DECODING DATA CENTERS:

Sustainability Due Diligence Across the Value Chain

Environmental Defense Fund

nuveen
A TIAA Company

This work is the result of a collaboration between Environmental Defense Fund and Nuveen.

March 2026

AUTHORS:

Environmental Defense Fund
Andrew Howell, CFA
Jenny Mandel
Leanna Tang

Nuveen
Sarah Wilson
Andre Shepley

CONTRIBUTORS:

Environmental Defense Fund
Jordan Faires
Sean Hackett
Maurice Hall
Lauren Johnson
Ted Kelly
Ian Lloyd
Nichole Saunders

Nuveen
Mark Cameron
Chloe Francisco
Adam Guerino
Graham Higgins
Ali Jones
Isha Sharma

Troop Nine Strategies
David Foye

TABLE OF CONTENTS

FOREWORD

By Sarah Wilson, Managing Director, Head of Climate Center of Excellence, Nuveen

Artificial intelligence is gripping investor attention as technology advances reshape global markets with high uncertainty about the industry's path. What has emerged already is a better understanding of the environmental externalities and implications for climate change of data centers' rapidly increasing energy use. As a longstanding responsible investor, Nuveen is committed to leadership and prudent stewardship in the face of heightened volatility and change.

Several months ago, we published A Sustainable Investor's Guide to AI, and with this publication we partnered with Environmental Defense Fund to delve deeper into the material environmental topics. Interrelated issues across energy intensity and emissions, water management, and community impacts and license to operate stand out as key risk factors.

At a systems level, we observe two competing effects to the energy transition and decarbonization: 1) increased energy and grid demand raising fossil fuel utilization rates, potentially delaying fossil plant retirements and catalyzing new gas infrastructure; and 2) elevated clean energy deployment, procurement, and investment supporting accelerated commercialization of emerging carbon-free technologies. Increasingly for hyperscalers, the ability to power compute is a competitive differentiator and scaling strategies will influence the balance of these two opposing forces.

For investors, there are tremendous opportunities, and risks — both idiosyncratic and systemic — in the AI boom. We encourage our peers and the investor ecosystem writ large to utilize this guide for environmental stewardship and risk mitigation and to elevate market expectations for investees throughout the diverse value chain.

INTRODUCTION AND KEY INSIGHTS

The rapid commercialization of artificial intelligence is reshaping businesses and introducing new environmental risk factors throughout the technology value chain. Global AI spending was well over $1 trillion in 2025, and could reach $2 trillion in 2026. For investors, AI is an infrastructure story playing out across a wide swath of sectors, geographies, and asset classes.

Beyond the competitive positioning of the hyperscale technology firms, investors must understand the physical footprint and supply-chain dependencies of the broader ecosystem: the electric utilities powering data centers, the chipmakers providing hardware, the real estate firms developing large-scale facilities, and the companies supplying commodities like cement and steel. Each has its own environmental risk factors, and each presents unique opportunities for sustainability-aligned value creation.

FIGURE 1:

Hyperscale technology companies are the tip of the proverbial iceberg

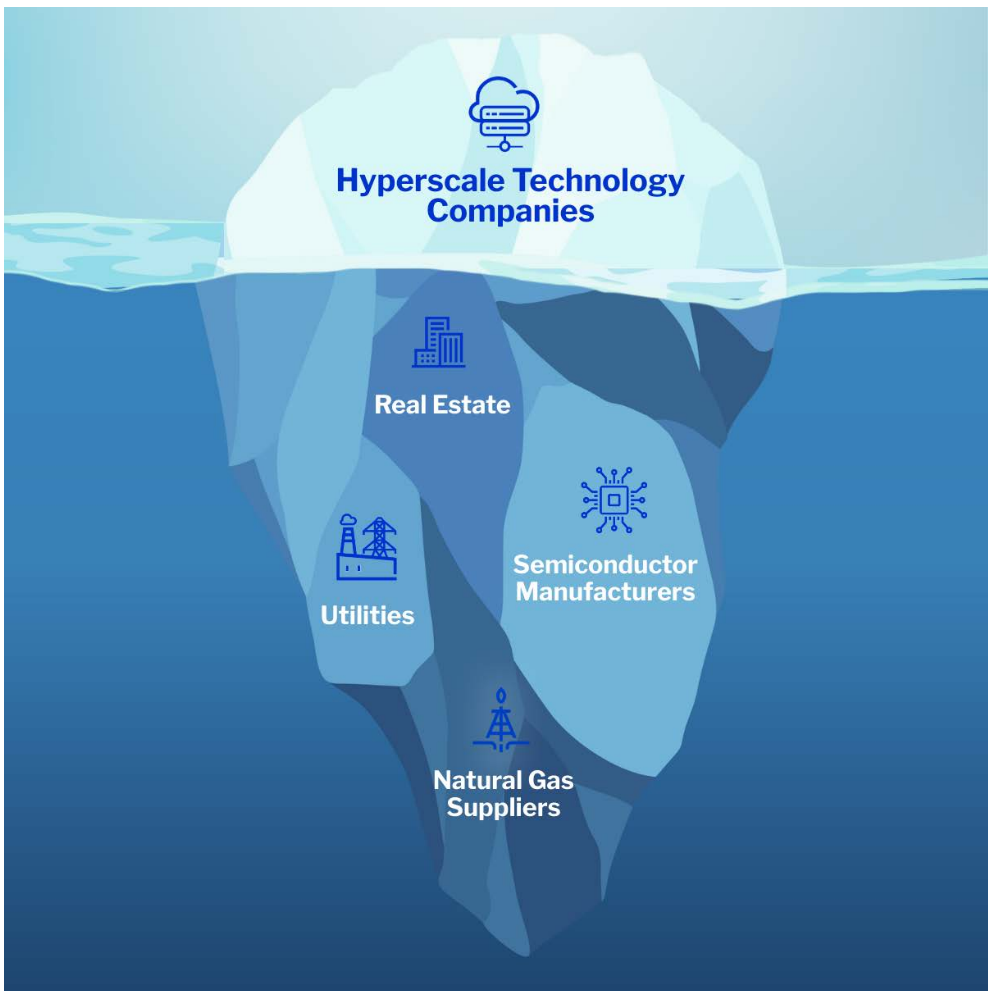

AI infrastructure has real implications for people and the environment. At the heart of data centers' environmental impacts are three over-arching drivers — energy use, water consumption, and local impacts. Data center energy use is driving soaring demand for electricity generation, reshaping the power system, increasing carbon emissions, and pressuring electricity prices with implications for other energy users and the pace of the transition to clean technologies. Water-intensive cooling technologies, alongside the even greater water needs of coal- and gas-fired power generation, can affect resource availability and quality. Social license to operate has become a critical issue for the industry and depends heavily on how data center operators, utilities, and other value chain actors engage with communities and address their concerns.

Business as usual puts investments and communities at risk. These dynamics underscore the critical importance of incorporating environmental performance into diligence activities and company engagement. Weak oversight can expose capital to cost overruns, delays, regulatory intervention, stranded assets, and reputational risk, and the linkages with social license have implications for individual projects and the industry as a whole. Investor insights into these factors could help to unlock better performance.

Investors are well positioned to talk to companies in the AI value chain about incorporating sustainability measures into their data center projects. As technology companies invest in capital-intensive new AI infrastructure, their shifting risk profiles call for due diligence that reflects environmental performance as core to financial performance. Understanding how companies plan, site, power, cool, and govern AI infrastructure is essential to assess critical dependencies on access to electricity, water, and land and mitigate these material sustainability risks.

Key Insights

Throughout this work, we find clear alignment between a science-based understanding of the environmental and social constraints shaping AI infrastructure development and investors' pursuit of risk-adjusted returns. Effective diligence requires asset-class-specific questions, engagement strategies, and expectations about where value can be protected or created.

Among the over 75 questions contained in this guide, three key themes emerge:

1. **The hyperscalers are central, but agency extends well beyond big tech.** Hyperscale technology companies play a central role in driving demand for AI infrastructure and setting expectations across the ecosystem. However, data center environmental and financial impacts extend far beyond these few firms' balance sheets — affecting utilities, real estate owners, semiconductor manufacturers, equipment suppliers, and materials producers. Investors need to look broadly to understand where risks accumulate and where leverage for improved practices reside.

2. **Efficiency should be considered across the value chain, especially with respect to microchips.** While efficiency improvements in chips, servers, and facilities have been substantial, absolute demand for electricity and infrastructure continues to grow rapidly and chip production remains among the most resource- and emissions-intensive segments of the value chain. In particular, understanding how chip efficiency gains, manufacturing constraints, and supply-chain practices interact with data center

expansion is essential to assessing both near-term bottlenecks and long-term sustainability risk. Efficiency also underpins capital planning, grid access, and asset upgrades and expansions, making this a core financial consideration rather than a purely technical one.

