stem

STEM.COM

WHITE PAPER · SOLAR-PLUS-STORAGE CONTROLS

How to Define Your EMS Requirements: From Cloud Optimization to OEM Integration

A requirements framework for developers, EPCs, and technical buyers scoping solar-plus-storage projects



Table of Contents

— Executive Summary 3

1 EMS definition & challenges 4

1.1 The term fragments across stakeholders, and hybrids widen the gap 4

1.2 The cost of an undefined boundary 4

1.3 A market scaling faster than buyer familiarity 5

2 The Energy Management Control Stack 6

2.1 Battery Management System 7

2.2 Unit Controller 7

2.3 Power Plant Controller 7

2.4 Site and Grid Interface 8

2.5 Cloud Optimization 8

3 A Central Decision: One Controller or Two 11

3.1 Whether PV and storage require separate controllers 12

3.2 The related case against reprogrammable plant control 14

4 The Optimization Decision 15

4.1 Where optimization actually resides 15

5 Integrated Procurement Versus Assembling the Layers 16

5.1 Purchasing an integrated package 17

5.2 Assembling the layers separately 17

6 A Scoping Framework 18

7 Conclusion 19

— References 19

— About Stem 20

Executive Summary

The term "energy management system" (EMS) has no shared definition across a renewable energy project. When asked what an EMS is, a controls engineer, an EPC project manager, and a project finance lead will each describe a different part of the control stack, and each will be partly correct. When the buyer has not drawn the definitional boundary before entering vendor conversations, then the vendor draws it; the vendor's boundary tends to favor its own product.

Battery storage deployment is accelerating across major markets globally, and solar-plus-storage is rapidly becoming the default project structure as storage costs fall and grid operators develop frameworks to accommodate co-located resources.[1, 2] Many of the organizations now scoping these hybrid projects are doing so for the first or second time, in markets where regulatory frameworks are evolving and vendor-neutral guidance is scarce.

This paper provides a practical framework for defining EMS requirements before vendor conversations begin. The central recommendation is direct: define which layers of the control stack a buyer is responsible for, and which you are buying, before you evaluate a single vendor.

© 2026 Stem, Inc.

SECTION 1

The Problem: "EMS" Is Not a Single Specification

1.1 The term fragments across stakeholders, and hybrids widen the gap

Different stakeholders approach a hybrid solar-plus-storage project with different priorities, and each tends to define the EMS in terms of the layer that matters most to their role.

- SCADA and controls engineers focus on communication protocols, control modes, and configuration.
- EPCs focus on delivery, integration, and commissioning timelines.
- Finance and asset management teams focus on availability and compatibility with the project's revenue model.

These perspectives are all valid; they accurately describe the functions that are most relevant to each role. The difficulty presented to buyers is that they describe different functions while using the same term.

A hybrid project compounds this definitional challenge. The PV side and the storage side of a project often involve different equipment, different OEMs, and in many organizations, different internal teams. Each side's stakeholders hold their own definition of an EMS, and a requirements document written from any single perspective often captures only one part of the control stack (the one most relevant to the most involved stakeholder) while leaving the rest as unstated assumptions. Those assumptions surface late, during integration or commissioning, when they are most expensive to resolve.[3]

1.2 The cost of an undefined boundary

A buyer who enters vendor evaluation without a defined requirements boundary is exposed to vendor-led scoping. Importantly, while the consequences of vendor-led scoping can be extremely detrimental to project timelines, the practice is rarely conducted in bad faith. A vendor naturally frames the EMS requirements in terms of what its platform does well, and absent a buyer-defined boundary, the vendor-led framing becomes the default specification. Two vendors framing the same hybrid project around their respective strengths will produce proposals that are difficult to compare, because they answer different questions. The remedy is for the buyer to define the question,[3] and doing so requires a clear view of the control stack rather than deep controls expertise.

© 2026 Stem, Inc.

1.3 A market scaling faster than buyer familiarity

The pace of hybrid deployment has outrun the accumulation of buyer expertise. Across the United States, Europe, and the Asia-Pacific region, solar-plus-storage is rapidly becoming the default project structure, driven by falling storage costs, grid operator incentives, and the practical advantages of co-location.[1, 2] Many of the organizations now scoping hybrid projects are doing so for the first or second time, in markets where regulatory frameworks are still evolving and vendor-neutral guidance is scarce.



