JABIL | PROCUREMENT & SUPPLY CHAIN SERVICES

AI-Ready Data Centers: The New Requirements for 100kW Racks

A guide to AI-ready data center requirements for high-density compute, liquid cooling, and power-constrained expansion.

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

Table of contents

1.0 Executive Summary 3

2.0 What Changed: AI Is Rewriting Data Center Physics 4

2.1 Sizing the Energy Footprint 5

2.2 When Power Becomes a Physical Constraint 6

3.0 From Air to Liquid: The Cooling Transition 7

3.1 Direct-to-Chip (DTC) Liquid Cooling 7

3.2 Immersion Cooling 7

3.3 Cooling Becomes an Operating Model

4.0 The Power Constraint 8

4.1 The Grid Bottleneck and the Interconnect Clock 8

4.2 Grid-to-Rack Requirements and the Modern Power Toolkit 9

5.0 The AI-Ready Bill of Materials: What Becomes Critical 10

5.1 What's Different About the AI-Ready BOM 11

5.2 The Integration Challenge 11

6.0 Beyond Power: Where Schedules Still Slip 12

6.1 Talent: the High-Voltage Skills Gap 12

6.2 Permitting: "License to Operate" Becomes a Schedule Variable 12

6.3 Geopolitics and Export Controls: Supply Volatility in Critical Components 12

7.0 Readiness Framework: Qualify > Source > Stage > Install > Maintain 13

7.1 Qualify 13

7.2 Source 13

7.3 Stage 14

7.4 Install & commission 14

7.5 Maintain & optimize 14

8.0 Partnering with Jabil 15

8.1 How Jabil Helps Data Center Operators Scale 15

8.2 Go deeper with the Data Centers Market Sector Report 15

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

1.0 Executive Summary

The race for AI-ready data centers, and AI-ready data center requirements, is well underway. Hyperscalers, colocation providers, enterprises, and increasingly governments, are moving in parallel to secure the inputs that determine speed to capacity: land, equipment, skilled labor, permitting headroom, and the chips that make AI workloads economically viable.

But the resource beneath every one of those constraints, and increasingly the factor that decides who can scale at all, **is energy.**

The market is colliding with a global power crunch, as aging electrical grids struggle to meet next-generation demand amid tightening pressure tied to carbon emissions and water usage. In core hubs, grid congestion is severe enough that new power connections can take 3–5 years, pushing developers toward alternative strategies and new geographies. Jabil's latest Data Center Market Sector report projects nearly 100 GW of added capacity needed by 2030, about 100 large nuclear reactors, or roughly a quarter-billion 400-W solar panels.

From here, the Energy Web traces the real dependencies across the grid-to-chip stack, where power access, infrastructure readiness, and operating realities ultimately shape how modern data centers get built, powered, and operated.

