The full-stack data center orchestrator

Data center delivery has become a multi-trillion-dollar coordination problem that the current ecosystem was not built to solve. The industry is in the midst of a historic buildout. Upcoming PwC research estimates a cumulative $5.1 trillion may be invested globally in these infrastructure projects in the five-year period ending 2030. Around $32 trillion could be spent building out the backbone of the digital economy over the next 25 years.

The question is not whether capital will flow into the sector, but whether that capital can be translated into reliable megawatt capacity that shows up as usable compute. Owners, developers, and hyperscalers need power, cooling, controls, and IT brought online as an integrated system, within a defined performance envelope, and repeatable across multiple builds. Yet even with strong specialists across the ecosystem, delivery still breaks at the interfaces: utility timelines drift from design and procurement reality, specifications get changed after long-lead orders have been placed, and “construction complete” fails to become “system accepted and operating.” When completion of a multi-billion-dollar data center campus slips by a quarter, the cost isn’t limited to carrying capital for longer. Delays can also force companies to forfeit the revenue and strategic commitments tied to the capacity that they have already sold forward.

This challenge points to the need for a single accountable orchestration layer, one designed to manage the seams across the delivery stack even if it does not own every scope item itself. PwC refers to this role as a full-stack orchestrator. It is not a traditional engineering, procurement, and construction (EPC) model, nor is it simply a program manager. It is a delivery platform that sets the standards, manages the integrated schedule, governs risk and change, and defines how acceptance is measured across the project. Its role is to make sure one supplier's delay or design change does not quietly disrupt the performance of the whole project.

The case for this delivery model rests on three claims:

- That the underlying need is large and urgent enough to justify a new delivery construct for the industry
- That a unified orchestration outperforms the bifurcated alternative
- That an orchestrator platform can capture durable value

In the sections that follow, we explore each claim, starting with why demand and delivery complexity have outgrown today's operating model, then examining why partial integration efforts still leave the highest-risk seams unmanaged, and finally showing how a full-stack orchestrator can compound advantage over time.

Why now? Data center demand is outpacing traditional delivery models

The scale of planned investment is only part of the story. What matters just as much is the form that demand is now taking. The market is moving away from isolated builds and toward multi-phase, multi-hundred-megawatt campus programs with committed delivery dates and limited room for delay.

Global data center peak power consumption stood at roughly 95 GW at the end of 2025. PwC projections forecast an increase in peak power consumption to 164 GW by end of 2030, driven by approximately 13% CAGR in data center power consumption attributable to AI-driven demand.

Hyperscalers are now purchasing hundreds of acres in single transactions, and each campus is typically built in phases. Within those campuses, roughly 65% to 75% of capital goes to compute hardware, 15% to 20% to power and energy infrastructure, and 10% to 15% to building and site work. Because the longest lead times often sit in the power and energy layer, infrastructure dependencies increasingly determine whether a project moves on schedule.

The four unsolved problems

The demand signal explains why the pressure on delivery has intensified. The next question is where projects are actually breaking down. The answer is not random execution failure. It is a set of recurring breakdowns at the points where responsibility passes from one party to another.

Four problem areas show up repeatedly across the delivery stack. Each sit at a critical interface, and in today's market no single party owns each from end to end.

- **The power grid:** Analysis from the **Lawrence Berkeley National Laboratory** demonstrates the average time projects spend in interconnection queues remains high in most, but not all, regions. The median project built in 2025 took 61 months from the interconnection request to commercial operations, compared to 36 months in 2015 and 22 months in 2008. At the same time, data center demand is on track to reach up to 12% of US electricity by 2028.
- **The equipment: Long-lead electrical equipment** is now shaping project schedules and cost in ways many delivery models were not built to handle. Order-to delivery lead times have stretched to about three years, with the cost of electrical equipment (switchgear, circuit breakers, etc.) increasing 20%-100%. Teams are often forced to place orders before specifications are fully approved. Order too late and the schedule slips. Order too early and downstream redesign and rework become far more likely.
- **The rack:** Power densities are climbing, 48 V architectures are standardizing, and cooling technology is shifting. Because power delivery, thermal design, and chip capabilities are co-evolving, a change in any one of them cascades into the others. Cooling architecture selected at design lock is routinely mismatched with the GPU generation that actually arrives.
- **The handoff:** Each OEM commissions its own scope, but no party governs the cross-system integrated test. Readiness is more complex than getting power to the building, because the power-to-compute integration is tightly coupled and largely unplanned in the traditional handover. The last mile from facility-accepted to cluster-producing-useful-GPU-hours then becomes an unbudgeted sinkhole.

