WHITE PAPER

From Build-To-Suit to Build-To-Performance Spec

Data Center Development in an AI-Driven Market

STREAM DATA CENTERS

The explosion of AI into the mainstream has driven massive changes in technology infrastructure and the data centers that support it. The industry is dealing with the massive demand for new AI data center capacity that has eclipsed the then-unprecedented post-COVID demand increase. Notwithstanding the evolving AI-driven requirements, sustainability remains a key requirement to serve our client base. In this white paper, learn how Stream is proactively evolving—from building to suit our customers' requirements to building to meet their performance specifications.

WHAT'S INSIDE

Introduction

Demand for all things digital drives innovation in applications, which drives innovation in technology infrastructure, which drives innovation in the data center. The pace of this innovation is exponential, and continues to accelerate. Three trends in particular have driven massive changes in technology infrastructure and the data centers that support it:

- First is the combination of unprecedented demand and constrained supply. While COVID-19 represented a huge step up in demand, accelerating growth has continued and today massive demand is the new normal. At the same time, supply chain delays and construction labor challenges have created new obstacles for the industry to overcome.
- Second, artificial intelligence, which represents a sea change from the technology developments that came before, has massively accelerated technological innovation—in large part because generative AI tools like ChatGPT are now accessible to anyone with an internet connection.
- Finally, there is an increasing focus on sustainability beyond energy efficiency (for example, targeting WUE and greenhouse gas emissions in addition to PUE). The next step in the evolution of our thinking about sustainability is to holistically consider the data center's environmental impact across its lifecycle.

At Stream Data Centers, we've been watching these trends develop and have worked hard to prepare ourselves to respond. We've evolved the data center as our customers have evolved their technology and the infrastructure they deploy. Developers like Stream used to deliver a lot of projects that were 'build to suit'—meaning we used our customers' plans to build them 'their' specified facilities. While we still are happy to deliver products that are built-to-suit, delivery models are changing.

Organizations that require significant capacity in a finite period of time are moving away from 'build me exactly what these plans specify' towards 'build your standard offering as long as it supports our technical performance specifications.' Building to performance spec gives us more flexibility in how we design and build the data center—allowing us, for example, to respond to supply chain constraints, or rising density—while always meeting or exceeding our customers' SLA and optimizing for speed and cost.

In this whitepaper, we detail the three most significant technology trends we face today, how those trends affect the data center, and how Stream is responding.

At Stream Data Centers, we've been watching these trends develop and have worked hard to prepare ourselves to respond. We've evolved the data center as our customers have evolved their technology and the infrastructure they deploy.

Trends	How the trends affect the data center	How Stream responds
Unprecedented demand and constrained supply	• Turnkey capacity is used up • Data center customers are turning to developers years earlier than they did pre-2020	• Proactive land development • Flexible deal structures • Configurable sustainable data center design
AI is now mainstream	• Power requirements are increasingly massive and density is rising • Closed-loop liquid cooling to the rack will be essential	• Configurable and innovative proprietary cooling design
Sustainability focus beyond PUE	• Municipalities are restricting data center water use • Customers are committing to reduce greenhouse gas emissions	• A shift toward WUE≈0 cooling systems • Carbon-focused data center design

Trends

Unprecedented Demand + Constrained Supply

Way back in 2022, Stream COO Michael Lahoud wrote about how COVID-19 created a "perfect storm" for the digital infrastructure industry: "Data volumes skyrocketed to unprecedented levels as pandemic lockdowns drove massive increases in web conferencing, video streaming, and online gaming."[1] While COVID-19 represented a huge step up in demand, accelerating growth has continued and today massive demand is the new normal.

New applications leveraging machine learning and artificial intelligence are further accelerating demand. The global datasphere—a measure of all the data created, captured, replicated, and consumed in a given year—is forecasted to nearly triple between 2021 and 2026.[2]

At the same time as demand growth has continued to accelerate, supply chain delays and construction labor challenges have created new obstacles for the industry to overcome.

Supply chains for infrastructure like generators, chillers, UPSs, and switchgear remain incredibly tight. The lead time for generators, as just one example, was considered 'bad' when it was 20-30 weeks; now it's 72-104+ weeks. A recent survey of data center operators by AFCOM revealed that 94% of respondents have had supply chain issues.[3] And, according to a survey by Uptime Institute, data center supply chain disruptions are getting worse.[4]

Where in the past warehouses might have been full of MEP infrastructure, now most developers have exhausted their warehoused infrastructure and are proactively ordering MEP equipment two years ahead. Power utilities pose an even bigger time-to-market issue now. Some are quoting *five years* to deliver additional capacity and infrastructure such as transformers and breakers.

[1] Michael Lahoud, **Navigating Today's "Perfect Storm" of Unprecedented Data Center Demand and Constrained Supply**, 31 May 2022.

[2] IDC Worldwide Global DataSphere Data Forecast, 2022-2026

[3] Network World, **Supply-chain constraints spike data-center outages**, 20 June 2023.

[4] Data Center Knowledge, **The State of Data Center Supply Chains**, 16 March 2023.

Artificial Intelligence is Now Mainstream

Widespread adoption of artificial intelligence (AI) has been 'just around the corner' for decades. (Researchers have been working on neural networks since the early 1970s.) Deep Blue beat Garry Kasparov in chess in 1997 but it wasn't until 14 years later that Watson won Jeopardy. Another decade later, and AI is driving cars, writing code, and making medical diagnoses. (And creating, including this whitepaper's cover image.)

