PHYSICAL CLIMATE RISK · DATA CENTRES

The Cloud Has a Physical Address: Climate Exposure in the Infrastructure Behind AI

Scientific Climate Ratings
An EDHEC Venture

ASSETS ANALYSED
1,000

COUNTRIES
34

OPERATORS
11

SCENARIOS
SSP2-4.5 and SSP5-8.5

HORIZONS
2035 and 2050

September 2026

EXECUTIVE SUMMARY

Scientific Climate Ratings analysed 1,000 data centres across 34 countries and 11 major operators to assess how physical climate hazards could affect the value and resilience of the infrastructure underpinning the AI economy. The analysis covers wind, flood, wildfire and heat under SSP2-4.5 and SSP5-8.5 scenarios to 2035 and 2050, translating asset-level hazard exposure into expected structural damage and operational disruption.

The central finding is that physical climate risk is driven much more by where a data centre is located than by who operates it. The dispersion across countries is substantially greater than across operators. By 2050, the average discounted cost of inaction ranges from 3.4% of asset value in Sweden to 20.6% in Japan, while the operator-level range is much narrower, from 4.3% for Meta to 9.6% for NTT. Facilities belonging to the same operator can therefore sit at opposite ends of the physical-risk spectrum.

Among the hazards that can currently be quantified financially, wind is the dominant portfolio-wide risk. It affects all 1,000 assets and generates an average annual expected asset-value impact of 0.65%, corresponding to a discounted cost of inaction of 5.51% by 2050. Wind differs from the other hazards because its impact is broad rather than concentrated: almost every asset carries some structural wind damage, with a median of 0.58% of asset value under SSP5-8.5 at 2050.

Flood is far more concentrated but can be much more severe at individual sites. Only 185 of the 1,000 facilities register modelled flood damage, producing a relatively small portfolio-wide annual expected impact of 0.06%. Yet among those exposed assets, average flood damage is around 0.33%, and the worst-placed site reaches 8.92% structural damage. This produces a highly asymmetric risk profile: little or no impact across most of the portfolio, but potentially very significant losses at a small number of facilities.

By comparison, wildfire is a very small direct financial risk in this particular data-centre sample. Only 19 assets register modelled wildfire damage, with an average annual expected impact of around 0.001% of asset value and a 2050 cost of inaction of 0.008%. This reflects the fact that relatively few facilities in the sample are located where wildfire intensity is sufficient to translate into direct physical damage.

Heat is widespread but produces only a small directly monetised impact in the current framework. All 1,000 facilities are exposed to heat, but heat is treated through operational disruption rather than structural damage, resulting in an average annual expected impact of 0.007%.

Importantly, some of the risks most frequently associated with data centres, particularly water stress and drought, are assessed at indicator level but are not yet translated into financial impact because sufficiently validated damage functions have not been incorporated. The hazard ranking should therefore be interpreted within that methodological scope.

Geographically, Japan stands out as the most exposed major market in the sample. All 22 Japanese facilities are rated F or G, with an average annual expected impact of 2.78% of asset value and an average 2050 cost of inaction of 20.6%. The Netherlands follows with an average annual impact of 1.85%, while Ireland reaches 1.47%. By comparison, the US, home to 616 of the 1,000 facilities analysed, averages 0.60%.

The study also shows that exposure is not the same as vulnerability. NTT's TK9 Data Center in Kōtō, Tokyo provides a concrete example. Before its documented flood adaptation and resilience measures are taken into account, the facility has a modelled annual value impact of 11.31% under SSP5-8.5 at 2050. Incorporating the measures already in place reduces that impact to 4.26%, a 62% reduction. Flood impact falls from 9.51% to 2.46%, while wind is unchanged at 1.78%. The facility remains highly exposed and retains a G rating, but substantially less of that exposure converts into financial loss.

The cumulative financial effect is substantial. At TK9, the discounted cost of inaction to 2050 falls from 65.4% to 29.5% of asset value once the existing flood-resilience measures are reflected, a reduction of 35.9 percentage points. This illustrates why climate-risk assessment cannot stop at hazard maps: resilience investments can materially alter the financial consequences of exposure even when the underlying location cannot be changed.

The broader implication is that the AI boom is becoming an infrastructure investment cycle as much as a technology cycle. Data centres are capital-intensive, long-lived and effectively immovable once constructed. Physical climate exposure therefore has direct implications for site selection, capital expenditure, insurance and financing. The cloud may be digital, but the economics of AI are becoming increasingly physical.

We may think of our data as living in the cloud, but the infrastructure that stores and processes it is firmly rooted on the ground. They are location-bound and unmovable. They depend on vast quantities of reliable energy and water. They must be able to withstand an increasingly volatile range of extreme weather events. As a result, they are exposed to the same physical climate risks as other critical infrastructure assets. This means the next constraint on the digital economy is increasingly physical, not digital.

One point of scope is worth stating upfront: the hazards we can price financially are not the hazards that dominate the public conversation about data centres. Heat, water stress and drought drive much of that debate, but only heat carries a financial figure here, and a small one; water stress and drought are assessed at indicator level and are not priced at all. That gap is itself a finding, and we set out why in the scope note accompanying Figure 6.

Computing infrastructure has evolved through four distinct eras. In the initial enterprise server room era, companies operated modest server closets on-site. From the early 2000s, colocation providers built shared facilities where multiple companies rented rack space, spreading the cost of power and cooling. The cloud era, beginning with AWS in the mid-2000s, let companies rent computing capacity rather than physical racks, hosted inside data centres owned by a small number of providers. Since the early 2010s, the largest of those providers have built ‘hyperscale’ campuses, which are dedicated, self-owned, gigawatt-class sites built purely to run their own services. The launch of consumer generative AI tools from late 2022 triggered a further step-change, and a build-out that is still accelerating.

FIGURE 1

Four eras, and then a step-change

Computing infrastructure has evolved through four distinct eras. The launch of consumer generative AI tools from late 2022 triggered a further step-change, and a build-out that is still accelerating.

Pre-2000

Enterprise server rooms

Companies operated modest server closets on-site.

Early 2000s

Colocation

Shared facilities where multiple companies rented rack space, spreading the cost of power and cooling.

Mid-2000s

Cloud

Companies rented computing capacity rather than physical racks, hosted by a small number of providers.

Since early 2010s

Hyperscale

Dedicated, self-owned, gigawatt-class campuses built purely to run providers' own services.

Late 2022 →

Generative AI era

Training and inference workloads are substantially more computationally intensive, raising power draw and the heat that must be dissipated.

