

Data Center Procurement Strategies in a Low-Vacancy Market

Unprecedented levels of colocation leasing have significantly decreased available space & power

By Kirk Killian

August 2024

© *Mission Critical Magazine*

KEYWORDS [colocation](#) / [absorption](#) / [vacancy](#) / [procurement](#)



Popular economist and Freakonomics author Steven Levitt wrote “If the price of a good goes up, the companies that make it figure out how to make more of it. Markets deal with supply and demand.”

The data center industry has evolved dramatically since 2021, with explosive growth in construction and occupancy of data centers. Recent shortages of available data center space are forcing enterprise users to rethink their planning and procurement strategies.

Construction boom is underway

The construction crane is the unofficial mascot for data center sites worldwide. More than 5 gigawatts of data center capacity are now under construction in the United States, a record pace and a sixfold increase from just two years ago. More than a quarter of this new capacity is in Northern Virginia, the world’s largest data center market. A similar construction surge is occurring worldwide, with more than 6 gigawatts announced for Europe and 13 gigawatts announced for Asia, although not all such projects will go forward as planned.

Demand outstrips supply

Absorption – the leasing of data center capacity – is also running at an all-time high, with more than 4.6 gigawatts leased in in the U.S. during 2023, a fivefold increase from 2021. Hyperscale users leased the vast majority of new data centers delivered in 2023, often committing to single chunks of 20 to 80 megawatts per contract – far larger than typically sought by other Fortune 500 enterprise customers. About 15 hyperscale customers represent more than 85% of recent worldwide data center leasing, including public cloud providers (Amazon Web Services, Google, Microsoft, & Alibaba), social media platforms (Facebook/Instagram-owner Meta, TikTok-owner ByteDance, Tencent, LinkedIn, and X/Twitter), consumer-facing companies (Apple, Uber, and Tesla), and large software platforms like Oracle.

A related phenomenon in data center absorption is the pre-leasing of data centers while still under construction; more than 60% of U.S. data centers to be completed through 2025 have been pre-leased. To illustrate, a hyperscale user might agree to occupy three adjacent 24-megawatt buildings to be constructed and delivered at the rate of one per year in 2024, 2025, and 2026. As a result, most data centers currently under construction are no longer available for occupancy by new enterprise customers.

Artificial intelligence providers and numerous enterprises across many industries which are deploying AI applications are also driving data center demand. AI's effectiveness and efficiency improve as the models incorporate ever-larger data sets into their programming, and these models require huge quantities of data center capacity. AI applications also work best in high-density data center deployments to take advantage of extremely low-latency networking between large clusters of power-hungry processors. As a result, those users prefer new data centers with the densest heat rejection capabilities.

Data center demand has also increased because some enterprise users are moving applications and data back from public cloud into colocation facilities, a process that has been termed “re-patriation” from cloud. These users have discovered that some of their applications don't perform well in public cloud, cloud costs are too high or have unexpected (and unbudgeted) price spikes, or audit and compliance opaqueness in the cloud doesn't meet regulatory requirements.

No room at the inn

Unprecedented levels of colocation leasing have significantly decreased the market vacancy rate – a measure of availability – of data centers everywhere. At year-end 2021, the overall data center vacancy rate in the United States was 7%; it now hovers around 2%. The vacancy rate is even lower in some markets like Northern Virginia with 1% vacancy. Vacancy rates in other large markets have dropped as well. Two years ago, an enterprise user seeking one megawatt of colocation space in Dallas had more than ten desirable options when the vacancy rate was 12%, but now only a few options exist with the local market vacancy rate at 3%. Leading international markets are tight, too. In Frankfurt – one of Europe's largest data center markets – the vacancy rate dropped from 12% three years ago to 6% today.



Construction of new data centers is booming worldwide as hyperscale and AI users fuel new facility demand.

Longer construction timelines have delayed the delivery and availability of new data centers. Obtaining property rezoning, construction permits, and critical systems equipment such as generators, transformers, and UPS systems can now take 18-36 months, compared to lead-times of just 6-12 months a few years ago.

Some communities which previously recruited and embraced new data center campuses and the abundant property taxes they generate are now tapping the brakes on approvals to reflect their constituents' objections to "huge windowless buildings" in their backyards.