© 2026 Stem, Inc.

SECTION 2

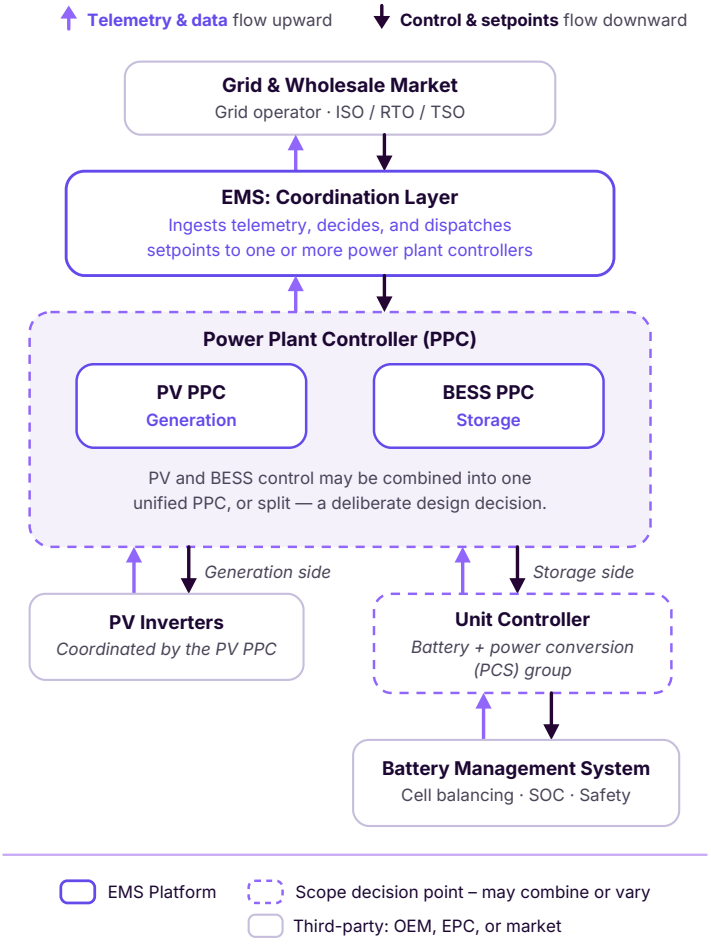
The Energy Management Control Stack

The most reliable way to define EMS requirements for a hybrid project is to stop treating “the EMS” as a single control product and instead recognize it as a standalone coordination layer sitting above a stack of distinct control components spanning both PV generation and storage. The EMS itself performs no control: it is the conductor coordinating the players below it, dispatching setpoints to one or more power plant controllers based on data it has ingested. Each of those layers performs a specific function, each can be owned by a different party, and the buyer's task is to decide which underlying layers fall within the scope of the system being purchased, and where the EMS’s coordinating role begins and ends.

FIGURE 1

The Energy Management Control Stack

How control flows through a solar-plus-storage plant. EMS coordinates the control stack below it; it does not perform control itself



© 2026 Stem, Inc.

2.1 Battery Management System

The Battery Management System (BMS) resides inside the battery and governs cell balancing, temperature, state of charge, and safety limits. The BMS is almost always provided by the battery OEM and is rarely included in a buyer's EMS scope. For requirements purposes, it is sufficient to recognize that this layer exists and that responsibility for it sits with the OEM. Because the BMS is tightly coupled to the specific cell chemistry and is integral to the battery's safety certification and warranty, the OEM must own it; a buyer generally cannot specify or alter this layer without voiding those protections, so recognizing its existence and ownership is all that requirements need to do.

The EMS conducts; it does not play

The EMS itself performs no control. It ingests data, decides, and dispatches setpoints to one or more power plant controllers.

2.2 Unit Controller

The Unit Controller coordinates a group of batteries and their associated power conversion equipment, translating higher-level commands — such as a plant-level real- or reactive-power setpoint, a charge/discharge instruction, or a curtailment order — into the device-level actions the hardware can execute, such as individual PCS or inverter setpoints. Whether buyers include this layer in their EMS definition varies considerably. Some treat it as part of the OEM's delivered package; others expect it within the EMS. Because the division is genuine and commonly encountered, this is the first layer where an explicit decision is valuable. An unstated assumption here commonly resurfaces as a costly coordination dispute during commissioning.

