Trusted Energy Interoperability Alliance

EXECUTIVE BRIEF

Scale VPPs by design

What VPP architecture actually costs at 10k, 50k, and 100k assets.

Contents

Founders

Executive summary

VPP operators and aggregators need to know whether their platform can reach 100k assets without a redesign.

The U.S. Department of Energy projects that domestic VPP capacity must roughly triple, from current scale toward 80 to 160 gigawatts by 2030, to meet 10 to 20 percent of peak demand.1 Meeting that target means coordinating millions of distributed energy resources at sub-second latency. Most platforms were not built for it.

This brief quantifies architectural costs at different commercial thresholds of enrolled assets, and frames the choice facing operators and aggregators today.

The figures are directional estimates grounded in DOE and Brattle Group analysis 2 and TEIA architecture research, not benchmarked client results. They are meant to make the decision clear, not to be audited line by line.

We explore architecture costs as fleets grow, comparing proprietary to open systems.

AI-driven virtual power plants now execute market transactions at machine speed. The architecture that carries them has to scale at the same pace.

The problem: Most platforms were not created to scale

Every VPP platform sits somewhere between two architectural poles, proprietary systems and ones that interoperate.

The blocker is rarely the device or the protocol. It is whether the architecture underneath can carry the fleet you are about to enroll.

An architectural divide between open and proprietary platforms is happening with VPPs. At low device counts the divide is invisible, because almost any architecture coordinates a few thousand assets. The divide only shows up at scale, which is exactly when it becomes expensive to fix.

Integration debt accrues on every release
Each OEM implements proprietary protocols, including OCPP variants, non-standard IEEE 2030.5 profiles, and bespoke HVAC interfaces. Without a standardized abstraction layer, every new manufacturer relationship requires custom integration code, and integration debt accrues on every firmware release.

Centralized control becomes a bottleneck
Monolithic control architectures centralize dispatch logic. That simplifies early design but creates a single coordination bottleneck. As concurrent connections climb, a centralized design can degrade from a couple hundred milliseconds toward a second or more at the P99, which is the point where ancillary-service revenue starts to slip away.

Security surface widens vendor by vendor
Proprietary integrations require per-integration security review. Across twenty or more manufacturers, maintaining a current security posture across inconsistent authentication models creates both escalating audit overhead and real incident exposure.

The cost of scaling

Open-architecture figures model platform cost as a per-asset operating rate that declines as automation absorbs growth.

Proprietary multipliers add the compounding cost of custom integration, operational overhead, and lost market access on top of that baseline.

At small volumes, the two models look deceptively similar. As asset counts climb, the proprietary system's costs compound until they become unsustainable.

The table below frames annual platform cost across three commercial thresholds.

Scale	Open (TEIA)	Proprietary	Cost and risk
10k assets	$0.5 to 1M/yr Integration overhead manageable	1 to 2x higher Custom API bloat; firmware drift	Begins to drag engineering velocity
50k assets	$2 to 3M/yrAutomation offsets growth	3 to 4x higher Audit backlog; latency degrading	Revenue risk: wholesale market exclusion
100k assets	$3 to 5M/yr Marginal cost declining	8 to 15x higher Failure cascades; resolution in hours	Platform operationally unsustainable

At 10k assets a proprietary approach is expensive but survivable. By 100k, incident cascades and re-integration make it hard to sustain under SLA-backed grid service agreements.

Proprietary architecture compounding cost

Three cost categories compound as device portfolios expand. Each is manageable on its own early on. Together they decide whether the economics of fleet growth ever work.

Each category is individually manageable. The risk is that they compound at the same time, exactly as the fleet you depend on for revenue reaches its largest scale.

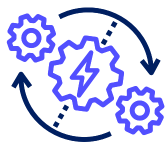

1. Integration engineering

At 10k assets across four to six manufacturers, integration is a known cost center. At 50k across fifteen to twenty manufacturers, it can consume a substantial share of engineering capacity. The trap is not the integration itself, it is the ongoing maintenance obligation that crowds out work on differentiated grid services. TEIA specifications define a common device abstraction layer, and organizations implementing TEIA-compliant adapters have reported onboarding cycles falling from four to twelve weeks per OEM to one to three days, though results vary by complexity.

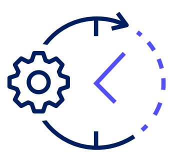

2. Operational overhead and latency

Frequency regulation and spinning reserve markets require sub-second response. A platform that cannot clear roughly 250 milliseconds at the P99 at 50k devices may be structurally excluded from ancillary-service revenue, often the highest-margin stream available to aggregators. Microservices architectures fan dispatch across independent domains and can hold sub-250-millisecond response at 100k devices and beyond under well-implemented designs.

