JANUARY
2026

Managing Sustainability Risk in Data Center Development

By Jeremy Gabe, PhD, and Phil Isaak

NAIOP RESEARCH FOUNDATION

About NAIOP

NAIOP, the Commercial Real Estate Development Association, is the leading organization for developers, owners and related professionals in office, industrial, retail and mixed-use real estate. NAIOP comprises some 21,000 members in North America. NAIOP advances responsible commercial real estate development and advocates for effective public policy. For more information, visit naiop.org.

The NAIOP Research Foundation was established in 2000 as a 501(c)(3) organization to support the work of individuals and organizations engaged in real estate development, investment and operations. The Foundation's core purpose is to provide information about how real properties, especially office, industrial and mixed-use properties, impact and benefit communities throughout North America. The initial funding for the Research Foundation was underwritten by NAIOP and its Founding Governors with an endowment established to support future research. For more information, visit naiop.org/foundation.

About the Authors

Jeremy Gabe, PhD, is the Capstone Advisors Endowed Professor of Real Estate at the University of San Diego (USD) Knauss School of Business. He teaches both undergraduate and graduate real estate courses in real estate finance, sustainability, market analysis and housing policy. Prior to joining USD, Gabe spent 15 years in the New Zealand real estate and architecture research community, consulting with governments and industry councils interested in the adoption of green building practices and green building assessment tools. His applied research addresses the value of urban design, adoption of sustainable building practices and housing affordability.

Phil Isaak, PE, P.Eng., DCDC, RCDD, SMIEE, is a data center infrastructure engineer and program/project manager with more than 20 years of experience. His engineering experience encompasses data center facilities, network infrastructure, IT platforms, and enterprise and cloud data center services. He has guided numerous clients, located in 28 countries across six continents, as they implement enterprise, colocation and cloud service data center strategies.

Disclaimer

Contents

Executive Summary

Technology firms are racing to acquire the data center capacity they need for new generative artificial intelligence (AI) models and related applications, opening opportunities for commercial real estate developers to participate in hyperscale data center development and increasing demand for colocation data centers. Large new data center campuses are going up across the U.S., and vacancy rates at existing data centers have fallen to new lows.

At the same time, the IT equipment (ITE) needed for AI training and inference, machine learning, large-scale data processing and image/video processing is more energy- and water-intensive than for traditional data center uses such as enterprise cloud computing. Most large technology firms have deferred their carbon reduction goals, but state and local authorities and utility companies have increased their scrutiny of the demands that new data center projects place on power and water resources. Local community concerns about undesirable effects from new data center construction remain a perennial challenge for developers. The NAIOP Research Foundation commissioned this report to examine best practices in sustainable data center development. The authors interviewed data center owners, operators, investors, developers and designers to evaluate how development teams are balancing three competing concerns: meeting market requirements, acquiring adequate power, and managing near-term and long-term sustainability risk—the liability that issues related to a project's real or perceived sustainability will result in a financial loss. Key findings from this research include:

- Developers should have a plan in place to address stakeholder concerns about a proposed project's use of power and water resources and the noise a completed data center will produce.
- Local market, regulatory and environmental conditions influence how development teams balance power and water efficiency.
- Some developers are turning to on-site generation to bypass delays from local utilities, but this strategy carries additional regulatory risks and is typically less cost-effective than sourcing power from the grid.
- New data centers are increasingly being built out with liquid-cooled ITE, which is more energy-efficient than air-cooled equipment and is often required to adequately cool state-of-the-art chipsets.
- Data centers face greater long-term financial risks associated with greenhouse gas emissions than other commercial properties due to their large power requirements. This risk can be mitigated by either sourcing low-carbon or carbon-free power or by developing a plan to reduce the property's direct or indirect emissions should the market or regulatory environment require it.

Introduction

A boom is underway in new data center development, driven by exponential growth in demand. JLL, Newmark and CBRE have estimated aggregate colocation data center absorption in 2024 to be around 3 gigawatts (GW) in U.S. primary markets, driving vacancy rates below 3 percent. The total stock of data centers in these primary markets is approximately 10 GW, meaning absorption over 2024 totaled about one-third of the existing stock. Capital is being reallocated to the asset class, which is underpinning the larger boom in AI-related investment. Analysts estimate that another 6 to 8 GW of colocation data center capacity is currently under construction as of early 2025, much of it in primary markets. In addition to the colocation market, large technology firms are announcing new hyperscale projects seemingly every month.

The sudden surge in demand and investment can feel analogous to a gold rush. The near-term financial incentives for new data center investment and development are clear, with demand growth coming from adoption of power-hungry generative AI and cryptocurrency "mining." Given rapid demand growth, investors and occupiers in the U.S. are prioritizing building out capacity over long-term sustainability objectives such as achieving carbon neutrality—at least for now. However, sustainability-related risks—including those related to the stability of the electric grid, decarbonization, water resource management, natural hazards and community integration—have not disappeared and are likely to grow more prominent when the market for new data centers begins to cool.

This report explores how developers and investors can maintain long-term financial returns by addressing emerging sustainability concerns in their feasibility and due diligence processes. As with any other boom in commercial real estate, such as for life sciences in the early 2020s, a point will come when winners emerge, the companies currently fueling rapid demand growth will rationalize their investments, and the pace of absorption will fall. As supply becomes less constrained, both occupiers and investors will become more sensitive to data centers' real or perceived sustainability risks. The authors expect that effectively managing the sustainability risks addressed in this report will allow data center developers and owners to navigate this transition.

This report identifies sustainability strategies that are likely to influence the data center market across a typical investment horizon (e.g., five to 10 years from now). Pioneering developers who have adopted these sustainability strategies provide a rich template for others interested in managing related investment risks.

Defining Data Center Sustainability

Like any other type of development, the degree to which a data center is or is not sustainable can be evaluated by its impact on the environment and local communities. However, in the context of commercial real estate, best practices in sustainable development relate to managing downside sustainability risk—the liability that a project's real or perceived sustainability will result in a financial loss. This report examines sustainability from the perspective of risk management in commercial real estate.

For example, consider one sustainability risk discussed later in this report: greenhouse gas emissions. Although there may not be a current financial benefit for data center developers to limit carbon emissions, data centers face "transition risk," the risk that future government regulations or market conditions will impose costs on owners of data centers that are tied to carbon emissions. In a worst-case scenario, a data center could become a stranded asset if the capital expenditures required to cure such obsolescence exceed its value.

Not all sustainability risks are climate change transition risks; some are more immediate concerns for a developer. Opposition to new data center construction often revolves around effects on electric utility rate payers, the local water supply, and air and noise pollution. If local officials are concerned about future negative outcomes for a community, a project may never receive the entitlements it needs to be completed. This report considers approaches to managing both near-term and long-term sustainability risks.

Research Methodology and Scope

The report's findings are based on unstructured interviews held in 2024 and 2025 with a broad range of stakeholders involved in the data center industry. Interview participants included major investment entities, developers, brokers, market analysts, colocation tenants, hyperscale tenants, architects, engineers, energy companies, start-ups and planning bodies. In addition, the authors conducted a thorough review of industry publications and industry standards associated with data center sustainability strategies.

Much of the report's discussion of sustainability risks and related strategies is framed from the perspective of colocation data center development and operations, as colocation data centers typify the broad range of considerations that apply to for-lease data centers. However, many of the report's observations about sustainability in colocation data centers also apply to enterprise and hyperscale data centers.

Measuring Efficiency and Sustainability in Data Centers

This section briefly outlines the major types of data center development and discusses how data center leases are typically structured, how energy and water efficiency are measured, and how greenhouse gas emissions are categorized and measured. This information provides important context for understanding the sustainability strategies identified in later sections of the report. Many concepts will be new to readers who are not already closely familiar with the data center market and related sustainability concerns. Those interested in a broader introduction to data centers can look to other publications, such as Best Practices in Data Center Development, published by NAIOP in 2025.

Common Definitions for Data Center Development

Colocation Data Center

A large, shared data center facility where multiple occupants lease dedicated space, power and cooling within the building. The occupants' demand is usually measured in kilowatts (kW) consumed by ITE, although larger colocation tenants lease a few megawatts (MW). The data center provider is responsible for the building, cooling, power delivery and physical security, and offers access to a diverse range of network carriers. Colocation facilities are often in urban areas to be close to end users of the data services provided by the IT systems, thus reducing network latency (time delay). Global connectivity is also a priority, so shore landing locations of undersea network cables are prized locations (e.g., Virginia and Ireland).

Hyperscale Data Center

An occupant, typically a major technology company, that builds or leases an entire large data center facility for its exclusive use is referred to as a hyperscaler, and the buildings it leases or owns are referred to as hyperscale data centers. A typical hyperscale facility is designed to deliver 50 to 150 MW of power for ITE.

