key Design Principles and Characteristics table

Two rows we would revise. *Coordination without prices* is addressed in Section 5: we do not think protocol-based coordination can be a standalone rule while the paper simultaneously claims benefits in electricity affordability, which is a price outcome. *Scale without singular operator* would be better stated as scale without a single point of control, because the operators remain and their objectives differ.

A32 “the smallest Grid 2 Allocation Domain would include a host utility and a flexible load customer”

Grid 2.0 Membership and Basis of Authority

A clean and useful definition. We would note that the minimal domain is also the minimal verifiable unit, which makes it the natural target for a shadow-mode pilot: one utility, one flexible load, one storage asset, one public record. That is the scope we would propose for the reference implementation in Section 8, and we see no reason the same methods could not later be applied at the Allocator layer itself, treating a set of Allocators as a single optimisation object.

A33 “Numerous announcements have been made about arrangements between utilities and flexible large loads”

Grid 2.0 Membership and Basis of Authority

We looked into this claim and it holds, with a caveat about which examples do what. The strongest documented case is Google’s demand-response agreements for machine-learning workloads with Indiana Michigan Power and the Tennessee Valley Authority, announced in August 2025 and since extended. Regulated flexible-access products now exist in tariff: ERCOT’s PCLR and SPP’s CHILLS.

A distinction worth preserving. Large flexible loads in ERCOT — including Bitcoin mining campuses such as Hut 8’s Vega facility in the Texas Panhandle, energised in mid-2025 — are genuinely interruptible, but their flexibility is commercial rather than contractual: they curtail because it pays, not because a connect-and-manage instrument obliges them to. Both are useful to Grid 2 and they are different evidence. We suggest RFC #2 separate contractual flexibility from economic flexibility, because only the first is a commitment a protocol can rely on.

A34 “the federated regional implementation diagram”

Extensible System-Wide Architecture, Figure 6

The nested visualisation communicates the scalability argument well. Our question about precedence between nested Allocators, at A27, is the one thing the figure raises without resolving.

A35 “These prior studies do not propose a unifying and automated protocol”

Industry Analysis Informing Grid 2.0

We found a tension worth naming. The RFC establishes the opportunity using studies that examine system-wide capacity potential, then narrows to data centres and interconnection speed, which is where the incumbent effort and the profit already are. If the analytical case is system-wide, the natural next move is to ask where the unserved whitespace is; if the strategic case is hyperscaler-first, that is a defensible go-to-market and should be stated as one. We are taking the smaller-scale route deliberately, starting with independent and developer-owned storage, and we think both routes are needed.

We would also note, since it bears on how the studies are used, that the Duke figures the RFC cites — 76 GW of new load at an average annual curtailment rate of 0.25% — come with conditions the authors have been explicit about, including that the analysis does not account for transmission constraints. Cited with its conditions it is a strong number; cited without them it invites the first reviewer to discount everything around it.

A36 “Differential Reliability is a possible path to ration scarce generation capacity”

Industry Analysis Informing Grid 2.0

We think PJM’s phrasing is better than the RFC’s own and would encourage the Task Force to adopt it. *Differential reliability* names the thing being allocated. *Proportional allocation* names an arithmetic operation and leaves open what it operates on. The first is a policy commitment that can be debated on its merits; the second is a formula that will be debated on its parameters.

A37 “the market approaches are equally diverse and complex to implement”

Industry Analysis Informing Grid 2.0

Agreed on the alignment with European and Australian work, with one framing note. The alignment is partly a consequence of a choice rather than a discovery: the RFC focuses on the network side and places market coordination outside its scope. That is a legitimate scoping decision and we would simply prefer it were presented as one, since it is the same decision that produces the coordination-without-prices principle we question in Section 5.

A38 “the Open Questions section”

Open Questions

This is the best section in the paper and it is unusual to publish it. Our substantive responses are consolidated in Section 9 rather than repeated here. We agree that Grid 1’s regulatory scaffolding is a genuine asset and that the standard should grow within it, and we would like to join the requisite technical studies the section invites.

A39 “Current Industry Approaches”

Current Industry Approaches table

Consolidated in Section 10, covering the FERC directives row, the energy markets row, the custom rate structures row, the off-grid loads row and IEEE 2030.5.

One general remark. Digitising a set of concepts and complying with existing standards is not technically difficult, and the question of why it has not been done before is worth answering directly. Our reading is that it was not previously possible rather than not previously attempted: the buffering assets that make best-efforts service tolerable, and the dispatchable large loads that make it worth doing, have only recently reached scale. The RFC makes this argument well in the Summary and then does not use it here, where it would do the most work.

A40 “Phase 0 — Shadow Mode”

Phased Adoption

The strongest idea in the deployment plan. Executing only what would have happened anyway removes every adoption objection at once, which is a rare property in a reform proposal, and it is the same sequence we follow ourselves. Our suggestion in Section 8 is simply that the shadow record be made evidentiary rather than merely logged, so the phase produces something a regulator can rely on.