3. Security surface and audit cost

Inconsistent device authentication is a documented source of vulnerability in VPP architectures, expanding the grid's attack surface through unverified endpoint commands.[3] Zero-trust architectures enforcing mutual TLS and certificate-based authentication at every service boundary eliminate the shared-secret and API-key exposure patterns common to legacy platforms.

Centralized policy enforcement means one audit scope, not twenty. An adapter layer also decouples firmware changes from core logic, and org-scoped authentication lets one platform serve multiple utilities natively rather than running separate instances per relationship.

What this means for executives

The structural question is whether your current architecture can reach 100k enrolled assets without a redesign. Three signals suggest it is worth a closer look. They are directional heuristics, not validated thresholds.

First, integration engineering consuming more than roughly 30 percent of engineering capacity on a quarterly basis. Second, new OEM onboarding routinely requiring more than four weeks of engineering time. Third, dispatch latency during peak events exceeding several hundred milliseconds at the P99 at current fleet size. Any one of these at present scale tends to get worse, not better, as the fleet grows.

A sequenced path, not a big-bang rewrite

Start with an architecture assessment that evaluates the current platform against TEIA scalability benchmarks, with a phased modernization roadmap and quantified cost and risk.

Adopt TEIA specifications for device abstraction next; existing protocol investments in IEEE 2030.5, OCPP, and OpenADR remain viable, because TEIA provides a security and interoperability layer across them rather than replacing them.

Then decouple device management, data processing, and grid service delivery incrementally, following migration patterns that avoid a single high-risk cutover.

Finally, set a zero-trust security baseline with mutual TLS and certificate-based authentication at service boundaries, which reduces ongoing audit cost while improving threat-detection coverage.

The economics of an open standard improve as you scale. The economics of a proprietary one degrade. The decision is easier to make before a redesign is the only option left.

Building for scale

The industry faces a fork. One path keeps scaling cost fragmented, rebuilt for every vendor and every audit. The other builds the platform once on an open standard designed to carry the fleet as it grows.

This pattern is familiar. The standards that made payment cards interoperable, and the certificate systems behind secure communication, all succeeded by separating a common trust and interoperability layer from the systems carrying the data. Distributed energy is reaching the same threshold, where the platforms that win are the ones that stopped renegotiating their architecture at every increment of scale.

The VPP market is growing into exactly the multi-vendor environments that reward this approach, with North American capacity already in the tens of gigawatts and DOE targeting 80 to 160 gigawatts by 2030.1 As the addressable market grows, so does the share of it that screens for architecture able to scale.

Operators that resolve cost, latency, and security now scale into that role. Those who defer keep paying the compounding cost of bespoke architecture, one integration and one audit at a time.

Open specifications, reference implementations, and implementation partners are available today at trusted-energy.org/specifications.

Sources

[1] U.S. Department of Energy, *Pathways to Commercial Liftoff: Virtual Power Plants* (2023, and 2025 Update).

[2] The Brattle Group, *Real Reliability: The Value of Virtual Power* (May 2023).

[3] *Cybersecurity challenges in renewable-dominated power grids: addressing vulnerabilities and resilience strategies*, Chalmers University Research Repository.

How we estimated cost of scaling chart

Open-architecture figures model annual platform cost as a per-asset operating rate that declines with scale as automation absorbs growth. Proprietary multipliers add three compounding components on top of that baseline: integration engineering, meaning per-OEM custom code and per-firmware maintenance; operational overhead, meaning incident handling and manual re-integration; and market-exclusion risk, meaning lost ancillary-service revenue when dispatch latency degrades. Inputs draw on DOE Liftoff[1] and Brattle Group VPP analysis[2] and TEIA architecture-pattern research. These are directional ranges, not benchmarked results; actual costs vary by organization, vendor mix, and deployment context.

Those who build a scalable architecture now define how the grid's next phase operates, instead of inheriting a platform they have to rebuild.

TEIA™ Trusted Energy Interoperability Alliance

Learn more at: trusted-energy.org
Contact us at: contact@trusted-energy.org
+1 408 616 1600

TEIA membership is available to all organizations active in the digital energy ecosystem. OEMs can join the partner program at no cost.

To learn more about TEIA standards, membership, or to request a technical briefing, visit **www.trusted-energy.org**