Source: Scientific Climate Ratings - The Cloud Has a Physical Address Report (September 2026)

As demand for AI computing has grown and diversified, so too have the physical demands placed on data centres. Training and inference workloads are substantially more computationally intensive than many previous cloud-era workloads, increasing power consumption and, critically, the amount of heat that must be dissipated. Higher-density computing therefore requires increasingly sophisticated and energy- and water-intensive cooling infrastructure, placing greater demands on the physical systems and resources that underpin each facility.

FIGURE 2

The asset and the network it depends on

Each facility is a chain of component systems inside a structural shell, and every link in that chain reaches beyond the fence line for power, water and connectivity.

OFF-SITE INFRASTRUCTURE

Beyond the operator's control

Electricity grid
Generation, transmission, delivery

Regional water supply
Basin availability, abstraction limits

Fibre routes
Terrestrial and subsea paths

STRUCTURAL SHELL

Must withstand extreme weather events

Power chain

Grid connection
▼
Switchgear and UPS
▼
On-site backup generation

Cooling system

Chillers and heat rejection
▼
Water loop
▼
Air handling in the hall

Server halls

Thousands of densely packed servers held within a narrow operating temperature

Data in and out

SITE BOUNDARY

NET EFFECT A single external shock to power, water, or connectivity can compromise a facility that remains structurally untouched.

OFF-SITE DEPENDENCIES

The wider electricity grid
Generation, transmission and delivery beyond the fence line. Outages, congestion and fuel constraints reach the power chain first.

A regional water supply
Basin-level availability for cooling. Drought, competing use and abstraction limits cap usable capacity.

The fibre routes carrying data in and out
Terrestrial and subsea paths. Cuts from natural hazards, construction or infrastructure failure isolate an intact site.

Source: Scientific Climate Ratings - The Cloud Has a Physical Address Report (September 2026)

Today, data centres are undeniably critical infrastructure assets on which society now depends as much as transport hubs and energy production facilities. Yet behind each of these assets sits a vulnerable network of component systems: a structural shell that must withstand extreme weather events, a power chain running from the grid connection to on-site backup generation; and a cooling system that keeps thousands of densely packed servers within a narrow operating temperature. Every part of this network is, in turn, dependent on off-site infrastructure: the wider electricity grid, a regional water supply, the fibre routes carrying data in and out. A single external shock to power, water, or connectivity can compromise a facility that remains structurally untouched. Their exposure to physical climate risks should therefore be treated as an issue of public interest.

While most of the debate about data centres and climate has focused on the associated emissions, far less attention has been paid to a different question, one that matters just as much: what physical risks are assets in a given location exposed to and how can these risks be mitigated?

Data centres are highly resource dependent and climate vulnerable

Data centres depend on continuous access to power, cooling, and connectivity. Even without the impacts of extreme weather and higher average temperatures, data centres are struggling to meet these operational resource demands. In Ireland, the national electricity grid is under unprecedented pressure: data centres now account for 23% of the country's metered electricity consumption, up from 5% in 2015[1]. Meanwhile in the USA, data centres consumed about 17bn gallons of water directly in 2023, projected to reach 38-73bn by 2028.[2]

FIGURE 3

The resource squeeze is already visible

Even without the impacts of extreme weather and higher average temperatures, data centres are struggling to meet their operational resource demands.

GRID CAPACITY

23%

of Ireland's metered electricity consumption is accounted for by data centres, up from 5% in 2015.

Source: CSO, 2025 data[1]

WATER TRAJECTORY

38-73bn

gallons of water consumed directly by US data centres by 2028, up from about 17bn in 2023. Excludes the water used to generate their electricity, which is an order of magnitude larger.

Source: LBNL, 2024[2]

SITING RISK

517

of 809 planned US data centres are in areas that experienced drought over the past year, against more than 60% of the country in drought that year.

Source: US Drought Monitor via Newsweek, 2026[3]

Chart: Scientific Climate Ratings - The Cloud Has a Physical Address Report (September 2026). Data as cited per panel.

A warming climate will only intensify resource scarcity. Heatwaves of unprecedented scale and severity are raising the cooling load of data centres and reducing cooling efficiency at a time when grids are already overstretched, while also impacting the productivity of workers who operate and maintain these centres. These cooling-intensive facilities are also inherently vulnerable to water availability constraints, yet 517 of the 809 data centres planned in the US are in areas that experienced drought over the past year - though more than 60% of the country was in drought during that period, so this is close to the base rate[3]. In addition, storms, high winds, and wildfires threaten the grid connections that the sites depend on, as well as causing extensive damage to the assets themselves. As such, data centre managers must make climate-informed decisions when planning and operating data centres to ensure minimum down-time and avoid asset abandonment.

Scientific Climate Ratings' analysis

To establish the extent of the physical climate risks faced by data centres, we assessed 1,000 individual assets spanning 34 countries and 11 operators. We selected these assets on the basis of operator relevance, data quality, and analytical coverage, rather than sampling the market at random.

FIGURE 4

Where the exposure sits

All 1,000 assets, coloured by Physical Exposure Rating and sized by modelled total value impact. Below, the same assets summarised by country. The United States dominates the sample at 616 assets and spans the full scale from A to G, while Japan, on 22 assets, is the only country where every asset rates F or G.