A41 "due to the urgency of related policy actions underway across the United States"

Beyond the Load Growth Challenge

We raised this as a contradiction in our working draft and, on re-reading, we think it is not one. Deferring to the urgency of US policy action while criticising a specific instrument's ability to produce uniformity are consistent positions. What creates the appearance of tension is the Current Industry Approaches table framing directives as a competing approach; fix that framing, as Section 10 proposes, and the passage reads naturally.

A42 "the two planned companion documents"

Beyond the Load Growth Challenge

We would like to contribute to the network-directed storage document. Anticipatory movement of energy in space and time is a network optimisation problem over a graph with physical constraints at every node, and it is the problem Wattness is built around.

On the renewables document, we would ask that *smarter separation of rapid access and long-term planning* be made concrete early, together with the mechanism by which renewable interconnections are actually accelerated, since it is the same separation the load-side paper makes and the mechanism should probably be the same one.

A43 "ready-made to help deliver the full system benefits of these innovations"

Beyond the Load Growth Challenge

We share the interest in hardware that changes the thermal and locational economics — co-location of compute with generation, heat recovery that returns waste heat to useful work, and power electronics that make routing decisions possible where today there are only flows.

One item is already in the RFC and deserves promotion. The Allocator row names dynamic line ratings as an input. Conductor rating is temperature and weather-dependent, and a domain that observes line conditions rather than assuming a static seasonal rating has materially more headroom to allocate, particularly in the hours when the constraint binds. If the Task Force is looking for a high-value, near-term input to standardise, we would nominate that one.

A44 "The G2TF was inspired by the Internet Engineering Task Force"

Development and Engagement Approach

The comparison is apt and the development principles are well chosen. We would note only that the IETF's authority rests on open participation and individual merit, and that the question of how a small contributor participates — and how rough consensus is judged when three named organisations support the work — will be asked. Publishing the membership and decision rules early would cost little and pre-empt the obvious critique.

A45 "Researchers and company labs actively modeling these dynamics"

Focused Stakeholder Outreach

This is the category we would fall into. We are also seeking validation and pilot opportunities with the other stakeholder groups named, and would welcome introductions where the Task Force thinks there is a fit.

A46 “Grid 2.0 is not intended to be specific to any market or nation”

Focused Stakeholder Outreach

Developed in Section 7. In brief: move the claim to the front and support it with a short market profile that an adopting region completes. Singapore — half-hourly dispatch and settlement, forty-eight periods per trading day — shows that the temporal nesting argument does not transplant without being re-derived.

Appendix B. Where The Three Layers Sit

Dimension	Grid 1 (Status Quo)	Grid 2.0 (Proposed)	Wattness Contribution
Access model	Firm at interconnection, for the life of the contract	Permission-based, tiered by service class, re-cleared each minute	Does not allocate access; decides what an asset does with the access it has
Protection	Planning margins plus manual curtailment	Automated proportional orchestration with failsafe reversion	Deterministic, physics-bounded agents whose decisions are reconstructible
Coordinating signal	Prices and bilateral contracts	Explicitly price-free protocol messages	Physics and marginal cost, with selective cryptographic audit
Unit of decision	Hourly to daily scheduling	One-minute allocation window	Offer curve and flexibility schedule, at the market’s own cadence
Treatment of asset state	Nameplate ratings and seasonal assumptions	Connection size and declared request	State of charge, state of health, thermal headroom entering the optimisation directly
Adoption path	Mandatory, tariff-driven	Voluntary, shadow then live	Shadow analysis, then pilot through an existing QSE, then product
Interface to the ISO	Native	New protocol stack, to be built	Ordinary offer curves, telemetry and settlement files, submittable today

Table 3. The three layers are complements. Grid 2.0 governs access; Wattness governs the decision taken within it; Grid 1 remains the binding constraint on both.

Appendix C. Sources

Every external factual claim in this response traces to one of the following. Where a claim rests on a secondary source rather than a primary document, we say so.

Federal action on large loads

- Secretary of Energy, Direction that FERC Initiate Rulemaking Procedures Regarding the Interconnection of Large Loads, under section 403 of the Department of Energy Organization Act (42 U.S.C. § 7173(a)), 23 October 2025. FERC docket RM26-4-000; the Commission's docket page states that large loads are generally defined as demand greater than 20 MW and that FERC committed to act by June 2026. `ferc.gov/rm26-4`
- FERC, "FERC Launches Aggressive Targeted Action to Speed Large Load Integration," news release, 18 June 2026. Six tailored show cause orders under section 206 of the Federal Power Act to PJM, MISO, SPP, CAISO, ISO-NE and NYISO; items E-7 to E-12, dockets EL26-67 through EL26-72; 30 days for an informational report on generation adequacy and 60 days to justify or reform tariffs; five named reform categories including providing new transmission services for flexible large loads.
- Secretary of Energy letter of 29 August 2025 invoking the same section 403 authority in a natural gas certificate proceeding (Docket Nos. PL18-1-000 and PL18-1-001). Cited here only to distinguish it from the large-load directive.