AI is part of a long history of digital innovation, but it represents a sea change from the technology developments that came before. It has massively accelerated technological innovation—in large part because generative AI tools like ChatGPT are now accessible to anyone with an internet connection. "For years, researchers have been on a tear, pumping out new AI models faster than they can be commercialized. But in 2022, a glut of new apps and programs have suddenly made these skills available to a general audience, and in 2023, as we continue scaling this new territory, things will start changing—fast."[5]

AI has gone mainstream, and the impact is reverberating everywhere—from the halls of Congress to corporate boardrooms. In a Deloitte study of corporate executives conducted at the end of 2020, nearly 80% of executives surveyed said they had implemented some form of AI-driven automation. Another 16% said they planned to do so within three years.[6]

AI impacts the data center, too, of course. "AI requirements, along with continued adoption of cloud services, are the main drivers of hyperscale expansion, driving record growth in the data center sector."[7] AI is expected to accelerate data center demand even more as cloud providers increasingly offer GPU-as-a-service offerings, AI companies offer new AI services, and enterprises race to adopt the technology.

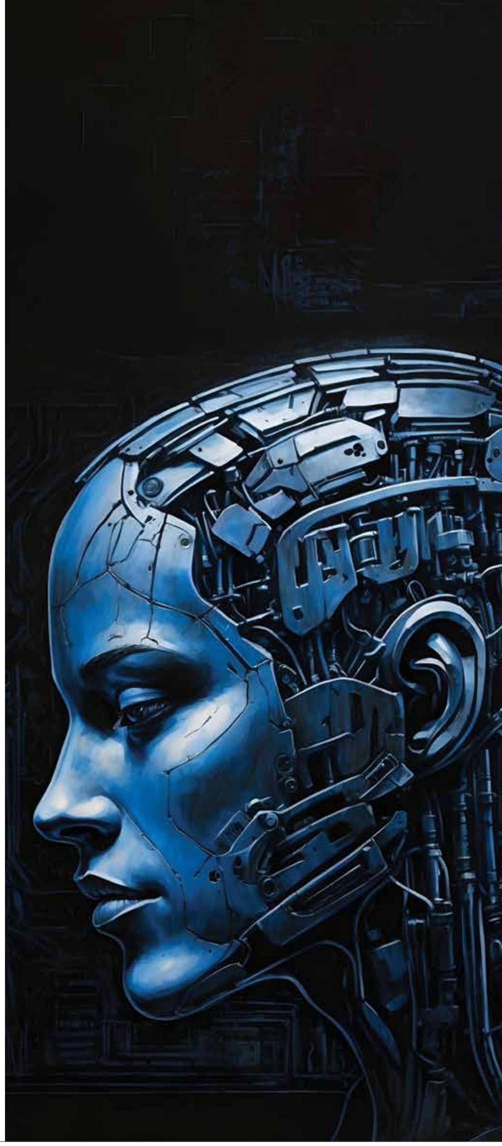

Exhibit 1: Worldwide AI Systems Spend Growing at 27% CAGR

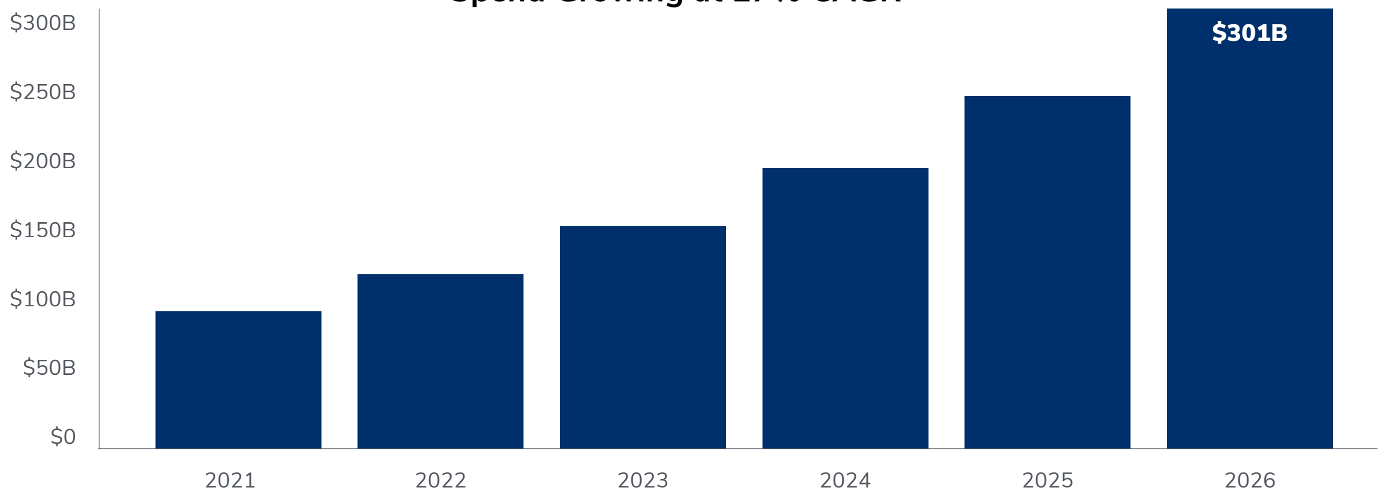

Source: IDC's Worldwide Artificial Intelligence Spending Guide, August 2022

---

[5] The Verge, **ChatGPT proves AI is finally mainstream—and things are only going to get weirder**, 8 December 2022.

[6] Deloitte, **Automation with intelligence**, 25 November 2020.

[7] JLL, **1H 2023 North American Data Center Report**, 24 August 2023.

Sustainability Focus Beyond PUE

For a long time, power usage effectiveness (PUE) was ***the*** data center sustainability metric. As data center users and operators became more comfortable bringing outside air into the facility, the thinking about energy efficiency evolved, and operators began to target—and reach—PUEs well under 1.5. The focus on PUE in data center sustainability efforts has driven significant improvements in energy efficiency. But as it only covers mechanical and electrical efficiency, PUE was never intended to be a broad measure of carbon impact, much less a holistic measure of sustainability.

Focusing on both PUE and WUE is a more holistic approach to data center sustainability.

The Phoenix metro area is only somewhat unique in that it is both a top data center market and a highly stressed watershed. Actually, about 20% of data centers draw water from moderately to highly stressed watersheds, according to researchers at Virginia Tech. Recently, attention has shifted to data centers' water use. Water usage effectiveness (WUE) is a new and important measure which, alongside PUE, does provide a more holistic view of a data center's environmental impact.