Why alliances are not solving those problems

If the four unsolved problems are concentrated at the seams, the obvious question is why the market has not already closed those gaps through partnerships and alliances. In reality, it has tried. The issue is not a lack of capable participants. It is that most of today's alliances expand coverage of the stack without creating a single operating model across it.

Six provider archetypes now claim meaningful portions of the delivery stack.

- Grid and substation OEMs own the equipment up to the fence line
- Industrial energy alliances deliver the distribution, UPS, and cooling plant inside it
- Behind-the-meter and energy-island packages from gas turbine OEMs to battery energy storage systems (BESS) integrators cover on-site generation and storage up to the substation
- EPCs and construction primes execute design-build with a fragmented subcontractor base
- Controls and digital platforms supply the building management systems (BMS), the engineering project management services (EPMS) and digital twin tooling.
- AI compute and cloud platforms reference architectures cover the right-hand side of the stack, from pod architecture through cluster bring-up

These alliances are an important step in the data center ecosystem. But operationalizing the synergies they describe requires cross-functional governance, joint commercial terms, and shared accountability for integrated MW. That model is something alliances have yet to build. As a result, the seams between archetypes are precisely where the four problems live, and the buyer is left to do the integration ad hoc, losing scalability and confidence in project timelines.

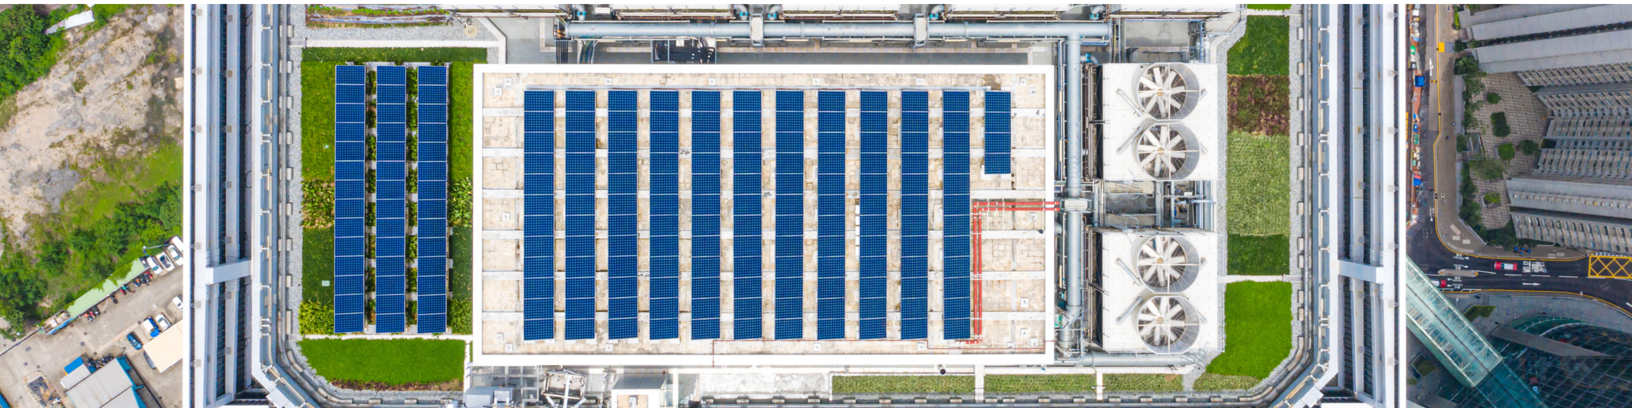

What the orchestrator is—a delivery platform for complex data center infrastructure

If alliances still leave the highest-risk seams unmanaged, then what is missing is not another specialist. It is a party responsible for the performance of the whole. That is the role of the full-stack orchestrator: to bring usable megawatt capacity online on time, within a defined performance envelope, and across repeated builds.

That accountability is not new. Sophisticated data center developers already carry it today by taking schedule risks, absorbing penalties, and tying their investment thesis to cost and delivery certainty. What has changed is the difficulty of fulfilling that role. As projects have become larger and more interdependent, delivery certainty has become harder to achieve through traditional coordination alone.

