



The Wild Wild West Of LEGO Datacenters

Everyone Says They're Modular, Do The Vendor Claims Hold Up? Zuck's Tents, AWS's Houdini, 60GW+ Modular Capacity Tracked, Full Vendor Landscape Mapping, Vertiv's 2x Content Uplift Per MW

NICOLAS BONTIGUI, ERIC (JUNQI) WEN, JEREMIE ELIAHOU ONTIVEROS, AND 3 OTHERS

JUL 30, 2026 • PAID



86



4

Share


The Labor Problem and Modularization to the Rescue

Today we dig into the world of datacenter construction, because how datacenters are built now bears little resemblance to how the industry has historically done it. Concrete walls arrive as finished panels, mechanical and electrical rooms arrive wired, and sometimes even entire data halls arrive on the back of a truck. Some of the largest datacenters in the world are increasingly assembled the same way you assemble your new Spider-Man LEGO set, only that the bricks weigh 50,000 pounds and are a tiny bit more complex. This is the world of modular construction.

From Hyperscaler to Colos to now even the AI labs, modular construction has become the default playbook for building fast. Our Modular Tracker, included in our [SemiAnalysis Industrials Model](#), tracks over 61GW of modular capacity and 1,000+ sites using some form of modularization or prefabrication strategy. Full breakdown by modular category and equipment type is included in the [Industrials Model](#). We estimate that modular penetration will reach 30%+ of total live capacity by the end of 2028.



semiAnalysis



Source: SemiAnalysis Industrials Model

Ultra-fast modular designs are increasingly the norm. Over a year ago, we were the first to call out Meta's drastic change to using "tent" buildings. As shown below, AWS is now rolling out at very large scale their own modular design codenamed "SAMDC".



人在草木间

更多数据请扫码






Source: [SemiAnalysis Industrials Model](#)

To understand the reason, we need to start looking at one of the structural bottlenecks that capitalist incentives alone cannot build past: labor.

Our recent articles have been a journey toward that bottleneck. In “[The Case for Space Datacenters](#)”, we showed the ceiling on terrestrial capacity. Last month, in “[Stop Saying Half of 2026 US Datacenter Capacity Is Canceled](#)”, we argued that most bottlenecks are misunderstood and solvable. Trade labor is an exception here, as you cannot quickly solve for a shortage of electricians and pipefitters. The race for that talent became a true constraint long ago, visible when operators like Crusoe pumped wages by 30% to bring talent to Abilene’s site, which required over 9,000 workers at its peak.



semi**analysis**



Source: [SemiAnalysis Industrials Model](#), US Census Bureau

Aiming to size the labor shortage trade by trade, we now also include the Labor Model as part of the [Industrials Model](#). It translates the state-by-state buildout from [our Datacenter Model](#) into hours of demand for every trade and sets them against reachable labor supply. To frame the problem before modularization enters the picture, the chart below is ex-modular construction, with labor demand curve assuming labor per GW stays roughly flat over the forecast period and doesn't yet reflect the benefits we'll cover later in this article. Reachable labor supply in each state, on the other hand, is affected by how much capacity is being built, and how much labor is being pulled, in other US states.



Source: [SemiAnalysis Industrials Model](#)

Electricians are a clear case, as they represent 30-40% over the total construction man hours in a datacenter project. The chart below shows an estimated electrician shortage emerging in 2027 driven by the huge mission-critical demand. On a state-by-state basis, the shortage is most acute where the buildout concentrates, like Texas and Ohio.

Source: [SemiAnalysis Industrials Model](#)

everything from wall panels to power rooms and cooling skids are built in parallel with the site and delivered as finished units. Besides, it does more than ease the labor crunch, promising large speed and time-to-build gains. And today, speed is revenue.

The shift is already underway, with Compass and Switch being the first operators to move parts of their datacenter builds offsite. Using some form of skidded solution for electrical equipment is now pretty standard for every operator. AWS's Project Houdini prefabricates the white-space buildout and collapses the time before servers go in from months to weeks. Meta is standing up fabric-clad "Tent"-like halls. And a wave of new entrants, both in the OEM and the System Integrator space, are building specifically around modular.

In this deep dive we rebuilt the modular case bottom-up against some of the speed and cost claims made by vendors like Vertiv or Schneider, finding that modular construction can compress the construction window by ~36%, or 7-9 months, and is ~8% cheaper on a Capex/MW basis. We also analyze how vendors like Vertiv are able to expand their value capture per project by offering the full stack solution, going from their historical ~\$3.5M/MW content to ~\$7M/MW with the modular solutions.

Source: [SemiAnalysis Industrials Model](#)

The problem, though, is that today everything seems to be modular and vendors, EPCs and Colocation providers many cases are describing entirely different things. To bring some order to this wild west, this article unpacks what modular actually means. We map the

For subscribers, we focus on the main beneficiaries and break down how each player is positioned, from the public names (FIX, STRL, PWR, VRT, SU, FLEX...) to private challengers such as Infra Partners, Bladerrum, and Faith Technologies, distilling the key insights from our recent Core Research subscriber notes on Comfort Systems: “*Modular Capex Is The Moat*” and Sterling Infrastructure: “*Winning Where It Counts: Quadrupled TAM via Texas, Pacific Northwest, and the Midwest; 2X Content per MW from CEC Attach; ~\$6B Run-Rate In View*”.

To start off, we’d like to thank QTS, EdgeConneX, Aligned Data Centers, Schneider Electric, Applied Digital, DG Matrix, Aran Industries, Karman Industries, Radiant, and Rajat Bhagat for their contributions and insights during the preparation of this deep dive.

The Modular Taxonomy

- **Prefabrication** is the broader concept: any part of the build manufactured offsite and delivered ready to install. It is a statement about where the work happened, not about the shape of the thing.
- **Modular** is narrower. It refers to the actual self-contained units (rooms, boxes, blocks) that ship complete and get bolted together on site. Every modular unit is prefabricated, but prefabrication may not be directly modular.

Source: [SemiAnalysis Industrials Model](#)

Hold onto that distinction, because it is the spine of everything below. From here on, we will walk through the landscape the way a datacenter is actually built up, then we will go into details on the taxonomy that makes up the modular market starting from the ground up.

Understanding The Datacenter Anatomy

At a high level, a datacenter can be seen as simple three stacks of layers: Site, Shell, and Systems.

At the bottom we have the site, or the physical land of the datacenter buildout. This is where grading, wiring, and foundation building take place. This layer cannot be modularized because you have to physically break ground on a parcel of land and pour foundations into it on the set up.

equipment and subsystem, including all the mechanical and electrical systems, sit.