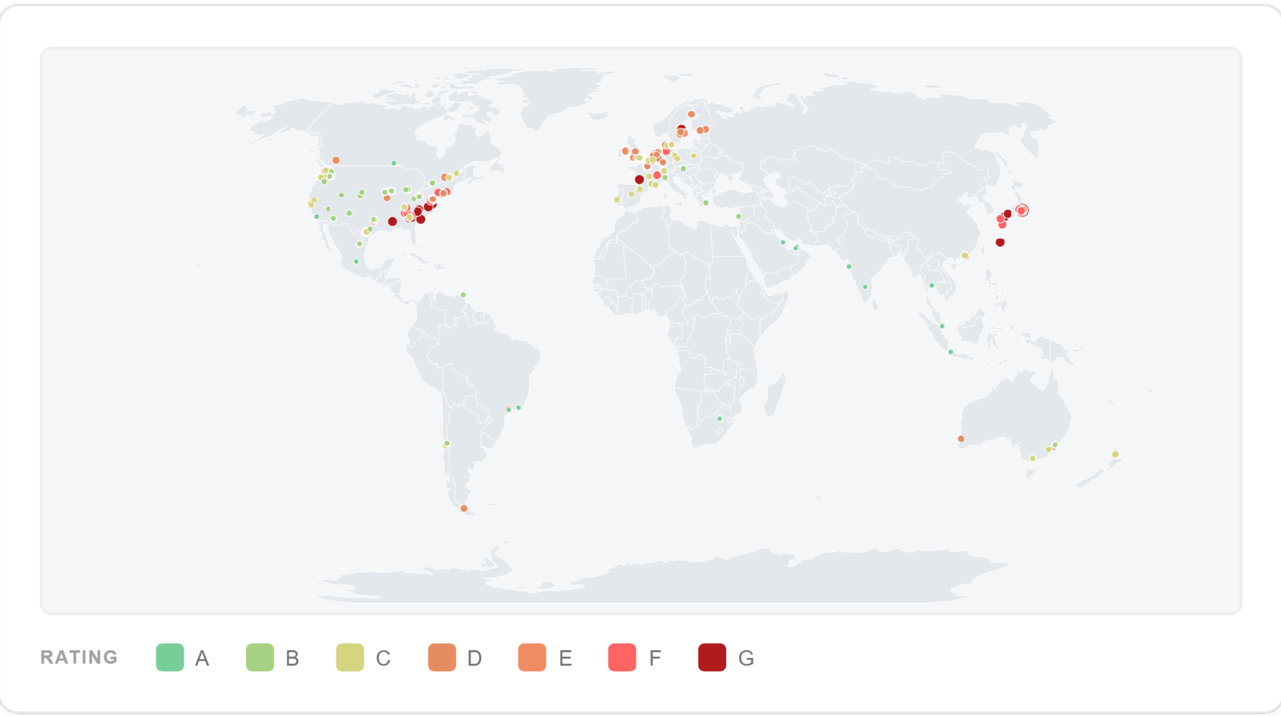

COUNTRY	ASSETS	AVERAGE RATING	RANGE	RATED F OR G	ANNUAL EXPECTED ASSET VALUE IMPACT
United States	616	C	A-G	1%	**0.60%**
Germany	63	D	C-D	0%	**0.80%**
Ireland	51	E	E-F	22%	**1.47%**
United Kingdom	39	C	C-D	0%	**0.81%**
France	32	D	C-F	6%	**0.86%**
Australia	30	C	B-C	0%	**0.43%**
Netherlands	26	F	D-G	69%	**1.85%**
Japan	22	F	F-G	100%	**2.78%**
Sweden	16	B	B-C	0%	**0.43%**
Canada	14	B	A-C	0%	**0.41%**

The ten countries holding the most assets, ordered by asset count; the remaining 24 countries are grouped. Average rating is the mean Physical Exposure Rating across the country's assets, expressed on the A-G scale. Range is the span of individual ratings present. Mean impact is the mean total value impact across that country's assets, as a percentage of asset value, SSP5-8.5 at 2050. The United States holds 616 of the 1,000 assets, so global averages are heavily US-weighted.

READING THE NUMBERS

An annual figure, not a one-off loss

These are annual expected values: the probability-weighted average loss in a typical year, as a share of tangible asset value. Not a single severe event, and not a forecast for one particular year.

A country mean near 1% therefore looks modest. But it recurs annually against assets built to run for decades and effectively impossible to relocate. It is the accumulation that matters to an owner or lender, and that is what the cost of inaction below measures.

Source: Scientific Climate Ratings - The Cloud Has a Physical Address Report (September 2026)

FIGURE 5

From asset location to Physical Exposure Rating

Scientific Climate Ratings' physical risk methodology, applied directly to 1,000 individual data centre assets.

1 Asset universe

1,000

individual data centre assets

Each asset's precise location and physical footprint is mapped before any hazard is applied.

2 Physical hazard models

Flood	Wind
Wildfire	Heat stress

Each model translates hazard intensity into an expected **damage or disruption metric**, calibrated with engineering damage functions and hazard-specific data.

3 Physical Exposure score

0-100

metrics combined into a single score

Benchmarked against a reference distribution, so results are comparable across assets, geographies and time.

4 Physical Exposure Rating

A → G

score converted to an A to G rating

One comparable grade per asset, per horizon.

RUN ACROSS

2035 **2050** horizons, under two **emissions pathways (SSP2-4.5 and SSP5-8.5)**

Source: Scientific Climate Ratings - The Cloud Has a Physical Address Report (September 2026)

Our approach starts at the level of the individual asset; we map each data centre's precise location and physical footprint, then assess it against Scientific Climate Ratings' physical hazard models for flood, wind, wildfire, and heat stress. Each model translates climate hazard intensity into an expected damage or disruption metric, using asset-type-specific engineering damage functions and location-specific hazard data. These metrics are then combined into a Physical Exposure score and converted into an A to G Physical

Exposure Rating, benchmarked against a reference distribution so that results are comparable across assets, geographies, and time. We ran this analysis across the 2035 and 2050 horizons under the two recognized emissions pathways, SSP2-4.5 and SSP5-8.5.

SCOPE

Why the analysis covers four hazards

Exposure to a broader set of physical hazards can be assessed in the underlying data. The analysis carries flood, wind, wildfire and heat through to financial impact because they combine material climate relevance with damage functions suitable for financial quantification.

1

1 - Material climate hazards
Flood, wind, wildfire and heat are among the major climate-related hazards responsible for substantial physical and economic impacts globally as identified in the UN Office for Disaster Risk Reduction's *The human cost of disasters: an overview of the last 20 years*.[4]

2

2 - Financially quantifiable
For these four hazards, peer-reviewed damage functions are available and sufficiently robust for our methodology, allowing physical exposure to be translated into financial impact.

Drought, precipitation, landslide, subsidence and earthquake are assessed at indicator level in the underlying data. They are not included in the financial-impact results presented here because sufficiently validated damage functions have not been incorporated into this analysis.

FIGURE 6

Four hazards, four mechanisms

Hazards reach a data centre by different routes, and the model treats them differently. Each card gives the average annual expected asset value impact, combining structural damage and operational disruption — the damage-only averages quoted later in the report are a component of these — alongside the cost of inaction that impact accumulates to by 2050. Averages are taken across all 1,000 assets, with assets not exposed to a hazard counted as zero rather than excluded, so the four figures are directly comparable. Wind and flood contribute through both damage and disruption; wildfire registers damage at a small number of sites only, and heat carries no damage mechanism, so its value impact is disruption alone. Hazard-level splits are apportioned from each asset's own hazard composition, so the four figures sum to the sample totals.

Wind

damage and disruption · 1,000 of 1,000 exposed

MECHANISM

Storms and high winds strike the building, rooftop cooling plant and external switchgear.

REACHES THE ASSET VIA

Direct damage, plus loss of the grid connection the site depends on.

Annual expected asset value impact

0.65%

Cost of inaction to 2050, discounted

5.51%

Wildfire

damage only · 19 of 1,000 exposed

MECHANISM

Fire front, radiant heat and smoke ingress reach the site and its transmission corridors.

