Hardware & Networking

Key Takeaways from Industry Updates on Data Center Demand

IT Hardware/ Telecom & Networking Equipment

Joseph Cardoso [AC]
(1-212) 622-9036
joseph.cardoso@jpmchase.com

Manmohanpreet Singh
(1-212) 622-4527
manmohanpreet.singh@jpmchase.com

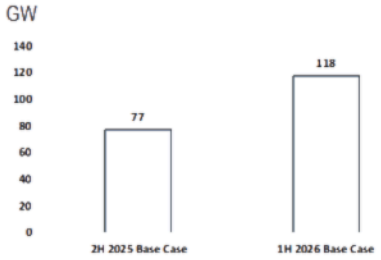
Marc Vitenzon
(1-212) 622-3342
marc.vitenzon@jpmchase.com

Akanksh Chauhan
(1-212) 622-0045
akanksh.chauhan@jpmorgan.com
J.P. Morgan Securities LLC

Over the prior week, several third-party analysts, *including BloombergNEF, Gartner, and Datacenter Hawk*, hosted webinars to discuss latest views on the data center demand outlook and project pipelines. *Key takeaways from the webinars are*:

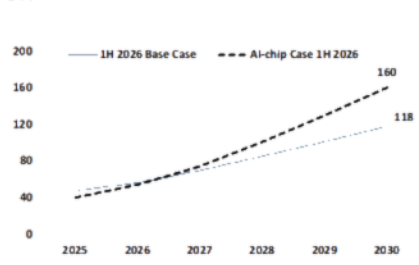
- **Demand for data center capacity continues to expand at a robust pace.** AI has reset the capex baseline, driving a clear shift in both the size and number of individual data center campuses needed to house advanced AI infrastructure and support continued innovation (*individual campuses now exceed 10 GW vs. traditional 1 GW*). Bloomberg's base case, built on current project timelines and buildout, implies U.S. data center power demand reaching 118 GW by the end of the decade, up from a prior (2H25) forecast of 77 GW. That upward revision is driven by a rapidly growing project pipeline but, according to the Bloomberg analysts, remains conservative. A more aggressive "AI-chip" case, meanwhile, based on AI server forecasts from semiconductor manufacturers and OEMs, points to demand of 160 GW by the end of the decade.

Figure 1: U.S. Data Center Demand Outlook 2030



Source: BloombergNEF.

Figure 2: Base Case vs. AI-Chip Case Outlook



Source: BloombergNEF.

- **Power availability is the primary constraint, driving push toward off-grid gas and alternative markets.** Under Bloomberg's base case of 118 GW of US data center demand by 2030, roughly 10 GW of incremental capacity is expected to be added each year, with ~80% carried on the grid and the remainder off-grid. Under the more aggressive "AI-chip" case, however, the grid would need to absorb 19 GW annually, a difficult target that would leave a much larger role for onsite gas generation. Power constraints are therefore beginning to dictate location preferences: off-grid gas campuses are scaling rapidly in West Texas, for example, while cheaper energy is increasing the appeal of international markets such as Northern Mexico.
- **Multi-component supply bottlenecks are driving high lead times, with real construction lagging behind proposed new interconnect.** For IT-related components, lead times vary from 3–12 months for AI Accelerators, 6–12 months for Optics, 24–36 months for Advanced Packaging, 30–36 months for HBM, and 36–48 months for DRAM/NAND. Lead times for some of the other products and services required to bring data centers online, however, run considerably longer, including 36–60 months for Heavy Electrical Equipment (*such as transformers, turbines, and switchgear*), 48–60 months for Skilled Trades (*electricians, refrigeration engineers, plumbers, and pipefitters*), and

See page 3 for analyst certification and important disclosures.

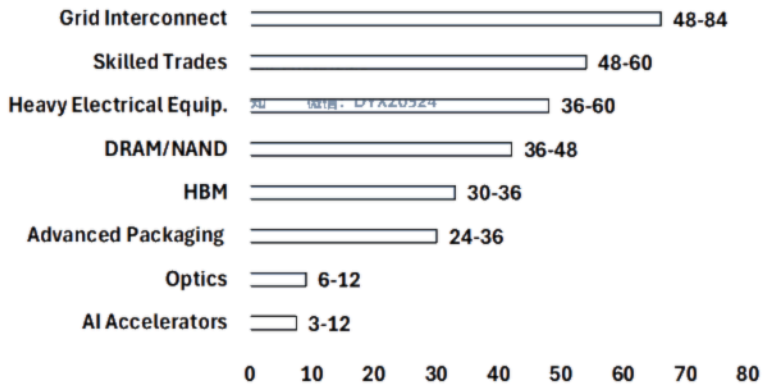


Joseph Cardoso [AC]
(1-212) 622-9036
joseph.cardoso@jpmchase.com

J.P.Morgan

48–84 months for Grid Interconnects.

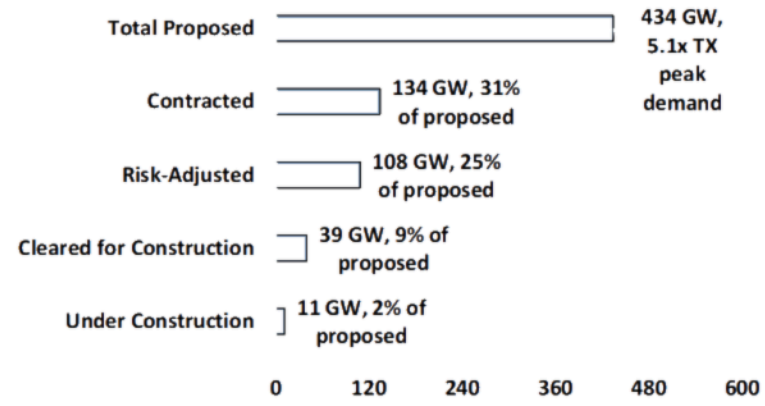
Figure 3: Lead Times in Months by Component



Note: The height of the bar represents the average lead times.
Source: Gartner.

- **Proposed power interconnect pipelines far exceed contracted and under-construction.** Gartner highlighted that total proposed power capacity in Texas, per ERCOT, is tracking closer to 434 GW, spanning all interconnection requests rather than data centers alone, or 5.1x the state's peak demand. Of that total, 134 GW (*31%*) is contracted, 39 GW (*9%*) is cleared for construction, and only 11 GW (*2%*) is actually under construction.

Figure 4: Power Interconnect Pipeline



Source: Gartner.

- **Data center builds run into rising local and state pushback.** New data center builds are increasingly facing pushback from local communities and state governments alike. This is evident in cancellations through June 2026 already surpassing the 2025 total, albeit against a significantly larger pipeline, with roughly $130 bn in projects blocked or delayed in 1Q26 alone. At the same time, as the pipeline expands rapidly, a rising share of projects is being handled by inexperienced developers, adding execution risk. Finally, were hyperscalers to shorten their assumptions for the useful life of infrastructure, demand for new data center buildings would decline, since new IT equipment can be installed in place of old within the same physical sites.