These facilities are often "build-to-suit," and thus owning the fit-out liability increases risk for investors. While hyperscalers often build and own their facilities, some sign long-term leases with data center developers that construct the building for them (see "Powered Shell" below). Following industrial location theory, hyperscalers may locate in rural areas to minimize land, labor and power costs or in urban areas if low latency or undersea cable access is critical for their services.

Enterprise Data Center

A data center facility that runs the services and applications of a single business. Most large enterprises, like national banks or retail chains, run a large variety of applications at a moderate volume. This results in a need for data centers with general purpose servers that can accommodate various types of applications. Enterprise data centers, also referred to as corporate data centers, can vary significantly in their designed characteristics, ranging from:

- A small room within an office building to single-purpose buildings either on a corporate campus or at a dedicated site.
- A couple of ITE cabinets to thousands of ITE cabinets.
- A computer room measuring from 100 square feet to hundreds of thousands of square feet.
- No redundancy in critical power, cooling or network redundancy to redundant systems plus redundant components within each system.
- Very low power densities to power densities that are near those of a hyperscale data center.

In addition, there is significant variation in the type of ITE needed to support a business's applications. It can be challenging to deploy this mix of systems within a colocation data center or to outsource the applications to a cloud service provider.

Powered Shell

A speculative data center development model, in which a developer constructs the building's core and shell, provides robust utility infrastructure (primarily via securing large amounts of power supply and network connectivity) and delivers it to the site. The tenant is then responsible for building out the interior data halls, including the specific electrical distribution, cooling systems and racks. A powered shell can be leased by a single hyperscale tenant or by a wholesale colocation provider that then subleases the fitted-out space.

Measuring Energy and Water Use in Modern Data Centers

Concerns about data center sustainability concentrate on the sector's natural resource consumption, particularly energy and water.

Energy

Power to run ITE is the primary service leased by a data center occupant. But operating ITE gives off heat, so additional energy is needed to transport and dispose of the heat. Data center occupants are sensitive to energy supply interruption, with various industry standards and guidelines that define levels of redundancy for critical systems. Levels of redundancy are described as:

- Single path: No redundancy provided.
- Single path with redundant components: Contains single points of failure.
- Concurrently maintainable: All components and systems can be maintained without scheduled downtime of IT systems; no single points of failure.
- Fault tolerant.

Some of the industry standards and guidelines that address data center redundancy levels include: ANSI/TIA-942, ANSI/BISCI 002, ISO/IEC 22237 series standards and and proprietary guidelines such as those published by The Uptime Institute.[1] To make supply more reliable, failure protections and redundant systems are installed.[2] As a result, removing redundancies is not a feasible way to improve energy efficiency in the data center industry. Finally, as with any other occupied building, energy is also consumed for lighting, security and space conditioning.

Operational energy efficiency is a key differentiator in the market for data centers. It is measured via a ratio, Power Use Effectiveness (PUE):

$$PUE = \frac{\textit{Total Site Energy Consumption}}{\textit{IT Equipment Energy Consumption}}$$

In this formula, ITE Energy Consumption (often called the "IT load") is the energy consumed directly by servers, data storage and network gear—the equipment for which tenants lease power. The numerator, Total Site Energy Consumption, includes the IT load plus all the energy required to support it, such as cooling systems, lighting and security. If all energy in the data center was consumed only by ITE, PUE would be 1.0. However, since energy is always required for general building systems and ITE cooling, PUE is always greater than 1.0. The more overhead energy consumption, the greater the resulting PUE.

Occupiers can find it difficult to make effective comparisons of different data centers' PUE ratios because owners do not consistently use standardized measurements for the ratio and can advertise a deceptively low PUE by narrowing the definition of "Total Site Energy."[3] Owners can be transparent about how they define PUE by following the ISO/IEC 30134-2 standard, which provides additional clarity regarding consistent measurement boundaries. ISO/IEC 30134-2 also clarifies that the PUE metric should not be used to compare different data centers, as there are too many variables affecting PUE to allow for direct comparison.

Analysts often assume a "typical" North American colocation facility runs at a PUE of approximately 1.5. Hyperscale facilities can be more efficient because a single occupant can streamline the cooling systems for its specific use, such as ITE that uses "direct-to-chip" liquid cooling, which is more energy-efficient than air-cooled ITE. Respondents report that as of the second quarter of 2025, best practice is for new hyperscale data centers to target a PUE of between 1.2 and 1.3.

Water

Water is often the selected medium to dispose of the heat created by a data center's IT load. The heat load per chip of the IT systems currently being installed in data centers is rapidly increasing due to technological development.

Similar to PUE, data center engineers have developed a water efficiency measure most commonly called Water Use Effectiveness (WUE).

$$WUE = \frac{\textit{Total Water Usage}}{\textit{IT Equipment Energy Consumption}}$$

Measured in liters per kilowatt-hour (L/kWh), WUE quantifies how much water a data center uses for cooling (and other uses, such as humidification) relative to the energy its ITE consumes. A lower WUE score indicates higher water efficiency. The WUE metric is now defined in the ISO/IEC 30134-9 standard.

The most common method for cooling ITE is a chilled water system with evaporative cooling. Within the data halls, air handling units pass hot exhaust air from servers over coils filled with cold water. This process warms the water, which is then pumped to a central chiller, which acts like a heat exchanger, transferring the heat from this closed loop of water to a second, separate loop of water. The second loop's water is then pumped outside to a cooling tower. Here, the water is exposed to ambient air, causing a portion of it to evaporate. This evaporation efficiently removes heat from the system, but evaporated water must be constantly replenished.

In addition to cooling, data centers in dry climates may need to use water for humidification systems. These systems maintain optimal moisture levels in the air to prevent the buildup of static electricity, which can damage ITE.

Balancing Water and Energy Efficiency

An interesting challenge in data center sustainability is the trade-off between water and power efficiency (WUE vs. PUE). Cooling methods that use little to no water often require more electricity, and vice versa. For example, a facility using air-cooling might achieve a "perfect" WUE of zero L/kWh but will have a higher (worse) PUE. Meanwhile, a facility using evaporative cooling will often have a lower PUE but will consume more water, resulting in a higher WUE. Because context matters, there is no single "ideal" WUE. Zero may be "perfect" from a water conservation perspective, but the optimal number depends on a facility's location, the climate, and the availability of water and power resources. In some locations, data centers will compete with agriculture for water resources, particularly groundwater aquifers.

The water-energy trade-off leads to a sustainability risk during site selection. Water-abundant locations are preferred if high PUE is a risk (e.g., liability for losses associated with decarbonization or high energy rates), while arid locations will require more energy to operate.

Since water sustainability concerns are broader than the narrow focus on consumption, new metrics, such as water stress (WS) and water usage impact (WUI), are emerging. A cooling design that consumes higher volumes of water on a site located in a low water stress area might be more desirable than a site located in a high water stress area with a cooling design that consumes little or no water. The WUE metric does not differentiate between these two scenarios. Developers and designers can use a new metric developed by The Green Grid to select a site and design a suitable cooling system for the site that results in an optimal PUE/WUE. The Green Grid's method for evaluating resource efficiency includes PUE, WUE, WS and WUI as inputs, and produces a combined energy and water usage effectiveness score.[4]

Greenhouse Gas Emissions

Greenhouse gas emissions are a classic example of a sustainability-related risk. In most U.S. markets, there is no direct price on carbon or regulations capping emissions, so the financial benefits to the developer or investor of reducing emissions are typically small. Data center developers face strong financial incentives to make a building's lighting and cooling systems as energy-efficient as possible. However, except in cases where an occupier requires that a data center directly source its energy from low-carbon or carbon-free generation, there are few incentives for developers to lower carbon emissions themselves.

By contrast, the capital and operational costs of mitigating carbon emissions are real.

The most obvious forward-looking risk is that a government may require a new project to limit emissions to secure needed entitlements, or it could initiate a price on carbon. This type of regulation is more common in Europe and Canada. Comparable regulations in the U.S. tend to be restricted to individual states or local jurisdictions. If the financial costs of compliance are too high in one location, data center development will displace to another location.

A less visible but more immediate transition risk in the U.S. is that investors and occupiers will eventually prefer that data centers limit carbon emissions. Consider a colocation tenant whose parent company is raising capital in a more competitive capital market based on a net-zero carbon strategy memo. The facility owner can expect questions on the facility's decarbonization strategy in any lease or management discussion. Many institutional investors already evaluate sustainability risks when acquiring new properties due to their very long investment horizons. Lenders and insurers also have a long-term interest in real estate, and many include greenhouse gas emissions and physical climate risks in their underwriting. Failure to decarbonize may make it more expensive in the future to retain or recruit new tenants, sell a property or raise capital. Market demand transitions generally happen slowly rather than abruptly, with a balance sheet effect of unexpectedly higher cost of capital and higher exit cap rates.