REACHES THE ASSET VIA

Damage to structures and to the grid connections the site depends on.

Annual expected asset value impact

0.001%

Cost of inaction to 2050, discounted

0.008%

Flood

damage and disruption · 185 of 1,000 exposed

MECHANISM

Water reaches ground-floor plant rooms, generators, fuel stores and cabling routes.

REACHES THE ASSET VIA

Damage to critical equipment and prolonged outage during recovery.

Annual expected asset value impact

0.06%

Cost of inaction to 2050, discounted

0.43%

Heat

disruption only · 1,000 of 1,000 exposed

MECHANISM

Heatwaves raise the cooling load and reduce cooling efficiency, while grids are already overstretched.

REACHES THE ASSET VIA

Operating conditions and worker productivity, not direct structural damage.

Annual expected asset value impact

0.007%

Cost of inaction to 2050, discounted

0.054%

Source: Scientific Climate Ratings - The Cloud Has a Physical Address Report (September 2026)

DEFINITION

What the cost of inaction measures

The cost of inaction is the total discounted cost of an asset's future climate impacts, from today to the chosen horizon. It is a discounted-cash-flow sum over the asset's evolving annual impact trajectory, not an annual average.

This is where an annual figure near 1% stops looking modest. Accumulated across the years an asset still has to run and discounted back to today, a recurring loss of that size becomes a substantial share of the asset's value, and that total is what an owner accepts by leaving the asset as it is.

FORMULA

$$COI = \Sigma_t \ impact(t) \ / \ (1 + r_t)^t$$

Summed from the base year to the horizon, where *impact(t)* is the actual annual value impact in year *t* and *r* is the discount rate. Both evolve year to year; neither is held constant.

Figures quoted in this report are the modelled cost-of-inaction values from the underlying data, discounted to the 2050 horizon. Discount rates come from the Scientific Climate Ratings and NGFS scenario framework: they sit near 10% and vary by year, sector and country rather than being a single flat rate. For these assets they are a simple unweighted average across three models and eight NGFS scenarios for the information and communication sector. The calculation runs on the total aggregated physical impact rather than hazard by hazard, and covers physical risk only: the full cost of inaction also carries a transition-risk component, absent for these assets, so the figures understate it. All are gross of any further resilience investment beyond the measures already in place.

One further point is worth making upfront: this is an analytical sample, not a full census of the global data centre market. The findings below should be read as a rigorous look at a representative set of assets, not as a definitive map of the entire sector's exposure.

What our analysis reveals

The scale of climate exposure faced by data centres is significant, but the distribution of risk across the portfolio is more revealing than any single headline number.

Across the 1,000 assets we assessed, their physical climate exposure is widespread and far from uniform. The operator comparison shows a varied distribution of A-G ratings (with A reflecting the lowest risk). NTT and Microsoft sit at the higher exposure end amongst the operators, with 9.8% and 9.6% of their data centre assets respectively assigned a ‘G’ (severe) exposure rating. At the country level the concentration is far sharper: 54% of assets in the Netherlands and 32% in Japan rate G, while no other country in the sample reaches 1%.

FIGURE 7

Operator exposure, 2050 horizon

Share of each operator's assets by Physical Exposure Rating, from A (lowest risk) to G (severe). Rows are sorted by the share rated G, then by the share rated F.