Source: [SemiAnalysis Industrials Model](#)

Considering that the site itself can not physically be moved, prefabrication strategies are focused on the other two layers, and we take them in that order, working from the outside in.

Modularizing the Datacenter Shell

The shell is the structure, walls, and roof that hold the datacenter up and keep the weather out. It generally follows either a frame-and-skin design, where the structure and cladding are separate, or a load-bearing panel design that combines both.

In a traditional build, both the skin and the frame need to go up on site. Crews break ground, pour the foundation, then form and cure concrete right where the building stands, one piece at a time. A modular shell starts the same way, on a poured foundation, but from there the structure and panels arrive as finished pieces from a factory, craned and bolted into position once arrived.



Source: [SemiAnalysis Industrials Model](#)

The time saving is evident from the graphic above. In the traditional cast-in-place buildout, each pour has to reach roughly 75% of its design strength before the next can go on top. Prefabricated structure sidesteps that wait.

However, the evolution did not stop with prefabricating the same conventional building. The larger gains now come from simplifying the building itself: moving from complex multistory facilities toward repeatable single-story halls, and then toward narrower, purpose-built structures.

Phase One: Precast Industrialized the Conventional Shell

The first phase is precast concrete described above. Instead of forming and curing the full structure in the field, panels are manufactured under controlled conditions, transported to the site, and craned onto a prepared foundation.

This is not new. Northern Virginia has used precast extensively for years because construction labor was already constrained. CloudHQ's two-story LC-2 facility in Ashburn is a representative example: its load-bearing shell supports long, column-free spans and enough structural load to place mechanical equipment on the roof.



Source: CloudHQ Datacenter Crogan

Nevertheless, the building still took roughly 18 to 20 months to deliver. Precast reduced field forming and curing, yet the underlying facility remained a large, multistory structure facility.

Tilt-up concrete follows a similar logic but casts the panels on the building slab rather than in a remote factory.

Source: Tilt up Panels at DPR Construction Ashburn Virginia



weather.

Phase Two: Simplifying the Building

The second phase is where design changes happen. In order to further speed up the time, the industry turn towards alternating the design for simplicity. These buildings use regular structural bays, fewer architectural features, and standardized exterior panels. Steel is often favored because the frame can be fabricated off site, shipped efficiently, and bolted together quickly across a large flat campus.

At the light end is the pre-engineered metal building, or PEMB built in three parts:

1. The primary frame serving as the structural skeleton
2. The secondary frame tying the main frames together, these are lighter steel member like roof purlins that span between the primary frames.
3. The skin, keep in mind this is different from the frame. They are the thin light weighted metal panel whose job is to protect the interior from extreme weather conditions

This is the fastest option since all parts arrive cut, punched, and labeled, a crew simply need to bolt them together on the site. Furthermore, a light steel structure needs far less material than concrete.

More demanding halls use structural steel - heavier, hot-rolled beams and columns fabricated off site and bolted into a rigid frame. It costs more than a light PEMB but supports wider spans, heavier loads, and more complex layouts.

QTS's Cedar Rapids campus shows the speed and scale this approach can unlock. The current 420 MW phase spans approximately 2.8 million square feet and uses roughly 28,000 tons of structural steel. QTS moved from groundbreaking to topping out in about five months, with the broader building delivered in approximately 11 months.

Source: [SemiAnalysis Industrials Model](#)

The exterior is then closed with prefabricated cladding, most commonly insulated metal panels. Similarly, we see this approach with Crusoe's Stargate campus in Abilene, each building used roughly 672 factory-made panels. The panels were fabricated in under 40 days and installed at approximately 15 to 20 per day, helping bring each building to a dried-in shell in under eight weeks.



Source: Crusoe Stargate Campus

The speed advantage therefore comes less from steel itself than from what it enables: simple single-story halls, repeatable structural bays, fewer field interfaces, and a supply chain that can be reproduced across markets. It can also reduce labor and structural material per MW compared with a dense multistory design.

The main trade-off is land. Single-story campuses require more acreage, but that is often acceptable in newer AI markets where land is cheaper and deployment speed matters more than maximizing MW per acre.

Phase Three: Purpose-Built Rapid-Deployment Shells

The third phase pushes simplification further by designing the enclosure around a more specific deployment model. Narrower, lighter structures can reduce the amount of conventional shell work and support faster repetition, although tighter optimization may leave less flexibility for future equipment or layout changes.

Meta's rapid-deployment structures at Prometheus campus in New Albany are the most visible extreme. The aluminum-framed, fabric-clad halls provide enclosure and weather



buildout was announced in July 2025.

Source: [SemiAnalysis Datacenter Model](#), Tent at Meta Prometheus New Albany

That does not mean Meta completed a full datacenter in nine months. The tents accelerate the enclosure, not utility interconnection, power, cooling, or commissioning. They also trade away some of the durability and long-term flexibility of a permanent concrete or steel building.

AWS is moving in a similar direction with its newest modular builds. Rather than treating the shell as a large generic warehouse, AWS is using narrower and more repeatable structures organized around the systems installed inside. The result is less building per MW, shorter structural spans, and fewer interfaces for field crews to assemble.



Source: AWS Multistory Design

The common thread is that shell modularity is increasingly about design simplification, not just prefabrication. Precast moved concrete production off site but largely preserved the conventional building. Standardized steel made single-story halls easier to repeat across markets. Purpose-built structures go further by reducing the size and complexity of the shell itself.

The cost savings thus come from removing floors, reducing structural complexity, decreasing labor on site, and repeating the same enclosure / supply chain across the campus. Once the shell is dried in, the larger modularization opportunity moves inside, to the power, cooling, and white-space systems that turn the enclosure into an operating datacenter.

Modularizing the Equipment and Subsystems

Equipment and subsystems are where most of the real modularization is happening, and the offerings span an enormous range, from a single piece of equipment all the way up to an entire building delivered ready to switch on. Previous deep dives already covered in big depth the anatomy of [Mechanical](#) and [Electrical systems](#). Besides, before we start naming categories, it helps to fix some vocabularies:

Source: SemiAnalysis Industrials Model

1. **Component:** The lowest level of form factor. It is a single piece of equipment manufactured in a factory.
2. **Skid:** First common modular form. It is a group of components mounted on an open frame. Instead of shipping each piece separately, the equipment is pre-arranged, configured, and shipped together as one package.
3. **Module:** A skid but in an enclosed space. It can go anywhere from a simple power room to a prefab mechanical/electrical room similar to a skid, but only it will become a module if you put walls and roof on top of it.
4. **Container:** A specific form of module using ISO container for packaging. An ISO container is built specifically for standard shipping dimension, which means it can travel anywhere on normal transport with a truck without limitation.
5. **Prefab Datacenter block:** Facility scale buildout that stitches multiple factory-built module into a much larger facility block. The intent of these block is to serve as an end-to-end datacenter buildout

This is like a ladder, from 1-5 increasingly factory integration and scope. If you look carefully, you may also be able to realize the first 4 levels are better known as the subsystem modularization. This is where the supplier delivers one part of the datacenter as a factory-built unit. That unit is arrived assembled and tested, but still has to be connected into the broader facility before it becomes useful.