2.3 Power Plant Controller

The Power Plant Controller (PPC) is the layer that causes a site of individual inverters and batteries to behave as a single coordinated plant at the point of interconnection. On the generation side, a PPC coordinates the PV inverters so that an array behaves as one large generator, regulating real and reactive power and curtailing output when instructed.[4] On the storage side, an equivalent function governs the battery's response. In a hybrid plant, these functions must work in concert because the grid responds to one plant at one point of interconnection regardless of how many resources sit behind it. This layer is where the central architectural decision of a hybrid project is made, and Section 3 examines it directly.

In our team’s field experience, fewer than a quarter of buyers include the PPC when they describe their EMS requirements despite its importance. Because the PPC is the layer that determines whether the plant meets its interconnection obligations as a single coordinated resource, leaving it unspecified is among the most consequential omissions a buyer can make; doing so cedes the defining architectural choice of the hybrid project to the vendor.

© 2026 Stem, Inc.

2.4 Site and Grid Interface

Interconnection and reliability standards govern how a resource must behave at the point of connection, and they vary significantly by region. In the United States, these include IEEE 1547, the bulk power system rules administered under NERC, and the developing IEEE 2800 standard for large inverter-based resources.[5] In Europe, equivalent requirements are set by national grid codes published by each transmission system operator, informed by ENTSO-E network codes. In Australia, the Australian Energy Market Operator (AEMO) publishes its own generator performance standards. Buyers usually include this layer in their EMS definition unless a specific requirement calls for a separate, dedicated device. In some regions, grid operators require a dedicated interface device — such as a Remote Intelligent Gateway (RIG) in U.S. CAISO-interconnected projects — that must be explicitly scoped and contracted rather than assumed to be part of the EMS.

A dedicated interface

Some regions require a separate device — like a Remote Intelligent Gateway (RIG) in CAISO-interconnected projects — that must be explicitly scoped, not assumed.

SITE & GRID INTERFACE TERMS

NERC / CIP

NERC (North American Electric Reliability Corporation) sets and enforces bulk power system reliability standards. CIP is NERC's cybersecurity framework governing how control systems are configured, accessed, and protected.

IEEE 1547 / IEEE 2800

IEEE 1547 governs how distributed energy resources connect to the grid; IEEE 2800 extends similar requirements to large inverter-based resources at the transmission level.

Remote Intelligent Gateway (RIG)

A device required by some grid operators that gives the utility direct visibility into and limited control over a plant's output. Requirements vary by region.

When these definitions matter:

The site and grid interface layer carries enforceable obligations set before you reach commissioning, and these apply regardless of what a vendor's platform includes by default.

2.5 Cloud Optimization

At the top of the control stack is software that determines when to charge, discharge, and export to capture the greatest value. In a hybrid plant, decisions regarding cloud optimization are more involved than in a storage-only project because the optimizer must coordinate two resources and address:

1. When to send PV directly to the grid
2. When to use PV to charge the battery
3. How to recapture energy that would otherwise be clipped
4. How to firm variable generation against a schedule or market signal.

© 2026 Stem, Inc.

Distributed generation projects frequently include optimization within their EMS definition. Utility-scale projects rarely do, because at that scale the function is commonly handled by a separate trading or scheduling operation. At utility scale, whether an optimizer is needed depends on the offtake structure: a merchant asset still needs an optimizer to determine dispatch, so the function is commonly handled by a separate trading or scheduling operation. A pure PPA asset typically does not need one because it does not address cloud optimization; the plant's setpoint is delivered directly via an on-site RTU, with no optimization layer required. This difference in where optimization sits, and whether it's needed at all is the most frequent source of misalignment between buyer expectations and vendor scope, and Section 4 treats it in detail.

The value of this layered view is that, when understood, it allows a buyer to convert a vague request into a precise one. A buyer may begin with the statement “We need an EMS,” but with an understanding of the layered view, their more effective statement becomes: “We require the unit controller through the site interface, coordinating both PV and storage; the OEM is delivering the battery management layer, and optimization will be handled by our trading operation.” Such a statement provides a strong foundation for a productive EMS vendor conversation.

© 2026 Stem, Inc.

