Comments of Donald C. Jackson on the Commissioner Workshop on Interoperability

Docket 24-FDAS-02, Load Flexibility Policy & Planning

Comments of Donald C. Jackson on the Commissioner Workshop on Interoperability of November 4, 2026

Donald C. Jackson, Grid Coordination grid-coordination.energy

With contributions from Bruce Nordman.

Summary

The grid has two problems and they are not the same problem. Balancing supply and demand across the state is an economics problem, answered by a price that changes hourly or faster. Protecting the transformer and the service drop serving one house is a physics problem, answered by a limit on that customer's whole connection, import and export. Both land on the bill. Today the customer receives neither signal directly.

What I ask the Commission to do.

What a flexible demand appliance standard should require

1. Require that each appliance have an OpenADR 3 Virtual End Node (VEN) of its own, implemented in the appliance rather than in a service it depends on, and test it with the appliance disconnected from the wider internet. The requirement should attach to receiving a price, not only to demand response participation as Section 1693(c)(2)(E) currently scopes it. Proposed rule text will be filed in Docket 24-FDAS-03.
2. Require that the OpenADR 3 server address, the Virtual Top Node (VTN) the appliance connects to, be configurable by the customer, and that the appliance first attempt discovery of a VTN on the customer's own network before attempting connection to a server beyond the premises. A customer-configured address should take precedence over both, and may itself be a service the customer chose, which is a different thing from the one the manufacturer pre-configured.
3. Require that every power and energy management capability an appliance implements be available locally to its owner, over an open, standard local protocol, and to an agent the owner authorizes, on terms no worse than those on which the manufacturer makes it available to itself or to any third party. A manufacturer may withhold functions that exist for safety, certification integrity, or compliance with interconnection requirements. It should not withhold from the owner a management function it uses itself or grants to others.
4. Provide that an appliance satisfies a flexible demand capability on its own. An appliance may additionally support working with an intermediary, and a customer remains free to add one, but compliance may not rest on it.

How the rule should be written

5. Write the capability once, in one common flexible demand appliance standard (FDAS) that states it for every appliance category and is the only place it is stated, leaving the appliance-specific rulemakings to their test methods and control vocabularies.
6. State that capability as the two signals, a price and a limit. Do not oblige an appliance to implement demand response operating schedules, which Section 1693(c)(2)(F)1 and Table C-1 item 7.5 currently require, while continuing to require reception of events carrying a price, an emissions intensity or a reliability condition.
7. Provide that the thermostat standard now in Express Terms marks which of its provisions are appliance-independent and which are specific to thermostats, and that the appliance-independent provisions are superseded by the common standard when one is adopted.
8. State the protocol criteria in the rule, meaning the properties a protocol must have to be eligible for the mandate rather than the name of any one protocol, name every protocol that meets them on adoption, and maintain the list where the Commission can amend it by Business Meeting approval rather than by a new rulemaking, as Section 1623(c)(5) of the Load Management Standards already provides for the statewide rate tool. Any party should be able to submit a further protocol with a showing that it meets the criteria, and the Commission should decide within a stated period and publish its reasons.
9. Where conformance rests on a third-party certification mark, require that the specification be published without paywall or membership condition, that the conformance test procedure be published on the same terms, and that a customer's ability to point the appliance at a server of their own choosing be one of the behaviors the test procedure actually verifies.

What the grid side should publish

10. Require utilities to publish prices over OpenADR 3, and have the MIDAS server obtain them as a client.

Capacity limits, which answer a different problem

A capacity limit is a property of the whole service connection, where everything above concerns what an individual appliance must do. The two are not variants of one requirement, and a rule that treats them as one will get both wrong.

11. State what each kind of signal obliges an appliance to do: a capacity limit on the service connection constrains all operation, including operation a customer or their authorized agent directs; a price is an input to optimization and never a command; within the limit the customer's instruction governs; and a third party acting on a customer's authorization has the customer's authority and no more.
12. Define how a capacity limit is represented, once, and require no appliance to hold the whole service connection under the limit. One listed power control system enforces it. The limit is delivered to that system, and to an energy management system planning the premises. It is not delivered to every appliance.

13. Require a utility that computes a capacity limit for a service connection to publish it to that connection's customer, over an open, standard protocol, authenticated by the customer's existing utility account credential. A utility already operating an OpenADR 3 Virtual Top Node to publish prices should publish limits over the same one rather than build a second interface. The obligation attaches where the limit exists; a utility that computes none has nothing to publish.
14. Provide that an existing utility account credential is sufficient to authenticate to a utility, and that conformance is not conditioned on program enrollment, a third-party contract, or a manufacturer-issued device certificate.

The findings these requests rest on.

1. **The draft thermostat standard published six days after this workshop never requires, in terms, that the open interface be in the appliance, and does not test it.** Its definitions of "connected device" and "interoperable device" point at the appliance, but the provision that carries the requirement names no actor, compliance rests on manufacturer attestation, and the demand response test passes a thermostat whose manufacturer relays events from its own cloud. **The fix is to say in Section 1693(c)(2)(A)2 what the definitions already imply, and to test it.** It matters because a customer whose only route to their own thermostat runs through a company they did not choose can be charged for passage, limited in what they may do with their own appliance, and cut off entirely when that company changes its terms or exits. The Commission specified what the thermostat must do with a signal and left unspecified where the thermostat gets the signal from.
2. **Most of what the thermostat draft requires applies to every appliance.** Retrieving a price on an interval, holding the software that turns it into a schedule, falling back to a stored copy when the network is gone: none of that is about thermostats. Drafted per appliance category, it will be written again for each one. The Commission has built one rule covering many parties before, in the Load Management Standards.
3. **Aggregation is a service, not an interface.** Commissioner McAllister asked twice about the role of the aggregator. The answer that protects both the aggregator and the customer is to require the interface at the appliance and let aggregation compete on top of it. A rule that assumes an aggregator in the path makes that position a toll booth; a rule that assumes none makes aggregation a product people choose. The toll is not the aggregator's alone. A manufacturer that holds the only route to its own installed base can charge for passage without ever calling itself an aggregator, and some would rather collect the toll than run the service.
4. **A customer who buys a home battery often cannot tell it when to charge or discharge, while an aggregator can.** California's virtual power plant results are a genuine achievement, and they were delivered through an architecture that does not reflect the Commission's own commitments on consumer-authorized control and freedom from proprietary restriction. Two things follow for appliance standards: a capability a manufacturer implements should be available to the owner, and a rule should say that a capacity limit binds while a price only informs.

5. **The MIDAS price server is built on the right principles and should acquire its data the same way it serves it.** Reading prices as an OpenADR 3 client removes a custom integration and leaves utilities operating the OpenADR 3 VTN service their own customers will need.
6. **Legacy bridging belongs in products, not in the mandate.** Schneider Electric described building protocol translation into its energy management system. That is the right division of labor and the CEC should codify it: the rule names what a new appliance must speak, and a product decides what else it chooses to speak.
7. **A certification mark is not interoperability.** The Mercury Consortium presented on Panel 2. Measured against the six tests set out in *Two Signals, Six Tests*, attached to these comments, its published EVSE specification fails four of the six. The response-time classes the Mark is advertised on appear in no published document, and its conformance test procedure has never been released.

1. The draft thermostat standard does not require the open interface in the appliance

Six days after this workshop, the Commission published draft Express Terms for flexible demand appliance standards for HVAC thermostats in Docket 24-FDAS-03, together with a draft staff report and a test method. A package of that scope was months in preparation. It is the Commission's considered position, not a first attempt, and nothing in it responds to anything said at this workshop.

It is the only place where the approach this workshop was convened to discuss has been written as enforceable rule text, so it shows what the Commission presently believes that approach requires.

The draft is a substantial improvement on anything that preceded it, but it is not good enough. Its requirements for what the appliance does with a signal pass every one of the six tests below. Its requirements for where the signal comes from fail the tests that decide who the customer deals with, beginning with whether the open interface reaches the appliance at all, which is the question this workshop was convened to address. The draft's own definitions point at the appliance; its operative provision and its test method do not. The draft is mature work rather than a rough cut, so the gap is a position the Commission currently holds rather than an oversight it has not yet caught.

At the workshop I offered six tests for any flexibility proposal, in the presentation *Two Signals, Six Tests*, which is attached to these comments. Applying them to the draft Express Terms:

PASSES. Test 1, does the grid send an objective, or operate the appliance? The thermostat computes its own operating schedule from MIDAS data, which the draft calls autoflex, and decides what to do within the consumer's Comfort Bounds. A demand response event overrides that schedule temporarily and remains constrained by the same bounds, so even the draft's dispatch path does not operate the appliance directly. Whether the dispatch path should exist at all is Section 2's question; on this test, as drafted, the grid sends information and the appliance decides.

FAILS. Test 2, does the open interface reach the appliance, or stop at the manufacturer's cloud? The draft never requires the interface to be in the appliance, and the only compliance step that asks is a manufacturer attestation no test procedure verifies. It does not permit the cloud either; it simply does not decide, and a rule that does not decide gets decided by whoever builds to it.

The definitions point at the appliance. Section 1693(c)(2)(A)1 requires each thermostat to meet the Section 1691 definition of a "connected device," which is a device that can "communicate via open standards with entities outside the device …" Section 1693(c)(2)(A)2 requires each thermostat to be an "interoperable device," defined as "a device utilizing a standardized, non-proprietary protocol to exchange information with external entities …" Both put the protocol in the device. Section 1693(c)(2)(D)1.a shows the Commission writing that obligation plainly when it chooses to: "[w]hen connectivity is enabled and a connection is available, the HVAC thermostat shall, at intervals not exceeding 26 hours, attempt to retrieve the most recent MIDAS dataset corresponding to the active RIN."

The operative provision does not. The sentence that carries the interoperability requirement reads "Interoperability shall be provided by complying with the Virtual End Node (VEN) requirements specified in the OpenADR 3.1 standard." It is passive and names no actor. Section 1693(c)(2)(E) provides that "[f]or participation in demand response programs and communication of demand response operating schedules, each HVAC thermostat shall meet the minimum specifications of OpenADR 3.1 and shall not rely solely on a manufacturer or third-party service" whose discontinuation would disable demand response. That scoping is itself the defect. The open protocol is required for the event path. The price path, in Section 1693(c)(2)(D)1, is "an application programming interface using open internet protocols" and names no protocol at all. "Shall meet the minimum specifications" does not say which component meets them. The "not rely solely" clause is a genuine protection against one failure mode, a manufacturer that discontinues the service, and the draft's stored static dataset is how a thermostat would satisfy it. It is not a requirement that the VEN be in the appliance: a thermostat whose VEN runs in its manufacturer's cloud, falling back to the stored dataset when that cloud is gone, complies. The draft contemplates the manufacturer's cloud in three places, Section 1693(c)(2)(D)4 and Table C-1 items 0.14 and 0.18, of which the first two, which concerns updates to the stored MIDAS static dataset.

And nothing tests it. The compliance path reaches the question only by attestation. Table C-1 item 0.3 requires that the "[m]anufacturer attests that the unit meets the Virtual End Node (VEN) requirements," and item 0.18 that the "[m]anufacturer attests that unit can operate with all connected flexible demand functionality without connecting to manufacturer cloud." That second attestation asks the right question. No test procedure in the Test Method verifies it, and the Test Method's own OpenADR verification accepts either an OpenADR Alliance certified product listing or, in its own words, a self-attestation statement the manufacturer provides, neither of which establishes where the VEN runs. An attested negative is not a tested requirement. Its demand response procedure directs the operator to "[c]onfigure the UUT to connect to the testing demand response OpenADR server based on manufacturer's instructions," and permits delivery "through another suitable method, such as a utility- or aggregator-hosted demand response automation server (DRAS)." A manufacturer that relays events from its own cloud passes, because the test observes the setpoint rather than the endpoint.

A manufacturer will build the cheapest thing that passes. Here that is a thermostat whose Virtual End Node lives in the manufacturer's cloud, because the operative provision never says otherwise and no test procedure looks. The requirement that would prevent it exists only if a reader assembles it from two definitions, and no laboratory does that.

The fix is smaller than the defect: say in Section 1693(c)(2)(A)2 what the definitions already imply, and test it with the appliance disconnected from the wider internet. That condition is what makes the test mean anything. A thermostat can present a local endpoint that forwards to its manufacturer's cloud and returns the answer, which satisfies a test run on a connected bench and defeats the requirement entirely. Proposed rule text will be filed in Docket 24-FDAS-03.

A shared VEN leaves the customer without an interface

The draft is silent on where the Virtual End Node lives, so a pooled one is permitted, and much of the market already works that way.

A VEN operated by an aggregator represents a fleet. A signal sent to it arrives at the aggregator, which then decides which appliances move, by how much, and whether a given customer's appliance participates at all. That is a consequence of pooling the devices, not a property of OpenADR 3. The customer's appliance holds no address of its own, so there is no interface the customer, or anyone the customer chooses, can reach.