ERCOT

- ERCOT, Trending Topic: New Batch Connection Process for Large Electricity Users, 18 June 2026.
- ERCOT Large Load Integration pages, including Section 23 Form W Part A for PCLR declaration of intent (to the interconnecting DSP or TSP by 10 July 2026, onward to ERCOT by 24 July 2026). `ercot.com/services/rq/large-load-integration`
- PGRR145 and NPRR1325 establish the Batch Zero framework, the Low Power Consumption and Maximum Power Consumption constructs, the Provisional Controllable Load Resource and the Withdrawal-Limited Private Use Network. ERCOT Board approval 1–2 June 2026; PUCT approval 18 June 2026; effective 10 July 2026. Adjusted Bid Caps cap a PCLR's energy bid curve based on constraint shadow prices and the PCLR's locational impact on those constraints, and ERCOT does not guarantee the requested maximum can ever be served in real time. Secondary sources consulted for the summary of NPRR1325 provisions: Willkie Farr & Gallagher, 16 June 2026; Sheppard Mullin, 23 June 2026; EPE Consulting.
- ERCOT market cadence: the day-ahead market gate closes at hour ending 10:00 CT on the day before the operating day; SCED produces locational prices at least every five minutes; real-time energy settles on fifteen-minute settlement intervals priced as the time-weighted average of the SCED prices within the interval.

Other market and regulatory sources

- SPP, High Impact Large Load integration and Conditional High Impact Large Load Service (CHILLS); FERC accepted the CHILLS tariff revisions effective 1 July 2026, Docket ER26-1323.
- PJM, *Powering Reliability Through Market Design*, special report, 6 May 2026, setting out Path A (stabilised markets), Path B (differential reliability) and Path C.
- PJM capacity and governance disputes: Commonwealth of Pennsylvania complaint at FERC, docket EL25-46-000, filed 30 December 2024; the resulting price collar settlement in docket ER25-1357; FERC proceedings on PJM governance and stakeholder processes during 2026.
- Texas Senate Bill 6, 89th Legislature (2025), and the associated PUCT rulemakings on large load curtailment obligations, demand reduction procurement and transmission cost allocation.
- Great Britain: the enduring Connect and Manage transmission access regime, published 2010 with implementation completed in February 2011; Ofgem's Connect and Manage monitoring reports document connection dates brought forward by several years alongside materially increased constraint costs. We cite

the qualitative finding; readers wanting annual figures should consult the Ofgem reports directly.

- Energy Market Company, National Electricity Market of Singapore: the real-time market clears and settles on half-hourly trading periods, forty-eight per trading day.

Standards and networking precedents

- IEEE 2030.5 (Smart Energy Profile 2.0): application-layer protocol over HTTPS with TLS, RESTful resource model, XML data model; adopted as the default smart-inverter protocol under California Rule 21 through the Common Smart Inverter Profile.
- IETF RFC 1633 (Integrated Services) and RFC 2475 (Differentiated Services architecture), for the per-flow-state versus stateless-class history and for traffic conditioning at the network edge.
- ATM Forum Traffic Management specification, for the CBR, VBR, ABR and UBR service categories and the ABR closed-loop rate feedback mechanism.
- Kubernetes documentation on Pod QoS classes (Guaranteed, Burstable, BestEffort) and eviction ordering under node resource pressure.
- On the manipulability of proportional division: X. Bei, B. Tao, X. Wu, S. Yang, "The Incentive Guarantees Behind Nash Welfare in Divisible Resources Allocation," arXiv:2308.08903 (WINE 2023), which bounds rather than eliminates the gain from misreporting under mechanisms of this family.
- FERC Order No. 2222, on the participation of distributed energy resource aggregations in RTO and ISO markets, referenced in Section 8 as a later-stage pathway.

Studies cited by RFC #1 that we relied on

- T. Norris, T. Profeta, D. Patiño-Echeverri, A. Cowie-Haskell, *Rethinking Load Growth: Assessing the Potential for Integration of Large Flexible Loads in US Power Systems*, Nicholas Institute, Duke University, February 2025. 76 GW of new load integrable at an average annual curtailment rate of 0.25%, rising to 98 GW at 0.5% and 126 GW at 1.0%, across the 22 largest US balancing authorities. The lead author has publicly emphasised that this is a first-order technical potential assessment that does not account for transmission constraints.
- R. Hledik, L. Lam, K. Peters, *The Untapped Grid: How a Better Utilized Power System Can Improve Energy Affordability*, The Brattle Group, March 2026, commissioned by the Utilize Coalition and GridLab. Improving system utilisation by 10% could reduce electricity rates by roughly 3.4% relative to current levels, all else equal, with modelled national consumer savings above $100 billion over a decade.
- Google, announcement of demand-response agreements for machine-learning workloads with Indiana Michigan Power and the Tennessee Valley Authority, 4 August 2025, subsequently extended to further utilities.
- Hut 8 Corp., energisation of the Vega facility in Oldham County, Texas (Panhandle, within ERCOT), 30 June 2025, cited only as an example of an economically flexible large load.