The last level and sometimes the fourth level, moves closer to a whole facility modularization. Here, the supplier is delivering a much larger portion of datacenter as an integrated product.



Source: [SemiAnalysis Industrials Model](#)

Subsystem Modularization

With the vocabulary in hand, we can now start climbing the ladder, and the natural place to begin is at the bottom, with subsystem modularization. This is the larger of the two families and where most of the market lives today. We start off with the grey space, and modular power block is where most people think of when discussing modular design.

Modular Power Blocks

Source: [SemiAnalysis Industrials Model](#)

of the most natural areas to modularized because the equipment lineup is well defined and lots of different pieces are needed to assemble the units. The electrical fit-out and commissioning take on average 5.5 - 16.7 months on a 50MW power hall.

To illustrate with an example, let's see Flex's modular power solution, built through the Anord Mardix unit below:

Source: Flex

Flex sells this in two versions, the power skid and the power module pod, which is the same lineup wrapped in a secure enclosure. Inside the room, you will find major electrical components laid out like a power train.

As you may have noticed, the module also contains the busway system like the IBAR feed above the building which bridges the power pod to the data hall. The CRAH units, or computer room air handler and fire protection system, also exist to help manage the air temperature inside the enclosed shell and protection against unexpected conditions.

By our own estimate, the power block is where the schedule payoff concentrates. Moving just the power scope into the factory, roughly ~26% of the build by content, gets a hall to IT-ready ~22% faster, ~13 months against ~16.7 for stick-build, and around ~5% cheaper per MW, largely by compressing the mechanical-and-electrical fit-out from ~5.5 months to ~2.5.

Lastly, within modular power blocks, a new subcategory is emerging, which we can define as "software defined" power routing blocks. Instead of packaging the conventional

connect grid, generation, storage, and DC loads through a common controlled power platform.

Modular Cooling Blocks and Prefabricated Cooling Infrastructure

Source: [SemiAnalysis Industrials Model](#)

Similar to a power module, the cooling block packages the datacenter's cooling loop into one integrated subsystem. At first glance, the case for modularizing cooling is weaker than for power, since there are simply fewer pieces to pre-assemble. But as primary and secondary loops grow more complex, the ability to add cooling capacity in repeatable, modular increments becomes far more attractive to operators.

Focusing on the TCS loop, the majority of modular efforts are centered on skids for CDUs. Airedale by Modine's skid-based CDU (for more on cooling systems, [read our Cooling deep dive here](#)) is an example of this category.



Source: Modine

The Airedale by Modine skid-based CDU is a 2 MW-class unit arrives on a pre-manufactured skid with everything included, from the red-and-silver cooling loop to the buffer tanks on the top right, and even the leak detection system built into the end of the skid. On site, all the crew needs to do is hook up the two loop connections and a power feed.

On the other hand, the value proposition for prefabricated cooling systems is stronger when looking at the secondary loop and the outdoor mechanical yard, where piping and other outdoor cooling infrastructure is prefabricated, significantly reducing the civil, piping, and controls work completed onsite.



Source: [SemiAnalysis Industrials Model](#)

When thinking about cooling equipment, keep in mind that much of it was originally designed for hospitals, campuses, and industrial process cooling rather than GW-scale datacenters. At that scale, large footprints and hundreds of co-located units can create issues such as hot-air recirculation and heat islanding.

Source: Vertiv's elevated chiller/cooling plant on a steel platform beside a prefab hall using mechanical yard infrastructure

The Vertiv installation above shows a conventional mechanical yard. The cooling plant sits outside the data hall as a separate piece of infrastructure, with chillers, pumps, piping, and supporting steel assembled around the building.

Some datacenter developers, like QTS, are now prefabricating most of the piping infrastructure, which allows for faster installation while maintaining high quality. That becomes especially valuable today, given the increasing piping requirements of dense liquid-cooled deployments.

A handful of new entrants are now even attacking the whole yard. Karman Industries' CO₂-based Heat Processing Unit (HPU) is a purpose-built unit borrowing SiC power electronics and permanent-magnet motors from EVs, and compact turbomachinery and advanced heat exchangers from aerospace. The HPU, which is configurable to each site, ships as an outdoor-rated NEMA skid at 4 to 5 times conventional power density and cuts yard footprint by 60 to 80%.



Source: Karman Industries

Other Modularized Options

Power and cooling are not the only parts of the datacenter that can be move into the factory. The same concept can be apply across the gray space such has the energy storage (BESS) system, water treatment skids, Fire safety system, and many more.

Many are not on the critical path for construction timeline, or they are small enough to be able to build on site. Moreover, where these systems do get modularized, they often ride along inside a large unit rather than shipping separately. For example, the Schneider EcoStruxure for example was built in with the fire protection system.

Factory Built White Space

Source: [SemiAnalysis Industrials Model](#)

The intention of the factory built white space module is to replace hand-built data hall and manual on-site wiring with a unit made factory product. This means operator can directly put compute in place without the need to figure out wiring and connection.

The whole package comes ready for the rack frames, and all the last mile connection points the rack need to operate. Think of it like a prepared envelope for compute, its organized where racks go, provide the cable and connection to how they receive power and extract heat, and include prefabricated power busway and technical water loop situated above the rack.

Source: Schneider Electric

Take Schneider's EcoStruxure Pod as an example. The image above shows the black cabinets forming the IT rack rows where products like Nvidia GPU servers would be

power to each rack, containment to capture hot air as some of the solution will still be air cooled, technical water loop to distribute liquid cooling that extracts heat, and cabling that connects each servers.

The interesting thing about this is the product design behind it. A factory white space must be able to serve different custom needs and therefore, suppliers like Schneider work with Nvidia to support more than 30+ reference designs. The buyer can essentially pick the specs that matches the chips it wants and gets a hall that is pre-coordinated with the matching required power and cooling module. In the coming section we will study in more detail how the design process takes place.