Which signals can be pooled, and which cannot, matters for what the Commission is considering next. A retail tariff price is a public attribute of the tariff, identical for every customer on it, so an intermediary holding one subscription on behalf of a hundred thousand customers receives exactly what each would have received alone. A price that varies by location is a different case and sits between the two. An operating envelope is different in kind, and is addressed in Section 5 rather than here: it is a property of one service connection rather than of a device, so the question of who may hold it is a different question from the one this section is about. The short version is that whoever holds a pooled VEN decides whether a limit aimed at one site is honored, and the customer whose transformer it protects is not a party to that decision.

My assessment is that the aggregation businesses operating in California today present one interface to the grid on behalf of a pooled fleet. That is the efficient design for distributing a tariff price, which is public and the same for everyone on it, and the wrong one for a per-customer limit.

This is not a hypothetical reading of the draft. It describes how residential thermostat demand response is already delivered in California, and how the firms that deliver it expect to continue delivering it. Renew Home's response to the Commission's Request for Information in Docket 24-FDAS-03 is cited here because it is an unusually clear statement of the model rather than an unusual one. It states that "VPPs, DERMs and device makers have developed various systems and APIs for running load shifting events on thermostat fleets," and that "[t]he Google Nest API controls a substantial fraction of smart thermostats enrolled in DR programs." On the protocol question it states that "[i]n theory" OpenADR "could support much of the needs for communicating between aggregators and devices, and it could support direct to device communication," while characterizing "direct to device being suboptimal, expensive" and OpenADR itself, while conceding that "OpenADR is designed to do these things," as "not yet widely used or properly designed for the real world implementation." It observes that "direct

integration with OEM cloud APIs could allow significant market share if those OEMs provided load shifting API support." And among the reasons it gives for manufacturers preferring the cloud, it is candid that cloud-based optimization "enables OEMs to capture some recurring revenue to defray other ongoing costs."

The point is about the rule rather than the filer, whose technical observations are reasonable and whose commercial reasoning is more candid than most. A rule that does not locate the VEN at the device will be read the way that filing describes, by every firm with an installed base to protect and a recurring revenue stream to defend.

There is also a jurisdictional consequence. Title 20 standards apply to appliances sold or offered for sale in California. A requirement satisfied by a manufacturer's cloud service is satisfied by something the Commission does not certify, cannot test on a sample, and cannot enforce against after sale. A manufacturer could change or withdraw that interface without affecting the certification of a single thermostat.

The fix is one sentence, and a test condition: the Virtual End Node required by Section 1693(c)(2)(A)2 shall be implemented in the HVAC thermostat, and conformance shall be demonstrated with the thermostat disconnected from any network path beyond the premises. Proposed rule text will be filed in Docket 24-FDAS-03.

FAILS. Test 3, can the customer choose who optimizes their home? Alongside the utility, the load-serving entity and the demand response aggregator, there is a fourth kind of party: a service that optimizes a customer's energy use on that customer's behalf, with no relationship to the grid and no program to enroll in. Such services exist today, and they are the natural beneficiary of a configurable server address.

The draft goes further than it is given credit for, and then stops short in the way that matters. Section 1694(b) requires a connected appliance to be "capable of receiving, acting upon, and responding to authorized remote requests via a communication link" that "schedule, shift, or curtail appliance operations," at minimum "the start or stop of operation" and "changes to equipment operation or schedule." Section 1691 defines "[a]uthorized entities" as "the owner of the equipment, the user of the equipment, or the owner's or user's delegated agent." The fourth party is contemplated, and the authority to direct the appliance is already written.

What is not written is where that communication link terminates or what protocol carries it. The draft names no local interface, no address the customer can set, and no protocol for this path, and the words local area network, energy management system and Matter appear nowhere in the Express Terms. A customer may authorize an agent to operate the thermostat and have no specified way for that agent to reach it. That is the same defect as the Virtual End Node: the obligation is written, and its location is not.

The consent provisions are good and they answer a different question. Third-party demand response requires express consumer authorization to opt in, the consumer may opt out of autoflex at any time, and may disable connectivity at any time. Every one of those is a power to refuse. The consumer may decline a program, switch off autoflex, or disconnect the thermostat from the network entirely. None of them is a power to choose: nothing in the draft lets a consumer decide which server the thermostat talks to, only whether it talks at all.

The draft requires a VEN and never says which VTN it connects to, or who decides. There is no consumer-configurable server address and no local network path, so an appliance that must speak an open protocol will speak it to whichever server its manufacturer selects. A requirement to speak an open protocol to an address the owner cannot change is a default that is a lock.

The capability is already in the standard the Commission chose, and the rule should reference it. OpenADR 3.1 specifies discovery of a local VTN using mDNS and DNS-SD, advertising the service type _openadr3._tcp., and states that "mDNS and DNS-SD support does not change the content of the OpenADR 3 standard; rather it occurs before any OpenADR communication occurs." A configurable address and local discovery therefore add no protocol work to a conforming device. Section 1693(c)(2)(A)2 should cite that discovery mechanism directly, so the obligation lands on a provision the standard already defines rather than on new rule text the Commission would have to write.

PARTIAL. Test 4, can the customer take the signals directly, or is a middleman required? The price passes and the demand response signal does not.

On price the draft is as good as anything California has produced. Retrieval occurs "through an application programming interface using open internet protocols," and transmitting a rate identifier "shall not constitute registration, authentication, device identification, persistent identity, session management, or telemetry transmission." That is the clearest statement in any California document that a device may receive a signal without entering a relationship.

On demand response the draft names no server and permits none to be chosen, so the question of whether an intermediary is required has no answer in the text. See Test 3.

PASSES. Test 5, is fail-safe behavior specified and mandated? Offline, the thermostat falls back to the most recently stored static dataset. If retrieval fails twice it continues on the last dynamic dataset, and after 48 hours reverts to the static one. On loss of communication the device must automatically attempt to rejoin. The static dataset is the right thing to fall back to, rather than merely an acceptable one: a time-of-use schedule is a published calendar rather than a real-time quantity, so a thermostat that has lost its connection can act on the stored copy and still be acting on the correct price.

PASSES. Test 6, does the customer who cannot respond end up better off? Autoflex is enabled by default at shipment. A customer who enrolls in no program, signs no contract and has no relationship with any third party gets a thermostat that responds to the stored MIDAS static dataset out of the box, and to live California prices and emissions as soon as they connect it, in both cases only within the Comfort Bounds the customer sets. Use of the dynamic dataset requires the consumer to opt in by enabling connectivity, which is the right place to put that choice. This is the most important design choice on the price path.

The draft specifies what the appliance must do with a signal, and not where the signal comes from. One defect accounts for every test it fails.

On what the appliance does, the Commission got everything right. The thermostat retrieves price data itself, stores it in its own memory, computes its own schedule, needs no enrollment, fails safe, and serves by default a customer who never engages. Every one of those requirements names the thermostat as the thing that must do it, so the decision ends up in the appliance

whatever server supplied the data. This is the architecture that makes classic demand response unnecessary for most purposes: a device holding a price and its owner's bounds does not need to be dispatched.

On where the signal comes from, it specifies almost nothing. The price source is an unnamed application programming interface, and the manufacturer's cloud is expressly permitted to supply the stored dataset. The demand response client's location is unstated and the server it connects to is unnamed. Read together, those silences permit precisely the architecture this workshop was convened to examine: an open protocol between a utility and an aggregator, and a proprietary one between the aggregator and the appliance the customer bought.

The Commission is one sentence away from closing it.

2. Write the capability once, not once per appliance

Every flexible appliance needs the same fact: what energy costs now and next. The limit on the connection is not an appliance signal. It goes to the device that enforces it and to an energy management system planning the premises, because a limit is usable only by something that can see the whole site's draw. There is nothing appliance-specific about either one.

The capability is a price and a limit. A rule written now should not build a third mechanism beside them.

The draft thermostat standard builds two paths and treats them as co-equal. Autoflex is computed in the device from a price, is on by default, and needs no counterparty. Third-party demand response is an event, delivered by a party, that the consumer opts into.

What those provisions get right is not the event; it is the bounding. Section 1693(c)(2)(F) requires express consumer authorization, constrains event setpoints to the consumer-defined Comfort Bounds, limits an override to the duration of the event, and restores the prior schedule automatically unless the consumer chooses otherwise. Those are the right protections for any party acting on a customer's behalf, and they should survive whatever the Commission decides about the event construct itself.

The architecture is the part to reconsider, and the word event does not mark the distinction. In OpenADR 3 an event is the carrier for everything, a price included. What matters is what it carries. An event carrying a price, an emissions intensity, or a reliability condition such as a Flex Alert is information, and a device holding that information and its owner's bounds already has what it needs to decide. A demand response operating schedule is different in kind: Section 1691 defines it as an operating schedule sent by an authorized entity "that adjusts how and when associated equipment operates," and that is the thing a price replaces.

The cost of leaving the two co-equal is that every later appliance standard inherits the shape: a price path and a dispatch path, each with its own server question and its own silence about where the client lives. Section 1 is what that cost looks like on the dispatch path.

The Commission is currently building a rule per appliance category, and the provisions examined in Section 1 show why that is the wrong unit. Section 1693(c)(2)(D)1.a, which obliges the appliance to retrieve MIDAS data on an interval, and Section 1693(c)(2)(B)1.e, which obliges it to hold the software that turns that data into a schedule, say nothing about thermostats. They

would apply word for word to a water heater, a clothes dryer, an electric vehicle charger or a home battery. On the present course they will be drafted again for each, in a separate proceeding, each one going stale on its own schedule.

There is a route from here that does not reopen anything. The thermostat standard is in Express Terms now, and its appliance-independent provisions are already drafted. Marking which provisions are appliance-independent, and providing that they are superseded by the common standard once adopted, asks staff to label rather than to redraft. The next appliance standard then references the common provisions instead of copying them, which is the whole of the saving.

The Commission has built the other shape before. The Load Management Standards are one rule covering every large investor-owned utility, the large publicly owned utilities and the large community choice aggregators, not a separate proceeding per entity.

An appliance needs two interfaces, and only one of them is the same for every appliance. The first is signal reception: receive a price, in a protocol written once for every appliance category. The second is functional control: be commanded by its owner, or by an agent its owner authorizes, in the appliance's own vocabulary. A thermostat has setpoints and comfort bounds, a charger has a current limit and a departure time, a battery has charge and discharge. Those vocabularies are device-specific and cannot be written once, which is a reason to put them in the appliance's appendix alongside its test method rather than a reason to write a separate rule per appliance.

Both are needed, and for different reasons. An appliance with signal reception alone optimizes itself but cannot be coordinated with the rest of the house. An appliance with functional control alone is something somebody else drives. Together, the appliance decides for itself by default and can be directed when the house needs coordinating, which is also how a power control system tells it what it may draw.

Recommendation. Open one common flexible demand appliance standard stating the capability for every appliance category, and provide that the capability is stated there and nowhere else. The appliance-specific rulemakings continue, limited to the test method and the control vocabulary for their appliance. That reaches thermostats, water heating, electric vehicle charging and home batteries at once, it answers the interoperability question in these comments once rather than once per appliance category, and it does not disturb a rulemaking already in Express Terms.

This is less regulation, not more. One capability written once replaces the same capability drafted again in every later appliance proceeding, and the protocols it would name are published and shipping today.

3. Aggregation is a service; a rule should not presume an intermediary

Commissioner McAllister returned to the role of the aggregator twice, and asked separately how community choice aggregators "can be a force for good and align with what we're all trying to do." Those are two questions, and the word they share hides the difference.

A demand response aggregator is an intermediary that operates customer equipment. What it aggregates is devices for dispatch, and it sits between the signal and the appliance. This is the actor these comments are about, and the one the Commissioner's question about the aggregator's role concerns.

A community choice aggregator is a load-serving entity. It procures energy, sets a tariff and bills a customer, in place of the incumbent utility. What it aggregates is customer load for procurement. The Commission's own draft Express Terms in Docket 24-FDAS-03 say so, defining "load-serving entity" as an electrical corporation, an electric service provider, or "a community choice aggregator as defined in Section 331.1 of the Public Utilities Code." A CCA is therefore on the same side of the line as a utility: a sender of signals, where a demand response aggregator is a layer between the customer and the signal. Nothing in what follows is specific to community choice aggregators. It applies to any retailer of electricity, which is the term used for this role in other markets, and California's term is the one that happens to contain the word.

The intermediary question

The structural answer is that **device aggregation should be a service a customer opts into, not a layer a customer must pass through.** The Commission has already written that principle into draft rule text: the thermostat Express Terms in Docket 24-FDAS-03 provide that third-party demand response functionality "shall require express consumer authorization to opt in." An opt-in is only meaningful when declining it still leaves the customer able to act on the signal.

When the open interface reaches the appliance, an aggregator earns a customer by being better at optimizing than the alternatives, and loses that customer the day it stops being better. The customer may also decline aggregation entirely and respond to the signal themselves. Competition runs on performance.

