


## Efficient AI: solutions for a resource-intensive era

Where the environmental case and investment case have converged for products and services that make AI more energy-efficient



## Lead contributors



Paolo Macri  
Head of Thematic Research



Luciano Lilloy  
Portfolio Manager



Barbara Heap  
Global Head of Portfolio Specialists

## Contents

Introduction	1
Part I - Where the AI build-out meets resource constraints	2
Part II - Where efficiency gains are delivered across the AI value chain	8
Implications for asset allocators	22
Important information	23

## Introduction

Artificial intelligence (AI) is the most capital-intensive technology shift of the past two decades. Its physical footprint and resource intensity are material at local, national and even global levels, and its scale means it cannot be overlooked by investors and capital allocators.

The implications are far-reaching and still evolving, creating both opportunities and risks. **In this paper, we focus on why only a portion of the AI build-out belongs within our Environmental Markets taxonomy and how we determine investment eligibility.**

Our observation is that, for AI and environmental markets (investing in products and services that deliver environmental solutions or improve resource efficiency), the environmental case and the investment case have converged.<sup>1</sup> Two factors support this.

**First**, the financial incentive, even imperative, to reduce the resource input cost per unit of compute is compelling. Every query run on a large language model (LLM), every new data centre and every training run adds to power and water demand. Efficiency, measured by performance per watt at the chip level, power usage effectiveness (PUE) at the facility level, and water and carbon intensity at the grid level, is now the dominant determinant of unit economics for AI workloads. Metrics for how efficiently inputs translate into intelligence, such as tokens per watt, are emerging, too.<sup>2</sup> Those companies that can improve efficiency (and lower costs for their customers) are likely to capture demand; companies that fail to improve may lose ground.

- 1 Environmental markets are the parts of the global economy where companies generate revenue from products and services that address environmental challenges – ranging from clean energy and water systems to resource efficiency and sustainable food
- 2 Given that metrics for how efficiently resources are translated into AI outputs remain relatively nascent, they are excluded from our analysis and are outside the scope of this paper. We expect to capture this important aspect of AI efficiency in the future.

**Second**, bottlenecks are emerging across the value chain. Early AI build-out was chip-constrained; the current stage is increasingly power- and grid-constrained. The IEA estimates that one in five planned data centre projects is at risk of delay, with new transmission lines taking four to eight years to build in advanced economies. The investable opportunity sits at each binding constraint, particularly among companies that improve efficiency and flatten the power demand curve.

This alignment does not extend across the whole supply chain. Only a portion stands up to scrutiny. Efficient AI earns its place in our taxonomy not because of sector allocation, but because specific products and services enable measurable efficiency gains.

**Part I** of this paper outlines how rising compute demand is reshaping pressure on energy and water resources, and how those pressures bring the AI build-out into the environmental markets frame.

**Part II** sets out how we classify efficiency-enhancing solutions across five clusters of the AI value chain (from design and manufacturing equipment, through silicon and memory, to interconnect, thermal management, data centres and cloud) and how we determine qualifying revenues for company eligibility.

AI carries both material risks and opportunities for capital allocation. This paper aims to help asset allocators looking for a disciplined way to translate booming AI capital expenditure (capex) into portfolio decisions consistent with environmental markets investment.



