Hyperscale Data Centres

Information sharing session for the Board

Released under the Official Information Act 1982

ELECTRICITY AUTHORITY
TE MANA HIKO

Purpose:

- To discuss:
  - what hyperscale datacentres are (briefly),
  - the datacentre industry (and recent developments) in New Zealand,
  - possible implications for reliability and affordability, and
  - potential gaps to address.

Released under the Official Information Act 1982

BACKGROUND AND CONTEXT



Released under the Official Information Act 1982

What are hyperscale data centres?

- >50MW
- Energy use is mainly ICT equipment to power AI workloads
- Growing fast
- Regulation moving slower than problems emerging in many places.

IEA scenarios of electricity demand from data centres

Global data centre electricity consumption (TWh). Four IEA scenarios for 2030 and 2035.

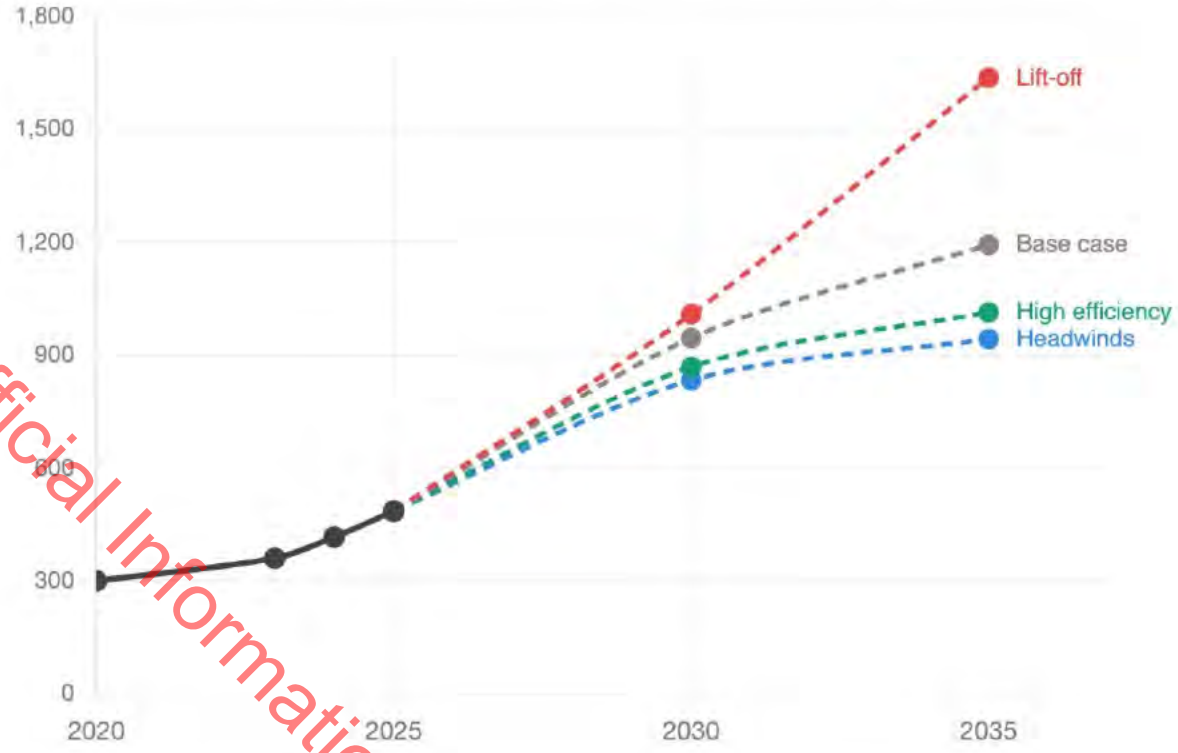

Note: 2035 figures are exploratory given the high level of uncertainty around data centre demand growth.
Data source: IEA, Key questions on Energy and AI (2026).

Released under the Official Information Act 1982

THE DATACENTRE INDUSTRY IN NEW ZEALAND

What we know currently



Released under the Official Information Act 1982

NZ portfolio

NZ has around **64 operational data centres**.
Average size is 227kW.

Total capacity:
2025: 104 to 125 MW
2026: 154 to 208MW

(Noting that capacity ≠ utilisation)

Estimates of total capacity vary:

Tech New Zealand (September 2025): **104MW**
BCG (November 2025): **125 MW**
NZ AI Infrastructure Index (July 2026): **154MW**
Research and Markets (June 2026): **208MW**

Microsoft, AWS and Datacom disclose no megawatts, so nobody can weight this by capacity.
And some figures that are disclosed (e.g., CDC) are contested.

Hence, we see estimate variation.

Released under the Official Information Act 1982

NZ portfolio

- **Around 4-6 could be called hyperscale:**

- CDC Silverdale 1 (Auckland)
- CDC Silverdale 2 (Auckland)
- CDC Hobsonville 1 (Auckland)

→ **80MW to 220MW** (contested)

- DCI AKL01 + AKL02 (Westgate + Albany, Auckland)

→ **~50MW**

- Microsoft New Zealand North (North and West Auckland)
- AWS Asia Pacific (Auckland)

→ **Likely tens of MW**, but not public

Released under the Official Information Act 1982

DCs were ~0.5% of NZ's total electricity consumption in 2024, increasing since 2021

Released under the Official Information Act 1982

Chart 13: Data centres (terajoules)



Source: EECA

What is driving interest in NZ as a potential data centre hub?

- Cool climate
- Political stability
- Connectivity
- Skilled labour
- Highly renewable grid
- Forward prices

Released under the Official Information Act 1982

Exhibit 9: New Zealand's industrial firmed renewable PPA price ranges are competitive with international peers

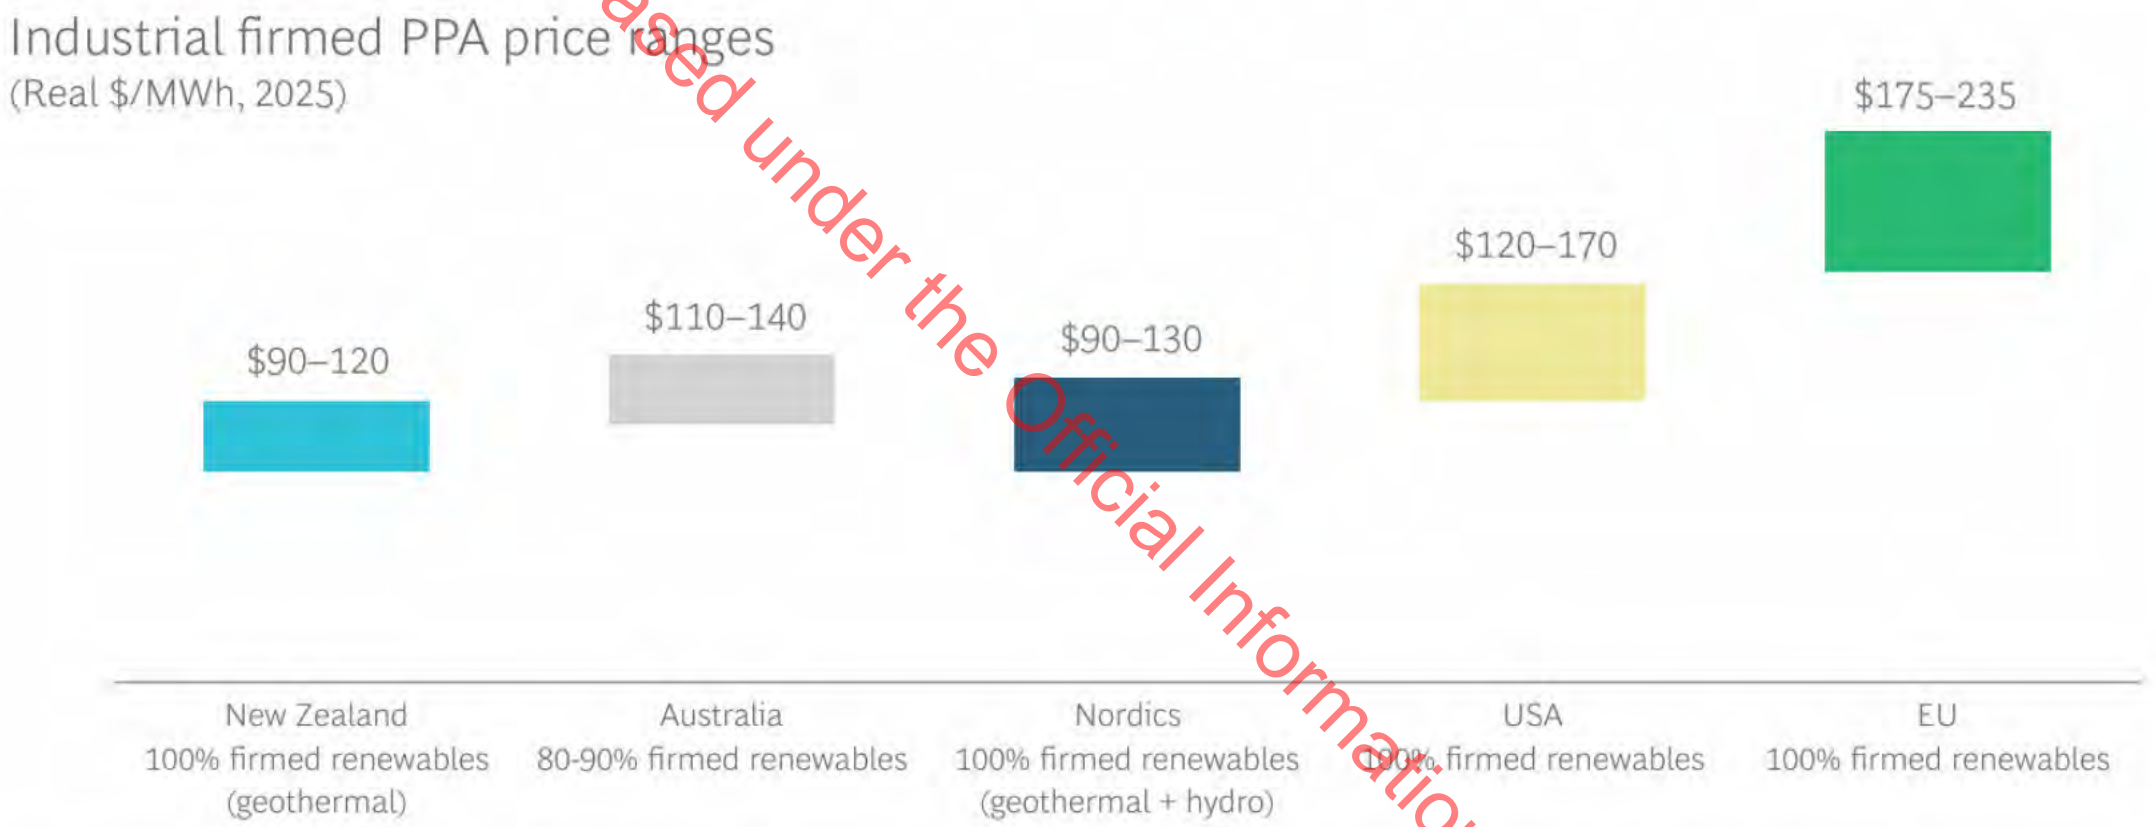

Note: Prices reflect energy-only costs for 10–12-year firm PPAs starting 2026–2028, exclusive of retail/wires. Methodologies benchmark against long-run baseload prices, geothermal LRMCs, and shaping premiums. ±10–20% caveats apply for basis risk, contract structure, hourly matching scope and local attribute rules. Currency conversions (USD, EUR, AUD) made 13 August 2025
Source: CSIRO; KYOS; OPIS; ASX Announcements; Lazard LCoE; Data Center Dynamics; Reuters; ATB/NREL; Eurelectric

Released under the Official Information Act 1982

From the pipeline (Transpower connection application data)