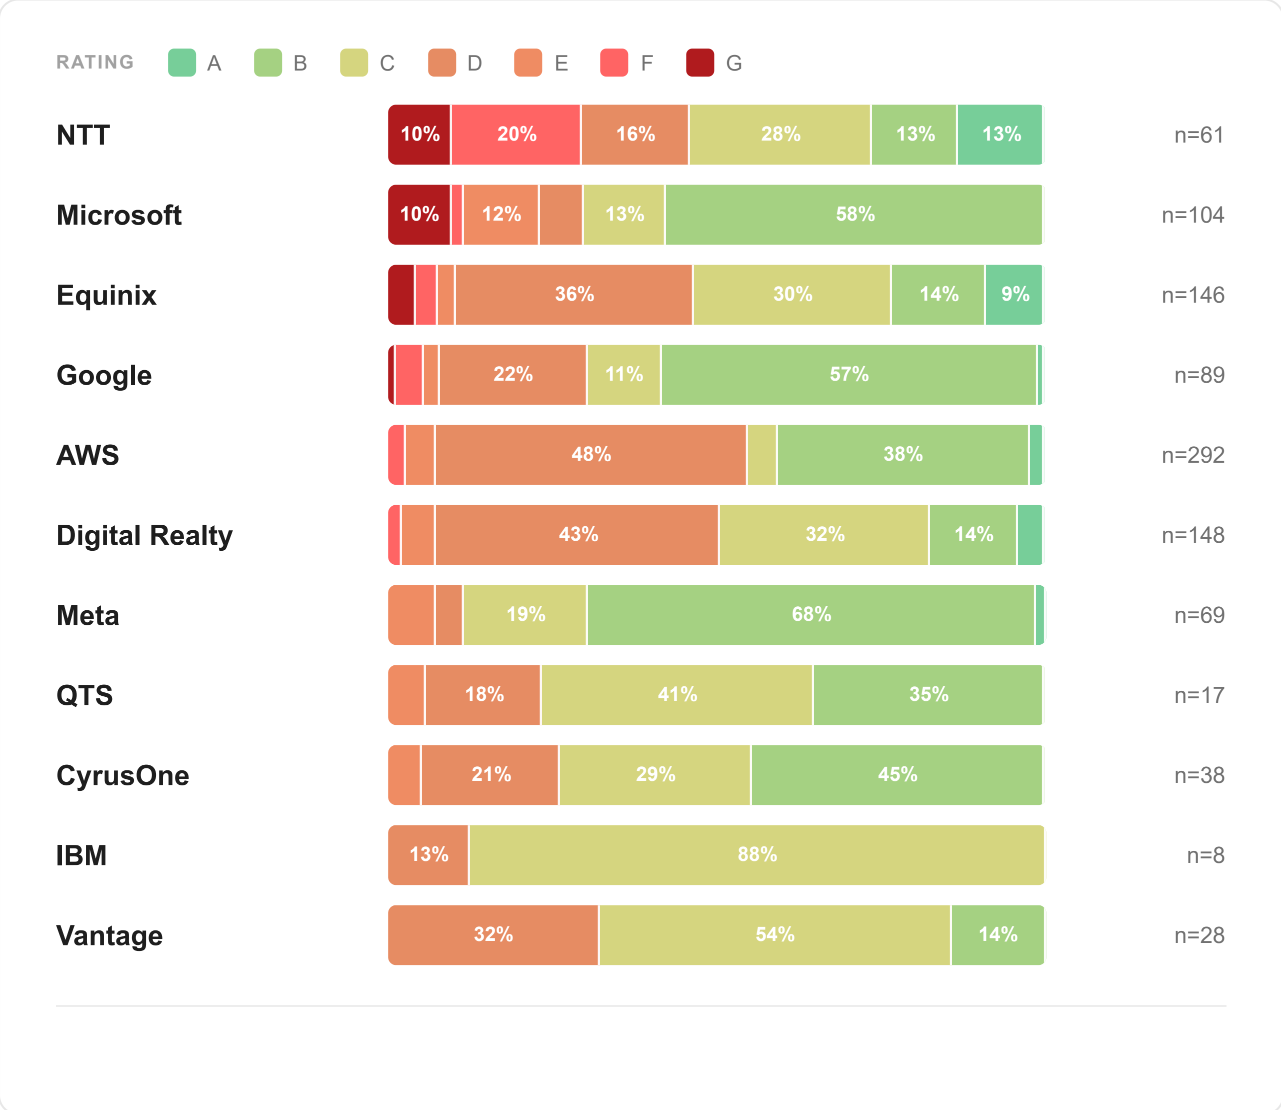

Source: Scientific Climate Ratings - The Cloud Has a Physical Address Report (September 2026)

The geographic variation is important because it reinforces a central finding of our analysis: climate exposure is driven more by where a data centre is located than by which company operates it. NTT holds both the single most exposed asset in the sample and assets near the bottom of the distribution, and Equinix appears among the five most exposed while also holding some of the least exposed. That is a striking illustration of the limitations of treating climate risk as an operator-level characteristic; two facilities belonging to the same company can sit at opposite ends of the physical-risk spectrum.

FIGURE 8

Same rating, different operators

The same view read the other way: all 1,000 assets sorted into their 2050 rating bucket and stacked by operator, with the range of modelled total value impact behind each rating letter. Most operators appear in most buckets. B is the largest bucket at 346 assets, and AWS supplies 32% of it; AWS's largest concentration is in D, where it holds 44%.

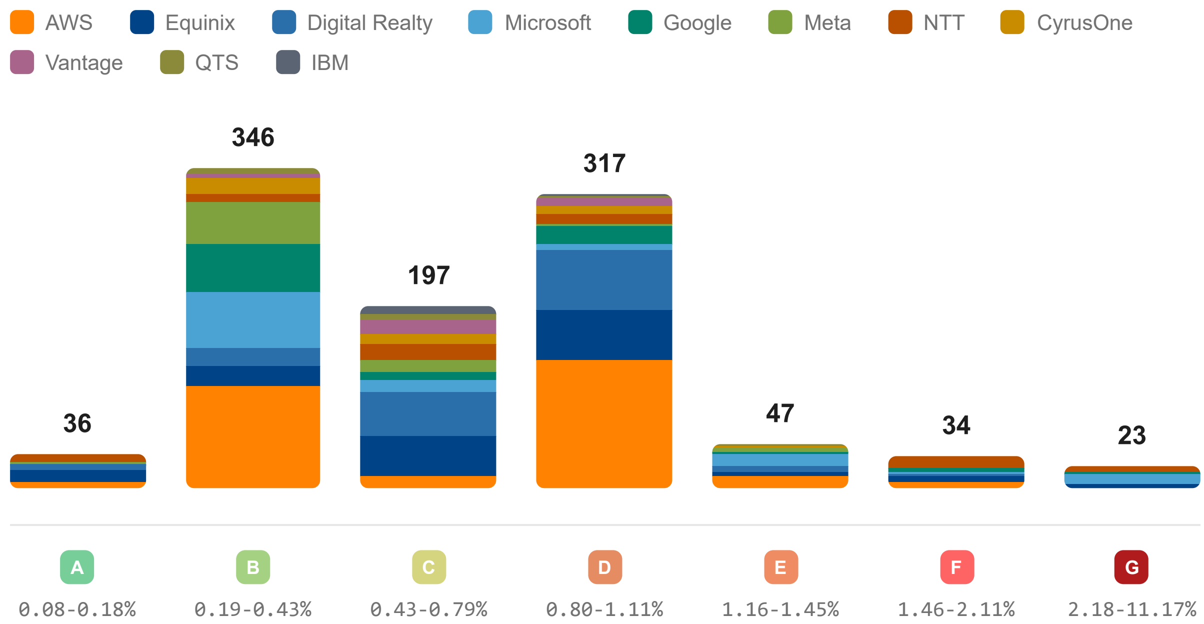

Column height is the number of assets in the bucket; the figure above each column is that count. Ranges are the minimum and maximum total value impact of assets carrying that rating at the 2050 horizon.

Source: Scientific Climate Ratings - The Cloud Has a Physical Address Report (September 2026)

The hazard-level results add another layer to this picture. The figures that follow isolate structural damage, which is one of the two routes by which a hazard reaches the value impact totals above; the other is operational disruption, the revenue effect of lost operating time. Separating the two matters because hazards do not split evenly between them: for the flood-exposed assets, damage accounts for 95% of the value impact, while heat carries no damage mechanism at all and reaches the totals through disruption alone. On damage, wind ranks highest, averaging 0.62% of asset value across all 1,000 assets, and it applies to every one of them. Flood ranks second on the same measure, averaging 0.06%, but is modelled for only 185 assets, or 18.5% of the sample. Wildfire is a distant third at under 0.01%, modelled for just 19 assets. Averaging across the full sample this way makes the hazards directly comparable, but it also means flood's low average reflects how few sites are exposed rather than how mild the hazard is where it does apply: among those 185 assets, flood damage averages 0.33%, their mean total value impact is 0.35%, and the worst-placed site reaches 8.9% of damage alone.

FIGURE 9

Structural damage: most assets sit near zero, a few sit far out

Structural damage only, as a percentage of asset value, across all 1,000 assets. Assets not exposed to a hazard count as zero rather than being excluded, so each bar describes the whole portfolio. Flood is the starkest case: 815 assets carry no flood damage at all, which puts both the median and the upper quartile at zero, yet one asset reaches 8.9%. Wind is the opposite, with no extreme cases but almost every asset carrying some damage.