3. **Clean energy additionality and natural gas methane emissions merit increased attention.** Many of the hyperscalers have ambitious clean energy goals, but accessing clean energy on the grid is not enough — significant new load from data centers will increase fossil fuel emissions unless companies ensure that load is matched with new carbon-free generation. Companies should thus focus on ensuring the additionality of their clean energy purchasing, while also signaling where emerging solutions like new clean firm technologies can play important roles. Conversely, a significant portion of the electricity serving AI infrastructure is supplied, at least in part, by natural gas-fired generation, making upstream methane emissions an important and often overlooked risk. As measurement technologies improve and scrutiny of climate claims increases, upstream methane leakage can materially affect the credibility of reported emissions and the real-world climate impact of power procurement decisions. Investors should therefore treat methane as a meaningful component of power-related emissions risk.

While this guide is primarily focused on the U.S. market, key elements of the AI value chain are global and many of the environmental impacts, and the strategies to manage their associated risks, apply globally. This guide is designed to help investors ask the right questions about environmental and social impacts across the AI value chain to manage these material financial factors.

INVESTOR ENGAGEMENT QUESTIONS

SECTOR: **TECHNOLOGY**

KEY COMPANIES INCLUDE: *Microsoft, Google, Apple, Meta, Oracle*

TOPIC: ENERGY AND EMISSIONS

For hyperscale technology companies, data center siting and design are critical to operational energy use and related impacts. Mismanagement of these elements can lead to project development delays, higher operating costs, and asset value impairment.

For new builds, firms should maximize zero-carbon energy and be able to demonstrate that they will shoulder incremental power-grid costs. The degree to which firms plan for efficiency and system flexibility can be a critical competitive advantage. From an environmental perspective, investors need to understand not only the emissions profile of a data center when market-based instruments (PPAs, unbundled RECs, etc.) are considered, but also time- and location-matched performance, pathways for future emissions reduction (if applicable), and community air pollution impacts from onsite power generation and backup selections. Where they rely on natural gas-fired electricity, either on-site (behind the meter) or from a utility, hyperscalers should assess and seek to mitigate methane emissions associated with their power supply through procurement choices, contractual standards, and engagement with utilities and/or suppliers. Policy and regulatory needs and priorities vary significantly between jurisdictions and can be critical enablers of clean energy adoption and grid optimization; companies should develop engagement strategies consistent with achieving their risk-aligned environmental performance goals.

1. What share of your current data center power needs is supplied by clean energy, and how will this change over the coming 5 years? How are you managing near-term emissions growth in the context of longer-run goals?
2. What energy and emissions metrics do you publicly disclose at the company and data center facility levels in the U.S., Europe, and other markets, and what are the most useful metrics for company- and facility-level insight?
3. What is your strategy for on-site versus utility-provided power supply, and how does that affect renewable energy use, backup power sources and emissions, local grid reliability, cost impacts to other utility ratepayers, and local land use impacts?
4. How do you incorporate flexibility into data center design and energy planning to meet changing energy and emissions requirements over the life of the asset?
5. How do you work with utilities to enable accurate forecasting and fair cost allocation for data center power needs, and to minimize associated emissions?
6. How do you partner with suppliers to facilitate emissions reductions within the value chain and, in the case of semiconductors, energy-efficient operations?

7. How does your energy procurement strategy balance investment in currently available, best-in-class energy deployment (e.g., solar with energy storage) with investment in emerging technologies that may not be immediately deployable (e.g., enhanced geothermal power or small modular nuclear reactors)?

8. As new data center development proposals face increasing community and political pressures related to electricity prices, how are you addressing energy affordability across your portfolio of existing assets and proposed new facilities?

9. Do any of your existing or planned data center loads rely on new coal- or natural gas-fired generation, or on contract structures that extend the operating life of existing fossil power plants? Have you observed increasing fossil-plant utilization rates and/or increasing grid carbon intensities in recent years where you have large load data centers? If so, provide details.

10. The GHG Protocol is considering revisions to Scope 2 accounting for procured electricity. Do you have a preferred accounting approach (e.g., 24-7 carbon-free energy, consequential accounting, etc.), and how would alternative approaches affect your electricity procurement strategy, reporting outcomes, and climate targets?

11. Where your data centers rely on natural gas-fired electricity, do you engage utilities (if grid-connected) or gas suppliers (for on-site generation) to improve methane performance, such as through the use of verified lower-emissions gas, and how does this engagement affect procurement or contracting decisions?

TOPIC: WATER USE

A meaningful share of existing and planned data center capacity is in moderately to highly water-stressed areas, and water availability and quality are growing factors in siting, permitting, and operational resilience. Commonly disclosed metrics do not provide adequate insight into the relevant risks and differentiation strategies. For investors, unmanaged water risk can translate into community opposition, permitting delays, operating restrictions, and higher costs for alternative water supply or cooling technologies.

1. How do you assess and incorporate regional water availability and local water stress into your data center design and siting processes?

2. Do you implement water-efficient designs at your existing and new build facilities? How do you balance the tradeoffs between energy and water inputs to your cooling systems?

3. Where applicable, do you acquire and retire water rights and allocations in the same basin as your planned facility withdrawals to meet or exceed ongoing usage and anticipated facility expansion?

4. What approaches are successful, and what key challenges do you face, in optimizing water use in leased facilities and sites where a third-party developer carried out initial water use planning?

5. In planning for and managing data centers across states or regions with elevated water stress, how do you incorporate water access and costs across asset lifetimes and how do you assess risks related to geographic concentration at the portfolio level?

6. What steps have you taken to support water monitoring programs for local water systems (local river basin or aquifer) and/or engaged with local water authorities to support long-term needs?

7. Is your water drawn from an established water agency/retailer or from an independent well or diversion, and what is its source? What share of water used is recycled, reclaimed, or non-potable, and how does this vary by facility?

8. How do you track and disclose water withdrawal and disposal at the facility and basin levels?

9. To what extent do you account for energy-related water use (also known as Scope 2 water use) in your water use and disclosure strategies?

TOPIC: TECHNOLOGY AND EFFICIENCY

Technology and efficiency choices affect data center cost, scalability, and environmental footprint. While chip and facility efficiency have improved over time, absolute electricity demand continues to rise, making server utilization, AI model efficiency, and hardware performance-per-watt critical determinants of total energy and cooling requirements. Technology decisions also affect lifecycle impacts: alongside electricity, the manufacturing of IT hardware can represent a significant share of data center value chain emissions, especially given short chip refresh cycles. Efficiency and utilization can materially influence unit economics of microchips, capacity needs, interconnection and expansion timelines, exposure to grid constraints, and embodied-carbon risk in capital planning.

1. How do chip design, AI model efficiency, and server utilization decisions affect downstream energy and water demand, and how are these factors integrated into technology planning? How do you work with chip suppliers to drive efficiency?

2. How do you manage hardware refresh cycles to balance performance gains with increased resource use and embodied emissions?

3. To what extent do efficiency and utilization gains reduce or defer new data center capacity, grid interconnections, or cooling infrastructure, and how are these effects reflected in capital planning and expansion timelines?

TOPIC: COMMUNITY ENGAGEMENT AND LOCAL IMPACTS

Community engagement and management of the local environmental, infrastructure, and economic impacts of data centers have become critical determinants of how data center projects are permitted, financed, and built. Community engagement is a core execution risk: inadequate consultation or failure to address local impacts can translate into zoning challenges, permitting delays, litigation, higher mitigation costs, and cancelled projects. Understanding how companies integrate community and infrastructure impacts into site planning, and the extent to which they engage communities, is essential to assessing the durability of growth plans, the stability of cash flows, and the potential for long-term social license to operate.

1. In your data center siting process, how do you incorporate potential community impacts such as infrastructure stress, water availability, land use, and air quality under normal and backup power conditions?

2. What is your process for consulting with communities around proposed data center infrastructure?

3. How do you work with communities, regulators, and utilities to ensure that costs associated with data center loads are not transferred to other electricity customers within shared service territories or grid regions, and that data center neighbors do not experience power quality impacts?

4. How do you assess and manage cumulative impacts in regions where multiple data centers are developed in close proximity, particularly with respect to energy demand, land use, and infrastructure expansion?

5. What information do you disclose publicly about facility-level resource use (such as electricity demand, on-site primary and back-up generation emissions, and water consumption) and associated infrastructure requirements to enable community engagement and planning?