**Lisa Beauvilain**

Global Head of Sustainability & Stewardship, Co-Head of Sustainability Centre





**Part I**  
Where the AI  
build-out meets  
resource  
constraints

## The AI build-out is environmentally consequential because of its sheer scale and density.

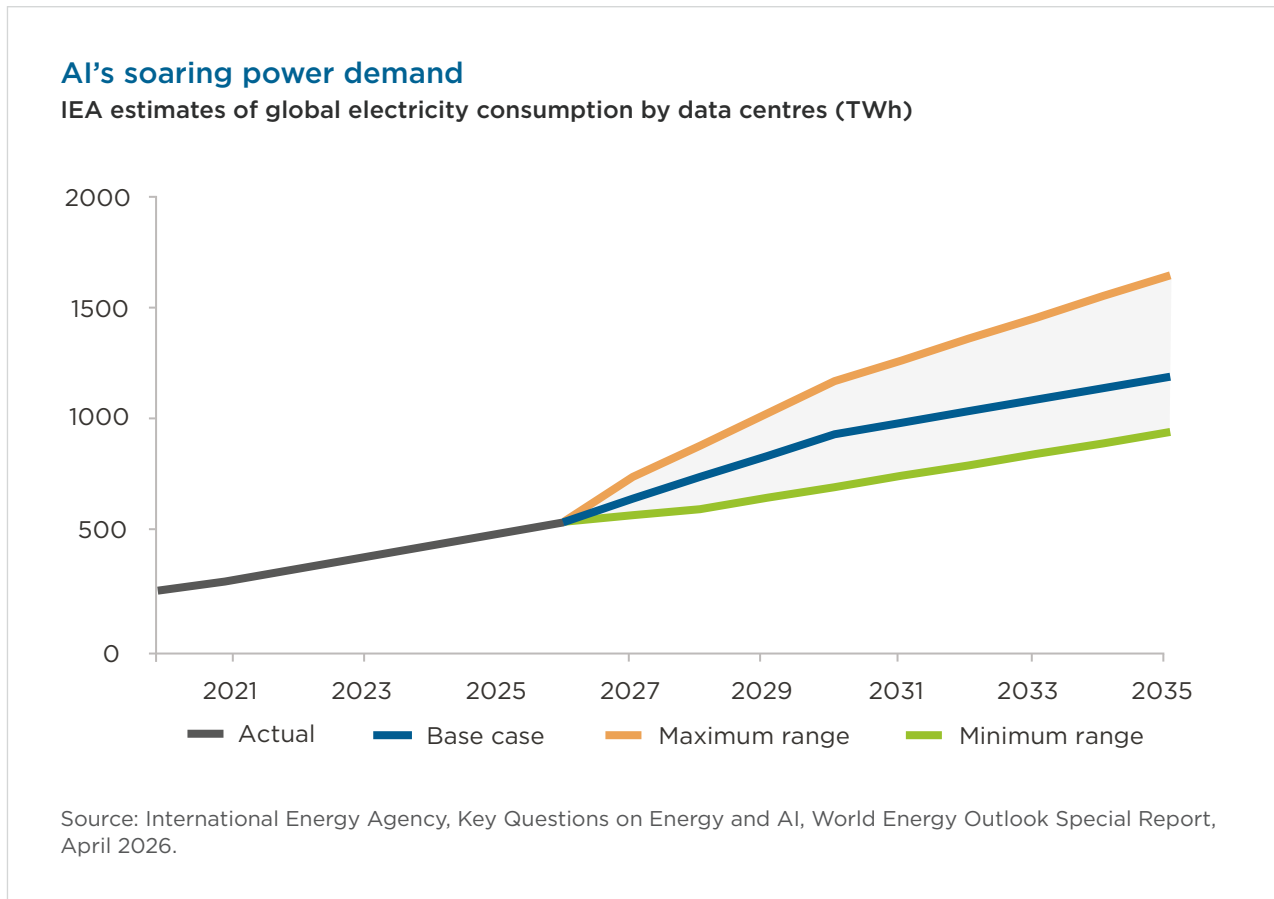
Energy is the core constraint: the electricity an AI campus draws is the master variable from which most of other resource-related pressures follow. Water use is largely a function of energy consumption and land for data centres is selected, first and foremost, for power availability.

Together, these interrelated pressures help define the macroeconomic context for the company-level analysis that we detail in Part II of this paper.

## Energy and the grid

Data centres accounted for around 1.5% of global electricity consumption in 2024.<sup>3</sup> The IEA projects that share will roughly double to about 3% by 2030, driven largely by AI (see chart below). Data centres already account for around 7% of electricity consumption in the US.<sup>4</sup>

This makes compute one of the fastest-growing new sources of energy demand in the world, and one of the few sectors where greenhouse gas (GHG) emissions are set to rise.



3 IEA, 2025: Energy and AI

4 International Data Center Authority, May 2026



Global and even national figures understates the strain on local grids. An AI-focused hyperscale facility – which are generally clustered in a handful of markets – can consume as much electricity as 100,000 homes.<sup>5</sup>

Hyperscale campuses, which now routinely occupy up to 300 acres (roughly 1km<sup>2</sup>), are clustered in a small number of hubs including Northern Virginia and Ireland, where data centres already account for more than 20% of electricity consumption.<sup>6</sup> Indeed, local grid pressures led to the imposition of a moratorium on new data centres in the Dublin area.


Not only do grid constraints threaten to undermine the expansion of AI there are also associated social and environmental implications. The power needs of massive complexes are driving up electricity bills. In areas of significant data centres activity, electricity costs have risen as much as 267% over five years.<sup>7</sup> Electricity constraints are also leading to the use of more fossil fuels. Data centre developers secure power by connecting to the grid, building out their own power supply (often renewable supply) and through multiple backup generators which typically are run on diesel fuel.<sup>8</sup>

The assumption that capacity can be sited wherever it is most economically efficient no longer holds in the markets where demand is most concentrated. AI infrastructure is now competing for grid capacity – which is finite in the short term – with industry, transport and heating, all of which are electrifying. Interconnection queues at major US grid operators now run into several years and many European operators have begun deferring new industrial connections altogether.<sup>9</sup>

The bottleneck in the AI build-out has migrated from chips to power and, increasingly, to the grid.

### The bottlenecks delaying AI rollout

Estimated new order deliveries (years)



Sources: IEA analysis based on data from Bloomberg (2026), IDC (2025), IEA (2025a), Refinitiv (2025), Reuters (2025), and Wood Mackenzie (2025). Based on 2025 as Year 0.

5 IEA, April 2025: Energy and AI

6 IEA, 2024

7 Bloomberg, 2025: How AI Data Centers Are Sending Your Power Bill Soaring

8 Martin, A., 12 November 2025: Data Centers' Use of Diesel Generators for Backup Power Is Commonplace—and Problematic. *Inside Climate News*

9 E3G, March 2026: Industrial electrification in the EU – Blocked by the grid?

## Water

The industry's thirst attracts less attention than its energy intensity, but is comparable in scale. Water and energy are, in effect, the same problem viewed from two angles.