Delays for power delivery also have pushed commissioning dates for some new data center projects several years into the future. Shortages of utility step-down transformers, transmission-voltage bottlenecks impeding delivery of power from generation sources to users, and challenges to new electricity generation from non-renewable sources have delayed many data center projects. The "NIMBYS" not-in-my-back-yard-syndrome impedes approvals of both data center campuses and utility infrastructure, including new power plants and transmission towers.



With limited availability of data center suites worldwide, users are scrambling to lease colocation capacity to meet exploding demand.

Prices increase everywhere

Historically low data center vacancy rates and extensive pre-leasing of projects now under construction have driven data center rental rates higher, mirroring inflation spikes across the economy. Data center rents in most major markets worldwide have risen 30-50% during the last three years.

Updated data center planning & procurement strategies

Despite challenging market supply conditions, enterprise users can embrace several strategies to bring flexibility, efficiency, and economy to their data center planning and procurement processes.

Start planning earlier

Since options are scarce and most new availabilities will not come online for several years, users should start the planning process for their next data center deployment 18-36 months before the target occupancy date. Advanced planning isn't an absolute pre-requisite, as we've assisted enterprises in numerous "last-minute" procurements, but those usually involve scrambling to match the best available facilities with specific project requirements.

Plan future data center needs in power ranges

Many enterprises don't accurately project their data center needs more than 3 years into the future, which is understandable given uncertainty around AI adoption, migrations to (and back from) public cloud, and organic application/data growth. However, wise enterprises should endeavor to forecast future needs in more detail to enable better procurement planning.

Prudent users should prepare a matrix of anticipated space and power needs for each of the next 5 to 8 years in three different ranges of low-, most likely-, and high-quantity estimates to guide procurement parameters. It's helpful to combine baseline colocation contract quantities with flexibility measures, including expansion rights, contraction rights, early termination options, and “spend shift” between occupancy costs and managed services. Users should also combine the initial contract terms with fixed-rate renewal options to control facility costs for years to come.

Enterprises can use detailed power usage projections to work cooperatively with data center providers to add more power and cooling in multiple contract phases when additional capacity is needed. For example, a user that needs 1,000 kW this year for a new data center deployment plus an additional 500 kW two years later can structure a power “ramp” in a facility that is tight on space and power now but expects to get additional capacity back from other customers or expand their facilities in the future. Accurate scoping decreases the tendency to over-commit for future growth of power usage, the largest source of contract overspending.



Electric utility service delays and critical systems equipment shortages are deferring new data center completions, decreasing available capacity and increasing rental rates.

“We put customers at the center of everything we do.” said Spencer Mullee, CEO of Centersquare, the newly rebranded combination of leading data center providers Evoque and Cxtera. “While current occupancy levels are at an all-time high throughout the industry, we offer workload appropriation exercises for our customers to ensure their applications are deployed in the most efficient way possible. We listen closely to our customers’ challenges and then build infrastructure around their business so it can scale and flex alongside them as not only the IT landscape changes, but also as their needs evolve. For most of our customers, deploying a hybrid mix of colocation and our on-demand interconnection and services helps them achieve their goals in the most efficient way possible.”

Densify

Many enterprises now deploy hardware more densely than ever, combining high-flow floor tiles and hot- or cold-aisle containment in a new data center suite. Increasing the initial deployment density enables easier future capacity expansions if more power and cooling become available. Some users planning low-density deployments of 5-10 kW cabinets today are eyeing the future to make sure their selected facility can cool entire pods averaging 15-25 kW per cabinet, plus support a liquid-to-the-cabinet solution if desired later. This can also reduce router costs and improve networking performance when an enterprise’s processing and storage hardware is placed closely together.

Consider “pre-loved” second generation data centers

Since most new data centers have been pre-leased by hyperscale users, many enterprise users are evaluating available suites recently vacated by previous occupants in facilities built 5-15 years ago. Most well-maintained data centers built within the last 15 years offer high reliability, and the major data center providers have good track records for critical systems maintenance.