The full-stack orchestrator is a response to that shift. It owns a small set of critical governance mechanisms and delivery artifacts, then uses them to coordinate a broad ecosystem of utilities, OEMs, EPCs, commissioning teams, and compute stakeholders. In other words, the orchestrator does not replace specialists. It makes their work add up to a predictable outcome.

The full-stack orchestrator is not a traditional EPC model. An EPC executes a defined scope against an approved design. The orchestrator operates at a different level, setting design standards, managing the integrated schedule across scopes, and governing the interfaces between EPCs, OEMs, utilities, and technology providers. The EPC remains essential, but as one participant in a broader delivery system rather than the system itself.

Nor is the orchestrator simply a program manager. Program managers typically advise, coordinate, and report. The orchestrator carries a higher level of accountability: it owns the risk model, the acceptance criteria, and the change-control process, so that when something breaks at a seam, it is responsible for resolving it rather than simply documenting it.

The five control plane artifacts

A credible full-stack orchestrator is defined by ownership of five artifacts that run across the delivery stack while suppliers and partners execute within capability rows.

Artifact	Core function	Typical current ownership
Reference design governance	A versioning and change-control protocol for reference designs, defining which version each build commits to at design lock, how mid-program updates are evaluated, and how accepted changes flow through to schedule, budget, business case, and investment thesis.	The largest self-builders run internal versioned libraries—about half of hyperscale capacity. The rest, including their own leased builds, stay fragmented across buyers, EPCs, and OEMs.
Integrated master schedule and lock-in gates	A single critical path that ties utility interconnect, long-lead procurement, construction sequencing, and commissioning readiness together. Lock-in gates define the point of no return for each major decision.	While self-builders maintain a master schedule, utility interconnection delays and customer design changes challenge effective management and end-to-end oversight.
Seam ownership model and change control	A documented map of each interface between suppliers, stages, and capability domains, with a named owner at each seam and a formal protocol for evaluating the cascade impact of any proposed change.	Owned by almost no one. A few build-to-suit contracts and name interfaces; cascade impact stays unassessed and reactive.
Acceptance criteria and commissioning dashboard	Objective, data-backed definitions of done at each level, from factory acceptance testing through integrated systems testing. A real-time readiness dashboard tracks status by system.	Commissioning specialists and self-builder internal teams go furthest. Integrated systems testing stays fragmented, and the facility-to-GPU-hours last mile is unowned.
Control tower	An integrated view of schedule, procurement (long-lead heatmap, release cadence, configuration baselines), and site configuration. The control tower is the orchestrator operating cockpit.	Controls and digital-twin platforms supply the tooling, but no party operates it as a cross-party cockpit. Absent in most programs.

The seam map—identifying the critical interfaces in data center delivery

The most critical points in the lifecycle are not within stages but at the transitions between them, where accountability shifts, information degrades, and assumptions go untested. Eight seams produce most of the schedule loss in a typical campus build.

Transition	What breaks	What the orchestrator does
Site to permitting	A site is selected on land quality without a permitting timeline assessment, leading to six to 12 months of foreseeable entitlement delays.	A standardized deliverability scorecard screens permitting complexity, AHJ engagement, and code-aware design variants before site control is taken.
Permitting to utility	Permits are secured without coordinating utility queue position, and permit conditions conflict with the interconnect design.	The integrated master schedule links permit conditions directly to interconnect design assumptions as mandatory deliverables.
Utility to design lock	Utility scope is finalized on voltage assumptions that subsequently change in design, or design is locked without a confirmed energization timeline.	The orchestrator identifies the design decisions contingent on unresolved utility scope and clearly documents the risk, allowing the remainder of the design to proceed on schedule.
Design lock to procurement	Specifications are still churning when 18-plus-month orders must be placed. Late orders slip the schedule; early orders create rework.	A release cadence protocol defines spec stability requirements for each long-lead category, and a long-lead heatmap tracks order readiness and flags misalignment early.