Measuring Emissions

Scope 1 emissions, also referred to as "direct emissions," are greenhouse gases generated on-site. For data centers, this includes the combustion of fossil fuels (e.g., diesel for backup generators) and fugitive emissions from the gradual leakage of refrigerants used in cooling systems. In jurisdictions that regulate Scope 1 emissions, the data center operator is directly liable for these emissions. Scope 1 emissions are calculated from (a) the import of fossil fuels onto the site, plus (b) a fraction of refrigerants present on-site, typically 1 percent per year, as a baseline gradual leakage rate. Sudden unexpected refrigerant leakages can also occur, adding considerable greenhouse gas emission liability.

Scope 2 "indirect" emissions originate from the production of purchased energy imported onto the site, mainly electricity, which is the most significant emission source for most data centers today. Emission liability is shared between the power producer and the data center.

Cloud Carbon Footprint offers services to calculate operational Scope 2 emissions for data center tenants and operators. The simple formula is as follows:

$$Total\ CO_2e = \left[U_{\text{compute}} \times E_{\text{conv}} \times PUE \times G_{\text{emissions}}\right] + E_{\text{embodied}}$$

In this formula, U_{compute} is the compute usage, measured in hours; E_{conv} is the energy conversion factor from compute usage to electricity consumption (kWh per hour of compute); PUE is the facility energy overhead factor described earlier; and $G_{\text{emissions}}$ is the grid, or imported electricity, emission factor (metric tons carbon dioxide equivalent per kWh). Two variables in the operational emissions equation represent decarbonization opportunities: PUE, where energy efficiency translates into lower emissions, and the electricity emission factor, which is determined by the source of the imported electricity. A zero-carbon energy source, such as solar, wind or nuclear, will have a $G_{\text{emissions}}$ of zero, while electricity taken from the local grid will have an emissions factor based on the average mix of generation fuels. These grid emissions factors, lagged by a few years, are published by the U.S. Environmental Protection Agency.

If a data center generates its own electricity, usually called a "behind the meter" strategy, any fossil fuels imported to the site for electricity generation are considered Scope 1 emissions. The amount of electricity generated on-site is deducted from the site's total electricity consumption for Scope 2 emission calculations.

Scope 3 emissions originate from the value chain of the data center operator. These are emissions from third-party sources related to the data center's entire life cycle. This includes, among many other things, the "embodied carbon" from manufacturing ITE, the development of data center infrastructure, transportation of supplies and employee commutes.

Typical Data Center Emissions Profile

For North American data centers, Scope 1 and 2 emissions are the primary sources of greenhouse gas emissions. Absent a decarbonized operational electricity source, the amount of emissions generated by a data center's power use far exceeds Scope 3 emissions. Respondents who own large U.S. colocation data center portfolios estimate that operational (Scope 1 and 2) carbon emissions represent 80 to 90 percent of the overall carbon emissions from a grid-connected data center over its lifetime.

However, once electricity is decarbonized, either through clean electricity generation or acquisition of clean energy credits, Scope 3 emissions become the primary source. Although Scope 3 emissions are someone else's Scope 1 and 2 emissions, ambitious data center operators are exploring how building materials selection and ITE selection can reduce Scope 3 emissions. More commonly, proactive data center operators measure Scope 3 emissions as a value chain risk assessment activity; should market demand or government regulations require a reduction in these emissions, vendors or customers with high emissions face business continuity risk.

Common Sustainability Strategies: Mitigating Physical Risk and Increasing Energy Efficiency

Developers and owners currently face industry standards and financial incentives to make data centers more sustainable in two ways: by making them resilient to physical risks such as those posed by extreme weather and natural disasters, and by making them more energy-efficient (reducing carbon emissions and costs from inefficiencies). Industry standards provide developers with a way to evaluate and mitigate potential losses from physical risks, while the leasing market for data centers financially rewards energy-efficient providers.

Physical Risk Mitigation and Climate Resilience

Physical risk evaluation is already an important step in data center site selection. An excellent resource to help guide physical risk assessments for data centers is the ANSI/BICSI 002 Data Center Design and Implementation Best Practices standard. The standard includes guidelines and recommended distances from "natural hazards," which include possible seismic or volcanic activity, wildfires, floodplains, areas with possible high winds, tornadoes or hurricanes, unstable ground or subsoil, and groundwater seepage. The standard also includes guidelines and recommended distances from "man-made" hazards, such as airports, buildings storing hazardous or highly flammable material, bodies of water with the potential of flooding, and transportation corridors such as railways and highways. The standard also provides recommendations regarding how close the data center should be to services such as a major highway, an airport (for rapid response of spare parts delivery) and first responders.

While data center industry standards stand out compared with other real estate sectors in managing physical hazards on-site, forward-looking developers are considering physical risks facing supporting infrastructure. For example, how secure is access for employees or the transport of diesel fuel to the site in the event of a long-term power outage? How resilient is the supporting fiber network? These considerations extend beyond the property line of the data center, reaching out to the outer boundaries of the community. Prudent developers will consider these questions when conducting due diligence for a site.

Energy Efficiency

Interviews and secondary sources indicate that common lease structures provide significant financial incentives for owners of data centers to invest in improving operational energy efficiency. This section focuses on lease terms and energy efficiency considerations for colocation data centers, although many of the observations also apply to single-tenant facilities.

Lease Structure Incentives

Leasing rates create strong financial incentives to design a colocation facility with a PUE ratio as close to 1.0 as possible. Colocation lease structures are similar to triple net leases. Tenants pay a per kW rate for their IT load and expense recoveries for the operating costs of the facility. The largest expense recovery is the energy cost of supporting infrastructure, the difference between total site energy consumption and ITE consumption. PUE thus directly affects leasing income, with many colocation operators determining the monthly lease payment as follows:

$$\text{Lease Payment} = (\text{Rate per kW} \times \text{kW of It Equipment} \times \text{PUE}) + \text{Other Recoveries}$$

The negotiated lease rate per kW is multiplied by the leased power, which is further multiplied by PUE for the landlord to recover the site overhead energy costs. Additional recoveries are site-specific and include security, cleaning, maintenance and other common amenities.

A differentiator between colocation vendors is how they determine the kW to be used in the lease rate. One method is to charge based on the quantity of circuits provided to the tenant and the kW capacity of each circuit. Another method is to install kWh metering on the circuits feeding the tenant's ITE. This is clearly the preferred method from the tenant's perspective, as it is directly related to the actual energy consumed rather than based on maximum capacity provided by the colocation vendor. It is recommended that energy meters comply with the IEC 62052 and IEC 62053 series, meeting accuracy-class 2 at a minimum, with accuracy-class 1 preferred. Colocation vendors that use class 2 or class 1 real-time energy meters to determine the kWh will be sought after (i.e., obtain a higher lease rate per kW), as they ensure the tenant is not incurring costs for energy not consumed.

While this net lease structure may seem to create the classic problem where lowering PUE (increasing energy efficiency) results in lower income for the landlord (who must spend capital to lower PUE), *the base lease rate per kW takes PUE into consideration.* Markets naturally allocate higher lease rates per kW to facilities that are more efficient, meaning owners' net operating income (NOI) is inversely related to PUE. High PUE means low NOI, and vice versa.

Consider a colocation facility with a market-average PUE of 1.5 and a negotiated $100 per kW lease rate. The landlord will collect $150/kW (=$100 x 1.5). $50/kW of that collection is a recovery of the overhead energy cost needed to operate the data center, so the landlord books an NOI of $100/kW.

If a competitor invests in energy efficiency and lowers PUE to 1.3, a colocation tenant could pay up to $115/kW as the base rental rate and still pay less overall relative to the PUE 1.5 facility. In this case, the PUE 1.3 landlord will collect slightly less than $150/kW ($115 x 1.3=$149.50). The PUE 1.3 landlord books an NOI of $115/kW.

On the other hand, a landlord operating in the same market that fails to invest in energy efficiency must discount the base rental rate—and realize an NOI below $100/kW—if its PUE is above 1.5.

The net lease structure could also theoretically motivate tenants with IT load timing flexibility to optimize consumption when overhead (i.e., cooling) costs are lowest (e.g., at night). Behavioral optimization that lowers site PUE could lower a tenant's total cost of occupancy without affecting the operator's NOI. Data center owners interested in lowering sustainability transition risks associated with their tenants' energy consumption could update PUE values annually based on measured values.[5]

However, convention generally fixes the PUE overhead charge, similar to a traditional real estate lease with an expense stop. A typical method of incorporating PUE into the lease terms is to determine a value based on a design PUE (dPUE). The dPUE is based on an estimated PUE calculated during the design phase of the data center and never changes throughout the defined lease term. While this is necessary for the first year a data center is in operation since there is no historic data, it might not be in the best interest of the data center owner since the tenant faces no additional costs or rewards for energy-related behavior.