A 100MW hyperscale facility can directly consume roughly 2.5bn litres a year for onsite cooling, equivalent to the needs of about 80,000 people.<sup>10</sup> It is important to understand though that a data centre's largest water footprint is usually indirect: the electricity it consumes is often generated, in large part, by thermoelectric plants that withdraw and evaporate substantial volumes of water.

That linkage creates a tension that the industry is only beginning to confront. Data centre operators site facilities where power is cheap and abundant, but those locations – often sun-rich regions largely powered by solar generation or regions where grids have spare thermal capacity – are frequently water-scarce. Optimising a campus for energy can therefore worsen its water exposure: hyperscale clusters in Arizona, Texas, Spain and Singapore sit in catchments that are already stressed, putting them in direct competition with agriculture and municipal consumers.

Adoption of water-efficient solutions can break the link between cooling and local freshwater withdrawal, and enhance operational resilience. Closed-loop and recycled-water cooling, the use of non-potable or reclaimed sources, and the shift to liquid cooling that recirculates (rather than evaporates), can all change the picture materially.

### AI growth is widely exposed to water scarcity

Average water stress index under a SSP3-7.0 scenario in the largest data centre markets



Source: Data as of August 2025. SSP3-7.0 refers to the Intergovernmental Panel on Climate Change's shared socioeconomic pathways (SSPs). SSP3-7.0 is a medium-to-high emissions scenario. The water stress metric is reported using a 0-100 range, and it describes the state of water availability (calculated on a decadal average) for the local water basin. We define high water stress as a value of 40 or higher, consistent with our previous published research (see Related Research), and the Aqueduct Water Risk Atlas and World Resource Institute (WRI) Low to medium water stress are values below 40.


<sup>10</sup> Impax analysis, based on figures from World Economic Forum, November 2024: Why circular water solutions are key to sustainable data centres



Two facilities with the same nominal water draw can impose very different burdens on a local catchment depending on how that water is sourced and returned. The efficiency question, in other words, is not only how much water is used, but how it is managed.

The same data centre campus that strains a regional grid also draws on a regional water table, competing for resources with other local energy and water consumers.

This dynamic means cost and security of supply are driving the adoption of resource-efficient solutions across the industry. We outline where in the value chain those responses translate into investable, demonstrable efficiency gains in Part II.

A photograph of a server room with rows of server racks. The scene is bathed in a blue light. A large, dark blue triangle is overlaid on the right side of the image, containing white text. The text discusses water-efficient solutions in the context of cooling and freshwater withdrawal.

Adoption of water-efficient solutions can break the link between cooling and local freshwater withdrawal, and enhance operational resilience.

A photograph of modern glass skyscrapers under a clear sky. A large, dark blue geometric shape, resembling a stylized arrow or a large 'V', is overlaid on the left side of the image. The text is positioned within this blue shape.

## **Part II**

Where efficiency gains are delivered across the AI value chain

Within the rapidly-expanding AI infrastructure universe, we are focused on companies that deliver efficiency gains – both measurably and materially.

Our framework is designed to identify these companies in a way that can be evidenced and quantified.

**The central question is: what percentage of a company’s revenue comes from products or services that demonstrably improve energy or resource efficiency above the market baseline?**

This metric is calculated differently depending on where in the AI value chain a company sits, but the principle is always the same: if we can evidence the efficiency mechanism and trace it to a specific revenue line, it counts.

Single-company narratives can be misleading unless we assess the full, interconnected AI value chain. Step-changes in efficiency only materialise when multiple components advance together (compute, memory, interconnect, power delivery and cooling). If any one of these lags, the expected system-level efficiency gains may not be realised.

**Our Environmental Markets taxonomy maps the full AI value chain into five clusters.** For each, it defines the environmental rationale, the counterfactual baseline and the eligibility rule for calculating qualifying revenues.

Using this framework, we aim to consistently capture the environmental markets relevance of company activities. Our goal is a repeatable, transparent framework that:

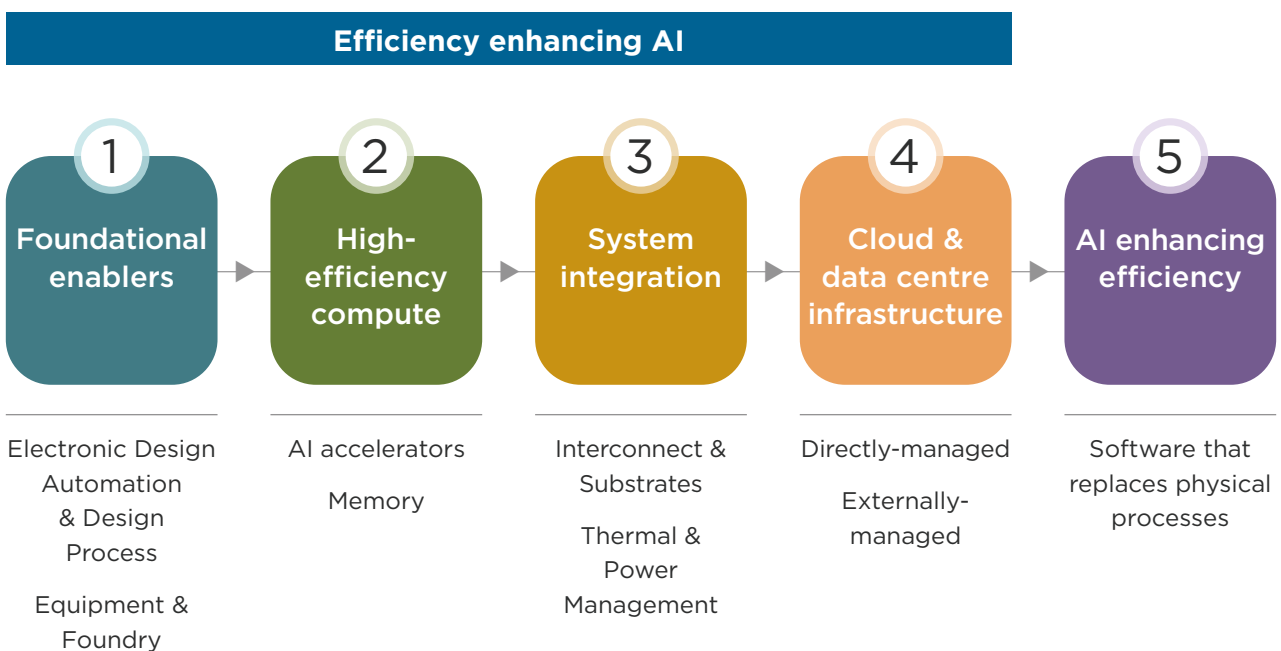
- Aligns with our Environmental Markets taxonomy
- Distinguishes between genuine enabling environmental contributions and ‘baseline technology exposure’
- Supports consistent eligibility and revenue attribution decisions across semiconductors, data centres, software and enabling AI infrastructure

## The five clusters

For every segment within each cluster, we detail the metric we use to quantify a company’s contribution to solving environmental challenges arising from AI adoption, and identify examples of companies leading the development of relevant products and services.

### Our taxonomy maps the AI value chain into five clusters

Within each cluster, only revenue from solutions that beat the prevailing market baseline qualifies



The background of the slide is a close-up, high-resolution image of a microchip. The chip's surface is covered in a complex grid of tiny, colorful components, including various colored squares (blue, green, yellow, red) and intricate patterns of circuitry. A large, semi-transparent purple arrow points from the left side of the frame towards the right, partially overlapping the microchip image. The text is positioned on the left side of the slide, within the purple arrow's path.

# **Cluster 1:** Foundational enablers

## In this first cluster of environmental solutions are the tools and processes behind device efficiency.

In modern semiconductors, transistors leak small amounts of current even when they are switched off. Given there can be billions of transistors on a single chip, this leakage accumulates into a significant energy drain, generating heat and consuming power.

Low-power chip design reduces energy leakage by as much as 50% in applications.<sup>11</sup> Advanced-node scaling (sub-5 nanometre (nm) line width) means more transistors on chips and faster computing. Chips can use less energy per calculation when electrons have shorter distances to travel, delivering material lifetime operational power savings that can outweigh their manufacturing footprint.

We divide this cluster into two segments:

**Electronic Design Automation (EDA) & Design Process** and **Equipment & Foundry**. Note that the efficiency gains from EDA software and advanced manufacturing equipment are captured downstream, in the second and fourth clusters, rather than here. This is because their processes are interlinked: you cannot build a sub-5nm chip without extreme ultraviolet lithography (EUV) lithography and you cannot optimise its power profile without power-aware EDA flows.

**1 EDA & Design Process.** The energy efficiency of a chip can be calculated before fabrication using advanced design tools to simulate and analyse their power consumption. EDA software tools are used by engineers to identify where energy is wasted and to adjust their design accordingly. They allow techniques such as clock-gating and power-gating – which switch off power until it is needed – and dynamic voltage-frequency scaling – which improves efficiency by allowing chips to adjust both voltage and clock speed based on workflow demand – to be built into chips at the design stage.<sup>12</sup>

**Metric:** Percentage of total company revenues attributed to the design of leading-edge semiconductors.

- **Company examples:** Among leading EDA software tools are PrimePower and Voltus, developed respectively by **Synopsys** and **Cadence**. The US-listed companies provide a range of software tools that simulate and validate electronic designs, helping modern AI chips to only use power when they are doing useful work.

**2 Equipment & Foundry.** Advanced processes such as EUV lithography – the process of chemically etching circuits onto chips – and cutting-edge packaging – the process of optimising chip configuration to improve performance and maximise space – are essential for next-generation semiconductors. The most advanced node chips (where line width is narrowest – and as low as 3nm) are significantly more energy efficient than larger nodes. It is therefore production of these leading-edge 3nm and 5nm-node chips that we deem to be contributing to making AI more efficient. Process equipment companies with broad exposure across all node generations do not therefore qualify for inclusion in our taxonomy on general revenues alone.


**Metric:** Percentage of total company revenues attributed to the manufacture of leading-edge semiconductors.

- **Company examples:** The EUV lithography machines developed by Dutch-listed **ASML** are the only equipment in the world that can etch circuits onto the smallest, sub-5nm chips. Correspondingly, **Taiwan Semiconductor Manufacturing Company (TSMC)** leverages that capability into high-volume chip production at the 3nm and 5nm nodes, production of which it dominates globally.<sup>13</sup>

<sup>11</sup> Cadence, 2024

<sup>12</sup> Without these techniques, the same transistor count would consume substantially more power: clock gating alone can reduce dynamic switching power by 20% to 30%, while dynamic voltage-frequency scaling can cut power by up to 40%, and power gating eliminates leakage in idle blocks. Across a realistic workload, these design-stage choices can reduce a chip's energy consumption by tens of percent relative to an equivalent design without them.