When the only interface is between the grid and an aggregator platform, the aggregator holds the customer's access to the signal. Switching costs the customer the signal itself. Competition runs on who holds the relationship.

Independent aggregators benefit from this. An aggregator that does not make devices has one hard problem: every manufacturer exposes a different interface, so supporting a new device means new engineering and a new commercial agreement with that manufacturer. One interface at the appliance replaces all of it, and replaces it with something no manufacturer can withdraw. Panel 2 made the point from the manufacturer's side. Speaking in the voice of an original equipment manufacturer, the interim chief executive of the Mercury Consortium said that what such a manufacturer wants is "a consistent way of talking to the thing that's going to control them," and that flexibility providers are looking for the same consistency in return.

An OEM that is also an aggregator will object, and that is the argument for requiring it. Where the manufacturer and the aggregator are the same company, a proprietary interface is not a cost to be eliminated. It is the thing that makes its installed base its own. An open local interface on that company's own product invites competitors to serve customers who have already bought its hardware. Faced with that, a vertically integrated manufacturer will rationally decline, and will keep offering its customers a single program operated by itself.

This is the whole case for a standard rather than an encouragement. The parties who would have to volunteer are precisely the parties who lose by volunteering. Voluntary adoption selects for the firms with least to protect, and the record shows it. AHRI 1380 permits a local interface, and conformance to it is self-declared: the standard provides that conformance "shall not be claimed or implied" unless a product meets its requirements, and names no certifying body. My assessment of the market is that essentially no product built to AHRI 1380 implements that local interface. The IEEE 2030.5 Common Smart Inverter Profile certifies servers, clients and aggregator clients, and its implementation guide places "[c]ommunications between the Aggregator/GFEMS and its managed DERs or communications within the DER" out of scope. So the profile standardizes the link from the utility to an aggregator or site manager, and says nothing about how either one then reaches the equipment it operates. The protocol can run on the inverter itself, and a small number of certified models do so, but that is the exception rather than the pattern. In each case the permissive option was available and almost nobody took it.

Pacific Gas and Electric's own emerging-technology program describes the resulting status quo. Reporting on a 2026 field study of standards-based smart home automation, it states that controlling customer devices "typically requires use of proprietary solutions that enable access to a single device or brand and are subject to device access and subscription fees," that including "multiple devices and brands into a VPP requires integration across multiple OEM platforms with custom implementations that must be managed and maintained," and that because such solutions "are subject to change or may be discontinued at the discretion of the OEMs," these factors "introduce the potential risk for enrolled devices to become stranded assets." A utility is describing the cost it absorbs and the risk its customers carry.

AHRI 1380 is the closest precedent for the defect identified in Section 1 of these comments, and it is an HVAC standard whose control point is the thermostat. A rule that permits the local interface without requiring it produces cloud-only products. The Commission does not need to assume bad faith on anyone's part. The opposite assumption is the one that matters: a company's officers and directors owe duties to the corporation and its shareholders, and are accountable to its owners for financial performance. Opening an installed base to competitors is hard to justify in those terms, and a manager who did it voluntarily would have to defend the choice against them. Expecting firms to volunteer is therefore expecting them to act against the interest they are accountable for advancing, which is not a thing that happens at scale. A mandate removes the conflict, because it applies to every competitor at once and asks no firm to give up ground the others keep.

Conversely, a standard that only standardizes the grid-to-aggregator link does not solve the problem the Commission named. The workshop schedule posted in this docket describes standards that "address different parts of the communication ecosystem but are not yet aligned into a consistent end-to-end framework." Alignment that stops at the aggregator's cloud leaves the last link, the one that actually moves load, as a per-vendor integration. Both links should be standard. A customer who wants an aggregator to operate their appliance should have that, and a customer who wants to operate it themselves, or to hire someone other than the manufacturer to, should have that too. Standardizing only the first forecloses the second.

A party outside the premises optimizes on less information. What the site's own generation is doing, what its storage holds, what else is drawing, and what its owner has asked for are all known locally and none of them is known to a server elsewhere. A party operating the appliance

from outside can recover them only by exporting the site's state to itself continuously, which is a heavier integration than the one it replaces and a disclosure the customer did not ask for. The aggregation results in California are real; this is not a claim about performance. It is a claim about where the information sits, and it is why an open interface at the appliance produces better answers rather than only fairer terms.

The CCA question

A CCA should publish signals rather than operate devices. It sells energy under a tariff and owns no customer equipment, which is both why it is well placed to be the force for good the Commissioner asked about and why device aggregation is the wrong business for it to be in. What it needs is a direct path to its own customers: the ability to publish its own prices and program signals to equipment the customer can point where they choose.

Without that path, a load-serving entity has to stand up or rent a device aggregation platform to reach customers it already bills, and inherits an intermediary's costs and device list in order to do what its tariff relationship should already allow. That is the position a CCA is in today, and it is an artifact of the missing path rather than a choice.

A CCA has reason to want this sooner than an incumbent does. CCAs compete for customers on what they can differentiate: cleaner supply, program design, local priorities. A price published directly to a customer's equipment is how those choices reach a thermostat instead of a brochure. A CCA is also smaller than the utility whose wires it uses, with less process between a decision and a deployment, so it can stand up a server of its own sooner and change it more readily. The alignment the Commissioner asked about is already there in the incentives. What is missing is the path.

A capacity limit is not the CCA's to publish. The constraint belongs to the party that owns the distribution equipment, which in a CCA's territory remains the incumbent utility, and the recommendation on envelopes below is addressed to that party.

Recommendation. Where the CEC requires a flexible demand capability, it should require it at the appliance over an open protocol the customer can point where they choose, and should neither require nor presume a device intermediary. Device aggregation then competes as a service, and load-serving entities, including CCAs, can reach their own customers directly.

4. A manufacturer should not withhold from the owner what it grants to an aggregator

The aggregation results in California are real and they are large. On July 29, 2025, from 7:00 p.m. to 9:00 p.m., several virtual power plant aggregators discharged portfolios of home batteries for an average output of 535 megawatts through the Emergency Load Reduction Program and the Commission's own Demand Side Grid Support program, across the service territories of all three large investor-owned utilities. Pacific Gas and Electric described it in its newsroom as "the largest test of its kind ever done in California—and maybe the world." Commissioner McAllister said of it:

> VPPs deliver for the grid! This test embodies California's innovation economy by applying digital tools to access and aggregate untapped electric resources - supporting grid reliability and helping manage electricity costs. We're all connected to the same grid, and small contributions from many individuals add up to hugely important, shared benefits. Even bigger, better things to come!

That assessment is right, and these comments share it. Hundreds of thousands of customer-owned devices, told what the grid needs, did something enormously valuable. Small contributions from many individuals did add up to shared benefits. This is grid coordination working at scale, and it is the strongest evidence California has that the approach this workshop is considering will deliver.

The architecture that produced it does not match what the Commission says it wants. That 535 megawatts was delivered through an arrangement that does not yet reflect several commitments this Commission and the Public Utilities Commission have made elsewhere, and the two have not yet been brought into line.

The commitment	Where it is stated
Consumer-authorized control	This workshop's notice, describing an approach that "enables consumer-authorized control"
Freedom from proprietary restriction	The Assigned Commissioner's and Administrative Law Judge's ruling of July 7, 2026 in R.21-06-017, asking "[w]hat hardware, software, or data sharing safeguards may be needed to ensure that customers can obtain a flexible connection without being restricted by proprietary arrangements?"
Open access and interoperability, technology and provider neutrality	Decision 26-10-054, as modified by Decision 26-11-006, adopting eleven guiding principles including "[o]pen access and interoperability" and "[t]echnology-neutral, performance-based"
No dependence on a single manufacturer's	The draft thermostat Express Terms, which

The commitment	Where it is stated
service	forbid a thermostat to “rely solely on a manufacturer or third-party service”

A customer whose battery participated in that July event did so through one manufacturer’s program, using one manufacturer’s software, on terms that manufacturer set. That is not a criticism of the program, which worked. It is an observation that the result was achieved without the architecture the Commission says it wants, and that the next increment of this capability should not have to be.

The same megawatts have a price-shaped route. An export price worth acting on, published to everyone on the tariff, reaches every battery on the system, including every one enrolled in nothing. Each decides for itself, through whatever energy management its owner chose, whether to export and how much. The grid gets the energy, the customer keeps the decision, and no manufacturer stands in the path. That is the same coordination delivered by a signal rather than by a command, and it is what these comments ask the Commission to make possible.

Saul Griffith put the open question more bluntly on the Volts podcast in December 2025, speaking about the Australian market:

> You have to ask yourself, why do I need aggregation? All the VPP does is charge you a fee so that somebody who can write software can aggregate your home with a whole bunch of other homes to trade the grid. I don’t believe that step is necessary. It adds a gatekeeper and it adds a rent seeker. If you had a well-designed electricity market that was fair, this stuff would just happen automatically at your smart meter or through your retailer or utility company. …

Aggregation delivers value today. Two features of the present arrangement contradict the commitments above, and both are within this Commission’s power to fix through appliance standards. A third observation about the structure of today’s arrangement follows them.

There is also a distributional point that favors the published signal on its own terms. Virtual power plant incentives are funded through rates, so every customer pays for them, including the customer who is not enrolled and cannot be. A price and a limit published to everyone are free to read, available on the same terms to everyone, and need no enrollment, no credit relationship and no particular brand of device. The price is identical for everyone on a tariff; the limit is identical in kind while differing in quantity from one connection to the next. Pacific Northwest National Laboratory’s DSO+T study found annual utility bills falling 14 to 16 percent for participating residential customers, and 10 percent for residential customers who stayed on a fixed-price tariff and did not participate at all, in its moderate renewable scenario.

The capability the owner does not have

A customer who buys a home battery owns the hardware outright. Whether they can tell it when to charge or discharge is a separate question, answered by the manufacturer rather than by any standard.

Taking the Tesla Powerwall, the most common battery in that test, as the example, the published record shows three things. The Commission should confirm all three with the manufacturer before relying on them.

There is no documented local interface. The gateway does expose an interface on the home network, and a customer can reach it. It is not published, and the tooling that uses it was reverse-engineered by third parties. A widely used open-source project for local Powerwall control, pypowerwall, documents local control of the operating mode, the backup reserve, grid charging and grid export, but implements it by writing to the gateway's internal configuration store, and records that from firmware version 25.10.0 Tesla no longer supports network routing to that endpoint to the gateway's own access-point address. A customer directing their own battery from their own network is therefore relying on undocumented behavior that a firmware update can remove, and that some firmware has already restricted.

Local control sets a standing preference, not a dispatch. Setting an operating mode or a reserve percentage states a standing preference. It is not the same as telling the battery to import or export a given amount now, which is what responding to a grid condition requires, and which is the capability an FDAS is ultimately about.

The documented remote interface withholds dispatch from the owner and grants it to aggregators. The manufacturer's cloud API is documented and supported, and it carries mode and reserve settings. It does not carry a discharge command. Third-party integrator documentation states that there is "no direct API command to force discharge a Powerwall," describes this as "a Tesla Fleet API limitation," and states that "only Virtual Power Plants (VPPs) have the ability to force discharge a Powerwall through their special API access."

An owner can reach the capability, but only outside anything the manufacturer documents. The same open-source project cited above implements a timed charge and discharge command against the gateway, and warns that it is undocumented and may void the warranty. So the capability exists in the hardware, is granted to aggregators through a supported interface, and is available to the owner of the asset only by means the manufacturer disclaims.

The architecture is explicit about who the customer is. The owner of the asset cannot command it to export. An aggregator can, through privileged access the manufacturer grants. The capability exists; it is reserved.

The Commission has already been asked the question in a neighboring proceeding. The Public Utilities Commission's ruling of July 7, 2026 in R.21-06-017 put the proprietary-restriction question quoted above to parties. The same question applies to flexible demand appliances, and this Commission is better placed to answer it, because it sets standards for the appliance itself.

Recommendation. An FDAS should require that every power and energy management capability an appliance implements be available to the owner of that appliance, and to an agent the owner authorizes, on terms no worse than those on which the manufacturer makes it available to itself or to any third party.

The condition also reaches the case where the manufacturer operates the virtual power plant itself. Tesla described its own position to this Commission in those terms, in comments on the

2026 Draft Demand Side Grid Support Program Guidelines, stating that its interest in the program includes “our role as an aggregator that has enrolled and dispatched tens of thousands of systems in the program, as an installer of solar and battery systems, as well as a manufacturer of the battery systems that are integral to the capacity services that DSGS is intended to deliver.” A capability granted to the manufacturer’s own aggregation business and withheld from the owner of the hardware is self-dealing in a market this Commission is being asked to open, and the customer who bought the battery is the only party denied what it can do.

The condition matters in both directions. A rule written only against capabilities a manufacturer grants to others would be satisfied by a manufacturer that grants the capability to nobody and reserves it entirely, which is the more complete exclusion of the two. A customer who has bought a battery outright should not discover that the manufacturer’s own software can direct an asset they own and they cannot.