SECTOR: **ELECTRIC UTILITIES**

KEY COMPANIES INCLUDE: *Duke Energy, Exelon, PSEG, Southern Company*

TOPIC: LOAD GROWTH, PLANNING, AND GENERATION MIX

Utilities face rising capital expenditure requirements, regulatory scrutiny, and political backlash where data center growth drives higher system costs or electricity prices. It also poses system-level planning challenges for U.S. utilities and grid operators, with implications for resource adequacy, capital needs, and customer affordability. In PJM, for example, the grid operator's 2025 long-term load forecast projects ~32 GW of peak load growth from 2024 to 2030 — roughly equivalent to the capacity added over the past 20 years — attributed predominantly to demand from data centers. For investors, it is important to understand how utilities are planning for and managing investments in expanded capacity, system decarbonization, and regulatory compliance, and how system changes account for evolving grid reliability and resilience needs.

Load growth can alter utilities' fuel mix and increase exposure to fuel price volatility as well as scrutiny of emissions and regulatory risk, particularly where incremental demand is served by natural gas or extended operation of existing fossil units. Methane emissions across the natural gas supply chain are increasingly recognized as a driver of climate impact and policy scrutiny. Utilities should minimize new gas buildouts or life extensions and use procurement standards and supplier engagement — such as adoption of measurement-based frameworks like the Oil and Gas Methane Partnership (OGMP) 2.0 — to manage methane-related transition, regulatory, and reputational risks embedded in gas-fired generation for on-site and grid power.

1. What impact has data center load growth had on your generation mix and system-wide emissions over the past 5 years? What are your projections for these impacts over the coming 5 years and 10 years?
2. To what extent is load growth driving new or proposed deployments of coal- or natural gas-fired power, or extending the life of power plants that would otherwise be retiring?
3. Where large load customers express a preference for carbon-free energy, how do you work with them to bring new clean energy supply online?
4. Do you have the regulatory authorities you need to ensure that data center customers cover the full costs associated with their energy services?
5. If you provide generation services, what methane measurement, reporting, and mitigation standards are applied to purchased natural gas? Are upstream suppliers required to meet a methane intensity threshold, such as the industry benchmark of less than 0.2%?
6. Where regulators are considering demand response, curtailment, or load-shaping mechanisms to manage data center interconnection, are you prepared to implement new requirements that may emerge?

TOPIC: WATER USE

While public attention often focuses on data centers' direct water use, water consumption associated with the electricity supplied to data centers is frequently larger and more systemically significant. In the U.S., thermoelectric power generation (encompassing fossil fuels, geothermal and nuclear) accounts for over 30% of total freshwater withdrawals, and thermal plants using wet cooling can withdraw tens of thousands of gallons of water per megawatt-hour generated. Drought, heat waves, and competing municipal and agricultural demands can constrain generation availability, increase operating costs, or force curtailments during peak periods; driving material risks related to resource planning, reliability, and regulatory exposure. Understanding how utilities incorporate water availability and competition into generation planning — and how data center-driven load affects water-intensive assets — is critical to assessing long-term system resilience, cost recovery, and the risk of stranded or underperforming assets.

1. How do data center load projections affect projected water use for electricity generation and cooling across your system?
2. How are water constraints or competing uses incorporated into generation planning and disclosed to customers?

TOPIC: RATEPAYER IMPACT AND COST

Rapidly rising residential electricity costs across much of the U.S., in part tied to data center load growth, are driving intense public debate about cost allocation and affordability in regulated utility systems. Serving large, geographically concentrated data centers often requires substantial new investments in generation, transmission, and distribution infrastructure, which may lead to cost recovery across the utility rate base. In some regions, utilities have proposed or received approval for tens of billions of dollars in new capital expenditure largely driven by projected data center demand — for example, Georgia

Power's recent resource plans include more than $15 billion in new generation investment, with regulators and consumer advocates debating whether residential customers would bear increased costs if data center growth underperforms expectations. For investors, cross-subsidization presents a regulatory and political risk: concerns about rising bills and equity impacts on low-income households can trigger regulatory intervention or litigation that affect cost recovery and earnings stability. Utilities need to work with their regulators to ensure that large load customers pay the costs associated with building or upgrading equipment to serve them, and that equity, affordability, and service quality improve or remain steady, rather than degrading, as a result.

1. How are infrastructure costs allocated between data center customers and the broader rate base? What regulatory or contract protections exist to ensure residential and small commercial customers are not cross-subsidizing large data center loads?

2. To what extent do regulators and the public have visibility into pending projects and large load growth projections? Do NDAs or other privacy shields limit public visibility into utility planning or regulatory agreements related to data centers?

SECTOR: **SEMICONDUCTORS**

KEY COMPANIES INCLUDE: *Nvidia, Intel, TSMC, ASML, Canon*

TOPIC: ENERGY USE AND DECARBONIZATION

The energy demands of AI-capable semiconductors span both manufacturing and operation, with significant emissions at each stage. Advanced fabrication plants (fabs) are among the most electricity-intensive industrial facilities globally, while the production of tools such as EUV lithography systems adds a substantial layer of embedded emissions upstream. As AI demand accelerates, both fabs and equipment suppliers face rapidly rising absolute energy use. Growth increasingly depends on access to large volumes of highly reliable power, the localized impacts and procurement risks of which may be obscured by use of market-based tools such as RECs and EACs. For investors, understanding how energy demand scales across the semiconductor value chain and the energy intensity of manufacturing and chip end use is essential to evaluating the feasibility and credibility of data center expansion as well as the carbon and water footprint of the end product.

1. What proportion of semiconductor manufacturing energy demand is met by clean electricity today, and how will total power demand and emissions evolve over the next five years as advanced capacity and EUV tool production scale?

2. Across major semiconductor manufacturing regions, to what extent are current and projected energy needs contributing to new coal- or gas-fired generation, delayed plant retirements, or increased fossil dispatch?

3. How do you partner with customers on chip power efficiency and thermal performance?

4. The GHG Protocol is considering revisions to Scope 2 accounting for procured electricity. How would alternative approaches affect your reporting outcomes and your customers' Scope 3 reporting?

TOPIC: PROCESS EMISSIONS

In addition to electricity use, semiconductor manufacturing generates process emissions from high-global warming potential (GWP) fluorinated gases used in etching, cleaning, and deposition. Process gases can account for ~30% of total greenhouse gas emissions from advanced semiconductor manufacturing, depending on chip design and abatement performance. While abatement technologies are widely deployed, their real-world effectiveness depends on operating conditions, maintenance, and scale. Transparent measurement, disclosure, and verification of fab-level process emissions are therefore financially material indicators of operational discipline and regulatory preparedness.

1. What high-global warming potential (GWP) gases are used in your manufacturing processes, and how are resulting process emissions measured, verified, and disclosed?
2. What abatement technologies do you use, and how effective are they as throughput increases?

TOPIC: WATER USE

Water is a critical production input for semiconductor manufacturing, used extensively in wafer cleaning, chemical processing, and cooling throughout the fabrication process. AI-capable chip manufacturing is particularly water-intensive: producing a single wafer can require around 2,000 gallons of water, and leading-edge fabs can consume as much as 5 million gallons of water per day. Fabs are capital-intensive, geographically fixed assets with long lifetimes, so exposure to local water stress, drought, or competing demands are potential constraints on permitting, output, and future expansion. Understanding water intensity trends, recycling and reclamation performance, and how local and basin-level water availability are incorporated into siting and capacity decisions is essential to assessing operational resilience and long-term asset viability.

1. To what extent are local and regional water stress incorporated into site selection and planning for manufacturing capacity, and how is planning for future water stress scenarios incorporated into those processes?
2. How do you track and publicly disclose water withdrawals and replenishment activities at the local level?

TOPIC: COMMUNITY ENGAGEMENT AND LOCAL IMPACTS

Like data centers themselves, semiconductor fabs and equipment manufacturing facilities are long-lived, capital-intensive industrial assets whose viability depends not only on technology and market demand, but also on physical resilience and social license. Facilities can stress local water systems, release air pollution from on-site energy systems with impacts on human health, and affect regional land use, transportation, and other infrastructure — especially where multiple facilities are geographically clustered. Semiconductor manufacturing also relies on large volumes of chemicals and ultra-pure water, creating wastewater and hazardous waste streams that must be carefully managed to avoid impacts on surrounding communities and ecosystems. Because these risks are often foreseeable but unevenly disclosed, investors benefit from understanding how companies assess local impacts, engage communities and regulators, and plan for cumulative regional constraints.

1. How are local community impacts including water consumption, wastewater disposal, air emissions, chemical use and waste management, land use, and infrastructure strain assessed for new or expanded fabs? Do these processes consider both facility-level and cumulative impacts?

2. How do you engage with communities and regulators before key decisions are made for new construction or major expansions?

3. How do you assess cumulative impacts on regional water supplies, infrastructure systems, and surrounding communities when multiple semiconductor facilities operate within the same geographic area?

4. What information do you disclose publicly about facility-level water use, air emissions, and wastewater management to enable community engagement and regulatory transparency?