Site Energy Efficiency Strategies

How are facilities lowering PUE? Multiple respondents commented on the energy cost savings resulting from a switch to liquid-cooled ITE.

The objective of all data center cooling systems is to transfer heat away from the ITE, either to the outdoor environment or to some other use. The heat capacity, density and heat transfer rate of water is better than air. Therefore, liquid-cooled ITE will consume less electrical energy, and if incorporated with a chiller and evaporative cooling tower, will also consume less water than air-cooled ITE. Two common methods for liquid-cooled ITE incorporate either "direct-to-chip" liquid cooling or immersion cooling, with direct-to-chip liquid cooling emerging as the preferred solution. Both are closed-loop systems for fluid flowing through the ITE, so neither method directly consumes water. If the data center cooling solution uses evaporative cooling, or any other system that consumes water to cool the computer room, liquid-cooled ITE will result in the data center consuming less energy and water than comparable air-cooled ITE. The appendix provides additional information on how air-cooled and liquid-cooled ITE is incorporated into facility cooling solutions.

ITE incorporated with immersion cooling technology has the fans removed, with the equipment placed into a nonconductive "bath." The liquid in the bath is circulated

FIGURE 1: Immersion Cooling

Photograph of an Ixora immersion cooling system.
Source: Ixora

through a heat exchanger connected to the building cooling system (e.g., chilled water). A variety of vendors provide immersion solutions, from large "racks" that are basically seven-foot cabinets laid on their sides and filled with fluid, to a typical rack-mounted chassis that servers are inserted into and is filled with fluid.

ITE that incorporates direct-to-chip liquid cooling does not have traditional air-cooled heat sinks mounted on top of the processors. Instead, a cold plate is mounted on top of each processor and fluid flows through pipes directly connected to the plates. The liquid flowing through the cold plates within the ITE is piped back to a central distribution unit, which is essentially a heat exchanger connected to the building cooling system.

FIGURE 2: Direct-to-Chip Liquid Cooling

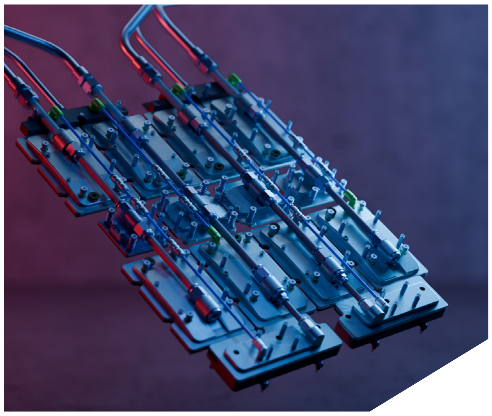

Photograph of a ZutaCore HyperCool Server Loop for NVIDIA HGX B300, an example of direct-to-chip liquid cooling that is currently on the market.
Source: ZutaCore

As chips for high-performance computing become more powerful (and densely packed), direct-to-chip liquid cooling is not just an energy-saving option; it is becoming the only viable solution to prevent damage from overheating. Most liquid-cooled ITE also has components that are still air-cooled. It is common for only high-power processors to be liquid-cooled. The challenge for colocation vendors is to design the correct mix of air-cooled and liquid-cooled capacity in the data center computer hall while providing flexibility to meet changing requirements as tenant needs vary over time.

Furthermore, decisions around air-cooled and liquid-cooled ITE are not typically made by data center facility designers or owners, which are not invited into the conversation in most cases. The tenant decides the cooling technology used by their ITE, and this can change over the lease term as that equipment is replaced at the end of its life cycle. Colocation vendors that have the flexibility to accommodate a variety of air-cooled and liquid-cooled ITE will typically be able to negotiate higher lease rates than vendors that have limited or no ability to accommodate liquid cooling. However, this flexibility comes with a higher investment cost for the developer and investor.

Other examples of design and operational methods being used to improve energy efficiency include:

- Specifying high-efficiency transformers and uninterruptable power supplies, evaluating the efficiency of the systems over wide utilization rates.
- Automating cooling systems with economization mode (e.g., free cooling[6]), taking advantage of lower energy requirements to cool the data center when weather conditions allow.
- Selecting efficient cooling technology for the regional climate zone and expected ITE heat load.
- For air-cooled ITE:
  - » Maintaining cold aisle ITE inlet temperatures[7] at the upper temperature recommended by ASHRAE.
  - » Requiring that colocation tenants use standardized IT cabinets to ensure proper containment can be consistently implemented. This helps to limit hot air in the hot aisle from recirculating with the cold aisle.
  - » Completely sealing the aisle containment system to restrict the hot air coming off the ITE from recirculating back to the cold aisle.
  - » Sealing gaps between adjacent IT cabinets and gaps between the bottom of IT cabinets and the floor. Many aisle containment solutions do not address these gaps between the hot and cold aisles.
  - » Requiring tenants to install blanking panels for all unoccupied rack units in the IT cabinets.

Participants also described a potential additional revenue source via "heat recovery," whereby waste heat from a data center is used to supply heating services to nearby facilities. This strategy is not yet popular in North America, but it is used in European industrial parks, where hot water is circulated to nearby office, flex industrial and warehouse buildings to reduce their energy consumption in winter. The benefit of this heat recovery can be measured with the ISO/IEC 30134-6 Energy Reuse Factor (ERF) metric.

Emerging Sustainability Risks

On top of what is already occurring in the sector, respondents also identified the most important sustainability risks over the next five to 10 years. These fall into three categories: community integration, decarbonization, and retrofitting data centers to liquid cooling. Data centers that address these risks are likely to be more competitive when it comes to securing discretionary entitlements (e.g., energy supply or building approvals) and attracting tenants in the event of a slowdown in demand.

Community Integration and Partnership

Data center development does not occur in a bubble. As in other types of commercial real estate development, local community members frequently voice a range of concerns related to new data center projects. Working to address these concerns and minimize unfavorable impacts on adjacent communities is emerging as a common social sustainability objective for new development. It is also the sustainability risk most likely to affect development in the near term rather than toward the end of a five- to 10-year investment horizon.

Effective developers should have a strategy in place to address community concerns at three key stages in the development process, when they:

- Pursue development approvals (e.g., zoning, site permits, construction permits).
- Seek power from a local utility or regulatory approval for on-site generation.
- Seek "last mile" network access.

Opposition to a new data center project most often revolves around three concerns:

- That the new data center will negatively impact the adjacent community, such as through noise and air pollution.
- That the electricity rates charged by a utility will increase to fund the generation and transmission capacity required by the new data center.
- That local employment created by a new data center after its construction will not be large enough to justify tax or zoning concessions granted to the project.

While zoning largely separates data center campuses from residential neighborhoods, they do attract noise complaints, particularly in Northern Virginia. Complaints report below-regulatory-threshold yet persistent "humming" noises, particularly when backup generators are running. To meet local sound ordinances and limit opposition to new projects, data center developers should proactively address these concerns through design interventions. An engineer interviewed for this report indicated that his firm now offers sound modeling as part of its data center design services to ensure that installed systems will be able to operate well below the constraints of local regulations. In addition, many data center developers build aesthetic and functional barriers around a new campus for security purposes. Clever landscape design that aesthetically integrates these barriers can reduce community concerns about the loss of open space.

Academic research from the Harvard Electricity Law Initiative supports concerns that local utility rates will eventually subsidize some fraction of new generation capacity to meet demand from data centers.[8] Developers can alleviate this concern by investing directly or indirectly in new capacity.

Network access is another critical service that must be procured. The appropriate question to answer is not just is fiber available, but rather whose fiber is at the site and how is it routed from the site to the regional ingress/egress nodes of major network service providers? If the site requires network service providers to build out the "last mile" of fiber infrastructure, the timeframe required to get approval, plan and construct the infrastructure can be lengthy, on a similar lead time to obtaining power capacity. However, developers can leverage these network investments to enhance internet and digital services in local communities, potentially helping them gain public support for a new project.

Finally, there is a trade-off between the pursuit of tax savings and community satisfaction with a new project. Developers are experienced at explaining how new development brings jobs, income and economic stability to a community. However, the report authors spoke with two public authorities, one in a large market and one in a small market, and both cautioned developers to avoid "asking for too much." In their experience, developers are typically asking for construction entitlements, public resources (power and water rights), a speedy application and review process, and local tax relief as compensation for investing private capital locally. The last of these—tax relief—is one of the most frequent reasons for community dissatisfaction. Data center development does bring construction jobs, but once in operation, the buildings are lightly staffed. In high-demand markets, such as Northern Virginia, tax relief is not available, and tax incentives are expected to become increasingly scarce in emerging data center markets.