However, enterprises should perform more extensive due diligence before choosing older facilities. Users should request a facility condition report discussing the projected remaining useful life of the major critical

systems components, maintenance program details, and a parts inventory plan for non-anticipated (e.g. emergency) replacements over the contract life. Many older colocation facilities also feature slightly higher Power Usage Effectiveness (PUE) ratios than new data centers, but providers will often reduce rental rates in older data centers to offset any electricity reimbursement cost differences.

Create location flexibility

Many enterprises have specific geographical requirements for new data center placements to benefit from low fiber latency to legacy IT installations, proximity to corporate IT staff, availability of cloud direct-connect circuits, comparatively lower electricity rates, and tax incentives. However, if data center capacity is not immediately available in the user's preferred locations, nearby secondary markets may meet the user's requirements with faster delivery dates and lower overall occupancy costs. Prudent users should develop meticulous scoping criteria to maximize selection flexibility as they evaluate multiple locations.



Enterprise users should deploy hardware densely to maximize expansion capabilities and utilize modern cooling designs including refrigerant to the cabinet.

Decide faster among fewer options

Because available data center space is more limited today than just two years ago, users are choosing among fewer prospective facilities, and the better options are being leased at a pace not seen since the dot-com era. In prior buying cycles, a creditworthy enterprise user could reasonably expect that most data center suites under consideration would still be available in the coming months. This allowed prospective customers to reserve a suite with a Letter of Intent, negotiate the contract during the subsequent 2-4 months, and then execute the contract after obtaining internal corporate approvals. Now, the best available suites may have multiple prospective customers competing for the space, and the preferred suites may be leased by competing users within 30-60 days after becoming available. Wise users should shorten their procurement cycle to secure the preferred data center options.

Stack tasks for procurement efficiency

The traditional process of requesting lengthy written proposals from data center providers, awaiting responses, and then negotiating against those proposals has evolved by necessity. Instead of conducting a formal, multi-staged Request for Proposal process over 4-8 months, many enterprises are now using a "short-form RFP" to research available facilities and their proposed pricing, followed by the written proposal and contract review phases once the best facility options have been determined. Experienced data center advisors can assist busy corporate IT and procurement teams in streamlining these tasks.

Conducting due diligence and selection in smaller phases typically accelerates the procurement process. Concurrently "stacking" some tasks can shave weeks from a project schedule, instead of performing tasks sequentially without starting the next task until the current task is completed. For example, a user might simultaneously evaluate suite delivery schedules, critical systems infrastructure, and telecom network options to all prospective facilities while also preparing build-out and occupancy cost budgets for the better options, rather than waiting to prepare budgets when all other information is at hand.

Seek multi-stage corporate approvals

Large corporate procurement departments are often slower than their smaller and nimbler counterparts. Astute enterprises are now obtaining preliminary project approvals from senior IT management prior to completing final approvals for signing new data center contracts.

In this two-step process, users summarize project goals, general facility and provider criteria, a preliminary budget, and a key tasks timeline to obtain preliminary internal approvals. The procurement team then ranks and down-selects the specific first- and second-choice data center options to negotiate key commercial and contractual terms, shaving months from the typical approval process.

In this accelerated process, legal counsel can rapidly provide a preliminary review of a data center provider's contract soon after the team has selected a desired suite. It's become crucial to begin contract review within a few business days of selecting the preferred facility option, with diligent pursuit of acceptable contract language soon thereafter.

A little help from my friends

Sophisticated enterprises also benefit from working with experienced data center procurement advisors who bring market knowledge and transactional experience to help users identify and select the best facilities, negotiate reduced costs, improve contract flexibility, and streamline the selection and procurement process. Advisors also assist in the contract review process, offering revisions to contract and business terms based upon hundreds of prior transactions, supplementing the attorneys' review of contract verbiage.



Limited availability of new data center suites has prompted some enterprise users to select 5-15 year old vacant facilities with lower occupancy costs.

KEYWORDS [colocation](#) / [absorption](#) / [vacancy](#) / [procurement](#)



Kirk Killian

Kirk Killian is President of Partners National Mission Critical Facilities and has advised major corporations in procuring data center facilities worldwide since 1999. He serves on the Mission Critical Magazine Technical Advisory Board, and assists clients in project planning, site selection, and contract negotiations. He can be reached with any questions at kkillian@pnmcf.com.