<sup>13</sup> International Trade Administration, December 2025



**Cluster 2:**  
High-efficiency  
compute

## In this second cluster of environmental solutions are the chips that enhance how much useful AI compute is delivered per watt of power consumed.

A standard computer processor (CPU) is built for flexibility. It needs to manage multiple tasks and so dedicates most of its power and circuitry to anticipating and managing unpredictable instructions. That sophistication is essential for general computing but is largely wasted on AI workloads, which are fundamentally different.

Training or running an AI model is essentially repeating the same mathematical operation billions of times in a predictable pattern. There is nothing to anticipate and nothing unpredictable to manage. Hardware that is better suited to AI-related processes – or indeed designed specifically to enable them – can yield significant gains in energy efficiency (and so financial savings for users).

We divide this cluster into two segments:

### **AI accelerators** and **Memory**.

**1 AI accelerators.** Purpose-built AI chips have a fundamental architectural advantage. Rather than being designed for flexible control logic, they can process millions of simple arithmetic operations in parallel workflow.

The result is a dramatic improvement in both performance and energy efficiency per unit of computation. These ‘AI accelerators’, such as Nvidia’s graphics processing units (GPUs) or Google’s Tensor Processing Units (TPUs), are one to three orders of magnitude more efficient per AI operation than a general-purpose server processor.<sup>14</sup>

The transformational scale of these advances means that standard performance benchmarks, such as clock speed, core count, or performance ratings, do not capture what actually matters for AI efficiency.

**Metric:** Our proprietary **AI Accelerator Efficiency Impax Score (AEIS)** ranks AI chips by how efficiently they use electricity. It measures one thing: how much AI computation a chip delivers per watt of energy consumed. It does this separately for training and inference – the two main types of AI workload – and combines them equally, since both matter in practice. To qualify for consideration, a chip must beat the current market average, not just outperform older hardware. As newer, more efficient chips enter the market and raise the average, the bar for eligibility rises automatically.

- **Company example: Nvidia** dominates the global market for the GPUs used in AI data centres.<sup>15</sup> The US-listed company’s chips for hyperscale data centres (which constitute approximately 40% of its data centre business) are eligible for inclusion as their AEIS score is superior to peers.<sup>16</sup>

**2 Memory.** High bandwidth memory (HBM) reduces the energy cost of moving data to a processor by as much as five times as much as dynamic random-access memory (DRAM).<sup>17</sup> HBM is essential for GPUs used in AI training and inference.

Conventional DRAM sits on a separate chip some distance away from the processes and energy is required to move data between the two. In contrast, HBM uses 3D-stacking where it sits on top of, or adjacent to, the processes and connects through vertical conductors called through-silicon-via (TSV) architecture. This means data arrives faster and travels a shorter distance.

**Metric:** Percentage of total company revenues attributed to the production of HBM.

**Company example: SK Hynix** is the largest producer of HBM, with a global market share of almost 60%.<sup>18</sup> The South Korean company is a leading supplier to Nvidia.

14 Newkirk, A.C. et al., 2025: Empirically-Calibrated H100 Node Power Models for Reducing Uncertainty in AI Training Energy Estimation. *Lawrence Berkeley National Laboratory and Carnegie Mellon University*

15 Pan, C. & Chen, L., 1 April 2026: Chinese chipmakers claim nearly half of local market as Nvidia’s lead shrinks. *Reuters*

16 Impax analysis, May 2026

17 Wevolver, 2025: HBM Memory: Complete Engineering Guide & Design Optimization 2025

18 Counterpoint, June 2026: Global DRAM and HBM Market Share: Quarterly

A photograph of a server room with a large green arrow pointing from the left towards a server rack. The server rack is filled with equipment and has several green lights illuminated. The room has a white floor and a ceiling with exposed pipes and lights.

# Cluster 3: System integration

In this third cluster of environmental solutions are companies whose products directly reduce resources wasted per unit of compute delivered.

An AI server rack – a cabinet holding multiple servers in a data centre – typically has densities up to 20 times that of conventional racks. These higher-density racks require significant amounts of power and can generate lots of heat, representing a material source of energy waste. Enhancing the energy efficiency of AI server racks and lowering heat losses can deliver meaningful financial savings for data centre operators, as well as lowering GHG emissions.

We divide this cluster into two segments: **Interconnect & Substrates** and **Thermal & Power Management**.

**1 Interconnect & Substrates.** In high-power applications like data centres, heat build-up at connectors is a major source of energy inefficiency. Three types of advanced components address this issue. First, high-performance connectors lower resistive losses at extreme current loads. Second, ceramic substrates used on advanced chips conduct heat away from components roughly 600 times more effectively than conventional printed circuit board material.<sup>19</sup> Third, coherent optical interconnect and active electrical cables cut power consumption in data transmission at campus and intra-rack scale respectively.

**Metric:** Data centre / AI exposure. Products that enable denser buildout without measurable efficiency gain are excluded.