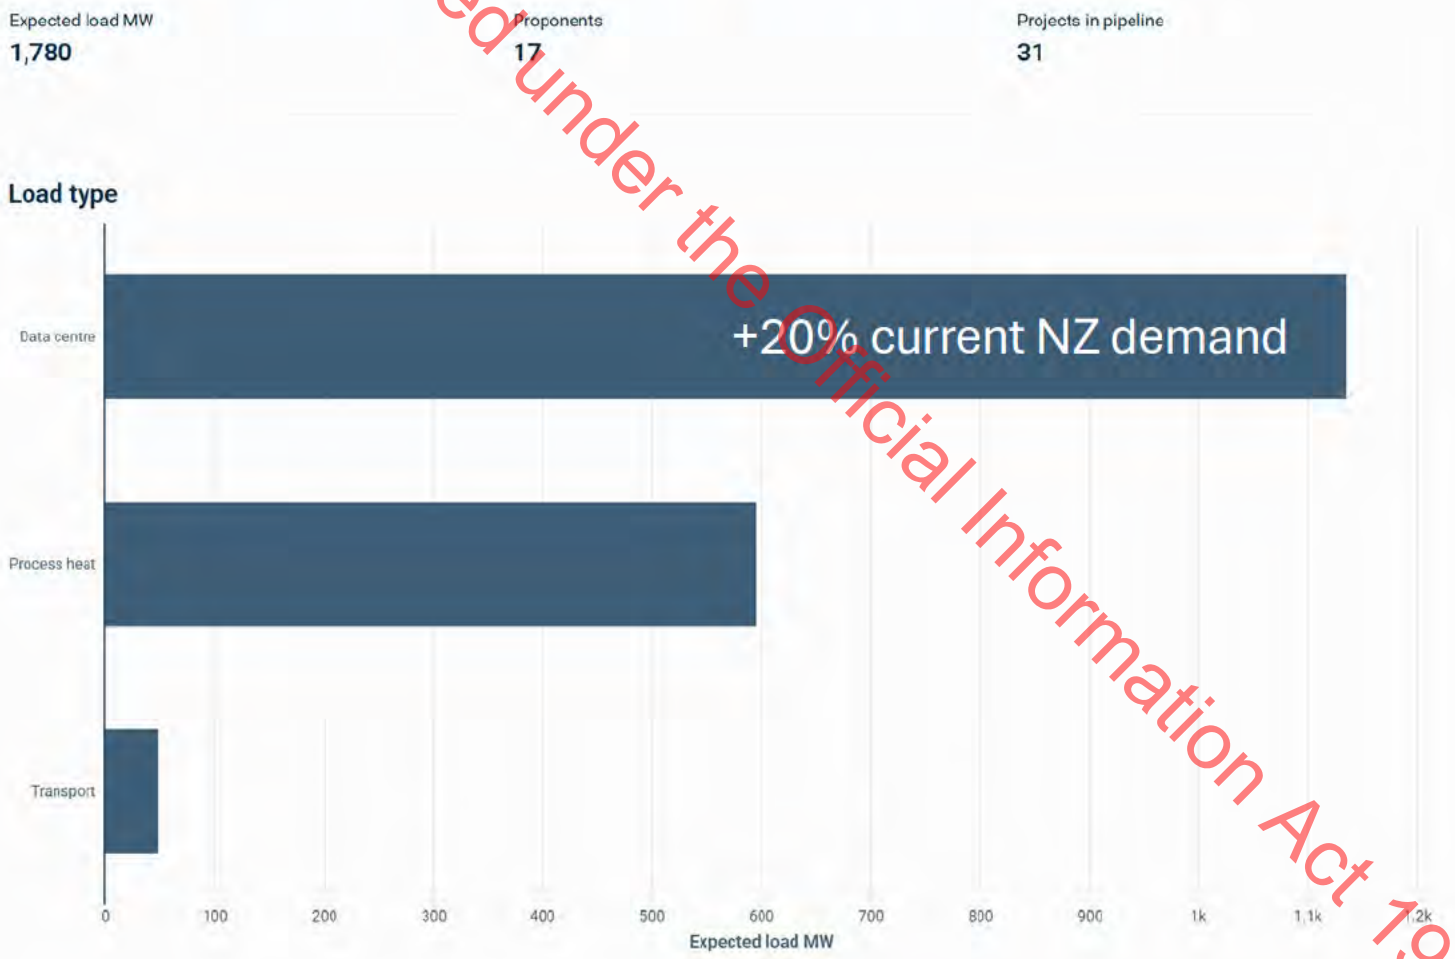

Released under the Official Information Act 1982

TWO EXAMPLES FROM THE PIPELINE



Released under the Official Information Act 1982

Let's consider…

Two data centre developments:

- CDC + Contact in Stratford
- Datagrid + Mercury in Invergcargill

And compare with one industrial:

- Tiwai (Potline 4)

Contact, CDC partner to explore data centre development

08:30 Mon 10th August 2026 • GENERAL

News / Business 23 July 2026, 10:52am

Mercury spends $53m onstake in Datagrid NZ

News / Business 27 May 2026, 3:56pm

Contact signs deal to get Line 4 potline at Tiwai Point aluminium smelter restarted

Released under the Official Information Act 1982

Impacts on the Power System are very different – not as simple as presented...





	MW generation	Annual GWh generation	MW load	Annual GWh load (at 75% utilisation)	Net GWh
Southland Wind Farm / PL4	380	**1210**	50	**438**	**+772**
Puke Kapo Hau Wind Farm / Datagrid	190	**550**	140	**920**	**-370**
Stratford solar farm / CDC datacentre	250	**312**	250	**1642**	**-1330**

Released under the Official Information Act 1982

NZ POLICY CONSIDERATIONS:

AFFORDABILITY AND RELIABILITY



Released under the Official Information Act 1982

Issue	Detail	Authority action underway (26/27)	Remaining gaps
Impact on wholesale electricity prices	Depends on demand and supply balance.	• Long-duration firming (LDF), • Improved stress testing • Updating the Official Conservation Campaign minimum weekly amount. • We are reviewing the market rules for hybrids (including data centres).	s 9(2)(g)(i)
Potential to create an energy gap	Data centre load growth could create a firm energy gap if it increases firm energy requirements without a corresponding increase in firm energy resources. Intermittent generation co-located with batteries (hybrids) can help to firm supply.		**LDF:** Staff will advise on options to strengthen incentives for LDF. Could be informed by consultation in October. Intersects with MBIE's work on dry year risk.

Released under the Official Information Act 1982

Issue	Detail	Authority action underway (26/27)	Remaining gaps
Impact on power system quality	Data centres can cause issues with: • Frequency and voltage due to site trippings • Power oscillations due to control system interactions • Scheduling and dispatch (load shifts quickly, difficult to forecast)	We are examining how overseas jurisdictions are managing the impact of data centres on system stability.	Assessment of obligations for large loads to support voltage and frequency. s 9(2)(g)(i)
Impact on network costs	Large load will share in the costs of existing network infrastructure. If demand growth exceeds current expectations, grid upgrades may be required to accommodate. Additional costs will be allocated according to net private benefits.	No relevant actions underway.	Staff do not assess any gaps at present.

Released under the Official Information Act 1982

Released under the Official Information Act 1982

ELECTRICITY AUTHORITY
TE MANA HIKO

DISCUSSION

Annex: Additional information

Released under the Official Information Act 1982

ELECTRICITY AUTHORITY
TE MANA HIKO

Key info (1/3)

Hyperscale: 50+ MW, with some reaching several hundred MW. (Compare with conventional DCs: <10-50 MW)

The sector has grown rapidly throughout the globe, and is expected to keep growing.

In some locations, DC's are a large % of total load (e.g., Ireland, ~20%), and expected growth (e.g., US, ~50%)

Several jurisdictions have started to regulate

(e.g., moratoriums, stricter connection criteria, firm + renewable + additionality requirements).

IEA scenarios of electricity demand from data centres

Global data centre electricity consumption (TWh). Four IEA scenarios for 2030 and 2035.

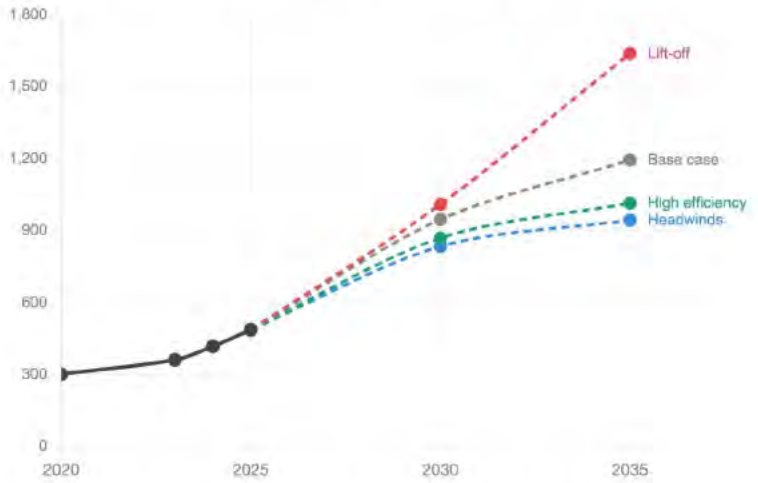

Note: 2035 figures are exploratory given the high level of uncertainty around data centre demand growth.
Data source: IEA, Key questions on Energy and AI (2026).

Irish metered electricity consumption, 2015 to 2024 (GWh)

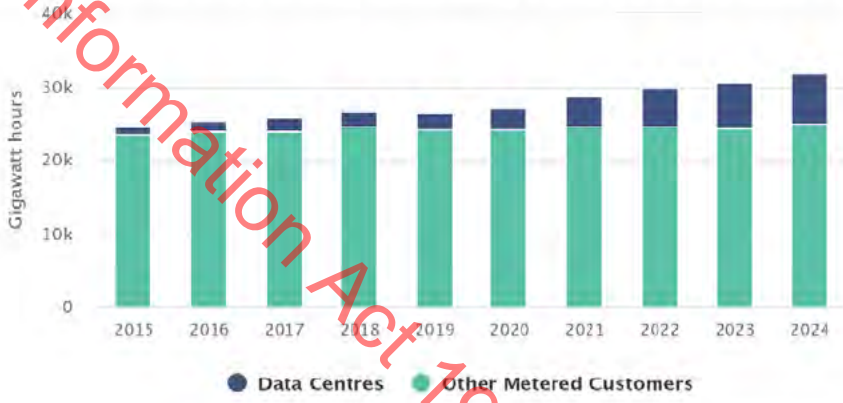

Source: CSO Ireland

Released under the Official Information Act 1982

Key info (2/3)

Hyperscale DC's typically take 2 to 4 years to develop.

Global power usage and compute is dominated by the US, China and the EU.

Geographical clustering is the norm.

Auckland, 34



Melbourne, 75



Dublin, 199

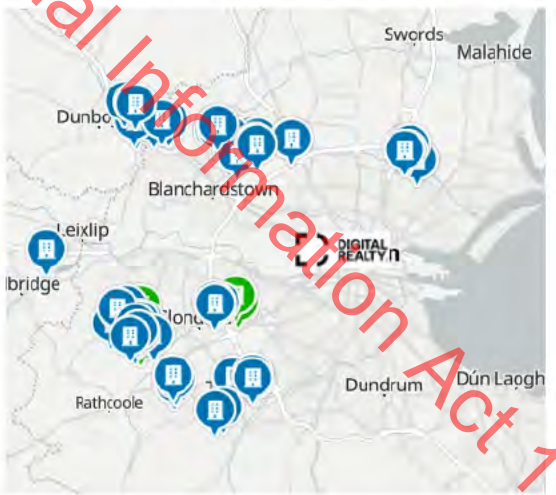

Ashburn, Virginia, 176



Released under the Official Information Act 1982

Key info (3/3)

Powered by a mix of renewables, gas, coal, and nuclear.

- Coal and gas due to 1) clustering in places with less renewable grids, and 2) self-supply driven by need for speed.

Virtually all electricity consumed becomes heat, but most heat is not re-used.

AI workloads have become more energy efficient, but consumption growth continues.

It is difficult to know what consumed electricity is actually used for. "AI" is a broad category, and data is commercially sensitive. Perhaps less than 10% is used for text queries.

Some data centres could be flexible, but in practice they are not inclined to be (99.999% uptime requirements).

Figure 5.4 ▷ Global electricity generation for data centres, 2020-2030

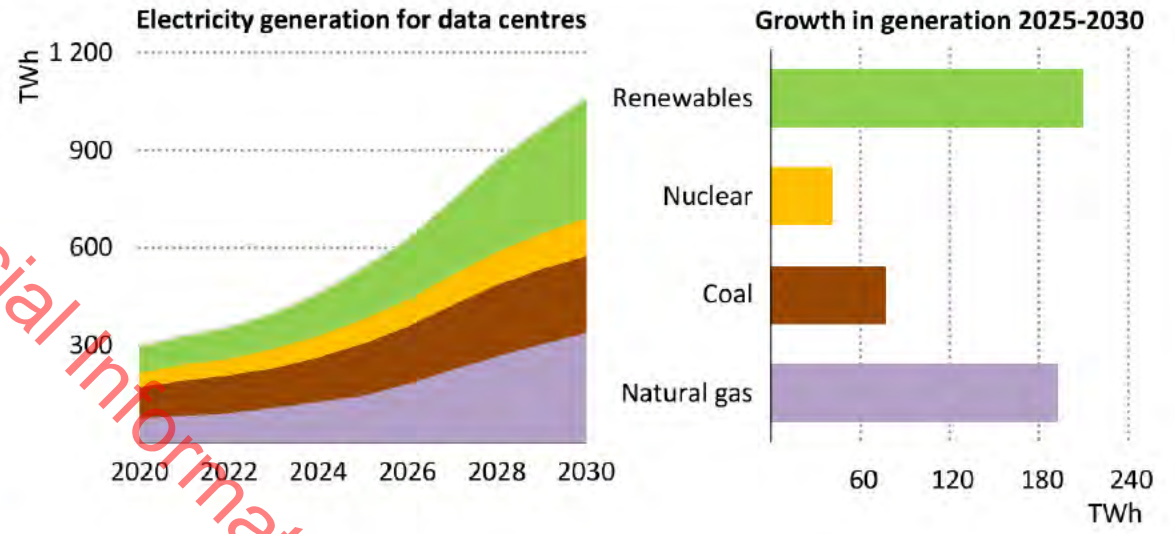

IEA. CC BY 4.0.

Released under the Official Information Act 1982

What's inside an individual server rack?

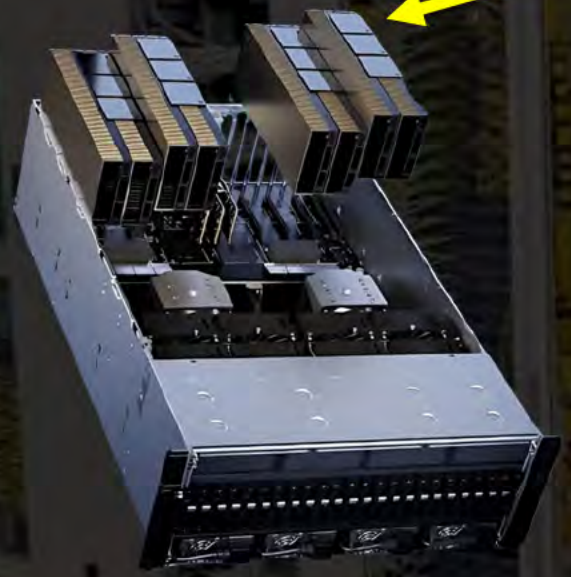

NVIDIA H100 (2022)



NVIDIA H200 (2023)

These chips ('GPUs') are the components inside HS-DCs that are used to power AI workloads,

and consume most of the electricity (60 to 90%)

and generate nearly all of the heat.