Decarbonization

In the current market for data centers in North America, speed to market and reliability are the dominant incentives. Occupiers find themselves in a competitive market with little opportunity to be selective about the source of a data center's energy supply. A future problem is that fossil fuel generation, particularly natural gas, is usually the best—and often only—option for rapid and reliable energy supply growth. Data centers consume "baseload" power, meaning energy demand is relatively constant throughout a typical day. Fossil fuels, nuclear fission, geothermal and hydroelectricity are reliable sources of consistent power at all times of day. Solar and wind energy supply varies over time and thus requires expensive batteries to distribute power evenly. Geothermal and hydropower are very location specific and may not be an option for most developers. Nuclear fission is widely used. However, the build out of additional new capacity paused following well-publicized accidents that occurred a decade or more in the past, such as the 2011 nuclear accident in Fukushima, Japan, that resulted from the failure of underlying infrastructure required for active cooling. Many established companies and startups are involved in nuclear fusion research, but knowledgeable interview participants all estimated that fusion technology is at least a decade away, not including regulatory review or permitting processes. Thus, it falls outside a typical developer's investment horizon. That leaves natural gas as the most accessible form of baseload power to fuel rapid growth.

The sustainability transition risk with natural gas is that the market will demand decarbonization at some point in the future. Although it is one of the less carbon-intensive fossil fuels, burning natural gas produces emissions. Those emissions fall under Scope 2 if the local utility invests in gas generation (the most common strategy) or Scope 1 if the data center burns gas to generate power "behind the meter" on-site (an emerging trend). Future carbon-conscious capital and tenants may find natural gas-powered facilities less attractive if the market places a stigma or a price on carbon.

OFF-SITE OR ON-SITE? One of the major debates underway is whether data center developers should obtain energy supply from the local utility or attempt to accelerate the project timeline by investing in on-site energy production. Developers in the largest market, Northern Virginia, are reporting more than five-year lead times for new grid supply in the PJM Interconnection, which is too long a wait to meet the current surge in demand. Creative data center developers are taking a more proactive approach to energy supply and investing in electricity generation as part of the development. Developers use two strategies, either investing in energy generation on-site “behind the meter” or investing in off-site energy generation to sell to the utility, thereby decreasing the utility’s lead time.

Behind-the-meter generation may be risky in most jurisdictions because “large-scale” power generators are regulated and must obtain retailing licenses, comply with Federal Energy Regulatory Commission (FERC) regulations, and obtain grid interconnection agreements that can take as long or longer than waiting for new grid capacity. Texas is one exception, where its grid (ERCOT) does not cross state lines and is therefore not subject to many FERC regulations.[9]

To avoid triggering time-consuming FERC regulatory oversight, new developments are buying smaller-scale gas turbines that are typical of full-redundancy backup generation. Essentially, the development doubles its investment in backup generation, intending to use one set of turbines as a primary energy source with the redundant set of turbines as a traditional backup. The exit strategy is to connect to the grid once power is available but use on-site turbines to speed up time to market. One problem with this strategy is that the gas turbine supply chain cannot keep up with demand, leading to similar yearslong waits for delivery and high costs that eat into returns. Another problem is that on-site gas turbine electricity generation is very noisy, potentially creating community conflict in suburban and urban settings. Other “quick” on-site installation options include hydrogen fuel cells, solar and battery storage, although these technologies lack the scale that gas turbines can provide. The largest planned fuel cell installation as of mid-2025 was an 80 MW industrial park project in Korea, and most solar power for data centers is obtained via an off-site power purchase agreement.

The authors spoke with facility operators who are planning or already investing for decarbonization. The dominant on-site strategy for rapid energy acquisition in North America is an option play: invest in gas generation today, with an option for decarbonization in the future. The dominant off-site strategy is a power purchase agreement (PPA) with an energy provider, where the grid operator serves as an intermediary in the transaction.

Option 1: Using Gas for On-site Hydrogen Fuel Cells

Fuel cells generate electricity through an electrochemical reaction, not combustion. In a pure hydrogen fuel cell, hydrogen and oxygen are combined to produce electricity and water, with no carbon dioxide emissions. The challenge to widespread adoption has been the lack of infrastructure for producing and distributing clean hydrogen safely and economically.

To bridge this gap, companies like Bloom Energy have successfully marketed fuel cell systems that run on piped natural gas. These systems convert natural gas into pure hydrogen, which is then consumed in the fuel cell. While this process still generates carbon dioxide, it is highly efficient. The on-site fuel cells can produce electricity with approximately 20 percent lower carbon emissions than a traditional natural gas power plant and eliminate losses from long-distance electricity transmission. The decarbonization option is a future where clean hydrogen is readily available to be piped directly to these fuel cells, converting them into a zero-emission power source and eliminating direct (Scope 1) carbon emissions.

However, the prevailing business model presents a significant barrier to this green transition. Suppliers typically install this equipment under a long-term PPA, often lasting 15 to 20 years. Under this “energy-as-a-service” contract, the data center agrees to buy power at a fixed price (competitive with current utility rates), but the fuel cell provider retains ownership and control of the assets. This arrangement effectively locks the data center into its current power source for decades, making it difficult to switch to clean hydrogen and fully decarbonize in the future, as that decision will remain with the fuel cell provider.

Option 2: Gas Today, Small Modular Nuclear Reactors in the Future

The promise of small modular nuclear reactors (SMRs) is a component of many decarbonization strategies. Inspired by over 60 years of proven performance in U.S. Navy submarines, these reactors are designed to be small enough for on-site power generation. While submarine reactors are typically 10 to 20 MW, many commercial startups are targeting a "sweet spot" of around 100 MW per reactor.

SMR start-up Oklo, for instance, is designing a 100 MW reactor that requires just three acres of space. SMR proponents claim that advanced designs can use existing nuclear waste stockpiles as fuel, consuming up to 90 percent of the material and reducing its long-term radioactivity. To address meltdown concerns, newer fission technology incorporates passive safety features. These include an automatic shutdown designed to make catastrophic failures in the event of power loss virtually impossible.

The first challenge for data center developers interested in installing new SMRs is federal regulatory approval. SMR developers are pursuing the newer U.S. Nuclear Regulatory Commission "Part 52" regulation, which combines construction and operating licenses, but navigating this process takes up to three years.[10] Most startups are targeting 2028 for final licensing, although outside energy experts see this timeline as ambitious. NuScale, an early entrant, had an SMR design approved under Part 52 in May 2025.

The second challenge is cost. SMRs currently lack the economies of scale and a robust supply chain needed to be cost-competitive. For comparison, the U.S. Energy Information Administration (EIA) conducted a cost survey of grid-scale energy generation technologies in 2023.[11] It reported combined-cycle gas generation to be between $800 and $1,000 per kW in capital costs. SMRs cost about 10 times more, estimated at just under $9,000 per kW. Operating costs are similar according to the EIA study, so significant reduction in capital cost is necessary to be competitive.

SMRs are a long-term decarbonization solution that will likely be available five to 10 years from now. To secure future orders, SMR developers are proposing an interim step for clients that need immediate power. Similar to the fuel cell market, companies like Oklo are offering on-site natural gas generation as a short-term bridge. The strategy is to provide power today, with the contractual option to replace the gas turbines with a zero-carbon SMR once it is delivered and commissioned, at which point the gas generators could serve as backup power.

Option 3: Gas with Carbon Capture and Storage in the Future

Gas turbine manufacturers, such as Solar Turbines, are promoting a decarbonization pathway for their customers using carbon capture and storage (CCS). This technology is designed to separate, or "scrub," carbon dioxide from the turbine's exhaust gas before it reaches the atmosphere. Once captured, the carbon dioxide is purified and compressed into a dense fluid. It can then be transported via pipeline and injected deep underground for permanent sequestration. Current technologies scrub around 90 percent of Scope 1 emissions.

While CCS is not a new concept, its cost has historically been a barrier to widespread adoption. The EIA report prices CCS as three times as expensive as combined-cycle gas generation without CCS. The CCS process is both capital-intensive and requires significant energy to operate. Data center developers considering on-site gas generation as a path to market and interested in sequestering carbon can ask suppliers whether their turbine technology is being designed to be easily retrofitted for CCS should innovation, cost reductions or market demand for decarbonization make it economically viable in the future.

Option 4: PPAs and Indirect Ownership of Clean Energy

With on-site solutions like hydrogen fuel cells, SMRs and CCS representing longer-term decarbonization options toward the end of a five- to-10-year investment horizon, how can a data center get to zero carbon emissions quickly? The primary strategy is to procure clean energy from off-site generation projects through mechanisms that grant the data center the rights to that zero-carbon power.