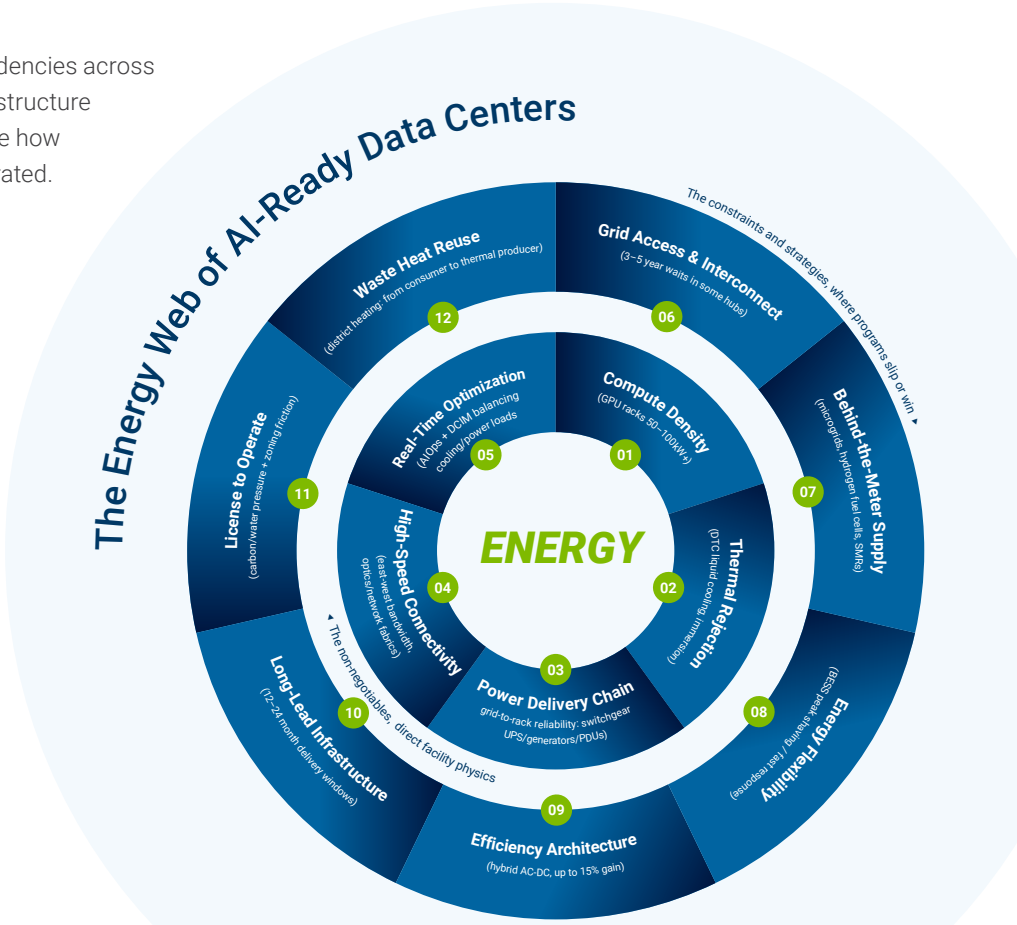

AI-ready scale is no longer gated by compute alone. Energy sits at the center. Power access, delivery, cooling, optimization, and infrastructure readiness now define speed to capacity.

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

2.0 What Changed: AI Is Rewriting Data Center Physics

AI workloads are changing AI-ready data center requirements at the level that matters most: physical density. As models scale and inference becomes continuous, the requirement shifts from "enough compute" to high-density compute delivered reliably, in clusters that can sustain performance without thermal throttling.

What changed, in practical terms:

- **Rack power density moved an order of magnitude. GPU-dense racks are designed to run massive accelerator clusters (e.g., NVIDIA H100s or AMD MI300s) in parallel. Standard enterprise racks typically draw 5-10 kW, while high-density racks consume 50 kW to over 100 kW per rack, setting new 100kW rack requirements for power and cooling. However, at the facility level, power demand is expressed in megawatts (MW), where 1 MW equals 1,000 kW. (see FIG 2).**
- **Cooling shifted from optimization to requirement.** In hyperscale AI clusters, air cooling can't handle the heat generated by dense GPU racks (often exceeding 100 kW), so coolant is piped directly to processors to maintain performance.
- **Demand is pulling capacity forward.** Jabil projects that global expansion of nearly 100 GW through 2030, effectively doubling today's footprint, and notes hyperscalers and major enterprises are pre-leasing up to 70% of new developments years before construction is completed.

The consequence is straightforward: once racks move into triple-digit kilowatts, thermal and power constraints stop being background conditions and become program-critical dependencies, reshaping what "ready" looks like for design, procurement, and deployment.

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

2.1 Sizing the Energy Footprint

Rack power density explains why requirements are changing, but data centers are planned and powered at the facility level in megawatts (MW), where 1 MW = 1,000 kW. Using the average U.S. household electricity consumption of roughly 10,500 kWh per year, a facility operating at 10 MW, 25 MW, or 100 MW of continuous IT load translates into household-scale equivalents that make the magnitude of demand easier to grasp.