Many HS or AI-focused data centres contain tens of thousands of these chips.

Released under the Official Information Act 1982

Putting Hyperscale Data Centre load into perspective…



=



One fridge sized server rack

~65 households (2027 projection)

Released under the Official Information Act 1982

What distinguishes hyperscale / AI datacentres?

Size and speed of sectoral growth.

Conventional:

<10-50 MW.

Often single buildings. Mostly ICT equipment.

Hyperscale:

50+ MW, with some reaching several hundred MW.

Often multiple buildings. More ICT equipment.

More power supply and cooling equipment, also.

How Meta rebuilt a data centre for AI



Source: Planet, FT research

Released under the Official Information Act 1982

Hyperscale datacentres typically take 2 to 4 years to develop

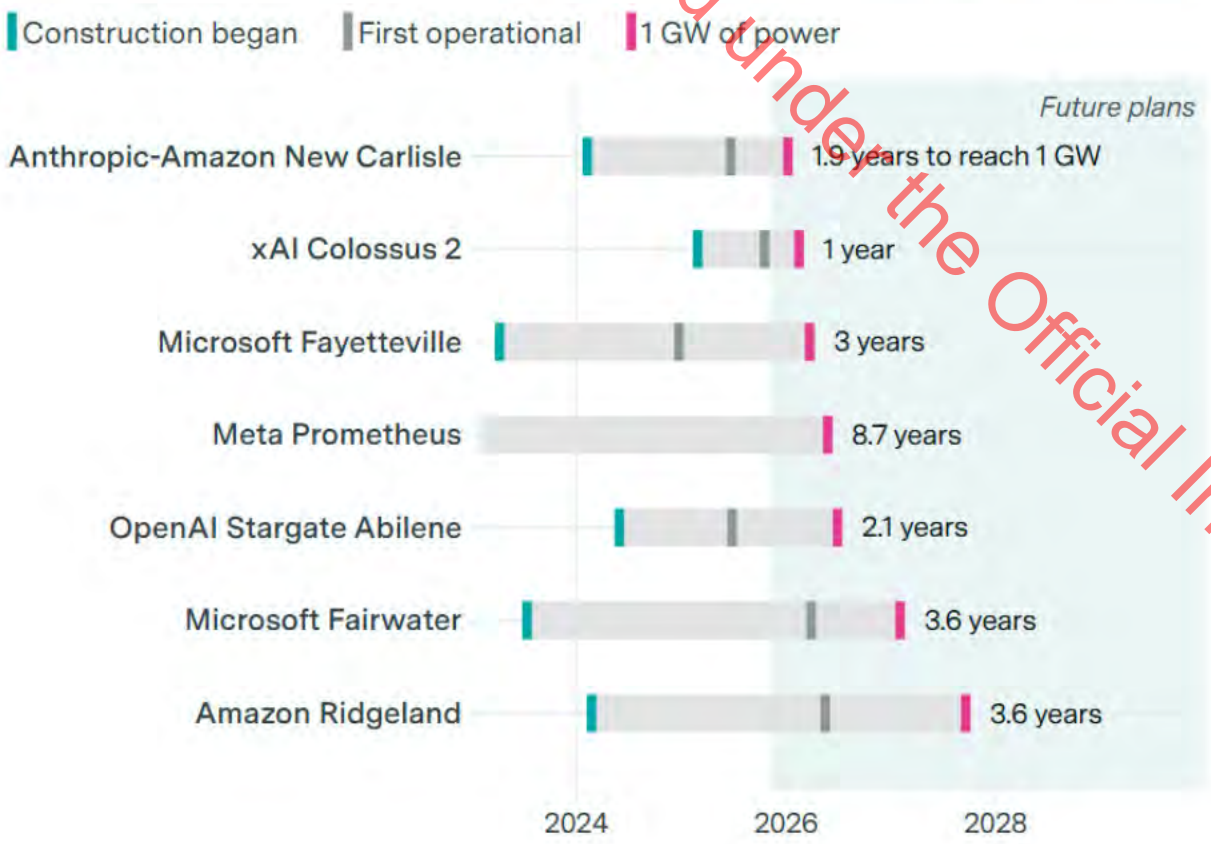

Timing includes:

- Site acquisition,
- Planning
- Consenting
- Connecting to the grid
- Construction
- Testing and commissioning

Released under the Official Information Act 1982

EPOCH AI | CC-BY

Global power usage and compute is dominated by the US, China and the EU

Released under the Official Information Act 1982

Share of electricity used by data centres, 2025

Data centre electricity consumption as a percentage of total electricity generation.

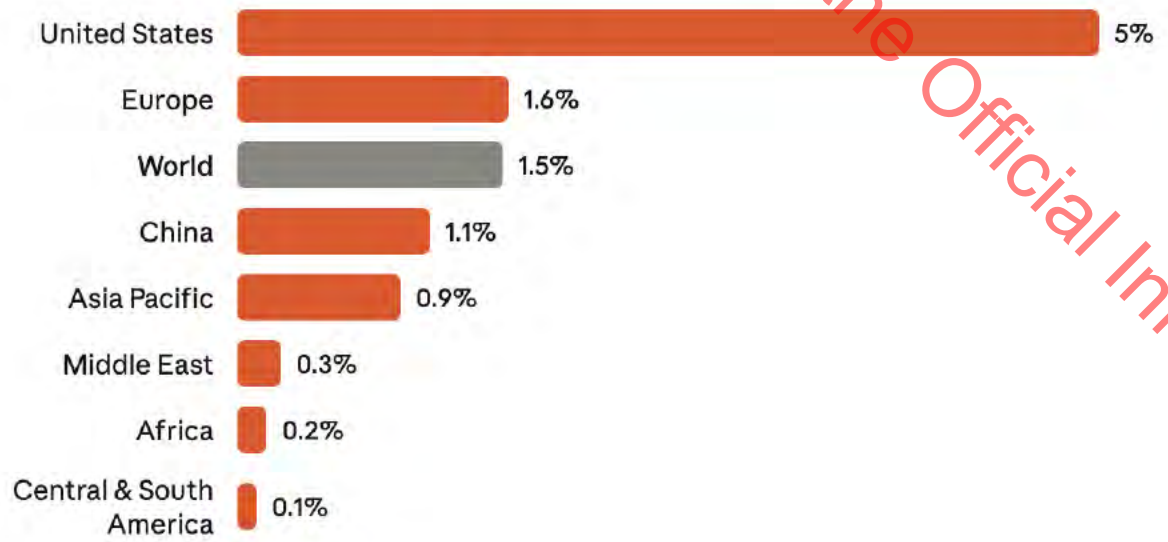

Note: Regional electricity totals are based on Ember (2026). Regional boundaries may not align exactly with IEA definitions.

Data source: Calculated based on IEA, Key questions on Energy and AI (2026) and Ember (2026).

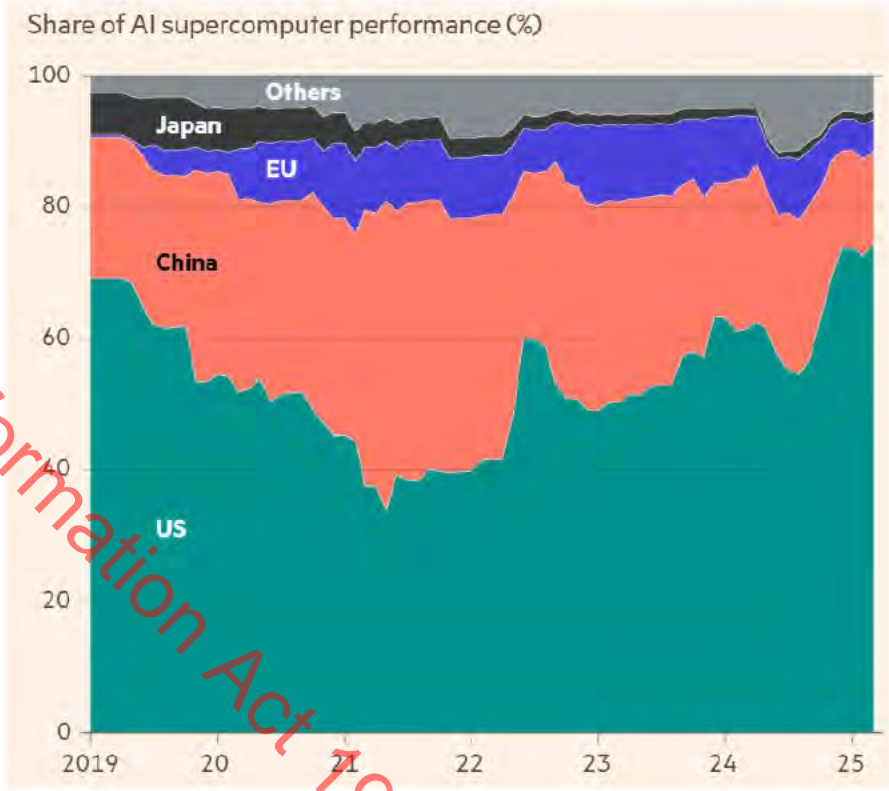

Geographical clustering is the norm, particularly for AI-focused DCs

Two reasons:

Within firms: low-latency GPU interconnection.

Across firms: shared infrastructure, resources, and skilled labour.

Auckland, 34

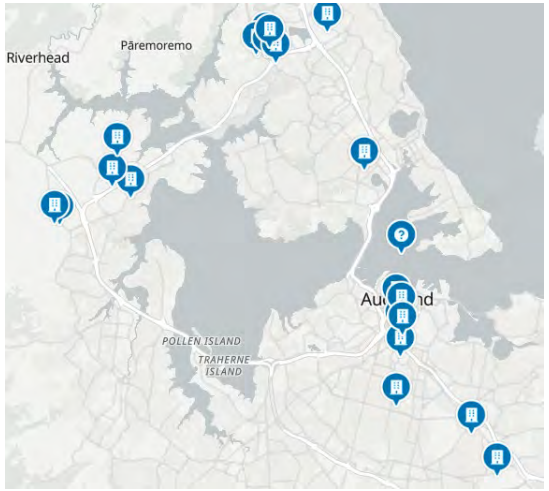

Melbourne, 75

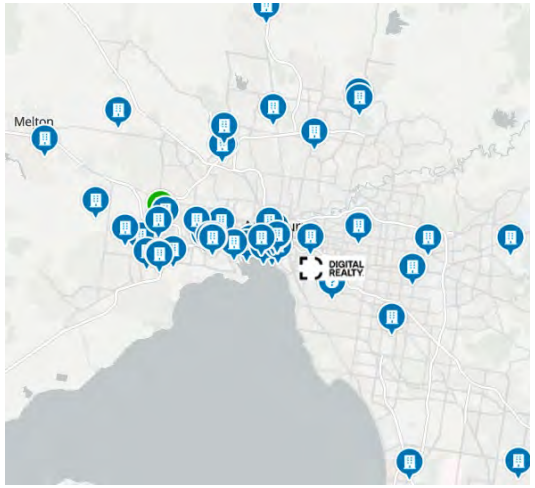

Dublin, 199

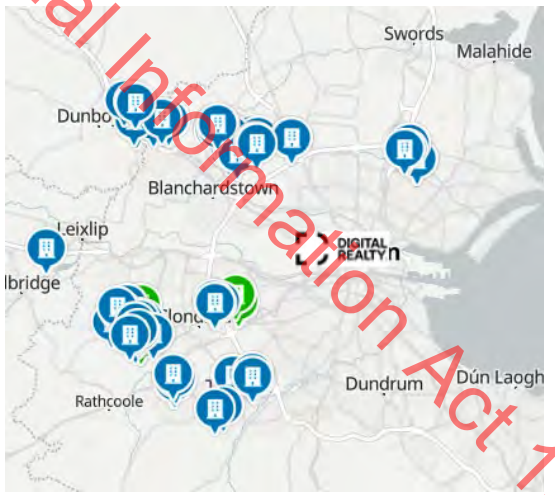

Ashburn, Virginia, 176



Released under the Official Information Act 1982

In the US, half of electricity demand growth in the next four years is expected to come from datacentres

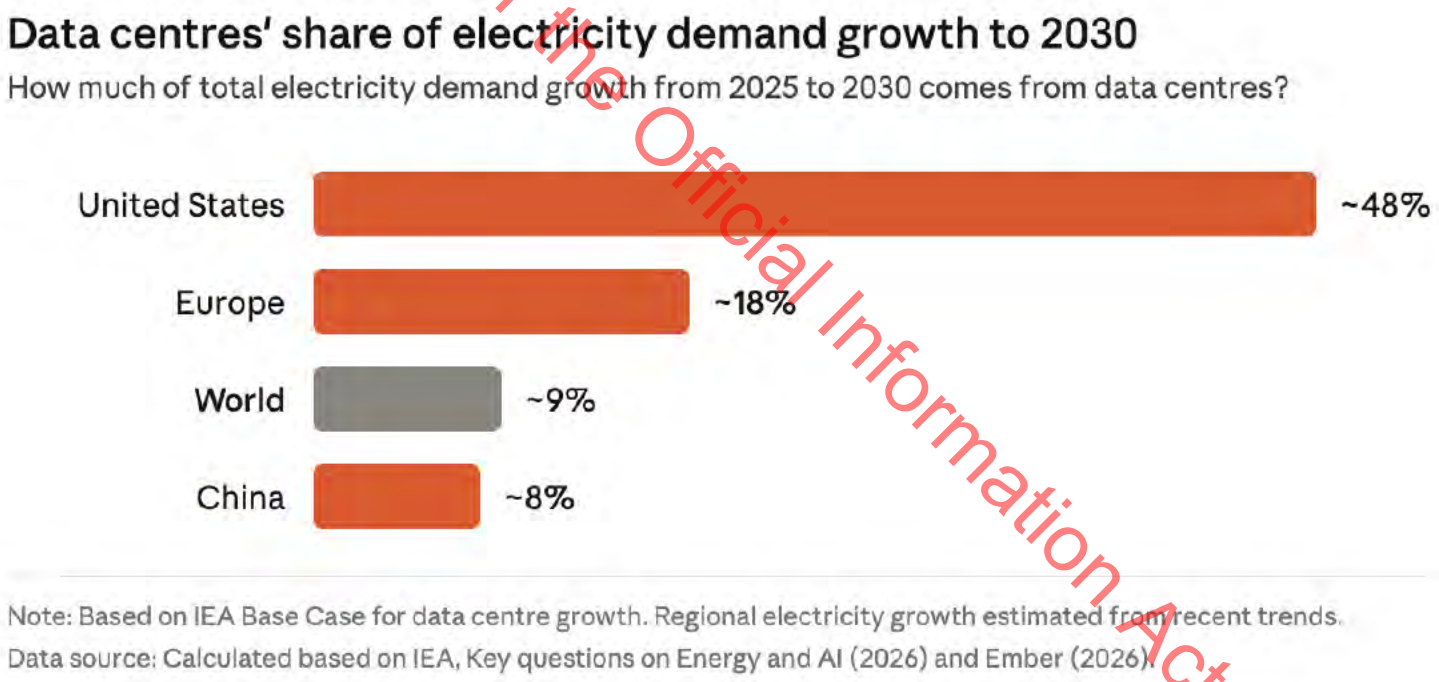

Released under the Official Information Act 1982

Data centres are expected to be a significant source of electricity demand growth, globally