Whole Facility Modularization

Whole Facility modularization is the literal datacenter-in-a-box model. Instead of delivering individual parts, the supplier delivers a complete or near-complete datacenter block.

Source: [SemiAnalysis Industrials Model](#)

Containerized Datacenters



Source: [SemiAnalysis Industrials Model](#)

Starting with a containerized datacenters, this is the 4th part of the form factor design where the datacenter itself is packaged into an ISO-style container or purpose-built weatherproof enclosure. Like we had discussed, the reason why this option exist is for the ease of transportation.

Source: Delta All-In-One Edge Solution

As the image shown above, you can see almost everything inside the box: the IT racks, the power equipment, the batteries, and even the cooling system.

datacenter quickly. For AI workloads, the buyer is usually not a compute startup chasing scale but an asset owner that needs low-latency inference at a fixed physical location.

Source: Flex's CrownPod Craned into site

The main limitation of this buildout is density. The same form factor that makes the unit portable and fast to deploy also fixes the layout. That's why suppliers are pushing beyond the containerized model toward all-in-one prefab datacenter blocks.

All-in-One Prefab Datacenter Block



Source: [SemiAnalysis Industrials Model](#)

The all-in-one prefab datacenter block is the more ambitious version of whole-facility modularization. Here the supplier delivers a larger facility block with more of the datacenter already integrated before delivery.





Source: Vertiv

Vertiv MegaMod shows what this can look like in practice. The structure effectively is a modular datacenter with the major systems packaged into one enclosure. The center of the block contains the IT racks, where servers are installed. Above and around the racks runs the fiber optics and cable-management pathways, along the perimeter are all the supporting infrastructure systems like cooling and power units.

The system's 1 MW reference design can stretch to approximately 26.5 meters long, 24 meters wide, and 4 meters high in dimensions. While the MegaMod plus version can extend to as much as 31 meters wide. You may be wondering, if it's this big how can it be shipped to the site? In reality, the structure needs to be broken into transportable prefabricated sections, shipped through a standard heavy-haul logistics truck, and then connected and commissioned as one.

Platform Modularization and Reference Designs

The last step in whole-facility modularization moves beyond any single vendor's block into a standardized reference design for the facility itself. In the same fashion the industry has reference designs for rack systems and CDUs, Nvidia now publishes one for the entire AI factory: Nvidia DSX. It was first unveiled as an Omniverse digital-twin blueprint at GTC

The DSX reference designs are validated AI factory architectures covering compute, networking, storage, hardware cluster design, and also the facilities side, including power, cooling and controls. Even civil, structural, and architectural design. The value proposition behind is that Nvidia's DSX Max-Q maximizes token per watt within a fixed power budget, and DSX Flex facilitates the connection the facility to grid services, dynamically adjusting power draw and orchestrating demand with hybrid onsite generation.

Source: Nvidia Vera Rubin DSX AI Factory

When deploying a DSX facility, through the Omniverse DSX Blueprint, an operator first builds a digital twin of the facility, simulates layouts, power topologies, thermal behavior, and operational policies in real time, and optimizes the design before construction begins, then reuses the same validated architecture across sites. For example, CoreWeave is already using DSX Air to build and test digital twins of its AI factories.

Besides, the whole DSX ecosystem includes pretty much all the supply chain: Cadence, Dassault Systemes, Eaton, Jacobs, Nscale, Phaidra, Procore, PTC, Schneider Electric, Siemens, Switch, Trane or Vertiv. Vertiv's OneCore, for example, packages power and cooling into standardized 12.5 MW pods that can be combined into larger AI-factory deployments

EdgeConneX estimates that a common design can advance a project to roughly a 30% to 60% permit set before site-specific localization is complete, allowing substantial off-site



Vendor landscape

If you have made it this far, you should have a working feel for the categories. You should also, maybe be a little buried in names. We have put a power module from Flex, a CDU from Airedale, a data-hall pod from Schneider, and an all-in-one datacenter from Vertiv all in front of you. So before going further, it helps to step back and lay the whole market out on a single map.

Source: [SemiAnalysis Industrials Model](#)

Our universe runs to over 80 players, and laying them out this way is useful because two patterns jump out right away.

2. The same names keep showing up across columns, because a vendor like Vertiv, Schneider, or Eaton sells a power module, a CDU, a white-space pod, and a whole block all at once.

Owning the Integration: Who actually does the modularization?

The vendor landscape above maps who builds each piece, but not how those pieces become a module or who is on the hook when they do. From the solution provider's point of view it has three answers:

1. At one end the operator holds both rights: it specifies the equipment, buys it directly, and hands it to an integrator purely for assembly.
2. In the middle sits the EPC- or integrator-led buildout, where the operator still sets the performance requirements but hires an EPC or integrator to source, coordinate, and build.
3. At the far end is the OEM-led model, where a vendor like Vertiv designs and sells its own stack as one finished product, as it does with the OneCore portfolio

Source: [SemiAnalysis Industrials Model](#)

Operator-Led Modularization

This is where the operator engineers the specification themselves, procures the equipment directly as owner-furnished gear, and hands it to an integrator purely for assembly.

inventory, and lead-time risk.

Source: [SemiAnalysis Industrials Model](#)

In a supply-constrained market that risk is sharp, since the operator is competing for scarce transformers and switchgear without a vendor's allocation leverage, unless it buys at enough scale to have that leverage of its own.

That is why operator-led modular is effectively confined to the largest hyperscalers. AWS, for example, engineers its own prefabricated data-hall skids under Project Houdini and procures the equipment directly, using Cupertino Electric as design partner.

Aligned is another example of operator-led modularization, although it relies on external integration partners for manufacturing capacity. Aligned defines the architecture, owner-furnishes the major components, and controls the commissioning and quality program, while its integration partners receive, store, assemble, and test the equipment across multiple factory locations. The partner provides the production footprint, but the modular system remains Aligned's design.

System-Integrator or EPC-Led Modularization

The EPC-led modularization includes system integrators or construction companies that take mostly third party equipment and convert it into a skid/module. The integrator does the assembly, installation, factory testing, enclosure, and is the party in charge of delivering the finished skid to the end customer.

infrastructure of Quanta's Cupertino Electric, and specialized modular integrators, like PCX, Nautilus, DXN, Infra Partners, Bladeroom, etc.

This type of modularization is vendor agnostic, which means that the customer keeps more control over the datacenter design, while the contractor moves part of the construction sequence into a prefab shop. The EPC buys the equipment, assembles it, wires it, pipes it, tests it, and ships it as completed construction scope.