From Megawatts to Neighborhoods

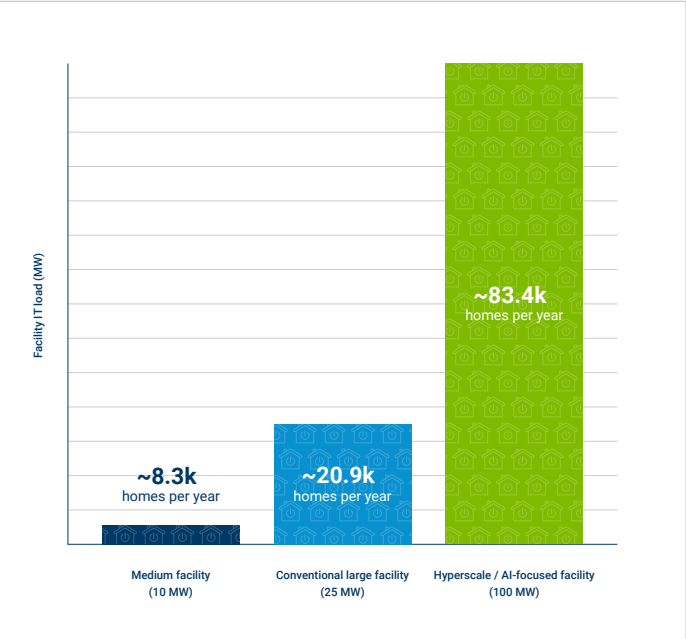

These household equivalents reflect continuous IT load only, the electricity consumed by compute and related IT equipment. Total facility electricity is higher once cooling and power distribution overhead are included (often captured through PUE), which is why a 100 MW hyperscale, AI-focused facility is sometimes described as consuming as much electricity annually as roughly 100,000 households on a facility-wide basis.

At 100 MW, the IT-load translation shown here is about 83,429 homes/year, roughly the number of households in a mid-sized U.S. city such as Tallahassee, Florida (about 85,311 households).

Current build pipelines are already pushing beyond megawatts into gigawatts, where 1 GW = 1,000 MW. At the outer edge, capacities of around 2,000 MW (2 GW) are under construction, and 5,000 MW (5 GW) are planned. This translates to roughly 1.7 million and 4.2 million homes' annual electricity use, about 16% more households than the City of Los Angeles and about 25% more than New York City, respectively.

Footnotes

1. Household baseline: Average U.S. household uses ~10,500 kWh/year (EIA).
2. Conversion: 1 MW = 1,000 kW. Homes/year assumes continuous load (24/7).
3. Homes equivalent formula: Homes/year = (Facility MW × 1,000 × 8,760) ÷ 10,500.
4. Rule of thumb: 1 MW of continuous IT load (24/7) ≈ 834 U.S. homes' annual electricity use, assuming ~10,500 kWh/home/year.
5. Scope note: This figure reflects IT load. Total facility electricity will be higher once cooling and power distribution overhead are included (PUE).
6. Reference points: IEA notes use 25 MW (conventional) and 100 MW (hyperscale); Uptime defines 5–19.9 MW as medium and 20 MW+ as large.

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

2.2 When Power Becomes a Physical Constraint

At this scale, energy demand stops being a planning assumption and becomes a set of physical requirements that determine what can be built, where it can be sited, and how fast it can come online. As loads move from Megawatts toward Gigawatts, the energy footprint turns into concrete constraints across the grid-to-chip stack, setting the ceiling for capacity, timelines, and execution.

Practical implications:

- **Facility power is the real unit of scale.** Rack density drives the shift, but delivery and siting decisions are made in MW (and increasingly GW), which sets the ceiling for expansion.
- **Bigger loads mean bigger hard infrastructure.** Higher IT loads force larger, more complex power delivery and thermal rejection systems, turning “power” into a tangible bill of materials, not an abstract requirement.
- **Energy footprint becomes a schedule driver.** At scale, the limiting factor is often the readiness of physical infrastructure, equipment availability, integration, and commissioning, not just space or demand.
- **Density changes the execution model.** Moving toward high-density compute tightens tolerances and raises the bar for repeatable deployment, spares planning, and operational readiness across a multi-vendor environment.

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

3.0 From Air to Liquid: The Cooling Transition

Thermal management has hit a physical threshold. As rack power loads rise, traditional air-based cooling becomes insufficient at scale, and liquid cooling moves from specialized use cases into standard operating reality for AI-ready environments.

Two liquid-cooling paths are emerging as the practical defaults:

3.1 Direct-to-Chip (DTC) Liquid Cooling

Direct-to-chip (DTC) liquid cooling is an advanced approach that uses cold plates mounted directly on high-heat components such as CPUs and GPUs, circulating coolant to remove heat at the source. DTC cooling has become an industry standard in 2026 for high-density AI clusters where rack power loads often exceed 100 kW.

For operators, this shift raises the readiness bar, requiring CDU and plumbing integration, qualified installation methods, leak-tight tolerances, and a service model built around fluids, fittings, and validated maintenance.