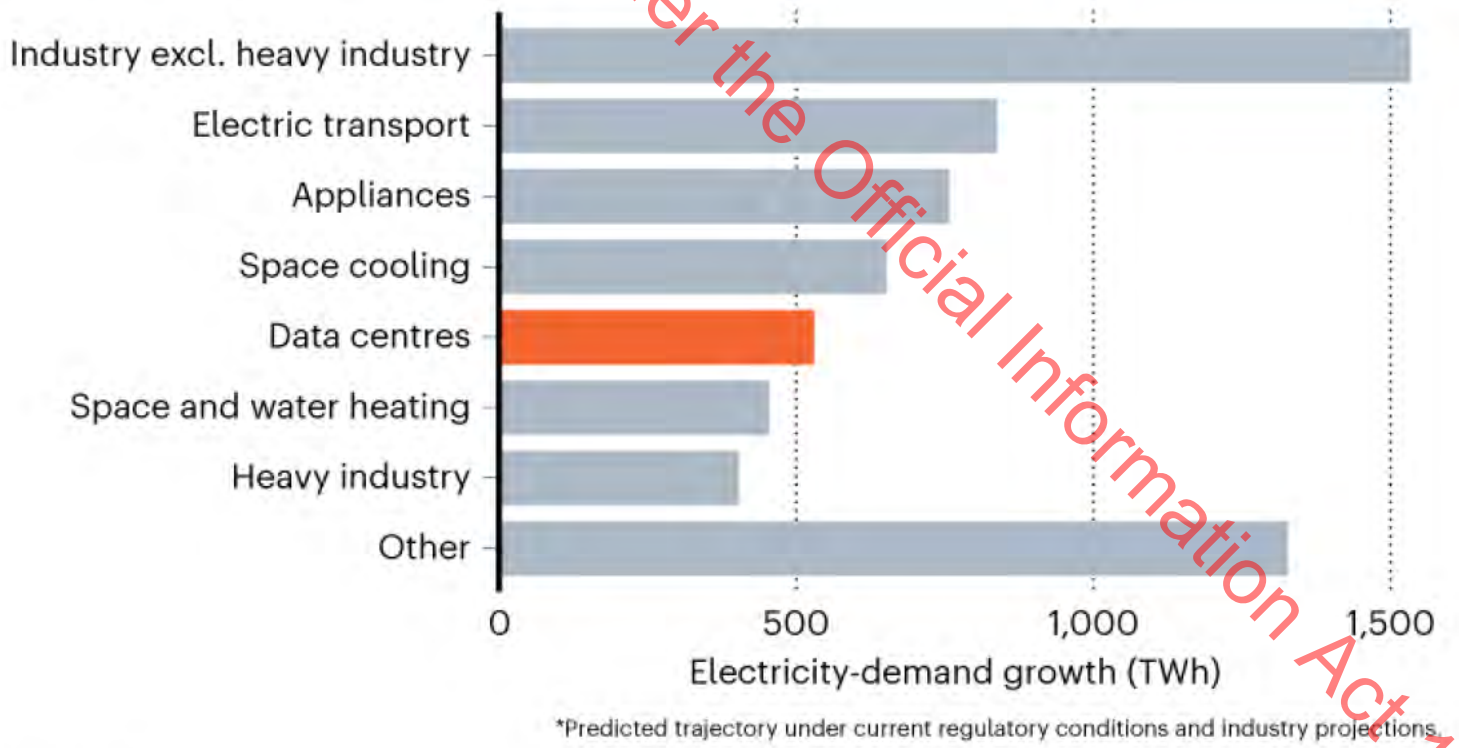

Released under the Official Information Act 1982

What powers them?

Released under the Official Information Act 1982

Figure 5.4 ▷ Global electricity generation for data centres, 2020-2030

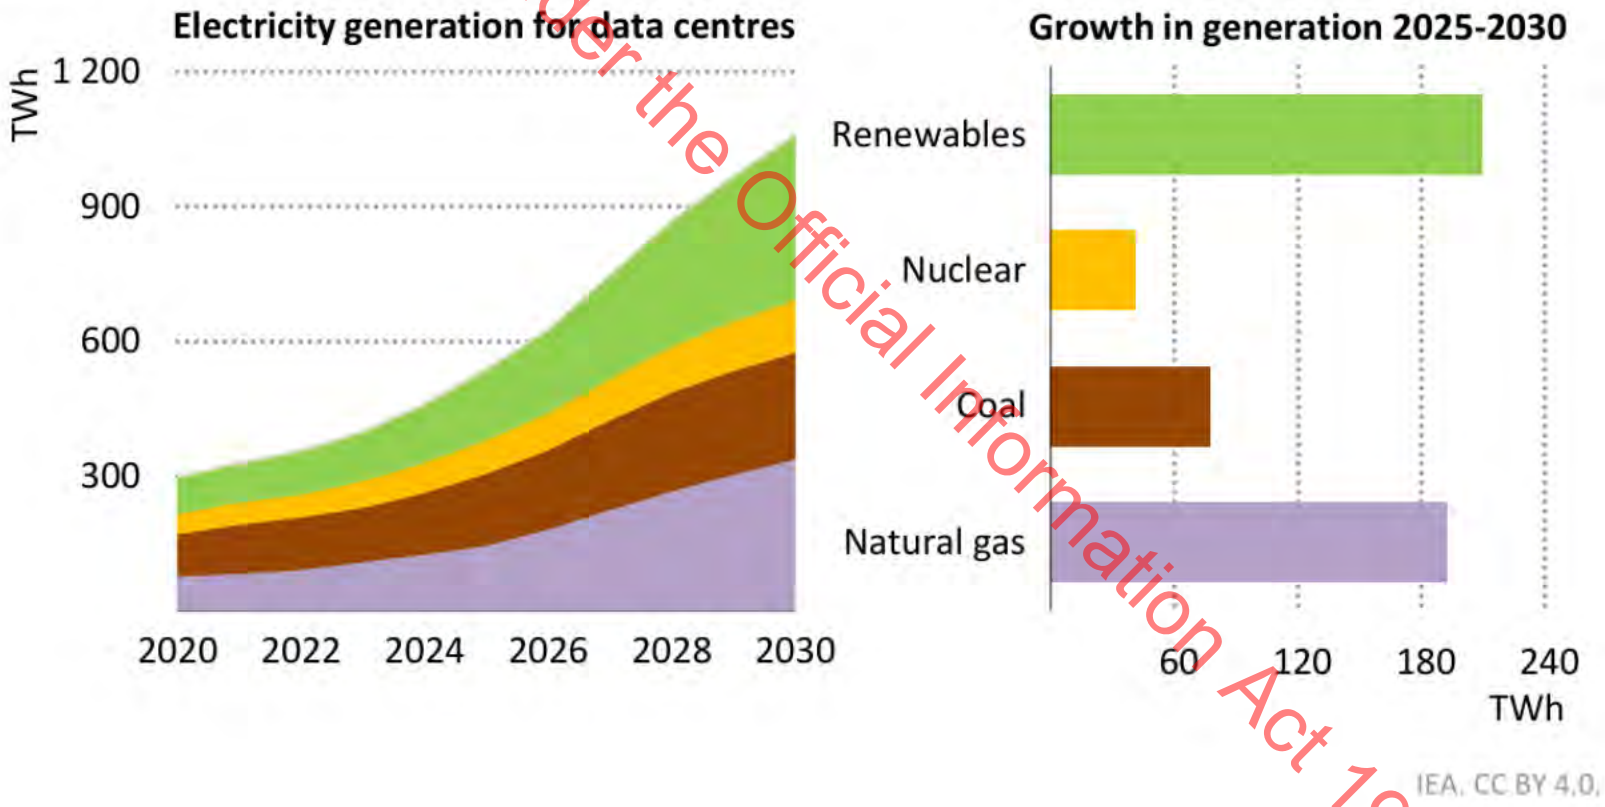

Why so much coal and gas?

Clustering in places with less-renewable grids.

- E.g., *Ireland*: 40-50% from gas; *USA* 55-60% from gas and coal.

AI companies also have a need for speed.

- They are racing to be the first to build superintelligence.
- Often it is quicker and easier to self-supply with gas peakers, than to wait for a grid connection, or firmed renewable PPAs to become available.

Released under the Official Information Act 1982

Most data centre waste heat is not re-used

- Virtually all electricity consumed by a data centre becomes heat
- Waste heat reuse is possible, but remains the exception rather than the rule.
- **~20-30** data centres globally are known to supply district-heating networks, compared with more than **1,000** hyperscale data centres worldwide.
- Typically requires a large nearby heat demand and, in many cases, district-heating infrastructure.
- And additional investment in heat recovery and distribution systems.



26 June 2024

Spark secures resource consent for data centre in world-first sustainable development in Auckland

Datagrid's 'waste heat' available for other Southland uses

Michael Fallow | The Southland Times
June 15, 2026 - 4:00am

Released under the Official Information Act 1982

AI workload efficiency has increased significantly

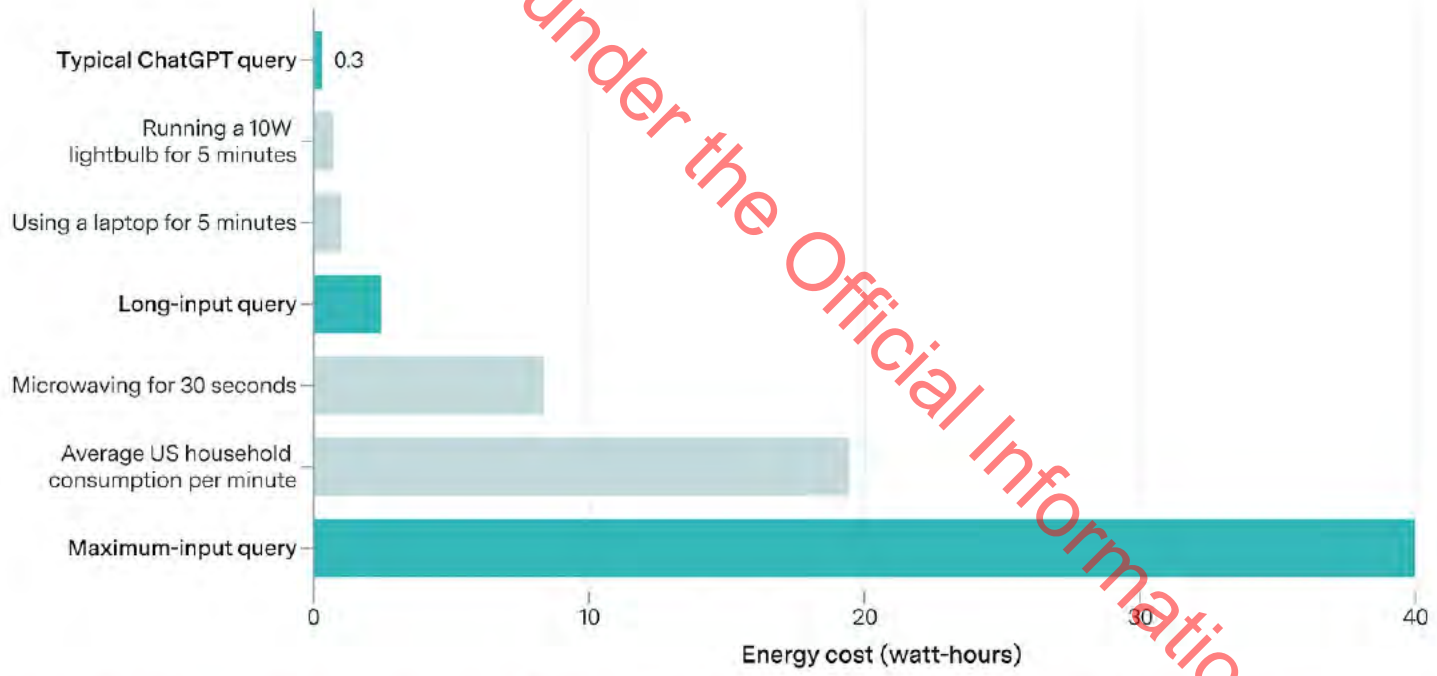

Pessimistic estimates of the energy usage of ChatGPT with GPT-4o across for different query lengths: typical (<100 words), long (~7,500 words), and maximum context length (~75,000 words), with an average response length of 400 words.

CC-BY

epoch.ai

Released under the Official Information Act 1982

If efficiency has improved, why do AI-focused DCs still use so much electricity?