Transition	What breaks	What the orchestrator does
Procurement to construction	Equipment arrives before the site is ready, or sequencing fails to protect commissioning prerequisites (clean rooms, cable pathways, isolation boundaries).	A constructability review at design lock and a sequencing plan explicitly protect commissioning readiness as a first-class constraint.
Construction to commissioning	The facility is complete but not IST ready. Each OEM claims its scope is done, but cross-system interfaces are unproven and integrated test cannot start.	A single IST governance protocol maps each OEM's field service into a unified test plan with readiness gates that the orchestrator owns.
Commissioning to operations	Documentation is incomplete, baseline telemetry is missing, and runbooks are untested. An unplanned gap appears between facility accepted and cluster producing useful GPU-hours.	Handover criteria covering runbooks, SLO baselines, spares, telemetry validation, and cluster bring-up smoke tests are part of the acceptance package.
Site operations to portfolio orchestration	Each campus is operated as a standalone facility. Cross-fleet opportunities to optimize power, workload scheduling, spare pools, and refresh cycles stay invisible to any party with authority to act.	A portfolio operating layer integrates real-time data across sites—grid, generation, cooling, GPU workload—and runs the optimization decisions out of reach at the single-campus level: demand response, workload-aware scheduling, coordinated maintenance, and synchronized hardware refresh.

Why a unified approach beats the bifurcated alternative

The industry is not stuck on a single answer. It is moving along a four-stage maturity ladder. Stage 1 (fragmented delivery) is the legacy approach in which the buyer is the de facto integrator. Stage 2 (alliance and partnership stacking) is today's reality: alliances cover more of the stack, but operationalizing the synergies requires a delivery model none of them has built. Stage 3 (power and technology infrastructure partnership platforms) is today's aspiration, with a power-side and a tech-side orchestrator splitting the stack between them. Stage 4 (unified full-stack orchestration) is the end-state model.

> "The orchestrator turns fragmented delivery into a governed system where complexity is managed, not absorbed by the owner."

Manas Satapathy
Principal, IPS and EUR Enterprise Function Strategy, Strategy& and PwC US

Delivery models are evolving toward integration with the highest-risk seams spanning where today's splits are drawn.

4 **End-state operating model**
Unified full-stack orchestration

One party drives repeatable execution due to ownership of reference designs, integrated master schedules, acceptance criteria, and develop and operate control towers across the full delivery stack, orchestrating EPCs, OEMs, and specialists.

3 **Today's aspiration**
Power and technology infrastructure partnership platforms

A power-side and tech-side orchestrator split the stack to provide a nearly integrated delivery model; Highest-risk seams—utility-to-design, electrical-to-cooling, construction-to-commissioning—sit exactly at the boundaries.

2 **Today's reality**
Alliance and partnership stacking

OEM alliances and "Energy Island" packages bundle complementary capabilities, but each solves a piece—operationalizing synergies requires an operating model that none of the alliances have built yet.

1 **Legacy approach**
Fragmented delivery

Developer, EPC, and OEMs operate on independent timelines with no integrated schedule, and when it works, it's because the buyer quietly acts as the integrator.

Where the bifurcated model breaks

The case for Stage 4 over Stage 3 is not theoretical. The most critical seams sit exactly at the boundary the split draws.

- **Utility interconnect to design lock:** Neither party owns the reconciliation when the one-line is evolving while utility scope is being finalized. Each orchestrator has a defensible reason to wait for the other to commit, and the gap is filled by buyer escalation rather than by governance.

- **Electrical distribution to cooling architecture:** At AI densities, power and cooling are co-designed, so a contractual interface specification is brittle the moment either side changes. At the current pace of the chip density roadmap, one side always changes.
- **Construction to commissioning:** Integrated systems testing is contested by design, because neither orchestrator has authority to impose a unified test plan on the other's equipment and OEMs. The result is the handoff failure mode described in the first section of this report.

Governing three seams across two independent orchestrators requires a coordination layer that, in practice, recreates most of the unified orchestrator artifacts without the authority, the data integration, or the feedback loops. That coordination layer is either under-resourced, because neither party wants to fund it, or it quietly becomes the unified orchestrator under a different name.

The compounding platform—how full-stack orchestration creates durable advantage

Delivery certainty is the starting point, not the full story. The deeper case for unified orchestration is that it creates a platform whose advantages build on one another over time. Each completed program improves the next one—not just operationally, but economically and strategically as well.

Reference designs that improve with each build: In a fragmented model, designs are often treated as project-specific outputs. In an orchestrated model, they become living, versioned assets that improve with use. Operational feedback can be incorporated into future builds, design choices can be standardized across sites, and changes in chip density or cooling requirements can be absorbed more systematically. As a result, the orchestrator that has built ten campuses operates from a fundamentally different reference design library than the one building its first.