A manufacturer may decline to implement a capability at all, and may withhold functions that exist for safety, for certification integrity, or for compliance with interconnection requirements. Those are real categories and a standard should say so. Dispatching a customer’s own battery within the customer’s own limits is in none of them.

A binding limit, an advisory price, and who decides inside the envelope

Flexibility policy is beginning to describe a capacity limit, a dynamic price, a third-party dispatch and a customer’s own instruction as though they were four signals competing for the same device, to be ranked against one another. They are not the same kind of thing, and ranking them is the wrong model.

A capacity limit binds. It exists to keep a service drop and a distribution transformer within their ratings, and nothing on the customer’s side of the meter is entitled to exceed it. It is not a signal weighed against other signals; it is the envelope everything else operates inside.

A price is advisory. It carries no authority over anything. It is information that lets whoever is deciding shift energy use in time, and it cannot command a device or raise a limit. A rule that puts a price in a precedence list alongside a limit has already made a category error.

A third-party dispatch is not a separate authority. A customer who authorizes an aggregator to dispatch a battery has delegated the customer’s own authority, bounded by the terms of the delegation. The aggregator inherits exactly what the customer gave it and nothing more, and its dispatch is subject to the same import and export limits as anything else the customer directs. There is no case in which an aggregator’s instruction outranks the envelope, because the aggregator is standing in the customer’s shoes.

So inside the envelope the customer decides. A customer’s instruction governs, whether given directly or through an agent they chose, until it would breach a limit that may not be overridden.

That resolves the cases that look like conflicts. A virtual power plant calling a discharge in an hour when a flexible service connection limits the site cannot have what it asked for; the limit was never optional. An aggregator that moves a water heater outside the hours or the temperature range its customer agreed to has exceeded its delegation, and the remedy is the terms of that delegation rather than a rule about precedence. Neither case needs a precedence table; both need

the standard to say which instrument is a constraint, which is information, and whose authority an agent is exercising.

The Commission has already drafted bounded delegation correctly, in the course of drafting something else. The thermostat Express Terms say what happens when a third party acts on a customer's appliance: a demand response event "[t]emporarily override[s] the active autoflex operating schedule or consumer-defined schedule solely for the duration of the event"; the thermostat "shall constrain temperature setpoints associated with demand response events to the consumer-defined Comfort Bounds"; and on conclusion it "shall automatically restore the operating mode and temperature schedule that were in effect immediately prior to initiation of the event unless the consumer selects otherwise".

Three sentences settle who may act, for how long, within whose limits, and what happens afterwards. The authority is the customer's, exercised by someone else, bounded by what the customer set. That is the shape the capacity case needs too.

Recommendation. An FDAS should state what each kind of signal obliges an appliance to do. It should provide that a capacity limit on the service connection constrains all operation, including operation a customer or the customer's authorized agent directs; that a price and comparable signals are inputs to optimization and never commands; that within the limit the customer's instruction governs; and that a third party acting on a customer's authorization has the customer's authority and no more, bounded by the terms of that authorization.

No appliance should be required to implement capacity logic. A limit on the service connection is held by one device: the power control system (PCS) listed to UL 3141, because the whole site's draw is visible only there. Two enforcers would conflict. Every other appliance can comply by responding to a price within its owner's bounds, and a rule that obliged a dishwasher to reason about a transformer would add cost without adding protection.

The limit goes to two places, not everywhere. The device that enforces it needs it. An energy management system planning the premises needs it, and is not thereby responsible for holding the site under it. An appliance that knows the site limit but not the site load cannot act on it, and should instead be told what it may draw, in its own vocabulary, by whichever of those holds the plan. Where that system publishes a local price, reflecting the tariff price together with on-site generation, storage state and the headroom it is managing, appliances optimize against that rather than against the tariff alone. So the Commission should define how a limit is represented, once, and say where it is delivered. Receiving the limit requires no listing. Enforcing it does.

That division also decides who carries the certification burden. One listed device is tested against the limit it enforces. Everything else is tested against a price it reads and bounds its owner set, which is the same test for a water heater, a heat pump and a dryer.

Often three companies must agree before one appliance moves

Today a utility signal commonly reaches a residential appliance by passing through an aggregator platform and then a manufacturer's cloud. Where those are separate businesses, which is the usual case, three commercial parties must hold agreements with one another, and all three must remain in business and keep their interfaces working, before one water heater changes what it is doing. Where the manufacturer runs its own aggregation, the chain is shorter by one

link and the customer's position is worse, because the single party now holds both the device and the route to it.

Every one of those links is a place where the chain breaks, and none of them is visible to the customer when it does. The arrangement also sets the price of entry for anyone new: a competitor must negotiate with manufacturers one at a time, which is why an aggregator's supported-device list reflects the agreements it has reached rather than the devices a customer owns.

An open local interface at the appliance does not abolish that chain. It makes it optional.

5. Require utilities to publish prices to MIDAS over OpenADR 3

The Load Flexibility Technical Lead described the MIDAS price server at the workshop. Two of the design choices she reported are rare, and I support both.

OpenADR 3 outbound, to the devices that read MIDAS. "We're working on OpenADR 3 compatibility. It's currently in alpha testing, and as soon as we get that dialed in[,] we'll make it … usable for everyone." A public price source speaking a published, royalty-free protocol is exactly what an appliance standard requiring an OpenADR 3 VEN needs on the other end.

No device registration, no tracking. "It's free to use, there's no cost to the utility, there's no cost to the user, there's no user device registration required. This is just a system to get prices. Everything in this is public information… We don't track who's using it."

That second one is the more significant. A great deal of effort across this industry goes into deciding who may receive a signal, and the MIDAS server simply publishes it. A device can act on California electricity prices without an account, a contract, a credential or a counterparty. The Commission should hold onto that property as the MIDAS server gains new capabilities, because it is far easier to keep than to recover.

The MIDAS server should read prices as an OpenADR 3 client

The MIDAS server acquires its data on a different basis than it serves it. As described at the workshop, "utilities upload their prices," through an interface specific to MIDAS, which means each utility builds and maintains a custom integration that serves one consumer of the data.

I urge the Commission to require utilities to publish their prices over OpenADR 3, and to have the MIDAS server obtain them as an OpenADR 3 client.

The custom integration disappears. A utility implements one published protocol rather than a MIDAS-specific upload, and the Commission maintains a client rather than an ingestion interface.

One protocol carries the data in both directions. The same specification the Commission is considering requiring in appliances would carry the price from the utility to MIDAS and from MIDAS to the appliance. Every participant implements the same thing once.

The utility ends up operating a VTN. A utility that publishes its prices over OpenADR 3 is operating an OpenADR 3 VTN. Having done that, it can serve its own customers directly from

the same infrastructure, at no additional protocol cost. The work is the same work; what changes is that it is reusable. The MIDAS server becomes one client among many rather than the only route by which a utility's prices reach anyone.

The same VTN serves the appliance requirement. An appliance required to carry a Virtual End Node needs a Virtual Top Node to reach, and a customer-configurable server address is worth little if the only server on offer is the manufacturer's. The VTN a utility operates to feed MIDAS is the one its own customers should be able to point their appliances at for a price, and the one over which it should publish the capacity limit to the device that enforces it. These are not separate obligations; they are uses of one piece of infrastructure.

One dependency should be pulled forward. Section 1623(d)(2) of the Load Management Standards requires each large investor-owned utility, within forty-five months of April 1, 2023, to offer every customer voluntary participation in a marginal cost-based rate, or, if no such rate has yet been approved, in a cost-effective load flexibility program. That falls on January 1, 2027. Section 1623.1 places the same obligation on each large community choice aggregator within fifty-one months, which falls on July 1, 2027.

An appliance that retrieves MIDAS data still has to determine which rate identifier applies to the premises where it is installed. That is the job of the single statewide standard tool required by Section 1623(c)(1), which must "[p]rovide the RIN(s) applicable to the customer's premise(s) to third parties authorized and selected by the customer" and "[b]e accessible in a digital, machine-readable format …" The tool was due to be submitted within eighteen months of April 1, 2023. The Commission has since extended that deadline to May 8, 2029 for nineteen load-serving entities, including all three large investor-owned utilities, in Orders 26-0312-03b and 26-0513-03a.

A thermostat shipped under the standards proposed in Docket 24-FDAS-03, which the draft staff report proposes to take effect on January 1, 2028, will therefore have a published price it can read and no supported statewide way to learn which price is its own. Requiring utilities to publish over OpenADR 3 helps here too, because the tariff a customer is on is something the utility knows and can serve directly to equipment that customer has authorized.

Operating envelopes point the same way

Whether the MIDAS server might also carry operating envelopes came up at the workshop, properly qualified as depending on how those envelopes are designed. That qualification is the right one, and the design dependency is decisive. An operating envelope, which is the capacity limit of Section 4 as the utility publishes it, differs from a price in ways that decide who can publish it.

It is per-customer. A price is a public attribute of a tariff, shared by everyone on that rate. An envelope is a limit on one service connection, derived from the capacity of the transformer and conductor serving it. There is no "program" to join; there are two numbers per customer, one for import and one for export.

It therefore requires authentication, which is not the same as enrollment. A customer's operating envelope is customer-specific, so the recipient has to be identified, and the MIDAS server's property of serving anyone anonymously cannot carry across. Identification is not the

obstacle it is sometimes made to be: the customer already has a billing relationship with the utility that holds the constraint. Authenticating to one's own utility with the credentials of that relationship adds no program to join, no third party to contract with, and no device certificate that only a manufacturer can provision.

An authorization requirement is the point at which an intermediary is most easily inserted, and a rule that requires authentication without saying what satisfies it will be met with a program to enroll in, an aggregator to sign with, or a private certificate authority.

And only the distribution operator holds it. The envelope exists to protect specific distribution equipment. The distribution operator is the only party that knows the constraint, the only party that can compute the limit, and the only party that can change it when conditions change.

And the limit has already been ordered. A UL 3141 power control system holds a whole site under a number, and Decision 26-02-025 of the Public Utilities Commission established a standard offer for flexible service connections, putting a scheduled per-customer import limit into a California tariff. That decision adopted static flexible service connections and deferred connections providing non-firm capacity, the limits that change with real-time or forecast grid conditions, to R.21-06-017. A per-customer limit therefore exists in California today, and the limit that responds to conditions is still being decided. What does not exist is a standard path by which it reaches the equipment in the house.

A proprietary one has been demonstrated. Pacific Gas and Electric's Demand Response Emerging Technologies program lab-tested a dynamic service rating sent from a vendor's fleet-management interface to a listed panel. The panel received the temporary limit and enforced it locally against the customer's own circuit priorities and always-on setting, and the report records that the function "operates locally with no dependency on cloud communications, other than for the customer to modify preferences." So the capability is proven and the architecture works. The path tested is one vendor's interface to that same vendor's panel, which is the per-manufacturer integration these comments object to everywhere else. The gap is in standardization rather than in capability, and it is the gap an OpenADR 3 requirement on utilities would close.

Recommendation. Where an FDAS requires an authenticated channel to a utility, it should provide that the credentials of the customer's existing utility account, or a credential the utility issues to that customer on request, are sufficient to establish it. It should further provide that conformance may not be conditioned on enrollment in a demand response program, a contract with a third party, or a device certificate obtainable only from the appliance manufacturer, and that a customer may authorize equipment of their choosing, including equipment they did not buy from the utility or the appliance manufacturer, to use that credential on their behalf.

So the envelope has to originate with the utility and be delivered to an authorized recipient. If utilities are already operating OpenADR 3 VTNs to publish prices, the envelope travels the same way, over the same protocol, from the party that owns the data, to a customer or an agent the customer has authorized.

That is the same conclusion from the other direction: **the destination is utility-operated OpenADR 3 service.** Prices get there first because they are public and easy. Envelopes follow because they must come from the same party, and because by then the infrastructure exists.

6. Mandate open, standard, local, IP-based protocols. Legacy bridging is a product decision

Schneider Electric described what its energy management system does: “there was a need, essentially, within the home for some fashion of a Rosetta Stone. You really have to have, essentially, a bridge between the different protocols. And so, as we have been developing our EMS, it’s one of the core technologies that we have worked on, is essentially protocol translation.”

The bridge belongs in the product and the protocol belongs in the standard. An energy management system that bridges Modbus, OCPP, BACnet and a dozen vendor APIs is doing valuable work, and a product that chooses to do it serves customers with equipment already on the wall. A rule should name what a new appliance must speak and leave the bridging to the products that find a market in it.

Modbus is the clearest case. It is a 1979 serial protocol with no discovery, no device model and no notion of identity, whose security was added later as a separate protocol rather than built in. There is an enormous installed base, particularly in inverters, and an EMS vendor that declines to speak it is leaving customers stranded. All of that is a reason for products to support Modbus. None of it is a reason for California to require a new appliance to offer it.