People are using AI services more…

- Jevon's Paradox: Making a resource more efficient increases consumption.
- *ChatGPT is now the fifth-most visited website in the world.*

..and for more computationally-intensive tasks, too.

- E.g., video generation, reasoning and agentic tasks.
- *"These kinds of tasks can consume **hundreds** or **thousands** of times more energy per query than simple text generation." – IEA*
- Chat GPT is used to generate 78 million images per day.

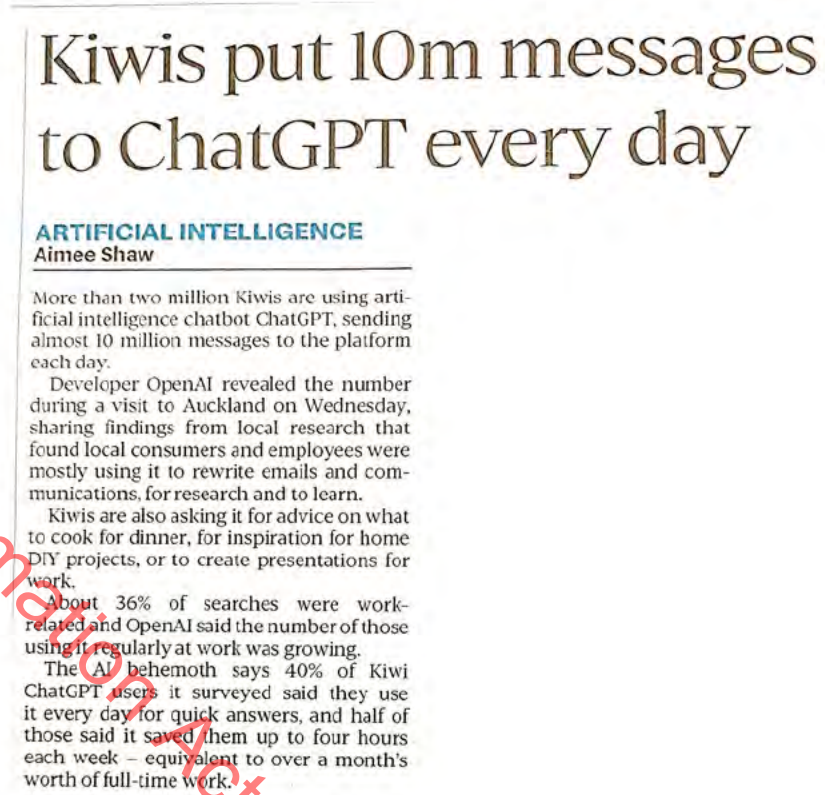

Kiwis put 10m messages to ChatGPT every day

ARTIFICIAL INTELLIGENCE
Aimee Shaw

More than two million Kiwis are using artificial intelligence chatbot ChatGPT, sending almost 10 million messages to the platform each day.

Developer OpenAI revealed the number during a visit to Auckland on Wednesday, sharing findings from local research that found local consumers and employees were mostly using it to rewrite emails and communications, for research and to learn.

Kiwis are also asking it for advice on what to cook for dinner, for inspiration for home DIY projects, or to create presentations for work.

About 36% of searches were work-related and OpenAI said the number of those using it regularly at work was growing.

The AI behemoth says 40% of Kiwi ChatGPT users it surveyed said they use it every day for quick answers, and half of those said it saved them up to four hours each week – equivalent to over a month's worth of full-time work.

Released under the Official Information Act 1982

“AI” is a broad category. What is the electricity actually used for?

Back-of-the-envelope estimate shows that this is a difficult question:

- **1Wh** of electricity for a large text query *(i.e., an overestimate for ChatGPT)*.
- Multiply by **10 billion queries a day** (*4x what ChatGPT reports)*
- Then, chatbot queries would use **3.65 TWh** per year.
- But, AI-focused data centres consumed **155 TWh** in 2025 (IEA).
- So, all of the world’s text queries may have accounted for ~2% of this electricity consumption.
- Doubling, would mean, 4%, then 8% or 10% - still a small fraction.

Released under the Official Information Act 1982

The point

- Unlike other industries (e.g., aluminium) it is difficult to identify exactly what hyperscale datacentre electricity consumption is used for.
- Even under assumptions that likely overstate text-query demand, text-only chatbot interactions appear unlikely to explain more than a small % of total AI-focused datacentre electricity consumption.

Released under the Official Information Act 1982

If not text queries, then what?

Again, it's hard to say with confidence.

Much of the relevant information is commercially sensitive and not publicly reported.

An educated guess:

- Difficult tasks = more power draw.
- So perhaps, **images, video, agents, and training**.
- Or, **enterprise and industrial/scientific use of AI.**

"The entire industry is systemically opaque [...] about how models are used."

– Stanford Institute for Human-Centred AI, December 2025,

2025 Foundation Model Transparency Index

Released under the Official Information Act 1982

Are datacentres flexible?

Possibly, DCs could:

- **Self-supply for short periods,**
- **Pay others to reduce their demand (i.e., VPPs),**
- **or curtail**

Examples of DR:

1. Google DR agreements in the US
2. 'Conductor' software: reduced AI workloads "by 25% for three hours" while maintaining QoS.
3. Some US utilities offer faster grid interconnection, in exchange for flexibility.

Though, in practice, DCs are not inclined to be flexible.

The industry standard for uptime is 99.999%. ~5 minutes of downtime per year.

Released under the Official Information Act 1982

Load profile simulations (April 2026)

(note: these are estimates, not actuals)

SOURCE: Vercellino et al. (2026). Measurement of Generative AI Workload Power Profiles for Whole-Facility Data Center Infrastructure Planning. *arXiv*

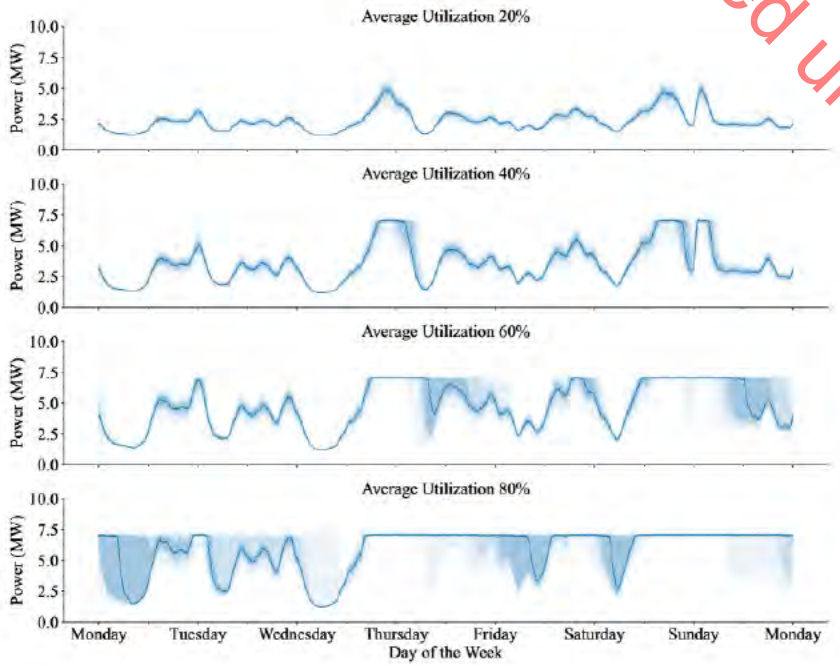

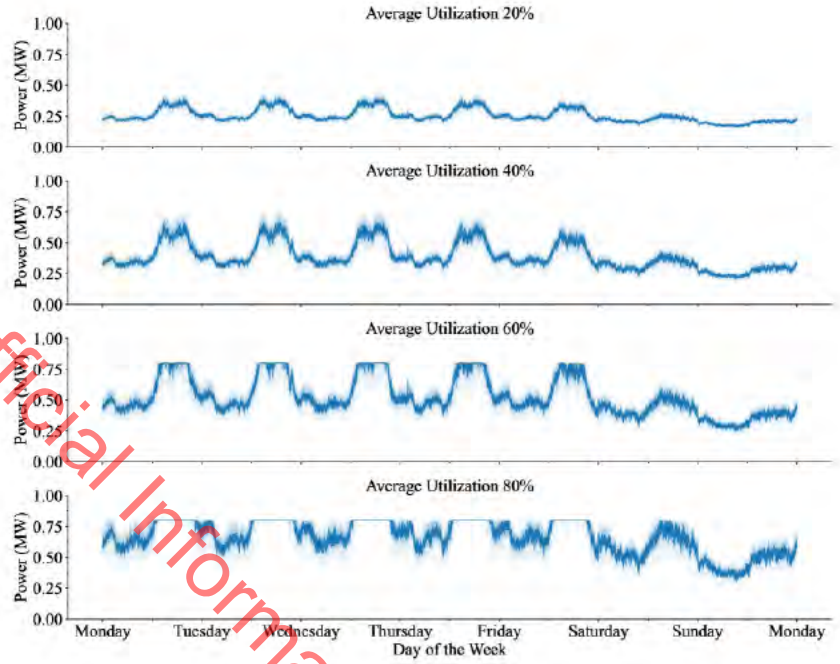

AI training workloads ('co-location')
– less predictable, but potentially easier to flex.
Though, timing of when to train models may also be commercially driven.

Customers using AI services in real time ('inference')
– more predictable, but harder to flex.
People expect near-instant service.

Released under the Official Information Act 1982

Of course, capacity =/= utilisation

- E.g., in the US in 2024, datacentres consistently used les have.
- Driven by global ramp up of new DC construction.

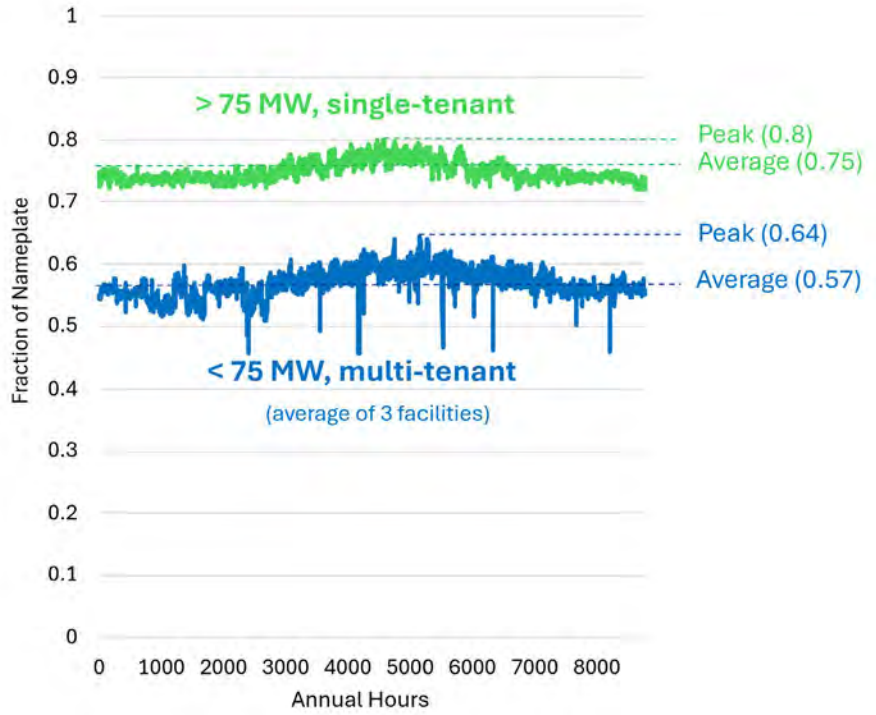

Though, utilisation rates are expected to increase in future.

- Will depend on demand for the services that DCs provide, and the number of them that are built around the world.

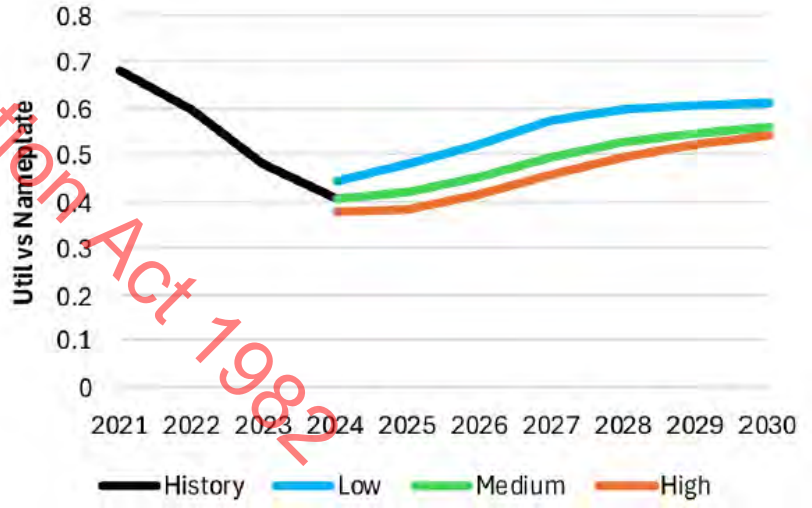

Released under the Official Information Act 1982

1) CDC + Contact

Released under the Official Information Act 1982

Contact has partnered with CDC to explore data centre development at its highly strategic Stratford site

Following the March 2026 closure of Contact's Taranaki Combined Cycle baseload gas plant (TCC), Contact is taking steps to leverage the unique combination of site resources and support growth in the Taranaki region

Contact and CDC will seek resource consent for a 250MW data centre[1]. It would be supported by co-located battery storage development. Current activities focus on consenting, site and infrastructure design, customer engagement and development planning.

Solar and battery development options



150MWac / 312GWh p.a. solar farm in consenting, with DC coupled batteries potentially providing up to 750MWh storage.[2]

500MW grid-scale battery capacity consented.

High-capacity fibre connection



Diverse high-capacity fibre, both terrestrial and subsea, provides connection to Auckland where New Zealand's international cables land.