FIGURE 2

Control Layer Ownership Matrix

Defining Responsibilities Across Project Configurations

Each row in the matrix below is a layer of the control stack; each column is a common project configuration. Ownership of each layer shifts meaningfully depending on project structure, which is why "we need an EMS" describes a different scope in each column.

Before entering vendor conversations, a buyer's first task is to identify which configuration most closely reflects their project and to examine every ownership boundary it implies — particularly any that have been assumed rather than deliberately assigned.

Control Stack Layer	Standalone Storage *Battery-only; no PV controller*	Hybrid, Integrated Control *PV + BESS under one EMS*	Hybrid, Split Control *PV + BESS owned separately*
Cloud Optimization *Dispatch & markets*	Buyer / Owner	**EMS Vendor**	Buyer / Owner
Site & Grid Interface *Interconnection & RIG*	**EMS Vendor**	**EMS Vendor**	EPC
PV Plant Controller *Solar (part of the PPC)*	*Not present*	**EMS Vendor**	Inverter OEM
BESS Plant Controller *Storage (part of the PPC)*	**EMS Vendor**	**EMS Vendor**	Battery OEM
Unit Controller *Battery + PCS group*	Battery OEM	**EMS Vendor**	Battery OEM
Battery Management *Inside the battery (BMS)*	Battery OEM	Battery OEM	Battery OEM

Within the EMS purchase Owned by another party

© 2026 Stem, Inc.

SECTION 3

A Central Decision: One Controller or Two

Critically, most disagreements about EMS scope are not errors. There are many sound ways to build a hybrid project, and an approach that appears wrong from one vantage point often underlies a different but equally valid design intention. For example, a merchant project carrying low leverage is reasonably built for the lowest capital cost. Meanwhile, a tolled project carrying higher leverage is reasonably designed to minimize operating cost over its life. Both reflect deliberate optimization for different financial structures.

Key term

Power Plant Controller (PPC): the layer that makes a site of inverters and batteries behave as one coordinated plant at the interconnection point.

At the control layer of a hybrid plant, however, one architectural decision carries real cost and risk, and a sound requirements framework should resolve it early. Early resolution is what lets a buyer hold every vendor to one comparable specification instead of accepting whichever architecture each vendor happens to prefer. The next two sections address the decisions where this leverage matters most.

PROJECT FINANCE TERMS

Merchant project:
A project that sells its output into the wholesale market at prevailing prices. Revenue is variable and depends on market conditions.

Tolled project:
A project under a long-term contract where a counterparty pays a fixed fee for the right to dispatch the asset. Revenue is predictable and contracted for years.

Leverage:
The share of a project funded by debt rather than equity. Lenders lend against predictable cash flow, so contracted (tolled) projects can carry more debt than merchant projects.

Capex (capital expenditure):
The upfront cost to build the project.

Opex (operating expenditure):
The ongoing cost to run and maintain the project over its life.

Where the definitions matter: *a merchant project funded largely by equity against uncertain revenue has reason to minimize upfront cost. A tolled project with secure, long-term revenue has reason to spend more upfront to lower operating costs across the asset's life. Different financial structures point to different engineering choices, and both can be correct.*

© 2026 Stem, Inc.

3.1 Whether PV and storage require separate controllers

A common assumption holds that a hybrid site needs one controller for the solar plant and a second for the battery, on the premise that the two are fundamentally different control problems. However, in most cases, they are not; a single PPC is capable of coordinating both PV and BESS at the point of interconnection. Buyers should note that splitting PV and BESS control into two independent controllers adds integration work, cost, and an additional interface at which faults can occur. In the typical case, this differentiation provides no corresponding benefit. It can make grid compliance harder because the plant must still present a single coordinated response at the point of interconnection regardless of how its controls are divided internally.

This decision deserves buyer attention precisely because the hybrid project has become the default structure. When a design proposes separate PV and storage controllers, the appropriate question should ask which specific requirement drives the separation. Some designs do have a legitimate reason, which is why the architecture allows for it. Frequently, though, no such requirement exists, and the separation is an inherited habit rather than a considered choice. Ultimately, buyer scoping requirements should treat unified coordination as the starting point and require any proposed split to be justified on its merits.



© 2026 Stem, Inc.