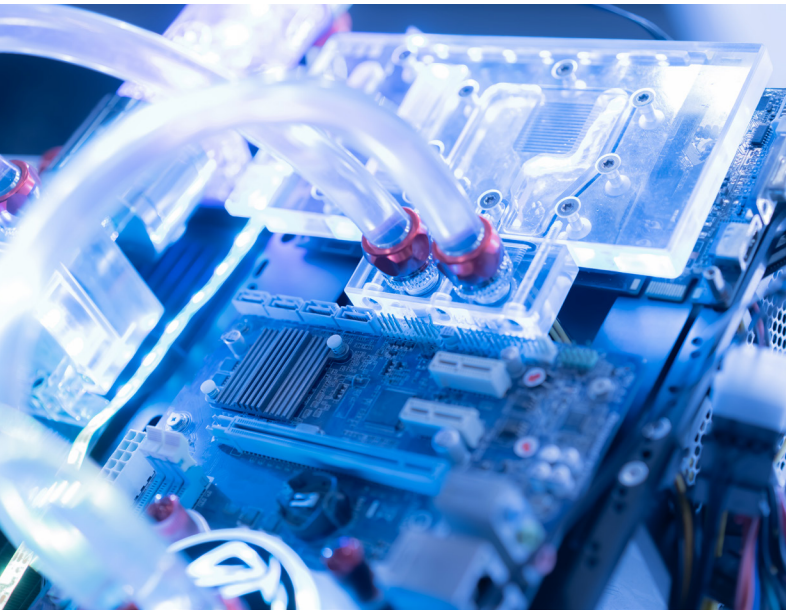

3.2 Immersion Cooling

Immersion cooling represents a different operational trade space. It submerges server components in a non-conductive dielectric fluid that absorbs heat and circulates it to heat exchangers, reducing internal fans and auxiliary power.

Currently, immersion is favored for extreme-density applications and "lights-out" modular data centers, and it **can reduce water consumption by up to 95%** compared to traditional evaporative cooling while improving Power Usage Effectiveness (PUE).

3.3 Cooling Becomes an Operating Mode

As rack power density rises, the shift from air to liquid changes what it takes to deploy at scale. Installation methods, commissioning rigor, serviceability, spares strategy, and supplier readiness move onto the critical path. In that sense, cooling becomes a prerequisite for predictable delivery and stable performance, not an optimization after the fact.

Looking further out, the most "futuristic" concepts share the same logic, maximize cooling by turning the environment into the heat sink.

- **Undersea designs** aim to reject heat directly into cold seawater through passive exchange, reducing reliance on energy-intensive chillers and airflow management.
- **Space-based concepts** apply a similar principle by dissipating heat via radiative cooling into the vacuum of space, paired with continuous solar energy, bypassing terrestrial grid and land constraints.

While these ideas are early, they reinforce the reality that as rack power density climbs, cooling capability becomes one of the defining limits on where and how compute can scale.

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

4.0 The Power Constraint

As AI pushes data centers into higher densities and larger footprints, power stops being a background utility and becomes the gating constraint. These data center power constraints shape site selection, build timelines, and the physical architecture required to deliver stable electricity to volatile compute loads.

4.1 The Grid Bottleneck and the Interconnect Clock

Grid congestion and capacity ceilings have become the defining bottleneck in data center expansion. In major hubs, interconnect timelines can stretch to 3–5 years, long enough to stall projects, force costly workarounds, or push developers toward new geographies. In supply terms, the constraint is increasingly no longer land or capital, it's access to utility-scale power, particularly in hubs including Northern Virginia, Dublin, and Singapore.

When interconnect becomes the pacing item, power shifts from a planning input to a set of physical requirements that must be sequenced and de-risked early. Equipment readiness, integration, and commissioning become schedule-critical alongside the grid connection itself. The result is a new "power readiness" imperative, treating grid access and the grid-to-rack infrastructure stack as front-end engineering decisions that determine time-to-capacity.