Water is an essential resource for the advancement of our digital society. It's a resource we must acknowledge as finite, and we must safeguard it accordingly—designing and operating data centers to minimize water consumption across the supply chain.

WUE represents an important evolution in how we think about the impact of the data center, because there are tradeoffs between PUE and WUE. All else equal, a data center with an evaporative air cooling system uses more water locally but less power than a data center with a closed-loop air cooling system. In most geographies, the best balance between water efficiency and energy efficiency is closed-loop liquid cooling to the rack.

As Google's Senior Vice President of Technical Infrastructure Urs Hölzle wrote, "There is no one-size-fits-all solution."[8] For example, given that renewable power generation typically uses less water than traditional generating methods, a data center using renewable energy for a closed-loop air cooling system could have a smaller total water footprint than a data center in the same location using an evaporative air cooling system.

The next step in the evolution of our thinking about sustainability is to holistically consider the data center's environmental impact *across its lifecycle*.

The next step in the evolution of our thinking about sustainability is to holistically consider the data center's environmental impact *across its lifecycle*.

Data center sustainability now encompasses energy efficiency, water efficiency, direct and indirect GHG emissions, waste, and impact on the local ecosystem. Paul Lin and Robert Bunger at Schneider Electric developed a framework[9] that includes those five categories, with 28 metrics "for data center operators in various stages of their sustainability journeys to take control of sustainability goals":

Energy

- Total energy consumption
- Power usage effectiveness (PUE)
- Total renewable energy consumption
- Renewable energy factor (REF)
- Energy reuse factor (ERF)
- Server utilization (ITEUsv)

[8] Urs Hölzle, **Our commitment to climate-conscious data center cooling**, 21 November 2022.

[9] Paul Lin and Robert Bunger, **Guide to Environmental Sustainability Metrics for Data Centers**, June 2023

GHG emissions

- Scope 1 GHG emissions
- Scope 2 location-based GHG emissions
- Scope 2 market-based GHG emissions
- Scope 3 GHG emissions
- Carbon usage effectiveness (CUE)
- Total carbon offsets
- Hourly renewable supply and consumption matching

Water

- Total site water usage
- Total source energy water usage
- Water usage effectiveness (WUE)
- Water replenishment
- Total water use in supply chain

Waste

- Total waste generated
- E-waste generated
- Battery waste generated
- Total waste diverted
- E-waste diverted
- Battery waste diverted

Local ecosystem

- Total land use
- Land use intensity
- Outdoor noise
- Mean species abundance (MSA)

How the Trends Affect the Data Center

All three trends have driven massive changes in technology infrastructure and the data centers that support it:

- First, the combination of unprecedented demand and constrained supply have used up turnkey data center inventory, and data center customers are turning to developers years earlier than they did pre-2020.
- Second, artificial intelligence, which represents a sea change from the technology developments that came before, is massively raising data center power requirements and density and forcing changes in how the data center is cooled.
- Finally, an increasing focus on sustainability beyond energy efficiency is driving new restrictions on data center water use and greenhouse gas emissions.

Turnkey Capacity is Used Up

Three years of unprecedented demand has eaten up most hyperscalers' existing data center inventory, and most third-party providers' off-the-shelf inventory too. In the first quarter of 2023, the average vacancy rate of the top 10 North American data center markets was 2.88% according to datacenterHawk, down from 8.3% just a year earlier.[10] Globally, twenty-four markets are now under 10% vacancy, according to Cushman & Wakefield, compared to just six markets in 2020.[11]

Both hyperscaler and third-party developers have been building, of course, but demand has been so massive, on top of supply chain constraints, that most new capacity gets absorbed before it's even commissioned.

"We've never been in a situation with such a supply and demand imbalance. A year-and-a-half ago, we had a healthy vacancy rate in most markets around the U.S. That meant we could pretty much find a home for all the requirements that are out there, and that has completely changed."

—CURT HOLCOMB

Executive Vice President and Co-lead of JLL's Data Center Markets Practice

[10] datacenterHawk, **1Q 2023 Data Center Market Recap** and **1Q 2022 Data Center Market Recap**.

[11] Cushman & Wakefield, **2023 Global Data Center Market Comparison** and **2022 Data Center Global Market Comparison**.