FIGURE 3

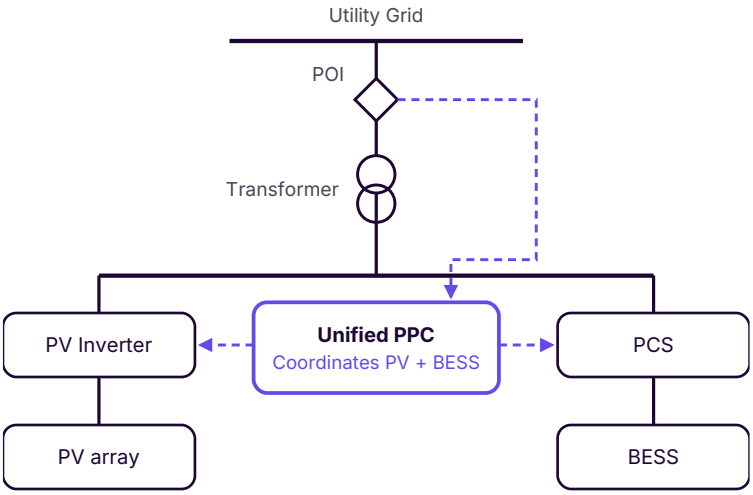
Unified vs. Split Plant Control

Two ways to coordinate PV and storage at a hybrid plant's point of interconnection — one controller for both assets, or a dedicated controller for each.

Unified Plant Control

A single controller meters the point of interconnection and commands both PV and storage.

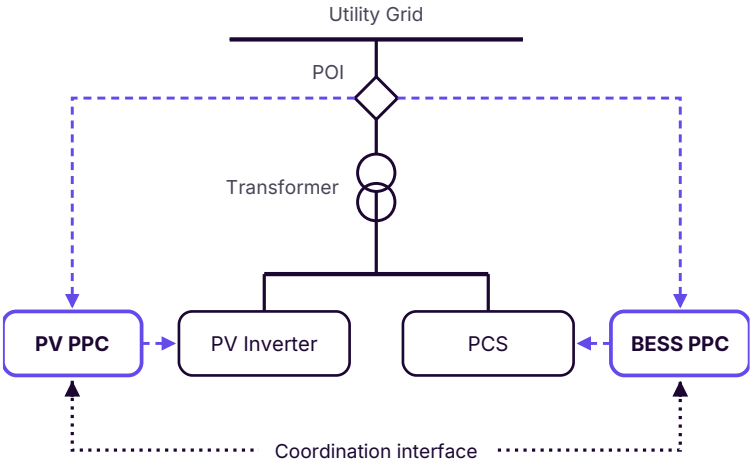
One vendor interface, one source of truth, and no negotiation between controllers — the simplest path to a compliant set point at the POI.

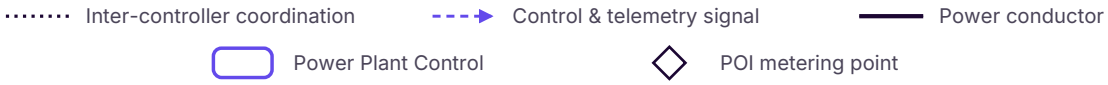


Split PV + BESS PPC

Dedicated controllers let each asset run its best-fit vendor.

But both share one point of interconnection and must coordinate with each other — an added interface to specify, test, and tune, and a path for conflicting set points if it is not.





Note: the split configuration shown in the diagram only functions when a top-level PPC and EMS sit above the two controllers to coordinate them, unless the asset is co-located and does not require POI control. Without that coordinating layer, dispatching PV and storage as separate controllers does not produce a single coherent plant response.

© 2026 Stem, Inc.

3.2 The related case against reprogrammable plant control

A second misconception common among experienced SCADA engineers is that the buyer must retain the ability to customize and reprogram the PPC. The expectation has an understandable origin: on earlier projects, particularly at utility scale, writing custom control code was standard practice, so the requirement was written as a need to reprogram the controller.

In nearly all cases now, this capability is unnecessary, and requesting it introduces risk. The control algorithms in a mature PPC are validated against applicable grid codes and, in many cases, tested against hardware-in-the-loop simulators before deployment. Rewriting such logic in the field reintroduces the faults and edge cases that validation is designed to eliminate. Field-rewritten logic has not passed the grid-code validation or hardware-in-the-loop testing the original code cleared, so faults that would have been caught in a controlled environment can instead first appear on a live, grid-connected plant — tripping it offline, breaching compliance at the point of interconnection, or damaging equipment.