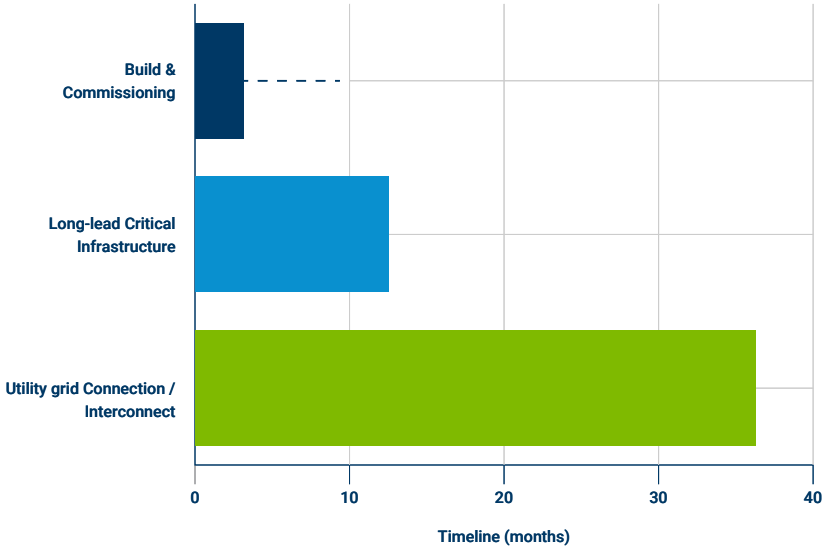

Lane	What this Covers	Analyst Insight	Timeline
Utility grid Connection / Interconnect	Utility connection / grid interconnect approvals and capacity availability	**Interconnect is the bottleneck.** In major hubs, new power connections can take 3–5 years, long enough to stall projects and drive site strategy.	**36–60 months / 3-5 years**
Long-lead Critical Infrastructure	Facility equipment; High-voltage transformers, backup generators, specialized liquid cooling manifolds	**Procurement sits on the critical path.** Delivery windows of 12–24 months for critical infrastructure keep equipment readiness tightly coupled to schedule.	**12–24 months**
Build & Commissioning	Build/commissioning activities (varies by site, delivery model, and approach)	**Construction is compressing, but variability remains.** Modular/prefabricated deployment is accelerating, enabling capacity in months rather than years; keep this lane labeled "months-scale" to avoid implying a fixed duration.	**Variable duration**

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

4.2 Grid-to-Rack Requirements and the Modern Power Toolkit

At AI scale, grid-to-rack power delivery becomes a build requirement with long lead times and commissioning risk. Delivery windows of 12–24 months for critical infrastructure, such as high-voltage transformers, backup generators, and specialized liquid-cooling manifolds, push procurement sequencing and integration readiness onto the critical path.

As a result, "power strategy" increasingly means supplementing the grid, not just connecting to it. Behind-the-meter options such as microgrids, hydrogen fuel cells, and SMRs are being pursued to bypass grid limits and protect time-to-capacity as developers shift toward regions with more available power headroom.

Within that physical constraint set, three architecture levers show up repeatedly as "toolkit" choices:

- **Battery Energy Storage Systems (BESS):** positioned as a shift away from diesel-only backup, with **millisecond-level response** to bridge fluctuations and **peak shaving** during high-tariff periods to lower operating cost.
- **Microgrids and SMRs:** framed as routes to power independence and carbon-neutral baseload; microgrids can enable facilities to **island** from the utility during emergencies or peak demand.
- **Hybrid AC-DC architectures:** implementing a native DC backbone (e.g., **800V DC**) can reduce conversion steps and deliver **up to 15% greater efficiency**, reducing wasted heat and lowering cooling burden.