SECTOR: **REAL ESTATE (DATA CENTER-RELATED INFRASTRUCTURE)**

KEY COMPANIES INCLUDE: *Equinix, CyrusOne, GDS Holdings*

Where facilities are developed for operation by another entity or owned and operated on behalf of customers that lease computing capacity, planning gaps and split incentives can lead to sub-optimal outcomes for facility performance. Investors can manage lifetime asset value by assessing processes for aligning incentives and considering long-term cost and resilience drivers.

In addition to the questions below, many of the questions relevant to technology companies will apply to data center real estate developers, as well.

TOPIC: POWER AND COOLING

Power delivery is integral to the selection and development of potential data center sites, with key decisions about renewable energy availability, on-site power requirements, and backup generation strategy determined by the availability of grid and land resources at a particular site. Clear demand signals from target buyers (for sites to be owned by technology companies) or customers (for sites to be fully or partially leased) are crucial for alignment among parties.

1. How are power needs for individual facilities planned, and how does this vary by geography? To what extent are facilities built with capacity for on-site renewable generation and energy storage, and how does the carbon intensity of local grid power factor into siting decisions?

2. How are anticipated costs related to utility upgrades incorporated into asset planning and valuation? What entity bears responsibility for ensuring that facility-related costs are borne by the data center operator or customer(s) and not divided across the utility ratebase?

3. To what extent is backup power planning reliant on diesel, natural gas, or other locally polluting fuels?
4. If a facility will use on-site power generation, how do you manage any potential transition to grid-connected power over time? Explain how power generation ownership is expected to change over the life of the asset and the cost and carbon intensity implications of those changes.
5. To what extent do you engage utilities regarding changes to the carbon intensity of grid emissions over time?
6. In planning for facility cooling technologies, how do you weigh system efficiency and energy and water tradeoffs?
7. How do assumptions about chip power density influence data center design and expansion?
8. How adaptable are facilities to future efficiency standards, cooling technologies, or changes in carbon intensity requirements?

TOPIC: WATER USE

Water availability is critical to site selection and development, as data center water usage for cooling and power generation can place significant strain on local resources. Developers should incorporate considerations around water stress, infrastructure, replenishment, and near- and long-term resilience, especially where rainfall and weather patterns are changing, and should engage early with local water providers and communities to inform decision-making. Transparent disclosures regarding water withdrawals, discharge, and consumption are also essential.

1. Do you assess local water scarcity when considering site locations? How do you incorporate cumulative regional water use and projected future water availability into decision-making?
2. How do you incorporate future facility water use into site planning?
3. To what extent do data center customers prioritize water resilience in their site assessment and metrics, and what demand signals are needed to drive strong performance on water management and replenishment?
4. How do you track and disclose absolute water withdrawals at the facility and basin levels, and what drought contingencies exist?
5. To what extent do you consider investments in local water infrastructure needs or water rights procurement and retirement in the facility planning process?

TOPIC: SITING, LAND USE, AND COMMUNITY ENGAGEMENT

When sites are planned and developed in advance of a customer commitment, it can be challenging to engage nearby communities regarding facility operating characteristics, energy and water needs, and operational commitments. In addition, it can be difficult to accurately assess facility clustering during early-stage project development due to lack of information on other early-stage planning processes and the uncertainty of project completion. Community engagement remains important to identify potential challenges and resource competition pressures, improve decision-making, and support project success and asset value.

1. How do land-use considerations such as zoning, proximity to residential areas or community resources, and ecological sensitivity factor into site selection and permitting?

2. How are local communities involved with site planning processes before key decisions are made?

3. How are the potential cumulative impacts of data center clustering evaluated, especially in light of uncertainty within any one project and limited transparency across regional early-stage development?

TOPIC: EMBEDDED CARBON

Within the data center value chain, small but significant percentages of carbon emissions are embedded within materials such as cement and steel, which are largely produced through highly carbon-intensive manufacturing processes. Although low-carbon cement and steel products exist, to date they have limited market penetration due to factors such as higher costs, small manufacturing scale, and limited performance data. In some cases, data center operators have piloted the use of these materials in building construction with dual goals of reducing embedded carbon in individual projects and accelerating commercialization of lower-carbon materials across the broader economy.

1. Has embodied carbon been holistically considered in the design, material selection, and construction decisions?

2. How do you quantify the decarbonization value of any investments in innovative technologies like low-carbon steel, cement, and mass timber, and what role do such investments play in your sustainability and procurement strategy?

3. How are recycling, secondary materials, and recovery addressed in your business strategy as an opportunity or risk?

4. How do circularity strategies factor into long-term production and capital allocation decisions?

OVER-ARCHING QUESTIONS

TOPIC: GOVERNANCE

Governance over environmental and social impacts associated with AI infrastructure is the foundation of sound risk management. Organizations that assign accountability at the executive and board level, integrate sustainability considerations into incentive structures, and embed oversight into core decision-making processes are generally better positioned to anticipate constraints, adapt to evolving regulatory and stakeholder expectations, and avoid value impairment.

1. Who is accountable at the executive and board level for managing energy, water, and community engagement risks?
2. What expertise do the board or relevant committees have in energy systems, water risk, or infrastructure development, and how frequently are these risks reviewed at the board level?
3. How is executive compensation aligned with prudent management of sustainability risks?

TOPIC: TARGETS AND DISCLOSURES

Transparent, consistent, and decision-useful disclosure is essential for investors and other stakeholders to assess environmental performance, execution risk, and long-term resilience. Data gaps limit companies' own ability to evaluate exposure to grid constraints, water stress, regulatory change, and reputational risk. Conversely, companies that invest in robust measurement and disclosure systems are often better positioned to leverage sustainability as a competitive strength. The act of collecting, analyzing, and reporting granular sustainability data can surface operational inefficiencies, identify emerging constraints, and inform capital allocation decisions. As such, disclosure is not merely a reporting exercise: it is a mechanism for organizational capability and preparedness.

1. What lifecycle emissions and water data do you collect internally, and what data do you provide to downstream customers?
2. What public data disclosures on energy, water, and emissions do you provide at the corporate/portfolio level and for individual facilities, and what voluntary or regulated standards do you align with?
3. How are environmental impacts of data center assets disclosed to investors at the facility or portfolio level?
4. What coalitions or standards development bodies do you participate in, if any, to support the development of comparable, specific, and decision-useful AI-related environmental disclosures?

ENVIRONMENTAL CHALLENGES AND CONSIDERATIONS

In assessing the environmental impacts of data center construction and operations, questions arise across three major topic areas: energy, water, and the local impacts experienced by communities near facilities. Each of these is discussed in detail below, along with considerations for managing them.

1. Energy: Major system constraint and emissions driver

Data centers have emerged as one of the fastest-growing sources of electricity demand globally, with power requirements ballooning as the industry shifts from traditional enterprise-scale facilities of around 10 MW to increasingly large-scale data center campuses with total planned capacities as high as 10 GW — roughly equivalent to the power demand of a large city or some countries (Figure 2). This growth has been geographically concentrated due to infrastructure, policy and other factors, and can occur on compressed timelines, outpacing lead times for new generation and transmission. While data center demand is impelling significant new clean energy and battery deployment, it is also being met by additional fossil fuel generation, particularly natural gas. This creates upward pressure on emissions even as renewable deployment accelerates.

FIGURE 2:

Data center power use

TWh/year of electricity consumption

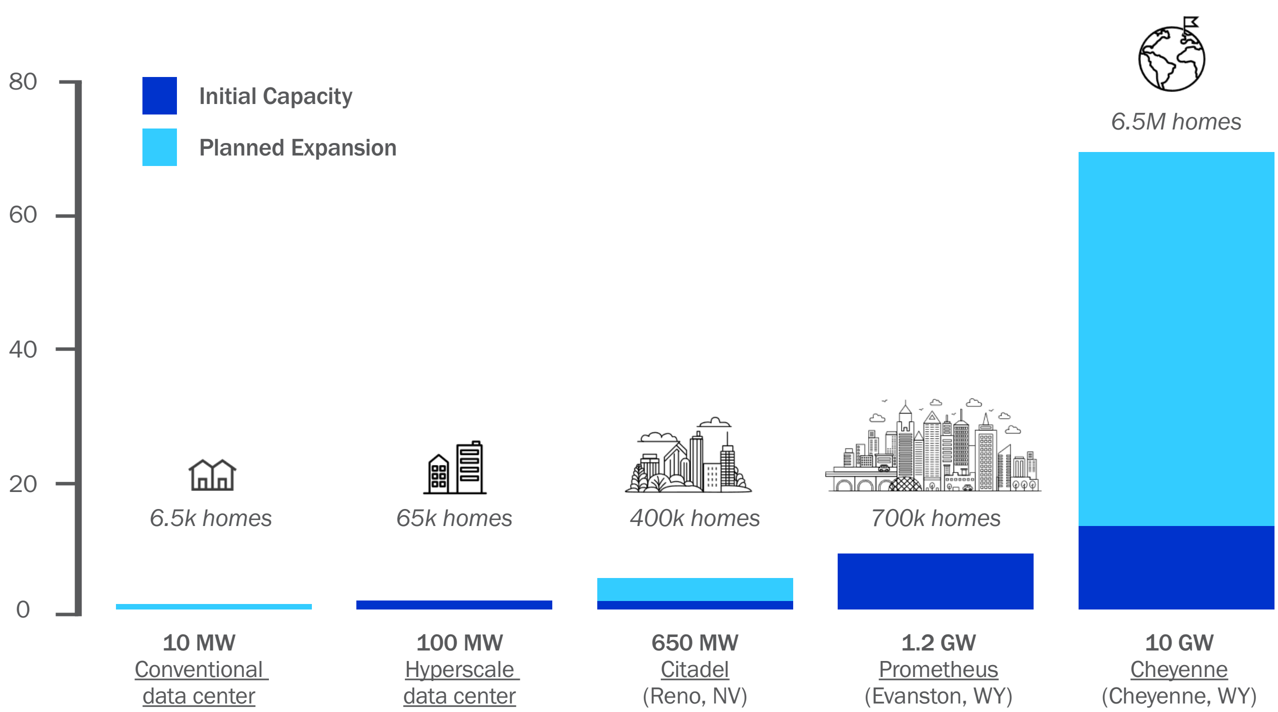

Growing coal and natural gas use are already driving up CO_2 emissions within the AI value chain through both higher utilization of existing fossil fuel plants and new construction, as seen in increases in location-based[1], year-over-year emissions from hyperscale tech firms in the latest reporting data. The IEA projects that natural gas and coal emissions will increase as they meet more than 40% of new data center electricity demand globally through 2030. In the near term, companies should limit their value chain CO_2 emissions by relying on new renewable energy supply wherever possible, and incorporate energy storage, demand flexibility, and other strategies that can help reduce or manage electricity use. In cases where companies rely on natural gas, and particularly where their demand leads to the construction of new plants, companies should consider the deployment of responsibly implemented carbon capture.

The CO_2 emissions associated with growing natural gas-fired power are further intensified by emissions of methane — a greenhouse gas with more than 80 times the warming power of carbon dioxide over 20 years — upstream from the generator. Methane is the main component of natural gas and is emitted via leaks and intentional venting at various points across the natural gas supply chain. Importantly, methane emissions within the natural gas supply chain vary significantly by basin, operator, and asset. Gas is not a uniform commodity from a climate perspective; high upstream leakage rates can greatly reduce or even eliminate any climate benefits of gas compared with other fossil fuels. Conversely, gas with demonstrably low methane intensity grounded in rigorous measurement and verification can significantly cut total emissions as the energy system transitions to cleaner sources.

Many technology companies have adopted clean energy targets that use market mechanisms such as power purchase agreements, energy attribute certificates or unbundled RECs to reduce market-based emissions associated with their grid electricity use. However, a heavy reliance on market-based instruments rather than new, local generation can lead to scrutiny over real-world emissions impacts and unaddressed pollution burdens. Over time, misalignment between data center expansion and power system decarbonization presents legal, regulatory and reputational risks that can impose constraints on asset life and future development.

Energy use by data centers can also present significant regulatory, reputational and other risks. Meta's Hyperion project in Richland Parish, Louisiana, is designed to deliver an initial 2 GW of computing capacity, with the potential to scale to 5 GW over the coming years. To support this load, Entergy Louisiana is constructing three dedicated natural gas power plants and additional transmission lines; Meta has also pledged to add at least 1.5 GW of renewable energy to the grid. This project has sparked significant debate regarding its impact on the regional power network: due to a lack of ratepayer protections, consumer advocates have voiced concerns that residents could eventually shoulder the costs of the new gas and transmission infrastructure if the data center's long-term energy contracts are not sustained. This highlights a core challenge of the "megacampus" era: these facilities are no longer mere customers of the grid, but rather the primary drivers of entire state-level energy strategies.

From a finance perspective, these factors can materialize as interconnection queue lending costs and project risk, contested utility rate cases for generation, transmission and

1 Under the Greenhouse Gas Protocol, location-based emissions are emissions calculated based on the fuel mix of the grid where the energy is used, while market-based emissions are calculated to reflect emissions from energy contracted elsewhere through market instruments such as power purchase agreements and energy attribute certificates.

distribution cost allocation, and stranded asset risk if data centers or the infrastructure constructed to serve them fail to materialize. Regulatory and ESG factors can weigh on projects even once completed, as over-reliance on fossil-fired generation creates climate transition and refinancing risk if regulators later disallow cost recovery or if carbon pricing accelerates. In addition, future methane-specific policy measures such as carbon border adjustments or methane fees could affect costs. This is particularly relevant for firms exposed to export markets such as the EU, where methane performance is increasingly tied to market access.

Strategies to reduce energy-related risks

- **System efficiency:** At the facility and system level, efficiency remains the first line of defense: improvements in server utilization, AI model efficiency and cooling design can significantly reduce total electricity requirements, with implications for initial capital outlay and operating costs.

- **Hardware efficiency:** A focus on lowering the emissions intensity of critical hardware, particularly microchips, can help manage lifecycle energy and emissions.

- **Operational flexibility:** Designing data centers and AI platforms to maximize demand and load flexibility through approaches such as load management, demand response, and peak shaving can reduce the need for new energy infrastructure, ease permitting and siting constraints, avoid costs, and support optionality related to future energy and operational constraints.

- **Utility coordination:** Better advance planning and coordination with utilities, including realistic load forecasts, phased interconnection, and transmission-aware siting, can reduce the risk of delays, congestion, and unanticipated infrastructure costs.

- **Renewable energy use:** Prioritizing procurement of new, local, and time-aligned carbon-free power, paired as appropriate with energy storage, can limit climate impacts, mitigate transition and reputational risk, and reduce exposure to fuel cost volatility. Hyperscalers that "bring their own" additional clean energy supply can simultaneously reduce affordability concerns.

- **Waste heat recovery:** Coordinating with an offtaker to utilize data center waste heat for district heating or other applications can improve system efficiency and reduce costs for both parties.

- **Natural gas:** Where increased natural gas use cannot be avoided, deploying highly efficient technology reduces associated risks. Where appropriate, responsible deployment of CCUS technology should be considered to mitigate carbon emissions.

- **Methane:** Where natural gas use is unavoidable for grid-supplied or onsite power, buyers (data center operators or utilities) can reduce emissions-related risks by engaging suppliers to source gas with demonstrably low methane intensity. Methane performance should ideally be based on direct measurement rather than generic emissions factors, consistent with the highest levels of transparency such as OGMP 2.0 Level 5 reporting. Buyers can also explore the use of differentiated gas or certified low-methane gas programs to incentivize improved methane performance across the supply chain.

- **Coal:** In some cases, increased electricity demand will be met with coal-fired power, significantly increasing data center carbon emissions. Companies should preference lower-emissions energy where possible through workload shifting. Where coal plant life extensions or new coal-fired generation are contemplated, companies should consider the associated energy transition and reputational risks, including whether coal-fired power increases utility rates, and the potential for impaired asset value from regulatory changes.

- **Community engagement:** Early and sustained engagement with local communities and regulators on power sourcing, infrastructure impacts, and cost allocation can help preserve social license, reduce permitting risk, and protect asset value over the long term.

- **Electricity costs:** Data center operators should assess how electricity loads associated with their facilities that interconnect onto the grid could affect local grid reliability and electricity costs for surrounding communities and work with utility regulators and utilities to take full responsibility for costs incurred to meet their needs for primary and backup power, as well as transmission and infrastructure investments needed for the grid to incorporate additional load. Energy affordability challenges, often tied in part to data center demand, are a significant driver of public concern and reputational risk for the industry, contributing to project risk, political and regulatory uncertainty, and the loss of social license to operate.

2. Water: Data centers can exacerbate local water stress

Many data centers have high water consumption due to the use of evaporative cooling for chip temperature management, and all facilities have significant additional indirect water use within the upstream value chain. Data centers that use water-based cooling systems can consume hundreds of millions to billions of gallons of water annually, reducing the availability of surface water or depleting groundwater stores. Water use associated with electricity generation is even greater and can further amplify local impacts. In 2023, power generation associated with U.S. data centers had an indirect water footprint of over 200 billion gallons, roughly equivalent to the annual water use of 2 million households. Semiconductor manufacturing is also reliant on water, with large fabrication facilities using up to 5 million gallons of water per day.