Existing generation



The operation of Contact's existing 200MW fast-start gas peaking assets at Stratford is not impacted by the data centre concept.

.

Available transmission capacity



350MW transmission capacity available within 1-2 years with minimal substation work.

600MW expected to be available within 2-3 years following planned upgrades.[3]

Land owned and under option



Existing site complemented by significant adjacent land options.

Able to support scalable renewable and load growth opportunities.

Long-term member of the community



50 years operating in the Stratford community.

Long-standing relationships with local stakeholders including councils and tangata whenua.

2) Datagrid + Mercury (Mahinerangi wind farm expansion)

Released under the Official Information Act 1982

Mercury to supply half AI factory power needs



Felicity Wolfe
Thu, Mar 19 2026



Mercury will supply about half Datagrid's power needs for its new "AI factory" planned for Southland, with announcements on further supply agreements "when the time is right".

Datagrid says its 140-megawatt power-purchase option agreement with Mercury provides a long-term hedge for its $3.5 billion data centre project.

The hyperscale facility, to be built north of Invercargill, will have a total load of 280 MW once completed in late 2028.

The 15-year agreement provides "price certainty" for 1.2 terawatt-hours each year, which the firm notes is about 3 per cent of national demand.

NEWS | MAHINERANGI

Fast-track approval for Mahinerangi wind expansion



Eamon Rood
Mon, Jul 6 2026



Mercury NZ has received final fast-track approval for its Mahinerangi wind farm in Otago.

Puke Kapo Hau will comprise up to 44 turbines with an installed capacity of up to 190 megawatts and generate about 550 gigawatt-hours per year - enough for about 78,000 homes.

The expert panel says it offers significant national and regional benefits by boosting renewable energy supply and construction job creation in a decision dated 3 July.

Mercury expects to invest between $500 million and $600 million on the project and begin construction by the end of the year, expected to take three years.

3) Tiwai Potline 4 + Contact (Southland Wind Farm / SWF)

Released under the Official Information Act 1982

A pathway is in place for NZAS to act as long-term customer underpinning Contact's TTWh+ Southland Wind Farm

Contact is advancing strategic partner identification and targeting mid-2027 for Final Investment Decision



Electricity stats

Up to
55 Turbines

>325MW
total capacity

Average annual output expected to be
>1,210GWh

Project updates



- **Consent approved** April 2026.
- **Specialised infrastructure advisor, Mafic, appointed** to run identification and selection process for a strategic partner for Contact's extensive wind pipeline.
- We have engaged a shortlist of **credible wind partners.**
- Non-binding letter of intent (LOI) signed with Rio Tinto for a PPA to support the **potential restart of 50MW Line 4 potline at NZAS.**

INTERNATIONAL SCAN



Released under the Official Information Act 1982

Several jurisdictions already have DC-specific regulations in place

Mainly focused around renewables, demand response, security of supply and managing the pace of growth



Ireland

- Moratorium on new DC connections between 2021 and 2025.
- Since 2025, DCs required to provide on-site or proximate capacity, equal to their max capacity. Of this, 80% must be domestic renewables.
- Supply options are de-rated, based on their capacity factor. E.g., gas turbine = 90%, wind farm = 10 or 20%.

USA, PJM (e.g., West-Virginia)

- Incentives-based: DCs encouraged (but not required) to bring new supply, or provide flex.
- Facilities that don't *may* be curtailed first during emergencies.

USA, New York State

- One year moratorium on DCs over 50MW from July 2026
- Will lift once state finalises new standards

Released under the Official Information Act 1982

Several jurisdictions already have DC-specific regulations in place

Mainly focused around renewables, demand response, security of supply and buying time

USA, ERCOT (Texas)

- Requirement-based: DCs must self-supply or flex during emergencies.
- Facilities that don't *will* be curtailed during emergencies.

Singapore

- DCs required to source at least 50% of their electricity from low or zero carbon sources.

Germany

- DCs required to source at least 50% of electricity from unsubsidised renewable energy.
- Will increase to 100% in 2027.



Released under the Official Information Act 1982

Some countries are considering DC-specific regulation



UK

- Considering whether data centres should be required to provide a minimum level of flexibility as a condition of connection

Australia

- Developing obligations to require DCs to underwrite their own new *firmed* power supply, pay their full share of connection costs, and be flexible.
- Developing govt. 'AI Office'

Denmark

- Moratorium on new connections in March 2026, due to an "explosion" of new requests
- (14GW of DC projects applied for connections, compared to 7 GW of peak electricity demand.)
- Now reviewing grid-access rules; considering prioritisation of grid capacity and stricter connection rules.

Released under the Official Information Act 1982

Ireland

Marginal cost effect:

- "Our modelling shows that the high, growing and inflexible nature of datacentres' electricity demand increases the number of hours in which gas sets the price in the Irish power system."

~$88NZD more per year for Irish households.

Analysis disputed by industry.

'The cost of data centres: Modelling the impact of data centres on electricity bills for Irish households'. Friends of The Earth Ireland, 2026

Released under the Official Information Act 1982

Capacity prices in PJM

Scarcity effect:

- PJM is the grid operator serving 14 mid-Atlantic and Midwest states.
- According to PJM's analysis, 63% of the 2025/2026 jump can be attributed to data centre load growth.

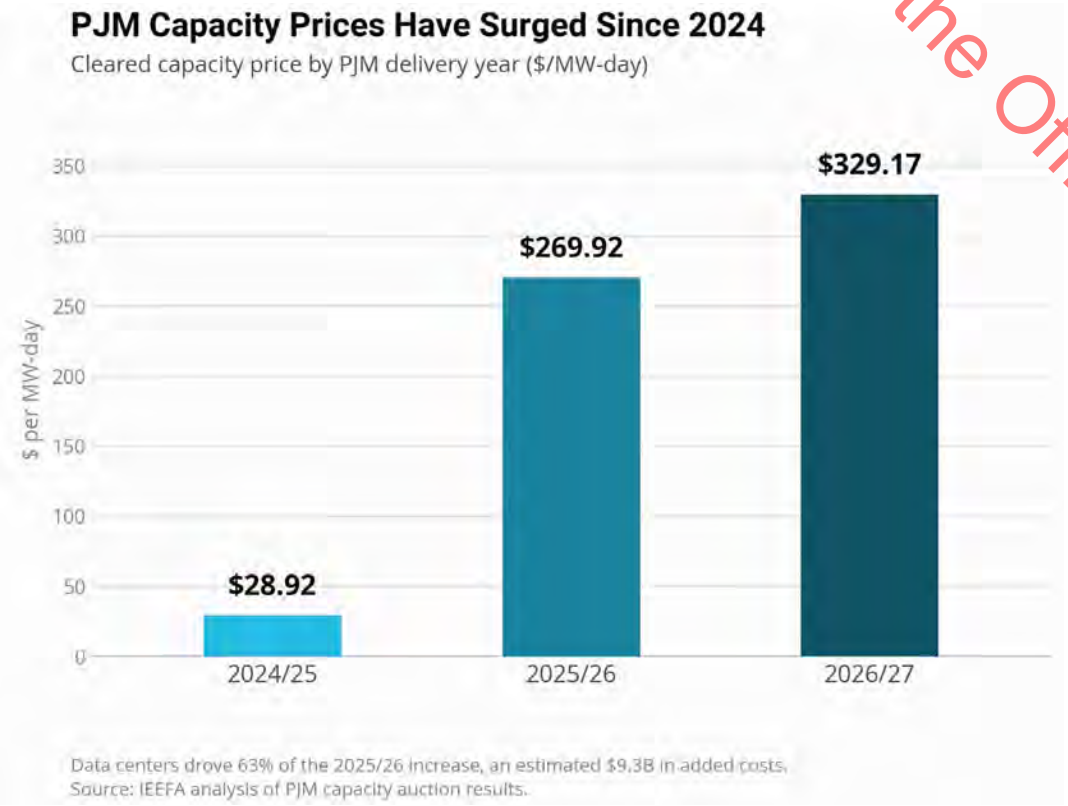

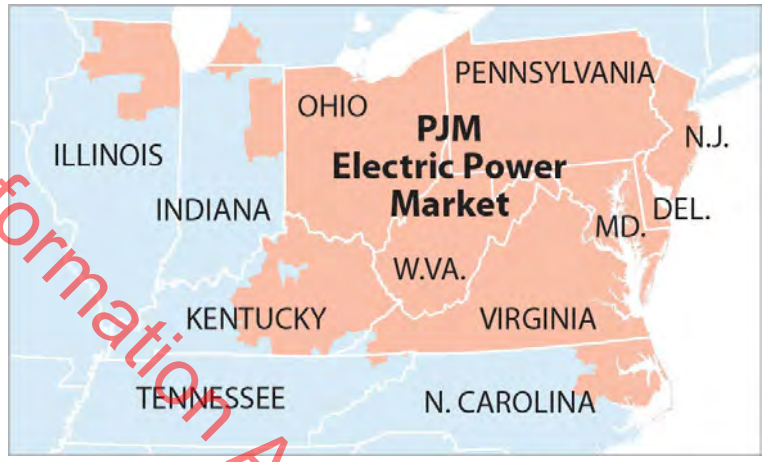

Released under the Official Information Act 1982

Cost sharing effect in three American States

- North Dakota, New Mexico and Nebraska.
- "Substantial commercial and industrial **growth enabled fixed costs to be spread over more load.**"

– Lawrence Berkeley National Laboratory, July 2026

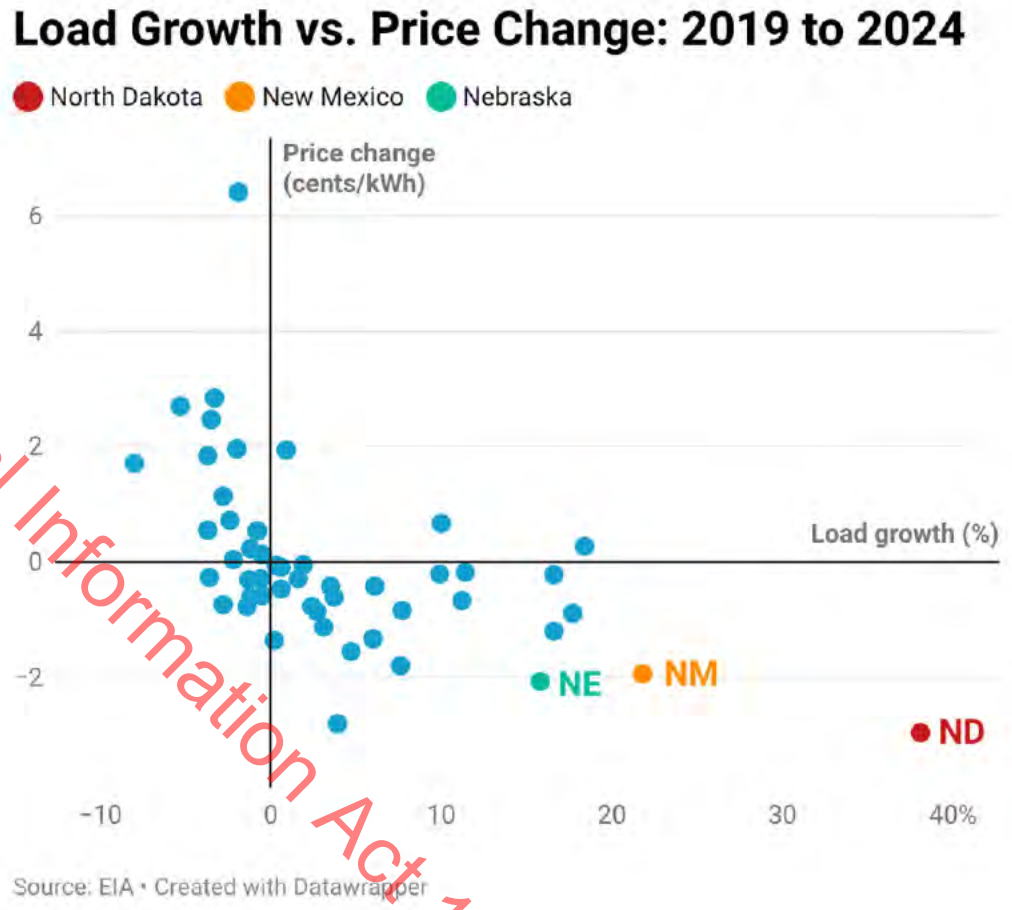

Released under the Official Information Act 1982

Datacentres face strong political and social license concerns

A uniquely bipartisan backlash against data centers shakes up the midterms

David Wright, Casey Tolan, and Isabelle Chapman for CNN

More than $31 million (NZ$52m) has been spent on political ads mentioning data centres so far this year, spiking last month, when more than 8 percent of all campaign spending on broadcast TV went to data centre or AI-related spots, according to tracking from the firm AdImpact. That's up significantly from less than 1 percent of ads last year.

Big tech fights back

AI companies and industry allies are also lining up to shift the political momentum before the midterms.

Tech giants like Meta, Google and Anthropic, a prominent venture capital firm, and top AI executives have donated more than $190 million to super PACs and nonprofits that have spent over $60 million on independent expenditures in the midterm elections so far.