Optimization that fragmented models cannot reach: Unified visibility across grid, on-site generation, battery state of charge, cooling thermal mass, and workload scheduling enables sophisticated demand response and workload-aware power planning at an industrial scale. Companies demonstrate this at single-site scale. An orchestrator operating across a multi-hundred-MW portfolio does it across the fleet. None of it is reachable when power and compute sit in separate organizations with separate data.

Portfolio economics: Procurement volume reserves manufacturing slots, standardizes specifications, and pools of spares. Commissioning protocols mature as IST duration and defect rates fall build over build, and hardware refresh becomes a planned line item ($500 to $800 million per year per 100 MW on a three-to-five-year cycle) rather than an unplanned shock.

Reducing OEM channel conflict: The most common objection to the full-stack orchestrator model is that major OEMs want direct relationships with hyperscaler buyers and view an orchestrator intermediary as a threat. In practice, the orchestrator does not need to displace the OEM's commercial relationship. It governs the interface instead. That means defining specifications, coordinating ordering cadence, setting acceptance evidence, and managing integrated systems testing, while the OEM continues to sell directly. For many OEMs, that can produce a better outcome than the fragmented status quo: more stable specifications, more predictable site readiness, and cleaner commissioning and warranty starts.

Taken together, these advantages translate into a defensible commercial position. In complex industrial buildouts, the orchestrator that owns standards, schedule, and seam integration typically captures 5% to 8% of total program spend. Applied to the $5.1 trillion of cumulative data center investment forecast through 2030, that puts the full-stack orchestration role at a $255 billion to $408 billion opportunity, captured through a blend of delivery fees per MW, managed procurement margin, and recurring service wrappers.

The path forward—how organizations can begin building orchestration now

The end state will not be realized overnight. A unified orchestrator depends on commercial models, partner relationships, and operating capabilities that no party has assembled at the scale the market now requires. The maturity ladder in Section 4 takes years to climb, not quarters.

In the meantime, the demand is not waiting. Hyperscalers and developers are committing to multi-billion-dollar campuses against fixed delivery dates today, against the current Stage 1 and Stage 2 ecosystem and all the failure modes it produces. Most are acting as the integrator themselves, standing up internal program teams and absorbing the cost when seams break.

The orchestration layer they require can begin now with the steps that produce immediate operational gains: integrated visibility into procurement and schedule risk across parties, formal interface ownership at the highest-risk seams, and disciplined change control against a defined reference design baseline.

The highest-leverage step organizations can take today is establishing a digital control tower that creates a single operating view across utilities, OEMs, EPCs, commissioning teams, and compute stakeholders. This is where integrated schedules become actionable, where design and procurement decisions are governed against common standards, and where seam ownership becomes visible in real time rather than reconstructed after delays occur.

Over time, that control tower evolves from a coordination mechanism into the orchestration layer itself—connecting delivery, operations, economics, and sustainability outcomes across an expanding portfolio. Organizations that begin building this capability now will not only reduce risk on today's programs; they will establish the operating model and data foundation required for full-stack orchestration at scale.

Buyers who begin orchestrating now build two advantages at once. They reduce delivery risk on the programs already under contract, and they build the operating muscle that compounds as the role matures from coordination toward full-stack governance across economics, engineering, and environmental performance. Waiting for the unified model to arrive is not a neutral choice; it is a choice to absorb the compounding cost of fragmentation for each quarter that passes.

Standing this up requires both an orchestration capability and a digital operating platform that can connect utilities, OEMs, EPCs, commissioning teams, and compute stakeholders around a common delivery model. The buyers that shape what orchestration looks like—by operating it, refining it, and integrating it into their delivery model—will likely be the ones who define the standard, rather than the ones who buy into it later.

Contact us

Manas Satapathy
Principal, IPS and EUR Enterprise Functional Strategy, Strategy& and PwC US
manas.satapathy@pwc.com

Matt Cardamone
Principal, Capital Projects and Infrastructure
matthew.j.cardamone@pwc.com

Hugh Le
Principal, Energy and Industrials, Strategy& US
hugh.le@pwc.com

Stephen Pillsbury
Principal, Strategy and Operations, PwC US
steve.pillsbury@pwc.com

Bridget McCarthy, Arsen Akopian, Keith Emert, and Tara Woram contributed to this report.