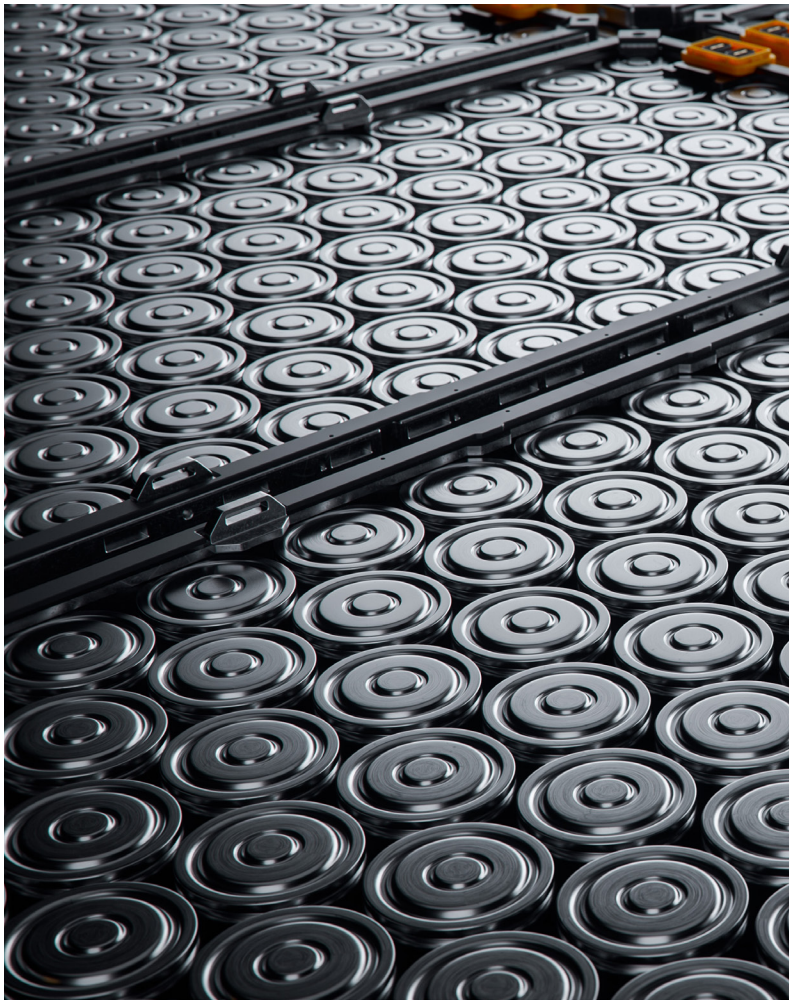

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

5.0 The AI-Ready Bill of Materials: What Becomes Critical

"AI-ready" turns the data center from a standalone IT build into an integrated facility-to-IT system. Performance and uptime now depend on how power, cooling, connectivity, and controls work together. As a result, the bill of materials (BOM) becomes a network of interconnected components that must be engineered and delivered as one system.

Compute & Racks

High-density compute drives new physical requirements at the rack and row level.

Cooling Loop

Liquid cooling hardware that moves heat reliably at scale (CDUs, manifolds, fluid, heat exchange).

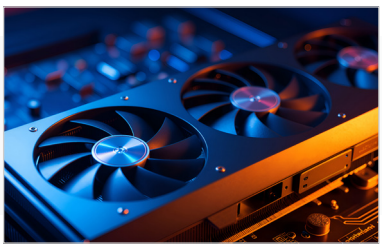

Power Chain (grid-to-rack)

Deliver, condition, and distribute stable power from utility to rack (switchgear > UPS > generators > PDUs).

Networking & Optics

High-throughput fabrics to prevent AI cluster bottlenecks (400–800G, 800G Ethernet).

Controls & Visibility

Monitoring + optimization layer that coordinates power and cooling in real time.

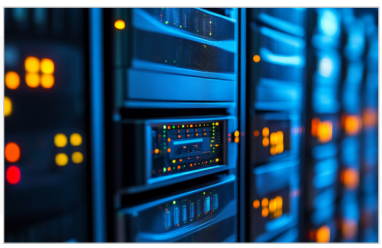

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

5.1 What's Different About the AI-Ready BOM

As AI requirements rise, "supporting infrastructure" is becoming a primary determinant of readiness and cost-to-serve. The power chain, in particular, is the uptime lifeline, UPS covers immediate failures, generators sustain long outages, and PDUs distribute electricity to racks.

At the same time, AI clusters are pushing east-west throughput higher inside the facility. Networking fabrics are now commonly measured in the hundreds of Gbps (e.g., 400–800 Gbps), with hyperscale environments moving toward advanced interconnects such as 800G Ethernet.

5.2 The Integration Challenge

At AI scale, the bill of materials behaves less like a checklist and more like a dependency web. Readiness is determined by how tightly coupled systems perform together, especially when density rises and tolerances tighten, so execution becomes as important as design.

What this means in practice:

- **Substitution becomes harder.** Critical components across thermal, power, and connectivity are less interchangeable once performance and reliability are engineered as a system.
- **Integration becomes the risk surface.** The highest failure and delay risk shifts to interfaces, liquid connections, power paths, and network fabrics, rather than individual parts.
- **Commissioning becomes a repeatable discipline.** Scaling requires standardized test methods and validation routines that reduce site-to-site variability.
- **Procurement becomes sequencing, not purchasing.** Delivery depends on the right parts arriving in the right order, with documentation, alternates, and validated install requirements.
- **Serviceability becomes part of the design.** Spares strategy, maintenance methods, and lifecycle coverage must match density and uptime expectations, or performance degrades over time.
- **The goal is industrialized deployment.** The competitive advantage comes from making high-density capacity deployable and supportable at scale, without turning each build into a bespoke integration project.

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

6.0 Beyond Power: Where Schedules Still Slip

Even with a viable power path, build programs can still lose time in three places that sit outside the interconnect timeline: **people, permits,** and **geopolitics**. These constraints show up as physical execution limits, who can install and operate mission-critical systems, what jurisdictions will approve, and what supply chains can reliably deliver.