Source: [SemiAnalysis Industrials Model](#)

Take Comfort Systems as an example. It is a MEP contractor, not an equipment maker, so it acts as the layer that procures and assembles the gear on the operator's behalf. The operator decides what equipment it wants and Comfort Systems does everything else. It runs that work through Environmental Air Systems and TAS Energy across 3.5+ million square feet of shop floor in Texas and North Carolina.

That concept is particularly attractive to hyperscalers. A big operator usually already knows exactly what equipment and design it wants, so it has no interest in buying someone else's fixed system. Working with EPC integrators like the kind for Comfort System, the operator keeps its own design and its own gear, and simply hands the building of it to a factory instead of a jobsite.

OEM-Led Modularization

Here the equipment vendor turns its own datacenter infrastructure stack into a repeatable module or platform. Although the product can still be configured for a specific site, the



Source: [SemiAnalysis Industrials Model](#)

Vertiv OneCore is a clear example. Instead of selling single equipment devices, Vertiv combines all the layers into a single modular platform. This platform integrates Vertiv's power, thermal, cooling, and IT infrastructure technologies inside a Vertiv's supplied steel shell.

Source: Vertiv

Vertiv to capture higher content by selling the entire stack end to end, with the TAM expanding up to ~\$7M/MW for some of their full-stack solutions.

The tradeoff is capacity and execution risk. Besides the fact that these companies are going up the value-chain toward market segments they were not previously involved in, OEM-led modularization can only scale as fast as the OEM's factory capacity, supplier base, and integration teams can support. Vertiv's modular solutions run lead times of over 12 months today. This is also pushing big OEMs, not only Vertiv but also companies like Schneider and Siemens, to be selective with capacity slot allocation, requiring certain capacity minimums and favoring bigger projects. As a result, operators or developers looking for smaller scale capacity are increasingly working with the System Integrators.

The Modularization Cycle


By now our readers should be familiar with the different forms of modular solutions and the players leading this transformation. They will also have noticed how different this looks compared to traditional datacenter construction, and will have many questions about the operational implications. Those are the topics we address in this section.

Source: [SemiAnalysis Industrials Model](#)

Stage 1: Solution Design and Simulation

Engineering the solution is the first step. In a field build, some of these decisions can change while construction is underway. In a modular build however, they need to be predetermined, frozen early, and repeatable since the box itself must come in a finish block.

single-line diagram (SLD) and layout, and complete the short-circuit, protection-coordination, and arc-flash studies in tools like ETAP and PSSE. These analyses determine how the system is sized and which components can be used. Because they depend on the full electrical path from the utility connection through the downstream equipment, they cannot be completed on an isolated skid. The facility design therefore has to come first, with the module designed as part of that broader system.



Source: Aran Industries

For 415/480VAC this work is well templated and repeatable. When considering all the implications of the 800VDC transition, it is not as well-templated. A handful of reference designs exist, but none are well baked yet, so the facility-level architecture is still being worked out design by design. This is also the part of the cycle now being automated. Companies like Aran Industries are building build custom software that plugs into ETAP, PSCAD, PSSE, Revit and the other design tools, compressing what is otherwise a multi-month (>2 months), multi-engineer electrical design process into hours of compute plus a single engineer reviewing the output.

Ownership of the Buildout

Once the design is set, ownership comes down to two questions:

Those decisions do not always sit with the same party. An operator may specify a component directly, or an integrator may select it and seek approval.

Source: Schneider's Factory showcasing different modules

Colocation operators are a clean exception. A wholesale colo is not bound to any single end-user's specification, and often does not even know who the tenant will be, so it is free to choose the equipment it wants and commit to it early, without waiting on anyone's sign-off.

Even then, local availability matters, because switchgear and transformers may carry 12 to 18-month lead times, while generators can require market-specific emissions controls. Customization can also reopen engineering and add roughly eight weeks.

Stage 2: Packages and Documentation

Once the design is frozen in Stage 1, what leaves the design stage is not a single drawing but a set of documentation packages, produced by different parties in different tools:

- An **issued-for-fabrication (IFF) package** with the shop drawings telling the factory how to build the skid
- An **issued-for-construction (IFC) package** telling the site how to receive, place and connect it
- A separate **permitting and commissioning documentation set**, proving the design will pass code and testing.



Source: [SemiAnalysis Industrials Model](#)

This stage is largely a documentation and paperwork exercise. Automating the Stage-1 design model lets these packages be generated from one source rather than be recreated by hand, removing the manual redrawing between parties (IFF and IFC serve different audiences, the factory and the field, and are issued separately; they are not the same document and do not drift against one another).

Stage 3: Assembling the module and Factory Testing

Once all the prework is completed, the module still has to be physically built. Think of this step as the production line, but in a much bigger scale of a datacenter. The assembly process starts with a base, a skid, or a frame that sets the foundation of the module. From there, equipments are layered in sequence, with each component placed into its designated area like how you will assemble a LEGO building.

Source: Flex's White Space Factory Line

inspection gate where the work is checked before the module advances. The second comes once the module is complete: the whole unit is powered up and run the way it will run on site to confirm its rated load and everything is in place properly. By the time it leaves the line, the module is cabled, labeled, sealed, and proven to work on its own.

Source: Vertiv Factory Acceptance Testing

Factory testing is critical. The industry often describes this through the 1-10-100 rule: a defect that costs \$1 to fix during design or assembly may cost \$10 once production begins and \$100 after the product has shipped. Furthermore, standardization also makes testing repeatable, so same FAT procedures can be run multiple times across production.

Stage 4: Delivering and installing the modules

Now, the module leaves the factory's controlled environment and runs into the uncontrolled area of site logistics. At the scale of today's AI campuses, bringing the first block online is no longer enough, because operators also have to think about time to the last megawatt. The critical questions become how many finished modules can reach the site, how many can be unloaded and set in parallel, and how effectively limited rigging and installation crews can move from one block to the next.

A closer look at: Transport and Logistics

In our conversations with datacenter developers, logistics appears as one of the main challenges equipment modularization presents. All these skids and modules are big and

closer to their customer's sites.

Federal law fixes the no-permit envelope at 102 inches wide and 80,000 pounds gross, leaving about 24 tons for the module on a standard deck. Microsoft's Azure Modular Datacenter and Schneider's Easy Modular ride inside a 40-foot ISO container at 96 inches wide, so they cross any state line or fly in a C-17 permit-free.

Past that threshold, you need a permit, although the truth is that permit cost has little impact and the real implication is on schedule. A standard oversize permit runs just \$15-100 a state, and even the line-haul, at \$12-14 a loaded mile, makes a 500 mile move only six to seven thousand dollars a trailer.