- **Company example: Amphenol** is a global leader in connectivity solutions, supplying components for applications across the automotive, industrial, data centre, consumer

electronics, aerospace and medical sectors. The US-listed company designs connectors with larger contact surface areas and that employ materials that dissipate heat more effectively, reducing thermal losses and preventing asset degradation over time.

**2 Thermal & Power Management.** Advanced systems are enhancing how much electricity actually reaches the IT load, rather than being lost as heat or conversion waste.

Liquid cooling systems enable facility power utilisation efficiency PUE to approach the range of 1.1 to 1.2 – a significant improvement on the baseline of 1.5 to 1.6 for air-cooled servers.<sup>20</sup> Meanwhile, wide-bandgap power semiconductors – chips that control, convert and regulate electrical power flowing through a system – achieve switching efficiencies above 97%. Finally, power management integrated circuits (PMICs) minimise waste in the final stage of server-level power delivery.

**Metric:** Data centre / AI exposure. Products that enable denser buildout without measurable efficiency gain are excluded.

- **Company example: Infineon** makes power semiconductors. It is also a leading producer of silicon carbide (SiC) and gallium nitride (GaN) chips, which can switch faster, handle higher voltages and generate less heat than traditional silicon chips. In data centres, where power efficiency is critical and energy costs are high, the German-listed company's products address improve power conversion efficiency and reduce waste.

<sup>19</sup> Aluminium nitride substrates (AlN) conduct heat away roughly 600 times more effectively than FR4. Source: NextPCB, 2026

<sup>20</sup> Schneider Electric, 2026. A PUE measure of 1.5, for example, means for every 100 watts of energy used in computations, 50 watts are consumer by cooling, power conversion, lighting and other overheads.

# Cluster 4: Cloud & data centre infrastructure

```
... Arcsinh, Arcsin, Denint, Beta, BetaNoncentral, BetaPrime,
... erfi, besseli, besseli)
...
... functions import erfinv
... import melferg
... intersection, FiniteSet
... where, density, variance, covariance, skewness, kurtosis, median,
... en, pspace, cdf, characteristic_function, moment_generating_function,
... ContinuousRV, sample, Arcsinh, Arcsin, Denint, Beta, BetaNoncentral, BetaPrime,
... ContinuousRV, Arcsin, Benini, Beta, BetaNoncentral, BetaPrime,
... Cauchy, Chi, ChiSquared, ChiNoncentral, Dagum, Erlang, Exponential,
... Exponential, ExponentialPower, FDistribution, Fisher2, Frechet, Gamma,
... GammaInverse, Gompertz, Gumbel, Kumaraswamy, Laplace, Levy, Logistic, LogCauchy,
...
... stats.crv import SingleContinuousSpace, SingleContinuousDomain
... stats.compound_rv import CompoundSpace
... j-stats.symbolic_probability import Probability
... py.testing.pytest import raises, XFAIL, slow, skip, ignore_warnings
... py.testing.pytest import raises, XFAIL, slow, ignore_warnings
... py.testing.randtest import verify_numerically as tn
...
... = S.Infinity
```

## In this fourth cluster of environmental solutions are operators and owners of the most efficient data centres that enable AI.

The extent to which cloud computing is less environmentally impactful than on-premise alternatives hinges on data centres' efficiency (both energy and water) and how they are powered.

We divide this cluster into two broad segments: **directly-managed facilities** and **externally-managed data centre real estate investment trusts** (REITs).

**1 Directly-managed facilities.** Entities that operate their own data centres can directly determine facilities' PUE, water usage effectiveness (WUE), hourly carbon-free energy (CFE) and renewable energy procurement.

Within this segment, we differentiate between hyperscalers and colocation operators.

**(i) Hyperscalers.** The efficiency case for hyperscale cloud operators rests on two distinct, but highly material factors: facility engineering and energy procurement. PUE is a metric for the former. Leading hyperscalers can achieve a PUE of approximately 1.1 to 1.2, compared to an industry average of 1.5 to 1.6.<sup>21</sup> This can represent hundreds of megawatts of continuous electricity draw at the largest campuses, depending on facility design. They achieve this through purpose-built campuses with custom cooling architectures, high-efficiency power distribution, waste heat recovery, and the ability to co-locate facilities in geographies with access to natural cooling. Additionally, hyperscalers are the largest direct buyers of new clean energy through corporate power purchase agreements (PPAs): over the past five years, about two-thirds of renewable PPAs signed in the US have been dedicated to data centre operations.<sup>22</sup>

**Metric:** Relative ranking, versus peers, based on our **proprietary Data Centre Efficiency Impax Score (DCEIS)** – see page 18. For hyperscalers, we only consider revenues generated from cloud and compute services to be eligible; advertising, e-commerce, and productivity software revenues are excluded.

- **Company example: Microsoft** currently boasts around one-fifth of the global market for cloud services – second only to Amazon.<sup>23</sup> Its Azure cloud platform generates around 28% of company revenues and is eligible for consideration as the US company's DCEIS score is above median, versus peers.<sup>24</sup> Oracle, another hyperscaler, is not eligible, for example, as its DCEIS is below that of peers.