6.1 Talent: the High-Voltage Skills Gap

Specialized labor is becoming a hard constraint, particularly high-voltage electricians and mechanical engineers, alongside broader shortages of electrical engineers, mechanical technicians, and facility managers needed to run mission-critical, high-density environments. That pressure is already reshaping workforce investment, with Google funding training for tens of thousands of U.S. electricians and committing $10M in grants to expand the talent pipeline as AI-driven power demand accelerates.

Strategic insight: Treat commissioning and operations labor as a capacity asset, standardize methods and lock critical partners early so labor availability doesn't become the bottleneck.

6.2 Permitting: "License to Operate" Becomes a Schedule Variable

Local policy can function like a cap on build velocity. Temporary moratoriums or strict caps have appeared in jurisdictions including Singapore, Amsterdam, and parts of Ireland, tied to carbon emissions and water usage, while permitting complexity can require design modifications that extend timelines. In practice, these constraints can persist even when demand is surging; Amsterdam's moratorium remains a limiting factor even as developers move forward with large-scale proposals in the region.

Strategic insight: Build "permit-ready" designs from day one, environmental performance and community constraints increasingly act as schedule insurance.

6.3 Geopolitics and Export Controls: Supply Volatility in Critical Components

Geopolitical tensions and export controls are adding sourcing risk to critical inputs such as scarce earth minerals and advanced memory/semiconductor components, contributing to price volatility and potential delays. Recent policy volatility around advanced AI accelerators underscores the risk: even after conditional pathways opened for certain high-end chips, export enforcement officials indicated that no Nvidia H200 chips had yet been delivered into China as of late February 2026, amid guardrails and heightened scrutiny. In parallel, diversion concerns remain live as U.S. officials have pointed to cases where advanced Nvidia chips were allegedly used in China despite restrictions, while China has also moved to limit foreign AI chips in some state-funded data center environments.

Strategic insight: Build "permit-ready" designs from day one, environmental performance and community constraints increasingly act as schedule insurance.

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

7.0 Readiness Framework: Qualify > Source > Stage > Install > Maintain

High-density AI programs move fastest when they're treated as a lifecycle system, not a one-time build. Modern facilities increasingly function as an integrated stack of IT infrastructure (servers, storage, networking) and facility infrastructure (power, cooling, security) that must be engineered and executed together.

The framework opposite translates that reality into a practical sequence, moving from upfront definition and supplier readiness to repeatable deployment and long-term operational stability.

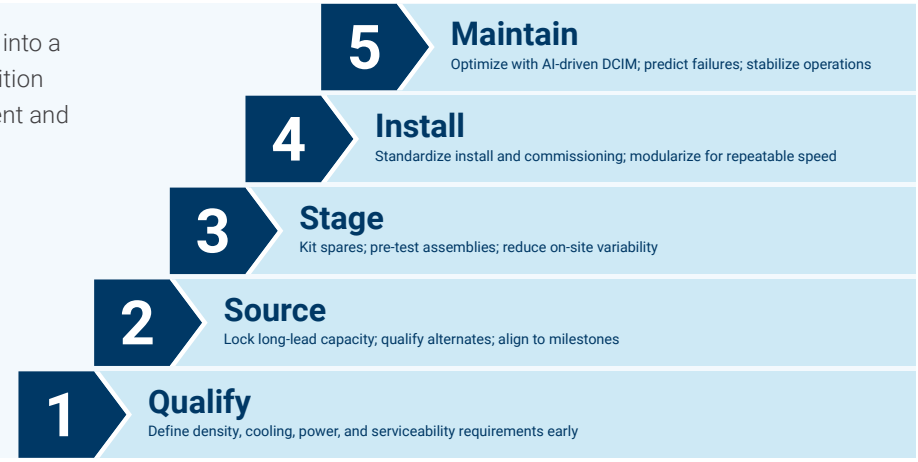

7.1 Qualify

Define "AI-ready" in physical terms, target rack density, cooling architecture (DTC vs. immersion), and a power approach that matches grid realities. Direct-to-chip liquid cooling is positioned as an industry standard in 2026 for clusters where rack loads often exceed **100 kW**, which raises the bar for mechanical tolerances and validated installation methods. When qualification is done well, interfaces and acceptance criteria are explicit early, so schedule risk gets designed out instead of surfacing during commissioning.