Source: Flex's Prefabricated modular solution ready to ship

On time, loads past roughly 16 feet become superloads, triggering a bridge-engineering review that runs 7 to 21 days per state and stacks toward months across a route, and escorts force restricted travel windows. Besides, the thresholds are not even uniform, as a module engineered to clear Virginia (superload at 18 feet or 250,000 pounds) can still trip Ohio's far lower 14-foot, 120,000-pound line.

Transport also becomes part of the reliability program. A module may experience greater mechanical stress on the road than during normal operation, particularly through vibration, braking, and loading. In one validation exercise, Aligned DC shipped a 3 MW module from Utah to Omaha and back with force loggers to measure the conditions it experienced in transit.

equipment on a single trailer creates an unacceptable insured loss no insurance company is willing to bet on. The risk is not theoretical and we have heard cases where truck carrying large UPS module tipped over on a road in West Virginia while traveling toward Northern Virginia, leading to large reimbursement.

Considering all these implications, operators are designing around the haul, like AWS engineering Houdini's skids onto low double-drop trailers to stay under bridges. Also Nautilus floated its datacenter 50 miles to the Port of Stockton on a barge, and Compass put a Schneider module factory next to its Red Oak campus. That said, some operators like DXN do manage to ship containers from their factories in Perth, Western Australia all the way to the US. These are, however, mostly smaller containers.

Source: [SemiAnalysis Industrials Model](#)

Once the module reaches the site, it must be lifted onto the prepared foundation, anchored, and connected to upstream power and downstream load. Even crane selection becomes part of the design process. A crawler crane's lifting capacity falls sharply as the load moves farther from the boom: a Manitowoc 18000 can lift roughly 600 tons at 7.3 meters, but only about 10 tons at 104 meters, so the module's weight, lifting points, and final landing position determine which crane is required and how much the lift will cost, typically

IT IS BUILT.

Stage 5: Site Commissioning

With the modules in place, it's time to commission them. Factory testing proves the individual unit, while commissioning proves the system under real conditions. Operators consistently call this the single biggest gap in the modular cycle, with a full site commissioning running 3 to 8 months end to end. While some vendors claim it can be shifted offsite, in practice the parts that matter can't be, because the major energy sources (the utility feed, the generators, the BESS) only meet on site.

The levels of commissioning

The industry generally runs a 6-level ladder, while some frameworks add a Level 0 design review at the front.



Source: SemiAnalysis Industrials Model

- **L1 - Factory Witness Test (the “red tag”):** each skid or module proven as a standalone unit at the manufacturer. This is the level factory testing covers, and the level modular front loads.
- **L2 - Delivery & installation verification:** the unit is received, set, anchored and inspected on site.
- **L3 - Pre-functional / startup (the “green tag”):** each system energized and started up on its own.
- **L4 - Functional performance testing (the “blue tag”):** each system run to the switchgear, the cooling loop.



transfer, a pump takeover, a black-start.

Source: Hioki Commissioning Steps

From L2 onward the work has to happen physically on site, because the sources being tested against like utility feed, generators, BESS are only present there.

To understand where time can actually be recovered, we identify two levers. First, parallelize commission modules and subsystems as they land rather than waiting for the whole site, so L2-L4 on the early modules overlap the delivery of later ones and only L5 has to wait for the full set. Second, reuse the plant as its own load bank: the power equipment and its battery storage can serve as the load bank during commissioning, so instead of trucking in rented gensets and resistive banks that leave when the test ends, the same gear that proves the plant stays on as part of it.

Source: Electrical Engineering Portal, commissioning substation and MV switchgear

Parallel Factory Commissioning

At scale the binding constraint is usually the Level 3 point-to-point work, because every PDU, control panel, and field device must be wired, addressed, named, and verified back to the BMS. Across hundreds of PDUs each exposing dozens to hundreds of points, this is the real commissioning bottleneck. To move faster, some big datacenter operators are running parallel factory commissioning tracks alongside the standard on-site process, completing equipment verification and control checks before shipment (repeating some tests post-transport) while the site team keeps delivery checks, live interconnection, and integrated-system testing.

This consideration also matters when studying aggressive project timelines claims. Some 6-9 month schedules are achieved partly by compressing or minimizing the commissioning process. That may accelerate initial turn-up, but it significantly shifts risk into operations.

Putting the Value Proposition of Modular Solutions to Test

All modular pitches mainly rest on three claims: (1) Speed to market; (2) Quality of Build and Safety; and (3) Total Cost of Ownership. The marketed numbers are huge figures, with Vertiv claiming MegaMod at up to 50% faster on module deployment and SmartRun at up to 85%, and Schneider claiming 60% faster and 13% lower first cost on power and cooling.

liquid cooled AI hall in the US. As headline, our numbers estimate a ~56% shorter construction window for modular construction when considering the full construction timeline, and a ~8% lower all-in cost. Let's break down these figures.

Source: [SemiAnalysis Industrials Model](#)

The Speed Advantage

Speed is undoubtedly the core vendor selling point. Following our previous taxonomy, the realized savings will depend a lot on which modular solutions you use in the build. To understand where the speed savings come from, we need to understand the construction timeline in commonly four blocks:

- **Groundworks:** 4-6 months, averaging about 5 months. Modularization changes little here beyond simplifying some pad and foundation work.
- **Structure and shell:** 2.5-4.5 months, with a midpoint of about 3.5 months.
- **Mechanical and electrical fit-out:** 7-11 months. Today's baseline already includes some preassembled electrical equipment, such as LV switchgear and UPS lineups.
- **Commissioning to IT-ready:** 3-8 months, with a single-hall midpoint of about 4.5 months. Longer timelines usually reflect phased, multi-hall handovers.

Added together, a conventional 50MW building runs about 18 to 24 months from groundbreaking to IT-ready, with permitting excluded from that clock. In addition to the

vertical works rather than at groundworks. Besides, permitting adds another 12 to 15 months that cannot overlap with construction, taking the all-in timeline to roughly ~30-35+ months for a stick build against ~24-30 for modular.

In the table below, which considers a scenario of an operator going full modular, modularization can compresses the construction window to about 12-18 months, around ~36% faster than a pure stick build and about ~30% faster than today's baseline that considers some modularization.

Source: [SemiAnalysis Industrials Model](#)

The savings scale with how much scope moves to the factory. An operator taking only MEP skids, midway between today's baseline and a full modular build, lands around 17 months. That schedule can be shortened even further when deploying all-in-one prefabricated block or containerized datacenters, which can take the building window all the way down to about 12 months. Today we can hear some claims even for under 12-month deliveries.