Released under the Official Information Act 1982

Similar concerns are felt in New Zealand

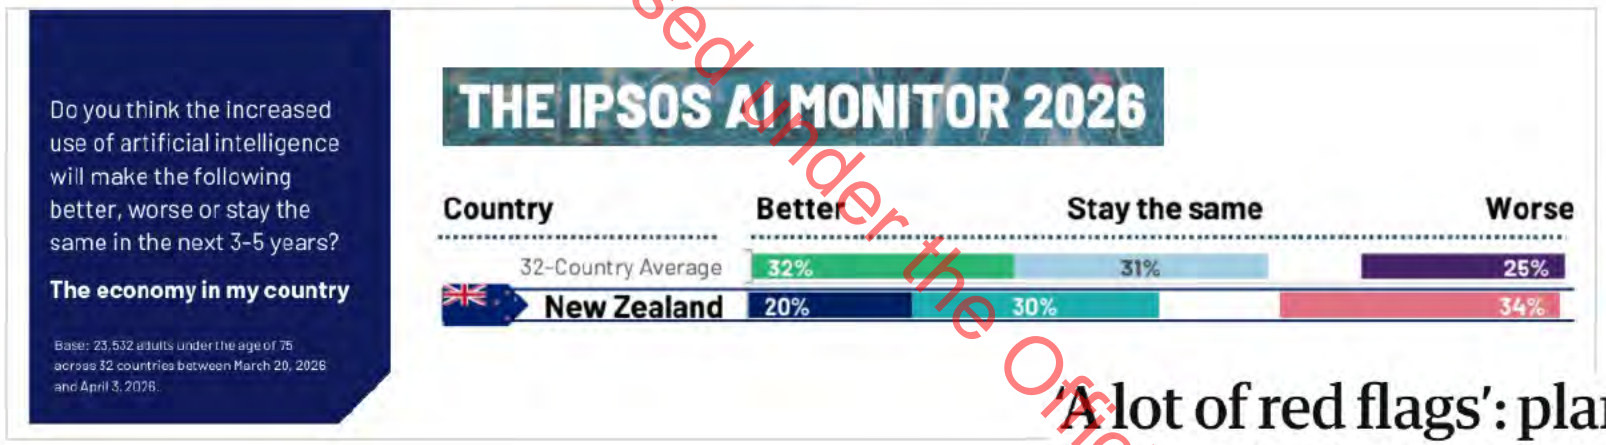

'A lot of red flags': plans for New Zealand's first AI datacentre spark concern as locals demand greater transparency

Plans to build a NZ$3.5bn datacentre in Makarewa in the country's south has drawn concern about electricity and water use, and potential noise pollution

Datagrid has approval to discharge air contaminants from up to 84 diesel back-up generators, to draw up to 604,800 litres of groundwater per day, discharge up to 5000 litres of treated wastewater a day and remove a nearby wetland, an Environment Southland regional council reports show.

Released under the Official Information Act 1982

Simeon Brown, Minister for Energy:
"*We're going to need these in our country. […] But I don't want them pushing up the power bills of ordinary Kiwis.*" (9/8/26)

Chris Bishop, Minister for Infrastructure:
"National will require new data centre projects to be accompanied by new, firmed generation." (17/8/26)

Released under the Official Information Act 1982

DC-specific policies of the main NZ political parties

National:

- "Will require an 'additionality test' new data centre projects must be accompanied by new, firmed generation, protecting households and businesses from higher power prices."

Greens:

- "Calling for a one-year moratorium on new AI data centres while proper rules are put in place to ensure they genuinely benefit New Zealanders and are not straining our resources."

Labour:

- Would pursue action "very similar" to what Australia is doing. Open to data centre investment, with a focus on appropriate regulation and system planning.

TOP, ACT, TPM, NZF:

- No data-centre specific policy identified (as at Sept 2026).

Released under the Official Information Act 1982

Framework for thinking about impact of datacentre demand growth on wholesale prices

	Increase	Unaffected	Decrease
1) Marginal cost effect *(i.e., what resource is setting the price)*	**Marginal Cost ↑** Higher-cost supply is required to meet demand more often and therefore sets the wholesale price more frequently than it otherwise would. Load may be met by high-cost existing capacity, or new capacity build.	**Marginal Cost ↔** Additional supply is broadly right-sized to DC load growth.	**Marginal Cost ↓** Lower-cost supply sets the price more often than it would have, because DC load growth induces investment in excess of its own requirements. Could be current requirements, or eventual.
2) Scarcity and risk premium effect *(i.e., how scarce is firm energy)*	**Scarcity Premium ↑** Firm energy becomes scarcer. Under-build of firm energy resources, and/or lack of DR from DCs.	**Scarcity Premium ↔** Firm energy scarcity is unchanged. Additional firming is broadly right-sized to DC load growth.	**Scarcity Premium ↓** Firm energy becomes less scarce. Over-build of firm energy resources, and/or sufficient DR from DCs.

Released under the Official Information Act 1982

Framework for thinking about impact of datacentre demand growth on network costs

The impact on network costs will depend on:

- Existing asset utilisation and
- whether grid upgrades are required.

	Increase	Unaffected	Decrease
1) Cost-sharing / load factor effect *(i.e., will new DCs share in existing fixed costs)*	*N/A or unlikely.*	*N/A or unlikely.*	**Network charges ↓** Large load will share in the existing fixed costs of network infrastructure, reducing such costs for other consumers. Effect would occur at any level of demand growth.
2) Grid upgrade effect *(i.e., will grid upgrades be necessary? and if so, who pays?)*	**Network charges ↑** Upgrades required to accommodate new load. Some costs may be allocated to existing consumers if they are assessed as receiving net private benefits*. E.g., avoided congestion, lower nodal price differentials, or improved reliability. However, may be seen as hollow – paying more for the same LoS. Because the large load would likely receive most of the benefits, most upgrade costs are expected to be allocated to that load. ** This effect would occur regardless of the source of load growth (e.g., electrification, population growth).*	**Network charges ↔** Upgrades not required. May occur if large load makes limited use of existing network capacity.	*N/A or unlikely.*

Released under the Official Information Act 1982

Wholesale prices

The impact on wholesale prices of datacentre demand growth will depend on:

- The scale and pace of demand growth, and
- what additional supply and firming resources are built or utilised, and how soon.
- **Two possible mechanisms:**
  - Marginal pricing (which generator sets the marginal price)
  - Risk premia (market expectations of scarcity).

Released under the Official Information Act 1982

Potential to create an energy gap

If:

- datacentre demand growth remains as expected (perhaps ~400 MW by 2035)
- generation expands (Transpower expects an additional 3.5 GW (5 TWh) of new generation in its connection queue to be delivered by 2029)

Then:

- Datacentres should not place direct upward pressure on electricity prices in the near term.

However, firming is an area of concern:
83% of projects in the generation pipeline are intermittent, while DC customers are typically inflexible. Rapid demand growth (from any large customers) for firm generation could outpace new generation and put upward pressure on prices.

Demand could also grow beyond what is currently expected:
Particularly if Government policy actively aims to establish NZ as a data centre hub. Or significant additional industrial electricity demand occurs.

Released under the Official Information Act 1982

Network costs

Released under the Official Information Act 1982

Existing asset utilisation could allow for wider cost sharing

Large new loads will share in the costs of existing network infrastructure, reducing costs for other users.

New assets

Further grid upgrades could be required in future to accommodate growth (from any source, not just data centres).

Current connection processes help to identify the size and timing of new demand, allowing coordination with (potentially large) network upgrades.

Transpower / EDB performance standards also provide them with incentives to manage the connection of large loads.

Beneficiary pays

Large loads will face cost-reflective (benefit- or service-based) network and connection charges. I.e, pay for their own connection.

- Large load thus has incentives to locate where there is network capacity (or planned), balancing this factor against its other objectives such as market access

This means that most costs will likely fall on new large loads. Though some would also fall on existing consumers – maintaining an existing level of service / reliability (vs. a counterfactual where no upgrade occurs) is a private benefit within the TPM.



Power system quality

AI training workloads can be highly volatile

- AI training workloads can ramp electricity demand up or down rapidly as computing jobs start, stop, or shift between clusters.
- Rapid changes in demand can contribute to voltage fluctuations, system strength issues, and oscillations.
- Large data centres may need technical performance requirements and information-sharing obligations to support secure system operation.

Batteries can help, but are not required

- On-site batteries (BESS) can help support system stability, and assist recovery following network disturbances.
- A key policy question is whether large data centres should be required to mitigate their power system impacts and support the grid during emergency conditions.

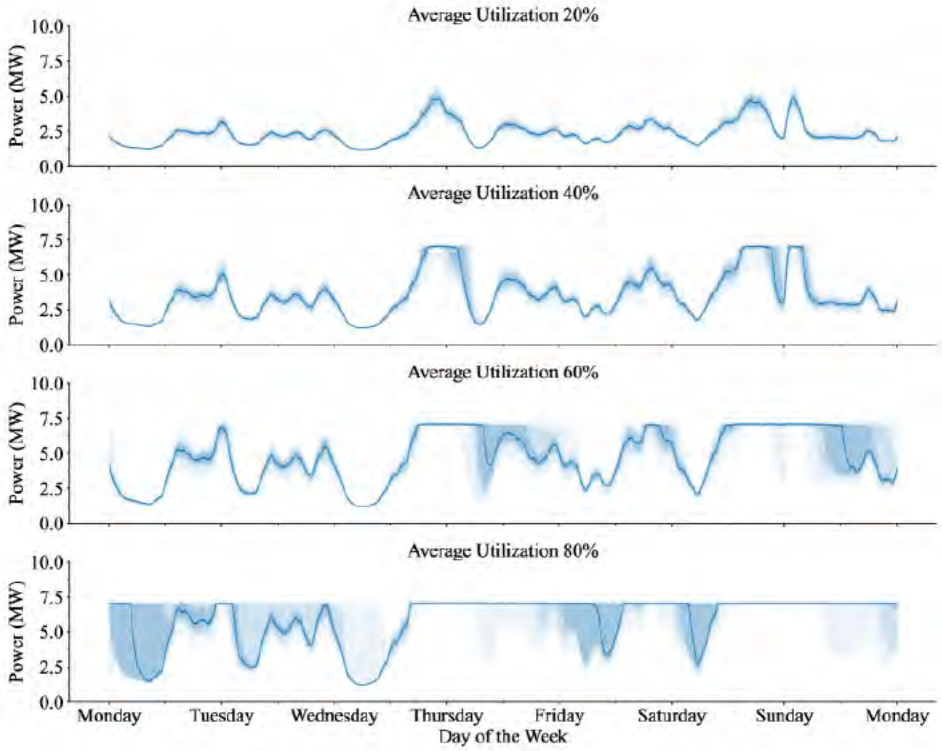

Indicative AI training workloads ('co-location') – less predictable, but potentially easier to flex. Though, timing of when to train models may also be commercially driven.

(note: these are estimates, not actuals)

SOURCE: Vercellino et al. (2026). Measurement of Generative AI Workload Power Profiles for Whole-Facility Data Center Infrastructure Planning. *arXiv*

Released under the Official Information Act 1982

An existing tool - Materially Large Contracts (MLC) provisions

A partial safeguard against cross-subsidization.

Aimed at preventing large customers (150MW+) being offered discounts that are effectively funded by other consumers.

Parties can apply for clearance from the Authority, if:

- the net value of the contract to the generator is positive, or
- on-selling by the large customer is allowed.

s 9(2)(g)(i)

Released under the Official Information Act 1982

Data centres and network pricing

Key points

This note addresses the question whether a data centre should pay for all of the new network investment it causes (rather than costs being shared based on net benefits). It touches on other topics for context, but the focus is on network pricing matters.

Under current pricing approaches and rules, new large data centres will pay for their own connection costs and a share in the costs of network services (for transmission: residual charge for historical assets, benefit-based allocation of newer grid assets). This reduces the share of existing network costs for all other customers.

But new large data centres may also drive upgrades of networks. Whether these then increase local network costs for other customers depends on an EDB's pricing policies, though charges are expected to be cost-reflective and subsidy free.

Transmission costs would be shared with other customers on a benefit-basis. Customers may object they must now pay a benefit-based share of costs that would not have been incurred without the data centre.

For example, they may object that without the data centre there would not have been congestion or reliability concerns the avoidance of which is presented as a benefit to them of new transmission investment.

However, this would also happen if local demand increased for other reasons. For instance, if a large plant electrified, population grew, or economic activity that raises electricity demand increased.

The purpose of the TPM is not to insulate or protect existing customers from the (near term) effects of economic or other societal changes, but to allocate costs for beneficial and efficient transmission investments that promote the efficient operation of the electricity industry, for the long term benefit of consumers.

Changing pricing rules to treat large data centres different to other sources of demand growth (allocating them the full cost of interconnection investment as some have suggested) would create uncertainty for investors in the electricity industry.

Treating data centres different to other customers when charging for network services also increases the risk of later disputes, including threats to exit (and leaving other customers to pay for the cost of white elephants) if customers do not get a special deal. Such

Released under the Official Information Act 1982

arguments will be harder to resist if charges are based on (disadvantageous) differential treatment.

s 9(2)(g)(i)

Released under the Official Information Act 1982

Context: Data centres in NZ

There is an international data centre investment boom driven by the AI race. Their actual and potential impacts on electricity costs and environment is raising questions on whether policy and regulatory settings are suited to the entry of very large electricity consumers.

There are already well over 50 data centres operating in NZ. However, hyperscale data centres would be quite different in size and impact.

Some project that data centre capacity will grow 3-6 times by 2035 to meet domestic demand and add export capacity. Transpower's Te Kanapu scenarios suggest data centres could account for between 2-6% of total demand in 2035.

This projected growth builds on New Zealand's advantages in renewable energy, cooler temperatures, water, and robust connection process compared to the rest of the world.

Location decisions depend on their purpose, supporting infrastructure, and latency:

- mostly near population centres: connectivity (fiber optics), low latency, access to skills
- near energy and cooling: data storage and AI training can be more remote (eg Datagrid).

Further context: high level implications for the electricity sector

This section seeks to paint a high-level picture of the types of issues that may arise, given size and international experience.

Data centres are energy-intensive large load. Competitive energy, grid and network costs will be key considerations to their investment case.

Hyperscale data centre investments do have incentives to secure generation contracts to manage their wholesale market price risk. Where new generation is required to support this (step change in) demand, PPAs etc can underwrite generation investment.

Even though as large load data centres have obvious incentives to avoid elevated energy costs, there are still near-term electricity wholesale price risks if supply lags as other markets have found. The market should be able to solve mismatches over time though given current projections of hyperscale data centres.