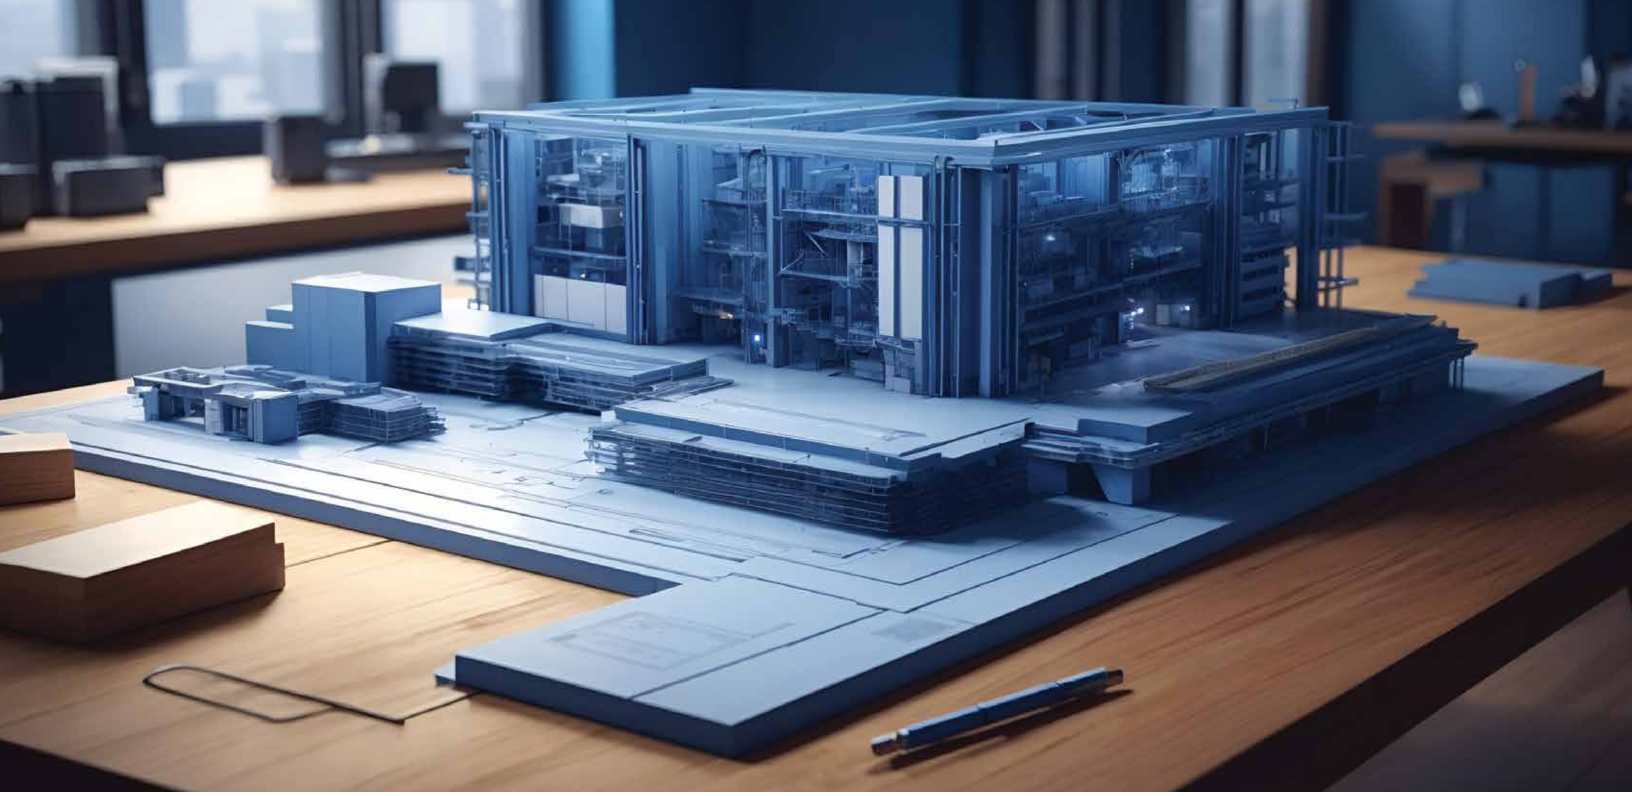

Data Center Customers are Turning to Developers Years Earlier than They Did Pre-2020

Four years ago, the market average for delivering new data centers was about 16 months. Often, a customer could visit a handful of providers in any major market, take their pick, and have new capacity online in months. Times have changed, of course—between the massive demand that is a new normal and supply chain constraints that continue to persist—and today the market average for time from leasing capacity to delivery is more than two years.[12]

As a result, hyperscalers—by far the largest consumers of data center capacity—are relying more heavily on third parties. "Outsourcing is sharply increasing as major hyperscalers are scrambling to find enough data center capacity globally to host their significant incoming workload activity," explained Credit Suisse analyst Sami Badri.

How did we get to a place where development is cool again? Rich Miller put it well: "The hyperscale operators are good at building data centers. But as their operations grow, their company-built campuses must also navigate a difficult construction environment, with shortages of both supplies and labor. That combination is prompting a greater reliance on developers specializing in data center construction."[13]

Power Requirements are Increasingly Massive and Density is Rising

As has always been the case, data center requirements are driven by the capabilities and constraints of the technology infrastructure inside the data center, and the demands of the workloads running on that gear. The increasingly widespread adoption of power-intensive workloads is driven by exponentially rising demand for more data, real-time data consumption, and real-time analytics—particularly as organizations embrace AI for an increasingly wide range of applications.

For example, a ChatGPT-4 query consumes 3-33 times more energy (3.6-36 kJ) than a Google search query (1.08 kJ).[14] To support those types of workloads, chipmakers are producing increasingly powerful chips. But in a shift, those increasingly powerful chips are also increasingly power-hungry. GPUs are around four times more power-dense than CPUs.[15] As infrastructure leaders look to pack more power into the same space, average data center densities are rising dramatically. And to support it all, critical loads are higher than ever.

---

[12] Credit Suisse, 2023 Outlook: The cloud has four walls, 18 January 2023.

[13] Data Center Frontier, **As Supply Chain Concerns Grow, Cloud Players Scoop Up Data Center Space**, 3 May 2022.

[14] **AI and its carbon footprint: How much water does ChatGPT consume?** 22 October 2023.

[15] **The energy challenge of powering AI chips**, 11 May 2023.

Chip performance is increasing—and so is power consumption.

For over a decade, chip performance increased while power consumption remained relatively flat. But since about 2018, rising performance has meant rising power consumption. More computation requires more watts. As just one example, the maximum power consumption of NVIDIA's latest GPU is 160% higher than that of the company's previous generation chips.[16]

AI workloads run on high-performance GPUs that are incredibly power intensive. By 2027 all the AI servers in the world, running at full capacity, would consume at least 85.4 terawatt-hours of electricity annually, according to new research.[17] Indeed, "The GPUs that are used to train generative AI models have large power draws," said Sean Graham, IDC's research director for cloud to edge data center trends.[18] That's raising data center power requirements and density and forcing changes in how the data center is cooled.