FIGURE 3: Energy Technology Transition Matrix

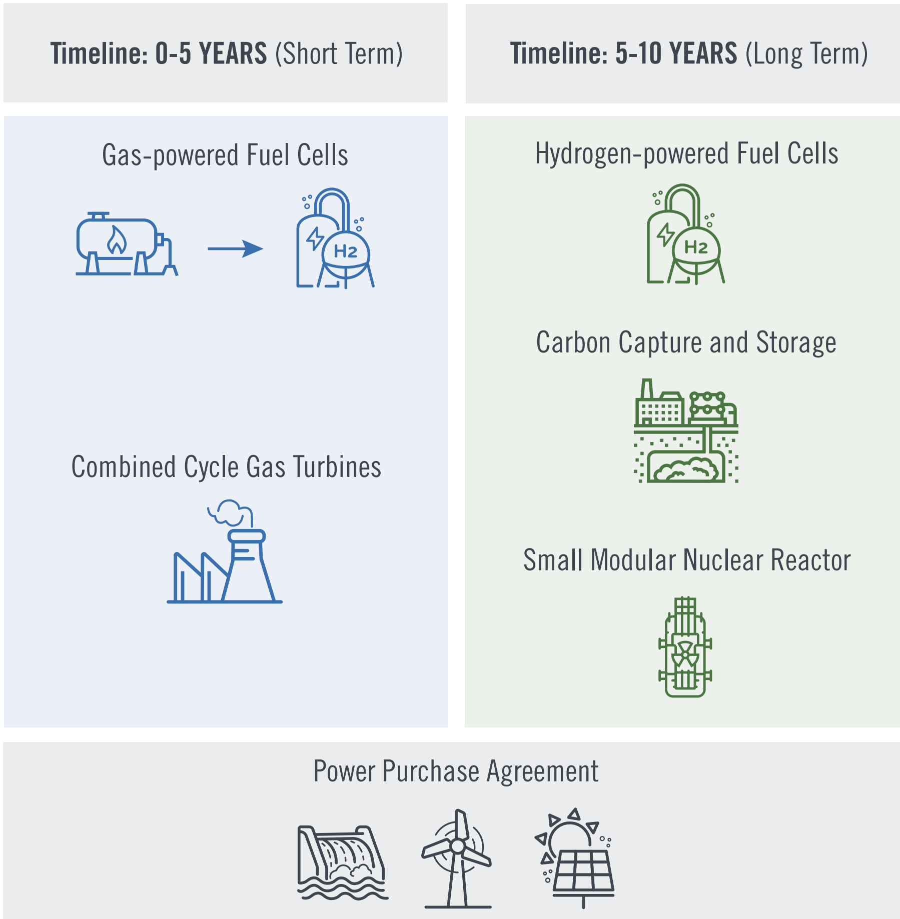

Recent examples of decarbonization in U.S. data centers via PPAs are easy to find. PPAs are contracts through which an energy user buys in bulk from an off-site energy supplier that will add new capacity to the grid. Meta has signed a PPA with Canadian energy firm Enbridge, which will develop a 600 MW solar farm in Texas and credit Meta with the renewable energy. Novva develops colocation data centers with renewable energy credits to attract sustainability-concerned occupants. Microsoft and Constellation Energy agreed to reopen the Three Mile Island nuclear fission plant in Pennsylvania, with Microsoft buying, via PPA, all 900 MW of new zero-carbon energy production for 20 years from 2028. Google recently partnered with energy supplier Broadwing Energy to support CCS technology installed in a 400 MW gas power plant in Illinois.

At a portfolio level, many data center owners are working to achieve 100 percent renewable energy matching their electricity demand. This involves sourcing clean energy equivalent to their total consumption through a variety of instruments, including:

- PPAs
- Green Tariffs: Purchasing clean energy directly from a local utility provider.
- Energy Attribute Certificates (EACs): The formal name for instruments like Renewable Energy Certificates (RECs) in North America, which represent the zero-carbon attributes of one megawatt-hour of electricity.

The environmental impact and credibility of these off-site credits vary. To ensure credible sustainability claims, best practices center on two key outcomes. First, the off-site clean energy from PPAs or EACs should be "additional" or "net new," meaning the data center's investment directly enables the creation of new clean energy capacity, rather than purchasing certificates from a pre-existing project. Second, the clean energy should be local, meaning it is delivered to the same grid from which the data center consumes power. This ensures the company is helping to decarbonize the specific electricity system it relies on. Therefore, a PPA for a new clean energy project on the same grid is considered the gold standard for off-site clean energy procurement, as it meets both criteria. In contrast, simply buying unbundled RECs from a distant grid fails to meet either principle.

Respondents frequently identified a caveat with the PPA strategy for new data center projects. The high demand for low-carbon PPAs has led to the low-hanging "affordable" fruit disappearing, leaving more expensive, risky and time-consuming options on the table, such as contracting with SMR startups or installing CCS technology on new gas turbines.

Retrofitting Existing Data Centers for Liquid Cooling

Colocation vendors respond to market demand, and while it is generally hyperscale data centers that are deploying liquid cooling throughout their ITE ecosystems, tenants of colocation vendors will not be far behind. Vendors that offer wholesale and retail colocation services are already trying to find liquid cooling solutions that work within their data centers. Some already offer these to their tenants. Data center operators are researching how liquid-cooled ITE can be incorporated into their existing computer hall cooling technologies when retrofitting an existing data center. Data centers able to accommodate liquid-cooled ITE can avoid the risk of obsolescence should tenants require those systems.

Retrofitting existing data centers to accommodate liquid-cooled ITE typically is not limited to the cooling methods. Since liquid cooling is more efficient, adopting it also results in the data hall space having greater capacity to support higher IT power loads than was the case with air-cooled ITE. The result is that owners also need to upgrade power distribution to provide power capacity that will match the data hall's increased cooling capacity.

Developing or retrofitting a data center to accommodate liquid cooling also requires a plan for integrating liquid piping systems with the center's internal infrastructure and floor design. While raised flooring within the ITE room historically was used for all but very small data centers, the trend over at least the past decade has been to design computer rooms without any raised floor; ITE racks and cabinets are placed on a floor slab, with all power and network infrastructure installed overhead. Piping is required for each rack or cabinet that has liquid-cooled ITE. Without a raised floor, this piping also needs to be installed overhead in a space that is typically already congested. Data center developers building new multitenant facilities should consider how piping to support liquid cooling will be implemented in the initial construction and modified in the future to accommodate varying tenant requirements. Opting for a raised floor in the ITE room is one design option that can ease future liquid cooling installations.

Conclusion

While data center occupiers currently face a supply-constrained market and a surge in demand for computing capacity, they have limited incentives or ability to be selective about a data center's sustainability. However, a large segment of data center occupiers has made long-term sustainability commitments. The authors expect that when the pace of demand growth slows and data center capacity becomes more abundant, many occupiers and investors will prioritize sustainability in their leasing and acquisition decisions. And in the near term, developers are finding they increasingly have to address local community concerns to secure needed entitlements.

This report identified five major sustainability concerns that developers and investors should plan for in a five- to 10-year investment horizon to maintain market competitiveness:

1. Physical risk resilience
2. Energy efficiency
3. Community integration and engagement
4. Decarbonization
5. Conversion from air-cooled to liquid-cooled ITE

Two of these concerns, physical risk resilience and energy efficiency, are managed through existing best practices that respond to current financial incentives and industry standards.

While most developers are experienced with community opposition to development, data centers bring new concerns that must be addressed to secure needed entitlements and access to water and grid electricity. Newer developments are mitigating noise and the loss of open space through engineering and boundary landscaping. Developers can address electricity cost concerns by investing in new capacity. Tax incentives are growing scarcer for new data center development in response to public opposition, but developers can still earn community goodwill by demonstrating how a project will contribute to the local economy and tax revenues.

Decarbonization risk is currently being exacerbated by the rapid growth in gas generation to meet shortest time-to-operation objectives. Developers can invest in on-site gas generation today while planning for decarbonization over the medium term by contracting with hydrogen fuel cell or SMR providers, or by preparing for the capability to install carbon sequestration capacity. For more immediate decarbonization, data center owners and occupants can invest in off-site clean energy, usually via a PPA, although clean energy PPAs are expected to become less plentiful and more costly over the coming years.

To mitigate obsolescence risk, the growth of high-capacity chipsets and processors will soon require that colocation operators offer liquid-cooled ITE options, which are more energy-efficient and also required for the most power-intensive equipment.

Data center developers that proactively address these sustainability risks will be well positioned to weather the risk of a slower data center market and the possibility of future local, state or federal regulations on their associated carbon emissions.

Numerous technologies are available to cool ITE. This appendix is not intended to address every possible solution. Rather, it highlights a few of the traditional methods and illustrates the difference between air-cooled and liquid-cooled IT cooling systems.