Source: SemiAnalysis Industrials Model

Digging deeper, during the MEP fit-out, an AI data hall absorbs about 12,000 field man-hours per MW, with electrical the dominant trade. A standard 50 MW hall concentrates 600,000 field hours into one building and stacks about 300 craft workers at peak. Relocating that MEP scope to the factory, on-site hours fall about ~63% to 4,500 per MW and licensed-electrician hours about 85%. The window itself compresses by less than the hours because the relocated work now runs in the factory in parallel with sitework. In addition, moving work into the factory also reduces dependence on weather and geographic constraints.

Source: SemiAnalysis Industrials Model

opportunity cost of the delay. In our calculations we consider as a rule of thumb that today, for CSPs, a megawatt of IT load generates on the order of \$12M to 15M of revenue a year, or about \$1-1.25M per IT MW per month. That said, in such a capacity constrained market, we are seeing new deals being made on much higher disclosed pays, see SpaceX and Anthropic deal, or [as we recently covered in our Meta newsletter post](#), revenues of \$50M per IT MW. Companies like Anthropic or OpenAI are making over \$50M/MW on API.

On the COGS side, the relevant cost is GPU depreciation, at an average \$30M/MW for an Nvidia cluster depreciated over five years, roughly \$500,000/MW per month. An idle month strands at least that depreciation, on GPUs already on the clock whether or not the hall is ready, so we value an earlier month to the owner-operator at about \$500k/MW, a conservative figure since the contribution margin on live compute is higher. For a wholesale colocation operator that never owns the GPUs, the value of an earlier month is just the lease it can now bill, about \$190,000/MW per month, roughly \$190/kW/mo.

Source: [SemiAnalysis Industrials Model](#)

Applying those unit economics to the roughly 8-month lead over a pure stick build, the owner-operator captures about \$200M undiscounted across the 50 MW hall, or ~\$4M/MW.

Of course, all of this is conditional on the building being the binding constraint. Go-live is the latest of three dates, building ready, power available, and GPUs delivered, and accelerating the building earns nothing if it was not the date that bound.

Compared to the realized speed benefits, capex savings are not meaningful. We estimate that same full modular 50 MW liquid cooled hall costs about ~\$1.1M/MW less all-in, ~\$13.5M/MW against ~\$14.6M/MW, just under an ~8% delta.

When looking at the all-in content/MW for the datacenter equipment, we can separate between hardware cost and service or installation cost. The hardware piece itself does not change much: switchgear, UPS, transformers, and other core hardware cost the same whether they are installed in a field-built room or integrated into a factory skid. The gross savings come from two places mainly, the labor portion and a shorter build time.

Source: SemiAnalysis Industrials Model

We estimate that moving MEP into the factory saves about \$0.6M/MW on construction services and \$0.5M/MW on installation. The wage gap is small, as BLS puts field construction at \$34/hour against \$33/hour in the factory, though overtime and site premiums lift the effective field electrician to roughly \$63/hour. The real lever here is higher factory throughput, dropping a full field hour of scope.

Second is the effect from shorter build time. Locking factory cost earlier and compressing the field period trims escalation, contingency, change orders and site general conditions.



Source: [SemiAnalysis Industrials Model](#)

Against those savings, modular carries some penalties, with double margin being the largest. A stick build marks up installed content once, through the general contractor and its subs. A modular build adds a separate module-vendor margin on top of the site integrator's. That layer compresses when the OEM also integrates on site, as Vertiv and Schneider do on their turnkey programs. Schneider's WP163 shows this dynamic, putting module hardware about 40% above traditional and netting to 13% lower first cost, after design and install labor savings are counted.



Source: Schneider

Two smaller penalties follow, including the module premium (chassis, bracing, extra interconnects, transport and craning), and from the vendor perspective, the additional factory burden.

Source: [SemiAnalysis Industrials Model](#)

The Quality and Predictability Advantage

Factory first-pass quality, the share of test and inspection points cleared without rework, is supposed to exceed 95% in modular solutions, against a field baseline of 60-70%, thanks to the on-factory fixed work instructions. The benefit is system predictability. Flex frames this as a design choice and engineers it through design failure-mode-and-effects analysis for

That said, recent conversations point to the opposite side, with operators and MEP contractors claiming that modular solutions have proven reliability issues and haven't lived up to their claims. This eventually means not only losing all the time savings gained initially, with a field team having to go on site, but more importantly, putting the precious hardware at risk.

Re-evaluating Vendors' Claims

Our ~36% time savings, or around 8 months, lands just above Flex's published 30%+ floor on whole projects. In order to do a fair comparison, we must consider that vendors' numbers usually come from narrower scopes or composite definitions, not the end-to-end construction timings. Vertiv's 85% SmartRun claim applies to overhead busway and containment, while MegaMod's 50% claim measures module deployment against on-site build. Schneider's 60% claim is for power-and-cooling modules.

Source: [SemiAnalysis Industrials Model](#)

What Operators and Developers Are Modularizing Today

Across the market, operators and developers are not converging on one modular strategy.

Hyper-scalers most frequently build their own data centers, vertically integrate capabilities, and engineer the datacenter design around their actual workload. That's why hyperscalers tend to show the innovator traits and have leading experimental footprint.

AWS

AWS shows the fastest scaling buildout currently, adding almost 3.9 GW of capacity to the end of 2025. Project Houdini is their largest internal attempt to modularization model. Rather than designing a modular building from scratch, AWS took its standard data-hall map and re-cut it into a factory-built skid roughly 45-feet in length and weighing about 2000 pounds that is transportable on double-drop trailers. This puts Houdini squarely in the “factor-built white-space” category.

Source: AWS Project Rainier Indiana

The approach cuts deployment from as much as 15 weeks to roughly 2–3 weeks and can eliminate more than 50,000 on-site electrician hours per module. The skids are built in Houston, Salt Lake City, and Topeka, with early deployments in Texas and South Bend, targeting ~25 weeks from construction start to the first server room.

Houdini is also notable for its approval model. Instead of relying only on site-level engineering sign-off, the factory process requires separate validation of the design and the



Source: [SemiAnalysis Industrials Model](#)

Meta

Meta has recently focused on how to get the building enclosed. Its “tents” datacenters, or rapid-deployment structures, are the clearest example of shell-level modularization in the taxonomy above. A lightweight structural frame supports a tensioned fabric membrane, creating a weather-protected enclosure much faster than a conventional steel-and-concrete building.