Exhibit 2: Today, rising performance means rising power consumption

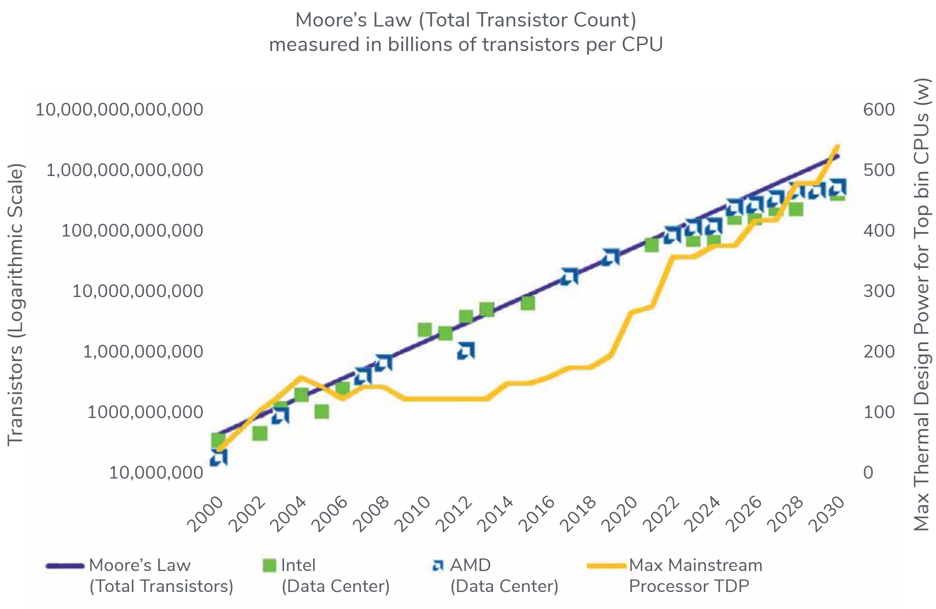

Source: Omdia research commissioned by Stream Data Centers

[16] Data Center Frontier, **Hotter Hardware: Rack Densities Test Data Center Cooling Strategies**, 8 May 2023.
[17] Alex de Vries, **The growing energy footprint of artificial intelligence**, 10 October 2023.
[18] Wall Street Journal, **Rising Data Center Costs Linked to AI Demands**, 13 July 2023.
[19] JLL, **2021 Year-End Data Center Outlook**, March 2022.
[20] IEA, **Data Centres and Data Transmission Networks**

As infrastructure leaders look to pack more power into the same space, average data center densities are rising dramatically.

For the most part, increasingly power-hungry technology infrastructure is being fit into the same data center footprint as previous generations. Average server rack density has been rising steadily. And while average air-cooled server deployments remain approximately 10 kW per rack, some more recent air-cooled deployments are reaching five times that level and liquid-cooled AI deployments can be as high as 50X the average air-cooled density. Because AI applications run at those much higher densities, high-density deployments will be the norm in three to five years with the next generation of technology infrastructure.

Supporting more power-hungry workloads, utility power requirements are increasingly massive.

By the end of 2021 there were more 20+ megawatt data center requirements in the market than 1 MW requirements.[19] Most data centers being built today are well over 30 MWs. Data centers are still increasingly efficient, but exponential growth in the workloads they support means significant growth in the size of data centers and the amount of electricity they consume in aggregate. According to the International Energy Agency, combined electricity use by Amazon, Microsoft, Google, and Meta more than doubled between 2017 and 2021, rising to around 72 TWh in 2021.[20]

Closed-Loop Liquid Cooling to the Rack Will be Essential

Increasingly massive power requirements and rising density mean more heat transfer rate per unit area (heat flux), which pushes the limits of air-based cooling. We are already seeing that heat fluxes for the most powerful processors are too high to manage with air cooling. Air is not nearly as effective a heat transfer medium as liquid, and at some point is unable to efficiently remove all the heat generated by high-power chips.

Supporting the levels of heat flux that will come with 100+ kW rack densities will require direct liquid cooling to the rack with a direct tie to a chilled water loop via cooling distribution units (CDUs). When it comes to heat transfer, water is fundamentally more efficient than air—because it is denser, has a higher specific heat capacity, and a lower thermal resistance. According to Dr. Tim Shedd at Dell, water to the rack should be able to support up to 1500 watts per socket.[21]

Municipalities are Restricting Data Center Water Use

The Uptime Institute cites California, Singapore, Spain, the United Arab Emirates, and Australia as cautionary examples of places where massive demand from data centers is stressing local water supplies—and governments are pushing back. In Spain, for example, "planning authorities are concerned over the use of water-intensive cooling and may restrict the use of systems that consume too much water. New builders may face legal challenges as they try to negotiate guaranteed water supplies."[22]

Likewise, in Phoenix, some local policymakers have denied new data center permits because of water use concerns.[22] As one MIT researcher explained it, "In Mesa, Arizona, where I spent six months researching the emergence of a desert data center hub, some politicians are now openly opposing the construction of data centers, framing the centers' water usage as inessential and irresponsible given resource constraints."[24]

Because data centers' water consumption puts pressure on local water supplies, it's time for designers and operators to consider water usage as seriously as energy usage and take steps to minimize water consumption in resource-constrained locations. Indeed, in Arizona and beyond, many data center operators are looking holistically at water and energy consumption across the entire supply chain, and optimizing resource usage overall based on the unique challenges and opportunities posed by particular locations. That can take the form of using less water to cool the data center, continuing to seek energy efficiency gains, and leveraging less water-intensive energy sources like solar and wind.

Many of the world's largest data center operators have recently made water sustainability commitments. Google has pledged to be "**water-positive**" by 2030—replenishing 120% of the water consumed across its offices and data centers.[25] Microsoft and Meta have also committed to being water positive by 2030.[26]

---

[21] R. Curtis, T. Shedd and E. B. Clark, "**Performance Comparison of Five Data Center Server Thermal Management Technologies**."

[22] Uptime Institute, **The gathering storm: Climate change and data center resiliency**, 16 February 2021.

[23] Chris Bair and Yvonne Deir, **Water Is Life**, 3 March 2022.

[24] MIT Press, **The Staggering Ecological Impacts of Computation and the Cloud**, 14 February 2022.

[25] Google, **Our commitment to water stewardship**, 9 September 2021.

[26] Microsoft, **Microsoft will replenish more water than it consumes by 2030**, 21 September 2020. Meta, **Restoring More Water Than We Consume by 2030**, 19 August 2021.