Air-cooled Systems

An air-cooled direct expansion (DX) system is one of the most common cooling technologies used, especially for small data centers. It is also used in applications where the data center has a focus on reducing water consumption required for operations.

Colocation vendors often implement a DX cooling system (also called computer room air conditioning or CRAC), which is similar to an air conditioning system in a home, except at larger scale and with enhanced environmental control. This cooling technology typically requires a lower initial investment but results in higher operating costs.

FIGURE A.1: Example of Air-cooled ITE with Direct Expansion Cooling

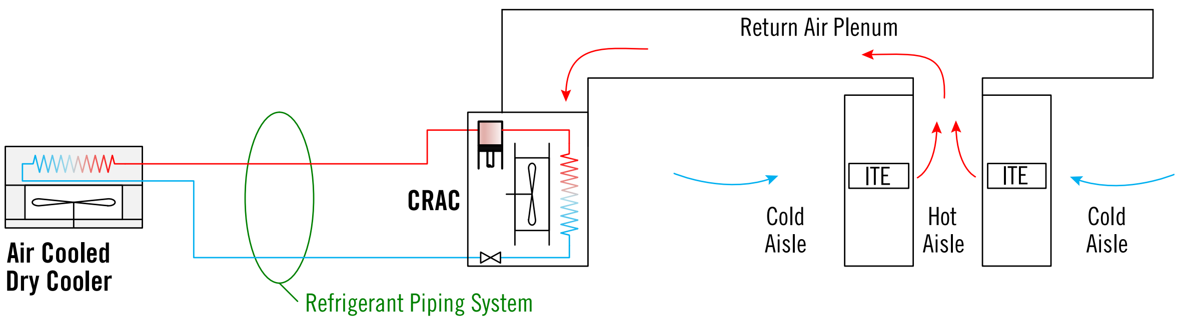

An evaporative cooling tower is another common cooling technology, especially for large data centers. However, with the increased interest in reducing water consumption, these towers are sometimes eliminated as a possible design option.

A chiller with an evaporative cooling tower to provide air conditioning and enable the dissipation of heat to the outdoors is more energy-efficient with respect to electricity consumption than a DX system. However, the cooling tower requires a constant supply of water to absorb the heat from the condenser water system fluid. Evaporation occurs, and the heat is dissipated in the form of water vapor to the outdoor environment. The condenser water system fluid is cooled and returned to the chiller to continue the cooling cycle. The other end of the chiller is connected to a facility water system that is connected to cooling coils. A computer room air handler (CRAH) blows warm air from the computer room onto the coils to cool the air and the warm water from the coils is cycled back to the chiller. A common range of cooling tower water consumption is between 0.5L/kWh and 2.0L/kWh, depending on the local environmental conditions and cooling system efficiency.

FIGURE A.2: Example of Air-cooled ITE with Chiller and Evaporative Cooling Tower

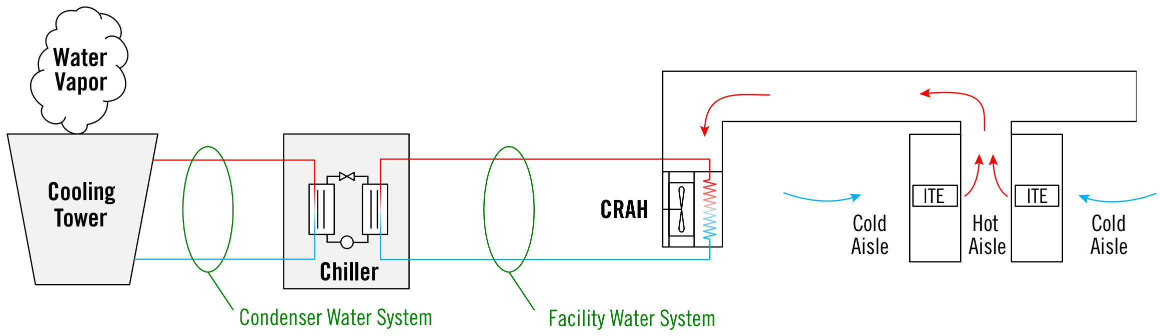

A similar system that does not require a supply of water for cooling incorporates an air-cooled dry cooler with a chiller. The elimination of water consumption comes with higher energy costs than a chiller/evaporative cooling tower combination. The combination of a chiller-based cooling system with an air-cooled dry cooler is often used for medium to large data centers. Interest in this method has increased in the pursuit to reduce water consumption.

FIGURE A.3: Example of Air-cooled ITE with Chiller and Air Cooled Dry Cooler

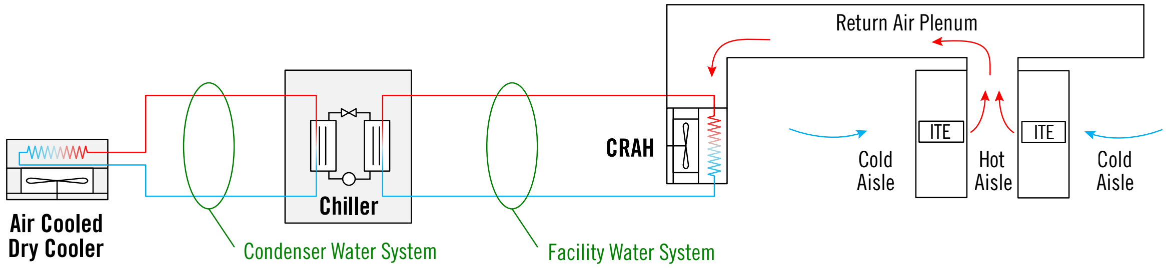

Heat wheel technology has been around for decades and is a proven efficient cooling technology in cooler climates. However, adoption in the U.S. has been limited. A heat wheel cooling system is energy-efficient and does not require a supply of water for cooling. The heat wheel is separated by a barrier, with half the wheel exposed to the air within the data center and the other half exposed to the outdoor environment. As the wheel rotates, it absorbs heat from the computer room, and as the wheel is exposed to the outdoor environment, the cooler outdoor air dissipates the heat to the outdoor environment. This system works only when the outdoor air temperature is below the desired computer room temperature. These systems are typically augmented by DX systems to “take over” the cooling function when the outdoor temperature is near or above the desired computer room temperature. The heat wheel system requires no water and very little electrical energy to operate.

FIGURE A.4: Example of Air-cooled ITE with Heat Wheel

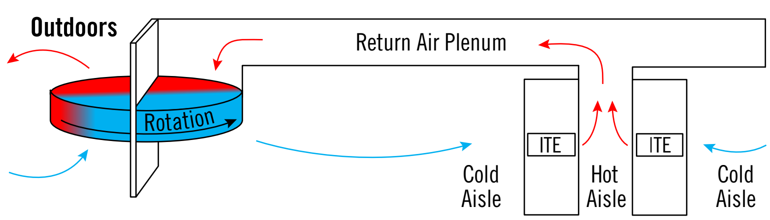

Liquid-cooled Systems

Liquid cooling typically becomes considered for applications where the rack power density is between 20 and 30 kW per rack. At 100 kW per rack, liquid cooling is necessary, as all other solutions become incapable of providing the required cooling.

Three main methods are used for ITE that is liquid-cooled. The first two methods described are used for direct-to-chip liquid-cooled ITE. On this type of ITE, the heat sinks are replaced by cold plates with piping to a cooling distribution unit (CDU). The CDU, cold plates and piping make up a closed-loop technology cooling system.

The system illustrated below has a CDU in each rack, with all the ITE within the rack piped to the CDU and the CDU connected to the facility water system.

FIGURE A.5: Example of Direct-to-chip Liquid-cooled ITE with Chiller and Evaporative Cooling Tower

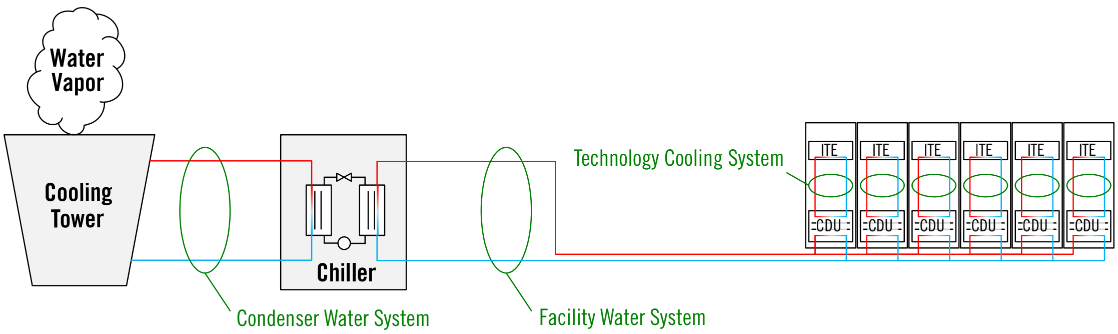

The system illustrated below, also used for direct-to-chip liquid-cooled ITE, has a CDU in the rack at the end of a row of racks. All the ITE within the racks is piped to the CDU at the end of the row. The CDU is connected to the facility water system.