An FDAS should name what a new product must speak, not what it may additionally speak. The criteria that make a protocol suitable for the mandate are testable:

- **Over IP**, which settles the radio question and costs nothing on a device that already ships Wi-Fi and a microcontroller
- **Specification published and obtainable**, so a test engineer can apply it
- **Carries power and energy management, control and reporting**, which is the functional requirement
- **A conformance test a third party can run**, without which compliance is the vendor’s own assertion
- **Reachable locally, with no vendor cloud required in the path**, which is what survives the manufacturer turning the cloud off

The Commission should name protocols, and should name every one that meets these criteria on the day the rule is adopted. A rule that states criteria and names nothing leaves each manufacturer to assert that its own protocol qualifies, and leaves the Commission adjudicating that one product at a time with no list to point at. Naming is what makes the requirement testable immediately. The Commission has already done this once, naming OpenADR 3.1 in the thermostat Express Terms, and that was the right call.

Name the standard, not the point release, and say which version applies. The Express Terms currently name “OpenADR 3.1.” A rule that names a point release makes a device built to a later version of the same standard non-conforming, and obliges the Commission to reopen the rule to permit an improvement it would welcome.

The Commission has already solved this in the same document, for a different artifact. Section 1693(c)(2)(B)1.a requires the stored MIDAS static dataset to be, “at the time of final factory

software image… the most recent version available from the Energy Commission," and provides that it "may be updated and overwritten." The version is not frozen in the rule; the obligation is to be current.

The same sentence works for a protocol. The rule should name **OpenADR 3**, and provide that an appliance shall conform to the most recent published version of that standard at the time of its final factory software image. The current version is 3.1, so a thermostat built today would conform to 3.1 and would carry the local discovery mechanism Section 1693(c)(2)(A)2 should cite. A thermostat built after the next release conforms to that one, without a rulemaking.

Anyone should be able to put a further protocol forward. The list belongs somewhere the Commission can add to without a new rulemaking, and the route to it should be open rather than by invitation. A manufacturer, a standards body, a utility or any other party should be able to submit a protocol with a showing that it meets each criterion above, including a published specification and a conformance test a third party can run. The Commission should decide within a stated period and publish the result with its reasons, so that an addition is visible and a refusal is reviewable. A petition route that anyone may use is also the answer to the objection that naming protocols picks winners: the list is open, and what closes it is failing the criteria rather than lacking a sponsor.

The Commission already works this way elsewhere. Section 1623(c)(5) of the Load Management Standards provides that changes to the statewide rate tool "shall be submitted to the Executive Director for approval," and that the Executive Director "shall submit any substantive changes to the Commission for approval at a Business Meeting." The draft thermostat Express Terms take the same approach to the MIDAS static dataset, requiring the stored copy to be "the most recent version available from the Energy Commission" and providing that it "may be updated and overwritten," rather than fixing a version in rule text. A list of qualifying protocols can be maintained on the same footing.

Protocols meeting these criteria ship today, so the requirement waits on nothing, and a California utility has tested the combination end to end. In the same study, Pacific Gas and Electric ran OpenADR 3 from the grid to a local controller in the home and Matter from that controller to the devices. Its conclusion on the grid side is that prices and demand response signals "were delivered securely without custom implementations," and that for "common, well-established use cases such as price and DR signaling," the design "supports strong interoperability between compliant, version-aligned OpenADR 3 VTNs and VENs."

The study is equally clear about what did not go smoothly, and the Commission should weigh it. Delivery of control events to the devices required customization, device setup required technical support, several participants hit router and connectivity problems, device state reporting was inconsistent across product types, and the larger loads that matter most for flexibility, naming electric vehicles, stationary batteries and heat pumps, had not yet adopted Matter. None of that bears on whether the protocol qualifies under the criteria above. All of it bears on how long the device side will take to mature, which is a reason to write the criteria now and let the appendix fill as products qualify. Matter is the clearest case: published by the Connectivity Standards Alliance, running over IP on the home network with no vendor cloud in the path, with a third-party certification program, and already carrying each of the signals this filing asks the Commission to require: price reception in its Commodity Price cluster; the service connection

limit in the PowerThreshold feature of its Meter Identification cluster, which reports the maximum instantaneous demand that may be distributed to a customer without risk of overload; and functional control in its Device Energy Management cluster. The first two were published in Matter 1.5 in November 2025. Matter itself ships in consumer products now, in the ecosystems households already own.

The thermostat Express Terms show the protocol half, and the gap

The draft Express Terms published in Docket 24-FDAS-03 on November 10 do most of this correctly. They require every thermostat to conform to the OpenADR 3.1 Virtual End Node requirements, without saying where that node runs, and they provide that a thermostat may not rely solely on a manufacturer or third-party service whose discontinuation would disable required demand response functionality. That second clause is the protection customers have lacked.

They also show the gap this comment is about: the draft requires a VEN and never says which VTN it connects to, or who decides. See Test 3 in Section 1.

Recommendation. Whatever protocol the Commission names in any FDAS, it should require that the server address be configurable by the customer and that the device attempt discovery of a local server. Without them, the protocol is open and the relationship is not.

7. A certification mark is not interoperability

The Mercury Consortium's interim CEO presented on Panel 2, describing a mark that will run "right the way across the flexibility stack" and that "will allow consumers to know that their devices will be able to be entered into different services."

That is a worthwhile objective and certification is a legitimate instrument. The caution is specific: **certification can only deliver interoperability when the specification and the test procedure are public.** Mercury's are not, and the gap between what is advertised and what is published is large enough that the Commission should not treat the Mark as evidence of the properties it is describing.

Mercury's EVSE specification set comprises three documents. Two are published:

- REQ: EVSE-001-2025-REQ, EVSE Functional Requirements v0.9a, May 2025, 15 pages
- APG: EVSE-001-2025-APG, EVSE Application Guide v0.9, May 2025, 13 pages, which describes itself as "an informative application guide associated with the Mercury Consortium technical requirements" (APG p.3)

What the published specification does well

The control surface is a limit on a schedule, and it is specified with real precision: up to 96 schedule entries, time resolution of "1 minute or better," limit resolution "configurable to at least 1% resolution," daily and weekly recurrence, actuation accuracy within 10% of the expected PWM duty cycle, and a rule that a dispatch schedule "shall take precedence over any Default Schedule that may be present" (REQ pp.9 to 11). This is the right shape for a capacity instruction and it is better specified than most.

Measured against the six tests

FAILS. Test 1, does the grid send an objective, or operate the appliance? The test asks where the instruction is addressed. In OCPP a charging profile travels over a per-charge-point connection, from a management system to one named charger at one premises. Mercury maps to OCPP 1.6-J, and the specification's own framing is a management system addressing a station it manages (REQ p.13), though the mapping subsections are presently headings only, "under review with the Open Charge Alliance (OCA) before its release." That is a command to a named appliance in a named house rather than an objective published to everyone on a tariff. A command is not itself the defect. A customer's own energy management system commanding a charger in that customer's house is the capability Section 4 asks the Commission to require. The defect is a command addressed to one named appliance from a server its owner did not choose.

FAILS. Test 2, does the open interface reach the appliance, or stop at the manufacturer's cloud? Mercury's own account of how the first devices were certified is that a working group "built our most recent specification to support existing OCPP architecture using the OEM's own cloud as a comms relay, so devices already shipped don't need a hardware swap to qualify." The manufacturer's cloud is the design point.

FAILS. Test 3, can the customer choose who optimizes their home? The released specification provides that the EVSE connects to "the dedicated server, as described by the OEM" (REQ p.13). A path to another server exists, and its conditions are the point: Section 3.3.3 opens it only on manufacturer failure or seven consecutive days of cloud unavailability. Section 3.3.2 requires a local override so a local command can regain control, which is a power to interrupt rather than a power to choose. In normal operation the customer cannot select a different server.

FAILS. Test 4, can the customer take the signals directly, or is a middleman required? See Tests 2 and 3.

PARTIAL. Test 5, is fail-safe behavior specified and mandated? The specification addresses behavior on loss of connectivity and names a schedule to fall back to, which is more than most. The local enrollment path of Section 3.3.3 is triggered by "7 consecutive days of unavailability of the OEM cloud" (REQ p.12). Seven days is the right order of magnitude for an insolvency and the wrong one for a grid event. In a state that runs public safety power shutoffs, a charger whose schedule arrives through one company's cloud holds only the last schedule that cloud delivered, for a week, before any alternative path opens.

SILENT. Test 6, does the customer who cannot respond end up better off? Not addressed, and arguably out of scope for a device specification.

The specification forecloses customer choice in its normative text

This is the finding the Commission should weigh most heavily, because it is not an omission. It is a positive statement about who selects the server.

> The EVSE shall have the ability to connect to the upstream monitoring and management system and establish communication based on the communication protocol options stated in Table 8. This action may include the connection to the dedicated server, **as described by the OEM**. (REQ p.13)

The customer is not a party to that sentence. The specification does provide a path to another coordination server, and its conditions are the point. Section 3.3.3, “Fail Safe and Stranded Assets,” mandates “a fallback functionality in the event of an EVSE manufacturer removing support for an EVSE, closing cloud services, or entering insolvency,” under which “[t]he consumer shall be able to switch to another flexibility program and enroll their EVSE with a DER Manager after one of the above events,” and “[t]he EVSE shall allow this function to be performed locally without any connection to the OEM cloud.” The switch “shall happen automatically,” and “[t]he trigger shall be 7 consecutive days of unavailability of the OEM cloud” (REQ pp.12 to 13).

That is a well-drafted stranded-asset protection and the Commission should note that Mercury wrote it. It is not customer choice, and the specification’s own architecture shows why.

Mercury contemplates the right counterparty. Describing who may control a charger, the requirements document states that the DER Manager’s “role can be taken by an aggregator, a Charge Point Operator (CPO), a utility, an OEM, a Home Energy Management System (HEMS) provider” (REQ p.4). A home energy management system is on that list. The specification does not confine control to a cloud.

It then makes that unreachable in normal operation. The charger “shall have the ability to connect to an IP-based network with internet access,” and on establishing communication, “[t]his action may include the connection to the dedicated server, as described by the OEM” (REQ p.13). No provision lets a customer or an installer direct a working charger to a different DER Manager. The local enrollment path of Section 3.3.3 is the only route to one, and it opens on the manufacturer’s failure or after seven days without its cloud.

So the customer who owns a charger and a home energy management system, and whose manufacturer’s cloud is working normally, may not connect the two, although the specification names their energy management system as a party entitled to control the charger. The capability is described and gated behind the manufacturer’s misfortune.

That is the opposite of what the Mark is sold as delivering. Mercury’s public pages promise that “[a] Mercury-certified device works with any Mercury-certified utility, so you know the charger, battery or heat pump you choose will fit your provider,” and that for drivers the confidence “shows up in a charger that works properly with whatever tariff or smart charging programme it’s plugged into.” A device whose server is “described by the OEM” works with the utilities and programs its manufacturer has arrangements with.

The protocol is OCPP 1.6, which was built for commercial charging

The normative binding is explicit and narrow:

> The implementation must be using the OCPP 1.6-J variant, that is using JSON over web sockets. (REQ p.13)

The specification names “OCPP 1.6 as the first international protocol selected by Mercury” (REQ p.13), and the normative reference is to the 2017 edition.

OCPP was designed for commercial and public charging networks, where a charge point operator runs a back office and manages a fleet of chargers it owns or contracts for. That is a real and well-served problem. It is not the residential problem. My assessment of the market is that the share of residential chargers sold in the United States that implement OCPP is very small, and that the share which also let the customer configure the server the charger connects to is smaller still, approaching zero. Those are two different questions, and the second is the one that matters for customer choice: a charger may speak OCPP and still speak it only to the address its manufacturer pre-configured. No public census settles either figure, because the Open Charge Alliance’s certified product register does not distinguish residential from commercial equipment. Local control from a server inside the home is in any case neither required by the protocol nor assumed by its architecture, which presumes a charge point management system rather than a household.

OpenADR appears zero times in the normative requirements document. It appears only in the non-normative application guide, and there only as future work: “an additional standard protocol for energy management, such as OpenADR 3.X, to unlock further compliance of OEMs devices and energy management programs in different geographies” (APG p.9). The document that establishes requirements does not mention it.

So a Mercury-certified EVSE is a charger that speaks a 2017 commercial fleet-management protocol to a server its manufacturer selects. Whatever else that is, it is not a path to the residential customer.

A second certification on top of one that already exists

OCPP conformance is already certified, by the body that publishes the protocol.

The Open Charge Alliance has run an OCPP 1.6 certification program since 2019. It publishes the test procedure and test case documents, maintains a public list of certified products, and accredits nine laboratory locations to perform the testing. Its Smart Charging option is **mandatory for a charge point management system and for a full charging-station certificate**, and optional for a subset certificate. It covers support for “all profile types, including stacking,” which is the charging-schedule and limit behavior that Mercury’s requirements document specifies.

Mercury’s own normative text acknowledges this. REQ p.13 describes “Consortium members who manufacture and/or integrate **OCPP-certified products** for consumer services and/or demand response programs,” and states that the Mercury-to-OCPP mapping “is under review with the Open Charge Alliance (OCA) before its release for further dissemination.”