Customers are Committing to Reduce Greenhouse Gas (GHG) Emissions

Much of the industry's focus is now on identifying GHG emissions and opportunities to reduce those, which encompasses energy efficiency in both mechanical and electrical systems. Scope 1 emissions are those the data center is directly responsible for—emissions produced by backup generators, for example. Scope 2 encompasses a data center's indirect GHG emissions, such as those created in the production of electricity used by the data center. Scope 3 emissions are the result of activities from assets not owned or controlled by the data center, but that the data center indirectly affects in its value chain—for example, the emissions generated in the production and transport of MEP gear.

Businesses of all kinds are paying more attention to these broader sustainability metrics. Most hyperscalers and third-party data center providers have committed to science-based emissions reductions targets—in line with limiting global warming to well-below 2°C above pre-industrial levels.[27] Many have announced plans to achieve net-zero emissions or even be carbon negative by 2030. In some cases, progress toward science-based targets is internally motivated, but in an increasing number of cases, it's being mandated if even implicitly—as in markets where public concerns about data centers' energy and water use have led policymakers to deny new construction permits.

Establishing a baseline for Scope 3 emissions, setting targets, and then achieving those targets is the most significant challenge in the industry today. Success requires collaboration across the value chain, which is one reason why Stream is proud to be a member of the **iMasons Climate Accord**, a "coalition united on carbon reduction in digital infrastructure" which has the mandate to "achieve global carbon accounting of digital infrastructure to influence market-based decisions and drive the industry to achieve carbon neutrality."

[27] Science Based Targets initiative (SBTi), **Companies Taking Action**

[28] Data Center Frontier, **Reports: "Tsunami" of AI Demand Drives Record Data Center Leasing**, 27 July 2023.

How We're Responding: From Build-To-Suit To Build-To-Performance Spec

At Stream Data Centers, we've been watching these trends develop and have worked hard to prepare ourselves to respond:

- In response to unprecedented demand and constrained supply, we expanded our dedicated site development team and continued to refine our configurable sustainable data center design. Because there are not many 'easy' sites left due to electrical utility capacity constraints and NIMBYism, among other challenges, we had to add development resources who were super focused on finding available utility resources vs. finding the perfect piece of dirt and simply connecting to the grid like you could do when data centers were smaller and power was less constrained.
- We walked in lockstep with clients preparing to support massive AI environments—in order to ensure that our standard product meets their evolving AI-driven requirements. For example, we developed a configurable and innovative proprietary cooling direct liquid cooling (DLC) design that supports both air cooling and liquid to the rack, with lower cost and less supply chain risk than the vendors who are providing their tenants off-the-shelf DLC solutions.
- In response to rising concern about water use in many geographies and our customers' commitments to sustainability, we shifted toward WUE≈0 cooling systems and carbon-focused data center design.

Instead of delivering facilities with a specified design, Stream is focused on configurable designs that meet our customers' performance requirements and use standard MEP equipment to minimize supply chain risk. We developed our design and construction standard to be configurable to customers' particular performance specs—to meet service level agreements (SLAs) around average density and PUE/WUE, for example—because that's how our customers deliver for their end users. Building to performance spec enables us to deliver the data center capacity our customers need, optimized for speed and cost.

Proactive Land Development

Proactive land development enables Stream to build new data centers quickly to meet unprecedented demand and ensure utility capacity to meet massive power requirements. At the beginning of the COVID-19 pandemic, Stream saw the emerging trend of high demand and constrained supply and set about expanding our dedicated site development team—we quadrupled the number of development professionals.

> "Huge interest in artificial intelligence is driving a land grab for data center real estate."[28]
>
> —RICH MILLER
> Founder and Editor at Large at Data Center Frontier

Proactive land development mitigates risk and speeds time to market.

When “scale” requirements were under five megawatts, the data center site selection process often involved finding and buying sites that were already fully developed, often in existing business parks. These ‘off-the-shelf’ sites were shovel-ready. They didn’t require bringing in high-voltage power lines and building new substations, moving or vacating existing easements or utilities, or mitigating complex environmental risks on greenfield land.

Having developed data centers since 1999, we remember the days when ideal sites—those where pre-development work could be accomplished efficiently and quickly—were relatively easy to come by. That’s no longer the case. In today’s environment, where campus scale requirements are often 500+ megawatts, developing sites that meet the massive size and infrastructure capabilities necessary to serve customer needs is a more complex undertaking.

Proactive land development is a critical component of Stream’s competitive advantage as it drives our ability to develop new data center capacity for our customers on time and without any surprises. It also de-risks our land banking strategy by ensuring sites are shovel-ready and can meet the tight timelines of our tenants. We control a large and growing number of properties across North America—allowing us to meet the disparate needs of sophisticated customers—all of which have been thoroughly vetted to ensure fast development.

Many public and private equity-backed data center providers are restricted on how much land banking they can do given relatively short-term shareholder demands for returns. At Stream, in contrast, because of the way we approach land development we’re able to secure capacity for our customers without overburdening the balance sheet. Working with our capital partners and leveraging our 25 years of successful data center development enables us to source ideal sites, get control of the sites, and then start the entitlement and utility interconnection processes, dramatically decreasing time to market once a customer is ready to contract.

[29] CBRE, **Global Data Center Trends 2023**, 14 July 2023.

Proactive land development ensures that utility power will be available on time.

In many markets, a shortage of available power is an even more challenging constraint than lack of available land, and it’s inhibiting growth of the data center market, according to CBRE.[29] (Not surprising, given that some power utilities are quoting five years to deliver additional capacity and infrastructure such as transformers and breakers.) Because validating utility capacity and managing the deployment of utility infrastructure has become the longest segment of the data center development process, by proactively securing and developing land in key areas (including securing utility power) we’re jumpstarting the development process and dramatically cutting time to market for our customers.

With an in-house team of deeply experienced location strategy and development experts, proprietary GIS technology, and a systematic and strategic approach, Stream is able to navigate the complexity and bring new cost-optimized capacity to market, quickly. We’re also ensuring the most ideal data center sites.