While this scale of water withdrawal may not be problematic where water is abundant year-round, it can have significant implications for local resources in areas with year-round or seasonal water stress. Around 40% of existing and planned data center capacity in the U.S. is located in areas experiencing high or extremely high water stress. Extreme heat events and drought can further drive up water and electricity consumption for cooling and may even require costly transitions to alternative cooling solutions. As a result, facilities in water-stressed regions can face higher costs, community opposition, restrictions on withdrawals, and limits on future expansion. In areas where rainfall patterns are changing due to climate or other factors, inflexible high-volume water use can reduce community resilience and adaptability. Because many companies emphasize efficiency metrics rather than absolute, basin-level withdrawals, these constraints may materialize abruptly, impairing data center asset performance or growth assumptions. There are also tradeoffs between water and energy consumption for cooling, so strategies to optimize water efficiency can ultimately worsen energy efficiency and increase indirect water usage.

In addition to withdrawal volumes, water quality presents a parallel and often underappreciated risk. Data centers that use water-based cooling systems generate wastewater streams that may contain biocides, anti-corrosion agents, concentrated minerals, and other treatment chemicals. In some cases, components used in cooling systems, firefighting foams, or building materials may contain per- and polyfluoroalkyl substances (PFAS), a class of persistent chemicals increasingly subject to federal and state regulation because of their environmental persistence and growing evidence of risks to human health. As regulatory scrutiny of PFAS intensifies, facilities may face higher treatment costs, discharge limitations, monitoring requirements, or liability exposure if contaminants enter municipal systems or local waterways.

Finally, water-related community concerns have already contributed to the delay or cancellation of projects in several markets, increasing execution risk for developers and tenants alike. For example, the Rock Creek Data Center Campus in Fort Worth, Texas has been delayed due in part to concerns around the project's compliance with city water regulations. Similarly, Google's data center campus in The Dalles, Oregon — where water consumption reached 434 million gallons as of 2024, roughly one-third of the entire city's water supply — has received significant media and community attention and pushback for its water use strategy.

Strategies to reduce water-related risks

- **System efficiency:** Reducing absolute data center water demand through advanced cooling designs, expanded chip operating temperature ranges, and air- or liquid-based alternatives can reduce exposure to scarcity and regulatory intervention.

- **Water sourcing:** Maximizing reuse and non-potable water sourcing, including reclaimed or recycled water, can mitigate competition with municipal and agricultural users.

- **Water stress:** Incorporating comprehensive water stress assessments into siting analysis, and avoiding high-risk regions where feasible, can reduce long-term constraints on operations and expansion.

- **Developer coordination:** Where data centers operators work with third party site developers, clear communication of water sensitivity requirements by technology company procurement teams can support suitable site selection.

- **Impact assessment:** Assessing potential impacts on host communities' water access and affordability when evaluating data center siting and cooling options can improve site selection processes and avoid potential community concerns.

- **Community engagement:** Engaging early with local communities and water authorities, with transparent disclosure of the potential quantities and sources of water withdrawals, can help identify opportunities to mitigate water stress, secure social license, and reduce the risk of delays, litigation, or adverse permitting outcomes.

- **Extreme weather resiliency:** Stress-testing operations under drought and heat scenarios can improve resilience and reduce the risk of forced curtailment during peak demand periods.

3. Local impacts: Failure to address community concerns poses risk

The local environmental and infrastructure impacts of data center development are experienced across household, community, and municipal systems, where large-scale infrastructure intersects with daily life and local resource management. In communities hosting dense or rapidly expanding data center clusters, residents report heavy construction truck traffic, persistent noise from cooling and power equipment, increased air pollution from on-site electricity generators, and land-use conflicts that reshape neighborhoods and local planning priorities. These impacts are felt directly as degraded air quality, disrupted sleep, altered community character, and heightened concern about long-term health and quality of life.

Data center facility design choices shape the local experience. Backup power generation from diesel generators — sometimes numbering in the dozens or hundreds per campus — can contribute to localized emissions during routine testing and emergency operations. Where data center campuses include on-site power generation, facility footprints can be dramatically larger and drive additional land use, air pollution, and other impacts. In addition to direct operational impacts, large data center campuses can reshape local infrastructure in ways that extend beyond facility boundaries, for example through the construction of new substations, transmission lines, and other grid equipment.

The accumulation of data center impacts across multiple, clustered facilities can be especially problematic for local communities, as environmental impacts are heightened yet responsibility and control are dispersed across multiple companies. Facility clustering is common due to underlying drivers like favorable regulatory environments, adequate fiber connectivity, and available power capacity. Local oversight and planning processes are often hindered by limited visibility into facility-level resource use. At the same time, because data centers are highly capital-intensive but relatively labor-light once operational, the scale of local infrastructure and environmental impacts may not be matched by proportional long-term employment or other gains. These dynamics can result in a concentration of infrastructure burdens in host communities even as the economic benefits of AI-enabled services accrue more broadly across users and regions.

Households may also experience indirect economic impacts far from any data center's physical footprint. Where utilities invest in new generation, transmission, or distribution infrastructure to serve large data center loads, household utility bills can rise without proportional economic benefits if costs are recovered broadly across ratepayers. Power generation and fossil fuel production can also drive significant impacts to nearby communities, including air pollution and land use changes from infrastructure expansion.

As the scale and pace of data center deployment have accelerated, concerns about local environmental, infrastructure, and cost impacts have prompted increased public scrutiny and engagement around proposed projects. Projects have been delayed, downsized, or modified as a result of resident and local government concerns; dozens of states and localities have also proposed legislation with data center-specific zoning, construction moratoria, mandatory reporting, or other requirements. According to Data Center Watch, 20 proposed data center projects, representing an estimated $98 billion in planned investment across 11 U.S. states, were delayed or halted during 2Q25, with many of those cases occurring due to local opposition or reflecting pullbacks in tax incentives linked to rising political risk.

For investors and developers, loss of social license represents both reputational risk and material execution risk. Community opposition can delay or block projects — over time, failure to address local impacts can constrain development across entire regions, leading to stranded project pipelines, reduced siting flexibility, and reduced asset values. In regions where data center development is highly concentrated, cumulative infrastructure demands and perceived imbalances between local burdens and economic benefits can further accelerate regulatory intervention, zoning restrictions, or permitting constraints. As a result, effective management of community and environmental impacts is increasingly becoming a core determinant of whether large-scale data center projects can proceed at the speed and scale investors expect.

Strategies to reduce local impact-related risks

- **Community engagement:** Early and transparent community engagement, before siting and design decisions are finalized, can surface concerns and reduce the likelihood of conflict-driven delays, particularly when engagement includes disclosure of expected on-site and off-site energy use, water requirements, and infrastructure needs.

- **Community benefits:** Aligning data center development with tangible local benefits such as workforce training and community investment can help secure the support of host communities, improving project viability and protecting asset value.

- **Pollution:** Limiting local air and noise impacts, including through cleaner backup power solutions, reduced generator run times, and noise mitigation, can lower health and nuisance risks for nearby communities and reduce regulatory and reputational exposure.

- **Cumulative impact assessment:** Addressing cumulative impacts, rather than evaluating projects in isolation, can help maintain trust in regions with high data center concentration.

- **Ratepayer protections:** Requiring fair power system cost allocation and preventing spillover onto local and regional ratepayers is important to preserve social license and reduce the risk of contested utility rate cases or political intervention.

- **Infrastructure coordination:** Planning grid infrastructure expansion transparently and coordinating with local governments on transmission corridors, substations, and land-use impacts can reduce community conflict and improve long-term siting flexibility.

- **Governance:** Embedding community considerations into governance and incentives signals that local impacts are treated as core execution risks, not externalities, and can improve community outcomes alongside financial performance.

METRICS AND DISCLOSURES

Information on data center environmental performance is available through public disclosures — largely in the form of company sustainability reports — and private disclosure to investors, both of which offer useful data but also reflect gaps that hinder due diligence.

Public disclosures vary by company and position in the value chain. The hyperscale technology firms and real estate firms developing data centers have tended to disclose the most relevant data; other sectors such as utilities and semiconductors typically provide fewer salient disclosures. Figure 3 summarizes some of the most commonly reported metrics from hyperscalers and real estate operators.

FIGURE 3:

Common data center metrics disclosed by hyperscale technology and real estate firms

Energy	Water	Emissions
Company energy use	Company water use	Company GHG emissions
Company electricity use	Data center water use, total	GHG emissions by type
Data center electricity use, total	Data center water use, facility-level	
Data center electricity use, facility-level	Data center WUE, company-wide	
Company renewable use	Data center WUE, facility-level	
Data center renewable use, total		
Data center renewable use, facility-level		
Data center PUE, company-wide		
Data center PUE, facility-level		

Key Takeaways:

- **Most disclosures are corporate-level, not asset-level.** Companies typically report total corporate GHG emissions, energy use, and electricity consumption, in line with prominent sustainability disclosure frameworks. These metrics are useful for tracking aggregate exposure and progress against targets, but they can obscure where impacts occur geographically and provide limited insight into execution risk, local constraints, or the environmental profile of new data center builds.