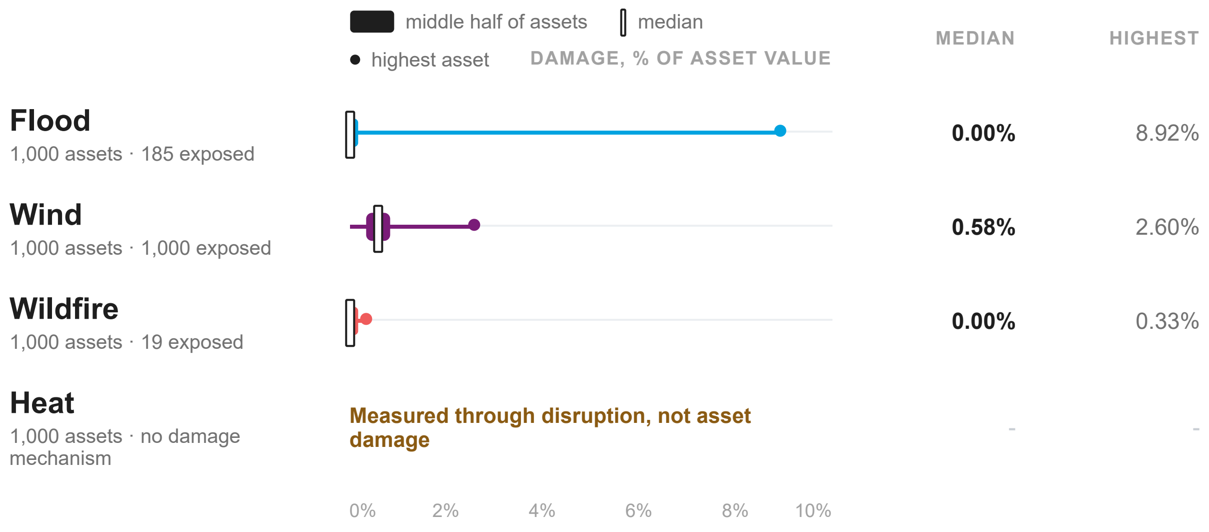

Structural damage as a percentage of asset value, SSP5-8.5 at the 2050 horizon, across all 1,000 assets, with assets not exposed to a hazard counted as zero. "Highest" is the single most exposed asset for that hazard, the value marked by the dot at the end of each bar. Operational disruption is excluded here; figure 6 gives each hazard's combined damage-and-disruption value impact. Heat impact is modelled for all 1,000 assets but has no structural damage mechanism, so it has no bar.

Source: Scientific Climate Ratings - The Cloud Has a Physical Address Report (September 2026)

Heat requires a different interpretation model again. The model does not assign a structural damage figure to heat: the absence of a value is not a zero, but reflects the fact that heat is treated as an operational hazard rather than a mechanism of direct physical damage to the structure. This distinction is particularly relevant for data centres, where extreme heat can affect cooling efficiency and operating conditions without necessarily producing the kind of direct asset damage captured for wind or flood.

The forward-looking analysis also reflects an important nuance. The 2035 and 2050 measures are look-ahead windows from today rather than sequential snapshots of the same asset's future state. In other words, the comparison should not automatically be read as a trajectory from 2035 to 2050. In fact, 93.5% of assets have an identical physical exposure rating at the two horizons. The implication is that much of the underlying location-specific exposure is already visible today, rather than emerging only in the longer-term scenario.

Carried forward, those annual figures compound into materially larger totals. On the same SSP5-8.5 pathway, discounted to 2050, the mean cost of inaction runs from 3.4% of asset value in Sweden to 20.6% in Japan, a sixfold spread driven almost entirely by location. The operator spread is far narrower, from 4.3% at Meta to 9.6% at NTT, which is the same finding the rating distributions showed: where an asset sits matters more than who runs it. The two charts below set out those totals, first by country and then by operator.

At the country level, however, the forward-looking financial picture shows meaningful differences in both the level and rate of change of exposure. Japan has the highest overall mean total-value impact in the sample at 2.78%, followed by the Netherlands at 1.85% and Ireland at 1.47%. The Netherlands is particularly notable because it has the steepest change between the two horizons, widening by 0.58 percentage points.

FIGURE 10

Country cost of inaction: to 2035 vs to 2050

One column per country; a blue dot marks the mean cost of inaction discounted to 2035, an orange dot the same figure to 2050. Japan leads on both, reaching 20.6% of asset value by 2050, six times Sweden's 3.4%. Limited to the ten countries with the most assets.

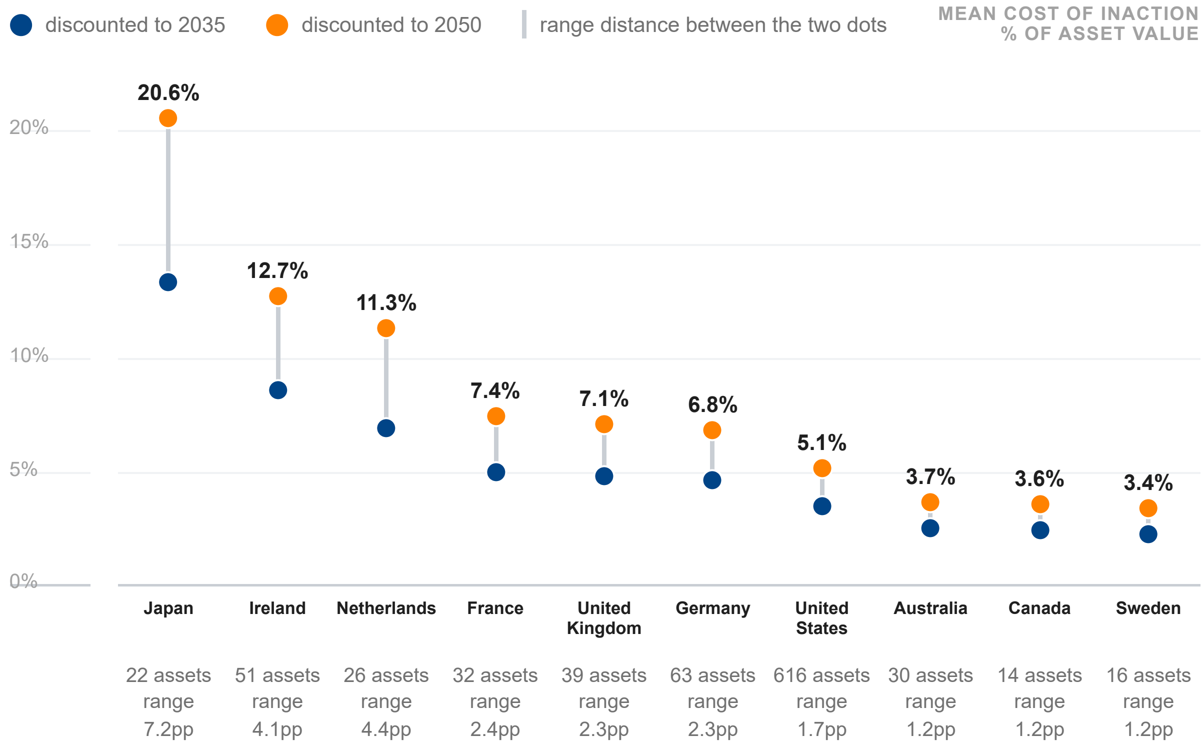

Mean cost of inaction per country, as a percentage of asset value, taken from the modelled cost-of-inaction figures in the underlying data and discounted using the Scientific Climate Ratings and NGFS rates, which sit near 10% and vary by year, sector and country. Columns are sorted by the 2050 figure; under each label is the asset count and the distance between the two dots in percentage points. The longer horizon carries more years of accumulated impact, so the 2050 figure is larger by construction. Figures cover physical risk only and are gross of any further resilience investment beyond the measures already in place.