**(ii) Colocation operators.** A colocation data centre is a shared facility where companies rent out rack space. The relevant comparison is therefore not the hyperscalers, but private server rooms or data centres that customers would otherwise operate. These private facilities are typically underutilised and lack the economies of scale in cooling and power infrastructure that purpose-built, professionally-managed colocation data centres can offer. Moving data to colocation centres that are shared by other customers provides consolidation benefits and gives access to specialists operator expertise.

**Metric: Data Centre Efficiency Impax Score (DCEIS),** (see page 18), benchmarked against private operators and peers rather than hyperscalers. The per-facility industry average PUE sits at roughly 1.6, weighed down by smaller and older sites, while purpose-built modern colocation is designed for 1.4 or below. That gap is the efficiency improvement the DCEIS is built to capture.<sup>25</sup>

**Company example: Equinix** runs a global network of roughly 260 data centres in more than 70 countries. Based on its DCEIS, all revenues are deemed to be eligible and the company is above peer average on DCEIS.

21 National Laboratory of the Rockies (formerly the National Renewable Energy Laboratory), 2026

22 IEA, 2026: Key Questions on Energy and AI

23 Synergy Research Group, November 2025

24 Impax analysis, June 2026. Azure figure is based on Microsoft figures for its 2025 financial year.

25 Uptime Institute, 2024



2

**Externally-managed data centre REITs.** Several pure-play REITs own data centre but delegate day-to-day facility operations to third-party managers. These REITs therefore do not directly control PUE, WUE, hourly CFE or renewable procurement.

**Metric:** Assets under management (AUM) certified under a Green Building Certification Standard (GBCS), namely LEED, BREEAM, BCA Green Mark or NABERS.<sup>26</sup> We look for companies to have at least 20% of AUM certified. Note that DCEIS is structurally inapplicable as a metric, given externally-managed data centre REITs do not control efficiency.


**Company example: Keppel Data Centre REIT** is a pure play, carrier-neutral data centre REIT based in Singapore that is benefitting from growth in data volume demand. The majority of its data centres have either LEED or BCA Green Mark certificates.

### Data Centre Efficiency Impax Score (DCEIS)


Metrics used to evaluate data centre sustainability performance

Metric	Description	Why it matters
<b>PUE</b> Power Usage Effectiveness	Measures how much total electricity entering a data centre reaches IT equipment, versus being lost to cooling, fans, power conversion and distribution.	The primary measure of operational energy efficiency. Directly reflects the carbon and cost impact of the physical infrastructure.
<b>WUE</b> Water Usage Effectiveness	Measures how many litres of water are consumed per kilowatt-hour of IT energy delivered.	Evaporative cooling in large data centres can consume enormous quantities of water, potentially imposing significant water costs on local communities, especially in drought-prone or water-scarce regions.
<b>REF</b> Renewable Energy Factor	The percentage of total electricity consumption covered by renewable energy procurement — including power purchase agreements, renewable energy certificates or direct ownership — over the course of a year.	Captures an operator’s commitment to decarbonising its energy supply.
<b>24x7 CFE</b> Hourly Clean Energy Matching	Like REF but, rather than matching energy on an annual basis, hourly matching requires clean energy to be available at every hour of operation.	Eliminates the paper accounting loophole of buying summer solar to offset winter coal. Rewards operators making real contributions to grid decarbonisation.
<b>Green Certification</b> DC Green Building Certification	The percentage of data centre floor space or capacity certified under green building standards such as LEED, BREEAM, Green Mark or ISO 14001/50001.	Signals asset-level design discipline, embodied carbon controls, commissioning rigour, and third-party verification — particularly meaningful for colocation operators.

26 LEED = Leadership in Energy and Environmental Design; BREEAM = Building Research Establishment Environmental Assessment Method; BCA = Building and Construction Authority; NABERS = National Australian Built Environment Rating System

A hand is pointing at a screen displaying code. A large green arrow points from the left towards the text. The background is a blurred image of a server room with blue and green lighting.

Our proprietary  
Data Centre  
Efficiency Index  
Score (DCEIS) is  
used to evaluate data  
centre sustainability  
performance



**Cluster 5:**  
AI enhancing  
efficiency

Unlike the first four clusters of environmental solutions, which are all focused on enhancing AI process efficiency, this fifth cluster is comprised of providers of software that can use AI to enhance efficiencies and improve environmental outcomes.

Here, we are looking for companies whose software demonstrably reduces the volume of physical activity required to achieve the same outcome.

The clearest examples of such solutions lie in three areas.

First, **building design software** can reduce costly rework and material waste by resolving problems digitally before construction begins.

Second, **product lifecycle management tools** can reduce the need for physical prototypes by simulating them instead.

Third, **logistics optimisation software** can find more efficient routes, directly cutting fuel consumption and GHG emissions.

In each case the environmental benefit comes not from doing the physical activity more cleanly, but from requiring less of it.

**Metric:** Product-level revenue alignment, based on a specific, measurable reduction in physical-world resource consumption.

- **Company example:** Industrial software provider **PTC** uses AI in its generative design tools across the full industrial lifecycle, including to optimise product designs for material use and energy consumption before manufacturing begins, predict equipment failures to reduce downtime and extend asset life, and analyse real-time factory data to identify inefficiencies in energy use and throughput.