7.2 Source

Secure capacity across a multi-vendor stack with the expectation that substitutions narrow once interfaces are locked. Long-lead items, high-voltage transformers, backup generators, and specialized liquid cooling manifolds, are cited with **12–24 month** delivery windows, so sourcing becomes a sequencing discipline as much as a supplier decision. The most resilient programs treat long-lead infrastructure as schedule anchors and align every upstream decision to commissioning milestones.

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

7.3 Stage

Reduce field variability by staging kits, spares, and pre-tested assemblies so installation becomes repeatable rather than improvisational. With critical components facing **12–24 month** delivery windows, staging also prevents the most common "silent delay", crews waiting on missing parts, incomplete documentation, or unvalidated assemblies. Strong staging turns resilience into something tangible: readiness to swap, repair, and recover without downtime.

7.4 Install & commission

Standardize installation and commissioning so each deployment is an iteration, not a reinvention. Modular and prefabricated approaches are positioned as accelerating deployment, enabling capacity in **months rather than years**, but the speed only holds when commissioning procedures and acceptance tests are consistent site to site. In practice, repeatable commissioning is the mechanism that turns design intent into predictable performance.

7.5 Maintain & optimize

Operate the facility as a closed-loop system where performance depends on continuous optimization, not static setpoints. AI-driven DCIM is described as autonomously optimizing power and cooling distribution based on workloads and predicting equipment failures weeks in advance, improving stability and efficiency in complex hyperscale environments. At higher densities, long-term readiness is proven in operations, telemetry, automation, and maintenance strategy become part of the physical design.

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.

8.0 Partnering with Jabil PSCS

AI-ready scale is being defined less by demand and more by physical readiness. As rack densities move into ultra-high-density territory and liquid cooling becomes standard for clusters that exceed 100kW per rack, the constraints shift from "design choices" to "execution realities," power availability, thermal rejection, network throughput, and repeatable commissioning all become inseparable.

8.1 How Jabil PSCS Helps Data Center Operators Scale

Jabil supports AI-era data center expansion through a **Grid-to-Chip portfolio** designed to bridge high-performance IT infrastructure and the facility systems required to power it, backed by engineering-led manufacturing and lifecycle execution.

Here's how Jabil is helping data center operators scale:

- **Lifecycle execution (Design-to-Dust™):** from early-stage prototyping through secure decommissioning.
- **Domestic capacity for hyperscaler scale:** investment in an advanced manufacturing facility in **Salisbury, North Carolina** to support cloud and AI data center infrastructure.
- **Facility-to-IT integration and turnkey power:** strengthened through the **January 2026 acquisition of Hanley Energy Group**, positioning Jabil to design, manufacture, and commission power solutions from the utility grid to the server rack.
- **Thermal at the next frontier:** advanced liquid cooling (DTC cold plates and CDUs) through **Mikros Technologies**, with development aimed at next-gen clusters approaching a **1 MW power density threshold**.
- **High-speed connectivity and energy systems:** leadership in **800G and 1.6T** optical transceivers/silicon photonics, plus BESS solutions for energy storage and microgrids; and sustainable lifecycle services including recovery via **Retronix**.

8.2 Go Deeper with the Data Centers Market Sector Report

For organizations scaling AI-ready data center infrastructure, repeatable performance requires more than point solutions, it requires an operating partner with proven scale across the grid-to-chip stack. Backed by 60+ years of supply chain expertise and 3,000+ supply chain experts, Jabil helps leaders move from insight to execution across power, thermal, high-speed connectivity, and lifecycle operations.

Jabil brings practitioner depth and global operating scale to help organizations execute with more resilience, control, and speed, without relying on one-off workarounds as requirements evolve. That scale spans 100+ locations across 25+ countries, enabling consistent execution and support across global availability zones.

With experience supporting 400+ leading brands and an ecosystem built on 38,000+ supplier relationships and $25B+ in annual procurement spend, Jabil is built to deliver outcomes that hold up under volatility.

Request the full Data Centers Market Sector Report to benchmark your grid-to-chip readiness.

This document is provided for informational purposes only. It is not intended to be used or construed as investment advice. The information contained in this document is not an offer to buy or sell any securities or other financial instruments. Any investment decisions made based on the information contained in this document are made at the sole discretion of the reader. The author(s) of this document are not investment advisors and are not registered with any regulatory authority. The author(s) of this document disclaims any liability for any direct or consequential loss arising from any use of this document or its contents.

©Jabil 2026. All Rights Reserved. Confidential and Proprietary.