- **Data center-specific disclosures are improving but remain inconsistent.** For many hyperscalers, data center operations represent one type of activity among many. Some firms separate data center electricity, renewable energy use, and emissions from those of other operations. However, a lack of transparency around the use and application of market-based instruments for electricity (e.g., RECs) and clean energy sourcing strategies can limit comparability.

- **The gold standard — facility-level electricity data — is rare.** Facility-level energy and renewables data allows investors to assess exposure to grid congestion, energy price volatility, water scarcity, community impacts, and the credibility of decarbonization strategies in specific markets. Few companies disclose this level of detail publicly.

- **PUE is widely reported, but insufficient on its own.** Power Usage Effectiveness (PUE) measures the ratio of total data center energy consumption to the energy used by IT equipment, indicating how much power is consumed by cooling, power delivery, and other overhead. It is the industry's most common energy efficiency metric and can be useful as a facility-level diagnostic tool. However, it does not capture the power efficiency of the computational work being undertaken, absolute electricity demand, carbon intensity of that electricity, geographic variation in cooling needs, or system-level impacts. In some situations, PUE can improve even as a facility's environmental footprint grows.

- **Water metrics lag energy metrics in maturity and consistency.** Total water usage — encompassing water withdrawals, discharge, and consumption — and Water Usage Effectiveness (WUE) disclosures are even less standardized than energy metrics, and facility-level water data disclosures are uncommon. However, without geographic context, water metrics provide limited insight into exposure to local water stress or permitting risk.

- **Tradeoffs between energy and water are often hidden.** Cooling strategies can shift environmental burdens between electricity systems and local water resources. Reporting PUE and WUE together — especially by region or water stress level, as well as, ideally, by facility — provides insight into how companies manage these tradeoffs, but few firms do so.

Limited steps have been taken toward standardizing data center-specific disclosures. The ISSB has outlined key disclosures around the environmental footprint of data center infrastructure that include total energy consumed, the percentage that is grid electricity, and the percentage that is renewable; total water withdrawn, total water consumed, and the percentage of each in regions with high or extremely high baseline water stress; and a description of how environmental considerations are integrated into strategic planning for data center needs. Though these data points provide value to investors, they are not sufficient to assess performance across this complex and wide-ranging industry. In addition, a range of stakeholder groups have come together to address the need for clear and consistent sustainability standards and metrics across the data center value chain. Some of these groups operate on an open basis, while others follow a membership model. However, the sector has yet to coalesce around a set of standard disclosures, making it more difficult for investors to compare and evaluate assets.

There has been a general reticence around project disclosures linked to concerns over competitive or proprietary information. However, this dynamic is shifting on the grounds of public and local community pushback on development and with requests for further transparency from an array of stakeholders. This may move the market towards more comprehensive public disclosures for specific environmental performance indicators, enhancing transparency for leased or privately managed assets, improving comparability across hyperscalers' strategy and environmental risk profiles, and allowing for differentiation based on sustainability performance.

ASSET CLASS CONSIDERATIONS

Different asset classes are exposed to the risks and opportunities associated with AI in distinct ways. Variations in capital source, vehicle structure, ownership rights, ownership concentration, and financial dynamics such as liquidity and diversification all affect the ways in which different classes of investors manage risk and encourage best practices.

A central distinction is between asset classes that primarily hold diversified, portfolio-level exposures and those that are concentrated in individual assets or projects. Public equity, credit, and diversified private market strategies are often more sensitive to system-wide factors such as federal policy, technology trends, public sentiment, and capital markets. Private equity, private credit, and real assets are more directly exposed to site-specific risks, including local permitting and power and water availability. Understanding how sustainability risks manifest at different points along this spectrum helps investors tailor their engagement strategies and align return expectations with the operational realities of data center growth.

Public Markets

Equity, Fixed Income, Securitized Markets

Public equity investors are exposed to data center growth across a broad and liquid set of companies spanning the value chain, including hyperscale technology firms, utilities, semiconductor manufacturers, data center REITs, and materials suppliers. This diversity allows investors not only to gain thematic exposure, but also to adjust portfolio positioning dynamically as sustainability-related risks and opportunities emerge —whether as headwinds, such as water scarcity or regulatory constraints, or as tailwinds, such as advances in carbon-free power and technological efficiency.

In public equities, influence is exercised primarily through engagement with management, proxy voting, and shareholder proposals. Investors should prioritize lines of inquiry that surface marginal impacts (for example, how incremental data center energy demand is served), governance and accountability mechanisms, and the credibility of forward-looking plans. Expectations around strategy, disclosure, target setting, and oversight have already contributed to improvement in operational practices among hyperscalers.

These engagement efforts can also shape outcomes beyond the public markets. Many data center operators are privately held, with capacity leased to public hyperscalers. Through their role as anchor tenants and counterparties, publicly listed companies can extend sustainability expectations to asset-level design, operations, and disclosure, creating spillover effects across the broader data center ecosystem

On the **public fixed income** side, a wave of upcoming debt financing by hyperscalers calls for investors to engage with issuers on environmental factors as they consider participation in particular offerings. For debt investors, risk containment takes center stage, rather than upside, as lenders stress-test an issuer's ability to absorb delays, cost overruns, or regulatory changes linked to power, water, or community impacts. Covenant structures, use-of-proceeds disclosures, and issuer engagement on capital allocation discipline are especially relevant where data center expansion materially changes the issuer's risk profile.

Looking ahead, issuance in **securitized markets** — including asset-backed securities and commercial mortgage-backed securities — is expected to grow, led by both public and private data center operators. These instruments embed the risk profile of hyperscale tenants and provide investors with a new lens through which to assess sustainability-related exposures. Some issuances are labeled as green bonds, with the use of proceeds often tied to improvements in PUE; this is the industry's most common energy efficiency metric but only captures a narrow dimension of sustainability performance.

Private Markets

Credit, Equity, Real Estate, Infrastructure, C-PACE

On the private markets side, exposures tend to be more concentrated and asset-specific characteristics take on a greater importance. For example, while community stakeholder management may affect the ability of a hyperscaler to execute strategy across its portfolio, community opposition to a specific project can have an immediate and disproportionate impact on the value of an individual asset. As a result, **private credit** strategies benefit from embedding sustainability considerations directly into underwriting, covenants, and monitoring. This may include conditions precedent related to permitting or infrastructure cost allocation, reporting requirements on power and water use, and covenants offering protection from material regulatory or community-related disruptions.

For **private equity,** sustainability due diligence should be tightly integrated into both underwriting and value-creation planning. The sustainability factors most relevant to asset value will vary depending on the nature of the exposure: for example, an investment in a data center development platform versus private credit financing of a single asset. In all cases, however, risks such as permitting delays, infrastructure constraints, and rising operating costs can be financially material, as they directly affect tenant demand, pricing power, and exit optionality — particularly where assets are perceived as environmentally or socially challenging. As owners of significant equity stakes and potentially board representation, PE owners are in a strong position to influence management's approach to sustainability and environmental impacts.

In **real estate,** there are key points of agency and transparency that serve as the basis for establishing sustainability expectations. For example, as a real estate owner leasing to a data center operator, the operator-owner relationship and contractual understanding may ensure or derail data disclosure, impacting efforts to improve performance.

When it comes to **infrastructure** investors, factors such as transmission upgrades, water rights, grid capacity, interconnection queue economics, asset durability, and stranded asset risk are highly material alongside other execution risks. This includes evaluating whether assets are adaptable to future efficiency standards, changes in grid carbon intensity, or tightening water constraints as well as whether cost-recovery mechanisms are robust under political and regulatory pressure. Engagement with operators, policymakers, and community stakeholders is often as important as engagement with asset managers themselves. Community approval and social license risk that can halt permitting and strand development pipelines are major concerns.

Commercial Property Assessed Clean Energy (C-PACE) financing allows building owners to fund energy, water, and resilience improvements through long-term, property-tax–assessed obligations that are typically senior in the capital stack and tied to the asset rather than the owner. In addition to potentially mobilizing efficiency-related investments, C-PACE deal originators have a structural opportunity to support sustainability, as the conversion of other forms of real estate assets to data centers can provide a less materials-intensive approach to new capacity build-out. The key consideration is whether financed measures meaningfully reduce absolute resource use and long-term risk, rather than simply improving efficiency ratios. Assessing baseline performance, alignment with broader grid and water constraints, and interaction with tenant-driven demand growth are critical. Well-structured C-PACE investments can enhance asset resilience and credit quality; poorly aligned projects risk underperformance if system-level constraints tighten faster than expected.