Source: Scientific Climate Ratings - The Cloud Has a Physical Address Report (September 2026)

FIGURE 11

Operator cost of inaction: to 2035 vs to 2050

Same format as the previous figure, by operator. The spread is far narrower than the country view: 4.3% to 9.6% by 2050, against 3.4% to 20.6% across countries. NTT stands clear at the top; the other ten operators sit between 4.3% and 6.3% on very different portfolio sizes.

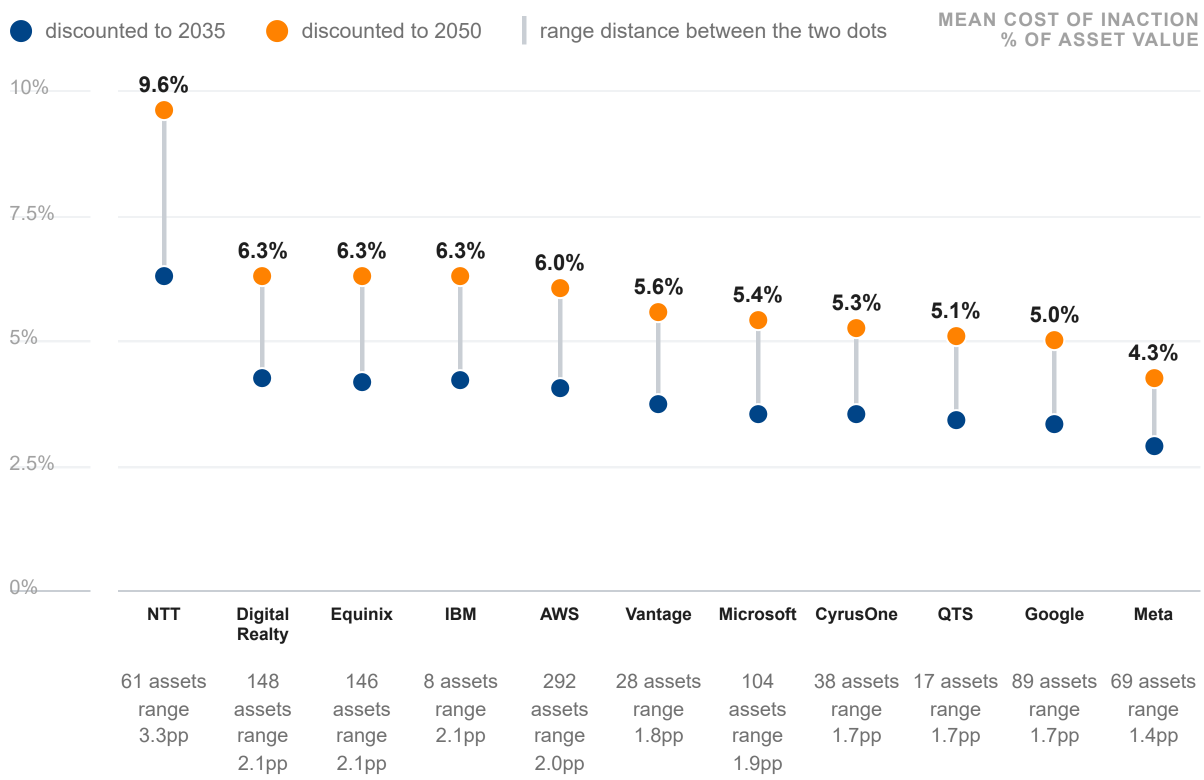

Identical to the previous figure, grouped by operator, and on the same basis: physical risk only, gross of further resilience investment. Note the vertical scale differs: it runs to 10% here against 22% for countries, so these columns are not directly comparable in height with the chart above.

Source: Scientific Climate Ratings - The Cloud Has a Physical Address Report (September 2026)

The extremes reinforce why averages should be treated carefully. The five highest-exposed assets are two NTT facilities and one Equinix facility at Urayasu on Tokyo Bay, an Equinix site at Union City (USA) and a Google facility at The Dalles (USA). At the other end of the distribution sit Amazon Web Services facilities in Bahrain and the United Arab Emirates, alongside NTT sites in Johannesburg and Bangalore and Equinix sites in Dubai and Abu Dhabi. The concentration of both high- and low-exposed assets within the same operators again points to location as a critical differentiator. The figure below takes the most exposed of those assets and asks a different question of it: what do adaptation and resilience (A&R) measures already in place change to the total impact on a single asset?

FIGURE 12

Deep dive: what A&R measures change at a single asset

Before its flood A&R measures are taken into account, NTT's TK9 Data Center in Kōtō, Tokyo is one of the most exposed assets in the sample, with a modelled value impact of 11.31%. With the A&R measures in place, that falls to 4.26%. The A&R measures address the asset's high flood exposure, representing almost all of this asset's exposure.[5]

G **NTT TK9 Data Center**
Kōtō, Tokyo · Japan · temperate zone

BEFORE A&R MEASURES	WITH A&R MEASURES	REDUCTION
11.31%	**4.26%**	**−62%**
17.1× the sample median	6.4× the sample median	of value impact

VALUE IMPACT BY HAZARD, BEFORE AND WITH MEASURES · SSP5-8.5 AT 2050

Legend: value impact with A&R measures; avoided by A&R measures

HAZARD	RT		BEFORE	AFTER	CHANGE
Flood	G		9.51%	**2.47%**	**−74%**
Wind	G		1.78%	**1.78%**	no change
Heat	F		0.01%	**0.01%**	no change
Wildfire	A	no impact modelled	-	-	-
Total			**11.31%**	**4.26%**	**−62%**

Each bar is one hazard's total value impact before A&R measures, scaled to the asset's 11.31% total; the solid portion is what remains with A&R measures in place and the hatched portion is what they avoid. Flood impact falls from 9.51% to 2.47% of asset value, a 74% reduction. Wind and heat are unchanged, since the A&R measures in place address flood only. The rating stays at G: the asset remains highly exposed, but a smaller share of that exposure converts into loss.