So for the protocol layer, a manufacturer seeking the Mercury Mark certifies twice. The same laboratory may perform both: Mercury's public account is that "DEKRA then tested the first three devices against it independently," and DEKRA Certification appears three times on the Open Charge Alliance's list of OCPP test labs.

What Mercury adds on top is performance grading: response time, telemetry interval and meter accuracy sorted into tiers. That is a reasonable thing to certify and the Open Charge Alliance does not do it.

Which produces the asymmetry the Commission should notice. **The half that duplicates an existing program is public. The half that is new is not.** The OCPP conformance that Mercury layers on top of has a published specification, a published test procedure and a public product register. The performance grading that is Mercury's actual contribution has parameters that appear in no published document, and a test procedure, EVSE-001-2025-TST, that has never been released.

The cost consequence follows. A second certification against a second specification, in a second laboratory engagement, is a real expense for a manufacturer and ultimately for a customer. That expense can be justified by what the second mark adds. It is harder to justify when the thing it adds cannot be inspected, and when the mark does not certify the one property a customer would most want, which is the ability to choose the service their charger talks to.

Four failures, one partial and one silence. The specification is competent at what it sets out to do, and does not set out to do what the six tests ask.

The Mark is advertised on parameters that appear in no published document

The specification does not carry a price. No price, tariff, cost or currency field appears in any of the requirement tables (REQ p.7, pp.11 to 12). The application guide puts price on the roadmap: "Part of this will be to define content generic to many or all devices such as price reception and response, and energy reporting" (APG p.10). It then assigns price behavior to the manufacturer by name, stating that Mercury "aims solely at standardizing the core functional behavior of EVSEs" and does not intend to limit OEM innovation including "the alignment with dynamic electricity pricing" (APG p.5). Mercury certification therefore says nothing about whether a charger can receive a price at all. That is a different matter from how a charger chooses to respond: two chargers responding differently to the same price is the system working as intended, while two chargers with no common way of being told the price is the gap.

The parameters the Mark is advertised on are not published. The released documents define Mercury Level 1, 2 and 3 with response times of under 3, 20 and 60 seconds (REQ p.6, APG p.7). Mercury's public pages advertise Class A, B and C at 1, 15 and 120 seconds. The strings "Class A," "Class B" and "Class C" appear nowhere in either released document, and "Level 1" appears on no public page. No public document maps one scheme onto the other. Meanwhile the third document in the set, EVSE-001-2025-TST, "Test procedure, detailing how compliance is verified," was "pending finalisation" at publication (REQ pp.3 to 4) and remains unpublished, while Mercury's /mercury-mark page states that "The class specification and conformance tests are published."

The released requirements document also directs the reader, normatively, to “the communication protocol options stated in Table 8” (REQ p.13). The document contains no Table 8; its last table is Table 7. A specification whose operative sentence on which protocols a certified product may speak points at a table that does not exist is not finished.

Recommendation. Where the CEC relies on third-party certification to demonstrate conformance, it should require three things of any mark, applied equally to all of them:

1. The specification and the conformance test procedure are both publicly available at no charge.
2. The parameters the mark is awarded on are the parameters the published specification defines.
3. The test procedure verifies that a customer can point the appliance at a server of their own choosing. A mark awarded to a device that connects to a server “as described by the OEM” certifies against customer choice rather than for it.

Conclusion

The Commission asked how these standards can work together, what their respective roles are, and what implementation pathways support dynamic signals and consumer-authorized control. Both signals should arrive over the Internet by an open standard protocol addressed to the customer. The price should reach the appliance, and the limit the device that enforces it, over an open, standard, local, IP-based protocol the customer can point where they choose. Everything else, including aggregation, certification and legacy bridging, is valuable work that should compete above that line rather than substitute for it.

Respectfully submitted,

Donald C. Jackson
Grid Coordination
grid-coordination.energy

Attachment A

Two Signals, Six Tests

The presentation given in the Interoperability Overview at the Commissioner Workshop on Interoperability, November 4, 2026.

Donald C. Jackson, Grid Coordination

GRID COORDINATION

Two signals.
Six tests.

The grid has two problems. Few proposals solve even one.

DC Jackson · grid-coordination.energy · November 2026

The grid has two problems, and they are not the same problem

BALANCE THE GRID

The entire grid, hour by hour

Supply must match demand, always.

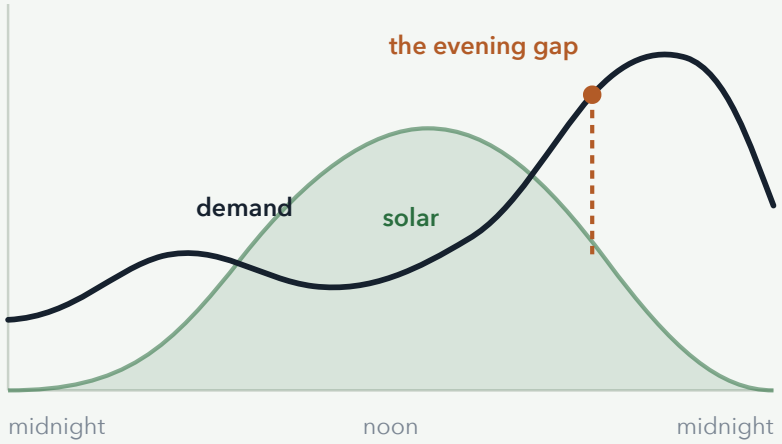

An economics problem. Move energy in time.

PROTECT THE TRANSFORMER

One block, right now

The transformers and wires have a fixed capacity.

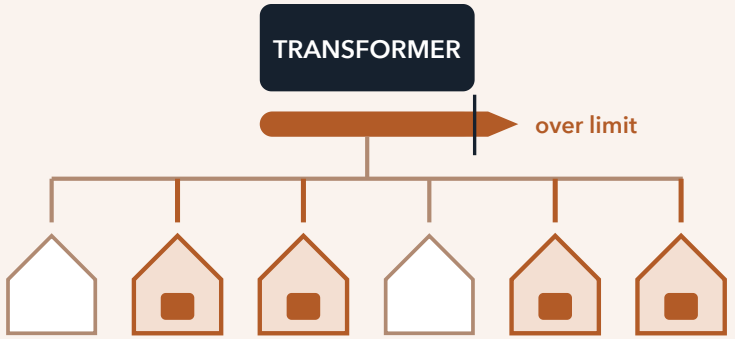

four of six charging at once

A physics problem. Cap power right now.

Balancing the grid is an **economics** problem. Protecting the transformer is a **physics** problem. **Both land on the bill.** Peak we cannot shift is generation and transmission somebody has to build; a transformer we cannot protect is a street somebody has to dig up. Ratepayers fund both.

Price balances the grid. It cannot see the transformer.

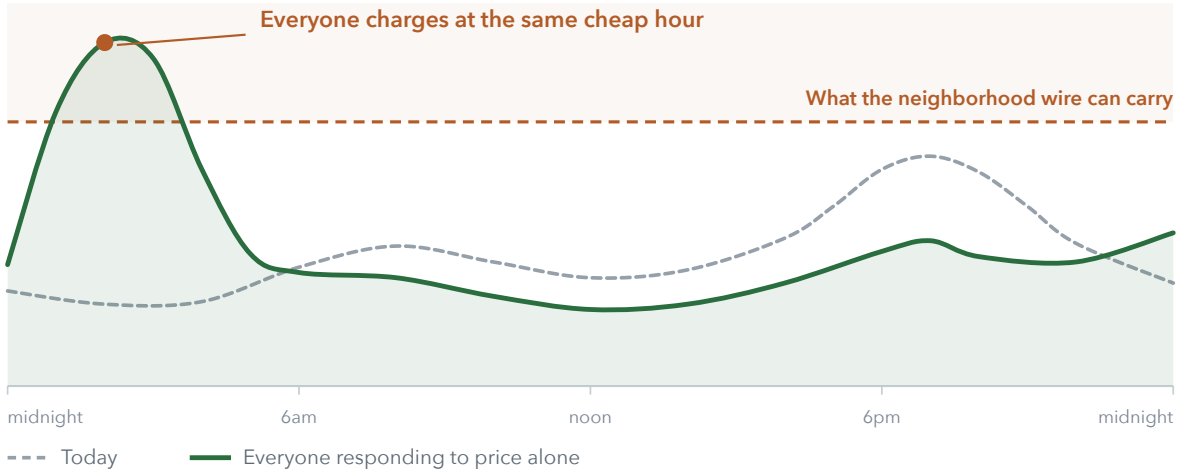

Highly dynamic pricing is the right answer for balancing the grid, and the most important thing California is doing on demand flexibility.

It should be pursued harder, not hedged. **Dynamic pricing reduces capacity problems. It cannot eliminate them**, because a price is shared by everyone on the same asset. It cannot stop five houses on one transformer responding to the same number at the same moment.

That is why we also need a capacity limit.

We have always had static capacity limits. Dynamic limits are now possible.

	STATIC	DYNAMIC
The size of the pipe	The service rating: the panel, the main breaker, the conductors back to the transformer	The power a connection may import, and may export, right now
It changes	Once, when somebody digs up the street	Every interval, with a schedule ahead and a default to fall back to
Who sets it	The electrical code, at installation	The distribution utility, from what the shared transformer and conductors can carry
Who enforces it	A breaker, and a **UL 3141** power control system (PCS)	The same **UL 3141** power control system
How it arrives	Built into the hardware	**This is the gap**

What makes them possible is already in place. A UL 3141 power control system holds a whole site under a number, and D.26-02-025 put a dynamic limit into a California tariff. **The limit needs to reach the house.**

Two signals: a price, and a capacity limit

	A DYNAMIC PRICE	A PER-CUSTOMER POWER LIMIT
What it is	What energy costs right now	How much power a connection may draw right now
Scale	The entire grid: hundreds of thousands of customers at once	The customers on one distribution transformer: five or six
What it solves	Balancing the grid	Protecting the transformer
What it does	Moves energy across time	Caps power, immediately
Without it	Nobody has a reason to shift	Nothing stops everyone shifting at once

These two can be adopted separately. Different signals, different timescales, different failure modes.

The limit is on the connection, not the device

One customer, one connection, one limit (import and export). A modern home has solar, a battery, one or two car chargers, a heat pump and a heat pump water heater, all at once. The limit belongs on the wire, not on any one appliance.

Today we enroll one device at a time. One thermostat program, one charger program, one inverter under IEEE 2030.5. A customer with all three is in three programs, none of which can see the other two.

> **Every appliance a customer adds makes device-by-device coordination worse.** Control each one separately and nobody is managing the total impact to the transformer and distribution grid.

TO CLOSE ON CAPACITY

One box enforces the limit. Everything else only needs the price.

Enforcing the limit is one box's job. The **UL 3141** power control system holds the whole site under the number, because that is the only place the whole site's draw is visible. Two enforcers would fight each other.

Every other appliance just reads a price. A water heater, a heat pump, a dryer, a dishwasher: receive a price stream and shift. No capacity logic, no new vocabulary, nothing about transformers.

An energy manager planning the house needs the limit too, and that is cheap. It schedules better for knowing it, and is not thereby responsible for holding the site under it. Every other appliance is told what it may draw, in its own vocabulary, by whichever of those holds the plan. Receiving requires no listing. Enforcing does.

Where there is an energy manager, it takes the price. Where there is not, the appliance takes it directly. Many homes will have no energy manager, especially in the near term, which is why the capability belongs in the appliance.

So the ask on an appliance vendor is two things, and they are the same for every product category. Respond to a price, and support an open local control protocol so the house can reach it.