FIGURE A.6: Second Example of Direct-to-chip Liquid-cooled ITE with Chiller and Evaporative Cooling Tower

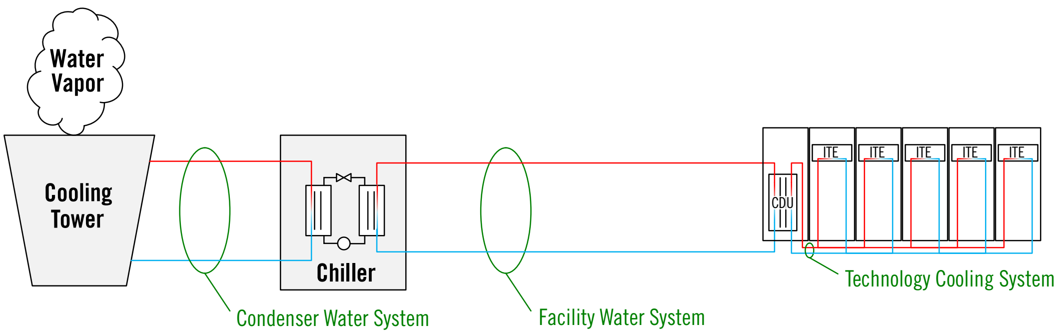

The third example of liquid-cooled ITE is an immersion system. The heat sinks and fans within the ITE are removed and replaced with heat transfer products suited for immersion in liquid. The entire ITE chassis is submerged within the “bath” of nonconductive fluid. The fluid absorbs the heat and cycles through the CDU connected to the facility water system.

FIGURE A.7: **Example of Immersion Liquid-cooled ITE with Chiller and Evaporative Cooling Tower**

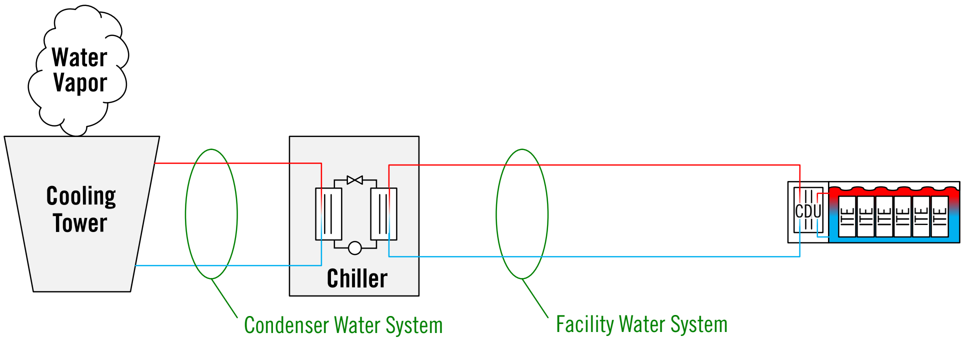

Note that all the liquid cooling solutions illustrated above show an evaporative cooling tower in combination with a chiller. However, an air-cooled dry cooler in combination with a chiller can also be used to eliminate the need for a supply of water to cool the ITE. An evaporative cooling tower solution consumes less electrical energy than an air-cooled dry cooler, but it also uses water to dissipate the heat.

Endnotes

[1] Each of these documents generally follows the above levels of redundancy descriptions. However, there are slight differences between the standards that can impact the design of critical systems. See NAIOP's Best Practices in Data Center Development for more. "Single path" is the lowest rating (less redundancy) while "fault tolerant" is the highest rating (most redundancy). Rents, therefore, are lowest for single path and highest for fault tolerant. However, interviews conducted for this research found that "concurrently maintainable" was the cost-reliability tradeoff most in demand by colocation tenants.

[2] The common infrastructure is a failure path protecting against loss of service in the event of a mains power failure and parallel redundancy within that failure path. The failure path, or paths, typically consists of one or more uninterruptable power supplies (UPS) and backup energy source (typically a generator). Both need to be sized to handle the entire facility's power load. The UPS is typically a large battery energy storage system that takes over immediately, within milliseconds of a mains power loss, with sufficient battery capacity to allow enough time for backup energy source to come online (usually measured in minutes).

[3] A common tactic is to exclude non-IT adjacent power loads from the numerator, such as the energy consumed by adjacent office space or exterior lighting, thereby making the facility appear more efficient than it is.

[4] See The Green Grid, "Whitepaper #93: Data Center Resource Effectiveness (DCRE) v1 Metric," February 17, 2025, https://www.thegreengrid.org/resources/library-and-tools/wp93-data-center-resource-effectiveness-dcre-metric and the associated DCRE v1 scoring calculator.

[5] The ISO/IEC 30134-2 PUE standard requires measurements over a coincident period of 12 months and should be reported at least each month based on a 12-month rolling period.

[6] Free cooling is when the data center cooling infrastructure takes advantage of the cooler outdoor air temperature while eliminating, or significantly reducing, mechanical cooling and associated energy consumption.

[7] The inlet temperature is the temperature of the air as it enters the ITE. Installing sensors at the inlet of ITE is typically impractical, so sensors are often installed on ITE cabinet doors at the top, bottom and middle of the doors; at the end of each row; and intermittently within each row.

[8] Ari Peskoe and Eliza Martin, "Extracting Profits from the Public: How Utility Ratepayers Are Paying for Big Tech's Power," Harvard Law School, March 2025, https://eelp.law.harvard.edu/extracting-profits-from-the-public-how-utility-ratepayers-are-paying-for-big-techs-power/.

[9] However, Texas also passed the Texas Responsible Artificial Intelligence Governance Act in 2025, which may dampen demand for hyperscale data centers in the state, though colocation demand is likely to remain unchanged. While it is too soon to know how Texas courts will interpret the law, the use of online photographs of people cannot be used in commercial AI models without the subject's consent. Since scraping social media is a common tactic to train AI models, the threat of litigation could entice hyperscalers to shop around for favorable regulations outside of Texas.

[10] Existing large-scale nuclear fission reactors are licensed under "Part 50," where construction and operation are separately assessed. Opinion was mixed as to whether the novel Part 52 will be quicker than the more familiar, and already licensed for construction, Part 50 process.

[11] U.S. Energy Information Administration, "Capital Cost and Performance Characteristics for Utility-Scale Electric Power Generating Technologies," January 2024, https://www.eia.gov/analysis/studies/powerplants/capitalcost/pdf/capital_cost_AEO2025.pdf.

“The work of the Foundation is absolutely essential to anyone involved in industrial, office, retail and mixed-use development. The Foundation’s projects are a blueprint for shaping the future and a road map that helps to ensure the success of the developments where we live, work and play.”

Ronald L. Rayevich, Founding Chairman
NAIOP Research Foundation

Recent NAIOP Research Foundation-Funded Research

Available at naiop.org/research

Office Space Demand Forecast, Fourth Quarter 2025
From Static to Strategic: AI’s Role in Next-Generation Industrial Real Estate (2025)
The NAIOP CRE Sentiment Index, Fall 2025
Industrial Space Demand Forecast, Third Quarter 2025
Economic Impacts of Commercial Real Estate in Canada, 2025 Edition
The NAIOP Market Monitor (2025)
Office Space Demand Forecast, Second Quarter 2025
Succession Planning for Commercial Real Estate Firms (2025)
Industrial Space Demand Forecast, First Quarter 2025
Economic Impacts of Commercial Real Estate, 2025 U.S. Edition
Reverse Logistics Strategies for the Post-pandemic Supply Chain (2024)
Recruiting, Training and Retaining Talent in the Real Estate Development Industry (2024)
Commercial Real Estate Terms and Definitions (2024)
Economic Impacts of Commercial Real Estate, 2024 U.S. Edition
Forging the Future: Manufacturing Growth and Its Effects on North American Industrial Markets (2024)

There are many ways to give to the Foundation and support projects and initiatives that advance the commercial real estate industry. If you would like to contribute to the Foundation, please contact Christopher Ware, vice president of partnerships and strategic initiatives, NAIOP, at 703-674-1419, or ware@naiop.org.

For information about the Foundation’s research, please contact Shawn Moura, PhD, executive director of the NAIOP Research Foundation, at 703-904-7100, ext. 117, or moura@naiop.org.

NAIOP | RESEARCH FOUNDATION

We're Shaping the Future

2355 Dulles Corner Boulevard, Suite 750
Herndon, VA 20171-3034

703-904-7100
naiop.org/researchfoundation