Here, there is an important distinction to be made between configuration and reprogramming. Engineers can adjust setpoints and tune parameters for site-specific conditions without touching the underlying closed-loop logic, and this kind of parameter-level configuration is routine and expected. Enabling a new control mode is different: even though it doesn't alter the underlying control logic, a new mode must still be tested and validated before it can be enabled — it cannot simply be switched on. Rewriting the underlying control logic itself goes a step further still, and it raises the probability of faults instead of reducing it; this layer stays with the platform vendor. The appropriate requirement should address configurability, parameter access, and validation evidence for any new control mode, rather than prioritizing the ability to rewrite control logic.

© 2026 Stem, Inc.

SECTION 4

The Optimization Decision

Among all the layers in the stack, optimization is the one buyers most often misplace, and it warrants separate treatment because the cost of misplacing it is high and because a hybrid plant places more demands on the layer.

4.1 Where optimization actually resides

Confusion about optimization is understandable; often, when a buyer states a need for "an EMS that optimizes revenue," they are describing a desired outcome and assuming the function resides within the system being purchased. Sometimes it does. Often it does not. Many platforms deliver the control stack up through the grid interface and then connect to a separate optimization layer, such as an in-house trading desk, a third-party scheduling service, or a managed service. In these cases, the platform enables optimization without directly performing it.

The distinction is highly relevant because the optimization layer is where the project's revenue model is executed, and for a hybrid project, that model is more complex. Coordinating PV and storage to capture energy that would otherwise be clipped, to firm generation, and to participate in wholesale electricity markets (whether administered by an ISO or RTO in the United States, a transmission system operator in Europe, or a market operator such as AEMO in Australia) requires a layer that understands both resources together.

Where the revenue model lives

The optimization layer is where the project's financial model is actually executed. A gap here does not show up in commissioning; it shows up in revenue.

Across major markets, regulators have moved to open wholesale electricity markets to storage participation. In the United States, FERC Order No. 841 requires grid operators to allow storage resources to compete for capacity, energy, and ancillary services, and Order No. 2222 extended comparable access to aggregated distributed resources.[6] In Great Britain, the grid operator has developed a range of balancing and ancillary service markets accessible to storage.[7]

MARKET PARTICIPATION TERMS

ISO / RTO
The entity that manages the transmission grid and administers wholesale electricity markets in a given region. Examples include CAISO, PJM, ERCOT, and MISO for the US. Equivalent bodies include National Grid ESO and Ofgem in Great Britain, ENTSO-E member transmission system operators across continental Europe, and AEMO in Australia.

Ancillary services
Grid services beyond energy delivery — including frequency regulation, spinning reserves, and voltage support — that storage is well-suited to provide and that often represent a significant share of project revenue.

Clipping / clipping recapture
When a solar array's peak output exceeds what its inverter can convert, the excess is lost. A co-located battery can absorb that excess, capturing value a standalone solar plant cannot.

When these definitions matter:
In a hybrid plant, these revenue streams depend on the PV and storage systems being coordinated together. They are not automatically available just because both resources are present on site.

© 2026 Stem, Inc.

In Australia, AEMO administers energy and frequency control ancillary service markets in which storage can participate.[8] The specific rules differ by market, but the commercial opportunity is consistent: a hybrid plant with access to the right optimization layer can capture value across multiple revenue streams simultaneously.

The risk of getting the optimization distinction wrong is direct: if the optimization layer is mis-scoped or lacks the interfaces to coordinate both resources, the stacked revenue streams the project's financial model assumes simply never materialize, and that shortfall is discovered only after the plant is earning.

4.2 The highest-priority scoping question

The first scoping decision to resolve, before any conversation moves to protocols and control modes, is whether optimization is within the scope of the EMS being purchased or is handled separately.

- **If it is in scope,** the buyer should confirm precisely what the platform performs independently for a hybrid plant and what it delegates.
- **If it is handled separately,** the buyer should confirm that the EMS provides documented, accessible interfaces for an external optimizer to coordinate both PV and storage.

Either arrangement is workable. An unexamined assumption is the condition to avoid, because the optimization layer is where the project's revenue model is executed and a gap there is discovered at the worst possible time: after the equipment is installed and the plant is energized, when interfaces are fixed and contracts are signed, so closing the gap means a costly retrofit or renegotiation while the plant is already expected to be generating revenue.