> **The capacity signal has to exist, be published, and reach the house.**

HOW IT WORKS TODAY

The problem with aggregators and manufacturer clouds
Three companies must agree before one appliance flexes its demand

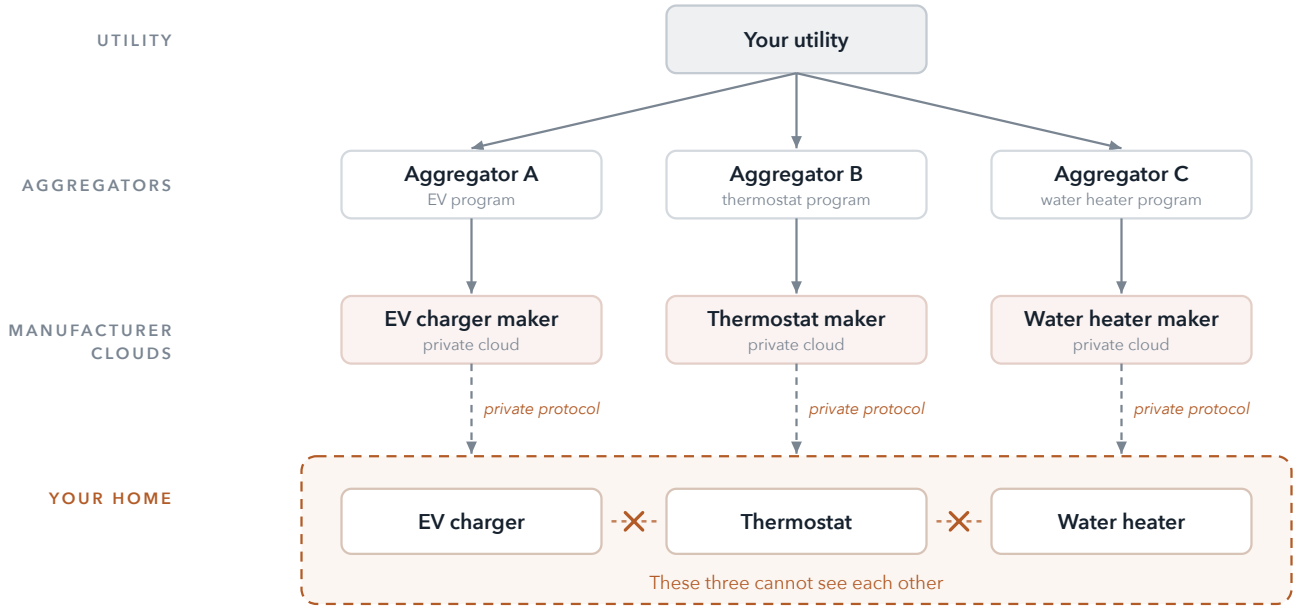

How can a customer's appliances and DERs be orchestrated when every one of them is controlled separately, and somewhere else?

And any of it can be switched off. On October 25, 2025 Google ended cloud support for first and second generation Nest thermostats: those homes lost every ability to respond to an outside signal, on a date the manufacturer chose.

HOW IT NEEDS TO WORK

Grid: Both signals, one open protocol, direct to the customer
Home: One open protocol for local appliance control and orchestration

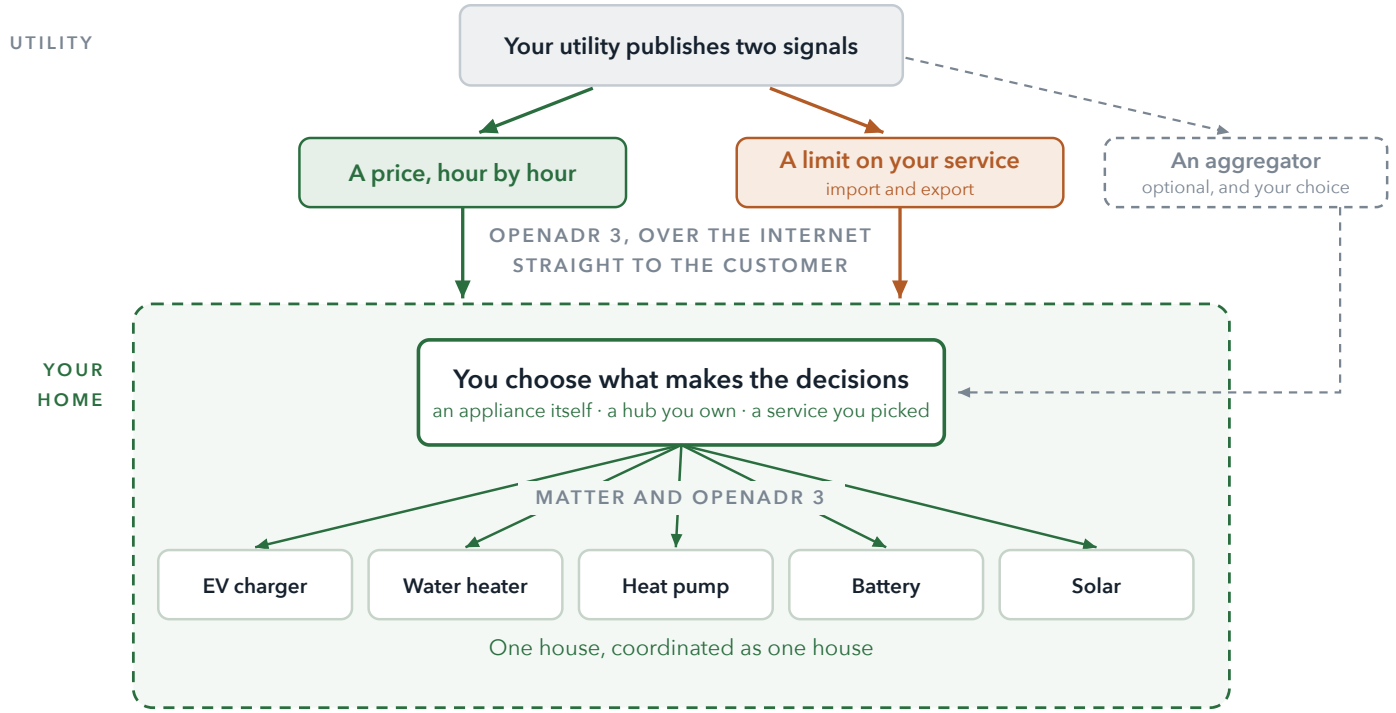

Most of this is running today, at pilot scale. A free, public server carries real hourly California prices, and free open-source software receives it and orchestrates the house: no account, no contract, no enrollment. PG&E is piloting dynamic capacity limits with AMI 2.0.

Six tests for any flexibility proposal

Rate designs, appliance standards, VPP programs, capacity and interconnection proposals. Every one of them will use the words flexible, smart, and customer choice. **Many of them also insert a new gatekeeper between the grid and the customer.**

1. Does the grid send the customer **an objective**, or does it **operate the appliance**?
2. Does the open interface reach **the appliance itself**, or stop at **the manufacturer's cloud**?
3. Can the customer choose **who optimizes their home**, including doing it themselves?
4. Can the customer take **the signals directly from their utility**, or is a **middleman required**?
5. Is **fail-safe behavior specified and mandated**: if the signal fails, is **a known default** held?
6. Does the customer **who cannot respond** end up better off, or worse?

TURNING TO THE CEC'S OWN RULEMAKINGS

FDAS should be one standard for all appliances, not one per appliance

We do not need OCPP for chargers, CTA-2045 for water heaters, AHRI 1380 for HVAC and IEEE 2030.5 for solar and storage: a protocol per appliance category, and a rulemaking behind each.

The Commission has built this shape before. The Load Management Standards are one rule covering every large IOU, not a separate proceeding per utility. FDAS should be one rule covering every flexible appliance, not a separate proceeding per appliance.

Every one of them needs the same fact: what energy costs now and next. The limit on the connection goes to the device that enforces it and to an EMS planning the house, not to every appliance. There is nothing water-heater-specific about either.

Grid to house: OpenADR 3. One open protocol over the Internet carrying both signals, with the server address a setting the customer controls.

House to appliance: a standard, open, local, IP protocol for power and energy management. Matter is the one that ships today in the ecosystems people already own.

This is less regulation, not more. Both protocols ship today. One capability written once, instead of a rulemaking per appliance forever, each going stale the day it is printed.

> Every appliance in the house is wired to one electrical code. Nobody wrote a separate code for water heaters, and the signals that tell them what to do should arrive the same way.

THE FDAS POOL CONTROLS RULING

It fails essentially all of these tests

THE SIX TESTS	VERDICT	POOL CONTROLS, TITLE 20 SECTIONS 1690 TO 1697
1. An objective, not operation	**Not met**	The rule fixes the hours instead: 9 a.m. to 3 p.m., nothing automatic 4 to 9 p.m.
2. Open to the appliance	**Barely**	Open standards required for the **consent** functions only
3. A server the customer chooses	**Not met**	No customer-configurable server address
4. Direct from the utility	**Not met**	A vendor-exclusive, cloud-only design complies
5. Specified fail-safe behavior	**Not met**	No signal to interrupt: the rule hardcodes a clock
6. The customer who cannot respond	**Met**	Answered, by brute force, for as long as those stay the right hours

Pool controls is adopted, and thermostats now have draft rule text.
HVAC, water heating, EV charging and home batteries are all flexible, and all of them land on the same transformer. Consolidate into a unified standard for all flexible demand appliances, one that passes all six tests.

What not to mandate

IF THE RULE SAYS	WHAT YOU WILL ACTUALLY GET
"Must support [named standard, 2019 edition]"	Compliance frozen at the state of the art in 2019, and a vendor incentive to do nothing further
"The manufacturer shall provide a grid interface"	A cloud interface. The appliance still speaks nothing open, and it dies when that cloud does
"Customers participate through a program administrator"	Every customer needs a third party's permission before they can help the grid
"Each inverter shall be controllable"	Five separate control paths into one house and nobody managing the total
"The appliance shall run between 9 a.m. and 3 p.m."	A load shape frozen into regulation. It cannot follow a price, cannot respect a limit, and cannot change when the grid does

Name the capability in the rule, and the standard in an appendix you can update without reopening the rule. Write it at the appliance, not at somebody's server.

Who needs to do what

CEC

One unified FDAS standard for all flexible demand appliances, instead of a rulemaking per appliance category.

Write the capability into it: read a price, and accept a draw limit from the house, over an open protocol the customer can repoint.

CPUC

Require utilities to publish the price and the capacity limit, via OpenADR 3. The limit ordered in D.26-02-025 has to reach the customer machine-readable, with no enrollment and no intermediary.

Extend the flexible connection tariff beyond PG&E and SCE.

Require every tariff filing to include a URPX rate plan, so any rate becomes a price stream.

Utilities

Publish the limit and the price, via OpenADR 3. A flexible connection the customer signs is worth nothing until the number reaches the equipment in the house.

Require both signals of your suppliers, in the next procurement you write.

Appliance OEMs

Both interfaces. Read a price, accept a draw limit from the house, and let the customer point the server somewhere else.

Ship it now. OpenADR 3 and Matter both do this today, and neither waits on a regulator.

> Every one of these fits inside a proceeding that is already open, and uses a standard that already ships.

BACKUP

Detail

Today, electrifying a home means a service upgrade

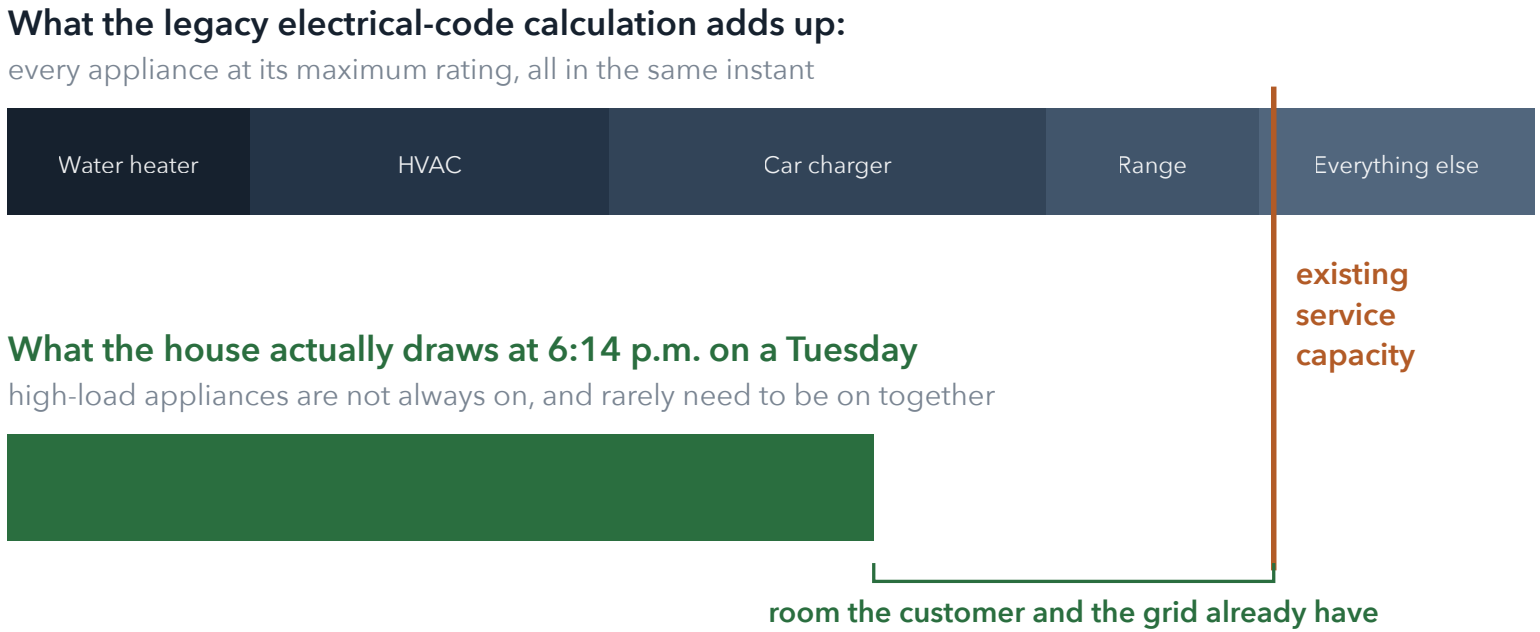

The top bar is a calculation. **The bottom bar is a measurement.**

We charge customers for a peak that does not happen, because we have no way to promise that it will not. And the upgrade answers only the customer's half of the problem: this house can electrify, and the shared transformer is handed more load to carry. Protecting that transformer means replacing it early. Both upgrades land in rates.