One of our differences is that many of our development team members honed their skills working for hyperscalers, who are the biggest buyers of data center capacity. Another differentiator is our proprietary Geographic Information System (GIS) platform that arms our team of experts to make data-driven decisions. Ours is a systematic and strategic approach informed by hyperscalers’ own processes.

Flexible Deal Structures

Another way we're responsive to our customers' particular needs is by offering flexible deal structures, including optionality around ownership, the ability to reserve capacity long term and the ability to contribute capital to a building or campus to help accelerate delivery and reduce operating expenses. Compared to a publicly traded REIT, real estate developers like Stream typically can offer more flexibility by crafting solutions that meet not just the technical performance specifications but also help meet more traditional 'business' goals for ownership, use of capital, etc.

Configurable Data Center Design

In 1908, Henry Ford proved the benefits of standardization. You likely know the adage, "Any color the customer wants, as long as it's black." Thanks in part to standardization, the Model T was the first widely affordable automobile.

At Stream, we've standardized our data center deployments across our portfolio with a design that meets the design and performance specifications of the world's largest data center users, and enables us to keep costs down and boost speed to market.

Having focused on a standardized MEP package, we are able to aggregate our equipment supply and then deploy where and when it's needed to meet demand across our portfolio. Standardization minimizes the risk that would otherwise come from procuring equipment ahead of demand. By procuring generators, for example, a year earlier than we would have pre-COVID, we are able to keep deployment timetables about where they were before, and we can do that with confidence because we can use the equipment anywhere.

Standardization can speed development time by 20x.

Our standardized design is configurable, so we can continue to adapt to meet our customers' performance specifications.

One of the potential downsides of standardization is rigidity, which doesn't serve well when working with very large users with very specific infrastructure requirements. The specific way we deploy our standard equipment for one customer is often different from the way we do it for another customer. At Stream we balance the benefits of standardization with agility by infusing configurability into our design and construction standard so we can be responsive to customer needs and deliver capacity relatively quickly.

The core of our design is standardized, with configurability built in so we can be responsive to customers' particular performance specs and deliver capacity relatively quickly. For example, a customer could decide to deploy UPS equipment, or not. They could deploy water to the rack, or leverage air cooling, or even require a combination of air and direct liquid cooling. Either way, the decision doesn't affect the long lead equipment we've already procured. Our customers can defer decisions like those until later in the process without extending the development timeline.

Stream's modus operandi enables real flexibility and adaptability with customers. Our role is to deliver what our customers need. We conform to them, but we're participative—we also bring to the table our deep expertise in serving the complex needs of other sophisticated customers. As a result, our customers are getting the capacity that meets their requirements, and they're getting it at the speed they need it.

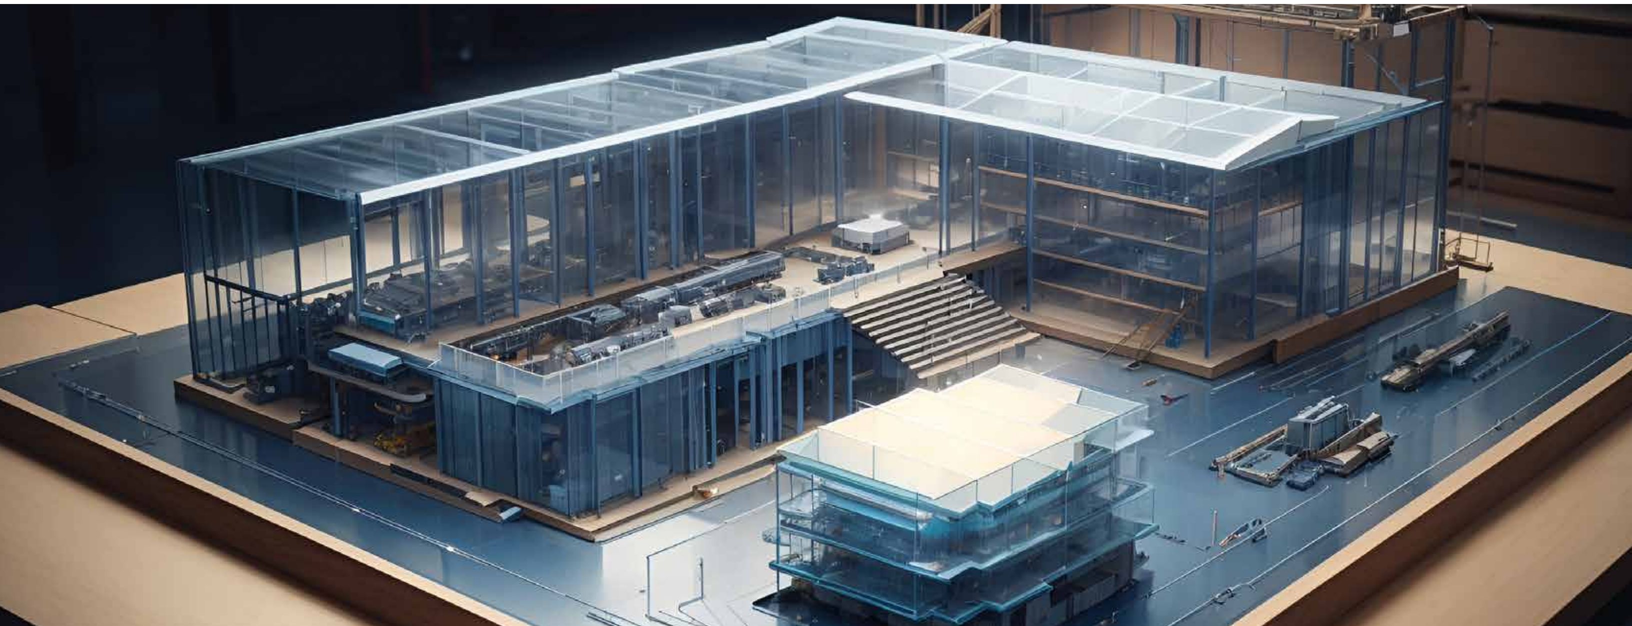

Case Study

How Stream Delivered Over 40 MW of Dedicated Data Center Capacity in 16 Months for a Hyperscale Tech Company With an Ambitious Vision

This story starts with the kind of scenario that keeps infrastructure leaders up at night: One of the world's largest technology companies announcing an ambitious new vision, requiring the very rapid deployment of hundreds of megawatts of new data center capacity. For the company's data center team, delivering the requisite capacity on time is not optional—no matter that off-the-shelf capacity is practically nonexistent even in major markets and development projects are taking over two years on average.