It is difficult to be definitive even about near-term impacts, given the ever-changing aggregate demand and supply mix (eg a data centre enters but timber mills close).

Data centres are inverter-based load, and grid disturbances could cause it to disconnect. Any data centre has incentives to manage its own power reliability requirements and outage risk (eg batteries). But sudden disconnection of one very large load or several smaller centres at the same time could cause grid instability and cascading outages with far wider consequences.

The EA has a future security and resilience work programme underway, which includes the consideration of such externalities.[1]

Data centres, like other large load, may also drive grid and local network upgrades. Note that this also happens when Fonterra or Kiwirail electrify their operations, for example, or as more people buy electric vehicles, or there is population growth.

Grid and local network supply is less responsive to step changes in demand, due to regulated planning, consenting, and funding cycles, and time to build infrastructure.

There is a risk that network capacity lags investment in data centres as has been the case in other countries, but:

- these are very large plant – existing connection processes will help identify demand and its timing, to help coordinate (potentially very large) grid/network upgrades with data centre investment
- Transpower's performance standards (code requirements, price-quality path incentives for quality of service) provide it with incentives to manage the connection of large load
- benefit- or service-based allocation of network cost means data centres face costs of their investment decisions (though with benefit-based transmission charges others might face costs too, in proportion to positive net private benefits, as discussed below)
- Transpower or distributors could develop capacity in anticipation, though at risk of stranded assets for other customers to pay if a data centre does not materialize. Given the potential size of upgrades, this puts a premium on the value of scrutiny of such transmission investment proposals.

[1] AEMC is proposing new grid standards for large data centre connections.

Released under the Official Information Act 1982

These caveats are not to minimise the potential impact of hyperscale data centres, but to put them in context of incentives and processes that seem likely to blunt some of the more extreme outcomes that can be imagined.

Network charges – cost-sharing of existing and new investments

New large data centres pay for their own connection costs.

They also reduce the share of *existing* network costs borne by other consumers, as data centres will pay a share of existing fixed infrastructure costs (note, for transmission this includes both residual charges and benefit-based charges).

However, the benefits (transfers) of reduced bills for existing customers may be offset by their share (if any) of the costs of grid and local network upgrades that are driven by data centres.

s 9(2)(g)(i)

Costs of upgrades and expansions are allocated to customers who benefit from these investments (benefit-based allocation of transmission charges, service-based allocations in distribution lines and connection charges). Charges are to be subsidy free. [2]

Transmission interconnection assets are allocated on a benefit-basis. When a grid upgrade serves to reduce elevated nodal prices or outage risks, 'caused' by data centre using up remaining capacity, then pre-existing consumers will be identified as among the beneficiaries of the upgrade. This would be reflected in their share of costs.

Yet the constraints, the impact on electricity prices and reliability, and the grid upgrade would not have happened without the data centre entering. Existing customers may thus ask in what sense they benefit from the upgrade and why they should contribute at all.

[2] .... Expand on what AEMC is proposing to require distribution companies to recover costs of transmission upgrades from data centres connecting at the distribution level to protect small customer...

Released under the Official Information Act 1982

However, this dynamic is not something unique to data centre investment. It happens whenever there is load growth that uses up network capacity for whatever reason, eg population growth, electrification by industrial plant, or rising EV uptake.

Any subsequent transmission investment is then a response to the new circumstances. And transmission charges should be allocated based on the benefits the investment provides, given these new demand and supply circumstances. Existing customers benefit from the upgrade because without it they would have faced reduced service levels. They may scrutinise the investment case and submit such benefits are not worth the charges.

The entry of new large load and its impacts on electricity prices and grid performance is just the economy in action. The impacts are signals that resources have become scarce. They provide information and incentives for responses by demand and supply. This is fundamental to allowing scarce resources to be allocated to their highest value use.

Network prices are to be cost reflective, so participants face the resource cost of their actions, and cost allocation is to be service- or benefit-based.

For interconnection costs, this generally does not mean that the marginal user should face the full cost of an interconnection investment as such investment tends to create capacity and benefits well beyond the marginal user. Allocating the full cost to them could then discourage efficient investment.

Durable, efficient network pricing settings for the long-term benefit of consumers

The current amount of funds supporting the AI race might suggest that data centres are insensitive to costs (and some may be willing to pay a premium for social licence).

There have been suggestions that data centres should pay the full cost of network costs driven by their entry, because they may be price incentive or are seen to impose costs on others.

But it is unlikely data centres will be insensitive to costs for long. Ultimately, data centres (like any other industrial demand for electricity) are attracted by competitive energy cost (inc network costs), access to talent and markets, and strong institutions.

Network charges are expected to recover costs efficiently, including by being cost reflective and subsidy free. They should not be designed to insulate or protect existing customers from the effects of economic or other societal changes, but to efficiently allocate costs of network investments that provide services with benefits greater than their costs.

Released under the Official Information Act 1982

Changing pricing rules to treat large data centres different to other sources of demand growth (allocating them the full cost of interconnection investment as some have suggested) would create uncertainty for investors in the electricity industry.

Differential treatment is also likely to increase the risk of later disputes, including threats to exit (and leaving other customers to pay for the cost of white elephants) if customers do not get a special deal. Such arguments will be harder to resist if charges are based on (disadvantageous) differential treatment.

The durable approach for network pricing therefore is to apply consistent principled and efficient pricing – charges that are cost reflective and subsidy free.

END

[Relevant content from this draft note has been included in draft slide pack]

Released under the Official Information Act 1982

From: Natalie Bartos
To: Caspian Leah
Subject: FW: Discussion with AEMO on large loads
Date: Wednesday, 26 August 2026 12:25:38 pm
Attachments: image001.png

UNCLASSIFIED - INTERNAL

Copy FYI

UNCLASSIFIED - INTERNAL

From: Natalie Bartos
Sent: Friday, 31 July 2026 12:35 pm
To: Andrew Marriott <Andrew.Marriott@ea.govt.nz>; Hamish McKinnon <hamish.mckinnon@ea.govt.nz>; Charlie Chrystall <Charlie.Chrystall@ea.govt.nz>; Chris Otton <Chris.Otton@ea.govt.nz>; Kirsty Hutchison <Kirsty.Hutchison@ea.govt.nz>; Doug Watt <Doug.Watt@ea.govt.nz>; Dean Yarrall <Dean.Yarrall@ea.govt.nz>
Cc: Hayden Glass <Hayden.Glass@ea.govt.nz>; Sheila Matthews <Sheila.Matthews@ea.govt.nz>
Subject: Discussion with AEMO on large loads

Hi team

Data centres are the new hot topic!! Below are my rough notes from an Ops Policy/FSR /SO catch up with AEMO yesterday on their experience with data centres. Feel free to forward internally to anyone who may be interested. Also useful to note that the FSR team is looking at AOPOs for large loads, and Ops Policy is considering large loads as part our hybrids work.

Thanks
Natalie

30/7/26 meeting with Tim George, Andrew Ley, John Deere from AEMO

Context in Australia

- Largest data centre is ~130 MW
- Currently have around 5 GW of data centres with applications to connect, with up to 1 GW per data centre. Forecasting another 20-30 GW over the next 5 years
- Data centres currently make up ~2-2.5% of demand in the NEM. Projected to reach ~10% by 2050
- However, there is a lot of uncertainty. Expect only 1 out of 7 applications to

Released under the Official Information Act 1982

proceed

- Current connections are primarily in metropolitan areas (Sydney, Melbourne)
- Issue: can build data centres faster than transmission or distribution

Australian Government

- Concerned about capability of system to meet data centre demand
- Want data centres to procure renewable energy to offset demand and provide firming capacity or flex demand. Oz is currently developing this policy (quickly!). Also need to consider network upgrades. Considering legislation to require data centres to pay for their own upstream network augmentation costs to facilitate their connection
- Situation is more extreme overseas – eg, Ireland, Singapore, USA (Virginia, Texas)
- In Ireland, >20% of demand is data centres. Some US states have ~25%. Have had to take more extreme measures to bring generation online

Requirements for data centres

- AEMO is working with AEMC. Biggest concern is contingency events on transmission which can cause voltage drop over a wide area
- Current data centres use phase-locked loops and will switch off (will block rectifiers and move to standby battery supply) when there is a fault. Will come back online when the fault has cleared. Can happen between 0.5s to 5s. Hard to manage when there is a lot of load coming back online.
- Fault ride through and recovery time are important considerations. The longer the recovery time, the bigger the impact on the system.
- Some jurisdictions (eg ERCOT) are specifying recovery times. Oz is trying to align with international practice (it helps when all tech has similar standards). Looking at recovery time of 0.5s. However, not all manufacturers are capable of this but expect they will build this capability in over time (ie, starting to recognise the need to produce more grid friendly tech)
- AEMO is focused on inverter-based loads (rather than data centres specifically). Not focused on things like smelters, because they are generally stable loads and don't create oscillations.
- AEMO is looking at requirements for voltage ride through, frequency ride through (this is generally not a problem) and requirements for certain sized loads to have instability modelling (as they have a propensity to oscillate). Also considering requirements to mitigate any adverse impacts on system strength.
- Current draft rule has 3 tiers:
    - 0-30 MW – no material impact, don't need to know too much.
    - 30 – 100 MW – some discretion for network service provider/distributor to require the load to follow standards (eg fault ride through). Can choose

Released under the Official Information Act 1982

which standards apply
    - >100 MW – all standards apply
- Target date for final determination is late Oct.
- Recently identified that current measure of system strength for large inverter-based loads doesn't work. There are various ways that system strength is provided in Oz. Don't have a clear view on what sys strength requirements should be, currently working through this.
- Also looking at standards for monitoring and communications equipment for connecting loads. Information sharing for efficient market management is key factor.
- Looking at regulations so that AEMO can have direct voice or other comms channels with operators of large loads to manage power system issues/security supply. Don't currently have issues with aluminium smelters and the like because these are typically stable loads. When shifting load, the smelter will contact the control centre in advance (ie, it is not happening within milliseconds)
- Looking for data centres to provide forewarning of significant demand shifts – to help with system adequacy, reserves, dispatch
- Fast demand ramping – concerns regarding speed at which data centres can shift demand. AI training facilities can have millisecond shifts in power consumption up and down. Not good for network. Internationally, limits have been imposed on IB loads and demand changes. Oz not sure on limitations for their system yet, thinking about it.
- Data centres also have significant back up gen/BESS – can get them to operate on back up if needed to support system security
- Data centres can be developed in modules (not a homogenous block). Could end up in a situation where each module has a different OEM and behaves differently. Experienced similar things with rooftop PV.

Data centres and BESS

- Data centres with BESS. There are some connection applications which have large BESS associated with a data centre. But these typically have less than an hour/ half an hour of energy storage. Not suitable for feeding back into the grid, just for taking over supply to the IT equipment (eg due to voltage depression)
- BESS arrangements create some issues for the data centre in the sense that they need to be registered, have limited capability (2hr). Considering requirements for data centres to put in BESS rather than diesel gen for back up supply.
- Seeing some applications with solar/data centre/BESS config.(with solar charging the BESS). In Oz, each device has its own performance standards and needs to be registered in its own right
- BESS connected to data centres could potentially provide FCAS, but it is not

Released under the Official Information Act 1982

mandated. BESS is unlikely to be suitable to provide FCAS (just because of how they are connected)

Lessons

- Look at NERC and ERCOT as good models, Ireland too
- ERCOT has a lot of renewables (wind and solar) in their system, has similar problems to NEM. When looking at access standards to apply to data centres, they were 5-6 months ahead of Oz, looking at specifying ride through. Requirements are sensible, reasonable, achievable and took into account current batch of OEMs.
- NERC doesn't have the same specificity as ERCOT, but does have interesting contributions in terms of generic models (PERC) – very useful for distribution side and potentially transmission. Takes into account performance standards
- Ireland is of interest because of high data centre load compared to max demand. Coming out with some requirements for back up generation to support the grid. Also have ride through requirements.

Resources

- AEMO | Forecasting Reference Group (FRG)
- Energy and Climate Change Ministerial Council meetings and communiques | energy.gov.au
- Expectations of data centres and AI infrastructure developers | Department of Industry Science and Resources
- Improving the NEM access standards – Package 2 | AEMC (The draft determination is a good starting point. Also has reference to ERCOT standards)
- Operational integration and visibility of large inverter-based loads | AEMC
- Security framework enhancements | AEMC
- draft-2026-gpsrr-report.pdf – this report contains risk assessment of large loads coming online
- AEMO | General Power System Risk Review
- ESIG-Large-Loads-Performance-Requirements-report-2026c.pdf (has a good outline of fault ride through requirements overseas)

Natalie Bartos (she/her)
Manager – Operations Policy
s 9(2)(a)

ELECTRICITY AUTHORITY
TE MANA HIKO

Level 7, AON Centre, 1 Willis Street

Released under the Official Information Act 1982

PO Box 10041, Wellington 6143, New Zealand
www.ea.govt.nz

Released under the Official Information Act 1982

s 9(2)(f)(iv)

Released under the Official Information Act 1982

New Zealand's existing regulatory framework already provides for connecting parties to meet the costs of connection assets and a share of 'core network' costs over time. The approach is set via the Transmission Pricing Methodology (TPM).

Under the TPM, a new connecting data centre will typically meet most, if not all, of the costs directly attributable to its own connection, as well as a share of interconnection grid costs. If spare network capacity is available, adding a new customer may spread existing costs across a larger group of customers so reduce charges for others. However, where a significant new customer drives additional grid investment, a portion of those costs may be allocated to other customers who are assessed as also expected to benefit from the grid upgrade.

So, under the existing TPM, a new large customer could either increase or decrease transmission charges for others depending on specific circumstances. Subject to Transpower modelling support, various scenarios could be tested to understand the quantum and balance of these offsetting factors.