STATIC CAPACITY, AND THE ONE THING THAT HAS CHANGED

How a house electrifies past the rating of its own wire

The static limit is three numbers that have to agree. The service conductors from the transformer, the main breaker, and the panel busbar. The house may not draw more, and a breaker makes sure of it.

A load calculation decides whether the new load may be installed at all. NEC Article 220 adds up connected load with fixed demand factors and no knowledge of how that house actually behaves. A home adding a heat pump, a heat pump water heater, an induction range and EV charging can exceed its service on paper while never approaching it in practice.

The UL 3141 power control system changed the answer. A listed PCS actively holds the site below the existing rating, so the load behind it comes out of the load calculation. The customer electrifies on the conductors and the service feed already there.

California put that in a tariff. D.26-02-025 excludes load managed behind a certified power control system from the connected load calculation.

> **The number the power control system holds the house under is still static.** It enforces what it was configured with. The hardware that respects a limit is listed, shipping, and recognized by code. Nothing outside the house can move the limit.

Test 1: send an objective, not a command

ADDRESSED TO
the water heater at 42 Main Street

TURN OFF NOW.

A COMMAND FROM OUTSIDE

ADDRESSED TO
42 Main Street

THIS HOUR: 41¢
NEXT HOUR: 30¢

YOUR CONNECTION
6 kW IMPORT UNTIL 6 PM

AN OBJECTIVE

Same kilowatts. **Only one of them has to know what is inside your house.**

A customer operating their own equipment, through a controller they chose, is exactly right. What does not belong at the grid boundary is a grid party or a manufacturer holding that control instead.

PASSES an hourly price and a connection limit, published to everyone.
FAILS a program that needs the ability to command a named appliance in a named house.

> **A price is a one-way number. A command is a key to somebody's house.**

Tests 2, 3 and 4: open all the way to the appliance

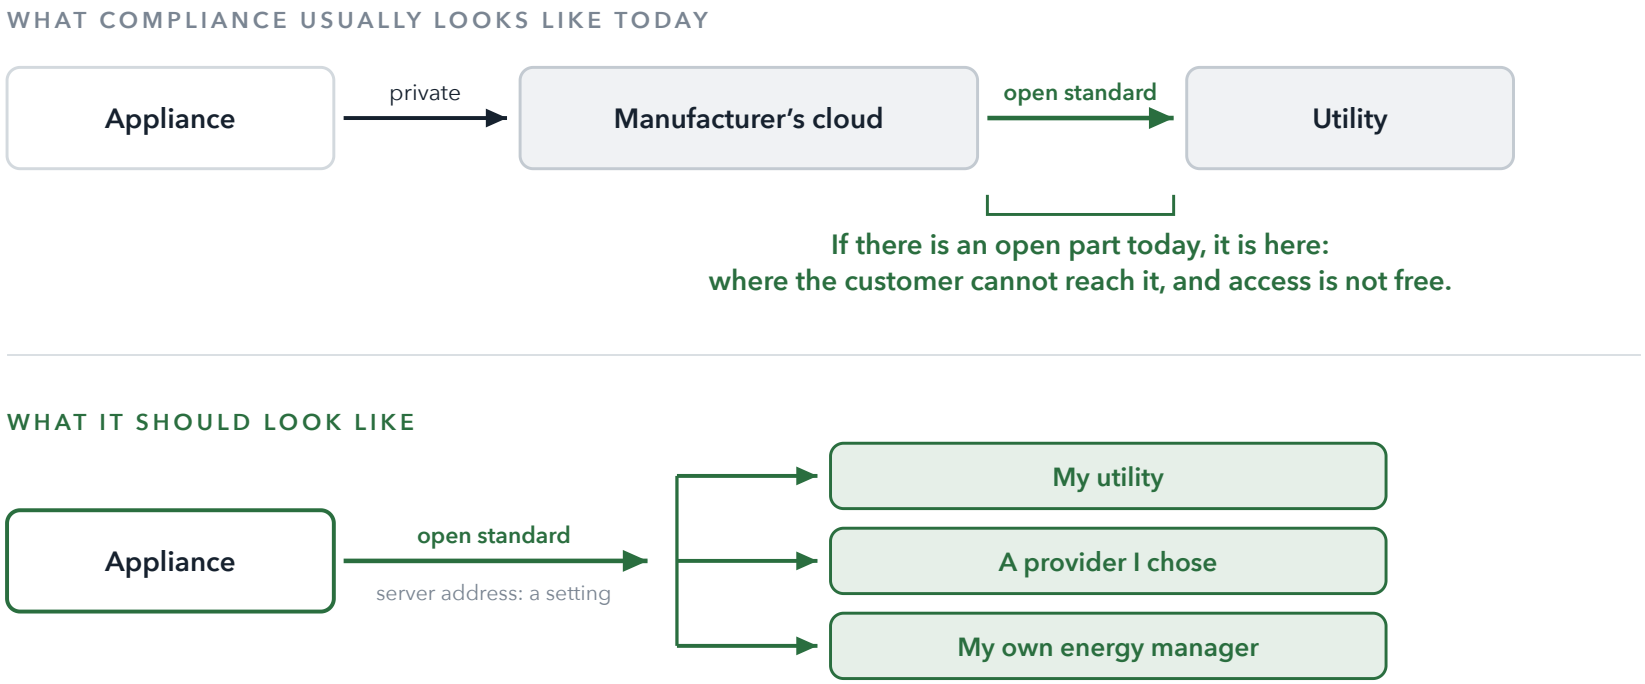

One setting. Three legitimate answers. **The customer picks.**

PASSES the server address is a setting the owner can change, and a tariff any customer can take without a third party.
FAILS an appliance that meets the standard only through a cloud the manufacturer can switch off.

Test 6: the customer who cannot respond

The first objection to hourly prices and per-customer limits is that the affluent optimize and everybody else gets rationed. It is the right question and it deserves the first answer, not the last.

	TO JOIN A PROGRAM TODAY	TO READ A PUBLIC SIGNAL
An enrollment	Required	Nothing
A credit relationship with a third party	Required	Nothing
A particular brand of device	Usually	Any device that can reach the network

Nobody should have to watch a price. The appliance watches the price. The customer sets a preference once.

Many customers who cannot shift are paying today for a grid sized to everyone else's worst hour. Capital spending that never enters rates is their benefit, and it is the largest one.

Dynamic pricing asks less of people than a virtual power plant does. Everyone already pays for VPP incentives through rates. PNNL's DSO+T study found participants **and non-participants** both saw bills fall 10 to 16 percent.

THE COMMISSION HAS ALREADY ASKED THE QUESTION

> "What hardware, software, or data sharing safeguards may be needed to ensure that customers can obtain a flexible connection without being restricted by proprietary arrangements?"

Assigned Commissioner and ALJ ruling, July 7, 2026, seeking comment on the Assigned Commissioner's Proposal on Flexible Service Connections. Grid modernization rulemaking R.21-06-017, Question 8 to parties.

A time-of-use rate is a price stream with fewer changes

The signal carries what energy costs now and over the coming hours. A schedule, not a single number. That forward view is what lets a house plan ahead.

Highly dynamic pricing is the destination, and nothing has to wait for it. A flat rate is a price stream with no changes at all. Same mechanism, same appliance, same software, whatever tariff the customer is on. The appliance never needs a tariff engine, only an interval consumer.

What that asks of the Commission: require every tariff filing to include a machine-readable rate plan. The open **Utility Rate Plan Exchange (URPX)** standard, hosted by LF Energy, is built for exactly this, and ships worked test cases built from published PG&E and Con Edison filings. I am modeling California residential rate plans in it, in the open, at **grid-coordination/urpx-rate-plans**.

And what it asks the Commission to retire: tariff attributes a price stream cannot carry. A **demand charge** is the clearest example: it prices a peak after the fact, so no device can act on it from the signal alone.

> Every customer can be served by the same mechanism today, whatever tariff they are on. The rate can get more dynamic later without touching a single appliance.

Where California actually stands

DRIVER	WHERE IT STANDS
CalFUSE	CPUC staff framework, June 2022, for a single universally available dynamic price signal. A white paper name, not a decision or a program
Load Management Standards	Large IOUs: an hourly, marginal cost-based rate for every customer class by January 1, 2027. Large CCAs: that rate or a qualifying program by July 1, 2027
Rate design	D.25-08-049, August 2025. Guidelines for designing demand flexibility rates; the rates themselves are still in testimony. Remaining issues moved to R.26-04-009, opened April 2026
Flexible service connections	D.26-02-025, February 10, 2026. PG&E and SCE to establish a standardized tariffed flexible connection. Load behind a certified power control system is excluded from the connected load calculation
Appliance standards	One adopted: pool controls, September 29, 2025. It mandates a preconfigured schedule (1693(b)(2)(C)(1)) and open standards for the **consent** functions (1693(b)(6)(B)), but no ability to receive a price, a signal or a limit, and no customer-configurable server. Thermostats (24-FDAS-03) have draft Express Terms as of November 10, 2026, requiring an OpenADR 3.1 Virtual End Node and forbidding sole reliance on a manufacturer's cloud, with no customer-configurable server. EV charging (24-FDAS-04) and home batteries (25-FDAS-01) are open dockets with no rule text
Legislation	AB 1787 held under submission in Senate Appropriations, August 13, 2026. The regulators are moving; the statute is not

Why some standards are hurting, not helping

CTA-2045, marketed as EcoPort. Washington, Oregon, Colorado and New York require the socket on new electric water heaters. Nothing above it was ever standardized, so every vendor's module speaks its own language upstream, and the module a customer has to buy and install runs well over $150. **The open interface stops at the socket. The standard is obsolete and should be retired, not written into new rules.**

AHRI 1380. The standard permits a local interface. **Virtually every AHRI 1380 product is cloud-only**, and AHRI certifies nothing against the standard.

IEEE 2030.5. CSIP certifies aggregators and utility servers as well as inverters, but puts everything below the aggregator or site management system **out of scope**, so it stops at the edge of the house. One inverter is still the wrong unit for a house with solar, a battery and two bidirectional chargers: control each separately and nobody manages the total. **What belongs there is one import and export limit for the whole connection.**

OCPP. Built for commercial and public EV charging. **Almost no** US residential chargers implement it, and controlling one from a local server is not required and almost never supported in shipping products. **There is no need for a protocol unique to EV chargers.**

Two open standards already support every one of these principles, and both ship today.

OpenADR 3 carries prices, events and limits from the utility to the customer's own equipment, **over the Internet**. The server address is the customer's setting.

Matter carries the same information to every device inside the house, **over the home network**.

WHERE EACH SIGNAL IS ACTUALLY DEFINED

Both signals already have a named home on both hops

	A PRICE	A CAPACITY LIMIT
Grid to customer	**OpenADR 3**	**OpenADR 3**
Inside the building	**Matter** and **OpenADR 3**. Matter carries it in the Commodity Price cluster	**Matter** and **OpenADR 3**. Matter carries it in the PowerThreshold feature of the Meter Identification cluster

Everything above is already published. **The gap is not in the standards. It is that nobody is sending the limit.**

The manufacturer's real objection, and the answer

What they say: the cost of the parts, and a ten-year support obligation on a device they no longer control.

The customer already paid for the parts. Every one of these appliances ships with a microcontroller and a Wi-Fi radio, on the customer's receipt. The customer still cannot talk to the thing they bought without going through the manufacturer.

What the closed path actually protects: the customer relationship, the customer's usage data, and a share of whatever that customer's load shifting earns.

None of those three are theirs. The usage data is private, and it belongs to the customer. The incentive is paid for the customer's forbearance, not for the manufacturer's software. And a customer-configurable server **reduces** the obligation they say worries them: they stop being the mandatory operator of a cloud service for the life of the appliance.

> The hard part was never the technology. It is that nobody is required to open the door.

EVERYTHING HERE IS PUBLIC AND FREE

Where to find all of this

This deck, and the one-page handout	grid-coordination.energy/presentations
The price server, with client tutorials	github.com/grid-coordination/price-server-user-guide
Home Assistant integration, in the HACS default repository	github.com/grid-coordination/openadr3-ven-hass
URPX rate plans for California, a public corpus of modeled tariffs	github.com/grid-coordination/urpx-rate-plans
Open-source libraries (Clojure and Python)	github.com/grid-coordination
My CEC docket comments and the California policy record	grid-coordination.energy/policy
OpenADR 3, the grid-to-customer standard	openadr.org
Matter, the in-home standard	csa-iot.org
URPX, machine-readable tariffs, hosted by LF Energy	urpx.org

> No account, no contract, no enrollment on any of it. Every link in this deck is live and clickable in the PDF.