**Building design software can reduce costly rework and material waste by resolving problems digitally before construction begins.**



# Implications for asset allocators

The AI build-out is generational in scale and uneven in environmental quality. The taxonomy set out in this paper does the work of separating the two: within each cluster, only revenues generated by solutions that demonstrably advance efficiency above the prevailing market baseline count towards a company's eligibility for our Environmental Markets strategies.

**We believe three implications follow for asset allocators.**

**First**, efficiency is now the dominant determinant of unit economics for AI workloads. Resource and energy costs are a large and rising share of the cost of compute. Companies that minimise resource intensity and optimise efficiency can lower costs for their customers. This confers a competitive advantage that allows them to capture and defend demand. The environmental case and the investment case converge here: the resource efficiency our taxonomy rewards is the same characteristic that underpins a durable competitive moat.

**Second**, the binding constraints of the AI build-out are where opportunities concentrate. As the bottleneck has migrated from chips to access to power and water, we expect value to accrue to those companies positioned across energy, water management, thermal management and grid-adjacent infrastructure, where immediate supply is scarce and slow to replicate. Allocators that index to compute demand alone risk overpaying for the layer where supply responds most readily to demand signals.

**Third**, our framework is dynamic, not static. Both our proprietary AEIS and DCEIS scores are relative, not absolute metrics; as technologies advance, the goalposts for inclusion move. Companies that fail to keep improving lose eligibility, so identifying the leaders at any given point demands continuous reassessment rather than a one-off screen.

We acknowledge that efficiency gains may lower the cost of compute and so invite more of it, keeping aggregate demand on its upward path. That is the trend, not a flaw in the approach: our framework rewards relative efficiency, not absolute restraint, and falling resource intensity per unit of output is the lever actually available to investors.

Maintaining exposure to the AI theme through the lens of environmental markets is an active discipline. It is one that we are committed to continue devoting resources to.



## Important information

**This material is intended for professional investors only. It provides general information only and does not constitute investment advice or a recommendation.**

This material is being provided for informational purposes only. The material is not intended for use by persons in jurisdictions which prohibit or restrict the distribution of the material and in certain countries, the material is provided upon specific request. Under no circumstances should the material, in whole or in part, be copied or redistributed without consent from Impax. The material does not constitute a distribution, an offer, an invitation, recommendation or solicitation to sell or buy any securities in any jurisdiction.

Impax Asset Management Group plc includes Impax Asset Management Ltd, Impax Asset Management (AIFM) Ltd., Impax Asset Management Ireland Ltd, Impax Asset Management LLC, Impax Asset Management GmbH, and Impax Asset Management (Hong Kong) Limited (together, "Impax").

Impax Asset Management Ltd, Impax Asset Management (AIFM) Ltd and Impax Asset Management LLC are registered as investment advisers with the U.S. Securities and Exchange Commission ("SEC"), pursuant to the Investment Advisers Act of 1940 ("Advisers Act"). Registration with the SEC does not imply a certain level of skill or training.

Impax Asset Management Ltd & Impax Asset Management (AIFM) are authorized and regulated by the Financial Conduct Authority. Impax Asset Management Ireland Ltd is authorised and regulated by the Central Bank of Ireland (Reference No: C186651). Impax Asset Management GmbH is authorised and regulated by the Federal Financial Supervisory Authority of Germany (BaFin).

In the United Kingdom, this material is a financial promotion and has been approved by Impax Asset Management Limited.

The statements and opinions expressed are those of the author, as of the date provided, and may not reflect current views. While the author has used reasonable efforts to obtain information from reliable sources, we make no representations or warranties as to the accuracy, reliability or

completeness of third-party information presented herein. Any forward-looking statements or forecasts are based on assumptions and actual results are expected to vary. References to specific securities are for illustrative purposes only and should not be considered as a recommendation to buy or sell. This material may not be relied upon as constituting any form of investment advice and prospective investors are advised to ensure that they obtain appropriate independent professional advice before making any investment decision.

### **Past performance does not guarantee future results.**

Nothing presented herein is intended to constitute investment advice and no investment decision should be made solely based on this information. Nothing presented should be construed as a recommendation to purchase or sell a particular type of security or follow any investment technique or strategy. Information presented herein reflects Impax Asset Management's views at a particular time. Such views are to change at any point and Impax Asset Management shall not be obligated to provide any notice. Any forward-looking statements or forecasts are based on assumptions and actual results are expected to vary. While Impax Asset Management has used reasonable efforts to obtain information from reliable sources, we make no representations or warranties as to the accuracy, reliability or completeness of thirdparty information presented herein. No guarantee of investment performance is being provided and no inference to the contrary should be made.

Impax is trademark of Impax Asset Management Group Plc. Impax is a registered trademark in the EU, US, Hong Kong, and Australia. © Impax Asset Management LLC, Impax Asset Management Limited and/or Impax Asset Management (Ireland) Limited. All rights reserved.Disclosures/Disclaimer




## **Efficient AI: solutions for a resource-intensive era**

How the environmental case and investment case have converged for products and services that make AI more energy-efficient

**IMPAX ASSET MANAGEMENT**

✉ [clientservices@impaxam.com](mailto:clientservices@impaxam.com)

 Impax Asset Management

[impaxam.com](https://www.impaxam.com)