At Prometheus in New Albany, Meta is using this strategy to pull shell construction forward

Source: Meta Tent at Prometheus New Albany

Meta has built six rapid deployment structure at this campus, with each roughly 125,000 in square feet dimension. To put speed savings into the picture, it's worth noting that the site's first five permanent buildings took Meta two to three years to complete, while the tents took a fraction of that, with eight standing by April 2026 from our satellite tracking since their announcement of the tent buildout in July of 2025.

This does not mean Meta built the complete datacenters in nine months, because the tent accelerate the shell but not the full facility. Also, the trade off on the other hand, is resilience. A fabric structure does not provide the same long-term durability or weather protection as a permeant steel or concrete shell as we had described above, it is also not built with intent to provide decade-long protection service.

GPU Clouds and Neoclouds

Crusoe

Crusoe participates in both the shell modularization and whole facility modularization taxonomy. At the Abilene Stargate campus, Crusoe accelerated the shell using structural steel and factory-made insulated metal panels. Working with Digital Building Components, each building used roughly 672 prefabricated panels. The panels were manufactured in under 40 days and installed at a rate of roughly 15 to 20 per day, allowing the building to reach a dried-in state in under eight weeks.



Source: Crusoe Stargate Abilene Campus

In addition, Crusoe's decided to vertically integrate the manufacturing process. Its 2022 acquisition of Easter-Owens brought modular datacenter and electrical-system manufacturing in-house, giving Crusoe greater control over design, supply chain, and production. The company is now expanding that capability through a dedicated Spark factory in Brighton, Colorado.

Each Spark unit are roughly one megawatt in scale and arrive substantially complete, placing Spark in the whole-facility modularization category.

The Redwood Materials deployment in Nevada shows how this model can scale. Grusoe initially installed four Spark units alongside a 12 MW microgrid and later announced an expansion to 24 units.

Hut 8

Hut 8 is buying the entire infrastructure stack. A former bitcoin miner with power and land positions, it is racing to convert those positions into leasable AI capacity, and the fastest way to do that is to buy a finished modular stack.

Source: Hut 8 Corpus Christi Beacon Point Platform

At its Beacon Point campus in Corpus Christi it runs Vertiv's OneCore to commercialize a 704MW IT lease.

The project is being designed around Nvidia's DSX reference architecture and delivered through a group of established counterparties. American Electric Power provides the utility relationship, Jacobs Solutions owns the EPCM scope, and Vertiv supplies the critical power and cooling infrastructure.

Nebius

Nebius takes a lighter approach to modularization by defining the facility around its own compute architecture and modularizing selected parts of the infrastructure such as power and cooling.

At the New Jersey campus, the facility is being built to Nebius's own design through its partnership with DataOne and is planned as a phased development expandable to 300 MW.

Source: NJ.com, Nebius/DataOne New Jersey

At Béthune, France, the company follows the same operating logic through a different construction method. The project reuses the former Bridgestone tire plant, avoiding part of the greenfield shell and permitting process. Azur Datacenter is responsible for the land, utility intake, construction, and physical plant, while Nebius focuses on the GPUs, racks, networking, and software.

Colocation Providers

Compass

Compass is the colo that has run the modular playbook longest. They industrialize the whole facility into a repeatable kit of parts, not just one subsystem. We estimate roughly 70 - 85% of each building is manufactured off-site and bolted together on site, standing up a building's framework and roof in 18 to 21 days. The kit spans the full stack:

- Shell: a rebar-free, fiber-reinforced precast shell, Compass casts itself from on-site batch plants

- Power block: a repeatable ~1.25 MW Schneider power center (Galaxy VX UPS, lithium-ion batteries, QED-2 switchgear)
- MV switchgear: modular medium-voltage switchgear skids under a Siemens deal for up to 1,500 units over five years

Source: Compass Datacenters in Red Oak

It runs this same kit across campuses, up to 360 MW at Red Oak, Texas and eight buildings across 1.8 million square feet at Goodyear, Arizona.

QTS

QTS's modular strategy is built around locking the design early, standardizing the interfaces between vendors, and purchasing critical equipment before a specific building needs it. The company maintains roughly 7 million square feet of warehouse capacity in Kansas for long-lead equipment, allowing UPS systems, switchgear, cooling equipment, and other owner-furnished components to be held in inventory rather than ordered after each customer signs.

The approach began with the Freedom Design, which keeps the building shell flexible but repeats the power architecture. Each factory-built pod combines a 1.5 MW UPS and switchgear package with a 2.25 MW generator, and the system scales in 1.5 MW increments. Freedom LC+ expands the same concept into cooling, with an architecture that can support either fully air-cooled or fully liquid-cooled deployments. QTS's newer rapid-deployment design pushes modularization further into the building itself, organizing capacity around repeatable 60 MW data-hall blocks.



facilities use tilt-up concrete, including Manassas, while larger campuses such as Cedar Rapids use structural steel.

Source: QTS Man1 Project Tilt-UP

The value is labor, safety, and schedule certainty. QTS estimates that labor pressure has increased datacenter construction costs by roughly 20–30% per MW. Moving to modular is the company is trend, and they begin by stocking long-lead equipment in advance removes another source of delay, giving QTS greater control over when the final megawatts can be commissioned.

QTS developed subsequently the “Rapids” design, which materially accelerates construction timelines. Two major AI companies have adopted this design at scale.



Source: SemiAnalysis Industrials Model

Aligned Data Centers

Aligned's modular strategy is centered on adaptability, keeping the core power and cooling architecture standardized while allowing the hall configuration to change as customer requirements and rack densities evolve. The approach began with core MEP infrastructure, including a 2 MW UPS container with integrated distribution switchgear and power. This approach is now extending into site conveyance through examples like prefabricated chilled-water assemblies. For customers able to coordinate early, the scope can extend even further into the white space through things like secondary fluid piping.



Source: Aligned Adaptive Modular Infrastructure

The other core element is Aligned's Adaptive Modular Infrastructure platform, which keeps the underlying power and cooling systems consistent while allowing the hall to shift between air, hybrid, and liquid cooling as rack densities evolve. The platform combines

- Delta³, Aligned's air-cooling system, supporting densities of up to roughly 50 kW per rack.
- DeltaFlow, its liquid-cooling platform, supporting densities above 350 kW per rack

Because the underlying chilled-water loop and facility interfaces remain consistent, the cooling mix can change without redesigning the entire hall. The components inside the modules remain standardized and do not change. An example is found in Project Caprock in Texas, 540 MW across six buildings and 1.65 million square feet.

This post is for paid subscribers

Subscribe

Already a paid subscriber? Sign in

← Previous



semi analysis

Substack is the home for great culture