COST OF INACTION, BEFORE AND WITH MEASURES

Discounted to 2035
~~38.8%~~ → **18.7%**
benefit of measures 20.1pp · −52%

Discounted to 2050
~~65.4%~~ → **29.5%**
benefit of measures 35.9pp · −55%

Discounted from 2026 at this asset's own Scientific Climate Ratings / NGFS rate of 10.2%. The A&R measures already in place halve the cost of inaction over both windows, from 65.4% to 29.5% of asset value to 2050 — a benefit of 35.9 percentage points at this single facility.

WHAT THE A&R MEASURES DO NOT REACH

Wind still contributes 1.78% and rates G. Flood A&R measures leave the asset's second hazard untouched, so residual exposure is now wind-led rather than flood-led.

OPERATOR CONTEXT

NTT's median asset sits at 0.79%, below the sample median. Its portfolio is not broadly exposed; its Japanese coastal sites are.

COUNTRY CONTEXT

Japan holds 22 assets in the sample, all of them rated F or G, with an average impact of 2.78% against 0.72% globally.

Source: Scientific Climate Ratings - The Cloud Has a Physical Address Report (September 2026)

The impact of A&R measures

TK9 is a concrete case of a more general point: an asset's exposure to a given hazard is not the whole story. It is true a data centre's location is fixed at the point of construction; an operator cannot move a facility away from a floodplain or a heat-stressed grid once ground has broken. But how that exposure translates into realised risk is not set in stone.

Cooling design, water strategy, and structural integrity are all operational choices, and each one can materially change how vulnerable a given asset is to the hazards it faces. Two facilities sitting in the same flood zone or the same heat corridor can carry very different risk profiles depending on whether one has invested in A&R measures.

This is the distinction between exposure and vulnerability that runs through this report and all our rating products: exposure describes what an asset is confronted with, while vulnerability describes how much of that exposure actually converts into downtime, damage, or abandonment. A&R measures are the mechanism by which operators narrow that gap. A high-exposure site with strong A&R measures can carry materially lower realised risk than its raw hazard picture would suggest: at TK9 the A&R measures already in place cut modelled value impact by 62% and halve the cost of inaction, without moving the asset off a G rating.

Taken together, the results suggest that the most useful way to think about physical climate risk in data centres is not as a simple ranking of operators. The more important question is where the highest-risk assets sit, which hazards drive that exposure, how much of the potential impact is concentrated in a relatively small number of facilities, and how much of it the A&R measures already in place remove.

The implications for stakeholders

The significance of physical climate exposure is amplified by the characteristics of data-centre assets. They are capital-intensive, have long operational lifetimes, and are effectively immovable once constructed. Decisions about location and design therefore have implications over decades rather than over a typical technology investment cycle.

That has practical consequences across four decision points.

CapEx planning

An asset's climate exposure should serve as a direct input into resilience spending, such as backup cooling capacity, on-site water storage, and hardened structures. This should be budgeted alongside the build itself rather than treated as a later retrofit.

Insurance terms

Underwriters are already differentiating pricing and, in some cases, availability by physical exposure. An asset's Physical Exposure Rating gives a forward-looking basis for that conversation rather than relying solely on historical loss data that predates the current pace of warming.

Financing conditions

Lenders assessing a data centre-backed facility can use asset-level exposure, not just operator-level averages, to price risk correctly, particularly given the variation within portfolios.

Site-selection strategy

For the next wave of build-out, exposure ratings are best used at the stage of choosing a site, when the cost of avoiding a hazard is lowest, rather than after ground has broken.

The AI build-out is therefore creating not only a technology investment cycle, but a major infrastructure investment cycle in which physical climate exposure can affect the cost, resilience, and long-term value of the assets being built.

What the results do-and do not-show

As with any asset-level climate analysis, the results should be interpreted within the scope of the data and methodology. The sample covers 1,000 data centres for the reasons described above and is intended to provide a rigorous view of physical climate risk across a large set of assets, rather than a complete census of the global market.

- **The analysis incorporates implemented A&R measures where they are publicly disclosed and can be evidenced.** Undisclosed measures may not be captured, meaning some individual sites could be more resilient than the available data indicates.
- **Operator-level comparisons depend on sample coverage.** Where only a small number of sites are represented, individual assets can materially influence an operator's aggregate result.
- **The cost-of-inaction figures presented here capture physical climate risk only.** The broader framework also incorporates transition risk, which is outside the scope of this study. The figures should therefore be read as a partial rather than comprehensive measure of climate-related financial impact.
- **Some systemic and second-order effects remain outside the current modelling framework**. These include cascading grid failures, supply-chain disruption, interdependencies between facilities and changes in insurance availability or pricing. Such effects could amplify the direct asset-level impacts estimated here.

The future of AI is also climate dependent

The AI build-out is already well underway, developing with unprecedented scale and speed. But this build-out is also running directly into a warming world and changing weather conditions, in ways that are already evident at the level of the individual asset.

The next constraint on the AI economy will not be chips or power alone. The continued operation of data centres will be dependent on climate-resilient design and financing and ensuring that these centres are built in places, and in ways, that can withstand the changing physical world.

IN CLOSING

“The cloud may be digital, but the risks it faces are increasingly physical.”

Scientific Climate Ratings The Cloud Has a Physical Address: Climate Exposure in the Infrastructure Behind AI

REFERENCES

1. Central Statistics Office (2026). Data Centres Metered Electricity Consumption 2025. Available at: https://www.cso.ie/en/releasesandpublications/ep/p-dcmec/datacentresmeteredelectricityconsumption2025/keyfindings/

2. Shehabi, A. et al. (2024). 2024 United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory. Available at: https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report_1.pdf

3. Newsweek (2026). Map shows data centers built in drought-hit areas, drawing on US Drought Monitor data. Available at: https://www.newsweek.com/map-data-centers-built-drought-hit-areas-11997520

4. UN Office for Disaster Risk Reduction (2020). The human cost of disasters: an overview of the last 20 years, 2000-2019.

5. NTT Communications. Tokyo No.9 Data Center facility specification. Available at: https://www.ntt.com/content/dam/nttcom/hq/jp/business/services/data-center/colocation/nexcenter/pdf/tokyo%20no9%20data%20center-E.pdf

6. Cover photograph by Tony Marinescu on Unsplash. Available at unsplash.com/photos/fImCPTZ026U