SECTION 5

Integrated Procurement Versus Assembling the Layers

The placement of the scope boundary interacts with a second decision: whether to purchase equipment from a factory integrator with much of the control stack pre-assembled, or to assemble the layers from separate suppliers. This decision is more important in a hybrid project, where the PV inverters and the battery system often come from different manufacturers.

PROCUREMENT TERMS

Factory integrator:
A battery OEM or system integrator that delivers the BESS as a pre-assembled package with some or all of the control stack already incorporated.

Division of Responsibility (DOR) matrix:
A procurement document that explicitly assigns ownership of each system component and interface to a specific party. Recommended by EPRI's Energy Storage Integration Council as a standard RFP supplement.[3]

When these definitions matter:
Without a DOR matrix, ownership of the seams between components is assumed rather than assigned, and assumed ownership is where gaps and disputes concentrate.

© 2026 Stem, Inc.

5.1 Purchasing an integrated package

Procuring an integrated package from a single OEM reduces the number of parties the buyer must coordinate. The considerations to evaluate in this case are transparency and data access. When several layers arrive within a single package, the buyer should confirm in writing what can be observed, what can be configured, and what data can be exported, before executing the contract. The questions that recur in the field cluster along the lines of:

1. How compliance is managed across multiple vendors and contracts
2. How warranty obligations are tracked and validated as equipment changes hands
3. How visibility is maintained into the layers the buyer does not directly control

In a hybrid project, these questions multiply because the generation and storage portions may carry separate warranties and separate compliance obligations.

5.2 Assembling the control layers separately

Assembling the control layers from separate suppliers provides greater control and flexibility at the cost of additional integration work and additional interfaces to manage. For buyers who choose this path, the timing of the EMS decision is more consequential than is generally appreciated. When the equipment is not pre-integrated, and particularly when PV and storage come from different OEMs, the EMS defines how the separate components communicate and how they are coordinated as one plant. Resolving the EMS decision early therefore shapes the equipment selection rather than being constrained by it.

> **Timing matters**
>
> Resolve the EMS decision early. When equipment is not pre-integrated, that decision shapes equipment selection rather than being constrained by it.

For the buyer, it's critical to understand that neither path is superior in all cases. The appropriate choice depends on the project's structure, the buyer's capacity to manage interfaces, and the value the buyer places on long-term flexibility across a multi-vendor hybrid plant. The avoidable error is making this choice implicitly and as a byproduct of which vendor was engaged first.

© 2026 Stem, Inc.

SECTION 6

A scoping framework: seven questions to resolve first

The purpose of the preceding sections is to allow a buyer to enter vendor conversations with the relevant boundaries already established. The following sequence of questions supports that outcome. A buyer should resolve them with the internal team first; the answers constitute the requirements.

- **Which layers of the stack are within scope, for both PV and storage?**
  Identify them explicitly, from the Battery Management System through optimization, on each side of the hybrid project. Mark which the OEMs will deliver, which are expected from the EMS vendor, and which the buyer will handle internally. This single step prevents the majority of downstream confusion.

- **Is this a hybrid design (single PPC coordinating PV and storage) or a co-located design (two separate assets sharing a point of interconnection but dispatched independently)?**
  a. Treat hybrid, unified coordination as the default, and require any proposed co-located split to be justified by a specific requirement.
  b. If a split is used, confirm a top-level EMS/Plant PPC still sits above both the PV PPC and BESS PPC to coordinate their combined response at the POI.

- **Is optimization in scope or handled separately?**
  Resolve this before any other technical question, and confirm the optimizer coordinates both PV and storage, whether it sits inside the EMS or connects to it.

- **What does the project's financial structure require?**
  A merchant project built for low capital cost and a tolled project built for low operating cost will have different requirements. The financial structure should inform the technical scope.

- **Integrated package or separate layers?**
  Decide deliberately. For an integrated package, secure transparency, data access, and warranty terms in writing for both generation and storage. For separate layers, resolve the EMS decision early, because it shapes the equipment selection.

- **What must be configurable, and how is it validated?**
  Specify the control modes and parameter access the site requires, and request evidence of validation against applicable grid codes. Do not specify the ability to rewrite control logic.