In this kind of environment it is the agile companies—whatever their size and purchasing power—that win. Stream was able to get the customer over 40 megawatts of dedicated data center capacity built and commissioned eight months faster than the current market average, enabling the company's ambitious new vision and its share of a multi-trillion dollar opportunity. How? It was collaboration, built-in adaptability, and flexibility over time alongside a commitment to be rigid when necessary to support the mission—and the committed, smart team members who were enabled to make decisions to drive the correct outcomes.

The experience makes the customer, and Stream, better for what comes next. We are comfortable with high degrees of transparency and collaboration. Sharing good news and bad news, exchanging ideas, and relying on smart people to power quick and accurate decisions. Market circumstances have forced even tighter collaboration and greater degrees of flexibility and we're going to continue to work that way, positioning us to lead in what comes next.

Read the full case study

Configurable and Innovative Proprietary Cooling Design

Configurability in our standardized design enables us to support what's next. For example, seeing the rise of applications leveraging AI and knowing it would dramatically raise densities, we foresaw the need for new cooling approaches. Yet there was no off-the-shelf solution that would accommodate today's air cooled densities and tomorrow's direct liquid cooled requirements without extensive rework and additional expense. So we developed a proprietary design and equipment that can easily support direct air or liquid-to-the-rack cooling.

Our configurable and innovative proprietary cooling design gives us the flexibility to support our customers' performance specs as they evolve. For many data center providers, raising density from 10-12 kW per rack to 400+ kW per rack would be untenable because it would require completely different technology—significantly increasing cost and complexity. At Stream, we're able to support that kind of evolution with less disruption and cost because our standardized design is configurable for what's next. (We'd love to share the details about our propriety design; **contact us**). There would be some downtime associated with reconfiguring a data hall for liquid cooling to the rack, of course, but not more than the customer would take to refresh their technology infrastructure.

One of the significant benefits of a configurable design is that it allows us to evolve as our customers' performance specs evolve. For the foreseeable future, at least, some customers will have 'traditional' average densities (10-12 kW per rack) mixed in with AI high density deployments. Our design supports both, mitigating environmental impact and reducing cost and complexity for tenants as they evolve their technology infrastructure.

A Shift Toward WUE≈0 Cooling Systems

Our design standard has also evolved as customers' thinking about sustainability has evolved. For example, our new design leverages closed loop chillers in most markets—more sustainable because it uses almost no water. This is one of the ways Stream is responding to customers' focus on sustainability beyond PUE and commitment to water efficiency. At Stream's new **Phoenix data center**, for example, we leverage closed-loop chillers to minimize our water usage. At full build, the campus will support up to 280 MW of critical load in 2 million square feet of data center space.

Carbon-Focused Data Center Design

The increasing focus on sustainability beyond PUE has driven us to consider the carbon impact of our data center design.

For example, in addition to supporting the high densities required by AI workloads, liquid cooling to the rack is also more sustainable—with reduced GHG emissions. By reducing the amount of energy required and opening more opportunities for economization, liquid cooling to the rack reduces Scope 2 emissions (indirect GHG emissions associated with the purchase of electricity). By reducing the amount of MEP infrastructure required, liquid cooling to the rack reduces Scope 3 emissions (indirect GHG emissions associated with the value chain; essentially, embedded carbon).

It's no doubt that our understanding of a data center's environmental impact will continue to evolve. Still, we must take action to mitigate our environmental impact, and we must work together to do that. At Stream, in alignment with our guiding principles, we're committed to supporting our customers' climate action goals, and to working with our partners and even our competitors to advance our understanding of how best to mitigate the environmental impact of the data center across the value chain.

Our tenants have some of the most forward-thinking, ambitious sustainability goals of anyone, and our data centers have had efficiency baked in right from the start—even before sustainability was our industry's leading motivator. Our land analysis and feasibility assessments have helped us build advantages from the ground up and keep us mindful from a biodiversity standpoint. Plus, our redevelopment projects help us shave off unnecessary carbon emissions wherever possible. We're quite practiced at reporting on waste, water, emissions and energy already—now we're pushing even farther.

Conclusion: Semper Gumby

Data center development has never been more dynamic. For the Marines, **Semper Gumby** means Always Flexible. In the military, the combination of flexibility when accomplishing a mission and rigidity when necessary is essential to successful (and ethical) operations. In today's environment, finding the appropriate level of flexibility is essential to designing, building, and operating facilities to meet or exceed our customers' SLA while optimizing for speed and cost.

Our customers' technology infrastructure is changing so fast it's a hard ask for them to lock in a particular decision two years before a new data center is scheduled to be commissioned. Being able to defer decisions, for example on cooling technology, until later gives our customers more freedom to match their data center deployments to their current and future performance specifications.

After a data center is operating, rapid technology changes continue to come to bear with each customer refresh cycle. A data center configurable to support the next generation of technology infrastructure is a data center better suited to support innovation. Responsiveness, adaptability, and agility position Stream—and by extension, our customers—to lead, whatever comes next.

Meet the Author

As Stream's Vice President of Solutions Architecture, **Mike Licitra** uses his experience at class-leading financial services and cloud organizations to help customers align their infrastructure strategy and business goals by optimizing design to maximize performance and minimize total cost of ownership. Learn more at **www.streamdatacenters.com/mike-licitra/**

We believe that deploying data center capacity should be a great experience.

streamdatacenters.com

STREAM DATA CENTERS