CONCLUSION

Across the AI value chain, environmental performance has emerged as a critical investment risk that calls for enhanced due diligence. For investors, how companies address data center energy and water use and manage their local impacts offers competitive insights. Companies that fail to anticipate these factors risk project development delays, rising operating costs, regulatory scrutiny, and reputational damage. Conversely, firms that proactively address energy constraints, manage water risk, and meaningfully engage with communities are more likely to improve project outcomes and support asset value.

While hyperscale technology companies play a central role, responsibility for the environmental footprint of data centers extends across the value chain. Utilities, semiconductor manufacturers, and real estate developers play key roles in how AI infrastructure is planned, built, powered, and governed. Understanding these interdependencies enables investors to identify bottlenecks, anticipate regulatory shifts, and encourage practices that support sustainability-aligned value creation.

In addition to the core environmental drivers, there are other noteworthy opportunities for investor engagement to influence data center development outcomes. These include calling for increased attention to system efficiency, ensuring that clean energy supplies are new and additional, and engaging natural gas suppliers to minimize upstream methane emissions. Asset-class-specific considerations — such as ownership structures, liquidity, leverage, and asset lifetimes — should also inform how investors approach due diligence and engagement.

As the rapid pace of data center buildout places growing demands on the environment and on communities hosting elements of the AI value chain, the financial sector has a meaningful role to play in guiding how this infrastructure evolves. Thoughtful diligence, informed engagement, and capital allocation that rewards responsible practices will best position investors to capitalize on the opportunities and manage the risks presented by the AI economy.

APPENDIX: INDUSTRY STRUCTURE

As of 2024, data centers were geographically concentrated in North America (44% of installed capacity), Asia Pacific (37%), and Europe (16%). The U.S. and China stand out as clear industry frontrunners, accounting for 42% and 25% of global data center capacity, respectively, in 2024.

Industry structure

A range of industries are relevant to the development, operation, and financing of data centers, each controlling different levers of risk and value creation. Upstream in the value chain, **semiconductor manufacturers** convert raw inputs, including silicon and rare-earth elements, into advanced chips through capital-intensive fabrication processes. Environmental impacts at this stage are primarily driven by electricity demand, water-intensive manufacturing steps, and the use of fluorinated process gases with high global-warming potential. These upstream supply chains also rely on the extraction and processing of critical minerals and materials, which is outside the scope of this review.

FIGURE 4:

AI value chain

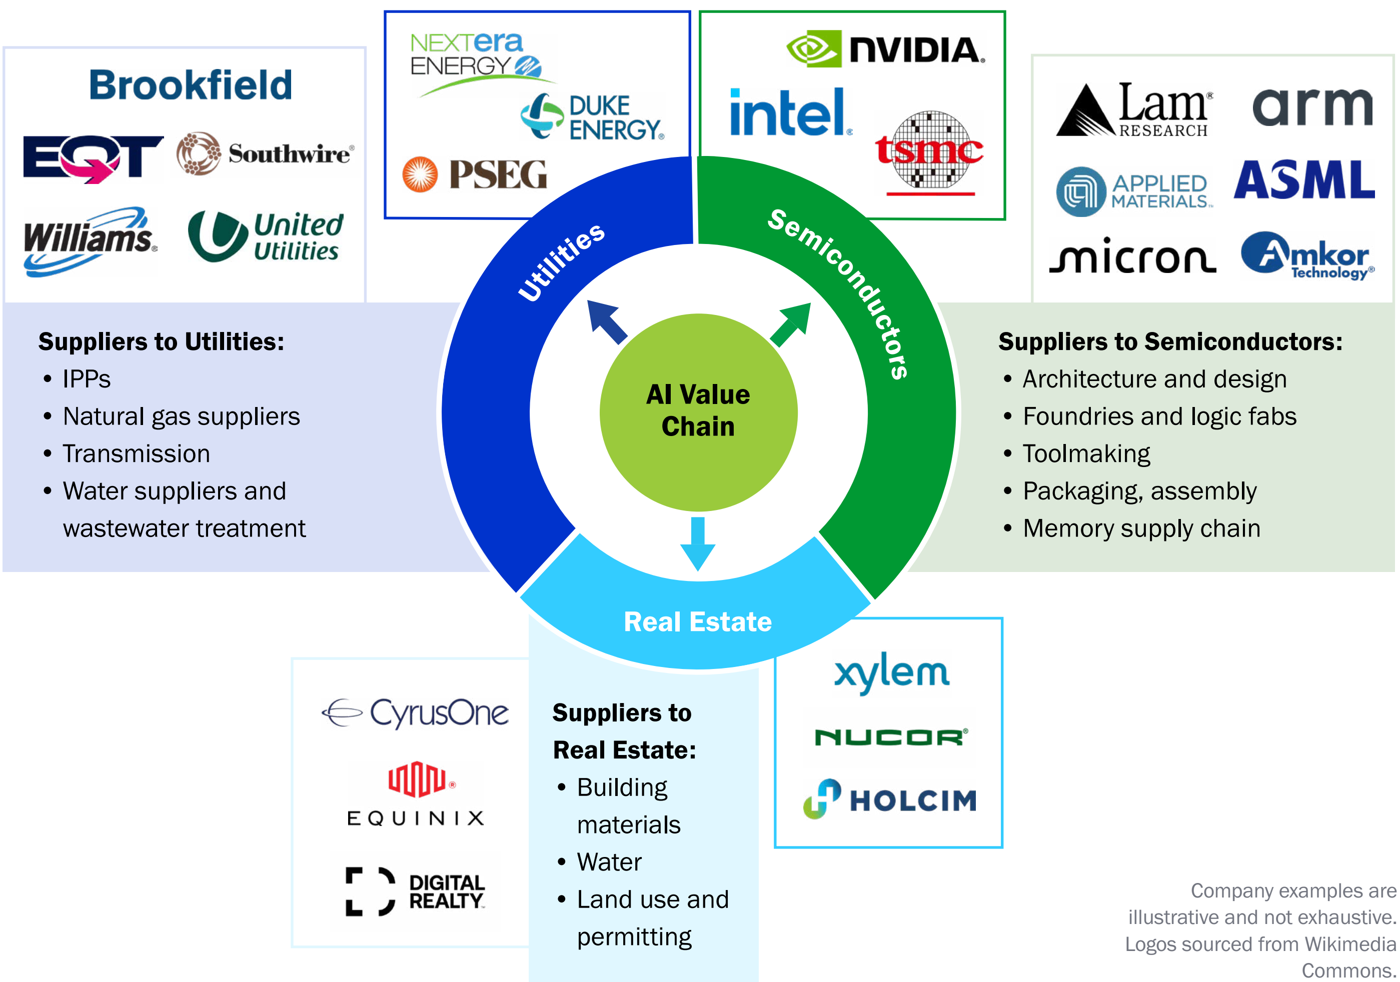

Company examples are illustrative and not exhaustive. Logos sourced from Wikimedia Commons.

Utilities and real estate developers sit at the physical foundation of data center deployment. Electric utilities must plan for large, geographically concentrated loads that can reshape generation mixes, transmission requirements, and rate structures, while managing reliability and affordability for existing customers. Data center real estate owners and developers make long-lived siting, design, and cooling decisions that shape asset exposure to grid constraints, water stress, and community impacts, influencing permitting outcomes, social license to operate, and long-term upgrade and expansion potential. For investors, these sectors present classic infrastructure risks, including capital intensity, execution risk, regulatory exposure, and long-term asset durability, further amplified by the pace and scale of AI-driven demand.

At the downstream end of the value chain, **technology firms** including the hyperscalers (Amazon, Alphabet/Google, Meta, Microsoft, and Oracle) drive demand for data centers through cloud computing and AI model development and deployment, shaping capacity expansion across the upstream industries. Even where technology companies do not directly own data center facilities, power generation assets, or semiconductor production plants, their procurement strategies, siting selections, and growth trajectories strongly influence where infrastructure is built, how it is powered, and which environmental and social risks are borne by communities and ratepayers.

Even further downstream, the businesses, governments, and individuals using AI have limited visibility into the electricity cost implications, emissions, water use, and localized environmental impacts of AI infrastructure. As this information becomes more widely available, poor practice can affect public perception, customer acceptance, and willingness to pay, all financially material value drivers for AI-enabled applications and investments.

However, data center sustainability remains a moving target as the industry rapidly evolves and norms and best practices for site design and construction emerge. The Innovation for Cool Earth Forum's 2025 Sustainable Data Centers Roadmap offers a deep look across the industry's energy and water use.

Environmental Defense Fund

nuveen
A TIAA Company