- **Has complexity been justified?**
  Each additional controller, interface, and custom function must be commissioned, maintained, and supported for the life of the project. The default should be the simplest design that meets the defined requirements, with the burden on the vendor to justify added complexity.

A buyer who enters vendor evaluation with clear answers to these questions is no longer subject to vendor-led scoping. The conversation shifts from the vendor defining the buyer's needs to the buyer defining those needs and assessing whether each vendor can meet them.

© 2026 Stem, Inc.

SECTION 7

Conclusion

The principal challenge in defining EMS requirements is not technical. It is the recognition that "EMS" denotes a stack of distinct control layers, that a solar-plus-storage plant spans those layers across two resources, and that the buyer's responsibility is to decide, in advance of vendor evaluation, which layers fall within the system being purchased and which do not. A buyer who establishes that boundary clearly, treats unified PV-and-storage coordination as the default, settles the optimization question early, resists unnecessary customization, and allows the project's financial structure to guide the technical scope will conduct a materially better procurement.

> **Defined up front, this work produces shorter vendor conversations, proposals that can be compared on equivalent terms, and decisions that hold up over the life of the project.**

References

1 International Energy Agency, "Battery Storage is Scaling Up and Taking on a Larger System Role," *Global Energy Review 2026*, IEA, Paris, France, Apr. 2026. Online. Available: https://www.iea.org/commentaries/battery-storage-is-scaling-up-and-taking-on-a-larger-system-role

2 R. Wiser, M. Bolinger, and W. Gorman, "Queued Up: Characteristics of Power Plants Seeking Transmission Interconnection," Lawrence Berkeley National Laboratory, Berkeley, CA, USA, 2024. Online. Available: https://emp.lbl.gov/queues

3 Electric Power Research Institute, "Energy Storage Technical Specification Template and Energy Storage Request for Proposal Guide," Energy Storage Integration Council, EPRI, Palo Alto, CA, USA, 2015. Online. Available: https://www.epri.com/pages/sa/epri-energy-storage-integration-council-esic

4 E. Muljadi, M. Singh, and V. Gevorgian, "Demonstration of Essential Reliability Services by a 300-MW Solar Photovoltaic Power Plant," National Renewable Energy Laboratory, Golden, CO, USA, Rep. NREL/TP-5D00-67799, 2016. Online. Available: https://docs.nrel.gov/docs/fy17osti/67799.pdf

5 B. Kroposki et al., "Fast Grid Frequency Support from Distributed Energy Resources," National Renewable Energy Laboratory, Golden, CO, USA, Rep. NREL/TP-5D00-71156, 2021. Online. Available: https://docs.nrel.gov/docs/fy21osti/71156.pdf

6 Federal Energy Regulatory Commission, "Electric Storage Participation in Markets Operated by Regional Transmission Organizations and Independent System Operators," Order No. 841, 162 FERC ¶ 61,127, Feb. 2018. Online. Available: https://www.ferc.gov/news-events/news/ferc-issues-final-rule-electric-storage-participation-regional-markets

7 National Energy System Operator, "Potential Electricity Storage Routes to 2050," NESO, Warwick, UK, 2023. Online. Available: https://www.neso.energy/document/273166/download

8 Australian Energy Market Operator, "Fast Frequency Response," AEMO, Melbourne, Australia, 2023. Online. Available: https://aemo.com.au/en/initiatives/major-programs/fast-frequency-response

About Stem, Inc.

Stem (NYSE: STEM) is the operating system for clean energy: where assets connect, sites are optimized, and value is captured.

Stem is a global leader in clean energy intelligence and the operating system for the energy transition: where assets connect, sites are optimized, and value is captured. Powered by decades of operational experience and purpose-built AI, our PowerTrack™ platform is the layer that sits between raw data and confident action, making every decision-maker smarter, every asset more productive, and every transition more certain.

From asset monitoring to portfolio optimization and control, PowerTrack delivers everything energy professionals need in one place. With projects managed across 55 countries, we have spent nearly 20 years empowering asset owners, operators, and developers to take command of what comes next, with the certainty to act, the trust to commit, and the intelligence to lead.

Stem, Inc.

1400 Post Oak Blvd, STE 560, Houston, TX 77056
866.303.5668 | sales@stem.com

© 2026 Stem, Inc.